

Internet Engineering Task Force (IETF)
Request for Comments: 5882
Category: Standards Track
ISSN: 2070-1721

D. Katz
D. Ward
Juniper Networks
June 2010

Generic Application of Bidirectional Forwarding Detection (BFD)

Abstract

This document describes the generic application of the Bidirectional Forwarding Detection (BFD) protocol.

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/rfc5882>.

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.

Table of Contents

1. Introduction	3
1.1. Conventions Used in This Document	3
2. Overview	3
3. Basic Interaction between BFD Sessions and Clients	4
3.1. Session State Hysteresis	4
3.2. AdminDown State	5
3.3. Hitless Establishment/Reestablishment of BFD State	5
4. Control Protocol Interactions	6
4.1. Adjacency Establishment	6
4.2. Reaction to BFD Session State Changes	7
4.2.1. Control Protocols with a Single Data Protocol	7
4.2.2. Control Protocols with Multiple Data Protocols	8
4.3. Interactions with Graceful Restart Mechanisms	8
4.3.1. BFD Fate Independent of the Control Plane	9
4.3.2. BFD Shares Fate with the Control Plane	9
4.4. Interactions with Multiple Control Protocols	10
5. Interactions with Non-Protocol Functions	10
6. Data Protocols and Demultiplexing	11
7. Multiple Link Subnetworks	11
7.1. Complete Decoupling	12
7.2. Layer N-1 Hints	12
7.3. Aggregating BFD Sessions	12
7.4. Combinations of Scenarios	12
8. Other Application Issues	13
9. Interoperability Issues	13
10. Specific Protocol Interactions (Non-Normative)	13
10.1. BFD Interactions with OSPFv2, OSPFv3, and IS-IS	14
10.1.1. Session Establishment	14
10.1.2. Reaction to BFD State Changes	14
10.1.3. OSPF Virtual Links	15
10.2. Interactions with BGP	15
10.3. Interactions with RIP	15
11. Security Considerations	16
12. References	16
12.1. Normative References	16
12.2. Informative References	16

1. Introduction

The Bidirectional Forwarding Detection [BFD] protocol provides a liveness detection mechanism that can be utilized by other network components for which their integral liveness mechanisms are either too slow, inappropriate, or nonexistent. Other documents have detailed the use of BFD with specific encapsulations ([BFD-1HOP] [BFD-MULTI] [BFD-MPLS]). As the utility of BFD has become understood, there have been calls to specify BFD interactions with a growing list of network functions. Rather than producing a long series of short documents on the application of BFD, it seemed worthwhile to describe the interactions between BFD and other network functions ("BFD clients") in a broad way.

This document describes the generic application of BFD. Specific protocol applications are provided for illustrative purposes.

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. Overview

The Bidirectional Forwarding Detection (BFD) specification defines a protocol with simple and specific semantics. Its sole purpose is to verify connectivity between a pair of systems, for a particular data protocol across a path (which may be of any technology, length, or OSI layer). The promptness of the detection of a path failure can be controlled by trading off protocol overhead and system load with detection times.

BFD is **not** intended to directly provide control protocol liveness information; those protocols have their own means and vagaries. Rather, control protocols can use the services provided by BFD to inform their operation. BFD can be viewed as a service provided by the layer in which it is running.

The service interface with BFD is straightforward. The application supplies session parameters (neighbor address, time parameters, protocol options), and BFD provides the session state, of which the most interesting transitions are to and from the Up state. The application is expected to bootstrap the BFD session, as BFD has no discovery mechanism.

An implementation SHOULD establish only a single BFD session per data protocol path, regardless of the number of applications that wish to utilize it. There is no additional value in having multiple BFD sessions to the same endpoints. If multiple applications request different session parameters, it is a local issue as to how to resolve the parameter conflicts. BFD in turn will notify all applications bound to a session when a session state change occurs.

BFD should be viewed as having an advisory role to the protocol or protocols or other network functions with which it is interacting, which will then use their own mechanisms to effect any state transitions. The interaction is very much at arm's length, which keeps things simple and decoupled. In particular, BFD explicitly does not carry application-specific information, partly for architectural reasons and partly because BFD may have curious and unpredictable latency characteristics and as such makes a poor transport mechanism.

