Master Course
Computer Networks
IN2097

Prof. Dr.-Ing. Georg Carle
Christian Grothoff, Ph.D.

Chair for Network Architectures and Services
Institut für Informatik
Technische Universität München
http://www.net.in.tum.de

Outline

- Recall last lecture
- Project status
- Internet Strucuture
- Network virtualisation

Internet-Shutdown

- http://www.turnofftheinternet.com/

Internet-Shutdown

- http://www.turnofftheinternet.com/shutdown.html

You have now safely shutdown the Internet.

turnofftheinternet.com

Registrant:
Cockos Incorporated
245 8th Street
San Francisco, CA 94103

Registrar: DOTSTER
Created on: 10-SEP-00
Expires on: 10-SEP-13
Last Updated on: 18-JUN-10

Administrative, Technical Contact:
Frankel, Justin
Cockos Incorporated
245 8th Street
San Francisco, CA 94103, US
707-726-2567

Domain servers in listed order:
NS1.COCKOS.COM
NS2.COCKOS.COM

- **About Cockos:** We lovingly craft the software that we would want to use.
  Past Experience: The employees of Cockos have created and launched many fine products and technologies prior to joining Cockos, including Gnutella, K-Meleon, Kaillera, NSIS, SHOUTcast, and Winamp.
- **Justin Frankel** (born 1978) is an American computer programmer best known for his work on the Winamp media player application and for inventing the gnutella peer-to-peer network. He's also the founder of Cockos Incorporated which creates music production and development software such as the REAPER digital audio workstation, the NINJAM **collaborative music tool** and the Jesusonic expandable effects processor. In 2002, he was named to the MIT Technology Review TR100 as one of the top 100 innovators in the world under the age of 35.[1] source: wikipedia
- **Christophe Thibault**
- **John Schwartz**

Internet Emergency Control by Whitehouse

- http://news.cnet.com/8301-13578_3-10320096-38.html

  Bill would give president emergency control of Internet
  28 August 2009
- http://www.politechbot.com/docs/rockefeller.revised.cybersecurity.draft.082709.pdf

- in the event of an immediate threat to strategic national interests involving compromised Federal Government or United States critical infrastructure information system or network

  (A) may declare a cybersecurity emergency; and

  (B) may, if the President finds it necessary for the national defense and security, and **in coordination with relevant industry sectors, direct the national response to the cyber threat and the timely restoration of the affected critical infrastructure information system or network**;

Electronic Frontier Foundation, 28 August 2009

- "The language has changed but it doesn't contain any real additional limits," EFF's Tien says. "It simply switches the more direct and obvious language they had originally to the more ambiguous (version)...The designation of what is a critical infrastructure system or network as far as I can tell has no specific process. There's no provision for any administrative process or review. That's where the problems seem to start. And then you have the amorphous powers that go along with it."
- Translation: If your company is deemed "critical," a new set of regulations kick in involving who you can hire, what information you must disclose, and when the government would exercise control over your computers or network.

Jena Longo, deputy communications director for the Senate Commerce Committee

The president of the United States has always had the constitutional authority, and duty, to protect the American people and direct the national response to any emergency that threatens the security and safety of the United States. The Rockefeller-Snowe Cybersecurity bill makes it clear that the president's authority includes securing our national cyber infrastructure from attack. The section of the bill that addresses this issue, applies specifically to the national response to a severe attack or natural disaster. This particular legislative language is based on longstanding statutory authorities for wartime use of communications networks. To be very clear, the Rockefeller-Snowe bill will not empower a "government shutdown or takeover of the Internet" and any suggestion otherwise is misleading and false. The purpose of this language is to clarify how the president directs the public-private response to a crisis, secure our economy and safeguard our financial networks, protect the American people, their privacy and civil liberties, and coordinate the government's response.

Project: VMs

Project Infrastructure

- 32 Virtual Machines created, hosted and made accessible
- Instructions available in svn under pub/README-vm.txt
- How to access your virtual machine

  You can log into your virtual machine using the OpenSSH private key named "id_rsa" in your team's SVN directory.
  If you're using OpenSSH type something like:

  ssh -i /path/to/mccn/team99/id_rsa root@mccn99.net.in.tum.de

  You have to replace "99" by your team number, obviously.
  If you prefer a different SSH implementation refer to its documentation on how to use OpenSSH private keys.
- How to install software on your virtual machine

  man 8 apt-get
- DNS

  The subdomain "t.mccn01.net.in.tum.de" is delegated to "mccn01.net.in.tum.de". The same goes for the other hosts.

Network Architectures

Link virtualization: ATM, MPLS

Asynchronous Transfer Mode: ATM

- 1990's/00 standard for high-speed networking
  - 155Mbps to 622 Mbps and higher
  - *Broadband Integrated Service Digital Network* architecture

- Goal: *integrated, end-end transport of carry voice, video, data*
  - meeting timing/QoS requirements of voice, video versus Internet best-effort model
  - "next generation" telephony: technical roots in telephone world
  - packet-switching (fixed length packets, called "cells") using virtual circuits, and label swapping

Datagram or VC network: why?

Internet

- data exchange among computers
  - “elastic” service, no strict timing requirements
- “smart” end systems (computers)
  - can adapt, perform control, error recovery
  - simple inside network, complexity at “edge”
- many link types
  - different characteristics
  - uniform service difficult

ATM

- evolved from telephony
- human conversation:
  - strict timing, reliability requirements
  - need for guaranteed service
- “dumb” end systems
  - telephones
  - complexity inside network

ATM architecture

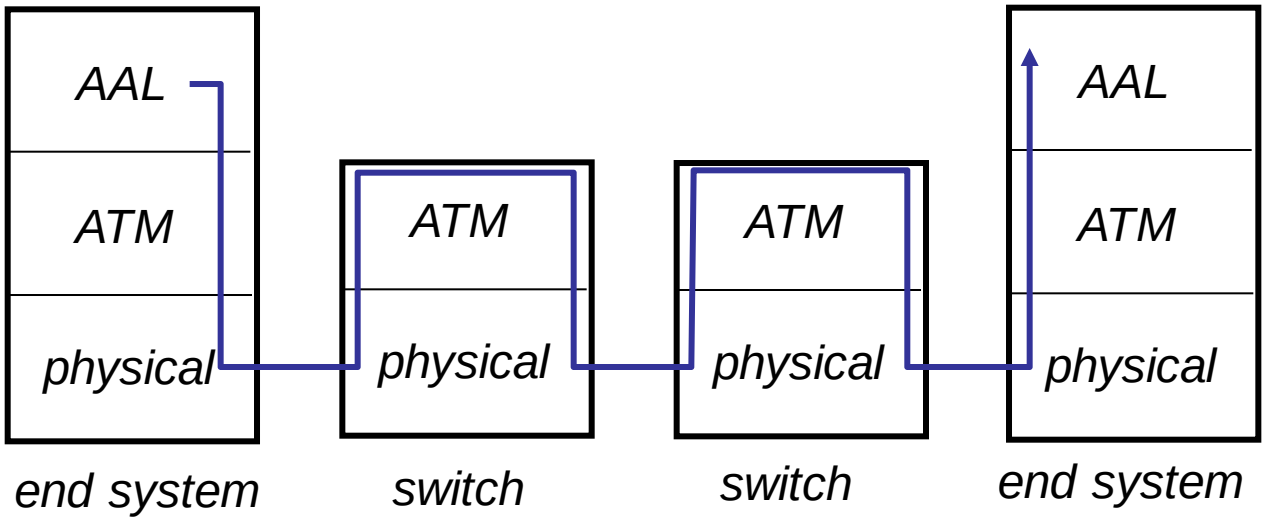

- adaptation layer: only at edge of ATM network
  - data segmentation/reassembly
  - roughly analogous to Internet transport layer
- ATM layer: "network" layer
  - cell switching, routing
- physical layer

ATM: Network or Link layer?

Vision: end-to-end transport:
"ATM from desktop to desktop"

- ATM *is* a network technology

Reality: used to connect IP backbone routers

- "IP over ATM"
- ATM as switched link layer, connecting IP routers

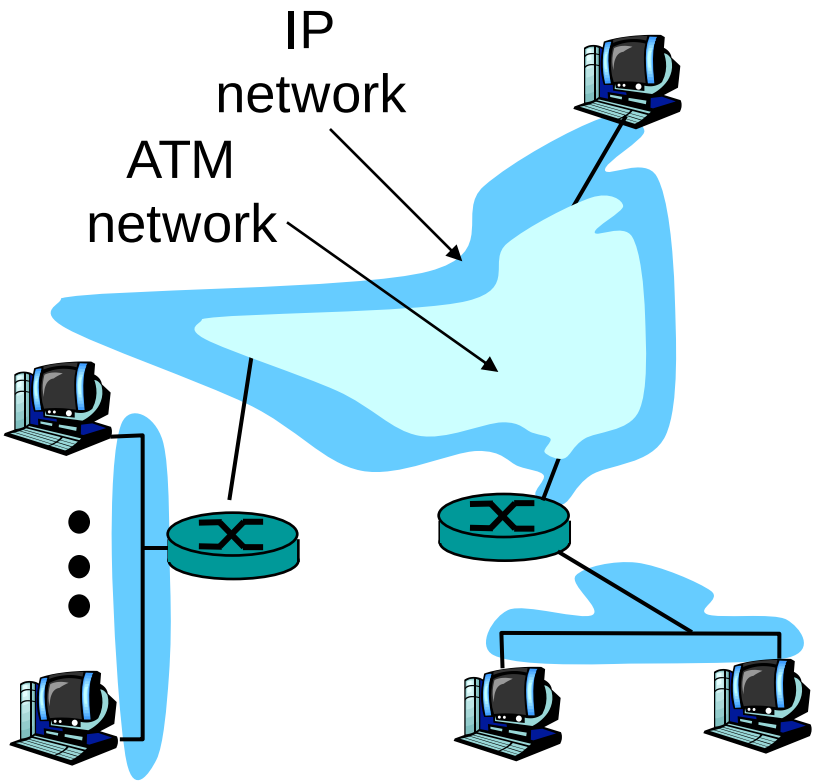

ATM Adaptation Layer (AAL)

- ATM **Adaptation Layer** (AAL): “adapts” upper layers (IP or native ATM applications) to ATM layer below
- AAL present **only in end systems**, not in switches
- AAL layer segment (header/trailer fields, data) fragmented across multiple ATM cells
  - analogy: TCP segment in many IP packets

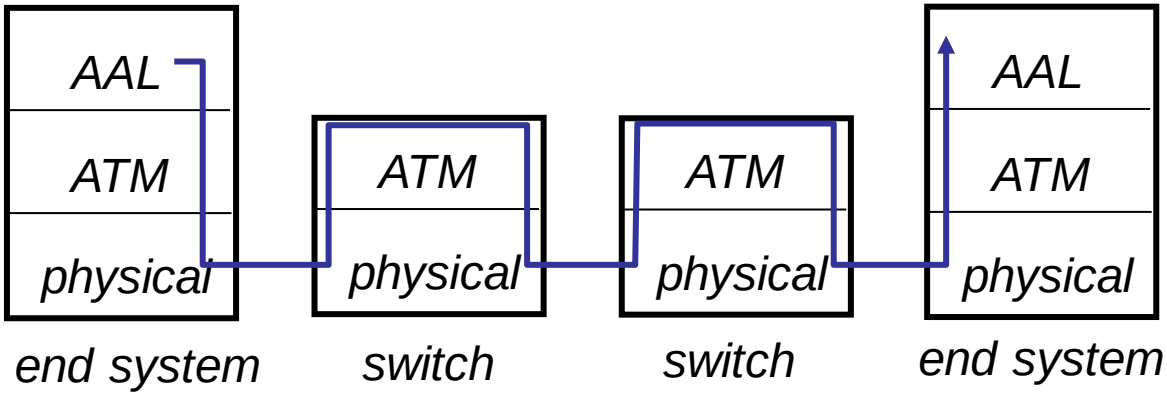

ATM Adaptation Layer (AAL) [more]

