



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

The Next Internet

Privacy in Internet Protocol Version 6 (IPv6)

Author: Kevin Scott
Date Submitted: 24 March 2004
Certification: GIAC Security Essentials Certifications (GSEC)
Practical Assignment, Version 1.4b, Option 1

Table of Contents

Abstract	ii
Introduction.....	1
Origin of IPv6	1
Capabilities of IPv6	2
Complications of IPv6	3
Privacy Within IPv6.....	4
Identification	4
Auto-Configuration	4
Solutions	5
DHCP or Manual Configuration	6
RFC 3401	6
Additional Considerations	7
Conclusion	8
References.....	1

Abstract

This paper addresses the aspect of privacy relating to Internet Protocol version 6 (IPv6). It analyzes both the security features implemented in IPv6 as well as privacy-relevant concerns of capabilities implemented within IPv6 such as automatic configuration. The purpose of this document is to give the reader some insight into the beneficial security aspects of the pending implementation of IPv6 as well as an understanding of the concern regarding privacy within the specifications of IPv6.

© SANS Institute 2004, Author retains full rights.

Introduction

Internet Protocol (IP) version 6 (IPv6) is a technical standard that defines how devices that use IPv6 will communicate across a single network or multiple interconnected networks (i.e. the Internet). This standard provides a common format to enable information to transit the myriad of dissimilar networks and devices that comprise both the public Internet and private intranets.

Addresses for devices that communicate via IPv6 are separated into two distinct segments. The network prefix segment of the address indicates the network where the device is connected. The other segment of the address consists of the interface identifier for the specific device. The interface identifier segment has the capability to remain static even if network prefix segment of the address changes (e.g., based on connecting to different network). Internet Protocol version 4 IPv4 did not allow for a static identifier when a device was connected to a different network.

The possibility that a static interface identifier within an address might be used to locate and track the movement and activities of an individual device is a singular concern. It is expected that the number and type of network-connected devices will continue to increase and more of them will be associated with individual users (e.g., cell phones, PDAs, etc.). The primary issue is the capability to uniquely identify a device in an IPv6 environment (and by association, its user) and allow the device to be traced across a network or networks based on a static interface identifier embedded within the IPv6 address.

Origin of IPv6

The utilization of a well-defined structure for network communications is the fundamental nucleus around which the Internet, as we now know it today, developed. The current foremost protocol for network-based communications is known as Internet Protocol version 4 (IPv4). IPv4 was originally developed by the Defense Advance Research Project Agency (DARPA) nearly three decades ago. When IPv4 was created it was originally utilized within a closed network used by academic institutions and Department of Defense (DoD) research facilities. The need for security mechanisms to protect devices and (network) applications from each other was not anticipated and so the original IPv4 specifications did not directly address the many facets of network security.

The continuing explosion in the use and commercialization of the Internet has highlighted several significant issues with IPv4. The Internet Engineering Task Force (IETF), in consultation with a wide variety of governmental and private institutions has developed IPv6 to address many of these issues. The most obvious difference between IPv6 and IPv4 is a significantly greater number of available Internet addresses.

Currently IPv4 allows for a theoretical maximum of 4,294,967,296 (2^{32}) unique addresses while IPv6 allows for a theoretical maximum of more than 340,000,000,000,000,000,000,000,000,000,000,000 (2^{128}). The theoretical maximum address capacity of both IPv4 and IPv6 is however, effectively limited by the actual number of bits used to identify individual devices as well as limitations with the size of routing tables. IPv4 has implemented address conservation techniques, such as Classless Inter-Domain Routing (CIDR) (Leiner, p. 1), and the use of Network Address Translation (NAT) to allow more devices to be connected to the Internet without requiring individual Internet addresses.

With the rapid expansion of the Internet the issue of security within IPv4 has become another major concern. Although many alternatives have been incorporated into IPv4 they are in many ways crippled by other aspects of IPv4. In particular, although support for IP Security (IPsec) in IPv4 is optional, support for IPsec within IPv6 is mandatory. IPsec consists of enhancements to original IP protocol which provide authenticity, integrity, confidentiality and access control to each IP packet.

The IETF began identifying the requirements for a new version of the internet protocol in 1993 (Bradner, p. 2). Development of what was then referred to as IPng (IP Next Generation) (Bradner, p. 5) continued with the issuance of a Request For Comments (RFC) for a wide variety of aspects envisioned for IPv6. The basic draft standards for IPv6 (RFCs 2460, 2461, 2462, 2463) were finalized in August 1998. The IETF continues to refine the specifications with more than 70 RFCs published (Hinden, p. 1) that constitute the IPv6 standard. Based on these specifications the fundamental features of IPv6 are well established. Although some aspects of IPv6 are still being developed manufacturers and developers are now offering a range of IPv6 compatible software and hardware products. In addition, within the United States federal government, the DoD Chief Information Officer has established the goal of transitioning all DoD networks to IPv6 by 2008.

