Global Information Assurance Certification Paper

Sans and MWC Security Essentials GSEC Practical Assignment Version 1.3

Security Features of Internet Protocol Version 6

By Robert G. Moulton
April 8, 2002

ABSTRACT

This paper examines the security features designed into the new Internet Protocol version 6, (IPv6) architecture. It looks at routing enhancements, new IP header structure, new authentication, and confidentiality capabilities, native IPSec support, discusses enhanced techniques such as VPN, and explains the new addressing scheme. It also provides a brief explanation of the current Internet Protocol's limitations as well as a comparison between the Internet Protocols' addressing schemes and Internet Protocol packet header formats because this is where the most significant changes take place and is fundamental to understanding the new IPv6 security enhancements. Understanding the new addressing scheme is also fundamental to understanding new protocol. The appendix provides a listing of the IPv6 architects who designed the protocol. This paper will be valuable to those in the IT field particularly the network administration arena who need to persuade management to finance an upgrade and also those tasked with implementing the new IPv6 architecture. Security features and enhancements are continually evolving and this paper is not an exhaustive resource. References are supplied at the end of this paper to help the reader further explore the topic. Although IPv6 features and technical aspects are explained by the Internet Engineering Task Force's request for comments (RFCs), some technical aspects of IPv6 are still evolving.

INTRODUCTION - Understanding the need for IPv6 security features

Security is becoming more important as businesses conduct transactions over the Web. IPv4 was never designed to be secure. It was originally designed for an isolated military network, and then adapted for a public educational & research network. IPv4 security features are retrofitted with many and varied solutions such as SSL, SHTTP, and IPSEC. Security was usually considered an issue addressed at the higher layers of the TCP/IP model, but with IPv6, security is implemented at the packet level as well. IPv4 lacks effective confidentiality and authentication mechanisms below the application layer. IPv6 addresses these shortcomings with two security services features employed as extension headers, "IPv6 Authentication Header" and the "Encrypted Security Payload (ESP) header"

New applications are more demanding and require guaranteed on-time delivery, guaranteed availability of bandwidth, and guaranteed secure transactions. This is difficult and burdensome to retrofit onto the base IPv4 technology and adds a lot of unnecessary overhead, slowing the network throughput and increasing users response times. IPv6 is delivers security without performance degradation and mandates the use of the IP Security (IPSec) standards.

IPv6 is a 128-bit version of the current 64-bit Internet protocol. These increase bits provides more information about the role of network, such as: server, workstation, private network, router, virtual private networks, etc. Also built into IPv6 are virtual private networking (VPN), Neighbor Discovery, Stateless Address Autoconfiguration, transmission of IPv6 over IPv4 Domains without explicit tunnels, and more.

IPv6 Security Enhancements

Routing Enhancements - No Need for NAT

Hastening the upgrade to IPv6 is the inefficient Network Address Translation routing scheme in use today. Network Address Translation is a mechanism for reducing the need for globally unique IP addresses. NAT allows an organization with addresses that are not globally unique to connect to the Internet by translating those addresses into globally routable address space. Corporate Intranets most commonly use this to hide the actual IP addresses. A proxy server interfaces the Internet and has the only globally unique address. None of this cumbersome and overhead intensive system is necessary with IPv6.

A new feature introduced with IPv6 is the Neighbor Discovery (ND) protocol. Used by both hosts and routers, ND discovers other nodes on the link and obtains their datalink-layer address. Hosts also use ND to find neighboring routers that are willing to forward packets. When a router or the path to a router fails, a host actively searches for functioning alternates. Besides detecting changed link-layer addresses, the protocol is also used by nodes to keep track of, which neighbors are reachable and which are not.

The IPv6 ND protocol is functionally the equivalent to a combination of the IPv4 protocols ARP, ICMP Router Discovery (RDISC), and ICMP Redirect (ICMP), although, IPv4 has no protocol or mechanism for unreachable neighbor detection.