Different versions of AAL layers, depending on ATM service class:

- AAL1: for CBR (Constant Bit Rate) services, e.g. circuit emulation
- AAL2: for VBR (Variable Bit Rate) services, e.g., MPEG video
- AAL5: for data (e.g., IP datagrams)

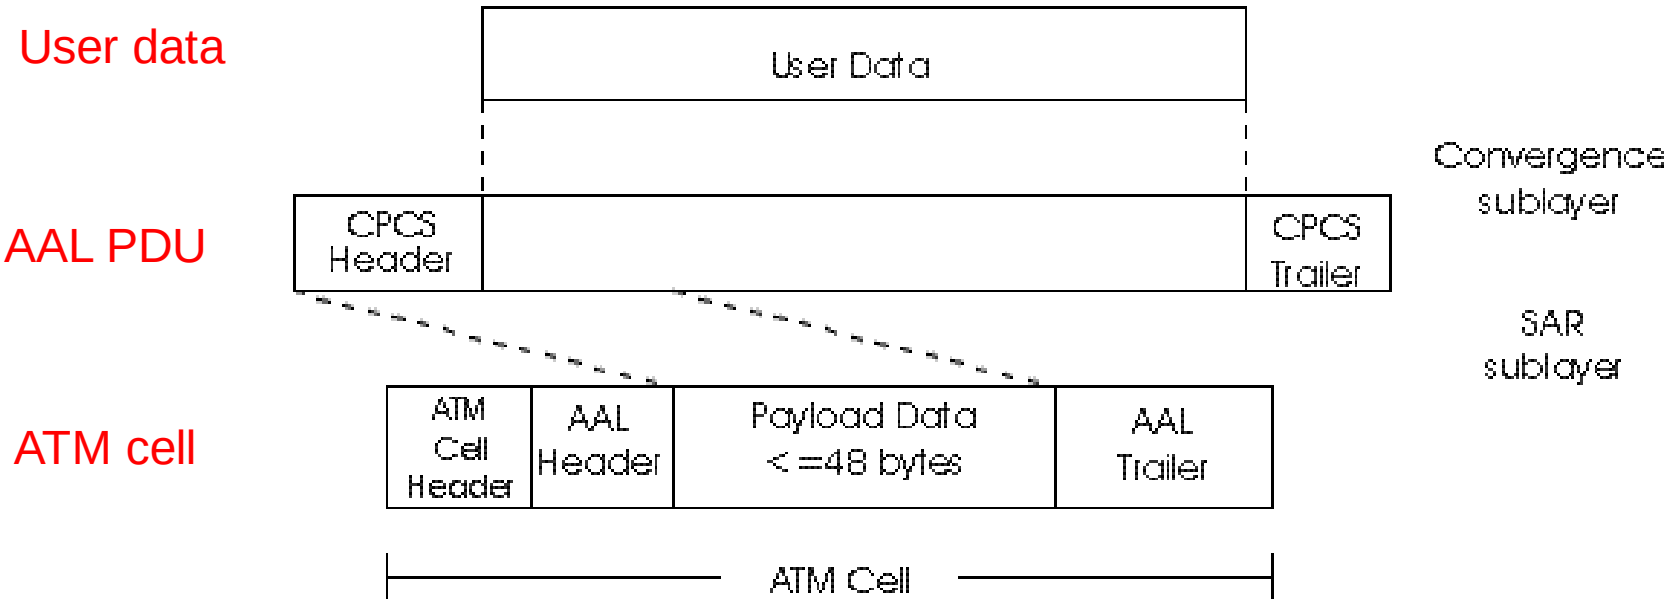

ATM Layer

Service: transport cells across ATM network

- analogous to IP network layer
- very different services than IP network layer
- possible Quality of Service (QoS) Guarantees

Network Architecture	Service Model	Guarantees ?				Congestion feedback
		Bandwidth	Loss	Order	Timing	
Internet	best effort	none	no	no	no	no (inferred via loss)
ATM	CBR	constant rate	yes	yes	yes	no congestion
ATM	VBR	guaranteed rate	yes	yes	yes	no congestion
ATM	ABR	guaranteed minimum	no	yes	no	yes
ATM	UBR	none	no	yes	no	no

ATM Layer: Virtual Circuits

- VC transport: cells carried on VC from source to destination
  - call setup, teardown for each call *before* data can flow
  - addressing of destination e.g. by E.164 number
  - each packet carries VC identifier (*not* destination ID)
  - label swapping: VC identifier may change along path
  - *every* switch on source-destination path maintains "state" for each passing connection
  - link, switch resources (bandwidth, buffers) may be *allocated* to VC: to get circuit-like perf.
- Permanent VCs (PVCs)
  - long lasting connections
  - typically: "permanent" route between to IP routers
- Switched VCs (SVC):
  - dynamically set up on per-call basis

ATM VCs

- Advantages of ATM VC approach:
  - QoS performance guarantee for connection mapped to VC (bandwidth, delay, delay jitter)
- Drawbacks of ATM VC approach:
  - Inefficient support of datagram traffic
  - one PVC between each source/destination pair does not scale
  - SVC introduces call setup latency, processing overhead for short lived connections

ATM Layer: ATM cell

- 5-byte ATM cell header
- 48-byte payload (Why?)
  - small payload ⇒ short cell-creation delay for digitized voice
  - halfway between 32 and 64 (compromise!)

Cell header

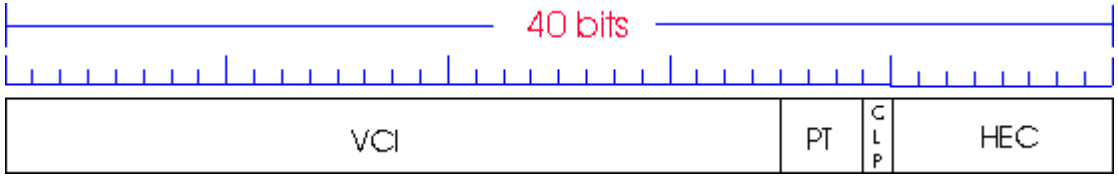

Cell format

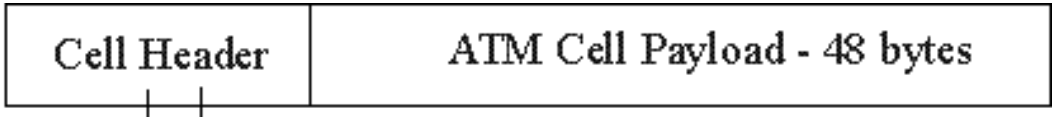

ATM cell header

- **VCI:** virtual channel ID
  - may *change* from link to link through network
- **PT:** Payload type: RM (resource management) vs. data cell
- **CLP:** Cell Loss Priority bit
  - CLP = 1 implies low priority cell, can be discarded if congestion
- **HEC:** Header Error Checksum
  - cyclic redundancy check

