

Internet Draft
[draft-irtf-rrg-ilnp-v4opts-02.txt](#)
Expires: 16 OCT 2012
Category: Experimental

RJ Atkinson
Consultant
SN Bhatti
U. St Andrews
April 16, 2012

IPv4 Options for ILNPv4
[draft-irtf-rrg-ilnp-v4opts-02.txt](#)

Status of this Memo

Distribution of this memo is unlimited.

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. Code Components extracted from this document must include Simplified BSD License text as described in [Section 4.e](#) of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other

documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/1id-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>

This document is not on the IETF standards-track and does not specify any level of standard. This document merely provides information for the Internet community.

This document is part of the ILNP document set, and has had extensive review within the IRTF Routing Research Group. 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.

Abstract

This document defines 2 new IPv4 options that are used only with ILNP for IPv4 (ILNPv4). ILNP is an experimental, evolutionary enhancement to IP. This document is a product of the IRTF Routing RG.

Table of Contents - ### to be updated

1. Introduction.....
2. IPv4 Options for ILNPv4.....
3. Security Considerations.....
4. IANA Considerations.....
5. References.....

1. INTRODUCTION

The Identifier Locator Network Protocol (ILNP) is an 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 ILNP Document Roadmap

The ILNP Architecture document [[ILNP-ARCH](#)] is the best place to start reading about ILNP. ILNP has multiple instantiations. [[ILNP-ENG](#)] discusses engineering and implementation aspects common to all instances of ILNP. This document discusses engineering and implementation details that are specific to ILNP for IPv4 (ILNPv4). [[ILNP-DNS](#)] describes new Domain Name System (DNS) resource records used with ILNP. [[ILNP-ICMPv4](#)] defines the ICMP Locator Update message used with ILNPv4. Other documents describe ILNP for IPv6 (ILNPv6) [[ILNP-ICMPv6](#)] [[ILNP-NONCE6](#)].

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](#). [[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 to 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.

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 [[ILNP-NONCE6](#)] performs for ILNPv6.

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|Version|IHL=12|Type of Service|                Total Length                |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                Identification                |Flags|      Fragment Offset      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|  Time to Live  |   Protocol   |                Header Checksum                |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                Source Locator (32 bits)                |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                Destination Locator (32 bits)                |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| OT=ILNPv4_ID  |   OL=5       |      Padding=0x0000      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                |                |
+                Source Identifier                +
|                |                |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                |                |
+                Destination Identifier                +
|                |                |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|OT=ILNPv4_NONCE|   OL=2       |      top 16 bits of nonce      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                |                |
+                lower 32 bits of nonce                +
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Figure 1: ILNPv4 header with ILNP ID option

Atkinson & Bhatti Expires in 6 months

[Page 4]

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

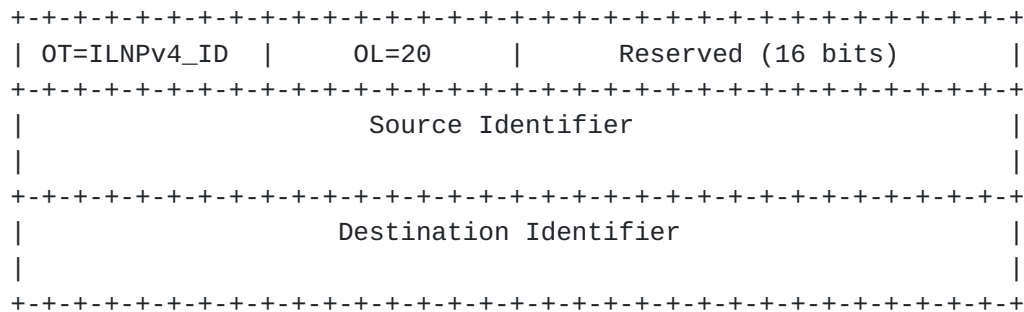


Figure 2: ILNP Identifier Option for IPv4

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 Reserved field is reserved for possible future use. For now, senders MUST initialise this field to all zero values, and receivers MUST ignore the contents of this field.

