A Technical Overview of the IPsec Internet Protocol Suite

John P. Pironti, CISSP
Genuity, Inc.
Principal Enterprise Solutions Architect
Principal Security Consultant

Version 1.0 – December 15, 2001

What is the IPsec Internet Protocol Suite?

The IPsec Internet protocol suite has revolutionized Internet Protocol (IP) security. The IPsec protocol suite utilizes cryptographic techniques to ensure data confidentiality, and digital signatures to authenticate the source of the data transmission. IPsec also brings a new level of interoperability to the Internet that never existed before. Because it is an IETF standard it does not rely on proprietary protocols or techniques to establish secure links between network nodes. By utilizing IPsec in virtual private networking solutions organizations can exchange sensitive data over public networks with the knowledge that the parties they are exchanging the data with are the intended receivers, that the data was kept confidential in transit, and that the data did not change during transmission.

IPsec is primarily used for Virtual Private Network applications at this time. This is due to the fact that it offers an organization the ability to leverage their existing Internet connectivity for secure communications with remote locations and business partners. Because of the interoperable nature of the IPsec protocol, business partners can communicate with each other without having to purchase identical equipment with each other. IPsec also provides strong authentication techniques, data integrity, and data confidentiality services that allow for a higher level of assurance to be achieved than most organizations currently have when performing these types of communication.

The way that the IPsec protocol suite achieves its higher level of assurance for data transport is through the use of multiple protocols including authenticated header (AH), encapsulated secure payload (ESP), and Internet key exchange (IKE). Each of these protocols can be used independently to provide specific heightened security capabilities, but when used together they create an extremely powerful suite of capabilities to ensure high levels of data security during data transport over public networks such as the Internet.

Key Elements of the IPsec Protocol Suite

Authenticated Header

One of the protocols with the IPsec protocol suite is the Authenticated Header (AH). The AH protocol provides data integrity, authentication, and optionally anti-replay capabilities. It does not provide data confidentiality. AH provides integrity protection for the payload and most of the IP header. AH will ensure that the fields that identify the source and destination of IP packet are valid and have not been modified during transmission. This is accomplished by using digital signature techniques or one-way hash functions. There are six parts that make up the AH:

- Next Header – Indicates what the higher-level protocol following the AH is (i.e. ESP, TCP).

- Payload Length – This is an 8 bit field that specifies the size of the AH.

- Reserved – This is a placeholder for a future function.

- Security Parameter Index (SPI) – SPI is a pseudo-random 32-bit number that specifies the security settings that are being used by the transmitter to communicate with the receiver. This includes the encryption algorithms that are being used, which encryption keys are being used, and the information about the validity period for these encryption keys.

- Sequence Number – The sequence number is a counting mechanism that increases incrementally each time a packet is transmitted using the parameters setup in the SPI. It identifies the packets and enumerates the number of times packets have been transmitted using the same SPI. The sequence number protects receiving nodes from replay attacks where an attacker will copy a packet and then resend it in an attempt to confuse the receiver.

- Authentication Data – This is the Integrity Check Value (ICV) for the packet. The originator will create a keyed one-way hash of the packet payload and attach this hash value to the packet as the authentication field. The IPsec standard specifies that the HMAC symmetric signature scheme along with hash algorithms SHA-1 or MD-5. The recipient can then validate that integrity of the payload data by hashing the payload data once it has been decrypted with the same one-way hash algorithm, which the originator used. If the two hash values are identical than the recipient can be confident that the data was not modified during the transmission. However, since the data was not encrypted this does not ensure the confidentiality of the payload data only the integrity.

Encapsulated Security Payload

The encapsulated security payload (ESP) is the portion of the IPsec protocol suite that addresses the confidentiality of the data that is being transmitted as well as offers authentication capabilities. ESP utilizes symmetric encryption techniques to encrypt the IP packet payload. The symmetric encryption algorithms that must be supported in order to be compliant to standard are DES, 3DES, RSA, CAST, and Blowfish. The ESP will not encrypt the IP header or information, which includes the information required for routing. It will only encrypt the packet payload, which will ensure the confidentiality of the data. There are six elements which make up the ESP which include:

- Security Parameter Index (SPI) – SPI is a pseudo-random 32-bit number that specify the security settings that are being used by the transmitter to communicate with the receiver. This includes the encryption algorithms that are being used, which encryption keys are being used, and the information about the validity period for these encryption keys.

