

IP Router Alert Considerations and Usage
draft-ietf-intarea-router-alert-considerations-09

Abstract

The IP Router Alert Option is an IP option that alerts transit routers to more closely examine the contents of an IP packet. Resource reSerVation Protocol (RSVP), Pragmatic General Multicast (PGM), Internet Group Management Protocol (IGMP), Multicast Listener Discovery (MLD), Multicast Router Discovery (MRD) and General Internet Signalling Transport (GIST) are some of the protocols that make use of the IP Router Alert Option. This document discusses security aspects and usage guidelines around the use of the current IP Router Alert Option. Specifically, it provides recommendation against using the Router Alert in the end-to-end open Internet as well as identify controlled environments where protocols depending on Router Alert can be used safely. It also provides recommendation about protection approaches for Service Providers. Finally it provides brief guidelines for Router Alert implementation on routers.

Status of this Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on February 3, 2012.

Copyright Notice

Copyright (c) 2011 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.

Table of Contents

1.	Terminology	3
1.1.	Conventions Used in This Document	3
2.	Introduction	4
3.	Security Concerns of Router Alert	6
4.	Guidelines for use of Router Alert	9
4.1.	Use of Router Alert End-to-End In the Internet (Router Alert in Peer Model)	9
4.2.	Use of Router Alert In Controlled Environments	10
4.2.1.	Use of Router Alert Within an Administrative Domain	10
4.2.2.	Use of Router Alert In Overlay Model	12
4.3.	Router Alert Protection Approaches for Service Providers	15
5.	Guidelines for Router Alert Implementation	17
6.	Security Considerations	18
7.	IANA Considerations	19
8.	Contributors	20
9.	Acknowledgments	21
10.	References	22
10.1.	Normative References	22
10.2.	Informative References	22
	Author's Address	24

1. Terminology

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

- o Fast path : Hardware or Application Specific Integrated Circuit (ASIC) processing path for packets. This is the nominal processing path within a router for IP datagrams.
- o Slow path : Software processing path for packets. This is a sub-nominal processing path for packets that require special processing or differ from assumptions made in fast path heuristics.
- o Next level protocol: the protocol transported in the IP datagram. In IPv4 [[RFC0791](#)], the next level protocol is identified by the IANA protocol number conveyed in the 8-bit "Protocol" field in the IPv4 header. In IPv6 [[RFC2460](#)], the next level protocol is identified by the IANA protocol number conveyed in the 8-bit "Next Header" field in the IPv6 header.

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 (RAO). In this document, we collectively refer to those as the IP Router Alert. The IP Router Alert Option is an IP option that alerts transit routers to more closely examine the contents of an IP packet.

Resource reSerVation Protocol (RSVP) ([RFC2205], [RFC3175], [RFC3209]), Pragmatic General Multicast (PGM) ([RFC3208]), Internet Group Management Protocol (IGMP) ([RFC3376]), Multicast Listener Discovery (MLD) ([RFC2710], [RFC3810]), Multicast Router Discovery (MRD) ([RFC4286]) and NSIS General Internet Signalling Transport (GIST) ([RFC5971]) are some of the protocols that make use of the IP Router Alert.

[Section 3](#) describes the security concerns associated with the use of the Router Alert Option.

[Section 4](#) provides guidelines for the use of Router Alert. More specifically, [Section 4.1](#) recommends that Router Alert not be used for end to end applications over the Internet, while [Section 4.2](#) presents controlled environments where applications/protocols relying on IP Router Alert can be deployed effectively and safely. [Section 4.3](#) provides recommendations on protection approaches to be used by Service Providers in order to protect their network from Router Alert based attacks.

Finally, [Section 5](#) provides generic recommendations for router implementation of Router Alert aiming at increasing protection against attacks.

The present document discusses considerations and practices based on the current specification of IP Router Alert ([RFC2113], [RFC2711]). Possible future enhancements to the specification of IP Router Alert (in view of reducing the security risks associated with the use of IP Router Alert) are outside the scope of this document. One such proposal is discussed in [[I-D.narayanan-rtg-router-alert-extension](#)] but at the time of this writing, the IETF has not adopted any extensions for this purpose.

