Enabling Practical IPsec Authentication for the Internet

Pedro J. Muñoz Merino, Alberto García-Martínez, Mario Muñoz Organero, and Carlos Delgado Kloos

Universidad Carlos III de Madrid, Department of Telematics Engineering,
Avda de la Universidad, 30
E-28911 Leganés (Madrid), Spain
{pedmume, alberto, munozm, cdk}@it.uc3m.es

Abstract. There is a strong consensus about the need for IPsec, although its use is not widespread for end-to-end communications. One of the main reasons for this is the difficulty for authenticating two end-hosts that do not share a secret or do not rely on a common Certification Authority. In this paper we propose a modification to IKE to use reverse DNS and DNSSEC (named DNSSEC-to-IKE) to provide end-to-end authentication to Internet hosts that do not share any secret, without requiring the deployment of a new infrastructure. We perform a comparative analysis in terms of requirements, provided security and performance with state-of-the-art IKE authentication methods and with a recent proposal for IPv6 based on CGA. We conclude that DNSSEC-to-IKE enables the use of IPsec in a broad range of scenarios in which it was not applicable, at the price of offering slightly less security and incurring in higher performance costs.

1 Introduction

The aim of IPsec ([1], [2], [3], [4]) is the provision of confidentiality, integrity and authenticity. IPsec defines two new types of headers at the IP level: AH (Authentication Header) that provides authenticity and integrity, and ESP (Encapsulating Security Payload) that provides data confidentiality. IPsec can be used in transport mode, in which all the transport and upper layer information is protected, but not the IP header, or tunnel mode, in which the whole IP packet is protected. Although VPNs make intensive use of tunnel mode IPsec between routers, IPsec has not reached widely adoption for protecting end-to-end communications. Furthermore, IPsec implementations like Oportunistic Encryption method proposed by the open source project FreeS/WAN [5] have not been widely adopted.

There are many reasons that explain the low acceptance of IPsec in end-to-end communications [6]. One major obstacle for the adoption of IPsec is the authentication mechanisms available. The traditional methods for authentication are Pre-Shared Key and Digital Signatures (for the rest of the paper the terms Digital Signatures or Certificates will be used to refer to the same authentication mechanism). The use of Pre-Shared Key is only feasible when the number of communicating hosts is small, because of the difficulties of distributing of the key to each pair of hosts. Therefore, this solution is not valid for the casual communications that occurs in the Internet. On

the other hand, the use of certificates requires a common Certificate Authority (CA) between two hosts, but at present no unique CA hierarchy has been adopted in the Internet. Therefore, this solution is not valid between two hosts that do not trust in a common CA. However, new possibilities for authentication have arisen such as the ones derived from the proper use of the Cryptographically Generated Addresses (CGA, [7]), although in this case IPv6 is required.

As it can be inferred from the previous paragraph, present authentication methods for IPsec can only be applied to hosts that conform to some restrictive conditions. In this paper, we try to solve the limited scope of present authentication methods for IPsec. We propose a new authentication method for IKE, DNSSEC-to-IKE that allows authentication for IPsec in a global basis through Internet. This is because the DNSSEC-to-IKE method only require from the hosts to have access to the DNSSEC security infrastructure, provided that a hierarchical DNSSEC infrastructure over Internet exists. With the inclusion of this authentication method one barrier to the usage of IPsec in casual communications could be removed.

DNSSEC-to-IKE is based on the use of DNSSEC security architecture and the ubiquitous reverse DNS infrastructure to store securely the public key that corresponds to a given IP address. The few changes required to IKE to support this mechanism guarantee easy deployment.

We perform a comparison between DNSSEC-to-IKE and state-of-the-art authentication mechanisms, to determine the conditions in which each mechanism can be applied, the security provided and the performance of the validation. We conclude that DNSSEC-to-IKE provides convenient authentication for casual communication between Internet hosts without requiring specific infrastructure or costly key distribution.

The remainder of this paper is organized as follows. In Section 2 there is an outline of the general framework for IPsec authentication; both the state-of-the-art authentication methods and the CGA authentication mechanism are described. Section 3 explains the DNSSEC-to-IKE authentication method. Section 4 makes a comparative analysis of the authentication methods in terms of applicability, security, and performance of the validation process. Finally, Section 5 is devoted to the conclusions.