The Source Identifier and Destination Identifier are unsigned 64-bit integers. [\[ILNP-ENG\]](#) 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 session, and MAY include this option in all packets of the session. It is RECOMMENDED to include this option in all packets of the session if packet loss higher than normal.

2.3 ILNP Nonce Option for IPv4

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|OT=ILNPv4_NONCE|   OL=8   |   NONCE (upper 16 bits)   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|
|   NONCE (lower 32 bits)
|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Figure 3: ILNP Nonce Option for IPv4

Notation for Figure 3:

OT: Option Type
OL: Option Length

This option contains a 48-bit ILNP Nonce. As noted in [ILNP-ARCH] and [ILNP-ENG], all ILNP Nonce values are unidirectional. This means, for example, that a typical TCP/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 [ILNP-ARCH]. Additional common security considerations are described in [ILNP-ENG]. 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 [ILNP-ENG] [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 IP Security is also in use. Deployments of ILNPv4 in high-threat environments SHOULD use IP Security 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.

[ID-IPv4-OPT-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

Subject to IESG Approval, consistent with the procedures of [\[RFC2780\]](#), IANA is requested to assign new IPv4 option values for the ILNPv4 Identifier option (ILNPv4_ID) and the ILNPv4 Nonce option (ILNPv4_NONCE).

The short name for ILNPv4_ID is "ID", for consistency with [\[ILNP-DNS\]](#), and the short name for ILNPv4_NONCE is "INONCE".

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

5. REFERENCES

This document has both Normative and Informational References.

5.1 Normative References

- [ILNP-ARCH] R. Atkinson and S. Bhatti, "ILNP Architecture", [draft-irtf-rrg-ilnp-arch](#), January 2012.
- [ILNP-ENG] R. Atkinson and S. Bhatti, "ILNP Engineering Considerations", [draft-irtf-rrg-ilnp-eng](#), January 2012.
- [ILNP-DNS] R. Atkinson and S. Bhatti, "DNS Resource Records for ILNP", [draft-irtf-rrg-ilnp-dns](#), January 2012.

- [ILNP-ICMPv4] R. Atkinson and S. Bhatti, "ICMP Locator Update message for ILNPv4", [draft-irtf-rrg-ilnp-icmpv4](#), January 2012.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC4301] S. Kent and K. Seo, "Security Architecture for the Internet Protocol", [RFC-4301](#), December 2005.

5.2 Informative References

- [ID-IPv4-OPT-FILTERING] F. Gont, R. Atkinson, and C. Pignatero, "Recommendations on Filtering of IPv4 Packets with IPv4 options", [draft-gont-opsec-ip-options-filtering](#), March 2012.
- [ILNP-NONCE6] R. Atkinson and S. Bhatti, "ILNPv6 Nonce Destination Option", [draft-irtf-rrg-ilnp-noncev6](#), January 2012.
- [ILNP-ICMPv6] R. Atkinson and S. Bhatti, "ICMPv6 Locator Update Message for ILNPv6", [draft-irtf-rrg-ilnp-icmpv6](#), January 2012.
- [RFC2780] S. Bradner & V. Paxson, "IANA Allocation Guidelines for Values in the Internet Protocol and Related Headers", [RFC 2780](#), March 2000.
- [RFC2827] P. Ferguson and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", [RFC-2827](#), May 2000.
- [RFC3704] F. Baker and P. Savola, "Ingress Filtering for Multihomed Networks", [RFC-3704](#), March 2004.
- [RFC4086] D. Eastlake 3rd, J. Schiller, and S. Crocker, "Randomness Requirements for Security", [RFC-4086](#), June 2005.

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.

RFC EDITOR NOTE

This section is to be removed prior to publication.

This document is written in English, not American. So English spelling is used throughout, rather than American spelling. This is consistent with existing practice in several other RFCs, for example [RFC-5887](#).

This document tries to be very careful with history, in the interest of correctly crediting ideas to their earliest identifiable author(s). So in several places the first published RFC about a topic is cited rather than the most recent published RFC about that topic.

AUTHOR'S ADDRESS

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

Expires: 16 OCT 2012

