

I2RS working group
Internet-Draft
Intended status: Standards Track
Expires: April 27, 2015

S. Hares
Huawei
October 24, 2014

An Information Model for Basic Network Policy
draft-hares-i2rs-bgp-compare-yang-00

Abstract

This document contains a comparison of two BGP yang models: draft-zhdankin-netmod-bgp-cfg-01 and draft-wang-i2rs-bgp-dm. The yang model in draft-zhdankin-netmod-bgp-cfg-01 is a model focused on configuration. The yang model in draft-wang-i2rs-bgp-dm-00 is focused on the status and ephemeral state changes needed for the I2RS interface. The conclusion of comparison is that there little overlap except the definitions of common bgp structures. The draft-wang-i2rs-bgp-dm-00 is necessary to fulfil a majority of the requirement drawn from the BGP use cases in the I2RS use cases (draft-i2rs-hares-usecase-reqs-summary).

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 27, 2015.

Copyright Notice

Copyright (c) 2014 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	3
3. BGP Yang Draft Comparison	4
4. Differences between the drafts	4
4.1. Overlap of configuration draft-zhdankin-netmod-bgp-01	7
4.2. Unique configuration for draft-zhdankin-netmod-bgp-01	7
4.3. Unique configuration for draft-wang-bgp-dm-00	8
5. Comparison of State needed versus BGP requirements	9
5.1. BGP-REQ 01	9
5.2. BGP-REQ 02	10
5.3. BGP-REQ 03	13
5.4. BGP-REQ 04	14
5.5. BGP-REQ 05	15
5.6. BGP-REQ 06	16
5.7. BGP-REQ 07	17
5.8. BGP-REQ 08	18
5.9. BGP-REQ 09	19
5.10. BGP-REQ 10	20
5.11. BGP-REQ 11	22
5.12. BGP-REQ 12	22
5.13. BGP-REQ 13	25
5.14. BGP-REQ 14	25
5.15. BGP-REQ 15	27
5.16. BGP-REQ 16	27
5.17. BGP-REQ 17	28
5.18. BGP-REQ 18	28
6. IANA Considerations	29
7. Security Considerations	29
8. Informative References	29
Author's Address	30

1. Introduction

The Interface to the Routing System (I2RS) provides read and write access to the information and state within the routing process within routing elements. The I2RS client interacts with one or more I2RS agents to collect information from network routing systems. The processing of collecting information at the I2RS agent may require the I2RS Agent to filter certain information, group pieces of

information, or perform actions on the I2rs collected information based on specific I2rs policies.

This draft is a comparison of the following two BGP yang models: [I-D.zhdankin-netmod-bgp-cfg], and [I-D.wang-i2rs-bgp-dm]. The comparison provides an overview of the differences, overlaps, and unique features of each yang model. The analysis also evaluates whether both models or a single model is necessary to satisfy the requirements for the BGP use cases found in the [I-D.hares-i2rs-usecase-reqs-summary].

This draft concludes each of these two drafts focuses on the purpose each draft was created for (configuration, I2rs status and ephemeral state). The drafts have little overlap outside common definitions for some of the BGP functions. Both drafts reference bgp policy outside the basic structures (Prefix-lists, and policy-groups). Both drafts have concepts of AFI level and BGP neighbor level features. The status and ephemeral state found in [I-D.wang-i2rs-bgp-dm] is necessary to fulfil the BGP use cases found in the [I-D.hares-i2rs-usecase-reqs-summary]. Configuration knobs in [I-D.zhdankin-netmod-bgp-cfg] were helpful to support these bgp use cases, but the lack of state would not allow support of these features. The recommendation is that the two drafts harmonize the group structures and continue as two separate drafts for their original purpose (config and I2RS).

One BGP requirement (BGP-REQ18) requires a re-computation of the local BGP tables after policies have been modified by the ephemeral interface. The exact methodology of this re-computation could be part of ephemeral soft re-configuration. However, the I2RS WG has not determine how ephemeral state acts. Neither draft has created mechanism to do the ephemeral state re-configuration which is wise since both the I2RS WG and netmod WG should develop a joint ephemeral re-configuration process.

The rest of this draft is the details so those who desire "sounds bytes" level reading may stop reading now.

2. Definitions and Acronyms

BGP - Border Gateway Protocol version 4

CLI: Command Line Interface

IGP: Interior Gateway Protocol

I2RS: Interface to (2) Routing system.

Information Model: An abstract model of a conceptual domain, independent of a specific implementations or data representation

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 Foo_INSTANCE in the text below.

NETCONF: The Network Configuration Protocol

3. BGP Yang Draft Comparison

The authors of [I-D.zhdankin-netmod-bgp-cfg] focused on the configuration aspect of BGP (98%+ configuration, 2% status). The [I-D.wang-i2rs-bgp-dm] is focused on the I2RS need for status with a small amount of specific I2RS configuration for I2RS needs (95% status, 2% config). The two drafts can be combined together, but guidance is needed from the netmod group as the I2RS state is read-write ephemeral.

Why is draft-wang-bgp-dm-00 mostly focused on routes, statistics, and state?

The use cases specified in [I-D.hares-i2rs-usecase-reqs-summary] demonstrate a need for the status found in [I-D.wang-i2rs-bgp-dm] which includes: bgp-local-rib, bgp-rib-in, bgp-rib-out and all kinds of statistics and state, such as protocol-status , bgp-rt-state-info, peer-state-info, max-prefix-rcv-limit, etc. The status is both a the AFI state and the BGP peer state (within the AFI).

Two versions of [I-D.zhdankin-netmod-bgp-cfg] have been released - a -00.txt and a -01.txt the second version (-01 version) only updated references to netmod and states on the yang modules. This draft will use the -01 version of [I-D.zhdankin-netmod-bgp-cfg],

The [I-D.wang-i2rs-bgp-dm] was released mistakenly as [I-D.hares-i2rs-bgp-dm]. A few corrections to the status fields were included in the [I-D.wang-i2rs-bgp-dm]. This draft uses [I-D.wang-i2rs-bgp-dm] for the comparison.