2 General Overview of the IPsec Authentication Framework

IPsec needs to establish a session security context in order to be used in AH and ESP headers. Although any set of protocols can be used for this functionality, the ones recommended for IPsec are: ISAKMP [8], [9], IKE [10] and OAKLEY [11].

In this work, the main mode was selected for IKE Phase 1. The main mode requires the six ISAKMP messages for each of the three authentication methods. The first two messages negotiate security parameters and the authentication interchange. The next two messages generate a shared Diffie-Hellman key. Finally, the last two messages authenticate the Diffie-Hellman key previously exchanged. Results can be easily extended for aggressive mode due to the similarity of the authentication process.

Among the three authentication methods described by IKE, we do not analyse in detail the Public Encryption method because it results in the same scalability

problems as the Pre-Shared Key one, since it is necessary for each host to know the public key of the rest of the hosts to which it communicates. In the following subsections the authentication model of the IKE protocol is briefly explained for Pre-Shared Key, Digital Signatures [10], and CGA.

2.1 IKE Authentication with Pre-shared Key

In the fifth and the sixth messages of the IKE exchange, the hosts can verify the identity of each other by checking the hashes received that are bound to the Pre-Shared Key agreed for the communication between these two hosts.

2.2 IKE Authentication with Certificates

When Digital Signatures (Certificates) are used, previously to the fifth message, each host can request a certificate to the correspondent host - the certificate is not needed if the host knows in advance the public key of the other host. In the request, a host indicates to the correspondent a list with the CAs in which it trusts. Next, in messages number 5 and 6, the hosts exchange an *Authorization Payload*, which is a signature with its private key. The messages also include a *CERT Payload* with certificates belonging to the CASs contained in the request. The host must check the authenticity of the certificates in order to validate the public key of the other host, and next the Authorization Payload is validated with the public key obtained from the certificate.

2.3 IKE Authentication with CGA

CGA, defined in [7], are IPv6 addresses that incorporate into the 64-bit interface identifier a cryptographic one-way hash of a public key and a prefix owned by the node, creating a binding between this public key and the resulting address. An enhancement to IKE to allow CGA-based authentication is described in [12] as a modification to the Digital Signatures method. In this case, the CERT payload contains the public key and all the parameters required to reconstruct the interface identifier of the address, and therefore validate the authenticity of the IPv6 address of the sending host. The Authorization Payload is signed with the private key associated to the public key of the CGA. Therefore, a host can be authenticated as the legitimate owner of an IPv6 address by checking the validity of the CGA and verifying the Authorization Payload signature with the public key received.

3 DNSSEC-to-IKE Authentication Method

We propose a new method for authentication through IKE, named DNSSEC-to-IKE. This new method can provide authentication to scenarios for which previous authentication schemes were not appropriate. We first present a brief overview of DNSSEC, from which the DNSSEC-to-IKE method derives its authentication infrastructure. Then we describe the DNSSEC-to-IKE method in detail, including the modifications required for the IKE exchange.

3.1 DNSSEC Overview

DNSSEC [13], [14], [15] defines a set of new registers for authenticating the information that is stored in the DNS. The added registers are four:

- DNSKEY, a public key that is associated to a DNS zone.
- RRSIG, a signature of any register of a specific DNS zone with its private key associated to the public key available in DNSKEY. For the rest of the paper we denote the signature of a register R as SIG(R).
- NSEC, which proves that some information does not exist in the DNS.
- DS, a hash of the DNSKEY register of a subzone.

Additionally, an IPSECKEY register for storing IPsec keying material associated to the name to be resolved in DNS [16] has also been added.

