Communications Security Establishment Canada
Centre de la sécurité des télécommunications Canada

CANADIAN CENTRE FOR
CYBER SECURITY

Security considerations for Internet Protocol version 6

Management

TLP:CLEAR

Canada

Foreword

This is an UNCLASSIFIED publication that has been issued under the authority of the Head of the Canadian Centre for Cyber Security (Cyber Centre). For more information, email, or phone our Contact Centre:

- contact@cyber.gc.ca
- (613) 949-7048 or 1-833-CYBER-88.

Effective date

This publication takes effect on October 10, 2025.

Revision history

Revision	Amendments	Date
1	First release.	October 10, 2025

D97-4/80-003-2025E-PDF
ISBN 978-0-660-78181-5

Overview

Exponential growth in the use of Internet-based technologies to deliver modern business services and applications is linked to the depletion of globally available Internet Protocol version 4 (IPv4) addresses. The Internet Protocol version 6 (IPv6) addressing scheme was designed by the Internet Engineering Task Force (IETF) to replace IPv4, and it offers significantly larger private and public address blocks to adequately support modern enterprise and non-enterprise needs. Deploying IPv6 endpoints alongside existing IPv4 infrastructure is emerging as a common strategy within enterprise networks. While IPv6 offers several security enhancements that IPv4 does not, running dual-stack architectures introduces new risks that must be appropriately managed.

To ensure its service architecture continues to evolve, the Government of Canada (GC) will need to design new network architectures and migrate existing digital infrastructure to support IPv6. As part of this strategy, IPv6-enabled services must be designed to securely co-exist alongside existing IPv4 infrastructure until an IPv6-only enterprise architecture emerges. While introducing IPv6 within GC infrastructure may have little or no direct impact on users and front-end services, GC departments must examine and assess the implications of IPv6 on their business services and security objectives.

This publication highlights critical security considerations for IPv6 deployments within GC networks. GC departments must design transition plans to support IPv6 addressing while ensuring operational and security risks are mitigated.

Table of contents

List of tables

1 Introduction

The GC relies on digital, inter-networked systems for delivering essential services to Canadians. Networking technologies continuously evolve due to the requirements of the digital infrastructure needed to support modern service connectivity. While the average Canadian user may not understand which Internet Protocol (IP) stack supports their services, the expectation is that GC digital service infrastructure should be able to process service requests from IPv6-enabled or IPv4-enabled devices. As GC networks and services are built to support the IPv6 technology stack, key stakeholders must assess the potential risks and impact of adopting the IPv6 protocol within the enterprise network, particularly security risks associated with implementing a dual-stack (IPv4 and IPv6) architecture.

Modern systems and applications have varying IPv6 protocol support; sometimes the protocol is available by default while other times, vendor-unique customizations are implemented, which can lead to interoperability challenges. These can expose enterprise networks to considerable security risks, increasing the likelihood for misconfigurations and gaps in security controls.

In 2012, the Treasury Board of Canada Secretariat (TBS) released the GC's IPv6 Adoption Strategy. The strategy was designed to provide a path for transitioning to IPv6-based architectures across GC systems and to ensure seamless and uninterrupted access to GC services within Canada and across the globe. In 2013, TBS released the IPv6 Network Equipment Procurement Guideline as a follow-up to the IPv6 Adoption Strategy. This guideline was meant to help GC departments understand the technical requirements when procuring network equipment (for example, routers, network monitoring devices, proxy servers, firewalls) to ensure IPv6 capabilities are evaluated as part of system procurement processes. However, neither the strategy nor the procurement guidelines adequately address security considerations for IPv6.

While sections of the existing departmental digital architecture may be capable of supporting IPv6, without a secured framework for implementation, security risks may be inadvertently introduced into the enterprise environment. Departments should not assume that enabling support for IPv6 can simply occur by flipping a switch.

