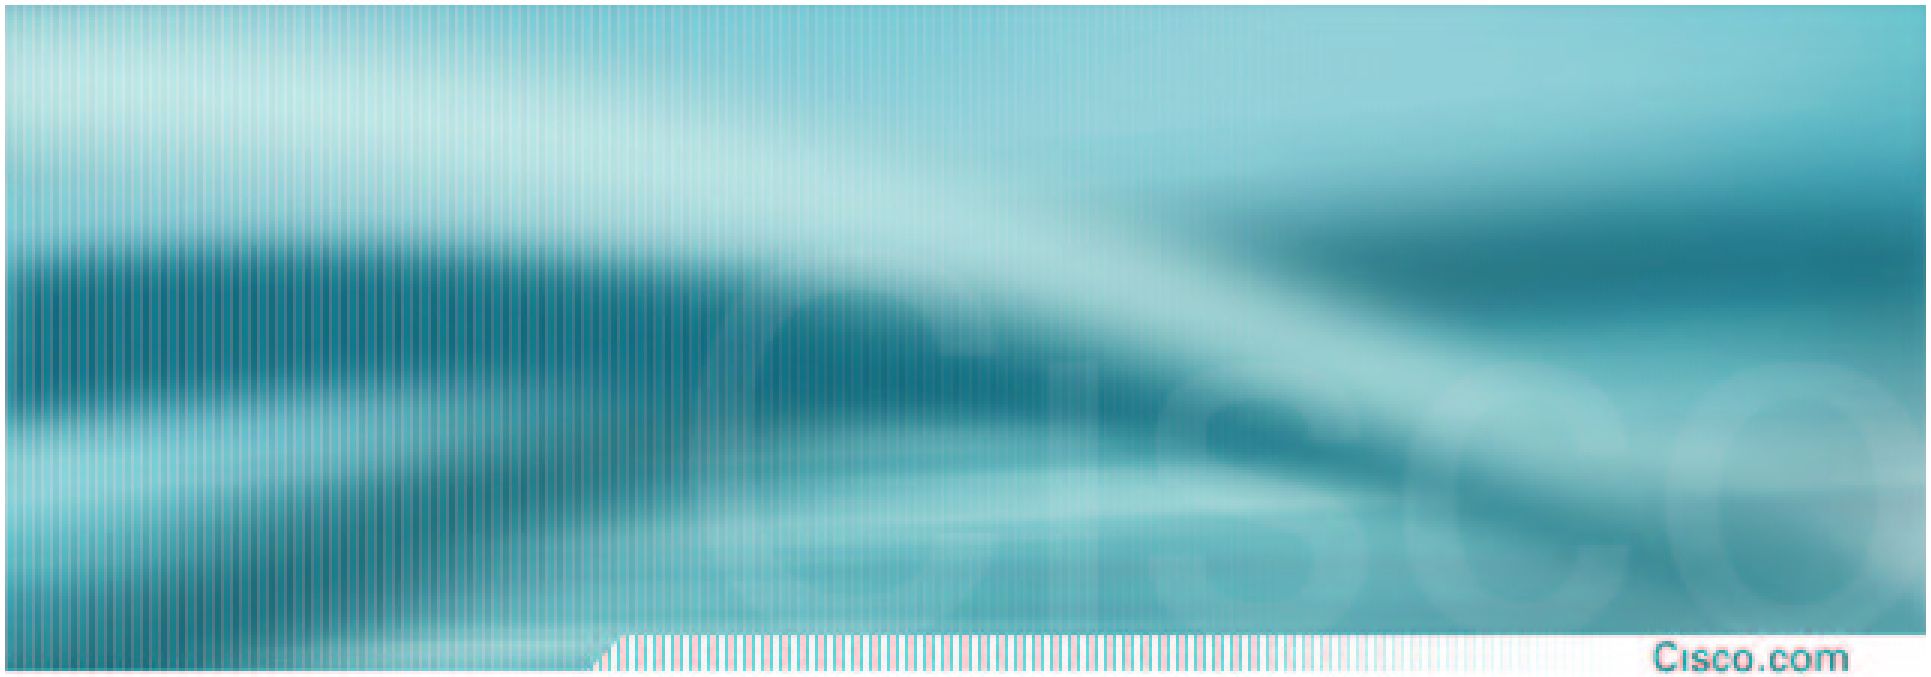

Traffic Engineering & HA for MPLS Networks

Francisco Bolaños (fbolanos@cisco.com)

Agenda

Cisco.com

Application 1: Increasing Bandwidth Inventory

- **Application 2: Minimizing Packet Loss**
- **Application 3: Optimizing the Core**
- **Deployment**

IP Routing and the Fish Problem

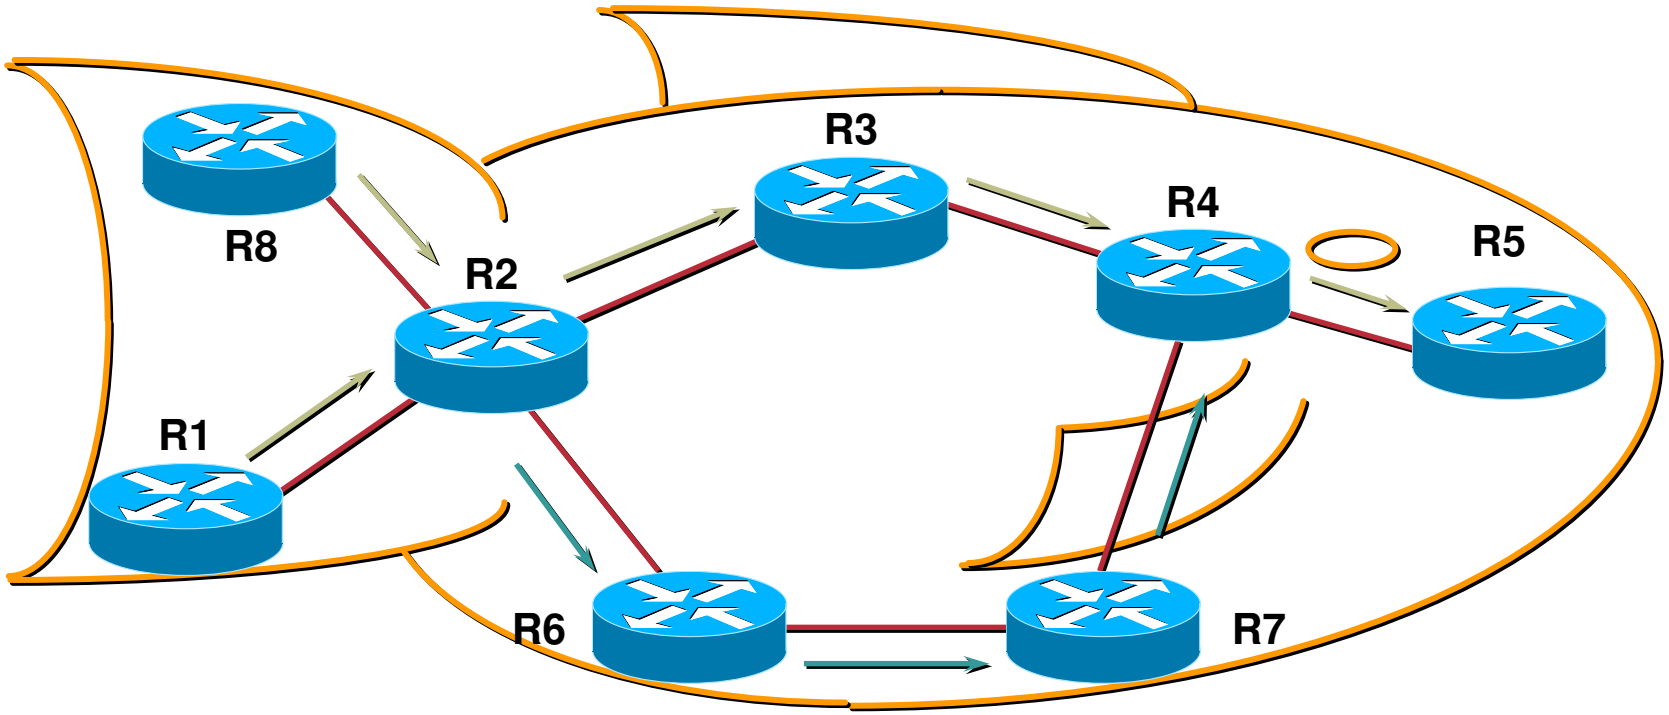

IP (Mostly) Uses Destination-Based Least-Cost Routing
Flows from R8 and R1 Merge at R2 and Become Indistinguishable
From R2, Traffic to R3, R4, R5 Use Upper Route

Alternate Path Under-Utilized

The Problem with Shortest-Path

Node	Next-Hop	Cost
B	B	10
C	C	10
D	C	20
E	B	20
F	B	30
G	B	30

- Some links are DS3, some are OC-3
- Router A has 40Mb of traffic for Route F, 40Mb of traffic for Router G
- Massive (44%) packet loss at Router B->Router E!
  - Changing to A->C->D->E won't help