- Sequence Number – The sequence number is a counting mechanism that increases incrementally each time a packet is transmitted using the parameters setup in the SPI. It identifies the packets and enumerates the number of times packets have been transmitted using the same SPI. The sequence number protects receiving nodes from replay attacks where an attacker will copy a packet and then resend it in an attempt to confuse the receiver.

- Payload Data – This is the data that is contained within the IP packet.

- Padding – Padding is used to prevent attackers from using sniffers to estimate that amount of data that is being transmitted in some encryption algorithms. Padding techniques utilize random data that can range from 0 – 255 bytes in length, and generally insert this random data after the valid payload.

- Pad Length – This field specifies the amount padding that is in place in a specific packet payload so it can be properly identified and stripped from the packet during the decryption process.

- Next Header – The next header field in the IP packet identifies the types of data that is being carried within the packet as well as the protocols are being utilized.

The SPI and the Sequence number are not encrypted, but are authenticated. The Payload Data, Padding, Pad Length, and Next Header are encrypted to ensure confidentiality. There is an optional field within the ESP as well, which is the authentication field. The authentication field contains the Integrity Check Value (ICV) and is calculated once the encryption function has been completed on the ESP.

The ICV is a digital signature, which is computed using the ESP (not including the authentication field itself). The originator will create a keyed one-way hash of the packet payload and attach this hash value to the packet as the authentication field. The IPsec standard specifies that the HMAC symmetric signature scheme along with hash algorithms SHA-1 or MD-5. The recipient can then validate that integrity of the payload data by hashing the payload data once it has been decrypted with the same one-way hash algorithm, which the originator used. If the two hash values are identical than the recipient can be confident that the data was not modified during the transmission.

There are also two different modes that ESP can operate in, transport and tunnel mode. In transport mode the packet payload is the only part of the packet that is encrypted which means that the original packet header left undisturbed. This has the advantage of reducing the overhead added to the packet size, and allowing the routing elements in the network to view the final destination address of the packet. This information can optionally can be used for QOS services and routing services by the routing elements that are encountered during packet transmission across the network. The downside to using transport mode is the exposure of the header data. An attacker could potentially perform traffic analysis on these packets during transmission and gain

insight about the transmission and gain insight about the activities that are being performed.

ESP tunnel mode encrypts the both the packet and the payload in a new IP packet. In this case all traffic is passed to the IPsec device and it then acts as a proxy element for this traffic. The IPsec device will perform all encryption and encapsulation activities without having to modify any of the other systems within the network. Tunnel mode also protects against traffic analysis since the attacker will only be able to decipher the tunnel end points, which will be the IPsec device. They will not be able to decipher the actual source address or destination address.

Internet Key Exchange

The Internet Key Exchange protocol (IKE) is the method used for public key exchange, secure association (SA) parameter negotiation, identification and authentication. IKE is actually a hybrid of three key management schemes. Internet Security Association and Key Management Protocol (ISAKAMP), Oakley, and SKEME. IKE operates creates an authenticated and secure tunnel between the originator and the receiver and then negotiates the security associations for IPsec. This process starts by the originator and the receiver authenticating themselves to each other to share keys with each other. In order to accomplish this the two parties must agree on a common authentication protocol through a negotiation process. The are two common methods for this negotiation:

- Pre-Shared Keys (shared secret) – The same key is pre-installed on both the originator and receiver's host. Both parties will then initiate an asymmetric key exchange with each other utilizing the Diffie-Hellman key agreement algorithm. This will allow both parties to establish an SA where they can transfer data securely and setup an initial set of parameters for data transfer via the SPI. Authentication will then be performed by both entities computing a one-way keyed hash of the pre-shared key and transmitting the resulting hash value to each other via the previously established SA. Both parties will then perform a one-hash operation using the previously agreed upon and utilized one-way hash algorithm and compare the output values. If the values are identical then both parties have the same secret which means that they have authenticated to each other. Once authentication has been completed and agreed up on by both parties the key exchange for the symmetric encryption algorithm can take place via the previously established SA.

