

Routing Area Working Group
Internet-Draft
Intended status: Informational
Expires: August 18, 2014

S. Hares
Hickory Hill Consulting
M. Chen
Huawei Technologies
February 14, 2014

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-02

Abstract

Software Defined Networks (SDN) provides a way to virtualize and abstract the network and present the virtual or abstract resources to third-party applications running in software. Applications can utilize a programmable interface to receive these virtual or abstract resources descriptions in a form that allows monitoring or manipulation of resources within the network. The Interface to the Routing System (I2RS) provides an interface directly to the routing System to monitor best paths to any destination or change routes in the routing information base (RIB) or MPLS Label Information Base (LIB). The I2RS interfaces may be combined with other interfaces to the forwarding plane (ForCES (RFC3746)), device configuration (NETCONF), or mid-level/peer-to-peer (ALTO, draft-ietf-alto-protocol) system to create these virtual pathways.

This document outlines how SDN networks can use the I2RS interface to implement an automated set of network services for the Virtual Connection on Demand (VCoD) and Virtual Network on Demand (VNoD). These systems provide service routing a better way to create paths within a hub and spoke environment, and provide service routing the ability to create pathways based on service.

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 August 18, 2014.

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. Background	3
3. Virtual Circuit on Demand	5
4. Virtual Network on Demand (VNoD)	8
5. Automated On Demand Networks	9
6. IANA Considerations	10
7. Security Considerations	11
8. References	11
8.1. Normative References	11
8.2. Informative References	11
Authors' Addresses	12

1. Introduction

The Interface to the Routing System (I2RS) architecture ([I-D.ietf-i2rs-architecture]) describes a mechanism where the distributed control plane can be augmented by an outside control plane through an open accessible programmatic interface. I2RS provides a "halfway point" between completely replacing the traditional distributed control planes and directly configuring devices via off-board processes.

This draft proposes a set of use cases using I2RS mechanisms to implement a Software Defined Network (SDN) to enact virtual connections and virtual networks as automated services. This document focuses on how I2RS would support two automated network services: Virtual Connection on Demand (VCoD) and Virtual Network on Demand (VNoD). The Virtual Connections on Demand (VCoD) and Virtual Network on Demand (VNoD) may be used within hub-spoke networks and

improved service routing. In the future, an application enabled SDN services may provides the Virtual Circuits (VCoD) and Virtual Networks on Demand (VNoD) for any type of network service.

This document contains a background section, a VCoD use case, a VNoD use case, and notes on future on-demand network services (connections and network). Those familiar with I2RS problem statement ([I-D.ietf-i2rs-problem-statement]), I2RS architecture ([I-D.ietf-i2rs-architecture]), and the concepts of Virtual Connections (VCs) or Virtual Networks (VNs) may wish to skip directly to the use cases.

SDN is a new to the Internet space. Each new adventure in Internet network services requires lots of use cases so that the IETF may determine the critical protocols to be developed.

2. Background

Applications and network layer flows have run independently since the Internet started in the late 1980s. Provisioning of network services and big flows has been done by service providers statically or with proprietary processes. Recently, new server and host technologies have increase application data traffic flows across the network. With the advent of Data Center providers and cloud services, applications life cycles have shortened to weeks rather than years. The need for fast automated provisioning of virtual network connections or quick provisioning of virtual private networks has increased.

Software Defined Networks have three areas of challenge to provide such quick network services: a) how to control the network flows, b) interfaces to networks, and c) how do calculate where these network flows go.

Network flows can be controlled at the forwarding device level or the network control plane level. Various programmatic interfaces have been proposed to provide control over individual forwarding devices. ForCES ([RFC3746]) provides mechanisms to replace the dynamic control plane processes on individual forwarding devices throughout a network with off box processes that interact with the forwarding tables on each device. Another example is NETCONF, which provides a fast and flexible mechanism to interact with device configuration and policy.

The trade-off with the device level approach to control flows has to do with benefits and challenges of having control systems off-board. The benefit of off-board control systems is that the calculation unit can be centralized. The challenge of the off-board control system has both technical challenges and deployment challenges. The

technical challenge is that off-board control systems may encounter time-delays and communication failure. The deployment issues concerns utilizing new protocols for this communication which may also have issues in deployment. The promised benefits of off-board devices are reduction in operational costs, improved scaling, control, and visibility. OpenFlow, for instance, provides a mechanism to replace the dynamic control plane processes on individual forwarding devices throughout a network with off box processes that interact with the forwarding tables on each device. Another example is NETCONF, which provides a fast and flexible mechanism to interact with device configuration and policy.

