

Network Working Group
Request for Comments: 3971
Category: Standards Track

J. Arkko, Ed.
Ericsson
J. Kempf
DoCoMo Communications Labs USA
B. Zill
Microsoft
P. Nikander
Ericsson
March 2005

SEcure Neighbor Discovery (SEND)

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2005).

Abstract

IPv6 nodes use the Neighbor Discovery Protocol (NDP) to discover other nodes on the link, to determine their link-layer addresses to find routers, and to maintain reachability information about the paths to active neighbors. If not secured, NDP is vulnerable to various attacks. This document specifies security mechanisms for NDP. Unlike those in the original NDP specifications, these mechanisms do not use IPsec.

Table of Contents

1.	Introduction.	3
1.1.	Specification of Requirements	4
2.	Terms	4
3.	Neighbor and Router Discovery Overview.	6
4.	Secure Neighbor Discovery Overview.	8
5.	Neighbor Discovery Protocol Options	9
5.1.	CGA Option.	10
5.1.1.	Processing Rules for Senders.	11
5.1.2.	Processing Rules for Receivers.	12
5.1.3.	Configuration	13
5.2.	RSA Signature Option.	14
5.2.1.	Processing Rules for Senders.	16
5.2.2.	Processing Rules for Receivers.	16
5.2.3.	Configuration	17
5.2.4.	Performance Considerations.	18
5.3.	Timestamp and Nonce Options	19
5.3.1.	Timestamp Option.	19
5.3.2.	Nonce Option.	20
5.3.3.	Processing Rules for Senders.	21
5.3.4.	Processing Rules for Receivers.	21
6.	Authorization Delegation Discovery.	24
6.1.	Authorization Model	24
6.2.	Deployment Model.	25
6.3.	Certificate Format.	26
6.3.1.	Router Authorization Certificate Profile.	26
6.3.2.	Suitability of Standard Identity Certificates	29
6.4.	Certificate Transport	29
6.4.1.	Certification Path Solicitation Message Format.	30
6.4.2.	Certification Path Advertisement Message Format	32
6.4.3.	Trust Anchor Option	34
6.4.4.	Certificate Option.	36
6.4.5.	Processing Rules for Routers.	37
6.4.6.	Processing Rules for Hosts.	38
6.5.	Configuration	39
7.	Addressing.	40
7.1.	CGAs.	40
7.2.	Redirect Addresses.	40
7.3.	Advertised Subnet Prefixes.	40
7.4.	Limitations	41
8.	Transition Issues	42
9.	Security Considerations	44
9.1.	Threats to the Local Link Not Covered by SEND	44
9.2.	How SEND Counters Threats to NDP.	45
9.2.1.	Neighbor Solicitation/Advertisement Spoofing.	45
9.2.2.	Neighbor Unreachability Detection Failure	46
9.2.3.	Duplicate Address Detection DoS Attack.	46

9.2.4. Router Solicitation and Advertisement Attacks .	46
9.2.5. Replay Attacks.	47
9.2.6. Neighbor Discovery DoS Attack	48
9.3. Attacks against SEND Itself	48
10. Protocol Values	49
10.1. Constants	49
10.2. Variables	49
11. IANA Considerations	49
12. References.	50
12.1. Normative References.	50
12.2. Informative References.	51
Appendices.	53
A. Contributors and Acknowledgments.	53
B. Cache Management.	53
C. Message Size When Carrying Certificates	54
Authors' Addresses.	55
Full Copyright Statements	56

1. Introduction

IPv6 defines the Neighbor Discovery Protocol (NDP) in RFCs 2461 [4] and 2462 [5]. Nodes on the same link use NDP to discover each other's presence and link-layer addresses, to find routers, and to maintain reachability information about the paths to active neighbors. NDP is used by both hosts and routers. Its functions include Neighbor Discovery (ND), Router Discovery (RD), Address Autoconfiguration, Address Resolution, Neighbor Unreachability Detection (NUD), Duplicate Address Detection (DAD), and Redirection.

The original NDP specifications called for the use of IPsec to protect NDP messages. However, the RFCs do not give detailed instructions for using IPsec to do this. In this particular application, IPsec can only be used with a manual configuration of security associations, due to bootstrapping problems in using IKE [19, 15]. Furthermore, the number of manually configured security associations needed for protecting NDP can be very large [20], making that approach impractical for most purposes.

The SEND protocol is designed to counter the threats to NDP. These threats are described in detail in [22]. SEND is applicable in environments where physical security on the link is not assured (such as over wireless) and attacks on NDP are a concern.

This document is organized as follows. Sections 2 and 3 define some terminology and present a brief review of NDP, respectively. Section 4 describes the overall approach to securing NDP. This approach involves the use of new NDP options to carry public key - based signatures. A zero-configuration mechanism is used for showing

address ownership on individual nodes; routers are certified by a trust anchor [7]. The formats, procedures, and cryptographic mechanisms for the zero-configuration mechanism are described in a related specification [11].

The required new NDP options are discussed in Section 5. Section 6 describes the mechanism for distributing certification paths to establish an authorization delegation chain to a trust anchor.

Finally, Section 8 discusses the co-existence of secured and unsecured NDP on the same link, and Section 9 discusses security considerations for SEcure Neighbor Discovery (SEND).

The use of identity certificates provisioned on end hosts for authorizing address use is out of the scope for this document, as is the security of NDP when the entity defending an address is not the same as the entity claiming that address (also known as "proxy ND"). These are extensions of SEND that may be treated in separate documents, should the need arise.

1.1. Specification of Requirements

In this document, several words are used to signify the requirements of the specification. These words are often capitalized. The key words "MUST", "MUST NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", and "MAY" are to be interpreted as described in [2].

2. Terms

Authorization Delegation Discovery (ADD)

A process through which SEND nodes can acquire a certification path from a peer node to a trust anchor.

Certificate Revocation List (CRL)

In one method of certificate revocation, an authority periodically issues a signed data structure called the Certificate Revocation List. This is a time-stamped list identifying revoked certificates, signed by the issuer, and made freely available in a public repository.

Certification Path Advertisement (CPA)

The advertisement message used in the ADD process.

Certification Path Solicitation (CPS)

The solicitation message used in the ADD process.

Cryptographically Generated Address (CGA)

A technique [11] whereby an IPv6 address of a node is cryptographically generated by using a one-way hash function from the node's public key and some other parameters.

Distinguished Encoding Rules (DER)

An encoding scheme for data values, defined in [12].

Duplicate Address Detection (DAD)

A mechanism assuring that two IPv6 nodes on the same link are not using the same address.

Fully Qualified Domain Name (FQDN)

A fully qualified domain name consists of a host and domain name, including the top-level domain.

Internationalized Domain Name (IDN)

Internationalized Domain Names can be used to represent domain names that contain characters outside the ASCII set. See RFC 3490 [9].

Neighbor Discovery (ND)

The Neighbor Discovery function of the Neighbor Discovery Protocol (NDP). NDP contains functions besides ND.

Neighbor Discovery Protocol (NDP)

The IPv6 Neighbor Discovery Protocol [7, 8].

The Neighbor Discovery Protocol is a part of ICMPv6 [6].

Neighbor Unreachability Detection (NUD)

A mechanism used for tracking the reachability of neighbors.

Non-SEND node

An IPv6 node that does not implement this specification but uses only the Neighbor Discovery protocol defined in RFCs 2461 and 2462, as updated, without security.

Nonce

An unpredictable random or pseudo-random number generated by a node and used exactly once. In SEND, nonces are used to assure that a particular advertisement is linked to the solicitation that triggered it.

Router Authorization Certificate

An X.509v3 [7] public key certificate using the profile specified in Section 6.3.1.

SEND node

An IPv6 node that implements this specification.

Router Discovery (RD)

Router Discovery allows the hosts to discover what routers exist on the link, and what subnet prefixes are available. Router Discovery is a part of the Neighbor Discovery Protocol.

Trust Anchor

Hosts are configured with a set of trust anchors to protect Router Discovery. A trust anchor is an entity that the host trusts to authorize routers to act as routers. A trust anchor configuration consists of a public key and some associated parameters (see Section 6.5 for a detailed explanation of these parameters).

3. Neighbor and Router Discovery Overview


The Neighbor Discovery Protocol has several functions. Many of these are overloaded on a few central message types, such as the ICMPv6 Neighbor Advertisement message. In this section, we review some of these tasks and their effects in order to better understand how the messages should be treated. This section is not normative, and if this section and the original Neighbor Discovery RFCs are in conflict, the original RFCs, as updated, take precedence.

The main functions of NDP are as follows:

- o The Router Discovery function allows IPv6 hosts to discover the local routers on an attached link. Router Discovery is described in Section 6 of RFC 2461 [4]. The main purpose of Router Discovery is to find neighboring routers willing to forward packets on behalf of hosts. Subnet prefix discovery involves determining which destinations are directly on a link; this information is necessary in order to know whether a packet should be sent to a router or directly to the destination node.
- o The Redirect function is used for automatically redirecting a host to a better first-hop router, or to inform hosts that a destination is in fact a neighbor (i.e., on-link). Redirect is specified in Section 8 of RFC 2461 [4].
- o Address Autoconfiguration is used for automatically assigning addresses to a host [5]. This allows hosts to operate without explicit configuration related to IP connectivity. The default autoconfiguration mechanism is stateless. To create IP addresses, hosts use any prefix information delivered to them during Router Discovery and then test the newly formed addresses for uniqueness. A stateful mechanism, DHCPv6 [18], provides additional autoconfiguration features.
- o Duplicate Address Detection (DAD) is used for preventing address collisions [5]: for instance, during Address Autoconfiguration. A node that intends to assign a new address to one of its interfaces first runs the DAD procedure to verify that no other node is using the same address. As the rules forbid the use of an address until it has been found unique, no higher layer traffic is possible until this procedure has been completed. Thus, preventing attacks against DAD can help ensure the availability of communications for the node in question.
- o The Address Resolution function allows a node on the link to resolve another node's IPv6 address to the corresponding link-layer address. Address Resolution is defined in Section 7.2 of RFC 2461 [4], and it is used for hosts and routers alike. Again, no higher level traffic can proceed until the sender knows the link layer address of the destination node or the next hop router. Note that the source link layer address on link layer frames is not checked against the information learned through Address Resolution. This allows for an easier addition of network elements such as bridges and proxies and eases the stack implementation requirements, as less information has to be passed from layer to layer.

