15-441 15-641 Computer Networking

Virtual Circuits, MPLS, VLANs

Peter Steenkiste

Fall 2015

www.cs.cmu.edu/~prs/15-441-F15

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
- Virtual LANs
  - How do they differ?

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.

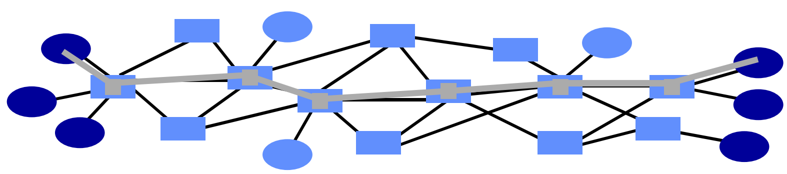

Circuit Switching

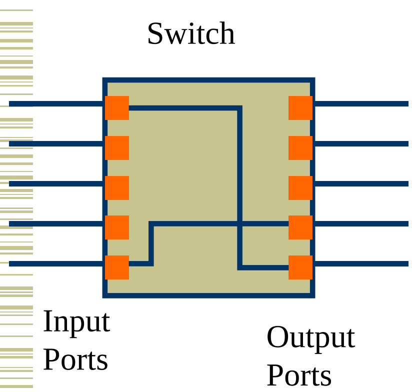

- Switch remembers how to forward data
  - No addresses!
- Many options
  - Between specific wires (circuit = wire)
  - Between timeslots (TDMA on each wire)
  - Between frequencies (FDMA on each wire)

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 set up when the VC is established
  - Can eue 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

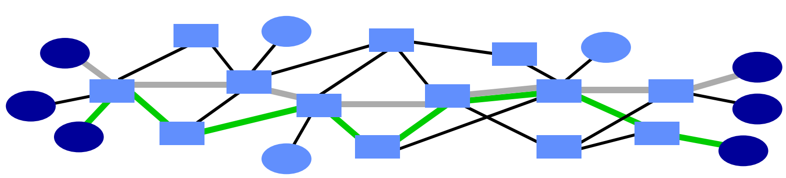

Virtual Circuits Versus Packet Switching

- Many similarities:
  - Forwarding based on “address” (VCID or dest address)
  - Statistical multiplexing for efficiency
  - Must have buffers space on switches
- Virtual circuit switching:
  - Uses short connection identifiers to forward packets
  - Switches know about the connections so they can more easily implement features such as quality of service
  - Switches are stateful: VC connection state cannot be lost
- Packet switching:
  - Use full destination addresses for forwarding packets
  - Can send data right away: no need to establish a connection first
  - Switches are stateless: easier to recover from failures
  - Adding QoS is hard

Virtual Circuit Forwarding

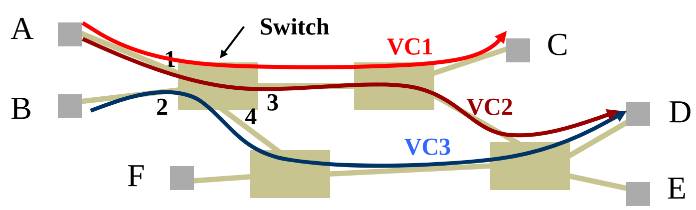

Address	Next Hop
VC1	3
VC2	3
VC3	4
VC4	?
VC5	?

- 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

8

VC versus Packets: Control over Path

Payload | VCI

Payload | Dst

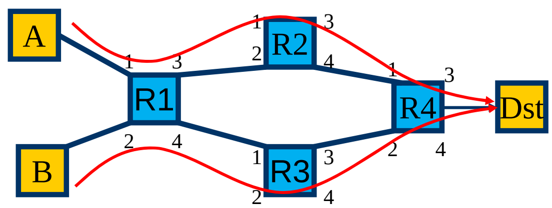

R1 packet forwarding table:

Dst R2

R1 VC table:

VC 1 R2

VC 2 R3

Different paths to same destination!

(useful for traffic engineering!)

How to Pick a VC Id?

