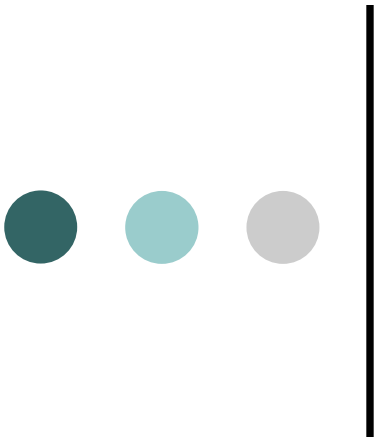

CS419: Computer Networks

Lecture 9: Mar 30, 2005

VPNs

VPN Taxonomy

CS419

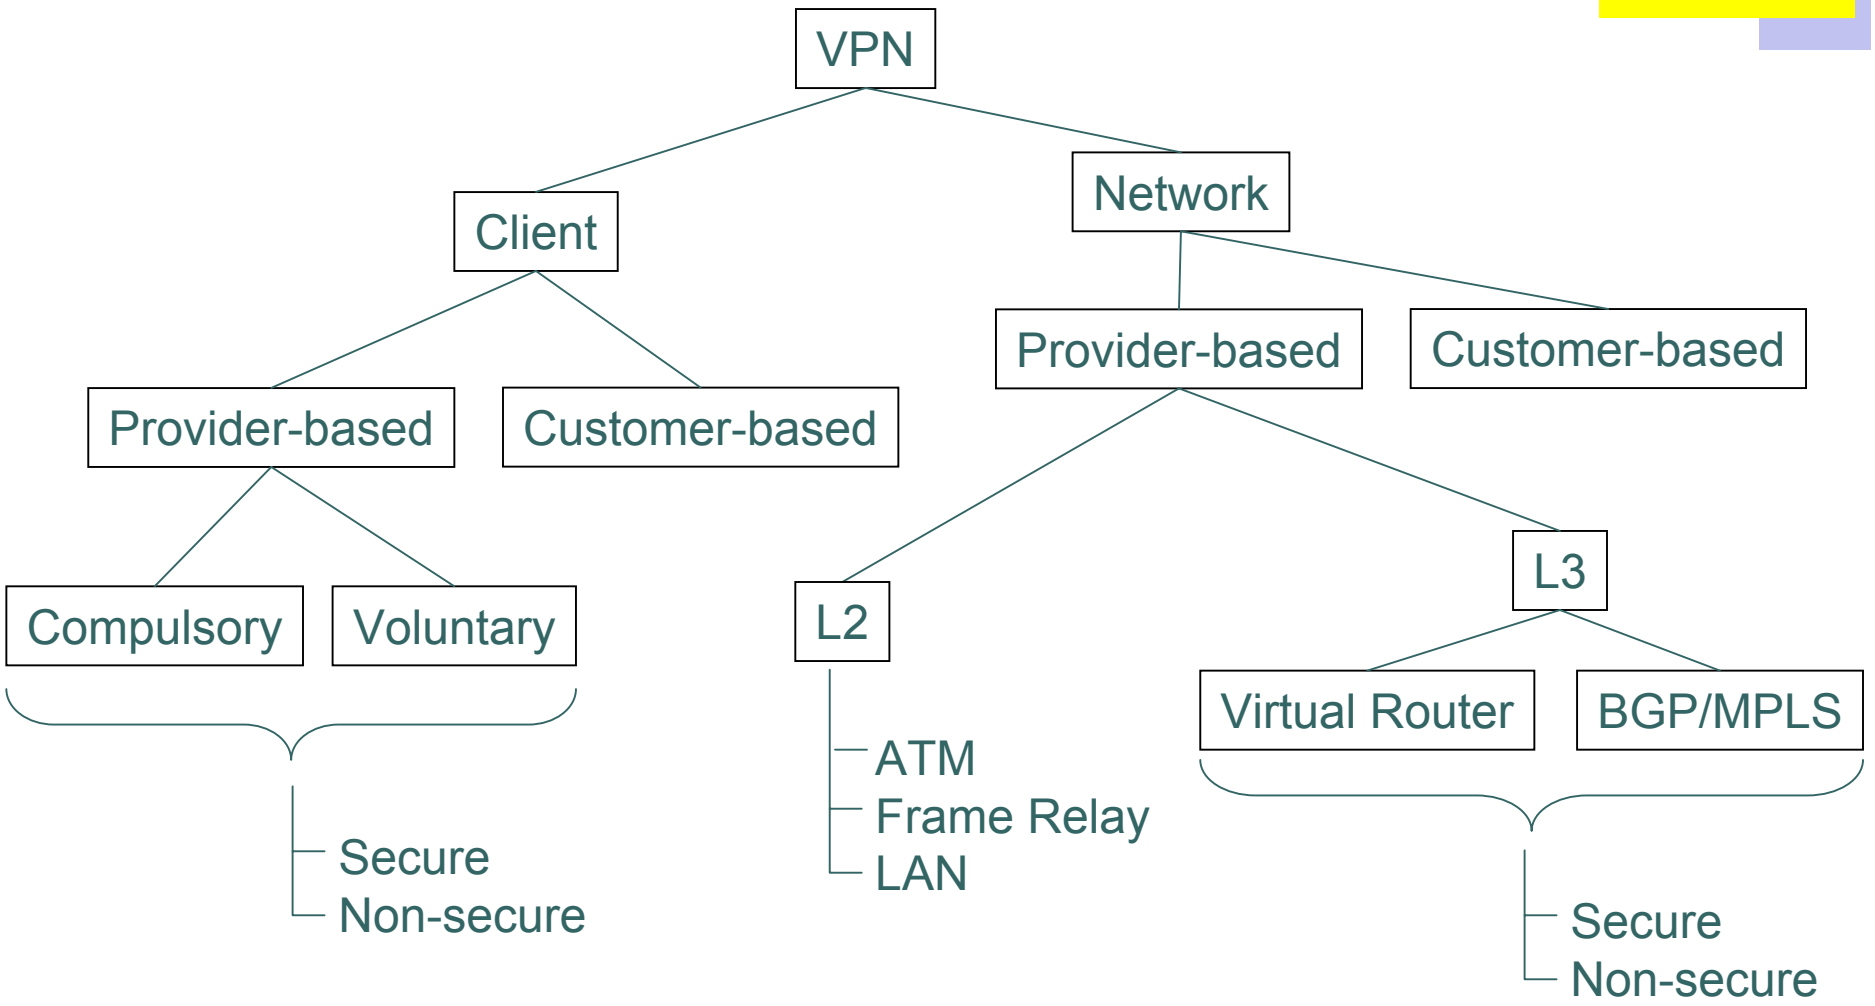

What is a VPN?

CS419

- Making a shared network look like a private network
- Why do this?
  - Private networks have all kinds of advantages
    - (we'll get to that)
  - But building a private network is expensive
    - (cheaper to have shared resources rather than dedicated)

History of VPNs

CS419

- Originally a telephone network concept
  - Separated offices could have a phone system that looked like one internal phone system
- Benefits?
  - Fewer digits to dial
  - Could have different tariffs
    - Company didn't have to pay for individual long distance calls
  - Came with own blocking probabilities, etc.
    - Service guarantees better (or worse) than public phone service

Original data VPNs

CS419

- Lots of different network technologies in those days
  - Decnet, Appletalk, SNA, XNS, IPX, …
  - None of these were meant to scale to global proportions
  - Virtually always used in corporate settings
- Providers offer virtual circuits between customer sites
  - Frame Relay or ATM
  - A lot cheaper than dedicated leased lines
- Customer runs whatever network technology over these
- These still exist (but being replaced by IP VPNs)

VPN Taxonomy

CS419

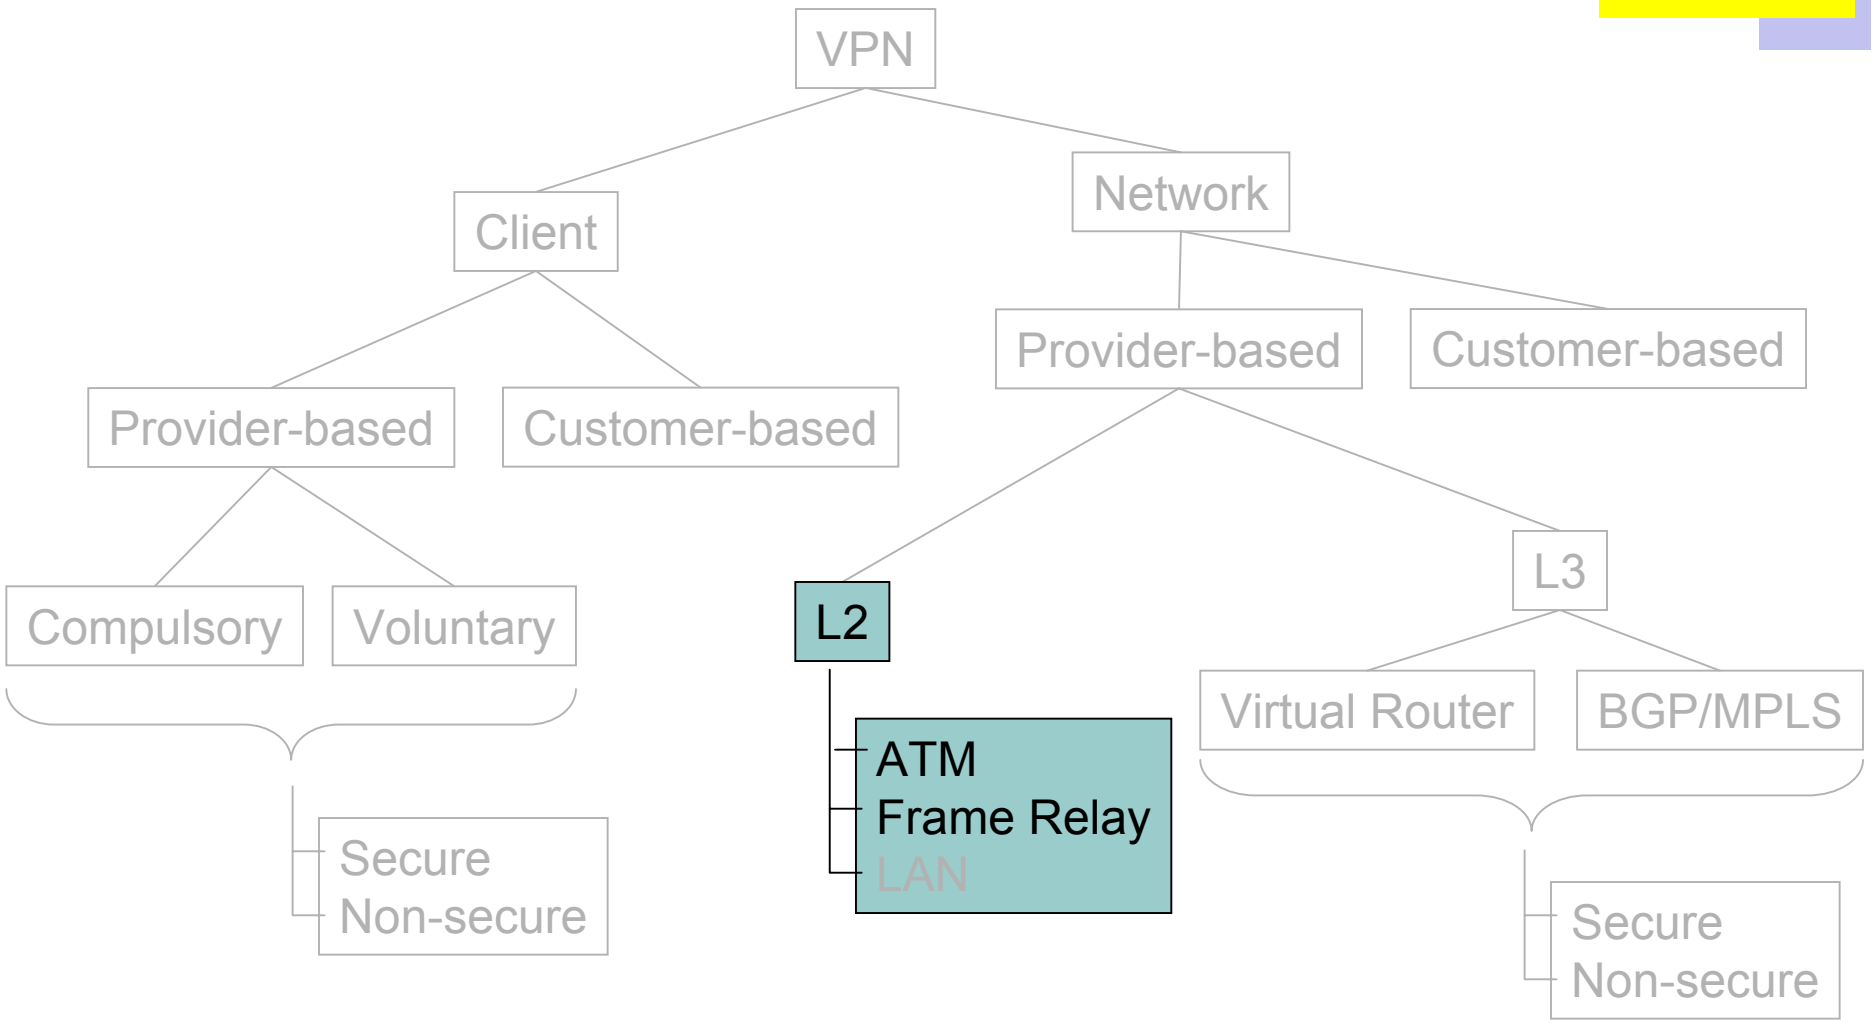

Advantages of original data VPNs

CS419

- Repeat: a lot cheaper than dedicated leased lines
  - Corporate users had no other choice
  - This was the whole business behind frame-relay and ATM services
- Fine-grained bandwidth tariffs
- Bandwidth guarantees
  - Service Level Agreements (SLA)
- “Multi-protocol”

Frame Relay VPN Example

CS419

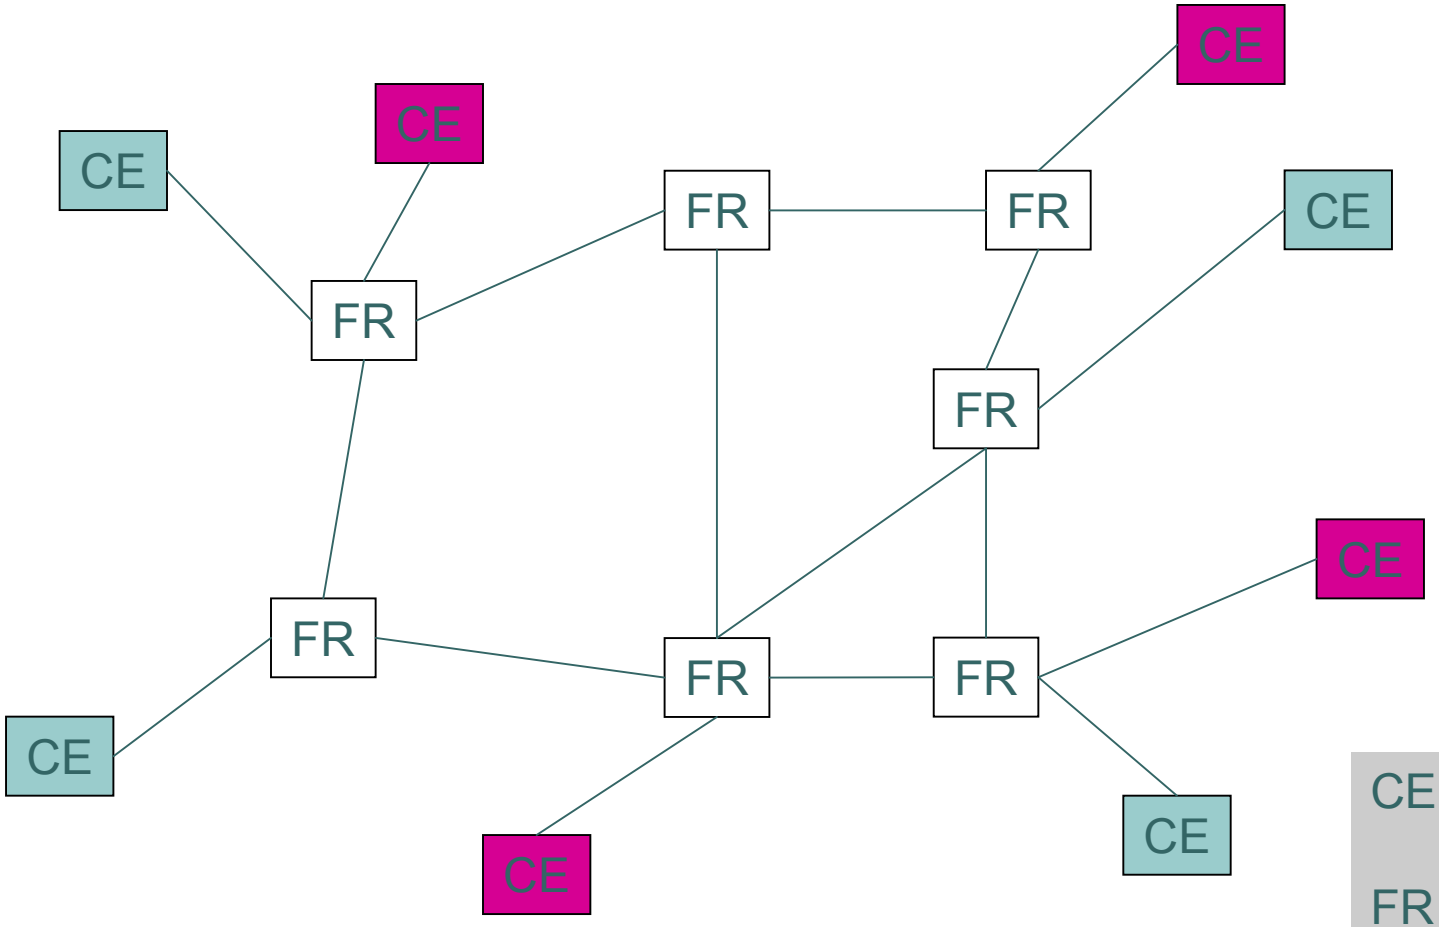

CE = Customer Equipment
FR = Frame Relay

Define circuits CE to CE (for given customer: purple)

CS419

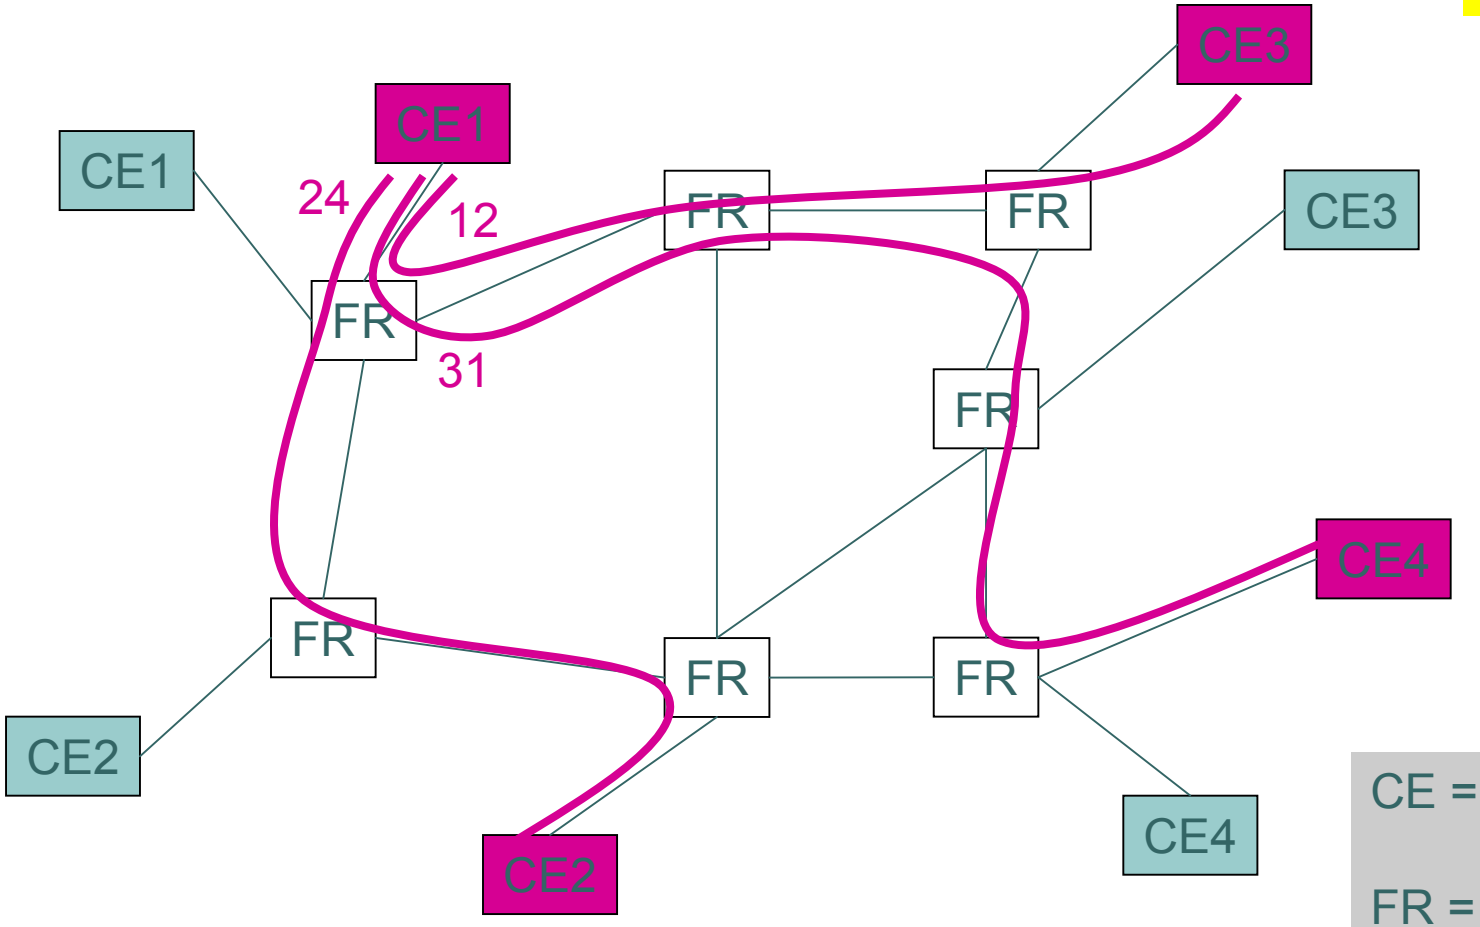

CE = Customer Equipment
FR = Frame Relay

Customer establishes routing tables (per protocol)

