



SR Hands-on v3

Frank Zhao


Yonzhao2@cisco.com, CSE, Arch

2019.7.10

Base on 652
SRv6 based on 662



Revise History

Rev	Date	Revise by	Comment
1.0	2019/02/03	Frank Zhao	only include hands-on brief
2.0	2019/04/16	Frank Zhao	hands-on based on XR 6.5.2 add some basic and highlight sr feature/solution detail configuration steps for each labs detail show information explanation for each labs
3.0	2019/07/10	Frank Zhao	add srv6 labs(<i>srv6 transport; control srv6 by 6vpe and DX4</i>), based on 6.6.2 update flexalgo labs



Segment Routing Update

SR Textbook Part II is now available!



SR Update



Clarence Filsfils • 1st
Cisco Fellow at Cisco Systems
1w

...

SR deployments keep rolling.

Etisalat and Cisco announce the first Self-Driving Internet Peering Network in the Region.

<http://bit.ly/2PoFXuE>



Etisalat and Cisco announce the first Self-Driving Internet Peering Network in the Region


emear.thecisconetwork.com

© 2018 Cisco 51 · 2 Comments



Clarence Filsfils • 1st
Cisco Fellow at Cisco Systems
4d

#SegmentRouting journey started 6 years ago with the first submission draft at IETF. Since then market traction has been phenomenal. As of today, over 100 deployments and still counting ... Adoption is covering all continents and market segments: Web, Service Provider and Enterprise; inherently multi-domain: Data Center, Access, Metro, Core. Innovation, product and deployment leadership across SR-MPLS and SRv6 would not have been possible without strong involvement from lead operators. A big thank you to them.



304 · 14 Comments



segment-routing.net

dCloud <https://dcloud.cisco.com/>



思科 Segment
Routing 中文网站




Segment Routing Basic

RFC 8402
- Jul 2018


- What is SR?
- Why need SR?
- What are the benefits of using SR?

Segment的的本质是嵌入在包头中的一条指令，沿途的包转发设备理解并执行该指令

在数据包中体现路径 Data



© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Confidential



Segment Routing标准化进展

- Segment Routing最重要的架构标准”Segment Routing Architecture”已经完成了标准化: RFC8402
- Segment Routing目前已经有8个标准的RFC文件, 另外有数量众多的技术草案预计在2019/2020年完成标准化
 - 25份工作组草案, 33份个人草案
 - [Segment Routing Architecture RFC 8402](#)
 - [Source Packet Routing in Networking \(SPRING\) Problem Statement and Requirements RFC 7855](#)
 - [Use Cases for IPv6 Source Packet Routing in Networking \(SPRING\) RFC 8354](#)
 - [Label Switched Path \(LSP\) Ping/Trace for Segment Routing Networks Using MPLS Dataplane RFC 8287](#)
 - [UDP Return Path for Packet Loss and Delay Measurement for MPLS Networks RFC 7876](#)
 - [IS-IS Traffic Engineering \(TE\) Metric Extensions RFC 7810](#)
 - [OSPF Traffic Engineering \(TE\) Metric Extensions RFC 7471](#)
 - [Packet Loss and Delay Measurement for MPLS Networks RFC 6374](#)

Agenda

- Hands-on Brief
- L3VPN in Segment Routing
- Interworking between LDP and SR
- TI-LFA in Segment Routing
- SRTE Solution in MPLS SR
- SR IGP Flexible Algorithms
- Egress Peer Engineering(EPE)
- SRv6 Solution



Hands-on Brief

dCloud Basic

Back

Info End Save Edit Extend Share

SR Workshop for Partner at ShangHai

Details Servers Resources 5d 07:41 0

IP Address: 198.18.133.252
Credentials:
username: Administrator
password: C1sco12345

Remote Desktop

L3-VLAN-1

Session Details


AnyConnect Credentials

Connect up to 16 devices to the session via Cisco AnyConnect.

Host	dcloud-chi-anyconnect.cisco.com
User	v260User1
Password	6bccca9

DNS Addresses

Open testbed



Login to XRv

The screenshot shows the VM Maestro interface with several windows open:


- Top Left Window:** Shows a network topology diagram with nodes: XTC/BGP Controller, P1, PE4, and AS10. A context menu is open on node P1, listing options: SSH..., Telnet..., Telnet over Web Socket..., Extract configuration(s), and Stop this node. The "Extract configuration(s)" option is highlighted.
- Top Right Window:** Shows a list of simulations under "guest". A context menu is open on the "Demo-SR-Hans-on-7gDp4" simulation, with "Extract configuration(s)" also highlighted.
- Middle Left Window:** Shows the Topology Palette with various tools like Select, Connect, Nodes, and Annotations.
- Middle Right Window:** Shows a detailed network diagram with nodes P1, PE4, P7, PE6, AS11, and AS10.
- Bottom Right Window:** Shows the Properties window for node PE6, displaying configuration details such as hostnames, IP addresses, and interfaces.
- Bottom Far Right Window:** Shows the Graph Overview window with a log of events.

Page Footer:

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Confidential


Lo0: 192.168.0.x; 2001::192:168:0:x
Conn: xy.1.1.x|y; 2001::xy:1:1:x|y

Topology



Lab1 - Segment Routing Basic

- Check reachable by ICMP
- Check ISIS database, verbose/ detail internal
- Check MPLS table, CEF
- Enable Segment Routing
- Check ISIS SR label table, Adj/Prefix info, ISIS Topology detail/SRGB, MPLS table/forwarding, CEF



General Checking

- show route isis; show route ipv6 isis
- show route 192.168.0.4 detail
- show cef 192.168.0.4/32 detail
- ping 192.168.0.4; ping 2001::192:168:0:4
- show isis database verbose internal P7
- show mpls for; show mpls label table

Enable SR at all router

```
router isis srte address-family ipv4 unicast segment-routing mpls
router isis srte address-family ipv4 unicast router-id lo0
router isis srte address-family ipv6 unicast segment-routing mpls
router isis srte interface Loopback0 address-family ipv4 unicast prefix-sid index x
router isis srte interface Loopback0 address-family ipv6 unicast prefix-sid absolute 1700x
```

Checking SR status and compare with old data

- show isis segment label table
- show mpls for
- show isis adj detail
- show isis topology det (SRGB)
- show mpls label table detail (SRGB)
- show isis database verbose internal P7
- show route 192.168.0.4 detail
- show cef 192.168.0.4/32 detail
- trace 192.168.0.4

```
RP/0/0/CPU0:PE6#sh isis segment-routing label table
Wed Apr 10 02:09:20.108 UTC
```

IS-IS srte IS Label Table	
Label	Prefix/Interface
16001	192.168.0.1/32
16004	192.168.0.4/32
16006	Loopback0
16007	192.168.0.7/32
17001	2001::192:168:0:1/128
17004	2001::192:168:0:4/128
17006	Loopback0
17007	2001::192:168:0:7/128

RP/0/0/CPU0:PE6#sh isis adj gigabitEthernet 0/0/0/0 detail						
Wed Apr 10 02:17:00.127 UTC						
IS-IS Level-2 adjacencies:						
System Id	Interface	SNPA	State	Hold	Changed	
IPv6						
BFD						
P1	Gi0/0/0/0	*PtoP*	Up	26	18:57:5	
None						
Area Address:	49					
Neighbor IPv4 Address:	16.1.1.1*					
Adjacency SID:	24000					
Non-FRR Adjacency SID:	24001					
Neighbor IPv6 Address:	fe80::f816:3eff:fe48:63af*					
Adjacency SID:	24004					
Non-FRR Adjacency SID:	24005					
Topology:	IPv4 Unicast					
Topology:	IPv6 Unicast					
Total adjacency count: 1						

```
RP/0/0/CPU0:PE6#sh mpls for
Wed Apr 10 02:10:23.104 UTC
```

Local Label	Outgoing Label	Prefix or ID	Outgoing Interface	Next Hop	Bytes Switched
16001	Pop	SR Pfx (idx 1)	Gi0/0/0/0	16.1.1.1	0
16004	16004	SR Pfx (idx 4)	Gi0/0/0/0	16.1.1.1	0
	16004	SR Pfx (idx 4)	Gi0/0/0/1	67.1.1.7	0
16007	Pop	SR Pfx (idx 7)	Gi0/0/0/1	67.1.1.7	0
17001	Pop	SR Pfx (idx 1001)	Gi0/0/0/0	fe80::f816:3eff:fe48:63af	\
					0
17004	17004	SR Pfx (idx 1004)	Gi0/0/0/0	fe80::f816:3eff:fe48:63af	\
	17004	SR Pfx (idx 1004)	Gi0/0/0/1	fe80::f816:3eff:fe25:8cfcd	\
					0
17007	Pop	SR Pfx (idx 1007)	Gi0/0/0/1	fe80::f816:3eff:fe25:8cfcd	\
					0
24000	Pop	SR Adj (idx 1)	Gi0/0/0/0	16.1.1.1	0
24001	Pop	SR Adj (idx 3)	Gi0/0/0/0	16.1.1.1	0
24002	Pop	SR Adj (idx 1)	Gi0/0/0/1	67.1.1.7	0
24003	Pop	SR Adj (idx 3)	Gi0/0/0/1	67.1.1.7	0
24004	Pop	SR Adj (idx 1)	Gi0/0/0/0	fe80::f816:3eff:fe48:63af	\
					0
24005	Pop	SR Adj (idx 3)	Gi0/0/0/0	fe80::f816:3eff:fe48:63af	\
					0
24006	Pop	SR Adj (idx 1)	Gi0/0/0/1	fe80::f816:3eff:fe25:8cfcd	\
					0
24007	Pop	SR Adj (idx 3)	Gi0/0/0/1	fe80::f816:3eff:fe25:8cfcd	\
					0

IS-IS TLV 扩展

- IS-IS为支持SR引入了下述(sub-)TLVs

• SR Capability sub-TLV (2)	IS-IS Router Capability TLV (242)
• Prefix-SID sub-TLV (3)	Extended IP reachability TLV (135)
• Prefix-SID sub-TLV (3)	IPv6 IP reachability TLV (236)
• Prefix-SID sub-TLV (3)	Multitopology IPv6 IP reachability TLV (237)
• Prefix-SID sub-TLV (3)	SID/Label Binding TLV (149)
• Adjacency-SID sub-TLV (31)	Extended IS Reachability TLV (22)
• LAN-Adjacency-SID sub-TLV (32)	Extended IS Reachability TLV (22)
• Adjacency-SID sub-TLV (31)	Multitopology IS Reachability TLV (222)
• LAN-Adjacency-SID sub-TLV (32)	Multitopology IS Reachability TLV (222)
• SID/Label Binding TLV (149)	

- 基于[draft-ietf-isis-segment-routing-extensions-02](#)的实施

```
RP/0/0/CPU0:PE6#sh mpls label table detail
```

```
Wed Apr 10 02:20:05.724 UTC
```

Table	Label	Owner	State	Rewrite
-------	-------	-------	-------	---------

0	0	LSD(A)	InUse	Yes
0	1	LSD(A)	InUse	Yes
0	2	LSD(A)	InUse	Yes
0	13	LSD(A)	InUse	Yes
0	16000	ISIS(A):srte (Lbl-blk SRGB, vers:0, (start_label=16000, size=8000))	InUse	No
0	24000	ISIS(A):srte (SR Adj Segment IPv4, vers:0, index=1, type=0, intf=Gi0/0/0/0, nh=16.1.1.1)	InUse	Yes
0	24001	ISIS(A):srte (SR Adj Segment IPv4, vers:0, index=3, type=0, intf=Gi0/0/0/0, nh=16.1.1.1)	InUse	Yes
0	24002	ISIS(A):srte (SR Adj Segment IPv4, vers:0, index=1, type=0, intf=Gi0/0/0/1, nh=67.1.1.7)	InUse	Yes
0	24003	ISIS(A):srte (SR Adj Segment IPv4, vers:0, index=3, type=0, intf=Gi0/0/0/1, nh=67.1.1.7)	InUse	Yes
0	24004	ISIS(A):srte (SR Adj Segment IPv6, vers:0, index=1, type=0, intf=Gi0/0/0/0, nh=fe80::f816:3eff:fe48:63af)	InUse	Yes
0	24005	ISIS(A):srte (SR Adj Segment IPv6, vers:0, index=3, type=0, intf=Gi0/0/0/0, nh=fe80::f816:3eff:fe48:63af)	InUse	Yes
0	24006	ISIS(A):srte (SR Adj Segment IPv6, vers:0, index=1, type=0, intf=Gi0/0/0/1, nh=fe80::f816:3eff:fe25:8cf0)	InUse	Yes
0	24007	ISIS(A):srte (SR Adj Segment IPv6, vers:0, index=3, type=0, intf=Gi0/0/0/1, nh=fe80::f816:3eff:fe25:8cf0)	InUse	Yes

```
RP/0/0/CPU0:PE6#sh isis topology detail systemid P7
```

```
Wed Apr 10 02:58:28.056 UTC
```

```
IS-IS srte Level 2 IPv4 Unicast IS Topology
```

```
P7 (192.168.0.7) [ucast 1 mcast <infinity>]
```

```
via 67.1.1.7, GigabitEthernet0/0/0/1, P7, SRGB Base: 16000, Weight: 0
```

```
parent PE6
```

```
adv 0 critical, 0 high, 1 medium, 1 low priority prefixes
```

```
RP/0/0/CPU0:PE6#sh isis database verbose internal P7
```

```
Wed Apr 10 02:31:14.488 UTC
```

```
IS-IS srte (Level-2) Link State Database
```

LSPID	LSP Seq Num	LSP Checksum	LSP Holdtime/Rcvd	ATT/P/OL
-------	-------------	--------------	-------------------	----------

P7.00-00	0x0000005f	0x043c	1192 /1199	0/0/0
----------	------------	--------	------------	-------

291

TLV code:1 length:2

Area Address: 49

TLV code:129 length:2

NLPID: 0xcc

NLPID: 0x8e

TLV code:132 length:4

IP Address: 192.168.0.7

TLV code:135 length:33

Metric: 1 IP-Extended 67.1.1.0/24

SubTLV code:4 length:1

Prefix Attribute Flags: X:0 R:0 N:0

Metric: 0 IP-Extended 192.168.0.7/32

SubTLV code:3 length:6

Prefix-SID Index: 7, Algorithm:0, R:0 N:1 P:0 E:0 V:0 L:0

SubTLV code:4 length:1

Prefix Attribute Flags: X:0 R:0 N:1

TLV code:137 length:2

Hostname: P7

TLV code:232 length:16

IPv6 Address: 2001::192:168:0:7

TLV code:237 length:61

Metric: 1 MT (IPv6 Unicast) IPv6 2001::67:1:1:0/120

SubTLV code:4 length:1

Prefix Attribute Flags: X:0 R:0 N:0

Metric: 0 MT (IPv6 Unicast) IPv6 2001::192:168:0:7/128
--

SubTLV code:3 length:6

Prefix-SID Index: 1007, Algorithm:0, R:0 N:1 P:0 E:0 V:0 L:0
--

SubTLV code:4 length:1

Prefix Attribute Flags: X:0 R:0 N:1

TLV code:229 length:4

MT: Standard (IPv4 Unicast)

MT: IPv6 Unicast

*In order easy check database,
only remain gi0/0/0 of P7 up.*

0/0/0

```
RP/0/0/CPU0:PE6#sh route 192.168.0.4 det
```

```
Wed Apr 10 02:59:33.542 UTC
```

```
Routing entry for 192.168.0.4/32
```

```
Known via "isis srte", distance 115, metric 2, labeled SR, type  
level-2
```

```
Installed Apr 10 02:31:07.059 for 00:28:26
```

```
Routing Descriptor Blocks
```

```
16.1.1.1, from 192.168.0.4, via GigabitEthernet0/0/0/0
```

```
Route metric is 2
```

```
Label: 0x3e84 (16004)
```

```
Tunnel ID: None
```

```
Binding Label: None
```

```
Extended communities count: 0
```

```
Path id:1 Path ref count:0
```

```
NHID:0x1(Ref:8)
```

```
Route version is 0xd (13)
```

```
Local Label: 0x3e84 (16004)
```

```
IP Precedence: Not Set
```

```
QoS Group ID: Not Set
```

```
Flow-tag: Not Set
```

```
Fwd-class: Not Set
```

```
Route Priority: RIB_PRIORITY_NON_RECURSIVE_MEDIUM (7) SVD Type
```

```
RIB_SVD_TYPE_LOCAL
```

```
Download Priority 1, Download Version 41
```

```
No advertising protos.
```

```
RP/0/0/CPU0:PE6#traceroute 192.168.0.4
```

```
Wed Apr 10 03:02:00.342 UTC
```

```
Type escape sequence to abort.
```

```
Tracing the route to 192.168.0.4
```

```
1 16.1.1.1 [MPLS: Label 16004 Exp 0] 119 msec 9 msec 0 msec  
2 14.1.1.4 19 msec * 19 msec
```

```
RP/0/0/CPU0:PE6#show cef 192.168.0.4/32 det
```

```
Wed Apr 10 03:00:55.596 UTC
```

```
192.168.0.4/32, version 41, labeled SR, internal 0x1000001 0x81 (ptr 0xa11f7a54)  
[1], 0x0 (0xa11da698), 0xa28 (0xa15d12ec)
```

```
Updated Apr 10 02:31:07.078
```

```
local adjacency 16.1.1.1
```

```
Prefix Len 32, traffic index 0, precedence n/a, priority 1
```

```
gateway array (0xa1116e04) reference count 3, flags 0x68, source rib (7), 0  
backups
```

```
[2 type 5 flags 0x8401 (0xa15ed5d8) ext 0x0 (0x0)]
```

```
LW-LDI[type=5, refc=3, ptr=0xa11da698, sh-ldi=0xa15ed5d8]
```

```
gateway array update type-time 1 Apr 10 02:31:07.078
```

```
LDI Update time Apr 10 02:31:07.078
```

```
LW-LDI-TS Apr 10 02:31:07.078
```

```
via 16.1.1.1/32, GigabitEthernet0/0/0/0, 7 dependencies, weight 0, class 0
```

```
[flags 0x0]
```

```
path-idx 0 NHID 0x0 [0xa189b1c0 0x0]
```

```
next hop 16.1.1.1/32
```

```
local adjacency
```

```
local label 16004 labels imposed {16004}
```

```
Load distribution: 0 (refcount 2)
```

Hash	OK	Interface	Address
------	----	-----------	---------


0	Y	GigabitEthernet0/0/0/0	16.1.1.1
---	---	------------------------	----------



L3VPN in Segment Routing


L3VPN in Segment Routing

- 初始状态：所有节点运行LDP, 而不运行SR
- 步骤1：所有节点都升级到支持SR
 - 没有特殊的顺序要求
 - 仍旧默认优先LDP标签压入
- 步骤2：所有PE都配置成优先SR标签压入
 - 没有特殊的顺序要求
- 步骤3：LDP被从网络节点中移除
 - 没有特殊的顺序要求
- 最终状态：所有节点运行SR , 而不运行LDP



Lab2 - L3VPN in Segment Routing

- Enable L3VPN in SR
- Check VRF/VPNv4/BGP/MPLS table/ CEF
- Check Data Plane, LFIB
- Enable/Check LDP Neighbor, and check MPLS Table
- Check Data Plane for SR/LDP
- Enable SR prefer on PE6, check RIB/CEF, trace at PE6 and PE4.



Enable L3VPN

```
route-policy RR_IN
  pass
end-policy
!
route-policy RR_OUT
  pass
end-policy

vrf A address-family ipv4 unicast
import route-target 1:1
export route-target 1:1
interface Loopback101
vrf A
  ipv4 address 172.16.1.6 255.255.255.255

router bgp 1
  bgp router-id 192.168.0.6
  address-family vpnv4 unicast
  neighbor 192.168.0.4
    remote-as 1
    description iBGP peer PE4
    update-source Loopback0
    address-family vpnv4 unicast
      route-policy RR_IN in
      route-policy RR_OUT out
      next-hop-self
    !
  vrf A
    rd 1:1
    address-family ipv4 unicast
    redistribute connected
  !
```

PE6

```
route-policy RR_IN
  pass
end-policy
!
route-policy RR_OUT
  pass
end-policy

vrf A address-family ipv4 unicast
import route-target 1:1
export route-target 1:1
interface Loopback101
vrf A
  ipv4 address 172.16.1.4 255.255.255.255

router bgp 1
  bgp router-id 192.168.0.4
  address-family vpnv4 unicast
  neighbor 192.168.0.6
    remote-as 1
    description iBGP peer PE6
    update-source Loopback0
    address-family vpnv4 unicast
      route-policy RR_IN in
      route-policy RR_OUT out
      next-hop-self
    !
  vrf A
    rd 1:2
    address-family ipv4 unicast
    redistribute connected
  !
```

PE4

Check VPN Info

- show route vrf A
- show bgp vpnv4 un sum
- show bgp vpnv4 un
- show bgp vpnv4 un label
- show bgp vrf A 172.16.1.4/32
- show cef vrf A 172.16.1.4/32
- trace vrf A 172.16.1.4
- show mpls label table
- show mpls for

```
RP/0/0/CPU0:PE6#sh route vrf A
Wed Apr 10 03:53:45.859 UTC
```

Codes: C - connected, S - static, R - RIP, B - BGP, (>) - Diversion path
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
i - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, su - IS-IS summary null, * - candidate default
U - per-user static route, o - ODR, L - local, G - DAGR, l - LISP
A - access/subscriber, a - Application route
M - mobile route, r - RPL, t - Traffic Engineering, (!) - FRR Backup
path

Gateway of last resort is not set

B 172.16.1.4/32 [200/0] via 192.168.0.4 (nexthop in vrf default), 00:16:46
L 172.16.1.6/32 is directly connected, 00:18:16, Loopback101

```
RP/0/0/CPU0:PE6#sh bgp vpnv4 un sum
Wed Apr 10 03:37:09.267 UTC
BGP router identifier 192.168.0.6, local AS number 1
.....
BGP is operating in STANDALONE mode.
```

