

I2RS working group
Internet-Draft
Intended status: Standards Track
Expires: January 7, 2016

S. Kini
Ericsson
S. Hares
L. Dunbar
Huawei
A. Ghanwani
R. Krishnan
Dell
D. Bogdanovic
Juniper Networks
J. Tantsura
R. White
Ericsson
July 6, 2015

Filter-Based RIB Information Model
draft-kini-i2rs-fb-rib-info-model-01

Abstract

This document defines an information model and a data model for the I2RS Filter-based Routing Information Base (RIB) Yang model. 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 January 7, 2016.

Copyright Notice

Copyright (c) 2015 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
2. Definitions and Acronyms	4
3. Filter-Based-RIB module	5
3.1. FB-RIB entries	7
3.2. Relationship between RB-RIB Rule Model and RIB Information Model	9
4. yang models	9
4.1. Filter-Based RIB types	9
4.2. fb-fib yang module	12
5. IANA Considerations	14
6. Security Considerations	14
7. References	15
7.1. Normative References:	15
7.2. Informative References	15
Authors' Addresses	16

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.

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 entry specifies matches on fields in a packet (which may include layer 2 fields, IP 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).

(Note: Filter-based RIBs (FB-RIBs) operate only on the interface the FB-RIB are configured on.)

If all matches fail, default action is to forward the packet using FIB entries that were programmed by the Routing Information Base (RIB) manager described in [I-D.ietf-i2rs-rib-info-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.

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), a list of interface the filter-based forwarding operates on.

The filter based-RIB will event-condition-action policy (ECA) rules based a set of policies specified by the I2RS FB-RIB. The filter-based RIB may utilize policies defined by common IETF policy modules or customer specific policies. The following policies are used in this version of the yang module:

- o Access lists (ACLs) [I-D.ietf-netmod-acl-model] (Note: The ACL Filters do not provide enough depth to support all use cases for filter-based FIBS. These filters are included for cases where this limited filters will work. In early deployments of the FB-FIB, these filters may be the only filters provided.),
- o Basic network filters [I-D.hares-i2rs-bnp-info-model]

The Filter-Based routing may provide many benefits, including better resource allocation, load balancing and QoS.

The I2RS use cases which benefit from Filter-Based Routing are:

- o Protocol independent Use cases and large flow use cases described in [I-D.hares-i2rs-usecase-reqs-summary].
- o the use cases of steering traffic to their designated service functions that are different than the packet's destinations, and
- o large flow use cases described in [I-D.hares-i2rs-usecase-reqs-summary]




Figure 1: Example of matching conditions for basic network filters

The action has to be egress port specific.

Figure 1: The network model structure

2. 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-bnp-info-model] allow grouping of policy by name. This name allow easier management of customer-based or provider based filters.

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.

3. Filter-Based-RIB module

A Filter-Based RIB (FB-RIB) contains an ordered set of filter routes where each filter-route is a match condition followed by an action. An FB-RIB is contained in a routing-instance that is defined in [I-D.ietf-i2rs-rib-info-model] and whose data modelling is done in [I-D.ietf-i2rs-rib-data-model]. An FB-RIB has a list of interfaces that is a subset of the list of interfaces in the routing-instance that it is contained in. An incoming packet on an interface belonging to a FB-RIB is first handled by the FIB programmed using that FB-RIB. If no match action succeeds, then the packet is forwarded using the FIB programmed using the RIB of that routing instance.

An ordered set of filters implies that the insertion of a filter route into a FB-RIB MUST provide the ability to insert a filter route at any specific position and delete of a filter-based route at a specific position. The ability to change a filter route at a specific position combines these two functions (delete an existing filter route rule and add a new policy rule).

Each FB-RIB is contained within a routing instance, but one routing instance (named by an INSTANCE_NAME) can contain multiple FB-RIBs. Each routing instance is associated with a set of interfaces, a router-id, and list of FB-RIBs. Each interface can be associated with at most one FB RIB.

The processing within the FB-RIB process within the routing system is expected to do the following:

- o When a packet successfully matches match term/entry in a filter-route, the corresponding rule-actions are applied.
- o If a packet does not match the match term/entry in the filter route, the filter route processing goes to the next term/entry in the order, and looks for a match, within the current filter or goes to the next filter in the list. This continues until either a filter route match term/entry is successfully matched, or no more filters in the list exists.

- o If no match has been found within list of filters in FB-RIB list, then the packet will be forwarded using the I2RS RIB specified by the FB-RIB if one exists. If no I2RS RIB is specified, the packet will be discarded.




