CS4450

Computer Networks:
Architecture and Protocols

Lecture 5
- Three Architectural Principles
- Design Goals

Rachit Agarwal

Announcements

- **Exam conflict:**
  - Today is the last day to announce your exam conflicts.
    - Exam 1: 09/21
    - Exam 2: 10/24
    - Exam 3: 12/05
  - Send an email to cs4450-staff@cornell.edu
  - Thank you to those who already sent us an email
  - We will send an email to all those who have a conflict
- Problem set 1 solutions released (on Ed discussions)
- Problem set 2 released (on course website)

Context for Today's Lecture

- So far, we have discussed several high-level concepts
  - Network sharing
  - End-to-end working of the Internet
  - Addressing, Routing, Switch/Router functionality, etc.
- And, have dived deep into several topics:
  - Circuit switching and packet switching (especially the "why")
  - Delays (transmission, propagation)
- **Problem set 2—first two questions—dive deeper into these. Goals:**
  - Problem 1: build a deep understanding of delays
  - Problem 2: build a deep understanding of circuit vs packet switching
  - Problem 3, 4, 5, 6: next 2 lectures
- **Today: Continue to lay the foundation for rest of the course**

Goals for Today's Lecture

- **Wrap up the three architectural principles:**
    - Layering
    - End-to-end principle
    - Fate Sharing principle
- **Design goals for computer networks:**
    - Eight of them

Quick recap

Recap: four fundamental problems!

- **Locating the destination**: Naming, addressing
    - Mapping of names to addresses using Domain Name System
- **Finding a path to the destination**: Routing
    - Distributed algorithm that computes and stores routing tables
- **Sending data to the destination:** Forwarding
    - Input queues, virtual output queues, output queues
    - Enablers: Packet header (address), and routing table (outgoing link)
- **Reliability:** Failure handling
    - Not much discussion, but **the** question: hosts or networks?

Recap: the final piece in the story — Host network stack

Of Sockets and Ports

- When a process wants access to the network, it opens a socket, which is associated with a port
- **Socket:** an OS mechanism that connects processes to the network stack
- **Port:** number that identifies that particular socket
- The port number is used by the OS to direct incoming packets

Recap: "Thinking" Network System Modularity

- Applications deal with data
- End-host network **stacks** move data from applications to the fabric
- Network **fabric** delivers data between **network stacks**
- **Network (stack + fabric)** delivers data **between applications**
- What is the **interface** between applications and network stacks?
  - **Sockets**
- What is the **interface** between network stacks and network fabric?
  - **Packet headers**
- The <u>right</u> way to think about sockets and packets

Network Modularity Decisions

- How to break system into modules?
  - Classic decomposition into tasks

- Where are modules implemented?
  - Hosts?
  - Routers?
  - Both?

- Where is state stored?
  - Hosts?
  - Routers?
  - Both?

Leads to three design principles

- How to break system into modules
    - **Layering**
- Where are modules implemented
    - **End-to-End Principle**
- Where is state stored?
    - **Fate-Sharing**

