



# Segment Routing - theory & usage

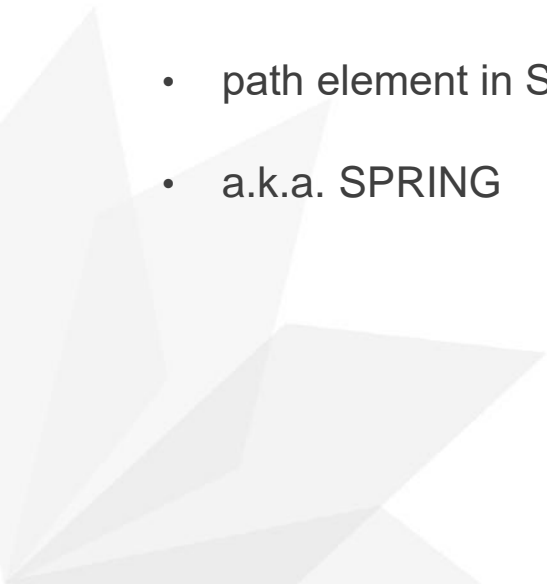
Tomáš Kubina / Orange Business Services  
tomas.kubina@orange.com

# Segment Routing - theory & usage



## What is Segment Routing

- source routing principle
  - source node selects path to destination node and put it into packet header
  - rest of nodes forward packet based on its encoded path
  - order of elements in path is important
- path element in SR world is called “segment”
- a.k.a. SPRING



# Segment Routing - theory & usage



## Why Segment Routing

- make network easier to operate & scale
  - no LDP
  - no RSVP (no tunnel signalling and states on each node)
  - keep IGP
- SRv6 removes also MPLS using only IPv6 address as path defining element – out of scope for this session
- ECMP by default
- resiliency using TI-LFA

# Segment Routing - theory & usage



## SID (Segment Identifier)

- Prefix-SID
  - globally significant in network
  - attached to IGP prefix
  - prefix that identifies router (loopback) => Node-SID
- Adjacency-SID
  - locally significant
  - identifies IGP unidirectional adjacency
- Binding-SID
  - locally or globally significant
  - TE usage mainly
  - allows to swap one SID by list of SIDs

# Segment Routing - theory & usage



## Use cases:

**SR-MPLS**

**SR-TE**

**TI-LFA**



# Segment Routing - theory & usage




## SR-MPLS

- most simple use case
- MPLS labels are used for encoding SIDs in data plane
- IGP distributes SIDs
- Prefix-SID assigned manually to achieve unique ID of node
  - SRGB
    - typically configured as SRGB starting value + index (ex.16000 + 11), absolute value also possible
- Adjacency-SID
  - SRLB
  - usually dynamically allocated per IGP adjacency
- Forwarding to router based on IGP best path using Prefix(Node)-SID only in packet by default

# Segment Routing - theory & usage

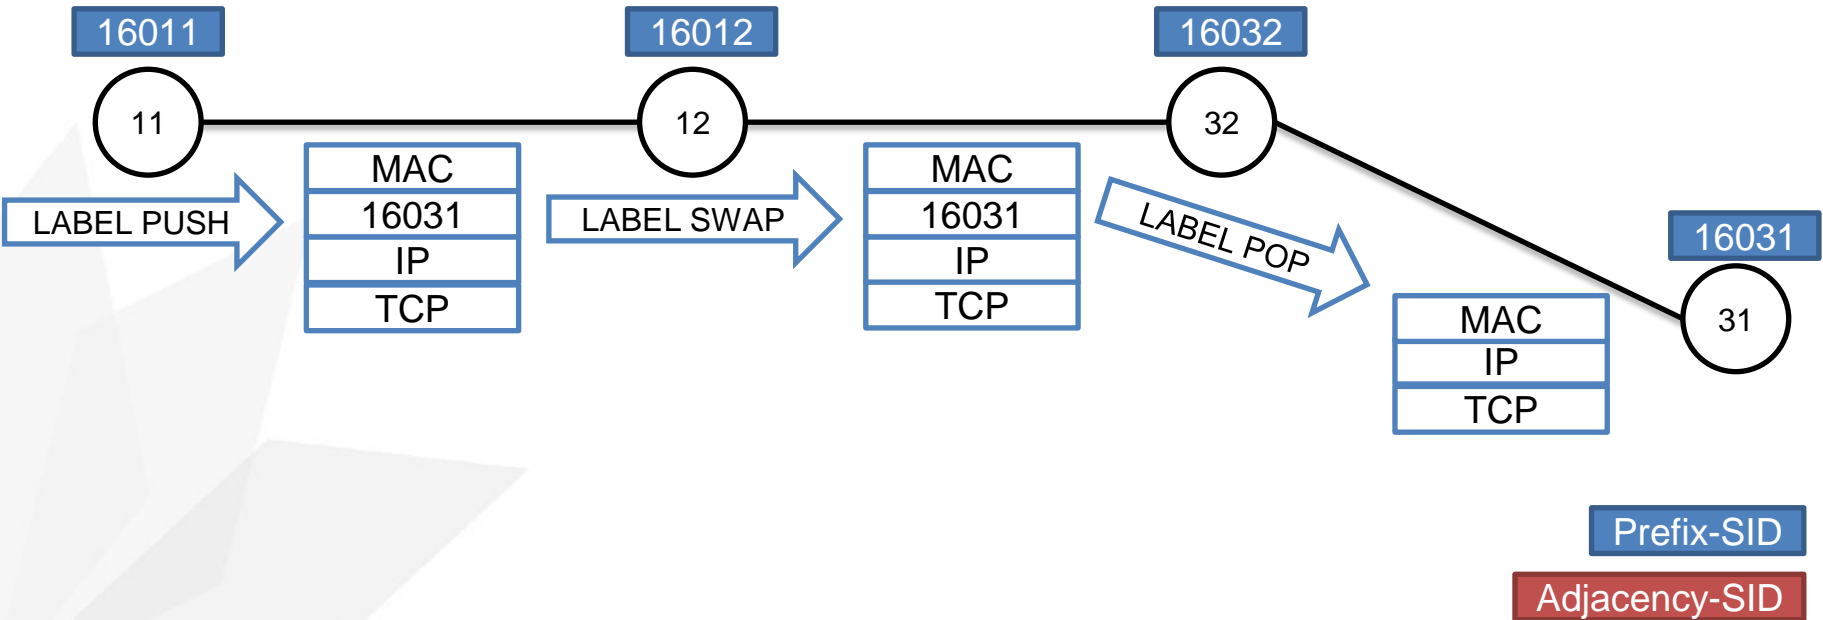


## SR-MPLS



# Segment Routing - theory & usage

## SR-MPLS





# Segment Routing - theory & usage



## SR-MPLS

```
PE-11#sh isis segment-routing global-block
```

```
Tag null:
```

```
IS-IS Level-2 Segment-routing Global Blocks:
```