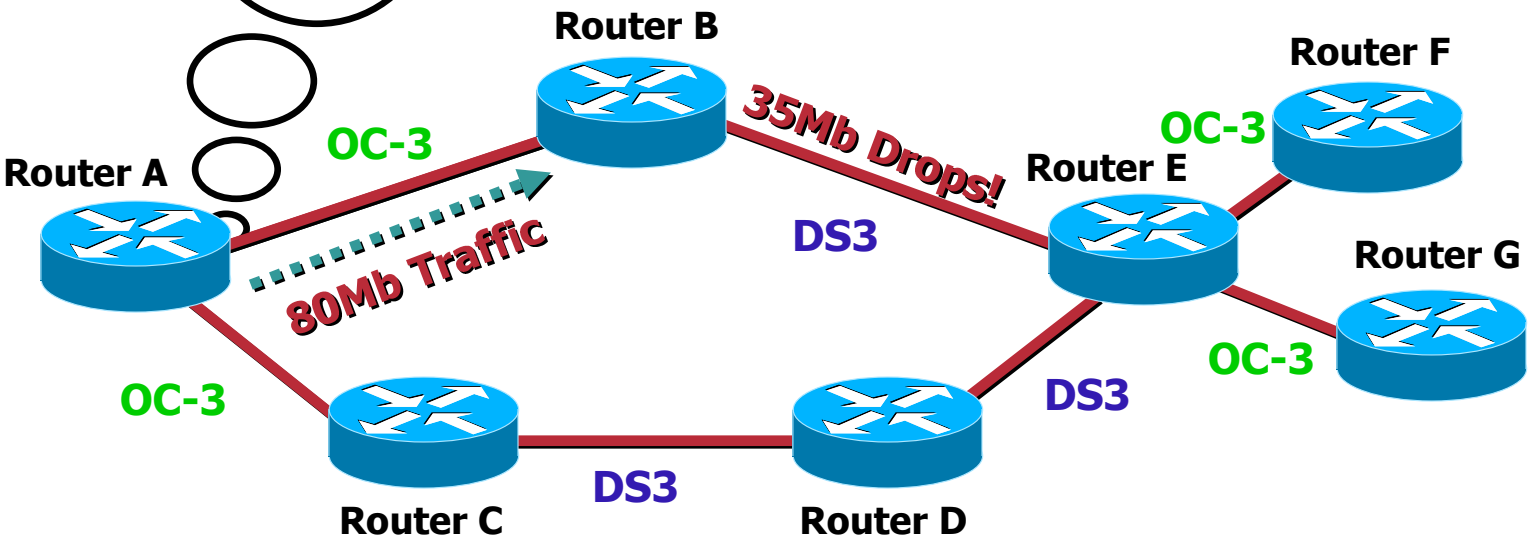

How MPLS TE Solves the problem

Node	Next-Hop	Cost
B	B	10
C	C	10
D	C	20
E	B	20
F	Tunnel 0	30
G	Tunnel 1	30

- Router A sees all links
- Router A computes paths on properties other than just shortest cost
- No link oversubscribed!

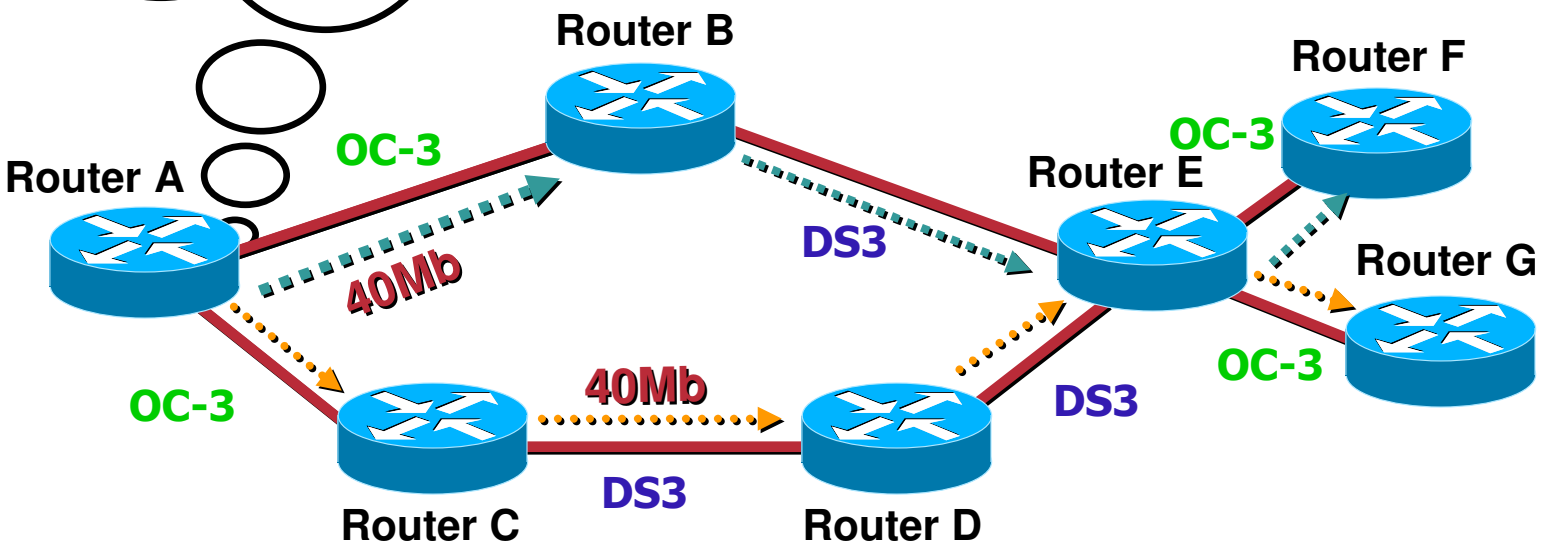

TE Fundamentals – "Building Blocks"

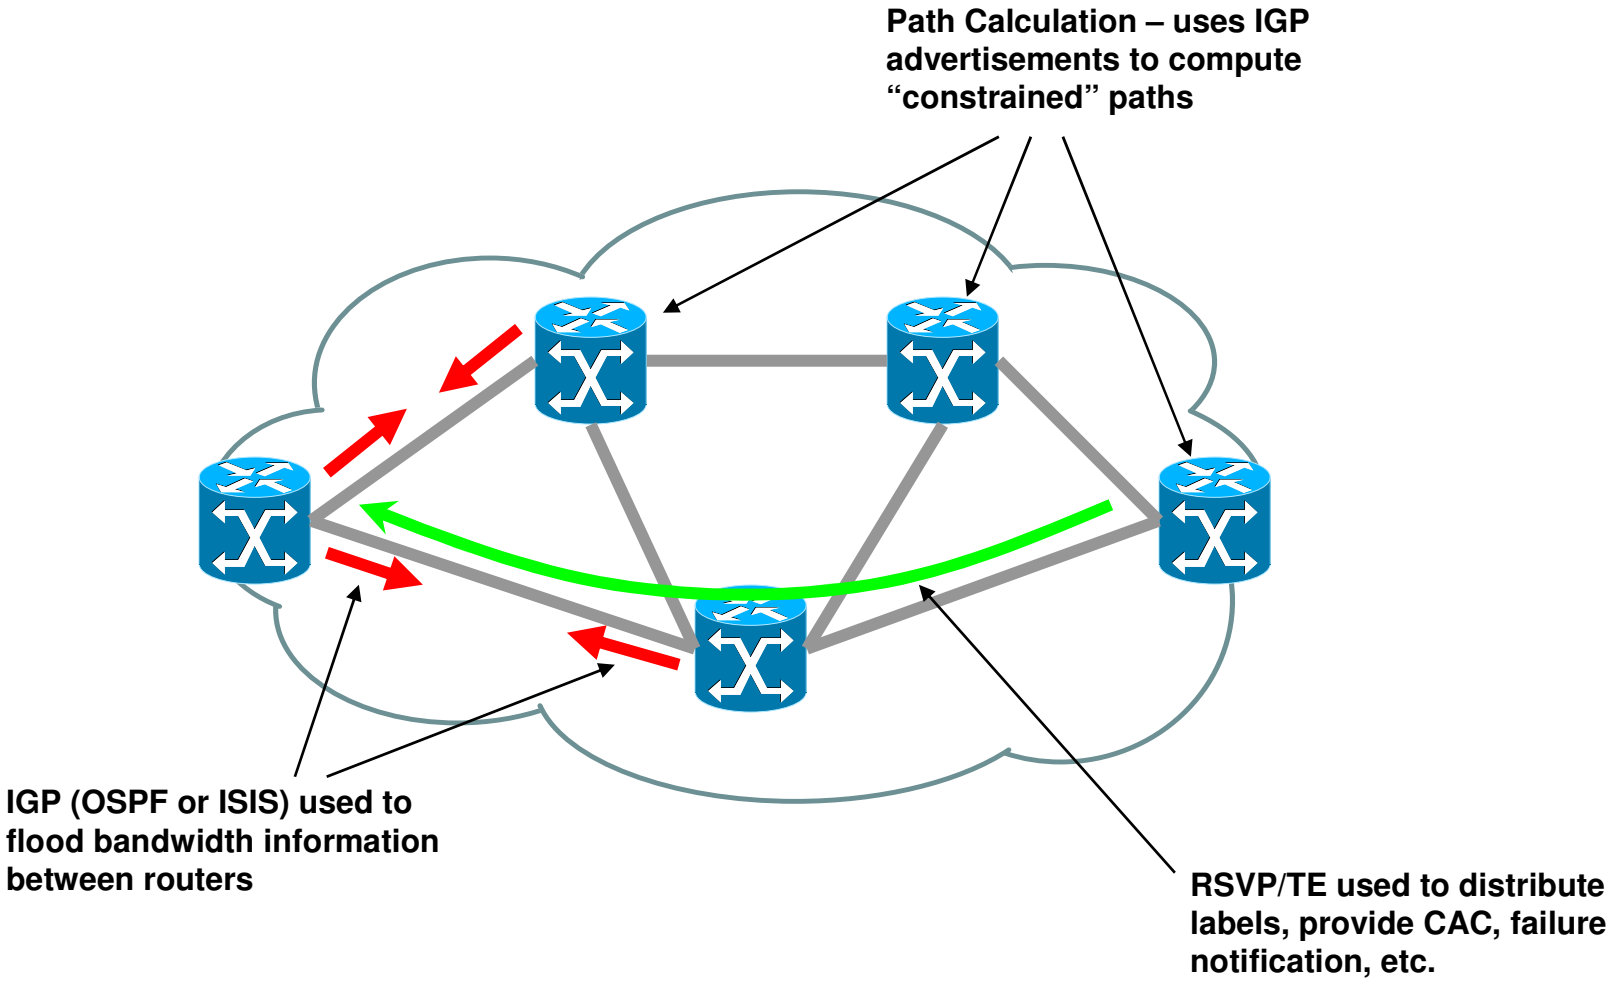

Path Calculation (PCALC)

Node	Next-Hop	Cost
B	B	10
C	C	10
D	C	20
E	B	20
F	Tunnel 0	30
G	Tunnel 1	30

- PCALC takes bandwidth, other constraints into account
- Paths calculated, resources reserved if necessary
- End result: Bandwidth used more efficiently!