It is important to remember that the interaction between BFD and its client applications has essentially no interoperability issues, because BFD is acting in an advisory nature (similar to hardware signaling the loss of light on a fiber optic circuit, for example) and existing mechanisms in the client applications are used in reaction to BFD events. In fact, BFD may interact with only one of a pair of systems for a particular client application without any ill effect.

3. Basic Interaction between BFD Sessions and Clients

The interaction between a BFD session and its associated client applications is, for the most part, an implementation issue that is outside the scope of this specification. However, it is useful to describe some mechanisms that implementors may use in order to promote full-featured implementations. One way of modeling this interaction is to create an adaptation layer between the BFD state machine and the client applications. The adaptation layer is cognizant of both the internals of the BFD implementation and the requirements of the clients.

3.1. Session State Hysteresis

A BFD session can be tightly coupled to its client applications; for example, any transition out of the Up state could cause signaling to the clients to take failure action. However, in some cases, this may not always be the best course of action.

Implementors may choose to hide rapid Up/Down/Up transitions of the BFD session from its clients. This is useful in order to support process restarts without necessitating complex protocol mechanisms, for example.

As such, a system MAY choose not to notify clients if a BFD session transitions from Up to Down state, and returns to Up state, if it does so within a reasonable period of time (the length of which is outside the scope of this specification). If the BFD session does not return to Up state within that time frame, the clients SHOULD be notified that a session failure has occurred.

3.2. AdminDown State

The AdminDown mechanism in BFD is intended to signal that the BFD session is being taken down for administrative purposes, and the session state is not indicative of the liveness of the data path.

Therefore, a system SHOULD NOT indicate a connectivity failure to a client if either the local session state or the remote session state (if known) transitions to AdminDown, so long as that client has independent means of liveness detection (typically, control protocols).

If a client does not have any independent means of liveness detection, a system SHOULD indicate a connectivity failure to a client, and assume the semantics of Down state, if either the local or remote session state transitions to AdminDown. Otherwise, the client will not be able to determine whether the path is viable, and unfortunate results may occur.

3.3. Hitless Establishment/Reestablishment of BFD State

It is useful to be able to configure a BFD session between a pair of systems without impacting the state of any clients that will be associated with that session. Similarly, it is useful for BFD state to be reestablished without perturbing associated clients when all BFD state is lost (such as in process restart situations). This interacts with the clients' ability to establish their state independent of BFD.

The BFD state machine transitions that occur in the process of bringing up a BFD session in such situations SHOULD NOT cause a connectivity failure notification to the clients.

A client that is capable of establishing its state prior to the configuration or restarting of a BFD session MAY do so if appropriate. The means to do so is outside of the scope of this specification.

4. Control Protocol Interactions

Very common client applications of BFD are control protocols, such as routing protocols. The object, when BFD interacts with a control protocol, is to advise the control protocol of the connectivity of the data protocol. In the case of routing protocols, for example, this allows the connectivity failure to trigger the rerouting of traffic around the failed path more quickly than the native detection mechanisms.

4.1. Adjacency Establishment

If the session state on either the local or remote system (if known) is AdminDown, BFD has been administratively disabled, and the establishment of a control protocol adjacency MUST be allowed.

BFD sessions are typically bootstrapped by the control protocol, using the mechanism (discovery, configuration) used by the control protocol to find neighbors. Note that it is possible in some failure scenarios for the network to be in a state such that the control protocol is capable of coming up, but the BFD session cannot be established, and, more particularly, data cannot be forwarded. To avoid this situation, it would be beneficial not to allow the control protocol to establish a neighbor adjacency. However, this would preclude the operation of the control protocol in an environment in which not all systems support BFD.

Therefore, the establishment of control protocol adjacencies SHOULD be blocked if both systems are willing to establish a BFD session but a BFD session cannot be established. One method for determining that both systems are willing to establish a BFD session is if the control protocol carries explicit signaling of this fact. If there is no explicit signaling, the willingness to establish a BFD session may be determined by means outside the scope of this specification.

If it is believed that the neighboring system does not support BFD, the establishment of a control protocol adjacency SHOULD NOT be blocked.

The setting of BFD's various timing parameters and modes are not subject to standardization. Note that all protocols sharing a session will operate using the same parameters. The mechanism for choosing the parameters among those desired by the various protocols

is outside the scope of this specification. It is generally useful to choose the parameters resulting in the shortest Detection Time; a particular client application can always apply hysteresis to the notifications from BFD if it desires longer Detection Times.

