

Segment Routing: Enterprise Use Cases

11 April 2024 Alessandra Falasca

Agenda

 Segment Routing Architecture
 Troubleshooting Simplification
 Traffic Protection with TI-LFA
 Low Latency SR-TE Policy



Segment Routing Architecture

Agenda

- Why Segment Routing
- SR Data Planes

From Unified MPLS to Segment Routing Protocol stack and network **simplification**



5 5

Unified "Stateless Fabric" for Service Creation





Unified underlay and overlay networks with segment routing and EVPN



E2E Cross-domain automation with model-driven programmability and streaming telemetry



Transform the CO into a data center to enable distributed service delivery and speed up service creation

6

7

Segment Routing

- Source Routing
 - the topological and service (NFV) path is encoded in packet header
- Scalability
 - the network fabric does not hold any per-flow state for TE or NFV
- Simplicity
 - automation: TILFA sub-50msec FRR
 - protocol elimination: LDP, RSVP-TE, VxLAN, NSH, GTP, ...
- End-to-End
 - DC, Metro, WAN

Simple extension to IGP



Simple extension to IS-IS or OSPF, automatically builds and maintains Segments
 Prefix Segment – A Shortest path to the related prefix
 Adjacency Segment – One hop through the related adjacency

Two dataplane instantiations



Network Evolution with SR-MPLS

Service Protocols





Segment Routing



- Source Routing paradigm
 - Stateless IP fabric !!!



Troubleshooting Simplification



IGP Prefix Segment

Shortest-path to the IGP prefix



IGP Prefix Segment

Shortest-path to the IGP prefix Global Segment



Global Segment

RP/0/R	P0/CPU0	: <mark>Node-3</mark> #sh	mpls forward	ling prefix 1.1	.1.6/32
Tue Ja	n 29 10:30):53.133 UT(C		
Local (Label	Outgoing Label	Prefix or ID	Outgoing Interface	Next Hop	Bytes Switcheo
16006	Рор	1.1.1.6/32	Te0/0/0/1	77.4.6.4	0

Traffic Protection with TI-LFA

Protect with automatic TI LFA FRR



Topology Independent LFA (TI-LFA) – Benefits

100%-coverage 50-msec link, node, and SRLG protection **Simple** to operate and understand

!IOS-XR

```
Router ospf 1
segment-routing mpls
fast-reroute per-prefix
fast-reroute per-prefix ti-lfa enable
!
```

! IOS-XE

```
Router ospf 1
segment-routing area 0 mpls
fast-reroute per-prefix
fast-reroute per-prefix ti-lfa enable
microloop avoidance segment-routing
```

TI-LFA protection Coverage

```
UOS-XR
RP/0/RSP0/CPU0:PE1#sh route 172.27.244.84
Fri Jun 23 15:01:00.947 MEST
Routing entry for 172.27.244.84/32
  Known via "ospf 1", distance 110, metric 31, labeled SR, type intra area
  Installed Jun 23 10:53:38.950 for 04:07:22
  Routing Descriptor Blocks
    172.16.1.2, from 172.16.0.6, via TenGigE0/0/0/10, Protected
      Route metric is 31
    172.16.1.26, from 172.16.0.6, via TenGigE0/0/0/8, Backup (TI-LFA)
      Repair Node(s): 172.16.0.41
     Route metric is 52
 No advertising protos.
```

Enabling TI-LFA under IS-IS interface (example)

```
router isis CORE
interface GigabitEthernet0/0/0/0
address-family ipv4 unicast
fast-reroute per-prefix ti-lfa
!
interface GigabitEthernet0/0/0/1
address-family ipv4 unicast
fast-reroute per-prefix ti-lfa
!
interface GigabitEthernet0/0/0/3
address-family ipv4 unicast
fast-reroute per-prefix ti-lfa
```

TI-LFA protection Coverage

RP/0/0/CPU0:XR-1#show isis fast-reroute summary											
IS-IS SR-AS-1 IPv4 Unicast FRR summary											
	Critica	l High	Medium	Low	Total						
	Priorit	y Priority	Priority	Priority							
Prefixes reachable in L2											
All paths protected	0	0	4	8	12						
Some paths protected	0	0	0	0	0						
Unprotected	0	0	0	0	0						
Protection coverage	0.00%	0.00%	100.00%	100.00%	100.00%						



Low Latency SR-TE Policy

An SR-TE policy is identified as an ordered list (head-end, color, end-point):

- Head-end Where the SR-TE policy is instantiated
- Color A numerical value that distinguishes between two or more policies to the same node pairs (Head-end – End point)
- End-point The destination of the SR-TE policy

Every SR-TE policy has a color value. Every policy between the same node pairs requires a unique color value.

