InformationWeek
:: analytics
Analytics.InformationWeek.com

June 2011
$99

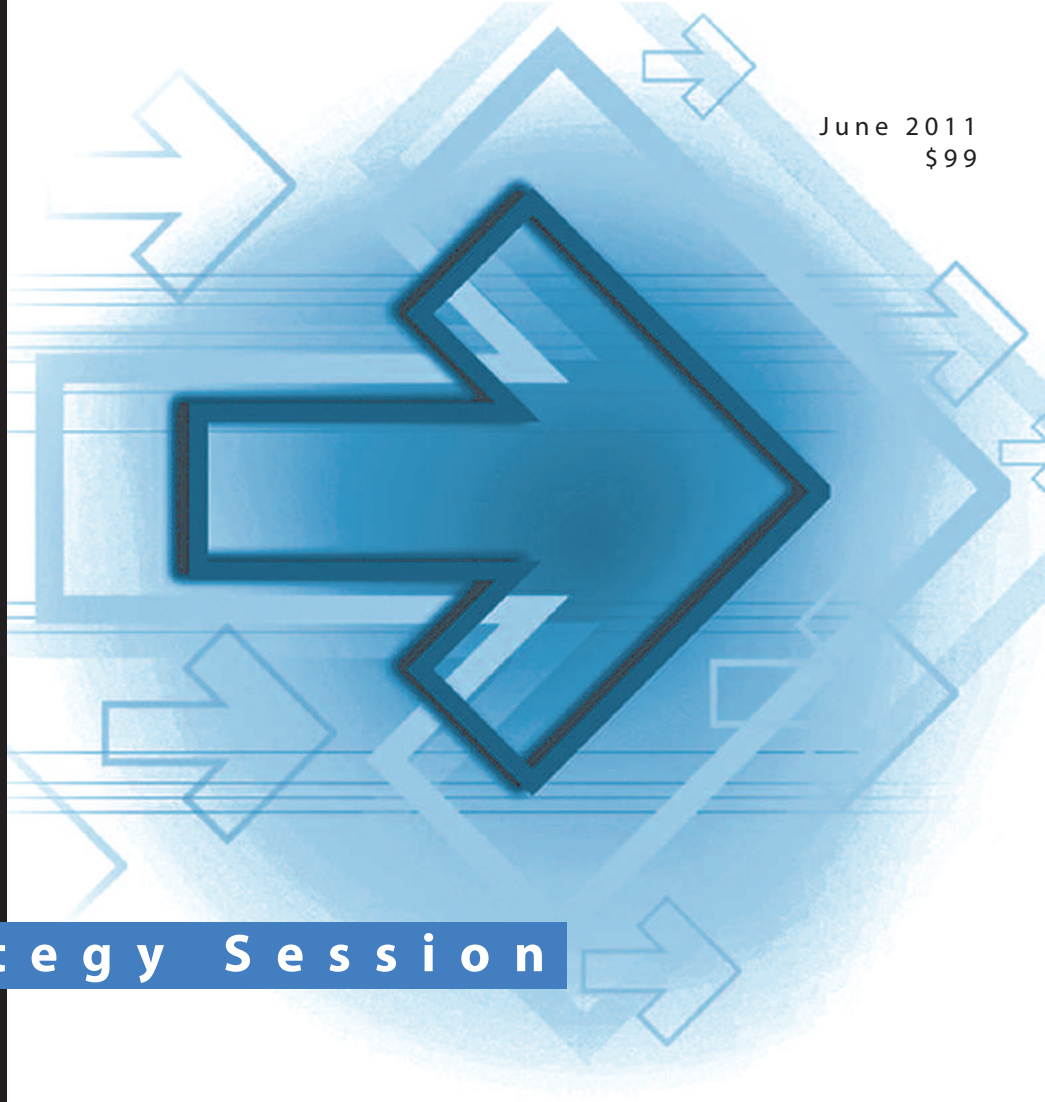

Strategy Session

IPv6 Security: Problem Child Or Opportunity to Improve?