4. Differences between the drafts

The [I-D.zhdankin-netmod-bgp-cfg] is composed by two parts: BGP Router Configuration and Prefix Lists. BGP Router Configuration contains including 3 parts: af-configuration , bgp-neighbors, and rpki-config (as shown in figure 1).

```

module: bgp
  +--rw bgp-router
  |   +--rw rpki-config
  |   +--rw af-configuration
  |   +--rw bgp-neighbors (98% config, 2% state)
  +--rw Prefix Lists

+--rw prefix-lists
+--rw prefix-list [prefix-list-name]
  +--rw prefix-list-name      string
  +--rw prefixes
  +--rw prefix [seq-nr]
    +--rw seq-nr              uint16
    +--rw prefix-filter
      +--rw (ip-address-group)?
      |   .....
      +--rw action              actions-enum
      +--rw statistics
      .....

```

Figure 1: draft-zhdankin-netmod-bgp-cfg structure

The [I-D.wang-i2rs-bgp-dm] is written with a status structure focus where manipulation of every concrete route through controlling policy and BGP attributes in different BGP address families. This does not contain the rpki-config. These drafts have ~5% overlap in status/configuration.

```

module: bgp
  +--rw bgp-protocols
  |   +--rw af-status
  |   |   +--rw bgp-local-rib
  |   |   +--rw bgp-neighbors (2% config, 98% state)
  |   |   |   +--rw policy-in
  |   |   |   +--rw policy-out
  |   |   |   +--rw peer-state
  |   |   |   +--rw bgp-rib-in
  |   |   |   +--rw-bgp-rib-out
  |   .....
  .....

module: ietf-pcim
  +--rw Policy-sets
  |   +--rw Policy-groups
  |   |   +--rw Policy-Rules
  |   .....
  .....

```

Figure 2 draft-i2rs-wang-bgp-dm-00

- o The focus is status-based based on AFI, but includes local-rib and BGP neighbor's policy (in and out), peer state, rib-in, and rib-out.
- o prefix list policy being covered inline ([I-D.zhdankin-netmod-bgp-cfg]) versus in the draft-hares-bnp-im-01 which uses the policy descriptions created by RFC3060, RFC3460, and other policy work. This methodology is being utilized by the OpenDaylight Group policy as well.
- o Redistribution being done inline ([I-D.zhdankin-netmod-bgp-cfg]) as a configuration. The draft-i2rs-wang-bgp-dm-00 did not configure af-configuration.
- o [I-D.zhdankin-netmod-bgp-cfg] claims to list the aspath path as well as the prefix configuration, but this section is missing in the draft. The example is the expression of as-path in draft-i2rs-wang-bgp-dm-00 is a actually string value of as-path which is one of attributes in BGP route rather not a Boolean value as used in [I-D.zhdankin-netmod-bgp-cfg].
- o The order of specifying the protocol elements is different in some cases due to the status versus configuration focus. For example, limiting maximum prefixes and maximum paths is done in a slightly different way. A second example is that community and cluster strings.

Below is an example of the AF-status structure found in draft-bgp-dm-00

```

module: bgp-protocol
  +--rw bgp-protocol
    +--rw bgp-ipv4-uni-instance
      |   +--rw bgp-local-rib
      |   +--rw bgp-peer-list
      |   +--rw bgp-rib-in
      |   +--rw bgp-rib-out
      .....
    +--rw bgp-evpn-instance
      |   +--rw bgp-local-rib
      |   +--rw bgp-peer-list
      |   +--rw bgp-rib-in
      |   +--rw bgp-rib-out

```

figure bgp-3

4.1. Overlap of configuration draft-zhdankin-netmod-bgp-01

The draft-zhdankin-netmod-bgp-01 and draft-wang-i2rs-bgp-dm-00 both contain at the protocol level:

```

module: bgp
  +--rw bgp router [bgp protocol]
    +--rw local-as-number?    unit32
    +--rw local-as-identifier inet:ip-address (zhdankin)
    +--rw router-id          inet:ip-address (wang)
    . . .
  figure bgp-3

```

The module contains at the peer level the association of a peer with an AS

[draft-zhdank-netmod-bgp-01]

```

+--rw bgp-neighbors* [AS-number]
  +--rw as-number
  +--rw (peer-address-type)?
  . . .
  +--rw prefix-list      prefix-list-ref
  +--default-action?    enumeration (permit/deny)

```

[[I-D.wang-i2rs-bgp-dm]]

```

+--rw bgp-peer-list* [bgp-peer-name]
  +--rw peer-session-address?
  . . .
  +--rw peer-name
  +--ro peer-type
  +--rw bgp-policy-in      policy-set-name
  +--rw bgp-policy-out    policy-set-name

```

figure bgp-5

4.2. Unique configuration for draft-zhdankin-netmod-bgp-01

[I-D.zhdankin-netmod-bgp-cfg] contains the unique configuration for RPKI, AF-configuration and BGP peers. A sample of the unique configuration for the AF-configuration is:

- o cost-communities
- o BGP-damping
- o route-aggregation - this is policy so we could easily add,

- o reflector clients
- o best external advertisement
- o aggregate timer (sending out aggregate times)
- o flags to compare router-id as part of bgp bestpath selection
- o MED-confederation
- o administrative distance (cisco)
- o packet forwarding over multiple paths
- o allowances for slow peers
- o BGP multi-path (ECMP peers)
- o external fail-over for peer
- o AS confederations
- o deterministic MEDs
- o Graceful-restart
- o BGP AS listener only
- o Logging of neighbor changes
- o Transport options for BGP
- o Removal of private AS

4.3. Unique configuration for draft-wang-bgp-dm-00

The following variables were included in [I-D.hares-i2rs-bgp-dm], but not in [I-D.zhdankin-netmod-bgp-cfg]:

- o protocol-status (ro) - needed for I2RS information
- o shutdown protocol - needed if I2RS is to shutdown bgp protocol,
- o AFI based local-rib
- o BGP-peer-status information - [I-D.zhdankin-netmod-bgp-cfg] has number of updates, but less status information in the draft.