Capabilities of IPv6

IPv6 has been designed to provide a myriad of features and capabilities aside from the significantly increased address space. IPv6 will have a simplified header containing a more efficient hierarchical address structure. To simplify the format, the width of an IPv6 address header is fixed to 40 bits. This compares to an IPv4 header that can be either 20, 40, or 60 bits wide. This header will also provide data integrity and data authentication for the entire IPv6 packet based on IPsec. In addition, IPv6 will offer improved support for IP header options and extensions, the ability to label traffic flows, stateless and stateful address configuration (i.e. auto-configuration or “plug and play”), improved Quality of Service (QoS) support, support for neighboring node interaction, and extensibility.

The stateless address configuration (i.e., auto-configuration) is one of the capabilities that will significantly ease the burden of administering addresses on a network. Devices will be able to disconnect from one network and connect to another nearly effortlessly. The capability for a network, and devices connected to it, to reconfigure as necessary will finally bring about the vision of “Plug and Play” for network devices.

Complications of IPv6

IPv6 is in use today and there are individuals already learning to exploit some of the capabilities of IPv6, particularly while many network administrators remain unaware of its presence on their own networks. The IPv6 standards continue to be updated to address new concerns as they are identified, just as the IPv4 standards continue to be updated. Many of the changes implemented in IPv6 are expected to directly and indirectly improve the security environment for devices connected to both public (i.e. the Internet) and private networks worldwide, however IPv6 cannot be considered a universal security remedy.

A significant complication with IPv6 will be the transition from IPv4. Many “legacy” systems, software, and hardware cannot operate on an IPv6 network. The legacy hardware and software currently in use will need to continue to have access to an IPv4 network. Several solutions to this problem have already been developed (e.g. 6to4, dual stack capable servers, etc.) to allow IPv4 and IPv6 networks to coexist, or at least communicate through a gateway device.

An associated problem is cost. The cost of transitioning to IPv6 will be dependent on the size and complexity of your architecture. It is expected that most enterprise environments will transition to IPv6 from the outside-in. That is, the external systems (e.g. routers, firewalls, web servers, intrusion detection systems, etc.) that may need to interface with both IPv6 and IPv4 will be changed first. Following the implementation of IPv6 on the externally connected systems the internal network will be converted on a phased basis. This allows for the cost of converting to IPv6 to be spread out over years as well as allow the administrators time to gain experience with IPv6. In the mean time, administrators will have to become proficient in both IPv4 and IPv6 environments. The transition period from IPv4 to IPv6 is likely to take a decade or longer before the majority of legacy hardware and software has been replaced.

Although IPv6 is designed to be more secure, some of the improvements also introduce other known vulnerabilities (not to mention those that remain unknown) that will require knowledgeable administrators to ensure the security of the network. The stateless address configuration is one of the improvements that introduces concerns associated with the privacy of devices (and users) on a network. In the end it must be said that with all its improvements, and complications, IPv6 in and of itself will not protect servers, devices, or applications from being misconfigured or inadequately administered. The

resulting vulnerabilities will continue to allow the disruption and exploitation of networks and systems by unauthorized (malicious) agents.

Privacy Within IPv6

Concerns about privacy within IPv6 are associated with several possible methods to create static interface identifiers associated with an IPv6 address for a specific device. The terms confidentiality and privacy can be misinterpreted. For the purposes of this paper confidentiality is defined as the ability to protect information from unauthorized entities. Privacy, in the context used here, is defined as the ability to protect or conceal one's identity from others (i.e. remain anonymous). IPv6 enables a substantial degree of confidentiality for information passed over a network however, the connectivity required to implement this can result in a loss of privacy based on the IPv6 addressing scheme used to identify the communicating devices.

Identification

This paper addresses the issue of identification as it relates to globally unique, constant (i.e., non-changing) identifiers used by devices connected to an IPv6 network. Address identifiers for devices can be created by three primary methods: manual configuration, Dynamic Host Configuration Protocol (DHCP), or IPv6 auto-configuration (plug and play). In addition, identification can occur on a network by utilizing a Domain Name System (DNS) name. DNS names do not frequently change and can be considered a form of globally unique, non-changing identification. Finally, a simple "cookie" supplied by a web server to a web browser can be used as a constant identifier even without a static address or DNS name. Auto-configuration within IPv6 will be the focus of this document since the other methods referred to currently exist within the IPv4 infrastructure and are outside the scope of this paper.

Device, or node, addresses are fundamental to the function of any network. Because they must be used to allow routing across a network addresses are difficult to conceal from other entities on the network. The privacy concerns created by the auto-configuration capability within IPv6 are based on the ability for IPv6 addresses to be generated from information unique to a device and that address can then be used to identify the device and (theoretically) trace its activity on a network. The utilization of static identifiers to generate IPv6 addresses using stateless address auto-configuration is the primary basis for the privacy concerns of individual users within an IPv6 environment.