IPv6 advocates have long touted the elimination of NAT and the return to a true peer-to-peer Internet. But IT pros who've come to see NAT as an essential network security element are worried, and they have some questions: What are the security implications of a wide-open, any-to-any connection model? Do some features and characteristics of IPv6 make it more secure? For that matter, are the security models we've used for most of our careers still valid for IPv6? We'll investigate.

By Jeff Doyle

Report ID: S2860611

TABLE OF CONTENTS

ABOUT US | ***InformationWeek Analytics'*** experienced analysts arm business technology decision-makers with real-world perspective based on a combination of qualitative and quantitative research, business and technology assessment and planning tools, and technology adoption Strategy Session gleaned from experience.

If you'd like to contact us, write to managing director **Art Wittmann** at *awittmann@techweb.com,* content director **Lorna Garey** at *lgarey@techweb.com* and research managing editor **Heather Vallis** at *hvallis@techweb.com.* Find all of our reports at *www.analytics.informationweek.com.*

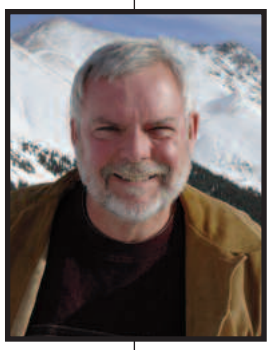

Jeff Doyle
Jeff Doyle and
Associates

Jeff Doyle specializes in IP routing protocols, MPLS and IPv6 and has designed or assisted in the design of large-scale IP service provider networks throughout North America, Europe, Japan, Korea, Singapore and the People's Republic of China. Over the past few years he has had extensive experience helping in the deployment of IPv6 in large networks around the world.

Jeff is the author of *CCIE Professional Development: Routing TCP/IP, Volumes I and II; OSPF and IS-IS: Choosing an IGP for Large-Scale Networks;* and is an editor and contributing author of *Juniper Networks Routers: The Complete Reference.* Jeff has presented numerous corporate seminars, and has spoken at Interop, NANOG, JANOG, APRICOT, IEEE/OWRA, QUESTnet and IPv6 Forum conferences worldwide. He is one of the founders of the Rocky Mountain IPv6 Task Force and is an IPv6 Forum Fellow.

IPv6's Problem Child: Security

We see security as a major stumbling block in enterprise migrations from IPv4 to IPv6. For starters, the code is mostly untested, and too few of our current network security products support IPv6, something the black hat community is banking on. And there's widespread confusion—the idea that some characteristics of IPv6 make it intrinsically more secure than IPv4 is, sadly, just plain false, and information security teams are largely in the dark on how to help their companies safely transition.

Consider the "NAT-bashing" slide, a fixture in IPv6 presentations. Presenters gleefully list all the reasons why Network Address Translation is evil and explain how an abundance of IPv6 addresses will finally let us eliminate what was always supposed to be a temporary address-conservation kludge and get back to a true peer-to-peer Internet. High-fives ensue.

Except, IT security professionals don't share the enthusiasm. Let's face it, IPv6 idealists can wave their fists at NAT all they want, but there are legitimate concerns about losing the ability to shield internal address schemes. No wonder, then, that among the *InformationWeek Analytics* sessions we presented at 2011 spring Interop conference, by far the most popular was our program on IPv6 with a focus on security. A quick show of hands revealed that most attendees are

Figure 1

Key Considerations for IPv6 Security

- Security means more than firewalls and ACLs. Ensure *all* your IP systems are ready for IPv6.
- Networking systems may process IPv6 in software. This is an opportunity for CPU depletion attacks.
- Many modern operating systems enable IPv6 by default. Do you know everywhere these OSes reside and how to secure them?
- IPv6 code is new. There have been security holes, and there will be more, so make sure you monitor. Remember, black hats are studying IPv6 closely.
- There are three legs to the security stool: tools, people and processes/policies. Budget sufficient time and money to update procedures and train your people.

still in the planning stages of their deployments, par for the course among companies we work with.

The good news is that they can take advantage of the lessons learned by telecom carriers and ISPs, which tend to be well along the road to IPv6. And while many conventional enterprise security models will need to change to work in a v6 network, the upgrade also provides opportunities for improvement—possibly even an outright reimagining of your security strategy.

