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

- Project status
- Internet Strucuture
- Network virtualisation

Talk announcement

Mon 28 Nov, 11:00-12:00 (s.t.), Room 00.07.014

Prof. Andy Hopper, Ph.D

Head of Informatics, University of Cambridge

Computing for the Future of the Planet

Abstract: Digital technology is becoming an indispensable and crucial component of our lives, society, and environment. A framework for computing in the context of problems facing the planet will be presented. The framework has a number of goals: an optimal digital infrastructure, sensing and optimising with a global world model, reliably predicting and reacting to our environment, and digital alternatives to physical activities. Practical industrial examples will be given as well as research goals.

Andy Hopper has pursued academic and industrial careers simultaneously. In the academic career he has worked in the Computer Laboratory and the Department of Engineering at Cambridge. In the industrial context he has co-founded a dozen spin-outs and start-ups, three of which floated on stock markets. He is currently Chairman of RealVNC and Ubisense plc.

First Feedback

- Many highly appropriate project plans :)
- team00; team07; team19: please contact me.
- Interesting: Gannt charts
- Interesting: different plans for communication among team members, including mobile phone, skype, instant messaging
- Interesting: different addditional tools: dropbox, git, google docs

⇒ Let us see what your final recommendations are

⇒ Let us see whether/how project plan correlates with outcome

Internet Structure

Internet structure: network of networks

- roughly hierarchical
- at center: "tier-1" ISPs (AT&T, Global Crossing, Level 3, NTT, Qwest, Sprint, Tata, Verizon (UUNET), Savvis, TeliaSonera), national/international coverage
  - treat each other as equals
  - can reach every other network on the Internet without purchasing IP transit or paying settlements

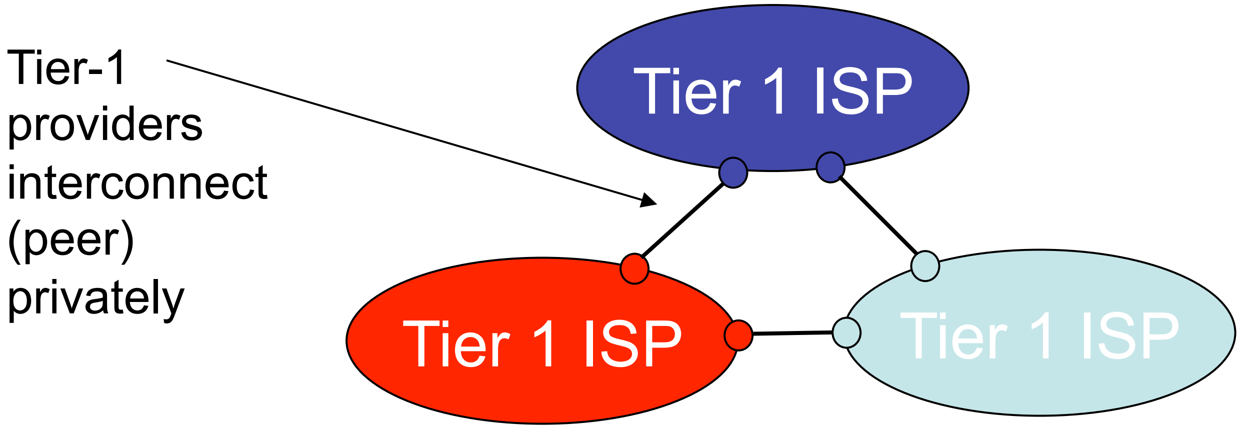

Tier-1 ISP: e.g., Sprint

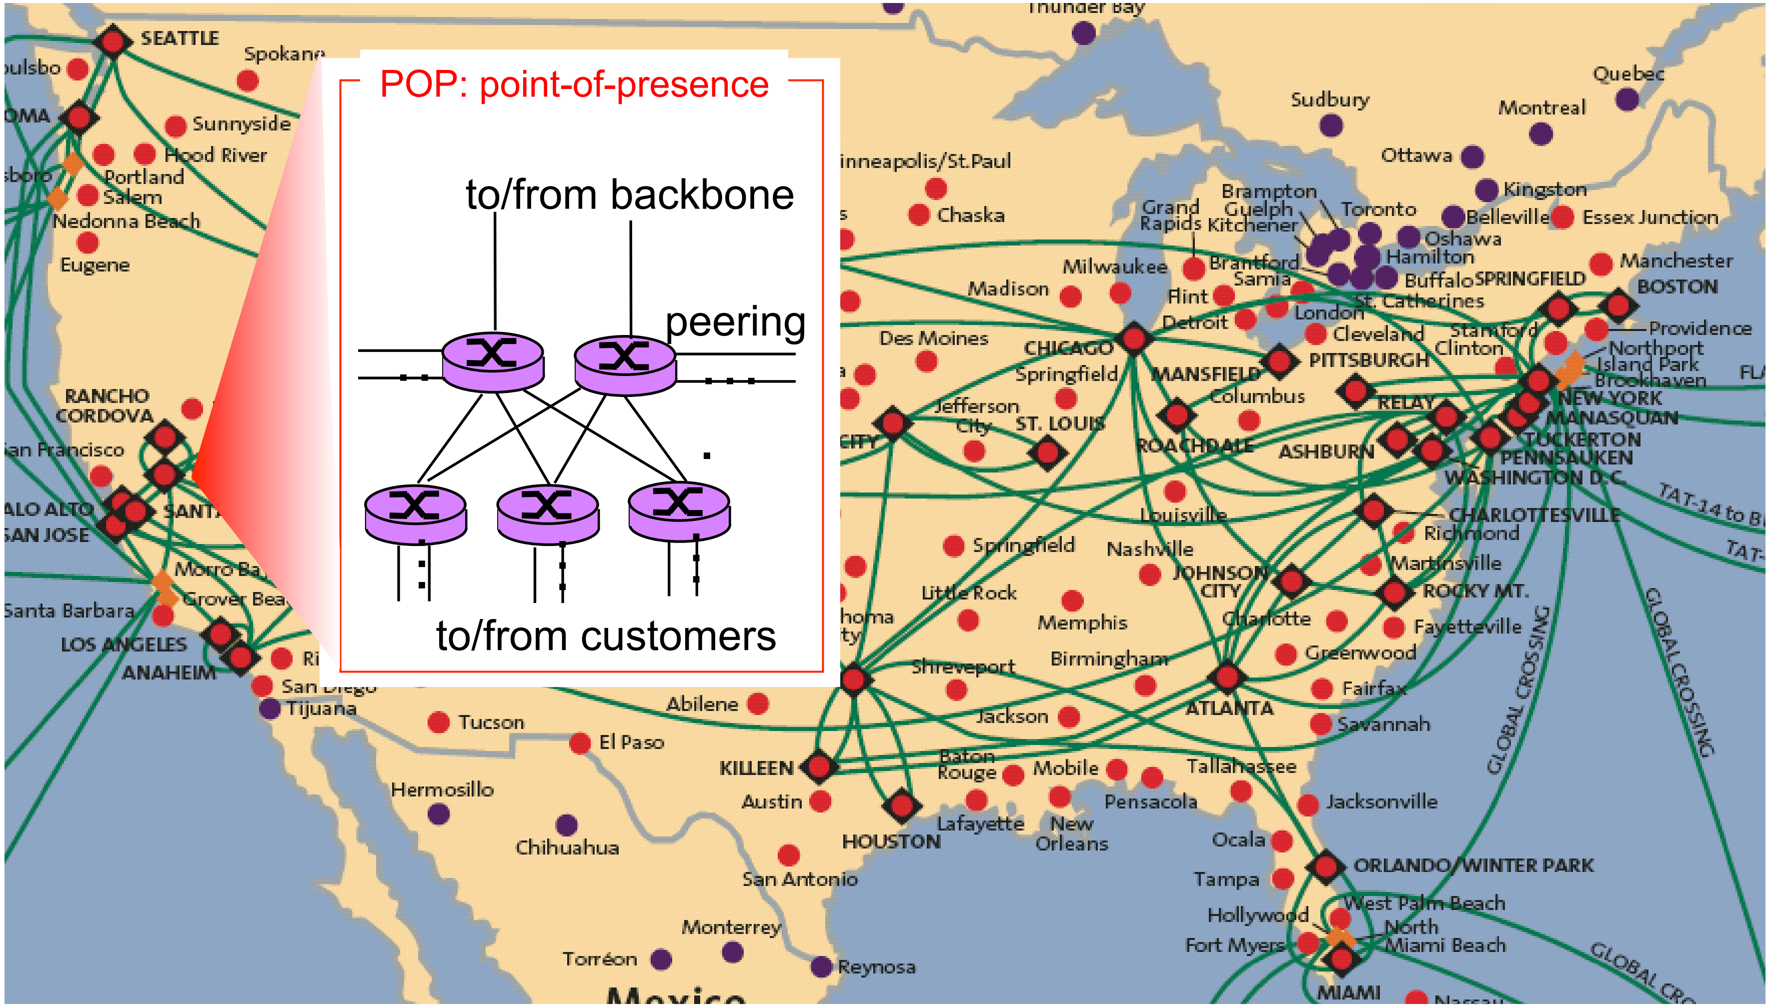

Internet structure: network of networks

- "Tier-2" ISPs: smaller (often regional) ISPs
  - Connect to one or more tier-1 ISPs, possibly other tier-2 ISPs

- Tier-2 ISP pays tier-1 ISP for connectivity to rest of Internet
- tier-2 ISP is *customer* of tier-1 provider

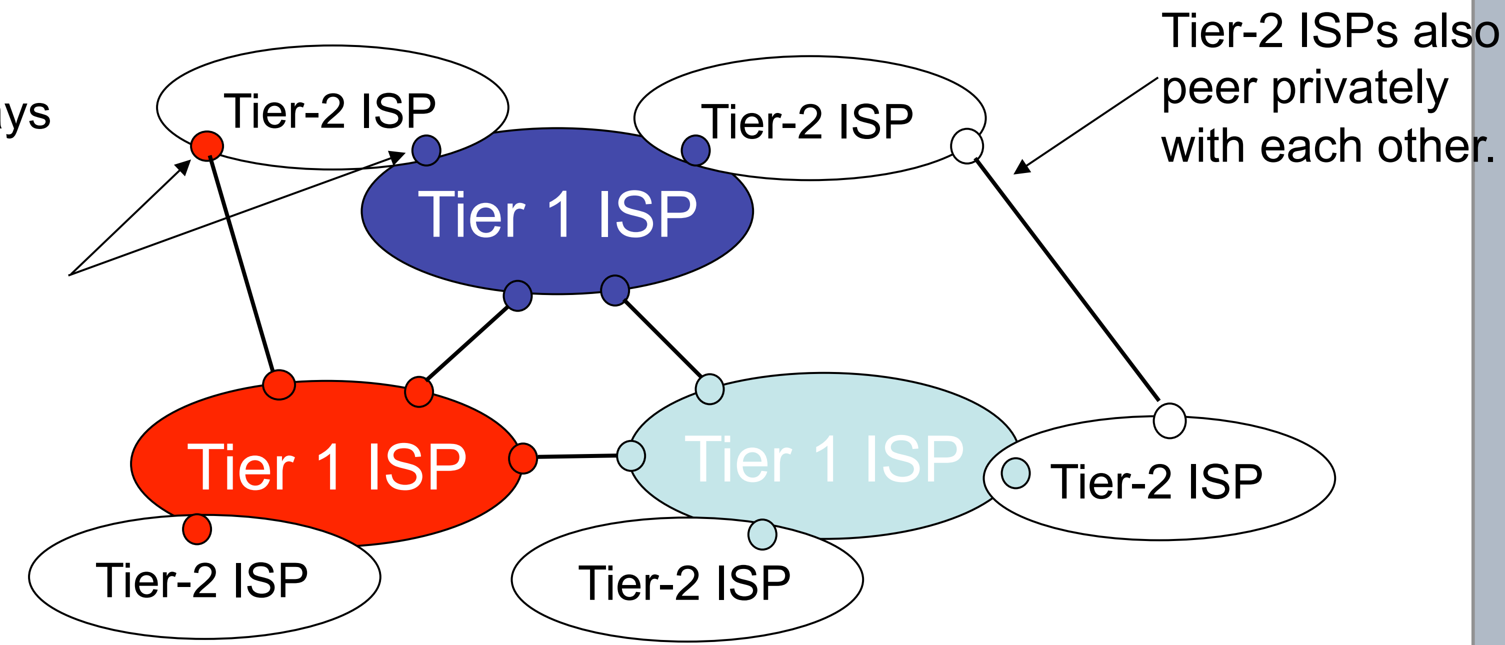

Internet structure: network of networks

- "Tier-3" ISPs and local ISPs
  - last hop ("access") network (closest to end systems)

Local and tier-3 ISPs are *customers* of higher tier ISPs connecting them to rest of Internet

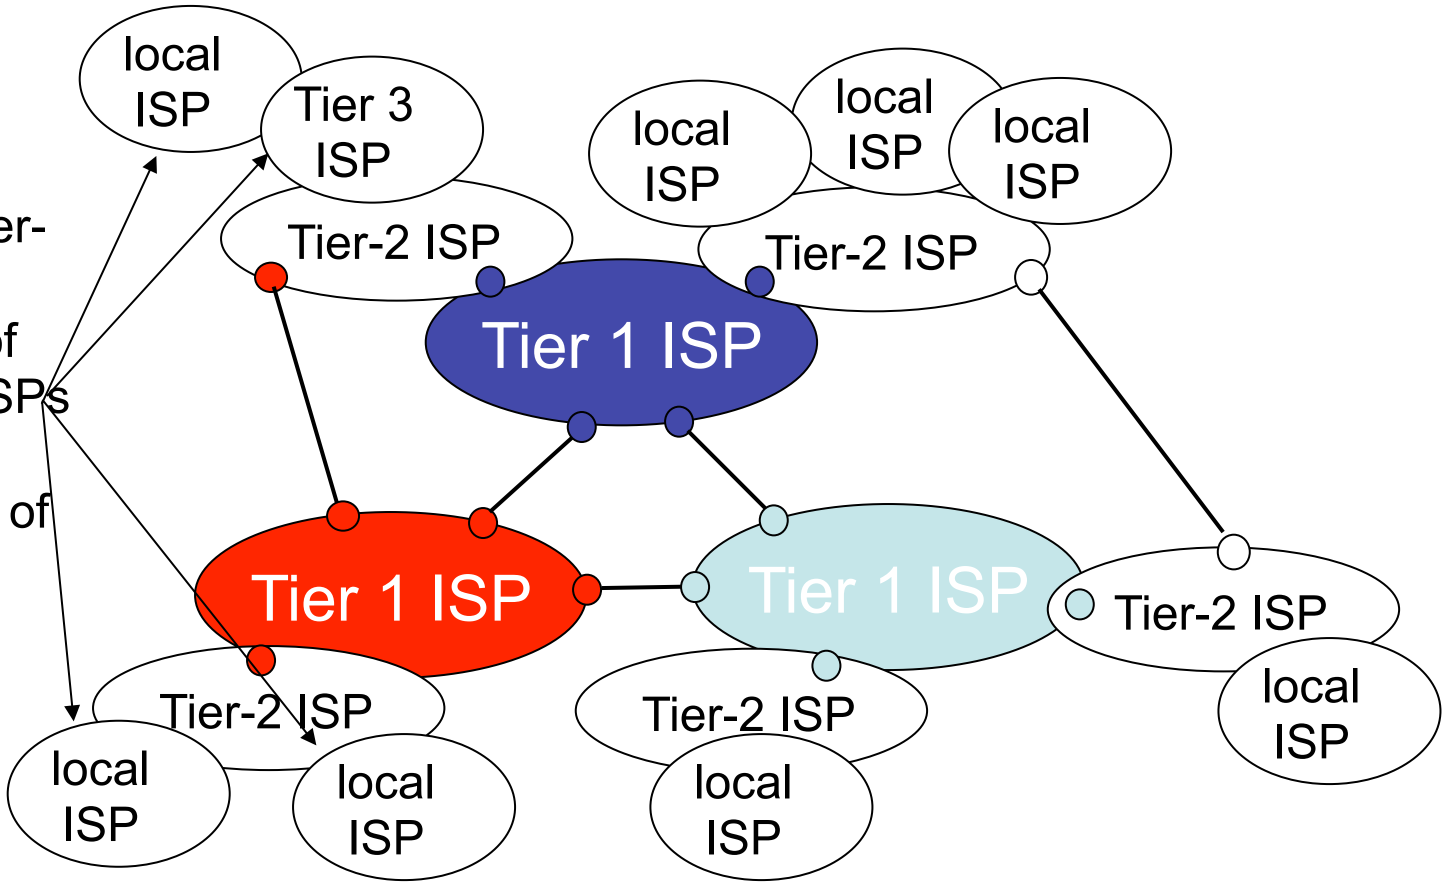

Internet structure: network of networks

- a packet passes through many networks!

