

i2rs
Internet-Draft
Intended status: Informational
Expires: September 16, 2016

S. Hares
Huawei
M. Chen
Huawei Technologies
March 15, 2016

Summary of I2RS Use Case Requirements
draft-ietf-i2rs-usecase-reqs-summary-02

Abstract

The I2RS Working Group (WG) has described a set of use cases that the I2RS systems could fulfil. This document summarizes these use cases. It is designed to provide requirements that will aid the design of the I2RS architecture, Information Models, Data Models, Security, and protocols.

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 16, 2016.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in [Section 4.e](#) of

the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

- [1. Introduction](#) [2](#)
- [2. Protocol Independent Use Case Requirements](#) [4](#)
- [3. BGP Use Case Requirements](#) [6](#)
- [4. IGP Use Cases](#) [8](#)
- [5. CCNE Use Cases](#) [10](#)
- [6. Topology Related Use Cases](#) [11](#)
 - [6.1. Virtual Connection Use Case Requirements](#) [11](#)
 - [6.2. Virtual Network Use Case Requirements](#) [11](#)
 - [6.3. Topology Use Case](#) [13](#)
 - [6.4. Virtual Topology Data Model](#) [17](#)
 - [6.5. Virtual Topology IP Data Model](#) [19](#)
 - [6.6. Virtual Topology Network Element](#) [19](#)
- [7. Requirements from SFC Use Cases](#) [20](#)
- [8. Requirements from Traffic Steering Use Cases](#) [22](#)
- [9. Requirements from MPLS TE Networks Use Cases](#) [22](#)
- [10. Requirements from MPLS LDP Networks Use Cases](#) [24](#)
- [11. Requirements from Mobile Backhaul Ues Cases](#) [25](#)
- [12. Requirements from Large Data Flows are](#) [27](#)
- [13. Large Data Collection Systems](#) [28](#)
- [14. CDNI](#) [30](#)
- [15. IANA Considerations](#) [31](#)
- [16. Security Considerations](#) [31](#)
- [17. References](#) [31](#)
 - [17.1. Normative References](#) [31](#)
 - [17.2. Informative References](#) [31](#)
- Authors' Addresses [34](#)

1. Introduction

The Architecture for the Interface to the Routing System [[I-D.ietf-i2rs-architecture](#)] allows for a mechanism where the distributed control plane can be augmented by an outside control plane through an open, accessible interface. This document summarizes the use case requirements for the I2RS client-I2RS Agent exchange found in the following documents:

- o Protocol Independent described in [[I-D.white-i2rs-use-case](#)]
- o BGP described in [[I-D.keyupate-i2rs-bgp-usecases](#)]
- o IGP protocols as described in [[draft-ietf-wu-i2rs-igp-usecases](#)]

- o Control of Forwarding Path by Central Control Network Element (CCNE) [[I-D.ji-i2rs-usecases-ccne-service](#)]
- o Virtual Connections and Virtual Networks described in [[I-D.hares-i2rs-use-case-vn-vc](#)]
- o Topology use cases [[I-D.amante-i2rs-topology-use-cases](#)]
- o Topology requirements [[I-D.medved-i2rs-topology-requirements](#)]
- o Service chaining described in [[I-D.bitar-i2rs-service-chaining](#)]
- o Traffic Steering described in [[I-D.chen-i2rs-ts-use-case](#)]
- o MPLS TE Networks described in [[I-D.huang-i2rs-mpls-te-usecases](#)]
- o MPLS LDP Networks described in [[I-D.chen-i2rs-mpls-ldp-usecases](#)]
- o Mobile BackHaul Use cases described in [[I-D.zhang-i2rs-mbb-usecases](#)]
- o Large Flows use case described in [[I-D.krishnan-i2rs-large-flow-use-case](#)]
- o Large Data Collection Systems Use cases described in [[I-D.swhyte-i2rs-data-collection-system](#)]
- o CDNI requesting routing [[I-D.shin-i2rs-usecases-cdni-request-routing](#)]

Each group of use cases is presented in its own document. Each use case is labeled with an identifier TTT-REQ-nn where TTT represents the type of use case. The abbreviations for TTT are:

- o PI - Protocol Independent
- o BGP - BGP
- o IGP - IGP protocols
- o CCNE - CCNE control of forwarding path
- o VCoD - Virtual Connections on Demand
- o VNoD - Virtual Networks on Demand
- o Topo - Topology Information

- o VT-TMD - Virtual Topology: Topology Data Model
- o VT-TDM-IP - Virtual Topology: Topology Data Mode for IP/MPLS
- o SFC - Service Chaining requirements
- o TS - Traffic Steering
- o MPLS-LDP - MLPS Topologies supported by LDP
- o MPLS-TE - MPLS-TE topologies
- o MBH - Mobile Back-Haul
- o L-Flow - Large Flows
- o L-Data - Large Data Collection
- o CDNI - CDNI networks

Each use case is also augmented with a notation signifying whether it is in or out of scope with regard to the current I2RS charter:

- o IC: In charter
- o OC: Out of charter
- o NA: not applicable to I2RS protocol, agent, client or models. Usually related to specific client-side app requirements.
- o ??: indicates this item needs additional classification aid from the WG.

In some cases a specific draft may be out of charter, but (sub)components of it's requirement set may be in charter. In charter. As such, (IC|OC|NA) designations may appear at the draft level, at the requirement level, or at the sub requirement level. In instances where designations do not appear at more specific level, the designation at the parent level should be considered to be inherited.

2. Protocol Independent Use Case Requirements

This is a summary of the I2RS requirements found in the Protocol Independent Use Cases described in: [[I-D.white-i2rs-use-case](#)] (IC):

- o PI-REQ01 (IC): The ability to monitor the available routes installed in the RIB of each forwarding device, including near

real time notification of route installation and removal. This information must include the destination prefix (NLRI), a table identifier (if the forwarding device has multiple forwarding instances), the metric of the installed route, and an identifier indicating the installing process.

- o PI-REQ02 (IC): The ability to install source and destination based routes in the local RIB of each forwarding device. This must include the ability to supply the destination prefix (NLRI), the source prefix (NLRI), a table identifier (if the forwarding device has multiple forwarding instances), a route preference, a route metric, a next hop, an outbound interface, and a route process identifier.
- o PI-REQ03 (IC): The ability to install a route to a null destination, effectively filtering traffic to this destination.
- o PI-REQ04(?): The ability to interact with various policies configured on the forwarding devices, in order to inform the policies implemented by the dynamic routing processes. This interaction should be through existing configuration mechanisms, such as NETCONF, and should be recorded in the configuration of the local device so operators are aware of the full policy implemented in the network from the running configuration.
- o PI-REQ05 (OC): The ability to interact with traffic flow and other network traffic level measurement protocols and systems, in order to determine path performance, top talkers, and other information required to make an informed path decision based on locally configured policy.
- o PI-REQ06 (IC): The ability to install destination based routes in the local RIB of each forwarding device. This must include the ability to supply the destination prefix (NLRI), a table identifier (if the forwarding device has multiple forwarding instances), a route preference, a route metric, a next hop, an outbound interface, and a route process identifier.
- o PI-REQ07 (IC): The ability to read the local RIB of each forwarding device, including the destination prefix (NLRI), a table identifier (if the forwarding device has multiple forwarding instances), the metric of each installed route, a route preference, and an identifier indicating the installing process.
- o PI-REQ08 (IC): The ability to read the tables of other local protocol processes running on the device. This reading action should be supported through an import/export interface which can present the information in a consistent manner across all protocol

implementations, rather than using a protocol specific model for each type of available process.