CS419

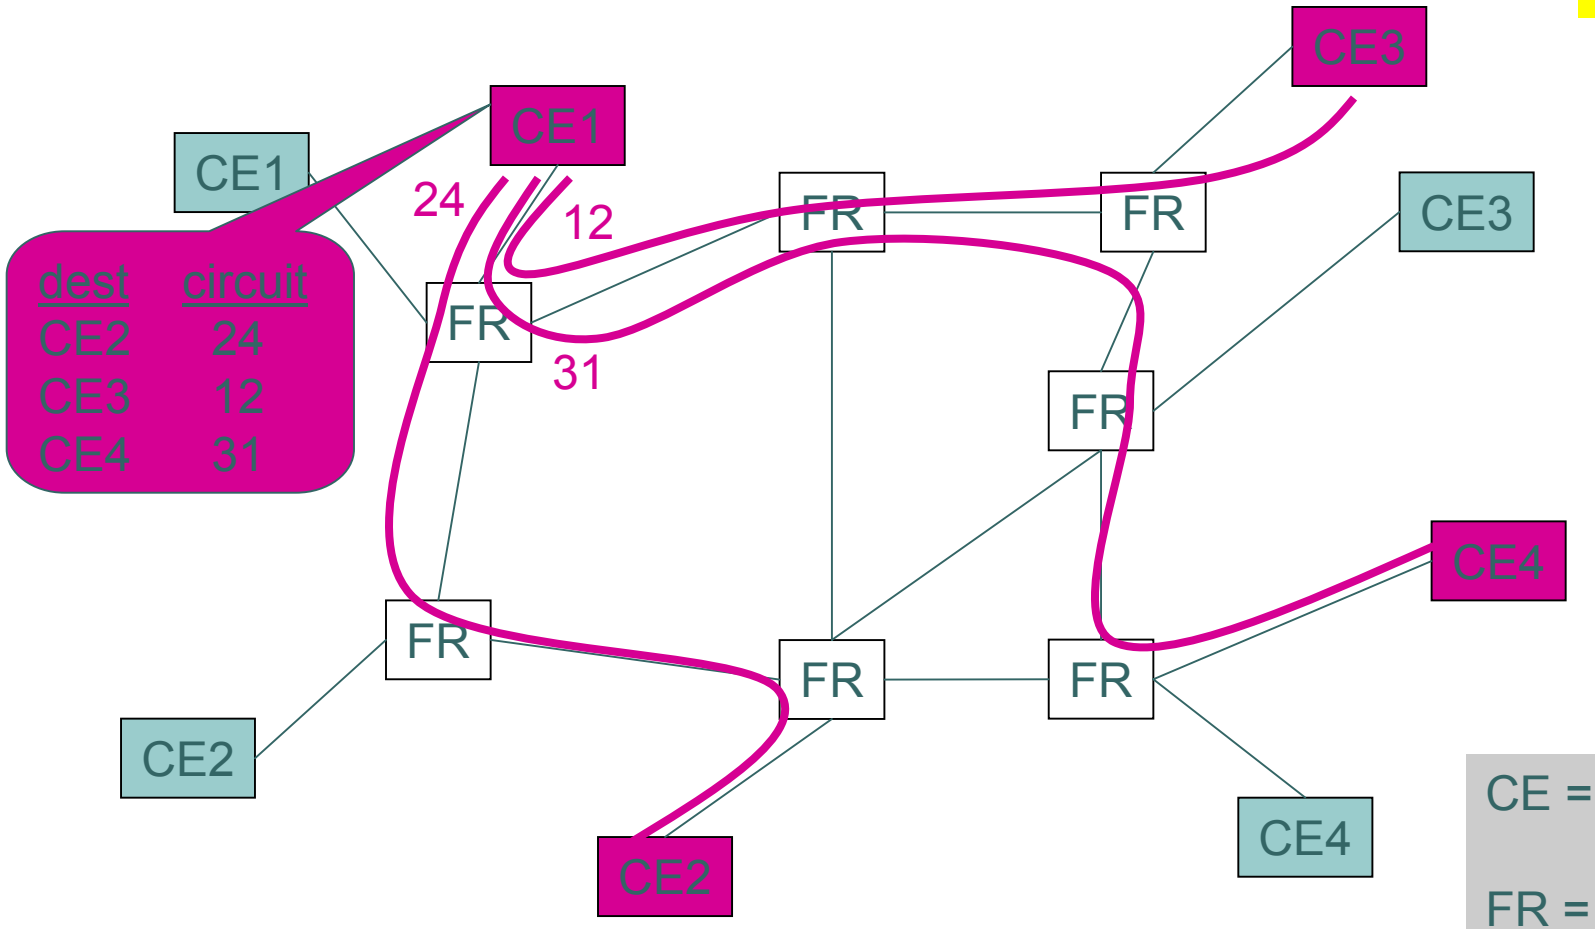

Provider provisions underlying network

CS419

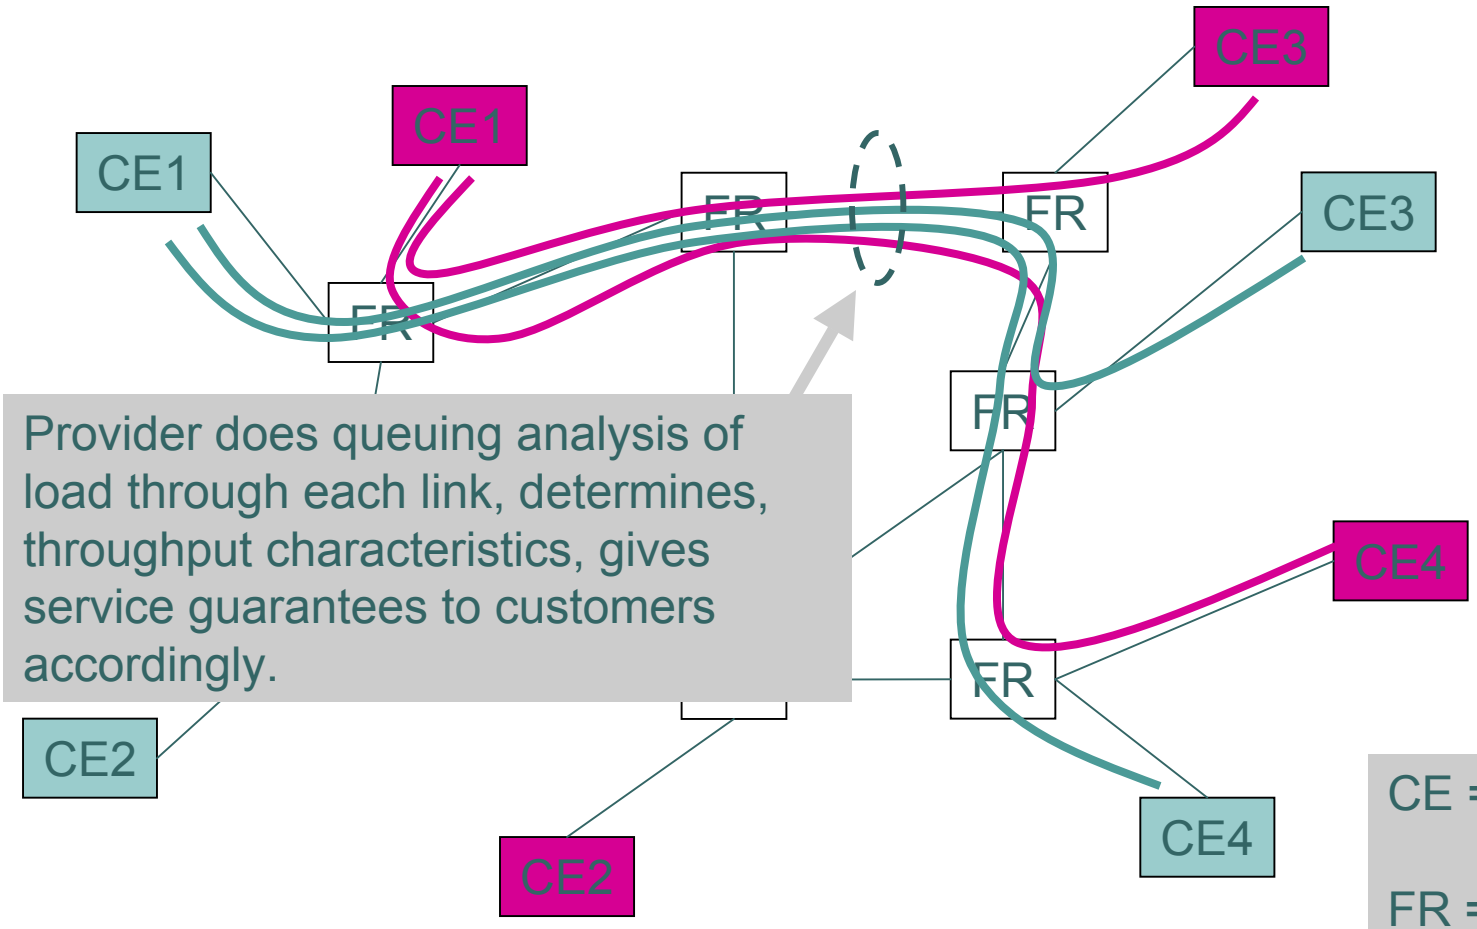

How has the world changed?

CS419

- Everything is IP now
  - Some old stuff still around, but most data networks are just IP
- So, why do we still care about VPNs???

IP VPN benefits

CS419

- IP not really global (private addresses)
  - VPN makes separated IP sites look like one private IP network
- Security
- Bandwidth guarantees across ISP
  - QoS, SLAs
- Simplified network operation
  - ISP can do the routing for you

Client VPNs

CS419

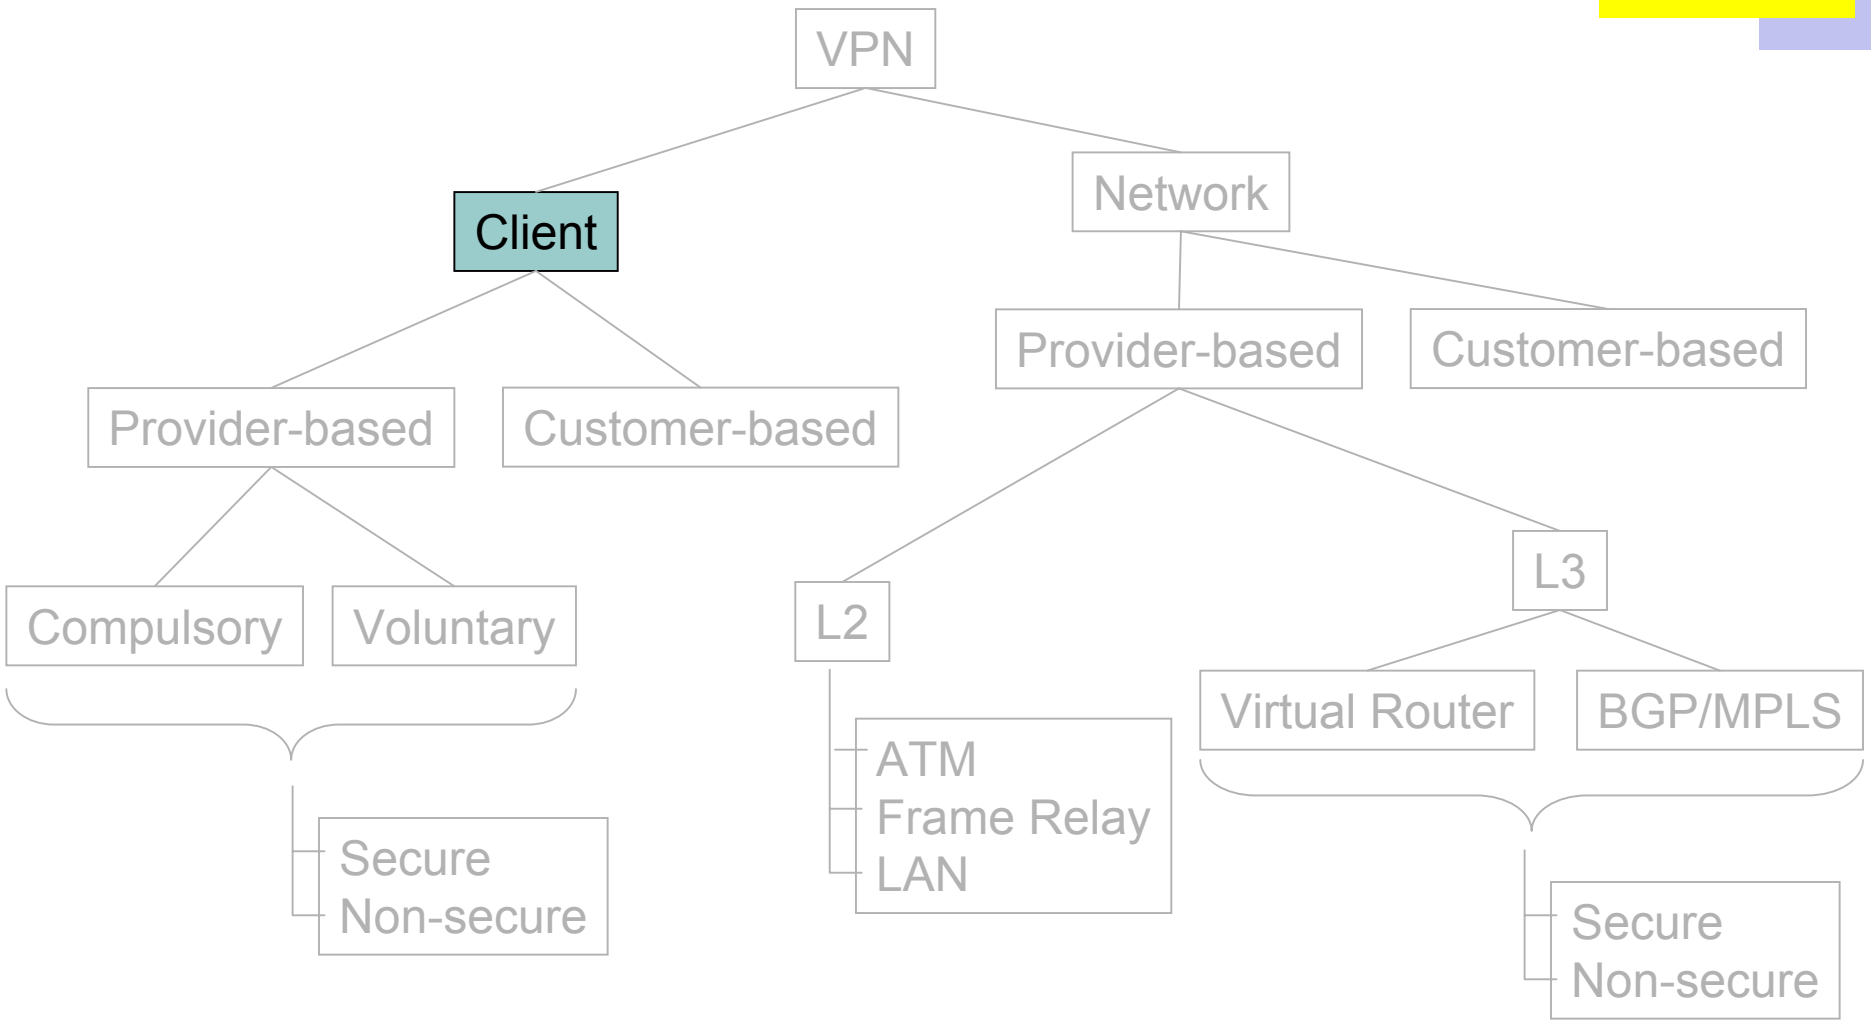

Client VPNs

CS419

- Solves problem of how to connect remote hosts to a firewalled network
  - Security and private addresses benefits only
  - Not simplicity or QoS benefits