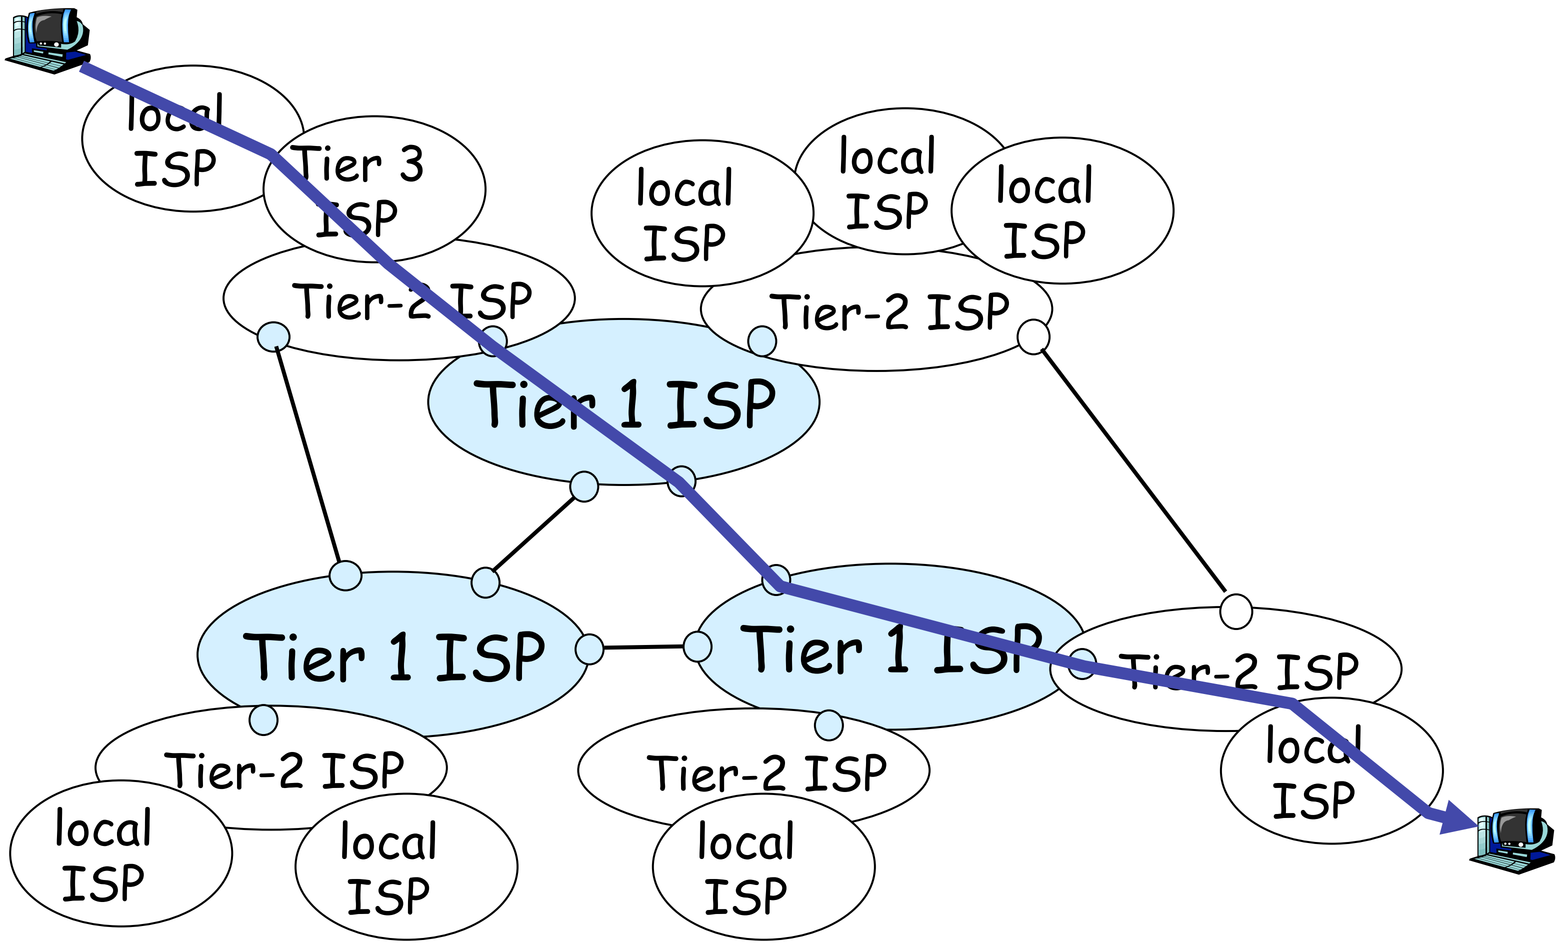

Internet Ecosystem

- >30,000 autonomous networks
- Networks with different
  - different roles and business type
    - stub networks
    - transit networks
    - content providers
  - Influenced by traffic patterns, application popularity, economics, regulation, ….
- Peering
  - bilateral contracts
  - Customer-provider, settlement-free peering, or in between
- Internet Exchange Points

