Network Working Group
Internet-Draft
Intended status: BCP
Expires: May 18, 2009

R. Rahman, Ed.
Cisco

November 14, 2008

IP Router Alert Considerations and Usage
draft-rahman-rtg-router-alert-considerations-00

Status of this Memo

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. This document discusses security aspects and common practices around the use of router alert and discusses consequences on the use of router alert by existing or new applications. Common practices in router alert implementation facilitating router protection are also discussed. Finally a possible enhancement to the current specification of Router Alert is presented for feedback.

Table of Contents

1. Terminology

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

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

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 option. RSVP ([RFC2205], [RFC3209]), PGM ([RFC3208]) and IGMP ([RFC3376]) are some of the protocols which make use of the IP Router Alert option. 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 option 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 also be attacked with IP packets destined to one of the router's local IP addresses.)

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

[RFC2711] mentions that limiting, by rate or some other means, the use of Router Alert option is a way of protecting against a potential attack. However, if rate limiting is used as a protection mechanism but if the granularity of the rate limiting is not fine enough to distinguish among router alert packet of interest from unwanted router alert packet, a router alert attack could still severely

degrade operation of protocols of interest that depend on the use of IP Router Alert. Section 4 discusses examples of protection mechanisms that may be available from a router implementation of the IP router alert.

In some environments where it was not possible to accurately and reliably distinguish between router alert packet of interest and unwanted router alert packets, operators have resorted to actively protecting themselves against externally generated router alert packets in ways that result in end to end router alert packets being (occasionally or systematically) dropped and/or ignored on a segment. The consequences of such practises on the use of router alert by existing or new applications are discussed in Section 3 of the present document. This section also discusses environments where IP router alert can be used effectively.

This document also discusses in Appendix A a possible enhancements to 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. The objective of this appendix is to solicit feedback on the question of whether an enhancement to the current router alert specification is justified, and if yes, how to enhance it.

2.1. 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].

3. Guidelines for use of Router Alert

3.1. Reliance on Router Alert by Applications

As mentioned earlier, some networks actively protect themselves against externally generated router alert packets. This may be by tunneling router alert packets [I-D.dasmith-mpls-ip-options] so that the router alert option is hidden through that network or it may be via mechanisms resulting in occasional (e.g. rate limiting) or systematic drop of router alert packets.

Also, application protocols are usually carried within IP by a transport protocol such as UDP, TCP, or SCTP. In these cases, the application protocol is not visible in the IP header and could not be determined without further "deep packet" inspection. In the event of

the Router Alert option being used for an application protocol carried in a transport protocol, the intercepted IP packet would be delivered to the transport protocol for processing in accordance with [RFC2113]. However, the behavior of the transport protocol in these circumstances is not defined, and may cause rejection of the packet for various reasons.

As a consequence, in the general case of open networks, applications can not safely rely on router alert packets being visible to all nodes on the path today, and importantly can not even rely on router alert packets being transported end to end.

[RFC2113] implies that the router may examine other fields of a received packet that contains the IP Router Alert option to decide whether that packet needs further processing, but no further advice is given. Examination of other fields of the received IP packet that carries the Router Alert to help determine what to do with the packet would result in implementation-specific behavior that is unpredictable to the sender of the packet. Therefore applications cannot depend on router alert interception involving inspection of fields in the packet outside the IP header.

Thus, creating an application or end-to-end protocol that uses the IP Router Alert is currently considered harmful and is strongly discouraged. A different mechanism should be used to decrease the risk of impacting existing protocols that use the IP Router Alert option.

3.2. When Consenting Adults Exchange IP Router Alert Packets

In some controlled environments, the network administrator can determine that IP router alert packets will only be received from trusted well-behaved devices or can establish that the protection

mechanisms discussed in the present document against the plausible RAO- based DoS attacks (e.g. RAO filtering and rate-limiting) are sufficient. In that case, an application relying on exchange and handling of RAO packets (e.g. RSVP) may be safely deployed within the controlled network. In other words, a network that feels appropriately protected against the risks associated with his environment, may decide to freely and openly partake in router alert message exchange with consenting entities. A private enterprise

network firewalled from the Internet may be an example of such controlled environment.