An example of how a host can obtain a register associated with a name when DNSSEC is used is the following: The querier host issues a request for a given Fully Qualified Domain Name to the root zone, for which it must know the $DNSKEY|_0$ public key. The resolver sends the query to a root server, and receives a response containing the $DNSKEY|_0$, $SIG|_0(DNSKEY_0)$, $NS|_0$ – Name Server -, $SIG|_0(NS|_0)$, $DS|_0$ and $SIG|_0(DS|_0)$ registers. All of these signatures use the $DNSKEY_0$ of the root zone, as it is referred by the notation $SIG|_0$, so the host, that previously knows the public key for this zone, can validate all the received registers ($DNSKEY|_0$, $NS|_0$ and $DS|_0$). The $NS|_0$ register allows the querier host to access to the server in charge of the subzone requested. Also, the $DS|_0$ register provides a hash of the $DNSKEY_1$ of the first subzone. In general, the resolver will communicate with the server referred in the $NS|_i$ register to obtain its $DNSKEY|_{i+1}$ and the information of the following subzone i+1. The hash of this $DNSKEY|_{i+1}$ is compared with the $DS|_i$ register obtained from the higher level i DNSSEC server. By this way, the $DNSKEY|_{i+1}$ public key of each subzone can be authenticated. The previous steps are repeated as many times as intermediate DNSSEC servers are between the root and the final local server of the host to be resolved.

3.2 Description of the DNSSEC-to-IKE Authentication Method

In this paper, we detail a method for conveying in IKE authentication information based on the use of the reverse DNS infrastructure along with the DNSSEC facilities to provide the public key corresponding to the IP address of a given host for which an IPSECKEY register has been configured, as required by IPsec authentication. The authentication information that is exchanged through IKE includes the chain of successive zone delegations until the leaf zone that corresponds to the IP address in the reverse DNS is reached, along with the signatures stored in the DNSSEC. Then a host can authenticate the IP-based identity for a corresponding host requiring only basic DNSSEC parameters such as the public key of the root zone. The correspondent node will use the root zone public key to validate the chain of zone keys until the public key of the other host is validated. Note that this validation process does not necessarily force accesses to the DNS, since all the information required can be conveyed in the IKE protocol.

The security provided by the DNSSEC-to-IKE method has to be analysed in administrative terms. It is important to note that IPsec authentication requires the

guarantee that a given host is the legitimate owner of a given IP address. Therefore, if an infrastructure is built to perform this IP identity proof, it should be assured that the infrastructure grants the rights to claim for the IP identity to the legitimate owner. The reverse DNS assigns a zone to each IP address by reversing the address in IP notation, and prepending the resulting string to the in-addr.arpa or ip6.arpa suffix, defining a correspondence between an IP address and a reverse DNS leaf zone. The RIRs (Regional Internet Registries[1]) guarantee that only the administration responsible for a range of addresses is assigned the management of the corresponding reverse DNS zone. Reverse DNS zone management is further delegated to clients as addresses are. The assignment of a stable address to a given host is also the responsibility of the administration in charge of the most specific address range. Then, through the appropriate coordination between the end-host administrator and the end-site administrator, the information stored in this reverse DNS infrastructure should correspond to the legitimate IP user.

Next we describe the details of the modifications of IKE to convey this information. Since a chain of successive DNSSEC authentication registers should be interchanged and a list of acceptable DNSSEC root servers for authentication should be suggested, then this is analogous to the Digital Signature authentication method but certificates and CAs are used instead of DNSSEC registers and DNSSEC root servers. For this reason, the IKE Digital Signature method can be taken as a starting point for the definition of the DNSSEC-to-IKE method. For the meaning of the *Certification Request Payload*, *CERT Payload* and *Authentication Payload* and its use inside the IKE Digital Signatures authentication method, the correspondent RFCs should be consulted [8], [10]. Next, we describe how these fields should be used for the DNSSEC-to-IKE method.

The validation process can be accelerated if the verification of some parts of the certificate chain could be omitted. For example, a host in the same network segment as the correspondent knows securely the public key in the DNSSEC for the segment, so there is no need to receive and validate the whole certificate chain beginning from the root zone. Therefore, a *Certification Request Payload* is used to suggest acceptable roots for the authentication, with the following different cases:

1) A Certification Request Payload is sent with a list of zones of the reverse DNS domain name corresponding to the IP address. In this case, the requester knows the DNSKEY public key for any of this zones, so the answer can include the security chain starting from any of this zones (preferably, the chain starting from the most specific zone should be sent to reduce the authentication payload).
2) A Certification Request Payload is sent with an empty zone name. In this case, a host is requesting all the authentication information of all the zones the other host holds, preferably starting from the root DNS zone to avoid the need for accessing to the DNS to perform the validation.
3) The Certification Request Payload is not sent. This occurs when the sender already has the public key of the other host, for example because it obtained it through DNSSEC.