According to the IPv6 protocol specification, IPv6 is prioritized over IPv4 by default. Although business enterprise applications may not use the IPv6 protocol, defined specification standards and vendor-implemented default configurations may allow communications with IPv6 link-local addresses. For example, Microsoft[1] does not recommend disabling IPv6 support on Windows operating environments even when not in use. To assess these and other issues, the Cyber Centre recommends that GC departments conduct a review of IPv6 network flows within their environment and address gaps that may exist within their network security monitoring tools before implementation. Enabling IPv6 traffic flows without adequate network visibility monitoring or appropriate network filtering protections may increase the enterprise attack surface and expose the network to additional security risks.

[1] Guidance for configuring IPv6 in Windows for advanced users: https://learn.microsoft.com/en-us/troubleshoot/windows-server/networking/configure-ipv6-in-windows

1.1 Internet Protocol version 6

IP is the primary communications protocol of the Internet; it specifies how network packets are to be transported across network boundaries. IP is a component of the network layer in the Open Systems Interconnection reference model, a framework for organizing communication protocols and sharing information over the public Internet.

IPv6 was designed to replace IPv4, with some enhancements in operational and security functions. Differences exist between IPv6 and IPv4 which have implications on network architecture designs. The IPv6 protocol standard is a 128-bit network addressing scheme, which provides a significantly wider address space compared to IPv4 (which uses a 32-bit network addressing scheme). By default, IPv6 is not backward compatible with IPv4, which may require network administrators to implement changes to existing network architectures.

Table 1 below outlines some of the differences between IPv4 and IPv6.

Table 1: Internet Protocol version 6 compared to Internet Protocol version 4

Protocol components	Internet Protocol version 4	Internet Protocol version 6
Address space and notation	• uses 32-bit address space, and therefore offers a limited address space for private and public use cases • address notation consists of numbers separated by a period, for example, 192.168.0.1	• uses 128-bit address space, and therefore allows up to 2^128 unique network addresses (approximately 340 trillion) • address notation consists of eight colon-separated hexadecimal values, for example, 2001:0DB8:0000:0000:0000:000A:09C0:00B4
Security functions	Protocol does not natively support authentication and security functions	Natively supports authentication, data integrity, and data confidentiality (for example, IP security (IPsec) support)
Types	Supports public and private static addressing to manage networks; however, address space is limited	• supports public routing and private static addressing to manage network devices • typical network address is composed of sections and identifiers (global routing prefix, local subnet identifier and interface identifier)
Address distribution	Autoconfiguration is not supported and would require static or Dynamic Host Configuration Protocol (DHCP) assignment of IP addresses	Allows autoconfiguration (stateless address configurations), easing the need for address assignment by a DHCP server. Autoconfiguration relies on router information for network addresses to access network services

1.2 Internet Protocol version 6 enhancements

The IPv6 specification standard proposed enhancements which were not previously available in IPv4. The following subsections provide additional information on the security enhancements.

1.2.1 Internet Protocol security support

IPsec is a suite of protocols that can be used for authentication, encryption, and integrity protections. While IPsec can be used as a retroactive extension in IPv4, for IPv6 it is supported as part of the standard. Note, IPsec is no longer mandatory in IPv6 as per RFC 8504[2].

1.2.2 Autoconfiguration

Autoconfiguration provides the ability for a node to self-assign its IPv6 network address based on the network prefix information advertised by the router. Stateless address autoconfiguration (SLAAC) is the mechanism by which this can be achieved.

1.2.3 Neighbor discovery

The neighbor discovery (ND) protocol replaces the address resolution protocol used in IPv4 networks, providing cryptographic options to secure discovery messages.

1.2.4 Dynamic host configuration protocol security

The dynamic host configuration protocol for IPv6 (DHCPv6) supports authentication (and encryption) of DHCP messages using IPsec, thus preventing eavesdropping and message intercept attacks.

1.2.5 Extension headers

IPv6 extension headers can be used to improve security, debugging, and management functions.

1.2.6 No broadcast addresses

The IPv6 standard abolished the use of broadcast addresses and adopted multicast addresses as the primary mechanism for group communications.

1.3 Problem statement

As enterprise networks evolve, IPv6 will inevitably need to be supported and managed. New network devices will likely support IPv6 and have it enabled by default, prioritizing its traffic flow in line with the specification standard. Deploying IPv6-enabled devices without proper understanding, adequate monitoring, hardening, and deployment of appropriate mitigation controls will increase the enterprise attack surface and expose the organization to significant risks.

