

Securing IPv6

Author: Ralph Wallace

Executive Summary

The intent of any organization's cybersecurity approach is to ensure the confidentiality, integrity and availability (CIA)* of IT assets and data. [Internet Protocol version 6 \(IPv6\)](#), based on lessons learned by the Internet Engineering Task Force (IETF) for Internet Protocol version 4 (IPv4), is designed to enhance an organization's network operations and provide a more secure environment in the hands of trained and experienced engineers. The [U.S. National Institute of Standards and Technology Guidelines for the Secure Deployment of IPv6 \(NIST 800-119\)](#) is the federal government standard for agencies to use in transitioning to IPv6. Noted in the foreword is the statement that IPv6 can be deployed just as securely as IPv4, but IPv6 has improved opportunities to enhance any organization's cybersecurity.

Appendix A of this document details the limitations of IPv4 and what was evident to the engineers who designed IPv6.

As noted above, IPv6 inherently has cybersecurity benefits, of which the dramatically significant increase in addresses is simply a part. The address structure (at 128 bits) affords a much cleaner and definitive packet structure, and the aggregation aspect supporting end-to-end interoperability throughout a routed network significantly supports Continuous Diagnostics and Mitigation (CDM) transparency objectives. CDM is a federal cybersecurity program operated under the Cybersecurity and Infrastructure Security Agency at the U.S. Department of Homeland Security.

Other elements observed through focused engineering and piloting break the IPv4 paradigm that a network device needs only one IP address to function, allowing multiple network interfaces and multiple subnets accessed by a single platform. (A recent open source National Security Agency circular notes this distinction for cybersecurity awareness and defense.) This directly supports network segmentation – a key element of the latest federal focus on Zero Trust Architecture – and increases the availability of various data streams at the network layer.

However, IP is still IP, and the devices in place today defending our IPv4 enterprise are still necessary to defend the IPv6 construct. The [defense-in-depth](#) architecture is an industry best practice and is implemented in production enterprises in several federal agencies as a federal IPv6 Task Force transition best practice. The approach is discussed in Scott Hogg's and Eric Vyncke's seminal Cisco Press book [IPv6 Security](#), as well as in [NIST 800-119](#), and encompasses a layered defensive posture. This architecture is in place already in many private and public enterprises but needs to be tooled to IPv6 specifics, and devices need to be appropriately configured.

*This term was coined by the U.S. Department of Defense in the 1980s.


When IPv6 was first suggested for implementation world-wide, many advocates – in an attempt to start the momentum of change – stated that IPv6 “is more secure.” The truth is that IPv6 can be more secure than an IPv4 enterprise if it is designed correctly and the devices are correctly configured. But what is “secure?” Someone can tighten an IPv4 network to be tighter than a drum. IPv6 makes it easier to tighten that drum.

Note: This white paper does not address performance and interoperability enhancements of IPv6 over IPv4, but aspects of end-to-end connectivity and address aggregation support this.

Inherent Opportunities

IPv6 addressing functionality and scoping facilitates the following:

- End-to-end transparency, due to no further need to use network address translation (NAT) or classless inter-domain routing (CIDR). This affords each [node](#) to be observed directly under a CDM environment. Removing NAT also:
 - Reduces application complexity (less code = more secure)
 - Reduces complexity of security devices
 - Eliminates fragmentation processing
 - Eliminates the need for everything to go through port 80/443
 - Improves forensics (ability to determine and define the source of an attack (CDM))
- Multiple subnets to be available on the same network interface. Each interface could hold up to 10 IPv6 global unicast addresses, effectively creating segmentation of multiple data paths and control plane packets. If one subnet is compromised, there is no need to shut down all subnets. Access can be managed with [Internet protocol security \(IPSec\)](#) with a separate [X.509 certificate](#) per subnet, allowing the administrator to issue one [certificate revocation list \(CRL\)](#) to stop traffic without impacting any other subnet attached to that interface.
- Significant increase in reconnaissance by attackers defense ([RFC 5157: IPv6 Implications for Network Scanning](#)). The much larger default 64-bit subnet address space of IPv6 makes traditional network (port) scanning techniques used by certain network worms or scanning tools ineffective. In-band or out-of-band network reconnaissance is typically the first step an attacker takes to identify assets to exploit. Reconnaissance attacks in an IPv6 environment differ dramatically from current IPv4 environments. Due to the size of IPv6 subnets, traditional IPv4 scanning techniques that would normally take seconds could take years on a properly designed IPv6 network. The figure below provides the time required to conduct the scanning.



