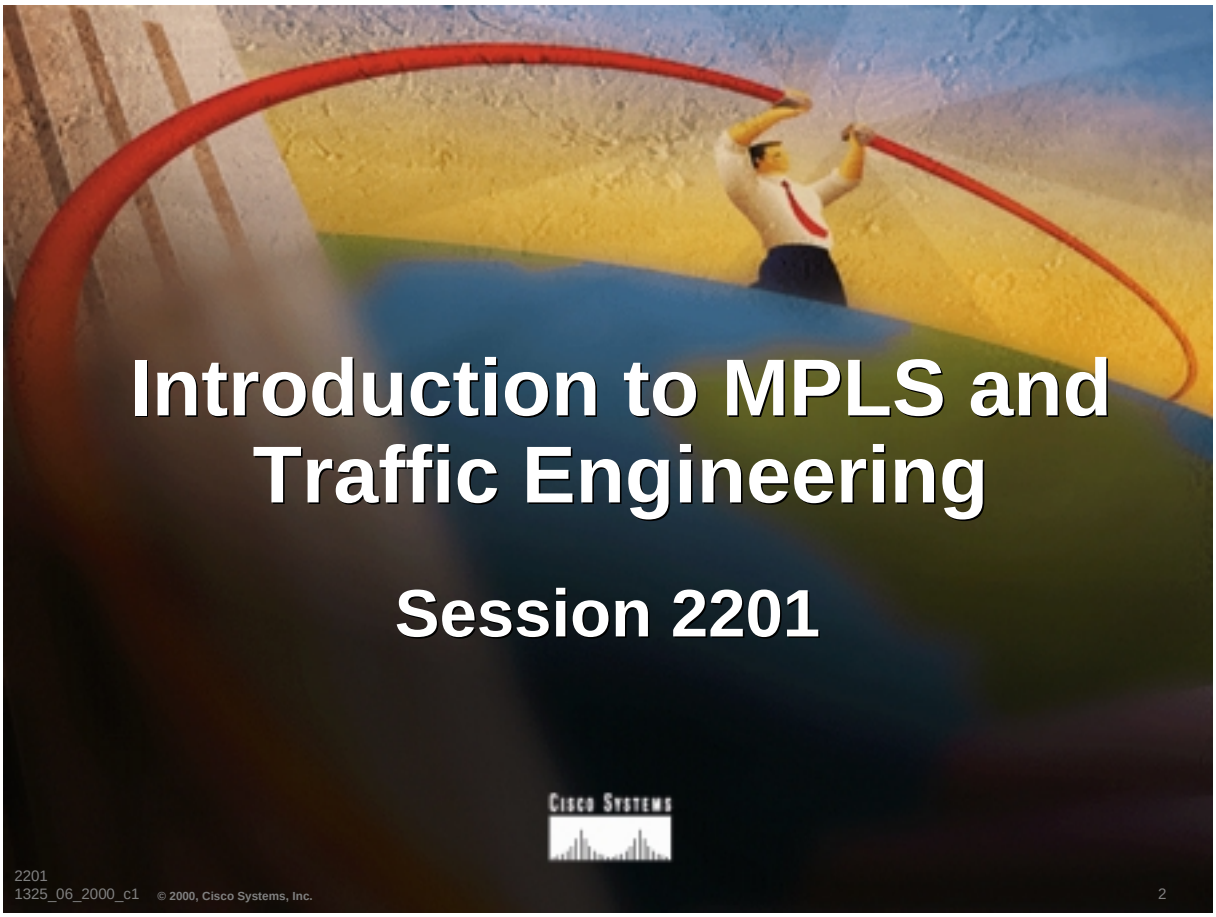

Topics

- **Motivations for MPLS**
- **MPLS Overview**
- **Applications**
- **Roadmap**

Why MPLS?

- **Integrate best of Layer 2 and Layer 3**
  - **Keep up with growth**
  - **Reduce operations costs**
  - **Increase reliability**
  - **Create new revenue from advanced IP services**
  - **Standards based**

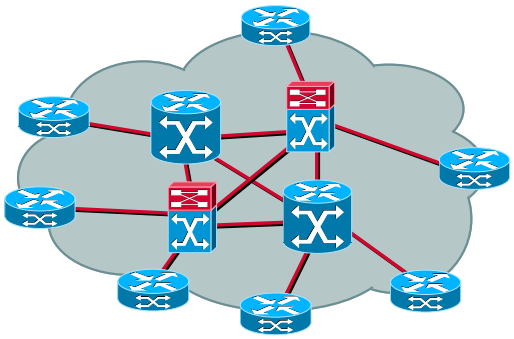

Key Cisco MPLS Solutions

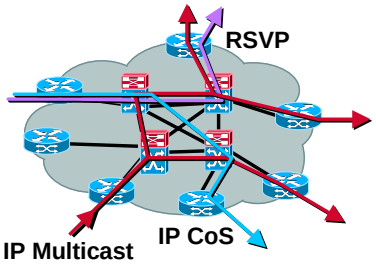

IP/ATM Integration

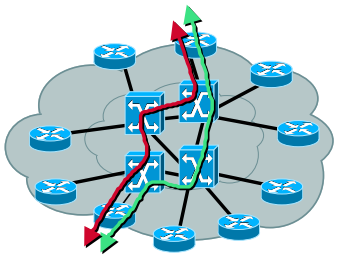

Traffic Engineering

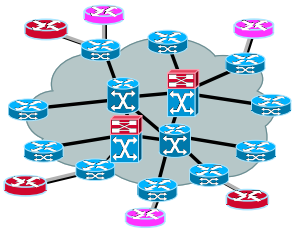

Internet Scale VPN/CoS

MPLS: Routing Scalability for IP over ATM

- **Internal routing scalability**
  - **Limited adjacencies**
- **External routing scalability**
  - **Full BGP4 support, with all the extras**
- **VC merge for very large networks**

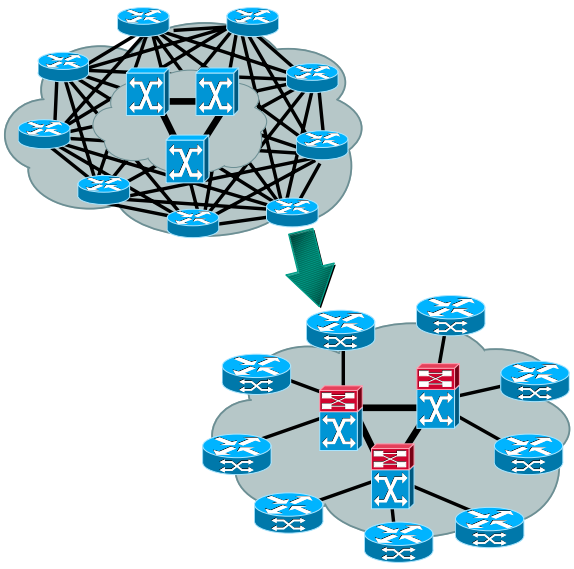

MPLS: End-to-End IP Services over ATM

- IP services directly on ATM switches
  - ATM switches support IP protocols directly
  - Avoids complex translation
- Full support for IP CoS, RSVP, IP multicast, future IP services

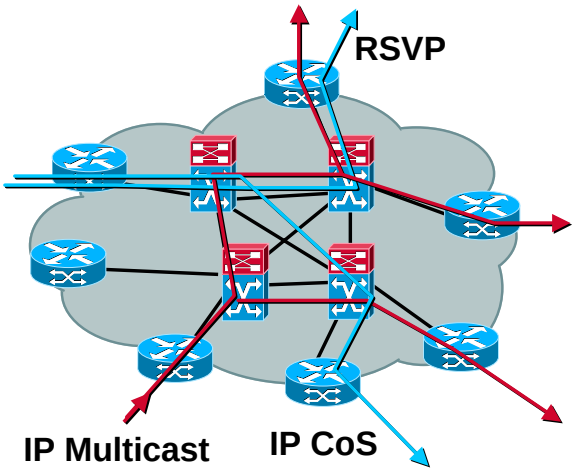

Benefits of MPLS Class of Service with ATM

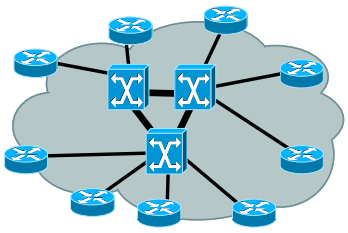

IP CoS over Standard ATM

- Allocate resources:
  - Per individual, edge-to-edge VCs
  - By kbps bandwidth
- Mesh of VCs to configure

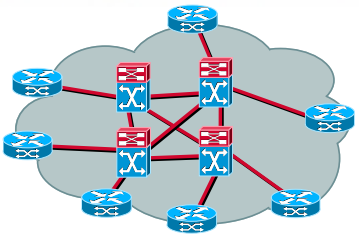

IP CoS with MPLS

- Allocate resources:
  - Per class, per link
  - By % bandwidth
- No VCs to configure
- Simpler to provision and engineer
- Even simpler with ABR

MPLS: Traffic Engineering

- **Characteristics**
  - High performance
  - Low overhead
  - End-to-end connectivity
- **Applications**
  - Constraint-based routing
  - Fast reroute
  - Guaranteed bandwidth
  - Frame/ATM transport
  - Control plane for ATM and OXCs

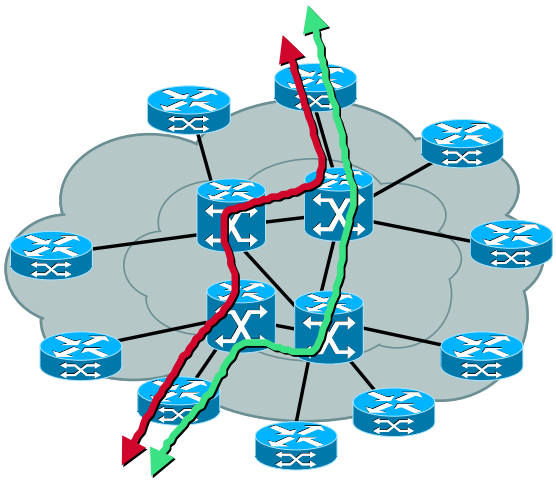

Motivations for Traffic Engineering

- **Links not available**
- **Economics**
- **Failure scenarios**
- **Unanticipated traffic**

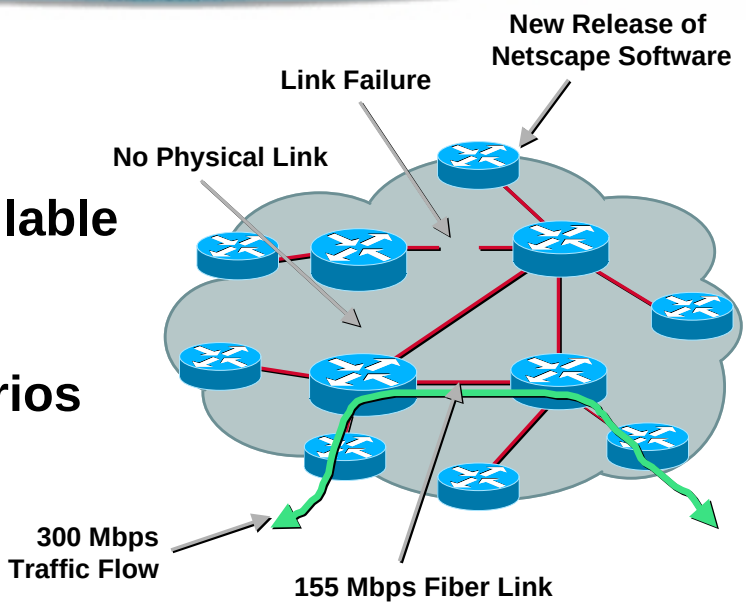

MPLS: Bringing Layer 2 Benefits to Layer 3

- **Traffic engineering**
  - Aligning traffic flows to resources
  - Optimize link utilization
- **Fast re-route**
  - Fast, local, link and node protection
- **Guaranteed bandwidth**
  - Hard end-to-end bandwidth and delay guarantees

