

Network Working Group
Request for Comments: 3609
Category: Informational

R. Bonica
MCI
K. Kompella
Juniper Networks
D. Meyer
Sprint
September 2003

Tracing Requirements for Generic Tunnels

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2003). All Rights Reserved.

Abstract

This document specifies requirements for a generic route-tracing application. It also specifies requirements for a protocol that will support that application. Network operators will use the generic route-tracing application to verify proper operation of the IP forwarding plane. They will also use the application to discover details regarding tunnels that support IP forwarding.

The generic route-tracing application, specified herein, supports a superset of the functionality that "traceroute" currently offers. Like traceroute, the generic route-tracing application can discover the forwarding path between two interfaces that are contained by an IP network. Unlike traceroute, this application can reveal details regarding tunnels that support the IP forwarding path.

1. Introduction

IP networks utilize several tunneling technologies. Although these tunneling technologies provide operators with many useful features, they also present management challenges. Network operators require a generic route-tracing application that they can use to verify the correct operation of the IP forwarding plane. The generic route-tracing application must be capable of detecting tunnels and revealing tunnel details. The application also must be useful in diagnosing tunnel faults.

Implementors also require a new protocol that will support the generic-route tracing application. This document specifies requirements for that protocol. It specifies requirements, primarily, by detailing the desired capabilities of the generic route-tracing application. A particular version of generic route-tracing application may implement some subset of the desired capabilities. It may also implement a superset of those capabilities. However, protocol designers are not required to consider the additional capabilities when designing the new protocol.

This document also specifies a few protocol requirements, stated as such. These requirements are driven by desired characteristics of the generic route-tracing application. Whenever a protocol requirement is stated, it is mapped to the desired characteristic of the route-tracing application.

2. Review of Existing Functionality

Currently, network operators use "traceroute" to trace through the forwarding path of an IP network. Section 3.4 of [RFC-2151] provides a thorough description of traceroute. Although traceroute is very reliable and very widely deployed, it is deficient with regard to tunnel tracing.

Depending upon tunnel type, traceroute may display an entire tunnel as a single IP hop, or it may display the tunnel as a collection of IP hops, without indicating that they are part of a tunnel.

For example, assume that engineers deploy an IP tunnel in an IP network. Assume also that they configure the tunnel so that the ingress router does not copy the TTL value from the inner IP header to outer IP header. Instead, the ingress router always sets the outer TTL value to its maximum permitted value. When engineers trace through the network, traceroute will always display the tunnel as a single IP hop, hiding all components except the egress interface.

Now assume that engineers deploy an MPLS LSP in an IP network. Assume also that engineers configure the MPLS LSP so that the ingress router propagates the TTL value from the IP header to the MPLS header. When engineers trace through the network, traceroute will display the LSP as a series of IP hops, without indicating that they are part of a tunnel.

3. Application Requirements

Network operators require a new route-tracing application. The new application must support all functionality that traceroute currently offers. It also must provide enhanced tunnel tracing capabilities.

The following list provides specific requirements for the new route-tracing application:

1) Support the notion of a security token as part of the tunnel trace request. The security token identifies the tracer's privileges in tracing tunnels. Network elements will use this security token to determine whether or not to return the requested information to the tracer. In particular, appropriate privileges are required for items (2), (3), (6), (8), (10), (13), and (14).

Justification: Operators may need to discover network forwarding details, while concealing those details from unauthorized parties.

2) Support in-line traces. An in-line trace reveals the path between the host upon which the route-tracing application executes and any interface in an IP network.

Justification: Operators need to discover how the network would forward a datagram between any two IP interfaces.

3) Support third-party traces. A third-party trace reveals the path between any two points in an IP network. The application that initiates a third-party trace need not execute upon a host or router that is part of the traced path. Unlike existing solutions [RFC-2151] [RFC-2925], the application will not rely upon IP options or require access to the SNMP agent in order to support third-party traces.

Justification: Operators need to discover how the network would forward a datagram between any two IP interfaces.

4) Support partial traces through broken paths or tunnels.

Justification: Operators need to identify the root cause of forwarding plane failures.

5) When tracing through a tunnel, either as part of an in-line trace or a third-party trace, display the tunnel either as a single IP hop or in detail. The user's request determines how the application displays tunnels, subject to the user having permission to do this.

Justification: As they discover IP forwarding details, operators may need to reveal or mask tunneling details.

6) When displaying a tunnel in detail, include the tunnel type (e.g., GRE, MPLS), the tunnel name (if applicable), the tunnel identifier (if applicable) and tunnel endpoint addresses. Also, include tunnel components and round trip delay across each component.

Justification: As they discover IP forwarding details, operators may need to reveal tunneling details.

7) Support the following tunneling technologies: GRE, MPLS, IPSEC, GMPLS, IP-in-IP, L2TP. Be easily extensible to support new tunnel technologies.

Justification: Operators will use the generic route-tracing application to discover how an IP network forwards datagrams. As many tunnel types may support the IP network, the generic route-tracing application must detect and reveal details concerning multiple tunnel types.

8) Trace through nested, heterogeneous tunnels (e.g., IP-in-IP over MPLS).

