



Deploying MPLS Traffic Engineering

PacNOG15

Nurul Islam Roman (nurul@apnic.net)

Agenda


- Technology Overview
- Bandwidth optimization
- TE for QoS
- Traffic Protection
- Inter-Domain Traffic Engineering
- General Deployment Considerations




Technology Overview

MPLS TE Overview

- Introduces **explicit routing**
- Supports **constraint-based routing**
- Supports **admission control**
- Provides **protection** capabilities
- Uses **RSVP-TE** to establish LSPs
- Uses **ISIS / OSPF extensions** to advertise link attributes



How MPLS TE Works



- Link information Distribution*
 - ISIS-TE
 - OSPF-TE
- Path Calculation (CSPF)*
- Path Setup (RSVP-TE)
- Forwarding Traffic down Tunnel
 - Auto-route (announce / destinations)
 - Static route
 - PBR
 - CBTS / PBTS
 - Forwarding Adjacency
 - Tunnel select

PacNOG15


© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

5

Link Information Distribution

- Additional link characteristics
 - Interface address
 - Neighbor address
 - Physical bandwidth
 - Maximum reservable bandwidth
 - Unreserved bandwidth (at eight priorities)
 - TE metric
 - Administrative group (attribute flags)
- IS-IS or OSPF flood link information
- All TE nodes build a TE topology database
- Not required if using off-line path computation




PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

6

Path Calculation



- TE nodes can perform constraint-based routing
- Tunnel head end generally responsible for path calculation
- Constraints and topology database as input to path computation
- Shortest-path-first algorithm ignores links not meeting constraints
- Tunnel can be signaled once a path is found
- Not required if using offline path computation

PacNOG15

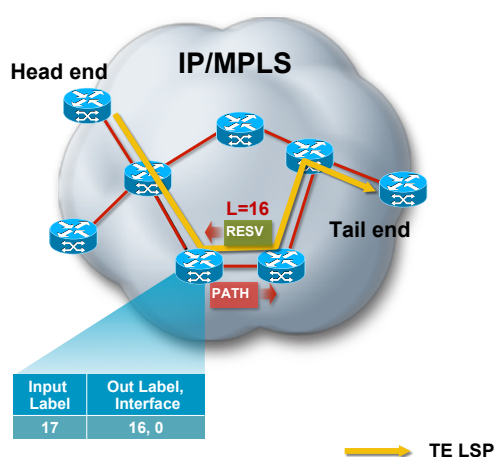
© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

7

TE LSP Signaling

- Tunnel signaled with TE extensions to RSVP
- Soft state maintained with **downstream PATH messages**
- Soft state maintained with **upstream RESV messages**
- New RSVP objects
 - LABEL_REQUEST (PATH)
 - LABEL (RESV)
 - EXPLICIT_ROUTE
 - RECORD_ROUTE (PATH/RESV)
 - SESSION_ATTRIBUTE (PATH)
- LFIB populated using RSVP labels allocated by RESV messages




PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

8

Traffic Selection



- Multiple traffic selection options
 - Auto-route
 - Static routes
 - Policy Based Routing
 - Forward Adjacency
 - Pseudowire Tunnel Selection
 - Class / Policy Based Tunnel Selection
- Tunnel path computation independent of routing decision injecting traffic into tunnel
- Traffic enters tunnel at head end

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

9

Configuring MPLS TE and Link Information Distribution Using IS-IS (Cisco IOS)

```

mpls traffic-eng tunnels
!
interface TenGigabitEthernet0/1/0
ip address 172.16.0.0 255.255.255.254
ip router isis
mpls traffic-eng tunnels
mpls traffic-eng attribute-flags 0xF
mpls traffic-eng administrative-weight 20
ip rsvp bandwidth 100000
!
router isis
net 49.0001.1720.1625.5001.00
is-type level-2-only
metric-style wide
mpls traffic-eng router-id Loopback0
mpls traffic-eng level-2
passive-interface Loopback0
!

```

Enable MPLS TE on this node

Enable MPLS TE on this interface

Attribute flags

TE metric

Maximum reservable bandwidth

Enable wide metric format and TE extensions (TE Id, router level)

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

1

Configuring MPLS TE and Link Information Distribution Using OSPF (Cisco IOS)

```

mpls traffic-eng tunnels
!
interface TenGigabitEthernet0/1/0
 ip address 172.16.0.0 255.255.255.254
 mpls traffic-eng tunnels
 mpls traffic-eng attribute-flags 0xF
 mpls traffic-eng administrative-weight 20
 ip rsvp bandwidth 100000
!
router ospf 100
 log-adjacency-changes
 passive-interface Loopback0
 network 172.16.0.0 0.0.255.255 area 0
 mpls traffic-eng router-id Loopback0
 mpls traffic-eng area 0
!

```

Enable MPLS TE on this node

Enable MPLS TE on this interface

Attribute flags

TE metric

Maximum reservable bandwidth

Enable TE extensions (TE router id and area)

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

11

Configuring MPLS TE and Link Information Distribution Using IS-IS (Cisco IOS XR)

```

router isis DEFAULT
 is-type level-2-only
 net 49.0001.1720.1625.5129.00
 address-family ipv4 unicast
 metric-style wide
 mpls traffic-eng level 2
 mpls traffic-eng router-id Loopback0
!
interface Loopback0
 passive
 address-family ipv4 unicast
!
interface TenGigE0/0/0/0
 address-family ipv4 unicast
!
!
rsvp
 interface TenGigE0/0/0/0
 bandwidth 100000
!
mpls traffic-eng
 interface TenGigE0/0/0/0
 admin-weight 5
 attribute-flags 0x8
!

```

Enable wide metric format and TE extensions (TE Id, router level)

Configuration mode for RSVP global and interface commands

Maximum reservable bandwidth

Configuration mode for MPLS TE global and interface commands

TE metric

Attribute flags

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

12

Configuring MPLS TE and Link Information Distribution Using OSPF (Cisco IOS XR)

```

router ospf DEFAULT
 area 0
  mpls traffic-eng
  interface Loopback0
    passive
  !
  interface TenGigE0/0/0/0
  !
  mpls traffic-eng router-id Loopback0
  !
  rsvp
  interface TenGigE0/0/0/0
  bandwidth 100000
  !
  mpls traffic-eng
  interface TenGigE0/0/0/0
  admin-weight 5
  attribute-flags 0x8
  !
  !

