Internet Engineering Task Force (IETF)
Request for Comments: 6398
BCP: 168
Updates: 2113, 2711
Category: Best Current Practice
ISSN: 2070-1721

F. Le Faucheur, Ed.
Cisco

October 2011

IP Router Alert Considerations and Usage

Abstract

The IP Router Alert Option is an IP option that alerts transit routers to more closely examine the contents of an IP packet. The Resource reSerVation Protocol (RSVP), Pragmatic General Multicast (PGM), the Internet Group Management Protocol (IGMP), Multicast Listener Discovery (MLD), Multicast Router Discovery (MRD), and General Internet Signaling 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, thereby updating RFC 2113 and RFC 2711. Specifically, it provides recommendations against using the Router Alert in the end-to-end open Internet and identifies controlled environments where protocols depending on Router Alert can be used safely. It also provides recommendations about protection approaches for service providers. Finally, it provides brief guidelines for Router Alert implementation on routers.

Status of This Memo

This memo documents an Internet Best Current Practice.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on BCPs is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc6398.

Table of Contents

1. Introduction

[RFC2113] and [RFC2711] define the IPv4 and IPv6 Router Alert Options (RAOs), respectively. In this document, we collectively refer to those options 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.

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

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.

This document discusses considerations and practices based on the current specifications of IP Router Alert ([RFC2113], [RFC2711]). Possible future enhancements to the specifications 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 [RAO-EXT], 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 options extension header. The hop-by-hop options 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 this document for the IPv6 Router Alert apply more generically to the concept of the IPv6 hop-by-hop options extension header. However, thoroughly addressing the broader concept of the

IPv6 hop-by-hop option would require additional material so as to cover additional considerations associated with it (e.g., the effectiveness of the attack could depend on how many options are included and on the range to which the option-type value belongs), so this is kept outside the scope of this document. A detailed discussion about security risks and proposed remedies associated with the IPv6 hop-by-hop option can be found in [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 the hop-by-hop and end-to-end options). The IPv4 Router Alert Option is one particular IPv4 option. Security concerns similar to those discussed in this document for the IPv4 Router Alert apply more generically to the concept of the IPv4 option. However, thoroughly addressing the security concerns of the broader concept of the IPv4 option is kept outside the scope of this 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 nonexistent.

2. Terminology

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

- 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.

- 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.

- 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.

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. 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 that 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 a 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 - Traffic Engineering (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 (1) they require the services of that transport protocol, (2) the protocol does not justify allocation of a scarce next level protocol value, or (3) 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 the 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 implementations 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 the 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 IP Router Alert packets of interest from unwanted IP Router Alert packets, an 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 the 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, with BGP peering, the resulting DoS attack vector is similar to or somewhat less serious than it would be with the Router Alert Option for a number of reasons, including the following:

- 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 address. However, we note that the effectiveness of such BGP filtering is dependent on proper security practices; poor BGP security practices (such as infrequent or nonexistent update of BGP peers' authentication keys) create vulnerabilities through which the BGP authentication mechanisms can be compromised.

- 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 the 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).

- 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 have been real-life occurrences of situations where incorrect information was propagated through the BGP system, causing 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 are 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 signaling flow.

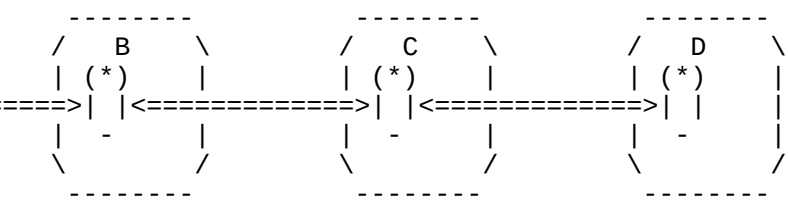

(*) closer examination of Router Alert Option datagrams

<==> 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. Similar security concerns would probably result whether this control plane access is provided through the Router Alert Option or provided by any other mechanism (e.g., deep packet inspection). 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 signaling with external uncontrolled devices through the open Internet, and therefore prevent deployment of end-to-end signaling in a Peer Model through their network (regardless of whether that signaling 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 a controlled environment. Such an environment is illustrated in Figure 2.

