# Segment Routing Topology Independent LFA (TI-LFA)

Clarence Filsfils
Kris Michielsen

# Topology Independent LFA (TI-LFA)

- Introduction to TI-LFA
- Simple, optimal and topology independent sub-50ms per-prefix protection
- Protects SR, LDP and IP traffic
- Examples of TI-LFA implementation

# Classic LFA Introduction

# Classic Loop Free Alternate Fast ReRoute (LFA FRR)

- Per-prefix LFA: Simple, automatic, local, sub-50msec fast reroute technique
- IGP pre-computes a backup path per primary path per IGP destination: per-path IP optimality
- The backup path is pre-installed in data plane
- Upon local failure, all the backup paths of the impacted destinations are enabled in a prefix-independent manner (<50msec loss of connectivity)

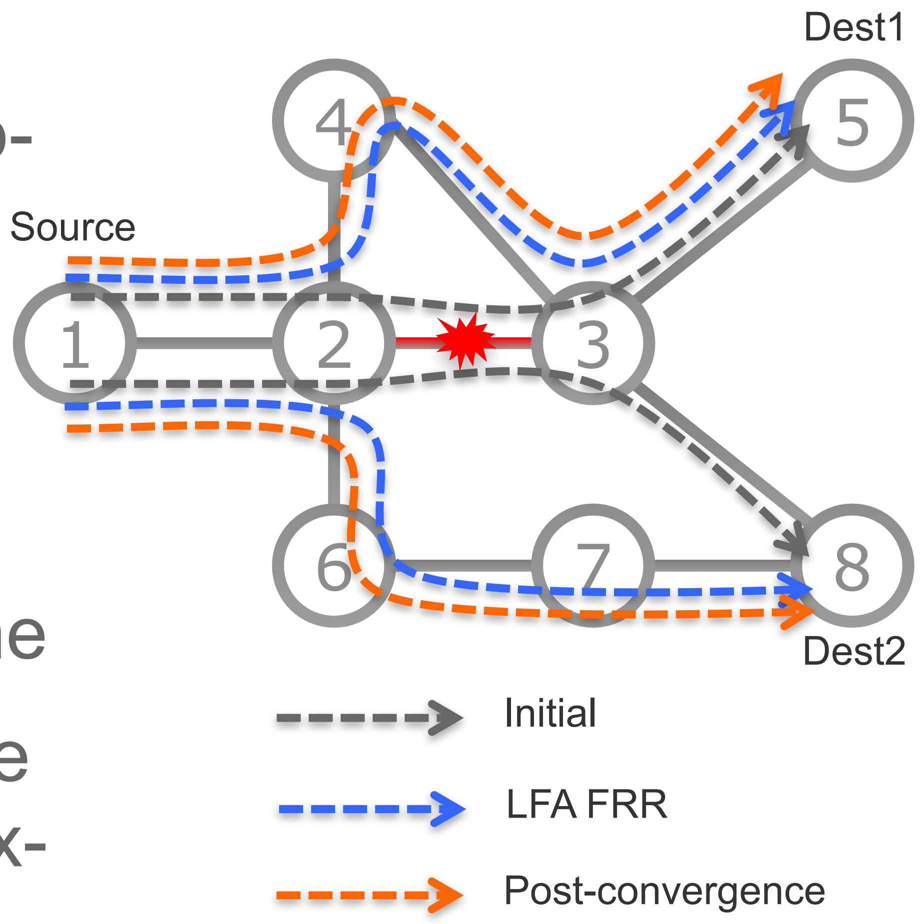

# Classic Per-Prefix LFA – disadvantages

- Classic LFA has disadvantages:
  - Incomplete coverage, topology dependent
  - Not always providing most optimal backup path

→ Topology Independent LFA (TI-LFA) solves these issues

# Classic LFA has partial coverage

- Classic LFA is topology dependent: not all topologies provide LFA for all destinations
  - Depends on network topology and metrics
  - E.g. Node6 is not an LFA for Dest1 (Node5) on Node2, packets would loop since Node6 uses Node2 to reach Dest1 (Node5)
    → Node2 does not have an LFA for this destination (no - - -> backup path in topology)

→ Topology Independent LFA (TI-LFA) provides 100% coverage

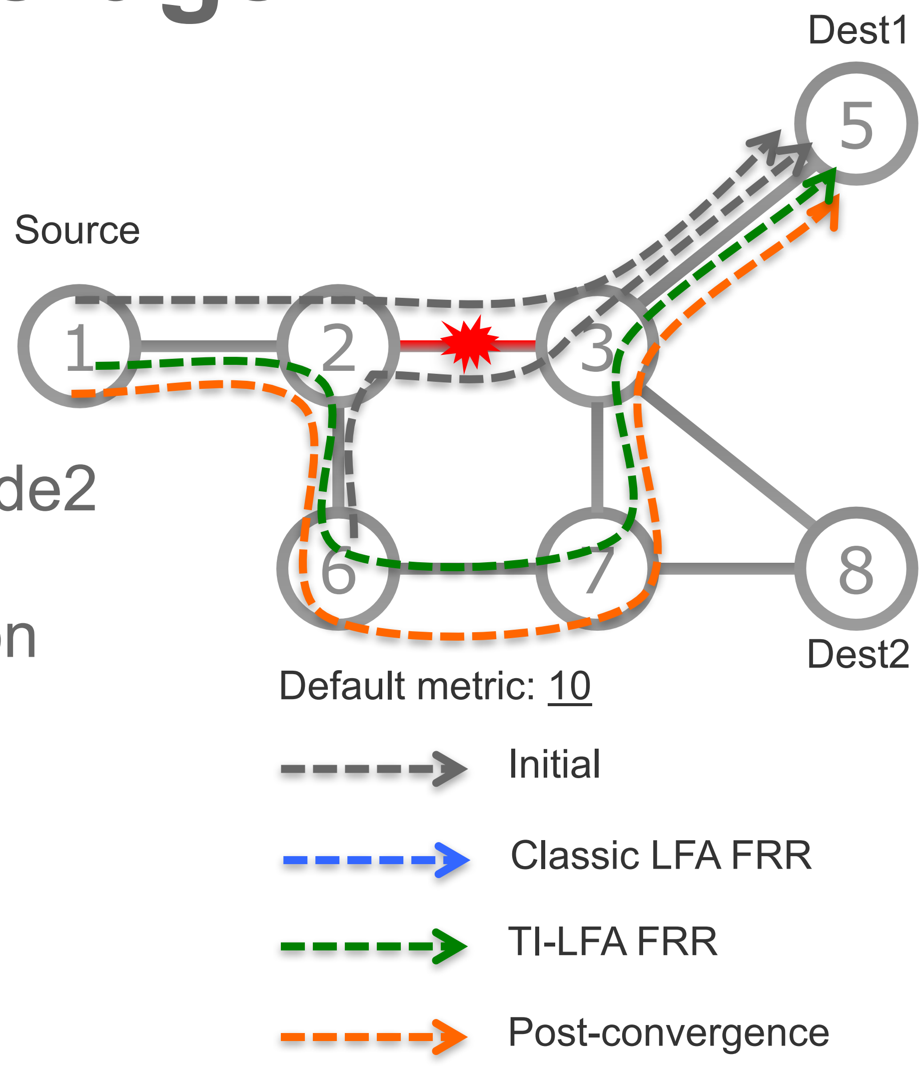


 

# Classic LFA may provide suboptimal backup path

- Classic LFA may provide a suboptimal FRR backup path:
  - This backup path may not be planned for capacity, e.g. P node 2 would use PE4 to protect a core link, while a common planning rule is to avoid using Edge nodes for transit traffic
  - Additional case specific LFA configuration would be needed to avoid selecting undesired backup paths
  - Operator would prefer to use the post-convergence path as FRR backup path, aligned with the regular IGP convergence

→ TI-LFA uses the post-convergence path as FRR backup path

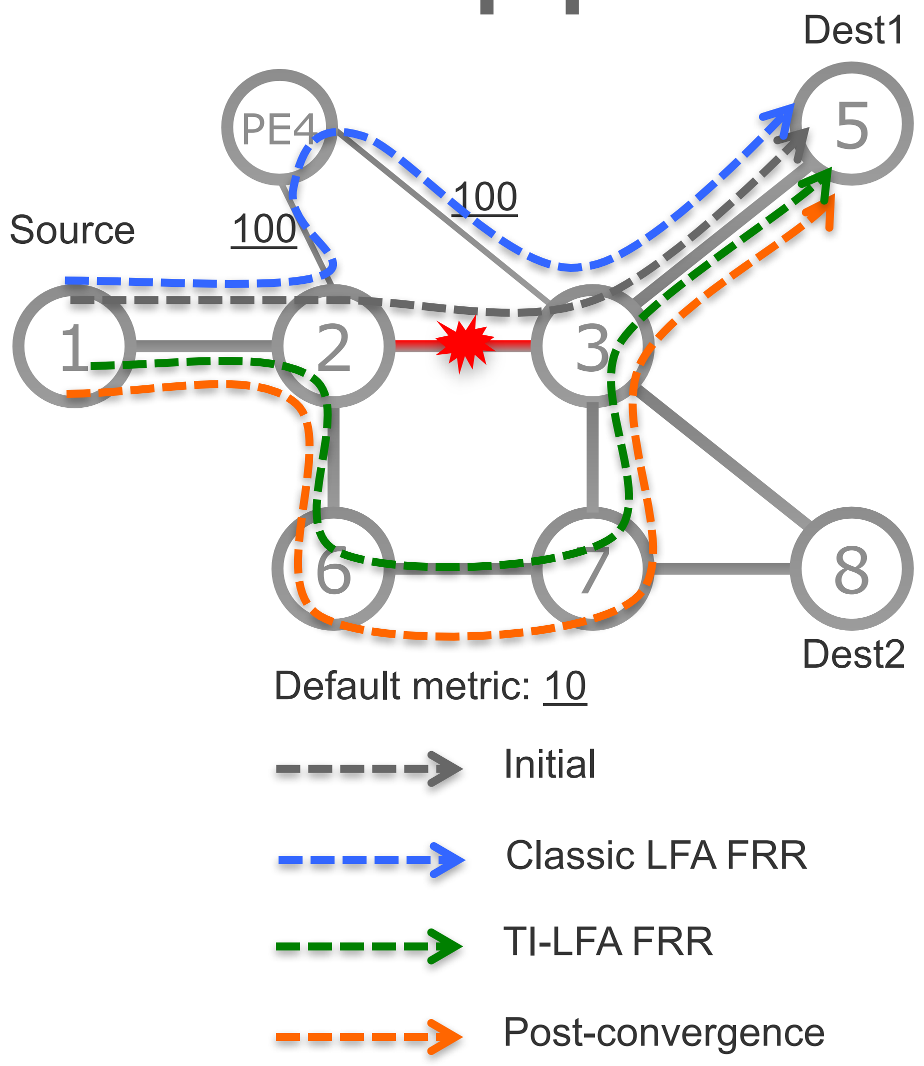

# Topology Independent LFA

# Topology Independent LFA (TI-LFA) – Benefits

- 100%-coverage 50-msec link and node protection
- Prevents transient congestion and suboptimal routing
  - leverages the post-convergence path, planned to carry the traffic
- Simple to operate and understand
  - automatically computed by the IGP
- Incremental deployment
  - also protects LDP and IP traffic

# TI-LFA uses Post-Convergence Path

- What is the more optimal and natural path upon a failure ?
  - The post-convergence path: the path that will be used after IGP has converged following a failure
  - Operator planned and dimensioned the post-convergence path to carry the traffic in the failure case
- Why have we never used it before SR ?
  - the post-convergence path may not be loop-free
- Thanks to SR, we can always use the post-convergence path
  - TI-LFA enforces the loop-free post-convergence path by encoding it as a list of segments

# TI-LFA uses Post-Convergence Path
## Optimality Benefit Example

- 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 ✗
- TI-LFA: Node2 switches all traffic destined to Node5 via high BW core links: OK! ✓

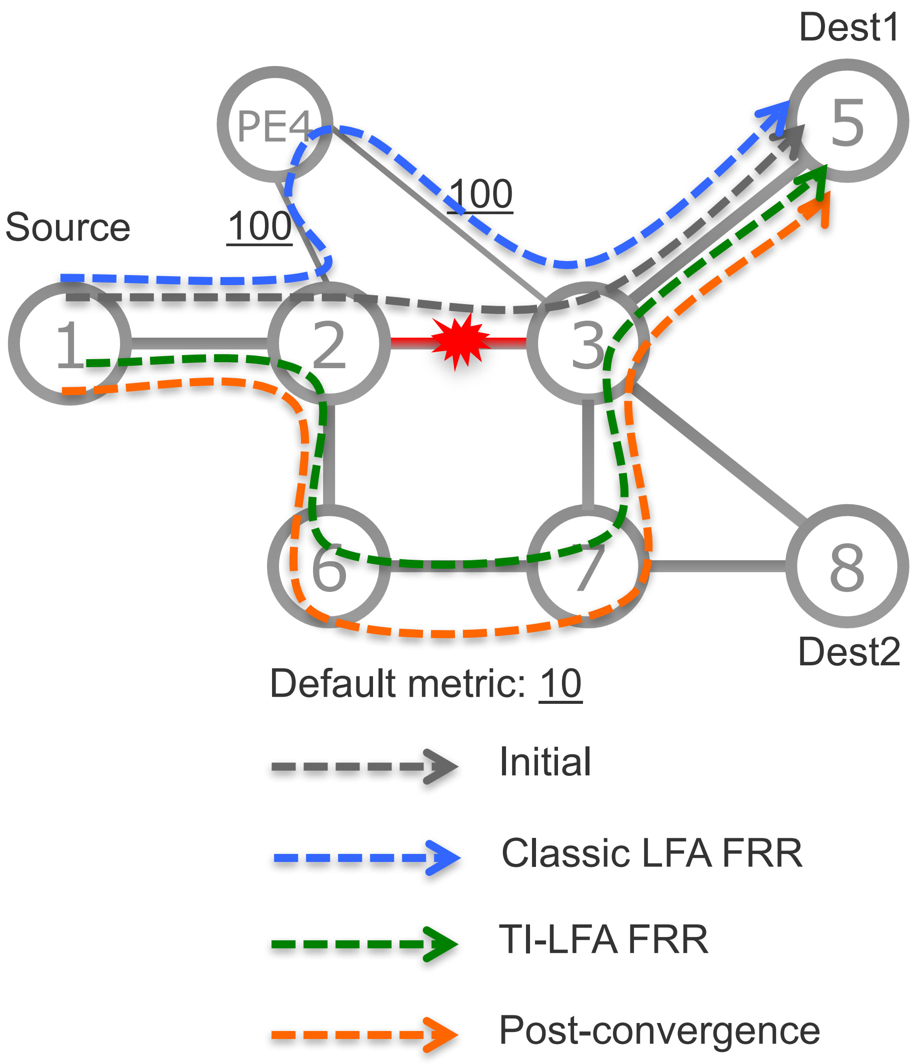


 

# Topology Independent LFA
## IOS XR 5.2.2 Implementation

- Only link protection, using at most two segments on the backup path
  - Gives 100% link protection coverage in a symmetric metric network*
- Backup path is computed by selecting PQ or P and Q nodes on the post-convergence path

* using same metric in both directions of each link

 

# Topology Independent LFA
## IOS XR 6.1.1 Implementation

- Provides 100% coverage for all failure types (link, node, SRLG) in all networks
  – Number of imposed segments only limited by the platform limit
- Generalizes the name of nodes on the backup path:
  – P-node: node reached via a Prefix-SID
  – Q-node: node reached via an Adj-SID

# Post-convergence path

- Enforcing loop-freeness on post-convergence path
  - Enforce a loop-free repair path by using a Segment List to explicitly steer traffic on this path
    - The protection path is expressed as a specific OIF plus a number of additional segments, up to the platform imposition limit

# Topology Independent LFA – High-level behavior

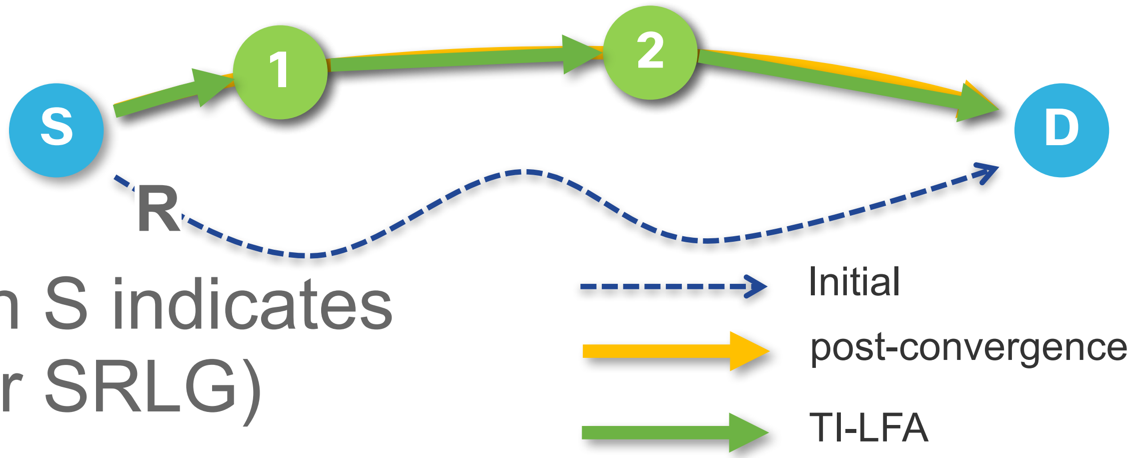


- Protect traffic from S to D
  - S has the shortest path to D
  - The outgoing interface configuration on S indicates the protected resource R (link, node, or SRLG)
  - Compute post-convergence SPT
    - >SPT calculation with protected resource R removed from topology
  - Build backup path segment list to tailor TI-LFA backup path along post-convergence path
- Repeat this for each destination D
  - Note: implementation is much more efficient than described above

 

# TI-LFA algorithm research

- The computation must be done per destination, every time the topology changes
- A lot of research has been done to ensure the post-convergence path computation and SID-list encoding is scalable
- The result of the research is applied in the implementation

# Post-Convergence Path – how many segments?

- Repair Segment List size
  - Symmetric metric network*, link protection:
    - > Proven: ≤ 2 segments
  - Asymmetric metric network, node, or SRLG protection:
    - > No theoretical bound
    - > In reality things are much simpler!

