

Internet Research Task Force (IRTF)
Request for Comments: 6744
Category: Experimental
ISSN: 2070-1721

RJ Atkinson
Consultant
SN Bhatti
U. St Andrews
November 2012

IPv6 Nonce Destination Option for the
Identifier-Locator Network Protocol for IPv6 (ILNPv6)

Abstract

The Identifier-Locator Network Protocol (ILNP) is an experimental, evolutionary enhancement to IP. ILNP has multiple instantiations. This document describes an experimental Nonce Destination Option used only with ILNP for IPv6 (ILNPv6). This document is a product of the IRTF Routing Research Group.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for examination, experimental implementation, and evaluation.

This document defines an Experimental Protocol for the Internet community. This document is a product of the Internet Research Task Force (IRTF). The IRTF publishes the results of Internet-related research and development activities. These results might not be suitable for deployment. This RFC represents the individual opinion(s) of one or more members of the Routing Research Group of the Internet Research Task Force (IRTF). Documents approved for publication by the IRSG are not a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6744>.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (http://trustee.ietf.org/license-info) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

This document may not be modified, and derivative works of it may not be created, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

- 1. Introduction2
 - 1.1. ILNP Document Roadmap3
 - 1.2. Terminology5
- 2. Syntax5
- 3. Transport Protocol Effects6
- 4. Location Changes7
- 5. Implementation Considerations7
 - 5.1. ILNP Communication Cache8
 - 5.2. Mode Indicator8
 - 5.3. IP Security8
- 6. Backwards Compatibility8
- 7. Security Considerations10
- 8. IANA Considerations12
- 9. References12
 - 9.1. Normative References12
 - 9.2. Informative References13
- 10. Acknowledgements14

1. Introduction

This document is part of the ILNP document set, which has had extensive review within the IRTF Routing RG. ILNP is one of the recommendations made by the RG Chairs. Separately, various refereed research papers on ILNP have also been published during this decade. So, the ideas contained herein have had much broader review than the IRTF Routing RG. The views in this document were considered controversial by the Routing RG, but the RG reached a consensus that the document still should be published. The Routing RG has had remarkably little consensus on anything, so virtually all Routing RG outputs are considered controversial.

At present, the Internet research and development community is exploring various approaches to evolving the Internet Architecture to solve a variety of issues including, but not limited to, scalability of inter-domain routing [RFC4984]. A wide range of other issues (e.g., site multihoming, node multihoming, site/subnet mobility, node mobility) are also active concerns at present. Several different classes of evolution are being considered by the Internet research and development community. One class is often called "Map and Encapsulate", where traffic would be mapped and then tunneled through the inter-domain core of the Internet. Another class being considered is sometimes known as "Identifier/Locator Split". This document relates to a proposal that is in the latter class of evolutionary approaches.

This document describes a new option for the IPv6 Destination Options header that is used with the Identifier-Locator Network Protocol for IPv6 (ILNPv6). ILNPv6 is an experimental protocol that is backwards compatible with, and incrementally upgradable from, IPv6. This option is ONLY used in ILNPv6 sessions and is never used with classic IPv6 sessions.

The Nonce Option for the IPv6 Destination Options Header that is described in this document provides two functions. First, it provides protection against off-path attacks for packets when ILNPv6 is in use. Second, it provides a signal during initial network-layer session creation that ILNPv6 is proposed for use with this network-layer session, rather than classic IPv6. This last function is particularly important for ensuring that ILNP is both incrementally deployable and backwards compatible with IPv6. Consequently, this option MUST NOT be used except by an ILNPv6-capable node.

Further, each Nonce value is unidirectional. Since packets often travel asymmetric paths between two correspondents, having separate Nonces for each direction limits the number of on-path nodes that can easily learn an ILNP session's nonce. So a typical TCP session will have two different nonce values in use: one nonce is used from Local Node to the Correspondent Node and a different nonce is used from the Correspondent Node to the Local Node.

