Workgroup: Network Working Group
Internet-Draft: draft-ietf-bfd-unsolicited-10
Updates: 9314 (if approved)
Published: 24 October 2022
Intended Status: Standards Track
Expires: 27 April 2023
Authors: E. Chen (Palo Alto Networks)
N. Shen (Zededa)
R. Raszuk (Arrcus)
R. Rahman (Graphiant)

Unsolicited BFD for Sessionless Applications

Abstract

For operational simplification of "sessionless" applications using Bidirectional Forwarding Detection (BFD), in this document we present procedures for "unsolicited BFD" that allow a BFD session to be initiated by only one side, and established without explicit per-session configuration or registration by the other side (subject to certain per-interface or global policies).

We also introduce a new YANG module to configure and manage "unsolicited BFD". The YANG module in this document is based on YANG 1.1 as defined in RFC 7950 and conforms to the Network Management Datastore Architecture (NMDA) as described in RFC 8342.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

Status of This Memo

Table of Contents

1. Introduction

The current implementation and deployment practice for BFD ([RFC5880] and [RFC5881]) usually requires BFD sessions be explicitly configured or registered on both sides. This requirement is not an issue when an application like BGP [RFC4271] has the concept of a "session" that involves both sides for its establishment. However, this requirement can be operationally challenging when the prerequisite "session" does not naturally exist between two endpoints in an application. Simultaneous configuration and coordination may be required on both sides for BFD to take effect. For example:

* When BFD is used to keep track of the "liveness" of the nexthop of static routes. Although only one side may need the BFD functionality, currently both sides need to be involved in

specific configuration and coordination and in some cases static routes are created unnecessarily just for BFD.

* When BFD is used to keep track of the "liveness" of the third-pary nexthop of BGP routes received from the Route Server [RFC7947] at an Internet Exchange Point (IXP). As the third-party nexthop is different from the peering address of the Route Server, for BFD to work, currently two routers peering with the Route Server need to have routes and nexthops from each other (although indirectly via the Router Server). These issues are also discussed in [I-D.ietf-idr-rs-bfd].

Clearly it is beneficial and desirable to reduce or eliminate unnecessary configurations and coordination in these "sessionless" applications using BFD.

In this document we present procedures for "unsolicited BFD" that allow a BFD session to be initiated by only one side, and established without explicit per-session configuration or registration by the other side (subject to certain per-interface or global policies).

Unsolicited BFD impacts only the initiation of BFD sessions. There is no change to all the other procedures specified in [RFC5880] such as, but not limited to, the Echo function and Demand mode.

With "unsolicited BFD" there is potential risk for excessive resource usage by BFD from "unexpected" remote systems. To mitigate such risks, several mechanisms are recommended in the Security Considerations section.

The procedure described in this document could be applied to BFD for Multihop paths [RFC5883]. However, because of security risks, this document applies only to BFD for single IP hops [RFC5881].

Compared to the "Seamless BFD" [RFC7880], this proposal involves only minor procedural enhancements to the widely deployed BFD itself. Thus we believe that this proposal is inherently simpler in the protocol itself and deployment. As an example, it does not require the exchange of BFD discriminators over an out-of-band channel before BFD session bring-up.

When BGP Add-Path [RFC7911] is deployed at an IXP using a Route Server, multiple BGP paths (when they exist) can be made available to the clients of the Router Server as described in [RFC7947]. The "unsolicited BFD" can be used in BGP route selection by these clients to eliminate paths with "inaccessible nexthops".

2. Procedures for Unsolicited BFD

With "unsolicited BFD", one side takes the "Active role" and the other side takes only the "Passive role" as described in [RFC5880], section 6.1.

Passive unsolicited BFD support MUST be disabled by default, and MUST require explicit configuration to be enabled. On the passive side, the desired BFD parameters SHOULD be configurable. The passive side MAY also choose to use the parameters that the active side uses in its BFD Control packets. The "My Discriminator", however, MUST be chosen to allow multiple unsolicited BFD sessions.

The active side starts sending the BFD Control packets as specified in [RFC5880]. The passive side does not send BFD Control packets initially, it sends BFD Control packets only after it has received BFD Control packets from the active side.

When the passive side receives a BFD Control packet from the active side with 0 as "Your Discriminator" and does not find an existing BFD session, the passive side MAY create a matching BFD session toward the active side, if permitted by local configuration and policy.

