

I2RS working group
Internet-Draft
Intended status: Standards Track
Expires: September 9, 2015

S. Kini
Ericsson
S. Hares
Huawei
A. Ghanwani
Dell
R. Krishnan
Brocade
Q. Wu
Huawei
D. Bogdanovic
Juniper Networks
J. Tantsura
R. White
Ericsson
March 8, 2015

Filter-Based RIB Information Model
draft-kini-i2rs-fb-fib-info-model-00

Abstract

This document defines an information model I2RS Filter based RIB (Routing information Model). Filter based forwarding matches fields in the IP header plus other higher layer packet information. These matches may be ordered. Matches may contain actions which could impact forward, such as setting a nexthop.

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 9, 2015.

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 Routing Information Model Overview [5](#)
 - [3.1.](#) Scope [6](#)
 - [3.2.](#) Generic Rules for Filter-Based RIBS [6](#)
- [4.](#) Filter-Based-RIB module [9](#)
 - [4.1.](#) FB-RIB entries [11](#)
 - [4.2.](#) FB-RIB Description [12](#)
 - [4.3.](#) Rules on Order Rule [13](#)
 - [4.4.](#) I2RS FB-RIB interaction with configured filter rules . . [15](#)
 - [4.5.](#) Relationship between RB-RIB Rule Model and RIB Information Model [15](#)
- [5.](#) L3 Match-Action Rules [16](#)
- [6.](#) Open issues [17](#)
- [7.](#) IANA Considerations [18](#)
- [8.](#) Security Considerations [18](#)
- [9.](#) References [18](#)
 - [9.1.](#) Normative References: [18](#)
 - [9.2.](#) Informative References [18](#)
- 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 generic information model for a I2RS filter based RIB (FB-RIB) and describes the I2RS interaction with routing filters within a routing element.

Filter based (FB) routing matches fields in the IP header plus other higher layer packet information. Filters with a match-action pair allow the filters to impact the forwarding of packets. Actions may impact forwarding or set something in the packet that will impact forwarding.

A Filter-Based RIB (Routing Information Base) contains a list of filters (match-action conditions) and a default RIB of the form found in [[I-D.ietf-i2rs-rib-info-model](#)]. The default RIB routes any packet not matched by the order list of filter rules. If any packet does not match filter, it is dropped.

Some drafts which provide models for match filters are the following:

- o Access lists (ACLs) [[I-D.ietf-netmod-acl-model](#)] (Note: This filter provides match-action filters),
- o routing filter policy based on filters for IP prefixes (IPv4, IPv6) (E.g. [[I-D.yan-rtgwg-routing-policy-yang](#)]) or ordered prefix lists (e.g. [[I-D.zhdankin-idr-bgp-cfg](#)]),
- o generic match-policy filters that support QoS filters (E.g. [[I-D.hares-i2rs-bnp-info-model](#)]).
- o routing filters that include BGP originated routes tracked by BGP attribute(asPath, BGP community, extended BGP community, RDs) or peer ([[I-D.shaikh-rtgwg-policy-model](#)]),
- o Filters passed in BGP for flows (E.g. [[RFC5575](#)]))

This generic model for filters aligns with the generic model for topology in providing a simple model that can be utilize for other filters. The abstract filter model utilizes a generic filter based model that can be applied for specific filters at each level. The default RIB specification for the FB-RIB uses the I2RS RIB Model.