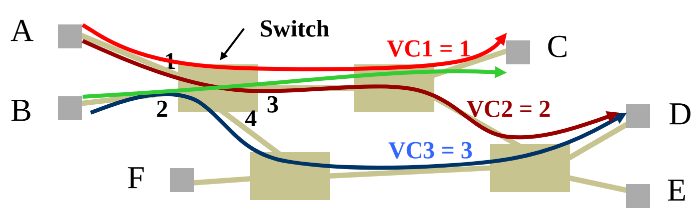

- When B establishes green virtual circuit, how does it know what VC ids are available?
- Even worse: every VC id may already be used on a link along the path to the destination
- Solution: VC id swapping

VC id Swapping

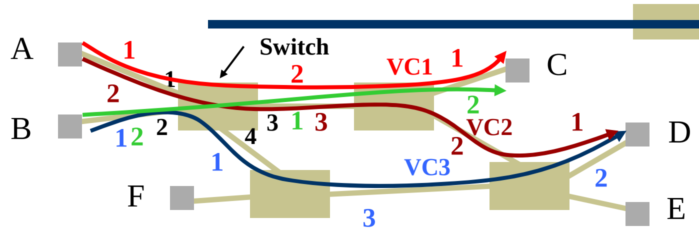

Address	Next Hop	Next id
VC1 = 1	3	2
VC2 = 2	3	3
VC3 = 1	4	1
VC4 = 2	3	1

- 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

11

Connections and Signaling

- Permanent vs. switched virtual connections (PVC/SVC)
  - 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

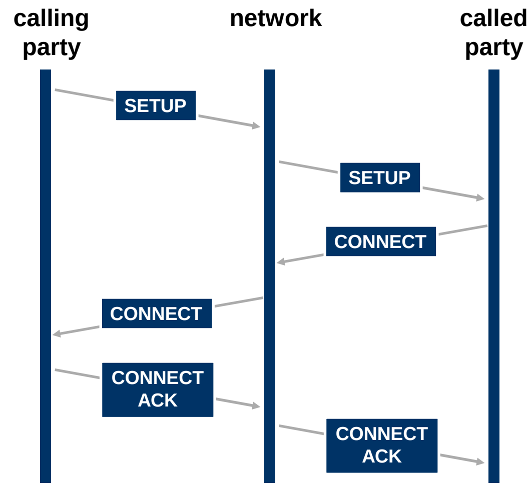

Virtual Circuits In Practice

- Asynchronous Transfer Mode - 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 many transit networks to provide traffic engineering, VPN support, 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.

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
- Virtual LANs
  - How do they differ?

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?
  - Voice only needs small “packets”
  - 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)

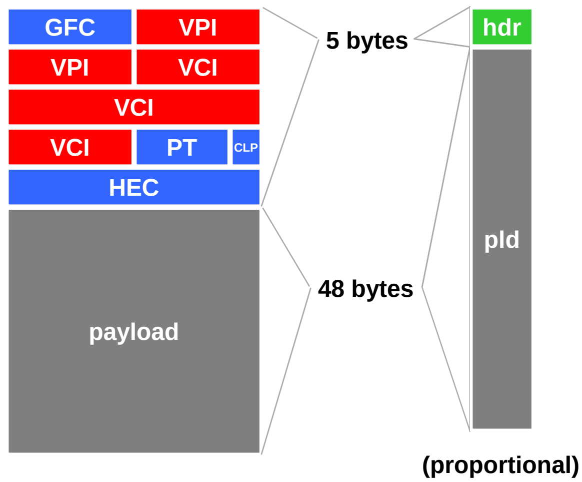

19

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

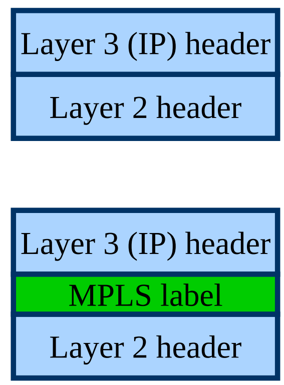

Some MPLS slides from H. Zhang

MPLS Vocabulary: LSP

- Label-switched path (LSP)
  - Simplex path through interior network
  - Forces packet along a specific path (set of routers)

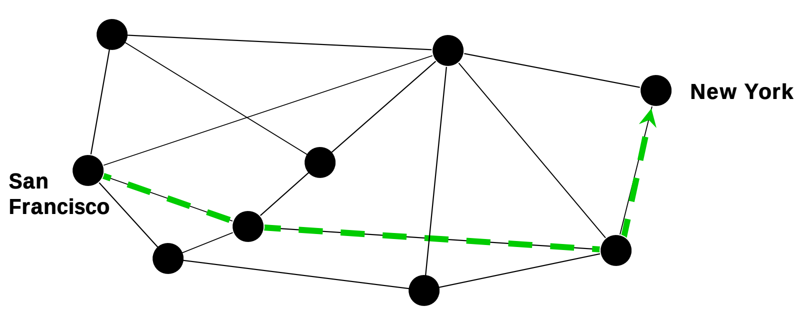

MPLS Vocabulary: LSR

- Label-switching router (LSR)
  - Performs LSP setup and MPLS packet forwarding
  - Label Edge Router (LER): LSP ingress or egress
  - Transit Router: swaps MPLS label, forwards packet