* using same metric in both directions of each link

- Orange use-case
  - 100% link protection
    - > 100% use ≤ 2 segments
  - 100% node protection (≤4 segments)
    - > 99.72% use ≤ 2 segments
    - > 0.24% use 3 segments
    - > 0.04% use 4 segments

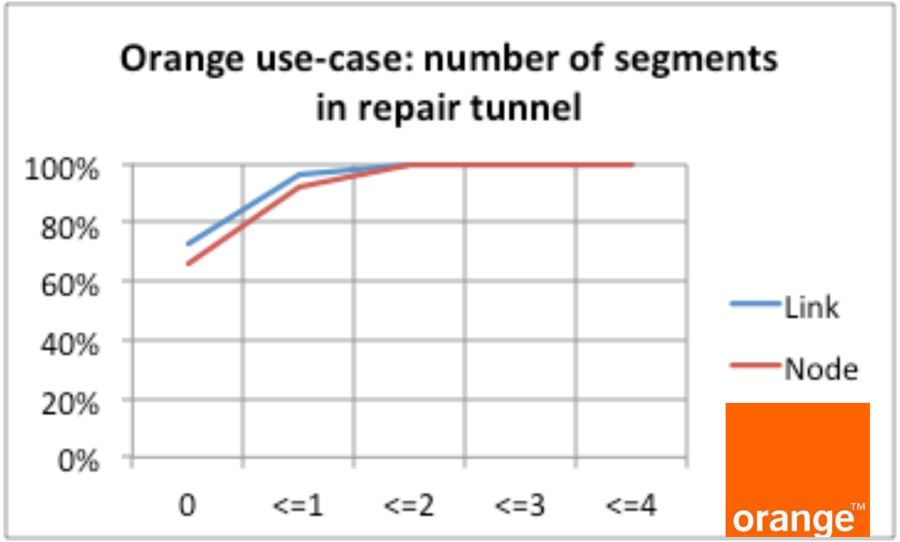


Ref. Orange @ MPLS/SDN WC 2014 – "Fast Reroute Approach Using Segment Routing"

# Do we need many SID's? No!

- TI-LFA analysis of 9 real Service Provider networks of various shapes and sizes
  - How many segments are required to steer traffic on the TI-LFA backup path for Link, Node, and SRLG protection?
  - Which percentage of destinations are protected with 0 segments, 1 segment, 2 segments, ...?
- Graphs show the percentage of destinations protected by TI-LFA using N segments for each protection type (Link, Node, SRLG)
  - Almost all destinations are protected by using 0 or 1 segment
  - Rarely 2 segments are required, and never are more than 4 segments needed

# Do we need many SID's? No!

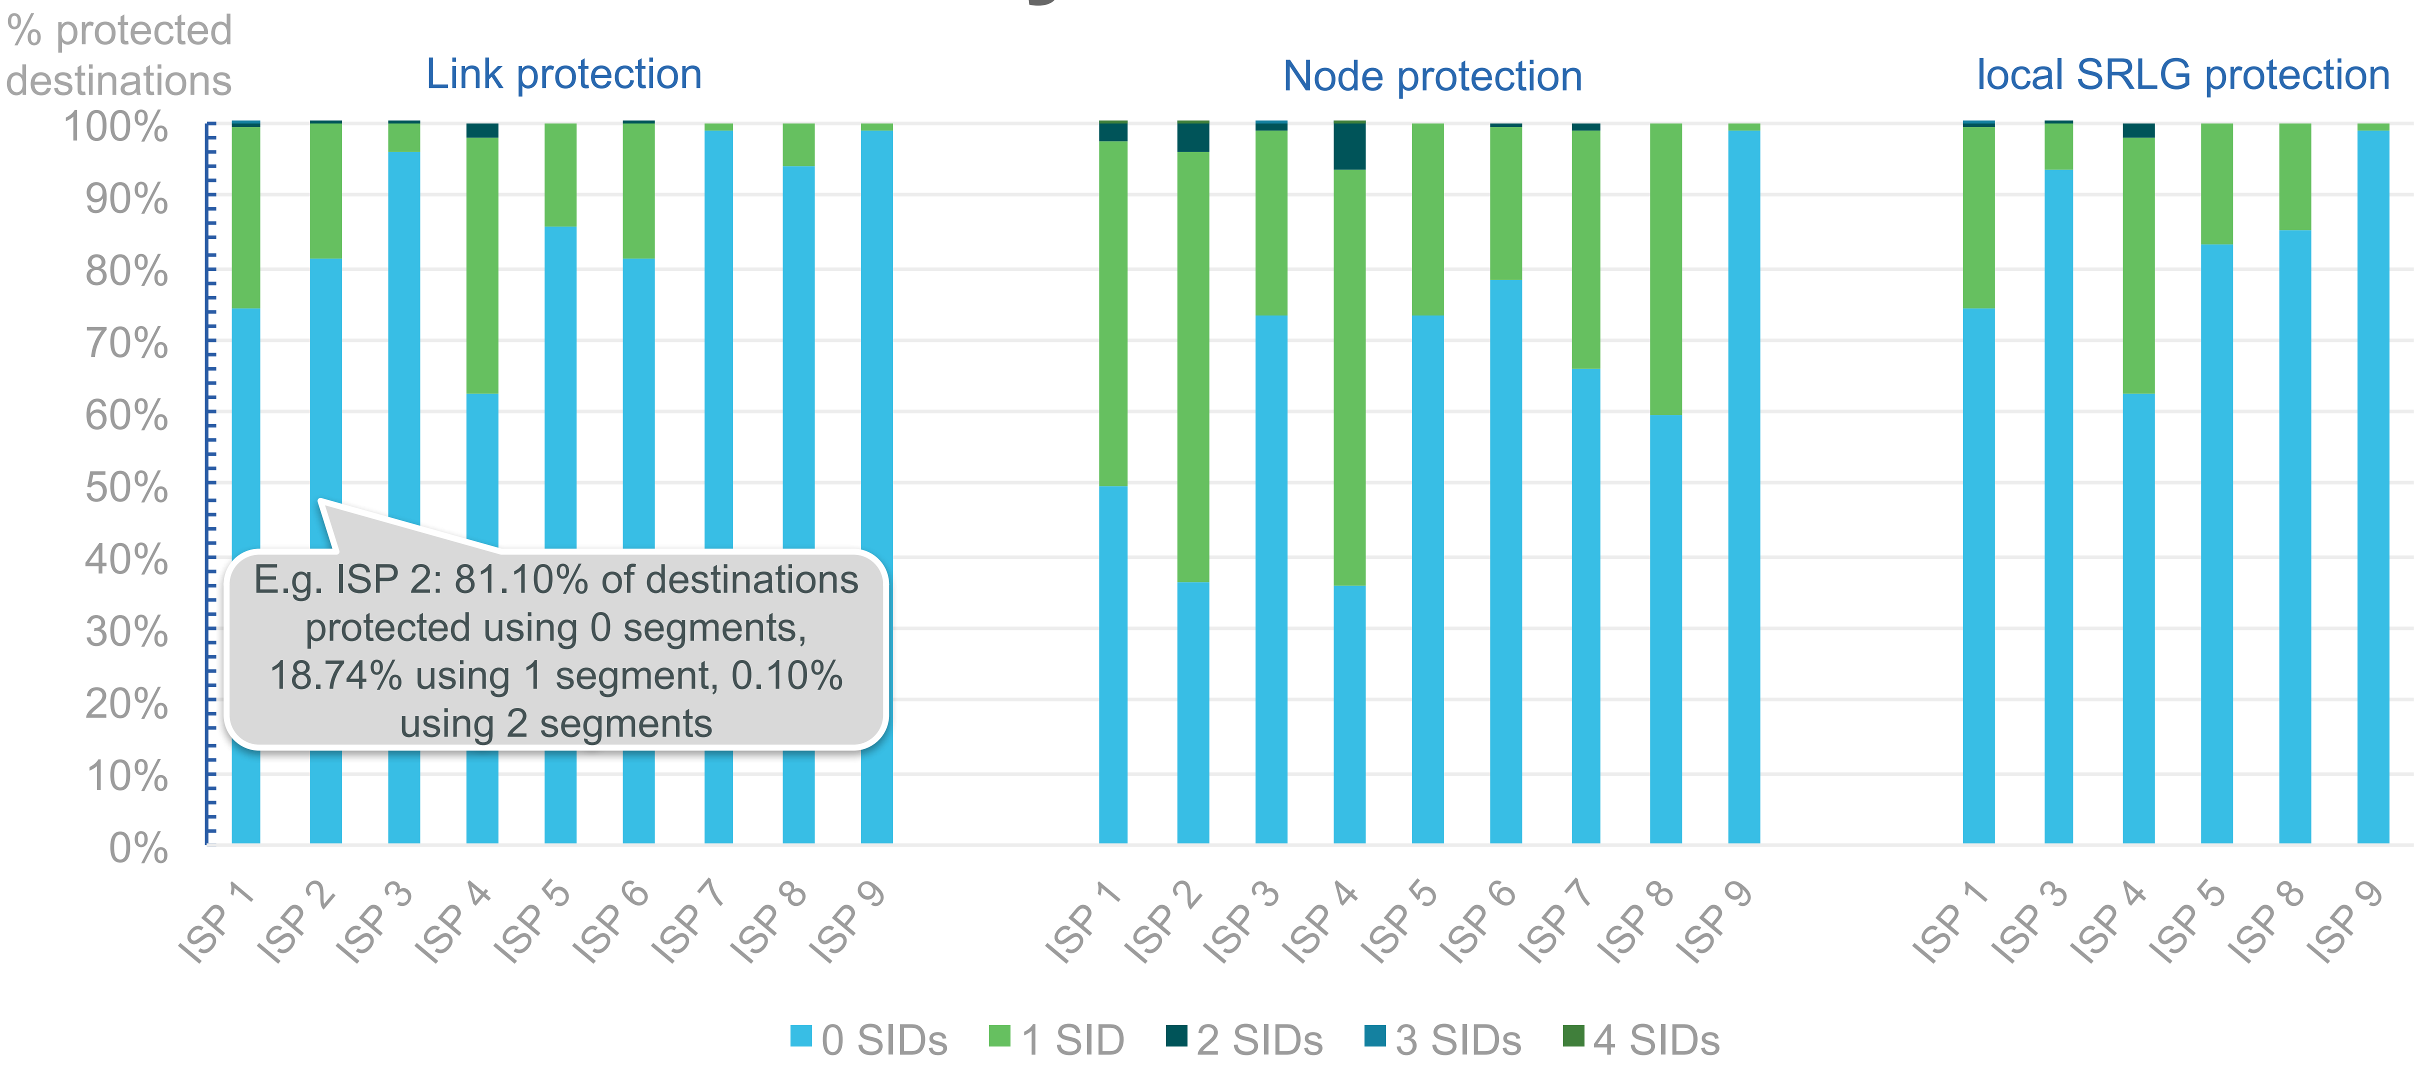

# Implementing Topology Independent LFA

# Topology Independent LFA – Configuration

- IS-IS configuration per interface

```
router isis 1
 interface GigabitEthernet0/0/0/2
  address-family ipv4|ipv6 unicast
   fast-reroute per-prefix
   fast-reroute per-prefix ti-lfa
```

Both commands are required for TI-LFA

- OSPF inheritance rules can be used
  - TI-LFA can be configured per interface, but can also be configured on instance or area level (inherited by lower levels)

```
router ospf 1
 fast-reroute per-prefix
 fast-reroute per-prefix ti-lfa
```

Both commands are required for TI-LFA

# TI-LFA for SR Traffic

# TI-LFA for SR Traffic Illustration

- This section illustrates the TI-LFA protection of Segment Routing based traffic to a Segment Routing destination
  - Assumption: Segment Routing is deployed on all nodes
- TI-LFA protection of LDP and IP traffic is covered in a following section

# Terminology – Reminder

- Terms from Remote LFA technology (RFC7490)
  - P-space(S,L): set of nodes reachable (using pre-convergence paths) from node S without using protected link L
    - Extended P-space(PLR,L): union of the P-spaces of the neighbors of PLR
  - Q-space(D, L): set of nodes that can reach (using pre-convergence paths) destination D without using protected link L

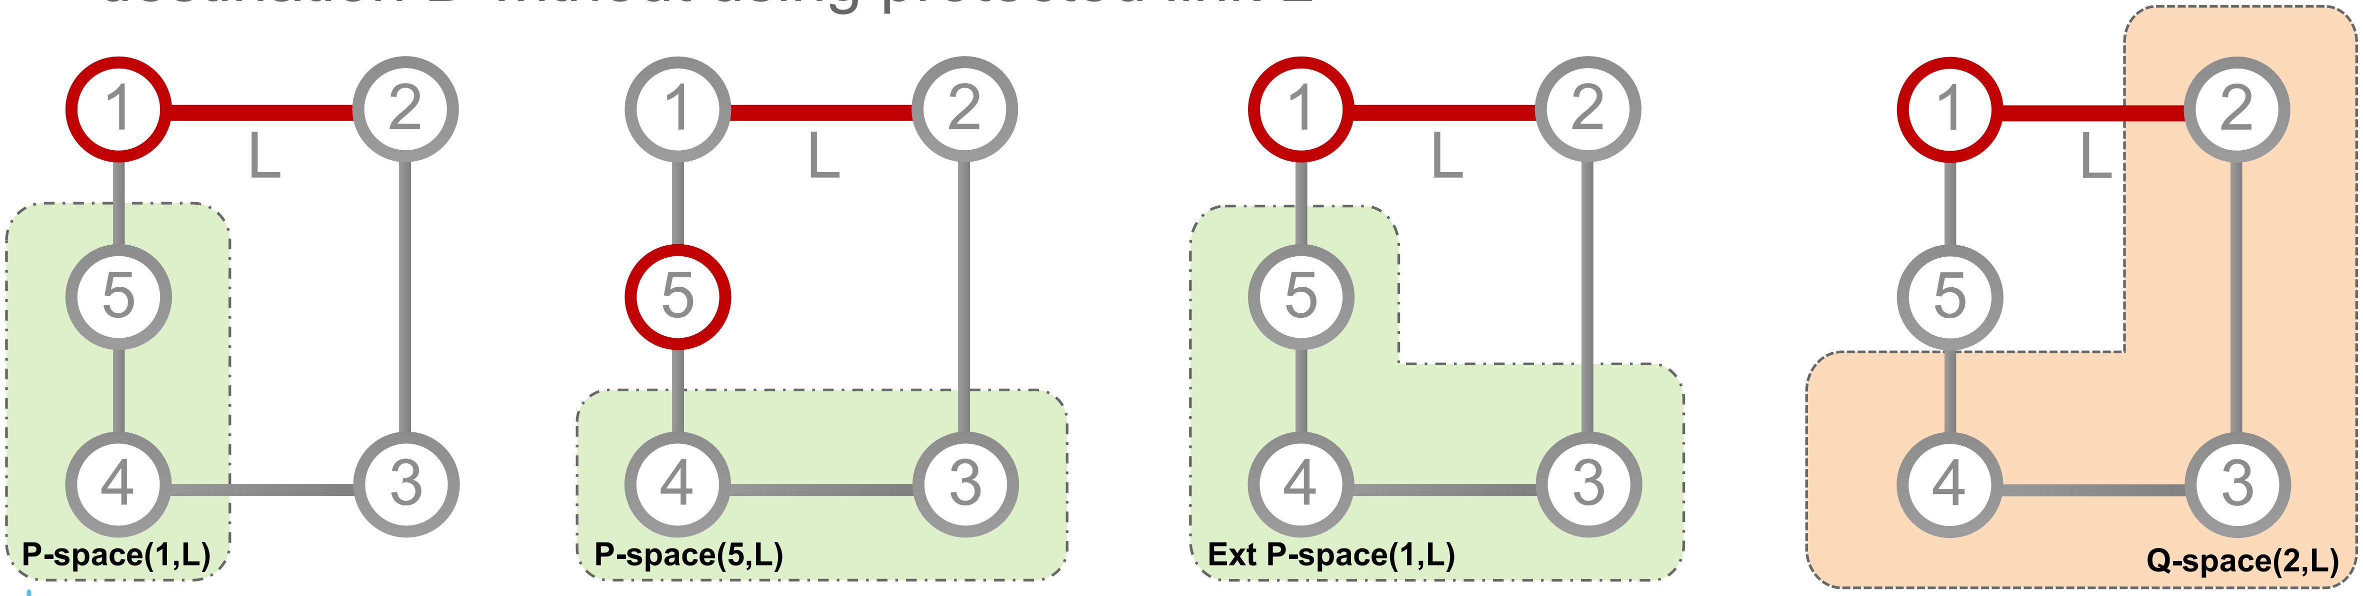


Default metric:10

# TI-LFA – zero-segment example

- For the destination Z, for the router R1, the primary link is R1R2. R1's TI-LFA computation for Z is:
  - Remove the primary link for Z (R1R2) and compute the SPF on the resulting topology. This gives us the post-convergence path from R1 to Z: <R5, R4, R3, R2>
  - R5 is in the P space (R1 can send a packet destined to R5 without any risk of having that packet flow back through the protected link R1R2)
  - R5 is in the Q space (R5 can send a packet to R2 without any risk of having this packet flow back through the protected link R1R2)
  - R5 is along the post-convergence path
  - Hence the TI-LFA backup computed by R1 for destination Z is "forward the packet to R5 without any additional segment"
- Note that this behavior is applied on a per-prefix basis and hence that for each prefix the primary link changes and the post-convergence path is computed accordingly together with the P and Q properties. The algorithm is proprietary (local behavior which is not in the scope of IETF standardization) and scales extremely well.

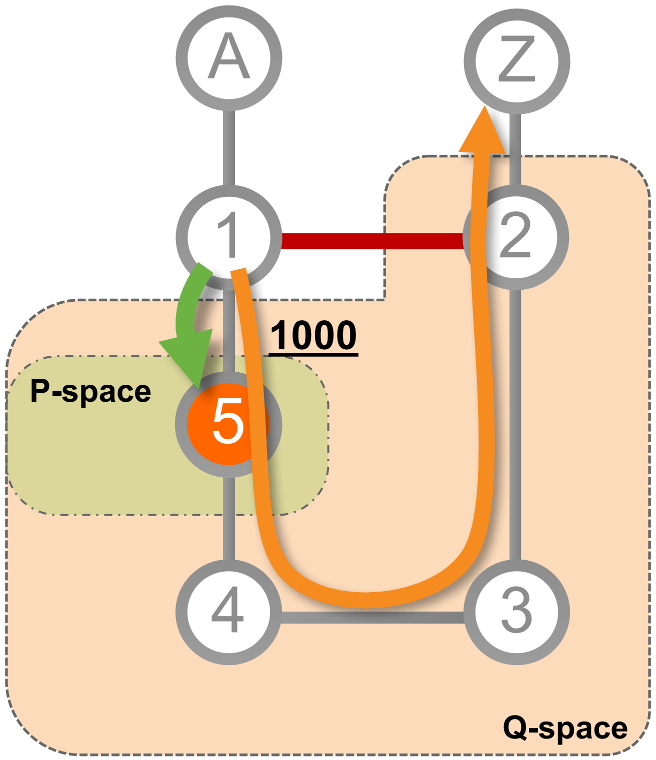


Default metric: 10

# TI-LFA – zero-segment example

- To steer packets on the TI-LFA backup path:
  “forward the packet to R5 without any additional segment”

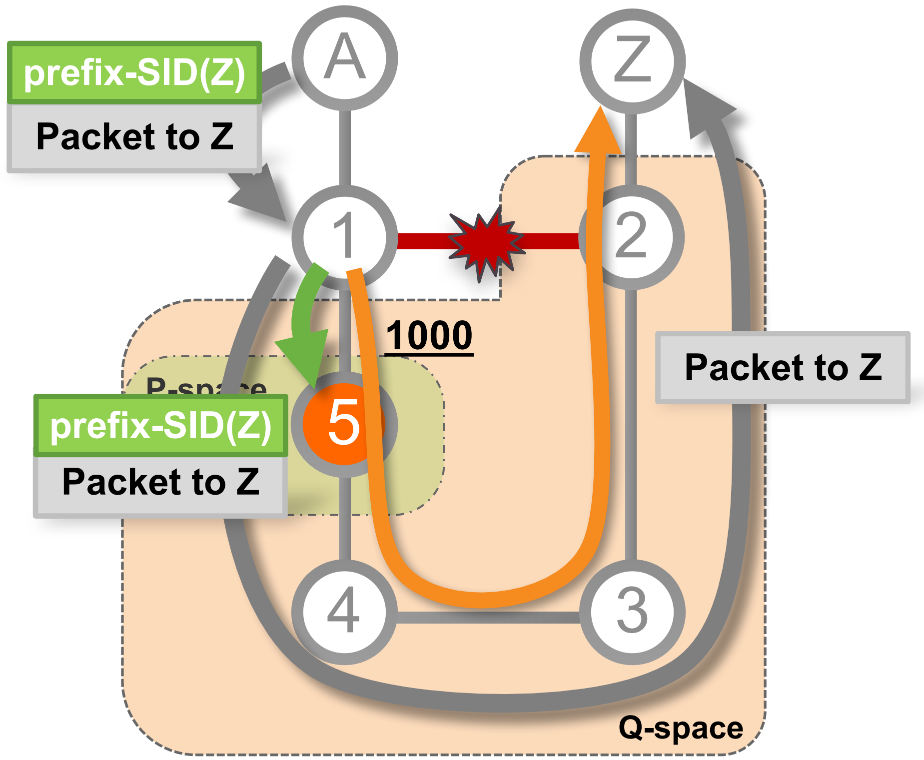


Default metric: 10

# TI-LFA – zero-segment example

- IS-IS calculated backup path

```
RP/0/0/CPU0:iosxrv-1#show isis ipv4 fast-reroute 1.1.1.6/32 detail

L2 1.1.1.6/32 [20/115] medium priority
     via 99.1.2.2, GigabitEthernet0/0/0/1, iosxrv-2, SRGB Base: 16000, Weight: 0
       FRR backup via 99.1.5.5, GigabitEthernet0/0/0/0, iosxrv-5, SRGB Base: 16000, Weight: 0
       P: No, TM: 1040, LC: No, NP: No, D: No, SRLG: Yes
     src iosxrv-6.00-00, 1.1.1.6, prefix-SID index 6, R:0 N:0 P:0 E:0 V:0 L:0
```

Zero-segment LFA

- OSPF calculated backup path

```
RP/0/0/CPU0:iosxrv-1#show ospf 1 routes 1.1.1.6/32 backup-path

Topology Table for ospf 1 with ID 1.1.1.1

Codes: O - Intra area, O IA - Inter area
       O E1 - External type 1, O E2 - External type 2
       O N1 - NSSA external type 1, O N2 - NSSA external type 2

O    1.1.1.6/32, metric 20
       99.1.2.2, from 1.1.1.6, via GigabitEthernet0/0/0/1, path-id 1
           Backup path:
              99.1.5.5, from 1.1.1.6, via GigabitEthernet0/0/0/0, protected bitmap 0000000000000001
              Attribues: Metric: 1040, Node Protect, SRLG Disjoint
```

