eecs589

Advanced Computer Networks

[CDP13] Cittadini, Di Battista, Patrignani, "MPLS Virtual Private Networks," in Haddadi and Bonaventure (eds.), *Recent Advances in Networking Research*, ACM SIGCOMM eBook Ch. 6, 2013

Datagram Networks

No call setup at network layer

Routers: no state about end-to-end connections

- no network-level concept of "connection"

Packets forwarded using destination host address

- packets between same source-destination pair may take different paths

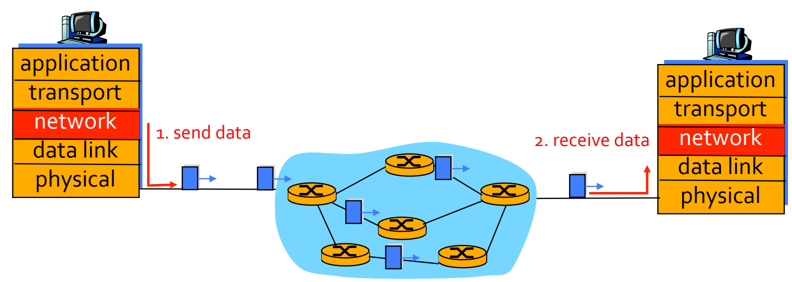

Virtual Circuits (VC)

Datagram network provides network-layer connectionless service

VC network provides network-layer connection-oriented service

Analogous to the transport-layer services, but:

- service: host-to-host
- implementation: in the core

"Source-to-destination path behaves much like a telephone circuit"

- Advantage: link, router resources (bandwidth, buffers) may be allocated to VC
  - QoS provisioning and performance guarantee (ideally)
- Disadvantages: network actions along source-to-destination path
  - call setup, teardown for each call before data can flow
  - each packet carries VC identifier
  - every router on source-destination path must maintain "state" for each passing connection

Virtual Circuits

Signalling protocol:

- used to setup, maintain, teardown VC
- e.g., LDP, ReSource reserVation Protocol (RSVP)

A VC consists of:

1. path from source to destination
2. VC numbers, one number for each link along path
3. entries in forwarding tables in routers along path

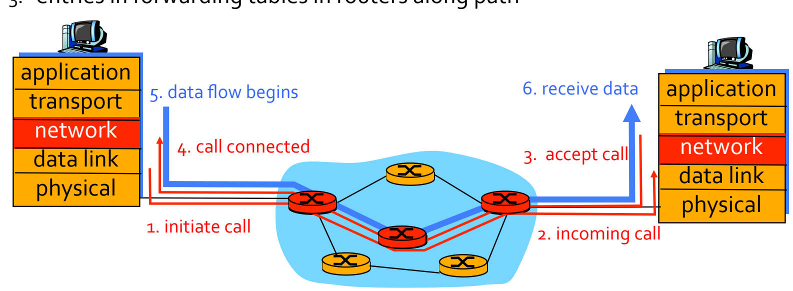

VC Forwarding Table

Packet belonging to VC carries a VC number

VC number must be changed for each link

New VC number obtained from forwarding table

Examples: MPLS, Frame-relay, ATM, PPP

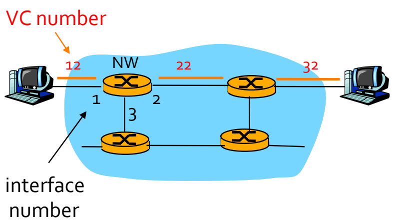

Forwarding table on router NW:

incoming interface	incoming VC#	outgoing interface	outgoing VC#
1	12	2	22
2	63	1	18
3	7	2	17
1	97	3	87
…	…	…	…

Routers maintain connection state information!

Datagram vs. VC Networks

Datagram network:

- destination address in packet determines next hop
- routes may change during session

Virtual circuit network:

- each packet carries tag (virtual circuit ID), tag determines next hop
- fixed path determined at call setup time, remains fixed through call
- routers maintain per-call state

Multi-Protocol Label Switching (MPLS)

Initial goal: speed up IP forwarding by using fixed length label (instead of IP address) to do forwarding

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

Network layer 3: IP
layer 2.5?: MPLS
Data Link layer 2: Ethernet, Frame relay, ATM, PPP, etc.
Physical layer 1

