

Cristina Nita-Rotaru



CS4700/5700: Network fundamentals

Internet architecture.



1: Internet architecture

Organizing network functionality

- ▶ **Networks are built from many components**
 - ▶ Networking technologies
 - ▶ Ethernet, Wifi, Bluetooth, Fiber Optic, Cable Modem, DSL
 - ▶ Network types (different characteristics, performance, designs)
 - ▶ Circuit switch, packet switch
 - ▶ Wired, Wireless, Optical, Satellite
 - ▶ Applications
 - ▶ Email, Web (HTTP), FTP, BitTorrent, VoIP
- ▶ **How does all work together?!**

A possible design

Web



Email




Bittorrent



VoIP



- This is a nightmare scenario
- Huge amounts of work to add new apps or media
- Limits growth and adoption



Ethernet



802.11



Bluetooth




Cellular

Internet architecture

What if application endpoints are on different type of networks

Bittorrent



Ethernet

Application endpoints
may not be on the
same media

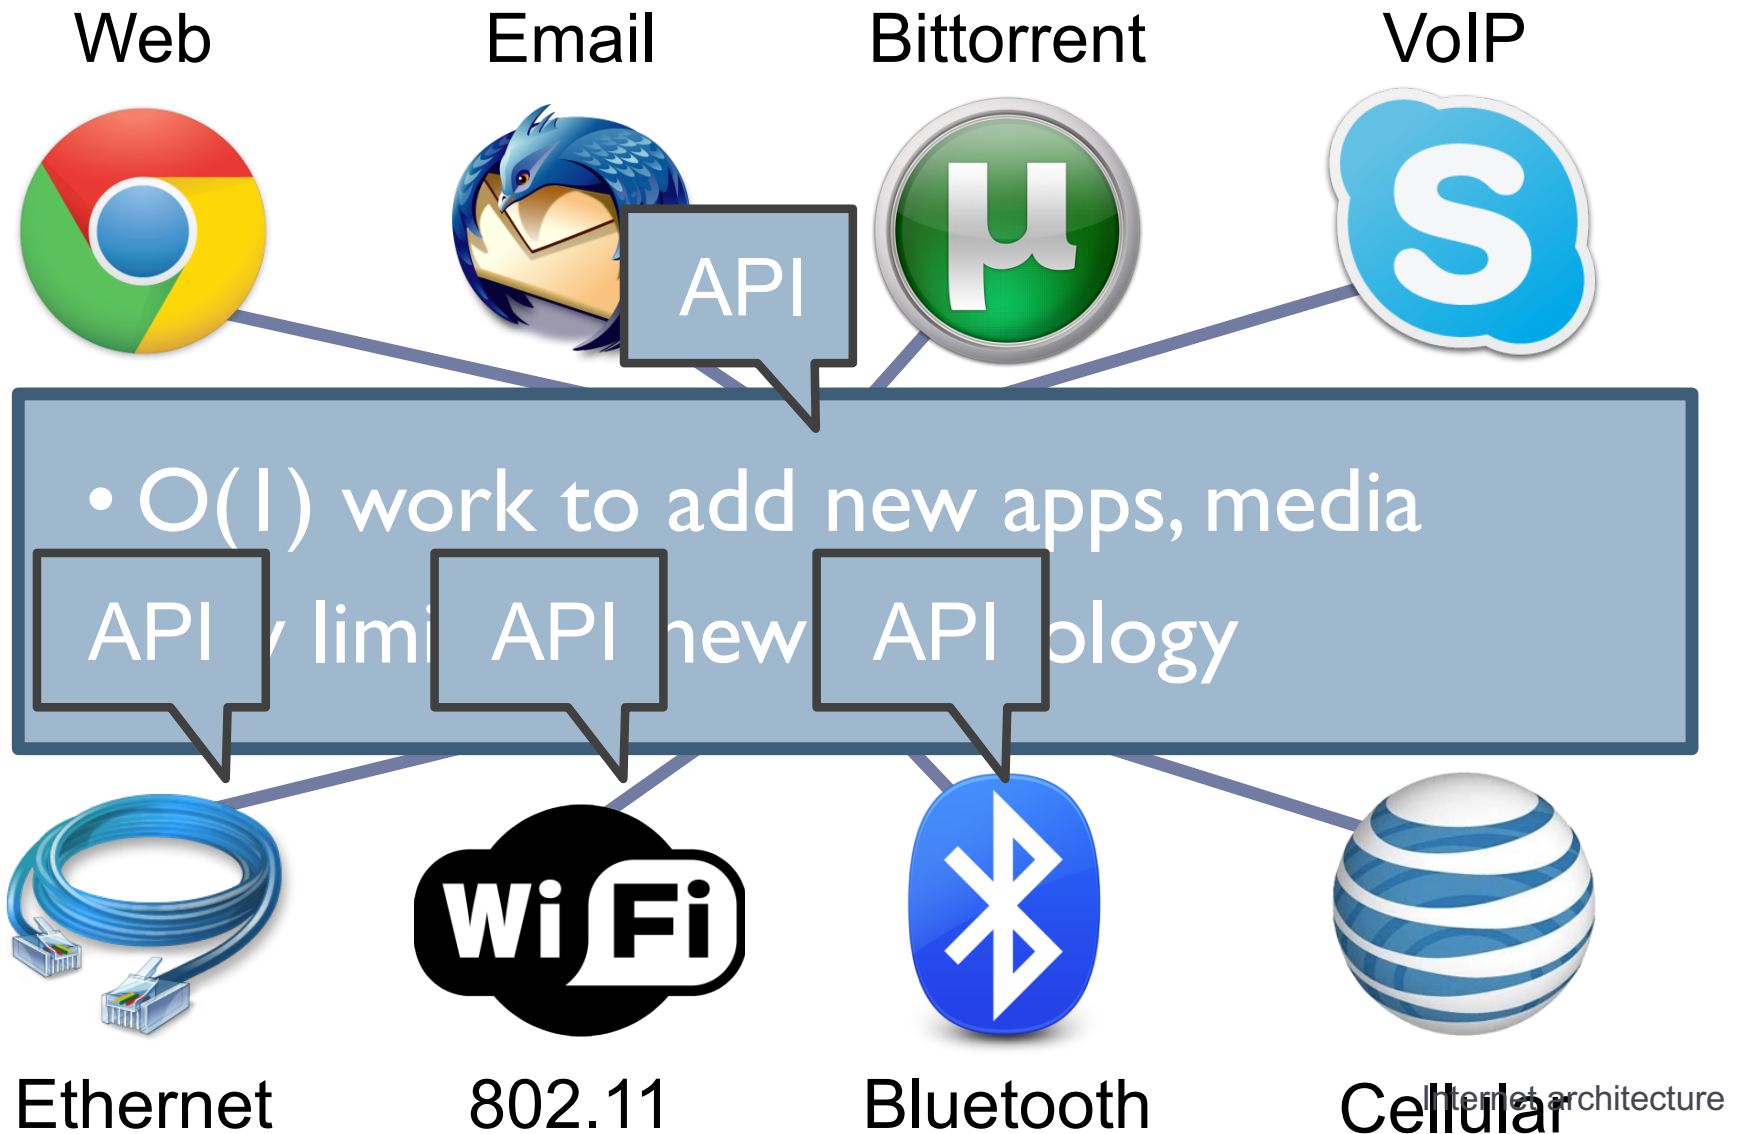
Bittorrent




802.11

Internet architecture

A better design: Use indirection




Layered Network Stack



- ▶ **Modularity**
 - ▶ Does not specify an implementation
 - ▶ Instead, tells us how to organize functionality
- ▶ **Encapsulation**
 - ▶ Interfaces define cross-layer interaction
 - ▶ Layers only rely on those below them
- ▶ **Flexibility**
 - ▶ Reuse of code across the network
 - ▶ Module implementations may change
- ▶ **Unfortunately, there are tradeoffs**
 - ▶ Interfaces hide information
 - ▶ Interfaces may hurt performance...


