

Internet Research Task Force (IRTF)
Request for Comments: 6746
Category: Experimental
ISSN: 2070-1721

RJ Atkinson
Consultant
SN Bhatti
U. St Andrews
November 2012

IPv4 Options for the Identifier-Locator Network Protocol (ILNP)

Abstract

This document defines two new IPv4 Options that are used only with the Identifier-Locator Network Protocol for IPv4 (ILNPv4). ILNP is an experimental, evolutionary enhancement to IP. 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/rfc6746>.

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. Introduction	2
1.1. Document Roadmap	3
1.2. Terminology	4
2. IPv4 Options for ILNPv4	4
2.1. ILNPv4 Packet Format	5
2.2. ILNP Identifier Option for IPv4	7
2.3. ILNP Nonce Option for IPv4	8
3. Security Considerations	8
4. IANA Considerations	9
5. References	9
5.1. Normative References	9
5.2. Informative References	10
6. Acknowledgements	11

1. Introduction

This document is part of the ILNP document set, and it 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 tunnelled 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.

The Identifier-Locator Network Protocol (ILNP) is a proposal for evolving the Internet Architecture. It differs from the current Internet Architecture primarily by deprecating the concept of an IP Address and instead defining two new objects, each having crisp syntax and semantics. The first new object is the Locator, a topology-dependent name for a subnetwork. The other new object is the Identifier, which provides a topology-independent name for a node.

1.1. Document Roadmap

This document describes a new IPv4 Nonce Option used by ILNPv4 nodes to carry a security nonce to prevent off-path attacks against ILNP ICMP messages and defines a new IPv4 Identifier Option used by ILNPv4 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) [RFC6744] 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.
- f) [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.
- 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. IPv4 Options for ILNPv4

ILNP for IPv4 (ILNPv4) is merely a different instantiation of the ILNP architecture, so it retains the crisp distinction between the Locator and the Identifier. As with ILNP for IPv6 (ILNPv6), when ILNPv4 is used for a network-layer session, the upper-layer protocols (e.g., TCP/UDP pseudo-header checksum, IPsec Security Association) bind only to the Identifiers, never to the Locators. As with ILNPv6, only the Locator values are used for routing and forwarding ILNPv4 packets.

However, just as the packet format for IPv4 is different from IPv6, so the engineering details for ILNPv4 are different also. Just as ILNPv6 is carefully engineered to be backwards-compatible with IPv6, ILNPv4 is carefully engineered to be backwards-compatible with IPv4.

Each of these options MUST be copied upon fragmentation. Each of these options is used for control, so uses Option Class 0.

Originally, these two options were specified to use separate IP option numbers. However, only one IP Option (decimal 158) has been defined for experimental use with properties of MUST COPY and CONTROL [RFC4727]. So these two options have been reworked to share that same IP Option number (158). To distinguish between the two actual options, the unsigned 8-bit field ILNPv4_OPT inside this option is examined.

It is important for implementers to understand that IP Option 158 is not uniquely allocated to ILNPv4. Other IPv4-related experiments might be using that IP Option value for different IP options having different IP Option formats.

2.1. ILNPv4 Packet Format

The Source IP Address in the IPv4 header becomes the Source ILNPv4 Locator value, while the Destination IP Address of the IPv4 header becomes the Destination ILNPv4 Locator value. Of course, backwards compatibility requirements mean that ILNPv4 Locators use the same number space as IPv4 routing prefixes.

ILNPv4 uses the same 64-bit Identifier, with the same modified EUI-64 syntax, as ILNPv6. Because the IPv4 address fields are much smaller than the IPv6 address fields, ILNPv4 cannot carry the Identifier values in the fixed portion of the IPv4 header. The obvious two ways to carry the ILNP Identifier with ILNPv4 are either as an IPv4 Option or as an IPv6-style Extension Header placed after the IPv4 header and before the upper-layer protocol (e.g., OSPF, TCP, UDP, SCTP).

Currently deployed IPv4 routers from multiple router vendors use packet forwarding silicon that is able to parse past IPv4 Options to examine the upper-layer protocol header at wire-speed on reasonably fast (e.g., 1 Gbps or better) network interfaces. By contrast, no existing IPv4-capable packet forwarding silicon is able to parse past a new Extension Header for IPv4. Hence, for engineering reasons, ILNPv4 uses a new IPv4 Option to carry the Identifier values. Another new IPv4 Option also carries a nonce value, performing the same function for ILNPv4 as the IPv6 Nonce Destination Option [RFC6744] performs for ILNPv6.

