

A Policy Look at IPv6: A Tutorial Paper

By John C. Klensin¹, 16 April 2002

The views expressed in this paper are those of the author and do not necessarily reflect the views of ITU or of its members.

1	Introduction: What is IPv6?.....	2
2	The Address Space Exhaustion Problem: A history	3
2.1	Total address space, networks and classes.....	3
2.2	Class-based addresses, large networks, and subnetting.	4
2.3	The advent of the personal computer and other surprises.....	5
2.4	Giving up on Classes: The introduction of CIDR.....	6
3	Relationship to topology	6
3.1	Analogies to the PSTN.....	7
3.2	Telephone numbers and the Domain Name System.....	7
3.2.1	Circuit identifiers and addresses	7
3.2.2	New technologies	9
3.3	Reviewing the end to end model.....	10
3.4	The 32 bit address space and exhaustion	11
3.4.1	Predictions about when we run out	11
3.4.2	Consequences and how soon?	12
4	Some proposed alternatives	14
4.1	Application gateways	14
4.2	NATs, VPNs, and private spaces.....	15
4.3	Implications for innovation and expansion.....	15
5	Network problems and IPv6.....	17
5.1	Solving the address space issue	17
5.2	The routing problem.....	17
5.3	Provider-dependency in addressing and multihoming.....	18
5.4	Security	19
6	Space allocation policy proposals	19
7	Deployment difficulties and models	21
7.1	Communication in a mixed-network environment	21
7.1.1	Dual-stack environments.....	21
7.1.2	Tunnels.....	21
7.1.3	Conversion gateways	22
7.2	Converting the network themselves	22

¹ PO Box 400197, Cambridge, MA 02140, USA. Tel: +1 617 491 5735, email: Klensin@jck.com

7.2.1	Conversion at the edges	22
7.2.2	Conversion at the core.....	23
7.2.3	Adoption by large “islands”.....	23
8	Potential roadblocks and solutions	23
8.1	Economical.....	23
8.2	Technical.....	24
8.3	Policy/political	25
9	Summary: Thinking about conversion	25

1 Introduction: What is IPv6?

IPv6 (Internet Protocol, version 6) was developed by the Internet Engineering Task Force (IETF), starting in 1993, in response to a series of perceived problems, primarily with exhaustion of the current, IP version 4 (IPv4), address space. It arose out of an evaluation and design process that began in 1990 and considered a number of options and a range of different protocol alternatives. The design process was essentially complete, and a protocol specified, in the first half of 1995, although refinement work continues². The current version of the specification was published, after considerable implementation experience had been obtained, at the end of 1998³. Controversy continues to this day about some of the choices, but there are no proposals for alternatives that are complete enough for a determination to be made about whether or not they are realistic. The principal motivation for the new protocol was the address space issue on which the balance of this paper focuses. However, a number of other changes were made in formats and the interpretation of data fields. Those changes are intended to make the network operate better in the long term and to expand options for the design of efficient protocols, but their presence makes transition more complex than it would have been with address space expansion alone⁴. The driving motivation for IPv6 is to solve an address space exhaustion

² The history of the selection process, and its general conclusions, are described in Bradner, S., and A. Mankin, “The Recommendation for the IP Next Generation Protocol”, RFC 1752, January 1995.

All “RFC” documents (the term originally referred to “request for comments” but has become simply the name for a publication series) are available online from a number of “mirror” locations. The official copies are located at <ftp://ftp.rfc-editor.org/in-notes/rfcNNNN.txt>, where NNNN is the RFC number.

³ Deering, S. and R. Hinden, “Internet Protocol, Version 6 (IPv6) Specification”, RFC 2460, December 1998.

⁴ A different, but complementary, view of some of the technical issues and challenges discussed in this paper may be found in Carmès, E., “The Transition to IPv6”, ISOC Member Briefing #6, Internet Society, January 2002. A more general discussion of issues facing the Internet as it becomes more critical to commerce and society, including sections that address some of the topics covered here, appears in Computer Science and Telecommunications Board (CSTB), National Research Council, *The Internet’s Coming of Age*, Washington, DC, USA: National Academy Press, 2001.

problem, but some communities have argued strongly that this problem does not exist or can be avoided by completely different approaches.

2 The Address Space Exhaustion Problem: A history

2.1 Total address space, networks and classes

While one would prefer to make a given error only once and then learn from it, a few design errors have been repeated multiple times with what is now the Internet. Most of these have involved underestimating the rate or total scale of network growth. The original ARPANET design assumed that there would never be more than a large handful of hosts and, hence, that permitting a total of 255 hosts (an eight-bit host address space) would be more than adequate. When the TCP/IP architecture for the Internet was designed as a replacement in the first half of the 1970s, a 32-bit address space was then believed to be adequate for all time. Since the Internet—and TCP/IP—are designed around the notion of a “network of networks”, rather than a single, seamless, network, that 32-bit address space was originally structured to permit a relatively small number of networks (roughly 256), with a large number of hosts (around 16 million) on each. It rapidly became clear that there would be a larger-than-anticipated number of networks, of varying sizes, and the architecture was changed to support three important “classes” of networks (there was a fourth class, but it is not relevant to this discussion): 128 Class A networks, each accommodating up to 16,777,215 hosts, 16,384 Class B networks, each accommodating up to 65,535 hosts, and around 4 million Class C networks, each accommodating up to 255 hosts. As one might have anticipated, Class C networks turned out to be too small for many enterprises, creating a heavy demand on Class B addresses, but those were larger than any but the very largest enterprises or networks needed.⁵

The distinction between networks and hosts on a network was, and remains, very important because Internet routing is closely tied to the separation of routing within a network and routing between networks. Using the division of an address into a network number and a host number, a given host can determine whether a packet is to be routed locally (on the same network), using some sort of “interior” protocol or whether it must be routed, typically through a gateway (although there are actually slight differences in meaning, the term “router” is often used interchangeably with “gateway” or, more precisely, “network-level gateway”), to another, “exterior”, network. Exterior routing protocols use only information about networks; they pay no attention to what goes on inside a network.

This network-based approach has very significant limitations as far as utilization of the address space is concerned. As indicated above, one doesn’t really have nearly 2³¹

⁵ For a more extensive discussion on this subject, and the evolution of “CIDR” addressing, see, for example, Chapter 9 of Huitema, Christian, *Routing on the Internet*, 2nd Ed. New Jersey: Prentice-Hall, 1999. The authoritative documents on CIDR are Fuller, V., T. Li, J. Yu, and K. Varadhan, “Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy”, RFC 1519, September 1993 and Rekhter, Y. and C. Topolcic, “Exchanging Routing Information Across Provider Boundaries in the CIDR Environment”, RFC 1520, September 1993.

(somewhat over $2 \cdot 10^9$) addresses with which to work. Any hierarchical addressing system would have similar problems with density of address usage. Different subdivision methods permit addresses to be used more or less densely, but no system for allocating network addresses can achieve perfect density. Instead, the entire Internet could accommodate a much smaller number of networks, the vast majority of them very small. Worse, because class boundaries were fixed, if a network contained up to 254 hosts, it could use a network of Class C addresses, but, when the number of hosts rose even by two or three more, it became necessary to either allocate an entire Class B address space, potentially tying up sixty thousand or more addresses that could not be used for other purposes, or to allocate multiple Class C networks. The latter could be problematic in designing local network topologies, but also threatened explosive routing table growth and routing complexity.