Key questions

- ▶ How do we divide functionality into layers?
 - ▶ Routing
 - Security
 - Fairness
 - And many more...
 - ▶ Congestion control
 - ▶ Error checking
- ▶ How do we distribute functionality across devices?
 - ▶ Example: who is responsible for security?




ISO OSI Model

OSI: Open Systems Interconnect Model




Layer features




- ▶ **Service**
 - ▶ What does this layer do?
- ▶ **Interface**
 - ▶ How do you access this layer?
- ▶ **Protocol**
 - ▶ How is this layer implemented?

Physical layer



- ▶ **Service**
 - ▶ Move information between two systems connected by a physical link
- ▶ **Interface**
 - ▶ Specifies how to send one bit
- ▶ **Protocol**
 - ▶ Encoding scheme for one bit
 - ▶ Voltage levels
 - ▶ Timing of signals
- ▶ **Examples: coaxial cable, fiber optics, radio frequency transmitters**

Data link layer



- ▶ **Service**

- ▶ Data framing: boundaries between packets
- ▶ Media access control (MAC)
- ▶ Per-hop reliability and flow-control


- ▶ **Interface**

- ▶ Send one **packet** between two hosts connected to the **same media**

- ▶ **Protocol**


- ▶ Physical addressing (e.g. MAC address)
- ▶ **Examples: Ethernet, Wifi, DOCSIS**

Network layer




- ▶ **Service**
 - ▶ Deliver packets across the network
 - ▶ Handle fragmentation/reassembly
 - ▶ Packet scheduling
 - ▶ Buffer management
- ▶ **Interface**
 - ▶ Send one packet to a specific destination
- ▶ **Protocol**
 - ▶ Define globally unique addresses
 - ▶ Maintain routing tables
- ▶ **Example: Internet Protocol (IP), IPv6**

Transport layer




- ▶ **Service**
 - ▶ Multiplexing/demultiplexing
 - ▶ Congestion control
 - ▶ Reliable, in-order delivery
- ▶ **Interface**
 - ▶ Send message to a destination
- ▶ **Protocol**
 - ▶ Port numbers
 - ▶ Reliability/error correction
 - ▶ Flow-control information
- ▶ **Examples: UDP, TCP**

Session layer



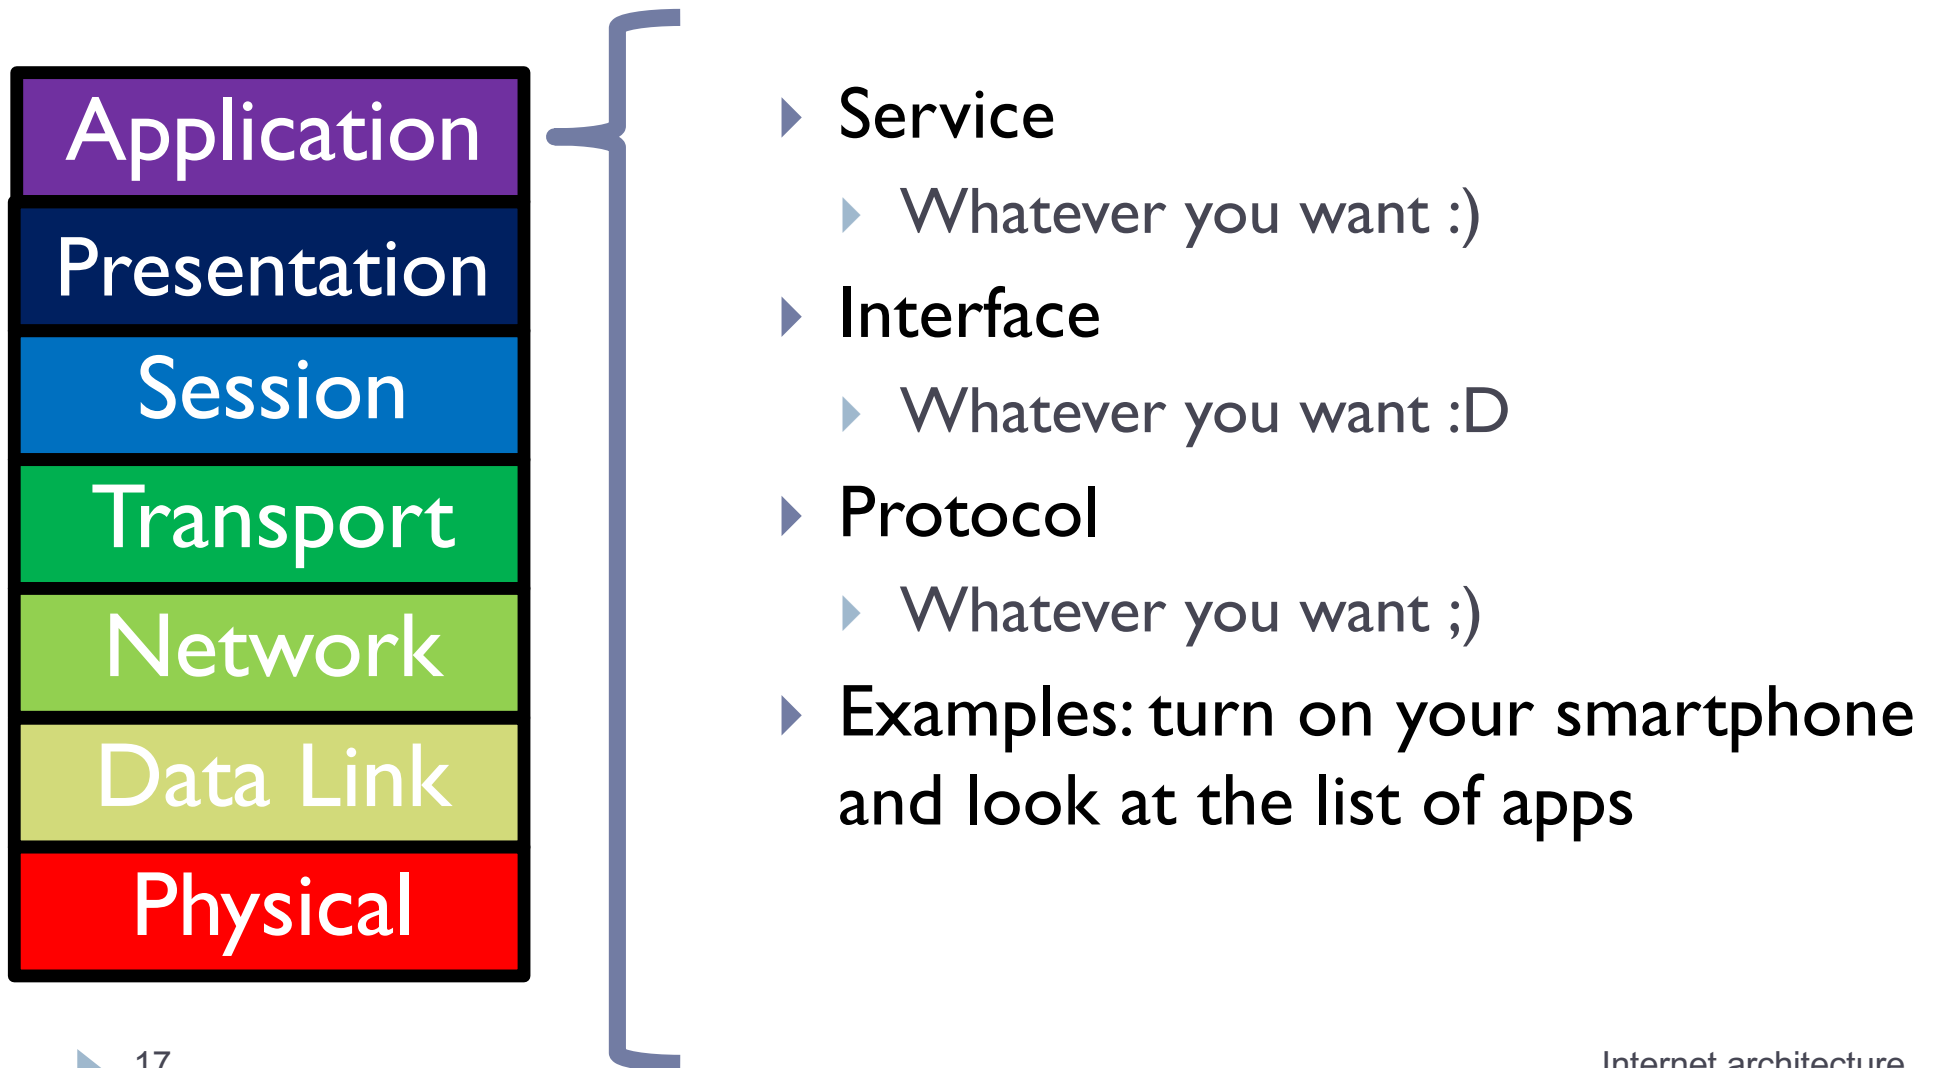
- ▶ **Service**
 - ▶ Access management
 - ▶ Synchronization
- ▶ **Interface**
 - ▶ It depends...
- ▶ **Protocol**
 - ▶ Token management
 - ▶ Insert checkpoints
- ▶ **Examples: none**

