

Internet Draft
draft-irtf-rrg-ilnp-noncev6-00.txt
Expires: 09 JUL 2012
Category: Experimental

RJ Atkinson
Consultant
S Bhatti
U. St Andrews
9 January 2012

IPv6 Nonce Destination Option for ILNPv6
draft-irtf-rrg-ilnp-noncev6-00.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. This document may not be modified, and derivative works of it may not be created, except to publish it 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/lid-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, which 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

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 RG.

Table of Contents

- 1. Introduction2
- 2. Syntax.....3
- 3. Transport Protocol Effects.....5
- 4. Location Changes.....5
- 5. Implementation Considerations.....6
- 6. Backwards Compatibility.....6
- 7. Security Considerations8
- 8. IANA Considerations9

9. References9

1. Introduction

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 an Identifier/Locator split is in use. Second, it provides a signal during initial IP session creation that the Identifier/ Locator Split operating mode is proposed for use with this session. This last function is particularly important for ensuring that the new Identifier/Locator Split operating mode is both incrementally deployable and backwards compatible with IPv6. Consequently, this option must not be used except by a node operating in the I/L Split (ILNP) mode.

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 a session's nonce. So a typical TCP session will have 2 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

The Identifier-Locator Network Protocol (ILNP), is described in the ILNP Architecture [ILNP-ARCH] document, which should be read first. ILNP can have multiple instantiations. [ILNP-ENG] discusses engineering and implementation aspects common to all ILNP instantiations. [ILNP-DNS] defines new Domain Name System (DNS) resource records for ILNP. [ILNP-ICMPv6] defines a new ICMPv6 Locator Update message for use with ILNPv6. Other documents describe ILNP for IPv4 (ILNPv4).

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 RFC 2119. [RFC 2119]

2. Syntax


The Nonce Option is carried within an IPv6 Destination Option Header. Section 4 of [RFC 2460] provides much more information on the various options and optional headers used with IPv6.

More than one option might be inside the IPv6 Destination Option Header, however at most 1 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 Option 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.

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



Next Header: 8-bit selector. Identifies the type of header immediately following the Destination Options header. Uses the same values as the IPv4 Protocol field [RFC 1700 et seq].

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 XXX, which is used to indicate the start of the Nonce Option.

- 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 used to prevent off-path attacks on an ILNP session. [RFC 4086] This field has variable length, with the length indicated by the Option Length field preceding it. Note that the overall 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, then ILNPv6 is in use for that communications session.

When a communications session is using ILNPv6, then the transport-layer pseudo-header calculations MUST 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 the session.

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

4. Location Changes

When a node has an unexpected 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.

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 ILNP Correspondent Cache's Set of Correspondent Locator(s) 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 session hijacking or session interference attacks.

Hence, the node that unexpectedly 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 3 round-trip times for each correspondent. (NB: 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 a session is using IP Security, then, of course, IP Security SHOULD continue to be used. Because IP Security for ILNP [[ILNP-ENG](#)] 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 multi-homing are functionally equivalent, 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 Correspondent Cache

As described in [[ILNP-ENG](#)], ILNP nodes maintain an ILNP Correspondent 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 IP and which IP sessions are using ILNP.

If a given transport-layer session is using ILNP, then an entry

corresponding to that session also will exist in the ILNP Correspondent Cache. Multiple transport-layer sessions between a given pair of nodes MAY share a single entry in the Correspondent Cache if they are similar except for details above the network-layer.

5.3 IP Security

Whether or not ILNP is in use, the IPsec subsystem MUST maintain an IPsec Security Association Database (SAD) and also MUST maintain information about which IPsec Selectors apply to traffic received by or sent from the local node. [RFC 4301] By combining the information in the IPsec SAD, of what IPsec Selectors apply, and information in the ILNP Correspondent Cache, an implementation has sufficient knowledge to apply IPsec properly to both received and transmitted packets.

6. Backwards Compatibility

If a node supports ILNP and the node wishes to be able to receive incoming new ILNP sessions, then that node's fully-qualified domain name SHOULD have one or more ID records and also one or more Locator (i.e. L64 or LP) records associated with it in the DNS.

