Segment Routing - Automated Steering

This document gives an example of how Automated Steering works in Segment Routing based on the below topology. Traffic flow will be examined from CE1 to CE2 based on policies set by the ingress PE (PE2). No custom SR Flexible-Algorithms are used.

What is Automated Steering?
Automated Steering (or AS) is a Segment Routing feature that enables the steering of a service route into a valid SR Policy based on its **next-hop** and **color**.

What is an SR Policy?
An SR Policy is essentially a SID list that is either explicitly defined, or dynamically created according to an optimisation criteria (e.g. lowest metric) and/or any optional restrictions (e.g. avoid certain links). SR Policies can be created explicitly, either through configuration or via PCE, or dynamically using ODN (On Demand Next-hop). This document explores explicitly create policies that create dynamic SID lists.

What is Color?
Color is a BGP opaque extended community defined in RFC 5512 used to mark routes. It is just a number but it is colloquially referenced using a color name. Don't confuse this with the term "link-color" which refers to the affinity bits on a TE link. Both are used here to illustrate the difference.

Network Topology

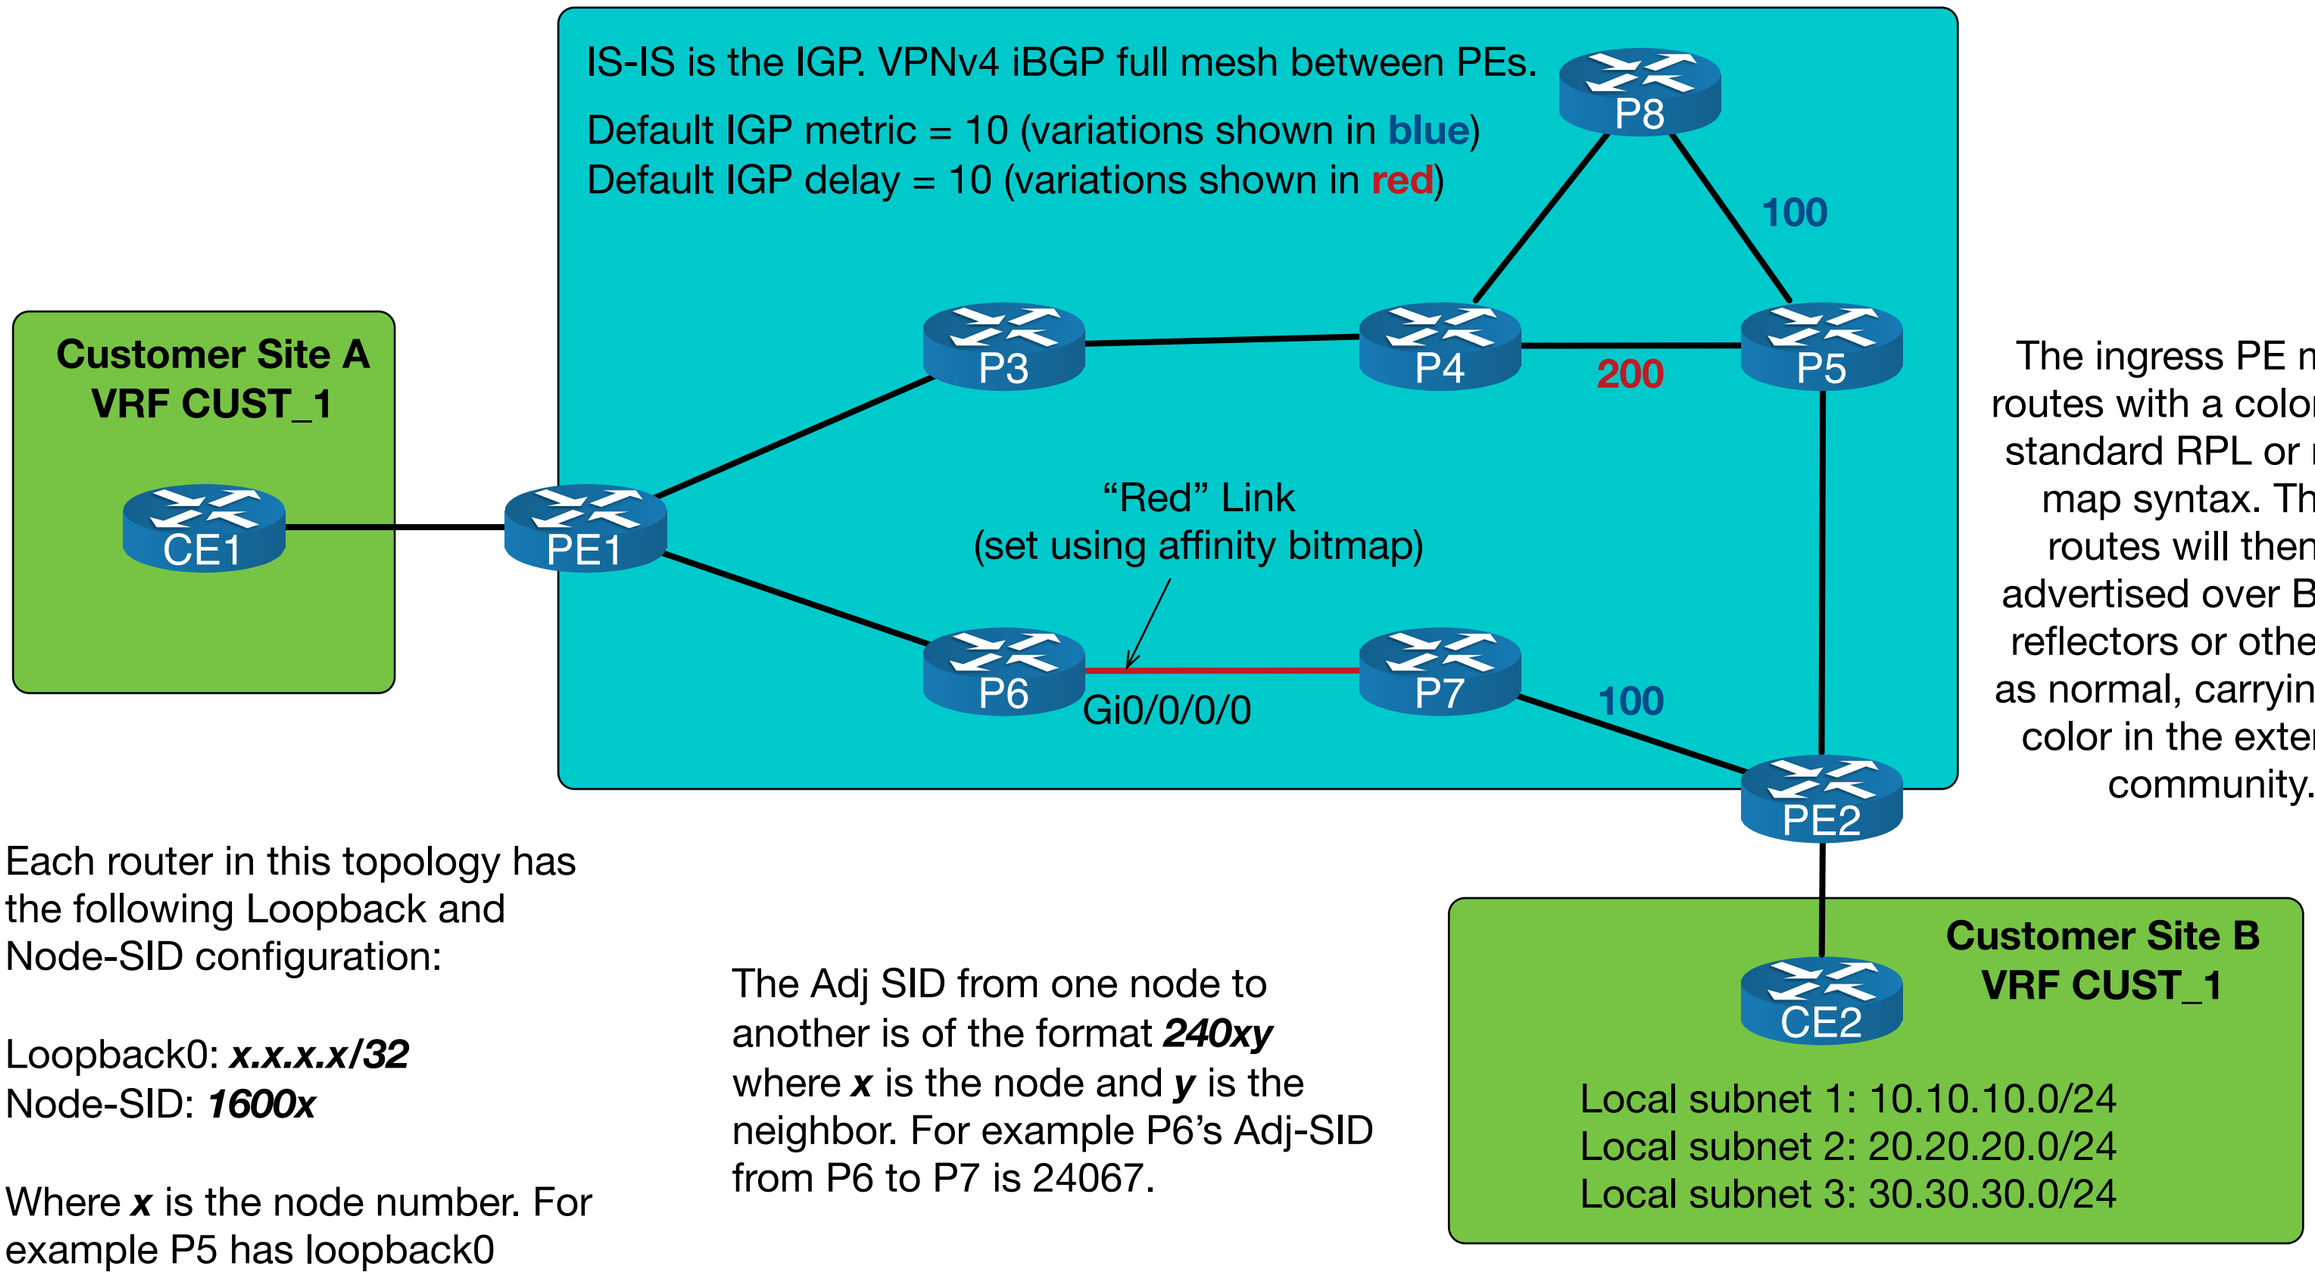

The ingress PE marks routes with a color using standard RPL or route-map syntax. These routes will then be advertised over BGP to reflectors or other PEs as normal, carrying their color in the extended community.

Each router in this topology has the following Loopback and Node-SID configuration:

Loopback0: ***x.x.x.x/32***
Node-SID: ***1600x***

Where ***x*** is the node number. For example P5 has loopback0 address 5.5.5.5/32 with corresponding Node-SID 16005.