A prime opportunity to see how all this works in the real world is World IPv6 Day, set for June 8, 2011. This is a milestone in the transition from IPv4 to IPv6, when companies including Akamai, Facebook, Google and Yahoo will offer their content over IPv6 for 24 hours. The event will provide us with valuable data on connectivity, ranging from simple system misbehavior to the amount of user traffic that will switch to IPv6 when content is available over the 64-bit version of IP. We're also fielding our first *InformationWeek Analytics* IPv6 Survey to see where our readers are on the migration curve. We'll share our results in an upcoming report.

The Perimeter Problem

One thing that quickly becomes clear when rolling out IPv6 is that network systems themselves are the easy part of the project, so much so that it's become accepted wisdom to start a deployment in the center and work outward. Difficulties present themselves in greater numbers as you make your way toward the network edge, where users are connected to services.

Envision a "core-to-edge" deployment strategy with your IPv6-enabled network in the middle, surrounded by concentric perimeters of services. Closest to the core are the services essential to the fundamental operation of the network: DNS, DHCP, NTP and the like. Around that perimeter are the services necessary to both manage the network and provide support services: Configuration management, change policy enforcement, monitoring, alarming and logging.

The outside perimeter is your security bulwark: Firewalls, ACLs, IDS/IPS and the policies enforced by them. The security perimeter not only protects the systems it surrounds, it also protects the users and services communicating across the network.

The implied order in which systems are tackled under this model—core to support systems to management to security—reflects the reality that security is often the last consideration in an

IPv6 deployment. And this is true in the inverse as well—while we've long had a wide variety of IPv6-capable routers and network support systems to choose from, security vendors have lagged the industry in adding IPv6 capabilities. If your company lists support for IPv6 among the must-have criteria when purchasing new security gear, you're ahead of the game—and likely frustrated that there isn't more available. Until recently, for example, relatively few firewalls had useful IPv6 capabilities, and there are still significant limitations.

The reason security is so often a latecomer in both deployment plans and in products is the same: It's a daunting challenge for all involved. In the case of vendors, they know IPv6 is coming, but they cannot make an internal business case for dedicating engineering resources until

