Network Address Translation
and the Private Internet Exchange

Given the potential proliferation of network address translation devices, it is not clear that IPng will secure sufficient following to attain market viability.[1]

Internet Address Depletion

The combination of explosive growth in TCP/IP networking and the long-standing practice of assigning globally unique IP addresses to all hosts on TCP/IP networks has resulted in rapid depletion of the available IP address space. Since a unique address is required for each host connected to the global Internet, this presents a serious problem for new enterprise connections. The *IP: next generation* (IPng) area of the Internet Engineering Task Force (IETF) is currently considering proposals for a long-term solution. However, the process of selecting the new standard, managing the transition, and finally achieving ubiquitous implementation will take several years, if indeed it can be accomplished at all. In the meantime, three primary strategies have emerged for maximizing the longevity of the current IP standard.

Address Allocation Guidelines and Private Internets

The traditional method of Internet address allocation split the address space into three classes of networks based on the number of hosts within them. For convenience, address classes were divided on 8-bit boundaries, allowing roughly 250 hosts on a Class C network, 64,000 on a Class B, and 16 million on a Class A. Unfortunately, this lack of granularity does not reflect the realities of enterprise networking. Many organizations have networks which fall somewhere between the Class C and B magnitudes. A network manager with 4,000 hosts, for example, faces the dilemma of using 16 Class C registrations or 1/16 of a Class B.

Because of the underutilization of address space within assigned Class B addresses, they are now nearly impossible to get. The Internet Registry will assign blocks of multiple Class C addresses to applicants who do not meet the IR requirements for a Class B allocation, which include a minimum of 4,096 hosts and the submission of a detailed network plan.

> The restrictions in allocation of Class B network numbers may cause some organizations to expend additional resources to utilize multiple Class C numbers. This is unfortunate, but inevitable if we implement strategies to control the assignment of Class B addresses. The intent of these guidelines is to balance these costs for the greater good of the Internet.[2]

Organizations have historically been assigned globally unique IP network addresses regardless of their intent to connect their private network to the Internet. Even those which do join the Internet usually only allow Internet access to a small percentage of the hosts on their stub domain. The ratio of hosts with direct Internet access to hosts without such connectivity is typically between 1:1,000 and 1:10,000 in large corporate networks.[3] The use of registered, globally unique IP addresses for such large numbers of hosts which don't need them has further exacerbated the address depletion problem.

Recently, the Internet Assigned Numbers Authority reserved three blocks of the address space for use by private networks: one Class A, 16 Class B, and 255 Class C network numbers. These addresses may be used on the enterprise LAN for hosts that will never have direct IP connectivity with external hosts. But the hosts within an enterprise network fall into three categories, not two: those which *sometimes* require Internet connectivity, as well as those which *always* and *never* need it.

Categories of Connectivity	
Level	*Examples*
Always	Email, FTP, World-Wide Web servers
Sometimes	User's workstation or PC
Never	Secure hosts, corporate database servers

The private network addressing scenario laid out in RFC 1597 relegates the connectivity needs of this middle category of hosts to "application layer relays", otherwise known as proxy servers.[4] But demand from end users for the direct Internet connectivity they need to run World-Wide Web browsers on their desktops is reaching a fever pitch, and the deployment of proxy servers adds another layer of complexity (and potential maintenance headaches) to the network.

Classless Inter-Domain Routing (CIDR)

CIDR, as described in a series of Internet RFCs† , is primarily aimed at increasing the efficiency (and reducing the size) of the Internet routing tables. This is being accomplished by a policy of allocating IP addresses in a way which allows routing information for multiple networks to be aggregated into a single routing table entry. Internet service providers are now being assigned contiguous blocks of the Class C address space, which are in turn reallocated to their new customers. The incorporation of variable-length netmask information into the routing protocols makes it possible for these multiple Class C networks to be served by a single routing table entry on the Internet. The designation of this mechanism as classless comes from the fact that it enables routing at intermediate levels between the traditional 8-bit boundaries of IP network classes.

One unfortunate side-effect of Classless Inter-Domain Routing is that, in order to maximize its effectiveness, existing domains may need to be renumbered, incurring a high administrative cost in the domains affected.