The Interface to Routing System (I2RS) interface provides an interface to all aspects of the routing system as a system. This interface allows the SDN approach to utilize the existing control plane software without changing it. The I2RS agent interacts with the control plane processes to monitor best paths to any destination and to interact with the routing information base (RIB) or MPLS label information base (LIB), and forward the information to the I2RS client. Applications associated with the I2RS client can compute where network flows should go, and then instruct I2RS agents in the appropriate nodes to change RIB or LIB routes to enact the changes to traffic flows.

This document describes a set of use cases which describe how automated creation of Virtual Connection on Demand (VCoD) and Virtual Networks On Demand (VNoD) based in SDN logic can be accomplished by using an interface to the routing system (I2RS). This document first examines the current use case for I2RS of improved hub-spoke routing and better service routing using VCoD (section 2), and VNoD (section 3). Secondly, this document examines the future I2RS use case of VCoD and VNoD for any network enabled by application or SDN processing.

A bit of context on abstract services may be useful as a background.

These abstract services (VC or VN) are logical services that can be mapped to specific services. For example, a virtual circuit may be mapped to a TE-LSP. These logical services provide a uniform abstract service model that allows applications to configure VC or VN services independent of the actual network technology implementing it.

There are several types of network services that can be considered as network services over which virtual connections or virtual networks can be created. These network services include: optical, Ethernet (VLAN and SPB), Internet Protocol (IP), Multi-protocol Label Switching (MPLS). Each of these networks can provide traffic

engineered paths, policy control (e.g. Access control Lists (ACLs)), security services, or some form of virtual LAN services (VLAN, VxLAN, L2/L3 VPN). The examples in this document focus on the transport and VPN related services that can be abstracted into Virtual Connection (VC) and Virtual Network (VN).

The use cases below leverage the SDN architecture and model and the I2RS Framework to implement Virtual Circuit on Demand (VCoD) and Virtual Network on Demand (VNoD).

Please note that this draft builds on the premise that SDN solutions can augment rather than replace traditional distributed control planes. Each use case is presented in its own section.

3. Virtual Circuit on Demand

The Virtual circuit on demand (VCoD) applications associates to I2RS client (or clients) which can communicate with the I2RS agent (or agents) which control the VCoD circuit's creation, deletion, modification, query for information or status changes. This information needs to include for this application network topology, interface statistics, available circuits per node, available bandwidth on circuits. Interface statistics might be required on a historical and instantaneous time basis. The circuit statistics might also need jitter, delay, and exit-point performance.

The virtual circuits may be obtained via RIB Informational Model (RIB IM) ([I-D.ietf-i2rs-rib-info-model]) from the interface list, or from the nexthop lists. Write access to set-up new interfaces is not clearly spelled out in the current version of the RIB IM, nor are the statistics (historical or time). This use case points out additional Information Models (IMs) that need to be added to the I2RS information models.

In the example topology below, the VCoD application's I2RS client communicates with I2RS agents to set-up virtual circuits from Edge 1 to Edge 2. The I2RS client communicates with I2RS Agent-1 on node 1, I2RS Agent-2 on node 2, I2RS Agent-3 on node 3, and I2RS Agent 4 on node 4 for to set-up the virtual circuit. The VCoD application contains the necessary logic to determine the pathway from Edge 1 to Edge 2.

A second option VCoD is to have an application communicate with two I2RS clients who cooperate to set-up the virtual connections between Edge 1 and Edge 2. Information passed between the two clients can be done via other IETF protocols (E.g. stateful PCE or ALTO).

