IPSec Basics

This section includes the following topics:

- Introduction to IPsec
- IKE, Internet Key Exchange
- IKE Authentication Methods (Manual, PSK, certificates)
- IPsec Protocols (ESP/AH)
- NAT Traversal

This chapter introduces IPsec, the method, or rather set of methods used to provide VPN functionality.

Introduction to IPsec

IPsec, Internet Protocol Security, is a set of protocols defined by the IETF, Internet Engineering Task Force, to provide IP security at the network layer.

An IPsec based VPN, such as Amaranten VPN, is made up by two parts:

- Internet Key Exchange protocol (IKE)
- IPsec protocols (AH/ESP/both)

The first part, IKE, is the initial negotiation phase, where the two VPN endpoints agree on which methods will be used to provide security for the underlying IP traffic. Furthermore, IKE is used to manage connections, by defining a set of Security Associations, SAs, for each connection. SAs are unidirectional, so there will be at least two SAs per IPsec connection. This is covered in greater detail in section 7.2, IKE, Internet Key Exchange.

The other part is the actual IP data being transferred, using the encryption and authentication methods agreed upon in the IKE negotiation. This can be accomplished in a number of ways; by using IPsec protocols ESP, AH, or a combination of both. These are explained in section 7.3, IPsec Protocols (ESP/AH).

The flow of events can be briefly described as follows:

- IKE negotiates how IKE should be protected
- IKE negotiates how IPsec should be protected
- IPsec moves data in the VPN

The following sections will describe each of these steps in detail.

IKE, Internet Key Exchange

This section describes IKE, the Internet Key Exchange protocol, and the parameters that are used with it.

Encrypting and authenticating data is fairly straightforward, the only things needed are encryption and authentication algorithms, and the keys used with them. The Internet Key Exchange protocol,

IKE, is used as a method of distributing these "session keys", as well as providing a way for the VPN endpoints to agree on how the data should be protected.

IKE has three main tasks:

Provide a means for the endpoints to authenticate each other

Establish new IPsec connections (create SA pairs)

Manage existing connections

IKE keeps track of connections by assigning a bundle of Security Associations, SAs, to each connection. An SA describes all parameters associated with a particular connection, including things like the IPsec protocol used (ESP/AH/both), the session keys used to encrypt/decrypt and/or authenticate/verify the transmitted data. An SA is, by nature, unidirectional, thus the need for more than one SA per connection. In most cases, where only one of ESP or AH is used, two SAs will be created for each connection, one describing the incoming traffic, and the other the outgoing. In cases where ESP and AH are used in conjunction, four SAs will be created.

IKE Negotiation

The process of negotiating session parameters consists of a number of phases and modes. These are described in detail in the below sections.

The flow of events can be briefly described as follows:

IKE Phase-1

- Negotiate how IKE should be protected

IKE Phase-2

- Negotiate how IPsec should be protected
- Derive some fresh keying material from the key exchange in phase-1, to provide session keys to be used in the encryption and authentication of the VPN data flow

Both the IKE and the IPsec connections have limited lifetimes, described both as time (seconds), and data (kilobytes). These lifetimes prevent a connection from being used too long, which is desirable from a cryptanalysis perspective.

The IPsec lifetime is generally shorter than the IKE lifetime. This allows for the IPsec connection to be re-keyed simply by performing another phase-2 negotiation. There is no need to do another phase-1 negotiation until the IKE lifetime has expired.

IKE Proposals

An IKE proposal is a suggestion of how to protect data. The VPN gateway initiating an IPsec connection, the initiator, will send a list of proposals, a proposal-list, suggesting different methods of how to protect the connection.

The connection being negotiated can be either an IPsec connection protecting the data flow through the VPN, or it can be an IKE connection, protecting the IKE negotiation itself.

The responding VPN gateway, upon receiving this proposal-list, will choose the most suitable proposal according to its own security policy, and respond by specifying which one of the proposal it has chosen.

If no acceptable proposal can be found, it will respond by saying that no proposal could be accepted, and possibly provide a reason why.

The proposals contain all parameters needed, such as algorithms used to encrypt and authenticate the data, and other parameters as described in section IKE Parameters.

IKE Phase-1 - IKE Security Negotiation

An IKE negotiation is performed in two phases. The first phase, phase-1, is used to authenticate the two VPN gateways or VPN Clients to each other, by confirming that the remote gateway has a matching Pre-Shared Key.

