INTERNET-DRAFT
Expiration: August 2001
File: draft-lindell-waypoint-00.txt

Bob Lindell
Bob Braden
USC/ISI

Waypoint - A Path Oriented Delivery Mechanism for IP based Control, Measurement, and Signaling Protocols

Status of this Memo

Abstract

This document describes the Waypoint path oriented delivery mechanism. Waypoint attempts to rationalize the packet interception problem that has been addressed by different mechanisms such as router alert or RSVP protocol number 46 intercept. It borrows concepts from prior mechanisms, including the hop by hop security model of RSVP. Waypoint strives to be complete, eliminating the need to reimplement common functionality in the higher layers of the signaling protocol stack.

1. Introduction

There are a broad range of protocols currently defined, or under development, for the Internet Protocol that require that ability to perform path oriented operations. Path oriented operations require the ability to process control information at every router or some defined subset of the routers along the end to end path traversed by a given IP routing path. For some path oriented operations the signaling is actually confined to a portion of the end to end path. On example is path oriented network management operations confined to a providers domain. Another example is the establishment of routing state within a given domain (e.g. traffic engineered tunnels). In this document we use the term end to end in this more general context. End to end will refer to either the entire data path or a well defined contiguous portion of the path.

We envision that a number of existing protocols, RSVP, RSVP-TE, and LDP could be layered upon Waypoint. These protocols, with end to end semantics and message reliability, most closely occupy the transport layer of the OSI reference model. Waypoint, on the other hand, is sandwiched somewhere in murky boundary between the transport and network layers. It is distinctly above IP, and as such, is expected to be implemented over the full range of IP protocols, both unicast and multicast IPv4 and IPv6 at the present time.

A major motivation for Waypoint is the belief that its existance will enable the creation of novel new path oriented control, signaling, and measurement protocols. Particularly appealing are the simple stateless measurement protocols. An example would be an enhanced path tracing protocol that was implemented on top of Waypoint rather than ICMP.

2. Current Shortcomings

There are a number of long term shortcomings to the current approaches of packet interception and the use of ICMP for measurement purposes. We will briefly highlight some of these issues in this section.

The current implementation of RSVP-TE grew out of an interest in reusing the existing RSVP protocol for MPLS tunnels. In the end, it required only the addition of a few new message object types and some additional processing rules. An unexpected result is that it seems to have reused the RSVP protocol number as well! This may have been deliberate in order to reuse the existing intercept mechanisms that may be limited to only a single protocol number on some router implementations. In the longer term, it would seem important for the clarity of implementations to separate different path oriented operations

with distinct protocol and/or port numbers.

Both the router alert option and the RSVP protocol number 46 intercept suffer from limited ability to control the interception process. Experience has demonstrated interest in the ability to tunnel a packet through a series of intermediate routers along a path. Similarly, for realistic deployment scenarios, some routers along a path will not implement a newly deployed service and it is desired that the packet be forwarded through to the next router along the path capable of processing the packet. Router alert and protocol 46 intercept requires IP encapsulation as the only method that will tunnel packets through a series of routers without intermediate processing. Router alert has another unwanted performance penalty. When a signaling packet is forwarded through a router that does not implement the path oriented service, it will likely be processed in the slower forwarding path due to the existance of an IP option in the packet. Basically, these current mechanisms have limited hop control and performance penalties over other approaches.

There are many variations of path oriented measurements that use ICMP. All of these approaches suffer substantially either from a feature perspective or measured results. It is well known that ICMP processing on most routers is not representative of the router performance, especially in the measured delay. Simple stateless path oriented measurement solutions using Waypoint would eliminate many of these flaws. Measurement packets could be properly timestamped at different time in the reception, processing, and transmission stages to more accurately represent the measured quantities of interest. The information gleamed from an ICMP response is also quite limited. At best, one is able to determine an IP address and some information about the type of router. Clearly more novel services could be created if there was an ability to perform path oriented operations coupled with freedom to control the payload contents.

By definition, IPSEC provides end to end security. Path oriented operations are therefore excluded from the use of IPSEC. To prevent each protocol from inventing their own security solution, it is important that the Internet architecture provide a comparable service to IPSEC for path oriented protocols.

3. Functional Description

The Waypoint protocol envisions a large collection of well defined path oriented protocols. To accommodate many services, Waypoint packets carry port fields in an identical fashion to the use of port numbers in the UDP and TCP protocols. Well known services will use globally assigned well known Waypoint port numbers.