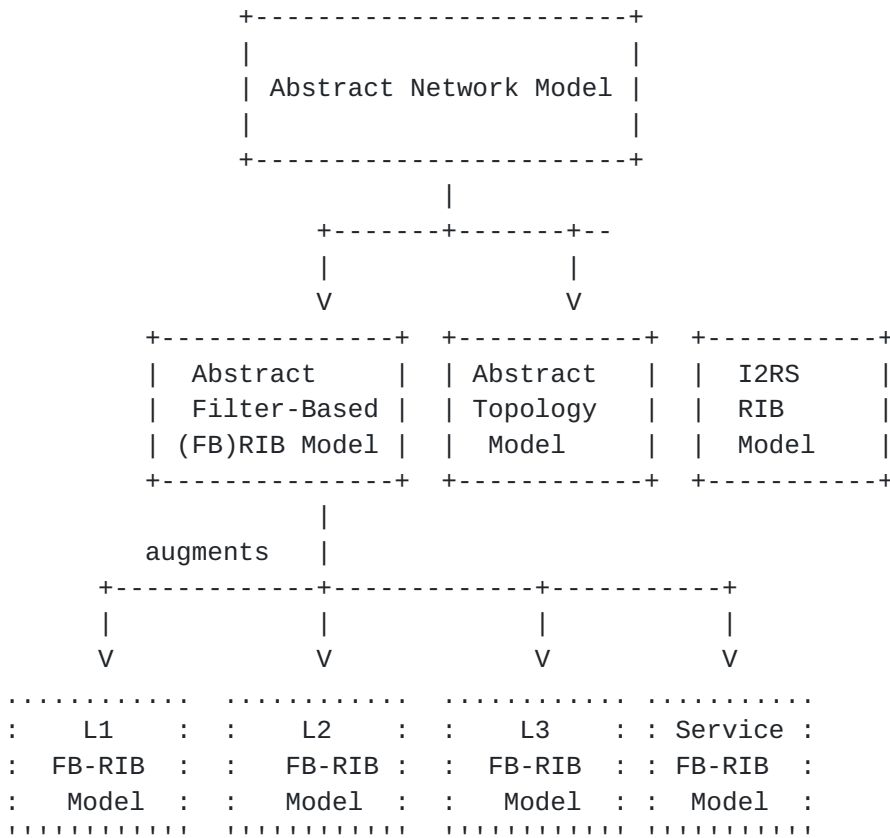


Figure 1: The network model structure

2. Definitions and Acronyms

CLI

Command Line Interface

FB Default RIB

The FB Default RIB is the default Routing Information Based use based for forwarding traffic for routes which do not match any FB-RIB Rule.

FB-RIB

Filter-Based Routing Information Base

IGP

IGP is an Interior Gateway Protocol

PCIM

Policy Core Information Model directly and indirectly the work of the PCIM Working Group.

Policy Rule

The PCIM framework defines a policy rule is often represented by "if Condition then action". The action may have set, modify, or notify actions. This draft uses the filters in [[I-D.hares-i2rs-bnp-info-model](#)], but policy can be used from a variety of filters.

Policy Group

The PCIM Framework defines policy groups as a group of policy rules into ordered and prioritized groups of policy.

Policy Set

The PCIM framework defines a the Policy set (specifically the PolicySetComponent) as an aggregation class that allows aggregation of Policy Groups and the nesting of Policy Groups under Policy set rules. The PolicySet rules include nesting policies and matching strategies (all-matching or first-match), priorities between rules, and roles. One of the roles that must be conditionally matched is the models denotation of "read-only" or "read-write" policy rules into ordered and prioritized groups of policy. The [[I-D.hares-i2rs-bnp-info-model](#)] suggests that non-nested policy groups may be sufficient for I2RS status and configuration work.

RIB IM

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

Routing instance

Routing Code often has the ability to spin up multiple copies of itself into virtual machines. Each Routing code instance or each protocol instance is denoted as N_INSTANCE in the text below.

3. Filter-Based Routing Information Model Overview

Filter based routing is a technique used to make packet routing decisions based on filter policies set by the network administrator. Routing decisions in a Filter-Based RIB (FB-RIB) are based on several criteria beyond destination address, such as application, IP protocol

used, identity of the end system, and even packet size. Policy actions are typically applied before applying QoS constraints since policy actions may override QoS constraint.

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: Protocol independent Use cases and large flow use cases described in [[I-D.hares-i2rs-usecase-reqs-summary](#)].

The Filter-based policies are specified in most routers/switches as an ordered set of rules. Each policy rule has a set of match conditions, and a set of actions which may include forwarding actions and QoS actions. This draft uses a generic description of filters rules described in [[I-D.hares-i2rs-bnp-info-model](#)], but other policy models could be used if they have the same characteristics.