Note that many control protocols assume full connectivity between all systems on multiaccess media such as LANs. If BFD is running on only a subset of systems on such a network, and adjacency establishment is blocked by the absence of a BFD session, the assumptions of the control protocol may be violated, with unpredictable results.

4.2. Reaction to BFD Session State Changes

If a BFD session transitions from Up state to AdminDown, or the session transitions from Up to Down because the remote system is indicating that the session is in state AdminDown, clients SHOULD NOT take any control protocol action.

For other transitions from Up to Down state, the mechanism by which the control protocol reacts to a path failure signaled by BFD depends on the capabilities of the protocol, as specified in the following subsections.

4.2.1. Control Protocols with a Single Data Protocol

A control protocol that is tightly bound to a single failing data protocol SHOULD take action to ensure that data traffic is no longer directed to the failing path. Note that this should not be interpreted as BFD replacing the control protocol liveness mechanism, if any, as the control protocol may rely on mechanisms not verified by BFD (multicast, for instance) so BFD most likely cannot detect all failures that would impact the control protocol. However, a control protocol MAY choose to use BFD session state information to more rapidly detect an impending control protocol failure, particularly if the control protocol operates in-band (over the data protocol).

Therefore, when a BFD session transitions from Up to Down, action SHOULD be taken in the control protocol to signal the lack of connectivity for the path over which BFD is running. If the control protocol has an explicit mechanism for announcing path state, a system SHOULD use that mechanism rather than impacting the connectivity of the control protocol, particularly if the control protocol operates out-of-band from the failed data protocol. However, if such a mechanism is not available, a control protocol timeout SHOULD be emulated for the associated neighbor.

4.2.2. Control Protocols with Multiple Data Protocols

Slightly different mechanisms are used if the control protocol supports the routing of multiple data protocols, depending on whether the control protocol supports separate topologies for each data protocol.

4.2.2.1. Shared Topologies

With a shared topology, if one of the data protocols fails (as signaled by the associated BFD session), it is necessary to consider the path to have failed for all data protocols. Otherwise, there is no way for the control protocol to turn away traffic for the failed data protocol (and such traffic would be black-holed indefinitely).

Therefore, when a BFD session transitions from Up to Down, action SHOULD be taken in the control protocol to signal the lack of connectivity for the path in the topology corresponding to the BFD session. If this cannot be signaled otherwise, a control protocol timeout SHOULD be emulated for the associated neighbor.

4.2.2.2. Independent Topologies

With individual routing topologies for each data protocol, only the failed data protocol needs to be rerouted around the failed path.

Therefore, when a BFD session transitions from Up to Down, action SHOULD be taken in the control protocol to signal the lack of connectivity for the path in the topology over which BFD is running. Generally, this can be done without impacting the connectivity of other topologies (since otherwise it is very difficult to support separate topologies for multiple data protocols).

4.3. Interactions with Graceful Restart Mechanisms

A number of control protocols support Graceful Restart mechanisms, including IS-IS [ISIS-GRACE], OSPF [OSPF-GRACE], and BGP [BGP-GRACE]. These mechanisms are designed to allow a control protocol to restart without perturbing network connectivity state (lest it appear that the system and/or all of its links had failed). They are predicated on the existence of a separate forwarding plane that does not necessarily share fate with the control plane in which the routing protocols operate. In particular, the assumption is that the forwarding plane can continue to function while the protocols restart and sort things out.

BFD implementations announce via the Control Plane Independent "C" bit whether or not BFD shares fate with the control plane. This

information is used to determine the actions to be taken in conjunction with Graceful Restart. If BFD does not share its fate with the control plane on either system, it can be used to determine whether Graceful Restart in a control protocol is NOT viable (the forwarding plane is not operating).

If the control protocol has a Graceful Restart mechanism, BFD may be used in conjunction with this mechanism. The interaction between BFD and the control protocol depends on the capabilities of the control protocol and whether or not BFD shares fate with the control plane. In particular, it may be desirable for a BFD session failure to abort the Graceful Restart process and allow the failure to be visible to the network.

4.3.1. BFD Fate Independent of the Control Plane

If BFD is implemented in the forwarding plane and does not share fate with the control plane on either system (the "C" bit is set in the BFD Control packets in both directions), control protocol restarts should not affect the BFD session. In this case, a BFD session failure implies that data can no longer be forwarded, so any Graceful Restart in progress at the time of the BFD session failure SHOULD be aborted in order to avoid black holes, and a topology change SHOULD be signaled in the control protocol.

