IPv6 Migration and Security Considerations

Paul Zoratti

East Carolina University

Introduction

Over the next several years, the number of devices connected to the Internet is expected to increase drastically. Recent estimates indicate that over 20 billion devices may be online by 2020 (Gartner Inc., 2017). Excluding devices relating to autonomous vehicle systems, most of these devices will be smart appliances for consumers, and electric meters and security cameras for businesses (Banking.com, 2017). For both the public and private sector, the demands of the applications poised to experience the most growth due to IoT devices cannot be satisfied by IPv4 alone. Currently, the 32-bit address scheme used by IPv4 can accommodate slightly over 4 billion unique addresses, which were completely allocated for the first time in 2011. By reclaiming and redistributing unused IP addresses, the Internet Assigned Numbers Authority (IANA) and the Regional Internet Registries (RIRs) have managed to prolong the life of IPv4, however the eventual exhaustion of the 32-bit address space is still inevitable. The proposed solution to this issue in scalability is IPv6. Despite being standardized in 1995 (Deering & Hinden, 1995), IPv6 is still available is less than 25% of global networks (Internet Society, 2017).

The advantages available in IPv6 are the result of several updates to the IPv4 protocol. By re-using and re-purposing components of the older protocol, IPv6 is able to operate in tandem with its predecessor. Changes like the larger address space, better integration of IPSec, the utilization of Neighbor Discovery (ND) instead of the insecure Address Resolution Protocol (ARP), and modifications to the packet header format provide improvements to security and network administration capabilities. In this paper, we will provide an overview of the current state of the migration from IPv4 to IPv6, describe the security improvements that IPv6 brings, and briefly discuss some of the unique security issues presented by IPv6.

Current State of Migration

Before we can discuss the progress that has been made in migrating from IPv4 to IPv6, we must first discuss how adoption is measured. Relevant data can be collected from the evidence that internet traffic leaves in routing protocols, domain names, and applications that are accessible from the public Internet using IPv6 and connections users are able to make. Being able to establish a connection to a web site or send an email using IPv6 proves all of the following are configured and working properly:

Client	Application, device, network, upstream provider
Server	Upstream provider, network, server, application

In this series of connections, the most common one preventing connectivity is the user or server network. One driver of this lack of connectivity at the network level is the fact that network providers must pay for IPv4 and IPv6 Access Point Names separately. In the past, with fewer applications utilizing IPv6, it was not cost effective for mobile providers to offer both IPv6 and IPv4 connectivity due to insufficient demand for IPv6 services from consumers. This demand for IPv4 over IPv6 is gradually beginning to shift as an increasing portion of IT departments are deciding that the cost of purchasing and maintaining IPv4 addresses is no longer worth the benefit. Similarly, for many organizations that already own large volumes of IPv4 addresses, it is becoming more cost effective to sell the IPv4 addresses and use the proceeds to fund the purchase of replacement IPv6 addresses. MIT is a prime example of an organization successfully employing such a strategy, selling over 8 million of the 14 million unused IPv4 addresses they once owned (McNamara, 2017). This hoarding of unused IPv4 addresses based on speculation of future usage is also a large driver in the depletion of addresses available for purchase.

According to data collected by Google, IPv6 (either native or via tunneling protocols like 6to4 and Teredo) routinely accounts for over 20% of their daily global traffic (Google, 2018). This rate of adoption varies widely by country, however. Cloud service provider Akamai estimates that IPv6 is responsible for over 40% of traffic originating in world leading countries like the United States (40.4%) and Belgium (46.4%), but only 20 countries have seen adoption rates that exceed 10% (Akamai, 2018). This disparity in global adoption rates indicates that IPv4 and IPv6 will have to continue to run in parallel for the foreseeable future. Despite this reality, the lack of ubiquitous IPv6 access should not prevent providers and developers from deploying and leveraging the myriad of benefits offered by IPv6.

In fact, continuing to delay adoption of IPv6 and only accommodating IPv4 traffic on a network is not an adequate solution to the dangers presented by IPv6. Many modern operating systems, including those running on phones and tablets, use IPv6 as the default method of communication and only use IPv4 as a fallback option. By making use of dynamic tunneling protocols (such as Toredo or 6in4), these devices are able to tunnel IPv6 connections through networks only supporting IPv4 traffic. Although these dynamic tunneling protocols are a necessity for the co-existence of IPv4 and IPv6, they provide malicious actors with a method of subjugating firewalls that are only configured to handle IPv4 traffic. Further compounding this issue is the prevalence of bring your own device programs and flexible work options that allow employees to connect remotely. Once devices are taken off of the IPv4 network, they are likely to connect to an IPv6 network at some point. This is particularly true with smartphones, tablets, and other devices that use cellular data due to the increasing availability of IPv6-capable LTE networks. If the controls and safeguards for employee devices only address IPv4 traffic, they will be inadequate to protect the same device on IPv6 capable networks. In the event that a device is compromised while on an IPv6 enabled network, there is the potential for further infiltration once that device reconnects to the IPv4 only enterprise network.

