IP Version 6	j h. woodyatt
Internet-Draft	Apple
Intended status: Standards Track	July 31, 2008
Expires: February 1, 2009	

TOC

Application Listener Discovery (ALD) for IPv6
draft-woodyatt-ald-03

Status of this Memo

Abstract

This document specifies the protocol used by IPv6 nodes comprising stateful packet filters to discover the transport addresses of listening applications (that is, application endpoints for which incoming traffic may be administratively prohibited).
Comments are solicited and should be sent to the author and the V6OPS Residential CPE Design Team mailing list at <v6ops-residential-cpe-design-team@external.cisco.com>.

Table of Contents

---

1. INTRODUCTION

TOC

In "Local Network Protection for IPv6" [RFC4864] (Van de Velde, G., Hain, T., Droms, R., Carpenter, B., and E. Klein, “Local Network Protection for IPv6,” May 2007.), IETF recommends 'simple security' capabilities for residential and small office gateways that prohibit, by default, all inbound traffic except those packets returning as part of locally initiated outbound flows. It further recommends "an easy interface which allows users to create inbound 'pinholes' for specific purposes such as online gaming."

In existing IPv4 gateways, where Network Address Translation (NAT) is commonly used for IPv4 network protection and firewalling, management applications typically provide an interface for manual configuration of pinholes. However, this method is unacceptably difficult for many non-technical Internet users, so most products in the market today also implement one or more automatic methods for creating pinholes.
These methods include:

*"NAT Port Mapping Protocol" [NAT-PMP] (Cheshire, S., Krochmal, M., and K. Sekar, "NAT Port Mapping Protocol (NAT-PMP)," November 2001.)

*"Internet Gateway Device (IGD)" standardized device control protocol of Universal Plug And Play [UPnP-IGD] (UPnP Forum, "Universal Plug and Play Internet Gateway Device Standardized Gateway Device Protocol," September 2006.)

The basic mechanism of these protocols is that applications notify the firewall of their expectation to receive inbound flows, and pinholes are opened accordingly. In the IPv4/NAT case, these protocols are also used for automatic creation of network address translator state in addition to packet filter state. In the IPv6 case, no network address translation is necessary, but packet filters still contain state and pinholes must still be created accordingly.
At present, no similar protocol exists for automatically notifying firewalls of the pinholes required by IPv6 endpoint applications. This document defines a method for making such notifications.
(NOTE: It is expected that this section will be revised once the concept presented in this document is well socialized in the Internet engineering and operations community.)

---

2. TERMINOLOGY

TOC

2.1. Requirements Language

TOC

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] (Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," March 1997.).
Paragraphs that begin with "EXPERIMENTAL:" describe how this protocol may be implemented using numbers assigned by IANA for experimental usage. Prior to publication of this document as a Request For Comments,

the RFC Editor is directed to delete all paragraphs that begin with this tag and all references to "Experimental Values in IPv4, IPv6, ICMPv4, ICMPv6, UDP, and TCP Headers" (Fenner, B., “Experimental Values In IPv4, IPv6, ICMPv4, ICMPv6, UDP, and TCP Headers,” November 2006.) [RFC4727].

2.2. Special Terms and Abbreviations

TOC

firewall: A node with the capability of administratively prohibiting the flow of packets between a protected "interior" region of the Internet and an "exterior" region.

flow initiation: The start of communications between two or more nodes in an application protocol, e.g. the TCP SYN packets that comprise the start of a telnet session, the UDP packets that start an NTP exchange, the first IPsec ESP packet for a new security parameter index (SPI), et cetera.

---

3. PROTOCOL OVERVIEW

TOC

This protocol solves a set of problems related to the interaction between applications awaiting reception of transport flow initiations (listeners) and IPv6 nodes comprising packet filtering network policy enforcement points (firewalls).
From the perspective of any given IPv6 node, the region of the Internet between itself and a given firewall is the 'interior' domain of that firewall. All other regions of the Internet are the 'exterior' from the perspective of the node. The ALD protocol is concerned only with the problems associated with listeners on nodes reachable only on the interior interfaces of firewalls in receiving transport flow initiations from nodes reachable only on exterior interfaces.
The ALD protocol defines methods for solving each of the following problems:

Listener Discovery:
How firewalls discover the transport protocols and addresses of applications awaiting reception of flow initiations.

Firewall Discovery: How nodes discover what firewalls to notify that applications are awaiting reception of transport flow initiations.

Firewall Reset Detection: How nodes discover that firewalls have been reset and now require nodes to restart their listener discovery functions.

Application Programming Interface: Extensions to the IPv6 API are defined to permit applications to be selective about how their transport endpoints are subjects of listener notification.

When nodes join network segments where one or more global scope address prefixes are advertised, they use a Firewall Discovery method to build or learn a list of firewalls to notify that applications are listening at specific unicast addresses. They send Firewall Solicitation messages to a specified destination address, which may be a multicast destination, and receive directed Firewall Advertisement messages in response.
Nodes send Listener Notification messages to firewalls to inform them of their expectations in receiving flow initiations. These messages are sent for each listener endpoint address in use, with retransmits as necessary. Firewalls send Listener Acknowledgment messages to squelch further retransmits.
It's important to recognize the notifications are not requests. Firewalls are under no obligation to change their behavior in response to receiving application listener notifications. Nodes are provided with no assurance that inbound flow initiations are or are not prohibited at firewalls in the network, whether advertised with ALD or not.
Every ALD message sent by a firewall includes a measurement of the elapsed time since their state was last reset. This is so nodes may recognize when it may be necessary to resend all its listener notifications. Firewalls periodically send announcements, but in general not at a frequency high enough that nodes may rely on the absence of them to detect the failure of a firewall.

