15-441 15-641 Computer Networking

Virtual Circuits, MPLS

Peter Steenkiste

Fall 2013

www.cs.cmu.edu/~prs/15-441-F13

Outline

- Circuit switching refresher
- Virtual Circuits - general
  - Why virtual circuits?
  - How virtual circuits? -- tag switching!
- Two modern implementations
  - ATM - teleco-style virtual circuits
  - MPLS - IP-style virtual circuits

Packet Switching

- Source sends information as self-contained packets that have an address.
  - Source may have to break up single message in multiple
- Each packet travels independently to the destination host.
  - Routers and switches use the address in the packet to determine how to forward the packets
- Destination recreates the message.
- Analogy: a letter in surface mail.

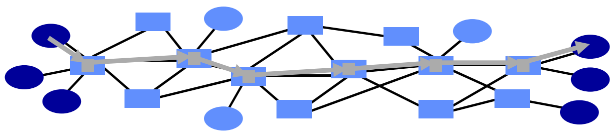

Circuit Switching

- Source first establishes a connection (circuit) to the destination.
  - Each router or switch along the way may reserve some bandwidth for the data flow
- Source sends the data over the circuit.
  - No destination address needed - routers know the path
- The connection is torn down.
- Example: traditional telephone network.

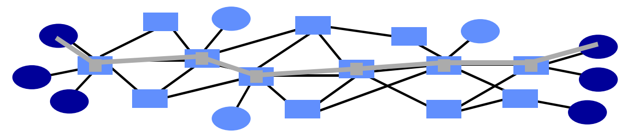

Circuit Switching

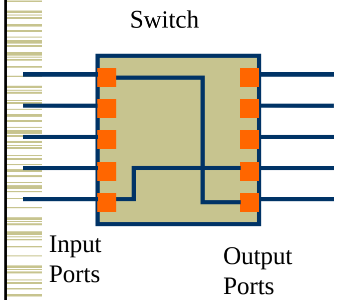

- Switch remembers how to forward data
  - No addresses!
- Many options
  - Between specific wires (circuit = wire)
  - Between timeslots (TDMA on each wire)
  - Between frequencies (FDMA on each wire)

Lots of Wires!

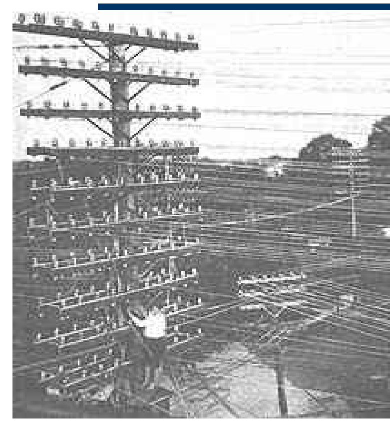

Circuit Versus Packet Switching

Circuit Switching

- Fast switches can be built relatively inexpensively
- Inefficient for bursty data
- Predictable performance (e.g. hard QoS)
- Requires circuit establishment before communication

Packet Switching

- Switch design is more complex and expensive
- Allows statistical multiplexing
- Difficult to provide QoS guarantees
- Data can be sent without signaling delay and overhead

Can we get the benefits of both?

Virtual Circuits

- Each wire carries many "virtual" circuits.
- Forwarding based on virtual circuit (VC) identifier
  - IP header: src, dst, etc.
  - Virtual circuit header: just "VC"
  - A path through the network is determined when VC is established
  - Use statistical multiplexing for efficiency
- Can support wide range of quality of service.
  - No guarantees: best effort service
  - Weak guarantees: delay < 300 msec, …
  - Strong guarantees: e.g. equivalent of physical circuit

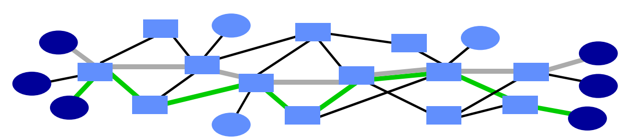

Packet Switching and Virtual Circuits: Similarities

- “Store and forward” communication based on an address.
  - Address is either the destination address or a VC identifier
- Must have buffer space to temporarily store packets.
  - E.g. multiple packets for some destination arrive simultaneously
- Multiplexing on a link is similar to time sharing.
  - No reservations: multiplexing is statistical, i.e. packets are interleaved without a fixed pattern
  - Reservations: some flows are guaranteed to get a certain number of “slots”