Process Speaker	RcvTblVer	bRIB/RIB	LabelVer	ImportVer	SendTblVer	StandbyVer			
	8	8	8	8	8	0			
Neighbor 192.168.0.4	Spk 0	AS 1	MsgRcvd 5	MsgSent 5	TblVer 8	InQ 0	OutQ 0	Up/Down 00:01:09	St/PfxRcd 1

```
RP/0/0/CPU0:PE6#sh bgp vpnv4 un
Wed Apr 10 03:38:35.471 UTC
BGP router identifier 192.168.0.6, local AS number 1
.....
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
Network          Next Hop          Metric LocPrf Weight Path
Route Distinguisher: 1:1 (default for vrf A)
*>i172.16.1.4/32    192.168.0.4        0     100      0 ?
*> 172.16.1.6/32      0.0.0.0        0           32768 ?
Route Distinguisher: 1:2
*>i172.16.1.4/32    192.168.0.4        0     100      0 ?
Processed 3 prefixes, 3 paths
```

```
RP/0/0/CPU0:PE6#sh bgp vpnv4 un labels
Wed Apr 10 03:43:25.241 UTC
BGP router identifier 192.168.0.6, local AS number 1
.....
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
Network          Next Hop          Rcvd Label   Local Label
Route Distinguisher: 1:1 (default for vrf A)
*>i172.16.1.4/32    192.168.0.4        24002      nolabel
*> 172.16.1.6/32      0.0.0.0        nolabel      24008
Route Distinguisher: 1:2
*>i172.16.1.4/32    192.168.0.4        24002      nolabel
Processed 3 prefixes, 3 paths
```

```
RP/0/0/CPU0:PE6#sh bgp vrf A 172.16.1.4/32
```

```
Wed Apr 10 03:45:54.621 UTC
```

```
BGP routing table entry for 172.16.1.4/32, Route Distinguisher: 1:1
```

```
Versions:
```

```
  Process          bRIB/RIB  SendTblVer
```

```
  Speaker           8          8
```

```
Last Modified: Apr 10 03:36:59.552 for 00:08:55
```

```
Paths: (1 available, best #1)
```

```
  Not advertised to any peer
```

```
  Path #1: Received by speaker 0
```

```
  Not advertised to any peer
```

```
  Local
```

```
    192.168.0.4 (metric 2) from 192.168.0.4 (192.168.0.4)
```

```
      Received Label 24002
```

```
      Origin incomplete, metric 0, localpref 100, valid, internal, best,  
      group-best, import-candidate, imported
```

```
      Received Path ID 0, Local Path ID 1, version 8
```

```
      Extended community: RT:1:1
```

```
      Source AFI: VPNv4 Unicast, Source VRF: default, Source Route
```

```
Distinguisher: 1:2
```

```
RP/0/0/CPU0:PE6#trace vrf A 172.16.1.4
```

```
Wed Apr 10 04:03:58.187 UTC
```

```
Type escape sequence to abort.
```

```
Tracing the route to 172.16.1.4
```

```
1 16.1.1.1 [MPLS: Labels 16004/24002 Exp 0] 19 msec 9 msec 9 msec
```

```
2 14.1.1.4 9 msec * 9 msec
```

```
RP/0/0/CPU0:PE6#sh cef vrf A 172.16.1.4/32 detail
```

```
Wed Apr 10 03:58:08.881 UTC
```

```
172.16.1.4/32, version 3, internal 0x5000001 0x0 (ptr 0xa11f80e4) [1], 0x0  
(0x0), 0x208 (0xa15d139c)
```

```
  Updated Apr 10 03:36:59.928
```

```
  Prefix Len 32, traffic index 0, precedence n/a, priority 3  
  gateway array (0xa1171d8) reference count 1, flags 0x2038, source rib (7),  
  0 backups
```

```
    [1 type 1 flags 0x48441 (0xa15ed558) ext 0x0 (0x0)]
```

```
    LW-LDI[type=0, refc=0, ptr=0x0, sh-ldi=0x0]
```

```
    gateway array update type-time 1 Apr 10 03:36:59.928
```

```
    LDI Update time Apr 10 03:36:59.928
```

```
    via 192.168.0.4/32, 3 dependencies, recursive [flags 0x6000]
```

```
    path-idx 0 NHID 0x0 [0xa1646058 0x0]
```

```
    recursion-via-/32
```

```
    next hop VRF - 'default', table - 0xe0000000
```

```
    next hop 192.168.0.4/32 via 16004/0/21
```

```
      next hop 16.1.1.1/32 Gi0/0/0/0      labels imposed {16004 24002}
```

```
Load distribution: 0 (refcount 1)
```

Hash	OK	Interface	Address
0	Y	Unknown	16004/0

```

RP/0/0/CPU0:PE6#sh mpls label table
Wed Apr 10 04:05:35.350 UTC
Table Label Owner State Rewrite
-----
0 0 LSD(A) InUse Yes
0 1 LSD(A) InUse Yes
0 2 LSD(A) InUse Yes
0 13 LSD(A) InUse Yes
0 16000 ISIS(A):srte InUse No
0 24000 ISIS(A):srte InUse Yes
0 24001 ISIS(A):srte InUse Yes
0 24002 ISIS(A):srte InUse Yes
0 24003 ISIS(A):srte InUse Yes
0 24004 ISIS(A):srte InUse Yes
0 24005 ISIS(A):srte InUse Yes
0 24006 ISIS(A):srte InUse Yes
0 24007 ISIS(A):srte InUse Yes
0 24008 BGP-VPNv4(A):bgp-default InUse Yes

```

```

RP/0/0/CPU0:PE6#sh mpls for
Wed Apr 10 04:05:58.939 UTC
Local Outgoing Prefix Outgoing Next Hop Bytes
Label Label or ID Interface Switched
-----
16001 Pop SR Pfx (idx 1) Gi0/0/0/0 16.1.1.1 0
16004 16004 SR Pfx (idx 4) Gi0/0/0/0 16.1.1.1 3978
16007 Pop SR Pfx (idx 7) Gi0/0/0/1 67.1.1.7 0
17001 Pop SR Pfx (idx 1001) Gi0/0/0/0 fe80::f816:3eff:fe48:63af \
0
17004 17004 SR Pfx (idx 1004) Gi0/0/0/0 fe80::f816:3eff:fe48:63af \
0
17007 Pop SR Pfx (idx 1007) Gi0/0/0/1 fe80::f816:3eff:fe25:8cf \
0
24000 Pop SR Adj (idx 1) Gi0/0/0/0 16.1.1.1 0
24001 Pop SR Adj (idx 3) Gi0/0/0/0 16.1.1.1 0
24002 Pop SR Adj (idx 1) Gi0/0/0/1 67.1.1.7 0
24003 Pop SR Adj (idx 3) Gi0/0/0/1 67.1.1.7 0
24004 Pop SR Adj (idx 1) Gi0/0/0/0 fe80::f816:3eff:fe48:63af \
0
24005 Pop SR Adj (idx 3) Gi0/0/0/0 fe80::f816:3eff:fe48:63af \
0
24006 Pop SR Adj (idx 1) Gi0/0/0/1 fe80::f816:3eff:fe25:8cf \
0
24007 Pop SR Adj (idx 3) Gi0/0/0/1 fe80::f816:3eff:fe25:8cf \
0
24008 Aggregate A: Per-VRF Aggr[V] A 728

```

Enable LDP

```
mpls ldp
router-id 192.168.0.6
log
hello-adjacency
neighbor
address-family ipv4 label local allocate for host-routes
interface GigabitEthernet0/0/0/0
interface GigabitEthernet0/0/0/1
```

PE6

```
mpls ldp
router-id 192.168.0.7
log
hello-adjacency
neighbor
address-family ipv4 label local allocate for host-routes
interface GigabitEthernet0/0/0/0
interface GigabitEthernet0/0/0/1
interface GigabitEthernet0/0/0/2
```

P7

```
mpls ldp
router-id 192.168.0.1
log
hello-adjacency
neighbor
address-family ipv4 label local allocate for host-routes
interface GigabitEthernet0/0/0/1
interface GigabitEthernet0/0/0/2
interface GigabitEthernet0/0/0/3
```

P1

```
mpls ldp
router-id 192.168.0.4
log
hello-adjacency
neighbor
address-family ipv4 label local allocate for host-routes
interface GigabitEthernet0/0/0/0
interface GigabitEthernet0/0/0/1
```

PE4

Check data plane again

- show mpls for
- show cef vrf A 172.16.1.4/32
- show mpls label table
- show mpls ldp bin 192.168.0.4/32
- trace 192.168.0.4
- trace vrf A 172.16.1.4
- trace mpls ipv4 192.168.0.4/32
- show route 192.168.0.4/32 det

```
RP/0/0/CPU0:PE6#show mpls label table
Wed Apr 10 04:40:14.028 UTC
```

Table	Label	Owner	State	Rewrite
0	0	LSD(A)	InUse	Yes
0	1	LSD(A)	InUse	Yes
0	2	LSD(A)	InUse	Yes
0	13	LSD(A)	InUse	Yes
0	16000	ISIS(A):srte	InUse	No
0	24000	ISIS(A):srte	InUse	Yes
0	24001	ISIS(A):srte	InUse	Yes
0	24002	ISIS(A):srte	InUse	Yes
0	24003	ISIS(A):srte	InUse	Yes
0	24004	ISIS(A):srte	InUse	Yes
0	24005	ISIS(A):srte	InUse	Yes
0	24006	ISIS(A):srte	InUse	Yes
0	24007	ISIS(A):srte	InUse	Yes
0	24008	BGP-VPNv4(A):bgp-default	InUse	Yes
0	24009	LDP(A)	InUse	Yes
0	24010	LDP(A)	InUse	Yes
0	24011	LDP(A)	InUse	Yes
0	24012	LDP(A)	InUse	Yes

```
RP/0/0/CPU0:PE6#sh mpls ldp bindings 192.168.0.4/32
```

```
Wed Apr 10 04:43:52.193 UTC
```

```
192.168.0.4/32, rev 10
```

```
Local binding: label: 24012
```

```
Remote bindings: (2 peers)
```

Peer	Label
192.168.0.1:0	24012
192.168.0.7:0	24004

```
RP/0/0/CPU0:PE6#sh mpls for
Wed Apr 10 04:38:20.336 UTC
```

Local Label	Outgoing Label	Prefix or ID	Outgoing Interface	Next Hop	Bytes Switched
16001	Pop	SR Pfx (idx 1)	Gi0/0/0/0	16.1.1.1	0
16004	16004	SR Pfx (idx 4)	Gi0/0/0/0	16.1.1.1	321
	16004	SR Pfx (idx 4)	Gi0/0/0/1	67.1.1.7	0
16007	Pop	SR Pfx (idx 7)	Gi0/0/0/1	67.1.1.7	0
17001	Pop	SR Pfx (idx 1001)	Gi0/0/0/0	fe80::f816:3eff:fe48:63af \	0
17004	17004	SR Pfx (idx 1004)	Gi0/0/0/0	fe80::f816:3eff:fe48:63af \	0
	17004	SR Pfx (idx 1004)	Gi0/0/0/1	fe80::f816:3eff:fe25:8cf \	0
17007	Pop	SR Pfx (idx 1007)	Gi0/0/0/1	fe80::f816:3eff:fe25:8cf \	0
24000	Pop	SR Adj (idx 1)	Gi0/0/0/0	16.1.1.1	0
24001	Pop	SR Adj (idx 3)	Gi0/0/0/0	16.1.1.1	0
24002	Pop	SR Adj (idx 1)	Gi0/0/0/1	67.1.1.7	0
24003	Pop	SR Adj (idx 3)	Gi0/0/0/1	67.1.1.7	0
24004	Pop	SR Adj (idx 1)	Gi0/0/0/0	fe80::f816:3eff:fe48:63af \	0
24005	Pop	SR Adj (idx 3)	Gi0/0/0/0	fe80::f816:3eff:fe48:63af \	0
24006	Pop	SR Adj (idx 1)	Gi0/0/0/1	fe80::f816:3eff:fe25:8cf \	0
24007	Pop	SR Adj (idx 3)	Gi0/0/0/1	fe80::f816:3eff:fe25:8cf \	0
24008	Aggregate	A: Per-VRF Aggr[V] A			728
24009	24006	192.168.0.9/32	Gi0/0/0/0	16.1.1.1	0
24010	Pop	192.168.0.1/32	Gi0/0/0/0	16.1.1.1	867
24011	Pop	192.168.0.7/32	Gi0/0/0/1	67.1.1.7	1291
24012	24012	192.168.0.4/32	Gi0/0/0/0	16.1.1.1	602
	24004	192.168.0.4/32	Gi0/0/0/1	67.1.1.7	0

```

RP/0/0/CPU0:PE6#sh cef 192.168.0.4/32 detail
Wed Apr 10 05:07:51.844 UTC
192.168.0.4/32, version 52, labeled SR, internal 0x1000001 0x81 (ptr
0xa11f7ae0) [1], 0x0 (0xa11da468), 0xa28 (0xa1a870b8)
Updated Apr 10 05:05:11.365
local adjacency 16.1.1.1
Prefix Len 32, traffic index 0, precedence n/a, priority 3
Extensions: context-label:16004
    gateway array (0xa11176c4) reference count 3, flags 0x68, source lsd
(5), 1 backups
        [2 type 5 flags 0x8401 (0xa15eda18) ext 0x0 (0x0)]
    LW-LDI[type=5, refc=3, ptr=0xa11da468, sh-ldi=0xa15eda18]
    gateway array update type-time 1 Apr 10 05:05:11.365
LDI Update time Apr 10 05:05:11.365
LW-LDI-TS Apr 10 05:05:11.365
    via 16.1.1.1/32, GigabitEthernet0/0/0/0, 13 dependencies, weight 0,
class 0 [flags 0x0]
        path-idx 0 NHID 0x0 [0xa189b1c0 0x0]
        next hop 16.1.1.1/32
        local adjacency
            local label 24012      labels imposed {24012}
        via 67.1.1.7/32, GigabitEthernet0/0/0/1, 11 dependencies, weight 0,
class 0 [flags 0x0]
        path-idx 1 NHID 0x0 [0xa189b270 0x0]
        next hop 67.1.1.7/32
        local adjacency
            local label 24012      labels imposed {24004}

```

```

RP/0/0/CPU0:PE6#show route 192.168.0.4/32 det
Wed Apr 10 04:50:59.534 UTC

Routing entry for 192.168.0.4/32
Known via "isis srte", distance 115, metric 2, labeled SR, type level-2
Installed Apr 10 04:36:20.984 for 00:14:38
Routing Descriptor Blocks
    16.1.1.1, from 192.168.0.4, via GigabitEthernet0/0/0/0
        Route metric is 2
        Label: 0x3e84 (16004)
        Tunnel ID: None
        Binding Label: None
        Extended communities count: 0
        Path id:2      Path ref count:0
        NHID:0x1(Ref:7)
    67.1.1.7, from 192.168.0.4, via GigabitEthernet0/0/0/1
        Route metric is 2
        Label: 0x3e84 (16004)
        Tunnel ID: None
        Binding Label: None
        Extended communities count: 0
        Path id:1      Path ref count:0
        NHID:0x2(Ref:5)
Route version is 0x11 (17)
Local Label: 0x3e84 (16004)
IP Precedence: Not Set
QoS Group ID: Not Set
Flow-tag: Not Set
Fwd-class: Not Set
Route Priority: RIB_PRIORITY_NON_RECURSIVE_MEDIUM (7) SVD Type
RIB_SVD_TYPE_LOCAL
Download Priority 1, Download Version 54
No advertising protos.

```

```
RP/0/0/CPU0:PE6#trace 192.168.0.4
Wed Apr 10 04:48:41.873 UTC
```

```
Type escape sequence to abort.
Tracing the route to 192.168.0.4
```

```
1 67.1.1.7 [MPLS: Label 24004 Exp 0] 19 msec 9 msec 9 msec
2 47.1.1.4 9 msec * 9 msec
```

```
RP/0/0/CPU0:PE6#!
```

```
RP/0/0/CPU0:PE6#trace vrf A 172.16.1.4
Wed Apr 10 04:49:50.988 UTC
```

```
Type escape sequence to abort.
Tracing the route to 172.16.1.4
```

```
1 16.1.1.1 [MPLS: Labels 24012/24002 Exp 0] 9 msec 9 msec 9 msec
2 14.1.1.4 9 msec * 9 msec
```

```
RP/0/0/CPU0:PE6#trace mpls ipv4 192.168.0.4/32
Wed Apr 10 04:54:47.998 UTC
```

```
Tracing MPLS Label Switched Path to 192.168.0.4/32, timeout is 2
seconds
```

```
Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
'L' - labeled output interface, 'B' - unlabeled output interface,
'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
'M' - malformed request, 'm' - unsupported tlvs, 'N' - no rx label,
'P' - no rx intf label prot, 'p' - premature termination of LSP,
'R' - transit router, 'I' - unknown upstream index,
'X' - unknown return code, 'x' - return code 0
```

```
Type escape sequence to abort.
```

```
0 67.1.1.6 MRU 1500 [Labels: 24004 Exp: 0]
L 1 67.1.1.7 MRU 1500 [Labels: implicit-null Exp: 0] 20 ms
! 2 47.1.1.4 60 ms
```

```
RP/0/0/CPU0:PE6#sh cef vrf A 172.16.1.4/32 detail
Wed Apr 10 05:05:26.714 UTC
172.16.1.4/32, version 5, internal 0x5000001 0x0 (ptr 0xa11f7bf8) [1], 0x0
(0x0), 0x208 (0xa15d1268)
Updated Apr 10 05:05:08.065
Prefix Len 32, traffic index 0, precedence n/a, priority 3
gateway array (0xa1116abc) reference count 1, flags 0x2038, source rib (7),
0 backups
[1 type 1 flags 0x48441 (0xa15ed318) ext 0x0 (0x0)]
LW-LDI[type=0, refc=0, ptr=0x0, sh-ldi=0x0]
gateway array update type-time 1 Apr 10 05:05:08.065
LDI Update time Apr 10 05:05:08.065
via 192.168.0.4/32, 3 dependencies, recursive [flags 0x6000]
path-idx 0 NHID 0x0 [0xa16455f4 0x0]
recursion-via-/32
next hop VRF - 'default', table - 0xe0000000
next hop 192.168.0.4/32 via 24012/0/21
next hop 16.1.1.1/32 Gi0/0/0/0 labels imposed {24012 24002}
next hop 67.1.1.7/32 Gi0/0/0/1 labels imposed {24004 24002}
```

```
Load distribution: 0 (refcount 1)
```

Hash	OK	Interface	Address
0	Y	Unknown	24012/0

Enable sr-prefer at PE6, then check data plane

```
router isis srte address-family ipv4 unicast segment-routing mpls sr-prefer PE6
```

- show mpls for
- show mpls label table
- show mpls ldp bin 192.168.0.4/32
- show route vrf A 172.16.1.4/32 det
- show cef vrf A 172.16.1.4/32 det
- show route 192.168.0.4/32 det
- show cef 192.168.0.4/32 flags
- trace vrf A 172.16.1.4/6

```

RP/0/0/CPU0:PE6#show cef 192.168.0.4/32 detail
Wed Apr 10 05:18:13.912 UTC
192.168.0.4/32, version 107, labeled SR, internal 0x1000001 0x83 (ptr
0xa11f7ae0) [1], 0x0 (0xa11da670), 0xa28 (0xa1a870e8)
    Updated Feb 10 03:06:27.329
    local adjacency 16.1.1.1
    Prefix Len 32, traffic index 0, precedence n/a, priority 1
        gateway array (0xa1117750) reference count 3, flags 0x68, source rib
(7), 1 backups
            [ 2 type 5 flags 0x8401 (0xa15ed498) ext 0x0 (0x0) ]
            LW-LDI[type=5, refc=3, ptr=0xa11da670, sh-ldi=0xa15ed498]
            gateway array update type-time 1 Apr 10 05:17:49.294
LDI Update time Apr 10 05:17:49.294
LW-LDI-TS Apr 10 05:17:49.294
    via 16.1.1.1/32, GigabitEthernet0/0/0/0, 11 dependencies, weight 0,
class 0 [flags 0x0]
        path-idx 0 NHID 0x0 [0xa189b1c0 0x0]
        next hop 16.1.1.1/32
        local adjacency
            local label 16004      labels imposed {16004}
    via 67.1.1.7/32, GigabitEthernet0/0/0/1, 9 dependencies, weight 0,
class 0 [flags 0x0]
        path-idx 1 NHID 0x0 [0xa189b270 0x0]
        next hop 67.1.1.7/32
        local adjacency
            local label 16004      labels imposed {16004}

Load distribution: 0 1 (refcount 2)

Hash OK Interface          Address
0   Y  GigabitEthernet0/0/0  16.1.1.1
1   Y  GigabitEthernet0/0/1  67.1.1.7

```

RP/0/0/CPU0:PE6#show cef 192.168.0.4/32 flags
Wed Apr 10 05:44:26.984 UTC
192.168.0.4/32, version 107, labeled SR, internal 0x1000001 0x83 (ptr
0xa11f7ae0) [1], 0x0 (0xa11da670), 0xa28 (0xa1a870e8)
leaf flags: owner locked, inserted
leaf flags2: LDP/SR merge requested,RIB pref over LSD,sr-pfx,
leaf ext flags:
PriChange,EXTERNAL_REACH_LC,L2TPV3_SPAN_DIAG_IFH_ENABLE,illegal-0x00000800,
Updated Feb 10 03:06:27.329
local adjacency 16.1.1.1
Prefix Len 32, traffic index 0, precedence n/a, priority 1
.....

RP/0/0/CPU0:PE6#traceroute vrf A 172.16.1.4
Wed Apr 10 05:21:46.367 UTC

PE6

Type escape sequence to abort.
Tracing the route to 172.16.1.4

1 16.1.1.1 [MPLS: Labels **16004/24002** Exp 0] 19 msec 9 msec 9 msec
2 14.1.1.4 9 msec * 9 msec

RP/0/0/CPU0:PE4#trace vrf A 172.16.1.6
Wed Apr 10 05:24:04.927 UTC

PE4

Type escape sequence to abort.
Tracing the route to 172.16.1.6

1 47.1.1.7 [MPLS: Labels **24000/24008** Exp 0] 69 msec 9 msec 9 msec
2 67.1.1.6 19 msec * 19 msec


*Even if changed sr-pre at head-end, not affect
remote PE's forwarding before removing LDP*



Interworking between LDP and SR


Interworking between LDP and SR

SR to LDP




- MPLS-2-MPLS, MPLS-2-IP
- IP-2-MPLS
- Ships In The Night – done at last lab
- Interworking


LDP to SR



SR via LDP




LDP via SR




Lab3 - Interworking between LDP and SR

- Shut other ports and only remain PE6, P1 and PE4
 - Ping test when remove SR config on P1
 - Check labels forwarding status
-
- Change config to: R6 SR, R4 LDP, P1 enable SR/LDP
 - Check labels allocate status; check global traffics from R4 -> R1 -> R6(MPLS) and R6 -> R1 -> R4(IP), VPN traffics failure, why?
 - Enable SRMS on R1, check forwarding status, at final, remove SRMS



SR/LDP(sr-pre) <-> LDP <-> LDP/SR

- show cef vrf A 172.16.1.4/32
- show cef 192.168.0.4/32
- show mpls for prefix 192.168.0.4/32
- show mpls ldp bin 192.168.0.4/32



```
int gi0/0/0/1 shutdown PE6  
int gi0/0/0/0 shutdown PE4
```

```
int gi0/0/0/3 shutdown  
no router isis srte address-family ipv4 unicast segment-routing mpls  
no router isis srte address-family ipv6 unicast segment-routing mpls  
no router isis srte interface Loopback0 address-family ipv4 unicast prefix-sid index 1  
no router isis srte interface Loopback0 address-family ipv6 unicast prefix-sid absolute 17001
```

```

RP/0/0/CPU0:PE6#show cef vrf A 172.16.1.4/32 detail
Wed Apr 10 06:26:04.243 UTC
172.16.1.4/32, version 5, internal 0x5000001 0x0 (ptr 0xa11f7bf8) [1], 0x0
(0x0), 0x208 (0xa15d1268)
Updated Apr 10 05:05:08.066
Prefix Len 32, traffic index 0, precedence n/a, priority 3
  gateway array (0xa1116abc) reference count 1, flags 0x2038, source rib
(7), 0 backups
  [1 type 1 flags 0x40441 (0xa15ed318) ext 0x0 (0x0)]
LW-LDI[type=0, refc=0, ptr=0x0, sh-ldi=0x0]
  gateway array update type-time 5 Apr 10 05:17:49.294
LDI Update time Apr 10 05:17:49.294
via 192.168.0.4/32, 3 dependencies, recursive [flags 0x6000]
  path-idx 0 NHID 0x0 [0xa16454dc 0x0]
  recursion-via-/32
next hop VRF - 'default', table - 0xe0000000
next hop 192.168.0.4/32 via 16004/0/21
  next hop 16.1.1.1/32 Gi0/0/0/0    labels imposed {24012 24002}

Load distribution: 0 (refcount 1)

Hash OK Interface Address
0 Y Unknown 16004/0

```

PE6

```

RP/0/0/CPU0:PE6#sh mpls ldp bindings 192.168.0.4/32
Wed Apr 10 06:30:02.237 UTC
192.168.0.4/32, rev 10
Local binding: label: 24012
  Remote bindings: (1 peers)
    Peer Label
    ----- -----
192.168.0.1:0 24012

```

PE6

RP/0/0/CPU0:PE6#sh cef 192.168.0.4/32 fla detail

Wed Apr 10 07:09:46.123 UTC

192.168.0.4/32, version 141, labeled SR, internal 0x100001 0x87 (ptr 0xa11f7ae0) [1], 0x0 (0xa11da670), 0xa28 (0xa15d14a4)

leaf flags: owner locked, inserted

leaf flags2: LDP/SR merge requested, **RIB pref over LSD,LDP/SR merge active,sr-pfx**, leaf ext flags:

PriChange,EXTERNAL_REACH_LC,L2TPV3_SPAN_DIAG_IFH_ENABLE,illegal-0x00000800, Updated Feb 10 03:06:27.328

local adjacency 16.1.1

Prefix Len 32, traffic index 0, precedence n/a, priority 15

gateway array (0xa11178f4) reference count 3, flags 0x68, source rib (7), 2 backups

[2 type 5 flags 0x8401 (0xa15edb18) ext 0x0 (0x0)]

LW-LDI[type=5, refc=3, ptr=0xa11da670, sh-ldi=0xa15edb18]

gateway array update type-time 1 Apr 10 06:03:16.006

LDI Update time Apr 10 06:03:16.006

LW-LDI-TS Apr 10 06:07:34.968

pl flags: mpls ext, imp, lsw

ldi flags: owner locked, added to pl, depth changed

via 16.1.1.1/32, GigabitEthernet0/0/0/0, 11 dependencies, weight 0, class 0 [flags 0x0]

path-idx 0 NHID 0x0 [0xa189b1c0 0x0]

next hop 16.1.1.1/32

local adjacency

local label 16004 labels imposed {24012}

*Auto mapping, not need
any extra configuration*

Load distribution: 0 (refcount 2)

Hash	OK	Interface	Address
0	Y	GigabitEthernet0/0/0/0	16.1.1.1

```
RP/0/0/CPU0:PE6#sh route 192.168.0.4/32 detail  
Wed Apr 10 06:32:21.567 UTC
```

PE6

```
Routing entry for 192.168.0.4/32  
Known via "isis srte", distance 115, metric 2, labeled SR, type level-  
2  
Installed Apr 10 06:07:34.949 for 00:24:46  
Routing Descriptor Blocks  
    16.1.1.1, from 192.168.0.4, via GigabitEthernet0/0/0/0  
        Route metric is 2  
        Label: None  
        Tunnel ID: None  
        Binding Label: None  
        Extended communities count: 0  
        Path id:1      Path ref count:0  
        NHID:0x1(Ref:6)  
Route version is 0x26 (38)  
Local Label: 0x3e84 (16004)  
IP Precedence: Not Set  
QoS Group ID: Not Set  
Flow-tag: Not Set  
Fwd-class: Not Set  
Route Priority: RIB_PRIORITY_NON_RECURSIVE_MEDIUM (7) SVD Type  
RIB_SVD_TYPE_LOCAL  
Download Priority 1, Download Version 141  
No advertising protos.
```

```
RP/0/0/CPU0:PE6#traceroute vrf A 172.16.1.4  
Wed Apr 10 07:24:24.483 UTC
```

```
Type escape sequence to abort.  
Tracing the route to 172.16.1.4
```

```
1 16.1.1.1 [MPLS: Labels 24012/24002 Exp 0] 19 msec 9 msec 0 msec  
2 14.1.1.4 9 msec * 0 msec
```

```
RP/0/0/CPU0:P1#sh mpls forwarding labels 24012  
Wed Apr 10 06:37:33.555 UTC
```

Local Label	Outgoing Label	Prefix or ID	Outgoing Interface	Next Hop	Bytes Switched
24012	Pop	192.168.0.4/32	Gi0/0/0/2	14.1.1.4	12590

```
RP/0/0/CPU0:P1#sh mpls ldp bin 192.168.0.4/32  
Wed Apr 10 07:21:30.584 UTC
```

192.168.0.4/32, rev 8	Local binding: label: 24012
Remote bindings: (2 peers)	
Peer	Label
192.168.0.4:0	ImpNull
192.168.0.6:0	24012

P1

P1

```

RP/0/0/CPU0:PE4#sh cef vrf A 172.16.1.6/32 detail
Wed Apr 10 07:27:21.831 UTC
172.16.1.6/32, version 3, internal 0x5000001 0x0 (ptr 0xa124d65c) [1], 0x0
(0x0), 0x208 (0xa16b123c)
Updated Apr 10 03:36:57.758
Prefix Len 32, traffic index 0, precedence n/a, priority 3
  gateway array (0xa114b638) reference count 1, flags 0x2038, source rib
(7), 0 backups
  [1 type 1 flags 0x40441 (0xa16cd3d8) ext 0x0 (0x0)]
LW-LDI[type=0, refc=0, ptr=0x0, sh-ldi=0x0]
  gateway array update type-time 5 Apr 10 04:35:55.416
LDI Update time Apr 10 04:35:55.416
via 192.168.0.6/32, 3 dependencies, recursive [flags 0x6000]
  path-idx 0 NHID 0x0 [0xa172642c 0x0]
  recursion-via-/32
next hop VRF - 'default', table - 0xe0000000
next hop 192.168.0.6/32 via 24007/0/21
  next hop 14.1.1.1/32 Gi0/0/0/1    labels imposed {24007 24008}

Load distribution: 0 (refcount 1)

Hash OK Interface Address
0 Y Unknown 24007/0

```

PE4

```

RP/0/0/CPU0:PE4#sh mpls ldp bin 192.168.0.6/32
Wed Apr 10 07:28:29.596 UTC
192.168.0.6/32, rev 9
  Local binding: label: 24007
  Remote bindings: (1 peers)
    Peer Label
    -----
    192.168.0.1:0 24007

```

PE4

```

RP/0/0/CPU0:PE4#show cef 192.168.0.6/32 fla de
Wed Apr 10 07:29:37.301 UTC
192.168.0.6/32, version 43, labeled SR, internal 0x1000001 0x81 (ptr
0xa124cf40) [1], 0x0 (0xa122f990), 0xa28 (0xa16b1370)
  leaf flags: owner locked, inserted
  leaf flags2: LDP/SR merge requested,sr-pfx, <<< no active
  leaf ext flags:
PriChange,EXTERNAL_REACH_LC,L2TPV3_SPAN_DIAG_IFH_ENABLE,illegal-0x00000800,
  Updated Apr 10 06:07:32.818
  local adjacency 14.1.1.1
  Prefix Len 32, traffic index 0, precedence n/a, priority 3
  Extensions: context-label:16006
  gateway array (0xa114b8f4) reference count 3, flags 0x68, source lsd (5),
1 backups
  [2 type 5 flags 0x8401 (0xa16cd798) ext 0x0 (0x0)]
LW-LDI[type=5, refc=3, ptr=0xa122f990, sh-ldi=0xa16cd798]
  gateway array update type-time 1 Apr 10 06:03:13.916
LDI Update time Apr 10 06:03:13.916
LW-LDI-TS Apr 10 06:03:13.916
  pl flags: mpls ext, imp, lsw

ldi flags: owner locked, added to pl, depth changed

via 14.1.1.1/32, GigabitEthernet0/0/0/1, 11 dependencies, weight 0, class
0 [flags 0x0]
  path-idx 0 NHID 0x0 [0xa189b320 0x0]
  next hop 14.1.1.1/32
  local adjacency
    local label 24007      labels imposed {24007}
```

Load distribution: 0 (refcount 2)

Hash	OK	Interface	Address
0	Y	GigabitEthernet0/0/0/1	14.1.1.1

```
RP/0/0/CPU0:PE4#sh route 192.168.0.6/32 detail  
Wed Apr 10 07:35:30.867 UTC
```

PE4

```
Routing entry for 192.168.0.6/32  
Known via "isis srte", distance 115, metric 2, labeled SR, type level-  
2  
Installed Apr 10 06:07:32.779 for 01:27:58  
Routing Descriptor Blocks  
 14.1.1.1, from 192.168.0.6, via GigabitEthernet0/0/0/1  
  Route metric is 2  
  Label: None  
  Tunnel ID: None  
  Binding Label: None  
  Extended communities count: 0  
  Path id:1      Path ref count:0  
  NHID:0x1(Ref:6)  
Route version is 0x21 (33)  
Local Label: 0x3e86 (16006)  
IP Precedence: Not Set  
QoS Group ID: Not Set  
Flow-tag: Not Set  
Fwd-class: Not Set  
Route Priority: RIB_PRIORITY_NON_RECURSIVE_MEDIUM (7) SVD Type  
RIB_SVD_TYPE_LOCAL  
Download Priority 1, Download Version 105  
No advertising protos.
```

```
RP/0/0/CPU0:P1#sh mpls for labels 24007  
Wed Apr 10 07:36:58.711 UTC
```

Local Label	Outgoing Label	Prefix or ID	Outgoing Interface	Next Hop	Bytes Switched
24007	Pop	192.168.0.6/32	Gi0/0/0/1	16.1.1.6	18549

```
RP/0/0/CPU0:P1#sh mpls ldp bindings 192.168.0.6/32  
Wed Apr 10 07:37:37.788 UTC
```

192.168.0.6/32, rev 7	Local binding: label: 24007
	Remote bindings: (2 peers)
	Peer Label
192.168.0.4:0	24007
192.168.0.6:0	ImpNull

```
RP/0/0/CPU0:PE4#traceroute vrf A 172.16.1.6  
Wed Apr 10 07:36:29.603 UTC
```

PE4

```
Type escape sequence to abort.  
Tracing the route to 172.16.1.6
```


```
1 14.1.1.1 [MPLS: Labels 24007/24008 Exp 0] 9 msec 9 msec 0 msec  
2 16.1.1.6 9 msec * 9 msec
```

P1

P1

SR <-> SR/LDP <-> LDP

- trace/ping 192.168.0.4
- ping vrf A 172.16.1.4 ?
- trace mpls ipv4 192.168.0.6/32
- show cef vrf A 172.16.1.4/32
- show isis segment-routing label ta



```
no mpls ldp  
PE6  
rollback configuration last 1 P1
```

```
no router isis srte address-family ipv4 unicast segment-routing mpls  
no router isis srte address-family ipv6 unicast segment-routing mpls  
no router isis srte interface Loopback0 address-family ipv4 unicast prefix-sid index 4  
no router isis srte interface Loopback0 address-family ipv6 unicast prefix-sid absolute 17004  
PE4
```

```

RP/0/0/CPU0:PE6#show cef vrf A 172.16.1.4/32 detail
Wed Apr 10 07:54:29.889 UTC
172.16.1.4/32, version 5, internal 0x5000001 0x0 (ptr 0xa11f7bf8) [1], 0x0
(0x0), 0x208 (0xa15d1268)
Updated Apr 10 05:05:08.065
Prefix Len 32, traffic index 0, precedence n/a, priority 3
  gateway array (0xa1116abc) reference count 1, flags 0x203a, source rib
(7), 0 backups
  [1 type 1 flags 0xc0441 (0xa15ed318) ext 0x0 (0x0)]
LW-LDI[type=0, refc=0, ptr=0x0, sh-ldi=0x0]
  gateway array update type-time 3 Apr 10 07:54:21.720
LDI Update time Apr 10 07:50:06.247
  via 192.168.0.4/32, 0 dependencies, recursive [flags 0x6000]
    path-idx 0 NHID 0x0 [0xa0d182d8 0x0]
    recursion-via-/32
  next hop VRF - 'default', table - 0xe0000000
  unresolved
  labels imposed {24002} <<< only bgp label

Load distribution: 0 (refcount 1)

Hash OK Interface Address
0 Y Unknown drop

```

PE6

```

RP/0/0/CPU0:PE6#sh mpls ldp bindings 192.168.0.4/32
Wed Apr 10 07:57:02.449 UTC
RP/0/0/CPU0:PE6#

```

PE6

PE6

```

RP/0/0/CPU0:PE6#sh cef 192.168.0.4/32 fla detail
Wed Apr 10 07:55:06.007 UTC
192.168.0.4/32, version 168, internal 0x1000001 0x0 (ptr 0xa11f7ae0) [1],
0x0 (0xa11daa30), 0x0 (0x0)
  leaf flags: owner locked, inserted

  leaf flags2:
  leaf ext flags:
  Updated Apr 10 07:50:06.248
  local adjacency 16.1.1.1
  Prefix Len 32, traffic index 0, precedence n/a, priority 1
    gateway array (0xa1117408) reference count 4, flags 0x0, source rib (7), 0
  backups
    [5 type 3 flags 0x8401 (0xa1182658) ext 0x0 (0x0)]
    LW-LDI[type=3, refc=1, ptr=0xa11daa30, sh-ldi=0xa1182658]
      gateway array update type-time 1 Apr 10 05:05:08.036
      LDI Update time Apr 10 05:05:08.036
      LW-LDI-TS Apr 10 07:50:06.248
        ldi flags: owner locked, added to pl, depth changed

        via 16.1.1.1/32, GigabitEthernet0/0/0/0, 9 dependencies, weight 0, class
0 [flags 0x0]
          path-idx 0 NHID 0x0 [0xa189b168 0x0]
          next hop 16.1.1.1/32 <<< no label info
          local adjacency

```

Load distribution: 0 (refcount 5)

Hash	OK	Interface	Address
0	Y	GigabitEthernet0/0/0/0	16.1.1.1

```

RP/0/0/CPU0:PE6#sh mpls ldp bindings 192.168.0.4/32
Wed Apr 10 07:57:02.449 UTC
RP/0/0/CPU0:PE6#sh route 192.168.0.4/32  detail
Wed Apr 10 07:57:19.858 UTC

Routing entry for 192.168.0.4/32
Known via "isis srte", distance 115, metric 2, type level-2
Installed Apr 10 07:50:06.228 for 00:07:13
Routing Descriptor Blocks
 16.1.1.1, from 192.168.0.4, via GigabitEthernet0/0/0/0
  Route metric is 2
  Label: None
  Tunnel ID: None
  Binding Label: None
  Extended communities count: 0
  Path id:1      Path ref count:0
  NHID:0x1(Ref:6)
Route version is 0x32 (50)
No local label
IP Precedence: Not Set
QoS Group ID: Not Set
Flow-tag: Not Set
Fwd-class: Not Set
Route Priority: RIB_PRIORITY_NON_RECURSIVE_MEDIUM (7) SVD Type
RIB_SVD_TYPE_LOCAL
Download Priority 1, Download Version 168
No advertising protos.

```

PE6

```

RP/0/0/CPU0:PE6#traceroute vrf A 172.16.1.4
Wed Apr 10 07:57:50.006 UTC
% FIB did not return Source Address
RP/0/0/CPU0:PE6#

```

```

RP/0/0/CPU0:P1#sh mpls for prefix 192.168.0.4/32
Wed Apr 10 07:59:47.737 UTC

```

Local Label	Outgoing Label	Prefix or ID	Outgoing Interface	Next Hop	Bytes Switched
24012	Pop	192.168.0.4/32	Gi0/0/0/2	14.1.1.4	4067

```

RP/0/0/CPU0:P1#sh mpls ldp bin 192.168.0.4/32
Wed Apr 10 08:00:05.806 UTC

```

192.168.0.4/32, rev 8	Local binding: label: 24012
	Remote bindings: (1 peers)
	Peer Label

	192.168.0.4:0 ImpNull <<< PE6 no ldp, so no label info

*Due to no SR label for 192.168.0.4,
forwarding issue from PE6 to PE4*

```

RP/0/0/CPU0:PE6#traceroute 192.168.0.4
Mon Jul 15 07:57:43.167 UTC

```

Type escape sequence to abort.
Tracing the route to 192.168.0.4

```

1 16.1.1.1 0 msec 0 msec 0 msec
2 14.1.1.4 0 msec * 0 msec

```

P1

P1

```

RP/0/0/CPU0:PE4#sh cef vrf A 172.16.1.6/32 detail
Wed Apr 10 08:10:14.275 UTC
172.16.1.6/32, version 3, internal 0x5000001 0x0 (ptr 0xa124d65c) [1], 0x0
(0x0), 0x208 (0xa16b123c)
Updated Apr 10 03:36:57.758
Prefix Len 32, traffic index 0, precedence n/a, priority 3
  gateway array (0xa114b638) reference count 1, flags 0x2038, source rib
(7), 0 backups
  [1 type 1 flags 0x40441 (0xa16cd3d8) ext 0x0 (0x0)]
LW-LDI[type=0, refc=0, ptr=0x0, sh-ldi=0x0]
  gateway array update type-time 5 Apr 10 04:35:55.416
LDI Update time Apr 10 04:35:55.416
  via 192.168.0.6/32, 3 dependencies, recursive [flags 0x6000]
    path-idx 0 NHID 0x0 [0xa172642c 0x0]
    recursion-via-/32
  next hop VRF - 'default', table - 0xe0000000
  next hop 192.168.0.6/32 via 24007/0/21
  next hop 14.1.1.1/32 Gi0/0/0/1      labels imposed {24007 24008}

Load distribution: 0 (refcount 1)

Hash OK Interface Address
0 Y Unknown 24007/0

```

PE4

```

RP/0/0/CPU0:PE4#sh mpls ldp bin 192.168.0.6/32
Wed Apr 10 08:10:51.962 UTC
192.168.0.6/32, rev 9
  Local binding: label: 24007
  Remote bindings: (1 peers)
    Peer Label
    ----- -----
    192.168.0.1:0 24007

```

PE4

PE4

```

RP/0/0/CPU0:PE4#show cef 192.168.0.6/32 fla de
Wed Apr 10 08:11:09.101 UTC
192.168.0.6/32, version 56, internal 0x1000001 0x0 (ptr 0xa124cf40) [1], 0x0
(0xa122f990), 0xa28 (0xa16b14fc)
  leaf flags: owner locked, inserted

  leaf flags2:
  leaf ext flags:
PriChange,EXTERNAL_REACH_LC,L2TPV3_SPAN_DIAG_IFH_ENABLE,illegal-0x00000800,
  Updated Apr 10 07:50:01.998
  local adjacency 14.1.1.1
  Prefix Len 32, traffic index 0, precedence n/a, priority 3
    gateway array (0xa114b1d8) reference count 6, flags 0x68, source lsd (5),
1 backups
  [3 type 5 flags 0x8401 (0xa16cd558) ext 0x0 (0x0)]
  LW-LDI[type=5, refc=3, ptr=0xa122f990, sh-ldi=0xa16cd558]
    gateway array update type-time 1 Apr 10 04:35:59.496
LDI Update time Apr 10 04:35:59.496
LW-LDI-TS Apr 10 07:50:01.998
  pl flags: mpls ext, imp, lsw

  ldi flags: owner locked, added to pl, depth changed

  via 14.1.1.1/32, GigabitEthernet0/0/0/1, 5 dependencies, weight 0, class
0 [flags 0x0]
    path-idx 0 NHID 0x0 [0xa189b320 0x0]
    next hop 14.1.1.1/32
    local adjacency
      local label 24007      labels imposed {24007}

Load distribution: 0 (refcount 3)

Hash OK Interface Address
0 Y GigabitEthernet0/0/0/1 14.1.1.1

```

```
RP/0/0/CPU0:PE4#sh route 192.168.0.6/32 detail  
Wed Apr 10 08:12:24.426 UTC
```

PE4

Routing entry for 192.168.0.6/32

Known via "isis srte", distance 115, metric 2, type level-2
Installed Apr 10 07:50:01.948 for 00:22:22
Routing Descriptor Blocks

14.1.1.1, from 192.168.0.6, via GigabitEthernet0/0/0/1

Route metric is 2

Label: None

Tunnel ID: None

Binding Label: None

Extended communities count: 0

Path id:1 Path ref count:0

NHID:0x1(Ref:6)

Route version is 0x2d (45)

No local label

IP Precedence: Not Set

QoS Group ID: Not Set

Flow-tag: Not Set

Fwd-class: Not Set

Route Priority: RIB_PRIORITY_NON_RECURSIVE_MEDIUM (7) SVD Type

RIB_SVD_TYPE_LOCAL

Download Priority 1, Download Version 132

No advertising protos.

```
RP/0/0/CPU0:PE4#traceroute 192.168.0.6  
Mon Jul 15 07:49:14.551 UTC
```

PE4

Type escape sequence to abort.

Tracing the route to 192.168.0.6

```
1 14.1.1.1 [MPLS: Label 24017 Exp 0] 9 msec 0 msec 0 msec  
2 16.1.1.6 0 msec * 9 msec
```

```
RP/0/0/CPU0:P1#sh mpls for labels 24007  
Wed Apr 10 08:13:36.550 UTC
```

Local Label	Outgoing Label	Prefix or ID	Outgoing Interface	Next Hop	Bytes Switched
24007	Pop	192.168.0.6/32	Gi0/0/0/1	16.1.1.6	1719

```
RP/0/0/CPU0:P1#sh mpls ldp bindings 192.168.0.6/32  
Wed Apr 10 08:14:11.028 UTC
```

192.168.0.6/32, rev 7

Local binding: label: 24007

Remote bindings: (1 peers)

Peer	Label
------	-------

192.168.0.4:0	24007 <<< no label info for PE6
---------------	---------------------------------

Auto mapping, not need any extra configuration

```
RP/0/0/CPU0:P1#sh route 192.168.0.6/32 detail  
Wed Apr 10 08:29:09.416 UTC
```

Routing entry for 192.168.0.6/32

Known via "isis srte", distance 115, metric 1, labeled SR, type level-2
Installed Apr 10 07:49:13.491 for 00:39:56
Routing Descriptor Blocks

16.1.1.6, from 192.168.0.6, via GigabitEthernet0/0/0/1

Route metric is 1

Label: 0x3 (3) <<< use this

Tunnel ID: None

Binding Label: None

Extended communities count: 0

Path id:1 Path ref count:0

NHID:0x3(Ref:2)

Route version is 0x15 (21)

Local Label: 0x3e86 (16006)

IP Precedence: Not Set

QoS Group ID: Not Set

Flow-tag: Not Set

Fwd-class: Not Set

Route Priority: RIB_PRIORITY_NON_RECURSIVE_MEDIUM (7) SVD Type

RIB_SVD_TYPE_LOCAL

Download Priority 1, Download Version 102

No advertising protos.


