MPLS in communication networks. (Part 3)

«Computer networks and telecommunications» (Additional chapters).

MPLS VPN Technology

Traffic Engineering Concepts

What Is Traffic Engineering?

- Term in common use in telephone voice network world
- Measures, models, and controls traffic to achieve various goals
- Provides an integrated approach to engineering traffic at Layer 3 (ISO/OSI)

What Is Traffic Engineering? Traffic Engineering Motivations

- Reduce the overall cost of operations by more efficient use of bandwidth resources
- Prevent a situation where some parts of a service provider network are overutilized (congested), while other parts remain underutilized
- Implement traffic protection against failures
- Enhance SLA in combination with QoS

Business Drivers for Traffic Engineering

- Routers always forward traffic along the least-cost route as discovered by IGP.
- Network bandwidth may not be efficiently utilized:
    - The least-cost route may not be the only possible route.
    - The least-cost route may not have enough resources to carry all the traffic.

Business Drivers for Traffic Engineering (Cont.)

- Lack of resources results in congestion in two ways:
  - When network resources themselves are insufficient to accommodate offered load
  - When traffic streams are inefficiently mapped onto available resources
- Some resources are overutilized while others remain underutilized.

Congestion Avoidance and Traffic Engineering

- Network congestion can be addressed by either:
  - Expansion of capacity or classical congestion control techniques (queuing, rate limiting, etc.)
  - Traffic Engineering (TE), if the problems result from inefficient resource allocation

The focus of TE is on congestion problems that are prolonged, not on short-term bursts

Congestion Avoidance and Traffic Engineering

Without the use of TE, all traffic can be redirected to the route, where there is not enough bandwidth - drops will begin.

The convergence rate of OSPF or ISIS even when using BFD Bidirectional Forwarding Detection is in the tens of «ms».

After that, the transport LSP must also be rebuilt.

It will not go unnoticed by subscribers.

Traffic Engineering with a Layer 2 Overlay Model

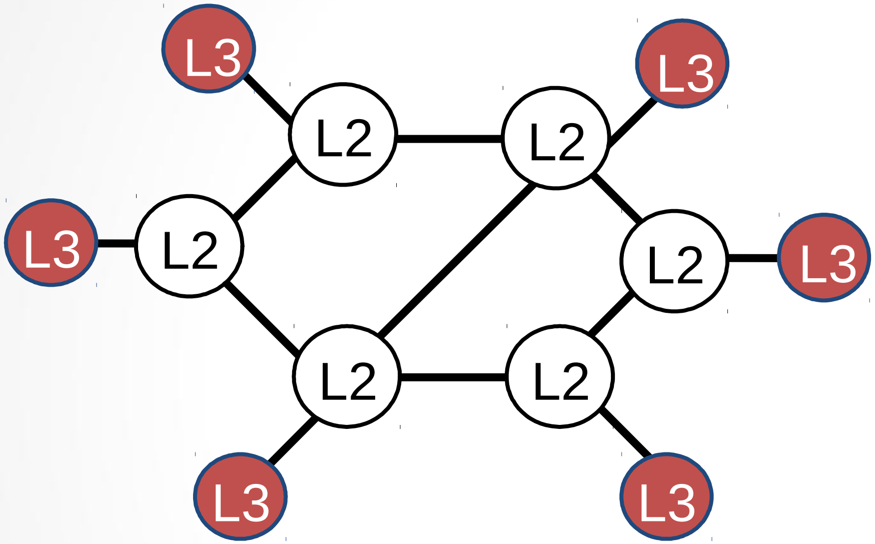

Physical

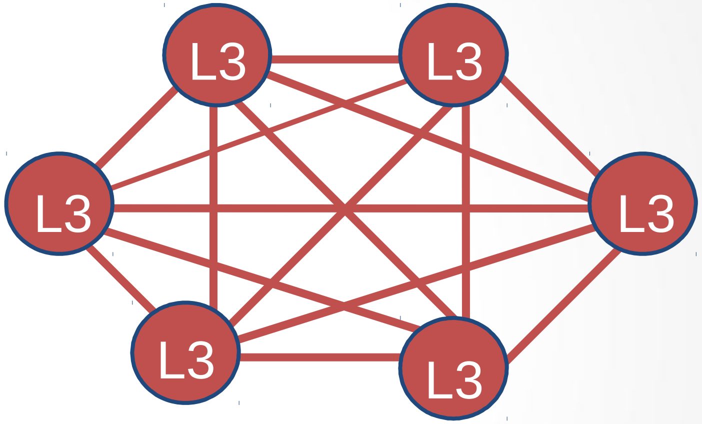

Logical

- The use of the explicit Layer 2 transit layer allows very exact control of how traffic uses the available bandwidth.
- Layer 3 at the edge sees a complete mesh.

Traffic Engineering with a Layer 2 Overlay Model (Cont.)

- Drawbacks of the Overlay Solution
  - – Extra network devices
  - – More complex network management:
    - Two-level network without integrated network management
    - Additional training, technical support, field engineering
  - – IGP routing scalability issue for meshes
  - – Additional bandwidth overhead (“cell tax”)
  - – No differential service (class of service)

Traffic Engineering with a Layer 3 Model

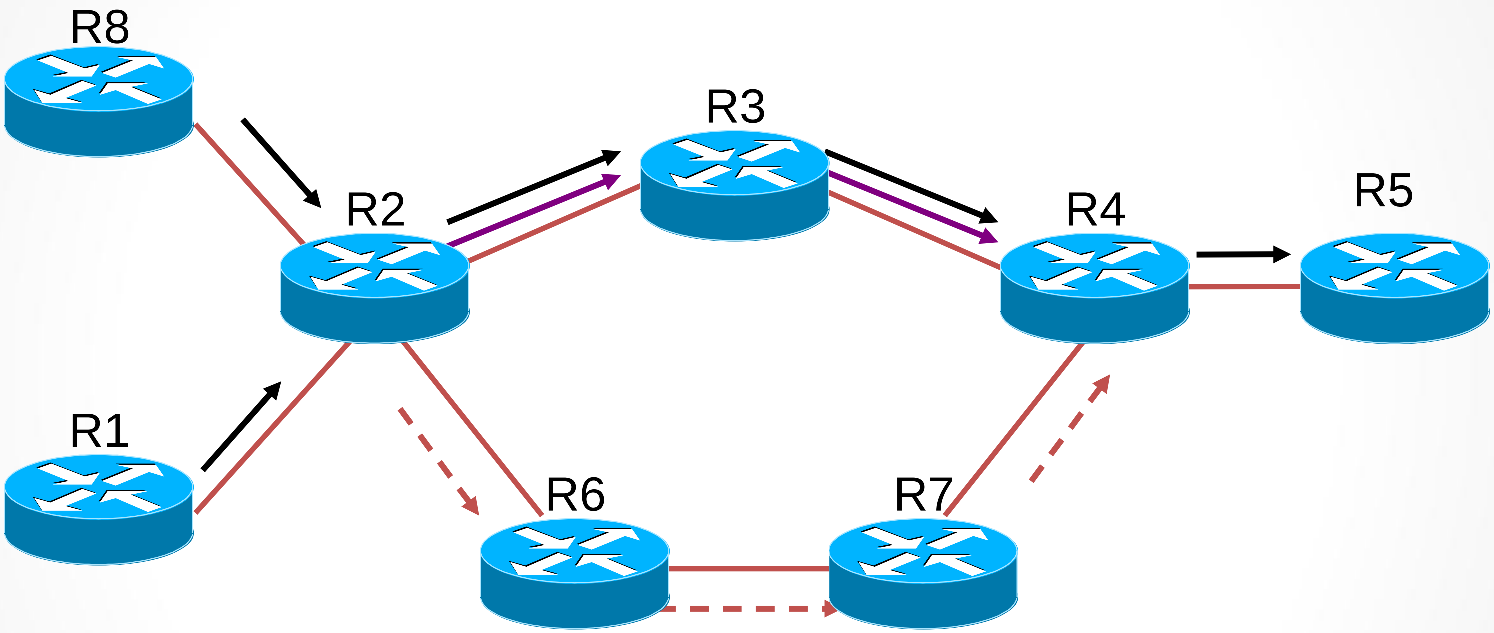

⟶ IP (mostly) uses destination-based least-cost routing. Flows from R8 and R1 merge at R2. From R2, traffic to R3, R4, and R5 uses the upper route.

⇢ The dashed arrow denotes an underutilized alternative path.

Traffic Engineering with a Layer 3 Model (Cont.)

- The current forwarding paradigm, centered around “destination-based,” is clearly inadequate:
  - Path computation based just on IGP metric is not enough.
  - Support for “explicit” routing (source routing) is not available.
  - Supported workarounds: static routes, policy routing.
  - Provide controlled backup and recovery.

Traffic Engineering with the MPLS-TE Model

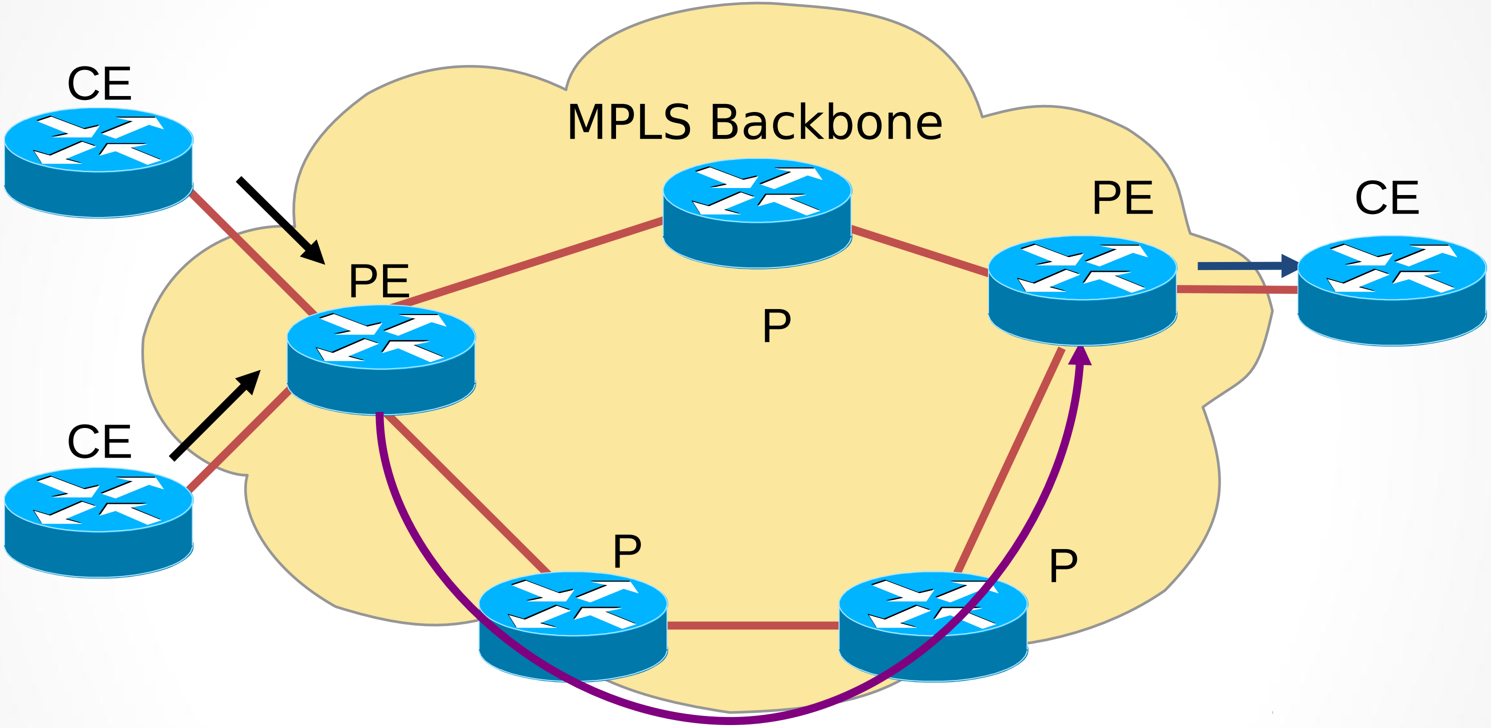

- Tunnel is assigned labels that represent the path (LSP) through the system.
- Forwarding within the MPLS network is based on labels (no Layer 3 lookup).

Traffic Engineering with the MPLS-TE Model (Cont.)

- The MPLS-TE LSPs are created by RSVP.
- The actual path can be specified:
  - Explicitly defined by the system administrator
  - Dynamically defined using the underlying IGP protocol

Traffic Engineering with the MPLS-TE Model

RSVP-Resource ReSerVation Protocol (RFC 2205) - 1993

The source node sends a special message in the RSVP Protocol format over the network before transmitting data that requires a certain non-standard quality of service (for example, constant bandwidth for video transmission). This message contains:

- type of information being transmitted
- bandwidth required.

It is transmitted between routers from the sending node to the destination address, and the sequence of routers in which you want to reserve a certain bandwidth is determined.

- When the router receives this message, it checks its resources.
- If the required bandwidth is achievable, the router configures the packet processing algorithm so that the specified bandwidth is always provided, and then sends the message to the next router along the path.
- In the absence, of the bandwidth the router rejects the request.

Traffic Engineering with the MPLS-TE Model

- The Path Packet reaches the recipient of the stream, who sends back a Resv message, confirming the allocation of resources throughout the path.
- The Original sender, having received Resv, understands that everything is ready for him, and he can send data.

MPLS TE Components

Traffic Tunnels: Concepts

- The concept of traffic tunnels (MPLS-TE tunnels) was introduced to overcome the limitations of hop-by-hop IP routing:
  - A tunnel is an aggregation of traffic flows that are placed inside a common MPLS label switched path.
  - Flows are then forwarded along a common path within a service provider network.

Traffic Tunnels: Concepts (Cont.)