Internet Exchange Point

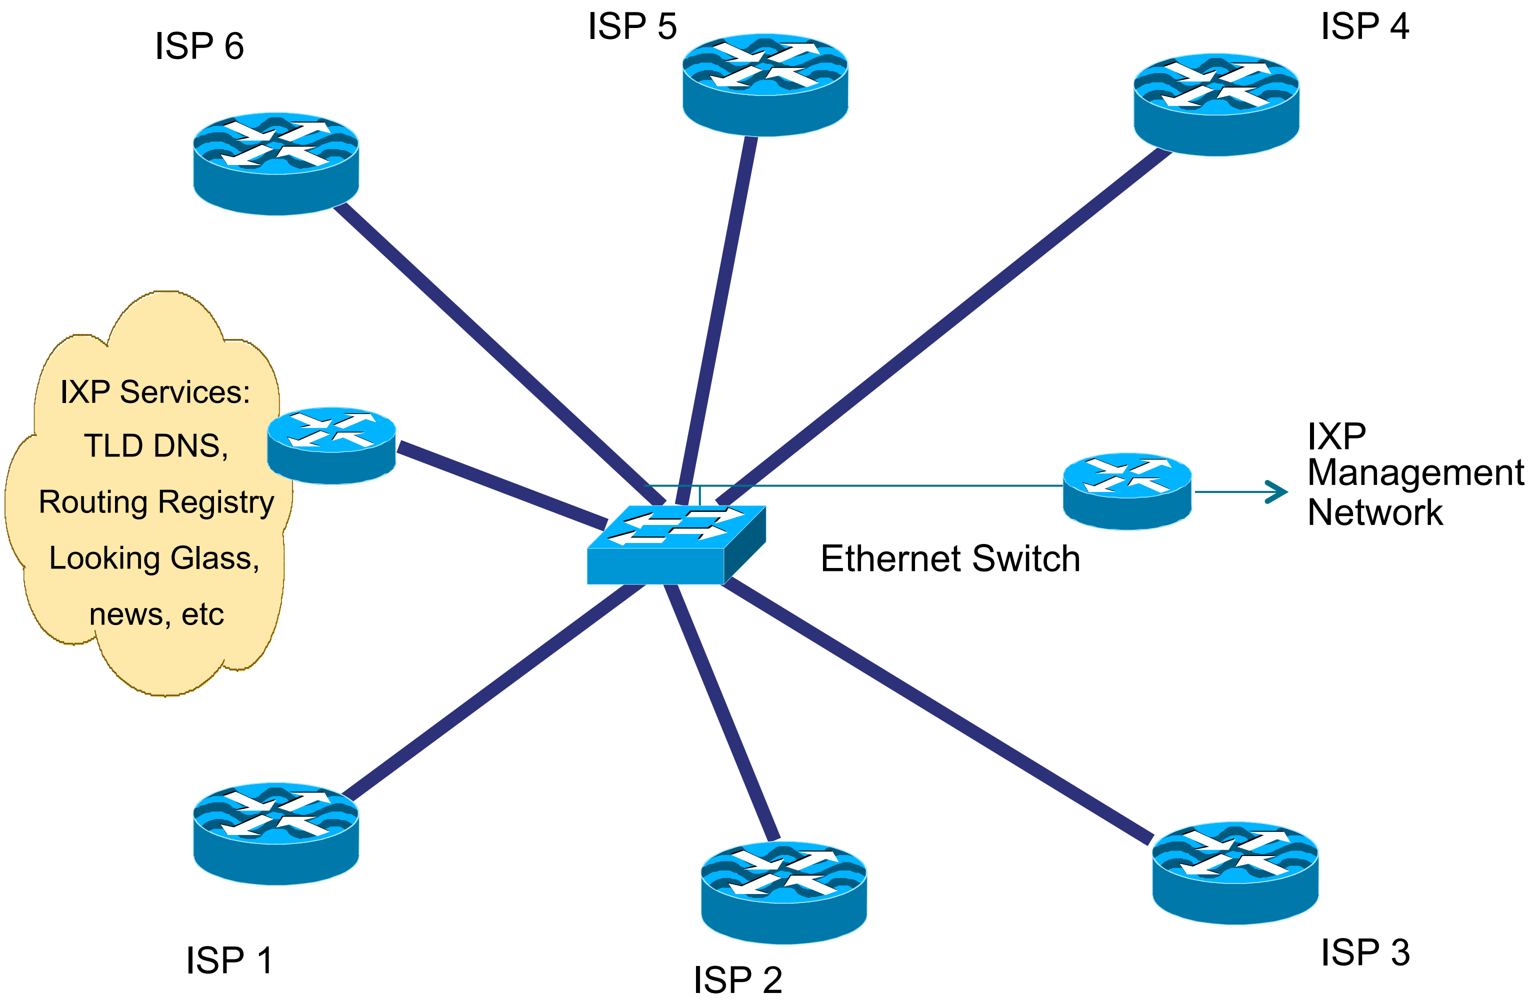

Cost of Peering at Internet Exchange Point

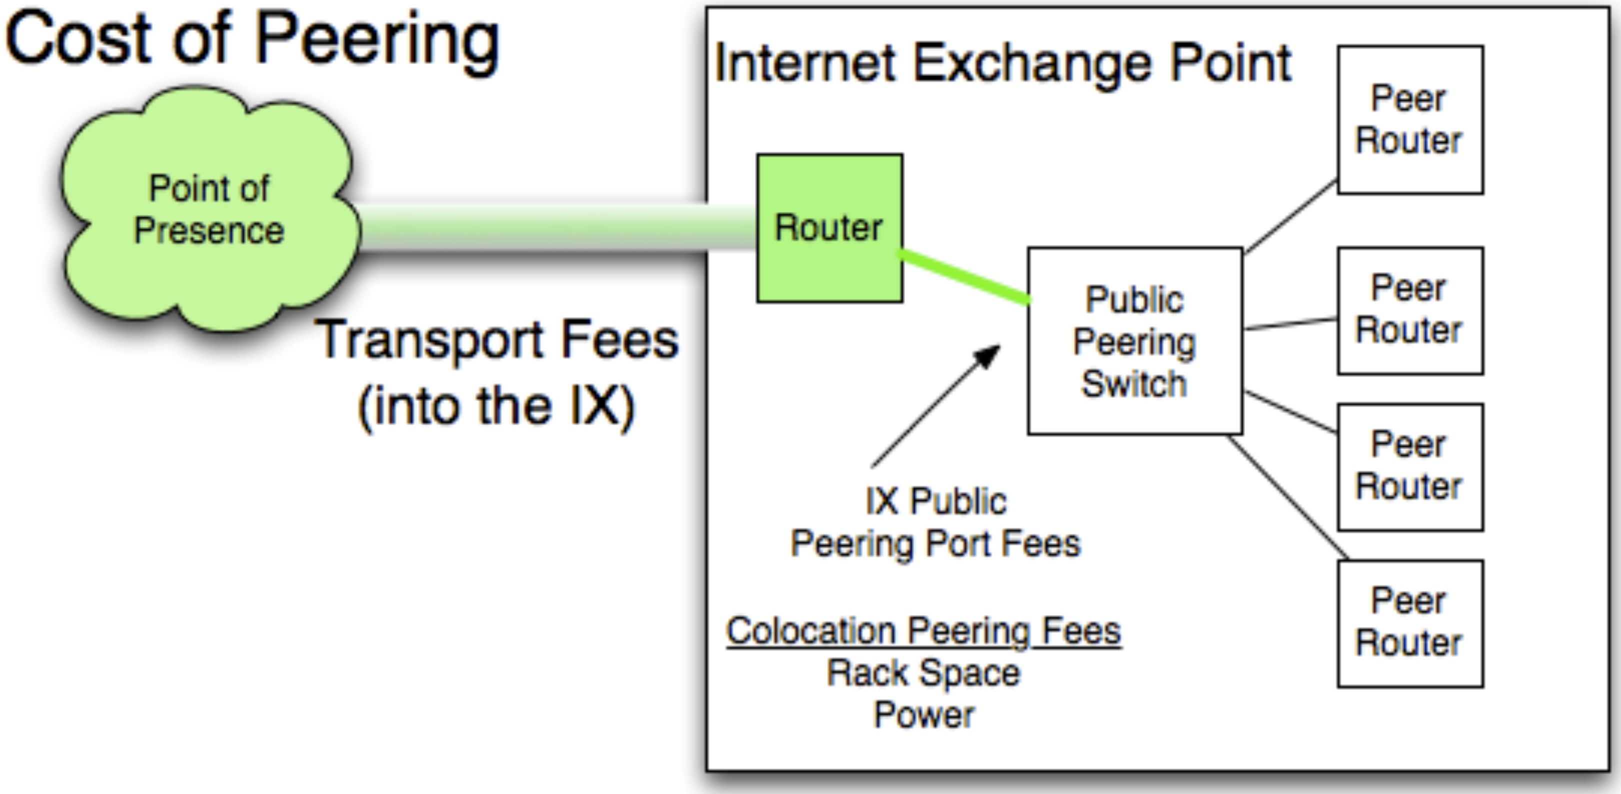

source: William B. Norton, „Internet Peering", http://drpeering.net/

ISP Peering Relations

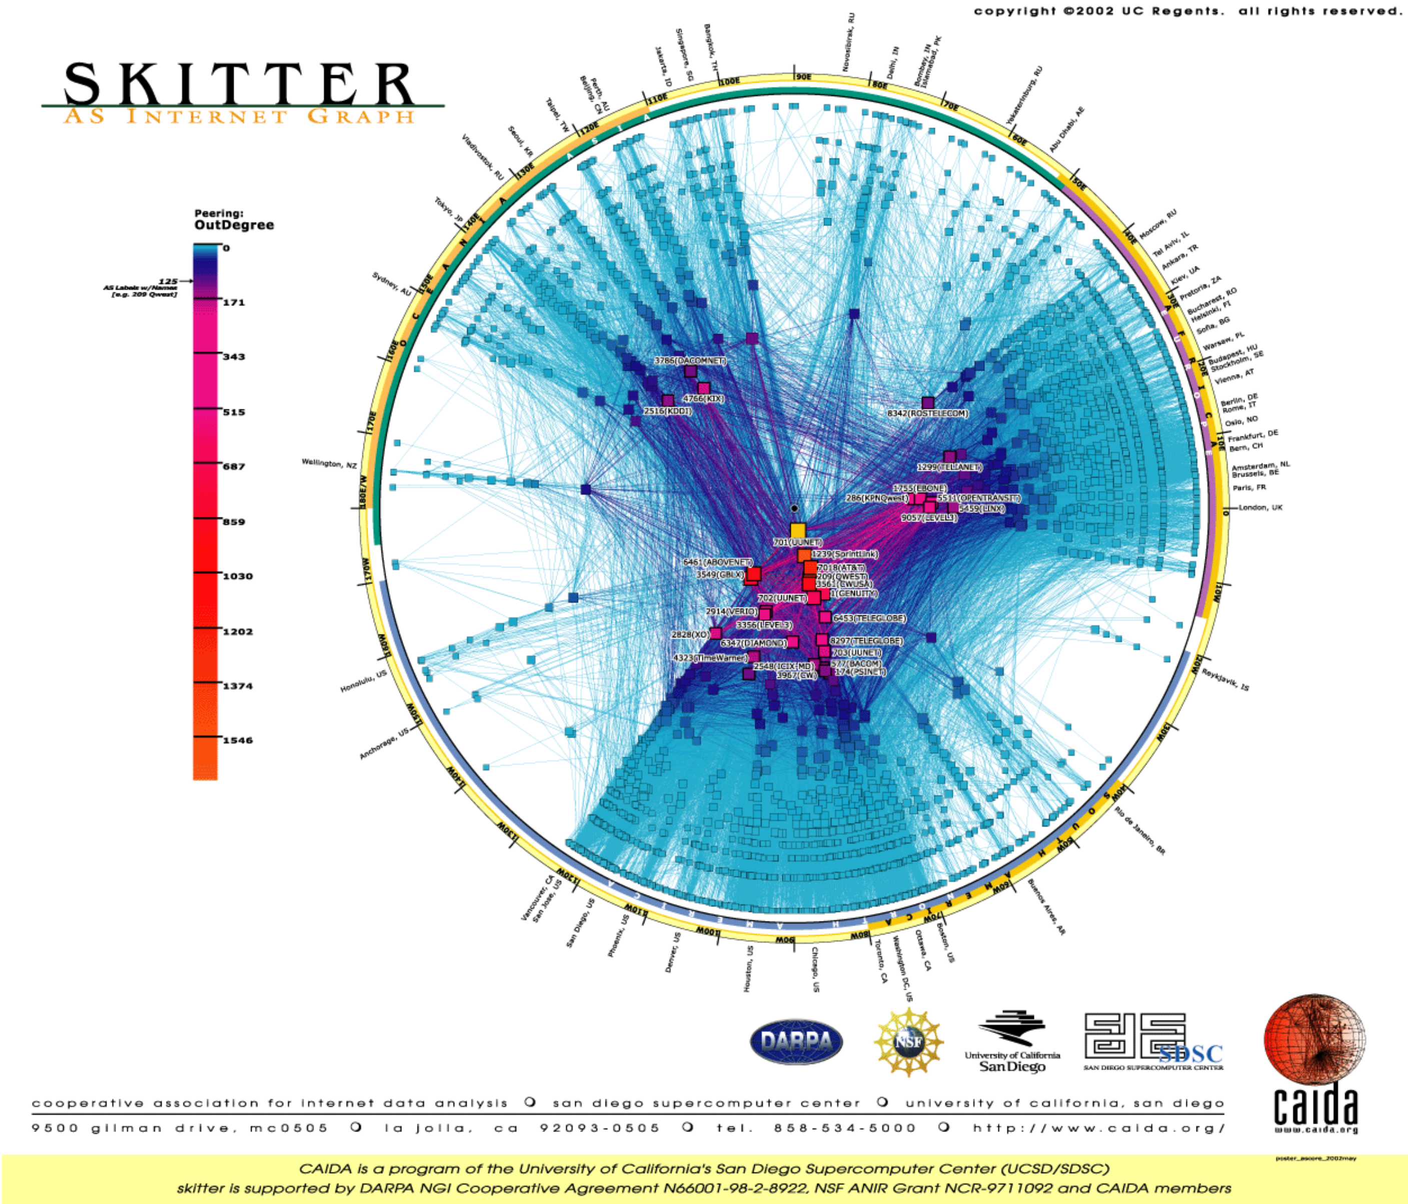

IPv4 vs. IPv6 Graphs

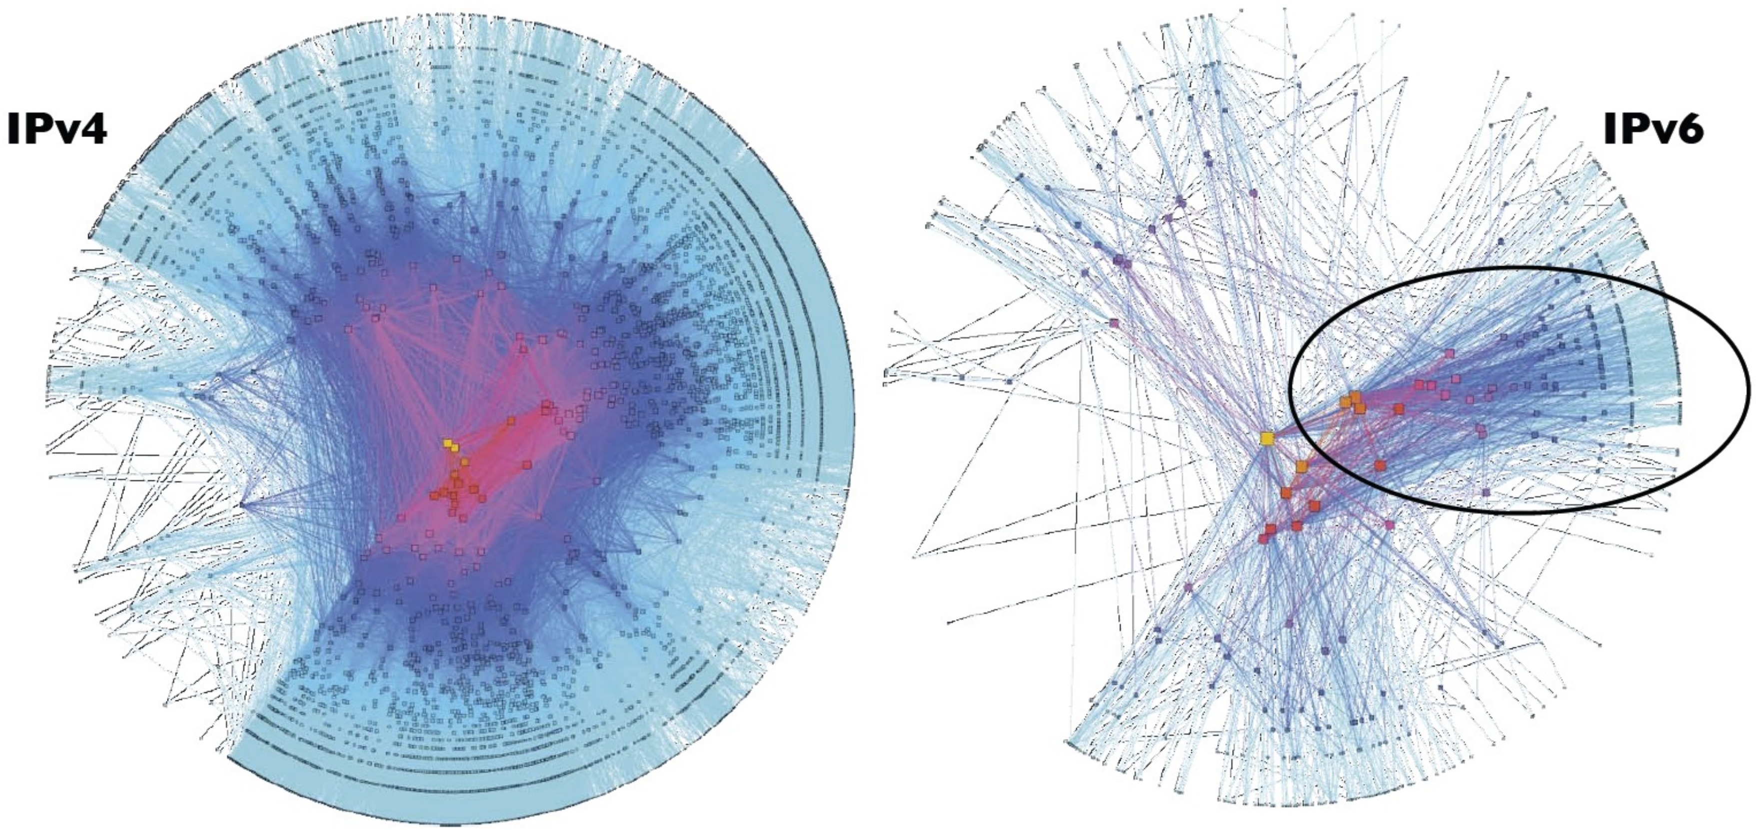

source: caida.org

AS Connectivity

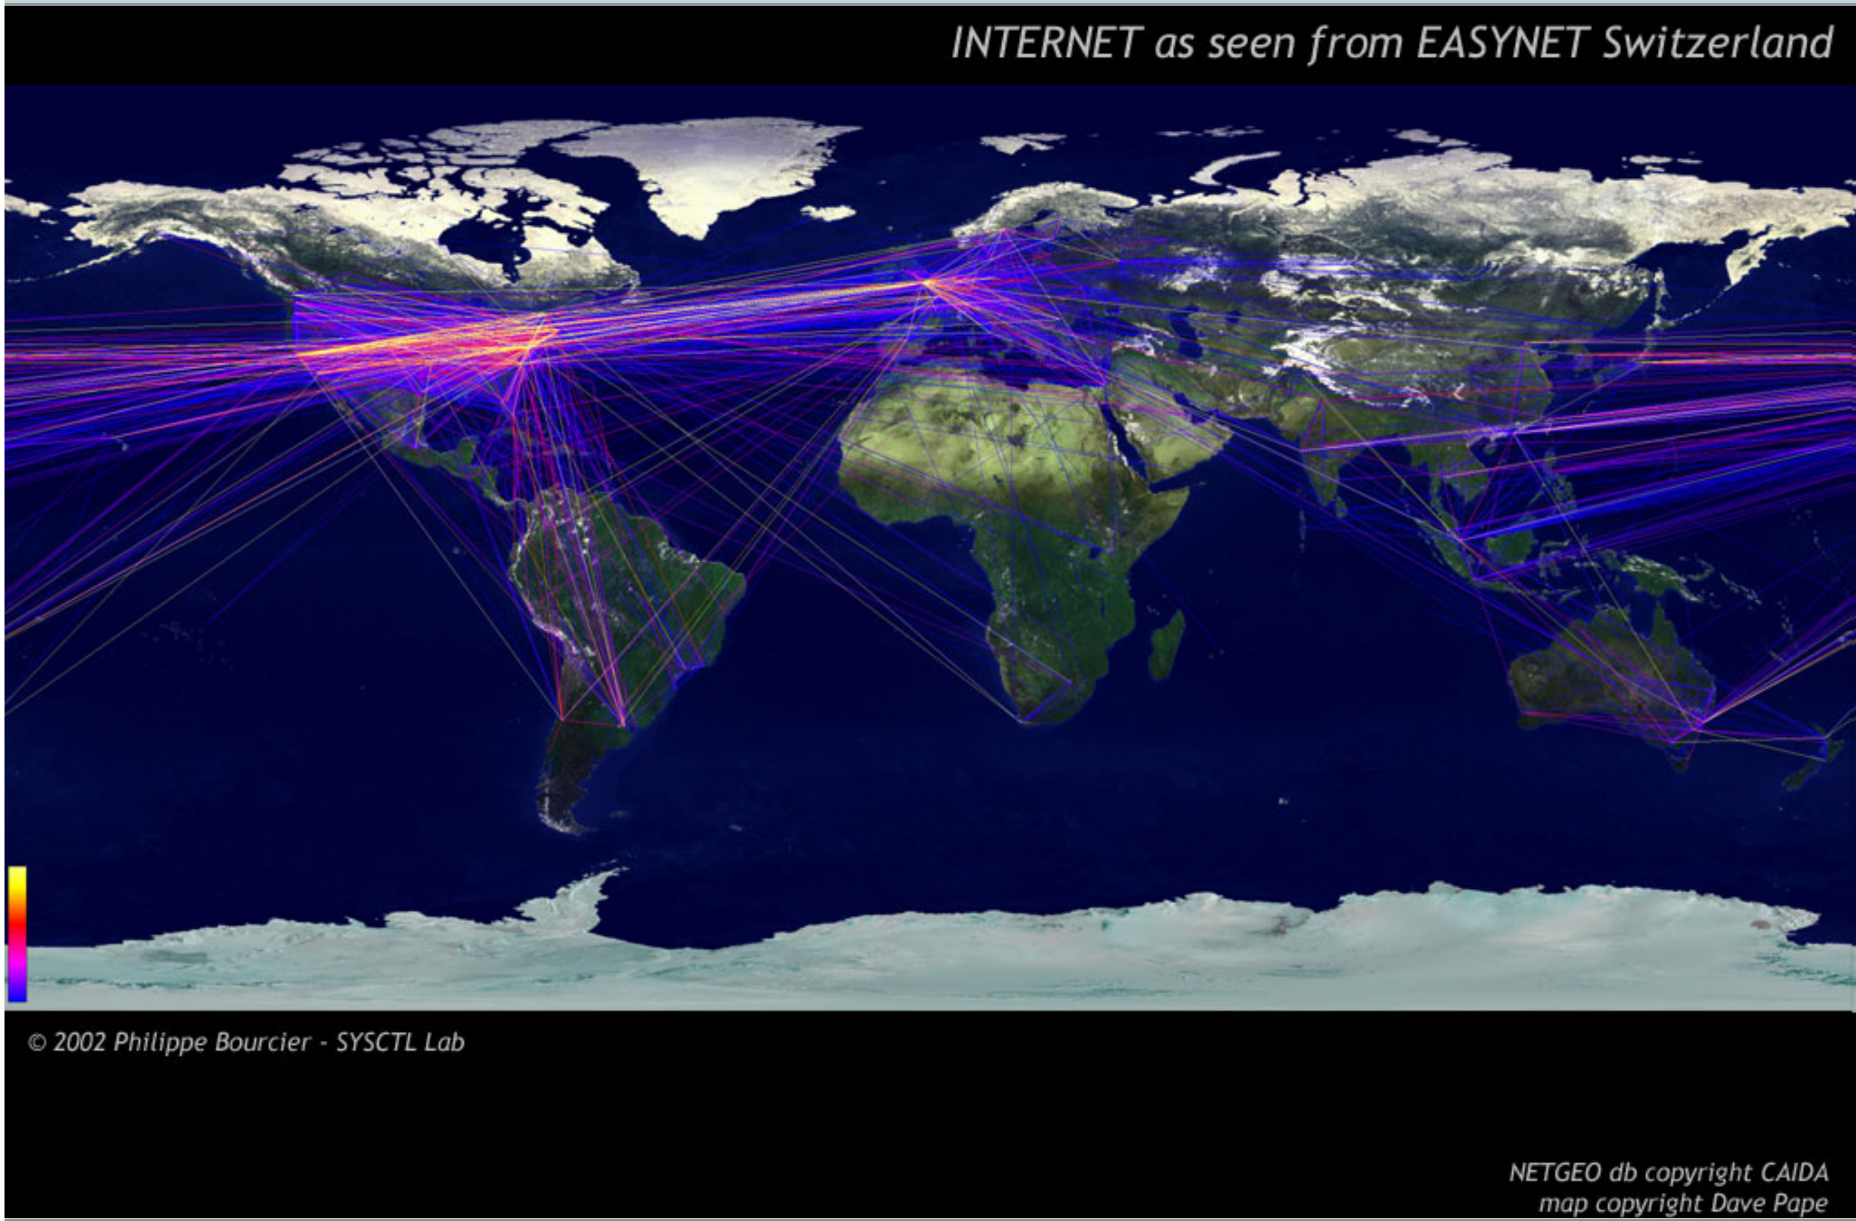

Chair for Network Architectures and Services – Prof. Carle
Department for Computer Science
TU München