```

Enable TE extensions on this area

TE router Id

Configuration mode for RSVP global and interface commands

Maximum reservable bandwidth

Configuration mode for MPLS TE global and interface commands

TE metric

Attribute flags

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

13

Configuring Tunnel at Head End (Cisco IOS)

```

interface Tunnell
 description FROM-ROUTER-TO-DST1
 ip unnumbered Loopback0
 tunnel destination 172.16.255.3
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng priority 5 5
 tunnel mpls traffic-eng bandwidth 10000
 tunnel mpls traffic-eng affinity 0x0 mask 0xF
 tunnel mpls traffic-eng path-option 5 explicit name PATH1
 tunnel mpls traffic-eng path-option 10 dynamic
 !
 ip explicit-path name PATH1 enable
 next-address 172.16.0.1
 next-address 172.16.8.0
 !

```

Destination (tunnel tail end)

TE tunnel (as opposed to GRE or others)

Setup/hold priorities

Signaled bandwidth

Consider links with 0x0/0xF as attribute flags

Tunnel path options (PATH1, otherwise dynamic)

Explicit PATH1 definition

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

14

Configuring Tunnel at Head End (Cisco IOS XR)

```
explicit-path name PATH1
index 1 next-address ipv4 unicast 172.16.0.4
index 2 next-address ipv4 unicast 172.16.0.7
index 3 next-address ipv4 unicast 172.16.4.2
```

```
!
interface tunnel-te1
description FROM-ROUTER-TO-DST1
ipv4 unnumbered Loopback0
priority 5 5
signalled-bandwidth 100000
destination 172.16.255.2
path-option 10 explicit name PATH1
path-option 20 dynamic
affinity f mask f
!
```

Explicit PATH1
definition

MPLS TE P2P tunnel

Setup/hold priorities

Signaled bandwidth

Destination (tunnel
tail end)

Tunnel path options
(PATH1, otherwise
dynamic)

Consider links with
0xF/0xF as attribute
flags

PacNOG15


© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

15

Characteristics of P2MP TE LSP

- Unidirectional
- Explicitly routed
- One head end, but **one or more** tail ends (destinations)
- **Same** characteristics (constraints, protection, etc.) for all destinations




PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

16

P2MP TE LSP Terminology



- **Head-end/Source:** Node where LSP signaling is initiated
- **Mid-point:** Transit node where LSP signaling is processed (not a head-end, not a tail-end)
- **Tail-end/Leaf/destination:** node where LSP signaling ends
- **Branch point:** Node where packet replication is performed
- **Source-to-leaf (S2L) sub-LSP:** P2MP TE LSP segment that runs from source to one leaf

PacNOG15


© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

17

P2MP TE LSP Path Computation

- Constrained Shortest Path First (CSPF) used to compute an adequate tree
- CSPF executed per destination
- TE topology database and tunnel constraints used as input for path computation
- Path constraints may include loose, included, excluded hops
- Same constraints for all destinations (bandwidth, affinities, priorities, etc.)
- Path computation yields explicit path to each destination
- No changes to OSPF/IS-IS TE extensions
- Static paths possible with offline path computation




PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

18

P2MP TE LSP Signaling



- Source sends unique PATH message per destination
- LFIB populated using RSVP labels allocated by RESV messages
- Multicast state built by reusing sub-LSP labels at branch points

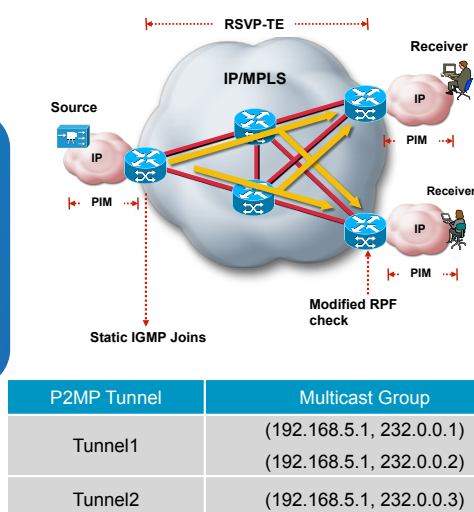
PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

19

P2MP TE LSP Traffic Selection IP Multicast



- One or more IP multicast groups mapped to a Tunnel
- Groups mapped via static IGMP join
- PIM outside of MPLS network
- Modified egress RPF check against TE LSP and tunnel head end (source address)
- Egress node may abstract TE LSP as a virtual interface (LSPVIF) for RPF purposes

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

20

Configuring P2MP Tunnel at Head End (Cisco IOS)

```

mpls traffic-eng destination list name P2MP-LIST-DST1
ip 172.16.255.1 path-option 10 explicit name PATH1
ip 172.16.255.2 path-option 10 dynamic
ip 172.16.255.3 path-option 10 dynamic
ip 172.16.255.4 path-option 10 dynamic
!
interface Tunnell
description FROM-ROUTER-TO-LIST-DST1
ip unnumbered Loopback0
ip pim passive
ip igmp static-group 232.0.0.1 source 192.168.5.1
ip igmp static-group 232.0.0.2 source 192.168.5.1
tunnel mode mpls traffic-eng point-to-multipoint
tunnel destination list mpls traffic-eng name P2MP-LIST-DST1
tunnel mpls traffic-eng priority 7 7
tunnel mpls traffic-eng bandwidth 1000
!

```

Destination list with one path-option per destination

Tunnel as passive PIM interface

Multicast groups mapped to tunnel

P2MP TE Tunnel

Destination list

Setup/hold priorities

Signaled bandwidth

Configuring RPF Check at P2MP Tunnel Tail End (Cisco IOS)

```

ip multicast mpls traffic-eng
ip mroute 192.168.5.1 255.255.255.255 172.16.255.5
!