This design nonetheless appeared reasonable for the early Internet, since the network was built around the assumption of relatively large, time-sharing, hosts, each with large numbers of users. If one thinks about an enterprise computing environment as consisting of a small number of mainframes, each with an address and with terminal devices attached to them that were not connected using internet or internet-like protocols, a Class C network with capacity for 254 hosts is likely to be more than adequate. Even if the norm were departmental computers, rather than centralized mainframes, and a few machines per department, only very few enterprises would anticipate more than 75 or 100 departments, and the class-based addressing system, primarily allocating from the large number of available Class Cs, still seemed reasonable.

2.2 Class-based addresses, large networks, and subnetting.

Another disadvantage of the “class”-based addressing system that became obvious fairly early was that some of the networks – all of the Class As and many of the Class Bs – became quite large and complex. Interior protocols that would work well for subsets of them would not work for the networks as a whole. Especially when the networks became very large (geographically or in number of hosts), one might actually wish for exterior-type protocols to route between components. This problem led, in the early 1980s, to the introduction of “subnetting”. Subnetting essentially provides for using the two-level network/host model within a network, so that one could now divide larger networks up into smaller ones and treat each as a separate (internal) network. This approach was particularly useful for enterprises with networks spanning very large areas, since it permitted multiple gateways and internal routing arrangements. But subnetting, like the change to Class-based addresses, was largely a response to routing issues – not enough networks in the first case and the need to subdivide networks in the second – rather than to concerns about address space exhaustion.

Despite these changes, it was generally assumed that any computer using TCP/IP, whether connected to the Internet or not, would be assigned a unique address. For connected machines, this was necessitated by the need to reach the machine and have it reach others (and more generally, by the “end-to-end principle”, discussed below). For machines and networks that were not connected, the term “yet” seemed applicable: a long-term trend emerged in which systems were built that were never

expected to be connected, only to have plans changed, resulting in a need to connect those networks. Renumbering was considered undesirable, less because of the problems associated with changing the address of a single host, but because of the need to simultaneously renumber all of the hosts on a network when it was connected: remember that a host handles packets bound for its own network somewhat differently than it does hosts that are on a different network or, more specifically, use a different network address. Thus renumbering is done in response to changes in routing, or typically requires routing adjustments.

2.3 The advent of the personal computer and other surprises

The appearance of small and inexpensive desktop computers changed all of this. Fairly quickly, the assumption that an enterprise or department would consist of a few large computers with attached terminals – with the terminals using a different protocol to communicate with the computers than the computers used to communicate with each other – evolved to a vision of networks as consisting of machines, interconnected with TCP/IP, and hence needing addresses whose numbers were roughly proportionate to the number of people, rather than the number of departments. Increasing modem speeds, combined with protocols that supported dialup use of TCP/IP with adequate authentication and management facilities for commercial use, made dialup networks and general home use of Internet connections plausible. As a result of this combination, Internet growth, spurred on by the introduction of the web and graphical interfaces to it, exploded. Several techniques were developed to reduce the rate of address space consumption below the rate of “Internet growth” (measured by the number of computers that were ever connected). The most important of these were

- (i) Dynamic assignment of addresses to dialup hosts, reducing the number of addresses toward the number of ports on dialup access servers, rather than the number of machines that might be connected.
- (ii) Increased use of “private” address space, i.e., the use of the same addresses in different locations on the assumption that the associated hosts would never be connected to the public Internet.⁶

These two approaches caused problems in edge cases, problems that presaged the requirement for a larger address space. Private address spaces required renumbering when the associated networks were connected (as anticipated some years earlier) and, worse, tended to “leak” into the public Internet when attempts were made to connect

⁶ The problem with privately-chosen address spaces was that they “leaked”, i.e., the addresses ended up appearing on the public Internet, where they conflicted with addresses privately chosen by others or with allocated and registered addresses. This resulted, in turn, in routing problems, some security risks, and, if nothing else, confusion. This “leakage” finally became a sufficiently large problem that the IETF and the regional registries decided to dedicate several address ranges to private use and recommend that Internet routers block these addresses if they appeared in the public network. The specification and some discussion appear in Rekhter, Y., B. Moskowitz, D. Karrenberg, G. J. de Groot, E. Lear, “Address Allocation for Private Internets”, RFC 1918, February 1996. But those private allocation spaces did not completely solve the problem, as discussed above.

the network through gateways that translated the addresses. Dynamic addressing, with a given host acquiring a different address each time it was connected to the network, worked fine but essentially prevented use of those hosts in environments in which other hosts needed to contact them (i.e., in peer-to-peer setups or as servers with client-server protocols). Patchwork solutions were developed in response to these problems, but they were not comprehensive in practice, or they introduced new or different problems. Those solutions, however, evolved into the alternatives to IPv6 discussed below.

2.4 Giving up on Classes: The introduction of CIDR

The final step in this sequence of changes to IPv4 addressing to better utilize the address space was the abandonment of the Classes and their fixed boundaries, replacing them with a classless system – Classless Inter-Domain Routing (see note 5). CIDR permitted the use of a variable-length network portion in the address, so that the remaining address space could be used more efficiently than the class-boundary network sizes permitted. An enterprise or network that needed, say, 500 addresses, could be allocated a network block with capacity for 511 or 1023 hosts, rather than requiring a full Class B network and “wasting” the remaining sixty-four thousand (or so) addresses. When very small networks became common a few years later, such as for home or small office networks using cable television or digital subscriber line (“DSL” or “xDSL”) connections, CIDR also permitted address allocations to be made in blocks considerably smaller than the original Class C (up to 255 host) ones.

At roughly the same time CIDR was proposed, the regional address registries adopted much more restrictive policies toward allocating space for those who requested it. The approval of CIDR reinforced this space-conserving trend, which some enterprises considered excessive at the time. Applicants were required to document plans for space utilization and to justify, in fairly specific terms, the amount of space they would need. The intent was not to prevent anyone from getting needed space, but to slow the rate of allocations and ensure that space was used as densely as possible.⁷ As discussed below, it is reasonable to assume that policies for conserving IPv4 space will become more aggressive as the available space is consumed.

3 Relationship to topology

To properly understand the Internet’s address space issues, it is probably useful to understand what the addressing system is and is not. In particular, an analogy is often drawn between internet addresses and telephone numbers, leading to discussions of number portability and alternate number routing. That analogy is, in most respects,

⁷ While the overall intent was the same, the regional registries adopted somewhat different policies to accomplish this goal. The instructions to the registries on the subject were contained in Gerich, E., ed. (IAB), “Unique Addresses are Good”, RFC 1814, June 1995. The high-level policies adopted appear in Hubbard, K, M. Kosters, D. Conrad, D. Karrenberg, J. Postel, “Internet Registry IP Allocation Guidelines”, RFC 2050, November 1996. The current policies of the RIPE NCC are probably representative and are implicit in their address space request “tips” memo, currently at <http://www.ripe.net/ripenncc/tips/tips.html>.

incorrect. Given its routing implications in a packet environment and binding to a particular interface on a particular machine, an Internet address can be more accurately compared to a circuit identifier in the public switched telephone network (PSTN) than to an (E.164) number. With some further adjustments because of the difference in character between circuit-switched and packet-switched networks, the Internet analogy to a PSTN telephone number is, for most purposes, more accurately a domain name. This point is fundamental, so bears repeating: an IP address is tied to routing information; it is not a name. A telephone number is a name and not a route. So a telephone number is more similar to a domain name than it is to an IP address.

