

Workgroup: SPRING
Internet-Draft:
draft-hr-spring-intentaware-routing-using-
color-03

Published: 23 October 2023

Intended Status: Informational

Expires: 25 April 2024

Authors: S. Hegde	D. Rao	
Juniper Networks Inc.	Cisco Systems	
J. Uttaro	A. Bogdanov	L. Jalil
Independent Contributor	BT	Verizon

Problem statement for Inter-domain Intent-aware Routing using Color

Abstract

This draft describes the scope, set of use-cases and requirements for a distributed routing based solution to establish end-to-end intent-aware paths spanning multi-domain packet networks. The document focuses on BGP given its predominant use in inter-domain routing deployments, however the requirements may also apply to other solutions.

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 <https://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 25 April 2024.

Copyright Notice

Copyright (c) 2023 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. [Introduction](#)
 - 1.1. [Objectives](#)
2. [Typical large scale network deployment scenarios](#)
 - 2.1. [5G access networks](#)
 - 2.2. [WAN networks for Content distribution](#)
 - 2.3. [Data Center Inter-connect Networks](#)
3. [Use Cases for Inter-domain Intent-based Transport](#)
 - 3.1. [Inter-domain Data Sovereignty](#)
 - 3.2. [Inter-domain Low-Latency Services](#)
 - 3.3. [Inter-domain Service Function Chaining](#)
 - 3.4. [Inter-domain Multicast Use cases](#)
4. [Deployment use cases](#)
 - 4.1. [Network Domains under different administration](#)
5. [Intent-Aware Routing Framework](#)
 - 5.1. [Intent](#)
 - 5.2. [Color](#)
 - 5.3. [Colored Service Route](#)
 - 5.4. [Intent-Aware Route using Color](#)
 - 5.5. [Service Route Automated Steering on intent-aware route using color](#)
 - 5.6. [Inter-Domain intent-aware routing using colors with SR Policy](#)
 - 5.7. [Motivation for a BGP-based intent-aware routing solution using colors](#)
 - 5.8. [BGP Intent-Aware Routing using Color](#)
 - 5.9. [Architectural consistency among intent-aware routing solutions using colors](#)
6. [Technical Requirements](#)
 - 6.1. [Intent Requirements](#)
 - 6.1.1. [Transport Network Intent Requirements](#)
 - 6.1.2. [VPN \(Service Layer\) Network Intent Requirements](#)
 - 6.1.3. [Multicast Intent Requirements](#)
 - 6.2. [Traffic Steering Requirements](#)
 - 6.3. [Deployment Requirements](#)
 - 6.3.1. [Multi-domain deployment designs](#)
 - 6.3.2. [Scalability Requirements](#)
 - 6.3.3. [Network Availability Requirements](#)

- [6.3.4. BGP Protocol Requirements](#)
- [6.3.5. OAM Requirements](#)
- [7. Backward Compatibility](#)
- [8. Security Considerations](#)
- [9. IANA Considerations](#)
- [10. Acknowledgements](#)
- [11. Co-authors](#)
- [12. Contributors](#)
- [13. References](#)
 - [13.1. Normative References](#)
 - [13.2. Informative References](#)
- [Authors' Addresses](#)

1. Introduction

Evolving trends in wireless access technology, cloud applications, virtualization, and network consolidation all contribute to the increasing demands being placed on a common packet network. In order to meet these demands, a given network will need to scale horizontally in terms of its bandwidth, absolute number of nodes, and geographical extent. The same network will need to extend vertically in terms of the different services and variety of intent that it needs to simultaneously support.

In order to operate networks with large numbers of devices, network operators organize networks into multiple smaller network domains. Each network domain typically runs an IGP which has complete visibility within its own domain, but limited visibility outside of its domain. Network operators will continue to use multiple domains to scale horizontally. In MPLS based networks BGP-LU (RFC8277) has been widely deployed for providing reachability across multiple domains.

The evolving network requirements (e.g. 5G, native cloud) in such a multi-domain network requires the establishment of paths that span multiple domains or AS's while maintaining specific transport characteristics or intent (e.g. bandwidth, latency). There is also a need to provide flexible, scalable, and reliable end-to-end connectivity for multiple services across the network domains.

1.1. Objectives

This document describes requirements for scalable, intent-aware reachability across multiple domains.

The base problem that it focuses on is the BGP-based delivery of an intent across several transport domains, however the requirements may also apply to other distributed solutions.

The problem space is then widened to include any intent (including Network Function Virtualization (NFV) chains and their location), any data plane and the application of intent-based routing to the Service/VPN routes.

It is intended that the requirements enable the design of technology and protocol extensions that address the widest application, while ensuring consistency and compatibility with existing deployed solutions.

2. Typical large scale network deployment scenarios

This section describes a few typical deployment scenarios that involve large-scale multi-domain network designs and use of various topology, IGP and BGP routing models. While the examples use specific types of deployments for illustration, neither the use-cases nor the network designs are limited to any particular provider deployment.

2.1. 5G access networks

Service Provider networks can contain many nodes distributed over a large geographic area. 5G networks can include as many as one million nodes, with the majority of those being radio access nodes. Radio and access nodes may be constrained by their memory and processing capabilities.

Such transport networks use multiple domains to support scalability. For this analysis, we consider a representative network design with four level of hierarchy: access domains, pre-aggregation domains, aggregation domains and a core. (See [Figure 1](#)). The separation of domains internal to the service provider can be performed by using either IGP or BGP.

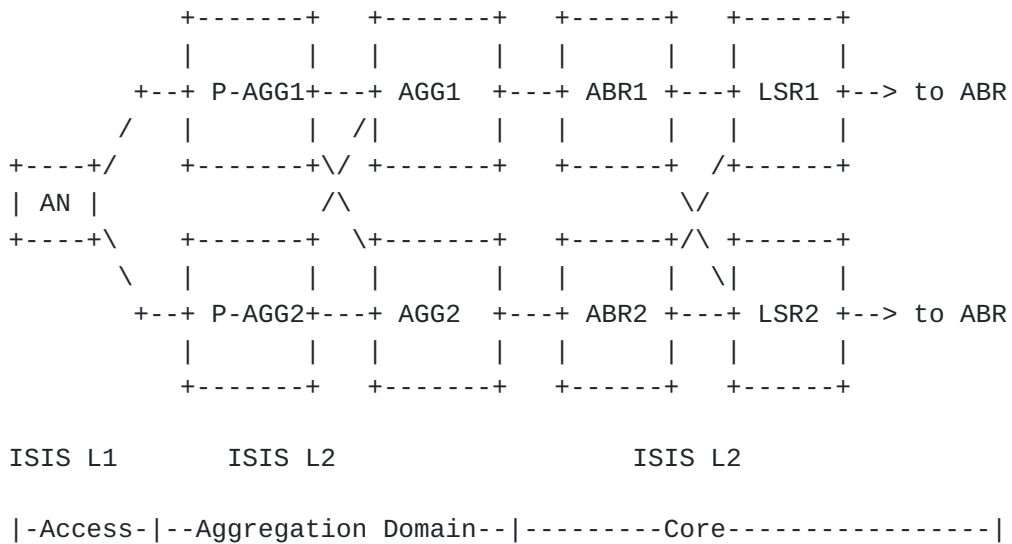


Figure 1: 5G network

5G networks support a variety of service use cases that may require end-to-end network slicing. In certain cases, the end-to-end connectivity requires the ability to forward over intent-aware paths, such as paths delivering low-delay. The inter-domain routing solution should support the establishment of end-to-end paths that address specific intent requirements, as well as support multiple such paths to address slicing requirements.

2.2. WAN networks for Content distribution

Networks built for providing delivery of content are geographically distributed by design to provide connectivity in multiple regions and sharing of data across regions.

As these WAN networks grow beyond several thousand nodes, they are divided into multiple IGP domains for scale and reliability. An illustration is provided in in [Figure 2](#).




Figure 2: Content distribution WAN Example

These large WAN networks often cross national boundaries. In order to meet data sovereignty requirements, operators need to maintain strict control over end-to-end traffic-engineered (TE) paths. A distributed inter-domain solution should be able to create highly constrained inter domain TE paths in a scalable manner.

Some deployments may use a controller to acquire the topologies of multiple domains and build end-to-end constrained paths. This approach can be scaled with hierarchical controllers. However, there is still a risk of a loss of network connectivity to one or more controllers, which could lead to a failure to satisfy the strict requirements of data sovereignty. The network should be able to have pre-established TE paths end-to-end that don't rely on controllers, to address these failure scenarios.

2.3. Data Center Inter-connect Networks

Distributed data centers are playing an increasingly important role in providing access to information and applications. Geographically diverse data centers are usually connected via a high speed, reliable and secure DC WAN core network.

One variation of a DCI topology is shown in [.Figure 3.](#)

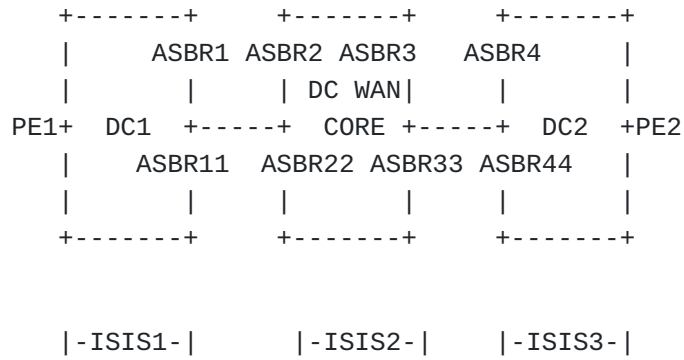


Figure 3: DCI Network

In many DC WAN deployments, applications require end-to-end path diversity and end-to-end low latency paths.

Another consideration in DC WAN deployments is the choice of encapsulation technologies. Some deployments use the same tunneling mechanism within the DC and DCI networks, while other deployments use different mechanisms in each. It is important for a solution to provide flexibility in choice of tunneling mechanisms across domains.