1.1. ILNP Document Roadmap

This document defines a new IPv6 Nonce Destination Option used by ILNPv6 nodes (1) to indicate to ILNP correspondent nodes (by inclusion within the initial packets of an ILNP session) that the node is operating in the ILNP mode and (2) to prevent off-path attacks against ILNP ICMP messages. This Nonce is used, for example, with all ILNP ICMPv6 Locator Update messages that are exchanged among ILNP correspondent nodes.

The ILNP architecture can have more than one engineering instantiation. For example, one can imagine a "clean-slate" engineering design based on the ILNP architecture. In separate documents, we describe two specific engineering instances of ILNP. The term "ILNPv6" refers precisely to an instance of ILNP that is based upon, and backwards compatible with, IPv6. The term "ILNPv4" refers precisely to an instance of ILNP that is based upon, and backwards compatible with, IPv4.

Many engineering aspects common to both ILNPv4 and ILNPv6 are described in [RFC6741]. A full engineering specification for either ILNPv6 or ILNPv4 is beyond the scope of this document.

Readers are referred to other related ILNP documents for details not described here:

- a) [RFC6740] is the main architectural description of ILNP, including the concept of operations.
- b) [RFC6741] describes engineering and implementation considerations that are common to both ILNPv4 and ILNPv6.
- c) [RFC6742] defines additional DNS resource records that support ILNP.
- d) [RFC6743] defines a new ICMPv6 Locator Update message used by an ILNP node to inform its correspondent nodes of any changes to its set of valid Locators.
- e) [RFC6745] defines a new ICMPv4 Locator Update message used by an ILNP node to inform its correspondent nodes of any changes to its set of valid Locators.
- f) [RFC6746] defines a new IPv4 Nonce Option used by ILNPv4 nodes to carry a security nonce to prevent off-path attacks against ILNP ICMP messages and also defines a new IPv4 Identifier Option used by ILNPv4 nodes.
- g) [RFC6747] describes extensions to Address Resolution Protocol (ARP) for use with ILNPv4.
- h) [RFC6748] describes optional engineering and deployment functions for ILNP. These are not required for the operation or use of ILNP and are provided as additional options.

1.2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2. Syntax


The Nonce Option is carried within an IPv6 Destination Options header. Section 4 of [RFC2460] provides much more information on the various options and optional headers used with IPv6.

More than one option might be inside the IPv6 Destination Options Header; however, at most, one Nonce Option exists in a given IPv6 packet.

A system that receives a packet containing more than one Nonce Option SHOULD discard the packet as "Authentication Failed" (instead of passing the packet up to the appropriate transport-layer protocol or to ICMP) and SHOULD log the event, including the Source Locator, Source Identifier, Destination Locator, Destination Identifier, upper-layer protocol (e.g., OSPF, TCP, UDP) if any, and transport-layer port numbers (if any), as a security fault in accordance with local logging policies.

As of this writing, IPv6 Destination Options headers, and the options carried by such headers, are extremely uncommon in the deployed Internet. So, it is expected that this Nonce Option commonly would be the only IPv6 Destination Option present in a given IPv6 packet. If a Common Architecture Label IPv6 Security Option (CALIPSO) label option [RFC5570] is also present in the same IPv6 Destination Options header, the CALIPSO Option SHOULD precede the Nonce Option. The Nonce Option SHOULD precede other possible options in the same IPv6 Destination Options header.

In the diagram below, we show not only the Nonce Option but also the IPv6 Destination Options header that carries the Nonce Option.



Next Header: 8-bit selector. Identifies the type of header immediately following the Destination Options header. This field uses the same values as the IPv4 Protocol field, as described in [RFC2460].

Hdr Ext Len: 8-bit unsigned integer. Length of the Destination Options header in 8-octet units, not including the first 8 octets.

Option Type: This contains the value 0x8B (139). This is the first octet of the Nonce Option itself.