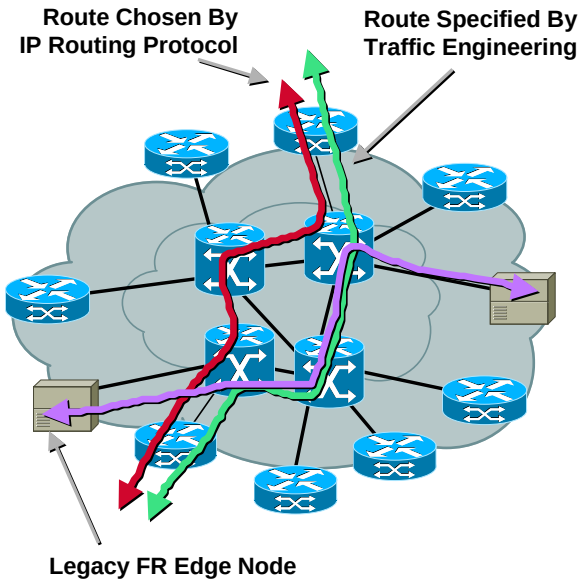

IP VPN Taxonomy

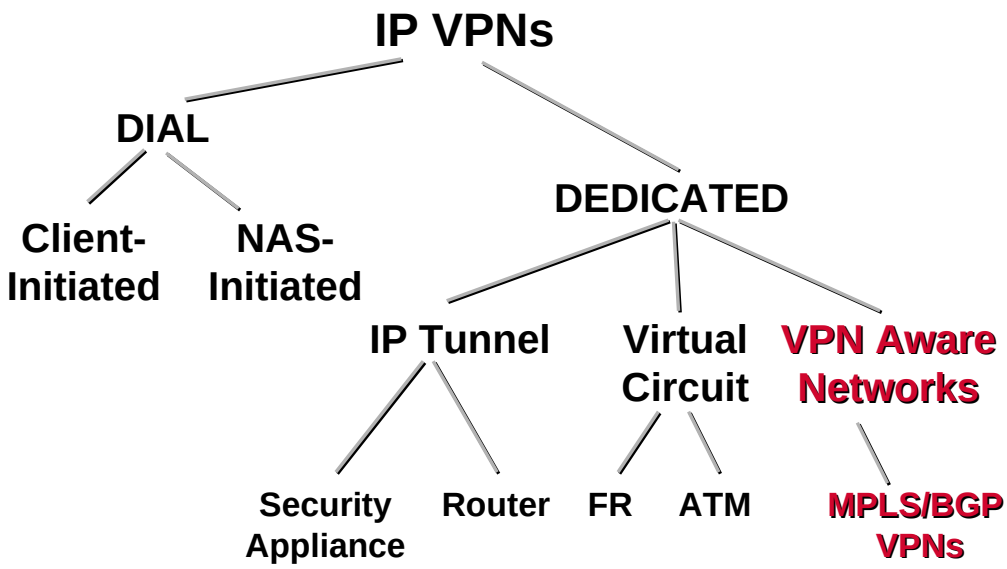

Cisco MPLS/BGP VPNs

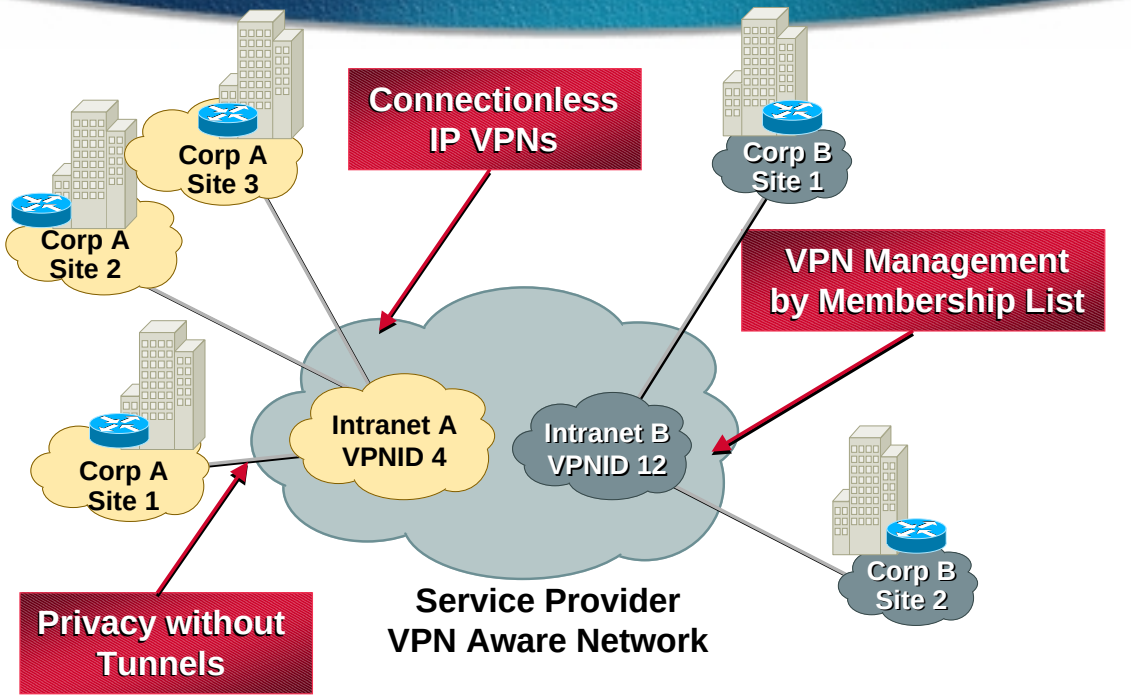

Benefits of MPLS/BGP VPNs

- Private, connectionless IP VPNs
- Outstanding scalability
- Customer IP addressing freedom
- Multiple QoS classes
- Secure support for intranets and extranets
- Simplified VPN Provisioning
- Support over any access or backbone technology

Connection-Oriented VPN Topology

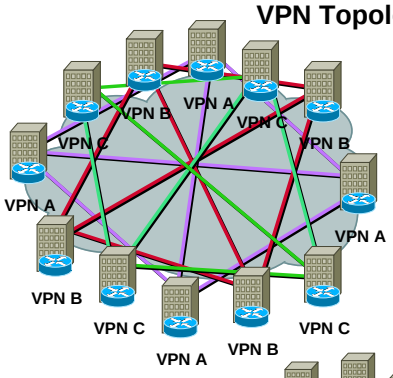

Connectionless VPN Topology

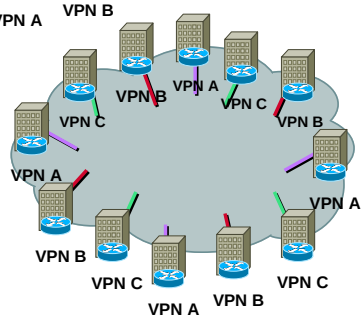

MPLS Benefits

	Benefits of MPLS
IP/ATM Integration	Shared Backbone for Economies of Scale Reduced Complexity for Lower Operational Cost Faster Time to Market for IP Services => More Revenue Use Best Technology => Lower Costs
Traffic Engineering	Traffic Eng. for Lower Trunk Costs and Higher Reliability Fast Reroute for Protection and Resiliency Guaranteed Bandwidth for Hard QoS Guarantees
MPLS BGP VPNs	New Revenue Opportunity for SPs Scalability for Lower Operational Costs and Faster Rollout L2 Privacy and Performance for IP

Topics

- Motivations for MPLS
- MPLS Overview
- Applications
- Roadmap

MPLS Concept

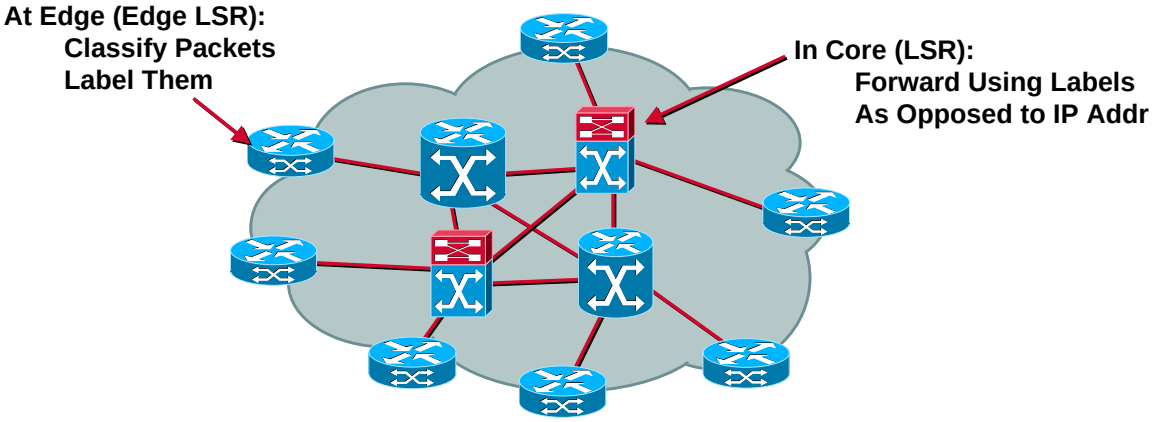

- Enable ATM switches to act as routers
- Create new IP capabilities via flexible classification