Why I2RS enabled solutions are necessary

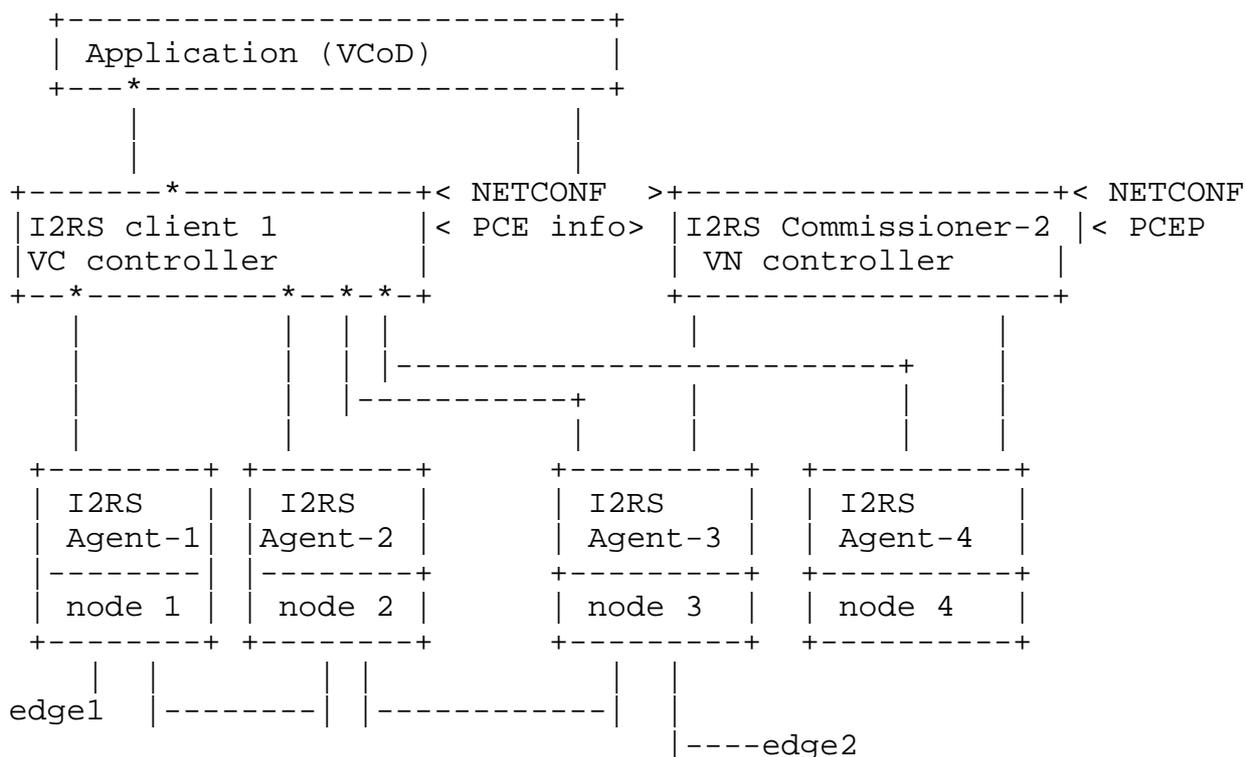
Past solutions in this area have included uses of device configuration across multiple nodes (SNMP or NETCONF based) with proprietary services combined with topology queries. The lack of coordinated responses to routing topology queries has created problems in quickly obtaining and configuring changes for Virtual Circuits. New algorithms can create better services in routing and switching. These algorithms include Fast-Reroute of RSVP or IGPs which aid the automatic re-establishment of some circuits, but the complexity of some of these algorithms increases cost within the network elements. It's often difficult to justify the added complexity in the database and algorithms of routing protocols to solve what is considered a point case.

While the set-up of these virtual circuits is possible with current technology, the lack of the I2RS-like framework makes VCoD network complex. With this support, VCoD may be able to reduce complexity on the individual nodes.

What's not in scope for I2RS

The means by which the VCoD application determines which I2RS client to associate with is outside the I2RS protocol and architecture. A list of virtual circuits per node may be queried from the RIB Informational Model's (RIB IM) ([I-D.ietf-i2rs-rib-info-model]) interface and nexthop lists. However, other means may be used to determine the possible interfaces on a node. For example, ALTO could inform the application which nodes have an I2RS Agent supporting the VCoD service, and SNMP/NETCONF could be used to determine which interfaces were configured.

Example Topology for Virtual Circuit on Demand (VCoD).



The following things need to be supported for this application:

- o I2RS Agents should provide the ability to read the virtual connection topology database for the technology supported. For optical, these are the optical connections and what node they connect to. For MPLS, this is virtual circuit available, and what nodes they connect to. For IP technologies, this could include the GRE tunnels and what interface it connects to. 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.
- o I2RS Agent should provide the ability to influence the configuration of a virtual circuit in a node.
- o I2RS Agent should provide monitor and provide statistics on the virtual connection to the I2RS client. 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.

What is needed in the RIB IM Model

The RIB IM model ([I-D.ietf-i2rs-rib-info-model] provides with each route an associated nexthop-list 0-N members. Each nexthop list is flagged with a protection preference (1 or 2), and a Load balance weight (1 to 99). If the host routes for all nodes in the topology exist within the RIB IM model's instantiation, then the nexthop provides the following information:

- o identifier for interface
- o egress interface (logical, virtual, or physical)
- o address of physical interface (IP address or MAC) plus RIB
- o tunnel encapsulation for interface (IP GRE, MPLS tunnel),
- o logical tunnel identifier
- o RIB name (for resolved look-ups)
- o flags for specialized look-ups (Discard packets, discard with error notification, receive)

The RIB IM model's primitives need to be expanded to include circuit type (p2p, mp2mp), optical connection information, and additional statistics per virtual circuit. The RIB IM model's instantiation within the protocol must provide an easy way to specify queries for this information.

