

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 4, 2009

A. Narayanan
F. Le Faucheur
D. Ward
R. Rahman
Cisco
March 3, 2009

IP Router Alert Option Extension
draft-narayanan-rtg-router-alert-extension-00

Status of this Memo

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

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/ietf/lid-abstracts.txt>.

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

This Internet-Draft will expire on September 4, 2009.

Copyright Notice

Copyright (c) 2009 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 in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Internet-Draft

Router Alert Option Extension

March 2009

Abstract

The IP Router Alert Option is an IP option that alerts transit routers to more closely examine the contents of an IP packet. RSVP, PGM and IGMP are some of the protocols which make use of the IP Router Alert option. The current specification for the IP Router Alert Option does not define mechanisms to facilitate discriminating across different users of Router Alert. As a result, networks using router Alert may have more security exposure than necessary and/or may unnecessarily block some transit Router Alert packets. This document describes new rules for the IP Router-Alert Option that aid routers to process these packets more selectively.

Table of Contents

1.	Terminology	3
1.1.	Conventions Used in This Document	3
2.	Introduction	4
3.	IP Router Alert Option Enhancement	6
4.	Security Considerations	9
5.	IANA Considerations	10
6.	Acknowledgments	11
7.	References	12
7.1.	Normative References	12
7.2.	Informative References	12
	Authors' Addresses	14

[1.](#) Terminology

For readability, this document uses the following loosely defined terms:

- o Slow path : Software processing path for packets
- o Fast path : ASIC/Hardware processing path for packets

[1.1.](#) Conventions Used in This Document

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. Introduction

[RFC2113] and [RFC2711] respectively define the IPv4 and IPv6 Router Alert Option. In this document, we collectively refer to those as the IP Router Alert. RSVP ([RFC2205], [RFC3209]), PGM ([RFC3208]) and IGMP ([RFC3376]) are but some of the protocols which make use of the IP Router Alert. Those protocols are used to support critical elements of the Internet infrastructure (e.g. RSVP-TE for traffic engineering within a service provider network) and as such they need to be protected.

IP datagrams carrying the IP Router Alert are usually examined in a router's "slow path" and an excess of such datagrams can cause performance degradation or packet drops in a router's "slow path". (Note that a router's "slow path" can potentially also be attacked with IP packets destined to one of the router's local IP addresses and requires corresponding security protection.)

[RFC4081] and [RFC2711] mention the security risks associated with the use of the IP Router Alert: flooding a router with bogus IP datagrams which contain the IP Router Alert would cause a performance degradation of the router's "slow path" and can also lead to packet drops in the "slow path".

[RFC2113] specifies no mechanism for identifying different users of IP Router Alert. As a result, many fast switching implementations of IP Router Alert punt most/all packets marked with IP Router Alert into the slow path. To protect against overloading routers which receive a large number of IP Router Alert packets that they do not

need to process, many router implementations limit the rate of packets punted into the slow path, but once again the lack of discrimination of different protocols may hamper the smooth functioning of protocols that depend on IP Router Alert. Further, some network operators actively protect routers from IP Router Alert packets by discarding these packets at the edge, which is undesirable for end-to-end operation of protocols carrying this option. Details on these issues and some recommendations on best practices are described in [[I-D.rahman-rtg-router-alert-considerations](#)]. The specification of an efficient, general-purpose, protocol-independent mechanism for discriminating between different applications would aid router implementations to more efficiently select the protocol messages they need to punt and locally process, while ignoring and forwarding in the fast path the messages that they do not need to see.

This document enhances the current specification of Router Alert to ensure that risks associated with unintentional interception of packets that are not of real interest to a given router are minimized

(if not eliminated) by facilitating identification in the fast path of the subset of packets with router alert that are of interest to the router. A key aspect of the proposal is to facilitate finer grain identification of router alert packets of interest versus unwanted router alert packets while only requiring inspection of the router alert header. In particular:

- o the enhancement allows the router to identify the application of packets marked with IP Router-Alert by simply examining the IP header, independent of any application packet format.
- o the enhancement allows router alert packets from different application protocols to be easily distinguished even if they share the same transport protocol (i.e. they have the same IP PID).
- o the enhancement allows router alert packets for the same application protocol but associated with different contexts (e.g. end to end RSVP vs internal RSVP-TE) to be easily distinguished.

Note that this mechanism does not prevent attacks of the form of bogus protocol messages which may be of interest to the router. More details on this are presented in [Section 4](#).