However since we do not want to publish too much of the negotiation in plaintext, we first agree upon a way of protecting the rest of the IKE negotiation. This is done, as described in the previous section, by the initiator sending a proposal-list to the responder. When this has been done, and the responder accepted one of the proposals, we try to authenticate the other end of the VPN to make sure it is who we think it is, as well as proving to the remote gateway that we are who we are.

Authentication can be accomplished through Pre-Shared Keys, certificates or public key encryption. Pre-Shared Keys is the most common authentication method today. PSK and certificates is supported by the Amaranten Firewall VPN module.

IKE Phase-2 - IPsec Security Negotiation

In phase two, another negotiation is performed, detailing the parameters for the IPsec connection.

In phase-2 we will also extract new keying material from the Diffie-Hellman key exchange in phase-1, to provide session keys to use in protecting the VPN data flow.

If PFS, Perfect Forwarding Secrecy, is used, a new Diffie-Hellman exchange is performed for each phase-2 negotiation. While this is slower, it makes sure that no keys are dependent on any other previously used keys; no keys are extracted from the same initial keying material. This is to make sure that, in the unlikely event that some key was compromised, no subsequent keys can be derived.

Once the phase-2 negotiation is finished, the VPN connection is established and ready for use.

IKE Parameters

There are a number of parameters used in the negotiation process.

Below is a summary of the configuration parameters needed to establish a VPN connection. We highly recommend learning what these parameters do before attempting to configure the VPN endpoints, since it is of great importance that both endpoints are able to agree on all of these parameters.

When installing two Amaranten Firewalls as VPN endpoints, this process is reduced to comparing fields in two identical dialog boxes. However, it is not quite as easy when equipment from different vendors is involved.

Below is a summary of the parameters involved in IKE negotiations, followed by detailed descriptions of all parameters.

Endpoint identification	Local and Remote networks/hosts
Tunnel/transport mode	Remote gateway
Main/aggressive mode	IPsec protocol (ESP/AH/both)
IKE encryption	IKE authentication
IKE DH group	IKE lifetime
PFS on/off/identities	IPsec DH group
IPsec encryption	IPsec authentication
IPsec lifetime	

Endpoint Identification

This is a piece of data representing the identity of the VPN gateway. What this is exactly, depends on the authentication method used. When Pre-Shared Keys are used, this is a piece of data, generally a hex-string or some kind of "pass phrase", identifying this VPN gateway. The remote gateway has to have the same PSK in order for the VPN gateways to authenticate each other.

Authentication using Pre-Shared Keys is based on the Diffie-Hellman algorithm.

Local and Remote Networks/Hosts

These are the subnets or hosts between which IP traffic will be protected by the VPN. In a LAN-to-LAN connection, these will be the network addresses of the respective LANs.

If roaming clients are used, the remote network will most likely be set to 0.0.0.0/0, meaning that the roaming client may connect from anywhere.

Tunnel / Transport mode

IPsec can be used in two modes, tunnel or transport.

Tunnel mode indicates that the traffic will be tunneled to a remote gateway, which will decrypt/authenticate the data, extract it from its tunnel and pass it on to its final destination. This way, an eavesdropper will only see encrypted traffic going from one of VPN endpoint to another. In transport mode, the traffic will not be tunneled, and is hence not applicable to VPN tunnels. It can be used to secure a connection from a VPN client directly to the security gateway, e.g. for IPsec protected remote configuration.

This setting will typically be set to "tunnel" in most configurations.

Remote Gateway

The remote gateway is the remote security gateway which will be doing decryption/authentication and pass the data on to its final location. This field can also be set to "none", forcing the Amaranten VPN to treat the remote address as the remote gateway. This is particularly useful in cases of roaming access, where the IP addresses of the remote VPN clients are not known beforehand. Setting this to "none" will allow anyone coming from an IP address conforming to the

"remote network" address discussed above to open a VPN connection, provided they can authenticate properly.

The remote gateway is not used in transport mode.

Main/Aggressive Mode

The IKE negotiation has two modes of operation, main mode and aggressive mode.

The difference between these two is that aggressive mode will pass more information in fewer packets, with the benefit of slightly faster connection establishment, at the cost of transmitting the identities of the security gateways in the clear.