System ID	<b>SRGB Base</b>	SRGB Range	<b>SID</b> Index	Strict-SID	SR-Cap	Strict
PE-21	<b>16000</b>	8000	<b>21</b>	None	Y	N
P-22	<b>16000</b>	8000	<b>22</b>	None	Y	N
PE-11	* <b>16000</b>	8000	<b>11</b>	None	Y	N
P-12	<b>16000</b>	8000	<b>12</b>	None	Y	N
PE-31	<b>16000</b>	1000	<b>31</b>	None	Y	N
P-32	<b>16000</b>	1000	<b>32</b>	None	Y	N

```
PE-11#
```

# Segment Routing - theory & usage



## SR-MPLS

```
PE-11#sh isis neighbors detail
```

```
Tag null:
```

System Id	Type	Interface	IP Address	State	Holdtime	Circuit Id
<b>PE-21</b>	<b>L2</b>	<b>Gi2</b>	10.11.21.2	UP	25	00

```
Area Address(es): 49.0421
```

```
SNPA: 5000.0003.0005
```

```
State Changed: 2d00h
```

```
Format: Phase V
```

```
Remote TID: 0
```

```
Local TID: 0
```

```
Interface name: GigabitEthernet2
```

```
Neighbor Circuit Id: 8
```

```
L(2) Adjacency SID Value: 18 f:0 b:0 v:1 l:1 s:0 p:0 weight:0
```

<b>P-12</b>	<b>L2</b>	<b>Gi1</b>	10.11.12.2	UP	26	01
-------------	-----------	------------	------------	----	----	----

```
Area Address(es): 49.0421
```

```
SNPA: 5000.0002.0000
```

```
State Changed: 2d00h
```

```
Format: Phase V
```

```
Remote TID: 0
```

```
Local TID: 0
```

```
Interface name: GigabitEthernet1
```

```
Neighbor Circuit Id: 1
```

```
L(2) Adjacency SID Value: 20 f:0 b:0 v:1 l:1 s:0 p:0 weight:0
```

```
PE-11#
```

# Segment Routing - theory & usage



## SR-MPLS

```
PE-11#sh isis topo
```

```
Tag null:
```

```
IS-IS TID 0 paths to level-2 routers
```

System Id	Metric	Next-Hop	Interface	SNPA
PE-21	10000	P-12	<b>Gi1</b>	5000.0002.0000
P-22	11000	P-12	<b>Gi1</b>	5000.0002.0000
PE-11	--			
P-12	5000	P-12	<b>Gi1</b>	5000.0002.0000
PE-31	15000	P-12	<b>Gi1</b>	5000.0002.0000
P-32	10000	P-12	<b>Gi1</b>	5000.0002.0000

```
PE-11#
```

# Segment Routing - theory & usage



## SR-MPLS

```
PE-11#sh mpls forwarding-table
```

Local Label	Outgoing Label	Prefix or Tunnel Id	Bytes Label Switched	Outgoing interface	Next Hop
16	No Label	1121/1[TE-Bind]	0	drop	
17	No Label	1131/1[TE-Bind]	0	drop	
18	Pop Label	10.11.21.2-A	0	Gi2	10.11.21.2
20	Pop Label	10.11.12.2-A	0	Gi1	10.11.12.2
<b>16012</b>	<b>Pop Label</b>	<b>192.168.12.1/32</b>	<b>0</b>	<b>Gi1</b>	<b>10.11.12.2</b>
<b>16021</b>	<b>16021</b>	<b>192.168.21.1/32</b>	<b>0</b>	<b>Gi1</b>	<b>10.11.12.2</b>
<b>16022</b>	<b>16022</b>	<b>192.168.22.1/32</b>	<b>0</b>	<b>Gi1</b>	<b>10.11.12.2</b>
<b>16031</b>	<b>16031</b>	<b>192.168.31.1/32</b>	<b>0</b>	<b>Gi1</b>	<b>10.11.12.2</b>
<b>16032</b>	<b>16032</b>	<b>192.168.32.1/32</b>	<b>0</b>	<b>Gi1</b>	<b>10.11.12.2</b>

```
A - Adjacency SID  
PE-11#
```

# Segment Routing - theory & usage



## SR-MPLS

```
PE-11#sh ip ro 192.168.31.1
```

```
Routing entry for 192.168.31.1/32
```

```
Known via "isis", distance 115, metric 15000, type level-2
```

```
Redistributing via isis
```

```
Last update from 10.11.12.2 on GigabitEthernet1, 00:03:30 ago
```

```
SR Incoming Label: 16031
```

```
Routing Descriptor Blocks:
```

```
* 10.11.12.2, from 192.168.31.1, 00:03:30 ago, via GigabitEthernet1, prefer-non-rib-labels, merge-labels
```

```
Route metric is 15000, traffic share count is 1
```

```
MPLS label: 16031
```

```
MPLS Flags: NSF
```

```
PE-11#sh ip cef 192.168.31.1 detail
```

```
192.168.31.1/32, epoch 2
```

```
sr local label info: global/16031 [0x1B]
```

```
1 RR source [no flags]
```

```
nexthop 10.11.12.2 GigabitEthernet1 label 16031-(local:16031)
```

```
PE-11#
```

```
PE-11#traceroute 192.168.31.1 source 192.168.11.1
```

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

```
Tracing the route to 192.168.31.1
```

```
VRF info: (vrf in name/id, vrf out name/id)
```

```
1 10.11.12.2 [MPLS: Label 16031 Exp 0] 19 msec 30 msec 15 msec
```

```
2 10.12.32.2 [MPLS: Label 16031 Exp 0] 522 msec 200 msec 325 msec
```

```
3 192.168.31.1 10 msec 16 msec 5 msec
```

```
PE-11#
```

# Segment Routing - theory & usage



## SR-MPLS - customer traffic over SR-MPLS ISP core

- CE behind PE-11, PE-21 and PE-31
- eBGP between CE and PE in GRT
- iBGP full mesh sessions between PEs (BGP RR would be used in normal setup)

```
PE-11#sh ip route bgp
```

```
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP  
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, m - OMP  
n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA  
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2  
ia - IS-IS inter area, * - candidate default, U - per-user static route  
H - NHRP, G - NHRP registered, g - NHRP registration summary  
o - ODR, P - periodic downloaded static route, l - LISP  
a - application route  
+ - replicated route, % - next hop override, p - overrides from PfR
```

```
Gateway of last resort is not set
```

```
192.13.1.0/32 is subnetted, 1 subnets  
B 192.13.1.1 [200/0] via 192.13.11.1, 2d16h  
192.13.2.0/32 is subnetted, 1 subnets  
B 192.13.2.2 [200/0] via 192.168.21.1, 16:16:05  
192.13.3.0/32 is subnetted, 1 subnets  
B 192.13.3.3 [200/0] via 192.168.31.1, 2d16h  
192.13.21.0/30 is subnetted, 1 subnets  
B 192.13.21.0 [200/0] via 192.168.21.1, 16:16:05  
PE-11#
```

# Segment Routing - theory & usage



## SR-MPLS - customer traffic over SR-MPLS ISP core

```
PE-11#sh ip cef 192.13.3.3 detail
192.13.3.3/32, epoch 2, flags [rib only nolabel, rib defined all labels]
  recursive via 192.168.31.1
    nexthop 10.11.12.2 GigabitEthernet1 label 16031-(local:16031)
PE-11#
```

```
### CE
CE-13#sh ip ro vrf vrf1 192.13.3.3
```

