

Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 5, 2015

M. Chen
S. Hares
Huawei Technologies Co., Ltd
July 4, 2014

I2RS Traffic Steering Use Case
draft-chen-i2rs-ts-use-case-01

Abstract

Traffic steering (TS) mechanisms are designed to improve the overall utilization of network resources in terms of bandwidth and services, and to optimize traffic flow in terms of latency and jitter through the network. Most existing TS systems require human intervention to monitor and manage the assignment of specific paths to specific flows or traffic types, reducing the effectiveness of TS. I2RS, as a near real time programmatic interface to the routing system, provides the ability to manage TS systems in a more dynamic and iterative way. This document describes several TS use cases for I2RS.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

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 January 5, 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. I2RS Requirements from the Traffic Steering Use Cases	3
3. Domain Exit Traffic Steering	4
4. End-to-end Traffic Steering	5
4.1. MPLS-TE	5
4.2. Segment Routing	6
5. IANA Considerations	7
6. Security Considerations	7
7. Acknowledgements	7
8. References	7
8.1. Normative References	7
8.2. Informative References	7
Authors' Addresses	8

1. Introduction

Interface to Routing System (I2RS) architecture [I-D.ietf-i2rs-architecture] defines a standard, programmatic interface to routing system. Through the I2RS interface, an entity external to the routing system can retrieve and program network states of the routing system, hence affecting packet forwarding behaviours.

Well designed Traffic Steering (TS) can improve the usage of network bandwidth resource, reduce the network congestion and satisfy the requirements (e.g., loss and delay) of the applications. Policy Routing (PR), MPLS Traffic Engineering (MPLS-TE) [RFC3209] are useful tools for traffic steering deployed in the field. Segment Routing (SR) leverages the source routing paradigm for packet forwarding, reducing the per flow state in transit devices to provide a scalable TR solution [I-D.filsfils-rtgwg-segment-routing]. A BGP specific use

case similar to the more general use cases considered here is described in section 3.4 of [I-D.keyupate-i2rs-bgp-usecases].

Most of the existing TS solutions need human intervention because they lack dynamic feedback which would facilitate programmatic adjustments to the policies impacting packet flows. This document describes use cases that leverage the I2RS interface to perform traffic steering dynamically and automatically.

2. I2RS Requirements from the Traffic Steering Use Cases

This section contains the list of I2RS requirements found in this draft. Each of these requirements has been given an ID number of TS-REQ-*nn* for ease of reference.

The requirements from the Traffic Steering use cases described in this document are:

TS-REQ01: The I2RS Client-Agent must be able to collect the topology (especially the exit links) and the traffic load of each link;

TS-REQ02: The I2RS Client-Agent must be able to read the local rib of each DC/Metro gateway and the policies deployed on each gateway;

TS-REQ03: The I2RS Client-Agent must be able to add or delete or modify the relevant rib items and relevant polices to steer the traffic as expected; and adjust traffic placement.

TS-REQ-04: The I2RS Client-Agent must have the ability to collect the LSP information either from the PCE or directly from network devices;

TS-REQ-05: 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;

TS-REQ-06: The I2RS Client-Agent must have the ability to read the rib information and relevant policies of each network node;

TS-REQ-07: collect the topology and segment information needed to help the I2RS client to compute the end-to-end path;

TS-REQ-08: read rib (especially the segment routing rib) information;

TS-REQ-09: add/delete/modify the segment rib, this finally determines how the traffic is forwarded.

Note: In each of these sections, the requirements match the technical specificaiton, but maybe listed of ascending numerical order.

3. Domain Exit Traffic Steering

Data Center (DC) fabrics and metro area networks often have more than one exit point. By enforcing relevant policies, it is possible to share the outbound traffic load among these exit points, in turn preventing a single link from becoming overloaded, and thus improving the experience of customers using these facilities. The figure below illustrates a typical network design with multiple exits.




Figure 1: DC Exit Traffic Steering

Router 1 and Router 2 represent the border leaf in a DC scenario, or the provider edge in a metro network. Routers A, B and C represent the fabric interconnect in a data center, or the network core in a metro network. Ideally, load would be shared between the outbound link at Router 1 and the outbound link at Router 2; equably if each link has the same capacity, or in proportion to the capacity of each link if they have different capacities. This requires that router 1 and router 2 know the load of each link so Routers A, B, and C, can dynamically steer traffic towards the correct exit point. Normally the load of each of these links is only available to Routers 1 and 2 locally; steering traffic to the correct exit requires manual monitoring and configuration on the part of the network operator.

I2RS introduces the concept of a feedback loop; the I2RS client can dynamically learn the utilization of the links at Routers 1 and 2, as well as the state of the routing tables at Routers A, B, and C, the topology of the network, and other information. Based on this information, the I2RS client can compute the changes necessary in the network to evenly balance the load between the two exit points, and inject the right information into the appropriate RIBs to make the necessary changes.

