

I2RS working group
Internet-Draft
Intended status: Standards Track
Expires: September 22, 2016

S. Hares
Huawei
S. Kini
Ericsson
L. Dunbar
Huawei
R. Krishnan
Dell
D. Bogdanovic
Juniper Networks
R. White
Linkedin
March 21, 2016

Filter-Based RIB Data Model
draft-hares-i2rs-fb-rib-data-model-03

Abstract

This document defines a data model to support the Filter-based Routing Information Base (RIB) Yang data models for I2RS. A routing system uses the Filter-based RIB to program FIB entries that process incoming packets by matching on multiple fields within the packet and then performing a specified action on it. The FB-RIB can also specify an action to forward the packet according to the FIB entries programmed using the RIBs of its routing instance.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on September 22, 2016.

Copyright Notice

Copyright (c) 2016 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	2
1.1. Definition of I2RS Filter Based RIB	3
2. Requirements Language	4
3. Definitions and Acronyms	4
4. High level Yang structure for the FB-RIB	5
4.1. Top Level Yang Structure for ietf-fb-rib	7
4.2. Filter-Based RIB structures	8
5. yang models	9
5.1. Filter-Based RIB types	9
5.2. FB-RIB	15
6. IANA Considerations	18
7. Security Considerations	18
8. References	18
8.1. Normative References:	18
8.2. Informative References	19
Authors' Addresses	20

1. Introduction

The Interface to the Routing System (I2RS) [I-D.ietf-i2rs-architecture] architecture provides dynamic read and write access to the information and state within the routing elements. The I2RS client interacts with the I2RS agent in one or more network routing systems.

This document provides a yang module for the I2RS filter Based Routing Information Base (FB-RIB) and describes the I2RS interaction with routing filters within a routing element. The informational model for the FB-RIB is in [I-D.kini-i2rs-fb-rib-info-model]

1.1. Definition of I2RS Filter Based RIB

Filter-based routing is a technique used to make packet forwarding decisions based on a filter that is matched to the incoming packets and the specified action. It should be noted that that this is distinct from the static routes in the RIB [I-D.ietf-i2rs-rib-info-model] where the routing is destination address based.

A Filter-Based RIB (Routing Information Base) is contained in a routing instance (defined in [I-D.ietf-i2rs-rib-info-model]). It contains a list of filters (match-action conditions) and a list of interfaces the filter-based forwarding operates on, and default RIB(s).

A Filter Based RIB uses packet forwarding policy. If packet reception is considered an event, then the I2RS Filter-based RIB uses a minimalistic Event-matchCondition-Action policy with the following characteristics:

event = packet/frame received,

match condition - match on field in frame/packet or circumstances relating to packet reception (e.g. time received),

action - modify packet and forward/drop packet.

A Filter-based RIB entry specifies match filters for the fields in a packet (which may include layer 1 to layer 3 header fields, transport or application fields) or size of the packet or interface received on. The matches are contained in an ordered list of filters which contain pairs of match condition-action (aka event-condition-action).

If all matches fail, default action is to forward the packet using Destination Based forward from the default RIB(s). The default RIBs can be:

- o created by the I2RS Routing Information Base (RIB) manager using the yang model described in: in [I-D.ietf-i2rs-rib-info-model], or
- o configured RIB created using static routes or [I-D.ietf-netmod-routing-cfg].
- o or static RIB created via static route yang model

Actions in the condition-action pair may impact forwarding or set something in the packet that will impact forwarding. Policy actions

are typically applied before applying QoS constraints since policy actions may override QoS constraint.

The Filter-Based RIB resides in ephemeral state as does the I2RS RIB and I2RS topology models.

2. Requirements Language

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 [RFC2119].

In this document, these words will appear with that interpretation only when in ALL CAPS. Lower case uses of these words are not to be interpreted as carrying RFC-2119 significance.