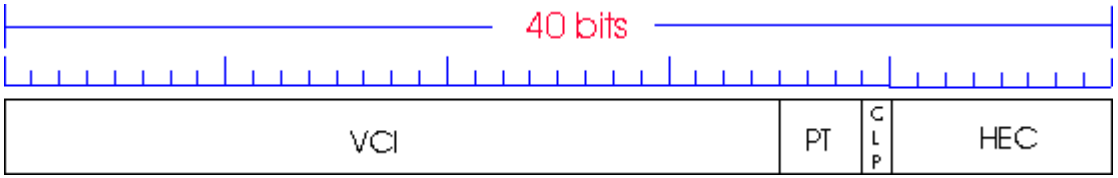

Virtual Circuit Switching

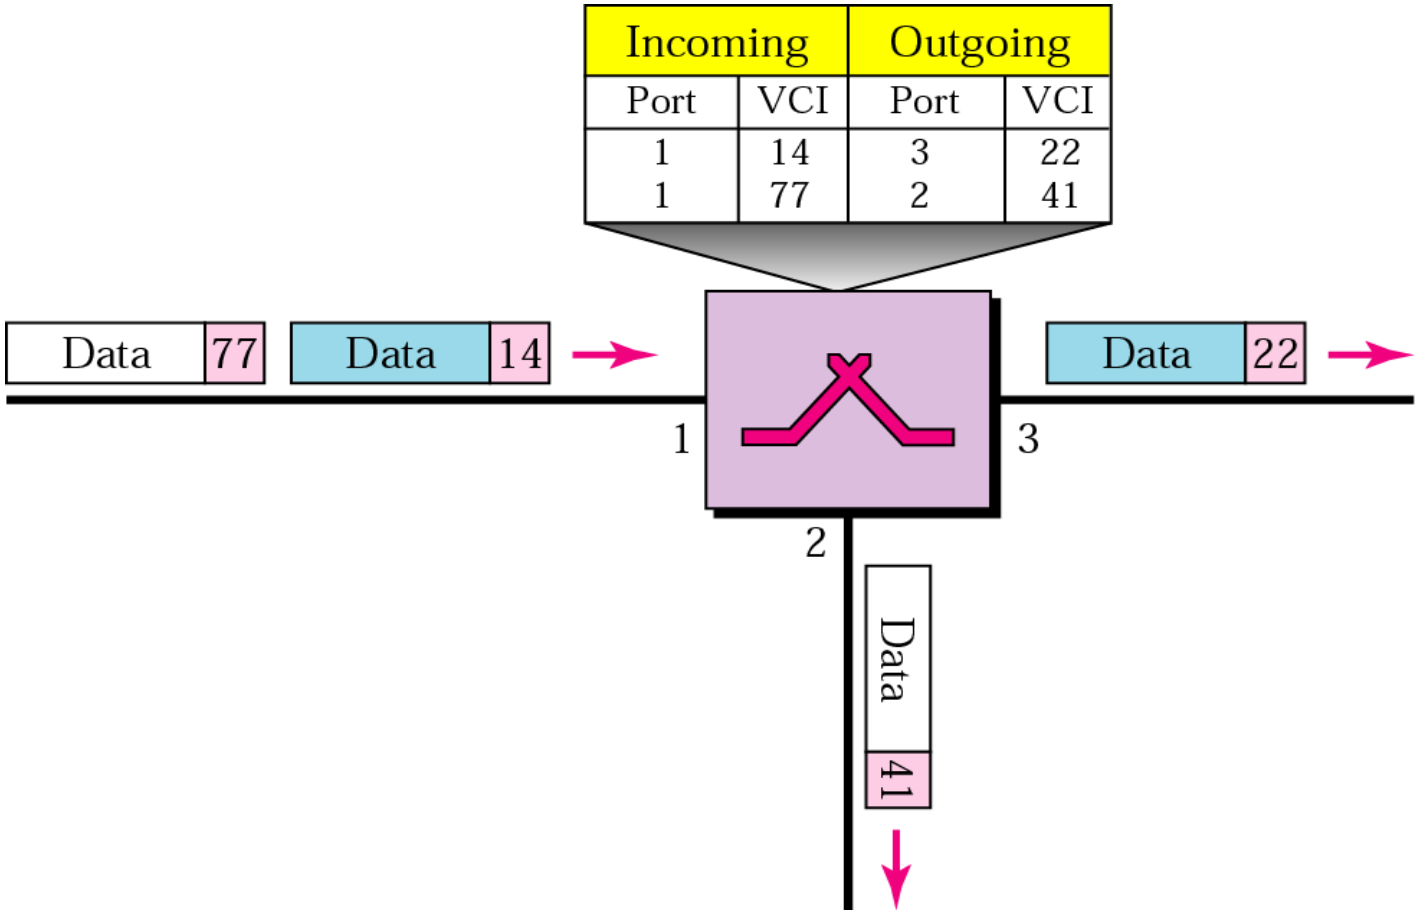

Multiplexing of Variable vs. Fixed Size Packets

- Multiplexing of variable size packets

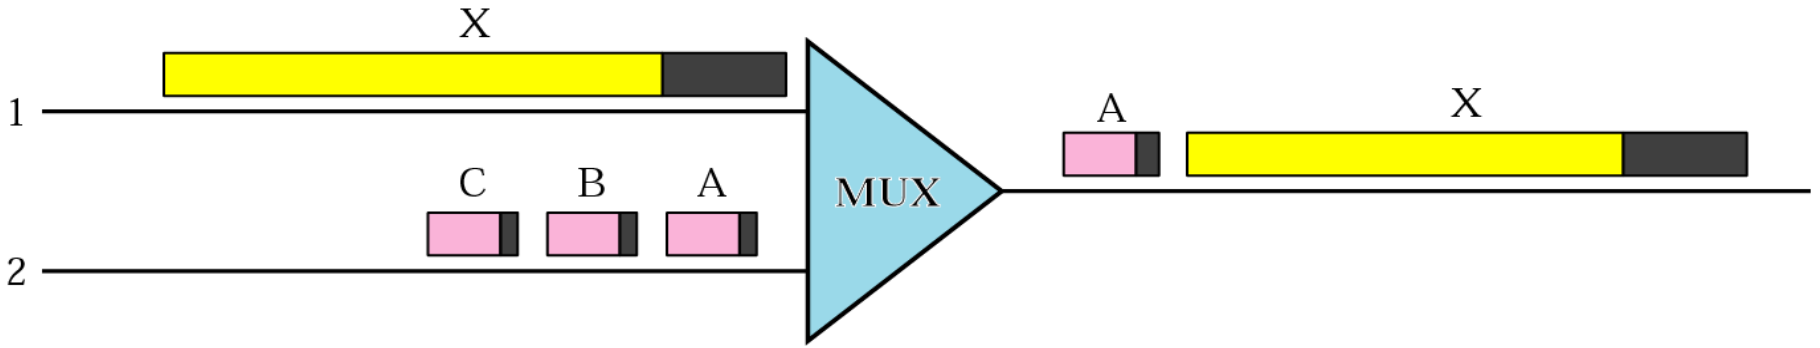

- ATM Multiplexing

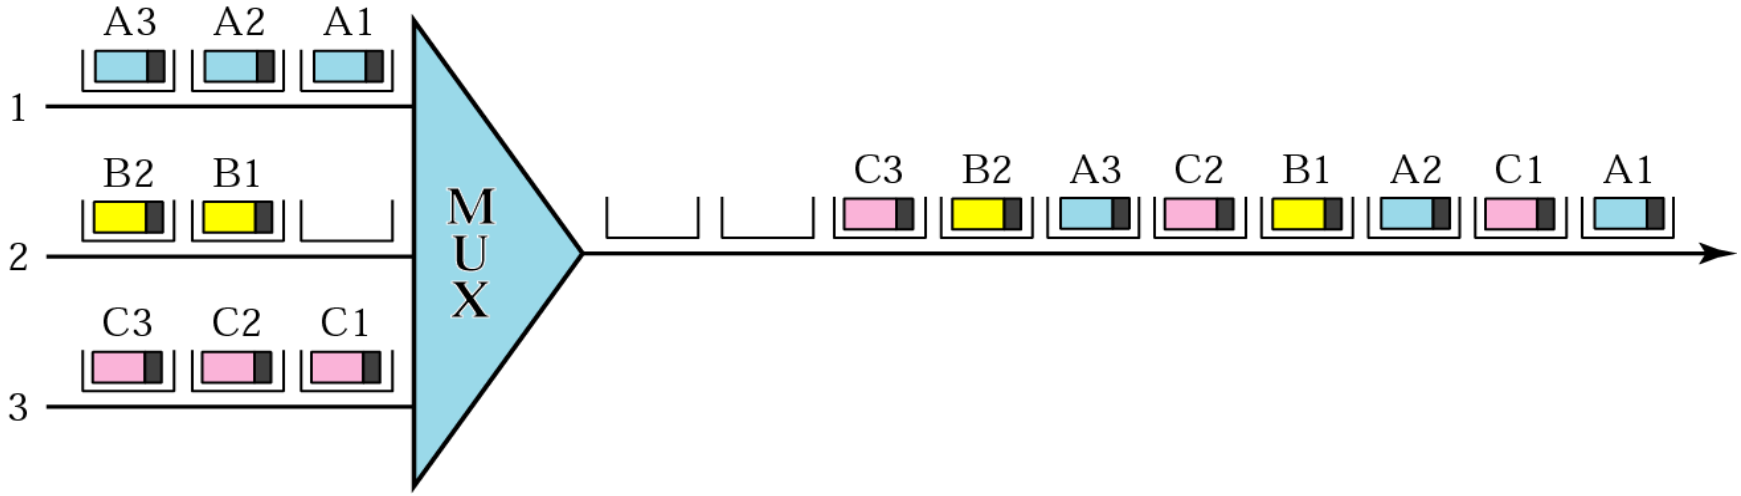

ATM Identifiers

- ATM Cell

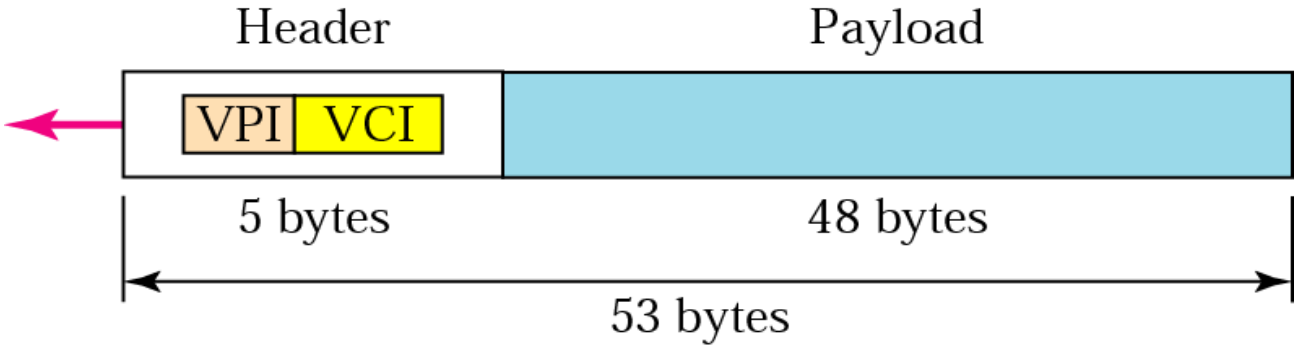

- Virtual Path Identifiers and Virtual Channel Identifiers

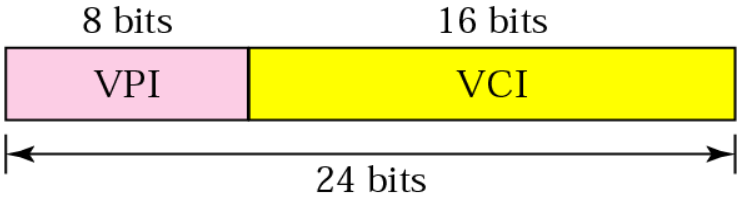

a. VPI and VCI in a UNI

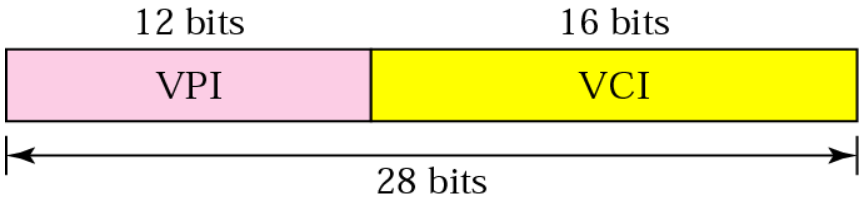

b. VPI and VCI in an NNI

(UNI: User-to-Network-Interface
NNI: Network-to-Network-Interface)

ATM Virtual Connections

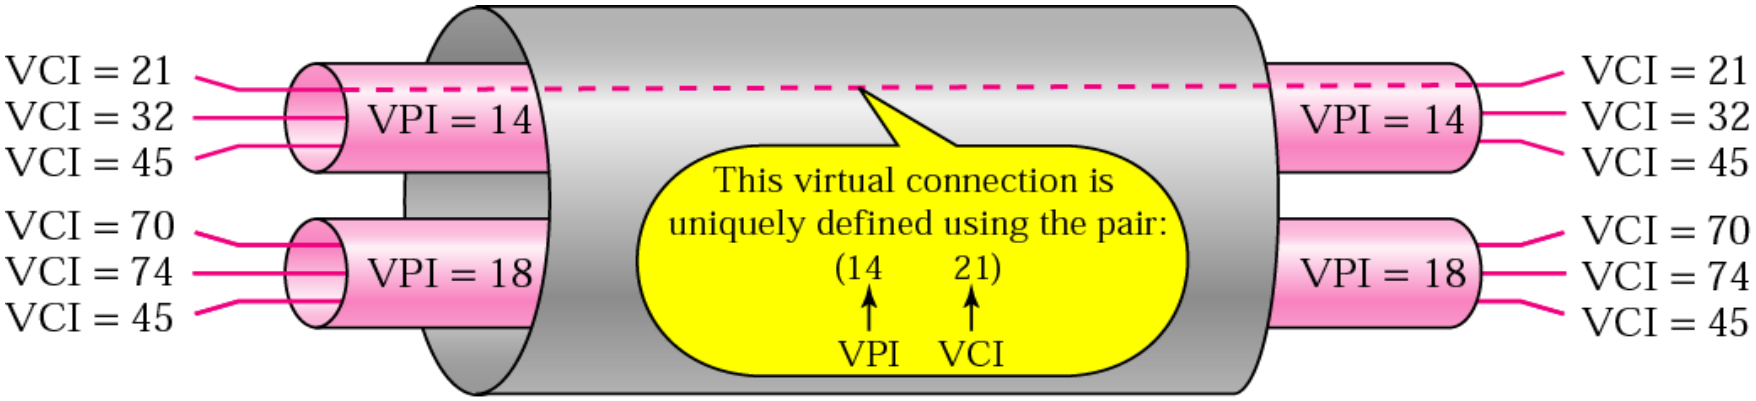

ATM Physical Layer

Physical Medium Dependent (PMD) sublayer

- **SONET/SDH**: transmission frame structure (like a container carrying bits);
  - bit synchronization;
  - bandwidth partitions (TDM);
  - several speeds:
  - OC3 = 155.52 Mbps
  - OC12 = 622.08 Mbps
  - OC48 = 2.45 Gbps
  - OC192 = 9.6 Gbps
- **TI/T3**: transmission frame structure (old telephone hierarchy): 1.5 Mbps/ 45 Mbps
- **unstructured**: just cells (busy/idle)
  - transmission of **idle cells** when no data cells to send

