15-441/641: Computer Networks
The Internet Model & Layering

15-441 Spring 2019

Profs Peter Steenkiste & **Justine Sherry**

Carnegie Mellon University

I have this problem

- I’m very bad at estimating time for lectures.
  - Maybe if I had been teaching this class since 1987 I’d know better
- I definitely though I was going to complete all that routing stuff last Thursday.
- Instead we finished it today (1/21).
- Here’s what I’d originally planned to discuss today — let’s see how far we make today :-)

What you know so far

- A network consists of nodes and links.
- Networks can be implemented using many mediums: fiber/light, copper/electricity, air/radio waves, string/knots…
- Data is transmitted in fixed-sized chunks called packets.
  - Packet headers (like Ethernet headers) “wrap” these packets with useful information, like the source and destination for the packet.
- You can calculate how long it takes for a packet to arrive at its destination using transmission delay and propagation delay.
- You know how a collection of nodes in a network routes a packet to its destination.

This all seems so basic!

- How could we possibly build a web service that scales the globe using such simple primitives?

Network Architecture

Network architecture is the design of a computer network. It is a framework for the specification of a network's physical components and their functional organization and configuration, its operational principles and procedures, as well as communication protocols used.

[Wikipedia]

Architectures

- The architecture of your house specifies its structure, from nails and boards up, through descriptions like “walls”, up to big ideas like “dining room”. It describes how all of these components fit together to make up your house.

- The architecture of your computer specifies its structure, starting from how transistors are connected, to create memory and circuits, all the way up to an ISA, BIOS and starting up an “operating system”.

- The architecture of the Internet similarly specifies how we go from electrons, light, flags or radio signals up to things like “Netflix streaming video service” or “Google search.”

The Internet Model

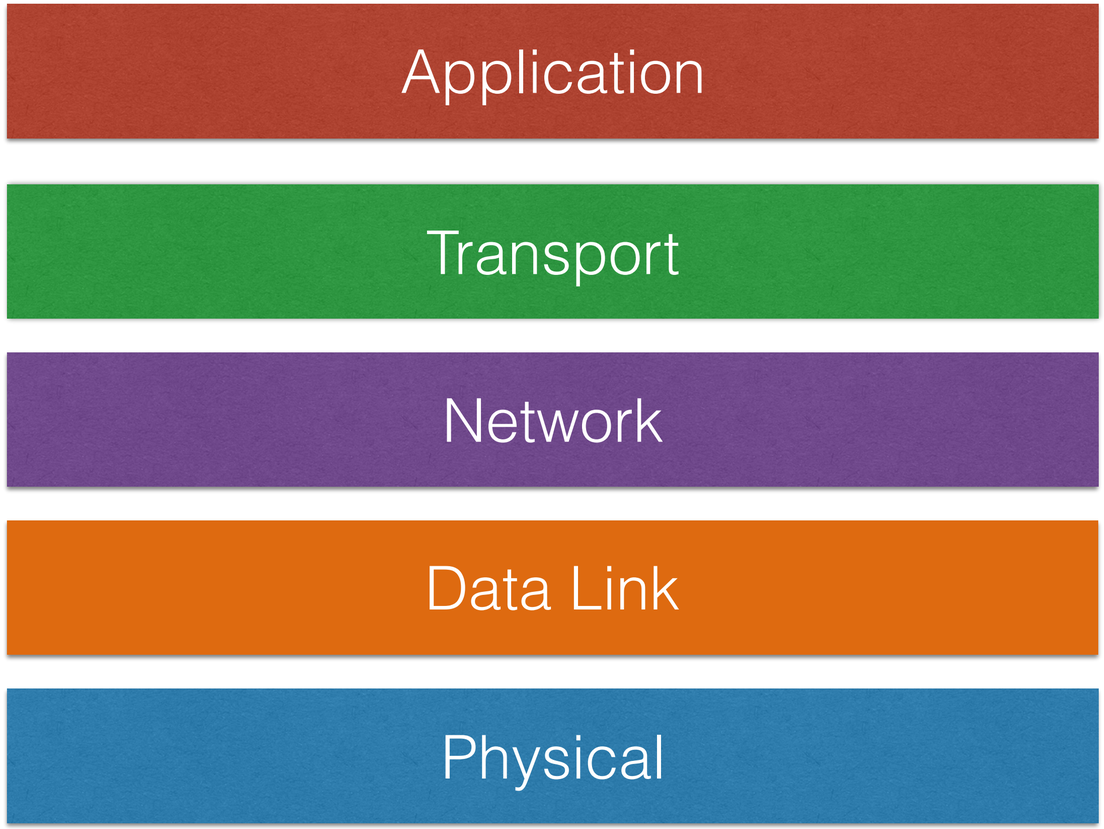

This video for no good reason…

… and the Internet has layers too!

Understanding the Layers of the Internet Model

Bottom Up:
Think of the
layers like APIs

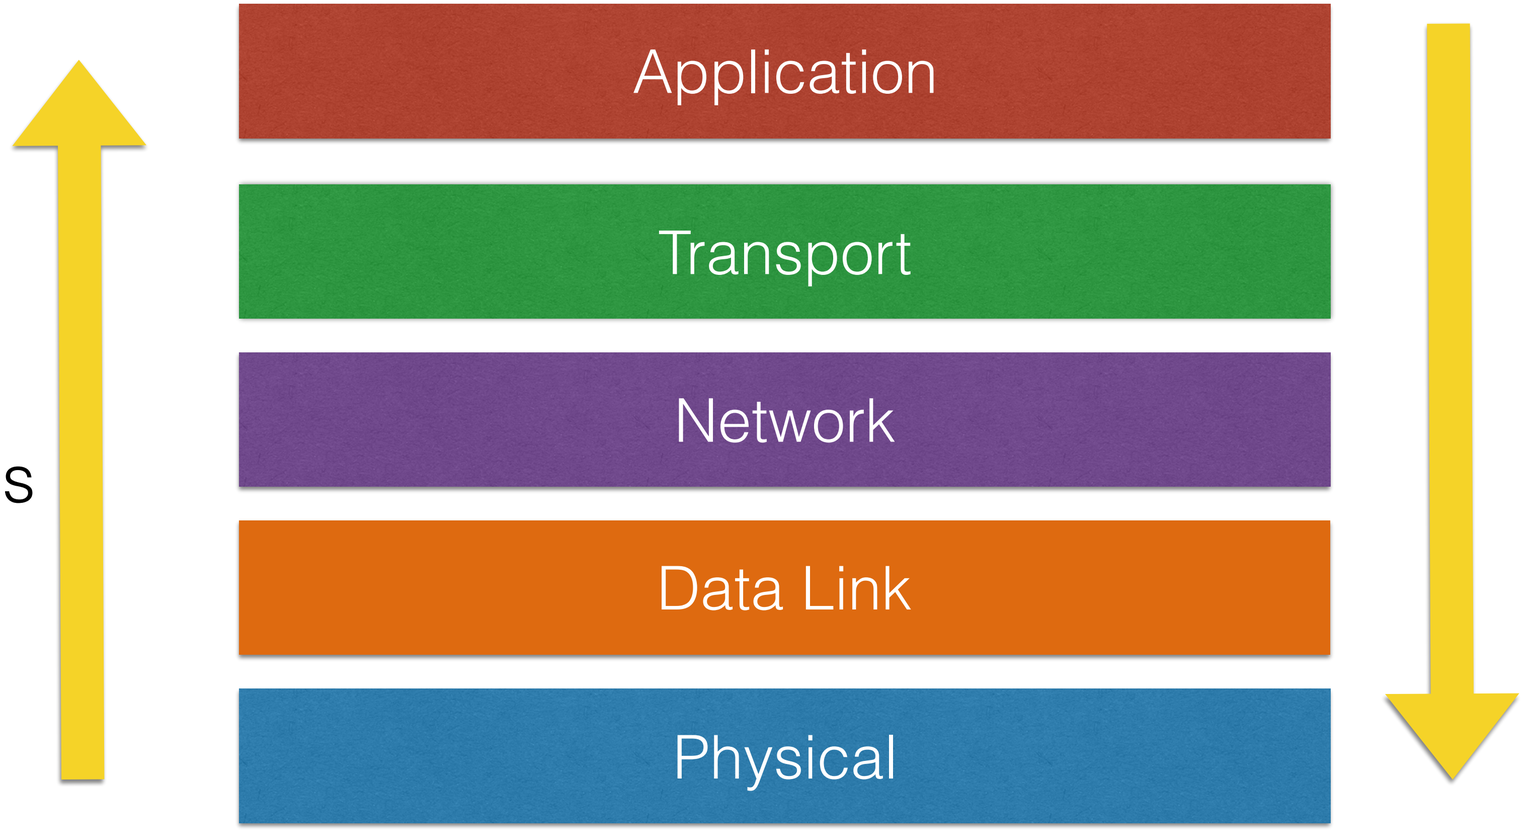

Top Down:
Encapsulation
Model

Understanding the Layers of the Internet Model

Bottom Up:
Think of the
layers like APIs

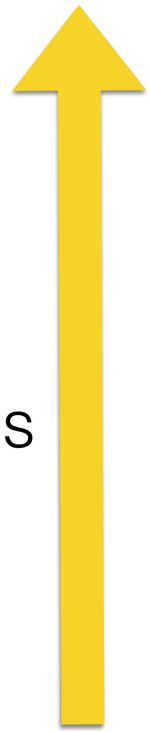

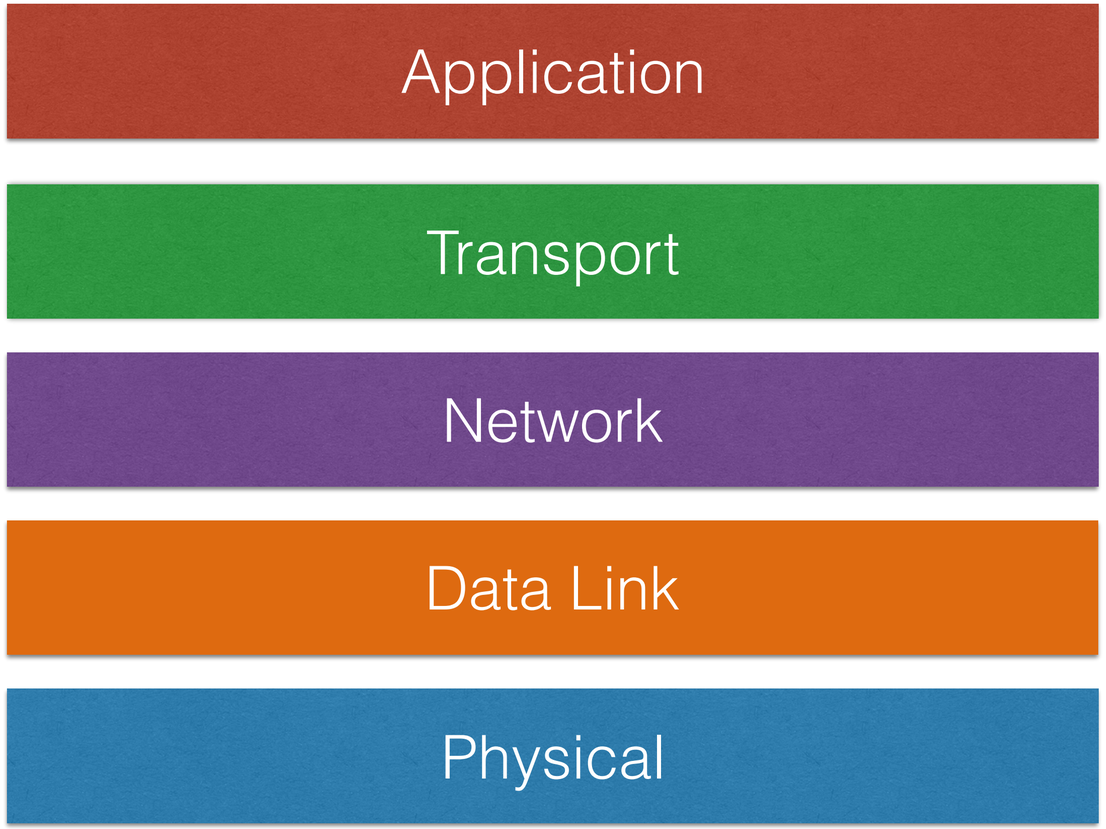

PHY: The Physical Layer (Layer 0)

- “The physical layer defines the means of transmitting raw bits... The bitstream may be grouped into code words or symbols and converted to a physical signal that is transmitted over a transmission medium." [Wikipedia]
- Takes as input: electricity, light, radio waves… any physical medium
- Provides an API to the “next up” layer that provides 1’s and 0’s

PHY: The Physical Layer (Layer 1)

0101010100100000010011111010101

Anyone *using* the PHY layer can just think about 1's and 0's — they don't have to worry about how the data is transmitted.

Physical Layer

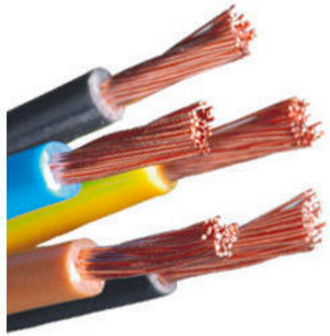

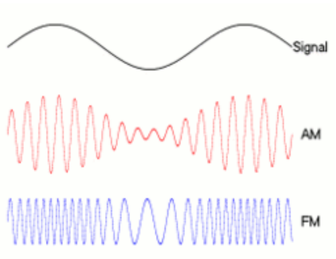

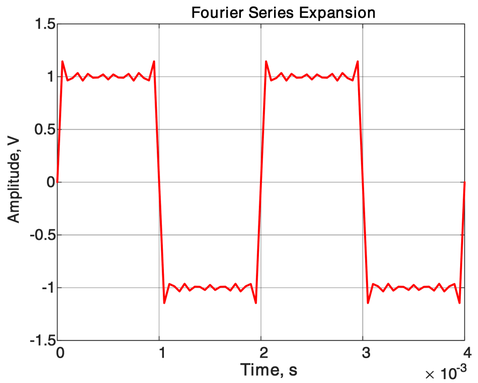

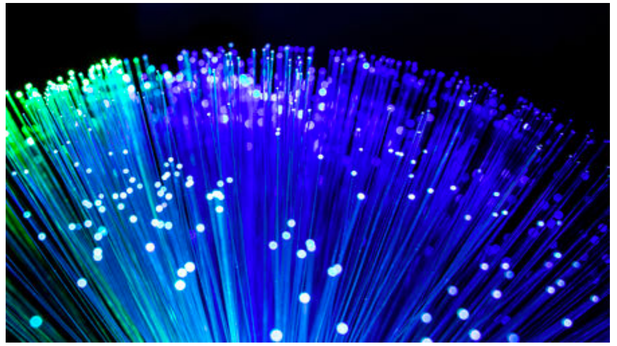

Systems Engineering Wisdom

“Modularity based on abstraction is the way things get done.”

Barbara Liskov
Turing Award Winner
+ von Neumann Medal, Computer Pioneer Award,…
+ Pretty much all the things.

