



Illustration by David Cowles

IP addresses: MAXED OUT

**AS THE WORLDWIDE
ONLINE POPULATION
EXPLODES AND AS NEW
APPLICATIONS EMERGE,
WE FACE AN INEVITABLE
IP-ADDRESS SHORTAGE.**

I OFTEN WONDER WHAT A STATE-OF-THE-ART PC might be like if it didn't have to also support the first MS-DOS program ever written. What if the companies that drive the Wintel world ripped up the blueprint? What if, with no constraints, their engineers created the best possible PC with no limits on processor options and no need to support legacy software? Sure, I would have to buy a

new system, replace all my software, and perhaps learn a new way of doing things. But maybe I'd be more productive. Maybe my system would never fall short of resources. Maybe I wouldn't need 256 Mbytes of memory. Maybe the revolutionary change would leave the business world better off.

Think about that hypothetical scenario. The PC has served us well, but surely it's not the best or most elegant architecture we're capable of devising. What's your vote—retrofit or replace?

Now, what if we ask the same question about the Internet? The Internet is actually older than the PC. People have subjected the Internet to far more retrofits and face-lifts than the PC has endured. Many argue that the Net needs a major overhaul to support the future demands we'll place on it.

Like the PC, the Internet isn't facing an immediate crisis. Despite some naysayers' claims, the Net won't just break down and stop working—at least not any day soon. Undoubtedly, the average user doesn't

even realize that the Internet faces significant obstacles. But obstacles there are.

Topping the list is a sort of housing shortage. Limited to a 32-bit address space, the Internet has a finite number of possible IP addresses. As the online population worldwide catches up with Europe and North America, we will run short. And that's despite stopgap measures, such as Network Address Translation (NAT). Worse, new applications, such as mobile Net access and home networking, could vastly increase the world's appetite for addresses. The Internet also faces problems in security, routing, and installation complexity. Together, these problems limit the Internet's potential.

As the new millennium dawns, the Internet's brain trust must decide how to upgrade the Net's foundation: the venerable TCP/IP infrastructure. The powers that be can continue a 20-year process of retrofits, some of which have been effective and some of which have merely mitigated problems in the short term. Or, they can make a more revolutionary change and replace the current Internet Protocol Version 4 (IPV4) with Internet Protocol Version 6 (IPV6), which has been in development for much of the past decade.

IPV6 offers a few significant advantages that simply cannot be spliced onto IPV4. Most significantly, IPV6 uses a 128-bit address space. Other examples include the end-to-end security feature IP Security (IPSec) and automatic or plug-and-play installation using node discovery and dynamic address assignment.

Other IPV6 features have been tacked onto IPV4 with varying degrees of success. For example, take IP Multicast, which allows a single server to broadcast an audio or a video stream to many nodes. Multicast is an absolute requirement if the Internet is to support a growing base of streaming content. Though IP Multicast originated as part of the IPV6 project, nearly half of the IPV4-based Internet can support it.

Then there's quality of service (QoS), the ability to give time-sensitive data, such as voice and video, an express ride to its destination. IPV6's creators included QoS from the start. A packet-flow label in the address header ensures reliable delivery of streaming data, enabling applications such as voice over IP (VoIP).

IPV4 retrofits provide QoS but force routers and other equipment to open and examine the actual data payload of packets. Moreover, companies making routers and switches support different QoS techniques (see sidebar "IPV6 resources" with the online version of this article at www.commvorgemag.com/commverge/issues/2000/200001/01cs.asp).

Sounds like a no-brainer; let's upgrade to IPV6. Unfortunately, logistic and technical roadblocks obstruct the path. For starters, the Internet is a huge entity. No one owns it or even has the power to control its direction. The Internet Engineering Task Force (IETF) shepherds the development of new technologies but can't dictate how they're deployed by the thousands of organizations that control little pieces of the Internet.

A complete upgrade to IPIETV6 would be what engineers wryly refer to as nontrivial. Every router, every PC, and every server owned by Internet-service-providers (ISPs), corporations, universities, governments, and individuals would have to be replaced or upgraded. Typically, upgrades happen gradually, the way that QoS and Multicast are seeping in. Individuals and businesses upgrade or replace network equipment and hosts one at a time. But IPV6 isn't backward-compatible with IPV4, and no smooth transition path exists.