4.3.2. BFD Shares Fate with the Control Plane

If BFD shares fate with the control plane on either system (the "C" bit is clear in either direction), a BFD session failure cannot be disentangled from other events taking place in the control plane. In many cases, the BFD session will fail as a side effect of the restart taking place. As such, it would be best to avoid aborting any Graceful Restart taking place, if possible (since otherwise BFD and Graceful Restart cannot coexist).

There is some risk in doing so, since a simultaneous failure or restart of the forwarding plane will not be detected, but this is always an issue when BFD shares fate with the control plane.

4.3.2.1. Control Protocols with Planned Restart Signaling

Some control protocols can signal a planned restart prior to the restart taking place. In this case, if a BFD session failure occurs during the restart, such a planned restart SHOULD NOT be aborted and the session failure SHOULD NOT result in a topology change being signaled in the control protocol.

4.3.2.2. Control Protocols without Planned Restart Signaling

Control protocols that cannot signal a planned restart depend on the recently restarted system to signal the Graceful Restart prior to the control protocol adjacency timeout. In most cases, whether the restart is planned or unplanned, it is likely that the BFD session will time out prior to the onset of Graceful Restart, in which case a topology change SHOULD be signaled in the control protocol as specified in Section 3.2.

However, if the restart is in fact planned, an implementation MAY adjust the BFD session timing parameters prior to restarting in such a way that the Detection Time in each direction is longer than the restart period of the control protocol, providing the restarting system the same opportunity to enter Graceful Restart as it would have without BFD. The restarting system SHOULD NOT send any BFD Control packets until there is a high likelihood that its neighbors know a Graceful Restart is taking place, as the first BFD Control packet will cause the BFD session to fail.

4.4. Interactions with Multiple Control Protocols

If multiple control protocols wish to establish BFD sessions with the same remote system for the same data protocol, all MUST share a single BFD session.

If hierarchical or dependent layers of control protocols are in use (say, OSPF and Internal BGP (IBGP)), it may not be useful for more than one of them to interact with BFD. In this example, because IBGP is dependent on OSPF for its routing information, the faster failure detection relayed to IBGP may actually be detrimental. The cost of a peer state transition is high in BGP, and OSPF will naturally heal the path through the network if it were to receive the failure detection.

In general, it is best for the protocol at the lowest point in the hierarchy to interact with BFD, and then to use existing interactions between the control protocols to effect changes as necessary. This will provide the fastest possible failure detection and recovery in a network.

5. Interactions with Non-Protocol Functions

BFD session status may be used to affect other system functions that are not protocol based (for example, static routes). If the path to a remote system fails, it may be desirable to avoid passing traffic

to that remote system, so the local system may wish to take internal measures to accomplish this (such as withdrawing a static route and withdrawing that route from routing protocols).

If it is known, or presumed, that the remote system is BFD capable and the BFD session is not in Up state, appropriate action SHOULD be taken (such as withdrawing a static route).

If it is known, or presumed, that the remote system does not support BFD, action such as withdrawing a static route SHOULD NOT be taken.

Bootstrapping of the BFD session in the non-protocol case is likely to be derived from configuration information.

There is no need to exchange endpoints or discriminator values via any mechanism other than configuration (via Operational Support Systems or any other means) as the endpoints must be known and configured by the same means.

6. Data Protocols and Demultiplexing

BFD is intended to protect a single "data protocol" and is encapsulated within that protocol. A pair of systems may have multiple BFD sessions over the same topology if they support (and are encapsulated by) different protocols. For example, if two systems have IPv4 and IPv6 running across the same link between them, these are considered two separate paths and require two separate BFD sessions.

This same technique is used for more fine-grained paths. For example, if multiple differentiated services [DIFFSERV] are being operated over IPv4, an independent BFD session may be run for each service level. The BFD Control packets must be marked in the same way as the data packets, partly to ensure as much fate sharing as possible between BFD and data traffic, and also to demultiplex the initial packet if the discriminator values have not been exchanged.

7. Multiple Link Subnetworks

A number of technologies exist for aggregating multiple parallel links at layer N-1 and treating them as a single link at layer N. BFD may be used in a number of ways to protect the path at layer N. The exact mechanism used is outside the scope of this specification. However, this section provides examples of some possible deployment scenarios. Other scenarios are possible and are not precluded.

