Liberté • Égalité • Fraternité
RÉPUBLIQUE FRANÇAISE

PREMIER MINISTRE

Secrétariat général
de la défense
et de la sécurité nationale

*Agence nationale de la sécurité
des systèmes d'information*

Paris, August 3, 2015

N° DAT-NT-003-EN/ANSSI/SDE/NP

Number of pages
(including this page): 17

Technical report

Recommendations for securing networks with IPsec

Targeted audience

Developers	
Administrators	✓
IT security managers	✓
IT managers	✓
Users	

Document Information

Disclaimer

This document, written by the ANSSI, presents the **"Recommendations for securing networks with IPsec"**. It is freely available at www.ssi.gouv.fr/ipsec. It is an original creation from the ANSSI and it is placed under the "Open Licence" published by the Etalab mission (www.etalab.gouv.fr). Consequently, its diffusion is unlimited and unrestricted.

This document is a courtesy translation of the initial French document **"Recommandations de sécurité relatives à IPsec pour la protection des flux réseau"**, available at www.ssi.gouv.fr/ipsec. In case of conflicts between these two documents, the latter is considered as the only reference.

These recommendations are provided as is and are related to threats known at the publication time. Considering the information systems diversity, the ANSSI cannot guarantee direct application of these recommendations on targeted information systems. Applying the following recommendations shall be, at first, validated by IT administrators and/or IT security managers.

Document contributors

Contributors	Written by	Approved by	Date
Cisco[1], DAT	DAT	SDE	August 3, 2015

Document changelog

Version	Date	Changelog
1.1	August 3, 2015	Translation from the original French document, version 1.1

Contact information

Contact	Address	Email	Phone
Bureau Communication de l'ANSSI	51 boulevard de La Tour-Maubourg 75700 Paris Cedex 07 SP FRANCE	communication@ssi.gouv.fr	01 71 75 84 04

1. This document is based on a courtesy translation provided by Cisco Systems Inc.

Contents

1 Introduction

Nowadays, IT systems generally adopt a distributed architecture. The different hardware and software building blocks are increasingly interconnected, not only among themselves, but also with remote systems and through the Internet. The rise of cloud computing and outsourcing accelerates this trend.

As these building blocks can be critical to an IT system, the data they exchange can also be highly critical. These streams include a lot of sensitive data (authentication data, confidential business information, industrial control systems commands...). The interception or alteration of such sensitive data by potentially malicious individuals are significant risks in a context where cyber attacks are increasingly numerous and sophisticated. The protection of these sensitive data streams is, therefore, paramount.

However, this issue is not always correctly addressed, and many sensitive network data streams are not as protected as they should be. IPsec is a set of secure communication protocols aiming at protecting network data streams. This technology is field proven, but often poorly mastered, and is still hardly used or poorly employed.

2 Overview of IPsec

IPsec encapsulates IP packets to provide data confidentiality, integrity and anti-replay protection at the network layer ("Internet" layer of the TCP/IP stack, "network" layer or layer 3 of the OSI model). IPsec is an IETF standard, defined through RFC 4301 to 4309. Several versions have followed and various additional elements have been defined. An inventory of all the documents defining IPsec and IKE related RFCs is given in RFC 6071.

A very large number of network devices, including routers and firewalls, implement IPsec. Similarly, many computer or smartphone operating systems natively support IPsec, usually allowing IPsec connections between those operating systems and/or equipments.

In many cases, the use of IPsec presents a favourable "costs to (security) benefits" ratio, as this technology is natively supported by most client systems and network equipments, therefore the investment is often limited. It is also a mature, well-understood and highly interoperable protocol. Its implementation can thus be done without excessive burden on network administration teams among different equipments.

3 Glossary

AH *Authentication Header* : part of the IPsec protocol, see 6.1.1.

ESP *Encapsulation Security Payload* : part of the IPsec protocol, see 6.1.2.

IKE *Internet Key Exchange* : key exchange protocol, see 6.4.3.

VPN *Virtual Private Network*.

IETF *Internet Engineering Task Force* : the Internet standards organization.

RFC *Request for comments* : documents from the IETF, such as standardization of protocols.

NAT *Network Address Translation*.

TLS *Transport Layer Security* : security protocol for the application layer.

