Tenable's SecurityCenter and IPv6

February 1, 2013

(Revision 1)

Table of Contents

Introduction

This document outlines some of the basic problems IPv6 was designed to fix and some of its other benefits when compared to the more commonly-known IPv4 protocol. It also provides information about how SecurityCenter 4.6 was developed in order to support IPv6 capable networks.

IPv6 Features

The following list summarizes some of the features of the IPv6 protocol:

- Use of scoped addresses and address selection
- More efficient forwarding
- Support for security and mobility

Use of Scoped Addresses and Address Selection

Unlike IPv4 addresses, IPv6 addresses have a “scope”, or a defined area, of the network over which they are unique and relevant. For example, IPv6 has a global address that is equivalent to the IPv4 public address and a unique local address that is roughly equivalent to the IPv4 private address. Typically, IPv4 routers do not distinguish between a public address and a private address and will forward a privately addressed packet over the Internet. An IPv6 router, on the other hand, is aware of the scope of IPv6 addresses and will never forward a packet over an interface that does not have the correct scope.

There are different types of IPv6 addresses with different scopes. When multiple IPv6 addresses are returned in a DNS name query, the sending node must be able to distinguish their types. When initiating communication, it uses a pair (source address and destination address) that is matched in scope and that is the most appropriate pair to use. For example, for a source and a destination that have been assigned both global (public) and link-local addresses, a sending IPv6 host would never use a global destination with a link-local source. IPv6 sending hosts include the address selection logic that is needed to decide which pair of addresses to use in communication. In addition to this logic, the address selection rules are configurable. This provides the ability to configure multiple addressing infrastructures within an organization. Regardless of how many types of addressing infrastructures are in place, the sending host always chooses the best set of addresses. In comparison, IPv4 nodes have no awareness of address types and can send traffic to a public address from a private address.

More Efficient Forwarding

IPv6 is a streamlined version of IPv4. Excluding prioritized delivery traffic, IPv6 has fewer header fields to process and fewer decisions to make in forwarding an IPv6 packet. Unlike IPv4, the IPv6 header is a fixed size (40 bytes), which allows routers to process IPv6 packets faster. Additionally, the hierarchical and efficient addressing structure of IPv6 global addresses means that there are fewer routes to analyze in the routing tables of organization and Internet backbone routers.

The massively large routing tables present on the Internet backbone today are an unintended side-effect of what was covered earlier regarding IPv4 address distribution. Since there was not a set structure for which addresses went where, there are now many disparate networks all over the world. The routing tables needed to support accurate packet delivery are very large, which causes unneeded latency. With IPv6 and a more systematic approach to address distribution based on region, these routing tables will shrink, causing a much more efficient way to route packets.

Support for Security and Mobility

IPv6 has been designed to support security (IPsec, AH and ESP header support required) and mobility (Mobile IPv6, optional). Although one could argue that these features are available for IPv4, they are only available on IPv4 as extensions, and therefore they have architectural or connectivity limitations that might not have been present if they had been part of the original IPv4 design. It is always better to design features into a product (or a protocol) rather than bolt them on later. Designing IPv6 with security and mobility in mind has resulted in an implementation that is a defined standard, has fewer limitations, and is more robust and scalable to handle the current and future communication needs of Internet users.

Problems Addressed by IPv6

The following list summarizes a few of the problems that the IPv6 protocol is designed to address:

- Address depletion
- International address allocation
- End-to-end communication

Address Depletion

On February 3, 2011, all unallocated Internet routable IPv4 addresses were officially exhausted, but this is not to say that IPv4 addresses are no longer being used. On the contrary, they are still very much in use today with the help of Network Address Translation (NAT). It is commonly known that IPv6 was designed to address the fact that we were going to run out of IPv4 addresses eventually. IPv6 solves this issue by providing an insurmountable number of addresses. The actual number of available IPv6 addresses that can be publicly allocated is 42 undecillion. That would be roughly 42 followed by thirty-six 0's. To put that number into perspective, it would take three times the age of the universe to scan all the IPv6 addresses in just a 48-bit subnet (which is what is presently allocated to individual sites) if you were scanning 1 million addresses per second, assuming that scientists are correct in estimating that the universe is about 13.7 billion years old.

International Address Allocation

The Internet was principally created by educational institutions and government agencies in the United States. In the early days, connected sites and even individual people received IPv4 address prefixes with little forethought. There is something to be said when Stanford University in California has more IPv4 addresses allocated to it than the entire country of China.

