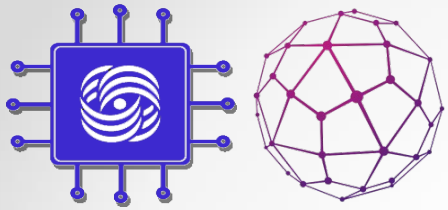

MPLS Technology

Traffic Engineering Concepts

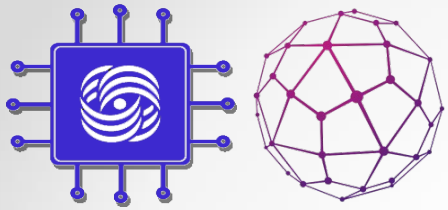

What Is Traffic Engineering?

- 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)

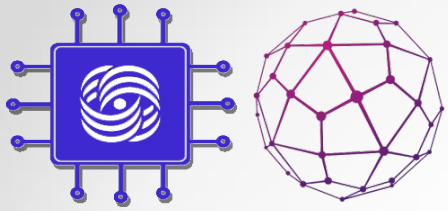

What Is Traffic Engineering? Traffic Engineering Motivations

- 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

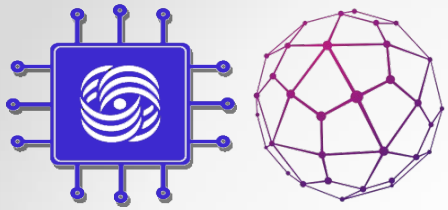

Business Drivers for Traffic Engineering

- 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.

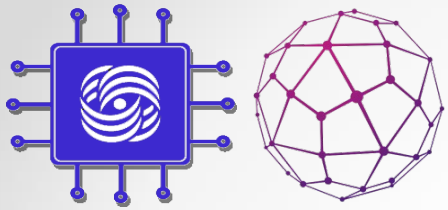

Business Drivers for Traffic Engineering (Cont.)

- 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.

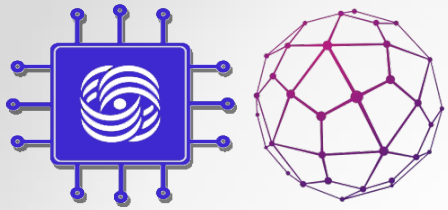

Congestion Avoidance and Traffic Engineering

- 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

Traffic Engineering with the MPLS-TE Model

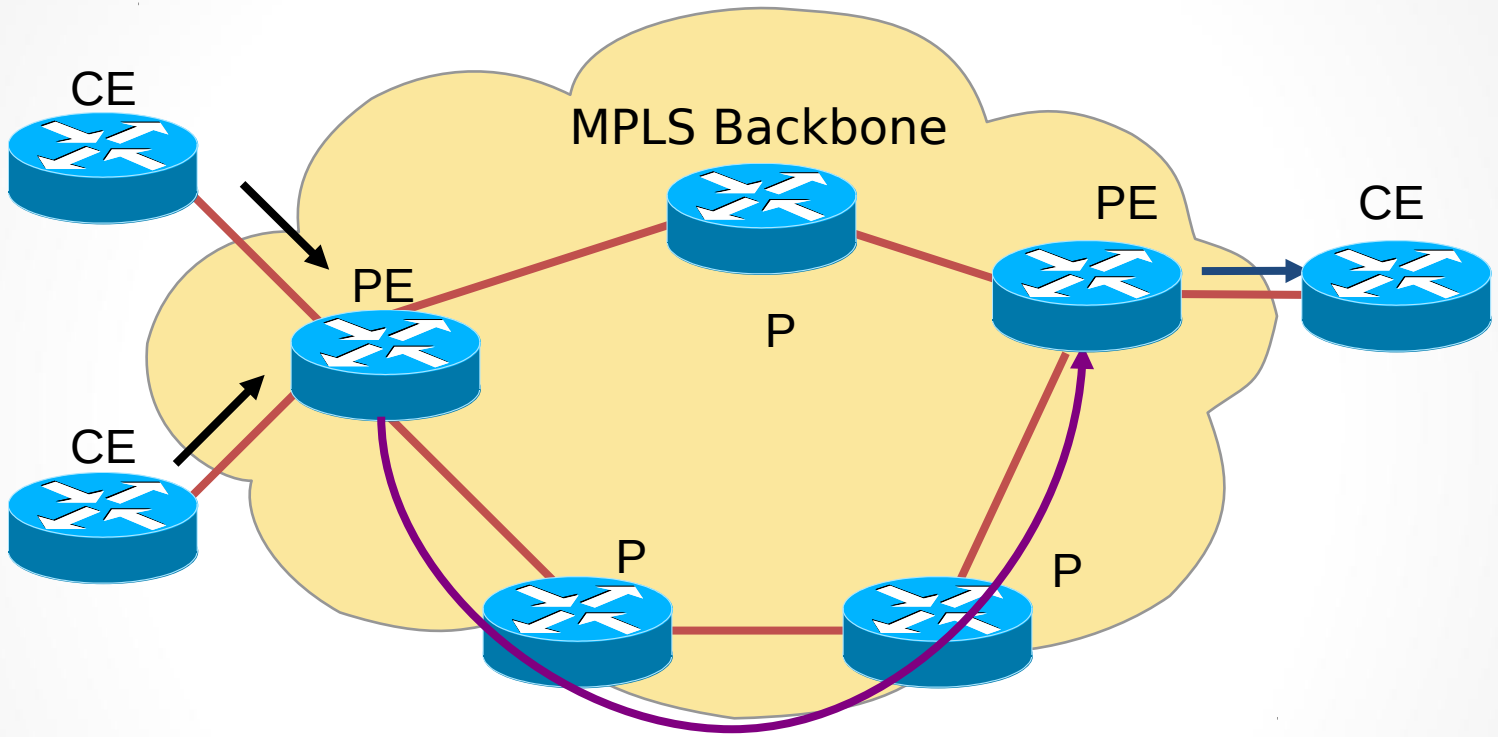

- Tunnel is assigned labels that represent the path (LSP) through the system.
- Forwarding within the MPLS network is based on labels (no Layer 3 lookup).

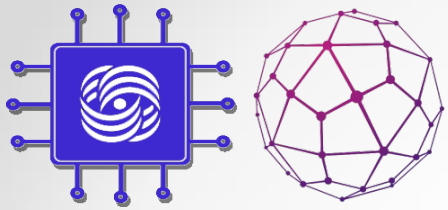

Traffic Engineering with the MPLS-TE Model (Cont.)

- 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

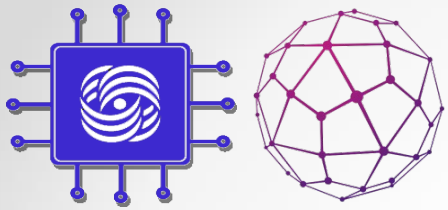

Traffic Engineering with the MPLS-TE Model (Cont.)

Data Plane

С точки зрения передачи данных TE несколько отличается от LDP. Жирным выделены отличия:

- **В TE-туннель трафик нужно поместить насильно, тогда как в LDP он попадает автоматически**
- Первый маршрутизатор навешивает внешнюю MPLS-метку (PUSH LABEL)
- Транзитные маршрутизаторы смотрят на какой интерфейс поступил пакет и значение метки и, поменяв её на новую согласно таблице меток, отправляют её в выходной интерфейс (SWAP LABEL)
- Предпоследний маршрутизатор снимает транспортную метку (POP LABEL, PHP — зависит от реализации и настроек)
- **В случае обрыва на пути трафик можно спасти путём перенаправления пакетов в заранее подготовленный туннель**.

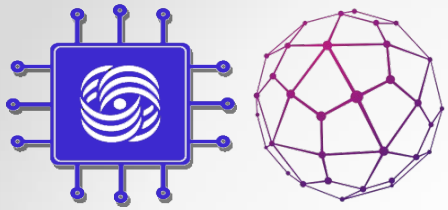

Traffic Engineering with the MPLS-TE Model (Cont.)

Control Plane

Терминология

LSP — Label Switched Path — любой путь через сеть MPLS.
RSVP LSP — соответственно LSP, построенный с помощью RSVP TE с учётом наложенных ограничений. Может также иногда называться CR-LSP — ConstRaint-based LSP.
Туннель - один или несколько MPLS LSP, соединяющих два LSR-маршрутизатора. Метка MPLS — это по сути туннельная инкапсуляция.
В случае LDP — каждый LSP — это отдельный туннель.
В случае RSVP туннель может состоять из одного или нескольких LSP: основной, резервный, best-effort, временный.
TE-туннель - туннель, построенный RSVP-TE.
TEDB — Traffic Engineering Data Base — тот же LSDB протоколов IS-IS/OSPF, но с учётом ресурсов сети, которые интересны модулю TE.
CSPF — Constrained Shortest Path First — расширение алгоритма SPF, которое ищет кратчайший путь с учётом наложенных ограничений.

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

Two LSPs are available on the network. The path LSRA->LSRB->LSRC->LSRD->LSRE is the primary LSP with an LSP ID of 2.

The path LSRA->LSRF->LSRG->LSRH->LSRE is the backup LSP with an LSP ID of 1024. The two LSPs form an MPLS TE tunnel with a tunnel ID of 100, and the tunnel interface is Tunnel1.

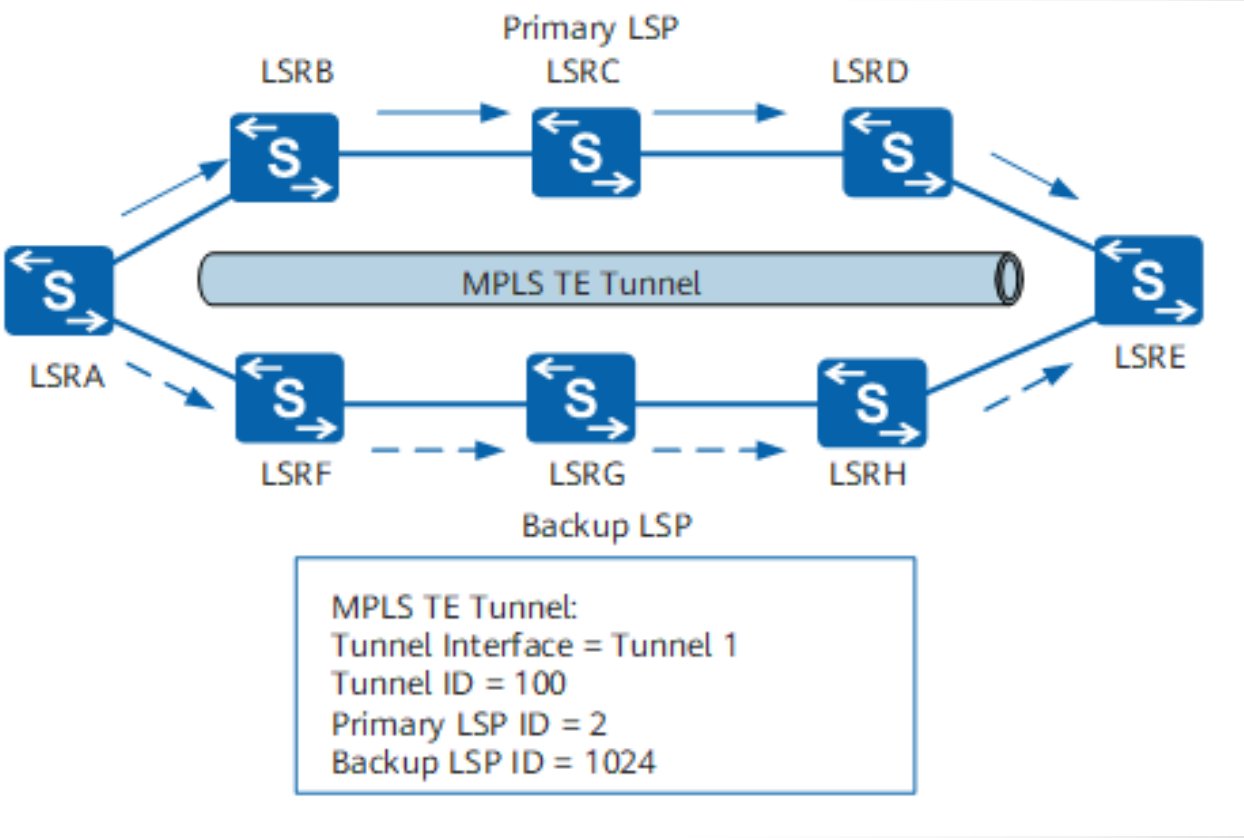

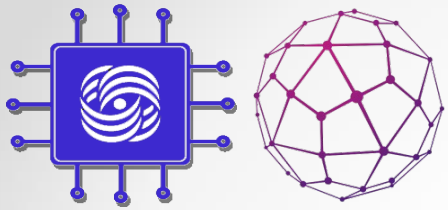

Traffic Engineering with the MPLS-TE Model (Cont.)

