The
Internet Protocol Journal

December 2000 Volume 3, Number 4

A Quarterly Technical Publication for Internet and Intranet Professionals

In This Issue

FROM THE EDITOR

Numerous technologies have been developed to protect or isolate corporate networks from the Internet at large. These solutions incorporate security, either end-to-end (IP security, or IPSec), or at the Internet/intranet border (firewalls). A third class of systems allows a range of IP addresses to be used internally in a corporate network, while preserving IP address consumption through the use of a *single* public address. This latter class of device is called a *Network Address Translator* (NAT), and while many Internet engineers consider NATs to be "evil," they are nonetheless very popular. Combining IPSec, NATs, and firewalls can be quite challenging, however. In our first article Lisa Phifer explains the problem and offers some solutions.

Successful network design is the result of many factors. In addition to the basic building blocks of routers, switches and circuits, network planners must carefully consider how these elements are interconnected to form an overall system with as few single points of failure as possible. In our second article, Valdis Krebs looks at how lessons learned from social network analysis can be applied to the design of computer networks.

The current Internet grew out of several government-funded research efforts that began in the late 1960s. Today, basic technology development as well as research into new uses of computer networks continues in many research "testbeds" all over the world. Bob Aiken describes the past, present and future state of network research and research networks.

The online subscription system for this journal will be up and running in January at **www.cisco.com/ipj**. In addition to offering a subscription form, the system will allow you to select delivery options, update your mailing and e-mail address, and much more. Please visit our Web site and give it a try. If you encounter any difficulties, please send your comments to **ipj@cisco.com**.

—Ole J. Jacobsen, Editor and Publisher
ole@cisco.com

The Trouble with NAT

by Lisa Phifer, Core Competence

Those who are implementing virtual private networks often ask whether it is possible to safely combine *IP Security* (IPSec) and *Network Address Translation* (NAT). Unfortunately, this is not a question with a simple "yes" or "no" answer. IPSec and NAT can be employed together in some configurations, but not in others. This article explores the issues and limitations associated with combing NAT and "NAT-sensitive" protocols like IPSec. It examines configurations that do not work, and explains why. It illustrates methods for using NAT and IPSec together, and discusses an emerging protocol that may someday prove more IPSec friendly.

This article builds upon "IP Security and NAT: Oil and Water?"[1] and "Realm-Specific IP for VPNs and Beyond"[2], works previously published by *ISP-Planet.*

What Is Network Address Translation?

NAT was originally developed as an interim solution to combat IPv4 address depletion by allowing globally registered IP addresses to be reused or shared by several hosts. The "classic" NAT defined by RFC 1631[3] maps IP addresses from one realm to another. Although it can be used to translate between any two address realms, NAT is most often used to map IPs from the nonroutable private address spaces defined by RFC 1918[4], shown below.

Class	Private Address Range
A	10.0.0.0 ... 10.255.255.255
B	172.16.0.0 ... 172.16.255.255
C	192.168.0.0 ... 192.168.255.255

These addresses were allocated for use by private networks that either do not require external access or require limited access to outside services. Enterprises can freely use these addresses to avoid obtaining registered public addresses. But, because private addresses can be used by many, individually within their own realm, they are nonroutable over a common infrastructure. When communication between a privately addressed host and a public network (like the Internet) is needed, address translation is required. This is where NAT comes in.

NAT routers (or NATificators) sit on the border between private and public networks, converting private addresses in each IP packet into legally registered public ones. They also provide transparent packet forwarding between addressing realms. The packet sender and receiver (should) remain unaware that NAT is taking place. Today, NAT is commonly supported by WAN access routers and firewalls—devices situated at the network edge.

NAT works by creating bindings between addresses. In the simplest case, a one-to-one mapping may be defined between public and private addresses. Known as static NAT, this can be accomplished by a straightforward, stateless implementation that transforms only the network part of the address, leaving the host part intact. The payload of the packet must also be considered during the translation process. The IP checksum must, of course, be recalculated. Because TCP checksums are computed from a pseudo-header containing source and destination IP address (prepended to the TCP payload), NAT must also regenerate the TCP checksum.

Figure 1: Static NAT

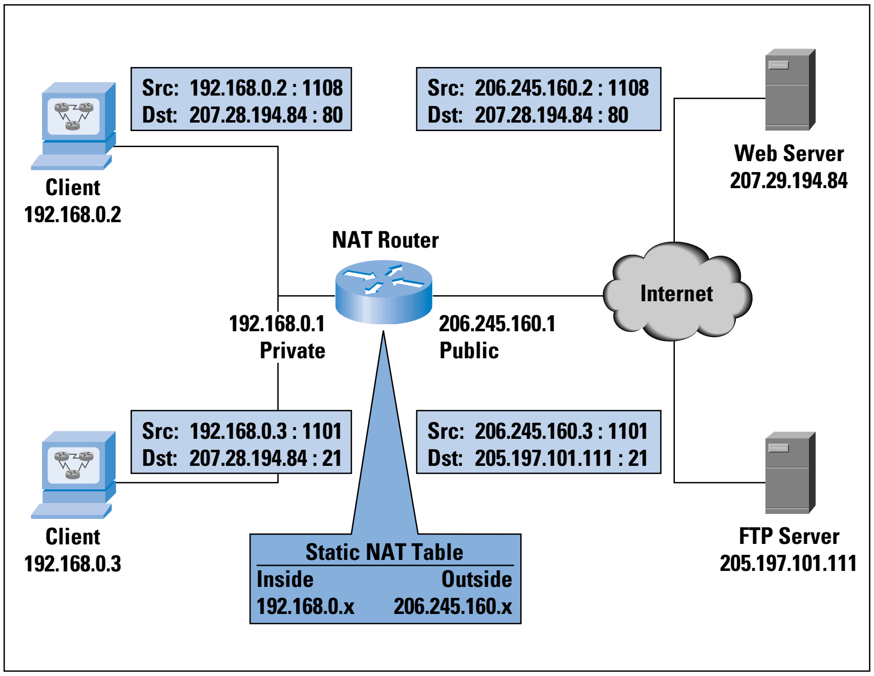

More often, a pool of public IP addresses is shared by an entire private IP subnet (dynamic NAT). Edge devices that run dynamic NAT create bindings "on the fly," building a NAT Table. Connections initiated by private hosts are assigned a public address from a pool. As long as the private host has an outgoing connection, it can be reached by incoming packets sent to this public address. After the connection is terminated (or a timeout is reached), the binding expires, and the address is returned to the pool for reuse. Dynamic NAT is more complex because state must be maintained, and connections must be rejected when the pool is exhausted. But, unlike static NAT, dynamic NAT enables address reuse, reducing the demand for legally registered public addresses.

Figure 2: Dynamic NAT

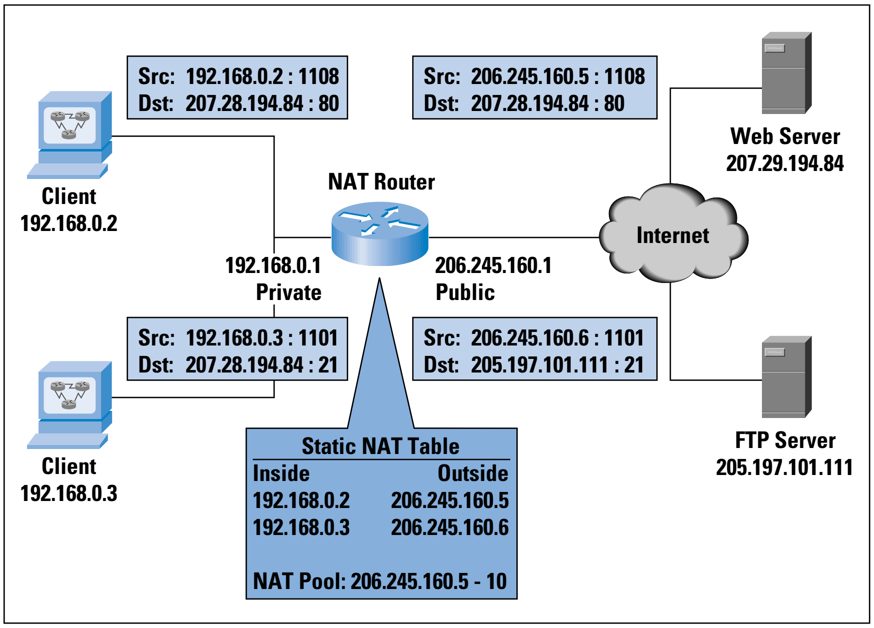

A variation of dynamic NAT known as *Network Address Port Translation* (NAPT) may be used to allow many hosts to share a single IP address by multiplexing streams differentiated by TCP/UDP port number. For example, suppose private hosts 192.168.0.2 and 192.168.0.3 both send packets from source port 1108. A NAPT router might translate these to a single public IP address 206.245.160.1 and two different source ports, say 61001 and 61002. Response traffic received for port 61001 is routed back to 192.168.0.2:1108, while port 61002 traffic is routed back to 192.168.0.3:1108.

Figure 3: NAPT

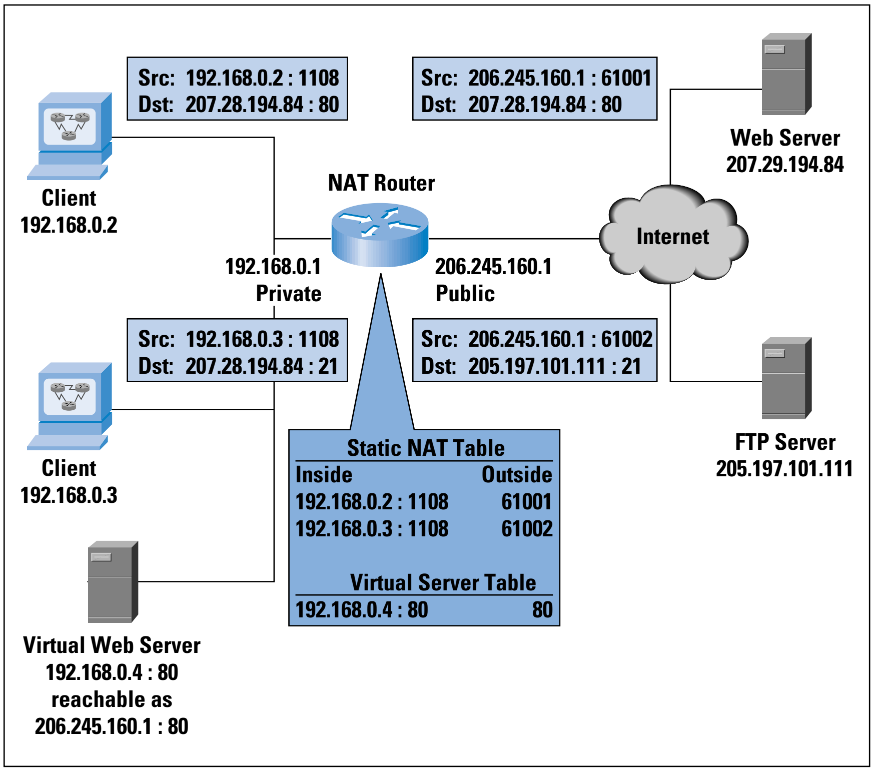

NAPT (masquerading) is commonly implemented on small Office/Home Office (SOHO) routers to enable shared Internet access for an entire LAN through a single public address. Because NAPT maps individual ports, it is not possible to "reverse map" incoming connections for other ports unless another table is configured. A virtual server table can make a server on a privately addressed DMZ reachable from the Internet via the public address of the NAPT router (one server per port). This is really a limited form of static NAT, applied to incoming requests.

In some cases, static NAT, dynamic NAT, NAPT, and even bidirectional NAT or NAPT may be used together. For example, an enterprise may locate public Web servers outside of the firewall, on a DMZ, while placing a mail server and clients on the private inside network, behind a NAT-ing firewall. Furthermore, suppose there are applications within the private network that periodically connect to the Internet for long periods of time. In this case:

- Web servers can be reached from the Internet without NAT, because they live in public address space.
- *Simple Mail Transfer Protocol* (SMTP) sent to the private mail server from the Internet requires incoming translation. Because this server must be continuously accessible through a public address associated with its *Domain Name System* (DNS) entry, the mail server requires static mapping (either a limited-purpose virtual server table or static NAT).
- For most clients, public address sharing is usually practical through dynamically acquired addresses (either dynamic NAT with a correctly sized address pool, or NAPT).
- Applications that hold onto dynamically acquired addresses for long periods could exhaust a dynamic NAT address pool and block access by other clients. To prevent this, long-running applications may use NAPT because it enables higher concurrency (thousands of port mappings per IP address).

Where is NAT used today? Outbound NAT is commonly employed by multihost residential users, teleworkers, and small businesses that share a single public IP for outbound traffic while blocking inbound session requests. In other words, small LANs connected via ISDN, *Digital Subscriber Line* (DSL), or cable modem.