(Note: Antecedents of this generic structure for filter/policy rules can be found in the IETF PCIM work ([[RFC3060](#)], [[RFC3460](#)], and [[RFC3644](#)]).)

[3.1.](#) Scope

A Filter-Based RIB (FB-RIB) information model can be considered in either a top-down view examining the filter policy which controls the RIBs or from a bottom-up view which considers the data plane. A top-down view considers how the I2RS client provides filters for what can be added to a FB-RIB. This draft takes a bottom-up approach and looks at just the routes being installed in the FB-RIB. The bottom-up view considers how routes link to forwarding data planes that must be supported. In this view, the match filters must consider IP [both IPv4 and IPv6], but may also consider MPLS and encapsulated protocols such as TCP [[RFC0793](#)], UDP [[RFC0768](#)], STCP [[RFC4960](#)], ICMP [[RFC0792](#)]. This draft takes the bottom-up viewpoint which looks at how the FB-RIB controls the data plane.

This provides a generic FB-FIB description in [section 4](#), and provide FB-FIB extension to cover the L3 IP filter covering IPv4 [[RFC0791](#)] and IPV6 [[RFC2460](#)] in [section 5](#).

[3.2.](#) Generic Rules for Filter-Based RIBS

Generic filter rules are described in [[I-D.hares-i2rs-bnp-info-model](#)]. The filter rules are included as list of groups of rules which in turn contain rules. This grouping hierarchy allows the ordering of all rules, and a logical group of filter rules based on a logical group (E.g. customers).

Within a particular order (E.g. Order 2), priority will establish the filter sequence within the order. If two priorities match, it is assumed the ordering of the filters do not impact the level

Each Rule within the Rule list has a rule-action match condition which is based on type. Type can be the "generic filter match-actions" or match actions specific to another type of policy (e.g. ACL rule match-action). For the generic filter match-actions has match field (bnp-term-match), action field (bnp-action), and a forwarding field (bnp-generic forwarding) as figure 1 shows.