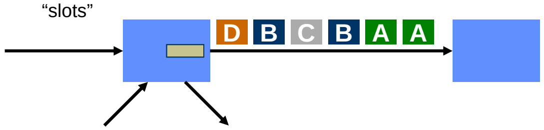

Packet switched vs. VC

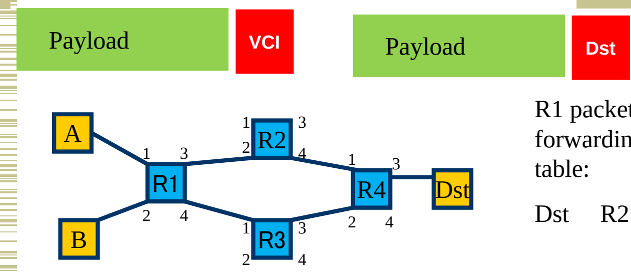

R1 packet forwarding table:

Dst R2

R1 VC table:

VC 1 R2

VC 2 R3

Different paths to same destination!

(useful for traffic engineering!)

Virtual Circuit Forwarding

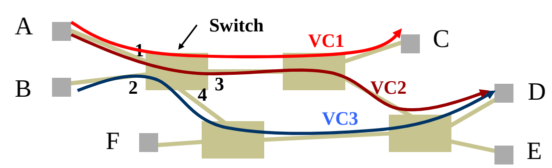

Address	Next Hop
VC1	3
VC2	3
VC3	4
VC4	?
VC5	?

- Address used for look up is a virtual circuit identifier (VC id)
- Forwarding table entries are filled in during signaling
- VC id is often shorter than destination address

How to Pick a VC Id?

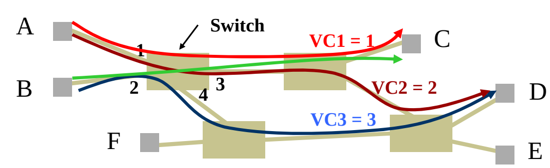

- When B establishes green virtual circuit, how does it know what VC ids are available?
- Even worse: every VC id may already be used on a link along the path to the destination
- Solution: VC id swapping

VC id Swapping

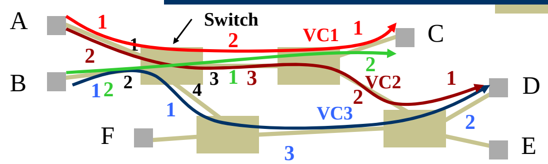

Address	Next Hop	Next id
VC1 = 1	3	2
VC2 = 2	3	3
VC3 = 1	4	1
VC4 = 2	3	1

- Look up is based on VC id in header + incoming port number
- Forwarding table specifies outgoing port and new VC id
- VC id conflicts can be resolved locally during signaling

Connections and Signaling

- Permanent vs. switched virtual connections (PVCs, SVCs)
  - static vs. dynamic. PVCs last “a long time”
    - E.g., connect two bank locations with a PVC
    - SVCs are more like a phone call
  - PVCs administratively configured (but not “manually”)
  - SVCs dynamically set up on a “per-call” basis
- Topology
  - point to point, point to multipoint, multipoint to multipoint
- Challenges: How to configure these things?
  - What VCI to use?
  - Setting up the path

SVC Connection Setup

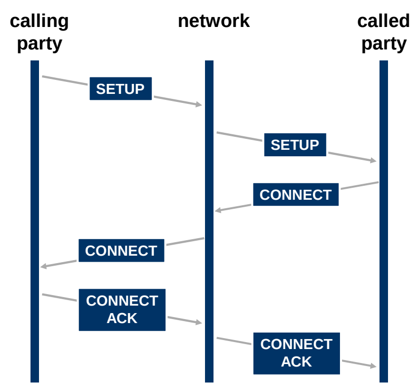

Virtual Circuits In Practice

- ATM: Teleco approach
  - Kitchen sink. Based on voice, support file transfer, video, etc., etc.
  - Intended as IP replacement. That didn’t happen. :)
  - Today: rarely used.
- MPLS: The “IP Heads” answer to ATM
  - Stole good ideas from ATM
  - Integrates well with IP
  - Today: Used inside some networks to provide VPN support, traffic engineering, simplify core.
- Other networks just run IP.
- Older tech: Frame Relay
  - Only provided PVCs. Used for quasi-dedicated 56k/T1 links between offices, etc. Slower, less flexible than ATM.