Zero-segment LFA

# TI-LFA – zero-segment example – IP cef entry

- SR IP-to-MPLS

```
RP/0/0/CPU0:iosxrv-1#show cef 1.1.1.6/32
1.1.1.6/32, version 143, internal 0x2000001 0x3 (ptr 0xa135b274) [1], 0x0 (0xa13268e4), 0xa28
(0xa16ae158)
 Updated Jul 24 16:06:04.340
 local adjacency 99.1.2.2
 Prefix Len 32, traffic index 0, precedence n/a, priority 1
   via 99.1.5.5, GigabitEthernet0/0/0/0, 13 dependencies, weight 0, class 0, backup [flags
0x300]
    path-idx 0 NHID 0x0 [0xa104e184 0x0]
    next hop 99.1.5.5
    local adjacency
     local label 16006      labels imposed {16006}
   via 99.1.2.2, GigabitEthernet0/0/0/1, 13 dependencies, weight 0, class 0, protected [flags
0x400]
    path-idx 1 bkup-idx 0 NHID 0x0 [0xa158c310 0x0]
    next hop 99.1.2.2
     local label 16006      labels imposed {16006}
```

Prefix-SID to destination

Zero-segment LFA
IP backup path

Primary path

# TI-LFA – zero-segment example – SR label entry

```
RP/0/0/CPU0:iosxrv-1#show mpls forwarding labels 16006 detail
Local  Outgoing    Prefix             Outgoing     Next Hop        Bytes
Label  Label       or ID              Interface                    Switched
------ ----------- ------------------ ------------ --------------- ------------
16006  16006       No ID              Gi0/0/0/1    99.1.2.2        0
     Updated Jul 24 16:06:04.341
     Path Flags: 0x400 [ BKUP-IDX:0 (0xa158c310) ]
     Version: 143, Priority: 1
     MAC/Encaps: 14/18, MTU: 1500
     Label Stack (Top -> Bottom): { 16006 }
     NHID: 0
     Packets Switched: 0

       16006       No ID              Gi0/0/0/0    99.1.5.5        0
     Updated Jul 24 16:06:04.341
     Path Flags: 0x300 [ IDX:0 BKUP, NoFwd ]
     Version: 143, Priority: 1
     MAC/Encaps: 14/18, MTU: 1500
     Label Stack (Top -> Bottom): { 16006 }
     NHID: 0
     Packets Switched: 0
```

Primary path

Zero-segment LFA
SR MPLS backup path

Prefix-SID to destination

# TI-LFA – single-segment example

- For the destination Z, for the router R1, the primary link is R1R2. R1's TILFA computation for Z is:
  - Remove the primary link for Z (R1R2) and compute the SPF on the resulting topology. This gives us the post-convergence path from R1 to Z: <R5, R4, R3, R2>
  - R4 is in the P space (R1 can send a packet destined to R4 without any risk of having that packet flow back through the protected link R1R2)
  - R4 is in the Q space (R4 can send a packet to R2 without any risk of having this packet flow back through the protected link R1R2)
  - R4 is along the post-convergence path
  - Hence the TILFA backup computed by R1 for destination Z is "forward the packet on interface to R5 and push the segment R4"
- Note that this behavior is applied on a per-prefix basis and hence that for each prefix the primary link changes and the post-convergence path is computed accordingly together with the P and Q properties. The algorithm is proprietary (local behavior which is not in the scope of IETF standardization) and scales extremely well

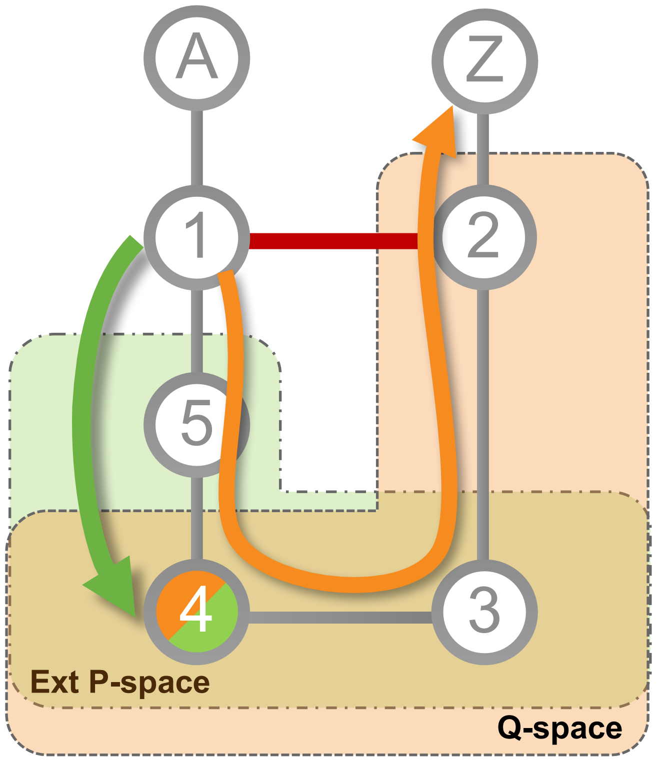


Default metric:10

# TI-LFA – single-segment example

- To steer packets on the TI-LFA backup path:
  “forward the packet on interface to R5 and push the segment <prefix-SID(R4)>”

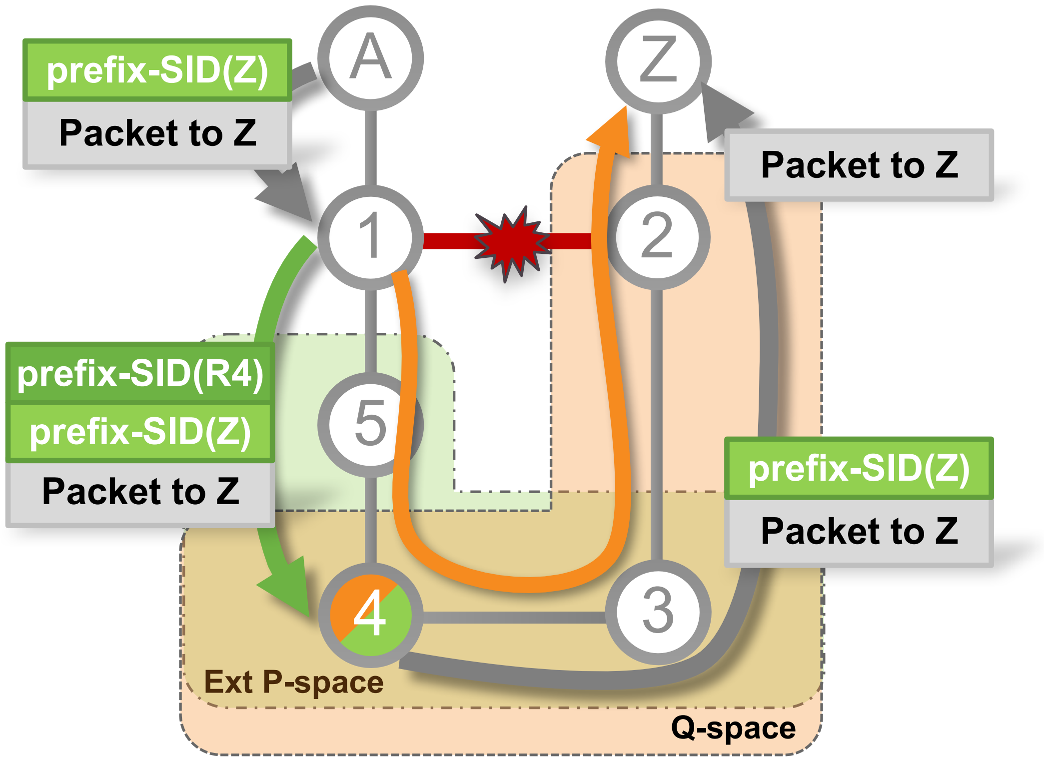


Default metric:10

# TI-LFA – single-segment example

```
RP/0/0/CPU0:iosxrv-1#show isis ipv4 fast-reroute 1.1.1.6/32 detail

L2 1.1.1.6/32 [20/115] medium priority
     via 99.1.2.2, GigabitEthernet0/0/0/1, iosxrv-2, SRGB Base: 16000, Weight: 0
       TI-LFA backup via iosxrv-4 (PQ) [1.1.1.4]
       via 99.1.5.5, GigabitEthernet0/0/0/0 iosxrv-5, SRGB Base: 16000
       Label stack [16004, 16006]
       P: No, TM: 50, LC: No, NP: No, D: No, SRLG: Yes
     src iosxrv-6.00-00, 1.1.1.6, prefix-SID index 6, R:0 N:0 P:0 E:0 V:0 L:0
```

Single-segment LFA (PQ)

Backup path label stack

```
RP/0/0/CPU0:ios#show ospf 1 routes 1.1.1.6/32 backup-path

Topology Table for ospf 1 with ID 1.1.1.1

Codes: O - Intra area, O IA - Inter area
       O E1 - External type 1, O E2 - External type 2
       O N1 - NSSA external type 1, O N2 - NSSA external type 2

O    1.1.1.6/32, metric 20
       99.1.2.2, from 1.1.1.6, via GigabitEthernet0/0/0/1, path-id 1
           Backup path: TI-LFA, P node: 1.1.1.4, Label: 16004
              99.1.5.5, from 1.1.1.6, via GigabitEthernet0/0/0/0, protected bitmap 0000000000000001
              Attribues: Metric: 50, SRLG Disjoint
```

Single-segment LFA

# TI-LFA – double-segment example

- For the destination Z, for the router R1, the primary link is R1R2. R1's TI-LFA computation for Z is:
  - Remove the primary link for Z (R1R2) and compute the SPF on the resulting topology. This gives us the post-convergence path from R1 to Z: <R5, R4, R3, R2>
  - R4 is in the P space (R1 can send a packet destined to R4 without any risk of having that packet flow back through the protected link R1R2)
  - R3 is in the Q space (R3 can send a packet to R2 without any risk of having this packet flow back through the protected link R1R2)
  - R4 and R3 are adjacent and along the post-convergence path
  - Hence the TI-LFA backup computed by R1 for destination Z is "forward the packet on interface to R5 and push the segments R4 and R4-R3"
- Note that this behavior is applied on a per-prefix basis and hence that for each prefix the primary link changes and the post-convergence path is computed accordingly together with the P and Q properties. The algorithm is proprietary (local behavior which is not in the scope of IETF standardization) and scales extremely well

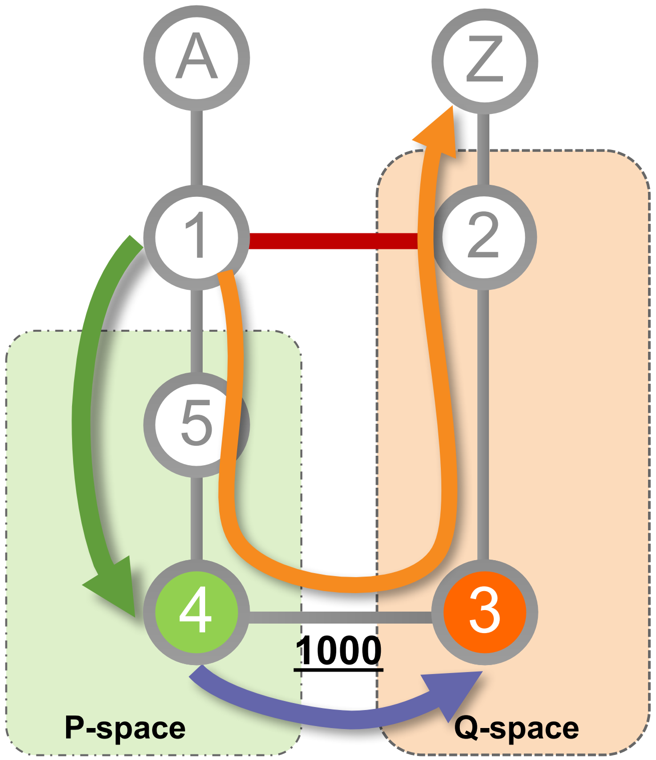


Default metric: 10

# TI-LFA – double-segment example

- To steer packets on the TI-LFA backup path: “forward the packet on interface to R5 and push the segments <prefix-SID(R4) and adj-SID(R4-R3)>”

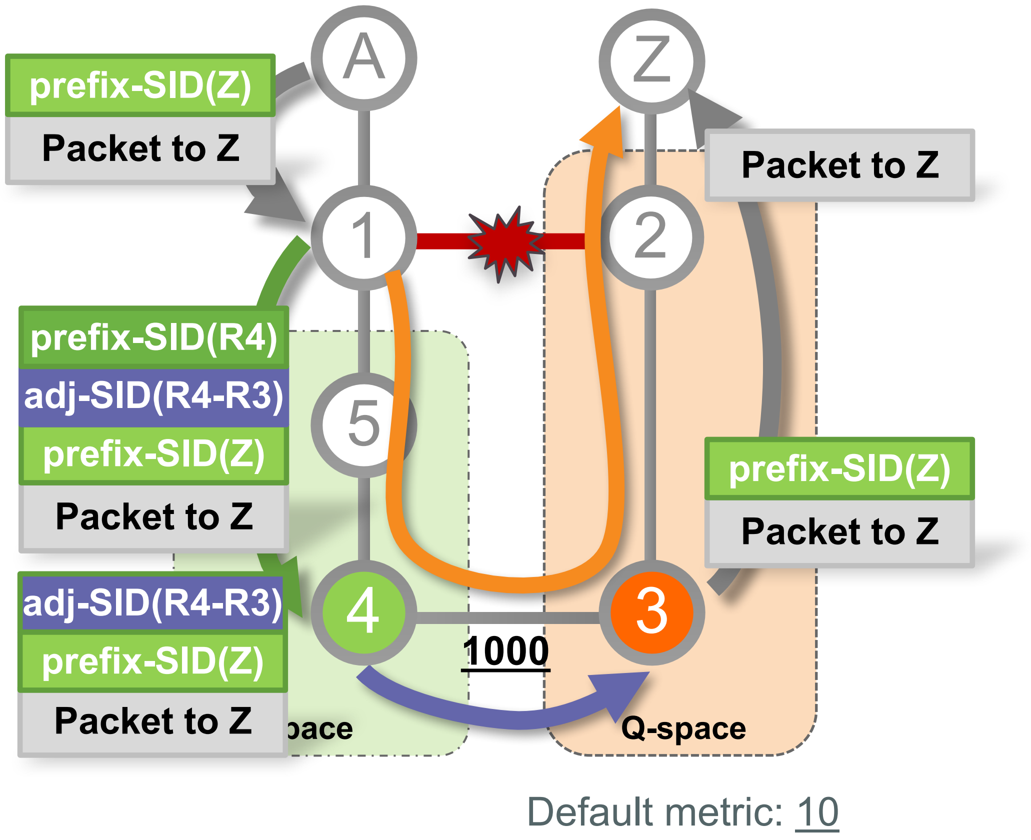

# TI-LFA – double-segment example

```
RP/0/0/CPU0:iosxrv-1#show isis ipv4 fast-reroute 1.1.1.6/32 detail

L2 1.1.1.6/32 [20/115] medium priority
     via 99.1.2.2, GigabitEthernet0/0/0/1, iosxrv-2, SRGB Base: 16000, Weight: 0
       TI-LFA backup via iosxrv-4 (P) [1.1.1.4], iosxrv-3 (Q) [1.1.1.3]
       via 99.1.5.5, GigabitEthernet0/0/0/0 iosxrv-5, SRGB Base: 16000
       Label stack [16004, 24000, 16006]
       P: No, TM: 1040, LC: No, NP: No, D: No, SRLG: Yes
     src iosxrv-6.00-00, 1.1.1.6, prefix-SID index 6, R:0 N:0 P:0 E:0 V:0 L:0
```

Double-segment LFA (P and Q)

Backup path label stack

```
RP/0/0/CPU0:ios#show ospf 1 routes 1.1.1.6/32 backup-path

Topology Table for ospf 1 with ID 1.1.1.1

Codes: O - Intra area, O IA - Inter area
       O E1 - External type 1, O E2 - External type 2
       O N1 - NSSA external type 1, O N2 - NSSA external type 2

O    1.1.1.6/32, metric 20
       99.1.2.2, from 1.1.1.6, via GigabitEthernet0/0/0/1, path-id 1
           Backup path: TI-LFA, P node: 1.1.1.4, Label: 16004, Q node: 1.1.1.3, Label: 24000
              99.1.5.5, from 1.1.1.6, via GigabitEthernet0/0/0/0, protected bitmap 0000000000000001
              Attribues: Metric: 1040, SRLG Disjoint
```

Double-segment LFA

 

# TI-LFA – double-segment example – IP cef entry

- SR IP-to-MPLS

```
RP/0/0/CPU0:iosxrv-1#show cef 1.1.1.6/32
1.1.1.6/32, version 236, internal 0x2000001 0x3 (ptr 0xa135b274) [1], 0x0 (0xa13266a4), 0xa28
(0xa16ae234)
 Updated Jul 24 16:06:04.341
 local adjacency 99.1.2.2
 Prefix Len 32, traffic index 0, precedence n/a, priority 1
   via 99.1.5.5, GigabitEthernet0/0/0/0, 19 dependencies, weight 0, class 0, backup (remote)
[flags 0x8300]
    path-idx 0 NHID 0x0 [0xa104e184 0x0]
    next hop 99.1.5.5, P-node 1.1.1.4, Q-node 1.1.1.3
    local adjacency
     local label 16006      labels imposed {16004 24000 16006}
   via 99.1.2.2, GigabitEthernet0/0/0/1, 19 dependencies, weight 0, class 0, protected [flags
0x400]
    path-idx 1 bkup-idx 0 NHID 0x0 [0xa158c310 0x0]
    next hop 99.1.2.2
     local label 16006      labels imposed {16006}
```

Prefix-SID to P node

Adj-SID from P to Q

Prefix-SID to destination

Double-segment LFA IP backup path

Primary path

# TI-LFA – double-segment example – SR label entry

```
RP/0/0/CPU0:iosxrv-1#show mpls forwarding labels 16006 detail
Local  Outgoing    Prefix             Outgoing     Next Hop        Bytes
Label  Label       or ID              Interface                    Switched
------ ----------- ------------------ ------------ --------------- ------------
16006  16006       No ID              Gi0/0/0/1    99.1.2.2        0
     Updated Jul 24 16:06:04.340
     Path Flags: 0x400 [ BKUP-IDX:0 (0xa158c310) ]
     Version: 236, Priority: 1
     MAC/Encaps: 14/18, MTU: 1500
     Label Stack (Top -> Bottom): { 16006 }
     NHID: 0
     Packets Switched: 0

       16004       No ID              Gi0/0/0/0    99.1.5.5        0
     Updated Jul 24 16:06:04.340
     Path Flags: 0x8300 [ IDX:0 BKUP, NoFwd ]
     Version: 236, Priority: 1
     MAC/Encaps: 14/26, MTU: 1500
     Label Stack (Top -> Bottom): { 16004 24000 16006 }
     NHID: 0
     Packets Switched: 0
```

# TI-LFA – Load-balancing on backup path

- TI-LFA makes use of the ECMP paths towards P or PQ node and from Q or PQ node to the destination
  - Natural property of prefix-SID
- Statistical load-balancing between equal cost P or PQ nodes
  - Per-destination selection of P or PQ node, based on hash function
- For disjoint P and Q nodes, statistical load-balancing between equal cost Q nodes
  - Per-destination selection of Q node, based on hash function

# TI-LFA – Load-balancing on backup path – Example

- Node2 provides TI-LFA protection for destinations 10 and 11 in case their primary link 2-6 fails
- TI-LFA calculates the same PQ Node4 for both destinations 10 and 11
- Two ECMP paths are available from PQ Node4 to destinations 10 and 11
  Traffic to both destinations is load-balanced over these paths