MPLS Encapsulation

Put an MPLS header in front of the packet

- MPLS header includes a label
- label switching between MPLS-capable routers

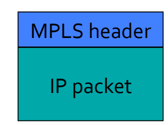

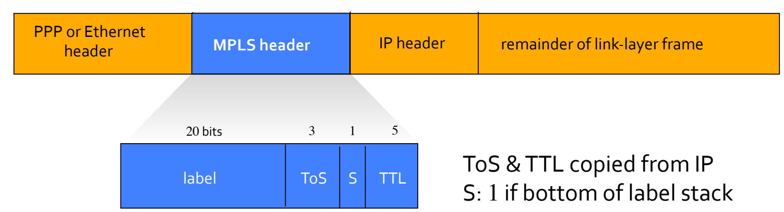

ToS & TTL copied from IP

S: 1 if bottom of label stack

Label Switching

Forwards packets to outgoing interface based only on label value (doesn't even look at IP address)

- MPLS forwarding table distinct from IP forwarding tables
- "downstream" MPLS router tells upstream neighbor the label it is using to identify a "flow"
- different labels used for each pair of MPLS routers
- a "flow" can range from a single connection to a pair of APs or aggregated APs, etc.

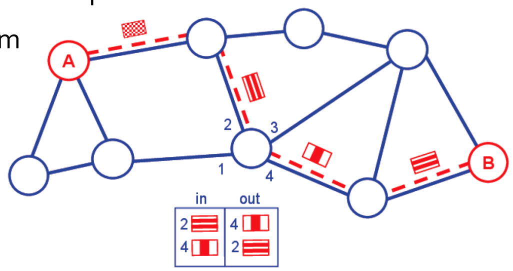

MPLS Forwarding Tables

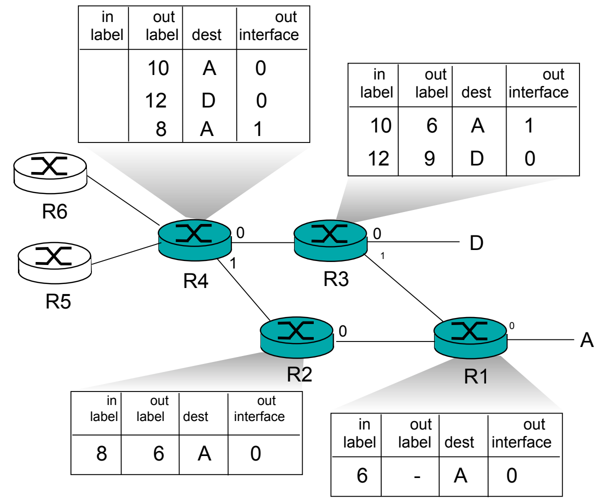

R4:

in label	out label	dest	out interface
	10	A	0
	12	D	0
	8	A	1

R3:

in label	out label	dest	out interface
10	6	A	1
12	9	D	0

R2:

in label	out label	dest	out interface
8	6	A	0

R1:

in label	out label	dest	out interface
6	-	A	0

MPLS Routers

Customer (edge) routers:

- doesn't speak MPLS, doesn't recognized labels at all
- speaks eBGP with MPLS routers on provider network to advertise APs
  - or statically configured with allocated APs

Provider routers:

- edge routers:
  - push (at ingress) or pop (at egress) label onto stack
  - forward IP packets to/from customer routers
- core routers:
  - swap (pop+push) label on top of stack
  - doesn't interact with customer routers

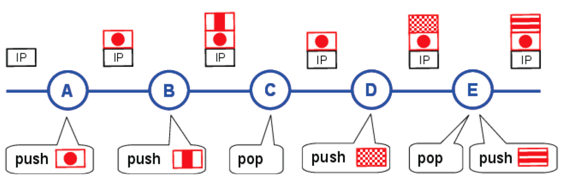

Label Distribution

Signaling protocol needed to set up forwarding

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

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

BGP-Free Backbone Core

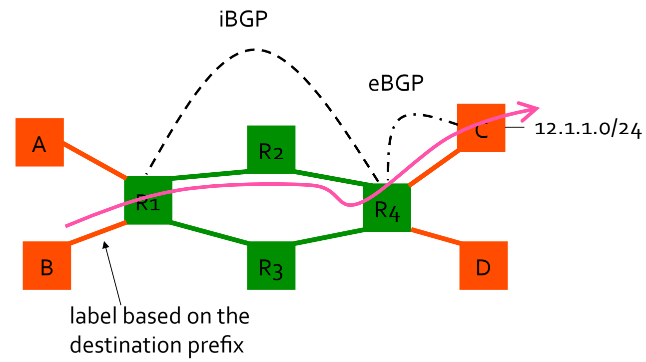

Routers R2 and R3 don't need to speak BGP

VPNs With Private Addresses

Why VPN?