[1] The RIRs have been delegated the responsability for managing the assignment of addresses, reverse DNS zones and autonomous system numbers, in their corresponding regions. Currently there are five RIRs: RIPE, ARIN, APNIC, LACNIC and AFRINIC.

For the each previous cases of Certification Request Payload, we have to consider the following CERT Payload formats:

1) If a Certification Request Payload is sent with a list of zones, then the CERT will contain only one block of DNSSEC-to-IKE information, with the information related to the zone with a more specific or complete name.
2) If the Certification Request Payload is sent with an empty zone name, then the CERT will contain only one block of DNSSEC-to-IKE information with the information correspondent to the zone with a shorter name.
3) If a Certification Request Payload is not sent, then the CERT will be empty.

A block of DNSSEC-to-IKE information contains the chain of successive DNSSEC registers required by a host that knows the DNSKEY of the zone of the reverse DNS name to authenticate to finally obtain the IPSECKEY of the host to be authenticated. Such a block is inserted into the CERT Payload and includes the chain of data composed by N pairs of information of this type:

$DNSKEY|_i$, $SIG|_i$ ($DNSKEY|_i$), $DS|_i$, $SIG|_i$ ($DS|_i$)

N is the number of DNSSEC successive subzones from a root zone to the final zone. All these zones are contained in the name stored in reverse DNS (for example the 4.3.2.1.in-addr.arpa name contains the following subzones: arpa, in-adrr.arpa, 1.in-addr.arpa and so on). Moreover, the block includes the DNSEC information of the public key of the IP address from the final zone:

$IPSECKEY|_N$, $SIG|_N$ ($IPSECKEY|_N$)

In the same way, a resolver DNSSEC host that knows the DNSKEY public key of the root zone can obtain the DNSSEC chain of authentication directly from the DNSSEC infrastructure instead of requesting it in the Certification Request Payload of IKE. However, including all the certification information in the IKE messages reduces the latency for host authentication.

The Authentication Payload includes a signature of the data used to authenticate the entity of the ID payload. In the DNSSEC-to-IKE, it is signed with the private key associated to the public key stored in the DNSSEC server ($IPSECKEY|_N$ register).

4 Comparative Analysis

In this section, a comparative analysis is performed between different authentication methods in terms of security, performance, and applicability. The methods considered for the analysis are Pre-Shared Key, Certificates, the IKE extensions for CGA-based authentication, and the DNSSEC-to-IKE method proposed in this paper.

4.1 Requirements

In this subsection, we discuss the requirements for applying each authentication method, so that it can be understood in which scenarios can be applied.

The Pre-Shared Key method requires the secure distribution of the agreed key before the communication.

The requirements for applying the Certificates method to authenticate a host are the following:

- A shared trusted CA in the hierarchy path of both the requesting host and the host to be authenticated
- Proper management for the Certification Authorities ensuring proper password security, revocation lists, etc.

The basic requirement for applying the CGA is that the authenticated address must be an IPv6 address.

The requirements for applying the DNSSEC-to-IKE authentication method are:

- The host to be authenticated must have an associated public key signed by its zone
- The requesting host must have at least the DNSKEY public key of a reverse DNS zone to which the host to be authenticated belongs, being the root zone a special case for this.

It may be required for the requester to implement resolver functions if the host to be authenticated does not include in the CERT payload the whole chain of signatures beginning at the root key, because the requester would need to use the DNS protocol to access to the keys of the upper zones in the hierarchy. However, this can be removed by requiring at the host to be authenticated to configure the whole chain to its reverse DNS name.

It is relevant to highlight that the deployment of an authentication mechanism for IPsec does not require additional infrastructure to the one naturally provided for securing the DNS, except for the configuration of the key of the reverse name corresponding to the end-host.

With these considerations, we can derive some conclusions:

CGA allows authentication between hosts that do not share any common information, and does not require any kind of infrastructure. These two valuable features are not available in the other authentication methods. Unfortunately, the IPv6 restriction limits drastically the number of practical scenarios in which the requirements are fulfilled, although it could be useful in the future if IPv6 is deployed.

Pre-Shared Key is limited to few hosts, so it cannot be considered as a solution for broad Internet support.