Router Example: Distributing Routing Information

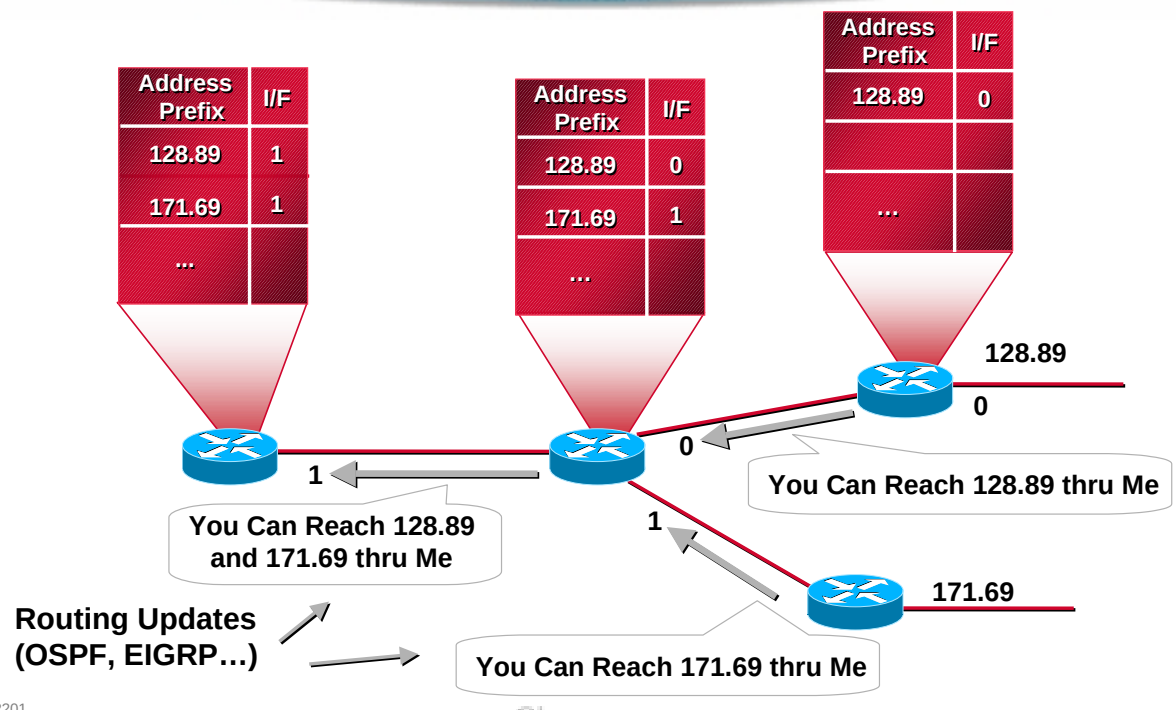

Router Example: Forwarding Packets

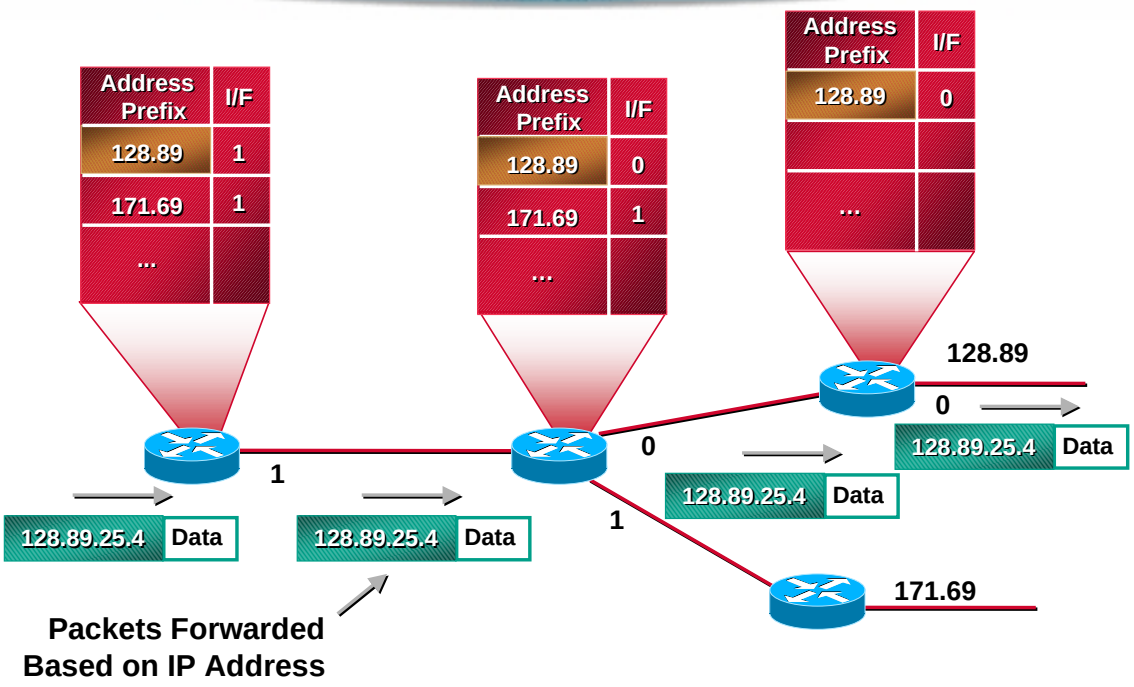

MPLS Example: Routing Information

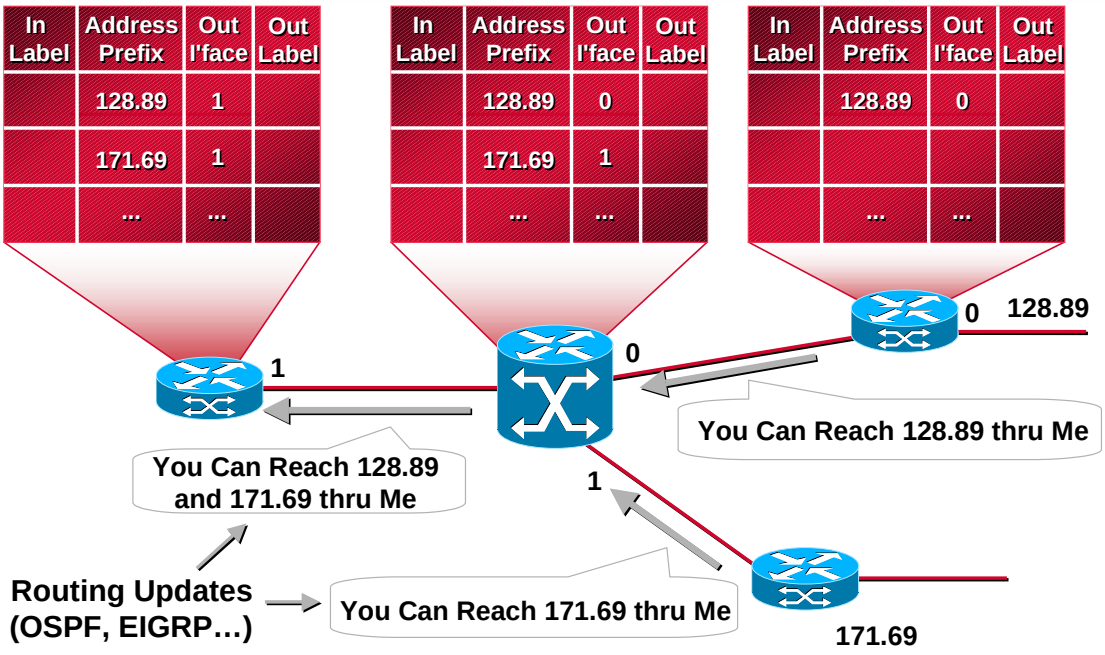

MPLS Example: Assigning Labels

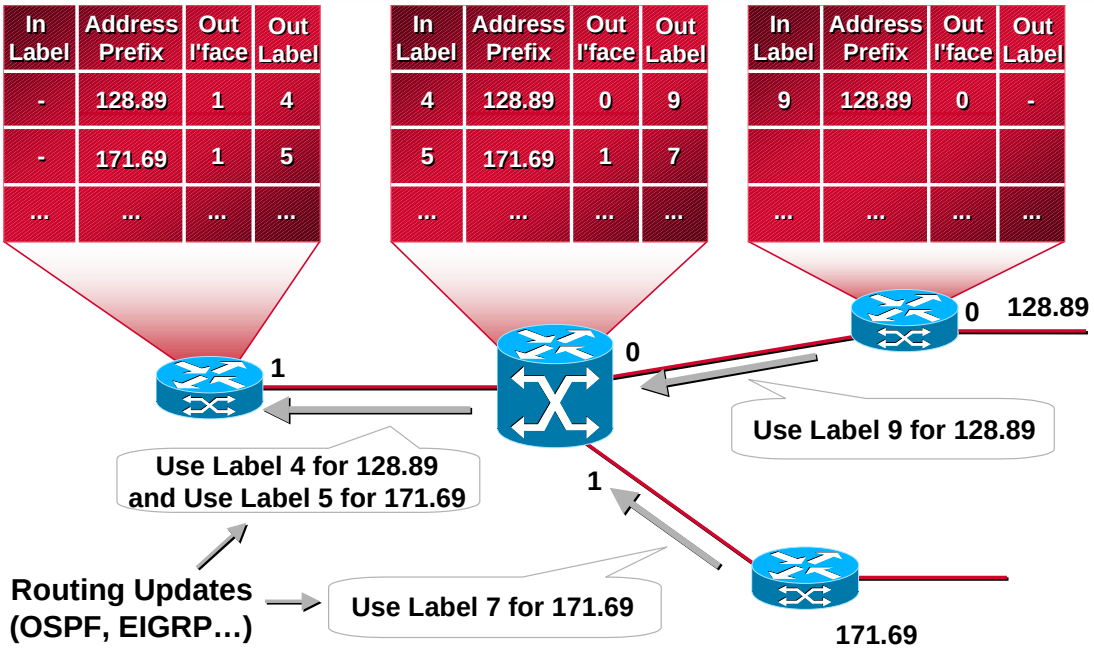

MPLS Example: Forwarding Packets

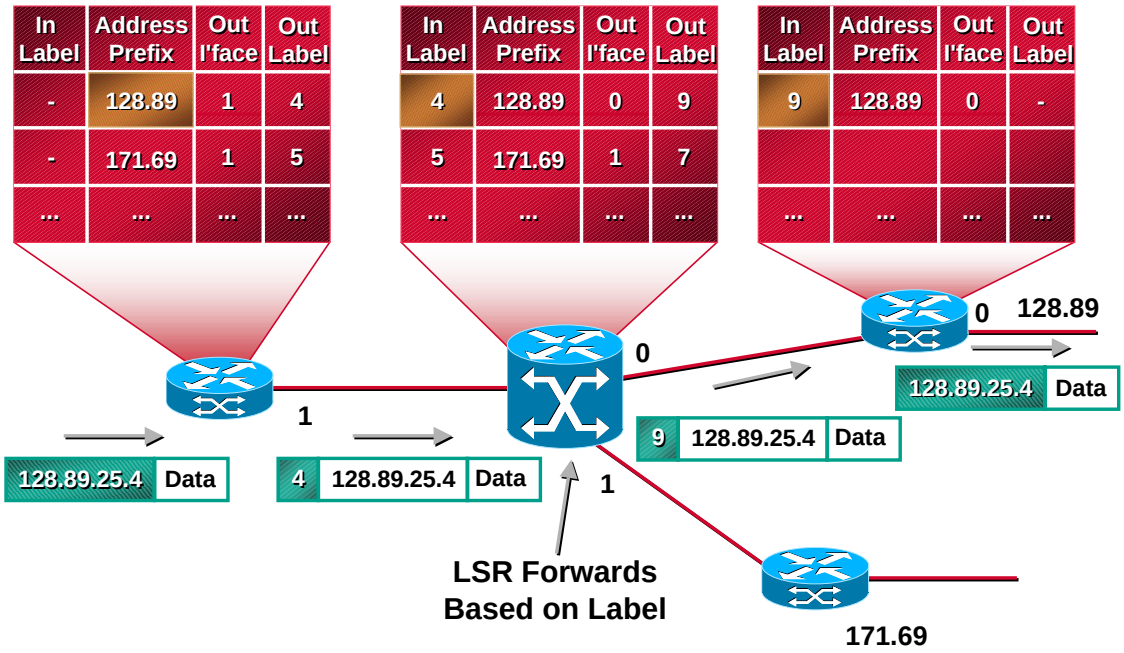

MPLS Example: More Details

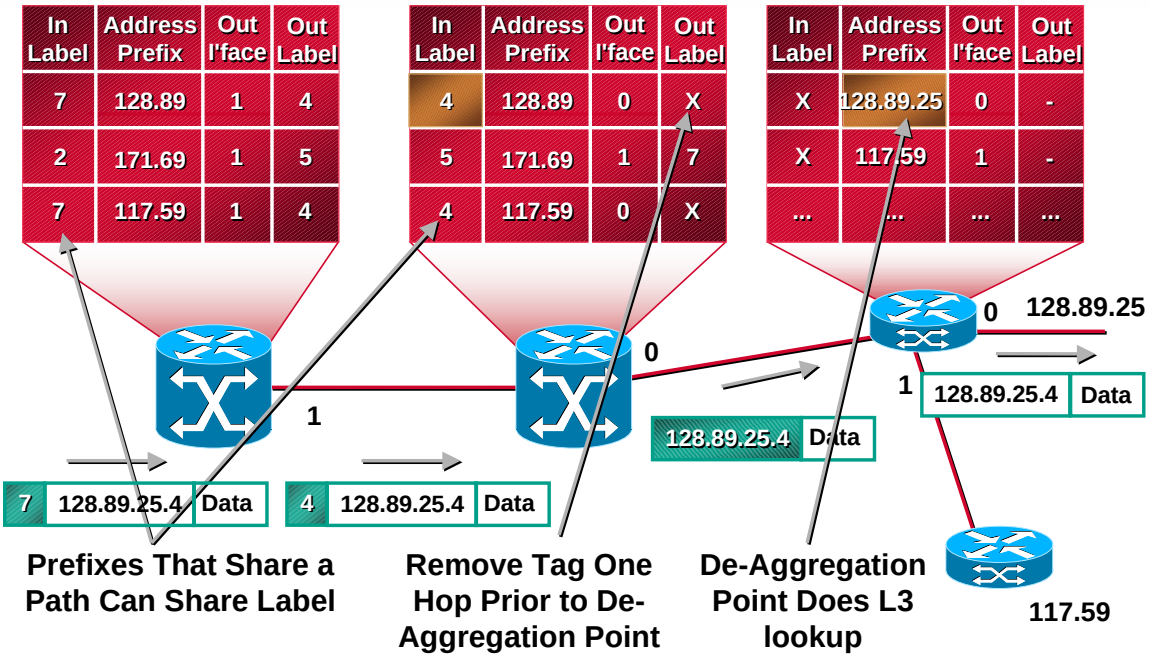

Encapsulations

ATM Cell Header

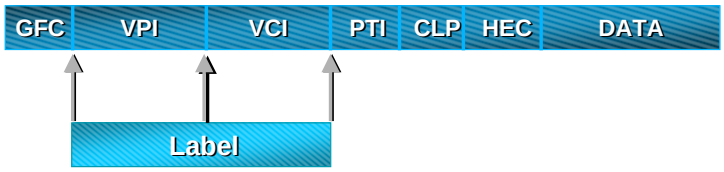

PPP Header (Packet over SONET/SDH)