- o PI-REQ09 (OC for some protocols): The ability to inject information directly into the local tables of other protocol processes running on the forwarding device. This injection should be supported through an import/export interface which can inject routing information in a consistent manner across all protocol implementations, rather than using a protocol specific model for each type of available process.
- o PI-REQ10 (OC): The ability to interact with policies and configurations on the forwarding devices using time based processing, either through timed auto-rollback or some other mechanism. This interaction should be through existing configuration mechanisms, such as NETCONF, and should be recorded in the configuration of the local device so operators are aware of the full policy implemented in the network from the running configuration.
- o PI-REQ-11 (IC) The ability to update the Local RIB with varying levels of checks on the route. These checks can be simply minimal reception checks (TLVs align correctly), all non-referential checks (do not do leafref, MUST, instance identifiers), do referential checks.

3. BGP Use Case Requirements

This is a summary of the requirements listed in [\[I-D.keyupate-i2rs-bgp-usecases\]](#) are (IC):

- o BGP-REQ01 (IC): 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
- o BGP-REQ02 (IC): 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) installed via the RIB-Info.
- o BGP-REQ03 (IC): I2RS client SHOULD be able to track via read/notifications all Traffic engineering changes applied via I2RS agents to BGP route processes in all routers in a network.

- o BGP-REQ04 (IC): I2RS Agents SHOULD support identification of routers as BGP ASBRs, PE routers, and IBGP routers.
- o BGP-REQ05 (IC): I2RS client-agent SHOULD support writing traffic flow specifications to I2RS Agents that will install them in associated BGP ASBRs and the PE routers.
- o BGP-REQ06 (IC): 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
- o BGP-REQ07 (IC): 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.
- o BGP-REQ08 (IC): 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.
- o BGP-REQ09 (IC): 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
- o BGP-REQ10 (IC): 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.
- o BGP-REQ11 (IC): 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).

- o BGP-REQ12 (IC): I2RS client SHOULD be able to request the I2RS agent to read installed BGP Policies.
- o BGP-REQ13 (IC): I2RS client SHOULD be able to instruct the I2RS Agent to write BGP Policies into the running BGP protocols and into the BGP configurations.
- o BGP-REQ14 (IC): 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.
- o BGP-REQ15 (IC): 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.
- o BGP-REQ16 (IC): 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
- o BGP-REQ17 (IC): The I2RS client via the I2RS agent SHOULD have the the ability to read the loc-RIB-in BGP table to discover overlapping routes, and determine which may be safely marked for removal.
- o BGP-REQ18 (IC): 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.

4. IGP Use Cases

This is a summary of the requirements listed in (ietf-draft-wu-ir2s-igp-usecases-00.txt) (OC):

- o IGP-REQ-01 (OC): I2RS Client/Agent SHOULD Be able to read/write the the unique IGP identification for router within an AS (router-id, system-id, or others). I2RS agents may notify the I2RS client of the detection of another router with the same unique ID.
- o IGP-REQ-02 (OC): I2RS Client SHOULD BE able to aid in IGP table reduction by actively monitoring IGP tables and by allowing changes to the IGP configuration in order to partition the IGPS and place ABRs and ASBRs. The I2RS Client/Agent exchange must allow for a rapid cycle of querying of IGP topology information

and downloading of a new protocol configuration or updating of IGP nexthops in RIBs and FIBs to rapidly switch to new temporary IGP topologies. These alternate topologies may be calculated by a application attached to the i2rs client and updated to the i2rs agent, or determined at the i2rs agent.

- o IGP-REQ-03 (OC): I2RS protocol and models should support Loop-Free Alternative (LFAs) [[RFC5286](#)] deployments in in pure IP and MPLS/LDP networks to provide single-point-failure protection for unicast traffic. This includes the configuration, monitoring of LFA changes, and letting off-line pre-computed paths for LFA backup of all links and prefixes in the network and calculating the protection coverage and recognizing optimization to be downloaded to appropriate devices via the I2RS interface (Client-Agent). Again, it is important to have deployment of changes followed by real-time feedback.
- o IGP-REQ-04 (OC): The I2RS programmatic interface SHOULD allow the balancing of both ECMP traffic flows and end-to-end traffic flows in the IGP. The I2RS SHOULD support monitoring of the dynamic traffic flow in the network, and the query of the maximum capacity of the network. This include the I2RS client's transmission to the I2RS agent of updated configuration after an off-line optimization to either spread traffic (across ECMP pathways) or aggregation of traffic onto a single path so the rest of the devices may power off saving power (and money).
- o IGP-REQ-05 (OC): The I2RS interface (protocol and data models) SHOULD use the subscription mechanism to filter the topology changes to interested events and use the publish mechanism to control the pace these events are notified. This filtering should protect the I2RS Client or even applications who depend on topology data from being drowned by massive original events or duplicate events from different sources
- o IGP-REQ-06 (OC): Since IGP protocol is essential to the whole network, the I2RS Clients SHOULD monitor about the protocol's running status before forwarding is impacted. Performance data can be collected through collecting static configuration and observing dynamic status. Static data includes the number of instances, interfaces, nodes in the network and etc. Dynamic data includes adjacency status, the number of entries in link-state database and in the routing table, the calculation status, the overload status, the graceful switch-over status, and others
- o IGP-REQ-07 (OC): The I2RS interface (protocol and IMs) should support a mechanism where the I2RS Clients can subscribe to the I2RS Agent's notification of critical node IGP events. For

example, link-state database or routing table is under the status of overflow or the overflow status is released, the calculation continues for a long time, the system is under graceful reboot.

- o IGP-REQ-08 (OC): The I2RS interface (protocol and IMs) should support the reporting of IGP statistic such as dropped packet statistics. These statistics will aid detection of network failures or security attacks.