When the passive side receives an incoming BFD Control packet on a numbered interfaces, the source address of that packet MUST belong to the subnet of the interface on which the BFD packet is received. The source address of the BFD Control packet SHOULD be validated against expected routing protocol peer addresses on that interface.

The passive side MUST then start sending BFD Control packets and perform the necessary procedure for bringing up, maintaining and tearing down the BFD session. If the BFD session fails to get established within certain specified time, or if an established BFD session goes down, the passive side SHOULD stop sending BFD Control packets and MAY delete the BFD session created until BFD Control packets are initiated by the active side again.

When an Unsolicited BFD session goes down, an implementation MAY retain the session state for a period of time. Retaining this state can be useful for operational purposes.

The "Passive role" may change to the "Active role" when a local client registers for the same BFD session, and from the "Active role" to the "Passive role" when there is no longer any locally registered client for the BFD session.

3. State Variables

This document defines a new state variable called Role.

bfd.Role

The role of the BFD session as per [RFC5880], section 6.1. Possible values are Active or Passive.

4. YANG Data Model

This section extends the YANG data model for BFD [RFC9314] to cover unsolicited BFD. The new module imports [RFC8349] since the "bfd" container in [RFC9314] is under "control-plane-protocol".

4.1. Unsolicited BFD Hierarchy

Configuration for unsolicited BFD parameters for IP single-hop sessions can be done at 2 levels:

* Globally, i.e. for all interfaces. This requires support for the "unsolicited-params-global" feature.
* For specific interfaces. This requires support for the "unsolicited-params-per-interface" feature.

For operational data, a new "unsolicited" container has been added for BFD IP single-hop sessions.

The tree diagram below uses the graphical representation of data models, as defined in [RFC8340].

```
module: ietf-bfd-unsolicited

  augment /rt:routing/rt:control-plane-protocols
            /rt:control-plane-protocol/bfd:bfd/bfd-ip-sh:ip-sh:
    +--rw unsolicited {bfd-unsol:unsolicited-params-global}?
       +--rw enabled?                          boolean
       +--rw local-multiplier?                 multiplier
       +--rw (interval-config-type)?
          +--:(tx-rx-intervals)
          |  +--rw desired-min-tx-interval?    uint32
          |  +--rw required-min-rx-interval?   uint32
          +--:(single-interval) {single-minimum-interval}?
             +--rw min-interval?              uint32
  augment /rt:routing/rt:control-plane-protocols
            /rt:control-plane-protocol/bfd:bfd/bfd-ip-sh:ip-sh
            /bfd-ip-sh:interfaces:
    +--rw unsolicited {bfd-unsol:unsolicited-params-per-interface}?
       +--rw enabled?                          boolean
       +--rw local-multiplier?                 multiplier
       +--rw (interval-config-type)?
          +--:(tx-rx-intervals)
          |  +--rw desired-min-tx-interval?    uint32
          |  +--rw required-min-rx-interval?   uint32
          +--:(single-interval) {single-minimum-interval}?
             +--rw min-interval?              uint32
  augment /rt:routing/rt:control-plane-protocols
            /rt:control-plane-protocol/bfd:bfd/bfd-ip-sh:ip-sh
            /bfd-ip-sh:sessions/bfd-ip-sh:session:
    +--ro role?   bfd-unsol:role
```