Presentation layer




- ▶ **Service**
 - ▶ Convert data between different representations
 - ▶ E.g. big endian to little endian
 - ▶ E.g. Ascii to Unicode
- ▶ **Interface**
 - ▶ It depends...
- ▶ **Protocol**
 - ▶ Define data formats
 - ▶ Apply transformation rules
- ▶ **Examples: none**

Application layer

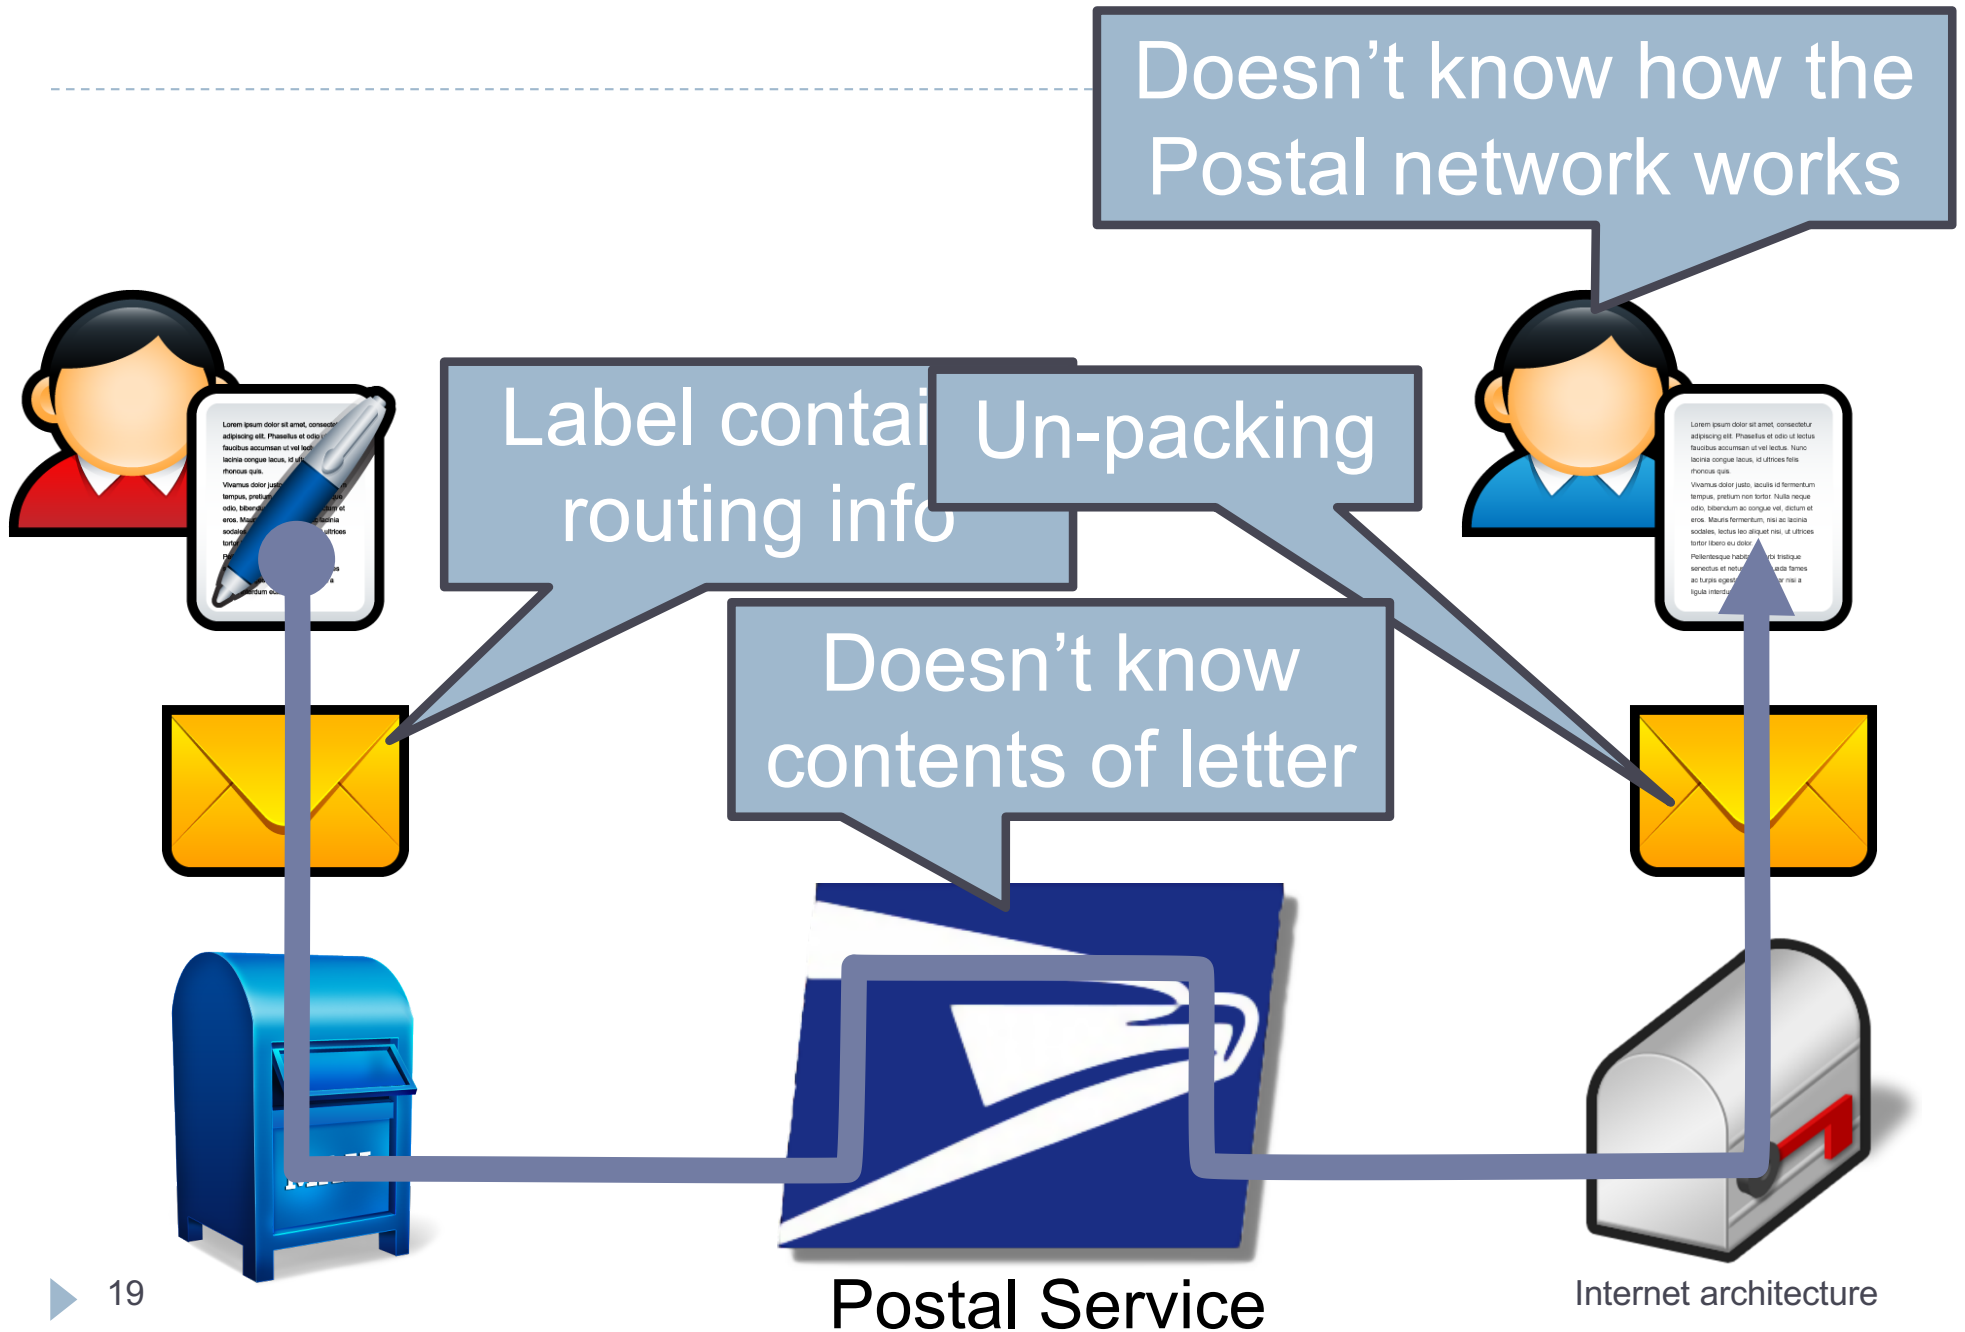


Encapsulation

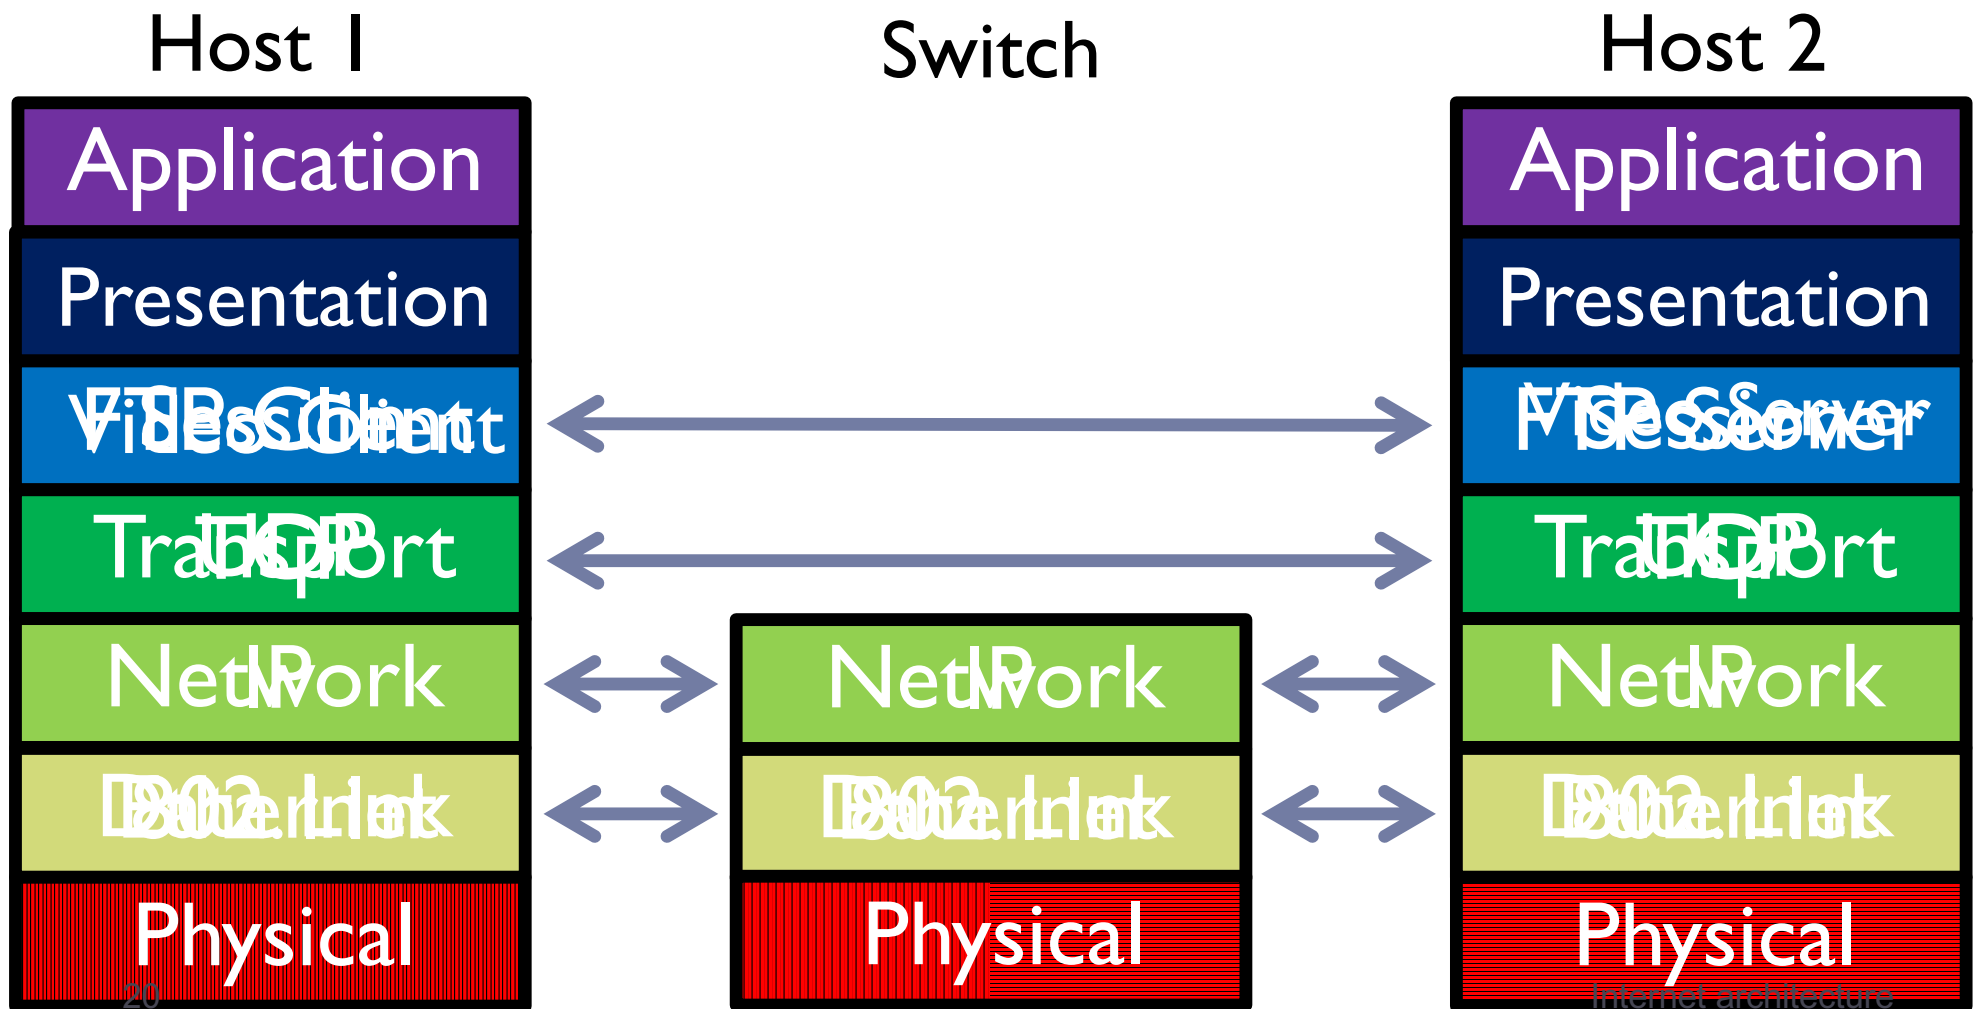
How does data move through the layers?




Real life analogy



Network stack in practice