Figure 2: Routing instance with FB-RIB

Policy definitions

ACL types:

```
Policy level access-lists
group level: access_lists: access-list-entries
rule level:  access_lists: access-list-entries:
              access-list-entry
```

BNP QOS

```
Policy level: bnp-eca: bnp-policy-set
group level:  bnp-eca: bnp-policy-set:rule-group-list:rule-group
rule level:   bnp-eca: bnp-policy-set:rule-group-list:rule-group
              policy-rule-list: policy-rule
```

Note: The ACL policy definitions do not provide sufficient depth for the I2RS Filter RIB, but are provided here for early implementations.

Figure 3

3.1. FB-RIB entries

The FB-RIB entries associated with each FB-RIB in a routing instance are:

instance-name (FB-FIB-instance-name)

Name of Routing instance

router-id (FB-RIB-router-id)

router id associated with the FB-RIB function of the Routing instance

Interface_list(FB-RIB-interface)

A list of interfaces that all of the FB-RIB RIBs operate over. This list must be a subset of the interface_list associated with the routing instance.

Default RIB

A RIB contained in the same routing instance that can be used to forward packets when the FIB entries in the FB-RIB list do not match the packets. This Default-RIB forwards based on destination based routing.

FB-RIB Order list of policy (FB-FIB-O-Filters

ordered list of filter rules of the form in
[I-D.hares-i2rs-bnp-info-model]

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

```

module: FB-RIB
+--FB-RIB-module
  +--rw FB-RIB-instance-name
  +--rw RB-RIB-router-id
    uses rt:router-id
  +--rw FB-RIB* [rib-name]
    |   +--rw rib-Name
    |   +--rw rib-afi
    | +--rw fb-rib-intf* if:inteface-ref
    |   +--rw I2RS-RIB
    |     |   +--RIB-name
    |     |     uses i2rs-rib:name
    |     +--rw fb-rib-status-info
    |     |   +--rw fb-rib-update-ref uint64
    |     +--rw fb-rib-Group*
    +--rw filter-type // for group
    +--rw order-number // for group
      + choice (filter-type)
        +-case: acl
          uses: acl: access_lists: access-list-entries
              // operational status augment to group
          augments: access_lists: access-list-entries
          uses fb-rib-group-order_status;

        // operational status augment to individual ACL
        augments: access_lists:access-list-entries:
                  access-list-entry
                  uses fb-rib-rule-order-status;

    +-case: bnp-eca Rules
      uses bnp-eca: bnp-policy-set
      augments bnp-eca:bnp-policy-set:group-list:group
        uses fb-rib-group-order_status
      augment bnp-eca:bnp-policy-set:group-list:group:rule
        uses fb-rib-rule-order_status

```

Figure 4: FB RIB Yang Structure

3.2. Relationship between RB-RIB Rule Model and RIB Information Model

The I2RS RIB module is described in [I-D.ietf-i2rs-rib-info-model] and [I-D.ietf-i2rs-rib-data-model]. The I2RS RIB contains a collection of RIBs with the following information per instance:

- o The set of interfaces indicates which interfaces are associated with this routing instance.
- o The RIBs specify how incoming traffic is to be forwarded based on destination (E.g. RIB and FB-RIB).
- o The routing parameters control the information in the RIBs.

A routing instance may have both an I2RS RIB modules and I2RS FB-FIB modules associated with it.

FB-RIB and RIB can not be used at the same time, which means:

- o If a router doesn't support filter-based routing, a router MUST use RIB and MUST not use FB-RIB.
- o If a router supports filter-based routing:
 - * FB-RIB is used
 - * Multiple FB-RIBs may exist within a routing instance
 - * An interface can be associated with at most one FB-RIB
 - * The Default RIB for a FB-RIB is used if several criteria beyond destination address is not matched.

4. yang models

4.1. Filter-Based RIB types

```
module fb-rib-types {  
  
  yang-version "1";  
  
  // namespace  
  namespace "urn:TBD1:params:xml:ns:yang:rt:i2rs:fbrib-types";  
  // replace with iana namespace when assigned  
  prefix "i2rs-fbrib-types";  
  
  // meta  
  organization
```

```
"TBD2"
```

```
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;
email: Jeff.Tantsura@ericsson.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 "2015-06-20" {
```

```
  description
```

```
    "I2RS Filter-Based RIB protocol ";
```

```
  reference "TBD";
```

```
}
```

```
typedef fb-rib-policy-type {
```

```
  type identityref {
```

```
    base "fb-rib-policy-type";
```

```
  }
```

```
  description
```

```
    "This type is used to refer to FB-RIB type";
```

```
  }
```

```
  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-oper-status-base";
    }
    description
    "This type is used to refer to FB-RIB type";
}

identity fb-rule-inactive {
    type identityref {
        base fb-rule-status;}
    description
    "policy rule is inactive";
}

identity fb-rule-active {
    type identityref {
        base fb-rule-status;}
    description
    "policy rule is active";
}

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

group-fb-rib-group-order-status {
    leaf group-order{ type uint16;
        description "default group order";
    }
    leaf group-refcnt {type uint16
        description "refcnt for this group";
    }
    leaf group installed {type uint16;