The following pieces are needed by I2RS:

- o At instance level, `bgp-instance-name`, `bgp-instance-create` (to create bgp process), `bgp-instance-type` (to specify what type to create),
- o At AFI level in local rib, `bgp_route_create` (to add bgp route for i2rs) and `bgp_state_info` for status updates.
- o At peer level, there must be maximum prefixes per peer (received and transmit), high water/low water prefix counts, and average number of prefixes;
- o Versions for instance publishing,
- o Details on path attributes: ASpath strings, community and extend-community attribute, cluster lists, aggregation, atomic aggregator;
- o Mechanisms for logging/publishing information

5. Comparison of State needed versus BGP requirements

5.1. BGP-REQ 01

BGP-REQ01 requirement is: I2RS client/agent exchange SHOULD support the read, write and quick notification of status of the BGP peer operational state on each router within a given Autonomous System (AS). This operational status includes the quick notification of protocol events that proceed a destructive tear-down of BGP session

[I-D.zhdankin-netmod-bgp-cfg] contains the following status related to each peer's bgp operational state.

```
module: bgp
```

```

+--bgp-router
+ . . .
+--rw bgp-neighbors
  +--rw peer-address
  | . . .
  +--rw bgp-neighbor-state
    +--rw bgp-peer-admin-status enumeration
    +--rw in-lastupdatetime
  Figure bgp-6

```

Conclusion: [I-D.zhdankin-netmod-bgp-cfg] does not provide support required by I2RS.

[I-D.wang-i2rs-bgp-dm] contains the following status related to each peer's BGP operational state:

```

module: bgp

+--bgp-protocol
+ . . .
+rw bgp-ipv4-uni-instance (af-instance)
  |--rw bgp-instance-name
  | . . .
  |--rw bgp-local-rib
  | . . .
  |--rw bgp-peer-list [bgp-peer-name]
  | . . .
  |--rw peer-state-info
  |   |--rw peer-current-state?   peer-state
  |   |--rw peer-last-state?     peer-state
  |   |--ro peer-down-reason
  |   |--ro error code?   enumeration
  |   |--ro (sub-error-code-type)?
  |   |   |--: (head-error-sub-code)
  |   |   |--ro head-error-sub-value? enumeration
  |   |   |--: (open-error-sub-code)
  |   |   |--ro open-error-sub-value? enumeration
  |   |   |--: (update-error-sub-code)
  |   |   |--ro-update-error-sub-value? enumeration
  |   |   |--: (route-refresh-error-sub-code)
  |   |   |--ro-route-refresh-error-sub-value? enumeration

```

figure bgp-7

Conclusion: [I-D.wang-i2rs-bgp-dm] provides for this requirement.

5.2. BGP-REQ 02

BGP-REQ02 requirement is: "I2RS client SHOULD be able to push BGP routes with custom cost communities to specific I2RS agents on BGP routers for insertion in specific BGP Peer(s) to aid Traffic engineering of data paths. These routes SHOULD be tracked by the I2RS Agent as specific BGP routes with customer cost communities. These routes (will/will not) be installed via the RIB-Info."

Status:

[I-D.zhdankin-netmod-bgp-cfg] supports configuration related to address family control of these features, but it does not have a route level support. The AFI family configuration is shown in context below:

```

module: bgp
  +--rw bgp-router
  |   +--rw local-as-number?      uint32
  |   +--rw local-as-identifier?  inet:ip-address
  |   +--rw rpki-config
  |   |   .....
  |   +--rw af-configuration
  |   |   +--rw ipv4
  |   |   |   +--rw mdt
  |   |   |   |   . . .
  |   |   |   +--rw multicast
  |   |   |   |   +--rw bgp
  |   |   |   |   |   +--rw bgp-af-config
  |   |   |   |   |   |   . . .
  |   |   |   |   |   |   +--rw bestpath
  |   |   |   |   |   |   |   +--case as-path:
  |   |   |   |   |   |   |   |   ignore-as-path boolean
  |   |   |   |   |   |   |   +--case compare-routerid
  |   |   |   |   |   |   |   |   |   ignore-routerid boolean
  |   |   |   |   |   |   |   + case cost-community
  |   |   |   |   |   |   |   |   ignore-cost-community Boolean
  |   |   |   |   |   |   |   . . .
  |   |   |   |   +--rw unicast
  |   |   |   |   |   +--rw bgp
  |   |   |   |   |   |   +--rw bgp-af-config
  |   |   |   |   |   |   |   . . .
  |   |   |   |   |   |   |   +--rw bestpath
  |   |   |   |   |   |   |   |   . . .
  |   |   |   |   |   |   |   |   +--case cost-community
  |   |   |   |   |   |   |   |   |   ignore-cost-community Boolean
  |   |   |   |   +--rw unicast
  |   |   |   |   |   +--rw bgp
  |   |   |   |   |   |   +--rw bgp-af-config
  |   |   |   |   |   |   |   . . .
  |   |   |   |   |   |   |   +--rw bestpath
  |   |   |   |   |   |   |   |   . . .
  |   |   |   |   |   |   |   |   +--case cost-community
  |   |   |   |   |   |   |   |   |   ignore-cost-community Boolean

```

(replicated for appropriate AFIs)
Figure bgp-8

Conclusion: This [I-D.zhdankin-netmod-bgp-cfg] does not adequately support the I2RS BGP REQ02. .

