Internet Engineering Task Force (IETF)
Request for Comments: 5883
Category: Standards Track
ISSN: 2070-1721

D. Katz
D. Ward
Juniper Networks

June 2010

Bidirectional Forwarding Detection (BFD) for Multihop Paths

Abstract

This document describes the use of the Bidirectional Forwarding Detection (BFD) protocol over multihop paths, including unidirectional links.

Status of This Memo

This is an Internet Standards Track document.

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 Internet Standards 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/rfc5883.

1. Introduction

The Bidirectional Forwarding Detection (BFD) protocol [BFD] defines a method for liveness detection of arbitrary paths between systems. The BFD one-hop specification [BFD-1HOP] describes how to use BFD across single hops of IPv4 and IPv6.

BFD can also be useful on arbitrary paths between systems, which may span multiple network hops and follow unpredictable paths. Furthermore, a pair of systems may have multiple paths between them that may overlap. This document describes methods for using BFD in such scenarios.

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 RFC 2119 [KEYWORDS].

2. Applicability

Please note that BFD is intended as an Operations, Administration, and Maintenance (OAM) mechanism for connectivity check and connection verification. It is applicable for network-based services (e.g. router-to-router, subscriber-to-gateway, LSP/circuit endpoints, and service appliance failure detection). In these scenarios it is required that the operator correctly provision the rates at which BFD is transmitted to avoid congestion (e.g link, I/O, CPU) and false failure detection. It is not applicable for application-to-application failure detection across the Internet because it does not have sufficient capability to do necessary congestion detection and avoidance and therefore cannot prevent congestion collapse. Host-to-host or application-to-application deployment across the Internet will require the encapsulation of BFD within a transport that provides "TCP-friendly" [TFRC] behavior.

3. Issues

There are three primary issues in the use of BFD for multihop paths. The first is security and spoofing; [BFD-1HOP] describes a lightweight method of avoiding spoofing by requiring a Time to Live (TTL)/Hop Limit of 255 on both transmit and receive, but this obviously does not work across multiple hops. The utilization of BFD authentication addresses this issue.

The second, more subtle, issue is that of demultiplexing multiple BFD sessions between the same pair of systems to the proper BFD session. In particular, the first BFD packet received for a session may carry

a Your Discriminator value of zero, resulting in ambiguity as to which session the packet should be associated. Once the discriminator values have been exchanged, all further packets are demultiplexed to the proper BFD session solely by the contents of the Your Discriminator field.

[BFD-1HOP] addresses this by requiring that multiple sessions traverse independent physical or logical links -- the first packet is demultiplexed based on the link over which it was received. In the more general case, this scheme cannot work, as two paths over which BFD is running may overlap to an arbitrary degree (including the first and/or last hop).

Finally, the Echo function MUST NOT be used over multiple hops. Intermediate hops would route the packets back to the sender, and connectivity through the entire path would not be possible to verify.

4. Demultiplexing Packets

There are a number of possibilities for addressing the demultiplexing issue that may be used, depending on the application.

4.1. Totally Arbitrary Paths

It may be desired to use BFD for liveness detection over paths for which no part of the route is known (or if known, may not be stable). A straightforward approach to this problem is to limit BFD deployment to a single session between a source/destination address pair. Multiple sessions between the same pair of systems must have at least one endpoint address distinct from one another.

In this scenario, the initial packet is demultiplexed to the appropriate BFD session based on the source/destination address pair when Your Discriminator is set to zero.

This approach is appropriate for general connectivity detection between systems over routed paths and is also useful for OSPF Virtual Links [OSPFv2] [OSPFv3].

4.2. Out-of-Band Discriminator Signaling

Another approach to the demultiplexing problem is to signal the discriminator values in each direction through an out-of-band mechanism prior to establishing the BFD session. Once learned, the discriminators are sent as usual in the BFD Control packets; no packets with Your Discriminator set to zero are ever sent. This method is used by the BFD MPLS specification [BFD-MPLS].

This approach is advantageous because it allows BFD to be directed by other system components that have knowledge of the paths in use, and from the perspective of BFD implementation it is very simple.

The disadvantage is that it requires at least some level of BFD-specific knowledge in parts of the system outside of BFD.

4.3. Unidirectional Links

Unidirectional links are classified as multihop paths because the return path (which should exist at some level in order to make the link useful) may be arbitrary, and the return paths for BFD sessions protecting parallel unidirectional links may overlap or even be identical. (If two unidirectional links, one in each direction, are to carry a single BFD session, this can be done using the single-hop approach.)

Either of the two methods outlined earlier may be used in the unidirectional link case, but a more general solution can be found strictly within BFD and without addressing limitations.

The approach is similar to the one-hop specification, since the unidirectional link is a single hop. Let’s define the two systems as the Unidirectional Sender and the Unidirectional Receiver. In this approach, the Unidirectional Sender MUST operate in the Active role (as defined in the base BFD specification), and the Unidirectional Receiver MUST operate in the Passive role.

In the Passive role, by definition, the Unidirectional Receiver does not transmit any BFD Control packets until it learns the discriminator value in use by the other system (upon receipt of the first BFD Control packet). The Unidirectional Receiver demultiplexes the first packet to the proper BFD session based on the physical or logical link over which it was received. This allows the receiver to learn the remote discriminator value, which it then echoes back to the sender in its own (arbitrarily routed) BFD Control packet, after which time all packets are demultiplexed solely by discriminator.

5. Encapsulation

The encapsulation of BFD Control packets for multihop application in IPv4 and IPv6 is identical to that defined in [BFD-1HOP], except that the UDP destination port MUST have a value of 4784. This can aid in the demultiplexing and internal routing of incoming BFD packets.

6. Authentication

By their nature, multihop paths expose BFD to spoofing. As the number of hops increases, the exposure to attack grows. As such, implementations of BFD SHOULD utilize cryptographic authentication over multihop paths to help mitigate denial-of-service attacks.

7. IANA Considerations

Port 4784 has been assigned by IANA for use with BFD Multihop Control.

8. Security Considerations

As the number of hops increases, BFD becomes further exposed to attack. The use of strong forms of authentication is strongly encouraged.

No additional security issues are raised in this document beyond those that exist in the referenced BFD documents.

9. References

9.1. Normative References

[BFD] Katz, D. and D. Ward, "Bidirectional Forwarding Detection", RFC 5880, June 2010.

[BFD-1HOP] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD) for IPv4 and IPv6 (Single Hop)", RFC 5881, June 2010.

[KEYWORDS] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

9.2. Informative References

[BFD-MPLS] Aggarwal, R., Kompella, K., Nadeau, T., and G. Swallow, "Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)", RFC 5884, June 2010.

[OSPFv2] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.

[OSPFv3] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, July 2008.

[TFRC] Floyd, S., Handley, M., Padhye, J., and J. Widmer, "TCP Friendly Rate Control (TFRC): Protocol Specification", RFC 5348, September 2008.

Authors' Addresses

Dave Katz
Juniper Networks
1194 N. Mathilda Ave.
Sunnyvale, CA 94089-1206
USA

Phone: +1-408-745-2000
EMail: dkatz@juniper.net

Dave Ward
Juniper Networks
1194 N. Mathilda Ave.
Sunnyvale, CA 94089-1206
USA

Phone: +1-408-745-2000
EMail: dward@juniper.net