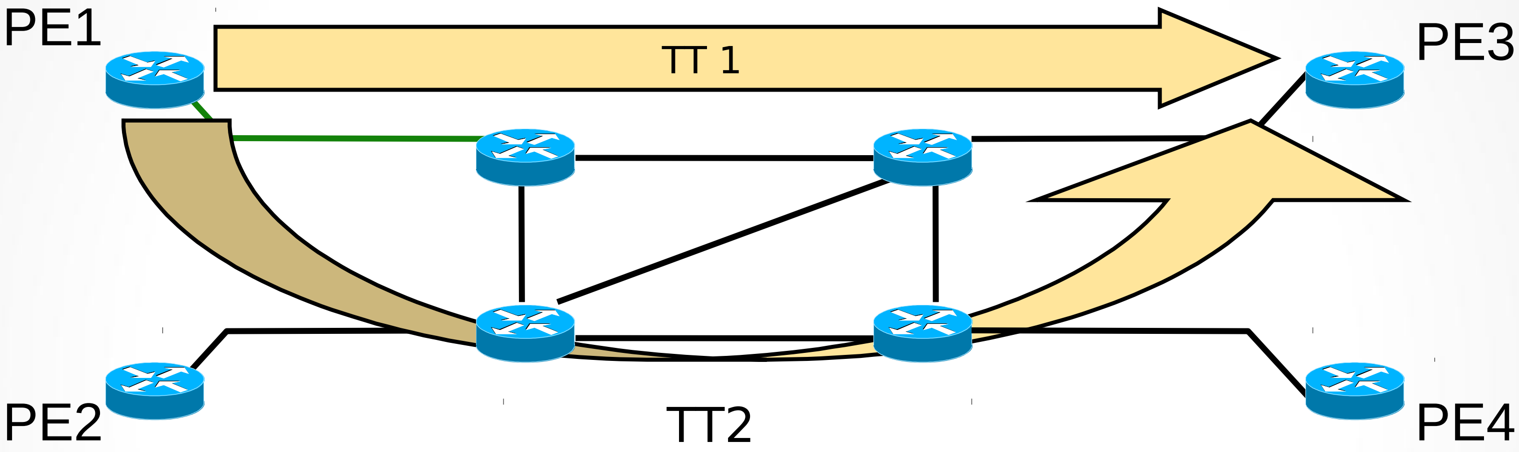

  - Unidirectional single class of service model encapsulates all of the traffic between an ingress and an egress router.
- **Different classes of service model assigns traffic into separate tunnels with different characteristics.**

Traffic Tunnels – Characteristics

- Traffic tunnels are routable objects (similar to ATM VCs).
- A traffic tunnel is distinct from the MPLS LSP through which it traverses:
  - In operational contexts, a traffic tunnel can be moved from one path onto another
- A traffic tunnel is assigned attributes influencing its characteristics.

Traffic Tunnels – Attributes

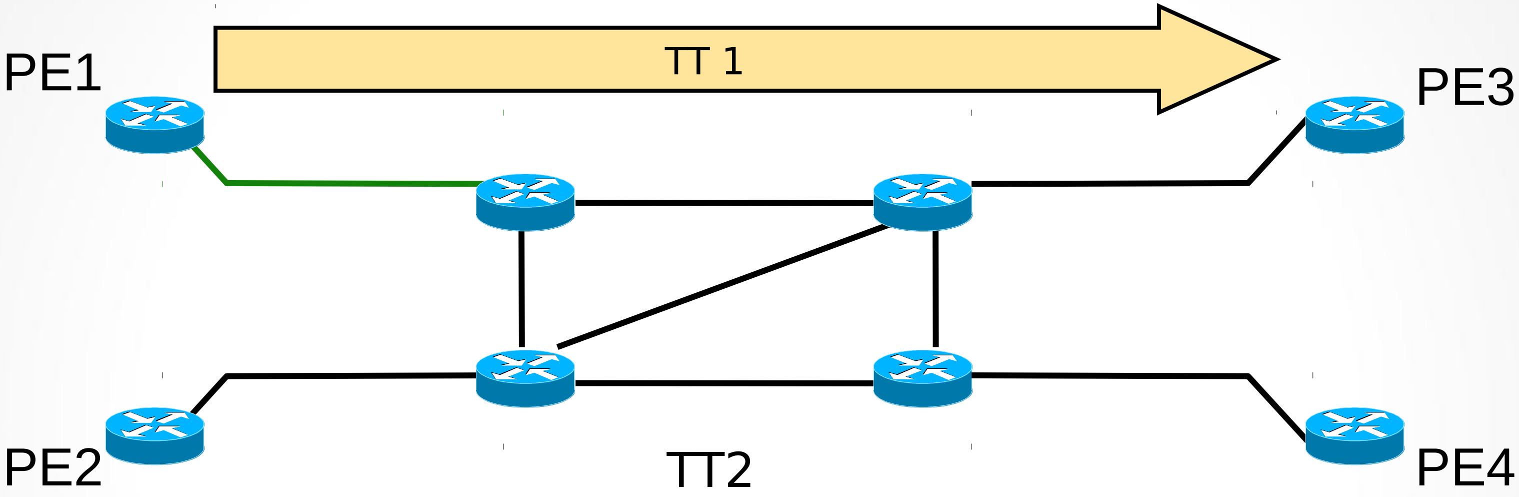

- Attributes are explicitly assigned to traffic tunnels through administrative action.
- A traffic tunnel is characterized by:
  - Its ingress and egress label switch routers
  - The forwarding equivalence class that is mapped onto it
  - A set of attributes that determine its characteristics

Traffic Tunnels – Attributes (Cont.)

- The administrator enters the relevant information (attributes) at the headend of the traffic tunnel:
  - Traffic parameter—resources required for tunnel (e.g., required bandwidth)
  - Generic path selection and management—path can be administratively specified or computed by the IGP
  - Resource class affinity—include or exclude certain links for certain traffic tunnels
  - Adaptability—should the traffic tunnel be reoptimized?
  - Priority and pre-emption—importance of a traffic tunnel and possibility for a pre-emption of another tunnel
  - Resilience—desired behavior under fault conditions

Network Links and Link Attributes

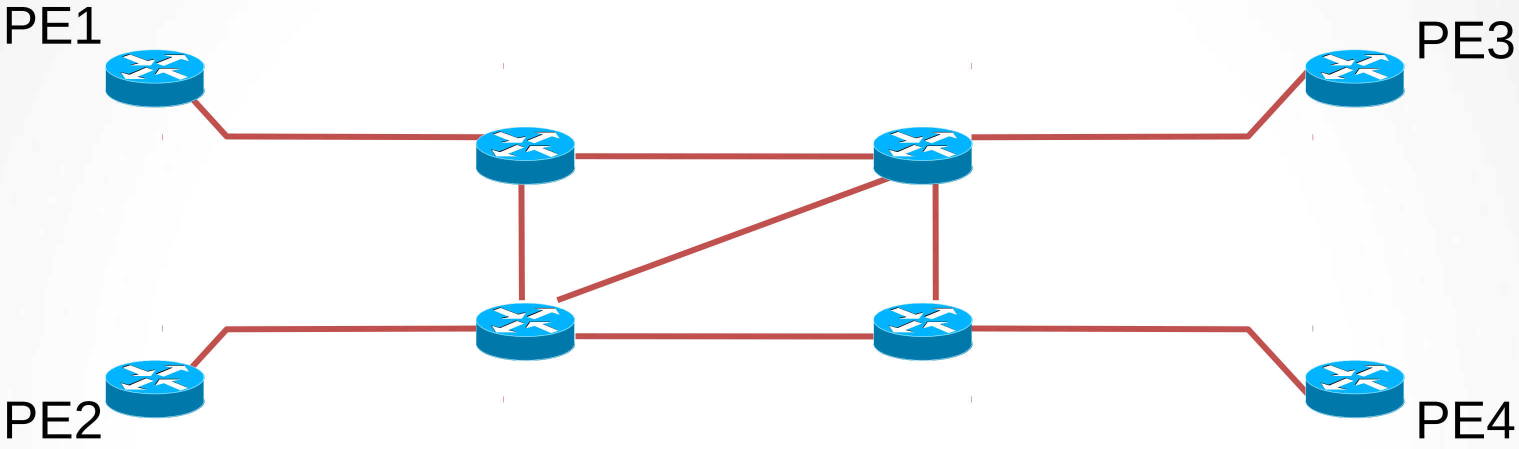

- Resource attributes (link availability) are configured locally on the router interfaces:
  - Maximum bandwidth
    - The amount of bandwidth available
  - Link affinity string
    - To allow the operator to administratively include or exclude links in path calculations
  - Constraint-based specific metric
    - Traffic engineering default metric

Constraint-Based Path Computation

- Constraint-based routing is demand-driven.
- Resource-reservation-aware routing paradigm:
  - Based on criteria including, but not limited to, network topology
  - Calculated at the edge of a network:
    - Modified Dijkstra’s algorithm at tunnel headend (CSPF [constrained SPF] or PCALC [Path Calculation]).
    - Output is a sequence of IP interface addresses (next-hop routers) between tunnel endpoints.

Constraint-Based Path Computation (Cont.)

- Constraint-based routing takes into account:
  - Policy constraints associated with the tunnel and physical links
  - Physical resource availability
  - Network topology state
- Two types of tunnels can be established across those links with matching attributes:
  - Dynamic—using the least-cost path computed by OSPF/IS-IS
  - Explicit—definition of a path by using OS configuration commands

Constraint-Based Path Computation (Cont.)

What is the best path from R1 to R6 with bandwidth of 30 Mbps?

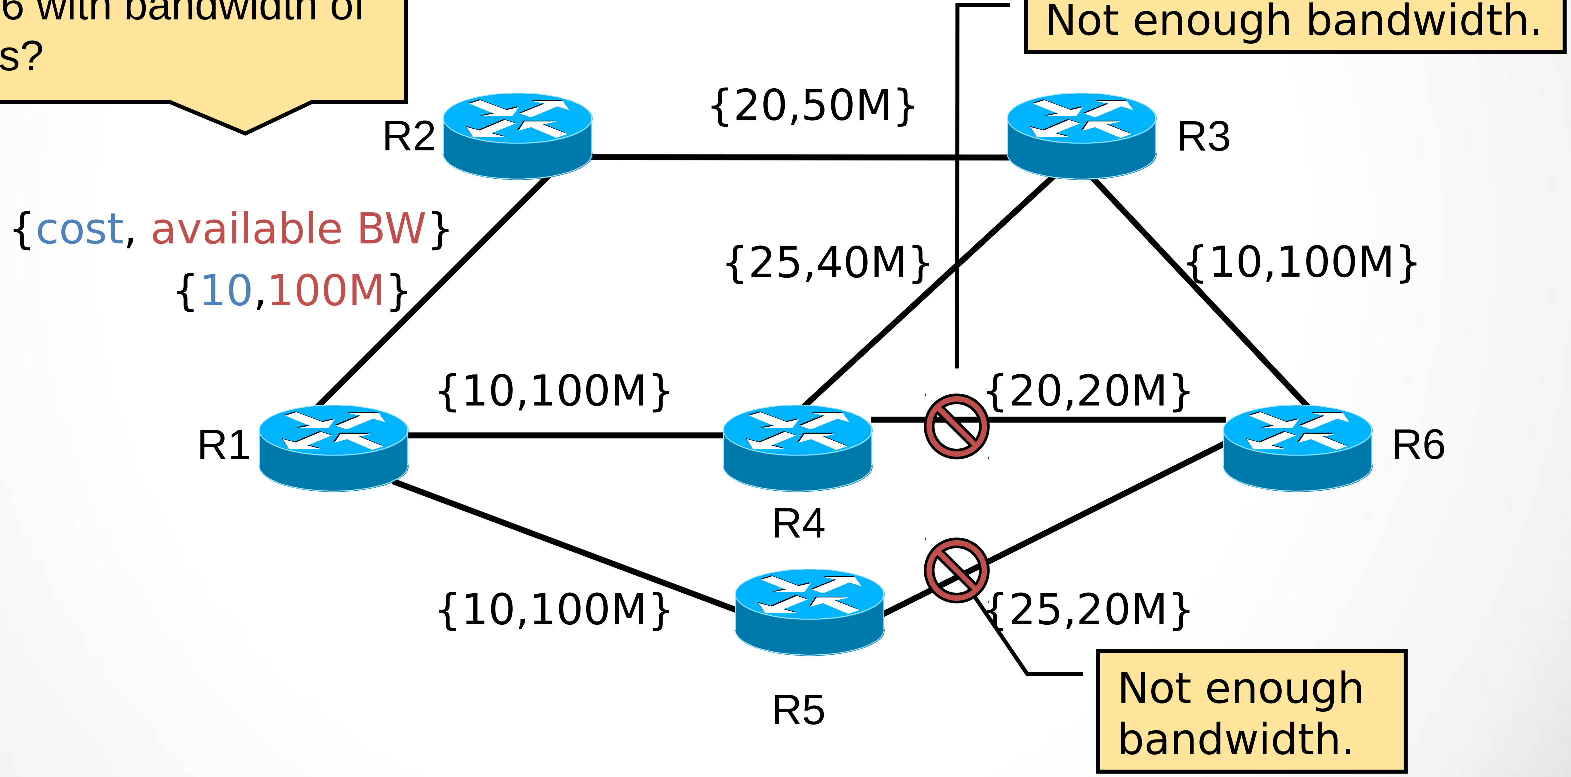

Not enough bandwidth.

Not enough bandwidth.

Constraint-Based Path Computation (Cont.)

Computed path for a dynamic constraint-based tunnel over the least-cost path.

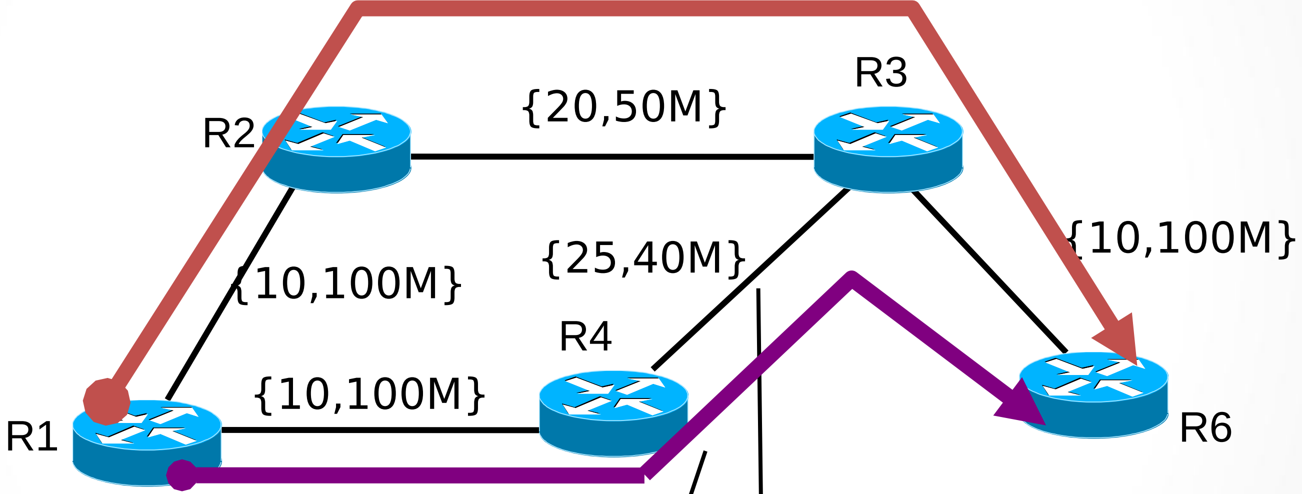

Administratively defined explicit path Tunnel is still possible over any eligible path.

Path has cost of 45, not the lowest cost.

Explicit and Dynamic Traffic Engineering Tunnels

Role of RSVP in Path Setup Procedures

- Once the path has been determined, a signaling protocol is needed:
  - To establish and maintain label switched paths (LSPs) for traffic tunnels
  - For creating and maintaining resource reservation states across a network (bandwidth allocation)
- The Resource Reservation Protocol (RSVP) was adopted by the MPLS workgroup of the IETF.