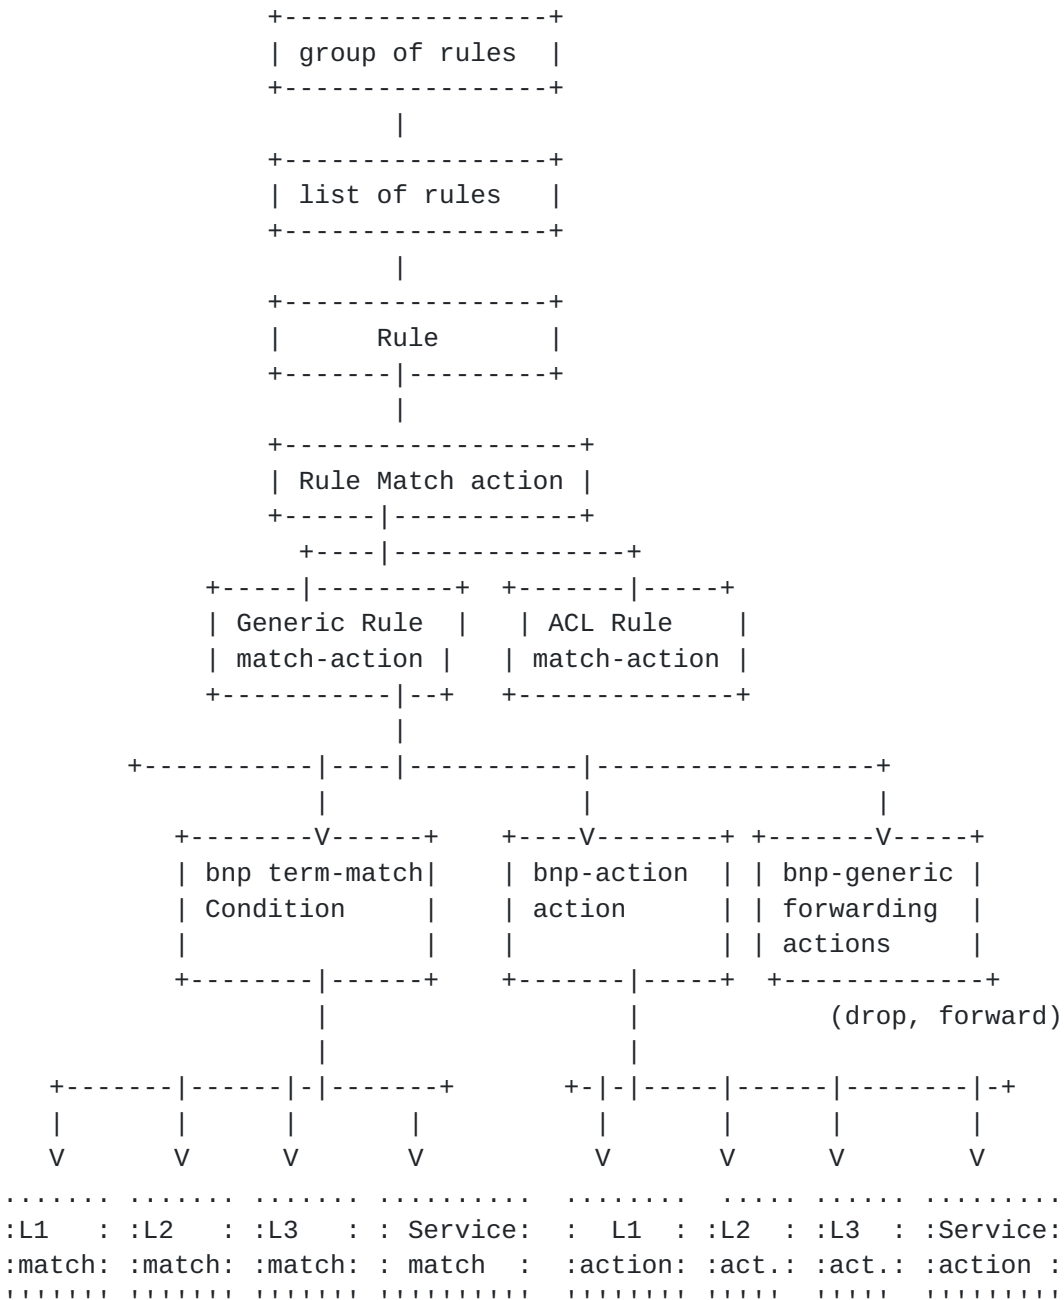


Figure 2

An example of this hierarchy is shown in figure 2:


```

Group
  Name: internal-nets
  Scope: L3-FB-RIB, R/W
  group-installer: v-netops
  rule-list
rule-1;
  name: v-netops-lan
  order: 1
  installer: v-netops
  status
    ro-status: active
    ro-rule-inactive-reason null
    ro-iule-installer: v-netops
  priority 1
  rule-match-act
    case:BNP-GENERIC-MATCH-ACTION
    Case:L3-Header
      term-match DEST-Header 192.200.1.* /24
      term-action:
        n-acts: 0
      term-forward: drop

  rule-2
  name:ICMP packets
  order: 2
  installer: v-netops
  status:
    ro-status: inactive
    ro-rule-inactive-reason: admin-inactive
  ro-installer-active-filter: (null)
  priority 3
  rule-match-act:
    Case:BNP-GENERIC_MATCH-ACTION
    Case:L3-Header
      term-match: ICMP-Type
      term-action:
        n-acts: 0
    term-forward: drop

```

Figure 3: Example structure

4. Filter-Based-RIB module

A Filter-Based RIB (FB-RIB) is an entity that contains an ordered set of filters (match/action conditions) and a default RIB of the form found in [[I-D.ietf-i2rs-rib-info-model](#)]. 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, a FB default-RIB, and list of FB-RIBs. Only some of the interfaces associated with a routing instance may be associated with a FB-RIB. Each interface can be associated with at most one FB RIB.