3. Use Cases for Inter-domain Intent-based Transport

The use cases for inter-domain intent-based packet transport described in this section are intended to provide motivation for the requirements that follow. They apply to all the different deployment scenarios described above.

3.1. Inter-domain Data Sovereignty

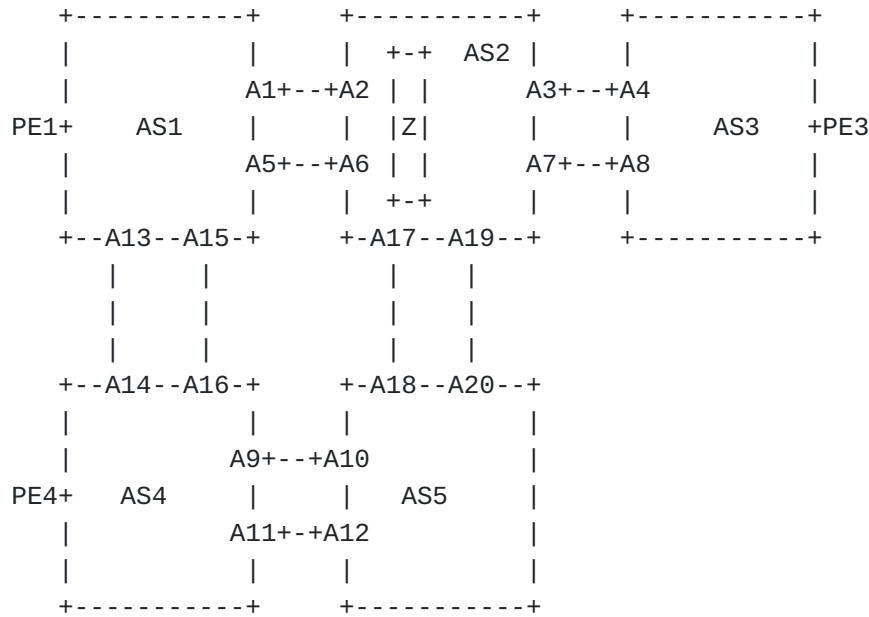


Figure 4: Multi domain Network

Figure 4 depicts an example of a WAN with multiple ASes, where each AS serves a continent. Certain traffic from PE1 (in AS1) to PE3 (in AS3) must not traverse country Z in AS2. However, all paths from AS1 to AS3 traverse AS 2. The inter-domain solution should provide end-to-end path creation that traverses AS 2 but avoids country Z.

In other networks, the domain to avoid may encompass an entire AS.

3.2. Inter-domain Low-Latency Services

Service provider networks running L2 and L3VPNs carry traffic for particular VPNs on low-latency paths that traverse multiple domains.

3.3. Inter-domain Service Function Chaining

RFC7665 defines service function chaining as an ordered set of service functions and automated steering of traffic through this set of service functions. There could be a variety of service functions such as firewalls, parental control, CGNAT etc. In 5G networks these functions may be completely virtualized or could be a mix of virtualized functions and physical appliances. It is required that the inter-domain solution caters to the service function chaining requirements. The service functions may be virtualized and spread across different data centers attached to different domains.

3.4. Inter-domain Multicast Use cases

Multicast services such as IPTV and multicast VPN also need to be supported across a multi-domain service provider network.

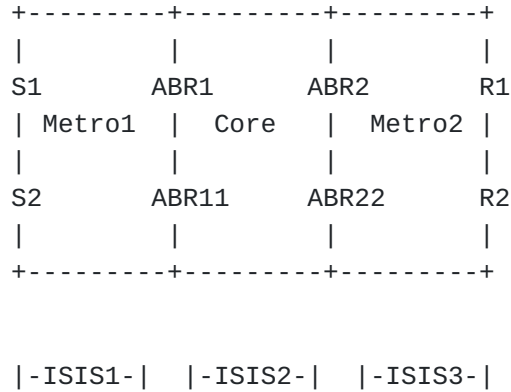


Figure 5: Multicast use cases

[Figure 5](#) shows a simplified multi-domain network supporting multicast. Multicast sources S1 and S2 are in a different domain from the receivers R1 and R2. The solution should support establishment of intent-aware multicast distribution trees (P tunnels) across the domains and steer customer multicast streams on it. It should maintain the scaling properties of a multi-domain architecture by avoiding leaking of RPF routing state into the IGP domains.

4. Deployment use cases

4.1. Network Domains under different administration




Figure 6: Networks with inconsistent intent mappings

In diagram Figure 5 above, AS1 and AS2 may be operating as closely coordinated but independent administrative domains, and still

require end-to-end paths across the two ASes to deliver services. This scenario could be a result of a merger. It is possible that AS1 and AS2 may have assigned different values for the same intent.

In some cases, organizations may continue to use option A or option B [RFC4364] style interconnectivity in which case the inter-domain solution should satisfy intent of the path on inter-domain links for the service prefixes. In other cases, organizations may prefer to use option C style connectivity from PE1 to PE2.

An inter-domain solution should provide effective mechanisms to translate intent across domains without requiring renumbering of the intent mapping.

5. Intent-Aware Routing Framework

This section describes the basic concepts, terminologies and architectural principles that define intent-aware routing and the protocols and technologies that currently support it. The goal of this section is to establish the requirement for consistency with existing deployed solutions and describe the framework for it.

The figure below is used as reference.



Figure 7: Intent-aware routing using color reference topology

5.1. Intent

Intent in routing may be any combination of the following behaviors:

- *Topology path selection (e.g. minimize metric, avoid resource)
- *NFV service insertion (e.g. service chain steering)
- *Per-hop behavior (e.g. QoS for 5G slice)

An intent-aware routed path may be within a single network domain or across multiple domains.

5.2. Color

Color is a 32-bit numerical value that is associated with an intent, as defined in [RFC9256]

5.3. Colored Service Route

An Egress PE E2 colors a BGP service (e.g. VPN) route V/v to indicate the particular intent that E2 requests for the traffic bound to V/v. The color (C) is encoded as a BGP Color Extended community [RFC9012].

5.4. Intent-Aware Route using Color

(C, E2) represents a intent-aware route to E2 which satisfies the intent associated with color C.

Multiple technologies already provide intent-aware paths in solutions that are widely deployed.

*SR Policy [RFC9256]

*IGP Flex-Algo [RFC9350]

In the context of large-scale SR-MPLS networks, SR Policy is applicable to both intra-domain and inter-domain deployments; whereas IGP Flex-Algo is better suited to intra-domain scenarios.

5.5. Service Route Automated Steering on intent-aware route using color

An ingress PE E1 automatically steers V-destined packets onto a intent-aware path bound to (C, E2). If several such paths exist, a preference scheme is used to select the best path: E.g. IGP Flex-Algo first, then SR Policy.

5.6. Inter-Domain intent-aware routing using colors with SR Policy

If E1 and E2 are in different domains, E1 may request an SR-PCE in its domain for a path to (C, E2). The SR-PCE (or a set of them) computes the end-to-end path and installs it at E1 as an SR Policy. The end-to-end intent-aware path may seamlessly cross multiple domains.

5.7. Motivation for a BGP-based intent-aware routing solution using colors

While the following requirements may be covered with an SR Policy solution, an operator may prefer a BGP-based solution due to:

- *Operational familiarity and expectation of incremental evolution from an existing Seamless-MPLS/BGP-LU inter-domain deployment [[I-D.ietf-mpls-seamless-mpls](#)]
- *Expectation of higher scale with BGP
- *Expectation of a familiar operational trust model between BGP domains (peering policy)

5.8. BGP Intent-Aware Routing using Color

A BGP Intent-Aware Routing solution signals intent-aware routes to reach a given destination (e.g. E2). (C, E2) represents a BGP hop-by-hop distributed route that builds an inter-domain intent-aware path to E2 for color C.

5.9. Architectural consistency among intent-aware routing solutions using colors

As seen above, multiple technologies exist that provide intent aware routing in a network. A BGP based solution must be compliant with the existing principles that apply to them.

A deployment model that provides consistency is as follows:

- *Service routes are colored using BGP Color Extended-Community to request intent [RFC9256]

- V/v via E, colored with C

- *Colored service routes are automatically steered on an appropriate intent-aware path using color

- V/v via E with C is steered via (E, C)

- (E, C) provided by any intent-aware technology or protocol

- *Intent-aware routes may resolve recursively via other intent-aware routes

- (E, C) via N recursively resolves via (N, C)

Here is a brief example that illustrates these principles.

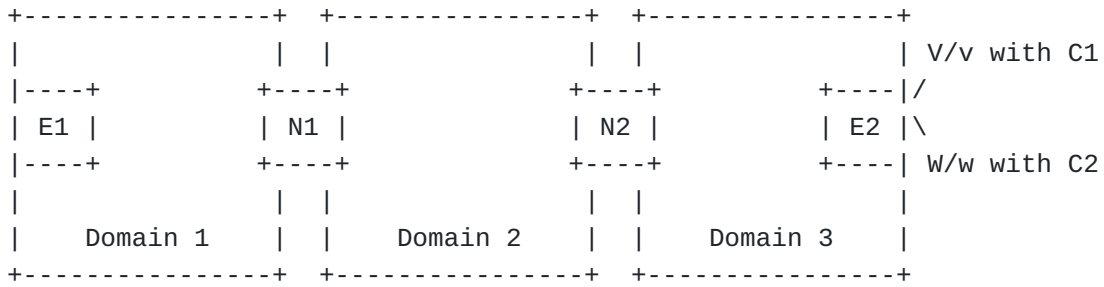


Figure 8: Inter-domain intent-aware routing using color reference topology

In the figure above, all the nodes are part of an inter-domain network under a single authority and with a consistent color-to-intent mapping:

*Color C1 is mapped to "low delay"

-Flex-Algo FA1 is mapped to "low delay" and hence to C1 in each domain

*Color C2 is mapped to "low delay and avoid resource R"

-Flex-Algo FA2 is mapped to "low delay and avoid resource R" and hence to C2 in each domain

E1 receives two BGP colored service routes from E2:

-V/v with BGP Color Extended community C1

-W/w with BGP Color Extended community C2

E1 has the following inter-domain intent-aware paths using color:

*(E2, C1) provided by BGP which recursively resolves via intra-domain intent-aware paths:

-(N1, C1) provided by IGP FA1 in Domain1

-(N2, C1) provided by SR Policy bound to color C1 in Domain2

*(E2, C2) provided by SR Policy

E1 automatically steers the received colored service routes as follows:

- V/v via (E2, C1) provided by BGP intent-aware route using color

- W/w via (E2, C2) provided by SR Policy

The example illustrates the benefits provided by leveraging the architectural principles:

- *Seamless co-existence of multiple intent-aware technologies, e.g. BGP and SR Policy

- V/v is steered on BGP intent-aware path

- W/w is steered on SR Policy intent-aware path

- *Seamless and complementary interworking between different intent-aware technologies

- V/v is steered on a BGP intent-aware path that is itself resolved within domain 2 onto an SR Policy bound to the color of V/v

- *Another benefit that can be extrapolated from the example is that intent-aware routes from different technologies may serve as alternative paths for the same intent.