```

```
        description "number of rules installed";
    }
}
```

4.2. fb-fib yang module

```
module fb-rib {

    yang-version "1";

    // namespace
    namespace "urn:TBD1:params:xml:ns:yang:rt:i2rs:fbrib";
    // replace with iana namespace when assigned
    prefix "i2rs-fb-rib";

    // import some basic inet types
    import ietf-inet-types { prefix inet; } // RFC6991
    import ietf-interfaces {prefix "if";}
    import i2rs-rib {prefix i2rs-rib; }
    import ietf-routing {prefix "rt"};
    import i2rs-fb-rib-types {prefix fb-rib-types}
    import access-control-list {prefix ietf-acl;}
    import bnp-eca-policy {prefix eca-pol;}
    // meta
    organization
        "TBD2"

    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;
         email: Jeff.Tantsura@ericsson.com;
        "

    description
        "This module describes a YANG model for the I2RS
         Filter-based RIB which is a protocol independent I2RS module.

    /// top level FB-RIB structure
    container routing-instance {
```

```
description
  "Configuration of an 'i2rs' pseudo-protocol instance
  consists of a list of ribs.";

  // name of instance
leaf name {
  description
    "A routing instance is identified by its name,
    INSTANCE_name. This MUST be unique across all routing
    instances in a given network device.";
  type string ;
  mandatory true;
}

//
list interface-list {
  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.";
  key "name";
  leaf name {
    type if:interface-ref;
    description
      "A reference to the name of a configured network layer
      interface.";
  }
}
container router-id {
  uses rt:router-id;
}

container i2rs-fb-rib {
  key "fb-rib-name";
  container fb-rib-name {
    leaf name {type string};
  }
  list interface-list {
  description
    "This represents the list of interfaces associated
    that this Filter-Based RIB runs on.";
  key "name";
  leaf name {
  }
}
```

```

type if:interface-ref;
description
  "A reference to the name of a configured network layer
  interface.";
}
}

    container fb-rib-group-list {
      description "lists of groups of ordered lists";
      leaf group-name {type string};
      container fb-group-status {
        uses fb-group-order_status;
      }
      case fb-rib-acl {
        uses acl:access_lists: access-list-entries;
        augments acl:access-lists:access-list-entries;
        uses fb-rib-group-order_status;
        augments acl:access-lists:access-list-entries;
      }
      case fb-eca-rules {
        uses bnp-eca:bnp-policy-sets
      }
      augments bnp-eca:bnp-policy-set:\
        rule-group-list:rule-group;
      uses fb-rib-group-order_status;
      augments bnp-eca:bnp-policy-set:\
        rule-group-list:rule-group:rule-list:rule
        uses fb-rib-rule-order_status
    }
    container fb-rib_status {
      leaf fb-rib-update-ref;
      type unit64;
      description "reference count for fb-rib";
    }
  }
}
}

```

5. IANA Considerations

TBD.

6. Security Considerations

TBD.

7. References

7.1. Normative References:

[I-D.hares-i2rs-bnp-info-model]

Hares, S., Wu, Q., Tantsura, J., and R. White, "An Information Model for Basic Network Policy and Filter Rules", draft-hares-i2rs-bnp-info-model-02 (work in progress), March 2015.

[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-09 (work in progress), March 2015.

[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-00 (work in progress), April 2015.

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

Bahadur, N., Folkes, R., Kini, S., and J. Medved, "Routing Information Base Info Model", draft-ietf-i2rs-rib-info-model-06 (work in progress), March 2015.

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

Bogdanovic, D., Sreenivasa, K., Huang, L., and D. Blair, "Network Access Control List (ACL) YANG Data Model", draft-ietf-netmod-acl-model-03 (work in progress), June 2015.

7.2. Informative References

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

Hares, S. and M. Chen, "Summary of I2RS Use Case Requirements", draft-hares-i2rs-usecase-reqs-summary-02 (work in progress), May 2015.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

Authors' Addresses

Sriganesh Kini
Ericsson

Email: sriganesh.kini@ericsson.com

Susan Hares
Huawei
7453 Hickory Hill
Saline, MI 48176
USA

Email: shares@ndzh.com

Linda Dunbar
Huawei
USA

Email: linda.dunbar@huawei.com

Anoop Ghanwani
Dell

Email: anoop@alumni.duke.edu

Ram Krishnan
Dell

Email: Ramkri123@gmail.com

Dean Bogdanovic
Juniper Networks
Westford, MA

Email: deanb@juniper.net

Jeff Tantsura
Ericsson

Email: jeff.tantsura@ericsson.com

Russ White
Ericsson

Email: russ@riw.us