3.1. Firewall Discovery

TOC

For the purposes of application listener discovery, firewalls have an "interior" subject to the policy requiring listeners to notify them,

and an "exterior" corresponding to the region of the Internet from which flow initiations are subject to administrative prohibitions. Nodes transmit Firewall Solicitation messages and receive Firewall Advertisement messages in acknowledgment. Firewall Advertisement messages inform nodes of firewalls that may prohibit flow initiations from exterior sources to the node.
A new neighbor discovery option is defined for use in Router Advertisements to specify the destination address and hop limit that nodes are expected to use when sending Firewall Solitation messages.

3.2. Listener Discovery

TOC

Nodes send Listener Notification messages to firewalls according to their policy requirements. These notifications inform firewalls of which nodes, protocols, and transport addresses are expecting to receive inbound flow initiations. Firewalls send Listener Acknowledgment messages in response to inform listeners how much time the application can expect receive flow initiations.
Nodes may notify firewalls that they expect to receive all inbound traffic, regardless of protocol or transport address. Alternatively, they can send notifications for narrower constraints on what to pass through to listening nodes.

3.3. Firewall Reset Detection

TOC

Firewalls periodically multicast Firewall Advertisement messages on their "interior" interfaces. Immediately after the state in a firewall resets, the transmit interval for these advertisements are very short, rapidly increasing thereafter.
Nodes receive Firewall Advertisements directly and compare the Elapsed Time Since Reset (ETSR) against the last value received in any previous message. Computing their own conservative estimates of the expected elapsed time, nodes are able to recognize when retransmitting their listener notifications might be necessary.

3.4. Application Programming Interface

TOC

Applications need not be written with specific awareness of listener discovery. Operating systems are implemented with default parameters suitable for all but the rarest of exceptions.

For example, nodes only inform firewalls about TCP sockets when they require transport address level notification and the node sets a TCP socket into the LISTENING state. Furthermore, the timing limits on notifications vary between temporary privacy addresses and permanently assigned addresses, i.e. a TCP socket bound to a temporary address will have a short binding time in the firewall compared to a TCP socket that binds to a permanent address.
Some extensions to the application programming interface are defined for those few applications that need them. These extensions allow applications to disable listener notification or override timing parameters on a case by case basis.

---

4. OPTION FORMATS

TOC

The need for nodes to proceed with firewall discovery is signaled by the presence of a Firewall Discovery option sent in Router Advertisement messages.

4.1. Firewall Discovery Router Advertisement Option

TOC

In Router Advertisements without the "other stateful configuration" flag set, the Firewall Discovery Option informs nodes of the destination address and hop limit for sending Firewall Solicitation messages.

Firewall Discovery Option

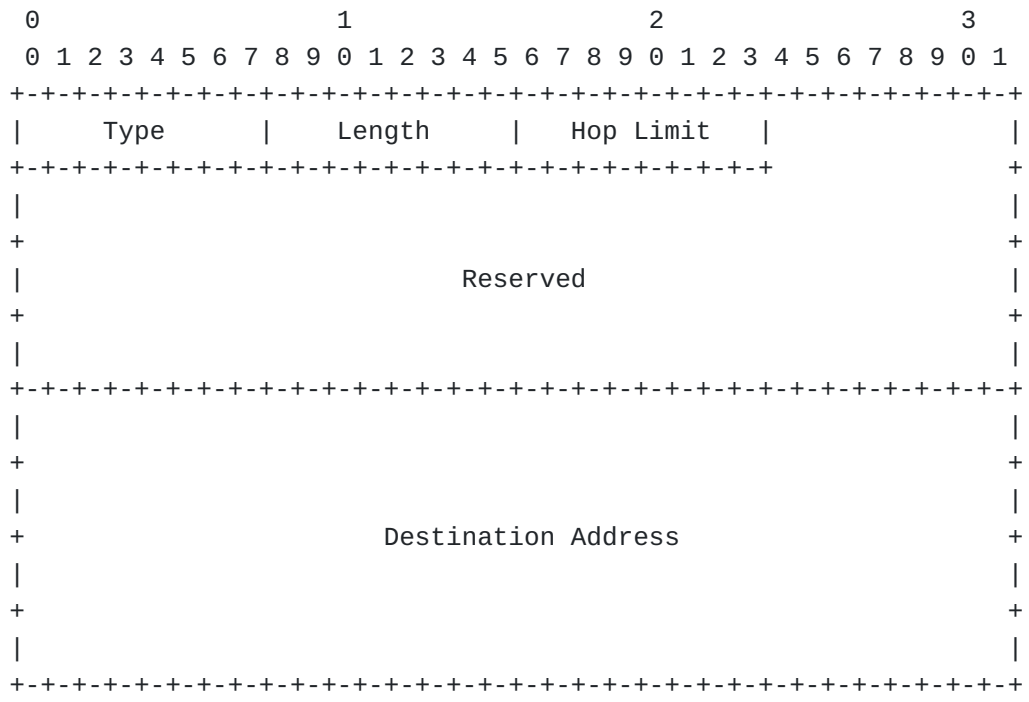

Type: TBD

Length: 4

Hop Limit: The hop limit nodes use to send Firewall Solicit messages.