The IPv6 base specification [RFC2460] defines the hop-by-hop option extension header. The hop-by-hop option header is used to carry optional information that must be examined by every node along a packet's delivery path. The IPv6 Router Alert Option is one particular hop by hop option. Similar security concerns to those discussed in the present document for the IPv6 Router Alert apply more generically to the concept of IPv6 hop-by-hop option extension

header. However, addressing the broader concept of IPv6 hop-by-hop option thoroughly would require additional material so as to cover additional considerations associated with it (such as the attacks effectiveness depending on how many options are included and on the range from to which the option-type value belongs, etc.), so this is kept outside the scope of the present document. A detailed discussion about security risks and proposed remedies associated with IPv6 hop-by-hop option can be found in [[I-D.krishnan-ipv6-hopbyhop](#)].

The IPv4 base specification [[RFC0791](#)] defines a general notion of IPv4 options that can be included in the IPv4 header (without distinguishing between hop-by-hop versus end-to-end option). The IPv4 Router Alert Option is one particular IPv4 option. Similar security concerns to those discussed in the present document for the IPv4 Router Alert apply more generically to the concept of IPv4 option. However, addressing the security concerns of the broader concept of IPv4 option thoroughly is kept outside the scope of the present document because it would require additional material so as to cover additional considerations associated with it (such as lack of option ordering, etc.), and because other IPv4 options are often blocked in firewalls and not very widely used, so the practical risks they present are largely non-existent.

3. Security Concerns of Router Alert

The IP Router Alert Option is defined ([RFC2113], [RFC2711]) as a mechanism that alerts transit routers to more closely examine the contents of an IP packet. [RFC4081] and [RFC2711] mention the security risks associated with the use of the IP Router Alert: flooding a router with bogus (or simply undesired) IP datagrams which contain the IP Router Alert could impact operation of the router in undesirable ways. For example, if the router punts the datagrams containing the IP Router Alert Option to the slow path, such an attack could consume a significant share of the router's slow path and could also lead to packet drops in the slow path (affecting operation of all other applications and protocols operating in the slow path) thereby resulting in a denial of service (DoS) ([RFC4732]).

Furthermore, [RFC2113] specifies no (and [RFC2711] specifies very limited) 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 (unless configured to systematically ignore or drop all Router Alert packets). However, some existing deployed IP routers can and do process IP packets containing the Router Alert Option inside the Fast Path.

Some IP Router Alert implementations are able to take into account the next level protocol as a discriminator for the punting decision for different protocols using IP Router Alert. However, this still only allows very coarse triage among various protocols using IP Router Alert for two reasons. First, the next level protocol is the same when IP Router Alert is used for different applications of the same protocol (e.g., RSVP vs. RSVP-TE), or when IP Router Alert is used for different contexts of the same application (e.g., different levels of RSVP aggregation [RFC3175]). Thus, it is not always possible to achieve the necessary triage in the fast path across IP Router Alert packets from different applications or from different contexts of an application. Secondly, some protocols requiring punting might be carried over a transport protocol (e.g., TCP or UDP) possibly because they require the services of that transport protocol, possibly because the protocol does not justify allocation of a scarce next level protocol value or possibly because not relying on a very widely deployed transport protocol is likely to result in deployment issues due to common middlebox behaviors (e.g. Firewalls or NATs discarding packets of "unknown" protocols). Thus, considering the next level protocol alone in the fast path is not sufficient to allow triage in the fast path of IP Router Alert packets from different protocols sharing the same transport protocol. Therefore, it is generally not possible to ensure that only the IP

Router Alert packets for next level protocols of interest are punted to the slow path while other IP Router Alert packets are efficiently forwarded (i.e., in fast path).