Layering

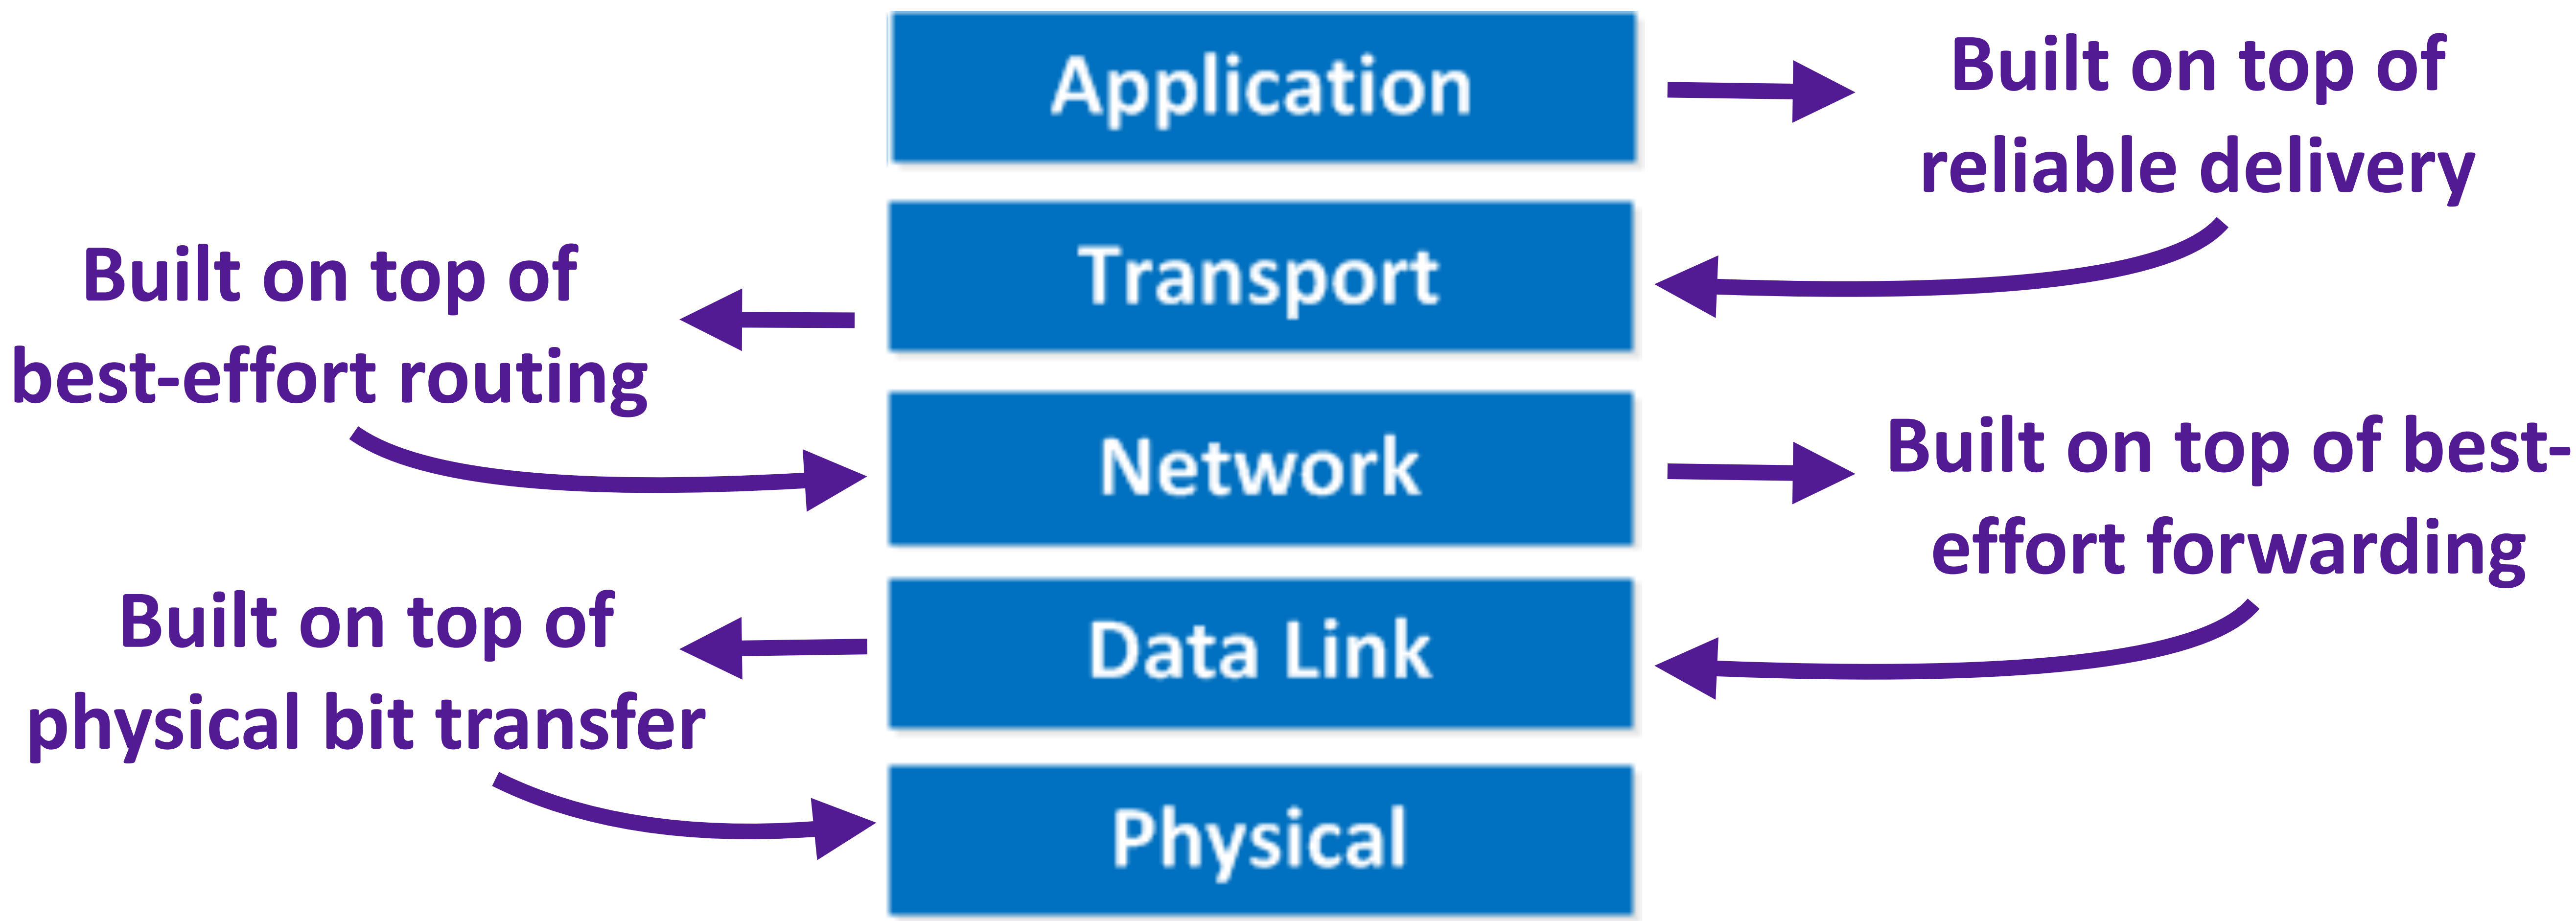

- **A kind of modularity**
  - Functionality separated into layers
  - Layer n **interfaces with only layer n-1 and layer n+1**
    - Hides complexity of surrounding layers

An end-to-end view of the layers

- Application: Providing network support for apps
- **Transport (L4):** (Reliable) end-to-end delivery
- **Network (L3):** Routing and forwarding across networks
- **Datalink (L2):** Forwarding within a local network
- **Physical (L1):** Bits on wire