[I-D.wang-i2rs-bgp-dm] does support per-route configuration tagging of route with customer community in local BGP rib, and per peer AdjRibIn and adjRibout

```

+--bgp-protocol
+ . . .
+rw bgp-ipv4-uni-instance (af-instance)
  +--rw bgp-instance-name
  | . . .
  +--rw afi
  +--rw safi
  +--rw bgp-local-rib
  | +--rw bgp-route-list* [ipv4-route ipv4-prefix-length]
  | | +--rw ipvr-route inet:ipv4-prefix
  | | +--rw ipv4-prefix-length uint8
  | | +--rw route-type? enumeration
  | | +--rw bgp-attribute-list*
  | | | +--rw bgp-origin?
  | | | . . .
  | | | +--rw bgp-extcommattr
  | | | | +--rw custom-community
  | | | | | +--rw valid boolean
  | | | | | +--rw insertion point uint8
  | | | | | +--rw community-id uint8
  | | | | | +--rw cost-id uint32
  | | | | +--rw useextcommsize uint16
  | | | | +--rw ulrefcount? uint16
  | | | | +--rw excmmattr-value string
  | | +--rw bgp-peer-list* (bgp-peer-name)
  | | . . .
  | | +--rw bgp-rib-in
  | | | +--rw bgp-rib-in-list* (ipv4 route
  | | | | ipv4-prefix-length route-distinguisher)
  | | | | +--rw ipv4-route inet_ipv4-prefix
  | | | | | +--rw bgp-attribute-list*
  | | | | | . . .
  | | | | | +--rw bgp-extcommattr
  | | | | | | +--rw custom-community
  | | | | | | +--rw valid boolean
  | | | | | | +--rw insertion point uint8
  | | | | | | +--rw community-id uint8
  | | | | | | +--rw cost-id uint32
  | | | | | | +--rw useextcommsize uint16
  | | | | | | +--rw ulrefcount? uint16
  | | | | | +--rw excmmattr-value string
  | | | | . . .
  | | | +--rw bgp-rib-out
  | | | | +--rw bgp-rib-out-list* (ipv4 route
  | | | | | ipv4-prefix-length route-distinguisher)
  | | | | | +--rw ipv4-route inet_ipv4-prefix
  | | | | | | +--rw bgp-attribute-list*
  | | | | | . . .

```

									+-rw bgp-extcommattr
									+-rw custom-community
									+-rw valid boolean
									+-rw insertion point uint8
									+-rw community-id uint8
									+-rw cost-id uint32
									+-rw useextcommsize uint16
									+-rw ulrefcount? uint16
									+-rw excmmattr-value string

Figure bgp-9

Conclusion: [I-D.wang-i2rs-bgp-dm] needed as well as [I-D.zhdankin-netmod-bgp-cfg].

5.3. BGP-REQ 03

BGP-REQ03 requirement is: "I2RS client SHOULD be able to track via read or notifications all Traffic engineering changes applied via I2RS agents to BGP route processes in all routers in a network."

Discussion on Requirement: Traffic engineering changes can include: ORFs (RFC5291, 5292), flows specifications (RFC5575, draft-ietf-), TE performance (draft-ietf-idr-te-pm-bgp-01), traffic-engineering state (draft-ietf-idr-te-lsp-distribution), and route target filtering. Additional input needs to be obtained from the i2rs WG on what constitutes traffic engineering.

Status:

[I-D.zhdankin-netmod-bgp-cfg] has the following potential configuration support:

- o on most af configurations: af-bgp_config container supports allowing the following features: add-path, best-external, aggregation timer, damping, propagating dmz-link-bw, and redistributing iBGP routes into IGP.
- o af rtfilters - AFI for rtfilters.

These features do not provide any of the traffic engineering input.

[I-D.wang-i2rs-bgp-dm]: has per route status tracking for the local-rib associated with each afi, and the virtual bgp AdjRibIn and BGP AdjRibOut for each peer. Each route has a route type that include route types for all valid NLRIs which includes: ipv4, ipv6, labeled ipv4, labeled ipv6, flows, link-state (ls) data, evpn, mvpn, vpls, l2vpn-singaling-nlri, rt-constraint, pw-route.

```

+--bgp-protocol
+ . . .
+rw bgp-ipv4-uni-instance (af-instance)
  +--rw bgp-instance-name
  | . . .
  +--rw afi
  +--rw safi
  +--rw bgp-local-rib
  | +--rw bgp-route-list* [ipv4-route ipv4-prefix-length]
  | | +--rw ipvr-route inet:ipv4-prefix
  | | +--rw ipv4-prefix-length uint8
  | | +--rw route-type? enumeration
  | | +--rw route-admin-distance uint16
  | | . . .
  | | +--rw bgp-attribute-list*
  | | | +--rw bgp-origin?

```

Figure bgp-10

What needs to be added: Once the I2RS WG specifies what traffic engineering flags to watch on the BGP routes at AFI local-rib level or the BGP-peer routes (AdjRibIn, AdjRibOut), then the [I-D.wang-i2rs-bgp-dm] can be augmented.

If the I2RS WG specifies configurations for traffic engineering at the AFI or BGP-peer level, these ephemeral status will either need to be added to draft-want-i2rs-bgp-dm-00 status or the peer.

5.4. BGP-REQ 04

BGP-REQ04 requirement is: "I2RS Agents SHOULD support identification of routers as BGP ASBRs, PE routers, and IBGP routers."

[I-D.zhdankin-netmod-bgp-cfg]: No status provides a summation of the BGP roles a BGP routing instance. The BGP-neighbor structure does not provide a role structure.

[I-D.wang-i2rs-bgp-dm]:

```

module: bgp-protocol
  +--rw bgp-protocol
    +--rw router-id?   inet:ip-address
    +--rw as-number?   uint32
    + . . .
    +--ro bgp-role?    enumeration
    +--rw af instance
      | +--rw bgp-local-rib
      | . . .
      | +--rw bgp-peer-list* [bgp-peer-name]
      | | +--rw bgp peer-session
      | | +--rw bgp-peer-name
      | | +--rw bgp-peer-type? enumeration

```

Figure bgp-11

The enumeration for `bgp-role` is `asbr`, `pe`, `ibgp`, `rr` where the role is a bit mask indicating that one or more peer has this role on the protocol instance.

The enumeration for `bgp-peer-type` is `asbr`, `ibgp`, `rr`, `rr-client`, `pe`, `ce`, `bgp-vendor-types`;

Conclusion: [I-D.wang-i2rs-bgp-dm] supports BGP-REQ 04

5.5. BGP-REQ 05

BGP-REQ05 is: I2RS client-agent SHOULD support writing traffic flow specifications to I2RS Agents that will install them in associated BGP ASBRs and the PE routers.

Status: BGP-REQ05 is the ability to read the roles within a `bgp` protocol instances at protocol level and at peer level, and to write routes with traffic flow specifications to AFI database and (optionally) `bgp-peer AdjRibOut`.