Link Attributes

- Total link bandwidth
- Maximum reservable bandwidth
- TE metric
- Shared risk link group (SRLG), a group of links that share a physical resource
- Link administrative group - A 32-bit vector that identifies link attributes, also called a link color. Each bit can be set to 0 or 1 by the network administrator. A link administrative group identifies an attribute, such as the link bandwidth or performance. A link administrative group can also be used for link management. For example, it can identify that an MPLS TE tunnel passes through a link or that a link is transmitting multicast services. The administrative group attribute must be used with the affinity attribute to control path selection.

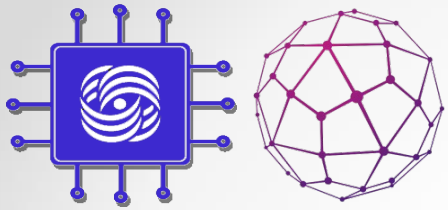

Traffic Engineering with the MPLS-TE Model (Cont.)

Процесс построения LSP:

- IGP собирает информацию (и заполняет TEDB):
  — о линиях и сетях,
  — о метриках,
  — о доступных ресурсах,
  — о характеристиках линий.
- RSVP-TE вызывает CSPF и передает ему ограничения, которые *могут* быть следующими:
  — требуемая полоса пропускания,
  — определённый путь или линии,
  — характеристики линии.
- Из запроса RSVP-TE CSPF берёт ограничения, а из TEDB — реальную информацию о сети. И вычисляет маршрут.
- Когда маршрут получен, RSVP-TE отправляет сообщение RSVP PATH конечному PE с запросом на резервирование ресурсов.
- Конечный PE возвращает сообщение RSVP RESV — так резервируются ресурсы на всём пути. Если RESV вернулся, RSVP LSP/туннель поднимается.

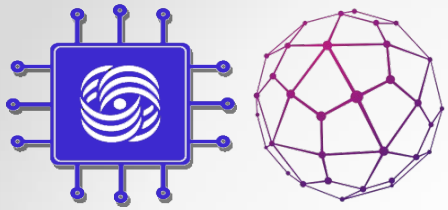

Traffic Engineering with the MPLS-TE Model

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.

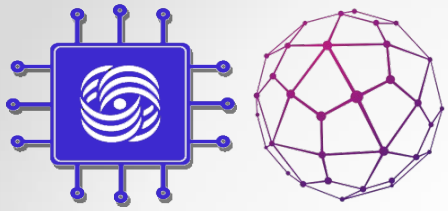

MPLS TE Components

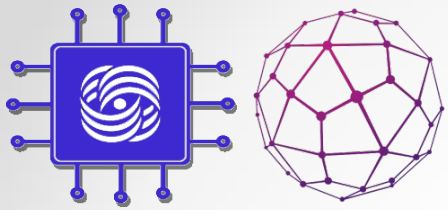

Traffic Tunnels: Concepts

- 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.)

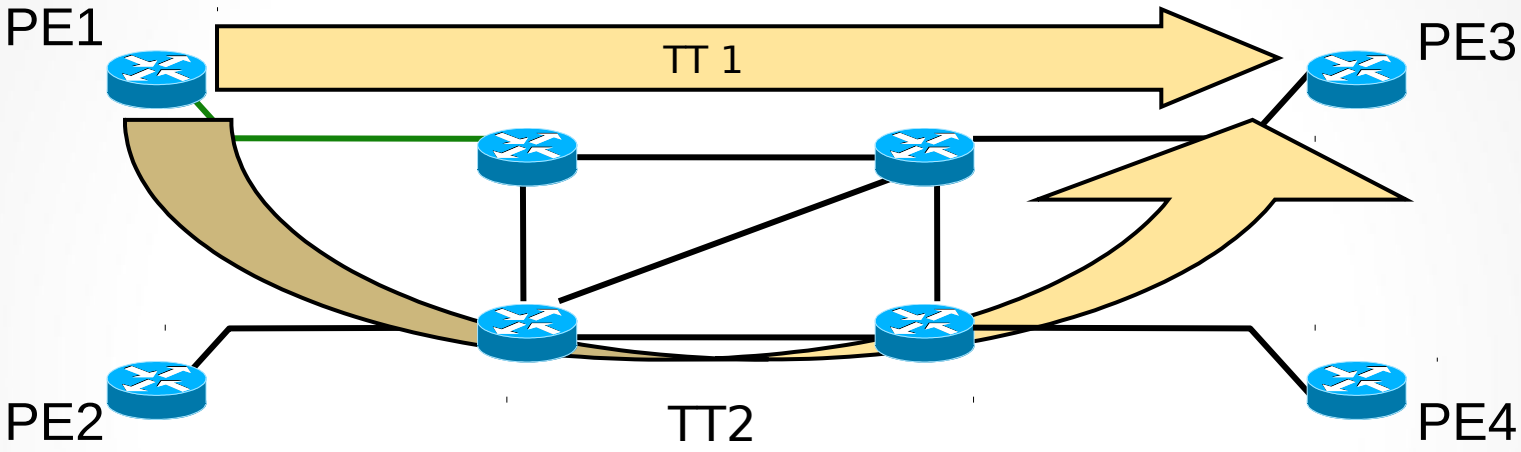

- Unidirectional single class of service model encapsulates all of the traffic between an ingress and an egress router.

- **Different classes of service model assigns traffic into separate tunnels with different characteristics.**

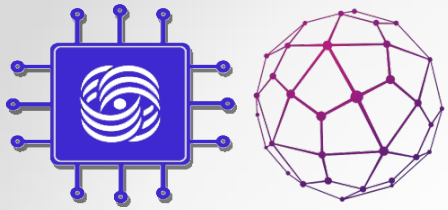

Traffic Tunnels – Characteristics

- 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

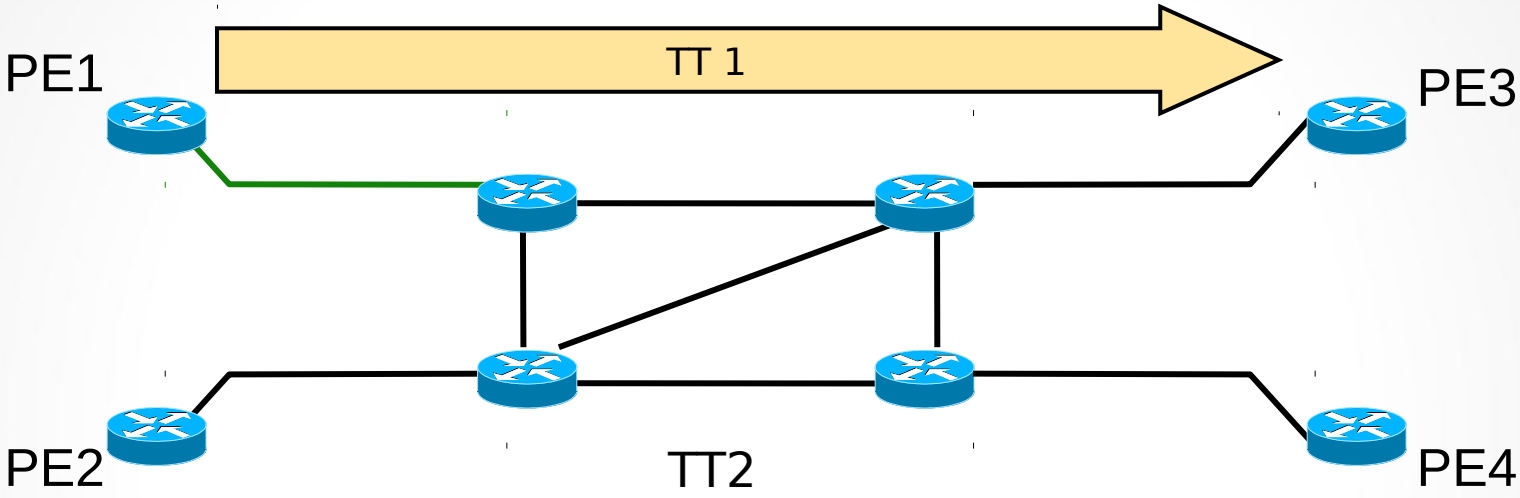

- Attributes are explicitly assigned to traffic tunnels through administrative action.
- A traffic tunnel is characterized by:
  - Its ingress and egress label switch routers
  - The forwarding equivalence class that is mapped onto it
  - A set of attributes that determine its characteristics

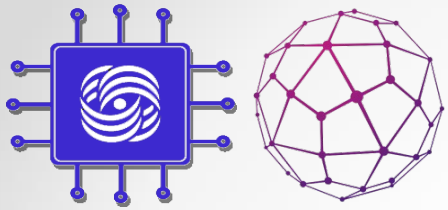

Traffic Tunnels – Attributes (Cont.)

- 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

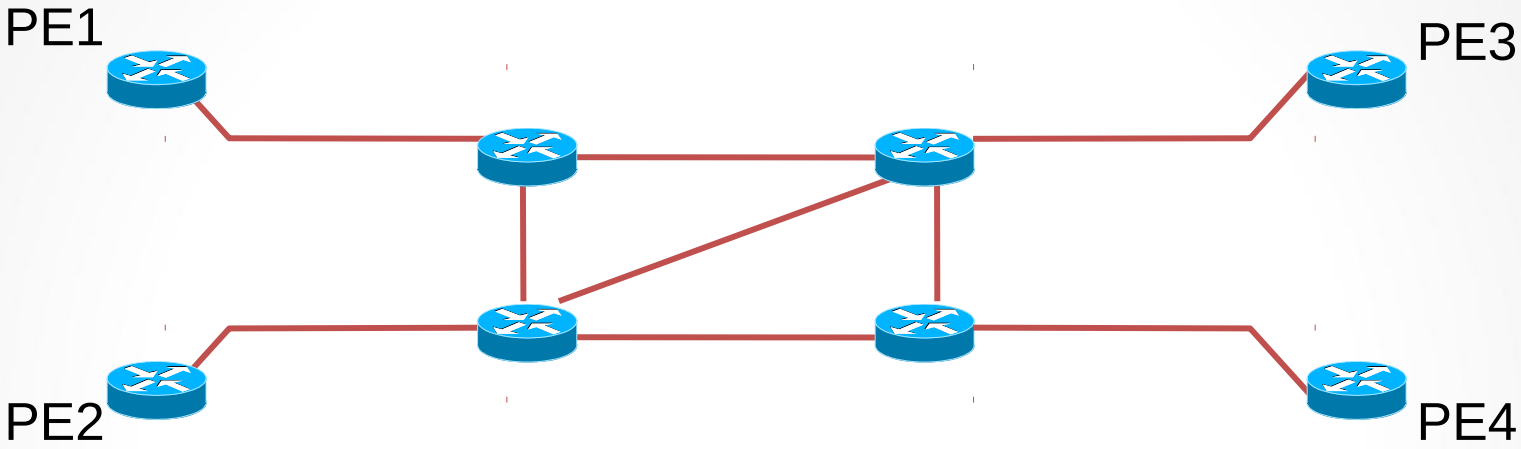

- Resource attributes (link availability) are configured locally on the router interfaces:
  - Maximum bandwidth
    - The amount of bandwidth available
  - Link affinity string
    - To allow the operator to administratively include or exclude links in path calculations
  - Constraint-based specific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

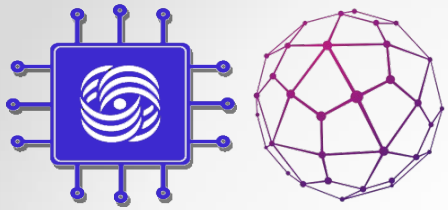

Constraint-Based Path Computation

- 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.)