LNK: The Data Link Layer (Layer 2)

0101010100100000010011111010101

Physical Layer

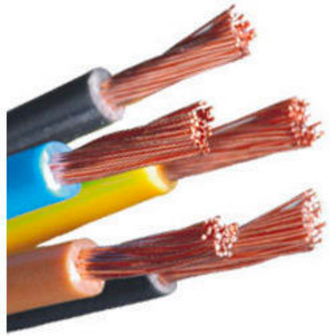

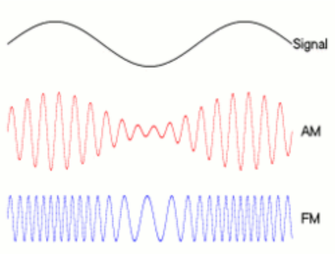

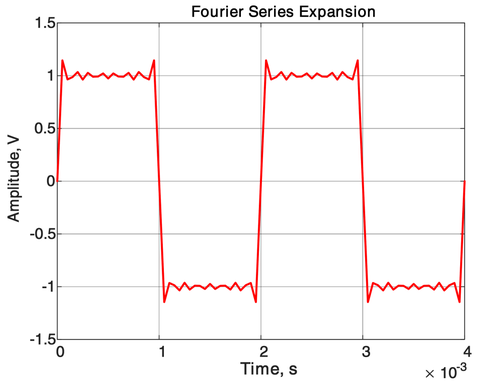

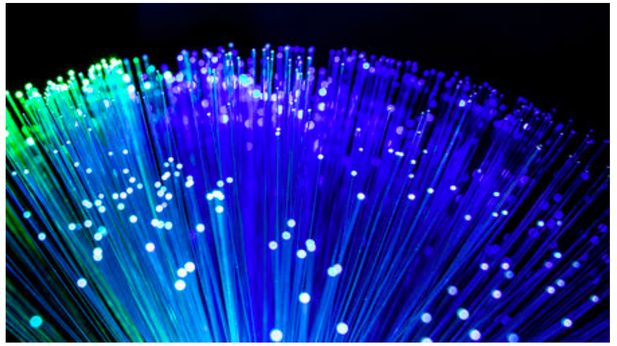

LNK: The Data Link Layer (Layer 2)

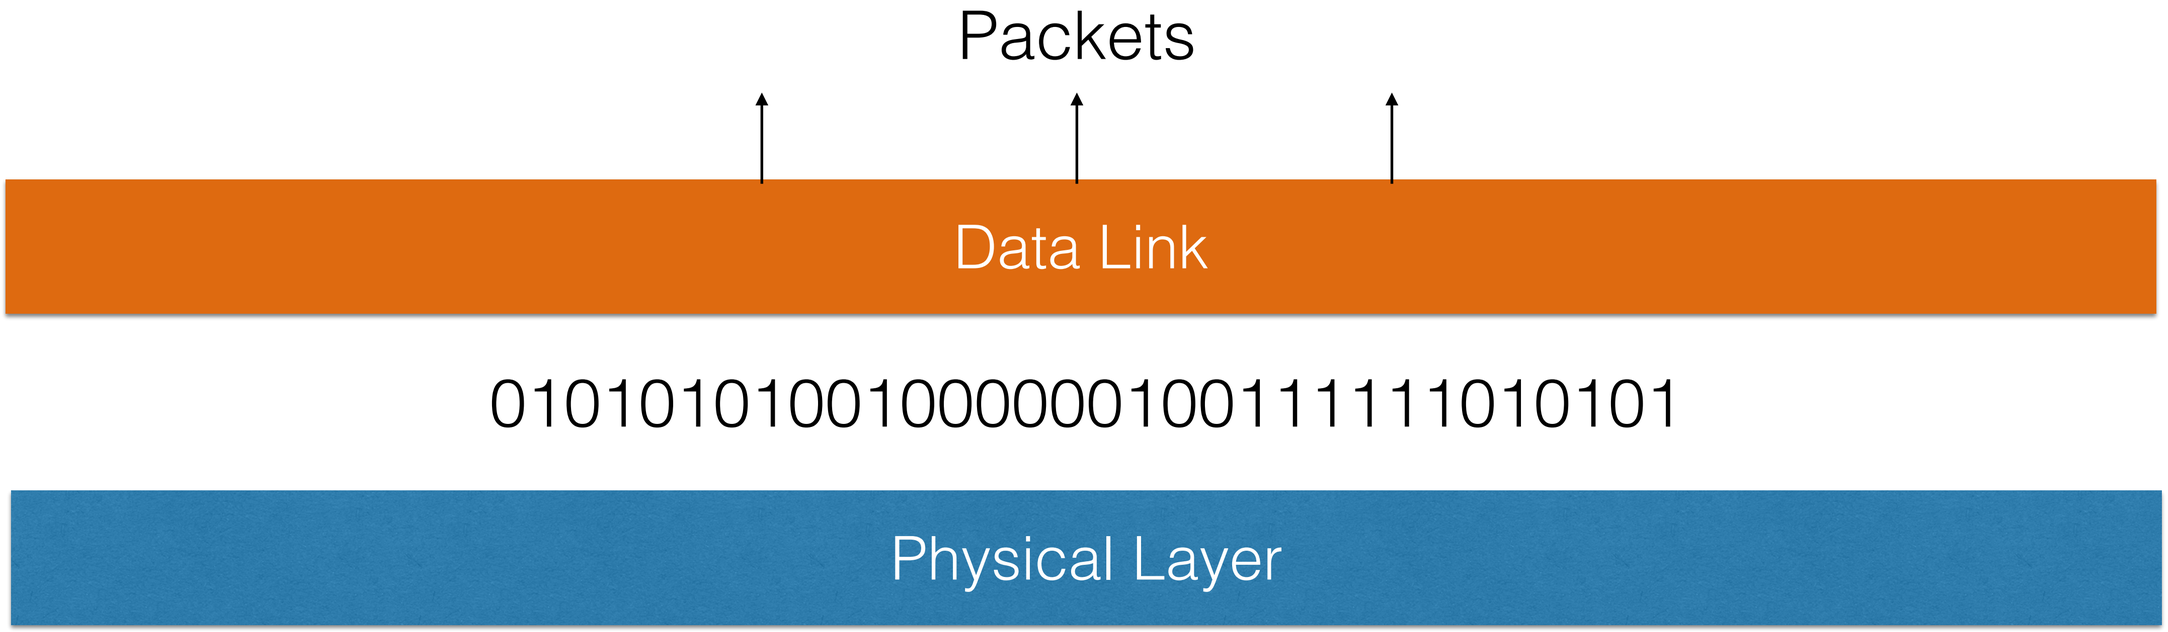

There’s a lot going on inside LNK

- **Ethernet is an example of a data link-layer protocol. Things you already know:**
  - Turns those 1’s and 0’s into packets.
  - Labels senders and receivers.
- **Other things the link layer does — we’ll talk more about this later in the semester.**
  - Allows everyone to talk to everyone else! It’s like we’re all on the same link.
  - Prevents multiple senders from transmitting at the same time
    - Think about wireless: if two transmitters transmit at the same time, they interfere with each other and the signal doesn’t get through.
  - Detects and prevents errors in packets — what if the bits are corrupted?
    - Solar flares, power surge, someone turned on a microwave…

There are a lot of link-layer protocols too

- The stuff you see in practice will almost alwa… …e version of Ethernet we learn in class.

Once again, important thing is that the layer above just sees nice, uncorrupted packets! Doesn't have to care how this is achieved.

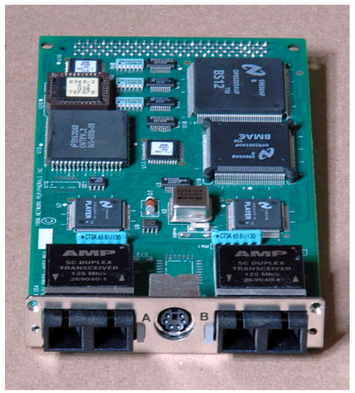

FDDI

Token Ring

PPP

Bluetooth®

L2CAP

LNK: The Data Link Layer (Layer 2)

Packets

Data Link

0101010100100000010011111110101010

Physical Layer

NET: The Network Layer (Layer 3)

Network / INTERNET

Packets

Data Link

0101010100100000010011111110101 01

Physical Layer


#FALLONTONIGHT
WELCOME, WELCOME, WELCOME,

The network layer is funny

- The structure of the data changes very little
  - You still get packets.
  - They just have different headers than they did at the Link Layer
- There aren’t multiple network layers.
  - There’s only one.
  - The Internet Protocol

NET: The Network Layer (Layer 3)

More Packets?

Network / INTERNET

Packets

Data Link

0101010100100000010011111110101 01

Physical Layer

NET vs LNK

- Both provide packet delivery
  - You know what packets are
    - They’re not packages
- So what is the difference between them?

History Time

How do we connect these very different networks?

Network 1
Ethernet Network
48-bit Addresses
1500 byte frames
1 Mbps

Network 2
Token Ring Network
16 bit addresses
17,000 byte frames
4 Mbps

Network 1
Ethernet Network
48-bit Addresses
1500 byte frames
100 Kbps

Answer: The Internet Protocol

- You frame your packets however you like.
  - Your addresses, your packet sizes, your headers…
- Inside of your packets, we will stick another header!
- This new header will contain global addresses: IP Addresses
- We'll divide the world into *switches* and *routers*
  - Switches will route on your, Link Layer (L2) Addresses
  - Routers will operate on IP (L3) Addresses