Network Address Translation (NAT)

The third (and most easily deployable) strategy for alleviating IP address depletion is Network Address Translation.[5] NAT is based on the concept of address reuse by private networks, and operates by mapping the reusable IP addresses of the stub domain to the globally unique ones required for communication with hosts on other networks. It would be difficult indeed to take full advantage of reusable addresses on a private network without employing NAT functionality.

It is also unlikely that many network managers will voluntarily incur the expense of renumbering their networks, as will eventually be necessary for full deployment of CIDR. The insertion of a Network Address Translator at the Internet connection point makes this a one-step operation, eliminating the need to visit each host on the corporate LAN to change its IP address.

† RFC 1467, RFC 1481, RFC 1517, RFC 1518, RFC 1519, and RFC 1520

Network Address Translation dovetails with both reusable addressing and CIDR, simplifying or eliminating many of the obstacles associated with the deployment of these initiatives. But NAT also provides simple solutions for a number of other network management problems.

Private Internet Exchange (PIX)

Currently, the only commercially available implementation of Network Address Translation is the Private Internet Exchange from Network Translation, Inc., which also incorporates features extending its functionality beyond the generic NAT device described in RFC 1631.

The Private Internet Exchange comes in a rack-mountable package and is equipped with two ethernet ports. In a typical installation, the inside (local) port is connected to the private network and the outside (global) port connects the PIX to the DMZ segment where the Internet router resides. Configuration is accomplished with the familiar ifconfig and route commands for each network interface. A global command specifies the virtual network number to which private host IP addresses will be mapped.

Dyanamic Address Allocation

Mapping between local and global addresses is done dynamically. When a host on the inside network initiates a connection to the outside world, PIX assigns it a globally unique IP number from a pool of available addresses. After a user-configurable timeout period during which there has been no activity on that connection, the translation table entry is removed, freeing the slot for use by another connection.

As mentioned previously, the number of hosts needing Internet connectivity from inside the corporate firewall is generally a very small percentage of the domain. But even fewer of them will require access simultaneously. Dynamic address allocation as implemented in PIX efficiently leverages a relatively small registered address space to serve the Internet connectivity needs of a much larger user population. In the course of expanding the private network, Internet access is available to the new hosts without reconfiguration.

Adaptive Application Security

Dynamic address translation is only enabled for connections intitiated from the inside network, and is port-specific. The translation for an outbound HTTP connection from a World-Wide Web client, for example, would only forward packets from the the external Web server which were destined for port 80 of the client machine. In the case of FTP, which uses an ephemeral port for its data connection, PIX takes note of the port number passively opened by the client's request and will only allow inbound FTP data for sessions which were initiated from inside the private network. Thus, the Private Internet Exchange provides the functionality of a proxy server without the extra administrative overhead and without the need for special client software.

Address translations may also be hard-wired for specific hosts on the inside network, such as SMTP or FTP servers, which need to accept connections from the outside world. Alternatively, these machines may be given a permanent registered IP address and placed on the DMZ network segment. Except for those internal hosts explicitly advertised, PIX conceals the architecture of the private network from the outside world.

Summary

The Private Internet Exchange offers a new degree of flexibility in network design. The private network may use any combination of address class and subnetting, or even take advantage of address reuse within the organization itself. Existing unregistered networks may be connected to the Internet in a few minutes, without having to change IP addresses on individual machines. Significant savings may also be realized from the reduction of administrative costs in the maintenance of PIX-connected networks.

Contact Information

Network Translation, Inc.
1901 Embarcadero Road
Palo Alto, CA 94303
Phone: (415) 494-NETS (6387)
Fax: (415) 424-9110
Email: info@translation.com

References

1. J. Curran, ''IPng White Paper on Market Viability,'' RFC 1669, BBN, August 1994.
2. E. Gerich, ''Guidelines for Management of IP Address Space,'' RFC 1466, Merit, May 1993.
3. E. Fleischman, ''A Large Corporate User's View of IPng,'' RFC 1687, Boeing Computer Services, August 1994.
4. Y. Reckhter, B. Moskowitz, D. Karrenberg, and G. de Groot, ''Address Allocation for Private Internets,'' RFC 1597, T.J. Watson Research Center, IBM. Corp., Chrysler Corp., RIPE NCC, March 1994.
5. K. Egevang and P. Francis, ''The IP Network Address Translator (NAT),'' RFC 1631, Cray Communications, NTT, May 1994.