The Adj SID from one node to another is of the format ***240xy*** where ***x*** is the node and ***y*** is the neighbor. For example P6's Adj-SID from P6 to P7 is 24067.

Example Policies

Goal	Color
Traffic from Customer site A to 10.10.10.10 should follow path with smallest IGP metric	**BLUE (20)**
Traffic from Customer site A to 20.20.20.20 should follow path with smallest delay metric	**GREEN (30)**
Traffic from Customer site A to 30.30.30.30 should follow path with smallest delay metric avoiding red links	**PURPLE (40)**

Segment Routing - Automated Steering

Configuration

This page shows IOS-XR CLI configuration to setup the policies shown on the previous page. PE1 and PE2 have a loopback to loopback VPNv4 iBGP session between to each other. The PE to CE routing protocol used is eBGP. To illustrate the RPL variations available, subnet1 and subnet2 are tagged with their colors inbound to PE2 and subnet3 is tagged outbound to PE1.

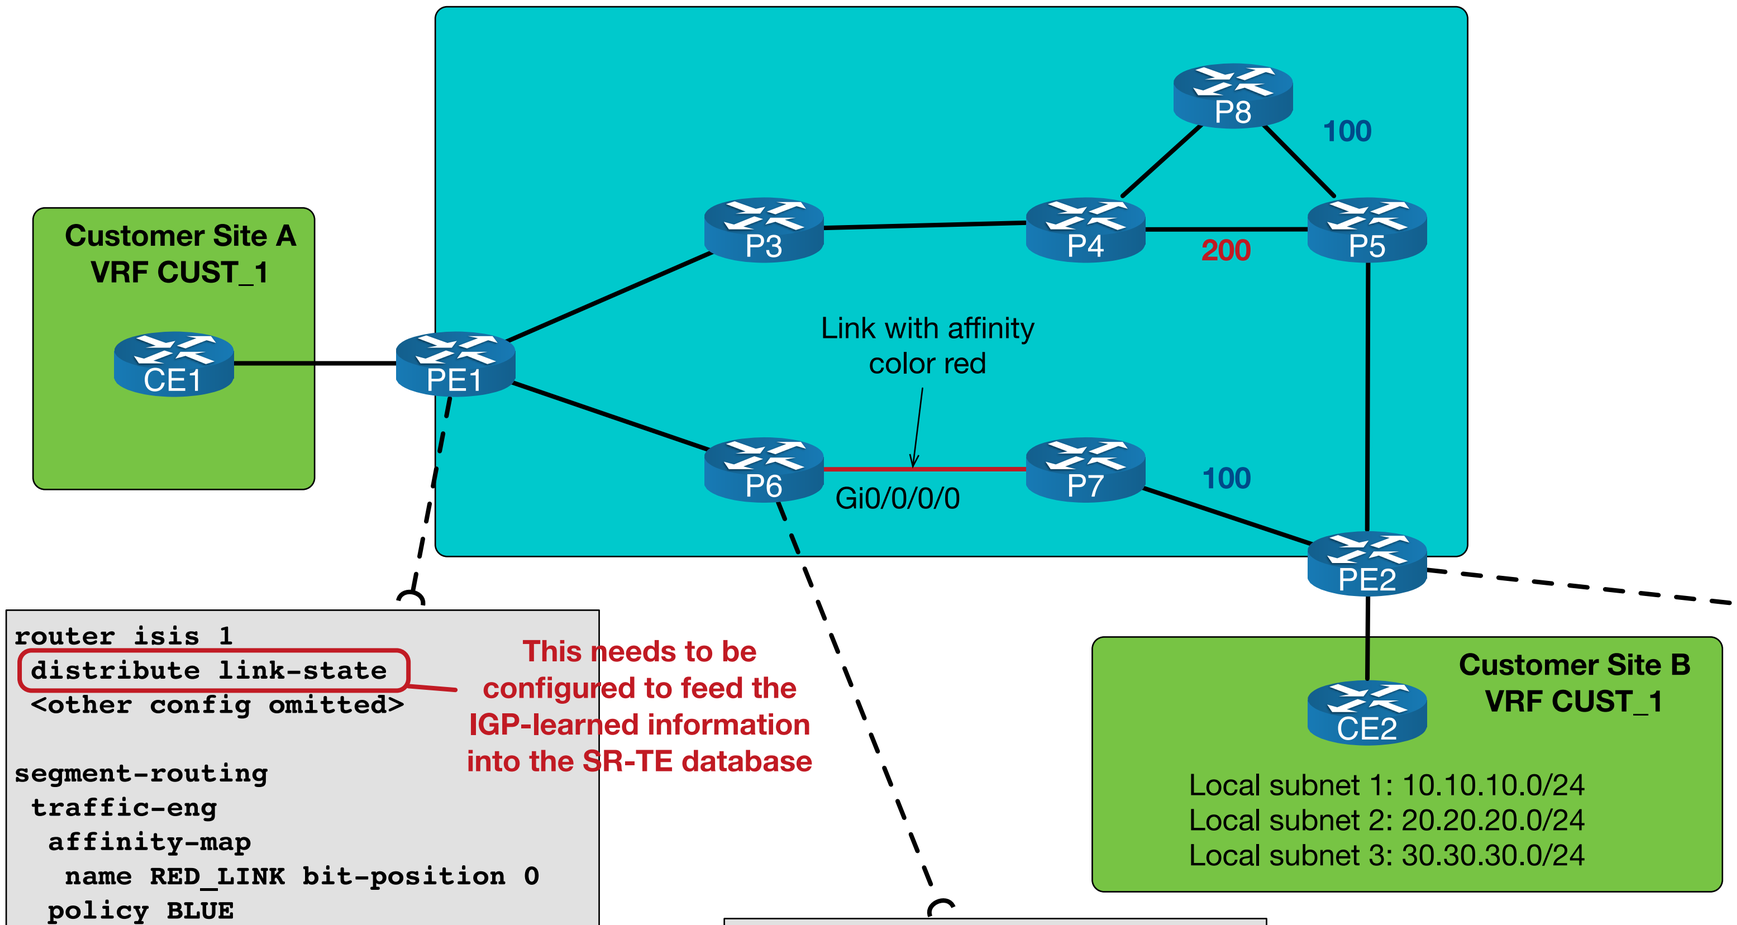


```
router isis 1
 distribute link-state
 <other config omitted>

segment-routing
 traffic-eng
  affinity-map
   name RED_LINK bit-position 0
  policy BLUE
   color 20 end-point ipv4 2.2.2.2
   candidate-paths
    preference 100
     dynamic
      metric
       type igp
  policy GREEN
   color 30 end-point ipv4 2.2.2.2
   candidate-paths
    preference 100
     dynamic
      metric
       type delay
  policy PURPLE
   color 40 end-point ipv4 2.2.2.2
   candidate-paths
    preference 100
     dynamic
      metric
       type delay
      contraints
       affinity exclude-any
         name RED_LINK
```

