International Journal of Scientific Research in Computer Science & Engineering
Technical Paper Vol-2, Issue-1 ISSN: 2320-7639

Adoption Ipv6: Security and Future

Abu Taha Zamani[1*], Javed Ahmad[2]

[1*]*Techno Global University, Shillong, Meghalaya, India, taha.abu@gmail.com*
[2]*Techno Global University, Shillong, Meghalaya, India, javedahmad.ali@gmail.com*

Available online at www.isroset.org

Received: 16 Jan 2014 Revised: 08 Feb 2014 Accepted: 20 Feb 2014 Published: 28 Feb 2014

Abstract-Internet Protocol version six (IPv6), the next generation Internet Protocol, exists sparsely in today's world. However, as it gains popularity, it will grow into a vital part of the Internet and communications technology in general. Many large organizations, including the Department of Defense, are working toward deploying IPv6 in many varied applications. This thesis focuses on the design and implementation issues that accompany a migration from Internet Protocol version four (IPv4) to IPv6 in the Monterey Security Enhanced Architecture (MYSEA). The research for this thesis consists of two major parts: a functional comparison between the IPv6 and IPv4 designs, and a prototype implementation of MYSEA with IPv6. The current MYSEA prototype relies on a subset of Network Address Translation (NAT) functionality to support the network's operation; and, due to the fact that IPv6 has no native support for NAT, this work also requires the creation of a similar mechanism for IPv6. This thesis provides a preliminary examination of IPv6 in MYSEA, which is a necessary step in determining whether the new protocol will assist with or detract from the enforcement of MYSEA policies.

Keywords- MYSEA, IPv4, IPv6, MLS, IP next Generation, Multilevel security, Network Address Translation

I. INTRODUCTION

In the Internet Protocol version six (IPv6), also known as the next generation Internet Protocol, lies the future of communications for networked computers and possibly the future of all telecommunications. Designed to augment and eventually replace the aging Internet Protocol version four (IPv4), the current standard, IPv6 stands in a position to replace the more than two-decade-old Internet Protocol (IP). The design of IPv6 likely contains improvements over the drawbacks of IPv4, some of which are causing concern among the community of Internet designers and engineers. Two examples of these trouble areas are the shrinking of the pool of available IP addresses, and the growth in size of routing tables stored on Internet routers. With time, the IP address space is becoming more and more stretched because of the unanticipated growth of the Internet. The growth of routing tables is attributable to inefficiencies of the initial IP addressing hierarchy. The web address cited in [PROBLEM] provides a synopsis of the history of the Internet's addressing troubles, and RFC 1752 [REC_IPng] provides a history of the birth of IPv6, including why it was developed. Additionally, new features in IPv6 may help to augment security and/or help IP to provide improved services. While IPv6 differs from IPv4, it is designed to perform the same basic functions as the original Internet Protocol. With this fact in mind, it is natural to hypothesize that the design of IPv6 improves on the original IP design while not adversely affecting the services it provides. The vastly larger address space and the native support for Internet Protocol Security (IPSEC) are two positive changes IPv6.

PURPOSE: There exist multiple reasons for performing this study. First of all, the Department of Defense (DoD) has committed itself to full deployment of Internet Protocol version six (IPv6) by the 2008 fiscal year [MEMO]. Secondly, the Internet is in the beginning stages of a transition to IPv6. Finally, new features in IPv6 have the potential to improve IP services in various applications. A clear determination of this potential is necessary before transitioning systems to IPv6.

The Monterey Security Enhanced Architecture (MYSEA) is a multilevel secure local area network (MLS LAN) that is designed to manage data at various levels of classification, and to allow untrusted commercial-off-the-shelf (COTS) client machines to securely access that data. This research specifically focuses on the design considerations of running MYSEA on an IPv6 network vice an IPv4 network. From a design perspective, it explores the areas in which IPv6 can assist in MYSEA's ability to enforce network policy.

In anticipation of making a transition to IPv6, it is necessary to analyze the costs and benefits of running MYSEA on an IPv6 network. Building MYSEA with native IPv6 functionality may even support and benefit the architecture more than IPv4. For a system like MYSEA to successfully complete a transition from IPv4 to IPv6, its designers and implementers must prepare early and understand any modifications this transition will demand. The research documented in this paper will provide the foundation of the work to build MYSEA in an IPv6 environment.