Bidirectional static NAT/NAPT combinations are typically used by enterprises that host services behind a masquerading firewall. NAT can also be employed by enterprises wishing to insulate themselves from *Internet Service Provider* (ISP) address changes, or by those wanting to obscure private network topology for security reasons.

NAT-Sensitive Protocols

Our need to conserve IPv4 addresses has prompted many to overlook the inherent limitations of NAT, recognized in RFC 1631 but deemed acceptable for a short-term solution.

As noted previously, NAT regenerates TCP checksums. This, of course, requires the TCP header containing the checksum to be visible (that is, not encrypted). If only the TCP payload is encrypted and immutable between the application source and destination (for instance, *Secure Shell Protocol* [SSH], *Secure Sockets Layer* [SSL]), then the checksum in the TCP header can be recalculated without a visible TCP payload. But if the TCP header is encrypted (for instance, IPSec transport mode), the TCP checksum field in the TCP header cannot be modified.

Furthermore, many application protocols carry IP addresses in an application-level protocol. In such cases, an *Application-Level Gateway* (ALG) is needed to complete the translation. For example:

- Many *Internet Control Message Protocol* (ICMP) packets (for instance, “Destination Unreachable”) carry embedded IP packets in ICMP payload. These require both address translation and checksum regeneration.
- A *File Transfer Protocol* (FTP) ALG is needed to rewrite IP addresses carried by FTP PORT and PASV control commands. In the IP header, these addresses are fixed-length words. Unfortunately, in the FTP protocol, these IP addresses are carried as human-readable, variable-length strings; rewriting can change the length of the TCP segment. If the segment is shortened, it can be padded. If the segment is lengthened, SEQ and ACK numbers must be transformed for the duration of the connection.
- Protocols like H.323 use multiple TCP connections or UDP streams to form “session bundles.” If all connections in the bundle originate from the same end system, an ALG may be avoided. But H.323 presents other challenges, including ephemeral ports and embedded, ASN.1-encoded IP addresses in application payload.
- *NetBIOS over TCP/IP* (NBT) can be challenging to translate correctly because packet-header information is placed in NetBIOS payload at inconsistent offsets, and many embedded IP addresses are exchanged during an NBT session. Fortunately, most companies do not let NBT beyond their firewall anyway.
- *Simple Network Management Protocol* (SNMP) packets also carry IP addresses that identify trap source and object instance. Perhaps more important, dynamic NAT makes it impossible to uniquely identify hosts by IP address; public addresses are transient and shared. Remote management of private hosts can thus be impeded by NAT.
- Obviously DNS, responsible for domain name/IP address mapping, is impacted by NAT. From simple query handling to zone transfers, a robust DNS ALG is defined by RFC 2694[9].

NAT-sensitive protocols such as Kerberos, X-Windows, remote shell, Session Initiation Protocol (SIP), and others are further described in the Internet Draft *"Protocol Complications with the IP Network Address Translation"*[12]. Another Internet Draft, *"NAT Friendly Application Design Guidelines"*[13], explains how new application protocols can integrate smoothly with NAT. But there are still cases where ALGs simply cannot "fix" packets modified by NAT.

Impact of NAT on IPSec

The IPSec *Authentication Header* (AH)[5] is an example. AH runs the entire IP packet, including invariant header fields such as source and destination IP address, through a message digest algorithm to produce a keyed hash. This hash is used by the recipient to authenticate the packet. If any field in the original IP packet is modified, authentication will fail and the recipient will discard the packet. AH is intended to prevent unauthorized modification, source spoofing, and man-in-the-middle attacks. But NAT, by definition, modifies IP packets. Therefore, AH + NAT simply cannot work.

Figure 4: NAT vs. AH (Transport Mode)

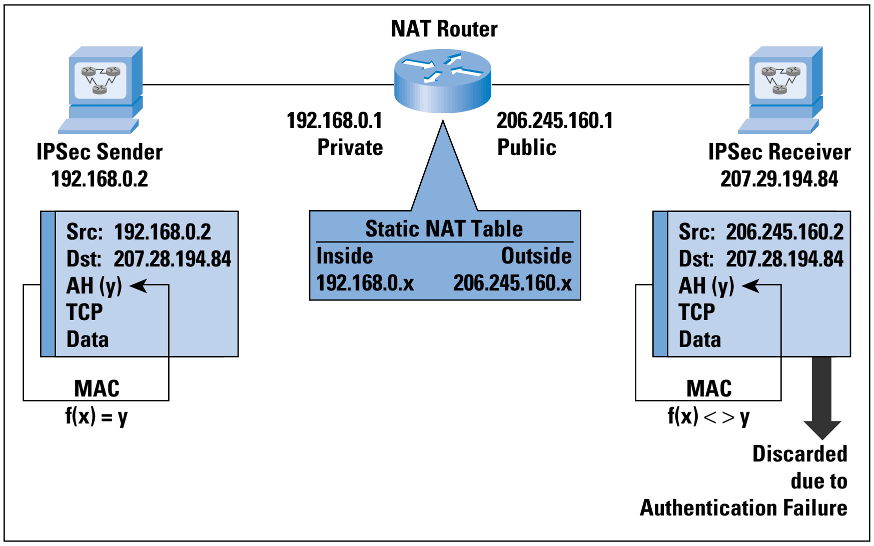

The IPSec *Encapsulating Security Payload* (ESP)[6] also employs a message digest algorithm for packet authentication. But, unlike AH, the hash created by ESP does not include the outer packet header fields. This solves one problem, but leaves others.

IPSec supports two "modes." Transport mode provides end-to-end security between hosts, while tunnel mode protects encapsulated IP packets between security gateways—for example, between two firewalls or between a roaming host and a remote access server. When TCP or UDP are involved—as they are in transport mode ESP—there is a catch-22. Because NAT modifies the TCP packet, NAT must also recalculate the checksum used to verify integrity. If NAT updates the TCP checksum, ESP authentication will fail. If NAT does not update the checksum (for example, payload encrypted), TCP verification will fail.

If the transport endpoint is under your control, you might be able to turn off checksum verification. In other words, ESP can pass through NAT in tunnel mode, or in transport mode with TCP checksums disabled or ignored by the receiver.

Figure 5: NAT vs. ESP (Transport Mode)

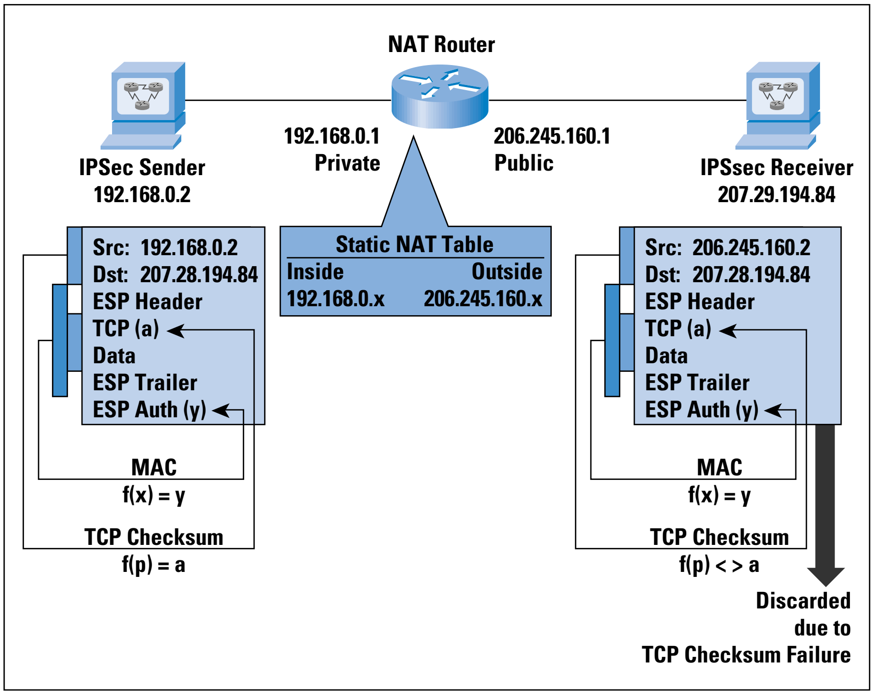

Static NAT Table

Inside	Outside
192.168.0.x	206.245.160.x

If we stick to ESP in tunnel mode or turn off checksums, there's still another obstacle: the *Internet Key Exchange* (IKE)[7]. IPSec-based *Virtual Private Networks* (VPNs) use IKE to automate security association setup and authenticate endpoints. The most basic and common method of authentication in use today is preshared key. Unfortunately, this method depends upon the source IP address of the packet. If NAT is inserted between endpoints, the outer source IP address will be translated into the address of the NAT router, and no longer identify the originating security gateway. To avoid this problem, it is possible to use another IKE "main mode" and "quick mode" identifier (for example, user ID or fully qualified domain name).

A further problem may occur after a *Security Association* (SA) has been up for awhile. When the SA expires, one security gateway will send a rekey request to the other. If the SA was initiated from the well-known IKE port UDP/500, that port is used as the destination for the rekey request. If more than one security gateway lies behind a NAPT router, how can the incoming rekey be directed to the right private IP address? Rekeys can be made to work by "floating" the IKE port so that each gateway is addressable through a unique port number, allowing incoming requests to be demultiplexed by the NAPT router.

Figure 6: NAT vs. IKE Rekey

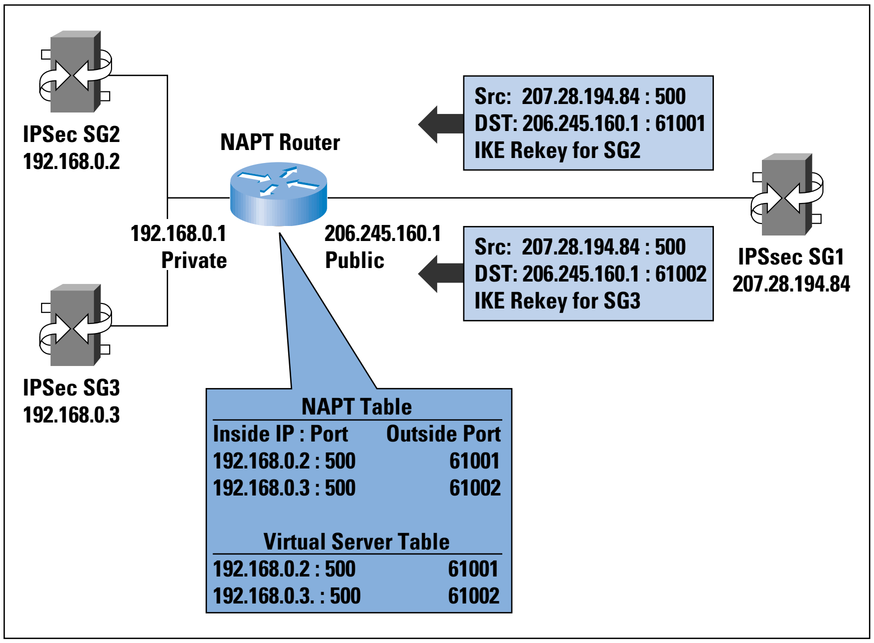

NAPT Table

Inside IP : Port	Outside Port
192.168.0.2 : 500	61001
192.168.0.3 : 500	61002

Virtual Server Table

192.168.0.2 : 500	61001
192.168.0.3. : 500	61002

At this point, two things should be clear: (1) it is possible to find a "flavor" of IPSec that will run through NAT, but (2) one must do so with great care and attention to detail. Recent Internet Drafts[12] [14] have recorded these problems for further consideration, and RFC 2709[10] describes a security model for running tunnel-mode IPSec through NAT.

One Solution: Avoid the Problem

By far the easiest way to combine IPSec and NAT is to completely avoid these problems by locating IPSec endpoints in public address space. That is, NAT before IPSec; don't perform IPSec before NAT. This can be accomplished in two ways:

- Perform NAT on a device located behind your IPSec security gateway; or
- Use an IPSec device that also performs NAT.