6. Technical Requirements

6.1. Intent Requirements

The BGP Intent-Aware routing solution must support the following intents bound to a color:

- *Minimization of a cost metric vs a latency metric

- Minimization of different metric types, static and dynamic

- *Exclusion/Inclusion of SRLG and/or Link Affinity and/or minimum MTU/number of hops

- *Bandwidth management

- *In the inter-domain context, exclusion/inclusion of entire domains, and border routers

*Inclusion of one or several virtual network function chains

-Located in a regional domain and/or core domain, in a DC

*Localization of the virtual network function chains

-Some functions may be desired in the regional DC or vice versa

Subsequent sections elaborate on these requirements.

6.1.1. Transport Network Intent Requirements

The requirements described in this document are mostly applicable to network under a single administrative domain that are organized into multiple network domains. The requirements are also applicable to multi-AS networks with closely cooperating administration.

The network diagram below illustrates the reference network topology used in this section

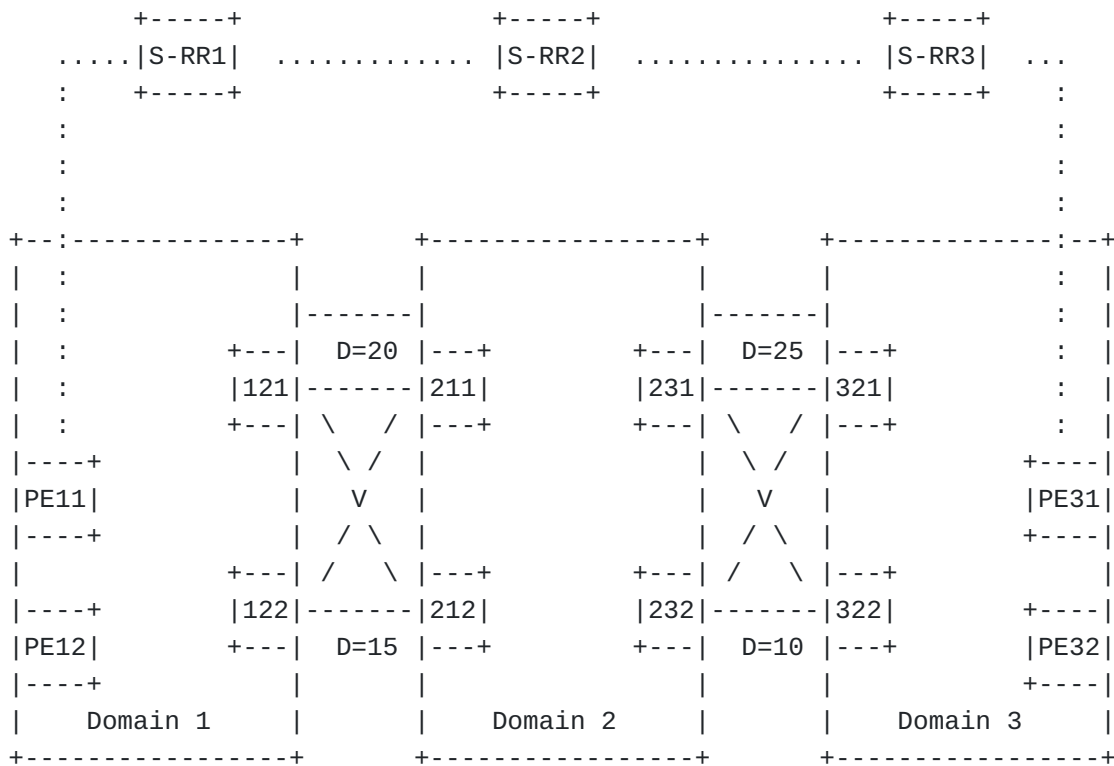


Figure 9: Transport Network Intent Requirements Reference Diagram

The following network design assumptions apply to the reference topology above, as an example:

- *Independent ISIS/OSPF SR instance in each domain.
- *eBGP peering link between ASBRs (121-211, 121-212, 122-211, 122-212, 231-321, 231-322, 232-321 and 232-322).
- *Peering links have equal cost metric.
- *Peering links have delay configured or measured as shown by "D".
D=50 for cross peering links.
- *The cross links between ASBRs share the same risk.
- *The top parallel link between 121-211 shares same risk with the link 122-212.
- *The top parallel link between 231-321 shares same risk with the link 232-322.
- *VPN service is running from PE31, PE32 to PE11, PE12 via service RRs (S-RRn in figure).

Intent-aware inter-domain routing information to end point E with intent C is represented using (C,E). The notation used is a representation of the intent-aware route using color, and does not indicate a specific protocol encoding.

The following sections illustrate requirements and provide detailed examples for several intent types.

6.1.1.1. Minimization of end-to-end metric

Various metric types can be advertised within an IGP domain and minimum metric paths can be computed within IGP domain, with Flex-Algo [RFC9350] for instance.

The BGP solution should allow the establishment of inter-domain intent-aware paths with low values of a metric type, accumulated over the end-to-end path.

In the reference topology of [Figure 9](#)

- Each domain has Algo 0 and Flex Algo 128
- Algo 0 is for minimum cost metric(cost optimized).

-Flex Algo 128 definition is for minimum delay (low latency).

*Cost Optimized end-to-end path

-Color C1 - Minimum cost intent.

-Intent-aware route for C1 sets up path(s) between PEs for end-to-end minimum cost.

-These paths traverse over intra-domain Algo 0 in each domain and account for the peering link cost between ASBRs.

-Example: PE11 learns (C1, PE31) intent-aware route via several equal paths:

- oOne such path is through FA0 to node 121, links 121-211, FA0 to 231, link 231-321, FA0 to PE31

- oAnother such path is through FA0 to node 122, link 122-212, FA0 to 232, link 232-322, FA0 to PE31.

- oPE11 may load-balance among these paths

-On PE11, VPN routes from PE31 colored with C1 are steered via (C1, PE31) intent-aware route.

*Latency Optimized End-to-end path

-Color C2 - Minimum latency intent.

-BGP Intent-aware route for C2 advertises path(s) between PEs for end-to-end minimum delay.

-These paths traverse over intra-domain Flex-Algo 128 in each domain and account for the peering link delay between ASBRs.

-Example: PE11 learns (C2, PE31) intent-aware route and best path is through FA128 to node 122, link 122-212, FA128 to 232, link 232-322, FA128 to PE31.

-On PE11, VPN routes from PE31 colored with C2 are steered via (C2, PE31) intent-aware route.

6.1.1.2. Exclusion/inclusion of link affinity

The Intent-aware BGP routing solution should allow the establishment of inter-domain paths that satisfy link affinity inclusion/exclusion constraints. The link affinity constraints should also be satisfied for inter-domain links, such as those between ASBRs.

Using the reference topology of Figure 7 for the example below:

*Color C3 - Intent to Minimize cost metric and avoid purple links

*Each domain has Flex Algo 129 and some links have purple affinity.

*Flex Algo 129 definition is set to minimum cost metric and avoid purple links (within domain).

*Peering cross links are colored purple by policy.

*BGP intent-aware route for C3 sets up paths between PEs for minimum end-to-end cost and avoiding purple link affinity.

*These paths traverse over intra domain Flex Algo 129 in each domain and accounts for peering link cost between ASBR and avoiding purple links.

*Example: PE11 learns (C3, PE31) intent-aware route via 2 paths.

-First path is through FA 129 to node 121, link 121-211, FA129 to 231, link 231-321, FA129 to PE31.

-Second path is through FA129 to node 122, link 122-212, FA129 to 232, link 232-322, FA129 to PE31.

*On PE11, VPN routes from PE31 colored with C3 are steered via (C3, PE31) intent-aware route.

6.1.1.3. Exclusion/inclusion of nodes

Support creating an inter-domain path that includes or excludes a certain set of nodes in each domain.

Mechanisms used to achieve the node inclusion/exclusion constraints within different domains should be independent.

For example, an RSVP-based domain may use link affinities to achieve node exclusion constraints, while an SR-based domain may use Flex-Algo, which natively supports excluding nodes.

The example below describes the details for [Figure 9](#)

*Color C4 - Intent to Minimize cost metric and avoid nodes

-Each domain has Flex Algo 129 and Flex-Algo 129 is not enabled on nodes 121,211,231,321

-Flex Algo 129 definition is set to minimum cost metric

*Intent-aware route for C4 sets up paths between PEs for minimum end-to-end cost and avoiding specific nodes.

*These paths traverse over intra domain Flex Algo 129 in each domain and accounts for peering link cost between ASBR and avoiding specific nodes.

*Example: PE11 learns (C4, PE31) intent-aware route via 1 path.

-The path is through FA129 to node 122, link 122-212, FA129 to 232, link 232-322, FA129 to PE31.