IP-Over-ATM

Classic IP only

- 3 “networks” (e.g., LAN segments)
- MAC (802.3) and IP addresses

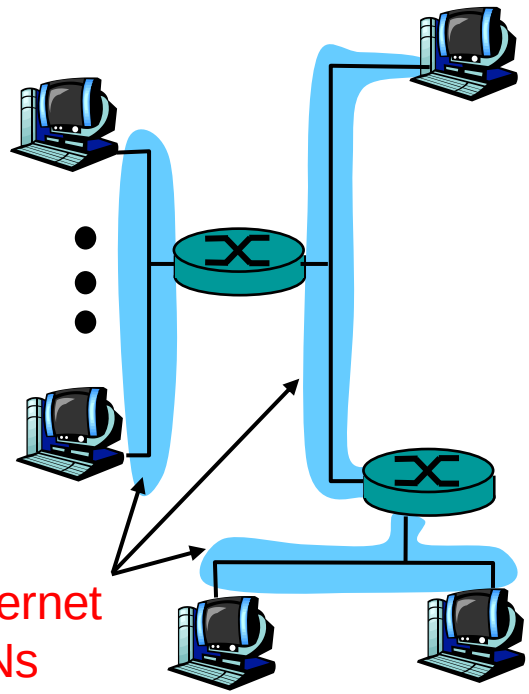

IP over ATM

- replace “network” (e.g., LAN segment) with ATM network
- ATM addresses, IP addresses

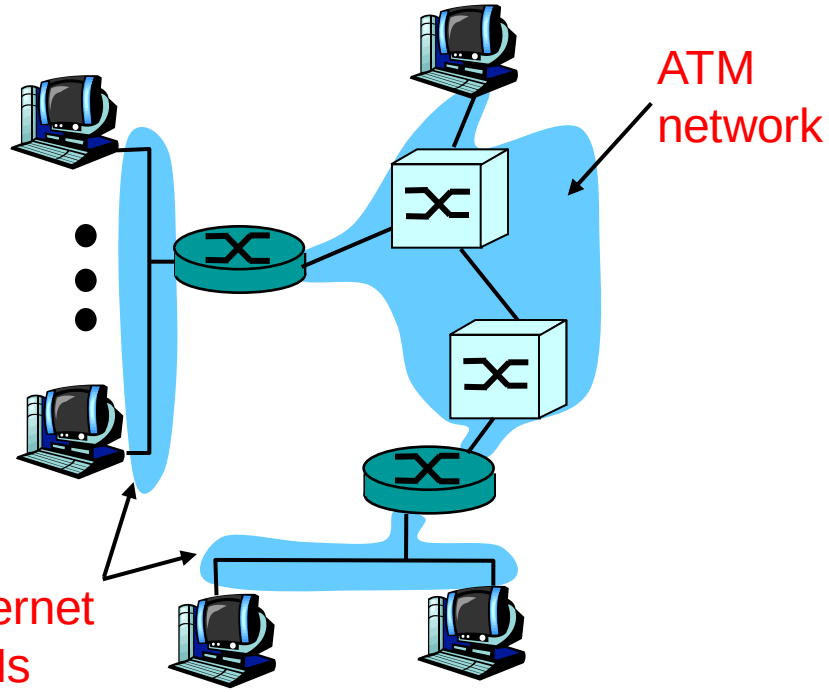

IP-Over-ATM

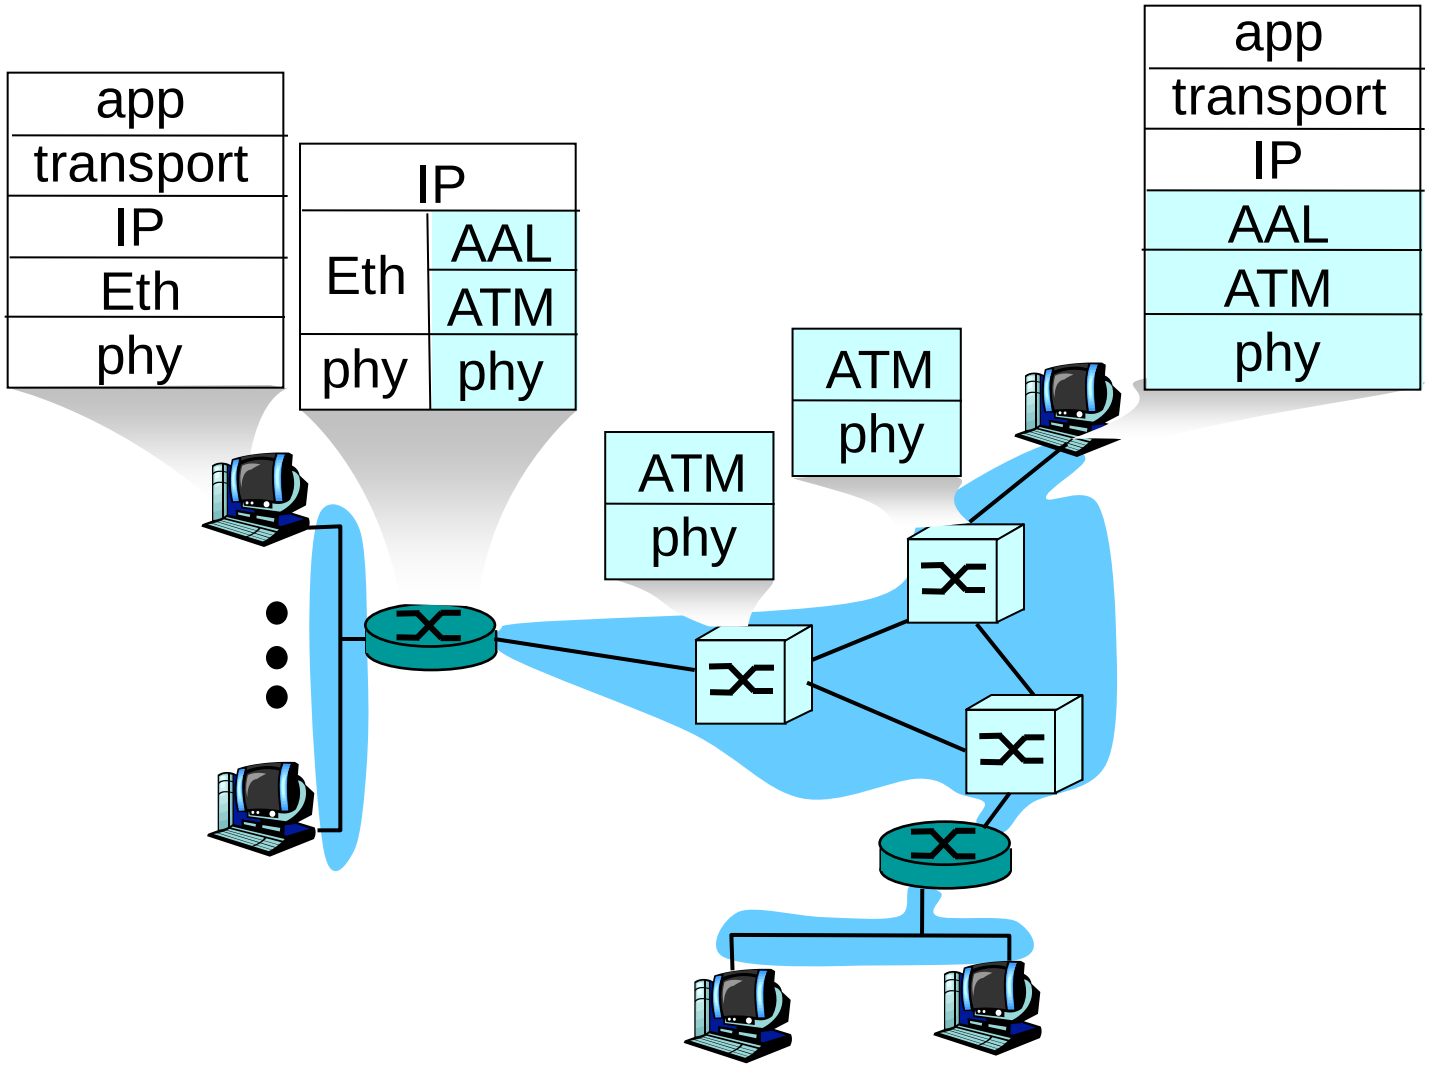
app
transport
IP
Eth
phy
IP
Eth
AAL
ATM
phy
phy
ATM
phy
ATM
phy
app
transport
IP
AAL
ATM
phy

Datagram Journey in IP-over-ATM Network

- at Source Host:
  - IP layer maps between IP, ATM destination address (using ARP)
  - passes datagram to AAL5
  - AAL5 encapsulates data, segments cells, passes to ATM layer
- ATM network: moves cell along VC to destination
- at Destination Host:
  - AAL5 reassembles cells into original datagram
  - if CRC OK, datagram is passed to IP

IP-Over-ATM

Issues:

- IP datagrams into ATM AAL5 PDUs
- from IP addresses to ATM addresses
  - just like IP addresses to 802.3 MAC addresses!
  - ARP server

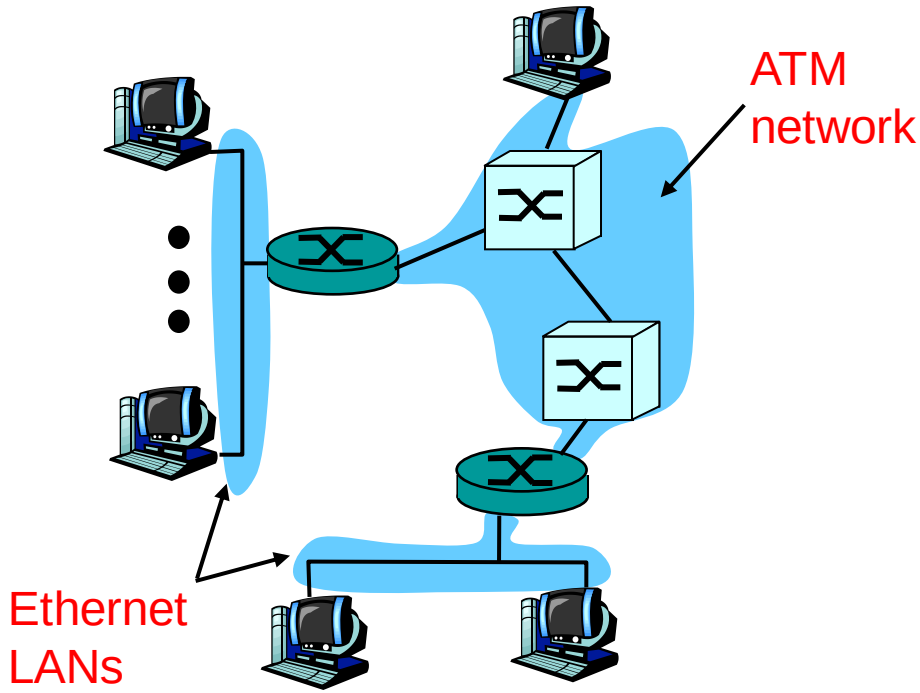

Multiprotocol label switching (MPLS)

- initial goal: speed up IP forwarding by using fixed length label (instead of IP address) to do forwarding
  - borrowing ideas from Virtual Circuit (VC) approach
  - but IP datagram still keeps IP address!