This work includes a review and comparison of the IPv4 and IPv6 designs. In addition, an IPv6 MYSEA prototype has also been developed. The MYSEA design requires functionality that is provided by network address

Corresponding Author: Abu Taha Zamani

translation (NAT) in IPv6; however, there currently are no – and there likely never will be any – NAT mechanisms defined or implemented for IPv6. This situation required either finding a replacement mechanism for NAT or implementing NAT in IPv6.

II. INETWORK ADDRESS TRANSLATION (NAT)

The development and deployment of NAT has come with many different benefits, and even some drawbacks. As explained in RFC 2663, "The term 'Network Address Translator' means different things in different contexts" [NAT_TERM]. The intent of this section is not to describe the many varieties, uses, advantages, and disadvantages of NAT; but merely to introduce the concept that it implements.

NAT Defined

Network Address Translation is a mechanism that allows nodes bearing private (unregistered) IP addresses to communicate in the global Internet by replacing the private addresses with public (globally unique) ones. The following paragraph illustrates the key ideas of NAT.

In a private network (using private IP addresses) that runs NAT, the border routers implement the NAT functionality. Normally, a border router will not forward any datagrams from an intranet into the Internet because they contain a private IP address as the source; however, a NAT router will simply swap the private address for a predetermined public address that conforms to the standard – either its own global address, or one from a pool of allocated valid addresses. After forwarding the modified datagram, the router maintains the address mapping so that it can map the reply packets to the substituted address. That is the basic function of NAT. Figure 5 and the example below it use the MYSEA architecture to illustrate how routers perform NAT.

INTERNET PROTOCOL VERSION SIX (IPV6)

IPv6 represents the next step in the evolution of a robust, flexible communications protocol that is intended to accommodate the communications and information sharing needs of the world. This section contains a summary of the IPv6 specification, [IP6]. The information herein focuses on IPv6 as it applies to MYSEA and with regard to IPv4. By no means does this section contain a comprehensive description of the protocol. For more details on IPv6 see [IP6].

General Changes to the IP Design

Note that the designers of IPv6 do not make any fundamental changes to the basic concept and functionality that the Internet Protocol intends to provide. IPv6 retains the same scope as IPv4, but the new design attempts to improve on the original design by making it simpler, yet more flexible, and no harder to implement. The following list, presented in [IP6], summarizes the intended changes from IPv4 to IPv6:

- Expanded addressing capability: The address size has increased from 32 to 128 bits. The new design also contains some changes to addressing schemes and address assignment that are beyond the scope of this discussion.
- Simplified header format: Discussed in the following section.
- Better support for extensions and options: The specification changes the encoding of IP header options, thereby increasing efficiency and flexibility, and easing the introduction of new options in the future.
- A flow labeling capability: A capability for labeling packets that belong to particular traffic flows for which a sender requests special handling.
- Privacy and authentication capabilities: IPv6 provides explicit extensions to support authentication, integrity, and confidentiality.

This list contains the intended changes from IPv4 to IPv6. Other significant changes in IPv6 include the assumption that every link in the Internet has an MTU of at least 1280 bytes. Also, only the originating node of a packet may perform fragmentation. The following sections will elaborate on the intended changes while providing an overview of IPv6.

IPv6 Headers

As previously stated, the format of the IPv6 header is a simplified version of the IPv4 header. Figure 6 illustrates the IPv6 header structure. As with the IPv4 header depiction, the numbers above the illustration represent a bit count, beginning with the number zero. The minimum size of an IPv6 header is 40 bytes, twice the size of the IPv4 header. The large size of the addresses almost necessitates simplifying and making the rest of the header smaller for the sake of conserving bandwidth.

The following list contains a brief description of each field in the header:

- Version: Current IP version
- Traffic Class: For use in distinguishing between classes or priorities of packets. This field is equivalent to the IPv4 TOS field.
- Flow Label: This field contains a label assigned to sequences of packets that require special handling by routers, such as a QoS specification.
- Payload Length: This field specifies the length, in bytes, of the payload, that is everything following the IPv6 header.
- Next Header: Specifies the type of header following the IPv6 header.
- Hop Limit: Performs the same function as the IPv4 TTL field. Each forwarding node decrements this value by one.
- Source Address: The 128-bit address of the originator of the packet.
- Destination Address: The address of the intended recipient of the packet.

IPv6 uses extension headers to encode optional information at the network layer, thereby adding to the

modularity of the IP design. These headers lie between the IPv6 header and the next layer protocol header in an IPv6 packet. Figure 7 illustrates the use of extension headers.