Justification: Operators will use the generic route-tracing application to discover how an IP network forwards datagrams. As nested, heterogeneous tunnels may support the IP network, the generic route-tracing application must detect and reveal details concerning nested, heterogeneous tunnels.

9) At the users request, trace through the forwarding plane, the control plane or both.

Justification: Operators need to identify the root cause of forwarding plane failures. Control plane information is sometimes useful in determining the cause of forwarding plane failure.

10) Support control plane tracing for all tunnel types. When tracing through the control plane, the hop ingress device reports hop details. The hop ingress device is the device that originates the hop.

Justification: Control plane information is available regarding all tunnel types.

11) Support tracing through forwarding plane for all tunnel types that implement TTL decrement (or some similar mechanism). When tracing through the forwarding plane, the hop egress device reports hop details. The hop egress device is the device that terminates the hop.

Justification: Forwarding plane information may not be available for tunnels that do not support TTL decrement.

12) Support tracing through the forwarding plane for all tunnel types that implement TTL decrement, regardless of whether the tunnel engages in TTL propagation. (That is, support tunnel tracing regardless of whether the TTL value is copied from an inner header to an outer header at tunnel ingress.)

Justification: Forwarding plane information is always available, regardless of whether the tunnel engages in TTL propagation.

13) When tracing through the control plane, display the MTU associated with each interface that forwards datagrams through the traced path.

Justification: MTU information is sometimes useful in identifying the root cause of forwarding and control plane failures.

14) When tracing through the forwarding plane, display the MTU associated with each interface that receives datagrams along the traced path.

Justification: MTU information is sometimes useful in identifying the root cause of forwarding and control plane failures.

15) Support partial traces through paths containing devices that do not provide protocol support for generic route tracing. When the application encounters such a device, it should inform the user and attempt to discover details regarding the next interface downstream.

Justification: The application must provide useful information even if the supporting protocol is not universally deployed.

4. Protocol Requirements

Implementors require a new protocol that supports the generic route-tracing application. This protocol reveals the path between two points in an IP network. When access policy permits, the protocol also reveals tunnel details.

4.1. Information Requirements

The protocol consists of probes and probe responses. Each probe elicits exactly one response. Each response represents a hop that contributes to the path between two interfaces. A hop can be either a top-level IP hop or lower-level hop that is contained by a tunnel.

Justification: Because the generic route-tracing application must trace through broken paths, the required protocol must use a separate response message to deliver details regarding each hop. The protocol must use a separate probe to elicit each response because the alternative approach, using the single probe with the IP Router Alert Option, is unacceptable. Many networks forward datagrams that specify IP options differently than they would forward datagrams that do not specify IP options. Therefore, the introduction of IP options would cause the application to trace a forwarding path other than the path that its user intended to trace.

4.2. Transport Layer Requirements

UDP should carry all protocol messages to their destinations. Other transport mechanisms may be considered when protocol details are specified.

Justification: Because the probe/response scheme described above is stateless, a stateless transport is required. Candidate transports included UDP over IP, IP and ICMP. ICMP was disqualified because carrying MPLS information in an ICMP datagram would constitute a layer violation. IP was disqualified in order to conserve protocol identifiers.

4.3. Stateless Protocol

The protocol must be stateless. That is, nodes should not have to maintain state between successive traceroute messages.

Justification: Statelessness is required to support scaling and to prevent denial of service attacks.

4.4. Routing Requirements

The device that hosts the route-tracing application must maintain an IP route to the ingress of the traced path. It must also maintain an IP route to the ingress of each tunnel for which it is requesting tunnel details. The device that hosts the tunnel tracing application need not maintain a route to any other device that supports the traced path.

All of the devices to which the route-tracing application must maintain a route must maintain a route back to the route-tracing application.

In order for the protocol to provide tunnel details, all devices contained by a tunnel must maintain an IP route to the tunnel ingress.

Justification: The protocol must be sufficiently robust to operate when tunnel interior devices do not maintain a route back to the device that hosts the route tracing application.

5. Security Considerations

A configurable access control policy determines the degree to which features described herein are delivered. The access control policy requires user identification and authorization.

The new protocol must not introduce security holes nor consume excessive resources (e.g., CPU, bandwidth). It also must not be exploitable by those launching DoS attacks or replaying messages.

6. Informative References

[RFC-2151] Kessler, G. and S. Shepard, "A Primer On Internet and TCP/IP Tools and Utilities", FYI 30, RFC 2151, June 1997.

[RFC-2925] White, K., "Definitions of Managed Objects for Remote Ping, Traceroute, and Lookup Operations", RFC 2925, September 2000.

7. Acknowledgements

Thanks to Randy Bush and Steve Bellovin for their comments.

8. Authors' Addresses

Ronald P. Bonica
MCI
22001 Loudoun County Pkwy
Ashburn, Virginia, 20147

EMail: ronald.p.bonica@mci.com

Kireeti Kompella
Juniper Networks, Inc.
1194 N. Mathilda Ave.
Sunnyvale, California 94089

EMail: kireeti@juniper.net

David Meyer

EMail: dmm@maoz.com

9. Full Copyright Statement

Copyright (C) The Internet Society (2003). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assignees.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS 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.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

