

I2RS working group
Internet-Draft
Intended status: Standards Track
Expires: April 30, 2015

S. Hares
Huawei
S. Kini
Ericsson
A. Ghanwani
Dell
R. Krishnan
Brocade
Q. Wu
Huawei
D. Bogdanovic
Juniper Networks
October 27, 2014

An Information Model for Basic Network Policy
draft-hareskini-i2rs-pbr-info-model-00

Abstract

This document defines the I2RS Policy-Based Routing (PBR) policy information model describing I2RS interactions with the PBR in a routing system. The PBR IM uses Policy Core Information Model (PCIM) framework ([RFC3060](#), [RFC3460](#), and [RFC3644](#)) to specify the ordered route list within the PBR RIB adapted to I2RS.

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 April 30, 2015.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

Internet-Draft

IM for policy

October 2014

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	3
3.	The Policy Based Routing Information Model Overview	4
3.1.	Scope	5
4.	PBR-RIB module	5
4.1.	PBR RIB	9
4.2.	PBR Rule Component	10
4.3.	I2RS PRB RIB interaction with PBR RIB	12
5.	Relationship between PBR Rule Model and RIB Information Model	13
6.	Discussion of I2RS related issues	14
7.	IANA Considerations	14
8.	Security Considerations	15
9.	Informative References	15
	Authors' Addresses	16

[1.](#) Introduction

The Interface to the Routing System (I2RS)

[[I-D.ietf-i2rs-architecture](#)] architecture calls out for 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 I2RS Policy-Based Routing (PBR) Information model defined in this document describes the I2RS interaction with PBR within a routing element.

The PBR requires an ordered list of policy. This PBR informational model uses the Policy Core Information Model (PCIM) framework as described in [[RFC3060](#)] with its extensions in [[RFC3460](#)] and QOS model in [[RFC3644](#)]. The adaptation of the PCIM model to I2RS use is

described in [[I-D.hares-i2rs-bnp-info-model](#)].

Internet-Draft

IM for policy

October 2014

[2.](#) Definitions and Acronyms

CLI

Command Line Interface

IGP

IGP is an Interior Gateway Protocol

Information Model

is an abstract model of a conceptual domain, independent of a specific implementations or data representation

MPLS

Multi-Protocol Label Switching.

NETCONF

The Network Configuration Protocol

PBR

Policy Based Routing.

PBR Default RIB

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

PBR-RIB

Policy Based Routing-Routing Information Base

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. The [[I-D.hares-i2rs-bnp-info-model](#)] provides

Hares, et al.

Expires April 30, 2015

[Page 3]

Internet-Draft

IM for policy

October 2014

examines of how ACLs, Prefix lists, and more complex BGP policy can be combined into a policy rule.

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 initial I2RS 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.

SNMP

The Simple Network Management Protocol

3. The Policy Based Routing Information Model Overview

Policy Based Routing (PBR) is a widely used term in the industry to describe a technique used to make packet forwarding decisions based on policies set by the network administrator. PBR enables network administrator to forward the packet based on other criteria than the destination address in the packet, which is used to lookup an entry in the routing table.

The PBR problem can be viewed as a resource allocation problem that incorporates business decision with routing. PBR may be used to

provide many benefits, including better resource allocation, load balancing and QoS.

Routing decisions in PBR 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 overrides QoS constraint.

The I2RS use cases which benefit from PBR are: Protocol independent Use cases and large flow use cases described in [\[I-D.hares-i2rs-usecase-reqs-summary\]](#)

The PBR 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. Since policy rules, groups of policy, and ordered sets of policy are used in other protocols (BGP or MPLS), these policy rules have been abstracted into a basic network policy instantiation of the PCIM ([\[RFC3060\]](#), [\[RFC3460\]](#), and [\[RFC3644\]](#)). This instantiation include in the ordered policy rule the references to other policy match-action conditions such as the ACLs ([\[I-D.bogdanovic-netmod-acl-model\]](#)), and

Prefix list ([\[I-D.zhdankin-netmod-bgp-cfg\]](#)).

[3.1.](#) Scope

A PBR IM can be considered in either a top-down view examining the policy which controls the data flow or from a bottom-up view which considers the data plane. A top-down view considers how policies control protocols (BGP or IGP (ISIS/OSPF)) transfer of routes to determine how data flows. The bottom-up view considers the 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 PBR RIB controls the data plane.