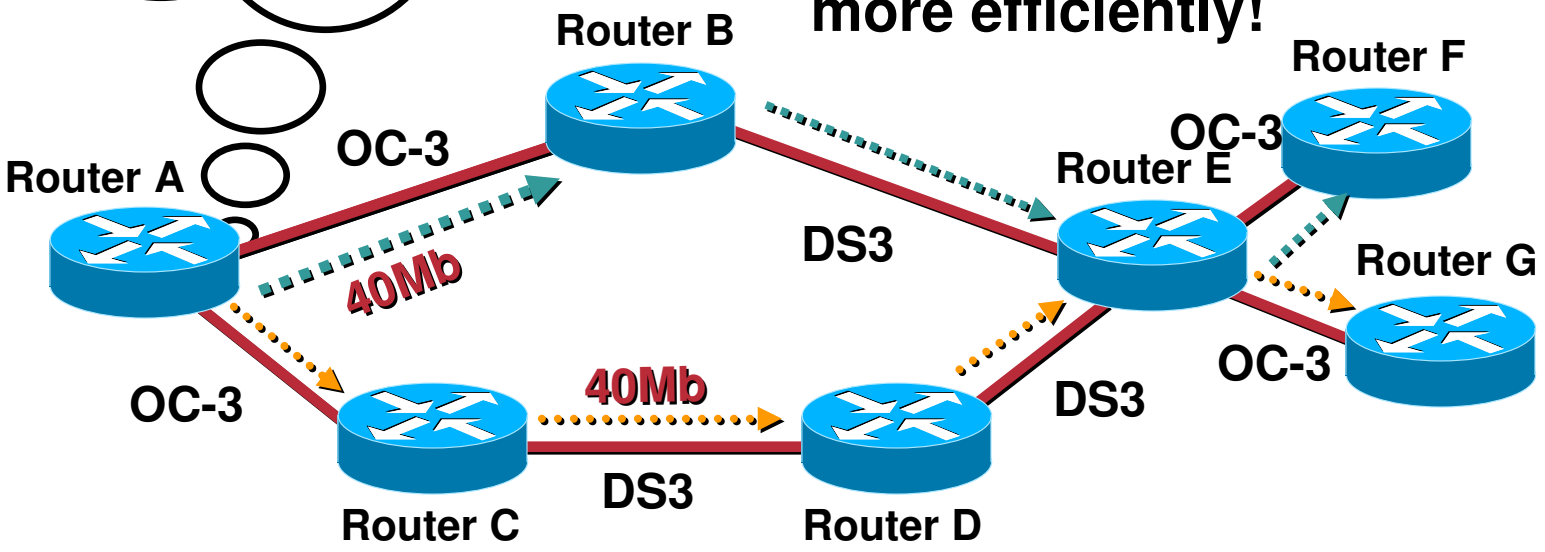

Path Setup

- PATH message: “Can I have 40Mb along this path?”
- RESV message: “Yes, and here’s the label to use”
- Labels are installed along each hop

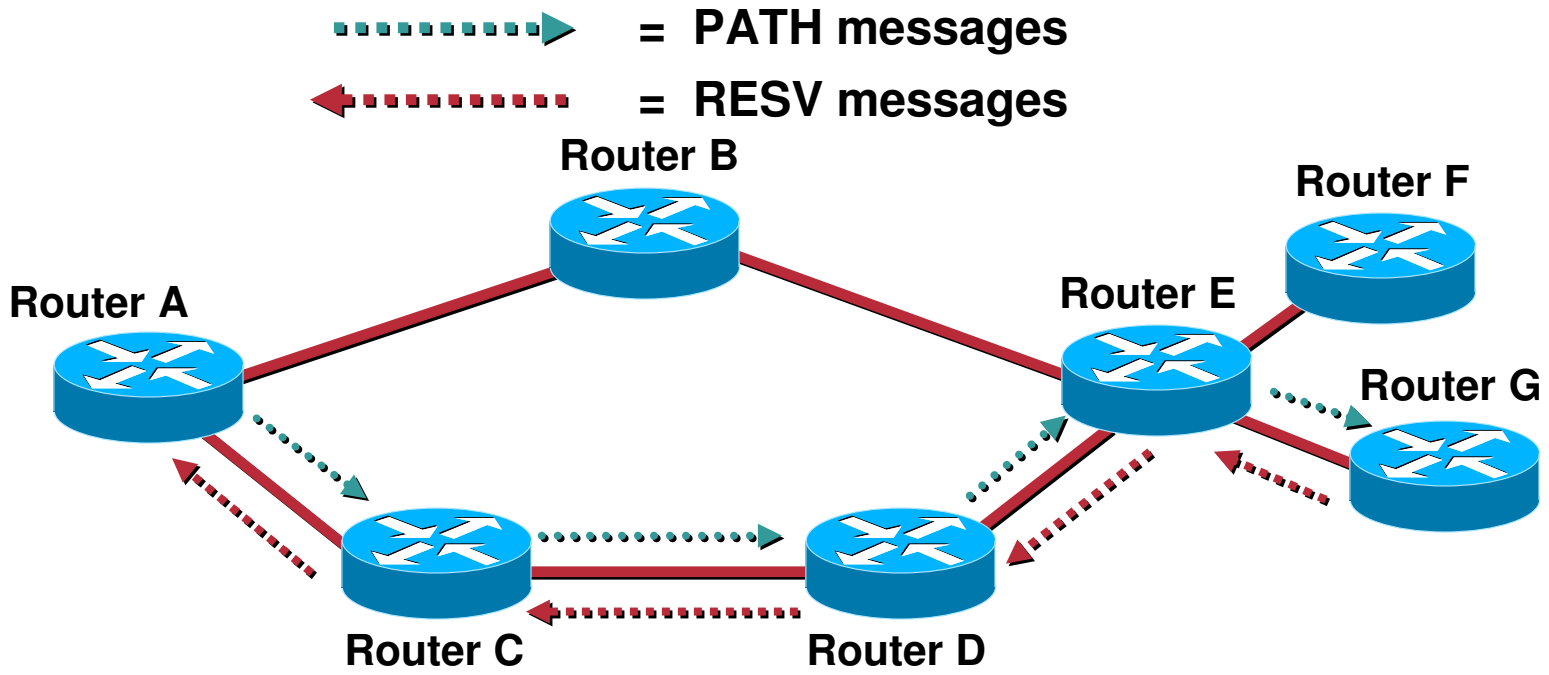

Auto-Route

Routing Table

Node	Next-Hop	Cost
B	B	10
C	C	10
D	C	20
E	B	20
F	B	30
G	Tunnel 1	30
H	Tunnel 1	40
I	Tunnel 1	40

- **Router A's routing table, built via auto-route**
- **Everything "behind" the tunnel is routed via the tunnel**

