



Interested in learning more about security?

SANS Institute InfoSec Reading Room

This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.

IP Security Protocol-based VPNs

IP Security Protocol (IPSec) defines a set of protocols and cryptographic algorithms for creating secure IP traffic sessions between IPSec gateways. Enterprises can no longer keep pace with ever changing technology. Flexibility, scalability, global reach, and security are critical network components that help to define an enterprise's success. IPSec Virtual Private Networks (VPNs) allow enterprises to leverage the speed, ubiquity, and flexibility of the Internet, while gaining security beyond that which is provided by ...

Copyright SANS Institute
Author Retains Full Rights

AD

A banner advertisement for FireEye. On the left is the FireEye logo, which consists of a stylized red and white eye with a flame-like shape above it, followed by the word "FireEye" in a bold, sans-serif font. To the right of the logo is a black background with white and red text. The text reads: "Protect critical data from the cyber theft pandemic." in white, with "Protect" in red. Below that, it says "Learn how in this FireEye white paper." in white, with "white paper" in yellow. On the far right of the banner is a small image of a man in a hard hat looking at a computer screen, with a yellow bird in a cage visible in the background.

Protect critical data from the
cyber theft pandemic.
Learn how in this FireEye **white paper**.

IP Security Protocol-based VPNs

Overview

IP Security Protocol (IPSec) defines a set of protocols and cryptographic algorithms for creating secure IP traffic sessions between IPSec gateways. Enterprises can no longer keep pace with ever changing technology; legacy networks such as ATM and Frame Relay are examples of such technologies that are not able to keep up with the technological advances. Flexibility, scalability, global reach, and security are critical network components that help to define an enterprise's success. IPSec Virtual Private Networks (VPNs) allow enterprises to leverage the speed, ubiquity, and flexibility of the Internet, while gaining security beyond that which is provided by legacy network technologies. There have been many improvements in the internet including quality of service, network performance and inexpensive technologies such as DSL. But, one of the most important advances has been in security. IPSec is one of the most complete, secure, and commercially available, standards-based protocols developed for transporting data.

VPN is a shared network where private data is segmented from other traffic so that only the intended recipient has access. A key aspect of data security is that the data flowing across the network is protected by encryption technologies. Private networks lack data security, which allows data attackers to tap directly into the network and read data. IPSec-based VPNs use encryption to provide data security, which increases the network's resistance to data tampering or theft. VPNs are used for Intranets, remote access and extranets.

What is IPSec and how does it provide security?

IPSec is an Internet Engineering Task Force (IETF) standard suite of protocols that provides data authentication, integrity, and confidentiality as data is transferred between communication points across IP networks. IPSec provides data security at the Packet level. It was designed to provide authentication (verifies that the packet received is actually from the sender), integrity (ensures that the contents of the packet did not change in transit), and confidentiality (conceals the message content through encryption).

IPSec contains Encapsulating Security Payload (ESP), provides confidentiality, authentication, and integrity. ESP provides all the encryption services. It also contains Authentication Header (AH), provides authentication and integrity, which protects against data tampering and unauthorized retransmission of packets. The last component it has, is the Internet Key Exchange (IKE), which provides key management and security association management. The aforementioned features of IPSec are accomplished through packet header modification.

IPSec introduces the concept of the Security Association (SA). An SA is a logical connection between two devices transferring data. An SA provides data protection for unidirectional traffic by using the defined IPSec protocols. SA operates using modes. A mode is the method in which the IPSec protocol is applied to the packet. IPSec can be used in tunnel mode or transport mode. Typically, the tunnel mode is used for gateway to gateway IPSec tunnel protection, while

transport mode is used for host to host IPSec tunnel protection. The transport mode encapsulates only the packet's payload, the IP header is not changed. After the packet is processed with IPSec, the new IP packet contains the old IP header (with source and destination IP addresses unchanged) and the processed IP packet payload. Transport mode does not shield the information in the IP header; therefore, an attacker can learn where the packet is coming from and where it is going to. The tunnel mode implementation encapsulates the entire IP packet. The entire packet becomes the payload of the packet that is processed with IPSec. A new IP header is created that contains the two IPSec gateway addresses. The gateways perform the encapsulation/de-capsulation on behalf of the hosts. Tunnel mode ESP prevents an attacker from analyzing data and deciphering it, as well as knowing who the packet is from and where it is going.

IPSec uses the Internet Key Exchange (IKE) protocol to facilitate and automate the SA setup and the exchange of keys between parties transferring data. Using keys ensures that only the sender and receiver of a message can access it. IPSec requires that keys be re-created, or refreshed, frequently so that the parties can communicate securely with each other. IKE manages the process of refreshing keys; however, a user can control the key strength and the refresh frequency. Refreshing keys on a regular basis ensures data confidentiality between sender and receiver. IKE works in a two-phase process. The first phase sets up the actual IKE SAs. The second phase sets up the secure data transmission channels, which are namely the IPSec SAs.