Figure 7: Combining IPSec and NAT

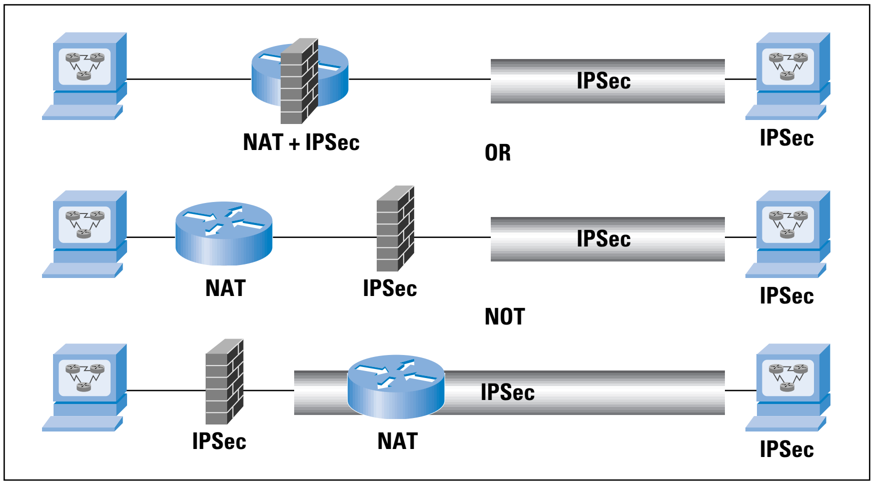

Many routers, firewalls, security gateways, and Internet appliances implement IPSec and NAT in the same box. These products perform outbound address translation before applying security policies; the order is reversed for inbound packets. A typical "any-to-any" security policy is easily specified with such a product. Granular policies can be a bit more difficult because filters are often based on IP address, and care must be taken to avoid overlapping filters.

If you cannot avoid translating IPSec-protected traffic midstream, limit use of IPSec to tunnel-mode ESP and design security policies with care. If you simply cannot NAT before IPSec or require transport-mode ESP, there may still be hope. The *Internet Engineering Task Force* (IETF) is now defining *Realm-Specific IP* (RSIP), an alternative that may someday prove kinder to IPSec.

What Is RSIP?

RSIP[16] leases public IP addresses and ports to RSIP hosts located in private addressing realms. Unlike NAT, RSIP does not operate in stealth mode and does not translate addresses on the fly. Instead, RSIP allows hosts to directly participate concurrently in several addressing realms. Although RSIP does require host awareness, it avoids violating the end-to-end nature of the Internet. With RSIP, IP payload flows from source to destination without modifications that cripple IPSec AH and many other NAT-sensitive protocols.

RSIP gateways are multihomed devices that straddle two or more addressing realms, just as NAT-capable firewalls and routers do today. When an RSIP-savvy host wants to communicate beyond its own private network, it registers with an RSIP gateway. The RSIP gateway allocates a unique public IP address (or a shared public IP address and a unique set of TCP/UDP ports) and binds the private address of the RSIP host to this public address. The RSIP host uses this public source address to send packets to public destinations until its lease expires or is renewed.

But the RSIP host cannot send a publicly addressed packet as-is; it must first get the packet to the RSIP gateway. To do this, the host wraps the original packet inside a privately addressed outer packet. This "encapsulation" can be accomplished using any standard tunneling protocol: IP-in-IP, the *Generic Routing Encapsulation* (GRE), or the *Layer 2 Tunneling Protocol* (L2TP). Upon receipt, the RSIP gateway strips off the outer packet and forwards the original packet across the public network, toward its final destination.

Figure 8: RSIP

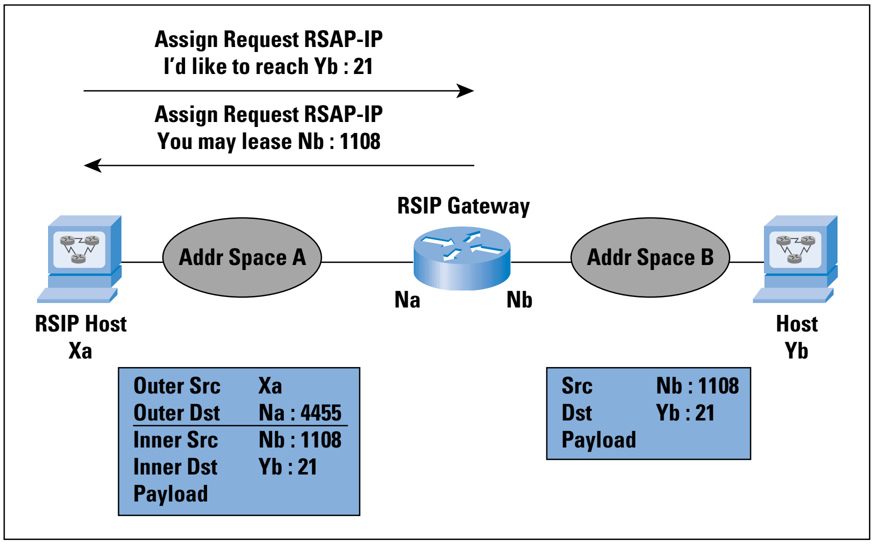

For simplicity, we talk about RSIP linking one private network to the public Internet, but RSIP can also be used to relay traffic between several privately addressed networks. An RSIP host can lease several different addresses as needed to reach different destinations networks. We've also focused on outgoing traffic, but an RSIP host can ask the RSIP gateway to "listen" and relay incoming packets addressed to a public IP and port.

Combining RSIP and IPSec

At first glance, RSIP sounds like a promising way for hosts to share public addresses while avoiding the pitfalls associated with applying NAT to IPSec traffic. But it turns out that RSIP extensions are needed to accommodate end-to-end IPSec[17].

Basic RSIP relies on unique port numbers to demultiplex arriving packets, but IPSec ESP encrypts port numbers. When several RSIP hosts use the same RSIP gateway to relay ESP, another discriminator is needed. Fortunately, every IPSec packet carries a unique *Security Parameters Index* (SPI), assigned during security association setup. Unfortunately, the SPI is guaranteed unique only for the responder. To enable demultiplexing, the tuple (SPI, protocol [AH or ESP], destination IP address) must also be unique at the initiating RSIP gateway.

A similar problem occurs during association setup with the IKE. IKE packets usually carry the well-known source port UDP/500. Using different source ports is the preferred solution, but if several RSIP hosts use the same RSIP gateway to relay IKE from port UDP/500, another discriminator is needed. Again, there is a convenient answer: every IKE packet carries the initiator cookie supplied in the first packet of an IKE session. The RSIP gateway can route IKE responses to the correct RSIP host using the tuple (initiator cookie, destination port [IKE], destination IP address). But rekeys may still be an issue.

To fix these problems, extensions have been proposed to allow RSIP hosts to register with an RSIP gateway for IPSec support, and allow hosts to request and receive unique SPI values along with leased IP addresses and ports.

Possible Applications for RSIP

RSIP specifications[16][17][18] are still at the Internet Draft stage. If and when RSIP matures, there may be a wide variety of applications:

- Residential power users and teleworkers with multihost LANs that share a single, publicly known IP address leased by an RSIP-enabled Internet appliance, DSL router, or cable modem;
- Small-to-midsize enterprise customers with dozens or hundreds of hosts, sharing a small pool of public IPs leased by an RSIP-enabled WAN access router or firewall;
- Multidwelling units (apartments, shared office buildings) with many private LANs, sharing public Internet access through an RSIP-enabled device;
- Hospitality networks (airports, hotels) where roaming hosts briefly lease the public IP(s) shared by the entire network;
- Remote access concentrators that use RSIP to lease private IP(s) to roaming corporate users that access the Internet via dynamically assigned public addresses; and
- Wireless devices (cell phones, personal digital assistants [PDAs]) that lease public IP(s) for "sticky sessions" that persist even when the mobile device moves from one location to another, updating its local access IP.

These scenarios, and the relationship of RSIP to IP multicast and differentiated services, are more fully explored in the RSIP framework[18].

Conclusion

Although NAT can be combined with IPSec and other NAT-sensitive protocols in certain scenarios, NAT tampers with end-to-end message integrity. RSIP—or whatever RSIP evolves into—may someday prove to be a better address-sharing solution for protocols that are adversely impacted by NAT. If RSIP fails to mature, another solution may be developed to broaden use of NAT with IPSec. Alternatives now under discussion within the IETF include UDP encapsulation and changes to IKE itself[14][15].

Despite its origin as a short-term solution, NAT is unlikely to disappear in the very near future. Until it does, understanding the relationship between NAT and IPSec and alternatives for safe combined deployment will remain an important aspect of VPN design.

References

[1] Phifer, L., “IP Security and NAT: Oil and Water?” *ISP-Planet,* June 15, 2000.
`http://www.isp-planet.com/technology/nat_ipsec.html`

[2] Phifer, L., “Realm-Specific IP for VPNs and Beyond,” *ISP-Planet,* June 23, 2000.
`http://www.isp-planet.com/technology/rsip.html`

[3] Egevang, K. and Francis, P., “The IP Network Address Translator (NAT),” RFC 1631, May 1994.

[4] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G.J., and Lear, E., “Address Allocation for Private Internets,” RFC 1918, February 1996.

[5] Kent, S. and Atkinson, R., “IP Authentication Header,” RFC 2402, November 1998.

[6] Kent, S. and Atkinson, R., “IP Encapsulating Security Payload (ESP),” RFC 2406, November 1998.

[7] Harkins, D. and Carrel, D., “The Internet Key Exchange (IKE),” RFC 2409, November 1998.

[8] Srisuresh, P. and Holdrege, M., “IP Network Address Translator (NAT) Terminology and Considerations,” RFC 2663, August 1999.

[9] Srisuresh, P., Tsirtsis, G., Akkiraju, P. and Heffernan, A., “DNS Extensions to Network Address Translators (DNS_ALG),” RFC 2694, September 1999.

[10] Srisuresh, P., “Security Model with Tunnel-Mode IPSec for NAT Domains,” RFC 2709, October 1999.

[11] Tsirtsis, G. and Srisuresh, P., “Network Address Translation-Protocol Translation (NAT-PT),” RFC 2766, February 2000.

[12] Srisuresh, P. and Holdrege, M., “Protocol Complications with the IP Network Address Translator,” Internet Draft, Work in Progress, July 2000.

[13] Senie, D., “NAT Friendly Application Design Guidelines,” Internet Draft, Work in Progress, July 2000.

[14] Aboba, B., “NAT and IPSec,” Internet Draft, Work in Progress, July 2000.

[15] Stenberg, M., Paavolainan, S., Ylonen, T., and Kivinen, T., “IPSec NAT-Traversal,” Internet Draft, Work in Progress, July 2000.

[16] Borella, M. and Lo, J., “Realm-Specific IP: Protocol Specification,” Internet Draft, Work in Progress, March 2000.

[17] Montenegro, G. and Borella, M., “RSIP Support for End-to-End IPSec,” Internet Draft, Work in Progress, March 2000.

[18] Borella, M., Lo, J., Grabelsky, D., and Montenegro, G., “Realm-Specific IP: Framework,” Internet Draft, Work in Progress, March 2000.

LISA PHIFER is vice president of Core Competence, Inc. **(`www.corecom.com`)**, a consulting firm specializing in Internet, network management, and security technologies. She earned her Master’s Degree in Computer Science from Villanova University. A Bellcore award recipient for her work in ATM network operations, Lisa has been involved in the design and deployment of networking protocols for over 18 years. She represented Bellcore and Unisys in several industry-standards organizations, and has participated in The Internet Security Conference (TISC) since its inception. Lisa consults, teaches, and writes about a variety of technologies, including caching, load balancing, DSL, ISDN, IPSec, PKI, OSS, and VPNs. Her monthly column on virtual private networking is published by *ISP-Planet.* E-mail: **`lisa@corecom.com`**

The Social Life of Routers
Applying Knowledge of Human Networks to the Design of Computer Networks

by Valdis Krebs

We often forget that computer networks are put in place to support human networks—person-to-person exchanges of information, knowledge, ideas, opinions, insights, and advice. This article looks at a technology that was developed to map and measure human networks—social network analysis—and applies some of its principles and algorithms to designing computer networks. And as we see more peer-to-peer (P2P) models of computer-based networks, the P2P metrics in human network analysis become even more applicable.

Social network analysts look at complex human systems as an interconnected system of nodes (people and groups) and ties (relationships and flows)—much like an internetwork of routers and links. Human networks are often unplanned, emergent systems. Their growth is sporadic and self-organizing[1]. Network ties end up being unevenly distributed, with some areas of the network having a high density of links and other areas of the network sparsely connected. These are called "small world networks"[2]. Computer networks often end up with similar patterns of connections—dense interconnectivity within subnetworks, and sparser connections uniting subnetworks into a larger internetwork.

Social network researchers and consultants focus on *geodesics*—shortest paths in the network. Many of today's social network algorithms are based on a branch of mathematics called *graph theory*. Social network scientists have concentrated their work, and therefore their algorithms, in the following areas:

- Individual node centrality within a larger network—network dependency and load upon individual routers
- Overall path distribution—good connectivity without excessive routing tables
- Improving communication flow within and between groups—designing better topologies
- Network patterns surrounding ego networks—strategies for analyzing and manipulating individual router connections
- Analyzing information flow behavior of client organization—how computer networks can support human networks

One of the methods used to understand networks and their participants is to evaluate the location of actors in the network. Measuring the network location is finding the *centrality* of a node[3]. All network measures discussed here are based on geodesics—the shortest path between any two nodes. We will look at a social network, called the *kite network,* that effectively shows the distinction between the three most popular centrality measures—the ABCs—*A*ctivity, *B*etweenness, and Closeness.