Some IP Router Alert implementations are able to take into account the value field inside the router alert option. However, only one value (zero) was defined in [RFC2113] and no IANA registry for IPv4 Router Alert values was available until recently ([RFC5350]). So this did not allow most IPv4 Router Alert implementation to support useful classification based on the value field in the fast path. Also, while [RFC2113] states that unknown values should be ignored (i.e. The packets should be forwarded as normal IP traffic), it has been reported that some existing implementations simply ignore the value field completely (i.e. Process any packet with an IPv4 Router Alert regardless of its option value). An IANA registry for further allocation of IPv4 Router Alert values has been introduced recently ([RFC5350]) but this would only allow coarse-grain classification, if supported by implementations. For IPv6, [RFC2711] states that "the value field can be used by an implementation to speed processing of the datagram within the transit router" and defines an IANA registry for these values. But again, this only allows coarse-grain classification. Besides, some existing IPv6 Router Alert implementations are reported to depart from that behavior.

[RFC2711] mentions that limiting, by rate or some other means, the use of IP 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 IP Router Alert packet of interest from unwanted IP Router Alert packet, a IP Router Alert attack could still severely degrade operation of protocols of interest that depend on the use of IP Router Alert.

In a nutshell, the IP router alert option does not provide a convenient universal mechanism to accurately and reliably distinguish between IP Router Alert packets of interest and unwanted IP Router Alert packets. This, in turn, creates a security concern when IP Router Alert Option is used, because, short of appropriate router implementation specific mechanisms, the router slow path is at risk of being flooded by unwanted traffic.

Note that service providers commonly allow external parties to communicate with a control plane application in their routers, such as with BGP peering. Depending on the actual environment and BGP security practices, the resulting DoS attack vector is similar, or somewhat less serious, with BGP peering than with Router Alert Option for a number of reasons that include:

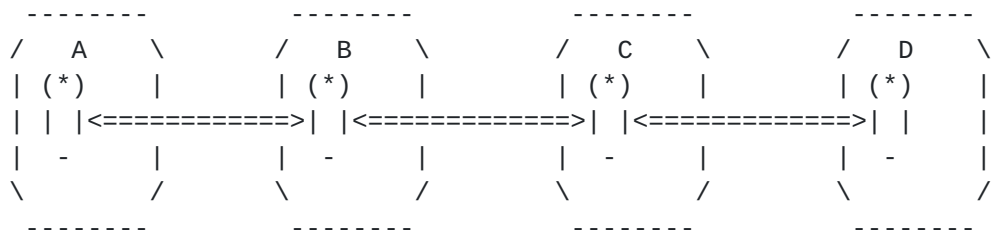
- o With BGP, edge routers only exchange control plane information with pre-identified peers and can easily filter out any control plane traffic coming from other peers or non-authenticated peers, while the Router-Alert option can be received in a datagram with any source address and any destination source. However, we note that effectiveness of such BGP filtering is dependent on proper security practices; poor BGP security practices (such as infrequent or inexistent update of BGP peers authentication keys) create vulnerabilities through which the BGP authentication mechanisms can be compromised.
- o with BGP Peering, the control plane hole is only open on the edge routers, and core routers are completely isolated from any direct control plane exchange with entities outside the administrative domain. Thus, with BGP, a DoS attack would only affect the edge routers, while with Router Alert Option, the attack could propagate to core routers. However, in some BGP environments, the distinction between edge and core routers is not strict, and many/most/all routers act as both edge and core routers; in such BGP environments, a large part of the network is exposed to direct control plane exchanges with entities outside the administrative domain (as it would be with Router Alert).
- o with BGP, the BGP policy control would typically prevent re-injection of undesirable information out of the attacked device, while with the Router-Alert option, the non-filtered attacking messages would typically be forwarded downstream. However, we note that there has been real life occurrences of situations where incorrect information was propagated through the BGP system, causing quite widespread problems.