- No need for dedicated network broadcasts. Specific language is in the request for comments (RFCs) to avoid Internet control message protocol version 6 (ICMPv6) broadcast amplification attacks.
- IPSec embedded in the IPv6 Address structure, allowing each respective host on the subnet to use the authentication header and/or the encapsulating payload.
- Neighbor discovery, the resolution of link layer (layer 2) addresses of each node on the subnetwork. This is restricted only to the respective subnet and is responsible for:
 - Determining the link layer addresses of other nodes
 - Maintaining L2 reachability state to neighbors
 - Finding available routers (and then globally routable IPv6 addresses)
 - Addressing configuration
 - Preventing duplicate IPv6 addresses

Path maximum transmission unit discovery ([RFC 8201](#)) implemented with ICMPv6 ([RFC 4443](#)) facilitates the following:

- Fragmentation only conducted at source host. Overlapping fragments are not allowed and must be filtered. Devices must drop reassembled packets that are less than 1,280 bytes and/or take too long to be re-assembled.

Note: ICMPv6 signatures should be incorporated into the intrusion detection system/intrusion prevention system (IDS/IPS) as well as Multicast Listener Discovery to identify nodes joining the network.

Architected and Configured Opportunities

Several security tools familiar to the IPv4 environment must be used for IPv6. Incorporating the defense-in-depth architecture, the configuration of each element has significantly different aspects due to the design differences between IPv6 and IPv4. This approach specifically establishes the perimeter, infrastructure and hosts as the defined layers of depth.

Perimeter

- **Enterprise “edge” routers:** [access control lists](#) (ACL). There are two types of ACLs:
 1. Traditional ACLs are supported on outbound and/or inbound traffic on layer 3 interfaces. IPv6 ACLs apply only to routed IPv6 packets.
 2. Port ACLs are supported only on inbound traffic on layer 2 interfaces. Port ACLs are applied to all packets entering the interface and may be configured to match only IPv6-specific packets.

ACLs must be complementary within the cybersecurity architecture supporting redundancy for all cybersecurity architecture elements, such as perimeter and endpoint firewalls.
- **Firewalls.** Because of the larger address space there is no need to use NAT with IPv6. IPv6 does not require end-to-end connectivity but facilitates end-to-end addressability when [global unicast addresses](#) are used. IPv6 network firewalls must support, for example, extension header chaining/processing.
 - **Perimeter.** The perimeter has distinct inbound and outbound rule sets for IPv6 and must also, in a dual stack environment, modify the IPv4 portion to defend against unwanted IPv6 tunnels in the IPv4 packets.
 - **Endpoint.** Each host must be prepared to defend the platform from attacks that may be attempted from within the enterprise.

- **Proxies.** A proxy is an intermediary situated between a requestor and responder of a transaction. A proxy that is in front of a group of origin servers, known as a reverse proxy or surrogate, may offer load balancing capability and hides the identities of those servers. Proxies provide many other types of services including user authentication, connection acceleration, redirect, request and response filtering, access logging, translation and transcoding, virus scanning and spyware removal.

For example, a proxy can accelerate secure socket layer (SSL) connections by offloading computation intensive cryptographic operations to the built-in crypto hardware. An IPv6 proxy needs to plug security holes created by the covert communications channels such as HTTP tunnels and secure port forwarders. A typical firewall, for instance, tends to open port 80 to allow for HTTP traffic. Spyware and Trojan horses punch through firewalls by exploiting this common default rule. An intelligent IPv6 proxy must examine the HTTP POST and CONNECT requests and appropriately determine whether to allow or deny such traffic according to the set policy rules.