This capability, in part, replaces the functionality of the variable-sized options field in the IPv4 header. Since all fields in the IPv6 header have a fixed-length, the IPv6 header has a truly static size. An IPv6 packet can contain zero or more extension headers. Following is the list of extension headers specified in [IP6]:

- Hop-by-Hop Options
- Routing
- Fragment
- Destination Options
- Authentication
- Encapsulating Security Payload (ESP)

Nodes that forward packets do not examine any of these headers, with the exception of the Hop-by-Hop Options header. Every node in a packet's path from source to destination always examines this header. The specification adds more structure to the use of the extension headers by setting a specific order in which to include them (see [IP6] for that order). The Routing header provides functionality similar to the Loose Source Routing, Strict Source Routing, and Record Route options in IPv4. The IPv4 section of this paper contains short descriptions of those routing options. Source nodes include a Fragment header with each fragment of a transmitted packet. The Destination Options header carries optional information that only the ultimate receiver of a packet inspects. Options following this header have a variable length. The specification currently defines two options dealing with padding. It also provides some initial structure – required values of high-order bits for unrecognized options – for option definitions, and it contains guidance for introducing new options. Finally, the Authentication and ESP headers provide authentication and encryption respectively. These two headers relate directly to IPsec, and they are discussed in the IPsec section.

3. Addressing Architecture

RFC 2373 [IP6 ADDR] is the primary resource for IPv6 addressing, and the majority of IPv6 addressing information resides in that document. The model for addressing in IPv6 closely resembles that of IPv4, except that it natively employs the concept of CIDR. The 128-bit address, the native use of CIDR, and a new IPv6 addressing model stand out the most.

Basic Differences from IPv4 Addressing

As stated above, IPv6 uses a classless addressing structure. While the hierarchy is classless, IPv6 still has various ranges of reserved IP addresses. The specification also defines the following three address types for IPv6: unicast, anycast, and multicast. A unicast address simply identifies a single interface and functions as a normal IP address, the same as IPv4 addresses. An anycast address identifies a set of interfaces (usually on different nodes), and the "nearest" one, according to the routing protocol, receives the so addressed packet. IPv4 has no inherent provision for anycast addresses. A multicast address also identifies multiple interfaces that normally lie on different machines. A packet destined for this type of address is accepted by all interfaces that share the address. The IPv6 multicast address overrides IPv4's broadcast capability, so there are no broadcast addresses in IPv6. Note that IPv4 does have a specified multicast capability which was developed after the initial IP addressing specification.

Finally, the format for representing an IPv6 address in text differs from the IPv4 format. While it is possible to represent an IPv6 address in bitwise or dotted decimal notation, it would be much harder for a human reader to interpret since an IPv6 address is eight times larger than an IPv4 address. Instead, the standard separates an IPv6 address into eight pieces, each one represented by a sixteen bit hexadecimal value. A colon separates each value. A common shorthand method for representing multiple sequential zeros is presented in Figure 8. Alternatively, one may specify the first ninety-six bits using hexadecimal values and then use the well-known IPv4 bitwise notation to represent the final thirty-two bits. This format is useful for representing IPv6 addresses that map directly into IPv4 addresses. The section on transition tools discusses this type of IPv6 address. There are other minor intricacies involved with representing these addresses, but this is the basic method. More information is contained in [IP6 ADDR].

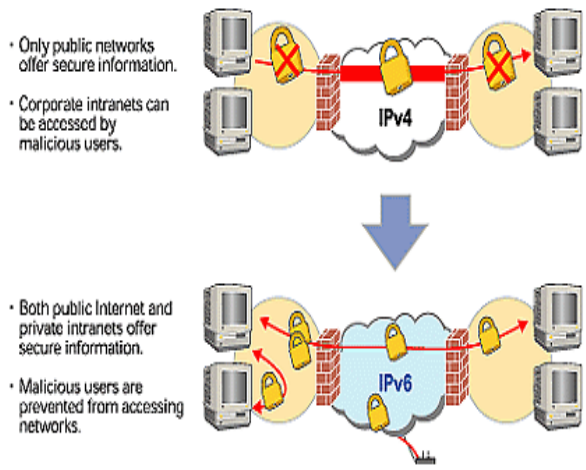

Fig-1 Comparisons of Ipv4 and Ipv6 security

III. SECURITY

As stated in the list of changes in IPv6, the protocol was designed with authentication and privacy capabilities. The Authentication and ESP extension headers provide these capabilities through the functions they perform, and these two headers are actually a part of the separately defined IP security architecture. This architecture is laid out and discussed in [IPSEC ARCH], and is briefly discussed in Section C. Based on the release dates of the RFCs, the security architecture existed before the IPv6 specification

was finalized. Therefore, since IPv6 incorporates IPsec into its design, it is accurate to state that IPv6 provides native support for authentication and confidentiality (encryption) of data. Section C introduces the IPsec architecture and describes the functions that the Authentication and ESP headers provide.