SSL *Secure Socket Layer* : obsolete version of TLS.

SC *Specialized Connection*.

MPLS *MultiProtocol Label Switching* : network protocol based on label switching, frequently used to provide "IP VPN" services.

RGS *Référentiel général de sécurité* : document available at `www.ssi.gouv.fr/rgs`.

MTU *Maximum Transmission Unit*: maximum packet size that can be sent or received on a network interface.

4 IPsec use cases

IPsec technology is primarily associated with VPN connections, which often rely on a public network such as the Internet. It is nevertheless important to note that this use of IPsec is far from being the only one.

4.1 Remote access to the IT system

Data streams between a mobile unit and its IT system must be protected. IPsec is often used to connect a remote workstation to a private network and is actually well suited for this use. The IPsec tunnel is usually installed between a client computer (a laptop or smartphone, using a VPN software client) and a network security device (a firewall or a VPN box).

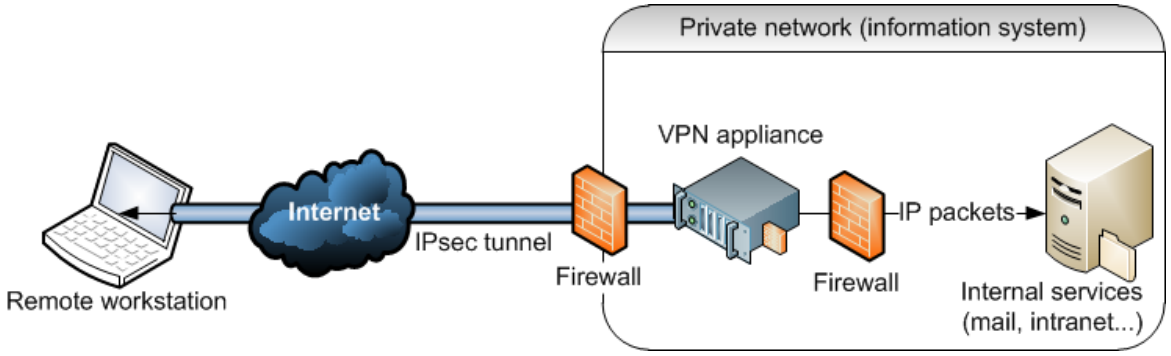

Figure 1: IPsec use case #1: secure remote access

The security provided by the IPsec tunnel is certainly more obvious for applications that do not rely on secure protocols (such as TLS), but the use of IPsec is highly recommended even for applications implementing security in the upper network layers. This is part of a defence-in-depth approach and it also allows to define and maintain a simple remote access policy to the IT system.

IPsec usually confers integrity and confidentiality protection of the data streams after an authentication phase. It complements protocols such as PPTP or L2TP. Its use as a protection layer for remote access to an IT system is thus highly recommended.

R1 - Trusted IPsec VPN products

Several software or hardware implementations of IPsec VPN were evaluated by the ANSSI and obtained a *security certification* or a *formal security approval*. Whenever a trusted IPsec VPN product is needed, priority shall be given to *formally approved security products*[2].

4.2 Connection of two remote sites

IPsec can also be set up to securely connect two remote sites in order to provide a secure connection between two local networks. Malicious activity consisting of accessing the link between these two networks, thus intercepting sensitive information or carrying out man-in-the-middle attacks is efficiently thwarted.

Using IPsec in this scenario, especially when the connection relies on a public network such as the Internet or on a leased link such as SC or MPLS VPN is highly encouraged. In this context, the

2. An up-to-date list of such products is available at www.ssi.gouv.fr/entreprise/qualifications/produits-qualifies-par-lanssi/les-produits/.

IPsec connection is usually configured between two specific IPsec devices or two edge firewalls of the IT systems. Note that recommendation R1 also applies to this use case.

