The Internet Protocol Journal

September 2002 Volume 5, Number 3

A Quarterly Technical Publication for Internet and Intranet Professionals

In This Issue

You can download IPJ back issues and find subscription information at: **www.cisco.com/ipj**

From The Editor

The *Internet Protocol Journal* (IPJ) does not have a marketing department. New subscribers learn about IPJ through our Web page, or perhaps by picking up a copy at an Internet conference or meeting such as the IETF. Word of mouth is perhaps the most effective "marketing tool." I was reminded of this in July when an article in IPJ was mentioned on the *SlashDot* Web site. Within a few days we received more than 900 new subscriptions, on the order of ten times the normal sign-up rate. I think this illustrates the power of the Web as a tool for information dissemination.

I am a big fan of visitor networks. Such networks, typically found in larger hotels, allow high-speed access to the Internet for a daily or weekly fee. Although most of the conferences and meetings I attend have purpose-built "terminal rooms," it is still nice to be able to work in your hotel room at speeds orders of magnitude better than what can be obtained with a dialup modem. Dory Leifer explains how visitor networks are designed and operated in our first article.

In a previous article we explored the basics of IEEE 802.11 wireless networking. Such networks are growing at an amazing rate. Reports about wireless network "wiretapping" are frequently found in the trade press. Gregory R. Scholz describes an architecture for securing wireless networks, using a variety of technologies and protocols.

Geoff Huston is back with another opinion piece, this time discussing the role of the *Internet Service Provider* (ISP) as a "common carrier." Many ISPs are finding themselves in the middle of disputes between customers, copyright owners, regulators and others. What role should an ISP play in this regard? Geoff provides some answers.

Please continue to provide your feedback to anything you read in this journal. Our "Letters to the Editor" section provides a sample of some of the correspondence we receive. As always, use **ipj@cisco.com** to contact us.

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

Visitor Networks

by Dory Leifer, DEL Communications Consulting

Visitor networks are LANs that are most often deployed in hotels, airports, cafés, college campuses, apartments, and other locations. They enable the public network access on an ad-hoc basis. Recently, 802.11 "hot spots" have gained increased attention; they represent one example of a visitor network.

Visitors attach devices such as a laptop or *personal digital assistant* (PDA) that they use only while traveling or, more often, they attach machines normally used in the office or home. These machines can be thought of as "visiting hosts."

This article explores some of the technical issues with IP visitor networks and considers practical options for service provider deployment on wired Ethernet and wireless networks. In exploring deployment options, the article focuses mainly on solutions that do not require client software on the visiting host. These clientless techniques are based on heuristics and, although they do not work effectively under all circumstances, they have proven to be quite useful in practice.

For this discussion, it is assumed that the service provided by the visitor network is for access in one location at a time. Therefore, the article does not address network hand-off for mobile clients that are moving from one network attachment point to another while attempting to maintain connectivity.

Traditional LANs vs. Visitor Networks

Traditional LANs have been well optimized for enterprise networks. They provide high bandwidth and an economical and universal method of delivering network connectivity. In comparison, visitor networks are a rather curious hybrid of a LAN and a public network, such as one used for dial-in network access. Their objective is to physically use LANs to deliver what has normally been considered a public network service: *universal access*.

In enterprise networks, traditional LANs are usually carefully administrated. Normally the connected hosts are owned and administrated by the same enterprise that operates the network. Hosts that are connected to the network are configured according to the designated protocol and address schemes. They are often configured for at least *Simple Mail Transfer Protocol* (SMTP), *Post Office Protocol* (POP), file, and print sharing. On visitor networks, the hosts are typically owned and configured by the visitors, while the service provider administrates the network.

This difference in administration creates a serious challenge for the visitor network. The network must support a wide range of configurations because they will differ from one visiting host to another. For example, if a host had previously been configured for a static IP address, that address is likely to be from a different subnet, perhaps from a private network that the visitor normally uses at the office. Even if a host gets some of its configuration from *Dynamic Host Configuration Protocol* (DHCP), *Domain Name System* (DNS) and SMTP servers may refer to addresses or names on a private network that are not reachable on the visitor network.

Traditional wired LANs normally span physically secure areas, so any person who has access to the Ethernet wall jack for the building can connect anything to the network. With a visitor network it may be undesirable to allow everyone access. For example, a visitor network deployed in a university library may be available only to students. Similar to public dial-in access, visitor networks often rely on authentication and authorization before granting service.

Whereas LANs are excellent at facilitating peer-to-peer services such as file and print sharing between connected hosts, visitor networks often attempt to minimize these direct interactions between visitors, instead establishing a set of services that the service provider itself offers or simply routing the IP packets off the LAN to an Internet Service Provider. Minimizing interactions between visitors is desirable because service providers will want to reduce the risk of a visitor's machine being attacked by another visitor. On some occasions, however, visitors who do trust each other may want to use the visitor network for file sharing, printing, or even network gaming.

Going Clientless

One of the most difficult choices for service providers deploying visitor networks is to decide whether or not to rely on the installation of specialized client software on the visiting host.

Client software allows specific network protocols to be passed between the client and the visitor network. Protocols such as *Point-to-Point Protocol over Ethernet* (PPPoE)[1], *Layer 2 Tunneling Protocol* (L2TP)[2], and Mobile-IP[3] support both authentication as well as IP tunneling to assist in routing and address assignment. On some wireless LANs and networks with high-end Ethernet switches, 802.1x (which will be discussed in more detail later) supports flexible authentication schemes and aids in data encryption[4]. Although these protocols implemented on the client can present a significant technical advantage for implementing visitor networks, they require at least some modification to the configuration on the visiting host.

The lowest common denominator for traveling laptops is a simple TCP/IP stack and a browser. If the service can accommodate the visitor with only these items, the visitor network becomes much more suitable to the broadest audience. Of course without authentication, tunneling, and client configuration available from client software, the visitor network must rely on a set of heuristics or, said by some, hacks, to perform its tricks. Subsequent sections of this article illustrate technically how a visitor network can operate without relying on the installation of client software.

The service provider may choose to distribute client software in a situation where the visitor may use the service repeatedly. In many other situations, however, it is not feasible. For example, the last thing that travelers want to find in a hotel room upon arriving at midnight and needing a network connection is a CD-ROM full of new software drivers to drop on their laptop before using the hotel's in-room Ethernet. Even if the provided software does nothing but change the configurations, such as select a Web proxy server, it may have negative consequences when the laptop is returned to the office. Such added steps could also discourage visitors from using the visitor network again.

Visitor Network Basics

There are no hard guidelines or standards on what constitutes a visitor network. However, numerous vendors are selling devices that operate with wired and wireless networks, and act as gateways between the visitor network and the traditionally routed infrastructure. The typical visitor experience proceeds as follows (this is essentially a clientless example): the visiting host would not require the installation of special software, and in many cases would not require configuration changes:

- The visiting host is physically attached to the network by connecting to a twisted-pair Ethernet port.
- Visitors open their browser and attempt to load any page with the *Hypertext Transfer Protocol* (HTTP).
- Regardless of the specified URL, the browser loads a default page that requests authentication or billing information.
- When authenticated, the visitors now have general Internet access.
- An accounting record describing a visitor's session is generated and processed by the service provider's billing system, resulting in a charge on either the visitor's account or a corporate account.

Visitor Gateways

Visitor networks can be implemented with a special-purpose device called a "visitor gateway." Figure 1 illustrates the basic functional schematic of an example device. (Unfortunately, just about every vendor selling these devices uses a different name. This article uses the term in a generic sense and not to refer to any company's particular product.)

Figure 1: Visitor Gateway

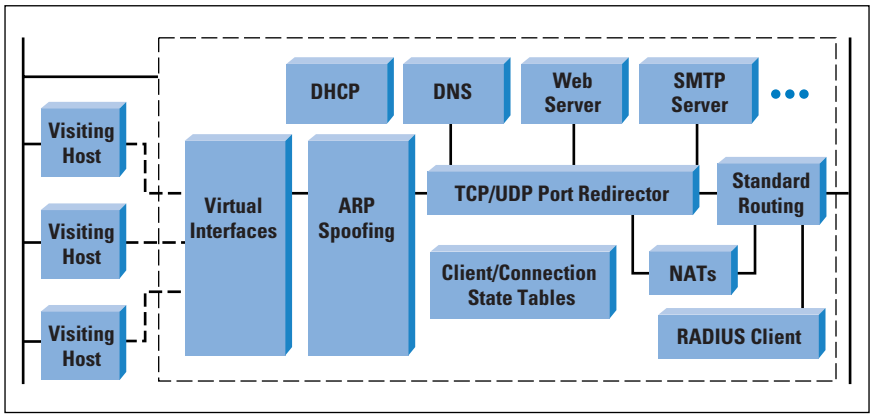

The visitor gateway sits between the LANs used to provide service to the visitors and a standard routed interface. Physically, a visitor gateway is a device that appears much like a router or firewall, with minimally two Ethernet interfaces.

Hybrid of NAS and LAN

The following sections focus on the visitor gateway, specifically its operational model, its handling of various Internet packet types, *virtual LANs* (VLANs), authentication, and accounting.