Figure 2

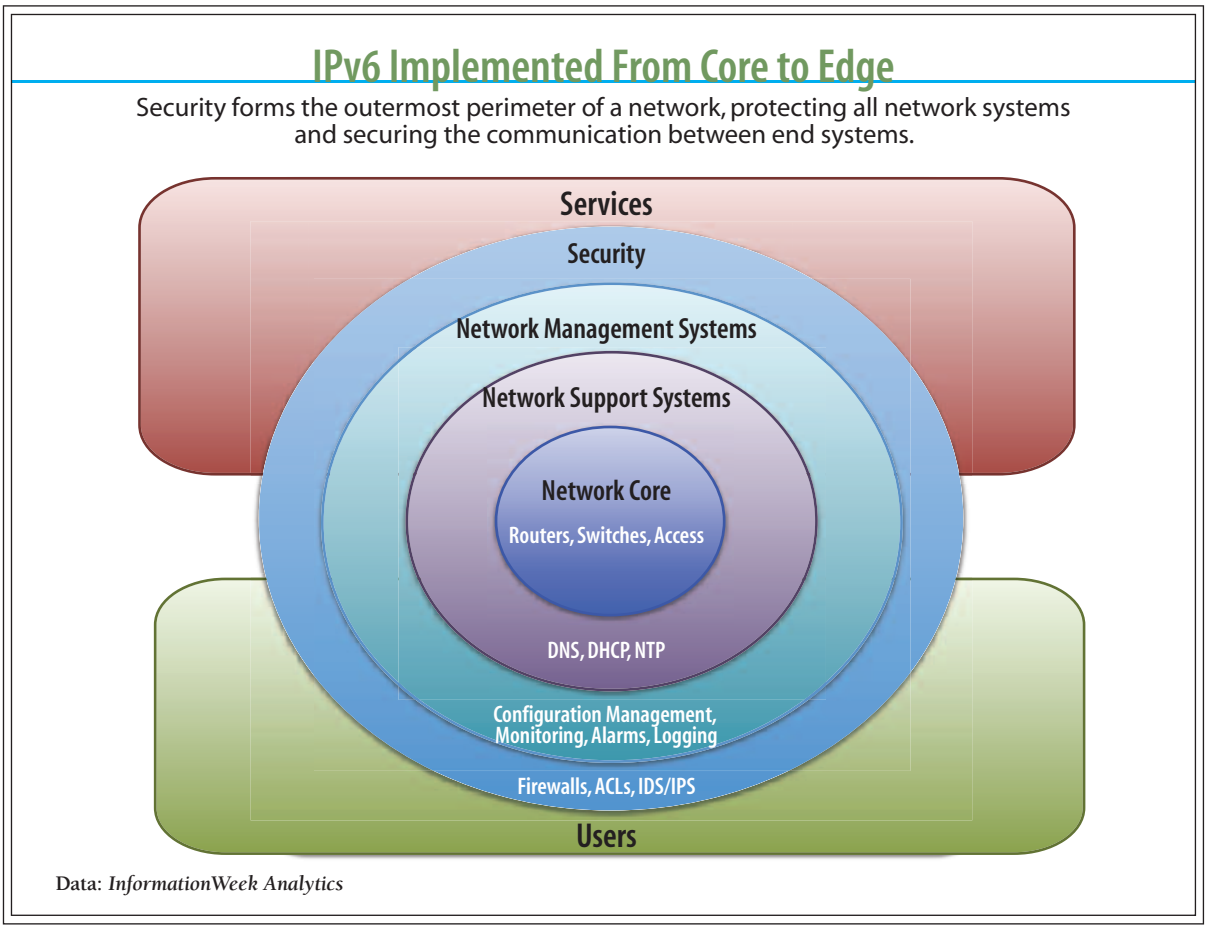

customers ask for support—and in a way that is connected to tangible sales. It's only been a few years that customers have gotten far enough into their IPv6 planning to start demanding that their security vendors add support. Factor in a typical 12-month development cycle, and we get a mixed bag, with most security vendors only recently providing decent IPv6 capabilities. A few are far along, a few haven't even begun.

Yet deploying any system you cannot properly secure is folly.

Security Mythology

An early and often-repeated myth about IPv6 is that it is more secure than IPv4. The misconception springs from the fact that support for authentication and encryption capabilities are integral to IPv6. But a capability called for in a specification does *not* necessarily translate into a real-world capability. The reality is that few IPv6 implementations provide "built-in" authentication and encryption, and end-to-end IPv6 sessions are not automatically secured. In fact, IPSec encryption has long been available with IPv4, albeit as an add on rather than an integral part of the base protocol, and it has not increased the security of IPv4. IPSec is seldom used for the simple reason that end-to-end encryption in and of itself is problematic enough—no matter whether it's in an IPv4 or an IPv6 network. The impediments to end-to-end security must be surmounted independently of the IP version used; when and if they are, the IP version won't matter.

Another piece of the IPv6 security myth stems from characteristics of the protocol that, while not directly security-related, do have security implications. For example, you'll regularly hear that IPv6's huge address space makes it immune to port scanning. By far the most common IPv6 subnet prefix is 64 bits, which supports up to 1.8×10^{19} individual addresses. Assuming a port scanner could "hit" one address per second, a scan of the entire address space of a 64-bit subnet would take over 584 billion years. That's an impressive stat, but it ignores the fact that smart subnet spies are already selective about the ports they scan and predictive about the IP addresses they target. Yes, port scanning is more problematic on a typical IPv6 subnet—both for ill-intentioned snoops and for your own security personnel—compared with most any IPv4 subnet. But stating that IPv6 is immune to scanning is just wrong.

