eecs 489 COMPUTER NETWORKS

Lecture 20:
MPLS, and VPN

Multi-Protocol Label Switching (MPLS)

Initial goal: speed up intra-domain IP forwarding by using circuit identifiers (fixed-length labels) instead of IP addresses

- borrow ideas from VC approach
  (but IP datagram still keeps IP address!)

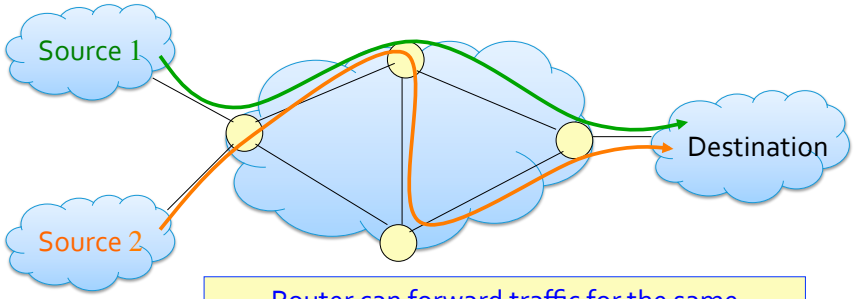

Router can forward traffic for the same destination on different interfaces/paths

Label Switching: Circuit Abstraction

Label-switched paths (LSPs):

- pre-compute a path for each "flow"
  - a "flow" can range from a single connection to a pair of APs or aggregated APs, etc.
- paths are "named" by the label at the path's entry point
- each MPLS router uses a different label to identify a flow
- "downstream" MPLS router tells upstream neighbor its label for each flow

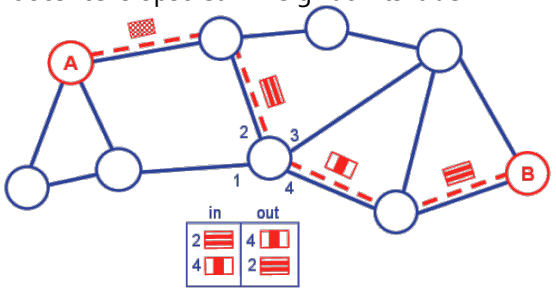

Label Swapping

At each hop, MPLS routers forward packets to outgoing interface based only on label value (doesn't even look at IP address)

- use label to determine outgoing interface
- replace incoming label with neighbor's label for the flow
- MPLS forwarding table distinct from IP forwarding tables

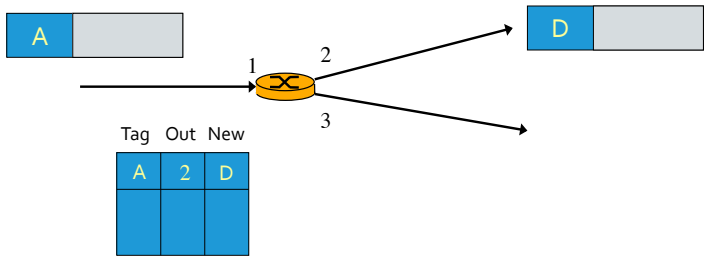

Label Distribution

Signaling protocol needed to set up forwarding

- responsible for disseminating signaling information
  - Label Distribution Protocol (LDP)
  - RSVP for Traffic Engineering (RSVP-TE)
- allows for forwarding along paths not otherwise obtained from IP routing (e.g., source-specific routing)
- must co-exist with IP-only routers

