INTERNET-DRAFT
<draft-nikander-ipng-address-ownership-00.txt>
Expires September 2001
Track: Informational

P. Nikander
Ericsson NomadicLab

February 2001

An Address Ownership Problem in IPv6

Status of this Memo

Abstract

This document seeks to clarify a number of problems related to authorization of changing local routing information in the current IPv6 architecture. This document does not (currently) cover actual routing protocols. Instead, in IPv6, there are a number of additional mechanisms that allow local routing information to be changed. Some mechanisms are meant to be used only locally, while someof them allow changes to be initiated from a remote location; in the latter case, IPsec is used to protect the relevant signalling messages. However, the current specifications are partially obscure about the actual authorization requirements that must be met in order to be actually secure. The purpose of this document is to clarify the situation, and foster understanding of the potential attacks and required countermeasures.

In this document, we first collect together and summarize the non-routing-protocol mechanisms that allow routing information to be changed. After that, we classify the mechanisms using a couple of orthogonal dimensions. Finally, we discuss the authorization requirements for the different mechanisms.

It is important to note that the security problems discussed in

this document become relevant only when we start to consider multiple security domains. As long as the mechanisms are used only within a single security domain, where all nodes are equally trusted, the problem does not exist. However, if several security domains are connected together, or if anything like "opportunistic IPsec", as promoted by John Gilmore, becomes reality, the problems discussed in this document will become very real.

An other way of expressing the scope of the problems is to say that the attacks can be characterized as insider attacks. In general, the IPsec architecture, as it stands today, is relatively good in keeping outsiders out. However, it is currently not nearly as good at dealing with attacks from within. In a way, when IPsec is used to protect application level traffic, the applications are assumed to take care of their specific protection needs, e.g., in the form of user names, passwords, and application or operating system access control lists. Unfortunately, when IPsec is used to protect traffic signalling, as discussed in this document, the controls do not seem to be adequate. Basically, this is an authorization problem.

Table of Contents

1. Introduction

The basic routing and packet forwarding mechanisms in IPv6 are supposed to be similar to those of IPv4. However, there are considerable differences in the mechanisms how the hosts and routers learn and alter the information used for constructing the actual routing tables entries and other related data structures. The differences are especially large when considering individual hosts, but there are also differences in the way the routers are supposed to be configured. When compared to the current practise in IPv4, the distribution and discovery of this routing information is planned to be highly dynamic in IPv6. While this dynamic approach makes network administration much easier, it is also a potential source of a number of security vulnerabilities. The purpose of this document is to clarify the current situation from the security point of view.

Particularly problematic are the cases where a remote node is able to establish exceptions to the default routing rules. The Mobile IPv6 Binding Update (BU) is a prime example of this; with a Binding Update, a remote node may inform any other node that all traffic sent to the remote node's home address should be currently sent to another address, the so called care-of-address. Unless properly authorized, this and other related mechanisms are potential sources of impersonation, man-in-the-middle, and denial-of-service attacks.

2. Mechanisms for Changing Routing Information

In this section, we discuss a number of mechanisms that allow modification of the effective routing information in IPv6 nodes. The focus on mechanisms that have effect on hosts, but some router related mechanisms are also discussed. It should be noted that not all of these mechanisms are supposed to make changes to the actual routing tables (e.g. Destination Cache) but that some specify the desired alternations in terms of additional data structures. The end effect, however, is always that the effective next hop of for a host or a prefix is changed.

2.1. ICMP Router Discovery

In IPv6, hosts are able to dynamically learn the identity and abilities of local routers dynamically [RFC2461]. Basically, the hosts listen to Router Advertisement messages, and alter their routing information as specified in Section 6.3.4. of RFC2461.

The specification describes the relevant information in terms of a Prefix List and a Default Router List. Basically, the Prefix List contains IPv6 address prefixes that are considered to be on-link, or directly reachable via a physical or pseudo interface. All other IPv6 address prefixes are considered to be off-link, requiring that a router must be selected to forward the packets. The Default Router List contains routers that are directly

reachable. The suggested way of using the Router List is to select routers from the Router List in a round robin fashion for new destinations.