- Public Key Cryptography – A key exchange using public key cryptography involves multiple steps. The first step in the process is for both parties to initiate an asymmetric key exchange utilizing the Diffie-Hellman key agreement algorithm. This will allow both parties to establish an SA where they can transfer data securely and setup an initial set of parameters for data transfer via the SPI. Once this SA has been established then both parties can transmit their public key data through this SA to each other. This will then allow for strong authentication

to take place if you are using a properly formatted X.509 certificate. Once the public keys have been exchanged and a new SA has been established the key exchange for the symmetric encryption algorithm can take place via the new SA.

The IKE protocol functions in two phases. The first phase is where the two entities that are attempting to communicate securely with each other setup a secure channel to negotiate security associations with each other. This phase does not take into account actual authentication of the two entities. In phase two of IKE the actual SA's are negotiated between the entities via the secure channel that was created in phase one.

There are also two modes specified in the IKE protocol, main mode and aggressive mode. Main mode performs the key exchange separately from the SA proposal activity in order to conceal the identity of the IKE agent. This provides for peer authentication of IKE agents. Aggressive mode reduces the number of messages involved with an IKE exchange, but does not conceal the IKE agent identities. The primary differences between the two modes are the amount of operations that take place during the IKE negotiation process. If you are establishing a secure tunnel with a recipient for the first time or in a highly volatile environment you would most likely want to use the main mode of operation. If you have already established SA relationships with a recipient before and you feel comfortable with their environment you may choose to use aggressive mode to reduce the number of processes involved with the IKE activity.

Security Association

The best way to ensure effective communication using authentication and encryption services is to have a mechanism that accounts for the individual elements of each connection. In the IPsec protocol suite this mechanism is known as the Security Association (SA). An SA contains all the data involved in communicated with another node securely. These elements include:

- The mode and keys of the authentication algorithm used in the AH.
- What protocols, encryption algorithm, and keys will be used to authenticate the communication.
- The encryption algorithm mode and keys used with the associated with the encryption algorithm for the ESP.
- Cryptographic synchronization information, including presence and size, to be used with the chosen cryptographic algorithm.
- Specification of the timing for key change.
- Validity period of the keys that are being utilized.
- Validity period of the individual SA that is being established.

- The mode and keys of the authentication algorithm used in the ESP.
- Source address for the SA that is being established.
- Sensitivity level descriptor.

The individual security associations that are used for communication paths that you establish to different nodes can be different for each node. This allows you the flexibility to have different levels of security based on the security policy criteria that you have setup with the entity that you are going to communicate with. An example where this can be used is when you are setting up an encrypted IPsec based Virtual Private Network (VPN) with various corporate offices that are located in multiple countries at the same time. Each country will have their own legal statutes that specify what types of encryption and what size encryption keys can be utilized within their borders. Using an encrypted IPsec VPN you can compensate for this fact by using individual SA's to establish different levels of encryption for each connection to your location. This will alleviate the problem of having to set all the connection settings to the same algorithm and least common denominator for key size, which was the traditional solution to this problem.

IPsec and Windows 2000

One of the great advances in the security architecture of Microsoft Windows 2000 was the addition of the IPsec protocol suite. Microsoft integrated IPsec with both Windows 2000 Domain and Active Directory services. The IPsec protocol suite is used in Windows 2000 to secure authentication and data communication between network nodes. In its current release at the time of this article (Service Pack 1) Windows 2000 only supports IPsec virtual private networks in a client-to-gateway and gateway-to-gateway configurations in conjunction with the use of the layer two tunneling protocol (L2TP).