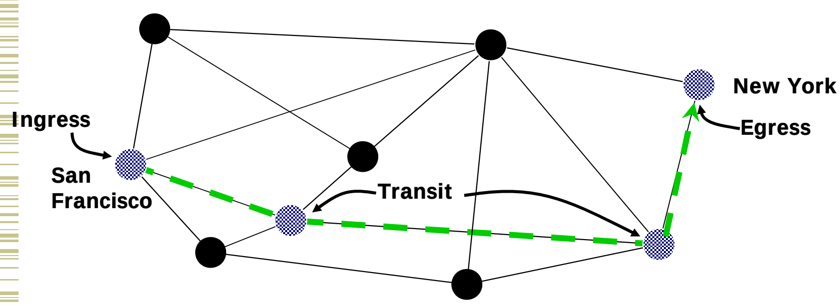

MPLS Header

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

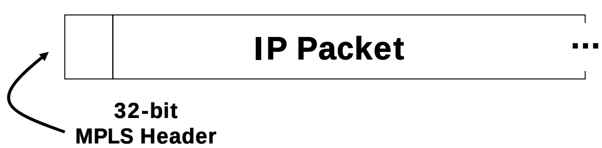

- IP packet is restored at end of LSP by egress router
  - TTL is adjusted, transit LSP routers count towards the TTL
  - MPLS is an optimization – does not affect IP semantics

MPLS Header

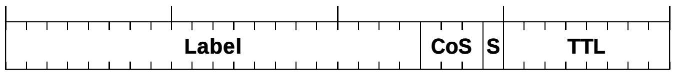

- Label
- Class of service
- Stacking bit
  - Remember me?
- Time to live
  - Decrement at each LSR, or
  - Pass through unchanged

Forwarding Equivalence Classes

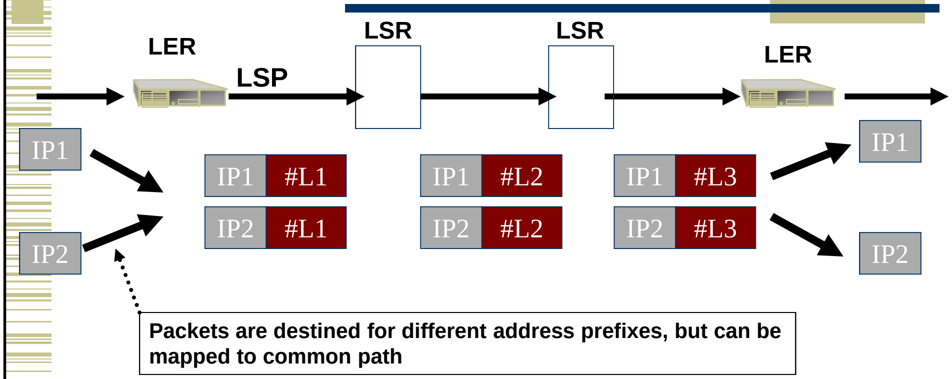

- FEC = "A subset of packets that are all treated the same way by a LSR"
- The concept of FECs provides for a great deal of flexibility and scalability
- Can be used to force flows of different "sizes" (e.g., Mbps) to follow certain paths through the network – more flexible than traditional routing

Establishing LSPs

- Use the Label Distribution Protocol (LDP) to establish paths based on IP forwarding tables
  - Simple
  - MPLS packets follow the same path as IP
- Explicitly route circuits
  - More work
  - Provides finer grain control over how traffic is distributed throughout the network
  - Important tool for traffic engineering

27

LSPs Driven by IP Routing

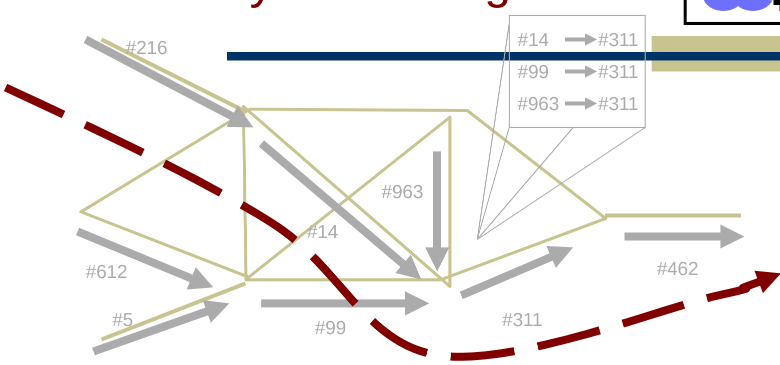

- A LSP is actually part of a tree from every source to that destination (unidirectional).
- A control protocol (e.g. Label Distribution Protocol, LDP) builds the tree based on the IP forwarding tables.

