IPv6 Security Guidance

Executive summary

Nearly all networked devices use the Internet Protocol (IP) for their communications. IP version 6 (IPv6) is the current version of IP and provides advantages over the legacy IP version 4 (IPv4). Most notably, the IPv4 address space is inadequate to support the increasing number of networked devices requiring routable IP addresses, whereas IPv6 provides a vast address space to meet current and future needs.

While some technologies, such as network infrastructure, are more affected by IPv6 than others, nearly all networked hardware and software are affected in some way as well. As a result, IPv6 has broad impact on cybersecurity that organizations should address with due diligence.

IPv6 security issues are quite similar to those from IPv4. That is, the security methods used with IPv4 should typically be applied to IPv6 with adaptations as required to address the differences with IPv6. Security issues associated with an IPv6 implementation will generally surface in networks that are new to IPv6, or in early phases of the IPv6 transition.

These networks lack maturity in IPv6 configurations and network security tools. More importantly, they lack overall experience by the administrators in the IPv6 protocol. Dual stacked networks (that run both IPv4 and IPv6 simultaneously) have additional security concerns, so further countermeasures are needed to mitigate these risks due to the increased attack surface of having both IPv4 and IPv6.

Introduction

Federal and Department of Defense networks are moving from legacy IPv4 to IPv6-only. During this transition, IPv4 will continue to be used, and many networks will operate dual stack (running both IPv4 and IPv6 protocols simultaneously) as an interim solution toward an IPv6-only end state. However, operating dual stack increases operational burden and the attack surface. System owners and administrators should implement cybersecurity mechanisms on both IP protocols to protect the network.

The network architecture and knowledge of those who configure and manage an IPv6 implementation have a big impact on the overall security of the network. As a result, the actual security posture of an IPv6 implementation can vary.

IPv6 security concerns and recommendations

To get a good start in implementing IPv6 networks and their potential security concerns, NSA recommends the following:

Auto-configuration

Stateless address auto-configuration (SLAAC) is an automatic method to self-assign an IPv6 address to a host. In some cases, such as for important servers, static addresses may be preferred, but allowing devices to automatically self-assign or request an IPv6 address dynamically is easier in most cases. In SLAAC, a host configures its own network address based on a network prefix received from a router. The assigned IPv6 address incorporates media access control (MAC) address information from the network interface and may allow for host identification via interface ID, network interface card, or host vendor. This leads to privacy concerns by linking movements to a specific device and deducing an individual associated with that equipment, as well as exposing the types of equipment used in a network.

NSA recommends assigning addresses to hosts via a Dynamic Host Configuration Protocol version 6 (DHCPv6) server to mitigate the SLAAC privacy issue. Alternatively, this issue can also be mitigated by using a randomly generated interface ID (*RFC 4941 – Privacy Extensions for Stateless Address Auto-configuration in IPv6*) [1] that changes over time, making it difficult to correlate activity while still allowing network defenders requisite visibility.

Automatic tunnels

Tunneling is a transition technique that allows one protocol to be transported, or tunneled, within another protocol. For example, a tunnel can be used to transport IPv6 packets within IPv4 packets. A network might use tunneling for its Internet connection, and some devices or apps might be designed to tunnel IPv6 traffic. Some operating systems will automatically establish an IPv6 tunnel when a client connects to a server, potentially causing an unwanted entry point to the host.

Unless transition tunnels are required, NSA recommends avoiding tunnels to reduce complexity and the attack surface. Configure perimeter security devices to detect and block tunneling protocols that are used as transition methods. In addition, disable tunneling protocols (6to4 [2], ISATAP [3], Teredo [4], etc.) on all devices where possible. Tunneling protocols can be allowed if they are required during a transition, but they should be limited to only approved systems where their usage is well understood and where they are explicitly configured.

Dual stack

A dual stack environment exists when devices run both IPv4 and IPv6 protocols simultaneously. This is a preferred method for staged IPv6 deployment, but it can be more expensive and tends to increase the attack surface. This approach provides a transition method to IPv6 because it allows devices to use IPv6 for communications that support IPv6 while maintaining the ability to use IPv4 for communications that do not support IPv6. As IPv6 deployments increase, a dual stack environment will transition to IPv6-focused operations by increasing the use of IPv6 and decreasing the use of IPv4.