Visitor gateways behave as a hybrid of a standard LAN and a *Network Access Server* (NAS). For illustration, one can compare the operation of the visitor gateway with the operation of a NAS. Like a NAS with individual modem ports, the visitor network gateway typically builds virtual port structures as new hosts are discovered on the connected LAN. These virtual interfaces are configured by the gateway to accommodate the IP addresses used and referred to by the visiting host. The visitor gateway may create a virtual port structure for every host based on its *Media Access Control* (MAC) address or VLAN identifier and treat every virtual interface as an independent subnet upon which the visiting host and the virtual interface of the visitor network are the only attachments. Think of the relationship as a logical point-to-point link.

Conversely, the NAS, using the *Point-to-Point Protocol* (PPP)[5] on a dial-in connection, has a significant advantage over the visitor gateway in this scenario. PPP allows the NAS to negotiate an acceptable IP address for the dial-in client, set the client's default gateway, and even in some cases configure the client's DNS. The NAS normally has at least *Password Authentication Protocol* (PAP) and *Challenge Handshake Authentication Protocol* (CHAP) for authentication. If the visiting host requests configuration through DHCP[6], the visitor network has an opportunity to assign private or public addresses that are mutually convenient for both parties. On the other hand, if the visiting host already has a static address configured for its native network, for example, then the visitor gateway must spoof or imitate the behavior of the configured subnet.

The appeal of PPP in the dial-in world led to the recent development of PPPoE for LANs. Although PPPoE has been used with service selection gateways to offer public *Digital Subscriber Line* (DSL), there has been little use of it on visitor gateways. This is likely to be true because of the lack of a ubiquitous client and the complexities of solving multilevel authentication and encryption involving the local link, local network, and private network. PPPoE certainly is worth future study for visitor networks.

ARP

Hosts learn Layer 2 MAC addresses using the *Address Resolution Protocol* (ARP). Although hosts and routers respond only when asked about the IP address of their interfaces or those on a proxy-ARP table, visitor gateways usually respond with their own MAC address to any ARP requests from the attached visiting hosts, effectively proxying for the host's default gateway (if one is configured). The visitor gateway can also configure the interface address of its virtual port based on the host's IP address. In this manner, the gateway auto-configures itself to accommodate the visitor, who can continue to use his/her configured address.

Used on a standard shared LAN, this technique only goes so far. If, for example, one host on the visitor network shared its default router configuration with the IP addresses of another host (not that uncommon for private network numbers), then when the first host attempted to get the MAC address of its default router, it would end up with two responses, one from the visitor gateway and one from the other host on the LAN.

TCP/UDP Port Redirector

The visitor gateway for each *Transmission Control Protocol* (TCP) and *User Datagram Protocol* (UDP) packet received from the visiting host decides whether to pass the packet through or direct it to a local service such as DNS, SMTP, or Web server. It makes this decision based on some configured policy from the service provider (such as to redirect all SMTP) and from authorization states of the visitors. For example, if the service provider wishes to charge visitors $10 for daily access at a hotel, the port redirector could reflect HTTP requests to the local Web server that would, in turn, present the option to the visitor. Subsequent HTTP requests presumably would always be passed transparently through the gateway to the intended address.

The operation of the redirector is fairly simple. It works as a backwards network address-port translator. Instead of modifying the source, it modifies the destination and then applies standard IP forwarding on the resulting packets.

DNS

Visitor gateways typically implement proxies for domain name service requests and channel all DNS requests from the visiting host through the proxy. This serves at a minimum to reflect DNS requests to a closer DNS server, a useful performance advantage if the visitor's configured DNS server is a considerable distance away. Of greater significance is that it allows general Internet access by the visitor even if the configured DNS server is on a private network, which is now unreachable because the visitor's laptop has been moved from the office.

Redirecting to a DNS server not of the visitor's choosing may work smoothly until the visitor attempts to resolve domain names known only to the real DNS server on the private network. There is, of course, a limit to how well you can hide reality.

One common problem encountered by visitor networks is with a Web proxy on a private network. If the visitor refers to a Web proxy by name, the visitor gateway may choose to respond, inventing an IP address for the proxy and then assuming, by itself, operation of the proxy function. This technique has to be used with some care because hosts often cache DNS responses; these are effectively convenient lies that could end up being carried as "dirty entries" on the visitor's machine for longer than intended.

Rewriting DNS queries and responses does open the opportunity for the service provider to "assume" (some may say "hijack") sites. This opens the door to the possibility that, for example, `yahoo.com` is resolved to an address that is not Yahoo but rather a Web site with an affiliation to the service provider. Although this is a policy and business issue for the service provider, it is likely to irritate quite a number of visitors and reduce the perceived value of the service.

NATs

Visitor network gateways frequently use *Network Address Translation* (NAT), and often with port translation, in order to conserve IP addresses by sharing a small address pool with a large number of visitor hosts. In addition, NAT is required by the gateway if the source address used by the visiting host is not routable by the rest of the network back to the visitor gateway. This is almost always the case when the visiting host is using a static preconfigured IP address from another network. The gateway may choose its application of NAT based on policy. For example, two visitors may be configured for DHCP but one is assigned a private "Net 10" (RFC 1918) address that is passed through a NAT while another is assigned a routable address. In practice this flexibility is useful for service in apartments where the visitors are expected to "visit" for months. The service provider may choose to offer tiered services, one with a routable address suitable for the customers to run servers, and another with a private address suitable only for outgoing connections (e-mail, HTTP, and so on).

VLANs

The visitor gateway—modeling its relationship with visiting hosts as a virtual point-to-point link—may attempt to ignore the fact that hosts are on a shared network. However, certain interactions between hosts are inevitable on a shared LAN. For example, if a visitor's Windows laptop is configured for file sharing with no security enabled, other visitors may see, or worse, have permission to write to, critical files.

Virtual LANs provide a solution for isolating individual clients. On a wired Ethernet, many modern Ethernet switches can be configured to implicitly treat each port as a member of a different VLAN. For example, port 1 could be on VLAN 11; port 2 on VLAN 12; and so on. The visitor gateway is connected to one or more "trunk" ports that are configured as a member of all VLANs. This effectively allows another level of addressing so the visitor gateway can individually address a single Ethernet network connected to a port. The VLAN switches then act as simple concentrators. If a visiting host attempts to broadcast or multicast, these frames end up only traveling to the gateway and are not seen by other visiting hosts.

Figure 2: VLAN Frame Format

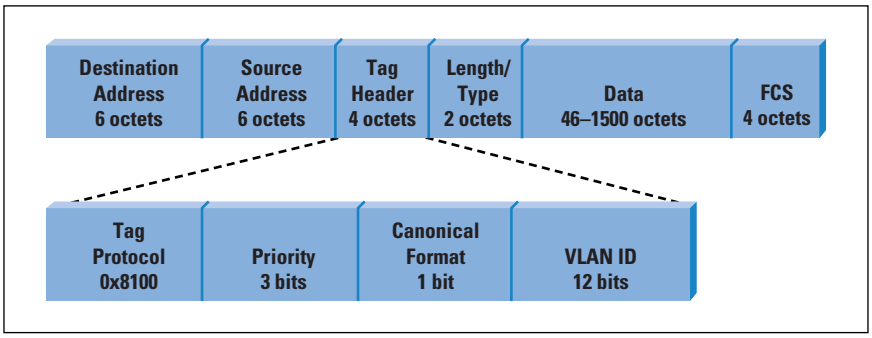

The VLAN frame format is shown in Figure 2. IEEE 802.1q defines the "tagging"[8]. The VLAN-enabled Ethernet switch can add the appropriate headers to standard Ethernet frames, and it forwards these through the trunk port. Optionally, another Ethernet switch concentrates the trunk traffic and attaches to the visitor gateway, as seen in Figure 3. One potential catch is that some Ethernet switches will not pass the oversized (maximum 1504 octet) VLAN frames; others attempt to be "overly aware" of the VLAN membership rules and insist on configuration of each of the VLANs, a challenging prospect if you are concentrating thousands of ports, each with a unique VLAN identifier.

Figure 3: VLAN Configuration

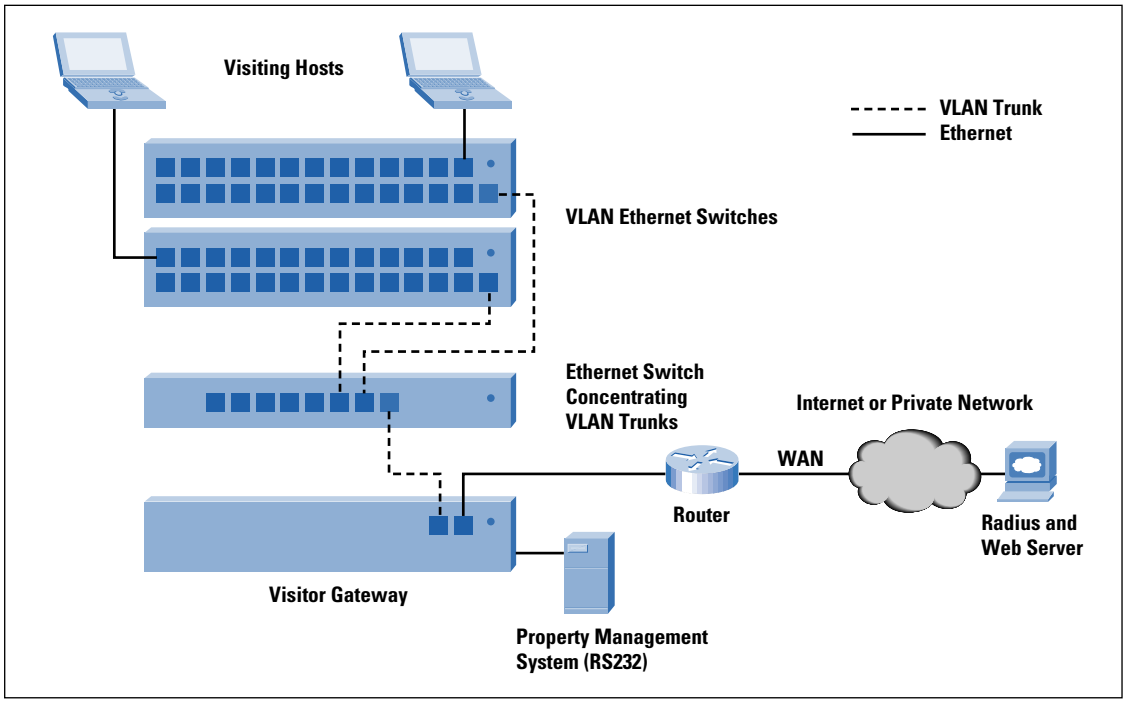

Some Ethernet switch vendors have implemented a nonstandard technique whereby broadcasts and multicasts are forwarded exclusively to a designated port, the theory being that if a host's broadcast and multicast frames do not get forwarded to other hosts, the hosts effectively will not "see" each other because they do not see ARP requests or higher-layer service advertisements. In some ways, this is simpler than using VLANs and provides some isolation over standard Ethernet networks.