Waypoint, using IP, delivers packets along an end to end path. Unlike other IP transport

protocols, it is not a strictly end to end protocol. Waypoint packets can be delivered to intermediate routers along the end to end path even though the destination address of the packet is not that of the router. This distinguishes Waypoint from most transport protocols and places the functionality somewhere in the murky boundary between the 3rd and 4th layer of the OSI protocol model.

Since Waypoint is not strictly end to end, the common functionality of IPSEC cannot be used with Waypoint. In the place of IPSEC, Waypoint defines its own secure transport functionality as a replacement service. Again, it is important that Waypoint preserve the existing IP functionality for the protocols which will be layered above Waypoint.

4. Protocol Packet Definition

The Waypoint protocol header appears immediately following the IP protocol header. The IP protocol number for Waypoint is ??. The header contains source and destination ports, a checksum field, and the length of the Waypoint header. The header length field delineates additional Waypoint header options for the payload. Waypoint header object are framed as type, length, value (TLV) objects. The only currently defined option is for security.

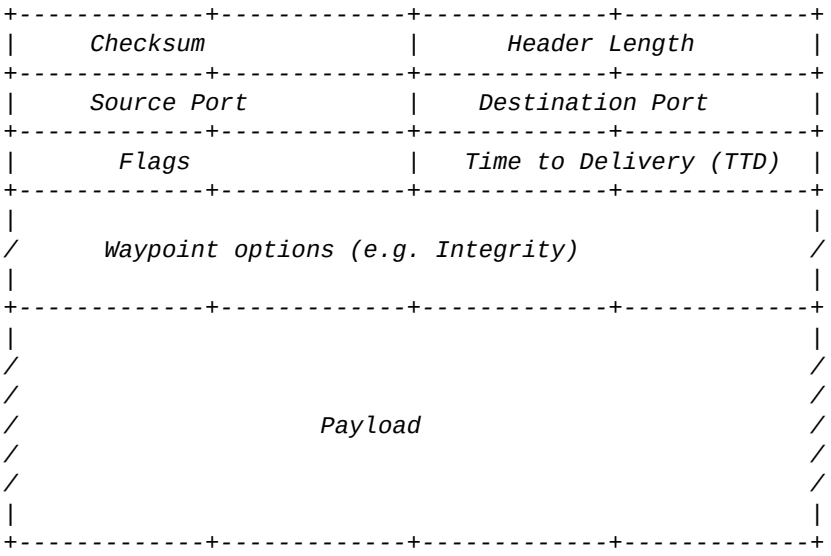

5. Protocol Processing Rules

Waypoint packets are forwarded along the end to end path towards the IP destination address for both unicast and multicast addresses. Unlike the typical IP data packet, it is likely that a Waypoint packet will be processed by an intermediate router that is part of the end to end forwarding path, but is not the destination address of the packet. How is this interception accomplished?

There are two primary solutions that have been implemented in the context of the RSVP protocol. One solution was to define a unique protocol number that is recognized by the router. Packets with this protocol number are not forwarded, but instead delivered to the RSVP protcol processing engine on the router. An alternative approach has been defined using IP options. A router alert IP option, when present in a packet, flags the packet for local delivery within the routers protocol processing engine.

In Waypoint, we define a more general concept for packet interception markings. Rather than a simple flag, Waypoint adopts the notion of a "time to delivery" (TTD) field. At each forwarding router, the TTD field is decremented, similar to the IP TTL field. When the TTD field is zero, the packet is intercepted by the router. It will be quite common to use a TTD field of 1 for many signaling protocols, but the ability to skip over a fixed number of intermediate routers provides the capability to "tunnel through" a sequence of routers when necessary.

One will notice that the TTD field is decremented in an identical fashion to the IP TTL field. Waypoint defines a flag that allows the TTL field to used as the TTD field. This seems like a useful option on many routers. IP routers are already programmed to decrement the IP TTL in the fast forwarding path and to send zero TTL packets to the slow path for an ICMP response to the sender. If the TTL is zero and the IP protocol number is Waypoint, this will cause a local delivery of the packet instead of an ICMP response. Alternatively, a router can use the TTD field in Waypoint, decrement the field (and update the checksum) during forwarding, and deliver packets with a zero TTD to the local protocol processing engine.

The destination port number of the Waypoint packet identifies the particular signaling service that will process a given packet. If there is no service associated with destination port number of a received packet, an ICMP response should be generated.