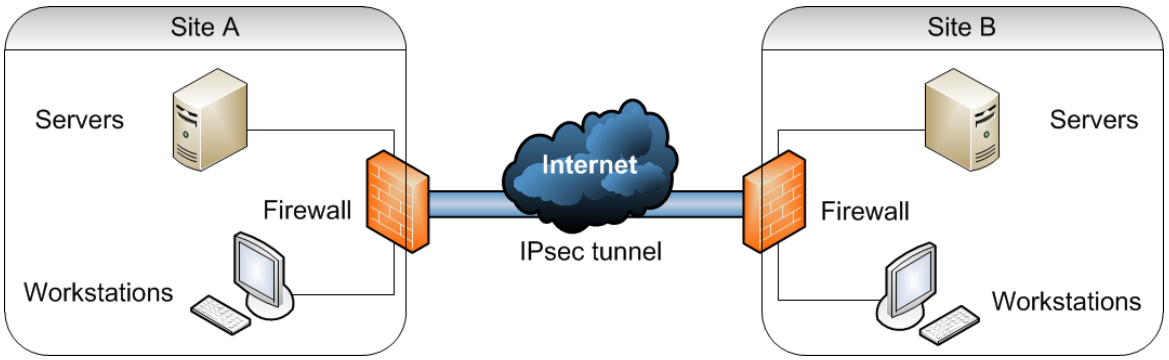

Figure 2: IPsec use case #2 : secure connection between two remote sites.

Note: in this example, the IPsec tunnel is created between two edge firewalls.

4.3 Protection against protocol weaknesses or software vulnerabilities

It may be necessary to cope with hardware or software components whose confidentiality, authentication or integrity mechanisms are either far from the state of the art or even non-existent and whose network communications are hence poorly protected. Malicious individuals having access to these data streams (because they transit through public networks or inadequately secured connections) could then intercept them or carry out man-in-the-middle attacks. Such scenarios include, but are not limited to the following communications

- between Industrial Control Systems (ICS) using poorly secured protocols;
- between application servers and database management systems;
- between distributed application components using proprietary protocols or poorly secured (or even unsecured) software buses;
- between clients and servers using insecure protocols (FTP, POP3, SMTP, HTTP, RDP, VNC. . .).

In many cases, network data streams must be protected by third-party solutions. In these cases, IPsec can be the ideal technology to address some weaknesses of higher level protocols. Its use is thus recommended for encapsulating network streams carrying information deemed sensitive, to provide the required protection. When terminal nodes do not support IPsec, additional network devices may be necessary, as presented in 3 and 4. Note that recommendation R1 is still applicable in this scenario and the residual risk lying in the unsecured end lines shall not be ignored.

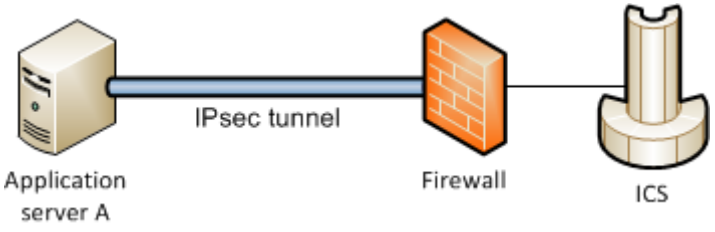

Figure 3: IPsec use case #3 : access to an ICS with an additional equipment.

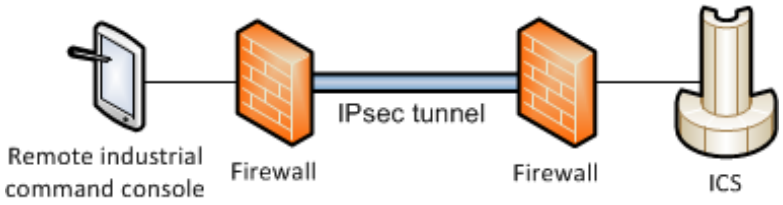

Figure 4: IPsec use case #4 : access to an ICS system with two additional equipments.

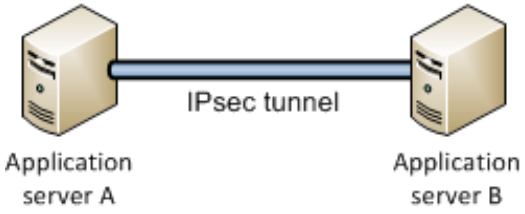

Figure 5: IPsec use case #5 : point to point IPsec access.

4.4 Defence-in-depth

As part of in-depth defence, IPsec is an efficient security measure for encapsulating protocols already secured by other mechanisms (or even an existing IPsec tunnel). As previously indicated, IPsec is inexpensive to implement in most cases and it can efficiently enhance the security level of an IT system.

