

Internet Engineering Task Force (IETF)
Request for Comments: 5881
Category: Standards Track
ISSN: 2070-1721

D. Katz
D. Ward
Juniper Networks
June 2010

Bidirectional Forwarding Detection (BFD)
for IPv4 and IPv6 (Single Hop)

Abstract

This document describes the use of the Bidirectional Forwarding Detection (BFD) protocol over IPv4 and IPv6 for single IP hops.

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/rfc5881>.

Copyright Notice

Copyright (c) 2010 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 Simplified BSD License.

1. Introduction

One very desirable application for Bidirectional Forwarding Detection (BFD) [BFD] is to track IPv4 and IPv6 connectivity between directly connected systems. This could be used to supplement the detection mechanisms in routing protocols or to monitor router-host connectivity, among other applications.

This document describes the particulars necessary to use BFD in this environment. Interactions between BFD and other protocols and system functions are described in the BFD Generic Applications document [BFD-GENERIC].

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. Applications and Limitations

This application of BFD can be used by any pair of systems communicating via IPv4 and/or IPv6 across a single IP hop that is associated with an incoming interface. This includes, but is not limited to, physical media, virtual circuits, and tunnels.

Each BFD session between a pair of systems MUST traverse a separate network-layer path in both directions. This is necessary for demultiplexing to work properly, and also because (by definition) multiple sessions would otherwise be protecting the same path.

If BFD is to be used in conjunction with both IPv4 and IPv6 on a particular path, a separate BFD session MUST be established for each protocol (and thus encapsulated by that protocol) over that link.

If the BFD Echo function is used, transmitted packets are immediately routed back towards the sender on the interface over which they were sent. This may interact with other mechanisms that are used on the two systems that employ BFD. In particular, ingress filtering [BCP38] is incompatible with the way Echo packets need to be sent. Implementations that support the Echo function MUST ensure that ingress filtering is not used on an interface that employs the Echo function or make an exception for ingress filtering Echo packets.

An implementation of the Echo function also requires Application Programming Interfaces (APIs) that may not exist on all systems. A system implementing the Echo function MUST be capable of sending

packets to its own address, which will typically require bypassing the normal forwarding lookup. This typically requires access to APIs that bypass IP-layer functionality.

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. Initialization and Demultiplexing

In this application, there will be only a single BFD session between two systems over a given interface (logical or physical) for a particular protocol. The BFD session must be bound to this interface. As such, both sides of a session MUST take the "Active" role (sending initial BFD Control packets with a zero value of Your Discriminator), and any BFD packet from the remote machine with a zero value of Your Discriminator MUST be associated with the session bound to the remote system, interface, and protocol.

4. Encapsulation

BFD Control packets MUST be transmitted in UDP packets with destination port 3784, within an IPv4 or IPv6 packet. The source port MUST be in the range 49152 through 65535. The same UDP source port number MUST be used for all BFD Control packets associated with a particular session. The source port number SHOULD be unique among all BFD sessions on the system. If more than 16384 BFD sessions are simultaneously active, UDP source port numbers MAY be reused on multiple sessions, but the number of distinct uses of the same UDP source port number SHOULD be minimized. An implementation MAY use the UDP port source number to aid in demultiplexing incoming BFD Control packets, but ultimately the mechanisms in [BFD] MUST be used to demultiplex incoming packets to the proper session.

BFD Echo packets MUST be transmitted in UDP packets with destination UDP port 3785 in an IPv4 or IPv6 packet. The setting of the UDP source port is outside the scope of this specification. The

destination address MUST be chosen in such a way as to cause the remote system to forward the packet back to the local system. The source address MUST be chosen in such a way as to preclude the remote system from generating ICMP or Neighbor Discovery Redirect messages. In particular, the source address SHOULD NOT be part of the subnet bound to the interface over which the BFD Echo packet is being transmitted, and it SHOULD NOT be an IPv6 link-local address, unless it is known by other means that the remote system will not send Redirects.