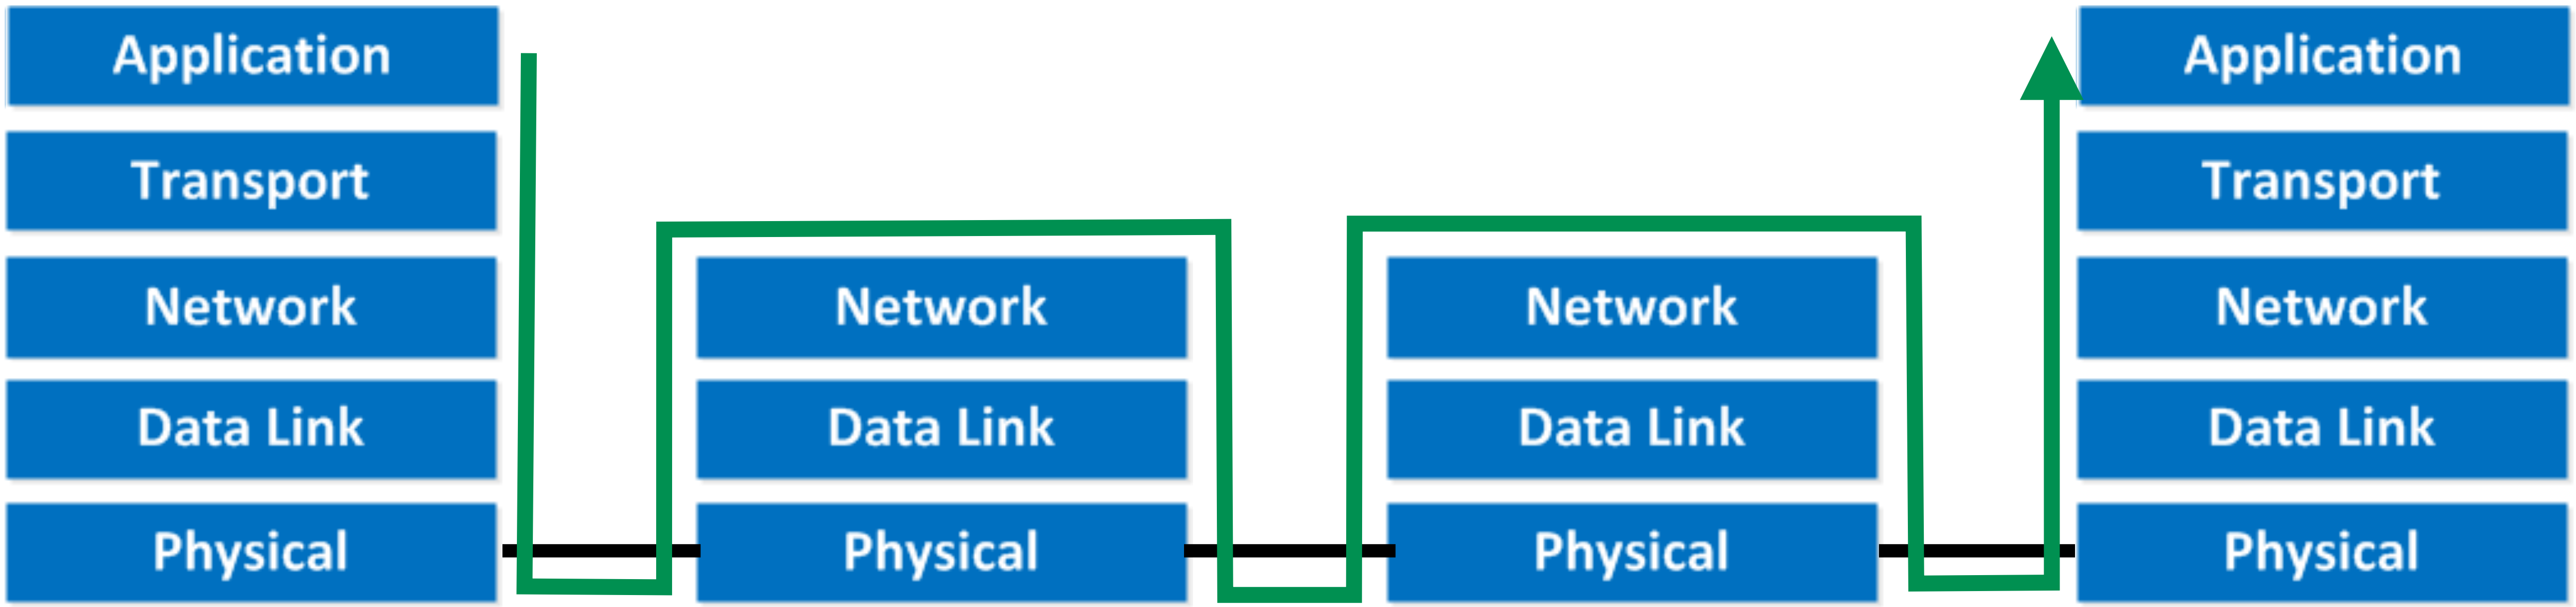

Why does the packet go all the way to network layer at each hop?

Questions?

Three Internet Design Principles

- How to break system into modules?
  - Layering

- Where are modules implemented?
  - **End-to-End Principle**

- Where is state stored?
  - Fate-Sharing

Distributing Layers across Network

- Layers are simple if only on a single machine
    - Just stack of modules interacting with those above/below
- But we need to implement layers across machines
    - Hosts
    - Routers/switches
- What gets implemented where? And why?