ND provides a multitude of improvements over the IPv4 set of protocols:

- Router Discovery is part of the base protocol set; there is no need for hosts to "snoop" the routing protocols.
- Router advertisements carry link-layer addresses; no additional packet exchange is needed to resolve the router's link-layer address.

- Router advertisements carry prefixes for a link; there is no need to have a separate mechanism to configure the netmask.
- Router advertisements enable address autoconfiguration.
- Routers can advertise a Maximum Transmission Unit (MTU) for hosts to use on the link, ensuring that all nodes use the same MTU value on links lacking a well-defined MTU.
- Address resolution multicasts are "spread" over 4 billion (2^32) multicast addresses greatly reducing address resolution related interrupts on nodes other than the target. Moreover, non-IPv6 machines should not be interrupted at all.
- Redirects contain the link-layer address of the new first hop. Therefore, separate address resolution is not needed upon receiving a redirect.
- Multiple prefixes can be associated with the same link. By default, hosts learn all on-link prefixes from router advertisements. However, routers may be configured to omit some or all prefixes from router advertisements. When this happens, hosts assume that destinations are off-link and send traffic to routers. A router can then issue redirects as appropriate.
- Unlike IPv4, the recipient of an IPv6 redirect assumes that the new next-hop is on-link. In IPv4, a host ignores redirects specifying a next-hop that is not on-link according to the link's network mask.
- Neighbor Unreachability Detection is part of the base significantly improving the robustness of packet delivery in the presence of failing routers, partially failing or partitioned links and nodes that change their link-layer addresses.
- Unlike ARP, Neighbor Discovery detects half-link failures (using Neighbor Unreachability Detection) and avoids sending traffic to neighbors with which two-way connectivity is absent.
- Unlike in IPv4 Router Discovery the Router Advertisement messages do not contain a preference field. The preference field is not needed to handle routers of different "stability"; the Neighbor Unreachability Detection will detect dead routers and switch to a working one.
- The use of link-local addresses to uniquely identify routers (see explanation of link-local addresses in appendix) makes it possible for hosts to maintain the router associations in case of the site renumbering to use new global prefixes.

Header Format Enhancement

IPv4 Header Format

Version	Header Length	Type of Service	Datagram Length
Datagram ID		Flags	Flag Offset
Time to Live	Protocol	Checksum	
Source IP Address			

Destination IP Address
Data Portion of the Datagram
Payload

IPv6 Header Format

Version	Traffic Class	Flow Label	
Payload Length		Next Header	Hop Limit
Source IP Address			
Destination IP Address			
Data Portion of the Datagram			
Payload			

The IETF doubled the basic IP header size while decreasing the number of fields within it. This new header format is the key to new IPv6 service capabilities (see Appendix B, IP Header Formats). Whereas the IPv4 Header contains 20 octets plus options of thirteen fields, including three flag bits, the IPv6 Header contains forty Octets plus only eight fields. This creates a better performing, more efficient protocol. The fixed Size IPv6 Header, unlike IPv4, options are not limited at 40 bytes. Fewer fields in basic header allow for faster processing of basic packets and unlike IPv4, there is no checksum. IPv6 utilizes a 64 Bit Alignment Header/Options which is efficient option processing. The option fields processed only when the option is present. Processing of most options is limited and performed only at the destination. Additional benefits are no fragmentation in the network, which means more router cycles available for forwarding packets, not resolving addresses.

Protocol data unit (PDU) is term that describes a chunk of data that could be of any one of many protocols. Each protocol refers to its PDUs differently. IP refers to its PDU as a datagram or packet. PDUs consist of a header at the start followed by the payload of data. As PDUs are processed through the layers of the OSI model, the fields in the header tell the network how to process it. The basic design of the IPv4 header restricts its ability to be efficiently routed and contains unnecessary field that cause wasted processing on each node it passes. (See IP header formats in appendix B). The IPv4 header, state-of-the-art when it was developed in 1975 is outdated. Technologies have evolved and better ways of doing things discovered. The IP header needs three simplifications to facilitate network security as well as efficiency. They are:

1. Assign a fixed format to all packet headers
2. Do away with the checksum field
3. Redesign the hop-by-hop routing system.