Auto-Configuration

Current network capable devices (i.e., Ethernet adaptors) are assigned unique 48-bit identification sequences known as embedded Institute of Electrical and Electronics Engineers (IEEE) identifiers (commonly referred to as Media Access Control (MAC) addresses). These MAC addresses are used by the stateless auto-configuration capability of IPv6 to create a unique, static address for the network interface of a device. The creation of this interface identifier does not require manual configuration or separate address-assignment (i.e., DHCP) servers.

When the interface identifier, which is the first 64-bits of the IPv6 address, is created from the 48-bit MAC address the MAC address must be expanded to create a 64-bit sequence. Simply put, the 48-bits from the MAC address are divided in half and a 16-bit value of “fffe” is placed in the middle to create a 64-bit address. This interface identifier derived from the network interface MAC address, which does not change, is then appended to the remainder of the IPv6 address. To create the complete 128-bit IPv6 address the 64-bits that need to be added to the interface identifier must be identified.

Auto-configuration allows a device on an IPv6 network to determine its network prefix by utilizing Neighbor Discovery to identify routers on the network. Neighbor Discovery allows a device to send a router solicitation message rather than wait for routers or other devices to initiate communication with the device. The router will respond with a router advertisement message that includes the network prefix. This network prefix includes the 64-bit network address for devices connected to that network (i.e., site-local and global-scope addresses). Finally, the 64-bit network prefix is combined with the 64-bit interface identifier to create the complete 128-bit IPv6 address.

As stated above, the privacy concern arises from the fact that when utilizing stateless auto-configuration, with an embedded IEEE identifier, the interface identifier portion of the devices address will remain static. Thus, if a network server were to log connection information then the source address of a device connecting to/through the server the interface identifier associated with that device would also be recorded. In such a case, the interface identifier embedded within an address could be used to track activities of a device (and an associated individual) no matter what network the device is connected to, or if a network were to be renumbered. Without some resolution to this privacy issue the ability for IPv6 to meet the intended goal of replacing the current IPv4 could be anticipated to meet strong resistance from a multitude of privacy advocacy groups worldwide.

Solutions

Privacy is a major issue for the IETF and although IPv6 can use unique identifiers as part of the network address, it is not a requirement of IPv6. (Deering, P. 1) In fact the IETF approved RFC3041, Privacy Extensions for Stateless Address Auto-configuration

in IPv6, in January 2001. The solution specified in RFC 3401 allows for continued use of the current IPv6 auto-configuration as well as the capability to create random interface identifiers to mitigate the privacy concerns about the IPv6 addressing scheme. Several solutions exist to mitigate the privacy concerns of using static interface identifiers within the IPv6 network address; these include: DHCP, manual configuration, stateless address auto-configuration based on original IPv6 standards, and stateless address auto-configuration utilizing a hashing algorithm defined in RFC 3401.

DHCP or Manual Configuration

One of the methods available to substitute a non-static identifier for a device into the network address within an IPv6 environment is DHCP. This method utilizes DHCP to assign addresses within a network. DHCP can be set to frequently “renew” a network devices lease of an IP address. The capability of DHCP to frequently change the identifier for a particular device cannot be assured however. The address a client gets from a DHCP server may change over time but most DHCP servers often return the same address to the same client for weeks or months. DHCP offers some resolution to the issue of privacy in an IPv6 network but not with any significant degree of assurance. Finally, although DHCP is also currently available in IPv4 networks, it will continue to be a viable alternative for managing address space within an IPv6 network.

Manual configuration of the interface identifier is another method. This may be acceptable for a single user on a home computer but for any larger environment the effort to continually change interface identifiers would quickly become impossible. Given the capabilities of either DHCP or manual configuration of a devices interface identifier the method identified in RFC 3401 represents the most effective solution to the privacy concerns of static identifiers associated with individual network devices.

RFC 3401

RFC 3041 was created to address the issue of privacy within the IPv6 network address by specifying a method to create modified interface identifiers. Changing the interface identifier would make it more difficult to examine IP addresses from separate transactions and determine which addresses are associated with the same device. RFC 3401 offers the capability to create new addresses utilizing a random number in place of the factory-assigned serial number. The approach specified in RFC 3401 provides a method to modify the interface identifier of a device by using an algorithm to generate proxy identifiers. The method of generating these random identifiers is to use an MD5 hash. To enhance the probability that different sequences of identifiers will be generated by different devices the pseudo-random sequence is based on a random component and the embedded IEEE identifier (if available). After a period of time a series of new interface identifiers will be generated and the preceding group of

addresses will be deprecated (this is done to obstruct reutilization of the same pseudo-random identifiers). (Narten, p. 7)

The rationale for selecting the MD5 hash mechanism centered on two factors; foremost was the capability for MD5 to satisfy the requirement to produce an acceptable degree of randomization. Even though many other algorithms could also meet this requirement, the secondary determinant for MD5 was that IPv6 capable devices must already utilize MD5 in support of IPsec and therefore the code already exists on IPv6 capable devices.