What gets implemented on Host?

- Bits arrive on wire, must make it up to application

- Therefore, all layers must exist at host!

What gets implemented on Router?

- Bits arrive on wire
    - Physical layer necessary
- Packets must be forwarded to next router/switch
    - Datalink layer necessary
- Routers participate in global delivery
    - Network layer necessary
- **Routers do not support reliable delivery**
    - Transport layer (and above) ***not*** supported
    - Why?

Visualizing what gets implemented where

- Lower three layers implemented everywhere
- Top two layers only implemented at hosts

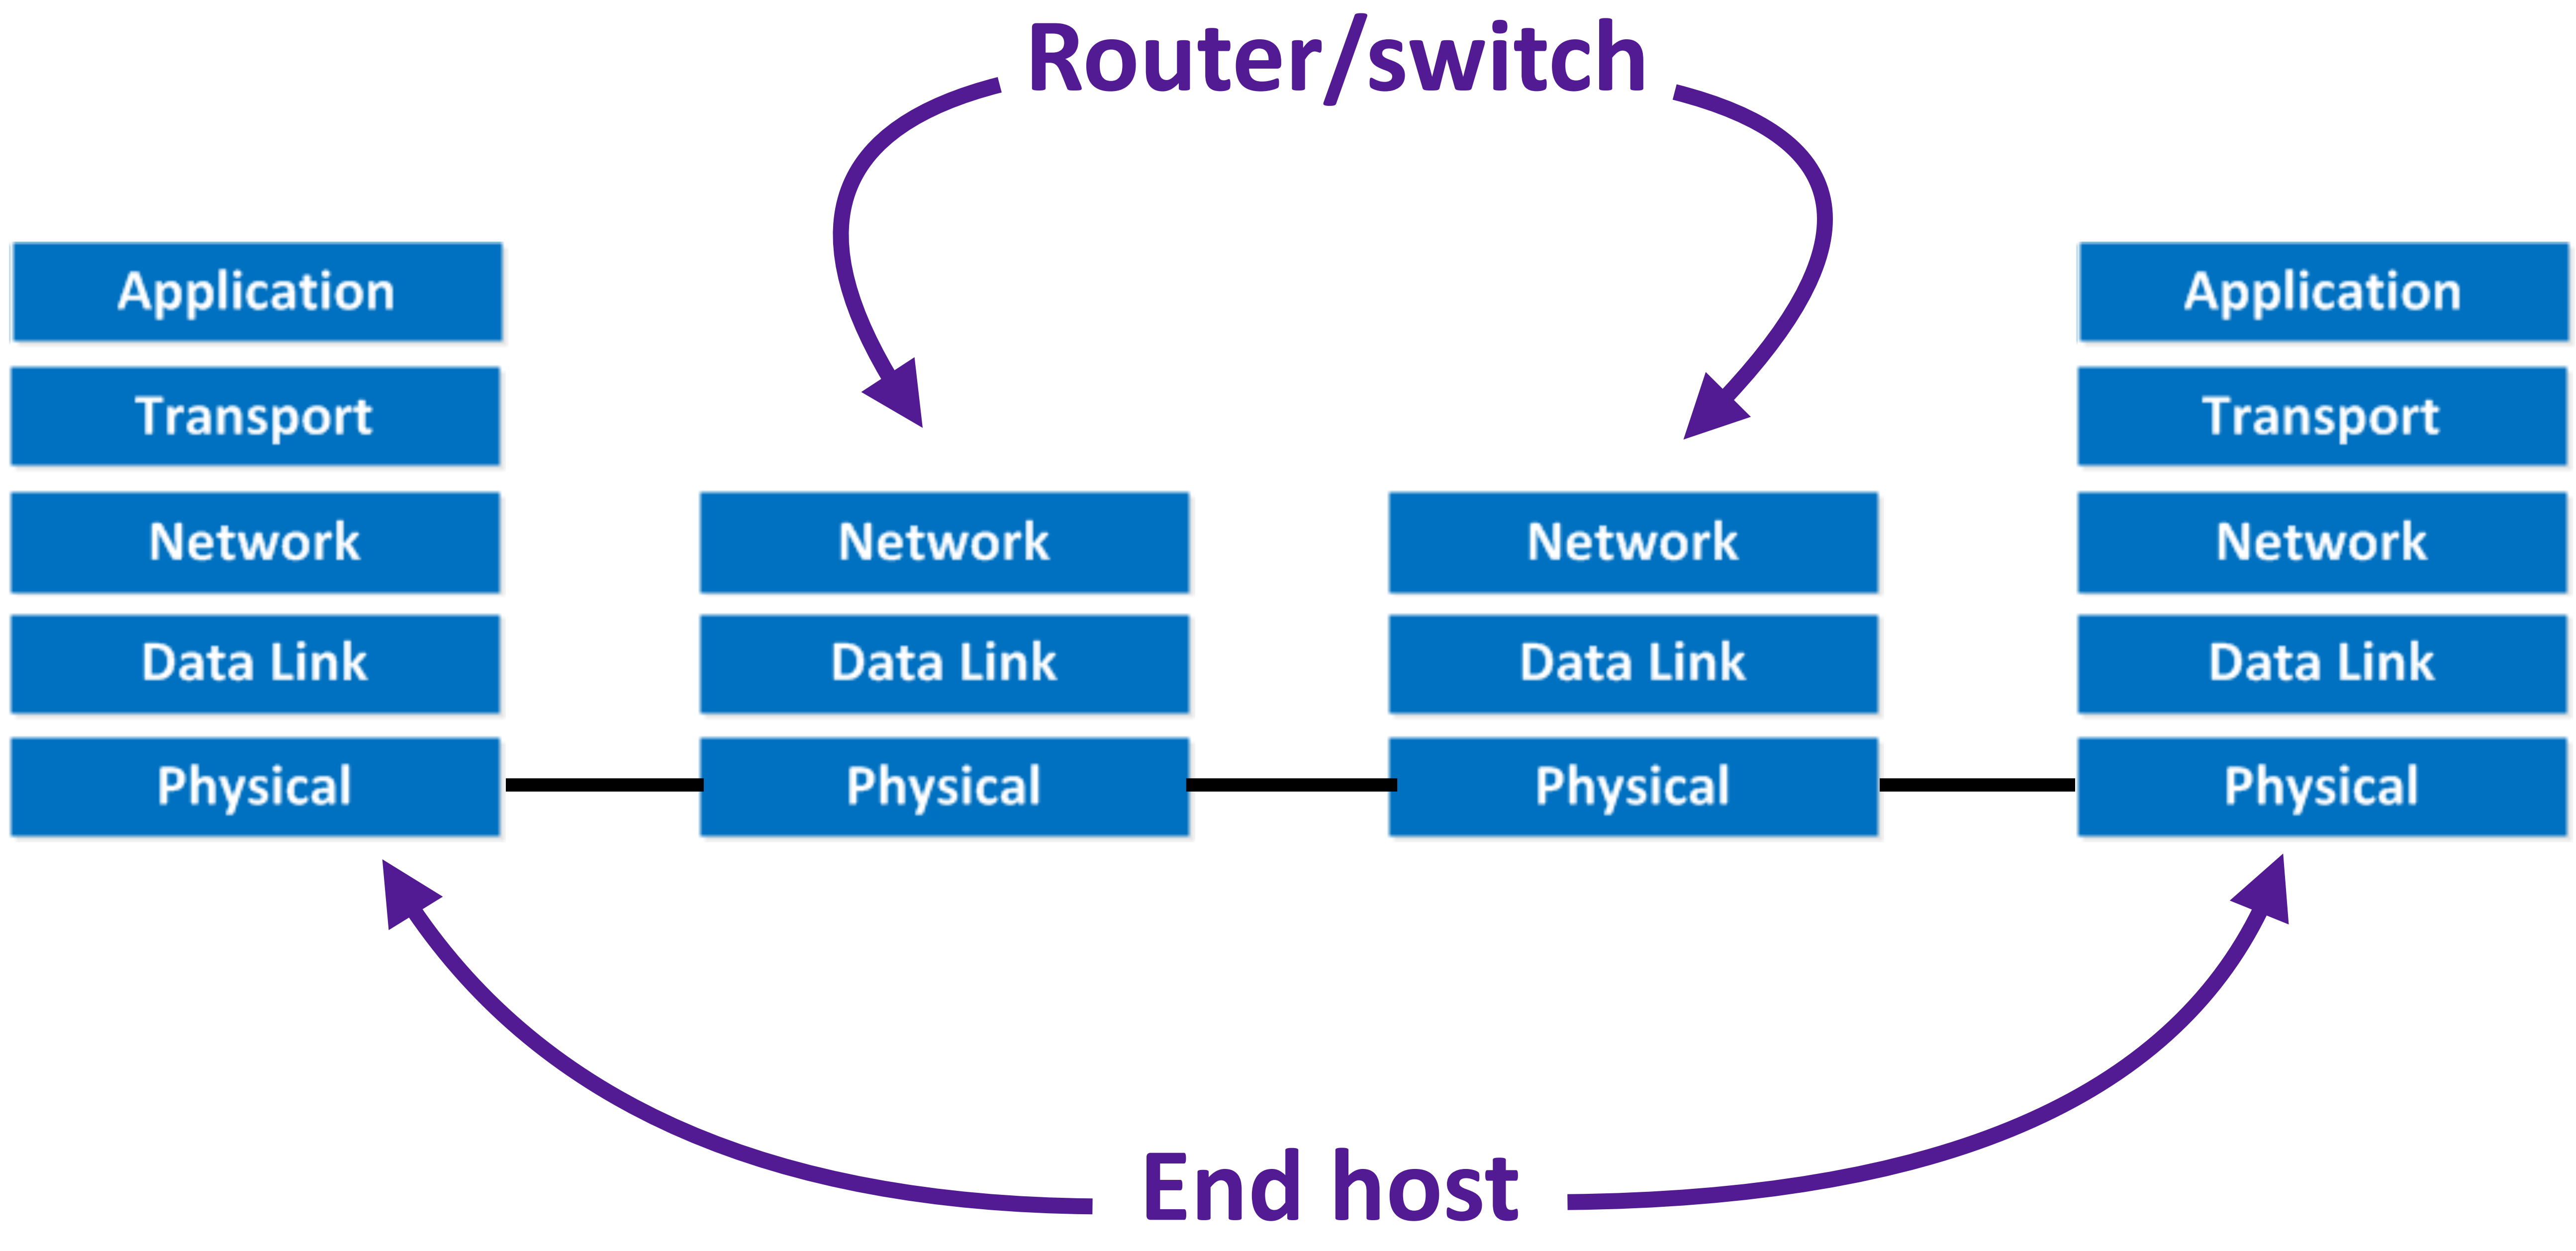

But why implemented this way?

- Layering doesn't tell you **what services each layer should provide**
- What is an effective division of responsibility between various layers?