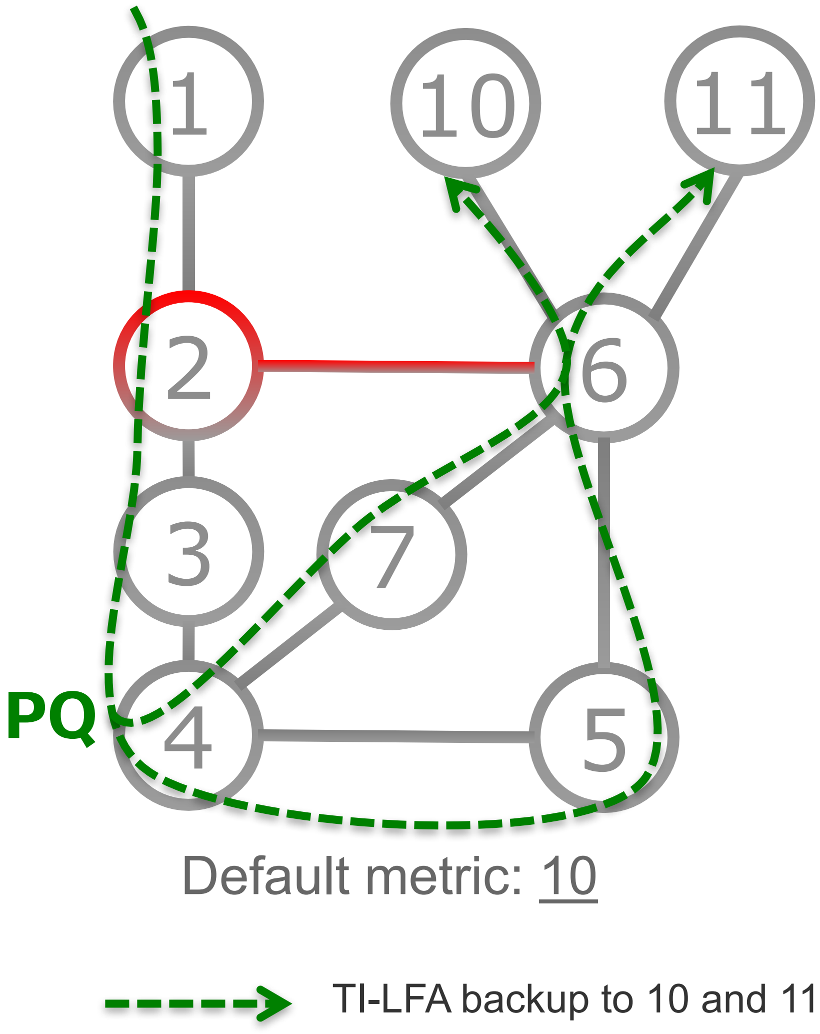

# TI-LFA – Load-balancing on backup path – Example

- Similar topology to previous slide, but links 4-7 and 4-5 have a high metric
- Node2 provides TI-LFA protection for destinations 10 and 11 in case their primary link 2-6 fails
- TI-LFA calculates 2 equally good pairs of P and Q nodes for both destinations 10 and 11:
  - (P:4, Q:5) and (P:4, Q:7)
- TI-LFA statistically load-balances (per destination) over these pairs, e.g.:
  - Destination 10, use (P:4, Q:7)
  - Destination 11, use (P:4, Q:5)

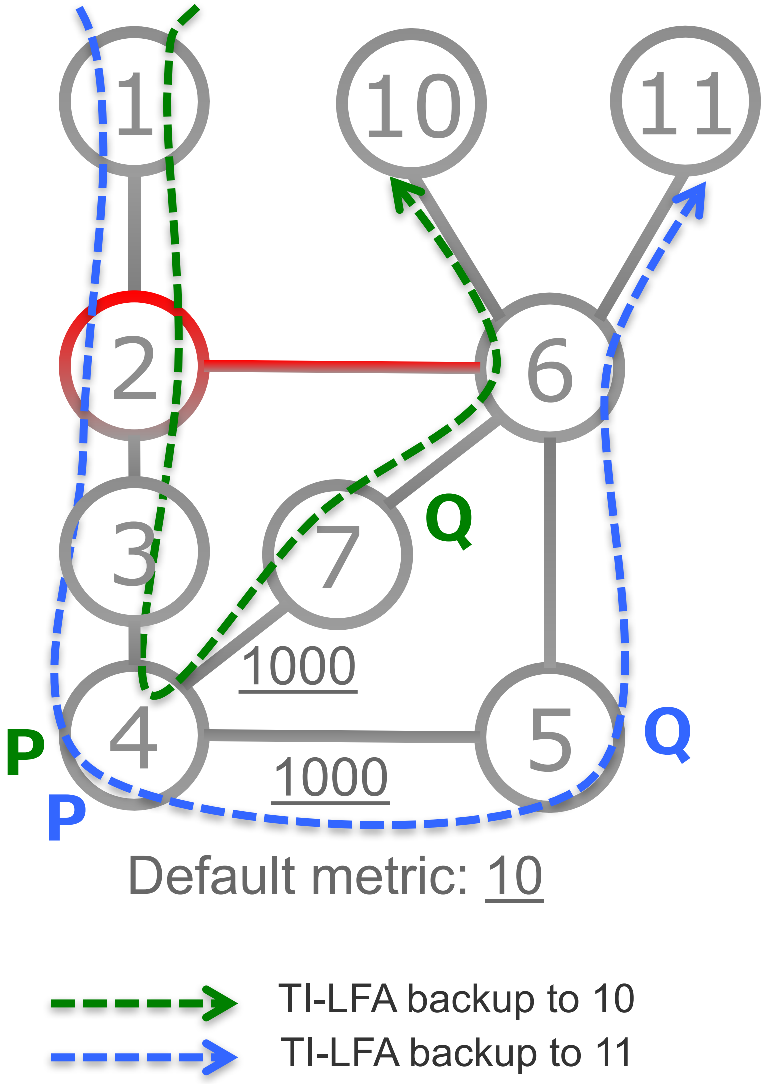

# TI-LFA for LDP Traffic

# Protecting LDP traffic with TI-LFA

- LDP traffic is protected with TI-LFA
  - 100% coverage
  - No targeted LDP sessions
  - Incremental deployment
    - SR can be deployed on islands within LDP network
  - Protection will leverage SR/LDP interworking functionality
    - FIB uses label merging functionality for backup path
  - Mapping server must advertise prefix-to-SID mappings for protected prefixes

# TI-LFA for LDP Traffic – Illustration

- The following slides illustrate the TI-LFA protection of LDP based traffic to a LDP only destination
  - Assumption: LDP and Segment Routing is deployed on all nodes of the TI-LFA backup path, except source and destination nodes
    - Other cases are covered after this
  - A prefix-SID for the LDP-only destination is advertised by the Mapping Server

# TI-LFA for LDP – zero-segment example

- Equivalent to the zero-segment case in the “TI-LFA for SR traffic” slides
  - Using LDP labels instead of SR labels
    - > LDP(X, Y) = LDP label advertised by NodeX for prefix Y
- To steer packets on the TI-LFA backup path: “forward the packet to R5 without any additional segment”

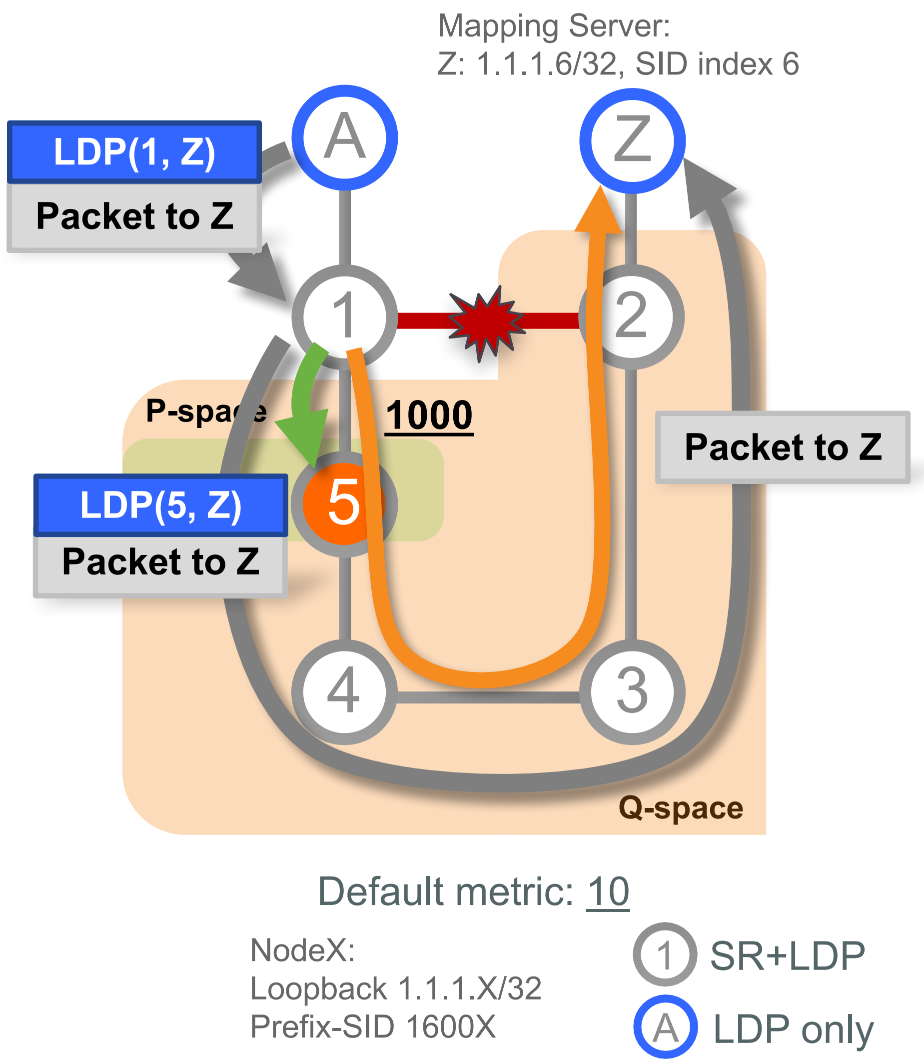

# TI-LFA for LDP – zero-segment example

- IS-IS calculated backup path

```
RP/0/0/CPU0:iosxrv-1#show isis ipv4 fast-reroute 1.1.1.6/32 detail

L2 1.1.1.6/32 [20/115] medium priority
     via 99.1.2.2, GigabitEthernet0/0/0/1, iosxrv-2, SRGB Base: 16000, Weight: 0
       FRR backup via 99.1.5.5, GigabitEthernet0/0/0/0, iosxrv-5, SRGB Base: 16000, Weight: 0
       P: No, TM: 1040, LC: No, NP: No, D: No, SRLG: Yes
     src iosxrv-6.00-00, 1.1.1.6
```

Zero-segment LFA

- OSPF calculated backup path

```
RP/0/0/CPU0:iosxrv-1#show ospf 1 routes 1.1.1.6/32 backup-path

Topology Table for ospf 1 with ID 1.1.1.1

Codes: O - Intra area, O IA - Inter area
       O E1 - External type 1, O E2 - External type 2
       O N1 - NSSA external type 1, O N2 - NSSA external type 2

O    1.1.1.6/32, metric 20
       99.1.2.2, from 1.1.1.6, via GigabitEthernet0/0/0/1, path-id 1
           Backup path:
              99.1.5.5, from 1.1.1.6, via GigabitEthernet0/0/0/0, protected bitmap 0000000000000001
              Attribues: Metric: 1040, Node Protect, SRLG Disjoint
```

Zero-segment LFA

# TI-LFA For LDP – zero-segment example – IP cef entry

- LDP IP-to-MPLS preferred (default)

```
RP/0/0/CPU0:iosxrv-1#show mpls ldp bindings 1.1.1.6/32
1.1.1.6/32, rev 24
        Local binding: label: 24008
        Remote bindings: (2 peers)
            Peer                Label
            -----------------   ---------
            1.1.1.2:0           24005
            1.1.1.5:0           24006
RP/0/0/CPU0:iosxrv-1#show cef 1.1.1.6/32
1.1.1.6/32, version 61, internal 0x2000001 0x1 (ptr 0xa135b274) [1], 0x0 (0xa1326b24), 0xa28
(0xa16ae12c)
 Updated Sep 20 21:01:03.119
 local adjacency 99.1.2.2
 Prefix Len 32, traffic index 0, precedence n/a, priority 3
   via 99.1.5.5, GigabitEthernet0/0/0/0, 14 dependencies, weight 0, class 0, backup [flags
0x300]
    path-idx 0 NHID 0x0 [0xa104e184 0x0]
    next hop 99.1.5.5
    local adjacency
     local label 24008      labels imposed {24006}
   via 99.1.2.2, GigabitEthernet0/0/0/1, 14 dependencies, weight 0, class 0, protected [flags
0x400]
    path-idx 1 bkup-idx 0 NHID 0x0 [0xa158c310 0x0]
    next hop 99.1.2.2
     local label 24008      labels imposed {24005}
```

LDP allocated local label 24008

LDP label to destination via 5

LDP label to destination

Zero-segment LFA IP backup path

Primary path

# TI-LFA for LDP – zero-segment example – LDP label entry

```
RP/0/0/CPU0:iosxrv-1#show mpls ldp bindings 1.1.1.6/32
1.1.1.6/32, rev 24
        Local binding: label: 24008
        Remote bindings: (2 peers)
            Peer                Label
            -----------------   ---------
            1.1.1.2:0           24005
            1.1.1.5:0           24006
RP/0/0/CPU0:iosxrv-1#show mpls forwarding labels 24008 detail
Local  Outgoing    Prefix             Outgoing     Next Hop        Bytes
Label  Label       or ID              Interface                    Switched
------ ----------- ------------------ ------------ --------------- ------------
24008  24005       1.1.1.6/32         Gi0/0/0/1    99.1.2.2        0
     Updated Jul 24 16:06:04.341
     Path Flags: 0x400 [ BKUP-IDX:0 (0xa158c310) ]
     Version: 57, Priority: 15
     MAC/Encaps: 14/18, MTU: 1500
     Label Stack (Top -> Bottom): { 24005 }
     NHID: 0
     Packets Switched: 0

       24006       1.1.1.6/32         Gi0/0/0/0    99.1.5.5        0
     Updated Jul 24 16:06:04.341
     Path Flags: 0x300 [ IDX:0 BKUP, NoFwd ]
     Version: 57, Priority: 15
     MAC/Encaps: 14/18, MTU: 1500
     Label Stack (Top -> Bottom): { 24006 }
     NHID: 0
     Packets Switched: 0
```

LDP allocated local label 24008

LDP label to destination via 5

Primary path

Zero-segment LFA LDP MPLS backup path

LDP label to destination

# TI-LFA for LDP – single-segment example

- Equivalent to the single-segment case in the “TI-LFA for SR traffic” slides
  - Using LDP labels instead of SR labels where possible
    - LDP(X, Y) = LDP label advertised by NodeX for prefix Y
- To steer packets on the TI-LFA backup path:
  “swap LDP(1, Z) with prefix-SID(Z), forward the packet on interface to R5 and push the label <LDP(5, 4)>”

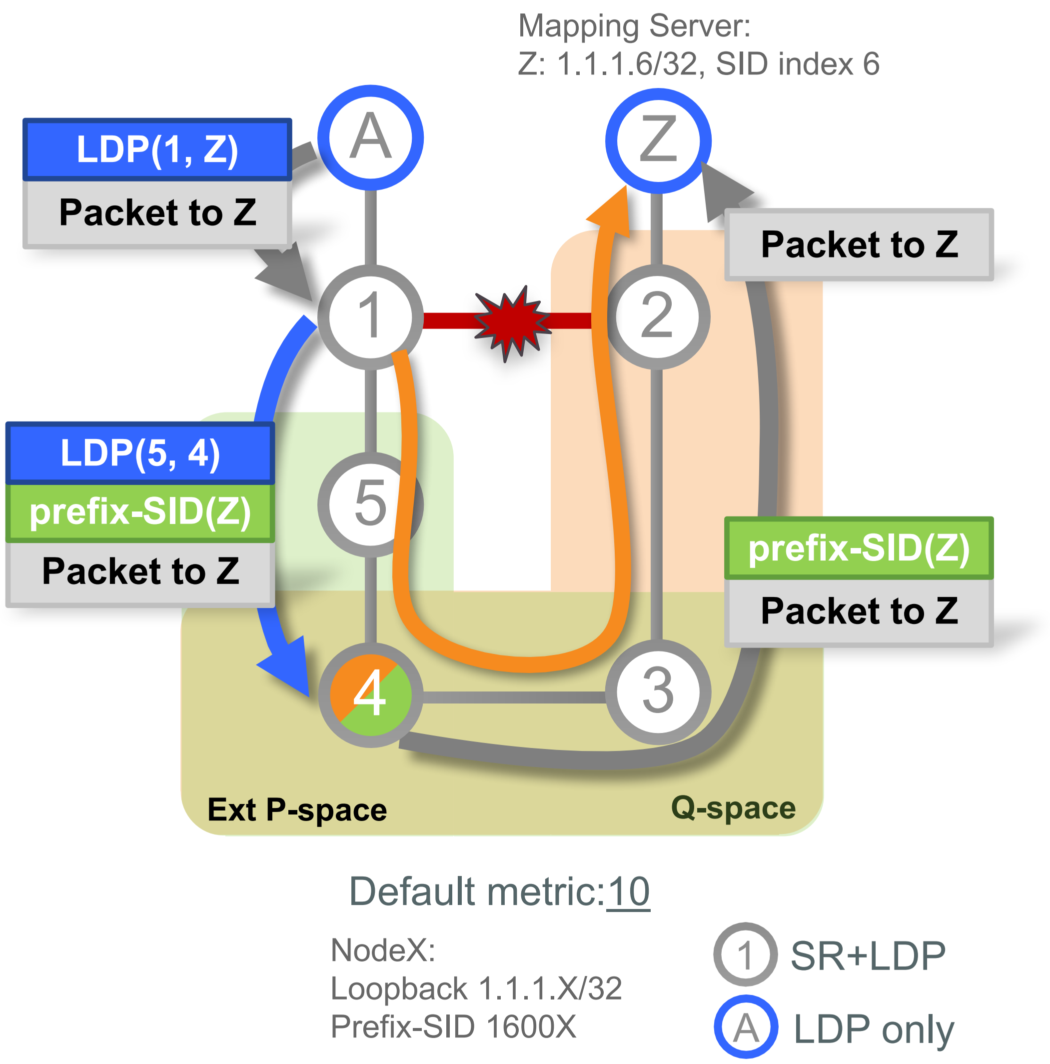




 

# TI-LFA for LDP – single-segment example

```
RP/0/0/CPU0:iosxrv-1#show isis ipv4 fast-reroute 1.1.1.6/32 detail

L2 1.1.1.6/32 [20/115] medium priority
     via 99.1.2.2, GigabitEthernet0/0/0/1, iosxrv-2, SRGB Base: 16000, Weight: 0
       TI-LFA backup via iosxrv-4 (PQ) [1.1.1.4]
       via 99.1.5.5, GigabitEthernet0/0/0/0 iosxrv-5, SRGB Base: 16000
       Label stack [16004, 16006]
       P: No, TM: 50, LC: No, NP: No, D: No, SRLG: Yes
     src iosxrv-6.00-00, 1.1.1.6
```

Single-segment LFA (PQ)

Backup path label stack

```
RP/0/0/CPU0:ios#show ospf 1 routes 1.1.1.6/32 backup-path

Topology Table for ospf 1 with ID 1.1.1.1

Codes: O - Intra area, O IA - Inter area
       O E1 - External type 1, O E2 - External type 2
       O N1 - NSSA external type 1, O N2 - NSSA external type 2

O    1.1.1.6/32, metric 20
       99.1.2.2, from 1.1.1.6, via GigabitEthernet0/0/0/1, path-id 1
           Backup path: TI-LFA, P node: 1.1.1.4, Label: 16004
              99.1.5.5, from 1.1.1.6, via GigabitEthernet0/0/0/0, protected bitmap 0000000000000001
              Attribues: Metric: 50, SRLG Disjoint
```

Single-segment LFA

# TI-LFA for LDP – single-segment example – IP cef entry

- LDP IP-to-MPLS preferred (default)

```
RP/0/0/CPU0:iosxrv-1#show cef 1.1.1.6/32
1.1.1.6/32, version 31, internal 0x2000001 0x5 (ptr 0xa135b274) [1], 0x0 (0xa1326ab8), 0xa28
(0xa16ae1dc)
 Updated Jul 24 16:06:04.340
 local adjacency 99.1.2.2
 Prefix Len 32, traffic index 0, precedence n/a, priority 15
   via 99.1.5.5, GigabitEthernet0/0/0/0, 21 dependencies, weight 0, class 0, backup [flags
0x300]
    path-idx 0 NHID 0x0 [0xa104e184 0x0]
    next hop 99.1.5.5, PQ-node 1.1.1.4
    local adjacency
     local label 24008      labels imposed {24007 16006}
   via 99.1.2.2, GigabitEthernet0/0/0/1, 21 dependencies, weight 0, class 0, protected [flags
0x400]
    path-idx 1 bkup-idx 0 NHID 0x0 [0xa158c310 0x0]
    next hop 99.1.2.2
     local label 24008      labels imposed {24005}
```