Customer has several geographically distributed sites
- wants a unique IP network connecting the sites
  - single IP addressing plan
  - virtual leased line
  - 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
    core states depend on topology, not number of customers

Customer with Multiple APs

Customer has several geographically distributed sites
- wants a unique IP network connecting the sites
  - multiple APs
  - virtual leased line
- uses eBGP to announce APs to edge routers

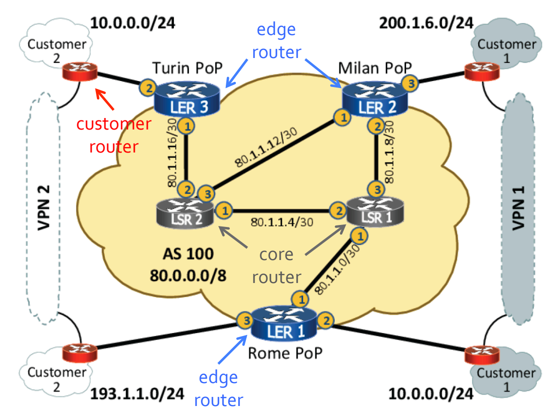

Provider Router Configuration

Edge routers:

- addressed by their virtual interfaces (to ensure global visibility)
- set up a Virtual Routing and Forwarding (VRF) table for each customer AP
- the VRF ID serves as the MPLS inner label for the VPN
- 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 MPLS inner labels
- runs iBGP to other edge routers to distribute VPN-AP reachabilities

Provider Router Configuration

Both edge and core routers:

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

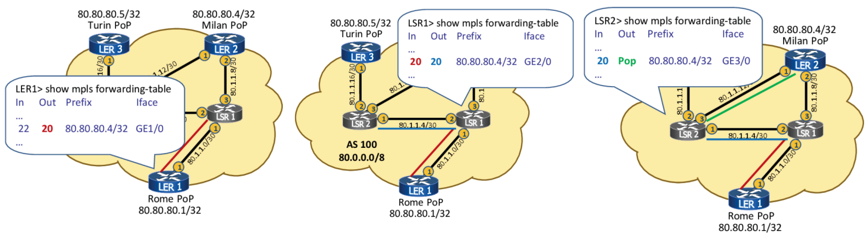

Packet Forwarding

Source customer router sends IP packet to edge router advertising destination AP

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

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

Edge egress router uses inner label to look up VRF and forward packet to customer router

Packet Forwarding

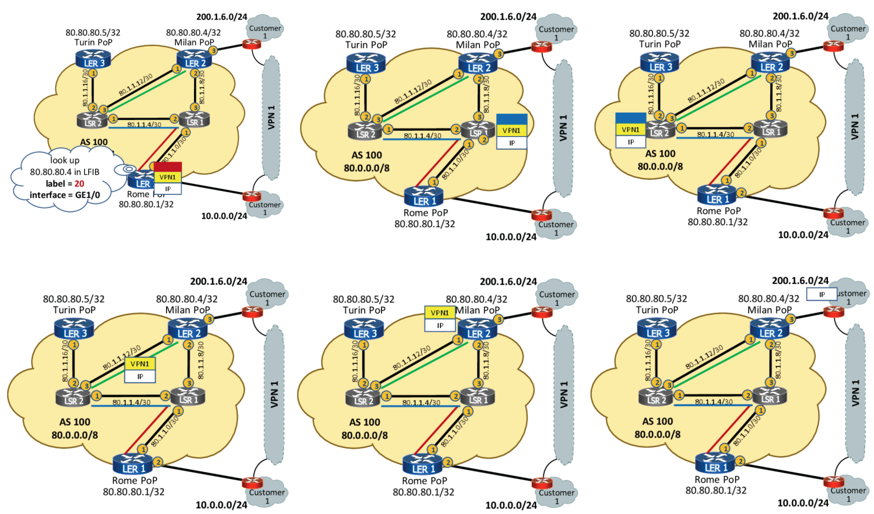

Advantages of MPLS VPN

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

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

The only manual configurations required are at the edge routers:

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