Client VPNs

CS419

- Solves problem of how to connect remote hosts to a firewalled network

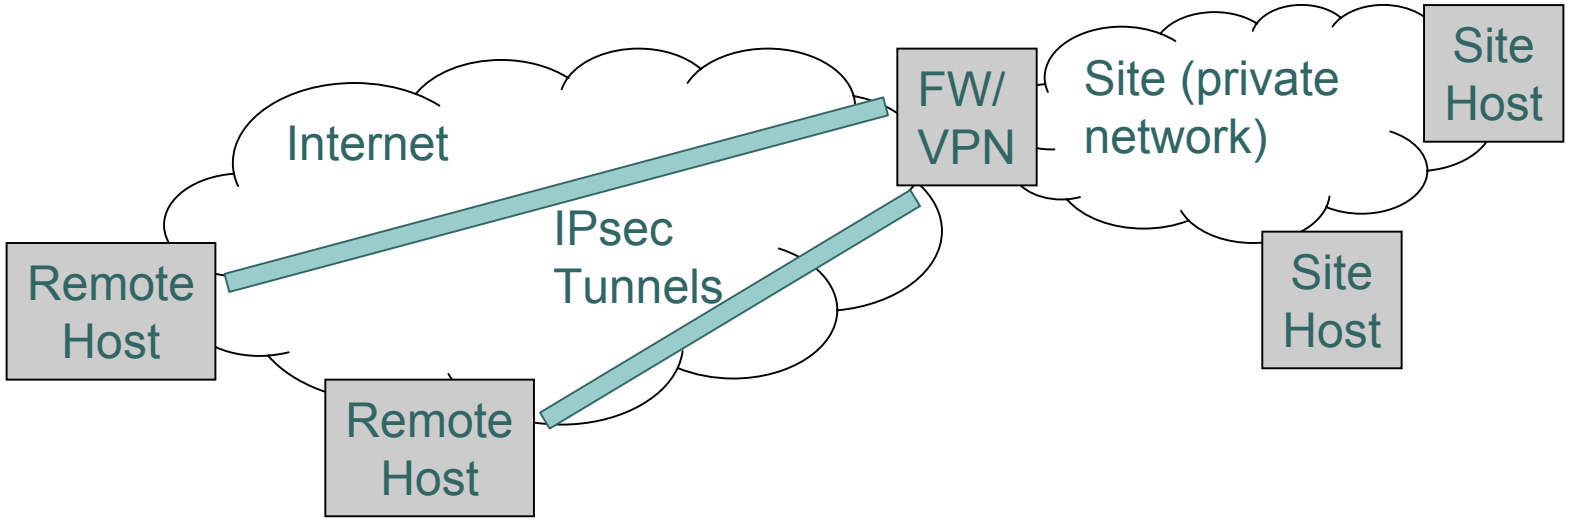

Client VPNs: Configuration

CS419

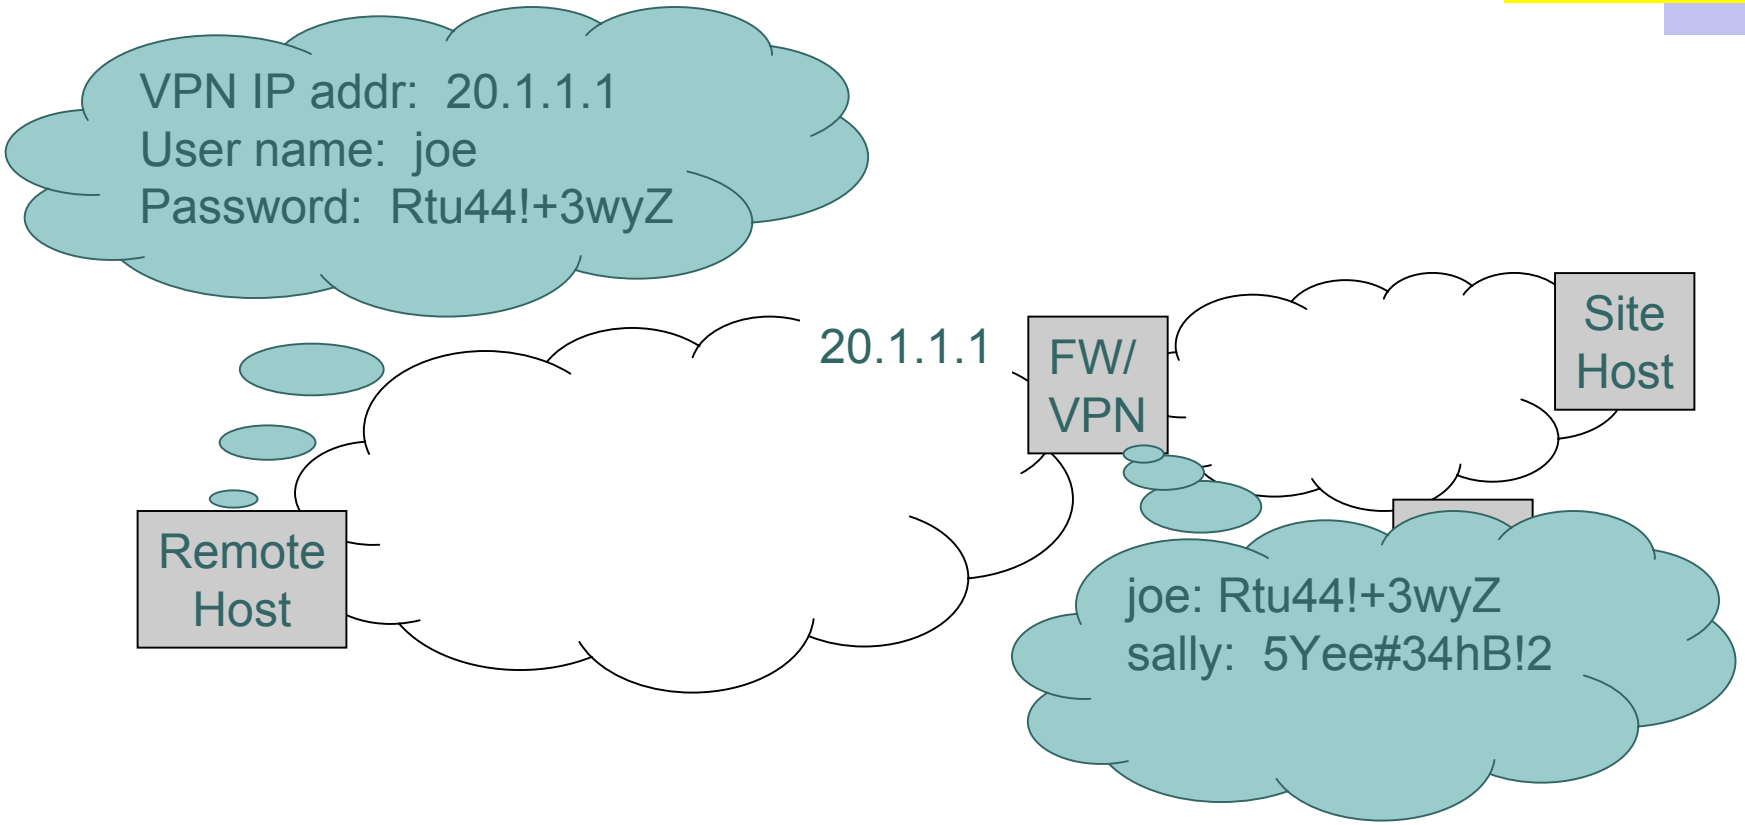

Client VPNs

CS419

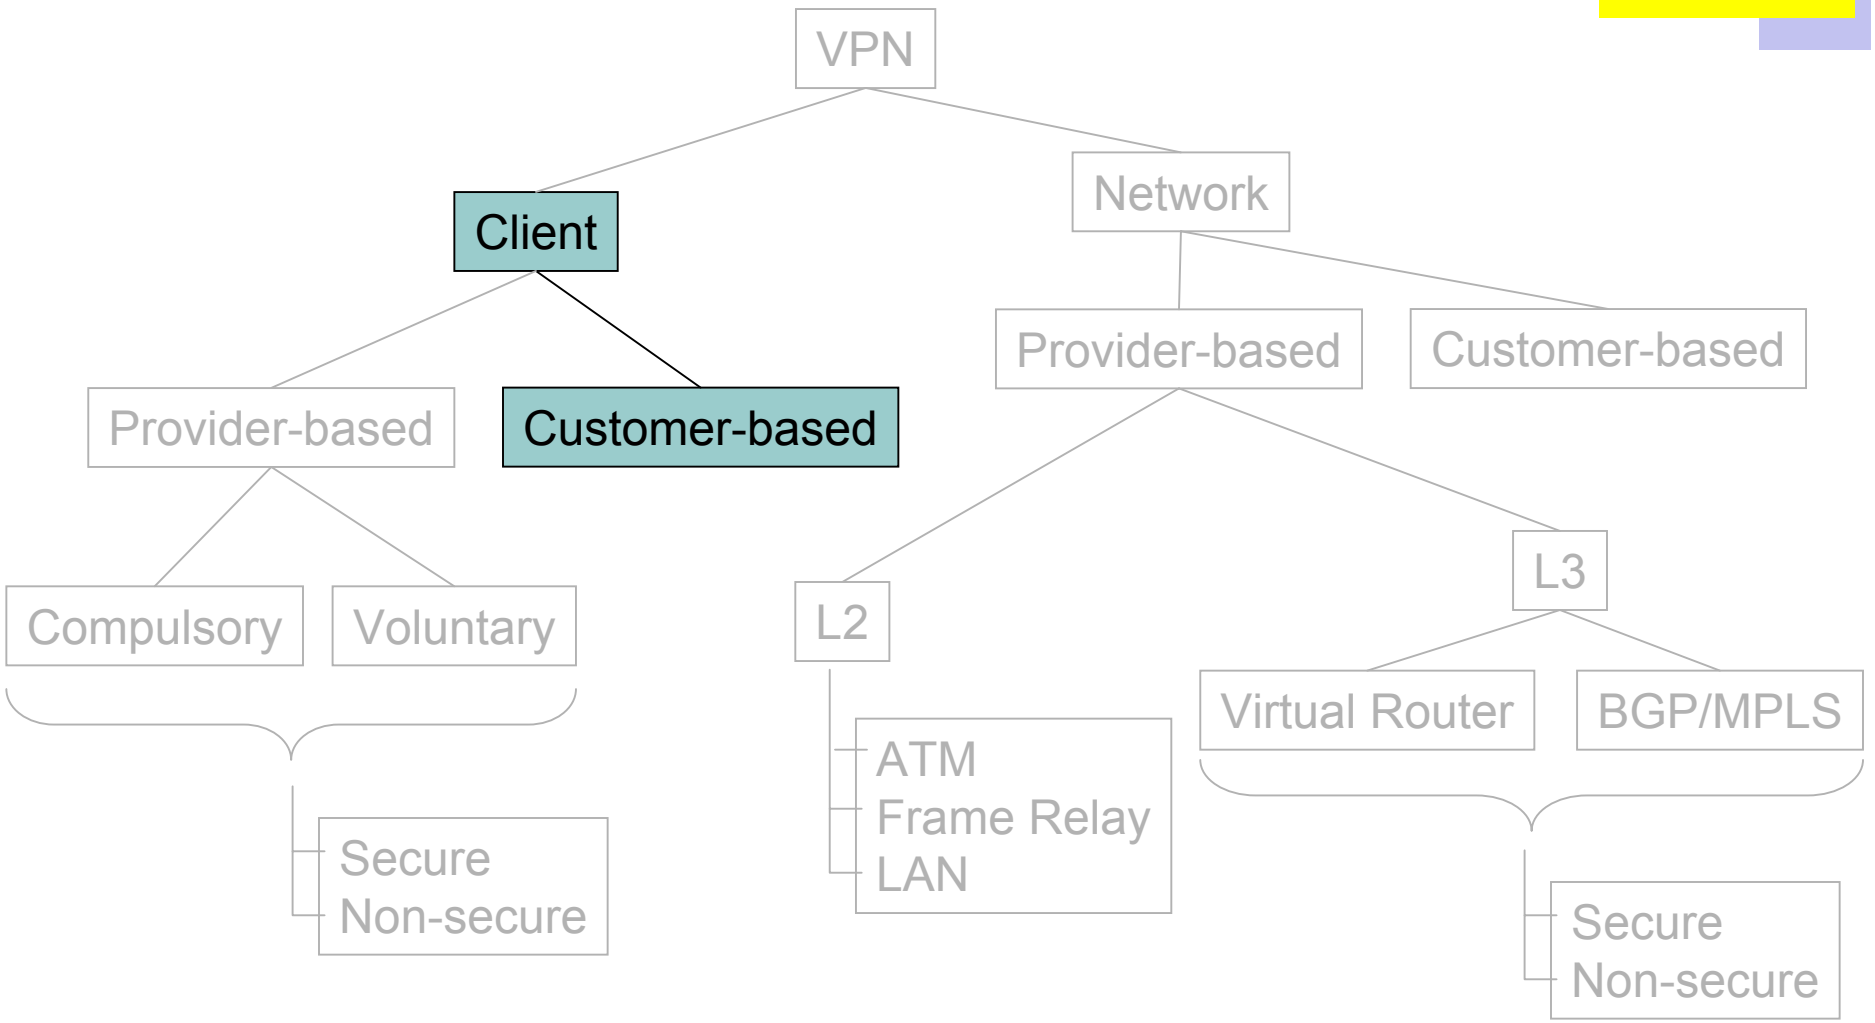

Client VPNs: Configuration

CS419

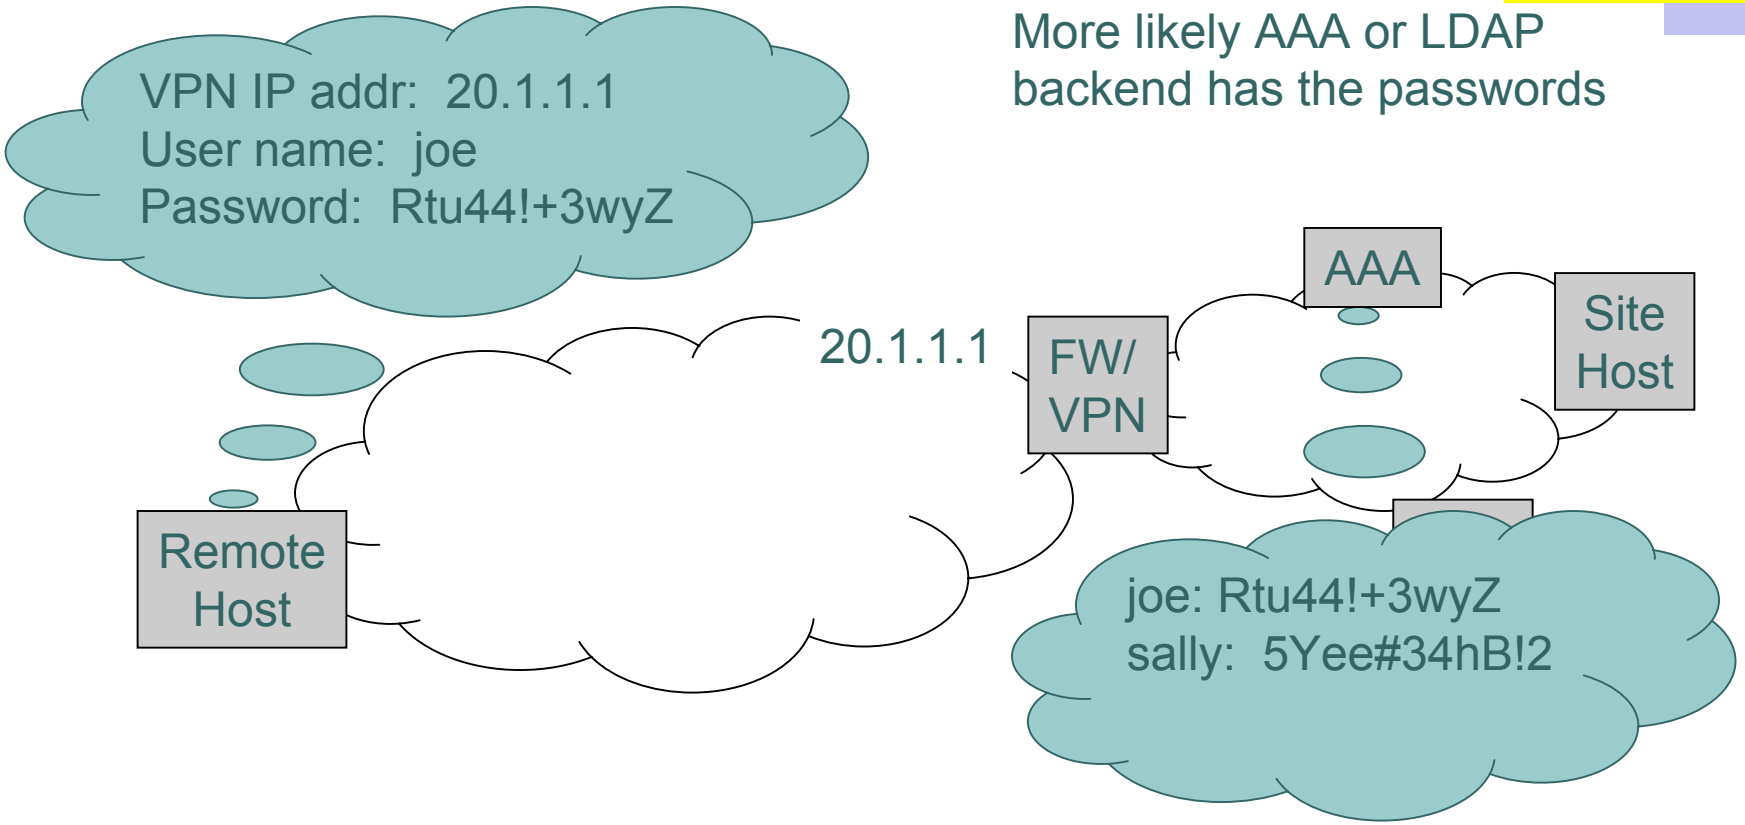

Client VPNs: Host gets local IP address

CS419

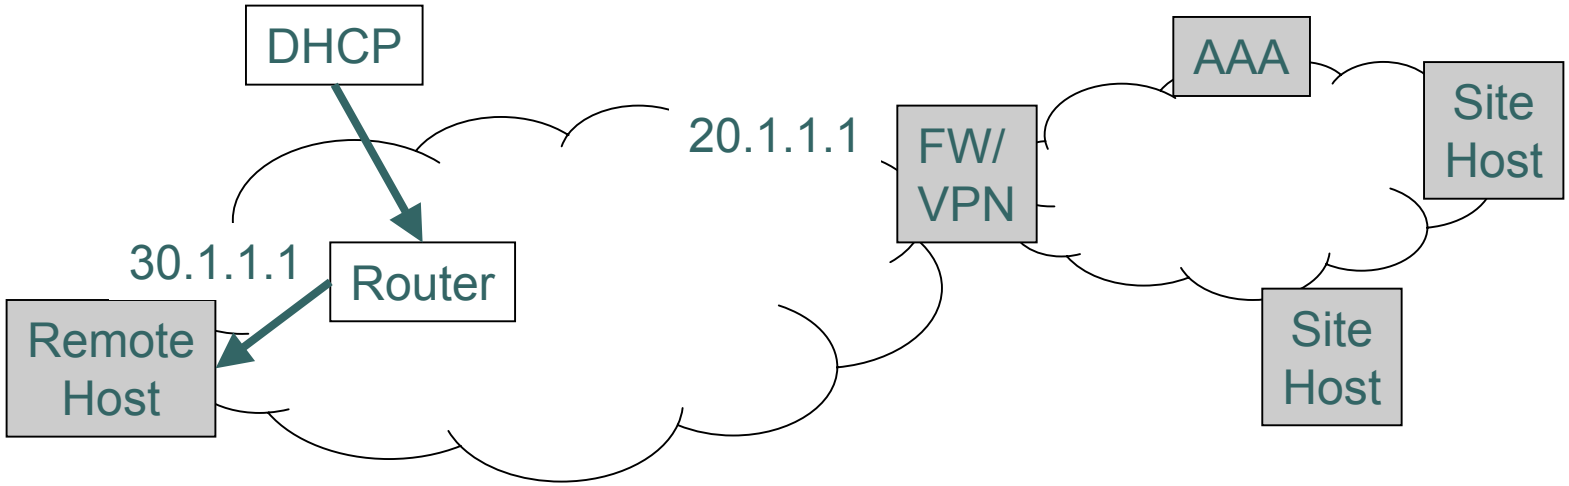

Client VPNs: Host connects to VPN

CS419

VPN authenticates remote host through backend database (RADIUS or LDAP)

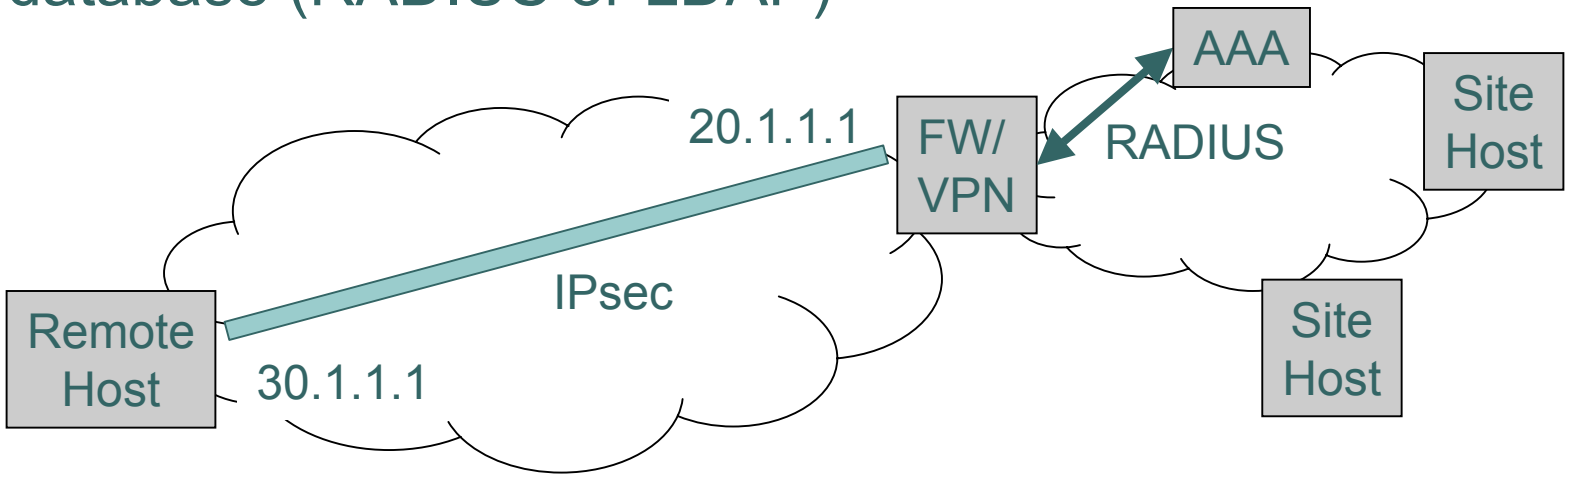

Client VPNs: VPN assigns site address

CS419

As proprietary enhancement to IPsec, or with PPP (over IPsec)

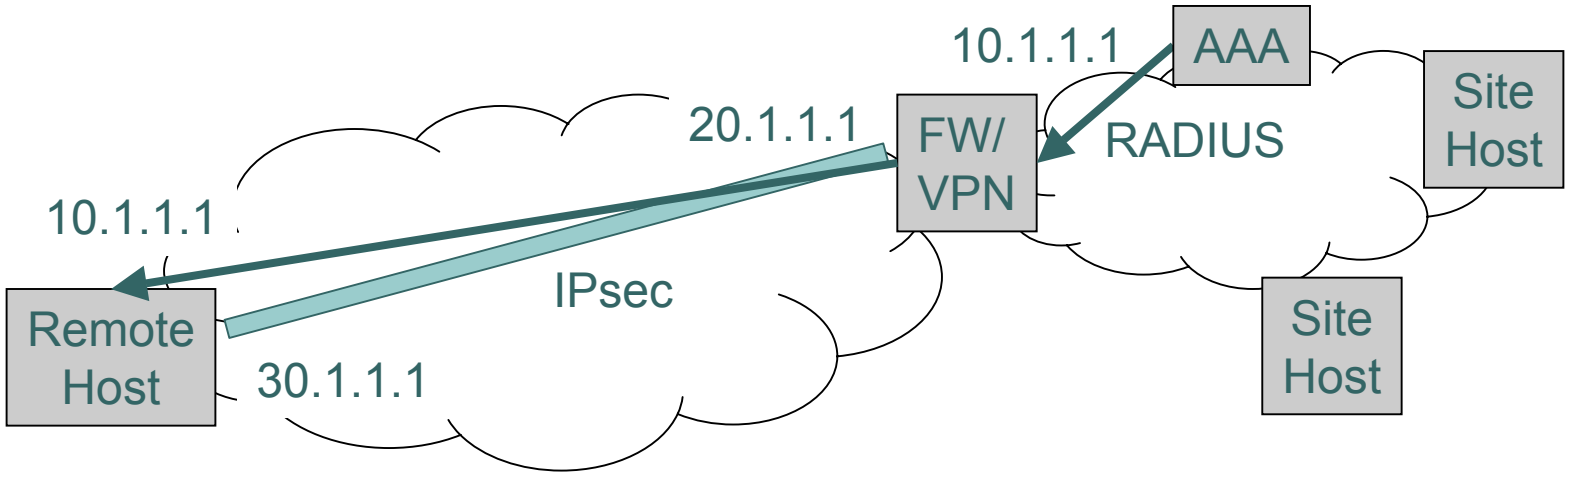

Client VPNs: Packets tunneled over IPsec

CS419

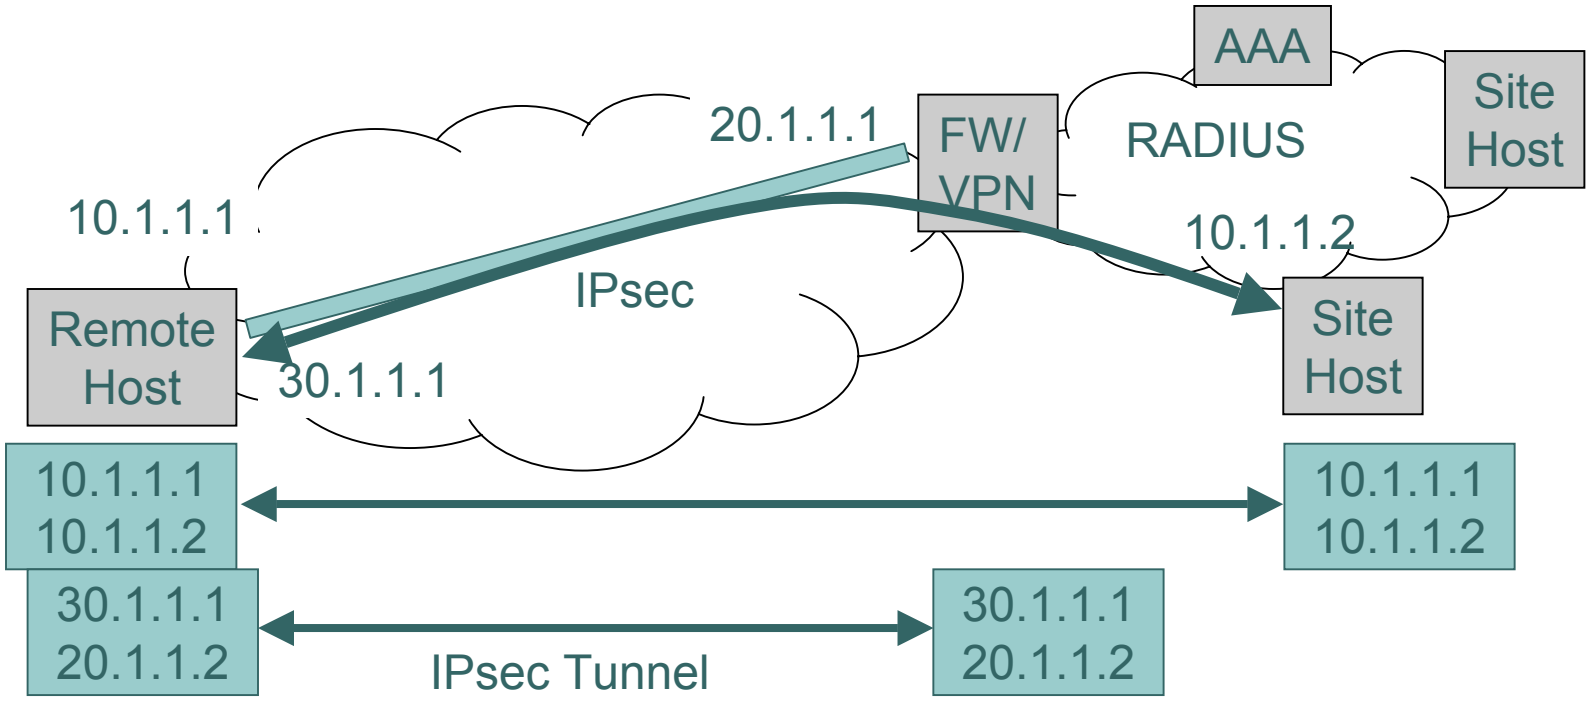

Client VPNs: Packets tunneled over IPsec

CS419

Some VPN clients smart enough to avoid sending non-VPN traffic through the VPN tunnel

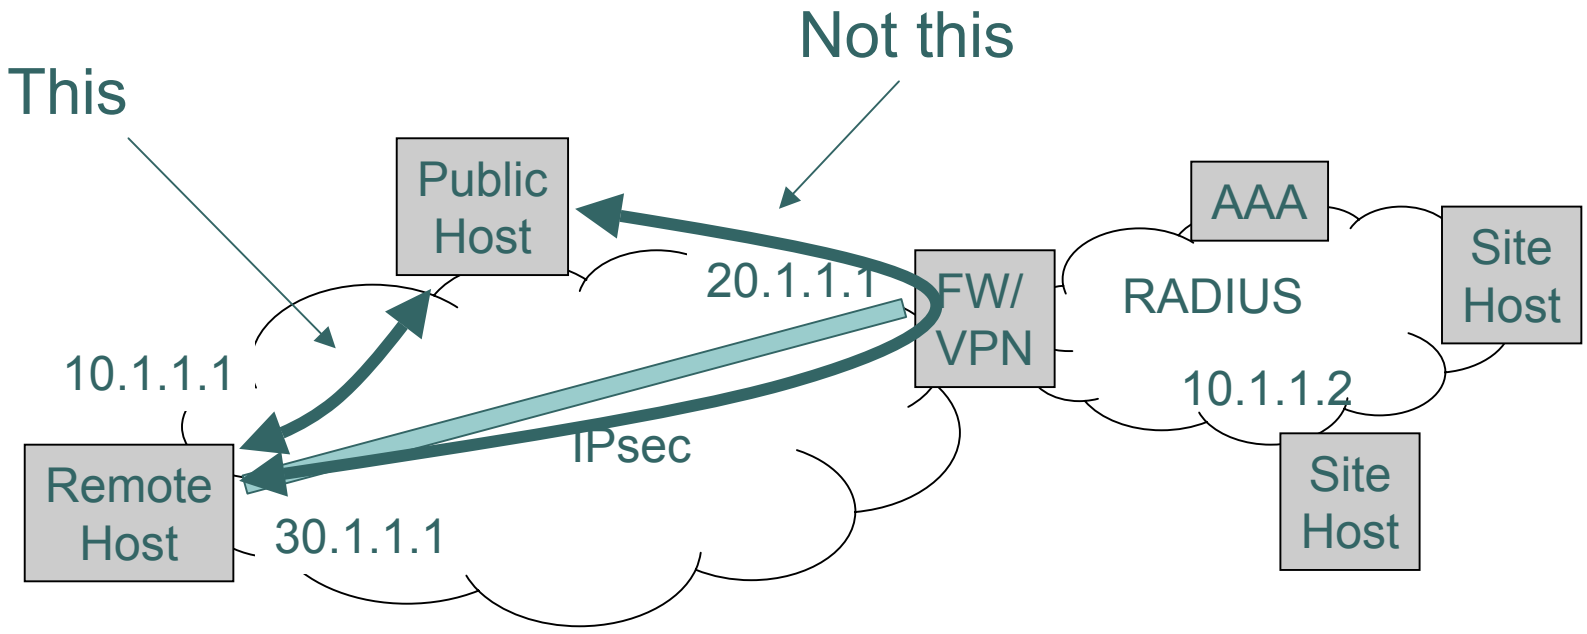

IPsec

CS419

- Two parts: Session Establishment (key exchange) and Payload
- IKE/ISAKMP is session establishment
  - Negotiate encryption algorithms
  - Negotiate payload headers (AH, ESP)
  - Negotiate policies
- Keying can be either:
  - Symmetric shared keys
  - Public keys (in certificates)
- Either way, a session key is negotiated by IKE

IPsec Payloads

CS419

- AH: Authentication Header
  - Authenticates each packet but doesn't encrypt
  - Has fallen out of favor (redundant and no more efficient)
- ESP: Encapsulating Security Payload
  - Encrypts (with authentication as side effect)

IPsec transmission modes: Transport or Tunnel mode

CS419

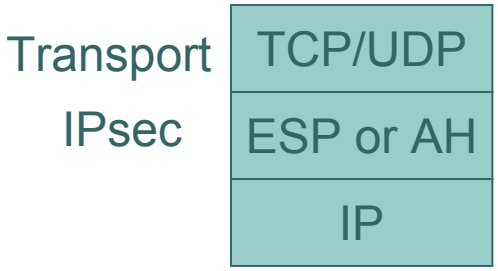

Transport mode. Used when IPsec tunnel is end-to-end. Operates over some of the IP fields, and doesn't work with NAT!