LDP label to PQ node

Prefix-SID to destination

Single-segment LFA
IP backup path

Primary path

# TI-LFA for LDP – single-segment example – LDP label entry

```
RP/0/0/CPU0:xrvr-1#show mpls ldp bindings 1.1.1.6/32 local
1.1.1.6/32, rev 24
        Local binding: label: 24008
RP/0/0/CPU0:iosxrv-1#show mpls forwarding labels 24008 detail
Local  Outgoing    Prefix             Outgoing     Next Hop        Bytes
Label  Label       or ID              Interface                    Switched
------ ----------- ------------------ ------------ --------------- ------------
24008  24005       1.1.1.6/32         Gi0/0/0/1    99.1.2.2        0
     Updated Jul 24 16:06:04.340
     Path Flags: 0x400 [ BKUP-IDX:0 (0xa158c310) ]
     Version: 31, Priority: 15
     MAC/Encaps: 14/18, MTU: 1500
     Label Stack (Top -> Bottom): { 24005 }
     NHID: 0
     Packets Switched: 0

       24007       1.1.1.6/32         Gi0/0/0/0    99.1.5.5        0
     Updated Jul 24 16:06:04.340
     Path Flags: 0x300 [ IDX:0 BKUP, NoFwd ]
     Version: 31, Priority: 15
     MAC/Encaps: 14/22, MTU: 1500
     Label Stack (Top -> Bottom): { 24007 16006 }
     NHID: 0
     Packets Switched: 0
```

# TI-LFA for LDP – double-segment example

- Equivalent to the single-segment case in the TI-LFA for SR traffic slides
  - Using LDP labels instead of SR labels where possible
    - > LDP(X, Y) = LDP label advertised by NodeX for prefix Y
- To steer packets on the TI-LFA backup path:
  "swap LDP(1, Z) with prefix-SID(Z), forward the packet on interface to R5 and push the labels <LDP(5, 4) and adj-SID(R4-R3)>"

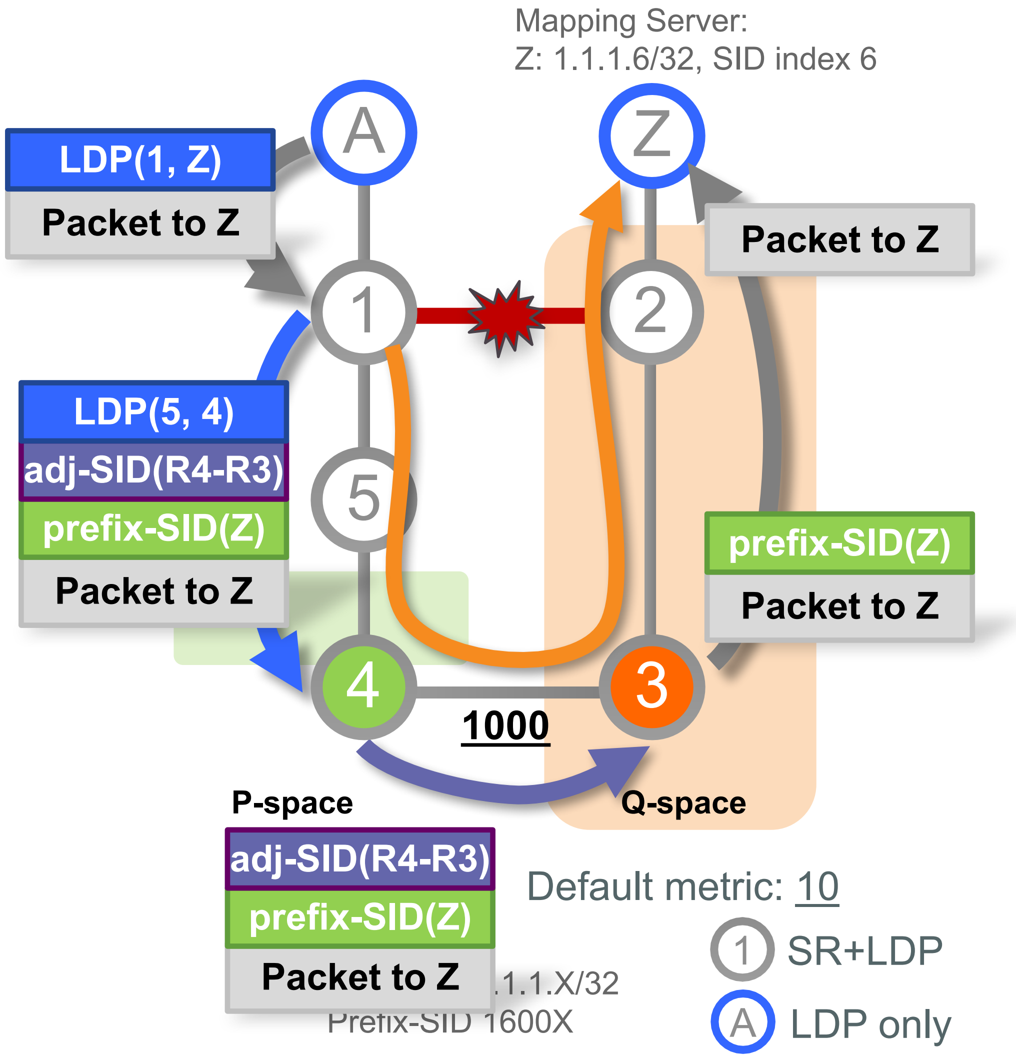


 

# TI-LFA for LDP – double-segment example

```
RP/0/0/CPU0:iosxrv-1#show isis ipv4 fast-reroute 1.1.1.6/32 detail

L2 1.1.1.6/32 [20/115] medium priority
     via 99.1.2.2, GigabitEthernet0/0/0/1, iosxrv-2, SRGB Base: 16000, Weight: 0
       TI-LFA backup via iosxrv-4 (P) [1.1.1.4], iosxrv-3 (Q) [1.1.1.3]
       via 99.1.5.5, GigabitEthernet0/0/0/0 iosxrv-5, SRGB Base: 16000
       Label stack [16004, 24000, 16006]
       P: No, TM: 1040, LC: No, NP: No, D: No, SRLG: Yes
     src iosxrv-6.00-00, 1.1.1.6
```

Double-segment LFA (P and Q)

Backup path label stack

```
RP/0/0/CPU0:ios#show ospf 1 routes 1.1.1.6/32 backup-path

Topology Table for ospf 1 with ID 1.1.1.1

Codes: O - Intra area, O IA - Inter area
       O E1 - External type 1, O E2 - External type 2
       O N1 - NSSA external type 1, O N2 - NSSA external type 2

O    1.1.1.6/32, metric 20
       99.1.2.2, from 1.1.1.6, via GigabitEthernet0/0/0/1, path-id 1
           Backup path: TI-LFA, P node: 1.1.1.4, Label: 16004, Q node: 1.1.1.3, Label: 24000
              99.1.5.5, from 1.1.1.6, via GigabitEthernet0/0/0/0, protected bitmap 0000000000000001
              Attribues: Metric: 1040, SRLG Disjoint
```

Double-segment LFA

# TI-LFA for LDP – double-segment example – IP cef entry

- LDP IP-to-MPLS preferred (default)

```
RP/0/0/CPU0:iosxrv-1#show cef 1.1.1.6/32
1.1.1.6/32, version 94, internal 0x2000001 0x5 (ptr 0xa135b274) [1], 0x0 (0xa1326638), 0xa28
(0xa16ae310)
 Updated Jul 24 16:06:04.341
 local adjacency 99.1.2.2
 Prefix Len 32, traffic index 0, precedence n/a, priority 15
   via 99.1.5.5, GigabitEthernet0/0/0/0, 21 dependencies, weight 0, class 0, backup [flags
0x300]
    path-idx 0 NHID 0x0 [0xa104e184 0x0]
    next hop 99.1.5.5, P-node 1.1.1.4, Q-node 1.1.1.3
    local adjacency
     local label 24008      labels imposed {24007 24000 16006}
   via 99.1.2.2, GigabitEthernet0/0/0/1, 21 dependencies, weight 0, class 0, protected [flags
0x400]
    path-idx 1 bkup-idx 0 NHID 0x0 [0xa158c310 0x0]
    next hop 99.1.2.2
     local label 24008      labels imposed {24005}
```

LDP label to P node

Adj-SID from P to Q

Prefix-SID to destination

Double-segment LFA IP backup path

Primary path

# TI-LFA for LDP – double-segment example – LDP label entry

```
RP/0/0/CPU0:iosxrv-1#show mpls ldp bindings 1.1.1.6/32 local
1.1.1.6/32, rev 24
        Local binding: label: 24008
RP/0/0/CPU0:iosxrv-1#show mpls forwarding labels 24008 detail
Local  Outgoing    Prefix             Outgoing     Next Hop        Bytes
Label  Label       or ID              Interface                    Switched
------ ----------- ------------------ ------------ --------------- ------------
24008  24005       1.1.1.6/32         Gi0/0/0/1    99.1.2.2        0
     Updated Jul 24 16:06:04.341
     Path Flags: 0x400 [ BKUP-IDX:0 (0xa158c310) ]
     Version: 87, Priority: 15
     MAC/Encaps: 14/18, MTU: 1500
     Label Stack (Top -> Bottom): { 24005 }
     NHID: 0
     Packets Switched: 0

       24007       1.1.1.6/32         Gi0/0/0/0    99.1.5.5        0
     Updated Jul 24 16:06:04.341
     Path Flags: 0x300 [ IDX:0 BKUP, NoFwd ]
     Version: 87, Priority: 15
     MAC/Encaps: 14/26, MTU: 1500
     Label Stack (Top -> Bottom): { 24007 24000 16006 }
     NHID: 0
     Packets Switched: 0
```

LDP allocated local label 24008

Primary path

Double-segment LFA LDP MPLS backup path

LDP label to PQ node

Prefix-SID to destination

Adj-SID from P to Q

cisco

 

# Protecting LDP traffic with TI-LFA – Restrictions

- Point of Local Repair (PLR) must be SR-capable
- Protected destination must have an associated prefix-SID
  - Advertised by itself if SR-capable
  - Advertised by mapping server if not SR-capable
- Single-segment protection: PQ node must be SR-capable
  - If not, use rLFA, which requires targeted LDP session
- Double-segment protection: P and Q nodes must both be SR-capable

# TI-LFA and SR/LDP Interworking

# TI-LFA and SR/LDP Interworking

- TI-LFA leverages SR/LDP interworking functionality if not all nodes on the backup path are SR enabled
  - See “SR/LDP Interworking” slides for interworking functionality

# TI-LFA and SR/LDP Interworking – Example

- Node1 and Node4 are SR and LDP enabled
- Destination Z and Node2, Node3 and Node5 are LDP only
- A Mapping Server advertises prefix-SID 16006 for Z's loopback (1.1.1.6/32)
- Node4 is PQ node for protection of destination Z on Node1 over primary link 1-2
  - "single-segment protection"
- Node1 and Node4 use SR/LDP interworking functionality to steer the packets on the TI-LFA backup path

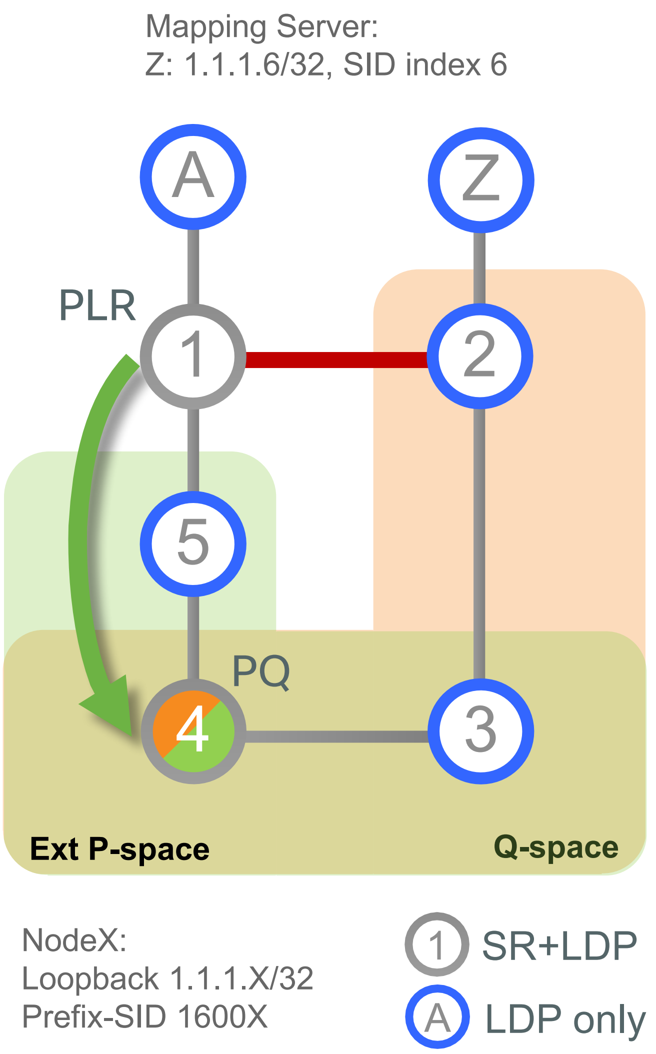

# TI-LFA and SR/LDP Interworking – Example

- On PLR Node1:

```
RP/0/0/CPU0:xrvr-1#show isis fast-reroute 1.1.1.6/32

L2 1.1.1.6/32 [20/115]
     via 99.1.2.2, GigabitEthernet0/0/0/2, xrvr-2, Weight: 0
       TI-LFA backup via xrvr-4 (PQ) [1.1.1.4]
       via 99.1.5.5, GigabitEthernet0/0/0/1 xrvr-5
       Label stack [None, 16006]
```

Backup via PQ node

IS-IS has no SR label to PQ node

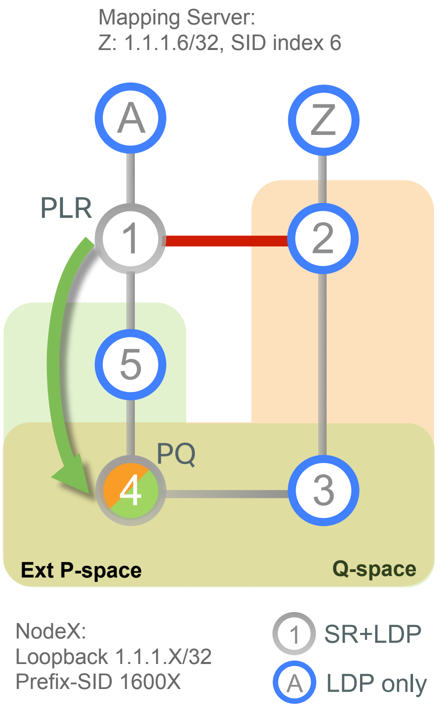


- IS-IS cannot provide an outgoing label for Node4 since the downstream neighbor Node5 is not SR enabled
- SR/LDP interworking will replace this missing label with the corresponding LDP label
- Reminder: prefix-SID for Z (16006) is advertised by Mapping Server

# TI-LFA and SR/LDP Interworking – Example

```
RP/0/0/CPU0:xrvr-1#show route 1.1.1.6/32 detail

Routing entry for 1.1.1.6/32
  Known via "isis 1", distance 115, metric 20, labeled SR(SRMS), type level-2
  Installed Sep 28 19:58:49.204 for 00:09:09
  Routing Descriptor Blocks
    99.1.5.5, from 1.1.1.6, via GigabitEthernet0/0/0/1, Backup (remote)
      Remote LFA is 1.1.1.4
      Route metric is 0
      Labels: 0x100001 0x3e86 (1048577 16006)
      Tunnel ID: None
      Binding Label: None
      Extended communities count: 0
      Path id:65                Path ref count:1
      NHID:0x2(Ref:7)
    99.1.2.2, from 1.1.1.6, via GigabitEthernet0/0/0/2, Protected
      Route metric is 20
      Label: None
      Tunnel ID: None
      Binding Label: None
      Extended communities count: 0
      Path id:1        Path ref count:0
      NHID:0x1(Ref:8)
      Backup path id:65
  Route version is 0x1b (27)
  Local Label: 0x3e86 (16006)
  <...>
```

Backup path

ISIS has no SR label to PQ node 0x100001 == unlabeled

SR/LDP interworking will replace the missing SR label with the corresponding LDP label

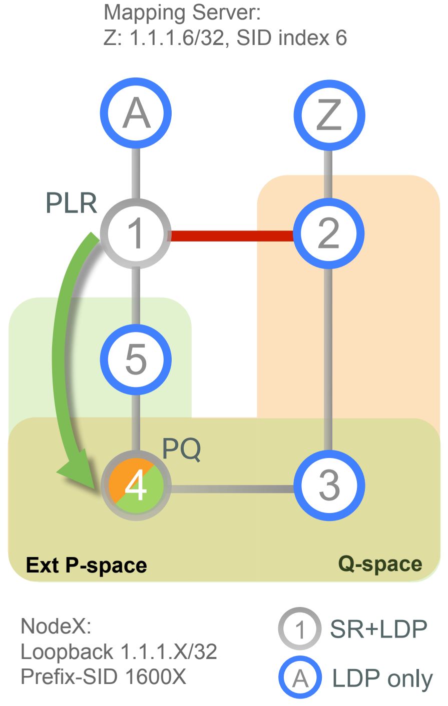

# TI-LFA and SR/LDP Interworking – Example

```
RP/0/0/CPU0:xrvr-1#show mpls ldp bindings 1.1.1.6/32
1.1.1.6/32, rev 38
        Local binding: label: 24017
        Remote bindings: (2 peers)
            Peer                Label
            -----------------   ---------
            1.1.1.2:0           24020
            1.1.1.5:0           24017
```

Local LDP label for destination prefix

```
RP/0/0/CPU0:xrvr-1#show mpls ldp bindings 1.1.1.4/32
1.1.1.4/32, rev 30
        Local binding: label: 24013
        Remote bindings: (2 peers)
            Peer                Label
            -----------------   ---------
            1.1.1.2:0           24017
            1.1.1.5:0           24011
```

LDP label advertised by Node5 for PQ Node4

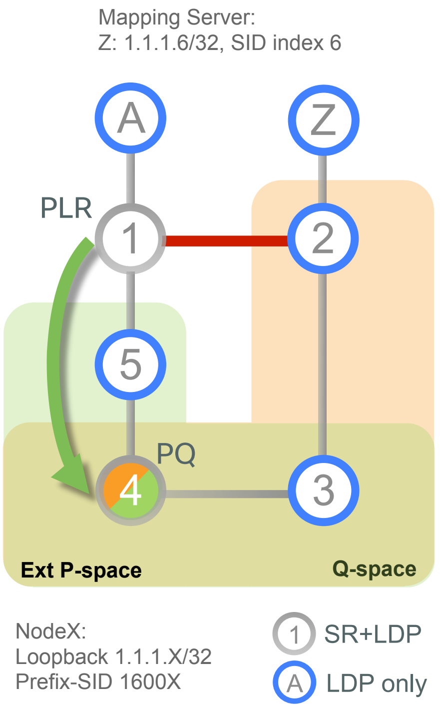


- LDP has received an outgoing label (24011) for Node4
- This LDP label will be used by SR/LDP interworking to replace the missing SR label to PQ Node4


 

# TI-LFA and SR/LDP Interworking – Example

```
RP/0/0/CPU0:xrvr-1#show cef 1.1.1.6/32
1.1.1.6/32, version 122, internal 0x1000001 0x5 (ptr 0xa13fa0f4) [1], 0x0 (0xa13dfb24), 0x
 local adjacency 99.1.2.2
 Prefix Len 32, traffic index 0, precedence n/a, priority 15
   via 99.1.5.5/32, GigabitEthernet0/0/0/1, 15 dependencies, weight 0, class 0, backup [flags 0x300]
    path-idx 0 NHID 0x0 [0xa0ed9640 0x0]
    next hop 99.1.5.5/32, PQ-node 1.1.1.4
    local adjacency
     local label 24017       labels imposed {24011 16006}
   via 99.1.2.2/32, GigabitEthernet0/0/0/2, 15 dependencies,                         ags 0x400]
    path-idx 1 bkup-idx 0 NHID 0x0 [0xa0c409dc 0x0]
    next hop 99.1.2.2/32
     local label 24017       labels imposed {24020}
```

Cef (IP2MPLS) Backup path

LDP label to PQ node

Prefix-SID to destination

```
RP/0/0/CPU0:xrvr-1#show mpls forwarding labels 24017 detail
Local  Outgoing    Prefix             Outgoing     Next Hop        Bytes
Label  Label       or ID              Interface
------ ----------- ------------------ ------------ -------------
24017  24020       1.1.1.6/32         Gi0/0/0/2    99.1.2.2
     Updated: Sep 28 19:57:35.169
     Path Flags: 0x400 [  BKUP-IDX:0 (0xa0c409dc) ]
     Version: 122, Priority: 15
     Label Stack (Top -> Bottom): { 24020 }
     NHID: 0x0, Encap-ID: N/A, Path idx: 1, Backup path idx: 0, Weight: 0
     MAC/Encaps: 14/18, MTU: 1500
     Packets Switched: 0

       24011       1.1.1.6/32         Gi0/0/0/1    99.1.5.5        0             (!)
     Updated: Sep 28 19:57:35.169
     Path Flags: 0x300 [  IDX:0 BKUP, N
     Version: 122, Priority: 15
     Label Stack (Top -> Bottom):  24011 16006
     NHID: 0x0, Encap-ID: N/A, Path idx: 0, Backup p
     MAC/Encaps: 14/22, MTU: 1500
     Packets Switched: 0
     (!): FRR pure backup
```