P1

P1

P1

Enable SRMS

- show segment-routing mapping-server prefix-sid-map ipv4
- show isis segment-routing prefix-sid-map active-policy
- show isis seg label ta
- show cef vrf A 172.16.1.4/32



```
segment-routing  
mapping-server  
prefix-sid-map  
address-family ipv4 192.168.0.4/32 44  
router isis srte address-family ipv4 unicast segment-routing prefix-sid-map advertise-local
```

```

RP/0/0/CPU0:PE6#show cef vrf A 172.16.1.4/32 detail
Wed Apr 10 08:54:49.211 UTC
172.16.1.4/32, version 5, internal 0x5000001 0x0 (ptr 0xa11f7bf8) [1], 0x0
(0x0), 0x208 (0xa15d1268)
Updated Apr 10 05:05:08.065
Prefix Len 32, traffic index 0, precedence n/a, priority 3
  gateway array (0xa1116abc) reference count 1, flags 0x2038, source rib
(7), 0 backups
  [1 type 1 flags 0x40441 (0xa15ed318) ext 0x0 (0x0)]
LW-LDI[type=0, refc=0, ptr=0x0, sh-ldi=0x0]
  gateway array update type-time 5 Apr 10 08:51:46.943
LDI Update time Apr 10 08:51:46.943
via 192.168.0.4/32, 3 dependencies, recursive [flags 0x6000]
  path-idx 0 NHID 0x0 [0xa1645568 0x0]
  recursion-via-/32
next hop VRF - 'default', table - 0xe0000000
  next hop 192.168.0.4/32 via 16044/0/21
    next hop 16.1.1.1/32 Gi0/0/0/0      labels imposed {16044 24002}

```

Load distribution: 0 (refcount 1)

Hash	OK	Interface	Address
0	Y	Unknown	16044/0

PE6

```

RP/0/0/CPU0:PE6#sh mpls ldp bindings 192.168.0.4/32
Wed Apr 10 07:57:02.449 UTC
RP/0/0/CPU0:PE6#

```

PE6

RP/0/0/CPU0:PE6#sh cef 192.168.0.4/32 fla detail
 PE6

```

Wed Apr 10 08:55:25.169 UTC
192.168.0.4/32, version 171, labeled SR, internal 0x1000001 0x83 (ptr
0xa11f7ae0) [1], 0x0 (0xa11daa80), 0xa28 (0xa15d1160)
  leaf flags: owner locked, inserted

  leaf flags2: LDP/SR merge requested,RIB pref over LSD,sr-pfx,
  leaf ext flags:
PriChange,EXTERNAL_REACH_LC,L2TPV3_SPAN_DIAG_IFH_ENABLE,illegal-0x00000800,
  Updated Apr 10 08:51:46.943
  local adjacency 16.1.1.1
  Prefix Len 32, traffic index 0, precedence n/a, priority 1
    gateway array (0xa1117dc) reference count 3, flags 0x68, source rib (7),
  0 backups
    [2 type 5 flags 0x8401 (0xa15eda58) ext 0x0 (0x0)]
  LW-LDI[type=5, refc=3, ptr=0xa11daa80, sh-ldi=0xa15eda58]
    gateway array update type-time 1 Apr 10 08:51:46.943
  LDI Update time Apr 10 08:51:46.943
  LW-LDI-TS Apr 10 08:51:46.943
  pl flags: mpls ext, imp, lsw

  ldi flags: owner locked, added to pl, depth changed

  via 16.1.1.1/32, GigabitEthernet0/0/0/0, 7 dependencies, weight 0, class
  0 [flags 0x0]
    path-idx 0 NHID 0x0 [0xa189b1c0 0x0]
    next hop 16.1.1.1/32
    local adjacency
      local label 16044      labels imposed {16044}

```

Load distribution: 0 (refcount 2)

Hash	OK	Interface	Address
0	Y	GigabitEthernet0/0/0/0	16.1.1.1

```
RP/0/0/CPU0:PE6#sh route 192.168.0.4/32 detail  
Wed Apr 10 08:57:40.170 UTC
```

PE6

Routing entry for 192.168.0.4/32

Known via "isis srte", distance 115, metric 2, labeled SR(SRMS), type level-2

Installed Apr 10 08:51:46.924 for 00:05:53

Routing Descriptor Blocks

16.1.1.1, from 192.168.0.4, via GigabitEthernet0/0/0/0

Route metric is 2

Label: 0x3eac (16044)

Tunnel ID: None

Binding Label: None

Extended communities count: 0

Path id:1 Path ref count:0

NHID:0x1(Ref:6)

Route version is 0x35 (53)

Local Label: 0x3eac (16044)

IP Precedence: Not Set

QoS Group ID: Not Set

Flow-tag: Not Set

Fwd-class: Not Set

Route Priority: RIB_PRIORITY_NON_RECURSIVE_MEDIUM (7) SVD Type

RIB_SVD_TYPE_LOCAL

Download Priority 1, Download Version 171

No advertising protos.

```
RP/0/0/CPU0:PE6#traceroute vrf A 172.16.1.4  
Wed Apr 10 08:58:54.515 UTC
```

PE6

Type escape sequence to abort.

Tracing the route to 172.16.1.4

```
1 16.1.1.1 [MPLS: Labels 16044/24002 Exp 0] 19 msec 9 msec 9 msec  
2 14.1.1.4 9 msec * 9 msec
```

```
RP/0/0/CPU0:P1#sh mpls for prefix 192.168.0.4/32  
Wed Apr 10 08:59:24.772 UTC
```

Local Label	Outgoing Label	Prefix or ID	Outgoing Interface	Next Hop	Bytes Switched
24012	Pop	192.168.0.4/32	Gi0/0/0/2	14.1.1.4	2546

Two data plane.


```
RP/0/0/CPU0:P1#  
RP/0/0/CPU0:P1#
```

```
RP/0/0/CPU0:P1#sh mpls for labels 16044
```

```
Wed Apr 10 08:59:44.281 UTC
```

Local Label	Outgoing Label	Prefix or ID	Outgoing Interface	Next Hop	Bytes Switched
16044	Pop	SR Pfx (idx 44)	Gi0/0/0/2	14.1.1.4	1462


After config SRMS, forwarding recovery.



TI-LFA in Segment Routing


TI-LFA in Segment Routing

- Protecting destination Node5 on Node2 against failure of link 2-3
- **Classic LFA:** Node2 switches all traffic destined to Node5 towards the edge node PE4
 - Low BW (high metric) links and an edge node are used to protect the failure of a core link
 - A common planning rule is to avoid Edge nodes for transit traffic
 - Classic LFA does not respect this rule X
- **TI-LFA:** Node2 switches all traffic destined to Node5 via high BW core links: OK! ✓



Lab4 - TI-LFA in Segment Routing

- Remove LDP and recovery links and shut P1-P7
- Adjust ISIS metric to 100 between P7 and PE4, then check route at PE6 and PE4
- **Adjust ISIS LSP Generate timer to 4s at P1**
- Shut link between R1 and R4, then check traffics drop status
- **Adjust ISIS LSP Update timer to 10s at PE6, then check Micro-loop between PE6 and P1**
- Enable TI-LFA on P1, Check backup path in RIB detail, ISIS fast route and LFIB



Set up Testbed

- show cef 192.168.0.4/6/32
- trace 192.168.0.4/6

```
no mpls ldp
router isis srte address-family ipv4 unicast segment-routing mpls
router isis srte address-family ipv6 unicast segment-routing mpls
router isis srte interface Loopback0 address-family ipv4 unicast prefix-sid index 4
router isis srte interface Loopback0 address-family ipv6 unicast prefix-sid absolute 17004
```

PE4

```
no mpls ldp
interface GigabitEthernet0/0/0/3 shutdown
router isis srte interface GigabitEthernet0/0/0/1 address-family ipv4 unicast metric 100
```

P1

```
interface GigabitEthernet0/0/0/2 shutdown
no interface GigabitEthernet0/0/0/0 shutdown
no interface GigabitEthernet0/0/0/1 shutdown
router isis srte interface GigabitEthernet0/0/0/1 address-family ipv4 unicast metric 100
```

P7

```

RP/0/0/CPU0:PE6#show cef 192.168.0.4
Sat Apr 13 09:52:15.577 UTC
192.168.0.4/32, version 45, labeled SR, internal 0x1000001 0x81 (ptr
0xa1415c84) [1], 0x0 (0xa13f8698), 0xa28 (0xa1757370)
Updated Apr 13 09:42:26.717
local adjacency 16.1.1.1
Prefix Len 32, traffic index 0, precedence n/a, priority 1
via 16.1.1.1/32, GigabitEthernet0/0/0/0, 7 dependencies, weight 0, class 0
[flags 0x0]
path-idx 0 NHID 0x0 [0xa189b1c0 0x0]
next hop 16.1.1.1/32
local adjacency
local label 16004      labels imposed {16004}
RP/0/0/CPU0:PE6#

```

```

RP/0/0/CPU0:PE4#sh cef 192.168.0.6/32
Sat Apr 13 09:53:55.320 UTC
192.168.0.6/32, version 47, labeled SR, internal 0x1000001 0x81 (ptr
0xa13fbbeb4) [1], 0x0 (0xa13de698), 0xa28 (0xa1757370)
Updated Apr 13 09:33:19.354
local adjacency 47.1.1.7
Prefix Len 32, traffic index 0, precedence n/a, priority 1
via 47.1.1.7/32, GigabitEthernet0/0/0/0, 7 dependencies, weight 0, class 0
[flags 0x0]
path-idx 0 NHID 0x0 [0xa189b3d0 0x0]
next hop 47.1.1.7/32
local adjacency
local label 16006      labels imposed {16006}
RP/0/0/CPU0:PE4#


```

```

RP/0/0/CPU0:PE6#traceroute 192.168.0.4
Sat Apr 13 09:52:59.064 UTC
Type escape sequence to abort.
Tracing the route to 192.168.0.4

1 16.1.1.1 [MPLS: Label 16004 Exp 0] 79 msec 9 msec 0 msec
2 14.1.1.4 9 msec * 19 msec

```



```

RP/0/0/CPU0:PE4#traceroute 192.168.0.6
Sat Apr 13 09:54:01.820 UTC

```

Type escape sequence to abort.
Tracing the route to 192.168.0.6

```



1 47.1.1.7 [MPLS: Label 16006 Exp 0] 29 msec 0 msec 0 msec
2 67.1.1.6 9 msec * 0 msec

```

Adjust ISIS LSP Generate timer at P1, shutdown test



- show cef 192.168.0.4/32
- show mpls for labels 16004 det
- ping vrf A 172.16.1.4 re 5000 time 1

```
router isis srte lsp-gen-interval maximum-wait 20000 initial-wait 4000
```




Adjust ISIS LSP Update timer at P1, check Micro-Loop

- show cef 192.168.0.4/32
- show mpls for labels 16004 det
- ping vrf A 172.16.1.4 re 5000 time 1



Enable Ti-FLA

- show isis fast-reroute summary
- show isis ipv4 fast-reroute 192.168.0.4/32 detail
- show cef 192.168.0.4/32
- show isis ad det
- show mpls for labels 16004



```
RP/0/0/CPU0:P1#sh isis fast-reroute summary
Sun Apr 14 05:29:28.559 UTC
IS-IS srte IPv4 Unicast FRR summary
```

	Critical Priority	High Priority	Medium Priority	Low Priority	Total
Prefixes reachable in L2					
All paths protected	0	0	3	2	5
Some paths protected	0	0	0	0	0
Unprotected	0	0	1	0	1
Protection coverage	0.00%	0.00%	75.00%	100.00%	83.33%

```
router isis srte interface GigabitEthernet0/0/0/1 address-family ipv4 unicast fast-reroute per-prefix
router isis srte interface GigabitEthernet0/0/0/1 address-family ipv4 unicast fast-reroute per-prefix ti-lfa P1
router isis srte interface GigabitEthernet0/0/0/2 address-family ipv4 unicast fast-reroute per-prefix
router isis srte interface GigabitEthernet0/0/0/2 address-family ipv4 unicast fast-reroute per-prefix ti-lfa
```

```
no router isis srte add ipv4 un spf-interval maximum-wait 20000 initial-wait 10000
```

P1

PE6

```

RP/0/0/CPU0:P1#sh isis fast-reroute
Sun Apr 14 04:52:40.780 UTC
.....
L2 47.1.1.0/24 [2/115]
    via 14.1.1.4, GigabitEthernet0/0/0/2, PE4, SRGB Base: 16000, Weight: 0
        Backup path: TI-LFA (link), via 16.1.1.6, GigabitEthernet0/0/0/1
PE6, SRGB Base: 16000, Weight: 0, Metric: 201
    P node: P7.00 [192.168.0.7], Label: 16007
    Prefix label: ImpNull
    Backup-src: P7.00
L2 67.1.1.0/24 [3/115]
    via 14.1.1.4, GigabitEthernet0/0/0/2, PE4, SRGB Base: 16000, Weight: 0
        Backup path: LFA, via 16.1.1.6, GigabitEthernet0/0/0/1, PE6, SRGB
Base: 16000, Weight: 0, Metric: 101
L2 192.168.0.4/32 [1/115]
    via 14.1.1.4, GigabitEthernet0/0/0/2, PE4, SRGB Base: 16000, Weight: 0
        Backup path: TI-LFA (link), via 16.1.1.6, GigabitEthernet0/0/0/1
PE6, SRGB Base: 16000, Weight: 0, Metric: 201
    P node: P7.00 [192.168.0.7], Label: 16007
    Q node: PE4.00 [192.168.0.4], Label: 24001
    Prefix label: ImpNull
    Backup-src: PE4.00
L2 192.168.0.6/32 [3/115]
    via 14.1.1.4, GigabitEthernet0/0/0/2, PE4, SRGB Base: 16000, Weight: 0
        Backup path: LFA, via 16.1.1.6, GigabitEthernet0/0/0/1, PE6, SRGB
Base: 16000, Weight: 0, Metric: 100
L2 192.168.0.7/32 [2/115]
    via 14.1.1.4, GigabitEthernet0/0/0/2, PE4, SRGB Base: 16000, Weight: 0
        Backup path: LFA, via 16.1.1.6, GigabitEthernet0/0/0/1, PE6, SRGB
Base: 16000, Weight: 0, Metric: 101
L2 192.168.0.9/32 [5000/115]
    via 19.1.1.9, GigabitEthernet0/0/0/0, SR-PCE, Weight: 0
    No FRR backup

```

```

RP/0/0/CPU0:P7#sh isis ad det
Sun Apr 14 11:52:24.304 UTC

IS-IS srte Level-2 adjacencies:
System Id      Interface      SNPA      State Hold Changed   NSF IPv4 IPv6
                                         BFD BFD
PE4            Gi0/0/0/1      *PtoP*      Up   27   01:15:53 Yes None None
    Area Address:          49
    Neighbor IPv4 Address: 47.1.1.4*
    Adjacency SID:        24000
    Non-FRR Adjacency SID: 24001
    Neighbor IPv6 Address: fe80::f816:3eff:fe81:1505*
    Adjacency SID:        24006
    Non-FRR Adjacency SID: 24007
    Topology:               IPv4 Unicast
    Topology:               IPv6 Unicast
PE6            Gi0/0/0/0      *PtoP*      Up   27   01:15:50 Yes None None
    Area Address:          49
    Neighbor IPv4 Address: 67.1.1.6*
    Adjacency SID:        24002
    Non-FRR Adjacency SID: 24003
    Neighbor IPv6 Address: fe80::f816:3eff:fe92:ff50*
    Adjacency SID:        24008
    Non-FRR Adjacency SID: 24009
    Topology:               IPv4 Unicast
    Topology:               IPv6 Unicast

Total adjacency count: 2

```

```

RP/0/0/CPU0:P1#sh mpls forwarding labels 16004 detail
Sun Apr 14 05:38:35.521 UTC
Local Outgoing Prefix Outgoing Next Hop Bytes
Label Label or ID Interface Switch
-----
16004 Pop SR Pfx (idx 4) Gi0/0/0/2 14.1.1.4 86912
    Updated: Apr 14 00:50:10.426
    Path Flags: 0x400 [ BKUP-IDX:0 (0xa19af210) ]
    Version: 452, Priority: 1
    Label Stack (Top -> Bottom): { Imp-Null }
    NHID: 0x0, Encap-ID: N/A, Path idx: 1, Backup path idx: 0, Weight: 0
    MAC/Encaps: 14/14, MTU: 1500
    Outgoing Interface: GigabitEthernet0/0/0/2 (ifhandle 0x00000080)
    Packets Switched: 1688

    16007 SR Pfx (idx 4) Gi0/0/0/1 16.1.1.6 0
(!)
    Updated: Apr 14 00:50:10.426
    Path Flags: 0xb00 [ IDX:0 BKUP, NoFwd ]
    Version: 452, Priority: 1
    Label Stack (Top -> Bottom): { 16007 24001 }
    NHID: 0x0, Encap-ID: N/A, Path idx: 0, Backup path idx: 0, Weight: 0
    MAC/Encaps: 14/22, MTU: 1500
    Outgoing Interface: GigabitEthernet0/0/0/1 (ifhandle 0x00000060)
    Packets Switched: 0
    (!): FRR pure backup

Traffic-Matrix Packets/Bytes Switched: 0/0

```

```

RP/0/0/CPU0:P1#sh isis adj gi0/0/0/2 det
Sun Apr 14 12:55:05.036 UTC

IS-IS srte Level-2 adjacencies:
System Id Interface SNPA State Hold Changed NSF IPv4 IPv6
PE4 Gi0/0/0/2 *PtoP* Up 28 02:18:33 Yes None None
BFD BFD

Area Address: 49
Neighbor IPv4 Address: 14.1.1.4*
Adjacency SID: 24000 (protected)
Backup label stack: [16007, 24001]
Backup stack size: 2
Backup interface: Gi0/0/0/1
Backup nexthop: 16.1.1.6
Backup node address: 192.168.0.4
Non-FRR Adjacency SID: 24001
Neighbor IPv6 Address: fe80::f816:3eff:fedb:730f*
Adjacency SID: 24008
Non-FRR Adjacency SID: 24009
Topology: IPv4 Unicast
Topology: IPv6 Unicast

Total adjacency count: 1

RP/0/0/CPU0:P1#show cef 192.168.0.4/32
Sun Apr 14 05:01:01.286 UTC
192.168.0.4/32, version 452, labeled SR, internal 0x1000001 0x81 (ptr 0xa1415fcc)
[1], 0x0 (0xa13f8850), 0xa28 (0xa195d0b8)
Updated Apr 14 00:50:10.427
local adjacency 14.1.1.4
Prefix Len 32, traffic index 0, precedence n/a, priority 1
via 16.1.1.6/32, GigabitEthernet0/0/0/1, 11 dependencies, weight 0, class 0,
backup (TI-LFA) [flags 0xb00]
path-idx 0 NHID 0x0 [0xa189b320 0x0]
next hop 16.1.1.6/32, Repair Node(s): 192.168.0.7, 192.168.0.4
local adjacency
    local label 16004 labels imposed {16007 24001}
via 14.1.1.4/32, GigabitEthernet0/0/0/2, 11 dependencies, weight 0, class 0,
protected [flags 0x400]
path-idx 1 bkpup-idx 0 NHID 0x0 [0xa19af210 0xa19af0f0]
next hop 14.1.1.4/32
local adjacency
    local label 16004 labels imposed {ImplNull}

```

More...we support now!

- Except Prefix-SID, Ti-LFA protect Adj-SID too, the scenario will also benefit when you include Adj-SID in your SRTE dynamic SID-list. [-> page 431](#)
- Ti-LFA protect IP traffics
- Ti-LFA protect LDP traffics
- Ti-LFA protect traffics in SR/LDP scenario
- Ti-LFA protect ECMP traffics
- Ti-LFA protect SR Policy
- Ti-LFA support Link, Node, SRLG, SRLG + Node, Remote SRLG



SRTE Solution in MPLS SR

简化、大规模和自动化的流量工程

SR 路径计算单元(SR-PCE)



基于意图的 SRTE SR 按需下一跳(ODN)




基于意图的延迟优化网络切片




多平面网络

由SR IGP Flex Algo提供支持



Lab5 - SRTE Solution in MPLS SR

- Recovery IGP metric for P1 and P7
- Deploy PCE/PCC in XTC and PE6, and enable BGP-LS between PE4 and XTC
- Set color 10 for ODN at PE4, ODN template at PE6, manual set latency and update to XTC via BGP-LS
- Set Explicit SID-List for PE6 → P1 → P7 → PE4
- Check SR Policy status and check path
- Change ODN from PCE to Local SR Policy, check Policy again.
- Enable Ti-LFA at PE6



Deploy PCE/PCC/BGP-LS

- show pce ipv4 peer
- show segment-routing traffic-eng pcc ipv4 peer
- show bgp link-state link-state summary
- show pce ipv4 topology

```
Router isis srte distribute link-state
router bgp 1
  address-family link-state link-state
  !
  neighbor 192.168.0.9
    remote-as 1
    description iBGP peer XTC      PE4
    update-source Loopback0
    address-family link-state link-state
```

```
seg tr pcc
source-address ipv4 192.168.0.6
pce address ipv4 192.168.0.9
password clear cisco
report-all
```

PE6

```
pce
  address ipv4 192.168.0.9
  rest
    user cisco
  !
  authentication basic
  !
  logging
    no-path
    fallback
  !
  password clear cisco
```

SR PCE

```
RP/0/0/CPU0:SR-PCE#sh pce ipv4 peer  
Sun Apr 14 13:18:16.052 UTC
```

PCE's peer database:

Peer address: 192.168.0.6

State: Up

Capabilities: Stateful, Segment-Routing, Update, Instantiation

```
RP/0/0/CPU0:PE6#show segment-routing traffic-eng pcc ipv4 peer  
Sun Apr 14 13:18:46.919 UTC
```

PCC's peer database:

Peer address: 192.168.0.9, Precedence: 255, (best PCE)

State up

Capabilities: Stateful, Update, Segment-Routing, Instantiation

```
RP/0/0/CPU0:SR-PCE#show pce ipv4 topology  
Sun Apr 14 13:19:35.046 UTC
```

PCE's topology database - detail:

Node 1

Host name: SR-PCE

ISIS system ID: 1921.6800.0009 level-2 ASN: 1

Node 2

Host name: P1

ISIS system ID: 1921.6800.0001 level-2 ASN: 1

Prefix SID:

ISIS system ID: 1921.6800.0001 level-2 ASN: 1 domain ID: 0

Prefix 192.168.0.1, label 16001 (regular), flags: N

SRGB INFO:

ISIS system ID: 1921.6800.0001 level-2 ASN: 1

SRGB Start: 16000 Size: 8000

Link[0]: local address 14.1.1.1, remote address 14.1.1.4

```
RP/0/0/CPU0:SR-PCE#show bgp link-state link-state summary  
Sun Apr 14 13:19:12.218 UTC
```

.....
BGP is operating in STANDALONE mode.

Process Speaker	RcvTblVer	bRIB/RIB	LabelVer	ImportVer	SendTblVer	StandbyVer			
192.168.0.4	52	52	52	52	52	0			
Neighbor	Spk	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	St/PfxRcd
192.168.0.4	0	1	37	9	52	0	0	00:06:53	51

```
RP/0/0/CPU0:PE4#show segment-routing traffic-eng ipv4 topology  
Sun Apr 14 15:46:26.522 UTC
```

SR-TE topology database

Node 1

TE router ID: 192.168.0.1

Host name: P1

ISIS system ID: 1921.6800.0001 level-2

ISIS system ID: 1921.6800.0001 level-2

Prefix SID:

Prefix 192.168.0.1, label 16001 (regular)

Prefix 192.168.0.1, label 16001 (regular)

Link[0]: local address 14.1.1.1, remote address 14.1.1.4

.....

Deploy PE4/CE

- show bgp vrf A sum
- show bgp vpnv4 un
- show bgp vrf A 192.168.0.10/32
- show bgp vrf A 192.168.0.11/32

```
router bgp 1
  vrf A
    neighbor 110.1.1.10
      remote-as 10
      address-family ipv4 unicast
        route-policy bgp_in in
        route-policy bgp_out out
    !
    neighbor 111.1.1.11
      remote-as 11
      address-family ipv4 unicast
        route-policy bgp_in in
        route-policy bgp_out out
```

```
route-policy bgp_in
  pass
end-policy
!
route-policy bgp_out
  pass
end-policy
!
interface GigabitEthernet0/0/0/2
  no ipv4 address
  vrf A
  ipv4 address 111.1.1.1 255.255.255.0
!
interface GigabitEthernet0/0/0/3
  no ipv4 address
  vrf A
  ipv4 address 110.1.1.1 255.255.255.0
!
```

```

RP/0/0/CPU0:PE4#sh bgp vpng4 un
Sun Apr 14 13:40:50.988 UTC
BGP router identifier 192.168.0.4, local AS number 1
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0 RD version: 0
BGP main routing table version 13
BGP NSR Initial initsync version 5 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
      Network          Next Hop          Metric LocPrf Weight Path
Route Distinguisher: 1:1
*>i172.16.1.6/32    192.168.0.6        0     100      0 ?
Route Distinguisher: 1:2 (default for vrf A)
*> 11.11.11.11/32   110.1.1.10        0         0 10 i
*       111.1.1.11        0         0 11 i
*> 110.1.1.0/24     0.0.0.0          0     32768 ?
*> 111.1.1.0/24     0.0.0.0          0     32768 ?
*> 172.16.1.4/32    0.0.0.0          0     32768 ?
*>i172.16.1.6/32    192.168.0.6        0     100      0 ?
*> 192.168.0.10/32  110.1.1.10        0         0 10 i
*> 192.168.0.11/32  111.1.1.11        0         0 11 i

Processed 8 prefixes, 9 paths

```

```

RP/0/0/CPU0:PE4#sh bgp vrf A sum
Sun Apr 14 13:36:37.596 UTC
BGP VRF A, state: Active
BGP Route Distinguisher: 1:2
VRF ID: 0x60000001
BGP router identifier 192.168.0.4, local AS number 1
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000010 RD version: 13
BGP main routing table version 13
BGP NSR Initial initsync version 5 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0

```

BGP is operating in STANDALONE mode.