*On PE11, VPN routes colored with C4 are steered via (C4, PE31) intent-aware route.

6.1.1.4. Diverse Paths

Support the creation of node- and link-diverse inter-domain paths.

The intra-domain portion of the end-to-end paths should make use of existing mechanisms for computing and instantiating diverse paths within a domain.

Inter-domain links (such as those connecting ASBRs) should also be taken into account for diverse inter-domain paths.

Support creation of inter-domain diverse paths that avoid shared risk links.

The example below describes the details for [Figure 8](#)

*Color C5 and C6 - Intent to create diverse paths avoiding common node, link and shared risk

-Each domain has SRLG aware diverse path built as below

-Domain 1: Color C5 -> PE11,121

-Color C6 -> PE12,122

-Domain 2: Color C5 -> 211,231

-Color C6 -> 212,232

-Domain 3: Color C5 -> 321,PE31

-Color C6 -> 322,PE32

-Shared risk among inter-domain links is as described in the topology description

- oIntent-aware diverse paths represented by C5 and C6 setup in each domain

- oLocal policies on inter-domain links to avoid common shared risk for intent C5 and C6

- oExample: PE11 learns (C5, PE31) intent-aware route via 1 path.

-The path is through PE11,121-211 (bottom link), 231-321 (bottom link), PE31

- oExample: PE12 learns (C6, PE32) intent-aware route via 1 path.

-The path is through PE12,122,212, 232,322, PE32

*On PE11, VPN routes colored with C5 are steered via (C5, PE31) Intent-aware route.

*On PE12, VPN routes colored with C6 are steered via (C6, PE32) intent-aware route.

6.1.1.5. Applicability of intent to a subset of domains

Support creation of paths with certain intents applicable to only a subset of domains.

No constraint specific state on internal nodes where intent is not applicable.

The example below describes the details for [Figure 9](#)

*Color C7 to exclude purple links

- Purple links exist only in domain 2

- Intra-domain Intent-aware paths in domain 2 via 211,231

- Intra-domain paths for C7 not created in Domain 1 and Domain 3

*On PE11, VPN routes colored with C7 are steered via (C7, PE31) intent-aware route.

- Intent-aware route (C7,PE31) uses best effort paths in Domain1 and Domain3

-Intent-aware route (C7,PE31) uses intra-domain intent-aware path C7 in Domain2

6.1.1.6. Exclusion/inclusion of domain

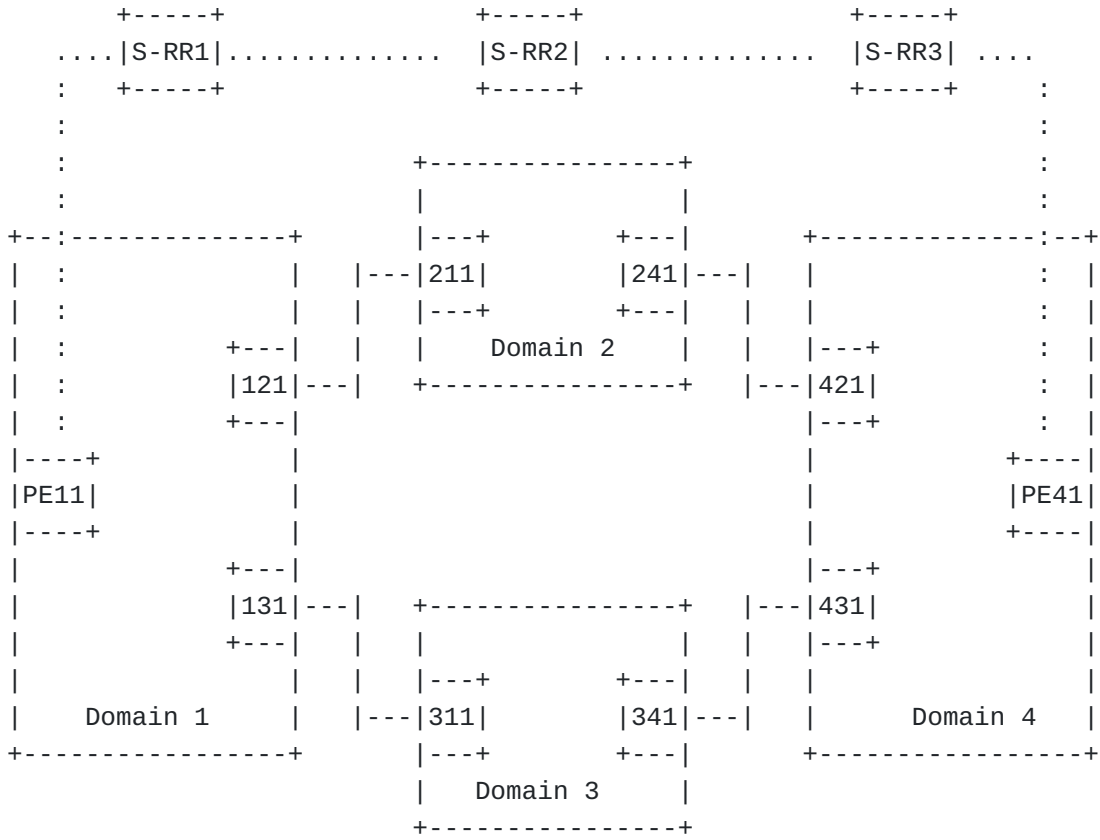


Figure 10: Domain Exclusion Diagram

Color C4 - Avoid sending selected traffic via Domain 3

*VPN routes advertised from PEs with Color C4

*Intent-aware route for Color C4 should only set up paths between PE11 and PE41 that exclude Domain 3

6.1.1.7. Virtual network function chains in local and core domains

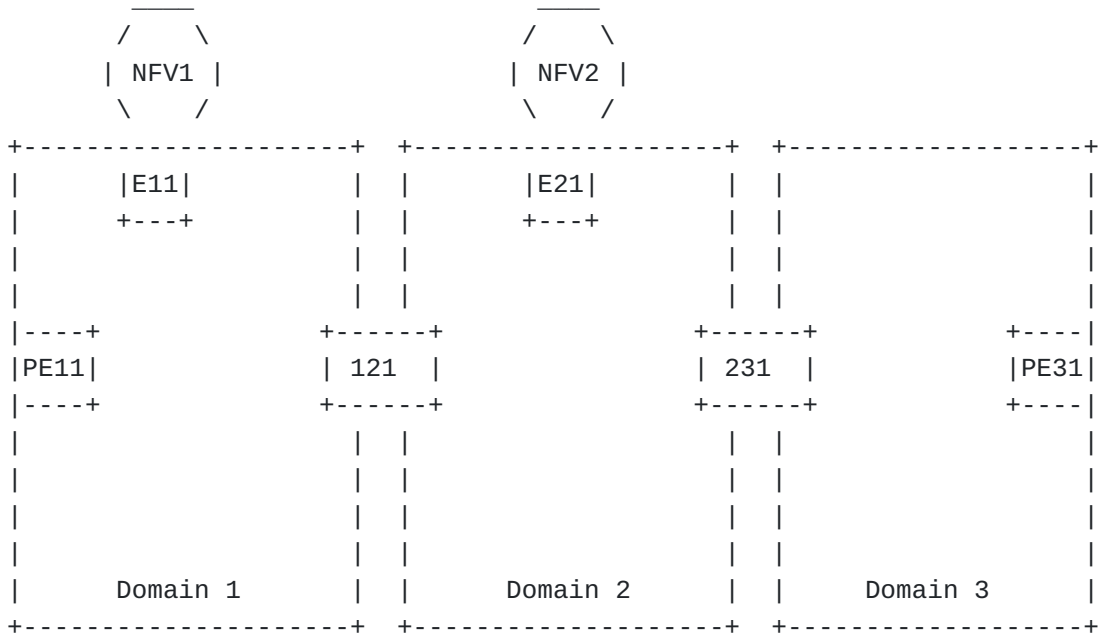


Figure 11: Transport NFV Diagram

*Color intent

- C5 - Routing via min-cost paths
- C6 - Routing via a local NFV service chain situated at E11
- C7 - Routing via a centrally located NFV service chain situated at E21

*Forwarding of packets from PE11 towards PE31:

- (C5, PE31) mapped packets are sent via nodes 121, 231 to PE31
- (C6, PE31) mapped packets are sent to E11 and then post-service chain, via 121, 231 to PE31
- (C7, PE31) mapped packets are sent via 121 to E21 and then post-service chain, via 231 to PE31

E11 and E21 MAY be involved in inter-domain signaling in order to send service traffic towards PEs in remote domains. Different functions may be collocated at the same network node. (For example, PE functionality and NFV attachment functionality may be collocated.)

6.1.2. VPN (Service Layer) Network Intent Requirements

This section describes requirements and reference use-cases for extending intent-aware routing to the VPN (Service) layer.

The solution should:

*Extend the signaling of intent awareness end-to-end to the customer domain: CE site to CE site across provider networks. Specific goals are to:

- Provide ability for a CE to select paths through specific PEs for a given intent

- oExample-1: Certain intent in transport not available via specific PEs

- oExample-2: Certain CE-PE connection does not support specific intent

- oExample-3: Customer Site access via certain CE node does not support specific intent. For instance, link connecting a specific CE to a DC hosting loss-sensitive service may have better quality than a link from another CE

- Provide ability for a CE node to send traffic indicating a specific intent (via suitable encapsulation) to the PE for optimal steering.

- oProvide ability for a PE node to apply filtering and other security mechanisms and authentication for the incoming encapsulated packets

- oProvide ability for a PE node to apply traffic policing and shaping mechanisms to the received encapsulated packets.

- oThe PE-CE link and the transport domains can be in different color domains.