This needs to be configured to feed the IGP-learned information into the SR-TE database

```
segment-routing
 traffic-eng
  affinity-map
   name RED_LINK bit-position 0
  !
  interface Gi0/0/0/0
   affinity name RED_LINK
```

Refers to the position in the affinity bitmap that RED_LINK represents

This name is locally significant

If multiple candidate paths are defined the highest preference will be chosen.

Each dynamic path is expressed in terms of an OPTIMIZATION OBJECTIVE (in this case lowest delay metric) and a CONSTRAINT (in this case avoid “red” links)

```
extcommunity-set opaque BLUE
 20
end-set
!
extcommunity-set opaque GREEN
 30
end-set
!
extcommunity-set opaque PURPLE
 40
end-set
!
route-policy COLOR-INBOUND-PREFIXES
 if destination in (10.10.10.0/24) then
  set extcommunity color BLUE
 endif
 if destination in (20.20.20.0/24) then
  set extcommunity color GREEN
 endif
  pass
end-policy
!
route-policy INTERNAL-COLOR
 if destination in (30.30.30.0/24) then
  set extcommunity color PURPLE
 endif
  pass
end-policy
!
router bgp 1
 bgp router-id 2.2.2.2
 address-family ipv4 unicast
 address-family vpnv4 unicast
 !
 neighbor 1.1.1.1
  description to CE2
  remote-as 1
  update-source Loopback0
  address-family ipv4 unicast
  address-family vpnv4 unicast
   route-policy INTERNAL-COLOR out
 !
 vrf CUST_1
  rd auto
  address-family ipv4 unicast
  !
  neighbor 172.16.1.2
   remote-as 2
   description to CE2
   address-family ipv4 unicast
    route-policy COLOR-INBOUND-PREFIXES in
```

Segment Routing - Automated Steering

This page examines the SID lists that are created as a results of the 3 policies covered in this document. Each SID in the resulting SID list represents a segment used to adhere to the desired policy.