The barriers to an easy transition have some industry insiders wondering whether IPV6 will ever be adopted,

whereas others insist that the question is when and not whether. The two sides disagree on whether IPV6 brings enough value to warrant the expense and time the transition would entail. Many believe IPV6 simply isn't revolutionary enough to justify the massive upheaval.

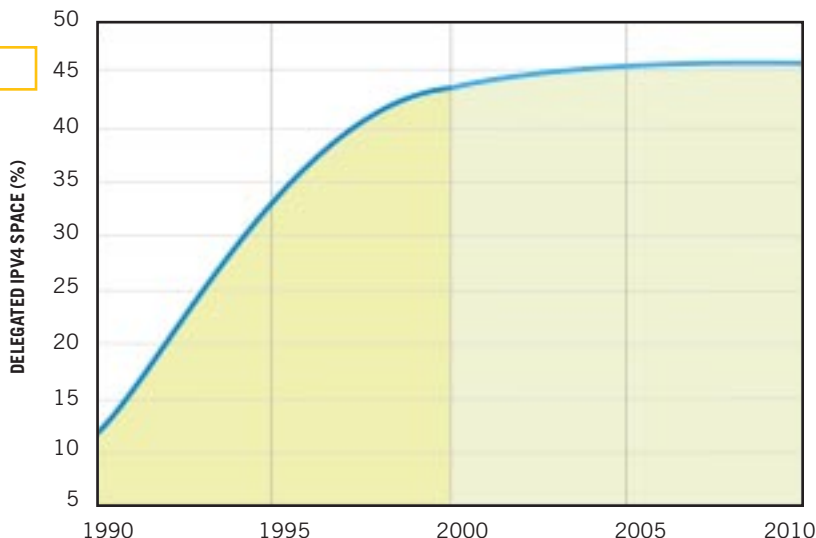
ADDRESS DEPLETION

Although IPV6 offers more than a larger address space, that feature appears to be the one most likely to force a change. IPV4 can theoretically address more than 4 billion nodes. IPV6 supports a number of addresses greater than 3 followed by 38 zeroes. Experts say that every person on earth could have a personal IPV6 network, which could in turn have a number of unique addresses represented by 20 digits.

But is that amount overkill? According to whom you ask, we've already occupied a quarter to a half of the available IPV4 address space. The lower estimates count routable addresses or, in other words, active addresses. The higher estimates count addresses that have been allocated by the various agencies in charge of allocation. Many of these addresses are essentially sitting on the shelf at ISPs, corporations, and other organizations. **Figure 1** depicts the ongoing address-depletion trend.

The current depletion curve seems to indicate that any serious address shortage is far in the future. The curve ramped sharply upward during much of the

Figure 1



This address-depletion curve shows that IP-address demand is leveling off, but technologies such as mobile Internet access and home networking could quickly turn up demand (courtesy Top Layer Networks).

1990s but started to flatten out toward the end of the decade.

So, are we really in danger? Ask Noel Chiappa, who invented the multiprotocol router while at the Massachusetts Institute of Technology (Cambridge, MA) and has worked on numerous IETF initiatives. “We’ve already run out of addresses,” he answers. “Just look at how many people are deploying NAT” he says. Chiappa, currently an independent network-architecture researcher, believes that IPV6 offers insufficient benefits to merit a change. He prefers a more radical approach to a next-generation network.

But the point in Chiappa’s comment is that the industry has already taken measures to prevent, or at least considerably delay, the address well from running dry. In the early 1990s, the Internet transitioned to a more granular routing architecture called CIDR (Classless Inter-Domain Routing—often pronounced “cider”). Before the development of—CIDR, all subnets were assigned based on 1- to 3-byte boundaries, which meant that small and large organizations often received far more addresses than they required (see “Connected: An Internet Encyclopedia” at www.freesoft.org/CIE/index.htm).

To further stretch the address space, the computer world has adopted NAT, in which a single host (with a single public IP address) conceals a host of other hosts. The hosts behind the gateway use private IP addresses; it doesn’t matter whether they duplicate addresses that exist elsewhere. NAT finds use in widely ranging scenarios, from connecting the private networks of huge corporations to allowing multiple PCs to share a home Internet connection. In its purest form, NAT transparently alters address headers as packets pass through the gateway.