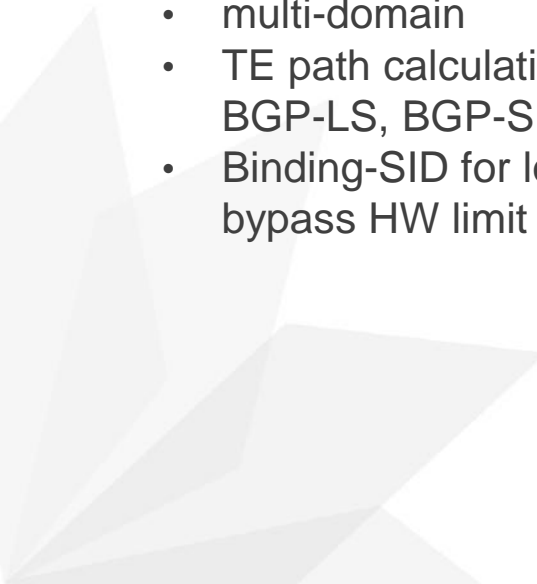
```
Routing Table: vrf1
Routing entry for 192.13.3.3/32
  Known via "bgp 65000", distance 20, metric 0
  Tag 421, type external
  Last update from 192.13.11.2 2d16h ago
  Routing Descriptor Blocks:
  * 192.13.11.2, from 192.13.11.2, 2d16h ago
    Route metric is 0, traffic share count is 1
    AS Hops 2
    Route tag 421
    MPLS label: none
```

```
CE-13#
CE-13#traceroute vrf vrf1 192.13.3.3 source 192.13.1.1
Type escape sequence to abort.
Tracing the route to 192.13.3.3
VRF info: (vrf in name/id, vrf out name/id)
 1 192.13.11.2 10 msec 48 msec 2 msec
 2 10.11.12.2 [MPLS: Label 16031 Exp 0] 25 msec 83 msec 14 msec
 3 * * *
 4 10.32.31.2 14 msec 22 msec 6 msec
 5 192.13.31.1 15 msec 18 msec *
CE-13#
```

# Segment Routing - theory & usage



## SR-TE (based on SR-MPLS)

- TE policy(path) = list of SIDs
  - no RSVP => no tunnel state along path => scaling
  - steering based on Prefix-SIDs and/or Adjacency-SIDs
  - using IGP extensions for TE metrics in SR-TE DB (metric, constraints, affinity)
  - multi-domain
  - TE path calculation on headend (dynamic or static) or centralized controller (PCE + BGP-LS, BGP-SRTE AFI)
  - Binding-SID for local identification of TE policy, stitching of SR paths across domains, bypass HW limit for label stack depth
- 



# Segment Routing - theory & usage



## SR-TE (based on SR-MPLS)

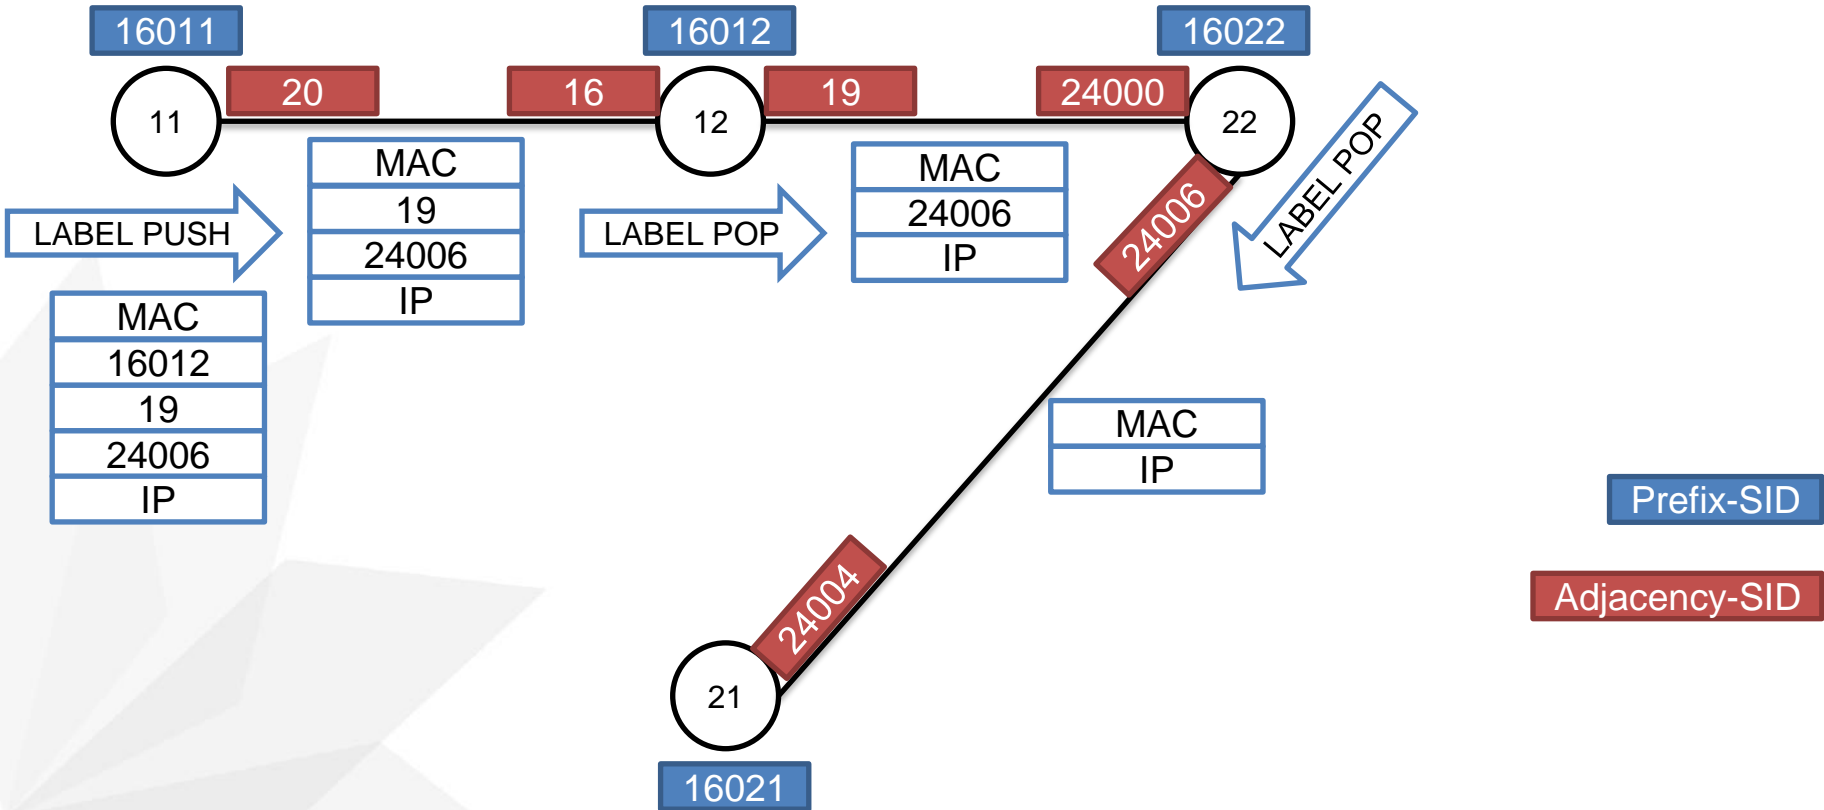
```
PE-11#sh ip cef 192.168.21.1 detail
192.168.21.1/32, epoch 2
  sr local label info: global/16021 [0x1B]
    1 RR source [no flags]
      nexthop 10.11.12.2 GigabitEthernet1 label 16021-(local:16021)
PE-11#traceroute 192.168.21.1
Type escape sequence to abort.
Tracing the route to 192.168.21.1
VRF info: (vrf in name/id, vrf out name/id)
  1 10.11.12.2 [MPLS: Label 16021 Exp 0] 38 msec 15 msec 2 msec
  2 10.12.21.2 94 msec 84 msec *
PE-11#

interface Tunnel1121
 ip unnumbered Loopback1
 shutdown
 tunnel mode mpls traffic-eng
 tunnel destination 192.168.21.1
 tunnel mpls traffic-eng autoroute announce
 tunnel mpls traffic-eng path-option 1 explicit name via_12_22 segment-routing
end

PE-11#sh ip explicit-paths name via_12_22
PATH via_12_22 (strict source route, path complete, generation 12)
  1: next-label 16012
  2: next-label 19
  3: next-label 24006
PE-11#
```