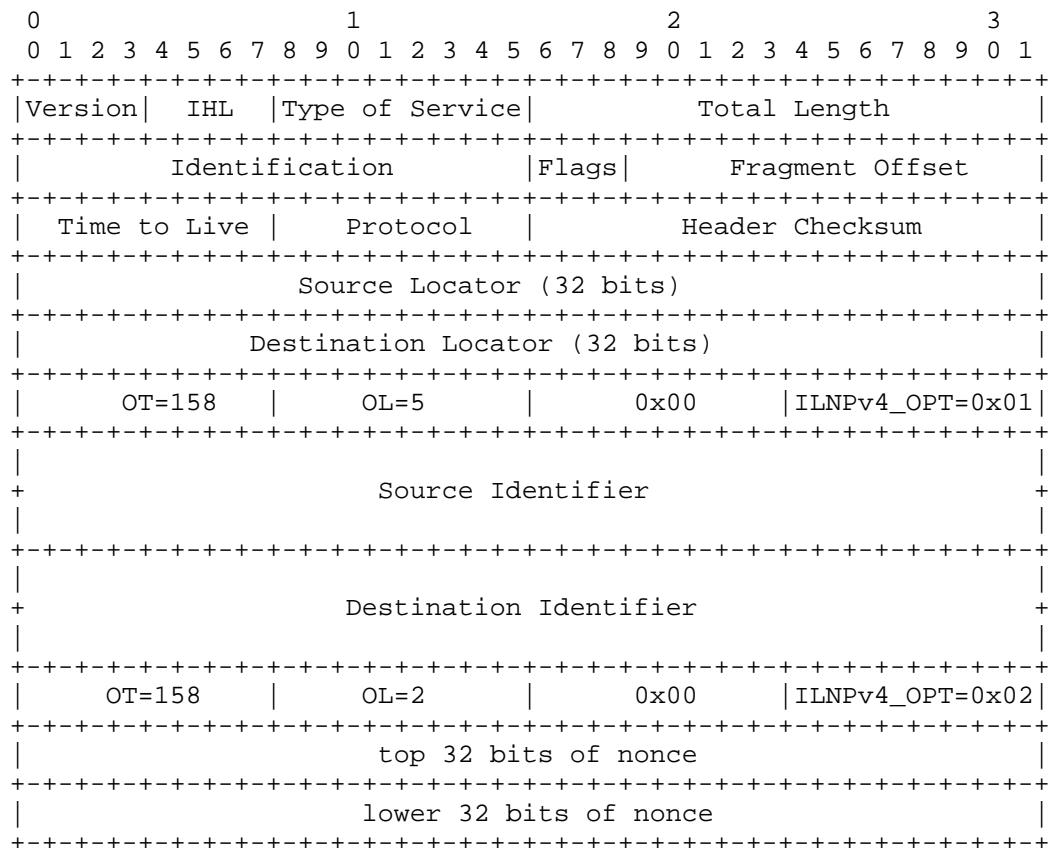


Figure 1: ILNPv4 Header with ILNP ID Option and ILNP Nonce Option

Notation for Figure 1:

IHL: Internet Header Length
 OT: Option Type
 OL: Option Length

2.2. ILNP Identifier Option for IPv4

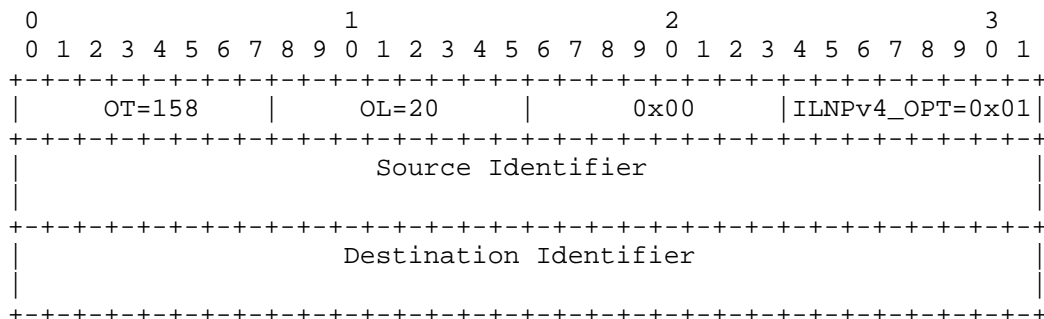


Figure 2: ILNP Identifier Option for IPv4

Notation for Figure 2:

OT: Option Type

OL: Option Length

RFC 791, Page 15 specifies that the Option Length is measured in words and includes the Option Type octet, the Option Length octet, and the option data octets.

The Source Identifier and Destination Identifier are unsigned 64-bit integers. [RFC6741] specifies the syntax, semantics, and generation of ILNP Identifier values. Using the same syntax and semantics for all instantiations of ILNP Identifiers simplifies specification and implementation, while also facilitating translation or transition between ILNPv4 and ILNPv6 should that be desirable in future.

This IP Option MUST NOT be present in an IPv4 packet unless the packet is part of an ILNPv4 session. ILNPv4 sessions MUST include this option in the first few packets of each ILNPv4 session and MAY include this option in all packets of the ILNPv4 session. It is RECOMMENDED to include this option in all packets of the ILNPv4 session if packet loss is higher than normal.