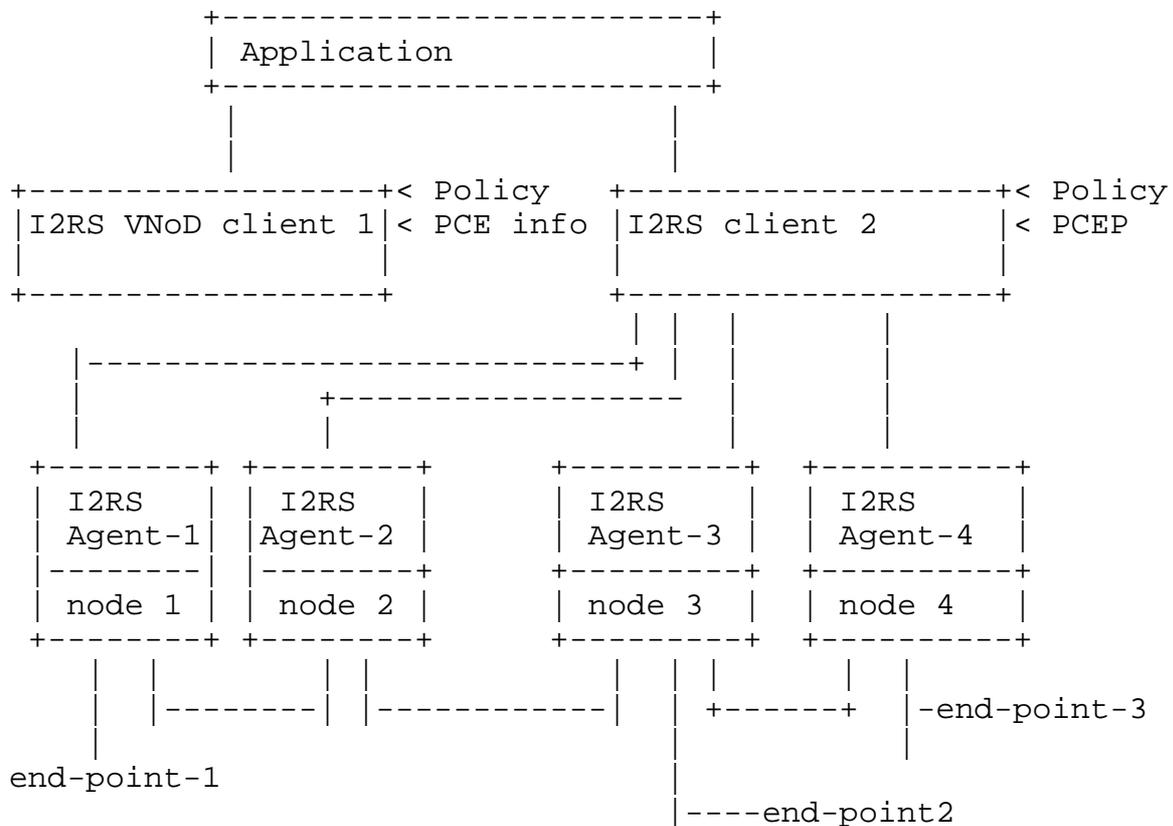
4. Virtual Network on Demand (VNoD)

Virtual Networks on Demand (VNoD) are simply extensions to the Virtual Connections on Demand concept. The I2RS client 2 is tasked to create a Virtual network instead of a single connection.

The example sequence would be that the application discovers the appropriate I2RS clients (I2RS VNoD client 1 and I2RS VNoD Client 2) which support VNoD via a protocol outside the I2RS framework (e.g. ALTO). The I2RS Client-2 works with the I2RS Agents 1-4 to set-up a virtual network. This involves the following:

- o gathering potential topology information (in order to create the network,
- o set-up the virtual network (via influencing configurations on node),
- o monitoring changes in topology (in order to potential failovers,

- o influencing changes to virtual network via configurations, and
- o removing the virtual network after the demand has expired.



This topology shares some configuration needs with the central membership computation for MPLS VPNs from (draft-white-i2rs-use-cases) but the mechanisms are not specific to MPLS VPNs.