- o Neighbor Unreachability Detection (NUD) is used for tracking the reachability of neighboring nodes, both hosts and routers. NUD is defined in Section 7.3 of RFC 2461 [4]. NUD is security sensitive, because an attacker could claim that reachability exists when in fact it does not.

The NDP messages follow the ICMPv6 message format. All NDP functions are realized by using the Router Solicitation (RS), Router Advertisement (RA), Neighbor Solicitation (NS), Neighbor Advertisement (NA), and Redirect messages. An actual NDP message includes an NDP message header, consisting of an ICMPv6 header and ND message-specific data, and zero or more NDP options. The NDP message options are formatted in the Type-Length-Value format.



4. Secure Neighbor Discovery Overview

To secure the various functions in NDP, a set of new Neighbor Discovery options is introduced. They are used to protect NDP messages. This specification introduces these options, an authorization delegation discovery process, an address ownership proof mechanism, and requirements for the use of these components in NDP.

The components of the solution specified in this document are as follows:

- o Certification paths, anchored on trusted parties, are expected to certify the authority of routers. A host must be configured with a trust anchor to which the router has a certification path before the host can adopt the router as its default router. Certification Path Solicitation and Advertisement messages are used to discover a certification path to the trust anchor without requiring the actual Router Discovery messages to carry lengthy certification paths. The receipt of a protected Router Advertisement message for which no certification path is available triggers the authorization delegation discovery process.

- o Cryptographically Generated Addresses are used to make sure that the sender of a Neighbor Discovery message is the "owner" of the claimed address. A public-private key pair is generated by all nodes before they can claim an address. A new NDP option, the CGA option, is used to carry the public key and associated parameters.

This specification also allows a node to use non-CGAs with certificates that authorize their use. However, the details of such use are beyond the scope of this specification and are left for future work.

- o A new NDP option, the RSA Signature option, is used to protect all messages relating to Neighbor and Router discovery.

Public key signatures protect the integrity of the messages and authenticate the identity of their sender. The authority of a public key is established either with the authorization delegation process, by using certificates, or through the address ownership proof mechanism, by using CGAs, or with both, depending on configuration and the type of the message protected.

Note: RSA is mandated because having multiple signature algorithms would break compatibility between implementations or increase implementation complexity by forcing the implementation of multiple algorithms and the mechanism to select among them. A second signature algorithm is only necessary as a recovery mechanism, in case a flaw is found in RSA. If this happens, a stronger signature algorithm can be selected, and SEND can be revised. The relationship between the new algorithm and the RSA-based SEND described in this document would be similar to that between the RSA-based SEND and Neighbor Discovery without SEND. Information signed with the stronger algorithm has precedence over that signed with RSA, in the same way that RSA-signed information now takes precedence over unsigned information. Implementations of the current and revised specs would still be compatible.


- o In order to prevent replay attacks, two new Neighbor Discovery options, Timestamp and Nonce, are introduced. Given that Neighbor and Router Discovery messages are in some cases sent to multicast addresses, the Timestamp option offers replay protection without any previously established state or sequence numbers. When the messages are used in solicitation-advertisement pairs, they are protected with the Nonce option.

5. Neighbor Discovery Protocol Options

The options described in this section MUST be supported.

5.1. CGA Option

The CGA option allows the verification of the sender's CGA. The format of the CGA option is described as follows:



Type

11

Length

The length of the option (including the Type, Length, Pad Length, Reserved, CGA Parameters, and Padding fields) in units of 8 octets.

Pad Length

The number of padding octets beyond the end of the CGA Parameters field but within the length specified by the Length field. Padding octets MUST be set to zero by senders and ignored by receivers.

Reserved

An 8-bit field reserved for future use. The value MUST be initialized to zero by the sender and MUST be ignored by the receiver.

CGA Parameters

A variable-length field containing the CGA Parameters data structure described in Section 4 of [11].

This specification requires that if both the CGA option and the RSA Signature option are present, then the public key found from the CGA Parameters field in the CGA option MUST be that referred by the Key Hash field in the RSA Signature option. Packets received with two different keys MUST be silently discarded. Note that a future extension may provide a mechanism allowing the owner of an address and the signer to be different parties.

Padding

A variable-length field making the option length a multiple of 8, containing as many octets as specified in the Pad Length field.

5.1.1. Processing Rules for Senders

If the node has been configured to use SEND, the CGA option MUST be present in all Neighbor Solicitation and Advertisement messages and MUST be present in Router Solicitation messages unless they are sent with the unspecified source address. The CGA option MAY be present in other messages.

A node sending a message using the CGA option MUST construct the message as follows:

The CGA Parameter field in the CGA option is filled according to the rules presented above and in [11]. The public key in the field is taken from the configuration used to generate the CGA, typically from a data structure associated with the source address. The address MUST be constructed as specified in Section 4 of [11]. Depending on the type of the message, this address appears in different places, as follows:

Redirect

The address MUST be the source address of the message.

Neighbor Solicitation

The address MUST be the Target Address for solicitations sent for Duplicate Address Detection; otherwise it MUST be the source address of the message.

Neighbor Advertisement

The address **MUST** be the source address of the message.

Router Solicitation

The address **MUST** be the source address of the message. Note that the CGA option is not used when the source address is the unspecified address.

Router Advertisement

The address **MUST** be the source address of the message.

5.1.2. Processing Rules for Receivers

Neighbor Solicitation and Advertisement messages without the CGA option **MUST** be treated as unsecured (i.e., processed in the same way as NDP messages sent by a non-SEND node). The processing of unsecured messages is specified in Section 8. Note that SEND nodes that do not attempt to interoperate with non-SEND nodes **MAY** simply discard the unsecured messages.

Router Solicitation messages without the CGA option **MUST** also be treated as unsecured, unless the source address of the message is the unspecified address.

Redirect, Neighbor Solicitation, Neighbor Advertisement, Router Solicitation, and Router Advertisement messages containing a CGA option **MUST** be checked as follows:

If the interface has been configured to use CGA, the receiving node **MUST** verify the source address of the packet by using the algorithm described in Section 5 of [11]. The inputs to the algorithm are the claimed address, as defined in the previous section, and the CGA Parameters field.

If the CGA verification is successful, the recipient proceeds with a more time-consuming cryptographic check of the signature. Note that even if the CGA verification succeeds, no claims about the validity of the use can be made until the signature has been checked.

A receiver that does not support CGA or has not specified its use for a given interface can still verify packets by using trust anchors, even if a CGA is used on a packet. In such a case, the CGA property of the address is simply left unverified.

5.1.3. Configuration

All nodes that support the verification of the CGA option MUST record the following configuration information:

minbits

The minimum acceptable key length for public keys used in the generation of CGAs. The default SHOULD be 1024 bits. Implementations MAY also set an upper limit for the amount of computation needed when verifying packets that use these security associations. The upper limit SHOULD be at least 2048 bits. Any implementation should follow prudent cryptographic practice in determining the appropriate key lengths.


All nodes that support the sending of the CGA option MUST record the following configuration information:

CGA parameters

Any information required to construct CGAs, as described in [11].

5.2. RSA Signature Option

The RSA Signature option allows public key-based signatures to be attached to NDP messages. The format of the RSA Signature option is described in the following diagram:



Type

12

Length

The length of the option (including the Type, Length, Reserved, Key Hash, Digital Signature, and Padding fields) in units of 8 octets.

Reserved

A 16-bit field reserved for future use. The value MUST be initialized to zero by the sender, and MUST be ignored by the receiver.

Key Hash

A 128-bit field containing the most significant (leftmost) 128 bits of a SHA-1 [14] hash of the public key used for constructing the signature. The SHA-1 hash is taken over the presentation used in the Public Key field of the CGA Parameters data structure carried in the CGA option. Its purpose is to associate the signature to a particular key known by the receiver. Such a key can either be stored in the certificate cache of the receiver or be received in the CGA option in the same message.

Digital Signature

A variable-length field containing a PKCS#1 v1.5 signature, constructed by using the sender's private key over the following sequence of octets:

1. The 128-bit CGA Message Type tag [11] value for SEND, 0x086F CA5E 10B2 00C9 9C8C E001 6427 7C08. (The tag value has been generated randomly by the editor of this specification.).
2. The 128-bit Source Address field from the IP header.
3. The 128-bit Destination Address field from the IP header.
4. The 8-bit Type, 8-bit Code, and 16-bit Checksum fields from the ICMP header.
5. The NDP message header, starting from the octet after the ICMP Checksum field and continuing up to but not including NDP options.
6. All NDP options preceding the RSA Signature option.

The signature value is computed with the RSASSA-PKCS1-v1_5 algorithm and SHA-1 hash, as defined in [13].