INTERNET PROTOCOL SECURITY (IPsec)

The Internet Protocol security architecture, better known as IPsec, has the capability to provide two essential functions to MYSEA. IPsec's encryption capabilities protect data flowing between the Trusted Path Extension and the server, and its authentication capabilities provide two-way authentication between those two nodes. The remainder of this section presents an overview of the IPsec design, framework, and its implementation in MYSEA. IPsec and all of its supporting concepts and operations are defined in multiple documents with a lot of intertwining information. This section attempts to capture the overall essence of IPsec without delving too deeply into the great amount of information defining it. The information provided here on IPsec is drawn from [IPsec, ISAKMP, and IKE]. Refer to [DOCMAP] for a listing of the documents pertaining to IPsec and a description of their interrelationships.

Design

IPsec is intended to provide a common set of security services for nodes on the Internet. These services are listed in the following sub-section. The major advantage of providing security services at the IP layer is that the services are available for IP traffic and all higher layer protocols [IPsec]. Since the Internet Protocol is standardized throughout the Internet, the IPsec services are universally available.

Goal

The design goal of IPsec aims to provide "interoperable, high quality, cryptographically-based security for IPv4 and IPv6" [IPsec]. IPsec provides the following services as described in [IPsec]:

- Access control
- Connectionless integrity
- Authentication
- Replay protection
- Data confidentiality
- Limited traffic flow confidentiality

In order to meet its goal and provide these services, IPsec relies on the AH and ESP headers as well as cryptographic key management protocols and procedures. Depending on user, application, and system requirements IPsec employs an appropriate set of protocols to provide security services requested by a user or application. While a default set of algorithms and protocols is defined to support interoperability in the Internet, IPsec is sufficiently flexible for groups of individuals to define and use their own sets of algorithms. Such flexibility is imperative for successful deployment of this protocol suite so it can provide all requested services while not interfering with the network and its usability.

How IPsec Provides Desired Services

First of all, note that the IPsec architecture does not cover the implementation of specific encryption algorithms and other protocols, but it assumes that their implementation is secure. The best-designed security algorithm or protocol can fail if poorly implemented; so, while algorithm implementations are beyond the scope of the architecture, it is important to recognize that they play a crucial role in the effectiveness of IPsec.

An IPsec implementation relies on a Security Policy Database (SPD) for direction on how to treat IP packets. Based on the security policy laid out in the SPD, packets are either provided with security services, discarded, or allowed to bypass IPsec altogether. On a single host, IPsec allows the system to specify security protocols, and then determines the algorithms and cryptographic keys that will facilitate the selected services. Once the services are selected, the cryptographic keys must be created on the desired machines.

IPsec uses symmetric (shared secret) keys and Security Associations (SA). A SA is a "simplex 'connection' that affords security services to the traffic" [IPsec] that it carries. IPsec relies on a separate mechanism for distributing the cryptographic keys and managing the SAs. The Internet Security Association and Key Management Protocol (ISAKMP), specified in [ISAKMP], presents a framework for managing security associations and cryptographic keys. ISAKMP does not define any specific methods for managing and distributing keys. Instead, it sets guidelines that all IPsec key management protocols must obey. With this method, IPsec can rely on any key management mechanism that is based on the ISAKMP template. The Internet Key Exchange (IKE), specified in [IKE], is an example of a public-key based approach for automatically distributing cryptographic keys. The keys may also be distributed manually or through some mechanism other than IKE. The distribution of keys, like encryption algorithms, is beyond the scope of [IPsec], so the design essentially assumes that effective key management and distribution methods are in use.

After key distribution, further communications between the involved nodes rely on the AH and ESP headers to provide the security services prescribed in the SPD. Both headers may provide connectionless integrity, data origin authentication, and an anti-replay service. The ESP can also provide confidentiality and limited traffic flow confidentiality.

IV. QOSS

One way that QoSS functionality can be provided to the network is through IPsec. As discussed in [QoSS], QoSS

functionality was added to OpenBSD's implementation of IPsec in IPv4.