Asynchronous Transfer Mode ATM

- Connection-oriented, packet-switched
  - (e.g., virtual circuits).
- Teleco-driven – fold the Internet into the telephone system
- Goals:
  - Handle voice, data, multimedia
  - Support both PVCs and SVCs
  - Replace IP. (didn't happen…)
- Important feature: Cell switching

ATM History

- Telephone companies supported voice telephony: 4 kHz analog, 64 kbs digital
- They provided lines for data networking
  - ISDN: 64 kbps and faster channels
  - T1 (1.544 Mbps)
  - T3 (44.736 Mbps)
- Wanted to become the primary service provider for data networking services
  - file transfer: bursty, many Mbps peak
  - database access: bursty, low latency
  - Multimedia: synchronized
  - Video: 6 MHz analog, 1.2-200 Mbps digital

Cell Switching

- Small, fixed-size cells
  [Fixed-length data][header]
- Why?
  - Efficiency: All packets the same
    - Easier hardware parallelism, implementation
  - Switching efficiency:
    - Lookups are easy -- table index.
  - Result: Very high cell switching rates.
  - Initial ATM was 155Mbit/s.
    - Ethernet was 10Mbit/s at the same time. (!)

ATM Features

- Fixed size cells (53 bytes).
- Virtual circuit technology using hierarchical virtual circuits (VP,VC).
- PHY (physical layer) processing delineates cells by frame structure, cell header error check.
- Support for multiple traffic classes by adaptation layer.
  - E.g. voice channels, data traffic
- Elaborate signaling stack.
  - Backwards compatible with respect to the telephone standards
- Standards defined by ATM Forum.
  - Organization of manufacturers, providers, users

The ATM Cell (UNI)

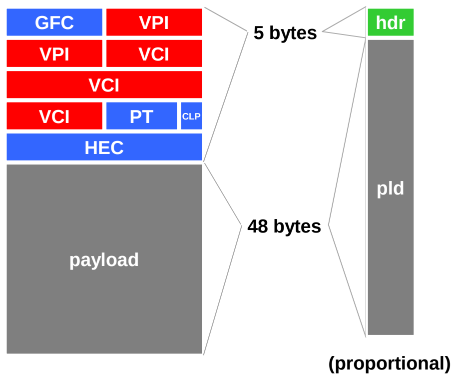

Virtual Path Trunking

- Virtual path is a bundle of VCs
  - All follow the same path through the network
  - No remapping of VC id in virtual paths – fast setup
- Allows aggregated resource management and fault recovery.

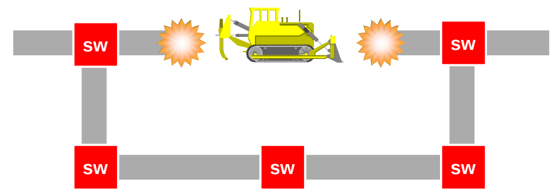

Why 53 Bytes?

- Small cells favored by voice applications
  - delays of more than about 10 ms require echo cancellation
  - each payload byte consumes 125 µs (8000 samples/sec)
- Large cells favored by data applications
  - Five bytes of each cell are overhead
- France favored 32 bytes
  - 32 bytes = 4 ms packetization delay.
  - France is 3 ms wide.
  - Wouldn't need echo cancellers!
- USA, Australia favored 64 bytes
  - 64 bytes = 8 ms
  - USA is 16 ms wide
  - Needed echo cancellers anyway, wanted less overhead
- Compromise

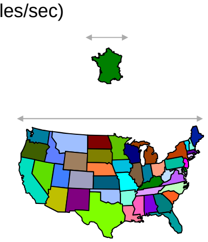

ATM Adaptation Layers

1	2	3	4	5
synchronous		asynchronous		
constant	variable bit rate			
connection-oriented			connectionless	

- **AAL 1: audio, uncompressed video**
- **AAL 2: compressed video**
- **AAL 3: long term connections**
- **AAL 4/5: data traffic**