Encapsulation, revisited



Hourglass design

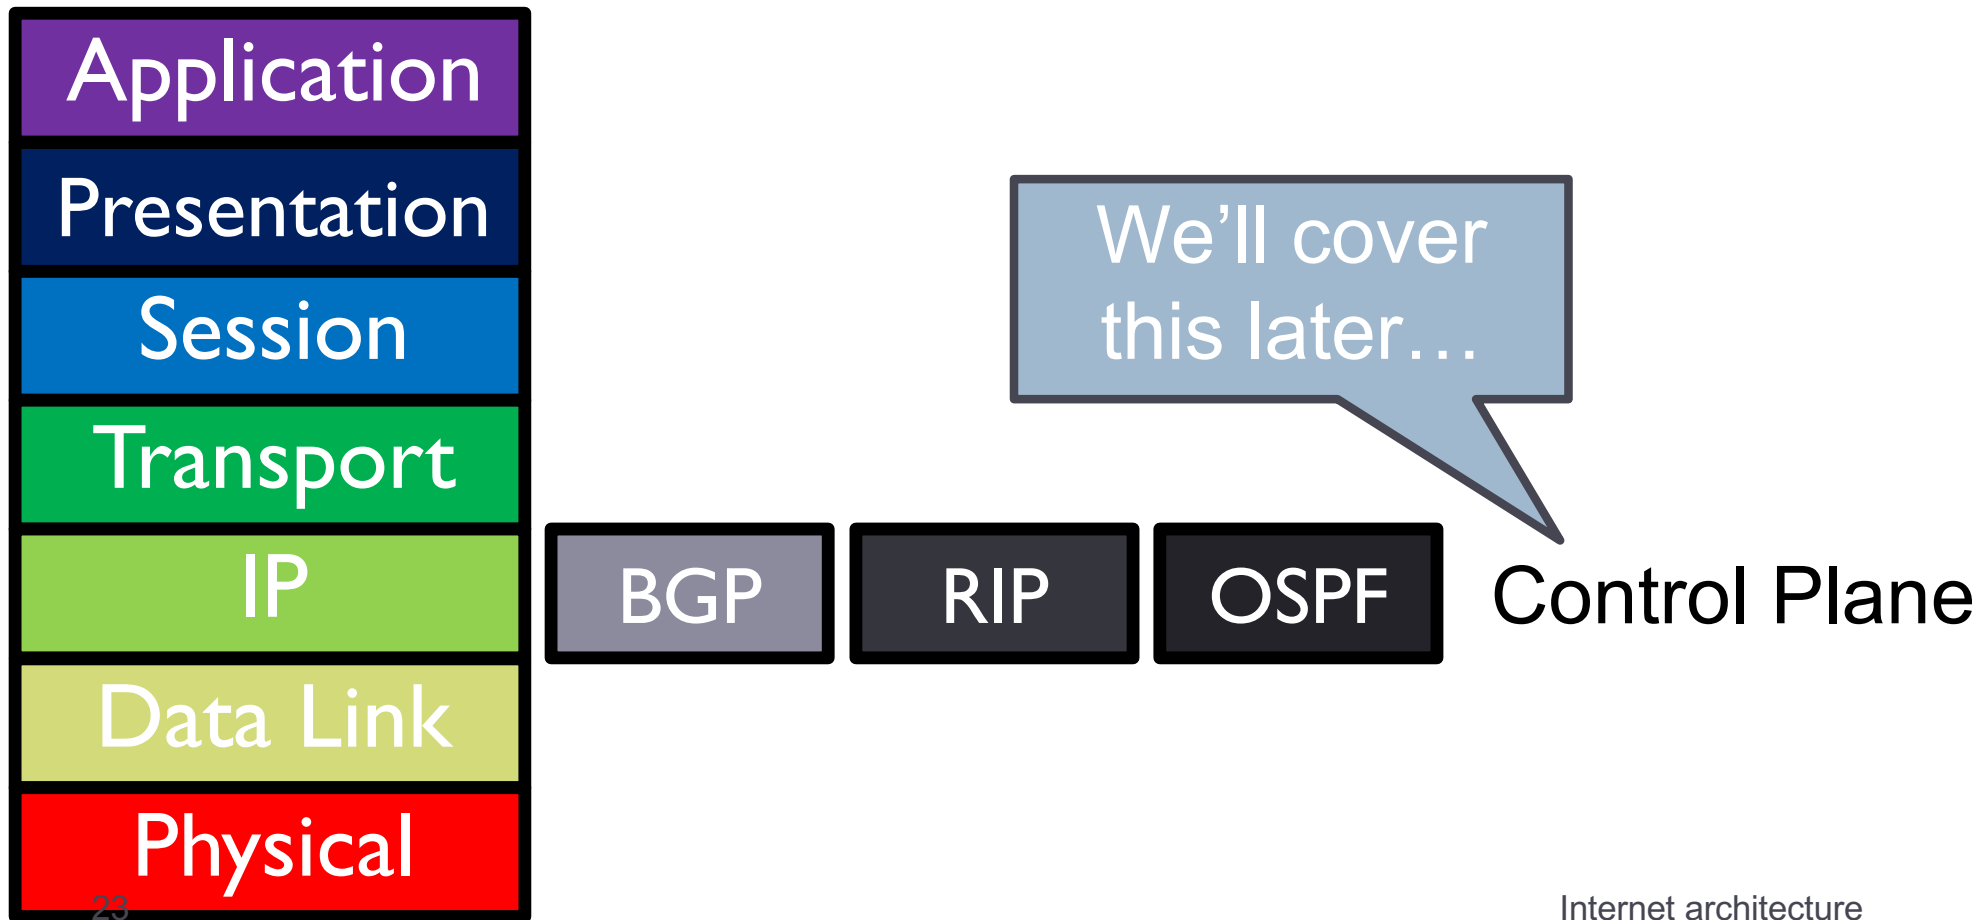
- One Internet layer means all networks interoperate
- All applications function on an
- Room for development above deploying IPv6...
- But, changing IP is insanely hard