5. CCNE Use Cases

The use cases in I2RS Use Cases for Control of the Forwarding Path by a Central Control Network Element (CCNE)

[[I-D.ji-i2rs-usecases-ccne-service](#)] indicate the following requirements for I2RS (OC):

- o CCNE-REQ-01 (IC): I2RS interface should support I2RS client running on a CCNE to be able to pull information from both the BGP RR and the PCE. This information can include: BGP topology information, BGP routes, BGP statistics, BGP Peer topologies, PCE topology information, and PCE state information. The I2RS Client's request for reading of the RR and PCE topology information needs to have timely and rapid response from the I2RS Agent.
- o CCNE-REQ-02 (IC for some constraints): I2RS client should be able to set resource constraints at the I2RS Agent, and receive status information on the setting of resource constraints.
- o CCNE-REQ-03 (IC for some constraints): I2RS interface should be able to set service goal value to CCNE.
- o CCNE-REQ-04 (OC): I2RS client should be able support information models that allow re-optimization traffic model at at CCNE .
- o CCNE-REQ-05 (IC): I2RS client should be able to receive notification at the CCNE, and be able to send status to the I2RS agent.
- o CCNE-REQ-06 (NA): I2RS client should work in parallel with traditional network management or OAM protocols sent to the general NE.
- o CCNE-REQ-07 (NA): I2RS clients should be able to to be light weight enough to be able to support running on a variety of devices (routers, centralized servers, or devices doing both).

6. Topology Related Use Cases

This section describes Topology or Virtual Topology related requirements the I2RS interface (protocol and information model (IM)) included in the following types of use cases:

- o Virtual Connections on Demand: VCoD-REQ
- o Virtual Networks on Demand: VNoD-REQ
- o Virtual Topology Information Topo-REQ
- o Virtual Topology Data Model: VT-TDM-REQ
- o Virtual Topology IP Data Model: VT-TDMIP-REQ
- o Virtual Topology Network Element: VT-NE-REQ (TMF-GEN-1)

6.1. Virtual Connection Use Case Requirements

- o VCoD-REQ01 (OC): I2RS Agents SHOULD provide the ability to read the virtual network topology database for the technology supported. For optical, these are the optical connections and what node they connect to, and the topologies created. For MPLS, this is virtual circuit available, what nodes they connect to, and the network topologies created. For IP technologies, this could include the GRE tunnels, what interface it connects to, and the topologies created. For Ethernet circuits this should involve circuit type (e.g, point-to-point (p2p) or point-to-multipoint (p2mp)) and what nodes it can reach, and the topologies created.
- o VCoD-REQ02 (OC): I2RS Agent SHOULD provide the ability to influence the configuration of a virtual circuit in a node.
- o VCoD-REQ03 (OC): I2RS Agent SHOULD provide monitor and provide statistics on the virtual connection to the I2RS client via a Read request or status Notification. The I2RS client can then determine if the connection falls below a quality level the application has requested. If the I2RS client does determine the circuit is below the required quality, it could create another circuit. The I2RS may choose to create the second virtual circuit, transfer flows, and then break the first circuit.


6.2. Virtual Network Use Case Requirements

The requirements for the Virtual Networks on Demand (VCoD) are:

- o VT-VN-REQ01 (IC): I2RS Agents SHOULD provide the ability to read the virtual network topology database for the technology supported to determine nodes and connections. For optical, these are the optical connections and what node they connect to, and the topologies created. For MPLS, this is virtual circuit available, what nodes they connect to, and the network topologies created. For IP technologies, this could include the GRE tunnels, what interface it connects to, and the topologies created. For Ethernet circuits this should involve circuit type (e.g, point-to-point (p2p) or point-to-multipoint (p2mp)) and what nodes it can reach, and the topologies created.
- o VNoD-REQ02 (IC): I2RS Agent SHOULD provide the ability to influence the configuration of a virtual circuit in a node.
- o VNoD-REQ03 (IC): I2RS Agent SHOULD provide monitor and provide statistics on the virtual connection to the I2RS client via a Read request or status Notification. The I2RS client can then determine if the connection falls below a quality level the application has requested. If the I2RS client does determine the circuit is below the required quality, it could create another circuit. The I2RS may choose to create the second virtual circuit, transfer flows, and then break the first circuit.
- o VNoD-REQ04 (IC): I2RS Agent SHOULD provide the ability to influence the configuration of a virtual network in a node.
- o VNoD-REQ05 (OC): I2RS Agent SHOULD provide the ability to report statistics on the network nodes and end-to-end traffic flows via read of status data or via notifications of status.
- o VNoD-REQ06 (IC): The I2RS protocol and RIB Informational Model (IM) must support logical tunnels of type MPLS as well as IP, GRE, VxLAN and GRE. Large Carrier networks utilize MPLS in a variety of forms (LDP, static MPLS TE, or dynamic TE LSPS created by RSVP-TE or CR-LDP).
- o VNoD-REQ07 (IC): I2RS SHOULD support Informational Models and features to allow MPLS technologies to create Hub-spoke topology and service routing in networks in Carriers, Enterprise, and Data Centers.
- o VNoD-REQ08 (IC): I2RS protocols, Information Models, and Data Models must be able to support Carriers using these MPLS technologies to support networks for Mobile BackHaul, on-demand MPLS overlays, and on-demand video conferencing networkings.

6.3. Topology Use Case

The requirements in [I-D.amante-i2rs-topology-use-cases] topology use cases focus around the architecture of topology manager, orchestration manager, and policy in the figure below (IC):



- o Topo-REQ-01 (IC): The Topology Manager Should be able to collect topological information via the I2RS Client-Exchange exchange from a variety of sources in a normalized topological model. These sources can be:

- * Live Layer IGP IGPs with information about the active topology such as the LSDB database or IGP updates,

- * The I2RS must enable the inventory system information to query for information about network components which are not visible to active L3. These systems can be active or simply invisible to the L3. Examples of this are L2 Ethernet switches or ROADMS.
- * Statistic Collection systems that provide traffic information, such as traffic demands or link utilizations.

(from [section 3.2](#))

- o Topo-REQ-02 (OC): Topology information is provided from Clients to high-layer applications via a northbound interface (such as ReST, Websockets, or XMPP).
- o Topo-REQ-03 (IC): Topology Manager should be able to collect and keep current topology information for multiple layers of the network: Transport, Ethernet and IP/MPLS, as well as information for multiple Layer 3 IGP areas and multiple Autonomous Systems (ASes). This information must contain cross-layer underlying Shared Risk Link Groups (SRLG) within transport or Ethernet layers. (from [section 3.2](#))
- o Topo-REQ-04 (OC): Topology manager be able to use I2RS Client-Agent protocol to collect dynamic inventory information from network elements. An example of these protocols are the Link Layer discovery protocols (LLDP, LMP, etc.) which automatically identify remote nodes and ports. (from [section 3.2](#))
- o Topo-REQ-05 (IC): I2RS Should enable the Policy manager to query and store the following types of policies:
 - * Policies that contain Logical identifier Numbering in order to correlate IP Prefixes to
 - + link based on link type (P-P, P-PE, or PE-CE),
 - + IGP Area
 - + L2 VLAN assignments
 - * Routing Configuration policies that correlate:
 - + OSPF area/ISIS Net-ID to Node (type)
 - + BGP node related policies (aggregation routes at node, max-prefix (per node), or AFI/SAFI per node)