Network Architectures

Link virtualization: ATM, MPLS

Technische Universität München

Virtualization of networks

- Virtualization of resources
  - powerful abstraction in systems engineering
- Computing examples of virtualisation popular for long time
  - virtual memory
  - virtual devices
  - Virtual machines: e.g., Java
  - IBM VM operation system (1960's/70's)
- Layering of abstractions
  - don't bother with details of the lower layer, only deal with lower layers abstractly

The Internet: Virtualizing Networks

- 1974: multiple unconnected networks
  - ARPAnet
  - data-over-cable networks
  - packet satellite network (Aloha protocol)
  - packet radio network

… differing in:

- addressing conventions
- packet formats
- error recovery
- routing

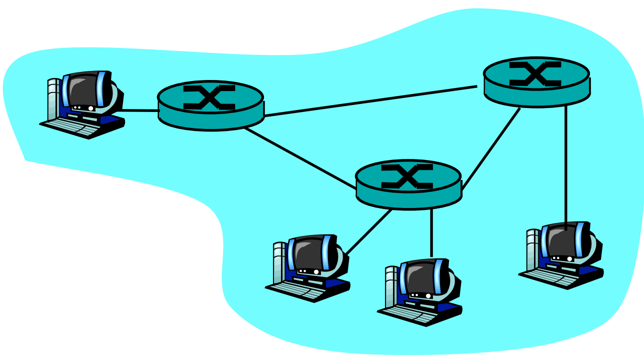

ARPAnet

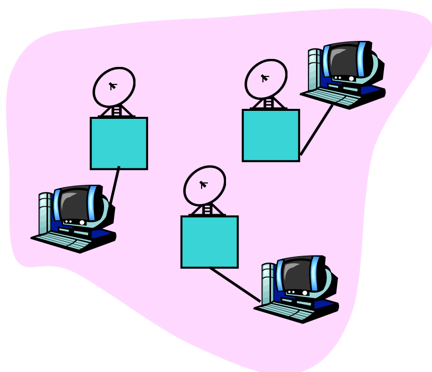

satellite net

"A Protocol for Packet Network Intercommunication",
V. Cerf, R. Kahn, IEEE Transactions on Communications,
May, 1974, pp. 637-648.

The Internet: Virtualizing networks

Internetwork layer (IP):
- addressing: internetwork appears as single, uniform entity, despite underlying local network heterogeneity
- network of networks

Gateway:
- "embed internetwork packets in local packet format or extract them"
- route (at internetwork level) to next gateway

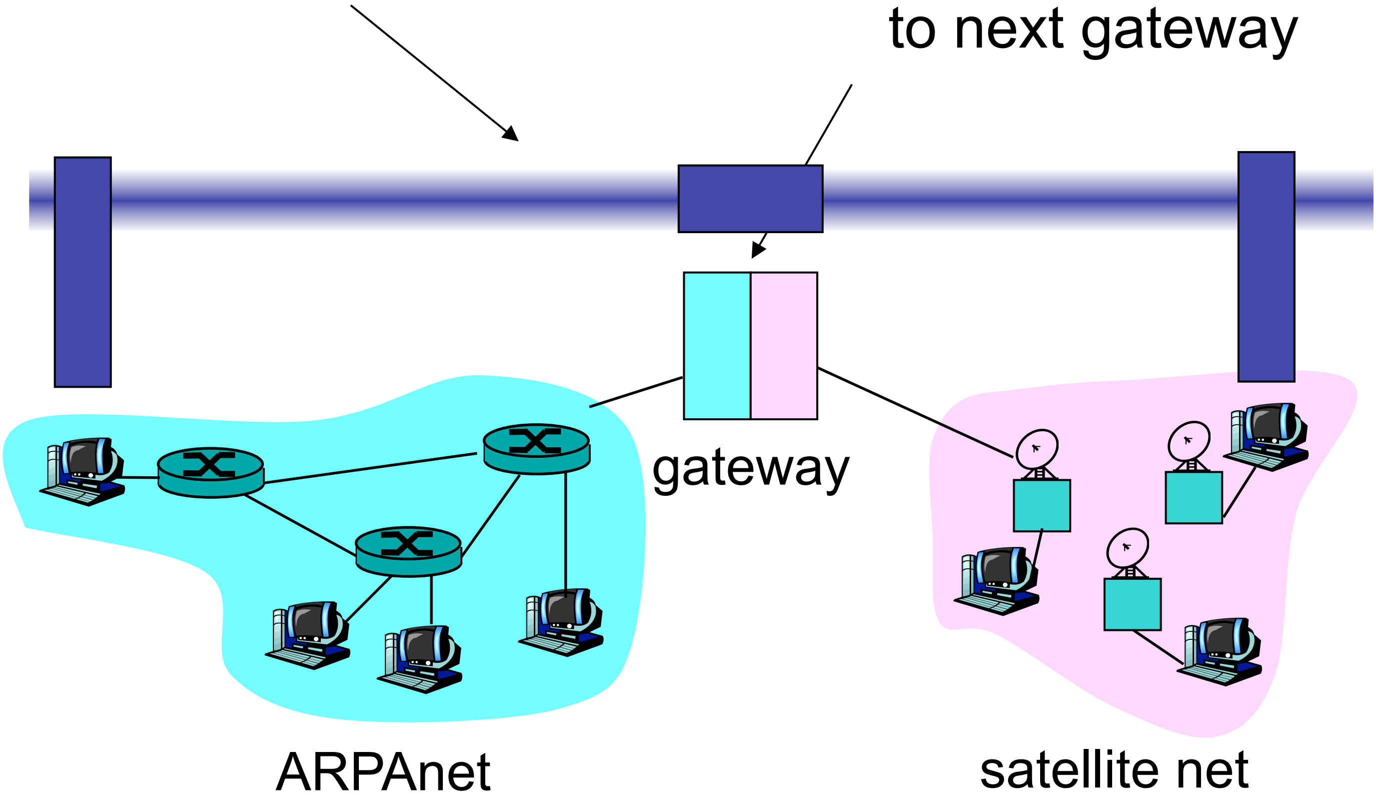

Cerf & Kahn's Internetwork Architecture

- What is virtualized?
- two layers of addressing: internetwork and local network
- new layer (IP) makes everything homogeneous at internetwork layer
- underlying local network technology
  - cable
  - satellite
  - 56K telephone modem
  - today: ATM, MPLS
- … "invisible" at internetwork layer.
  Looks like a link layer technology to IP!

ATM and MPLS

- ATM, MPLS separate networks in their own right
  - different service models, addressing, routing from Internet
- Viewed by Internet as logical link connecting IP routers
  - just like dialup link is really part of separate network (telephone network)
- ATM, MPLS: of technical interest in their own right

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

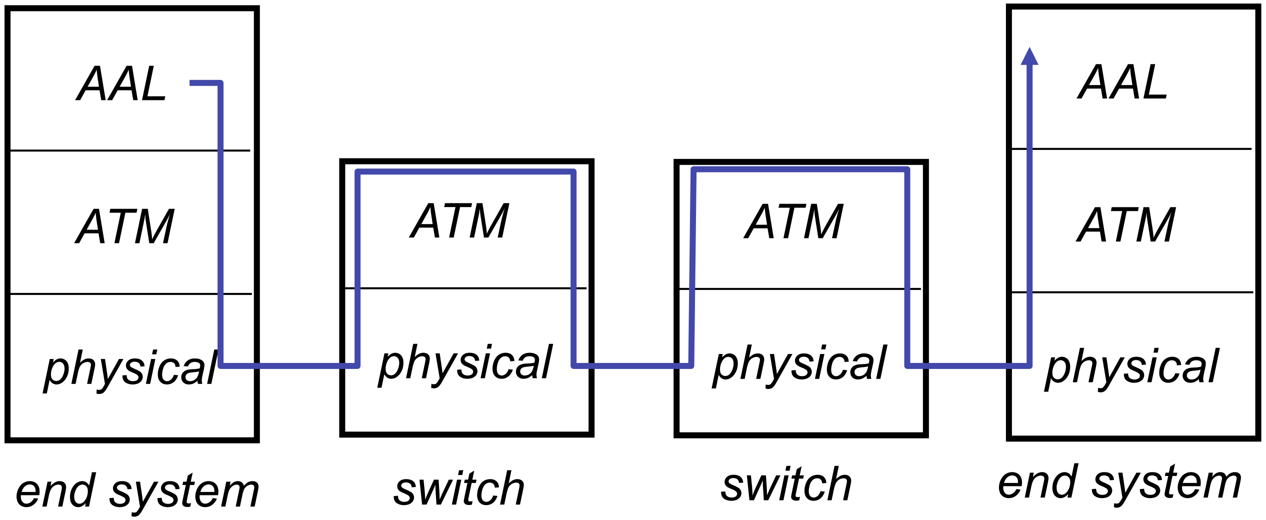

- adaptation layer: only at edge of ATM network
  - data segmentation/reassembly
  - roughly analogous to Internet transport layer
- ATM layer: "network" layer
  - cell switching, routing
- physical layer

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

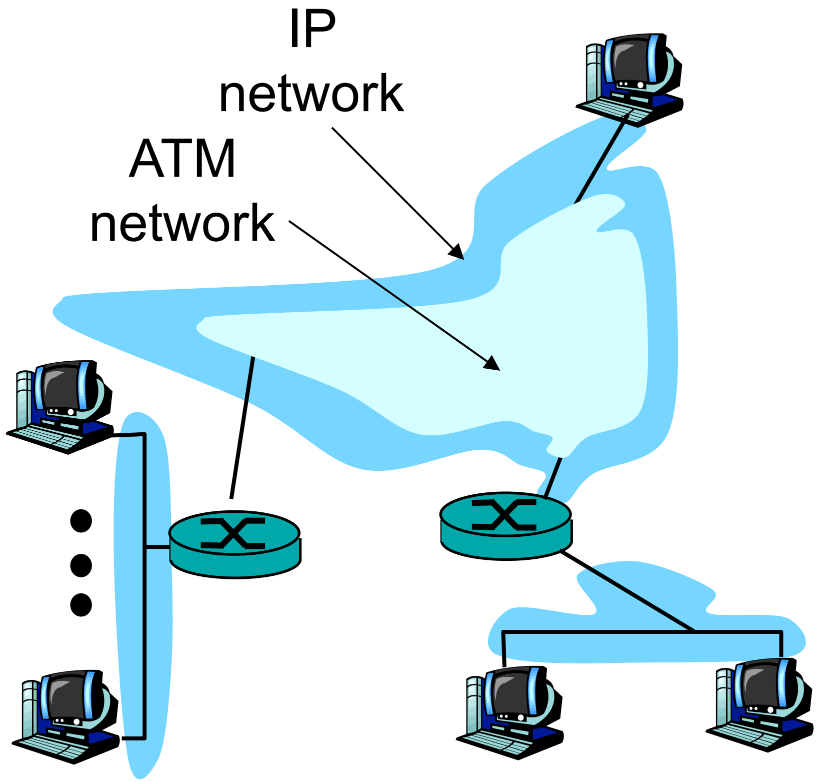

ATM Adaptation Layer (AAL)

- 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

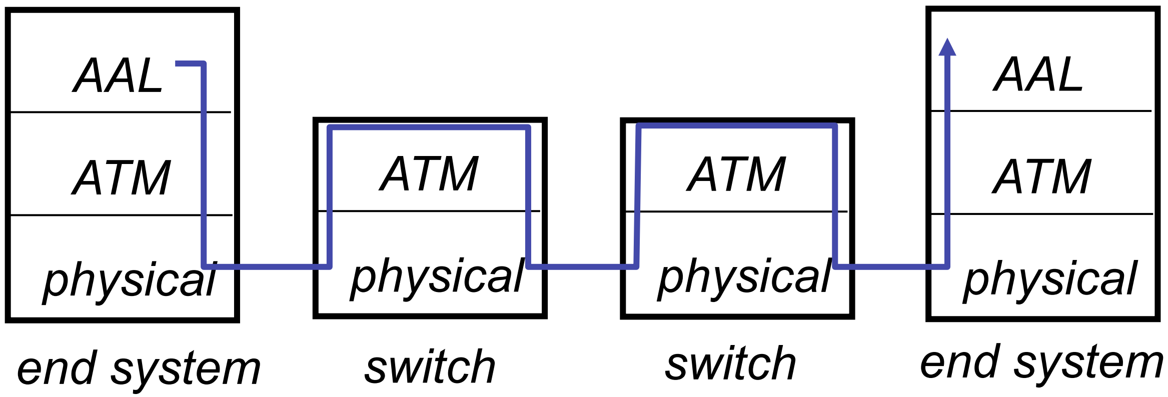

ATM Adaptation Layer (AAL) [more]

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)