When a Router Advertisement is received, the host updates its Prefix List and Default Router List. Basically, this allows the sender of the Router Advertisement to create new local routes and new default routes. Consequently, a simple way of launching a denial-of-service attack is to claim that a specific prefix is on-link even if it is not. As a result, the host will try to learn the link-level address of the destination trough Neighbor Discovery, and fail. More advanced attacks are easy to imagine.

A basic level of security is achieved by requiring that Router Advertisement messages must have been locally sent (RFC2461 Section 6.1.2.). Additionally, IPsec MAY be used to protect Router Advertisements. Unfortunately, such a practice seems to be quite hard in reality; see [ICMP-IPSEC] for relevant discussion.

2.2. ICMP Redirect

The purpose of the ICMP Redirect [RFC2461] is to inform a host that a particular destination is better reachable through a different router or that the destination happens to be on-link. In effect, a Redirect creates a host specific routing table entry, specifying the next hop for the specific address.

A basic level of security is achieved by requiring that Redirect messages must be have been locally sent, the IP source address of the Redirect is the same as the current default router, the destination address is not a multicast address, and that the target address is a link local address of a local router or that the target address is the destination address, informing that the destination is actually at the local link (RFC2461, Section 8.1). Additionally, IPsec MAY be used to protect the Redirect messages.

2.3. Generic and IPsec Tunneling

RFC2473 specifies a generic tunneling mechanism for IPv6. Currently there are no automatic mechanisms defined for creating Generic Tunnels. Thus, all Generic Tunnels are supposed to be created though administrative actions. However, at the endpoints, a Generic Tunnel is handled as a standard interface. Thus, it is possible to send local ICMP packets through a tunnel to the other tunnel endpoint. Together with Router Advertisement and ICMP Redirect, this may cause unwanted behaviour unless care is taken to screen the packets received from the tunnel.

RFC2401 specifies IPsec tunnel mode. Opposed to Generic Tunneling, IPsec tunnels may be automatically created as a result of IKE negotiation. This creates some potential vulnerabilities. First,

care must be taken that IPsec tunnels are only created to authorized destination subnets; a tunnel creates a routing entry for the specified subnet, leading to similar attacks outlined above. Second, care must be taken to specify the local IKE policies in such a way that the tunnel cannot be used as a source of launching other attacks outlined in this document. Specifically, the SPD rules should be configured in such a fashion that undesired ICMP messages are not accepted from the tunnel exit point. In a way, this latter problem is just one specific case of a larger problem. That is, the IP nodes should be able to handle received ICMP packets differently depending on which interface (physical or virtual) they were received through.

Thus, from the routing information point of view, the ability to automatically create IPsec tunnels with IKE seems to create a new potential vulnerability. Depending on the way policy management is implemented in the IKE, it may be possible to create tunnels that inappropriately divert traffic. That is, if the implementation does not sufficiently bind the credentials of IKE Phase 1 with the client identifiers presented in IKE Phase 2, it may be possible to divert IP traffic to a "wrong" tunnel.

2.4. Router Renumbering

RFC2894 specifies a method for instructing a number of routers to be dynamically configured and reconfigured. All Router Renumbering (RR) packets must be authenticated with IPsec. The specification suggests that the relevant SAs ans SPD entries are created manually, and that the SPD entries explicitly allowing RR ICMP messages. If implemented properly, the suggested practice seems to be appropriate.

However, care must be taken as there seems to exist at least two related potential vulnerabilities. First, unless the default policy for _other_ SAs and related SPDs is to drop RR ICMP messages, it may be possible to make the router to accept unauthorized RRs. In particular, the default set of SPD and SAD entries defined in Section 7 of RFC2894 does not seem to make any distinction between SAs that are authorized to send RRs and other SAs. This may be sufficient if the router does not have Security Associations with any other nodes but those authorized to send RRs. However, if that is not the case, or if the router may create SAs dynamically with IKE, the policy is not sufficient. The alternate policy of explicitly specifying SPD and SAD for each management station and/or trusted routers' unicast addresses, together with a separate default SPD/SAD entry discarding the rest of traffic, seems to be sufficient.