With IPv6, public address prefixes are assigned to regional Internet registries, which then assign address prefixes to other ISPs and organizations based on justified need. This new address allocation practice ensures that address prefixes will be distributed globally based on regional connectivity needs, rather than by historical origin. The business benefit to organizations across the globe is that they can rely on having available public IPv6 address space, without the extraordinarily high costs we see today when obtaining IPv4 public addresses.

End-to-End Communication

With IPv6, NATs are no longer necessary to conserve public address space, and the problems associated with mapping addresses and ports disappear for developers of applications and gateways. More importantly, end-to-end communication is restored by using packets containing addresses that do not change while in transit. This functional restoration has immense value when one considers the emergence of peer-to-peer telephony, video, and other real-time collaboration technologies for personal communications, and that of the next wave of devices that are connected to the Internet, which includes many types of peer-to-peer devices such as mobile phones.

By restoring global addressing and end-to-end connectivity, IPv6 has no barrier to new applications that are based on ad hoc connectivity and peer-based communication. The business benefit for software developers is easier development of peer-based applications to share information, music, and media, or to collaborate without having to work around the NAT translation barrier.

IPv4 and IPv6 Comparison

IPv4	IPv6
Source and destination addresses are 32 bits (4 bytes) in length.	Source and destination addresses are 128 bits (16 bytes) in length.
IPsec header support is optional.	IPsec header support is required.
No identification of packet flow for prioritized delivery handling by routers is present within the IPv4 header.	Packet flow identification for prioritized delivery handling by routers is present within the IPv6 header using the Flow Label field.

Fragmentation is performed by the sending host and at routers, slowing router performance.	Fragmentation is performed only by the sending host.
Has no link-layer packet-size requirements, and must be able to reassemble a 576-byte packet.	Link layer must support a 1280-byte packet and be able to reassemble a 1500-byte packet.
Header includes a checksum.	Header does not include a checksum.
Header includes options.	All optional data is moved to IPv6 extension headers.
ARP uses broadcast ARP Request frames to resolve an IPv4 address to a link-layer address.	ARP Request frames are replaced with multicast Neighbor Solicitation messages.
Internet Group Management Protocol (IGMP) IGMP is used to manage local subnet group memberships.	IGMP is replaced with Multicast Listener Discovery (MLD) messages.
ICMP Router Discovery is used to determine the IPv4 address of the best default gateway and is optional.	ICMPv4 Router Discovery is replaced with ICMPv6 Router Solicitation and Router Advertisement messages, and it is required.
Broadcast addresses are used to send traffic to all nodes on a subnet.	There are no IPv6 broadcast addresses. Instead, a link-local scope all-nodes multicast address is used.
Must be configured either manually or through DHCP for IPv4.	Does not require manual configuration or DHCP for IPv6.
Uses host address (A) resource records in the Domain Name System (DNS) to map host names to IPv4 addresses.	Uses AAAA records in the DNS to map host names to IPv6 addresses.
Uses pointer (PTR) resource records in the IN-ADDR.ARPA DNS domain to map IPv4 addresses to host names.	Uses pointer (PTR) resource records in the IP6.ARPA DNS domain to map IPv6 addresses to host names.

IPv6 Addressing

The most obvious distinguishing feature of IPv6 is its use of much larger addresses. The size of an address in IPv6 is 128 bits, a bit-string that is four times longer than the 32-bit IPv4 address. A 32-bit address space (IPv4) allows for 4,294,967,296 possible addresses. A 128-bit address space (IPv6) allows for up to 340,282,366,920,938,463,463,374,607,431,768,211,456 possible addresses.

Here is an example of an IPv6 address:

1999:DB8:0:2F3B:2AA:FF:FE28:9C5A

It is easy to get lost in the vastness of the IPv6 address space. The unthinkably large 128-bit IPv6 address that is assigned to an interface on a typical IPv6 host is composed of a 64-bit subnet prefix and a 64-bit interface identifier (a 50-50 split between subnet space and interface space).

1999:DB8:0:2F3B – Subnet space
AA:FF:FE28:9C5A – Host address

The 64 bits of subnet prefix leaves enough addressing room to satisfy the addressing requirements of the levels of Internet service providers (ISPs) between your organization and the backbone of the Internet and the addressing needs of your organization. The initial idea for the 64 bits of interface identifier was to map to the hosts link-layer media access control (MAC) addresses. However, there are ways to dynamically allocate a host address to the interface using a form of DHCP.