The fixed size IPv6 header unlike IPv4's variable-length concept is borrowed from the asynchronous transfer mode (ATM) architecture. ATM is a streamlined high-speed protocol that is unencumbered with overhead processing requirements. ATM transmits its data in fixed size chunks called cells. The fixed size of the cells simplifies the processing required at each node on the network and thereby supports high-speed transmission. IPv6's fixed size packet header simplifies processing by every node in the path. A uniformly sized header reduces the need for a header length field and the processing it takes to determine when the entire packet is received.

Authentication and Confidentiality Enhancement

Authentication guarantees a message's origin and that its content has not been altered. Confidentiality guarantees that the message cannot be read by a third party. IPv6 provides these security features, standardized and mandated by the specification. All implementations of IPv6 must offer them. One benefit of the new technology is that current applications do not need to be modified to use IPv6 security enhancements.

The "IPv6 Authentication Header", is an extension header that provides authentication and integrity to IP datagrams ensuring that the sender and receiver are who they claim to be. While the authentication header design calls for it to be algorithm-independent and will support many different authentication techniques, the use of MD5 is proposed to help ensure interoperability within the worldwide Internet. MD5 is the latest in a series of techniques in which the function takes an arbitrary-length message and transforms it into a fixed-length quantity. MD stands for message digest, which is simply a hash function. The MD5 value is entered into the authentication header at the source assuring that the packet was sent by the originating source. If an intermediary were to modify the packet in route to its destination, the destination would receive an error message stating that the value of the MD5 and the packet's content do not match. This eliminates a significant portion of network attacks, including host masquerading attacks. From a network design standpoint, the placement of authentication at the Internet layer can help provide host origin authentication to those upper layer protocols, services and applications that currently lack meaningful protections. In addition, this mechanism is exportable outside the United States because it only provides authentication and integrity, and specifically does not provide confidentiality with sophisticated encryption. The authentication header does not transform the data in the packet. It is in the clear visible to any unscrupulous individual with a network sniffer.

The "Encrypted Security Payload Header" also referred to, as the "IPng Encapsulating Security Header" is a second extension header provided with IPv6. This feature provides integrity and confidentiality to IPv6 datagrams. It is simpler than some similar security protocols such as SP3D, ISO, or NLSP but remains flexible and algorithm-independent.

To achieve interoperability within the global Internet, the use of DES CBC[1] is being used as the standard algorithm for use with the IPng Encapsulating Security Header. The ESP header is always the last header in the chain of IPv6 extension headers. More correctly, it is the last header to remain visible once encryption is applied.

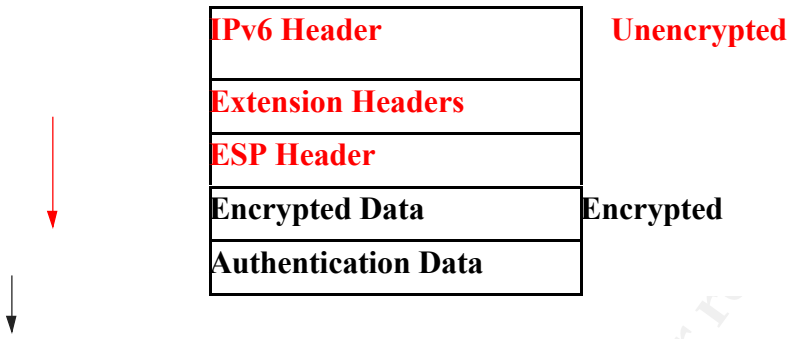

An encrypted packet using the ESP header

The ESP header was redesigned in 1997 revision of the Internet security standards. It now includes a checksum and a sequence number. The sequence number serves the same functions as the authentication header. A source produced hash function protects the receiver from intermediary tampering commonly called a replay attack. The authentication checksum carried after the encrypted data is used to protect the receiver from a type of network attack where the data is modified or truncated. An exact match of the resolved authentication data assures accuracy of the data and the sequence number.