Network 1
Ethernet Network
48-bit Addresses
1500 byte frames
1 Mbps

Network 2
Token Ring Network
16 bit addresses
17,000 byte frames
4 Mbps

Network 1
Ethernet Network
48-bit Addresses
1500 byte frames
100 Kbps

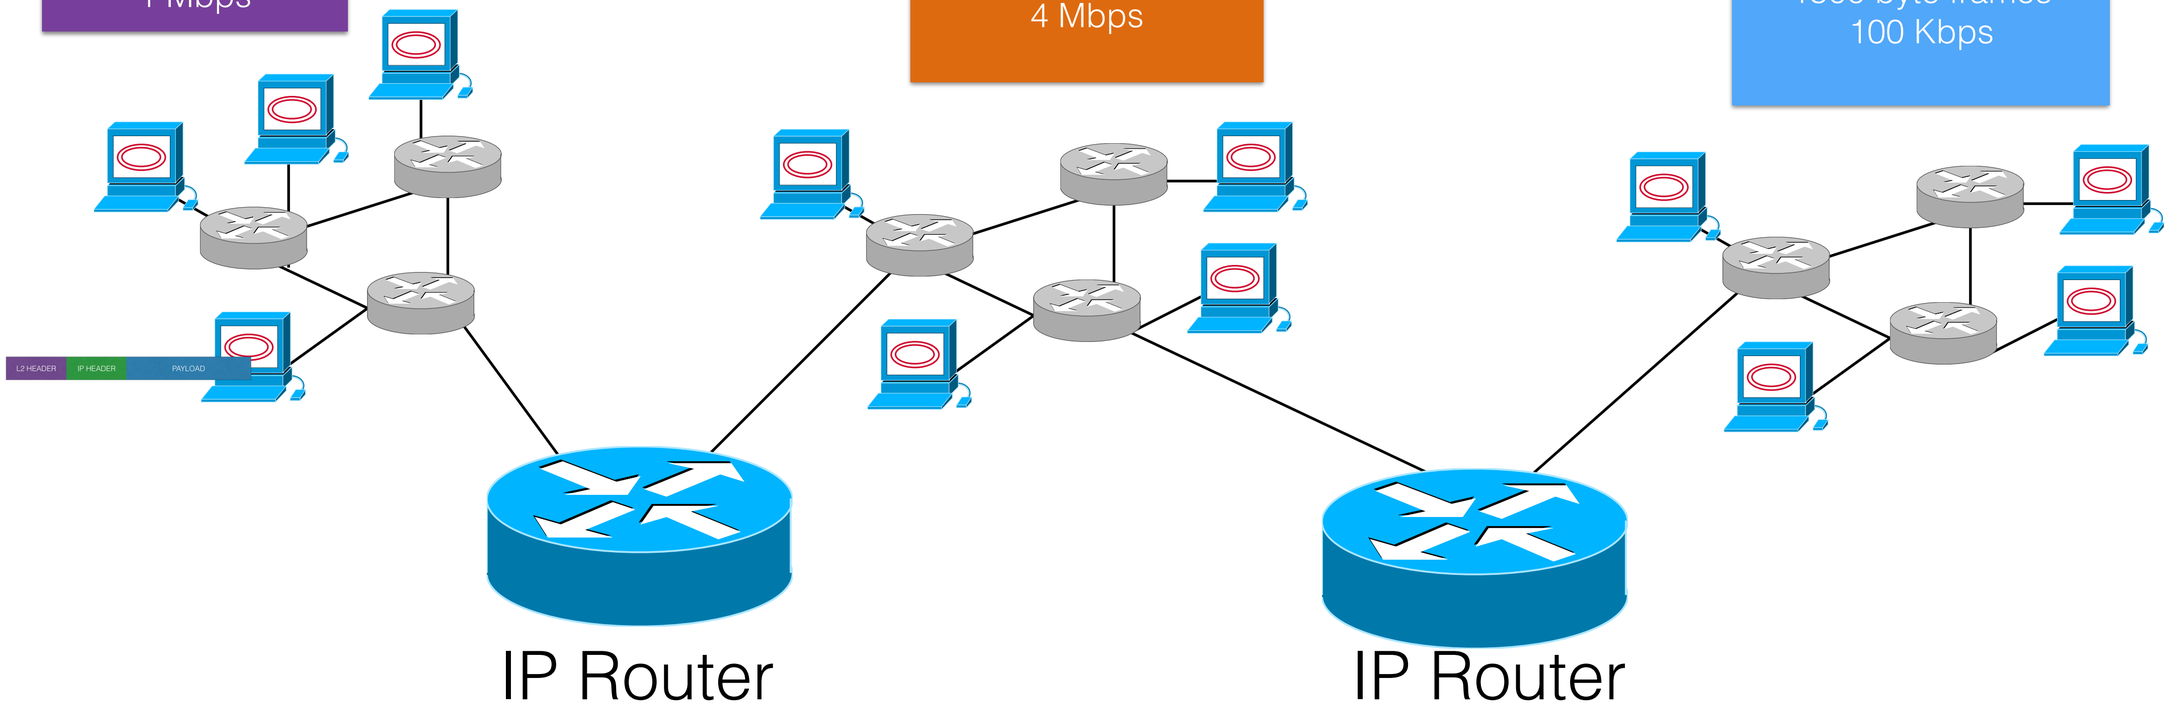

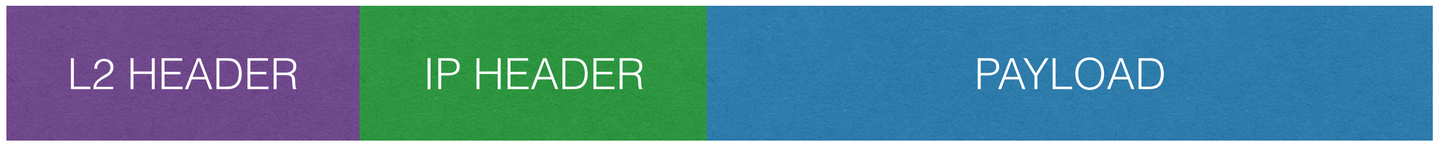

Network 1
Ethernet Network
48-bit Addresses
1500 byte frames
1 Mbps

Network 2
Token Ring Network
16 bit addresses
17,000 byte frames
4 Mbps

Network 1
Ethernet Network
48-bit Addresses
1500 byte frames
100 Kbps

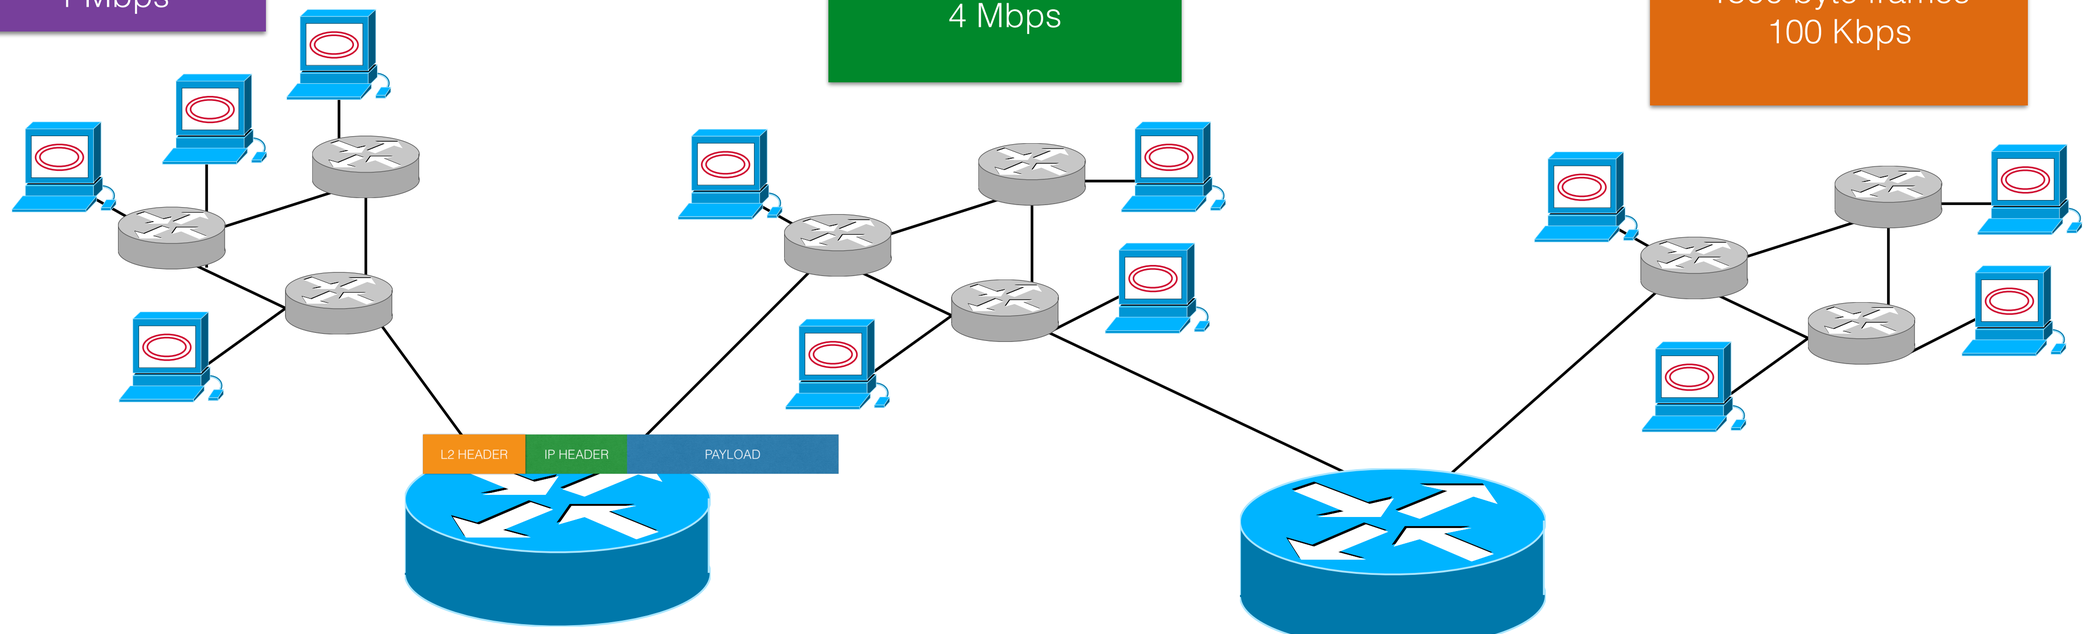

IP Router

IP Router

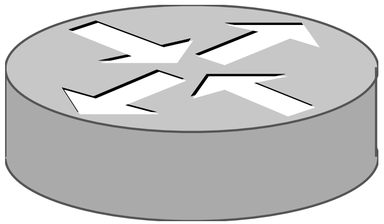

Switches look at the Layer-2 Header and route according to the Layer-2 Address (e.g, MAC addresses)

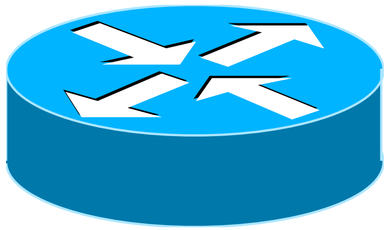

Routers strip off the Layer-2 Header and look at the Layer 3 Header. They route using Layer 3 (i.e. IP) addresses.

The Narrow Waist

- This design means that everyone has to agree to use IP.
  - All of the network have to use it in the same way.
- So there can be lots of L2 designs… but there is only one L3 design.
- This is why IP is often called “The Narrow Waist” of the Internet.

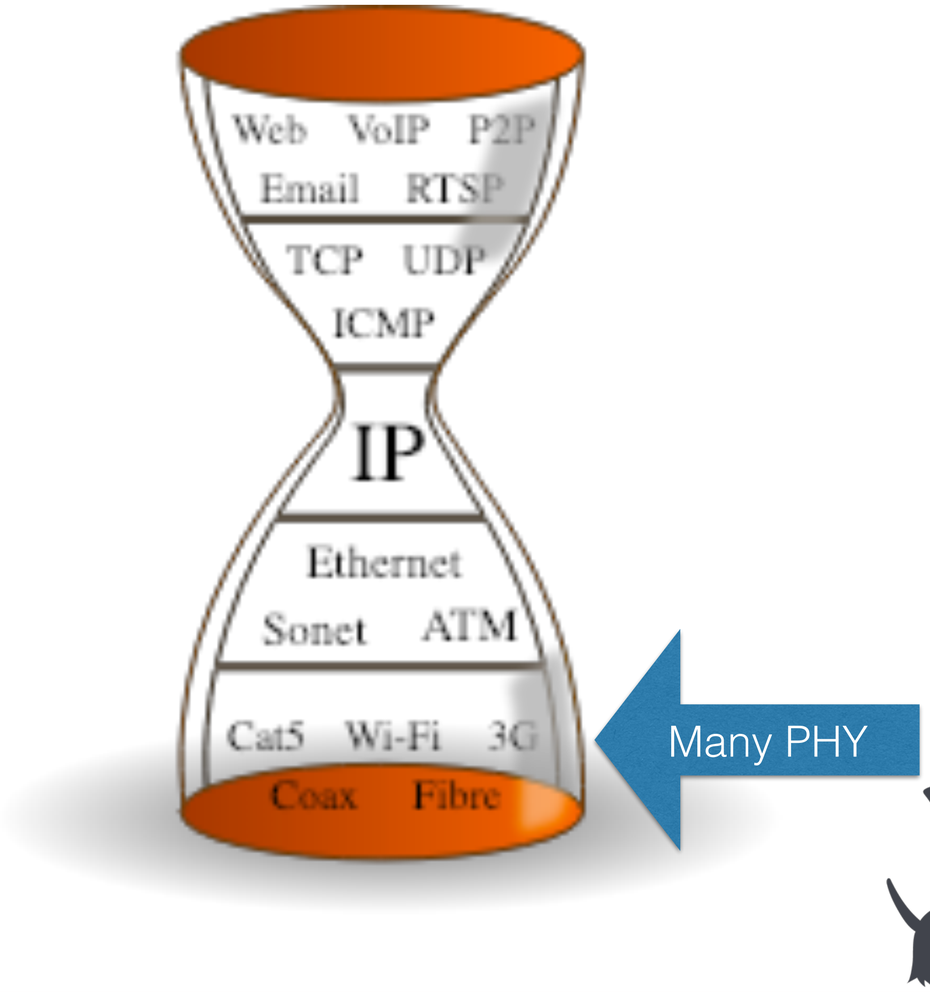

The Narrow Waist

- This design means that everyone has to agree to use IP.
  - All of the network have to use it in the same way.
- So there can be lots of L2 designs… but there is only one L3 design.
- This is why IP is often called “The Narrow Waist” of the Internet.

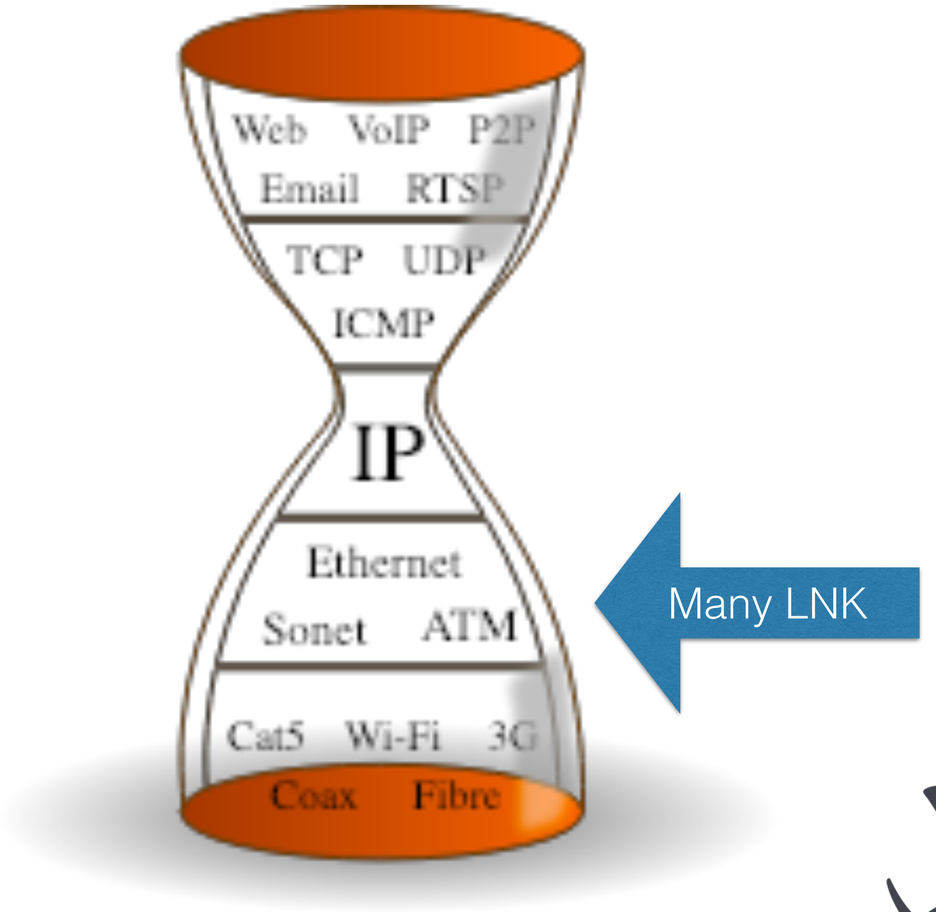

The Narrow Waist

- This design means that everyone has to agree to use IP.
  - All of the network have to use it in the same way.
- So there can be lots of L2 designs… but there is only one L3 design.
- This is why IP is often called “The Narrow Waist” of the Internet.

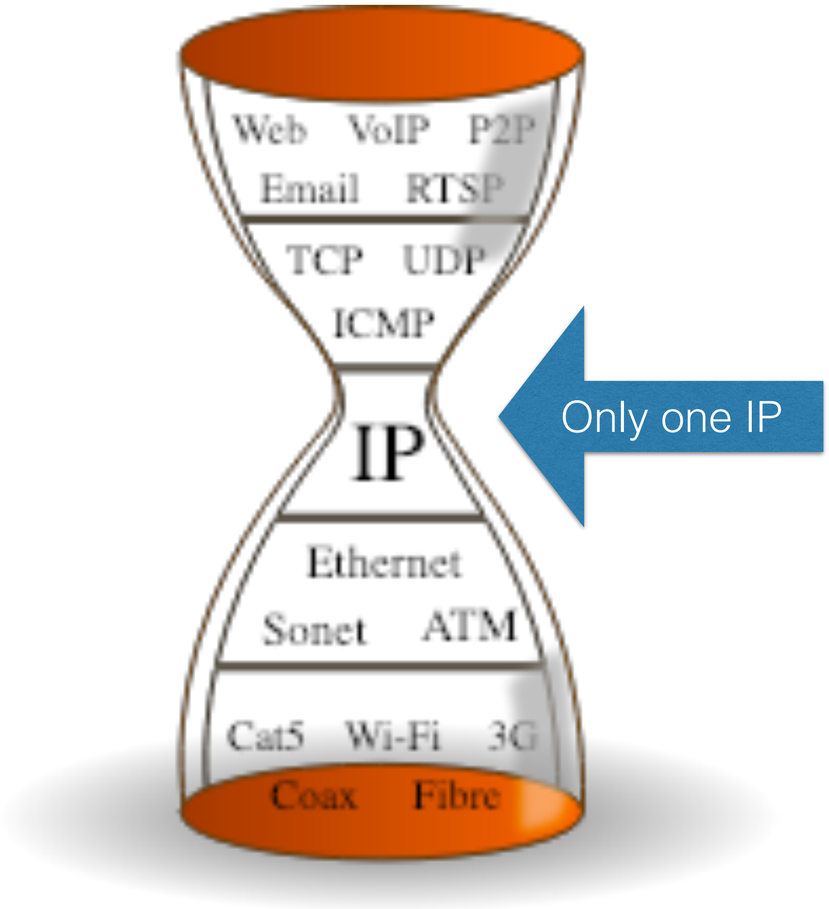

Coping with other differences between networks

- Some networks are higher speed than others. IP will just drop packets if a fast network sends packets too quickly into a slow one.
- Some networks guarantee no loss. IP doesn't even try — since some networks can't guarantee loss-free, in order delivery, the Internet doesn't try to guarantee loss-free, in order delivery.
  - This is called *best effort service.*
- Some networks have bigger packets than others.
  - IP supports "fragmentation" — splitting one packet into many smaller ones — to make up for this.

Professor Steenkiste will tell you lots more about IP shortly!

NET: The Network Layer (Layer 3)

Globally Routable Packets

Network / INTERNET

Packets

Data Link

01010101001000000100111111010101

Physical Layer

The Transport Layer (Layer 4)

Transport

Globally Routable Packets

Network / INTERNET

Packets

Data Link

01010101001000000100111111010101

Physical Layer

The Need for the Transport Layer

- The IP layer doesn't give any guarantees!
  - Packets can arrive out of order.
  - They can be lost.
- The transport layer makes up for this. The transport layer is implemented in code running on end hosts.
  - It <u>can</u> establish a *connection* between two endpoints.
  - It <u>can</u> re-transmit data that is lost.
  - It <u>can</u> put packets back together in the right order.

But it doesn't have to!

Choose your own Transport

TCP:
Reliable
In-Order
“Byte Stream”

UDP:
No guarantees

Why would we want both of these?

The Transport Layer (Layer 4)

You're pretty familiar with sockets and getting data across the network in-order — so I'll leave transport be for now.

Connection

Transport

Globally Routable Packets

Network / INTERNET

Packets

Data Link

0101010100100000010011111110101 01

Physical Layer

The Application Layer (Layer 7)

Application

Connection

Transport

How did we get to 7?

Globally Routable Packe

Network / INTERNET

Packets

Data Link

01010101001000000100111111010101

Physical Layer


California's Handgun Vote · Ray Mancini: a Death in the Ring
RollingStone
MICHAEL JACKSON
Life as a Man
Eurovision Germany 1983
SCARFACE
IDCA
1983
The Future Isn't What It Used To Be
1983
NBC SATURDAY MORNING WE
STARTS SATURDAY SEPT. 17TH
NEW
NEW
SUSAN SARANDON
FOME DE VIVER
THE HUNGER
HOJE
METRO 1
IPIRANGA 2
GEMINI 2
TOP CINE
IBIRAPUERA 2
MORUMBI 2
SPLENDID
SERRADOR 2
INDAIA
RollingStone
PRINCE'S HOT ROCK

1983: The OSI Model

- Imagining the future of networking, the International Telecommunications Union (ITU) proposed the OSI model.
- It had 7 layers.