```

Enable IPv4 multicast over P2MP TE LSP

Tunnel source (172.16.255.5) as next-hop for IP Multicast source (192.168.5.1) RPF check

Configuring P2MP Tunnel at Head End (Cisco IOS XR)

```

interface tunnel-mte1
  ipv4 unnumbered Loopback0
  destination 172.16.255.129
  path-option 10 explicit name PATH1
  path-option 20 dynamic
  destination 172.16.255.130
  path-option 10 dynamic
  priority 0 0
  signalled-bandwidth 100000
  node-capability label-switched-multicast
  multicast-routing
  address-family ipv4
  interface tunnel-mte1
  enable
  interface all enable
  !
router igmp
  vrf default
  interface tunnel-mte1
  static-group 232.0.0.1 192.168.5.1
  static-group 232.0.0.2 192.168.5.1
  !
  !

```

- MPLS TE P2MP tunnel
- Destination with path-option list
- Destination with single path-option
- Setup/hold priorities
- Signaled bandwidth
- Enable MPLS multicast
- Enable multicast forwarding over tunnel-mte1
- Multicast groups mapped to tunnel-mte1

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

23

Configuring RPF Check at P2MP Tunnel Tail End (Cisco IOS XR)

```

multicast-routing
  address-family ipv4
  core-tree-protocol rsvp-te
  static-rpf 192.168.5.1 32 mpls 172.16.255.3
  interface all enable
  !
  !

```

- Enable IPv4/v6 multicast over P2MP TE LSP
- Tunnel source (172.16.255.3) as next-hop for IP Multicast source (192.168.5.1) RPF check

PacNOG15

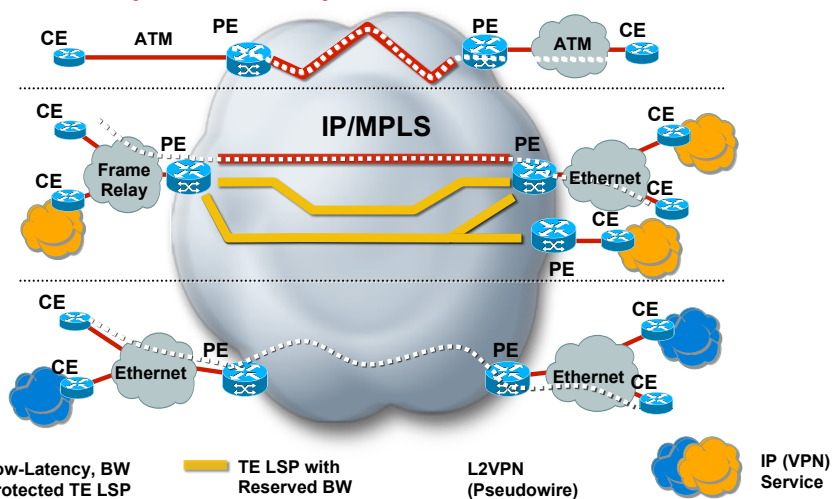
© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

24

MPLS TE Integration with Network Services

A TE LSP provides transport for different network services




PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

25

MPLS TE Deployment Models



PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

26



Bandwidth optimization


Planned Bandwidth Optimization

Traffic Matrix

	R1	R2	R3	R4	R5	R6
R1	4	7	1	5	4	5
R2	2	2	4	7	2	3
R3	1	2	9	5	5	5
R4	9	1	4	1	3	1
R5	3	7	9	2	7	7
R6	6	3	5	4	9	12


+

Physical Topology



=


Tunnel mesh to satisfy traffic matrix



- Tries to optimize underlying physical topology based on traffic matrix
- Key goal is to avoid link over/under utilization
- On-line (CSPF) or off-line path computation
- May result in a significant number of tunnels
- Should not increase your routing adjacencies

Traffic Matrix Measurement

- Interface counters on unconstrained tunnels
- Interface MIB
- MPLS LSR MIB
- NetFlow
 - NetFlow BGP Next Hop
 - MPLS-Aware NetFlow
 - Egress/Output NetFlow
- BGP policy accounting
 - Communities
 - AS path
 - IP prefix



PacNOG15


© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

29

AutoTunnel Mesh

- Mesh group: LSRs to mesh automatically
- Membership identified by
 - Matching TE Router ID against ACL
 - IGP mesh-group advertisement
- Each member automatically creates tunnel upon detection of a member
- Tunnels instantiated from template
- Individual tunnels not displayed in router configuration




PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

30

Auto Bandwidth



- Dynamically adjust bandwidth reservation based on measured traffic
- Optional minimum and maximum limits
- Sampling and resizing timers
- Tunnel resized to largest sample since last adjustment
- Actual resizing can be subject to adjustment threshold and overflow/underflow detection

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

31

Configuring AutoTunnel Mesh (Cisco IOS)

```

mpls traffic-eng tunnels
mpls traffic-eng auto-tunnel mesh
!
interface Auto-Templat1
ip unnumbered Loopback0
tunnel destination mesh-group 10
tunnel mode mpls traffic-eng
tunnel mpls traffic-eng autoroute announce
tunnel mpls traffic-eng path-option 10 dynamic
tunnel mpls traffic-eng auto-bw frequency 3600
!
router ospf 16
log-adjacency-changes
mpls traffic-eng router-id Loopback0
mpls traffic-eng area 0
mpls traffic-eng mesh-group 10 Loopback0 area 0
passive-interface Loopback0
network 172.16.0.0 0.0.255.255 area 0
!

```

Enable Auto-tunnel Mesh

Tunnel template

Template cloned for each member of mesh group 10

Dynamic (CSPF) path to each mesh group member

Tunnels will adjust bandwidth reservation automatically

Advertise mesh group 10 membership in area 0

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

32

Configuring AutoTunnel Mesh (Cisco IOS XR)

```

mpls traffic-eng
  auto-tunnel mesh
    group 10
      attribute-set 10
      destination-list DST-RID-ACL
    !
    tunnel-id min 1000 max 2000
  !
  attribute-set auto-mesh 10
    autoroute announce
    auto-bw collect-bw-only
  !
!