4.2. Unsolicited BFD Module

```
<CODE BEGINS> file "ietf-bfd-unsolicited@2022-10-24.yang"

module ietf-bfd-unsolicited {

  yang-version 1.1;

  namespace "urn:ietf:params:xml:ns:yang:ietf-bfd-unsolicited";

  prefix "bfd-unsol";

  // RFC Ed.: replace occurences of YYYY with actual RFC numbers
  // and remove this note

  import ietf-bfd-types {
    prefix "bfd-types";
    reference
      "RFC 9314: YANG Data Model for Bidirectional Forwarding Detection
      (BFD)";
  }

  import ietf-bfd {
    prefix "bfd";
    reference
      "RFC 9314: YANG Data Model for Bidirectional Forwarding Detection
       (BFD)";
  }

  import ietf-bfd-ip-sh {
    prefix "bfd-ip-sh";
    reference
      "RFC 9314: YANG Data Model for Bidirectional Forwarding Detection
      (BFD)";
  }

  import ietf-routing {
    prefix "rt";
    reference
      "RFC 8349: A YANG Data Model for Routing Management
       (NMDA version)";
  }

  organization "IETF BFD Working Group";

  contact
    "WG Web:   <https://datatracker.ietf.org/wg/bfd/>
     WG List:  <rtg-bfd@ietf.org>

     Editors:  Enke Chen (enchen@paloaltonetworks.com),
               Naiming Shen (naiming@zededa.com),
               Robert Raszuk (robert@raszuk.net),
```

```
              Reshad Rahman (reshad@yahoo.com)";

  description
    "This module contains the YANG definition for BFD unsolicited
     as per RFC YYYY.

     Copyright (c) 2021 IETF Trust and the persons
     identified as authors of the code.  All rights reserved.

     Redistribution and use in source and binary forms, with or
     without modification, is permitted pursuant to, and subject
     to the license terms contained in, the Simplified BSD License
     set forth in Section 4.c of the IETF Trust's Legal Provisions
     Relating to IETF Documents
     (http://trustee.ietf.org/license-info).

     This version of this YANG module is part of RFC YYYY; see
     the RFC itself for full legal notices.";

  reference "RFC YYYY";

  revision 2022-10-24 {
    description
      "Initial revision.";
    reference
      "RFC YYYY: Unsolicited BFD for Sessionless Applications.";
  }

  /*
   * Feature definitions
   */
  feature unsolicited-params-global {
    description
      "This feature indicates that the server supports global
       parameters for unsolicited sessions.";
    reference
      "RFC YYYY: Unsolicited BFD for Sessionless Applications.";
  }

  feature unsolicited-params-per-interface {
    description
      "This feature indicates that the server supports per-interface
       parameters for unsolicited sessions.";
    reference
      "RFC YYYY: Unsolicited BFD for Sessionless Applications.";
  }

  /*
   * Type Definitions
   */
```

```
typedef role {
  type enumeration {
    enum active {
      description "Active role";
    }
    enum passive {
      description "Passive role";
    }
  }
  description "Role";
  reference
    "RFC5880: Bidirectional Forwarding Detection (BFD),
     Section 6.1";
}

/*
 * Augments
 */
 augment "/rt:routing/rt:control-plane-protocols/"
       + "rt:control-plane-protocol/bfd:bfd/bfd-ip-sh:ip-sh" {
   if-feature bfd-unsol:unsolicited-params-global;
   description
     "Augmentation for BFD unsolicited parameters";
   container unsolicited {
     description
       "BFD unsolicited top level container";
     leaf enabled {
       type boolean;
       default false;
       description
         "BFD unsolicited enabled globally for IP single-hop.";
     }
     uses bfd-types:base-cfg-parms;
   }
 }

 augment "/rt:routing/rt:control-plane-protocols/"
       + "rt:control-plane-protocol/bfd:bfd/bfd-ip-sh:ip-sh/"
       + "bfd-ip-sh:interfaces" {
   if-feature bfd-unsol:unsolicited-params-per-interface;
   description
     "Augmentation for BFD unsolicited on IP single-hop interface";
   container unsolicited {
     description
       "BFD IP single-hop interface unsolicited top level
        container";
     leaf enabled {
       type boolean;
       default false;
```

```
        description
          "BFD unsolicited enabled on this interface.";
       }
       uses bfd-types:base-cfg-parms;
     }
   }

  augment "/rt:routing/rt:control-plane-protocols/"
        + "rt:control-plane-protocol/bfd:bfd/bfd-ip-sh:ip-sh/"
        + "bfd-ip-sh:sessions/bfd-ip-sh:session" {
    description
      "Augmentation for BFD unsolicited on IP single-hop session";
      leaf role {
        type bfd-unsol:role;
        config false;
        description "Role.";
      }
  }
}


<CODE ENDS>
```

5. IANA Considerations

This document registers the following namespace URI in the "IETF XML Registry" [RFC3688]:

URI: urn:ietf:params:xml:ns:yang:ietf-bfd-unsolicited

Registrant Contact: The IESG.

XML: N/A; the requested URI is an XML namespace.