The Certificate based method can be applied for all the hosts that trust a common CA. Consequently, its applicability is restricted, since a common CA infrastructure is not currently available for authenticating any pair of hosts in the Internet.

Finally, the DNSSEC-to-IKE method enables the possibility of authenticating any pair of hosts in the Internet, since the DNSSEC is expected to provide an Internet-wide trust infrastructure. Although the DNSSEC infrastructure has not been fully deployed, recent advances have been made for the specific support of reverse DNS zones: the reverse DNS zones depending from RIPE (www.ripe.net) are signed since January 2006. The deployment of DNSSEC for reverse DNS zones by the RIRs would enable access to the reverse DNS with the authentication being provided by a small number of keys (as many as RIRs).

4.2 Security

The specific risks for the authentication method based in Pre-Shared Key depend on the security of the key distribution process, on the hash functions and on the private keys exchanged, apart from the risks associated to the management of the private key in the host. Pre-Shared Key authentication does not depend on the trust of a third party.

The security of the Certificates method is based on several aspects: the trust model used for issuing certificates from a Certification authority to another or to a final user, the security provided by the private keys of the CAs, the private keys of the entities to be authenticated, the revocation lists, and the appropriate management of the private keys in all the CAs and the authenticated host. In this authentication mechanism, the user must trust in one or several third parties.

For CGA-based authentication, the weakest security element is imposed by the limitation to 64 bits in the interface identifier of the resulting address. An attacker can try to generate private/public key pairs until the 64-bit hash containing the public key equals to the legitimate one. CGA provides some means of making more difficult this attack through a SEC parameter contained in the IPv6 address that imposes an additional condition to the CGA structure, requiring the last 16*SEC bytes of a hash different from the one used to obtain the address to be 0. Then, an attacker willing to hijack the CGA identity requires $O(2^{59+16 * SEC})$, ranging SEC from 0 to 7. The condition expressed by the SEC parameter affects only the time to generate the CGA, not the time required for validating its identity. Additionally, the security of the public key used for the hash, and proper management of the private key are relevant for the protection provided by this authentication method.

The elements affecting the security of the DNSSEC-to-IKE mechanism are the trust model for delegating reverse DNS domains to the administrations responsible for the corresponding addresses, the strength of the public key of the host and of the DNS zones, and the security of the private keys used by each element in the authentication chain. Additionally, several signatures generated by the private keys are public, so some attacks to the private keys used for signing are enabled.

As a conclusion, the security offered by the presented authentication methods greatly varies. The Pre-Shared Key authentication method is the most robust, since the key used can be as strong as required, and the distribution method can be devised to involve a few entities, therefore reducing risks. Certificates are also robust, allowing proper selection of the key length, although any of the several entities involved could be compromised, leading to a security disruption. DNSSEC-to-IKE is similar to Certificates, but it does not allow revocations, providing lower security. Finally, CGA provide the weakest security, because the interface identifier, that is the main cryptographic token for the authorisation process, is limited to 64 bits, although the security can be increased by proper configuration of the SEC parameter.

4.3 Performance Time

In this section we estimate the computing time that is necessary for authenticating the other host in IKE. For the meaning of the different parameters the IKE RFC [10] should be consulted. D symbolizes decryption and V validation.

The performance time for checking the authenticity of a host with Pre-Shared Key is T and it can be calculated, being prf a pseudorandom function,(in case no prf is negotiated in IKE, then prf corresponds to the negotiated hash in IKE), as follows:

T = T1+T2+T3+T4+T5+T6
T1 = T[SKEYD] = T[prf(pre-shared-key, Ni_b | Nr_b)]
T2 = T[SKEYD_d] = T[prf(SKEYID, g^{xy} | CKY-I | CKY-R | 0)]
T3 = T[SKEYD_a] = T[prf(SKEYID, SKEYID_d | g^{xy} | CKY-I | CKY-R | 1)]
T4 = T[SKEYD_e] = T[prf(SKEYID, SKEYID_a | g^{xy} | CKY-I | CKY-R | 2)]
T5 = T[HASH_I] = T[prf(SKEYID, g^{xi} | g^{xr} | CKY-I | CKY-R | SAi_b | IDii_b)]
T6 = T[D(payloads)]

T1 is the time to compute SKEYD which is a string derived from secret key material that is only known by the two hosts (this time is calculated in a slightly different way in Pre-Shared-Key and Certificates). T2, T3 and T4 are the times to compute different keys that are used in ISAKMP. T5 is the time to compute a specific hash. T6 corresponds to the time to decrypt the payloads of messages 5 and 6 of the IKE exchange. T2, T3, T4, T5 and T6 also appear in the estimation of the authentication cost for the rest of the methods.