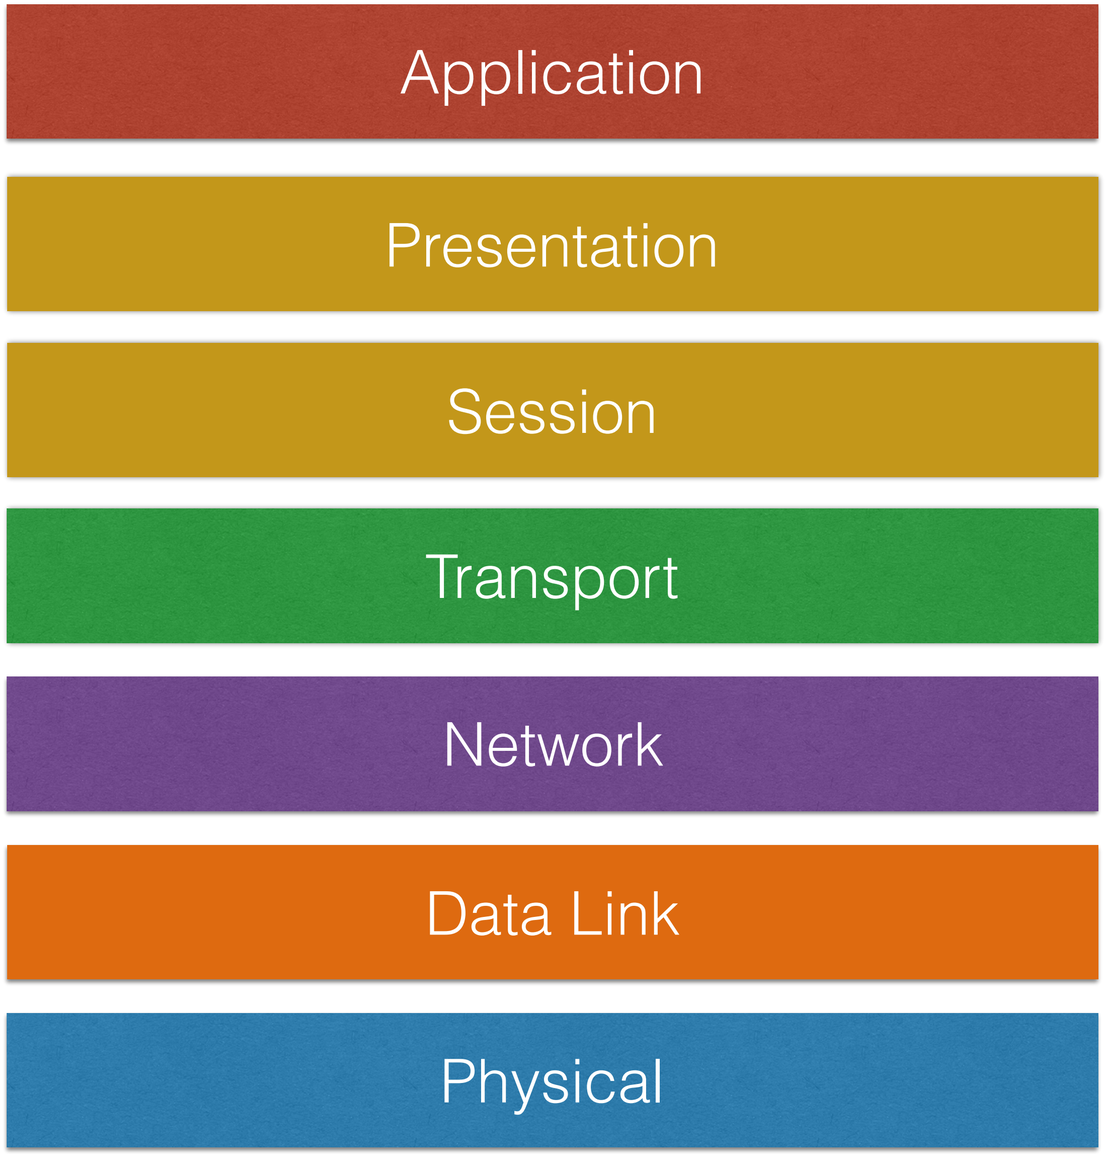

1983: The OSI Model

- … the Session and Presentation layers didn't really work out in practice.
- We don't really talk about them anymore.
- Byeeeeeeeee

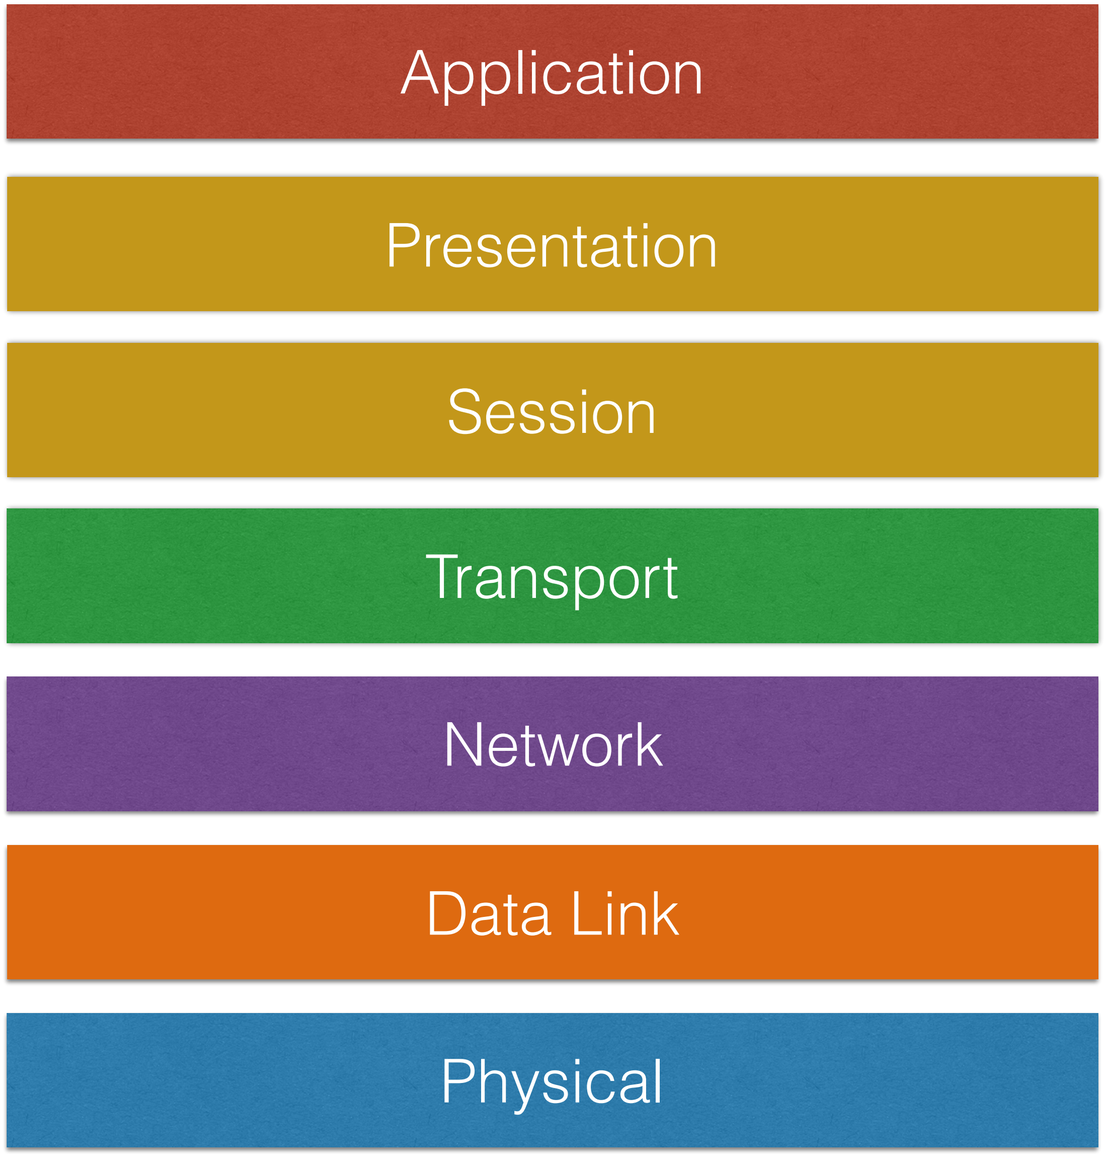

Back to Layer 7

Application

Connection

Transport

Globally Routable Packets

Network / INTERNET

Packets

Data Link

01010101001000000100111111010101

Physical Layer

It’s your application!

NETFLIX

HTTP
LISO WEB SERVER

Understanding the Layers of the Internet Model

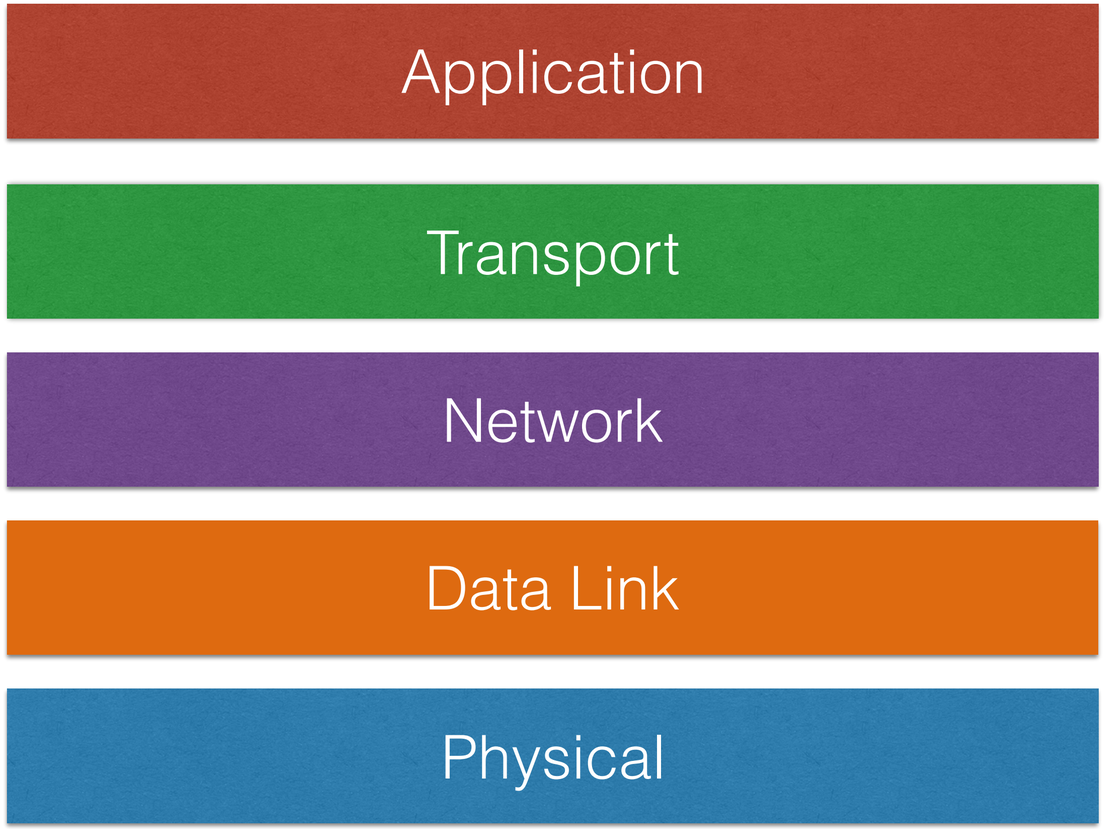

Top Down:
Encapsulation
Model

Understanding the Layers of the Internet Model

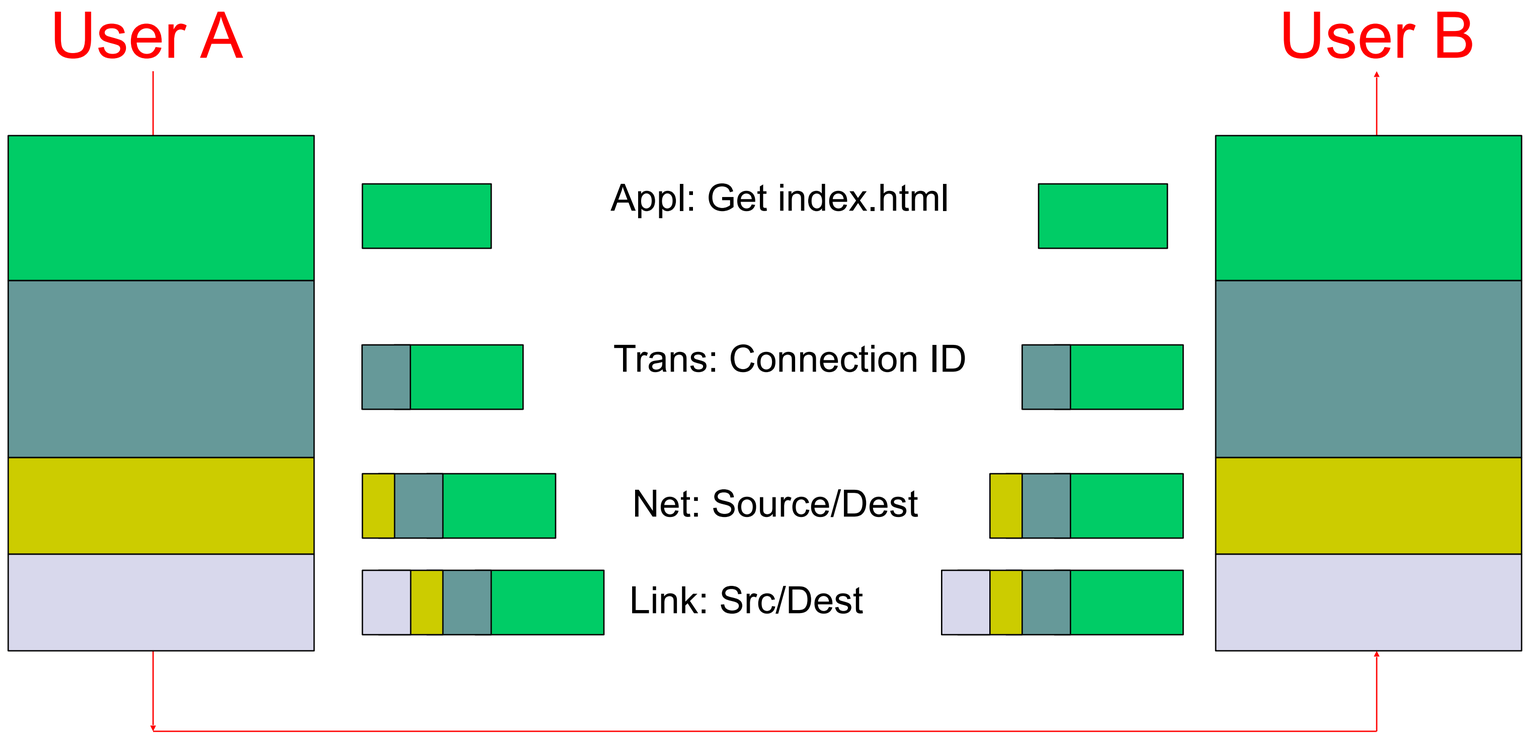

Understanding the Layers of the Internet Model

Application sends a chunk of data (bytestream) down to transport layer.