4. Guidelines for use of Router Alert

4.1. Use of Router Alert End-to-End In the Internet (Router Alert in Peer Model)

Because of the security concerns associated with Router Alert discussed in [Section 3](#), network operators SHOULD actively protect themselves against externally generated IP Router Alert packets. Because there is no convenient universal mechanisms to triage between desired and undesired router alert packets, network operators currently often protect themselves in ways that isolate them from externally generated IP Router Alert packets. This might be achieved by tunneling IP Router Alert packets [[RFC6178](#)] so that the IP Router Alert Option is hidden through that network, or it might be achieved via mechanisms resulting in occasional (e.g., rate limiting) or systematic drop of IP Router Alert packets.

Thus, applications and protocols SHOULD NOT be deployed with a dependency on processing of the Router Alert Option (as currently specified) across independent administrative domains in the Internet. Figure 1 illustrates such a hypothetical use of Router Alert end-to-end in the Internet. We refer to such a model of Router Alert Option use as a "Peer Model" Router Alert Option use, since core routers in different administrative domains would partake in processing of Router Alert Option datagrams associated with the same signalling flow.



(*) closer examination of Router Alert Option datagrams

<==> flow of Router Alert Option datagrams

Figure 1: Use of Router Alert End-to-End in the Open Internet (Router Alert in Peer Model)

While this recommendation is framed here specifically in the context of router alert, the fundamental security risk that network operators want to preclude is to allow devices/protocols that are outside of their administrative domain (and therefore not controlled) to tap into the control plane of their core routers. Whether this control

plane access is provided through router alert option or would be provided by any other mechanism (e.g. Deep packet inspection) probably results in similar security concerns. In other words, the fundamental security concern is associated with the notion of end to end signaling in a Peer Model across domains in the Internet. As a result, it is expected that network operators would typically not want to have their core routers partake in end-to-end signalling with external uncontrolled devices through the open Internet, and therefore prevent deployment of end to end signalling in a Peer model through their network (regardless of whether that signalling uses Router Alert or not).

4.2. Use of Router Alert In Controlled Environments

4.2.1. Use of Router Alert Within an Administrative Domain

In some controlled environments, such as within a given Administrative Domain, the network administrator can determine that IP Router Alert packets will only be received from trusted well-behaved devices or can establish that specific protection mechanisms (e.g., RAO filtering and rate-limiting) against the plausible RAO-based DoS attacks are sufficient. In that case, an application relying on exchange and handling of RAO packets (e.g., RSVP) can be safely deployed within the controlled network. A private enterprise network firewalled from the Internet and using RSVP reservations for voice and video flows might be an example of such controlled environment. Such an environment is illustrated in Figure 2.



(*) closer examination of Router Alert Option datagrams

<==> flow of Router Alert Option datagrams

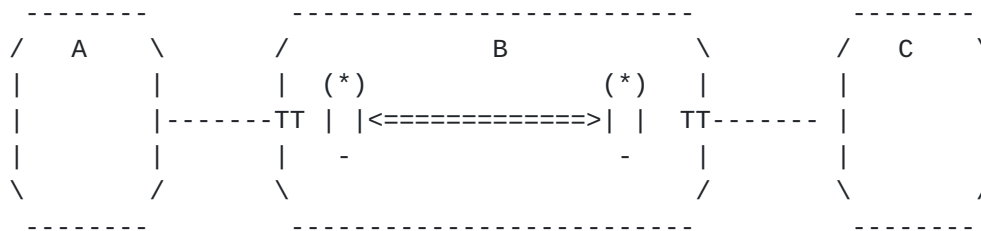
FW Firewall

Figure 2: Use of Router Alert Within an Administrative Domain

In some controlled environments, several Administrative Domains have a special relationship whereby they cooperate very tightly and

effectively operate as a single trust domain. In that case, one domain is willing to trust another with respect to the traffic injected across the boundary. In other words, a downstream domain is willing to trust that the traffic injected at the boundary has been properly validated/filtered by the upstream domain. Where it has been established that such trust can be applied to router alert option packets, an application relying on exchange and handling of RAO packets (e.g., RSVP) can be safely deployed within such a controlled environment. The entity within a company responsible for operating multimedia endpoints and the entity within the same company responsible for operating the network might be an example of such controlled environment. For example, they might collaborate so that RSVP reservations can be used for video flows from endpoints to endpoints through the network.