5 Comparison with TLS

It is not uncommon to see IPsec compared to TLS [3]. These two technologies ensure confidentiality, integrity and authentication of data streams. However, there are several important differences in IPsec favour.

TLS is a higher level protocol than IPsec in the network stack, above the transport layer carried by TCP. TLS is often used to secure higher level protocols : HTTPS is, for example, HTTP secured by TLS. However this protocol competes with IPsec on a specific use case, namely secure remote access, and is known as "SSL VPN". An "SSL VPN" encapsulates a network stream into a TLS session and some "SSL VPN" solutions even rely on a browser to avoid installing a specific "SSL VPN" client on remote equipments.

The primary drawback is that the TLS security-related operations are performed in userland, using a client process. These operations (and the secrets they handle) are thus much more vulnerable to attacks than with IPsec where critical operations take place within the kernel space or in dedicated processes. This assertion is especially true when the VPN client relies on the browser, a software with a considerable attack surface, even from a remote point of view.

Moreover, several cryptographic factors tend to favour IPsec. On the one hand, IPsec allows a wider use of modern algorithms, recommended in many best practices, considering both the coverage by the standard and the availability of software implementations. Furthermore, the use of cryptographic primitives in IPsec is more compliant with the best practices: for example, IPsec is based on the "Encrypt-then-MAC" method, considered safer than the "MAC-then-Encrypt" method employed by TLS.

Finally, "SSL VPN" is a tweak from TLS initial intended use [4] and it is not ideal. The encapsulation of packets from the network layer in the application layer leads to an "external" TCP header with no correlation to any "internal" TCP header leading to suboptimal functioning of the network congestion control mechanisms.

R2

In the use cases outlined above, IPsec shall be preferred to TLS

Note: TLS is perfectly suited for securing a particular application protocol such as (but not limited to) HTTPS (HTTP over TLS) or IMAPS (IMAP over TLS). This use is complementary and fully compatible with the use of IPsec.

3. The word SSL refers strictly speaking to older and now obsolete versions of the TLS protocol. Nevertheless, many documents refer to SSL to describe TLS to the extent that it eclipses the use of this term.

4. TLS is in principle not intended to secure connections between remote sites but rather secure communications between a user and a service. Although the distinction is sometimes weak, it is really a different approach.

6 How IPsec works

IPsec, by its complexity, is often partially understood or not well mastered. The configuration choices, including the default ones, are not always appropriate, leading to a level of security lower than expected.

6.1 IPsec security services

Security services provided by IPsec are based on two different protocols that are the core of the IPsec technology:

- AH: "Authentication Header" (protocol #51), standardized by RFC 4302 in its latest version;
- ESP: "Encapsulating Security Payload" (protocol #50), standardized by RFC 4303 in its latest version.

Both protocols can be used independently or, more rarely, in combination.

6.1.1 AH : integrity and authentication of IP packets

The AH protocol, less frequently used than ESP, ensures the integrity of IP packets, and, combined with IKE (see section 6.4.3), the authentication of the peers. In other words, AH firstly ensures that exchanged packets have not been altered during transit, and secondly guarantees the identity of a packet's sender. It also provides protection against replayed packets.

It's worth noting that AH does not protect the confidentiality of the data exchanged. The data is not encrypted and sent as clear text, or more exactly in the same format as if the link were not protected with IPsec (encryption can be implemented in higher protocol layers, e.g., using TLS).

The integrity check is performed on the whole IP packets including the headers except variable header fields such as DSCP, ECN, TTL fields, "Flags", fragmentation offset, and checksum. In particular, the source and destination addresses are part of the protected data. A modified packet is thus considered corrupted, leading to an incompatibility between AH and the address translation mechanisms.

RFCs defining IPsec consider optional the support of AH by equipment implementing IPsec, while the support of ESP is mandatory. Usually, the AH protocol is considered less secure than ESP and obsolete; its implementation is hardly necessary.

R3

IPsec shall be used with the ESP protocol. Although AH has no structural vulnerabilities, its use is not recommended.