3. Definitions and Acronyms

CLI

Command Line Interface

FB-RIB

Filter-Based Routing Information Base

FB-Route

The policy rules in the filter-based RIB are prescriptive of the Event-Condition-Action form which is often represented by "if Condition then action".

Policy Group

Policy Groups are groups of policy rules. The groups of policy in the basic network policy [I-D.hares-i2rs-pkt-eca-data-model] allow grouping of policy by name. This structure allow easier management of customer-based or provider based filters, but does not change the policy-rules list.

RIB IM

RIB Informational Model (RIB IM) [I-D.ietf-i2rs-rib-info-model]

Routing instance

A routing instance, in the context of the FB-FIB is a collection of RIBs, interfaces, and routing parameters. A routing instance

creates a logical slice of the router and allows different logical slices; across a set of routers; to communicate with each other.

4. High level Yang structure for the FB-RIB

There are three levels in the Filter-Based RIB (FB-RIB) structure:

- o a global FB-RIB structures,
- o the common structure of the FB-RIB, and
- o the groupings that make up the FB-RIB

All structures have two types: configuration/ephemeral state and operational state.

This yang model describes three types of FB-RIBS: configuration, I2RS, and BGP Flow Specification. The configuration FB-RIB yang module is config state ("config true" and "ephemeral false") and survives a reboot. The I2RS FB-RB yang model is reboot ephemeral ("config true" and "ephemeral true;"). The BGP Flow Specification Filter-Based RIB stores policy which is received by the BGP peers. The BGP configuration contains a flow-specification as part of its configuration, and the ability to export this flow specification to other BGP peers. The BGP local configuration with the flow specification is consider ("config true"), and it is possible to be ephemeral ("ephemeral true") or local configuration ("ephemeral false"). The BGP flow specifications received from peers are derived state, just like other BGP derved state.

Configuration RIBS




Figure 3: Routing instance with three types of Filter-FIB lists

The following section provides the high level yang structure diagrams for the following levels of structures for both config/ephemeral state and operationa.

- o ietf-fb-rib - contains filter-based RIBS for config, I2RS FB-RIB, and BGP Flow Specification.
- o fb-rib - that contains the structures for the filter-based grouping
- o fb-rib-types - that contains the structures for groupings within the filter-based RIBS

These structures are contained within the yang section in this draft.

The packet-reception ECA policy yang module is contained in the draft [I-D.hares-i2rs-pkt-eca-data-model].

For those who desire more information regarding the logic behind the I2RS Filter-Based RIB, please see the Informational Model at: [I-D.kini-i2rs-fb-rib-info-model].

4.1. Top Level Yang Structure for ietf-fb-rib

The Top-level Yang structure for a global FB-RIB types (similar to acl) is not defined for filter-based RIBS. The I2RS Filter-Based RIB should be defined under this structure under a routing instance. The three things under this RIB would be: configured Filter-Based RIB (aka Policy routing), I2RS reboot Ephemeral Filter-Based RIB, and BGP Flow Specification's Filter-Based RIB. All of these RIBs have similar actions.

There are two types top-level structures for ietf-fb-ribs: config and operational state.

The Top-level Yang structure for a global configuration of Filter-Based RIBs are:

```
Augments rt:logical-network-elements:\
    :logical-network-element:network-instances: \
        network-instance
```

```
ietf-fb-rib module
  +--rw ietf-fb-rib
    +--rw default-instance-name string
    +--rw default-router-id rt:router-id
    +--rw config-fb-ribs
      if-feature "config-filter-based-RIB";
      uses fb-ribs;
    +--rw i2rs-fb-ribs
      if-feature "I2RS-filter-based-RIB";
      uses fb-rib-t:fb-ribs;
    +--rw bgp-fs-fb-ribs
      if-feature "BGP-FS-filter-based-RIB";
      uses fb-rib-t:fb-ribs;
```