In some environments, the network administrator can reliably ensure that IP router alert packets from any untrusted device (e.g. from external routers) are prevented from entering a trusted area (e.g. the internal routers). For example, this may be achieved by ensuring that routers straddling the trust boundary (e.g. edge routers) always encapsulate those packets (without Router Alert) through the trusted area (as discussed in [I-D.dasmith-mpls-ip-options]). In such environments, the risks of DOS attacks through the IP router alert vector is removed in the trusted area (or greatly reduced) even if IP router alert is used inside the trusted area (say for RSVP). Thus an application relying on Router Alert may be safely deployed within the trusted area. In other words, the network protects itself from the risks of partaking in IP Router-Alert exchange with strangers but feels free to exchange router alert messages among trusted parties. A Service Provider running RSVP-TE in his network may be an example of such protected environment.

When a controlled environment requires RAO packet exchange across his routers for some application (e.g. RSVP) and transits via a Service Provider network (that is not part of the controlled environment), the administrator of the controlled environment needs to ensure with the Service Provider that the router alert messages will not be dropped when transiting the Service Provider network. In other words, the network ought to ensure that another network that needs to be involved in exchange of router alert packet is consenting.

Since some existing applications would benefit from end-to-end transport of router alert packets, it is desirable that a Service Provider protects his network from attacks based on router alert using mechanisms that minimize dropping of end to end router alert packets. For example, using protection mechanisms such as those described in Section 4 a Service Provider can safely protect operation of a protocol depending on router alert within his network (e.g. RSVP-TE) while at the same time safely transporting router alert packets carrying another protocol that may be used end to end. As another example, using mechanisms such as those discussed in [I-D.dasmith-mpls-ip-options] a Service Provider can safely protect

operation of a protocol depending on router alert within his network

(e.g. RSVP-TE) while at the same time safely transporting router alert packets carrying another protocol that may be used end to end (e.g. RSVP IPv4/6).

Where the Service Provider cannot transport end to end router alert packets over his network (i.e. they choose to drop them), it is desirable that the Service Provider carry these packets through their network and remove the IP router alert option from the IP header on ingress/receipt of the packet. This ensures that at least the packet will make it through the Service Provider network allowing the packet to be intercepted via other means than the IP router alert.

Where the Service Provider does not ensure transport of router alert packets for an end to end protocol of interest to a Service Provider user, the user may remove the IP router alert option from the IP header before sending the packet to the Service Provider and may then intercept the packet on the other side using some other interception technique (and then possibly restore the IP Router alert option).

The authors of this document are seeking feedback on whether some of the practices discussed in this section ought to be elevated to BCP requirements or recommendations.

4. Example Protection Mechanisms in a Router Alert Implementation

Implementations of Router Alert on routers generally include mechanisms for protection against the associated security risk. This section provides examples of behaviors that may be supported by router implementations to help protect the router. The authors are seeking feedback about whether such mechanisms (or a subset/superset of those) ought to be elevated to BCP requirements and recommendations instead of simply described as examples.

4.1. Handling Packets Carrying the Router Alert Option

- a router implementation may elect to perform packet inspection to see whether they carry the Router Alert option in the "fast path". This avoids punting all packets to the slow path when the router is interested in some router alert packets.

- a router implementation may elect to not send to the "slow path" IP packets carrying the Router Alert option unless router alert interception is explicitly enabled on the router (or interface). This avoids punting any router alert packet when the router is not interested in any of those.

- a router implementation may elect to not send to the "slow path" IP packets carrying the Router Alert option unless there is at least one protocol explicitly enabled on the router (or interface) which is defined to use the IP Router Alert option. This avoids punting any router alert packet when the router is not interested in any of those.

- a router implementation may elect to allow configuration of which protocol is "of interest" for the Router Alert option interception (on router or interface level). The router implementation may then elect to not send packets carrying the Router Alert option to the "slow path" unless the payload protocol carried in the packet is configured as "protocol of interest". This avoids punting router alert packet carrying a protocol in which the router is not interested.

4.2. Applying Rate Limiting

- a router implementation may elect to support (in the fast path) rate limiting of the number of Router Alert IP datagrams (e.g. at router or interface level) which go to the "slow path". The benefits of rate limiting is described in [RFC2711].

- a router implementation may elect to support (in the fast path) separate rate limiting per payload protocol. This allows one

protocol relying on router alert to be protected from DOS attacks using router alert with a different protocol.

4.3. Router Alert in Congested Systems

- a router implementation may elect to support selective dropping of packets carrying the Router Alert option (rather than pass them to the "slow path") in preference to dropping other control plane packets, in the face of control plane congestion. This protects other control plane protocols from router alert attacks.