6.1.2 ESP : confidentiality, integrity, and authentication of IP packets

The ESP protocol provides confidentiality and integrity of IP packets and, used with IKE (see section 6.4.3), authentication of the peers. It also provides protection against replay attacks. One can also set up ESP to provide integrity and authentication without encryption, which suits almost all the use cases where AH was used, justifying the withdrawal of the latter.

Some implementations even allow the use of confidentiality without an integrity checking mechanism. This use is also obsolete and shall be avoided. The suppression of the integrity service has no benefit (the integrity computation/check performance cost is hardly significant compared to the encryption/decryption performance cost) and exposes the user to numerous known and realistic

attacks.

R4

The ESP privacy service shall never be used without activating the integrity control mechanism.

As opposed to the AH protocol, only the "payload" is protected with the ESP protocol (*i.e.* the contents of the IP packet and not its headers). There is, therefore, no fundamental incompatibility with the address translation mechanisms. However, it is necessary to take a number of measures to ensure interoperability between ESP and address translation mechanisms.

6.2 Transport and tunnel modes of operation

Regardless of the choice between AH and ESP, IPsec has two modes of operation called "tunnel mode" and "transport mode". Tunnel mode produces the service expected in most cases.

In transport mode, the data associated with AH or ESP are directly inserted into the original IP packet (*i.e.* the one that would have been sent in the absence of IPsec). The resulting IP packet contains an AH or ESP packet itself containing the initial payload of the packet (e.g., a TCP segment). It may be noted that the initial IP header must be amended: its protocol field should be set to 50 (resp. 51) for ESP (resp. AH) instead of 6 (TCP) or 17 (UDP) for example. The encapsulated protocol, previously present in the original IP header, is now indicated in the added ESP or AH header.

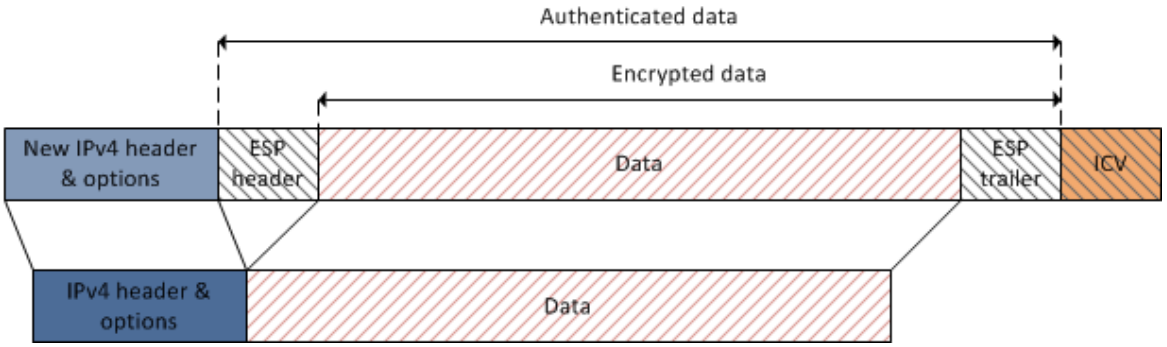

Figure 6: Using ESP in transport mode. ICV stands for "Integrity Check Value", the value used by the integrity checking mechanism

Figure 7: Using AH in transport mode

In tunnel mode, a new IP packet is generated to include the AH or ESP packet containing itself the unchanged original IP packet. In this mode, there are always two IP headers. The outer header is used for routing the packet and the inner header, which can be encrypted when using ESP with

the privacy service, is only processed by the recipient (of the outer packet). The inner header will be ignored by the network equipment located between the transmitter and receiver. This provides a "tunnel" through the network, in the same way as protocols such as IPIP (RFC 2003) or GRE (RFC 2784).