User data

AAL PDU

ATM cell

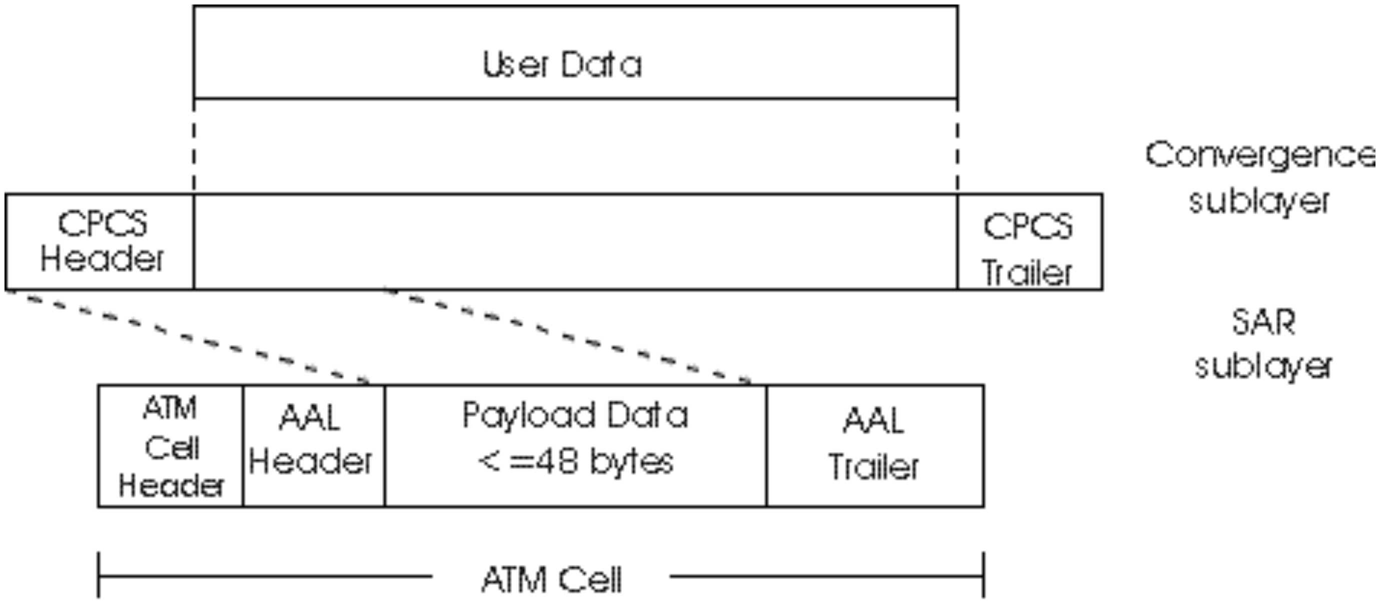

ATM Layer

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 ? Bandwidth	Loss	Order	Timing	Congestion feedback
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

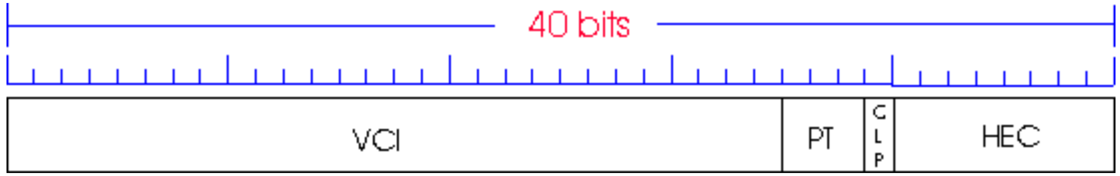

Cell format

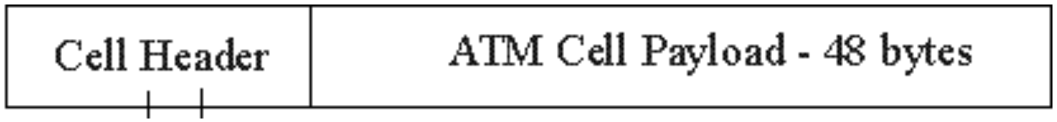

ATM cell header

- **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

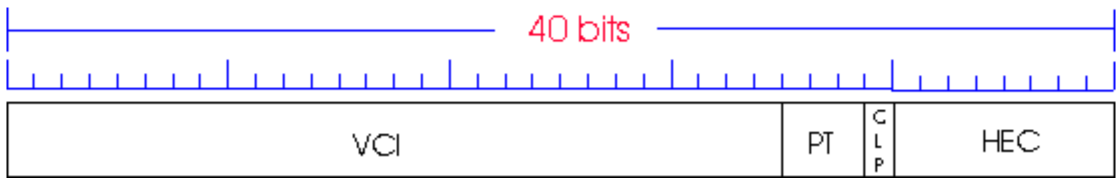

Virtual Circuit Switching

Incoming		Outgoing	
Port	VCI	Port	VCI
1	14	3	22
1	77	2	41

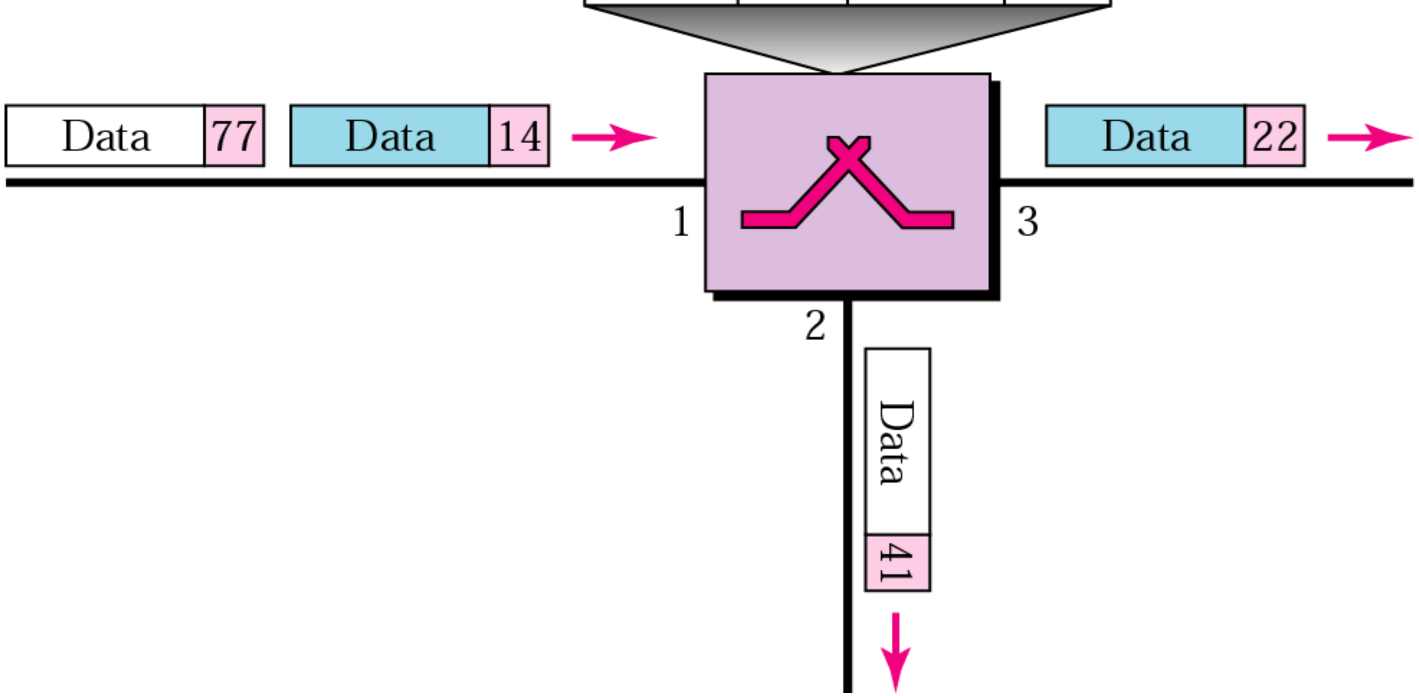

Multiplexing of Variable vs. Fixed Size Packets

- Multiplexing of variable size packets

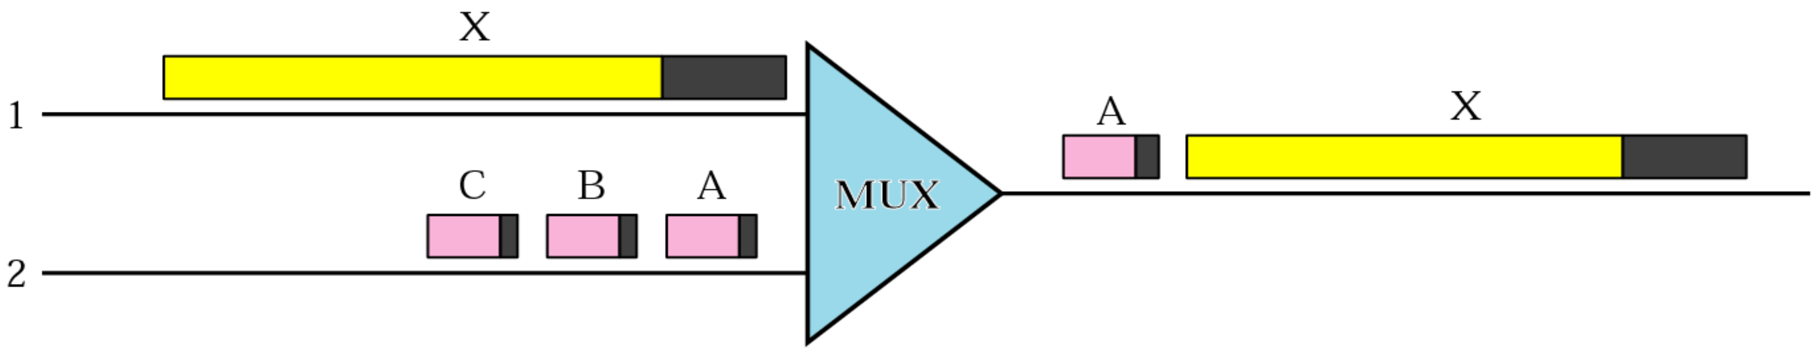

- ATM Multiplexing

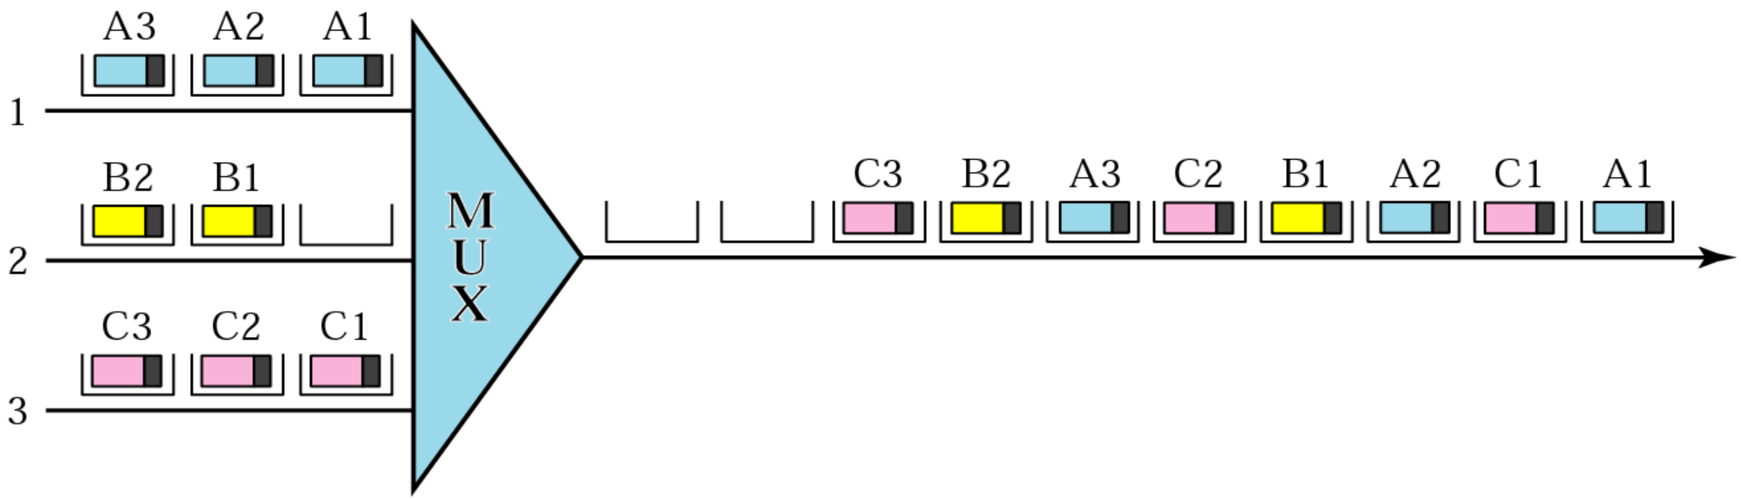

ATM Identifiers

- ATM Cell

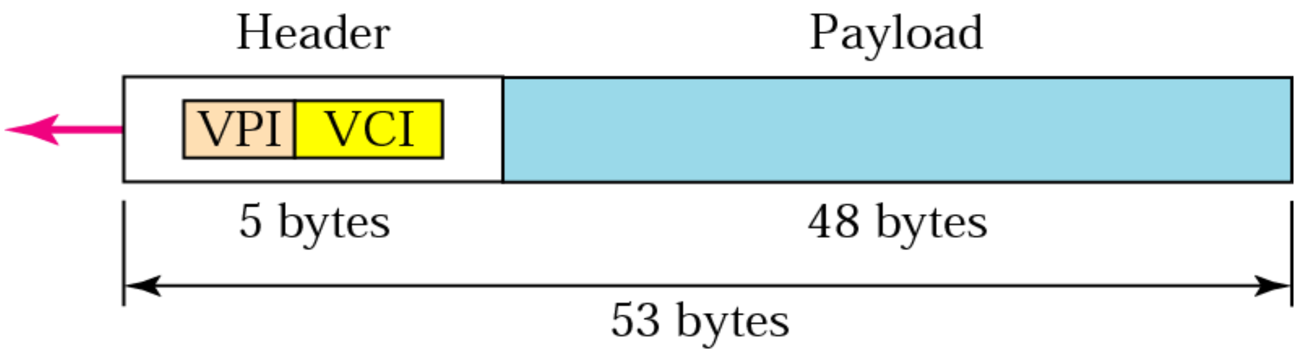

- Virtual Path Identifiers and Virtual Channel Identifiers

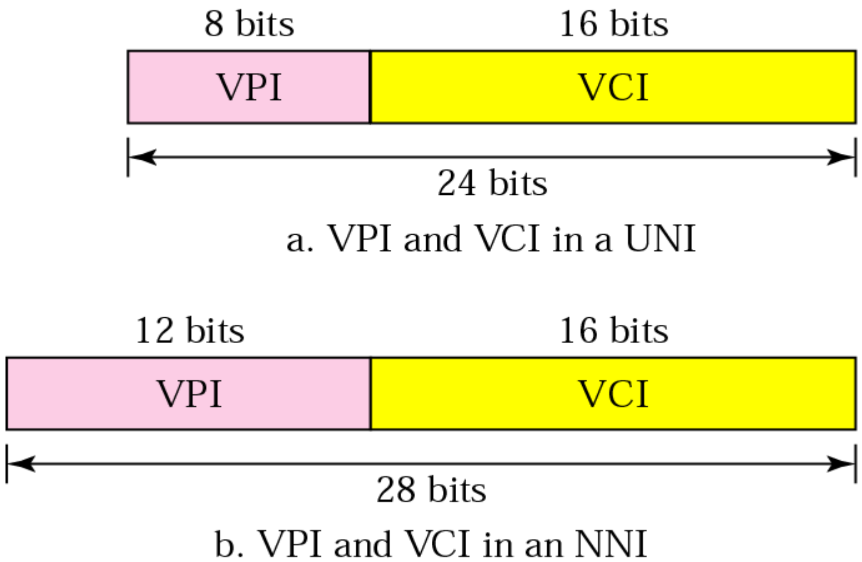

(UNI: User-to-Network-Interface
NNI: Network-to-Network-Interface)