Role of RSVP in Path Setup Procedures

The goal of RSVP-TE is the same as that of LDP - to distribute the labels between the LSR and compile the resulting LSP from the recipient to the sender.

RSVP TE allows you:

- to build a primary and backup LSP,
- reserve resources on all nodes,
- detect network accidents,
- build pre-workarounds,
- do fast traffic redirection,
- avoid channels that physically pass through the same path.

- LSP - unidirectional, resources will be reserved only in one direction.

Forwarding Table Modifications

- IP routing is separate from LSP routing and does not see internal details of the LSP.
- The traffic has to be mapped to the tunnel:
  - Static routing—the static route in the IP routing table points to an LSP tunnel interface.
  - Policy routing—the next-hop interface is an LSP tunnel
  - Forwarding adjacency—the tunnel is announced as a point-to-point link to all other routers within an area
  - Autoroute—SPF enhancement:
    - The headend sees the tunnel as a directly connected interface (for modified SPF only).
    - The default cost of a tunnel is equal to the shortest IGP metric regardless of the used path.

Constraint-Based Path Computation

Constraint-Based Path Computation

- Constraint-based path computation provides several resource attributes to control LSP path determination.
  - Link resource attributes that provide information on the resources of each link.
  - Traffic tunnel attributes characterize the traffic tunnel.

MPLS-TE Link Resource Attributes

- Maximum bandwidth
- Maximum reservable bandwidth
- Link resource class
- Constraint-based specific link metric

MPLS-TE Link Resource Attributes: Maximum Allocation Bandwidth

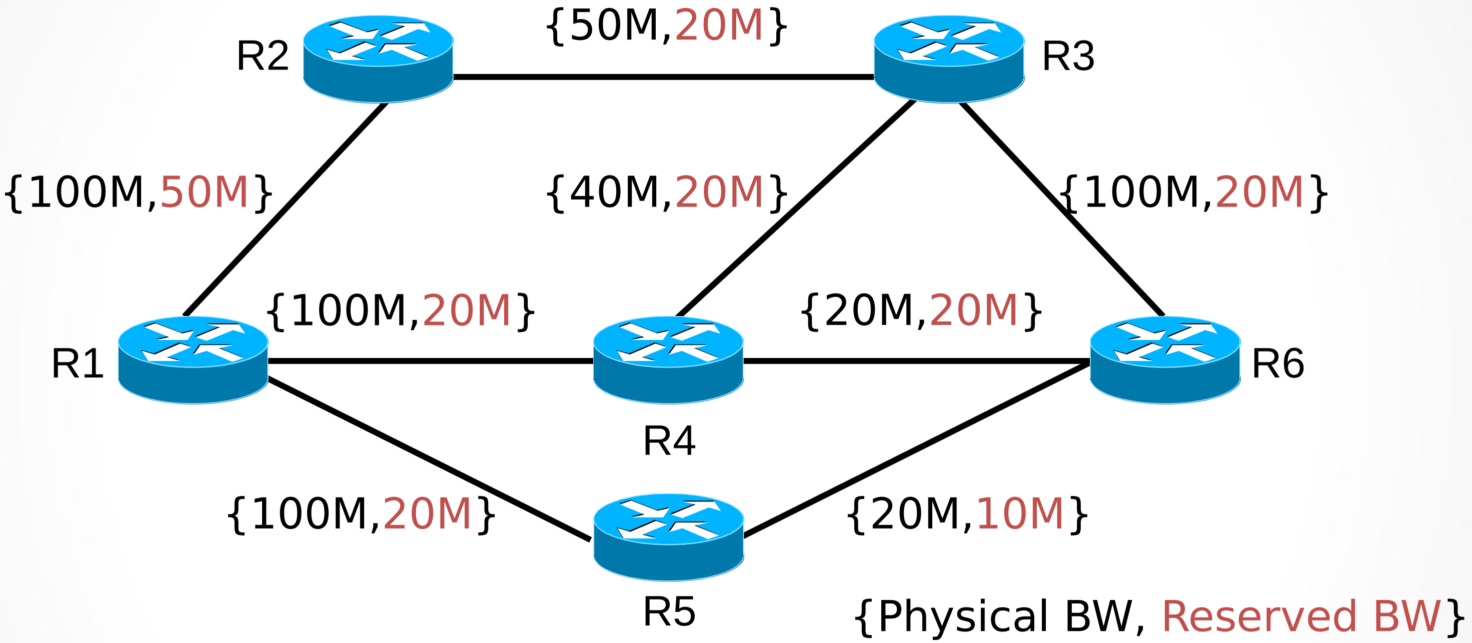

- Maximum bandwidth: the maximum bandwidth that can be used on this link in this direction (physical link)
- Maximum reservable bandwidth: The maximum amount of bandwidth that can be reserved in this direction on this link

MPLS-TE Link Resource Attributes: Link Resource Class

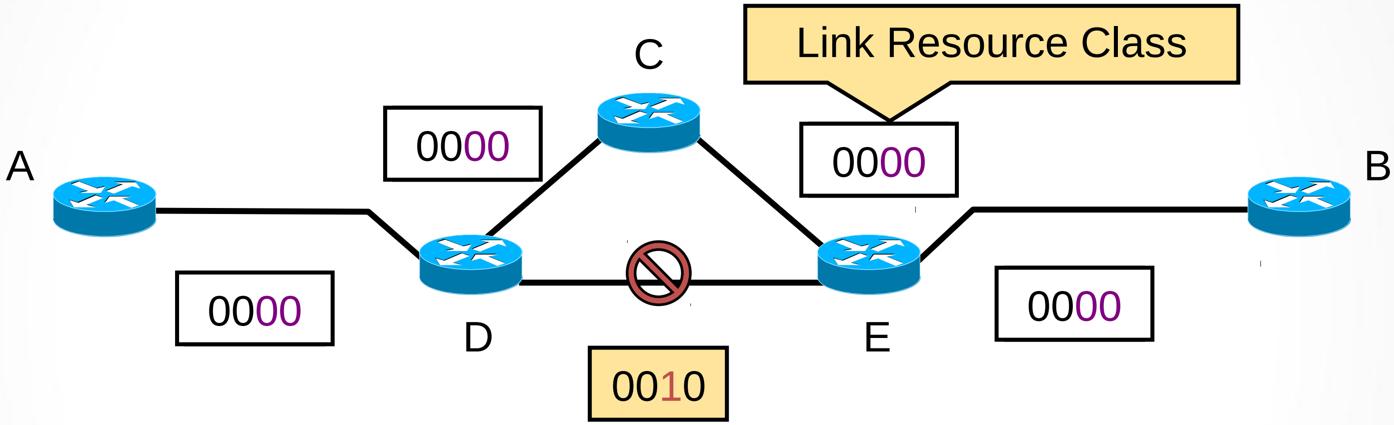

– Link is characterized by a 32-bit resource class attribute.

– Associated with a traffic tunnel in order to include or exclude certain links into or from the path of the traffic tunnel.

MPLS-TE Link Resource Attributes: Constraint-Based Specific Link Metric

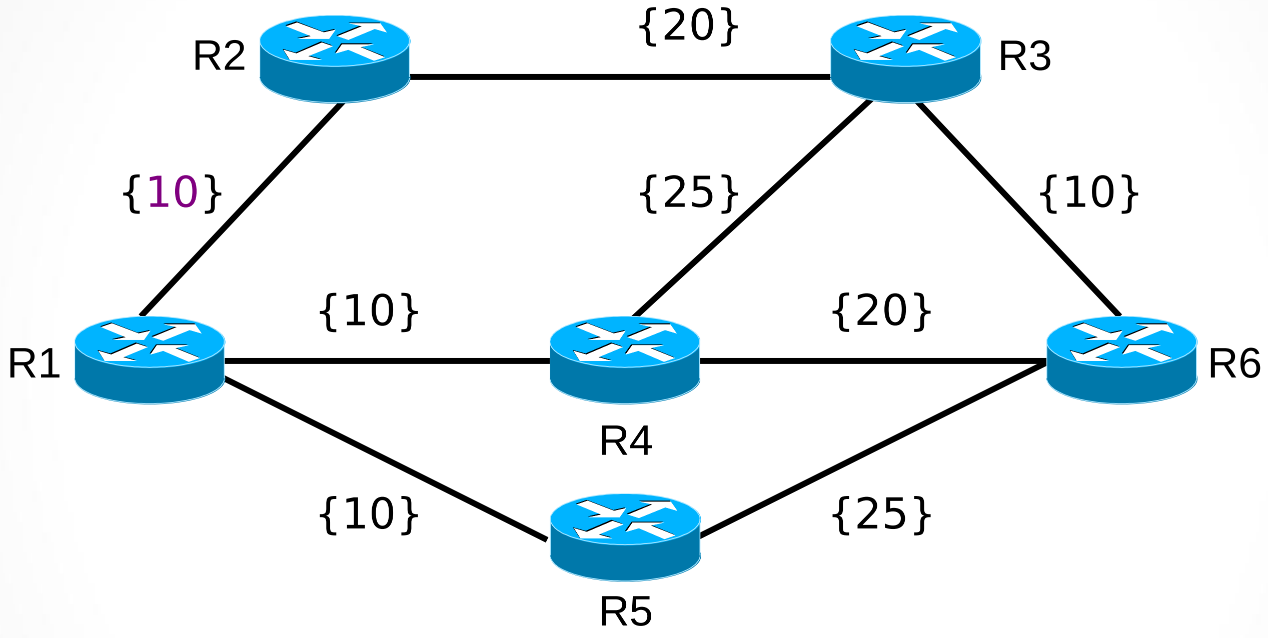

- This metric is administratively assigned to present a differently weighted topology to traffic engineering SPF calculations:
  - » Administrative weight (TE metric)

MPLS-TE Tunnel Attributes

- – Traffic parameter
- – Generic path selection and management
- – Tunnel resource class affinity
- – Adaptability
- – Priority
- – Pre-emption
- – Resilience

MPLS-TE Tunnel Attributes (Cont.)

- Traffic parameter:
  - Indicates the resource requirements (for example, bandwidth) of the traffic tunnel
- Generic path selection and management:
  - Specifies how the path for the tunnel is computed:
    - Dynamic LSP — Constraint-based computed paths based on a combination of bandwidth and policies
    - Explicit LSP — administratively specified off line (typically using CLI)

MPLS-TE Tunnel Attributes (Cont.)

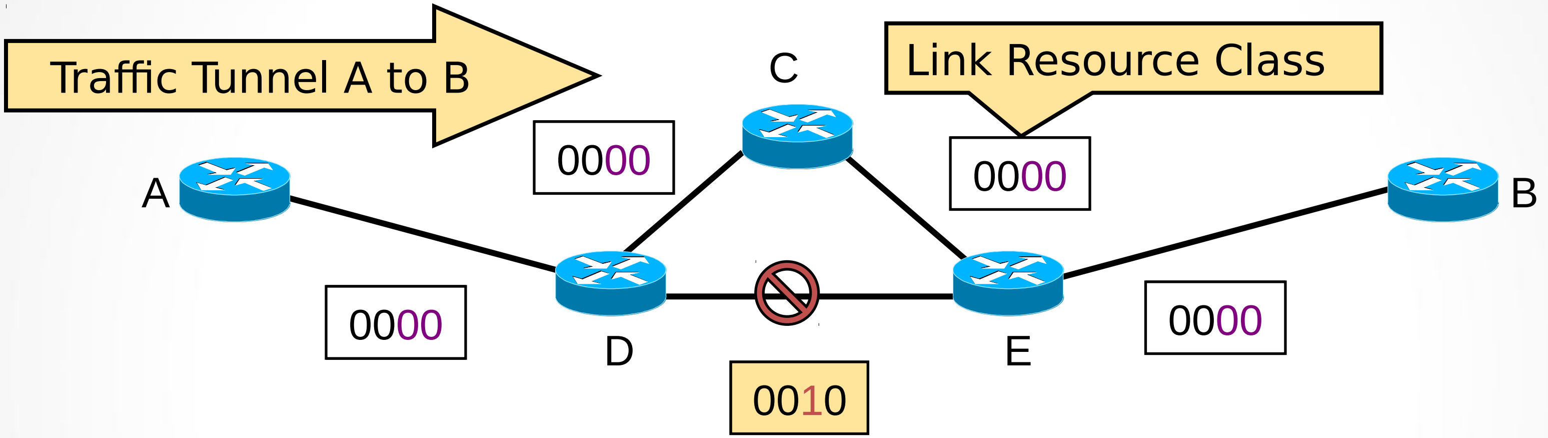

- Tunnel Resource Class Affinity:
  - The properties that the tunnel requires from internal links:
    - 32-bit resource class affinity bit string + 32-bit resource class mask
  - Link is included in the constraint-based LSP path when the tunnel resource affinity string or mask matches the link resource class attribute.

MPLS-TE Tunnel Attributes (Cont.)

- Adaptability:
  - If reoptimization is enabled, then a traffic tunnel can be rerouted through different paths by the underlying protocols:
    - Primarily due to changes in resource availability
- Priority:
  - Relative importance of traffic tunnels
  - Determines the order in which path selection is done for traffic tunnels at connection establishment and under fault scenarios:
    - Setup priority: Priority for taking a resource
- Pre-emption:
  - Determines whether another traffic tunnel can pre-empt a specific traffic tunnel:
    - Hold priority: Priority for holding a resource

MPLS-TE Tunnel Attributes (Cont.)

- Resilience:
  - Determines the behavior of a traffic tunnel under fault conditions:
    - Do not reroute the traffic tunnel
    - Reroute through a feasible path with enough resources
    - Reroute through any available path regardless of resource constraints

Implementing TE Policies with Affinity Bits

- Link is characterized by the link resource class
  - Default value of bits is 0
- Tunnel is characterized by:
  - Tunnel resource class affinity
    - Default value of bits is 0
  - Tunnel resource class affinity mask
    - (0=do not care, 1=care)
    - Default value of the tunnel mask is 0x0000FFFF

Implementing TE Policies with Affinity Bits (Cont.)

Setting a link bit in the lower half drives all tunnels off the link, except those specially configured.