In our humble opinion, a basic problem in RFC2894 is the failure of explicitly classifying the Security Associations into ones trusted as sources of RRs and ones not trusted so. Especially in the IPv6

world, the source address does not necessarily have any security relevance [Nikander01]. From the security point of view, the sender of the RR must be authorized as well as authenticated. This is especially true if an untrustworthy attacker could pass an authentication challenge, but still mount an attack against the router by sending RR's for which it is not authorized.

The second related potential vulnerability may be created through a a use of dynamic SAs, i.e., if IKE and/or a future multicast key agreement protocol is used. Unless the management protocol requires proper authorization when creating new SAs, the result may be SAs that inappropriately allow RRs to be applied.

2.5. Reversing Routing Header

In Section 8.4 of RFC2460, the cases when a Routing Header may be reversed are specified. As the specification stands today, if a received packet has been authenticated and integrity protected with IPsec, the reply packets may carry a reversed routing header.

The intention of the specification seems to be that reversed routing headers are never cached. Thus, they do not directly change the permanent routing information within the host. However, depending of the actual host implementation, a number of vulnerabilities may be exposed.

First, an otherwise passive attacker may effectively turn itself into an active one by replying sending packets with authenticated routing headers. That is, a passive attacker is able to eavesdrop communication, but it may be inable to effectively block it. However, using authenticated routing headers, it may be able to send forged packets in such a way that the replies are sent back directly to it. As a result, the original destination will not receive these replies, making it much harder for it to detect the attack.

Second, depending on the way the routing header handling and IPsec interact with each other at the end host, authenticated routing headers may be used to make a remote host to act as a cryptographic oracle (see Appendix B). Other attacks may also be possible.

2.6. Mobile IPv6

In Mobile IPv6 [MIP6], a Mobile Node (MN) may send a Binding Update (BU) to any host it communicates with. Customarily, the other host is called the Correspondent Node (CN). Basically, the BU creates a temporary routing table entry specifying that all traffic sent to the MN's so called home address is sent to its so called care-of-address (CoA). In Mobile IPv6 terminology, such routing table entries are called Binding Cache entries. The BU must be authenticated and integrity protected with AH.

A vulnerability may be created if the CN has separate AH SAs with several independent nodes. Unless the CN properly checks that the MN sending a BU is actually authorized to specify new routing information for the claimed home address, any host (that has an SA with the CN) may be able to create arbitrary new Binding Cache entries.

At the IPsec level, the proper way to address the situation is to record the home addresses for each SA and to make sure that Binding Updates are only accepted for the recorded home address. The current situation effectively depends on the implementation. The specification requires that the Home Address Destination Option is inserted in the packet _before_ the AH header. This, in a typical implementation the Home Address Destionation Option is processed before AH, thereby making sure that the AH processing finds the home address in the IP header source address field. Respectively, a Binding Update is placed on a separate Destination Option header inserted _after_ the AH header, and therefore typically processed after AH processing.

Now, if the IPsec implementation is strict in its policy filtering, and if the SPD rule associated with the SA only allows packets from a single specified source address, then the rules in Section 8.2. of [MIP6] make sure that the Binding Update actually applies to the address specified in the SPD rule. That is, since the Home Address Destination Option is processed before AH, the AH processing passes or drops the packet based on the home address, only passing packets whose home address equals with the allowed source address specified in the SPD Entry. As a result, the BUs will be processes only in packets whose Home Address Destionation Option has positively matched with the expected source address, as specified in the SPD.

The vulnerability is exposed if the SPD filtering is not performed correctly, or if the SPD rule allows more than a single source address. In the latter cases, the hosts passed by the SPD rule may modify the Binding Update entries for each other.

The larger problem becomes apparent once we consider how the SAs and SPDs are created. As long as Mobile IPv6 is used within a single administrative domain, the problem does not become apprarent. However, as a part of the Mobile IPv6 motivation, it is noted that at least a substantial factor of IPv6 hosts are assumed to be mobile. Therefore, it would be important that Mobile IPv6 can be used between any CN and MN. Indeed, in Section 2. of [MIP6] it is stated that "integration of Route Optimization functionality allows direct routing from any correspondent node to any mobile node, without needing to pass through the mobile node's home network and be forwarded by its home agent, ..."