In the forwarding table, SR/LDP interworking has automatically replaced the missing SR label to PQ Node4 with the corresponding LDP label

MPLS2MPLS Backup path

LDP label to PQ node

Prefix-SID to destination

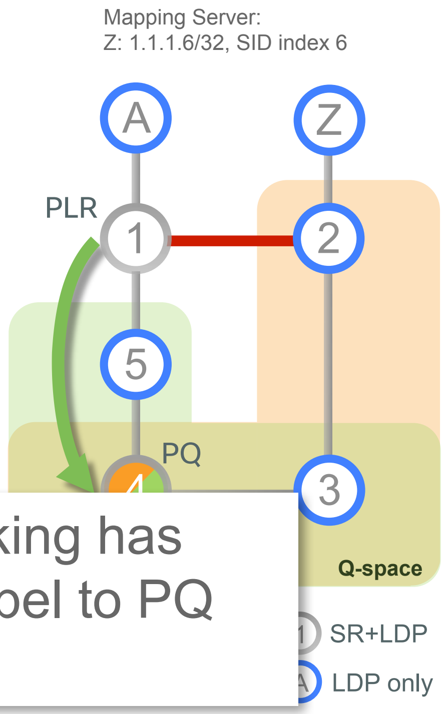

# TI-LFA and SR/LDP Interworking – Example

- On PQ node Node4:

```
RP/0/0/CPU0:xrvr-4#show mpls ldp bindings 1.1.1.6/32
1.1.1.6/32, rev 38
        Local binding: label: 24017
        Remote bindings: (2 peers)
            Peer                Label
            -----------------   ---------
            1.1.1.3:0           24017
            1.1.1.5:0           24017
```

LDP label advertised by Node3 for destination Z

```
RP/0/0/CPU0:xrvr-4#show mpls forwarding labels 16006
Local  Outgoing    Prefix             Outgoing     Next Hop        Bytes
Label  Label       or ID              Interface                    Switched
------ ----------- ------------------ ------------ --------------- --------
16006  24017       No ID              Gi0/0/0/1    99.3.4.3        0
```

In lbl: 16006, out lbl: LDP label for destination Z

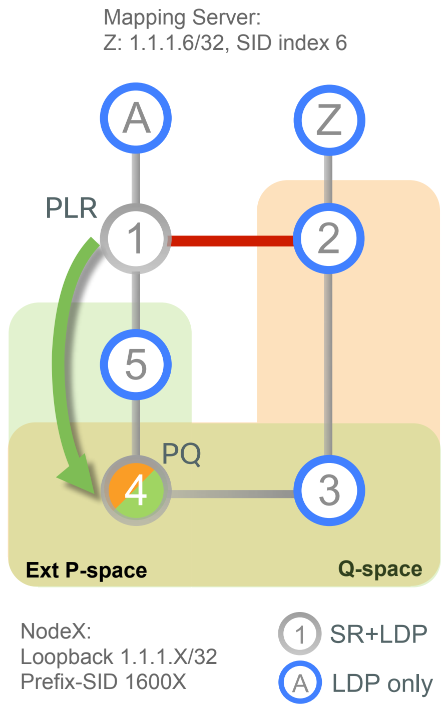


- To go from PQ Node4 to destination Z via non-SR nodes, SR/LDP interworking is used to automatically connect the prefix-SID of Z (16006) to the outgoing LDP label for Z (24017)
- Reminder: prefix-SID for Z (16006) is advertised by Mapping Server

# TI-LFA and SR/LDP Interworking – Example

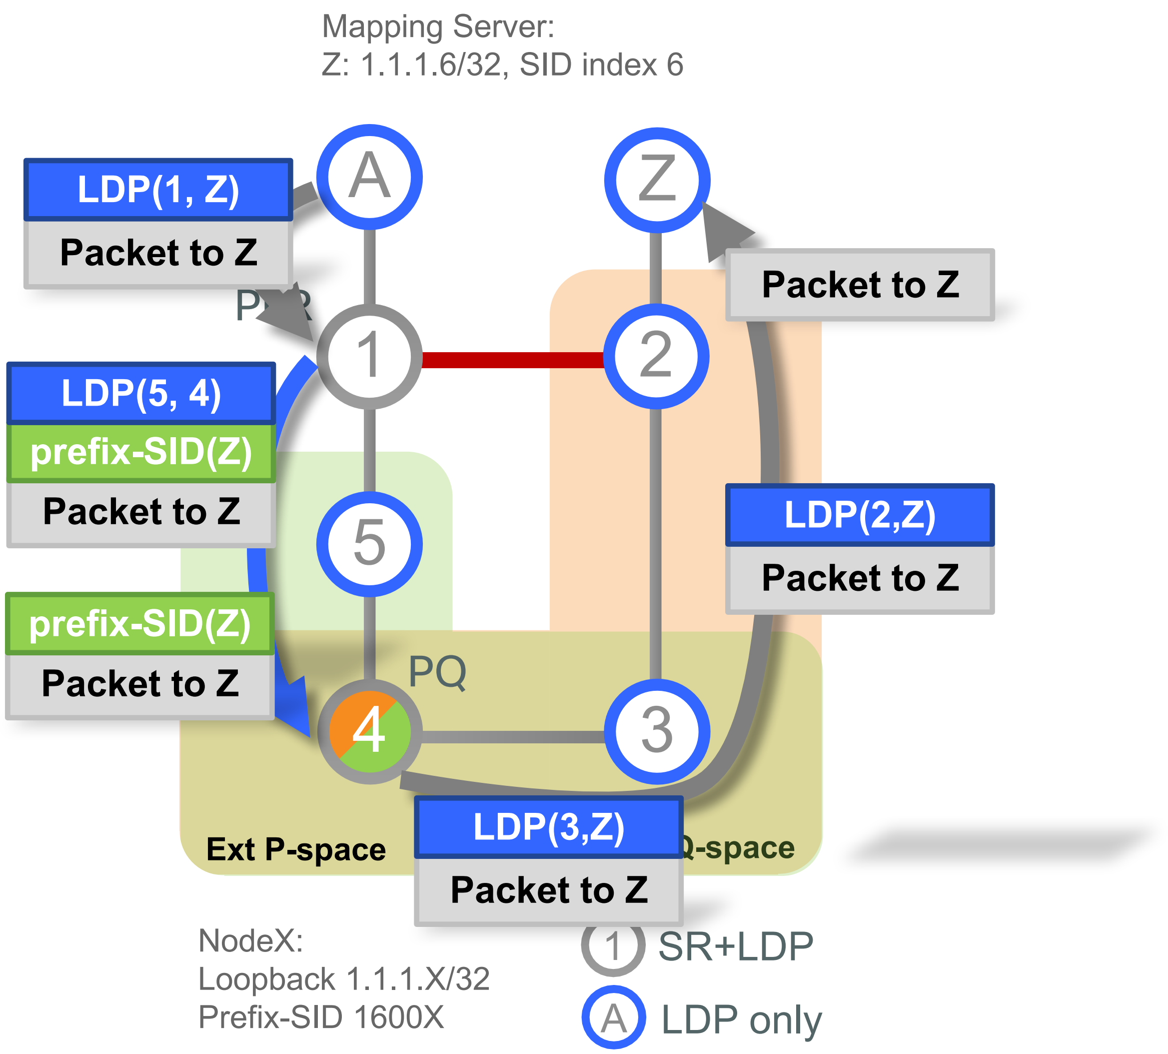

# TI-LFA for IP Traffic

# Protecting IP traffic with TI-LFA

- TI-LFA as of IOS XR 5.2.2 provides protection for IP traffic destined to D if D has also an associated Prefix-SID
  - Prefix-SID of D may be advertised by Mapping Server
- As of IOS XR 6.1.1, TI-LFA protects traffic destined to D even if D does not have an associated Prefix-SID
  - Globally configure **`segment-routing`** for this functionality

```
segment-routing
!
```

  - Note: This functionality applies to “pure” unlabeled IP traffic (no SR and no LDP labels). If a destination has an LDP label, then a Prefix-SID must be advertised for the prefix by Mapping Server, see section “TI-LFA for LDP traffic”

# Protecting IP traffic with TI-LFA – Example

- Node1-5 are SR enabled
- Source A and destination Z are IP only
- Node4 is PQ node for protection of destination Z on Node1 over primary link 1-2
  - “single-segment protection”

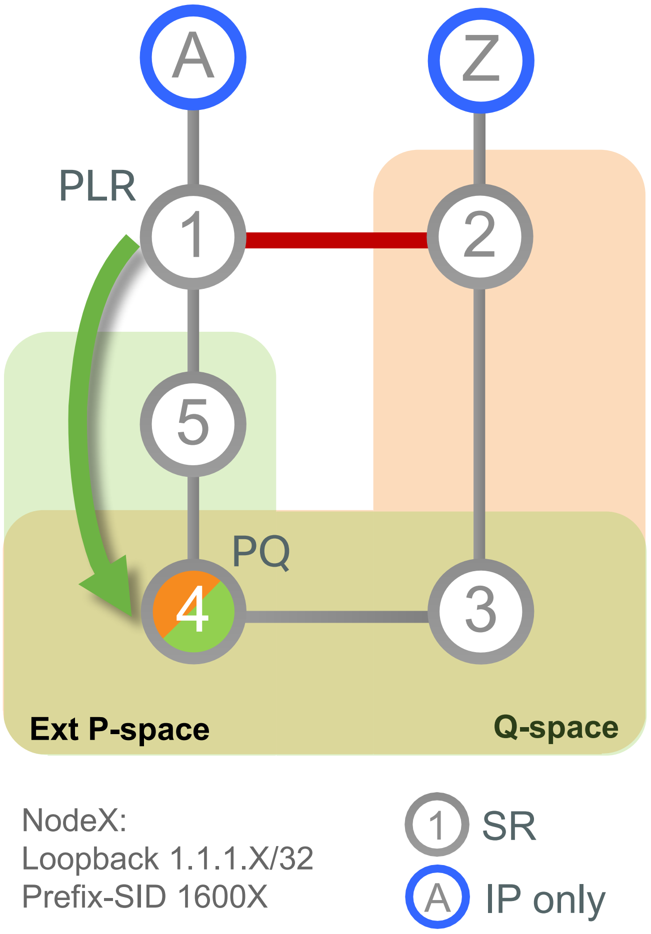

# Protecting IP traffic with TI-LFA – Example

- On PLR Node1:

```
RP/0/0/CPU0:xrvr-1#show isis fast-reroute 1.1.1.6/32

L2 1.1.1.6/32 [20/115]
     via 99.1.6.6, GigabitEthernet0/0/0/2, xrvr-6, SRGB Base: 16000, Weigh
         Backup path: TI-LFA (link), via 99.1.5.5, Gi0/0/0/1 xrvr-5, SRGB
           P node: xrvr-4.00 [1.1.1.4], Label: 16004
           Prefix label: None
```

Backup via PQ node

no SR label for destination

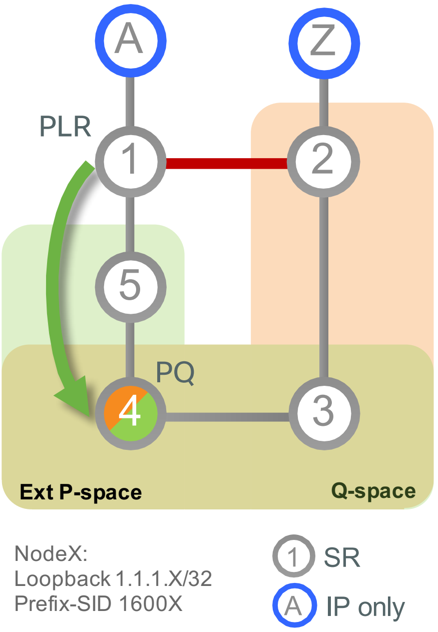


- Since prefix 1.1.1.6/32 has no Prefix-SID label, IS-IS only provides a label (16004) to reach PQ Node4 on the backup path
- From Node4, the traffic to 1.1.1.6 follows the shortest IGP path to NodeZ; the traffic to 1.1.1.6 on this path is unlabeled

# Protecting IP traffic with TI-LFA – Example

```
RP/0/0/CPU0:xrvr-1#show route 1.1.1.6/32 detail

Routing entry for 1.1.1.6/32
  Known via "isis 1", distance 115, metric 20, type level-2
  Installed Jul 27 11:47:10.385 for 01:31:51
  Routing Descriptor Blocks
    99.1.5.5, from 1.1.1.6, via GigabitEthernet0/0/0/1, Backup (TI-LFA)
      Repair Node(s): 1.1.1.4
      Route metric is 40
      Label: 0x3e84 (16004)
      Tunnel ID: None
      Binding Label: None
      Extended communities count: 0
      Path id:65               Path ref count:1
      NHID:0x1(Ref:12)
    99.1.2.2, from 1.1.1.6, via GigabitEthernet0/0/0/2, Protected
      Route metric is 20
      Label: None
      Tunnel ID: None
      Binding Label: None
      Extended communities count: 0
      Path id:1        Path ref count:0
      NHID:0x2(Ref:12)
      Backup path id:65
  Route version is 0x2 (2)
  No local label
  <...>
```

Backup path

Prefix-SID label to PQ Node4

Prefix 1.1.1.6/32 has no label on the primary path

IGP does not have a local label for this prefix

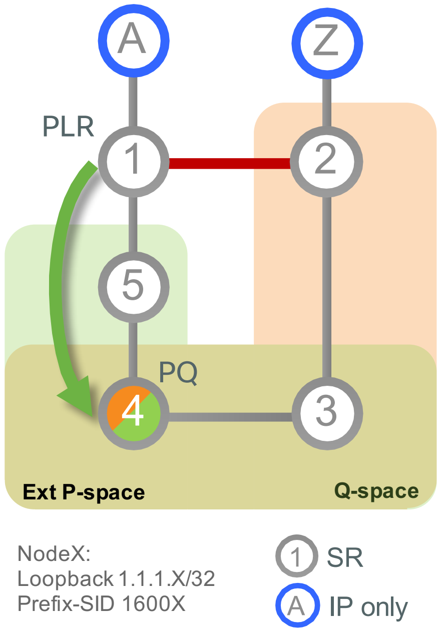

# Protecting IP traffic with TI-LFA – Example

```
RP/0/0/CPU0:xrvr-1#show cef 1.1.1.6/32
1.1.1.6/32, version 73, internal 0x1000001 0x5 (ptr 0xa12fdf34) [1], 0x0 (0xa12e1c68), 0
 Updated May 29 11:54:52.680
 local adjacency 99.1.2.2
 Prefix Len 32, traffic index 0, precedence n/a, priority 15
   via 99.1.5.5/32, GigabitEthernet0/0/0/1, 12 dependencies, weight 0, class 0, backup [flags 0x300]
    path-idx 0 NHID 0x0 [0xa165aac4 0x0]
    next hop 99.1.5.5/32
    local adjacency
     local label 28108       labels imposed {16004}
   via 99.1.2.2/32, GigabitEthernet0/0/0/2, 12 dependencies, weight 0, class 0, protected [flags 0x400]
    path-idx 1 bkup-idx 0 NHID 0x0 [0xa182742c 0x0]
    next hop 99.1.2.2/32
     local label 28108       labels imposed {None}
```

cef (IP2MPLS) Backup path

Prefix-SID to PQ node

No label imposition on primary path

Dynamically allocated local label

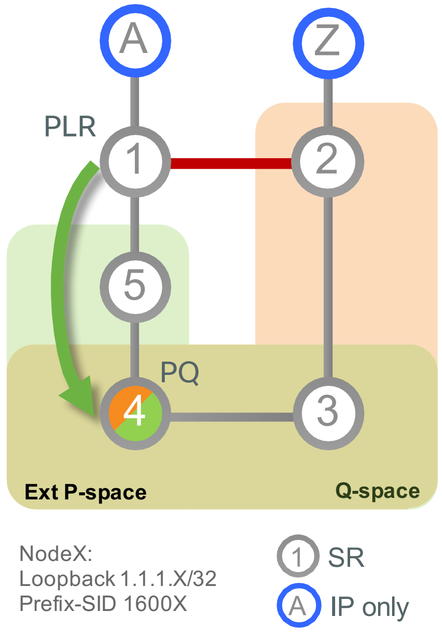


 

# TI-LFA Node and SRLG protection

# Protecting link, node, SRLG

- User specifies which resource needs to be protected: Link, Node, or SRLG
  - Node and SRLG protection imply link protection
  - Link protection may also protect node and/or SRLG (guaranteed or de facto)
- One repair path (per destination) is installed in FIB, repair path does not depend on actual failure
  - The protection of the configured type is activated by link/BFD failure
- Failure Statistics indicate that node failures are very rare
  - Enabling node protection is not always better, even for link failures the node failure post-convergence path is followed