Option Length: This indicates the length in 8-bit octets of the Nonce Value field of the Nonce Option. This value must be selected so that the enveloping IPv6 Destination Option complies with the IPv6 header alignment rules. Common values are 4 (when the Nonce Value is 32 bits) and 12 (when the Nonce value is 96 bits).

Nonce Value: An unpredictable cryptographically random value [RFC4086] used to prevent off-path attacks on an ILNP session. This field has variable length, with the length indicated by the Option Length field preceding it. Note that the overall IPv6 IPv6 Destination Option MUST comply with IPv6 header alignment rules. Implementations MUST support sending and receiving 32-bit and 96-bit Nonce values.

3. Transport Protocol Effects

When the initial packet(s) of an IPv6 session contain this Nonce Destination Option, ILNPv6 is in use for that network-layer session. (NOTE: Backwards compatibility and incremental deployment are discussed in more detail in Section 6 below.)

When a network-layer session is using ILNPv6, the transport-layer pseudo-header calculations MUST set to zero the high-order 64-bits ("Locator" or "Routing Prefix") of each IPv6 address. This has the effect that the transport-layer is no longer aware of the topological network location of either node in that transport-layer session.

The preceding rule applies not only to unicast ILNPv6 sessions but also to multicast or anycast ILNPv6 sessions.

4. Location Changes

When a node has a change in its Locator set that causes all previously valid Locators to become invalid, the node MUST send an ICMP Locator Update message (containing the Nonce Option with the appropriate nonce value) to each of its correspondents [RFC6740] [RFC6743].

In the deployed Internet, packets sometimes arrive at a destination out of order. A receiving node MUST drop a packet arriving from a correspondent if the Source Locator of the received packet is not in the receiving node's Identifier-Locator Communication Cache's (ILCC's) Set of Correspondent Locators UNLESS that packet contains a Nonce Option with the appropriate nonce value for that Source Identifier and Destination Identifier pair. This is done to reduce the risk of ILNP session hijacking or ILNP session interference attacks.

Hence, the node that has had all previously valid Locators become invalid MUST include the Nonce Option with the appropriate nonce value in all packets (data or otherwise) to all correspondents for at least three round-trip times (RTTs) for each correspondent. (N.B. An implementation need not actually calculate RTT values; it could just use a fixed timer with a time long enough to cover the longest RTT path, such as 1 minute.) This "gratuitous authentication" ensures that the correspondent can authenticate any received packet, even if the ICMP Locator Update control message arrives and is processed AFTER some other packet using the new Source Locator(s). If an ILNP session is using IPsec, then, of course, IPsec SHOULD continue to be used even if one or more participating nodes change location. Because IP Security for ILNP [RFC6741] binds only to the Identifiers, and not to the Locators in the packet, changes in Locator value have no impact on IP Security for ILNP sessions.

As mobility and multihoming are functionally equivalent for ILNP, this section applies equally to either situation and also to any other situation in which a node's set of Locators might change over time.

5. Implementation Considerations

Implementers may use any internal implementation they wish, PROVIDED that the externally visible behaviour is the same as this implementation approach.

5.1. ILNP Communication Cache

As described in [RFC6741], ILNP nodes maintain an Identifier-Locator Communication Cache (ILCC) that contains several variables for each correspondent. The ILNP Nonce value is an important part of that cache.

5.2. Mode Indicator

To support ILNP, and to retain needed incremental deployability and backwards compatibility, the network layer needs a (logical) mode bit in the Transport Control Block (or equivalent for one's implementation) to track which IP sessions are using traditional IPv6 and which IP sessions are using ILNPv6.

If a given transport-layer session is using ILNP, then an entry corresponding to the network-layer components of that transport-layer session also will exist in the ILNP Communication Cache. Multiple transport-layer sessions between a given pair of nodes normally share a single entry in the ILNP Communication Cache, provided their network-layer details (e.g., Identifiers, Nonces) are identical. Because two different ILNP nodes at two different locations might share the same Identifier value, it is important for an ILNP implementation to use the ILNP Nonce values to distinguish between different ILNP nodes that happen to be using the same (or some of the same) Identifier value(s).