Thus, since MIP6 Binding Updates are assumed to be used between any CN and MN, it must be possible to create IPsec SAs between any CN and MN. That may not be so easy. However, at least in theory, the ability of running IKE between any nodes may be achieved through various means, e.g., by establishing a global PKI, by using "opportunistic IPsec", etc. Unfortunately, as describe in Section 4. in detail, even a PKI is not necessarily enough in order to make sure that the SPD entries have the right Home Address recorded.

2.7. Inverse Neighbour Discovery

[TBD.]

2.8. SCTP and address sets

[TBD.]

3. Classifications of the Mechanism

In this section, we classify the above discussed mechanisms. First, in Section 3.1. we discuss the locality of the possible origins of the packets. In Section 3.2. we classify the mechanisms based on the extend of the effect of the mechanisms. Here the extent denotes whether the mechanism effects a host's idea of whole subnets or just single hosts. Finally, Section 3.3. contains a table summarizing the classification

3.1. Locality of origination

The ICMP Router Discovery and Redirect mechanisms are designed to be available only to nodes on-link. However, inappropriate tunneling may change the situation.

Generic Tunnels may currently only be created through administrative actions. Thus, the source of origin is meant to be a human administrator.

IPsec tunnels may be created automatically through IKE. In theory, the peer node may be anywhere in the Internet.

Router Renumbering messages may be issued by any authorized node. Consequently, the origin may potentially be anywhere in the Internet. However, the obvious intent is to use RR only within a single security domain.

Authorized Routing Headers may be used by anyone having an SA, and SAs may be created with IKE with anyone in the Internet. Thus, the potential scope of the origin is again the whole Internet. Furthermore, they would be useful between security domains.

Mobile IPv6 Binding Updates are _meant_ to be sent by any Mobile Node anywhere in the Internet, and having their home address

anywhere in the Internet. Thus, the locality of the source is the whole Internet both in the actual BU level and in the level of SA creation. Clearly, there is a strong need to use BUs between security domains.

Inverse ND and SCTP are TBD.

3.2. Extent of effect

Router Discovery messages change hosts' idea of what prefixes are local and what not. Thus, in the extreme, it may be used to claim that all globally routable addresses are in fact local and on-link. In the other end, Router Discovery may be used to claim that a single address (/128 prefix) is on-link.

Redirect messages always affect routing for single hosts.

Tunneling mechanisms may be specified for various prefix lengths, ranging, at least in theory, from /0 to /128.

Router Renumbering changes the routers' idea of prefixes.

Routing Header does not actually create permanent routing information. It only affects the immediate reply packet.

A Mobile IPv6 Binding Update changes the routing information of a single home address.

Inverse ND and SCTP are TBD.

3.3. Classification summary

\ Extent / Origin \	Any prefix in routers	Any prefix in hosts	Single dest only	Reply packet
Admin only	Generic Tunnels			
On-link *)		Router Discovery	Redirect	
Internet, limited domains	Router Renumber, IPsec tunnels	IPsec tunnels		
Internet, multiple domains			MIP6 BU	Routing Header

*) The mechanisms are designed to be available only on-link. However, interaction with tunnels may make them available from anywhere in the Internet.

4. Authorization requirements

In this section, we discuss the authorization requirements for the mechanisms covered above. In this section we only cover the cases where there are several security domains involved. In the case of on-link mechanisms this means, in practice, that some of the nodes are assumed to be potentially hostile towards the others. There clearly are situations where this assumption may be appropriate, e.g., public wireless LAN access.

Due to our exclusion of mechanisms intented to be used within a single domain, or to be configured only manually, we only cover ICMP Router Discovery and Redirect, inter-domain IPsec tunneling, Mobile IPv6 Binding Updates, and reversable Routing Headers.

Our focus is on clarifying the requirements for dynamically creating IPsec Security Associations for various purposes.

4.1. ICMP Router Discovery and Redirect

TBD. See also [ICMP-IPSEC].

4.2. Inter-domain tunneling, Binding Updates, and Routing Headers

In the case of creating inter-domain tunnels, processing Mobile IPv6 Binding Updates, and possibly even when reversing properly authenticated and integrity protected Routing Headers, it becomes important to check that the operation is properly authorized. Let us consider the authorization requirements case-by-case.

