

Network Working Group  
Internet Draft  
Intended status: Proposed Standard

Expires: July, 2010

D. Katz  
Juniper Networks  
D. Ward  
Juniper Networks  
January 5, 2010

BFD for Multihop Paths  
draft-ietf-bfd-multihop-09.txt

### Status of this Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/lid-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>

### Copyright Notice

Copyright (c) 2009 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 BSD License.

Internet Draft

BFD for Multihop Paths

January, 2010

## Abstract

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

## 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].

## 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.

## 2. Applicability

Please note that BFD is intended as a connectivity check/connection verification OAM mechanism. 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 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 which 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 done 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 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 increase, the exposure to attack grows. As such, implementations of BFD SHOULD utilize cryptographic authentication over multihop paths to help mitigate denial-of-service attacks.

## Normative References

- [BFD] Katz, D., and Ward, D., "Bidirectional Forwarding Detection", [draft-ietf-bfd-base-10.txt](#), January, 2010.
- [BFD-1HOP] Katz, D., and Ward, D., "BFD for IPv4 and IPv6 (Single Hop)", [draft-ietf-bfd-v4v6-1hop-11.txt](#), January, 2010.
- [KEYWORD] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [RFC 2119](#), March 1997.

## Informative References

- [BFD-MPLS] Aggarwal, R., Kompella, K., et al, "BFD for MPLS LSPs", [draft-ietf-bfd-mpls-07.txt](#), June, 2008.
- [OSPFv2] Moy, J., "OSPF Version 2", [RFC 2328](#), April 1998.
- [OSPFv3] Coltun, R., et al, "OSPF for IPv6", [RFC 2740](#), December 1999.
- [TFRC] Floyd, S., et al, "TCP Friendly Rate Control (TFRC): Protocol Specification", [RFC 5348](#), September, 2008.

## 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.

## IANA Considerations

Port 4784 has been assigned by IANA for use with this protocol.

## Authors' Addresses

Dave Katz  
Juniper Networks  
1194 N. Mathilda Ave.  
Sunnyvale, California 94089-1206 USA  
Phone: +1-408-745-2000  
Email: dkatz@juniper.net

Dave Ward  
Juniper Networks  
1194 N. Mathilda Ave.  
Sunnyvale, California 94089-1206 USA  
Phone: +1-408-745-2000  
Email: dward@juniper.net

## Changes from the previous draft

An applicability section was added. All other changes are editorial in nature.

This document expires in July, 2010.