In some environments, the network administrator can reliably ensure that 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 might be achieved by ensuring that routers straddling the trust boundary (e.g., edge routers) always encapsulate those packets (without setting IP Router Alert -or equivalent- in the encapsulating header) through the trusted area (as discussed in [[RFC6178](#)]). 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-TE). Thus an application relying on IP Router Alert can be safely deployed within the trusted area. A Service Provider running RSVP-TE within his network might be an example of such protected environment. Such an environment is illustrated in Figure 3.



(*) closer examination of Router Alert Option datagrams

<==> flow of Router Alert Option datagrams

TT Tunneling of Router Alert Option datagrams

Figure 3: Use of Router Alert Within an Administrative Domain

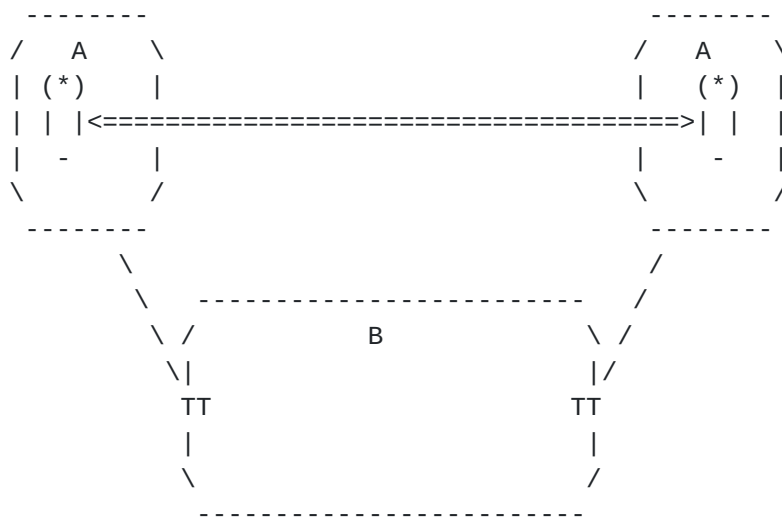
4.2.2. Use of Router Alert In Overlay Model

In some controlled environment:

- o the sites of a network A are interconnected through a service provider network B
- o the service provider network B protects itself from IP Router Alert messages without dropping those when they transit over the transit network (for example using mechanisms discussed in [[RFC6178](#)])

In such controlled environment, an application relying on exchange and handling of RAO packets (e.g., RSVP) in the network A sites (but not inside network B) MAY be safely deployed. We refer to such a deployment as a use of Router Alert in a Water-Tight Overlay.

"Overlay" because Router Alert Option datagrams are used in network A on top of, and completely transparently to, network B. "Water-Tight" because router alert option datagrams from A cannot leak inside network B. A private enterprise intranet realised as a Virtual Private Network (VPN) over a Service Provider network, and using RSVP to perform reservations within the enterprise sites for voice and video flows might be an example of such controlled environment. Such an environment is illustrated in Figure 4.



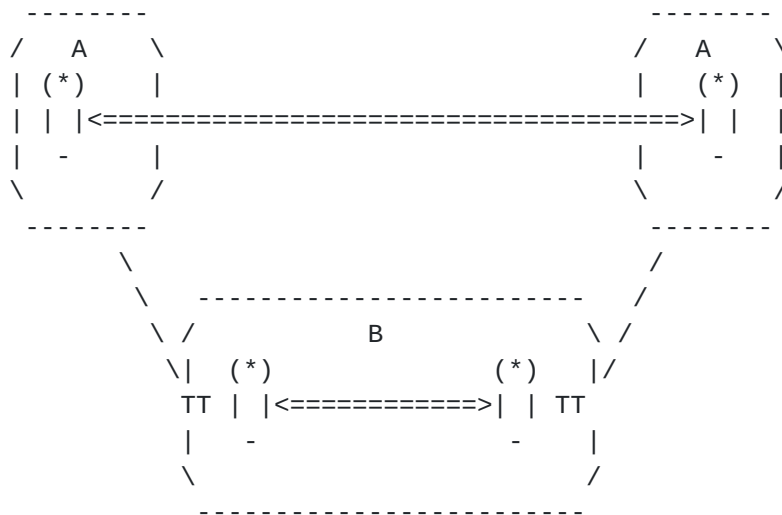
(*) closer examination of Router Alert Option datagrams