Process	RcvTblVer	bRIB/RIB	LabelVer	ImportVer	SendTblVer	StandbyVer			
Speaker	13	13	13	13	13	0			
Neighbor	Spk	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	St/PfxRcd
110.1.1.10	0	10	8	10	13	0	0	00:04:16	2
111.1.1.11	0	11	7	11	13	0	0	00:04:03	2

```
RP/0/0/CPU0:PE4#sh bgp vrf A 192.168.0.10/32
Sun Apr 14 13:41:33.816 UTC
BGP routing table entry for 192.168.0.10/32, Route Distinguisher: 1:2
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          11        11
    Local Label: 24010
Last Modified: Apr 14 13:32:25.859 for 00:09:07
Paths: (1 available, best #1)
  Advertised to CE update-groups (with more than one peer):
    0.2
  Path #1: Received by speaker 0
  Advertised to CE update-groups (with more than one peer):
    0.2
  10
    110.1.1.10 from 110.1.1.10 (192.168.0.10)
      Origin IGP, metric 0, localpref 100, valid, external, best, group-best,
import-candidate
      Received Path ID 0, Local Path ID 1, version 11
      Extended community: RT:1:1
      Origin-AS validity: (disabled)
RP/0/0/CPU0:PE4#
```

```
RP/0/0/CPU0:PE4#sh bgp vrf A 192.168.0.11/32
Sun Apr 14 13:41:41.605 UTC
BGP routing table entry for 192.168.0.11/32, Route Distinguisher: 1:2
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          13        13
    Local Label: 24011
Last Modified: Apr 14 13:32:39.859 for 00:09:01
Paths: (1 available, best #1)
  Advertised to CE update-groups (with more than one peer):
    0.2
  Path #1: Received by speaker 0
  Advertised to CE update-groups (with more than one peer):
    0.2
  11
    111.1.1.11 from 111.1.1.11 (192.168.0.11)
      Origin IGP, metric 0, localpref 100, valid, external, best, group-best,
import-candidate
      Received Path ID 0, Local Path ID 1, version 13
      Extended community: RT:1:1
      Origin-AS validity: (disabled)
```

Deploy SR Policy by ODN with latency

- show bgp vpng4 un
- show bgp vrf A 192.168.0.10/11/32
- show segment-routing traffic-eng policy detail
- show performance-measurement interfaces

PM support on ASR9k

```
performance-measurement interface gi0/0/0/0 delay-measurement advertise-delay 500  
performance-measurement interface gi0/0/0/1 delay-measurement advertise-delay 50 PE6
```

```
performance-measurement interface gi0/0/0/1 delay-measurement advertise-delay 500  
performance-measurement interface gi0/0/0/3 delay-measurement advertise-delay 50 P1  
performance-measurement interface gi0/0/0/2 delay-measurement advertise-delay 50
```

```
performance-measurement interface gi0/0/0/1 delay-measurement advertise-delay 500  
performance-measurement interface gi0/0/0/0 delay-measurement advertise-delay 50 P7  
performance-measurement interface gi0/0/0/2 delay-measurement advertise-delay 50
```

```
performance-measurement interface gi0/0/0/0 delay-measurement advertise-delay 500  
performance-measurement interface gi0/0/0/1 delay-measurement advertise-delay 50 PE4
```

```
segment-routing  
traffic-eng  
on-demand color 10  
dynamic  
pcep  
!  
metric  
type latency PE6
```

```
extcommunity-set opaque odn  
10  
end-set  
route-policy bgp_color_in  
if destination in (192.168.0.11/32)  
then  
set extcommunity color odn  
endif  
pass  
end-policy
```

```
router bgp 1  
vrf A  
neighbor 111.1.1.11  
address-family ipv4 unicast  
route-policy bgp_color_in in  
! PE4
```

```
RP/0/0/CPU0:PE6#show segment-routing traffic-eng policy detail
Sun Apr 14 15:08:15.839 UTC
```

SR-TE policy database

Color: 10, End-point: 192.168.0.4

Name: srte_c_10_ep_192.168.0.4

Status:

Admin: up Operational: up for 00:05:27 (since Apr 14
15:02:48.872)

Candidate-paths:

Preference: 200 (BGP ODN) (shutdown)

Requested BSID: dynamic

Dynamic (invalid)

Preference: 100 (BGP ODN) (current)

Requested BSID: dynamic

Dynamic (pce 192.168.0.9) (valid)

Metric Type: LATENCY, Path Accumulated Metric: 150

16007 [Prefix-SID, 192.168.0.7]

16001 [Prefix-SID, 192.168.0.1]

16004 [Prefix-SID, 192.168.0.4]

LSPs:

LSP[0]:

LSP-ID: 1 policy ID: 3 (current)

Local label: 24010

State: Programmed

Binding SID: 24012

Attributes:

Binding SID: 24012

Forward Class: 0

Steering BGP disabled: no

IPv6 caps enable: yes

```
RP/0/0/CPU0:PE6#sh bgp vpnv4 un rd 1:1
```

Sun Apr 14 15:06:41.206 UTC

.....

Status codes: s suppressed, d damped, h history, * valid, > best
i - internal, r RIB-failure, S stale, N Nexthop-discard

Origin codes: i - IGP, e - EGP, ? - incomplete

Network	Next Hop	Metric	LocPrf	Weight	Path
---------	----------	--------	--------	--------	------

Route Distinguisher: 1:1 (default for vrf A)

*>i11.11.11.11/32	192.168.0.4	0	100	0	10 i
*>i110.1.1.0/24	192.168.0.4	0	100	0	?
*>i111.1.1.0/24	192.168.0.4	0	100	0	?
*>i172.16.1.4/32	192.168.0.4	0	100	0	?
*> 172.16.1.6/32	0.0.0.0	0		32768	?
*>i192.168.0.10/32	192.168.0.4	0	100	0	10 i
*>i192.168.0.11/32	192.168.0.4 C:10				
		0	100	0	11 i

Processed 7 prefixes, 7 paths

```
RP/0/0/CPU0:PE6# sh bgp vrf A 192.168.0.11/32
```

Sun Apr 14 15:07:28.893 UTC

BGP routing table entry for 192.168.0.11/32, Route Distinguisher: 1:1

Versions:

Process	bRIB/RIB	SendTblVer
---------	----------	------------

Speaker	39	39
---------	----	----

Last Modified: Apr 14 15:02:48.873 for 00:04:40

Paths: (1 available, best #1)

Not advertised to any peer

Path #1: Received by speaker 0

Not advertised to any peer

11

192.168.0.4 C:10 (bsid:24012) (metric 2) from 192.168.0.4 (192.168.0.4)

Received Label 24011

Origin IGP, metric 0, localpref 100, valid, internal, best, group-best, import-candidate, imported

Received Path ID 0, Local Path ID 1, version 20

Extended community: Color:10 RT:1:1

SR policy color 10, up, registered, bsid 24012, if-handle 0x000000d0

Source AFI: VPNv4 Unicast, Source VRF: default, Source Route Distinguisher: 1:2

```
RP/0/0/CPU0:PE6#sh performance-measurement interfaces gigabitEth 0/0/0/0
```

```
Sun Apr 14 15:14:54.682 UTC
```

```
0/0/CPU0
```

```
Interface Name: GigabitEthernet0/0/0/0 (ifh: 0x40)
```

```
Delay-Measurement : Enabled  
Local IPV4 Address : 16.1.1.6  
Local IPV6 Address : 2001::16:1:1:6  
Local MAC Address : fa16.3ef7.2b4e  
Primary VLAN Tag : None  
Secondary VLAN Tag : None  
State : Up
```

```
Delay Measurement session:
```

```
Session ID : 33554433
```

```
Last advertisement:
```

```
Advertised at: Apr 14 2019 14:21:21.583 (3213.130 seconds ago)
```

```
Advertised reason: Advertise delay config
```

```
Advertised delays (uSec): avg: 500, min: 500, max: 500, variance: 0
```

```
Next advertisement:
```

```
Check scheduled in 4 more probes (roughly every 120 seconds)
```

```
No probes completed <<< support on asr9k, not xrv
```

```
RP/0/0/CPU0:SR-PCE#show bgp lin lin
```

```
[E][L2][I0x0][N[c1][b0.0.0.0][s1921.6800.0006.00]][R[c1][b0.0.0.0][s1921.6800.0001.00]][L[i16.1.1.6][n16.1.1.1]]/696 detail  
Fri Jul 12 14:56:32.368 UTC
```

```
.....  
192.168.0.4 (metric 5001) from 192.168.0.4 (192.168.0.4)  
Origin IGP, localpref 100, valid, internal, best, group-best  
Received Path ID 0, Local Path ID 1, version 16  
Link-state: MSD: Type 1 Value 10, Local TE Router-ID:  
192.168.0.6 Remote TE Router-ID: 192.168.0.1, metric: 1  
ADJ-SID: 24001(30), Link Delay: 500 us Flags: 0x00  
Min Delay: 500 us Max Delay: 500 us Flags: 0x00, Delay:  
Variation: 0 us
```

```
RP/0/0/CPU0:PE6#sh isis database PE6 internal verbose
```

```
Fri Jul 12 14:40:40.373 UTC
```

```
IS-IS srte (Level-2) Link State Database
```

LSPID	LSP Seq Num	LSP Checksum	LSP Holdtime/Rcvd	ATT/P/OL
PE6.00-00	* 0x00000008	0xc9ac	1104 /*	0/0/0
532				

```
.....
```

```
TLV code:22 length:112
```

```
Metric: 1 IS-Extended P1.00
```

```
SubTLV code:6 length:4
```

```
Interface IP Address: 16.1.1.6
```

```
SubTLV code:8 length:4
```

```
Neighbor IP Address: 16.1.1.1
```

```
SubTLV code:33 length:4
```

```
Link Average Delay: 500 us
```

```
SubTLV code:34 length:8
```

```
Link Min/Max Delay: 500/500 us
```

```
SubTLV code:35 length:4
```

```
Link Delay Variation: 0 us
```

```
SubTLV code:15 length:2
```

```
Link Maximum SID Depth:
```

```
Subtype: 1, Value: 10
```

```
SubTLV code:31 length:5
```

```
ADJ-SID: F:0 B:0 V:1 L:1 S:0 P:0 weight:0 Adjacency-sid:24001
```

```
RP/0/0/CPU0:PE6#traceroute sr-mpls nil-fec policy name srte_c_10_ep_192.168.0.4  
Sun Apr 14 15:17:54.760 UTC
```

```
Tracing MPLS Label Switched Path with Nil FEC for SR-TE Policy
```

```
srtc_c_10_ep_192.168.0.4, timeout is 2 seconds
```

```
.....
```

```
Type escape sequence to abort.
```

Need to enable "mpls oam" first, then do sr trace

```
0 67.1.1.6 MRU 1500 [Labels: 16001/16004/explicit-null Exp: 0/0/0]
```

```
L 1 67.1.1.7 MRU 1500 [Labels: implicit-null/16004/explicit-null Exp: 0/0/0] 10 ms
```

```
L 2 17.1.1.1 MRU 1500 [Labels: implicit-null/explicit-null Exp: 0/0/0] 20 ms
```

```
! 3 14.1.1.4 40 ms
```

Troubleshooting for “No path”

- show segment-routing traffic-eng policy private
- show pce lsp private last 1
- debug pce

```
RP/0/0/CPU0:SR-PCE#debug pce ?
  config      Configuration related events
  cspf        CSPF related events
  cspf-internal  CSPF internal related events
  db          Database related events
  error       Errors
  filter      Filter debug messages
  netconf     Netconf related events
  path        Path computation events
  pcep        PCEP events
  policy      Policy related events
  rest        REST related events
  rest-api    rest-api related events
  segment-routing Segment-routing related events
  topology    Topology events
  verify      LSP verification events
```

Deploy SR Policy by Explicit path

```
segment-routing
  traffic-eng
    segment-list static
      index 10 mpls label 16001
      index 20 add ipv4 192.168.0.7
      index 30 mpls label 16004
    !
    policy static
      color 20 end-point ipv4 192.168.0.4
      candidate-paths
        preference 100
        explicit segment-list static
```

PE6

```
RP/0/0/CPU0:PE6#sh seg traffic-eng policy name srte_c_20_ep_192.168.0.4
Sun Apr 14 15:53:47.642 UTC

SR-TE policy database
-----
Color: 20, End-point: 192.168.0.4
Name: srte_c_20_ep_192.168.0.4
Status:
Admin: up Operational: down for 00:00:07 (since Apr 14 15:53:40.303)
Candidate-paths:
  Preference: 100 (configuration)
  Name: static
  Requested BSID: dynamic
Explicit: segment-list static (invalid)
Weight: 1, Metric Type: TE
Attributes:
  Forward Class: 0
  Steering BGP disabled: no
  IPv6 caps enable: no
```

Why?

```
RP/0/0/CPU0:PE6#traceroute sr-mpls nil-fec policy name srte_c_20_ep
Sun Apr 14 16:03:18.113 UTC
.....
Type escape sequence to abort.
```

```
0 16.1.1.6 MRU 1500 [Labels: 16007/16004/explicit-null Exp: 0/0/0]
L 1 16.1.1.1 MRU 1500 [Labels: implicit-null/16004/explicit-null Exp: 0/0/0] 0 ms
L 2 17.1.1.7 MRU 1500 [Labels: implicit-null/explicit-null Exp: 0/0/0] 0 ms
! 3 47.1.1.4 1 ms
```

```
RP/0/0/CPU0:PE6#show segment-routing traffic-eng ipv4 topology
Sun Apr 14 15:46:13.513 UTC
RP/0/0/CPU0:PE6#
```

distribute link-state

Check Local SR Policy with ODN and enable Ti-LFA

```
segment-routing
  traffic-eng
    on-demand color 10
    dynamic
      no pcep

  router isis srte interface GigabitEthernet0/0/0/1 address-family
  ipv4 unicast fast-reroute per-prefix

  router isis srte interface GigabitEthernet0/0/0/1 address-family
  ipv4 unicast fast-reroute per-prefix ti-lfa

  router isis srte interface GigabitEthernet0/0/0/0 address-family
  ipv4 unicast fast-reroute per-prefix

  router isis srte interface GigabitEthernet0/0/0/0 address-family
  ipv4 unicast fast-reroute per-prefix ti-lfa
```

PE6

```
RP/0/0/CPU0:PE6#sh segment-routing traffic-eng forwarding policy
```

```
Sun Apr 14 16:20:32.132 UTC
```

Color	Endpoint	Segment List	Outgoing Label	Outgoing Interface	Next Hop	Bytes Switched	
10	192.168.0.4	dynamic	16007	Gi0/0/0/0	16.1.1.1	0	(!)
			16001	Gi0/0/0/1	67.1.1.7	0	
20	192.168.0.4	static	16007	Gi0/0/0/0	16.1.1.1	0	
			16001	Gi0/0/0/1	67.1.1.7	0	(!)

```
RP/0/0/CPU0:PE6#sh segment-routing traffic-eng policy name
srte_c_10_ep_192.168.0.4 d
Sun Apr 14 16:14:38.357 UTC

SR-TE policy database
-----
Color: 10, End-point: 192.168.0.4
  Name: srte_c_10_ep_192.168.0.4
  Status:
    Admin: up Operational: up for 01:11:49 (since Apr 14 15:02:48.872)
    Candidate-paths:
      Preference: 200 (BGP ODN) (current)
        Requested BSID: dynamic
        Dynamic (valid)
          Metric Type: LATENCY, Path Accumulated Metric: 150
            16007 [Prefix-SID, 192.168.0.7]
            16001 [Prefix-SID, 192.168.0.1]
            16004 [Prefix-SID, 192.168.0.4]
      Preference: 100 (BGP ODN)
        Requested BSID: dynamic
        Dynamic (pce 192.168.0.9) (valid)
          Metric Type: LATENCY, Path Accumulated Metric: 150
            16007 [Prefix-SID, 192.168.0.7]
            16001 [Prefix-SID, 192.168.0.1]
            16004 [Prefix-SID, 192.168.0.4]
LSPs:
  LSP[0]:
    LSP-ID: 2 policy ID: 3 (current)
    Local label: 24014
    State: Programmed
    Binding SID: 24012
Attributes:
  Binding SID: 24012
```

We can compute path by head-end, but not only controller!

More...we support now!


- BGP IPv4 over SRTE (+ODN/AS)
- BGP IPv6 over SRTE (+ODN/AS)
- BGP IPv6 (6PE) over SRTE (+ODN/AS)
- VPNv4 over SRTE (+ODN/AS)
- VPNv6 (6VPE) over SRTE (+ODN/AS)
- EVPN VPWS Single-homing over SRTE (+ODN/AS)
- VPWS/VPLS (LDP) over SRTE



SR IGP Flexible Algorithms


SR IGP Flexible Algorithms

- ISIS 链路度量 10
- 链路颜色
- 所有节点
- Algo 128 使流量在
- 链路颜色
- ISIS link metric 10
- Link colors shown **1G / 10G / 100G / 400G**
- All nodes support Algo 0 & 128
- Algo 128: minimize IGP while traversing links of 100G or more (**exclude brown and red**)
- Per-link colors flooded in IGP



Lab6 - SR IGP Flexible Algorithms

- Set all node Algo 128, then binding with color 10, compare currently latency path with original latency paths
- Due to the link issue between P1 and P7, so exclude the link from Algo 128, check status, then rollback
- As right diagram, Set Alog 0(all nodes), 128(PE6/P1/PE4),129(PE6/P7/PE4); Alog 0 default, Algo 128 delay, Algo 129 Low cost.
- Deploy SR Policy with ODN for Alog 0(color 50),128(color 100),129(color 110)
- Set color at PE6, 11.11.11.11 with color 50; 192.168.0.10 with color 100; 192.168.0.11 with color 110;
- Check SR Policy at PE6
- Capture packets at G0/0/0/0 of PE6 and check label info



```

router isis srte
flex-algo 128
metric-type delay
advertise-definition
!
interface Loopback0
address-family ipv4 unicast
prefix-sid algorithm 128 absolute 1680x

```

ALL

```

segment-routing
traffic-eng
on-demand color 10
dynamic
sid-algorithm 128
!
```

PE6

```

RP/0/0/CPU0:PE6#show segment-routing traffic-eng policy detail
Sun Apr 14 15:08:15.839 UTC

```

Original latency path, Page 74

```

SR-TE policy database
-----
Color: 10, End-point: 192.168.0.4
Name: srte_c_10_ep_192.168.0.4
Status:
Admin: up Operational: up for 00:05:27 (since Apr 14
15:02:48.872)
Candidate-paths:
Preference: 200 (BGP ODN) (shutdown)
Requested BSID: dynamic
Dynamic (invalid)
Preference: 100 (BGP ODN) (current)
Requested BSID: dynamic

```

```

Dynamic (pce_192.168.0.9) (valid)
Metric Type: LATENCY, Path Accumulated Metric: 150
16007 [Prefix-SID, 192.168.0.7]
16001 [Prefix-SID, 192.168.0.1]
16004 [Prefix-SID, 192.168.0.4]

```

```

RP/0/RP0/CPU0:PE6#sh segment-routing traffic-eng policy color 10
Sun Jul 14 01:50:00.738 UTC

```

SR-TE policy database

7.0.1 Engineer version

```

Color: 10, End-point: 192.168.0.4
Name: srte_c_10_ep_192.168.0.4
Status:
Admin: up Operational: up for 00:10:23 (since Jul 14 01:39:37.015)
Candidate-paths:
Preference: 200 (BGP ODN) (active)
Requested BSID: dynamic
Constraints:
Prefix-SID Algorithm: 128
Maximum SID Depth: 10
Dynamic (valid)
Metric Type: LATENCY, Path Accumulated Metric: 150
16804 [Prefix-SID: 192.168.0.4, Algorithm: 128]
Preference: 100 (BGP ODN)
Requested BSID: dynamic

```

```

RP/0/RP0/CPU0:PE6#sh mpls for labels 16804
Sun Jul 14 01:52:24.255 UTC

```

Local Label	Outgoing Label	Prefix or ID	Outgoing Interface	Next Hop	Bytes Switched
16804	16804	SR Pfx (idx 804)	G10/0/0/1	67.1.1.7	0
	16804	SR Pfx (idx 804)	G10/0/0/0	16.1.1.1	0

```

RP/0/RP0/CPU0:PE6#traceroute sr-mpls nil-fec policy name
srte_c_10_ep_192.168.$
Sun Jul 14 01:54:24.357 UTC
.....
```

Type escape sequence to abort.

```

0 67.1.1.6 MRU 1500 [Labels: 16804/explicit-null Exp: 0/0]
L 1 67.1.1.7 MRU 1500 [Labels: 16804/explicit-null Exp: 0/0] 64 ms
L 2 17.1.1.1 MRU 1500 [Labels: implicit-null/explicit-null Exp: 0/0] 54 ms
! 3 14.1.1.4 21 ms

```

Exclude the link between P1 and P7

rollback configuration last 1

```
router isis srte
  affinity-map issue bit-position 1
  flex-algo 128
    affinity exclude-any issue
  !
  interface GigabitEthernet0/0/0/2 | 3
    affinity flex-algo issue
  !
```

```
RP/0/0/CPU0:PE6#sh mpls for labels 16804
Sun Jul 14 07:36:40.282 UTC
```

Local Label	Outgoing Label	Prefix or ID	Outgoing Interface	Next Hop	Bytes Switched

16804	16804	SR Pfx (idx 804)	Gi0/0/0/0	16.1.1.1	0
	16804	SR Pfx (idx 804)	Gi0/0/0/1	67.1.1.7	0

```
RP/0/0/CPU0:PE6#traceroute sr-mpls nil-fec policy name srte_c_10_ep_192.168.0.4
Sun Jul 14 07:47:00.530 UTC
```

```
Tracing MPLS Label Switched Path with Nil FEC for SR-TE Policy
srte_c_10_ep_192.168.0.4, timeout is 2 seconds
```

.....

Type escape sequence to abort.

```
 0 67.1.1.6 MRU 1500 [Labels: 16804/explicit-null Exp: 0/0]
L 1 67.1.1.7 MRU 1500 [Labels: implicit-null/explicit-null Exp: 0/0] 10 ms
! 2 47.1.1.4 1 ms
```

TLV code:22 length:179

.....

Metric: 1 IS-Extended P1.00
SubTLV code:6 length:4
Interface IP Address: 17.1.1.7
SubTLV code:8 length:4
Neighbor IP Address: 17.1.1.1
SubTLV code:33 length:4
Link Average Delay: 50 us
SubTLV code:34 length:8
Link Min/Max Delay: 50/50 us
SubTLV code:35 length:4
Link Delay Variation: 0 us
SubTLV code:16 length:9
Application Specific Link Attributes:
L flag: 0, SA-Length 1, UDA-Length 0
Standard Applications: FLEX-ALGO
SubTLV code:14 length:4
Ext Admin Group:
0x00000002
SubTLV code:15 length:2
Link Maximum SID Depth:
Subtype: 1, Value: 10
SubTLV code:31 length:5
ADJ-SID: F:0 B:0 V:1 L:1 S:0 P:0 weight:0 Adjacency-sid:24001

TLV code:242 length:37

Router Cap: 192.168.0.7, D:0, S:0
SubTLV code:2 length:9

Segment Routing: I:1 V:1, SRGB Base: 16000 Range: 8000

SubTLV code:19 length:3

SR Algorithm:

Algorithm: 0

Algorithm: 1

Algorithm: 128

SubTLV code:23 length:2

Node Maximum SID Depth:

Subtype: 1, Value: 10

SubTLV code:26 length:10

Flex-Algo Definition:

Algorithm: 128 Metric-Type: 1 Alg-type: 0 Priority: 128

SubTLV code:1 length:4

Enable Two Planes by Flex-Algo

```
router isis srte      PE4/6
  flex-algo 129
    advertise-definition
  !
  interface Loopback0
    address-family ipv4 unicast
      prefix-sid algorithm 129
      absolute 1690x
    !