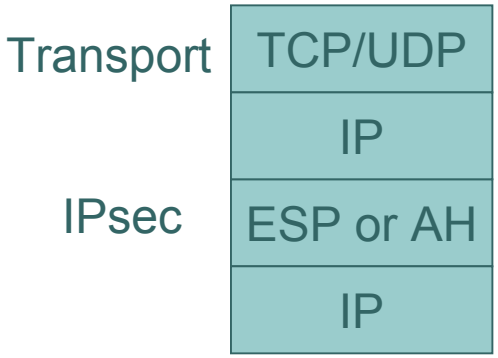

Tunnel mode. Used when IPsec tunnel not end-to-end. Hides the IP identity of endpoints. Operates over inner IP fields…can work with NAT.

AH header format

CS419

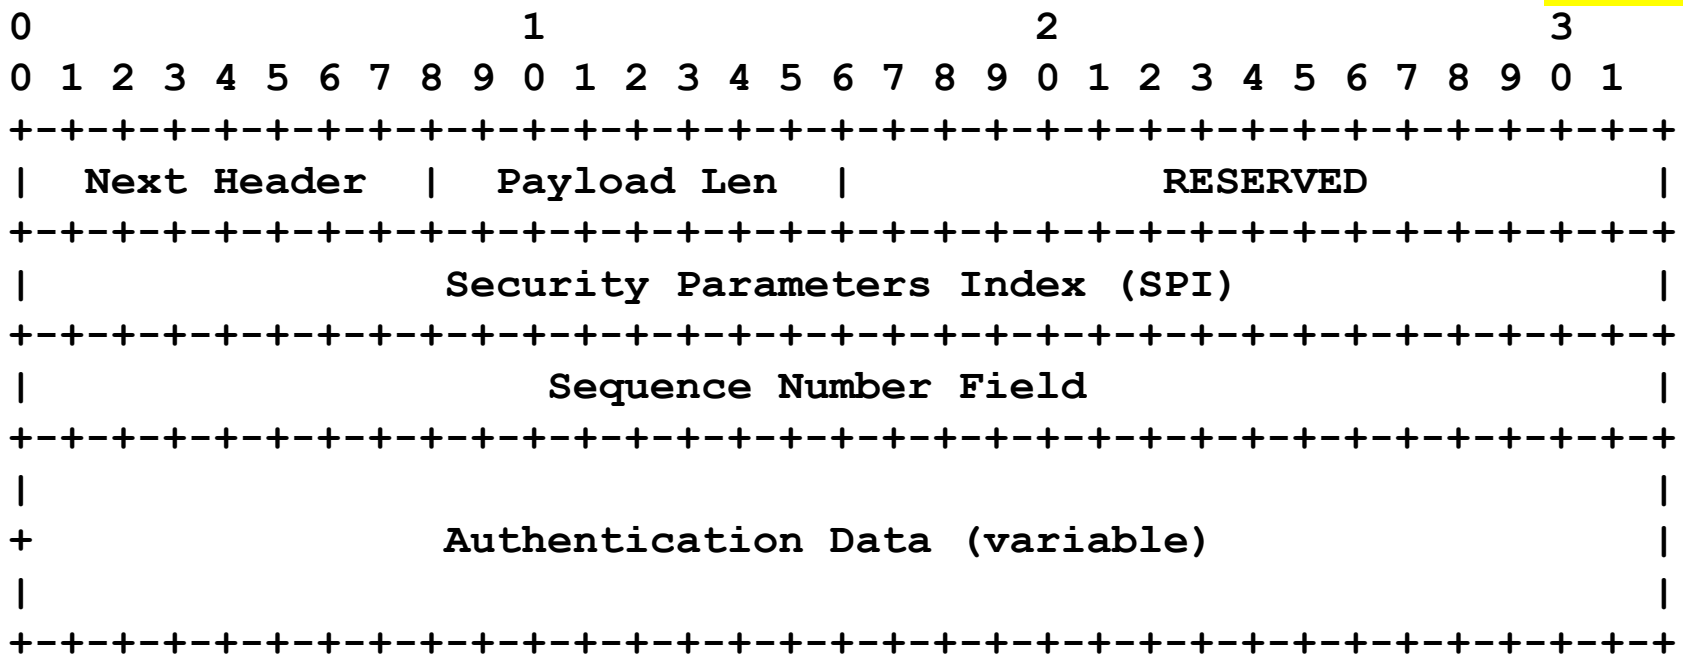

ESP header format

CS419

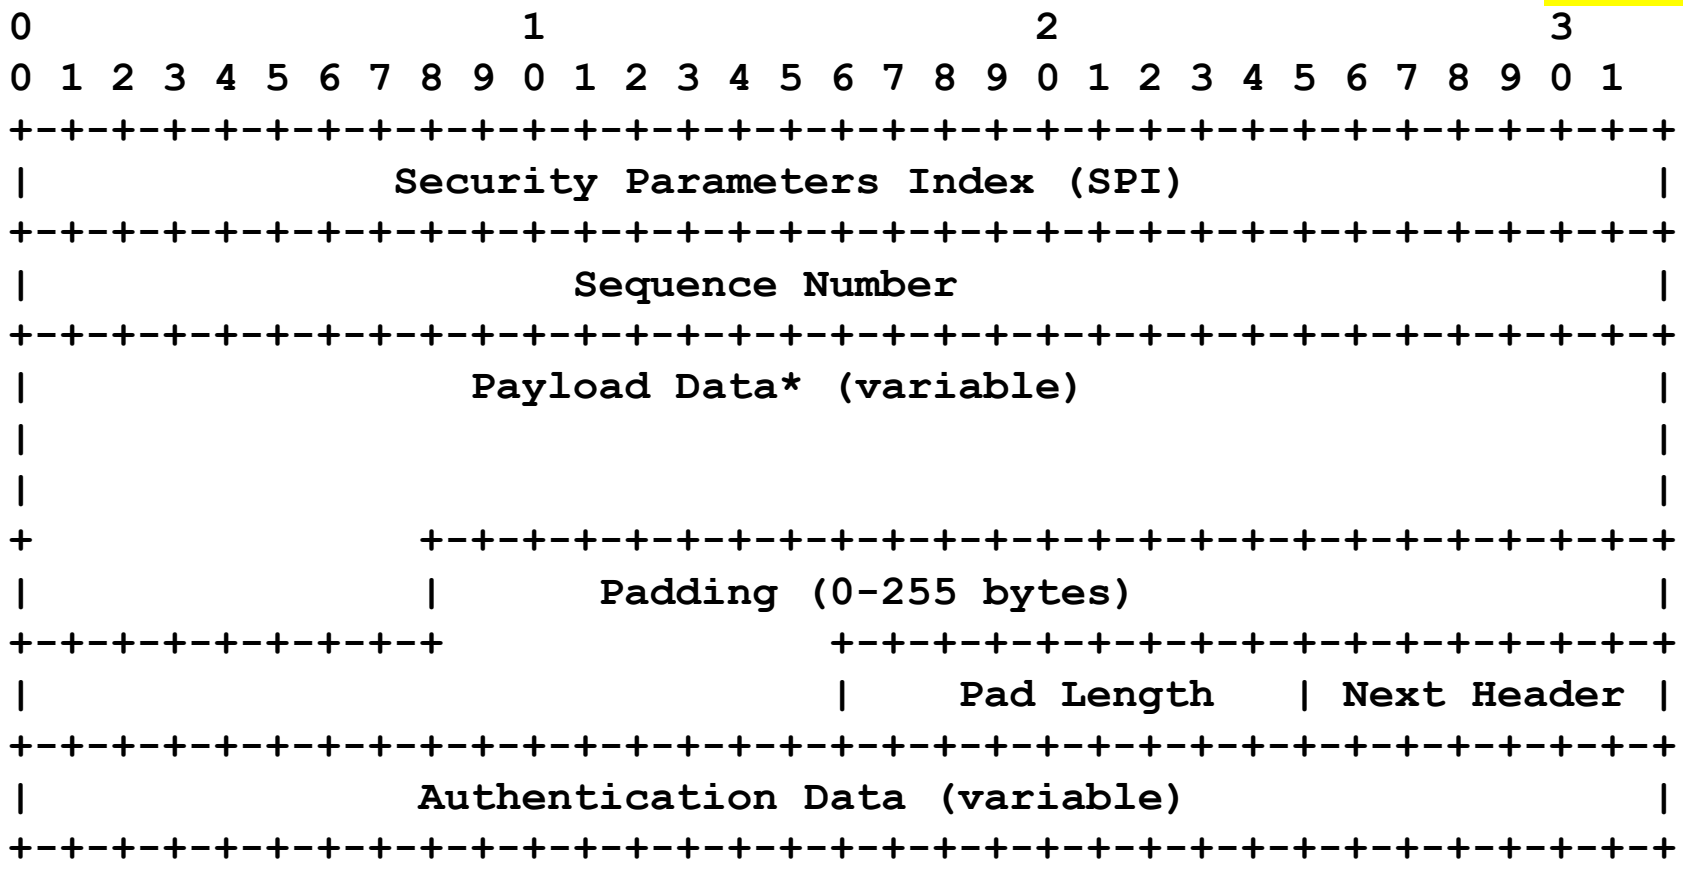

New IPsec transmission modes

CS419

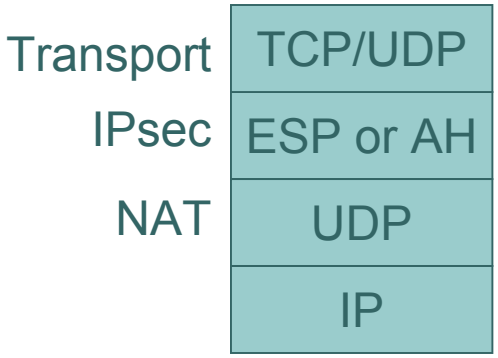

Extra layer of UDP allows IPsec to work over NAT.

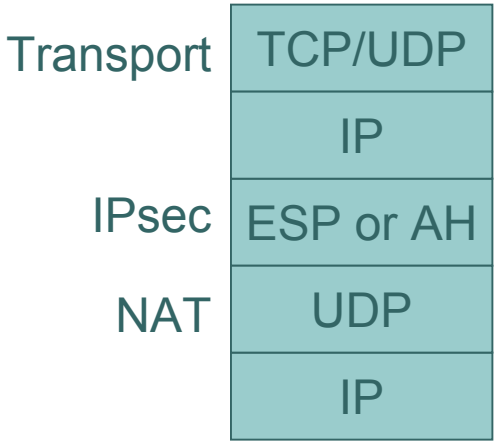

Client VPNs

CS419

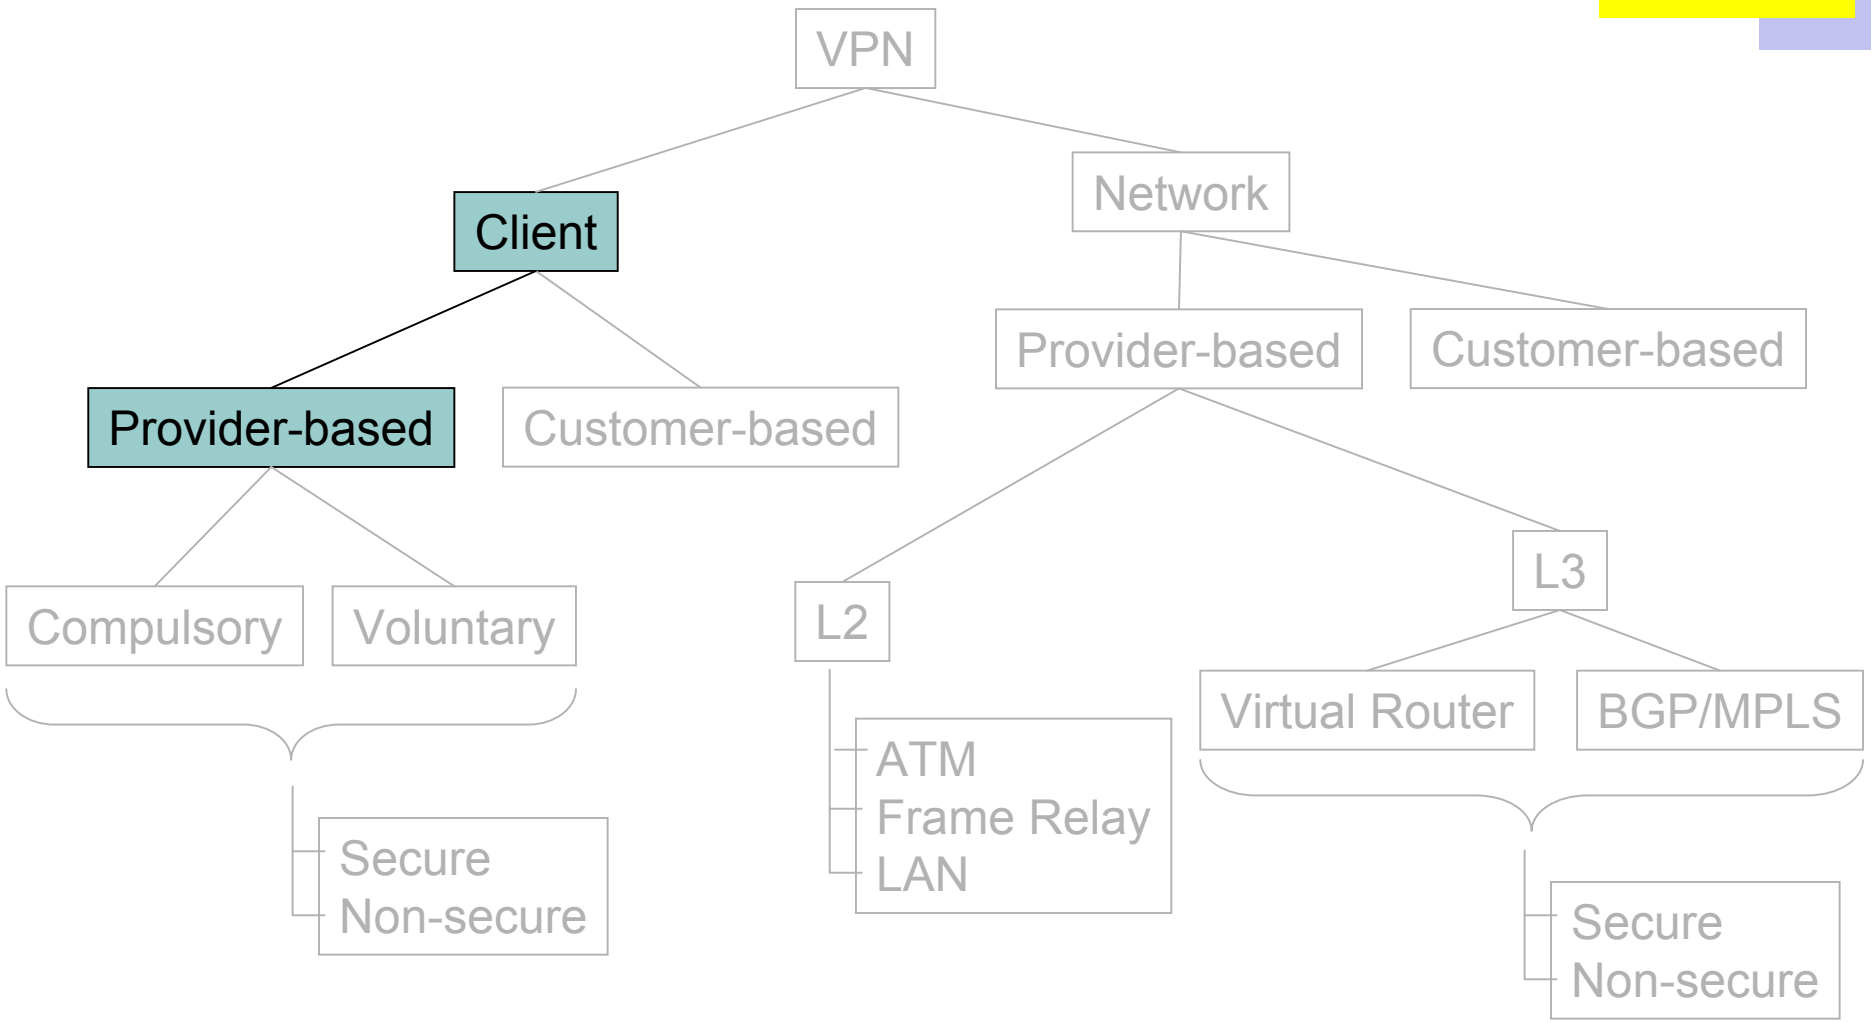

Client VPNs: Host gets local IP address

CS419

2. If PPP, AAA tells Access Router to tunnel user to VPN. (If not PPP, Access Router uses local configuration.)

3. Tunnel established (or packets forwarded over pre-established tunnel)

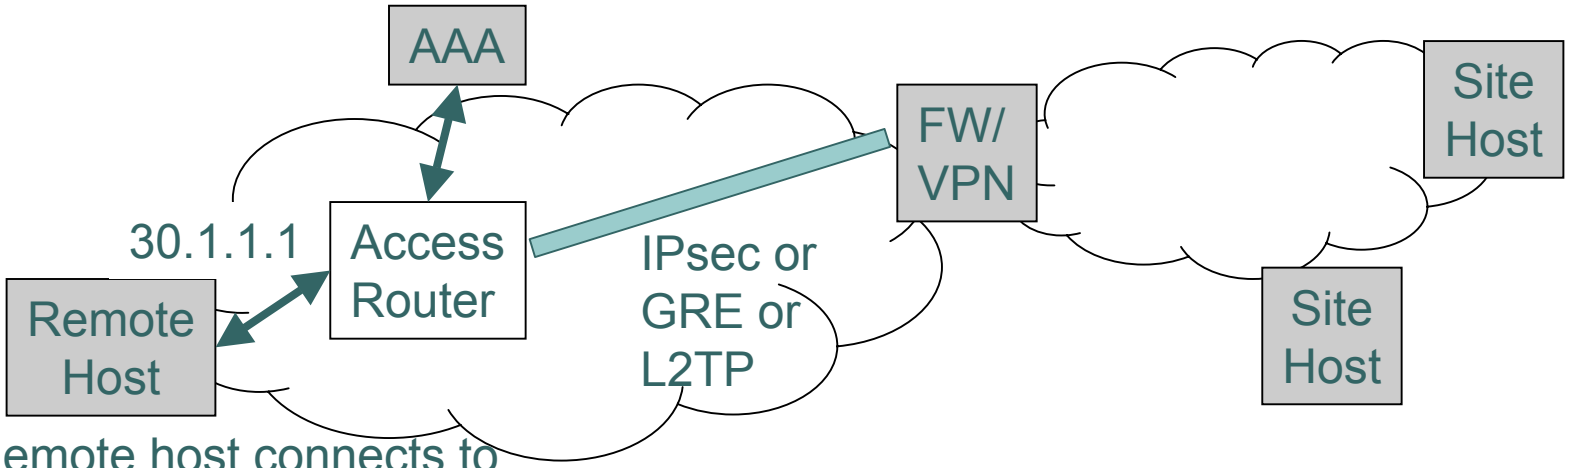

1. Remote host connects to Internet (dialup-PPP or PPPoE (cable) or DSL)

Compulsory if Access Router forces tunnel, voluntary if user requests it (through certain NAI). NAI = “user@domain”

Provider-based client VPNs

CS419

- Used for instance when enterprise pays for employee access, wants it to go through enterprise network
  - I know Cisco did this
  - But never used that much
    - Business model didn't take off
  - Used even less now
    - In part because VPN client comes with windows OS???
- The tunneling technology commonly used for roaming dialup though