<==> flow of Router Alert Option datagrams

TT Tunneling of Router Alert Option datagrams

Figure 4: Use of Router Alert In Water-tight Overlay

In the controlled environment described above, an application relying on exchange and handling of RAO packets (e.g. RSVP-TE) in the service provider network B (but not in network A) can also be safely deployed simultaneously. Such an environment with independent, isolated, deployment of router alert in overlay at two levels is illustrated in Figure 5.



(*) closer examination of Router Alert Option datagrams

<==> flow of Router Alert Option datagrams

TT Tunneling of Router Alert Option datagrams

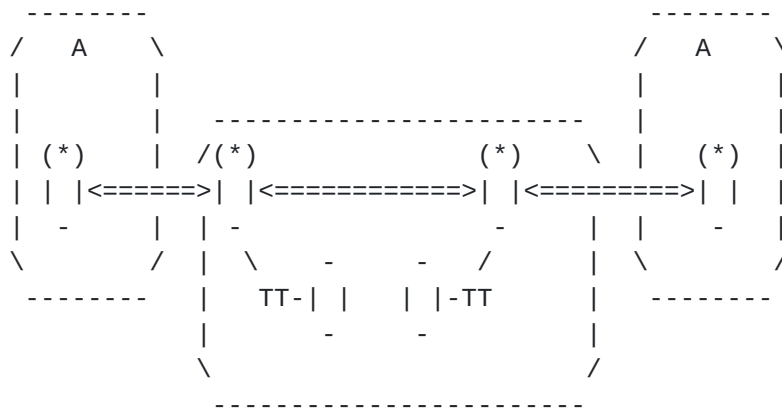
Figure 5: Use of Router Alert In Water-tight Overlay at Two Levels

In some controlled environment:

- o the sites of a network A are interconnected through a service provider network B
- o the service provider B processes router alert packets on the edge routers and protect these edge routers against RAO based attacks using mechanisms such as (possibly per port) RAO rate limiting and filtering
- o the service provider network B protects its core routers from Router Alert messages without dropping those when they transit over the transit network (for example using mechanisms discussed in [[RFC6178](#)])

In such controlled environment, an application relying on exchange and handling of RAO packets (e.g., RSVP) in the network A sites and in network B Edges (but not in the core of network B) can be safely deployed. We refer to such a deployment as a use of Router Alert in a Leak-Controlled Overlay. "Overlay" because Router Alert Option datagrams are used in network A on top of, and completely transparently to, network B core. "Leak-Controlled" because router alert option datagrams from A leak inside network B's B edges but not

inside network B's core. A private enterprise intranet, whose sites are interconnected through a Service Prover network, using RSVP for voice and video within network A sites as well as on Network B's edge to extend the reservation onto the attachment links between A and B (as specified in [\[RFC6016\]](#)) might be an example of such controlled environment. Such an environment is illustrated in Figure 4.



(*) closer examination of Router Alert Option datagrams

<==> flow of Router Alert Option datagrams

TT Tunneling of Router Alert Option datagrams

Figure 6: Use of Router Alert In Leak-Controlled Overlay

4.3. Router Alert Protection Approaches for Service Providers

[Section 3](#) discusses the security risks associated with the use of the IP Router Alert and how it opens up a DoS vector in the router control plane. Thus, a Service Provider **MUST** implement strong protection of his network against attacks based on IP Router Alert.