This draft considers match and action filters for the data-planes using IP (both IPv4 [[RFC0791](#)] and IPV6 [[RFC2460](#)]).

[4.](#) PBR-RIB module

A PBR-RIB is an entity that contains an ordered set of policy routes and is analogous to a RIB defined in [[I-D.ietf-i2rs-rib-info-model](#)]. An ordered set of policy routes implies that the insertion into a PBR-RIB must allow for inserting of a PBR route at any specific position and deleting a route at a specific position. The ability to


change a policy rule at a specific position combines these two functions (deleting an existing policy rule and adding a new policy rule).

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

Packets arriving on an interface associated with a PBR-RIB will be forwarded based on a PBR-RIB in the list or PBR Default RIB (if no matches occur). The policy processing within the PBR process within

the routing system is expected to do the following:

- o When a packet successfully matches a PBR Match term/entry, the corresponding policy-actions are applied.
- o If a packet does not match a PBR match term/entry, the PBR 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 PBR match term/entry is successfully matched, or no more filters in the list exists.
- o If no match has been found within the PBR filter list, then the packet will be forwarded using the PBR Default-RIB if one exists. If no PBR Default-RIB is specified, the packet will be discarded.



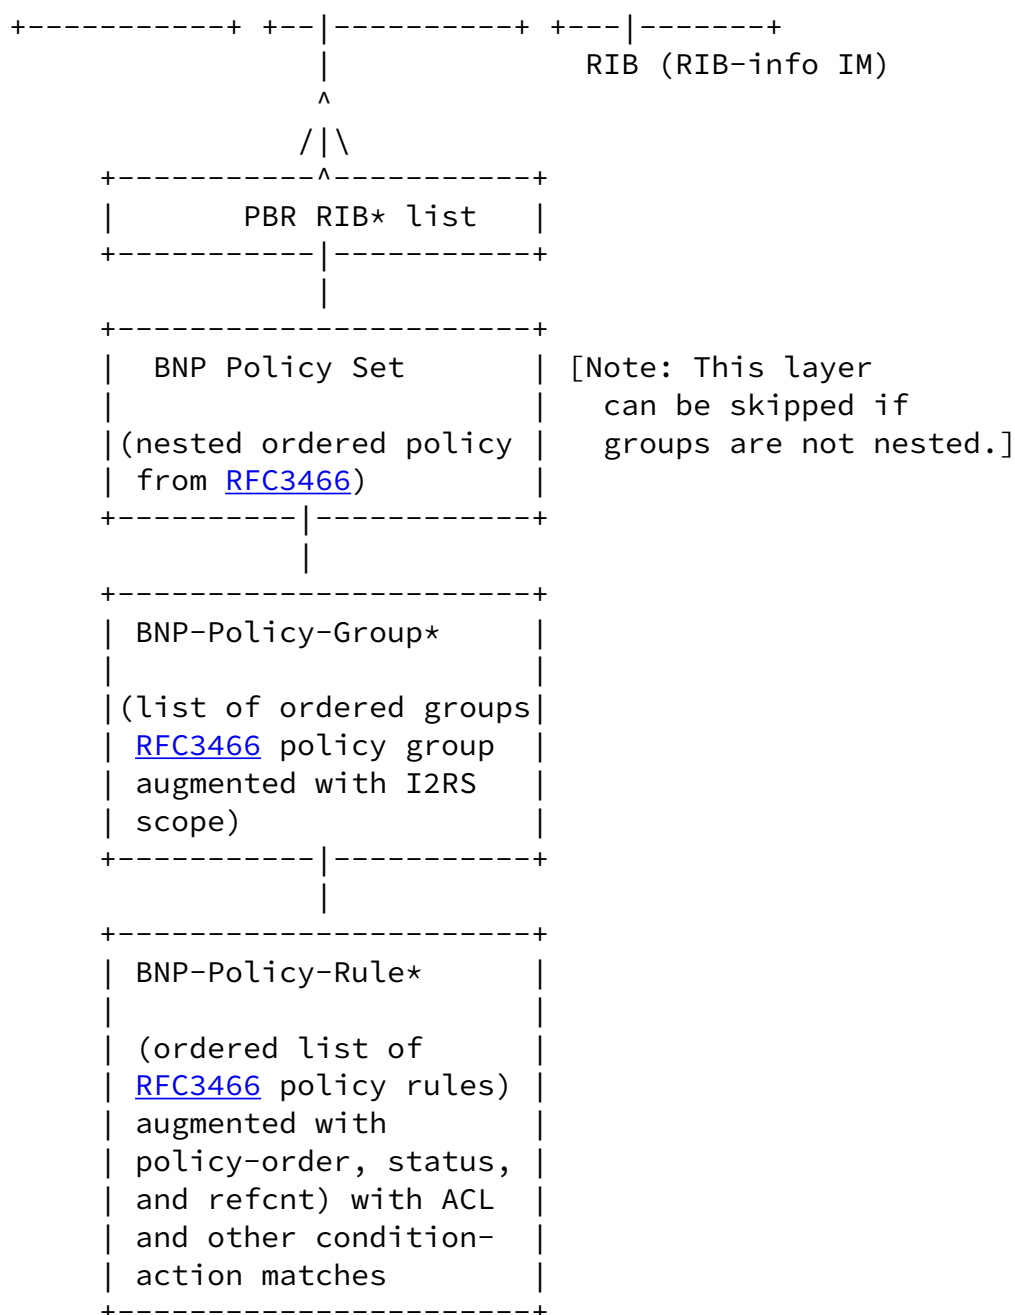


Figure 1: Routing instance with PBR RIB

The PBR entries associated with each PBR in a routing instance are:

Name of Routing instance

pbr-router-id

router id associated with the PBR function of the Routing instance

Interface_list

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

PBR Default RIB

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

PBR-RIB* list

list of PBR-RIBs

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

```
module: PBR
  +--PRB-RIB-module
    +--rw pbr-instance-name
    +--rw pbr-router-id  uint32
    +--rw pbr-interface*
    |   +--rw pbr-interface interface-ref-id