Packets arriving on an interface associated with a FB-RIB will be forwarded based on filters in a FB-RIB or the FB-RIB Default RIB (if no matches occur). 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 the FB-RIBs on the FB-RIB list, then the packet will be forwarded using the Default-RIB specified by the FB-RIB if one exists. If no Default-RIB is specified, the packet will be discarded.

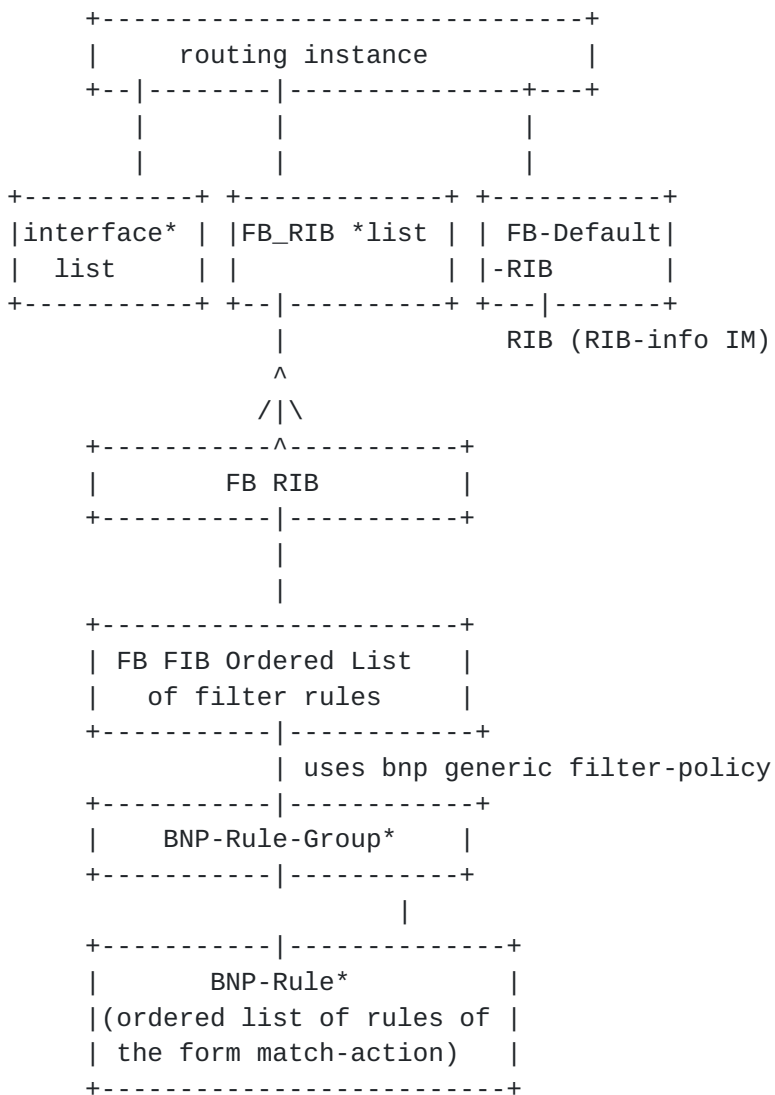


Figure 4: Routing instance with FB-RIB

4.1. 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-0-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 uint32
      +--rw FB-RIB-interface*
        | +--rw FB-RIB-interface interface-ref-id
        +--rw FB-Default-RIB rib-ref
    +--rw FB-RIB
      +--rw FB-RIB-Name
        +--rw FB-RIB-AFI
        +--rw FB-RIB-intf*
        +--rw FB-FIB-status-info
        | +--rw fb-rib-update-ref uint64
        +--rw FB-RIB-Ordered-Filters
        uses bnp-policy for filters
        augments /nt:bnp-generic-rules/rule-group/
  
```

Figure 4: FB RIB Yang Structure

4.2. FB-RIB Description

Each FB-RIB has the following:

- o FB-RIB-Name - Name identifier for FB- RIB
- o FB-RIB-AFI - AFI Supported by the FB-RIB

- o FB-RIB-intf* - Interface FB-FIB operates on. Note that an interface can be associated with at most one FB-RIB. For example interfaces eth1 and eth2 can be associated to FB-RIB, but these two interfaces cannot be connected to any other FB-RIB.
- o FB-RIB-Status-info - status at RIB level which includes number of times since reconfiguration this FB-RIB has been updated.
- o FB-RIB-Ordered-Filters contains list of rule groups
 - * Each rule-group is indexed by group name contains:
 - + group-name (string)
 - + status-info which contains two elements:
 - group status (installed, active or inactive).
 - inactive reason (null, policy-conflict, unsupported).
 - group-installer-identity (string)
 - + group order (uint16)
 - + ordered rule list

4.3. Rules on Order Rule

This section provides a short description of the generic filter policy rule's condition-action from [[I-D.hares-i2rs-bnp-info-model](#)] which is used by the FB-RIB.

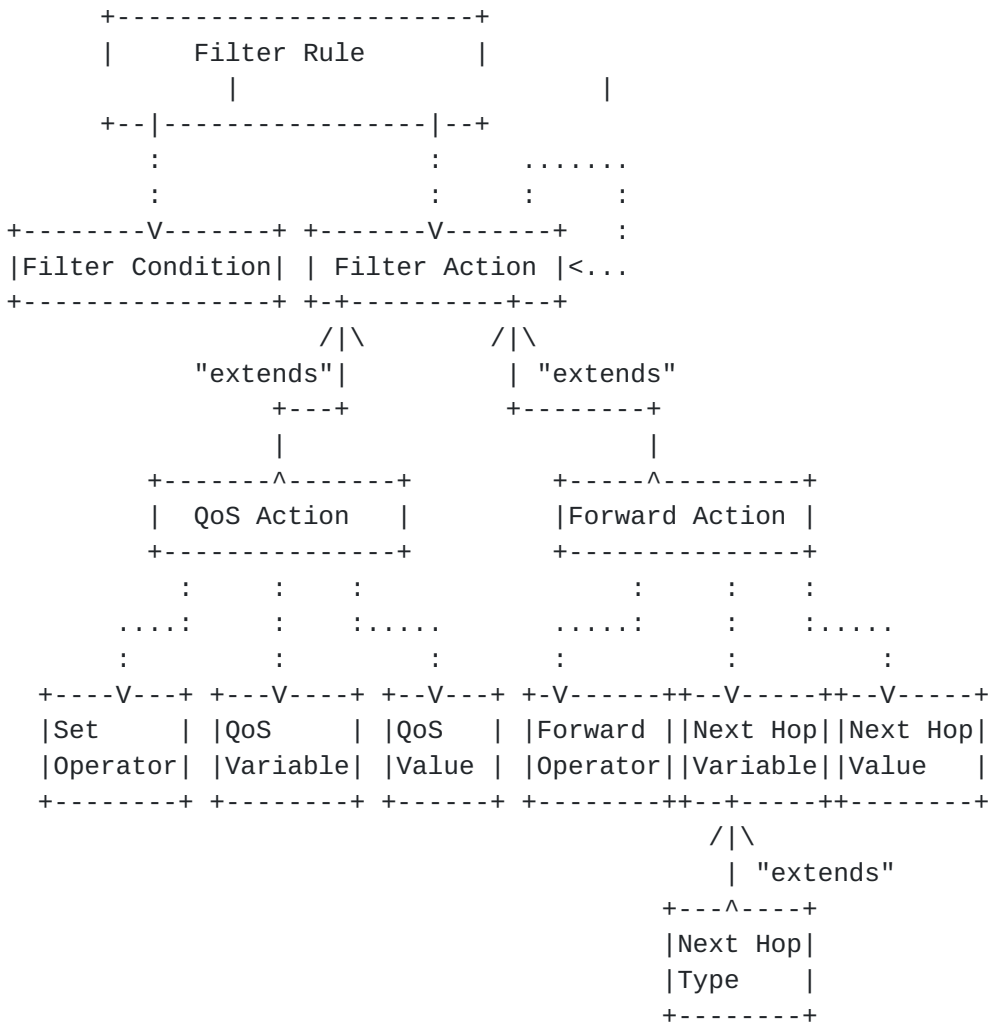


Figure 5: Filter Actions for FB-RIB

The policy/filter rule contains the following:

- o rule-ref - ordered id number for the policy rule
- o rule-status-info - status on the policy rule that contains the following:
 - * rule-status - installed, active, or inactive.
 - * rule-inactive-reason - can be null, policy-conflict, i2rs-rule-supersedes, unsupported)