Reserved: This field is unused. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

Destination Address: The destination address for nodes to use when sending Firewall Solicit messages.

Routers MUST NOT send Router Advertisements containing the Firewall Discovery option if the "other stateful configuration" flag is set. Likewise, nodes MUST NOT process the Firewall Discovery Option unless the "other stateful configuration" flag is set in the Router Advertisement that contains it.
Routers MUST NOT send Router Advertisements with more than one Firewall Discovery Option present. If nodes receive such Router Advertisements, then nodes MUST NOT process any of the Firewall Discovery Options.
Nodes that process Firewall Discovery Options in Router Advertisements MUST NOT send any Firewall Solicitation messages from any addresses in the advertised prefixes except to the specified destination address, and with the specified hop limit.
Nodes receiving Router Advertisements with the "other stateful configuration" flags not set, and without a Firewall Discovery Option

present, MAY send Firewall Solicitation messages from the advertised prefixes to any address and with any hop limit.
EXPERIMENTAL: The type value 253 is defined in section 5.1.3 of "Experimental Values in IPv4, IPv6, ICMPv4, ICMPv6, UDP, and TCP Headers" (Fenner, B., “Experimental Values In IPv4, IPv6, ICMPv4, ICMPv6, UDP, and TCP Headers,” November 2006.) [RFC4727] for use with experimental protocols. Operation of ALD in experimental mode requires the four octet code 0x6161706c be inserted between the Length and Hop Limit fields, and the size of the Reserved field to be reduced by four octets to keep the destination address aligned. Experimental Firewall Discovery Options, i.e. those described in this paragraph, MUST NOT be processed unless the type value is 253 and the four octet code is present in the required position.

5. MESSAGE FORMATS

TOC

ALD is a sub-protocol of ICMPv6, that is, ALD message types are a subset of the set of ICMPv6 messages, and ALD messages are all identified in IPv6 packets by a preceding Next Header value of 58. ALD messages all have the same Type value, [TBD, assigned by IANA], and their function is differentiated by the Code value.
This document defines the formats for ALD messages with the following Code values:

ALD Message Codes

Code	Description	Reference
1	Firewall Solicitation	Section 5.1 (Firewall Solicitation)
2	Firewall Advertisement	Section 5.2 (Firewall Advertisement)
3	Listener Notification	Section 5.4 (Listener Notification)
4	Listener Acknowledgment	Section 5.5 (Listener Acknowledgment)

Table 1

All other Code values are reserved for future use. Nodes MUST NOT send messages containing them.
Firewalls MUST NOT prohibit the flow of ALD messages from their exterior to their interior.

TOC

5.1. Firewall Solicitation

Nodes send Firewall Solicitation messages to request firewalls to respond with directed Firewall Advertisement messages. They are sent periodically to the destination addresses specified in any Firewall Discovery Options received in Router Advertisements for networks they join.

Firewall Solicitation

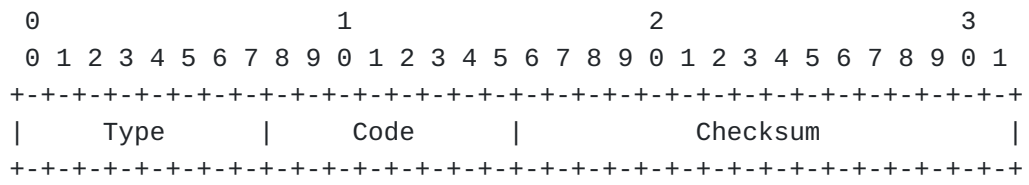

Type: TBD. Assigned by IANA to ALD messages.

Code: 1.

Checksum: ICMPv6 checksum.

EXPERIMENTAL: Nodes operating in experimental mode MAY send the Experimental Firewall Solicitation message, i.e. the same message except with type value 100 as defined in "Internet Control Message Protocol (ICMPv6)" (Conta, A., Deering, S., and M. Gupta, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification," March 2006.) [RFC4443] for use in experimental protocols, and the four octet code 0x6161706c appended after the checksum. Nodes MUST NOT send Experimental Firewall Solicitation messages to destination addresses received in the regular Firewall Discovery Option.

5.2. Firewall Advertisement

TOC

Firewalls send Firewall Advertisement messages to notify listeners reachable on their interior interfaces that inbound flow initiations to a specific prefix are subject to policy enforcement.