This field starts after the Key Hash field. The length of the Digital Signature field is determined by the length of the RSA Signature option minus the length of the other fields (including the variable length Pad field).

Padding

This variable-length field contains padding, as many bytes long as remain after the end of the signature.

5.2.1. Processing Rules for Senders

If the node has been configured to use SEND, Neighbor Solicitation, Neighbor Advertisement, Router Advertisement, and Redirect messages MUST contain the RSA Signature option. Router Solicitation messages not sent with the unspecified source address MUST contain the RSA Signature option.

A node sending a message with the RSA Signature option MUST construct the message as follows:

- o The message is constructed in its entirety, without the RSA Signature option.
- o The RSA Signature option is added as the last option in the message.
- o The data to be signed is constructed as explained in Section 5.2, under the description of the Digital Signature field.
- o The message, in the form defined above, is signed by using the configured private key, and the resulting PKCS#1 v1.5 signature is put in the Digital Signature field.

5.2.2. Processing Rules for Receivers

Neighbor Solicitation, Neighbor Advertisement, Router Advertisement, and Redirect messages without the RSA Signature option MUST be treated as unsecured (i.e., processed in the same way as NDP messages sent by a non-SEND node). See Section 8.

Router Solicitation messages without the RSA Signature option MUST also be treated as unsecured, unless the source address of the message is the unspecified address.

Redirect, Neighbor Solicitation, Neighbor Advertisement, Router Solicitation, and Router Advertisement messages containing an RSA Signature option MUST be checked as follows:

- o The receiver MUST ignore any options that come after the first RSA Signature option. (The options are ignored for both signature verification and NDP processing purposes.)
- o The Key Hash field MUST indicate the use of a known public key, either one learned from a preceding CGA option in the same message, or one known by other means.

- o The Digital Signature field MUST have correct encoding and MUST not exceed the length of the RSA Signature option minus the Padding.
- o The Digital Signature verification MUST show that the signature has been calculated as specified in the previous section.
- o If the use of a trust anchor has been configured, a valid certification path (see Section 6.3) between the receiver's trust anchor and the sender's public key MUST be known.

Note that the receiver may verify just the CGA property of a packet, even if, in addition to CGA, the sender has used a trust anchor.

Messages that do not pass all the above tests MUST be silently discarded if the host has been configured to accept only secured ND messages. The messages MAY be accepted if the host has been configured to accept both secured and unsecured messages but MUST be treated as an unsecured message. The receiver MAY also otherwise silently discard packets (e.g., as a response to an apparent CPU exhausting DoS attack).

5.2.3. Configuration

All nodes that support the reception of the RSA Signature options MUST allow the following information to be configured for each separate NDP message type:

authorization method

This parameter determines the method through which the authority of the sender is determined. It can have four values:

trust anchor

The authority of the sender is verified as described in Section 6.3. The sender may claim additional authorization through the use of CGAs, but this is neither required nor verified.

CGA

The CGA property of the sender's address is verified as described in [11]. The sender may claim additional authority through a trust anchor, but this is neither required nor verified.

trust anchor and CGA

Both the trust anchor and the CGA verification is required.

trust anchor or CGA

Either the trust anchor or the CGA verification is required.

anchor

The allowed trust anchor(s), if the authorization method is not set to CGA.

All nodes that support sending RSA Signature options MUST record the following configuration information:

keypair

A public-private key pair. If authorization delegation is in use, a certification path from a trust anchor to this key pair must exist.

CGA flag

A flag that indicates whether CGA is used or not. This flag may be per interface or per node. (Note that in future extensions of the SEND protocol, this flag may also be per subnet prefix.)

5.2.4. Performance Considerations

The construction and verification of the RSA Signature option is computationally expensive. In the NDP context, however, hosts typically only have to perform a few signature operations as they enter a link, a few operations as they find a new on-link peer with which to communicate, or Neighbor Unreachability Detection with existing neighbors.


Routers are required to perform a larger number of operations, particularly when the frequency of router advertisements is high due to mobility requirements. Still, the number of required signature operations is on the order of a few dozen per second, some of which can be precomputed as explained below. A large number of router solicitations may cause a higher demand for performing asymmetric operations, although the base NDP protocol limits the rate at which multicast responses to solicitations can be sent.

Signatures can be precomputed for unsolicited (multicast) Neighbor and Router Advertisements if the timing of the future advertisements is known. Typically, solicited neighbor advertisements are sent to the unicast address from which the solicitation was sent. Given that the IPv6 header is covered by the signature, it is not possible to precompute solicited advertisements.

5.3. Timestamp and Nonce Options

5.3.1. Timestamp Option

The purpose of the Timestamp option is to make sure that unsolicited advertisements and redirects have not been replayed. The format of this option is described in the following:



Type

13

Length

The length of the option (including the Type, Length, Reserved, and Timestamp fields) in units of 8 octets; i.e., 2.

Reserved

A 48-bit field reserved for future use. The value MUST be initialized to zero by the sender and MUST be ignored by the receiver.


Timestamp

A 64-bit unsigned integer field containing a timestamp. The value indicates the number of seconds since January 1, 1970, 00:00 UTC, by using a fixed point format. In this format, the integer number of seconds is contained in the first 48 bits of the field, and the remaining 16 bits indicate the number of 1/64K fractions of a second.

Implementation note: This format is compatible with the usual representation of time under UNIX, although the number of bits available for the integer and fraction parts may vary.

5.3.2. Nonce Option

The purpose of the Nonce option is to make sure that an advertisement is a fresh response to a solicitation sent earlier by the node. The format of this option is described in the following:



Type

14

Length

The length of the option (including the Type, Length, and Nonce fields) in units of 8 octets.

Nonce

A field containing a random number selected by the sender of the solicitation message. The length of the random number MUST be at least 6 bytes. The length of the random number MUST be selected so that the length of the nonce option is a multiple of 8 octets.

5.3.3. Processing Rules for Senders

If the node has been configured to use SEND, all solicitation messages **MUST** include a Nonce. When sending a solicitation, the sender **MUST** store the nonce internally so that it can recognize any replies containing that particular nonce.

If the node has been configured to use SEND, all advertisements sent in reply to a solicitation **MUST** include a Nonce, copied from the received solicitation. Note that routers may decide to send a multicast advertisement to all nodes instead of a response to a specific host. In such a case, the router **MAY** still include the nonce value for the host that triggered the multicast advertisement. (Omitting the nonce value may cause the host to ignore the router's advertisement, unless the clocks in these nodes are sufficiently synchronized so that timestamps function properly.)

If the node has been configured to use SEND, all solicitation, advertisement, and redirect messages **MUST** include a Timestamp. Senders **SHOULD** set the Timestamp field to the current time, according to their real time clocks.

5.3.4. Processing Rules for Receivers

The processing of the Nonce and Timestamp options depends on whether a packet is a solicited advertisement. A system may implement the distinction in various ways. Section 5.3.4.1 defines the processing rules for solicited advertisements. Section 5.3.4.2 defines the processing rules for all other messages.

In addition, the following rules apply in all cases:

- o Messages received without at least one of the Timestamp and Nonce options **MUST** be treated as unsecured (i.e., processed in the same way as NDP messages sent by a non-SEND node).
- o Messages received with the RSA Signature option but without the Timestamp option **MUST** be silently discarded.
- o Solicitation messages received with the RSA Signature option but without the Nonce option **MUST** be silently discarded.
- o Advertisements sent to a unicast destination address with the RSA Signature option but without a Nonce option **SHOULD** be processed as unsolicited advertisements.

- o An implementation MAY use some mechanism such as a timestamp cache to strengthen resistance to replay attacks. When there is a very large number of nodes on the same link, or when a cache filling attack is in progress, it is possible that the cache holding the most recent timestamp per sender will become full. In this case, the node MUST remove some entries from the cache or refuse some new requested entries. The specific policy as to which entries are preferred over others is left as an implementation decision. However, typical policies may prefer existing entries to new ones, CGAs with a large Sec value to smaller Sec values, and so on. The issue is briefly discussed in Appendix B.
- o The receiver MUST be prepared to receive the Timestamp and Nonce options in any order, as per RFC 2461 [4], Section 9.

5.3.4.1. Processing Solicited Advertisements

The receiver MUST verify that it has recently sent a matching solicitation, and that the received advertisement contains a copy of the Nonce sent in the solicitation.

If the message contains a Nonce option but the Nonce value is not recognized, the message MUST be silently discarded.

Otherwise, if the message does not contain a Nonce option, it MAY be considered an unsolicited advertisement and processed according to Section 5.3.4.2.

If the message is accepted, the receiver SHOULD store the receive time of the message and the timestamp time in the message, as specified in Section 5.3.4.2.

5.3.4.2. Processing All Other Messages

Receivers SHOULD be configured with an allowed timestamp Delta value, a "fuzz factor" for comparisons, and an allowed clock drift parameter. The recommended default value for the allowed Delta is `TIMESTAMP_DELTA`; for fuzz factor `TIMESTAMP_FUZZ`; and for clock drift, `TIMESTAMP_DRIFT` (see Section 10.2).

To facilitate timestamp checking, each node SHOULD store the following information for each peer:

- o The receive time of the last received and accepted SEND message. This is called `RDlast`.
- o The time stamp in the last received and accepted SEND message. This is called `TSlast`.

An accepted SEND message is any successfully verified Neighbor Solicitation, Neighbor Advertisement, Router Solicitation, Router Advertisement, or Redirect message from the given peer. The RSA Signature option MUST be used in such a message before it can update the above variables.