# Protecting link, node, SRLG

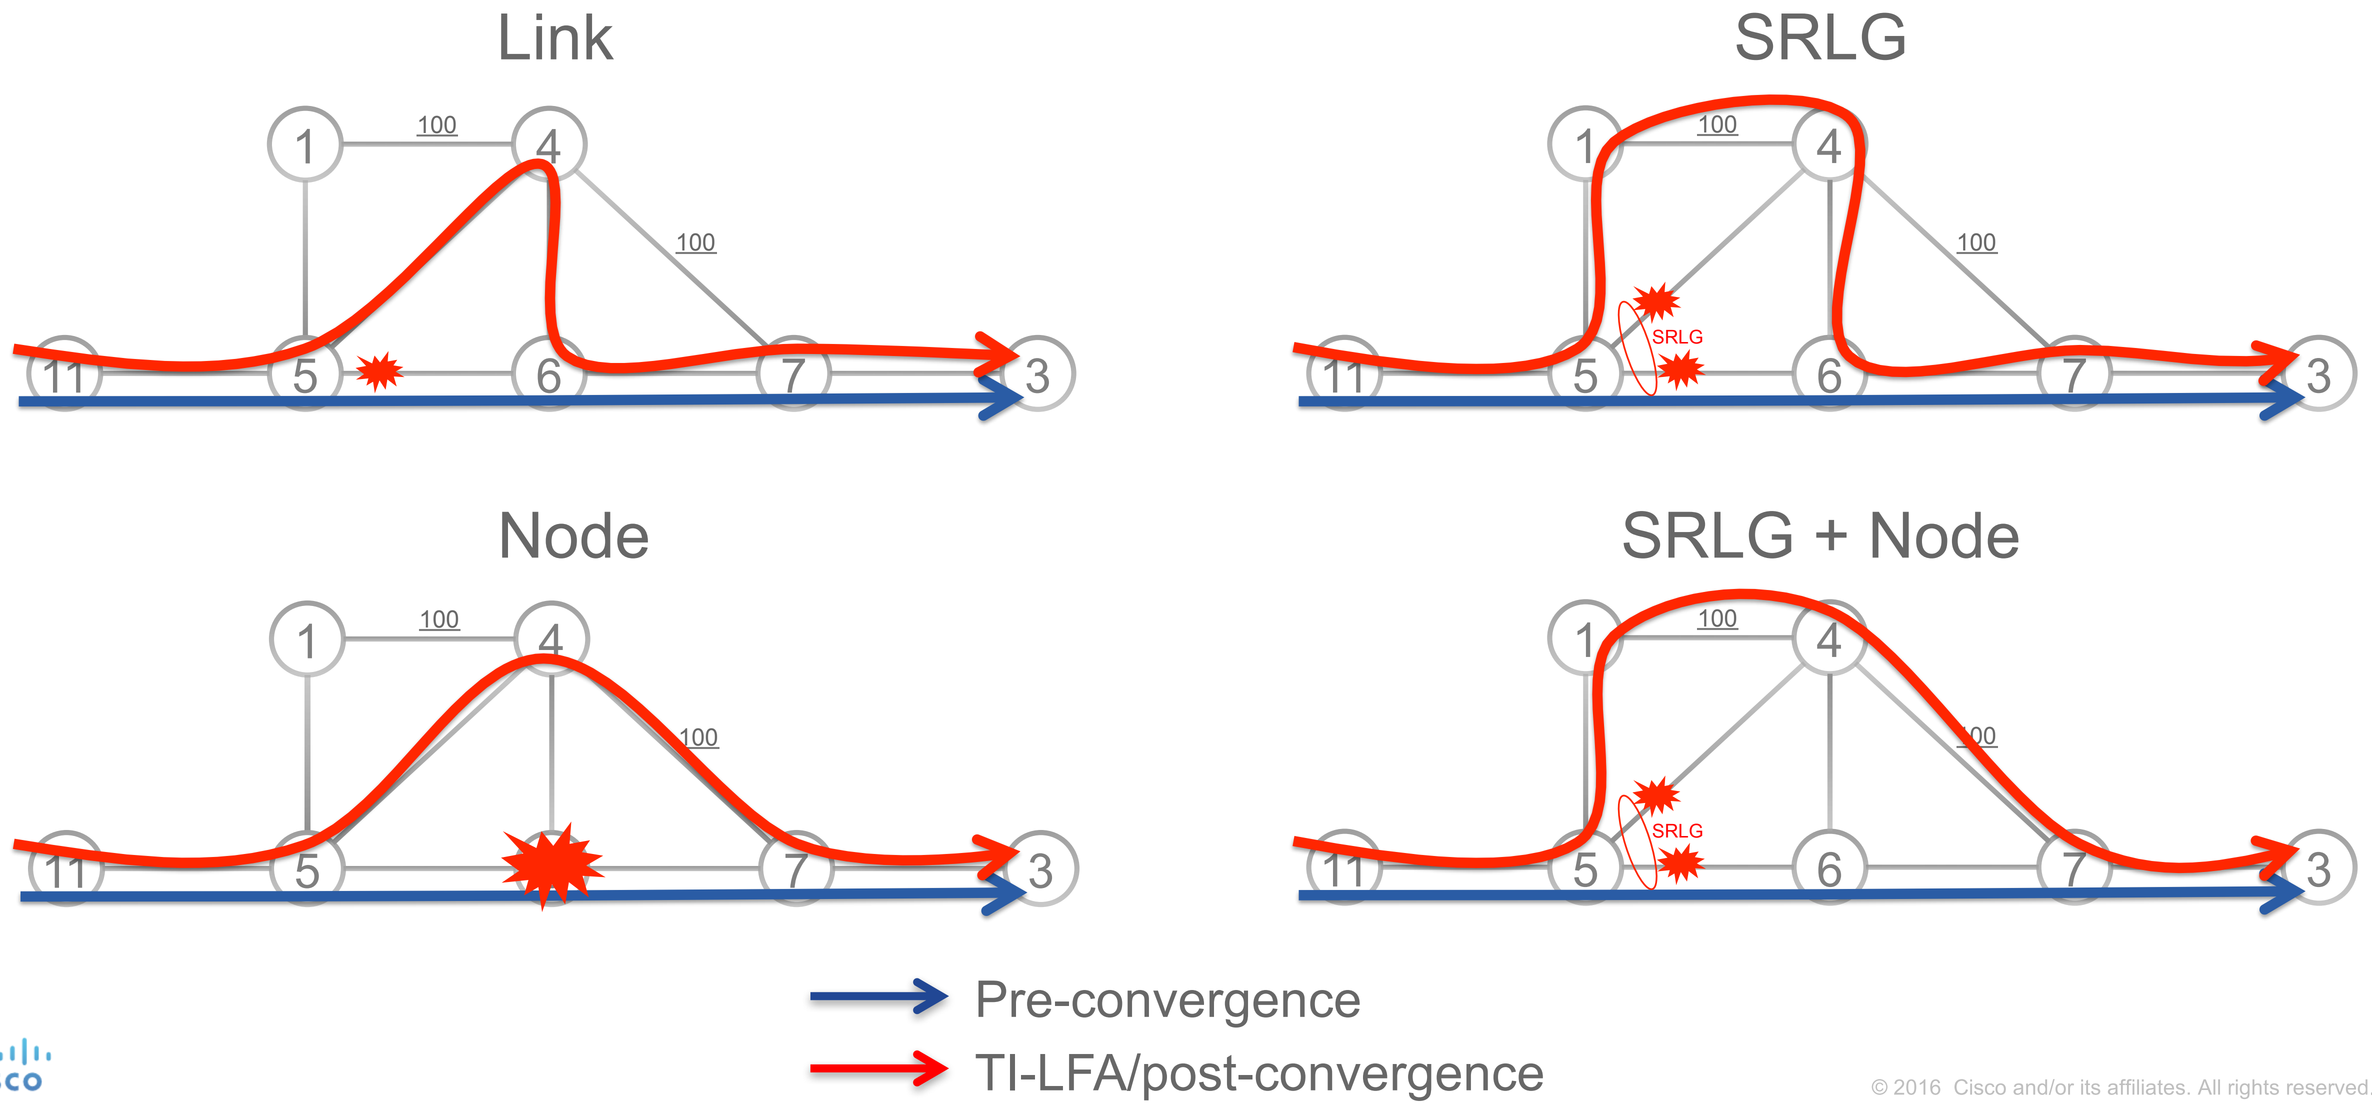


 

# Guaranteed or de-facto node protection

- Link protection may also provide guaranteed node protection
  - The link protecting repair path is the same as the node protecting repair path

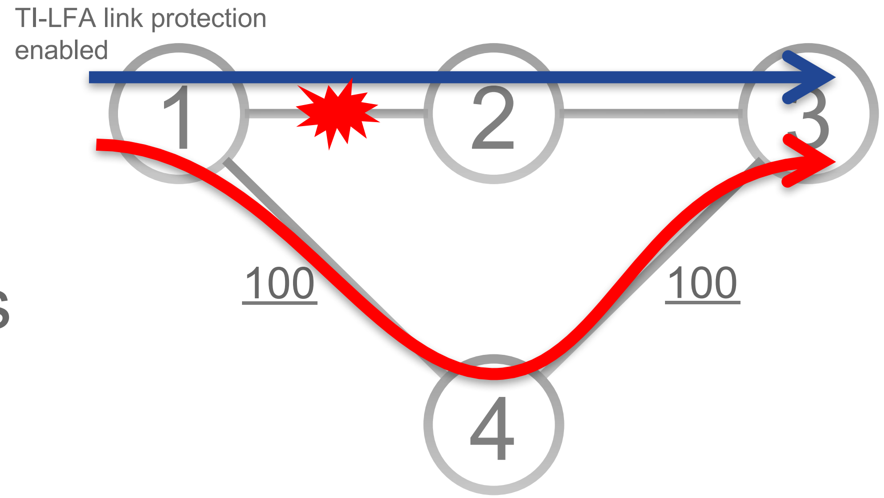


- Link protection can provide de-facto node or SRLG protection
  - Multiple link protecting repair paths can provide protection in case of node failures

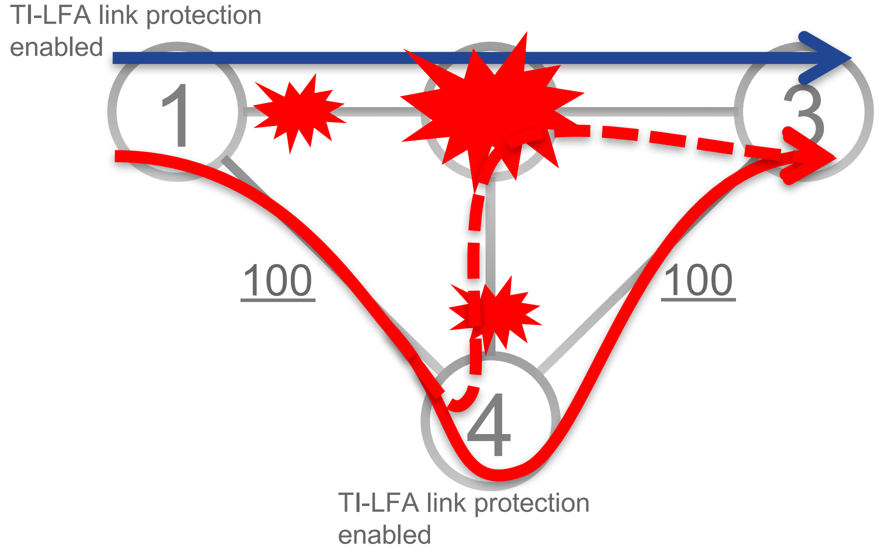

# TI-LFA Node and SRLG protection

- From IOS XR 6.1.1 TI-LFA provides protection for Node and local SRLG Failures

# Topology Independent LFA – Configuration

- TI-LFA uses tie-breaker configuration to specify preference order of protection modes: link, node, SRLG
- Protection mode preference has interface scope
  - ISIS: also configurable per instance address-family
  - OSPF: also configurable per instance and per area
- Link protection always has lowest preference, not configurable
- OSPF, ISIS: most preferred has highest index

# Topology Independent LFA – Configuration

- IS-IS configuration

Node5:

```
router isis 1
 interface GigabitEthernet0/0/0/1
  point-to-point
  address-family ipv4 unicast
   fast-reroute per-prefix
   fast-reroute per-prefix tiebreaker node-protecting index 200
   fast-reroute per-prefix tiebreaker srlg-disjoint index 100
   fast-reroute per-prefix ti-lfa
```

Node + SRLG + Link possible

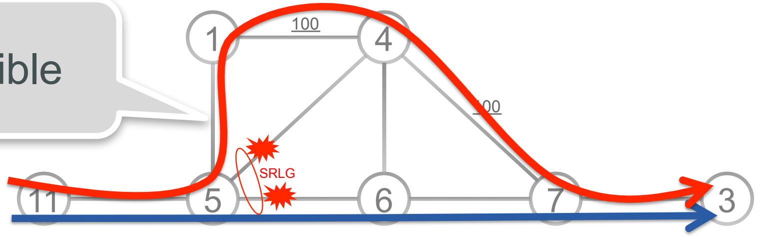


Node and SRLG configured, Link is default

```
RP/0/0/CPU0:xrvr-5#show isis fast-reroute 1.1.1.3/32 detail

L2 1.1.1.3/32 [120/115] medium priority
     via 99.5.6.6, GigabitEthernet0/0/0/1, xrvr-6, SRGB Base: 16000, Weight: 0
         Backup path: TI-LFA (node+srlg), via 99.1.5.1, GigabitEthernet0/0/0/0 xrvr-1, SRGB Base: 16000, Weight: 0
           P node: xrvr-4.00 [1.1.1.4], Label: 16004
           Q node: xrvr-7.00 [1.1.1.7], Label: 24047
           Prefix label: 16003
       P: No, TM: 310, LC: No, NP: No, D: No, SRLG: No
     src xrvr-3.00-00, 1.1.1.3, prefix-SID index 3, R:0 N:1 P:0 E:0 V:0 L:0
```

Node + SRLG (+ Link) protection installed

# Node xor SRLG protection

- If Node protection is possible and local SRLG protection is possible but not concurrently, then the tiebreaker preference index indicates which one to select
  - if node-protecting index > srlg-disjoint index then node-protecting else srlg-disjoint

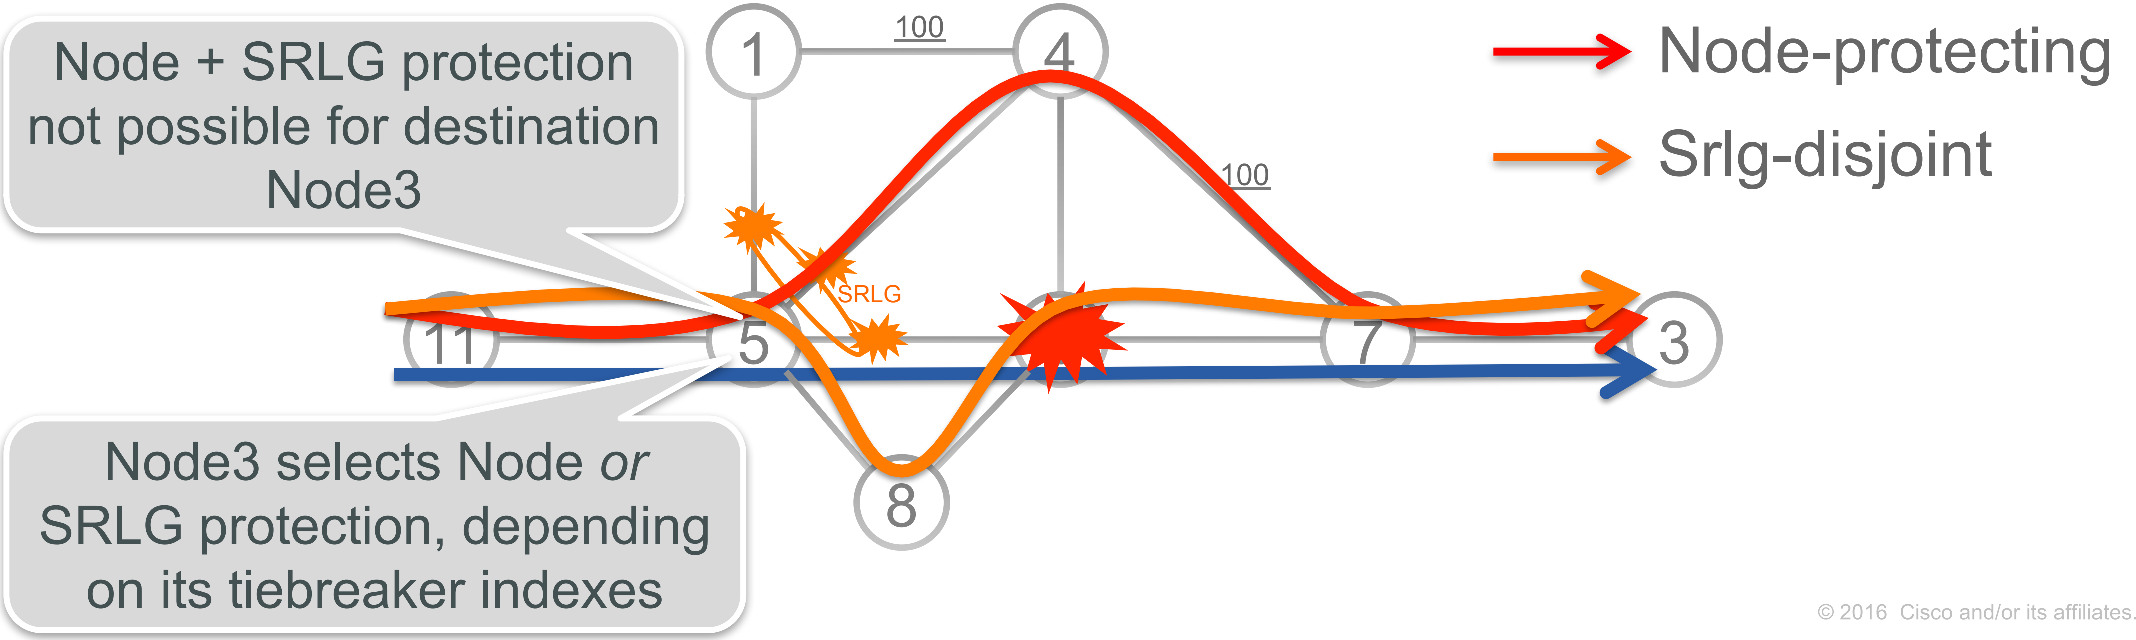


 

# Topology Independent LFA – Configuration

- IS-IS configuration

SRLG + Link possible,
Node not possible

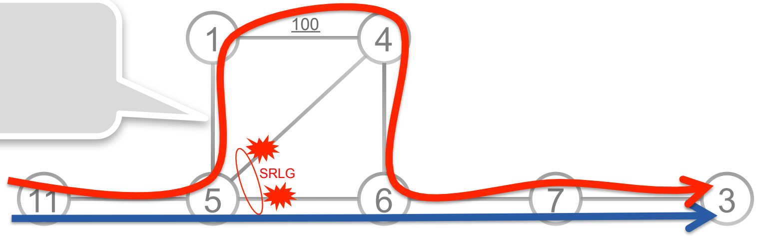


Node5:

```
router isis 1
 interface GigabitEthernet0/0/0/1
  point-to-point
  address-family ipv4 unicast
   fast-reroute per-prefix
   fast-reroute per-prefix tiebreaker node-protecting index 200
   fast-reroute per-prefix tiebreaker srlg-disjoint index 100
   fast-reroute per-prefix ti-lfa
```

Node and SRLG configured, Link is default

```
RP/0/0/CPU0:xrvr-5#show isis fast-reroute 1.1.1.3/32 detail

L2 1.1.1.3/32 [120/115] medium priority
     via 99.5.6.6, GigabitEthernet0/0/0/1, xrvr-6, SRGB Base: 16000, Weight: 0
         Backup path: TI-LFA (srlg), via 99.1.5.1, GigabitEthernet0/0/0/0 xrvr-1, SRGB Base: 16000, Weight: 0
           P node: xrvr-4.00 [1.1.1.4], Label: 16004
           Prefix label: 16003
       P: No, TM: 230, LC: No, NP: No, D: No, SRLG: No
     src xrvr-3.00-00, 1.1.1.3, prefix-SID index 3, R:0 N:1 P:0 E:0 V:0 L:0
```

SRLG (+ Link) protection installed

# Topology Independent LFA – Configuration

- IS-IS configuration

Node + Link possible,
SRLG not possible

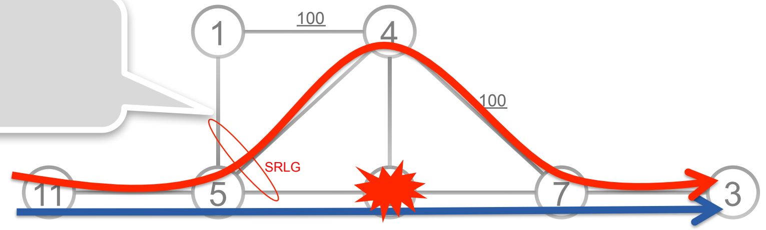


Node5:

```
router isis 1
 interface GigabitEthernet0/0/0/1
  point-to-point
  address-family ipv4 unicast
   fast-reroute per-prefix
   fast-reroute per-prefix tiebreaker node-protecting index 200
   fast-reroute per-prefix tiebreaker srlg-disjoint index 100
   fast-reroute per-prefix ti-lfa
```

Node and SRLG configured, Link is default

```
RP/0/0/CPU0:xrvr-5#show isis fast-reroute 1.1.1.3/32 detail

L2 1.1.1.3/32 [120/115] medium priority
     via 99.5.6.6, GigabitEthernet0/0/0/1, xrvr-6, SRGB Base: 16000, Weight: 0
         Backup path: TI-LFA (node), via 99.4.5.4, GigabitEthernet0/0/0/2 xrvr-4, SRGB Base: 16000, Weight: 0
           P node: xrvr-4.00 [1.1.1.4], Label: ImpNull
           Q node: xrvr-7.00 [1.1.1.7], Label: 24047
           Prefix label: 16003
       P: No, TM: 210, LC: No, NP: No, D: No, SRLG: No
     src xrvr-3.00-00, 1.1.1.3, prefix-SID index 3, R:0 N:1 P:0 E:0 V:0 L:0
```

Node (+ Link) protection installed

 

# Topology Independent LFA – Configuration

- IS-IS configuration

Node5:

```
router isis 1
 interface GigabitEthernet0/0/0/1
  point-to-point
  address-family ipv4 unicast
   fast-reroute per-prefix
   fast-reroute per-prefix tiebreaker node-protecting index 200
   fast-reroute per-prefix tiebreaker srlg-disjoint index 100
   fast-reroute per-prefix ti-lfa
```

Link possible,
Node or SRLG not possible

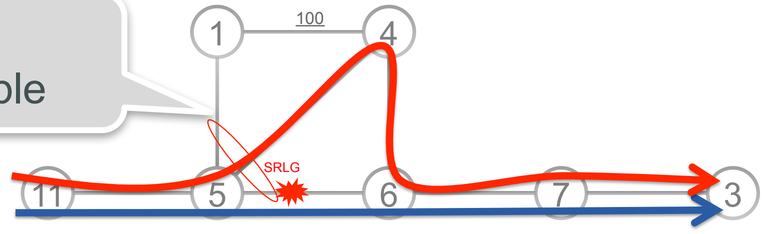


Node and SRLG configured, Link is default

```
RP/0/0/CPU0:xrvr-5#show isis fast-reroute 1.1.1.3/32 detail

L2 1.1.1.3/32 [120/115] medium priority
     via 99.5.6.6, GigabitEthernet0/0/0/1, xrvr-6, SRGB Base: 16000, Weight: 0
       FRR backup via 99.4.5.4, GigabitEthernet0/0/0/2, xrvr-4, SRGB Base: 16000, Weight: 0, Metric: 130
       P: No, TM: 130, LC: No, NP: No, D: No, SRLG: No
     src xrvr-3.00-00, 1.1.1.3, prefix-SID index 3, R:0 N:1 P:0 E:0 V:0 L:0
```