IP over ATM Switched VCs

- When sending IP packets over an ATM network, set up a VC to destination.
  - ATM network can be end to end, or just a partial path
  - ATM is just another link layer
- Virtual connections can be cached.
  - After a packet has been sent, the VC is maintained so that later packets can be forwarded immediately
  - VCs eventually times out
- Properties.
  - Overhead of setting up VCs (delay for first packet)
  - Complexity of managing a pool of VCs
  - + Flexible bandwidth management
  - + Can use ATM QoS support for individual connections (with appropriate signaling support)

LAN Emulation

- Motivation: making a non-broadcast technology work as a LAN.
  - Focus on 802.x environments
- Approach: reuse the existing interfaces, but adapt implementation to ATM.
  - MAC - ATM mapping
  - multicast and broadcast
  - bridging
  - ARP
- Example: Address Resolution "Protocol" uses an ARP server instead of relying on broadcast.
- A lot of hard work!

IP over ATM Permanent VCs

- Establish a set of "ATM pipes" that defines connectivity between routers.
- Routers simply forward packets through the pipes.
  - Each statically configured VC looks like a link
- Properties.
  - Some ATM benefits are lost (per flow QoS)
  - + Flexible but static bandwidth management: can force flows along specific paths
  - + No set up overheads

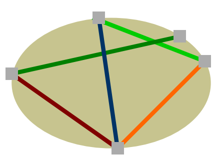

ATM Discussion

- Vision: ATM is a replacement for IP.
  - Could carry both traditional telephone traffic (CBR circuits) and other traffic (data, VBR)
  - Simple switching core: forwarding based on VC identifiers
  - Better than IP, since it supports QoS, traffic engineering
- Reality: Traffic engineering benefits were attractive
  - Fast VCI lookup became less critical over time
- But: Complex technology.
  - Signaling software is very complex
  - Technology did not match people's experience with IP
    - supporting connection-less service model on connection-based technology is painful
    - deploying ATM in LAN is complex (e.g. broadcast)
  - With IP over ATM, a lot of functionality is replicated

MPLS

- Multi-Protocol Label Switching
- Bringing virtual circuit concept into IP
- Driven by multiple forces
  - QoS
  - Traffic engineering
  - High performance forwarding
  - VPN

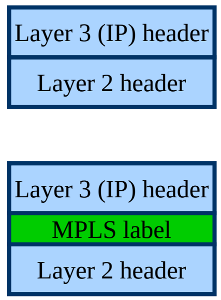

Some MPLS slides from H. Zhang

MPLS Vocabulary: LSP

- Label-switched path (LSP)
  - Simplex path through interior network

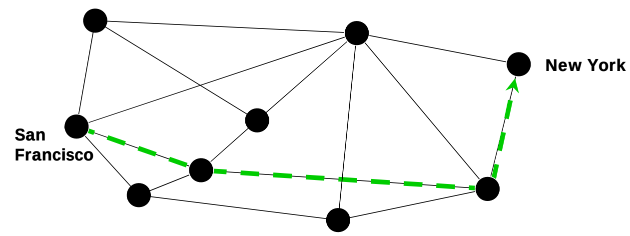

MPLS Vocabulary: LSR

- Label-switching router (LSR) performs
  - MPLS packet forwarding
  - LSP setup

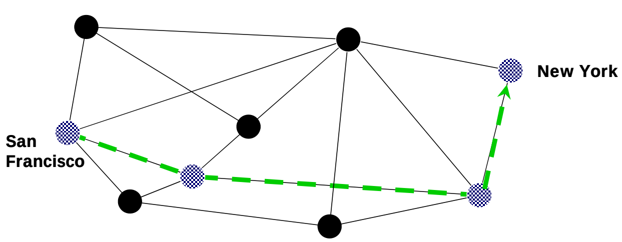

MPLS Vocabulary: LER

- Label Edge Router (LER)
  - Ingress and egress node of LSP
  - Packet enters and leaves the LSP