This model[4] was first developed by David Krackhardt, a leading researcher in social networks.

Activity

Figure 1 shows a simple social network. A link between a pair of nodes depicts a bidirectional information flow or knowledge exchange between two individuals. Social network researchers measure network activity for a node by using the concept of *degrees*—the number of direct connections a node has.

In this human network, Diane has the most direct connections in the network, making hers the most active node in the network with the highest degree count. Common wisdom in personal networks is "the more connections, the better." This is not always so. What really matters is where those connections *lead to*—and how they connect the otherwise unconnected![5] Here Diane has connections only to others in her immediate cluster—her clique. She connects only those who are already connected to each other—does she have too many redundant links?

Figure 1: Human Network

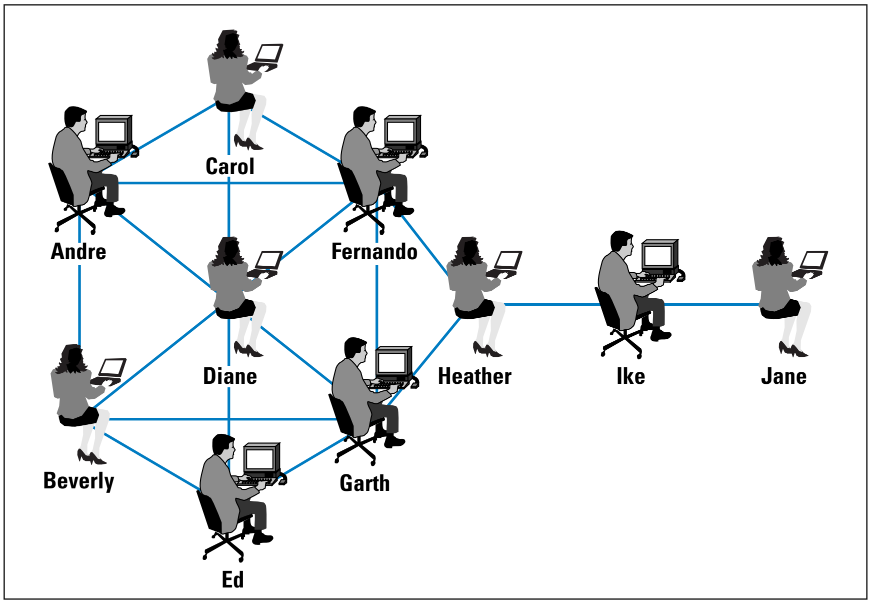

Betweenness

While Diane has many direct ties, Heather has few direct connections—fewer than the average in the network. Yet, in may ways, she has one of the best locations in the network—she is a boundary spanner and plays the role of broker. She is *between* two important constituencies, in a role similar to that of a border router. The good news is that she plays a powerful role in the network, the bad news is that she is a single point of failure. Without her, Ike and Jane would be cut off from information and knowledge in Diane's cluster.

Closeness

Fernando and Garth have fewer connections than Diane, yet the pattern of their ties allow them to *access* all the nodes in the network more quickly than anyone else. They have the shortest paths to all others—they are *close* to everyone else. Maximizing closeness between *all* routers improves updating and minimizes hop counts. Maximizing the closeness of only one or a few routers leads to counterproductive results, as we will examine below.

Their position demonstrates that when it comes to network connections, quality beats out quantity. Location, location, location—the golden rule of real estate also works in networks. In real estate it is geography—your physical neighborhood. In networks, it is your virtual location determined by your network connections—your network neighborhood.

Network Centralization

Individual network centralities provide insight into the individual's location in the network. The relationship between the centralities of all nodes can reveal much about the overall network structure. A very centralized network is dominated by one or a few very central nodes. If these nodes are removed or damaged, the network quickly fragments into unconnected subnetworks. Highly central nodes can become critical points of failure. A network with a low centralization score is not dominated by one or a few nodes—such a network has no single points of failure. It is resilient in the face of many local failures. Many nodes or links can fail while allowing the remaining nodes to still reach each other over new paths.

Average Path Length in Network

The shorter the path, the fewer hops/steps it takes to go from one node to another. In human networks, short paths imply quicker communication with less distortion. In computer networks, the signal degradation and delay is usually not an issue. Nonetheless, a network with many short paths connecting all nodes will be more efficient in passing data and reconfiguring after a topology change.

Average Path Length is strongly correlated with Closeness throughout the network. As the closeness of all nodes to each other improves (average closeness), the average path length in the network also improves.

Internetwork Topology

In the recent network design book, *Advanced IP Network Design*[6], the authors define a well-designed topology as the basis of a well-behaved and stable network. They further propose that "three competing goals must be balanced for good network design":

- Reducing hop count
- Reducing available paths
- Increasing the number of failures the network can withstand

Our social network algorithms can assist in measuring and meeting all three goals.

- Reducing the hop count infers minimizing the average path length throughout the network—maximize the closeness of all nodes to each other.
- Reducing the available paths leads to minimizing the number of geodesics throughout the network.
- Increasing the number of failures a network can withstand focuses on minimizing the centralization of the whole network.

On the following pages we examine various network topologies and evaluate them using social network measures while remembering these three competing goals of network design.

The models we examine do *not* cover hierarchical structures—with Core, Distribution, and Access layers—found in networks of hundreds or thousands of routers. We examine flat, nonhierarchical topologies such as those found in smaller internetworks, area subnetworks, or within core backbones. The topologies we model are the most commonly used—Star, Ring, Full Mesh, and Partial Mesh. We compute the social network measures on each of the topologies and discuss how the various measures help us meet the competing goals discussed above.

Star Topology

The Star topology, shown in Figure 2, has many advantages—but one glaring fault. The advantages include ease of management and configuration for the network administrators. For the Star, the three competing goals delineate as follows:

- *Reducing hop count:* The short average path length (1.75) throughout the network meets this goal well. Any router can *reach* any other router in two steps or less.
- *Reducing available paths:* The fact that there are a minimum number of possible available paths (56) to reach all other nodes—will not overload the routing tables, nor cause delays during routing table updates. It takes only seven bidirectional links to create the available paths.

- *Reducing network failures:* The network fails miserably if Router A goes down. Also, any link failure isolates the attached router—there are no multiple paths to reach each router.

Router A is not only a single point of failure—it is also a potential bottleneck—it will likely become overburdened with packet flows and routing updates as more routers are added in the star structure.

Router A receives the top score (1.000) in Activity, Betweenness, and Closeness. As a result, the network is very centralized around Router A from the perspective of all measures.

Figure 2: Routers in Star Topology

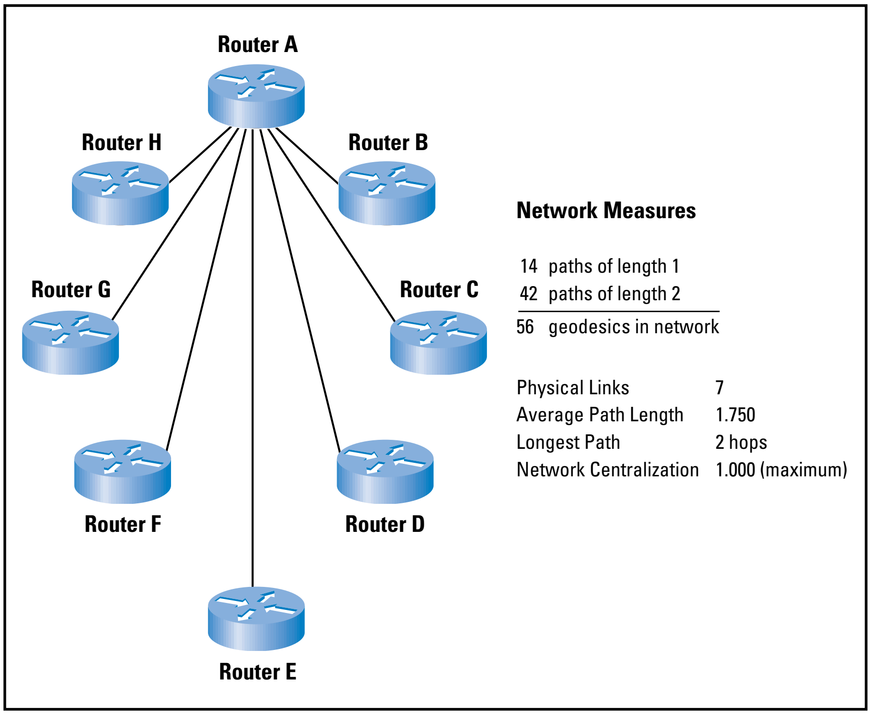

Ring Topology

The Ring topology, shown in Figure 3, is an improvement over the Star. It has some of the same advantages, but does not eliminate all of the drawbacks of the Star. The advantages include ease of management and configuration for the network administrators—adding another router is very simple. Unlike the Star topology, the Ring provides some redundancy and, therefore, eliminates the single point of failure—all nodes have an alternate path through which they can be reached. Yet it is still vulnerable to both link and router failures. For the Ring, the three competing goals delineate as follows:

- *Reducing hop count:* The average path length of 2.5 is quite long for a small network of eight nodes. Some routers (that is, A and E) require four steps to reach each other! Many ring physical layers hide this complexity from the IP layers in order to make those hops invisible to routing protocols.

- *Reducing available paths:* This configuration has more geodesics (64) than Star, yet not significantly more to overload the routing tables, nor cause delays during table updates.
- *Reducing network failures:* Even though network centralization is at the minimum (no node is more central than any other), this network reaches failure quickly because of its weak redundancy. The Ring topology can withstand one link failure or one router failure and still keep a contiguous network. Two simultaneous failures can cause unreachable segments because of the lack of redundancy.

Most modern ring technologies such as *Synchronous Optical Network* (SONET) or the Cisco *Dynamic Packet Transport Protocol* (DPT) add a measure of redundancy by running a dual ring that heals itself if a link gets cut. The network "wraps" to avoid the downed line and operates at lower speed. A two-hop path can become a six-hop path if a single link fails. This can cause network congestion if the original dual ring was being used for data in all directions.

Figure 3: Routers in Ring Topology

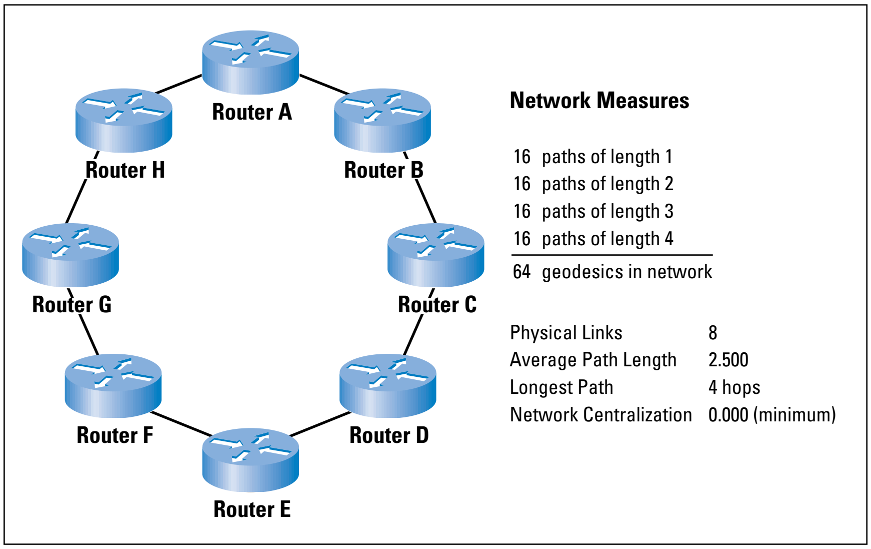

Full Mesh Topology

The Full Mesh topology has several big advantages and several faults. The advantages include short path length (one hop) to all other routers and maximum resilience to failure if links or routers start failing. The disadvantages revolve around the complexity created by this topology. For the Full Mesh, the three competing goals delineate as follows:

- *Reducing hop count:* The shortest path length possible is attained for all routes—all nodes can reach each other in one hop.
- *Reducing available paths:* There are a minimum number of possible available paths (56) to reach all other nodes. The routing entries will not overload the routing tables, nor cause delays during routing table updates.

- *Reducing network failures:* The network is not dependent upon any single node (network centralization = 0.000). This configuration represents the most robust topology available—chances are very slim that the number of failures necessary to fragment the network will actually occur within the same time period.

The disadvantages of the Full Mesh topology all focus on one glaring fault—there are too many physical links. If the routers are far apart, the link costs can quickly become prohibitively expensive because adding routers creates a geometrical explosion in links required—soon the routers do not have enough ports to support this topology. Administering the system and keeping an up-to-date topology map becomes more and more complex as routers are added. The network in Figure 4 has 28 two-way links. Double the routers, in a full mesh topology, and the link count increases by a factor greater than 4.