BGP-REQ 4 showed that the [I-D.wang-i2rs-bgp-dm] supports the identification of BGP ASBR and PE routers at a peer level. It also has a quick summary of the roles of BGP routers at the protocol level. [I-D.wang-i2rs-bgp-dm] specifies a `route-type` variable within each route in the AFI local-rib or the BGP Peers routes (`AdjRibIn`, `AdjRibOut`), and this `route-type` includes a enumeration variable for flows.

[I-D.wang-i2rs-bgp-dm]

```

module: bgp protocol
  +--bgp-protocol
  + . . .
  +rw bgp-ipv4-uni-instance (af-instance)
    +--rw bgp-instance-name
    | . . .
    +--rw afi
    +--rw safi
    +--rw bgp-local-rib
    | +--rw bgp-route-list* [ipv4-route ipv4-prefix-length]
    | | +--rw ipv4-route inet:ipv4-prefix
    | | +--rw ipv4-prefix-length uint8
    | | +--rw route-type? enumeration

```

Figure bgp-12

Conclusion: [I-D.wang-i2rs-bgp-dm] supports BGP-REQ-05.

5.6. BGP-REQ 06

BGP-REQ06 requirement is: "I2RS Client SHOULD be able to track flow specifications installed within a IBGP Cloud within an AS via reads of BGP Flow Specification information in I2RS Agent, or via notifications from I2RS agent."

Status: As section 3.5.5 on BGP-REQ04 shows [I-D.wang-i2rs-bgp-dm] supports the tracking of flow-specification routes associated with the local-rib for a AFI or a BGP Peer.

[I-D.wang-i2rs-bgp-dm]:


```

module: bgp protocol
  +--bgp-protocol
  + . . .
  +rw bgp-ipv4-uni-instance (af-instance)
    +--rw bgp-instance-name
      |
      | . . .
      +--rw bgp-local-rib
      |   +--rw bgp-rib-in-list* (ipv4 route ipv4-prefix-length
      |     |   +--rw ipv4-route inet:ipv4-prefix
      |     |   +--rw ipv4-prefix-length uint8
      |     |   +--rw route-type? enumeration
      |     | . . .
      |     +--rw bgp-peer-list* (bgp-peer-name)
      |       +--rw bgp-peer-session address
      |       +--rw bgp-peer-name
      |       +--rw bgp-peer-type
      |       +--rw bgp-rib-in
      |         +--rw bgp-rib-in-list* (ipv4 route ipv4-prefix-length
      |           |   +--rw ipv4-route inet:ipv4-prefix
      |           |   +--rw ipv4-prefix-length uint8
      |           |   +--rw route-type? enumeration

```

Figure bgp-13

5.7. BGP-REQ 07

BGP-REQ07 requirement is: I2RS client-agent exchange SHOULD support the I2RS client being able to prioritize and control BGP's announcement of flow specifications after status information reading BGP ASBR and PE router's capacity. BGP ASBRs and PE routers functions within a router MAY forward traffic flow specifications received from EBGp speakers to I2RS agents, so the I2RS Agent SHOULD be able to send these flow specifications from EBGp sources to a client in response to a read or notification.

Discussion: The I2RS WG needs to provide additional input on what status information is key to track for the BGP ASBR and PE router's capacity.

Status:

[I-D.zhdankin-netmod-bgp-cfg] has prefix-lists which can allow or deny prefixes based on the NLRI family. This feature allows the control of routes via the flow specification NLRI. However peer status does not provide an easy way to detect BGP ASBR or PE status, or the number of routes.

[I-D.wang-i2rs-bgp-dm] has the ability to specify flexible policy via policy-sets for inbound and outbound policy that can filter based on prefix or any match sequence within the route and peer. This draft also provides a data model that tracks track which peers are ASBR and PEs at the peer level via `bgp-peer-type` and at the protocol level via `bgp-role` (as described above). This draft also specifies administrative distance on route structures in the per AFI `bgp-local-rib` or in the peers routes per AFI.

```

module: bgp protocol
  +--bgp-protocol
  + . . .
  +rw bgp-ipv4-uni-instance (af-instance)
    +--rw bgp-instance-name
      |
      | . . .
      +--rw bgp-local-rib
        +--rw bgp-rib-in-list* (ipv4 route ipv4-prefix-length
          |
          | +--rw ipv4-route inet:ipv4-prefix
          | +--rw ipv4-prefix-length uint8
          | +--rw route-type? enumeration
          | +--rw route-admin-distance uint32
          | +--rw route-rpki-origin-validity
          | | +--rw rt-rpki-origin-valid Boolean
          | | +--rw rt-rpki-ref rpki-validity-ref

```

figure bgp-14

5.8. BGP-REQ 08

BGP-REQ08 requirement is: "I2RS Client SHOULD be able to read BGP route filter information from I2RS Agents associated with legacy BGP routers, and write filter information via the I2RS agent to be installed in BGP RR. The I2RS Agent SHOULD be able to install these routes in the BGP RR, and engage a BGP protocol action to push these routers to ASBR and PE routers."

Discussion: The router filter information is determined to be the prefix-filters or policy-filters associated with BGP routes found in the AFI based `bgp-local-rib` or peer's structures (`AdjRibIn`, `AdjRibout`).

Status:

[I-D.zhdankin-netmod-bgp-cfg] has prefix-lists which can allow or deny prefixes based on the NLRI. This yang model does not provide an easy way to detect peers as taking on the BGP RR, ABSR, and PE. (See section 3.3 and yang model for the prefix-list descriptions).

[I-D.wang-i2rs-bgp-dm] has the ability to specify flexible policy via policy-groups or policy sets for inbound and outbound policy that can filter based on prefix or any match sequence within the route and peer. This draft also provides a data model that tracks track which peers are ASBR, PEs, and RR at the peer level via `bgp-peer-type` and at the protocol level via `bgp-role`. (Please see draft-ietf-i2rs-hares-bnp-info-model and draft-ietf-hares-i2rs-bnp-dm for full description).