3. IP Router Alert Option Enhancement

We propose an extension to the specification and processing behaviour of the IP RAO header. [RFC2113] specifies a 2-octet value in the IP RAO option field. [RFC5350] specifies creation of an IANA registry for managing this 2-octet value, and proposes IPv4 RAO usage as follows:

Value	Description	Reference
0	Router Shall Examine Packet	RFC2113
1-32	Aggregated Reservation Nesting Level	RFC3175

33-65502	Available for assignment by the IANA	RFC5350	
65503-65534	Available for experimental use	RFC5350	
65535	Reserved	RFC5350	
+-----+			

An IANA-maintained IPv6 RAO registry is specified in [[RFC2711](#)] and clarified in [[RFC5350](#)]. The current IPv6 RAO usage is:

Value	Description	Reference	
+-----+			
0	Multicast Listener Discovery	RFC2711	
1	RSVP	RFC2711	
2	Active Networks	RFC2711	
3	Reserved	RFC5350	
4-35	RSVP Aggregation	RFC3175	
36-65535	Available for assignment by the IANA	RFC2711	
+-----+			

We propose to extend the definition of IP Router-Alert Option values. The 2-octet Option Value field will now be used to identify the protocol and context from an IP RAO perspective. For IANA assignment purposes, this value will be split into two fields as follows:

+-----+	
service selector (10 bits)	context selector (6 bits)
+-----+	

The service selector will be assigned to different applications by IANA and the context selector will be specific to the application protocol. Service selector 0 is reserved for backward compatibility and service selector 1023 is reserved for experimental use. Depending on requirements, a single protocol or application may be assigned multiple service selectors. All currently assigned option

values of IPv4 and IPv6 RAO have service selector 0. The experimental use range is extended to be 65472-65534.

All new protocols using IP RAO MUST request allocations of new service and context selector values, and MUST follow this format for the Option Value in the IP header. Additionally, new service and context selector values will be allocated for legacy protocols using IP RAO. Existing implementations of these legacy protocols SHOULD be updated to use the new selector values. The use of new option values for legacy protocols SHOULD be configurable by the user, and SHOULD be on by default. However, extensions to legacy protocols that require new option values MUST follow the new option format. For the purposes of this requirement, "legacy protocols" are defined as those already standardized in the IETF as using IP Router-Alert - specifically:

- o RSVP - IPv4: 0-32, IPv6: 1, 4-35 ([[RFC2205](#)],[[RFC3175](#)])
- o RSVP/TE - IPv4: 0, IPv6: 1 ([[RFC3209](#)])
- o IGMPv2 - IPv4: 0 ([[RFC2236](#)])
- o IGMPv3 - IPv4: 0 ([[RFC3376](#)])
- o MLDv1 - IPv6: 0 ([[RFC2710](#)])
- o MLDv2 - IPv6: 0 ([[RFC3810](#)])
- o Multicast Router Discovery - IPv4: 0, IPv6: unspecified ([[RFC4286](#)])
- o PGM - IPv4: 0 ([[RFC3208](#)])
- o Active Network - IPv6: 2 (cited in [[RFC2711](#)])

Correspondingly, "new protocols" are those for which the use of IP Router-Alert has not yet been standardized.

For service selector 1-1022, the value of the service and context

selector fields MUST be assigned in a manner such that the content of the IP RAO option is sufficient to determine whether a packet is of

interest to a node, with a reasonable level of granularity. For example, having the [\[RFC3175\]](#) aggregate reservation nesting level in the context selector allows P routers to quickly separate out RSVP messages for aggregate vs. end-to-end flows. Or, a separate context (or service) selector for RSVP-IPv4 vs. RSVP-TE sessions allows nodes to efficiently ignore one session type while processing another. Also, service and context selectors SHOULD be assigned the same for IPv4 and IPv6 versions of the same application.

Fast path switching implementations SHOULD first look at the service selector and context selector fields to determine whether they wish to select a packet with IP RAO for local processing. A table of in-use service and context selector values can be looked up during packet switching to determine whether the packet is to be locally processed. Thus, for packets marked with service selectors 1-1022, the value of the IP RAO value field is sufficient to rapidly determine whether the packet may be forwarded unmodified or whether it should be examined further for local processing. If the protocol/context selector of the packet does not match those that are of interest to currently running protocols, the router SHOULD forward these packets unmodified in the fast path. By extension, if no applications that use IP Router-Alert are currently running on the router, the router SHOULD forward all packets with IP Router-Alert Option in the fast path, unmodified.

The service selector 0 is reserved for backwards compatibility. For packets marked with service selector 0, the packet MUST be examined further to determine whether it is of local interest, in compliance with current protocol requirements. The context selector may be ignored for these packets.

All the practices described in [\[I-D.rahman-rtg-router-alert-considerations\]](#) regarding protecting router control plane resources from attacks based on IP RAO, and protecting different protocols using IP RAO from each other, continue to apply in this context.