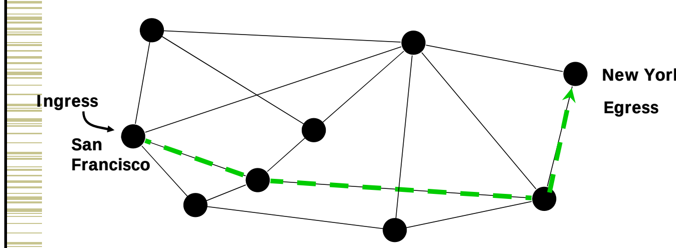

MPLS Vocabulary: Transit

- Transit router
  - Zero or more transit routers
  - Swaps MPLS label
  - Sends traffic to next hop in LSP

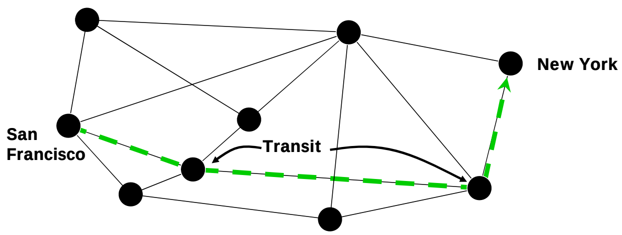

MPLS Header

- IP packet is encapsulated in MPLS header and sent down LSP

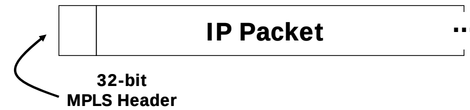

- IP packet is restored at end of LSP by egress router
  - TTL is adjusted, transit LSP routers count towards the TTL
  - MPLS is an optimization – does not affect IP semantics

MPLS Header

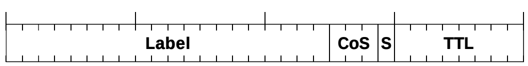

- Label
- Class of service
- Stacking bit
  - Remember me?
- Time to live
  - Decrement at each LSR, or
  - Pass through unchanged

Forwarding Equivalence Classes

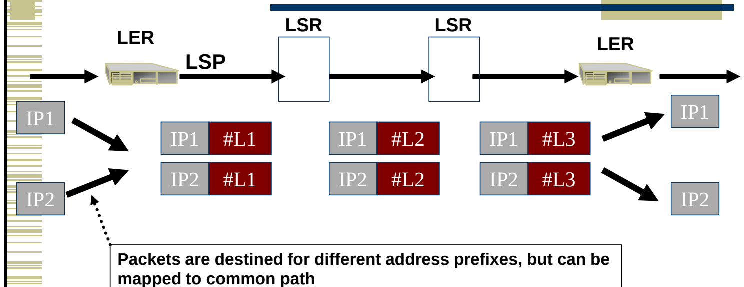

- FEC = "A subset of packets that are all treated the same way by a router"
- The concept of FECs provides for a great deal of flexibility and scalability
- In conventional routing, a packet is assigned to a FEC at each hop (i.e. L3 look-up), in MPLS it is only done once at the network ingress.

LSPs Driven by Routing

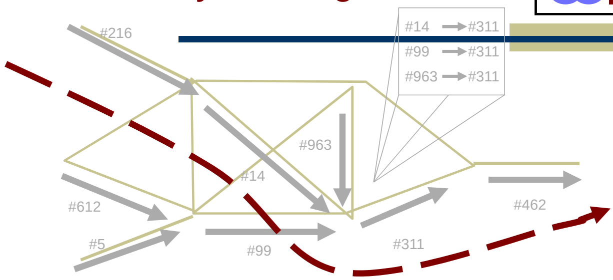

- A LSP is actually part of a tree from every source to that destination (unidirectional).
- Control protocol (e.g. LDP) builds that tree using existing IP forwarding tables to route the control messages.