```

```
router isis srte      P7
  no flex-algo 128
  flex-algo 129
    advertise-definition
  !
  interface Loopback0
    address-family ipv4 unicast
      no prefix-sid algorithm 128
      absolute 1680x
      prefix-sid algorithm 129
      absolute 1690x
    !

```

```
RP/0/0/CPU0:PE6#sh mpls forwarding labels 16804
Sun Jul 14 06:40:53.151 UTC
Local Outgoing Prefix          Outgoing      Next Hop      Bytes
Label  Label   or ID          Interface     Substituted
----- -----
16804  16804    SR Pfx (idx 804)  Gi0/0/0/0  16.1.1.1      0
RP/0/0/CPU0:PE6#sh mpls forwarding labels 16904
Sun Jul 14 06:40:56.011 UTC
Local Outgoing Prefix          Outgoing      Next Hop      Bytes
Label  Label   or ID          Interface     Substituted
----- -----
16904  16904    SR Pfx (idx 904)  Gi0/0/0/1  67.1.1.7      0
```

```
RP/0/0/CPU0:PE6#sh isis database PE6 internal verbose
Sun Jul 14 06:41:36.538 UTC
.....
  TLV code:135 length:67
.....
    Metric: 0           IP-Extended 192.168.0.6/32
    SubTLV code:3 length:6
      Prefix-SID Index: 6, Algorithm:0, R:0 N:1 P:0 E:0 V:0 L:0
      SubTLV code:3 length:6
        Prefix-SID Index: 806, Algorithm:128, R:0 N:1 P:0 E:0 V:0 L:0
        SubTLV code:3 length:6
          Prefix-SID Index: 906, Algorithm:129, R:0 N:1 P:0 E:0 V:0 L:0
          SubTLV code:4 length:1
            Prefix Attribute Flags: X:0 R:0 N:1
            SubTLV code:11 length:4
              Source Router ID: 192.168.0.6
.....
  TLV code:242 length:49
    Router Cap: 192.168.0.6, D:0, S:0
    SubTLV code:2 length:9
      Segment Routing: I:1 V:1, SRGB Base: 16000 Range: 8000
    SubTLV code:22 length:9
      SR Local Block: Base: 15000 Range: 1000
    SubTLV code:19 length:4
      SR Algorithm:
        Algorithm: 0
        Algorithm: 1
        Algorithm: 128
        Algorithm: 129
      SubTLV code:23 length:2
        Node Maximum SID Depth:
          Subtype: 1, Value: 10
      SubTLV code:26 length:4
        Flex-Algo Definition:
          Algorihm: 128 Metric-Type: 1 Alg-type: 0 Priority: 128
      SubTLV code:26 length:4
        Flex-Algo Definition:
          Algorihm: 129 Metric-Type: 0 Alg-type: 0 Priority: 128
.....
```

Deploy SR Policy with flexalgo

```
RP/0/0/CPU0:PE6#sh segment-routing traffic-eng policy tabular
Mon Apr 15 01:36:36.407 UTC
```

```
SR-TE policy database
```

Color	Endpoint	Admin State	Oper State	Binding SID
20	192.168.0.4	up	up	24013
10	192.168.0.4	up	up	24016

*Color 10 -> odn with latency
Color 20 -> static sr policy*

```
segment-routing
traffic-eng
on-demand color 10
dynamic
no sid-algorithm
!
!
on-demand color 50
dynamic
metric
type igrp
!
!
on-demand color 100
dynamic
metric
type igrp
!
sid-algorithm 128
!
!
on-demand color 110
dynamic
metric
type latency
!
sid-algorithm 129
!
```

Set color at PE4 that face to CE

```
RP/0/0/CPU0:PE6#sh bgp vpng4 un rd 1:1
Mon Apr 15 01:47:43.641 UTC
.....
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
Network          Next Hop          Metric LocPrf Weight Path
Route Distinguisher: 1:1 (default for vrf A)
*->i11.11.11.11/32  192.168.0.4 C:50
                                         0    100      0 10 i
*>i110.1.1.0/24     192.168.0.4      0    100      0 ?
*>i111.1.1.0/24     192.168.0.4      0    100      0 ?
*>i172.16.1.4/32    192.168.0.4      0    100      0 ?
*> 172.16.1.6/32    0.0.0.0          0            32768 ?
*>i192.168.0.10/32   192.168.0.4 C:100
                                         0    100      0 10 i
*>i192.168.0.11/32   192.168.0.4 C:110
                                         0    100      0 11 i
Processed 7 prefixes, 7 paths
RP/0/0/CPU0:PE6#
RP/0/0/CPU0:PE6#show segment-routing traffic-eng policy tabular
Mon Apr 15 01:48:50.206 UTC

SR-TE policy database
-----
Color        Endpoint Admin State Oper State Binding SID
-----  
20          192.168.0.4 up    up      up      24013  
110         192.168.0.4 up    up      up      24010  
100         192.168.0.4 up    up      up      24014  
50          192.168.0.4 up    up      up      24015
```

```
extcommunity-set opaque alo0
  50
end-set
!
extcommunity-set opaque alo128
  100
end-set
!
extcommunity-set opaque alo129
  110
end-set
!
route-policy bgp_color_in
  if destination in (192.168.0.10/32) then
    set extcommunity color alo128
  elseif destination in (192.168.0.11/32) then
    set extcommunity color alo129
  elseif destination in (11.11.11.11/32) then
    set extcommunity color alo0
  endif
  pass
end-policy
!
router bgp 1
  vrf A
    neighbor 110.1.1.10
      address-family ipv4 unicast
        route-policy bgp_color_in in
we can set color at PE6's ingress direction

```

```
RP/0/0/CPU0:PE6#sh segment-routing traffic-eng policy name
srte_c_100_ep_192.168.0.4
Mon Apr 15 02:08:41.685 UTC
```

SR-TE policy database

```
Color: 100, End-point: 192.168.0.4
Name: srte_c_100_ep_192.168.0.4
Status:
  Admin: up Operational: up for 00:21:51 (since Apr 15
01:46:50.275)
```

Candidate-paths:
 Preference: 200 (BGP ODN) (**current**)
 Requested BSID: dynamic

Constraints:
 Prefix-SID Algorithm: 128
 Dynamic (valid)
16804 [Prefix-SID: 192.168.0.4, Algorithm: 128]

Preference: 100 (BGP ODN)
 Requested BSID: dynamic
 Constraints:
 Prefix-SID Algorithm: 128
 Dynamic (pce 192.168.0.9) (valid)
 Metric Type: IGP, Path Accumulated Metric: 2
16004 [Prefix-SID, 192.168.0.4]

Attributes:
 Binding SID: 24014
 Forward Class: 0
 Steering BGP disabled: no
 IPv6 caps enable: yes

```
RP/0/0/CPU0:PE6#sh segment-routing traffic-eng policy name
```

```
srte_c_110_ep_192.168.0.4
Mon Apr 15 02:07:12.581 UTC
```

SR-TE policy database

```
Color: 110, End-point: 192.168.0.4
Name: srte_c_110_ep_192.168.0.4
Status:
  Admin: up Operational: up for 00:20:24 (since Apr 15 01:46:48.265)
Candidate-paths:
  Preference: 200 (BGP ODN) (current)
    Requested BSID: dynamic
    Constraints:
      Prefix-SID Algorithm: 129
      Dynamic (valid)
16904 [Prefix-SID: 192.168.0.4, Algorithm: 129]
  Preference: 100 (BGP ODN)
    Requested BSID: dynamic
    Constraints:
      Prefix-SID Algorithm: 129
      Dynamic (pce 192.168.0.9) (valid)
        Metric Type: LATENCY, Path Accumulated Metric: 150
          16007 [Prefix-SID, 192.168.0.7]
          16001 [Prefix-SID, 192.168.0.1]
          16004 [Prefix-SID, 192.168.0.4]
```

Attributes:
 Binding SID: 24010
 Forward Class: 0
 Steering BGP disabled: no
 IPv6 caps enable: yes

```
RP/0/0/CPU0:PE6#sh mpls for labels 16804
Mon Apr 15 02:12:11.060 UTC
```

Local Label	Outgoing Label	Prefix or ID	Outgoing Interface	Next Hop	Bytes Switched
16804	16804	SR Pfx (idx 804)	Gi0/0/0/0	16.1.1.1	0

```
RP/0/0/CPU0:PE6#sh mpls for labels 16904
```

Local Label	Outgoing Label	Prefix or ID	Outgoing Interface	Next Hop	Bytes Switched
16904	16904	SR Pfx (idx 904)	Gi0/0/0/1	67.1.1.7	0

```
RP/0/0/CPU0:PE6#sh segment-routing traffic-eng policy name  
srte_c_50_ep_192.168.0.4  
Mon Apr 15 02:11:06.395 UTC
```

SR-TE policy database

```
Color: 50, End-point: 192.168.0.4  
Name: srte_c_50_ep_192.168.0.4  
Status:  
Admin: up Operational: up for 00:24:16 (since Apr 15 01:46:50.212)  
Candidate-paths:  
Preference: 200 (BGP ODN) (current)  
Requested BSID: dynamic  
Dynamic (valid)  
Metric Type: IGP, Path Accumulated Metric: 2  
16004 [Prefix-SID, 192.168.0.4]  
Preference: 100 (BGP ODN)  
Requested BSID: dynamic  
Dynamic (pce 192.168.0.9) (valid)  
Metric Type: IGP, Path Accumulated Metric: 2  
16004 [Prefix-SID, 192.168.0.4]  
Attributes:  
Binding SID: 24015  
Forward Class: 0  
Steering BGP disabled: no  
IPv6 caps enable: yes
```

```
RP/0/0/CPU0:PE6#sh mpls for labels 16004  
Mon Apr 15 02:14:26.641 UTC
```

Local Label	Outgoing Label	Prefix or ID	Outgoing Interface	Next Hop	By
16004	16004	SR Pfx (idx 4)	Gi0/0/0/0	16.1.1.1	18064
	16004	SR Pfx (idx 4)	Gi0/0/0/1	67.1.1.7	0

```
RP/0/0/CPU0:PE6#traceroute sr-mpls nil-fec policy name srte_c_20_ep_192.168.0.4  
Mon Apr 15 02:15:30.487 UTC
```

```
0 16.1.1.6 MRU 1500 [Labels: 16007/16004/explicit-null Exp: 0/0/0]  
L 1 16.1.1.1 MRU 1500 [Labels: implicit-null/16004/explicit-null Exp: 0/0/0] 0 ms  
L 2 17.1.1.7 MRU 1500 [Labels: implicit-null/explicit-null Exp: 0/0/0] 0 ms  
! 3 47.1.1.4 10 ms
```

```
RP/0/0/CPU0:PE6#traceroute sr-mpls nil-fec policy name srte_c_50_ep_192.168.0.4  
Mon Apr 15 02:15:36.746 UTC
```

```
0 67.1.1.6 MRU 1500 [Labels: 16004/explicit-null Exp: 0/0]  
L 1 67.1.1.7 MRU 1500 [Labels: implicit-null/explicit-null Exp: 0/0/0] 0 ms  
! 2 47.1.1.4 20 ms
```

```
RP/0/0/CPU0:PE6#traceroute sr-mpls nil-fec policy name srte_c_100_ep_192.168.0.4  
Mon Apr 15 02:15:52.735 UTC
```

```
0 16.1.1.6 MRU 1500 [Labels: 16804/explicit-null Exp: 0/0]  
L 1 16.1.1.1 MRU 1500 [Labels: implicit-null/explicit-null Exp: 0/0/0] 0 ms  
! 2 14.1.1.4 10 ms
```

```
RP/0/0/CPU0:PE6#traceroute sr-mpls nil-fec policy name srte_c_110_ep_192.168.0.4  
Mon Apr 15 02:15:57.055 UTC
```

```
0 67.1.1.6 MRU 1500 [Labels: 16904/explicit-null Exp: 0/0]  
L 1 67.1.1.7 MRU 1500 [Labels: implicit-null/explicit-null Exp: 0/0/0] 10 ms  
! 2 47.1.1.4 10 ms
```


More...we support now!

- BGP IPv4 over SRTE with Flex Algo
- BGP IPv6 over SRTE with Flex Algo
- BGP IPv6 (6PE) over SRTE with Flex Algo
- VPNv4 over SRTE with Flex Algo
- VPNv6 (6VPE) over SRTE with Flex Algo
- EVPN VPWS Single-homing over SRTE with Flex Algo



Egress Peer Engineering(EPE)


Egress Peer Engineering(EPE)



1. Go to AS6's route via P1 at default
2. Check BMP/Netflow, and find top10/5 mapping for prefix and egress port
3. Base on the mapping table, controller sent the route(e.g 1.1.1.0/24) with higher LP and P1 EPE label to head by BGP LU)
4. If link issue between R5 and R6, then R5 send withdraw to R3
5. BGP controller find the issue due to not found in per BGP neighbor BMP DB, so withdraw the route that sent to R1 at former.

Lab7 - Egress Peer Engineering(EPE)

- Remove to global scenario from VRF. And enable eBGP Multipath
- Enable EPE at PE4 and check EPE Label
- Option 1
 - Set up BGP-LU between BGP Controller and head-end PE6, sent route by higher LP, check traffics behavior at PE4
 - Change EPE Label at controller, check again
 - Change color at controller, control Intranet path
- Option 2
 - Set SR Policy with explicit path that include EPE label / SID Descriptor
- Recovery to LDP scenario, double check whether EPE solution workable too



Change to Global Scenario

```
router bgp 1
address-family ipv4 unicast
network 192.168.0.6/32
!
neighbor 192.168.0.4
address-family ipv4 unicast
route-policy RR_IN in
route-policy RR_OUT out
next-hop-self
!
no address-family vpnv4 unicast
!
```

PE6

```
interface Loopback101
no vrf A
no ipv4 address
ipv4 address 172.16.1.4 255.255.255.255
!
interface GigabitEthernet0/0/0/2
no vrf A
no ipv4 address
ipv4 address 111.1.1.1 255.255.255.0
!
interface GigabitEthernet0/0/0/3
no vrf A
no ipv4 address
ipv4 address 110.1.1.1 255.255.255.0
!
```

PE4

```
router bgp 1
bgp bestpath as-path multipath-relax
address-family ipv4 unicast
maximum-paths ebgp 16
redistribute connected
!
neighbor 110.1.1.10
remote-as 10
address-family ipv4 unicast
route-policy bgp_color_in in
route-policy bgp_out out
!
neighbor 111.1.1.11
remote-as 11
address-family ipv4 unicast
route-policy bgp_color_in in
route-policy bgp_out out
!
no vrf A
!
neighbor 192.168.0.6
address-family ipv4 unicast
route-policy RR_IN in
route-policy RR_OUT out
next-hop-self
!
no address-family vpnv4 unicast
!
```

PE4

```

RP/0/0/CPU0:PE6#sh bgp ipv4 un
Mon Apr 15 03:23:22.098 UTC
BGP router identifier 192.168.0.6, local AS number 1
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000000 RD version: 14
BGP main routing table version 14
BGP NSR Initial initsync version 14 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
Network      Next Hop          Metric LocPrf Weight Path
*>i11.11.11/32 192.168.0.4 C:50
*>i14.1.1.0/24 192.168.0.4      0    100      0 10 i
*>i47.1.1.0/24 192.168.0.4      0    100      0 ?
*>i110.1.1.0/24 192.168.0.4     0    100      0 ?
*>i111.1.1.0/24 192.168.0.4     0    100      0 ?
*>i172.16.1.4/32 192.168.0.4     0    100      0 ?
*>i192.168.0.4/32 192.168.0.4     0    100      0 ?
*> 192.168.0.6/32 0.0.0.0        0            32768 i
*>i192.168.0.10/32 192.168.0.4 C:100
*>i192.168.0.11/32 192.168.0.4 C:110
                                         0    100      0 11 i

Processed 10 prefixes, 10 paths

```

```

RP/0/0/CPU0:PE4#sh route 11.11.11.11
Mon Apr 15 03:24:13.064 UTC

Routing entry for 11.11.11.11/32
  Known via "bgp 1", distance 20, metric 0
  Tag 10, type external
  Installed Apr 15 03:18:37.497 for 00:05:35
  Routing Descriptor Blocks
    110.1.1.10, from 110.1.1.10, BGP external, BGP multi path
      Route metric is 0
    111.1.1.11, from 111.1.1.11, BGP external, BGP multi path
      Route metric is 0
      No advertising protos.
RP/0/0/CPU0:PE4#sh bgp 11.11.11.11/32
Mon Apr 15 03:24:22.054 UTC
BGP routing table entry for 11.11.11.11/32
Versions:
  Process      bRIB/RIB  SendTblVer
  Speaker           2           2
Last Modified: Apr 15 03:18:37.859 for 00:05:44
Paths: (2 available, best #1)
  Advertised IPv4 Unicast paths to update-groups (with more than one peer):
    0.3
  Advertised IPv4 Unicast paths to peers (in unique update groups):
    192.168.0.6
  Path #1: Received by speaker 0
  Advertised IPv4 Unicast paths to update-groups (with more than one peer):
    0.3
  Advertised IPv4 Unicast paths to peers (in unique update groups):
    192.168.0.6
  10
    110.1.1.10 from 110.1.1.10 (192.168.0.10)
      Origin IGP, metric 0, localpref 100, valid, external, best, group-best, multipath
      Received Path ID 0, Local Path ID 1, version 2
      Extended community: Color:50
      Origin-AS validity: (disabled)
  Path #2: Received by speaker 0
  Not advertised to any peer
  11
    111.1.1.11 from 111.1.1.11 (192.168.0.11)
      Origin IGP, metric 0, localpref 100, valid, external, multipath
      Received Path ID 0, Local Path ID 0, version 0

```

Enable EPE

```
router bgp 1
  neighbor 110.1.1.10
    egress-engineering
  !
  neighbor 111.1.1.11
    egress-engineering
```

PE4

```
RP/0/0/CPU0:PE4#sh segment-routing traffic-eng ipv4 topology isis
hostname PE4
Mon Apr 15 03:50:12.987 UTC
```

SR-TE topology database

Node 2
.....

```
Link[4]: local address 110.1.1.1, remote address 110.1.1.10
  Local node:
    BGP router ID: 192.168.0.4
  Remote node:
    BGP router ID: 192.168.0.10
  Metric: IGP 0, TE 0
  Bandwidth: Total link 0, Reservable 0
  Admin-groups: 0x00000000
  Adj SID: 24012 (epe)
```

```
Link[5]: local address 111.1.1.1, remote address 111.1.1.11
  Local node:
    BGP router ID: 192.168.0.4
  Remote node:
    BGP router ID: 192.168.0.11
  Metric: IGP 0, TE 0
  Bandwidth: Total link 0, Reservable 0
  Admin-groups: 0x00000000
  Adj SID: 24013 (epe)
```

PE6 have the data?

```
RP/0/0/CPU0:PE4#show bgp egress-engineering
Mon Apr 15 03:47:30.818 UTC
```

```
Egress Engineering Peer Set: 110.1.1.10/32 (15636480)
  Nexthop: 110.1.1.10
  Version: 2, rn_version: 2
  Flags: 0x00000006
  Local ASN: 1
  Remote ASN: 10
  Local RID: 192.168.0.4
  Remote RID: 192.168.0.10
  Local Address: 110.1.1.1
  First Hop: 110.1.1.10
  NHID: 5
  IFH: 0x80
  Label: 24012, Refcount: 3
  rpc_set: 166a77d0, ID: 1
```

```
Egress Engineering Peer Set: 111.1.1.11/32 (1563666c)
  Nexthop: 111.1.1.11
  Version: 3, rn_version: 3
  Flags: 0x00000006
  Local ASN: 1
  Remote ASN: 11
  Local RID: 192.168.0.4
  Remote RID: 192.168.0.11
  Local Address: 111.1.1.1
  First Hop: 111.1.1.11
  NHID: 6
  IFH: 0x60
  Label: 24013, Refcount: 3
  rpc_set: 1670ef58, ID: 2
```

Option 1 – BGP Controller

```
router bgp 1
neighbor 192.168.0.9
remote-as 1
description iBGP peer XTC
update-source Loopback0
address-family ipv4 labeled-unicast
!
```

PE6

```
extcommunity-set opaque frank
 50
end-set
!
route-policy frank
  if destination in (11.11.11.11/32) then
    set next-hop 192.168.0.4
    set local-preference 1000
    set extcommunity color frank
    set label 24012
  endif
end-policy
```

Controller

```
interface Loopback1
  ipv4 address 11.11.11.11 255.255.255.255
!
router bgp 1
  address-family ipv4 unicast
    network 11.11.11.11/32
    allocate-label all
  !
  neighbor 192.168.0.6
    remote-as 1
    update-source Loopback0
    address-family ipv4 labeled-unicast
      route-policy frank out
    !
!
```

Controller

```

RP/0/0/CPU0:PE6#sh bgp ipv4 un
Mon Apr 15 07:07:40.916 UTC
BGP router identifier 192.168.0.6, local AS number 1
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000000 RD version: 16
BGP main routing table version 16
BGP NSR Initial initSync version 14 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
      Network          Next Hop          Metric LocPrf Weight Path
* i11.11.11.11/32    192.168.0.4 C:50
                                         0     100      0 10 i
*:>i                 192.168.0.4 C:50
                                         0     1000      0 i
*:>i14.1.1.0/24      192.168.0.4
                                         0     100      0 ?
*:>i47.1.1.0/24      192.168.0.4
                                         0     100      0 ?
*:>i110.1.1.0/24     192.168.0.4
                                         0     100      0 ?
*:>i111.1.1.0/24     192.168.0.4
                                         0     100      0 ?
*:>i172.16.1.4/32    192.168.0.4
                                         0     100      0 ?
*:>i192.168.0.4/32   192.168.0.4
                                         0     100      0 ?
*> 192.168.0.6/32    0.0.0.0
                                         0           32768 i
*:>i192.168.0.10/32  192.168.0.4 C:100
                                         0     100      0 10 i
*:>i192.168.0.11/32  192.168.0.4
                                         0     100      0 11 i

Processed 10 prefixes, 11 paths

```

```

RP/0/0/CPU0:PE6#sh bgp ipv4 labeled-unicast labels
Mon Apr 15 07:12:56.464 UTC
BGP router identifier 192.168.0.6, local AS number 1
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000000 RD version: 16
BGP main routing table version 16
BGP NSR Initial initSync version 14 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
      Network          Next Hop          Rcvd Label      Local Label
*:>i11.11.11.11/32    192.168.0.4 C:50
                                         24012          nolabel

Processed 1 prefixes, 1 paths

```

```

RP/0/0/CPU0:AS10#sh int gi0/0/0/0 | i packets input
Mon Apr 15 07:17:20.495 UTC
1017 packets input, 115219 bytes, 0 total input drops

```

Change EPE label

```
route-policy frank
  if destination in (11.11.11.11/32) then
    set next-hop 192.168.0.4
    set local-preference 1000
    set extcommunity color frank
    set label 24013
  endif
end-policy
```