```

Enable Auto-tunnel Mesh

Mesh group 10

Attribute set to use

ACL matching matching TE router ids associated with mesh

Range of dynamically created tunnel interfaces

Attribute set definition

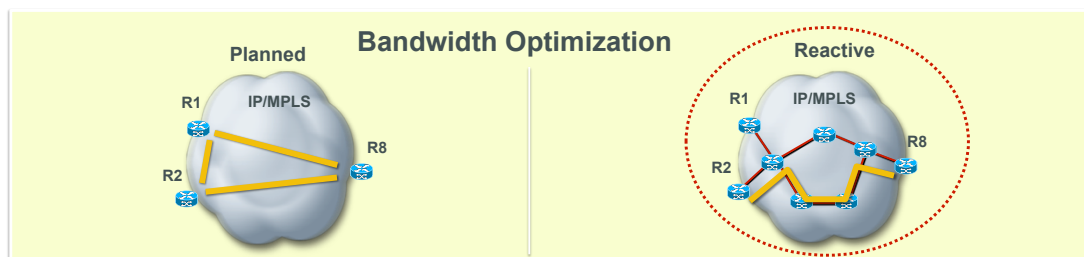
PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

33

Reactive Bandwidth Optimization



- Selective deployment of tunnels when highly-utilized links are identified
- Generally, deployed until next upgrade cycle alleviates congested links

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.


Cisco Public

34




TE for QoS

Motivations




- Point-to-point SLAs
- Admission control
- Integration with DiffServ architecture
- Increased routing control to improve network performance


MPLS TE and DiffServ Deployment Models



- A solution when:
 - No differentiation required
 - Optimization required
- Limit **link** load to **actual link** capacity
- No notion of traffic classes



- A solution when:
 - Differentiation required
 - Optimization required
- Limit **class** capacity to **expected class** load
- Limit **class** load to **actual class** capacity for one class




- A solution when:
 - Strong differentiation required
 - Fine optimization required
- Limit **class** capacity to **expected class** load
- Limit **class** load to **actual class** capacity for at least two classes

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

DiffServ-Aware Traffic Engineering



- Enables per-class traffic engineering
- IS-IS or OSPF flood link information (as usual)
- Per-class unreserved bandwidth on each link
- New RSVP object (CLASSTYPE)
- Nodes manage link bandwidth using a bandwidth constraint model
- Two models defined
 - Maximum Allocation Model (MAM)
 - Russian Doll Model (RDM)
- Unique class definition and constraint model throughout network
- Two classes (class-types) in current implementations

PacNOG15


© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

38

Maximum Allocation Model (MAM)

- BW pool applies to one class
- Sum of BW pools may exceed MRB
- Sum of total reserved BW may not exceed MRB
- Current implementation supports BC0 and BC1



PacNOG15


© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

39

Russian Dolls Model (RDM)

- BW pool applies to one or more classes
- Global BW pool (BC0) equals MRB
- BC0..BCn used for computing unreserved BW for class n
- Current implementation supports BC0 and BC1



PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

40

Configuring DS-TE Classes and Bandwidth Constraints (Cisco IOS)

RDM

```
mpls traffic-eng tunnels
mpls traffic-eng ds-te mode ietf
mpls traffic-eng ds-te te-classes
  te-class 0 class-type 1 priority 0
  te-class 1 class-type 1 priority 1
  te-class 2 class-type 1 priority 2
  te-class 3 class-type 1 priority 3
  te-class 4 class-type 0 priority 4
  te-class 5 class-type 0 priority 5
  te-class 6 class-type 0 priority 6
  te-class 7 class-type 0 priority 7
!
interface TenGigabitEthernet0/1/0
 ip address 172.16.0.0 255.255.255.254
 mpls traffic-eng tunnels
 ip rsvp bandwidth rdm bc0 155000 bc1 55000
!
```

Enable IETF DS-TE

Explicit TE-Class definition

RDM bandwidth constraints

MAM

```
mpls traffic-eng tunnels
mpls traffic-eng ds-te mode ietf
mpl traffic-eng ds-te bc-model mam
!
interface TenGigabitEthernet0/1/0
 ip address 172.16.0.0 255.255.255.254
 mpls traffic-eng tunnels
 ip rsvp bandwidth mam max-reservable-bw 155000 bc0 100000 bc1 55000
!
```

Enable IETF DS-TE and use default TE-Class definition

Enable MAM

MAM bandwidth constraints

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

41

Configuring DS-TE Tunnel (Cisco IOS)

```
interface Tunnel1
 description FROM-ROUTER-TO-DST1-CT0
 ip unnumbered Loopback0
 no ip directed-broadcast
 tunnel destination 172.16.255.3
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng priority 5 5
 tunnel mpls traffic-eng bandwidth 100000 class-type 0
 tunnel mpls traffic-eng path-option 10 dynamic
!
interface Tunnel2
 description FROM-ROUTER-TO-DST1-CT1
 ip unnumbered Loopback0
 no ip directed-broadcast
 tunnel destination 172.16.255.3
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng priority 0 0
 tunnel mpls traffic-eng bandwidth 50000 class-type 1
 tunnel mpls traffic-eng path-option 10 dynamic
!
```

Signal Tunnel1 with CT0 (priority and CT must match valid TE-Class)

Signal Tunnel2 with CT1 (priority and CT must match valid TE-Class)

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

42

Configuring DS-TE Classes and Bandwidth Constraints (Cisco IOS XR)

RDM

```

rsvp
 interface TenGigE0/0/0/0
  bandwidth rdm bc0 155000 bc1 55000
  !
mpls traffic-eng
 interface TenGigE0/0/0/0
  !
 ds-te mode ietf
 ds-te te-classes
  te-class 0 class-type 1 priority 0
  te-class 1 class-type 1 priority 1
  te-class 2 class-type 1 priority 2
  te-class 3 class-type 1 priority 3
  te-class 4 class-type 0 priority 4
  te-class 5 class-type 0 priority 5
  te-class 6 class-type 0 priority 6
  te-class 7 class-type 0 priority 7
  !
  !