LAN MAC Label Header

Label Header for Packet Media

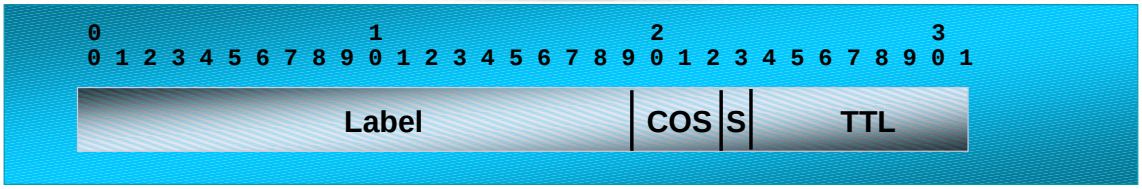

Label = 20 bits
COS = Class of Service, 3 Bits
S = Bottom of Stack, 1 Bit
TTL = Time to Live, 8 Bits

- Can be used over Ethernet, 802.3, or PPP links
- Uses two new ether types/PPP PIDs
- Contains everything needed at forwarding time
- One word per label

ATM MPLS Example: Routing Information

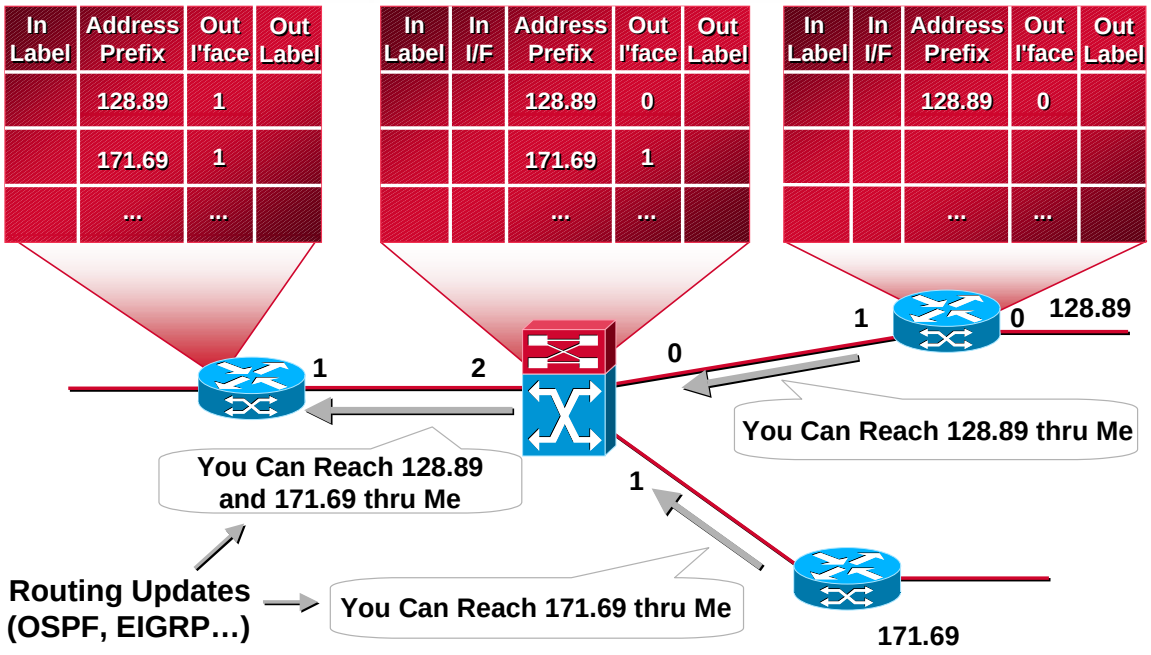

ATM MPLS Example: Requesting Labels

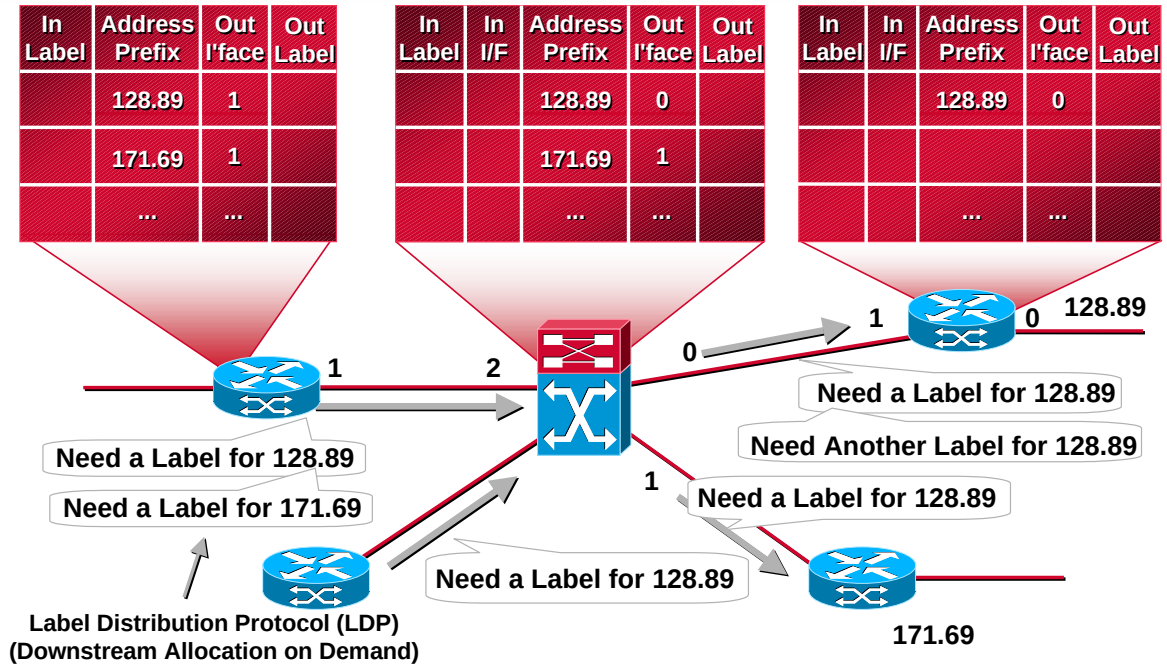

ATM MPLS Example: Assigning Labels

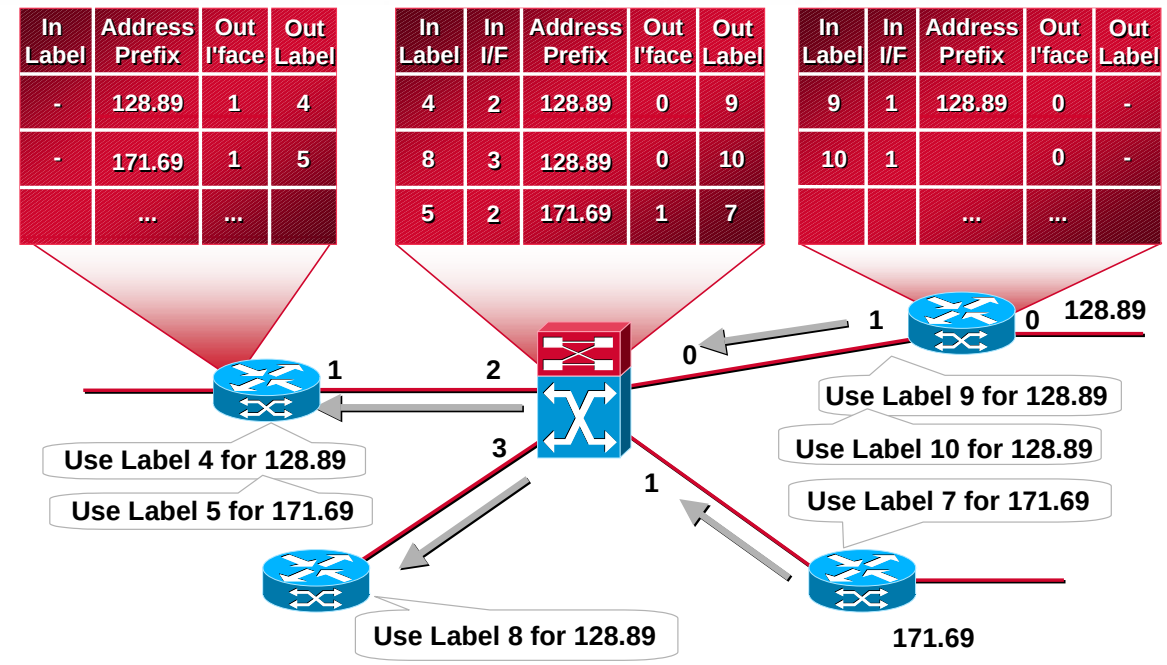

ATM MPLS Example: Packet Forwarding

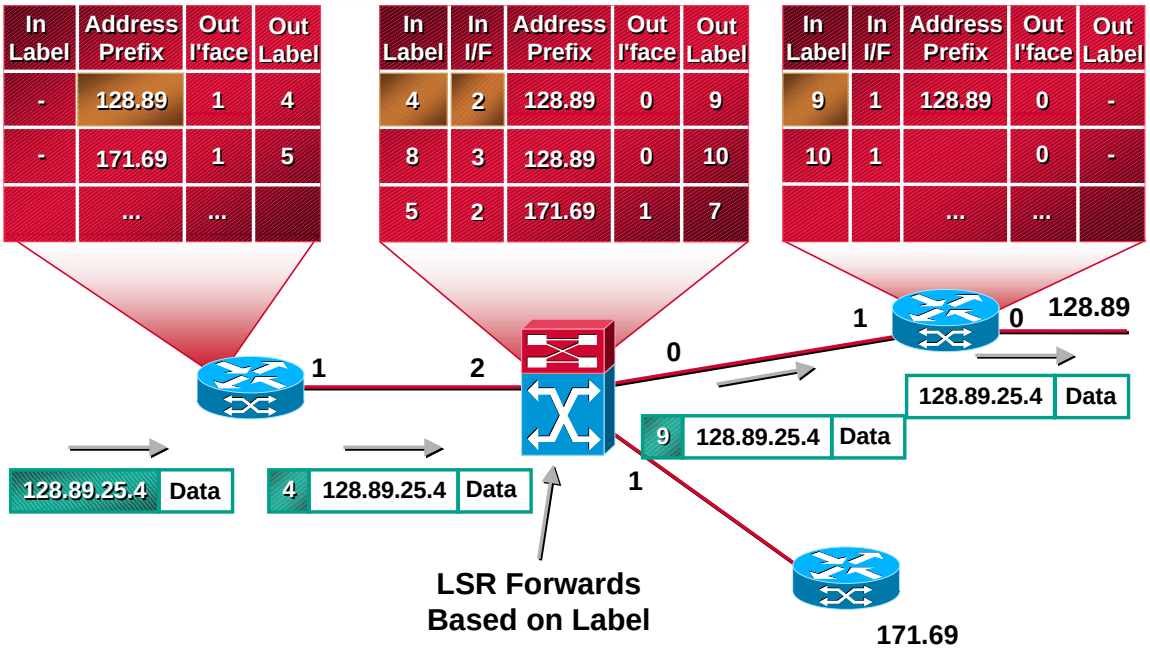

Why Multiple Labels with ATM?