Figure 5: configuration state

The Top-level Yang structure for a global operational state of Filter-Based RIBs are:

```
Augments rt:logical-network-elements:\
  :logical-network-element:network-instances: \
    network-instance

ietf-fb-rib module
  +--rw ietf-fb-rib-opstate
  +--rw default-instance-name string
  +--rw default-router-id rt:router-id
  +--rw config-fb-rib-opstate
    if-feature "config-filter-based-RIB";
    uses fb-rib-t:fb-ribs-oper-status;
  +--rw i2rs-fb-rib-opstate {
    if-feature "I2RS-filter-based-RIB";
    uses fb-rib-t:fb-ribs-oper-status;
  +--rw bgp-fs-fb-rib-opstate
    if-feature "BGP-FS-filter-based-RIB";
    uses fb-rib-t:fb-ribs-oper-status;
```

Figure 5: operational state

4.2. Filter-Based RIB structures

The Top-level yang structures at the Filter-Based RIB level have two types: configuration and operational state.

The Top-level Yang structure for the FB-RIB types is:


```

module: fb-rib-types:
+--rw fb-ribs
  +--rw fb-rib* [rib-name]
    |   +--rw rib-name string
    |   |   rw fb-type identityref / ephemeral or not
    |   +--rw rib-afi rt:address-family
    |   +--rw fb-rib-intf* [name]
    |   |   +--rw name string
    |   |   +--rw intf if:interface
    |   +--rw default-rib
    |   |   +--rw rt-rib rt:routing:routing-instance:name
    |   |   +--rw config-rib string; // config rib name
    |   |   +--rw i2rs-rib:routing-instance:name
    |   |   +--rw i2rs-rib string; //ephemeral rib name
    |   |   +--rw bgp-instance-name string
    |   |   +--rw bgp-rib string //session ephemeral
    |   +--rw fb-rib-refs
    |   |   +--rw fb-rib-update-ref uint32 /count of writes
    |   +--rw instance-using*
    |   |   device:networking-instance:networking-instance-name
    |   +--use pkt-eca:pkt-eca-policy-set

```

Figure 6: FB RIB Type Structure

High Level Yang

```

+--rw fb-ribs-oper-status
  +--rw fb-rib-oper-status* [fb-rib-name]
    uses pkt-eca:pkt-eca-opstate

```

5. yang models

5.1. Filter-Based RIB types

```

<CODE BEGINS> file "ietf-fb-rib-types@2016-02-09.yang"
module ietf-fb-rib-types {

  yang-version "1";

  // namespace
  namespace "urn:ietf:params:xml:ns:yang:ietf-fb-rib-types";
  prefix "fb-rib-t";
  import ietf-interfaces {prefix "if";}
  import ietf-routing {prefix "rt";}
  import ietf-pkt-eca-policy {prefix "pkt-eca";}

  // meta
  organization