```

RDM bandwidth constraints

Enable IETF DS-TE

Explicit TE-Class definition

MAM

```

rsvp
 interface TenGigE0/0/0/0
  bandwidth mam max-reservable-bw 155000 bc0 100000 bc1 55000
  !
mpls traffic-eng
 interface TenGigE0/0/0/0
  !
 ds-te mode ietf
 ds-te bc-model mam
  !
  !

```

MAM bandwidth constraints

Enable IETF DS-TE and use default TE-Class definition

Enable MAM

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

43

Configuring DS-TE Tunnels (Cisco IOS XR)

```

interface tunnel-te1
 description FROM-ROUTER-TO-DST1-CT0
 ipv4 unnumbered Loopback0
 priority 5 5
 signalled-bandwidth 100000 class-type 0
 destination 172.16.255.2
 path-option 10 dynamic
 !
interface tunnel-te2
 description FROM-ROUTER-TO-DST1-CT1
 ipv4 unnumbered Loopback0
 priority 0 0
 signalled-bandwidth 50000 class-type 1
 destination 172.16.255.2
 path-option 10 dynamic
 !

```

Signal tunnel-te1 with CT0 (priority and CT must match valid TE-Class)

Signal tunnel-te2 with CT1 (priority and CT must match valid TE-Class)

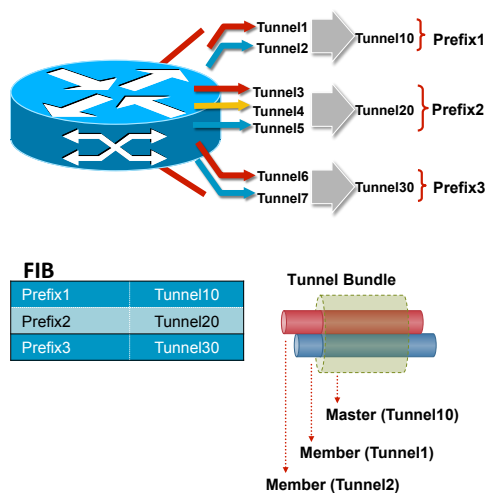
PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

44

Class-Based Tunnel Selection: CBTS



- EXP-based selection between **multiple tunnels** to **same destination**
- Local mechanism at head-end (no IGP extensions)
- Tunnel master bundles tunnel members
- Tunnel selection configured on tunnel master (auto-route, etc.)
- Bundle members configured with EXP values to carry
- Bundle members may be configured as default
- Supports VRF traffic, IP-to-MPLS and MPLS-to-MPLS switching paths

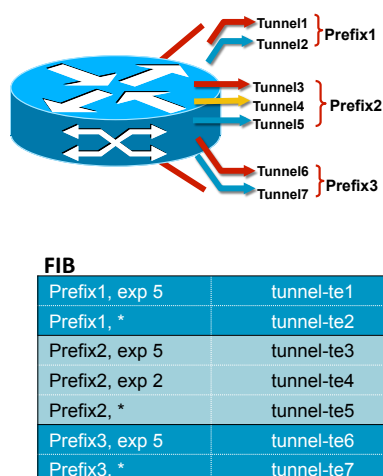
PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

45

Policy-based Tunnel Selection: PBTS



- EXP-based selection between **multiple tunnels** to **same destination**
- Local mechanism at head-end
- Tunnels configured via `policy-class` or `forwarding-class` with EXP values to carry
- No IGP extensions
- Supports VRF traffic, IP-to-MPLS and MPLS-to-MPLS switching

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

46

Configuring CBTS (Cisco IOS)

```

interface Tunnel1
 ip unnumbered Loopback0
 tunnel destination 172.16.255.2
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng bandwidth 50000 class-type 1
 tunnel mpls traffic-eng path-option 10 dynamic
 tunnel mpls traffic-eng exp 5
!
interface Tunnel2
 ip unnumbered Loopback0
 tunnel destination 172.16.255.2
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng bandwidth 100000 class-type 0
 tunnel mpls traffic-eng path-option 10 dynamic
 tunnel mpls traffic-eng exp default
!
interface Tunnel10
 ip unnumbered Loopback0
 tunnel destination 172.16.255.2
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng exp-bundle master
 tunnel mpls traffic-eng exp-bundle member Tunnel1
 tunnel mpls traffic-eng exp-bundle member Tunnel2
!
ip route 192.168.0.0 255.255.255.0 Tunnel10
!

```

Tunnel1 will carry packets with MPLS EXP 5

Tunnel2 will carry packets with MPLS EXP other than 5

Tunnel10 defined as bundle master with Tunnel2 and Tunnel1 as members

CBTS performed on prefix 192.168.0.0/24 using Tunnel10

Configuring PBTS (Cisco IOS XR)

```

interface tunnel-te1
 ipv4 unnumbered Loopback0
 autoroute announce
 signalled-bandwidth 10000
 destination 172.16.255.2
 policy-class 5
 path-option 10 explicit name PATH1
 path-option 20 dynamic
!
interface tunnel-te2
 ipv4 unnumbered Loopback0
 autoroute announce
 signalled-bandwidth 50000
 destination 172.16.255.2
 path-option 10 explicit name PATH2
 path-option 20 dynamic
!