When using aggressive mode, some configuration parameters, such as Diffie-Hellman groups, and PFS, can not be negotiated, resulting in a greater importance of having "compatible" configurations on both ends.

IPsec Protocols

The IPsec protocols describe how the data will be processed. The two protocols to choose from are AH, Authentication Header, and ESP, Encapsulating Security Payload.

ESP provides encryption, authentication, or both. However, we do not recommend using encryption only, since it will dramatically decrease security.

More on ESP inESP (Encapsulating Security Payload).

AH only provides authentication. The difference from ESP with authentication only is that AH also authenticates parts of the outer IP header, for instance source and destination addresses, making certain that the packet really came from who the IP header claims it is from.

More on AH inAH (Authentication Header).

Note: Amaranten Firewall does not support AH.

IKE Encryption

This specifies the encryption algorithm used in the IKE negotiation, and depending on the algorithm, the size of the encryption key used.

The algorithms supported by Amaranten VPN are:

- AES
- Blowfish
- Twofish
- Cast128
- 3DES
- DES

DES is only included to be interoperable with other older VPN implementations. Use of DES should be avoided whenever possible, since it is an old algorithm that is no longer considered secure.

IKE Authentication

This specifies the authentication algorithm used in the IKE negotiation.

The algorithms supported by Amaranten VPN are:

- SHA1
- MD5

IKE DH (Diffie-Hellman) Group

This specifies the Diffie-Hellman group to use when doing key exchanges in IKE.

The Diffie-Hellman groups supported by Amaranten VPN are:

- DH group 1 (768-bit)
- DH group 2 (1024-bit)
- DH group 5 (1536-bit)

The security of the key exchanges increase as the DH groups grow larger, as does the time of the exchanges.

IKE Lifetime

This is the lifetime of the IKE connection.

It is specified in time (seconds) as well as data amount (kilobytes). Whenever one of these expires, a new phase-1 exchange will be performed. If no data was transmitted in the last "incarnation" of the IKE connection, no new connection will be made until someone wants to use the VPN connection again.

PFS

With PFS disabled, initial keying material is "created" during the key exchange in phase-1 of the IKE negotiation. In phase-2 of the IKE negotiation, encryption and authentication session keys will be extracted from this initial keying material. By using PFS, Perfect Forwarding Secrecy, completely new keying material will always be created upon re-key. Should one key be compromised, no other key can be derived using that information.

PFS can be used in two modes, the first is PFS on keys, where a new key exchange will be performed in every phase-2 negotiation. The other type is PFS on identities, where the identities are also protected, by deleting the phase-1 SA every time a phase-2 negotiation has been finished, making sure no more than one phase-2 negotiation is encrypted using the same key.

PFS is generally not needed, since it is very unlikely that any encryption or authentication keys will be compromised.

IPsec DH Group

This is a Diffie-Hellman group much like the one for IKE. However, this one is used solely for PFS.

IPsec Encryption

The encryption algorithm to use on the protected traffic.

This is not needed when AH is used, or when ESP is used without encryption.

The algorithms supported by Amaranten VPN are:

- AES
- Blowfish
- Twofish
- Cast128
- 3DES
- DES

IPsec Authentication

This specifies the authentication algorithm used on the protected traffic.

This is not used when ESP is used without authentication, although it is not recommended to use ESP without authentication.

The algorithms supported by Amaranten VPN are:

- SHA1
- MD5

IPsec Lifetime

This is the lifetime of the VPN connection. It is specified in both time (seconds) and data amount (kilobytes). Whenever either of these values is exceeded, a re-key will be initiated, providing new IPsec encryption and authentication session keys. If the VPN connection has not been used during the last re-key period, the connection will be terminated, and re-opened from scratch when the connection is needed again.

IKE Authentication Methods (Manual, PSK, certificates)

Manual Keying

The "simplest" way of configuring a VPN is by using a method called "manual keying". This is a method where IKE is not used at all; the encryption and authentication keys as well as some other parameters are directly configured on both sides of the VPN tunnel.

Note: Amaranten Firewall does not support Manual Keying.

Advantages

Since it is very straightforward it will be quite interoperable. Most interoperability problems encountered today are in IKE. Manual keying completely bypasses IKE and sets up its own set of IPsec SAs.