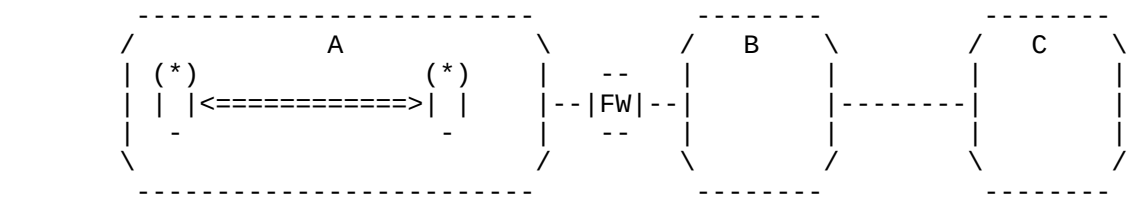

(*) closer examination of Router Alert Option datagrams

<==> flow of Router Alert Option datagrams

FW: Firewall

Figure 2: Use of Router Alert within an Administrative Domain - Private Enterprise Network Firewalled from the Internet and Using RSVP Reservations

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 a 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 are removed (or greatly reduced) in the trusted area 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 its network might be an example of such a protected environment. Such an environment is illustrated in Figure 3.

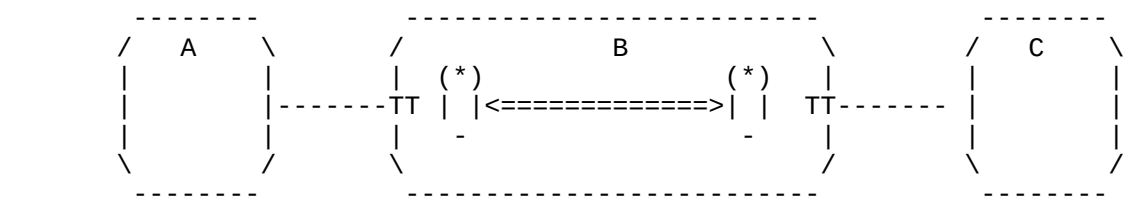

(*) closer examination of Router Alert Option datagrams

<==> flow of Router Alert Option datagrams

TT: Tunneling of Router Alert Option datagrams

Figure 3: Use of Router Alert within an Administrative Domain - Service Provider Running RSVP-TE within Its Network

4.2.2. Use of Router Alert in Overlay Model

In some controlled environment:

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

In such a 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) can 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; and "Water-Tight", because Router Alert Option datagrams from network A cannot leak inside network B. A private enterprise intranet realized 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 a controlled environment. Such an environment is illustrated in Figure 4.

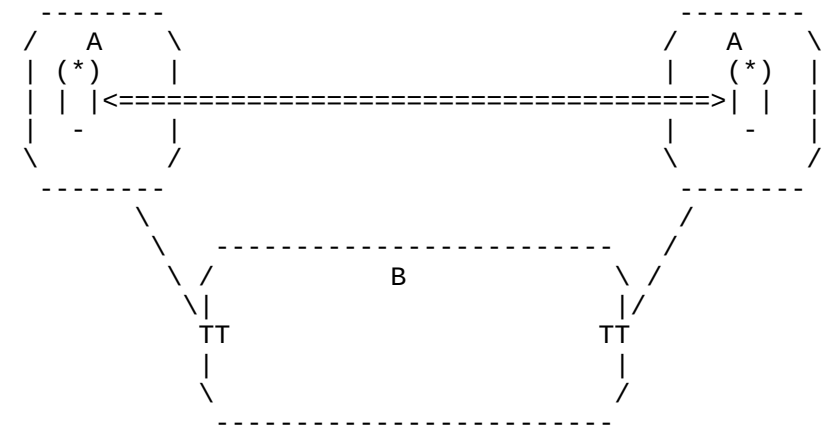

(*) closer examination of Router Alert Option datagrams

<==> 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.