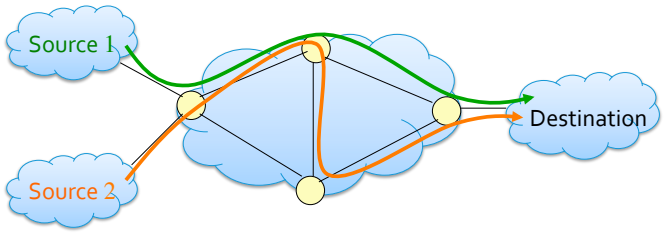

MPLS Encapsulation

Put an MPLS header in front of IP packet

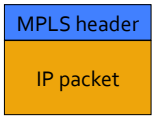

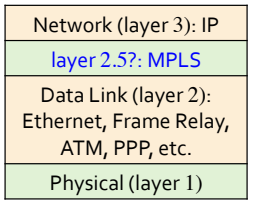

- MPLS header includes a label

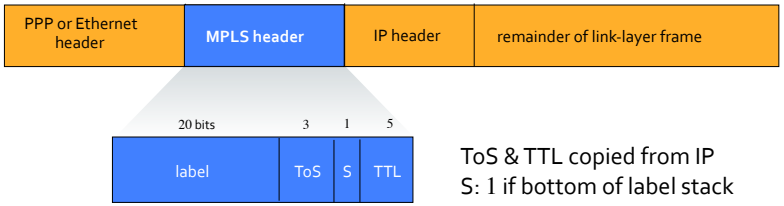

ToS & TTL copied from IP
S: 1 if bottom of label stack

BGP-Free Backbone Core

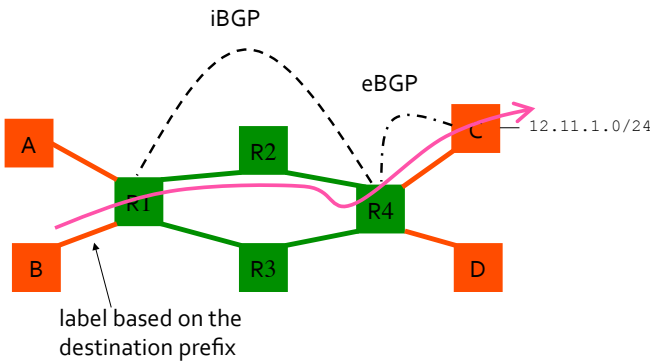

Routers R2 and R3 don't need to speak BGP

VPNs With Private Addresses

Why VPN?

Customer has several geographically distributed sites

- wants private communications over the public network
- wants a unique IP network connecting the sites
  - single IP addressing plan
  - virtual leased line connecting the sites
  - guaranteed quality of service

Providers have overprovisioned backbones

- want to sell pseudo-wires (leased lines) that allow for increased backbone utilization
- want technology that has
  - low configuration and maintenance costs
  - is scalable to the number of customers, i.e., core states depend on topology, not number of customers

Recall: Customer-based VPN

Encrypt packets at network entry and decrypt at exit

Eavesdropper cannot snoop the data
or determine the real source and destination

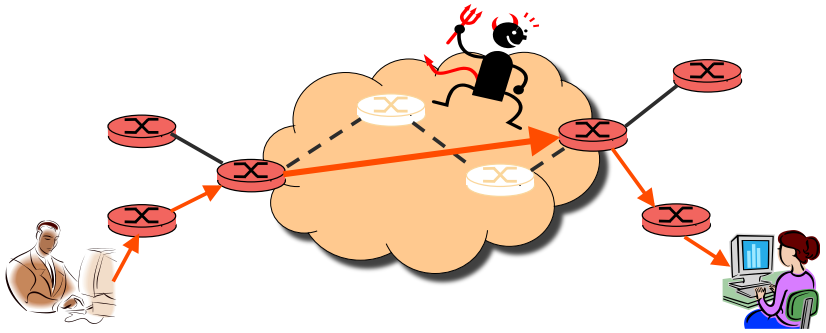

Network VPNs

Customer based:
- customer buys own equipment, configures IPSec tunnels across the global Internet, manages addressing and routing
- ISP plays no role
- customer has more control over security and ISP choices, but requires skills