Network VPNs

CS419

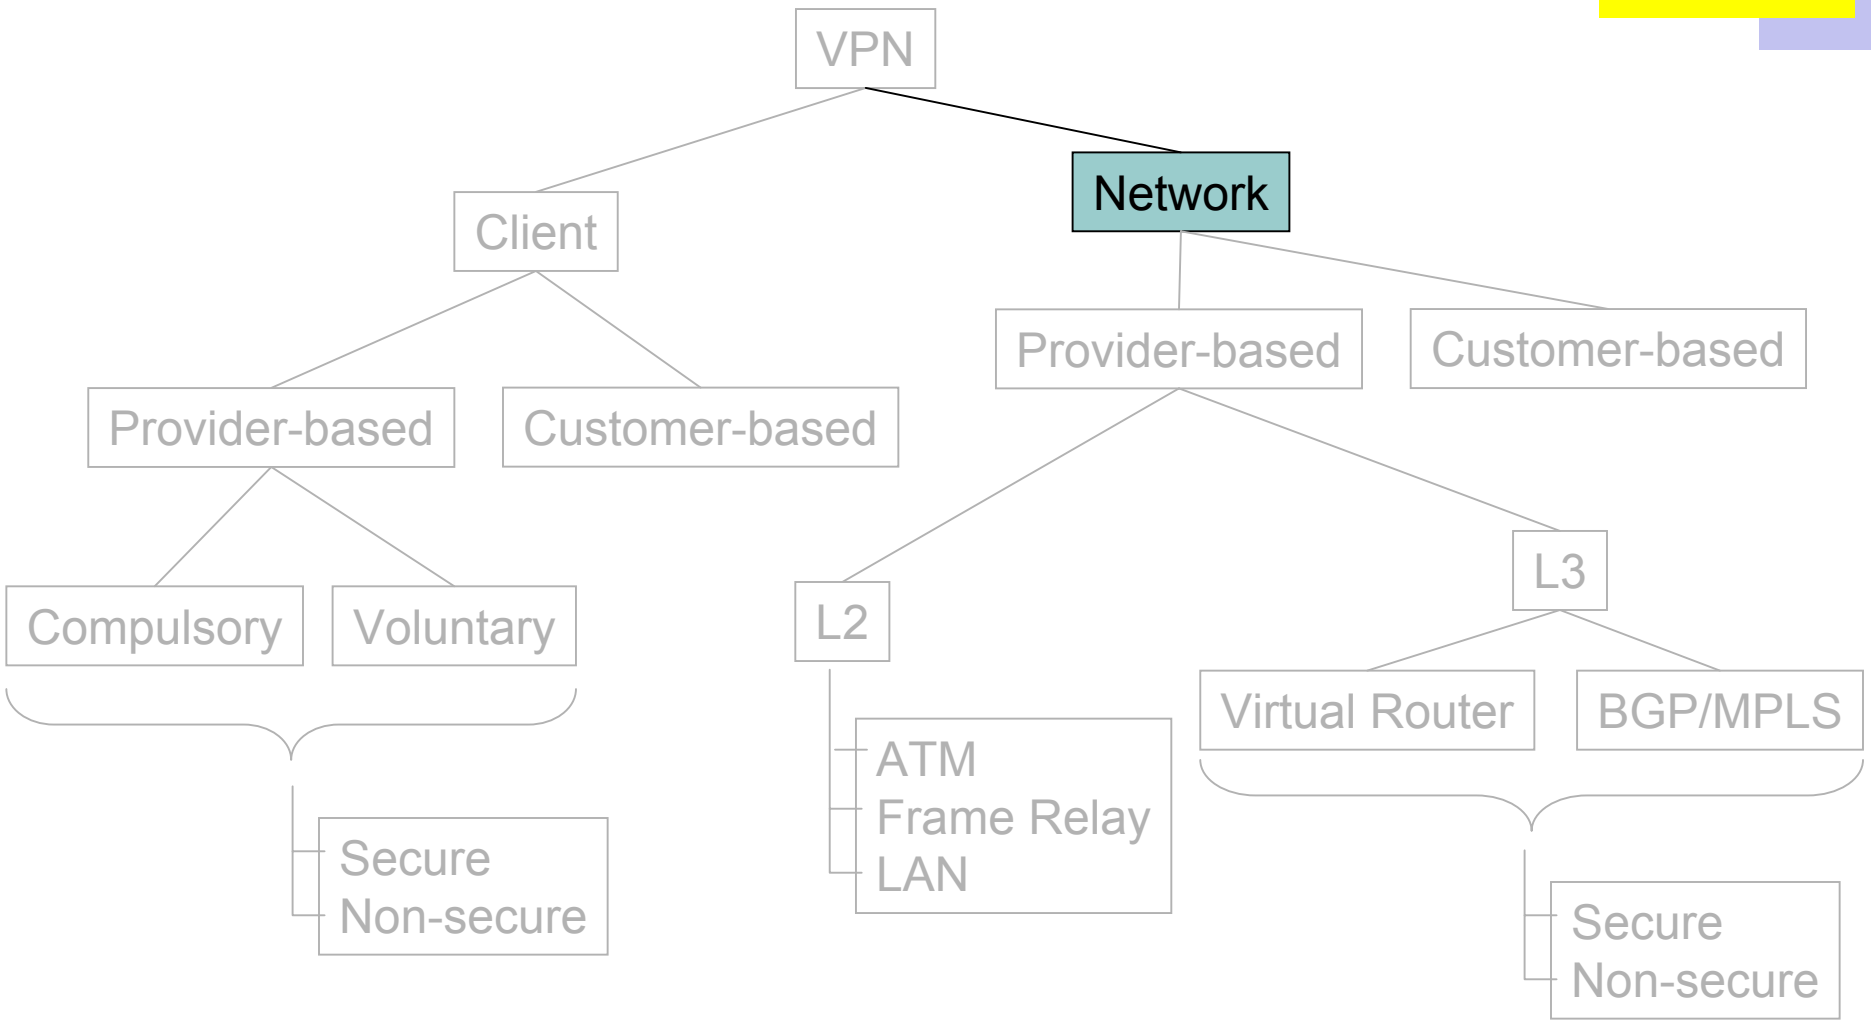

Reiterate network VPN benefits

CS419

- Makes separated IP sites look like one private IP network
- Security
- QoS guarantees
- Simplified network operation

Customer-based Network VPNs

CS419

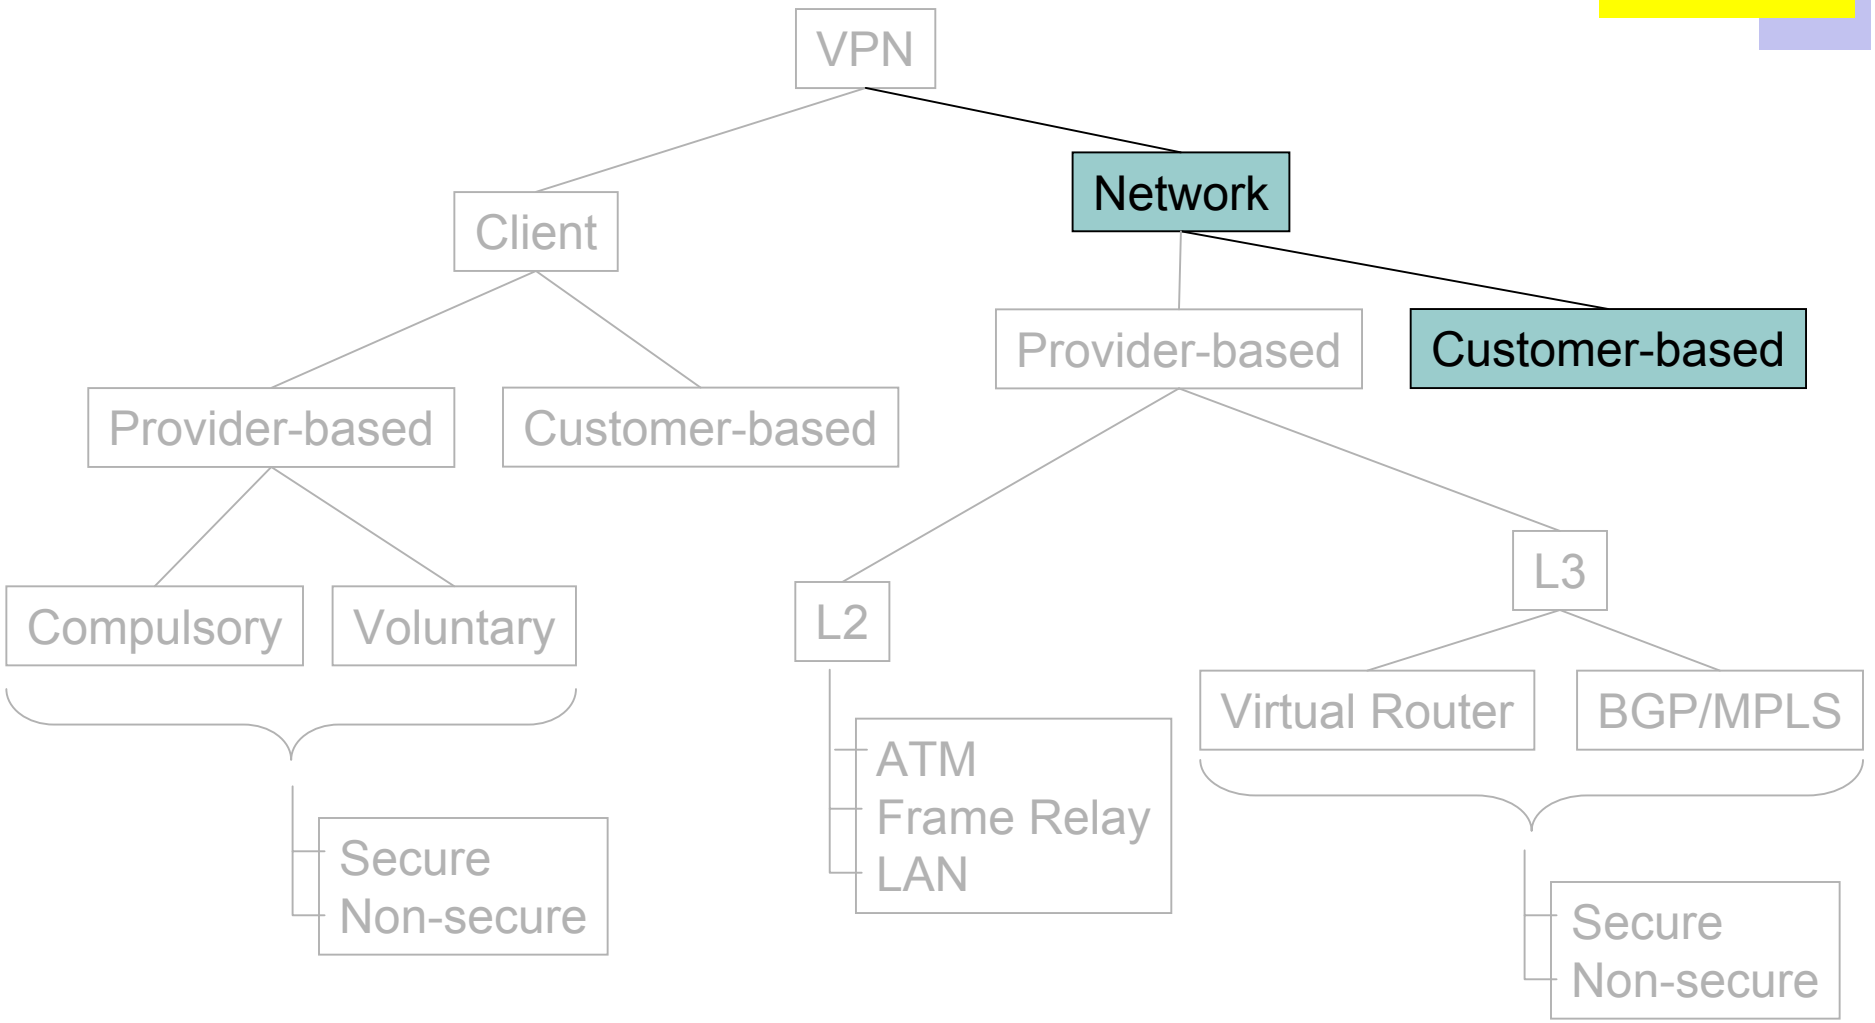

Customer-based Network VPNs

CS419

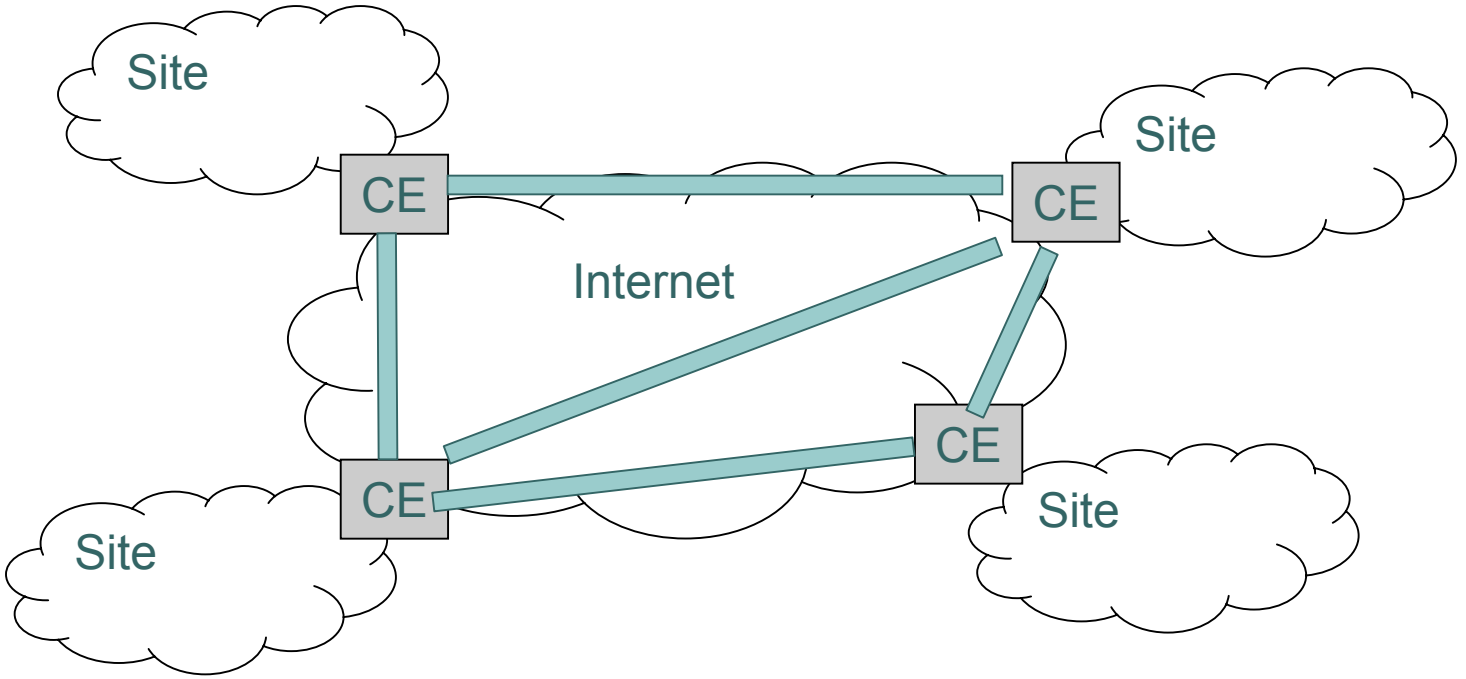

Customer buys own equipment, configures IPsec tunnels over the global internet, manages addressing and routing. ISP plays no role.

Customer-based Network VPNs

CS419

- Great for enterprises that have the resources and skills to do it
  - Large companies
- More control, better security model
  - Doesn’t require trust in ISP ability and intentions
  - Can use different ISPs at different sites
- But not all enterprises have this skill

Provider-based Network VPNs (aka Provider Provisioned: PPVPN)

CS419

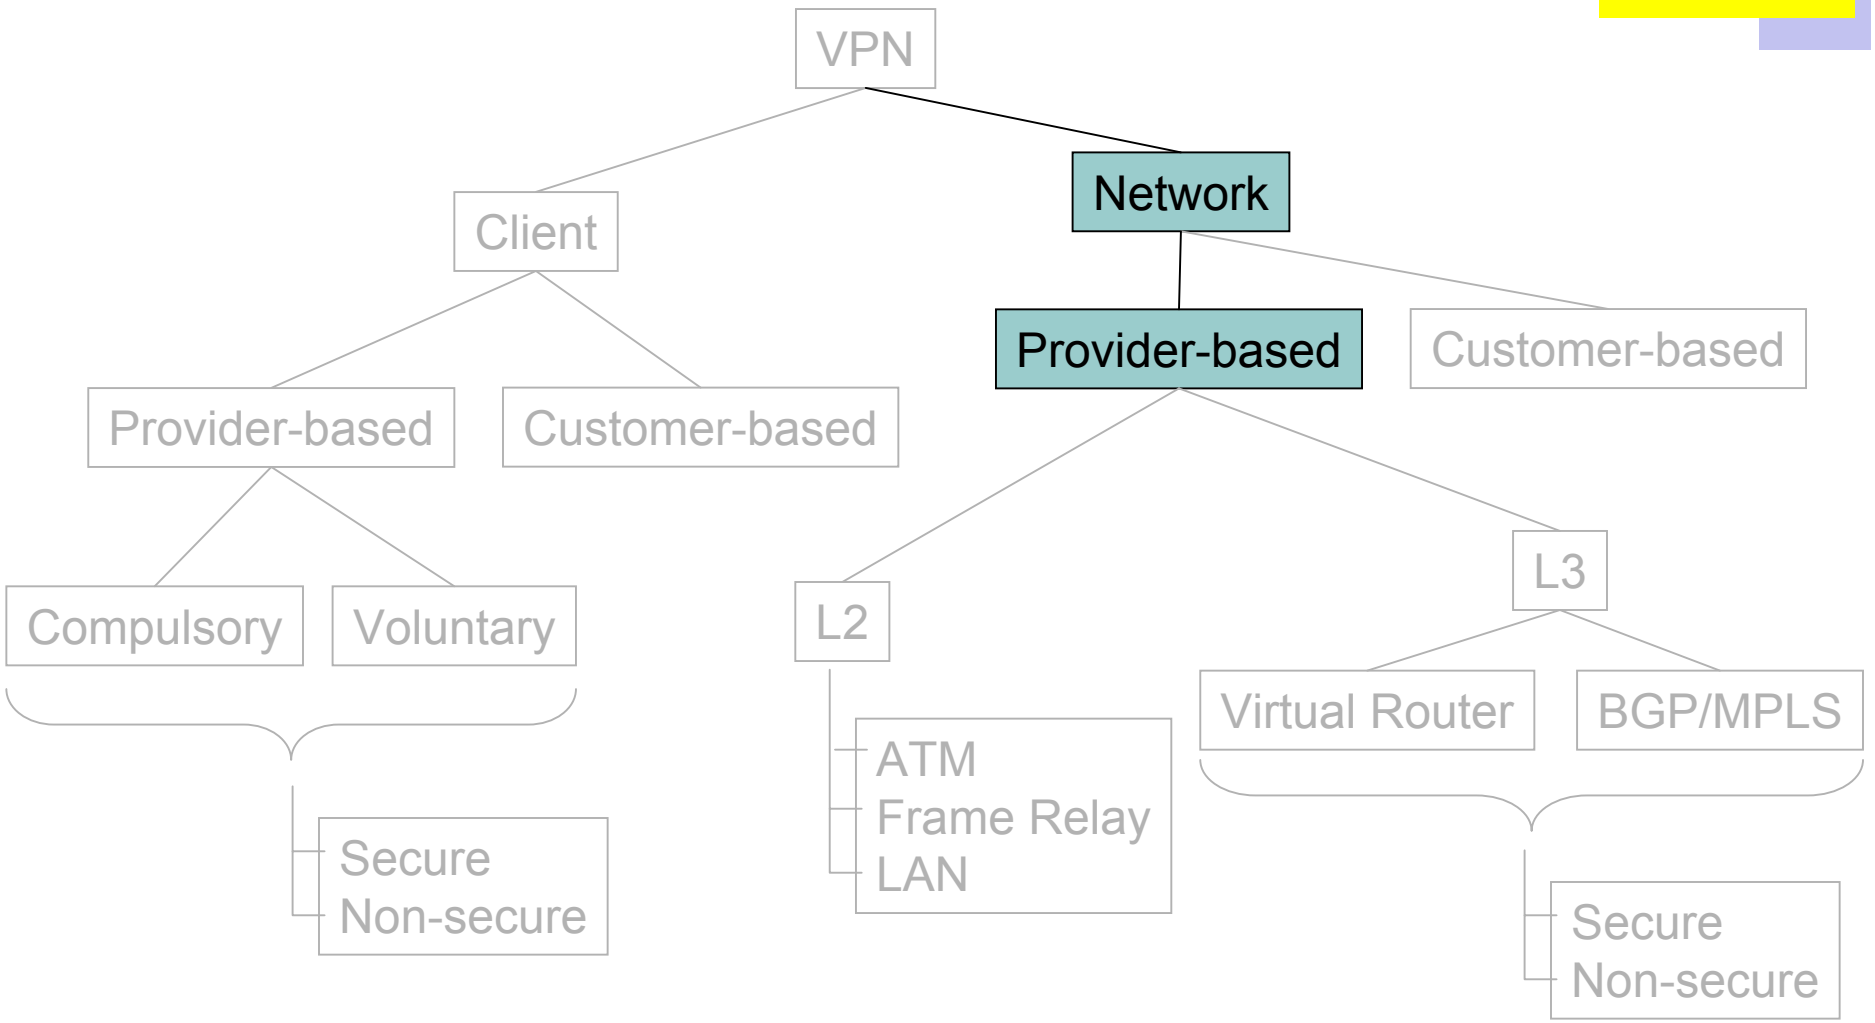

Provider-based Network VPNs

CS419

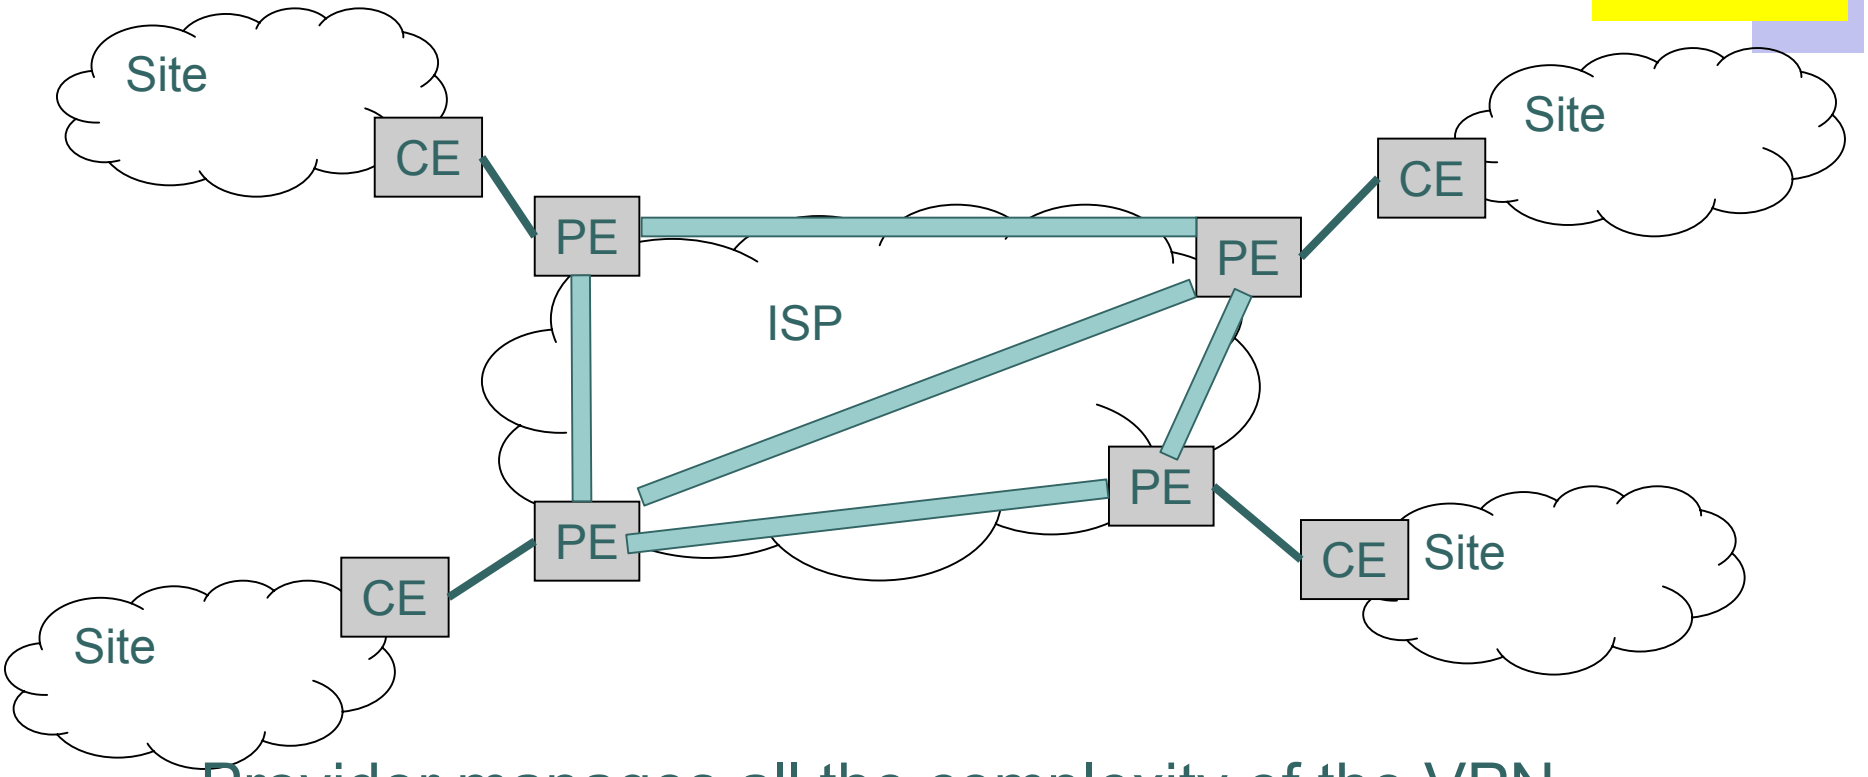

Provider manages all the complexity of the VPN.
Customer simply connects to the provider equipment.