Figure 4: Routers in Full Mesh Topology

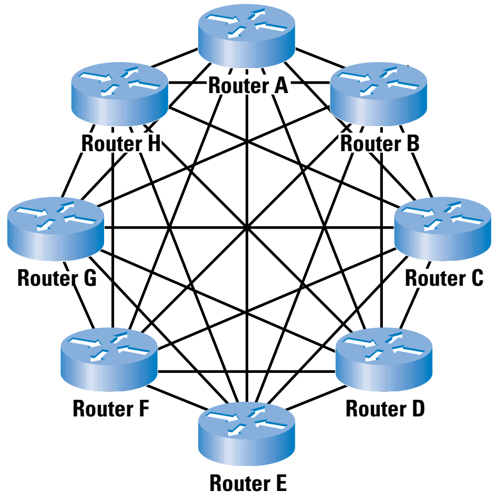

Network Measures

56 paths of length 1
56 geodesics in network

Physical Links	28
Average Path Length	1.000
Longest Path	1 hop
Network Centralization	0.000 (minimum)

Partial Mesh Topology

The Partial Mesh topology is quite different. It is the most difficult to build—there is no simple rule to follow (rule for Star: connect everyone to Router A; rule for Full Mesh: connect everyone to everyone). If built incorrectly, the partial mesh layout can have many of the disadvantages of the former topologies without many of the benefits. If built correctly, the opposite is true—more advantages, fewer disadvantages.

Building a successful partial mesh topology is where the interactive use of our social network measures really comes into play. The design below evolved after several iterations. With every iteration the average path length dropped until it appeared to reach a plateau where no further changes lowered the hop count without noticeably increasing the number of physical links. For the Partial Mesh, the three competing goals delineate as follows:

- *Reducing hop count:* The short average path length (1.667) throughout the network meets this goal well. Any router can *reach* any other router in two steps or less. Path length is less than that for the Star and Ring topologies.
- *Reducing available paths:* The number of available paths in the network (72) is the highest among all topologies, though not significantly more than the Ring topology. As the number of nodes in a network increases, this could become a problem—the average path length vs. path count trade-off needs to be closely monitored.
- *Reducing network failures:* Network centralization (0.000) is the same as for the Full Mesh topology—no router, nor link, is more important than any other. As nodes or links are removed from this network, it does not fragment quickly. Chances are slim that the number of failures necessary to fragment the network will actually occur within the same time period. Although we optimized our network centralization for this small "toy" network, we cannot expect this for most real networks. Yet, the goal remains to keep this metric as small as possible.

This topology in Figure 5 was built starting with a Ring topology—a simple architecture. A link was added and the network was remeasured. Was this structure better than the previous? If so, the current structure was kept and another link was added and the network was remeasured. This iterative process was continued until no further improvements happened after several changes. This process does not guarantee an *optimum* solution, yet it quickly converges on a *good* solution—even large networks improve quickly with just a few added links.

Figure 5: Routers in Partial Mesh Topology

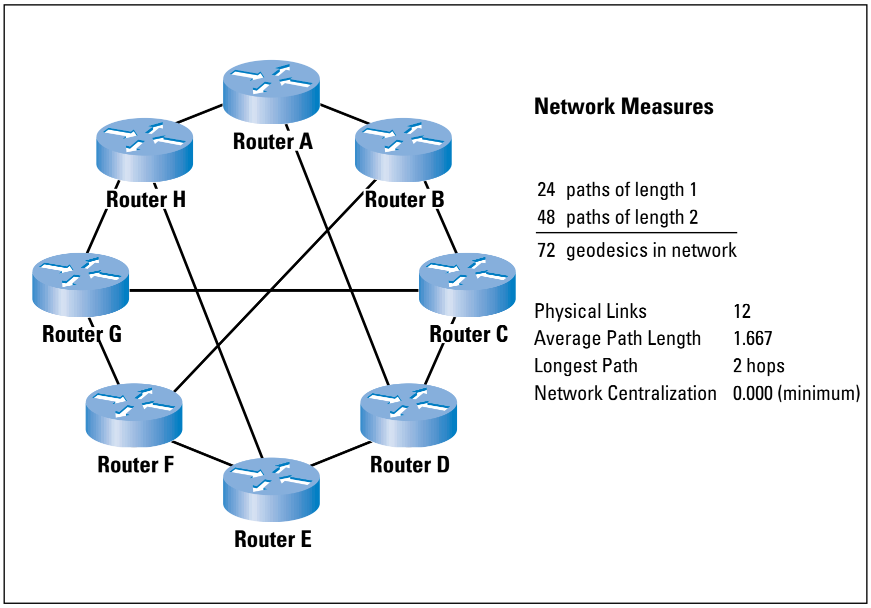

Network Measures

24 paths of length 1
48 paths of length 2
72 geodesics in network

Physical Links	12
Average Path Length	1.667
Longest Path	2 hops
Network Centralization	0.000 (minimum)

A quirky aspect of networks is that sometimes you can subtract by adding—add a link to a network and reduce the average path length. The opposite also works, sometimes. You can add by subtracting—remove a tie and watch the average hop count grow. Yet, you never know for certain what effect adding or removing a link will have—it is neither a linear nor a local phenomenon. The size and direction of these changes depend upon the existing topology of the network and the location of the added or removed tie. It is key to have a model that allows quick what-if calculations.

Let's experiment with removing random ties—a situation similar to links between routers failing. If we remove the link between Router A and Router H in Figure 5, the number of geodesics in the network increases from 72 to 76, and the average path length increases to 1.815. Yet, removing a different link, G to F, reduces the the number of geodesics in the network from 72 to 66, while the average path length increases only to 1.727. If we are concerned about too many paths in the network, we can remove another link, B to C. This further decreases the number of shortest paths to 60, while reducing physical links to 10. This is very near the 56 paths in the very efficient star topology. Whereas the star is very vulnerable because of its single point of failure, this partial mesh, with the two links removed, is still robust. While the number of geodesics drops, the average path length creeps up slightly to 1.80 with the removal of the second link. Figure 5 has no paths greater than two hops. With the two links (G to F, B to C) removed, we now have 8 geodesics of three hops, while at the same time 12 fewer geodesics to load into routing tables, and two fewer physical links. It is a constant trade-off.

NSFnet Backbone

The NSFnet Backbone network, shown in Figure 6, connected the supercomputing centers in the USA in 1989. It is a partial mesh design that functions as a real-life example to test our social network algorithms.

Figure 6: NSFnet in 1989

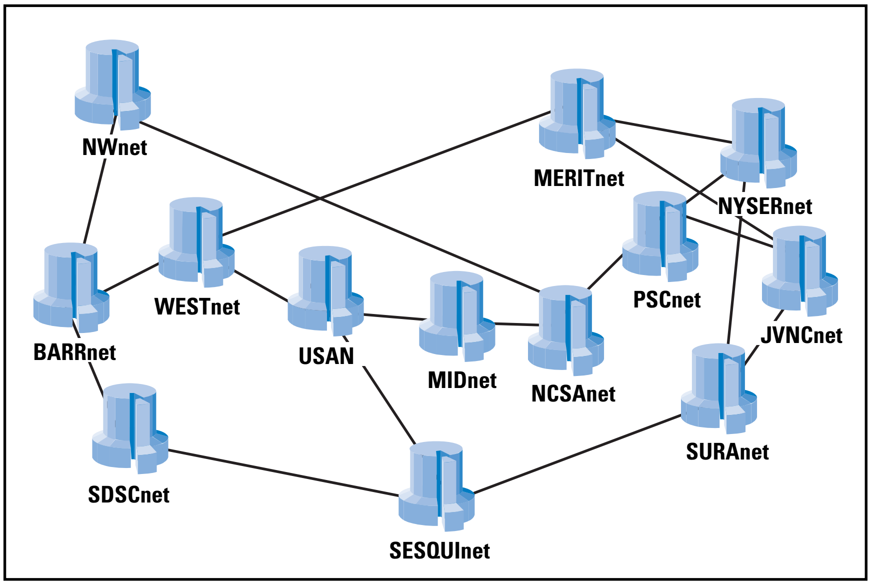

We remember our three competing goals for good internetwork design.

- Reducing hop count: average path length in steps/hops
- Reducing available paths: total geodesics in the network
- Increasing the number of failures the network can withstand: network centralization

What happens to these goals as we experience failures in the links or the nodes of the network? Table 1 shows the base metrics for Figure 6 and then shows what happens to the metrics, and our three goals, when five different failures occur.

Table 1: Possible Link and Node Failures

Scenario	Number of Geodesics in the Network	Network Centralization	Longest Path (hops)	Average Path Length (hops)
Original Design (Figure 6)	200	0.062	4	2.370
1) Node failure: NCSA	180	0.208	5	2.689
2) Node failure: MID	180	0.083	4	2.489
3) Node failure: JVNC	148	0.046	4	2.324
4) Link failure: NCSA–PSC	230	0.167	6	2.974
5) Link failure: USAN–MID	212	0.123	5	2.660
6) Link failure: MERIT–JVNC	192	0.069	4	2.458

The most damaging was link failure 4—the link failure between NCSA and PSC. This link is between two of the most central nodes in the network. If the flows between nodes are distributed somewhat evenly, then this link is one of the most traveled in the network.

The least damaging is node failure 3—the node failure at JVNC. In fact, this failure improved most metrics! By removing this node from the network, the number of network paths drops significantly, network centralization decreases, path length decreases slightly, and the longest path is still four hops.

The original NSFnet topology design is very efficient. I tried two different strategies to improve the network. The first strategy involved moving existing links to connect different pairs of routers. No obviously better topology was found by rearranging links among the routers. I was not able to find a better design that reduced both the number of geodesics and the average path length without significantly increasing the number of physical links in the network.

The second strategy is counter-intuitive, yet often networks respond well to this approach. It is the "subtracting by adding" approach described above. By adding new links in the right place in the network, we not only reduce the distance between nodes, we also decrease the number of geodesics in the network.

Because the NSFnet nodes had a maximum limit of three direct neighbors, I started connecting the nodes of Degree = 2. Options 1 through 3 show the various combinations and their effect on the total network. The improvements are minimal, yet each option offers specific strengths.

Option 2 offers more improvements than the others.

- The longest geodesic was reduced to three hops.
- The average path length was reduced throughout the network.
- The number of paths for the routers to remember was reduced slightly.
- Network centralization did not increase enough to noticeably affect the number of failures the network could withstand.

Table 2: Possible Network Improvements

Scenario	Number of Geodesics in the Network	Network Centralization	Longest Path (hops)	Average Path Length (hops)
Original Design (Figure 6)	200	0.062	4	2.370
Option 1 (add link: SDSC–MID)	202	0.071	4	2.287
Option 2 (add link: NW–DSC)	198	0.074	3	2.273
Option 3 (add link: NW–MID)	202	0.050	4	2.356

The improvement in Option 2 (add link: NW–SDSC) was actually implemented in the 1991 version of NSFnet—an excellent example of the "subtracting by adding" network dynamic. Networks are complex systems. How the network responds to change is based on the distribution and pattern of connections throughout the network.

Conclusion

In the real world we may not have the flexibility to experiment with our network model as we have with these examples. There will be more constraints. The information flows in your organization may require that specific pairs of routers have direct links—even if those connections would not be recommended by the algorithms we have been examining. Yet, when we have our "must-have" connections in place, we can experiment with the placement of the remaining connections using these social network metrics to indicate when we are getting close to a robust, yet efficient topology.

Given "initial conditions," social network methods can model our computer networks and suggest link changes[7] to form an effective topology that has a short average hop count, not too many paths, and just enough redundancy.

References

[1] Krebs V., "Visualizing Human Networks," *Release 1.0*, Esther Dyson's Monthly Report, February 1996.

[2] Watts D., Strogatz S., "Collective Dynamics of Small World Networks," *Nature,* 4 June 1998.

[3] Freeman L., "Centrality in Social Networks: A Conceptual Clarification," *Social Networks,* No. 1, 1979.

[4] Krackhardt D., "Assessing the Political Landscape: Structure, Cognition, and Power in Organizations," *Administrative Science Quarterly,* No. 35, 1990, page 351.

[5] Burt, Ronald S., *Structural Holes—The Social Structure of Competition,* ISBN 0674843711, Harvard University Press, 1992.

[6] Retana, A., Slice, D., White, R., *Advanced IP Network Design,* ISBN 1578700973, Cisco Press, 1999.

[7] Hagen G., Discussions with fellow network researcher, Guy Hagen, regarding combinatorial algorithms and models for recommending changes to improve the overall topology of a network.