- **IDS/IPS.** IDS/IPSs play an important role in looking for network exploits and, in the case of the IPS, acting based on rules such as shutting off access. The extension header structure for IPv6 and the change in fragmentation, along with a multitude of tunneling options, are several aspects of IPv6 that impact IDS/IPS.
- **Deep Packet Inspection (DPI).** DPI can be an important tool to gain visibility into IPv6 tunneled packets. DPI is often a component of a router, firewall, IDS/IPS – or even quality of service – implementation but can also exist as a separate function/device.

Your IDS/IPS/DPI must:

- Identify/block tunnels (Teredo, Proto-41, GRE, etc.)
- Include multiple levels of encapsulation
- Detect link-local attacks
- Detect domain name system queries
 - A/AAAA over IPv4/IPv6
 - Perhaps block AAAA (Quad-A DNS resource record) queries when they are not expected
 - Detect known IPv6 vulnerabilities
 - Detect firewall misconfiguration and unexpected protocols

Infrastructure

- **Routers.** Router advertisements are essential for hosts to determine where they are in the enterprise. Therefore, these advertisements are targets for “man-in-the-middle” attacks, spoofing and assuming the main router’s identity. To guard against this, the implementation of [RFC 6105: IPv6 Router Advertisement Guard](#) is strongly advised (following [RFC 7113: Implementation Advice for IPv6 Router Advertisement Guard](#) and implementation of ACLs, as noted earlier).
- **Domain Name Server (DNS).** DNS is also a target for spoofing. The DNS security extensions described in [RFC 4034: Resource Records for the DNS Security Extensions](#) and [RFC 4035: Protocol Modifications for the DNS Security Extensions](#), and implemented per IETF best current practice [RFC 9364: DNS Security Extensions \(DNSSEC\)](#), are now available for internal enterprise implementation, although they were first proposed for external-facing websites.
- **Dynamic Host Control Protocol version 6 (DHCPv6).** DHCPv6 is a target for spoofing as well. Protection against these spoofing attempts is [RFC 7610: DHCPv6-Shield: Protecting against Rogue DHCPv6 Servers](#), which is a valuable mechanism for protecting hosts connected to a switched network against rogue DHCPv6 servers.

Endpoints

- **Firewall.** Each host must be prepared to defend the platform from attacks that may be attempted from within the enterprise. Endpoint firewalls must have a layered ruleset for “least to most restrictive” given the location of their connection (e.g., on-premises, telework, remote location).
- **Operating System.** Each host platform (e.g., client, server, mainframe, BYOD) has security settings internal to the operating system that supports cybersecurity requirements.

What Can Be Done Today: A Checklist

The cybersecurity checklist below is from NIST 800-119, with actions defined according to engineering analysis. (The checks indicate appropriate actions to take; an arrow indicates action to be researched.)

- ✓ Apply an appropriate mix of different types of IPv6 to limit access and knowledge of IPv6-addressed environments.
- ✓ Use automated address management tools to avoid manual entry of IPv6 addresses, which are prone to error because of their length.
- ✓ Develop a granular Internet Control Message Protocol for IPv6 (ICMPv6) filtering policy for the enterprise. Ensure that ICMPv6 messages essential to IPv6 operation are allowed but others are blocked (firewall rule sets, ACLs).
- ✓ Identify capabilities and weaknesses of network protection devices in an IPv6 environment (accomplished and defined with our defense-in-depth security architecture specifying perimeter, infrastructure and host defensive postures).
- ✓ Enable controls that might not have been used in IPv4 due to a lower threat level during initial deployment, implementing default deny access control policies, routing protocol security, etc. (firewall rule sets, ACLs and IPv6 signatures within the Sourcefire IDS/IPS).
- ✓ Pay close attention to the security aspects of transition mechanisms such as tunneling protocols (e.g., firewall rule sets, ACLs and IPv6 signatures within the IDS/IPS).
- ✓ Ensure IPv6 routers, packet filters, firewalls and tunnel endpoints enforce multicast scope boundaries and that multicast listener discovery packets are not inappropriately routable (firewall rule sets, ACLs and IPv6 signatures within the IDS/IPS).
- ✓ Use IPSec to authenticate and provide confidentiality to assets that can be tied to a scalable trust model. An example is access to human resources assets by internal employees that use an organization’s public key infrastructure to establish trust (available for use; design to be created based on appropriate trust model and to be tested).