3.1 Analogies to the PSTN

For many years, the primary mechanism in the PSTN for mapping from the names of people to telephone numbers (and thence to routing information and circuits) has been a “white pages” service, supplemented by operator services. Although there have been local exceptions, there has never been a successful global, or globally interoperable, “white pages” service for the Internet. That situation is largely due to the competitive and regulatory environments of telephony services, which are still largely national, with international service being an additional service, priced at a premium and negotiated between carriers or countries on a bilateral basis.

By contrast, Internet services have been largely international from inception. While prices differ from one locale to another, there is essentially no pricing difference between national and international services.

These differences have had a number of implications, one of which has been additional confusion about the role of Internet domain names vis-à-vis host addresses. The confusion has sometimes been increased by the fact that IP address formats and usability are independent of the physical media by which the associated hosts are connected to the network: they are not physical-layer addresses. The next section discusses some of the additional issues associated with these distinctions.

3.2 Telephone numbers and the Domain Name System

The Internet’s domain name system (DNS), like telephone numbers, provides a certain amount of portability of reference. One can retain a name and have the underlying address (or circuit) change and can even, at least in principle, change transport mechanisms – e.g., from wireline to wireless – with the same number. But IPv6 introduces something of a new twist, since, with its deployment in the network, a DNS name may be associated with one or more IPv4 addresses, one or more IPv6 addresses, or any combination of them. If both types of addresses are returned, the initiating system will normally choose to communicate with IPv6, preferring addresses on networks which it can reach most directly, but other options are possible if required by performance or policy considerations.

3.2.1 Circuit identifiers and addresses

As mentioned above, historically the closest analogy to an IP address in the PSTN is a circuit identifier. However the analogy is not exact: the IP address does not directly

represent a physical-layer entity. In addition, IP addresses are potentially visible to users and user-level applications. By contrast, circuit identifiers are not only invisible to the user but would normally be of no value if obtained (e.g., one cannot place a telephone call using a circuit number). And both systems have evolved somewhat in recent years, yielding different models of what their practical lowest-level identifiers actually represent.

3.2.1.1 Fixed addresses in some networks

With dedicated, “permanent”, IP attachments, as with conventional, wireline, telephone systems, an address identifies a particular terminal or host and conveys a good deal of the information needed to access it. The information is not volatile: while it can change, the nature of changes is that they typically scheduled to occur over relatively long periods of time, and can be planned for and adjustments made on a fairly leisurely basis. For the telephone system, this is true even when number portability is introduced – changes do not happen overnight and without warning. In the Internet case, such addresses are typically configured into the machine itself: a change in address requires reconfiguring the machine in, depending on the operating system, a more or less significant way. Renumbering such machines, whether within IPv4 or IPv6 space, or from an IPv4 address to an IPv6 one, involves often-significant per-machine costs. IPv6 installations are expected to put more reliance on dynamic (really centralized-server-based or server-less automatic configuration) allocation of addresses than has been typical in IPv4 networks, partially to reduce these problems (the facilities to do this were not available when IPv4 was deployed).

From a routing standpoint, these addresses are completely stable: their network portions can safely be used as routing information, if the network’s location changes with regard to the topology, the address will normally change as well.

3.2.1.2 Long-timeframe variable addresses in other networks

Other types of addresses are more volatile and less predictable by the user, but still involve relatively long periods of stability. Internet address assignments that are dynamic in nature, but bound to particular equipment, tend to be much longer lived than those that are likely to differ over very short times and be linked to transient phenomena, such as dialup attachments.

Dynamic host configuration protocol (DHCP) assignments using DSL or cable “modems” or hardware (“MAC”) addresses typically have these properties. The addresses are usually stable for many months, but can be reassigned, over a fairly short period, by a change in the tables of the supplying Internet service provider. However, unless those tables are linked to the DNS records for the user site, the change will need to be reported to the user and the DNS changes made to correspond. Problems often occur during the period between the address change and the DNS changes. Finally, it is usually undesirable for the provider to operate the DNS entries for their customers – normally a requirement if address changes are to be immediately reflected in the DNS – since it prevents the administrators of the local area network (LAN) managing their own name spaces.

From a routing standpoint, these addresses are the same as the fixed ones – routing table updates occur on a per-minute or per-hour basis (or less) and these addresses are very long-lived by comparison to the network's ability to reflect routing changes.

3.2.1.3 Short-timeframe variable

Other types of addresses actually are dynamic, even as the Internet measures time. A mobile telephone that is actually moving between locations requires routing information that is dependent on more than its address, or requires rapidly-changing addressing or a different form of addressing. For many purposes, dynamic address assignment for dialup connections, in which each connection is likely to acquire a different address than the host had on the previous connection, raises similar problems. Of course, rapidly changing addresses also have advantages (a least as seen from some perspectives), e.g., they make usage more private and the user somewhat harder to trace.

IPv6 will make enough address space available to expand the options in these areas; with IPv4, optimality is, of necessity, usually defined in terms of address space conservation.

3.2.2 New technologies

Just as telephone numbers have evolved from a tight binding to a circuit to a number of portability ideas and then to wireless access, there are some notions about mobility and reuse of IP addresses and their use independent of their apparent routing properties. Each has some unfortunate side effects – a globally accessible mobile phone number can lead to calls at unfortunate times of day if one is in a time zone far removed from one's normal location and a caller doesn't know that, but can also be very helpful. With IP addresses, there are specific mobility techniques that are not relevant to this discussion and a mechanism, called "network address translation" (NAT) that permits the addresses used within a particular network to be different from the global, routable, addresses by which that network and its hosts are addressed from the broader Internet.

This translation capability can serve several purposes. If the addresses on a network need to be renumbered due to a change in connectivity (whether voluntary or imposed by a provider), address translation may permit the "new" addresses to appear externally while avoiding reconfiguration of the machines on the network (proper use of DHCP, as suggested in section 3.2.1.2, is almost always a better solution to this problem). If there is an address shortage, address translation and port remapping may be used to permit several hosts on a LAN to share a single globally accessible address, thereby helping to conserve global address space, but with some costs and tradeoffs (see sections 4.2 and 4.3).

NATs, and address translations more generally, are discussed more extensively in subsequent sections.

3.3 Reviewing the end to end model

The design of the Internet depends critically on what is known as the end-to-end architecture and the associated “hourglass model”⁸. The key principles that impact IPv6 considerations can be summarized as:

- (i) From an addressing and protocol standpoint, any host on the Internet should be able to identify and access any other host through a uniform set of protocols and global addresses. Of course, security considerations may restrict what, if anything, can be done with that access, but the principle remains important.
- (ii) The network itself is “dumb”, that is, it is insensitive to the applications and protocols being run on it above the IP level. Changes are not required to the network to add new applications or terminal device capabilities.
- (iii) By consequence, the terminal devices are assumed to be as “smart” and capable as needed to support the applications that users intend to run. Uniformity of terminal device capabilities around the network is not necessary or assumed.
- (iv) The architecture of applications above the IP layer is dependent on that layer only; applications do not depend on, nor are they aware of, the physical media, interfaces, or interconnection devices being used. Consequently, introducing a new medium, or new link-level, has no impact on applications.
- (v) The network does not change the content or other properties of traffic flowing through it, i.e., it is transparent to that traffic.

This combination of characteristics is actually one of the key differences (other than the obvious “circuit” and “packet” designs) between the design of the Internet and that of the PSTN. The latter is built around an assumption of an extremely dumb terminal with almost no local capability. Those terminals are also assumed to be fairly uniform as far as the network is concerned – additional capability in the terminal does not permit more sophisticated signaling into the network. Of course, ISDN and mobile systems change this model somewhat – the computational, routing, and interface capabilities of even 2.5G cellular devices considerably exceed that of a PSTN desk phone and are equivalent to some Internet devices – but, for the base PSTN, the characteristic remains.