VALDIS E. KREBS leads his own management consulting firm—orgnet.com He holds an undergraduate degree in Mathematics & Computer Science and a graduate degree in Human Resources. Since 1988 he has applied organizational network analysis to improve knowledge work within and between Fortune 500 firms such as IBM, Lucent, TRW, and supported consulting firms such as Ernst & Young, PricewaterhouseCoopers, and Booz-Allen-Hamilton. In addition to knowledge networks, he has applied these methodologies to mapping, measuring, and molding strategic alliances, communities of interest, emergent structures on the WWW, and internetworks. His work has been referenced in many publications, including the *Wall Street Journal, Entrepreneur, Training, PC Magazine, ZDNet, Corporate Leadership Council's Best Practices Reports, Knowledge Management, Across the Board, Business Week, HR Executive, Personnel Journal, FORTUNE,* and Esther Dyson's influential information industry newsletter, *Release 1.0.* He writes a regular column, "Working in the Connected World," for the *IHRIM Journal.* His Web site is at: **www.orgnet.com** and his e-mail is: **valdis@orgnet.com**

New Frontiers for Research Networks in the 21st Century

by Robert J. Aiken, Cisco Systems, Inc.

A famous philosopher, Yogi Berra, once said, "Prediction is hard. Especially the future."[1] In spite of this sage advice, we will still make an attempt at identifying the frontiers for research networks. By first examining and then extrapolating from the evolution and history of past research networks, we may be able to get an idea about the frontiers that face research networks in the future. One of the initial roles of the research network was to act as a testbed for network research on basic network protocols, mostly focusing on network Layers 1 through 4 (that is, the physical, data link, network, transport, and network management layers), but also including basic applications such as file transport and e-mail. During the early phases of the Internet, the commercial sector could not provide the network infrastructure sought by the research and education communities. Consequently, research networks evolved and provided backbone and regional network infrastructures that provided production-quality access to important research and education resources such as supercomputer centers and collaboratories[2]. Recent developments show that most research networks have moved away from being testbeds for network research and have evolved into production networks serving their research and education communities. It's time to make the next real evolutionary step with respect to research networks, and that is to shift our research focus toward maximizing the most critical of resources—*people*.

Given the growth and maturity of commercial service providers today, there may no longer be a pressing technical need for governments to continue to support pan-national backbone networks, or possibly even production-like national infrastructures, for Internet-savvy countries. Since commercially available *Virtual Private Networks* (VPNs) can now easily support many of the networked communities that previously required dedicated research networks, government and other supporting organizations can now support their research and education communities by providing the funding for backbone network services much as it does for telephony, office space, and computing capabilities; that is, as part of their research award. However, there may be valid social, political, and long-term economical reasons for continuing the support for such networks. For instance, a nation may decide that in order to ensure its economic survival in the future it wishes to accelerate the deployment and use of Internet technologies among its people, and thus the nation may decide to subsidize national research networks. In addition, it should be noted that VPNs often recreate the "walled" separation of communities, a scenario that was previously accomplished through the hard multiplexing of circuits.

But, in order to make technical advances in the e-economy, governments should now focus on supporting the evolution of intelligent and adaptable edge and access networks. These, in turn, will support the *Ubiquitous Computing* (UC) and persistent presence environments that will soon be an integral part of our future Internet-based economies.

The United States's recently expanded *National Science Foundation* (NSF)[3] research budget and the *Defense Advanced Projects Agency's* (DARPA's)[4] prior support of middleware research are good examples of moving in the right direction. The Netherland's Gigaport[5] project, which incorporates network and application research as well as an advanced technology access and backbone network infrastructure, is a good example of how visionary research networks are evolving.

Just as Internet technologies and network research have matured and evolved, so should the policies concerning the support of research networks. Policies need to be developed to again encourage basic network research and the development of new technologies. In addition, research networks need to encourage and accentuate new network capabilities in edge networks, on campus infrastructures, and in the end systems to support the humans in these new environments. This article focuses mainly on the future of research networks in e-developed nations; but, this is not to diminish the need or importance for e-developed nations to help encourage the same development in network-challenged nations.

Context and Definitions

Before delving into our discussion, we first need to define a few terms. These definitions will not only aid in our discussion, but may also help to highlight the role and function of various types of research networks. The most important terms to define are "network research" and "research network," both of which often get interchanged during discussions concerning policy, funding, and technology.

In this article, the term "network research" means long-term basic research on network protocols and technologies. The many types of network research can be categorized into three classes. The first category covers research on network transport infrastructure and generally includes research on the *Open System Interconnection* (OSI) Model Layers 1 through 4 (that is, the physical, data link, network, and transport layers) as well as research issues relating to the interconnection and peering of these layers and protocols. We will refer to this class of research as "transport services."

The second class consists of research covering what can nominally be referred to as "middleware"[6]. Middleware basically includes many of the services that were originally identified as network Layers 4 through 6. Layer 4 is included because of the need for interfaces to the network layer (sockets, TCP, and so on).

In addition, it nominally includes some components, such as e-mail gateways or directory services, which are normally thought of as being network applications, but which have subcomponents that may also be included in middleware. Given that the definition of middleware is far from an exact science, we shall say that middleware depends on the existence of the network transport services and supports applications.

The third area covers research on the real applications (for example, e-commerce, education, health care, and so on), network interfaces, network applications (for example, e-mail, Web, file transfer, and so on), and the use of networks and middleware in a distributed heterogeneous environment. Applications depend on both the middleware and transport layers. Advanced applications include *Electronic Persistence Presence* (EPP) and UC. EPP, or e-presence, describes a state of a person or application as always being "on the network" in some form or another. The concept of session-based network access will no longer apply. EPP assumes that support for UC and both mobile and nomadic networking exists. UC refers to the pervasive presence of computing and networking capabilities throughout all of our environments; that is, in automobiles, homes, and even on our bodies.

A "research network," on the other hand, is a production network; that is, one aspiring to the goal of 99.99999-percent "up time" at Layers 1 through 3, which supports various types of domain-specific application research. This application research is most often used to support the sciences and education, but can also be used in support of other areas of academic and economic endeavor. These networks are often referred to as *Research Networks* (RNs) or *Research and Education* (R&E) *Networks*. In this article, we further classify these RNs based on their general customer base. *Institutional Research Networks* (IRNs) support universities, institutes, libraries, data warehouses, and other "campus"-like networks. *National Research Networks* (NRNs)[7], such as the Netherland's Gigaport or Germany's DFN networks, support IRNs or affinity-based networks. *Pan National Research Networks* (PNRNs) interconnect and support NRNs. An example of a couple of current production PNRNs are Dante's Ten-155 and the NORDUNET[8] networks. In this article we will also classify the older *National Science Foundation Network's* (NSFNET's), *very-high-performance Backbone Network Service* (vBNS), CANARIE's CA*NET 3[9], and the Internet 2[10] Abilene networks as PNRNs because in terms of scale and policy they address the same issues of interconnecting a heterogeneous set of regionally autonomous networks (for example, NSFNET's regionals and Internet 2's Gigapops) as do the PNRNs.

A hybrid state of RN also exists. When we introduce one or more advanced technologies into a production system, we basically inject some amount of chaos into the system. The interplay between the new technologies and other existing technologies at various levels of the infrastructure, as well as scaling issues, can cause unanticipated results.

Research quality systems engineering and design is then required to address these anomalies. An example of this phenomenon is the problem encountered with ATM cell discard and its effect on TCP streams and subsequent retransmissions (that is, early packet discard and partial packet discard). The term *Virtual Private Network* (VPN) is used in this article in the classical sense; that is, a network tunneled within another network (for example, IP within IP, ATM virtual circuits [VCs], and so on), and it is not necessarily a security-based network VPN. *Acceptable Use Policy* (AUP) refers to the definition of the type of traffic or use that is allowed on a network infrastructure. *Conditions of Use* (COU) is basically another version of AUP.

Background

During the early phases of the evolution of research networks and the Internet, national research networks were building and managing backbone networks because there was a technical reason to do so. Governments supported these activities, because at the time the commercial sector Internet Service Providers (ISPs) could not do it and the expertise to do so resided within the R&E community. Much of the research or testing of this time still focused on backbone technologies as well as aggregation networks and architectures. Research networks started out by supporting longer-term risky network research and quickly evolved to support shorter-term no-risk production infrastructure.

The research during the *Advanced Research Projects Agency Network* (ARPANET) and early NSFNET phases of the Internet focused on basic infrastructure protocols and technologies. Now commodity services, these services are both easily and cost-effectively available from the commercial sector. We have come a long way since then. Except for a few universities and research centers, the commercial sector now dominates R&D in the backbone technology space. Commercially provided VPNs can now cost-effectively support most of the requirements of the R&E communities. Given the current domination of R&D in backbone technologies by the commercial sectors, as well as the need to address true end-to-end services, it is time that network research and research networks realign their focus onto the research and development of end-system and campus and edge network technologies. Most of the intelligence of the network (for example, *Quality of Service* [QoS], security, content distribution and routing, and so on) will live at the edges, and in some way will be oblivious to the backbone service over which it will operate. In addition, in order for applications to be able to make use of this network, intelligent RNs need to be able to provide the middleware and services that exist between the application and the transport systems. The real future for most RNs is in helping to analyze and identify, not necessarily run and manage, advanced network infrastructures for their R&E communities.

One of the problems faced by the R&E community is how to obtain support from their governments and other supportive organizations (both for-profit and nonprofit). In attempts to support advanced applications and end-user research, organizations and governments may be convinced into supporting RNs, which end up providing commodity services and competing with the commercial sector. One reason that this can occur is that governments often wish to see results very quickly in order to justify their support of the research community; but, by doing so they drive the recipient researchers and research network providers to focus on short-term results and abandon basic long-term research. This pressure from the supporting organizations can also force researchers to compete in a space—that is, transport layers—for which industry may be better suited and adapted in both scale and time. Another issue facing today's research networks is that many of the R&E community, who once would endure downtime and assume some risk in trade for being part of an experimental network, are now demanding full production-quality services from those same R&E networks. Subsequently, the RNs are then being precluded from aggressively pursuing and using really advanced technologies that may pose a risk. And finally, many times research networks, science communities, and researchers claim they are doing network research, when in reality they are not, because they wish to have decent network connectivity, and they assume that this is the only way to get funding and support for good network connectivity with which to support their real research objectives. All of these issues have driven RNs at all levels into difficult positions. RNs need to be able to again take risks if they are to push the envelope in adopting new technology. Likewise, it is also valid to provide production-quality network transport services to support research for middleware, network application (for example collaborative technologies), and R&E application (for example, medical, sciences, education, and so on) research. All of these requirements need to be addressed in the manner most expedient and cost-effective to the government or organization providing the support.

All research carries with it a certain amount of risk. There is theoretical and experimental research. Some research is subject to validation; some is *retrospective*—for example, examining packet traces to verify the existence of nonlinear synchronization—but some is *prospective* and involves reprogramming network resources, and any reprogramming is susceptible to bugs. The amount of risk often depends on the area of research undertaken. The lower down in the network structure that one performs experimental research, the more difficult it is to support this research and still maintain a production-like environment for the other researchers and applications; yet we need to provide support for all levels of experimental research, as described in MORPHNET[11]. The ideal environment would support applications that could easily migrate from a production network to one prototyping recent network research, and then back again if the experiment fails. Recent advances in optical networking show promise in realizing this goal, but many technical and policy-based challenges are yet to be addressed.

ARPANET and Early NSFNET Phase: 1980s

The ARPANET, one of the many predecessors of today's Internet, was a research project run by researchers as a sandbox where they could develop and test many of the protocols that are now integral components of the Internet. Because this was a research network that supported network research, there were times the network would "go down" and become unavailable. Although that was certainly not the goal, it was a reality when performing experimental network research. This was acceptable to all involved and allowed for the quick "research-to-production" cycle, now associated with the Internet, to develop. The management of the network with respect to policy was handled by the *Internet Activities Board* (IAB), which has since been renamed the *Internet Architecture Board,* and revolved around the actual use of the network as a research vehicle. The research focused mainly on Layers 1 through 4, and application research was secondary and used to demonstrate the underlying technologies.

At the end of the 1980s, the Internet and its associated set of protocols rapidly gained speed in deployment and use among the research community. This started the major shift away from research networks supporting experimental network protocols toward RNs supporting applications via production research networks; for example, the mission agencies' (that is, those agencies whose mission was fairly well focused in a few scientific areas) networks at the *Department of Energy* (DoE) (ESnet[12]) and NASA (NSInet). At the same time, the NSFNET was still somewhat experimental with the introduction and use of "home-grown" T1 and T3 routers, as well as with pioneering research on peering and aggregation issues associated with the hierarchical NSFNET backbone. It also focused on issues relating to the interconnection of the major agency networks and international networks at the *Federal Internet Exchanges* (FIXes), as well as the policy landscape of interconnecting commercial e-mail (MCIMail) with the Internet. The primary policy justification for supporting these networks (for example ESnet, NSInet, NSFNET) in the late 1980s was to provide access to scarce resources, such as supercomputer centers, although the NSFNET still supported network research, albeit on peering and aggregation.