MPLS Builds on Standard IP

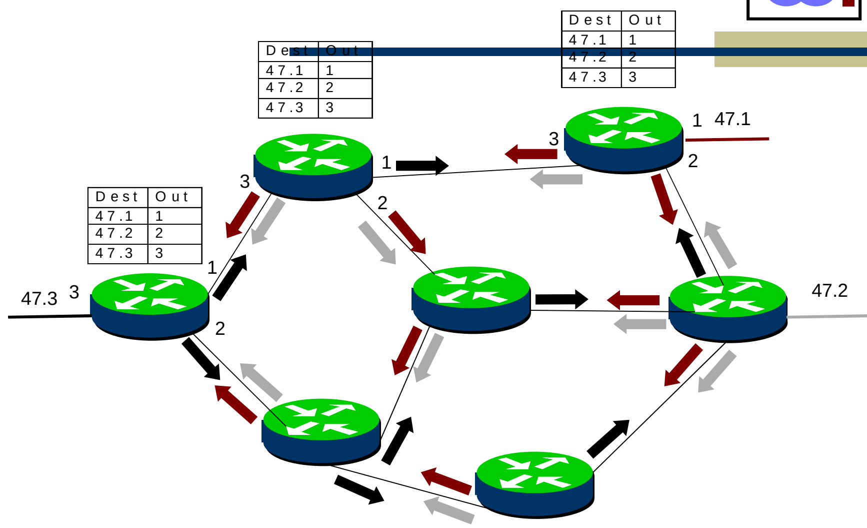

Dest	Out
47.1	1
47.2	2
47.3	3

Dest	Out
47.1	1
47.2	2
47.3	3

Dest	Out
47.1	1
47.2	2
47.3	3

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

Label Switched Path (LSP)