Oversimplifying somewhat, pressures on address space and responses to them pose three serious challenges to this traditional Internet model. The introduction of NATs to conserve address space implies that there is no longer a global end-to-end model,

⁸ A more extensive, and excellent recent discussion on the hourglass architecture and its implications was presented by Steve Deering at IETF 51; the presentation can be found online at <http://www.ietf.org/proceedings/01aug/slides/plenary-1/index.html>. A slightly different version appears in *The Internet’s Coming of Age* report, *op. cit.*

but, instead, that there are isolated hosts and LANs that can be accessed only through intermediaries. IPv6 itself is a disruptive technology, since it changes the IP layer and, inevitably, interfaces to it in ways that do require altering applications or applications services. And most of the transition strategies for a network containing a mix of IPv4 and IPv6 systems (see section 7.1) imply that some hosts will require more complex arrangements to reach others than the simplicity that the end-to-end architecture contemplates⁹.

3.4 The 32 bit address space and exhaustion

3.4.1 Predictions about when we run out

Since the beginning of the last decade, predicting the date on which the IPv4 address space will be exhausted has been a popular sport, with a variety of estimates deriving from different assumptions and calculations. Under the most pessimistic of those estimates, we would have run out of addresses already. Some people observe that address exhaustion has not occurred and infer that there is no risk of ever running out. Early in the process that developed IPv6, the IETF's working group on the subject estimated address exhaustion between 2005 and 2011, and some people believe we are still on that schedule. Others see a longer lifetime for the IPv4 space (largely because of CIDR and NAT), and still others continue to predict that catastrophe is right around the corner.

More realistically, we should understand that any estimate of an exhaustion date is made on the assumption that neither major technology nor policy changes will occur or, at best, with only the most tentative estimates of the impact of such changes. Looking backward, those early, most pessimistic, estimates did not fully account for some of the effects of CIDR, nor of provider refusal to route small blocks¹⁰, nor of more restrictive regional registry allocation policies, nor of heavy use of NAT and private addresses, each of which has had a significant impact on the rate at which addresses have been allocated and consumed. Although there have been some

⁹ *The Internet's Coming of Age* report, *op. cit.*, discusses end to end transparency issues in more detail. Some of the issues remain controversial; the introduction of many types of functions into the core network to mediate different functions (performance, security, efficiency of access to selected materials, etc.) can be seen as threats to the model even in the absence of IPv6 or NAT concerns. See, e.g., Carpenter, B., "Internet Transparency", RFC 2775, February 2000 and Kaat, M., "Overview of 1999 IAB Network Layer Workshop", RFC 2956, October 2000. A summary of the challenges and an argument that it is time to rethink important aspects of the model appears in M. Blumenthal and D. Clark. Rethinking the design of the Internet: The end to end arguments vs. the brave new world. To appear in *ACM Trans. Internet Technology*. Also to appear in *Communications Policy in Transition: The Internet and Beyond*. B. Compaine and S. Greenstein, eds. MIT Press, Sept. 2001. See also the presentation by Deering, referenced immediately above.

¹⁰ After CIDR was introduced, a few ISPs adopted policies, to protect their own routing tables and those of the Internet more broadly, that restricted the routes they would accept from other ISPs to networks they considered large enough. The net effect of this was that very small networks ("long prefixes") could not be reliably routed across the entire Internet. Those policies, in turn, further increased pressure to aggregate addresses into larger, ISP-dependent, blocks. But those policies were adopted spontaneously by specific ISPs and were never mandated by either the registries nor the IETF.

extensions to NAT to expand the range of situations in which the technique can be applied, there are no further major technical tricks in the IETF queue that would reduce address space pressure. Applications could come along that would dramatically increase it. And it is sensible to anticipate that the regional registries will apply increasingly restrictive allocation policies as the available remaining space shrinks.

In addition, there are very large blocks of addresses tied up in Class A and Class B blocks allocated in the early days of the Internet. Organizations that hold those blocks have little or no incentive to renumber out of them, and are typically very sensitive to the considerable costs of doing so. But it appears fairly obvious that, at some point, if there is enough pressure on the address space, economics will shift sufficiently to force exactly that renumbering. For example, several of the “legacy” Class A address blocks are held by private universities in the USA. Those institutions tend to be fairly insensitive to low-level economic pressures, but are not insensitive to legislative or significant economic ones (in, e.g., units of buildings or long-term endowment)¹¹. One can easily imagine consortia formed to make financial offers for address space or mounting pressure for legislation if available address space becomes limited enough to impede significant new commercial ventures.

3.4.2 Consequences and how soon?

True address space exhaustion – with no additional space available for new hosts or new applications – would be a nightmare of the worst sort and, effectively, the end of the Internet, at least as a growing entity, since the time to deploy new protocols or applications would be significant, at least unless a significant portion had been converted to IPv6 (or, in principle, some other strategy) already. However, the expectation of changes in technology and policy, predicted long-term difficulties if NATs must be nested to preserve space or if additional NAT-resistant applications or work styles are deployed, and economic or other changes that might free up space that is poorly utilized by today’s standards makes it unrealistic to make policy decisions based on the date at which the address space will be exhausted. Instead, it is useful to examine more likely outcomes if an address space expansion strategy is not deployed and globally accessible addresses continue to be necessary.

- (i) The regional registries gradually shift toward allocation policies that require ever-stronger justifications for additional space and, in a potential reversal from today’s stated policies, require explanations and justifications of why private space cannot be used. Their allocations from IANA (ICANN)¹² also gradually tightened up, requiring the registries

¹¹ Stanford University did voluntarily try to set an example: it renumbered out of its original Class A allocation and returned the space to the general Internet community, but that action has, so far, been unique.

¹² IANA, the Internet Assigned Numbers Authority, was the original registry for all Internet protocol identifiers and the top-level allocation source for addresses and domain names. When the IANA function was moved out of the Information Sciences Institute of the University of Southern California and privatized, it passed to ICANN, the Internet Corporation for Assigned Names and Numbers.

themselves to prepare stronger justifications for allocations of space from the remaining pools. It is likely that governmental or other pressures on ICANN could introduce distortions in the policies and allocations to the registries in the interests of some perception of fairness.

- (ii) A serious and expensive market opens up to recover underutilized space from legacy allocations, primarily in Class A and B space. Some sites and ISPs discover that the “network real estate” they occupy in address space has a higher market value than their business (based on any reasonable capital valuation of the latter) and they, or at least their server functions, are forced off of the network. This is likely to lead to additional pressures for legislative intervention in allocation policy decisions, but such interventions can ultimately only influence the price structure associated with mandatory renumbering and reallocation, or can allocate scarcity, but cannot create more address space.

There are some further complexities associated with the notion of a market in address space, even in large blocks (small blocks can encounter the practical restrictions discussed in note 10). For example, as a security measure (one that has proven important in resisting attacks in the past), some ISPs will only accept address and routing announcements from sources that are registered with the appropriate registry as authorized to utilize that space. This implies that a transfer, if made between an address-holding party and a party that wishes to acquire the space, may be useless unless the registries agree to it. Proposed new security policies and protocol modifications would reinforce this trend. On the other hand, if the available IPv4 address space reaches the end game contemplated by this section, it is hard to imagine the registries being able to adopt and retain overly restrictive policies relative to approval of transfers that extend the useful life of the address space¹³.