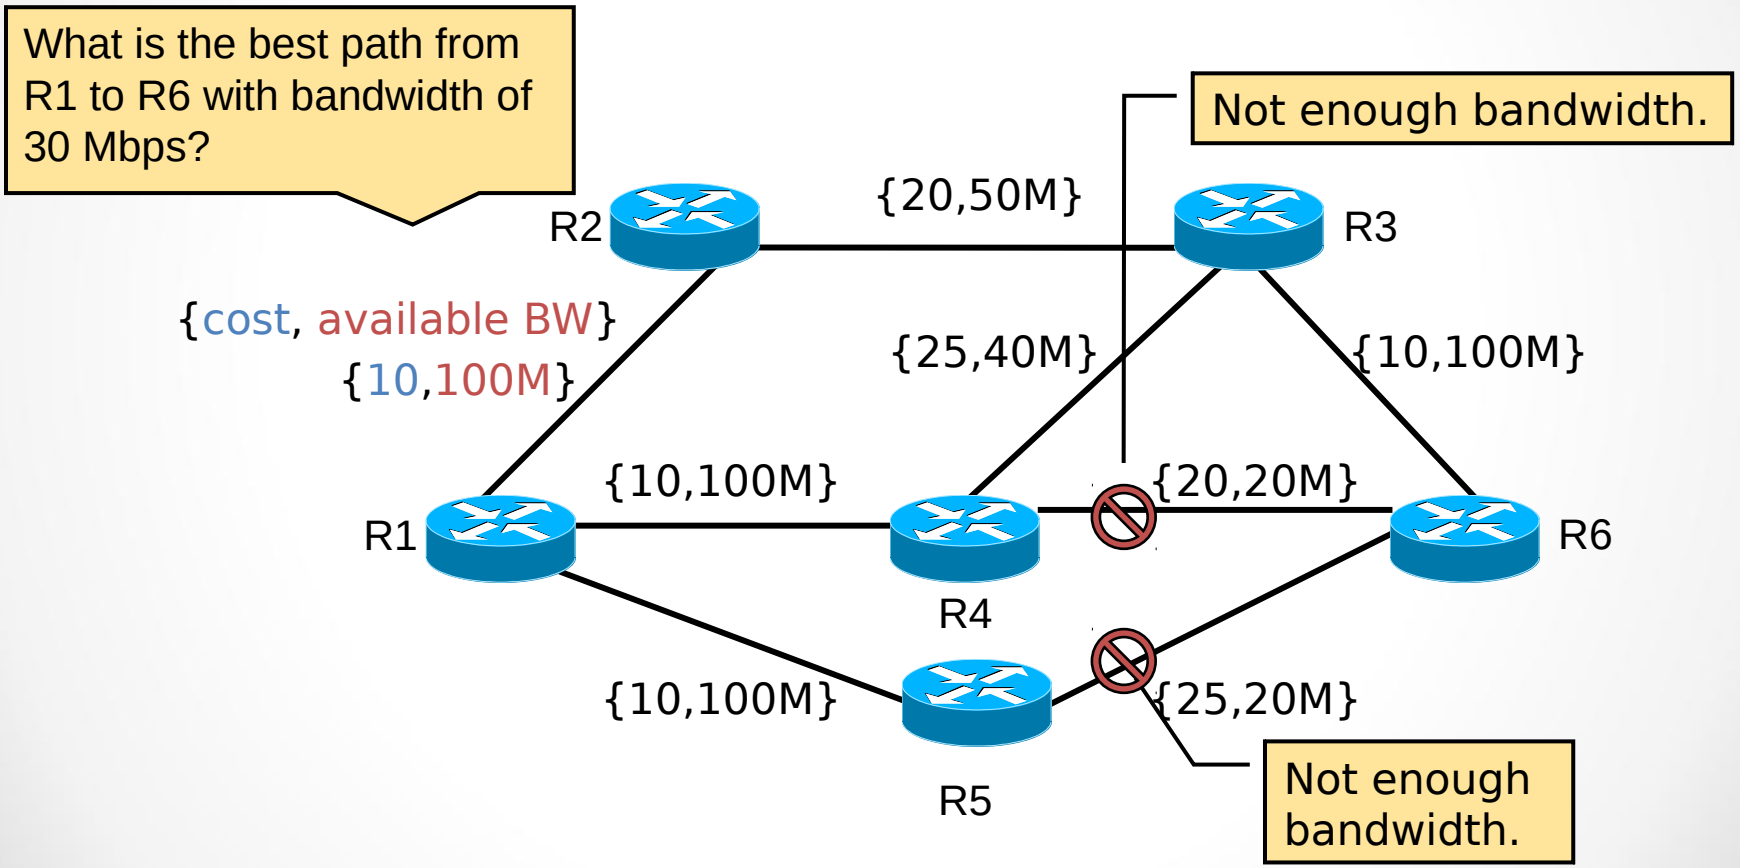

Constraint-Based Path Computation (Cont.)

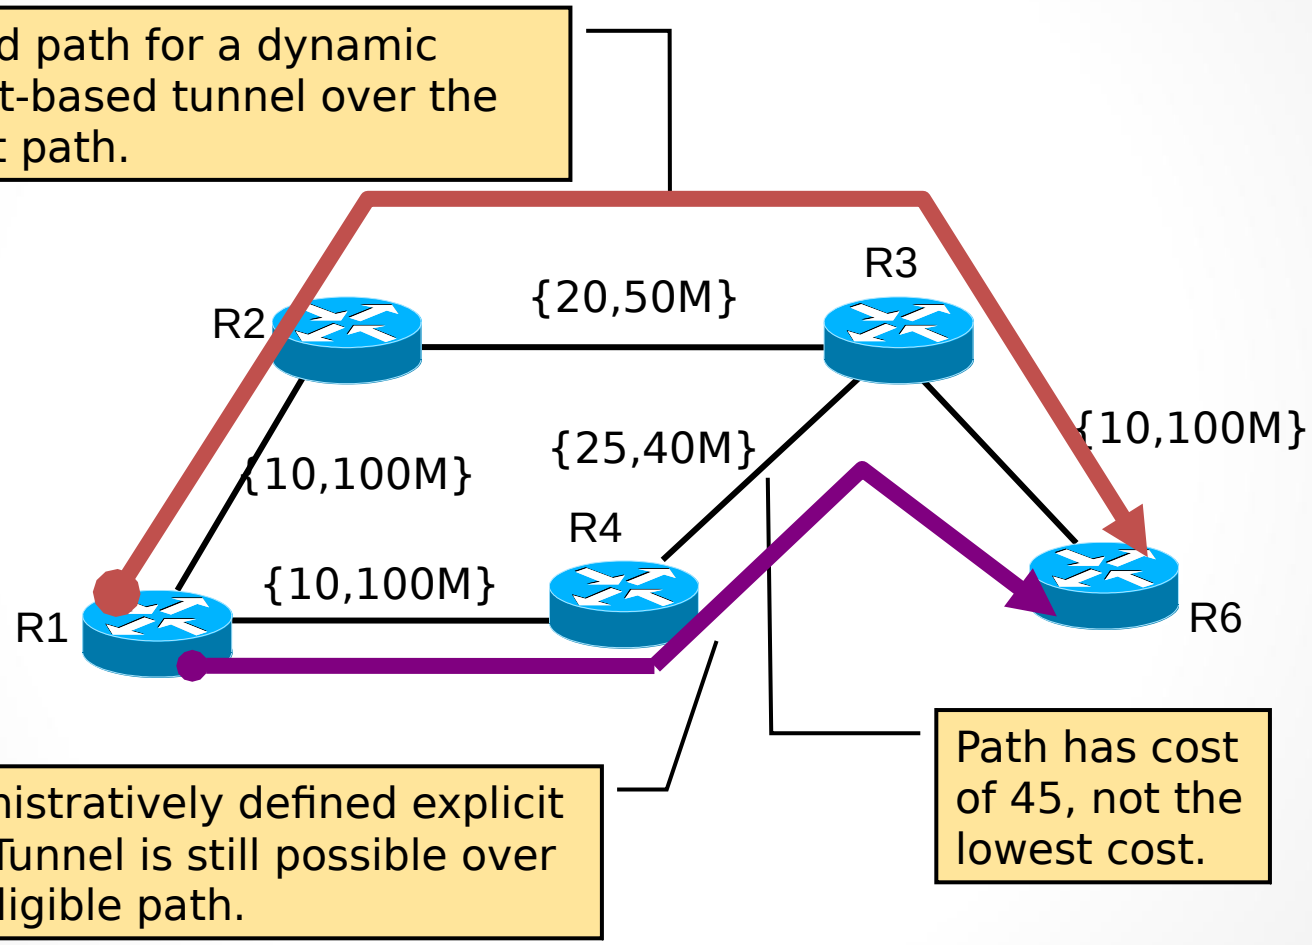

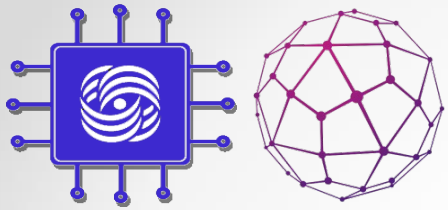

Role of RSVP in Path Setup Procedures

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.

Role of RSVP in Path Setup Procedures

RSVP TE is very closely related to dynamic routing protocols

- We can use, only protocols based on link-state algorithms, i.e. OSPF and ISIS.
- OSPF and ISIS are expanding by introducing new elements. In OSPF - new type of LSA-Opaque LSA, in ISIS - new TLV IS Neighbor and IP Reachability.
- A special modification of the SPF — CSPF (Constrained Shortest Path First) algorithm is used to calculate the path between Ingress LSR and Egress LSR.

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

Hop-by-Hop Path Setup with RSVP

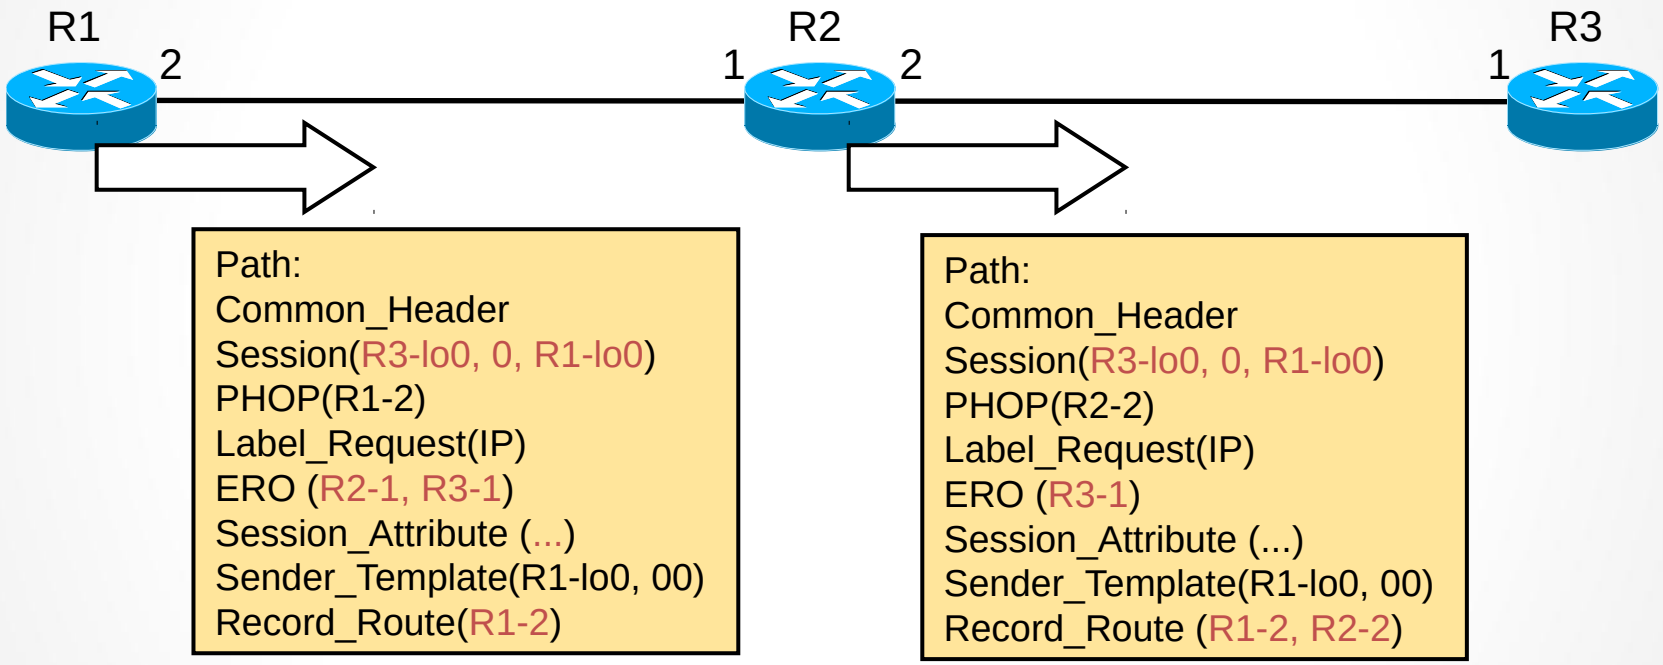

Hop-by-Hop Path Setup with RSVP (Cont.)