Increased Address Space

Instead of the 32-bit addresses used by IPv4, IPv6 uses 128-bit IP addresses. This increases the total number of unique IP addresses from around 4 billion (2^{32}) to over 340 undecillion (2^{128}). This larger address space satisfies the primary requirement for the replacement of IPv4 (namely, the ability to allocate unique IP addresses to large quantities of networked devices) and also provides many secondary benefits over its predecessor.

The abundance of unique IP addresses in IPv6 also removes the need for address conservation techniques such as the deployment of Network Address Translation (NAT) devices, which has the effect of simplifying network administration duties considerably and potentially allowing for full end-to-end connectivity. At first, it may appear that the network loses many of the secondary benefits provided by NAT such as reduced exposure to individual nodes, masking of the internal network topology, and IP masquerading capabilities. Fortunately, many of these risks can be mitigated by implementing firewall rules that ensure connections between devices outside the network and those inside network are only possible when the connection is initiated by the node within the network perimeter. NAT also represents a single point of failure within the network; any packet that crosses the perimeter of such a network will necessarily need to traverse the NAT device. This makes NAT devices a prime target for DoS attacks, as compromising one of these nodes can have severe impact on the network as a whole (Winemiller et al., 2012). Although some of these attacks can be negated in smaller networks by implementing rate-limiting functionality on the NAT device, either by limiting the total number of translations or by setting a maximum number of translations that can be made by an individual host, these solutions fail to scale to larger enterprise networks as they have the potential to prohibit legitimate users from establishing new connections.

The 128-bit address space also changes the dynamics of conducting port scanning in an IPv6 network. Port scanning is one of the first methods used by actors to perform reconnaissance and gather information about a network prior to an attack and is a valuable tool in discovering potential vulnerabilities. Researchers at the University of Michigan have built a tool that is capable of scanning the entirety of the public IPv4 address space for a given port in less than 45 minutes using nothing more than consumer-available hardware (Durumeric, Wustrow, & Halderman, 2013). In practice, most IPv4 subnets are /24 (up to 256 hosts), which allow comprehensive scans to be conducted with relative ease. By contrast, the minimum recommended subnet (and the required subnet size to utilize Stateless Address Autoconfiguration) for an IPv6 LAN is /64 (up to 18 x 10^{18} hosts). For this reason, while IPv4 network administrators must be concerned with whether or not a subnet has enough unique IP addresses to accommodate their given volume of end users, IPv6 network administration places more focus on the total number of subnets that are available for allocation. Although it currently isn’t feasible to conduct comprehensive scans of subnets at that scale, the systems used for managing and configuring those subnets make such scans unnecessary.

Several methods exist for configuring IPv6 addresses: Dynamic Host Configuration Protocol Version 6 (DHCPv6), StateLess Address AutoConfiguration (SLAAC), and manual.

Due to scalability issues inherent to manual solutions, we will not consider their implications in this discussion; however, it is worth noting that they tend to face the same challenges as those found in DHCPv6. Chief among those challenges is the fact that, by default, DHCP servers will assign a sequential IP address to each individual interface on the network. Attackers that are aware of this caveat can then drastically reduce the scope of their scans to start at lower IPs and work their way up through the sequence of allocated addresses. In implementations that do not use DHCP, SLAAC is often used for network configuration. In SLACC, a unique identifier (also known as an Interface Identifier, or IID) is generated for each individual interface on the network by inserting the string “0xfffe” into the MAC address of the interface. This presents several issues. First, this means that attackers can search for the “0xfffe” string in the IID to reduce the scope of their scans. Also, because the first portion of a MAC addresses is its Organizationally Unique Identifier (OUI), attackers can further limit their searches to known OUIs. In addition to allowing attackers to reduce their search space, the use of OUIs also allows attackers to launch device-specific attacks once a connection has been established, rather than having to perform further reconnaissance to determine the type of device they have found. Another consequence of the practice of constructing the IID using static information about the device itself is that the IID for a particular device will not vary over time. This allows for analysis of the activity of particular hosts within the network, which is a compromise to privacy. Even more serious is the fact that this IID can then potentially be used to track the activity of the host across multiple networks. For example, when a device connects to the public Internet from several different access points, the IID remains constant. Finally, similar to the sequential numbering issue found in DHCP implementations, it is possible that devices purchased in bulk (a common practice in most enterprises when acquiring laptops, smartphones, servers, and other IT resources) will have sequential MAC addresses due to logistical factors. As with DHCP configured networks, this allows attackers to tailor their scans to search for adjacent IP addresses once a responsive address has been identified. Although security through obscurity is not an adequate solution on its own, limiting the capabilities of port scans can be a valuable component for a defense-in-depth strategy. As a result, several options exist to obfuscate patterns in IP address allocation in the IPv6 space.