Controller

```
RP/0/0/CPU0:PE4#sh mpls for labels 24013 detail
Mon Apr 15 14:13:10.789 UTC
Local Outgoing Prefix Outgoing Next Hop Byt
Label Label or ID Interface Hop Swi
-----
24013 Pop No ID Gi0/0/0/2 111.1.1.11 0
  Updated: Apr 15 13:59:47.674
  Path Flags: 0x6000 [ ]
  Version: 29, Priority: 3
  Label Stack (Top -> Bottom): { Imp-Null }
  NHID: 0x0, Encap-ID: N/A, Path idx: 0, Backup path idx: 0, Weight: 0
  MAC/Encaps: 14/14, MTU: 1500
  Outgoing Interface: GigabitEthernet0/0/0/2 (ifhandle 0x00000060)
  Packets Switched: 0
```

```
RP/0/0/CPU0:PE6#sh bgp ipv4 un labels
Mon Apr 15 14:12:08.714 UTC
BGP router identifier 192.168.0.6, local AS number 1
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000000 RD version: 16
BGP main routing table version 16
BGP NSR Initial initsync version 2 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
Network          Next Hop        Rcvd Label      Local Label
*>i11.11.11.11/32   192.168.0.4 C:50           24013          nolabel
Processed 1 prefixes, 1 paths
```

```
RP/0/0/CPU0:AS11#sh int gi0/0/0/0 | i packets input
Mon Apr 15 14:16:06.246 UTC
  1001 packets input, 114073 bytes, 0 total input drops
RP/0/0/CPU0:AS11#
```

Change SLA to low latency by diff color

```
extcommunity-set opaque frank  
 10  
end-set
```

Controller

```
RP/0/0/CPU0:PE6#sh segment-routing traffic-eng policy name  
srte_c_10_ep_192.168.0.4  
Mon Apr 15 14:23:39.846 UTC
```

SR-TE policy database

```
Color: 10, End-point: 192.168.0.4  
  Name: srte_c_10_ep_192.168.0.4  
  Status:  
    Admin: up  Operational: up for 00:02:22 (since Apr 15  
14:21:17.406)  
  Candidate-paths:  
    Preference: 200 (BGP ODN) (current)  
    Requested BSID: dynamic  
    Dynamic (valid)  
      Metric Type: LATENCY, Path Accumulated Metric: 150  
        16007 [Prefix-SID, 192.168.0.7]  
        16001 [Prefix-SID, 192.168.0.1]  
        16004 [Prefix-SID, 192.168.0.4]
```

```
.....  
Attributes:  
  Binding SID: 24014  
  Forward Class: 0  
  Steering BGP disabled: no  
  IPv6 caps enable: yes
```

```
RP/0/0/CPU0:PE6#sh bgp ipv4 un  
Mon Apr 15 14:21:26.325 UTC  
BGP router identifier 192.168.0.6, local AS number 1  
BGP generic scan interval 60 secs  
Non-stop routing is enabled  
BGP table state: Active  
Table ID: 0xe0000000  RD version: 18  
BGP main routing table version 18  
BGP NSR Initial initSync version 2 (Reached)  
BGP NSR/ISSU Sync-Group versions 0/0  
BGP scan interval 60 secs  
  
Status codes: s suppressed, d damped, h history, * valid, > best  
              i - internal, r RIB-failure, S stale, N Nexthop-discard  
Origin codes: i - IGP, e - EGP, ? - incomplete  
Network          Next Hop            Metric LocPrf Weight Path  
* i11.11.11.11/32 192.168.0.4 C:50  
                                         0     100      0 10 i  
*:>i 192.168.0.4 C:10  
                                         0   1000      0 i  
*:>i14.1.1.0/24 192.168.0.4  
*:>i47.1.1.0/24 192.168.0.4  
*:>i110.1.1.0/24 192.168.0.4  
*:>i111.1.1.0/24 192.168.0.4  
*:>i172.16.1.4/32 192.168.0.4  
*:>i192.168.0.4/32 192.168.0.4  
*> 192.168.0.6/32 0.0.0.0  
*:>i192.168.0.10/32 192.168.0.4 C:100  
*:>i192.168.0.11/32 192.168.0.4 C:110  
Process 10 prefixes, 11 paths
```

Option 2 – Head end

```
router bgp 1
neighbor 192.168.0.6
  address-family link-state link-state
!
!
router bgp 1
  address-family link-state link-state
!
  neighbor 192.168.0.4
    address-family link-state link-state
!
segment-routing
  traffic-eng
    segment-list test-epe
      index 5 address ipv4 192.168.0.4
      index 10 address ipv4 110.1.1.10
!
policy epe
  color 50 end-point ipv4 192.168.0.4
  candidate-paths
    preference 100 <<< change to 300 for testing
    explicit segment-list test-epe
!
```

PE4

PE6

```
RP/0/0/CPU0:PE6#sh segment-routing traffic-eng ipv4 topology isis hostname PE4
Mon Apr 15 14:50:45.595 UTC

SR-TE topology database
-----
Node 2
  TE router ID: 192.168.0.4
  Host name: PE4
  BGP router ID: 192.168.0.4
  ISIS system ID: 1921.6800.0004 level-2
  ISIS system ID: 1921.6800.0004 level-2
  BGP router ID: 192.168.0.4
.....
Link[4]: local address 110.1.1.1, remote address 110.1.1.10
  Local node:
    BGP router ID: 192.168.0.4
  Remote node:
    BGP router ID: 192.168.0.10
    Metric: IGP 0, TE 0
    Bandwidth: Total link 0, Reservable 0
    Admin-groups: 0x00000000
    Adj SID: 24012 (epe)

Link[5]: local address 111.1.1.1, remote address 111.1.1.11
  Local node:
    BGP router ID: 192.168.0.4
  Remote node:
    BGP router ID: 192.168.0.11
    Metric: IGP 0, TE 0
    Bandwidth: Total link 0, Reservable 0
    Admin-groups: 0x00000000
    Adj SID: 24013 (epe)
```

```
RP/0/0/CPU0:PE6#sh segment-routing traffic-eng policy color 50 detail
Mon Apr 15 15:26:15.769 UTC
```

SR-TE policy database

Color: 50, End-point: 192.168.0.4

Name: srte_c_50_ep_192.168.0.4

Status:

Admin: up Operational: up for 00:36:33 (since Apr 15 14:49:42.189)

Candidate-paths:

Preference: 300 (configuration) (current) (reoptimizing)

Name: epe

Requested BSID: dynamic

Explicit: segment-list test-epe (valid)

Weight: 1, Metric Type: TE

16004 [Prefix-SID, 192.168.0.4]

24012 [Adjacency-SID, 110.1.1.1 - 110.1.1.10]

Preference: 200 (BGP ODN)

Requested BSID: dynamic

Dynamic (invalid)

Metric Type: IGP, Path Accumulated Metric: 2

Preference: 100 (BGP ODN)

Requested BSID: dynamic

Dynamic (pce) (invalid)

```
RP/0/0/CPU0:PE6#ping 11.11.11.11 source lo0 repeat 1000
Mon Apr 15 15:28:07.771 UTC
Type escape sequence to abort.
Sending 1000, 100-byte ICMP Echos to 11.11.11.11, timeout is 2 seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Success rate is 100 percent (1000/1000), round-trip min/avg/max = 1/6/29 ms
```

```
RP/0/0/CPU0:AS10#sh int gi0/0/0/0 | i packets input
Mon Apr 15 15:28:16.969 UTC
2163 packets input, 238947 bytes, 0 total input drops
RP/0/0/CPU0:AS10#!
```

Recovery to LDP and check EPE solution

```
RP/0/0/CPU0:PE6#sh segment-routing traffic-eng policy tabular
Mon Apr 15 15:48:48.296 UTC
```

SR-TE policy database

Color	Endpoint	Admin State	Oper State	Binding SID
20	192.168.0.4	up	up	24016
100	192.168.0.4	up	up	24028
50	192.168.0.4	up	up	24017
110	192.168.0.4	up	up	24030

```
RP/0/0/CPU0:PE6#sh cef 192.168.0.1/32
Mon Apr 15 15:53:37.337 UTC
192.168.0.1/32, version 170, labeled SR, internal 0x1000001 0x81 (ptr
0xa1415ae0) [1], 0x0 (0xa13f8440), 0xa28 (0xa1a512c8)
Updated Apr 15 15:40:26.771
local adjacency 16.1.1.1
Prefix Len 32, traffic index 0, precedence n/a, priority 3
Extensions: context-label:16001
```

```
    via 16.1.1.1/32, GigabitEthernet0/0/0/0, 14 dependencies, weight 0
class 0, protected [flags 0x400]
path-idx 0 bkp-idx 1 NHID 0x0 [0xa1bb7210 0xa1bb70f0]
next hop 16.1.1.1/32
local label 24009      labels imposed {ImplNull}
    via 67.1.1.7/32, GigabitEthernet0/0/0/1, 14 dependencies, weight 0
class 0, backup (Local-LFA) [flags 0x300]
path-idx 1 NHID 0x0 [0xa189b270 0x0]
next hop 67.1.1.7/32
local adjacency
local label 24009      labels imposed {24013}
```

```
RP/0/0/CPU0:PE6#sh bgp ipv4 un
Mon Apr 15 15:47:53.460 UTC
Status codes: s suppressed, d damped, h history, * valid, > best
i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
Network          Next Hop          Metric LocPrf Weight Path
*>i11.11.11.11/32 192.168.0.4 C:50
```

*>i14.1.1.0/24	192.168.0.4	0	100	0	10	i
*>i47.1.1.0/24	192.168.0.4	0	100	0	?	
*>i110.1.1.0/24	192.168.0.4	0	100	0	?	
*>i111.1.1.0/24	192.168.0.4	0	100	0	?	
*>i172.16.1.4/32	192.168.0.4	0	100	0	?	
*>i192.168.0.4/32	192.168.0.4	0	100	0	?	
> 192.168.0.6/32	0.0.0.0	0			32768	i
*>i192.168.0.10/32	192.168.0.4 C:100	0	100	0	10	i
*>i192.168.0.11/32	192.168.0.4 C:110	0	100	0	11	i

Processed 10 prefixes, 10 paths

```
RP/0/0/CPU0:P7#sh mpls for labels 24013
```

Local Label	Outgoing Label	Prefix or ID	Outgoing Interface	Next Hop	Bytes Switched
24013	Pop	192.168.0.1/32	Gi0/0/0/2	17.1.1.1	1764

Refer to LDP config, please reference Page
24

```
RP/0/0/CPU0:PE6#sh cef 11.11.11.11/32 detail
```

```
Mon Apr 15 15:59:19.313 UTC
```

```
11.11.11.11/32, version 156, internal 0x5000001 0x0 (ptr 0xa14163a0) [1], 0x0 (0x0), 0x0 (0x0)
```

```
Updated Apr 15 15:40:02.092
```

```
Prefix Len 32, traffic index 0, precedence n/a, priority 4
```

```
gateway array (0xa134fe6c) reference count 1, flags 0x2010, source ri  
[1 type 3 flags 0x48441 (0xa13a07d8) ext 0x0 (0x0)]
```

```
LW-LDI[type=0, refc=0, ptr=0x0, sh-ldi=0x0]
```

```
gateway array update type-time 1 Apr 15 15:40:02.092
```

```
LDI Update time Apr 15 15:40:02.092
```

```
Level 1 - Load distribution: 0
```

```
[0] via 192.168.0.4/32, recursive
```

```
via local-label 24017, 3 dependencies, recursive [flags 0x6000]
```

```
path-idx 0 NHID 0x0 [0xa17cdc3c 0x0]
```

```
recursion-via-label
```

```
next hop via 24017/1/21
```

```
Load distribution: 0 (refcount 1)
```

Hash	OK	Interface	Address
0	Y	srte_c_50_ep_192.168.0.4	point2point

```
RP/0/0/CPU0:PE6#traceroute 11.11.11.11 source lo0
```

```
Mon Apr 15 16:01:55.342 UTC
```

```
Type escape sequence to abort.
```

```
Tracing the route to 11.11.11.11
```

```
1 67.1.1.7 [MPLS: Labels 16004/24012 Exp 0] 9 msec 0 msec 9 msec
```

```
2 47.1.1.4 [MPLS: Label 24012 Exp 0] 0 msec 0 msec 9 msec
```

```
3 110.1.1.10 9 msec * 9 msec
```

```
RP/0/0/CPU0:PE6#sh segment-routing traffic-eng policy name  
srte_c_50_ep_192.168.0.4
```

```
Mon Apr 15 16:02:47.429 UTC
```

```
SR-TE policy database
```

```
-----
```

```
Color: 50, End-point: 192.168.0.4
```

```
Name: srte_c_50_ep_192.168.0.4
```

```
Status:
```

```
Admin: up Operational: up for 00:22:48 (since Apr 15 15:39:59.083)
```

```
Candidate-paths:
```

```
Preference: 300 (configuration) (current)
```

```
Name: epe
```

```
Requested BSID: dynamic
```

```
Explicit: segment-list test-epe (valid)
```

```
Weight: 1, Metric Type: TE
```

```
16004 [Prefix-SID, 192.168.0.4]
```

```
24012 [Adjacency-SID, 110.1.1.1 - 110.1.1.10]
```

```
Preference: 200 (BGP ODN)
```

```
Requested BSID: dynamic
```

```
Dynamic (invalid)
```

```
Last error: unresolved first label (16004)
```

```
Metric Type: IGP, Path Accumulated Metric: 2
```

```
Preference: 100 (BGP ODN)
```

```
Requested BSID: dynamic
```

```
Dynamic (pce) (invalid)
```

```
Metric Type: IGP, Path Accumulated Metric: 2
```

```
Attributes:
```

```
Binding SID: 24017
```

```
Forward Class: 0
```

```
Steering BGP disabled: no
```



```
IPv6 caps enable: yes
```



SRv6 Solution

SRv6 Solution

Network Functions	IPv6
Reachability	IPv6 Header
Engineered Load Balancing	IPv6 Header
VPN	IPv6 Header
Traffic Engineering	IPv6 Header
Source Routing	IPv6 Header
Service Chaining	IPv6 Header



SRv6 Segment格式: Locator+Function

- SRv6 SIDs are 128-bit addresses
 - Locator: most significant bits are used to route the segment to its **parent node**
 - Function: least significant bits identify the **action** to be performed on the parent node
 - Argument [optional]: Last bits can be used as a local function argument
- SIDs have to be **specifically enabled** as such on their parent node
 - A local address **is not** by default a local SID
 - A local SID does not have to be associated with an interface

<i>Locator</i>	<i>Function</i>	<i>Argument</i>
1111 : 2222 : 3333 : 4444 : 5555 : 6666 : 7777 : 8888		

```

segment-routing
  srv6
    encapsulation
      source-address 2001:1111:1::1
    !
    locators
      locator PE_GZ_1
        prefix fabb:eced:1::/64
    !

```

locator configurable



```

RP/0/RP0/CPU0:PE-GZ-1#show segment-routing srv6 sid
Sat Mar 23 15:02:49.360 UTC

```

*** Locator: 'PE_GZ_1' ***

SID	Function	Context	Owner	State	RW
fabb:eced:1: 0:1 ::	End (PSP)	'default':1	sidmgr	InUse	Y
fabb:eced:1: 0:40 ::	End.X (PSP)	[Gi0/0/0/0, Link-Local]	isis-1	InUse	Y
fabb:eced:1: 0:41 ::	End.DT4	'red'	bgp-1	InUse	Y

function auto generated

Lab8 - SRv6 Overlay and Underlay

- Adjust VPP1/2 config script, then restart VPP, and check whether reachable between VPP1/2 and LXC1/2; Change IPv6 from IPv4 on PE1/2/3; Enable SRv6 and allocate locator, and deliver locator in IGP; Enable SR Policy between VPP1 and VPP2, paths are follow:
 - Lxc-1 -> VPP1 -> PE1 -> PE2 -> PE3 -> VPP2 -> Lxc-2
 - Lxc-2 -> VPP2 -> PE3 -> PE1 -> VPP1 -> Lxc-1


BB: SRv6 + IGP/IPv6 (SRv6 Transport)
Cu: IPv6 traffics

- Enable SR MPLS and 6vPE; Change SRH at VPP, only include VPP1/2's locator; then control SRv6 traffics by SR MPLS by 6vPE solution

BB: IPv4 SR MPLS, 6vPE
Cu: IPv6 traffics

- Remove SR MPLS, add IPv4 address at gi0/0/0/0 of PE1, and gi0/0/0/2 of PE3, shutdown link between PE1 and PE3; then enable SRv6 L3VPN

BB: SRv6 + VPNv4 (DX4)
Cu: IPv4 traffics



Linux SRv6实战（第三篇）多云环境下
Overlay(VPP) 和Underlay整合测试

```
#hostnamectl set-hostname frank-vpp1
```

```
[root@frank-vpp1 cisco]# more /opt/vpp-interface-up.txt
set interface state GigabitEthernet0/5/0 up
set interface ip address GigabitEthernet0/5/0 2001:1a::1/126
ip route add ::/0 via 2001:1a::2
set interface state GigabitEthernet0/4/0 up
set interface ip address GigabitEthernet0/4/0 10.0.0.1/24
loopback create-interface
set interface ip address loop0 fc00:1::1/128
set interface state loop0 up
```

```
[root@frank-vpp1 vpp]# vppctl show interface addr
```

```
GigabitEthernet0/4/0 (up):
```

```
    L3 10.0.0.1/24
```

```
GigabitEthernet0/5/0 (up):
```

```
    L3 2001:1a::1/64
```


```
local0 (dn):
```

```
loop0 (up):
```

```
    L3 fc00:1::1/64
```

```
[root@frank-vpp1 vpp]#
```

```
[root@frank-vpp1 vpp]# vppctl
```



```
vpp# ping 10.0.0.2
```

```
116 bytes from 10.0.0.2: icmp_seq=1 ttl=64 time=.3462 ms
116 bytes from 10.0.0.2: icmp_seq=2 ttl=64 time=.1119 ms
```

```
[root@frank-vpp1 vpp]# ls /etc/vpp/startup.conf*
/etc/vpp/startup.conf  /etc/vpp/startup.conf.bak
[root@frank-vpp1 vpp]# diff -ud startup.conf startup.conf.bak
--- startup.conf 2019-07-09 03:28:30.712907659 +0000
+++ startup.conf.bak 2019-07-09 03:26:56.345628976 +0000
@@ -5,7 +5,6 @@
    full-coredump
    cli-listen /run/vpp/cli.sock
    gid vpp
-   startup-config /opt/vpp-interface-up.txt
}


api-trace {
@@ -81,7 +80,7 @@
    # default data-size 2048
#
-dpdk {
+## dpdk {
## Change default settings for all interfaces
# dev default {
## Number of receive queues, enables RSS
@@ -105,8 +104,6 @@
## Whitelist specific interface by specifying PCI address
# dev 0000:02:00.0
-   dev 0000:00:04.0
-   dev 0000:00:05.0

## Blacklist specific device type by specifying PCI vendor:device
## Whitelist entries take precedence
@@ -148,7 +145,7 @@
## Disables UDP / TCP TX checksum offload. Typically needed for
use
## faster vector PMDs (together with no-multi-seg)
# no-tx-checksum-offload
-
```

```
#hostnamectl set-hostname frank-vpp2
```

```
[root@frank-vpp2 cisco]# more /opt/vpp-interface-up.txt
set interface state GigabitEthernet0/4/0 up
set interface ip address GigabitEthernet0/4/0 2001:2c::1/126
ip route add ::/0 via 2001:2c::2
set interface state GigabitEthernet0/5/0 up
set interface ip address GigabitEthernet0/5/0 10.0.1.1/24
loopback create-interface
set interface ip address loop0 fc00:2::1/128
set interface state loop0 up
```

```
[root@frank-vpp2 cisco]# systemctl restart vpp
[root@frank-vpp2 cisco]# vppctl show interface addr
GigabitEthernet0/4/0 (up):
  L3 2001:2b::1/126
GigabitEthernet0/5/0 (up):
  L3 10.0.1.1/24
local0 (dn):
loop0 (up):
  L3 fc00:2::1/128
[root@frank-vpp2 cisco]# vppctl
```



```
vpp# ping 10.0.1.2
116 bytes from 10.0.1.2: icmp_seq=2 ttl=64 time=.1829 ms
116 bytes from 10.0.1.2: icmp_seq=3 ttl=64 time=.1496 ms
[root@frank-vpp2 cisco]#
```

```
[root@frank-vpp2 vpp]# ls /etc/vpp/startup.conf*
/etc/vpp/startup.conf  /etc/vpp/startup.conf.bak
[root@frank-vpp2 vpp]# diff -ud startup.conf startup.conf.bak
--- startup.conf 2019-07-09 04:35:17.966905509 +0000
+++ startup.conf.bak 2019-07-09 04:33:28.686905509 +0000
@@ -5,7 +5,6 @@
  full-coredump
  cli-listen /run/vpp/cli.sock
  gid vpp
- startup-config /opt/vpp-interface-up.txt
}

api-trace {
@@ -81,7 +80,7 @@
  # default data-size 2048
#
-dpdk {
+## dpdk {
    ## Change default settings for all interfaces
    # dev default {
    ## Number of receive queues, enables RSS
@@ -105,8 +104,6 @@
    ## Whitelist specific interface by specifying PCI address
    # dev 0000:02:00.0
-      dev 0000:00:04.0
-      dev 0000:00:05.0
+
## Blacklist specific device type by specifying PCI vendor:device
## Whitelist entries take precedence
@@ -148,7 +145,7 @@
    ## Disables UDP / TCP TX checksum offload. Typically needed for
use
    ## faster vector PMDs (together with no-multi-seg)
    # no-tx-checksum-offload
-
```

```
interface Loopback0
  ipv6 address fc00:a::1/128
!
interface GigabitEthernet0/0/0/0
  no ipv4 address
  ipv6 address 2001:1a::2/126
!
interface GigabitEthernet0/0/0/1
  no ipv4 address
  ipv6 address 2001:12::1/126
!
interface GigabitEthernet0/0/0/2
  no ipv4 address
  ipv6 address 2001:13::1/126
!
router isis 1
  no address-family ipv4 unicast
  address-family ipv6 unicast
    metric-style wide
!
interface Loopback0
  no address-family ipv4 unicast
  address-family ipv6 unicast
!
!
interface GigabitEthernet0/0/0/1
  no address-family ipv4 unicast
  address-family ipv6 unicast
    metric 1
!
!
interface GigabitEthernet0/0/0/2
  no address-family ipv4 unicast
  address-family ipv6 unicast
    metric 1
!
!
```

PE1

```
interface Loopback0
  ipv6 address fc00:c::1/128
!
interface GigabitEthernet0/0/0/0
  no ipv4 address
  ipv6 address 2001:13::3/126
!
interface GigabitEthernet0/0/0/1
  no ipv4 address
  ipv6 address 2001:23::3/126
!
interface GigabitEthernet0/0/0/2
  no ipv4 address
  ipv6 address 2001:2c::2/126
!
router isis 1
  no address-family ipv4 unicast
  address-family ipv6 unicast
    metric-style wide
!
interface Loopback0
  no address-family ipv4 unicast
  address-family ipv6 unicast
!
!
interface GigabitEthernet0/0/0/0
  no address-family ipv4 unicast
  address-family ipv6 unicast
    metric 1
!
!
interface GigabitEthernet0/0/0/1
  no address-family ipv4 unicast
  address-family ipv6 unicast
    metric 1
!
!
```

PE3

```
interface Loopback0
  ipv6 address fc00:b::1/128
!
interface GigabitEthernet0/0/0/0
  no ipv4 address
  ipv6 address 2001:12::2/126
!
interface GigabitEthernet0/0/0/1
  no ipv4 address
  ipv6 address 2001:23::2/126
!
router isis 1
  no address-family ipv4 unicast
  address-family ipv6 unicast
    metric-style wide
!
interface Loopback0
  no address-family ipv4 unicast
  address-family ipv6 unicast
!
!
interface GigabitEthernet0/0/0/0
  no address-family ipv4 unicast
  address-family ipv6 unicast
    metric 1
!
!
interface GigabitEthernet0/0/0/1
  no address-family ipv4 unicast
  address-family ipv6 unicast
    metric 1
!
!
```

PE2

Cisco C

```
segment-routing
srv6
encapsulation
source-address fc00:a::1
!
locators
 locator PE1
 prefix fabb:eced:1::/64
!
!
!
```