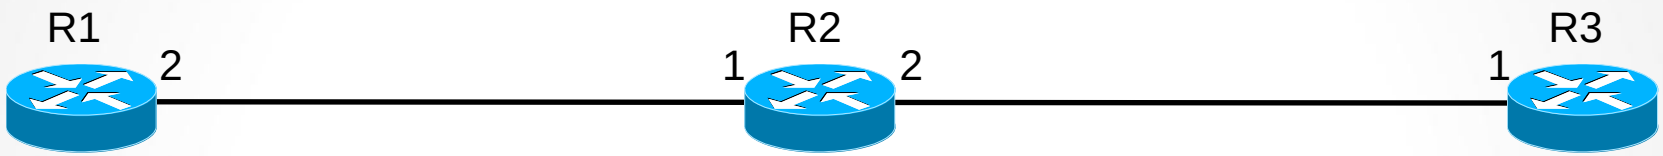

Path State:
Session(R3-lo0, 0, R1-lo0)
PHOP(R2-2)
Label_Request(IP)
ERO ()
Session_Attribute (...)
Sender_Template(R1-lo0, 00)
Record_Route (R1-2, R2-2, R3-1)

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

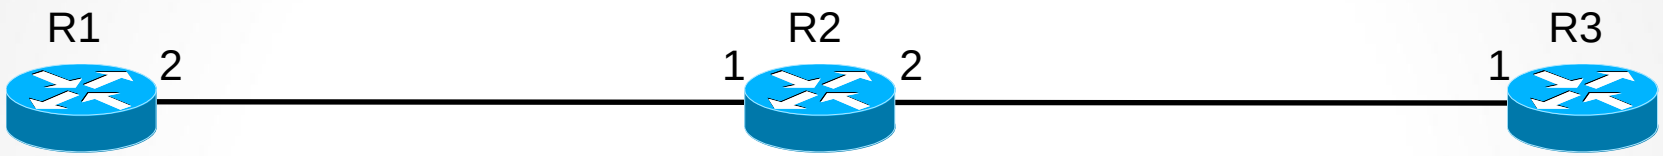

Resv:
Common_Header
Session(R3-lo0, 0, R1-lo0)
PHOP(R2-1)
Sender_Template(R1-lo0, 00)
Label=25
Record_Route(R2-1, R3-1)

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.)

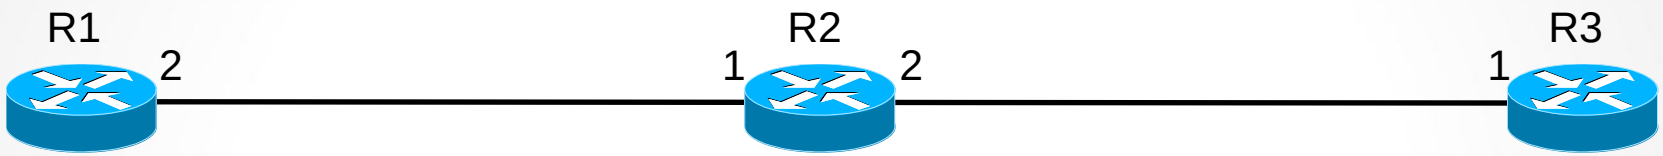

Resv state:
Session(R3-lo0, 0, R1-lo0)
PHOP(R2-1)
Sender_Template(R1-lo0, 00)
Label=5
Record_Route(R1-2, R2-1, R3-1)

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

```
143…  874…  1.1.1.1          4.4.4.4      RSVP   294 PATH Message. SESSION: IPv4-LSP, Destination 4.4.4.4,
143…  874…  10.0.15.5        10.0.15.1    RSVP   142 RESV Message. SESSION: IPv4-LSP, Destination 4.4.4.4,

▷ Frame 14326: 294 bytes on wire (2352 bits), 294 bytes captured (2352 bits) on interface 0
▷ Ethernet II, Src: aa:bb:cc:00:01:10 (aa:bb:cc:00:01:10), Dst: aa:bb:cc:00:05:00 (aa:bb:cc:00:05:00)
▷ Internet Protocol Version 4, Src: 1.1.1.1, Dst: 4.4.4.4
◢ Resource ReserVation Protocol (RSVP): PATH Message. SESSION: IPv4-LSP, Destination 4.4.4.4, Short Call ID 0, Tunnel
   ▷ RSVP Header. PATH Message.
   ▷ SESSION: IPv4-LSP, Destination 4.4.4.4, Short Call ID 0, Tunnel ID 4, Ext ID 1010101.
   ▷ HOP: IPv4, 10.0.15.1
   ▷ TIME VALUES: 30000 ms
   ▷ EXPLICIT ROUTE: IPv4 10.0.15.5, IPv4 10.0.25.5, IPv4 10.0.25.2, ...
   ▷ LABEL REQUEST: Basic: L3PID: IPv4 (0x0800)
   ▷ SESSION ATTRIBUTE: SetupPrio 7, HoldPrio 7, SE Style,  [To Linkmeup_R4]
   ▷ SENDER TEMPLATE: IPv4-LSP, Tunnel Source: 1.1.1.1, Short Call ID: 0, LSP ID: 43.
   ▷ SENDER TSPEC: IntServ, Token Bucket, 1000000 bytes/sec.
   ▷ ADSPEC
```

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):
    - Triggers IGP information distribution when resource thresholds are crossed

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

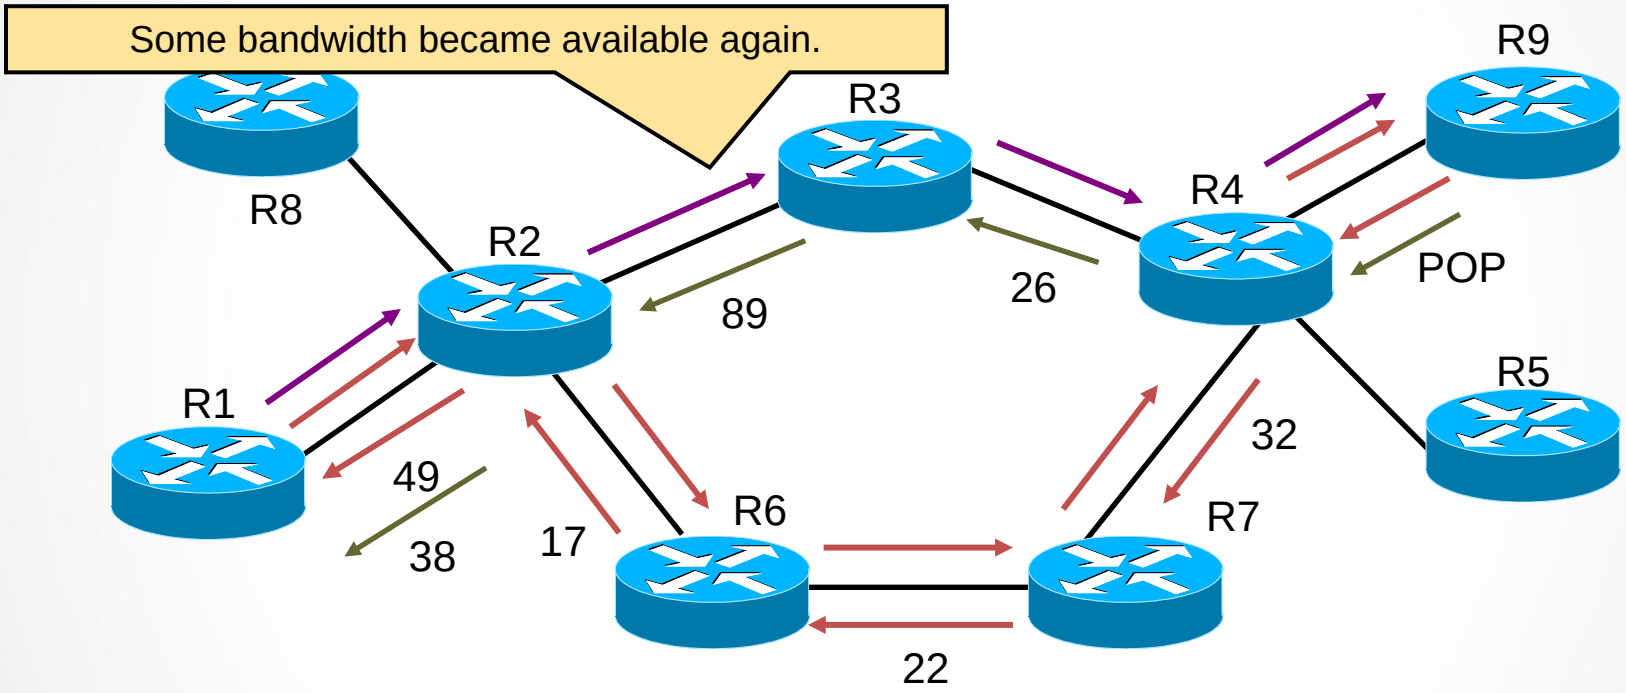

Current Path (ERO = R1 -> R2 -> R6 -> R7 -> R4 -> R9).

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

- The Goal
  - Repair at the headend of the tunnel in the event of failure of an existing LSP:
    - IGP or RSVP alarms the headend.
  - New path for LSP is computed, and eventually a new LSP is signaled.
  - Tunnel interface goes down if there is no LSP available for 10 sec.

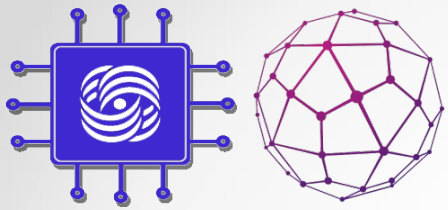

Constraint-Based Path Computation

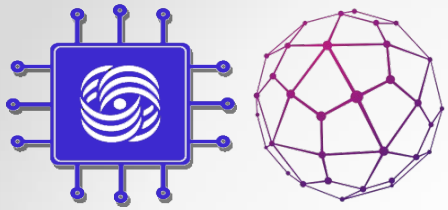

CSPF

CSPF must be aware of constraints and of the available resources on the nodes of the entire network.
The input data - the restrictions specified in the tunnel and the network topology — (the topology contains information about available resources in addition to prefixes and metrics).
Routers communicate with each other through OSPF and ISIS messages not only basic information, but also characteristics of lines, interfaces, etc.

OSPF introduced 3 additional LSA types:

- Type 9 — link-local scope
- Type 10 — area-local scope
- Type 11 — AS scope

Opaque (for OSPF) - special types of LSA, not taken into account in the OSPF. They can be used by any other protocols for their needs. TE uses them to build its topology (it is called TED-Traffic Engineering Database).
ISIS works the same way. New messages: IS-IS TLV 22 (Extended IS Reach), 134 (Traffic Engineering router ID), 135 (Extended IP reach).