5. Automated On Demand Networks

Automated On-Demand networks becomes a reasonable technology within a network by utilizing the I2RS architecture. While automated on-demand circuit provisioning and de-provisioning is possible now, the effort to configure and reconfigure nodes to provide the Automatic On-Demand circuits can be difficult. With I2RS, the I2RS client can instruct the I2RS Agents within a network to create On-Demand circuits and then remove the circuits returning the network to its configured state. With I2RS enhanced monitoring capability, the monitoring needed for these state changes is incorporated within the I2RS framework.

The current scope for these Automated On-Demand Circuits in the IETF's I2RS working group's charter is limited to hub-spoke networks and service routing. This section discusses the progress on the I2RS against the use cases, and proposes additional additional Automated On-Demand Circuits.

Current Status of the Automated On-Demand Functionality

Both the hub-spoke network and service network may include a centralized control network element such as [I-D.ji-i2rs-usecases-ccne-service]. These centralized control network elements may use I2RS access to individual node's RIB information via the I2RS RIB Information Model (IM) ([I-D.ietf-i2rs-rib-info-model]), or obtain full network topology information from other protocols (BGP Route Reflector, PCE ([RFC4655]), or ALTO [I-D.bernstein-alto-topo]). With the recent inclusion of ISIS link-state information into BGP TLVs via [I-D.ietf-idr-ls-distribution], all of these sources can provide centralized service can provide topology maps at the AS and IGP level.

I2RS Information Models (IM) are being proposed which can store:

- o Network Topologies (IM) [I-D.medved-i2rs-topology-im], and
- o Service Topologies IM) [I-D.wu-i2rs-IM-service-topo].

Needed Future On-Demand Networks

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). These MPLS technologies can be used to create Hub-spoke topology and service routing in networks in Carriers, Enterprise, and Data Centers. The RIB IM supports logical tunnels of type MPLS as well as IP, GRE, VxLAN and GRE.

Carriers using these MPLS technologies also use these MPLS and IP networks to support networks for Mobile BackHaul, on-demand MPLS overlays, and on-demand video conferencing networkings.

6. IANA Considerations

This document includes no request to IANA.

7. Security Considerations

This document has no security issues as it just contains use cases.

8. References

8.1. Normative References

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

8.2. Informative References

[I-D.bernstein-alto-topo]

Bernstein, G., Yang, Y., and Y. Lee, "ALTO Topology Service: Uses Cases, Requirements, and Framework", draft-bernstein-alto-topo-00 (work in progress), October 2013.

[I-D.ietf-alto-protocol]

Alimi, R., Penno, R., and Y. Yang, "ALTO Protocol", draft-ietf-alto-protocol-25 (work in progress), January 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-01 (work in progress), February 2014.

[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-00 (work in progress), August 2013.

[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-01 (work in progress), October 2013.

[I-D.ietf-idr-ls-distribution]

Gredler, H., Medved, J., Previdi, S., Farrel, A., and S. Ray, "North-Bound Distribution of Link-State and TE Information using BGP", draft-ietf-idr-ls-distribution-04 (work in progress), November 2013.

- [I-D.ji-i2rs-usecases-ccne-service]
Ji, X., Zhuang, S., and T. Huang, "I2RS Use Cases for Control of Forwarding Path by Central Control Network Element (CCNE)", draft-ji-i2rs-usecases-ccne-service-00 (work in progress), October 2013.
- [I-D.keyupate-i2rs-bgp-usecases]
Patel, K., Fernando, R., Gredler, H., and S. Amante, "Use Cases for an Interface to BGP Protocol", draft-keyupate-i2rs-bgp-usecases-00 (work in progress), March 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-01 (work in progress), August 2013.
- [I-D.wu-i2rs-IM-service-topo]
Wu, Q., Hares, S., and X. Guan, "An Information Model for Network Topologies", ID draft-medved-i2rs-topology-im-01, October 2003.
- [I.D.medved-i2rs-topology-im]
Medved, J., Bahadur, N., Clemm, A., and H. Ananthakrishnan, "An Information Model for Network Topologies", ID draft-medved-i2rs-topology-im-01, October 2003.
- [RFC3746] Yang, L., Dantu, R., Anderson, T., and R. Gopal, "Forwarding and Control Element Separation (ForCES) Framework", RFC 3746, April 2004.
- [RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, August 2006.

Authors' Addresses

Susan Hares
Hickory Hill Consulting
7453 Hickory Hill
Saline, MI 48176
USA

Email: shares@endzh.com

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

Email: mach.chen@huawei.com