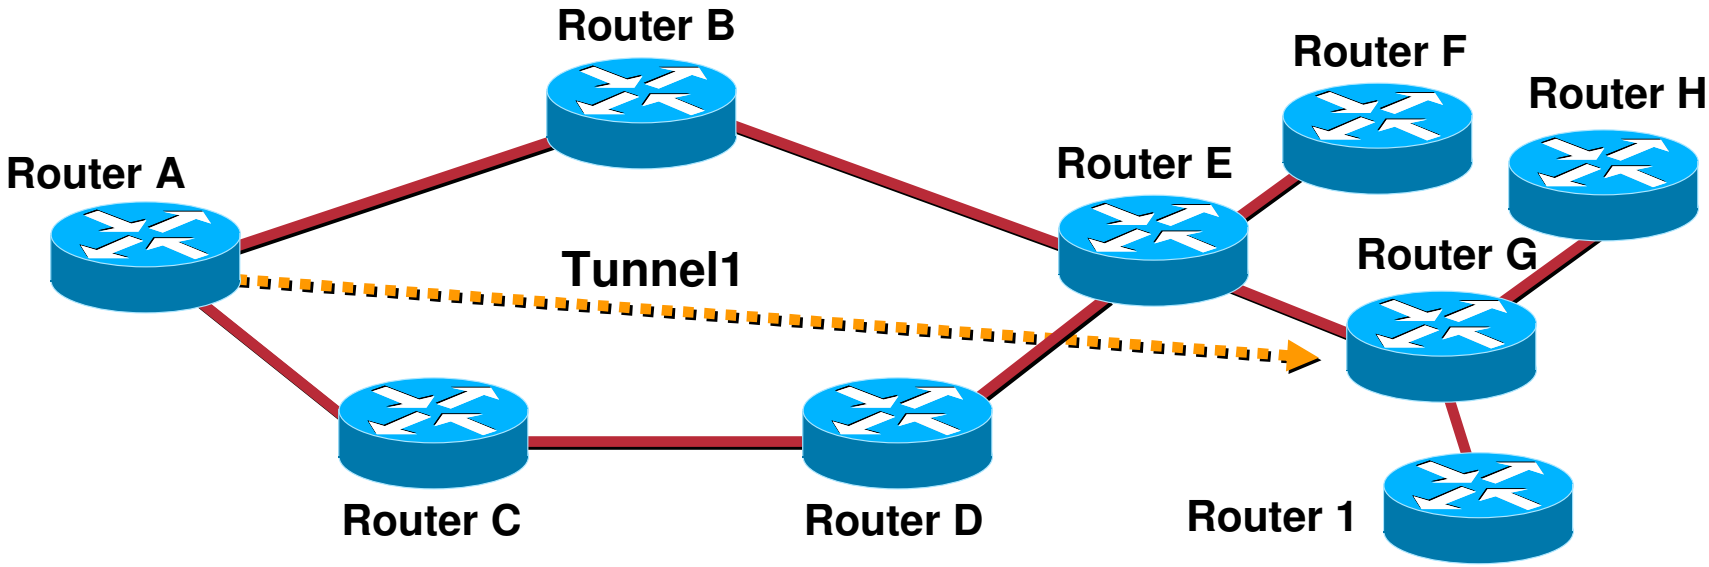

Static Routing

Routing Table

Node	Next-Hop	Cost
B	B	10
C	C	10
D	C	20
E	B	20
F	B	30
G	B	30
H	Tunnel 1	40
I	B	40

- Router H is known via the tunnel
- Router G is not routed to over the tunnel, even though it's the tunnel tail!

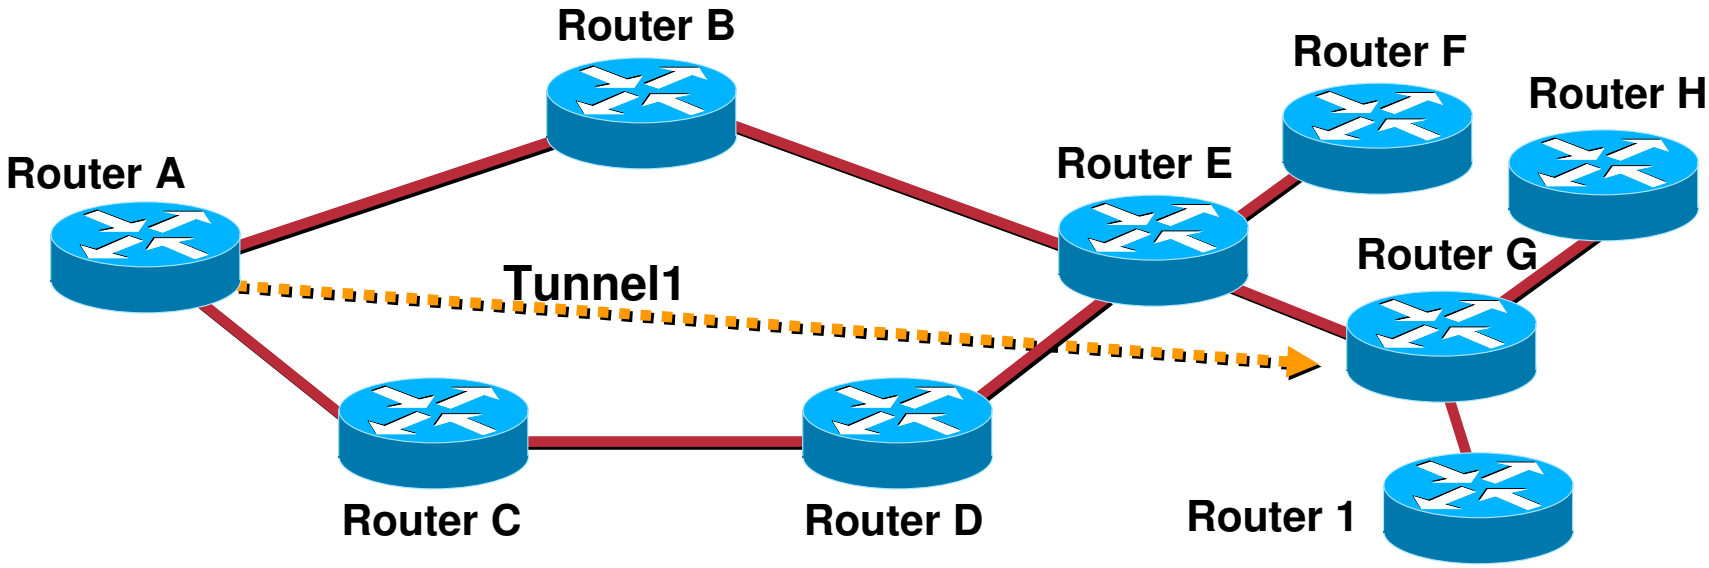

Policy Routing

Routing Table

Node	Next-Hop	Cost
B	B	10
C	C	10
D	C	20
E	B	20
F	B	30
G	B	30
H	B	40
I	B	40

- **Routing table isn't affected by policy routing**
- **Require 'set interface tunnel' within PBR to work**

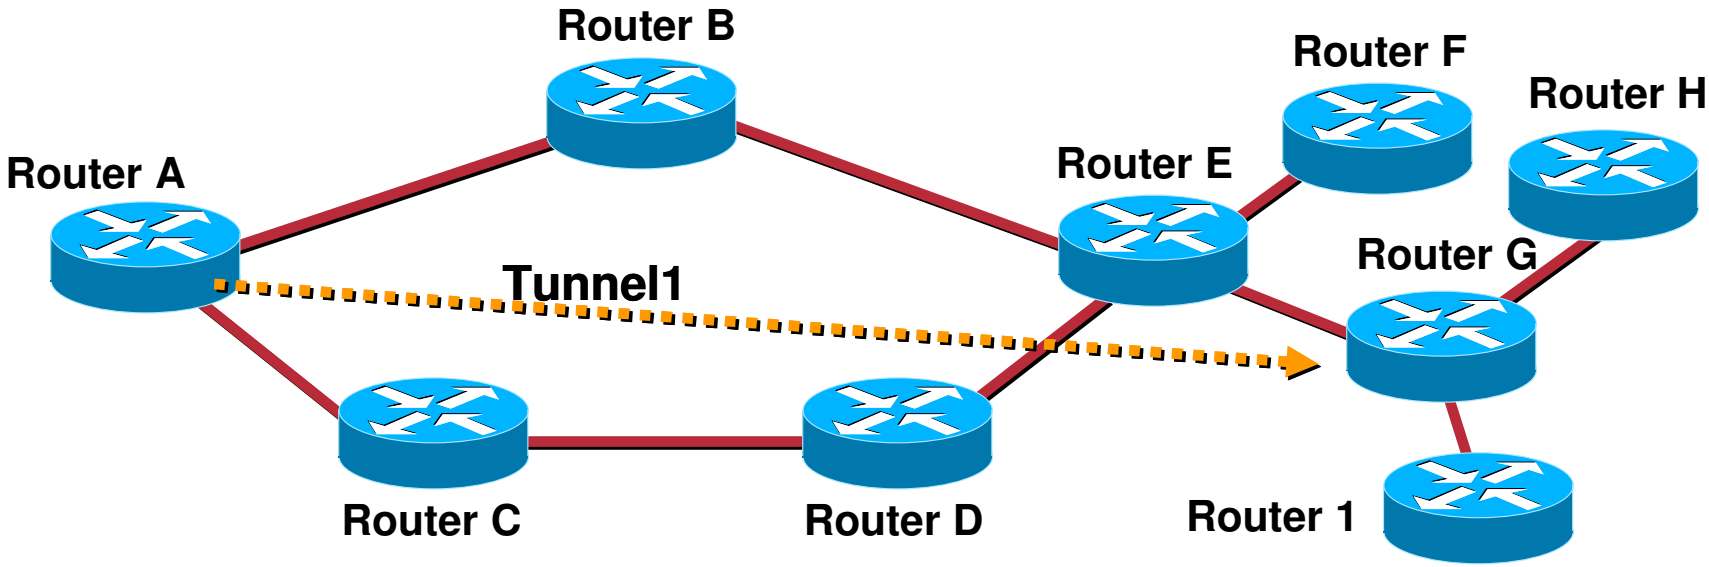

Agenda

Cisco.com

- Application 1: Increasing Bandwidth Inventory
- Application 2: Minimizing Packet Loss
- Application 3: Optimizing the Core
- Deployment