1.4 Threat context

This guidance is intended for systems operating at the UNCLASSIFIED, PROTECTED A, and PROTECTED B levels. In general, the Cyber Centre recommends that GC departments and agencies conduct a threat and risk assessment within the context of

[2] IPv6 Node Requirements: https://datatracker.ietf.org/doc/html/rfc8504

their business needs before partial or full-scale adoption of IPv6. As organizations consider threat sources that may exploit IPv6 vulnerabilities, the Cyber Centre assesses that unsophisticated threat actors (Td3) may target device misconfiguration errors and unintentionally exposed devices to infiltrate networks and maximize their criminal operations. Cybercrime groups and financially motivated cyber threat actors (Td4 and Td5) may target IPv6-related device vulnerabilities and design implementation weaknesses[3]. State-sponsored actors (Td6 and above), in addition to lower-level tactics, may target IPv6 protocol specification weaknesses and system integration vulnerabilities to achieve larger strategic objectives. Mitigations to address state-sponsored advanced threats are considered out of scope for this guidance.

Below are some identified potential threat events (attacks) that could be applicable within IPv6 environments:

1.4.1 Protocol tunneling

Threat actors may encapsulate network packets within another protocol or create multiple tunnels through a network device to evade detection controls. For example, network devices may allow a malicious actor to embed unauthorized IPv6 packets within IPv4 tunnels to evade or bypass network filtering controls. Additionally, threat actors may launch spoofing attacks utilizing tunnel injection techniques, i.e. where a threat actor forges a valid encapsulated packet (based on partial knowledge of the tunnel endpoints and the encapsulation protocol).[4]

1.4.2 Distributed denial-of-service attacks

Threat actors may utilize IPv6 protocol capabilities such as multicast messages or extension headers to launch distributed denial-of-service (DDoS) attacks to overwhelm network defence systems. For example, a threat actor can use spoofed IPv6 link-layer multicast messages to launch a denial-of-service attack on a target source address.

1.4.3 Command and control

Threat actors may leverage IPv6 enhancements (extension headers or others) to embed and communicate control signals or beacons through a compromised network. Globally accessible and larger address space blocks make IPv6 attractive for threat actors to deploy command and control channels.

1.4.4 Network device misconfigurations

Threat actors may exploit network device misconfigurations or inconsistencies when perimeter gateway access control filters are not properly implemented. Threat actors may exploit network devices which expose unconfigured IPv6 interfaces by default to bypass network security controls.

[3] National Cyber Threat Assessment 2025-2026: https://www.cyber.gc.ca/en/guidance/national-cyber-threat-assessment-2025-2026
[4] RFC 9099: Operational Security Considerations for IPv6 Networks: https://www.rfc-editor.org/rfc/rfc9099.html

1.4.5 Network service discovery

IPv6 multicast service discovery messages (e.g. multicast DNS (mDNS)[5] or Link-Local Multicast Name Resolution (LLMNR)[6]) can be spoofed or crafted to redirect endpoints to an attacker-controlled infrastructure. Also, threat actors may leverage IPv6's default protocol capabilities (such as Neighbor Discovery) to support reconnaissance operations (for example, extracting sensitive network device information) which can then be used to target vulnerabilities.

[5] Multicast DNS - https://datatracker.ietf.org/doc/html/rfc6762
[6] Link-Local Multicast Name Resolution (LLMNR) - https://www.rfc-editor.org/rfc/rfc4795.html

2 Security considerations

The security model required for enterprise network architectures to support IPv6-enabled devices is different from traditional IPv4 implementations. This section highlights cyber security considerations and recommended actions to mitigate risks associated with the use of IPv6 within an enterprise network. IPv6 transition plans must consider the impact on business services and the organization's security posture.

2.1 Migration risks

Enabling IPv6 alters an organization's network communications and security monitoring requirements. Accordingly, a systematic approach considering transition plans, interoperability risks and future operational requirements is highly recommended. Organizations may have both IPv4 and IPv6 deployed over their transition period. As such, it is crucial to consider whether existing network security infrastructures can support IPv6. Management must ensure that IPv6 transition plans adhere to change management processes. Security program policies and procedures at the organization level may require updates as necessary.