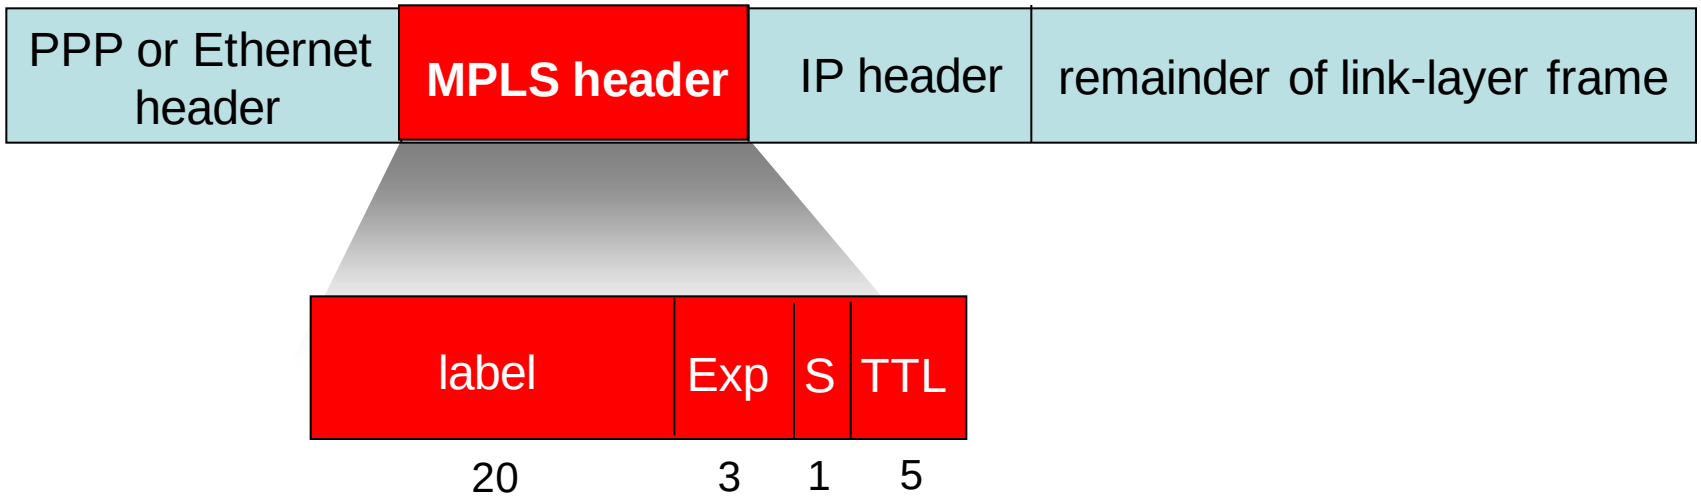

MPLS capable routers

- a.k.a. label-switched router
- forwards packets to outgoing interface based only on label value (don’t inspect IP address)
  - MPLS forwarding table distinct from IP forwarding tables
- signaling protocol needed to set up forwarding
  - Label Distribution Protocol (LDP)
  - RSVP-TE
- forwarding possible along paths that IP alone would not allow (e.g., source-specific routing)
- MPLS supports traffic engineering
- must co-exist with IP-only routers

MPLS forwarding tables

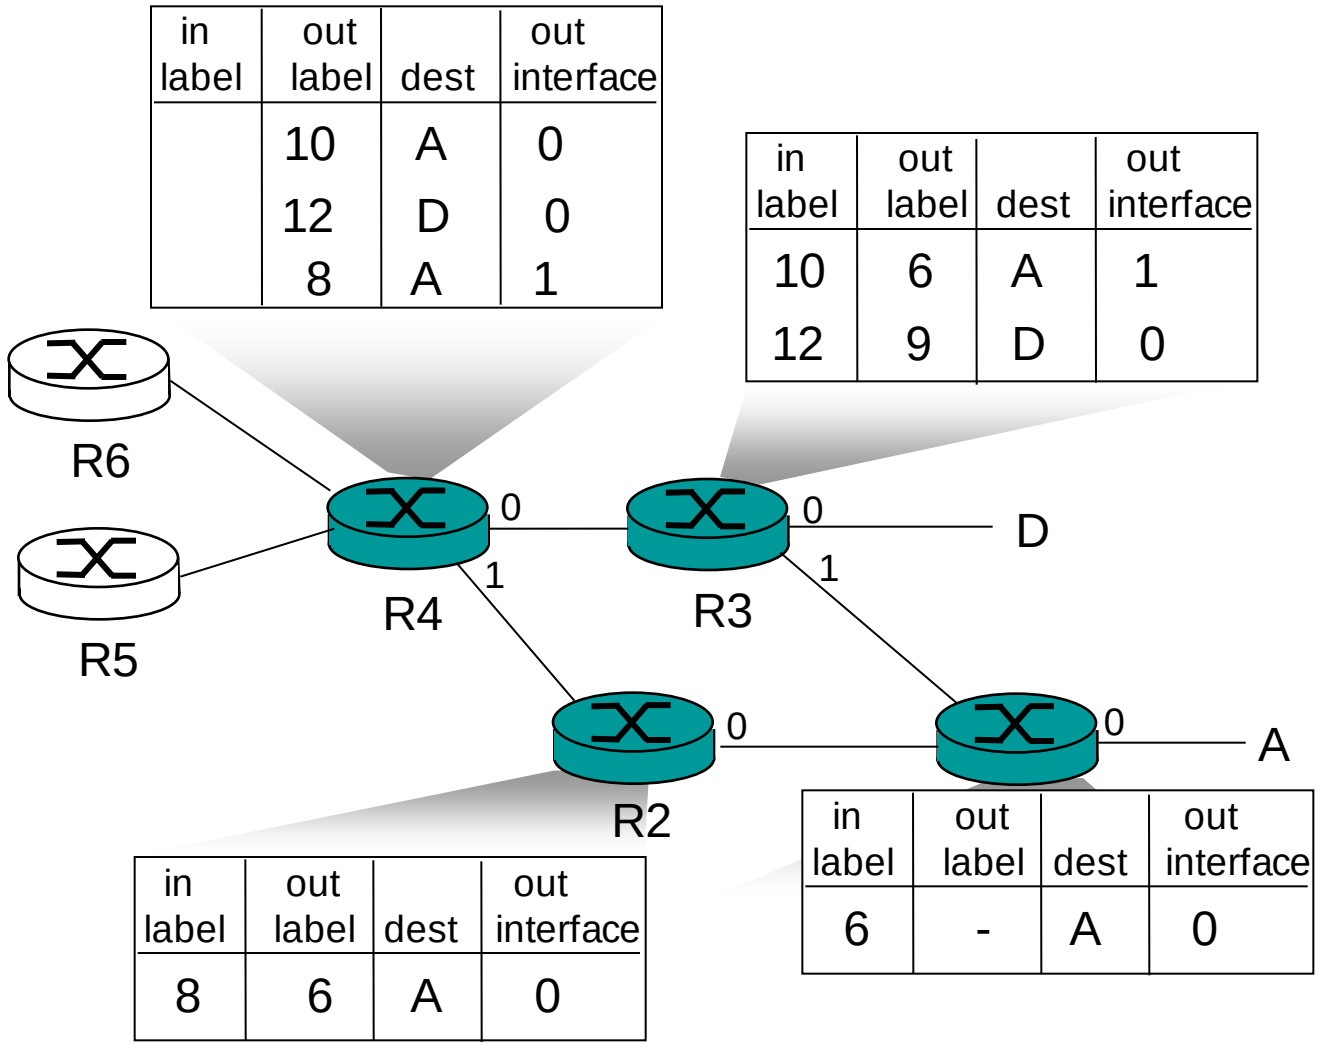

Virtual Private Networks

Virtual Private Networks (VPN)

VPNs

Networks perceived as being private networks by customers using them, but built over shared infrastructure owned by service provider (SP)

- Service provider infrastructure:
  - backbone
  - provider edge devices
- Customer:
  - customer edge devices (communicating over shared backbone)