Link Protection

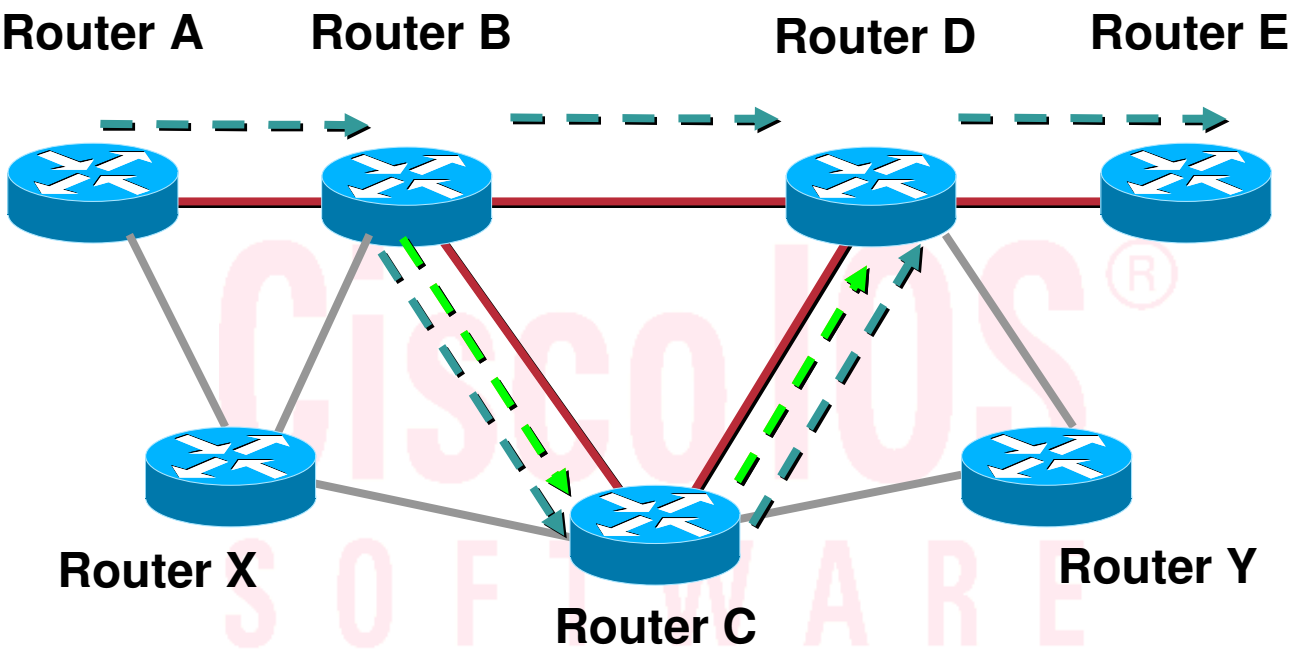

- Primary Tunnel: A -> B -> D -> E
- BackUp Tunnel: B -> C -> D (Pre-provisioned)
- Recovery = ~50ms

*Actual time varies – well below 50ms in lab tests, can also be higher

Node Protection

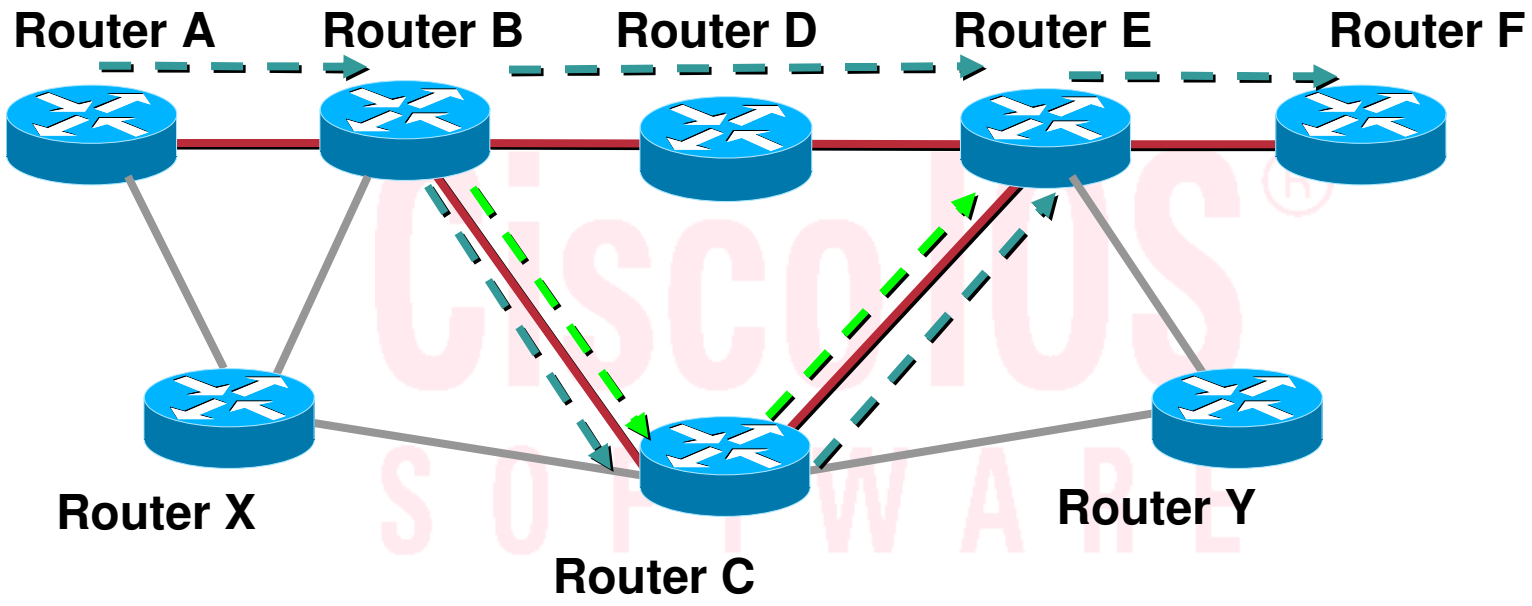

- Primary Tunnel: A -> B -> D -> E -> F
- BackUp Tunnel: B -> C -> E (Pre-provisioned)
- Recovery = ~100ms

What is Bandwidth Protection?

Subscribers want bandwidth & services from point A to B for Voice & Video traffic. They don't care what happens in the network – HOW it is offered by a Service Provider is secondary.

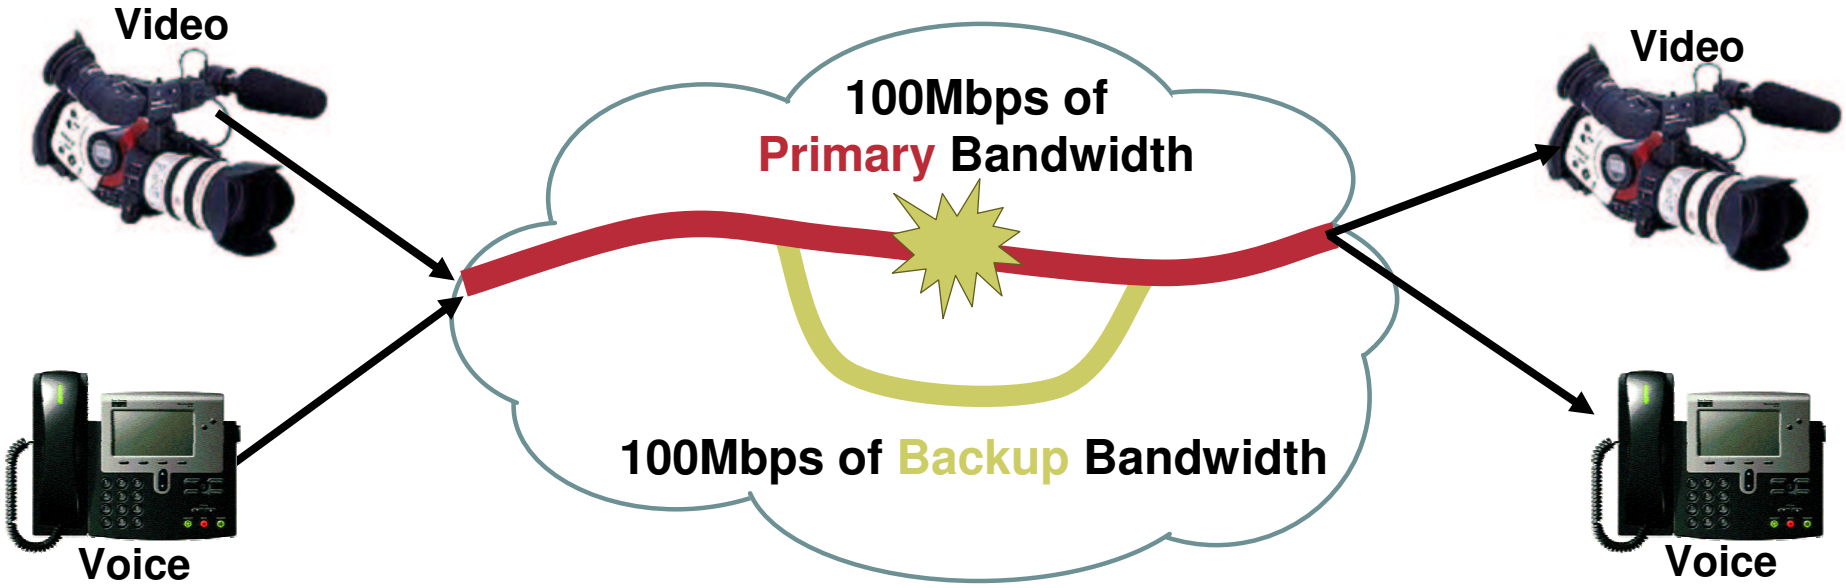

Bandwidth Protection is NOT a new problem – but using MPLS we have a new paradigm to provide a solution