Management should identify the target objective, transition timelines, and migration paths. Security control policies that manage audit and monitoring, interconnection requirements, device identification and authentication, boundary protection, and managed interfaces may require updates. In general, the Cyber Centre recommends using the risk management framework detailed in IT security risk management: A lifecycle approach (ITSG-33) to identify and manage related information system security risks.

2.2 Procurement and testing

IPv6 transition and implementation plans must be aligned with the organization's procurement strategy. Procurement of assets with networking functions should be assessed for IPv6 support. The National Institute of Standards and Technology (NIST) and the University of New Hampshire (UNH) InterOperability Laboratory have developed an assessment and testing program that can assist with functional evaluation of IPv6 products. This program maintains a product registry of IPv6 devices and applications that have been tested against the technical requirements of the United States Government IPv6 standards profile (USGv6-r1 Profile)[7] for performance and conformance. The Cyber Centre recommends that organizations consider products on the USGv6 program registry as part of their procurement strategy. Organizations should review a product's Supplier Declaration of Conformity (SDoC), which documents IPv6 capability claims. Additionally, organizations should test the network infrastructure's capability to support IPv6-only deployment scenarios.

2.3 Target architecture

The target architecture for the adoption of IPv6 must fall within an acceptable level of residual risk (risk tolerance) for the organization. The Cyber Centre recommends a target architecture plan that ultimately leads to an IPv6-only network infrastructure end-state. While dual-stack architectures (IPv4/IPv6) might be an obvious transition choice, the Cyber Centre

[7] United States Government (USGv6-r1) Profile: https://www.nist.gov/programs-projects/usgv6-program/usgv6-revision-1

recommends designing the transition plan with the goal of an IPv6-only end-state architecture. A single-stack (IPv6-only) end-state architecture simplifies network management and security monitoring, as well as a reduction in the overall operational costs.

2.4 Legacy applications

Legacy applications may lack native support for IPv6, making them incapable of processing IPv6 packet data. This can be particularly complicated with critical business applications with no mechanisms to support IPv6. When IPv6 is enabled, legacy applications or security controls that rely on hard-coded IPv4 addresses as hostnames may be impacted. If adequate traffic translation mechanisms are not implemented, IPv6-only endpoints may be prevented from connecting to services that are only IPv4-aware and vice-versa. The Cyber Centre recommends that organizations assess the impact of their transition plans on their software applications.

The Happy Eyeballs[8] Version 2: Better Connectivity Using Concurrency algorithm is an IETF-proposed standard for managing how system applications can initiate and process asynchronous Domain Name System (DNS) queries on dual-stack hosts. The algorithm allows web applications to switch seamlessly between IPv4 and IPv6 networks. Network administrators should therefore test business applications for IPv6-based capabilities. While the Happy Eyeballs algorithm offers the benefit of managing switches between IPv4 and IPv6, it may also mask network problems. Hence, successfully connecting to an application may not be an indication of a clean bill of health on the IPv4 or IPv6 networks in a dual-stack environment.

2.5 Unauthorized tunnels

Organizations should implement network security controls to detect and block the use of unauthorized IPv6 transition tunnels. Transition tunnels are techniques used to transport IPv6 packets over IPv4 network infrastructure. IPv6 tunnels can be manual or automatic tunnels, such as those provided by Teredo, 6to4[9], or Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)[10]. Teredo is an automatic tunneling protocol designed by Microsoft, and it uses User Datagram Protocol (UDP) to tunnel IPv6 packets over IPv4 networks. The IETF designed "6to4" to provide automatic IPv6-over-IPv4 tunneling to interconnect IPv6 networks, while ISATAP is used to transmit IPv6 packets between dual-stack nodes on an IPv4 network. While these techniques and protocols may offer benefits, particularly during the transition phase, transporting IPv6 packets over IPv4 infrastructure can have security implications. These tunneling applications can be used to bypass network filtering policies. Organizations should implement mechanisms to block the use of default, automatic tunnels on end-user and perimeter devices (firewalls and edge routers). The Cyber Centre recommends using tunnel-aware security solutions. On network edge devices such as firewalls, organizations should deny by default all UDP outbound traffic and implement exceptions for authorized services only[11].