Figure 4: Switch Multicast Blocking

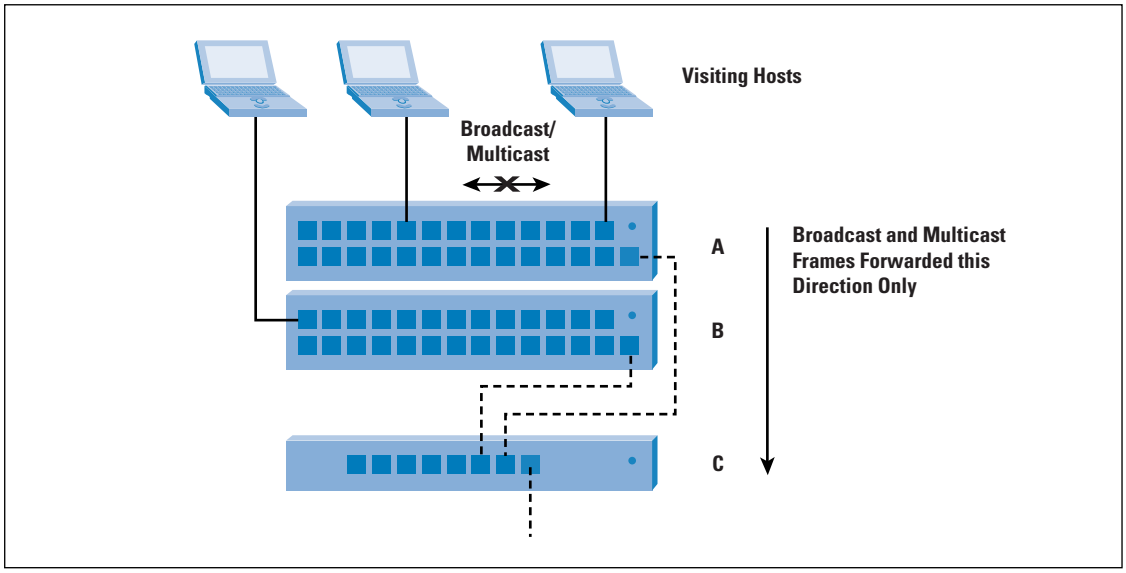

Combinations of these switches with normal ones can lead to some interesting frame forwarding scenarios. For example, as seen in Figure 4, Ethernet switches A and B are each connected to Ethernet switch C. Visiting hosts are attached to the ports on A and B. A and B are designed to have the forwarding restriction described, but C is a normal switch. This means that a broadcast from a visitor connected to A will not be seen by other visitors on A (by nature of the restriction) but it will be forwarded "upstream" to C, which will then forward it to B. Because B received it coming from the upstream, it will forward it to all the visiting hosts on B, causing the isolation technique to fail. A, B, and C all need to have the forwarding restriction.

Web-Based Authentication and Policy

Visitor networks often avail themselves of the one reliable way to converse with a human without additional client software: the Web browser. By selectively reflecting HTTP requests to the local gateway, the gateway can perform or facilitate several operations:

- Authenticate users with traditional username/password—The visitor gateway may, in turn, use a *Remote Access Dial-In User Service* (RADIUS)[7] authentication request to validate the user.
- Provide links within a "walled garden"—sites that can be visited without authentication—These sites are implemented with either a Web proxy inside the gateway or access control lists effective on the individual visiting host's virtual interface.
- Gather and validate credit card information through third-party credit card processing Web sites
- Offer visitors Web pages they can use to subscribe to services or to change service parameters

Using the browser can have a significant advantage, even over installed client software. The browser allows a conversation with a human user instead of a software client. This affords the network provider a wide variety of options, such as dealing politely with an authentication rejection, providing additional troubleshooting help, or confirming "conditions of use" before the user accepts charges. It is also a place for offering the user other products and services through Web links.

A central repository for visitor policy and configuration is especially important when a large number of gateways are deployed in disparate physical locations. An interesting option for visitor gateways is for them to learn policy by participating in the exchange of HTTP between the visiting host and an external Web server. The visitor gateway can piggyback the origin and state of a visiting host in a URL and refer the visitor's browser to a Web site. This origin information when presented to a service selection application running in a provider's data or operation center allows the application to determine which gateway the visitor is attached to as well as the visitor's virtual port identification and MAC address.

With the origin information, the service selection gateway can present the visitor with any number of billing, quality of service, or IP addressing options that apply to his/her connection. When the service selection application needs to affect the policy information stored in the visitor gateway, it can use a similar piggyback technique in the return direction.

Accounting

Finding an easy-to-deploy accounting method is crucial for service providers to generate accurate billing. The visitor gateway may send RADIUS accounting records in response to connections and disconnections made by visiting hosts. Disconnections can be determined by *Simple Network Management Protocol* (SNMP) traps from the physical layer devices or by repeated interval polling of the visiting host using ARPs or pings. Because RADIUS has been widely deployed by service providers for dial-in or other networks, it is very possible that the existing accounting system would be able to support the visitor gateway if it, too, offers RADIUS.

In hotels, accounting information can be sent directly to the hotel's *Property Management System* (PMS), causing users to see an access charge on their folio. This is normally accomplished by connecting a standard low-speed serial interface between the visitor gateway and the PMS. The visitor gateway posts the charges by exchanging records with the PMS. A simple record format is used to identify a room and associated charge. Although the format and exchange protocols are usually simple, they are rarely standard. Interfacing to a PMS may require the vendor of the visitor gateway to pay a license fee to the company selling the PMS before it can implement a PMS protocol. Additionally, after implementation, the visitor gateway vendor may need to go through certification for each PMS to which the gateway will be connected. Even if the equipment vendor pays the license, service providers are rarely free to go to a hotel and attach to the hotel PMS—often the service provider is shocked that the hotel insists that they be reimbursed for "interface license fees" charged by the PMS vendor to "enable the protocol."

The 802.1x Standard and Wireless LANs

Techniques of implementing visitor networks using wireless LANs (WLANs) have been both widely publicized and debated. Wireless 802.11 "hot spots" and the like have been the subject of great publicity because these WLANs are so convenient and cost-effective to deploy that they allow service providers to economically deploy them in areas that would be impractical to serve with wired networks. However, WLANs continue to be the topic of great debate because they have been plagued by the lack of compatibility and weaknesses in security architectures.

The 802.1x standard, recently ratified by the IEEE, holds the best promise in offering a standard authentication scheme for LANs. The 802.1x standard operates with client software. In one sample scenario, the visiting host, also known as the "supplicant," receives an *Extensible Authentication Protocol* (EAP) request/identity message from the visitor network via an Ethernet switch, a WLAN access point, or a visitor gateway, any of which function as the "authenticator." The authenticator then relays the client's identification to an authentication server. The server then decides if the supplicant is to be allowed access and responds appropriately to the authenticator.

With WLANs, the *Wired Equivalent Privacy* (WEP) keys can be loaded as part of the exchange so the client and access points can operate without manual key selection. WEP has been used for several years as a method of encrypting user data over the air interface. Without WEP (or even with it, as we have seen), anyone with a laptop and a receiver can spy on the exchanged traffic[10, 11].

Microsoft ships an 802.1x client in the standard distribution of Windows XP, an important move forward in making the protocol universal. Other software vendors are shipping or have announced product for older versions of Windows, Macintoshes, Linux, and some PDAs. The 802.1x standard client implementations, however, may need firmware support on the host adapters, and support may never be available on a large number of 802.11 cards already deployed. Furthermore, all 802.1x standards are not alike because they may implement different authentication schemes. Microsoft's current implementation uses the *Extensible Authentication-Transport Level Security* (EA-TLS) protocol, which requires a *Public Key Infrastructure* (PKI)[9]. Some critics contend that this creates additional deployment burdens on organizations with small networks. If a common provider, such as Boingo or T-Mobile, provides the visitor network in "hot spots" (that is, cafés and airports), the PKI requirement should not be an issue.