# Segment Routing - theory & usage

## SR-TE (based on SR-MPLS)



# Segment Routing - theory & usage



## SR-TE (based on SR-MPLS)

```
PE-11#sh ip cef 192.168.21.1 detail
192.168.21.1/32, epoch 2
  1 RR source [no flags]
    nexthop 192.168.21.1 Tunnel1121
PE-11#traceroute 192.168.21.1
Type escape sequence to abort.
Tracing the route to 192.168.21.1
VRF info: (vrf in name/id, vrf out name/id)
 1 10.11.12.2 [MPLS: Labels 19/24006 Exp 0] 25 msec 32 msec 14 msec
 2 10.12.22.2 [MPLS: Label 24006 Exp 0] 31 msec 22 msec 4 msec
 3 10.21.22.1 24 msec 29 msec *
```

```
PE-11#sh mpls traffic-eng tunnels Tunnel1121 detail | i Segment|Binding
Binding SID: 16
Segment-Routing Path Info (isis level-2)
Segment0[Node]: 192.168.12.1, Label: 16012
Segment1[Link]: 10.12.22.1 - 10.12.22.2, Label: 19
Segment2[ - ]: Label: 24006
```

```
PE-11#show mpls forwarding-table interface Tunnel1121 detail
Local      Outgoing  Prefix          Bytes Label  Outgoing  Next Hop
Label      Label     or Tunnel Id   Switched     interface
16         Pop Label  1121/1[TE-Bind] 0             Tu1121     point2point
MAC/Encaps=14/22, MRU=1496, Label Stack{19 24006}, via Gi1
5000000200005000000100008847 0001300005DC6000
No output feature configured
```

```
PE-11#
PE-11#sh mpls forwarding-table labels 16012
Local      Outgoing  Prefix          Bytes Label  Outgoing  Next Hop
Label      Label     or Tunnel Id   Switched     interface
16012     Pop Label  192.168.12.1/32 0             Gi1        10.11.12.2
```

```
PE-11#
```

# Segment Routing - theory & usage



## SR-TE (based on SR-MPLS)

```
P-12#sh mpls forwarding-table labels 19
```

Local Label	Outgoing Label	Prefix or Tunnel Id	Bytes Switched	Label	Outgoing interface	Next Hop
19	Pop Label	10.12.22.2-A	3226		Gi3	10.12.22.2

```
A - Adjacency SID
```

```
P-12#
```

```
RP/0/RP0/CPU0:P-22#sh mpls forwarding labels 24006
```

```
Sat Sep 5 07:47:22.156 UTC
```

Local Label	Outgoing Label	Prefix or ID	Outgoing Interface	Next Hop	Bytes Switched
24006	Pop	SR Adj (idx 1)	Gi0/0/0/0	10.21.22.1	26356

```
RP/0/RP0/CPU0:P-22#
```

# Segment Routing - theory & usage



## SR-TE (based on SR-MPLS)

- BGP-LS to transport IGP info to remote controller (mostly PCE)

```
CE-13#sh ip bgp link-state link-state nLRI [T4][L2][I0x0][N[c421][b0.0.0.0][s0000.0000.0021.00]][P[p192.168.21.1/32]]
BGP routing table entry for [T4][L2][I0x0][N[c421][b0.0.0.0][s0000.0000.0021.00]][P[p192.168.21.1/32]], version 6
Paths: (1 available, best #1, table link-state link-state)
  Not advertised to any peer
  Refresh Epoch 1
  Local
  10.12.13.1 from 10.12.13.1 (192.168.12.1)
    Origin IGP, metric 0, localpref 100, valid, internal, best
    LS Attribute: Metric: 0, PFX-SID: 21(40/0), Extended IGP flags: 0x20,
      Source Router ID: 192.168.21.1
    rx pathid: 0, tx pathid: 0x0
    Updated on Sep 2 2020 14:11:45 UTC
CE-13#
```

# Segment Routing - theory & usage



## SR-TE (based on SR-MPLS)

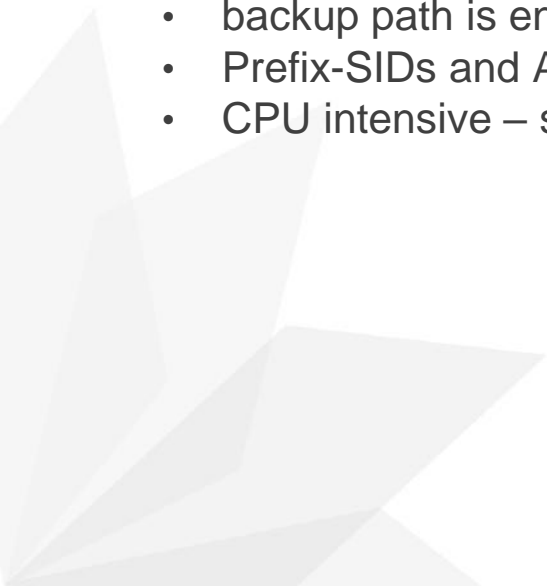
```
CE-13#sh ip bgp link-state link-state n1RI
[E][L2][I0x0][N[c421][b0.0.0.0][s0000.0000.0102.00]][R[c421][b0.0.0.0][s0000.0000.0022.00]][L[i10.12.22.1][n10.12.22.2][t0x0000]]
BGP routing table entry for
[E][L2][I0x0][N[c421][b0.0.0.0][s0000.0000.0102.00]][R[c421][b0.0.0.0][s0000.0000.0022.00]][L[i10.12.22.1][n10.12.22.2][t0x0000]], version 28
Paths: (1 available, best #1, table link-state link-state)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    10.12.13.1 from 10.12.13.1 (192.168.12.1)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      LS Attribute: Local TE Router-ID: 192.168.12.1,
        Remote TE Router-ID: 192.168.22.1, admin-group: 0x00000000,
        max-link-bw (kbits/sec): 1000000,
        max-reserv-link-bw (kbits/sec): 0,
        max-unreserv-link-bw (kbits/sec): 0 0 0 0 0 0 0,
        TE-default-metric: 6000, metric: 6000, ADJ-SID: 19(30)
      rx pathid: 0, tx pathid: 0x0
      Updated on Sep 2 2020 14:11:45 UTC
CE-13#
```