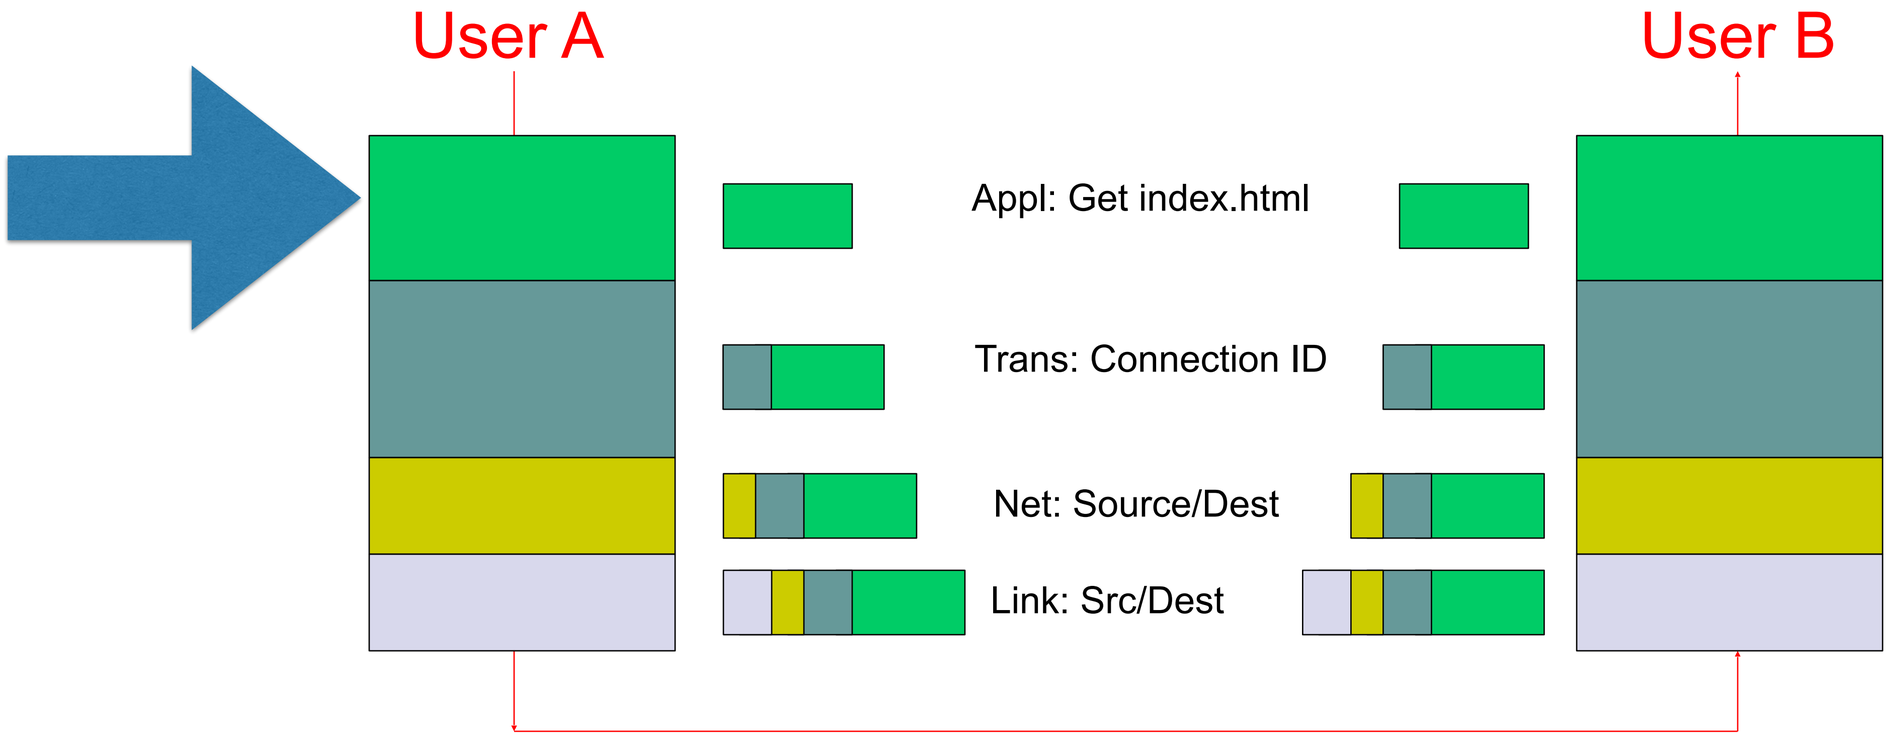

Understanding the Layers of the Internet Model

Transport layer chops the data into pieces, adds a connection ID & header.

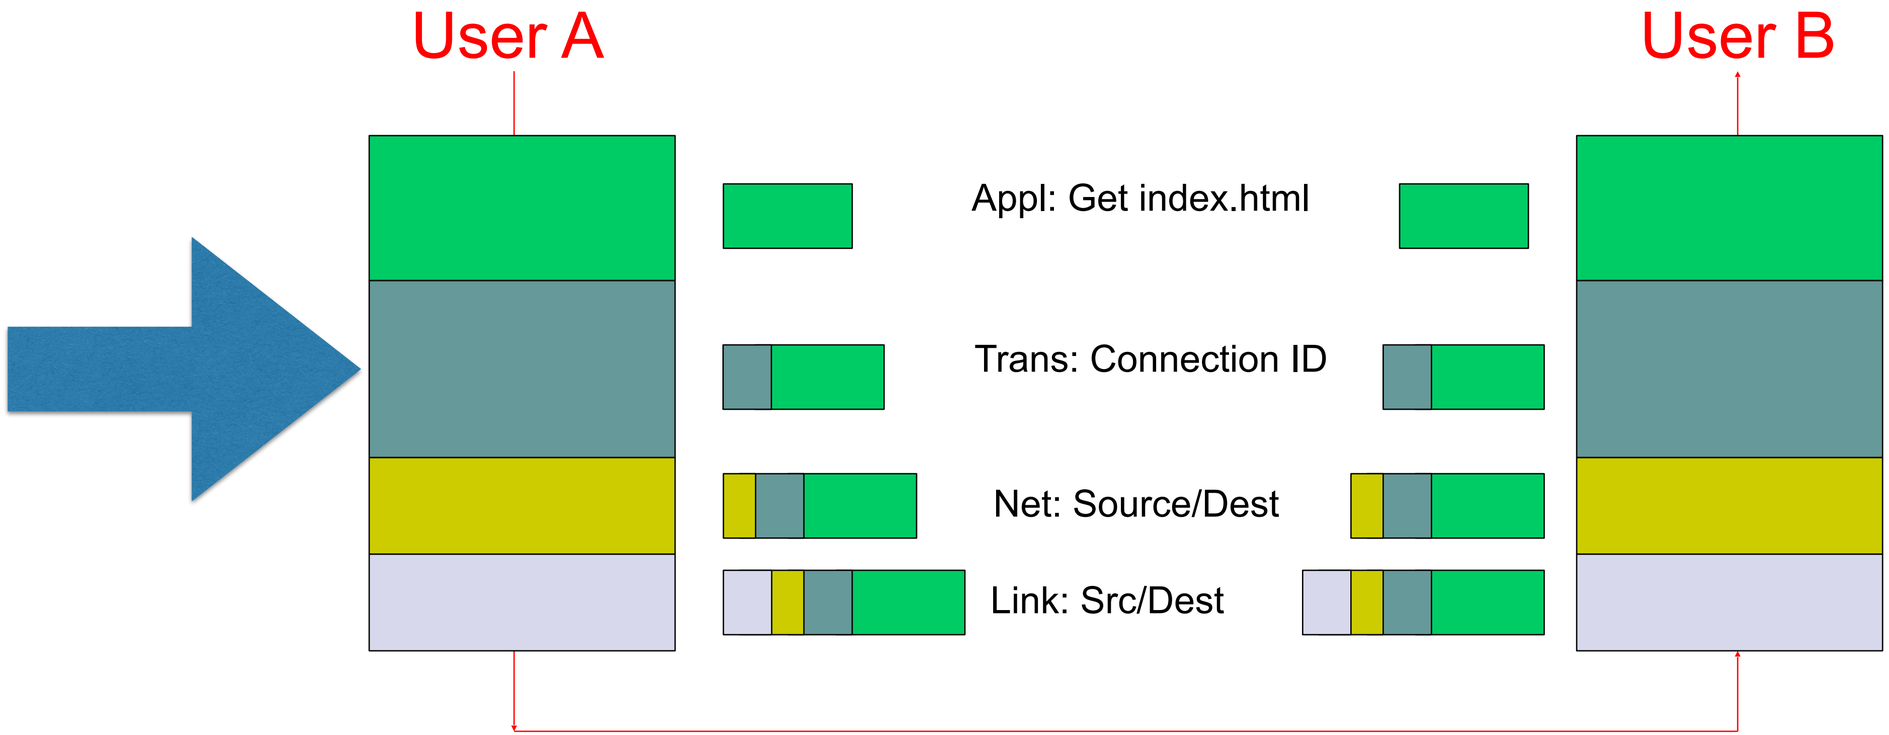

Understanding the Layers of the Internet Model

Transport layer hands down to IP layer where packet gets an IP address.

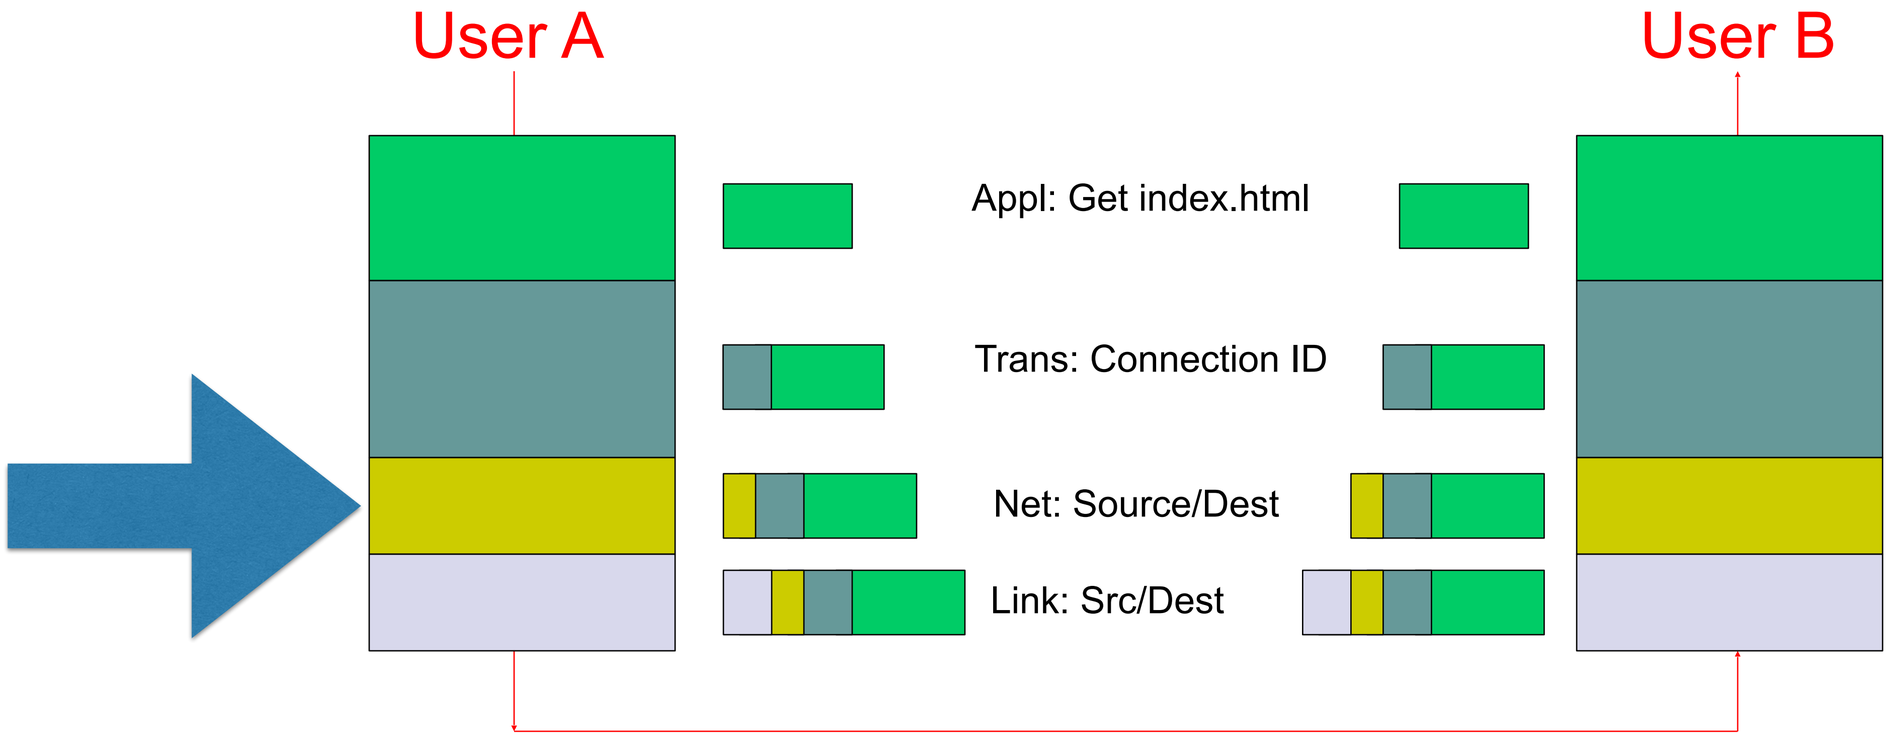

Understanding the Layers of the Internet Model

Packet goes down to link layer where packet gets, eg. an Ethernet header.

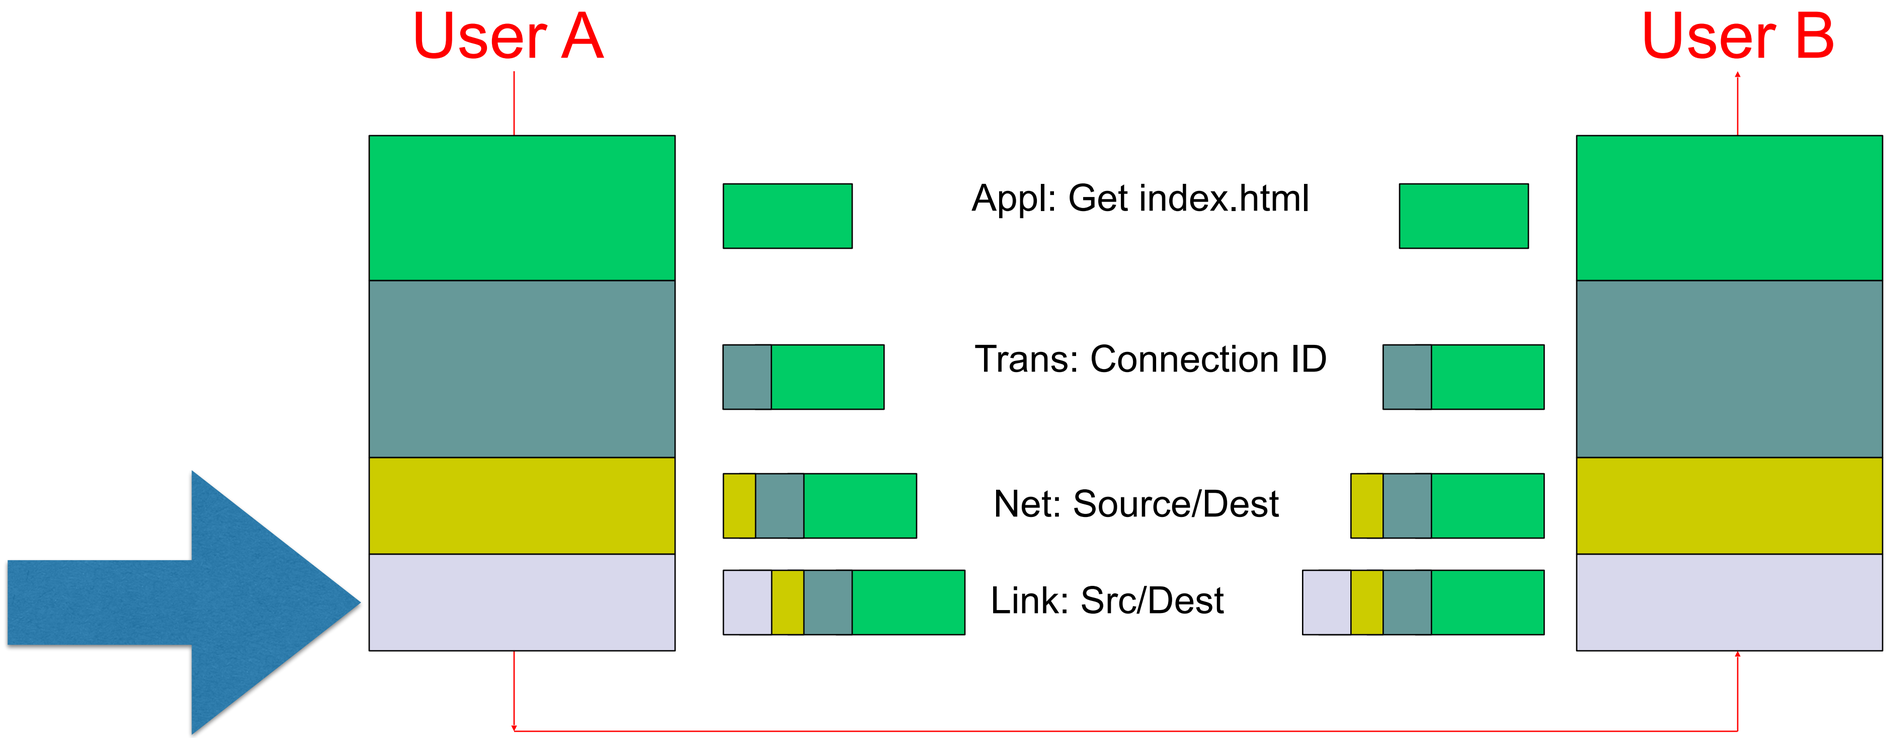

Understanding the Layers of the Internet Model

Packet goes down to link layer where packet gets, eg. an Ethernet header.