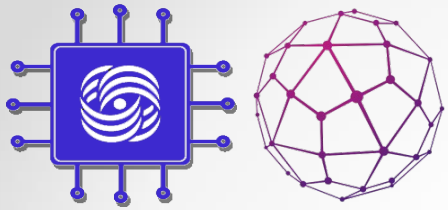

Constraint-Based Path Computation

- 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

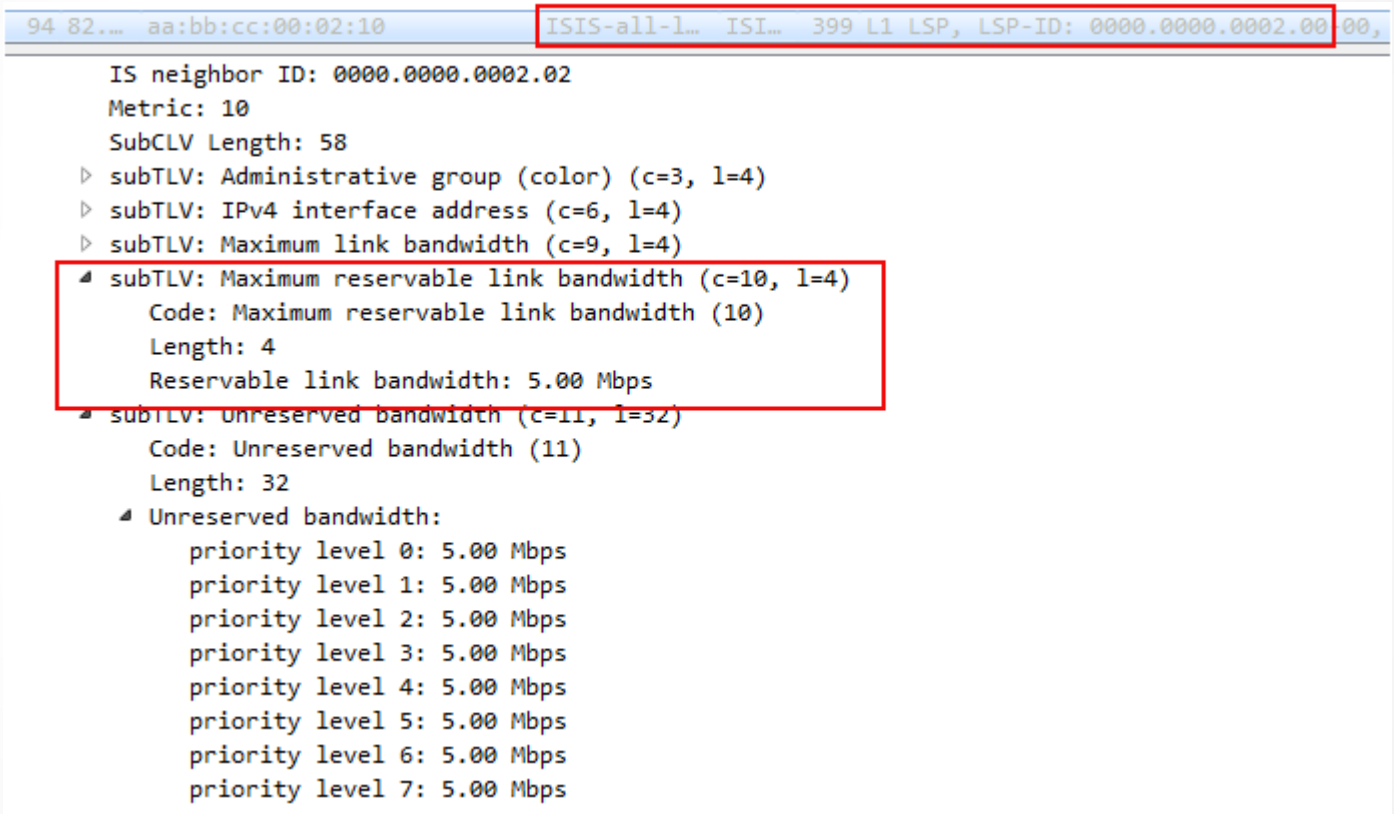

MPLS-TE Link Resource Attributes: Maximum Allocation Bandwidth

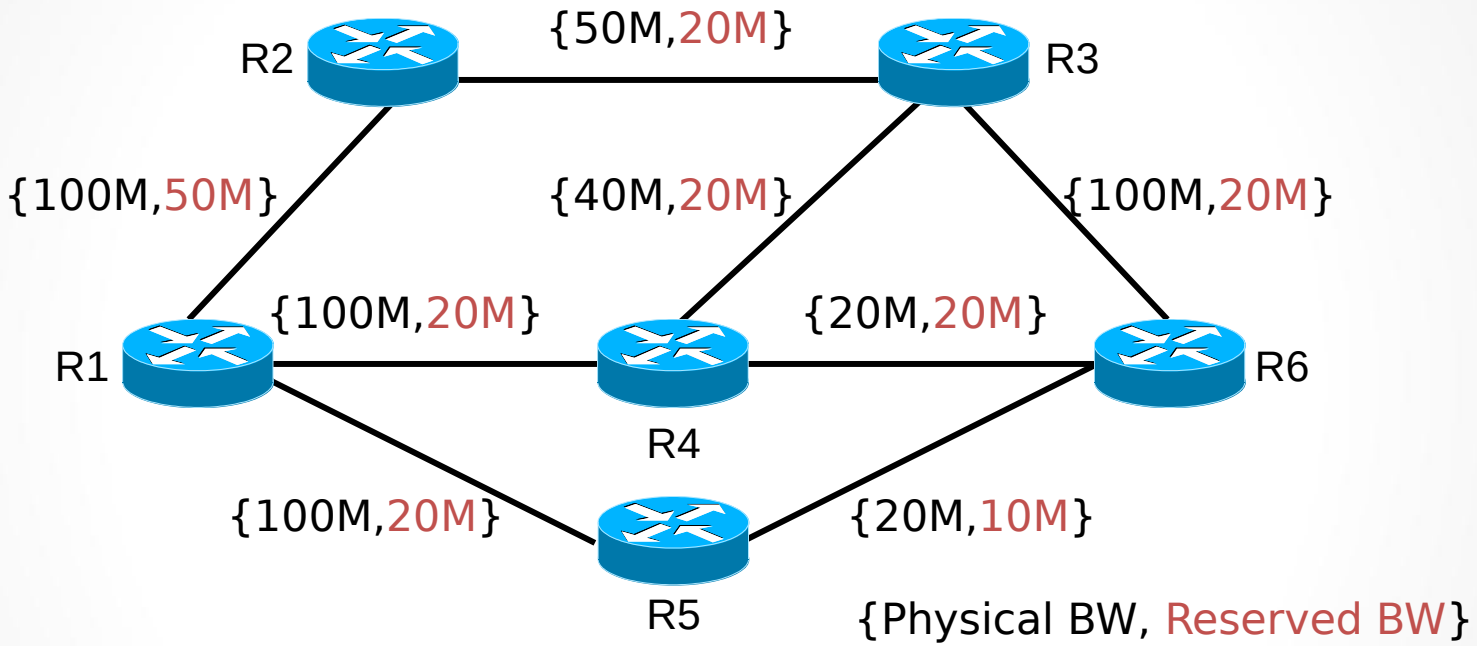

- Maximum bandwidth: the maximum bandwidth that can be used on this link in this direction (physical link)
- Maximum reservable bandwidth: The maximum amount of bandwidth that can be reserved in this direction on this link

MPLS-TE Link Resource Attributes: Link Resource Class

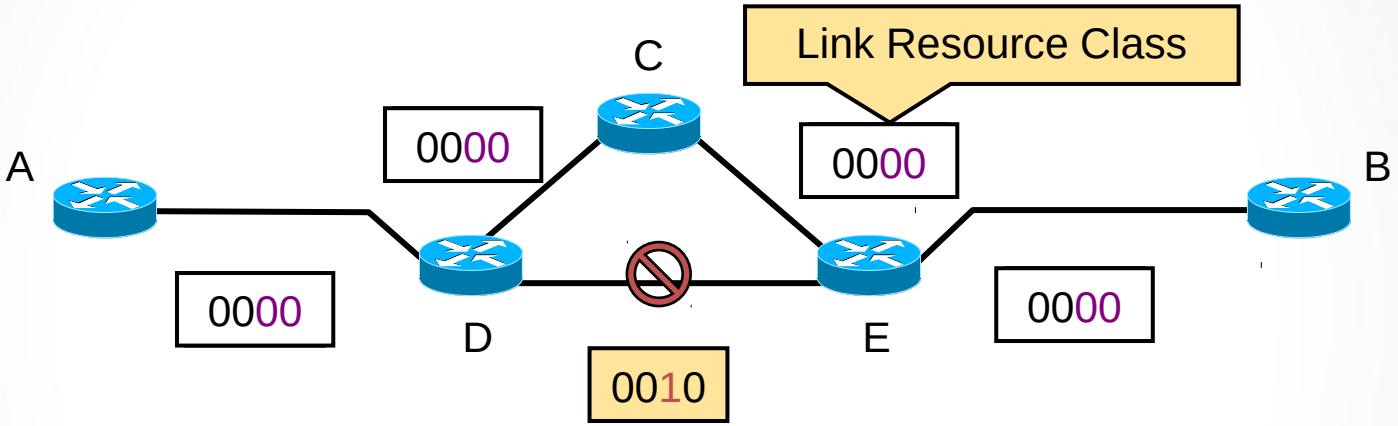

- Link is characterized by a 32-bit resource class attribute.
- Associated with a traffic tunnel in order to include or exclude certain links into or from the path of the traffic tunnel.

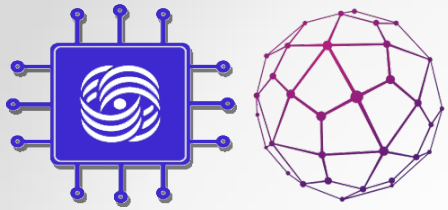

MPLS-TE Link Resource Attributes

Administrative Group (у Juniper: admin-group, у Cisco: Attribute-Flag) — это атрибут физического интерфейса, который может описать 32 его дискретных характеристики.
Какой бит из 32 за что отвечает решает оператор.
Например:
первый бит = 1 означает, что это оптика
второй бит = 1 означает, что это РРЛ
третий бит = 0 означает, что это линия в сторону сети доступа, а 1 — магистральный интерфейс.
четвёртый бит = 1 означает, что это аренда
пятый бит = 1 означает, что это Балаган-Телеком
шестой бит = 1 означает, что это Филькин-Сертификат
седьмой бит = 1 означает, что это канал через интернет без гарантий.
...
десятый бит = 1 означает, что полоса пропускания меньше 500 Мб/с
.....

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

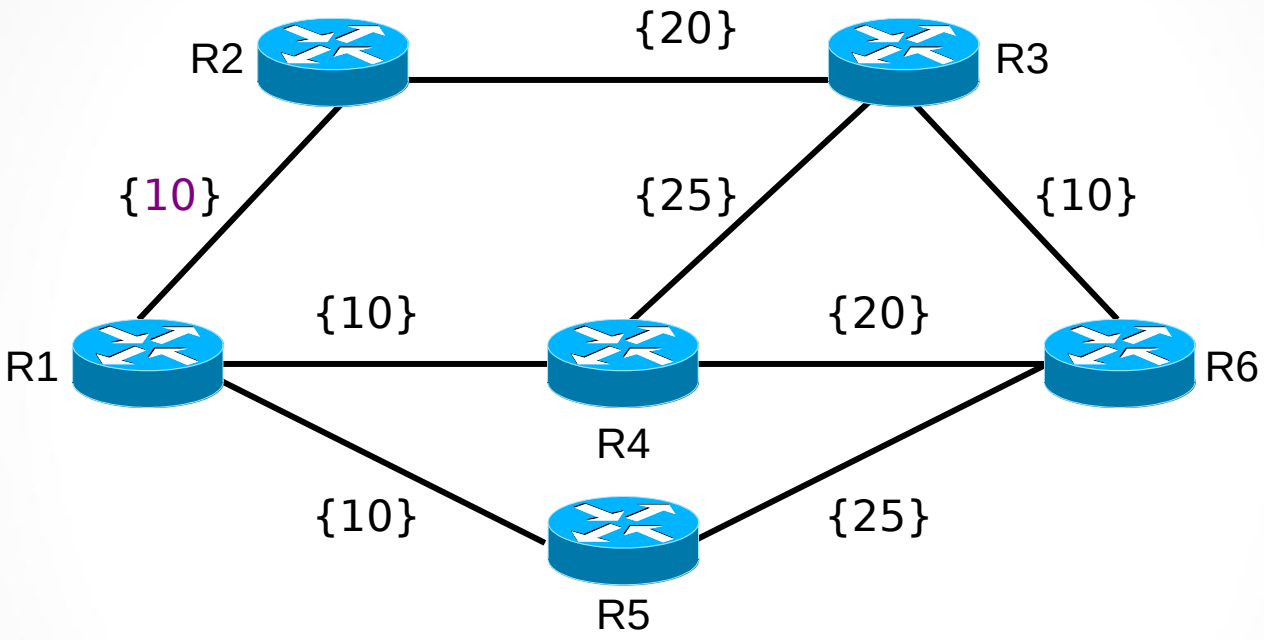

- This metric is administratively assigned to present a differently weighted topology to traffic engineering SPF calculations:
  - » Administrative weight (TE metric)

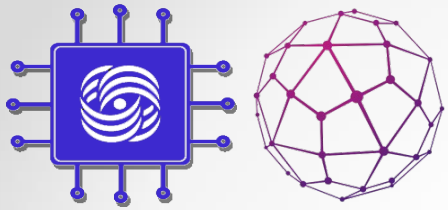

MPLS-TE Tunnel Attributes

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

MPLS-TE Tunnel Attributes (Cont.)