The best way to look at the net result of these dismal scenarios is that one does not want to be the organization or enterprise that requests the last block of available IPv4 space, or even any blocks of space after things become significantly tight.

¹³ A further policy complication may be worth noting. Some people in the Internet community have seriously proposed that criteria for allocating IPv4 space should be considerably relaxed, and relaxed immediately. They suggest that the current restrictive model is distorting Internet growth and retarding important developments (such as extensive use of attachments by small networks to more than one ISP and protocol and security strategies that require multiple addresses per host). A subset of them suggest that the side effect of a more relaxed allocation strategy – rapid exhaustion of the IPv4 address space – would actually be an advantage, because it would force rapid movement to IPv6.

4 Some proposed alternatives

4.1 *Application gateways*

Most of the plausible suggestions for continuing to use the IPv4 packet and addressing for an extended period, potentially indefinitely, depend on localizing addresses and then using the same addresses in multiple locations around the network.

The classical version of this approach involves placing one or more applications servers at the boundary of a network, using public addresses for those servers. Those servers then accept all inbound traffic, convert it as necessary, and transfer it, as needed, to the “inside” network. With this approach, it is not necessary that the “inside” network be using the same addressing arrangements as the “outside” one and, indeed, it is not necessary that it be running TCP/IP at all. It is believed that the first instances of this approach were developed when protocol conversion – for the network protocols, the applications, or both – was required, e.g., when one network was operating with SNA (originally, “system network architecture”, an IBM proprietary protocol suite whose original versions focused on access to mainframes) and the other with TCP/IP or when email arrived using an Internet protocol such as the “simple mail transfer protocol”, better known just as SMTP, and email on the “inside” network used a proprietary protocol such as cc:mail® or MSMail®¹⁴.

This “application gateway” approach had (and has) a few disadvantages. For example:

- (i) It was necessary to locate the conversion (“gateway”) servers at the boundary of the “inside” network, rather than inside it. Under many circumstances, this caused security and operational problems.
- (ii) While some applications and protocols, such as email, lend themselves to a “receive, store, convert, forward” architecture, the approach is hostile to anything that requires real-time communication at the packet level.
- (iii) Hosts on the “inside” were subject to the same restrictions as those on the “outside”, i.e., they could communicate only via the applications gateways, not directly.
- (iv) Since the gateways needed to accept the application traffic, understand it, and potentially convert it to other formats, any application, or application option, that was not preprogrammed into the gateway could not be accepted or accessed at all. Hence, the presence of such gateways tends to create huge impediments to deploying new type of applications.

For situations in which the network technology on both sides of the boundary was TCP/IP, Network Address Translation (“NAT”s or “NAT boxes”) was developed.

¹⁴ cc:mail and MSMail are registered trademarks of IBM and Microsoft Corporation, respectively.

NATs solved some of the problems with application gateways used alone, but not others.

NAT advocates also turned a few of these problems into virtues as the notion of security firewalls evolved. For example, if an “inside” network was invisible to an outside one, and could be accessed only in terms of specific, permitted, protocols, that provided a certain amount of protection against malicious actions, especially if other protocols were involved.

4.2 NATs, VPNs, and private spaces

NAT arrangements, or, more often, firewalls with address translation and tunneling capabilities, have been used to design and build virtual private networks, with the tunnels used to interconnect LANs that share a private address space. Such networks have proven quite useful to many enterprises. However, inevitable changes in business relationships, mergers, and sometimes spin-offs, have resulted in a need to consolidate previously independent private address spaces. That, in turn, has forced sometimes-painful renumbering when the addresses used have overlapped. If the networks involved were using IPv6 rather than IPv4, the much greater address space available would permit use of public addresses (even if the sites were administratively blocked and could not be accessed from the public network) so that these conflicts would not arise.

4.3 Implications for innovation and expansion

There are a few major difficulties with NAT-based approaches, difficulties that may or may not be issues for a given network and configuration. To state the issue positively, if the entire network “behind” a NAT consists of machines that are used only as clients of other Internet services, and there is only one NAT between that network and the uniquely-addressed Internet core, it is unlikely that problems will be encountered. Indeed, arrangements of that type are used in thousands of locations today and are essentially required when, e.g., a cable model or DSL provider refuses to provide more than one address to a subscriber who actually has a home network, or makes excessive charges for additional addresses. Even in those “client-only” setups, it can be somewhat more difficult to diagnose problems that occur, since the addresses “inside” the local network are not visible from other Internet locations, and do not correspond to external addresses. In practice, this has not been a significant problem for most NAT users, possibly because few ISPs provide useful diagnostic help for problems within user networks even when public addresses are used.

On the other hand, there are many ways to violate the “client-only, no nested NATs”, assumption. Some of the more significant examples include:

- (i) The assumption that machines within the user network only support client functions impedes the use of any type of server or peer-to-peer function of the inside network that are intended to be accessed from the outside one. The Internet was designed for what are today called peer-to-peer functions and some protocols, including the famous Napster music exchange one, depend heavily on peer-to-peer functions. Operation of a home web site or

external access to home control capabilities, typically require server functions on the home network. Remote access to computers on the local network also depends on either peer-to-peer or server functions. As an important example, the two protocols that support the use of telephony-like voice over the Internet – the SIP¹⁵ protocol and the Internet adaptation of H.323 – do not work through NATs unless the NATs are extended with special modifications for those protocols. Finally, it is difficult or impossible to predict future protocol evolution – a new protocol could easily arise that requires peer-to-peer or server functionality on networks that are now behind NATs, requiring operators of those networks to switch to non-NAT (public address) or more complex arrangements and producing a rapid increase in demand for public addresses.

Newer NAT architectures do provide for servers on the network behind the NAT by doing what is often called “port mapping”. With this facility, the NAT box appears to be a single server to the external network and provides internal connections to support particular protocols, with each chosen protocol mapping to a single internal server. This approach is fairly satisfactory (although having servers on the internal LAN increases the need for external problem diagnosis mentioned above) but limits the LAN to a single server for a given protocol or requires complex internal proxy functions. For example, a LAN that requires a primary and backup mail server to ensure availability or performance is extremely difficult to configure with a NAT.

A few (non-standard) protocols operate by trying to hunt for a port on which to operate and will operate on any port that supports their services. Since a NAT must typically be configured for each port and protocol for which services are to be provided to the external Internet, use of NATs would potentially make such protocols almost impossible to use.

- (ii) The user will typically have no control over NAT nesting functions. If the supplying ISP decides to create a NAT setup rather than provide public addresses to any of its customers, then the port mapping of server functions in the local LANs may become essentially unusable since the external NAT would be required to designate a single internal address (NAT or otherwise) for each protocol to be supported.
- (iii) IP-level security protocols, notably the specific one called “IPSec”, require end to end addresses and do not function properly in NAT environments in which either or both endpoint hosts use private addresses. Where appropriate, this problem can be circumvented by assuming that hosts on the local LAN are trusted and operating IPSec only between the network boundary and the remote host or network. However, this requires a somewhat more sophisticated NAT box than many of those now available,

¹⁵ The name “SIP” is derived from “Session Initiation Protocol”, a term that no longer accurately reflects the uses or function of that protocol.

since it must establish its own security associations, and will not work satisfactory in a nested NAT environment in which the “outer” NAT is not controlled or trusted by the owner of the protected LAN.

Of course, each of these disadvantages might be considered as advantages if the operator of the NAT wanted to prevent use of server functions, IP-based security, or unanticipated protocols on the local LAN¹⁶.