MPLS Builds on Standard IP

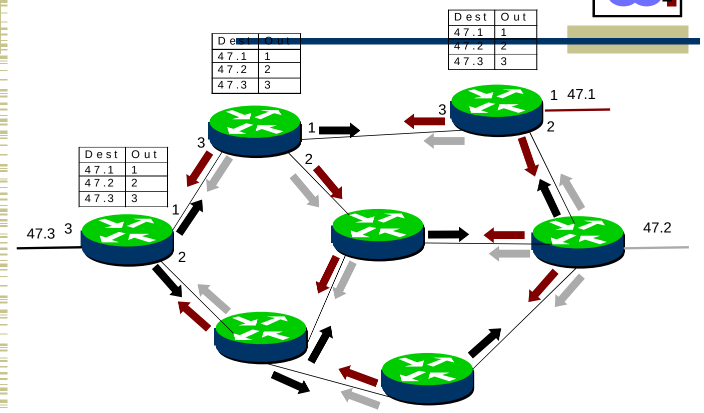

- **Destination based forwarding tables as built by OSPF, IS-IS, RIP, etc.**

IP Forwarding – Hop-by-Hop Control

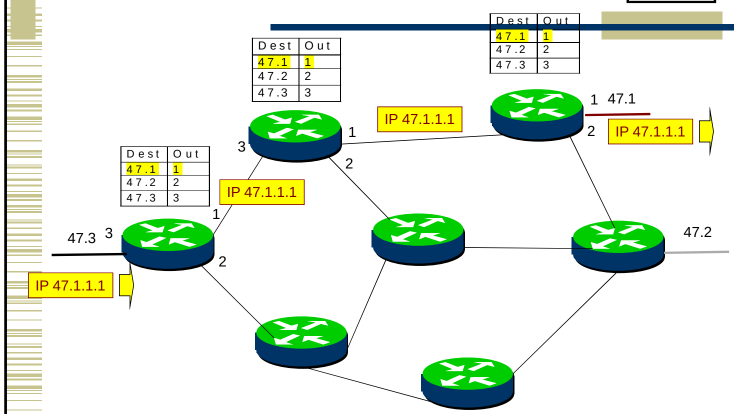

Label Switched Path (LSP)

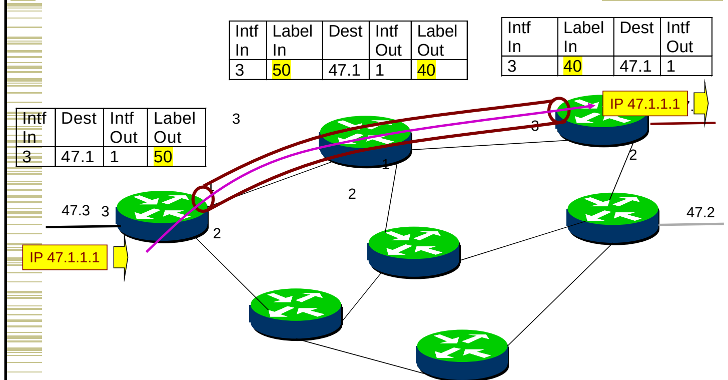

Explicityly Routed - ER-LSP

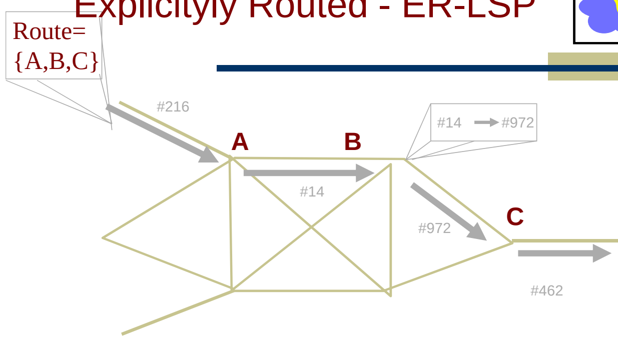

ER-LSP follows route that source chooses. In other words, the control message to establish the LSP (label request) is ***source routed***.

Explicitly Routered LSP - Example

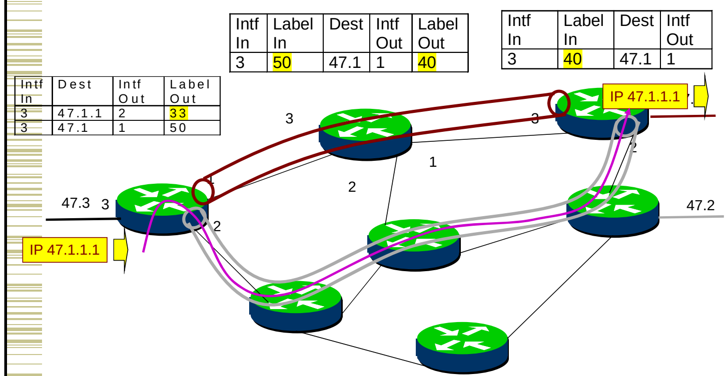

Protocol Comparison

	Forwarding	Control Protocols
Ethernet	**Dest MAC address** **Exact match**	**Learning** **Spanning tree**
IP	**Dest IP address** **Longest prefix match**	**Routing protocol**
TDM	**Time slot, exact match** **Time Slot Exchange (TSE)**	**E2E signaling protocol** **Routing protocol**
ATM	**Label, exact match** **Label swapping**	**E2E signaling protocol** **Routing protocol**
MPLS	**Label, Dest IP Address**	**Flexible signaling** **Routing Protocol**