VPN Reference Architecture

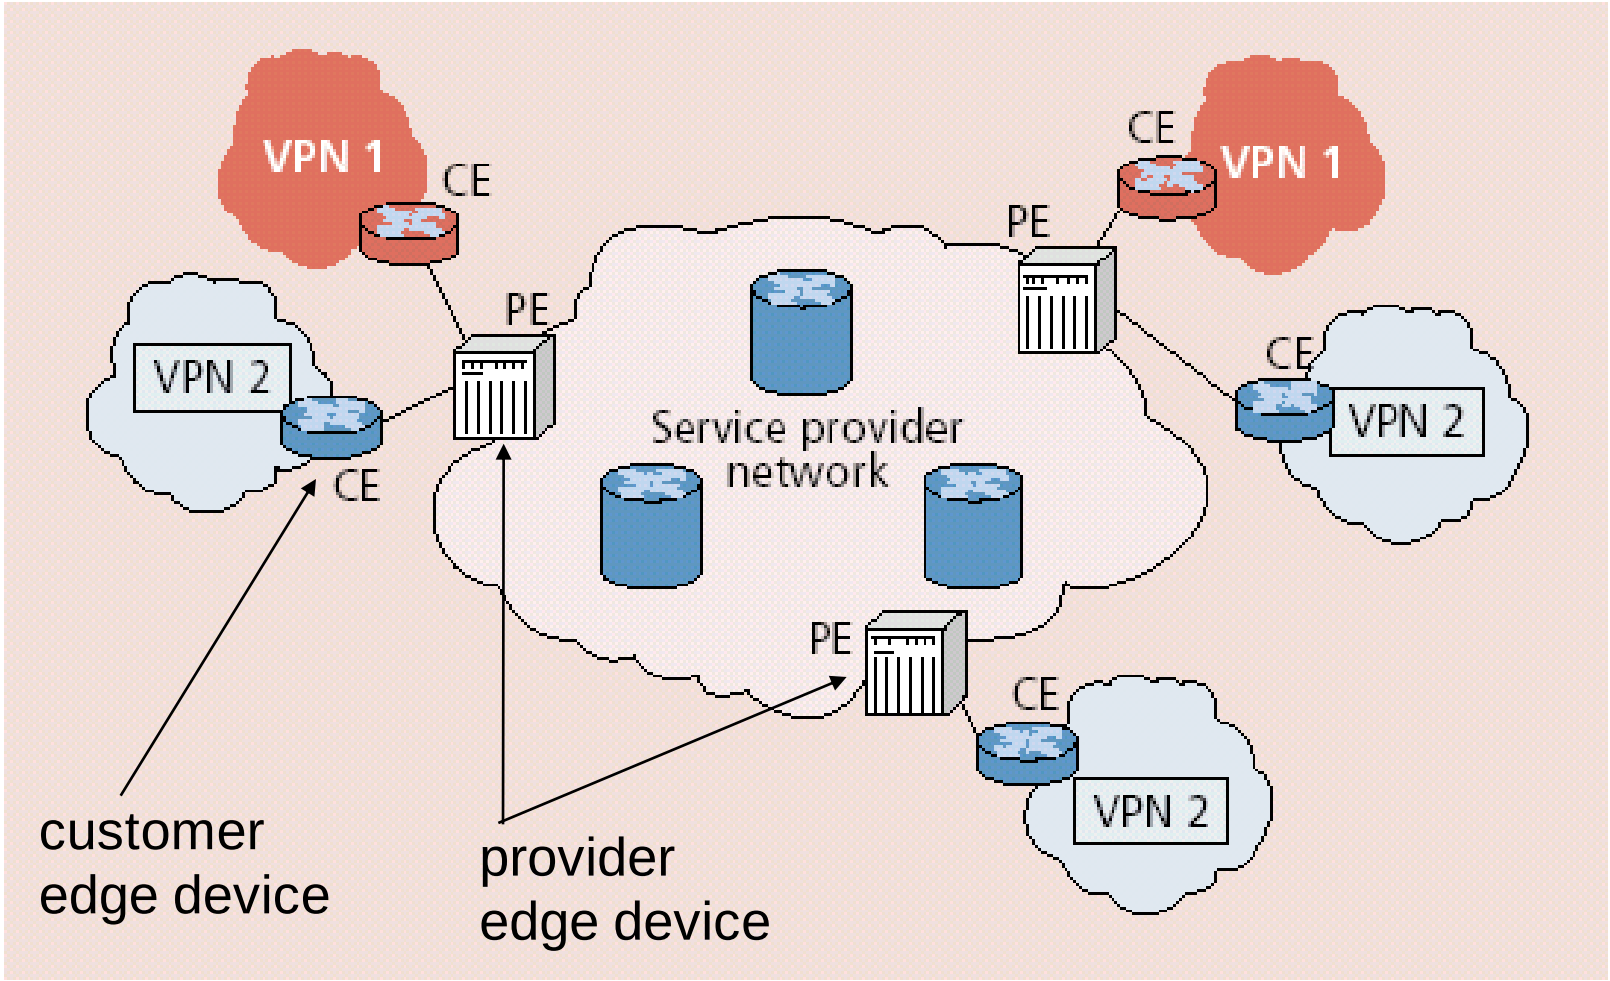

VPNs: Why?

- Privacy
- Security
- Works well with mobility (looks like you are always at home)
- Cost
  - many forms of newer VPNs are cheaper than leased line VPNs
  - ability to share at lower layers even though logically separate means lower cost
  - exploit multiple paths, redundancy, fault-recovery in lower layers
  - need isolation mechanisms to ensure resources shared appropriately
- Abstraction and manageability
  - all machines with addresses that are “in” are trusted no matter where they are