Scenario 1: Backup Bandwidth Sharing

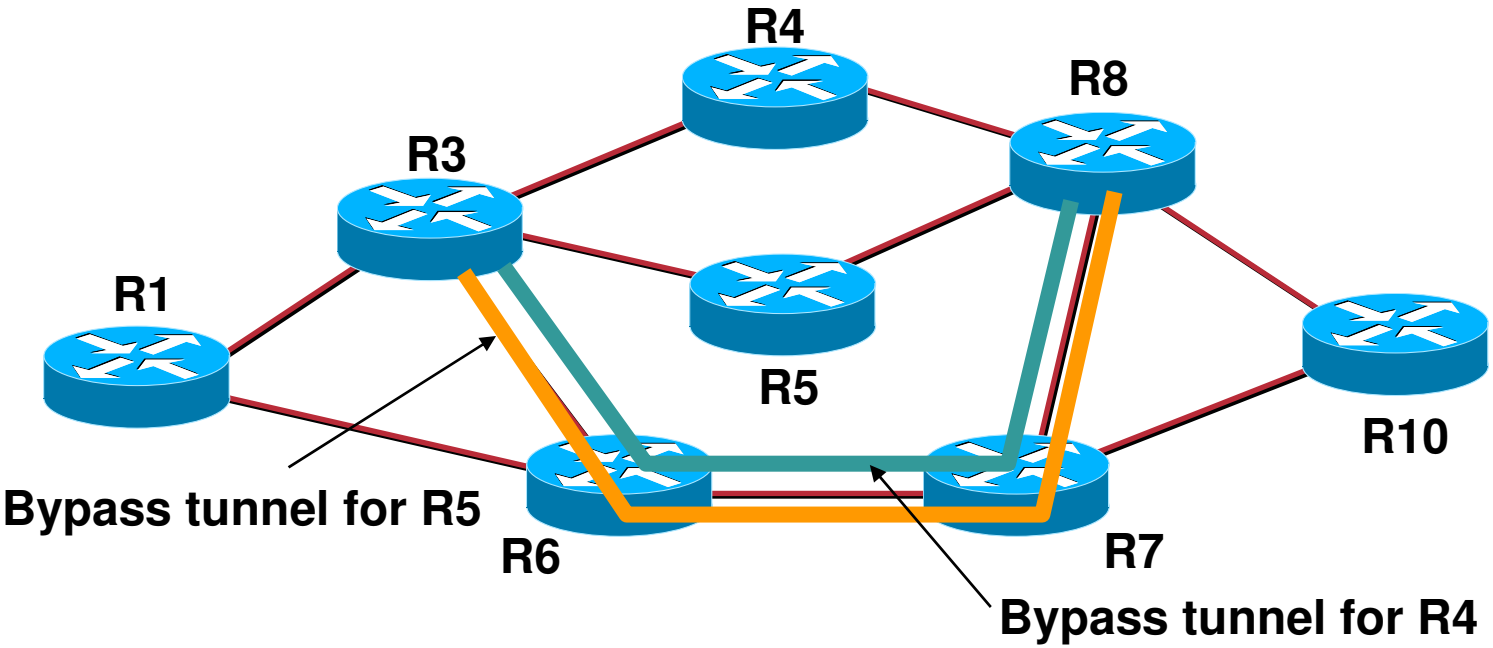

- Only need to allocate enough BW on R3-R6-R7-R8 to protect for a single node failure – "N:1" protection

Scenario 2: Backup Bandwidth Sharing

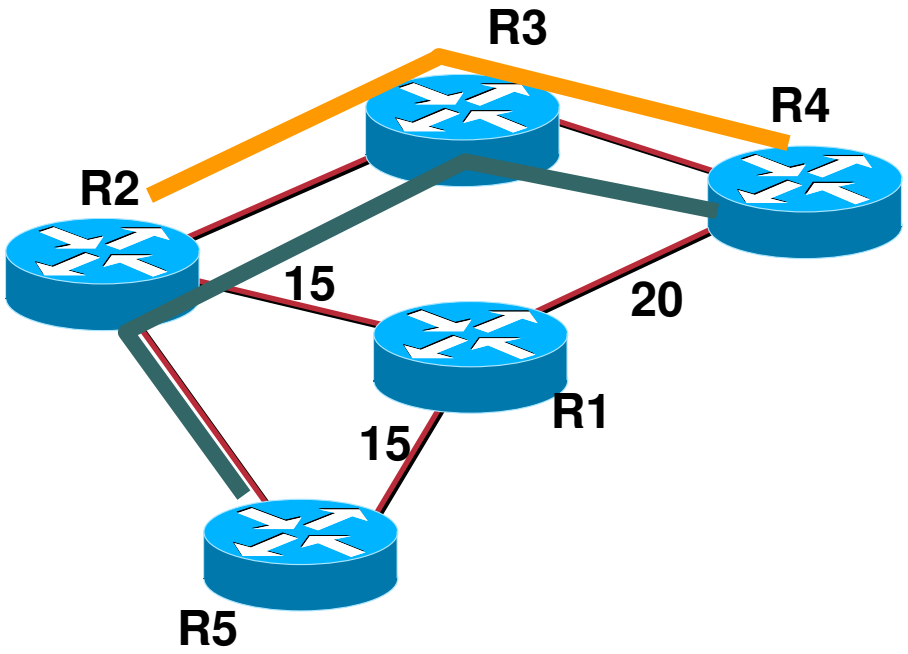

- Backup tunnels R5-R2-R3-R4 and R2-R3-R4 protect R1
- Naïve approach – each tunnel needs capacity 15
- Shared approach – allocate 20Mbps on R2-R3 and R3-R4; 15 Mbps on R5-R2

MPLS TE FRR – Backup Labels

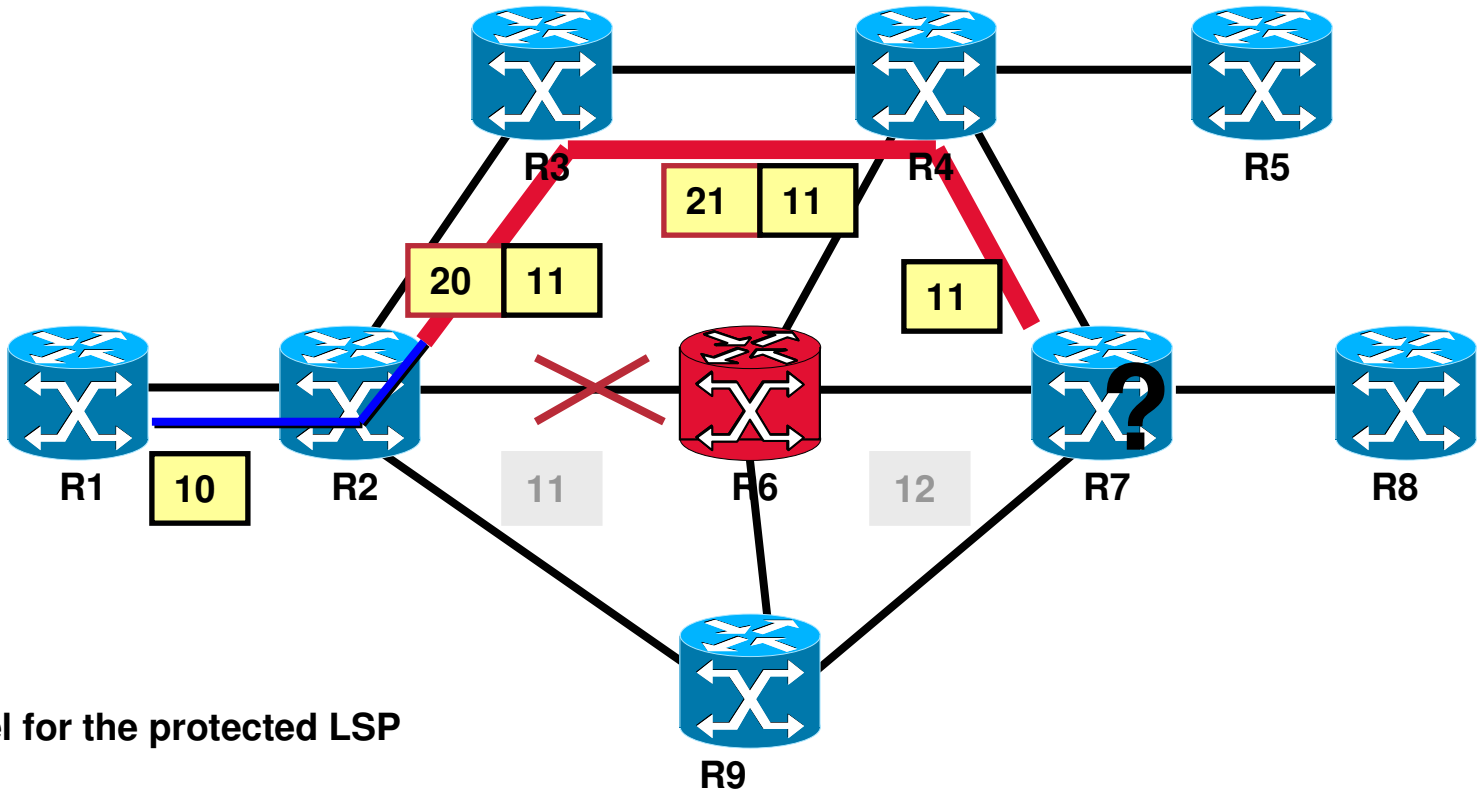

- **The Point of Local Repair learns the label to use from the Record Route Object object carried in the Resv message when the reroutable LSP is first established – Global Label Allocation must be used.**

Agenda

Cisco.com

- Application 1: Increasing Bandwidth Inventory
- Application 2: Minimizing Packet Loss
- Application 3: Optimizing the Core
- Deployment

What is DiffServ aware Traffic Engineering?