```

```
"IETF";
```

```
contact
```

```
  "email: shares@ndzh.com;  
    email: sriganesh.kini@ericsson.com  
    email: cengiz@packetdesign.com  
    email: ivandean@gmal.org  
    email: linda.dunbar@huawei.com;  
    email: russ@riw.com;  
  ";
```

```
description
```

```
"This module describes a YANG model for the I2RS  
Filter-based RIB Types.  These types  
specify types for the Filter-Based RIB.
```

```
  Copyright (c) 2015 IETF Trust and the persons identified as  
  the document authors.  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).";
```

```
revision "2016-02-09" {
```

```
  description
```

```
    "Filter-Based RIB protocol ";  
    reference "draft-hares-i2rs-fb-rib-data-model-01";
```

```
}
```

```
typedef fb-rib-type-def {
```

```
  type identityref {  
    base "fb-rib-type";
```

```
  }
```

```
  description
```

```
    "This type is used to refer to  
    source of Filter-Based RIB:  
    configuration, I2RS, Flow-Spec.";
```

```
}
```

```
  identity fb-rib-type {
```

```
    description
```

```
      "This type is used to refer to  
      source of Filter-Based RIB:  
      configuration, I2RS, Flow-Spec.";
```

```
    }

    identity fb-rib-config-type {
        base fb-rib-type;
        description
            "config Filter-Based RIB";
    }

    identity fb-rib-i2rs-ephemeral-type {
        base fb-rib-type;
        description
            "I2RS Reboot ephemeral Filter-Based RIB";
    }

    identity fb-rib-BGP-FS-type {
        base fb-rib-type;
        description
            "BGP Flow Specification Filter-Based RIB";
    }

typedef fb-rib-policy-type-def {
    type identityref {
        base "fb-rib-policy-type";
    }
    description
        "This type is used to refer to FB-RIB type";
}

identity fb-rib-policy-type {
    description
        "Types of filter-based policies
        acl and eca";
}

identity fb-rib-acl {
    base fb-rib-policy-type;
    description
        "filter based policy based on access-lists";
}

identity fb-bnp-eca-rules {
    base fb-rib-policy-type;
    description
        "filter based policy based on qos forwarding rules";
}

typedef fb-rules-status {
    type identityref {
```

```
        base "fb-rule-opstat";
        }
        description
        "This type is used to refer to FB-RIB type";
    }

    identity fb-rule-opstat {
        description
        "operational statuses for filter rules
        inactive and active";
    }

    identity fb-rule-inactive {
        base fb-rule-opstat;
        description
        "policy rule is inactive";
    }

    identity fb-rule-active {
        base fb-rule-opstat;
        description
        "policy rule is active";
    }

    grouping fb-rib-rule-order-status {
    leaf statement-order {
        type uint16;
        description "order identifier";
    }
    leaf statement-oper_status {
        type fb-rules-status;
        description "status of rule";
    }
    description "filter-rib
    policy rule order and status";
    }

    grouping fb-rib-group-order-status {
    leaf group-refcnt {
        type uint16;
        description "refcnt for this group";
    }
    leaf group-installed {
        type uint32;
        description "number of rules installed";
    }
    leaf group-matches {
        type uint64;
    }
    }
```

```
        description "number of matches by all
          rules in group";
      }
      description "fb-rib group list order
        and status info.";
    }

grouping fb-rib-updates {
  leaf fb-rib-update-ref {
    type uint64;
    description
      "number of updates to this FB RIB
        since last reboot";
  }
  description "FB-RIB update info";
}

grouping default-fb-rib {
  // configuration instance for default RIB
  leaf config-instance {
    type rt:routing-instance-ref;
    description "instance name";
  }
  leaf config-rib {
    type string;
    description "name of config default RIB";
  }
  //I2RS default instance for default RIB
  leaf i2rs-instance-name {
    type string;
    description "I2RS instance name";
  }
  leaf i2rs-rib-name {
    type string;
    description "name of default I2RS RIB";
  }
  leaf bgp-instance-name {
    type string;
    description "name of bgp instance";
  }

  leaf bgp-fs-rib-name {
    type string;
    description "name of BGP
      flow specification default RIB";
  }
  description "default RIB for forwarding
    if the policy match";
}
```

```

}

grouping fb-ribs {
  list fb-rib {
    key fb-rib-name;
    leaf fb-rib-name {
      type string;
      mandatory true;
      description "RIB name";
    }
    uses rt:address-family;
    leaf fb-type {
      type fb-rib-type-def;
      description "type of RIB
        list: config, I2RS reboot
        ephemeral, BGP Flow Specification
        ephemeral. ";
    }
    list fb-rib-intf {
      key "name";
      leaf name {
        type if:interface-ref;
        description
          "A reference to the name of a
            configured network layer
            interface.";
      }
      description "This represents
        the list of interfaces
        associated with this routing instance.
        The interface list helps constrain the
        boundaries of packet forwarding.
        Packets coming on these interfaces are
        directly associated with the given routing
        instance. The interface list contains a
        list of identifiers, with each identifier
        uniquely identifying an interface.";
    }
    uses default-fb-rib; // defaults ribs
    uses fb-rib-updates; // write refs to this RIB
    list instance-using {
      key instance-name;
      leaf instance-name {
        type string;
        description
          " name of instance using this fb-rib
            rt:routing-instance";
      }
    }
  }
}