Native IPSec Support Enhancement

IPSec is a requirement of the IETF's specification for IPv6. It is an enhancement to the Internet Protocol that provides encryption and authentication at the transport layer (layer 3 of the OSI model).[2] Users and application are generally unaware of the IPSec security that takes place throughout the network. Often even unaware, that IPSec is tunnelling their data through insecure networks.

The way IPSec works is with IPv6 extension headers. The authentication header, AH, the encapsulation security payload header, ESP, and the Internet key exchange, IKE, perform the authentication, encapsulation, encryption and functions collectively or individually. It is possible to only use only authentication by adding the AH header. While this does not protect the data with encryption, it does use less overhead and therefore process faster.

---

[1] IEFTF Network Working Group http://www.ietf.org/rfc/rfc2405.txt

[2] Napier, Duncan Administering Linux IPSec Virtual Private Networks http://www.samag.com/documents/s=4072/sam0203c/sam0203c.htm

The diagram below shows the AH header placement in an IP packet.

IPv6 Header or Original IPv4 Header	AH	TCP Header + Data

Hash Checksum

The authentication header will protect the entire packet even the part before the AH header by verifying that the sender is who he or she says they are. AH is made up of a 96-bit field that contains the number of the next header, the length of the authenticated payload, 16 reserved bits, a 32-bit security parameter index, SPI and a 32-bit sequence number. The sequence number field was added in 1997 to stop a security vulnerability referred to as the "replay attack" in which the hacker obtains a copy of a valid authenticated packet and replays it. The AH carries a hash value as the authentication information, typically a message digest algorithm, MD5, or secure hash algorithm, SHA. Although the AH does not perform encryption, it should be sufficient to prevent address spoofing attacks and prevent hackers from stealing connections.

The ESP header may be used in conjunction with or without the AH header because it provides both encryption and authentication. The encryption technique is symmetric (a system where the same key is used to encrypt and decrypt a message). A 3DES block cipher is generally used. The encryption key is shared between users using IKE. The diagram "An encrypted packet using the ESP header" above demonstrates an encrypted packet composition. Note that the entire packet is not encrypted. In the unencrypted ESP header there is a hash function that provides authentication for the entire packet.

IKE does more than simple key exchange. It authenticates the devices at both ends of VPN tunnels, determines the optimal encryption and authentication algorithms to use in a session, and generates and manages encryption keys. Although IKE is the current IPSec method of key exchange, the IETF is devising a new protocol that is more flexible, could lead to equipment that is more interoperable, supports better security, and is easier to configure.[3]. As IKE establishes a session, it negotiates the encryption key handling and renewal parameters for the connection. Authentication can be in the form of passwords, digital signatures, or RSA cryptographic keys that guarantee the users' identities. IKE uses the Diffie-Hellman method also called exponential key agreement protocol of key agreement. Diffie-Hellman allows two users to exchange a secret key over an insecure medium without any prior secrets. The protocol uses two public parameters prime p and generator g. Parameter p is a prime number and parameter g is an integer with a value less than p, and also has the following property: For every number n between 1 and p-1 inclusive, there is a power k of g such that $n = g^k \bmod p$.

[3]NetworkFusionNews http://www.nwfusion.com/news/2001/1210ike.html

Example[4]: Alice and Bob want to share a secret key using the Diffie-Hellman key agreement protocol. They proceed as follows:

1. Alice generates a random private value *a*
2. Bob generates a random private value *b*. Both *a* and *b* are drawn from the set of integers {1, …, *p*-2}. Then they derive their public values using parameters *p* and *g* and their private values. Alice's public value is $g^a \bmod p$ and Bob's public value is $g^b \bmod p$.
3. They exchange their public values.
4. Alice computes $g^{ab} = (g^b)^a \bmod p$
5. Bob computes $g^{ba} = (g^a)^b \bmod p$
6. Since $g^{ab} = g^{ba} = k$, Alice and Bob now have a shared secret key *k*.