End-to-end Principle

If a function can completely and correctly be implemented only with the knowledge and help of the application standing at the endpoints of the communication system,

then providing that function as a feature of the communication system itself is not possible.

Sometimes providing an incomplete version of that function as a feature of the communication system itself may be useful as a performance enhancement.

End-to-end Principle: an example

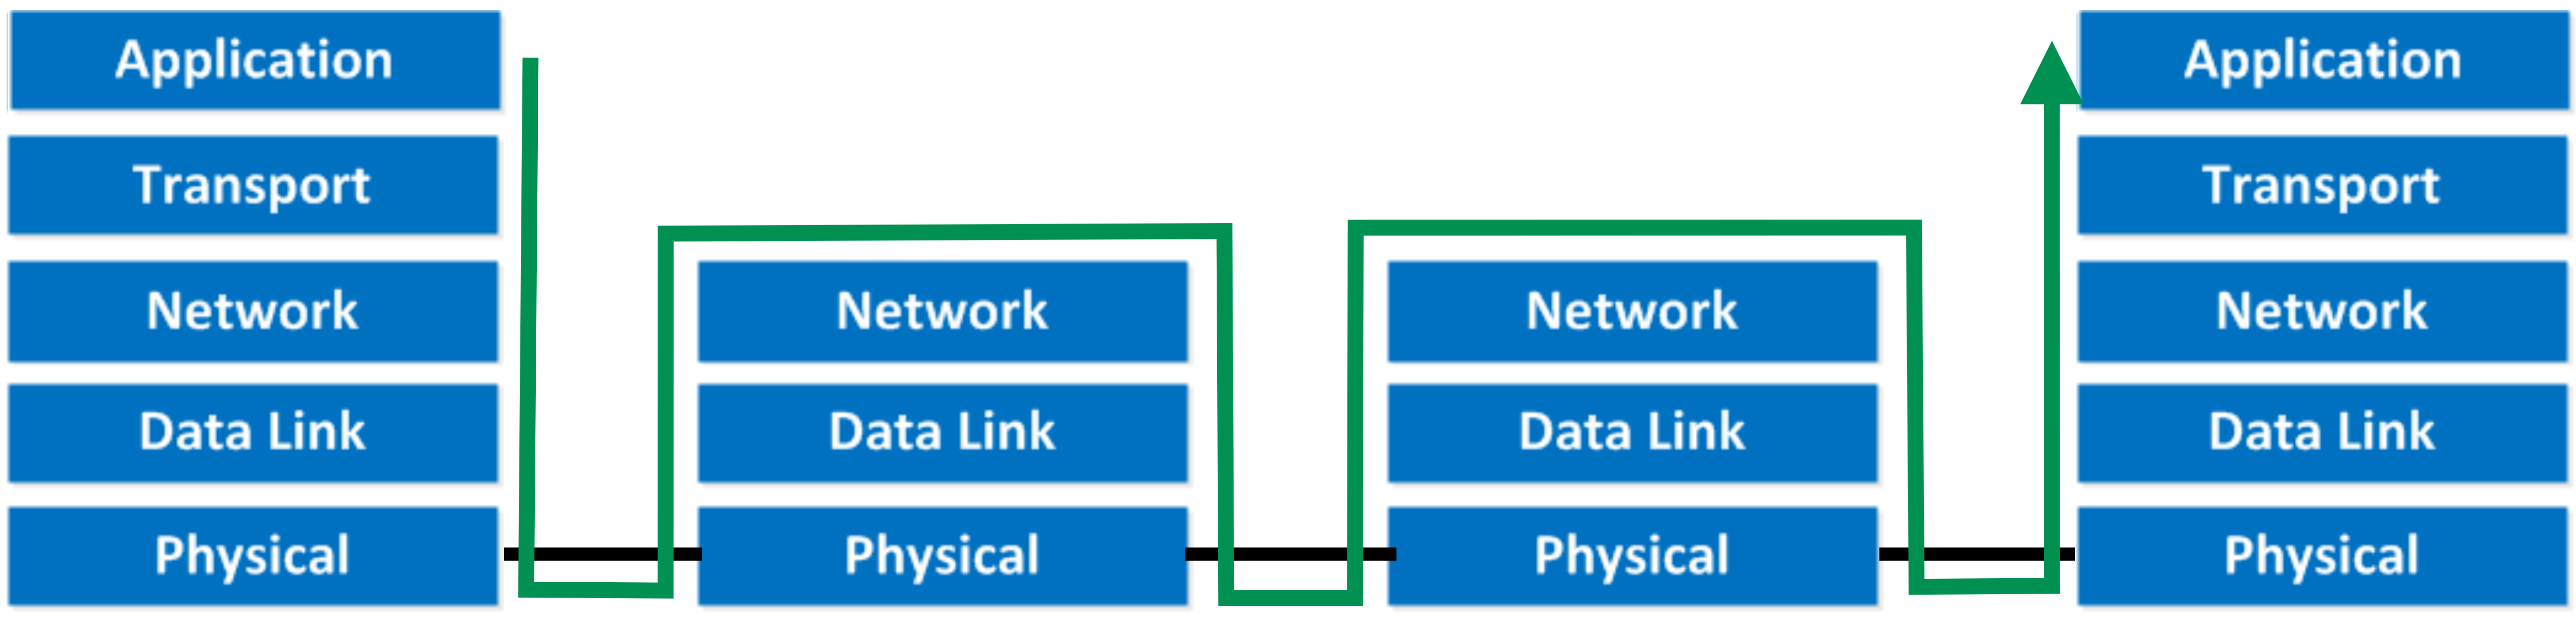

- **Suppose each link layer transmission is reliable**
    - Does that ensure end-to-end (application-to-application) reliability?
- **Suppose network layer is reliable**
    - Does that ensure end-to-end (application-to-application) reliability?