BFD Echo packets MUST be transmitted in such a way as to ensure that they are received by the remote system. On multiaccess media, for example, this requires that the destination datalink address corresponds to the remote system.

The above requirements may require the bypassing of some common IP layer functionality, particularly in host implementations.

5. TTL/Hop Limit Issues

If BFD authentication is not in use on a session, all BFD Control packets for the session MUST be sent with a Time to Live (TTL) or Hop Limit value of 255. All received BFD Control packets that are demultiplexed to the session MUST be discarded if the received TTL or Hop Limit is not equal to 255. A discussion of this mechanism can be found in [GTSM].

If BFD authentication is in use on a session, all BFD Control packets MUST be sent with a TTL or Hop Limit value of 255. All received BFD Control packets that are demultiplexed to the session MAY be discarded if the received TTL or Hop Limit is not equal to 255. If the TTL/Hop Limit check is made, it MAY be done before any cryptographic authentication takes place if this will avoid unnecessary calculation that would be detrimental to the receiving system.

In the context of this section, "authentication in use" means that the system is sending BFD Control packets with the Authentication bit set and with the Authentication Section included and that all unauthenticated packets demultiplexed to the session are discarded, per the BFD base specification.

6. Addressing Issues

Implementations MUST ensure that all BFD Control packets are transmitted over the one-hop path being protected by BFD.

On a multiaccess network, BFD Control packets MUST be transmitted with source and destination addresses that are part of the subnet (addressed from and to interfaces on the subnet).

On a point-to-point link, the source address of a BFD Control packet MUST NOT be used to identify the session. This means that the initial BFD packet MUST be accepted with any source address, and that subsequent BFD packets MUST be demultiplexed solely by the Your Discriminator field (as is always the case). This allows the source address to change if necessary. If the received source address changes, the local system MUST NOT use that address as the destination in outgoing BFD Control packets; rather, it MUST continue to use the address configured at session creation. An implementation MAY notify the application that the neighbor's source address has changed, so that the application might choose to change the destination address or take some other action. Note that the TTL/Hop Limit check described in section 5 (or the use of authentication) precludes the BFD packets from having come from any source other than the immediate neighbor.

7. BFD for Use with Tunnels

A number of mechanisms are available to tunnel IPv4 and IPv6 over arbitrary topologies. If the tunnel mechanism does not decrement the TTL or Hop Limit of the network protocol carried within, the mechanism described in this document may be used to provide liveness detection for the tunnel. The BFD authentication mechanism SHOULD be used and is strongly encouraged.

8. IANA Considerations

Ports 3784 and 3875 were assigned by IANA for use with the BFD Control and BFD Echo protocols, respectively.

9. Security Considerations

In this application, the use of TTL=255 on transmit and receive, coupled with an association to an incoming interface, is viewed as supplying equivalent security characteristics to other protocols used in the infrastructure, as it is not trivially spoofable. The security implications of this mechanism are further discussed in [GTSM].

The security implications of the use of BFD authentication are discussed in [BFD].

The use of the TTL=255 check simultaneously with BFD authentication provides a low overhead mechanism for discarding a class of unauthorized packets and may be useful in implementations in which cryptographic checksum use is susceptible to denial-of-service attacks. The use or non-use of this mechanism does not impact interoperability.

10. References

10.1. Normative References

- [BFD] Katz, D. and D. Ward, "Bidirectional Forwarding Detection", RFC 5880, June 2010.
- [BFD-GENERIC] Katz, D. and D. Ward, "Generic Application of Bidirectional Forwarding Detection (BFD)", RFC 5882, June 2010.
- [GTSM] Gill, V., Heasley, J., Meyer, D., Savola, P., Ed., and C. Pignataro, "The Generalized TTL Security Mechanism (GTSM)", RFC 5082, October 2007.
- [KEYWORDS] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

10.2. Informative References

- [BCP38] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 2000.
- [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