Disadvantages

It is an old method, which was used before IKE came to be, and is thus lacking all the functionality of IKE. This method therefore has a number of limitations, such as having to use the same encryption/authentication key always, no anti-replay services, and it is not very flexible. There is also no way of assuring that the remote host/gateway really is the one it says it is.

This type of connection is also vulnerable for something called "replay attacks", meaning a malicious entity which has access to the encrypted traffic can record some packets, store them, and send them to its destination at a later time. The destination VPN endpoint will have no way of telling if this packet is a "replayed" packet or not. Using IKE eliminates this vulnerability.

Pre-Shared Keying, PSK

Pre-shared keying is a method where the endpoints of the VPN "share" a secret key. This is a service provided by IKE, and thus has all the advantages that come with it, making it far more flexible than manual keying.

Advantages

Pre-Shared Keying has a lot of advantages over manual keying. These include endpoint authentication, which is what the PSKs are really for. It also includes all the benefits of using IKE. Instead of using a fixed set of encryption keys, session keys will be used for a limited period of time, where after a new set of session keys are used.

Disadvantages

One thing that has to be considered when using Pre-Shared Keying is key distribution. How are the Pre-Shared keys distributed to remote VPN clients and gateways? This is a major issue, since the security of a PSK system is based on the PSKs being secret. Should one PSK be compromised in some way, the configuration will need to be changed to use a new PSK.

Certificates

Each VPN gateway has its own certificate, and one or more trusted root certificates.

The authentication is based on several things:

- That each endpoint has the private key corresponding to the public key found in its certificate, and that nobody else has access to the private key.
- That the certificate has been signed by someone that the remote gateway trusts.

Advantages

Added flexibility. Many VPN clients, for instance, can be managed without having the same pre-shared key configured on all of them, which is often the case when using pre-shared keys and roaming clients. Instead, should a client be compromised, the client? certificate can simply be revoked. No need to reconfigure every client.

Disadvantages

Added complexity. Certificate-based authentication may be used as part of a larger public key infrastructure, making all VPN clients and gateways dependent on third parties. In other words, there are more things that have to be configured, and there are more things that can go wrong.

IPsec Protocols (ESP/AH)

The IPsec protocols are the protocols used to protect the actual traffic being passed through the VPN. The actual protocols used, and the keys used with them are negotiated by IKE.

There are two protocols associated with IPsec, AH and ESP. These are covered in the sections below.

AH (Authentication Header)

AH is a protocol used for authenticating a data stream. It uses a cryptographic hash function to produce a MAC from the data in the IP packet. This MAC is then transmitted with the packet, allowing the remote gateway to verify the integrity of the original IP packet, making sure the data has not been tampered with on its way through the Internet.

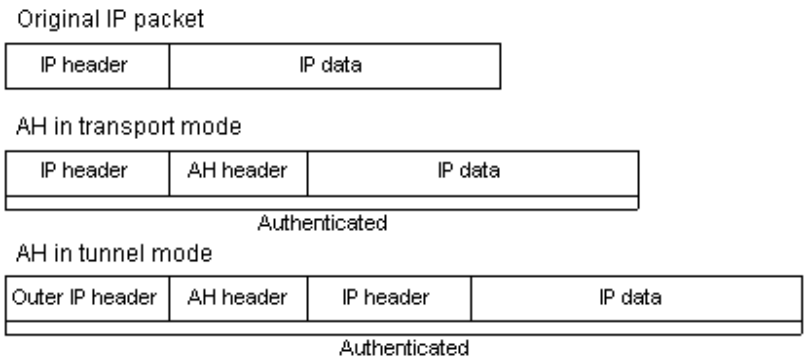

Apart from the IP packet data, AH also authenticates parts of the IP header.

The AH protocol inserts an AH header after the original IP header, and in tunnel mode, the AH header is inserted after the outer header, but before the original, inner, IP header.

Note: Amaranten Firewall does not support AH.

ESP (Encapsulating Security Payload)

The ESP protocol is used for both encryption and authentication of the IP packet. It can also be used to do either encryption only, or authentication only.