Global Unicast Addressing

When you break down the address into smaller pieces, the address actually becomes quite simple. To understand how the address is written out, consider the following address:

1999:DB8:0:2F3B:2AA:FF:FE28:9C5A

In any globally unique IPv6 address, the first 48 bits of the address are what makes a network unique from any other and are highlighted in blue below:

1999:DB8:0:2F3B:2AA:FF:FE28:9C5A

This portion of the address, called the "Global Routing Prefix", is given to a company, organization, or other type of entity by their ISP. In your organization, this portion of the global address will never change and after some time will become as familiar to you as any or your current IPv4 addresses.

The next portion of the address is called the Site ID. The Site ID is used within an organization's site to identify subnets within the site. The size of this field is 16 bits. The Site ID portion of the address is highlighted in red below:

1999:DB8:0:2F3B:2AA:FF:FE28:9C5A

The organization's site can use these 16 bits within its site to create 65,536 subnets (the equivalent of a public IPv4 Class A network). The routing structure of the organization's network is not visible to the ISP.

The remainder of the IPv6 address is referred to as the Host ID.

Local-use Unicast Addresses

Local-use unicast addresses do not have a global scope and can be reused. There are two types of local-use unicast addresses:

- Link-local addresses – Used between on-link neighbors and for Neighbor Discovery processes.
- Site-local addresses –Used between nodes communicating with other nodes in the same organization.

Link-local Addresses

IPv6 link-local addresses are used by nodes when communicating with neighboring nodes on the same link. For example, on a single-link IPv6 network with no router, link-local addresses are used to communicate between hosts on the link. IPv6 link-local addresses are similar to IPv4 link-local addresses that use the 169.254.0.0/16 prefix. The scope of a link-local address is the local link. A link-local address is required for some Neighbor Discovery processes and is always automatically configured, even in the absence of all other unicast addresses.

Link-local addresses always begin with FE80. With the 64-bit interface identifier, the prefix for link-local addresses is always FE80::/64. An IPv6 router never forwards link-local traffic beyond the link.

Site-local Addresses

Site-local addresses are equivalent to the IPv4 private address space (10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16). For example, private intranets that do not have a direct, routed connection to the IPv6 Internet can use site-local addresses without conflicting with global addresses. Site-local addresses are not reachable from other sites, and routers must not forward site-local traffic outside the site. Site-local addresses can be used in addition to global addresses. The scope of a site-local address is the site. Unlike link-local addresses, site-local addresses are not automatically configured and must be assigned either through stateless or stateful address auto-configuration.

The first 10 bits are always fixed for site-local addresses, beginning with **fe80::/10**. After the 10 fixed bits is a 54-bit Subnet ID field that provides 54 bits with which you can create subnets within your organization. You can have a flat subnet structure, or you can divide the high-order bits of the Subnet ID field to create a hierarchical and efficient routing infrastructure that can be easily summarized. After the Subnet ID field is a 64-bit Interface ID field that identifies a specific interface on a subnet.

In 2003, the use of site-local addresses was deprecated from the IPv6 standard, although you are still free to use them. The reason for the deprecation was mainly due to the ambiguity of the actual site-local addresses.

6to4 Addresses

6to4 is an address assignment and router-to-router, host-to-router, and router-to-host automatic tunneling technology that is used to provide unicast IPv6 connectivity between IPv6 sites and hosts across the IPv4 Internet. 6to4 treats the entire IPv4 Internet as a single link.

6to4 addresses are based on the prefix 2002:WWXX:YYZZ::/48 (in which WWXX:YYZZ is the colon hexadecimal representation of w.x.y.z, a public IPv4 address). An example of a 6to4 address is:

2002:836b:1:25:2aa:ff:fe53:ba63

In this example, 836b:1 is the colon hexadecimal version of 131.107.0.1. 6to4 allows you to assign global IPv6 addresses within your organization and to reach locations on the IPv6 Internet without requiring that you obtain a direct connection to the IPv6 Internet or an IPv6 global address prefix from an ISP.

However, for 6to4 tunneling to work, edge routers must be configured on the perimeter of your network and the destination network ready to encapsulate the IPv6 traffic appropriately.

Teredo Addresses

Teredo is an IPv6 transition technology that allows automatic IPv6 tunneling between hosts that are located on the IPv4 Internet, even when those hosts are behind one or more IPv4 NATs. Teredo resolves the issues related to the lack of 6to4 functionality in modern-day Internet edge devices and multilayered NAT configurations by tunneling IPv6 packets between hosts.

Teredo addresses are based on the prefix 2001::/32. These address prefixes are used to create global IPv6 addresses for IPv6/IPv4 nodes that are connected to the IPv4 Internet, even when they are located behind NATs. An example of a Teredo address is:

2001::ce49:7601:2cad:dfff:7c94:fffe

When using Teredo tunneling in your transition to IPV6, remember that IPv6 traffic that is tunneled with Teredo is not subject to the IPv4 packet filtering function of typical firewalls.