 - * rule-installer - the entity that installed rule.
- o match-filter - ordered match field for FB-RIB route entry which contains:

- * order- order number in match sequence
- * match-term - contains matches for filters for different packets based on L1, L2, L3, transport, or service level.
- * rule-action* - An ordered list of policy actions that includes the following:
 - + n-acts - number of actions
 - + Actions: set values in one or more of the following:
 - + forwarding-actions - which includes
 - std-forwarding - (enumeration) forwarding packet
 - o Drop_Packet - drop packet
 - o Drop_Packet_ICMP - dropping packet with ICMP unreachable sent
 - o Forward_Packet_specific - send to specific next hop
 - o Forward_Packet_default - forward based on FB-RIB Default RIB

4.4. I2RS FB-RIB interaction with configured filter rules

The I2RS client-agent pair process within a routing process to add ephemeral these changes to the filter State so that

FB-RIB-rules(running) = FB-RIB-config + FB-Rules-I2RS-ephemeral

The I2RS ephemeral state will not survive a reboot of the machine. Upon a reboot, the I2RS client must reload the I2RS Agent with the I2RS FB-RIB state lost in the reboot.

Writing I2RS FB-rules to permanent configuration may be desirable. This has not been considered in this version of this draft.

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

The RIB in a router with I2RS is the following:

running RIB = configured-RIB + routes-installed-from-protocols + I2RS-ephemeral-state

As described in [[I-D.ietf-i2rs-rib-info-model](#)], the I2RS ephemeral RIB information in routing instance contains a collection of RIBs, interfaces, and routing parameters including the following:

- 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.

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.

5. L3 Match-Action Rules

Layer 3 match might contain the following:

- o IPv4 header match with one or these fields: IPv4 source address, IPv4 destination address, IPv4 Protocol, IPv4 TOS/DSCP field, IPv4 ICMP field, and the length of the packet. These matches can be exact matches, longest prefix matches for addresses, or range matches for values in TOS/DSCP field, ICMP field or length of packet.
- o IPv6 header match with one or more match of IPv6 source address, IPv6 destination address, IPv6 Traffic class (DSCP), IPv6 Flow label, IPv6 payload length, IPv6 next-header, hop-limit. These matches can be exact matches, longest prefix matches for addresses, or range matches.

Layer 3 Actions might set values in:

- o In IPv4 packets set values in any of the following fields: IPv4 source address, IPv4 destination address, IPv4 Protocol, IPv4 TOS/DSCP field, IPv4 ICMP field or the length of the packet. (Please note that hardware data plane forwarders may only be able to set TOS/DSCP while software data plane forwarders may be able set additional fields.)
- o In IPv6 packets set values in any of the following fields: IPv6 source address, IPv6 destination address, IPv6 Protocol value, IPv6 Flow, or IPv6 packet length.

Layer 3 Forwarding can augment the basic to forward via tunnels.

6. Open issues

This section record the issues with the initials of the person who recorded it.

Forwarding per interface (JMH)

- The authors believe the forwarding per interface is covered by the attachment of a FB-RIB to interface-list.

Centralized or Distributed filter policy Strategy (JMH)

The authors believe this structure can be used by either centralized or distributed forwarding for configuration or the I2RS ephemeral data structure

policy database-enforcement points architecture (JMH)

The authors believe this yang modules describes the filters which provides a specific enforcement of forwarding policy. The wider constraints of how filter policy is stored as filter rules or groups of filters rules can be done as the generic network policy as described in [[I-D.hares-i2rs-bnp-info-model](#)] or other policy. Other forms of policy rule filter sets can be used.

policy rule conflicts (JMH)

Detection of filter rule conflicts are done by the agent module receiving the filters from configuration or ephemeral I2RS stream. The filter can be reject or installed and rejected from active use due to conflicts at either a group level or the filter rule level. At the policy group level the group-policy-status-info contains a status of installed, active, or installed-inactive. If the status is inactive the group-policy-inactive-reason can indicate policy-conflicts. The policy-rule has a similar status (policy-rule-

status-info with policy-rule-status and policy-rule-inactive-reason).

7. IANA Considerations

This draft includes no request to IANA.

8. Security Considerations

TBD.

9. References

9.1. Normative References:

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

Hares, S. and Q. Wu, "An Information Model for Basic Network Policy", [draft-hares-i2rs-bnp-info-model-01](#) (work in progress), October 2014.

[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-info-model]