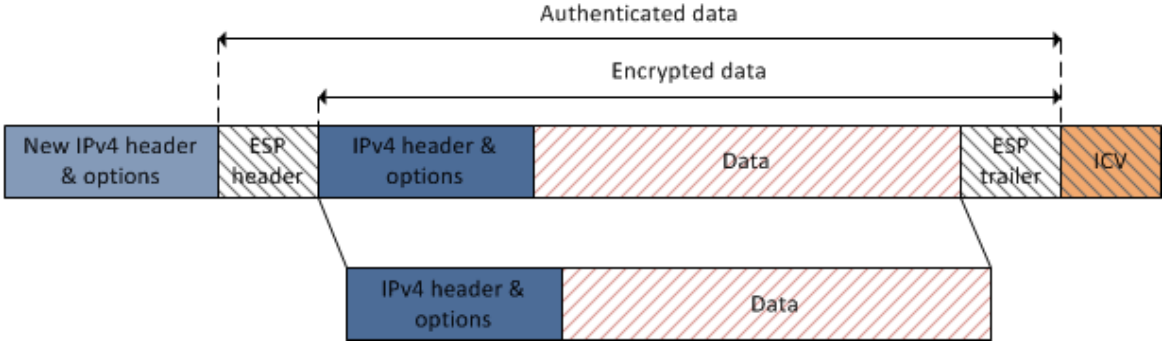

Figure 8: Using ESP in tunnel mode. ICV stands for "Integrity Check Value", the value used by the integrity check mechanism

Figure 9: Using AH in tunnel mode. The ICV is included in the "AH" header

Tunnel mode is well suited for remote access to a private network through a public network. It allows hiding the internal addressing plan (often non routable) from the public network. IPsec is used on the public network between the client and a gateway that extracts the inner IP packet and injects it into the private network (and vice versa for the opposite direction of communication).

Naturally, the use of tunnel mode results in larger packets than transport mode for the same payload because of the IP headers duplication. Network resource consumption is, therefore, increased. In particular, care should be taken that the actual MTU of the tunnel is the MTU minus the size of the meta-information added by IPsec link (new IP header, ESP header and trailer, checksum). The size of these meta-data varies depending on the particular cryptographic parameters, but it is common that the "overhead" due to the tunnel is 50 to 100 bytes. In some cases, mechanisms for automatic MTU configuration mismanage the situation. It is, therefore, relatively common, when using IPsec, to manually limit the maximum size of the clear IP packets to ensure that, once encapsulated, they do not reach a size requiring fragmentation.

One should note that, in tunnel mode, the integrity check offered by the AH protocol not only bears on the inner packet but also on the external IP header. This behaviour is of course incompatible with NAT mechanisms where the outside header is modified (see section 6.5). On the contrary, the ESP integrity check only covers the inner packet and thus allows modifications of the external header.

6.3 Security Policy (SP)

In IPsec, a "Security Policy" (SP) defines an unidirectional link [5] and the following parameters:
- mandatory, optional or non-use of IPsec;
- use of tunnel or transport mode;
- use of AH or ESP.

All of the SPs are collected into a "Security Policy Database" (SPD). Much like a firewall, the SPs are intended to specify allowed and banned network streams.

R5

Security Policys allowing "voluntary" or "optional" use of IPsec shall be avoided because security cannot be ensured ("downgrade attack" may apply). When security is required for a given link, "mandatory" use shall be defined

R6

Although no technical limitations exist, asymmetric SPs for a given bidirectional link should not be used in order to provide a uniform security policy between the outbound link and the inbound link.

6.4 Establishing an IPsec connection

Cryptographic mechanisms used for integrity protection or confidentiality rely on one or more keys. These elements must be shared by different hosts using IPsec. Two different approaches can be set up: manually setting keys on each peer or using the "Internet Key Exchange" (IKE) mechanism to negotiate these keys on demand.

6.4.1 Security Association (SA)

Every unidirectional link is associated with a context called "Security Association" (SA) where the following characteristics reside:
- the source and destination peers;
- the IPsec mode (transport or tunnel) and the protocols (AH or ESP) in use;
- the cryptographic algorithms used;
- the keys associated with these algorithms.

Each SA is associated with a validity period and an unique integer called the "Security Parameter Index" (SPI), identifying the SA within the "Security Association Database" (SAD). The AH and ESP headers consistently indicate the SPI associated with the SA used. The first elements (end peers, mode and protocol) are controlled by the SP in place: a system shall never have SAs violating its SP. Cryptographic parameters (algorithms, key size) can be manually set or negotiated by the IKE protocol (see section 6.4.3). The options are configured by the administrator.

R7

It is advisable to manually determine the algorithms and key sizes used (see below for choices) and only use IKE for key exchange. If these items cannot be configured, algorithm negotiation should only be allowed on a small subset of algorithms.