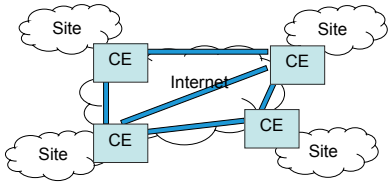

Provider based:
- provider manages all the complexity of the VPN, usually with MPLS
- customer simply connects to the provider equipment

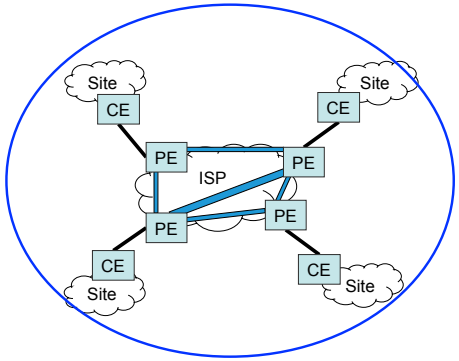

Types of MPLS Routers

Customer edge (CE) routers:
- do not speak MPLS, do not recognize labels at all
- speak eBGP with MPLS routers on provider network to advertise APs
  - or statically configured with allocated APs

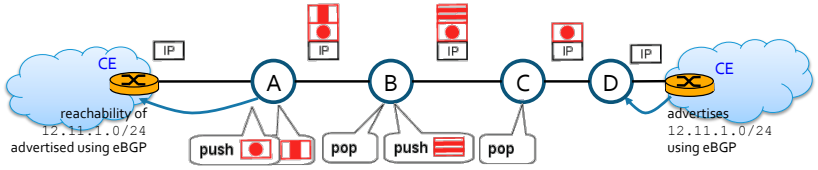

MPLS Routers

Provider routers:
- provider edge (PE) : routers A and E
  - push (at ingress) or pop (at egress) label onto stack
  - forward IP packets to/from customer routers
- core (P) : routers B, C, and D
  - swap (pop+push) label on top of stack
  - doesn't interact with customer routers

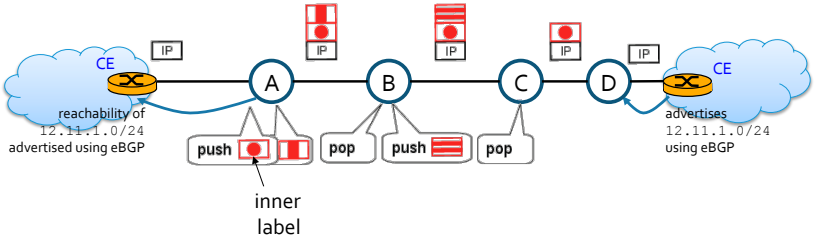

Provider-based VPN

Layer 3 BGP/MPLS VPNs (RFC2547)

- provides isolation: multiple logical networks over a single, shared physical infrastructure
- uses BGP to exchange routes
  - eBGP to announce APs to PE routers
- MPLS to forward traffic
  - tunneling: P core routers don't have to do routing, just label switching

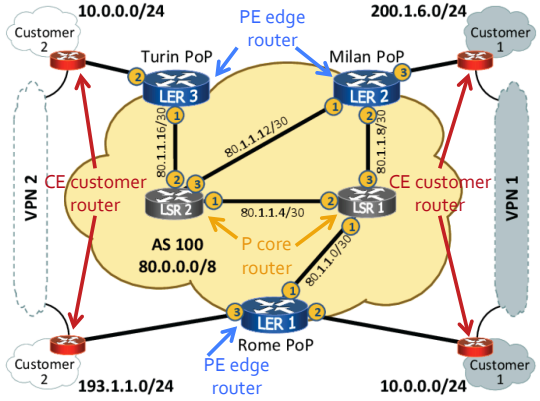

High-Level Overview of Operation

IP packets arrive at provider edge (PE) router

Destination IP looked up in "virtual" forwarding table

- there are multiple such tables, one per customer

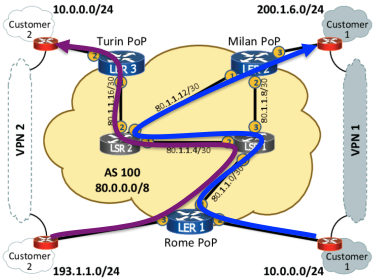

Datagram sent to customer's network using tunneling (i.e., an MPLS label-switched path)