5.3. IP Security

Whether or not ILNP is in use, the IPsec subsystem MUST maintain an IPsec Security Association Database (SAD) and MUST maintain information about which IPsec Selectors apply to traffic received by or sent from the local node [RFC4301]. By combining the information in the IPsec SAD, of what IPsec Selectors apply, and the information in the ILCC, an implementation has sufficient knowledge to apply IPsec properly to both received and transmitted packets.

6. Backwards Compatibility

This option MUST NOT be present in an IPv6 packet unless the packet is part of an ILNPv6 session. As is explained below in more detail, the presence or absence of this option from the initial packets of a new IPv6 session is an important indication of whether the session is using classic IPv6 or ILNPv6.

ILNPv6 nodes MUST include this option in the first few packets of each ILNPv6 session, MUST include this option in all ICMP messages generated by endpoints participating in an ILNPv6 session, and MAY

include this option in any and all packets of an ILNPv6 session. It is recommended that this option be included in all packets of the ILNPv6 session if the packet loss for that session is known to be much higher than normal.

If a node supports ILNP and the node wishes to be able to receive incoming new ILNP sessions, then that node's FQDN SHOULD have one or more Node Identifier (NID) records and also one or more Locator (e.g., L64 or LP) records associated with it in the DNS [RFC6742].

When a host ("initiator") initiates a new IP session with a correspondent ("responder"), it normally will perform a DNS lookup to determine the address(es) of the responder. A host that has been enhanced to support the Identifier/Locator Split operating mode SHOULD look for Node Identifier ("NID") and Locator ("L64") records in any received DNS replies. DNS servers that support Identifier and Locator (i.e., L64 or LP) records might include them (when they exist) as additional data in all DNS replies to DNS queries for DNS A or AAAA records associated with a specified DNS FQDN.

If the initiator supports ILNP, and from DNS data learns that the responder also supports ILNP, then the initiator SHOULD attempt to use ILNP for new sessions with that responder. In such cases, the initiator MUST generate an unpredictable, cryptographically random, ILNP Nonce value, MUST store that ILNP Nonce value in the local ILCC, and MUST include the ILNP Nonce Destination Option in its initial packet(s) to the responder. The IETF has provided advice on generating cryptographically random numbers, such as this nonce value [RFC4086].

If the responder supports ILNP and receives initial packet(s) containing the ILNP Nonce Destination Option, the responder will thereby learn that the initiator supports ILNP and the responder also will use ILNP for this new IP session.

If the responder supports ILNP and receives initial IP packet(s) NOT containing the Nonce Destination Option, the responder will thereby learn that the initiator does NOT support ILNP and the responder will use classic IPv6 for this new IP session.

If the responder does not support ILNP and receives initial packet(s) containing the ILNP Nonce Destination Option, the responder MUST drop the packet and MUST send an ICMP "Parameter Problem" error message back to the initiator [RFC4443]. Indeed, it is not expected that this behaviour will need to be coded into non-ILNP nodes, as this is the normal behaviour for nodes receiving unknown option headers.

If the initiator EITHER does not receive a response from the responder in a timely manner (e.g., within the applicable TCP timeout for a TCP session), and does not receive an ICMP Unreachable error message for that packet, OR receives an ICMP Parameter Problem error message for that packet, then the initiator infers that the responder is not able to support ILNP. In this case, the initiator should try again to create the new IP session, but this time use classic IPv6 and hence MUST NOT include the ILNP Nonce Destination Option.

7. Security Considerations

The ILNPv6 Nonce Destination Option is used ONLY for ILNPv6 sessions, because this option is part of the backwards compatibility and incremental-deployment approach for the Identifier-Locator Network Protocol (ILNP). This option MUST NOT be used with classic IPv6 sessions.