Firewalls Advertisement

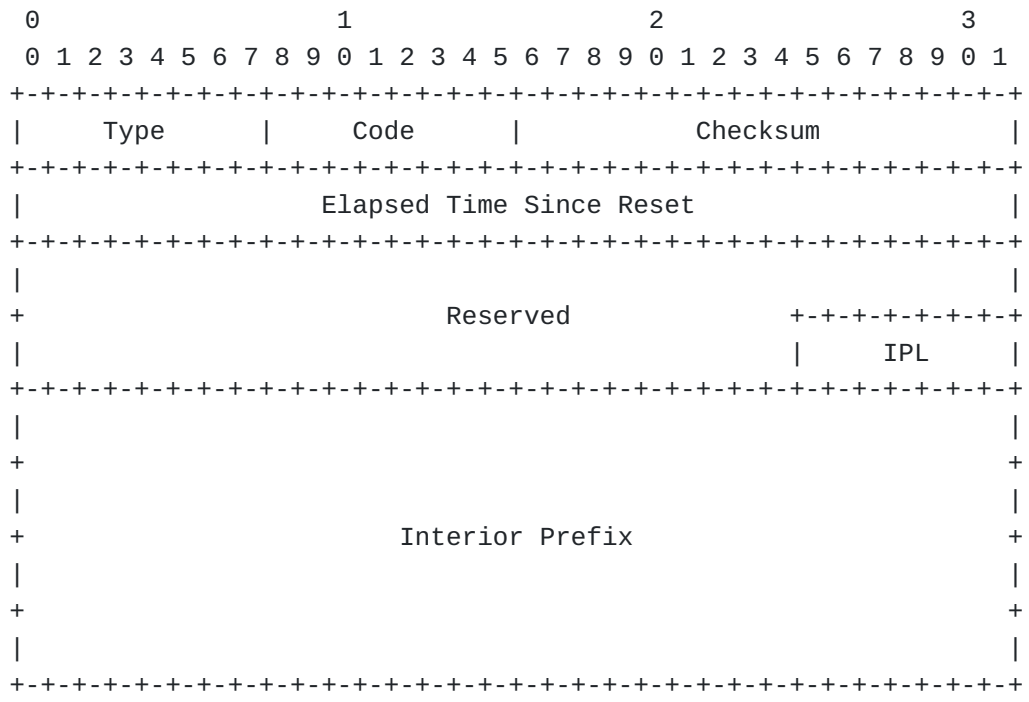

Type: TBD. Assigned by IANA to ALD messages.

Code: 2.

Checksum: ICMPv6 checksum.

Elapsed Time Since Reset: Number of elapsed seconds since the firewall state was last reset.

IPL: The length of the interior prefix. Values less than 48 are reserved. Senders MUST NOT use them, and receivers MUST NOT process any messages that contain them. (Note: the width of this field is seven bits.)

Reserved: This field is unused. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

Interior Prefix: The IPv6 address prefix on the interior subject to the firewall policy.

Starting when a firewall begins operating on the interior prefix from its reset state, it MUST periodically send Firewall Advertisement messages on all its interfaces where the interior prefix is reachable using a Hop Limit of 255 to the organizational scope All Nodes multicast address, FF08::1. The time interval between multicast transmissions MAY be of any duration. The recommended period is every two seconds for the first ten seconds after the state is reset,

followed by a doubling of the interval for every transmission thereafter until the interval reaches a maximum of one hour.
EXPERIMENTAL: Firewalls operating in experimental mode MAY send Experimental Firewall Advertisement messages, i.e. the same message except with type value 100 as defined in "Internet Control Message Protocol (ICMPv6)" (Conta, A., Deering, S., and M. Gupta, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification," March 2006.) [RFC4443] for use in experimental protocols and the four octet code 0x6161706c inserted between the Checksum and Elapsed Time Since Reset fields. These are sent to the organizational scope "any private experiment" multicast destination address, i.e. FF08::114, instead of the All Nodes address. Nodes MUST NOT send Experimental Firewall Advertisement messages to any other multicast destination.

5.3. Listener Address Specifier

TOC

Listener Notification and Listener Acknowledgment messages (see below) each contain Listener Address Specifier elements. These are structured data that describe the transport layer component of a listener address that firewalls are expected to filter, e.g. TCP and UDP ports, etc. As a general rule, this protocol number is expected to match the upper-layer-protocol of the outer-most IPv6 header (including all its extension headers). See "Internet Protocol, Version 6" (Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification," December 1998.) [RFC2460] for details.
The first octet of any Listener Address Specifier is an Internet protocol number, which serves as the type discriminator for a variant subtype of Listener Address Specifier elements.
Nodes MUST NOT send Listener Address Specifiers with protocol numbers assigned for identifying IPv6 extension headers.

5.3.1. All Protocols Listener Address Specifier

TOC

Nodes notify firewalls that inbound flow initiations are expected by sending a Listener Notification message with the All Protocols Listener Address Specifier. This is a single octet with all zero bits, followed by a reserved field of three octets.

All Protocols Listener Address Specifier

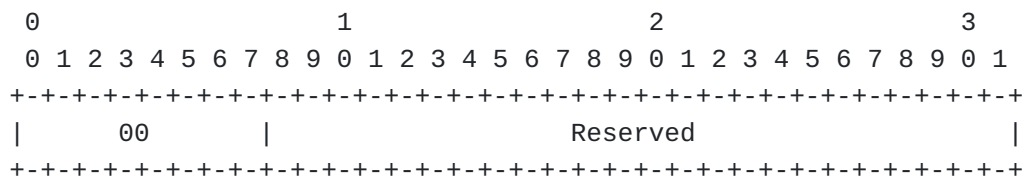

Reserved: This field is unused. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

Note: the value of zero is used here for specifying all protocols, even though it is used in IPv6 for specifying hop-by-hop options.

5.3.2. All Specific Protocol Listener Address Specifier

TOC

Nodes notify firewalls that all inbound flow initiations for a specific upper-layer protocol are expected by sending a Listener Notification message with an All Specific Protocol Listener Address Specifier. This is a single octet with the protocol number, followed by three octets of zeroes.