In addition, the NSFNET was first in pioneering research on network measurement and characterization, leading to today's *Cooperative Association for Internet Data Analysis* (CAIDA) as well as to Surveyor installations on Abilene. As researchers became dependent on the network to support their research, the ability to introduce new and risky technologies into the network became more difficult, as shown by the second-phase T3 router upgrade for the NSFNET when many researchers vehemently complained about any "downtime."

At this time, there were still no commercial service providers from which to procure IP services to connect the numerous and varied sites of the NSFNET and other research networks. Hence there were still valid technical reasons for NRNs and R&E networks to exist and provide backbone services.

The policy decisions affecting the interconnection of the agency networks at the FIXes, as well as engineering international interconnectivity, were loosely coordinated by an ad hoc group of agency representatives called the *Federal Research Internet Coordinating Committee* (FRICC). The FRICC became the *Federal Networking Council* (FNC) in the early 1990s, and then became the *Large-Scale Network* (LSN) working group by the mid-1990s.

The FNC wisely left the management of the Internet protocols to the IAB, the *Internet Engineering Task Force* (IETF), and the *Internet Engineering Steering Group* (IESG); however, the FNC did not completely relinquish its responsibility, as evidenced by its prominent role in prodding the development of *Classless Interdomain Routing* (CIDR) and originating the work that led to new network protocols (for example, IPv6).

The Next-Generation NSFNET: Early 1990s

During the early 1990s, the Internet evolved and grew larger. It could no longer remain undetected on the government policy radar screen. Many saw the NSFNET and agency networks as competing with commercial *Internet Service Providers* (ISPs). Because of the charters of the agencies of the U.S.-based RNs (for example NSF, DoE, NASA), all traffic crossing their networks had to adhere to their respective AUPs. These AUPs prohibited any "commercial entity-to-commercial entity traffic" to use a U.S. government supported network as transit. In addition, the demand for generic Internet support for all types of research and education communities became much stronger, and at the same time there was growing support among the U.S. Congress and Executive branches to end the U.S. Federal Government support of the U.S. Internet backbone.

In response to these pressures and the responses to a NSF draft "New NSFNET" proposal, the NSF elected to get out of the business of being the Internet backbone within the United States. This policy change was the nexus for the design of the vBNS, *Network Access Points* (NAPs), and *Routing Arbiter* (RA) described in the ABF paper[13] by early 1992. The vBNS was meant to provide the NSF supercomputer sites a research network that was capable of providing the high-end network services required by the sites for their Metacenter, as well as to provide the capability for their researchers to perform network research because the centers were still the locus for network expertise. The NAPs were designed to enhance the AUP free interconnectivity of both commercial and R&E ISPs and to further evolve the interconnection of the Internet started by the FIXes and the *Commercial Internet eXchange* (CIX).

The research associated with NRNs is already evolving from dealing with mainly IP and transport protocol research to research addressing the routing and peering issues associated with a highly interconnected mesh of networks. Research was an integral part of the NAP and RA design, but it was now focused on peering of networks as opposed to the transport layer protocols themselves. Although this network was not official until 1995, commercial prototype AUP free NAPs (for example, MAE-EAST) immediately sprang up and hastened the transition to a commercial network. The network was transformed from a hierarchical network topology to a decentralized and distributed peer-to-peer model. It no longer existed for the sole purpose of connecting a large aggregation of R&E users to supercomputer centers and other "one-of-a-kind" resources. The NAPs and the "peering" advances associated with the NAPs constituted a very crucial step for the success of applications such as the *World Wide Web* (WWW) and the subsequent commercialization of the Internet because they provided the required seamless interconnected infrastructure. Although some ISPs, for example UU-NET and PSInet, were quickly building out their infrastructure at that time, there still existed the need for PNRNs to act as brokers for acquiring and managing end-to-end IP services for their R&E customer base; it would not be much longer, however, before the ISPs had the necessary infrastructure in place to do this themselves.

The Internet 2 Phase: 1996–2000

The transition to the vBNS, NAP, and RA architecture became official early in 1995 and, as a result, the United States university community lost its government-subsidized production backbone. NSF-supported regionals had lost their support years earlier, and many had already transitioned to become commercial service providers, and the NSF "connections" program for tier 2 and lower schools persisted because it was felt (policy wise) that it was still valid to support such activities. The result of this set of affairs led to the creation of the Internet 2. Many of the top research universities in the United States felt that the then-current set of ISPs could not affordably provide adequate end-to-end services and bandwidth for the academic community's perceived requirements. As a result, the NSF decided to again support production-quality backbone network services for an elite set of research institutions. This was clearly a policy decision by NSF that had support from the U.S. Congress and Executive branches of government, even though in the early 1990s both Congress and the Executive branches were fairly vocal about not supporting such a network.

The initial phase was to expand to the vBNS and connect hundreds of research universities. The vBNS again changed from a research network, connecting a few sites and focusing on network and Metacenter research, back into a production research network. The vBNS is soon eclipsed by the OC-48 Abilene network. Gigapops, which are localized evolutions of NAPs, are used to connect the top R&E institutions to the Internet 2 backbones (that is, vBNS and Abilene).

These backbones were subject to COU as a way to restrict the traffic to that in direct support of R&E, much like the NSFNET was subject to its AUP.

The ISPs who complained so bitterly about unfair competition in the early 1990s no longer cared, because they had more business than they could handle in selling to corporate customers. An ironic spin on this scenario is that the business demands placed on the commercial ISPs by the late 1990s drove them to aggressively adopt new technologies to remain competitive. Not only were they willing to act as testbeds, they paid for that privilege since it gave them a competitive edge. The result is that in a lot of cases regarding the demonstration and testing of backbone-class technologies, the R&E community was time-wise behind the commercial sector. This situation is further aggravated by the fact that many, but not all, backbone network-savvy R&E folks went to work in industry. Another side effect of this transition is the loss of available network monitoring data. The data used by CAIDA, *The National Laboratory for Applied Network Research* (NLANR), and other network monitoring researchers had been gathered at the FIXes where most traffic used to pass. With the transition to a commercially dominated infrastructure, meaningful data becomes harder to obtain. In addition, as a result of the COU of the Internet 2 network, and the type of applications it supports (for example, trying to set bandwidth speed records), the traffic passing over its networks can no longer be assumed to be representative Internet data, and its value in this regard is diminished.

Another milestone is reached. ISPs have grown or merged so that they are offering both wide- and local-area network services, and anyone can now easily acquire national and international IP and transport services. The deployment and use of VPNs allows the commercial service providers (SPs) to provide and support various acceptable policy networks with differing AUP/COU on the same infrastructure. The technical need for most PNRNs or NRNs to exist to fulfill this function fades away. Researchers should now be able to specify wide-area network support as a line item in their research proposal budgets, just as they do for telephony and computing support. Most governments do not support separate research "Plain Old Telephone Service" (POTS) networks so that researchers can talk with one another. They provide funding in the grants to allow the researchers to acquire this from the commercial sector. However, valid technical reasons for selectively supporting some research networks still exist. A prime example is the CA*Net 3 network in Canada, which has been extremely aggressive in the adoption and use of preproduction optical networking technologies and infrastructure and has been instrumental in advancing our knowledge on this area.

During this evolution of research networks capabilities, network research is also going through its own evolution. DARPA starts focusing its research on optics, wireless, mobility, and network engineering as part of its Next-Generation Internet program. In addition, the research moves up the food chain of network layers. DARPA and DoE start supporting research on middleware. Globus[14], along with Legion[15], Condor[16], and POLDER[17], are major middleware research efforts that become the main impetus for GRIDs; and although they are focused mainly on seeking the holy grail of distributed computing, many of the middleware services they are developing are of value in a broader research and infrastructure context. The focus of network research and research networks now starts moving away from backbone transport services to research on advanced collaboratory, ubiquitous computing, mobile, nomadic, and EPP environments.

The policy management of the Internet now becomes an oxymoron and reflects the completion of the transition of the Internet to a distributed commercial Internet. Many organizations are now vying for a say in how the Internet evolves. Even the IETF is suffering from its own success. It now faces many of the same political challenges the ITU faced, that is, some commercial companies now try to affect the standards process for their own benefit by introducing standards contributions and only later disclosing the fact that they have filed patents on the technology in question. It is now much more difficult to make policy decisions regarding the future of Internet protocols, technologies, and architectures.

Future Frontiers

UC and EPP are the paradigm shifts at the user level that are already drastically altering our concept and understanding of networks. The scale, number, and complexity of networks supporting these new applications will far exceed anything we have experienced or managed in the past. Users will "be on the net" all the time, either as themselves or indirectly through agents and "bots." They will be mobile and nomadic. There will be "n" multiple instances of a user active on a network at the same time, and not necessarily from the same logical or geographical location. The frontiers associated with this new focus are many times more complex from a systems integration level than any work we have done in the past with backbone networks. This new frontier will provide new technical challenges at the periphery of the network; that is, the intelligent access and campus networks necessary to support these new environments. EPP and UC will drastically affect our research networks and application environments, much as the Web and its protocols drastically changed Internet and traffic patters in the 1990s.

The frontiers faced by research networks of the future will depend upon many technical and sociopolitical factors on a variety of levels. The sociopolitical frontiers can be divided into two different classes, one for e-developed nations who have already gone through the learning process of building an Internet-based infrastructure, and another for the e-challenged nations who still face the challenges of building a viable network transport infrastructure. The developed nations need to now grapple with how they can encourage the next evolutionary phase of their Internet-based economies. Because of the fast evolution of technology, the technical need for subsidizing transport-based network infrastructure is no longer the pressing need it was in the 1990s. The future research network will most likely be nothing more than a VPN based on a commercial ISP "cloud" service that interconnects researchers. The *High Energy Physicists* (HEPs) have already proved that life as a VPN-based affinity group overlaid on production network services is a viable solution to providing for their network requirements. The *High-Energy Physics Network* (HEPnet)[18] is a virtual set of users and network experts using ESnet and other ISP VPN-based network services to support the HEP scientists. Although we still have some technical challenges associated with backbone network technology (for example, optics), there are now only a very small number of institutions and organizations capable of working with industry and making substantial contributions in this area.

The new technical challenges that need to be addressed now include how to build and deploy intelligent edge and campus networks, content delivery and routing, mobile/nomadic/wireless access to the Internet, and the support for both UC and EPP. The latter two require major advancements and will require a whole bevy of middleware that is both network aware and an integral component of an intelligent network infrastructure. This includes, but is not limited to, directories, locators, presence servers, call admission control services, self-configuring services, mobility, media servers, policy servers, bandwidth brokers, intrusion-detection servers, accounting, authentication, and access control. IRNs and RNs can contribute to our knowledge and growth of these new areas by acting as leaders in areas that tend to be more difficult for the commercial sector to address, for instance, the development and deployment of advanced end-to-end services that operate over one or more ISP-provided clouds. Examples include interdomain bandwidth broker services, multi *Public Key Infrastructure* (PKI) trust models, defining multisite policies and schemas for directory-based policy services, and developing scalable naming conventions.

In order for policy makers to make informed decisions on the evolution and support of Internet technologies and architectures, they will need access to a generic mix of real backbone network data. There still exists a dire need at this point for such data. Innovative solutions that respect the privacy and business concerns of all types of ISPs and RNs, while at the same time making available "scrubbed" data, need to be developed. In addition, with the new focus on edge and metro networks, we might be able to shift our monitoring attentions to this area as well in order to better understand traffic demands and patterns on these scales of networks. Network monitoring is only one of the challenges facing us.

As the scale and complexity of networks grows, even at the pico and body area network level, we will need to develop new techniques to support network modeling, simulation, and experimentation. The University of Utah is developing a test facility[19] comprising a large number of networked processors, the network equivalent of a supercomputer center, to be used experimentally in the design and development of new transport layer protocols.

Summary

"Being on the net" will change our way of doing e-everything, and the evolution of the underlying infrastructure will need to change in order to support this paradigm shift. The intelligence of the network will not only move to the periphery, but even beyond, to the personal digital assistant and body area network. Therefore, it is important that the goals and focus of the research networks also evolve. Leave the R&D associated with backbone networks mainly with the commercial sector because this is their raison d'etre. The research networks of the future will be mostly VPNs, with a few exceptions, as noted earlier in this article. Research networks need to focus on the new technologies at the periphery as well as the middleware necessary to support the advanced environments that will soon be commonplace. Many research networks will themselves become virtual, for example, HEPnet, providing expertise but not necessarily a network service.

Policy makers must adapt to address not only these substantial technical and architectural changes but also second-order policy issues such as security and privacy and how to ensure that we don't end up with a bifurcated digital economy of e-savvy and e-challenged communities.

E-developed nations have already been through the technology learning curve of implementing and deploying a transport infrastructure. The e-challenged nations, with respect to network infrastructure, still face these same challenges, and they have the benefit of taking advantage of the knowledge of the nations who have successfully made the transition. In order to speed up the deployment of Internet technologies and infrastructure in the e-challenged nations, it may be best to first create technologically educated people and then to provide them an economic and social environment where they can apply their knowledge and build the infrastructure. E-savvy nations should help by providing the "know-how." The *North Atlantic Treaty Organization* (NATO) has a joint program with the *Trans-European Research and Education Networking Association* (TERENA) to provide for the instruction of Eastern European nations on the use and deployment of Internet technology (that is, how to configure and manage routers).