Added security is provided by re-keying the session periodically. This requires that a new authentication key be negotiated at frequent time intervals.

IPSec inserts the encrypted data payload and appropriate headers into an existing IP packet thereby allowing IPSec to travel through any IP network. This is the primary reason for its ubiquitous use. All IP networks even Ipv4 can pass IPSec traffic without modifying the underlying network infrastructure. IPSec has both hardware and software implementations and is supported by all major vendors. Although all vendors are not necessarily 100 percent compatible due in part to different interpretations of the RFCs.

IP is Enhanced by Better VPN support

A portion of IPv6's appeal is from its autoconfiguration capability. Automatic address configuration allows an IP address to be dynamically assigned. Although that was possible under the IPv4 architecture, IPv6 Autoconfiguration makes things even simpler by automating the entire IP assignment and detection process throughout the entire network from network interface cards (NICs) to hubs, routers, switches, and gateways. Corporations are taking advantage of autoconfiguration of routers to create secure virtual private networks (VPNs), which in turn is facilitating automation of business-to-business e-commerce. In addition, IPv6's routing features support the use of multiple IP addresses on a single physical address, which means that instead of spending an inordinate amount of time tweaking network address translation (NAT) tables and domain name server (DNS) entries, you can simultaneously deploy and support IP addresses for Internet connection and your VPN.

Enhanced Network Numbering Scheme

[4] RSA Laboratories http://www.rsasecurity.com/rsalabs/faq/3-6-1.html

The increased address space that IPv6 offers is the most obvious enhancement over IPv4. While Ipv4 is based on a 32-bit wide address, IPv6 uses a 128-bit. This enlarged address space is the reason that insecure protocols such as NAT do not have to be used anymore. The 128-bit address provides for full, unconstrained IP connectivity for today's IP-based machines as well as upcoming mobile devices like PDAs and cell phones. All of these mobile devices will benefit from full IP access through general packet radio service (GPRS) and Universal Mobile Telecommunications System UMTS.

The following is a diagram to help explain how the IPv4 32 bit address is determined.

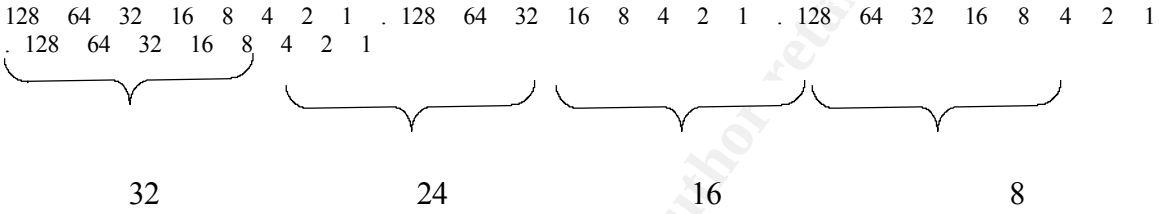

The figure below contains an overview of how IPv4 uses the 32 addressing bits to establish network classes, Network addresses, and host addresses.

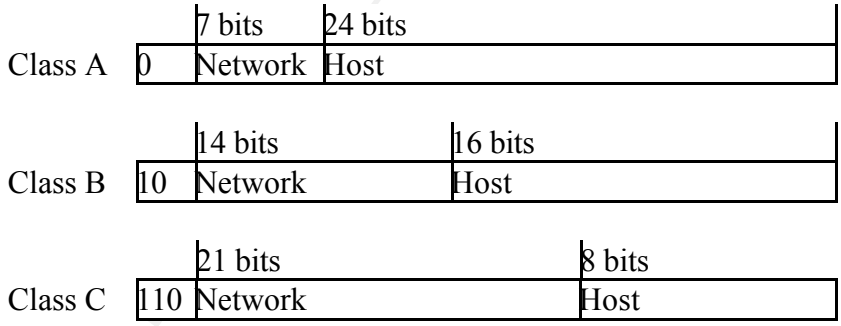

IPv4 network and host identifier in the IPv4 32 bit-addressing scheme

IPv6 has 128 address bits broken down into eight sixteen-bit subsets.