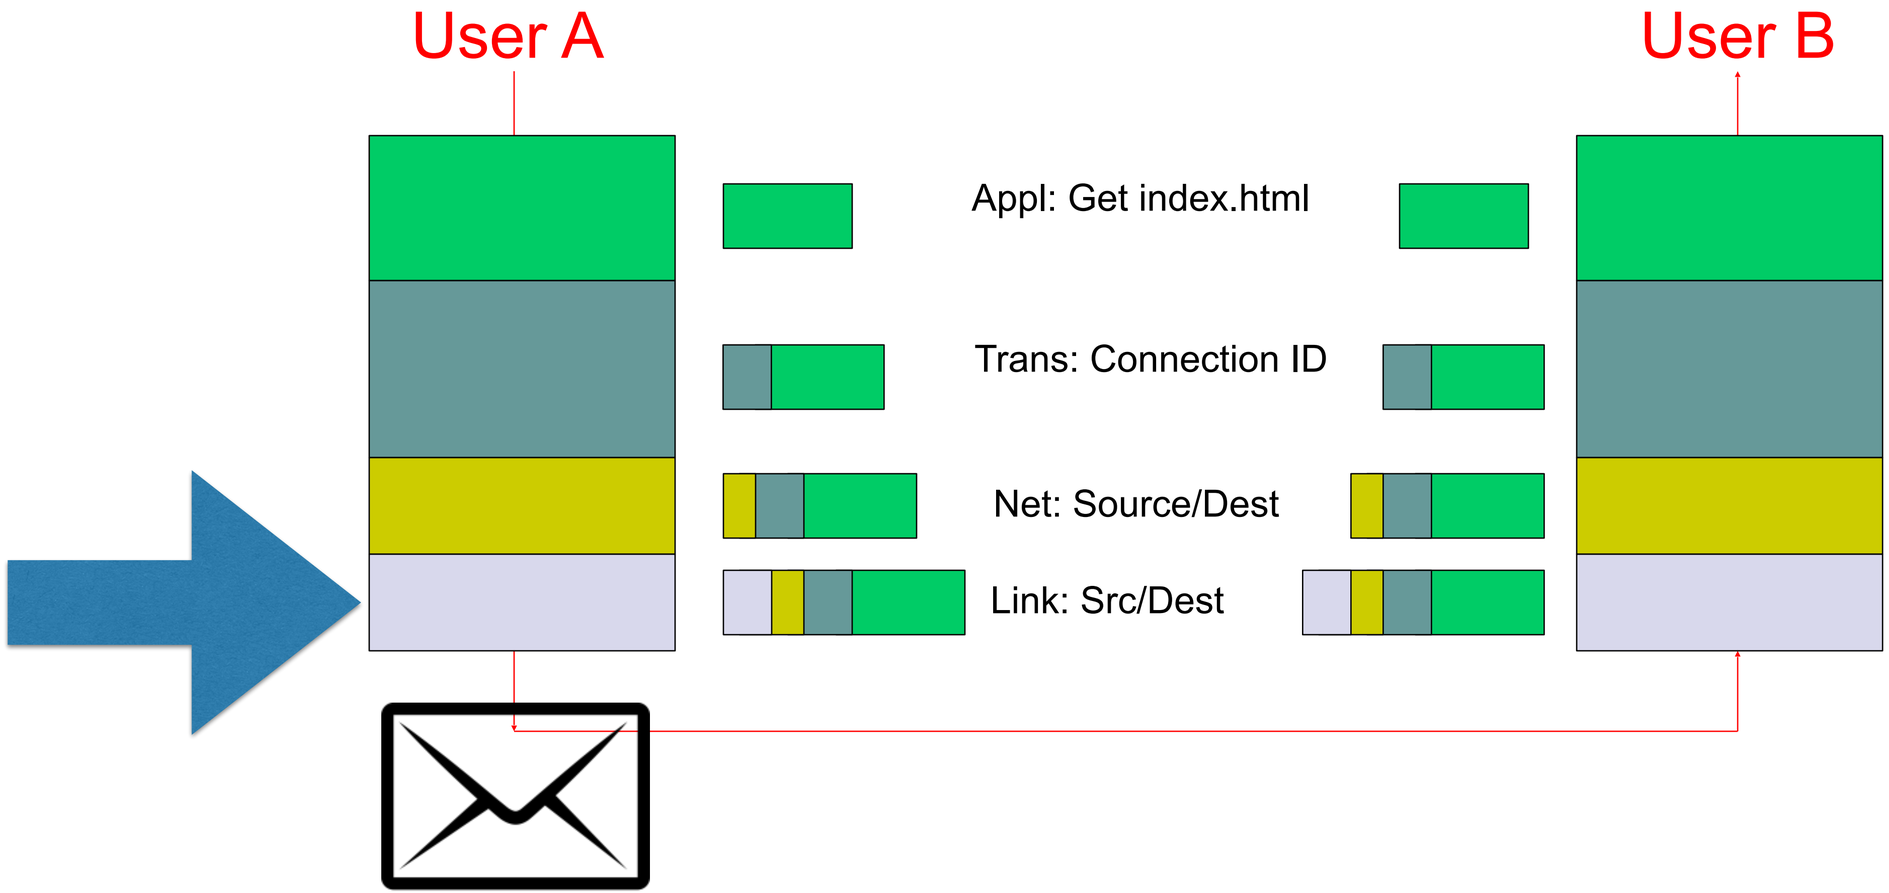

Understanding the Layers of the Internet Model

Strip off link layer header.

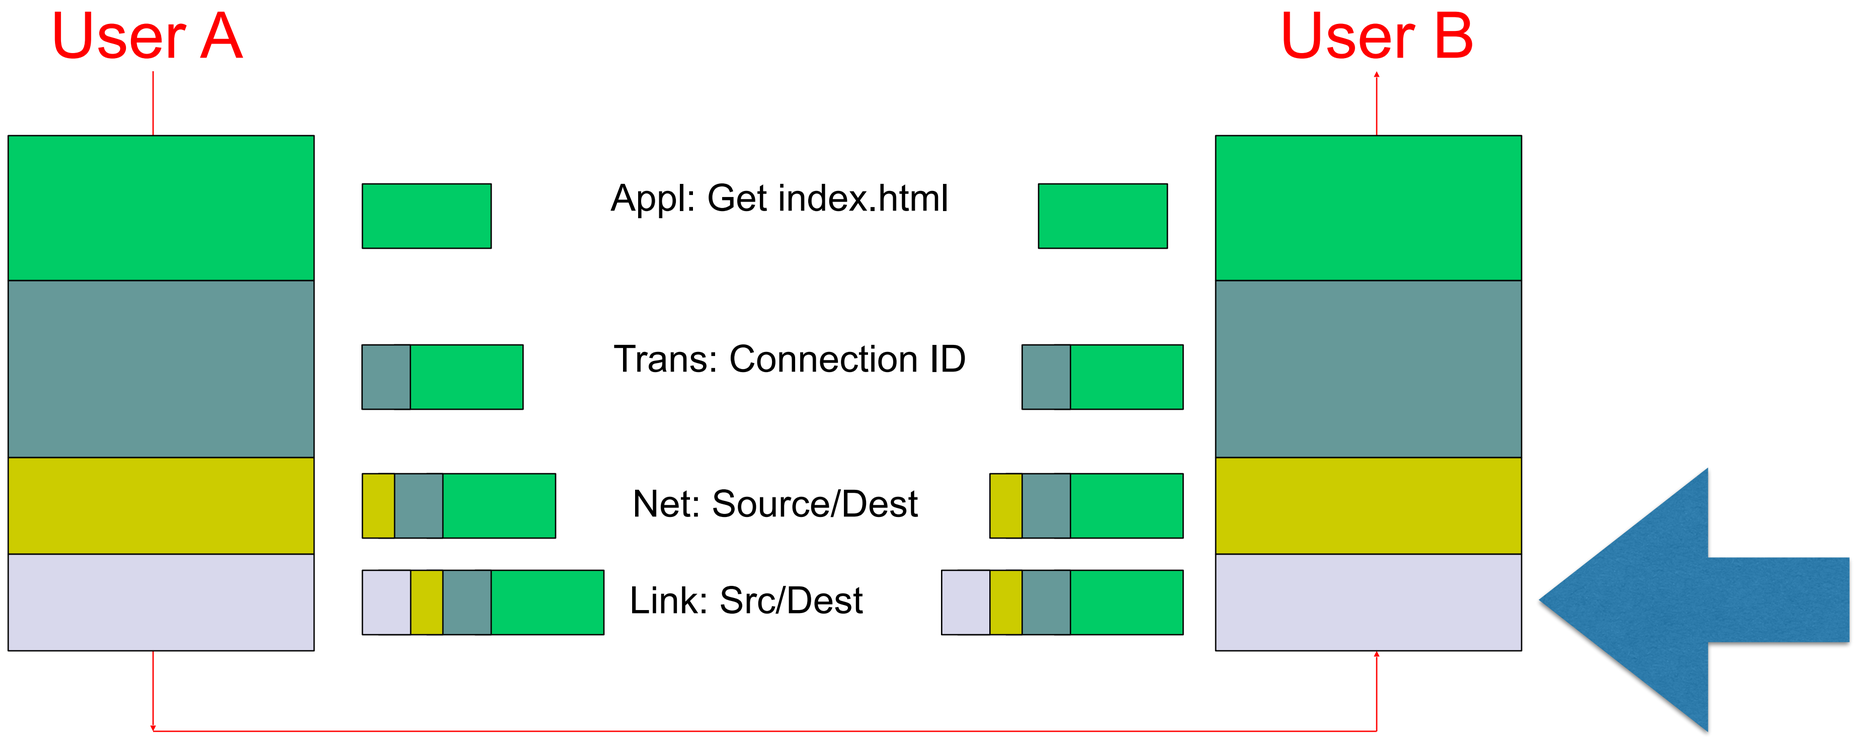

Understanding the Layers of the Internet Model

Strip off IP header.

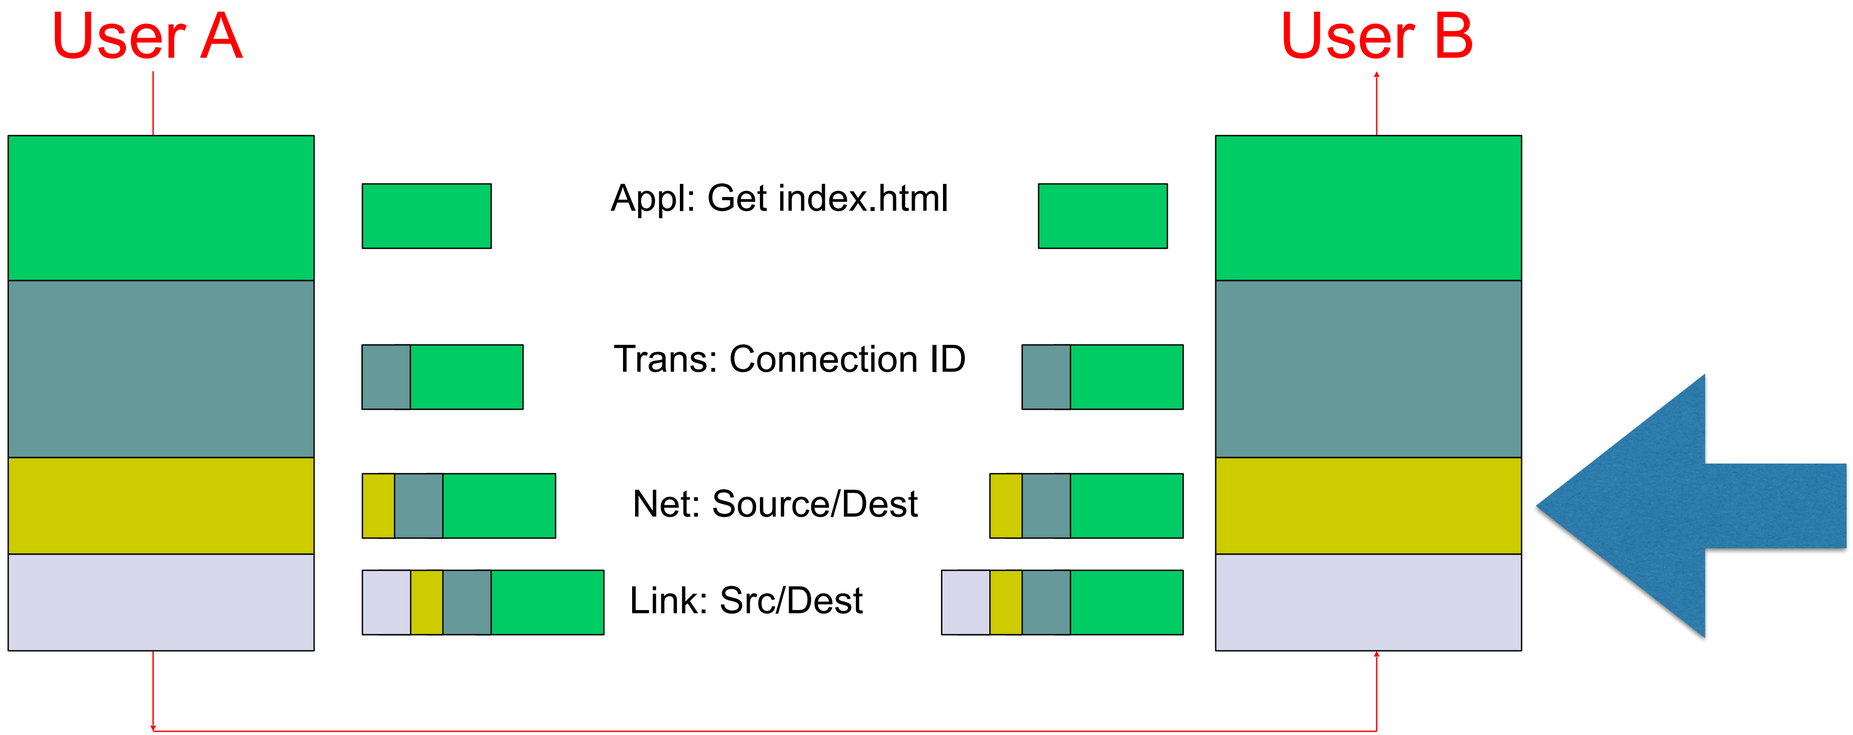

Understanding the Layers of the Internet Model

Strip off transport header, potentially put packets back in order.