# Segment Routing - theory & usage



## TI-LFA (based on SR-MPLS)

- improvement of LFA and rLFA
- alternative to MPLS-TE FRR
- 100% link and node protection – details out of scope for this session
- backup paths computed automatically by IGP
- backup path is encoded into packet header in case of link/node failure and is loop-free
- Prefix-SIDs and Adjacency-SIDs used to steer traffic
- CPU intensive – some optimization exists if too expensive



# Segment Routing - theory & usage



## TI-LFA (based on SR-MPLS)

```
PE-11#sh ip route 192.168.31.1
Routing entry for 192.168.31.1/32
  Known via "isis", distance 115, metric 15000, type level-2
  Redistributing via isis
  Last update from 10.11.12.2 on GigabitEthernet1, 00:05:34 ago
  SR Incoming Label: 16031
  Routing Descriptor Blocks:
  * 10.11.12.2, from 192.168.31.1, 00:05:34 ago, via GigabitEthernet1, prefer-non-rib-labels, merge-labels
    Route metric is 15000, traffic share count is 1
    MPLS label: 16031
    MPLS Flags: NSF
    Repair Path: 10.11.21.2, via GigabitEthernet2
```

PE-11#

```
PE-11#sh isis fast-reroute summary
```

Tag null:

Microloop Avoidance Enable State: enabled for segment-routing microloop avoidance

IPv4 Fast-Reroute Protection Summary:

Prefix Counts:	Total	Protected	Coverage
High priority:	0	0	0%
Normal priority:	16	13	81%
Total:	16	13	81%

PE-11#



# Segment Routing - theory & usage



## TI-LFA (based on SR-MPLS)

```
PE-11#sh ip cef 192.168.31.1/32 detail
192.168.31.1/32, epoch 2
  sr local label info: global/16031 [0x1B]
  1 RR source [no flags]
  nexthop 10.11.12.2 GigabitEthernet1 label [16031|16031]-(local:16031)
  repair: attached-nexthop 10.11.21.2 GigabitEthernet2
  nexthop 10.11.21.2 GigabitEthernet2 label 16031-(local:16031), repair
PE-11#
```

```
PE-11#traceroute 192.168.31.1 source 192.168.11.1
Type escape sequence to abort.
Tracing the route to 192.168.31.1
VRF info: (vrf in name/id, vrf out name/id)
 1 10.11.12.2 [MPLS: Label 16031 Exp 0] 19 msec 10 msec 5 msec
 2 10.12.32.2 [MPLS: Label 16031 Exp 0] 127 msec 212 msec 418 msec
 3 192.168.31.1 236 msec 32 msec 436 msec
```

```
PE-11#
```

FAILURE of 10.12.32.2

```
PE-11#sh ip cef 192.168.31.1/32 detail
192.168.31.1/32, epoch 2
  sr local label info: global/16031 [0x1B]
  1 RR source [no flags]
  nexthop 10.11.12.2 GigabitEthernet1 label [16031|16031]-(local:16031)
  repair: attached-nexthop 10.11.21.2 GigabitEthernet2
  nexthop 10.11.21.2 GigabitEthernet2 label 16031-(local:16031), repair
PE-11#
```

```
PE-11#traceroute 192.168.31.1 source 192.168.11.1
Type escape sequence to abort.
Tracing the route to 192.168.31.1
VRF info: (vrf in name/id, vrf out name/id)
 1 10.11.12.2 [MPLS: Label 16031 Exp 0] 45 msec 13 msec 17 msec
 2 10.12.22.2 [MPLS: Label 16031 Exp 0] 38 msec 16 msec 4 msec
 3 192.168.31.1 31 msec 21 msec 3 msec
```

```
PE-11#
```

# Segment Routing - theory & usage



## Resources

<https://www.segment-routing.net/tutorials/2016-09-27-segment-routing-introduction/>

<https://www.segment-routing.net/tutorials/2016-09-27-segment-routing-mpls-data-plane/>

<https://www.segment-routing.net/tutorials/2016-09-27-segment-routing-global-block-srgb/>

<https://www.segment-routing.net/tutorials/2016-09-27-segment-routing-igp-control-plane/>

<https://www.segment-routing.net/tutorials/2017-03-06-segment-routing-traffic-engineering-srte/>

<https://www.segment-routing.net/tutorials/2016-09-27-topology-independent-lfa-ti-lfa/>

<https://www.sanog.org/resources/sanog34/SANOG34-Tutorials-spring-for-service-providers-v1.3.pdf>



# Segment Routing - theory & usage



Thank you



# Segment Routing - theory & usage



**Extra slides**



# Segment Routing - theory & usage



## SR-MPLS on IOS-XR

```
RP/0/RP0/CPU0:P-22#sh isis segment-routing label table
Sat Sep  5 06:11:44.064 UTC
```

```
IS-IS 421 IS Label Table
Label          Prefix/Interface
-----
16011          192.168.11.1/32
16012          192.168.12.1/32
16021          192.168.21.1/32
16022          Loopback1
16031          192.168.31.1/32
16032          192.168.32.1/32
RP/0/RP0/CPU0:P-22#
```

# Segment Routing - theory & usage



## SR-MPLS on IOS-XR

```
RP/0/RP0/CPU0:P-22#sh isis adjacency detail
Sat Sep  5 06:12:46.961 UTC
```

```
IS-IS 421 Level-2 adjacencies:
```

System Id	Interface	SNPA	State	Hold	Changed	NSF	IPv4 BFD	IPv6 BFD
P-12	Gi0/0/0/1	*PtoP*	Up	23	2d16h	Yes	None	None
Area Address: 49.0421								
Neighbor IPv4 Address: 10.12.22.1*								
Adjacency SID: 24000								
Non-FRR Adjacency SID: 24001								
Topology: IPv4 Unicast								
P-32	Gi0/0/0/2	*PtoP*	Up	23	3d22h	Yes	None	None
Area Address: 49.0421								
Neighbor IPv4 Address: 10.22.32.1*								
Adjacency SID: 24002								
Non-FRR Adjacency SID: 24003								
Topology: IPv4 Unicast								
PE-31	Gi0/0/0/3	*PtoP*	Up	24	3d10h	Yes	None	None
Area Address: 49.0421								
Neighbor IPv4 Address: 10.22.31.1*								
Adjacency SID: 24004								
Non-FRR Adjacency SID: 24005								
Topology: IPv4 Unicast								