- Intent-aware routing may extend further into customer networks towards the network edge, closer to applications that originate traffic

- oApplications or hosts that do not participate in routing can indicate the intent desired for an emitted traffic flow by setting appropriate DSCP values in the packet header. This enables the ingress intent-aware routing device to perform the necessary classification and traffic steering (Section 6.2)

*Support intent aware routing for multiple service (VPN) interworking models

-IBGP and Inter-AS Option C models are inherently supported since they natively extend from PE to PE. Additional models to be supported are:

oInter-AS Option A

oInter-AS Option B

oGW based interworking (L3VPN, EVPN)

-Co-existence with legacy PEs and CEs in a L3VPN network

oIntent-aware routing capable PEs co-exist with other PEs that are not capable

oIntent-aware routing capable PEs simultaneously interact with both capable CEs and legacy CEs

The network diagram below illustrates the reference network topology used in this section for VPN Intent-aware routing using Color

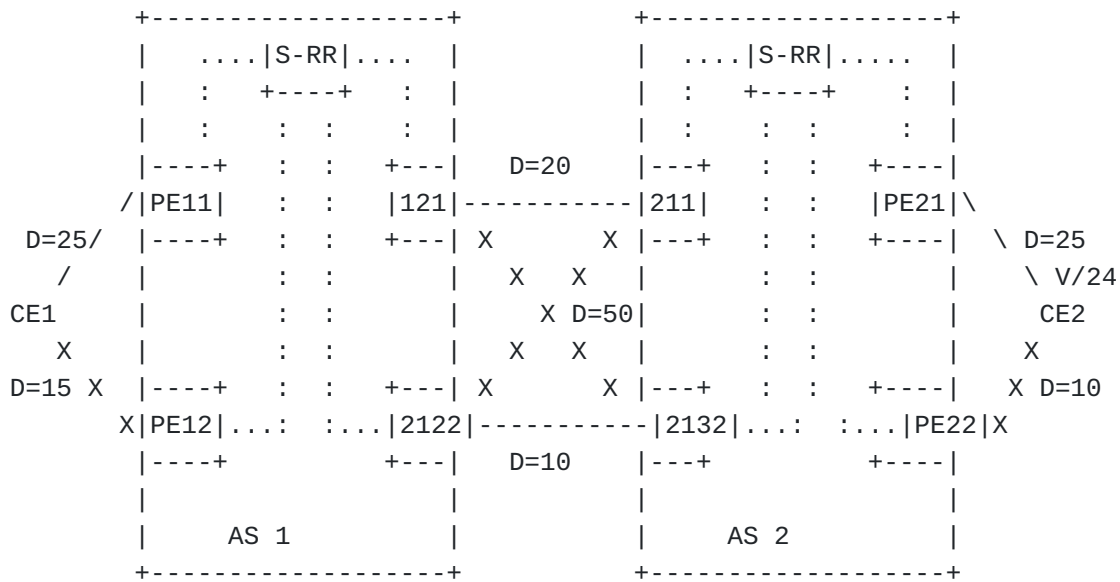


Figure 12: VPN (Service) intent routing reference topology

The following network design assumptions apply to the reference topology above, as an example:

- *eBGP peering link between VPN ASBRs 121-211, 121-212, 122-211, 122-212

- *VPN service runs between PEs in each AS via service RRs to local VPN ASBRs. Between ASBRs, its VPN IAS-Option-B i.e. next hop self.

- *CE1 is dual homed to PE11,PE12. CE2 is dual homed to PE21, PE22.

- *Peering links have equal cost metric

- *Peering links have delay configured or measured as shown by "D"

The following sections illustrate a few examples of intent use-cases applicable to VPN routes.

6.1.2.1. Minimization of end-to-end metric

This use-case extends the transport use-case from Minimization of end-to-end metric section to further to establish e2e paths with low values of a metric type between CEs attached to different PEs, additionally taking the metrics on the PE-CE links and inter-ASBR links into account.

- *In the reference topology of VPN service intent topology, each AS has Flex Algo 0 and 128. Flex Algo 0 is for minimumcost metric (cost optimized) while Flex Algo 128 definition is for minimum delay (low latency)

- *Cost Optimized end-to-end (CE-CE) path

 - Color C1 - Minimum cost intent.

 - On CE1, flows requiring cost optimized paths to V/24 are steered over (C1, V/24) intent-aware route using color.

 - oThis needs BGP intent-aware route between PE-CE for V/24 prefix and color C1 awareness.

 - oIt also needs BGP VPN Intent-aware route between PEs and ASBRs for V/24 prefix with VPN RD and color C1 awareness (C1, RD:V/24)

oCE1 may learn (C1, V/24) route through several equal cost paths. For example:

oOne path is through link CE1-PE11, FA0 to 121, link 121-211, FA0 to PE21 and link PE21-CE2.

oAnother such path is through CE1-PE12, FA0 to node 122, link 122-212, FA0 to PE22, link PE22-CE2.

oCE1 may load-balance among these paths

*Latency optimized end-to-end (CE-CE) path

-Color C2 - Minimum latency intent

-On CE1, flows requiring low latency paths to prefix V/24 are steered over (C2, V/24) intent-aware route using color.

oThis needs BGP intent-aware route between PE-CE for V/24 prefix and color C2 awareness.

oIt also needs BGP VPN intent-aware route between PEs and ASBR for V/24 prefix with VPN RD and color C2 awareness

oPaths traverse over intra-domain Flex Algo 128 in each AS and accounts for inter ASBR link delays and PE-CE link delays.

oCE1 learns (C2, V/24) BGP intent-aware best route using color through link CE1-PE12, FA128 to 122, link 122-212, FA128 to PE22 and link PE22-CE2 between PE-CE for V/24 prefix and color C2 awareness.

6.1.2.1.1. Exclusion/inclusion of link affinity

*Color C3 - Intent to minimize cost metric and avoid purple links

*In the reference topology of Figure 6 Each AS has Flex Algo 129 and some links have purple affinity. Flex Algo 129 definition is set to minimum cost metric and avoid purple links (within AS). ASBR cross links are colored purple by policy. Bottom PE-CE links are colored purple as well by policy

*On CE1, flows requiring minimum cost path avoiding purple links to V/24 are steered over (C3, V/24) BGP intent-aware route using color

CE1 learns (C3, V/24) route through link CE1-PE11, FA129 to 121, link 121-211, FA129 to PE21 and link PE21-CE2.

*From PE11:

- [V1/24, C1] mapped packets are sent via S1, S2 and then routed to PE31, CE2
- [V2/24, C2] mapped packets are sent via S3 and then routed to PE32, CE3

6.1.3. Multicast Intent Requirements

This section describes the intent requirements for the multicast traffic. In principle all the intent requirements described in section [Section 6.1.1](#) can apply to multicast traffic as well. The intent requirements as currently seen from actual usecases have been listed here.

- *Ability to create multicast distribution trees that minimize latency metric
- *Ability to create multicast distribution trees that avoid nodes located in certain geographical region
- *Ability to create multicast distribution trees that provide bandwidth guarantees within as well as across domains
- *Ability to create multicast distribution trees that use subset of the topology

6.2. Traffic Steering Requirements

Traffic arriving at an ingress PE for a colored service route gets steered into an intent-aware path to the egress PE. [Section 5](#) illustrates the automated steering mechanism, driven through Color Extended Community in the service route.

- *Flexible traffic steering is required, with support for different types:
 - Per-Destination Steering: Incoming packets are steered based on the destination address of the packets
 - Per-Flow Steering: Incoming packets are steered based on the destination address of the packets and additional fields in the packet header
 - oDSCP for IPv4/IPv6 packets and EXP for MPLS packets
 - o5-tuple IP flow (Source address, destination address, source port, destination port and protocol fields).

-The Per-Flow Steering enables different flows for the same destination to be steered into different paths – for example, one flow into an intent-aware path and another into a best-effort path; or two different flows steered into paths of two different intents. Section 8.6 of RFC 9256 describes the operation of per-flow steering in detail.

*When no path that fulfills the desired intent is available:

-An option of ordered fallback should be supported

o via one or more alternative intents; or via a best-effort path.

-An option of not using a fallback path for the service route should also be supported.

-Fallback scheme per service route should be supported

o Fallback schemes should be decoupled from primary. For example, different service routes using same primary but different fallback schemes

-An indication that the route followed a less preferred path due to fallback may be given to a CE by modifying/adding suitable BGP attribute through policy.

*Above steering mechanisms should be supported for any service, including L2/L3 VPNs and Internet/global routing.

6.3. Deployment Requirements

The solution must support the representative deployment designs and associated deployment requirements described in the following sub sections.

6.3.1. Multi-domain deployment designs

This section describes four different ways that multi-domain networks could be organized. This is a representation of most common deployments and not an exhaustive coverage.

6.3.1.1. Multiple IGP domains within a single AS, inter-connected at border nodes




Figure 14: Transport Multiple Domains Network Diagram

The above diagram shows three different IGP domains, Domain1, Domain2 and Domain3 inter-connected at the ABRs 121,122,231,232.

This single-AS network uses I-BGP sessions, with ABRs acting as inline route reflectors to PEs.

Note that the IGP design included here and in other models below is illustrative. In practice, there may be multiple areas/levels or multiple IGP instances.