Think about the difficulty of deploying IPv6...

Fiber, Coax, Twisted Pair, Radio, ...

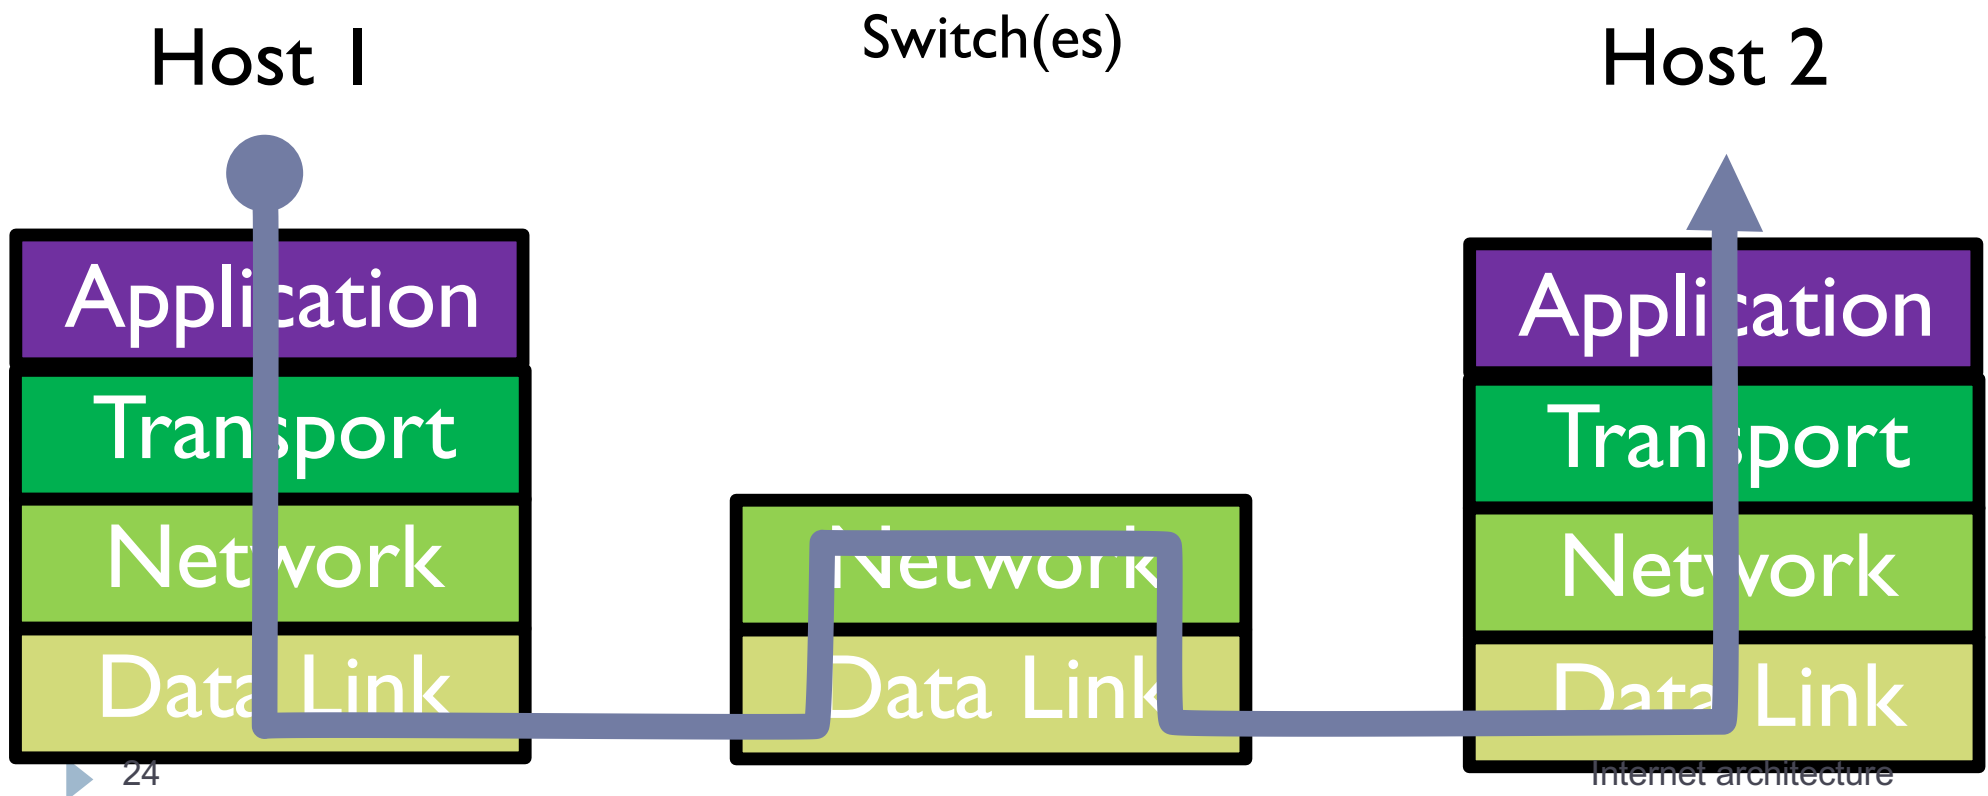
Orthogonal planes

Control plane: How Internet paths are established



Orthogonal planes

Data plane: How data is **forwarded** over Internet paths



Reality check

- ▶ The layered abstraction is very nice
- ▶ Does it hold in reality?