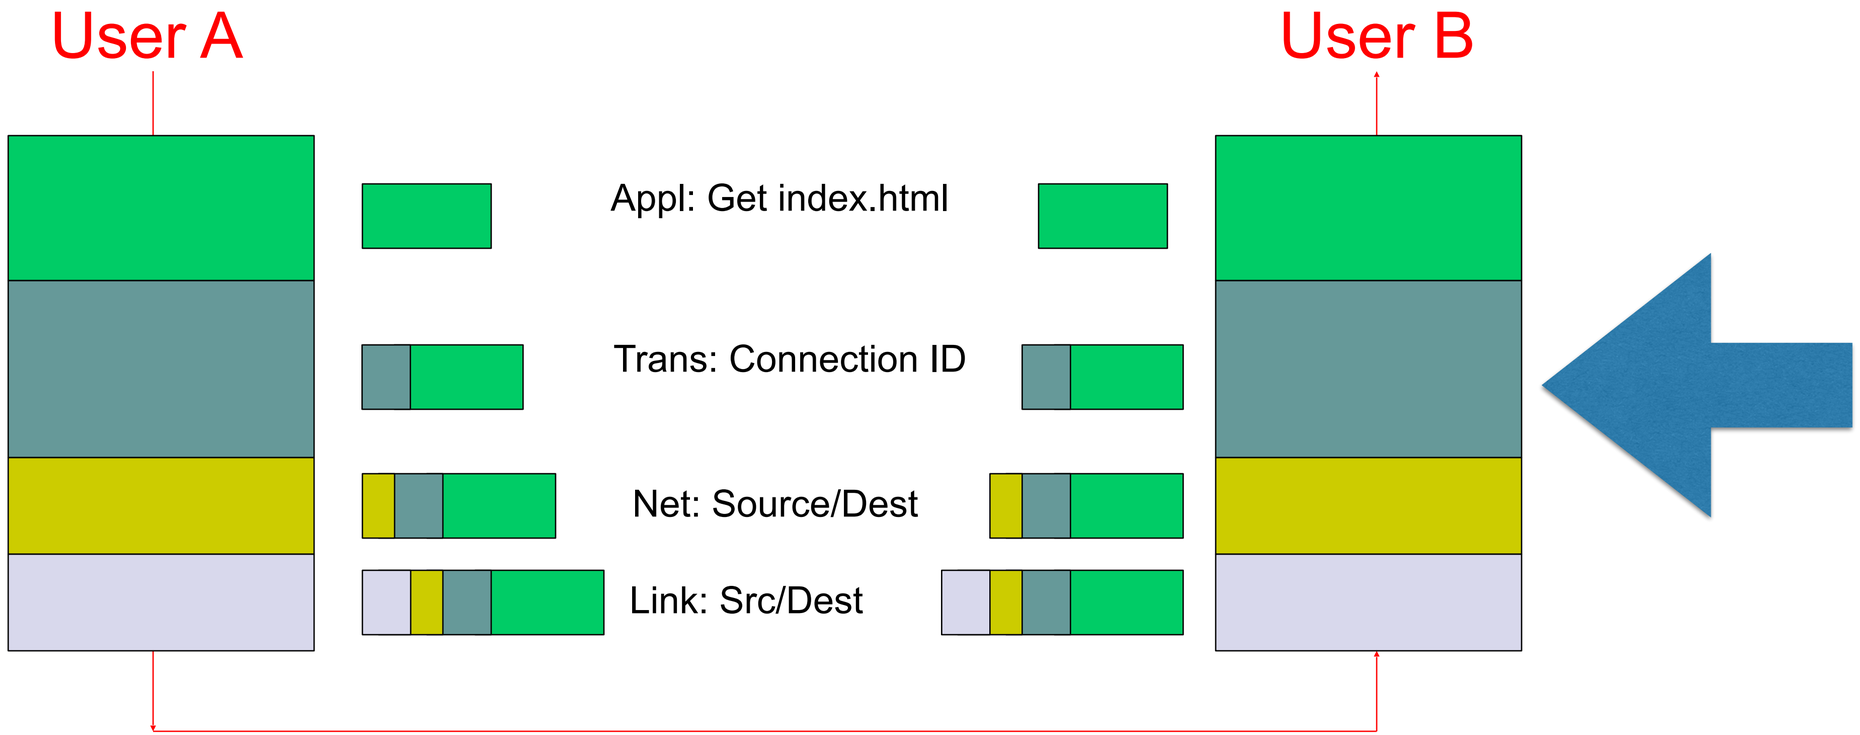

Understanding the Layers of the Internet Model

Application finally reads the data that was sent.

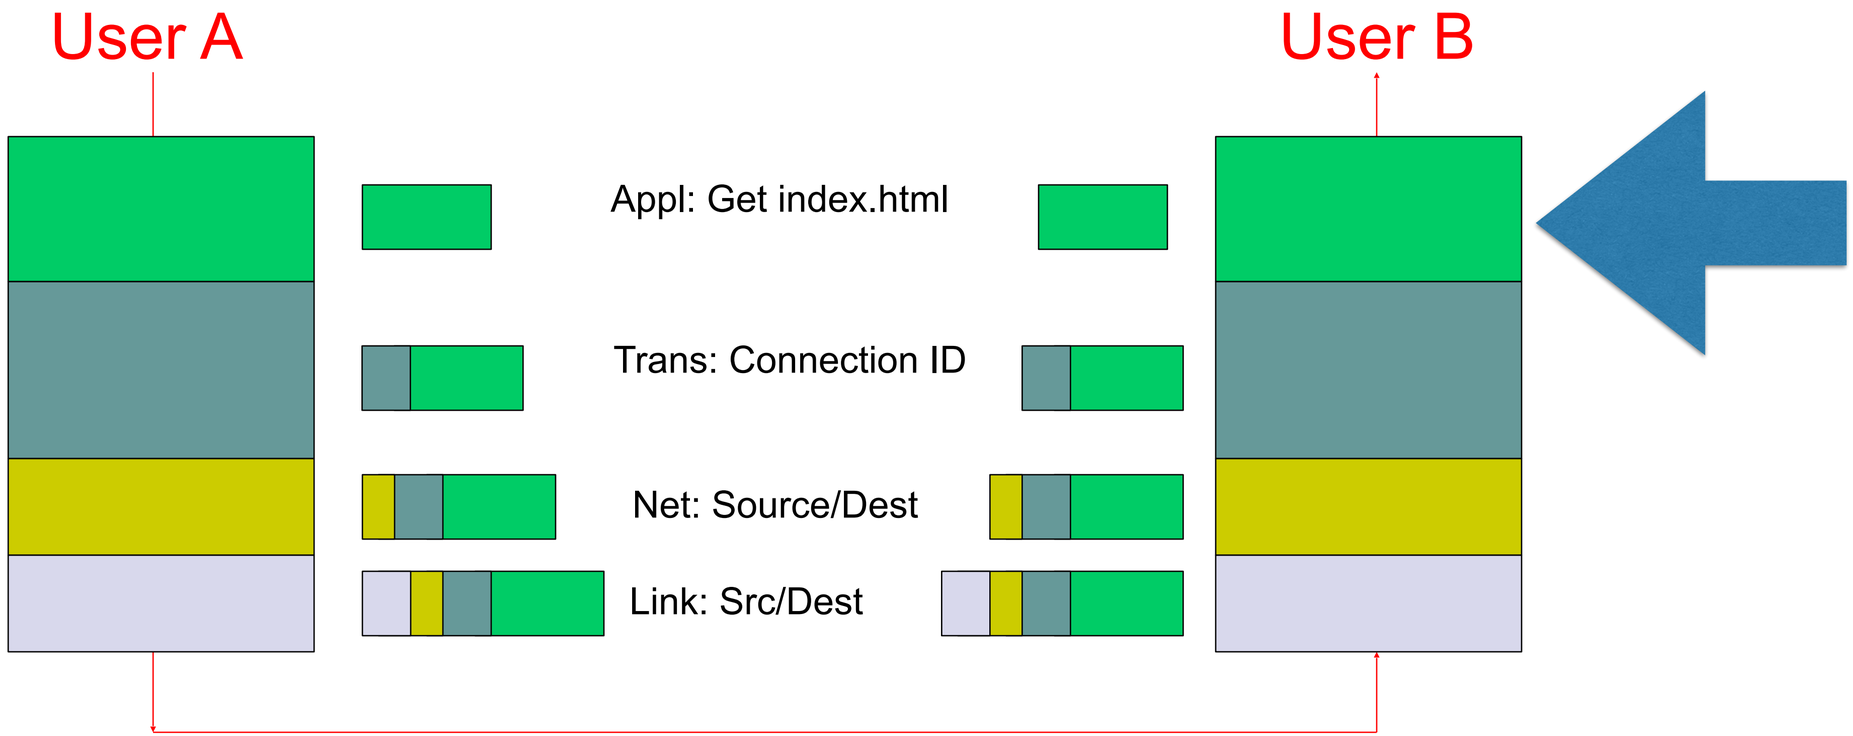

So what goes in what layer?

J. H. Saltzer, D. P. Reed, and D. D. Clark. 1984. End-to-end arguments in system design. ACM Trans. Comput. Syst. 2, 4 (November 1984), 277-288.

We were still trying to understand this “layering” business. What are the layers. And what do we do with them?

“Careful File Transfer”

At host A the file transfer program calls upon the file system to read the file from the disk, where it resides on several tracks, and the file system passes it to the file transfer program in fixed-size blocks chosen to be disk-format independent.

A

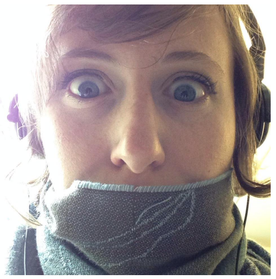

File System

Program

Network

B

File System

Program

Network

“Careful File Transfer”

At host A the file transfer program calls upon the file system to read the file from the disk, where it resides on several tracks, and the file system passes it to the file transfer program in fixed-size blocks chosen to be disk-format independent.

A

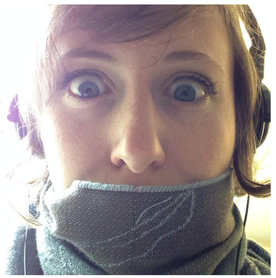

File System

Program

Network

Alternate *ACM* SIG Proceedings Paper in LaTeX Format*

[Extended Abstract]†

B

File System

Program

Network

“Careful File Transfer”

At host A the file transfer program calls upon the file system to read the file from the disk, where it resides on several tracks, and the file system passes it to the file transfer program in fixed-size blocks chosen to be disk-format independent.

A

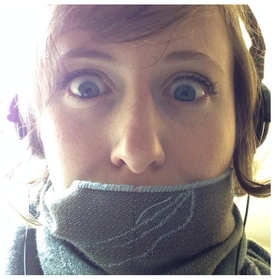

B

File System

Program

Network

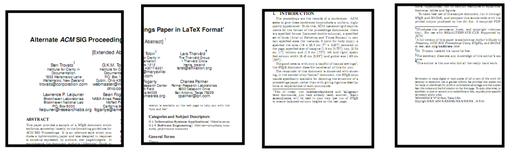

File System

Program

Network

“Careful File Transfer”

Also at host A the file transfer program asks the data communication system to transmit the file using some communication protocol that involves splitting the data into packets. The packet size is typically different from the file block size and the disk track size.

A

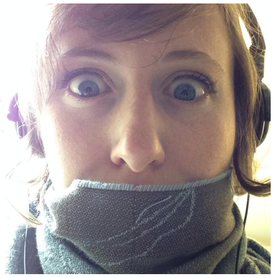

B

File System

Program

Network

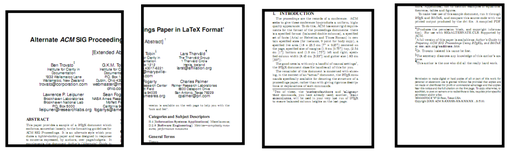

File System

Program

Network

“Careful File Transfer”

Also at host A the file transfer program asks the data communication system to transmit the file using some communication protocol that involves splitting the data into packets. The packet size is typically different from the file block size and the disk track size.

A

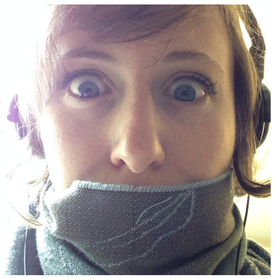

File System

Program

Network

B

File System

Program

Network

“Careful File Transfer”

The data communication network moves the packets from computer A to computer B.

A

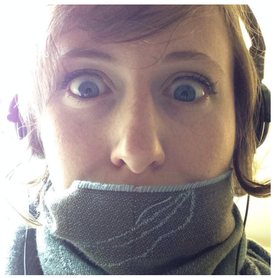

B

File System

Program

Network

File System

Program

Network

“Careful File Transfer”

At host B a data communication program removes the packets from the data communication protocol and hands the contained data on to a second part of the file transfer application, the part that operates within host B.

A

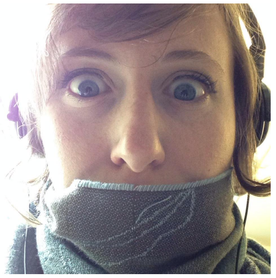

File System

Program

Network

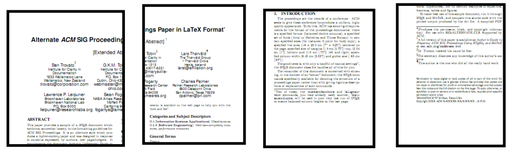

B

File System

Program

Network

“Careful File Transfer”

At host B, the file transfer program asks the file system to write the received data on the disk of host B.

A

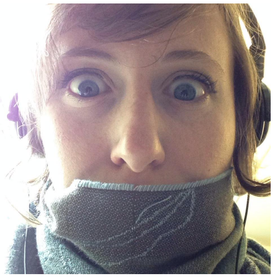

File System

Program

Network

Alternate *ACM* SIG Proceedings Paper in LaTeX Format

[Extended Abstract]

B

File System

Program

Network

What if Zeeshan later reads the file and find it is corrupted? What could have gone wrong?

The file, though originally written correctly onto the disk at host A, if read now may contain incorrect data, perhaps because of hardware faults in the disk storage system.

The software of the file system, the file transfer program, or the data communication system might make a mistake in buffering and copying the data of the file, either at host A or host B.

The hardware processor or its local memory might have a transient error while doing the buffering and copying, either at host A or host B.

The communication system might drop or change the bits in a packet, or lose a packet or deliver a packet more than once.

Either of the hosts may crash part way through the transaction after performing an unknown amount (perhaps all) of the transaction.

How do we re-design our system to make sure the file doesn’t get corrupted?

B

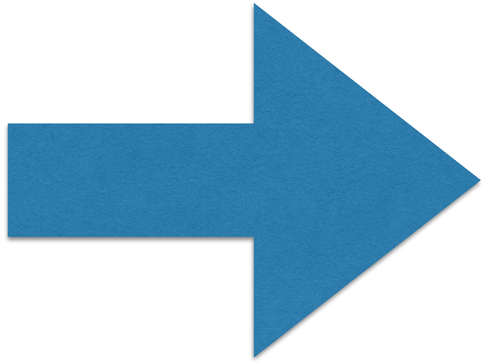

File System

Program

Network

B

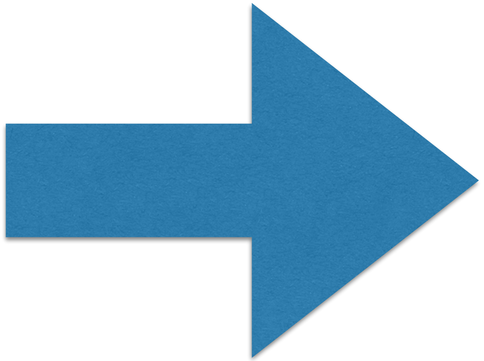

File System

Program

Network

B

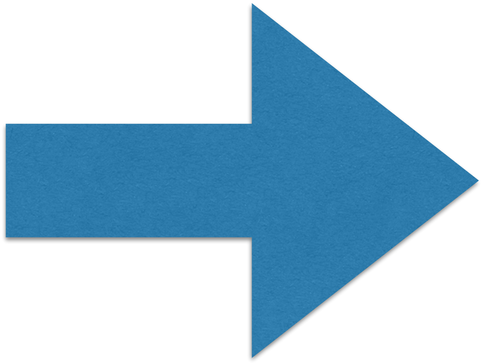

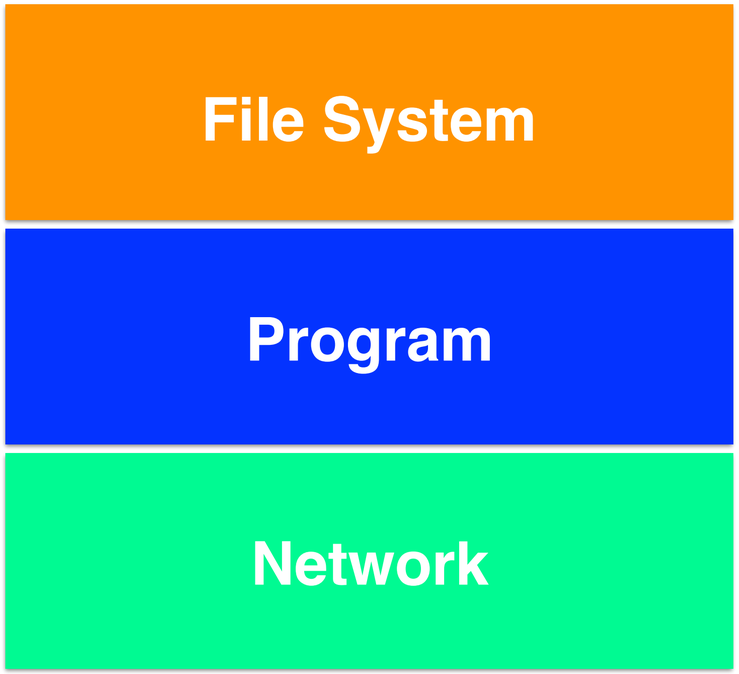

The End-to-End Argument

[If] the function in question can completely and correctly be implemented with the knowledge and help of the application standing at the endpoints of the communication system:

[Then] providing that questioned function as a feature of the communication system [or lower layer] is not possible.