Receivers SHOULD then check the Timestamp field as follows:

- o When a message is received from a new peer (i.e., one that is not stored in the cache), the received timestamp, TSnew, is checked, and the packet is accepted if the timestamp is recent enough to the reception time of the packet, RDnew:

$$-\text{Delta} < (\text{RDnew} - \text{TSnew}) < +\text{Delta}$$

The RDnew and TSnew values SHOULD be stored in the cache as RDlast and TSlast.

- o If the timestamp is NOT within the boundaries but the message is a Neighbor Solicitation message that the receiver should answer, the receiver SHOULD respond to the message. However, even if it does respond to the message, it MUST NOT create a Neighbor Cache entry. This allows nodes that have large differences in their clocks to continue communicating with each other by exchanging NS/NA pairs.
- o When a message is received from a known peer (i.e., one that already has an entry in the cache), the timestamp is checked against the previously received SEND message:

$$\text{TSnew} + \text{fuzz} > \text{TSlast} + (\text{RDnew} - \text{RDlast}) \times (1 - \text{drift}) - \text{fuzz}$$

If this inequality does not hold, the receiver SHOULD silently discard the message. If, on the other hand, the inequality holds, the receiver SHOULD process the message.

Moreover, if the above inequality holds and TSnew > TSlast, the receiver SHOULD update RDlast and TSlast. Otherwise, the receiver MUST NOT update RDlast or TSlast.

As unsolicited messages may be used in a Denial-of-Service attack to make the receiver verify computationally expensive signatures, all nodes SHOULD apply a mechanism to prevent excessive use of resources for processing such messages.

6. Authorization Delegation Discovery

NDP allows a node to configure itself automatically based on information learned shortly after connecting to a new link. It is particularly easy to configure "rogue" routers on an unsecured link, and it is particularly difficult for a node to distinguish between valid and invalid sources of router information, because the node needs this information before communicating with nodes outside of the link.

As the newly-connected node cannot communicate off-link, it cannot be responsible for searching information to help validate the router(s). However, given a certification path, the node can check someone else's search results and conclude that a particular message comes from an authorized source. In the typical case, a router already connected beyond the link can communicate if necessary with off-link nodes and construct a certification path.

The Secure Neighbor Discovery Protocol mandates a certificate format and introduces two new ICMPv6 messages used between hosts and routers to allow the host to learn a certification path with the assistance of the router.

6.1. Authorization Model

To protect Router Discovery, SEND requires that routers be authorized to act as routers. This authorization is provisioned in both routers and hosts. Routers are given certificates from a trust anchor, and the hosts are configured with the trust anchor(s) to authorize routers. This provisioning is specific to SEND and does not assume that certificates already deployed for some other purpose can be used.

The authorization for routers in SEND is twofold:

- o Routers are authorized to act as routers. The router belongs to the set of routers trusted by the trust anchor. All routers in this set have the same authorization.
- o Optionally, routers may also be authorized to advertise a certain set of subnet prefixes. A specific router is given a specific set of subnet prefixes to advertise; other routers have an authorization to advertise other subnet prefixes. Trust anchors may also delegate a certain set of subnet prefixes to someone (such as an ISP) who, in turn, delegates parts of this set to individual routers.

Note that while communicating with hosts, routers typically also present a number of other parameters beyond the above. For instance, routers have their own IP addresses, subnet prefixes have lifetimes, and routers control the use of stateless and stateful address autoconfiguration. However, the ability to be a router and the subnet prefixes are the most fundamental parameters to authorize. This is because the host needs to choose a router that it uses as its default router, and because the advertised subnet prefixes have an impact on the addresses the host uses. The subnet prefixes also represent a claim about the topological location of the router in the network.

Care should be taken if the certificates used in SEND are also used to provide authorization in other circumstances; for example, with routing protocols. It is necessary to ensure that the authorization information is appropriate for all applications. SEND certificates may authorize a larger set of subnet prefixes than the router is authorized to advertise on a given interface. For instance, SEND allows the use of the null prefix, which might cause verification or routing problems in other applications. It is RECOMMENDED that SEND certificates containing the null prefix are only used for SEND.

Note that end hosts need not be provisioned with their own certified public keys, just as Web clients today do not require end host provisioning with certified keys. Public keys for CGA generation do not need to be certified, as these keys derive their ability to authorize operations on the CGA by the tie to the address.

6.2. Deployment Model

The deployment model for trust anchors can be either a globally rooted public key infrastructure or a more local, decentralized deployment model similar to that currently used for TLS in Web servers. The centralized model assumes a global root capable of authorizing routers and, optionally, the address space they advertise. The end hosts are configured with the public keys of the global root. The global root could operate, for instance, under the Internet Assigned Numbers Authority (IANA) or as a co-operative among Regional Internet Registries (RIRs). However, no such global root currently exists.

In the decentralized model, end hosts are configured with a collection of trusted public keys. The public keys could be issued from various places; for example, a) a public key for the end host's own organization, b) a public key for the end host's home ISP and for ISPs with which the home ISP has a roaming agreement, or c) public keys for roaming brokers acting as intermediaries for ISPs that don't want to run their own certification authority.

This decentralized model works even when a SEND node is used both in networks that have certified routers and in networks that do not. As discussed in Section 8, a SEND node can fall back to the use of a non-SEND router. This makes it possible to start with a local trust anchor even if there is no trust anchor for all possible networks.

6.3. Certificate Format

The certification path of a router terminates in a Router Authorization Certificate that authorizes a specific IPv6 node to act as a router. Because authorization paths are not a common practice in the Internet at the time of this writing, the path MUST consist of standard Public Key Certificates (PKC, in the sense of [8]). The certification path MUST start from the identity of a trust anchor shared by the host and the router. This allows the host to anchor trust for the router's public key in the trust anchor. Note that there MAY be multiple certificates issued by a single trust anchor.

6.3.1. Router Authorization Certificate Profile

Router Authorization Certificates are X.509v3 certificates, as defined in RFC 3280 [7], and SHOULD contain at least one instance of the X.509 extension for IP addresses, as defined in [10]. The parent certificates in the certification path SHOULD contain one or more X.509 IP address extensions, back up to a trusted party (such as the user's ISP) that configured the original IP address block for the router in question, or that delegated the right to do so. The certificates for the intermediate delegating authorities SHOULD contain X.509 IP address extension(s) for subdelegations. The router's certificate is signed by the delegating authority for the subnet prefixes the router is authorized to advertise.

The X.509 IP address extension MUST contain at least one `addressesOrRanges` element. This element MUST contain an `addressPrefix` element containing an IPv6 address prefix for a prefix that the router or the intermediate entity is authorized to route. If the entity is allowed to route any prefix, the IPv6 address prefix used is the null prefix, `::/0`. The `addressFamily` element of the `IPAddrBlocks` sequence element MUST contain the IPv6 Address Family Identifier (0002), as specified in [10], for IPv6 subnet prefixes. Instead of an `addressPrefix` element, the `addressesOrRange` element MAY contain an `addressRange` element for a range of subnet prefixes, if more than one prefix is authorized. The X.509 IP address extension MAY contain additional IPv6 subnet prefixes, expressed as either an `addressPrefix` or an `addressRange`.

A node receiving a Router Authorization Certificate MUST first check whether the certificate's signature was generated by the delegating authority. Then the client SHOULD check whether all the addressPrefix or addressRange entries in the router's certificate are contained within the address ranges in the delegating authority's certificate, and whether the addressPrefix entries match any addressPrefix entries in the delegating authority's certificate. If an addressPrefix or addressRange is not contained within the delegating authority's subnet prefixes or ranges, the client MAY attempt to take an intersection of the ranges/subnet prefixes and to use that intersection. If the resulting intersection is empty, the client MUST NOT accept the certificate. If the addressPrefix in the certificate is missing or is the null prefix, ::/0, the parent prefix or range SHOULD be used. If there is no parent prefix or range, the subnet prefixes that the router advertises are said to be unconstrained (see Section 7.3). That is, the router is allowed to advertise any prefix.

The above checks SHOULD be done for all certificates in the path. If any of the checks fail, the client MUST NOT accept the certificate. The client also has to perform validation of advertised subnet prefixes as discussed in Section 7.3.

Hosts MUST check the subjectPublicKeyInfo field within the last certificate in the certificate path to ensure that only RSA public keys are used to attempt validation of router signatures. Hosts MUST disregard the certificate for SEND if it does not contain an RSA key.

As it is possible that some public key certificates used with SEND do not immediately contain the X.509 IP address extension element, an implementation MAY contain facilities that allow the prefix and range checks to be relaxed. However, any such configuration options SHOULD be switched off by default. The system SHOULD have a default configuration that requires rigorous prefix and range checks.

The following is an example of a certification path. Suppose that isp_group_example.net is the trust anchor. The host has this certificate:

Certificate 1:

Issuer: isp_group_example.net
Validity: Jan 1, 2004 through Dec 31, 2004
Subject: isp_group_example.net
Extensions:
 IP address delegation extension:
 Prefixes: P1, ..., Pk
 ... possibly other extensions ...
 ... other certificate parameters ...

When the host attaches to a link served by router_x.isp_foo_example.net, it receives the following certification path:

Certificate 2:

Issuer: isp_group_example.net
Validity: Jan 1, 2004 through Dec 31, 2004
Subject: isp_foo_example.net
Extensions:
 IP address delegation extension:
 Prefixes: Q1, ..., Qk
 ... possibly other extensions ...
 ... other certificate parameters ...

Certificate 3:

Issuer: isp_foo_example.net
Validity: Jan 1, 2004 through Dec 31, 2004
Subject: router_x.isp_foo_example.net
Extensions:
 IP address delegation extension:
 Prefixes R1, ..., Rk
 ... possibly other extensions ...
 ... other certificate parameters ...