Tunnel Affinity: bits = 0000, mask = 0011

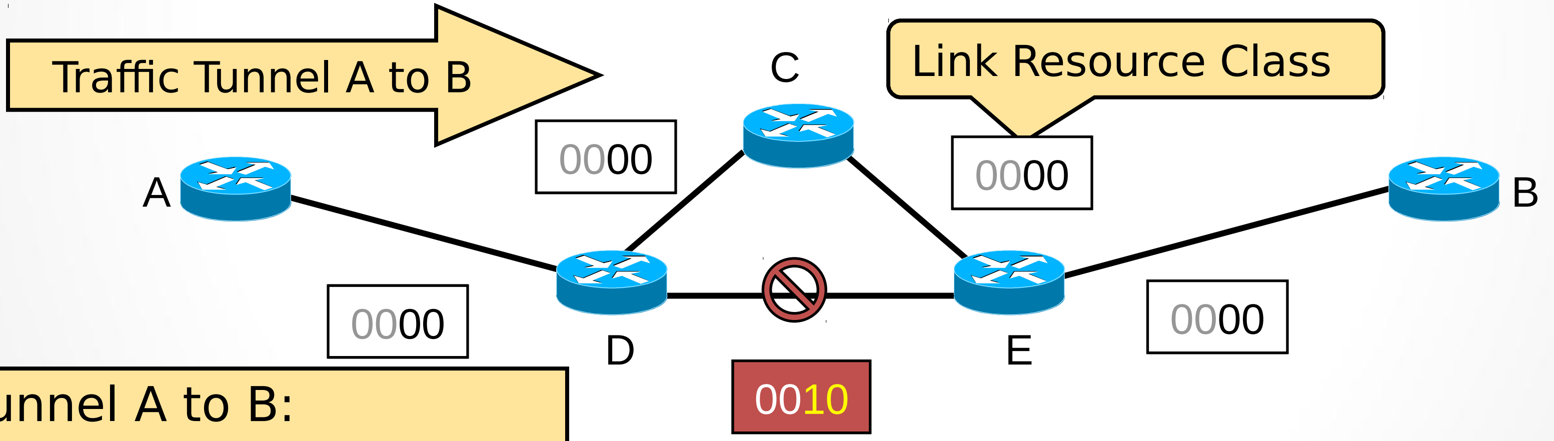

Tunnel A to B:

Only ADCEB is possible.

Using Affinity Bits to Avoid Specific Links

Implementing TE Policies with Affinity Bits (Cont.)

A specific tunnel can then be configured to allow all links by clearing the bit in its affinity attribute mask.

Tunnel Affinity: bits = 0000, mask = 0001

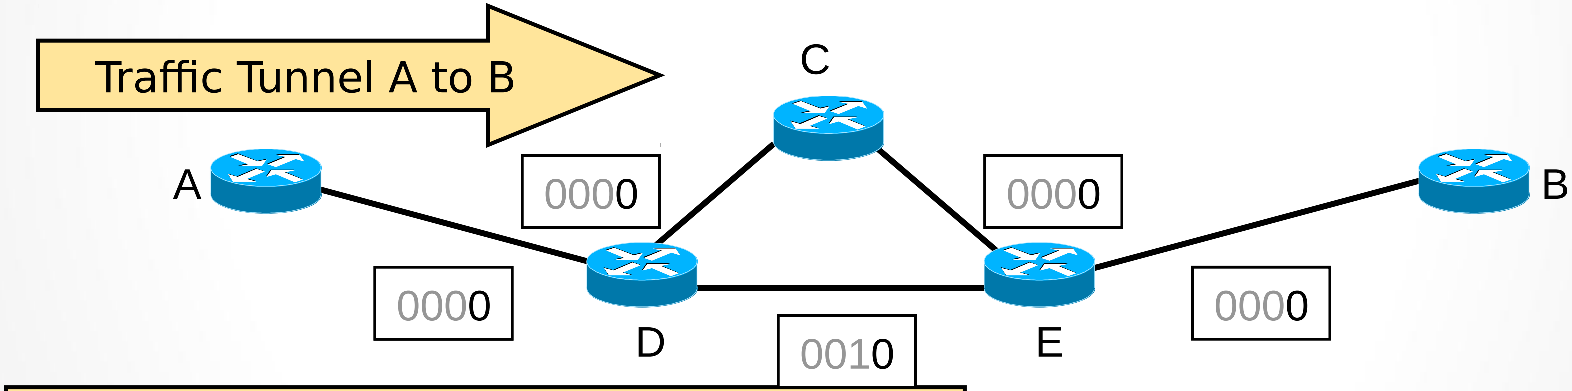

Tunnel A to B:

Again, ADEB and ADCEB are possible.

Using the Affinity Bit Mask to Allow all Links

Implementing TE Policies with Affinity Bits (Cont.)

A specific tunnel can be restricted to only some links by turning on the bit in its affinity attribute bits.

Tunnel Affinity: bits = 0010, mask = 0011

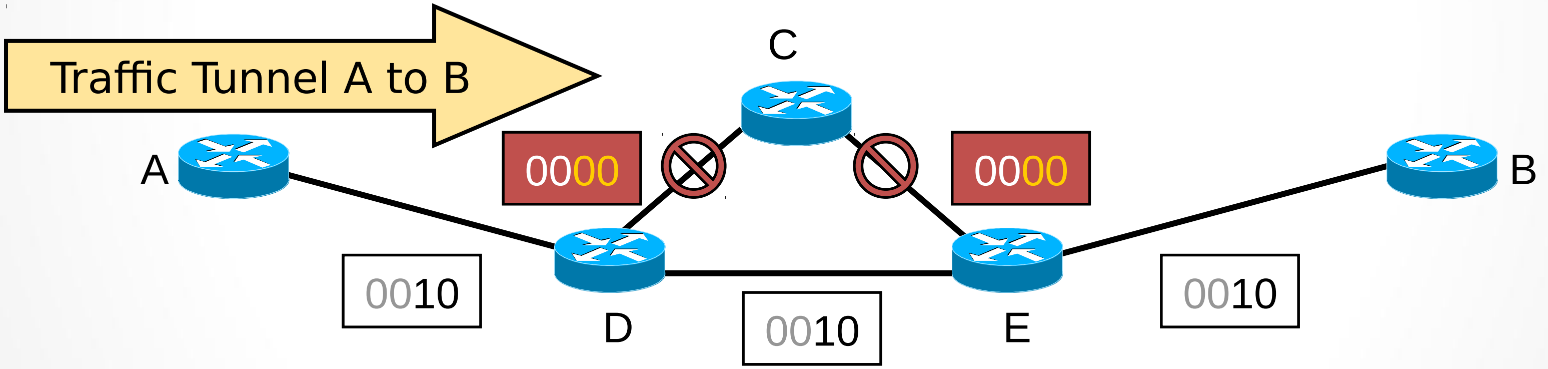

Tunnel A to B:

ADEB is possible.

Using Affinity Bits to Dedicate Links to Specific Purposes

Propagating MPLS-TE Link Attributes with Link-State Routing

- IGP resource flooding takes place in the following situations:
  - Link-state changes
  - Resource class of a link changes:
    - Manual reconfiguration
    - Amount of available bandwidth crosses one of the preconfigured thresholds
  - Periodic (timer-based):
    - A node checks attributes; if they are different, it floods its update status
  - On LSP setup failure

Constraint-Based Path Computation

- When establishing a tunnel, the edge routers have knowledge of both network topology and link resources within its area:
  - Two methods for establishing traffic tunnels:
    - Static and dynamic path setup
  - In both cases the result is an explicit route expressed as a sequence of interface IP addresses (for numbered links) or TE router IDs (for unnumbered links) in the path from tunnel endpoints.
  - RSVP is used to establish and maintain constraint-based label switched paths for traffic tunnels along an explicit path.

Constraint-Based Path Computation (Cont.)

Constraint-Based Path Selection

- Path selection:
  - CBR uses its own metric (administrative weight, or TE cost; by default equal to the IGP cost)—used only during constraint-based computation
  - In case of equal cost, select the path with:
    - The highest minimum bandwidth
    - The smallest hop count
    - If everything else fails, then pick a path at random
- LSP path setup—an explicit path is used by RSVP to reserve resources and establish LSP path
- Final result: unidirectional MLPS-TE tunnel, seen only at the headend router

Constraint-Based Path Computation (Cont.)

Request by tunnel:

From R1 to R6; Priority 3, BW = 30 Mbps

Resource Affinity: bits = 0010, mask = 0011

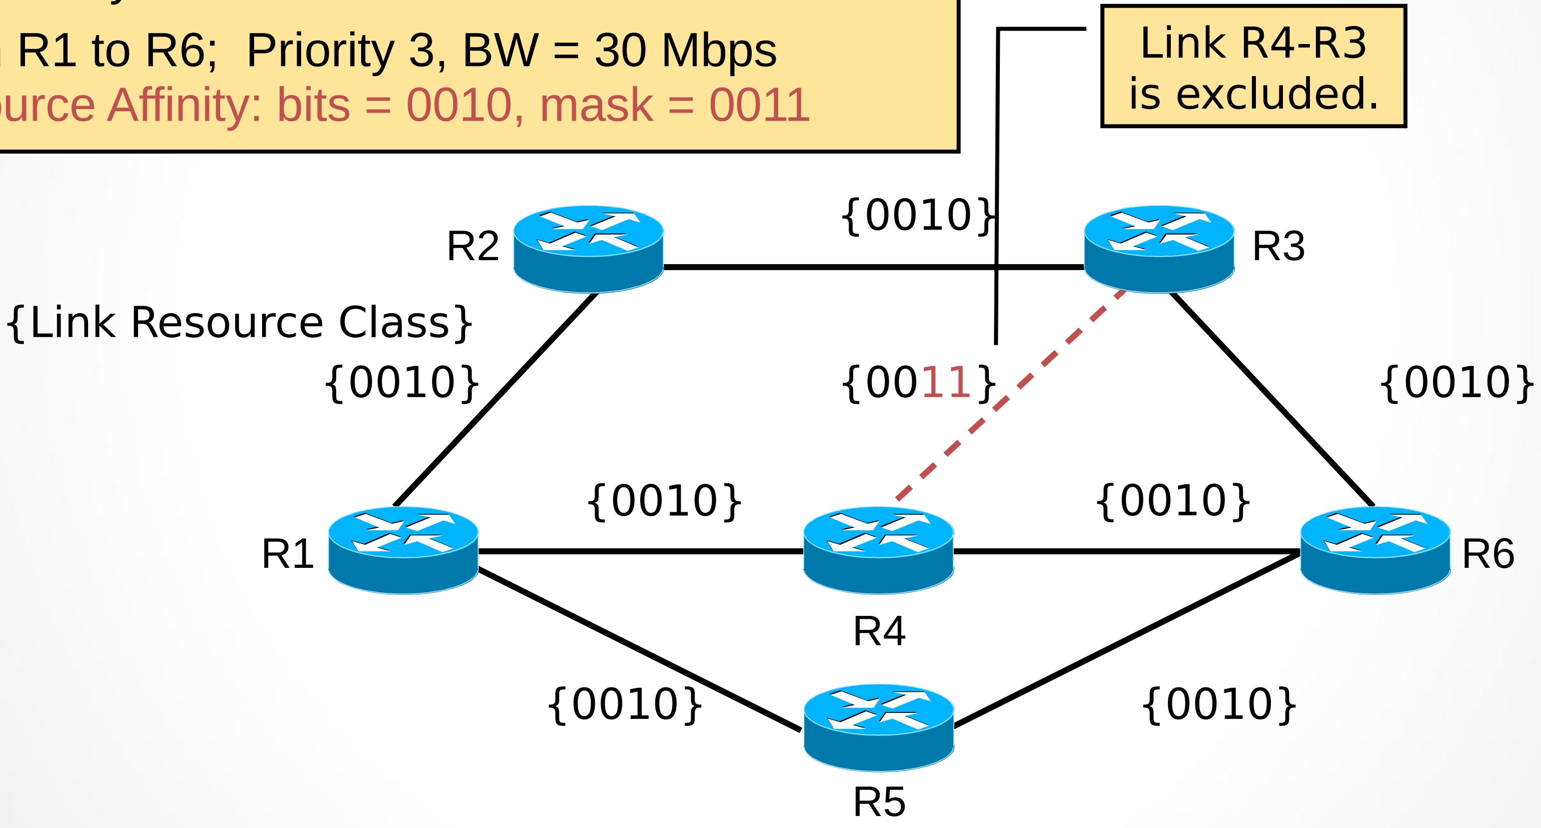

Path Selection Considering Policy Constraints

Constraint-Based Path Computation (Cont.)

Request by tunnel:

From R1 to R6; Priority 3, BW = 30 Mbps

Resource Affinity: bits = 0010, mask = 0011

Not enough bandwidth

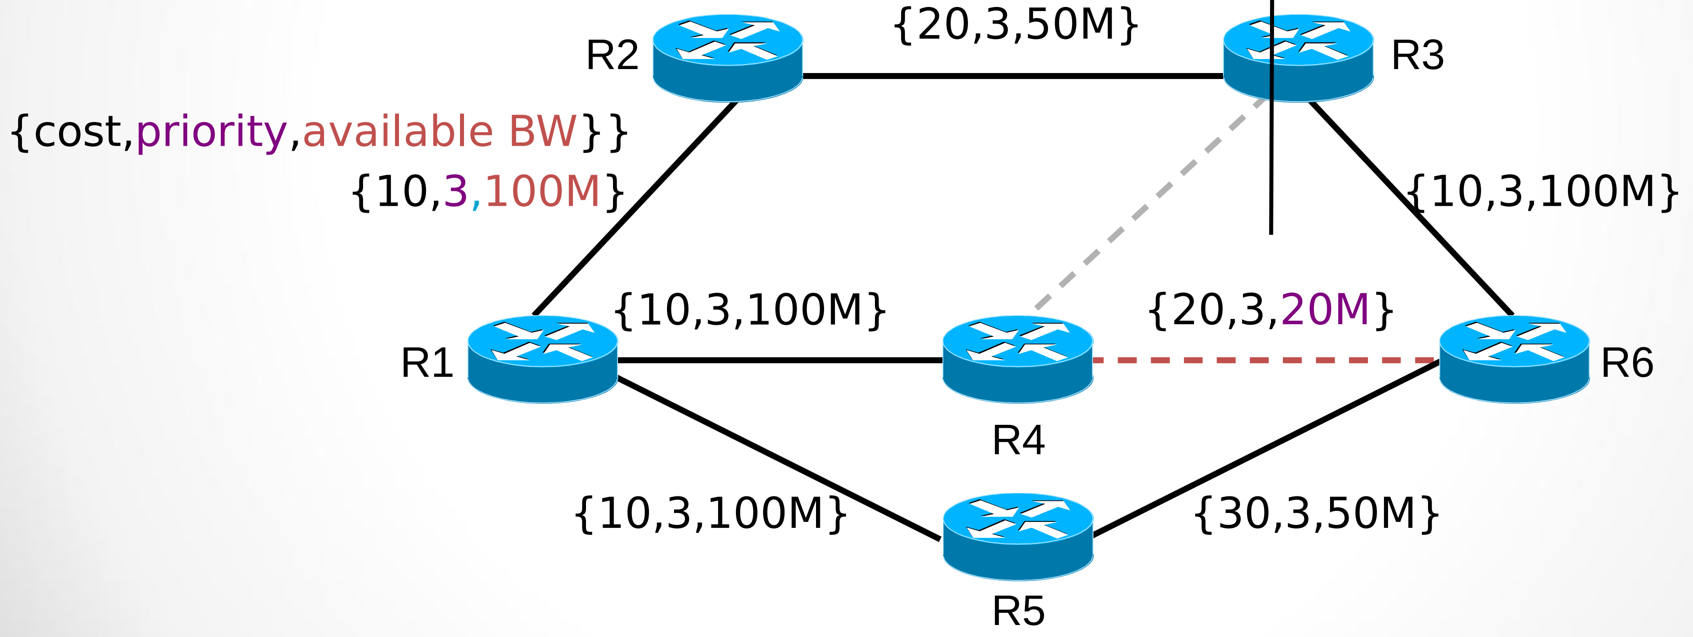

Path Selection Considering Available Resources

Constraint-Based Path Computation (Cont.)