Same provider equipment used for multiple customers

CS419

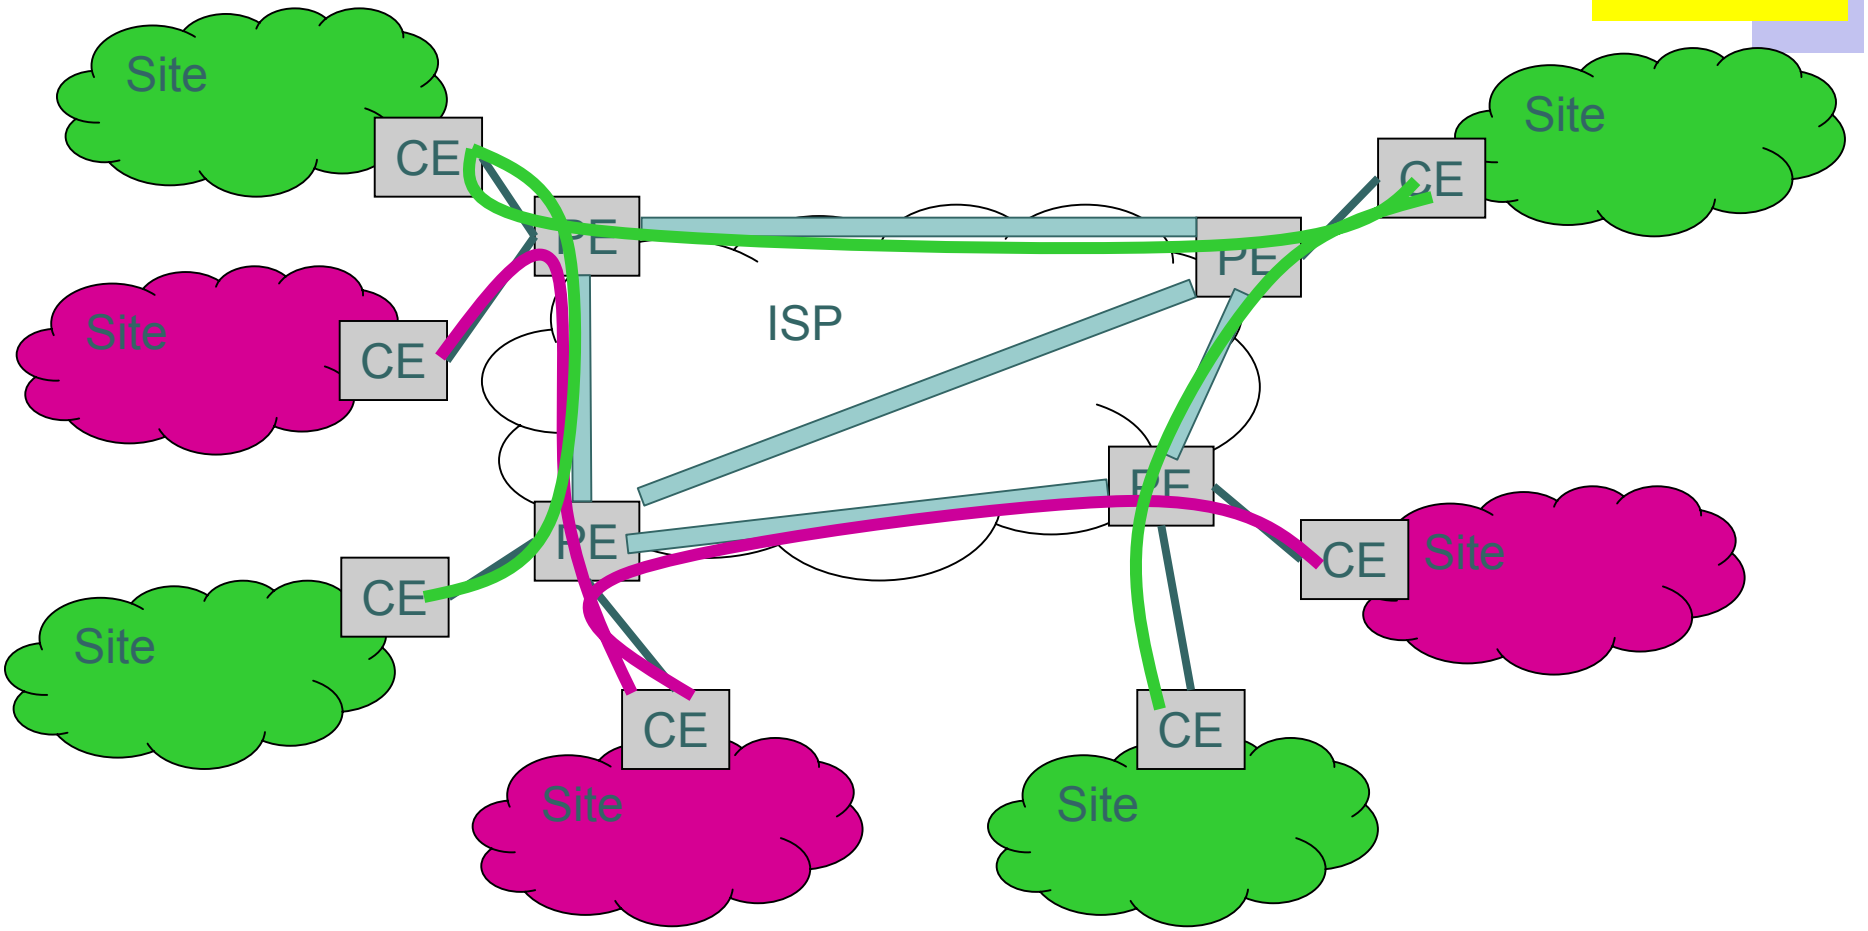

Model for customer

- Attach to ISP router (PE) as though it was one of your routers
- Run routing algorithm with it
  - OSPF, RIP, BGP
- PE will advertise prefixes from other sites of same customer

Various PPVPN issues

CS419

- Tunnel type?
  - IPsec (more secure, more expensive)
  - GRE etc.
- How to discover which customer is at which PE?
  - Don’t want PEs without given customer to participate in routing for that customer
- How to distinguish overlapping private address spaces

BGP/MPLS VPNs (RFC2547)

CS419

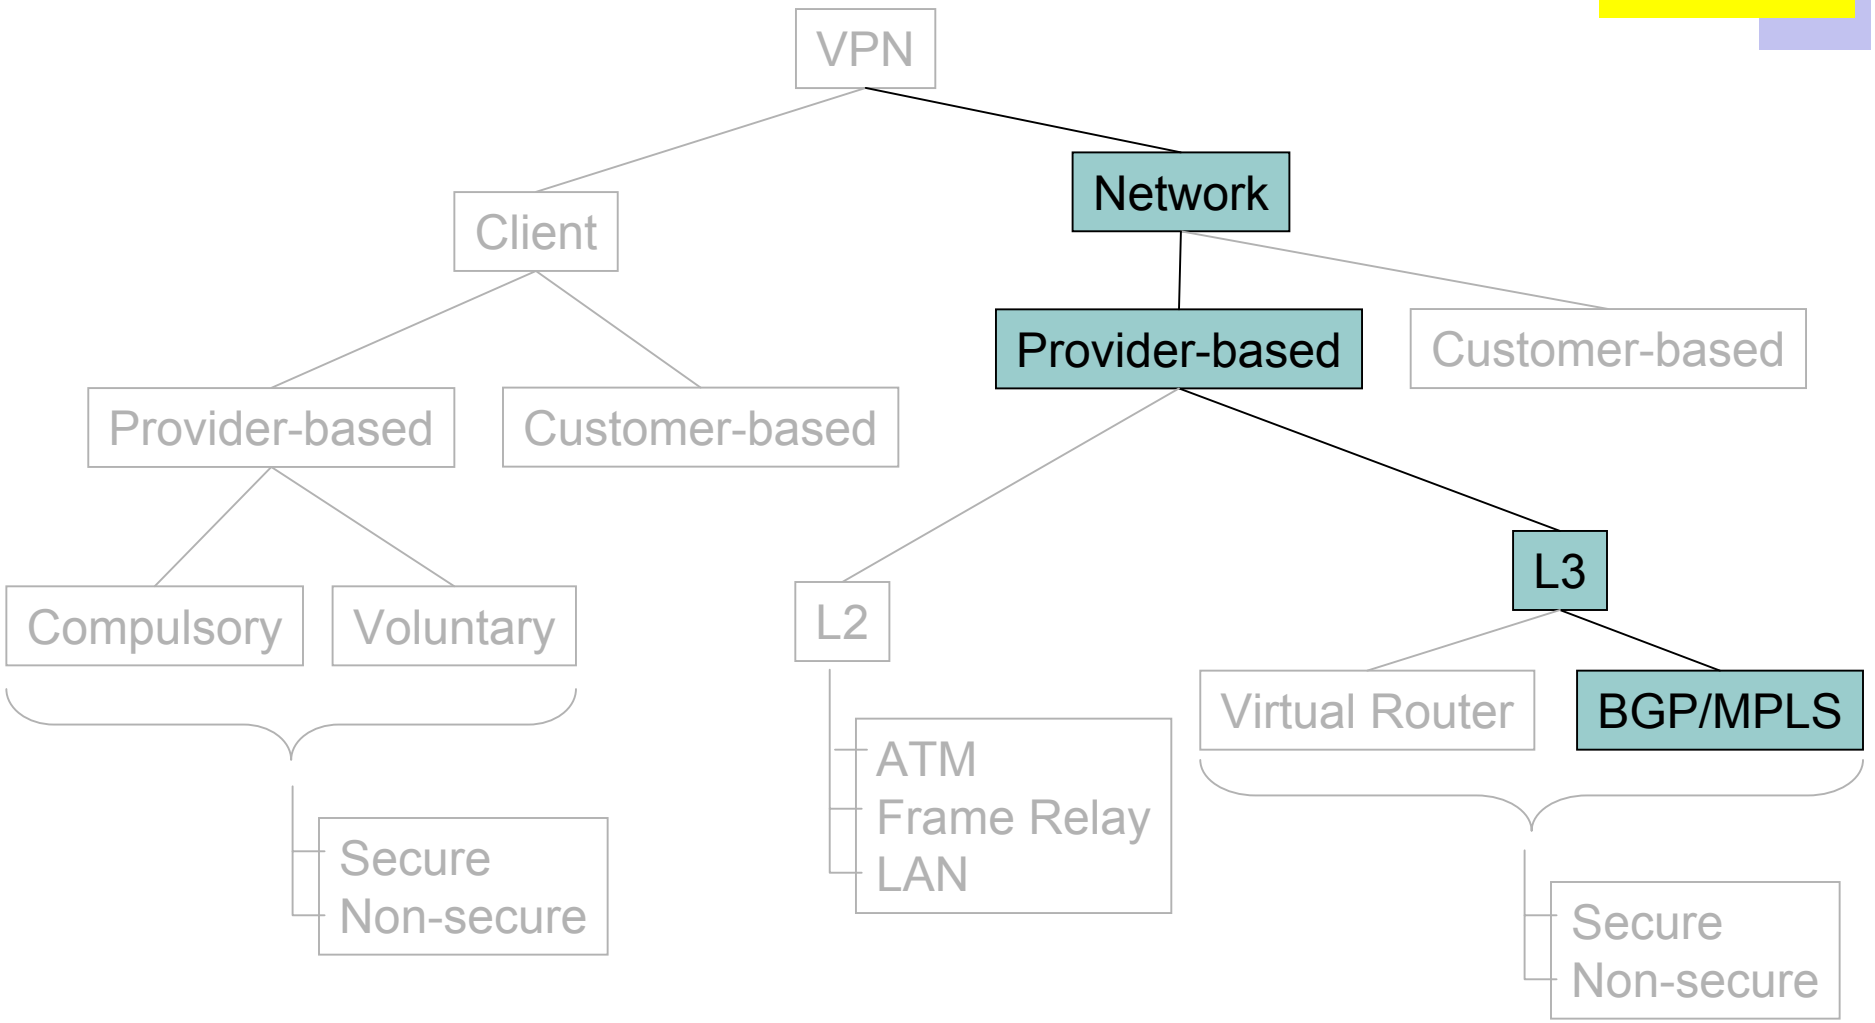

BGP/MPLS VPNs (RFC2547)

CS419

- Cisco invention
  - Leverage Cisco's investment in both BGP and MPLS (Multi-Protocol Label Switching)
- What is MPLS?
  - Link-layer technology
    - Tags like circuit switching
    - But with some IP awareness
  - How Cisco killed Epsilon
  - Initially marketed as high performance switching
  - Later became "traffic engineering" and VPN