5. In fact, one distinguishes the A → B link and the B → A link, which may have two different policies.

Note: in this case, the security of the link is determined by the weakest algorithm among those accepted; it is, therefore, mandatory that all options are consistent with the security policy of the organization.

6.4.2 Manual key setting

Algorithms and keys can be manually set on each device. This method is often referred as "Manual keying" and it must not be confused with the "Pre-Shared Key" IKE method of authentication discussed in the next chapter.

Manual keying is strongly discouraged. It requires a meticulous set up, the key being ideally different for each host pair. Moreover, key renewal is in practice hard enough to be consistent both with good cryptographic practices (key wear) and IT Security (crypto period) [6]. In addition, it does not benefit from "Perfect Forward Secrecy" property presented in section 6.6. Finally, it does not allow the implementation of cryptographic authentication mechanisms.

Ultimately, manual key setting should be reserved for test procedures, diagnoses, or very specific systems that have been subject to a detailed security analysis, where the key life cycle is especially well managed and where different keys are provided for different streams.

6.4.3 Internet Key Exchange (IKE)

The IKE protocol, currently in version 2 (), allows two IPsec peers to dynamically negotiate the cryptographic algorithms and keys of a SA.

R8

Use of IKE version 2 is strongly recommended.

6.4.3.1 A two-phase protocol

IKE is divided into two distinct phases. In a first phase, a secure channel (encrypted and authenticated) is created between the two peers. In a second phase, this channel is used to negotiate the various parameters of the inbound and the outbound SAs.

The first phase also uses cryptographic algorithms, which are not necessarily the same as those defined in the SA. The secure channel parameters negotiated during the initial phase and used to protect the second phase are sometimes referred to as ISAKMP SA [7], or even IKE SA, as opposed to IPsec SA which are the SAs negotiated during the second phase and used to protect "useful" traffic.

The first IKE version allowed two different modes for phase 1, called "main mode" and "aggressive mode". The latest requires fewer messages than the first, but it does not hide the identity of the peers to a potential attacker eavesdropping on the network. Phase 2 used the so-called "quick mode". IKE version 2 first phase is similar to IKEv1 main mode, but no longer uses this terminology. However, one can still find products that call the first phase "main mode" and the second phase, "quick mode".

6.4.3.2 Peers authentication

During the first phase, peers authentication can either be done with a shared secret (PSK: "Pre-Shared Key") or with asymmetric cryptography mechanisms, such as RSA. In this latter case, the use

6. There is indeed a distinction between the "key wear" describing a mathematical property that weakens the cryptographic system beyond a certain amount of encrypted information and the "crypto period", denoting a key maximum period of use based on organizational considerations aiming at limiting the impact of a data compromise.

7. ISAKMP is a "protocol framework" used to define the IKE protocol.

of a Public Key Infrastructure (PKI) to certify the public keys of all the participants avoids to deploy them on all hosts.

One should favour the use of a PKI, which drastically simplifies the system administration: adding a new host or revoking a compromised key is easier (no additional intervention on the deployed equipments). In other modes, it can be tricky to react with due diligence to a key compromise, for example, in case of equipment theft.

Usually, PSK mode should be avoided in production systems and confined to testing environments or diagnostic operations. If it were necessary as an exceptional measure, a good general practice for shared secrets is to take care that the entropy is sufficient enough to make a brute force attack difficult. More information on the subject can be found in the RGS published by ANSSI. Less than 100 of bits entropy is currently considered as a risky choice.

R9

In general, it is strongly discouraged to use manual keying or pre-shared key (PSK) to establish an IPsec link. Mechanisms based on asymmetric cryptography shall be preferred. In particular, PKI based solutions, which allows for rapid revocation of compromised keys, especially in the case where a unit is lost, shall be preferred. Exceptions to this recommendation may only be made after a rigorous security study.

6.4.3.3 SP negotiation

IKE also allows SP negotiation. In most cases, all SP parameters are known in advance and this negotiation is of little interest. However, this mechanism makes sense when mobile devices are involved. In this case, the IP address of the mobile client is not known beforehand: it is assigned by the network hosting the mobile workstation. Adapting this parameter from the SP on the fly is thus very useful.