When the three certificates are processed, the usual RFC 3280 [7] certificate path validation is performed. Note, however, that when a node checks certificates received from a router, it typically does not have a connection to the Internet yet, and so it is not possible to perform an on-line Certificate Revocation List (CRL) check, if necessary. Until this check is performed, acceptance of the certificate MUST be considered provisional, and the node MUST perform a check as soon as it has established a connection with the Internet through the router. If the router has been compromised, it could interfere with the CRL check. Should performance of the CRL check be disrupted or should the check fail, the node SHOULD immediately stop using the router as a default and use another router on the link instead.

In addition, the IP addresses in the delegation extension MUST be a subset of the IP addresses in the delegation extension of the issuer's certificate. So in this example, R1, ..., Rs must be a subset of Q1, ..., Qr, and Q1, ..., Qr must be a subset of P1, ..., Pk. If the certification path is valid, then router_foo.isp_foo_example.com is authorized to route the prefixes R1, ..., Rs.

6.3.2. Suitability of Standard Identity Certificates

As deployment of the IP address extension is, itself, not common, a network service provider MAY choose to deploy standard identity certificates on the router to supply the router's public key for signed Router Advertisements.

If there is no prefix information further up in the certification path, a host interprets a standard identity certificate as allowing unconstrained prefix advertisements.

If the other certificates contain prefix information, a standard identity certificate is interpreted as allowing those subnet prefixes.

6.4. Certificate Transport


The Certification Path Solicitation (CPS) message is sent by a host when it wishes to request a certification path between a router and one of the host's trust anchors. The Certification Path Advertisement (CPA) message is sent in reply to the CPS message. These messages are kept separate from the rest of Neighbor and Router Discovery to reduce the effect of the potentially voluminous certification path information on other messages.

The Authorization Delegation Discovery (ADD) process does not exclude other forms of discovering certification paths. For instance, during fast movements, mobile nodes may learn information (including the certification paths) about the next router from a previous router, or nodes may be preconfigured with certification paths from roaming partners.

Where hosts themselves are certified by a trust anchor, these messages MAY also optionally be used between hosts to acquire the peer's certification path. However, the details of such usage are beyond the scope of this specification.

6.4.1. Certification Path Solicitation Message Format

Hosts send Certification Path Solicitations in order to prompt routers to generate Certification Path Advertisements.



IP Fields:

Source Address

A link-local unicast address assigned to the sending interface, or to the unspecified address if no address is assigned to the sending interface.

Destination Address

Typically the All-Routers multicast address, the Solicited-Node multicast address, or the address of the host's default router.

Hop Limit

255

ICMP Fields:

Type

148

Code

0

Checksum

The ICMP checksum [6].

Identifier

A 16-bit unsigned integer field, acting as an identifier to help match advertisements to solicitations. The Identifier field MUST NOT be zero, and its value SHOULD be randomly generated. This randomness does not have to be cryptographically hard, as its purpose is only to avoid collisions.

Component

This 16-bit unsigned integer field is set to 65,535 if the sender seeks to retrieve all certificates. Otherwise, it is set to the component identifier corresponding to the certificate that the receiver wants to retrieve (see Sections 6.4.2 and 6.4.6).

Valid Options:

Trust Anchor


One or more trust anchors that the client is willing to accept. The first (or only) Trust Anchor option MUST contain a DER Encoded X.501 Name; see Section 6.4.3. If there is more than one Trust Anchor option, the options beyond the first may contain any type of trust anchor.

Future versions of this protocol may define new option types. Receivers MUST silently ignore any options they do not recognize and continue processing the message. All included options MUST have a length greater than zero.

ICMP length (derived from the IP length) MUST be 8 or more octets.

6.4.2. Certification Path Advertisement Message Format

Routers send out Certification Path Advertisement messages in response to a Certification Path Solicitation.



IP Fields:

Source Address

A link-local unicast address assigned to the interface from which this message is sent. Note that routers may use multiple addresses, and therefore this address is not sufficient for the unique identification of routers.

Destination Address

Either the Solicited-Node multicast address of the receiver or the link-scoped All-Nodes multicast address.

Hop Limit

255

ICMP Fields:

Type

149

Code

0

Checksum

The ICMP checksum [6].

Identifier

A 16-bit unsigned integer field, acting as an identifier to help match advertisements to solicitations. The Identifier field **MUST** be zero for advertisements sent to the All-Nodes multicast address and **MUST NOT** be zero for others.

All Components

A 16-bit unsigned integer field, used to inform the receiver of the number of certificates in the entire path.

A single advertisement **SHOULD** be broken into separately sent components if there is more than one certificate in the path, in order to avoid excessive fragmentation at the IP layer.

Individual certificates in a path **MAY** be stored and used as received before all the certificates have arrived; this makes the protocol slightly more reliable and less prone to Denial-of-Service attacks.

Examples of packet lengths of Certification Path Advertisement messages for typical certification paths are listed in Appendix C.

Component

A 16-bit unsigned integer field, used to inform the receiver which certificate is being sent.

The first message in an N-component advertisement has the Component field set to N-1, the second set to N-2, and so on. A zero indicates that there are no more components coming in this advertisement.

The sending of path components **SHOULD** be ordered so that the certificate after the trust anchor is sent first. Each certificate sent after the first can be verified with the previously sent certificates. The certificate of the sender comes last. The trust anchor certificate **SHOULD NOT** be sent.

Reserved

An unused field. It **MUST** be initialized to zero by the sender and **MUST** be ignored by the receiver.

Valid Options:

Certificate

One certificate is provided in each Certificate option to establish part of a certification path to a trust anchor.

The certificate of the trust anchor itself SHOULD NOT be sent.

Trust Anchor


Zero or more Trust Anchor options may be included to help receivers decide which advertisements are useful for them. If present, these options MUST appear in the first component of a multi-component advertisement.

Future versions of this protocol may define new option types. Receivers MUST silently ignore any options they do not recognize and continue processing the message. All included options MUST have a length that is greater than zero.

The ICMP length (derived from the IP length) MUST be 8 or more octets.

6.4.3. Trust Anchor Option

The format of the Trust Anchor option is described in the following:



Type

15

Length

The length of the option (including the Type, Length, Name Type, Pad Length, and Name fields), in units of 8 octets.

Name Type

The type of the name included in the Name field. This specification defines two legal values for this field:

- | | |
|---|------------------------|
| 1 | DER Encoded X.501 Name |
| 2 | FQDN |

Pad Length

The number of padding octets beyond the end of the Name field but within the length specified by the Length field. Padding octets MUST be set to zero by senders and ignored by receivers.

Name

When the Name Type field is set to 1, the Name field contains a DER encoded X.501 Name identifying the trust anchor. The value is encoded as defined in [12] and [7].

When the Name Type field is set to 2, the Name field contains a Fully Qualified Domain Name of the trust anchor; for example, "trustanchor.example.com". The name is stored as a string, in the DNS wire format, as specified in RFC 1034 [1]. Additionally, the restrictions discussed in RFC 3280 [7], Section 4.2.1.7 apply.

In the FQDN case, the Name field is an "IDN-unaware domain name slot", as defined in [9]. That is, it can contain only ASCII characters. An implementation MAY support internationalized domain names (IDNs) using the ToASCII operation; see [9] for more information.


All systems MUST support the DER Encoded X.501 Name. Implementations MAY support the FQDN name type.

Padding

A variable-length field making the option length a multiple of 8, beginning after the previous field ends and continuing to the end of the option, as specified by the Length field.

6.4.4. Certificate Option

The format of the certificate option is described in the following:



Type

16

Length

The length of the option (including the Type, Length, Cert Type, Pad Length, and Certificate fields), in units of 8 octets.

Cert Type

The type of the certificate included in the Certificate field. This specification defines only one legal value for this field:

1 X.509v3 Certificate, as specified below

Reserved

An 8-bit field reserved for future use. The value MUST be initialized to zero by the sender and MUST be ignored by the receiver.

Certificate

When the Cert Type field is set to 1, the Certificate field contains an X.509v3 certificate [7], as described in Section 6.3.1.

Padding

A variable length field making the option length a multiple of 8, beginning after the ASN.1 encoding of the previous field [7, 15] ends and continuing to the end of the option, as specified by the Length field.

6.4.5. Processing Rules for Routers

A router MUST silently discard any received Certification Path Solicitation messages that do not conform to the message format defined in Section 6.4.1. The contents of the Reserved field and of any unrecognized options MUST be ignored. Future, backward-compatible changes to the protocol may specify the contents of the Reserved field or add new options; backward-incompatible changes may use different Code values. The contents of any defined options that are not specified to be used with Router Solicitation messages MUST be ignored, and the packet processed in the normal manner. The only defined option that may appear is the Trust Anchor option. A solicitation that passes the validity checks is called a "valid solicitation".

Routers SHOULD send advertisements in response to valid solicitations received on an advertising interface. If the source address in the solicitation was the unspecified address, the router MUST send the response to the link-scoped All-Nodes multicast address. If the source address was a unicast address, the router MUST send the response to the Solicited-Node multicast address corresponding to the source address, except when under load, as specified below. Routers SHOULD NOT send Certification Path Advertisements more than MAX_CPA_RATE times within a second. When there are more solicitations, the router SHOULD send the response to the All-Nodes multicast address regardless of the source address that appeared in the solicitation.