This document registers the following YANG module in the "YANG Module Names" registry [RFC6020]:

Name: ietf-bfd-unsolicited

Namespace: urn:ietf:params:xml:ns:yang:ietf-bfd-unsolicited

Prefix: bfd-unsol

Reference: RFC YYYY

6. Acknowledgments

Authors would like to thank Acee Lindem, Greg Mirsky, Jeffrey Haas, Raj Chetan, Tom Petch, Henning Rogge, Mahesh Jethanandani, Gyan Mishra and John Scudder for their review and valuable input.

7. Security Considerations

7.1. BFD Protocol Security Considerations

The same security considerations and protection measures as those described in [RFC5880] and [RFC5881] apply to this document. In addition, with "unsolicited BFD" there is potential risk for excessive resource usage by BFD from "unexpected" remote systems. To mitigate such risks, the following measures are mandatory:

- Limit the feature to specific interfaces, and to single-hop BFD with "TTL=255" [RFC5082].
- Apply "policy" to allow BFD packets only from certain subnets or hosts.
- Deploy the feature only in certain "trustworthy" environment, e.g., at an IXP, or between a provider and its customers.
- Use BFD authentication, see [RFC5880]. In some environments, e.g. when using an IXP, BFD authentication can not be used because of the lack of coordination into the operation of the two endpoints of the BFD session. In other environments, e.g. when BFD is used to track the next hop of static routes, it is possible to use BFD authentication: this comes with the extra cost of configuring matching key-chains at the two endpoints. If BFD authentication is used, the Meticulous Keyed SHA1 mechanism SHOULD be used.

7.2. YANG Module Security Considerations

The YANG module specified in this document defines a schema for data that is designed to be accessed via network management protocols such as NETCONF [RFC6241] or RESTCONF [RFC8040]. The lowest NETCONF layer is the secure transport layer, and the mandatory-to-implement secure transport is Secure Shell (SSH) [RFC6242]. The lowest RESTCONF layer is HTTPS, and the mandatory-to-implement secure transport is TLS [RFC8446].

The NETCONF access control model [RFC8341] provides the means to restrict access for particular NETCONF or RESTCONF users to a preconfigured subset of all available NETCONF or RESTCONF protocol operations and content.

There are a number of data nodes defined in this YANG module that are writable/creatable/deletable (i.e., config true, which is the default). These data nodes may be considered sensitive or vulnerable in some network environments. Write operations (e.g., edit-config)

to these data nodes without proper protection can have a negative effect on network operations. These are the subtrees and data nodes and their sensitivity/vulnerability:

/routing/control-plane-protocols/control-plane-protocol/bfd/ip-sh / unsolicited:

- *data node "enabled" enables creation of unsolicited BFD IP single-hop sessions globally, i.e. on all interfaces. See Section 7.1.
- *data nodes local-multiplier, desired-min-tx-interval, required-min-rx-interval and min-interval all impact the parameters of the unsolicited BFD IP single-hop sessions.

/routing/control-plane-protocols/control-plane-protocol/bfd/ip-sh / interfaces/interface/unsolicited:

- *data node "enabled" enables creation of unsolicited BFD IP single-hop sessions on a specific interface. See Section 7.1.
- *data nodes local-multiplier, desired-min-tx-interval, required-min-rx-interval and min-interval all impact the parameters of the unsolicited BFD IP single-hop sessions on the interface.

Some of the readable data nodes in this YANG module may be considered sensitive or vulnerable in some network environments. It is thus important to control read access (e.g., via get, get-config, or notification) to these data nodes. These are the subtrees and data nodes and their sensitivity/vulnerability:

/routing/control-plane-protocols/control-plane-protocol/bfd/ip-sh / sessions/session/role: access to this information discloses the role of the local system in the creation of the unsolicited BFD session.

8. References

8.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC3688] Mealling, M., "The IETF XML Registry", BCP 81, RFC 3688, DOI 10.17487/RFC3688, January 2004, <https://www.rfc-editor.org/info/rfc3688>.

[RFC5082] Gill, V., Heasley, J., Meyer, D., Savola, P., Ed., and C. Pignataro, "The Generalized TTL Security Mechanism (GTSM)", RFC 5082, DOI 10.17487/RFC5082, October 2007, <https://www.rfc-editor.org/info/rfc5082>.

[RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, DOI 10.17487/RFC5880, June 2010, <https://www.rfc-editor.org/info/rfc5880>.