In cases where such a negotiation mechanism is not necessary, a static configuration is preferable to maintain control of the security policy (if allowed by the equipments). If the SPs are negotiated with IKE, one should ensure through a filtering policy that the system is never in a state where it exchanges data in clear text. In practice, network filtering rules should be set to block all traffic not using IKE, ESP, and (if applicable) protocols such as ICMP, required for network operations.

R10

Static configuration of the SP will be favoured whenever possible. Otherwise, one shall ensure the absence of communications in clear text through a filtering policy.

6.5 Using IPsec behind NAT

Using an address translation system (NAT) in conjunction with IPsec can lead to several problems. Firstly, AH cannot be used as the integrity check covers the IP headers and it would become invalid when the source or the destination address is changed.

In addition, certain very common NAT mechanisms require the ability to change the TCP or UDP ports. If the protocol transported by IP is ESP, which has no port, such a mechanism will not work. One shall, in this case, use the “NAT-Traversal” (NAT-T) method, which consists in encapsulating the IKE and ESP traffic into UDP datagrams using the default port 4500.

Finally, some non-standard extensions may be incompatible with the use of NAT or require the use of NAT-T even in cases where the change of ports is not necessary.

R11

If NAT is required for a given IPsec link, one shall enable the NAT-Traversal mechanism.

6.6 PFS: Perfect forward Secrecy

The PFS (Perfect Forward Secrecy) property is the characteristic of some cryptographic protocols guaranteeing that an attacker who recorded the exchange of data at a given moment and who was able to obtain cryptographic secrets at a later date is not able to decipher those records.

PFS establishes "narrow" time windows in the sense that the impact of a possible attack cannot (under certain assumptions) extend to windows that have already been closed. This is a risk mitigation measure. This property is obtained (for IPsec) by using a key exchange mechanism called "Ephemeral Diffie-Hellman" (EDH) or its elliptic curve variant (ECDHE).

The granularity degree of protection (the time window) is the session time defined during the IKE key exchange. Thus, each key exchange guarantees that previous exchanges are permanently protected even if subsequent secret is compromised. This is, among other things, the reason why a crypto-period both in time and data volume shall be defined.

This property is only available when using IKE; manual keying will never provide such a mechanism. The granularity of this property can be improved on some devices by using a second key exchange in phase 2 (rather than deriving all the keys from those negotiated in phase 1). This exchange will be renewed several times over the lifetime of a SA. This practice has the effect of shortening the "session" mentioned above, by renewing the cryptographic keys more frequently. One can usually associate a time delay and a maximum amount of data exchanged over a SA, the first quota reached causing the keys renewal.

R12

If available, one shall activate the PFS property in IKEv2 second phase (a.k.a "quick mode") using a Diffie-Hellman key exchange or its elliptic curve variant.

R13

It is recommended to force the periodic renewal of the keys, e.g. every hour and every 100 GB of data, in order to limit the impact of a key compromise.

Note: the exact definition of the renewal period (crypto period) is an organizational decision and relies on the risk analysis of a specific deployment.

6.7 Parameters selection

The recommended choices of cryptographic algorithms and key lengths are covered in the RGS (appendix B1 "Rules and recommendations for the selection and dimensions of cryptographic mechanisms" available at www.ssi.gouv.fr/rgs)[8].

R14

It is strongly discouraged to use the MD5 hash function, DES encryption, RSA keys smaller than 2048 bits or ECDSA keys smaller than 200 bits.

8. The RGS is a security baseline from the French government which defines security criteria for on-line services of the government, local authorities and public establishments

R15

It is discouraged to use 3DES, SHA-1, or ECDSA with keys smaller than 256 bits, if safer alternatives exist such as AES (AES-128 or AES-256), SHA-2 (SHA-224, SHA-256, SHA-384 or SHA-512) or ECDSA with keys of at least 256 bits.

R16

Care must be taken with the Diffie-Hellman groups employed. Groups 1 or 2, very frequently proposed by default, are no longer of an acceptable size. Groups with modulus length being at least 2048 bits (*e.g.* group 14 or group 15) or even groups defined by elliptic curves over prime fields greater or equal to 256 bits (*e.g.* ecp256, aka group 19 or ecp384bp, aka group 29) shall be preferred.