In DHCPv6 implementations, it is recommended that the DHCP server be configured to start allocating addresses somewhere other than the first address in the available space (typically ::1) (Gont & Chown, 2016). If possible, it is also advised that random addressing be used to introduce sparseness into the allocated address space. This scenario highlights the aforementioned scalability issues with manual implementations; DHCPv6 makes this pseudo-random assignment a viable option, while manual implementations struggle to accommodate such a scheme. Finally, DHCPv6 is also capable of using privacy extensions to further obfuscate IP address allocation.

In SLAAC implementations, alternative methods exist to create interface identifiers that do not depend on the MAC address of the interface, called Semantically Opaque Interface Identifiers (Gont, 2014). These random, temporary IIDs, also known as temporary addresses, can help complicate the task of information gathering and traffic analysis within a network. The standard IPv6 IIDs are still used for incoming connections (and thus remain static), while the temporary addresses are used to establish outgoing connections. This approach has its own set of benefits and drawbacks; although the use of temporary addresses can indeed make eavesdropping more difficult, it also introduces increased complexity to network administration tasks like event logging, identity and access management, and troubleshooting/root cause analysis. It is also important to be aware that because the traditional IID is still used as a static address for incoming connections, some traffic analysis and tracking will continue to be possible. Likewise, the traditional IID can still be used for network scanning purposes to identify active nodes. Because these Semantically Opaque Interface Identifiers complicate network administration tasks and fail to fully eliminate the vulnerabilities of traditional IIDs, many organizations have been reluctant to allow the use of temporary addresses within their enterprise networks.

IPSec

The IP Security suite, or IPSec, is a collection of protocols that is meant to provide flexible end-to-end encryption and support both data authentication and data integrity assurance. Originally IPSec was designed to be a mandatory component of IPv6, but has since been downgraded to merely a recommended function. Due to changes that had to be made in order for IPSec to work in IPv4 networks, IPv4 is not able to utilize the full scope of benefits offered by IPSec in IPv6 implementations. One example of this is the way that Authentication Header (AH) is implemented in each network configuration. AH provides protection against option-insertion, payload manipulation, and alteration of non-mutable header fields when implemented in an IPV4 network. However, AH also protects header-insertion and several additional header fields when used in an IPv6 network. IPv6 also has built-in extension headers that can be utilized for IPSec, but because of its usage in the IPv4 space, IPSec is still often included in the packet payload to provide full end to end encryption at the application level.

There are also several vulnerabilities that arise due to the extension headers used by IPv6. The Hop-by-Hop option header is an example of an extension header that can easily be re-purposed for malicious activity. This particular header is processed by every node that views the header of the packet. Because of the flexible nature of many of the extension headers in IPv6, they can be filled with large or unlimited amount of data. If an attacker were to fill the Hop-by-Hop header with massive amounts of data and random type identifiers, they could essentially launch a simple Denial of Service (DoS) attack. While this example references the Hop-by-Hop header specifically, similar vulnerabilities exist in other extension headers as well. As a result, these headers and their associated risks and attack vectors must be dealt with by all devices responsible for facilitating network connections (firewalls, switches, routers, etc.).

Neighbor Discovery

Neighbor Discovery Protocol (NDP) is a protocol in IPv6 that handles the responsibilities previously controlled by Address Resolution Protocol (ARP) and Internet Control Message Protocol (ICMP). NDP utilizes five new ICMPv6 packet types (Router Solicitation, Router Advertisement, Neighbor Solicitation, Neighbor Advertisement, and Redirect) to gather connection information like domain name servers, local connection configurations, and other connection data. This allows hosts to find routers, determine next-hop nodes, and discover when a neighbor has become unreachable.