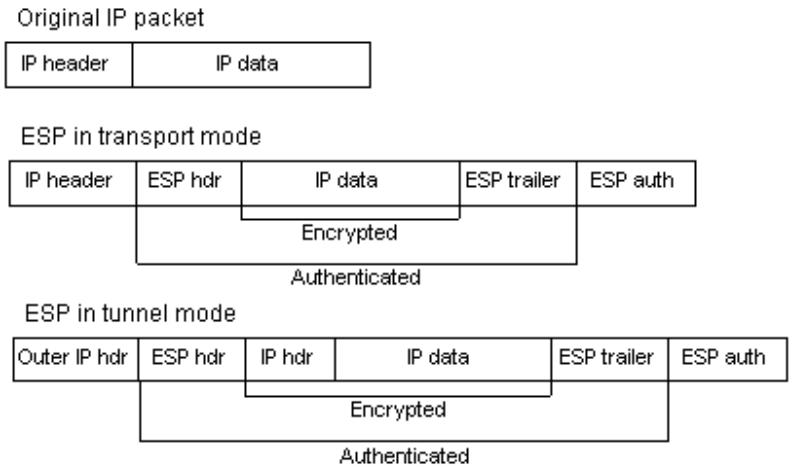

The ESP protocol inserts an ESP header after the original IP header, in tunnel mode, the ESP header is inserted after the outer header, but before the original, inner, IP header.

All data after the ESP header is encrypted and/or authenticated. The difference from AH is that ESP also provides encryption of the IP packet. The authentication phase also differs in that ESP only authenticates the data after the ESP header; thus the outer IP header is left unprotected.

NAT Traversal

There is one big problem with the IKE and IPsec protocols. That problem is NAT. The IKE and IPsec protocols were not designed to work through NATs. Because of this, something called "NAT traversal" has evolved. NAT traversal is an add-on to the IKE and IPsec protocols that makes them work when being NATed.

In short, NAT traversal is divided into two parts:

- Additions to IKE that lets IPsec peers tell each other that they support NAT traversal, and the specific versions of the draft they support.
- Changes to the ESP encapsulation. If NAT traversal is used, ESP is encapsulated in UDP, which allows for more flexible NATing.

Below is a more detailed description of the changes made to the IKE and IPsec protocols.

NAT traversal is only used if both ends has support for it. For this purpose, NAT traversal aware VPNs send out a special "vendor ID", telling the other end that it understand NAT traversal, and which specific versions of the draft it supports.

NAT detection: Both IPsec peers send hashes of their own IP addresses along with the source UDP port used in the IKE negotiations. This information is used to see whether the IP address and source port each peer uses is the same as what the other peer sees. If the source address and port have not changed, then the traffic has not been NATed along the way, and NAT traversal is not necessary. If the source address and/or port has changed, then the traffic has been NATed, and NAT traversal is used.

Once the IPsec peers have decided that NAT traversal is necessary, the IKE negotiation is moved away from UDP port 500 to port 4500. This is necessary since certain NAT devices treat UDP packet to port 500 differently from other UDP packets in an effort to work around the NAT

problems with IKE. The problem is that this special handling of IKE packets may in fact break the IKE negotiations, which is why the UDP port used by IKE has changed.

Another problem NAT traversal resolves is that the ESP protocol is an IP protocol. There is no port information like in TCP and UDP, which makes it impossible to have more than one NATed client connected to the same remote gateway and the same time. Because of this, ESP packets are encapsulated in UDP. The ESP-UDP traffic is sent on port 4500, the same port as IKE when NAT traversal is used. Once the port has been changed all following IKE communications are done over port 4500. Keepalive packets are also being sent periodically to keep the NAT mapping alive.

NAT traversal drafts supported by Amaranten firewall:

- draft-ietf-ipsec-nat-t-ike-00
- draft-ietf-ipsec-nat-t-ike-01
- draft-ietf-ipsec-nat-t-ike-02
- draft-ietf-ipsec-nat-t-ike-03

NAT traversal configuration

Most of the NAT traversal functionality is completely automatic. It is enabled when necessary.

There are a few things to keep in mind though:

Initiating firewall:

- No special configuration needed.

Responding firewall:

- On responding firewalls, the *remote gateway* field is used as a filter on the source IP of recevied IKE packets. This should be set to allow the NATed IP address of the initiator.
- Individual pre-shared keys can not be used where multiple clients connecting to one remote gateway gets NATed out through the same address. Having the same pre-shared key on all clients will work. However, this is not recommended. The preferred way is to used certificates instead.