```

```

        description "instances using
            this fb-rib";
    }
    // ordered rule list + group list
    uses pkt-eca:pkt-eca-policy-set;

    description "Configuration of
        an filter-based rib list";
    }
    description "fb-rib group";
}

grouping fb-ribs-oper-status {
    list fb-rib-oper-status {
        key fb-rib-name;
        leaf fb-rib-name {
            type string;
            description "rib name";
        }
    }
    uses pkt-eca:pkt-eca-opstate;
    description "Configuration of
        an filter-based rib list";
    }
    description "list of FB-FIB operational
        status";
}

}

```

<CODE ENDS>

5.2. FB-RIB

```

<CODE BEGINS> file "ietf-fb-rib@2016-02-09.yang"
module ietf-fb-rib {
    yang-version "1";

    // namespace
    namespace "urn:ietf:params:xml:ns:yang:ietf-fb-rib";
    // replace with iana namespace when assigned
    prefix "fb-rib";

    // import some basic inet types
    import ietf-yang-types {prefix "yang";}
    import ietf-fb-rib-types { prefix "fb-rib-t";}
}

```

```
// meta
organization
  "IETF";

contact
  "email: sriganesh.kini@ericsson.com
    email: cengiz@packetdesign.com
    email: anoop@ieee.duke.edu
    email: ivandean@gmail.org
    email: shares@ndzh.com;
    email: linda.dunbar@huawei.com;
    email: russ@riw.com;
  ";

description
  "This Top level module describes a YANG model for the I2RS
    Filter-based RIB which is an global protocol independent FB RIB modu
    le.";

  revision "2016-02-09" {
    description "initial revision";
    reference "draft-hares-i2rs-fb-rib-data-model-01";
  }

  feature config-filter-based-RIB {
description
  "This feature means that a node support
    config filter-based rib.";
  }

  feature I2RS-filter-based-RIB {
description
  "This feature means that a node support
    I2RS filter-based rib.";
  }

  feature BGP-FS-filter-based-RIB {
description
  "This feature means that a node support
    BGP FS filter-based rib.";
  }

  container ietf-fb-rib {
    presence "top-level structure for
      configuration";
    leaf default-instance-name {
      type string;
      mandatory true;
    }
    description
      "A routing instance is identified by its name,
```



```
    INSTANCE_name.  This MUST be unique across all routing
    instances in a given network device.";
  }
  leaf default-router-id {
    type yang:dotted-quad;
    description "Default router id";
  }
  container config-fb-rib {
    if-feature config-filter-based-RIB;
    uses fb-rib-t:fb-ribs;
    description "config filter-based RIB";
  }

  container i2rs-fb-rib {
    if-feature I2RS-filter-based-RIB;
    uses fb-rib-t:fb-ribs;
    description "bgp-fs filter-based RIB";
  }
  container bgp-fs-fb-rib {
    if-feature BGP-FS-filter-based-RIB;
    uses fb-rib-t:fb-ribs;
    description "bgp fs filter-based RIB";
  }
  description "fb-rib augments routing instance";
}

container ietf-fb-rib-opstate {
  presence "top-level structure for
  op-state";
  config "false";
leaf default-instance-name {
  type string;
  mandatory true;
description
  "A routing instance is identified by its name,
  INSTANCE_name.  This MUST be unique across all routing
  instances in a given network device.";
}
  leaf default-router-id {
    type yang:dotted-quad;
    description "Default router id";
  }
  container config-fb-rib-opstate {
    if-feature config-filter-based-RIB;
    uses fb-rib-t:fb-ribs-oper-status;
    description "config filter-based RIB";
  }
  container i2rs-fb-rib-opstate {
```

```
        if-feature I2RS-filter-based-RIB;
        uses fb-rib-t:fb-ribs-oper-status;
        description "bgp-fs filter-based RIB";
    }
    container bgp-fs-fb-rib-opstate {
        if-feature BGP-FS-filter-based-RIB;
        uses fb-rib-t:fb-ribs-oper-status;
        description "bgp fs filter-based RIB";
    }
    description "fb-rib augments routing instance";
}
}
```