[8] Happy Eyeballs Version 2: Better Connectivity Using Concurrency: https://datatracker.ietf.org/doc/html/rfc8305
[9] RFC 6343 Advisory Guidelines for 6to4 Deployment: https://datatracker.ietf.org/doc/html/rfc6343
[10] Wikipedia: RFC 5214 Intra-site Automatic Tunnel Addressing Protocol (ISATAP): https://en.wikipedia.org/wiki/ISATAP
[11] RFC 9099: Operational Security Considerations for IPv6 Networks: https://www.rfc-editor.org/rfc/rfc9099.html

2.6 Default configurations

Modern operating systems (OS) and network devices will most likely support IPv6 and, due to the standard's requirements, this may be enabled by default. In addition, critical system functions may also require IPv6 to be enabled. For example, Microsoft does not recommend disabling IPv6 support on Windows OS devices, even when not in use [5]. To understand and assess related risks, organizations should proactively review the default status of IPv6 on their devices. Stay aware of risks associated with default configurations and design monitoring and preventative controls to mitigate those risks. For example, the 6to4 tunneling protocol is enabled by default on Windows servers when an interface is assigned a public IPv4 address. The tunnel assigns and dynamically registers an IPv6 address on the network[12]. If not monitored, this exposes the network to considerable risks. Organizations should implement mechanisms to drop unauthorized IPv6 traffic flows. To mitigate threats associated with IPv6 traffic transiting the network undetected, the Cyber Centre recommends proactive host-based monitoring for IPv6 network communications, even when the network interface is disabled. Detection of unauthorized network traffic should be investigated.

2.7 Unauthorized Internet Protocol version 6 traffic flows

Lack of visibility into IPv6 traffic flows represents a considerable risk on the network. Organizations with no approved use for IPv6 traffic should ensure IPv6 traffic flows are filtered on network edge routers and firewalls according to their network policies. A network that has deployed IPv6 should only allow IPv6 traffic that is permitted by policy, with access control lists (ACL) allowing only authorized flows and protocols and blocking all others by default. When IPv6 is being deployed, depending on the business case, a threat and risk assessment (TRA) may be required to identify and mitigate associated risks. In some cases, it may be infeasible to fully disable IPv6 functionality even with no business use case. For example, Microsoft does not recommend disabling IPv6 on Windows as some components require it to function properly. The Cyber Centre recommends a risk assessment to identify operational and security protections that could mitigate associated risks. In general, we recommend disabling IPv6 except where there is an approved operational need for its use on the enterprise network [5].

2.8 Monitoring and management tools

Network management and monitoring tools require substantial updates to manage and support IPv6 network traffic. Network security monitoring and reporting tools, such as an intrusion detection and prevention system (IDPS), log aggregation (via a security information and event management (SIEM) system), vulnerability scanners, and patch management tools, must support IPv6 protocols to ensure ongoing compliance with organizational security policies. The Cyber Centre recommends that organizations prioritize testing for different network monitoring scenarios (dual-stack and IPv6-only) as part of their IPv6 transition strategy. In addition, tailored test cases should be developed to validate support for IPv6 for software and business service development-related activities.

[12] Guidance for configuring IPv6 in Windows for advanced users: https://learn.microsoft.com/en-us/troubleshoot/windows-server/networking/configure-ipv6-in-windows

2.9 Addressing scheme

A robust IPv6 addressing scheme increases the security of the network, while providing the flexibility to support business services and mitigate information leakage threats. Organizations should consider their network's current state architecture as well as future needs when selecting an IPv6 addressing plan. Considering the sophistication and interdependencies of modern networks and applications, an addressing plan which supports a phased and incremental approach to IPv6 is recommended. An IP address management (IPAM) system is essential for effectively managing the addressing plan. Organizations should consider supported business applications and security policies when selecting an addressing scheme. The addressing plan can also be used to enhance an organization's security posture, as a foundational means for separating networks, while enforcing the zero-trust principles of network segmentation and segregation. If considering Unique Local Addresses (ULAs), they must be generated following approved pseudorandom algorithms and should be filtered at the network boundaries and not exposed beyond the internal network. While ULAs offer some benefits in IPv6 deployments, we would not recommend their use in dual-stack environments. For ULAs to be effective in dual-stack deployments, the address selection policy table precedence and label values may need to be updated on all devices on the network, introducing additional operational complexities and complicating network management and security processes.