5.9. BGP-REQ 09

BGP-REQ09: I2RS client(s) SHOULD be able to request the I2RS agent to read BGP routes with all BGP parameters that influence BGP best path decision, and write appropriate changes to the BGP Routes to BGP and to the RIB-Info in order to manipulate BGP routes.

Discussion: Best-path is considered when policy evaluation is the same. The best path could be considered based on `origin-as`, `as-path`, `router-id`, `cost-community`, `igp-metric`, `med-confed`, `missing-as-med`, `rpki-origin` validity. This requirement needs to be refined to specify an initial set of BGP parameters that influence BGP best path decisions.

Status:

[I-D.zhdankin-netmod-bgp-cfg] configures the parameters that influence BGP bestpath decisions. However, this draft does not provide these parameters within each `bgp-route`.

[I-D.wang-i2rs-bgp-dm] provides the per route status on `origin-as`, `as-path`, `router-id`, `cost-community`, `igp-metric`, `MED`, and `rpki-origin` validity. This route status is on the AFI level `local-rib` and the per peer routes (`AdjRibIn`, `AdjRibOut`).

```

module: bgp protocol
  +--bgp-protocol
  + . . .
  +rw bgp-ipv4-uni-instance (af-instance)
    +--rw bgp-instance-name
      |
      | . . .
      +--rw bgp-local-rib
        |
        | +--rw bgp-rib-in-list* (ipv4 route ipv4-prefix-length
        | |
        | | +--rw ipv4-route inet:ipv4-prefix
        | | +--rw ipv4-prefix-length uint8
        | | +--rw route-type? enumeration
        | | +--rw route-admin-distance uint32
        | | +--rw route-rpki-origin-validity
        | | |
        | | | +--rw rt-rpki-origin-valid Boolean
        | | | +--rw rt-rpki-ref rpki-validity-ref

```

figure bgp-15

5.10. BGP-REQ 10

BGP-REQ10 requirement is: I2RS client SHOULD be able instruct the I2RS agent(s) to notify the I2RS client when the BGP processes on an associated routing system observe a route change to a specific set of IP Prefixes and associated prefixes. Route changes include: 1) prefixes being announced or withdrawn, 2) prefixes being suppressed due to flap damping, or 3) prefixes using an alternate best-path for a given IP Prefix. The I2RS agent should be able to notify the client via publish or subscribe mechanism.

Discussion: RFC5277 indicates that a netconf-filter looks for specific data value and data item. Therefore, the I2RS client must specify the whether the data is in the AFI based local-rib or the BGP (AdjRibIn, AdjRibOut) and the specific values. The specific values indicated by BGP-REQ-10 are prefixes with: announce flags, withdrawn flags, flap-dampened suppression flags, on-best-path-external or rejected due to best-path external. This comparison assumes the RFC5277 paths can be made to work for the ephemeral storage.

Sorting of these filters into critical or normal status requests is not considered in this comparison as adding this upon a non-existent definition of ephemeral services seems futile.

[I-D.zhdankin-netmod-bgp-cfg] configures the parameters that influence BGP bestpath decisions or flap damping. However, this draft does not provide these parameters within each bgp-route.

[I-D.wang-i2rs-bgp-dm]:

```

+--rw ipv4-route                inet:ipv4-prefix
+--rw ipv4-prefix-length        uint8
+--rw bgp-route-type?           enumeration
+--rw bgp-attribute-list
    ...
+--ro bgp-rt-state-in
  +--ro rib-current-state?      rib-state-def
  +--ro rib-last-state?         rib-state-def
  +--ro rib-rejected-reason?    enumeration
  +--ro not-preferred-reason?   enumeration
. . .

typedef rib-state-def {
  type enumeration {
    enum "active";
    enum "in-active";
    enum "primary";
    enum "backup";
    enum "suppressed (flap dampened)";
    enum "suppressed non-flap dampen";
    enum "active on alternate best path"
  }
}

Leaf not-preferred-reason {
  Type enumeration {
    enum "peer-address";
    enum "router-id";
    enum "cluster-list-length";
    enum "igp-metric";
    enum "peer-type";
    enum "origin";
    enum "weight-or-preferred-value";
    enum "local-preference";
    enum "route-type";
    enum "as-path-length";
    enum "med";
    enum "flap-dampened route"; [new]
    enum "not-this-path-prefix-uses-alternate-best-path"; [new]
    enum "overlapping-route-marked-to-remove"; [new] (BGP-REQ17)
  }
}

```

Figure bgp 16

Conclusion: [I-D.wang-i2rs-bgp-dm] status is needed for this BGP-REQ10.

5.11. BGP-REQ 11

BGP-REQ11 requirement is the "I2RS client SHOULD be able to read BGP route information from BGP routers on routes in received but rejected from ADJ-RIB-IN due to policy, on routes installed in ADJ-RIB-IN, but not selected as best path, and on route not sent to IBGP peers (due to non-selection)."

Discussion: As discussed in BGP-REQ10, RFC5277 indicates that a netconf-filter looks for specific data value and data item. Therefore, the I2RS client must specify the whether the data is in the AFI based local-rib, or the BGP (AdjRibIn, AdjRibOut) and the specific values

Status:

[I-D.zhdankin-netmod-bgp-cfg] configures policy that indicates what routes are filtered, but it does not provide the status parameter on each BGP route.

[I-D.wang-i2rs-bgp-dm]: Shows that the status can be tracked in the AFI based bgp local-rib and the per AFI per peer AdjRibIn and AdjRib out.

Conclusion: [I-D.wang-i2rs-bgp-dm] status is needed for BGP-REQ10.

5.12. BGP-REQ 12

BGP-REQ12 requirement is: the "I2RS client SHOULD be able to request the I2RS agent to read installed BGP Policies."

Discussion: BGP policies can be inbound filters, ACLs, or route maps. Three yang drafts take different approaches to the filters: draft-bogdanovic-netmod-acl-model-02, [I-D.zhdankin-netmod-bgp-cfg], and draft-hares-2rs-bnp-dm-01.