SecurityCenter 4.6 and IPv6

The following list summarizes the major IPv6 features in SecurityCenter (SC) 4.6:

- Perform an IPv6 host discovery scan on your network
- Passively detect IPv6 hosts on your network
- Scan a single IPv6 host
- Scan a whole range of IPv6 addresses
- Scan Assets defined with IPv6 addresses

IPv6 scanning is not available with Nessus 5.0.2 on the Windows platform.

From an engineering perspective, the goal for IPv6 within SecurityCenter was to treat it as much like IPv4 as possible. All of the functions and features available for IPv4 addressed hosts were designed to be present for IPv6 hosts. For the most part, this has been accomplished, but there are some aspects of the IPv6 protocol that must be considered:

- Scanning of link-local addresses

- Maximum scanning range
- Vulnerability correlation

Scanning of Link-local Addresses

SecurityCenter will technically allow you to scan link-local addresses. As mentioned earlier, link-local addresses cannot be routed past their local network segment. This means that the scanner has to sit on the local segment where the host to be scanned is located. In practice, there is no way that a scanner can see which link-local devices are on its specific local network. There is also no clear way to split link-local addresses into scan zones. What this essentially means is that scanning link-local addresses is not practical.

It is also important to understand how the new IPv6 functionality in PVS plays a part. As long as there is a repository with the link-local address range defined, PVS can be configured to analyze those packets. It is very important to note that the collection of these link-local addresses will count against an IP address count in your license.

Maximum Scanning Range

Scanning IPv6 addresses by range is supported, but is not encouraged. This document has covered the vastness of IP addresses that will be available, even for individual organizations. It should be assumed that the very large majority of these addresses will never be used, and that attempting to scan that large of an empty address space will waste time and resources. SecurityCenter implements a maximum target scan range of 2^24 IP addresses, which is fairly close to an equivalent IPv4 Class A network. To scan even that large of a range will be a waste of time and resources for most organizations.

There are a number of methods to define IPv6 ranges or hosts in a scan policy. The first (and easiest) is by leveraging the passive discovery of IPv6 hosts with PVS and the Dynamic Asset List functionality built into SecurityCenter. This is essentially an automated way to define IPv6 assets that can then be used in an active scan policy. A more manual way to create assets is to log into an external asset management server and copy and paste objects into an Asset List. This, of course, is not the most optimal method due to the lack of synchronization between asset management servers and SecurityCenter. It is also possible to utilize a new SC 4.6 feature that allows the ability to import host objects from an LDAP directory for use in creating Asset Lists. Finally, host names of machines known to be IPv6 capable can also be scanned.

Vulnerability Correlation

With the introduction of IPv6 support, vulnerability correlation may become somewhat more difficult. Multiple IP addresses for individual hosts will be seen, and there is no easy way to correlate those disparate addresses into one actual host. However, there are a few ways to potentially help with this correlation. One option is to sort SecurityCenter views by MAC address. Regardless of the number of IP addresses assigned to a NIC, the NIC itself will only have a single MAC address. The other option is to sort by host name, since the host name is generally unique across IPv4 and IPv6 networks.

For more information about the IPv6 capabilities available in SecurityCenter 4.6, see the “SecurityCenter 4.6 User Guide” available on the Tenable Support Portal.

References

- Davies, J. “*Understanding IPv6, Second Edition*.” Redmond, Washington: Microsoft Press. 2008.
- Huitema, C. *“Deprecating Site Local Addresses.”* http://tools.ietf.org/rfc/rfc3879.txt. IETF, September 2008.
- Huston, G. “*IPv4 Address Report*.” http://www.potaroo.net/tools/ipv4/index.html. November 2012.
- Welsh, C. “*Just how many IPv6 addresses are there*?” http://rednectar.net/2012/05/24/just-how-many-ipv6-addresses-are-there-really/ May 2012.

About Tenable Network Security

Tenable Network Security, the leader in Unified Security Monitoring, is the source of the Nessus vulnerability scanner and the creator of enterprise-class, agentless solutions for the continuous monitoring of vulnerabilities, configuration weaknesses, data leakage, log management, and compromise detection to help ensure network security and FDCC, FISMA, SANS CSIS, and PCI compliance. Tenable's award-winning products are utilized by many Global 2000 organizations and Government agencies to proactively minimize network risk. For more information, please visit http://www.tenable.com/.

GLOBAL HEADQUARTERS

Tenable Network Security
7021 Columbia Gateway Drive
Suite 500
Columbia, MD 21046
410.872.0555
www.tenable.com

tenable
network security