The headend router has two possible paths with a total cost of 40: R1 – R2 – R3 – R6 and R1 – R5 – R6, both offering at least 50 Mbps (minimim bandwidth). Because of the smaller hop count, R1 – R5 – R6 is selected.

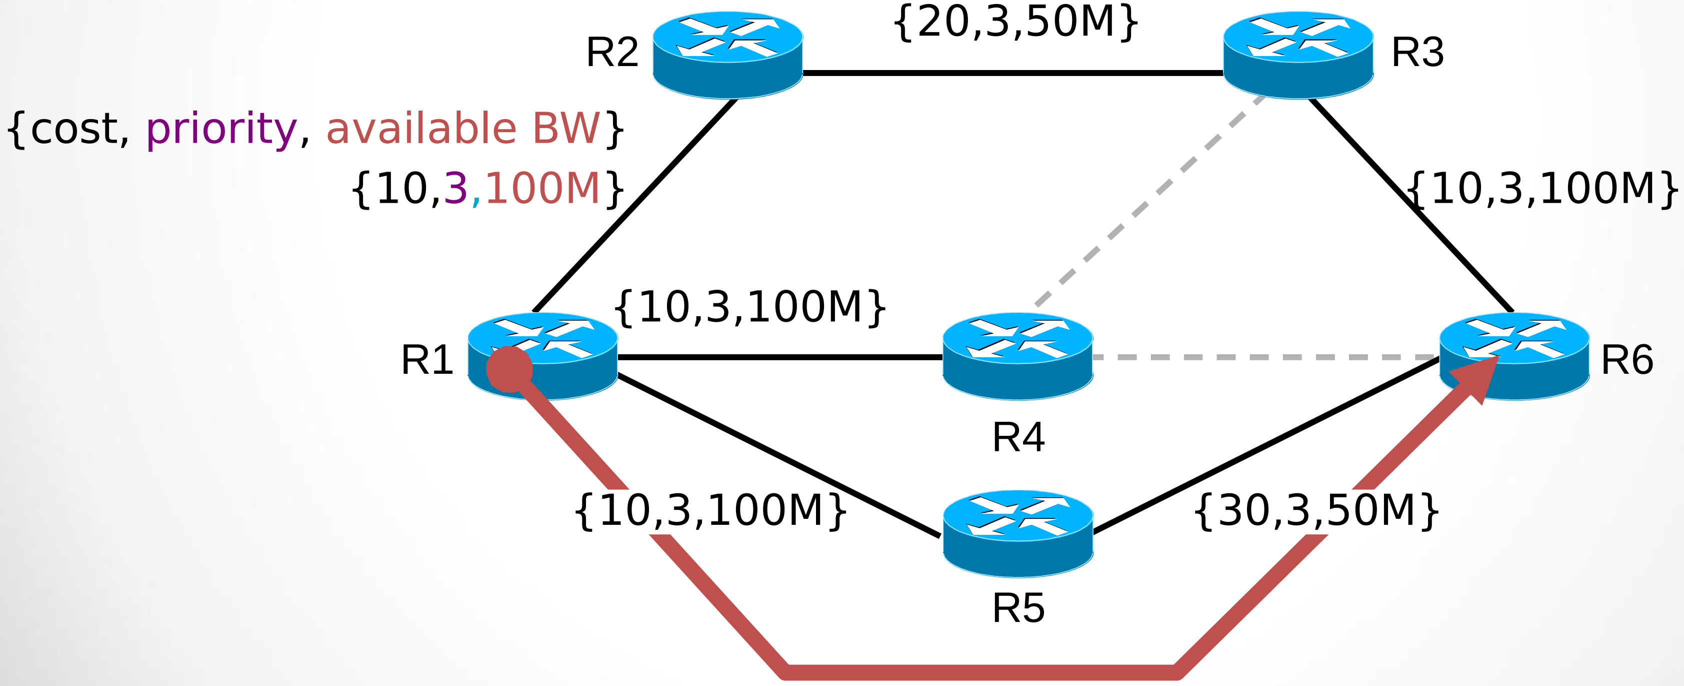

Selecting the Best Path

Path Setup and Maintenance

Path Setup

- LSP path setup is initiated at the headend of a tunnel.
- The route (list of next-hop routers) is either:
  - Statically defined
  - Computed by CBR
- The route is used by RSVP to:
  - Assign labels
  - Reserve bandwidth on each link
- Tunnel attributes that affect path setup:
  - Bandwidth
  - Priority
  - Affinity attributes

RSVP Usage in Path Setup

- RSVP makes resource reservations for both unicast and multicast applications:
  - RSVP provides support for dynamic membership changes and automatic adaptation to routing changes.
  - RSVP sends periodic refresh messages to maintain the state along the reserved path.
  - RSVP sessions are used between routers, not hosts.

Hop-by-Hop Path Setup with RSVP

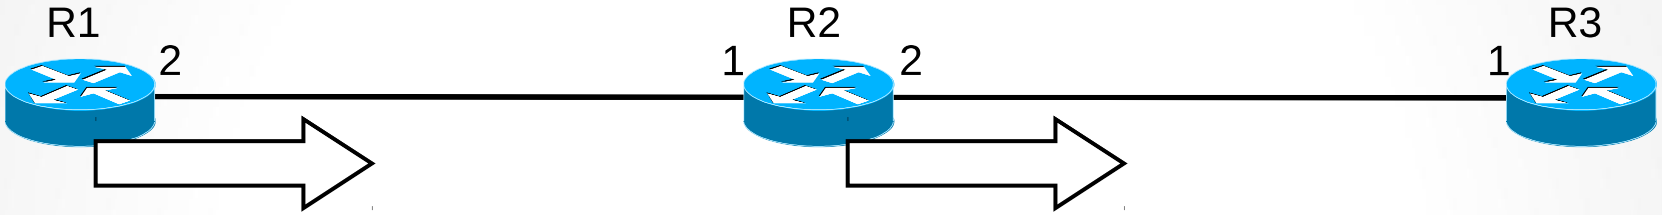

Path:
Common_Header
Session(R3-lo0, 0, R1-lo0)
PHOP(R1-2)
Label_Request(IP)
ERO (R2-1, R3-1)
Session_Attribute (...)
Sender_Template(R1-lo0, 00)
Record_Route(R1-2)

Path:
Common_Header
Session(R3-lo0, 0, R1-lo0)
PHOP(R2-2)
Label_Request(IP)
ERO (R3-1)
Session_Attribute (...)
Sender_Template(R1-lo0, 00)
Record_Route (R1-2, R2-2)

Hop-by-Hop Path Setup with RSVP (Cont.)

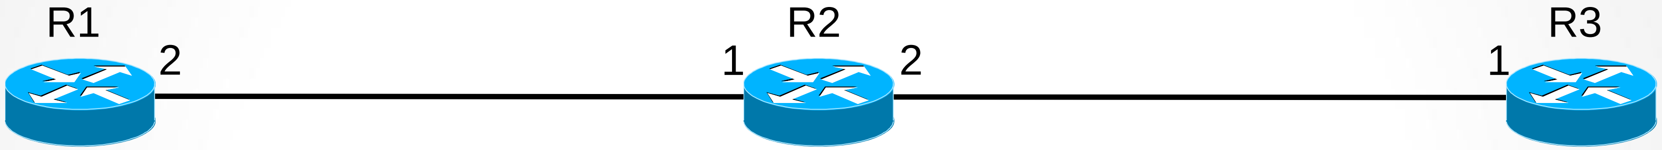

Path State:
Session(R3-lo0, 0, R1-lo0)
PHOP(R2-2)
Label_Request(IP)
ERO ()
Session_Attribute (...)
Sender_Template(R1-lo0, 00)
Record_Route (R1-2, R2-2, R3-1)

Hop-by-Hop Path Setup with RSVP (Cont.)

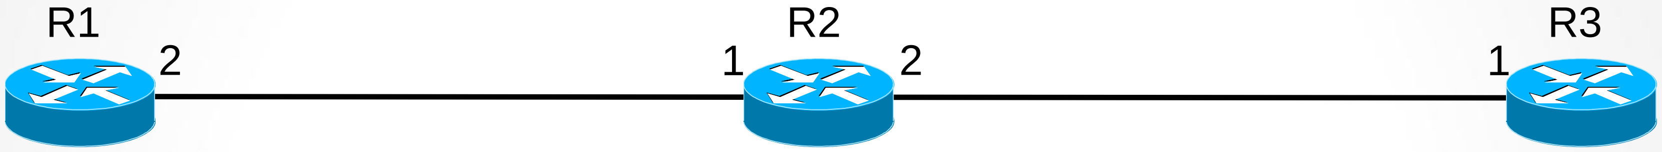

Resv:
Common_Header
Session(R3-lo0, 0, R1-lo0)
PHOP(R2-1)
Sender_Template(R1-lo0, 00)
Label=25
Record_Route(R2-1, R3-1)

Resv:
Common_Header
Session(R3-lo0, 0, R1-lo0)
PHOP(R3-1)
Sender_Template(R1-lo0, 00)
Label=POP
Record_Route(R3-1)

Hop-by-Hop Path Setup with RSVP (Cont.)

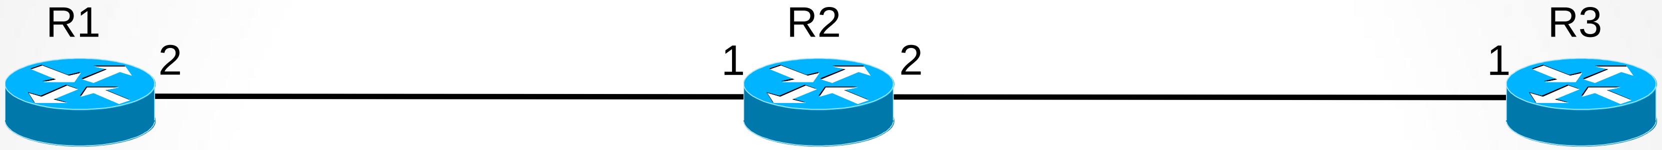

Resv state:
Session(R3-lo0, 0, R1-lo0)
PHOP(R2-1)
Sender_Template(R1-lo0, 00)
Label=5
Record_Route(R1-2, R2-1, R3-1)

Tunnel and Link Admission Control

- Invoked by RSVP Path message:
  - Determines if resources are available
  - If bandwidth is not available:
    - Link-level call admission control (LCAC) says no to RSVP
    - PathErr message is sent
  - If bandwidth is available, this bandwidth is put aside in a waiting pool (waiting for the Resv message):

Tunnel and Link Admission Control (Cont.)

- Pre-emption
  - The process of LSP path setup may require the pre-emption of resources.
  - LCAC notifies RSVP of the pre-emption.
  - RSVP sends PathErr or ResvErr or both for the pre-empted tunnel.

Path Reoptimization

- Problem: Some resources become available, which results in a nonoptimal path of traffic tunnels
- Solution: Reoptimization:
  - A periodic timer checks for the most optimal path
  - If a better LSP seems to be available:
    - The device attempts to signal the better LSP
    - If successful, replaces the old and inferior LSP with the new and better LSP

Path Reoptimization (Cont.)

Nondisruptive Rerouting — Reoptimization

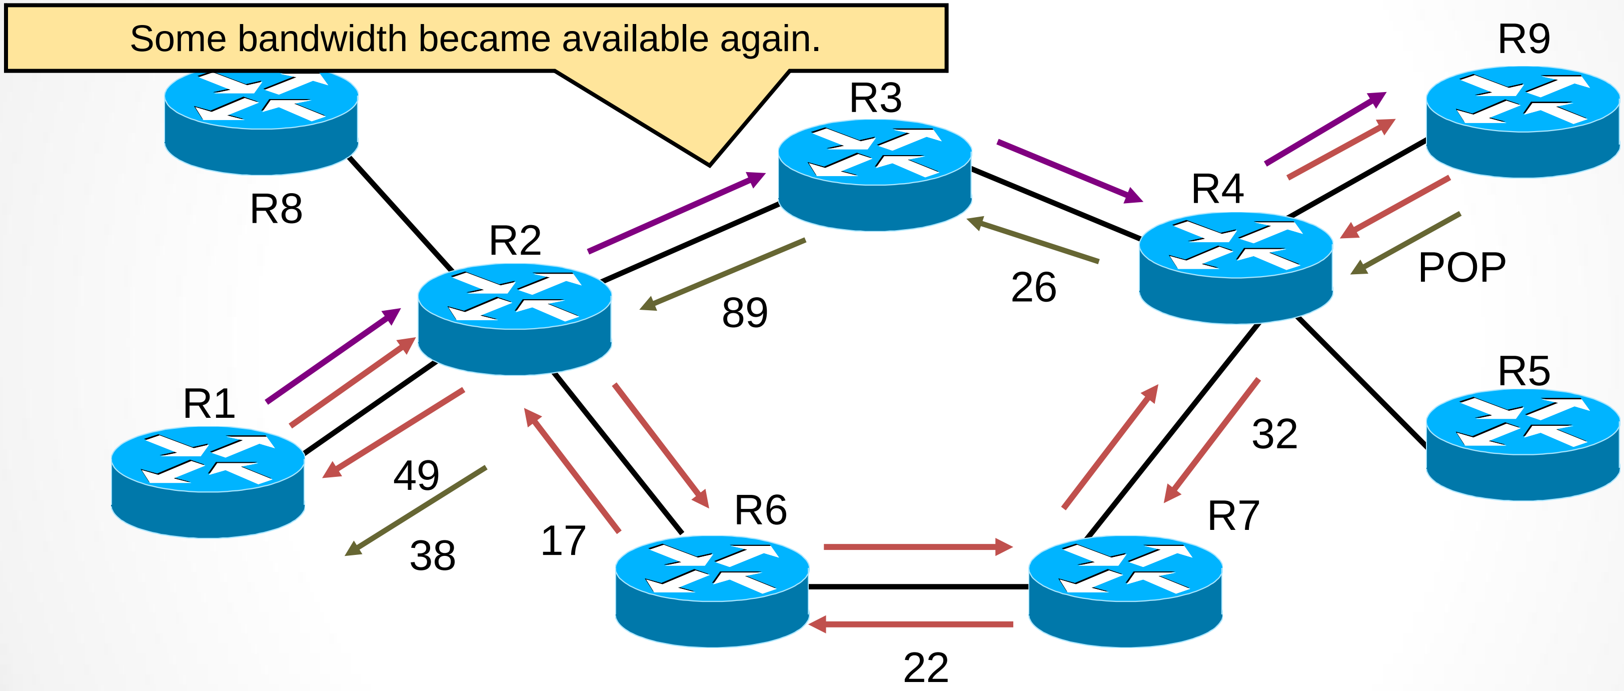

Some bandwidth became available again.