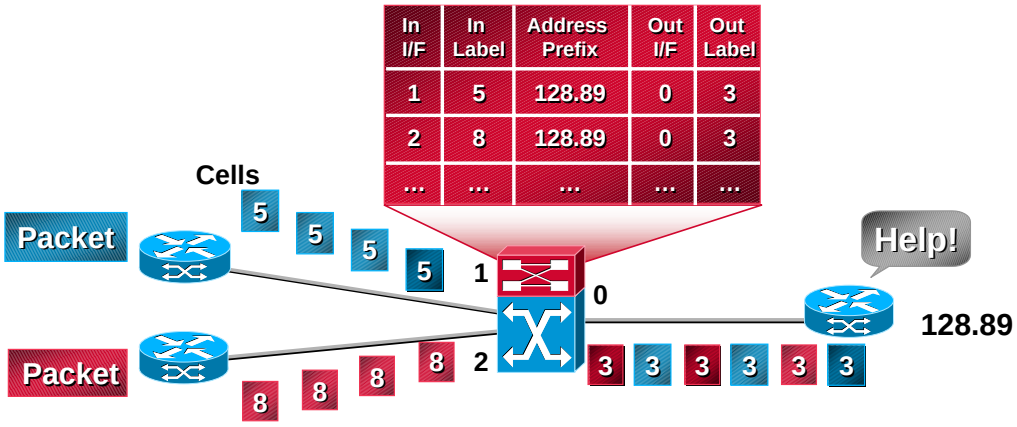

- **If didn't allocate multiple labels**
  - **Cells of different packets would have same label (VPI/VCI)**
  - **Egress router can't reassemble packets**

Multiple Labels

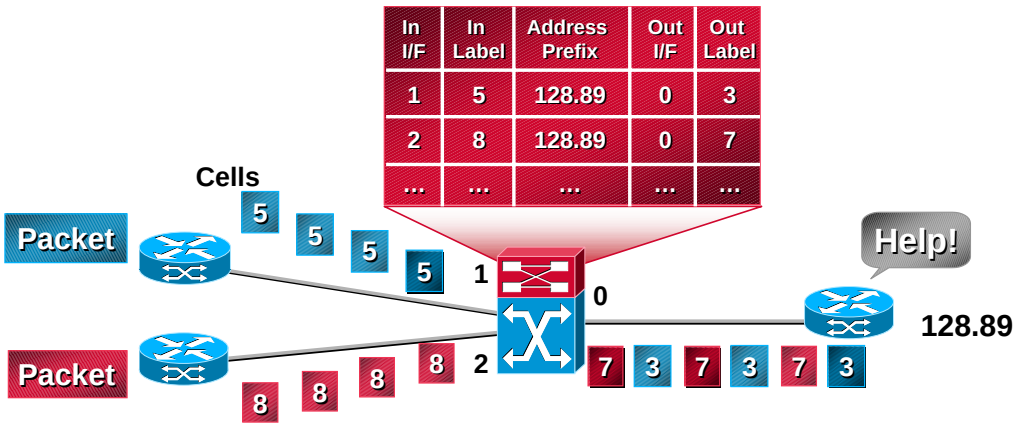

- **Multiple labels enable edge router to reassemble packets correctly**

VC Merge

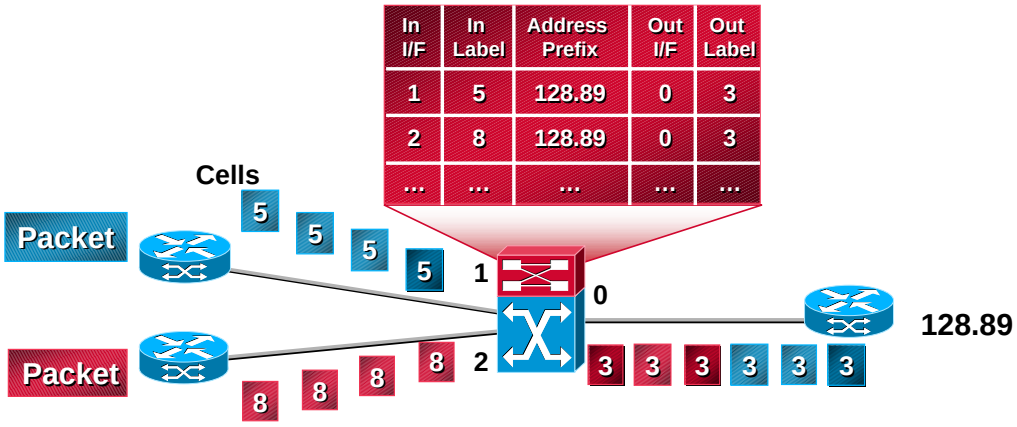

- **With ATM switch that can merge VCs**
  - **Can reuse outgoing label**
  - **Hardware prevents cell interleave**
  - **Fewer labels required**
  - **For very large networks**

Advanced MPLS

- Basic MPLS: destination-based unicast
- Many additional options for assigning labels
- The key: separation of routing and forwarding

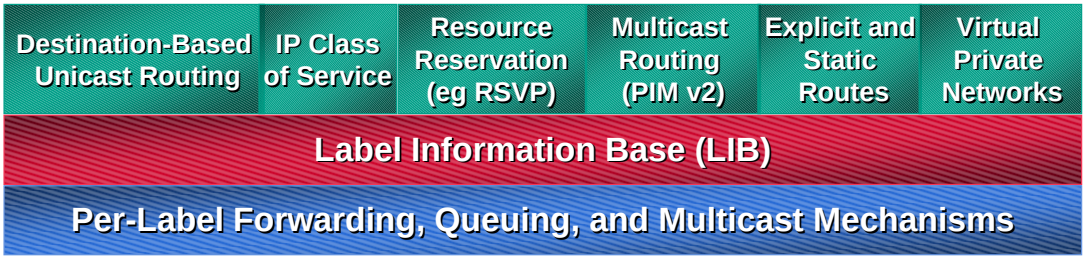

Building VPNs with MPLS

- Constrained distribution of routing information
  - Routes are only communicated to routers that are members of a VPN
- VPN-IP addresses
  - Supports overlapping address spaces
- Multiprotocol Label Switching (MPLS)
  - Labels used to define VPNs
  - Labels used to represent VPN-IP addresses
- Peer model
  - Simplifies routing for end customers

MPLS VPN Example

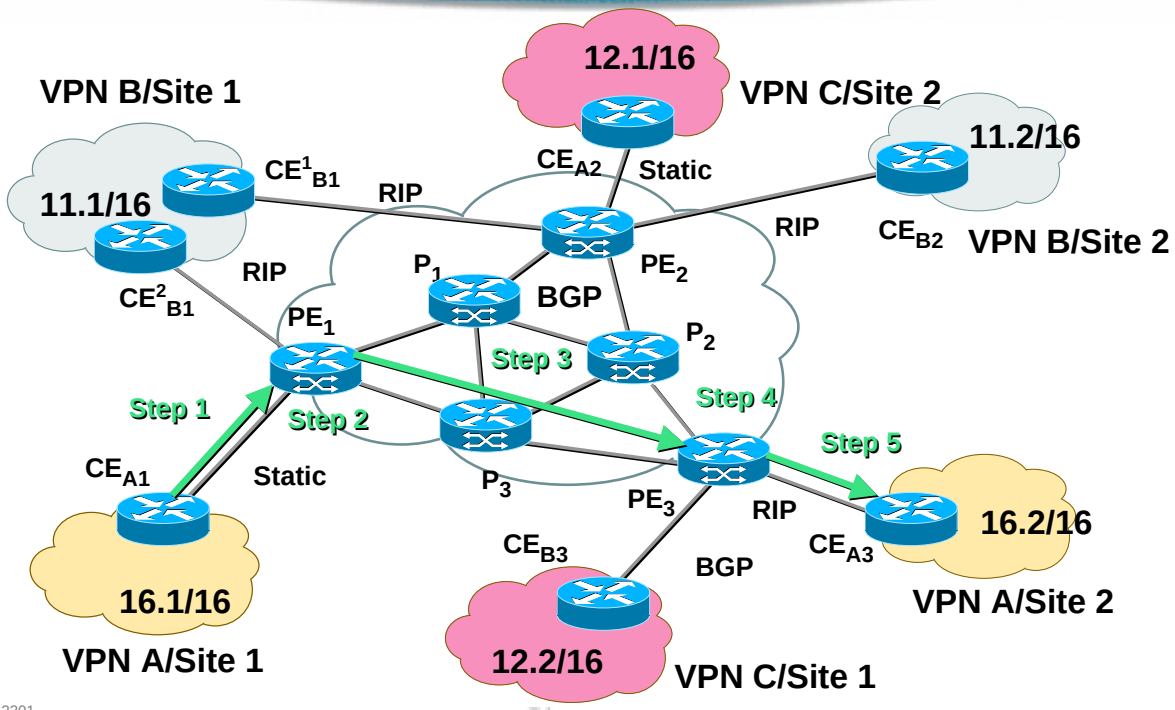

Routing Information Distribution

Step 1: From site (CE) to service provider (PE)

E.g. via RIP, OSPF, static routing, or BGP

Step 2: Export to provider's BGP at ingress PE

Step 3: Within/across service provider(s) (among PEs):

E.g. via BGP

Step 4: Import from provider's BGP at egress PE

Step 5: From service provider (PE) to site (CE)

E.g. via RIP, or static routing, or BGP

Packet Forwarding

- IP packet received on sub-interface
- Sub-interfaced configured with VPN ID
- BGP binds labels to VPN-IP routes
- LDP binds labels to IGP routes and defines CoS
- Logically separate forwarding information base (FIB) for each VPN