6.3.1.2. Multiple IGP domains within a single AS, with iBGP between border nodes

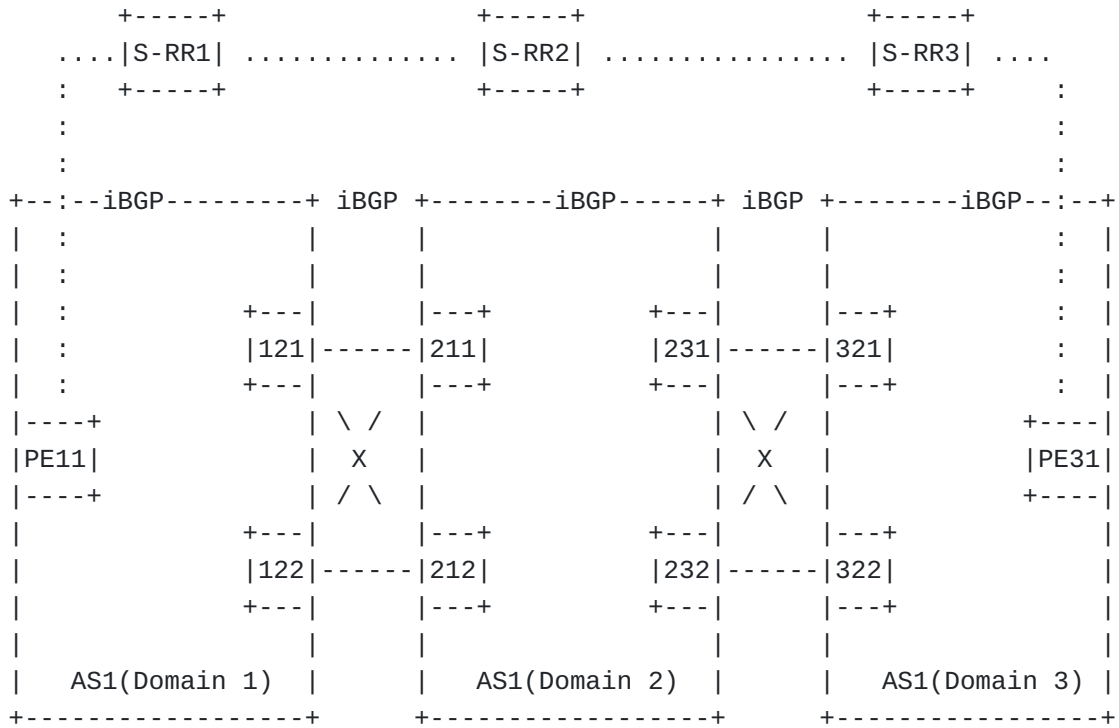


Figure 15: Transport Multiple Domains with iBGP Network Diagram

The above diagram shows a single AS1 with three different IGP domains, Domain1, Domain2, and Domain3.

121,122,211,212,231,232,321,322 are border nodes for the IGP domains and they participate in only one IGP domain.

In this design, domain inter-connect is via iBGP peering links between Area border nodes. ABRs act as inline route reflectors to PEs.

6.3.1.3. Multiple ASes inter-connected with E-BGP between border nodes

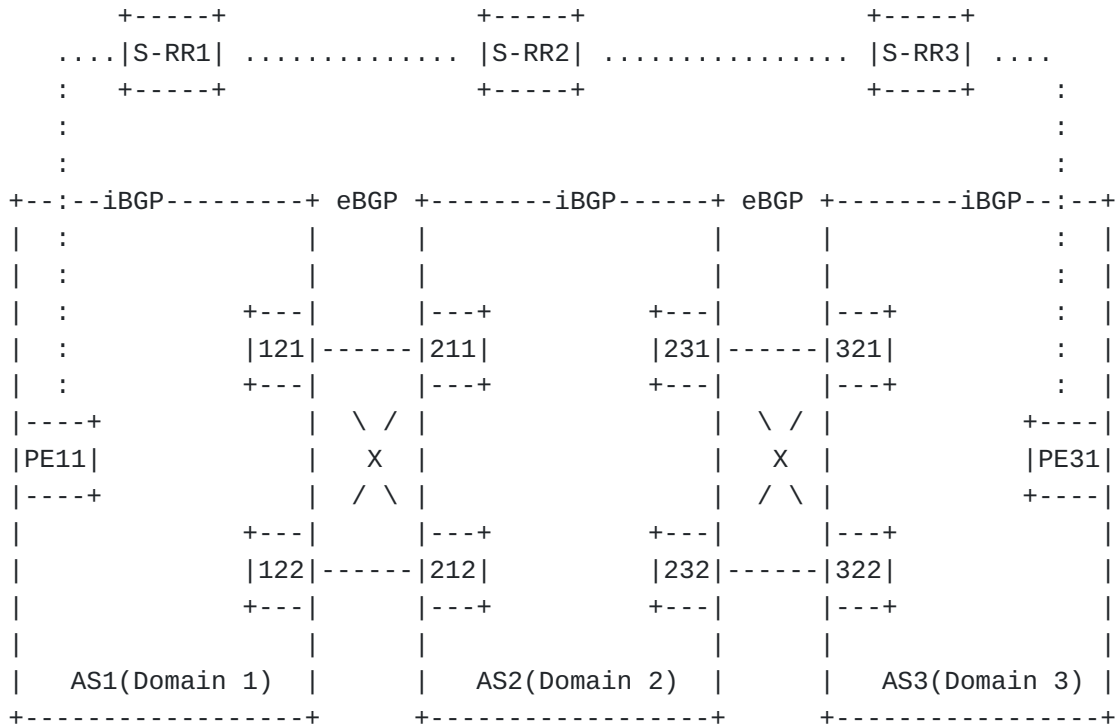


Figure 16: Transport Multiple Domains with eBGP Network Diagram

The above diagram shows three different ASes (AS1, AS2 and AS3.) 121,122, 211, 212, 231,232, 321,322 are border nodes between the ASes.

In this design, domain inter-connect is via eBGP peering links between AS border nodes. The ASBR also runs I-BGP sessions with other ASBRs or RRs in the same AS.

6.3.1.4. Multiple sites with same AS connected via different core AS




Figure 17: Transport Multiple Domains with same AS Network Diagram

121,122,231,232 belong to AS2 only. AS1 and AS2 domains may run multi-instance IGP or different levels/areas.

This topology uses I-BGP sessions to some clients and E-BGP sessions to other nodes. When an RR is used between PEs in AS1 and ABRs in AS2, it will have iBGP sessions to clients in same AS and e-BGP sessions to nodes in other AS.

6.3.1.5. AS Confederations

BGP confederations [RFC 5065] allows the division of a public AS into multiple sub-ASes, usually with private identifiers. The solution should support BGP based intent-aware paths within the sub-AS or across the sub-ASes of the confederation, in any of the network designs described in sections 5.4.1.1 to section 5.4.1.4.

6.3.1.6. Transport Technologies

6.3.1.6.1. Unicast transport

The solution must support the following:

*End-to-end paths crossing transport domains that use different technologies and encapsulations, such as:

-LDP-MPLS

-RSVP-TE-MPLS

-SR-MPLS

-SRv6

-SR-TE (MPLS and SRv6)

-IGP Flex-Algo (MPLS and SRv6)

-Native IPv4/IPv6 forwarding (networks without MPLS enabled)

*Note:

-All MPLS/SR-MPLS deployments may be IPv4/IPv6 or dual-stack

-SR-TE includes color-only and other policies as defined in [RFC9256]

*Interworking between domains with different encapsulations (e.g. SR-MPLS and SRv6)

*Different transport encapsulations simultaneously within a domain, for co-existence and migration

6.3.1.6.2. Multicast transport

A routing solution for end-to-end intent-aware paths should support multicast as well as unicast. An End-to-end multicast path crossing multiple transport domains may use different encapsulation mechanisms, such as:

*mLDP

*RSVP-TE P2MP

*SR-MPLS Tree SIDs

*SRv6 Tree SIDs

*Native IPv4/IPv6 forwarding (networks without MPLS enabled)

6.3.1.7. Co-existence, compatibility and interworking with existing intent-aware routing solutions

The BGP intent-aware routing solution MUST be compliant with the intent-aware routing framework described in [Section 5](#). Specifically,

*It MUST support service routes using Color Extended-Community to request intent as defined in [\[RFC9256\]](#)

*It MUST support automated steering of colored service routes on a BGP intent-aware path using color

*Intent-aware routes MAY resolve recursively via other intent-aware routes provided by any solution

6.3.1.8. Co-existence and Interworking with BGP-LU

BGP-LU [\[RFC8277\]](#) is widely deployed to provide inter-domain best-effort connectivity across different domains. The BGP intent-aware routing solution should support:

*Establishment of best-effort paths by using a color to represent best-effort intent, to avoid the need to deploy both technologies

*Co-existence of inter-domain BGP-LU and BGP intent aware routing in a network

*Support interworking of BGP-LU and BGP intent-aware network domains.

6.3.1.9. Domains with different intent granularity

All domains in a network may not support the same number and granular definition of colors. However, the maximum granularity of colors should be provided for end to end paths that are set up for steering of a colored service route, with mapping from a more granular color to a less granular color where needed.

6.3.1.10. Domains with non-congruent Color-to-intent Mappings

As illustrated in [Section 4.1](#), network domains under different administrative control may assign different colors to represent the same intent.

A color domain represents a collection of one or more network (IGP/BGP) domains with a single, consistent set of color-to-intent mappings.

Color for a given intent may need to be re-mapped across a color domain boundary. The solution should support efficient color re-mapping for intent-aware routes that are propagated to a different color domain.

6.3.1.11. Color-to-intent coordination

It may be useful to have a few well-defined common intents with well-known color values assigned that can be used across multiple operator networks. Such a scheme can provide a consistent service definition to customers that use paths for such intents from multiple operators.

This scheme does not preclude operators from defining intents with similar characteristics and assigning other color values.

It may additionally be useful to have a reserved range of color values for requirements that may arise in future.

However, it should also be noted that color assignments have been used in customer deployments for a while now. Coordination and care will be necessary for defining a range that does not conflict with current deployments.

6.3.1.12. Co-existence with alternative solutions

Section 5 describes co-existence and interworking of the BGP intent aware routing solution with other existing intent-aware solutions.

Controller based approaches or other distributed TE solutions can also address the use-cases in this document.

The intent-aware routing solution should coexist with such alternative solutions.

- *It should allow traffic to use paths created by an alternative solution.

- *It should allow part of the inter-domain path to be created by an alternative solution.

- *The routing solution may be used to provide backup paths for a primary path created by an alternative solution, or vice versa.