End-to-end Principle: lets read again

If a function can completely and correctly be implemented only with the knowledge and help of the application standing at the endpoints of the communication system,

then providing that function as a feature of the communication system itself is not possible.

Sometimes providing an incomplete version of that function as a feature of the communication system itself may be useful as a performance enhancement.

End-to-end Principle (Interpretation)

Assume the condition (IF) holds. Then,

- **End-to-end implementation**
  - Correct
  - Generalized, and simplifies lower layers
- **In-network implementation**
  - Insufficient
  - May help — or hurt — performance

End-to-end Principle (Interpretation)

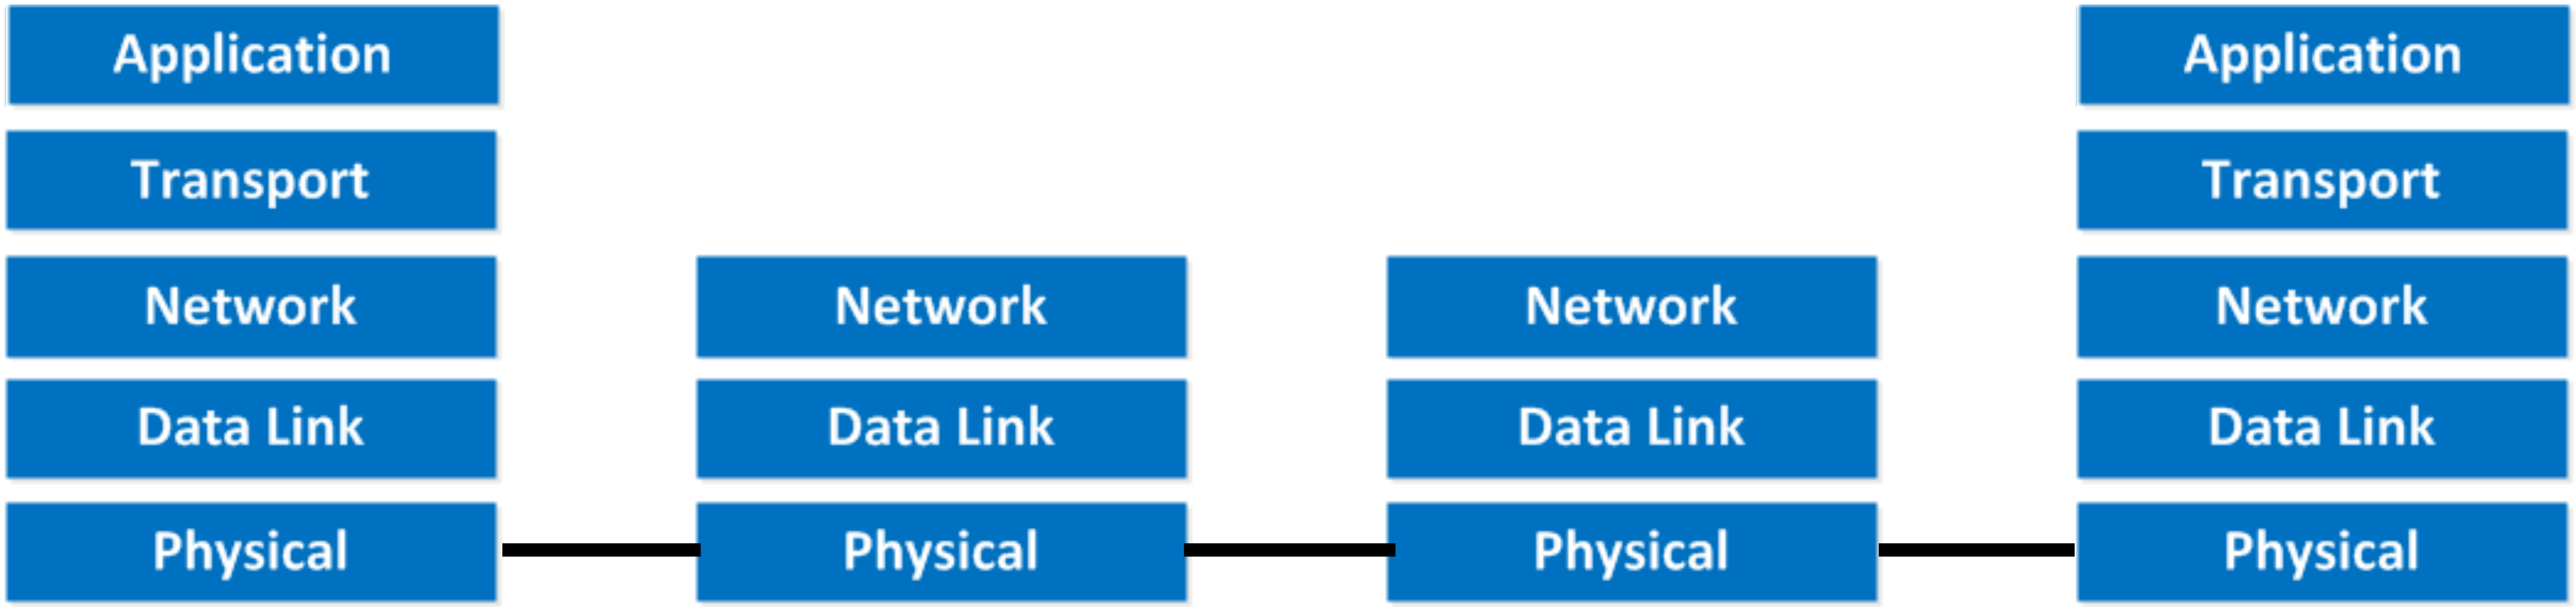

What does the end mean?