The performance time T for checking the authenticity of a host with Digital Signatures (certificates) is calculated as:

T = T1+T2+T3+T4+T3+T4+T5+T6+T7+T8
T1 = T[SKEYID] = T[prf(Ni_b | Nr_b, g^{xy})]
T7 = V(SIGN) +M
T8 = N * T[check one certificate] = N * (V [certificate] + hash[certificate]+M)

T2, T3, T4, T5 and T6 have the same meaning as in the Pre-Shared-Key method while T1 is calculated a bit different. T7 is the time for validating the signed message that is on the Authorization Payload. T8 represents the time for checking the certificates of all the CAs that chain from the common trusted root. N is the number of CAs in this path, including the common trusted root. V[certificate] is the time to validate the encrypted hash of the certificate using the public key of a CA. hash[certificate] is the time to compute the hash of the certificate and M to check if it matches with the decrypted before. The cost expressed in T8 is also incurred when the host does not receive the certificates from the correspondent host but directly from the CAs.

The performance time for checking the authenticity of a host with CGA can be calculated as:

T = T1+T2+T3+T4+T3+T4+T5+T6+T7+T8
T1 = T[SKEYID] = T[prf(Ni_b | Nr_b, g^{xy})]
T7 = V(SIGN) +M
T8 = hash1(public key, parameters) + hash2(public key, parameters) + M

The first seven expressions are the same as in the Certificate method. The difference for CGA is shown in T8, which represents the time required to validate the CGA address checking the association to the public key that has been received. Then, it represents the time for performing the hash1 operation as it is defined in the RFC for CGAs [7], and the hash2 operation if SEC is not 0, and matching the first result with the interface identifier of the address.

The performance time for checking the authenticity of a host with the DNSSEC-to-IKE method is computed as follows:

T = T1+T2+T3+T4+T5+T6+T7+T8
T1 = T[SKEYID] = T[prf(Ni_b | Nr_b, g^{xy})]
T7 = V(SIGN) +M
T8 = N * T[check registers of one server of the chain]= N *T[(V[DNSKEY] + hash[DNSKEY] + M + V[DS] + hash[DS] + M)]
T9 = T[check final registers] = T[V[IPSECKEY] + hash[IPSECKEY] + M]

The first seven expressions are analogous to the method with certificates. T8 represents the time for validating the DNSKEY and DS registers that form the chain of authenticity if N zones are considered. T9 is the time for validating the final IPSECKEY register.

We have measured the CPU time required for the validation with each authentication method in a Pentium 4, 2.1 GHz with Windows XP. For each authentication method, different authentication cases were considered. For each authentication case, the different IKE messages were generated, considering typical payload values. The execution time is the medium time between all the considered cases. For each case, the time is calculated adding the different execution times of each cryptographic operation performed in the PC. It was used for IKE: Diffie-Hellman group 2 (1024 bits), DES as encryption algorithm, HMAC-SHA-1 as hash function and also as prf function. X.509 has been used for certificates and RSA-1024 as digital signature algorithm. The results are shown in Table 1.

Table 1. Time required for the different authentication methods

Authentication method	Execution time (in microseconds)
Pre-Shared-Key	25
CGA	223
Certificates with N=1 (1 CA)	396
Certificates with N=4 (4 CAs)	966
DNSSEC with N=1 (1 intermediate zone)	762
DNSSEC with N=4 (4 intermediate zone)	1870

From the expressions shown above and the times presented in Table 1, we can derive some conclusions. The fastest method is the Pre-Shared Key because only five hash operations and a decryption have to be computed. The second fastest one is CGA, that adds the cost a signature and two hash calculation to the previous case. The most expensive methods are Certificates and DNSSSEC-to-IPsec, and the validation time depends on the number of hierarchical elements between a root and the final authority for both CAs and DNSSEC-to-IKE. Being the number of intermediate elements in the hierarchy the same, then certificates will execute about two times faster than DNSSEC because in DNSSEC it is necessary to perform two validation operations per each hierarchical element, while for certificates only one is required.