BGP Updates

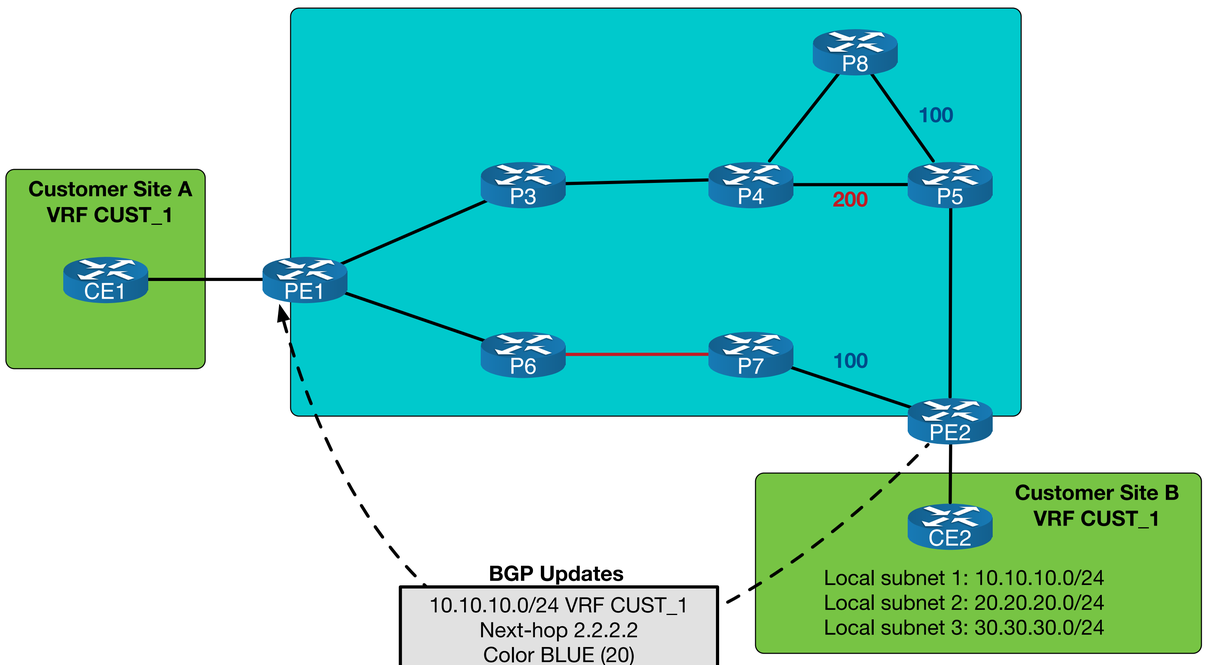

BGP Updates

10.10.10.0/24 VRF CUST_1 Next-hop 2.2.2.2 Color BLUE (20)
20.20.20.0/24 VRF CUST_1 Next-hop 2.2.2.2 Color GREEN (30)
30.30.30.0/24 VRF CUST_1 Next-hop 2.2.2.2 Color PURPLE (40)

BLUE Policy Traffic Flow (Smallest IGP Metric)

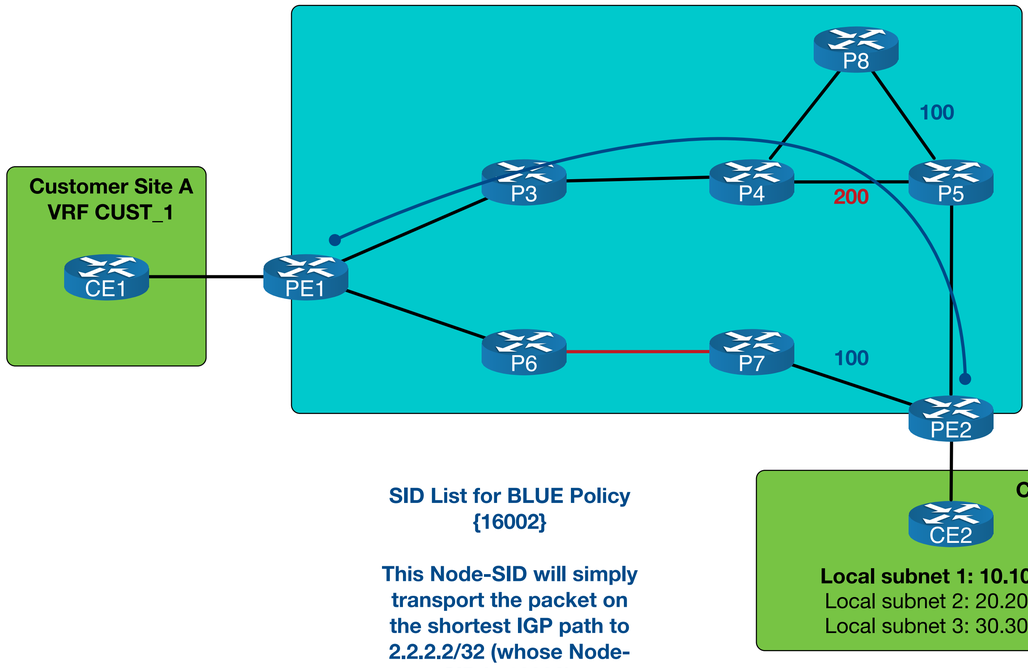

SID List for BLUE Policy
{16002}

This Node-SID will simply transport the packet on the shortest IGP path to 2.2.2.2/32 (whose Node-SID is 16002).

Label Stack:

PE1 - P3	P3-P4	P4-P5	P5-PE2 (assuming default PHP)	PE2-CE2
16002	16002	16002	VPN LABEL	PAYLOAD
VPN LABEL	VPN LABEL	VPN LABEL	PAYLOAD	
PAYLOAD	PAYLOAD	PAYLOAD		

GREEN Policy Traffic Flow (Smallest Delay Metric)