In an advertisement, the router SHOULD include suitable Certificate options so that a certification path can be established to the solicited trust anchor (or a part of it, if the Component field in the solicitation is not equal to 65,535). Note also that a single advertisement is broken into separately sent components and ordered in a particular way (see Section 6.4.2) when there is more than one certificate in the path.

The anchor is identified by the Trust Anchor option. If the Trust Anchor option is represented as a DER Encoded X.501 Name, then the Name must be equal to the Subject field in the anchor's certificate. If the Trust Anchor option is represented as an FQDN, the FQDN must be equal to an FQDN in the subjectAltName field of the anchor's certificate. The router SHOULD include the Trust Anchor option(s) in the advertisement for which the certification path was found.

If the router is unable to find a path to the requested anchor, it SHOULD send an advertisement without any certificates. In this case, the router SHOULD include the Trust Anchor options that were solicited.

6.4.6. Processing Rules for Hosts

A host MUST silently discard any received Certification Path Advertisement messages that do not conform to the message format defined in Section 6.4.2. The contents of the Reserved field, and of any unrecognized options, MUST be ignored. Future, backward-compatible changes to the protocol MAY specify the contents of the Reserved field or add new options; backward-incompatible changes MUST use different Code values. The contents of any defined options not specified to be used with Certification Path Advertisement messages MUST be ignored, and the packet processed in the normal manner. The only defined options that may appear are the Certificate and Trust Anchor options. An advertisement that passes the validity checks is called a "valid advertisement".

Hosts SHOULD store certification paths retrieved in Certification Path Discovery messages if they start from an anchor trusted by the host. The certification paths MUST be verified, as defined in Section 6.3, before storing them. Routers send the certificates one by one, starting from the trust anchor end of the path.

Note: Except to allow for message loss and reordering for temporary purposes, hosts might not store certificates received in a Certification Path Advertisement unless they contain a certificate that can be immediately verified either to the trust anchor or to a certificate that has been verified earlier. This measure is intended to prevent Denial-of-Service attacks, whereby an attacker floods a host with certificates that the host cannot validate and overwhelms memory for certificate storage.

Note that caching this information, and the implied verification results between network attachments for use over multiple attachments to the network, can help improve performance. But periodic certificate revocation checks are still needed, even with cached results, to make sure that the certificates are still valid.

The host SHOULD retrieve a certification path when a Router Advertisement has been received with a public key that is not available from a certificate in the hosts' cache, or when there is no certification path to one of the host's trust anchors. In these situations, the host MAY send a Certification Path Solicitation message to retrieve the path. If there is no response within CPS_RETRY seconds, the message should be retried. The wait interval for each subsequent retransmission MUST exponentially increase, doubling each time. If there is no response after CPS_RETRY_MAX seconds, the host abandons the certification path retrieval process. If the host receives only a part of a certification path within CPS_RETRY_FRAGMENTS seconds of receiving the first part, it MAY in

addition transmit a Certification Path Solicitation message with the Component field set to a value not equal to 65,535. This message can be retransmitted by using the same process as for the initial message. If there are multiple missing certificates, additional CPS messages can be sent after getting a response to first one. However, the complete retrieval process may last at most CPS_RETRY_MAX seconds.

Certification Path Solicitations SHOULD NOT be sent if the host has a currently valid certification path from a reachable router to a trust anchor.

When soliciting certificates for a router, a host MUST send Certification Path Solicitations either to the All-Routers multicast address, if it has not selected a default router yet, or to the default router's IP address, if a default router has already been selected.

If two hosts want to establish trust with the CPS and CPA messages, the CPS message SHOULD be sent to the Solicited-Node multicast address of the receiver. The advertisements SHOULD be sent as specified above for routers. However, the exact details are outside the scope of this specification.

When processing possible advertisements sent as responses to a solicitation, the host MAY prefer to process those advertisements with the same Identifier field value as that of the solicitation first. This makes Denial-of-Service attacks against the mechanism harder (see Section 9.3).

6.5. Configuration

End hosts are configured with a set of trust anchors in order to protect Router Discovery. A trust anchor configuration consists of the following items:

- o A public key signature algorithm and associated public key, which may optionally include parameters.
- o A name as described in Section 6.4.3.
- o An optional public key identifier.
- o An optional list of address ranges for which the trust anchor is authorized.

If the host has been configured to use SEND, it SHOULD possess the above information for at least one trust anchor.

Routers are configured with a collection of certification paths and a collection of certificates containing certified keys, down to the key and certificate for the router itself. Certified keys are required for routers so that a certification path can be established between the router's certificate and the public key of a trust anchor.

If the router has been configured to use SEND, it should be configured with its own key pair and certificate, and with at least one certification path.

7. Addressing

7.1. CGAs

By default, a SEND-enabled node SHOULD use only CGAs for its own addresses. Other types of addresses MAY be used in testing, in diagnostics, or for other purposes. However, this document does not describe how to choose between different types of addresses for different communications. A dynamic selection can be provided by an API, such as the one defined in [21].

7.2. Redirect Addresses

If the Target Address and Destination Address fields in the ICMP Redirect message are equal, then this message is used to inform hosts that a destination is, in fact, a neighbor. In this case, the receiver MUST verify that the given address falls within the range defined by the router's certificate. Redirect messages failing this check MUST be treated as unsecured, as described in Section 7.3.

Note that base NDP rules prevent a host from accepting a Redirect message from a router that the host is not using to reach the destination mentioned in the redirect. This prevents an attacker from tricking a node into redirecting traffic when the attacker is not the default router.

7.3. Advertised Subnet Prefixes

The router's certificate defines the address range(s) that it is allowed to advertise securely. A router MAY, however, advertise a combination of certified and uncertified subnet prefixes. Uncertified subnet prefixes are treated as unsecured (i.e., processed in the same way as unsecured router advertisements sent by non-SEND routers). The processing of unsecured messages is specified in Section 8. Note that SEND nodes that do not attempt to interoperate with non-SEND nodes MAY simply discard the unsecured information.

Certified subnet prefixes fall into the following two categories:

Constrained

If the network operator wants to constrain which routers are allowed to route particular subnet prefixes, routers should be configured with certificates having subnet prefixes listed in the prefix extension. These routers SHOULD advertise the subnet prefixes that they are certified to route, or a subset thereof.

Unconstrained

Network operators that do not want to constrain routers this way should configure routers with certificates containing either the null prefix or no prefix extension at all.

Upon processing a Prefix Information option within a Router Advertisement, nodes SHOULD verify that the prefix specified in this option falls within the range defined by the certificate, if the certificate contains a prefix extension. Options failing this check are treated as containing uncertified subnet prefixes.

Nodes SHOULD use one of the certified subnet prefixes for stateless autoconfiguration. If none of the advertised subnet prefixes match, the host SHOULD use a different advertising router as its default router, if one is available. If the node is performing stateful autoconfiguration, it SHOULD check the address provided by the DHCP server against the certified subnet prefixes and SHOULD NOT use the address if the prefix is not certified.

7.4. Limitations

This specification does not address the protection of NDP packets for nodes configured with a static address (e.g., PREFIX::1). Future certification path-based authorization specifications are needed for these nodes. This specification also does not apply to addresses generated by the IPv6 stateless address autoconfiguration from a fixed interface identifiers (such as EUI-64).

It is outside the scope of this specification to describe the use of trust anchor authorization between nodes with dynamically changing addresses. These addresses may be the result of stateful or stateless address autoconfiguration, or may have resulted from the use of RFC 3041 [17] addresses. If the CGA method is not used, nodes are required to exchange certification paths that terminate in a certificate authorizing a node to use an IP address having a particular interface identifier. This specification does not specify the format of these certificates, as there are currently only a few

cases where they are provided by the link layer, and it is up to the link layer to provide certification for the interface identifier. This may be the subject of a future specification. It is also outside the scope of this specification to describe how stateful address autoconfiguration works with the CGA method.

The Target Address in Neighbor Advertisement is required to be equal to the source address of the packet, except in proxy Neighbor Discovery, which is not supported by this specification.

8. Transition Issues

During the transition to secured links, or as a policy consideration, network operators may want to run a particular link with a mixture of nodes accepting secured and unsecured messages. Nodes that support SEND SHOULD support the use of secured and unsecured NDP messages at the same time.

In a mixed environment, SEND nodes receive both secured and unsecured messages but give priority to secured ones. Here, the "secured" messages are those that contain a valid signature option, as specified above, and "unsecured" messages are those that contain no signature option.

A SEND node SHOULD have a configuration option that causes it to ignore all unsecured Neighbor Solicitation and Advertisement, Router Solicitation and Advertisement, and Redirect messages. This can be used to enforce SEND-only networks. The default for this configuration option SHOULD be that both secured and unsecured messages are allowed.

A SEND node MAY also have a configuration option whereby it disables the use of SEND completely, even for the messages it sends itself. This configuration option SHOULD be switched off by default; that is, SEND is used. Plain (non-SEND) NDP nodes will obviously send only unsecured messages. Per RFC 2461 [4], such nodes will ignore the unknown options and will treat secured messages in the same way that they treat unsecured ones. Secured and unsecured nodes share the same network resources, such as subnet prefixes and address spaces.

SEND nodes configured to use SEND at least in their own messages behave in a mixed environment as explained below.

SEND adheres to the rules defined for the base NDP protocol, with the following exceptions:

- o All solicitations sent by a SEND node MUST be secured.