When deploying a dual stack network, organizations should implement IPv6 cybersecurity mechanisms that achieve parity with their IPv4 mechanisms or better. For any security mechanism implemented for IPv4, a corresponding security mechanism should be implemented for IPv6, with the IPv6 mechanism addressing any differences for IPv6. [5] For example, firewall rules that filter higher level protocols (such as TCP or UDP) should be applied to both IPv6 and IPv4 protocols. Many modern network security mechanisms support both IPv4 and IPv6, although administrators should verify specific product compatibility. Also, other transition mechanisms, such as tunneling and translation, should be avoided at this step in the transition strategy as they add transport and cybersecurity complexities.

Hosts with multiple IPv6 addresses

Unlike IPv4, multiple network addresses are commonly assigned to an interface in IPv6. Multiple addresses create a wider attack surface than with a single address. Generating filtering rules or access control lists (ACLs) can be a challenge. It also requires firewalls and intermediate security devices to be aware of all of the addresses in order to be effective.

To mitigate this concern, carefully review ACLs to ensure they deny all traffic by default, so only traffic from authorized addresses are permitted through the firewalls and other security devices. Ensure all traffic is logged, and review the logs on a regular basis to ensure the allowed traffic matches the organization's policies.

IPv6 education

A successfully secured IPv6 network requires, at a minimum, a fundamental knowledge of the differences between the IPv4 and IPv6 protocols and how they operate. The lack of this knowledge could lead to IPv6 misconfigurations. Misconfigured IPv6-enabled devices (resulting from an error in the configuration) could introduce vulnerabilities, making the devices more prone to compromise.

Learning the IPv6 protocol and knowing how to configure IPv6 effectively are the most critical things to protect and enhance IPv6 security on a network. NSA recommends ensuring all network administrators have received the proper training and education to adequately administer IPv6 networks.

Additional considerations

While there are convincing reasons to transition from IPv4 to IPv6, security is not the main motivation. Security risks exist in IPv6 and will be encountered, but they should be mitigated with a combination of stringently applied configuration guidance and training for system owners and administrators during the transition. In addition to the potential security issues previously described, what follows is a list of additional considerations to secure IPv6 networks:

Use split domain name system (Split DNS)

The Domain Name System (DNS) has been expanded for IPv6 with a new AAAA record that provides IPv6 addresses in addition to the A record that provides IPv4 addresses.

Therefore, a dual stack DNS implementation may need to support both A and AAAA records. Due to SLAAC and other mechanisms, sensitive information could be included in the AAAA records for internal hosts. Split DNS uses two separate DNS servers created for the same domain, one for the external network and one for the internal network. The goal of split DNS, as opposed to a single DNS, is to increase security and privacy by not inadvertently exposing sensitive information in a DNS record from the internal network to the external network. NSA recommends implementing split DNS, for both IPv4 and IPv6 networks.

Filter IPv6 traffic (boundary protection)

IPv6 traffic should be filtered according to the organization's network policies. A network that has not yet deployed IPv6 should block all IPv6 at the network border, including any IPv6 that is tunneled in IPv4. A network that has deployed IPv6 should only allow IPv6 traffic that is permitted by policy, with ACLs allowing authorized flows and protocols and blocking all others by default. Although the IPv6 filtering policy may be based on an existing IPv4 policy, the IPv6 policy should reflect IPv6-specific issues. In addition, the filtering policy should reflect that Internet Control Message Protocol for IPv6 (ICMPv6) is more fundamental to IPv6 communications than the corresponding ICMP for IPv4. Specific ICMPv6 messages, such as neighbor discovery and router advertisement, may need to be permitted even if the corresponding message in ICMP for IPv4 is blocked. [6] [7]

Protect the local link

IPv6 defines network functions that operate on the local link. This includes link-layer address resolution, router discovery, and stateless auto-configuration of addresses. [8] [9] Compared to IPv4, local-link operations for IPv6 are more complex and provide more attack surface. Therefore, any relevant mitigations (i.e., Router Advertisement (RA) Guard [10] [11] to protect against rogue RA messages, Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Shield [12] to protect against rogue DHCPv6 servers) provided by switches and routers should be considered.

Avoid network address and protocol translation

IPv6-only networks will likely implement translation, such as NAT64/DNS64 (Network Address Translation between IPv6 hosts and IPv4 servers and synthesizing DNS AAAA records from A records) or 464XLAT (translation between IPv4 private addresses, IPv6

addresses, and IPv4 global addresses), to communicate with other networks that do not yet support IPv6. As dual stack and IPv6-only deployments increase, translation use will decrease, and eventually, the translation functions will no longer be used and can be removed.