End-to-end Principle (Three things to know)

- **Everyone knows what it is**
    - So, you must!
- **Everyone believes it**
    - So, you must!
- **Nobody knows what it means**
    - So, it is okay if you feel so too.

Questions?

Three Internet Design Principles

- How to break system into modules?
  - Layering

- Where are modules implemented?
  - End-to-End Principle

- **Where is the state stored?**
  - **Fate-sharing**

Fate-Sharing

- **Note that the end-to-end principle relied on "fate-sharing"**
  - Invariants only break when endpoints themselves break
  - Minimize the dependence on other network elements

- This should dictate placement of state

- What does **state** mean?
  - Network stacks store **state**:
    - socket/port related
    - (If reliability implemented at hosts) Reliability related
    - …
  - Network routers store **state**:
    - Routing tables (across networks)
    - Forwarding state (within a local network)
    - ….

General Principle: Fate-Sharing

- When storing state in a distributed system, colocate it with entities that rely on that state

- Only way failure can cause loss of the critical state is if the entity that cares about it also fails …
  - … in which case it doesn’t matter

- Often argues for keeping network state at end hosts rather than inside routers
  - E.g., packet switching rather than circuit switching

Questions?

Decisions and their Principles

- How to break system into modules
  - **Dictated by layering**
- Where modules are implemented
  - **Dictated by End-to-End Principle**
- Where state is stored
  - **Dictated by Fate Sharing**

From Architecture to Design:

Design Goals

David Clark

- Wrote a paper in 1988 that tried to capture why the Internet turned out as it did
- It described an ordered list of priorities that informed the decision
- What do you think those priorities were?

Internet Design Goals (Clark '88)

- **Connect existing networks**
- Robust in face of failures
- Support multiple types of delivery services
- Accommodate a variety of networks
- Allow distributed management
- Easy host attachment
- Cost effective
- Allow resource accountability

#1: Connect Existing Networks

Want one protocol that could be used to connect any pair of (existing) networks

- Different networks may have different needs
    - For some: reliable delivery more important
    - For others: performance more important
    - **But there is one need that every network has: connectivity**
- The Internet Protocol (IP) is that unifying protocol
    - All (existing) networks must be able to implement it

#2: Robust in Face of Failures

As long as network is not partitioned, two hosts should be able to communicate (eventually)

- Must **eventually recover** from failures

- Very successful in the past; unclear how relevant now
    - **Availability** is becoming increasingly important than **recovery**

#3: Support Multiple Types of Delivery Services

Different delivery services (applications) should be able to co-exist

- Already implies an application-neutral framework
- Build lowest common denominator service
    - **Again: connectivity**
    - Applications that need reliability may use it
    - Applications that do not need reliability can ignore it
- **This isn’t as obvious as it seems…**
    - What would applications in 2050 need?

Questions?

#4: Variety of Networks

Must be able to support different networks with different hardware

- **Incredibly successful!**
  - Minimal requirements on networks
  - No need for reliability, in-order, fixed size packets, etc.
  - A result of aiming for lowest common denominator

- **Again: Focus on connectivity**
  - Let networks do specific implementations for other functionalities
  - Automatically adapt: WiFi, LTE, 3G, 4G, 5G ….

#5: Decentralized Management

No need to have a single “vantage” point to manage networks

- Both a curse and a blessing
    - Important for easy deployment
    - Makes management hard today

- Recent efforts have improved management of individual networks
    - But no attempt to manage the Internet as a whole…
    - What might make this complex?

#6: Easy Host Attachment

The mechanism that allows hosts to attach to networks must be made as easy as possible, but no easier

- Clark observes that cost of host attachment may be higher because hosts had to be smart

- But the administrative cost of adding hosts is very low, which is probably more important
    - Plug-and-play kind of behavior…

- And now most hosts are smart for other reasons
    - So the cost is actually minimal…

#7: Cost Effective

Make networks as cheap as possible, but no cheaper

- Cheaper than circuit switching at low end
- More expensive than circuit switching at high end
- Not a bad compromise:
    - Cheap where it counts (low-end)
    - More expensive for those who can pay…

#8: Resource Accountability

Each network element must be made accountable for its resource usage

- Failure!

Internet Motto

“We reject kings, presidents and voting. We believe in rough consensus and running code.”

- - David Clark

Real Goals

- **Build something that works**
- Connect existing networks
- Robust in face of failures
- Support multiple types of delivery service
- Accommodate a variety of networks
- Allow distributed management
- Easy host attachment
- Cost effective
- Allow resource accountability

Questions to think about

- What goals are missing from this list?
    - **Suggestions?**

- What would the resulting design look like?

Some of the missing issues

- **Performance**
- Security
    - Resilience to attacks (denial-of-service)
    - Endpoint security
    - Tracking down misbehaving users
- Privacy
- Availability
- Resource sharing (fairness, etc.)
- ISP-level concerns
    - Economic issues of interconnection

Questions?

Next lecture

- Beginning of “Design of computer networks”

- Start with Layer 1 and Layer 2
  - Physical bits (very little)
  - Local best-effort forwarding
  - Lots of interesting aspects
  - Lots of group activities
  - …