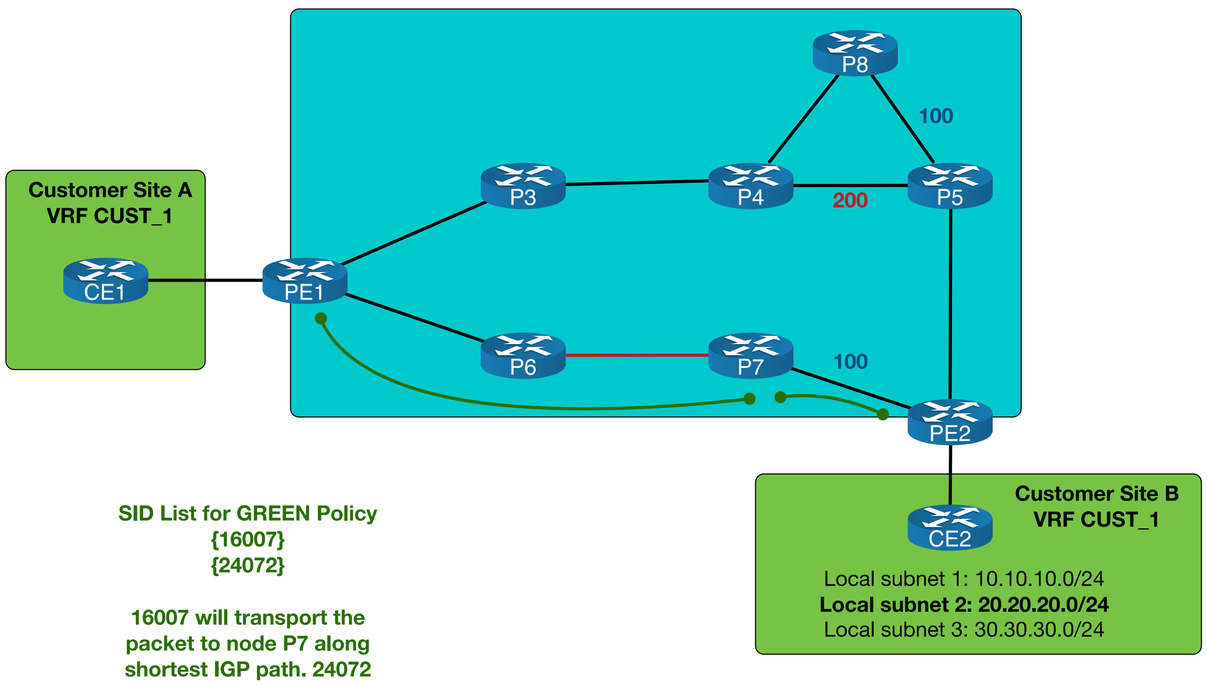

SID List for GREEN Policy
{16007}
{24072}

16007 will transport the packet to node P7 along shortest IGP path. 24072 corresponds to the Adj-SID between P7 and PE2.

Label Stack:

PE1 - P6	P6 - P7 (assuming default PHP)	P7-PE2	PE2-CE2
16007	24072	VPN LABEL	PAYLOAD
24072	VPN LABEL	PAYLOAD	
VPN LABEL	PAYLOAD		
PAYLOAD			

PURPLE Policy Traffic Flow (Smallest Delay Metric Avoiding red links)

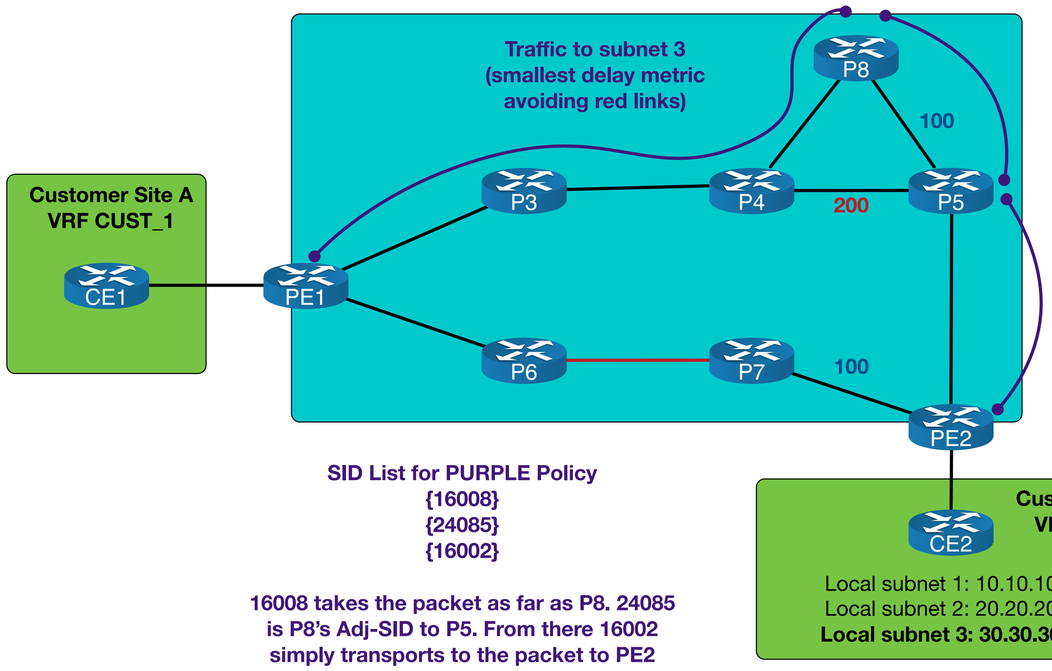

SID List for PURPLE Policy
{16008}
{24085}
{16002}

16008 takes the packet as far as P8. 24085 is P8's Adj-SID to P5. From there 16002 simply transports to the packet to PE2 using the lowest IGP metric.

Label Stack:

PE1 - P3	P3 - P4	P4 - P8 (assuming default PHP)	P8 - P5	P5 - PE2 (assuming default PHP)	PE2-CE2
16008	16008	24085	16002	VPN LABEL	PAYLOAD
24085	24085	16002	VPN LABEL	PAYLOAD	
16002	16002	VPN LABEL	PAYLOAD		
VPN LABEL	VPN LABEL	PAYLOAD			
PAYLOAD	PAYLOAD				

by Steven Crutchley

Segment Routing - Automated Steering

Output - GREEN Policy

This page shows various CLI output from PE1 as it relates to the GREEN policy which specifies the lowest delay metric.