6.3.2. Scalability Requirements

6.3.2.1. Scale Requirements

*Support a massive scaled transport network

-Number of Remote PE's: $\geq 300k$

-Number of Colors C: ≥ 5

*Support a scalable MPLS dataplane solution

*Constraints that need to be addressed:

-Typical inter-domain MPLS network designs (e.g. Seamless-MPLS) build hop-by-hop stitched MPLS LSPs towards every PE in the network. For the scale above, the number of forwarding entries required to represent each remote PE for each color will exceed the 1M MPLS label space limit.

-PE and transit nodes may be devices with low FIB capacity.

-Additionally, they may also have constraints on packet processing (e.g, label ops, number of labels pushed)

*To address these constraints:

-The solution must support hierarchy in the forwarding plane E.g. via a label stack or a list of segments, such that no single node needs to support a data-plane scaling in the order of (Remote PE * C)

-The solution should minimize state on border nodes in order to reduce label and FIB resource consumption, while taking into account packet processing constraints.

*Support ability to abstract the topology and network events from remote domains - for scale, stability and faster convergence.

-E.g. contain the control plane propagation of a failure event for an ABR within its attached upstream domain.

*Support an Emulated-PULL model for the BGP signaling

PE nodes may be devices with limited CPU and memory. The state on a PE should be restricted to transport endpoints that it needs for service steering.

BGP Signaling is natively a PUSH model.

For comparison, the SR-PCE solution natively supports a PULL model: when PE1 installs a VPN route V/v via (C, PE2), PE1 requests its serving SR-PCE to compute the SR Policy to (C, PE2). I.e. PE1 does not learn unneeded SR policies.

Emulated-PULL refers to the ability for a BGP node PE1 to "subscribe" to (C, PE2) route such that only paths for (C, PE2) are signaled to PE1.

The requirements for an Emulated-PULL solution are as follows:

- *The subscription and related filtering solution must apply to any BGP node.

- *For transport routes, this means

- Ability for a node (e.g. PE/ABR/ASBR) to signal interest for routes of specific colors.

- Ability for a node (e.g. ABR/ASBR) to propagate the subscription message.

- PEs may choose to only learn routes that they need - e.g. remote VPN endpoints (PEs/VPN ASBRs) or transit nodes (ABRs/transport ASBRs).

- ABR/ASBRs also only learn and propagate routes for which nodes within the local domain have expressed interest.

- The requirements for VPN routes will be updated in the future version of the document.

- *Automation of the subscription/filter route

- Similar to the SR-PCE solution, when an ingress PE1 installs VPN V/v via (C, PE2), PE1 originates its subscription/filter route for (C, PE2).

- *Efficient propagation and processing of subscription/filter routes.

- Ability to summarize the endpoints and thus request a number of endpoints for a particular intent in a single subscription route.

- *The solution may be optional for networks that do not have the large scaling requirements.

6.3.2.2. Scale Analysis

This section will be updated in the future revision of the document.

6.3.3. Network Availability Requirements

*A BGP intent-aware routing solution should provide high network availability for typical deployment topologies, with minimum loss of connectivity in different network failure scenarios.

*The network failure scenarios, applicable technologies and design options described in [[I-D.ietf-mpls-seamless-mpls](#)] should be used as a reference.

*In the Seamless-MPLS reference topology in section 5.4.1.1 :

- Failure of intra-domain links should limit loss of connectivity (LoC) to under 50ms. E.g., PE11 to a P node (not shown), 121 to a P node in Domain1 or Domain2)

- Failure of an intra-domain node (P node in any domain) should limit LoC to under 50ms

- Failure of an ABR node (e.g. 121, 231) should limit LoC to under 1sec, or under 50ms depending on the network deployment scenario.

- Failure of a remote PE node (e.g. PE31) should limit LoC to under 1sec, or under 50ms depending on the network deployment scenario and specific service failover requirements

*In the Inter-AS Option C VPN reference topology in Section 5.4.1.3:

- Failure of intra-domain links should limit LoC to under 50ms. E.g., PE11 to a P node (not shown), 121 to a P node in Domain1 or Domain2)

- Failure of an intra-domain node (P node in any domain) should limit LoC to under 50ms

- Failure of an ASBR node (e.g. 121, 211) should limit LoC to under 1sec, or under 50ms depending on the network deployment scenario.

- Failure of a remote PE node (e.g. PE31) should limit LoC to under 1sec, or under 50ms depending on the network deployment scenario and specific service failover requirements

- Failure of an external link (e.g. 121-211) should limit LoC to under 1sec, or under 50ms depending on the network deployment scenario.

*The solution should explore and describe additional techniques and design options that are applicable to further improve handling of the failure cases listed above.

6.3.4. BGP Protocol Requirements

This section summarizes the key protocol requirements that should be addressed by the intent-aware BGP routing solution. While the context for several requirements has been discussed earlier in the document, this section emphasizes aspects pertinent to the protocol design.

The solution should support the following:

*Signaling and distribution of different Intent-aware routes to reach a participating node, e.g. a PE. Intent must be indicated by the notion of a Color as defined in [[RFC9256](#)]

- Signal different instances of a prefix, one route per color

- Signal intent (color) associated with each route

- At any BGP hop, allow propagating the best path selected for each route, or additional paths

- Generate routes sourced from IGP-FA, SR-TE Policies, RSVP-TE and BGP-LU from a domain

*Path selection for Intent-aware routes

- Accumulation of intent specific metric at each BGP hop and compare the accumulated metric across all received paths at intermediate hops and at an ingress PE.

- Ability to load balance among multiple received paths at intermediate BGP hops and at an ingress PE

- Backup path installation for fast convergence at intermediate BGP hops and at an ingress PE

*Validation of received paths

- Resolvability of next-hop in control plane

- Availability of encapsulation in data plane

*Next-hop resolution for BGP Intent-aware route

- Flexibility to use different intra-domain and inter-domain mechanisms, both intent-aware and traditional

 - oIGP-FA, SR-TE, RSVP-TE, IGP, BGP-LU etc.

- Recursive resolution over other BGP Intent-Aware routes

- Recursive resolution via alternative color or best-effort paths when a particular intent is not available in a domain

*Flexible, efficient, extensible protocol definition

- As an example for context, currently deployed mechanisms such as BGP-LU (RFC 8277) were designed for MPLS, hence only signal per prefix label(s) in NLRI. However, RFC9012 and RFC8669 have described extensions to BGP to signal multiple encapsulations, though in BGP attributes. The target deployments for intent-aware routing need to support additional transport as described in section 6.3.1.6.1. In addition, they also need to support a significantly higher targeted scale as described in scaling requirements.

- Hence, the protocol definition should

 - oSupport efficient signaling of different transport encapsulations

 - oSupport efficient signaling multiple encapsulations for co-existence and migration between encapsulations

 - oAccommodate efficiency of processing and future extensibility

*Separation of transport and VPN service semantics

- Allow for different route distribution planes or processing for service vs transport routes

*Signaling across domains with different color mappings for a given intent

6.3.5. OAM Requirements

OAM in each domain should be function independently. This allows for more flexible evolution of the network.

Basic MPLS OAM mechanisms described in [RFC8029] should be supported for MPLS based solutions deployments. Extensions defined in [RFC8287] should be supported.

Mechanisms described in [RFC 9259] should be supported for SRv6 based deployments.

End-to-end ping and traceroute procedures should be supported.

The ability to validate the path inside each domain should be supported.

Statistics for inter-domain intent-based transport paths should be supported on a per intent-aware path basis on the ingress PE nodes and as needed on egress and border nodes.

7. Backward Compatibility

This section will be updated in the future version of the document.

8. Security Considerations

This section will be updated in the future version of the document.

9. IANA Considerations

This section will be updated in the future version of the document.

10. Acknowledgements

The authors would especially like to thank Joel Halpern for his guidance on the collaboration work that has produced this document and feedback on many aspects of the problem statement.

We would like to thank Daniel Voyer, Robert Raszuk, Kireeti Kompella, Ron Bonica, Krzysztof Szarkowicz, Julian Lucek, Ram Santhanakrishnan, Stephane Litkowski, Andrew Alston for discussions and inputs.

We also express our appreciation to Hannes Gredler, Simon Spraggs, Jose Liste and Jiri Chaloupka for discussions that have helped provide input to the problem statement.

Many thanks to Colby Barth, John Scudder, Kamran Raza, Kris Michelson, Huaimo Chen for their review and valuable suggestions.

Many thanks to Qassim Badat for valuable inputs on multicast requirements.

11. Co-authors

1. Srihari Sangli

Juniper Networks Inc.

ssangli@juniper.net

2. Swadesh Agrawal

Cisco Systems

swaagraw@cisco.com

3. Clarence Filsfils

Cisco Systems

cfilsfils@cisco.com

4. Ketan Talaulikar

Cisco Systems

ketan.ietf@gmail.com

5. Keyur Patel

Arrcus, Inc

keyur@arrcus.com

6. Bruno Decraene

Orange

bruno.decraene@orange.com

7. Xiaohu Xu

China Mobile

13910161692@qq.com

8. Arkadiy Gulko

EdwardJones

arkadiy.gulko@edwardjones.com

9. Mazen Khaddam

Cox communications

mazen.khaddam@cox.com

10. Luis M. Contreras

Telefonica

luismiguel.contrerasmurillo@telefonica.com

11. Dirk Steinberg

Lapishills Consulting Limited

dirk@lapishills.com

12. Jim Guichard

Futurewei

james.n.guichard@futurewei.com

13. Wim Henderickx

Nokia

wim.henderickx@nokia.com

14. Chris Bowers

Independent Contributor

12. Contributors

1. Kaliraj Vairavakkalai

Juniper Networks

kaliraj@juniper.net

2. Jeffrey Zhang

Juniper Networks

zzhang@juniper.net

13. References

13.1. 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, <<https://www.rfc-editor.org/info/rfc2119>>.

13.2. Informative References

[I-D.dskc-bess-bgp-car]

Rao, D., Agrawal, S., Filsfils, C., Steinberg, D., Jalil, L., Su, Y., Decraene, B., Guichard, J., Talaulikar, K., Patel, K., Wang, H., and J. Uttaro, "BGP Color-Aware Routing (CAR)", Work in Progress, Internet-Draft, draft-dskc-bess-bgp-car-05, 6 July 2022, <<https://datatracker.ietf.org/doc/html/draft-dskc-bess-bgp-car-05>>.