# Segment Routing - theory & usage



## SR-MPLS on IOS-XR

```
PE-21          Gi0/0/0/0          *PtoP*          Up    27    3d10h    Yes None None
  Area Address:          49.0421
  Neighbor IPv4 Address: 10.21.22.1*
  Adjacency SID:        24006
  Non-FRR Adjacency SID: 24007
  Topology:              IPv4 Unicast
```

```
Total adjacency count: 4
RP/0/RP0/CPU0:P-22#
```

```
RP/0/RP0/CPU0:P-22#sh isis topo
Sat Sep  5 06:15:47.501 UTC
```

```
IS-IS 421 paths to IPv4 Unicast (Level-2) routers
System Id      Metric  Next-Hop      Interface      SNPA
PE-21          5000    PE-21         Gi0/0/0/0     *PtoP*
P-22           --
PE-11          11000   P-12          Gi0/0/0/1     *PtoP*
P-12           6000    P-12          Gi0/0/0/1     *PtoP*
PE-31          5000    PE-31         Gi0/0/0/3     *PtoP*
P-32           5000    P-32          Gi0/0/0/2     *PtoP*
RP/0/RP0/CPU0:P-22#
```

# Segment Routing - theory & usage



## SR-MPLS on IOS-XR

```
RP/0/RP0/CPU0:P-22#sh mpls forwarding
```

```
Sat Sep 5 06:16:35.741 UTC
```

Local Label	Outgoing Label	Prefix or ID	Outgoing Interface	Next Hop	Bytes Switched
16011	16011	SR Pfx (idx 11)	Gi0/0/0/1	10.12.22.1	0
16012	Pop	SR Pfx (idx 12)	Gi0/0/0/1	10.12.22.1	0
16021	Pop	SR Pfx (idx 21)	Gi0/0/0/0	10.21.22.1	832975
16031	Pop	SR Pfx (idx 31)	Gi0/0/0/3	10.22.31.1	815251
16032	Pop	SR Pfx (idx 32)	Gi0/0/0/2	10.22.32.1	750813
24000	Pop	SR Adj (idx 1)	Gi0/0/0/1	10.12.22.1	0
24001	Pop	SR Adj (idx 3)	Gi0/0/0/1	10.12.22.1	0
24002	Pop	SR Adj (idx 1)	Gi0/0/0/2	10.22.32.1	0
24003	Pop	SR Adj (idx 3)	Gi0/0/0/2	10.22.32.1	0
24004	Pop	SR Adj (idx 1)	Gi0/0/0/3	10.22.31.1	0
24005	Pop	SR Adj (idx 3)	Gi0/0/0/3	10.22.31.1	0
24006	Pop	SR Adj (idx 1)	Gi0/0/0/0	10.21.22.1	22935
24007	Pop	SR Adj (idx 3)	Gi0/0/0/0	10.21.22.1	0

```
RP/0/RP0/CPU0:P-22#
```



# Segment Routing - theory & usage



## SR-MPLS on IOS-XR

```
RP/0/RP0/CPU0:P-22#sh ip cef 192.168.32.1
Sat Sep  5 06:17:41.253 UTC
192.168.32.1/32, version 386, labeled SR, internal 0x1000001 0x81 (ptr 0xe0298b0) [1], 0x0 (0xe1ec268), 0xa20
(0xe5f02e8)
Updated Sep  1 19:26:38.503
remote adjacency to GigabitEthernet0/0/0/2
Prefix Len 32, traffic index 0, precedence n/a, priority 1
  via 10.22.32.1/32, GigabitEthernet0/0/0/2, 10 dependencies, weight 0, class 0 [flags 0x0]
    path-idx 0 NHID 0x0 [0xf1041d0 0xf104fe0]
    next hop 10.22.32.1/32
    remote adjacency
      local label 16032      labels imposed {ImplNull}
RP/0/RP0/CPU0:P-22#sh ip cef 192.168.11.1
Sat Sep  5 06:17:59.666 UTC
192.168.11.1/32, version 513, labeled SR, internal 0x1000001 0x81 (ptr 0xe029710) [1], 0x0 (0xe1ebe68), 0xa28
(0xe5f03a8)
Updated Sep  2 14:01:53.230
remote adjacency to GigabitEthernet0/0/0/1
Prefix Len 32, traffic index 0, precedence n/a, priority 1
  via 10.12.22.1/32, GigabitEthernet0/0/0/1, 6 dependencies, weight 0, class 0 [flags 0x0]
    path-idx 0 NHID 0x0 [0xf104ec0 0x0]
    next hop 10.12.22.1/32
    remote adjacency
      local label 16011      labels imposed {16011}
RP/0/RP0/CPU0:P-22#
```

# Segment Routing - theory & usage



## SR-MPLS on JUNOS

```
demo@PE-31> show isis database detail | match "(Sequence|Index)"
PE-21.00-00 Sequence: 0x459, Checksum: 0x4aaf, Lifetime: 930 secs
  IPV4 Index: 21
    Start Index : 0, Size : 8000, Label-Range: [ 16000, 23999 ]
P-22.00-00 Sequence: 0x48e, Checksum: 0xb6a, Lifetime: 724 secs
  IPV4 Index: 22
    Start Index : 0, Size : 8000, Label-Range: [ 16000, 23999 ]
PE-11.00-00 Sequence: 0x47f, Checksum: 0xf204, Lifetime: 426 secs
  IPV4 Index: 11
    Start Index : 0, Size : 8000, Label-Range: [ 16000, 23999 ]
P-12.00-00 Sequence: 0x489, Checksum: 0x868e, Lifetime: 821 secs
  IPV4 Index: 12
    Start Index : 0, Size : 8000, Label-Range: [ 16000, 23999 ]
PE-31.00-00 Sequence: 0x48e, Checksum: 0x22ed, Lifetime: 686 secs
  IPV4 Index: 31
    Start Index : 0, Size : 1000, Label-Range: [ 16000, 16999 ]
P-32.00-00 Sequence: 0x495, Checksum: 0xdb44, Lifetime: 1139 secs
  IPV4 Index: 32
    Start Index : 0, Size : 1000, Label-Range: [ 16000, 16999 ]

demo@PE-31>
```

# Segment Routing - theory & usage



## SR-MPLS on JUNOS

```
demo@PE-31> show isis adjacency detail
```

```
P-32
```

```
Interface: ge-0/0/0.0, Level: 2, State: Up, Expires in 25 secs  
Priority: 0, Up/Down transitions: 1, Last transition: 3d 10:33:22 ago  
Circuit type: 2, Speaks: IP, IPv6  
Topologies: Unicast  
Restart capable: Yes, Adjacency advertisement: Advertise  
IP addresses: 10.32.31.1  
Level 2 IPv4 Adj-SID: 16
```

```
P-22
```

```
Interface: ge-0/0/1.0, Level: 2, State: Up, Expires in 20 secs  
Priority: 0, Up/Down transitions: 1, Last transition: 3d 10:32:51 ago  
Circuit type: 2, Speaks: IP  
Topologies: Unicast  
Restart capable: Yes, Adjacency advertisement: Advertise  
IP addresses: 10.22.31.2  
Level 2 IPv4 Adj-SID: 17
```

```
demo@PE-31>
```