5 Conclusions and Future Work

In this paper we have presented a new authentication method for IPsec, DNSSEC-to-IKE for the inclusion of DNSSEC based authentication in the IKE exchange. This new method provides authentication for the use of IPsec between two end hosts in many situations that were not possible with previous authentication methods if the appropriate DNSSEC infrastructure exists.

The DNSSEC-to-IKE method is a variation from the IKE Digital Signatures authentication specification. The security of the information in which the DNSSEC-to-IKE authentication relies is provided by applying DNSSEC to the reverse DNS infrastructure. The IKE exchange contains a chain of authenticating elements that rely on the key of a zone known for both parties, being in the worst case the DNS root zone key, or a limited number of keys such as the ones provided by the Regional Internet Registries. It should be noted that the deployment of an authentication mechanism for DNSSEC-to-IKE does not require additional infrastructure to the one naturally provided for securing the reverse DNS.

We have presented a comparative analysis of the DNSSEC-to-IKE method with traditional IPsec authentication methods and the recently proposed CGA-based authentication in terms of applicability, security, and computing cost. DNSSEC-to-IKE is the only method that enables practical authentication for any two previously unrelated hosts in the current IPv4 Internet by taking advantage of the security infrastructure that is being built now for securing the reverse DNS (for example, through the signature of the reverse DNS zones depending from the RIRs). CGA provides similar features for the IPv6 Internet at a very low performance cost, although offering weaker security. In scenarios in which PKI is practical, the Certificate based authentication provides better security support than DNSSEC-to-IKE by means of revocation lists, and also provide better performance for the authentication process. Finally, the Pre-Shared Key method can only be applied to sets containing a limited number of hosts, although it provides very good security and excellent performance.

As future work, the proposed method can be tested in different scenarios and it can be integrated with different IPsec implementations. In addition, an algorithm can be proposed to execute in each host in order to determine for each situation in real time what is the best authentication method to use.

Acknowledgement. This work has been partially supported by the Programa Nacional de Tecnologías de la Información y de las Comunicaciones, MEC-CICYT project MOSAIC-LEARNING TSI2005-08225-C07-01 and 02 and by the IMPROVISA project TSI2005-07384-C03-02.

References

1. Kent, S., Atkinson, R.: Security Architecture for the Internet Protocol. RFC 2401, (1998)
2. Kent, S., Atkinso, R.: IP Authentication Header. RFC 2402, (1998)
3. Kent, S., Atkinson, R.: IP Encapsulating Security Payload (ESP). RFC 2406, (1998)
4. Thayer, R., Doraswamy, N., Glenn, R., IP Security Document Roadmap, RFC 2411 (1998)
5. FreeS/WAN Project, http://www.freeswan.org/

6. Ionnadis, J.: Why don't we still have IPsec, dammit. NDSS 2003, (2003)
7. Aura, T.: Cryptographically Generated Addresses (CGA). RFC 3972, (2005)
8. Maughan, D., Schertler, M., Schneider, M., Turner, J.: Internet Security Association and Key Management Protocol (ISAKMP). RFC 2408
9. Piper, D.: The Internet IP Security Domain of Interpretation for ISAKMP. RFC 2407, (1998)
10. Harkins, D., Carrel, D.: The Internet Key Exchange (IKE). RFC 2409, (1998)
11. Orman, H.: The OAKLEY Key Determination Protocol. RFC 2412, (1998)
12. Laganier, J.: Using IKE with IPv6 Cryptographically Generated Address. Internet Draft, (2003)
13. Arends, R., Austein, R., Larson, M., Massey, D., Rose, S.: Protocol Modifications for the DNS Security Extensions, RFC 4035 (2005)
14. Arends, R., Austein, R., Larson, M., Massey, D., Rose, S.: Resource Records for the DNS Security Extensions. RFC 4034, (2005)
15. Arends, R., Austein, R., Larson, M., Massey, D., Rose, S.: DNS Security Introduction and Requirements. RFC 4033, (2005)
16. Richardson, M.: A Method for Storing IPsec Keying Material in DNS. RFC 4025, (2005)