ATM Virtual Connections

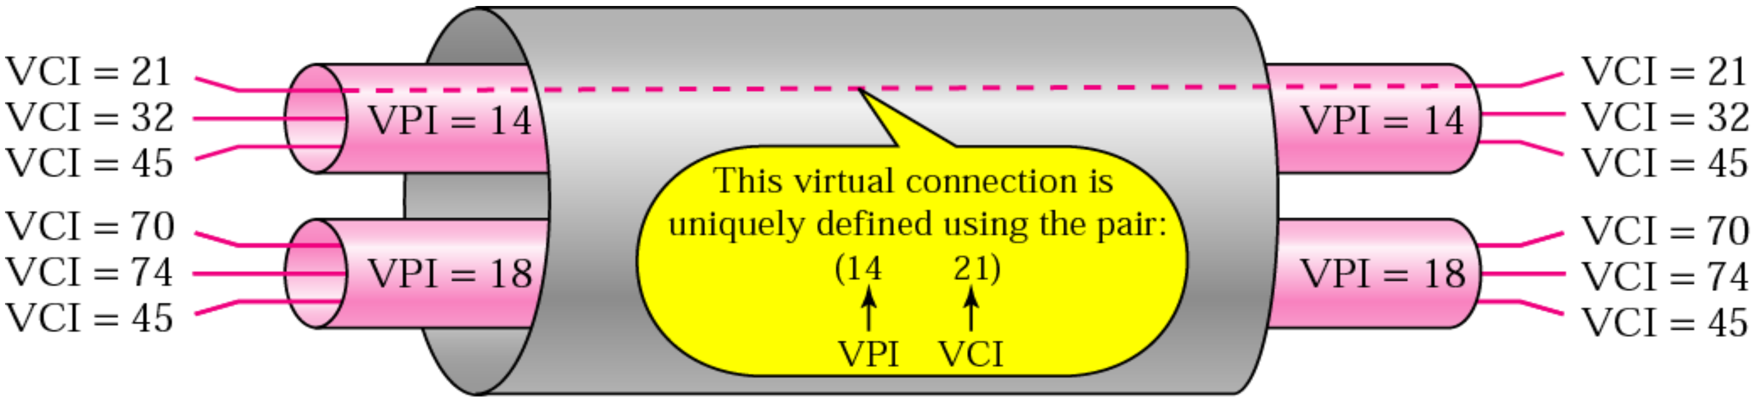

ATM Physical Layer

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

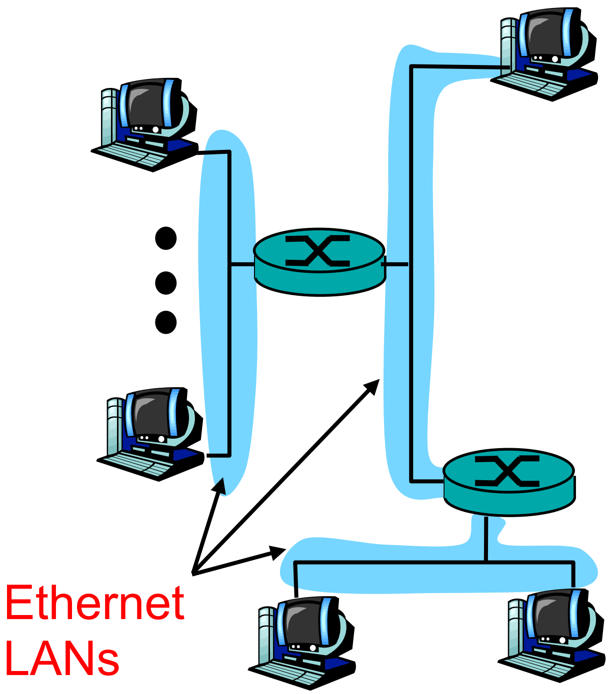

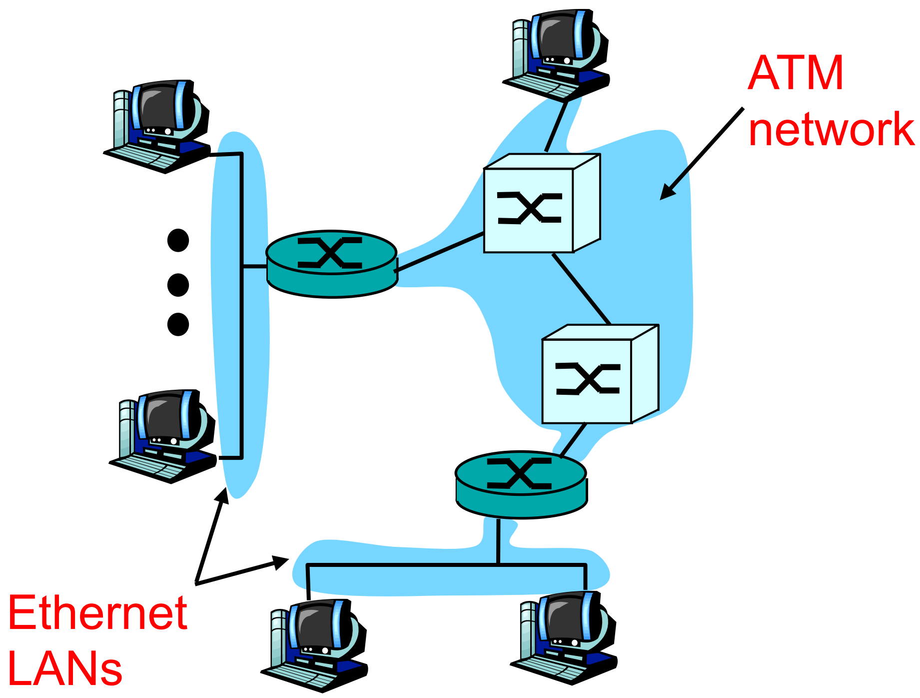

IP-Over-ATM

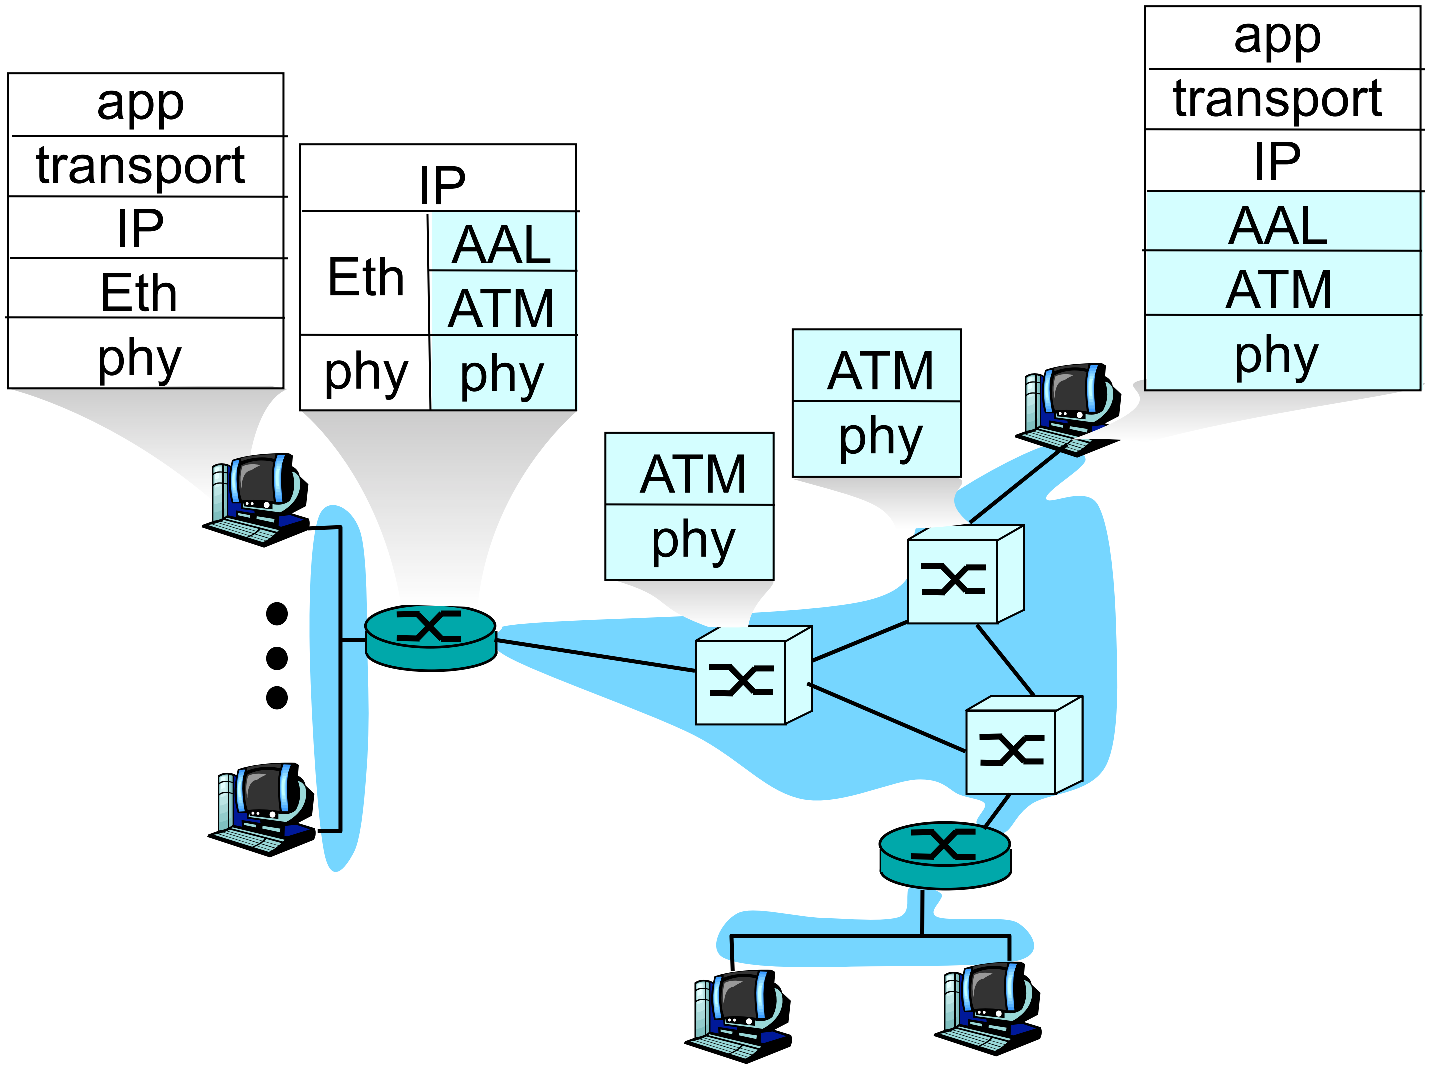
app
transport
IP
Eth
phy
IP
Eth
AAL
ATM
phy
phy
ATM
phy
ATM
phy
app
transport
IP
AAL
ATM
phy

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

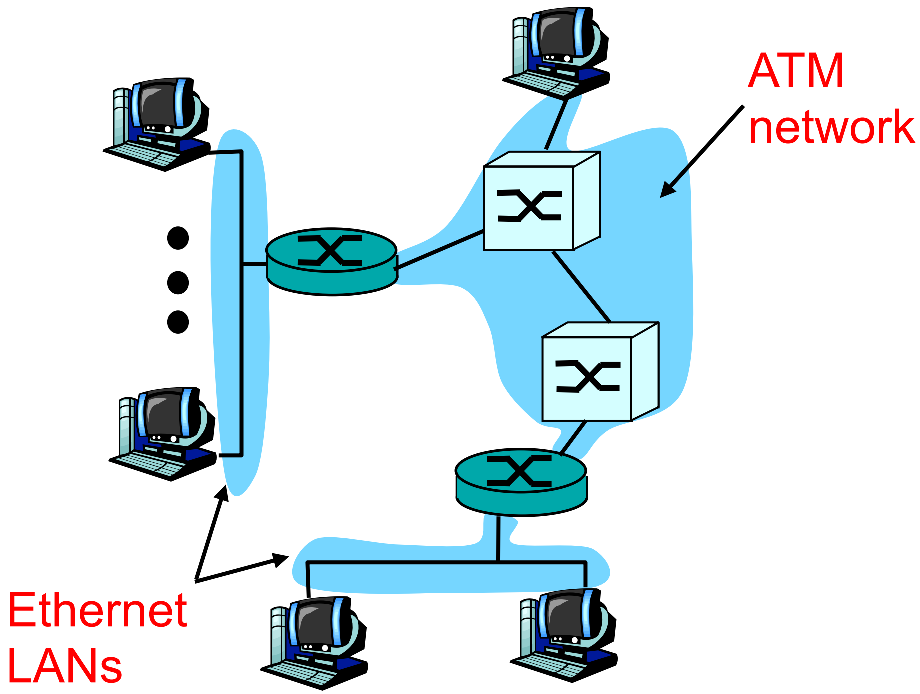

Multiprotocol label switching (MPLS)

- 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!

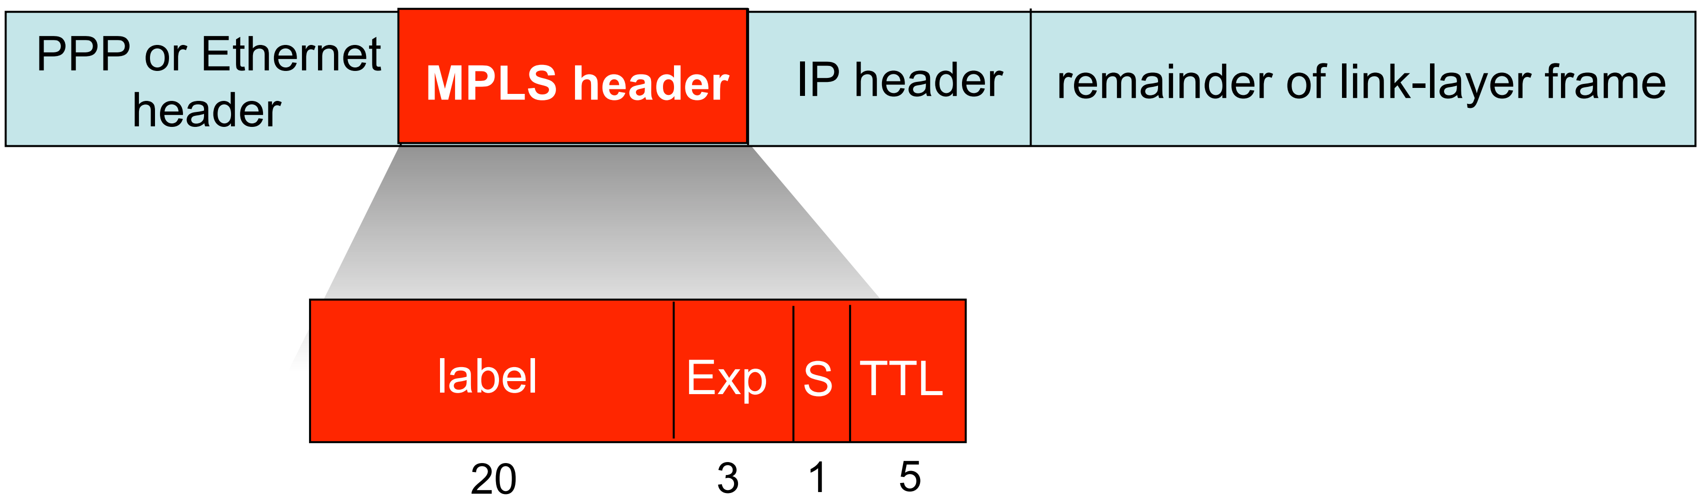

MPLS capable routers

- 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

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

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

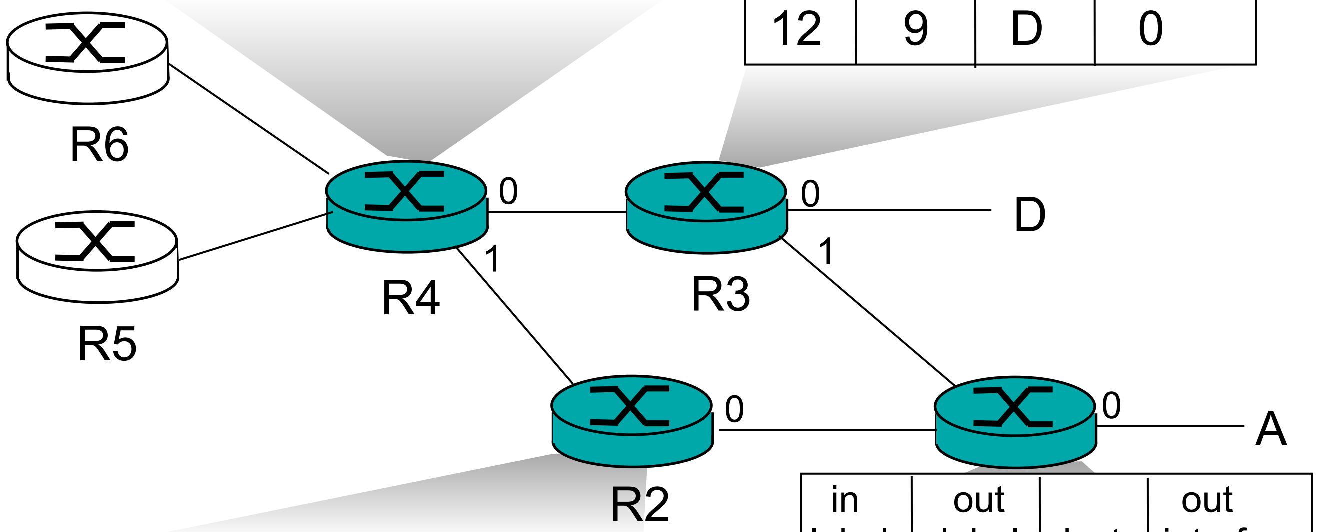

in label	out label	dest	out interface
8	6	A	0

in label	out label	dest	out interface
6	-	A	0

Chair for Network Architectures and Services – Prof. Carle
Department for Computer Science
TU München