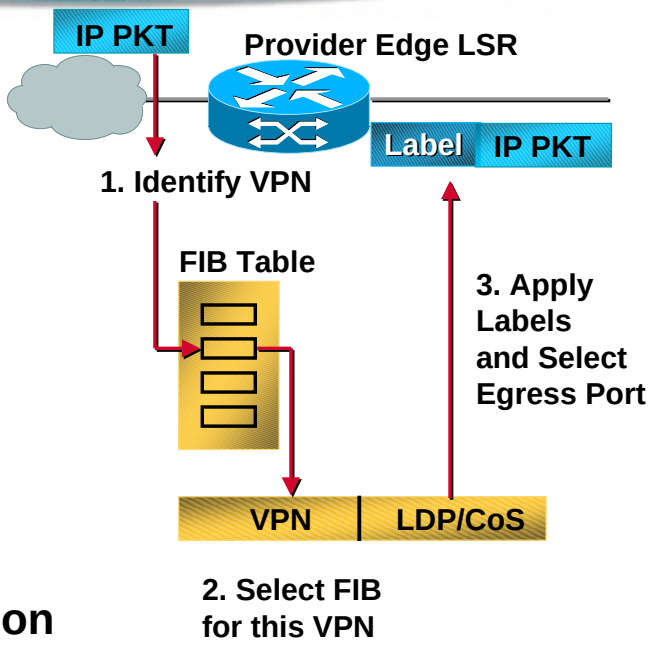

MPLS VPN Example

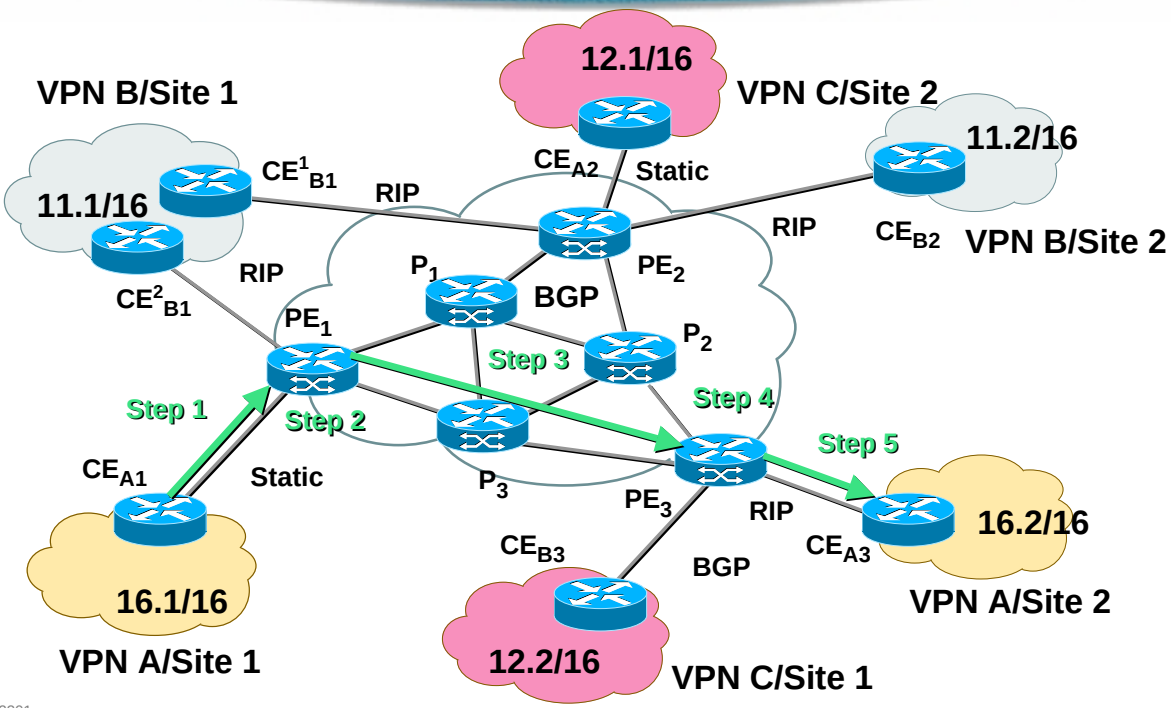

Explicit Routing

- Traffic engineering requires the capability to specify a path
- Voice networks, Frame Relay, ATM are explicitly routed at connection setup
- But IP uses hop-by-hop destination-based routing

The "Fish" Problem

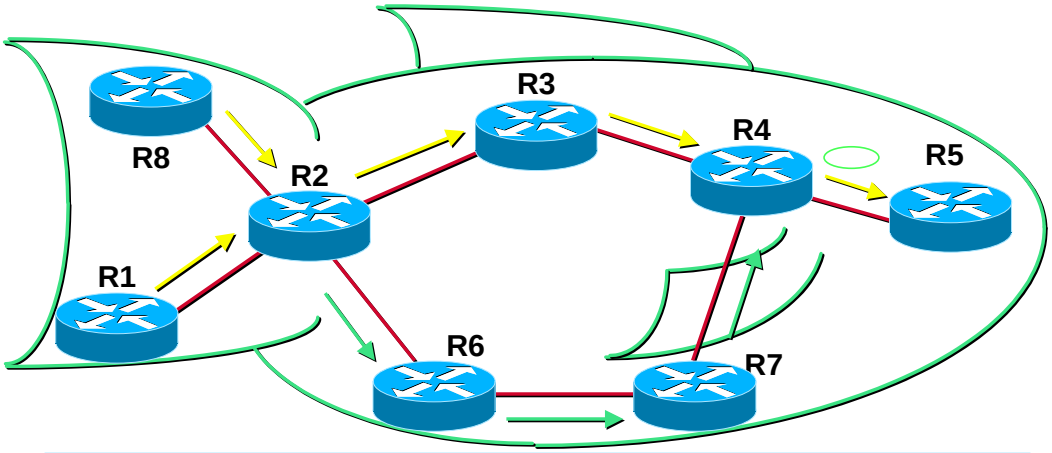

IP Uses Shortest Path Destination-Based Routing

Shortest Path May Not Be the only path

Alternate Paths May Be under-Utilized while the Shortest Path Is over-Utilized

An LSP Tunnel

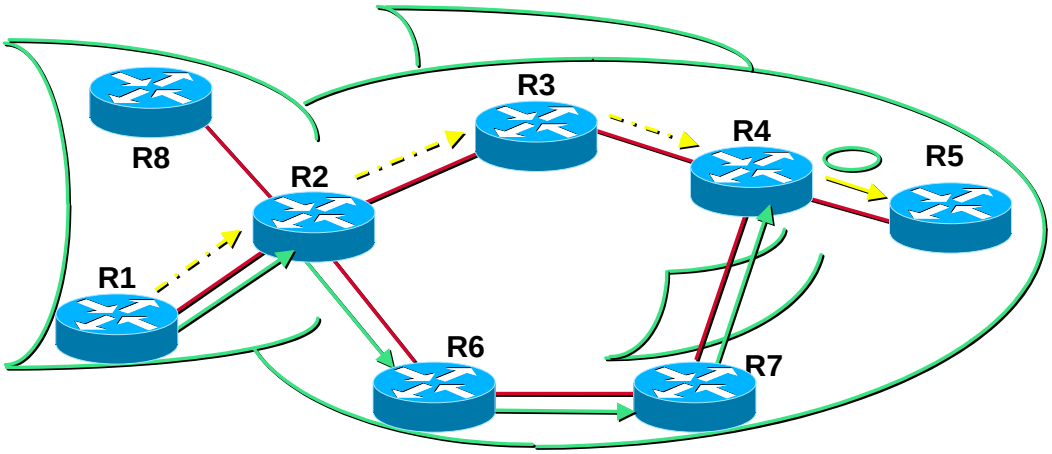

Labels, Like VCIs Can Be Used to Establish Virtual Circuits

Normal Route R1->R2->R3->R4->R5

Tunnel: R1->R2->R6->R7->R4

Traffic Engineering

- **Provides**
  - **Constraint-based routing**
    - **Similar to PNNI routing**
  - **Control of traffic engineering**
  - **Path selection**
  - **Tunnel setup**

Basic Traffic Engineering

- LSP tunnels used to steer traffic

  (Termed traffic engineering or TE tunnels)
- Represent inter-POP traffic as flows in bits/sec
- Determine bandwidth requirements for tunnels between POP pairs
- Automated procedures route and setup the inter-POP TE tunnels

TE Components

(1) Information distribution

Distributes constraints pertaining to links

Available bandwidth is just one type of constraint

(2) Path selection algorithm

Selects paths that obey the constraints

TE Components (Cont.)

(3) Route setup

Uses RSVP for signaling LSPs

(4) Link admission control

Decides which tunnels may have resources

(5) Traffic engineering control

Establishes and maintains tunnels

(6) Forwarding data

System Block Diagram

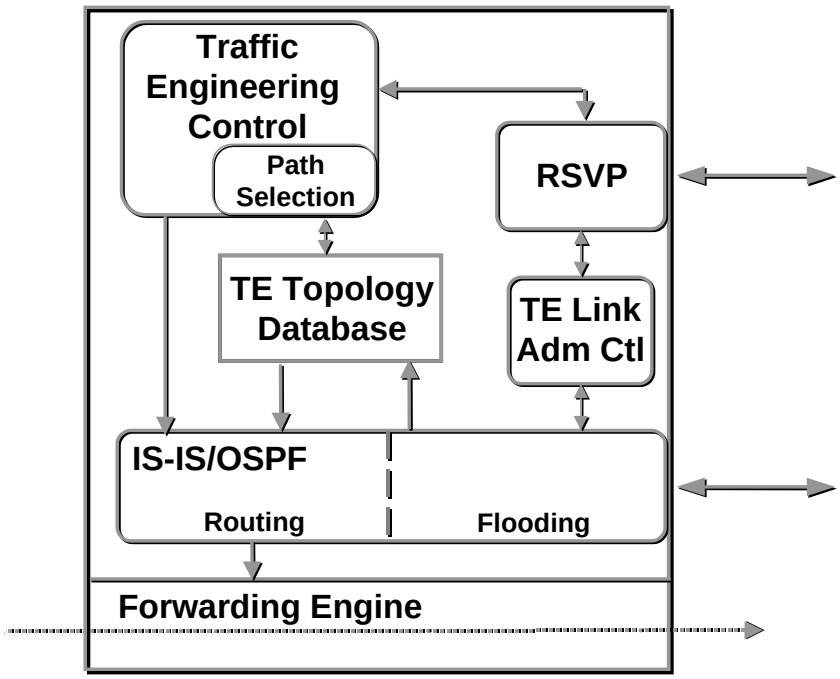

LSP Tunnel Setup

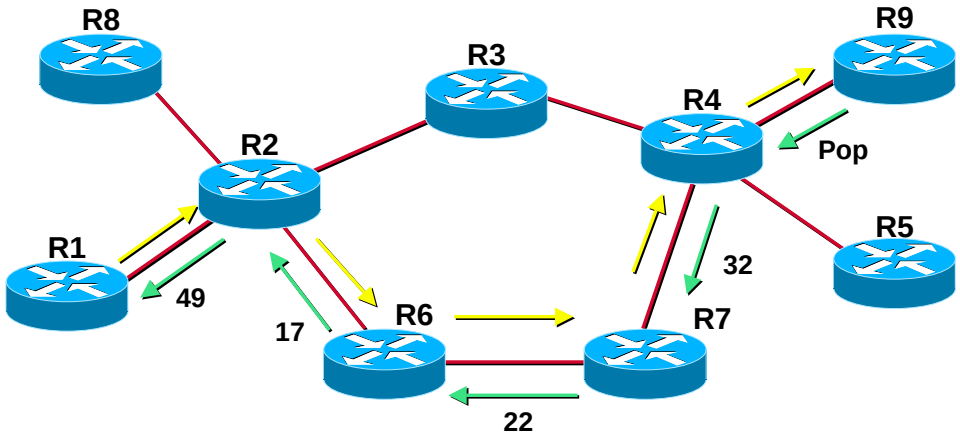

Setup: Path (R1->R2->R6->R7->R4->R9) Tunnel ID 5, Path ID 1

Reply: Communicates Labels and Label Operations
Reserves Bandwidth on Each Link

Rerouting to an Alternate Path

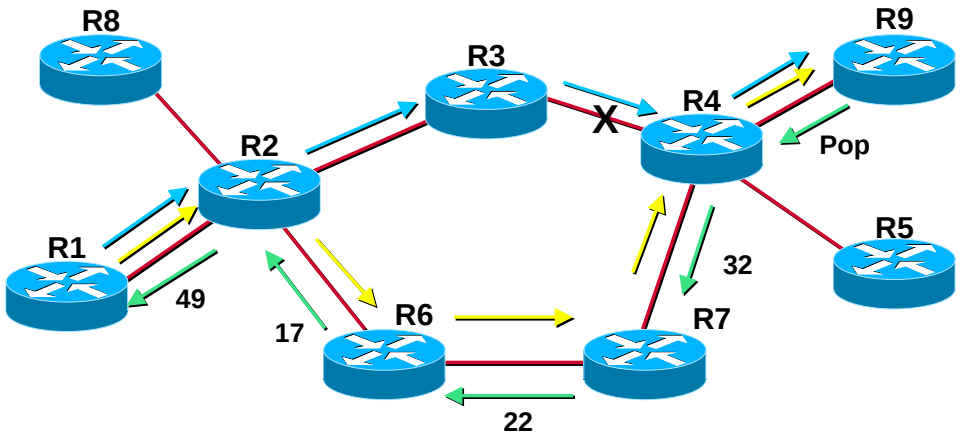

Setup: Path (R1->R2->R3->R4->R9) Tunnel ID 5, Path ID 2

Until R9 Gets New Path Message, Current Resv Is Refreshed

Bridge and Roll

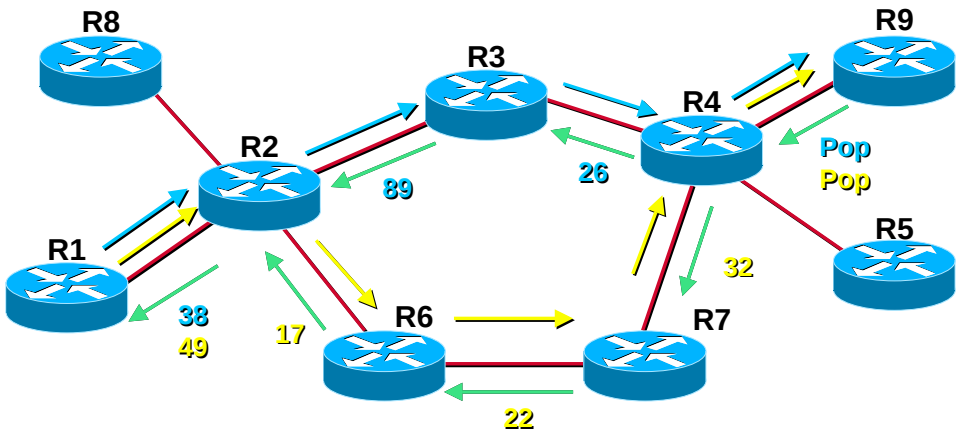

Resv: Allocates Labels for Both Paths
Reserves Bandwidth Once Per Link

PathTear Can then Be Sent to Remove Old Path and Release Resources

Assigning Traffic to Tunnels

- **Automatic assignment based on IGP**
- **Modified SPF calculation**

  When the endpoint of a tunnel is reached, the next hop to that node is set to the tunnel interface