As discussed in [Section 4.2.2](#) some applications can benefit from the use of IP Router Alert packets in an Overlay model (i.e. Where Router Alert packets are exchanged transparently on top of a Service Provider). Thus, a Service Provider protecting his network from attacks based on IP Router Alert **SHOULD** use mechanisms that avoid (or at least minimize) dropping of end to end IP Router Alert packets (other than those involved in an attack).

For example, if the Service Provider does not run any protocol depending on IP Router Alert within his network, he might elect to simply turn-off punting/processing of IP Router Alert packet on his routers; this will ensure that end-to-end IP Router Alert packet

transit transparently and safely through his network.

As another example, using protection mechanisms such selective filtering and rate-limiting (that [Section 5](#) suggests be supported by IP Router Alert implementations) a Service Provider can protect the operation of a protocol depending on IP Router Alert within his network (e.g., RSVP-TE) while at the same time transporting IP Router Alert packets carrying another protocol that might be used end to end. Note that the Service Provider might additionally use protocol specific mechanisms that reduce the dependency on Router Alert for operation of this protocol inside the Service Provider environment; use of RSVP refresh reduction mechanisms ([\[RFC2961\]](#)) would be an example of such mechanisms in the case where the Service Provider is running RSVP-TE within his network since this allows refresh of existing Path and Resv states without use of the IP Router Alert Option.

As yet another example, using mechanisms such as those discussed in [\[RFC6178\]](#) a Service Provider can safely protect the operation of a protocol depending on IP Router Alert within his network (e.g., RSVP-TE) while at the same time safely transporting IP Router Alert packets carrying another protocol that might be used end to end (e.g., IPv4/IPv6 RSVP). We observe that while tunneling of Router Alert Option datagrams over an MPLS backbone as discussed in [\[RFC6178\]](#) is well understood, tunneling Router Alert Option datagrams over an non-MPLS IP backbone presents a number of issues (and in particular for determining where to forward the encapsulated datagram) and is not common practice at the time of writing this document.

As a last resort, if the SP does not have any means to safely transport end to end IP Router Alert Option packets over his network, the SP can drop those packets. It must be noted that this has the undesirable consequence of preventing the use of the Router Alert Option in the Overlay Model on top of this network, and therefore prevents users of that network from deploying a number of valid applications/protocols in their environment.

5. Guidelines for Router Alert Implementation

A router implementation of IP Router Alert Option SHOULD include protection mechanisms against Router Alert based DoS attacks appropriate for their targeted deployment environments. For example, this can include ability on an edge router to "tunnel" IP Router Alert Option of received packets when forwarding those over the core as discussed in [[RFC6178](#)]. As another example, although not always available from current implementations, new implementations MAY include protection mechanisms such as selective (possibly dynamic) filtering and rate-limiting of IP Router Alert Option packets.

In particular, router implementations of IP Router Alert Option SHOULD offer the configuration option simply to ignore the presence of "IP Router Alert" in IPv4 and IPv6 packets. As discussed in [Section 4.3](#), that permits IP Router Alert packets to transit a network segment without presenting an adverse operational security risk to that particular network segment, provided the operator of that network segment does not ever use the IP Router Alert messages for any purpose.

If an IP packet contains the IP Router Alert Option, but the next level protocol is not explicitly identified as a protocol of interest by the router examining the packet, the behavior is not explicitly defined by [[RFC2113](#)]. However, the behavior is implied and, for example, the definition of RSVP in [[RFC2205](#)] assumes that the packet will be forwarded using normal forwarding based on the destination IP address. Thus, a router implementation SHOULD forward within the "fast path" (subject to all normal policies and forwarding rules) a packet carrying the IP Router Alert Option containing a next level protocol that is not a protocol of interest to that router. The "not punting" behavior protects the router from DoS attacks using IP Router Alert packets of a protocol unknown to the router. The "forwarding" behavior contributes to transparent end to end transport of IP Router Alert packets (e.g., to facilitate their use by end to end application).