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 Identifier ("ID") and Locator ("L64") records in any received DNS replies. DNS servers that support ID and Locator (i.e., L64 or LP) records SHOULD 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 Fully-Qualified Domain Name (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 nonce value, MUST store that value in the local correspondent cache, 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. [RFC 4086]

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. [[RFC 4443](#)]

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 also does not receive an ICMP Unreachable error message for that packet, OR if the initiator 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 IP Security is in use. [[CA-1995-01](#)] [[RFC 4301](#)] 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 for a communications session, 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 IP Security also is in use for that session. Note that certain ICMP messages, for example a "Path Too Big" message, might be generated by transit devices that are not aware of the ILNP Nonce in use for that session and hence are not able to include the ILNP Nonce. Again, this also is 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 [ILNP-ENG] [RFC 4301]. 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 [ILNP-ICMPv6].

As a performance optimisation, it is suggested that when both the Nonce Option and IP Security are present in a packet and the Nonce Option has not been encrypted, that the Nonce Option value be checked for validity before beginning IP Security 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 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 a session at a different sensitivity level [RFC 5570]. Further, multi-level secure systems SHOULD use different nonce values for sessions with different Sensitivity Levels [RFC 5570].

In all cases, the Nonce Value MUST be unpredictable and cryptographically random. [RFC 4086] provides concrete advice on how to generate a suitable nonce value.

As this is an option within the IPv6 Destination Option Header,

Atkinson & Bhatti Expires in 6 Months

[Page 9]

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 (D)DOS attack vector on backbone routers.

8. IANA Considerations

IANA is requested to assign a new IPv6 Destination Option Type value (replacing XXX, in [Section 2](#) above).

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

Further, if an end system receives a packet containing this option, but does not recognise the option, the end system MUST drop the received packet. In that case, if and only if the Destination Address was NOT an IPv6 multicast address the receiving end system MUST send an ICMP Parameter Problem, Code 2, message to the packet's Source Address, pointing to the unrecognised Option Type.

9. References

9.1. 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 and Implementation 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-ICMPv6] R. Atkinson & S. Bhatti, "ICMP Locator Update message for ILNPv6", [draft-irtf-rrg-ilnp-icmpv6](#), January 2012.
- [RFC 2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#),

March 1997.

- [RFC 2460] S. Deering & R. Hinden, "Internet Protocol Version 6 Specification", [RFC 2460](#), December 1998.
- [RFC 4301] S. Kent & K. Seo, "Security Architecture for the Internet Protocol", [RFC 4301](#), December 2005.
- [RFC 4443] A. Conta, S. Deering, M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for IPv6 Specification", [RFC 4443](#), March 2006.

9.2. Informative References

- [8+8] M. O'Dell, "8+8 - An Alternate Addressing Architecture for IPv6", Internet-Draft, October 1996.
- [GSE] M. O'Dell, "GSE - An Alternate Addressing Architecture for IPv6", Internet-Draft, February 1997.
- [CA-1995-01] US CERT, "CERT Advisory 1995-01", Pittsburgh, PA, USA, 1995.
- [RFC 4086] D. Eastlake 3rd, J. Schiller, & S. Crocker, "Randomness Requirements for Security", [RFC 4086](#), June 2005.
- [RFC 5570] M. StJohns, R. Atkinson, and G. Thomas, "Common Architecture Label IPv6 Security Option (CALIPSO)", [RFC-5570](#), July 2009.

ACKNOWLEDGEMENTS

Steve Blake, Noel Chiappa, Steve Hailes, Joel Halpern, Mark Handley, Volker Hilt, Tony Li, and Yakov Rehkter (in alphabetical order) provided review and feedback on earlier versions of ILNP documents. Steve Blake provided an especially thorough review of the entire ILNP document set, which led to significant improvements in this document.

Authors' Addresses:

RJ Atkinson
Consultant
San Jose, CA

95125 USA

rja.lists@gmail.com

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

Expires: 09 JUL 2012