It's also naïve to assume that there's safety in ignorance, that because most network engineers are not yet familiar with IPv6, the bad guys won't be either. In fact, the opposite is true. There

are black hats out there who see IPv6 as a once-in-a-lifetime opportunity. Security systems might leave holes not available in IPv4 networks, operating systems might have new vulnerabilities, and security personnel might overlook an exploitable IPv6 feature. A quick Web search reveals that there is already a hardware store full of public-domain tools available for the attacker to add to his box, both IPv6-enabled versions of familiar IPv4 tools and a few specifically created to attack IPv6 networks. Fewer than 40% of the 1,084 respondents to our 2011 *InformationWeek Analytics* Strategic Security Survey have formal SDLC processes, so odds are you can't depend on flaws being discovered before code is deployed.

Still, it's up to you to ensure that your systems are protected and that your security personnel are educated. The best place to start identifying potential vulnerabilities is to understand some key differences between IPv6 and IPv4. Here are seven areas to watch:

Neighbor Discovery Protocol (NDP): This protocol is essential to the operation of IPv6. It replaces several functions performed by separate protocols under IPv4, such as ARP, router dis-

Figure 3

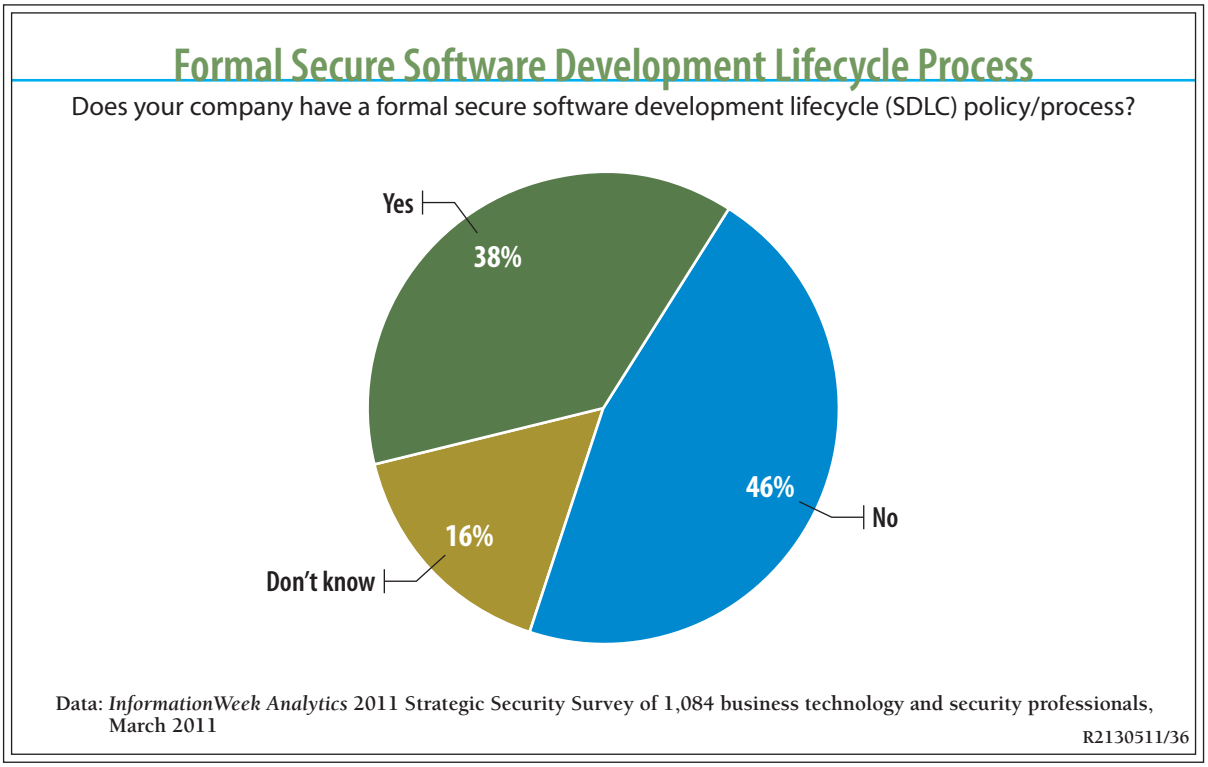

covery and redirects; it also enables new functions for IPv6, such as Stateless Address Autoconfiguration, Duplicate Address Detection and Neighbor Unreachability Detection. But NDP presents a range of exploits for an attacker who can gain local access to a subnet. For instance:

- He can use **Duplicate Address Detection** mechanisms to convince neighbors on the subnet that their addresses are not legitimate.
- He can use spoofed **NDP Neighbor Solicitation** and **Neighbor Advertisement** messages to redirect traffic away from legitimate hosts for either simple denial of service or to steal data.
- He can use spoofed **NDP Router Advertisement** messages to redirect traffic away from a legitimate default gateway or cause hosts to think the default gateway has timed out.
- He can use spoofed **NDP Router Advertisement** messages to cause all hosts on a subnet to change their subnet addresses or their subnet parameters.
- He can manipulate flags in spoofed **Router Advertisement** messages to cause hosts to look for a fake DHCPv6 server—from which he can change addresses or give the hosts illegitimate DNS addresses.

There are tools to help protect NDP, such as Cisco's RA-Guard feature or an authenticated version of NDP called Secure Neighbor Discovery (SEND), but availability and use of such protections are limited.

ICMPv6: ICMP is a favorite attack vector for denial-of-service and CPU attacks, and guarding against ICMP message floods is a fundamental security best practice. But IPv6 is more dependent on ICMP than is IPv4; simply blocking all ICMP messages at security checkpoints can break some IPv6 functions. It's therefore important to understand those dependencies and how to adapt your security policies to them by, for example, rate-limiting certain ICMPv6 messages rather than blocking them entirely.

Fragmentation: Fragmentation attacks are another old favorite that might be given a new spin by IPv6. Unlike with IPv4, in IPv6, routers do not fragment packets. Instead, the spec requires the originating end system either to test the MTUs along a path to a destination and fragment accordingly or to fragment all packets exceeding 1280 bytes—the smallest MTU an IPv6 inter-

face is allowed to support. IPv6 fragmentation attacks can look much like they do in IPv4, with overlapping fragment offsets, incomplete fragments, fragmentation buffer overloads and firewall bypasses, but immature IPv6 operating system code can increase vulnerabilities.

Extension Headers: IPv6 economizes its default header by eliminating optional fields, such as IPv4's Flags, Fragment Offset and IP Options. Instead, when an optional capability, such as fragmentation, source routing, encryption or authentication, is required, an applicable extension header is inserted between the default IPv6 header and the packet payload. The default IPv6 header is more efficient because it is always the same size, and the ability to add a capability in the future by defining a new extension header makes the protocol adaptable. But attackers can abuse extension headers in a number of ways. For example:

- Adding unknown extension header types or misordering the headers can confuse or crash a poorly written operating system.
- A cleverly written extension header might be used to create a covert channel into a network.
- Long chains of extension headers can be used to impact firewall or end system performance.