PANACEA OR POISON?

But NAT is far from a panacea. In fact, it’s almost pure evil in the eyes of some Internet purists. NAT violates one of the Net’s basic tenets: that the source and destination hosts control communications between themselves. With NAT, all data passes through an intermediary, and



Global Impact: Fred Baker heads the Internet Engineering Task Force.

NAT critics see that single point of failure as potential trouble. NAT can also cause huge logistical problems. For example, if you try to merge two private networks that previously existed behind NAT gateways, you may have to contend with duplicate addresses.

Web access and e-mail get along fine with NAT, but other applications fare less well. Many newer and innovative applications actually place addresses inside the data payload. Moreover, some applications, such as security and videoconferencing, use encrypted data. A NAT box can easily prevent such applications from working properly. In fact, IPV6’s IPsec can’t be retrofitted to IPV4, due to the lack of a direct end-to-end connection.

“The NAT problem is unbounded,” says Fred Baker, IETF Chairman and Cisco fellow.

NAT might not even be up to the task of preserving address space. Should mo-

bile Internet access take off, it’s unclear whether NAT would be effective—especially when cellular hand-offs take place as subscribers roam among different cells and service providers. Initiatives in development now dynamically reuse IP addresses as mobile devices connect and disconnect from the Internet. But as Baker points out, “You will still need a huge amount of addresses just to support all of the subscribers connected at any specific instance.”

Moreover, NAT can’t head off escalating address requirements arising from the move to broadband connections. With dial-up modems, ISPs need IP addresses for only a fraction of their subscriber bases, because the ISPs dynamically assign an address when a subscriber dials in. With an always-on broadband connection, such as a digital subscriber line or cable modem, each subscriber needs at least one—and potentially more—unique address. If home networking and automation take off, each subscriber may need a block of addresses (see sidebar “The connected home” with the online version of this article at www.commvergemag.com/commverge/issues/2000/200001/01cs.asp).

Finally, new applications can eat up globs of address space. For example, IP Multicast needs a unique public address for every audio/video stream. The Multicast standard reserves 268,435,200 IP addresses. Who knows what application will emerge next month or next year to devour even more?


WHERE TO GO

Most Internet experts agree that NAT is, at best, an interim solution to a future address shortage. But even the staunch proponents of IPV6 can’t reach a consensus on how to proceed.

The next big retrofit to IPV4 could be realm-specific IP (RSIP, typically pronounced “are-sip”). According to your perspective, RSIP could be the start of a transition to IPV6 or a way to delay a transition while a more radical alternative is developed.

RSIP solves many of NAT’s problems (with, of course, a few adverse side effects) but does no more than NAT to handle address

Figure 2



Most observers believe that IPV6 will work its way inward starting from clouds on the edges of the IPV4-based Internet, eventually replacing the Net’s IPV4 core.

depletion. Whereas NAT is transparent to the host, RSIP requires that each host participate in the process by which many hosts share one public IP address. In doing so, RSIP restores the ability to make end-to-end connections, thereby enabling applications such as IPSec and real-time video delivery.

In an RSIP scenario, the host makes a request to the RSIP gateway and retrieves instructions on how to address and format its outgoing packets. Still in development, RSIP could be deployed over the next year. RSIP only works when the host behind the RSIP gateway instigates a session, but some say it can be extended to allow an external host to initiate a session.

Mike Borella, senior architect at 3Com and one of the leaders in the IETF RSIP development effort, believes the technology can lead the transition to IPv6. Most transition strategies include a mixture of three concepts:

- Dual stacks, in which hosts or routers have both IPv6 and IPv4 protocol stacks and can work on either type of packet.
- Translation, which works like a NAT gateway, except that the gateway translates an IPv6 address to an IPv4 address.
- Tunneling, in which IPv6 packets are encapsulated within IPv4 packets, so that IPv4 networks can connect IPv6 clouds.