All Specific Protocol Listener Address Specifier

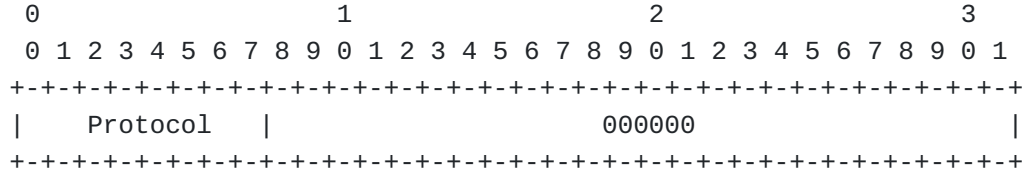

Protocol: The upper-layer protocol number.

Nodes MUST NOT send All Specific Protocol Listener Address Specifier elements with protocol numbers reserved for IPv6 header extensions in the Protocol field.
Nodes MUST NOT send All Specific Protocol Listener Address Specifier elements with 255 in the Protocol field.

5.3.3. Encapsulating Security Payload Listener Address Specifier

TOC

Nodes notify firewalls of that inbound IP Encapsulating Security Payload (ESP) flows (Kent, S., “IP Encapsulating Security Payload (ESP),” December 2005.) [RFC4303] are expected by sending a Listener

Notification message with the Encapsulating Security Payload Listener Address Specifier. This is a single octet with the ESP protocol number in it, followed by a reserved field of three octets.

Encapsulating Security Payload Listener Address Specifier

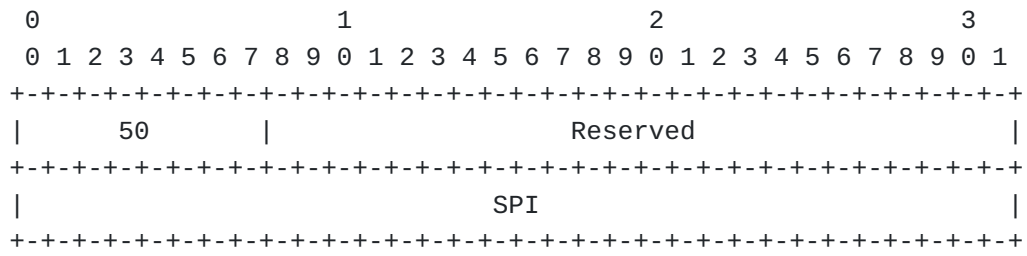

Reserved: This field is unused. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

SPI: Security Parameter Index for inbound flow.

An ESP Listener Address Specifier with a value of all zero octets in the SPI field is equivalent to the All Specific Protocol Listener Address Specifier with the ESP protocol number in the Protocol field.

5.3.4. TCP Listener Address Specifier

TOC

Nodes notify firewalls that inbound Transmission Control Protocol (TCP) connections (Postel, J., "Transmission Control Protocol," September 1981.) [RFC0793] are expected by sending a Listener Notification message with the TCP Listener Address Specifier. This is a single octet with the TCP protocol number in it, followed by a reserved octet, followed by the TCP port number for the application endpoint.

TCP Listener Address Specifier

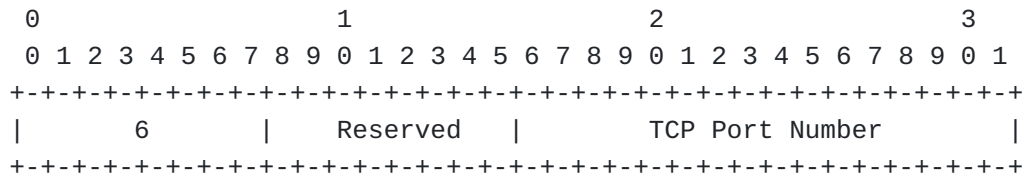

Reserved: This field is unused. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

TCP Port Number: The TCP port for the application endpoint.

A value of zero in the TCP Port Number field indicates all TCP flows. This is identical to the All Specific Protocol Listener Address Specifier for TCP.

5.3.5. UDP Listener Address Specifier

TOC

Nodes notify firewalls that inbound User Datagram Protocol (UDP) flow initiations (Postel, J., "User Datagram Protocol," August 1980.) [RFC0768] are expected by sending a Listener Notification message with the UDP Listener Address Specifier. This is a single octet with the UDP protocol number in it, followed by a reserved octet, followed by the UDP port number for the application endpoint.

UDP Listener Address Specifier

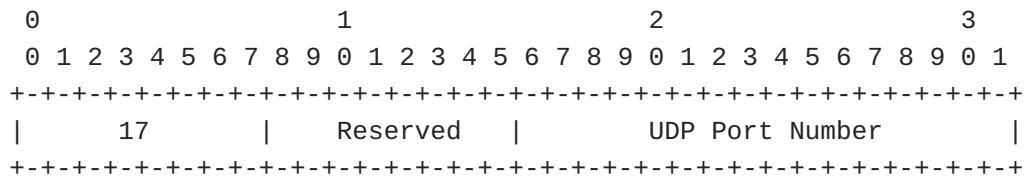

> **Reserved:** This field is unused. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.
>
> **UDP Port Number:** The UDP port for the application endpoint.

A value of zero in the UDP Port Number field indicates all UDP flows. This is identical to the All Specific Protocol Listener Address Specifier for UDP.

5.3.6. SCTP Listener Address Specifier

TOC

Nodes notify firewalls that inbound Stream Control Transport Protocol (SCTP) flow initiations (Stewart, R., "Stream Control Transmission Protocol," September 2007.) [RFC4960] are expected by sending a Listener Notification message with the SCTP Listener Address Specifier. This is a single octet with the SCTP protocol number in it, followed by a reserved octet, followed by the SCTP port number for the application endpoint.