- + Security policies - with ACLs or rate-limits
 - + Network Component access policies (for management)
- (from [section 3.3](#))
- o Topo-REQ-06 (OC): I2RS should enable a orchestration manager attached to an I2RS client to communicate with I2RS agents into order to stitch together End-to-end services for network bandwidth optimization, load balancing, and Class-of-Service with point services (Firewall or NAT) within the end-to-end service). The orchestration manager should also be able to immediately schedule any of these resources via the I2RS-Client I2RS agent exchange. (from [section 3.4](#))
 - o Topp-REQ-07 (OC): The I2RS exchange should enable a statistics collector to collect statistics from the routing function of the network nodes and archive and aggregate the statistics into a statistics warehouse. Statistics must be given and stored in a normalized form. Metadata must be stored with the statistics. (from [section 4.1.1.2](#)) (Editor: there is some suggestion of periodic reports)
 - o Topo-REQ-08 (IC): I2RS Client-I2RS agent exchange must be provide enough interoperability that the Topology manager, Policy manager, and inventory systems can be available from different vendors
 - o Topo-REQ-09 (IC): TE tunnels must be able to be created by the exchange between the I2RS client and the I2RS agent. (from [section 4.1.1](#))
 - o Topo-Req-10 (NA): I2RS must provide a common and up-to-date normalized view of the topologies that that support security auditing, and IP/MPLS Provisioning (L2/L3) which includes:
 - * Identifying Service PE's in all markets/cities where the customer has identified they want service,
 - * Identifying one or more existing Servies PE's in each city with connectivity to the access network(s) (e.g.: SONET/TDM) used to deliver the PE-CE tail circuits to the Service's PE),
 - * Obtain via query/notification the available capacity on Services PE in both the PE-CE access interface and its uplinks to terminate the tail circuit
 - * Providing the context in I2RS for an iterative query mechanism needed by I2RS client attached to the the Topology to narrow

down the scope of resources to the set of Services PEs with the appropriate uplink bandwidth and access circuit capability plus capacity to realize the requested VPN service.

(from [section 4.1.2](#))

- o Topo-REQ-11 (NA): The VPN application attached to the I2RS client should be able to hand the I2RS Client a candidate list of Service PE's and associated access circuits to set up a Customer's VPN service into the network. (from [section 4.1.3](#)) [Editor's note This request shares requirements with VCoD-REQ-01.]
- o Topo-REQ-12 (NA): The Topology Manager associated with the I2RS client must be able to use the normalized view of the network to set up additional queries (or notification publications) to provide an accurate and comprehensive picture in order a) diagnose faults/failures, and b) augment the network with additional services, and c) provide network topology maps for different purposes. (from [section 4.1.3](#))
- o Topo-REQ-13 (IC): The I2RS client-agent exchange and informational models should support a Virtual Network Topology (VNT) comprise of one or more LSPPS and lower layer resources. The VNT of MPLS must be able to link lower layer resources with the higher layer, and present a normalize form the the PCE as defined [[RFC5623](#)].
- o Topo-REQ-14 (OC): The I2RS client-agent protocol and models should support the use of a PCE to compute MPLS-TE paths within an "domain" (IGP area), or across multiple "domains" (multi-area AS, multiple ASes") as specified in [[RFC4655](#)]. This means the PCE Informational model should support:
 - * enhanced computation in the single IGP domain
 - * cross-AS path computation based on the multiple entrance of exit points from an AS,
 - * linking multiple PEs in multiple domains together, and
 - * synchronization of TED associated with the PCE to the topology manager (via I2RS client/messages), and
 - * sending read/writes to the head-end-nodes

([section 4.3](#))

- o Topo-REQ-15 (OC): the I2RS protocol and Information models should support the ALTO ([[RFC5693](#)]) generation of abstract network

topology models and the APIs it support over web-service API. The ALTO abstract network topology comes in two forms: Network Map (based prefix-to PID mapping), and Cost map. The ALTO map is automatically generated from BGP and IGP data which the ALTO server queries from the network and makes available to applications via web-service API. (from [section 4.4](#))

6.4. Virtual Topology Data Model

The [[I-D.medved-i2rs-topology-requirements](#)] specifies the following Topology Data Model requirements (IC):

VT-TDM-REQ1 (IC): The topology data model MAY be able to describe topology and characteristics of the following layers:

- * Optical DWDM (optional) (OC),
- * Optical OTN (optional) (OC),
- * L2 (Aggregated links, L2 topologies) (IC),
- * IP/MPLS (IC),
- * VPNs (IC), and
- * Services (such as cloud services, or CDNs).

VT-TDM-REQ2 (IC): The topology data model MUST support multiple Autonomous System deployments.

VT-TDM-REQ3 (IC): The I2RS topology data model must support include topology information from multiple Administrative Domains or multiple elements into a single common format.

VT-TDM-REQ4 (IC): The I2RS topology data model MUST be able to convey enough information so that an I2RS client can correlate topologies in different layers and multiple Autonomous Systems.

VT-TDM-REQ5 (NA): The topology data model MUST support multi-layer group of elements as a means of coalescing different SFF Nodes and links into a network layers from various layers. For example, links with IPv4 addresses might represent Layer 3 of the network topology while links with Ethernet MAC addresses might represent Layer 2.

VT-TDM-REQ6 (IC): The topology model should allow association between components of different layers. For example, Layer 2 port

may have several IPv4/IPv6 interfaces. The Layer-2 port and the IPv4/IPv6 interfaces would have an association.

VT-TDM-REQ7 (NA): The topology model MUST represent both inactive and active topologies in the topology Data base. Inactive topologies may include new line cards, ports in down state, etc.

VT-TMF-DM-REQ8 (NA): The topology data model MUST be hierarchical and MUST support summarization of sub-topologies. Topology summarization and creation of abstract topologies can be provided by either by the application associated with the I2RS client, or by the I2RS Agent prior to transmission to the I2RS client.

VT-TDM-REQ9 (IC): The topology data model MUST be able to describe abstract topologies. Abstract topologies can contain real and abstract nodes and real and abstract links. An abstract topology MAY be used by a provider to describe characteristics of a transit network (bandwidth, delay, protection, etc.)

VT-TDM-REQ10 (OC): The topology data model MUST support dynamic data, such as link and node utilizations (perhaps as optional attributes).

VT-TDM-REQ11a (??): The topology data model MUST allow I2RS client-agent to be able to identify and query for the path between two nodes.

VT-TDM-REQ11b (OC): The topology data model should support the I2RS Client requesting the I2RS Agent to trace the path at all network layers that participate in the delivery of packets between two nodes. This trace MAY involve either an I2RS Agent information trace or the I2RS Agent requesting the routing function trace the path at multiple levels (L3/L2.5/L2/L1)

VT-TDM-REQ12 (IC): The topology data model MUST support multiple BGP Autonomous Systems and multiple IGP areas. Support for multiple administrative domains is for further study.

VT-TDM-REQ13 (IC): The topology data model MUST be human-friendly, i.e. not SNMP MIBs, but something much more analogous to YANG models.