```


tunnel-te1 will carry packets with MPLS EXP 5

tunnel-te2 will carry packets with MPLS EXP other than 5 (default tunnel)



Traffic Protection


Traffic Protection Using MPLS TE Fast Re-Route (FRR)



- Sub-second recovery against node/link failures
- Scalable 1:N protection
- Greater protection granularity
- Cost-effective alternative to 1:1 protection
- Bandwidth protection

FRR Link Protection Operation

- Requires **pre-signalled next-hop** (NHOP) backup tunnel
- Point of Local Repair (PLR) swaps label and pushes backup label
- Backup terminates on Merge Point (MP) where traffic re-joins primary
- Restoration time expected under ~50 ms



PacNOG15


© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

51

FRR Node Protection Operation

- Requires **pre-signalled next-next-hop** (NNHOP) backup tunnel
- Point of Local Repair (PLR) swaps **next-hop label** and pushes backup label
- Backup terminates on Merge Point (MP) where traffic re-joins primary
- Restoration time depends on failure detection time



PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

52

Configuring FRR (Cisco IOS)

Primary Tunnel

```
interface Tunnel1
  description FROM-ROUTER-TO-DST1-FRR
  ip unnumbered Loopback0
  tunnel destination 172.16.255.2
  tunnel mode mpls traffic-eng
  tunnel mpls traffic-eng bandwidth 20000
  tunnel mpls traffic-eng path-option 10 dynamic
  tunnel mpls traffic-eng fast-reroute
!
```

Indicate the desire for local protection during signaling

Backup Tunnel

```
interface Tunnel1
  description NNHOP-BACKUP
  ip unnumbered Loopback0
  tunnel destination 172.16.255.2
  tunnel mode mpls traffic-eng
  tunnel mpls traffic-eng path-option 10 explicit name PATH1
!
interface TenGigabitEthernet1/0/0
  ip address 172.16.192.5 255.255.255.254
  mpls traffic-eng tunnels
  mpls traffic-eng backup-path Tunnel1
  ip rsvp bandwidth
!
```

Explicitly routed backup to 172.16.255.2 with zero bandwidth

Use Tunnel1 as backup for protected LSPs through TenGigabitEthernet1/0/0

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

53

Configuring FRR (Cisco IOS XR)

Primary Tunnel

```
interface tunnel-te1
  description FROM-ROUTER-TO-DST1-FRR
  ipv4 unnumbered Loopback0
  signalled-bandwidth 30000
  destination 172.16.255.2
  fast-reroute
  path-option 10 dynamic
!
```

Indicate the desire for local protection during signaling

Backup Tunnel

```
interface tunnel-te1
  description NHOP-BACKUP
  ipv4 unnumbered Loopback0
  destination 172.16.255.130
  path-option 10 explicit name PATH1
!
mpls traffic-eng
  interface TenGigE0/0/0/0
    backup-path tunnel-te 1
  !
!
```

Explicitly routed backup to 172.16.255.130 with zero bandwidth

Use tunnel-te1 as backup for protected LSPs through TenGigE0/0/0/0

PacNOG15


© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

54

Bidirectional Forwarding Detection Trigger for FRR

- FRR relies on quick PLR failure detection
- Some failures may not produce loss of signal or alarms on a link
- BFD provides light-weight neighbor connectivity failure detection
- Preferred over RSVP Hellos



PacNOG15


© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

55

Bandwidth Protection

- Backup tunnel with associated bandwidth capacity
- Backup tunnel may or may not actually signal bandwidth
- PLR will decide best backup to protect primary
 - nhop/nnhop
 - backup-bw
 - class-type
 - node-protection flag



PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.


Cisco Public

56

AutoTunnel: Primary Tunnels

What's the Problem?

- FRR can protect TE Traffic
- No protection mechanism for IP or LDP traffic
- How to leverage FRR for all traffic?
- What if protection desired without traffic engineering?



PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.


Cisco Public

57

AutoTunnel: Primary Tunnels

What's the Solution?

Forward all traffic through a one-hop protected primary TE tunnel



- Create protected one-hop tunnels on all TE links

Priority	7/7
Bandwidth	0
Affinity	0x0/0xFFFF
Auto-BW	OFF
Auto-Route	ON
Fast-Reroute	ON
Forwarding-Adj	OFF
Load-Sharing	OFF

- Tunnel interfaces not shown on router configuration
- Configure desired backup tunnels (manually or automatically)

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.


Cisco Public

58

AutoTunnel: Primary Tunnels

Why One-Hop Tunnels?

- CSPF and SPF yield same results (absence of tunnel constraints)
- Auto-route forwards all traffic through one-hop tunnel
- Traffic logically mapped to tunnel but no label imposed (imp-null)
- traffic is forwarded as if no tunnel was in place



PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

59

Configuring AutoTunnel Primary Tunnels (Cisco IOS)

```
mpls traffic-eng tunnels
mpls traffic-eng auto-tunnel primary onhop
mpls traffic-eng auto-tunnel primary tunnel-num min 900 max 999
!
```

Enable auto-tunnel
primary

Range for tunnel
interfaces

PacNOG15


© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

60

AutoTunnel: Backup Tunnels What's the Problem?

- MPLS FRR requires backup tunnels to be preconfigured
- Automation of backup tunnels is desirable



PacNOG15


© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

61

AutoTunnel: Backup Tunnels What's the Solution?

Create backup tunnels automatically as needed



- Detect if a primary tunnel requires protection and is not protected
- Verify that a backup tunnel doesn't already exist
- Compute a backup path to NHOP and NNHOP excluding the protected facility
- Optionally, consider shared risk link groups during backup path computation
- Signal the backup tunnels


PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

62

AutoTunnel: Backup Tunnels What's the Solution? (Cont.)



- Backup tunnels are preconfigured

Priority	7/7
Bandwidth	0
Affinity	0x0/0xFFFF
Auto-BW	OFF
Auto-Route	OFF
Fast-Reroute	OFF
Forwarding-Adj	OFF
Load-Sharing	OFF
- Backup tunnel interfaces and paths not shown on router configuration

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

63

Configuring AutoTunnel Backup Tunnels (Cisco IOS)

```

mpls traffic-eng tunnels
mpls traffic-eng auto-tunnel backup nhop-only
mpls traffic-eng auto-tunnel backup tunnel-num min 1900 max 1999
mpls traffic-eng auto-tunnel backup timers removal unused 7200
mpls traffic-eng auto-tunnel backup srlg exclude preferred
!

```

Enable auto-tunnel
backup (NHOP
tunnels only)

Range for tunnel
interfaces

Tear down unused
backup tunnels

Consider SRLGs
preferably

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

64

Configuring AutoTunnel Backup Tunnels (Cisco IOS XR)

```

!
ipv4 unnumbered mpls traffic-eng Loopback 0
!
mpls traffic-eng
 interface GigabitEthernet0/0/0/0
  auto-tunnel backup
  exclude srlg preferred
  nhop-only
!
!
auto-tunnel backup
 timers removal unused 7200
 tunnel-id min min 1900 max 1999
!
!