Figure 4

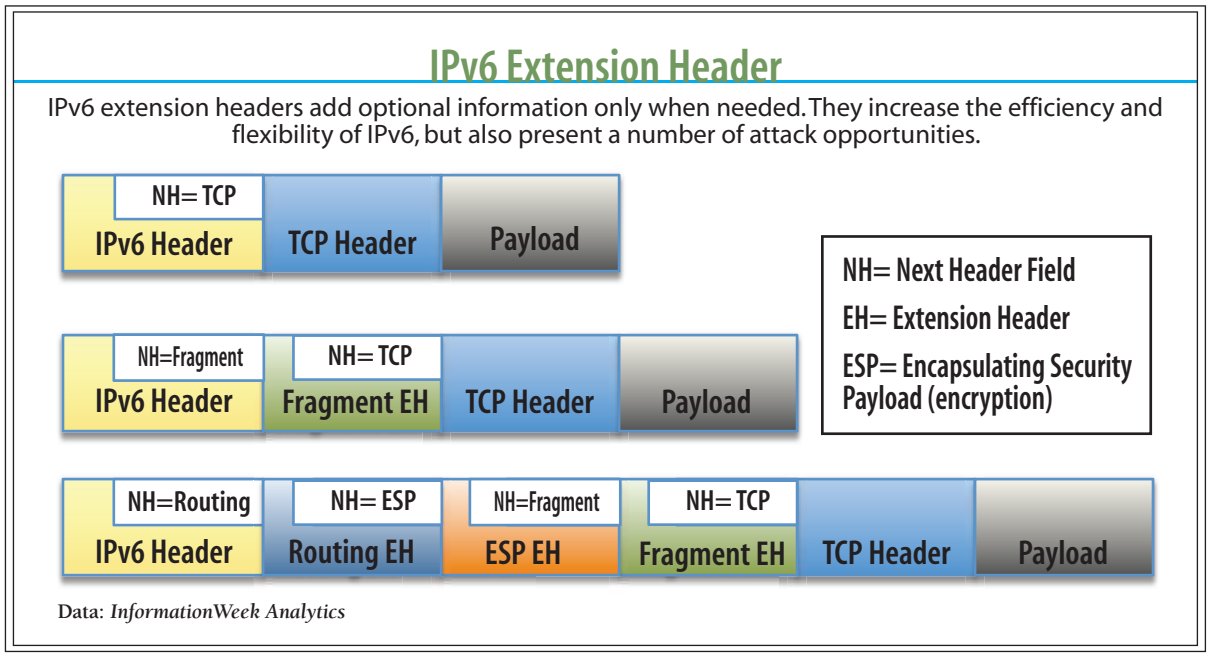

And finally, extension headers present a problem that has nothing to do with malicious activity: Firewalls or ACLs looking for fields in an upper-layer header can pay a performance penalty when they have to sort through a string of extension headers to find the beginning of the upper-layer payload. Cisco provides a good overview of extension headers and the performance problems they might present.

Flow Labels: The Flow Label field is the only field in the default IPv6 header that has no analogous function in the IPv4 header. Intended to enable efficient processing of microflows for improved service classification, mainstream network systems do not yet use it. An intentionally miswritten flow label value could potentially be used to create a covert channel.

Automatic Tunnels: Automatic tunneling mechanisms such as 6to4, ISATAP and Teredo, supported by most host operating systems, are used to create IPv6 connectivity over an IPv4-only network or network segment. But they can also be used to create an unsecured channel into and out of a network, and most lack a means of authentication. Some operating systems not only support multiple automatic tunneling mechanisms but enable them by default, allowing a careless user or a compromised system to bypass some security protections. And as already discussed, once a single host on a subnet is compromised, there is a whole menu of mischief from which to choose.

Large-Scale NAT: Also called Carrier-Grade NAT (CGN), LSN is not a part of the IPv6 specification, but it is often associated with IPv6 transitional architectures. LSNs allow network operators to centralize their public IPv4 address pools, to extend their useful lives by multiplexing more IPv4 flows to each address. These centralized NATs—often single points of failure for tens of thousands of end systems—present an attractive target for CPU or address pool depletion attacks.

Beyond Black Hats

Security goes beyond protecting your network from nefarious attackers. You must also guard against actions or effects that can bring down portions of the network as effectively as any DoS attack. In the case of IPv6, there are two dissimilar nonmalicious threats to watch for.

First, don't assume that because you get a certain performance level from a network system running IPv4 that you will realize the same performance when you add IPv6. A router that

processes and forwards IPv4 packets in hardware might process IPv6 packets in software. A firewall's CPU might slow significantly when it processes IPv6, particularly if extension headers are involved. Talk to your vendors about the possible performance impacts of IPv6, and when possible, test performance yourself in a lab.

The other major nonmalicious threat to your IPv6 network—in fact, the most serious threat most enterprises will face—is a lack of training. From the very different address format (do your operators know how to work with hexadecimal numbers?) to the key protocol differences between IPv4 and IPv6, to differences in configuration commands, your network operators and engineers need to be prepared *before* you turn IPv6 over to them. Budget for training during your earliest planning efforts, from formal classroom instruction to self-guided lessons to informal brown-bag sessions. Then, involve your engineering staff in predeployment lab testing, and have a strategy in place to deploy IPv6 incrementally, with a set period to evaluate any unforeseen effects before moving to the next stage. This is the best way to mitigate the operational risks associated with deploying any new technology.