5.1. Security Options Processing

If a security option is present in the packet, it is processed before the packet is delivered to the signaling service. Currently, Waypoint defines a hop by hop packet integrity option that provides functionality similar to the IPSEC AH header. Because packets are processed at intermediate routers, the key exchange and sharing rules of IPSEC, which are end to end, cannot be applied to Waypoint. We instead adopt the hop by hop integrity solution developed for the RSVP protocol.

Waypoint, unlike RSVP, needs to address confidentiality, as well as authentication. The RSVP integrity solution will be enhanced to include confidentiality for the Waypoint design.

6. Simple Waypoint Examples

The simplest control plane, path oriented, services are measurement rather than signaling operations. This is expected since signaling protocols generally have additional complexity to handle all of the special cases due to errors or faults. Also, signaling protocols usually maintain state and the state maintenance can add complexity. The services mentioned in this section are all stateless.

There is a large class of measurement services, all basically similar, that trace out an end to end path using an expanding ring search with ICMP. Examples include traceroute, pathchar, and mercator[2]. All of these tools attempt to provide as much information that can be gleaned with ICMP responses; IP addresses, operating systems types, and link round trip times.

It is also well known that measurements based on ICMP responses are flawed. Many router implementations assign a low priority to the task of ICMP responses. Measured round trip times can have excessive delay and high variability.

A traceroute service, available on a well known Waypoint port of every router would be an extremely useful service for the Internet. It could provide a more robust service: a complete list of all IP addresses along a path and accurate round trip delays. Removing the limitation of ICMP functionality would allow a Waypoint based implementation to timestamp the response and reply to accurately determine delay.

A pathchar implementation, based on Waypoint, would include link capacity information rather than relying on the tedious task of attempting to determine the capacity based on very noisy measurement data.

7. Complex Signaling Protocols

There are a number of reasonably complex signaling protocols that are in use or being proposed for use in the Internet. RSVP signals for end to end QoS needs along a path. RSVP-TE is used for the set up of traffic engineered tunnels. Another modifications of RSVP is being proposed for optical switch path configuration.

At the heart of all of these protocols, there is a need to deliver control packets at every, or nearly every router along the path. Current mechanisms, such as router alert, provide no ability to separate out signaling packets for different services. As an example, both RSVP and RSVP-TE use IP protocol 46. A router which support both RSVP and RSVP-TE concurrently would have to analyze the packet contents to separate out which packets are being used for which protocol. In this context, Waypoint, with a defined port space, provides a cleaner alternative to the router alert option.

Waypoint does not address all of the common functionality between various signaling protocols. This may include soft state managment, interfaces to routing, and message reliability mechanisms. It is believed that this common functionality, at the transport layer, may lend itself to organization into a set of reusable building blocks. Waypoint only strives to provide common functionality at the intermediate layer between network and transport.

8. Conclusion

Waypoint provides an elementary deliver mechanism for both simple and complex path oriented control, measurement, and signaling protocols. It differs from current mechanisms, such as router alert, in a number of important areas. First, it does not require the use of IP options, which may add additional processing expense on some routers. It provides hop by hop security, enabling signaling packets to have similar security features to IP data packets which can use IPSEC. Unlike the family of RSVP protocols, it provides a distinct port addresses for each new protocol. The time to deliver (TTD) field provides increased delivery control enabling protocols to "tunnel through" a series of routers along a path.

It is believed that the implementation of Waypoint would be straightforward and of low overhead for most router implementations. The ability to use the TTL field as the TTD field should make Waypoint more compatible with existing IP forwarding implementations and only require simple modifications to the ICMP message generation path.

9. References

[1] Braden, R., Ed., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSer-Vation Protocol (RSVP) -- Version 1 Functional Specification". RFC 2205, September 1997.

[2] Govindan R., Tangmunarunkit H., "Heuristics for Internet Map Discovery", Proc IEEE Infocom 2000, Tel Aviv, Israel

[3] Atkinson, R., and S. Kent, "Security Architecture for the Internet Protocol", RFC 2401, November 1998.

[4] Maughan, D., Schertler, M., Schneider, M., and J. Turner, "Internet Security Association and Key Management Protocol (ISAKMP)", RFC 2408, November 1998.

[5] Kent, S., and R. Atkinson, "IP Authentication Header", RFC 2402, November 1998.

[6] Kent, S., and R. Atkinson, "IP Encapsulating Security Payload (ESP)", RFC 2406, November 1998.

10. Security Considerations

To be completed.

11. Authors' Addresses

Bob Lindell
USC Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292
Phone: (310) 448-8727
Email: lindell@ISI.EDU

Bob Braden
USC Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292
Phone: (310) 448-9173
Email: braden@ISI.EDU