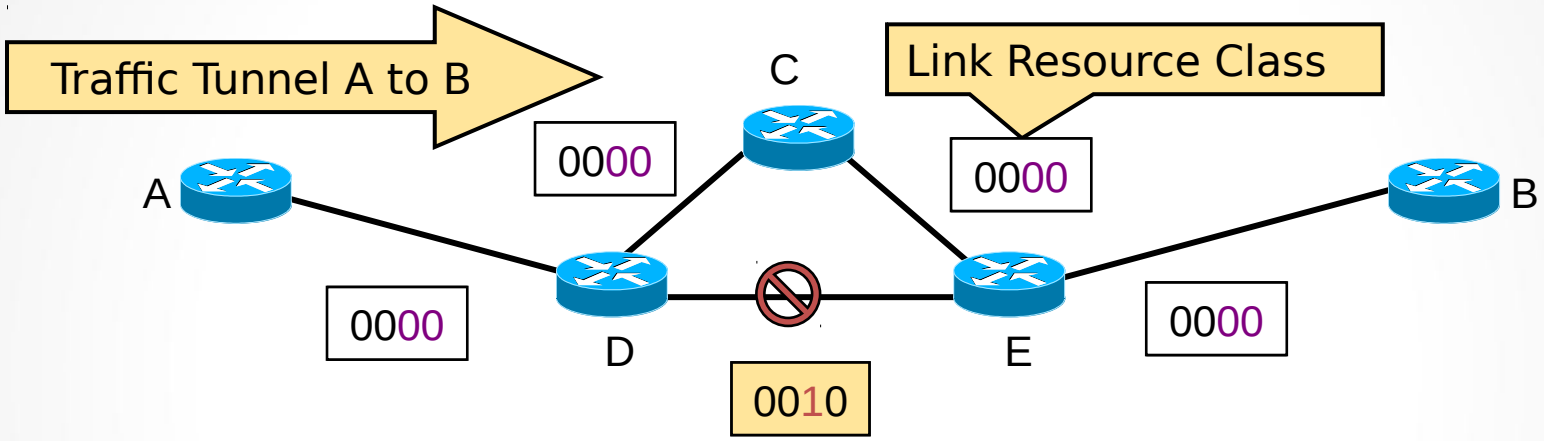

- Tunnel Resource Class Affinity:
  - The properties that the tunnel requires from internal links:
    - 32-bit resource class affinity bit string + 32-bit resource class mask
  - Link is included in the constraint-based LSP path when the tunnel resource affinity string or mask matches the link resource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

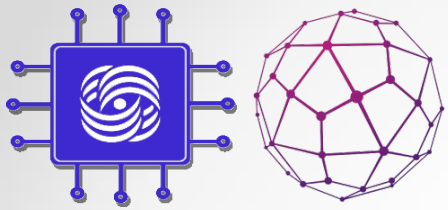

MPLS-TE Tunnel Attributes (Cont.)

- 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
- Pre-emption:
  - Determines whether another traffic tunnel can pre-empt a specific traffic tunnel:

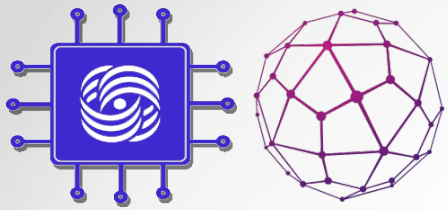

MPLS-TE Tunnel Attributes (Cont.)

- 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

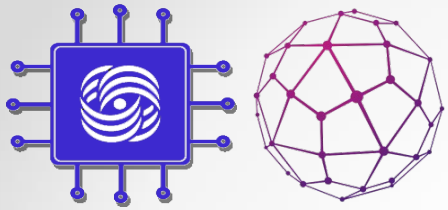

Implementing TE Policies with Affinity Bits

- 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

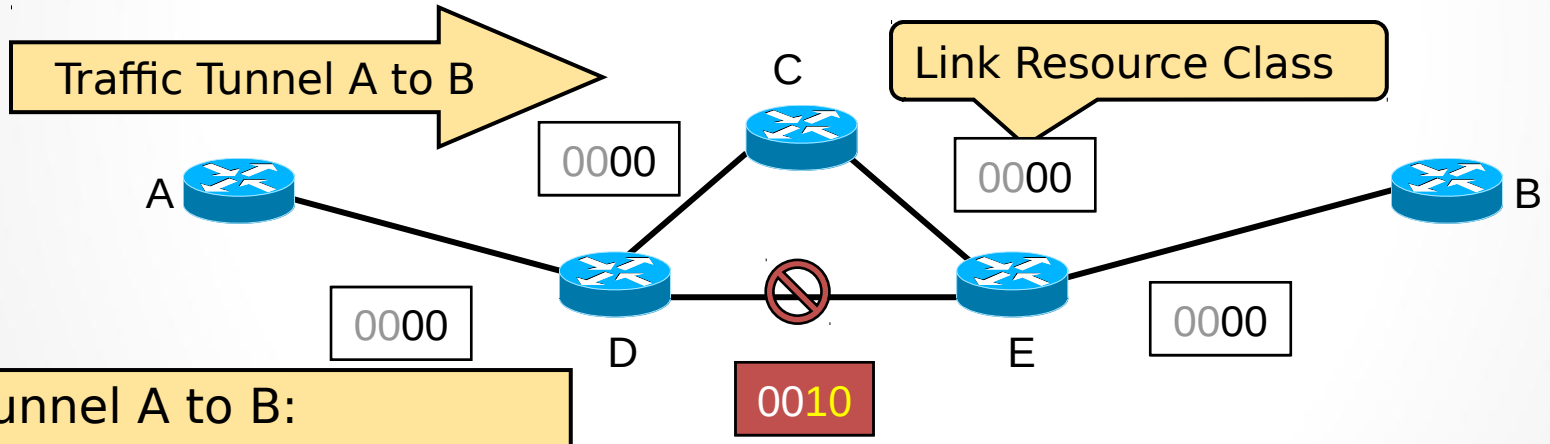

Tunnel A to B:

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

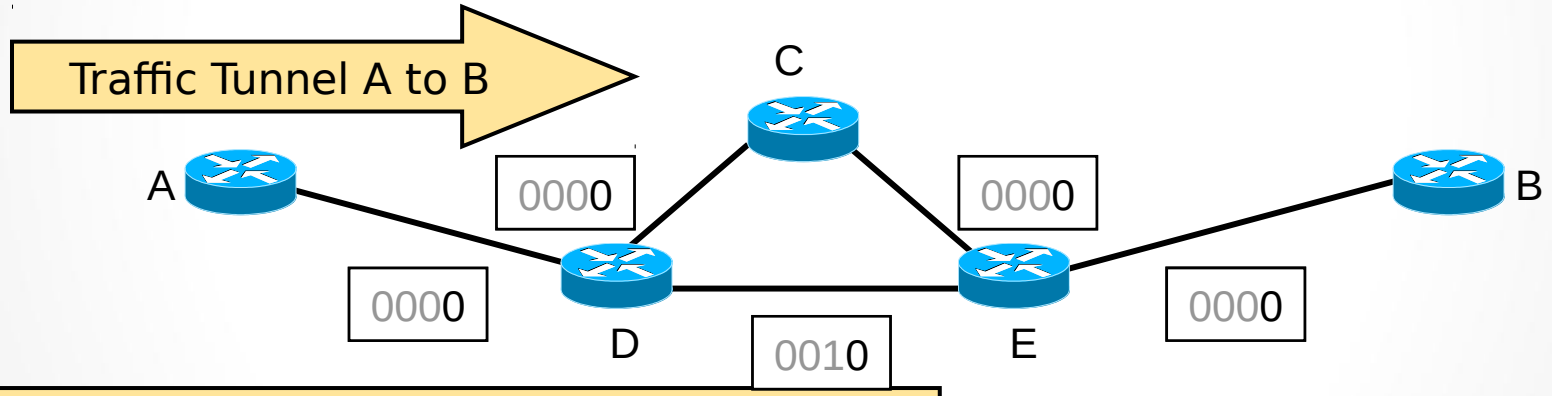

Tunnel A to B:

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

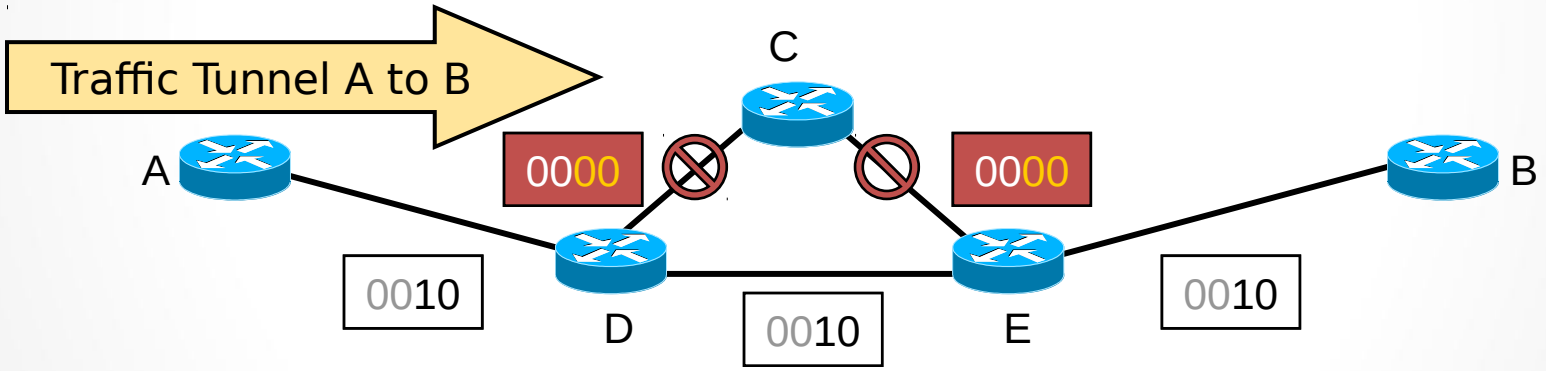

Tunnel A to B:
ADEB is possible.

Using Affinity Bits to Dedicate Links to Specific Purposes

Constraint-Based Path Computation (Cont.)

Request by tunnel:

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

Resource Affinity: bits = 0010, mask = 0011

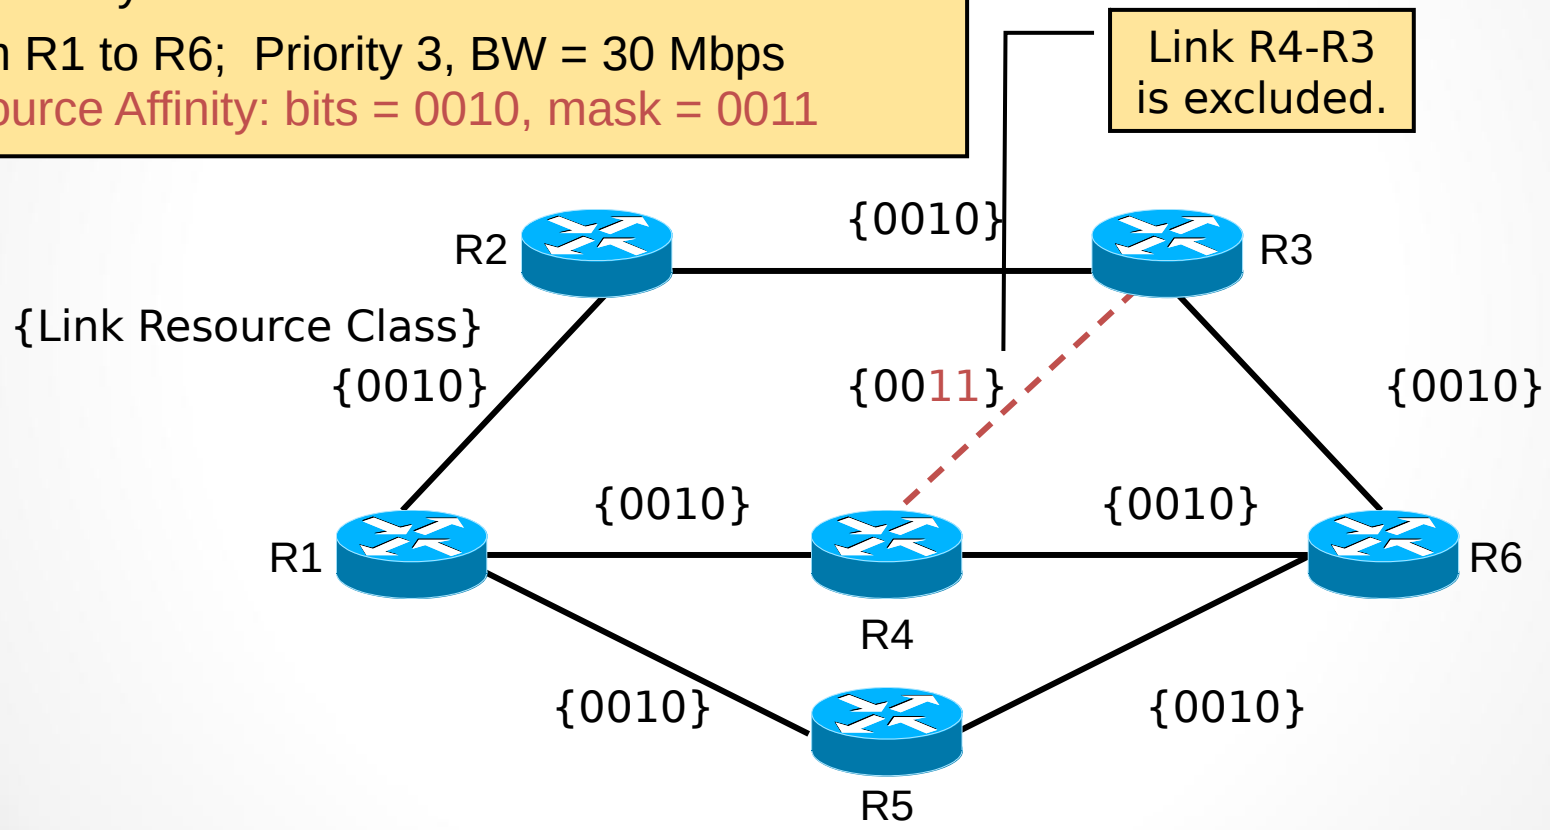

Path Selection Considering Policy Constraints

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

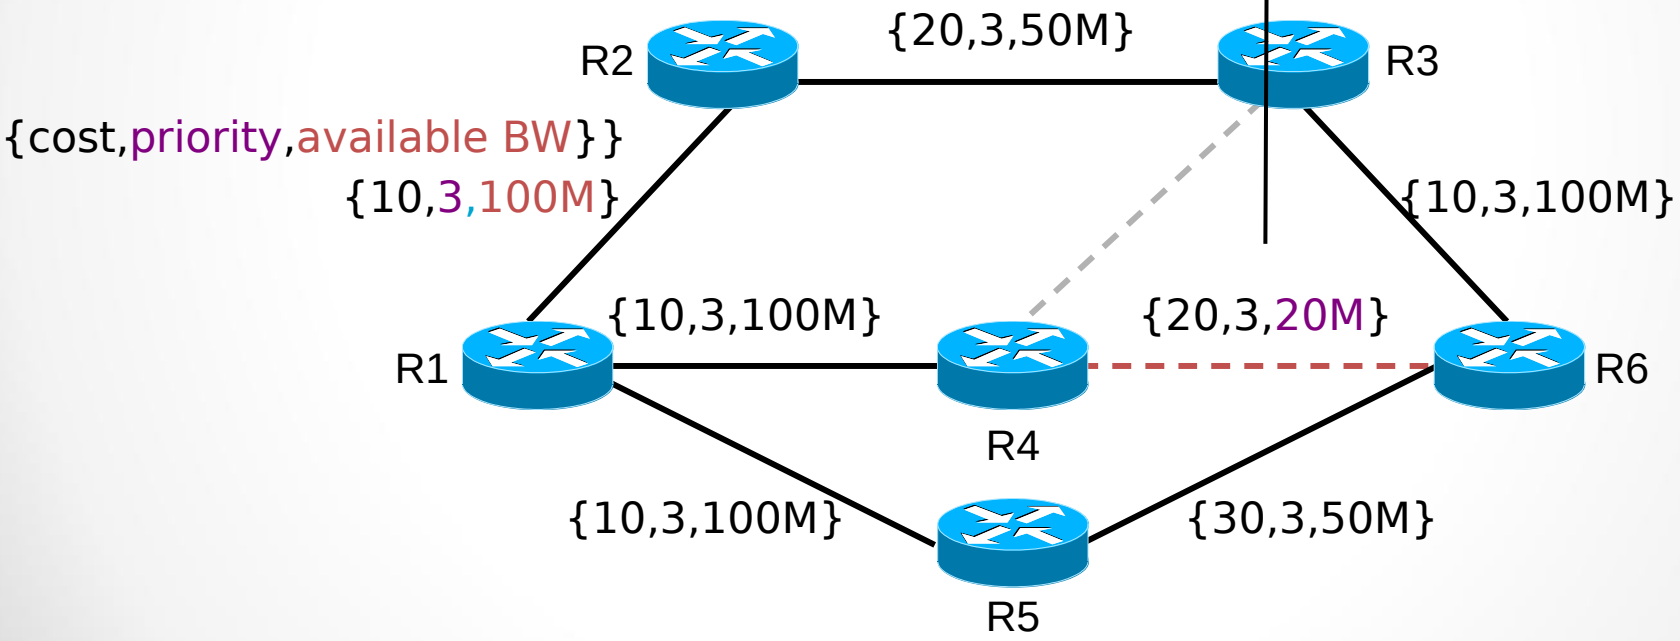

Path Selection Considering Available Resources

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.

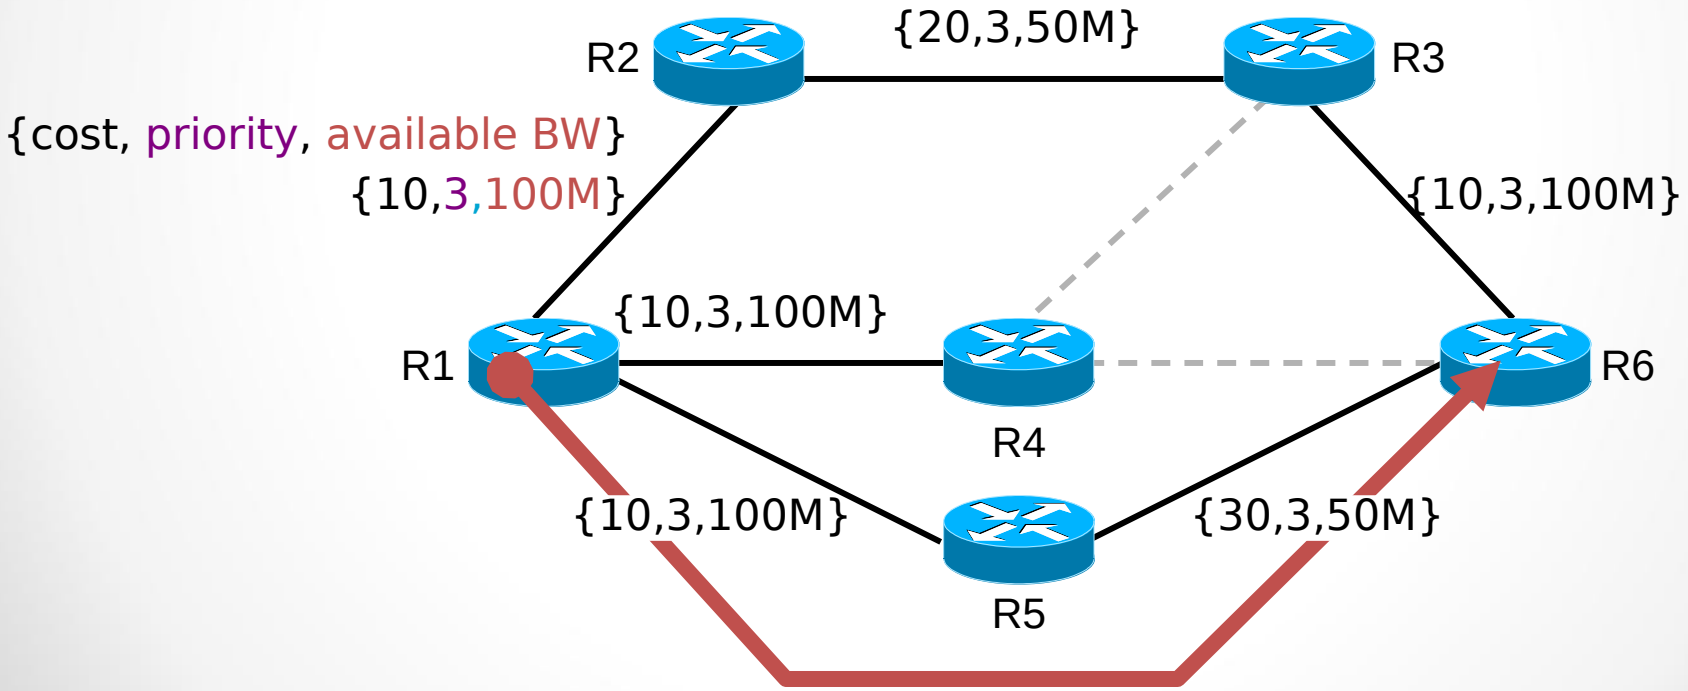

Selecting the Best Path

Traffic Flow Modifications

- In contrast to LDP LSP, on which traffic runs by default, we need to direct traffic in TE-tunnels.

- Static route
- PBR
- IGP Shortcut
- Tunnel-policy - применяется для перенаправления исключительно трафика VPN в туннели.
  То есть в режиме настройки VPN (не важно, L2 или L3) указывается какой туннель должен быть использован.

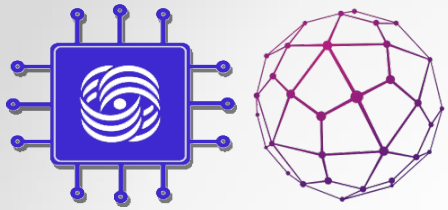

IGP Shortcut

- Этот способ наиболее распространённый и поддерживается почти всеми производителями.
- Маршрутизатор рассматривает туннель, как виртуальный интерфейс. И через этот интерфейс удалённые маршрутизаторы словно бы непосредственно подключены к локальному.

- С помощью IGP Shortcut мы вынуждаем протокол маршрутизации на Ingress LSR рассматривать туннель как обычную линию — Egress LSR будто бы подключен непосредственно. А соответственно и все сети, находящиеся за Egress LSR, будут доступны через туннель.
  Таким образом всё, чьей точкой назначения является этот маршрутизатор, или узлы за ним, будет отправлено в туннель. В том числе и VPN-пакеты.

Tunnel management methods

- MPLS te path metric
- The restriction on the bandwidth
- Explicit-Path
- The priorities of tunnels

Tunnel management methods

– The restriction on the bandwidth

- **Offline Bandwidth** - A method that uses a static setting of the required bandwidth value
- **Auto-Bandwidth** - This method involves tracking the tunnel load over a period of time and then adapting the reservation.

- Adjust Interval — the time during which the router monitors traffic and tracks peaks.
- Adjust Threshold-the threshold after which RSVP overwrites the reservation.

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).

