OPEN ISSUE
Deployment of IPv6

More than 15 years have elapsed since the shortage of IPv4 address space became serious enough to warrant proposals for a new version of IP. The original IPv6 specification is now more than 10 years old. IPv6-capable host operating systems are now widely available and the major router vendors offer varying degrees of support for IPv6 in their products. Yet the deployment of IPv6 in the Internet can only be described as embryonic. It is worth wondering when deployment is likely to begin in earnest, and what will cause it.

One reason why IPv6 has *not* been needed sooner is because of the extensive use of NAT (network address translation, described earlier in this chapter). As providers viewed IPv4 addresses as a scarce resource, they handed out fewer of them to their customers, or charged for the number of addresses used; customers responded by hiding many of their devices behind a NAT box and a single IPv4 address. For example, it is likely that most home networks with more than one IP-capable device have some sort of NAT in the network to conserve addresses. So one factor that might drive IPv6 deployment would be applications that don't work well with NAT. While client-server applications work reasonably well when the client's address is "hidden" behind a NAT box, peer-to-peer applications fare less well. Examples of applications that would work better without NAT and would therefore benefit from more liberal address allocation policies are multiplayer gaming and IP telephony.

Obtaining blocks of IPv4 addresses has been getting more difficult for years, and this is particularly noticeable in countries outside the United States. As the difficulty increases, the incentive for providers to start offering IPv6 addresses to their customers also rises. At the same time, for existing providers, offering IPv6 is a substantial additional cost, because they don't get to stop supporting IPv4 when they start to offer IPv6. This means, for example, that the size of a provider's routing tables can only increase initially, because they need to carry all the existing IPv4 prefixes plus new IPv6 prefixes.

At the moment, IPv6 deployment is happening primarily in research networks. A few service providers are starting to offer it (often with some incentive from national governments). It seems hard to imagine that the Internet can continue to grow indefinitely without IPv6 seeing some more significant deployments, but it also seems likely that the overwhelming majority of hosts and networks will be IPv4-only for several more years at least.

FURTHER READING

Not surprisingly, there have been countless papers written on various aspects of the Internet. Of these, we recommend two as must reading: The paper by Cerf and Kahn is the one that originally introduced the TCP/IP architecture and is worth reading just for its historical perspective; the paper by Bradner and Mankin gives an informative overview on how the rapidly growing Internet has stressed the scalability of the original architecture, ultimately resulting in the next generation IP. The paper by Paxson describes a study of how routers behave in the Internet. It also happens to be a good example of how researchers are now studying the dynamic behavior of the Internet. The final paper discusses multicast, presenting the approach to multicast originally used on the MBone.

- Cerf, V., and R. Kahn. "A Protocol for Packet Network Intercommunication." *IEEE Transactions on Communications* COM-22(5):637–648, May 1974.
- Bradner, S., and A. Mankin. "The Recommendation for the Next Generation IP Protocol." *Request for Comments* 1752, January 1995.
- Paxson, V. "End-to-End Routing Behavior in the Internet." *SIGCOMM '96*, pp. 25–38, August 1996.
- Deering, S., and D. Cheriton. "Multicast Routing in Datagram Internetworks and Extended LANs." *ACM Transactions on Computer Systems* 8(2):85–110, May 1990.

Beyond these papers, Perlman gives an excellent explanation of routing in an internet, including coverage of both bridges and routers [Per00]. Also, the book by Lynch and Rose gives general information on the scalability of the Internet [Cha93]. Some interesting experimental studies of the behavior of Internet routing are presented in Labovitz et al. [LAAJ00].

Many of the techniques and protocols developed to help the Internet scale are described in RFCs: Subnetting is described in Mogul and Postel [MP85], CIDR is described in Fuller and Li [FL06], RIP is defined in Hedrick [Hed88] and Mogul and Postel [MP94], OSPF is defined in Moy [Moy98], and BGP-4 is defined in Rekhter et al. [RLH06]. The OSPF specification, at over 200 pages, is one of the longer RFCs around, but also contains an unusual wealth of detail about how to implement a protocol. A collection of RFCs related to IPv6 can be found in Bradner and Mankin [BM95], and the most recent IPv6 spec is by Deering and Hinden [DH98]. The reasons to avoid IP fragmentation are examined in Kent and Mogul [KM87] and the Path MTU discovery technique is described in Mogul and Deering [MD90]. Protocol Independent Multicast (PIM) is described in Deering et al. [DEF+96] and [EFH+98]. MSDP is described in