The draft-bogdanovic-netmod-acl-model-02 takes the approach of extending the firewall ACLs, and suggests that proprietary methods be used to extend this to the ranges needed for BGP policy. The index to the ACLS is the rule-name. For a single prefix accept/deny the generic ACL policy may be sufficient. Prefix ranges or the ability to set custom cost communities or other extended communities must use a proprietary vendor's model.

The [I-D.zhdankin-netmod-bgp-cfg] (10/1/2014) suggests prefix-list matches that should provide prefix matches an index of route ID number (uint16). The prefix matches can be either ip-address, prefix, host address, or prefix-range. The only actions taken are

accept or deny the prefix. The usage statistics contains only the "hit count" for the usage of the prefix mask.

The [I-D.wang-i2rs-bgp-dm] (10/23/2014) provides a policy link to the Basic Network Policy (draft-hares-i2rs-bnp-info-model/draft-hares-i2rs-bnp-dm-01) which provides ordered list of policy rules that can provide sequences of match, actions (accept/deny, set variable, and modify). A group of these policy rules is called a policy group which can be named.

This model uses the policy definitions concept is also found in the Policy Core Information Model (PCIM) (RFC3060) and the Quality of Service (QoS) Policy Information Model (QPIM)(RFC3644) and policy based routing. ACL-based policy (draft-bogdanovic-netmod-acl-model-02), and prefix-list policy (accept/deny) can be one of the policy rule extensions. The [I-D.zhdankin-netmod-bgp-cfg] can provide a second policy rule extension.

The section below provides the specific modeling parameters for each draft.

draft-bogdanovic-netmod-acl-model-02

```

+--rw access-list-entry* [rule-name]
  +--rw rule-name          string
  +--rw matches
    +--rw (ace-type)?
      +--:(ace-ip)
        +--rw source-port-range
          +--rw lower-port    inet:port-number
          +--rw upper-port?   inet:port-number
        +--rw destination-port-range
          +--rw lower-port    inet:port-number
          +--rw upper-port?   inet:port-number
          +--rw dscp?         inet:dscp
          +--rw ip-protocol?  uint8
          +--rw (ace-ip-version)?
            +--:(ace-ipv4)
              +--rw destination-ipv4-address?
                inet:ipv4-prefix
              +--rw source-ipv4-address?
                inet:ipv4-prefix
            +--:(ace-ipv6)
              +--rw destination-ipv6-address?
                inet:ipv6-prefix
              +--rw source-ipv6-address?
                inet:ipv6-prefix
              +--rw flow-label?  inet:ipv6-flow-label

```

```

| | +--:(ace-eth)
| |   +--rw destination-mac-address?
| |       yang:mac-address
| |   +--rw destination-mac-address-mask?
| |       yang:mac-address
| |   +--rw source-mac-address?
| |       yang:mac-address
| |   +--rw source-mac-address-mask?
| |       yang:mac-address
| +--rw input-interface?                string
| +--rw absolute
|   +--rw start?          yang:date-and-time
|   +--rw end?           yang:date-and-time
|   +--rw active?        boolean
|   +--rw actions
| +--rw (packet-handling)?
|   +--:(deny)
|   | +--rw deny?        empty
|   +--:(permit)
|   | +--rw permit?     empty
+--ro ace-oper-data
+--ro match-counter?    ietf:counter64

```

Figure bgp-17

[I-D.zhdankin-netmod-bgp-cfg] (10/1/2014)

```

+--rw prefix-lists
+--rw prefix-list [prefix-list-name]
+--rw prefix-list-name    string
+--rw prefixes
+--rw prefix [seq-nr]
+--rw seq-nr uint16
+--rw prefix-filter
+--rw (ip-address-group)?
| (cases
+--rw action actions-enum
+--rw statistics
..

```

Figure bgp-18

[I-D.wang-i2rs-bgp-dm]


```
module: bgp_protocol
+--rw bgp-protocol
  + af config
    +--rw bgp-policy-in  policy-group-ref
    +--rw bgp-policy-out  policy-group-ref
```

Figure bgp-19

draft-bnp-i2rs-bnp-dm-00

(To be Completed after bnp drafts)

Figure bgp-20

5.13. BGP-REQ 13

BGP-REQ13 requirement is: the "I2RS client SHOULD be able to instruct the I2RS Agent to write BGP Policies into the running BGP protocols and into the BGP configurations."

Discussion: BGP-REQ 13 indicates that the policy changes supported by BGP-REQ 12 must be able to operate in the running configuration. The I2RS and netmod groups are discussing the definition of ephemeral. Two definitions are being discussed - patch and pull-up configuration as described below:

running = config + ephemeral patches [patch]

running = config (copy) + ephemeral additions (pull-up)

Either definition implies that the I2RS changes can alter the policy based on the bgp configuration and I2rs model.

The writing of changes from I2RS to the configuration was specifically ignored. Writing of specific configuration options from the ephemeral store to the running configuration can initially be done by the I2RS client writing via the configuration interface to the datastore.

Data models needed: The Policy data models for filter or filter and action are the same as in BGP-REQ13.

5.14. BGP-REQ 14

BGP-REQ14 requirement is: the "I2RS client-agent SHOULD be able to read BGP statistics associated with Peer, and to receive notifications when certain statistics have exceeded limits. An example of one of these protocol statistics is the max-prefix limit."

Discussion: BGP-REQ01 examines the peer connectivity state. BGP-REQ14 examines BGP peer statistics. [I-D.zhdankin-netmod-bgp-cfg] provides statistics on the number of updates received or sent, but it does not have statistics on the max-prefix exceeded. It does provide a limit for maximum-prefix per peer.

[I-D.wang-i2rs-bgp-dm] has statistics on the number of updates received or sent, number of routes received or sent plus maximum prefix high-water and low-water. This draft also has the limits copied into the status fields for easy reading.