PE1

```
segment-routing
srv6
encapsulation
source-address fc00:b::1
!
locators
 locator PE2
 prefix fabb:eced:2::/64
!
!
```

PE2

```
segment-routing
srv6
encapsulation
source-address fc00:c::1
!
locators
 locator PE3
 prefix fabb:eced:3::/64
!
!
```

PE3

```
RP/0/RP0/CPU0:PE1#sh route ipv6
Tue Jul  9 05:44:38.751 UTC

.....
i L2 fc00:b::1/128
[115/1] via fe80::f816:3eff:fea0:a2c0, 00:33:54, GigabitEthernet0/0/0/1
i L2 fc00:c::1/128
[115/1] via fe80::f816:3eff:fe82:53fe, 00:35:39, GigabitEthernet0/0/0/2
```

```
RP/0/RP0/CPU0:PE1#show segment-routing srv6 locator
```

```
Tue Jul  9 05:45:42.787 UTC
```

Name	ID	Prefix	Status
PE1*	2	fabb:eced:1::/64	Up

```
RP/0/RP0/CPU0:PE1#
```

```
RP/0/RP0/CPU0:PE1#show segment-routing srv6 sid
```

```
Tue Jul  9 05:45:48.986 UTC
```

```
*** Locator: 'PE1' ***
```

SID	Function	Context	Owner	State	RW
fabb:eced:1:0:1::	End (PSP)	'default':1	sidmgr	InUse	Y

```
RP/0/RP0/CPU0:PE1#sh isis database internal verbose PE1 | i "END SID"
```

```
Tue Jul  9 05:53:53.696 UTC
```

*Have locator and
SID, but not in IGP*

```

router isis 1
address-family ipv6 unicast
segment-routing srv6
locator PE1
!
!
```

PE1

```

router isis 1
address-family ipv6 unicast
segment-routing srv6
locator PE2
!
!
```

PE2

```

router isis 1
address-family ipv6 unicast
segment-routing srv6
locator PE3
!
!
```

PE3

```

RP/0/RP0/CPU0:PE1#sh route ipv6
Tue Jul  9 06:04:13.041 UTC
.....
i L2 fabb:eced:2::/64
[115/2] via fe80::f816:3eff:fea0:a2c0, 00:00:30, GigabitEthernet0/0/0/1
i L2 fabb:eced:3::/64
[115/2] via fe80::f816:3eff:fe82:53fe, 00:06:13, GigabitEthernet0/0/0/2
i L2 fc00:b::1/128
[115/1] via fe80::f816:3eff:fea0:a2c0, 00:40:44, GigabitEthernet0/0/0/1
i L2 fc00:c::1/128
[115/1] via fe80::f816:3eff:fe82:53fe, 00:42:30, GigabitEthernet0/0/0/2

```

```

RP/0/RP0/CPU0:PE1#
RP/0/RP0/CPU0:PE1#show segment-routing srv6 locator
Tue Jul  9 06:11:07.939 UTC

```

Name	ID	Prefix	Status
PE1*	3	fabb:eced:1::/64	Up

```

RP/0/RP0/CPU0:PE1#
RP/0/RP0/CPU0:PE1#show segment-routing srv6 sid
Tue Jul  9 06:11:15.916 UTC

```

*** Locator: 'PE1' ***

SID	Function	Context	Owner	State	RW
fabb:eced:1:0:1::	End (PSP)	'default':1	sidmgr	InUse	Y
fabb:eced:1:0:40::	End.X (PSP)	[Gi0/0/0/1, Link-Local]	isis-1	InUse	Y
fabb:eced:1:0:41::	End.X (PSP)	[Gi0/0/0/2, Link-Local]	isis-1	InUse	Y

```
RP/0/RP0/CPU0:PE1#
```

```
RP/0/RP0/CPU0:PE1#sh isis database internal verbose | i "END SID"
```

```
Tue Jul  9 08:43:05.484 UTC
```

```

END SID: fabb:eced:1:0:1:: End (PSP)
END SID: fabb:eced:1:0:1:: End (PSP)
END SID: fabb:eced:2:0:1:: End (PSP)
END SID: fabb:eced:3:0:1:: End (PSP)

```


Add follow script to “/opt/vpp-interface-up.txt” in VPP1/2:

```
set sr encaps source addr fc00:1::1  
sr localsid address fabb:1a::1a behavior end.dx4 GigabitEthernet0/4/0 10.0.0.2  
sr policy add bsid fabb:1a::11:1 next fabb:eced:1:0:1:: next fabb:eced:2:0:1:: next fabb:eced:3:0:1:: next fabb:2b::2b encaps  
sr steer l3 10.0.0.0/24 via bsid fabb:1a::11:1
```

VPP1

```
set sr encaps source addr fc00:2::1  
sr localsid address fabb:2b::2b behavior end.dx4 GigabitEthernet0/5/0 10.0.1.2  
sr policy add bsid fabb:2b::22:1 next fabb:eced:3:0:1:: next fabb:eced:1:0:1:: next fabb:1a::1a encaps  
sr steer l3 10.0.1.0/24 via bsid fabb:2b::22:1
```

VPP2



The cmd “sr encaps source” cmd must at first, otherwise not workable...
That maybe a VPP bug

```

vpp# show sr policies
SR policies:
[0].- BSID: fabb:1a::11:1
Behavior: Encapsulation
Type: Default
FIB table: 0
Segment Lists:
[0].- < fabb:eced:1:0:1::, fabb:eced:2:0:1::, fabb:eced:3:0:1::,
      fabb:2b::2b > weight: 1
-----
vpp# show sr localsids
SRv6 - My LocalSID Table:
=====
Address: fabb:1a::1a
Behavior: DX4 (Endpoint with decapsulation and IPv4 cross-connect)
Iface: GigabitEthernet0/4/0
Next hop: 10.0.0.2
Good traffic: [0 packets : 0 bytes]
Bad traffic: [0 packets : 0 bytes]
-----
vpp# show sr steering-policies
SR steering policies:
Traffic SR policy BSID
L3 10.0.1.0/24 fabb:1a::11:1

```

VPP1

```

vpp# show sr policies
SR policies:
[0].- BSID: fabb:2b::22:1
Behavior: Encapsulation
Type: Default
FIB table: 0
Segment Lists:
[0].- < fabb:eced:3:0:1::, fabb:eced:1:0:1::, fabb:1a::1a > weight: 1
-----
vpp# show sr localsids
SRv6 - My LocalSID Table:
=====
Address: fabb:2b::2b
Behavior: DX4 (Endpoint with decapsulation and IPv4 cross-connect)
Iface: GigabitEthernet0/5/0
Next hop: 10.0.1.2
Good traffic: [0 packets : 0 bytes]
Bad traffic: [0 packets : 0 bytes]
-----
vpp# show sr steering-policies
SR steering policies:
Traffic SR policy BSID
L3 10.0.0.0/24 fabb:2b::22:1

```

VPP2

```
route add -net 10.0.0.0/24 gw 10.0.1.1 dev eth1
```


```
route add -net 10.0.1.0/24 gw 10.0.0.1 dev eth1
```

```
router static
address-family ipv6 unicast
fabb:1a::1a/128 2001:1a::1
!
!
```

```
router static
address-family ipv6 unicast
fabb:2b::2b/128 2001:2c::1
!
!
```

Backbone: SRv6 + IGP/IPv6

Customer: IPv6 traffics



```

vrf A
address-family ipv6 unicast
import route-target 1:1
export route-target 1:1
!
interface GigabitEthernet0/0/0/0
no ipv6 address 2001:1a::2/126
vrf A
ipv6 address 2001:1a::2/126
!
interface GigabitEthernet0/0/0/1
ipv4 address 12.1.1.1 255.255.255.0
interface GigabitEthernet0/0/0/2
ipv4 address 13.1.1.1 255.255.255.0
!
router static
no address-family ipv6 unicast
vrf A
address-family ipv6 unicast fabb:1a::1a/128 2001:1a::1
!
router isis 1
address-family ipv4 unicast
metric-style wide
router-id Loopback0
segment-routing mpls
!
interface Loopback0
address-family ipv4 unicast
metric 1
prefix-sid index 1
!
interface GigabitEthernet0/0/0/1
address-family ipv4 unicast
metric 1
!
interface GigabitEthernet0/0/0/2
address-family ipv4 unicast
metric 1
!
```

PE1

```

router bgp 1
address-family ipv4 unicast
address-family vpng6 unicast
!
neighbor 192.168.0.3
no address-family ipv4 unicast
address-family vpng6 unicast
!
vrf A
rd 1:1
address-family ipv6 unicast
no redistribute connected
redistribute static
!
```

Confidential

```

vrf A
address-family ipv6 unicast
import route-target 1:1
export route-target 1:1
!
interface GigabitEthernet0/0/0/2
no ipv6 address 2001:2c::2/126
vrf A
ipv6 address 2001:2c::2/126
!
interface GigabitEthernet0/0/0/1
ipv4 address 23.1.1.3 255.255.255.0
interface GigabitEthernet0/0/0/0
ipv4 address 13.1.1.3 255.255.255.0
!
router static
no address-family ipv6 unicast
vrf A
address-family ipv6 unicast fabb:2b::2b/128 2001:2c::1
!
router isis 1
address-family ipv4 unicast
metric-style wide
router-id Loopback0
segment-routing mpls
!
interface Loopback0
address-family ipv4 unicast
metric 1
prefix-sid index 3
!
interface GigabitEthernet0/0/0/0
address-family ipv4 unicast
metric 1
!
interface GigabitEthernet0/0/0/1
address-family ipv4 unicast
metric 1
!
```

PE3

```

router bgp 1
address-family ipv4 unicast
address-family vpng6 unicast
!
neighbor 192.168.0.1
no address-family ipv4 unicast
address-family vpng6 unicast
!
vrf A
rd 2:2
address-family ipv6 unicast
no redistribute connected
redistribute static
!
```

6vPE

```

interface GigabitEthernet0/0/0/0
ipv4 address 12.1.1.2 255.255.255.0
!
interface GigabitEthernet0/0/0/1
ipv4 address 23.1.1.2 255.255.255.0
!
router isis 1
address-family ipv4 unicast
metric-style wide
router-id Loopback0
segment-routing mpls
!
interface Loopback0
address-family ipv4 unicast
metric 1
prefix-sid index 2
!
!
interface GigabitEthernet0/0/0/0
address-family ipv4 unicast
metric 1
!
!
interface GigabitEthernet0/0/0/1
address-family ipv4 unicast
metric 1
!
```

PE2

SR MPLS Active, 6vPE Active

- For detail SR status verify, please reference page 22, Lab3.

```
RP/0/RP0/CPU0:PE1#sh route vrf A ipv6
Fri Jul 12 02:57:44.123 UTC
```

Codes: C - connected, S - static, R - RIP, B - BGP, (>) - Diversion path
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
i - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, su - IS-IS summary null, * - candidate default
U - per-user static route, o - ODR, L - local, G - DAGR, l - LISP
A - access/subscriber, a - Application route
M - mobile route, r - RPL, t - Traffic Engineering, (!) - FRR Backup
path

Gateway of last resort is not set

C 2001:1a::/126 is directly connected,
 00:41:52, GigabitEthernet0/0/0/0

L 2001:1a::2/128 is directly connected,
 00:41:52, GigabitEthernet0/0/0/0

S fabb:1a::1a/128
 [1/0] via 2001:1a::1, 00:41:52

B fabb:2b::2b/128
 [200/0] via ::ffff:192.168.0.3 (nexthop in vrf default), 00:11:49

```
RP/0/RP0/CPU0:PE1#
```

```
RP/0/RP0/CPU0:PE1#sh bgp vpnv6 un
Fri Jul 12 03:09:39.404 UTC
BGP router identifier 192.168.0.1, local AS number 1
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0    RD version: 0
BGP main routing table version 14
BGP NSR Initial initSync version 5 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
BGP scan interval 60 secs
```

Status codes: s suppressed, d damped, h history, * valid, > best
i - internal, r RIB-failure, S stale, N Nexthop-discard

Origin codes: i - IGP, e - EGP, ? - incomplete

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 1:1 (default for vrf A)					
*> fabb:1a::1a/128	2001:1a::1	0		32768	?
*>ifabb:2b::2b/128	192.168.0.3	0	100	0	?
Route Distinguisher: 2:2					
*>ifabb:2b::2b/128	192.168.0.3	0	100	0	?

Processed 3 prefixes, 3 paths

Update SID List in VPP1/2 config script

```
sr policy add bsid fabb:1a::11:1 next fabb:eced:1:0:1:: next fabb:eced:2:0:1:: next fabb:eced:3:0:1:: next fabb:2b::2b encaps
```

sr policy add bsid fabb:1a::11:1 next fabb:2b::2b encaps

```
sr policy add bsid fabb:2b::22:1 next fabb:eced:3:0:1:: next fabb:eced:1:0:1:: next fabb:1a::1a encaps
```

sr policy add bsid fabb:1a::11:1 next fabb:2b::2b encaps

```
cisco@lxc-1:~$ ping 10.0.1.2
PING 10.0.1.2 (10.0.1.2) 56(84) bytes of data.
64 bytes from 10.0.1.2: icmp_seq=1 ttl=62 time=119 ms
64 bytes from 10.0.1.2: icmp_seq=2 ttl=62 time=106 ms
64 bytes from 10.0.1.2: icmp_seq=3 ttl=62 time=117 ms
64 bytes from 10.0.1.2: icmp_seq=4 ttl=62 time=115 ms
64 bytes from 10.0.1.2: icmp_seq=5 ttl=62 time=44.5 ms
^C
```

RP/0/RP0/CPU0:PE1#sh mpls for Fri Jul 12 05:30:56.206 UTC					
Local Label	Outgoing Label	Prefix or ID	Outgoing Interface	Next Hop	Bytes Switched

-	16002	Pop SR Pfx (idx 2)	Gi0/0/0/1	12.1.1.2	0
	16003	Pop SR Pfx (idx 3)	Gi0/0/0/2	13.1.1.3	20325
	24000	Unlabelled fabb:1a::1a/128[V]	Gi0/0/0/0	2001:1a::1	0
	24001	Pop SR Adj (idx 1)	Gi0/0/0/2	13.1.1.3	0
	24002	Pop SR Adj (idx 3)	Gi0/0/0/2	13.1.1.3	0
	24003	Pop SR Adj (idx 1)	Gi0/0/0/1	12.1.1.2	0
	24004	Pop SR Adj (idx 3)	Gi0/0/0/1	12.1.1.2	0

Control SRv6 traffics by SR MPLS

- Same as lab5, page 73 ; only different is SRv6 traffics

```
segment-routing
traffic-eng
on-demand color 10
dynamic
!
metric
type latency
!
performance-measurement
interface GigabitEthernet0/0/0/1
delay-measurement
advertise-delay 50
!
!
interface GigabitEthernet0/0/0/2
delay-measurement
advertise-delay 500
!
```

PE1

```
extcommunity-set opaque odn
10
end-set
!
route-policy color_in
set extcommunity color odn
end-policy
!
router bgp 1
neighbor 192.168.0.1
address-family vpnv6 unicast
route-policy color_in_out
!
!
!
performance-measurement
interface GigabitEthernet0/0/0/0
delay-measurement
advertise-delay 500
!
!
interface GigabitEthernet0/0/0/1
delay-measurement
advertise-delay 50
!
```

PE4

```
performance-measurement
interface GigabitEthernet0/0/0/0
delay-measurement
advertise-delay 50
!
!
interface GigabitEthernet0/0/0/1
delay-measurement
advertise-delay 50
```

PE2

Backbone: IPv4 SR MPLS, 6vPE
Customer: IPv6 traffics



10	12.463842	10.0.0.2	10.0.1.2	IC
▶ Frame 6: 146 bytes on wire (1168 bits), 146 bytes captured (1168 bits)				
▶ Ethernet II, Src: fa:16:3e:16:1f:88 (fa:16:3e:16:1f:88), Dst: fa:16:3e:a0:a2:00 (fa:16:3e:a0:a2:00)				
▶ MultiProtocol Label Switching Header, Label: 16003, Exp: 0, S: 0, TTL: 62				
▶ MultiProtocol Label Switching Header, Label: 24002, Exp: 0, S: 1, TTL: 62				
▶ Internet Protocol Version 6, Src: fc00:1::1, Dst: fabb:2b::2b (fc00:1::1, fabb:2b::2b)				
▶ Internet Protocol Version 4, Src: 10.0.0.2, Dst: 10.0.1.2 (10.0.0.2, 10.0.1.2)				
▶ Internet Control Message Protocol				


```
RP/RP0/CPU0#sh segment-routing traffic-eng forwarding policy tabular
Fri Jul 12 08:27:12.786 UTC
```

Color	Endpoint	Segment List	Outgoing Label	Outgoing Interface	Next Hop	Bytes Switched	Pure Backup
10	192.168.0.3	dynamic	16003	Gi0/0/0/1	12.1.1.2	2688	

Add IPv4 address and enable DX4 (SRv6 L3VPN)

```
vrf A
address-family ipv4 unicast
import route-target
1:1
!
export route-target
1:1
!
!
!
interface GigabitEthernet0/0/0/0
ip4 address 172.16.1.1 255.255.255.0
!
router bgp 1
address-family vpnv4 unicast
segment-routing srv6
locator PE1
!
!
neighbor fc00:c::1
remote-as 1
description iBGP peer PE3
update-source Loopback0
address-family vpnv4 unicast
!
!
no neighbor 192.168.0.3
no vrf A
vrf A
rd 1:1
address-family ipv4 unicast
segment-routing srv6
alloc mode per-ce
!
redistribute connected
```

```
vrf A
address-family ipv4 unicast
import route-target
1:1
!
export route-target
1:1
!
!
!
interface GigabitEthernet0/0/0/2
ip4 address 172.16.2.1 255.255.255.0
!
router bgp 1
address-family vpnv4 unicast
segment-routing srv6
locator PE3
!
!
neighbor fc00:a::1
remote-as 1
description iBGP peer PE1
update-source Loopback0
address-family vpnv4 unicast
!
!
no neighbor 192.168.0.1
no vrf A
vrf A
rd 1:1
address-family ipv4 unicast
segment-routing srv6
alloc mode per-ce
!
redistribute connected
```



Backbone: SRv6 + VPNv4
Customer: IPv4 traffics

```

RP/0/RP0/CPU0:PE1#sh bgp vpng4 un
Fri Jul 12 10:22:11.531 UTC
BGP router identifier 192.168.0.1, local AS number 1
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0 RD version: 0
BGP main routing table version 6
BGP NSR Initial initsync version 6 (Not Reached)
BGP NSR/ISSU Sync-Group versions 0/0
BGP scan interval 60 secs

```

```

Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
  Network          Next Hop          Metric LocPrf Weight Path
Route Distinguisher: 1:1 (default for vrf A)
*> 172.16.1.0/24      0.0.0.0            0        32768 ?
*>i172.16.2.0/24    fc00:c::1            0       100      0 ?

```

```

RP/0/RP0/CPU0:PE1#show segment-routing srv6 sid
Fri Jul 12 10:42:10.382 UTC

```

*** Locator: 'PE1' ***

Please compare with page 113

SID	Function	Context	Owner	State	RW
fabb:eced:1:0:1::	End (PSP)	'default':1	sidmgr	InUse	Y
fabb:eced:1:0:40::	End.X (PSP)	[Gi0/0/0/1, Link-Local]	isis-1	InUse	Y
fabb:eced:1:0:42::	End.DT4	'A'	bgp-1	InUse	Y

```

RP/0/RP0/CPU0:PE1#sh cef vrf A 172.16.2.0/24 detail
Fri Jul 12 10:39:43.689 UTC
17.2.0/24, version 3, SRv6 Transit, internal 0x5000001 0x0 (ptr 0xe2cb69c)
[1], 0x0 (0xe499fa8), 0x0 (0xf3a01a8)
Updated Jul 12 10:20:22.821
Prefix Len 24, traffic index 0, precedence n/a, priority 3
gateway array (0xf34a0a8) reference count 1, flags 0x10, source rib (7), 0
backups
[2 type 3 flags 0x441 (0xe3b4150) ext 0x0 (0x0)]
LW-LDI[type=3, refc=1, ptr=0xe499fa8, sh-ldi=0xe3b4150]
gateway array update type-time 1 Jul 12 10:20:22.820
LDI Update time Jul 12 10:27:08.798
LW-LDI-TS Jul 12 10:27:08.798

Level 1 - Load distribution: 0
[0] via fabb:eced:3::/128, recursive
via fabb:eced:3::/128, 3 dependencies, recursive [flags 0x6000]
path-idx 0 NHID 0x0 [0xe02e29c 0x0]
next hop VRF - 'default', table - 0xe0800000
next hop fabb:eced:3::/128 via fabb:eced:3::/64
SRv6 T.Encaps.Red SID-list {fabb:eced:3:0:42::}
```

Load distribution: 0 (refcount 2)

Hash	OK	Interface	Address
0	Y	GigabitEthernet0/0/0/1	remote

```
RP/0/RP0/CPU0:PE1#ping vrf A 172.16.2.1
Fri Jul 12 10:49:39.225 UTC
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.2.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 54/59/64 ms
```

Due to T.Encaps.Red, remove last SRH, so no SRH

	4 1.519468	172.16.1.1	172.16.2.1	ICMP
►	Frame 4: 154 bytes on wire (1232 bits), 154 bytes captured (1232 bits)			
►	Ethernet II, Src: fa:16:3e:16:1f:88 (fa:16:3e:16:1f:88), Dst: fa:16:3e:a0:a2:c0 (fa:16:			
▼	Internet Protocol Version 6, Src: fc00:a::1, Dst: <u>fabb:eced:3:0:42::</u>			
	0110 = Version: 6			
► 0000 0000 = Traffic Class: 0x00 (DSCP: CS0, ECN: Not-E			
 1110 1000 0110 0100 1110 = Flow Label: 0xe864e			
	Payload Length: 100			
	Next Header: <u>IPIP (4)</u>			
	Hop Limit: 255			
	Source: fc00:a::1			
	Destination: fabb:eced:3:0:42::			
►	Internet Protocol Version 4, Src: 172.16.1.1, Dst: 172.16.2.1			
►	Internet Control Message Protocol			

TI-LFA for SRv6 Locator


- Recovery link between PE1 and PE3, and enable Ti-LFA on PE1

```
router isis 1
interface GigabitEthernet0/0/0/1
address-family ipv6 unicast
fast-reroute per-prefix
fast-reroute per-prefix ti-lfa
!
!
interface GigabitEthernet0/0/0/2
address-family ipv6 unicast
fast-reroute per-prefix
fast-reroute per-prefix ti-lfa
```

*SRv6 TILFA only work
on NCS5500 now.*

```
RP/0/RP0/CPU0:PE1#sh isis ipv6 fast-reroute fabb:eced:3::/64 det
Fri Jul 12 13:09:59.526 UTC

L2 fabb:eced:3::/64 [2/115] low priority
    via fe80::f816:3eff:fea0:a2c0, GigabitEthernet0/0/0/1, PE3, SRGB Base: 16000,
Weight: 0
    Backup path: LFA, via fe80::f816:3eff:fea0:a2c0, GigabitEthernet0/0/0/1,
PE2, SRGB Base: 16000, Weight: 0, Metric: 3
        P: No, TM: 3, LC: No, NP: No, D: No, SRLG: Yes
        src PE3.00-00, fc00:c::1
```



Official support from 6.6.1

More...we support now!

- L3VPN (VPNv4) Support
- L3VPN v4 QoS Marking + Priority Queuing
- SRv6 OAM with SRH (ping, trace)
- ISIS TI-LFA (NCS5500)
- ISIS for SRv6 Control Plane / Transport
- SRV6 onXRv9K

https://www.cisco.com/c/en/us/td/docs/iosxr/ncs5500/segment-routing/66x/b-segment-routing-cg-ncs5500-66x/b-segment-routing-cg-ncs5500-66x_chapter_011.html