VT-TDM-REQ14 (IC): The data model SHOULD support topology abstraction, allowing clients that consume topology information in a constrained manner. For example, a client wishing to view only interfaces and nodes present in a sub-graph of the Layer 3 topology should be able to specify an interest in this subset of

information rather than having to read out and parse through the entire set of links and nodes.

6.5. Virtual Topology IP Data Model

The [[I-D.medved-i2rs-topology-requirements](#)] specifies the following requirements for the Virtual Topology IP Data Model's IP/MPLS links and topologies (IC):

- o VT-TDM-IP-REQ1 (IC): The I2RS topology data model for the IP/MPLS layer MUST support both link topology and prefixes,
- o VT-TDM-IP-REQ2 (IC): The I2RS agent may import topology information from the routing processes, IGP process, BGP-LS information, or management processes.
- o TM-DM-IP-REQ3 (IC): The I2RS SFC Data model must support links that are IP/MPLS with the following attributes:
 - * local and Remote anchor node IDs (Router ID, AS#, Area ID, MT topology),
 - * metrics,
 - * admin group,
 - * max bandwidth links
 - * unreserved/utilized bandwidth
 - * link-protection type
 - * MPLS protocol mask
 - * link prefix
 - * link characteristics (BW, Delay, error rate)
 - * Link Description, and
 - * Link-specific timers (Hello and Holddown).

6.6. Virtual Topology Network Element

The [[I-D.medved-i2rs-topology-requirements](#)] specifies the following requirements (IC):

- o VT-NE-01 (IC): Each network element should contain an inventory data base which should be a definitive source of information with respect to the physical HW and Logical, logically significant identifiers (E.g. VLANs). The I2RS client should be able to import data from this DB into the I2RS Node IM or SFC IM.
- o VT-NE-02 (IC): The inventory DB of the network element should be augmented with the physical properties associated with the ports/interfaces that are directly connected to the device (BW, media type). The I2RS client should be able to import data from this augmented DB into the I2RS Node IM or SFC IM.
- o NE-3 (NA): The I2RS client may write information into the NE inventory data base via the Network-element Data Model that the network element may not be able to learn on its own. This information may include the physical location (address), rack/bay information.

7. Requirements from SFC Use Cases

The SFC use case document in [[I-D.bitar-i2rs-service-chaining](#)] suggests that the following requirements (OC):

SFC-Use-REQ01 (IC):Address

has the following address requirements:

- * IP address
- * service-node tuple (service node IP address, Host system address)
- * host-node tuple (hosting system IP-address, system internal identifier)

SFC-Use-REQ02 (IC):Supported Service Types

SHOULD include: NAT, IP Firewall, Load balancer, DPI, and others

SFC-Use-REQ03 (IC):Virtual contexts

SHOULD include:

- * Maximum Number of virtual contexts supported
- * Current number of virtual contexts in use
- * Number of virtual contexts available

- * Supported Context (VRF)

SFC-Use-REQ04 (IC): Customers currently on node

SFC-Use-REQ05 (IC): Customer Support Table (per customer ID)

- * Customer-id
- * List of supported Virtual Contexts

SFC-Use-REQ06 (OC): Service Resource table

which includes:

- * index: Comprised of service node, virtual context, service type
- * service bandwidth capacity
- * supported packet rate (packets/second)
- * supported bandwidth (kps)
- * IP Forwarding support: specified as routing-instance(s), RIBs, Address-families supported
- * Maximum RIB-size (WG Note: problematic)
- * Maximum Forward Data Base size (WG Note: problematic)
- * Maximum Number of 64 bit statistics counters for policy accounting
- * Maximum number of supported flows for services (WG Note: problematic)

SFC-Use-REQ07 (IC): Virtual Network Topology (VNT)

which includes:

- * number of access points to which service topology applies
- * topology of access points

8. Requirements from Traffic Steering Use Cases

The requirements from the Traffic Steering use case described in [[I-D.chen-i2rs-ts-use-case](#)] are (OC):

- o TS-REQ01 (IC): The I2RS Client-Agent must be able to collect the topology (especially the exit links) and the traffic load of each link;
- o TS-REQ02 (IC): The I2RS Client-Agent must be able to read the local rib of each DC/Metro gateway and the policies deployed on each gateway;
- o TS-REQ03 (IC): The I2RS Client-Agent must be able to add or delete or modify the relevant rib items and relevant policies to steer the traffic as expected; and adjust traffic placement.
- o TS-REQ-04 (IC): The I2RS Client-Agent must have the ability to collect the LSP information either from the PCE or directly from network devices;
- o TS-REQ-05 (OC): The I2RS Client-Agent must have the ability to collect the traffic matrix of the network, this is used to help the I2RS client to determine how to adjust the traffic placement;
- o TS-REQ-06 (IC): The I2RS Client-Agent must have the ability to read the rib information and relevant policies of each network node;
- o TS-REQ-07 (OC): collect the topology and segment information needed to help the I2RS client to compute the end-to-end path;
- o TS-REQ-08 (OC): read rib (especially the segment routing rib) information;
- o TS-REQ-09 (??): add/delete/modify the segment rib, this finally determines how the traffic is forwarded.

9. Requirements from MPLS TE Networks Use Cases

These are the requirements from the Traffic Steering use case described in [[I-D.huang-i2rs-mpls-te-usecases](#)] (OC):

- o MPLS-TE-REQ-01 (OC): Network programming software managing the static CR-LSP devices may incorporate an I2RS Client along with a path calculation entity, a label management entity, and a bandwidth management entity. The I2RS Client should be able to

communicate the static configuration to the network nodes, and monitor the status of the CR-LSPs.

- o MPLS-TE-REQ-02 (OC): The I2Client should be able to synchronously send the configuration for all of the network nodes from egress node to ingress node via the I2RS Agents attached to each node, and be able to delay the final ingress node configuration until all the I2RS AGents on all other nodes toward the egress have denoted a successful path set-up.
- o MPLS-TE-REQ-03 (OC): MPLS TE defines abundant constraints such as explicit path, bandwidth, affinity, SRLG, priority, hop limit, and others. The I2RS Client Agent exchange should be able to signal concurrent local path calculation could obtain an optimized result and allow more services to be held in a TE network. The I2RS Agent should be able to trigger a global concurrent re-optimization at a specific time on multiple nodes by communicating with each node's I2RS agent.
- o MPLS-TE-REQ-04 (NA): The I2RS client should be able to manually calculate a re-optimization of the the MPLS TE network and send the new constraints including the calculated path to each node via the I2RS agent with an indication to re-signal the TE LSPs with make-before-break method.
- o MPLS-TE-REQ-05 (OC): With I2RS, the node's I2RS agent should be able to send to an I2RS client a status notification that not enough resources exist for a back up LSP and TE tunnel. Upon receiving this notification the I2RS client should be able to trigger concurrent calculation for the failed path calculation of the backup LSP or TE tunnel and send the updated paths to I2RS agents with a command to re-signal the TE LSPS with make-before-break Method.
- o MPLS-TE-REQ-06 (NA): With I2RS, upon receipt the failure notification from an I2RS Agent, the I2RS client would create a global concurrent optimization to handle the failure event. This would occur by the I2RS client signalling the I2RS agents on all nodes to: a) trigger a new concurrent calculation of the backup LSP or TE tunnel via failed path calculation, and b) re-signal updates to the TE LSPs process with a make-before-break method.
- o MPLS-TE-REQ-07 (NA): Upon receiving a signal an upgrade event signal (from operator), the I2RS client could calculate another path for the affected TE tunnels to deviate traffic away from the resource being upgraded, and then send the request to I2RS agents on the appropriate nodes to move the traffic. After the upgrade completes, the I2RS client can simply remove I2RS configurations

causing the traffic to revert to the original path. Or, the I2RS can re-optimize the TE tunnels for another pathways (E.g. as a part of a sequence of upgrades).