No.



Firewalls

- Analyze application layer headers



Transparent Proxies

- Simulate application endpoints within the network



NATs


- Break end-to-end network reachability

From layers to eating cake

- ▶ IP gives us best-effort datagram forwarding
 - ▶ So simple anyone can do it
 - ▶ Large part of why the Internet has succeeded
 - ▶ ...but it sure isn't giving us much
- ▶ Layers give us a way to **compose** functionality
 - ▶ Example: HTTP over TCP for Web browsers with reliable connections
- ▶ ...but they do not tell us where (in the network) to implement the functionality

Where to place functionality

- ▶ How do we distribute functionality across devices?
 - ▶ Example: who is responsible for security?

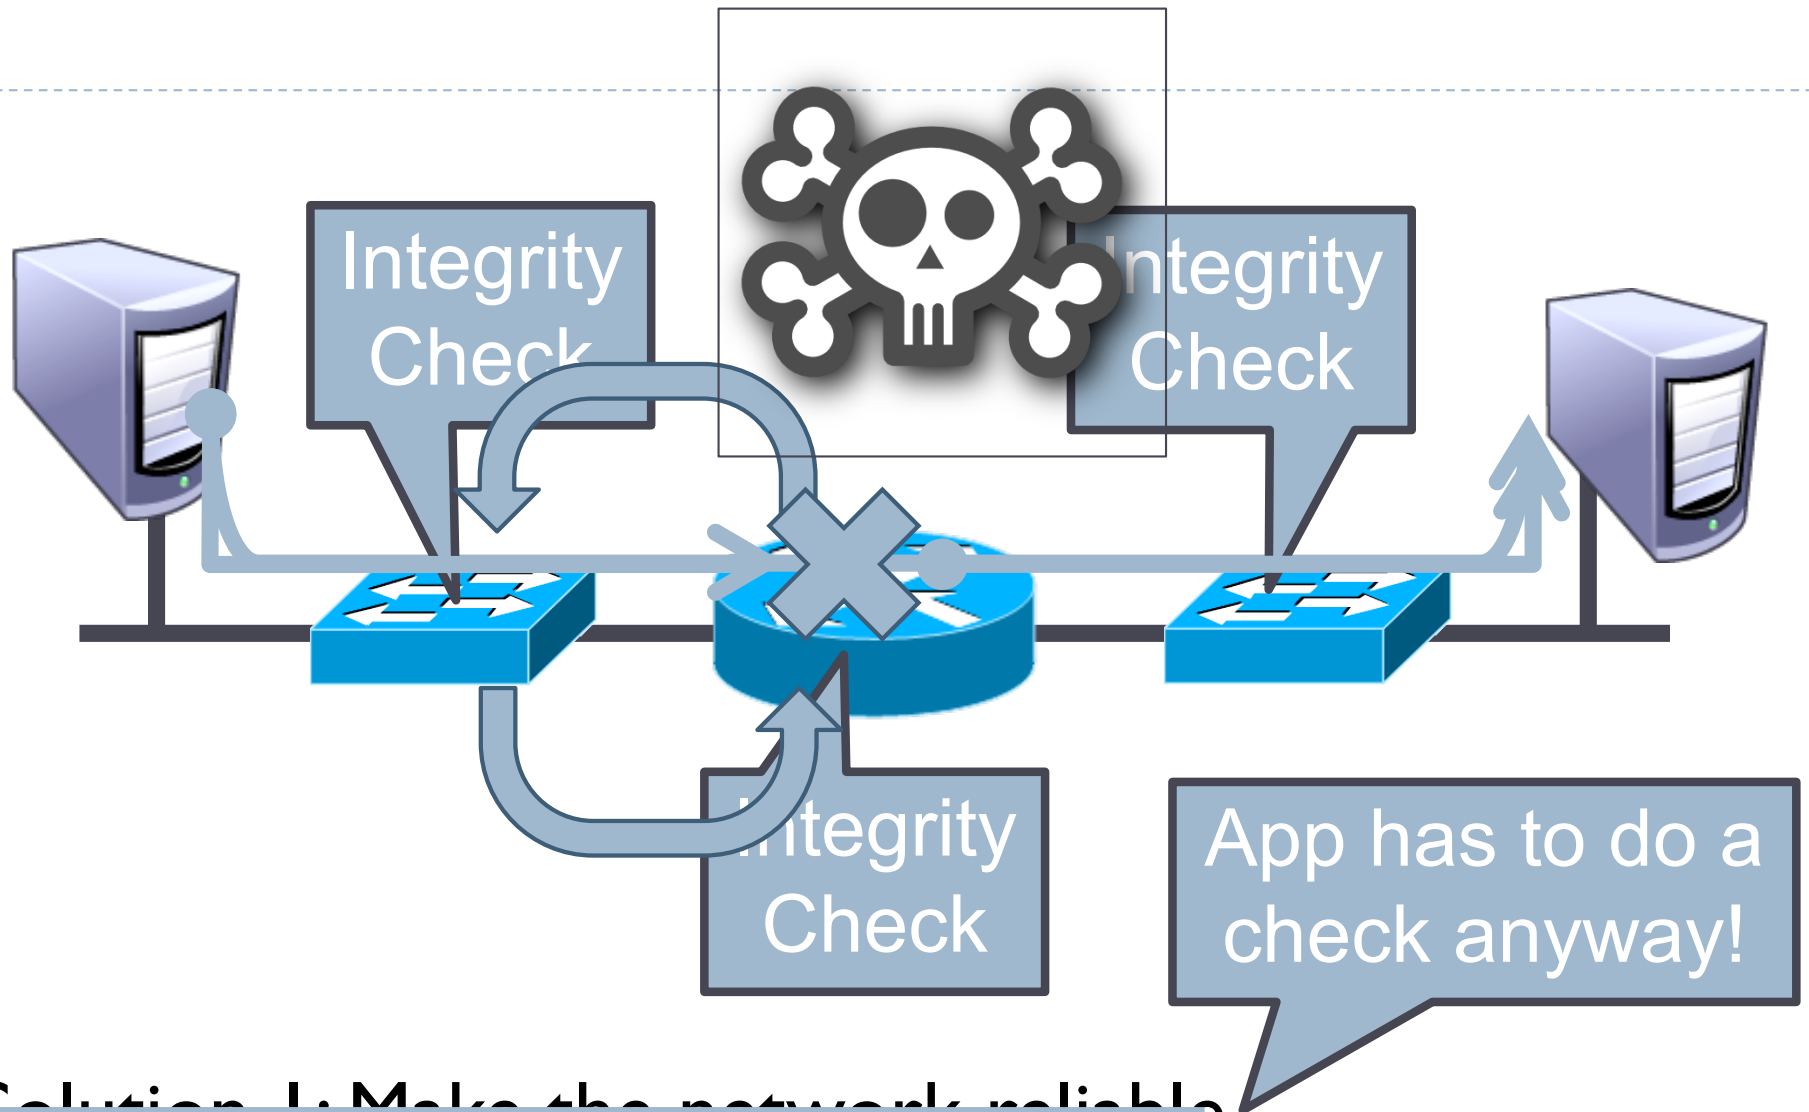


- “The End-to-End Arguments in System Design”
 - ▣ Saltzer, Reed, and Clark
 - ▣ The Sacred Text of the Internet
- ▶ ²⁷ Endlessly debated by researchers and engineers Internet architecture

Basic observation

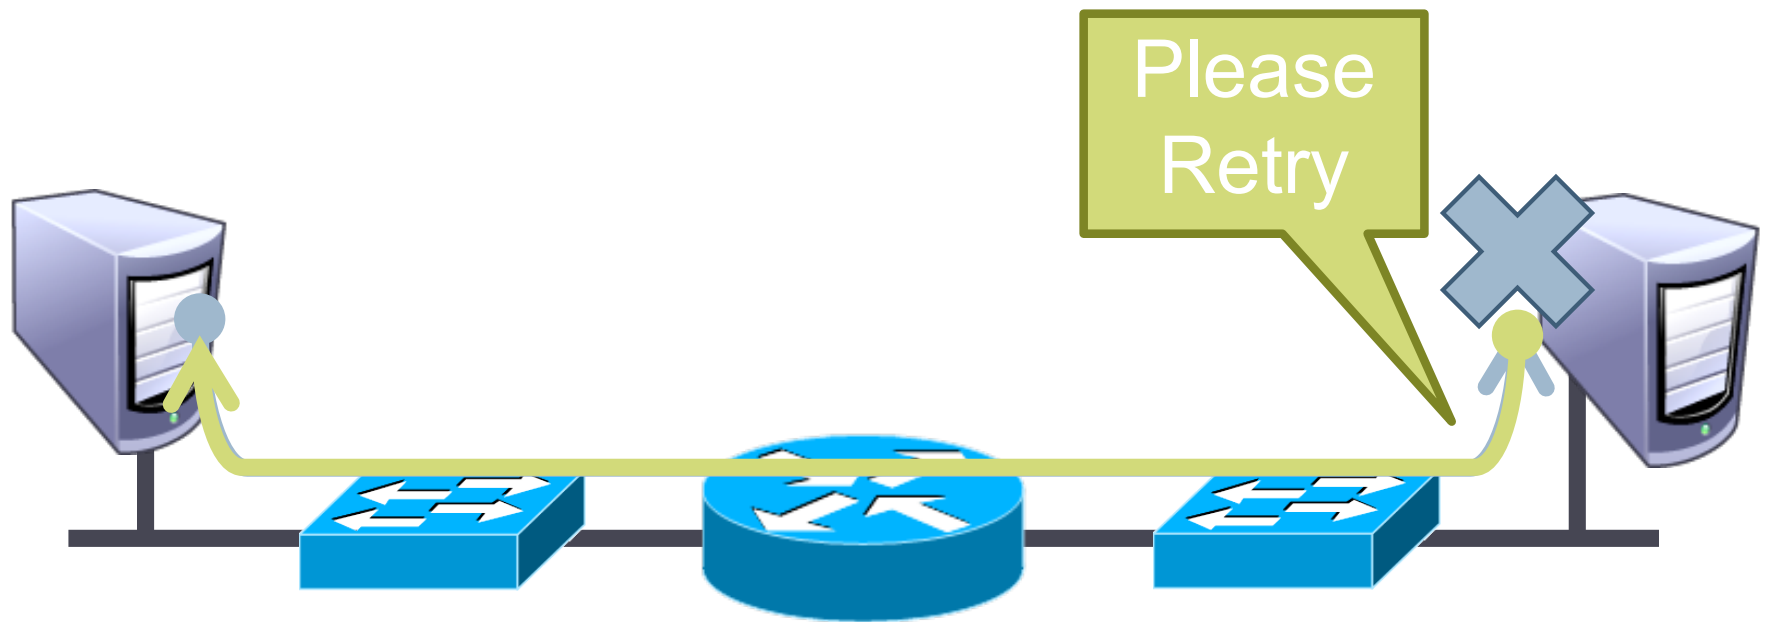
- ▶ **Some applications have end-to-end requirements**
 - ▶ Security, reliability, etc.
- ▶ **Implementing this stuff inside the network is hard**
 - ▶ Every step along the way must be fail-proof
 - ▶ Different applications have different needs
- ▶ **End hosts...**
 - ▶ Can't depend on the network
 - ▶ Can satisfy these requirements without network level support

Example: Reliable File Transfer



- ❑ ~~Solution 1: Make the network reliable~~
- ❑ Solution 2: App level, end-to-end check, retry on failure

Example: Reliable file transfer



Full functionality can be built at App level

- ~~Solution 1: Make the network reliable~~