# Segment Routing - theory & usage



## SR-MPLS on JUNOS

```
demo@PE-31> show isis spf brief
IS-IS level 1 SPF results:
 0 nodes
```

```
IS-IS level 2 SPF results:
```

Node	Metric	Interface	NH	Via	SNPA
PE-11.00	15000	ge-0/0/0.0	IPV4	P-32	50:0:0:8:0:4
P-12.00	10000	ge-0/0/0.0	IPV4	P-32	50:0:0:8:0:4
PE-21.00	10000	ge-0/0/1.0	IPV4	P-22	50:0:0:4:0:6
P-32.00	5000	ge-0/0/0.0	IPV4	P-32	50:0:0:8:0:4
P-22.00	5000	ge-0/0/1.0	IPV4	P-22	50:0:0:4:0:6
PE-31.00	0				

```
6 nodes
```

```
demo@PE-31>
```

# Segment Routing - theory & usage



## SR-MPLS on JUNOS

Routing table: default.mpls

MPLS:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		dscd	50	1	
0	user	0		rtbl	1	3	
0(S=0)	user	0		rtbl	3	4	
1	user	0		recv	49	2	
2	user	0		rtbl	2	3	
2(S=0)	user	0		rtbl	3	4	
13	user	0		recv	49	2	
16	user	0	10.32.31.1	Pop	588	3	ge-0/0/0.0
16(S=0)	user	0	10.32.31.1	Pop	589	3	ge-0/0/0.0
17	user	0	10.22.31.2	Pop	590	3	ge-0/0/1.0
17(S=0)	user	0	10.22.31.2	Pop	594	3	ge-0/0/1.0
16011	user	0	10.32.31.1	Swap 16011		612	2 ge-0/0/0.0
16012	user	0	10.32.31.1	Swap 16012		613	2 ge-0/0/0.0
16021	user	0	10.22.31.2	Swap 16021		595	2 ge-0/0/1.0
16022	user	0	10.22.31.2	Pop	590	3	ge-0/0/1.0
16022(S=0)	user	0	10.22.31.2	Pop	594	3	ge-0/0/1.0
16032	user	0	10.32.31.1	Pop	588	3	ge-0/0/0.0
16032(S=0)	user	0	10.32.31.1	Pop	589	3	ge-0/0/0.0

# Segment Routing - theory & usage



## SR-MPLS on JUNOS

```
demo@PE-31> show route forwarding-table destination 192.168.11.1
```

```
Routing table: default.inet
```

```
Internet:
```

```
Enabled protocols: Bridging,
```

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
192.168.11.1/32	user	1	10.32.31.1	Push	16011	608	3 ge-0/0/0.0

```
demo@PE-31> show route forwarding-table destination 192.168.22.1
```

```
Routing table: default.inet
```

```
Internet:
```

```
Enabled protocols: Bridging,
```

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
192.168.22.1/32	user	0	10.22.31.2	ucst	599	2	ge-0/0/1.0

# Segment Routing - theory & usage



## SR-TE (based on SR-MPLS) on IOS-XR

```
RP/0/RP0/CPU0:PE-21#sh running-config segment-routing
Sat Sep  5 12:17:47.096 UTC
segment-routing
traffic-eng
 interface GigabitEthernet0/0/0/0
  affinity
   name red
  !
  metric 10
 !
 interface GigabitEthernet0/0/0/1
  affinity
   name green
  !
  metric 10
 !
 interface GigabitEthernet0/0/0/2
  affinity
   name black
  !
```

```
metric 100
!
segment-list via_22_32_12
 index 1 mpls label 16022
 index 2 mpls label 16032
 index 3 mpls label 16012
 index 4 mpls label 16011
!
policy to_192.168.11.1
 color 10 end-point ipv4 192.168.11.1
 autoroute
  metric constant 5
  include ipv4 192.168.11.1/32
!
candidate-paths
 preference 10
  explicit segment-list via_22_32_12
```

# Segment Routing - theory & usage



## SR-TE (based on SR-MPLS) on IOS-XR

```
RP/0/RP0/CPU0:PE-21#show segment-routing traffic-eng policy
SR-TE policy database
```

```
-----
Color: 10, End-point: 192.168.11.1
Name: srte_c_10_ep_192.168.11.1
Status:
  Admin: up Operational: up for 2d22h (since Sep  2 13:20:58.230)
Candidate-paths:
  Preference: 10 (configuration) (active)
  Name: to_192.168.11.1
  Requested BSID: dynamic
  Explicit: segment-list via_22_32_12 (valid)
  Weight: 1, Metric Type: TE
    16022
    16032
    16012
    16011
Attributes:
  Binding SID: 24007
  Forward Class: Not Configured
  Steering BGP disabled: no
  IPv6 caps enable: yes
```

```
RP/0/RP0/CPU0:PE-21#
```



# Segment Routing - theory & usage



## SR-TE (based on SR-MPLS) on IOS-XR

```
RP/0/RP0/CPU0:PE-21#sh route 192.168.11.1  
Sat Sep  5 12:22:44.258 UTC
```

```
Routing entry for 192.168.11.1/32
```

```
Known via "isis 421", distance 115, metric 10001, labeled SR, label redistrib non FIB, type level-2
```

```
Installed Sep  2 14:01:51.446 for 2d22h
```

```
Routing Descriptor Blocks
```

```
192.168.11.1, from 192.168.11.1, via srte_c_10_ep_192.168.11.1
```

```
Route metric is 10001
```

```
No advertising protos.
```

```
RP/0/RP0/CPU0:PE-21#
```

```
RP/0/RP0/CPU0:PE-21#traceroute 192.168.11.1 source 192.168.21.1
```

```
Sat Sep  5 12:23:12.994 UTC
```

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

```
Tracing the route to 192.168.11.1
```

```
 1 10.21.22.2 [MPLS: Labels 16032/16012/16011 Exp 0] 129 msec  69 msec 104 msec  
 2 10.22.32.1 [MPLS: Labels 16012/16011 Exp 0] 350 msec 300 msec 370 msec  
 3 10.12.32.1 [MPLS: Label 16011 Exp 0] 109 msec  62 msec  75 msec  
 4 10.11.12.1 89 msec * 64 msec  
RP/0/RP0/CPU0:PE-21#
```