- o MPLS-TE-REQ-08 (OC): I2RS agents can notify I2RS Clients of impending or existing MPLS TE overload conditions that might cause TE LSP rejections. This overload conditions include: due to CPU, memory, LSP label space, or LSP numbers.
- o MPLS-TE-09 (IC): Automatic bandwidth adjustment applications can also be linked to the I2RS clients need to monitor the traffic on TE tunnels in order to provide traffic analysis. The I2RS client should be able to read the TE Tunnel topology and the bandwidth analysis in order to automatically calculate a new path for the TE tunnel if it is needed. The I2RS Client also needs to be able to the I2RS agents in the nodes to install the new TE Tunnels with the make-before-break option.
- o MPLS-TE-REQ-10 (IC): With I2RS, the node failure or link failure can be part of the notification stream sent by an I2RS Agent to an I2RS Client on a centralized server gathering information.
- o MPLS-TE-REQ-11 (IC): The I2RS client can notify the I2RS agents on specific nodes (or devices) to re-signal TE LSPs one by one if there is a resource dependency.
- o MPLS-TE-REQ-12 (IC): The I2RS Client can gather the TE LSPs' state from I2RS Agents on all nodes in order to coordinate such handling of LSP resources.
- o MPLS-TE-REQ-13 (OC): The I2RS Clients collecting information from I2RS Agents can be arranged in a hierarchy to provide scaling of collections. An application hosting an I2RS client collecting information from I2RS Agents on nodes can have an I2RS Agent that reports combined information to a single location.

10. Requirements from MPLS LDP Networks Use Cases

These are the I2RS requirements for the MPLS LDP use case described in [[I-D.chen-i2rs-mpls-ldp-usecases](#)]:

- o MPLS-LDP-REQ-01 (IC): The I2RS Client-agent exchange should allow the distribution of the configuration for PWE3, MPLS LDP and associated protocols to be distributed from a central location where the global PWE3 provisioning information could be stored. The I2RS Client-Agent exchange should also be able to push the configuration of the local LDP LSR ID and peer addresses to set up the targeted session to the pseudowire endpoints.

- o MPLS-LDP-REQ-02 (IC): When an the end-user wants to disable IPoMPLS (IP over MPLS) application on a L2VPN/PW Targeted LDP session, the I2RS Client-I2RS agent should be able to set type of application over the established LDP session. In this way LDP speaker can only advertise to its peer the application data which the user is interested in.
- o MPLS-LDP-REQ-03 (OC): The I2RS Agent notifications should allow an I2RS client to subscribe to a stream of state changes regarding the LDP sessions or LDP LSPs from the I2RS Agent. Specifically it is important that LDP session is tract for sessions state coming up or going down. The I2RS Client-I2RS Agent exchange should allow additional queries to the AAgent to determine a) why the service is invalid, b) calculating whether an alternate path should be switched to, and c) determining how to switch to other links or nodes in order to recover from the link failure or node failure.
- o MPLS-LDP-REQ04 (IC): The I2RS interface provides way to monitor and control the limited resources on these access devices. The I2RS client should be able to instruct the I2RS agent in each of these devices to set the maximum number of LDP LSPs in each device prior to enabling LDP on the devices. The I2RS client should also be able to enable a notification service on each device with a with a warning threshold. Once the number of LDP LSPs reaches the threshold, the I2RS agent will send a notification message to the I2RS client. Often the I2RS client will be associated a network management agent that can determine what next steps need to be done based on policy or operator input.

11. Requirements from Mobile Backhaul Ues Cases

Mobile BackHaul Use cases described in [[draft-ietf-zhang-mbb-usecases-01](#)] are:

- o MBH-REQ-01 (OC): The I2RS client-agent communication can distribute position-critical changes to IGP nodes using this global knowledge to quicken changes to support traffic during failures or traffic overloads. To enable this feature, the I2RS Clients-Agent communication needs to pass information on which IGP process or Level or Area the given node and links belong to.
- o MBH-REQ-02 (OC): I2RS must allow operators to use of I2RS clients to distribute time-critical changes in configuration to I2RS agents associated with each routing node. This feature will simplify and automate configuration and monitoring of a mobile backhaul network to allow it to readily adapt to changing network sizes (and scales) and radio applications.

- o MBB-REQ-03 (OC): I2RS Clients-Agent communication needs to pass information on:
 - * T-LDP configurations and status;
 - * BGP peer configurations, peer topologies and status;
 - * BGP-based LSP topologies and status;
 - * Reset VPN topologies, and per node configurations;
- o MBB-REQ04 (IC): Route policy enforcement in mobile backhaul networks needs to be more dynamic and flexible than the current methods take hours (or even days) to configure route policy across a network. The I2RS interface must provide a programmatic way to configure (both policy and device) and monitor thousands of devices individually whose configuration is based on the devices role (such as ASRSs in one AS, ASBRs between ASs and other service-touch nodes).
- o MBB-REQ-05 (NA): I2RS clients should be able to contact I2RS agents on nodes to query role-based information from the network status. After collecting the status, the I2RS client can develop the BGP policies based on role information and push the BGP policies to the I2RS agents that would load the alternate policies into the network device. The I2RS Agents loading the alternate policies could then send status back to the I2RS Client.
- o MBH-REQ06 (??): I2RS clients can provide centralized control of many network devices via the I2RS Client-Agent communication. The I2RS programmatic interface can automate the collection and analysis of each device's capability so that the centralized I2RS client could calculate the optimal LSP path and distribute the configuration to individual devices. Automation of the collection of device capability should be available as query, notification, or a published stream.
- o MBH-REQ07 (NA): While the I2RS RIB Information Model [\[\[I-D.ietf-i2rs-rib-info-model\]\]](#) provides for routes with tunnels or MPLS LSP, the features defined in this model are not sufficient to configure both types of LSPs needed for the VPN technology in mobile backhaul networks. Additional I2RS Informational models need to be created to support these features.
- o MBH-REQ08 (NA): The hierarchical protection architecture in mobile backhaul network offer high network reliability and more flexibility to meet the various needs of the tunnels and services. The I2RS interface in this use case is needed to automate the

configuration and monitoring so that tunnel protection and service protection interwork in a flexible and reliable manner.