Recall this frame-relay traffic engineered L2 VPN…

CS419

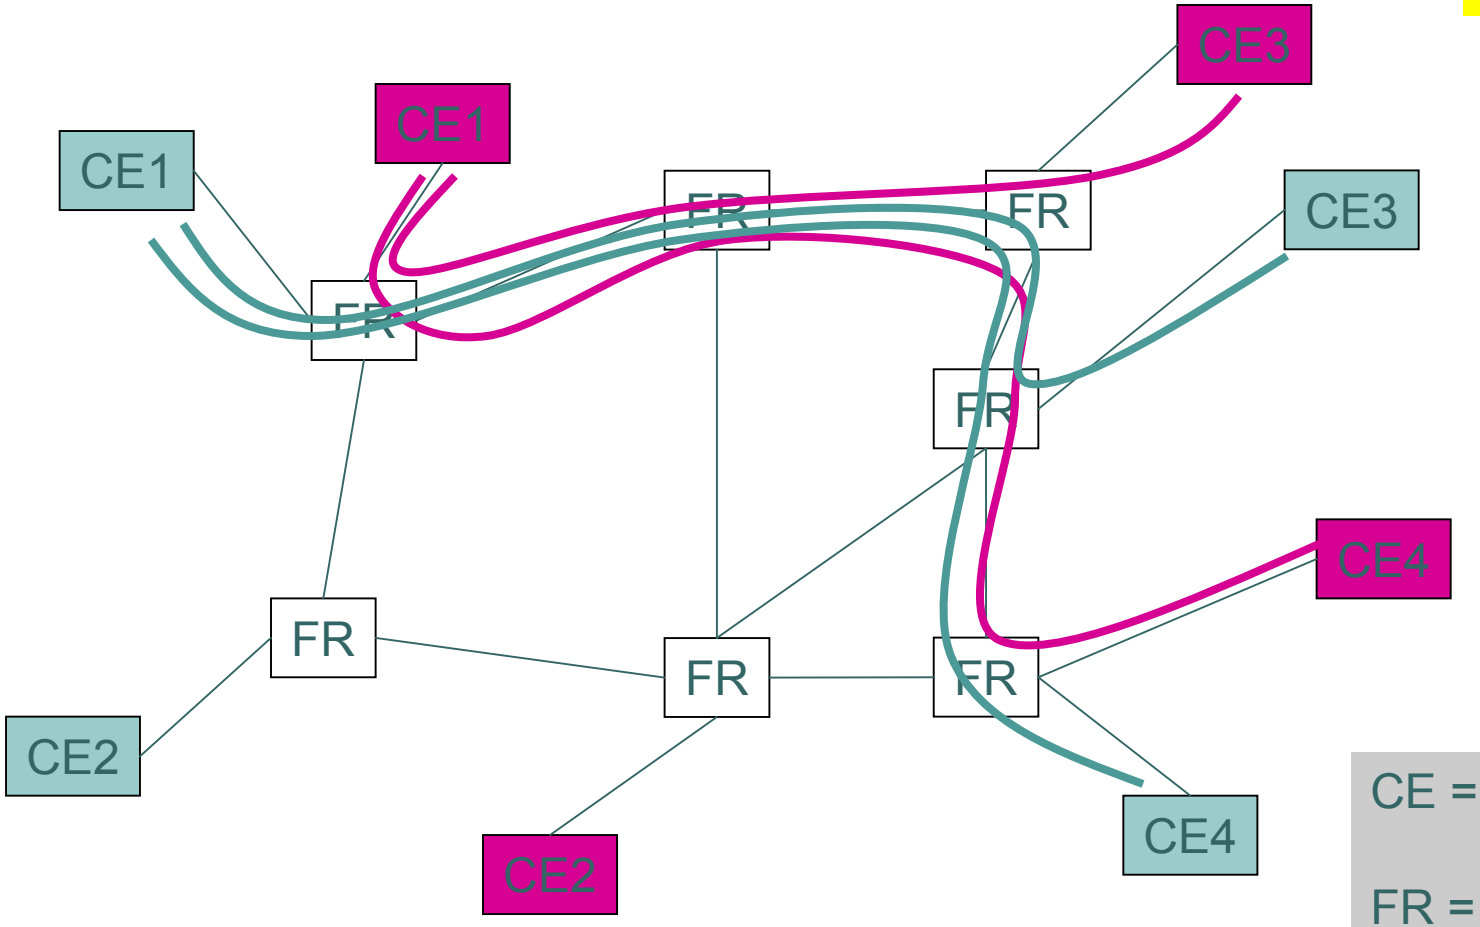

ISPs historically used L2 networks in their core

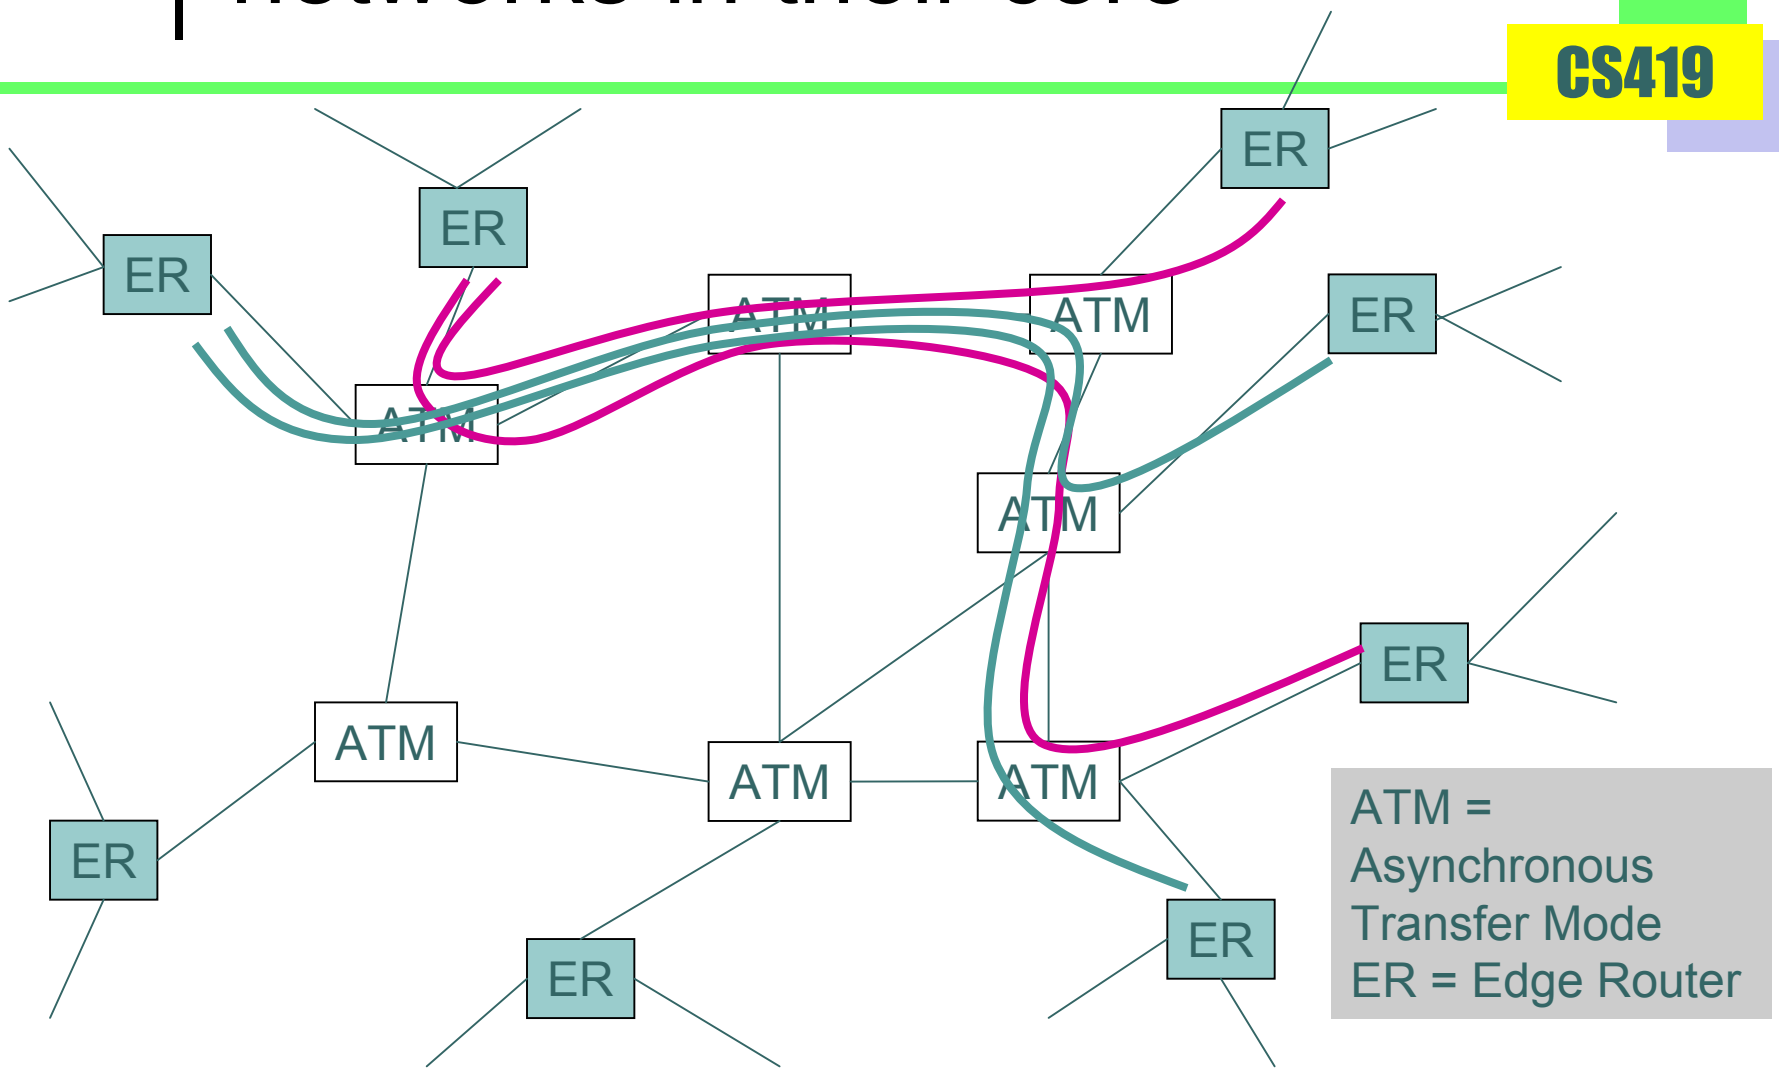

Logically, ISPs were structured like this

CS419

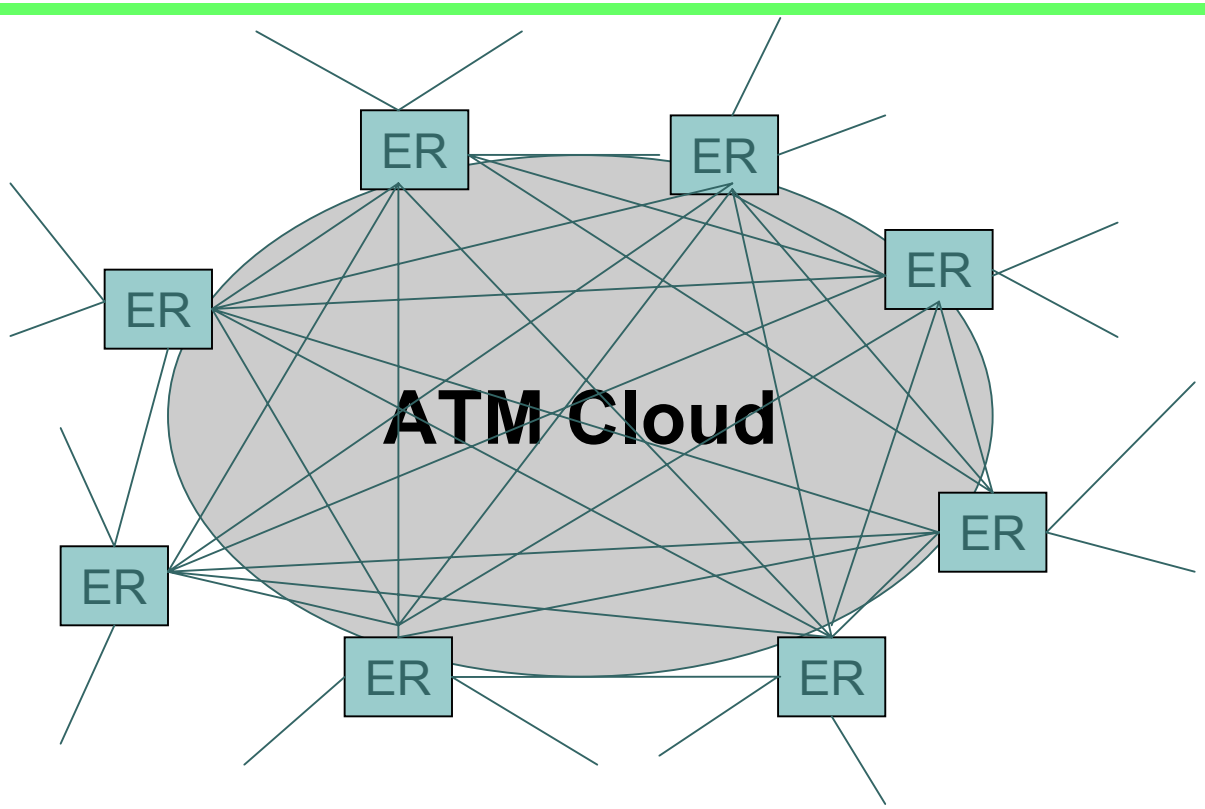

Every router was “adjacent” to every other

Why L2 (ATM)?

CS419

- ATM was, at least until 4-5 years ago, faster than IP forwarding
- ATM switches were better matched to the underlying SONET transmission links
- It was easier to traffic engineer based on virtual circuits than based on destination IP address
- IP wasn't the only network protocol

But there were problems…

CS419

- ISPs had 100's of routers, each of which logically had a link to all others
  - Was difficult to manage and run routing over all of these logical links
  - Scaled poorly
- Basic idea of MPLS was to elevate ATM intelligence to L3, while doing switching at L2!
  - Epsilon business model…

MPLS tried to get the best of both worlds

CS419

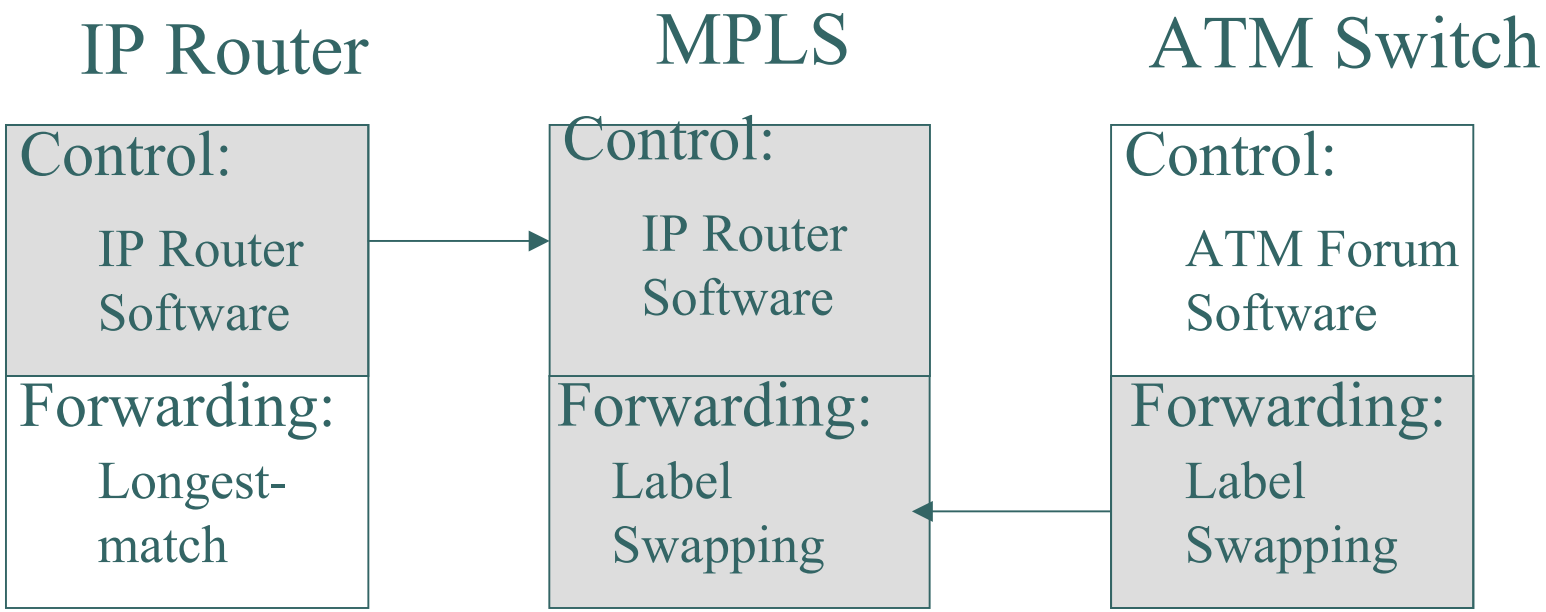

MPLS Operation

1a. Routing protocols (e.g. OSPF-TE, IS-IS-TE) exchange reachability to destination networks

1b. Label Distribution Protocol (LDP) establishes label mappings to destination network

4. LER at egress removes label and delivers packet

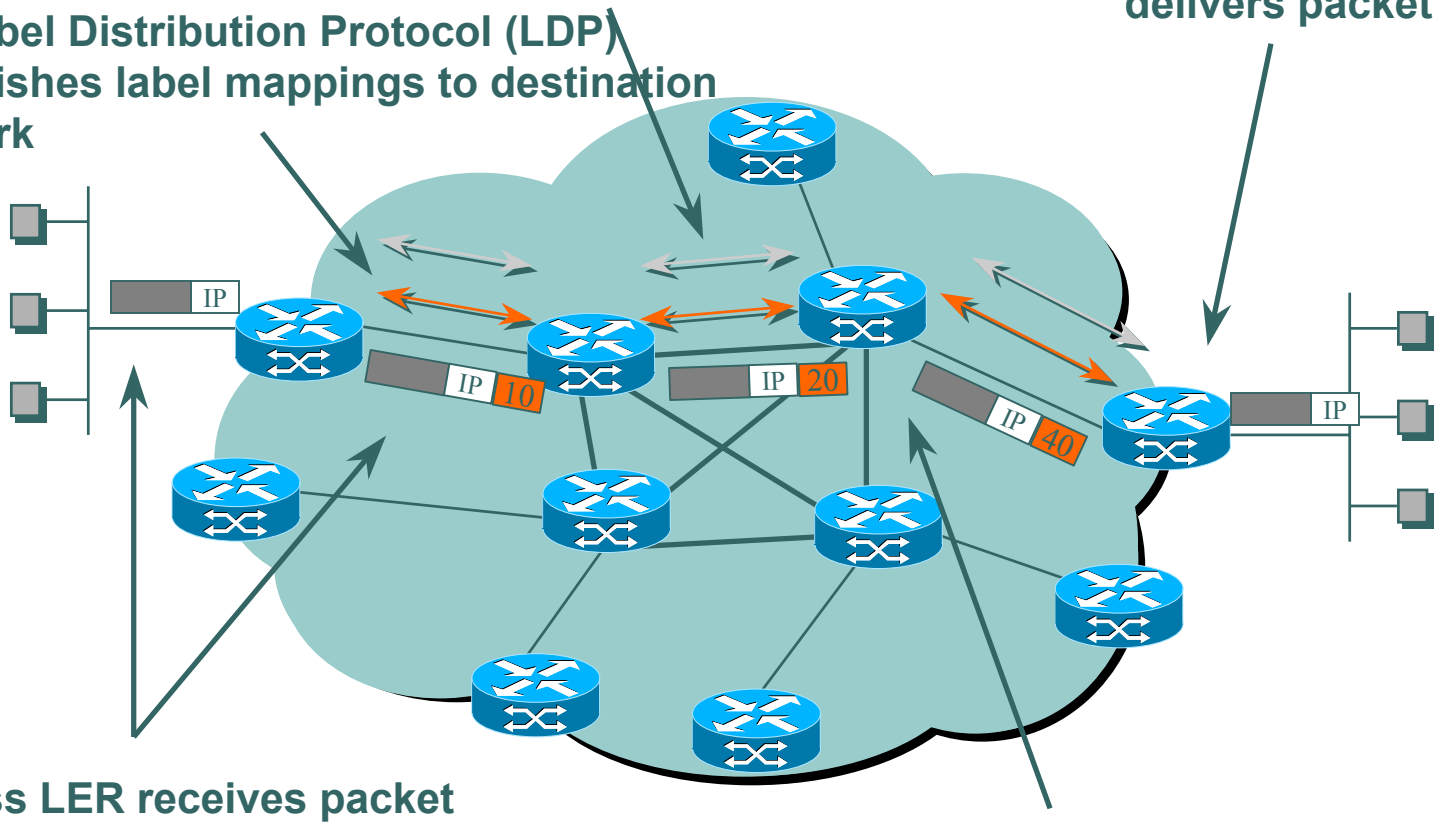

2. Ingress LER receives packet and “label”s packets

3. LSR forwards packets using label swapping

Original business model failed

CS419

- Simple reason:
  - People figured out how to make IP fast…as fast as ATM
- MPLS spent a long time looking for a reason to exist
  - Finally found it in MPLS-BGP PPVPNs

Basic difficulty with PPVPN: private addresses

CS419

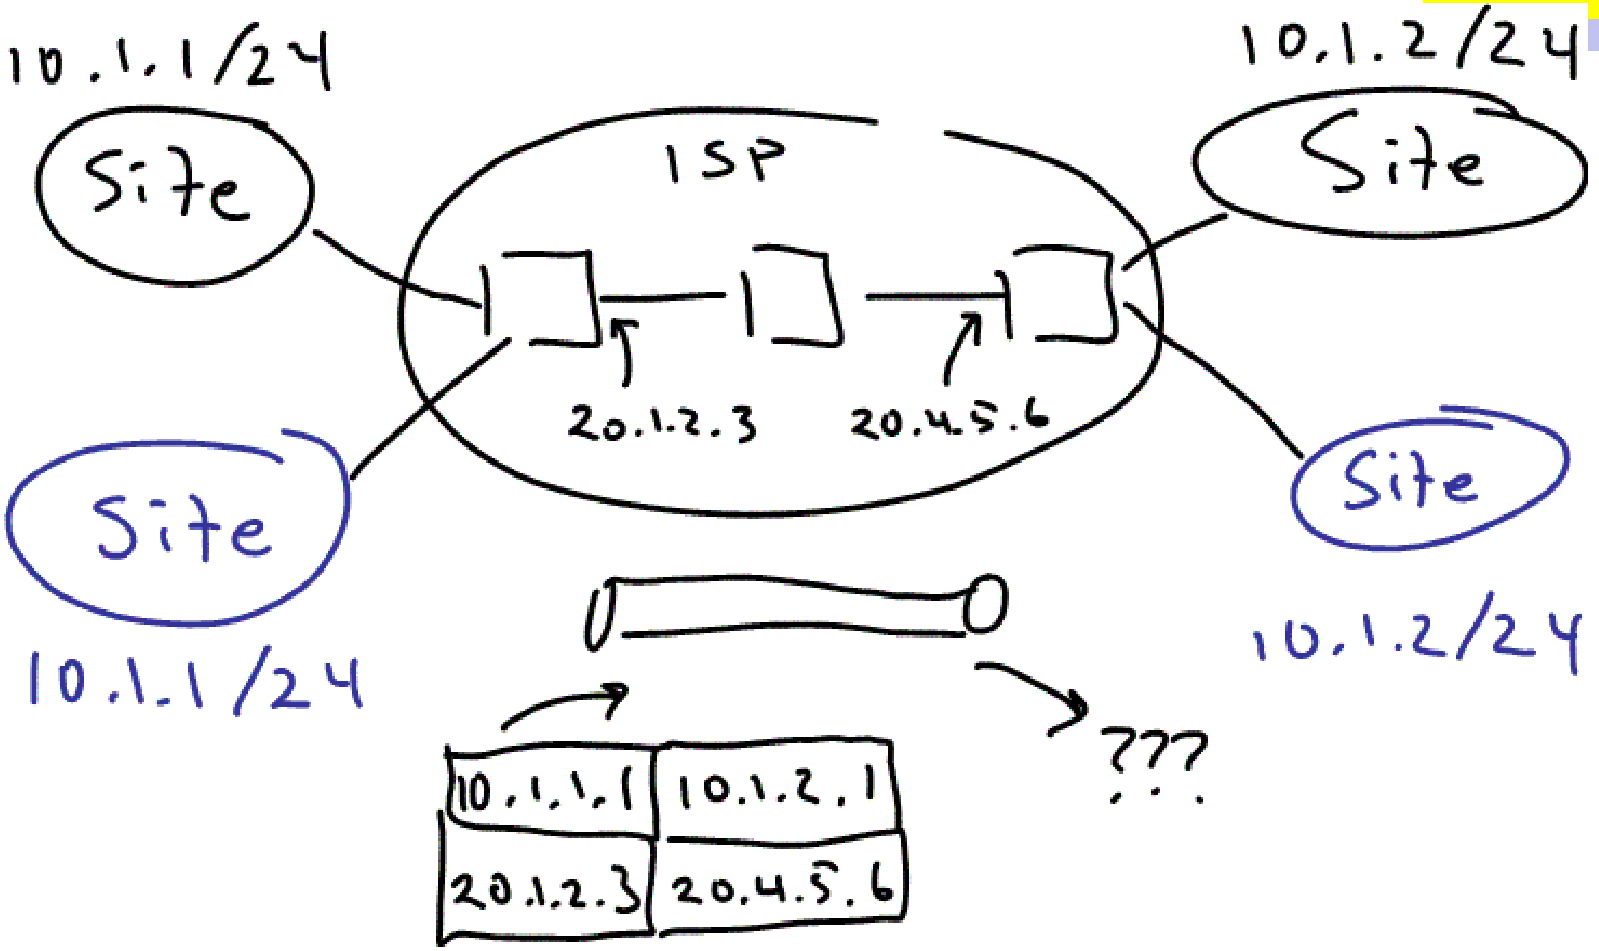

MPLS Label identifies VPN

CS419

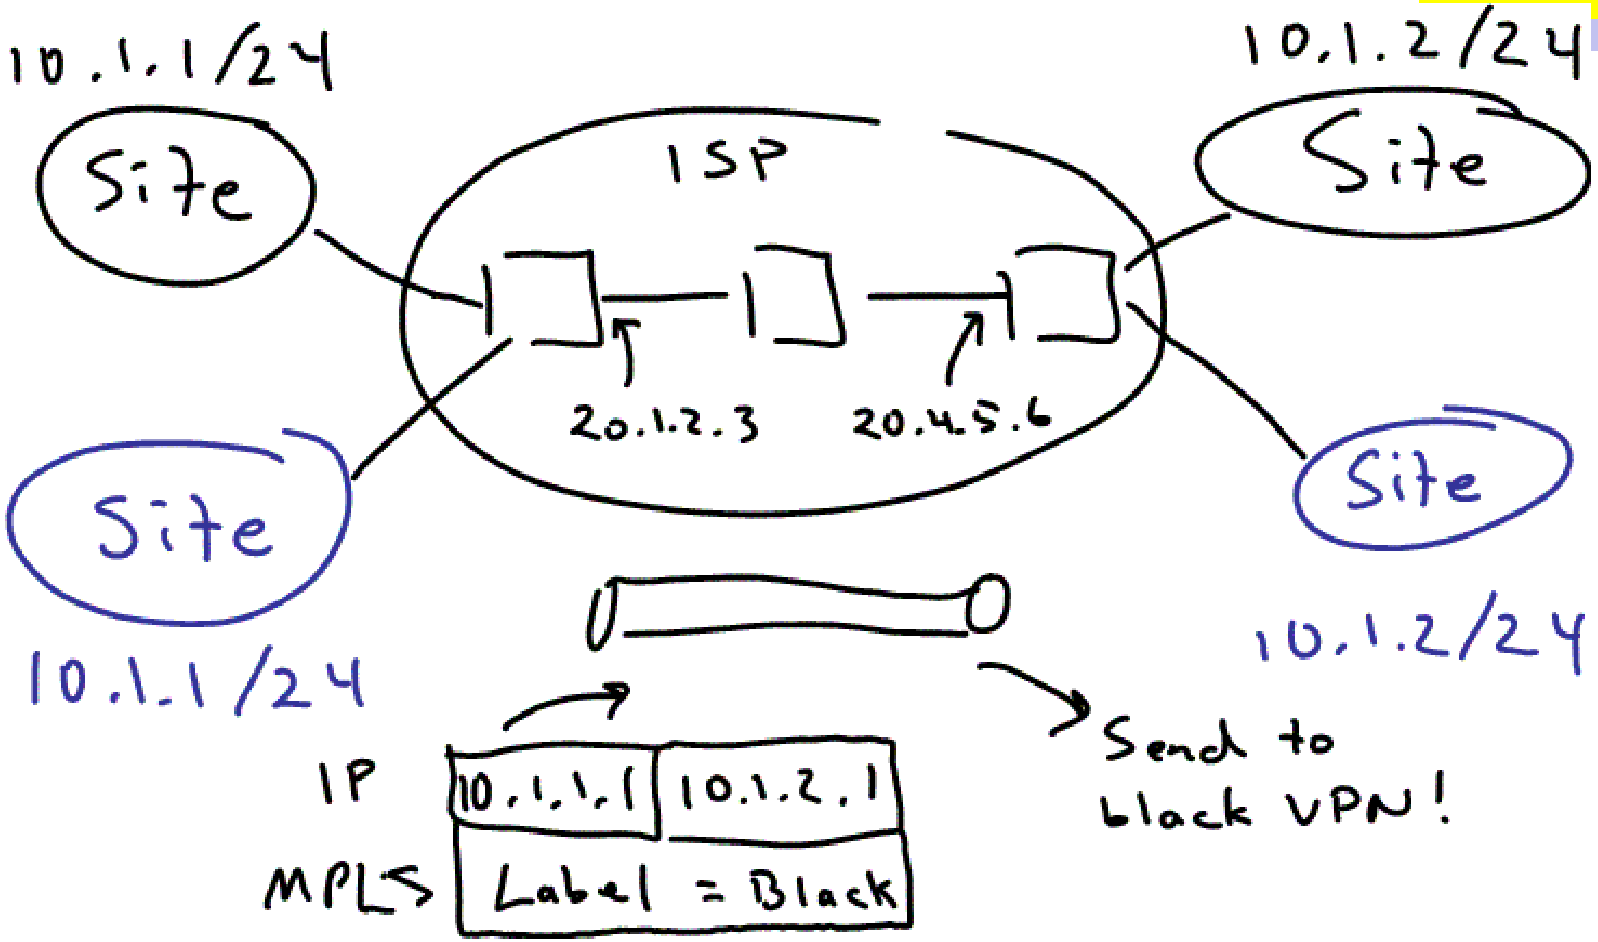

How BGP/MPLS VPNs work

CS419

- BGP updates normally carry a set of IP prefixes in the routing path
- With MPLS VPN, they carry a VPN identifier, and an MPLS tag
  - VPN identifier distinguishes overlapping address
  - MPLS tag says how to encapsulate customer's IP over MPLS
- Within MPLS, the tag both routes the packet and identifies the customer
- Tunnels are typically not secure
  - Customer assumes provider links are physically secure

A few more MPLS details

CS419

- Headers are stackable
- Uses variant of RSVP for establishing label values
- Also used these days for Traffic Engineering
  - Because can route on source and dest
  - Allows per-customer Service Level Aggrements