The ILNPv6 Nonce Destination Option only seeks to provide protection against off-path attacks on an IP session. Ordinary IPv6 is vulnerable to on-path attacks unless IPsec is in use [CA-1995-01] [RFC4301]. This option exists to provide non-cryptographic protection for ILNP sessions, protection equivalent to the security of IP sessions that do NOT use IPsec.

When ILNPv6 is in use, the ILNP Nonce Destination Option MUST be included in any ICMP control messages (e.g., ICMP Unreachable, ICMP Locator Update) sent by participants in that ILNPv6 session, even if IPsec also is in use for that ILNPv6 session. Note that certain ICMP messages, for example, a "Packet Too Big" message, might be generated by transit devices that are not aware of the ILNP Nonce in use for that ILNPv6 session; hence, they are not able to include the ILNP Nonce. Again, this is also true of classic IPv6 in the same operational situations, so this does not create a new security issue.

For ILNPv6 sessions, any ICMP control messages received from a participant in that ILNPv6 session that lack a Nonce Destination Option MUST be discarded as forgeries. This security event SHOULD be logged in accordance with local security logging policies, including details of the received packet (i.e., Source Locator, Source Identifier, Destination Locator, Destination Identifier, upper-layer protocol (e.g., TCP, UDP, OSPF) if any, transport-layer port numbers if any, and the date and time the packet was received).

For ILNPv6 sessions, ICMP control messages received from a participant in that ILNPv6 session that have a Nonce Destination Option, but do NOT have the correct nonce value inside the Nonce Destination Option, MUST be discarded as forgeries. This security event SHOULD be logged as described above.

Of course, longer nonce values provide greater resistance to random guessing of the nonce value. However, ILNPv6 sessions operating in higher risk environments SHOULD also use the cryptographic authentication provided by IP Security for ILNP [RFC6741] [RFC4301]. Use of IP Security for ILNP for an ILNPv6 session does not eliminate the need for the ILNPv6 Nonce Option to be included as described here or as described in [RFC6743].

As a performance optimisation, it is suggested that when both the Nonce Option and IPsec are present in a packet and the Nonce Option has not been encrypted the Nonce Option value be checked for validity before beginning IPsec processing. This minimises the ability of an off-path attacker to force the recipient to perform expensive cryptographic computations on received control packets.

For environments with data at differing Sensitivity Levels operating over common infrastructure (e.g., when the IPv6 CALIPSO is deployed), it is recommended that the ILNP Nonce Option be encrypted by using ESP Transport-Mode or ESP Tunnel-Mode in order to reduce the covert channel bandwidth potential created by the Nonce Option and to prevent a node at one Sensitivity Level from attacking an ILNP session at a different Sensitivity Level [RFC5570]. Further, Multi-Level Secure (MLS) systems SHOULD use different nonce values for ILNP sessions having different Sensitivity Levels [RFC5570]. Also, an MLS implementation of ILNP will also store the Sensitivity Label information associated with each ILNP session in the implementation's ILCC. When the Nonce Option and the CALIPSO Option are present in the same IPv6 Destination Options header, the CALIPSO Option SHOULD appear before the Nonce Option.

In all cases, the ILNP Nonce Value MUST be unpredictable and cryptographically random. [RFC4086] provides concrete advice on how to generate a suitable nonce value.

As this is an option within the IPv6 Destination Options header, rather than an option within the IPv6 Hop-by-Hop Option Header, the presence of this option in an IPv6 packet ought not disturb routers along the path an IP packet containing this option happens to travel. Further, many deployed modern IP routers (both IPv4 and IPv6) have been explicitly configured to ignore all IP Options, even including the "Router Alert" option, when forwarding packets not addressed to the router itself. Reports indicate this has been done to preclude use of IP Options as a (Distributed) Denial-of-Service attack vector on backbone routers.

As the Nonce is used in the checksum of all Authentication Header (AH) protected packets, as an implementation hint, it would seem sensible to include the Nonce value from the ILCC for that ILNP session.

8. IANA Considerations

Consistent with the procedures of [RFC2780], IANA has assigned a new IPv6 Destination Option Type value of 0x8B.