- o MBB-REQ09 (OC): The I2RS architecture (client-agent) should allow the two features for network monitoring naturally in its basic modes:
 - * allow a combination of multi-layer network monitor tools with exact detection parameters to be configured on the network device
 - * Facilitate the reporting the detection result as notification or publication stream

It is important the result of these features allow the outages and traffic congestion or discards to be detected real-time with I2RS Client(s) in each node, and the detection result will be reported to the I2RS agents to get the exact status of the network.

12. Requirements from Large Data Flows are

Each of these requirements has been given an an ID number of L-Flow-nn for ease of reference.

The requirements from the Large Data Flows use case described in [[I-D.krishnan-i2rs-large-flow-use-case](#)] are (IC):

L-Flow-REQ-01 (IC): For redirecting large flows to a specific component, a PBR entry should be programmable for the flow with its nexthop that identifies the specific LAG or ECMP component.

L-Flow-REQ-02 (IC): For adjusting the weights used to distribute traffic across components of the LAG or ECMP, I2RS should provide a programmable mechanism should be provided that identifies ECMP entries and is able to associate weights that can be programmed for each of the components. To do this in a scalable fashion, it would be useful to have the notion of an ECMP nexthop that is used by multiple routes

L-Flow-REQ-03 (IC): The I2RS interface (protocol/IMs) should allow for a globally optimal path is programmed in the IP network using hop-by-hop PBR rules. These PBR rules may include:

- * Being able to adjust the weights of the ECMP table for different nexthops should be adjusted to factor the large flows
- * Being able to address an ECMP group, so that all routes sharing an ECMP group are addressed together.

- * the ability to program PBR entries at the edge LSR, and
- * the ability to program new LSPs in the network.

L-Flow-REQ-04 (OC): The I2RS protocol should be able to invoke the link aggregation IEEE 802.1AX Marker Protocol via the I2RS protocol. This is useful during a period of rebalancing occurs before flows are moved.

L-Flow-REQ-05 (IC): The I2rs protocol should allow Quality of Service (QoS) actions such as rate-limiting, re-marking, or discarding can be performed on the flows based on configured policies and nexthop redirection actions to be programmed, and to be programmed independently of of each other.

L-Flow-REQ-06 (IC): Once a large flow has been detected, I2RS must be used to modify the forwarding tables in the router to:

- * In the case of large flow load balancing, be able to redirecting the large flow to a particular member with the LAG or ECMP group and readjusting the weights of the other members to account for the large flow
- * In the case of DDoS mitigation, the action involves rate limiting, remarking or potentially discarding the large flow in question.

13. Large Data Collection Systems

The requirements from the Large Data Collection Systems Use cases described in [[draft-swhyte-i2rs-data-collection-system](#)] are (OC):

L-Data-REQ-01 (OC): I2rs must be able to collect large data set from the network with high frequency and resolution with minimal impact to the device's CPU and memory.

L-Data-REQ-02 (IC): I2RS must be able to use a database model where the data on the network node must be able to be described in the I2RS exchange as the data plus the structure of the data. The I2RS management system consumes and understand the data only after it consumes and understand the database model or has been trained by vendor published model

L-Data-REQ-03 (IC): I2RS should use a pub-sub model which allows scaling plus push or pull of data.

L-Data-REQ-04 (IC): I2RS should support capability negotiation to inform a subscriber of the options for publication of data. The options include transport, security, and error handling.

L-Data-REQ-05 (IC): The I2RS data transfer should be format agnostic. This means the publisher and subscriber may agree upon XML, JSON, MTL, protobufs or any other format.

L-Data-REQ-06 (IC): I2RS Transports must be able to be chosen by a I2RS Client-I2RS Agent pair. An I2RS Client-I2RS Agent pair should be allowed to negotiate the transport options from a list of options.

L-Data-REQ-07 (IC): The I2RS interface (protocol and IMs) should allow a subscriber to select portions of the data model.

L-Data-REQ-08 (IC): The I2RS interface (protocol and IMs) should allow for multiple publish subscriptions at a time.

L-Data-REQ-09 (IC): Timestaps should be associated with data that requires it. Not all data will require a time stamp. Additional time stamps may be added.

L-Data-REQ-10 (IC): The I2RS should support the query and "introspection" of the data model. The Introspections provides support for data verification, easier inclusion in legacy data, and easier merging with data streams.

L-Data-REQ-11 (IC): After the I2rs Client-Agent have exchanged capabilities, a database model, and filters used to select elements of the model to subscribe to, the framework should support a standard way to register for all the data desired, using whatever capabilities were advertised by the node. Once registration is complete, the control channel can be closed. Ensuring subscriptions are correct, complete, and replicated or not, is up to the overall system and not the agent on the network node.

L-Data-REQ-12 (IC): The I2RS interface should support user subscriptions to data with the following parameters:

- * push of data synchronously or asynchronously via registered subscriptions
- * pull data off in a one-shot pull or in multiple sequences
- * provide dynamic subscriptions that can be setup via IPFIX feed

- * support of subscriber and consumer I2RS Client-agent pairs
- * allow remapping of a node's databases

L-Data-REQ-13 (IC): The I2RS interface must handle and report errors that occur with data subscription, stale data, repeated transport failures, and other (yet unknown) errors

14. CDNI

The requirements from the Content Delivery Network Interaction described in [[I-D.shin-i2rs-usecases-cdni-request-routing](#)] are (OC):

- o CDNI-REQ-01 (OC): The I2RS interface should support two CDNI functionalities [[I-D.ietf-cdni-framework](#)]:
 - * Request Routing Interface - Footprint and Capabilities Advertisement; the asynchronous advertisement of footprint and capabilities by a dCDN that allows a uCDN to decide whether to redirect particular user requests to that dCDN via the ALTO protocol; and
 - * Request Routing Interface - Redirection; the synchronous operation of actually redirecting a user request via I2RS manipulation of the routing plane.
- o CDNI-REQ-02 (OC): The I2RS (Protocol and IM) should provide facilities to enable the query/response of information from an ALTO services in a node routing functions so that the upstream CDN provider can select a proper downstream CDN provider for a given end user request.
- o CDNI-REQ-03 (OC): I2RS (protocol and IM) should provide facilities to enable I2RS can help the upstream CDN provider to redirect a content request message to a downstream CDN provider for a given end user request as with the following features:
 - * The uCDN relays this message between I2RS Clients and I2RS agents with content distribution metadata, and queries the dCDN whether user request message can be delivered. This query can have multiple dCDN that the user message can be delivered to.
 - * the I2RS agent associated with the dCDN delivery requests indicating which dCDN (if any) the user message can be delivered to.
 - * Allow dCDN to be managed to deliver content by having the messages to signal back to the uCDN the (destination (?)) IP

address for the content, on the dCDN, and the pathway between the uCDN for surrogate deliver via the dCDN of user data. Part of this management is the passing of URL of the surrogate in dCDN (for HTTP Redirection to be transmitting) back from the dCDN to the uCDN so the uCDN can inform the end user.