Bahadur, N., Folkes, R., Kini, S., and J. Medved, "Routing Information Base Info Model", [draft-ietf-i2rs-rib-info-model-05](#) (work in progress), January 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-02](#) (work in progress), March 2015.

9.2. 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-01](#) (work in progress), October 2014.

- [I-D.shaikh-rtgwg-policy-model]
Shaikh, A., Shakir, R., D'Souza, K., and C. Chase,
"Routing Policy Configuration Model for Service Provider
Networks", [draft-shaikh-rtgwg-policy-model-00](#) (work in
progress), January 2015.
- [I-D.yan-rtgwg-routing-policy-yang]
Yan, G. and S. Zhuang, "Yang Data Model for Routing
Policy", [draft-yan-rtgwg-routing-policy-yang-00](#) (work in
progress), December 2014.
- [I-D.zhdankin-idr-bgp-cfg]
Alex, A., Patel, K., Clemm, A., Hares, S., Jethanandani,
M., and X. Liu, "Yang Data Model for BGP Protocol", [draft-
zhdankin-idr-bgp-cfg-00](#) (work in progress), January 2015.
- [RFC0768] Postel, J., "User Datagram Protocol", STD 6, [RFC 768](#),
August 1980.
- [RFC0791] Postel, J., "Internet Protocol", STD 5, [RFC 791](#), September
1981.
- [RFC0792] Postel, J., "Internet Control Message Protocol", STD 5,
[RFC 792](#), September 1981.
- [RFC0793] Postel, J., "Transmission Control Protocol", STD 7, [RFC
793](#), September 1981.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6
(IPv6) Specification", [RFC 2460](#), December 1998.
- [RFC3060] Moore, B., Ellesson, E., Strassner, J., and A. Westerinen,
"Policy Core Information Model -- Version 1
Specification", [RFC 3060](#), February 2001.
- [RFC3460] Moore, B., "Policy Core Information Model (PCIM)
Extensions", [RFC 3460](#), January 2003.
- [RFC3644] Snir, Y., Ramberg, Y., Strassner, J., Cohen, R., and B.
Moore, "Policy Quality of Service (QoS) Information
Model", [RFC 3644](#), November 2003.
- [RFC4960] Stewart, R., "Stream Control Transmission Protocol", [RFC
4960](#), September 2007.

[RFC5575] Marques, P., Sheth, N., Raszuk, R., Greene, B., Mauch, J.,
and D. McPherson, "Dissemination of Flow Specification
Rules", [RFC 5575](#), August 2009.

Authors' Addresses

Sriganesh Kini
Ericsson

Email: sriganesh.kini@ericsson.com

Susan Hares
Huawei
7453 Hickory Hill
Saline, MI 48176
USA

Email: shares@ndzh.com

Anoop Ghanwani
Dell

Email: anoop@alumni.duke.edu

Ram Krishnan
Brocade

Email: ramk@Brocade.com

Qin Wu
Huawei
101 Software Avenue, Yuhua District
Nanjing, Jiangsu 210012
China

Email: bill.wu@huawei.com

Dean Bogdanovic
Juniper Networks
Westford, MA

Email: deanb@juniper.net

Jeff Tantsura
Ericsson

Email: Jeff Tantsura jeff.tantsura@ericsson.com

Russ White
Ericsson

Email: russw@riw.us