- Used when there exist multiple diverse links
- Create TE tunnels on a Per-Class basis
- One TE Tunnel for Voice, another for Data

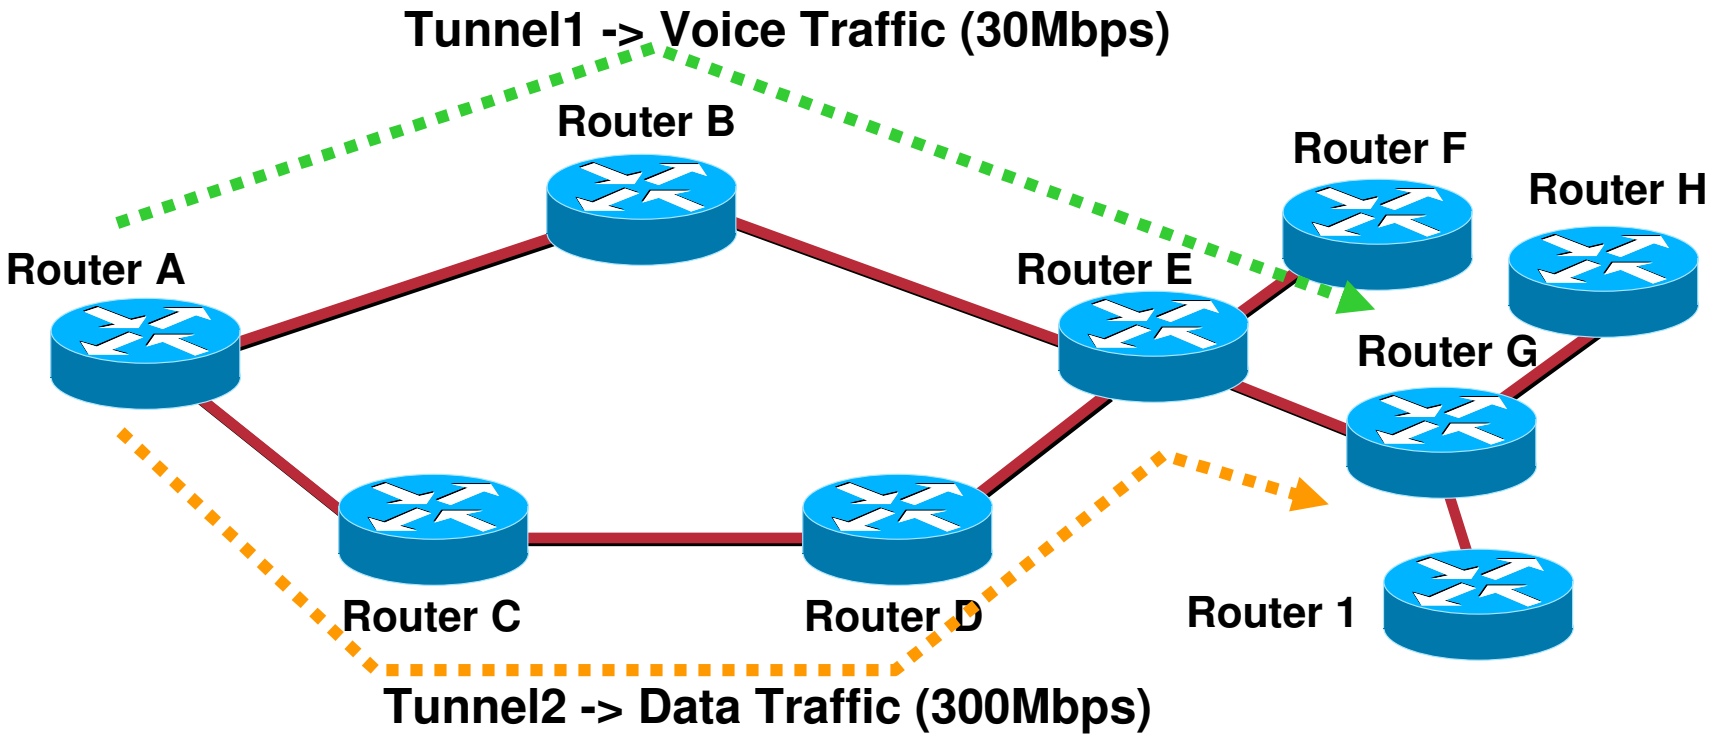

Do I need DS-TE in my network?

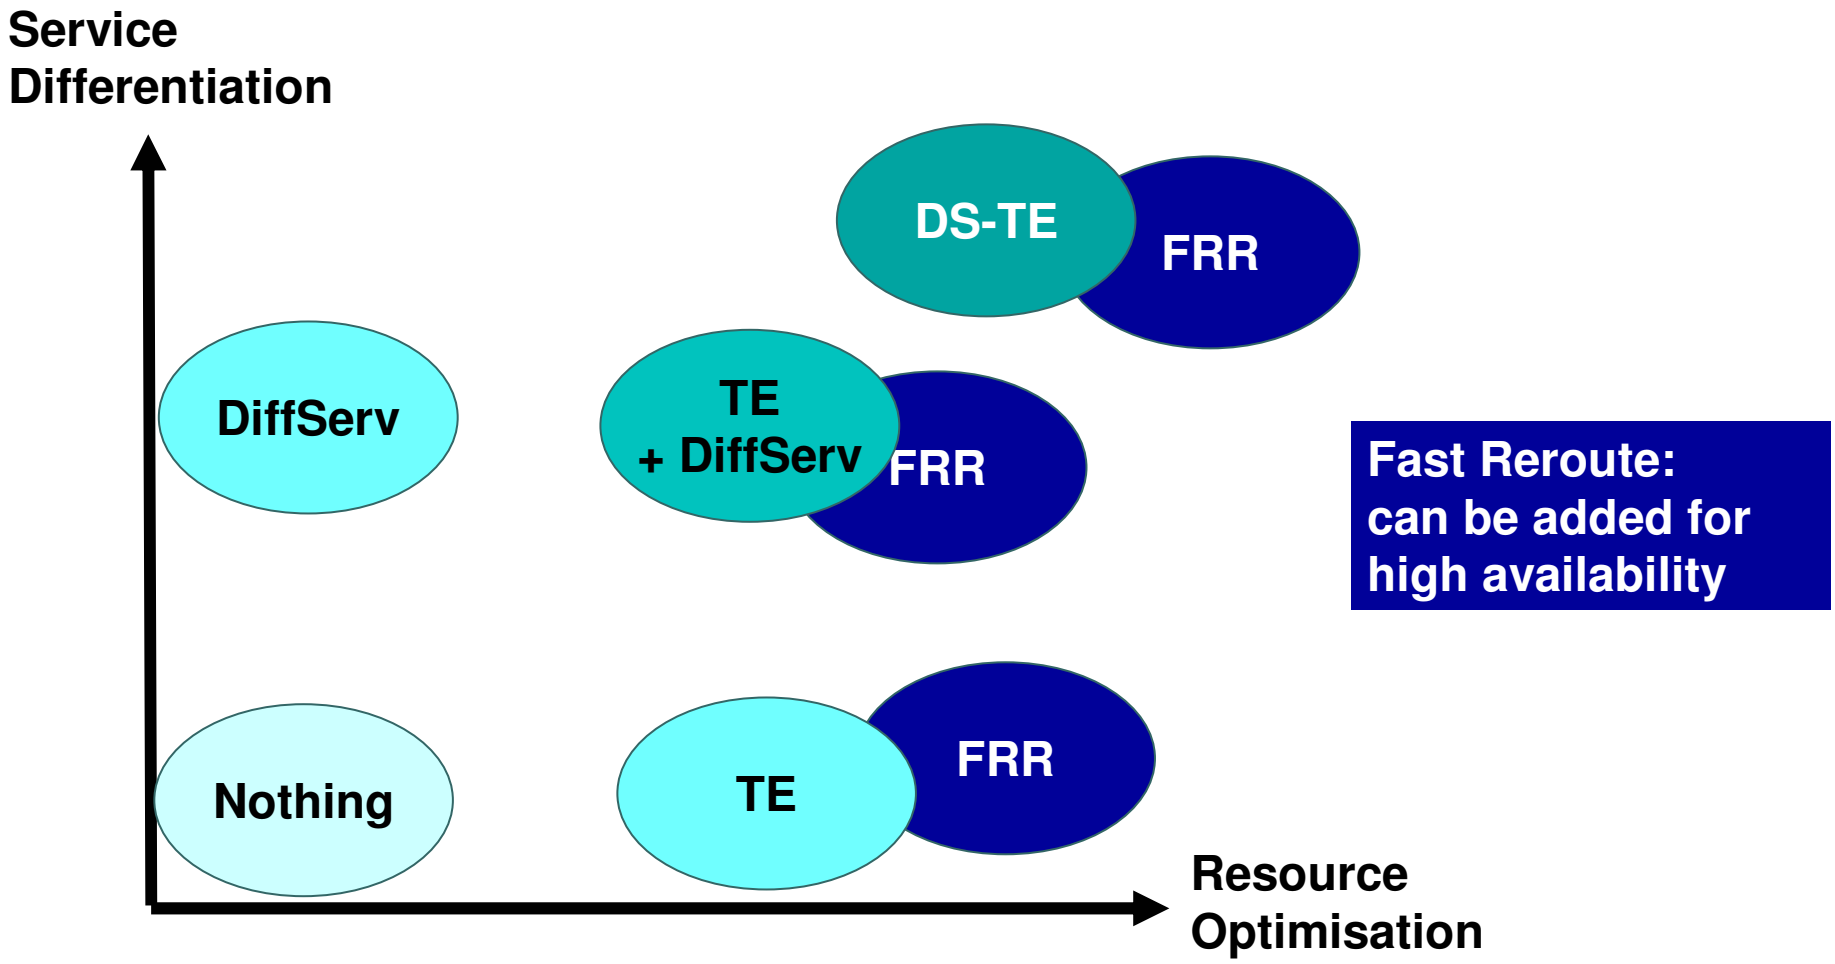

Agenda

Cisco.com

- MPLS Fundamentals
- Application 1: Increasing Bandwidth Inventory
- Application 2: Minimizing Packet Loss
- Application 3: Optimizing the Core
- Deployment