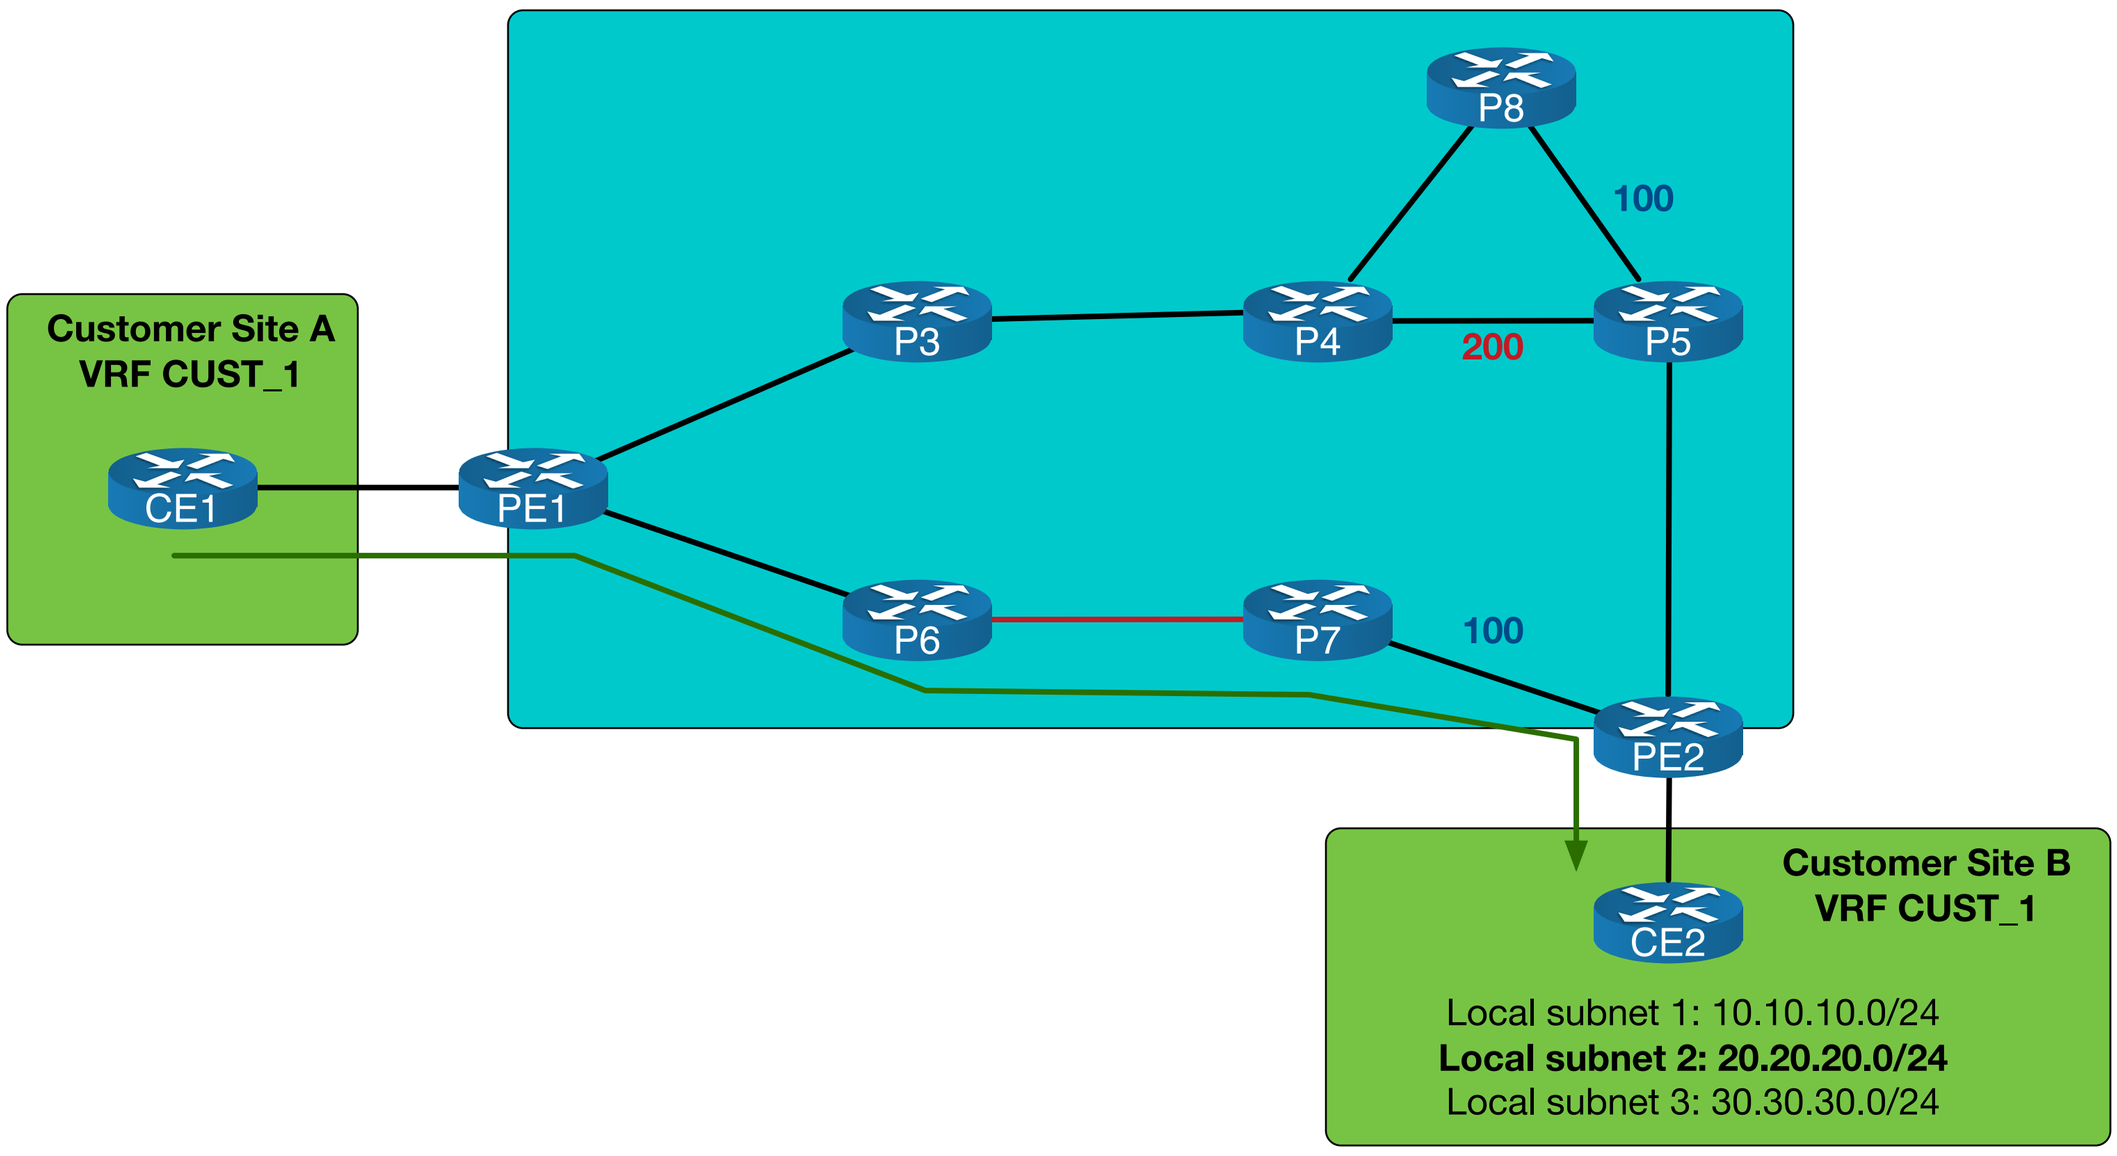


```
RP/0/0/CPU0:PE1# show bgp vrf CUST_1 20.20.20.0/24
BGP routing table entry for 20.20.20.0/24, Route Distinguisher: 1.1.1.1:0
Versions:
  Process bRIB/RIB SendTblVer
  Speaker 7 7
Last Modified: Aug 2 12:01:11.830 for 01:03:00
Paths: (1 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Not advertised to any peer
  2
    2.2.2.2 C:30 (bsid:40002) (metric 40) from 2.2.2.2 (2.2.2.2)
      Received Label 92221 Origin IGP, metric 0, localpref 100, valid, internal, best,
group-best, import-candidate, imported
      Received Path ID 0, Local Path ID 1, version 7
      Extended community: Color:30 RT:1:1 SR policy color 30, up, registered, bsid 40002,
if-handle 0x00000410

      Source AFI: VPNv4 Unicast, Source VRF: default, Source Route Distinguisher: 2.2.2.2:0
```