It may also be worth stressing a point that is implicit in the comments above. When NATs are nested, or need to support server functions, or involve other situations similar to the examples above, they often become quite difficult to configure properly and maintain, even for experts. That cost should not be ignored or underestimated; for some organizations, it may be a primary motivator for considering an early transition to IPv6.

5 Network problems and IPv6

5.1 Solving the address space issue

The obvious solution to a shortage of available address space is more address space. That, of course, isn't the only solution – others (most discussed elsewhere in this document) include reducing pressures on the address space by becoming more selective about which resources actually need addresses and determining ways to reuse addresses in different parts of the network. Nonetheless, while it provides some other benefits, IPv6 is essentially a “larger address space” solution. If the growth of the Internet – or, more specifically, the unique-address-using Internet – were to stop tomorrow, IPv6 would probably be unnecessary¹⁷.

5.2 The routing problem

Many people have asserted that the key capacity issue with the Internet is not a shortage of addresses, but growth of the routing tables, i.e., the information that non-default routers must have and process in order to route packets to their destinations. IPv6 does not address routing problems in any direct way. IPv6 is designed to permit easier renumbering of hosts and local networks to accommodate changes in topology, but there is a great deal of skepticism about whether those facilities are going to be

¹⁶ A more extensive treatment of the implications and risks of using NAT setups appears in Hain, T., ed. (IAB), “Architectural Implications of NAT”, RFC 2993, November 2000. D. Senie discusses ways to design applications that work better with NATs than some others in “Network Address Translator (NAT)-Friendly Application Design Guidelines”, RFC 3235, January 2002, but, of course, most LAN administrators and users have little control over the design of the applications they would like to operate.

¹⁷ “Probably” is, unfortunately, a necessary qualification here. As discussed above, some potential applications, features, and styles of network design are incompatible with NAT arrangements. No one actually knows how many hosts are “hidden” behind NAT arrangements today. It is possible that an abrupt conversion of all of them to public address space in order to accommodate NAT-hostile applications or features would essentially exhaust the existing IPv4 address space.

significantly useful in practice. Debate continues, in the IETF and elsewhere, about whether the DNS mechanisms designed to complement the renumbering facilities should be retained.

For the last several years, routing table growth has been dealt with by brute force: availability of processing capability and memory in routers has managed to keep up with table growth, which slowed considerably after CIDR was introduced, and slowed further when some key ISPs started refusing to route traffic for very small, independent, networks across their backbones. There is some evidence that growth rates have increased again in the last year, probably due to increased use of “multihomed” connections (connections of a host, LAN, or subnet to more than one provider). While less readily measurable than table sizes, the length of (clock) time needed to make complete routing computations (or, more precisely, to have those computations converge) has also been rising, threatening a situation in which routing changes arrive at a higher rate than can be accommodated¹⁸.

IPv6 may help with the routing table problems simply by giving us an opportunity to organize the address space in a more optimal way (from a routing standpoint), without having to deal with more than a decade of legacy allocations that did not follow a scheme we would consider appropriate today (see the next section). Also, most new research and developments in the routing area assume the larger address space of IPv6 and the associated flexibilities; if we reach a point at which major changes in Internet routing algorithms are needed, those changes are much more likely to be practical in IPv6 than in IPv4.

Even with IPv6, some decisions about addressing and routing may have profound implications. For example, an ISP (or multiple-ISP consortium) with global reach might prefer to have either a single prefix and a fairly complex routing policy or multiple prefixes that reflect a more regional network organization. Even if the first decision were made, it is likely that it would want to allocate addresses within its network according to some topological structure so that interior routing could be easily optimized.

5.3 Provider-dependency in addressing and multihoming

The key to CIDR’s role as a routing table conservation strategy is that, to the degree possible, blocks of addresses would be allocated to ISPs, which would then allocate sub-blocks of them to their customers. If this worked perfectly, a single address block could be used to represent any given major ISP. Any other site or network would need to know only how to route to that network. And, once traffic reached that ISP, it would know how to reach its customers. Those customers, who might be yet other ISPs, would draw address space from the large allocations made to their parent ISPs and would reassign parts of it to their own customers or organizational locations. Consequently, smaller sites, and even smaller ISPs, would not be permitted to obtain space directly from regional registries and would be required to obtain it from their

¹⁸ For a more extensive discussion of these issues, see, among others, the references mentioned in footnote 5.

upstream ISPs. This model could not work perfectly for a number of reasons, and created many (predictable) resentments. A change in providers would mean renumbering, since the addresses belonged to the providers. A site that wished to connect to multiple providers (so-called “multihoming”) for additional robustness or efficiency would either require provider-independent address space or would need to somehow “punch holes” in the CIDR block of one provider or the other (with equal, or worse, damage to routing table compactness). And, of course, much of the address space was already allocated to sites and networks who had, as discussed below, little incentive to disrupt their networks and give up flexibility in the general interest of Internet efficiencies that they wouldn’t experience directly.

IPv6 may permit some rationalization of this situation, both by providing a way to start over on address allocations (without the legacy effects) and by providing somewhat cleaner mechanisms for allocations and renumbering.

5.4 Security

One of the characteristics of IPv6 that was controversial when the protocol was approved was the requirement that all implementations provide support for end-to-end security (privacy and integrity) of the data stream. That provision recognized the fact that security provisions to prevent intermediate sites or networks from tampering with, or eavesdropping on, packets in transit are increasingly important in today’s Internet and that such provisions become much more useful when supported by all sites. But, while it guarantees that the facilities are available, IPv6 does not intrinsically provide more security than IPv4: the IP-layer security (“IPSec”) tools proposed for IPv6 were modified and designed to support both IPv4 and IPv6¹⁹, and considerable configuration and key management are needed to make the security provisions work effectively with either protocol stack²⁰. Nonetheless, by reducing or eliminating the need for NAT, IPv6 should facilitate use and deployment of end-to-end IPv6.

6 Space allocation policy proposals

The question of how the registries should allocate IPv6 address space to those who request it remains an ongoing controversy, at least at the detail level²¹. Some points seem clear:

¹⁹ The base architecture for IPSec is described in Kent, S. and R. Atkinson, “Security Architecture for the Internet Protocol”, RFC 2401, November 1998.

²⁰ A similar situation applies with Quality of Service (QoS) capabilities: while the ability to accommodate them was important to the design of IPv6, the protocols that have been designed and standardized will operate in either IPv4 or IPv6 environments.

²¹ The general principles that the registries are expected to follow are specified in IAB and IESG, “IAB/IESG Recommendations on IPv6 Address Allocations to Sites”, RFC 3177, September 2001. Two of the three regional registries have accepted those recommendations, and discussion is in progress in the third. However, that document specifies the sizes of blocks to be allocated, rather than

- (i) Allocations will continue to be made to the regional registries and then further allocated to ISPs and larger (in terms of network requirements) enterprises user organizations, as they are today with IPv4. This differs from the telephony model in which numbers are allocated on a national basis, but, as with the address system itself, reflects routing necessities. The regional registries themselves will develop the detailed allocation rules, as has been the case for IPv4²².
- (ii) If allocations are made on the same basis as they have been with IPv4, i.e., by requiring plans and justification for a very high density of use of the space requested, the perception will be that IPv6 space is as difficult to obtain as IPv4 space, even though it will last a good deal longer. Such scarcity will impede the development of applications and approaches that require multiple addresses per host, many small-scale (“ubiquitous”) IP-connected and globally accessible computers, and other address-consuming applications.
- (iii) If allocations are made on an “anyone gets whatever they ask for” basis, we could, within a reasonably short period, be facing the same address space exhaustion and routing table problems we have with IPv4. As mentioned elsewhere, one of the attractions of IPv6 is that it permits us to start the allocation process over and in a way that is consistent with longer-term and more rational planning.
- (iv) Some allocation requests and proposals have been made on a basis that make no sense given the constraints imposed by routing of network addresses. E.g., the Internet is not the PSTN, nor is it the connection-oriented OSI network represented by E.121 addresses. While the management advantages of consolidating such addresses are obvious, each set has routing implications in its own context. The number of countries in which Internet routing and connectivity to the outside exactly parallels that of the other networks is very small, if not zero; the number in which that situation not only prevails today but can be guaranteed to continue is even fewer.