VLAN Introduction

- VLANs logically segment switched LANs based on organization or function, independent of their physical location in the network
  - Devices on a VLAN share their own (private) LAN
  - Form their own IP subnet
- Offers many benefits:
  - Performance: limits broadcast messages to the VLAN – improves scalability
  - Security: isolates VLAN – VLANs connected by routers with smarter filtering capabilities
  - Management: manage network topology without changing the physical topology

VLAN Example

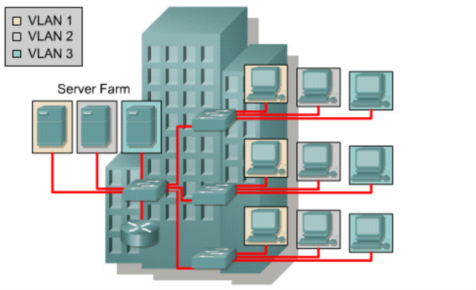

- A switch creates a broadcast domain
- VLANs help manage broadcast domains
- VLANs can be defined on port groups, users, or protocols
- LAN switches and network management software provide a mechanism to create VLANs

VLAN Logical Topology

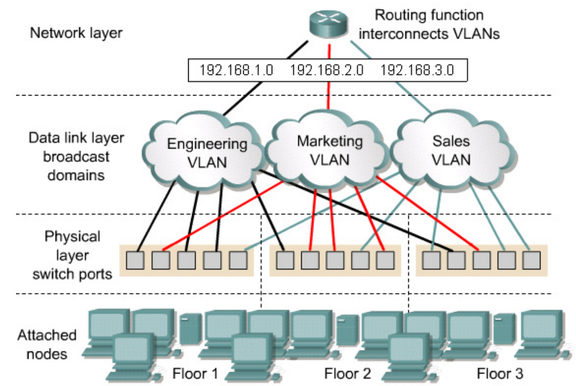

VLAN Types

- VLANs are implemented by switches
- VLAN memberships can be controlled by a switch in different ways, based on:
  - Port: incoming ports are tagged with VLAN ID
  - MAC address: switch has (MAC, VLAN ID) table
  - Protocol: switch as (protocol, VLAN ID) table
- The frame headers are encapsulated or modified to insert a VLAN ID
  - Is inserted by first switch before forwarding packet
  - Removed by last switch before forwarding to the device

Example: 802.1Q Standard for VLANs over Ethernet

- A 32 bit VLAN header is inserted after the MAC addresses

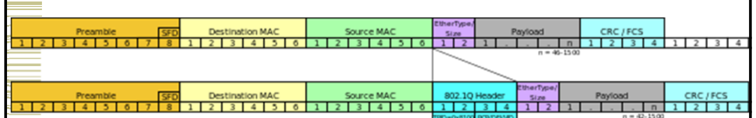

- Header consists of
  - Tag Protocol Identifier (16b): single value that marks frame as a VLAN frame
  - Control bits (4b): mostly priority
  - VLAN Identifier (12b): identifies VLAN

Take Home Points

- Costs/benefits/goals of virtual circuits
- Cell switching (ATM)
  - Fixed-size pkts: Fast hardware
  - Packet size picked for low voice jitter. Understand trade-offs.
  - Beware packet shredder effect (drop entire pkt)
- Tag/label swapping - basis for most VCs.
  - Makes label assignment link-local. Understand mechanism.
- MPLS - IP meets virtual circuits (links)
  - Used for VPNs, traffic engineering, reduced core routing table sizes
  - Management of ISPs at layer 3
- Virtual LANs – manage LANs in software
  - Simplifies management of edge networks at layer 2
  - Set up by manager based on organizational structure – no tag swapping