  Nodes downstream of the tunnel inherit the tunnel interface as their next hop

  (Encountering a node with its own tunnel replaces the next hop)

Topology with Tunnel

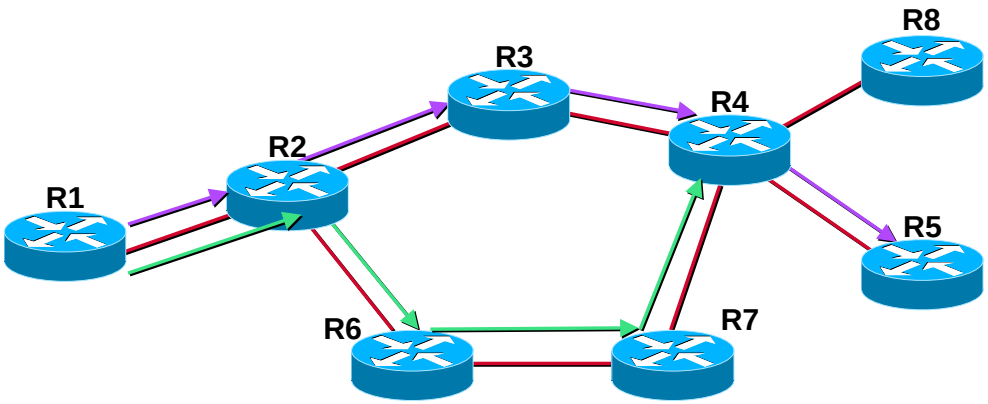

Tunnel1: Path (R1->R2->R6->R7->R4)

Tunnel2: Path (R1->R2->R3->R4->R5)

Normal Dijkstra, Except Tunnel Interfaces Used when Tunnel Tail Is Encountered

Forwarding Tree

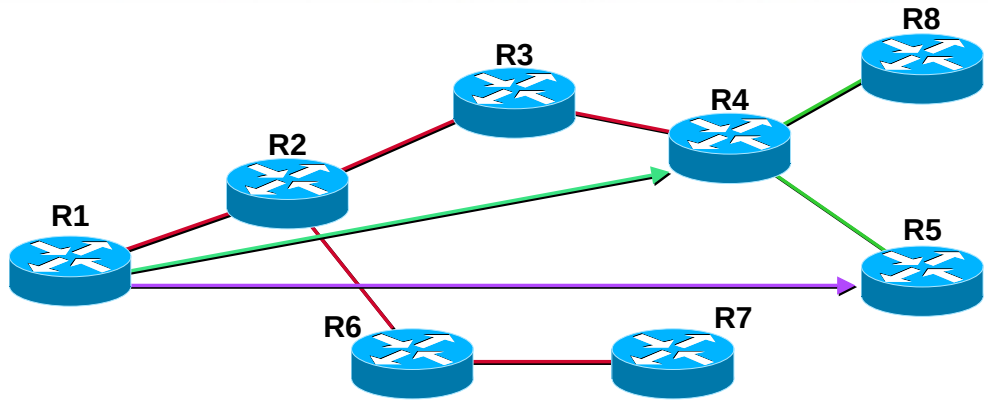

Tunnel1 Tunnel2

R4 and R8 Have Tunnel1 Interface as Next Hop; R5 Has Tunnel2

Fast Reroute

- Goal—match Sonet restoral times—50 ms
- Locally patch around lost facilities
- Strategies
  - Alternate tunnel (1->1 mapping)
  - Tunnel within tunnel (n->1 mapping)

Fast Reroute

- Labels are carried in a stack, making it possible to nest tunnels
- RSVP has a notion of PHOP, allowing the protocol to be independent of the back channel
- A tunnel can use another tunnel as a tunnel hop to enable fast reroute

Nested Tunnels—Outer

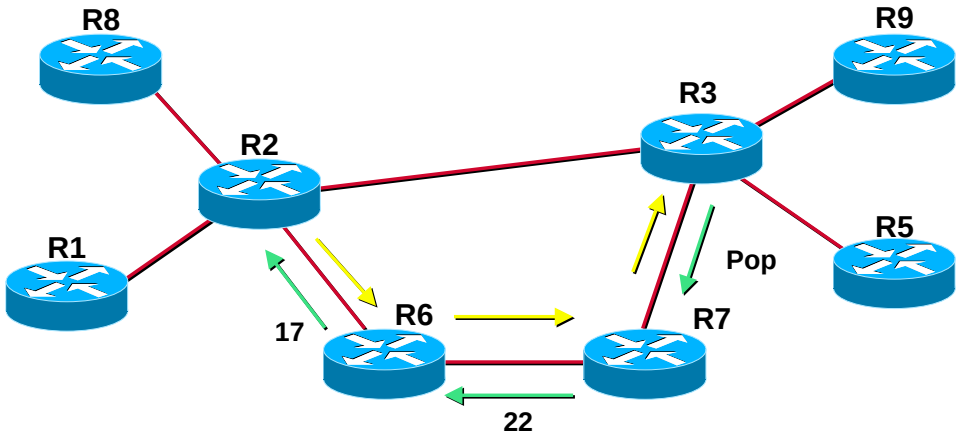

Setup: Path (R2->R3->R4) Session 5, ID 2

Labels Established on Resv Message

Nested Tunnels—Inner

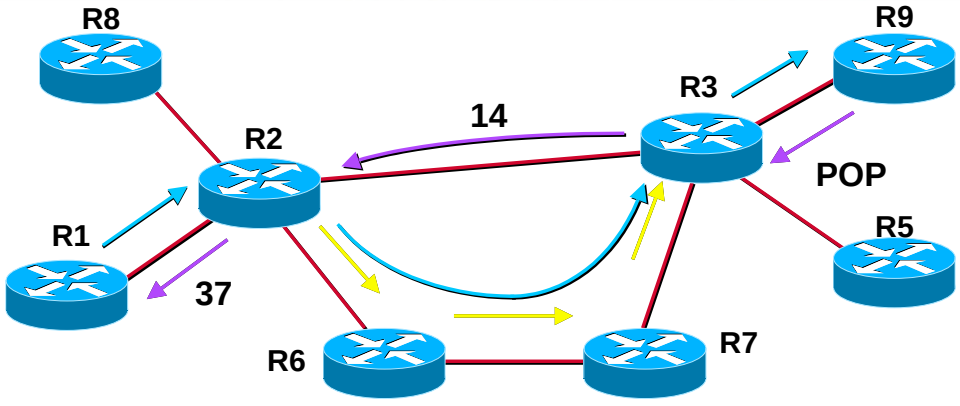

Setup: Path (R1->R2->R4->R9) Path Message Travels on Tunnel from R2 to R4

R4 Send Resv Message Directly to R2

Nested Tunnels—Operation

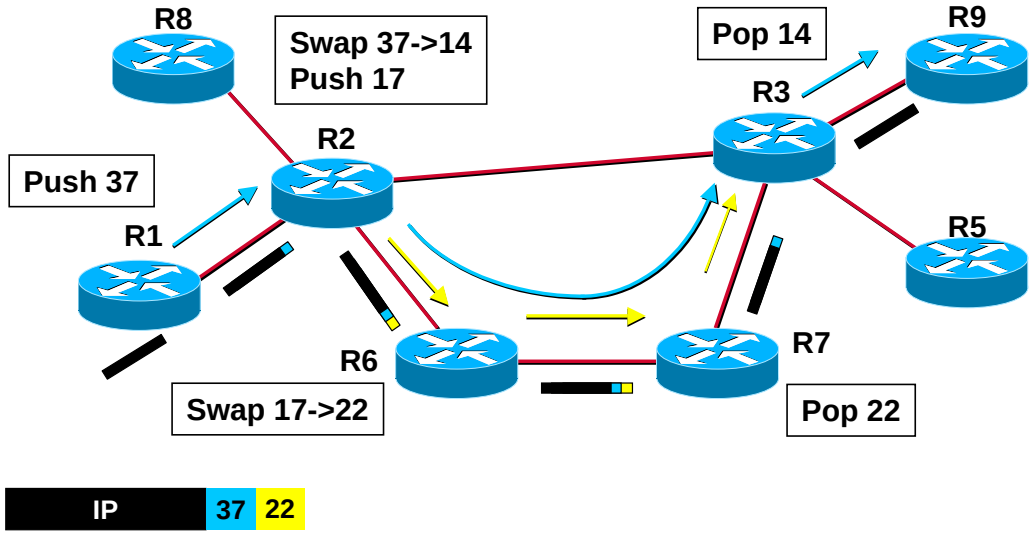

Nested Tunnels—Fast Reroute

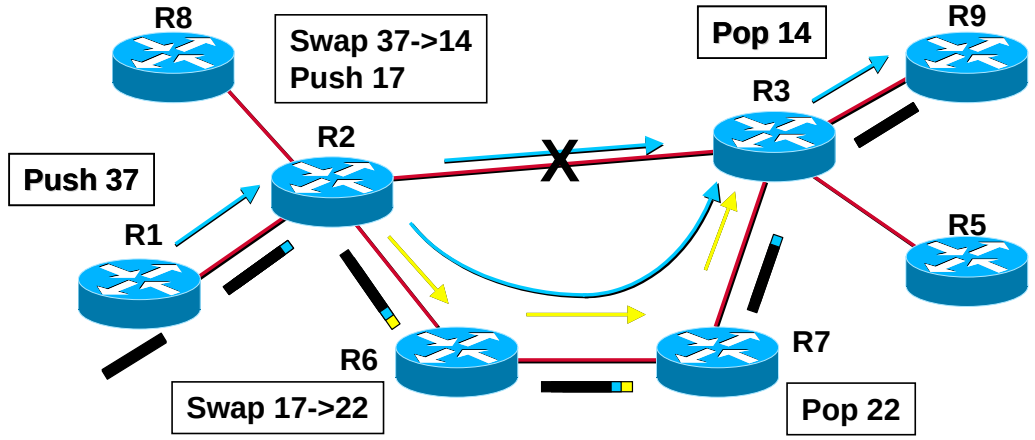

On Failure of Link from R2 -> R3, R2 Simply Changes the Outgoing Interface and Pushes on the Label for the Tunnel to R3

Conclusions: MPLS Fundamentals

- Based on the label-swapping forwarding paradigm
- As a packet enters an MPLS network, it is assigned a label based on its Forwarding Equivalence Class (FEC)
  - As determined at the edge of the MPLS network
- FECs are groups of packets forwarded over the same Label Switched Path (LSP)