Performance Measurement (PM)

- The PM for link delay uses the IP/UDP packet format defined in RFC 5357 (TWAMP-Light) for probes
- -The feature is required on all links over the path
- -P2P configuration on IGP is required for flooding the value.
- For link bundles, the hashing function may select a member link for forwarding but the reply may come from the remote line card on a different member link of the bundle.
- -For one-way delay measurement, clocks should be synchronized on two endpoint nodes of the link using PTP.

OSPF

Value Sub-TLV

- 27 Unidirectional Link Delay
- 28 Min/Max Unidirectional Link Delay
- 29 Unidirectional Delay Variation

OSPF

- RFC 7471 (OSPF Traffic Engineering (TE) Metric Extensions)
- Used to advertise extended TE metrics e.g. link delay (in usec)



For Color100 we have 3 available paths:

1. R-MPLS-PE-AREA6-1 (16061)- R-MPLS-P-1 (16001)- R-MPLS-PE-AREA7-1 (16071)

2. R-MPLS-PE-AREA6-1 (16061) - R-MPLS-PE-AREA6-2 (16062) - R-MPLS-P-2 (16002) - R-MPLS-PE-AREA7-2 (16072) - R-MPLS-PE-AREA7-1 (16071)

3. R-MPLS-PE-AREA6-1 (16061) - R-MPLS-PE-AREA6-2 (16062) - R-MPLS-P-2 (16002) - R-MPLS-P-1 (16001) - R-MPLS-PE-AREA7-1 (16071)

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

SR-TE R-MPLS-AREA6-1

! IOS-XR

```
Router ospf 65
 distribute link-state
 router-id 10.13.0.11
 segment-routing mpls
 segment-routing sr-prefer
 address-family ipv4 unicast
 area O
  !mpls traffic-eng
  interface Loopback0
  prefix-sid absolute 16061
   fast-reroute per-prefix
   fast-reroute per-prefix ti-lfa enable
interface TenGigabitEthernet0/0/1/1
interface TenGigabitEthernet0/0/1/0
```

show segment-routing traffic-eng policy color 100

```
segment-routing
 global-block 16000 23999
 traffic-enq
 policy SRTE SACOI 61 71
   color 100 end-point ipv4 10.13.0.13
    candidate-paths
    preference 100
      dynamic
       metric
        type latency
performance-measurement
 interface TenGigabitEthernet0/0/1/1
  delay-measurement
  next-hop ipv4 10.13.1.74
 interface TenGigabitEthernet0/0/1/0
  delay-measurement
```

```
next-hop ipv4 10.13.1.49
```

SR-TE R-MPLS-AREA6-1

```
!IOS-XR
interface TenGigE0/0/1/2.400 l2transport
 encapsulation dotlq 400
 rewrite ingress tag pop 1 symmetric
mtu 1518
12vpn
 pw-class PW CLASS SACOI 61 71
 encapsulation mpls
   ipv4 source 10.13.0.11
  preferred-path sr-te policy SRTE SACOI 61 71 fallback disable
```

SR-TE R-MPLS-AREA7-1

! IOS-XE

```
segment-routing mpls
set-attributes
address-family ipv4
sr-label-preferred
exit-address-family
!
global-block 16000 23999
!
connected-prefix-sid-map
address-family ipv4
10.13.0.13/32 absolute 16071 range 1
exit-address-family
```

```
router ospf 65
router-id 10.13.0.13
segment-routing area 0 mpls
segment-routing mpls
fast-reroute per-prefix ti-lfa
microloop avoidance segment-routing
distribute link-state
```

segment-routing traffic-eng policy SRTE SACOI 71 61 color 100 end-point 10.13.0.11 candidate-paths preference 100 constraints segments dataplane mpls dynamic metric type delay performance-measurement delay-measurement performance-measurement interface Ten0/4/1 delay-measurement next-hop ipv4 10.13.1.78 interface TenGigabitEthernet0/4/2 delay-measurement next-hop ipv4 10.13.1.57

RP/0/RSP1/CPU0:R-MPLS-PE-AREA6-1#sh performance-measurement sessions Thu Apr 4 13:41:50.922 CET

Transport type : Interface

Measurement type : Delay Measurement

Interface name : TenGigE0/0/1/0

Nexthop : 10.13.1.49

Delay Measurement session:

Session ID : 1

Last advertisement:

Advertised at: Apr 04 2024 09:46:46.113 (14104.829 seconds ago)

Advertised reason: First advertisement

Advertised delays (uSec): avg: 234, min: 233, max: 234, variance: 1

Next advertisement:

Threshold check scheduled in 2 more probes (roughly every 120 seconds) Aggregated delays (uSec): avg: 233, min: 233, max: 233, variance: 0 Rolling average (uSec): 233

Current Probe: Started at: Apr 04 2024 13:41:23.566 (27.376 seconds ago) Packets Sent: 10, received: 10 Measured delays (uSec): avg: 233, min: 233, max: 233, variance: 0 Next probe scheduled at: Apr 04 2024 13:41:53.499 (in 2.557 seconds) Next burst packet will be sent in 2.557 seconds Burst packet sent every 3.0 seconds Responder IP : 10.13.1.49 Number of Hops : 1 RP/0/RSP1/CPU0:R-MPLS-PE-AREA6-1#sh performance-measurement sessions !!second interface

Transport type : Interface

Measurement type : Delay Measurement

Interface name : TenGigE0/0/1/1

Nexthop : 10.13.1.74

Delay Measurement session:

Session ID : 2

Last advertisement:

Advertised at: Apr 04 2024 09:57:31.617 (13459.325 seconds ago) Advertised reason: First advertisement Advertised delays (uSec): avg: 10, min: 9, max: 10, variance: 1

Next advertisement:

Threshold check scheduled in 3 more probes (roughly every 120 seconds) No probes completed Rolling average (uSec): 9

Current Probe: Started at: Apr 04 2024 13:41:38.736 (12.206 seconds ago) Packets Sent: 5, received: 5 Measured delays (uSec): avg: 9, min: 9, max: 9, variance: 0 Next probe scheduled at: Apr 04 2024 13:42:08.661 (in 17.719 seconds) Next burst packet will be sent in 2.719 seconds Burst packet sent every 3.0 seconds Responder IP : 10.13.1.74 Number of Hops : 1 RP/0/RSP1/CPU0:R-MPLS-PE-AREA6-1#sh segment-routing traffic-eng policy Thu Apr 4 14:00:14.404 CET SR-TE policy database

Color: 100, End-point: 10.13.0.13 Name: srte_c_100_ep_10.13.0.13 Status: Admin: up Operational: up for 00:01:22 (since Apr 4 13:58:30.921) Candidate-paths: Preference: 100 (configuration) (active) Name: SRTE_SACOI_61_71 Requested BSID: dynamic Maximum SID Depth: 10 Dynamic (valid) Metric Type: LATENCY, Path Accumulated Metric: 172 16062 [Prefix-SID, 10.13.0.12] 16072 [Prefix-SID, 10.13.0.14] 16071 [Prefix-SID, 10.13.0.13] Attributes: Binding SID: 24187 Forward Class: Not Configured Steering labeled-services disabled: no Steering BGP disabled: no IPv6 caps enable: yes

RP/0/RSP1/CPU0:R-MPLS-PE-AREA6-1#sh segment-routing traffic-eng policy Thu Apr 4 14:21:14.582 CET SR-TE policy database

Color: 100, End-point: 10.13.0.13 Name: srte c 100 ep 10.13.0.13 Status: Admin: up Operational: up for 00:22:23 (since Apr 4 13:58:30.921) Candidate-paths: Preference: 100 (configuration) (active) Name: SRTE_SACOI_61_71 Requested BSID: dynamic Maximum SID Depth: 10 Dynamic (valid) Metric Type: LATENCY, Path Accumulated Metric: 299235 16062 [Prefix-SID, 10.13.0.12] 16002 [Prefix-SID, 10.13.0.2] 16001 [Prefix-SID, 10.13.0.1] 16071 [Prefix-SID, 10.13.0.13] Attributes: Binding SID: 24187 Forward Class: Not Configured Steering labeled-services disabled: no Steering BGP disabled: no IPv6 caps enable: yes

```
R-MPLS-PE-AREA7-1#sh segment-routing traffic-eng policy all
Name: SRTE_SACOI_71_61 (Color: 100 End-point: 10.13.0.11)
Owners : CLI
Status:
Admin: up, Operational: up for 00:00:20 (since 04-04 13:59:10.284)
Candidate-paths:
```

Preference 100 (CLI):

Dynamic (active)

Metric Type: DELAY, Path Accumulated Metric: 239

```
16061 [Prefix-SID, 10.13.0.11]
```