2.10 Multi-addressing support

A single IPv6 interface can hold multiple addresses.[13] For example, an interface loopback address, a link local address, a unique local address, or a globally routable address. By default, a network interface with IPv6 is assigned a link local address. Multiple addresses offer both security and operational benefits; however, this can make it difficult to enforce network monitoring and filtering policies, particularly if filtering policies are not robust enough. This represents an increased threat surface and may allow threat actors to evade network traffic detection rules. The Cyber Centre recommends that system administrators implement restrictions on unauthorized changes to IPv6 addresses and ensure that monitoring controls are in place to prevent and detect changes. To mitigate the threat of malicious actors evading network security policies, implement deny-by-default policies to ensure traffic to and from an interface is blocked on network boundaries except for traffic that is explicitly allowed by the organization's network security policies.

2.11 Dynamic Host Configuration Protocol for Internet Protocol version 6

Most enterprise networks rely on Dynamic Host Configuration Protocol (DHCP) for distributing IP addressing information across the network. For IPv6, DHCP version 6 (DHCPv6)[14] supports both stateless and stateful addressing for network devices. Like the traditional DHCP protocol, DHCPv6 is susceptible to a variety of attacks such as malicious intercept, spoofing, and DDoS attacks. For enterprise networks deploying DHCPv6, the Cyber Centre recommends protecting DHCP network messages by using IPsec with encryption [6]. The Cyber Centre further recommends enabling authentication mechanisms between the DHCP servers, relay hosts, and client endpoints. Organizations should also implement additional protections such as DHCPv6 Guard[15] to block malicious DHCP reply and advertisement messages from unauthorized

[13] RFC 7934 Host Address Availability Recommendations: https://datatracker.ietf.org/doc/html/rfc7934.txt
[14] RFC 8415: Dynamic Host Configuration Protocol for IPv6 (DHCPv6): https://www.rfc-editor.org/rfc/rfc8415
[15] Cisco: DHCP – DHCPv6 Guard: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr_dhcp/configuration/15-sy/dhcp-15-sy-book/ip6-dhcpv6-guard.pdf

network devices. Organizations should consider DHCPv6 failover[16] capabilities to provide high-availability and protect against denial-of-service (DoS) attacks.

2.12 Address autoconfiguration protections

The IPv6 protocol specification allows devices to self-assign network addresses (i.e., interface identifiers (IIDs)) using the Stateless Address Autoconfiguration (SLAAC) protocol. As outlined in NIST SP 800-119 Guidelines for the secure deployment of IPv6 (PDF), SLAAC relies on network information received from the router and the device MAC address and can allow threat actors to track IPv6 endpoints. The Cyber Centre recommends disabling the use of SLAAC, particularly if a public addressing model is implemented. However, if an approved use case exists for SLAAC, the Cyber Centre recommends enabling SLAAC privacy extensions (which generate temporary IPv6 addresses) for external communications outside the enterprise network (for example, with the Internet or third-party networks). Enabling DHCPv6 temporary addressing can also provide the same protections as SLAAC privacy extensions. Please note that certain endpoints may not support DHCPv6, such as devices running on the Android OS, and may require self-configured SLAAC addressing as their only autoconfiguration option. In those scenarios, organizations should enable DHCPv6 Address Registration as a mechanism for SLAAC devices to inform the DHCPv6 server[17] of the address they self-generated. However, note that this may not provide visibility into auto-configured devices that don't support address registration or maliciously choose not to inform the DHCPv6 server. Organizations should implement network security controls to identify, manage and authorize network links with autoconfigured addresses.