[RFC5881] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD) for IPv4 and IPv6 (Single Hop)", RFC 5881, DOI 10.17487/RFC5881, June 2010, <https://www.rfc-editor.org/info/rfc5881>.

[RFC6020] Bjorklund, M., Ed., "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)", RFC 6020, DOI 10.17487/RFC6020, October 2010, <https://www.rfc-editor.org/info/rfc6020>.

[RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", RFC 6241, DOI 10.17487/RFC6241, June 2011, <https://www.rfc-editor.org/info/rfc6241>.

[RFC6242] Wasserman, M., "Using the NETCONF Protocol over Secure Shell (SSH)", RFC 6242, DOI 10.17487/RFC6242, June 2011, <https://www.rfc-editor.org/info/rfc6242>.

[RFC8040] Bierman, A., Bjorklund, M., and K. Watsen, "RESTCONF Protocol", RFC 8040, DOI 10.17487/RFC8040, January 2017, <https://www.rfc-editor.org/info/rfc8040>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8340] Bjorklund, M. and L. Berger, Ed., "YANG Tree Diagrams", BCP 215, RFC 8340, DOI 10.17487/RFC8340, March 2018, <https://www.rfc-editor.org/info/rfc8340>.

[RFC8341] Bierman, A. and M. Bjorklund, "Network Configuration Access Control Model", STD 91, RFC 8341, DOI 10.17487/RFC8341, March 2018, <https://www.rfc-editor.org/info/rfc8341>.

[RFC8349] Lhotka, L., Lindem, A., and Y. Qu, "A YANG Data Model for Routing Management (NMDA Version)", RFC 8349, DOI 10.17487/RFC8349, March 2018, <https://www.rfc-editor.org/info/rfc8349>.

[RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <https://www.rfc-editor.org/info/rfc8446>.

[RFC9314] Jethanandani, M., Ed., Rahman, R., Ed., Zheng, L., Ed., Pallagatti, S., and G. Mirsky, "YANG Data Model for Bidirectional Forwarding Detection (BFD)", RFC 9314, DOI 10.17487/RFC9314, September 2022, <https://www.rfc-editor.org/info/rfc9314>.

8.2. Informative References

[I-D.ietf-idr-rs-bfd] Bush, R., Haas, J., John Scudder, G., Nipper, A., and C. Dietzel, "Making Route Servers Aware of Data Link Failures at IXPs", Work in Progress, Internet-Draft, draft-ietf-idr-rs-bfd-09, 21 September 2020, <https://www.ietf.org/archive/id/draft-ietf-idr-rs-bfd-09.txt>.

[RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <https://www.rfc-editor.org/info/rfc4271>.

[RFC5883] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD) for Multihop Paths", RFC 5883, DOI 10.17487/RFC5883, June 2010, <https://www.rfc-editor.org/info/rfc5883>.

[RFC7880] Pignataro, C., Ward, D., Akiya, N., Bhatia, M., and S. Pallagatti, "Seamless Bidirectional Forwarding Detection (S-BFD)", RFC 7880, DOI 10.17487/RFC7880, July 2016, <https://www.rfc-editor.org/info/rfc7880>.

[RFC7911] Walton, D., Retana, A., Chen, E., and J. Scudder, "Advertisement of Multiple Paths in BGP", RFC 7911, DOI 10.17487/RFC7911, July 2016, <https://www.rfc-editor.org/info/rfc7911>.

[RFC7947] Jasinska, E., Hilliard, N., Raszuk, R., and N. Bakker, "Internet Exchange BGP Route Server", RFC 7947, DOI 10.17487/RFC7947, September 2016, <https://www.rfc-editor.org/info/rfc7947>.

[RFC8342] Bjorklund, M., Schoenwaelder, J., Shafer, P., Watsen, K., and R. Wilton, "Network Management Datastore Architecture (NMDA)", RFC 8342, DOI 10.17487/RFC8342, March 2018, <https://www.rfc-editor.org/info/rfc8342>.

Authors' Addresses

Enke Chen
Palo Alto Networks

Email: enchen@paloaltonetworks.com

Naiming Shen
Zededa

Email: naiming@zededa.com

Robert Raszuk
Arrcus
2077 Gateway Place
San Jose, CA 95110
United States of America

Email: robert@raszuk.net

Reshad Rahman
Graphiant
Canada

Email: reshad@yahoo.com