The predominant theory concludes that a transition will move inward from the individual hosts (the PCs and servers) on the “edges” of the Internet (**Figure 2**). Adding dual stacks in those hosts will be relatively painless because it’s a software upgrade. You can already buy IPv6 stacks for most Unix boxes. Microsoft, however, is dragging its feet on providing IPv6 support for Windows; it will first arrive in Windows 2000 (Windows NT). But at some point, the IPv6 stack upgrade will become essentially free, allowing organizations to create IPv6 clouds all around the edges of the Internet.

Unfortunately, upgrading routers to support IPv6 is less straightforward. Many routers use hardware-based schemes to speed packet processing. These will have to be replaced—a process many in the IETF expect to take more than a decade.

Of the ways to connect IPv6 clouds to the IPv4 infrastructure, Borella thinks, the RSIP concept offers the best scenario.

“IPV6 OFFERS AN OPPORTUNITY TO VASTLY EXPAND THE ADDRESS BASE, TO ALLOW FOR MORE FLEXIBLE ADDRESS ASSIGNMENT POLICIES, MULTICASTING DESIGN, AND SECURITY FEATURES TO BECOME BASIC COMPONENTS OF THE EVOLVING INTERNET.”



A dual-stack RSIP gateway could support end-to-end connections for hosts in the IPv6 cloud, whether the remote host speaks IPv4 or IPv6. The scenario would still require the RSIP gateway to perform tunneling to link two IPv6 hosts across what for some time will remain an IPv4 Internet. And the RSIP gateway could integrate NAT functions as well, so that RSIP can be deployed across connected hosts one at a time.

A TOUGH SELL

Others, such as Chiappa, see NAT, RSIP, and other retrofits as ways to delay the transition to IPv6. Chiappa points out that ISPs and corporate MIS managers—whose support would speed a brick-by-brick transition to IPv6—have no incentive to make the change.

“If 5 or 10% of the managers upgrade to IPv6,” he asks, “would they be better off?” And he’s probably right in assuming the answer would be “No.” You couldn’t run IPv6-enabled applications in such a situation, except in cases in which you knew for sure that both ends of a connection had been upgraded.

Chiappa also points out that IPv6 doesn’t solve the biggest problem ISPs and MIS managers face. Routing-table entries have topped 70,000, and IPv6 doesn’t solve that problem, he claims. According to the IETF’s Baker, however, IPv6, does bring some benefits to routing (see the online version of this story at www.commvorgemag.com for Baker’s take on routing issues.) Chiappa believes technologies such as multiprotocol label switching (MPLS), which the IETF is working on, could provide the answer. With MPLS, routers at the edge of the In-

ternet attach tags or labels to packets, which allows routers on the major backbones to forward the packets without retrieving an address from a routing table. So far, MPLS has been seen mostly as a way to improve the operation of backbones, and companies such as AT&T and MCI WorldCom are already deploying it. But Chiappa thinks it could be an important piece of an alternative to IPv6.

Even Chiappa concedes, however, that a 32-bit address space is too small in the long term. Considering that IPv6 took more than five years to develop, most industry luminaries believe there’s no viable alternative to IPv6.

Borella from 3Com believes movement toward IPv6 will come suddenly; he just doesn’t know when. Perhaps, he suggests, an application such as Internet services for cell phones will explode overnight in a region such as Asia, jump-starting the process.

Advocates of IPv6 are also—finally—making an effort to promote the technology. In mid-1999, industry leaders formed the IPv6 Forum to educate everyone from MIS managers to end users. Vint Cerf, senior vice president for Internet architecture and technology at MCI WorldCom and a pioneer of the Internet, serves as honorary chairman of the group.

“IPv6 offers an opportunity to revisit design of the basics of Internet protocol,” Cerf says. “To vastly expand the address base, to allow for more flexible address-assignment policies, multicasting design, and security features to become basic components of the evolving Internet. IPv6 also offers features that facilitate the deployment of new qualities of service, which are needed to serve an expanding array of applications.”

From a more pragmatic perspective, Jim Bound, senior member of the technical staff at Compaq and co-chairman of the IPv6 Deployment Group claims, “IPv6 is the only solution to maintaining the principles of end-to-end networking as we approach the next millennium.” □

AUTHOR’S BIOGRAPHY

Maury Wright, who used to cover computers and multimedia for EDN, is now executive editor of CommVerge. You can reach him at 1-858-748-6785, fax 1-858-679-1861, e-mail maury-wright@home.com.