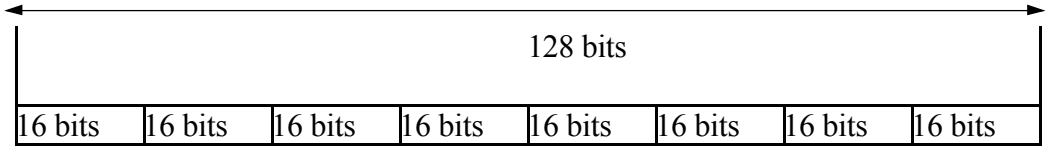

IPv6 address format

A typical IPv6 address looks like this: 3FFE:0301:DEC1::0A00:2BFF:FE36:701E

Prefix Interface ID

3FFE:0301:DEC1::0A00:2BFF:FE36:701E

Each of the eight 16-bit segments is identified by four hexadecimal digits delineated by colons. A double colon may be used to represent a segment with a value of zero. The IPv6 address may include the legacy IPv4 address as the least significant 32 bits and will be one of the following three types or scopes from rfc2373[5]

1. **Global**

3 bits	13 bits	8 bits	24 bits	16 bits	64 bits
FP	TLA ID	Res	NLA	SLA ID	Interface ID

Where

FP	Format Prefix (3 bit) for Aggregatable Global Unicast Addresses
TLA ID	Top-Level Aggregation Identifier
RES	Reserved for future use
NLA ID	Next-Level Aggregation Identifier
SLA ID	Site-Level Aggregation Identifier
INTERFACE ID	Interface Identifier

2. **Site-Local**

10 bits	38 bits	16 bits	64 bits
1111111011	0	Subnet ID	Interface ID

Site-Local addresses are used for addressing inside of a site without the need for a global prefix. Routers do not forward any packets with site-local source or destination addresses outside of the site.

3. **Link-Local**

10 bits	54 bits	64 bits

[5] RFC2373, IP Version 6 Addressing Architecture 7/1998 http://www.landfield.com/rfcs/rfc2373.html

1111111010	0	Interface ID

Link-Local addresses are used for addressing on a single link for purposes such as auto-address configuration, neighbor discovery, or when no routers are present. Routers do not forward any packets with link-local source or destination addresses to other links.

The figure below is an example of what the Microsoft Windows 2000 client sees when requesting his IPv6 interface information. The IPv6 Link-local addresses have the prefix FE80::/64. Microsoft's configuration uses Link-local as its default setting for all IPv6 addresses. The link local addresses is displayed on Windows 2000 Operating Systems by typing the IPv6if command and looking for an interface with a link-level address of the form aa-bb-cc-dd-ee-ff as shown in figure 3. The preferred address is the link-local address for the interface.

An Example of the Windows 2000 IPv6 if command

```
C:\>IPv6 if
        Interface 4 (site 1): Local Area Connection
        uses Neighbor Discovery
        link-level address: 00-10-5a-aa-20-a2
        preferred address fe80::210:5aff:feaa:20a2, infinite/infinite
        multicast address ff02::1, 1 refs, not reportable
        multicast address ff02::1:ffaa:20a2, 1 refs, last reporter
        link MTU 1500 (true link MTU 1500)
        current hop limit 128
        reachable time 43500ms (base 30000ms)
        retransmission interval 1000ms
        DAD transmits 1
        Interface 3 (site 1): 6-over-4 Virtual Interface
        uses Neighbor Discovery
        link-level address: 10.0.0.2
        preferred address fe80::a00:2, infinite/infinite
        multicast address ff02::1, 1 refs, not reportable
        multicast address ff02::1:ff00:2, 1 refs, last reporter
        link MTU 1280 (true link MTU 65515)
        current hop limit 128
        reachable time 34000ms (base 30000ms)
        retransmission interval 1000ms
        DAD transmits 1
        Interface 2 (site 0): Tunnel Pseudo-Interface
        does not use Neighbor Discovery
        link-level address: 0.0.0.0
        preferred address ::10.0.0.2, infinite/infinite
        link MTU 1280 (true link MTU 65515)
        current hop limit 128
        reachable time 0ms (base 0ms)
        retransmission interval 0ms
        DAD transmits 0
        Interface 1 (site 0): Loopback Pseudo-Interface
        does not use Neighbor Discovery
        link-level address:
        preferred address ::1, infinite/infinite
```

```
link MTU 1500 (true link MTU 1500)
current hop limit 1
reachable time 0ms (base 0ms)
retransmission interval 0ms
DAD transmits 0
```