    +--rw PBR-Default-RIB
    +--rw PBR-RIB
      +--rw PBR-RIB-Name
      +--rw PBR-RIB-AFI
      +--rw PBR-RIB-intf*
      +--rw PBR-status-info
      |   +--rw pbr-update-ref uint64
      +--rw PBR-Ordered-Route-Policy
        +--rw pbr-group-policy* [group-policy-ref]
        |   +--rw group-policy-ref  uint16
        |   +--rw group-policy-name string
        |   +--ro group-policy-status-info
        |   |   +--ro group-policy-status
        |   |   +--ro group-policy-inactive-reason
        |   +--rw policy-rule* [policy-rule-ref]
```

```

|      +--rw policy-rule-ref
|      +--ro policy-rule-status-info
|      |  +--ro policy-rule-status enumeration
|      |  +--ro policy-rule-inactive-reason
|      +--rw pbr-match-filter* [nr-policy-match]
|          +--rw pbr-match-term
|              +--rw pbr-match-condition
|                  +--rw nr-policy-match
|                  +--rw pbr-ipv4-matches
|                  +--rw pbr-ipv6-matche
|                  +--rw pbr-transport-matches
|                  +--rw pbr-combo-operator
|              +--rw pbr-rule-action
|                  +--rw pbr-QoS-acts [nbr-act]
|                  +--rw npbr-act
|                  +--rw set-in-ipv4-packet
|                  |  ...
|                  +--rw set-in-ipv6-packet
|                  |  ...
|                  +--rw set-vendor
|                  |  . . .
|          +--rw pbr-forwarding-actions
|          +--rw pbr-std-fwd enumeration
|          +--rw pbr-vendor-fw enumeration
+--rw pbr-policy-set[policy-set-name]
+--rw policy-set-name
+  . . .

```

Figure 2: PBR RIB Yang Structure

[4.1.](#) PBR RIB

Each PBR RIB has the following:

- o PBR-RIB-Name - Name identifier for PBR RIB
- o PBR-RIB-AFI - AFI Supported by the PBR RIB
- o PBR-RIB-intf* - Interface PBR operates on. Note that an interface can be associated with at most one PBR RIB. For example interfaces eth1 and eth2 can be associated to PBR_RIB, but these two interfaces cannot be connected to any other PBR RIB.
- o PBR-Status-info - status at PBR RIB level which includes number of times since reconfiguration this PBR has been updated.