SCTP Listener Address Specifier

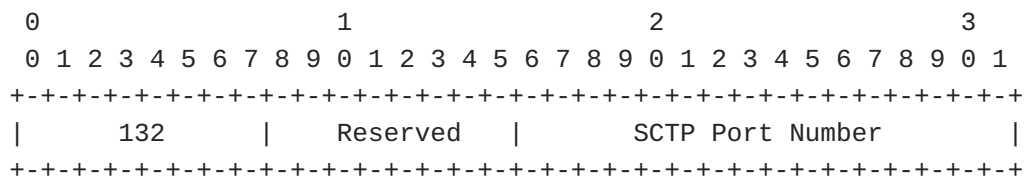

Reserved: This field is unused. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

UDP Port Number: The SCTP port for the application endpoint.

A value of zero in the SCTP Port Number field indicates all SCTP flows. This is identical to the All Specific Protocol Listener Address Specifier for SCTP.

5.3.7. DCCP Listener Address Specifier

TOC

Nodes notify firewalls that inbound Datagram Congestion Control Protocol (DCCP) flow initiations (Kohler, E., Handley, M., and S. Floyd, “Datagram Congestion Control Protocol (DCCP),” March 2006.) [RFC4340] are expected by sending a Listener Notification message with the DCCP Listener Address Specifier. This is a single octet with the DCCP protocol number in it, followed by a reserved octet, followed by the DCCP port number for the application endpoint.

DCCP Listener Address Specifier

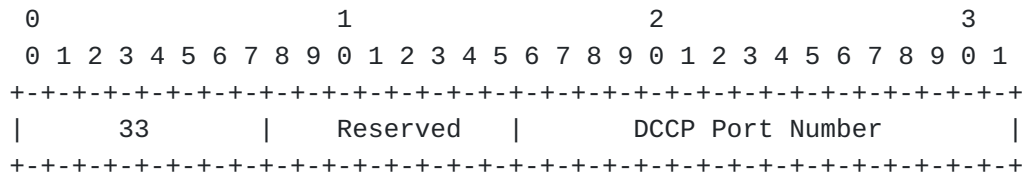

Reserved: This field is unused. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

UDP Port Number: The DCCP port for the application endpoint.

A value of zero in the DCCP Port Number field indicates all DCCP flows. This is identical to the All Specific Protocol Listener Address Specifier for DCCP.

TOC

5.4. Listener Notification

When a node expects to receive inbound flows from the exterior of a firewall, it MAY send a Listener Notification message to signal that inbound flow initiations should not be prohibited.

Listener Notification

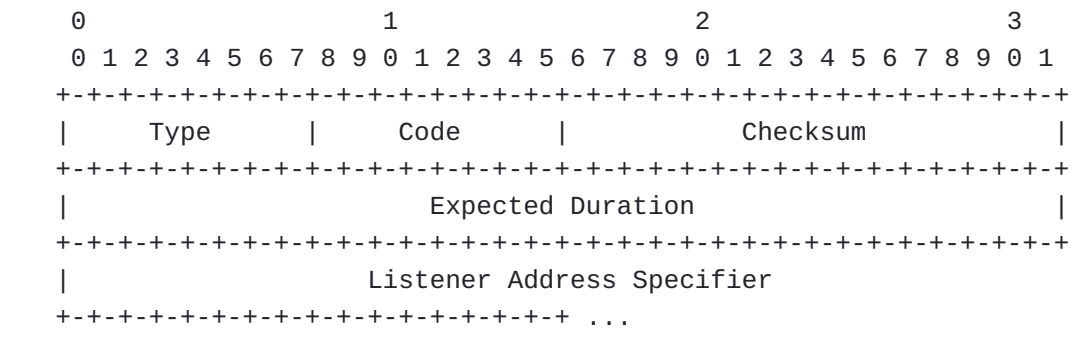

Type: TBD. Assigned by IANA to ALD messages.

Code: 3.

Checksum: ICMPv6 checksum.

Expected Duration: The number of seconds the application expects to be listening.

Listener Address Specifier: Describes the transport address of the application listener. See Section 5.3 (Listener Address Specifier).

Nodes MUST NOT send Listener Notification messages on any network to any destinations other than the unicast source addresses from which they receive Firewall Advertisement messages after joining the network. EXPERIMENTAL: Nodes operating in experimental mode MAY send the Experimental Listener Notification message, i.e. the same message except with type value 100 as defined in "Internet Control Message Protocol (ICMPv6)" (Conta, A., Deering, S., and M. Gupta, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification," March 2006.) [RFC4443] for use in experimental protocols and the four octet code 0x6161706c inserted between the Checksum and Expected Time Interval fields. Nodes MUST NOT send Experimental Listener Notification messages to destination addresses after receiving any regular Firewall Advertisement messages from the same source address.

TOC

5.5. Listener Acknowledgment

Firewalls send Listener Acknowledgment messages in response to receiving Listener Solication messages from nodes.