The Nonce Option MUST NOT change in transit and MUST be included in IP Authentication Header calculations.

Further, if an end system receives an IPv6 packet containing this option, but does not recognise this option, the end system MUST discard the packet and, regardless of whether or not the received packet's Destination Address was a multicast address, send an ICMPv6 Parameter Problem, Code 2 ("Unrecognised IPv6 Option Encountered"), message to the received packet's Source IPv6 Address, pointing to the unrecognised Option Type.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.
- [RFC2780] Bradner, S. and V. Paxson, "IANA Allocation Guidelines For Values In the Internet Protocol and Related Headers", BCP 37, RFC 2780, March 2000.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.
- [RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 4443, March 2006.
- [RFC6740] Atkinson, R. and S. Bhatti, "Identifier-Locator Network Protocol (ILNP) Architectural Description", RFC 6740, November 2012.

- [RFC6741] Atkinson, R. and S. Bhatti, "Identifier-Locator Network Protocol (ILNP) Engineering and Implementation Considerations", RFC 6741, November 2012.
- [RFC6743] Atkinson, R. and S. Bhatti, "ICMPv6 Locator Update Message", RFC 6743, November 2012.

9.2. Informative References

- [CA-1995-01] US CERT, "CERT Advisory CA-1995-01 IP Spoofing Attacks and Hijacked Terminal Connections", Pittsburgh, PA, USA, 1995.
- [RFC4086] Eastlake 3rd, D., Schiller, J., and S. Crocker, "Randomness Requirements for Security", BCP 106, RFC 4086, June 2005.
- [RFC4984] Meyer, D., Ed., Zhang, L., Ed., and K. Fall, Ed., "Report from the IAB Workshop on Routing and Addressing", RFC 4984, September 2007.
- [RFC5570] StJohns, M., Atkinson, R., and G. Thomas, "Common Architecture Label IPv6 Security Option (CALIPSO)", RFC 5570, July 2009.
- [RFC6742] Atkinson, R., Bhatti, S. and S. Rose, "DNS Resource Records for the Identifier-Locator Network Protocol (ILNP)", RFC 6742, November 2012.
- [RFC6745] Atkinson, R. and S. Bhatti, "ICMP Locator Update Message for the Identifier-Locator Network Protocol for IPv4 (ILNPv4)", RFC 6745, November 2012.
- [RFC6746] Atkinson, R. and S. Bhatti, "IPv4 Options for the Identifier-Locator Network Protocol (ILNP)", RFC 6746, November 2012.
- [RFC6747] Atkinson, R. and S. Bhatti, "Address Resolution Protocol (ARP) Extension for the Identifier-Locator Network Protocol for IPv4 (ILNPv4)", RFC 6747, November 2012.
- [RFC6748] Atkinson, R. and S. Bhatti, "Optional Advanced Deployment Scenarios for the Identifier-Locator Network Protocol (ILNP)", RFC 6748, November 2012.

10. Acknowledgements

Steve Blake, Stephane Bortzmeyer, Mohamed Boucadair, Noel Chiappa, Wes George, Steve Hailes, Joel Halpern, Mark Handley, Volker Hilt, Paul Jakma, Dae-Young Kim, Tony Li, Yakov Rehkter, Bruce Simpson, Robin Whittle, and John Wroclawski (in alphabetical order) provided review and feedback on earlier versions of this document. Steve Blake provided an especially thorough review of an early version of the entire ILNP document set, which was extremely helpful. We also wish to thank the anonymous reviewers of the various ILNP papers for their feedback.

Roy Arends provided expert guidance on technical and procedural aspects of DNS issues.

Authors' Addresses

RJ Atkinson
Consultant
San Jose, CA 95125
USA

EEmail: rja.lists@gmail.com

SN Bhatti
School of Computer Science
University of St Andrews
North Haugh, St Andrews
Fife KY16 9SX
Scotland, UK

EEmail: saleem@cs.st-andrews.ac.uk