Windows 2000 integrates the IPsec protocol suite within its Active Directory and Domain services by providing policy distribution for the IPsec entities through these mechanisms. IPsec policies are distributed to domain members through Windows 2000 group policy. Because this provides local configuration policies as well, a host is not required to be a member of a domain to take advantage of this capability. Once the IPsec policies have been disseminated to the hosts within a particular network they can take advantage of the IPsec protocol suite security capabilities. The hosts will utilize the IKE protocol to negotiate SA's with each other and setup secure communication channels. The keys and SA's will automatically refresh themselves according to the policy that has been installed on the host.

The cryptographic algorithms that are supported within Windows 2000 are a subset of those required to be compliant to the IPsec standard (www.ietf.org/rfc/RFC2401.txt). The supported encryption algorithms are 56bit DES and 168bit 3DES, integrity algorithms are SHA-1 and MD5, and key exchange algorithm

is Diffie-Hellman. The 168bit 3DES algorithm can be used within the continental United States and any countries in which the US has certified it to be exportable to. The corresponding country must allow for the import of the 168bit 3DES algorithm as well. In order to compensate for those countries where 168bit 3DES is not allowed, you can normally use 40bit DES, but as always it is suggested that you check with the local government to confirm this.

Windows 2000 does not come natively equipped with the code for 168bit 3DES. You have to acquire the high encryption pack from Microsoft to enable this function. Microsoft will only allow individuals from authorized countries to download the upgrade software, so you can be somewhat assured that you are authorized to use the 168bit 3DES if you are allowed to download the upgrade.

One of the more interesting parts of the IPsec implementation with Windows 2000 is its integration with L2TP. IPsec is widely utilized in VPN solutions because of its security features as well as its interoperability capabilities. When setting up Windows 2000 in a client-to-gateway or gateway-to-gateway VPN Microsoft utilizes the L2TP protocol as a wrapper for the IPsec protocol suite. The reason for this is to allow for transport of IPsec data through gateways that do not support the IPsec protocol suite but do support L2TP. While this may be an advantage for these gateways, it does force these configurations not to be compliant to the IPsec standard (RFC 2401). This is due to the fact that the packet has been modified after the IPsec device has processed it.

Final thoughts

The IPsec protocol suite has filled the much needed security gap that exists in the TCIP/IP protocol. It not only promises to provide heightened levels of security, but also allows for unsurpassed levels of interoperability. It can promise these capabilities because it is an accepted Internet standard (www.ietf.org/rfc/RFC2401.txt). It is no surprise that security vendors and software vendors alike have rushed to implement this new protocol into their software. The integration of IPsec within Microsoft Windows 2000 was a natural fit. The press and the hackers of the world have long since branded Windows NT as an insecure operating system with a lot to be desired. Microsoft Windows 2000 promised to be a big step towards changing that branding, and with the integration of the IPsec protocol suite it is my opinion that they are on their way.

IPsec does have some problems that still need to be overcome as well. The most criticized portion of the IPsec protocol suite is the IKE protocol. This is due to its complexity. It does not lend itself to be easily integrated into web phones, personal data assistants, and IP enabled devices. The question that has to be asked in this scenario is how much security is required for these applications? IPsec and IKE allow for a highly secure data transmission capability, but not every situation needs such a high level of security. The flexibility of IPsec allows a user to tailor the security capabilities to their specific needs, but still requires numerous operations to take place in order to be properly utilized.

It is always important to remember that the software is only as secure as the administrator who configures it. Be sure to carefully review all of your configuration settings and trust relationships. You may now know whom you are communicating with and that you are communicating with them securely, but you still do not necessarily know what they are communicating to you. It is still important to only allow users access to the resources they absolutely need access to. It is also important to educate users about the capabilities and limitations of the new levels of security, which they are now utilizing. If you do not you chance your users becoming complacent with their choices of how they secure their data since they may believe you have created the magic bullet for them and are taking care of these issues. Unfortunately, no matter how many technological safeguards you have in place, you still cannot prevent a user from making a mistake.

Further Reading and Information Resources:

www.ietf.org/rfc/RFC2401.txt
www.ietf.org/rfc/RFC2409.txt
www.microsoft.com/WINDOWS2000/library/howitworks/security/sectech.asp
www.ip-sec.com