7.1. Complete Decoupling

The simplest approach is to simply run BFD over the layer N path, with no interaction with the layer N-1 mechanisms. Doing so assumes that the layer N-1 mechanism will deal with connectivity issues in individual layer N-1 links. BFD will declare a failure in the layer N path only when the session times out.

This approach will work whether or not the layer N-1 neighbor is the same as the layer N neighbor.

7.2. Layer N-1 Hints

A slightly more intelligent approach than complete decoupling is to have the layer N-1 mechanism inform the layer N BFD when the aggregated link is no longer viable. In this case, the BFD session will detect the failure more rapidly, as it need not wait for the session to time out. This is analogous to triggering a session failure based on the hardware-detected failure of a single link.

This approach will also work whether or not the layer N-1 neighbor is the same as the layer N neighbor.

7.3. Aggregating BFD Sessions

Another approach would be to use BFD on each layer N-1 link and to aggregate the state of the multiple sessions into a single indication to the layer N clients. This approach has the advantage that it is independent of the layer N-1 technology. However, this approach only works if the layer N neighbor is the same as the layer N-1 neighbor (a single hop at layer N-1).

7.4. Combinations of Scenarios

Combinations of more than one of the scenarios listed above (or others) may be useful in some cases. For example, if the layer N neighbor is not directly connected at layer N-1, a system might run a BFD session across each layer N-1 link to the immediate layer N-1 neighbor and then run another BFD session to the layer N neighbor. The aggregate state of the layer N-1 BFD sessions could be used to trigger a layer N BFD session failure.

8. Other Application Issues

BFD can provide liveness detection for functions related to Operations, Administration, and Maintenance (OAM) in tunneling and pseudowire protocols. Running BFD inside the tunnel is recommended, as it exercises more aspects of the path. One way to accommodate this is to address BFD packets based on the tunnel endpoints, assuming that they are numbered.

If a planned outage is to take place on a path over which BFD is run, it is preferable to take down the BFD session by going into AdminDown state prior to the outage. The system asserting AdminDown SHOULD do so for at least one Detection Time in order to ensure that the remote system is aware of it.

Similarly, if BFD is to be deconfigured from a system, it is desirable not to trigger any client application action. Simply ceasing the transmission of BFD Control packets will cause the remote system to detect a session failure. In order to avoid this, the system on which BFD is being deconfigured SHOULD put the session into AdminDown state and maintain this state for a Detection Time to ensure that the remote system is aware of it.

9. Interoperability Issues

The BFD protocol itself is designed so that it will always interoperate at a basic level; asynchronous mode is mandatory and is always available, and other modes and functions are negotiated at run time. Since the service provided by BFD is identical regardless of the variants used, the particular choice of BFD options has no bearing on interoperability.

The interaction between BFD and other protocols and control functions is very loosely coupled. The actions taken are based on existing mechanisms in those protocols and functions, so interoperability problems are very unlikely unless BFD is applied in contradictory ways (such as a BFD session failure causing one implementation to go down and another implementation to come up). In fact, BFD may be advising one system for a particular control function but not the other; the only impact of this would be potentially asymmetric control protocol failure detection.

10. Specific Protocol Interactions (Non-Normative)

As noted above, there are no interoperability concerns regarding interactions between BFD and control protocols. However, there is enough concern and confusion in this area so that it is worthwhile to provide examples of interactions with specific protocols.

Since the interactions do not affect interoperability, they are non-normative.

10.1. BFD Interactions with OSPFv2, OSPFv3, and IS-IS

The two versions of OSPF ([OSPFv2] and [OSPFv3]), as well as IS-IS [ISIS], all suffer from an architectural limitation, namely that their Hello protocols are limited in the granularity of their failure detection times. In particular, OSPF has a minimum detection time of two seconds, and IS-IS has a minimum detection time of one second.

BFD may be used to achieve arbitrarily small detection times for these protocols by supplementing the Hello protocols used in each case.

10.1.1. Session Establishment

The most obvious choice for triggering BFD session establishment with these protocols would be to use the discovery mechanism inherent in the Hello protocols in OSPF and IS-IS to bootstrap the establishment of the BFD session. Any BFD sessions established to support OSPF and IS-IS across a single IP hop must operate in accordance with [BFD-1HOP].