<CODE ENDS>

6. IANA Considerations

TBD

7. Security Considerations

A I2RS RIB is ephemeral data store that will dynamically change traffic paths set by the routing configuration. An I2RS FB-RIB provides dynamic Event-Condition-Action policy that will further change the operation of forwarding by allow dynamic policy and ephemeral RIBs to alter the traffic paths set by routing configuration. Care must be taken in deployments to use the appropriate security and operational control to make use of the tools the I2RS RIB and I2RS FB-RIB provide.

8. References

8.1. Normative References:

[I-D.hares-i2rs-pkt-eca-data-model]

Hares, S., Wu, Q., and R. White, "Filter-Based Packet Forwarding ECA Policy", draft-hares-i2rs-pkt-eca-data-model-02 (work in progress), February 2016.

[I-D.ietf-i2rs-architecture]

Atlas, A., Halpern, J., Hares, S., Ward, D., and T. Nadeau, "An Architecture for the Interface to the Routing System", draft-ietf-i2rs-architecture-13 (work in progress), February 2016.

[I-D.ietf-i2rs-rib-data-model]

Wang, L., Ananthakrishnan, H., Chen, M., amit.dass@ericsson.com, a., Kini, S., and N. Bahadur, "A YANG Data Model for Routing Information Base (RIB)", draft-ietf-i2rs-rib-data-model-05 (work in progress), March 2016.

[I-D.ietf-i2rs-rib-info-model]

Bahadur, N., Kini, S., and J. Medved, "Routing Information Base Info Model", draft-ietf-i2rs-rib-info-model-08 (work in progress), October 2015.

[I-D.ietf-netmod-acl-model]

Bogdanovic, D., Koushik, K., Huang, L., and D. Blair, "Network Access Control List (ACL) YANG Data Model", draft-ietf-netmod-acl-model-07 (work in progress), March 2016.

[I-D.ietf-netmod-routing-cfg]

Lhotka, L. and A. Lindem, "A YANG Data Model for Routing Management", draft-ietf-netmod-routing-cfg-21 (work in progress), March 2016.

[I-D.kini-i2rs-fb-rib-info-model]

Kini, S., Hares, S., Dunbar, L., Ghanwani, A., Krishnan, R., Bogdanovic, D., and R. White, "Filter-Based RIB Information Model", draft-kini-i2rs-fb-rib-info-model-03 (work in progress), February 2016.

8.2. Informative References

[I-D.acee-rtgwg-yang-rib-extend]

Lindem, A. and Y. Qu, "YANG Data Model for RIB Extensions", draft-acee-rtgwg-yang-rib-extend-00 (work in progress), October 2015.

[I-D.ietf-i2rs-usecase-reqs-summary]

Hares, S. and M. Chen, "Summary of I2RS Use Case Requirements", draft-ietf-i2rs-usecase-reqs-summary-02 (work in progress), March 2016.

[RFC2119]

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

Authors' Addresses

Susan Hares
Huawei
7453 Hickory Hill
Saline, MI 48176
USA

Email: shares@ndzh.com

Sriganesh Kini
Ericsson

Email: sriganesh.kini@ericsson.com

Linda Dunbar
Huawei
USA

Email: linda.dunbar@huawei.com

Ram Krishnan
Dell

Email: Ramkri123@gmail.com

Dean Bogdanovic
Juniper Networks
Westford, MA

Email: ivandean@gmail.org

Russ White
Linkedin

Email: russ@riw.us