- Solution 2: App level, end-to-end check, retry on failure

Conservative interpretation

“Don’t implement a function at the lower levels of the system unless it can be completely implemented at this level” (Peterson and Davie)

Basically, unless you can completely remove the burden from end hosts, don’t bother

Radical interpretation


- ▶ Don't implement anything in the network that can be implemented correctly by the hosts
- ▶ Make network layer absolutely minimal
- ▶ Ignore performance issues

Moderate interpretation

- ▶ Think twice before implementing functionality in the network
- ▶ If hosts can implement functionality correctly, implement it at a lower layer only as a performance enhancement
- ▶ But do so only if it does not impose burden on applications that do not require that functionality...
- ▶ ...and if it doesn't cost too much \$ to implement

Another example: Anonymity

- ▶ Should we implement this in the network?
- ▶ How about at the endpoints?



Reality check, again

- ▶ Layering and E2E principals regularly violated



Firewalls



Transparent Proxies




NATs

- Conflicting interests
 - ▣ Architectural purity
 - ▣ Commercial necessity

Takeaways

- ▶ **Layering for network functions**
 - ▶ Helps manage diversity in computer networks
 - ▶ Not optimal for everything, but simple and flexible
- ▶ **Narrow waist ensures interoperability, enables innovation**
- ▶ **E2E argument (attempts) to keep IP layer simple**
- ▶ **Think carefully when adding functionality into the network**



Reading

- ▶ J. Saltzer, D. Reed, and D. Clark, "End-to-end Arguments in System Design". ACM Transactions on Computer Systems (TOCS), Vol. 2, No. 4, 1984, pp. 195-206.