```

Source interface for backup tunnels

Protect interface with dynamically created backup tunnels

Consider SRLGs preferably

Create NHOP backup tunnel only

Tear down unused backup tunnels

Range for tunnel interfaces

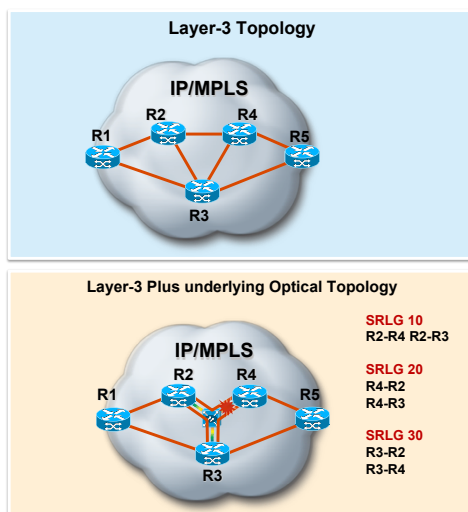
PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

65

Shared Risk Link Group (SRLG)



- Some links may share same physical resource (e.g. fiber, conduit)
- AutoTunnel Backup can force or prefer exclusion of SRLG to guarantee diversely routed backup tunnels
- IS-IS and OSPF flood SRLG membership as an additional link attribute

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

66

Configuring SRLG (Cisco IOS)

```
interface TenGigabitEthernet0/1/0
ip address 172.16.0.0 255.255.255.254
mpls traffic-eng tunnels
mpls traffic-eng srlg 15
mpls traffic-eng srlg 25
ip rsvp bandwidth
!
interface TenGigabitEthernet1/0/0
ip address 172.16.0.2 255.255.255.254
mpls traffic-eng tunnels
mpls traffic-eng srlg 25
ip rsvp bandwidth
!
```

Interface member of
SRLG 15 and 25

Interface member of
SRLG 25

Configuring SRLG (Cisco IOS XR)


```
srlg
interface POS0/4/0/0
value 15
value 25
!
interface POS0/4/0/1
value 25
!
!
```

Interface member of
SRLG 15 and 25

Interface member of
SRLG 25

What About Path Protection?

- Primary and standby share head and tail, but expected to be diversely routed
- Generally higher restoration times compared to local protection
- Doubles number of TE LSPs (1:1 protection)
- May be an acceptable solution for restricted topologies (e.g. rings)
- Cisco IOS
 - Separate path option sequences for primary and standby
 - Explicit paths only
 - No path diversity
- Cisco IOS XR
 - Single path-option sequence for primary and standby
 - Explicit and dynamic paths
 - Automatic path diversity (node-link, node, link)



PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

69

Configuring Path Protection (Cisco IOS)

```
interface Tunnel1
 ip unnumbered Loopback0
 tunnel mode mpls traffic-eng
 tunnel destination 172.16.255.2
 tunnel mpls traffic-eng autoroute announce
 tunnel mpls traffic-eng path-option 10 explicit name PATH1
 tunnel mpls traffic-eng path-option 20 explicit name PATH2
 tunnel mpls traffic-eng path-option protect 10 explicit name PPATH1
 tunnel mpls traffic-eng path-option protect 20 explicit name PPATH2
!
```

Standby path to be
used for PATH1

Standby path to be
used for PATH2

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

70

Configuring Enhanced Path Protection (Cisco IOS)

```
mpls traffic-eng path-option list name PATH-LST
```

```
path-option 10 explicit name PE1-P3-P4-PE2
path-option 20 explicit name PE1-P5-P6-PE2
path-option 30 explicit name PE1-P7-P8-PE2
```

```
!
```

```
interface Tunnell
```

```
ip unnumbered Loopback0
```

```
tunnel mode mpls traffic-eng
```

```
tunnel destination 172.16.255.2
```

```
tunnel mpls traffic-eng autoroute announce
```

```
tunnel mpls traffic-eng path-option 10 explicit name PE1-P1-P2-PE2
```

```
tunnel mpls traffic-eng path-option protect 10 list name PATH-LST
```

```
!
```

→ List of standby paths

→ Use path list to protect primary path

Configuring Path Protection (Cisco IOS XR)

```
interface tunnel-tel
```

```
description FROM-ROUTER-TO-DST1
```

```
ipv4 unnumbered Loopback0
```

```
signalled-bandwidth 100000
```

```
destination 172.16.255.2
```

```
affinity f mask f
```

```
path-protection
```

```
path-option 10 explicit name PATH1
```

```
path-option 20 explicit name PATH2
```


```
path-option 30 dynamic
```

```
!
```

→ Signal an acceptable (node-link, node, link diverse) standby TE LSP based on path-option sequence

P2MP TE LSP Traffic Protection

- No new protocol extensions to support FRR
- Protection requirement applies to all destinations
- P2P LSP as backup tunnel for a sub-LSP
- No changes to label stacking procedure
- Only link protection supported
- Head-end protection requires path redundancy (live-standby / live-live)



PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

73

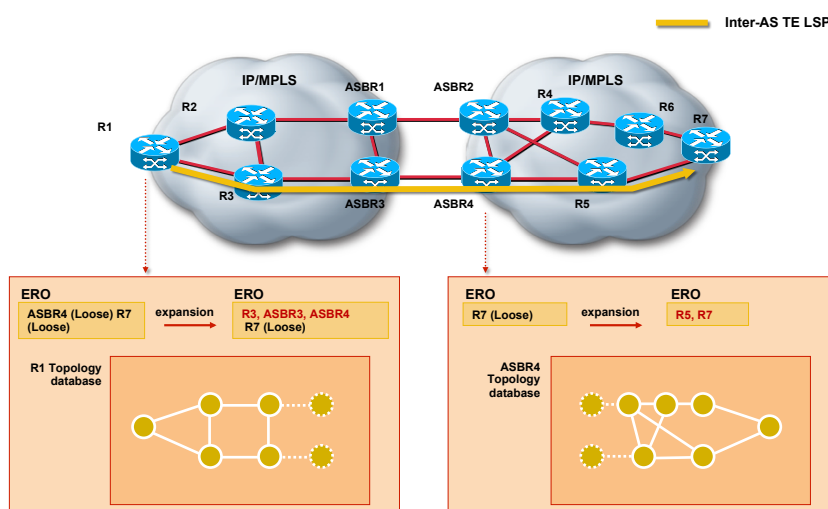


Inter-Domain Traffic Engineering

Inter-Domain Traffic Engineering: Introduction

- Domain defined as an IGP area or autonomous system
- Head end lacks complete network topology to perform path computation in both cases
- Two path computation approaches
 - Per-domain (ERO loose-hop expansion)
 - Distributed (Path Computation Element)

Per-Domain Path Computation Using ERO Loose-hop Expansion



Configuring Inter-Area Tunnels (Cisco IOS)

```

mpls traffic-eng tunnels
!
interface Tunnel1
 ip unnumbered Loopback0
 no ip directed-broadcast
 tunnel destination 172.16.255.7
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng path-option 10 explicit name LOOSE-PATH
!
ip route 172.16.255.7 255.255.255.255 Tunnel1
!
ip explicit-path name LOOSE-PATH enable
 next-address loose 172.16.255.3
 next-address loose 172.16.255.5
!

```

Loose-hop path

Static route mapping
IP traffic to Tunnel1

List of ABRs as loose
hops

Configuring Inter-Area Tunnels with Autoroute Destinations (Cisco IOS)

```

interface Tunnel1
 ip unnumbered Loopback0
 tunnel mode mpls traffic-eng
 tunnel destination 172.16.255.7
 tunnel mpls traffic-eng autoroute destination
 tunnel mpls traffic-eng path-option 10 explicit name LOOSE-PATH
!
ip explicit-path name LOOSE-PATH enable
 next-address loose 172.16.255.3
 next-address loose 172.16.255.5
!

```

Create static route to
tunnel destination
(172.16.255.7)

Loose-hop path

List of **ABRs** as
loose hops

Configuring Inter-Area Tunnels (Cisco IOS XR)

```
explicit-path name LOOSE-PATH
index 1 next-address loose ipv4 unicast 172.16.255.129
index 2 next-address loose ipv4 unicast 172.16.255.131
!
interface tunnel-te1
description FROM-ROUTER-TO-DST3
ipv4 unnumbered Loopback0
destination 172.16.255.2
path-option 10 explicit name LOOSE-PATH
!
router static
address-family ipv4 unicast
172.16.255.2/32 tunnel-te1
!
```


List of ABRs as loose hops

Loose-hop path


Static route mapping IP traffic to tunnel-te1

Distributed Path Computation using Path Computation Element

Backward Recursive PCE-based Computation (BRPC)



Inter-Domain TE – Fast Re-route



- Same configuration as single domain scenario
- Support for node-id sub-object required to implement ABR/ASBR node protection
- Node-id helps point of local repair (PLR) detect a merge point (MP)

Inter-Domain TE Take into Account before Implementing

- Semantics of link attributes across domain boundaries
- Semantics of TE-Classes across domain boundaries for DS-TE
- Auto-route destinations creates a static route to tunnel destination and facilitates traffic selection
- Auto-route announce not applicable for traffic selection




General Deployment Considerations

Should RSVP-TE and LDP be Used Simultaneously?

- Guarantees forwarding of VPN traffic if a TE LSP fails
- May be required if full mesh of TE LSPs not in use
- Increased complexity

How Far should Tunnels Span?



- **PE-to-PE Tunnels**
 - More granular control on traffic forwarding
 - Larger number of TE LSPs
- **P-to-P Tunnels**
 - Requires IP tunnels or LDP over TE tunnels to carry VPN traffic (deeper label stack)
 - Fewer TE LSPs
 - May be extended with PE-P tunnels


PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

85

MPLS TE on Link Bundles



- Different platforms support different link bundles
 - Ethernet
 - POS
 - Multilink PPP
- Bundles appear as single link in topology database
- Same rules for link state flooding
- LSP preemption if bundle bandwidth becomes insufficient
- Configurable minimum number of links to maintain bundle active
- Bundle failure can act as trigger for FRR


PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

86

Scaling Signaling (Refresh Reduction)



- RSVP soft state needs to be refreshed periodically
- Refresh reduction extensions use message Identifier associated with Path/Resv state
- Summary Refresh (SRefresh) message refreshes state using a message_id list
- SRefresh only replaces refresh Path/Resv messages

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

87

Configuring Refresh Reduction (Cisco IOS)

```

mpls traffic-eng tunnels
!
interface TenGigabitEthernet0/1/0
 ip address 172.16.0.0 255.255.255.254
 mpls traffic-eng tunnels
 ip rsvp bandwidth 100000
!
router ospf 100
 log-adjacency-changes
 passive-interface Loopback0
 network 172.16.0.0 0.0.255.255 area 0
 mpls traffic-eng router-id Loopback0
 mpls traffic-eng area 0
!
ip rsvp signalling refresh reduction
!

```

Enable refresh reduction

* Enabled by default in Cisco IOS XR

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

88



Summary

Summary

- Technology Overview
 - Explicit and constrained-based routing
 - TE protocol extensions (OSPF, ISIS and RSVP)
 - P2P and P2MP TE LSP
- Bandwidth optimization
 - Planned (full mesh, auto-tunnel)
 - Reactive
- TE for QoS
 - DS-TE (MAM, RDM)
 - CBTS
- Traffic Protection
 - Link/node protection (auto-tunnel)
 - Bandwidth protection
- Inter-Domain Traffic Engineering
 - Inter-Area
 - Inter-AS (Authentication, policy control)
- General Deployment Considerations
 - MPLS TE and LDP
 - PE-to-PE vs. P-to-P tunnels
 - TE over Bundles
 - Scaling signaling