- For inter-domain tunnels, the tunnel endpoint must know that the other end is authorized for the addresses that it claims to be reachable through the tunnel. That is, unless otherwise manually configured, the other end must "prove" that it "owns" the addresses that it wants to receive through the tunnel.

  For example, if the remote endpoint is a router, and the actual route to the 3ffe:200:8:3f01::/64 goes through it, it indeed is authorized to request all addresses within 3ffe:200:8:3f01::/64 to be tunneled to it. That is, whether the packets are sent directly or tunneled does not have any impact on the effective routing.

- For Mobile IPv6 Binding Updates, the CN must know that the MN is authorized to the change routing information for its home address. In other words, it must "prove" that it "owns" its home address.

- For reversible Routing Headers, the reversal is safe only if the replying node knows that the SA used to protect the Routing Header is authorized to specify alternate routing information for the final destination. That is, when creating the SA, the peer should "prove" that it "owns" the address to whom alternate routing information may be specified. Even though this practice does not seem to be absolutely necessary, it would certainly stop the attacks outlined in Section 2.5.

In each of the cases, we clearly should check the proper "ownership" of an IP address or prefix. Furthermore, in the case of Mobile IPv6 and Routing Headers this check must be made in two distinct phases: First, when creating the SA, it must be recorded what operations the SA is authorized for. Second, when processing a BU or reversing a Routing Header, it must be checked that the SA is actually authorized for the operation. (For inter-domain tunnels, the second phase is already built in to the IPsec policy filtering mechanism.)

The hardest problem is faced when considering how to create the SAs between two arbitrary security domains. This is discussed next.

4.3. Proving address ownership when dynamically creating SAs

Let us consider the situation where two previously unrelated nodes want to create an IPsec SA with IKE. If they have some common point of reference, such as an X.509v3 CA, they are able to do so subject to their local policy definitions. However, such an SA does not provide any assurance about the honesty and competency levels of the nodes. It only allows packets to be exchanged between the nodes in a secure manner.

Clearly, an IPsec SA created in such a manner is insufficiently authorized for the purposes discussed elsewhere in the document. That is, the nodes cannot trust each other for changing their internal routing information unless otherwise authorized by the local configuration. For example, if one of the nodes is a Mobile Node, the Correspondent Node must not accept Binding Updates from it since it has now way of actually knowing that the claimed home address really belongs to the Mobile Node.

An other way of expressing the situation is the following. In order to accept messages (or mechanisms) that alter internal routing information, the receiving node must know that the originator of each specific message (or application of a mechanism) is authorized to perform the specific suggested change. In other words, the receiving node must know that the originating know is authorized to alter routing information for the specified address or address prefix. One way of expressing this is to say that the originator must "own" the address or address prefix.

Currently there exists no specified mechanism for proving address ownership in Internet-wide scale. Proposing solutions goes beyond the scope of this document.

5. Security and Privacy Considerations

This document discusses security throught the document. Some other related privacy issues are briefly covered in [Nikander01].

6. Intellectual Property Right Notice

For Ericsson policy on IPR issues, see the Ericsson General IPR statement for IETF, http://www.ietf.org/ietf/IPR/ERICSSON-General

Acknowledgements

We want to express our thanks to Michael Thomas and Jari Arkko for their fruitful comments in the initial phases of this work.

References

[RFC2401] S. Kent, R. Atkinson, "Security Architecture for the Internet Protocol, RFC 2401, Internet Engineering Task Force, November 1998.

[RFC2460] S. Deering and R. Hindeng, "Internet Protocol, Version 6 (IPv6) Specification, RFC2460, December 1998.

[RFC2461] T. Narten, E. Nordmark, W. Simpson, "Neinghor Discovery for IP Version 6 (IPv6), RFC2461, December 1998.

[RFC2473] A. Conta and S. Deering, "Generic Packet Tunneling in IPv6 Specification", RFC2473, December 1998.

[RFC2894] M. Crawford, "Router Renumbering for IPv6", RFC2894, August 2000.