The intended result of RFC 3401 is to not change the basic methodology of how addresses are generated via stateless address auto-configuration. This methodology is designed to create a series of addresses from a sequence of randomized interface identifiers so that it would be very difficult for someone to identify whether two different interface identifiers belong to the same device.

Several factors concerning the creation of multiple interface identifiers had to be addressed within RFC 3401 and some assumptions were made. For example, it was assumed that the device could be configured to give preference to randomized temporary addresses as opposed to public addresses. This would require all connections initiated by the device use a randomized temporary address. In addition it is assumed that specific applications, based on their individual requirements, would be able to require either a randomized temporary address or a public address. (Narten, p. 8)

Additional Considerations

The impacts of utilizing randomized temporary addresses must also be considered. One factor to be evaluated is how often interface identifiers should be changed. This will primarily depend on the privacy concerns of the environment in which a device is used and the local policies of the user community. It must be noted that the need to rapidly change addresses can impact the device attempting to communicate using randomized temporary addresses. The device can be expected to experience additional processing overhead related to having a significant number of addresses associated with the device. Therefore, it can be anticipated that changing addresses frequently will degrade the overall network performance of a device. (Narten, p. 12)

Another potential problem with devices that regularly change their interface identifier is that it becomes more difficult for network administrators to isolate problems to individual devices. The need to monitor and manage the network may conflict with the effort to protect privacy and it may be that it is decided not to utilize temporary addresses within a network.

Auto-configuration of addresses makes the setup of an IPv6 network comparatively easy. In addition the capability to allow mobile users to seamlessly integrate into the network is significantly enhanced with stateless auto-configuration. .

Conclusion

IPv6 is a new, stable version of Internet Protocol that implements a significant number of performance and security improvements. The implementation of IPv6 has already begun and will gain momentum as more systems become IPv6 capable and more administrators become familiar with managing and IPv6 network. The increased security, enhanced capabilities, and ease of configuration should encourage many network administrators to migrate to IPv6 as soon as possible.

The capability to uniquely identify a device on a network and potentially trace its activity and/or its location across multiple networks has generated significant concern. A majority of the privacy concerns with IPv6 have been effectively resolved following the approval of RFC 3401. RFC 3401 provides extensions for stateless address auto-configuration in the IPv6 protocol and thereby enables the capability to maintain the privacy of devices connected to a network. Knowledge of and implementation of effective IPv6 security practices will allow for the advantages of IPv6 to be fully utilized by network administrators worldwide while reducing the privacy concerns for those connected to their network.

The implementation of IPv6 will create new opportunities to do much greater things with networks than have been done in the past. The capability to have devices automatically configure themselves based on the network they are connected to is only one element of what is possible. Additional capabilities can be expected in the form of RFCs. As developers and vendors gain proficiency with IPv6 they are bound to create new functions and capabilities that IPv6 was designed to be flexible enough to accommodate.

References

Leiner, Barry *et al.* "A Brief History of the Internet." Version 3.32. 10 December 2003. URL: <http://www.isoc.org/internet/history/brief.shtml> (17 February 2004).

Majstor, Franjo. "Does IPv6 protocol solve all security problems of IPv4?". October 2003. URL: http://www.employees.org/~franjo/papers/IPv6_security_paper.zip (12 February 2004).

Bradner, S. and Mankin A. "Network Working Group; Request for Comments: 1752". January 1995. URL: <http://www.ietf.org/rfc/rfc1752.txt?number=1752> (30 January 2004)

Hinden, Robert. "IPng Current Specifications". 21 September 2001. URL: <http://playground.sun.com/pub/ipng/html/specs/specifications.html> (2 March 2004).

Narten, T. (IPv6 Network Working Group). "Request for Comments: 3041". January 2001. URL: <ftp://ftp.isi.edu/in-notes/rfc3041.txt> (30 January 2004)

Warfield, M. "Security Implications of IPv6". 10 June 2003. URL: <http://documents.iss.net/whitepapers/IPv6.pdf> (28 February 2004)

Deering, S. "Statement on IPv6 Address Privacy". 6 November 1999. URL: <http://playground.sun.com/pub/ipng/html/specs/ipv6-address-privacy.html> (12 March 2004)

Särs, Camillo. "Address Assignment and Management in the IPv6 Environment". 15 April 1996. URL: <http://www.niksula.cs.hut.fi/~ged/Internetworking/> (12 March 2004)