Network Working Group
Internet Draft

D. Katz
Juniper Networks
D. Ward
Cisco Systems

Expires: January, 2006

July, 2005

BFD for Multihop Paths
draft-ietf-bfd-multihop-03.txt

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with Section 6 of 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/1id-abstracts.html

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

Copyright Notice

Copyright (C) The Internet Society (2005). All Rights Reserved.

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. Issues

There are two 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 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.)

3. Demultiplexing Packets

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

3.1. 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].

3.2. Out-of-band Discriminator Signalling

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 BFD's perspective 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.

3.3. Unidirectional Links

Unidirectional links are classified as multihop paths because the return path (which must 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 unidirection 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.

4. 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 <TBD>. This can aid in the demultiplexing and internal routing of incoming BFD packets.

5. Authentication

By their nature, multihop paths expose BFD to spoofing. Implementations of BFD SHOULD utilize 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-03.txt, July, 2005.

[BFD-1HOP] Katz, D., and Ward, D., "BFD for IPv4 and IPv6 (Single Hop)", draft-ietf-bfd-v4v6-1hop-03.txt, July, 2005.

[BFD-MPLS] Aggarwal, R., and Kompella, K., "BFD for MPLS LSPs", draft-ietf-bfd-mpls-01.txt, February, 2005.

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

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

[OSPFv3] Coltun, R., et al, "OSPF for IPv6", RFC 2740, December 1999.

Security Considerations

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

IANA Considerations

An additional UDP port number needs to be allocated for use with multihop BFD.

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
Cisco Systems
170 W. Tasman Dr.
San Jose, CA 95134 USA
Phone: +1-408-526-4000
Email: dward@cisco.com

Changes from the previous draft

The only substantive change was the requirement for a new UDP port number for use with multihop BFD. All other changes were editorial in nature.

IPR Disclaimer

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at http://www.ietf.org/ipr.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-

ipr@ietf.org.

Full Copyright Notice

Copyright (C) The Internet Society (2005).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM 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.

This document expires in January, 2006.