On Ethernet switches, 802.1x implemented directly on the switches may be adequate if the policy for visitor access is relatively simple. For example, if users on a particular network are all trusted employees working for the same business, the work of the authentication/authorization scheme is then to determine whether or not to allow someone to access the network, simply "port on" or "port off." A more sophisticated approach would allow users to be classified as belonging to a set of classes. On some switches, 802.1x would allow each port to assume a set of VLAN memberships. For example, VLAN 120 would allow unrestricted Internet access, VLAN 119 would restrict access to a set of Web servers, and VLAN 118 would restrict access further to only an authentication server. The authentication system using 802.1x would direct the switch port configuration.

In practice, the control required by visitor networks needs to be far more flexible, and perhaps should be left to the visitor gateway. The gateway, as diagrammed in Figure 1, can control the routing system as well as higher-level protocol proxies based on policy. Besides, leaving the authentication behind the switches allows network implementors the flexibility of using virtually any Ethernet switch, or even other media such as Ethernet framing over xDSL.

The switch-based 802.1x approach, however, may have a significant advantage over the visitor gateway in that after the authentication is out of the way, the Ethernet switch can switch traffic simply at full speed without additional per-packet overhead.

Security Concerns—Better Just to Bootstrap?

Visitor networks are particularly vulnerable to hacking and snooping by virtue of their physical locations, especially if serviced by WLANs. Unfortunately, security is one of the few things that a service provider cannot deliver to visitors without their explicit cooperation and participation. The service providers face a difficult choice to either stay out of the solution or attempt to deliver adequate security through client configuration or special software distribution. The answer is difficult to determine; however, at least two factors to consider are whether the network is wired or wireless, and what the expectations of the visitors will be.

Weaknesses in WEP commonly offered on wireless LAN products have been very well publicized[10,11]. These weaknesses involve the encryption protocols and the fact that most implementations use manually configured keys. The latter is of little use on a visitor network because the network provider would need to disclose the same keys to everyone. Better proprietary systems have been deployed using PKI, and 802.1x is also a possibility. WEP may be replaced by much stronger *Advanced Encryption Standard* (AES) in *Offset Codebook* (OCB) mode as part of the IEEE 802.1i working group[12]. No solution has been both standardized and universally deployed. The lack of a standard and universal solution to replace WEP requires that the service provider who chooses another form of security customize a wireless solution. They may need to distribute specialized client software and/or restrict their service to supporting a set of wireless cards and drivers.

Simple Ethernet switches can provide some isolation between ports, but the learning bridge algorithms they use are designed to efficiently deliver Ethernet frames, not provide a secure service. With many switches, it takes one frame with a sham source MAC address to convince the switch to spill someone else's traffic onto the wrong port. "Man in the middle" attacks are often trivial after a visiting host is tricked into sending its traffic somewhere else; the opportunities of doing this to another machine on the same LAN are abundant.

As an end user of a visitor network, trusting an unfamiliar service provider in an unknown environment is a fundamentally insecure process. So, why not let the visitor network provide the basic IP connectivity in order to bootstrap the connection, and then let the visitors themselves implement the security on top? One reason is that unsuspecting users getting hacked at their favorite hotel chain does not bode well for the hotel if the incidents end up in the press. Guests probably feel pretty secure using the hotel phone for a dial-in network connection without any encryption; many also feel secure locking the door with the sliding chain.

One reasonable compromise is matching the security of a dial-in connection. A wired Ethernet, assuming that it cannot be easily coaxed to spill traffic between ports, could present an acceptable risk level. On the other hand, a poorly protected wireless network is like a hotel door without a lock.

If the visitor network offers no protection, then the burden is placed completely on the visitors to implement their own end-to-end security. Using *Virtual Private Network* (VPN) software that implements *IP Security* (IPSec) is one possibility. Unfortunately, even that is not always straightforward, given the complexities with using protocols such as IPSec over NATs[13]. Other protocols such as *Transport Layer Security* (TLS) and *Secure Shell* (SSH), which operate above the network layer, may be a better option. In addition, several proprietary VPN protocols are designed to tunnel through NATs. Those without any security solution could compromise not only their personal data but also the security of their employer's networks.

Any long-term security solution is going to demand proper client configuration and compatible software. Ultimately, development of standards and client sophistication will make this possible, but in the meantime, we will need to choose between ease of connecting to an insecure network and dealing with the potential multiple layers of authentication and encryption before gaining access. Sadly, faced with this choice and looking forward to a 7 a.m. meeting, the trusty hotel phone and modem jack on the laptop might look pretty inviting.

Summary

Visitor networks allow service providers to provide access in public places. These networks can be implemented in a way that either may or may not require specialized client software on the visiting host. Client software allows service providers to more carefully control the behavior of the visiting host but, at the same time, may limit the user base to those who have the software installed.

Visitor networks often rely on a visitor gateway to perform functions generally not required on a traditional LAN. The gateway, which shares certain characteristics with a NAS, is responsible for routing, address assignment, translation, TCP/UDP redirection, authentication, accounting, and affecting policy.

The visitor gateway exchanges packets with the visiting hosts via LANs. On Ethernet, VLANs are often best suited to visitor networks because they allow the gateway to address each client separately providing the greatest level of isolation compared to other Ethernet options.

WLANs represent an important advance toward the universal deployment of visitor networks in "hot spots." However, the lack of a common and effective solution may force service providers to choose between ease of access and security. Visitors may choose to implement a VPN or security scheme on top of the raw IP access offered by the visitor network.

References

[1] L. Mamakos, K. Lidl, J. Evarts, D. Carrel, D. Simone, R. Wheeler, "A Method for Transmitting PPP over Ethernet (PPPoE)," RFC 2516, February 1999.

[2] W. Townsley, A. Valencia, A. Rubens, G. Pall, G. Zorn, B. Palter, "Layer 2 Tunneling Protocol, L2TP," RFC 2661, August 1999.

[3] S. Glass, T. Hiller, S. Jacobs, C. Perkins, "Mobile IP Authentication, Authorization, and Accounting Requirements," RFC 2977, October 2000.

[4] IEEE Standards for Local and Metropolitan Area Networks: Port-Based Network Access Control, IEEE Std 802.1X-2001, June 2001.

[5] W. Simpson, "The Point-to-Point Protocol (PPP)," STD 51, RFC 1661, July 1994.

[6] R. Droms, "Dynamic Host Configuration Protocol," RFC 2131, March 1997.

[7] A. Rubens, W. Simpson, S. Willens, C. Rigney, "Remote Authentication Dial-In User Service (RADIUS)," RFC 2058, January 1997.

[8] IEEE standard for local and metropolitan area networks: Virtual Bridged Local Area Networks, IEEE Std 802.1Q-1998.

[9] B. Adoba, D. Simon, "PPP EAP TLS Authentication Protocol," RFC 2716 (experimental), October 1999.

[10] N. Borisov, I. Goldberg, D. Wagner, "Intercepting Mobile Communications: The Insecurity of 802.11," **http://www.isaac.cs.berkeley.edu/isaac/wep-faq.html**

[11] E. Danielyan, "IEEE 802.11," *The Internet Protocol Journal,* Volume 5, Number 1, March 2002.

[12] D. Whiting, R. Housley, "AES Encryption & Authentication Using CTR Mode with CBC-MAC," Status of Project IEEE 802.11i, July 2002, **http://grouper.ieee.org/groups/802/11/Documents/DocumentHolder/2-001.zip**

[13] B. Adoba, "IPSec-NAT Compatibility Requirements," Internet-Draft, **http://www.ietf.org/internetdrafts/draft-ietf-ipsec-nat-reqts-01.txt**, March 1, 2002.

DORY LEIFER is a principal with DEL Communications Consulting, Inc. He had co-founded PublicPort in 1998, an Ann Arbor, Michigan, startup that developed one of the first visitor network gateways. After Tut Systems acquired PublicPort, he held a director of marketing position with Tut until 2001. Leifer spent 11 years with the University of Michigan and Merit Network and during that time contributed to the *Internet Engineering Task Force* (IETF). He has taught tutorials in access technologies for various seminars and tutorials, including NetWorld+Interop. He holds a B.S. in Computer Science from Rensselaer Polytechnic Institute and an M.S.E. in Industrial and Operations Engineering from the University of Michigan. Leifer currently resides in the San Francisco Bay Area and can be reached by e-mail at **leifer@del.com**

An Architecture for Securing Wireless Networks

by Gregory R.Scholz, Northrop Grumman Information Technology

Wireless networks are described as both a boon to computer users as well as a security nightmare; both statements are correct. The primary purpose of this article is to describe a strong security architecture for wireless networks. Additionally, the reader should take from it a better understanding of the variety of options available for building and securing wireless networks, regardless of whether all options are implemented. The security inherent with IEEE 802.11 wireless networks is weak at best. The 802.11 standard provides only for *Wired Equivalent Privacy,* or WEP, which was never intended to provide a high level of security[1]. For an overview of 802.11 and WEP, see reference[2]. Wireless networks can, however, be highly secure using a combination of traditional security measures, open standard wireless security features, and proprietary features. In some regard, this is no different than traditional wired networks such as Ethernet, IP, and so on, which have no security built in but can be highly secure. The design described here uses predominantly Cisco devices and software. However, unless explicitly stated to be proprietary, it should be assumed that a described feature is either open standard or, at least, available from multiple vendors.