# Segment Routing - theory & usage



## SR-TE (based on SR-MPLS) on JUNOS

```
demo@PE-31> show configuration protocols source-packet-
routing
preference 1;
segment-list via_32_12_22 {
  hop1 ip-address 10.32.31.1;
  hop2 label 16012;
  hop3 label 16022;
  hop4 label 16021;
}
segment-list via_22_21 {
  hop1 ip-address 10.22.31.2;
  hop2 label 16021;
  hop3 label 16011;
}
source-routing-path to_192.168.21.1 {
  to 192.168.21.1;
  binding-sid 1000999;
  metric 10;
```

```
  primary {
    via_32_12_22 weight 5;
  }
}
source-routing-path to_192.168.11.1 {
  to 192.168.11.1;
  binding-sid 1000990;
  metric 10;
  primary {
    via_22_21 weight 5;
  }
}
demo@PE-31>
```

# Segment Routing - theory & usage



## SR-TE (based on SR-MPLS) on JUNOS

```
demo@PE-31> show spring-traffic-engineering lsp detail
```

```
Name: to_192.168.11.1
```

```
To: 192.168.11.1
```

```
State: Up, Outgoing interface: ge-0/0/1.0
```

```
SR-ERO hop count: 3
```

```
Hop 1 (Strict):
```

```
NAI: IPv4 Adjacency ID, 0.0.0.0 -> 10.22.31.2
```

```
SID type: None
```

```
Hop 2 (Strict):
```

```
NAI: None
```

```
SID type: 20-bit label, Value: 16021
```

```
Hop 3 (Strict):
```

```
NAI: None
```

```
SID type: 20-bit label, Value: 16011
```

```
Name: to_192.168.21.1
```

```
To: 192.168.21.1
```

```
State: Up, Outgoing interface: ge-0/0/0.0
```

```
SR-ERO hop count: 4
```

```
Hop 1 (Strict):
```

```
NAI: IPv4 Adjacency ID, 0.0.0.0 -> 10.32.31.1
```

```
SID type: None
```

```
Hop 2 (Strict):
```

```
NAI: None
```

```
SID type: 20-bit label, Value: 16012
```

```
Hop 3 (Strict):
```

```
NAI: None
```

```
SID type: 20-bit label, Value: 16022
```

```
Hop 4 (Strict):
```

```
NAI: None
```

```
SID type: 20-bit label, Value: 16021
```

```
Total displayed LSPs: 2 (Up: 2, Down: 0)
```

```
demo@PE-31>
```

# Segment Routing - theory & usage



## SR-TE (based on SR-MPLS) on JUNOS

```
demo@PE-31> show route 192.168.11.1
```

```
inet.0: 27 destinations, 35 routes (27 active, 0 holddown, 0 hidden)  
+ = Active Route, - = Last Active, * = Both
```

```
192.168.11.1/32    *[L-ISIS/14] 2d 22:24:53, metric 15001  
                  > to 10.32.31.1 via ge-0/0/0.0, Push 16011  
                  [IS-IS/18] 2d 22:24:53, metric 15001  
                  > to 10.32.31.1 via ge-0/0/0.0
```

```
inet.3: 5 destinations, 7 routes (5 active, 0 holddown, 0 hidden)  
+ = Active Route, - = Last Active, * = Both
```

```
192.168.11.1/32    *[SPRING-TE/1] 3d 15:16:39, metric 10  
                  > to 10.22.31.2 via ge-0/0/1.0, Push 16011, Push 16021(top)  
                  [L-ISIS/14] 2d 22:24:53, metric 15001  
                  > to 10.32.31.1 via ge-0/0/0.0, Push 16011
```

```
demo@PE-31>
```

# Segment Routing - theory & usage



## SR-TE (based on SR-MPLS) on JUNOS

```
demo@PE-31> show route 192.168.21.1
```

```
inet.0: 27 destinations, 35 routes (27 active, 0 holddown, 0 hidden)  
+ = Active Route, - = Last Active, * = Both
```

```
192.168.21.1/32    *[L-ISIS/14] 3d 16:33:13, metric 10000  
                  > to 10.22.31.2 via ge-0/0/1.0, Push 16021  
                  [IS-IS/18] 3d 16:33:14, metric 10000  
                  > to 10.22.31.2 via ge-0/0/1.0
```

```
inet.3: 5 destinations, 7 routes (5 active, 0 holddown, 0 hidden)  
+ = Active Route, - = Last Active, * = Both
```

```
192.168.21.1/32    *[SPRING-TE/1] 3d 16:41:22, metric 10  
                  > to 10.32.31.1 via ge-0/0/0.0, Push 16021, Push 16022, Push 16012(top)  
                  [L-ISIS/14] 3d 16:33:13, metric 10000  
                  > to 10.22.31.2 via ge-0/0/1.0, Push 16021
```

```
demo@PE-31>
```

# Segment Routing - theory & usage



## SR-TE (based on SR-MPLS) on JUNOS

```
CE-13#traceroute vrf vrf11 192.13.2.2 source 192.13.3.3
Type escape sequence to abort.
Tracing the route to 192.13.2.2
VRF info: (vrf in name/id, vrf out name/id)
 1 192.13.31.2 8 msec 9 msec 2 msec
 2 * * *
 3 10.12.32.1 [MPLS: Labels 16022/16021 Exp 0] 25 msec 14 msec 5 msec
 4 10.12.22.2 [MPLS: Label 16021 Exp 0] 28 msec 22 msec 7 msec
 5 10.21.22.1 28 msec 15 msec 11 msec
 6 192.13.21.1 10 msec 16 msec *
CE-13#
```

# Segment Routing - theory & usage



## TI-LFA (based on SR-MPLS) on IOS-XR

```
RP/0/RP0/CPU0:PE-21#sh isis fast-reroute 192.168.31.1/32
Sat Sep  5 13:04:35.168 UTC
```

```
L2 192.168.31.1/32 [10000/115]
  via 10.21.22.2, GigabitEthernet0/0/0/0, P-22, SRGB Base: 16000, Weight: 0
  Backup path: LFA, via 10.12.21.1, GigabitEthernet0/0/0/1, P-12, SRGB Base: 16000, Weight: 0, Metric: 15000
RP/0/RP0/CPU0:PE-21#
```

```
RP/0/RP0/CPU0:PE-21#sh ip cef 192.168.31.1
Sat Sep  5 13:05:21.655 UTC
192.168.31.1/32, version 387, labeled SR, internal 0x1000001 0x81 (ptr 0xe0293a0) [3], 0x0 (0xe1ebba8), 0xa28
(0xf65d1c8)
Updated Sep  5 13:03:43.816
remote adjacency to GigabitEthernet0/0/0/0
Prefix Len 32, traffic index 0, precedence n/a, priority 1
  via 10.21.22.2/32, GigabitEthernet0/0/0/0, 10 dependencies, weight 0, class 0, protected [flags 0x400]
    path-idx 0 bkup-idx 1 NHID 0x0 [0xf56d190 0x0]
    next hop 10.21.22.2/32
      local label 16031      labels imposed {16031}
  via 10.12.21.1/32, GigabitEthernet0/0/0/1, 12 dependencies, weight 0, class 0, backup (Local-LFA) [flags 0x300]
    path-idx 1 NHID 0x0 [0xf104410 0x0]
    next hop 10.12.21.1/32
      remote adjacency
        local label 16031      labels imposed {16031}
RP/0/RP0/CPU0:PE-21#
```