15. IANA Considerations

This document makes no request of IANA.

16. Security Considerations

Routing information is very critical and sensitive information for the operators. I2RS should provide strong security mechanism to protect the routing information that it could not be accessed by the un-authorized users. It should also protect the security and integrity protection of the routing data.

17. References

17.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC3746] Yang, L., Dantu, R., Anderson, T., and R. Gopal, "Forwarding and Control Element Separation (ForCES) Framework", [RFC 3746](#), DOI 10.17487/RFC3746, April 2004, <<http://www.rfc-editor.org/info/rfc3746>>.

17.2. Informative References

- [I-D.amante-i2rs-topology-use-cases] Medved, J., Previdi, S., Lopez, V., and S. Amante, "Topology API Use Cases", [draft-amante-i2rs-topology-use-cases-01](#) (work in progress), October 2013.
- [I-D.bitar-i2rs-service-chaining] Bitar, N., Heron, G., Fang, L., ramki, r., Leymann, N., Shah, H., and W. Haddad, "Interface to the Routing System (I2RS) for Service Chaining: Use Cases and Requirements", [draft-bitar-i2rs-service-chaining-01](#) (work in progress), February 2014.

[I-D.chen-i2rs-mpls-ldp-usecases]

Chen, X. and Z. Li, "Use Cases for an Interface to LDP Protocol", [draft-chen-i2rs-mpls-ldp-usecases-00](#) (work in progress), October 2013.

[I-D.chen-i2rs-ts-use-case]

Chen, M. and S. Hares, "I2RS Traffic Steering Use Case", [draft-chen-i2rs-ts-use-case-01](#) (work in progress), July 2014.

[I-D.hares-i2rs-use-case-vn-vc]

Hares, S. and M. Chen, "Use Cases for Virtual Connections on Demand (VCoD) and Virtual Network on Demand (VNoD) using Interface to Routing System", [draft-hares-i2rs-use-case-vn-vc-03](#) (work in progress), July 2014.

[I-D.huang-i2rs-mpls-te-usecases]

Huang, T., Li, Z., and S. Hares, "Use Cases for an Interface to MPLS TE", [draft-huang-i2rs-mpls-te-usecases-02](#) (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-13](#) (work in progress), February 2016.

[I-D.ietf-i2rs-problem-statement]

Atlas, A., Nadeau, T., and D. Ward, "Interface to the Routing System Problem Statement", [draft-ietf-i2rs-problem-statement-10](#) (work in progress), February 2016.

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

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

[I-D.ietf-sfc-problem-statement]

Quinn, P. and T. Nadeau, "Service Function Chaining Problem Statement", [draft-ietf-sfc-problem-statement-13](#) (work in progress), February 2015.

[I-D.ji-i2rs-usecases-ccne-service]

Ji, X., Zhuang, S., Huang, T., and S. Hares, "I2RS Use Cases for Control of Forwarding Path by Central Control Network Element (CCNE)", [draft-ji-i2rs-usecases-ccne-service-02](#) (work in progress), July 2014.

[I-D.keyupate-i2rs-bgp-usecases]

Patel, K., Fernando, R., Gredler, H., Amante, S., White, R., and S. Hares, "Use Cases for an Interface to BGP Protocol", [draft-keyupate-i2rs-bgp-usecases-04](#) (work in progress), July 2014.

[I-D.krishnan-i2rs-large-flow-use-case]

ramki, r., Ghanwani, A., Kini, S., McDysan, D., and D. Lopez, "Large Flow Use Cases for I2RS PBR and QoS", [draft-krishnan-i2rs-large-flow-use-case-04](#) (work in progress), April 2014.

[I-D.lapukhov-bgp-routing-large-dc]

Lapukhov, P., Premji, A., and J. Mitchell, "Use of BGP for routing in large-scale data centers", [draft-lapukhov-bgp-routing-large-dc-06](#) (work in progress), August 2013.

[I-D.medved-i2rs-topology-requirements]

Medved, J., Previdi, S., Gredler, H., Nadeau, T., and S. Amante, "Topology API Requirements", [draft-medved-i2rs-topology-requirements-00](#) (work in progress), February 2013.

[I-D.shin-i2rs-usecases-cdni-request-routing]

Shin, M. and S. Lee, "CDNI Request Routing with I2RS", [draft-shin-i2rs-usecases-cdni-request-routing-00](#) (work in progress), July 2014.

[I-D.swhyte-i2rs-data-collection-system]

Whyte, S., Hines, M., and W. Kumari, "Bulk Network Data Collection System", [draft-swhyte-i2rs-data-collection-system-00](#) (work in progress), October 2013.

[I-D.white-i2rs-use-case]

White, R., Hares, S., and A. Retana, "Protocol Independent Use Cases for an Interface to the Routing System", [draft-white-i2rs-use-case-06](#) (work in progress), July 2014.

[I-D.zhang-i2rs-mbb-usecases]

Zhang, L., Li, Z., Liu, D., and S. Hares, "Use Cases of I2RS in Mobile Backhaul Network", [draft-zhang-i2rs-mbb-usecases-01](#) (work in progress), February 2014.

[RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", [RFC 4655](#), DOI 10.17487/RFC4655, August 2006, <<http://www.rfc-editor.org/info/rfc4655>>.

- [RFC5212] Shiomoto, K., Papadimitriou, D., Le Roux, JL., Vigoureux, M., and D. Brungard, "Requirements for GMPLS-Based Multi-Region and Multi-Layer Networks (MRN/MLN)", [RFC 5212](#), DOI 10.17487/RFC5212, July 2008, <<http://www.rfc-editor.org/info/rfc5212>>.
- [RFC5286] Atlas, A., Ed. and A. Zinin, Ed., "Basic Specification for IP Fast Reroute: Loop-Free Alternates", [RFC 5286](#), DOI 10.17487/RFC5286, September 2008, <<http://www.rfc-editor.org/info/rfc5286>>.
- [RFC5623] Oki, E., Takeda, T., Le Roux, JL., and A. Farrel, "Framework for PCE-Based Inter-Layer MPLS and GMPLS Traffic Engineering", [RFC 5623](#), DOI 10.17487/RFC5623, September 2009, <<http://www.rfc-editor.org/info/rfc5623>>.
- [RFC5693] Seedorf, J. and E. Burger, "Application-Layer Traffic Optimization (ALTO) Problem Statement", [RFC 5693](#), DOI 10.17487/RFC5693, October 2009, <<http://www.rfc-editor.org/info/rfc5693>>.

Authors' Addresses

Susan Hares
Huawei

Email: shares@ndzh.com

Mach Chen
Huawei Technologies
Huawei Bld., No.156 Beiqing Rd.
Beijing 100095
China

Email: mach.chen@huawei.com