2.3. ILNP Nonce Option for IPv4

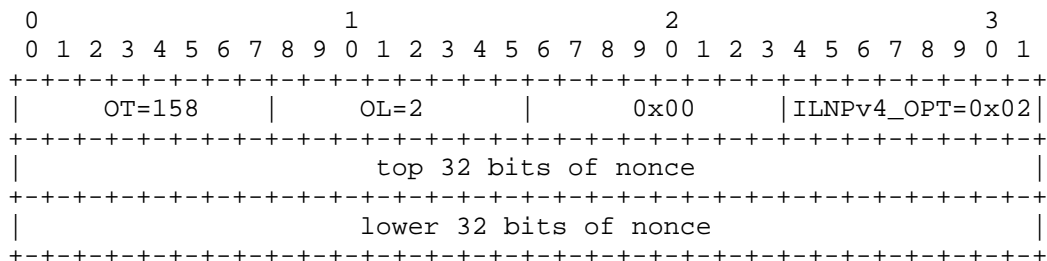


Figure 3: ILNP Nonce Option for IPv4

Notation for Figure 3:

OT: Option Type
OL: Option Length

This option contains a 64-bit ILNP Nonce. As noted in [RFC6740] and [RFC6741], all ILNP Nonce values are unidirectional. This means, for example, that when TCP is in use, the underlying ILNPv4 session will have two different NONCE values: one from Initiator to Responder and another from Responder to Initiator. The ILNP Nonce is used to provide non-cryptographic protection against off-path attacks (e.g., forged ICMP messages from the remote end of a TCP session).

Each NONCE value MUST be unpredictable (i.e., cryptographically random). Guidance to implementers on generating cryptographically random values is provided in [RFC4086].

This IP Option MUST NOT be present in an IPv4 packet unless the packet is part of an ILNPv4 session. ILNPv4 nodes MUST include this option in the first few packets of each ILNP session, MUST include this option in all ICMP messages generated by endpoints participating in an ILNP session, and MAY include this option in all packets of an ILNPv4 session.

3. Security Considerations

Security considerations for the overall ILNP Architecture are described in [RFC6740]. Additional common security considerations are described in [RFC6741]. This section describes security considerations specific to ILNPv4 topics discussed in this document.

If the ILNP Nonce value is predictable, then an off-path attacker might be able to forge data or control packets. This risk also is mitigated by the existing common practice of IP Source Address filtering [RFC2827] [RFC3704].

IP Security for ILNP [RFC6741] [RFC4301] provides cryptographic protection for ILNP data and control packets. The ILNP Nonce Option is required in the circumstances described in Section 3, even if IPsec is also in use. Deployments of ILNPv4 in high-threat environments SHOULD use IPsec for additional risk reduction.

This option is intended to be used primarily end-to-end between a source node and a destination node. However, unlike IPv6, IPv4 does not specify a method to distinguish between options with hop-by-hop behaviour versus end-to-end behaviour.

[FILTERING] provides general discussion of potential operational issues with IPv4 options, along with specific advice for handling several specific IPv4 options. 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 (D)DoS attack vector on backbone routers.

4. IANA Considerations

This document makes no request of IANA.

If in the future the IETF decided to standardise ILNPv4, then allocation of two unique Header Option values to ILNPv4, one for the Identifier option and one for the Nonce option, would be sensible.

5. References

5.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.
- [RFC4727] Fenner, B., "Experimental Values In IPv4, IPv6, ICMPv4, ICMPv6, UDP, and TCP Headers", RFC 4727, November 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.
- [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 Version 4 (ILNPv4)", RFC 6745, November 2012.
- [RFC6747] Atkinson, R. and S. Bhatti, "Address Resolution Protocol (ARP) Extension for the Identifier-Locator Network Protocol Version 4 (ILNPv4)", RFC 6747, November 2012.

5.2. Informative References

- [FILTERING] Gont, F., Atkinson, R., and C. Pignataro, "Recommendations on filtering of IPv4 packets containing IPv4 options", Work in Progress, March 2012.
- [RFC2780] Bradner, S. and V. Paxson, "IANA Allocation Guidelines For Values In the Internet Protocol and Related Headers", BCP 37, RFC 2780, March 2000.
- [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 2000.
- [RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, March 2004.
- [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.
- [RFC6743] Atkinson, R. and S. Bhatti, "ICMP Locator Update Message for the Identifier-Locator Network Protocol Version 6 (ICMPv6)", RFC 6743, November 2012.
- [RFC6744] Atkinson, R. and S. Bhatti, "IPv6 Nonce Destination Option for the Identifier-Locator Network Protocol Version 6 (ILNPv6)", RFC 6744, November 2012.

[RFC6748] Atkinson, R. and S Bhatti, "Optional Advanced Deployment Scenarios for the Identifier-Locator Network Protocol (ILNP)", RFC 6748, November 2012.

6. 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

EMail: rja.lists@gmail.com

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

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