VPN: logical view

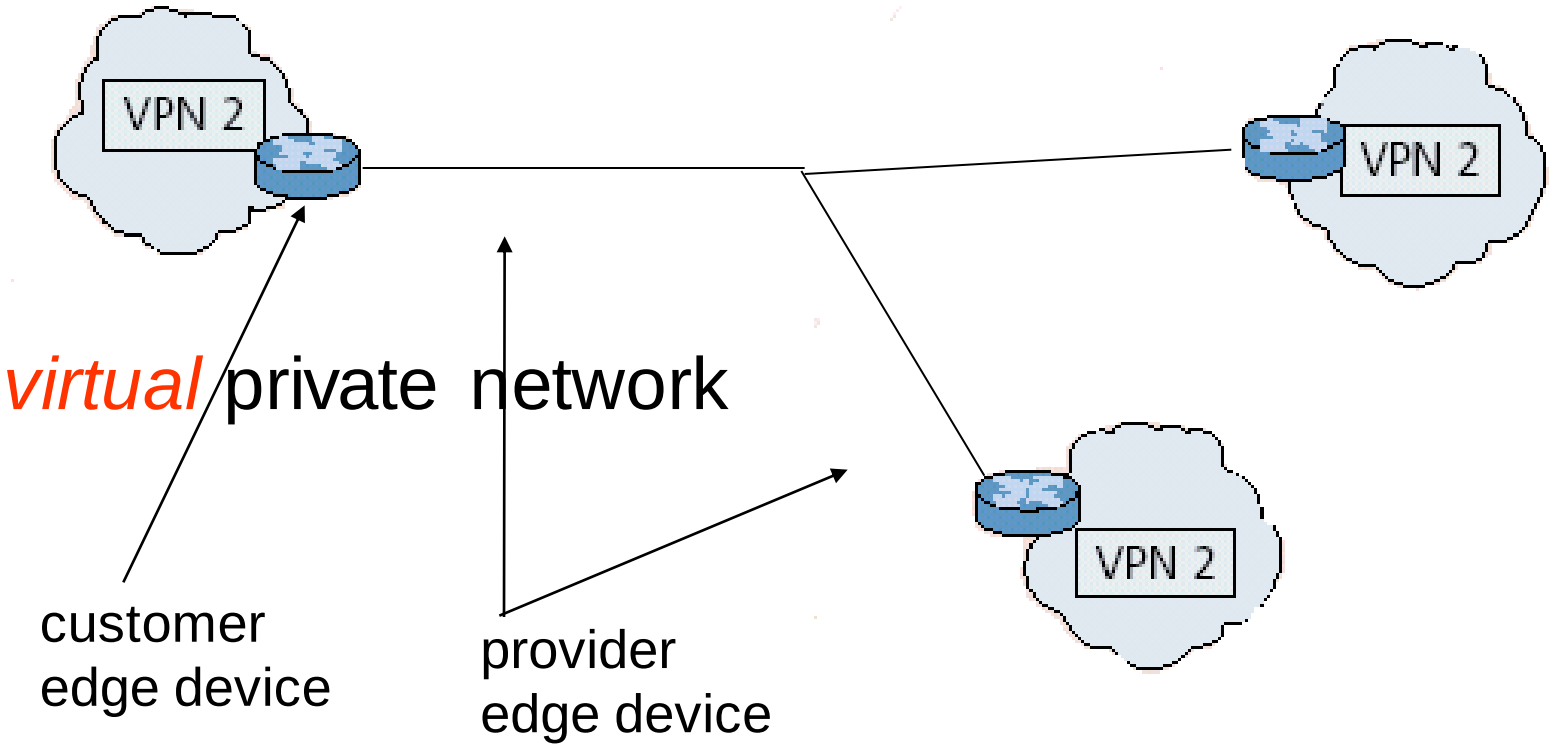

Leased-Line VPN

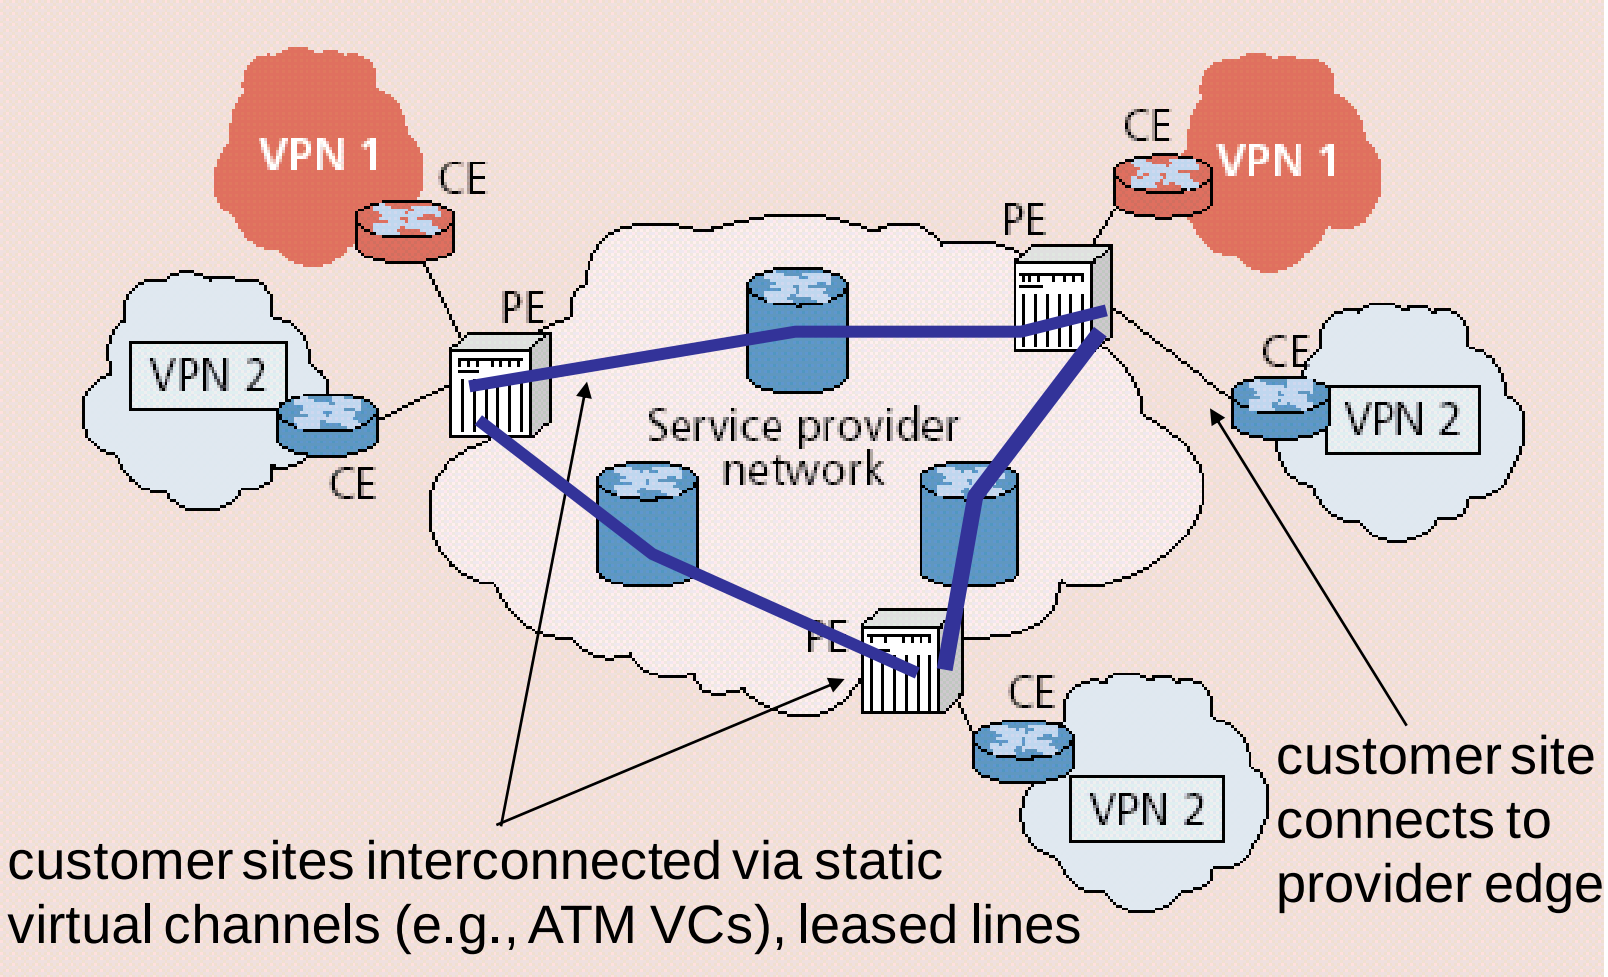

Customer Premise VPN

- all VPN functions implemented by customer

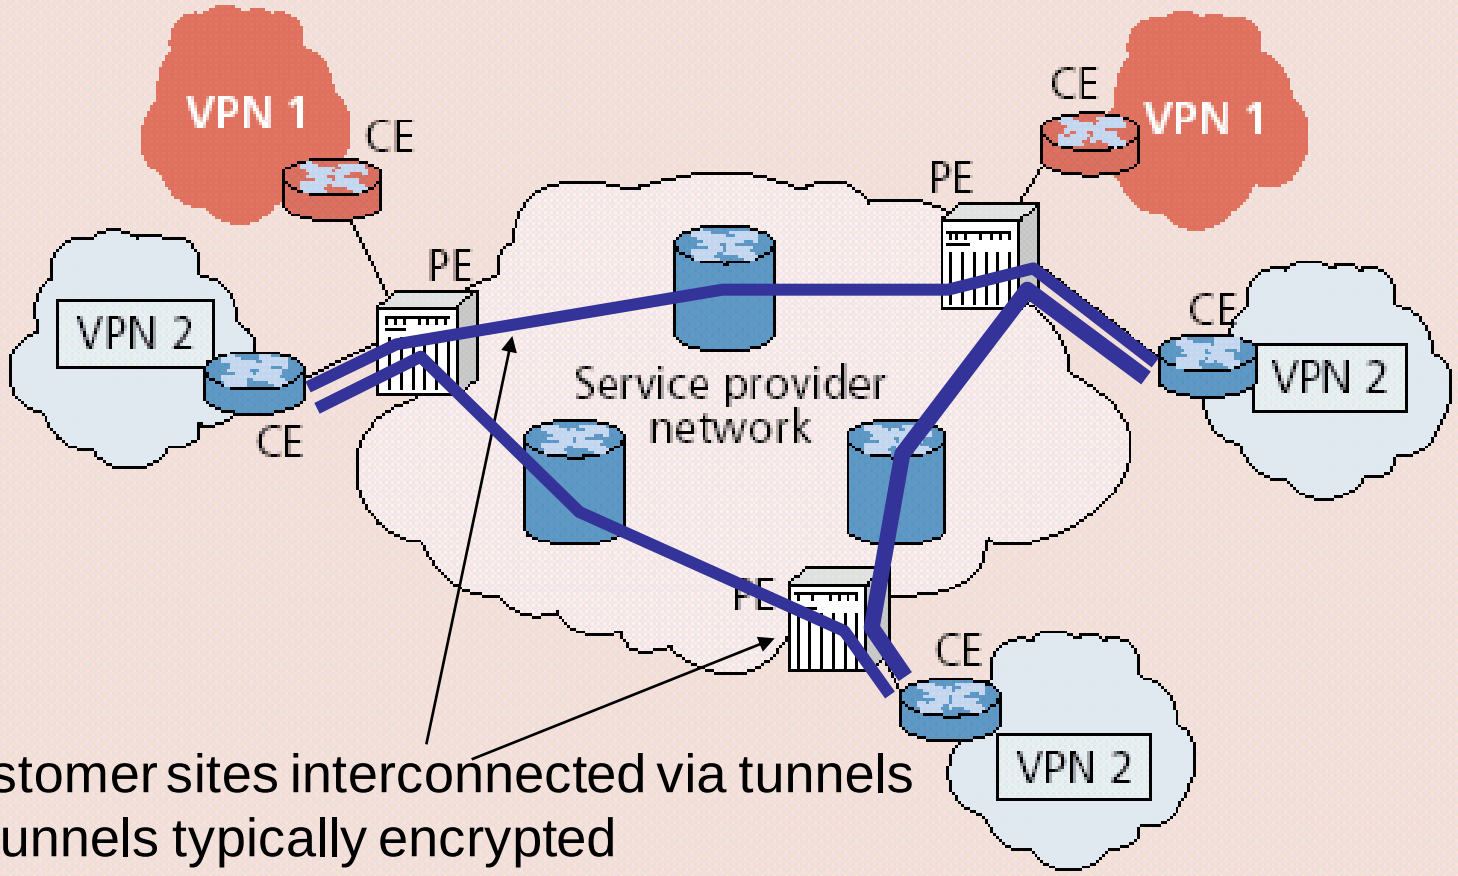

customer sites interconnected via tunnels

- tunnels typically encrypted
- Service provider treats VPN packets like all other packets

Variants of VPNs

- Leased-line VPN
  - configuration costs and maintenance by service provider: long time to set up, manpower
- CPE-based VPN
  - expertise by customer to acquire, configure, manage VPN
- Network-based VPN
  - Customer routers connect to service provider routers
  - Service provider routers maintain separate (independent) IP contexts for each VPN
    - sites can use private addressing
    - traffic from one VPN cannot be injected into another