IPv6 Security Concerns

Although IPv6 provides many new features and enhancements there is concern that hackers could exploit these and compromise the network. An example[6] of this is provided by Richard Rager of www.penguinman.com. Richard Rager is a security consultant in Columbus, Ohio with a broad background in Internet Technologies. He is concerned that because IPv6 can be sent through IPv4 networks, this might allow someone outside a protected network to use these new protocols and gain information about the private network from the networks multi-homed edge router and its firewall. An example he offers is the ease of an email server that is used from both outside and inside the private network. He states "These internal internets (any thing goes zones or DMZ) are a need for your Internet presence and can cause a security risk if some one can redirect the resource to tell them about your internal network." In addition, he states that most routers now have IPv6 enabled and no one is using the security feature of IPSec, so their networks are vulnerable. The vulnerability happens when the transmission of IPv6 over IPv4 domains do not use explicit tunnels, neighbor discovery, and generic packet tunnelling. The attack sequence that exploits this condition would proceed as follows:

1. First, it uses the IPv6. Protocols of sending IPv6 over IPv4 domains without explicit tunnels.
2. Then use the neighbor discovery to find other routers in your network.
3. The next step is to use generic packet tunnelling to come inside the network from a trusted router, thereby compromising the network.

Conclusion

The security features built into IPv6 are necessary and overdue. The possibility or perhaps probability that it will be exploited is not a huge reason for concern because as with all technologies, it will continue to improve with each attack. A vulnerability will be exploited, a fix implemented and the network will then be a little more secure. I view the security features of IPv6 as another layer to add to the security in depth architecture. My only criticism of it is although the IPv6 architects designed the security features to be user and application unaware, I believe that the user needs a security status of the network that provides him or her with that "warm fuzzy feeling "that the network is secure.

[6]Rager, Richard IPv6 and network security http://penguinman.com/ipv6.html

APPENDIX A THE Architects of IPv6

The Internet Engineering Task Force (IETF) is the agency responsible for IPv6 development. It is comprised of four groups:

1. The Internet Society (ISOC) and its Board of Trustees is a professional society that is concerned with the growth and evolution of the worldwide Internet, how it is used, and its social, political, and technical issues. The ISOC Trustees are responsible for approving appointments to the IAB from among the nominees submitted by the IETF nominating committee.

2. The Internet Architecture Board (IAB) is the technical advisory group of the ISOC. It is chartered to provide oversight of the architecture of the Internet and its protocols, and to serve, in the context of the Internet standards process. The IAB is responsible for approving appointments to the IESG from among the nominees submitted by the IETF nominations committee.

3. The Internet Engineering Steering Group (IESG) is responsible for technical management of IETF activities and the Internet standards process. As part of the ISOC, it administers the process according to the rules and procedures that have been ratified by the ISOC Trustees. The IESG is directly responsible for the actions associated with entry into and movement along the Internet "standards track", including final approval of specifications as Internet Standards.