Customer needs

Customer needs range from highly secure applications containing financial or confidential medical information to convenience for the public "hot spot" needing access to the Internet. The former requires multiple layers of authentication and encryption that ensures a hacker will not be able to successfully intercept any usable information or use the wireless network undetected. The latter requires little or no security other than policy directing all traffic between the wireless network and the Internet. Security is grouped into two areas: maintaining confidentiality of traffic on the wireless network and restricting use of the wireless network. Some options discussed here provide both, whereas others provide for a specific area of security.

The level of security required on the wireless network is proportional to the skill set required to design it. However, the difficulty of routine maintenance of a secure wireless network is highly dependant on the quality of the design. In most cases, routine maintenance of a well-designed wireless network is accomplished in a similar manner to the existing administrative tasks of adding and removing users and devices on the network. It is also assumed that security-related services such as authentication servers and firewall devices are available on the wired network to control the wireless network traffic.

It is not necessarily the case that one can see the user or device attempting to use the wireless network. This is the most alarming part of wireless network security. In a wired network, an unauthorized connected host can often be detected by link status on an access device or by actually seeing an unknown user or device connected to the network. The term "inside threat" is often used to refer to authorized users attempting unauthorized access. This is the inside threat because they exist within the boundaries that traditional network security is designed to protect. Wireless hackers must be considered more dangerous than traditional hackers and the inside threat combined because if they gain access, they are already past any traditional security mechanisms. A wireless network hacker does not need to be present in the facility. This new inside threat may be outside in the parking lot. *War Driving*[3] is the new equivalent to the traditional war dialing. All that is required to intercept wireless network communications is to be within range of a wireless access point inside or outside the facility.

Physical Wireless Network

In a highly secure environment, a best practice is to have the wireless access points connect to a wired network physically or logically separate from the existing user network. This is accomplished using a separate switched network as the wireless backbone or with a *Virtual LAN* (VLAN) that does not have a routing interface to pass its traffic to the existing wired network. This network terminates at a *Virtual Private Network* (VPN) device, which resides behind a firewall. In this manner, traffic to and from the wireless network is controlled by the firewall policy and, if available, filters on the VPN device. The VPN device will not allow any traffic that is not sent through an encrypted tunnel to pass through, with the exception of directed authentication traffic described later. With this model, the wireless clients can communicate among themselves on the wireless network, but there is no access to internal network resources unless fully encrypted from the wireless client to the VPN. This design may be further secured by configuring legitimate wireless-enabled devices to automatically initiate a VPN tunnel at bootup and by enabling a software firewall on the devices that does not allow communication directly with other clients on the local wireless subnet. In this manner, all legitimate communication is encrypted while traversing the wireless network and must be between authenticated wireless clients and internal network resources.

Authentication

Many security measures available relate to access controlled through individual user authentication. Authentication can be accomplished at many levels using a combination of methods. For example, Cisco provides *Lightweight Extensible Authentication Protocol* (LEAP)[4] authentication based on the IEEE 802.1x[5] security standard. LEAP uses *Remote Authentication Dial-In User Service* (RADIUS)[6] to provide a means for controlling both devices and users allowed access to the wireless network.

Although LEAP is Cisco proprietary, similar functionality is available from other vendors. Enterasys Networks, for example, also uses RADIUS to provide a means for controlling *Media Access Control* (MAC) addresses allowed to use the wireless network. With these features, the access points behave as a kind of proxy, passing credentials to the RADIUS server on behalf of the client. When these features are properly deployed, access to the wireless network is denied if the MAC address of the devices or the username does not match an entry in the authentication server. The access points in this case will not pass traffic to the wired network behind them. For security, the authentication server should be placed outside the local subnet of the wireless network. The firewall and VPN devices must allow directed traffic between the access points and the authentication server further inside the network and only to ports required for authentication. This design protects the authentication server from being attacked directly.

In addition to authenticating users to the wireless network, the VPN authentication and standard network logon can be used to control access further into the wired network. In this solution, the VPN client has the ability to build its tunnel prior to the workstation attempting its network logon, but after the device has been allowed on the wireless network. After the tunnel is built, specific rules on the VPN and the firewall allow the traditional network logon to occur. A robust VPN solution also treats the users differently based on the group to which they are assigned. Different IP address ranges are assigned to each group, allowing highly detailed rules to be created at the firewall controlling access to internal network resources based on user or group needs. The policy on the firewall must be as specific as possible to restrict access to internal resources to only those clients for whom it is necessary. Building very specific policy for users' access will also allow an *Intrusion Detection System* (IDS) to better detect unauthorized access attempts.

Encryption

LEAP also provides for dynamic per-user, per-session WEP keys. Although the WEP key is still the 128-bit RC4 algorithm proven to be ineffective in itself[7], LEAP adds features that maintain a secure environment. Using LEAP, a new WEP key is generated for each user, every time the user authenticates to use the wireless network. Additionally, using the RADIUS timeout attribute on the authentication server, a new key is sent to the wireless client at predetermined intervals. The primary weakness of WEP is due to an algorithm that was easy to break after a significant number of encrypted packets were intercepted. With LEAP, the number of packets encrypted with a given key can be tiny compared to the number needed to break the algorithm.

When using LEAP for user and device authentication, WEP encryption is automatically enabled and cannot be disabled. However, if added security is needed, a VPN, as described earlier, can provide any level of encryption desired. Using a VPN as the bridge between the wired and wireless network is recommended regardless of the underlying vendor or technology used on the wireless network. *IP Security* (IPSec) is a proven, highly secure encryption algorithm available in VPNs. By requiring all wireless network traffic to be IPSec encrypted to the VPN over the WEP-encrypted 802.11 Layer 2 protocol, any data passed to and from wireless clients can be considered secure. All traffic is still susceptible to eavesdropping, but will be completely undecipherable.

Aside from WEP and LEAP, some vendors provide other forms of built-in security. Symbol Technologies' Spectrum24 product provides Kerberos encryption when combined with a Key Distribution Center. Kerberos is more lightweight than IPSec and, therefore, may be better suited to certain applications such as IP phones or low-end *personal digital assistants* (PDAs). Other methods of automating the assignment and changing of WEP keys are also available, such as Enterasys' Rapid-Rekey[8]. Wireless vendors have realized that security has become of critical importance and most, if not all, are working on methods for conveniently securing wireless networks. When available, most vendors seemingly prefer to use open-standard, interoperable security mechanisms with proprietary security being additionally available.

Bringing it all together

Numerous options are available to secure a wireless network. A highly secure design will include, at a minimum, an authentication server such as RADIUS, a high-level encryption algorithm such as IPSec over a VPN, and access points that are capable of restricting access to the wireless network based on some form of authentication. When all the security options are tied together, the wireless network requires explicit authentication to allow a device and the user on the wireless network, the traffic on the wireless network is highly encrypted, and traffic directed to internal network resources is controlled per user or group by an access policy at the firewall or in the VPN.

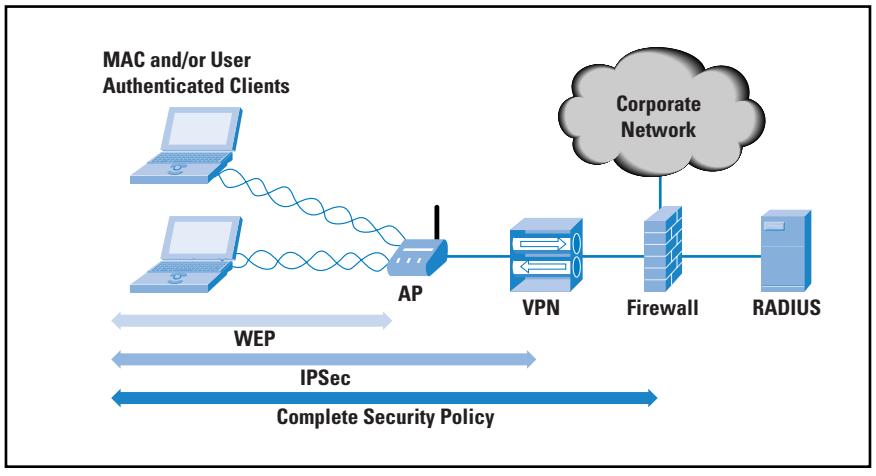

Figure 1: A Highly Secure Wireless Network

There is no substitute for experience and research when designing a network security solution. Using network security and design experience to exploit available technologies can further increase security of a wireless network. For example, grouping users into IP address ranges based on access requirements allows firewall access policy to help restrict unnecessary access. This can be accomplished using *Dynamic Host Configuration Protocol* (DHCP) reservations, assigning per-user or -group IP address ranges to the VPN tunnels or statically assigning addresses. Using a centralized accounts database for all authentication helps avoid inadvertently allowing an account that has been disabled in one part of the network to access resources through the wireless network. To use an existing user database for authentication while providing for dynamic WEP keys, use a LEAP-enabled RADIUS server that has the ability to query another server for account credentials. As with most network designs, a solid understanding of the available technologies is paramount to achieving a secure environment.

Utilizing all the security described in this article would yield the following design. When a device first boots up, it receives an IP address within a specified range on a segregated portion of the network. This IP range is based on the typical usage of the device and is most useful for machines dedicated to specific applications. As a user attempts to log onto a wireless device, a RADIUS server authenticates both the MAC address and the username of the device. If the user authentication is successful, access is granted within the wireless network. In order for traffic to leave the wireless network to access other network resources, a VPN tunnel must be established. Again, the IP address assigned to the tunnel can be controlled based on individual user authentication to help enforce access policy through the firewall. When the tunnel is established, firewall access policy will restrict access to resources on the network. Most, if not all, of the authentications required may be automated to use a user's existing network logon and transparently complete each authentication. This is not the most secure model, but it would be as secure as any single signon environment.