Virtual Private Networks

Technische Universität München

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

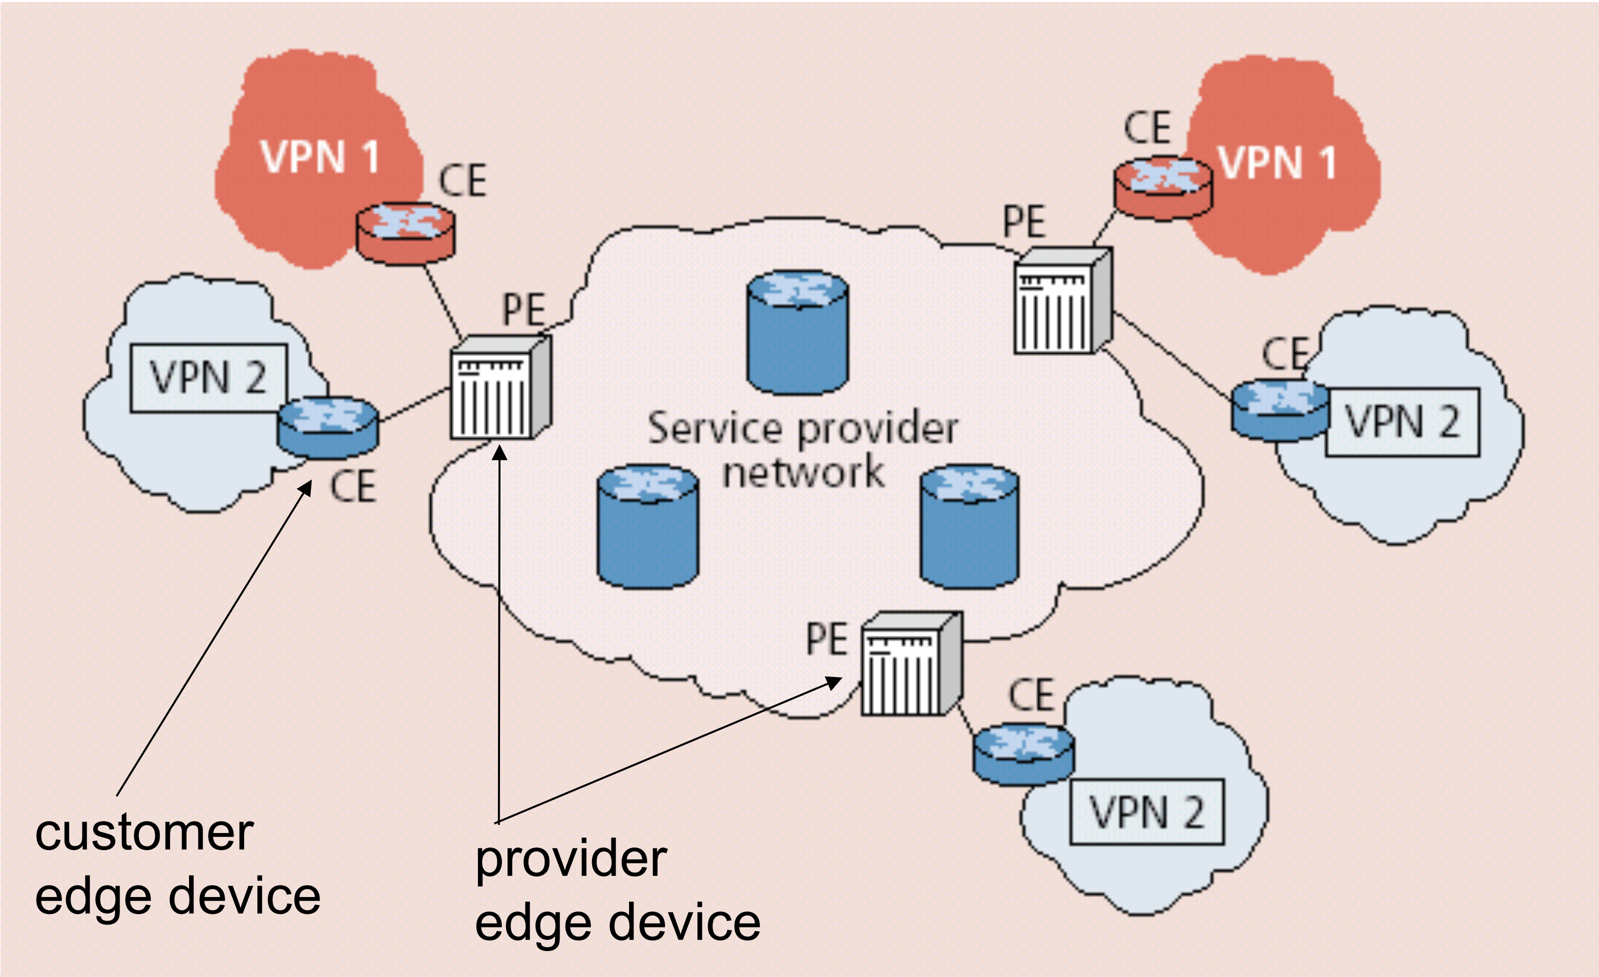

VPNs: Why?