2.13 Dual-stack environments

Dual stacking is a cost-attractive proposition for organizations, allowing the use of existing IPv4 infrastructure alongside IPv6. However, the need to maintain the IPv4 infrastructure while onboarding new IPv6 networks can increase the management burden and the attack surface. Dual-stack networks pose additional security concerns due to the use of multiple IP stacks, which increases the attack surface and requires additional security controls to mitigate associated risks. Dual-stacked host endpoints in particular present higher security challenges. Endpoint controls must have addressing controls for both IPv4 and IPv6 addressing schemes, which introduces additional complexity. The Cyber Centre recommends that organizations consider restricting host endpoints to single IP stack solutions (IPv4-only or IPv6-only). Limiting dual-stack architectures to transition mechanisms, switches, routers, or network gateways will help reduce the attack surface. Organizations should ensure that network or application firewalls are aware of and capable of filtering both IPv4 and IPv6 network packets.

The IPv6 specification standard establishes precedence rules which govern dual-stack interfaces. According to the IETF's request for comments (RFC) 6724 publication, default address selection for Internet Protocol Version 6, configured default policies may prioritize specific address groups over others, thereby leading to network operational complexities. This can have operational and security implications within dual-stack networks. Network and security administrators should be aware of address-selection precedence values deployed within their network environment. Administrators should also review and approve address-selection policies and ensure they are aligned with their network security objectives. Network security

[16] RFC 8156: DHCPv6 Failover Protocol: https://www.rfc-editor.org/rfc/rfc8156
[17] Registering Self-Generated IPv6 Addresses Using DHCPv6: https://datatracker.ietf.org/doc/rfc9686/

devices, including firewalls, edge routers, and gateways, should implement filtering policies to prevent unauthorized inbound or outbound IPv4 and IPv6 traffic.

In dual-stack DNS environments, A records (used to map domain names to IPv4 addresses) and AAAA records (used to map domain names to IPv6 addresses) are crucial for maintaining services. For Internet-exposed networks, the Cyber Centre recommends that organizations use separate DNS infrastructure for internal and external IPv4 and IPv6 networks (also known as split DNS architecture). This is to ensure the stability of system applications, increase security, and preserve the privacy of enterprise network data. For more information on split DNS architecture, read the NSA's Internet Protocol Version 6 Security Guidance.

2.14 Protection of data and management planes

Network administrative communications for IPv6 environments should be protected against eavesdropping, sniffing, and similar threats. The Cyber Centre recommends separating the management plane from the data plane using mechanisms such as virtual local area network (VLAN) separation or firewall filtering. ACLs, Intrusion Prevention Systems (IPS), and layer-2 filtering should also be used to protect the network management plane devices. For higher sensitivity networks, physical and cryptographic separation is highly recommended, for example, separation of management and data planes. The Cyber Centre further recommends that organizations use IPsec to protect IPv6 communications. Only CSE-approved cryptographic algorithms should be used, as indicated in the Cyber Centre's publication Cryptographic algorithms for UNCLASSIFIED, PROTECTED A, and PROTECTED B information (ITSP.40.111). Control plane protocols for IPv6 networks include ND, DHCP, Border Gateway Protocol (BGP), Network Time Protocol (NTP), and others. Organizations should consider implementing network filtering security controls. These controls will prevent control plane messages from inadvertently leaking information and disable or block unauthorized control plane protocols.

2.15 Neighbor discovery messages

Neighbor discovery (ND)[18] in the IPv6 specification is similar to the Address Resolution Protocol (ARP) used by IPv4. ND is used to manage crucial IPv6 capabilities such as address autoconfiguration, address resolution, duplicate address detection and others. However, the ND protocol is susceptible to several attacks[19] and can also be used by threat actors to perform address spoofing or poisoning attacks. The Cyber Centre recommends implementing network products which support cryptographic protections for ND such as Secure Neighbor Discovery (SEND)[19] . Cryptographic signatures generated through SEND are used to validate and verify ND messages, protecting against address spoofing attacks. Enabling IPsec can help secure ND messages. It is also advisable to filter ND messages (i.e., Internet Control Message Protocol version 6 (ICMPv6)) on external network boundary gateways except those required for IPv6 network connectivity. Please refer to RFC 4890[20] for guidance on filtering ICMPv6 messages in firewalls.