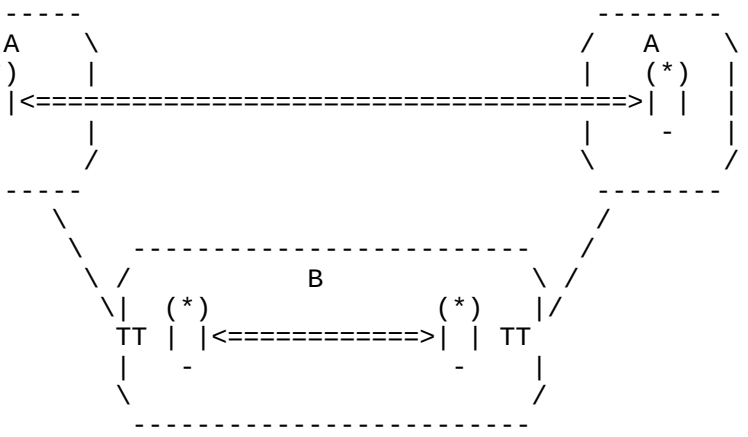

(*) closer examination of Router Alert Option datagrams

<==> 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:

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

In such a controlled environment, an application relying on exchange and handling of RAO packets (e.g., RSVP) in the network A sites and in network B's 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's core; and "Leak-Controlled", because Router Alert Option datagrams from network A leak inside network B's edges but not inside network B's core. A private enterprise intranet, whose sites are interconnected through a service provider 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 networks A and B (as specified in [RFC6016]), might be an example of such a controlled environment. Such an environment is illustrated in Figure 6.

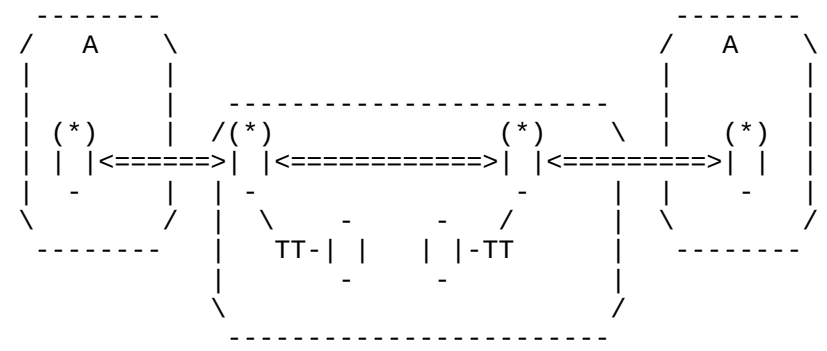

(*) closer examination of Router Alert Option datagrams

<==> 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 its 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 its network from

attacks based on IP Router Alert SHOULD use mechanisms that avoid (or at least minimize) the 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 its network, it might elect to simply turn off punting/processing of IP Router Alert packets on its routers; this will ensure that end-to-end IP Router Alert packets transit transparently and safely through its network.

As another example, using protection mechanisms such as selective filtering and rate limiting (which 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 its 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 its network, since this allows the refresh of existing Path and Resv states without the 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 its 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 a non-MPLS IP backbone presents a number of issues (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 service provider does not have any means to safely transport end-to-end IP Router Alert Option packets over its network, the service provider 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 that 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 the IP Router Alert Option SHOULD include protection mechanisms against Router-Alert-based DoS attacks as appropriate for their targeted deployment environments. For example, this can include the ability of an edge router to "tunnel" received IP Router Alert Option packets when forwarding those packets 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 the IP Router Alert Option SHOULD offer the configuration option to simply 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 applications).

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 a value that is not relevant for that router (e.g., nesting levels of aggregated RSVP reservation [RFC5350]).

6. Security Considerations

This document expands the security considerations of [RFC2113] and [RFC2711], which define the IPv4 and IPv6 RAOs, respectively, by discussing security risks associated with usage of the current IP Router Alert Option and associated practices. See [RFC4081] for additional security considerations.

7. Contributors

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

Reshad Rahman
Cisco Systems
rrahman@cisco.com

David Ward
Juniper Networks
dward@juniper.net

Ashok Narayanan
Cisco Systems
ashokn@cisco.com

Adrian Farrel
OldDog Consulting
adrian@olddog.co.uk

Tony Li
Cisco Systems
tony.li@tony.li

8. Acknowledgments

The editor and contributors 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 is borrowed from a similar discussion in [RFC5971].

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.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 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.

9.2. Informative References

[IPv6-HOPBYHOP] Krishnan, S., "The case against Hop-by-Hop options", Work in Progress, October 2010.

[RAO-EXT] Narayanan, A., Le Faucheur, F., Ward, D., and R. Rahman, "IP Router Alert Option Extension", Work in Progress, March 2009.

[RFC2205] Braden, R., Ed., 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., Ed., and L. Costa, Ed., "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., Ed., Rescorla, E., Ed., 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