In lieu of subsidizing networks in these nations, NATO and TERENA are providing the basic knowledge that these people need to build, manage, and evolve their own networks and infrastructure. This should be the model to consider for e-developing nations. This is not to diminish the challenges of building network infrastructure in some areas where there is no such infrastructure, and perhaps in some of these areas working with other utility infrastructure providers might advance this cause.

Disclaimer

The ideas, comments, and projections proffered in this article are the sole opinions of the author, and in no way represent or reflect official or unofficial positions or opinions on the part of Cisco Systems, Inc. This article is based on my experience designing and managing operational international research networks, as well as being a program manager for network research, during the formative years of the Internet (that is, my tenure as a program manager for the United States Government's National Science Foundation and the Department of Energy), and my recent experience within Cisco working with next-generation Internet projects and managing its University Research Program. Many of the examples that I cite in this work are based on the development and deployment of the U.S.-based Internet and research networks, although the lessons learned in the United States may also be illuminating elsewhere.

Gratitude

I would like to thank my friend and colleague, Dr. Stephen Wolff, of the Office of the CTO, Cisco Systems Inc., for many good suggestions with respect to improving the content and presentation of this article; but, mostly for his good-humored authentication of my history and facts.

References

[0] This article was presented at the third Global Research Village Conference organized jointly by the Organization for Economic Cooperation and Development (OECD) and the Netherlands in Amsterdam, December 6–8, 2000.

[1] This is also attributed to the famous Physicist Niels Bohr.

[2] Wulf, William A. 1988. "The National Collaboratory—A white paper," Appendix A. In "Towards a National Collaboratory," Unpublished report of a National Science Foundation invitational workshop. Rockefeller University, New York, March 17–18, 1989.

[3] `http://www.nsf.gov/`

[4] `http://www.darpa.mil/`

[5] `http://www.gigaport.nl/`

[6] `Draft-aiken-middleware-reqndef-01.txt`, Internet Draft, Work in Progress, May 1999, `http://www.anl.gov/ECT/Public/research/morphnet.html`

[7] See **http://www.dante.org/** and **http://www.terena.nl/** for full lists of European research networks.

[8] **http://www.nordu.net/**

[9] **http://www.canarie.ca/**

[10] **http://www.internet2.org/**

[11] "Architecture of the Multi-Modal Organizational Research and Production Heterogeneous Network (MORPHnet)," Aiken, et al, ANL-97/1 technical report, and 1997 Intelligent Network and Intelligence in Networks Conference.
http://moat.nlanr.net/Papers/iinren.ps

[12] **http://www.es.net/**

[13] "NSF Implementation Plan for an Interagency Interim NREN," (aka Architecture for vBNS, NAPs and RAs), Aiken, Braun, and Ford, GA A21174, May 1992.

[14] **http://www.globus.org/**

[15] **http://www.cs.virginia.edu/~legion/**

[16] **http://www.cs.wisc.edu/condor/**

[17] **http://www.science.uva.nl/projects/polder/**

[18] **http://www.hep.net/hepnrc.html**

[19] **http://www.cs.utah.edu/flux/testbed/**

ROBERT J. AIKEN has an MS in Computer Science from Temple University. He is the Manager of the Cisco University Research Program. Prior to joining Cisco, Bob was the network and security research program manager for DoE's HPCC program and Next-Generation Internet (NGI) initiative. He was a program manager at the National Science Foundation (NSF), and with colleagues Peter Ford and Hans-Werner Braun coauthored the conceptual design and architecture of the second-generation National Science Foundation Network (NSFNET) (vBNS, Network Access Points [NAPs], and the Routing Arbiter [RA]), which enabled the commercialization of the then-U.S.-federally supported Internet. Before his NSF tenure, he served as DoE's ESnet program manager and was the creator and manager of the ESnet Network Information and Services group. Prior to his career in networking, Bob was responsible for managing supercomputers and coding their operating systems. His academic experience includes being an Assistant Professor of Computer Science at Hood College in Maryland, an adjunct Professor at California State University, Hayward, and the Manager of Technology Services at Gettysburg College in Pennsylvania. E-mail: **raiken@cisco.com**

Book Review

Intrusion Detection

Network Intrusion Detection—An Analyst's Handbook, by Stephen Northcutt, ISBN 0735708681, New Riders Publishers, 1999.

Network security and the ability to detect intrusion attempts has become extremely important in today's networks, regardless of size. I was looking for a book that would get technical on the details in these matters. Laura Chappell, the guru of packet-level information (**`www.packet-level.com`**), recommended this book to me. I should have realized what I was getting into at that point. I purchased the book, which was a bit expensive for its size at $39.99, and eagerly began reading it.

Mr. Northcutt starts out with a good discussion on how Kevin Mitnick conducted his famous attack. The book presents some very good information on a variety of topics, intermixed with personal observations and opinion. This made for an enjoyable read. If you are considering getting an *Intrusion Detection System* (IDS), then this book will provide you with some valuable insight and guidelines to consider from a recognized industry expert in this field. Mr. Northcutt is affiliated with The *System Administration, Networking, and Security* (SANS) *Institute* (**`www.sans.org`**).

Be aware that this book is not for the faint of heart. You will dive into the depths of packets and intrusion detection rather quickly, and never look back. This is both good and bad. I prefer an easy-to-read technical book, but the level of technical knowledge required to make sense of many of the examples is rather extensive. This includes how the many trace examples are presented in rather specialized fashion; in addition, the touted "detailed" explanations varied in usefulness quite a bit.

The book was marketed as a training aid; however, I suspect most readers need to be quite experienced to benefit from it. I admit I had to read many sections more than once in order to grasp the finer points being conveyed. I am confident that many readers have already echoed this sentiment to the author and publisher, since the second edition of this book was published in September 2000 and the page count has doubled, with only a modest price increase. I put it on my Christmas list!

—*Tom Thomas, Mentor Technologies Group*
`tothomas@mentortech.com`

Would You Like to Review a Book for IPJ?

We receive numerous books on computer networking from all the major publishers. If you've got a specific book you are interested in reviewing, please contact us and we will make sure a copy is mailed to you. The book is yours to keep if you send us a review. We accept reviews of new titles, as well as some of the "networking classics." Contact us at **`ipj@cisco.com`** for more information.

Call for Papers

The Internet Protocol Journal (IPJ) is published quarterly by Cisco Systems. The journal is not intended to promote any specific products or services, but rather is intended to serve as an informational and educational resource for engineering professionals involved in the design, development, and operation of public and private internets and intranets. The journal carries tutorial articles ("What is…?"), as well as implementation/operation articles ("How to…"). It provides readers with technology and standardization updates for all levels of the protocol stack and serves as a forum for discussion of all aspects of internetworking.

Topics include, but are not limited to:

- Access and infrastructure technologies such as: ISDN, Gigabit Ethernet, SONET, ATM, xDSL, cable, fiber optics, satellite, wireless, and dial systems
- Transport and interconnection functions such as: switching, routing, tunneling, protocol transition, multicast, and performance
- Network management, administration, and security issues, including: authentication, privacy, encryption, monitoring, firewalls, trouble-shooting, and mapping
- Value-added systems and services such as: Virtual Private Networks, resource location, caching, client/server systems, distributed systems, network computing, and Quality of Service
- Application and end-user issues such as: e-mail, Web authoring, server technologies and systems, electronic commerce, and application management
- Legal, policy, and regulatory topics such as: copyright, content control, content liability, settlement charges, "modem tax," and trademark disputes in the context of internetworking

In addition to feature-length articles, IPJ will contain standardization updates, overviews of leading and bleeding-edge technologies, book reviews, announcements, opinion columns, and letters to the Editor.

Cisco will pay a stipend of US$1000 for published, feature-length articles. Author guidelines are available from Ole Jacobsen, the Editor and Publisher of IPJ, reachable via e-mail at **`ole@cisco.com`**

Fragments

New Top-Level Domains

On November 16, 2000 The board of directors of the *Internet Corporation for Assigned Names and Numbers,* (ICANN) announced its selections for registry operators for new top level domains. The applications selected for further negotiation are the following:

`.aero`	Societe Internationale de Telecommunications Aeronautiques SC, (SITA)
`.biz`	JVTeam, LLC
`.coop`	National Cooperative Business Association, (NCBA)
`.info`	Afilias, LLC
`.museum`	Museum Domain Management Association, (MDMA)
`.name`	Global Name Registry, LTD
`.pro`	RegistryPro, LTD

The ICANN staff will now work through the end of the year to negotiate registry agreements with the applicants selected. The proposed schedule for completion of negotiations is December 31, 2000. The negotiated registry agreements must then be approved by the board of directors. Following that approval, the ICANN board will forward its recommendations to the U.S. Department of Commerce for implementation. For more on the history of ICANN's new TLD application process, please see **`http://www.icann.org/tlds/`** Multimedia archives of the annual meeting can be reviewed at **`http://cyber.law.harvard.edu/icann/la2000/`**

ICANN is a technical coordination body for the Internet. Created in October 1998 by a broad coalition of the Internet's business, technical, academic, and user communities, ICANN is assuming responsibility for a set of technical functions previously performed under U.S. government contract by IANA and other groups. Specifically, ICANN coordinates the assignment of the following identifiers that must be globally unique for the Internet to function: Internet domain names, Internet Protocol address numbers, and protocol parameter and port numbers. In addition, ICANN coordinates the stable operation of the Internet's root server system. As a non-profit, private-sector corporation, ICANN is dedicated to preserving the operational stability of the Internet; to promoting competition; to achieving broad representation of global Internet communities; and to developing policy through private-sector, bottom-up, consensus-based means. ICANN welcomes the participation of any interested Internet user, business, or organization. See **`http://www.icann.org`**

ISOC Launches Platinum Membership Level

The *Internet Society* (ISOC) is pleased to announce its *Platinum Sponsorship Program,* The Platinum program, which is in addition to and distinct from ISOC's standard organizational membership categories, provides interested organizations with the ability to designate support for specific areas of ISOC's work.

The initial participants, who also helped define the program, included Cisco, IBM, Microsoft, Nortel, RIPE NCC and SoftComca.com. APNIC has since joined the list of Platinum sponsors. Platinum level sponsors contribute $100,000 annually, with non-profit organizations eligible for funding at half that amount.

The Platinum program was initially developed to bolster support for the standards activities of ISOC, specifically ISOC's support of the *Internet Engineering Task Force* (IETF). Recently the program was expanded beyond Standards to include the three remaining areas of ISOC activities: Education & Training, Public Policy, and Member Services. As a result, participants in the Platinum program can now earmark their contribution for any of these four functional areas, or choose to allocate support for multiple areas, should they so desire.

ISOC is dependent upon individual and organizational members for its funding. ISOC believes that allowing contributors to designate where their money will be spent through the Platinum program enhances the Society's ability to undertake activities in these four areas, and, at the same time, provides an attractive support option for many organizations. ISOC will provide a report on the use of funds to each Platinum-Level sponsor at the end of each year. More information on the Platinum-Level Support Program can be found at:
`http://www.isoc.org/isoc/membership/platinum.shtml`

More information on ISOC's standard membership categories is available from: **`http://www.isoc.org/orgs/benefits.shtml`**

100 Million Internet Hosts

The Internet reached 100,000,000 hosts on 2 November 2000, according to John S. Quarterman, founder of Matrix.Net, a provider of Internet performance, measurement and intelligence. From its humble beginnings of 4 sites in the western United States in December 1969, the Internet has now reached over 150 countries and is nearly pole to pole. "This is an impressive achievement," said Quarterman. "We have been tracking the growth and development of the Internet for this entire decade. If this kind of growth continues, we will hit 1,000,000,000 hosts in 2006." For more information, see **`http://www.matrix.net/`**

This publication is distributed on an "as-is" basis, without warranty of any kind either express or implied, including but not limited to the implied warranties of merchantability, fitness for a particular purpose, or non-infringement. This publication could contain technical inaccuracies or typographical errors. Later issues may modify or update information provided in this issue. Neither the publisher nor any contributor shall have any liability to any person for any loss or damage caused directly or indirectly by the information contained herein.

The Internet Protocol Journal

Ole J. Jacobsen, Editor and Publisher

The Internet Protocol Journal is published quarterly by the Chief Technology Office, Cisco Systems, Inc.
www.cisco.com
Tel: +1 408 526-4000
E-mail: ipj@cisco.com

Cisco Systems

The Internet Protocol Journal, Cisco Systems
170 West Tasman Drive, M/S SJ-10/5
San Jose, CA 95134-1706
USA

ADDRESS SERVICE REQUESTED