Achieving this requires the I2RS Client-Agent to have the following abilities:

- o TS-REQ01: be able to collect the topology (especially the exit links) and the traffic load of each link;
- o TS-REQ02: be able to read the local rib of each DC/Metro gateway and the policies deployed on each gateway;
- o TS-REQ03: be able add or delete or modify the relevant rib items and polices to steer the traffic as expected.

4. End-to-end Traffic Steering

4.1. MPLS-TE

In MPLS-TE, traffic is placed into a Label Switched Path (LSP), guiding that traffic through a specific set of nodes and edges to traverse the network. Each LSP represents a (potentially) different path through the network, allowing different streams to take a different path to reach the same destination device. In placing LSPs, the network operator effectively steers traffic through the network.

LSP computation and optimization can be performed by Path Computation Element (PCE) [RFC4655], which uses the Path Computation Element Communication Protocol (PCEP) [RFC5440] as a southbound interface. Since the Path Computation Client (PCC) is embedded in the network devices, it can actively request path computation or just receive the path establishment direction from the stateful PCE [I-D.ietf-pce-stateful-pce] and then establish the path.

PCE and I2RS architectures contain similar functionalities. In the end-to-end traffic steering scenario, these similar architectures provide complementary functions. Since the PCE is mainly for path computation, traffic placement and steering are out of the scope of PCE, and I2RS can be used in these aspects.

In order to support traffic placement and steering, the I2RS (I2RS client-agent exchange) is required to support:

- o TS-REQ-04: The ability to collect the LSP information either from the PCE or directly from network devices;
- o TS-REQ-05: 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: The ability to read the rib information and relevant policies of each network node;
- o TS-REQ-03: The ability to add/delete/modify rib items and relevant policies hence to adjust traffic placement and steer as desired.

4.2. Segment Routing

Segment Routing (SR) supports end-to-end traffic steering by directly encoding the path and forwarding instruction information into the packet header, steering packets through an explicit route without per-flow states maintained in the transit nodes. This provides a scalable way to perform Traffic Engineering (TE).

In an SR capable network, there are two types of state. One is the segment information that is supplied to the path computation entity for path computation. It can be obtained either from the IGP based Segment advertisement [I-D.psenak-ospf-segment-routing-extensions] [I-D.previdi-isis-segment-routing-extensions], or through the unified BGP interface [draft-chen-idr-segment-distribution]. The other is the SR RIB information that is computed and installed by the path computation entity.

The path computation entity can be either the control plane of the ingress node of a path or an external controller (e.g., PCE or SDN controller). Since this is an I2RS use case, this document mainly talks about the latter (Controller based SR). The controller can be an I2RS client or an application running on an I2RS client. In either case, the I2RS interface should have two capabilities. One is to collect the segment information, the other is able to read and write the SR RIB state.

To support segment routing, the I2RS (client-agent exchange) is required to support the following abilities:

- o TS-REQ-07: collect the topology and segment information needed to help the I2RS client to compute the end-to-end path;
- o TS-REQ-05: collect the network traffic matrix needed to help the I2RS client to compute the optimized path;
- o TS-REQ-08: 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.

5. IANA Considerations

This document makes no request of IANA.

Note to RFC Editor: this section may be removed on publication as an RFC.

6. Security Considerations

This document does not introduce new security issues.

7. Acknowledgements

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.filsfils-rtgwg-segment-routing]

Filsfils, C., Previdi, S., Bashandy, A., Decraene, B., Litkowski, S., Horneffer, M., Milojevic, I., Shakir, R., Ytti, S., Henderickx, W., Tantsura, J., and E. Crabbe, "Segment Routing Architecture", draft-filsfils-rtgwg-segment-routing-01 (work in progress), October 2013.

[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-04 (work in progress), June 2014.

[I-D.ietf-pce-stateful-pce]

Crabbe, E., Minei, I., Medved, J., and R. Varga, "PCEP Extensions for Stateful PCE", draft-ietf-pce-stateful-pce-09 (work in progress), June 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-01 (work in progress), February 2014.

[I-D.previdi-isis-segment-routing-extensions]

Previdi, S., Filsfils, C., Bashandy, A., Gredler, H., Litkowski, S., and J. Tantsura, "IS-IS Extensions for Segment Routing", draft-previdi-isis-segment-routing-extensions-05 (work in progress), February 2014.

[I-D.psenak-ospf-segment-routing-extensions]

Psenak, P., Previdi, S., Filsfils, C., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPF Extensions for Segment Routing", draft-psenak-ospf-segment-routing-extensions-05 (work in progress), June 2014.

[RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.

[RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, August 2006.

[RFC5440] Vasseur, JP. and JL. Le Roux, "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, March 2009.

Authors' Addresses

Mach(Guoyi) Chen
Huawei Technologies Co., Ltd
Q14 Huawei Campus, No. 156 Beiqing Road, Hai-dian District
Beijing 100095
China

Email: mach.chen@huawei.com

Susan Hares
Huawei Technologies Co., Ltd
7453 Hickory Hill
Saline, MI 48176
USA

Email: shares@ndzh.com