4.4. Handling Unknown Payload Protocols

If an IP packet contains the Router Alert option, but the payload protocol is not explicitly identified as a Payload of interest by the router examining the packet, the behavior is not defined by [RFC2113]. However, the definition of RSVP in [RFC2205] assumes that the packet will be forwarded using normal forwarding based on the destination IP address.

- a router implementation may elect to forward within the "fast path" (subject to all normal policies and forwarding rules) a packet carrying the Router Alert option containing a payload that is not a payload of interest to that router. The "not passing" behavior protects the router from DOS attacks using router alert packets of a protocol unknown to the router. The "forwarding" behavior contributes to transparent end to end transport of router alert packets (e.g. to facilitate their use by end to end application).

5. Security Considerations

This document discusses security risks associated with current usage of the IP Router Alert Option and associated practices.

6. IANA Considerations

None.

7. Contributors

The contributors to this document (in addition to the editors) are:

- David Ward:
  - Cisco Systems
  - wardd@cisco.com
- Francois Le Faucheur:
  - Cisco Systems
  - flefauch@cisco.com
- Ashok Narayanan:

    * Cisco Systems

    * ashokn@cisco.com

  o Adrian Farrell:

    * OldDog Consulting

    * adrian@olddog.co.uk

  o Tony Li:

    * tony.li@tony.li

8. Acknowledgments

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

9. References

9.1. 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.

9.2. 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.

[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.

[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.

[RFC4081] Tschofenig, H. and D. Kroeselberg, "Security Threats for Next Steps in Signaling (NSIS)", RFC 4081, June 2005.

[RFC5350] Manner, J. and A. McDonald, "IANA Considerations for the IPv4 and IPv6 Router Alert Options", RFC 5350, September 2008.

Appendix A. A New Filtering Mechanism to Select IP RAO Packets of Interest

This appendix discusses a possible enhancement to 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:

- the proposal allows router alert packets from different application protocols to be easily distinguished even if they share the same transport protocol (i.e. the have the same IP PID).

- the proposal 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.

The objective of this appendix is to solicit feedback on the question of whether an enhancement to the current router alert specification is justified, and if yes, how to enhance it.

[RFC2113] specifies no mechanism for identifying different users of IP RAO, with the result that many fast switching implementations punt most/all packets marked with IP RAO into the slow path. It is desirable for fast switching implementations to easily identify which packets marked with IP RAO are actually of interest to local protocols, so that other packets marked with IP RAO may be efficiently forwarded. In the past, router alert implementations have also looked at the IP PID [RFC0791] as a discriminator for different protocols using IP RAO. However, this has two drawbacks. The first is that messages with the same IP PID may represent different protocol operations for IP RAO processing (e.g. RSVP vs. RSVP-TE), or even different contexts (e.g. different levels of RSVP aggregation [RFC3175]), and it is desirable to distinguish these in the fast path. The second drawback is that IP PID values are a scarce resource, and it is likely that new IP datagram protocols will

be assigned UDP port numbers rather than IP PIDs. Any such future protocol which desires to use IP RAO would require additional checks in the fast path to select out the correct packets for local processing. To solve these problems, 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 a unified IPv4/IPv6 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	
65503-65534	Available for experimental use	
65535	Reserved	

We propose the following change to IP RAO processing:

o The following 2-octet field will now be used to identify the protocol and context from an IP RAO perspective. For IANA assignment purposes, this field will be split into two octets

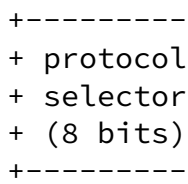

The protocol selector will be assigned for major protocols by IANA and the context selector will be specific to the protocol. Protocol selector 0 is reserved for backward compatibility and protocol selector 255 is reserved for experimental use. New protocol using IP RAO MUST allocate and use new protocol selector and context selector values. For protocol selector 1-254, the value of the protocol 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 selector for RSVP-IPv4 vs. RSVP-TE sessions allows nodes to efficiently ignore one session type while processing another.

- Fast path switching implementations SHOULD use this field to determine whether they wish to select a packet with IP RAO for local processing. A table of in-use protocol/context selector values can be looked up during packet switching to determine whether the packet is to be locally processed. For packets marked with protocol selectors 1-254, 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 punted to the "slow path" for local processing.

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

- All the requirements 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. The protocol selector and context selector fields MAY be used to differentiate between these protocols.

- Protocol and context selector values will be allocated for existing users of IP RAO as well (e.g. RSVP, IGMPv2 and PGM).

Author's Address

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

Email: rrahman@cisco.com

Full Copyright Statement

OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at http://www.ietf.org/ipr.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.