In summary, whatever strategies for IPv6 address allocation are considered or adopted, the relationship between the network portion of those addresses (the address “prefix”) and the Internet’s current and likely future routing architectures is such that allocations and address assignment must follow the routing topology and neither physical or political geography nor the topology of other telecommunications networks.

details about the criteria to be used to accept applications for the allocations – those criteria will still need to be established and interpreted by the registries.

²² See notes 7 and 21 for additional discussion and references.

7 Deployment difficulties and models

Any process for conversion to IPv6 requires mechanisms for operating the network with a mixture of IPv4 and IPv6 hosts and subnetworks. This section describes three possible mixed-protocol network approaches and then discusses the models by which conversion may occur. While the issues identified here are important, and will continue to be important as long as IPv4 hosts remain attached to the network, it is useful to understand that native IPv6 operation is no longer an experiment: the “6bone” (and IPv6 network running on top of the IPv6 one) now reaches almost a thousand sites in 56 countries and native IPv6 connectivity was available at the IETF meeting in March 2002.

7.1 Communication in a mixed-network environment

There are three major models for operation of a network that involves both IPv4 and IPv6 protocols. Even the most optimistic of IPv6 proponents now believe that there will be some IPv4-only machines on the Internet for a very long time, so at least some of these scenarios are important both for conversion and for long-term operations.

7.1.1 Dual-stack environments

In a dual stack environment, both IPv4 and IPv6 capability is available on every machine. When a pair of machines is capable of communicating using IPv6 (as noted for a remote machine by the presence of IPv6 information in its DNS entry), they do so, perhaps tunneling across intervening IPv4-only networks; when one or the other is not, they communicate using IPv4.

When IPv6 was originally defined, and the period of transition was expected to be shorter, the dual-stack model was assumed to be the primary transition mechanism. However, a transition based on dual-stack tends to assume an edge-first transition process (see section 7.2 below), while other techniques may be more useful if IPv4 machines and LANs are to be connected to IPv6 networks.

7.1.2 Tunnels

As with many other network technologies, it is possible to embed IP packets “inside” other IP packets, so that part of the network is aware of carrying only the “outside” packet, with encapsulation occurring at one boundary and unpacking occurring at another. This technique, called “tunneling” in many of its applications, is often used for security (the “inside” packet may be encrypted to hide its content and source and destination information) or to support some mobility and rerouting strategies, but it can also be used to embed IPv6 packets inside IPv4 ones and vice versa. In particular, two hosts or networks that are utilizing IPv6 can communicate with each other across an IPv4 network if the routers or other devices at the boundaries between the two network types encapsulate outgoing IPv6 packets in IPv4 ones and the corresponding devices unpack and discard the IPv4 “wrappers”, leaving the IPv6 ones at the far end. Conversely, a similar technique can be used to permit a pair of IPv4 networks to communicate across an IPv6 infrastructure (although there may be good alternatives for that case that do not require tunneling).

7.1.3 Conversion gateways

Finally, one can actually convert packets from one format into the other. This can, of course, be done at the applications level, viewing the activity as a protocol conversion one, but techniques have been proposed for reversibly embedding IPv4 addresses and packet structure within IPv6 ones that do not require tunneling. For the medium term, these techniques are probably the preferred ones, where feasible, for communication between IPv4 and IPv6 networks and endpoints. However, they do present some of the same issues as described above for NATs, although typically for NATs with the same number of addresses on the “inside” and “outside” networks and more obvious address mappings. E.g., they should be much less problematic for problem diagnosis (from either side) than a multiple-inside-address NAT.

7.2 *Converting the network themselves*

The techniques discussed above outline mechanisms by which IPv4 and IPv6 networks can work interoperably during a period in which the Internet contains some instances of each. From a policy and economic standpoint, perhaps a more interesting question is how IPv6 might actually be introduced on a large-scale basis. Two models were proposed early in IPv6 development: conversion occurring first from the edges and conversion from the core of the network (the backbone ISPs) outward. Both of these depend on “tipping” behavior after some critical mass is achieved. Conversion from the edges of the network has historically been considered the more likely although some opinion has been shifting toward the alternative of conversion from the core of the network. A third model – adoption by large connectivity “islands” – has recently appeared and is the primary model of IPv6 introduction today.

7.2.1 Conversion at the edges

In the edge conversion model, individual enterprises or LANs decide to convert, perhaps because they are feeling the address space restrictions for one reason or another – the perceived difficulty of obtaining IPv4 addresses typically being a major factor -- and see too many disadvantages associated with classical NAT setups. These types of conversions have occurred already, but have been limited by, among other factors, the lack of availability of production-quality IPv6 “stacks” in popular desktop operating systems. For many LANs, NAT-based technologies provide a clear alternative to this type of conversion, especially when combined with security approaches that include NAT-capable firewalls and “DMZ” strategies that are installed anyway.

Another edge-conversion option, which has not proven popular, involves running “dual stack” on both machines within the LAN and on the routers at its edges.

Until the bulk of the network, and especially provider-ISPs, converts to IPv6, operating a LAN with IPv6 requires either dual-stack or some type of address and protocol translation at the LAN boundaries to communicate with other sites. On the other hand, as more “edge” customers convert, the presumption is that they would pressure their ISP suppliers to offer IPv6 service – or would select ISPs on the basis

of IPv6 service being available to avoid doing conversions themselves – and that the implied economic forces would gradually spread IPv6 across the network.

7.2.2 Conversion at the core

The core-based conversion model suggests that larger ISPs may tire of the costs and problems of operating IPv4, and using IPv4 addresses, on their core networks. They would switch those networks over to IPv6, tunneling the IPv4 packets they receive from their customers. In this approach, it is reasonable to assume that those ISPs would notice that it was less expensive to offer IPv6 service to customers (over their IPv6 network), since they would not need to provide conversion or tunneling services, and that they would eventually reflect this in discounts to customers presenting and receiving IPv6 packets at their routers. They might also conclude it was in their interest to offer more attractive peering arrangements to other ISPs who were able to handle IPv6 traffic. As the cost differentials increased, customers would be motivated to convert.

The economic forces, and the relatively small number of core/ backbone/ “tier 1” ISPs probably make the tipping properties of core-based models better than those of edge-based ones.

7.2.3 Adoption by large “islands”

An additional model was not anticipated in the early guesses about conversion scenarios but may turn out to be a major driver of IPv6. If one is anticipating putting TCP/IP into a major new area – geographical, enterprise, or technological – for the first time, it may make sense to go to IPv6 initially, rather than installing IPv4 and having to undergo conversion costs. The decision to use IPv6 for 3G wireless and active work on it in China and Japan seem to reflect these “island” factors as well as assumptions about increasing difficulties (or costs) in obtaining large quantities of IPv4 address space.