→ Current Path (ERO = R1 -> R2 -> R6 -> R7 -> R4 -> R9).

→ New Path (ERO = R1 -> R2 -> R3 -> R4 -> R9)—shared with current path and reserved for both paths.

← Until R9 gets new Path message, current Resv is refreshed—PathTear can then be sent to remove old path (and release resources).

Path Rerouting: Link Failure (Cont.)

- Link failure – What happens:
- (Example: One link along a dynamic tunnel LSP path goes down.)
  - RSVP PathTear causes the headend to flag LSP as dead
  - RSVP session is cleared
  - PCALC triggered:
    - No alternative path:
      - Headend sets the tunnel down
    - Alternative path found:
      - New LSP directly signaled
      - Adjacency table updated for the tunnel interface

Assigning Traffic to Traffic Tunnels

Traffic Flow Modifications

- CBR used to find the path for an LSP tunnel.
- IP routing does not see internal details.
- Tunnels can be used for routing only if they are explicitly specified:
  - Static route in the IP routing table points to a selected LSP tunnel interface.
  - Advertise the tunnel to IP using autoroute.
  - Policy routing—the next-hop interface is an LSP tunnel.

Static Routing on R1 Pointing to Tunnel Interfaces (T1 and T2) for R4 and R5

Topology

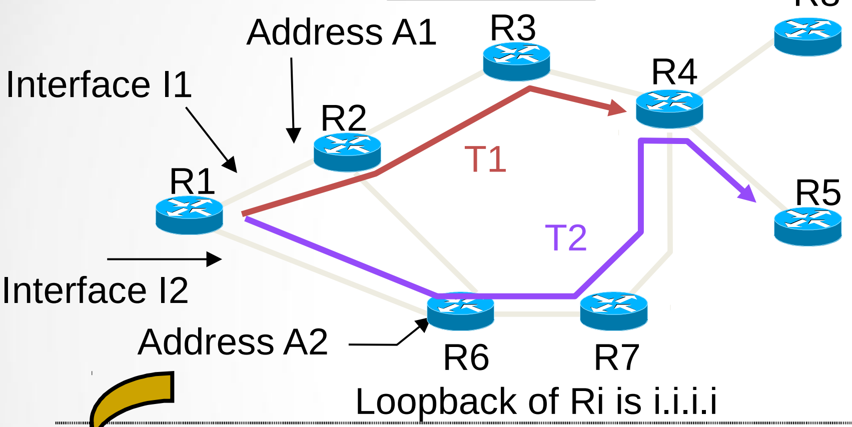

Loopback of Ri is i.i.i.i

Routing Table

Dest	Out Intf	Next Hop	Metric
2.2.2.2	I1	A1	1
3.3.3.3	I1	A1	2
4.4.4.4	T1	R4	3
5.5.5.5	T2	R5	4
6.6.6.6	I2	A2	1
7.7.7.7	I2	A2	2
8.8.8.8	I1 I2	A1 A2	4 4

Shortest-Path Tree

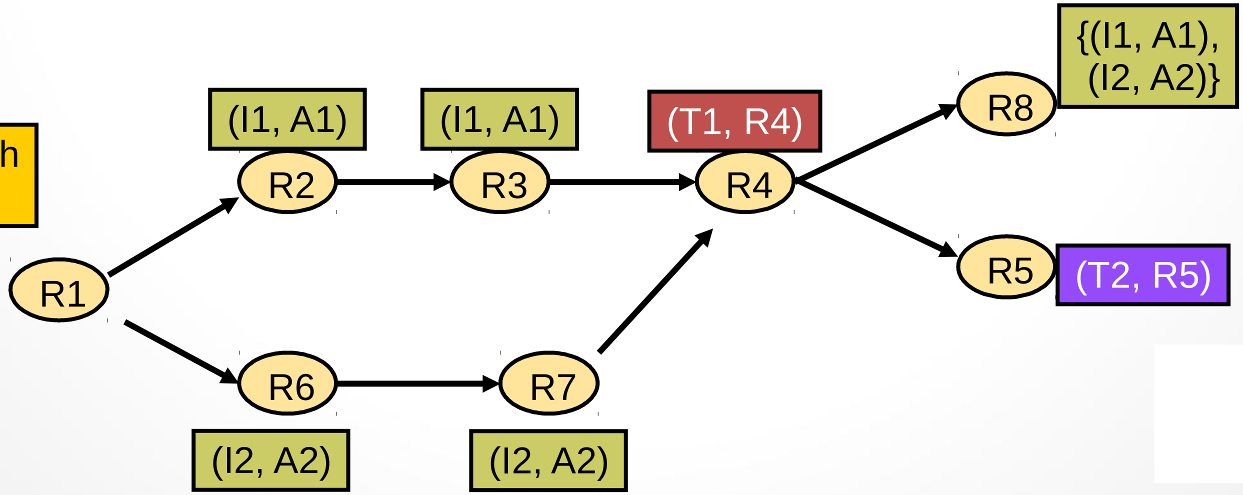

IP Forwarding Database Modification with Autoroute

- Autoroute feature enables the headend to see the LSP as a directly connected interface:
  - Only for the SPF route determination, not for the constraint-based path computation.
  - All traffic directed to prefixes topologically behind the tunnel endpoint (tailend) is forwarded onto the tunnel.
- Autoroute affects the headend only; other routers on the LSP path do not see the tunnel.

Forwarding Adjacency

- Mechanism for:
  - Better intra- and inter-POP load balancing
  - Tunnel sizing independent of inner topology
- Allows the announcement of established tunnel via link-state (LSP) announcements

Forwarding Adjacency (Cont.)

Traffic flow without Forwarding Adjacency

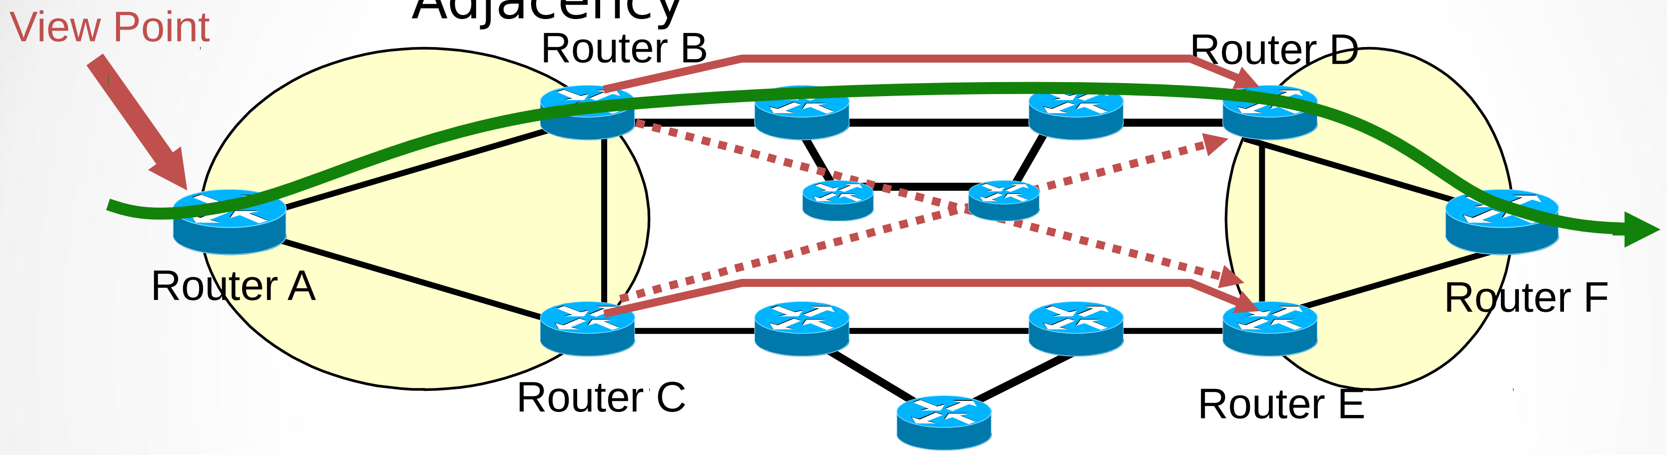

- **Tunnels created and announced to IP with autoroute with equal cost to load-balance.**
  - All the POP-to-POP traffic exits via the routers on the IGP shortest path:
    - No load balancing
    - All traffic flows on tunnel: A B D F

Forwarding Adjacency (Cont.)

Traffic flow without Forwarding Adjacency

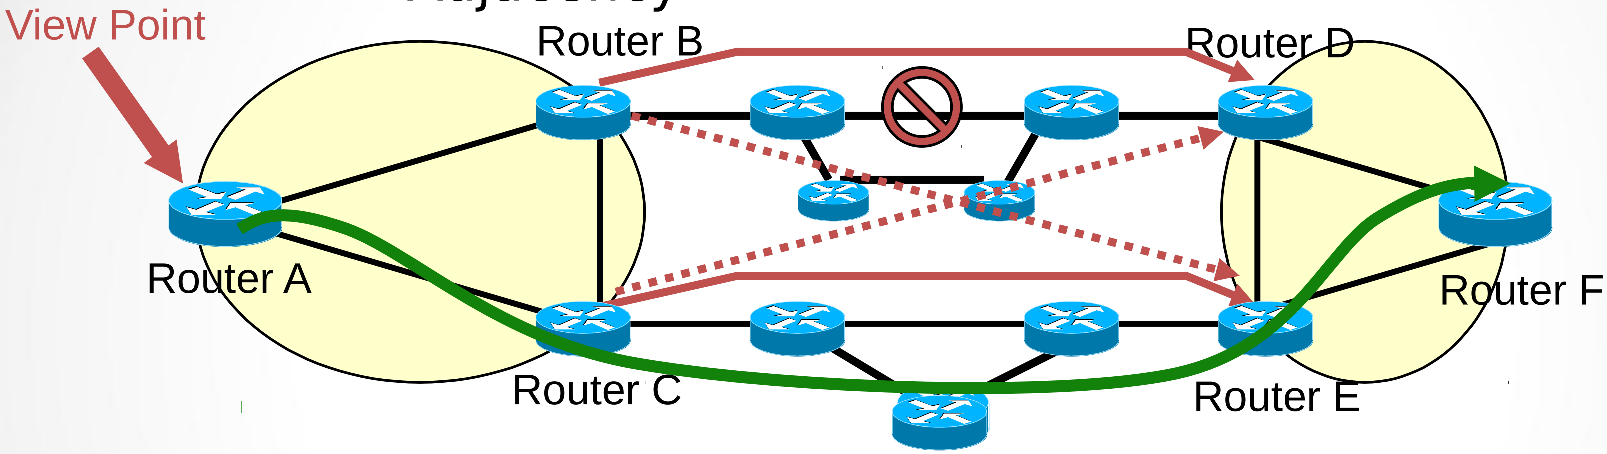

- All the POP-to-POP traffic exits via the routers on the IGP shortest path.
- Change in the core topology does affect the load balancing in the POP:
  - Normal state: All traffic flows A B D F
  - Link failure: All traffic flows A C E F

Forwarding Adjacency (Cont.)

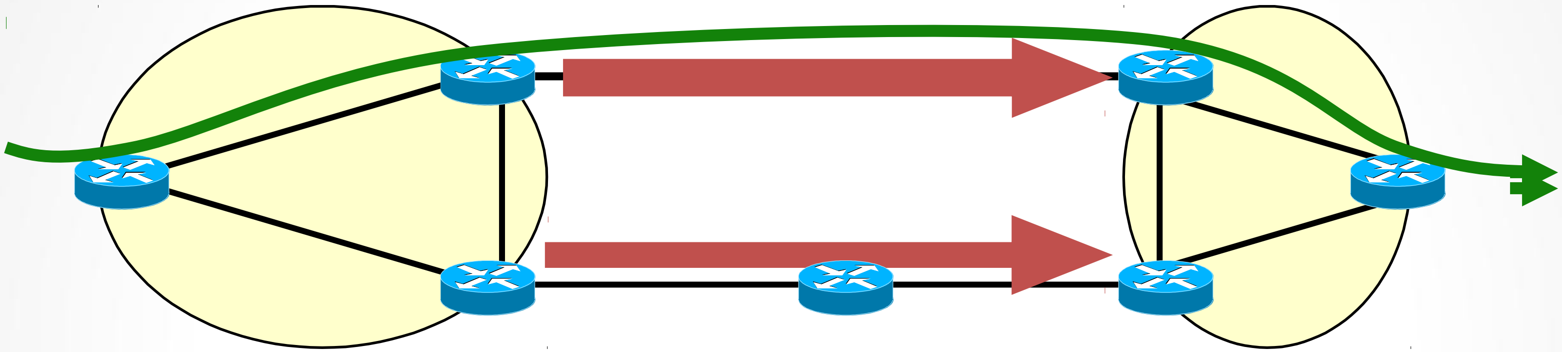

- POP to POP traffic is better load balanced:
  - In the POP: The two core routers are used
  - In the core: At least, two tunnels are used
  - As long as the IGP metric for a path with the FA (e.g. 25) is shorter than the FA-free path (e.g.. 30)
- Inner Topology does not affect Tunnel Sizing:
  - Change in the core topology does not affect the load balancing
    in the POP

Advanced MPLS-TE Link Protection

Improving Convergence Time

The search for an alternative path and its signaling takes too long and has a negative impact on packet forwarding.

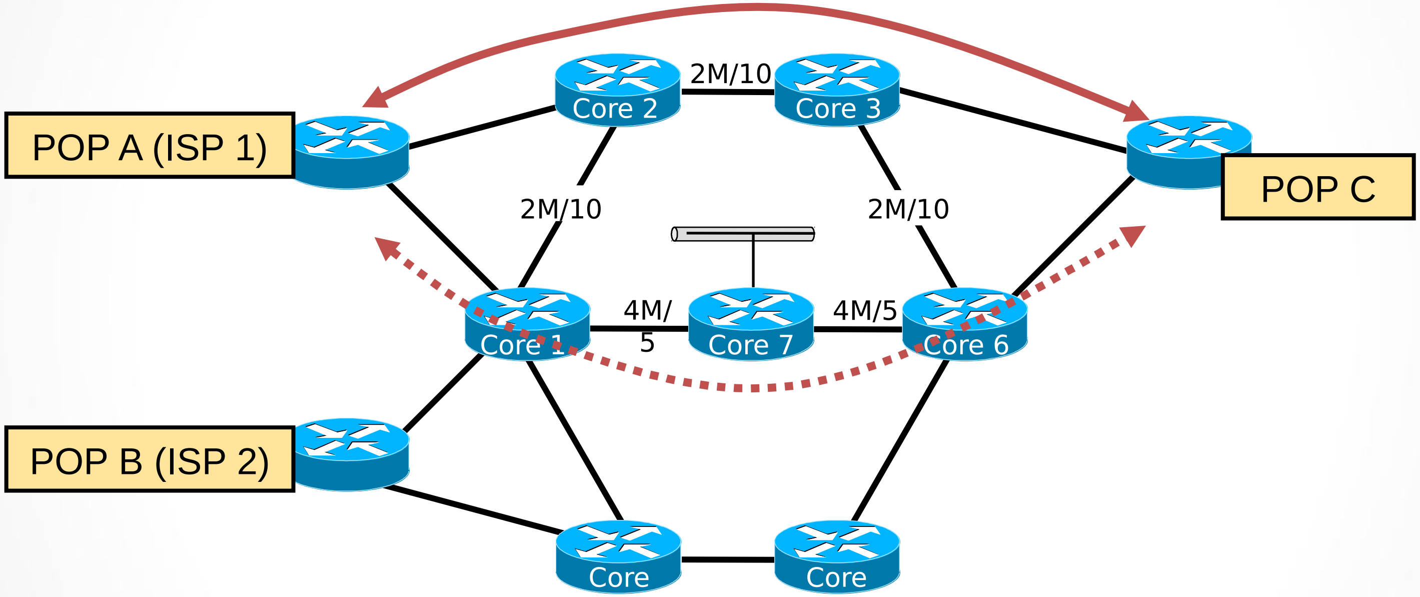

Solution with two pre-established tunnels to the same destination:

- One tunnel could be configured as a backup to another tunnel.
- LSP for the secondary tunnel is presignaled and available if the first tunnel fails.
- Double reservation can be avoided with a “make-before-break” mechanism.