4. The IETF itself is divided into eight functional areas. They are: Applications, Internet, IP: Next Generation, Network Management, Operational Requirements, Routing, Security, Transport and User Services. Each area has one or two area directors. The area directors, along with the IETF/IESG Chair, form the IESG. Fred Baker is the current IETF/IESG chair. Each area has several working groups. A working group is a group of people who work under a charter to achieve a certain goal. That goal may be the creation of an Informational document, the creation of a protocol specification, or the resolution of problems in the Internet. Most working groups have a finite lifetime. Once a working group has achieved its goal, it disbands. As in the IETF, there is no official membership for a working group. Unofficially, a working group member is somebody who is on that working group's mailing list; however, anyone may attend a working group meeting. Areas may also have Birds of a Feather (BOF) sessions. They generally have the same goals as working groups, except that they have no charter and usually only meet once or twice. BOFs are often held to determine if there is enough interest to form a working group[7].

The work group that is controlling the development of IPV6 is called IPng. The next

generation of the Internet Protocol (IPv6) is intended to support Internet traffic for many years into the future by providing enhancements over the capabilities of the existing IPv4 service. This working group will produce specifications for the core functionality of that service. The working group shall carry out the recommendations of the IPng Area Directors as outlined at the July 1994 IETF and in ``The Recommendation for the IP Next Generation Protocol," Internet-Draft, (draft-ipng-recommendation-00.txt), September 1994.[8] Members of corporations that are leaders in the manufacturing IPv6 products interestingly enough head the work group.
Chairmen of the work group are:

- Bob Hinden, Chief Technical Officer (CTO) of Nokia.
- Steve Deering, Technical Leader at Cisco Systems

Internet Area Directors for the work group are:

- Thomas Narten, TCP/IP strategist with IBM Corporation and Adjunct Assistant Professor in the Duke Computer Science Department.
- Erik Nordmark, TCP/IP Engineer at Sun Microsystems, Inc.

References

Callon, R. and Haskin, D. “Routing Aspects of IPv6 Transition”, RFC2185 September 1997.

Crawford, M. “Transmission of IPv6 Packets over Ethernet Networks”, RFC2464, December 1998

Deering, Steve, and Robert Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC2460, December 1998.

Gilligan, R. and Nordmark E. “Transition Mechanisms for IPv6 Hosts and Routers”, RFC1933, April 1996

Grossetete, Patrick and McNealis, Martin. "Cisco Systems Statement of Direction IP Version 6", June 2000, http://www.cisco.com/warp/public/732/IPv6/IP_Vers6_SD_0622.pdf

Huitema, Christian. IPv6 The New Internet Protocol. Upper Saddle River New Jersey: Prentice Hall Incorporated, Second Edition, 1997

[7] The Tao of IETF Web Site. http://www.ietf.org/tao.html#What_Is_IETF

[8] Hinden, Bob IETF IPng Work Group Chairman

IETF Network Working Group http://www.ietf.org/rfc/rfc2405.txt

Kaeo, Merike. Designing Network Security. Indianapolis: Macmillan Publishing, 1999

Loshin, Pete. IPv6 Clearly Explained. San Francisco: Morgan Kaufman Publishers, Incorporated, 1999

Miller, Mark. Implementing IPv6. Foster City California: MBT Books, Second Edition, 2000

Napier, Duncan Administering Linux IPSec Virtual Private Networks http://www.samag.com/documents/s=4072/sam0203c/sam0203c.htm

Narten, T, Simpson W. and Nordmark, E" Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.

Narten, T, and Thomson S. 1971 "IPv6 Stateless Address Autoconfiguration", RFC1971, August 1996

NetworkFusionNews http://www.nwfusion.com/news/2001/1210ike.html

Rager, Richard IPv6 and network security http://penguinman.com/IPv6.html

RFC2373, IP Version 6 Addressing Architecture 7/1998 http://www.landfield.com/rfcs/rfc2373.html

RSA Laboratories http://www.rsasecurity.com/rsalabs/faq/3-6-1.html

The IPv6 Forum Web Site, "The New Internet: Internet for Everyone Quality, Mobility, Security", http://www.IPv6forum.com/

The Tao of IETF Web Site. http://www.ietf.org/tao.html#What_Is_IETF