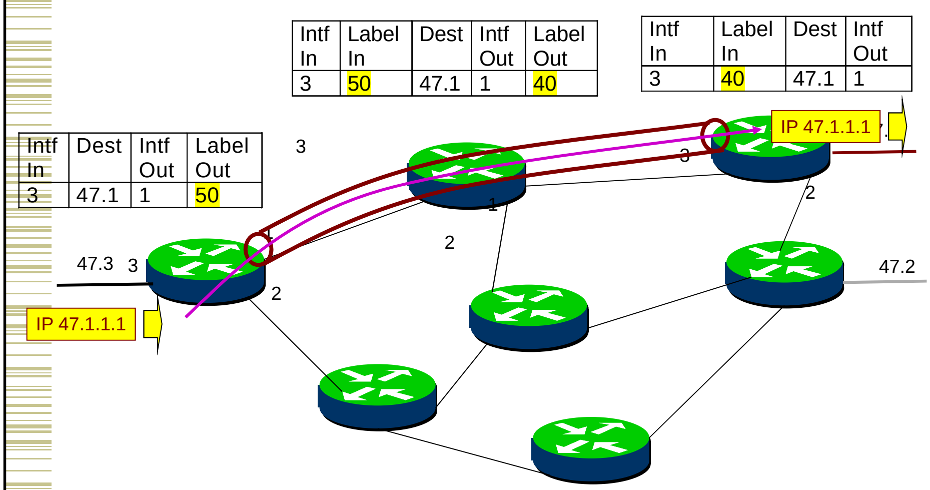

Intf In	Dest	Intf Out	Label Out
3	47.1	1	50

Intf In	Label In	Dest	Intf Out	Label Out
3	50	47.1	1	40

Intf In	Label In	Dest	Intf Out
3	40	47.1	1

Explicityly Routed - ER-LSP

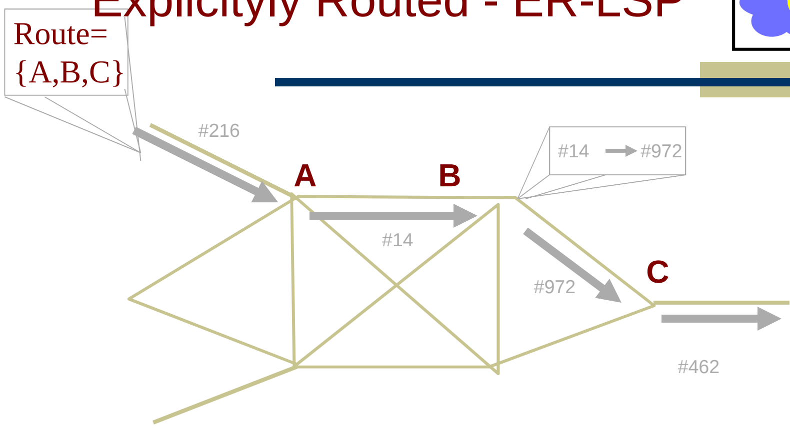

ER-LSP follows route that source chooses. In other words, the control message to establish the LSP (label request) is ***source routed***.

Explicitly Routed LSP - Example

Intf In	Dest	Intf Out	Label Out
3	47.1.1	2	33
3	47.1	1	50

Intf In	Label In	Dest	Intf Out	Label Out
3	50	47.1	1	40

Intf In	Label In	Dest	Intf Out
3	40	47.1	1

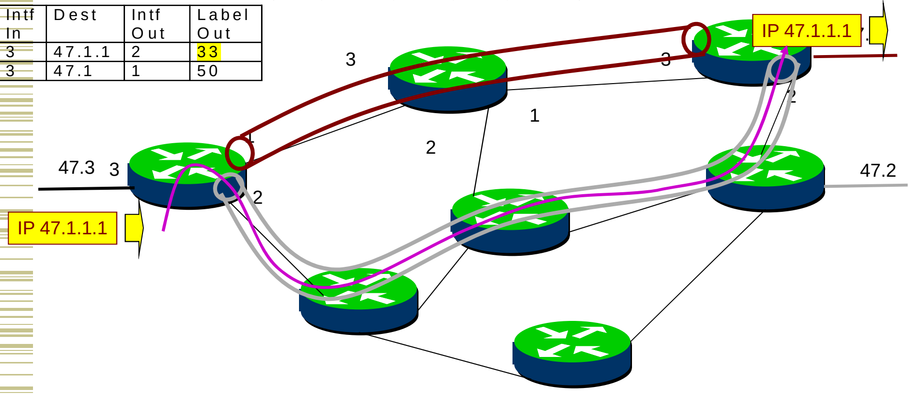

Outline

- 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
- Virtual LANs
  - How do they differ?

VLAN Introduction

- VLANs logically segment switched LANs (layer 2!) 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

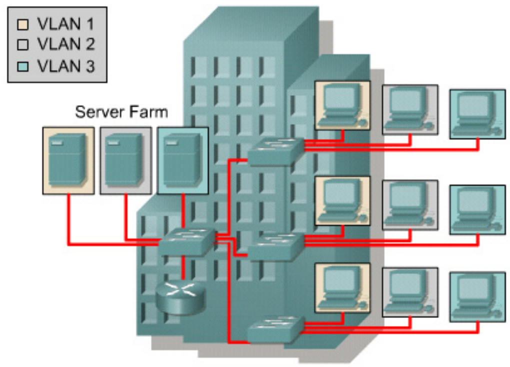

- A switch creates a broadcast domain
- VLANs help manage broadcast domains
- VLANs can be defined on port groups, users, or protocols
- LAN switches and network management software provide a mechanism to create VLANs

35

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 destination device

Example: 802.1Q Standard for VLANs over Ethernet

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

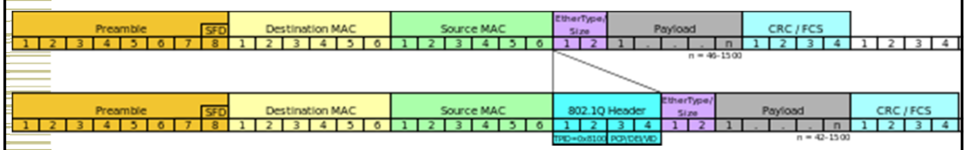

- Header consists of
  - Tag Protocol Identifier (16b): single value that marks frame as a VLAN frame
  - Control bits (4b): mostly priority
  - VLAN Identifier (12b): identifies VLAN

37

Take Home Points

- Costs/benefits/goals of virtual circuits
- Cell switching (ATM)
  - Fixed-size pkts: Fast hardware – limited benefit
  - Packet size picked for low voice jitter. Understand trade-offs.
  - Complex technology – not a good fit for IP
- 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