[18] RFC 4861 Neighbor Discovery for IP version 6 (IPv6): https://datatracker.ietf.org/doc/html/rfc4861
[19] RFC 3971 SEcure Neighbor Discovery (SEND): https://datatracker.ietf.org/doc/html/rfc3971
[20] RFC 4890 Recommendations for Filtering ICMPv6 Messages in Firewalls: https://datatracker.ietf.org/doc/html/rfc4890

2.16 Address translation risks

Address translation and tunneling of IPv4 over IPv6 and vice versa can introduce additional risk concerns. Translation devices can be a single point of failure and therefore high-availability and redundancy protections should be included as part of their architecture whenever they are deployed. Translation interfaces also force the termination of security mechanisms such as IPsec and Domain Name System Security Extensions (DNSSEC).

Network Address Translation-Protocol Translation (NAT-PT)[21] is a common translation mechanism that allows IPv4-only devices to communicate with IPv6-only devices. The Cyber Centre does not recommend using NAT-PT to communicate between IPv6-only networks via an IPv4 backbone or vice versa because of availability and end-to-end security concerns. Organizations can consider NAT64 (Stateful[22] Network Address Translation for IPv6-only clients to reach IPv4 servers) alongside DNS64 (a mechanism for synthesizing DNS AAAA records from A records)[23] or 464XLAT (combination of stateful[22] and stateless[24] translation for IPv4 connectivity across IPv6-only networks.)[25]

2.17 Zero trust architecture

Zero trust architecture (ZTA) is built on the foundational security principle of eliminating implicit trusts within the enterprise network. Zero trust assumes no inherent trust for resources and thus requires each resource (application, device, user, and network interface) to be uniquely identified, authenticated, and authorized.

The IPv6 standard provides foundational capabilities for the implementation of zero trust. These capabilities include an expanded address space, multiple addresses per interface, and IPsec header extensions for source authentication, data integrity and data confidentiality.

A multiple addressing strategy can be used to identify devices, interfaces or applications on the network, providing foundational support for micro-segmentation. This makes micro-segmentation easier, allowing traffic flows to be managed through fine-grain network access control lists.

Additionally, organizations can leverage IPv6 extension headers by enabling IPsec to achieve secure end-to-end IP communications. Enabling IPsec provides interface authentication and end-to-end confidentiality and integrity protections for data and control messages on the network.

2.18 Technical and operational depth

The lack of technical understanding and operational expertise in IPv6 represents a significant challenge for many organizations. Few network engineers possess detailed knowledge of the IPv6 specification standards. To build the technical competencies required for the future, organizations should invest in training networking and security professionals

[21] RFC 4966 Reasons to Move the Network Address Translator - Protocol Translator (NAT-PT) to Historic Status: https://datatracker.ietf.org/doc/rfc4966/
[22] RFC 6146 Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers: https://datatracker.ietf.org/doc/html/rfc6146
[23] DNS64: DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers: https://datatracker.ietf.org/doc/html/rfc6147
[24] RFC 7915 Stateless IP/ICMP Translation Algorithm: https://datatracker.ietf.org/doc/html/rfc7915
[25] 464XLAT: Combination of Stateful and Stateless Translation: https://datatracker.ietf.org/doc/html/rfc6877

on IPv6 and its capabilities. Organizations should also develop expertise by exploring IPv6 capabilities within dedicated network labs and/or limited pilot deployments. Organizations are encouraged to strengthen technical competencies and capabilities required to ensure network performance, address network design and operational issues, and architect security requirements.

3 Conclusion

The challenges associated with limited IPv4 addresses are likely to increase. IPv6 is designed to address these issues and offer additional security benefits. Modern networking stacks prioritize IPv6, conforming with the IETF specification standard, and organizational network enterprise strategies must be updated to manage associated risks. Traditional security controls that were built around IPv4 addressing, such as monitoring capabilities for example, will require updates and re-alignment. The Cyber Centre strongly recommends that GC organizations undertake proactive and informed actions to securely design and scope their IPv4 to IPv6 transition plans in line with Cyber Centre recommendations.