Listener Acknowledgment

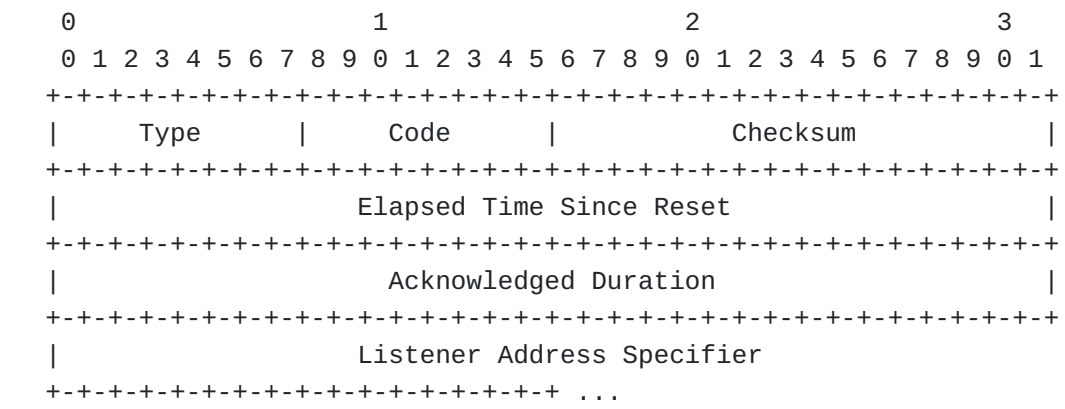

Type: TBD. Assigned by IANA to ALD messages.

Code: 4.

Checksum: ICMPv6 checksum.

Elapsed Time Since Reset: Number of elapsed seconds since the firewall state was last reset.

Acknowledged Duration: The number of seconds the firewall acknowledges the node will be listening.

Listener Address Specifier: Describes the transport address of the application listener. See Section 5.3 (Listener Address Specifier).

Firewalls MUST NOT transmit Listener Acknowledgment messages except in response to received Listener Notification messages.
Firewalls MUST NOT transmit Listener Acknowledgment messages with an Acknowledged Duration greater than the Expected Duration in the corresponding Listener Notification message.
After receiving a Listener Acknowledgment message, nodes MUST NOT transmit Listener Notification messages with a non-zero Requested Lifetime and the same Listener Address Specifier unless the Requested Lifetime is less than seven eighths (87.5%) of the Granted Lifetime value.
EXPERIMENTAL: Firewalls operating in experimental mode MAY respond to Experimental Listener Notification messages with the Experimental Listener Acknowledgment message, i.e. the same message except with type value 100 as defined in "Internet Control Message Protocol (ICMPv6)" (Conta, A., Deering, S., and M. Gupta, “Internet Control Message

Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification," March 2006.) [RFC4443] for use in experimental protocols and the four octet code 0x6161706c inserted between the Checksum and Elapsed Time Since Reset fields.

---

6. APPLICATION PROGRAMMING INTERFACE

TOC

[This section needs to be expanded to discuss how ALD functions are related to the operation of the conventional socket layer interface, i.e. how Listener Notifications are emitted when TCP sockets are put into and taken out of the LISTENING states, etc. Additional socket options for advanced usage may also be necessary here. Specific description of behavior for sockets in O_NONBLOCK mode should be defined.]

Existing programming interfaces, e.g. the widely used BSD sockets API, are sufficient for most applications. When TCP endpoints bound to global addresses transition into the LISTENING state, firewalls can be notified automatically with Listener Notification messages. Similar methods can be used for all other transport protocols.

Some applications require finer control over whether and how to notify firewalls of their listeners. This document recommends extensions to system configuration, interface control messages and socket options to meet their needs.

6.1. Normal Behavior of IPv6 Sockets

TOC

Applications using the BSD listen(2) function to place a TCP socket into the LISTENING state MAY be blocked while ALD notifies the appropriate firewalls. If the socket descriptor is opened with O_NONBLOCK or is otherwise marked as non-blocking, then listen(2) MAY return EINPROGRESS to indicate that ALD has not yet received Listener Acknowledgment messages from all appropriate firewalls. It MAY be possible to select(2) for completion by checking the socket for writing.

Applications using the BSD bind(2) function with UDP sockets MAY be blocked while ALD notifies the appropriate firewalls. If the socket descriptor is opened with O_NONBLOCK or is otherwise marked as non-blocking, then bind(2) MAY return EINPROGRESS to indicate that ALD has not yet received Listener Acknowledgment messages from all appropriate firewalls. It MAY be possible to select(2) for completion by checking the socket for writing.

Implementations of SCTP and DCCP are expected to implement similar methods of plumbing up ALD operations to the application layer.

If an application binds to specific interface addresses, then Listener Notification messages MAY be sent only to those firewalls with matching interior prefixes.
If a node receives a Listener Acknowledgment with an address specification that indicates the firewall has already discovered the application listener, then transmitting a Listener Notification MAY be skipped. If no ALD messages are necessary, then the application MUST receive the same service from the bind(2) and listen(2) system functions as when ALD is not operating.

6.2. Extensions to BSD Socket Interface

TOC

A new system configuration variable of boolean type, net.inet6.icmp6.ald_enabled, MAY be available on nodes to control whether ALD is enabled. The recommended default value is TRUE.
A new interface flag, IFF_NOALD MAY be available for disabling ALD on a per-interface basis. The recommended default if for the flag not to be set. The ifconfig(8) utility MAY provide the "-ald" parameter for controlling this option.
A new socket option of boolean type, IPV6_ALD_ENABLED MAY be used to control whether ALD is to be used on a per-socket basis. The default value for is recommended to be TRUE unless net.inet6.icmp6.ald_enabled is FALSE or the socket has already been bound to an interface address for which the interface has the IFF_NOALD flag set.