Summary

A secure wireless network is possible using available techniques and technologies[8] [9] [10]. After researching needs and security requirements, any combination of the options discussed here, as well as others not discussed, may be implemented to secure a wireless network. With the right selection of security measures, one can ensure a high level of confidentiality of data flowing on the wireless network and protect the internal network from attacks initiated through access gained from an unsecured wireless network. At a minimum, consider the current level of network security and ensure that the convenience of the wireless network does not undermine any security precautions already in place in the existing infrastructure.

References

[1] "Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," IEEE Standard 802.11, 1999 Edition.

[2] "802.11," Edgar Danielyan, *The Internet Protocol Journal,* Volume 5, Number 1, March 2002.

[3] "War Driving," Andrew Woods, **http://www.personaltelco.net/index.cgi/WarDriving**, last viewed August 11, 2002.

[4] "Cisco Aironet® Product Overview," Cisco Systems, **http://www.cisco.com/univercd/cc/td/doc/product/wireless/airo_350/350cards/pc350hig/pc_ch1.htm,** last viewed August 11, 2002.

[5] "IEEE Standard for Local and Metropolitan Area Networks—Port-Based Network Access Control," IEEE Standard 802.1X, 2001.

[6] "Remote Authentication Dial-In User Service," C. Rigney, S. Willens, A. Rubens, and W. Simpson, IETF RFC 2865, June 2000.

[7] "Security of the WEP Algorithm," Nikita Borisov, Ian Goldberg, and David Wagner, **http://www.isaac.cs.berkeley.edu/isaac/wep-faq.html,**, last viewed August 11, 2002.

[8] "802.11 Wireless Networking Guide," Enterasys Networks, June 2002, **http://www.enterasys.com/support/manuals/hardware/4042_08.pdf,** last viewed August 11, 2002.

[9] "Wireless LAN Security in Depth," Sean Convery and Darrin Miller, Cisco Systems, **http://www.cisco.com/warp/public/cc/so/cuso/epso/sqfr/safwl_wp.htm,** last viewed August 11, 2002.

[10] "Making IEEE 802.11 Networks Enterprise-Ready," Arun Ayyagari and Tom Fout, Microsoft Corporation, May 2001, **http://www.microsoft.com/windows2000/docs/wirelessec.doc,** last viewed August 11, 2002.

GREGORY SCHOLZ holds a BS in Computer and Information Science from the University of Maryland. Additionally, he has earned a number of certifications from Cisco and Microsoft as well as vendor-neutral certifications, including a wireless networking certification. After serving in the Marine Corps for six years as an electronics technician, he continued his career working on government IT contracts. Currently he works for Northrop Grumman Information Technology as a Network Engineer supporting Brook Army Medical Center, where he performs network security and design functions and routine LAN maintenance. He can be reached at: **gscholz@wireweb.net**

Opinion: The ISP—The Uncommon Carrier

by Geoff Huston, Telstra

There is a long-standing role in the communications industry where a provider of public carriage services undertakes the role of a *common carrier*. What's so special about the role of a common carrier, and why is this role one that is quite uncommon in the *Internet Service Provider* (ISP) world?

Side comment: There once was a time when you could not trust the messenger. There once was a time when not only did you pay to have your message sent, but you paid to *receive* messages. And there was no guarantee that the message would not be read by the messenger. The contents of your note could have been used to determine how much the receiver should pay for the message. Your message could have been copied and sold to other parties. If you can't trust the messenger, then communications becomes a risky business.

The Messenger

Throughout history the position of a messenger has been a mixed blessing. To be the bearer of bad news was not an enviable role, and rather than being rewarded for the effort of delivering the message, the messenger might have been in dire straits, given the level of wrath of the recipient. The option of reading the message before delivering it could be seen as a personal survival strategy, as well as being a prudent business move—bad news could be discarded immediately, whereas good news could have the potential of extracting a higher delivery fee from the recipient. Although this scenario would have been good for the messenger, such a mode of operation was not beneficial to all. For the parties attempting to use the messenger service, message delivery could be a very haphazard affair. The message might or might not get delivered, the delivery time was variable, as was the cost of delivery, and if the message itself was intended to be a secret, then one could confidently anticipate that the messenger would compromise this secrecy.

The Common Carrier

For a communications network to be truly useful, numerous basic attributes must be maintained. These include predictability, so that a message passed to a communications carrier is delivered reliably to the intended recipient. Integrity is also necessary, because a message must not be altered by the carrier in any way. Privacy is also an essential attribute, because the message must not be divulged to any party other than the intended recipient, nor should even the existence of the message be made known to any other party. And above all there must be a solid foundation for trust between the carrier and the clients of the service. So in this form of social contract, what does the carrier get in return?

Apart from payment for the service, the carrier is absolved from liability regarding the content of the messages, and from the actions of the customers of the service. This form of social contract is the basis for the status of a common carrier.

It may have taken some time, but this role is well understood by the public postal network. And as many national postal operators encompassed the role of national telephone carrier, the common carrier role has been an integral part of the public telephone network.

The ISP's Role

But in the world of the ISP the position of common carrier is very uncommon indeed.

There once was a time when folk did not need to encrypt their letters nor speak in scrambled code to undertake a private conversation. The assumption, made law in many countries, was that the entity entrusted with public communications, the common carrier, was barred from deliberately inspecting the contents of the plain transmission, and various dire penalties were in place if a public carrier's employees or agents divulged anything they may have learned by virtue of being public carriers. Various measures were put in place to execute interception and monitoring, but these measures required due process and reference to some law enforcement agency and also the judiciary to ensure that the rights of the public user were adequately safeguarded.

The issues of the role of a common carrier and the current role of an ISP are clearly seen when looking at the reactions to unsolicited commercial e-mail, or spam. Every day ISPs receive strident demands of the form: "One of your users is sending unsolicited messages—disconnect him now!" Internet users are, in effect, holding the ISP responsible for the actions of its customers. A similar expectation of the ISP's responsibility for the actions of its customers is seen in response to various forms of hacking, such as port scanning. Similar messages are sent to ISPs, demanding the immediate disconnection of those customers who are believed to be originating such malicious attacks. From a small set of complaining messages some years back, the volume of such demands for ISP action is now a clamor that is impossible for any ISP to ignore.

What should the ISP do? Many responsible ISPs see it as appropriate to conduct an investigation in response to such complaints. ISPs often include provisions in their service contracts with their customers to allow them to terminate the service if they believe that their investigation substantiates the complaints on the basis of a breach of contract. When disconnected, such customers are often blacklisted by the ISP to ensure that they cannot return later and continue with their actions. Surely this is an appropriate response to such antisocial actions?

This may be the case, but it is not necessarily consistent with the role of the ISP as a common carrier. A common carrier is not a law enforcement agency, nor is it an agent of the judiciary. It may be entirely appropriate for a common carrier to investigate, under terms of strict privacy, a customer's activities and inspect the contents of traffic passed across the network if it has reasonable grounds to suspect that the integrity of the network itself is under threat. Equally, it is probably inappropriate for a common carrier to extend the scope of such investigations on the basis of external allegations of activities that are not related to the integrity of the service itself.

The assumption that an ISP is, in some way, responsible for the actions of its customers has been extended further in some countries, such that the ISP is, in part, responsible for the content carried over its network, including content that originates with a customer of its service. This expectation that ISPs should actively control and censor content passed across their network is not just an expectation of some Internet users. This expectation appears in numerous legislative measures enacted in many countries. The *Communications Decency Act* in the United States legislature is an example of such an expectation of the active role of the ISP in controlling content passed across its network.

Who Will You Call?

Perhaps the issue here is one of expediency. Where can a user direct a complaint after receiving yet another piece of unsolicited, and possibly highly offensive, e-mail, apart from the ISP of the sender of the message? Where else can users direct a complaint after being the subject of yet another port scan of their system, but to the ISP? And what else can an ISP do in response? The ISP often has little choice but to investigate such complaints in good faith, and take corrective action if the complaint is substantiated. In the absence of any effective regulatory framework that would allow such investigations to be undertaken by an appropriate external agency, the ISP is in a difficult position.

Whereas it may be the correct common carrier position to disclaim all responsibility for the actions of its customers together with the content passed across its network, to ignore such complaints marks the ISP as a haven for such antisocial activities. Adopting such a position often has a negative impact on the ISP's ability to interconnect with other ISPs, because ISPs also tend to hold each other responsible for the actions of their customers and the content passed across their network. ISPs tend to avoid extending interconnection services to those ISPs that disclaim any such responsibility. So the expedient response is for the ISP to assume some level of responsibility for its customers and the content of its network and act accordingly.

But short-term expedient measures should not be confused with long-term effective solutions. The problem with these short-term responses lies in the uniquely privileged position of the carrier. Even rudimentary forms of data mining of each customer's communications patterns and the content of their communications can yield vast quantities of valuable information. Such information can allow a carrier to discriminate between customers, compromise the integrity of the customer's use of the network, and actively censor the content passed across the network. Positions of privilege without accompanying checks and balances are readily abused. There is already the widespread expectation and acceptance that an ISP has the ability and duty to inspect network content and monitor customers' activities with respect to various forms of anti-social and often malicious activities. But how can checks and controls be enforced such that the information gained through such monitoring activities is not used for other purposes? Such monitoring is not without cost, and the option of recouping some revenue to balance this expenditure by regarding this information as a business asset is always present. The regulatory impost of a common carrier role is intended to be an economically efficient response to this issue. The common carrier role is intended to reduce the social power of public carriers and protect the public's open, uncensored, and equal access to the carrier's services.