Tactical TE Deployment

Cisco.com

Requirement: Need to handle scattered congestion points in the Network
Solution: Deploy MPLS TE on only those nodes that face congestion

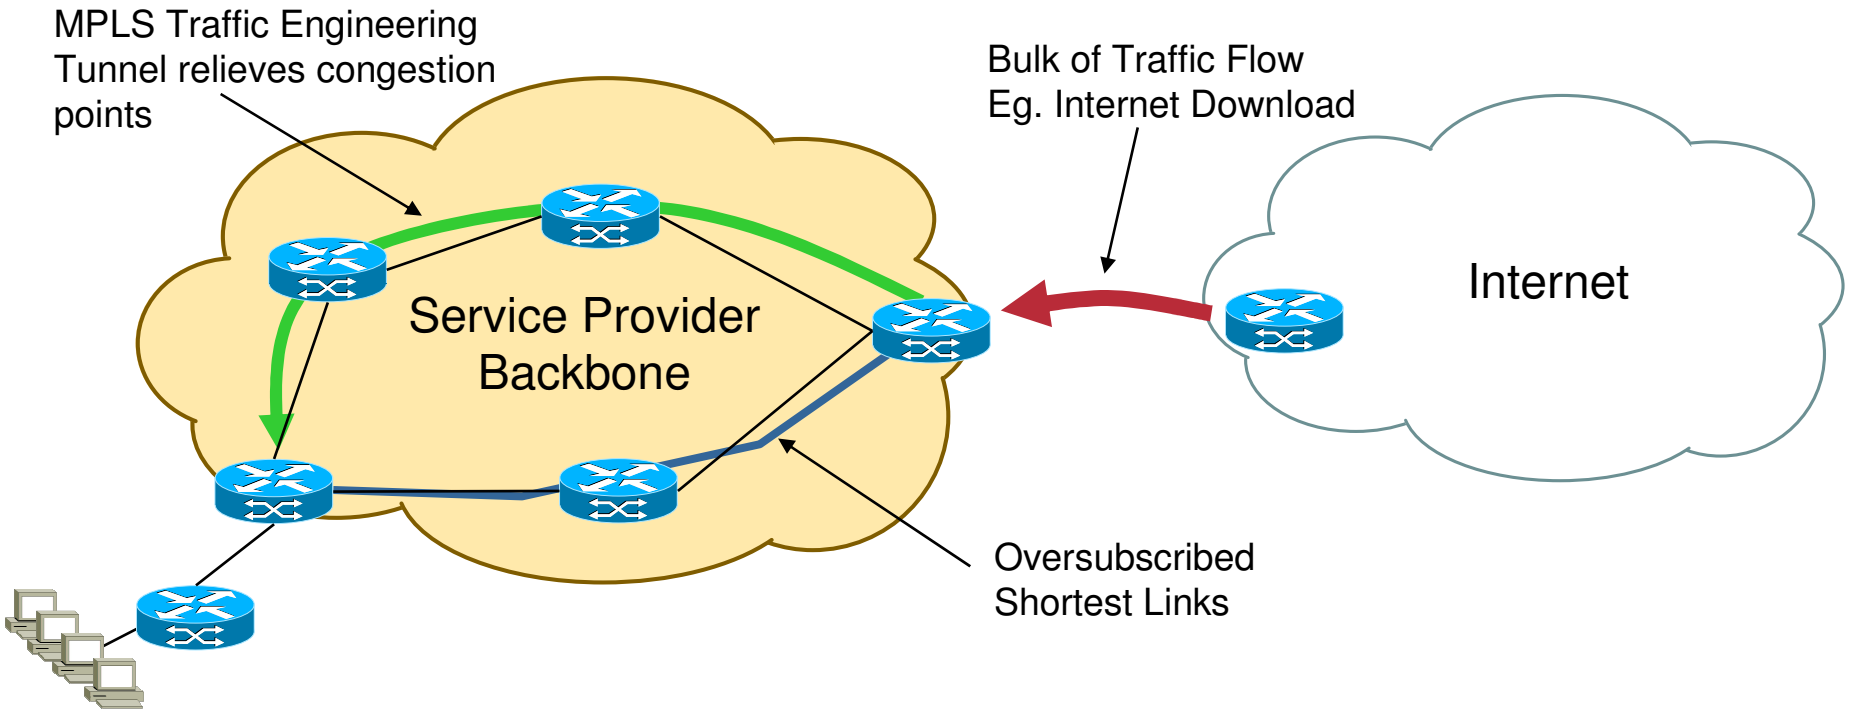

1-Hop TE Deployment

Cisco.com

Requirement: Need protection only – minimize packet loss. Lots of Bandwidth in the core

Solution: Deploy MPLS Fast Reroute for less than 50ms failover time with 1-Hop Primary TE Tunnels and Backup Tunnel for each

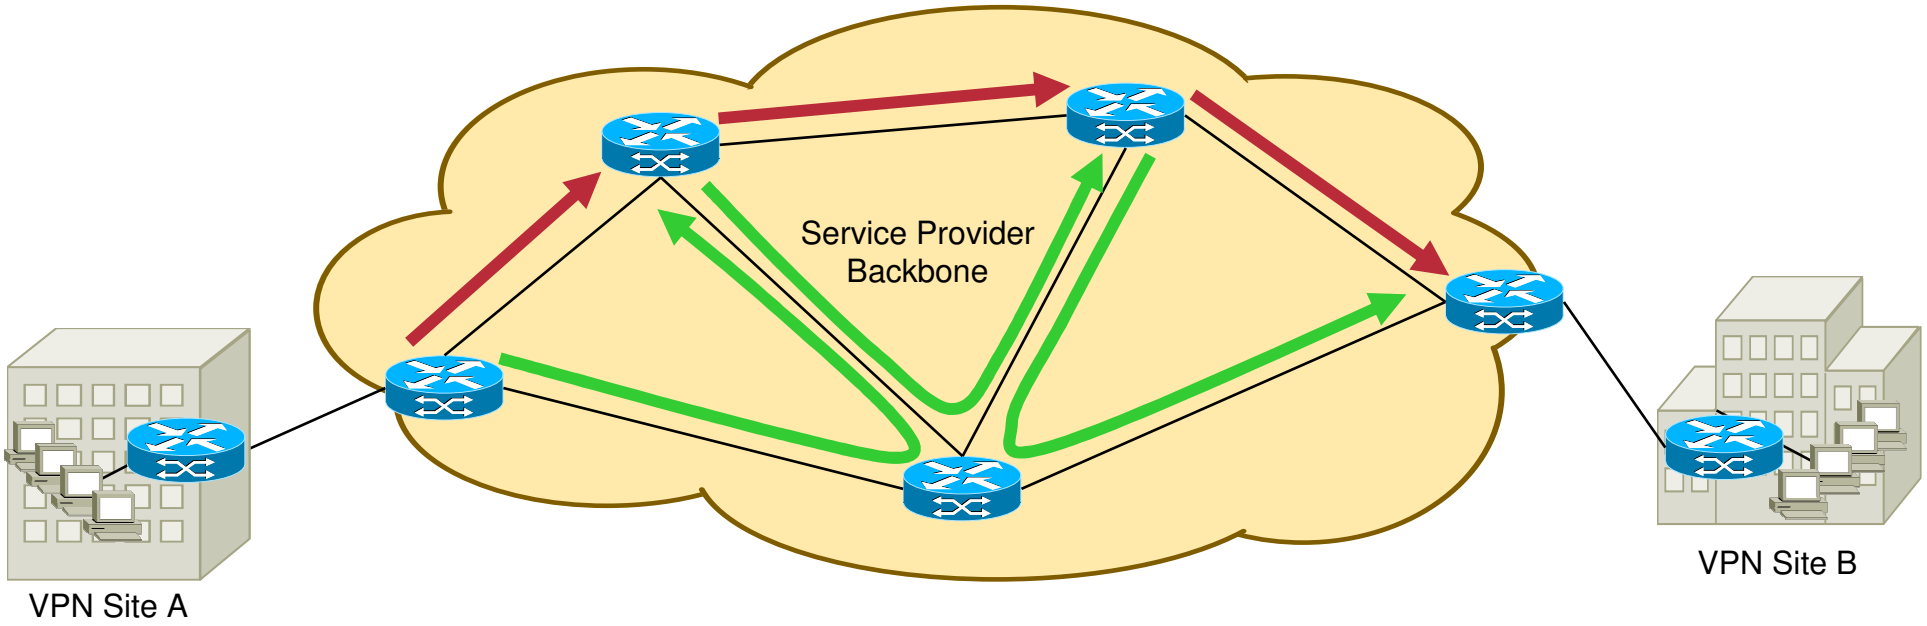

Virtual Leased Line Deployment

Cisco.com

Requirement: Need to create dedicated point-to-point circuits with bandwidth guarantees – Virtual Leased Line (VLL)

Solution: Deploy MPLS TE (or DS-TE) with QoS. Forward traffic from L3 VPN or L2 VPN into a TE Tunnel. Unlike ATM PVCs, use 1 TE Tunnel for multiple VPNs creating a scalable architecture