- o Unsolicited advertisements sent by a SEND node MUST be secured.
- o A SEND node MUST send a secured advertisement in response to a secured solicitation. Advertisements sent in response to an unsecured solicitation MUST be secured as well, but MUST NOT contain the Nonce option.
- o A SEND node that uses the CGA authorization method to protect Neighbor Solicitations SHOULD perform Duplicate Address Detection as follows. If Duplicate Address Detection indicates that the tentative address is already in use, the node generates a new tentative CGA. If after three consecutive attempts no non-unique address is generated, it logs a system error and gives up attempting to generate an address for that interface.

When performing Duplicate Address Detection for the first tentative address, the node accepts both secured and unsecured Neighbor Advertisements and Solicitations received in response to the Neighbor Solicitations. When performing Duplicate Address Detection for the second or third tentative address, it ignores unsecured Neighbor Advertisements and Solicitations. (The security implications of this are discussed in Section 9.2.3 and in [11].)

- o The node MAY have a configuration option whereby it ignores unsecured advertisements, even when performing Duplicate Address Detection for the first tentative address. This configuration option SHOULD be disabled by default. This is a recovery mechanism for cases in which attacks against the first address become common.
- o The Neighbor Cache, Prefix List, and Default Router list entries MUST have a secured/unsecured flag that indicates whether the message that caused the creation or last update of the entry was secured or unsecured. Received unsecured messages MUST NOT cause changes to existing secured entries in the Neighbor Cache, Prefix List, or Default Router List. Received secured messages MUST cause an update of the matching entries, which MUST be flagged as secured.
- o Neighbor Solicitations for the purpose of Neighbor Unreachability Detection (NUD) MUST be sent to that neighbor's solicited-nodes multicast address if the entry is not secured with SEND.

Upper layer confirmations on unsecured neighbor cache entries SHOULD NOT update neighbor cache state from STALE to REACHABLE on a SEND node if the neighbor cache entry has never previously been REACHABLE. This ensures that if an entry spoofing a valid SEND

host is created by a non-SEND attacker without being solicited, NUD will be done with the entry for data transmission within five seconds of use.

As a result, in mixed mode, attackers can take over a Neighbor Cache entry of a SEND node for a longer time only if (a) the SEND node was not communicating with the victim node, so that there is no secure entry for it, and (b) the SEND node is not currently on the link (or is unable to respond).

- o The conceptual sending algorithm is modified so that an unsecured router is selected only if there is no reachable SEND router for the prefix. That is, the algorithm for selecting a default router favors reachable SEND routers over reachable non-SEND ones.
- o A node MAY adopt a router sending unsecured messages, or a router for which secured messages have been received but for which full security checks have not yet been completed, while security checking is underway. Security checks in this case include certification path solicitation, certificate verification, CRL checks, and RA signature checks. A node MAY also adopt a router sending unsecured messages if a router known to be secured becomes unreachable, but because the unreachability may be the result of an attack it SHOULD attempt to find a router known to be secured as soon as possible. Note that although this can speed up attachment to a new network, accepting a router that is sending unsecured messages or for which security checks are not complete opens the node to possible attacks. Nodes that choose to accept such routers do so at their own risk. The node SHOULD, in any case, prefer a router known to be secure as soon as one is made available with completed security checks.

9. Security Considerations

9.1. Threats to the Local Link Not Covered by SEND

SEND does not provide confidentiality for NDP communications.

SEND does not compensate for an unsecured link layer. For instance, there is no assurance that payload packets actually come from the same peer against which the NDP was run.

There may not be cryptographic binding in SEND between the link layer frame address and the IPv6 address. An unsecured link layer could allow nodes to spoof the link layer address of other nodes. An attacker could disrupt IP service by sending out a Neighbor Advertisement on an unsecured link layer, with the link layer source address on the frame set as the source address of a victim, a valid

CGA address and a valid signature corresponding to itself, and a Target Link-layer Address extension corresponding to the victim. The attacker could then make a traffic stream bombard the victim in a DoS attack. This cannot be prevented just by securing the link layer.

Even on a secured link layer, SEND does not require that the addresses on the link layer and Neighbor Advertisements correspond. However, performing these checks is RECOMMENDED if the link layer technology permits.

Prior to participating in Neighbor Discovery and Duplicate Address Detection, nodes must subscribe to the link-scoped All-Nodes Multicast Group and the Solicited-Node Multicast Group for the address that they are claiming as their addresses; RFC 2461 [4]. Subscribing to a multicast group requires that the nodes use MLD [16]. MLD contains no provision for security. An attacker could send an MLD Done message to unsubscribe a victim from the Solicited-Node Multicast address. However, the victim should be able to detect this attack because the router sends a Multicast-Address-Specific Query to determine whether any listeners are still on the address, at which point the victim can respond to avoid being dropped from the group. This technique will work if the router on the link has not been compromised. Other attacks using MLD are possible, but they primarily lead to extraneous (but not necessarily overwhelming) traffic.

9.2. How SEND Counters Threats to NDP

The SEND protocol is designed to counter the threats to NDP, as outlined in [22]. The following subsections contain a regression of the SEND protocol against the threats, to illustrate which aspects of the protocol counter each threat.

9.2.1. Neighbor Solicitation/Advertisement Spoofing

This threat is defined in Section 4.1.1 of [22]. The threat is that a spoofed message may cause a false entry in a node's Neighbor Cache. There are two cases:

1. Entries made as a side effect of a Neighbor Solicitation or Router Solicitation. A router receiving a Router Solicitation with a Target Link-Layer Address extension and the IPv6 source address unequal to the unspecified address inserts an entry for the IPv6 address into its Neighbor Cache. Also, a node performing Duplicate Address Detection (DAD) that receives a Neighbor Solicitation for the same address regards the situation as a collision and ceases to solicit for the address.

In either case, SEND counters these threats by requiring that the RSA Signature and CGA options be present in these solicitations.

SEND nodes can send Router Solicitation messages with a CGA source address and a CGA option, which the router can verify, so that the Neighbor Cache binding is correct. If a SEND node must send a Router Solicitation with the unspecified address, the router will not update its Neighbor Cache, as per base NDP.

2. Entries made as a result of a Neighbor Advertisement message.
SEND counters this threat by requiring that the RSA Signature and CGA options be present in these advertisements.

Also see Section 9.2.5, below, for discussion about replay protection and timestamps.

9.2.2. Neighbor Unreachability Detection Failure

This attack is described in Section 4.1.2 of [22]. SEND counters it by requiring that a node responding to Neighbor Solicitations sent as NUD probes include an RSA Signature option and proof of authorization to use the interface identifier in the address being probed. If these prerequisites are not met, the node performing NUD discards the responses.

9.2.3. Duplicate Address Detection DoS Attack

This attack is described in Section 4.1.3 of [22]. SEND counters this attack by requiring that the Neighbor Advertisements sent as responses to DAD include an RSA Signature option and proof of authorization to use the interface identifier in the address being tested. If these prerequisites are not met, the node performing DAD discards the responses.

When a SEND node performs DAD, it may listen for address collisions from non-SEND nodes for the first address it generates, but not for new attempts. This protects the SEND node from DAD DoS attacks by non-SEND nodes or attackers simulating non-SEND nodes, at the cost of a potential address collision between a SEND node and a non-SEND node. The probability and effects of such an address collision are discussed in [11].

9.2.4. Router Solicitation and Advertisement Attacks

These attacks are described in Sections 4.2.1, 4.2.4, 4.2.5, 4.2.6, and 4.2.7 of [22]. SEND counters them by requiring that Router Advertisements contain an RSA Signature option, and that the signature is calculated by using the public key of a node that can

prove its authorization to route the subnet prefixes contained in any Prefix Information Options. The router proves its authorization by showing a certificate containing the specific prefix or an indication that the router is allowed to route any prefix. A Router Advertisement without these protections is discarded.

SEND does not protect against brute force attacks on the router, such as DoS attacks, or against compromise of the router, as described in Sections 4.4.2 and 4.4.3 of [22].

9.2.5. Replay Attacks

This attack is described in Section 4.3.1 of [22]. SEND protects against attacks in Router Solicitation/Router Advertisement and Neighbor Solicitation/Neighbor Advertisement transactions by including a Nonce option in the solicitation and requiring that the advertisement include a matching option. Together with the signatures, this forms a challenge-response protocol.

SEND protects against attacks from unsolicited messages such as Neighbor Advertisements, Router Advertisements, and Redirects by including a Timestamp option. The following security issues are relevant only for unsolicited messages:

- o A window of vulnerability for replay attacks exists until the timestamp expires.

However, such vulnerabilities are only useful for attackers if the advertised parameters change during the window. Although some parameters (such as the remaining lifetime of a prefix) change often, radical changes typically happen only in the context of some special case, such as switching to a new link layer address due to a broken interface adapter.

SEND nodes are also protected against replay attacks as long as they cache the state created by the message containing the timestamp. The cached state allows the node to protect itself against replayed messages. However, once the node flushes the state for whatever reason, an attacker can re-create the state by replaying an old message while the timestamp is still valid. Because most SEND nodes are likely to use fairly coarse-grained timestamps, as explained in Section 5.3.1, this may affect some nodes.

- o Attacks against time synchronization protocols such as NTP [23] may cause SEND nodes to have an incorrect timestamp value. This can be used to launch replay attacks, even outside the normal window of vulnerability. To protect against these attacks, it is

recommended that SEND nodes keep independently maintained clocks or apply suitable security measures for the time synchronization protocols.

9.2.6. Neighbor Discovery DoS Attack

This attack is described in Section 4.3.2 of [22]. In it, the attacker bombards the router with packets for fictitious addresses on the link, causing the router to busy itself by performing Neighbor Solicitations for addresses that do not exist. SEND does not address this threat because it can be addressed by techniques such as rate limiting Neighbor Solicitations, restricting the amount of state reserved for unresolved solicitations, and clever cache management. These are all techniques involved in implementing Neighbor Discovery on the router.