- 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

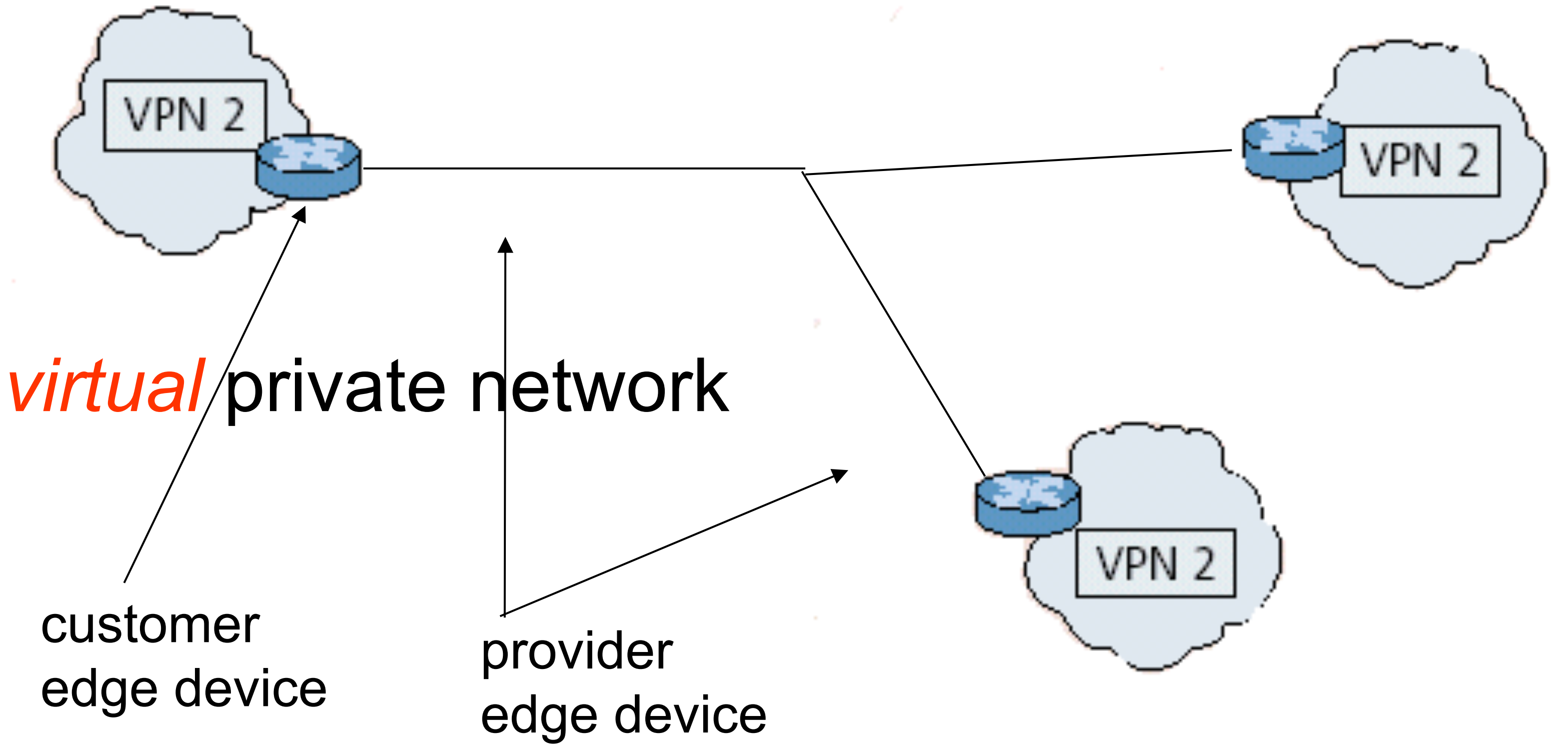

Leased-Line VPN

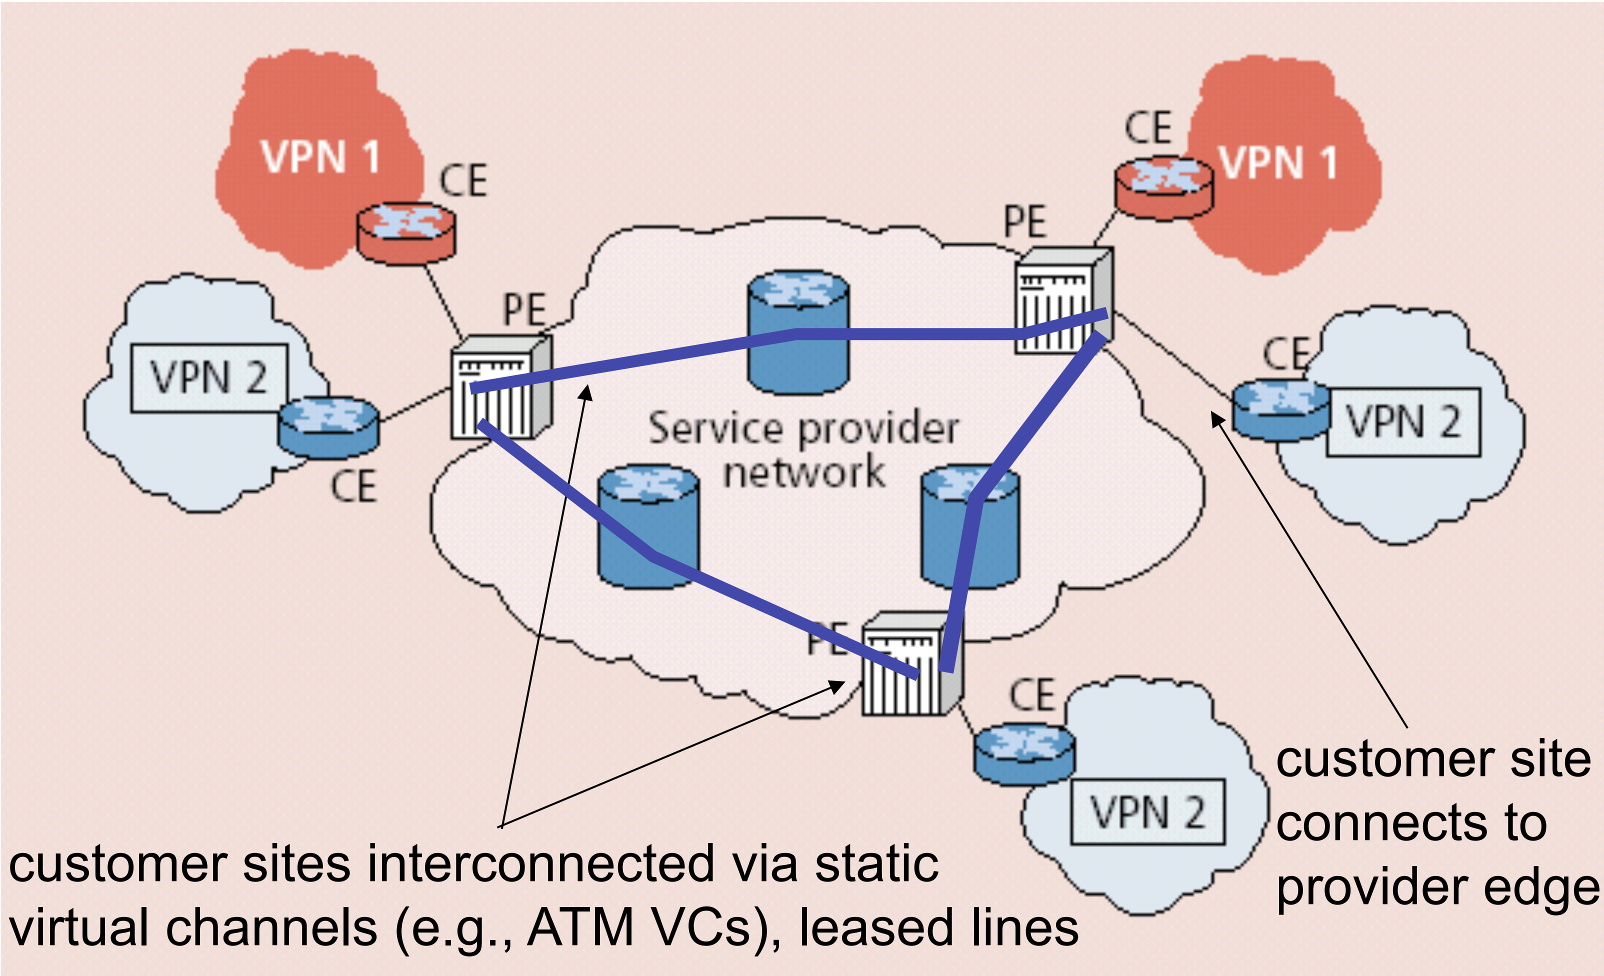

Customer Premise VPN

- all VPN functions implemented by customer

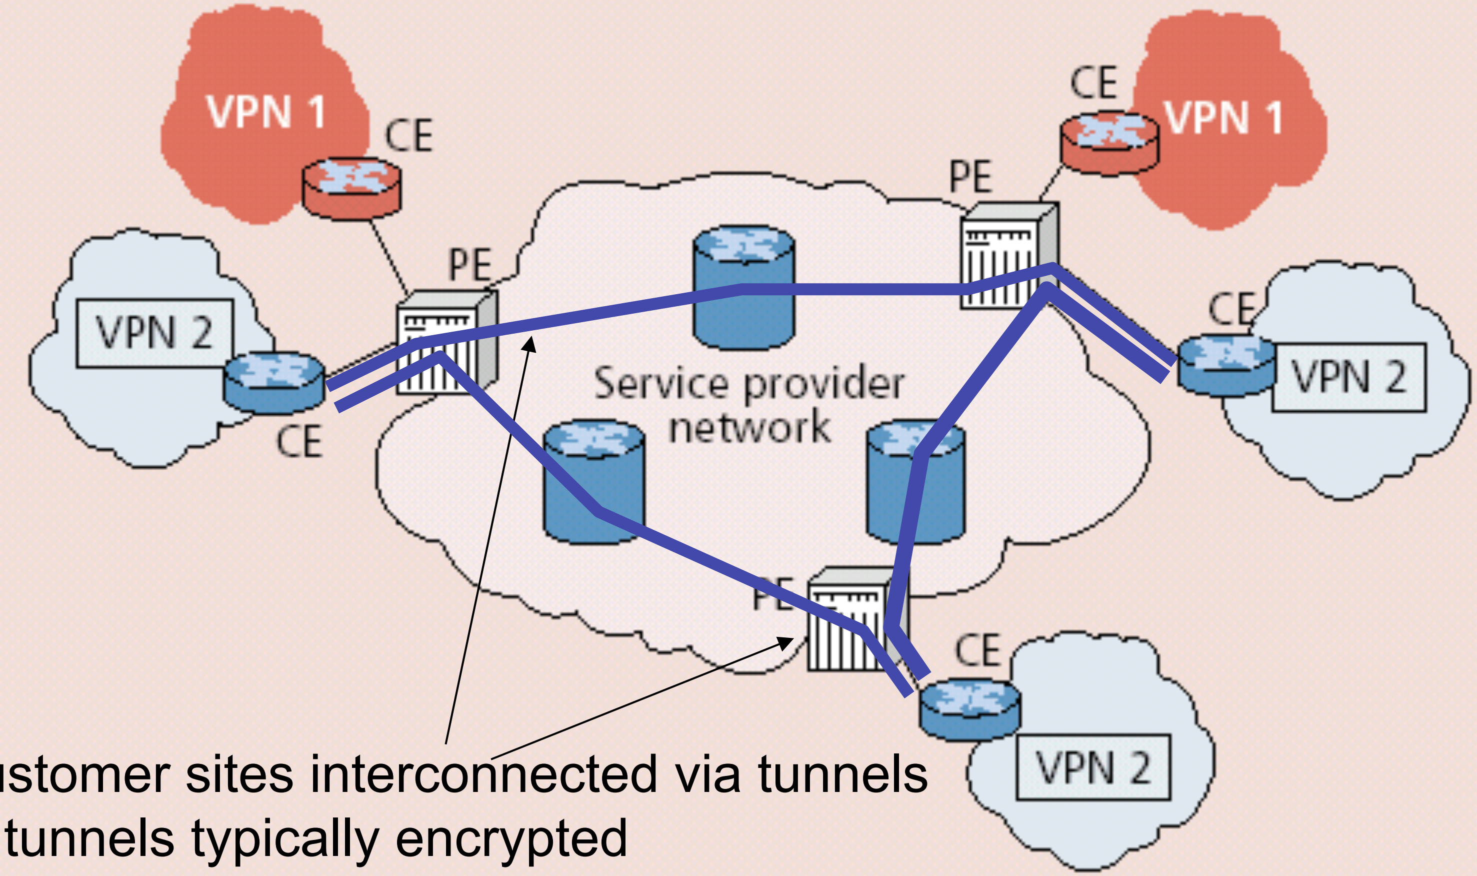

customer sites interconnected via tunnels

- 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

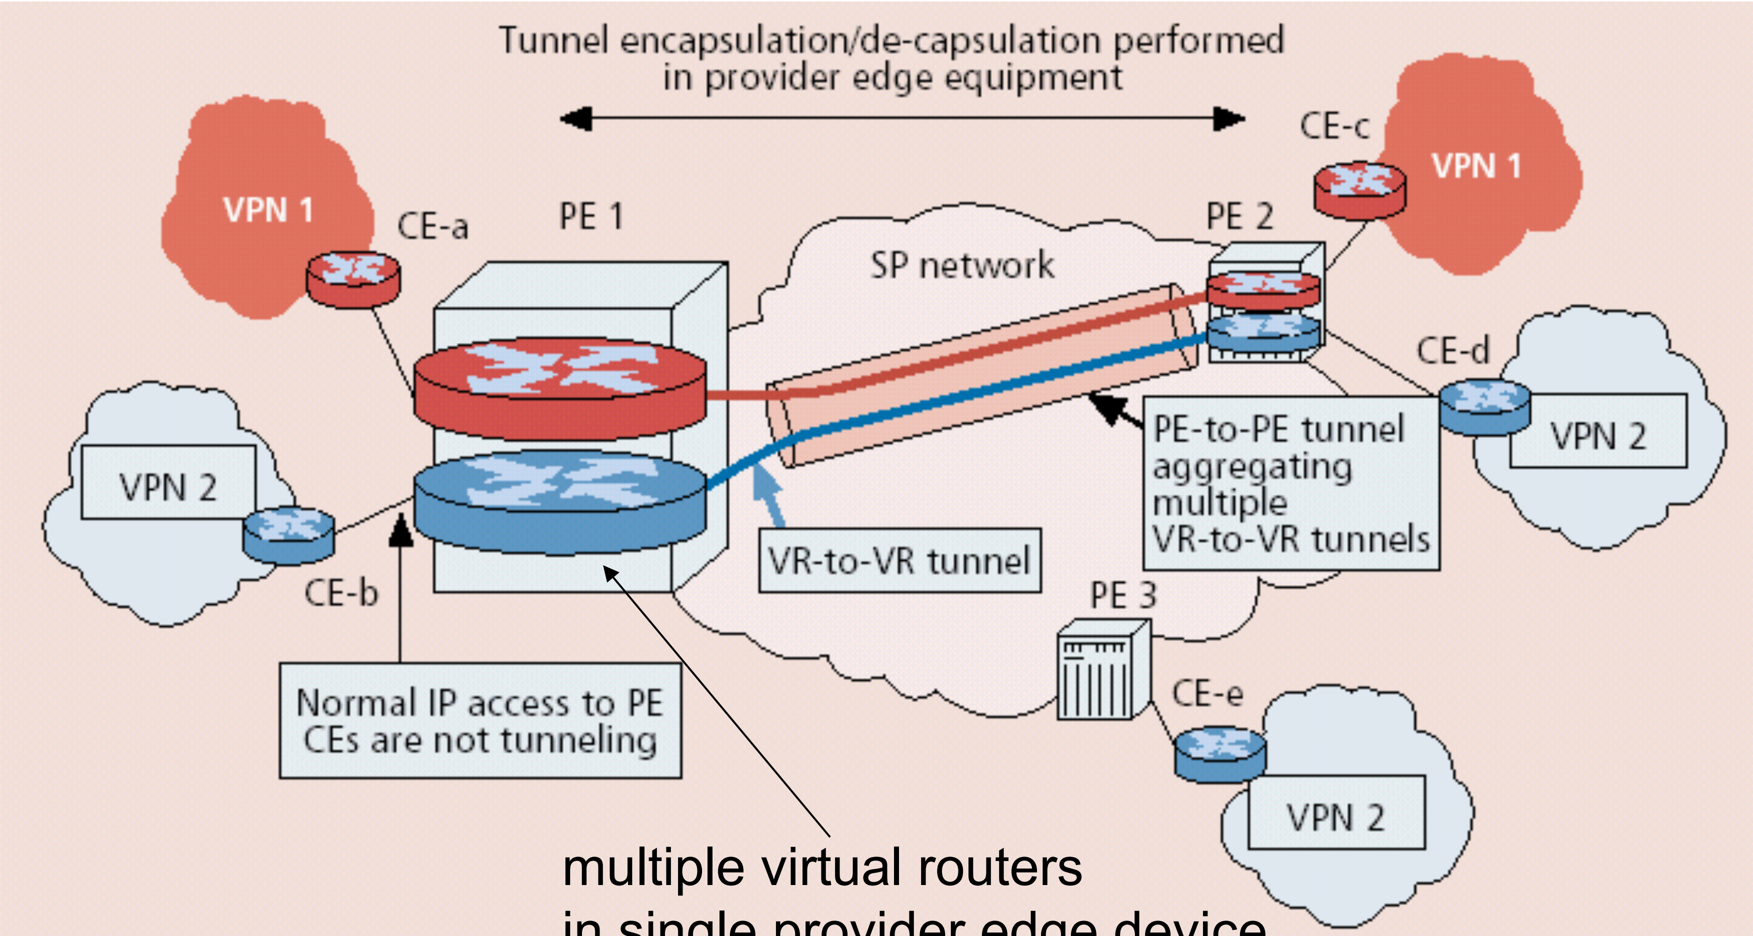

multiple virtual routers
in single provider edge device

Tunneling

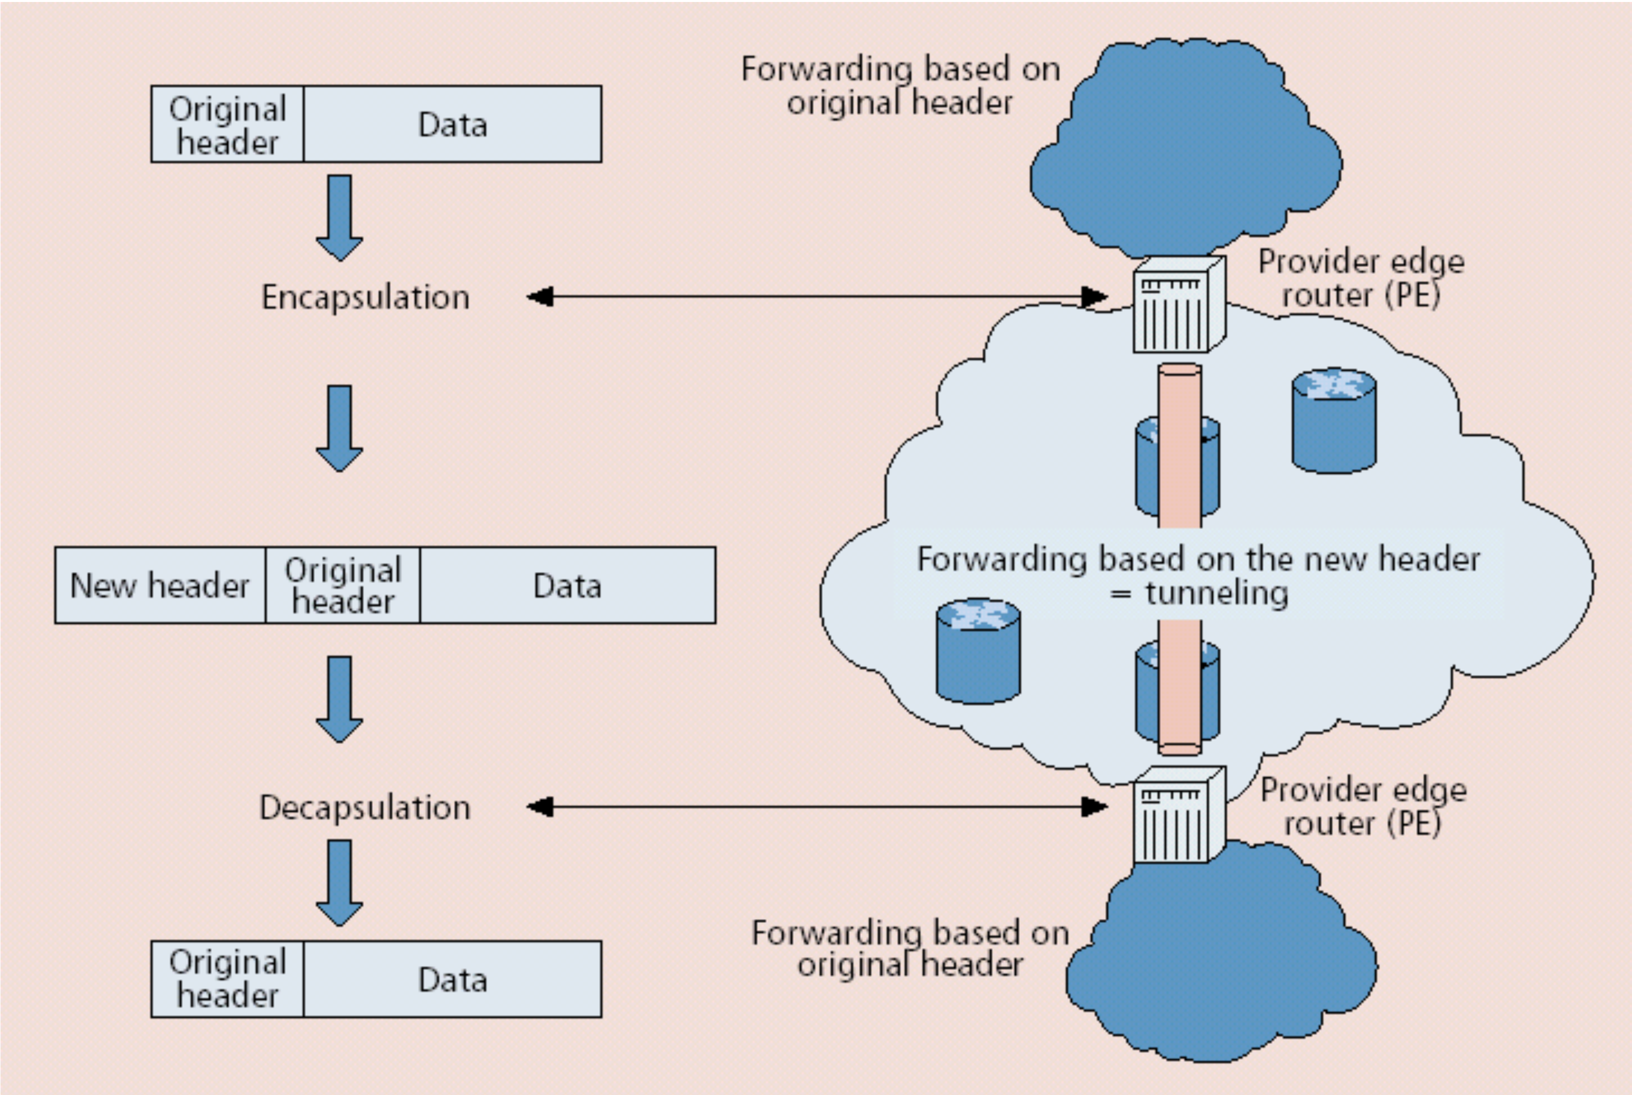

Original header
Data
Encapsulation
New header
Original header
Data
Decapsulation
Original header
Data
Forwarding based on original header
Provider edge router (PE)
Forwarding based on the new header = tunneling
Provider edge router (PE)
Forwarding based on original header

MPLS-based VPN