[MIP6] David B. Johnson and Charles Perkings, "Mobility Support In IPv6", work in progress, draft-ietf-mobileip-ipv6-13.txt

[ICMP-IPSEC] Jari Arkko et al., "Manual SA Configuration for IPv6 Link Local Messages", work in progress, draft-arkko-manual-icmpv6-sas-00.txt

[Nikander01] Pekka Nikander, "IPv6 Source Addresses Considered Harmful," unpublished manuscript submitted for publication, February 2001. Available from the author on request.

Authors' Addresses

Pekka Nikander
Ericsson Research NomadicLab
phone: +358-40-721 4424
email: Pekka.Nikander@nomadiclab.com

Appendix A. An address "stealing" attack example

Let us consider a server serving a number of future mobile terminals. The actual nature of the server is not important. The terminals and the server communicate with Mobile IPv6. The server is meant to be open, i.e., to serve any hosts using Mobile or non-mobile IPv6.

Following the cryptographic tradition, let us call the server Alice. To simplify the situation, let us have only two clients: Bob, who is honest, and Mallory, who is malicious and whose aim is to "steal" or at least disturb communication with Alice and Bob. It is important to notice that Mallory has selected Bob as his target, and he attempts to perform his attack in such a way that neither Alice or Bob are aware of the attack; the result of a succesfull attack is that Mallory controls, at least to a degree, all communication between Alice and Bob.

Now, if we assume that the MIP6 implementation that Alice uses fully conforms to the standards but is simple minded, the following attack would work.

1. Mallory contacts Alice and creates a pair of AH SAs with her.

2. Using the SAs created in Step 1, Mallory sends a MIP6 Binding Update to Alice, claiming that his Home Address is that of Bob's and that Bob (the Home Address) is currently visiting at Mallory's (care-of-address).

3. Since the Binding Update is protected with AH (see MIPv6 Sec 4.4 and 8.2), Alice accepts the Binding Update. (Note that this is a mistake at Alice's part. However, giving Alice the required knowledge to make the right decision is hard. For more discussion, see [address-ownership-problem].

4. As part of the Binding Update processings (sec 8.3), Alice creates a new Binding Cache entry telling that all future communications to Bob's Home Address should be sent to the address Mallory gave (the CoA).

Now, let us assume that Bob now wants to create a TCP connection with Alice. Thus, he sends a TCP SYN packet, using his home address, to Alice. Alice receives this packet normally, and passes it to TCP. Alice's TCP creates a new TCB and sends a SYN ACK packet, using Bob's home address as the destination address.

5. Now, as a part of output processing (MIP6 sec 8.9), Alice checks its Binding Cache, find a Binding matching the destination address, and adds a Routing Header to route the packet through Mallory to Bob.

6. Mallory receives the SYN ACK packet from Alice.

7. Mallory has now a number of options to further fool Bob.

   a) Mallory may choose to claim Bob that Alice is a mobile node herself, and currently at the location where Mallory is. To do this, he creates a pair of IPsec SAs with Bob (or uses existing ones), and sends a Binding Update claiming that Alice is currently at his address. (This assumes, of course, that Bob makes the same mistake Alice made above at step 3.)

   b) Mallory may choose to claim Bob that Alice is currently (better) reachable through him. To do this, he replaces the routing header that Alice inserted with his own, and uses an existing AH SA (or creates a new) between himself and Bob to protect the routing header to get replies back, as specified in RFC2460 section 8.4. The real difference between this and the a) alternative is that Bob does not insert a Home Address destination option or a Binding Update, but relies on Bob reversing the Routing Header since it is protected with AH.

8. Independent on whether Mallory decides use Binding Updates or Routing Headers, he further has two options on how to handle the data stream in the future

   a) Mallory may decide to act as a man-in-the-middle, passing data between Alice and Bob, and modifying it at need. Furthermore, if the IPsec implementation Alice and Bob are using is simple minded enough, he _may_ be able to fool Alice that she is securely (i.e. with IPSec) talking with Bob, and fool Bob that he is securely talking with Alice.

   b) Mallory may decide to play Alice to Bob, and completely terminate Bob's session with Alice.

This ends the attack description; the only purpose of it is to be illustrate the problem.