- o PBR-Ordered-Route-Policy contains two sub-elements:

- * pbr-group-policy - group policy list indexed by group-policy-ref number. Policy group contains a reference number (group-policy-ref), name, status-info, and a list of policy-rules. the group policy status can be one of the following: installed, active, inactive, I2RS-active, and I2RS-inactive). The inactive reason can be one of the following: null, poicy-conflict, i2rs-supersedes, unsupported).
- * pbr-policy-set - policy set identified by name

Initially, it is expected the simply group policy list will be sufficient. (See [[I-D.hares-i2rs-bnp-info-model](#)] for an examples of the policy rules can contain ACL policy, Prefix-list policy, and more complex (match/set) policy.)

[4.2.](#) PBR Rule Component

A PBR policy rule used by has the following general architecture.

Internet-Draft

IM for policy

October 2014

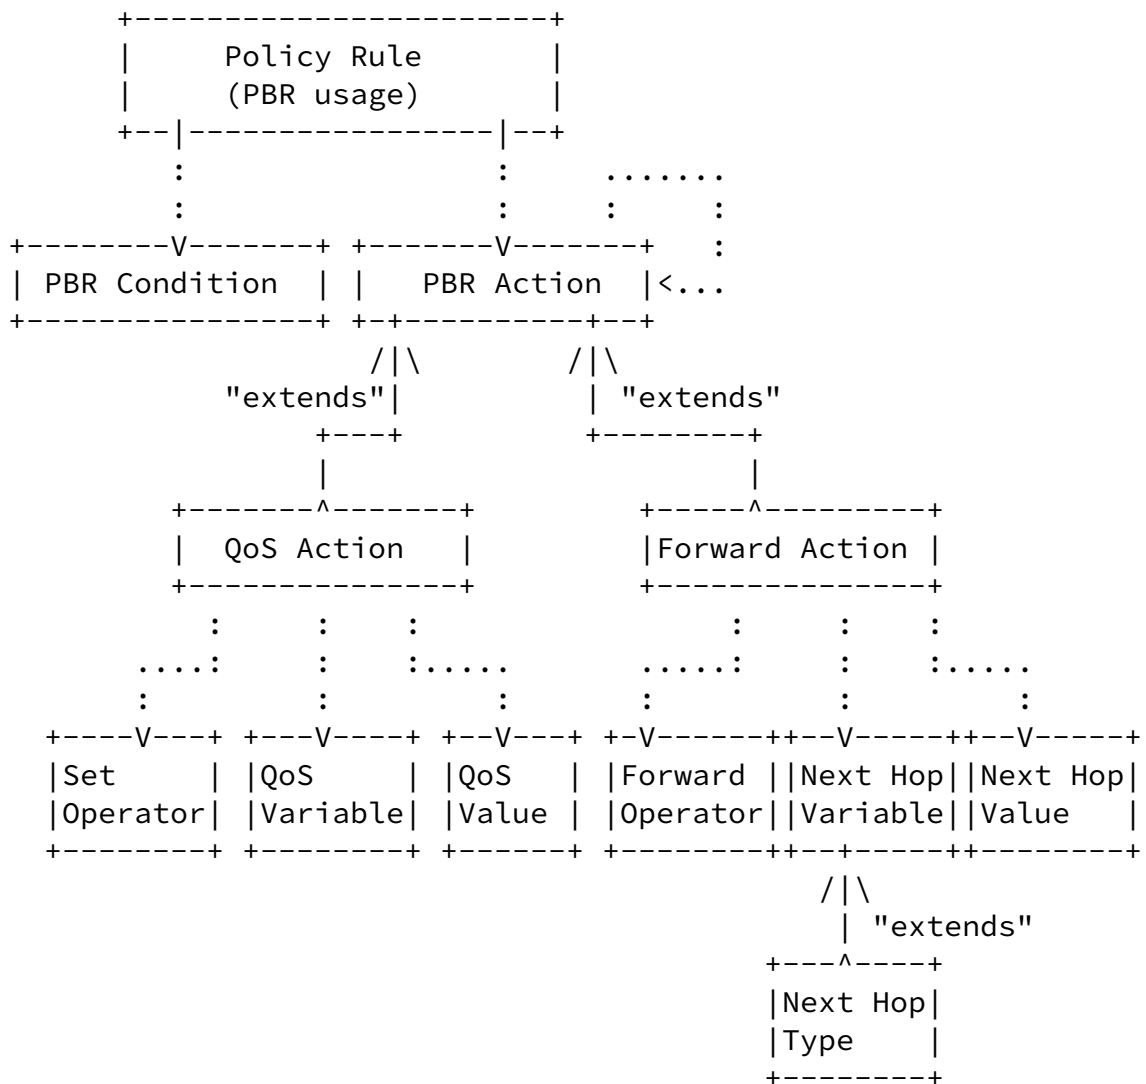


Figure 3: Policy Rules for PBR routing

The policy-rule contains the following:

- o PBR-match-filter - ordered PBR match field for a route entry which contains either:
 - * nr-policy-match - order number in match sequence
 - * pbr-ipv4-matches - one or more matches of 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.
 - * pbr-ipv6-matches - 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.
- * pbr-transport-matches - one or more matches in source port or destination port
 - * pbr-combo-operator - logical OR or logical AND that combine matches in one match filter.
- o pbr-rule-action* - An ordered list of policy actions that includes the following:
 - * npbr-acts - order number in action sequence
 - * Actions: set values in one or more of the following:
 - + IPv4 packets in 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.)
 - + IPv6 packets in IPv6 source address, IPv6 destination address, IPv6 Protocol value, IPv6 Flow, or IPv6 packet

length.

- * pbr-forwarding-actions - which includes
 - + pbr-std-forwarding - (enumeration) forwarding packet
 - Drop_Packet - drop packet
 - Drop_Packet_ICMP - dropping packet with ICMP unreachable sent
 - Forward_Packet_specific - send to specific next hop
 - Forward_Packet_default - forward based on PBR Default RIB
 - + pbr-vendor-fwd - Vendor specific action

[4.3.](#) I2RS PRB RIB interaction with PBR RIB

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

PBR-running = PBR-config + PBR-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 PBR RIB state lost in the reboot.

The PBR RIB module must allow both the I2RS client-agent to read the PBR IM as a query or as a notification stream. The pbr-update-ref parameter of the PBR-status-info provides an update count for the PBR configuration to indicate if the PBR has been updated with additions or deletions of the PBR policy rules. This provides the I2RS interface a quick way to check for changes by other entities to the PBR route list.

[5.](#) Relationship between PBR 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 PBR-RIB).
- o The routing parameters control the information in the RIBs.

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

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

[6.](#) Discussion of I2RS related 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 PBR to interface-list.

Centralized or Distributed Policy Strategy (JMH)

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

policy database-enforcement points architecture (JMH)

The authors believe this yang modules describes the PBR which provides a specific enforcement of forwarding policy. The wider constraints of how policy groups are stored, administered or distributed should be engaged at a higher layer. The authors note the Policy-Group project in OpenDaylight has an architecture for policy enforcement that renders the results to a particular instantiation in nodes. One such instantiation could be the I2RS policy.

policy rule conflicts (JMH)

Detection of policy rule conflicts are done by the policy module receiving the configuration or ephemeral I2RS stream. The policy can be reject or installed and rejected from active use due to conflicts at either the policy group level or the policy 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. Informative References

- [I-D.bogdanovic-netmod-acl-model]
Bogdanovic, D., Sreenivasa, K., Huang, L., and D. Blair, "Network Access Control List (ACL) YANG Data Model", [draft-bogdanovic-netmod-acl-model-02](#) (work in progress), October 2014.
- [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-00](#) (work in progress), September 2014.
- [I-D.hares-i2rs-usecase-reqs-summary]
Hares, S., "Summary of I2RS Use Case Requirements", [draft-hares-i2rs-usecase-reqs-summary-00](#) (work in progress), July 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-05](#) (work in progress), July 2014.
- [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-03](#) (work in progress), May 2014.
- [I-D.zhdankin-netmod-bgp-cfg]
Alex, A., Patel, K., and A. Clemm, "Yang Data Model for BGP Protocol", [draft-zhdankin-netmod-bgp-cfg-01](#) (work in progress), October 2014.
- [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., Ellessen, 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.

Authors' Addresses

Susan Hares
Huawei
7453 Hickory Hill
Saline, MI 48176
USA

Email: shares@ndzh.com

Sriganesh
Ericsson

Email: sriganesh.kini@ericsson.com

Anoop Ghanwani
Dell

Email: anoop@alumni.duke.edu

Internet-Draft

IM for policy

October 2014

Ram Krishnan
Brocade

Email: ramk@Brocade.com

Qin Wu
Huawei
Beijing
China

Email: Bill.Wu@huawei.com

Dean Bogdanovic
Juniper Networks
Westford, MA

Email: deanb@juniper.net

Hares, et al.

Expires April 30, 2015

[Page 17]