4. Security Considerations

This document describes an efficient mechanism for router implementations to identify packets marked with the IP Router-Alert Option but which are not of interest to this router, and forward them unprocessed.

It is important to note that the use of this extension does not change in any way the security properties of the IP Router-Alert Option. Specifically, no claim is made of enhancing the security of IP Router-Alert Option usage. An attacker can always consume excess resources on a router's control plane and/or slow path by sending it bogus packets with IP RAO protocol/context selector values that are of interest to the router. However, the network operator now has the option to selectively suppress incoming IP RAO packets at the edge for protocols they are using in their network, while still permitting other applications with IP RAO to transit efficiently across their network. For example, a network operator could choose to suppress incoming IP RAO packets at the edge corresponding to RSVP/TE if they are using RSVP/TE in their network, but still transit end-to-end IPv4 RSVP sessions efficiently.

5. IANA Considerations

This document requires an extension to the IP RAO IANA registry established in [[RFC5350](#)]. IP Router-Alert Option values will be assigned as described in [Section 3](#).

[6.](#) Acknowledgments

We would like to thank Dave Oran, Magnus Westerlund, John Scudder, Ron Bonica, Ross Callon, and Alfred Hines for their comments.

[7.](#) References

[7.1.](#) Normative References

- [RFC0791] Postel, J., "Internet Protocol", STD 5, [RFC 791](#), September 1981.
- [RFC2113] Katz, D., "IP Router Alert Option", [RFC 2113](#), February 1997.
- [RFC2711] Partridge, C. and A. Jackson, "IPv6 Router Alert Option", [RFC 2711](#), October 1999.

[7.2.](#) Informative References

- [I-D.dasmith-mpls-ip-options]
Jaeger, W., Mullooly, J., Scholl, T., and D. Smith,
"Requirements for Label Edge Router Forwarding of IPv4
Option Packets", [draft-dasmith-mpls-ip-options-01](#) (work in
progress), October 2008.
- [I-D.rahman-rtg-router-alert-considerations]
Rahman, R., "IP Router Alert Considerations and Usage",
[draft-rahman-rtg-router-alert-considerations-00](#) (work in
progress), November 2008.

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2205] Braden, B., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", [RFC 2205](#), September 1997.
- [RFC2236] Fenner, W., "Internet Group Management Protocol, Version 2", [RFC 2236](#), November 1997.
- [RFC2710] Deering, S., Fenner, W., and B. Haberman, "Multicast Listener Discovery (MLD) for IPv6", [RFC 2710](#), October 1999.
- [RFC3175] Baker, F., Iturralde, C., Le Faucheur, F., and B. Davie, "Aggregation of RSVP for IPv4 and IPv6 Reservations", [RFC 3175](#), September 2001.
- [RFC3208] Speakman, T., Crowcroft, J., Gemmell, J., Farinacci, D., Lin, S., Leshchiner, D., Luby, M., Montgomery, T., Rizzo, L., Tweedly, A., Bhaskar, N., Edmonstone, R., Sumanasekera, R., and L. Vicisano, "PGM Reliable Transport

Protocol Specification", [RFC 3208](#), December 2001.

- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", [RFC 3209](#), December 2001.
- [RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", [RFC 3376](#), October 2002.
- [RFC3810] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", [RFC 3810](#), June 2004.
- [RFC4081] Tschofenig, H. and D. Kroeselberg, "Security Threats for Next Steps in Signaling (NSIS)", [RFC 4081](#), June 2005.
- [RFC4286] Haberman, B. and J. Martin, "Multicast Router Discovery", [RFC 4286](#), December 2005.

- [RFC4782] Floyd, S., Allman, M., Jain, A., and P. Sarolahti, "Quick-Start for TCP and IP", [RFC 4782](#), January 2007.
- [RFC5350] Manner, J. and A. McDonald, "IANA Considerations for the IPv4 and IPv6 Router Alert Options", [RFC 5350](#), September 2008.

Narayanan, et al. Expires September 4, 2009 [Page 13]

Internet-Draft Router Alert Option Extension March 2009

Authors' Addresses

Ashok Narayanan
Cisco Systems
1414 Mass Ave
Boxborough, MA 01719
USA

Email: ashokn@cisco.com

Francois Le Faucheur

Cisco Systems
Greenside, 400 Avenue de Roumanille
Sophia Antipolis, 06410
France

Email: flefauch@cisco.com

David Ward
Cisco Systems

Email: dward@cisco.com

Reshad Rahman
Cisco Systems
2000 Innovation Dr.
Kanata, Ontario K2K 3E8
Canada

Email: rrahman@cisco.com