The first phase includes these tasks:

1. The two parties negotiate the encryption and authentication algorithms to use in the IKE SAs.
2. The two parties authenticate each other using a predetermined mechanism, such as pre-shared keys or digital certificates.
3. A shared master key is generated by the Diffie-Hellman Public Key Algorithm within the IKE framework for the two parties. The master key is also used in the second phase to derive IPSec keys for the SAs.

The second phase includes these tasks:

1. The two parties negotiate the encryption and authentication algorithms to use in the IPSec SAs.
2. The master key is used to derive the IPSec keys for the SAs. Once the SA keys are created and exchanged, the IPSec SAs are ready to protect user data between the two VPN gateways.

Many of the details associated with processing IP traffic in an IPSec implementation are largely a local matter, not subject to standardization. However, some external aspects of the processing must be standardized, to ensure interoperability and to provide a minimum management capability that is essential for productive use of IPSec.

The following describes a general model for processing IP traffic relative to security associations, in support of these interoperability and functionality goals. The model described below is nominal; compliant implementations need not match details of this model as presented, but the external behavior of such implementations must be mappable to the externally observable characteristics of this model.

There are two nominal databases in this model: the Security Policy Database and the Security Association Database. The former specifies the policies that determine the disposition of all IP traffic inbound or outbound from a host, security gateway, or BITS or BITW IPsec implementation. The latter database contains parameters that are associated with each (active) security association. This section also defines the concept of a Selector, a set of IP and upper layer protocol field values that is used by the Security Policy Database to map traffic to a policy, i.e., an SA (or SA bundle).

Each interface for which IPsec is enabled requires nominally separate inbound vs. outbound databases (SAD and SPD), because of the directionality of many of the fields that are used as selectors. Typically there is just one such interface, for a host or security gateway (SG). Note that an SG would always have at least 2 interfaces, but the "internal" one to the corporate net, usually would not have IPsec enabled and so only one pair of SADs and one pair of SPDs would be needed. On the other hand, if a host had multiple interfaces or an SG had multiple external interfaces, it might be necessary to have separate SAD and SPD pairs for each interface.

Ultimately, a security association is a management construct used to enforce a security policy in the IPsec environment. Thus an essential element of SA processing is an underlying Security Policy Database (SPD) that specifies what services are to be offered to IP datagrams and in what fashion. The form of the database and its interface are outside the scope of this specification. However, this section does specify certain minimum management functionality that must be provided, to allow a user or system administrator to control how IPsec is applied to traffic transmitted or received by a host or transiting a security gateway.

The SPD must be consulted during the processing of all traffic (INBOUND and OUTBOUND), including non-IPsec traffic. In order to support this, the SPD requires distinct entries for inbound and outbound traffic. One can think of this as separate SPDs (inbound vs. outbound). In addition, a nominally separate SPD must be provided for each IPsec-enabled interface.

An SPD must discriminate among traffic that is afforded IPsec protection and traffic that is allowed to bypass IPsec. This applies to the IPsec protection to be applied by a sender and to the IPsec protection that must be present at the receiver. For any outbound or inbound datagram, three processing choices are possible: discard, bypass IPsec, or apply IPsec. The first choice refers to traffic that is not allowed to exit the host, traverse the security gateway, or be delivered to an application at all. The second choice refers to traffic that is allowed to pass without additional IPsec protection. The third choice refers to traffic that is afforded IPsec protection, and for such traffic the SPD must specify the security services to be provided, protocols to be employed, algorithms to be used, etc.

For every IPsec implementation, there **MUST** be an administrative interface that allows a user or system administrator to manage the SPD. Specifically, every inbound or outbound packet is subject to processing by IPsec and the SPD must specify what action will be taken in each case. Thus the administrative interface must allow the user (or system administrator) to specify the security processing to be applied to any packet entering or exiting the system, on a packet by packet basis. In a host IPsec implementation making use of a socket interface, the SPD may not need to be consulted on a per packet basis, but the effect is still the same. The management interface for the SPD **MUST** allow creation of entries consistent with the aforementioned selectors, and **MUST** support (total) ordering of these entries. It is expected that through the use of wildcards in various selector fields, and because all packets on a single UDP or TCP connection will tend to match a single SPD entry, this requirement will not impose an unreasonably detailed level of SPD specification. The selectors are analogous to what are found in a stateless firewall or filtering router and which are currently manageable this way.

In host systems, applications may be allowed to select what security processing is to be applied to the traffic they generate and consume. Means of signaling such requests to the IPsec implementation are outside the scope of this standard. However, the system administrator must be able to specify whether or not a user or application can override default system policies. Note that application specified policies may satisfy system requirements, so that the system may not need to do additional IPsec processing beyond that needed to meet an application's requirements. The form of the management interface is not specified by this document and may differ for hosts vs. security gateways, and within hosts the interface may differ for socket-based vs. BITS implementations. However, this document does specify a standard set of SPD elements that all IPsec implementations must support.