Transitioning QoSS Capabilities to IPv6 in MYSEA

Implementing the QoSS capabilities in IPv6 will potentially involve changing the source code that implements it in IPv4. Since the concept was created and developed under IPv4, it is possible that some of the program code depends on peculiarities of that protocol. Such a situation would simply require "porting" those sections of code into conformance with IPv6. Otherwise, given the fact that IPsec is designed to function in either an IPv4 or an IPv6 environment, the QoSS additions to an IPsec implementation should be a transparent issue when switching protocols.

IPV4 VERSUS IPV6

Based on the above summaries of the IPv4 and IPv6 protocols, this section presents a comparison of the two designs. While some broad issues are addressed, this comparison primarily focuses on the issues that affect MYSEA. It seeks to pinpoint portions of the IPv6 design, if any, that could detract from the basic functionality that MYSEA aims to provide.

THE IPV4-TO-IPV6 TRANSITION

Over the last few years, a point of division has grown among the engineers and architects of the Internet. On one side of the debate stand those who believe that the shrinking address space of IPv4 (along with other concerns such as the size of routing tables) is not a significant problem. Opposing them are those who believe that IPv6 is the only option for the Internet's future communications protocol. Many among the IPv6 proponents believe that the immensely larger IPv6 address space will allay the world's IP address space concerns, and that the new protocol will greatly contribute to the advent of mobile and pervasive computing.

The obvious question arising from this debate is "who is right?" A potential answer could be that neither side is exclusively correct. As stated in [MECHS], "the Internet will need [both IPv4 and IPv6] compatibility for a long time … and perhaps indefinitely." Considering this possibility, it becomes clear that there is a need for mechanisms to allow seamless communication between nodes using either protocol. Therefore, this section does not seek to argue for one side or the other, but merely presents facts about current work intended to prepare the Internet for the use of IPv6. These transition mechanisms could also positively impact the use of MYSEA in an IPv6 environment.

V. CONCLUSION OF THE COMPARISON

It appears that the IPv6 design attempts to increase the overall modularity of the IP design. From the header to the extension headers to the aggregately addressing hierarchy, the specifications for IPv6 appear to focus on modularizing the design while minimizing interdependencies of those modules. In general, modularity is good because it increases the flexibility of the design. Just as IPv6's modular header design makes it easier to define new options, modular components increase the ease of modifying single components without affecting the entire design.

Based on its design and its comparison with the IPv4 design, the conclusion is that IPv6 can at the least provide the same unaltered services as IPv4. Furthermore, IPv6 could possibly improve the efficiency and security of those services. Changes involving the addressing structure and the default MTU have the potential to provide added efficiency across the network; and the simplification of the design coupled with the fact that IPsec is part of the design can provide more assurance of security. IPv6's monumental address space should do away with the necessity for performing NAT in the Internet; however, because its address hiding functionality is fundamental to the MYSEA design, that functionality must be implemented in an IPv6 version of MYSEA.

References

[1]. IPv6 Standards Profile Released," September 19, 2008.http://gcn.com/articles/2008/09/19/ipv6-standards-profile-released.aspx

[2]. IPv6 Wiki for Transition Managers https://max.omb.gov/community/x/EhPVI

[3]. NIST IPv6 Deployment Monitor http://fedv6-deployment.antd.nist.gov/

[4]. NIST SP800-119 "Guidelines for the Secure Deployment of IPv6, December 2010 http://csrc.nist.gov/publications/nistpubs/800-119/sp800-119.pdf

[5]. "NIST Special Publication 500-267: A Profile for IPv6 in the U.S. Government – Version 1.0", July 2008 http://www.antd.nist.gov/usgv6/usgv6-v1.pdf

[6]. NIST USGv6 Deployment Test Suite http://www.antd.nist.gov/usgv6/

[7]. NTIA IPv6 Web-page and Resources http://www.ntia.doc.gov/category/ipv6

[8]. "OMB: Agencies met IPv6 deadline," July 1, 2008. http://fcw.com/Articles/2008/07/01/OMB-Agencies-met-IPv6-deadline.aspx

[9]. Planning Guide/Roadmap Toward IPv6 Adoption within the U.S. Government http://www.cio.gov

[10]. Practical Guide on Federal Service Oriented Architecture, June 2008. http://www.whitehouse.gov/omb/E-Gov/pgfsoa

[11]. http://www.cio.gov/documents/IPv6MemoFINAL.pdf