Attributes:

Binding SID: 228

Allocation mode: dynamic

State: Programmed

SR-TE R-MPLS-AREA6-2: Tshoot policy down

!IOS-XR

```
Router ospf 65
 distribute link-state
 router-id 10.13.0.12
 segment-routing mpls
 segment-routing sr-prefer
 address-family ipv4 unicast
 area O
 mpls traffic-eng
  interface Loopback0
  prefix-sid absolute 16062
   fast-reroute per-prefix
   fast-reroute per-prefix ti-lfa enable
 interface GigabitEthernet0/0/0/0
```

```
segment-routing
global-block 16000 23999
 traffic-enq
 policy SRTE SACOI 62 32
   color 100 end-point ipv4 10.13.0.18
    candidate-paths
    preference 100
      dynamic
       metric
        type latency
performance-measurement
 interface GigabitEthernet0/0/0/0
  delay-measurement
  next-hop ipv4 10.13.1.94
```

show segment-routing traffic-eng policy color 100

Why the policy is Down?

Last error: unresolved first label 16032

```
RP/0/RSP1/CPU0:R-MPLS-PE-AREA6-2#show segment-routing traffic-eng policy state down
Thu Mar 28 13:45:01.207 CET
SRTTE policy database
Color: 100, End-point: 10.13.0.18
 Name: srte c 100 ep 10.13.0.18
  SLatus:
    Admin: up Operational: down for 01:03:15 (since Mar 28 12:41:26.159)
  Candidate-paths:
    Preference: 100 (configuration)
     Name: SRTE SACOI 62 32
      Requested BSID: dynamic
       Maximum SID Depth: 10
      Dynamic (invalid)
      Last error: unresolved first label (16032)
       Metric Type: DATENCY, Path Accumulated Metric: 293
          16032 [Prefix-SID, 10.13.0.18]
  Attributes:
    Forward Class: 0
    Steering labeled-services disabled: no
    Steering BGF disabled: no
    1Fv6 caps enable: no
```

```
Policy Down : troubleshooting
 SR-TE R-MPLS-AREA6-2
                                        show mpls forwarding
                                        show mpls forwarding labels 16032
                                        detail
                                        show cef 10.13.0.18/32 detail
RP/0/RSP1/CPU0:R-MPLS-PE-AREA6-2#show mplowrwargment-routing traffic-eng
                                        policy
Tue Apr 2 11:25:43.332 CET
Local Outgoing Prefix
                     Outgoing Next Hop
                                               Bytes
Label Label or ID
                       Interface
                                              Switched
              SR Pfx (idx 1) Te0/0/1/0
16001 16001
                                    10.13.1.53
                                                 0
              SR Pfx (idx 1) Te0/0/1/1 10.13.1.73
      16001
                                                 0
             SR Pfx (idx 2) Te0/0/1/0 10.13.1.53
16002 Pop
                                                803428
[deleted output]
16061 Pop
             SR Pfx (idx 61) Te0/0/1/1 10.13.1.73
                                                 799204
16071 16071
              SR Pfx (idx 71) Te0/0/1/0 10.13.1.53
                                                 0
```

RP/0/RSP1/CPU0:R-MPLS-PE-AREA6-2#show mpls forwarding labels 16032 detail

Tue Apr 2 11:26:02.685 CET lempty output

We found a misconfiguration for prefix sid : on the router we had a typo 16132

REFERENCES

Performance measurement configuration guidelines

https://www.cisco.com/c/en/us/td/docs/iosxr/ncs5500/segment-routing/72x/b-segment-routing-cg-ncs5500-72x/ configure-performance-measurement.html#id_118505

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/seg_routing/configuration/xe-17/segrt-xe-17-book/m-sr-performancemeasurement.html

Configure SR TE policy Cisco IOS-XR

https://www.cisco.com/c/en/us/td/docs/iosxr/ncs560/segment-routing/79x/b-segment-routing-cg-79x-ncs560/m-configure-sr-te-policynew-cli-ncs5xx.pdf

SR MPLS – Performance Monitoring

https://www.segment-routing.net/images/tutorials/sr-mpls-pm-rev1g-km1.pdf



Segment Routing Part I by Clarence Filsfils (Author), Kris Michielsen (Author), Ketan Talaulikar (Author)

Segment Routing Part II – traffic Engineering by Clarence Filsfils (Author), Kris Michielsen (Author), Francois Clad (Author), Daniel Voyer (Author)