Appendix A: Limitations of IPv4

IPv4 ([RFC 791](#)) was designed over 30 years ago for a relatively small number of users. At that time, it seemed unlikely that personal computing technology would become as widespread as it is today in the United States and worldwide. The rapid, universal adoption and growth of personal computing technologies, including IP networking, were unforeseen in 1981. At that time, the Internet was used almost exclusively by scholars and researchers, and IPv4’s 4.3 billion theoretically available addresses were considered more than sufficient.

As a result of growing Internet use, IPv4’s address capacity could not meet the demand. In practice, the supply of available IPv4 addresses has been limited since the early 1990s. Previously, an organization could apply for and receive more IPv4 addresses than it could justify. However, because of regulatory advances, IP address allocations are now bound by strict policies that include formal justification to a regional Internet registry.

During the 1990s, address allocation policies, along with address reuse and restriction technologies, were adopted to conserve IPv4 addresses. Technologies widely adopted in response to the constrained supply of IPv4 addresses are NAT [[RFC 3022](#)] and CIDR [[RFC 4632](#)].

NAT essentially makes private IPv4 addresses (also known as non-routable addresses) at least partially functional on the global Internet. Despite their adaptation for other uses, private IPv4 addresses were designed for testing and other non-production purposes and never intended for use on the Internet. Nevertheless, a NAT-capable router positioned at an organization's boundary can connect an entire network of privately addressed nodes within the organization to the Internet via a single routable IP address. This technology saves IPv4 address space because nodes bearing private addresses are essentially on the Internet but do not have globally unique IP addresses.

Nevertheless, this address conservation technology can actually defeat certain aspects of the design intent of IPv4: network layer end-to-end security, peer-to-peer (host-to-host) connectivity and interoperability. A host using private addressing behind a NAT device cannot have a full peer-to-peer relationship with another host via the Internet or backbone enterprise network using globally unique addressing. This is because NAT does not allow communication sessions to be initiated from globally addressed nodes to privately addressed nodes.


NAT traversal technologies are available to work around some of these barriers. They typically work in one of two ways:

1. By maintaining stateful address lookup tables and redirecting inbound traffic to appropriate private addresses
2. By employing application layer gateways that listen for specific port numbers and redirect traffic according to pre-configured parameters.

Neither of these approaches to NAT traversal lends itself to scalability or guarantees compatibility with all forms of NAT, not to mention the efforts put into each of these workarounds. In addition, neither approach lends itself to dynamic configuration when, for example, hosts move or networks are renumbered.

Another limitation of IPv4 is that its design favored interoperability over security and did not contain features that protected the confidentiality, integrity or availability of communications. For example, IPv4 could not cryptographically protect data from eavesdropping or manipulation, and IPv4 did not provide a method for endpoints to authenticate each other.

Over time, the open nature of IPv4 was increasingly a target of exploitation. The multi-path nature of the Internet, which was designed for high availability, also allows multiple attack vectors for a variety of threats. As a response, new technologies were added to IPv4 to provide needed security functionality. With IPv6, these features were designed into the new protocol as mandatory components.



Ralph Wallace is program director and IPv6 lead at Aptive Resources. He also chairs the U.S. IPv6 Council (IPv6 Forum).

To reach Ralph, send an email to ralph.wallace@aptiveresources.com.