[I-D.zhdankin-netmod-bgp-cfg] (10/1/2014)

```

module: bgp
+ ....
+--rw bgp-neighbors
|   +--rw bgp-neighbor [as-number]
|   |   +--rw as-number                               uint32
|   |   +--rw (peer-address-type)?
|   |   |   ....
|   |   +--rw prefix-list?                           prefix-list-ref
|   |   +--rw default-action?                        actions-enum
|   |   +--rw af-specific-config
|   |   +--ro bgp-neighbor-state
|   |   |   +--ro bgp-peer-admin-status;
|   |   +--ro bgp-neighbor-statistics
|   |   |   +--ro nr-in-updates
|   |   |   +--rw nr-out-updates

```

[I-D.wang-i2rs-bgp-dm]

```

+--rw bgp-peer-list* [bgp-peer-name]
  +--rw peer-session-address
    |   +--rw local-ipv4-address?    inet:ipv4-prefix
    |   +--rw remote-ipv4-address?  inet:ipv4-prefix
  +--rw bgp-peer-name                string
  +--ro bgp-peer-type?                enumeration
  +--ro bgp-peer-create?              enumeration
  +--rw bgp-policy-in                 policy-group-ref
  +--rw bgp-policy-out                 policy-group-ref
  +--rw peer-state-info
    |   +--ro peer-current-state?    peer-stat
    |   +--ro peer-last-state?       peer-state
    |   +--ro peer-down-reason
    |     |   +--ro error-code?      enumeration
    |     |   +--ro sub-error-code
    |     |     |   ...
  +--ro peer-transmit-update-cnt?     uint64
  +--ro peer-recived-update-cnt       uint64
  +--ro peer-received-route-cnt?      uint64
  +--ro peer-send-route-cnt?          uint64
  +--rw max-prefix-rcv-limit?         uint64
  +--rw max-prefix-xmt-limit?         uint64
  +--ro peer-prefix-high?             uint64
  +--ro peer-prefix-low?              uint64

```

Conclusion: Full support of BGP-REQ 14 requires the [I-D.wang-i2rs-bgp-dm] draft.

5.15. BGP-REQ 15

BGP-REQ15 requirement is: the "I2RS client via the I2RS agent MUST have the ability to read the loc-RIB-In BGP table that gets all the routes that the CE has provided to a PE router."

Discussion: The [I-D.zhdankin-netmod-bgp-cfg] provides no indication of a local-rib or a local-RIB-in associated with a BGP peer. The [I-D.wang-i2rs-bgp-dm] provides a local-rib per AFI, and a local-RIB-IN (AdjRIBIn and AdjRIBOut) associated with each peer. The route table format is presented in figure bgp-9.

Conclusion: [I-D.wang-i2rs-bgp-dm] is necessary for this requirement.

5.16. BGP-REQ 16

BGP-REQ16 requirement is: the "I2RS client via the I2RS agent MUST have the ability to install destination based routes in the local RIB of the PE devices. This must include the ability to supply the destination prefix (NLRI), a table identifier, a route preference, a

route metric, a next-hop tunnel through which traffic would be carried."

Discussion: If this refers to the I2RS LOC-RIB, then both drafts have policy which could redistribute I2RS routes.

If BGP-REQ 16 refers to a BGP local-rib per AFI and the bgp-peer based bgp-rib-in (AdjRibIn) and bgp-rib-out (AdjRibOut). As this previous discussion indicates, the [I-D.zhdankin-netmod-bgp-cfg] does not have a local rib-in.

The document [I-D.wang-i2rs-bgp-dm] describes a per AFI bgp local-rib and the per peer bgp-rib-in (AdjRIBIn) and bgp-rib-out (AdjRibOut). The routes in these RIBs fall under a table identifier structure and have a destination prefix (NLRI), route administrative preference, route local-reference, and a next-hop. However, [I-D.wang-i2rs-bgp-dm] does not have a tunnel based definition for the next-hop. This would need to be added.

Conclusion: Additions to [I-D.wang-i2rs-bgp-dm] may need to be made to fulfill this requirement.

5.17. BGP-REQ 17

BGP-REQ17 requirement is: the "I2RS client via the I2RS agent SHOULD have the ability to read the loc-RIB-in BGP table to discover overlapping routes, and determine which may be safely marked for removal."

Discussion: As discussed in BGP-REQ15 and BGP-REQ16, draft [I-D.zhdankin-netmod-bgp-cfg] does not have a local-RIB-In BGP table. [I-D.wang-i2rs-bgp-dm] has a BGP local-rib per AFI and a per BGP Peer bgp-rib-in. As described in BGP REQ10, this each route may set a "remove overlapping route" status flag.

Conclusion: [I-D.wang-i2rs-bgp-dm] supports BGP-REQ 17.

5.18. BGP-REQ 18

BGP-REQ18 requirement is the "I2RS client via the I2RS Agent SHOULD have the ability to modify filtering rules and initiate a re-computation of the local BGP table through those policies to cause specific routes to be marked for removal at the outbound eBGP edge."

Discussion: This feature requires that I2RS should be able to do a re-computation of policies. This re-computation of policies is part of a soft-reconfig which [I-D.zhdankin-netmod-bgp-cfg] allows by

configuration. However, the I2RS client will need a parameter to be set to do a reconfigure.

Neither [I-D.zhdankin-netmod-bgp-cfg] or [I-D.wang-i2rs-bgp-dm] have this feature.

Conclusion: This feature needs to be added to final model

6. IANA Considerations

This draft includes no request to IANA.

7. Security Considerations

TBD

8. 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-bgp-dm]

Wang, L., Hares, S., and S. Zhuang, "An I2RS BGP Data Model", draft-hares-i2rs-bgp-dm-00 (work in progress), October 2014.

[I-D.hares-i2rs-info-model-service-topo]

Hares, S., Wu, W., and X. Guan, "An Information model for service topology", draft-hares-i2rs-info-model-service-topo-01 (work in progress), July 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.wang-i2rs-bgp-dm]
Wang, L., Hares, S., and S. Zhuang, "An I2RS BGP Data Model", draft-wang-i2rs-bgp-dm-00 (work in progress), October 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.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [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.
- [RFC5511] Farrel, A., "Routing Backus-Naur Form (RBNF): A Syntax Used to Form Encoding Rules in Various Routing Protocol Specifications", RFC 5511, April 2009.

Author's Address

Susan Hares
Huawei
7453 Hickory Hill
Saline, MI 48176
USA

Email: shares@ndzh.com