Network-based Layer 3 VPNs

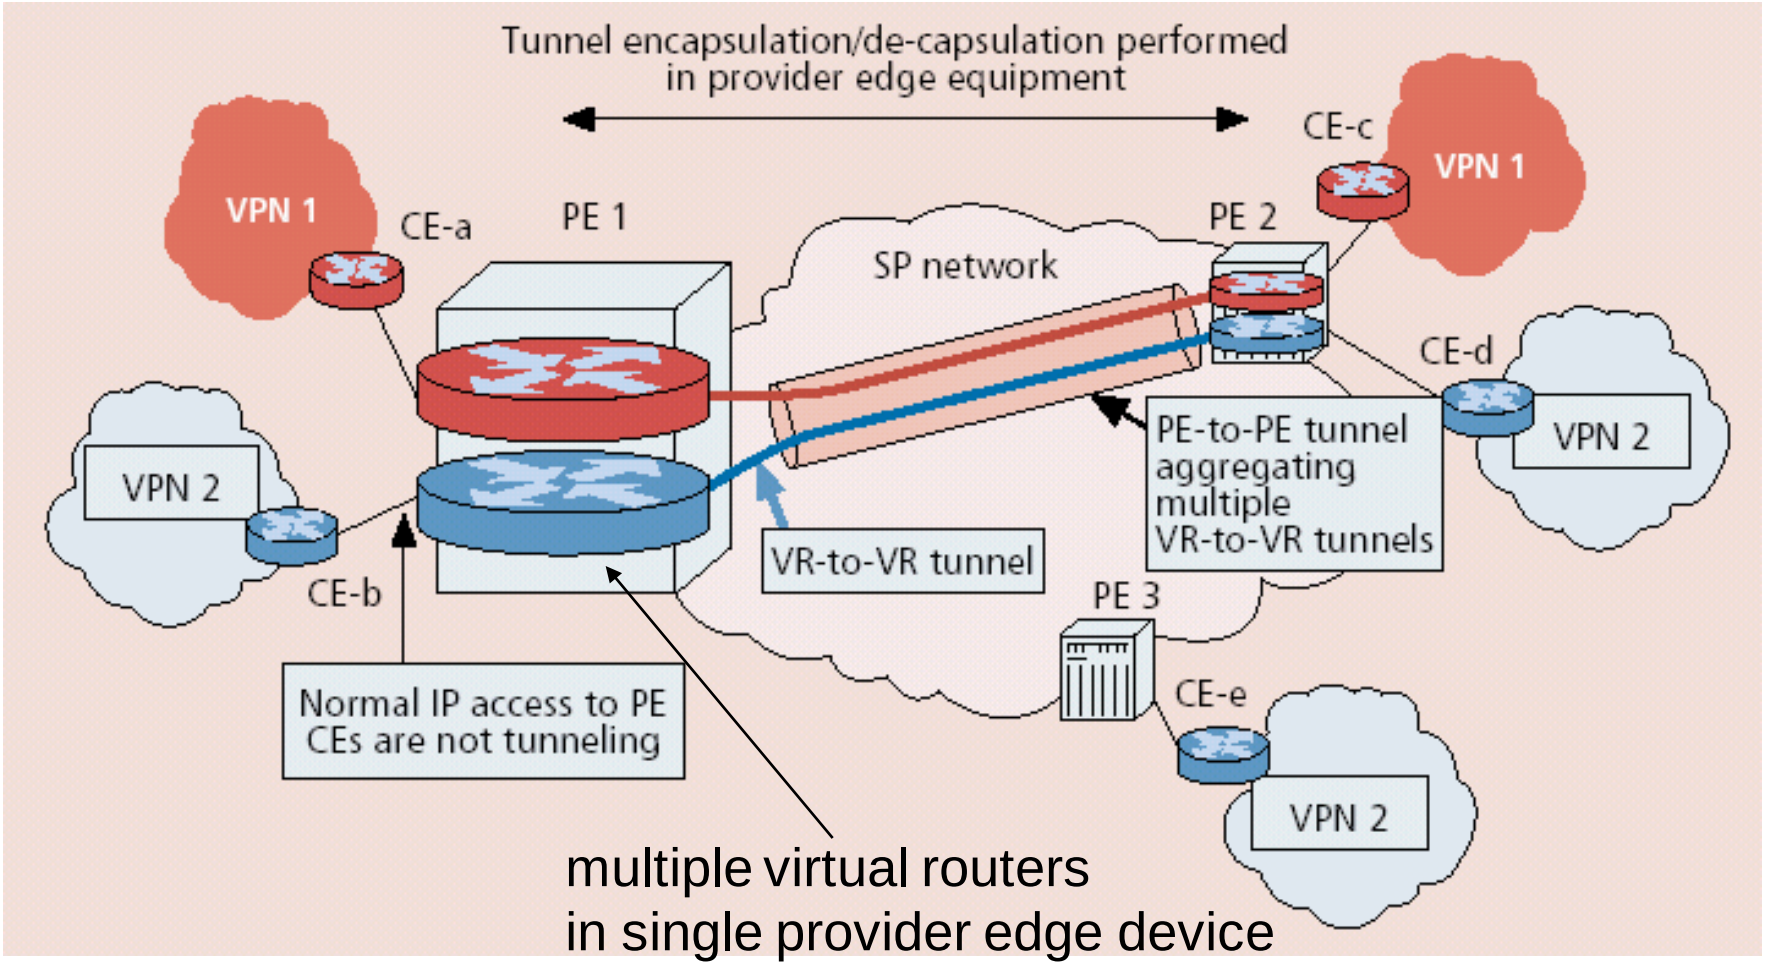

Tunneling

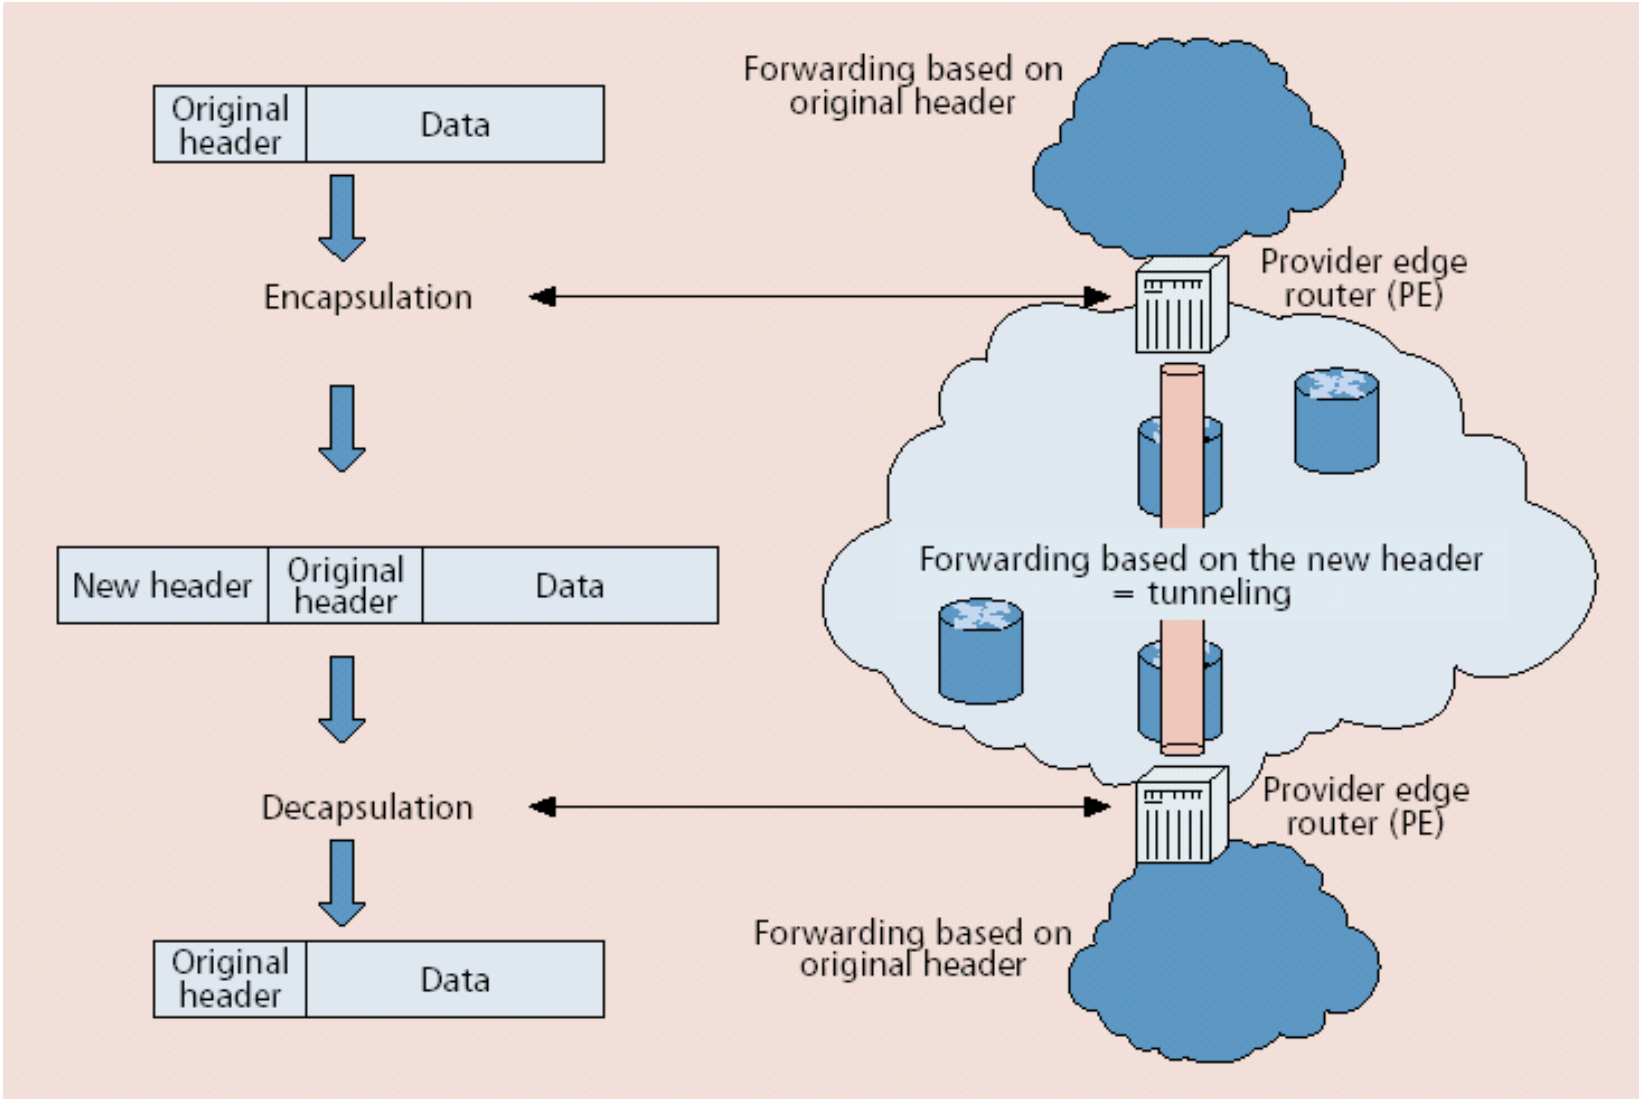
Original header
Data
Encapsulation
New header
Original header
Data
Decapsulation
Original header
Data
Forwarding based on original header
Provider edge router (PE)
Forwarding based on the new header = tunneling
Provider edge router (PE)
Forwarding based on original header
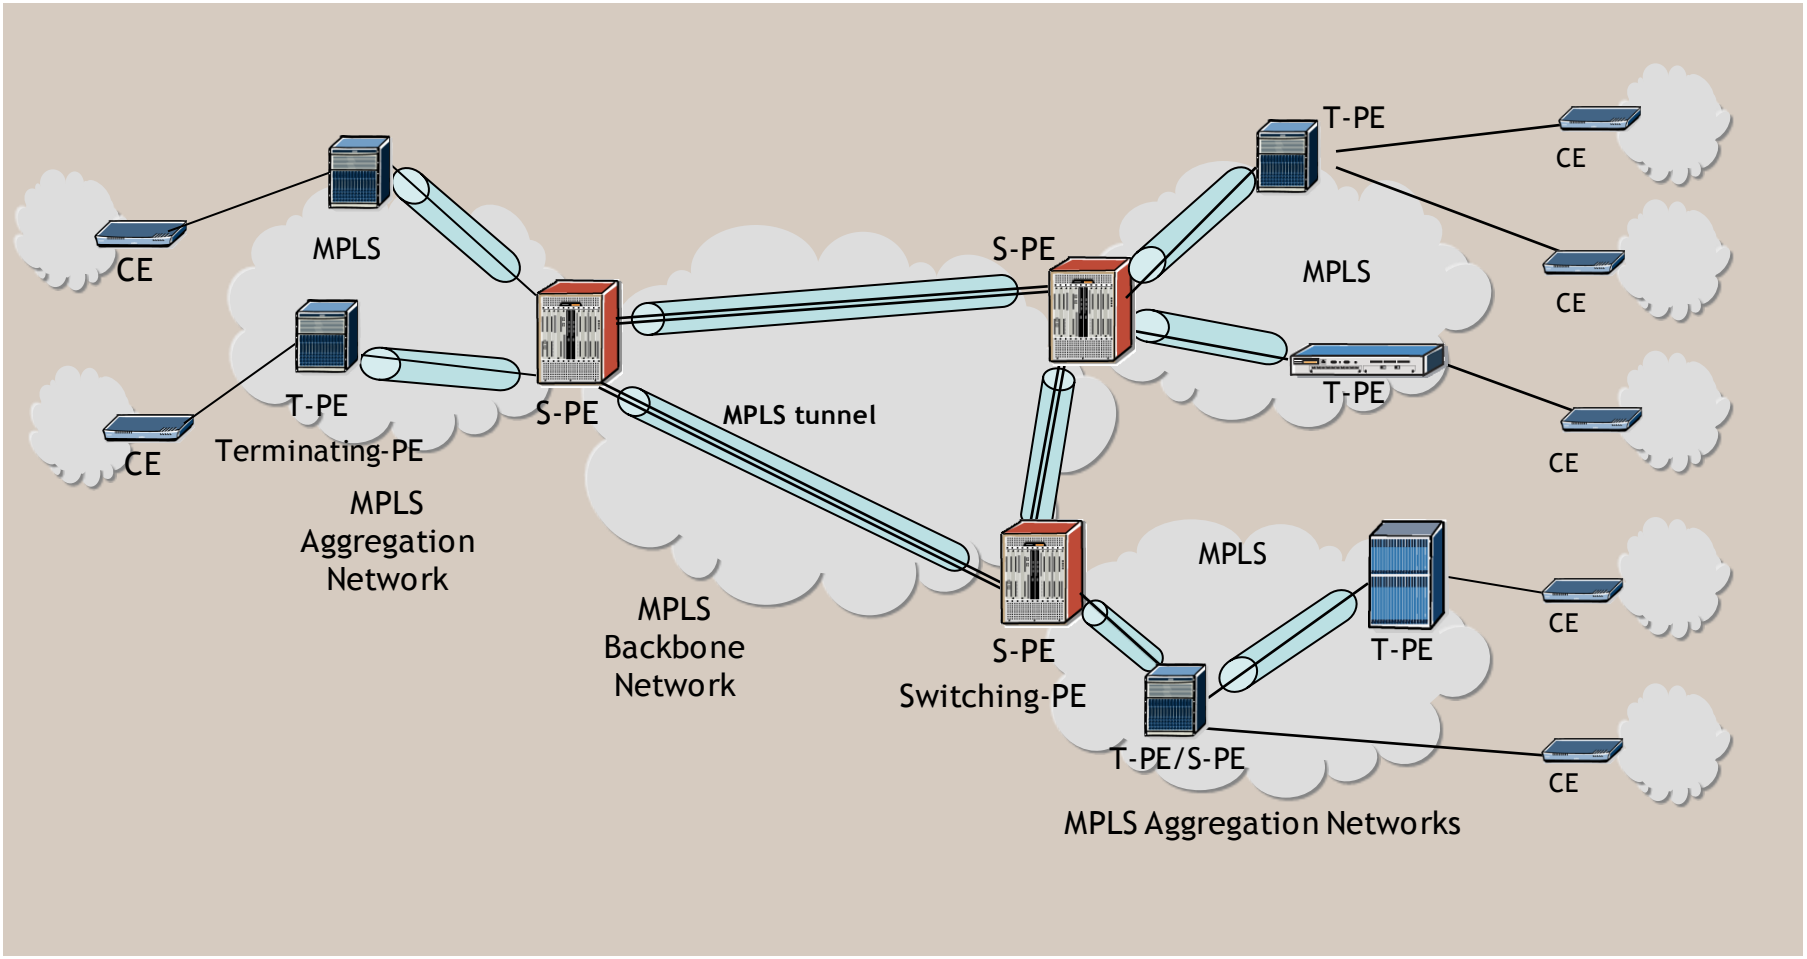
MPLS-based VPN