Drawbacks of Parallel Tunnels

Path Protection with Preconfigured Tunnels

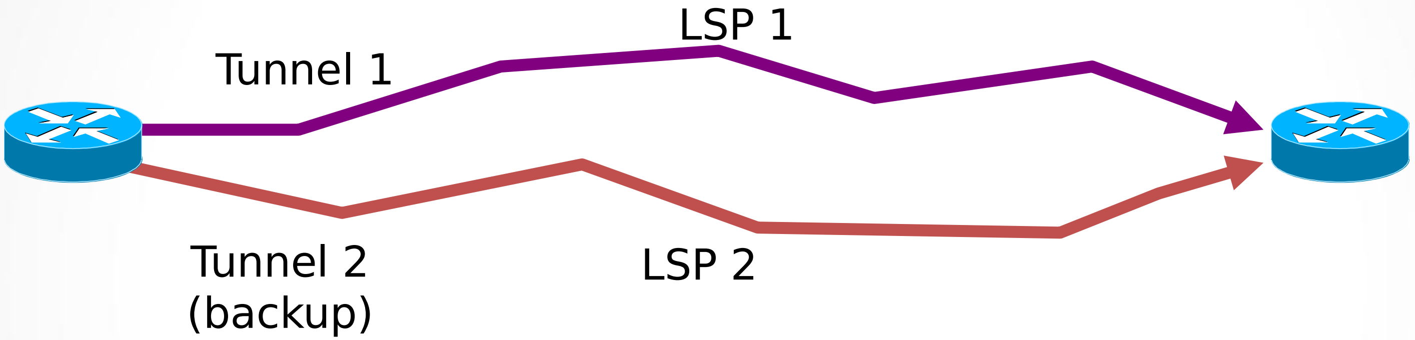

- Preconfigured tunnels speed up recovery by moving the traffic on a preinstalled LSP as soon as the headend learns the primary LSP is down.
- Drawbacks:
  - Backup tunnel allocates labels and reserves bandwidth over the entire path
  - Double counting of reservations via RSVP over the entire path

Fast Reroute: Case Study

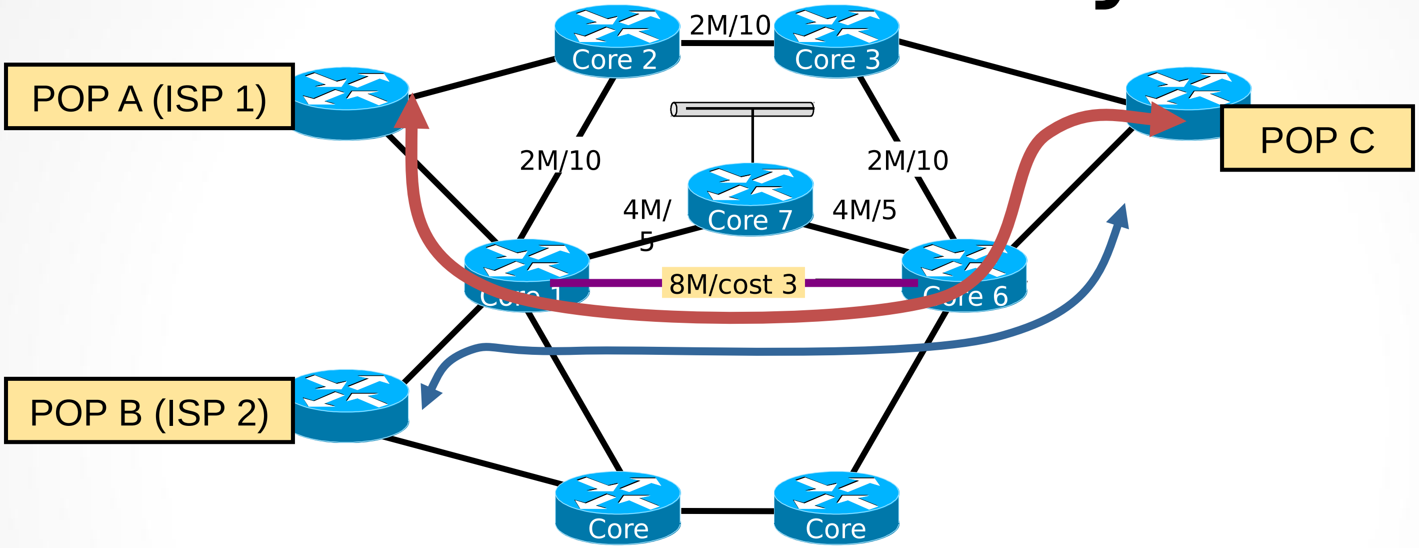

- The company decided to retain only dynamic tunnels. A new high- speed link was introduced between Core 1 and Core 6 to influence CBR and native path selection and speed up transport across the network.
- The new high-speed link is now heavily used by traffic tunnels and may cause a serious disruption.

Fast Reroute

- Fast Reroute allows for temporary routing around a failed link or a failed node while the headend is rerouting the LSP:
  - Controlled by the routers with preconfigured backup tunnels around the protected link or node (link or node protection).
  - The headend is notified of the failure through the IGP and through RSVP.
  - The headend then attempts to establish a new LSP that bypasses the failure (LSP rerouting).

Link Protection with FRR

Link Protection for Core 1 – Core 6 Link

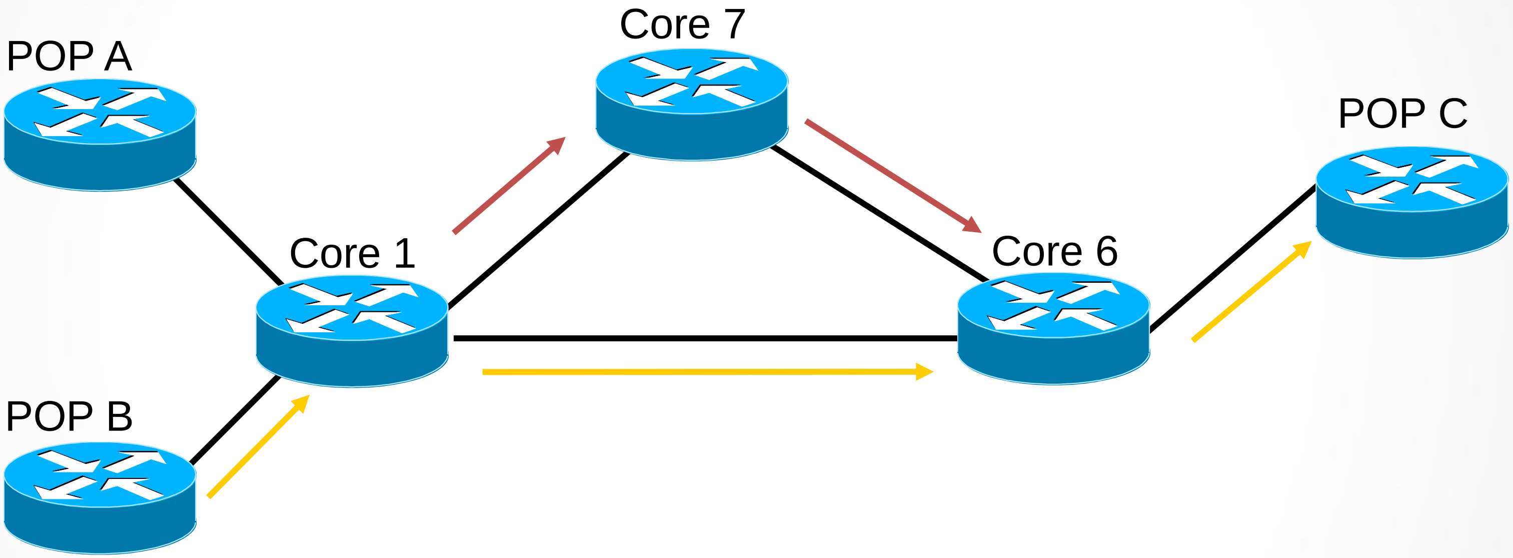

→ End-to-end tunnel onto which data normally flows

→ Bypass (backup) static tunnel to take in the event of a failure

Link Protection with FRR (Cont.)

- “Link Down” Event
  - The router, realizing the link is down:
    - Issues an IGP advertisement
    - Issues an RSVP message with session attribute flag 0x01=ON
      (do not break the tunnel; you may continue to forward packets during the reoptimization)
  - In the event of a failure, an LSP is intercepted and locally rerouted using a backup tunnel.
    - Original LSP nested within protection LSP
    - Minimum disruption of LSP flow
      (under 50 ms - time to detect and switch)
  - The headend is notified by RSVP PathErr and by IGP
    - Special flag in RSVP PathErr (reservation in place) indicates that the path states must not be destroyed, so the LSP flow is not interrupted.
    - The headend of the tunnel smoothly re-establishes the tunnel along a new route.

Node Protection with FRR

Node Protection for Core 5

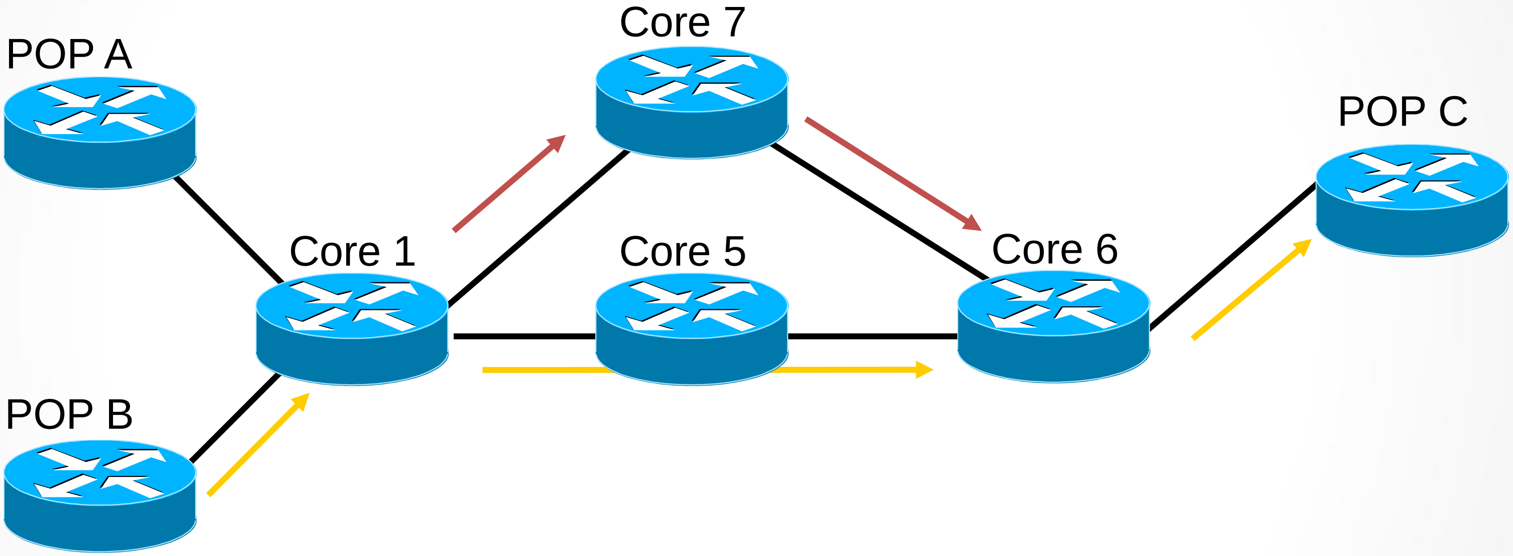

End-to-end tunnel onto which data normally flows

Bypass (backup) static tunnel to take in the event of a failure

Node Protection with FRR (Cont.)

- Router node fails; the router detects this failure by an “interface down” notification.
  - It switches LSPs going out that interface onto their respective backup tunnels (if any).
- RSVP hellos can also be used to trigger Fast Reroute.
  - Messages are periodically sent to the neighboring router.
  - If no response is received, hellos declare that the neighbor is down.
  - Causes any LSPs going out that interface to be switched to their respective backup tunnels.

QoS in MPLS Applications

MPLS-TE with a Best-Effort Network

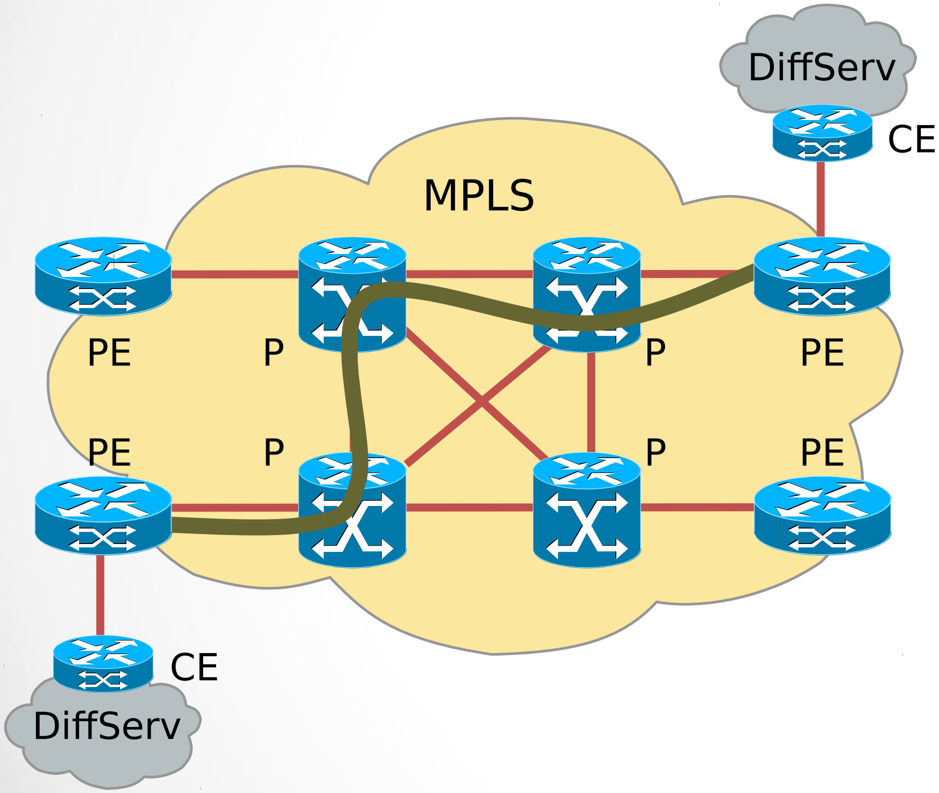

- MPLS-TE defines the path that packets follow to meet constraints (bandwidth).
- LSRs advertise a single available bandwidth via IGP.
- All packets receive best-effort service.