As with the other components of IPv6, there are vulnerabilities inherent to NDP that must be considered prior to deployment. Unfortunately, in order for NDP to function, each node on the local network must have some degree of trust in the other nodes. This means that, should a malicious actor get access to the LAN, they could use NDP to conduct information gathering, forward packets, or even perform man in the middle attacks. One area of particular concern is the ability for an attacker to generate rogue Router Advertisements. This can allow a malicious node to masquerade as the default router (Ouseph, Chandavarkar, 2016). Secure Neighbor Discovery (SeND) attempts to mitigate some of these vulnerabilities within the Neighbor Discovery protocol itself, however, lack of widespread implementation of SeND means that it cannot be relied on implicitly, and other mitigation strategies (like smaller subnets or using immutable NDP entries) must be employed.

Conclusion

The purpose of this paper is to provide an overview of the current state of the migration to IPv6, identify security improvements in IPv6, and explain some of the major security issues and their associated mitigation strategies. IPv6 adoption has been slow at best and is still non-existent in some regions of the world. As new nodes continue to be incorporated into the global internet of connected devices, the migration to IPv6 becomes increasingly urgent in order to provide sufficient individual IP addresses for each node. Some of the primary inhibitors to more widespread adoption of IPv6, excluding cost, are issues with security that are not relevant in IPv4 deployments. While the 128 bit address space, IPsec, and decoupling from NAT services offer many benefits, they also introduce several novel vulnerabilities and attack vectors to be used by malicious actors. Regardless of the individual implementation, due to the interoperability of IPv4 and IPv6, these risks must be accounted for whether or not a given network plans to fully support IPv6 traffic itself. Although IPv4 will not be fully deprecated anytime in the near future, the benefits of implementing IPv6 now outweigh the costs. Finally, the consequences of not hardening public facing networks against IPv6 traffic can be disastrous for legacy IPv4 networks.

References

Akamai (2018, April 4). State of the Internet: IPv6 Adoption Visualizaiton. Retrieved from https://www.akamai.com/uk/en/about/our-thinking/state-of-the-internet-report/state-of-the-internet-ipv6-adoption-visualization.jsp

Banking.com. (2017, March 9). The IoT explosion and what it means for payments. Retrieved from http://www.fitech.com/news/the-iot-explosion-and-what-it-means-for-payments/

Deering, S., & Hinden, R. (1995, December). "Internet Protocol, Version 6 (IPv6) Specification", RFC 1883, DOI 10.17487/RFC1883. Retrieved from https://www.rfc-editor.org/info/rfc1883

Durumeric, Z., Wustrow, E., & Halderman, J. (2013, August). “ZMap: Fast Internet-Wide Scanning and Its Security Applications.” ZMap: Fast Internet-Wide Scanning and Its Security Applications | USENIX, Proceedings of the 22nd USENIX Security Symposium, Aug. 2013. Retrieved from zmap.io/paper.pdf.

Gartner Inc. (2017, February 7). Forecast: Internet of Things - Endpoints and Associated Services [Press Release]. Retrieved from https://www.gartner.com/newsroom/id/3598917

Gont, F., & Chown, T. (2016, March). “Network Reconnaissance in IPv6 Networks”, RFC 7707, DOI 10.17487/RFC7707. Retrieved from https://www.rfc-editor.org/info/rfc7707

Gont, F. (2014, April). “A Method for Generating Semantically Opaque Interface Identifiers with IPv6 Stateless Address Autoconfiguration (SLAAC)”, RFC 7217, DOI 10.17487/RFC7217. Retrieved from https://www.rfc-editor.org/info/rfc7217

Google (2018, April 6). IPv6 Adoption Statistics. Retrieved from https://www.google.com/intl/en/ipv6/statistics.html#tab=ipv6-adoption&tab=ipv6-adoption

Internet Society. (2017, May 25). State of IPv6 Deployment 2017. Retrieved from https://www.internetsociety.org/resources/doc/2017/state-of-ipv6-deployment-2017/

McNamara, P. (2017, April 21). MIT selling 8 million coveted IPv4 addresses; Amazon a buyer. *NetworkWorld.com.* Retrieved from https://www.networkworld.com/article/3191503/internet/mit-selling-8-million-coveted-ipv4-addresses-amazon-a-buyer.html

Ouseph, C. & Chandavarkar, B. (2016, May 20). "Prevention of MITM Attack Caused by Rogue Router Advertisements in IPv6". IEEE International Conference on Recent Trends in Electronics, Information, & Communication Technology. DOI: 10.1109/RTEICT.2016.7807969

Winemiller, N., Hartpence, B., Johnson, D., & Mishra, S. (2012). NAT Denial of Service: An Analysis of Translation Table Behavior on Multiple Platforms. *The 2012 International Conference on Security and Management.* Rochester, New York: Rochester Institute of Technology. Retrieved from http://scholarworks.rit.edu/cgi/viewcontent.cgi?article=1756&context=other