The SPD contains an ordered list of policy entries, each policy entry is keyed by one or more selectors that define the set of IP traffic encompassed by this policy entry. These define the granularity of policies or SAs. Each entry includes an indication of whether traffic matching this policy will be bypassed, discarded, or subject to IPsec processing. If IPsec processing is to be applied, the entry includes an SA (or SA bundle) specification, listing the IPsec protocols, modes, and algorithms to be employed, including any nesting requirements. For example, an entry may call for all matching traffic to be protected by ESP in transport mode using 3DES-CBC with an explicit IV, nested inside of AH in tunnel mode using HMAC/SHA-1. For each selector, the policy entry specifies how to derive the corresponding values for a new Security Association Database from those in the SPD and the packet (Note that at present, ranges are only supported for IP addresses; but, wildcarding can be expressed for all selectors):

A- Use the value in the packet itself.

This will limit use of the SA to those packets which have this packet's value for the selector even if the selector for the policy entry has a range of allowed values or a wildcard for this selector.

B- Use the value associated with the policy entry.

If this were to be just a single value, then there would be no difference between (b) and (a). However, if the allowed values for the selector are a range (for IP addresses) or wildcard, then in the case of a range, (b) would enable use of the SA by any packet with a selector value within the range not just by packets with the selector value of the packet that triggered the creation of the SA. In the case of a wildcard, (b) would allow use of the SA by packets with any value for this selector.

For example, suppose there is an SPD entry where the allowed value for source address is any of a range of hosts (192.168.2.1 to 192.168.2.10). And suppose that a packet is to be sent that has a source address of 192.168.2.3. The value to be used for the SA could be any of the sample values below depending on what the policy entry for this selector says is the source of the selector value:

source for the example of
value to be new SAD
used in the SA selector value

- a. packet 192.168.2.3 (one host)
- b. SPD entry 192.168.2.1 to 192.168.2.10 (range of hosts)

Note that if the SPD entry had an allowed value of wildcard for the source address, then the SAD selector value could be wildcard (any host). Case (a) can be used to prohibit sharing, even among packets that match the same SPD entry

Conclusion

We've come a long way from what we need for security, through why it was so hard to get before, to how IPSec makes things so much easier. Large IP networks are remarkable entities. Almost like living things, they tend to grow and evolve with organizations that run them – developing into larger and larger tangles of nodes, until no single person has a way of knowing, nor any hope of ever fixing out, just what the network's dimensions are. This is not, in itself, a bad thing. It is part of the beauty of IP that it is possible in the first place. But it does make keeping track of security difficult. As noted in the introduction, this is why the Internet, and large IP networks in general, have not in the past been particularly safe places for sensitive data and communications.

That is unfortunate, because large IP networks are remarkable resources. The security weaknesses, however understandable from a historical perspective, can be terribly limiting in terms of the scope of those networks' usefulness. It is a tricky thing adding security to this jungle. A challenging thing, probably most challenging of all, you do not want to do anything that might limit the net's remarkable flexibility. Large IP nets are chaotic, hazy, fluid things, but that is part of what makes them valuable and powerful. It is challenging, but not impossible, to add a level of security to this mix without damaging what is already there. And the IPSec suite's designers have answered that challenge.

References

<http://www.openbsd.org/faq/faq13.html>
<http://www.idi.ntnu.no/~runhan/project/report-html>
<http://people.cs.uct.ac.za/~idavies/Security/node45.html>
<http://secinf.net/info/unix/lasg/ipsec>
<http://www.aciksystem.org.tr/as98s/ciliz/tsld029.htm>
http://www.e-businessworld.com/crd_users_66512.html
http://www.prenhall.com/books/ptr_0130118982.html
<http://www.softpro.com/0-13-011898-2.html>

© SANS Institute 2001, Author retains full rights



Upcoming SANS Training

[Click Here for a full list of all Upcoming SANS Events by Location](#)

| | | | |
|---|-------------------------------|------------------------------------|-------------------|
| SANS London 2009 | London, United Kingdom | Nov 28, 2009 - Dec 06, 2009 | Live Event |
| SANS WhatWorks in Incident Detection Summit 2009 | Washington, DC | Dec 09, 2009 - Dec 10, 2009 | Live Event |
| SANS CDI East 2009 | Washington, DC | Dec 11, 2009 - Dec 18, 2009 | Live Event |
| SANS WhatWorks in Data Leakage Prevention and Encryption Summit 2010 | New Orleans, LA | Jan 07, 2010 - Jan 12, 2010 | Live Event |
| SANS Security East 2010 | New Orleans, LA | Jan 10, 2010 - Jan 18, 2010 | Live Event |
| SANS AppSec 2010 and WhatWorks in AppSec Summit | San Francisco, CA | Jan 29, 2010 - Feb 05, 2010 | Live Event |
| SANS Phoenix 2010 | Phoenix, AZ | Feb 14, 2010 - Feb 20, 2010 | Live Event |
| SANS Tokyo 2010 Spring | Tokyo, Japan | Feb 15, 2010 - Feb 20, 2010 | Live Event |
| SANS Geneva CISSP at HEG 2009 Autumn | OnlineSwitzerland | Nov 23, 2009 - Nov 28, 2009 | Live Event |
| SANS OnDemand | Books & MP3s Only | Anytime | Self Paced |