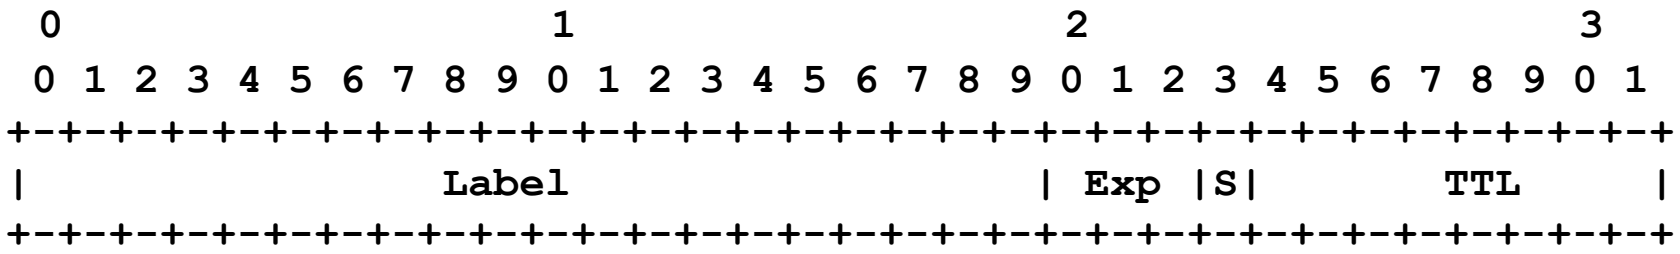

Virtual Router based L3 VPNs

CS419

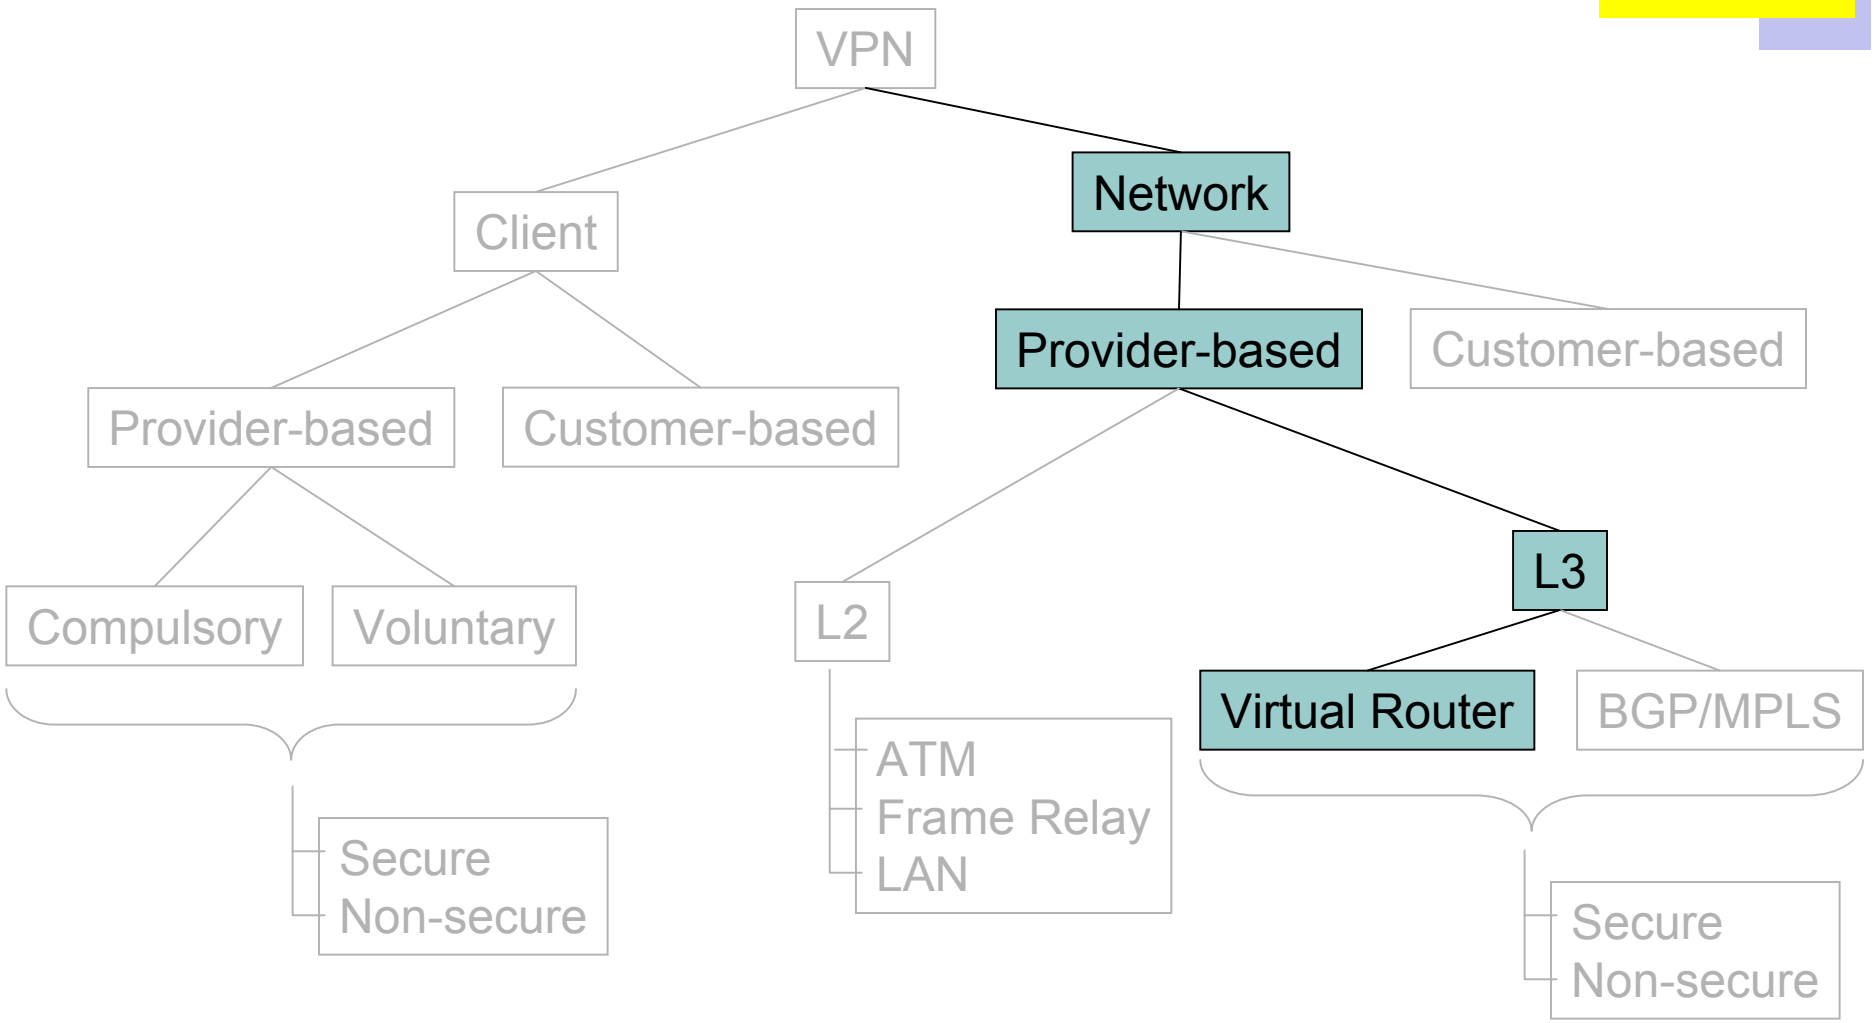

Virtual Router based L3 VPNs

CS419

- BGP/MPLS gave Cisco a huge advantage
  - Because Cisco was the BGP and MPLS expert
- Competitors' counter argument:
  - No need to couple routing technology with tunneling technology…they are separate issues
  - Simpler to use virtual routers

What is a virtual router (VR)?

CS419

- Separate logical router within a single physical router
  - Runs its own routing algorithm
  - Has its own FIB (Forwarding Information Base)
- Basic idea: Incoming tunnel identifies which VR is intended
  - If GRE, then GRE key field
  - If IPsec, then IPsec SPI field
  - If L2TP, then L2TP key field
- This is how overlapping addresses are distinguished

VR approach has discovery issues

CS419

- No standard way to configure tunnels and discover which PEs attach to which customers
  - All manually configured (via management system)
- Various proposals exist
  - Via BGP, OSPF, DNS, an LDAP database, and even IP multicast

Layer 2 LAN VPNs

CS419

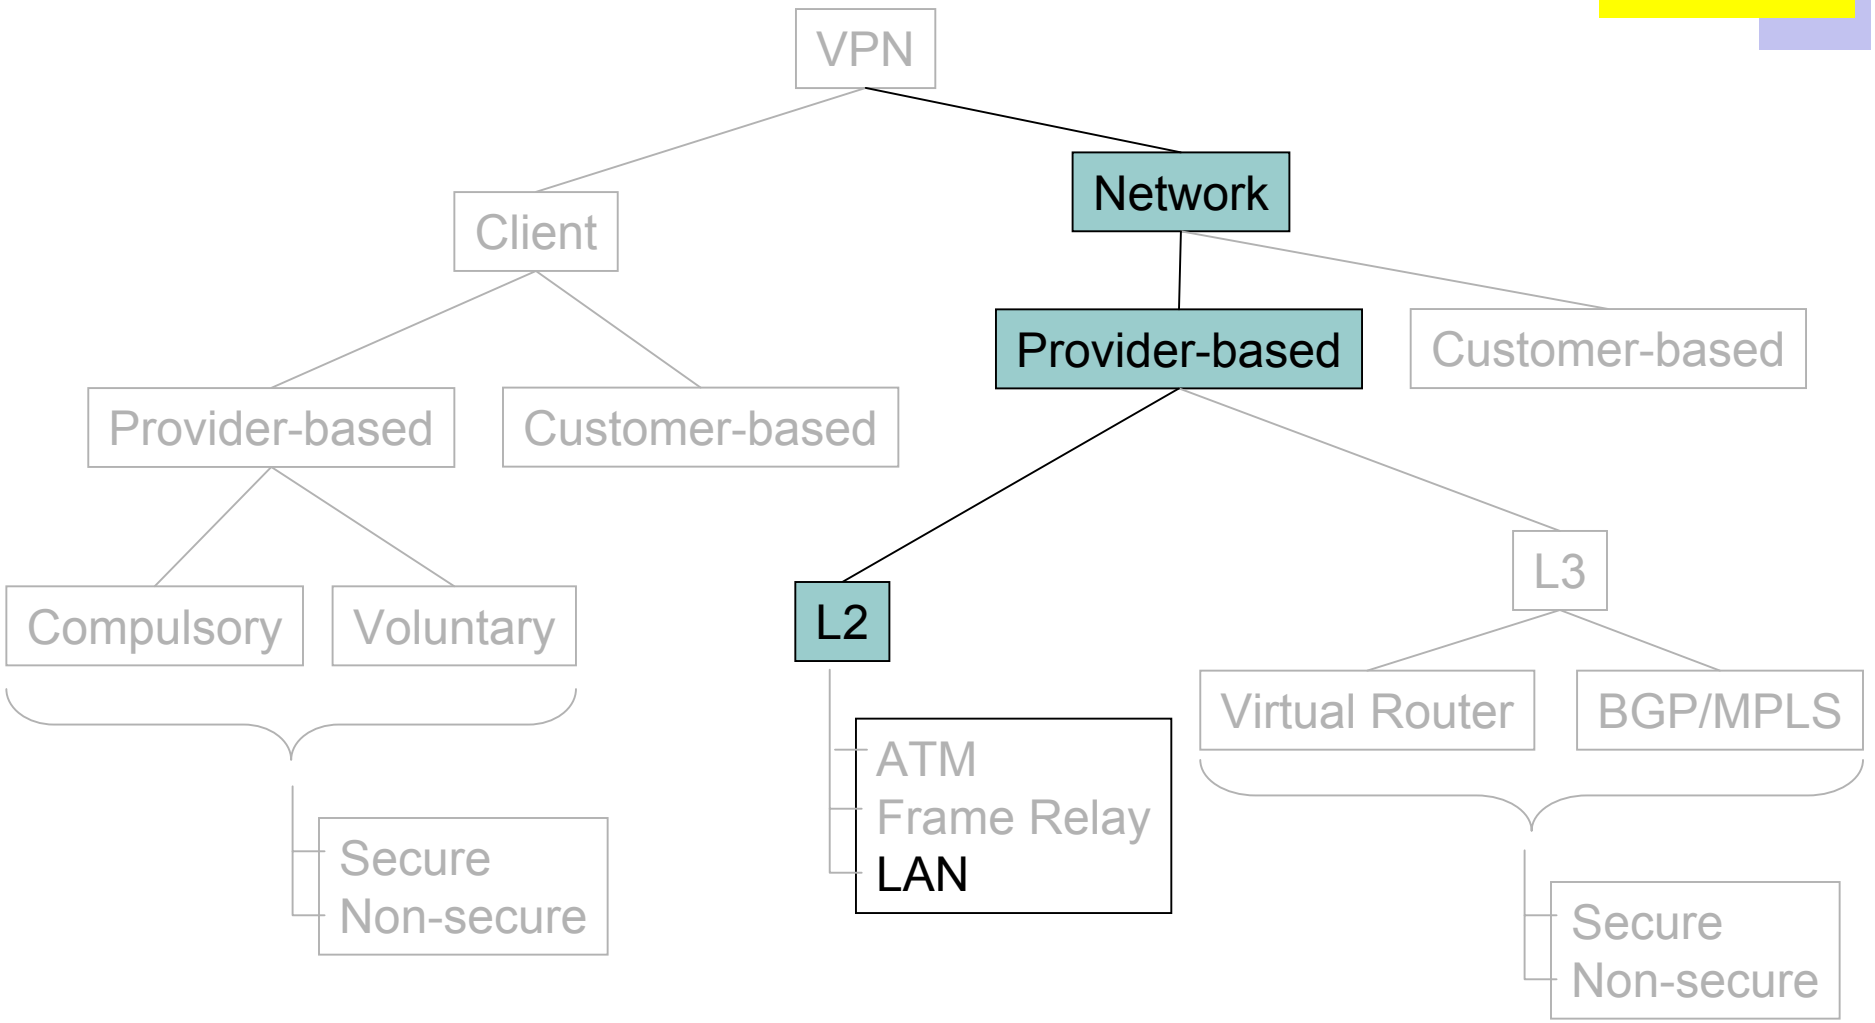

Layer 2 LAN VPNs

CS419

- Model is for PE to look like LAN to CE
- CE broadcast over LAN reaches only other CEs of the same customer
  - Thus customer can run OSPF over LAN in standard way
  - Supports multicast
  - Multi-protocol
- Uses VLAN (Virtual LAN) tags to distinguish customers
- Advantages over FR and ATM are:
  - Ethernet is more common interface
  - Supports broadcast/multicast

What is a VLAN?

CS419

- A “virtual LAN”: makes a single physical look like multiple LANs
- Virtual LAN and priority capabilities are provided by 802.1Q/p:
  - a VLAN tag is provided by 802.1Q to identify VLAN membership
    - Limited to 4096 VLANs – this is a potential scalability issue
  - the VLAN tag has a 3-bit priority field that allows 8 possible service classes (matches DiffServ’s 8 possible classes)

Why VLANs?

CS419

- LAN scalability:
  - limits broadcast domains (limits broadcast storms);
  - also limits multicast, chatty protocols, etc., reducing overall network traffic.
- Network efficiencies: traffic flows from different VLANS can be segregated
- Allows non-physical grouping of nodes that share similar resources
- Allows easy changing of LAN membership
- Reduces the amount of level 3 (IP) routing
- Security: limits snooping; authentication required (via GVRP) to join VLAN

More to the point

CS419

- Ethernet has gotten very fast
  - GigE common
  - 10gig Ethernet coming (optical)
- We can put much more on an Ethernet, so we need to segregate
- These days, site networks are composed of ethernet switches and VLANs, not routers and subnets!

Typical site configuration (from Cisco)

CS419

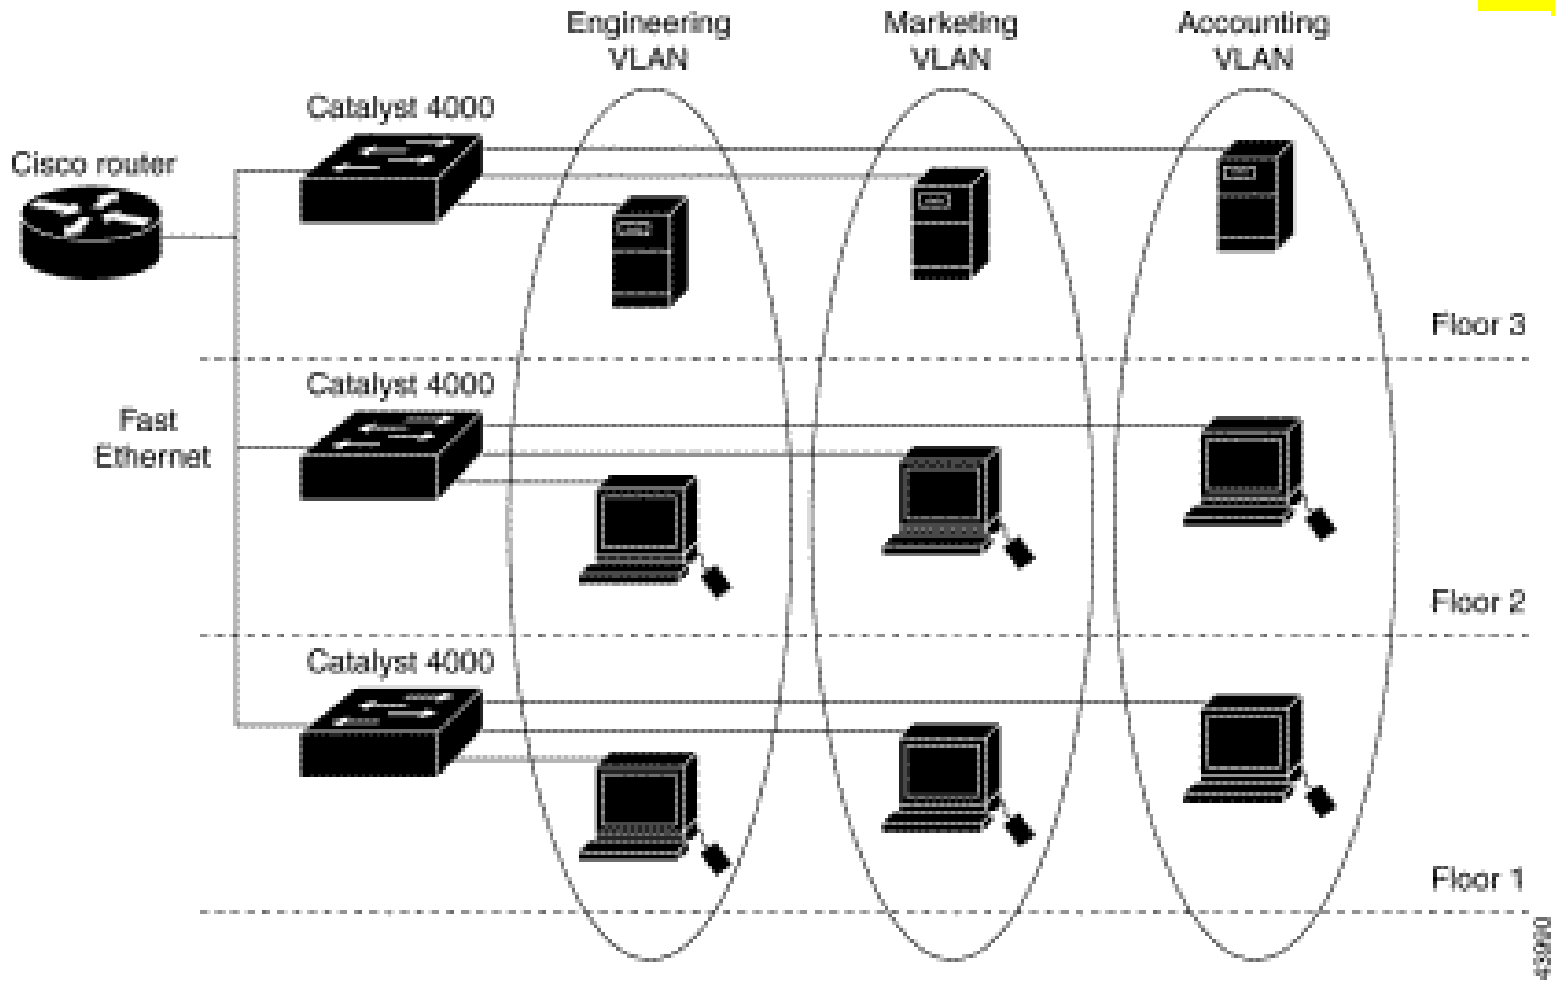

VLAN Header

CS419

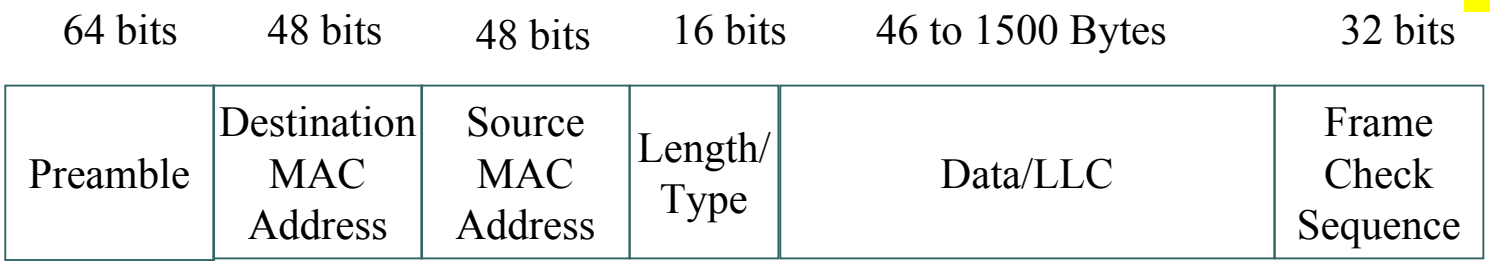

Original Ethernet Frame Structure

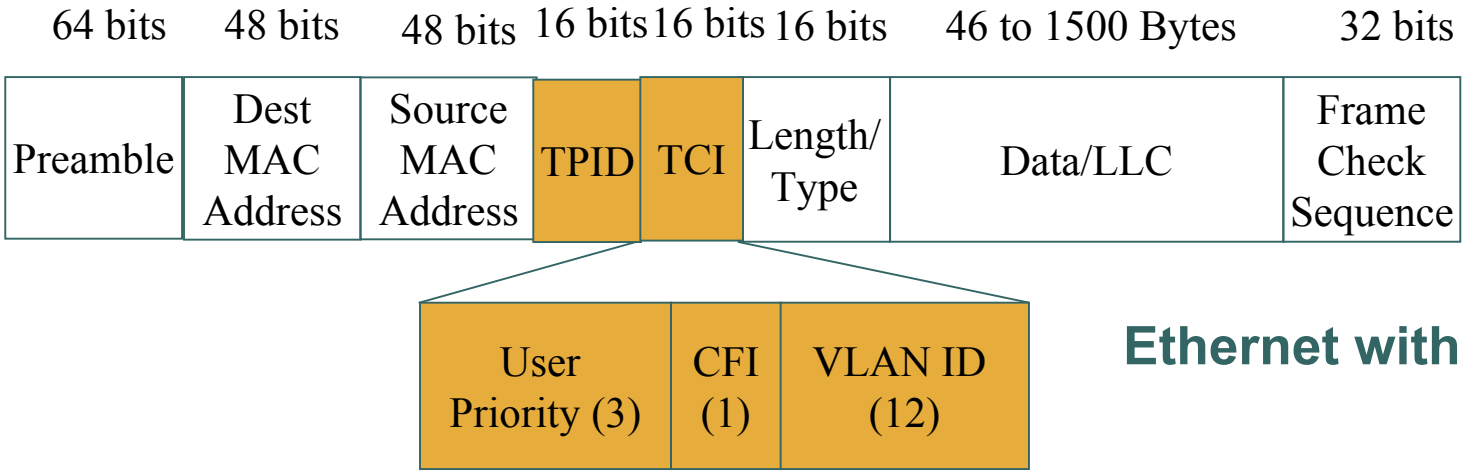

Ethernet with VLAN

Meta-Point: Its all about tunnels!

CS419

- In this lecture we saw a lot of tunnels
  - IPsec, MPLS, GRE, L2TP
- I said before that the Internet has two ways to scale:
  - Hierarchy and caching
- It has a third way:
  - Tunnels!

Tunnels are scalable

CS419

- Tunnels prevents the “middle” from having to know details of the “edge”
  - But in a manner that is more flexible than hierarchy
  - Hierarchy forces a structure from the middle (top)
  - Tunnels “cut through” the middle transparently
- Tunnels have been introduced piecemeal
  - We still don’t have a coherent architecture for them . . .