---

7. IANA CONSIDERATIONS

TOC

This memo includes several requests to IANA, which need to be gathered into this section accordingly.
All drafts are required to have an IANA considerations section (see the update of RFC 2434 (Narten, T. and H. Alvestrand, “Guidelines for Writing an IANA Considerations Section in RFCs,” March 2008.) [I-D.narten-iana-considerations-rfc2434bis] for a guide). If the draft does not require IANA to do anything, the section contains an explicit statement that this is the case (as above). If there are no requirements for IANA, the section will be removed during conversion into an RFC by the RFC Editor.

TOC

8. SECURITY CONSIDERATIONS

The author has not yet given sufficient consideration to security for writing an adequate security considerations section. Some readers have expressed concerns about spoofing. The author thinks protecting unicast ALD messages with IPsec Authenticated Header is the appropriate method for addressing such issues. An argument might be entertained for protecting the privacy of Listener Notification and Acknowledgment messages, and the author likewise believes IPsec Encapsulating Security Payload is the appropriate method for that. Key exchange for such security mechanisms should be specified by this document if IETF consensus regards addressing these considerations as essential.
All drafts are required to have a security considerations section. See "Guidelines for Writing RFC Text on Security Considerations" (Rescorla, E. and B. Korver, “Guidelines for Writing RFC Text on Security Considerations,” July 2003.) [RFC3552] for a guide.

---

9. References

TOC

9.1. Normative References

TOC

[RFC0768]	Postel, J., “User Datagram Protocol,” STD 6, RFC 768, August 1980 (TXT).
[RFC0793]	Postel, J., “Transmission Control Protocol,” STD 7, RFC 793, September 1981 (TXT).
[RFC2119]	Bradner, S., “Key words for use in RFCs to Indicate Requirement Levels,” BCP 14, RFC 2119, March 1997 (TXT, HTML, XML).
[RFC2460]	Deering, S. and R. Hinden, “Internet Protocol, Version 6 (IPv6) Specification,” RFC 2460, December 1998 (TXT, HTML, XML).
[RFC4303]	Kent, S., “IP Encapsulating Security Payload (ESP),” RFC 4303, December 2005 (TXT).
[RFC4340]	Kohler, E., Handley, M., and S. Floyd, “Datagram Congestion Control Protocol (DCCP),” RFC 4340, March 2006 (TXT).
[RFC4443]	Conta, A., Deering, S., and M. Gupta, “Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification,” RFC 4443, March 2006 (TXT).
[RFC4727]	Fenner, B., “Experimental Values In IPv4, IPv6, ICMPv4, ICMPv6, UDP, and TCP Headers,” RFC 4727, November 2006 (TXT).

[RFC4960]	Stewart, R., "Stream Control Transmission Protocol," RFC 4960, September 2007 (TXT).

9.2. Informative References

TOC

[I-D.narten-iana-considerations-rfc2434bis]	Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs," draft-narten-iana-considerations-rfc2434bis-09 (work in progress), March 2008 (TXT).
[NAT-PMP]	Cheshire, S., Krochmal, M., and K. Sekar, "NAT Port Mapping Protocol (NAT-PMP)," November 2001.
[RFC3552]	Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations," BCP 72, RFC 3552, July 2003 (TXT).
[RFC4864]	Van de Velde, G., Hain, T., Droms, R., Carpenter, B., and E. Klein, "Local Network Protection for IPv6," RFC 4864, May 2007 (TXT).
[UPnP-IGD]	UPnP Forum, "Universal Plug and Play Internet Gateway Device Standardized Gateway Device Protocol," September 2006.

Appendix A. Change Log

TOC

A.1. draft-woodyatt-ald-02 to draft-woodyatt-ald-03

TOC

*Adjusted column widths in XML source.

A.2. draft-woodyatt-ald-01 to draft-woodyatt-ald-02

TOC

*Fixed spelling errors.

*Local Network Protection is now [RFC4864] (Van de Velde, G., Hain, T., Droms, R., Carpenter, B., and E. Klein, "Local Network Protection for IPv6," May 2007.).

*Fix some bugs related to [RFC2119] (Bradner, S., “Key words for use in RFCs to Indicate Requirement Levels,” March 1997.) compliance.

*SCTP is now [RFC4960] (Stewart, R., “Stream Control Transmission Protocol,” September 2007.).

A.3. draft-woodyatt-ald-00 to draft-woodyatt-ald-01

TOC

*Added geeky cross-references for TCP and UDP.

*Simplified description of ICMPv6 checksum field descriptions.

*Changed the All Protocols Listener Address Specifier to use zero instead of 41, so that IPv6-in-IPv6 is eligible for specification.

*Added the SPI field to the ESP Listener Address Specifier.

*Added a note about zero UDP and TCP port numbers in the associated Listener Address Specifiers.

*Added Listener Address Specifiers for SCTP and DCCP.

*Added the All Specific Protocol Listener Address Specifier element and changed the associated requirements language to allow nodes to send them, and to explicitly disallow protocol numbers corresponding to IPv6 header extensions and the reserved protocol number.

Author's Address

TOC

	james woodyatt
	Apple Inc.
	1 Infinite Loop
	Cupertino, CA 95014
	US
Email:	jhw@apple.com

Full Copyright Statement

TOC

Intellectual Property