A.2. Attack analysis

Even an superficial analysis on the attack description reveals the basic problem: Alice is using an "untrustworthy" SA to accept Binding Updates concerning Bob, and Bob is using equally "untrustworthy" SA to verify Binding Updates or Routing Headers. If we look at a little bit closer to steps 3. and 7. above, we can describe the problem as an

authorization problem.

First, in step 3., the SA that Mallory has created with Alice should NOT be authorized to create a Binding Cache entry for Bob's home address. Similarily, in step 7a., the SA that Mallory has created with Bob should NOT be authorized to create a Binding Cache entry for Alice. Still, in step 7b., the SA that Mallory has created with Bob should NOT be authorized to accept the Routing Header as a reversible Routing Header (RFC2460 sec 8.4.).

Appendix B. A crypto oracle attack example

As mentioned in Section 2.5., depending on the details of the implementation, reversable routing headers may be misused so that a remote host will act as a cryptographic oracle. The attack described is by no means new; a similar kind of attack may, under certain conditions, to be easily launched by a local host without needed routing headers. The noteworthy issue here is that the routing headers allow a remote host to launch such an attack.

For the purposes of this example, a cryptoraphic oracle is an entity that encrypts or integrity protects a given piece of code. To an attacer, the main benefit of an oracle is to allow the attacker to use selected plain text based cryptoanalysis.

Let us consider a situation where Mallory, a Malicious node, wants to cryptanalyse traffic between Alice and Bob. To ease the analysis task, his aim is to use Alice as an cryptographic oracle, thereby makeing his cryptanalysis task easier. However, since Mallory is not local to Alice nor Bob, nor does sit on the path between them, he has to use some other means in order to fool Alice to perform cryptographic operations on his behalf. In this example, we show how the IPv6 Routing Header, under a specific set of conditions, may allow him to do so. It should be understood that this specific example is not an exhaustive study of the problem, but simply a way to illustrate the dangers of not performing proper authorization on all IPsec Security Associations.

In the beginning, Alice and Bob communicate using an IPsec SA. For illustrative purposes, we may assume that this SA is an ESP SA that includes both confidentiality and integrity protection. To allow such communication, Alice has an IPsec SPD Entry specifying that all data to be sent to Bob must be protected with the given SA.

To launch his attack, Mallory first sets up an AH SA with Alice. Taking advantage of the lack of proper authorization checks in many of the current IKE implementations, he establishes an unidirectional SA which looks like being one from Bob to Alice. That is, the SPD entry in the SA specifies that the source address in the received packets must be that of Bob's.

In the second step, Mallory creates a packet that has a Routing Header, AH header, an UDP header directing the packet to the UDP echo service, and a payload that he wants Alice to protect with the existing ESP SA between Alice and Bob. The Routing Header claims that the packet has been originated by Bob, but that is currently being routed through Mallory.

When Alice receives Mallory's packet, she notices that there is a routing header but that the packet is already arriving in its final destination (i.e. she herself). Therefore she processes the packet normally, verifying the integrity and authentication. Since the packet is authenticated, and appears to be originally from Bob, it passes the host local packet filters and is passed to UDP. Depending on the implementation, the packet may be annotated with metadata containing the authenticated routing header. As a final step of the receival process, UDP passes the packet to the echo service.

The echo service simply echos back the UDP payload. In some implementation dependent way, the underlying UDP or IP layer is aware that the packet is a reply to the packet above. Thus, they reverse the authenticated routing header (found in the metadata), creating a routing header that routes the packet to Bob through Mallory. The resulting packet is passed to the IPsec layer. Because the final destination is Bob, the IPsec layer decides that the packet must be protected with the ESP SA that exists between Alice and Bob. The resulting ESP protected packet is sent to the first hop in the route, i.e., Mallory.

Thus, as a result, Mallory receives an encrypted packet that contains an easily predictable UDP header and the payload that it originally sent. In other words, Mallory is efectively using Alice as a cryptographic oracle crypting any plaintext that Mallory chooses.

The example above depends on many implementation details and on some configuration choices. The specific attack is easy to block by modifying some of these. Unfortunately, it seems like that there exists many different cases where similar kinds of attacks are possible. The only real way of preventing these kinds of attacks is to perform proper authorization when creating the SA between Mallory and Alice, and to enforce the authorized permissions in the SPD level.