[I-D.dskc-bess-bgp-car-problem-statement]

Rao, D., Agrawal, S., Filsfils, C., Decraene, B., Steinberg, D., Jalil, L., Guichard, J., Talaulikar, K., Patel, K., and W. Henderickx, "BGP Color-Aware Routing Problem Statement", Work in Progress, Internet-Draft, draft-dskc-bess-bgp-car-problem-statement-05, 26 May 2022, <<https://datatracker.ietf.org/doc/html/draft-dskc-bess-bgp-car-problem-statement-05>>.

[I-D.filsfils-spring-sr-policy-considerations]

Filsfils, C., Talaulikar, K., Król, P. G., Horneffer, M., and P. Mattes, "SR Policy Implementation and Deployment Considerations", Work in Progress, Internet-Draft, draft-filsfils-spring-sr-policy-considerations-09, 24 April 2022, <<https://datatracker.ietf.org/doc/html/draft-filsfils-spring-sr-policy-considerations-09>>.

[I-D.hegde-rtgwg-egress-protection-sr-networks] Hegde, S., Lin, W., and S. Peng, "Egress Protection for Segment Routing (SR) networks", Work in Progress, Internet-Draft, draft-hegde-rtgwg-egress-protection-sr-networks-02, 2 March 2022, <<https://datatracker.ietf.org/doc/html/draft-hegde-rtgwg-egress-protection-sr-networks-02>>.

[I-D.hegde-spring-node-protection-for-sr-te-paths]

Hegde, S., Bowers, C., Litkowski, S., Xu, X., and F. Xu, "Node Protection for SR-TE Paths", Work in Progress, Internet-Draft, draft-hegde-spring-node-protection-for-sr-te-paths-07, 30 July 2020, <<https://>

datatracker.ietf.org/doc/html/draft-hegde-spring-node-protection-for-sr-te-paths-07>.

[I-D.hegde-spring-seamless-sr-architecture]

Hegde, S., Bowers, C., Xu, X., Gulko, A., Bogdanov, A., Uttaro, J., Jalil, L., Khaddam, M., and A. Alston, "Seamless Segment Routing Architecture", Work in Progress, Internet-Draft, draft-hegde-spring-seamless-sr-architecture-00, 22 February 2021, <<https://datatracker.ietf.org/doc/html/draft-hegde-spring-seamless-sr-architecture-00>>.

[I-D.ietf-idr-performance-routing]

Xu, X., Hegde, S., Talaulikar, K., Boucadair, M., and C. Jacquenet, "Performance-based BGP Routing Mechanism", Work in Progress, Internet-Draft, draft-ietf-idr-performance-routing-03, 22 December 2020, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-performance-routing-03>>.

[I-D.ietf-idr-segment-routing-te-policy]

Previdi, S., Filsfils, C., Talaulikar, K., Mattes, P., and D. Jain, "Advertising Segment Routing Policies in BGP", Work in Progress, Internet-Draft, draft-ietf-idr-segment-routing-te-policy-25, 26 September 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-segment-routing-te-policy-25>>.

[I-D.ietf-lsr-flex-algo-bw-con]

Hegde, S., Britto, W., Shetty, R., Decraene, B., Psenak, P., and T. Li, "Flexible Algorithms: Bandwidth, Delay, Metrics and Constraints", Work in Progress, Internet-Draft, draft-ietf-lsr-flex-algo-bw-con-07, 26 September 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-lsr-flex-algo-bw-con-07>>.

[I-D.ietf-mpls-seamless-mpls]

Leymann, N., Decraene, B., Filsfils, C., Konstantynowicz, M., and D. Steinberg, "Seamless MPLS Architecture", Work in Progress, Internet-Draft, draft-ietf-mpls-seamless-mpls-07, 28 June 2014, <<https://datatracker.ietf.org/doc/html/draft-ietf-mpls-seamless-mpls-07>>.

[I-D.ietf-pce-segment-routing-policy-cp]

Koldychev, M., Sivabalan, S., Barth, C., Peng, S., and H. Bidgoli, "PCEP extension to support Segment Routing Policy Candidate Paths", Work in Progress, Internet-Draft, draft-ietf-pce-segment-routing-policy-cp-12, 24

July 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-pce-segment-routing-policy-cp-12>>.

[I-D.ietf-rtgwg-segment-routing-ti-lfa]

Litkowski, S., Bashandy, A., Filsfils, C., Francois, P., Decraene, B., and D. Voyer, "Topology Independent Fast Reroute using Segment Routing", Work in Progress, Internet-Draft, draft-ietf-rtgwg-segment-routing-ti-lfa-11, 30 June 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-rtgwg-segment-routing-ti-lfa-11>>.

[I-D.kaliraj-idr-bgp-classful-transport-planes]

Vairavakkalai, K., Venkataraman, N., Rajagopalan, B., Mishra, G. S., Khaddam, M., Xu, X., Szarecki, R. J., Gowda, D. J., Yadlapalli, C., and I. Means, "BGP Classful Transport Planes", Work in Progress, Internet-Draft, draft-kaliraj-idr-bgp-classful-transport-planes-17, 30 June 2022, <<https://datatracker.ietf.org/doc/html/draft-kaliraj-idr-bgp-classful-transport-planes-17>>.

[I-D.voyer-pim-sr-p2mp-policy] Voyer, D., Filsfils, C., Parekh, R., Bidgoli, H., and Z. J. Zhang, "Segment Routing Point-to-Multipoint Policy", Work in Progress, Internet-Draft, draft-voyer-pim-sr-p2mp-policy-02, 10 July 2020, <<https://datatracker.ietf.org/doc/html/draft-voyer-pim-sr-p2mp-policy-02>>.

[I-D.zzhang-bess-bgp-multicast]

Zhang, Z. J., Giuliano, L., Patel, K., Wijnands, I., Mishra, M. P., and A. Gulko, "BGP Based Multicast", Work in Progress, Internet-Draft, draft-zzhang-bess-bgp-multicast-03, 29 October 2019, <<https://datatracker.ietf.org/doc/html/draft-zzhang-bess-bgp-multicast-03>>.

[RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, DOI 10.17487/RFC3630, September 2003, <<https://www.rfc-editor.org/info/rfc3630>>.

[RFC3906] Shen, N. and H. Smit, "Calculating Interior Gateway Protocol (IGP) Routes Over Traffic Engineering Tunnels", RFC 3906, DOI 10.17487/RFC3906, October 2004, <<https://www.rfc-editor.org/info/rfc3906>>.

[RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <<https://www.rfc-editor.org/info/rfc4271>>.

- [RFC4272] Murphy, S., "BGP Security Vulnerabilities Analysis", RFC 4272, DOI 10.17487/RFC4272, January 2006, <<https://www.rfc-editor.org/info/rfc4272>>.
- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, DOI 10.17487/RFC4364, February 2006, <<https://www.rfc-editor.org/info/rfc4364>>.
- [RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, DOI 10.17487/RFC5305, October 2008, <<https://www.rfc-editor.org/info/rfc5305>>.
- [RFC6952] Jethanandani, M., Patel, K., and L. Zheng, "Analysis of BGP, LDP, PCEP, and MSDP Issues According to the Keying and Authentication for Routing Protocols (KARP) Design Guide", RFC 6952, DOI 10.17487/RFC6952, May 2013, <<https://www.rfc-editor.org/info/rfc6952>>.
- [RFC7311] Mohapatra, P., Fernando, R., Rosen, E., and J. Uttaro, "The Accumulated IGP Metric Attribute for BGP", RFC 7311, DOI 10.17487/RFC7311, August 2014, <<https://www.rfc-editor.org/info/rfc7311>>.
- [RFC7471] Giacalone, S., Ward, D., Drake, J., Atlas, A., and S. Previdi, "OSPF Traffic Engineering (TE) Metric Extensions", RFC 7471, DOI 10.17487/RFC7471, March 2015, <<https://www.rfc-editor.org/info/rfc7471>>.
- [RFC7911] Walton, D., Retana, A., Chen, E., and J. Scudder, "Advertisement of Multiple Paths in BGP", RFC 7911, DOI 10.17487/RFC7911, July 2016, <<https://www.rfc-editor.org/info/rfc7911>>.
- [RFC8570] Ginsberg, L., Ed., Previdi, S., Ed., Giacalone, S., Ward, D., Drake, J., and Q. Wu, "IS-IS Traffic Engineering (TE) Metric Extensions", RFC 8570, DOI 10.17487/RFC8570, March 2019, <<https://www.rfc-editor.org/info/rfc8570>>.
- [RFC9012] Patel, K., Van de Velde, G., Sangli, S., and J. Scudder, "The BGP Tunnel Encapsulation Attribute", RFC 9012, DOI 10.17487/RFC9012, April 2021, <<https://www.rfc-editor.org/info/rfc9012>>.
- [RFC9256] Filsfils, C., Talaulikar, K., Ed., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture",

RFC 9256, DOI 10.17487/RFC9256, July 2022, <<https://www.rfc-editor.org/info/rfc9256>>.

[RFC9350] Psenak, P., Ed., Hegde, S., Filsfils, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", RFC 9350, DOI 10.17487/RFC9350, February 2023, <<https://www.rfc-editor.org/info/rfc9350>>.

Authors' Addresses

Shraddha Hegde (Editor)
Juniper Networks Inc.
Exora Business Park
Bangalore 560103
KA
India

Email: shraddha@juniper.net

Dhananjaya Rao (Editor)
Cisco Systems
United States of America

Email: dhrao@cisco.com

James Uttaro
Independent Contributor

Email: juttaro@ieee.org

Alex Bogdanov
BT

Email: alex.bogdanov@bt.com

Luay Jalil
Verizon

Email: luay.jalil@verizon.com