Link protection installed

# Topology Independent LFA – Configuration

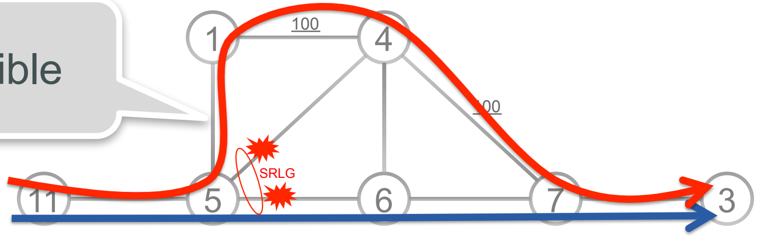


Node + SRLG + Link possible

- OSPF configuration

Node5:

```
router ospf 1
 area 0
  interface GigabitEthernet0/0/0/1
   network point-to-point
   fast-reroute per-prefix
   fast-reroute per-prefix ti-lfa enable
   fast-reroute per-prefix tiebreaker node-protecting index 200
   fast-reroute per-prefix tiebreaker srlg-disjoint index 100
```

Node and SRLG configured, Link is default

```
RP/0/0/CPU0:xrvr-5#show ospf routes 1.1.1.3/32 backup-path
Codes: O - Intra area, O IA - Inter area
       O E1 - External type 1, O E2 - External type 2
       O N1 - NSSA external type 1, O N2 - NSSA external type 2

O    1.1.1.3/32, metric 121
       99.5.6.6, from 1.1.1.3, via GigabitEthernet0/0/0/1, path-id 1
           Backup path: TI-LFA, Repair-List: P node: 1.1.1.1         Label: 3
                                             Q node: 1.1.1.4         Label: 24014
                                             Q node: 1.1.1.7         Label: 24047
             99.1.5.1, from 1.1.1.3, via GigabitEthernet0/0/0/0, protected bitmap 0000000000000001
             Attributes: Metric: 311, Node Protect, SRLG Disjoint
```

Node + SRLG (+ Link) protection installed

# Topology Independent LFA – Configuration

- OSPF configuration

SRLG + Link possible,
Node not possible

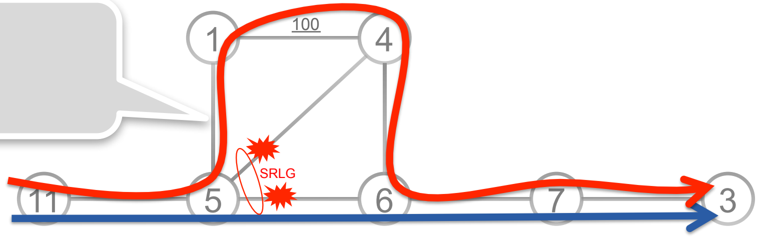


Node5:

```
router ospf 1
 area 0
  interface GigabitEthernet0/0/0/1
   network point-to-point
   fast-reroute per-prefix
   fast-reroute per-prefix ti-lfa enable
   fast-reroute per-prefix tiebreaker node-protecting index 200
   fast-reroute per-prefix tiebreaker srlg-disjoint index 100
```

Node and SRLG configured, Link is default

```
RP/0/0/CPU0:xrvr-5#show ospf routes 1.1.1.3/32 backup-path
Codes: O - Intra area, O IA - Inter area
       O E1 - External type 1, O E2 - External type 2
       O N1 - NSSA external type 1, O N2 - NSSA external type 2

O    1.1.1.3/32, metric 121
       99.5.6.6, from 1.1.1.3, via GigabitEthernet0/0/0/1, path-id 1
           Backup path: TI-LFA, Repair-List: P node: 1.1.1.1         Label: 3
                                             Q node: 1.1.1.4         Label: 24014
              99.1.5.1, from 1.1.1.3, via GigabitEthernet0/0/0/0, protected bitmap 0000000000000001
              Attributes: Metric: 231, SRLG Disjoint
```

SRLG (+ Link) protection installed

# Topology Independent LFA – Configuration

Node + Link possible,
SRLG not possible

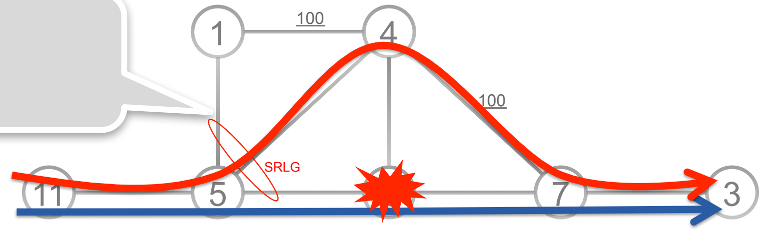


- OSPF configuration

Node5:

```
router ospf 1
 area 0
  interface GigabitEthernet0/0/0/1
   network point-to-point
   fast-reroute per-prefix
   fast-reroute per-prefix ti-lfa enable
   fast-reroute per-prefix tiebreaker node-protecting index 200
   fast-reroute per-prefix tiebreaker srlg-disjoint index 100
```

Node and SRLG configured, Link is default

```
RP/0/0/CPU0:xrvr-5#show ospf routes 1.1.1.3/32 backup-path
Codes: O - Intra area, O IA - Inter area
       O E1 - External type 1, O E2 - External type 2
       O N1 - NSSA external type 1, O N2 - NSSA external type 2

O    1.1.1.3/32, metric 121
       99.5.6.6, from 1.1.1.3, via GigabitEthernet0/0/0/1, path-id 1
           Backup path: TI-LFA, Repair-List: P node: 1.1.1.4         Label: 3
                                             Q node: 1.1.1.7         Label: 24047
              99.4.5.4, from 1.1.1.3, via GigabitEthernet0/0/0/2, protected bitmap 0000000000000001
              Attributes: Metric: 211, Node Protect,
```

Node (+ Link) protection installed

# Topology Independent LFA – Configuration

- OSPF configuration

Link possible,
Node or SRLG not possible

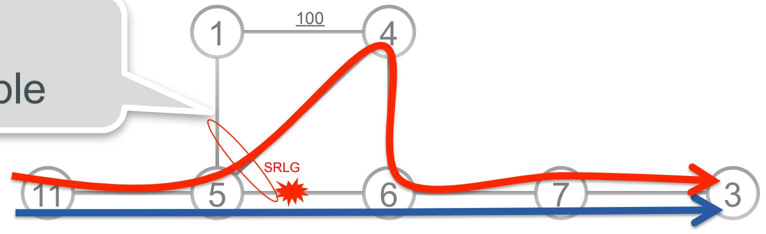


Node5:

```
router ospf 1
 area 0
  interface GigabitEthernet0/0/0/1
   network point-to-point
   fast-reroute per-prefix
   fast-reroute per-prefix ti-lfa enable
   fast-reroute per-prefix tiebreaker node-protecting index 200
   fast-reroute per-prefix tiebreaker srlg-disjoint index 100
```

Node and SRLG configured, Link is default

```
RP/0/0/CPU0:xrvr-5#show ospf routes 1.1.1.3/32 backup-path
Codes: O - Intra area, O IA - Inter area
       O E1 - External type 1, O E2 - External type 2
       O N1 - NSSA external type 1, O N2 - NSSA external type 2

O    1.1.1.3/32, metric 121
       99.5.6.6, from 1.1.1.3, via GigabitEthernet0/0/0/1, path-id 1
           Backup path:
              99.4.5.4, from 1.1.1.3, via GigabitEthernet0/0/0/2, protected bitmap 0000000000000001
              Attributes: Metric: 131,
```

Link protection installed

# Inheritance – ISIS

- TI-LFA tiebreakers can be configured under the instance Address-Family (AF)
  - These tiebreakers are the default for that AF, applied on an interface when NO tiebreaker is configured under that interface AF
  - Tiebreaker configurations under instance and interface are not merged

# Inheritance – ISIS – Example

```
router isis 1
 address-family ipv4 unicast
   fast-reroute per-prefix tiebreaker node-protecting index 200
   fast-reroute per-prefix tiebreaker srlg-disjoint index 100
 !
 interface GigabitEthernet0/0/0/1
  point-to-point
  address-family ipv4 unicast
   fast-reroute per-prefix
   fast-reroute per-prefix ti-lfa
```

Use these tiebreakers on Gi0/0/0/1

```
router isis 1
 address-family ipv4 unicast
   fast-reroute per-prefix tiebreaker node-protecting index 200
   fast-reroute per-prefix tiebreaker srlg-disjoint index 100
 !
 interface GigabitEthernet0/0/0/1
  point-to-point
  address-family ipv4 unicast
   fast-reroute per-prefix
   fast-reroute per-prefix tiebreaker node-protecting index 200
   fast-reroute per-prefix ti-lfa
```

Only Node-protecting (and link-protecting) considered on Gi0/0/0/1

 

# Inheritance – ISIS – Example

```
router isis 1
 address-family ipv4 unicast
  fast-reroute per-prefix tiebreaker node-protecting index 200
  fast-reroute per-prefix tiebreaker srlg-disjoint index 100
 !
 interface GigabitEthernet0/0/0/1
  point-to-point
  address-family ipv4 unicast
   fast-reroute per-prefix
   fast-reroute per-prefix tiebreaker default
   fast-reroute per-prefix ti-lfa
```

Only link-protecting considered on Gi0/0/0/1

# Microloop avoidance



 

# What is a microloop?

- Microloops are a day-one IP drawback
- IP hop-by-hop routing may induce microloop at any topology transition
  - Link up/down, metric up/down
- E.g. Microloops can occur after failure of link 6-7

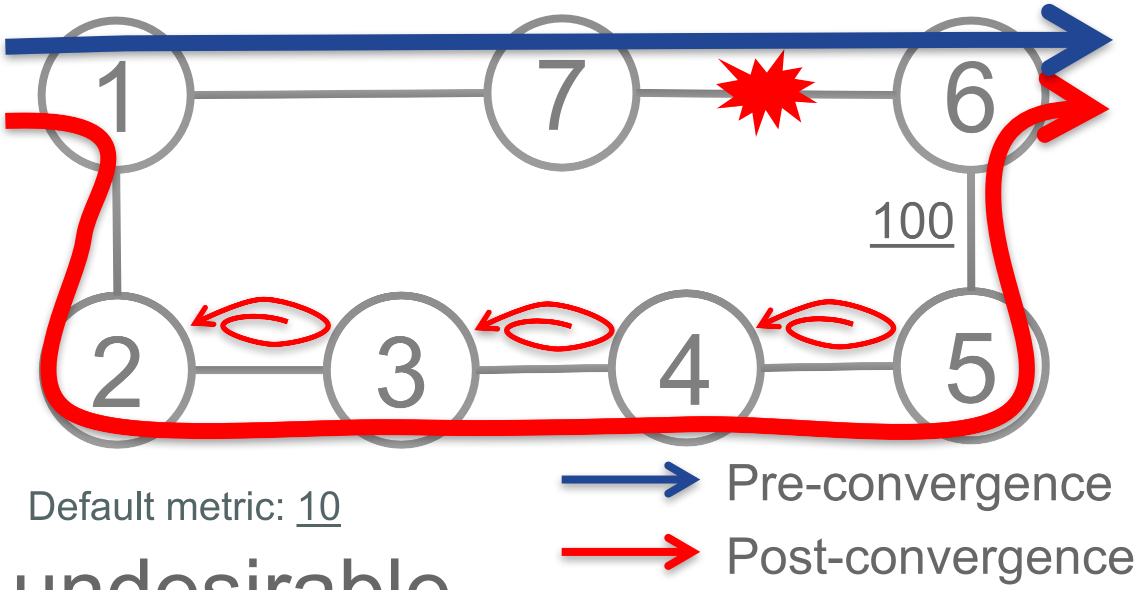


- Microloops can increase packet loss, which is especially undesirable when FRR is used.

# SR microloop avoidance

- Prevent any microloop upon an isolated convergence due to
  - link up/down event
  - metric increase/decrease event
- If multiple back-to-back convergences, fall back to native IP convergence
- Configuration:

```
router isis 1
 address-family ipv4 unicast
  microloop avoidance segment-routing
```

```
router ospf 1
 microloop avoidance segment-routing
```

# SR microloop avoidance – workflow

- TI-LFA protection kicks in on Node7, repairing the traffic to Node6 via Node5 and link Node5-Node6
- All nodes are notified of the topology change due to the failure
- E.g. Node1 computes the post-convergence SPT and detects possible microloops on the post-convergence paths for any destination, such as Node6
- If microloops are possible on the post-convergence path for a destination, then a SID-list is constructed to steer the traffic to that destination loop-free over the post-convergence path; in this example: <Prefix-SID(5), Adj-SID(5-6)> for destination Node6

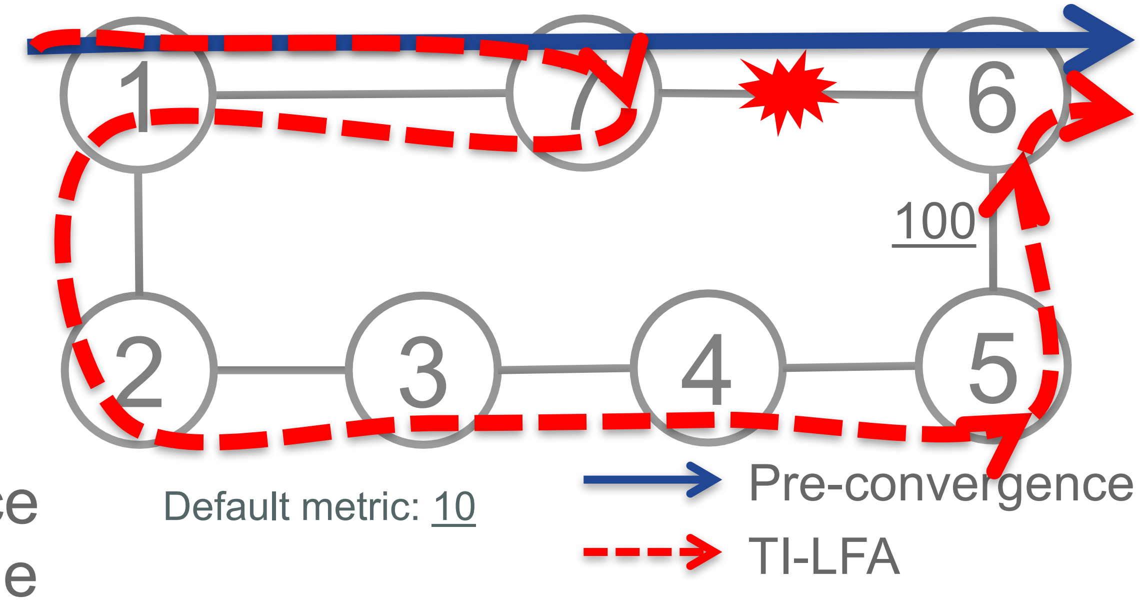

# SR microloop avoidance – workflow

- IGP on Node1 updates the forwarding table and installs the SID-list imposition entries for those destinations with possible microloops, such as destination Node6
  - Node1 imposes SID-list <Prefix-SID(5), Adj-SID(5-6)> on packets to Node6

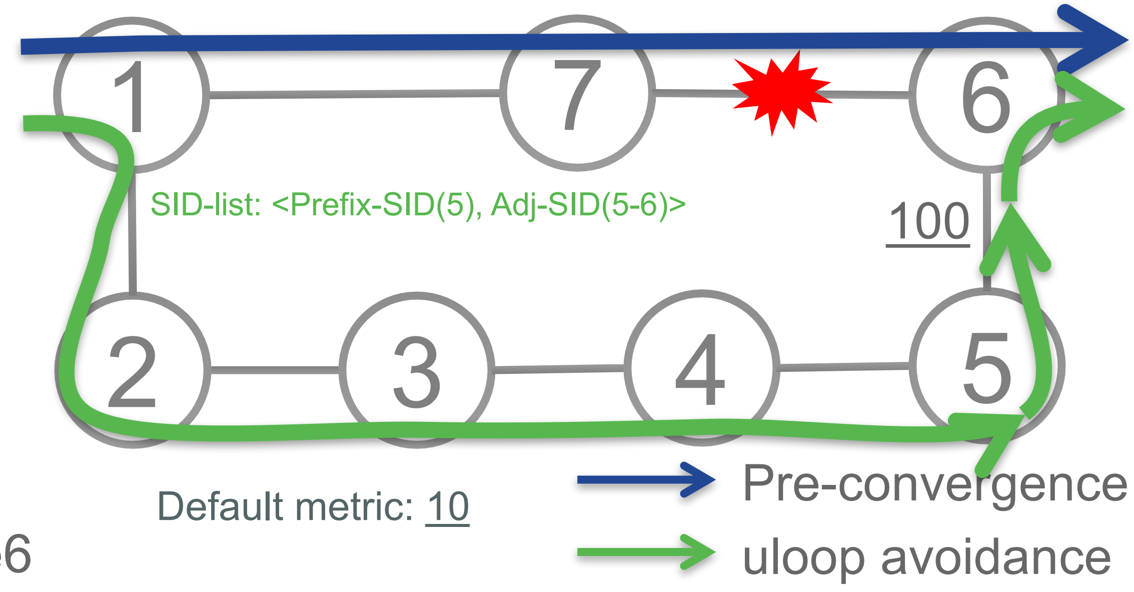


- All nodes converge and update their forwarding tables, using SID-lists where needed

# SR microloop avoidance – workflow

- Some time later, the new topology is applied and no more microloops are expected
- IGP updates the forwarding table, removing the microloop avoidance SID-lists; traffic now natively follows the post-convergence path

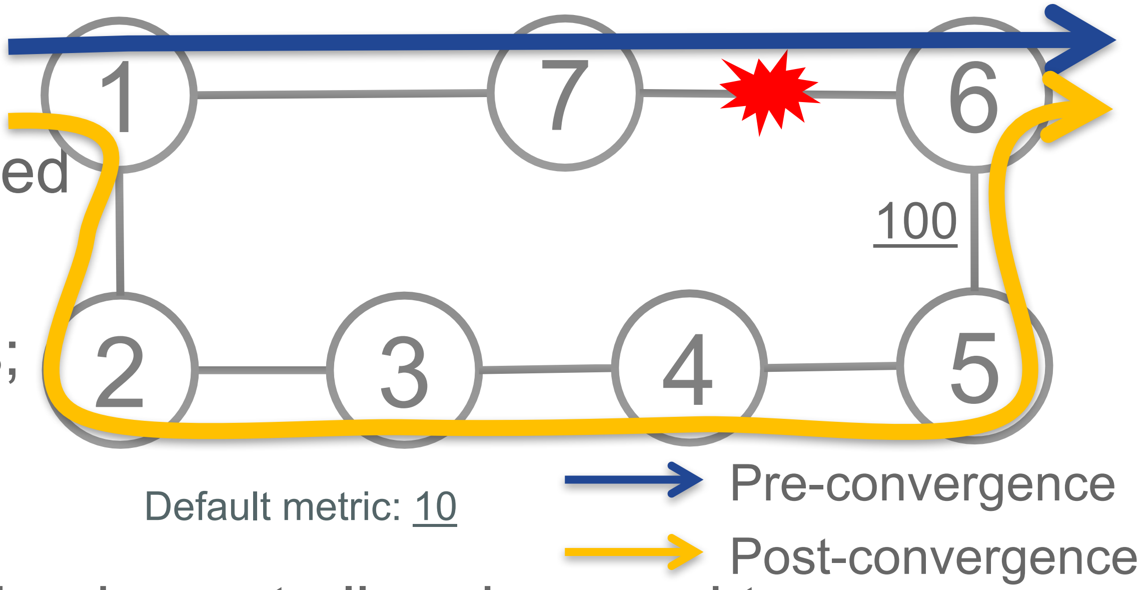


- Note: SR microloop avoidance is a local behavior, not all nodes need to implement it to get the benefits
- There is incremental benefit for each node that has it implemented
  - E.g. if only Node1 has SR microloop avoidance, then e.g. traffic entering Node2 (not from Node1) to Node6 would still see microloops
  - When enabling SR microloop avoidance on Node2, then e.g. traffic entering Node3 (not from Node2) to Node6 would still see microloops, etc.

# Using TE infrastructure

# Using TE infra for TI-LFA backup path

- IGP automatically instantiates an SRTE Policy if the label stack of the repair path exceeds 3 labels, for all protection modes (link, node, srlg)
  - Label stack size limited only by platform label imposition limit
- TE provides the infrastructure to impose a label stack
  - IGP provides the outgoing interface and label stack to TE
  - TE instantiates SRTE Policy
- No other involvement of TE
  - E.g. not required to distribute TE link attributes
- MPLS PIE is required to use TE
  - IGP still provides repair path with ≤3 labels without MPLS PIE

 

# Required TE configuration

```
ipv4 unnumbered mpls traffic-eng Loopback0
```

Provide default source address for SRTE Policy

- Optional

```
mpls traffic-eng
 auto-tunnel p2p
  tunnel-id min 1000 max 1999
```

Specify range of tunnel-te interface numbers to use

Visit us:

cisco.com

segment-routing.net

Acknowledgements:

Ahmed Bashandy
Robert Hanzl
Steven Luong
Stefano Previdi
Peter Psenak

Thank you.