BGP Color Community

Locally significant Binding SID used to present this policy. This will appear in CEF and the label forwarding table as the local label.

Standard IGP metric (not delay metric)

The standard VPNv4 Label for this prefix

```
RP/0/0/CPU0:PE1# show segment-routing traffic-eng policy

Color: 30, End-point: 2.2.2.2
  Name: srte_c_30_ep_2.2.2.2
  Status:
    Admin: up Operational: up for 00:02:54 (since Aug 2 12:04:31.523)
  Candidate-paths:
    Preference: 100 (configuration) (active)
      Name: GREEN
      Requested BSID: dynamic
      Dynamic (valid)
        Metric Type: delay, Path Accumulated Metric: 30
          16007 [Prefix-SID, 7.7.7.7]
          24072 [Adjacency-SID, 99.2.7.7 99.2.7.2]
  Attributes:
    Binding SID: 40002
    Forward Class: 0
    Steering BGP disabled: no
    IPv6 caps enable: yes
```

This is auto-generated based on the color number and destination.

Accumulated Delay Metric (not IGP metric)

Calculated SID list for this policy

```
RP/0/0/CPU0:PE1# show cef vrf CUST_1 20.20.20.0/24
20.20.20.0/24, version 1, internal 0x5000001 0x0 (ptr 0xa14f440c) [1], 0x0 (0x0), 0x208 (0xa16ac190)
Updated Aug 2 12:04:40.728
Prefix Len 24, traffic index 0, precedence n/a, priority 3
  via local-label 40002 , 3 dependencies, recursive [flags 0x6000]
   path-idx 0 NHID 0x0 [0xa171651c 0x0]
   recursion-via-label
   next hop VRF 'default', table 0xe0000000
   next hop via 40002/0/21
     next hop srte_c_30_ep_2.2.2.2 labels imposed {ImplNull 92221}
```

```
RP/0/0/CPU0:PE1# traceroute vrf CUST_1 20.20.20.1

 Type escape sequence to abort.
 Tracing the route to 20.20.20.1
 1 99.1.6.6 [MPLS: Labels 16007/24072/92221 Exp 0] 19 msec 9 msec 9 msec
 2 99.6.7.7 [MPLS: Labels 24072/92221 Exp 0] 9 msec 10 msec 9 msec
 3 99.2.7.2 [MPLS: Label 92221 Exp 0] 9 msec 9 msec 9 msec
 4 172.16.1.2 9 msec 9 msec 9 msec
```

The ImplNull in this output indicates that the local BSID should be referenced (40002). This will in turn cause the calculated SID-List to be place on to the packet (with 92221, the VPN label, as the bottom of stack label)