10.1.2. Reaction to BFD State Changes

The basic mechanisms are covered in Section 3 above. At this time, OSPFv2 and OSPFv3 carry routing information for a single data protocol (IPv4 and IPv6, respectively) so when it is desired to signal a topology change after a BFD session failure, this should be done by tearing down the corresponding OSPF neighbor.

IS-IS may be used to support only one data protocol, or multiple data protocols. [ISIS] specifies a common topology for multiple data protocols, but work is under way to support multiple topologies. If multiple topologies are used to support multiple data protocols (or multiple classes of service of the same data protocol), the topology-specific path associated with a failing BFD session should no longer be advertised in IS-IS Label Switched Paths (LSPs) in order to signal a lack of connectivity. Otherwise, a failing BFD session should be signaled by simulating an IS-IS adjacency failure.

OSPF has a planned restart signaling mechanism, whereas IS-IS does not. The appropriate mechanisms outlined in Section 3.3 should be used.

10.1.3. OSPF Virtual Links

If it is desired to use BFD for failure detection of OSPF Virtual Links, the mechanism described in [BFD-MULTI] MUST be used, since OSPF Virtual Links may traverse an arbitrary number of hops. BFD authentication SHOULD be used and is strongly encouraged.

10.2. Interactions with BGP

BFD may be useful with External Border Gateway Protocol (EBGP) sessions [BGP] in order to more rapidly trigger topology changes in the face of path failure. As noted in Section 4.4, it is generally unwise for IBGP sessions to interact with BFD if the underlying IGP is already doing so.

EBGP sessions being advised by BFD may establish either a one-hop [BFD-1HOP] or a multihop [BFD-MULTI] session, depending on whether or not the neighbor is immediately adjacent. The BFD session should be established to the BGP neighbor (as opposed to any other Next Hop advertised in BGP). BFD authentication SHOULD be used and is strongly encouraged.

[BGP-GRACE] describes a Graceful Restart mechanism for BGP. If Graceful Restart is not taking place on an EBGP session, and the corresponding BFD session fails, the EBGP session should be torn down in accordance with Section 3.2. If Graceful Restart is taking place, the basic procedures in Section 4.3 apply. BGP Graceful Restart does not signal planned restarts, so Section 4.3.2.2 applies. If Graceful Restart is aborted due to the rules described in Section 4.3, the "receiving speaker" should act as if the "restart timer" expired (as described in [BGP-GRACE]).

10.3. Interactions with RIP

The Routing Information Protocol (RIP) [RIP] is somewhat unique in that, at least as specified, neighbor adjacency state is not stored per se. Rather, installed routes contain a next hop address, which in most cases is the address of the advertising neighbor (but may not be).

In the case of RIP, when the BFD session associated with a neighbor fails, an expiration of the "timeout" timer for each route installed from the neighbor (for which the neighbor is the next hop) should be simulated.

Note that if a BFD session fails, and a route is received from that neighbor with a next hop address that is not the address of the neighbor itself, the route will linger until it naturally times out

(after 180 seconds). However, if an implementation keeps track of all of the routes received from each neighbor, all of the routes from the neighbor corresponding to the failed BFD session should be timed out, regardless of the next hop specified therein, and thereby avoiding the lingering route problem.

11. Security Considerations

This specification does not raise any additional security issues beyond those of the specifications referred to in the list of normative references.

12. References

12.1. 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.
- [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.
- [BFD-MULTI] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD) for Multihop Paths", RFC 5883, June 2010.
- [KEYWORDS] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

12.2. Informative References

- [BGP] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.
- [BGP-GRACE] Sangli, S., Chen, E., Fernando, R., Scudder, J., and Y. Rekhter, "Graceful Restart Mechanism for BGP", RFC 4724, January 2007.
- [DIFFSERV] Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, December 1998.

- [ISIS] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", RFC 1195, December 1990.
- [ISIS-GRACE] Shand, M. and L. Ginsberg, "Restart Signaling for IS-IS", RFC 5306, October 2008.
- [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.
- [OSPF-GRACE] Moy, J., Pillay-Esnault, P., and A. Lindem, "Graceful OSPF Restart", RFC 3623, November 2003.
- [RIP] Malkin, G., "RIP Version 2", STD 56, RFC 2453, November 1998.

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