RFC 4942 is a good place to start studying the security implications of IPv6.

New Code, New Bugs

IPv6 implementations almost always mean running new code or code that has not yet undergone extensive production vetting. A router vendor might have supported OSPFv2 for almost 20 years, but OSPFv3 for IPv6 is still going to be new—and possibly buggy—code. Has your firewall vendor added IPv6 support within the last couple of years? Then assume there are surprises waiting for you in the implementation. This is not an indictment of sloppy development work; it is a fact that no vendor can catch all of the bugs in its code. We all depend on extensive production deployments to reveal problems. Yet worldwide, IPv6 is still in its early stages of use, meaning that even IPv6 implementations that were written years ago may be getting their first large-scale field tests.

The more you're aware of the risk of new code, the better prepared you are to test and verify before putting the code into production.

Even standards bodies are occasionally guilty of overlooking security risks. Two infamous examples of early oversights in IETF specifications were an IPv6 source routing vulnerability

that opened the possibility of amplification attacks and firewall bypasses, and an ICMPv6 vulnerability that allowed ping-pong attacks on point-to-point links. Both vulnerabilities were well-known in IPv4 and had long been corrected in earlier standards, but were simply overlooked in initial IPv6 specifications.

And while these mistakes have been corrected by new RFCs, you need to assume that some operating systems in your network have been upgraded or patched to incorporate the new code—which brings us right back to awareness, vendor communication and testing. Be sure you're aware of identified IPv6 security vulnerabilities and their fixes, talk to your vendors to ensure you have their latest patches, and test before deployment whenever possible. Monitoring IETF IPv6 working groups and NANOG discussions can help, although it is time-consuming. US-CERT (www.us-cert.gov/index.html) can be a good single source of information and is probably already monitored by your security team.

Don't Panic

This litany of security concerns might make you nervous about IPv6, but it shouldn't. The great majority of IPv6 security issues differ little, if at all, from IPv4. And in reality, the most dangerous security threats are to your data and your applications, not to your networking protocols. Common sense and information are, as with all facets of network security, your friends.

The best place to start is with a good inventory and assessment that includes:

- Security systems and components;
- Overall security architecture;
- Security policies and practices; and
- Team experience and knowledge.

Then look at your IPv6 deployment plan. How are your security systems and practices affected? What new requirements do your plans impose on them? From here, create a security plan and a detailed requirements list. If there are capabilities missing in your network, talk to your vendors. Understand their development roadmaps, and be sure they understand your plans.

If there isn't a good match within the timeframes you need, look at what other vendors have available and weigh risks, costs and goals. But *don't* compromise. Don't deploy something you cannot secure, and don't leave security to the later stages of your planning.

New Opportunities

The gradual transition from IPv4 to IPv6 is a major evolutionary event in IP networking. It's also, like death and taxes, unavoidable unless retirement is in your near-term plans. And although IPv6 presents some new security challenges, none is insurmountable given the right preparation.

In fact, smart CIOs are looking at the transition as an opportunity. Are your security practices, policies and systems all that you want them to be? If not, an IPv6 deployment can be the perfect opportunity to assess your network security situation and improve or replace your current architecture and practices.

The transition to IPv6 is also an opportunity for us as a community to reconsider the way security is practiced in general. Are firewalls and intrusion detection sufficient protection for our networks? All of the respondents to our 2011 Strategic Security Survey use firewalls, and 93% have intrusion detection/prevention systems in place. But physical walls have never by themselves been effective deterrents to invasion by either individuals or armies. Why should we assume that building digital walls around our networks is effective on its own? It's time to consider new ways of securing individual devices and sessions. For example, consider "personalizing" the security of individual devices and individual sessions by moving to a model of end-to-end authentication and encryption (personal privacy), creating "zones of protection" around individual hosts and servers (personal security), and adopting improved algorithms for threat analysis and interdiction (proactive policing).