Other than using NAT64/DNS64 [13] [14] or 464XLAT [15] for IPv6-only networks, address translation should generally not be used. In particular, many IPv4 networks use NAT, specifically NAT44, to translate between internal and external addresses. On the other hand, IPv6 networks should instead use global addresses on all systems that require external communications and non-routable addresses inside the network. If unique local addresses [16] are used on internal systems, any system that requires external communications should also have a global address.

Plan for IPv6 stumbling blocks

As with all network changes, new security issues or variations of existing ones, will arise during the transition to IPv6. As described in this cybersecurity information sheet, addressing the issues up front in IPv6 implementation plans, configuration guidance, and appropriate training of administrators will aid organizations to avoid security pitfalls during the transition and to leverage IPv6 benefits properly.▪

Works cited

[1] Narten, T., Draves, R., Krishna S. (2007), Privacy Extension for Stateless Address Autoconfiguration in IPv6, RFC 7123. https://datatracker.ietf.org/doc/html/rfc4941

[2] Carpenter, B. and Moore, K. (2001), Connection of IPv6 Domains via IPv4 Clouds, RFC 3056. https://datatracker.ietf.org/rfc/rfc3j056.html

[3] Templin, F., Gleeson, T., and Thaler, D. (2008), Intra-Site Automatic Tunnel Addressing Protocol (ISATAP), RFC 5214. https://datatracker.ietf.org/doc/rfc5214

[4] Huitema, C. (2006), Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs), RFC 4380. https://datatracker.ietf.org/doc/html/rfc4380

[5] National Institute of Standards and Technology (NIST) (2010), SP 800-119 Guidelines for the Secure Deployment of IPv6. https://nvlpubs.nist.gov/nistpubs/legacy/sp/nistspecialpublication800-119.pdf

[6] Davies, E. and Mohacsi, J. (2007), Recommendations for Filtering ICMPv6 Messages in Firewalls, RFC 4890. https://datatracker.ietf.org/doc/rfc4890/

[7] Chittimaneni, K., Chown, T., Howard, L., Kuarsingh, V., Pouffary, Y., and Vyncke, E. (2014), Enterprise IPv6 Deployment Guidelines, RFC 7381. https://datatracker.ietf.org/doc/rfc7381/

[8] Narten, T., Nordmark, E., Simpson, W., and Soliman, H. (2007), Neighbor Discovery for IP version 6 (IPv6), RFC 4861. https://datatracker.ietf.org/doc/rfc4861

[9] Thomson, S., Narten, T., and Jinmei, T. (2007), IPv6 Stateless Address Autoconfiguration, RFC 4862. https://datatracker.ietf.org/doc/rfc4862/

[10] Levy-Abegnoli, E., Van de Velde, G., Popoviciu, C., and Mohacsi, J. (2011), IPv6 Router Advertisement Guard, RFC 6105. https://datatracker.ietf.org/doc/html/draft-ietf-v6ops-ra-guard

[11] Gont, F. (2014), Implementation Advice for IPv6 Router Advertisement Guard (RA-Guard), RFC 7113. https://datatracker.ietf.org/doc/rfc7113
[12] Gont, F., Liu, W., and Van de Velde, G. (2015), DHCPv6-Shield: Protecting against Rogue DHCPv6 Servers, RFC 7610. https://datatracker.ietf.org/doc/rfc7610
[13] Bagnulo, M., Matthews, P., and van Beijnum, I. (2011), Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers, RFC 6146. https://datatracker.ietf.org/doc/rfc6146
[14] Bagnulo, M., Sullivan, A., Matthews, P., and van Beijnum, I. (2011), DNS64: DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers, RFC 6147. https://datatracker.ietf.org/doc/rfc6146
[15] Mawatari, M., Kawashima, M., and Byrne, C., 464XLAT: Combination of Stateful and Stateless Translation, RFC 6877, 2013. https://datatracker.ietf.org/doc/rfc6877/
[16] Hinden, R. and Haberman, B., Unique Local IPv6 Unicast Addresses, RFC 4193, 2005. https://www.rfc-editor.org/rfc/rfc4193

Disclaimer of endorsement

Purpose

This document was developed in furtherance of NSA's cybersecurity missions, including its responsibilities to identify and disseminate threats to National Security Systems, Department of Defense, and Defense Industrial Base information systems, and to develop and issue cybersecurity specifications and mitigations. This information may be shared broadly to reach all appropriate stakeholders.

Contact

Report Feedback / General Cybersecurity Inquiries: CybersecurityReports@nsa.gov
Defense Industrial Base Inquiries and Cybersecurity Services: DIB_Defense@cyber.nsa.gov
Media Inquiries / Press Desk: 443-634-0721, MediaRelations@nsa.gov