8 Potential roadblocks and solutions

Many of the issues covered in this section have been discussed, in context, above, so it may be considered a review of key issues.

8.1 *Economical*

Conversion to IPv6 will not be either painless or inexpensive. Remaining with IPv4, especially where networks are expected to grow or new applications will be deployed, will not be painless or inexpensive either. Economic incentives to convert will occur if the costs of IPv4 connectivity rise significantly above those of IPv6 connectivity (most likely with the “core-first” model discussed in section 7.2.2), if a conversion is forced by technical factors, or if the costs of obtaining required address space become excessive. The economic risk of delay is that, while differential connectivity costs are likely to rise slowly, if at all, the others drivers could occur with catastrophic suddenness, requiring a conversion with little planning time once the catastrophe becomes apparent. Such conversions are almost always significantly more expensive

than those that can be planned well in advance and carried out on a convenient schedule.

As long as the alternatives to IPv6, including the use of NATs to avoid needing more IPv4 address space and the use of a combination of address translation and protocol conversion to continue use of IPv4 even after significant fractions of the network have converted, are acceptable, the temptation to avoid converting will be very high. On first analysis, waiting would seem to have no consequences, the IPv4 environment will be more familiar than the IPv6 one, and, for at least the near future, all applications will support IPv4 and some may not support IPv6. Viewed this way, one might reasonably avoid converting until forced to do so. On the other hand, especially if the network continues to grow (both locally and globally), conversion costs later will almost certainly exceed those of an early conversion.

Finally, as discussed above, large enterprises or countries that are considering getting on the Internet for the first time, or that are considering very high growth rates in the foreseeable future, may want to consider early adoption of IPv6. Doing so would avoid the costs of adopting and deploying an IPv4 system and then having to convert or upgrade it, either to move to IPv6 along with the rest of the world or to obtain additional address space after serious scarcity sets in.

8.2 *Technical*

The major technical obstacle to conversions to IPv6 today is the lack of tested, production-quality software at all levels. Backbone network and very large enterprise conversions are impeded because, while router vendors are shipping IPv6 stacks, their high-performance, hardware-assisted packet forwarding implementations are still in the future. Similarly, for end-user systems and local servers, while experimental software is available, Microsoft is not yet shipping an IPv6 stack that they consider production-quality (it is generally assumed that, when a production-quality stack is available it will be supplied only for Windows XP and not for earlier versions). The situation for Linux and most versions of UNIX is somewhat better – good-quality IPv6 stacks are available and have been in use for most of these platforms for a few years. Conversions of applications software has, naturally, lagged the stacks with, again, the UNIX-derived platforms generally being further along than the Microsoft environments and more popular and obvious applications (e.g., web browsers) taking the lead over more obscure ones.

These obstacles are likely to be overcome with time, although it is important to note that vendors and users are, at present, trapped in a deadly embrace: from the user standpoint, the obstacle to conversion is a lack of vendor-supplied software and facilities even to the degree needed to test, build pilot implementations, and make firm deployment plans. From the vendor standpoint, those facilities are not being implemented and deployed because there is no strong user demand for them. Similar problems occur from the standpoint of ISPs: customers who perceive themselves as unable to construct realistic pilot and test programs do not place demands on their ISPs for native and optimized IPv6 services; the absence of ISPs aggressively offering IPv6 services persuades many potential customers that they should defer planning conversions.

Enterprises and organizations that do decide to convert do, fortunately, find that tunneling systems to link to other IPv6 environments across the IPv4 backbones are readily available (and that there have been many years of experience with them) and that protocol conversion and address mapping gateways are readily available.

There also are some technical risks that impact the economic considerations. To give only two examples: if a network is using a NAT arrangement to avoid asking for additional IPv4 address space, and a new protocol is deployed (or an older one becomes of interest) that will not work behind a NAT, an immediate network reconfiguration may be required, either to accommodate at least some hosts “outside” the NAT arrangement, to obtain and install a larger IPv4 address pool, or to go to IPv6. If new protocols come along that are of interest and that are, for whatever reason, supported only under IPv6, the enterprise will be faced with a choice between IPv6 support and access to the capabilities of that protocol. And, finally, as operating system, software vendors, and router manufacturers shift to IPv6, history leads us to predict rising maintenance and support costs for IPv4 versions of their systems.

All of these factors suggest that, while there may or may not be advantages of being an early adopter of IPv6, once significant conversions begin, one does not want to be the last, or even in the last group, to convert – at least to an environment that can successfully interwork with IPv6 and, more likely, to IPv6 itself.

8.3 *Policy/political*

With the exception of the “islands” discussed in section 7.2.3, most of the policy issues associated with IPv6 track the economic and technical issues discussed in the previous two sections. The islands are an interesting technical and policy challenge: is it better to ahead of much of the world on IPv6 adoption or to adopt a more conservative plan, presumably with the expectation of later conversion? And, if one takes a conservative approach, will sufficient IPv4 address space be available to permit fulfilling reasonable projections about required addresses? The latter consideration was presumably important in driving third-generation wireless planning toward IPv6; similar considerations would presumably apply to any other applications – systems requiring fairly large per-residence home networks with peer-to-peer capability would be one possible example – that could be predicted to require very large amounts of globally-accessible address space. Finally, there are important features of IPv6 that this document does not address²³. If some of them become important to some future application, there will be an important question as to whether to implement that application less efficiently with IPv4 (assuming that is possible) or to implement it for IPv6 alone.

9 Summary: Thinking about conversion

Available globally addressable space on the IPv4 Internet is decreasing. It is difficult to measure the rate of decrease, and even one or two very large-scale applications that

²³ See the RFCs cited in notes 2 and 3.

require global address space could exhaust most of the space that can be allocated without disruption to existing users and applications. Even an expansion of dedicated Internet connections in China or India to the density now seen in several Scandinavian countries, if done using IPv4, could substantially exhaust the remaining IPv4 address space.

This implies, in turn, that, in thinking about conversion, there are two conditions under which IPv6 can safely be ignored:

- (i) One concludes that, despite network growth and the technical factors and risks associated with the Internet's running out of space or having to request an allocation after free space has become very rare, that an address space crisis will never occur or is sufficiently far in the future that one can safely ignore that possibility.
- (ii) One concludes that one's supplying ISP will not make a conversion to IPv6 and subsequently either discontinue IPv4 service or offer it only at a prohibitive price. One could, of course, reach this conclusion after negotiating a long-term agreement with that ISP or assuming that ISPs will always be available that offer IPv4 service (and, if necessary, conversion and translation services to reach remote systems that are IPv6-only).

If neither of those assumptions applies, then the important questions are not about whether to convert to IPv6, but how and when to make the conversion. For the enterprise or other LANs, the major choices of "how" are whether to continue with an IPv4 network and address translation and protocol conversion gateways or to move directly to IPv6 with either a dual-stack environment on individual machines or address translation/ protocol conversion gateway capability from IPv6 to IPv4. The "when" question is even more complicated. In general, it will be less disruptive to convert at one's own convenience, rather than being forced into it by inability to reach and communicate with critical customers or correspondents or by pressing external economic considerations (unacceptable costs of additional IPv4 space or high ISP prices for IPv4 services). Also in general, if one waits, software may become more stable and per-station conversion costs may decrease somewhat as experience accumulates. On the other hand, if the patterns of growth in the Internet and Internet usage within organizations continue, the number of stations will rise over time, probably enough to more than wipe out any per-station savings.