It is often said that the road to hell is paved with the best of intentions—that the ultimate outcome of the solution is potentially far worse than the immediate problem being addressed. The ultimate outcome of erosion of the common carrier role is that public users of a public communications service can confidently expect their communications to be monitored, potentially stored and cross referenced, and possibly later acted on.

Public Policy

Today the short-term expedient measures abound. There is enormous pressure on ISPs from both the Internet's user base and numerous legislatures to take an active position of being responsible—and liable, for the content on the networks and the actions of their clients. If left unchecked, this will have severe longer-term consequences for free speech, basic personal privacy, and uncensored, nondiscriminatory, universal access to the Internet. And when the user base comes to recognize the debased value of such a compromised communications system, they will inevitably look to other means of communication that have retained their essential integrity as a common carriage service.

Perhaps it is time for the debate regarding the role and responsibilities of an ISP to be placed on the agenda of public policy makers. Perhaps it is time to recognize that ISPs are indeed common carriers, and that they have a clearly bounded set of responsibilities with respect to both content and the actions of clients of the service.

Perhaps it is time to consider how best to enforce social norms on the Internet without compromising the basic integrity of the carrier as a neutral party to the content being carried across the network. Perhaps it is time to recognize that in this domain the Internet is not entirely novel, and what we have learned from a rich history of carriage provision in society has direct relevance to the Internet today.

The Internet is simply too valuable a communications service to have its long-term potential as a universal communications service mindlessly destroyed on the altar of short-term expediency.

Disclaimer: I am by profession neither a lawyer nor a public policy maker. However, by virtue of working in the ISP industry, I have an increasing level of interest in the activities of these folk, for the reasons outlined above. I should also note that personal opinion comes in many forms. The above is one such form.

GEOFF HUSTON holds a B.Sc. and a M.Sc. from the Australian National University. He has been closely involved with the development of the Internet for the past decade, particularly within Australia, where he was responsible for the initial build of the Internet within the Australian academic and research sector. Huston is currently the Chief Scientist in the Internet area for Telstra. He is also a member of the Internet Architecture Board, and is the Secretary of the APNIC Executive Committee. He is author of *The ISP Survival Guide,* ISBN 0-471-31499-4, *Internet Performance Survival Guide: QoS Strategies for Multiservice Networks,* ISBN 0471-378089, and coauthor of *Quality of Service: Delivering QoS on the Internet and in Corporate Networks,* ISBN 0-471-24358-2, a collaboration with Paul Ferguson. All three books are published by John Wiley & Sons.
E-mail: **gih@telstra.net**

Letters to the Editor

ENUM

Ole,

I was looking at the June 2002 issue of *The Internet Protocol Journal,* and noticed what might be a misprint. In the story on ENUM, the next-to-last paragraph on page 21 has a sentence reading:

> North America has the .164 country code of "1," implying that under ENUM there is a single DNS domain for ENUM, namely **`1.e164.arpa.`**

I suspect it should read "... there is a single DNS domain for North America..." or something like that. (The ".164" should probably also be "E.164"—you don't refer to it as just ".164" elsewhere in the article.)

A more substantive comment on Marshall Rose's BEEP article in the same issue: It was a good overview, but I would have liked to see a mention of which application protocols are likely to use BEEP (assuming that none has already) in the near future. The middle of page 11 explains why the IETF thinks this is a good idea and why new application protocols need BEEP, but it was hard to tell whether it actually is being actively considered for use by any IETF working group.

Overall, I liked the issue, and particularly Peter Salus's review of Padlipsky's book—I came across it in the late 1980s, and actually met Michael sitting in a hallway at one of the Interop conferences before they got too big for Silicon Valley and I stopped attending. I still remember some of his cartoons and slogans (e.g., something to the effect "... the ITU is planning to have an 11-layer model because it's a sacred number in Bali..."). I've also found the articles in some of the other recent issues of the IPJ—e.g., the articles on wireless LANs (particularly the discussion of security issues) and code signing/mobile code in the March 2002 issue—very helpful, and have pointed colleagues to them.

Best wishes.

—Eric M. Berg
Managing Director,
Technology Forecast Publications
PricewaterhouseCoopers Technology Centre
`Eric.Berg@us.pwcglobal.com`

Geoff Huston responds:

While we all try hard to eliminate various errors in manuscripts prior to publication, there are always a few author-mishaps that manage to sneak past the eagle eyes of the editor, and this is one of them.

The offending sentence should read:

> North America has the .E164 country code of "1," implying that under ENUM there is a single DNS domain for ENUM in North America, namely **`1.e164.arpa.`**

Thanks for pointing this out.

—Geoff

More about ENUM

Ole,

In the June 2002 issue of IPJ (Volume 5, Number 2), Geoff Huston wrote an interesting article about ENUM. The technical side of ENUM (using DNS to map E164 numbers to services) seems rather straightforward. But its implications on both technical and social issues are much more complex and (in my opinion) interesting. I am not an expert on the subject, but I'd like to share a few thoughts about this. First, two technical issues come to mind.

The first one is about the use of the *Domain Name System* (DNS). The DNS has been very successful as a distributed replicated database of hostname-to-IP address (and reverse) mappings. Will it be able to handle gracefully all the stuff people intend to put in it? This is not certain, as shown by ICANN's cautious attitude concerning the creation of new Top Level Domains. Content Distribution Networks, for example, often use lots of domain names with short TTLs, reducing the effectiveness of DNS caching (Geoff mentions this caching issue for ENUM). After all, DNS stands for "Domain Name System," not "General Purpose Infinitely Scalable Distributed Dynamic Database."

The second issue is about the status of addresses and names in the Internet. Simplifying things, we can say the following happens when somebody wants to access an Internet service with an E.164 number: The E.164 number is translated into a DNS name, and a DNS lookup gives back an URI. If the URI is a simple URL, the domain name in the URL is DNS-looked-up for an IP address, and then packets are sent to that IP address. If the URI is not a simple URL (such as a URN), some other resolving process implying the DNS occurs anyway.

That makes two levels of indirection, but, moreover, creates an "interesting" situation: IP addresses are "addresses," i.e., network-friendly identifiers, whose structure is tied to the network topology.

Such identifiers are not user friendly, so user-friendly identifiers called "names" have been created, and a "domain name system" set up to translate names into addresses. E.164 numbers are really telephone addresses. They are tied to the telephone network topology and are surely not user friendly. There are no user-friendly names in the telephone system.

The strange thing is that with ENUM, E.164 numbers are not linked anymore to the network topology, but rather become names intended for user usage. In a sense, they even are "meta names," since they translate to DNS names (that translate to addresses). But they obviously have not become user-friendly in the process.

I must admit I oversimplify a bit since I don't distinguish between names and addresses identifying level 3 (network) resources (i.e., hosts) and those identifying level 7 (application) resources (e-mails, Web pages, etc.), but this doesn't invalidate the idea.

Addresses are what the network needs, and names are what the users need. This brings me to the politics aspects of ENUM: who administers/controls/owns the namespace? A namespace is only partly technical; defining a namespace includes defining how and by whom the namespace is operated. The DNS is technically a big success, but the politics side is controversial, as shown by domain-name disputes or the setting up of alternative domain-name systems. It seems that social aspects are often more difficult to deal with than technical issues are to solve.

When I was studying networking we were taught how the technical differences between the Internet and the telephone network took their roots into a fundamental difference of culture. Now that the Internet culture seems to have won on the technical aspect (IP over broadband ISDN), wouldn't it be a strange outcome for the Internet namespace to be owned by telephone companies?

To conclude, I think this ENUM stuff shows that the Internet community really needs to work on the namespace issue, to ensure a technically and socially sound namespace for the Internet.

—Christophe Deleuze, Ph.D.
R&D Senior Engineer
ActiVia Networks
Christophe.Deleuze@ActiVia.net

Book Review

Carrier-Scale IP Networks

Carrier-Scale IP Networks: Designing and Operating Internet Networks, edited by Peter Willis, ISBN 0-85296-982-1, The Institute of Electrical Engineers, London, United Kingdom, 2001

My heart jumped when I saw the nondescript brown box, about the thickness of a book, sitting by the receptionist. It was finally here! I had waited almost two months in great anticipation for this book to show up. Was it going to be the all-encompassing handbook for the network designers, operators, and managers in large-scale IP environments? The first few lines in the text indicated that it just might be: "The aim of this book is to give the reader an understanding of all the aspects of designing, building and operating a large global IP network."

The definition of "large-scale" as given by the author and for the purposes of this review follows: Provides services for millions of end users, high-speed (greater than 100 Mbps) transit services, and is reliable, scalable, and manageable.

One thing to keep in mind is the way this book was constructed. The 16 chapters had 29 authors. Almost all authors came from some area of British Telecom (BT) and all were subject matter experts in the chapter they wrote. The 16 chapters are grouped roughly into four sections: Designing and building IP networks, transmission and access networks, operations, and development of future networks. Sadly, all of this is squeezed into 293 pages.