To Use Level 3 BGP/MPLS VPN

Two steps needed:

1. set up the VPN
2. forward packets on the VPN

Identifying a BGP/MPLS VPN

Three things are needed to identify a BGP/MPLS VPN

1. inner label: a way for the provider edge (PE) routers at each end of a VPN to associate a VPN with its owner's customer edge (CE) router
2. VPN-APs: a way for the customer's address prefixes (APs) to be advertised by BGP
   - the issue is: since customers can use private address ranges (`10/8`, `172.16/12`, and `192.168/16`), how to differentiate the same private address range that has been chosen and used by different customers?
3. outer label: the MPLS labels used by provider's core (P) routers to identify a VC

Setup: Inner Label

Provider-edge (PE) routers:

- set up a Virtual Routing and Forwarding (VRF) table for each customer AP
  - the VRF ID serves as the inner label for the VPN

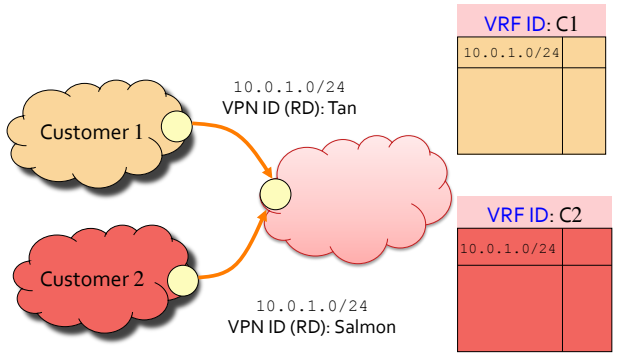

Setup: VPN-APs

Provider-edge (PE) routers:

- use Multi-Protocol BGP's Route Distinguisher (RD) as the VPN ID to differentiate the same APs of different customers
  - use MP-BGP to announce VPN-APs reachability, along with their inner labels
  - runs iBGP to other edge routers to distribute VPN-AP reachabilities

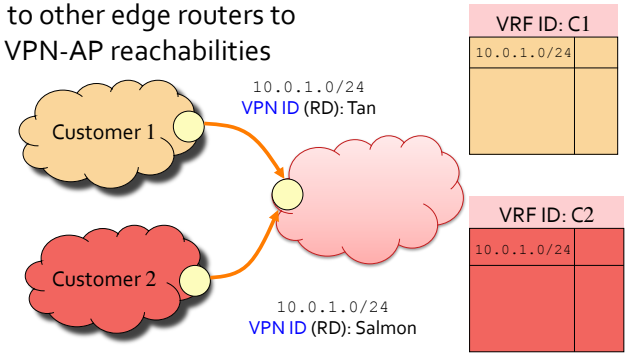

Setup: Outer Label

Both provider-edge (PE) and core (P) routers:

- run MPLS
- use LDP (Label Distribution Protocol) to set up outer labels for forwarding
  - the PE router advertising a customer AP (i.e., the "destination" or egress router) initiates LDP to distribute labels

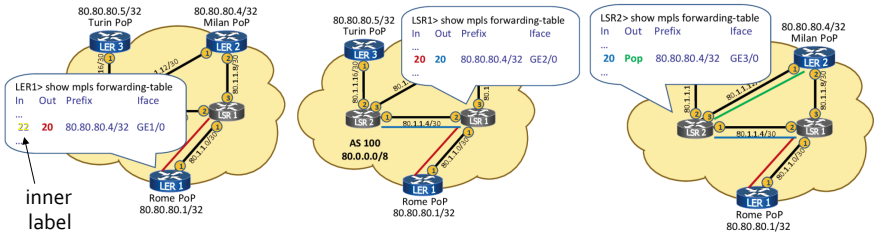

To use Level 3 BGP/MPLS VPN

Two steps are needed to use a level 3 BGP/MPLS VPN:

1. Set up the VPN
2. Forward packets on the VPN

Forwarding in BGP/MPLS VPNs

Step 1: packet arrives from CE router at PE router's incoming interface

- look up customer's VRF to determine egress PE and inner label (Label I)

Step 2: egress PE lookup, add corresponding outer label (Label O, also at customer's VRF)

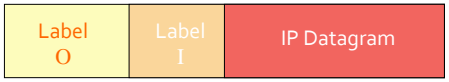

Forwarding

Ingress PE router encapsulates IP packet in MPLS with outer and inner labels

Two-label stack is used for packet forwarding

- top label indicates next-hop P router (outer label)
- second label indicates outgoing CE interface / VRF (inner label)

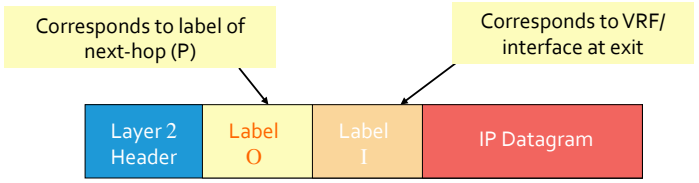

Forwarding on BGP/MPLS VPNs

Source CE router sends IP packet to ingress PE router that advertises destination AP

Ingress PE router looks up egress PE router's virtual interface address and the inner label for destination AP, then encapsulates IP packet in MPLS with outer and inner labels

Core P routers along the path swap outer labels
Penultimate core P router pop outer label only

Egress PE router uses inner label to look up VRF and forward packet to customer CE router

Packet Forwarding

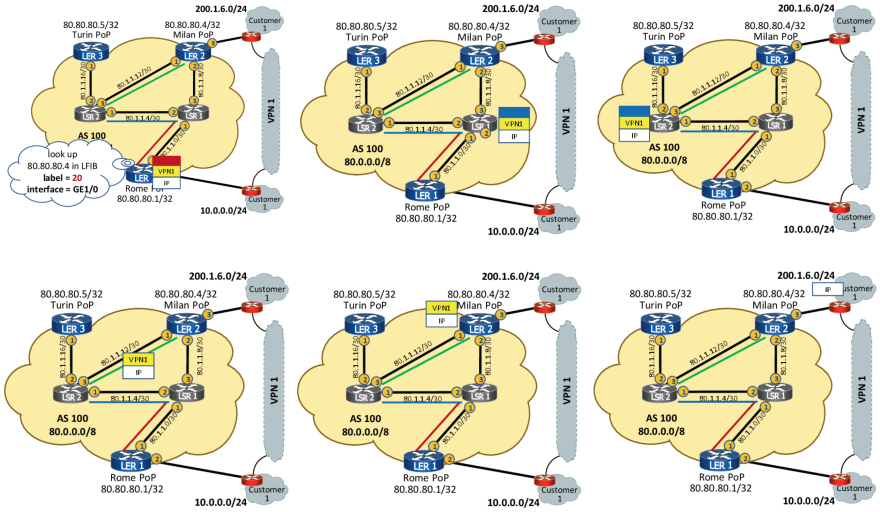

Advantages of MPLS VPN

Customer's adding or changing APs does not require manual configuration at provider

Core P routers do not need to know customer's CE routers or APs ⇒ forwarding tables only need to scale to number of edge PE routers, not number of customers, APs, or VPNs

The only manual configurations required are at the edge PE routers:

- VRF ID and customer's CE router's IP address
- MP-BGP Route Distinguisher as VPN ID

Status of MPLS

Deployed in practice

- BGP-free backbone/core
- Virtual Private Networks
- Traffic engineering

Challenges

- protocol complexity
- configuration complexity
- difficulty of collecting measurement data