9.3. Attacks against SEND Itself

The CGAs have a 59-bit hash value. The security of the CGA mechanism has been discussed in [11].

Some Denial-of-Service attacks remain against NDP and SEND itself. For instance, an attacker may try to produce a very high number of packets that a victim host or router has to verify by using asymmetric methods. Although safeguards are required to prevent an excessive use of resources, this can still render SEND non-operational.

When CGA protection is used, SEND deals with the DoS attacks by using the verification process described in Section 5.2.2. In this process, a simple hash verification of the CGA property of the address is performed before the more expensive signature verification. However, even if the CGA verification succeeds, no claims about the validity of the message can be made until the signature has been checked.

When trust anchors and certificates are used for address validation in SEND, the defenses are not quite as effective. Implementations SHOULD track the resources devoted to the processing of packets received with the RSA Signature option and start selectively discarding packets if too many resources are spent. Implementations MAY also first discard packets that are not protected with CGA.

The Authorization Delegation Discovery process may also be vulnerable to Denial-of-Service attacks. An attack may target a router by requesting that a large number of certification paths be discovered for different trust anchors. Routers SHOULD defend against such attacks by caching discovered information (including negative

responses) and by limiting the number of different discovery processes in which they engage.

Attackers may also target hosts by sending a large number of unnecessary certification paths, forcing hosts to spend useless memory and verification resources on them. Hosts can defend against such attacks by limiting the amount of resources devoted to the certification paths and their verification. Hosts SHOULD also prioritize advertisements sent as a response to solicitations the hosts have sent about unsolicited advertisements.

10. Protocol Values

10.1. Constants

Host constants:

CPS_RETRY	1 second
CPS_RETRY_FRAGMENTS	2 seconds
CPS_RETRY_MAX	15 seconds

Router constants:

MAX_CPA_RATE	10 times per second
--------------	---------------------

10.2. Variables

TIMESTAMP_DELTA	300 seconds (5 minutes)
TIMESTAMP_FUZZ	1 second
TIMESTAMP_DRIFT	1 % (0.01)

11. IANA Considerations

This document defines two new ICMP message types, used in Authorization Delegation Discovery. These messages must be assigned ICMPv6 type numbers from the informational message range:

- o The Certification Path Solicitation message (148), described in Section 6.4.1.
- o The Certification Path Advertisement message (149), described in Section 6.4.2.

This document defines six new Neighbor Discovery Protocol [4] options, which must be assigned Option Type values within the option numbering space for Neighbor Discovery Protocol messages:

- o The CGA option (11), described in Section 5.1.

- o The RSA Signature option (12), described in Section 5.2.
- o The Timestamp option (13), described in Section 5.3.1.
- o The Nonce option (14), described in Section 5.3.2.
- o The Trust Anchor option (15), described in Section 6.4.3.
- o The Certificate option (16), described in Section 6.4.4.

This document defines a new 128-bit value under the CGA Message Type [11] namespace, 0x086F CA5E 10B2 00C9 9C8C E001 6427 7C08.

This document defines a new name space for the Name Type field in the Trust Anchor option. Future values of this field can be allocated by using Standards Action [3]. The current values for this field are

- 1 DER Encoded X.501 Name
- 2 FQDN

Another new name space is allocated for the Cert Type field in the Certificate option. Future values of this field can be allocated by using Standards Action [3]. The current values for this field are

- 1 X.509v3 Certificate

12. References

12.1. Normative References

- [1] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, November 1987.
- [2] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [3] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 2434, October 1998.
- [4] Narten, T., Nordmark, E. and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.
- [5] Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 2462, December 1998.

- [6] Conta, A. and S. Deering, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 2463, December 1998.
- [7] Housley, R., Polk, W., Ford, W. and D. Solo, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 3280, April 2002.
- [8] Farrell, S. and R. Housley, "An Internet Attribute Certificate Profile for Authorization", RFC 3281, April 2002.
- [9] Faltstrom, P., Hoffman, P. and A. Costello, "Internationalizing Domain Names in Applications (IDNA)", RFC 3490, March 2003.
- [10] Lynn, C., Kent, S. and K. Seo, "X.509 Extensions for IP Addresses and AS Identifiers", RFC 3779, June 2004.
- [11] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, March 2005.
- [12] International Telecommunications Union, "Information Technology - ASN.1 encoding rules: Specification of Basic Encoding Rules (BER), Canonical Encoding Rules (CER) and Distinguished Encoding Rules (DER)", ITU-T Recommendation X.690, July 2002.
- [13] RSA Laboratories, "RSA Encryption Standard, Version 2.1", PKCS 1, November 2002.
- [14] National Institute of Standards and Technology, "Secure Hash Standard", FIPS PUB 180-1, April 1995, <<http://www.itl.nist.gov/fipspubs/fip180-1.htm>>.

12.2. Informative References

- [15] Harkins, D. and D. Carrel, "The Internet Key Exchange (IKE)", RFC 2409, November 1998.
- [16] Deering, S., Fenner, W. and B. Haberman, "Multicast Listener Discovery (MLD) for IPv6", RFC 2710, October 1999.
- [17] Narten, T. and R. Draves, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 3041, January 2001.
- [18] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C. and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.

- [19] Arkko, J., "Effects of ICMPv6 on IKE and IPsec Policies", Work in Progress, March 2003.
- [20] Arkko, J., "Manual SA Configuration for IPv6 Link Local Messages", Work in Progress, June 2002.
- [21] Nordmark, E., Chakrabarti, S. and J. Laganier, "IPv6 Socket API for Address Selection", Work in Progress, October 2003.
- [22] Nikander, P., Kempf, J., and E. Nordmark, "IPv6 Neighbor Discovery (ND) Trust Models and Threats", RFC 3756, May 2004.
- [23] Bishop, M., "A Security Analysis of the NTP Protocol", Sixth Annual Computer Security Conference Proceedings, December 1990.

Appendix A. Contributors and Acknowledgments

Tuomas Aura contributed the transition mechanism specification in Section 8. Jonathan Trostle contributed the certification path example in Section 6.3.1. Bill Sommerfeld was involved with much of the early design work.

The authors would also like to thank Tuomas Aura, Bill Sommerfeld, Erik Nordmark, Gabriel Montenegro, Pasi Eronen, Greg Daley, Jon Wood, Julien Laganier, Francis Dupont, Pekka Savola, Wenxiao He, Valtteri Niemi, Mike Roe, Russ Housley, Thomas Narten, and Steven Bellovin for interesting discussions in this problem space and for feedback regarding the SEND protocol.

Appendix B. Cache Management

In this section, we outline a cache management algorithm that allows a node to remain partially functional even under a cache-filling DoS attack. This appendix is informational, and real implementations SHOULD use different algorithms in order to avoid the dangers of a mono-cultural code.

There are at least two distinct cache-related attack scenarios:

1. There are a number of nodes on a link, and someone launches a cache filling attack. The goal here is to make sure that the nodes can continue to communicate even if the attack is going on.
2. There is already a cache-filling attack going on, and a new node arrives to the link. The goal here is to make it possible for the new node to become attached to the network, in spite of the attack.

As the intent is to limit the damage to existing, valid cache entries, it is clearly better to be very selective in throwing out entries. Reducing the timestamp Delta value is very discriminatory against nodes with a large clock difference, as an attacker can reduce its clock difference arbitrarily. Throwing out old entries just because their clock difference is large therefore seems like a bad approach.

It is reasonable to have separate cache spaces for new and old entries, where when under attack, the newly cached entries would be more readily dropped. One could track traffic and only allow reasonable new entries that receive genuine traffic to be converted into old cache entries. Although such a scheme can make attacks harder, it will not fully prevent them. For example, an attacker could send a little traffic (i.e., a ping or TCP syn) after each NS

to trick the victim into promoting its cache entry to the old cache. To counter this, the node can be more intelligent in keeping its cache entries than it would be just by having a black/white old/new boundary.

Distinction of the Sec parameter from the CGA Parameters when forcing cache entries out -- by keeping entries with larger Sec parameters preferentially -- also appears to be a possible approach, as CGAs with higher Sec parameters are harder to spoof.

Appendix C. Message Size When Carrying Certificates

In one example scenario using SEND, an Authorization Delegation Discovery test run was made with a certification path length of 4. Three certificates are sent by using Certification Path Advertisement messages, as the trust anchor's certificate is already known by both parties. With a key length of 1024 bits, the certificate lengths in the test run ranged from 864 to 888 bytes; the variation is due to the differences in the certificate issuer names and address prefix extensions. The different certificates had between 1 and 4 address prefix extensions.

The three Certification Path Advertisement messages ranged from 1050 to 1,066 bytes on an Ethernet link layer. The certificate itself accounts for the bulk of the packet. The rest is the trust anchor option, ICMP header, IPv6 header, and link layer header.

Authors' Addresses

Jari Arkko
Ericsson
Jorvas 02420
Finland

EMail: jari.arkko@ericsson.com

James Kempf
DoCoMo Communications Labs USA
181 Metro Drive
San Jose, CA 94043
USA

EMail: kempf@docomolabs-usa.com

Brian Zill
Microsoft Research
One Microsoft Way
Redmond, WA 98052
USA

EMail: bzill@microsoft.com

Pekka Nikander
Ericsson
Jorvas 02420
Finland

EMail: Pekka.Nikander@nomadiclab.com

Full Copyright Statement

Copyright (C) The Internet Society (2005).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