Conclusions: MPLS Main Ideas

- Separate forwarding information (label) from the content of IP header
- Single forwarding paradigm (label swapping)—multiple routing paradigms
- Multiple link-specific realizations of the label swapping forwarding paradigm
- Flexibility of forming Forwarding Equivalence Classes (FECs)
- Forwarding hierarchy via label stacking

Topics

- Motivations for MPLS
- MPLS Overview
- Applications
- Roadmap

Application: Multiservice ATM Backbone with IP

- MPLS provides
  - Scalable IP routing
  - Advanced IP services
  - Internet scale VPNs
- Benefits
  - Lower operations costs
  - Keep up with Internet growth
  - New revenue services
  - Multiservice backbone
  - Faster time to market

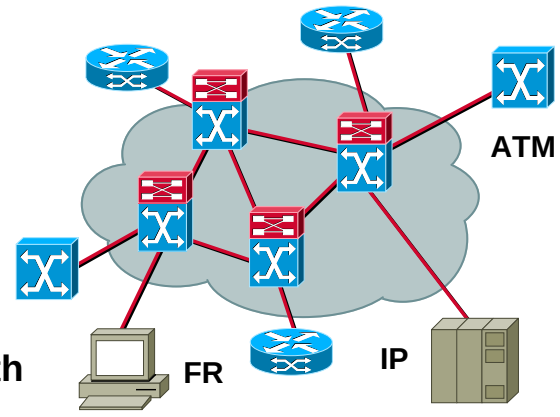

Application: Packet over SONET/SDH IP Backbone

- MPLS provides
  - Isolation of backbone from BGP
  - Traffic engineering
  - Guaranteed bandwidth
  - Internet scale VPNs
  - FR/ATM over MPLS
- Benefits
  - Improved line utilization
  - Increased reliability
  - Convergence
  - New revenue services

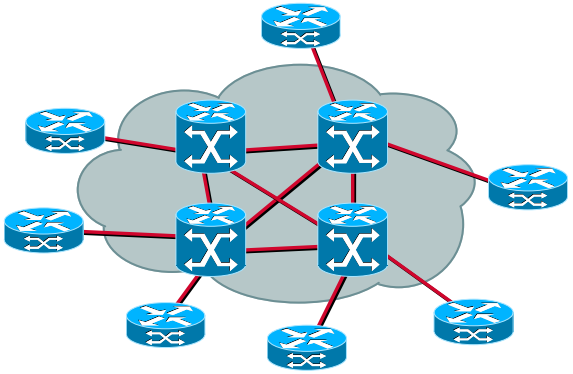

Application: Mixed POS/ATM Backbone

- MPLS provides
  - Tight integration of routers and ATM switches
  - End-to-end IP services
  - Internet scale VPNs
- Benefits
  - Network design flexibility
  - Transition to IP router backbone
  - Faster time to market

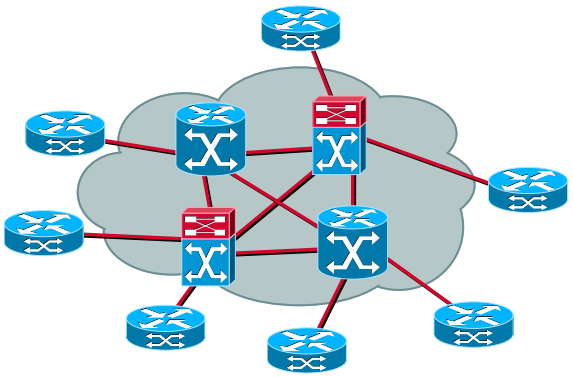

Applications: Enterprise Backbone

- **MPLS provides**
  - **Scalability**
  - **IP services**
  - **Traffic engineering**
- **Benefits**
  - **Flexibility**
  - **Reduced complexity for lower cost**

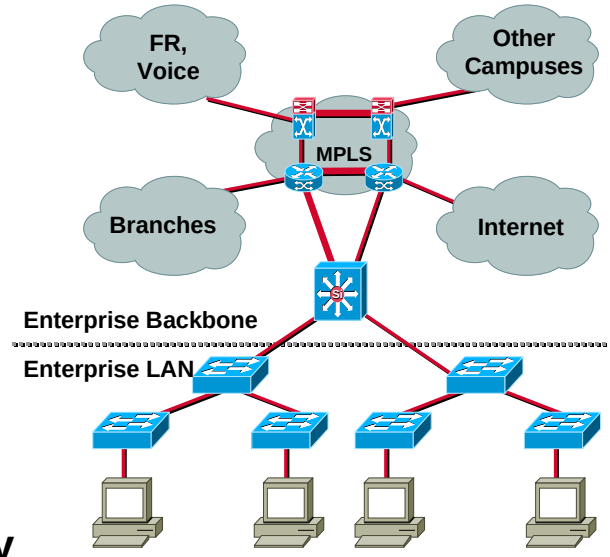

Topics

- **Motivations for MPLS**
- **MPLS Overview**
- **Applications**
- **Roadmap**

Leadership MPLS Solutions

- IP and ATM integration
- MPLS traffic engineering
- MPLS VPNs with integrated QoS

Leadership MPLS Solutions

- MPLS VPN management
- MPLS connection services

MPLS Platform Support

All Available Today!

BPX 8650

Catalyst 8540

Cisco 3600, 2600

BPX 8680

MGX 8850

Cisco 4500, 4700

Cisco 7200

LS1010

GSR 12000

Cisco 7500

Building on Open Standards

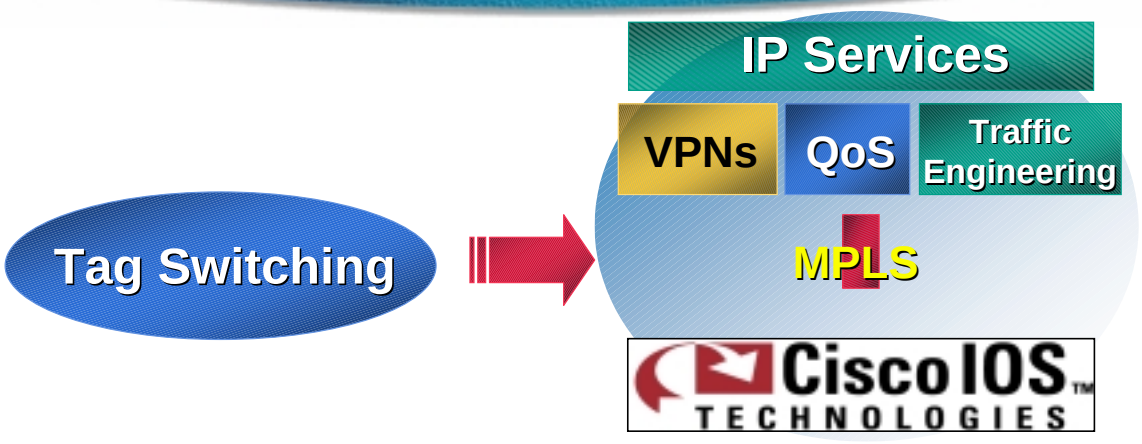

- MPLS is based on Cisco's tag switching
- Cisco is using MPLS as the basis for developing support for new value-added IP services
- Expect IETF ratification of the 12 MPLS RFCs in summer 2000

MPLS: The Cisco Advantage

- Industry IP leadership
- Most advanced MPLS solutions
- Broadest range of platforms supported in the industry today
- MPLS solutions deployed in real world production networks
- Standards-based solutions

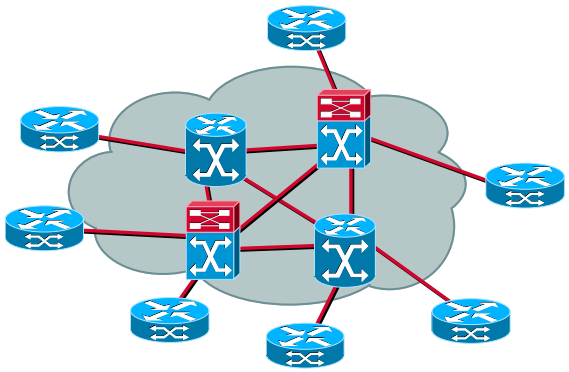

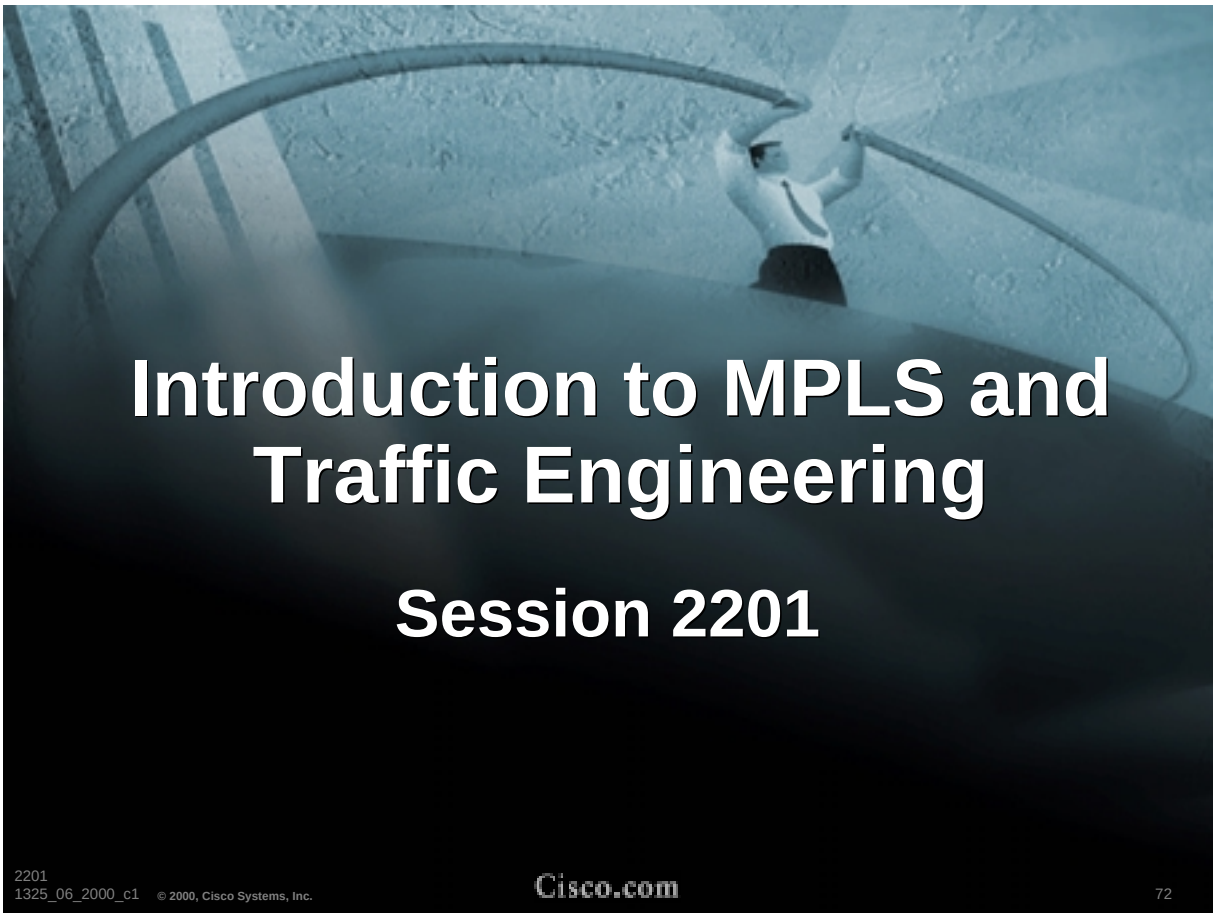

Introduction to MPLS and Traffic Engineering

Session 2201

CISCO SYSTEMS

EMPOWERING THE INTERNET GENERATION℠