Designing and building IP networks

For the reader new to designing and building large-scale IP networks, the first few chapters are gold. For the reader already experienced in this area, it may bring back nostalgic feelings for the good old days of exponential growth. A lot of ground is covered, including the obligatory overview of IP, sufficient enough to give a nontechnical person the key concepts of IP routing, but can by skipped by those with even basic knowledge in this area. The examples given throughout this chapter (and the rest of the book) come directly from the design of BT's and Concert's backbone. A whole chapter, "The Art of Peering," not to be mistaken for an excellent paper of the same name[1], gives excellent key concepts in peering. Some coverage is even given to the logistics and difficulties in building points of presence globally, going so far as to mention earthquake bracing for equipment bays.

The next set of chapters give the reader detail about the transmission network (for some, be prepared to think *Synchronous Optical Network* [SONET] when you read *Synchronous Digital Hierarchy* [SDH]), and access networks, including various forms of broadband, wireless, dial, and satellite.

The technical information was squeezed into these chapters, not enough for a good technical treatise, but enough to give readers good grounding in a technology that is unfamiliar to them. The coverage was closer to being marketing material. These chapters alone are not enough to bring those new to the field up to speed if they are to design or operate such a network.

BT opened itself up and gave us a view into the operations of its network. Individuals who have worked in an environment like this will find something familiar. We get to see how BT structures the people, processes, and technologies. This is something that is not usually open to inspection by people outside of an organization. Planning and developing the operations side of the house is a difficult job. These chapters may give a kick-start to those coming into such a role.

I was disappointed with the two final chapters. Of course anything listed as being "the future" will one day become the present, but I digress. These two chapters seem like the odd couple that just did not fit with the rest of the chapters. The first chapter is on Traffic Engineering. It is really a primer on *Multiprotocol Label Switching Traffic Engineering* (MPLS TE). The second chapter covers *Virtual Private Networks* (VPNs), both the MPLS and *IP Security* (IPSec) types.

Recommendation

The authors set out with a lofty goal, and did not quite hit the mark. This book would be appropriate for someone trying to get a feel for what goes on inside of a carrier-scale network. People already in the business would be better served by just paying attention to what goes on around them.

Perhaps a small focused group could set out to create a book (or should I say tome) covering the elements of design, the foundation of support, and the basics of management. Something timeless is required here, independent of the protocol du jour, to develop the next generation of competent netheads.

—Kris Foster
kris.foster@telus.com

[1] "The Art of Peering: The Peering Playbook," William B. Norton, Equinix

Fragments

Stephen Wolff receives Postel Service Award

In June 2002, Internet pioneer Stephen Wolff was honored by the *Internet Society* (ISOC) for his significant contributions on behalf of the Internet. A founding member of the ISOC, Wolff is considered one of the "fathers of the Internet" and was directly involved with its development and evolution.

Wolff received the *Postel Service Award,* named for Dr. Jonathan B. Postel, an Internet pioneer and head of the organization that administered and assigned Internet names, protocol parameters, and *Internet Protocol* (IP) addresses. He was the primary architect behind what has become the *Internet Corporation for Assigned Names and Numbers* (ICANN), the successor organization to his work. The recipient of the award receives a $20,000 cash honoraria.

"We are pleased to recognize Steve with the Postel Award," said ISOC President/CEO Lynn St.Amour, "especially as his contributions are well know to ISOC, having previously been commended by ISOC's board for helping transform the Internet from an activity serving the particular goals of the research community to a worldwide enterprise which has energized scholarship and commerce in dozens of nations."

The 1994 commendation from the ISOC board also states that "The personal leadership of Dr. Wolff, often under conditions of public controversy, has been an indispensable ingredient in surmounting a daunting array of technical, operational and economic challenges. His extraordinary commitment to the growth and success of the Internet reflect the highest standard of service to the networking community and command our respect and admiration."

As Director of the Division of Networking and Communications Research and Infrastructure at the US National Science Foundation, he was responsible for NSNET, the *National Research and Education Network* (NREN), and for NSF's support of basic research in networking and communications. While at the NSF he was among the founders of the interagency and international research networking management and advisory structure whose descendants today include the Large-scale Networking (LSN) working group and the PITAC.

Wolff left the federal government and joined Cisco Systems, Inc. in 1995, where he works in the University Research Program—Cisco's program supporting academic investigators with unrestricted grants for research on computer networks.

Wolff was educated at Swarthmore College, Princeton University, and Imperial College. He taught electrical engineering at the Johns Hopkins University for ten years and subsequently spent fifteen years leading a computing- and network-related research group at the U.S. Army Research Laboratory. In 1983 he took a sabbatical half-year as a Program Director in the Mathematics Division of the U.S. Army Research Office.

ISOC is a not-for-profit membership organization founded in 1991 to be the international focal point for global cooperation and coordination in the development of the Internet. Through its current initiatives in support of education and training, Internet standards and protocol, and public policy, ISOC has played a critical role in ensuring that the Internet has developed in a stable and open manner. For 10 years ISOC has run international network training programs for developing countries which have played a vital role in setting up the Internet connections and networks in virtually every country that has connected to the Internet. For more information, please visit: **`http://www.isoc.org/`**

ISOC to Run `.org`?

Recently ICANN posted a preliminary Staff Report on the selection of a new registry operator to assume responsibility on January 1, 2003 for the **`.org`** registry. The report, which is subject to public comment and comment by all the bidders before being submitted for approval to the ICANN Board of Directors, recommends that the Board select the *Internet Society* (ISOC) as the successor registry operator for the **`.org`** registry, currently operated by VeriSign.

This preliminary report follows an extensive bid solicitation and evaluation process that was launched last April. Eleven bids were received in response to a Request for Proposals. These bids were analyzed and evaluated by three evaluation teams that operated independently of each other.

"We received eleven very strong and thoughtful proposals," noted Stuart Lynn, President of ICANN. "We appreciate the response of the institutions behind these proposals. The ISOC proposal was the only one that received top ranking from all three evaluation teams. On balance, their proposal stood out from the rest." Lynn also emphasized the openness and transparency of the solicitation and evaluation process.

Two evaluation teams focused on technical issues: one from Gartner, Inc., an international consulting and research organization that specializes in information technologies, and the other a team mainly composed of CIOs of major universities. Another team was provided by ICANN's *Non Commercial Domain Name Holders* constituency; the NCDNHC team focused on the effectiveness of the proposals to address the particular needs of the **`.org`** registry. The staff report integrates these evaluations and other factors into the preliminary recommendation.

ISOC is an international not-for-profit organization of over 6,000 individual and 150 organizational members with chapters in over 100 countries. It provides leadership in addressing issues that confront the future of the Internet, as well as being a home for the *Internet Engineering Task Force* (IETF) and the *Internet Architecture Board* (IAB).

In operating the **.org** registry, ISOC will team with Afilias, an operating registry that recently launched the **.info** *top level domain* (TLD) that was authorized by ICANN as one of seven new TLDs over this past year.

"Afilias will provide ISOC with the necessary experience at operating a large registry," said Lynn. "The **.info** registry already houses about 1 million domain names, which is on a scale that approaches the much older **.org** registry."

ICANN is re-assigning the **.org** registry under a revised agreement among ICANN, VeriSign, and the U.S. Department of Commerce that was signed in May 2001. Under that agreement, VeriSign was permitted to keep its registrar business, NSI (that it was obligated to sell under the prior agreements) provided that it agreed to relinquish **.org** at the end of December 2002, and subject to other provisions of the revised agreements. As part of those revised agreements, VeriSign agreed to endow the new operator with US$ 5 million to help fund operating costs, provided that the new operator was a not-for-profit organization.

Following an open and transparent process, ICANN has posted all eleven applications online together with all supplemental material and community comments received. The preliminary staff report and the evaluations are posted at:
http://www.icann.org/tlds/org/preliminary-evaluation-report-19aug02.htm.

Applicants and any member of the community are invited to send comments on the preliminary report and evaluations by e-mail to:
org-eval@icann.org

Upcoming Events

The *IETF* will meet in Atlanta, Georgia, USA, November 17–21, 2002.
http://www.ietf.org/meetings/meetings.html

ICANN will meet in Shanghai, China, October 27–31, 2002.
http://www.icann.org/meetings/

The next *Asia Pacific Regional Internet Conference on Operational Technologies* (APRICOT) will take place February 19–28, 2003 in Taipei, Taiwan. **http://apricot2003.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

Editorial Advisory Board

Dr. Vint Cerf, Sr. VP, Internet Architecture and Technology
WorldCom, USA

Dr. Jon Crowcroft, Marconi Professor of Communications Systems
University of Cambridge, England

David Farber
The Alfred Fitler Moore Professor of Telecommunication Systems
University of Pennsylvania, USA

Peter Löthberg, Network Architect
Stupi AB, Sweden

Dr. Jun Murai, Professor, WIDE Project
Keio University, Japan

Dr. Deepinder Sidhu, Professor, Computer Science &
Electrical Engineering, University of Maryland, Baltimore County
Director, Maryland Center for Telecommunications Research, USA

Pindar Wong, Chairman and President
VeriFi Limited, Hong Kong

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, Cisco Systems, and the Cisco Systems logo are registered trademarks of Cisco Systems, Inc. in the USA and certain other countries. All other trademarks mentioned in this document are the property of their respective owners.

Copyright © 2002 Cisco Systems Inc. All rights reserved. Printed in the USA.

The Internet Protocol Journal, Cisco Systems
170 West Tasman Drive, M/S SJ-7/3
San Jose, CA 95134-1706
USA

ADDRESS SERVICE REQUESTED

PRSRT STD
U.S. Postage
PAID
Cisco Systems, Inc.