Fast Reroute

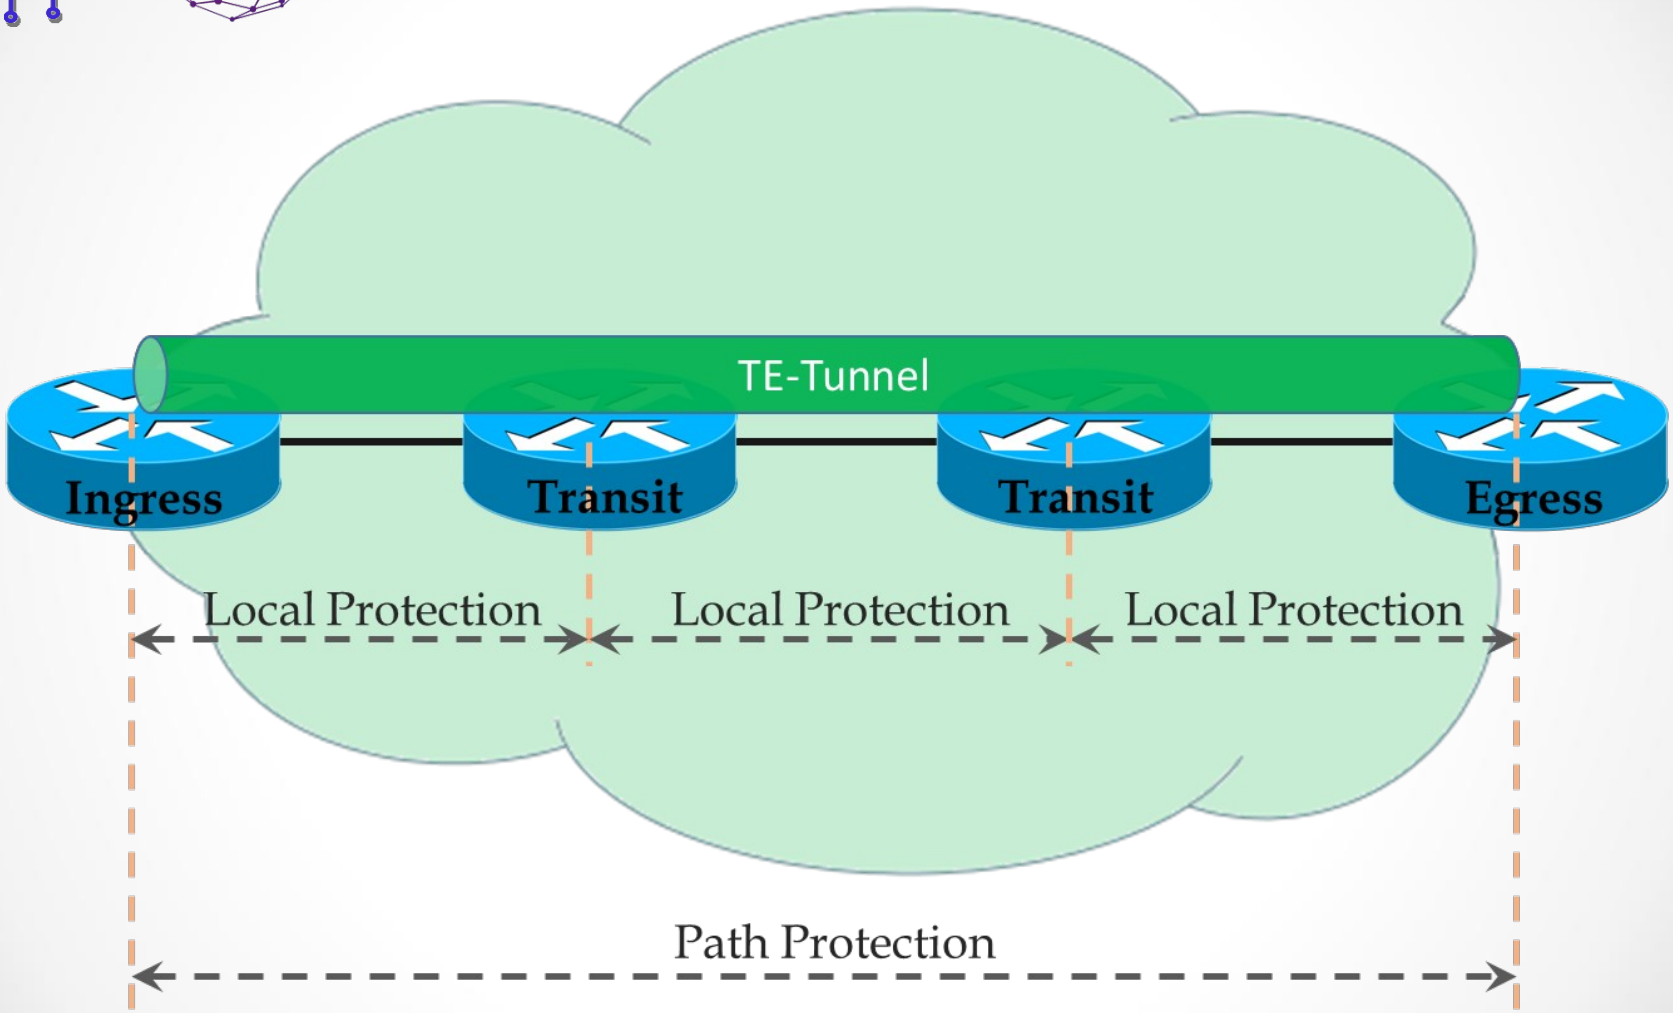

Fast Reroute

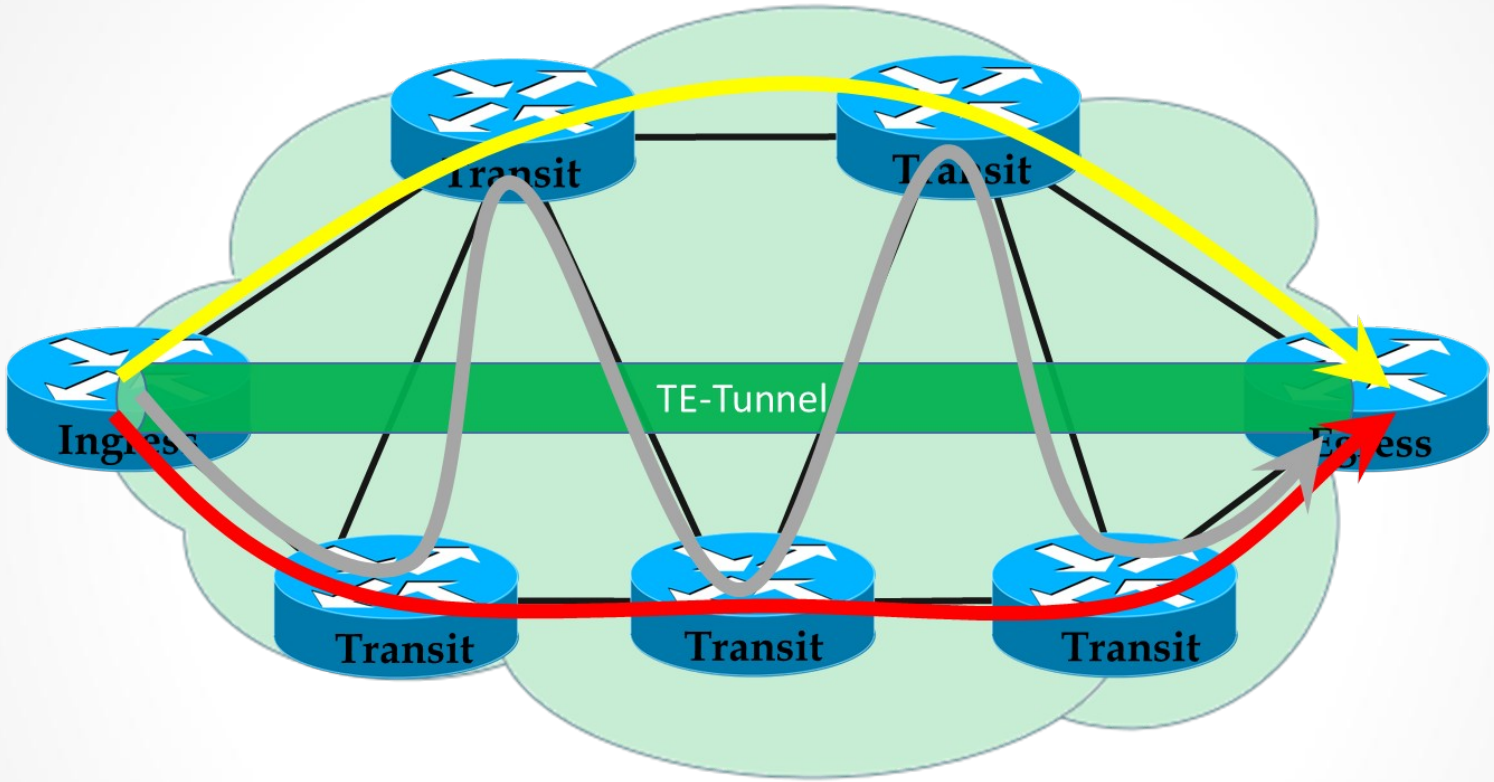

Primary CR-LSP
Основной

Secondary CR-LSP
Запасной

Best-Effort CR-LSP
Как получится

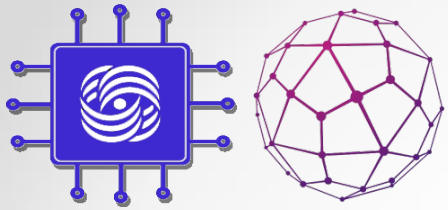

Fast Reroute

Primary — это основной LSP, который и будет использоваться для передачи трафика.
Secondary — запасной LSP. Если Ingress PE узнаёт от падении основного — он переводит трафик на запасной.
Последний в свою очередь тоже может быть:
Standby — всегда наготове: путь заранее вычислен и LSP сигнализирован. То есть он сразу готов подхватить трафик. Может также называться Hot-standby.
Non-standby — путь заранее вычислен, но LSP не сигнализирован. При падении основного LSP Ingress LSP сначала с помощью RSVP-TE строит запасной, потом пускает в него трафик. Зато полоса не простаивает зарезервированная. Может называться Ordinary.
Best Effort — если основной и запасной пути сломались или не могут быть удовлетворены условия, то RSVP-TE построит хоть как-нибудь без резервирования ресурсов.

Fast Reroute

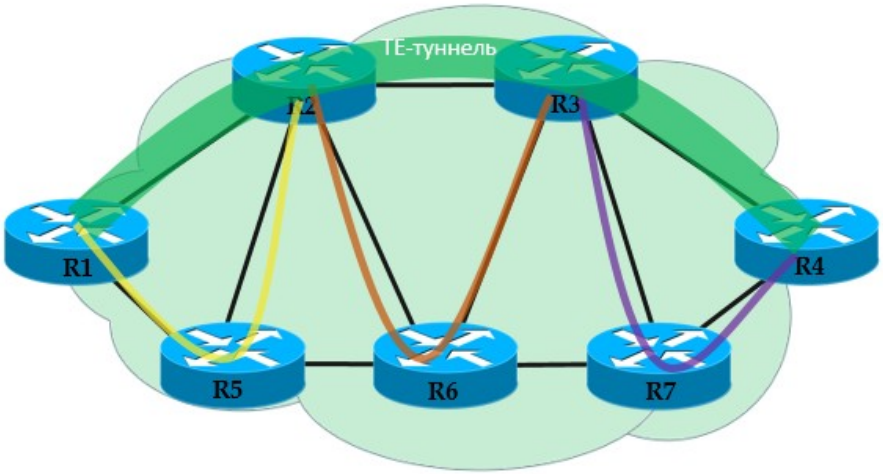

а) Защита линии (Link Protection)

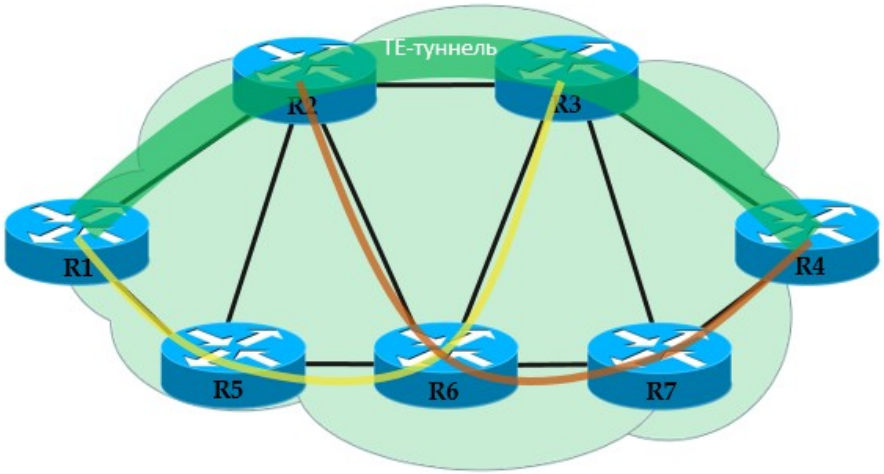

б) Защита узла (Node Protection)

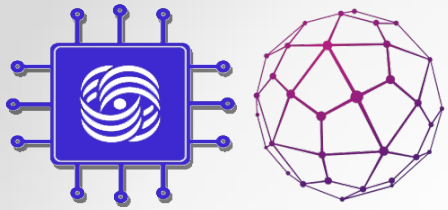

Fast Reroute

Терминология

PLR — Point of Local Repair — Это узел, который инициирует защиту линии или узла. Может быть Ingress PE или любой транзитный P, но не Egress PE
MP — Merge Point — точка схода, куда приходит защитный туннель. Любой транзитный P или Egress PE, но не Ingress PE.
Primary LSP или **Protected LSP** — исходный LSP, который требует защиты.
Bypass LSP — защитный LSP.
NHOP — Next Hop — следующий после PLR узел в Primary LSP.
NNHOP — Next Next Hop — соответственно следующий узел после Next Hop.

Fast Reroute

FRR Link Protection

Задача FRR — спасти пакеты, которые уже передаются, уведя их на Bypass LSP.

Когда PLR замечает, что линия, через которую лежит транзитный LSP, упала, он перенаправляет трафик. Не Ingress PE, а именно тот узел, на котором произошёл обрыв. Падение линка фиксируется по падению интерфейса или BFD-сессии.

Чтобы так быстро перенаправить пакеты, Bypass LSP должен быть построен заранее.

Каждый узел по ходу Primary LSP ищет, как обойти падение следующего линка и падение следующего узла. То есть он запускает полный механизм построения LSP: CSPF до MP (NHOP для случая падения линка и NNHOP для случая падения узла) Отправляет по просчитанному пути RSVP PATH с запросом резервирования. Получает RSVP RESV, если резервирование удалось.

Туннели строятся автоматически и не отображаются в конфигурации. Но в остальном это обычные туннели.
Технически, обычный FRR требует ручной настройки всех Bypass-туннелей. То есть оператор сам должен предусмотреть все места, где может что-то сломаться и настроить резервирование.

Существует механизм AutoTunnel, который инструктирует каждый узел на пути Primary LSP самостоятельно и автоматически рассчитывать Bypass-туннели. Включается он на Ingress LSR командой

Fast Reroute