Similarly, an implementation MAY support selective forwarding within "the fast path" (subject to all normal policies and forwarding rules) or punting of a packet with the IP Router Alert Option, based on the Value field of the Router Alert Option. This would allow router protection against DoS attacks using IP Router Alert packets with value that is not relevant for that router (e.g. Nesting levels of Aggregated RSVP Reservation [[RFC5350](#)]).

6. Security Considerations

This document discusses security risks associated with current usage of the IP Router Alert Option and associated practices. This document expands the security considerations of [[RFC2113](#)] and [[RFC2711](#)], which defined the RAO, to discuss security risks associated with current usage of the IP Router Alert Option and associated practices. See [[RFC4081](#)] for additional security considerations.

7. IANA Considerations

None.

8. Contributors

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

- o Reshad Rahman:
 - * Cisco Systems
 - * rrahman@cisco.com
- o David Ward:
 - * Juniper Networks
 - * dward@juniper.net
- o Ashok Narayanan:
 - * Cisco Systems
 - * ashokn@cisco.com
- o Adrian Farrel:
 - * OldDog Consulting
 - * adrian@olddog.co.uk
- o Tony Li:
 - * tony.li@tony.li

9. Acknowledgments

We would like to thank Dave Oran, Magnus Westerlund, John Scudder, Ron Bonica, Ross Callon, Alfred Hines, Carlos Pignataro, Roland Bless, Jari Arkko and Ran Atkinson for their comments. This document also benefited from discussions with Jukka Manner and Suresh Krishnan. The discussion about use of the value field in the IPv4 Router Alert borrowed from a similar discussion in [[RFC5971](#)].

10. References

10.1. Normative References

- [RFC0791] Postel, J., "Internet Protocol", STD 5, [RFC 791](#), September 1981.
- [RFC2113] Katz, D., "IP Router Alert Option", [RFC 2113](#), February 1997.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [RFC 2460](#), December 1998.
- [RFC2711] Partridge, C. and A. Jackson, "IPv6 Router Alert Option", [RFC 2711](#), October 1999.
- [RFC5350] Manner, J. and A. McDonald, "IANA Considerations for the IPv4 and IPv6 Router Alert Options", [RFC 5350](#), September 2008.

10.2. Informative References

- [I-D.krishnan-ipv6-hopbyhop]
Krishnan, S., "The case against Hop-by-Hop options", [draft-krishnan-ipv6-hopbyhop-05](#) (work in progress), October 2010.
- [I-D.narayanan-rtg-router-alert-extension]
Narayanan, A., Faucheur, F., Ward, D., and R. Rahman, "IP Router Alert Option Extension", [draft-narayanan-rtg-router-alert-extension-00](#) (work in progress), March 2009.
- [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.
- [RFC2710] Deering, S., Fenner, W., and B. Haberman, "Multicast Listener Discovery (MLD) for IPv6", [RFC 2710](#), October 1999.
- [RFC2961] Berger, L., Gan, D., Swallow, G., Pan, P., Tommasi, F., and S. Molendini, "RSVP Refresh Overhead Reduction Extensions", [RFC 2961](#), April 2001.

- [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.
- [RFC4732] Handley, M., Rescorla, E., and IAB, "Internet Denial-of-Service Considerations", [RFC 4732](#), December 2006.
- [RFC5971] Schulzrinne, H. and R. Hancock, "GIST: General Internet Signalling Transport", [RFC 5971](#), October 2010.
- [RFC6016] Davie, B., Le Faucheur, F., and A. Narayanan, "Support for the Resource Reservation Protocol (RSVP) in Layer 3 VPNs", [RFC 6016](#), October 2010.
- [RFC6178] Smith, D., Mullooly, J., Jaeger, W., and T. Scholl, "Label Edge Router Forwarding of IPv4 Option Packets", [RFC 6178](#), March 2011.

Author's Address

Francois Le Faucheur (editor)
Cisco Systems
Greenside, 400 Avenue de Roumanille
Sophia Antipolis 06410
France

Phone: +33 4 97 23 26 19
Email: flefauch@cisco.com