MPLS-TE with a DiffServ Network

- MPLS-TE defines packet path but not packet scheduling.
- LSRs advertise a single available bandwidth via IGP.
- Packets are scheduled at every hop according to EXP marking regardless of LSP.

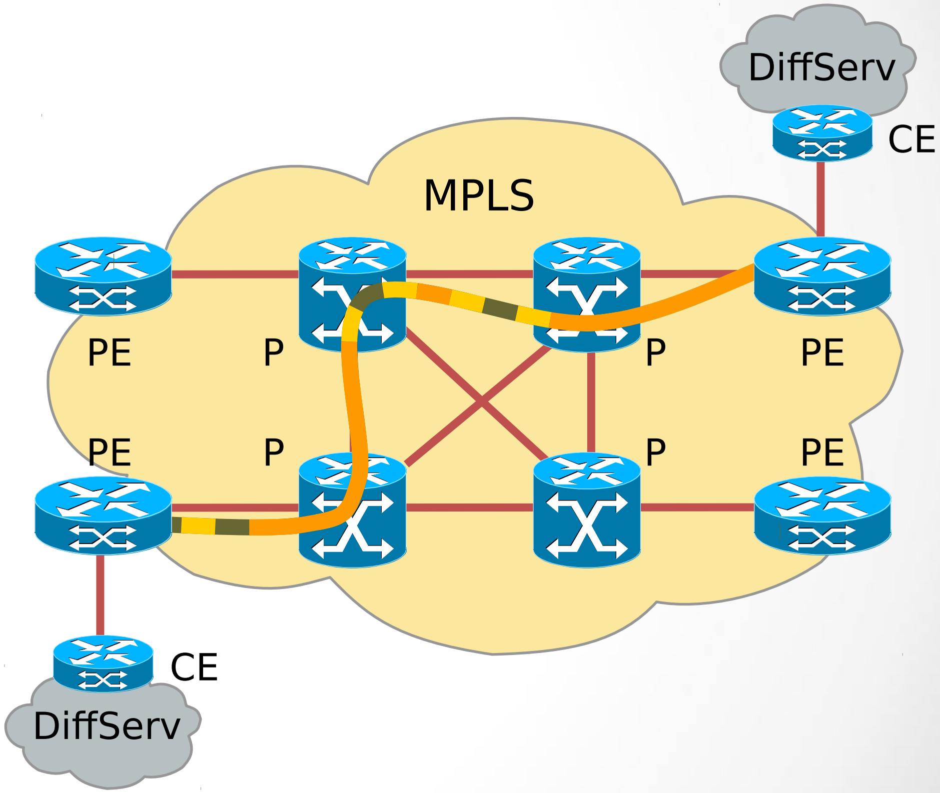

MPLS DS-TE with a DiffServ Network

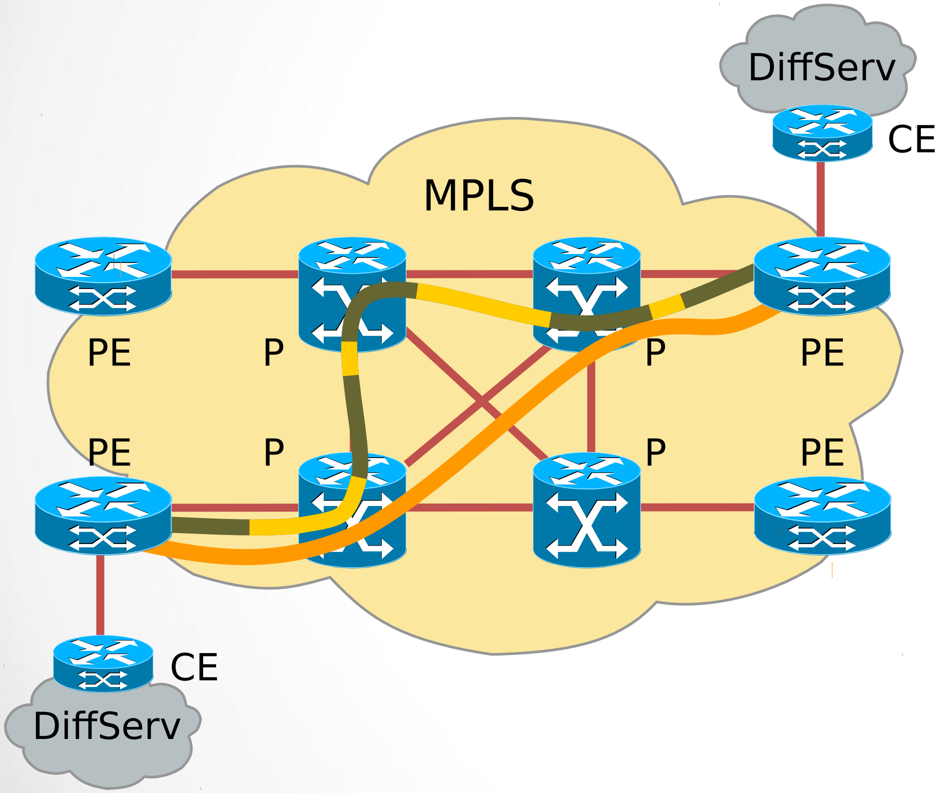

- LSRs advertise multiple available bandwidths (currently two) via IGP.
- Bandwidth carving in data and control plane needs to be provisioned.
- Packets should enter tunnel based on expected QoS.
- Packets are scheduled at every hop according to EXP marking regardless of LSP.

DiffServ-aware TE (DS-TE)

- **Traditional MPLS TE** reserves resources for each node along the MPLS TE tunnel to ensure QoS, but cannot use one TE tunnel to provide differentiated services.
- **The Diff-Serv model** can reserve resources for a single node, but cannot guarantee the QoS over an entire path.
- **MPLS DS-TE** uses the Class Type (CT) so that MPLS TE can allocate resources based on the type of traffic and provide differentiated services. To provide differentiated services, DS-TE divides the LSP bandwidth into one to eight parts, each part corresponding to one Class of Service (CoS). A set of bandwidth of an LSP or a group of LSPs with the same CoS are called a CT.

QoS-Enabled MPLS VPNs

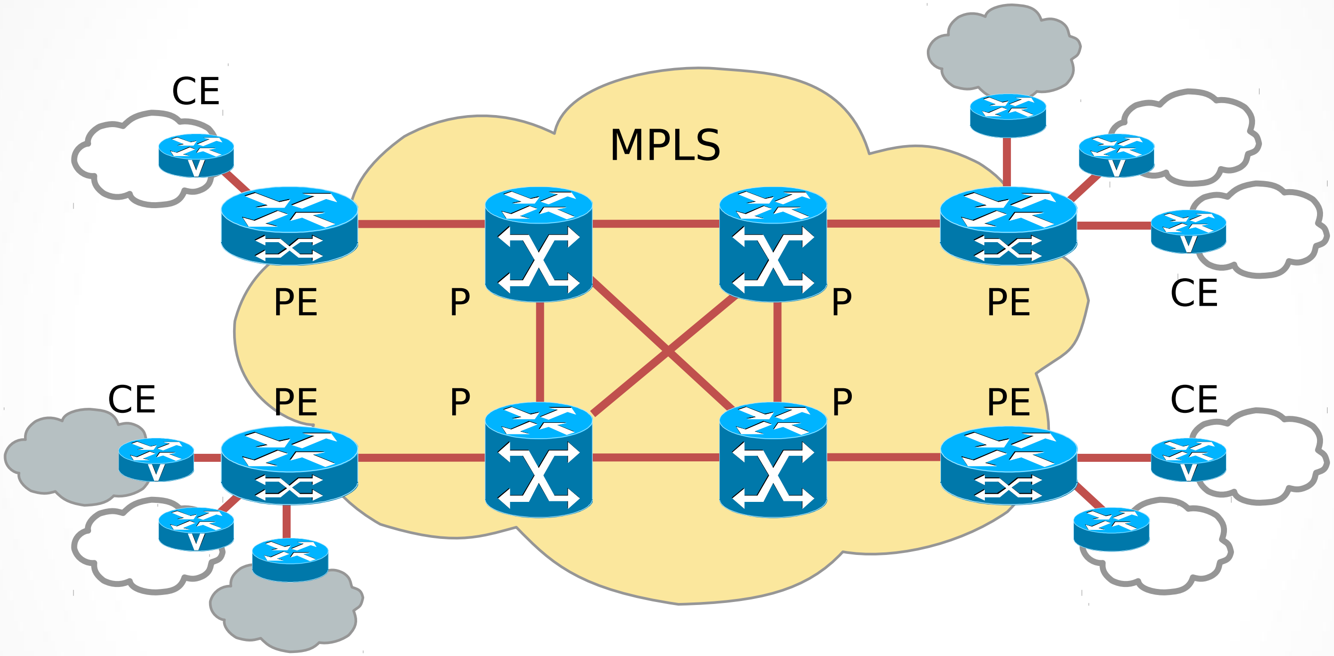

- MPLS VPN QoS models:
  - Point-to-cloud
  - Point-to-point

QoS-Enabled MPLS VPNs (Cont.)

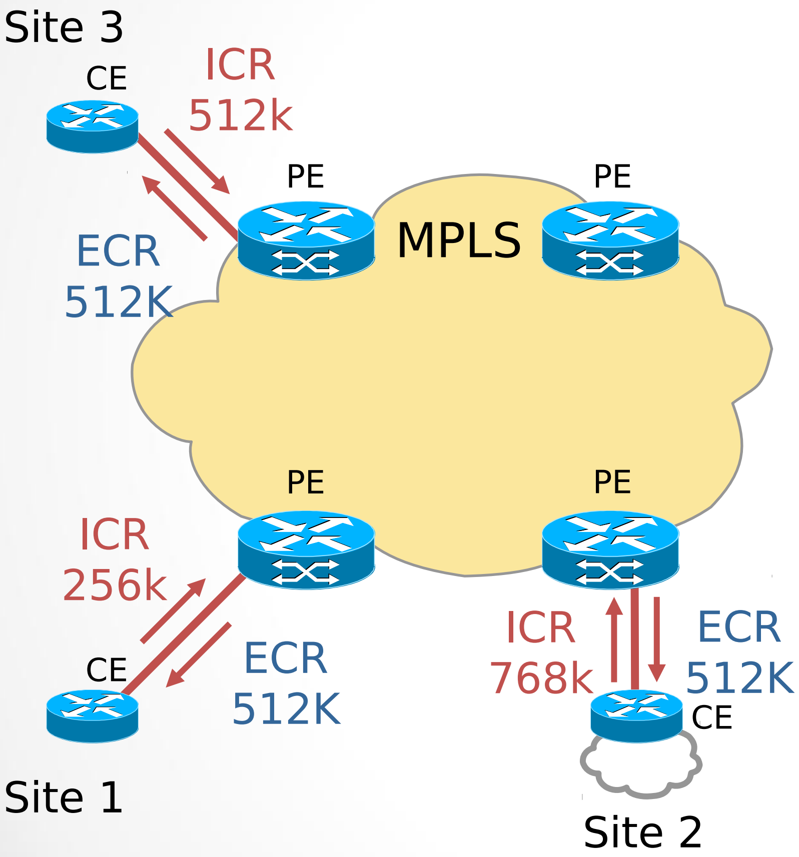

ECR–Egress Committed Rate
ICR–Ingress Committed Rate

Point-to-Cloud Connection

- Per-VPN QoS policies at the edge.
- Same MPLS QoS policies for packets of all VPNs in the core.
- QoS can be implemented with point-to-network guarantees.
- QoS can also be implemented with point-to-point.

QoS-Enabled MPLS VPNs (Cont.)

- Point-to-Point Connection
- Point-to-point (site-to-site) guarantees.
- DS-TE is required to offer hard point-to-point guarantees.
- Point-to-network and point-to-point model are not mutually exclusive.

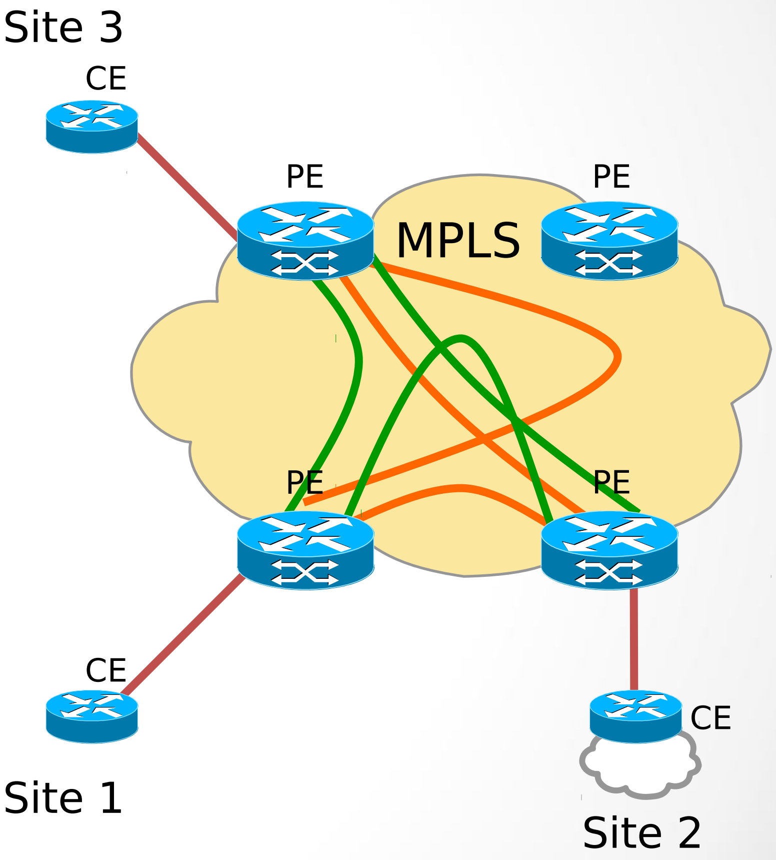
Site 3
CE
PE
PE
MPLS
PE
PE
CE
Site 1
CE
Site 2