[However], sometimes an incomplete version of the function provided by the communication system may be useful as a performance enhancement.

Let's say we had a perfectly reliable network

A

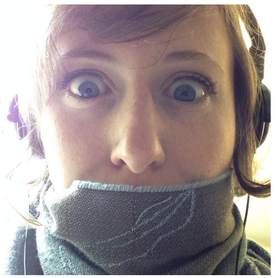

B

Would that solve our reliability problem?

The file, though originally written correctly onto the disk at host A, if read now may contain incorrect data, perhaps because of hardware faults in the disk storage system.

The software of the file system, the file transfer program or the data communication system might make a mistake in buffering and copying the data of the file, either at host A or host B.

The hardware processor or its local memory might have a transient error while doing the buffering and copying, either at host A or host B.

The communication system might drop or change the bits in a packet, or lose a packet or deliver a packet more than once.

Either of the hosts may crash part way through the transaction after performing an unknown amount (perhaps all) of the transaction.

Would that solve our reliability problem?

The file, though originally written correctly onto the disk at host A, if read now may contain incorrect data, perhaps because of hardware faults in the disk storage system.

The software of the file system, the file transfer program or the data communication system might make a mistake in buffering and copying the data of the file, either at host A or host B.

The hardware processor or its local memory might have a transient error while doing the buffering and copying, either at host A or host B.

The communication system might drop or change the bits in a packet, or lose a packet or deliver a packet more than once.

Either of the hosts may crash part way through the transaction after performing an unknown amount (perhaps all) of the transaction.

Well, that wasn’t very helpful…

“End to End Check and Retry”

A

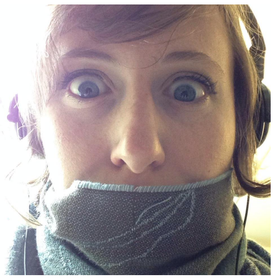

File System

Program

Network

B

File System

Program

Network

“End to End Check and Retry”

A

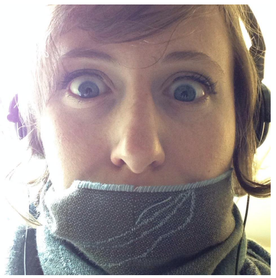

Read file and its checksum from disk.
Verify file + checksum.
Send File AND Checksum.


Alternate *ACM* SIG Proceedings Paper in LaTeX Format*
[Extended Abstract]

B

File System

Program

Network

File System

Program

Network

“End to End Check and Retry”

A

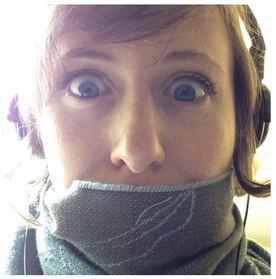

File System

Program

Network

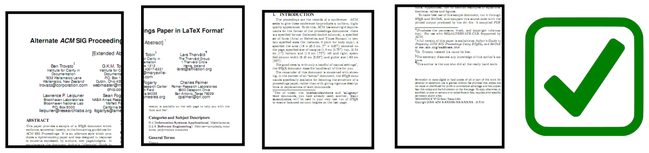

B

File System

Program

Network

“End to End Check and Retry”

A

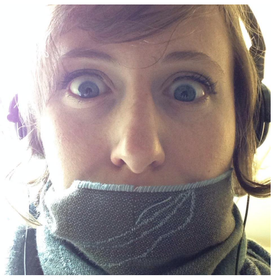

File System

Program

Network

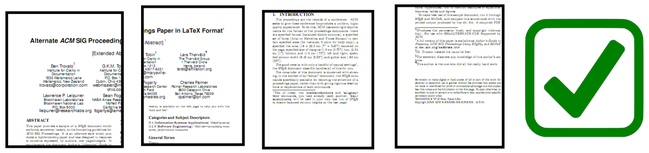

B

File System

Program

Network

“Careful File Transfer”

The data communication network moves the packets from computer A to computer B.

A

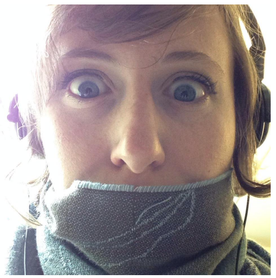

B

File System

Program

Network

File System

Program

Network

“End to End Check and Retry”

A

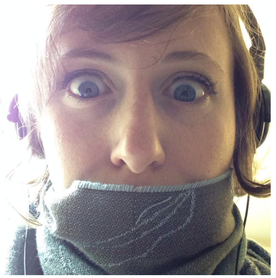

File System

Program

Network

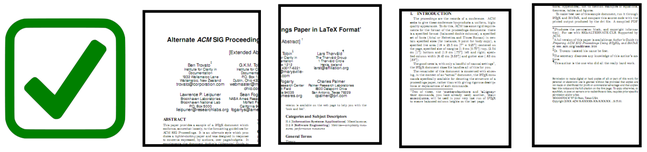

B

File System

Program

Network

“End to End Check and Retry”

A

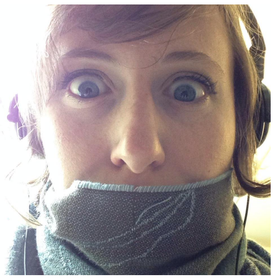

File System

Program

Network

Write file and checksum to disk.
Then read back and double-check that checksum + file verify.

Alternate *ACM* SIG Proceedings Paper in LaTeX Format

[Extended Abstract]

B

File System

Program

Network

“End to End Check and Retry”

A

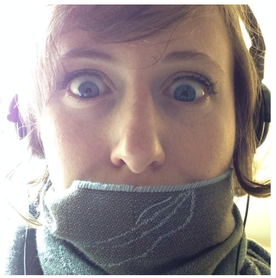

File System

Program

Network

If Checksum doesn’t match?
Just ask Justine to re-send.
(ie, try all over again!)

Alternate *ACM* SIG Proceedings Paper in LaTeX Format

[Extended Abstract]

B

File System

Program

Network

Would that solve our reliability problem?

The file, though originally written correctly onto the disk at host A, if read now may contain incorrect data, perhaps because of hardware faults in the disk storage system.

The software of the file system, the file transfer program, or the data communication system might make a mistake in buffering and copying the data of the file, either at host A or host B.

The hardware processor or its local memory might have a transient error while doing the buffering and copying, either at host A or host B.

The communication system might drop or change the bits in a packet, or lose a packet or deliver a packet more than once.

Either of the hosts may crash part way through the transaction after performing an unknown amount (perhaps all) of the transaction.

Lesson: If you can do it at the "higher" layer, don't bother implementing it at a lower layer.

Other places to apply E2E in Networks

- Encryption
- First-in-first-out ordering
- Duplicate message surpression
- Multi-message transactions

The End-to-End Argument

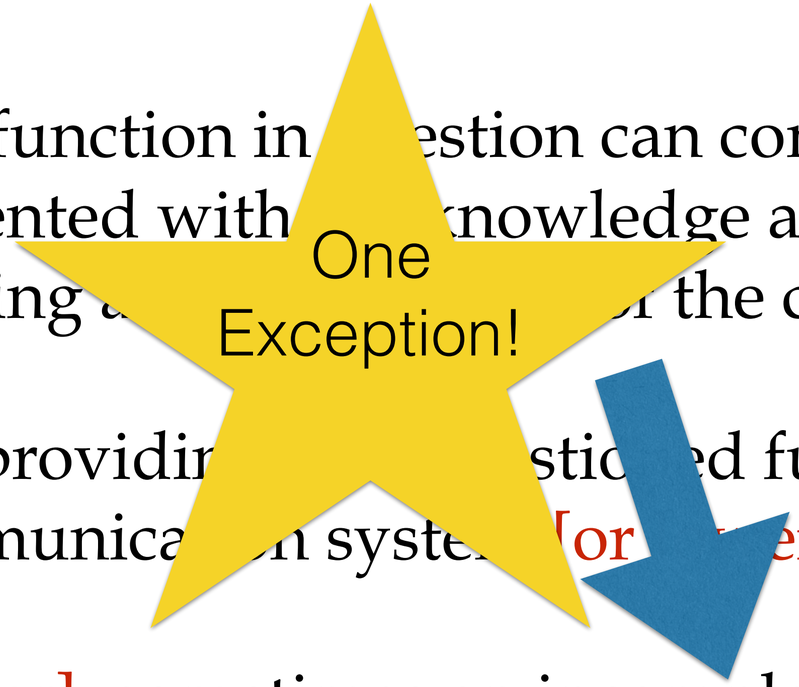

[If] the function in estion can completely and correctly be implemented with knowledge and help of the application standing a the communication system:

[Then] providi sti ed function as a feature of the communica n syste [or er layer] is not possible.

[However], sometimes an incomplete version of the function provided by the communication system may be useful as a performance enhancement.

What if 90% of my loss really was happening at the network layer?

A

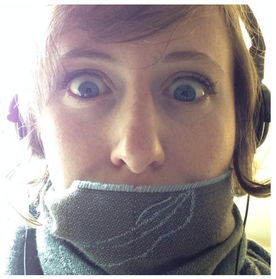

File System

Program

Network

As a performance optimization, you might want to implement it in the lower layer anyway (redundantly).

B

File System

Program

Network

“End to End Check and Retry” + A Reliable Network

A

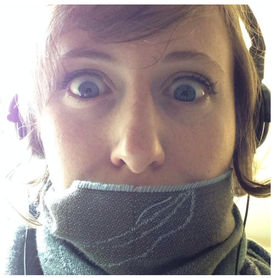

B

Anyone have any other examples where this plays out?

The “Strong” End-to-End Argument

It’s not just a waste of time to put non-essential functionality at lower layers: *it’s actually harmful*.

“End to End Check and Retry” + A Reliable Network

A

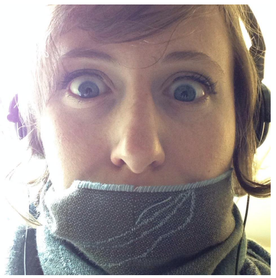

B

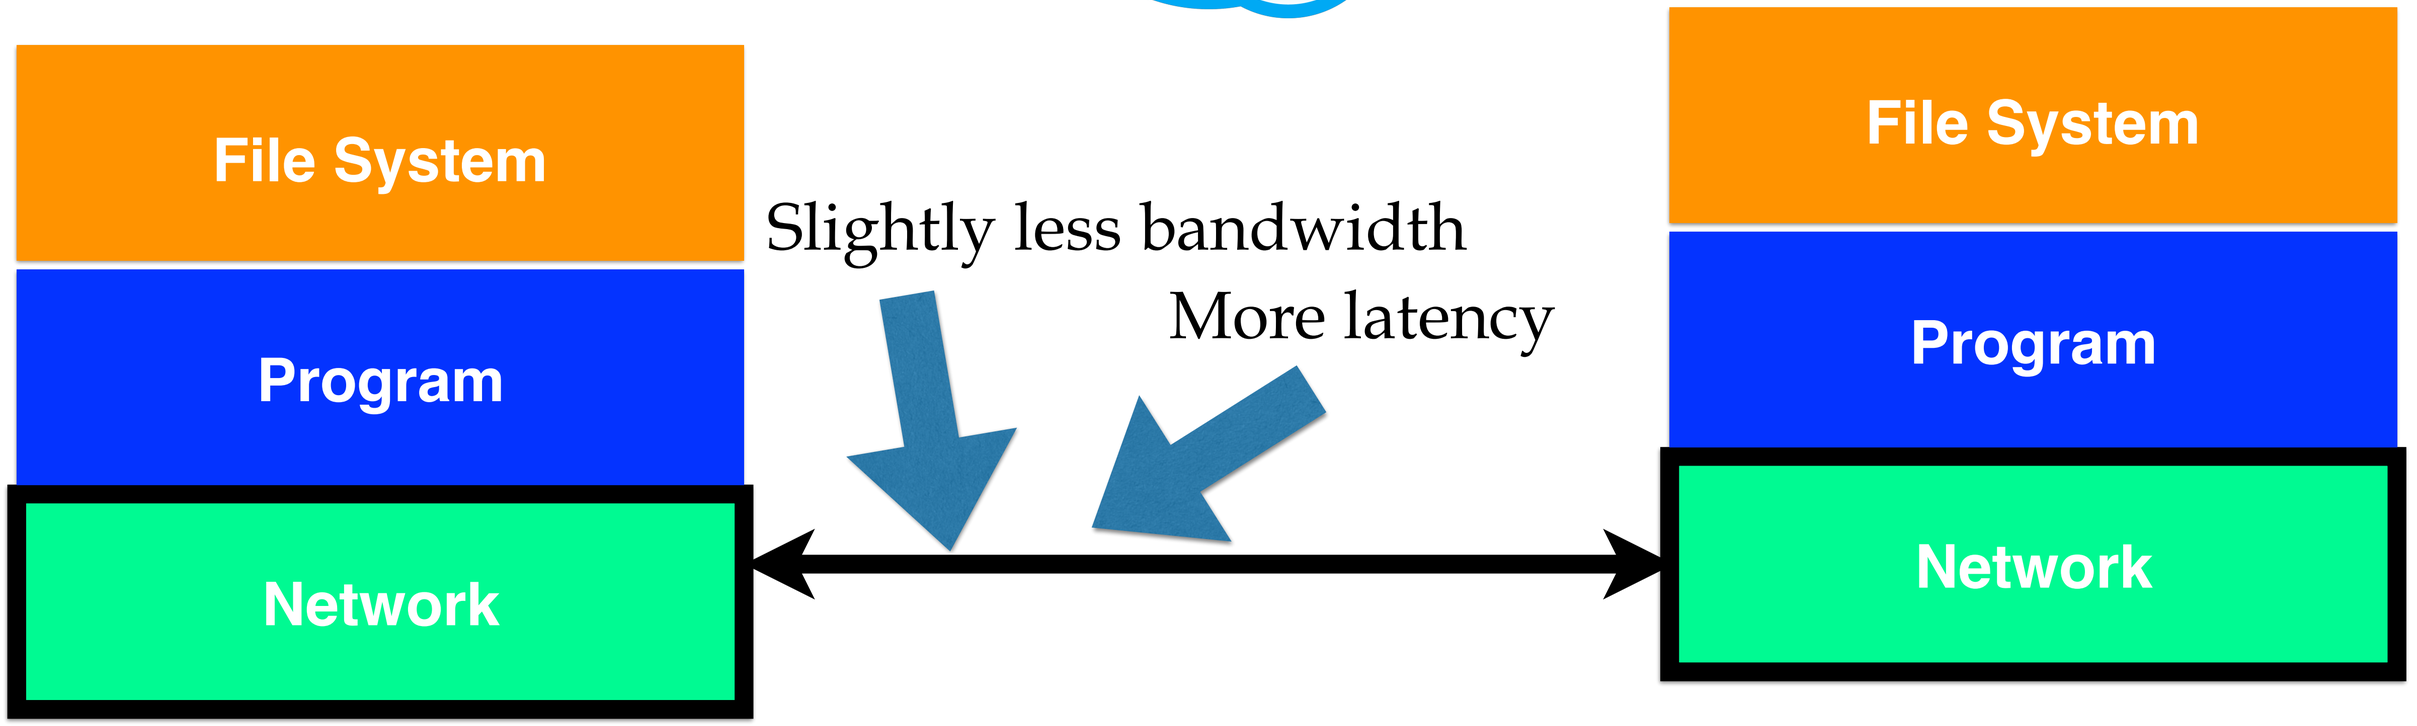

Some applications may be *constrained* by the new functionality.

Firewalls and Intrusion Detection

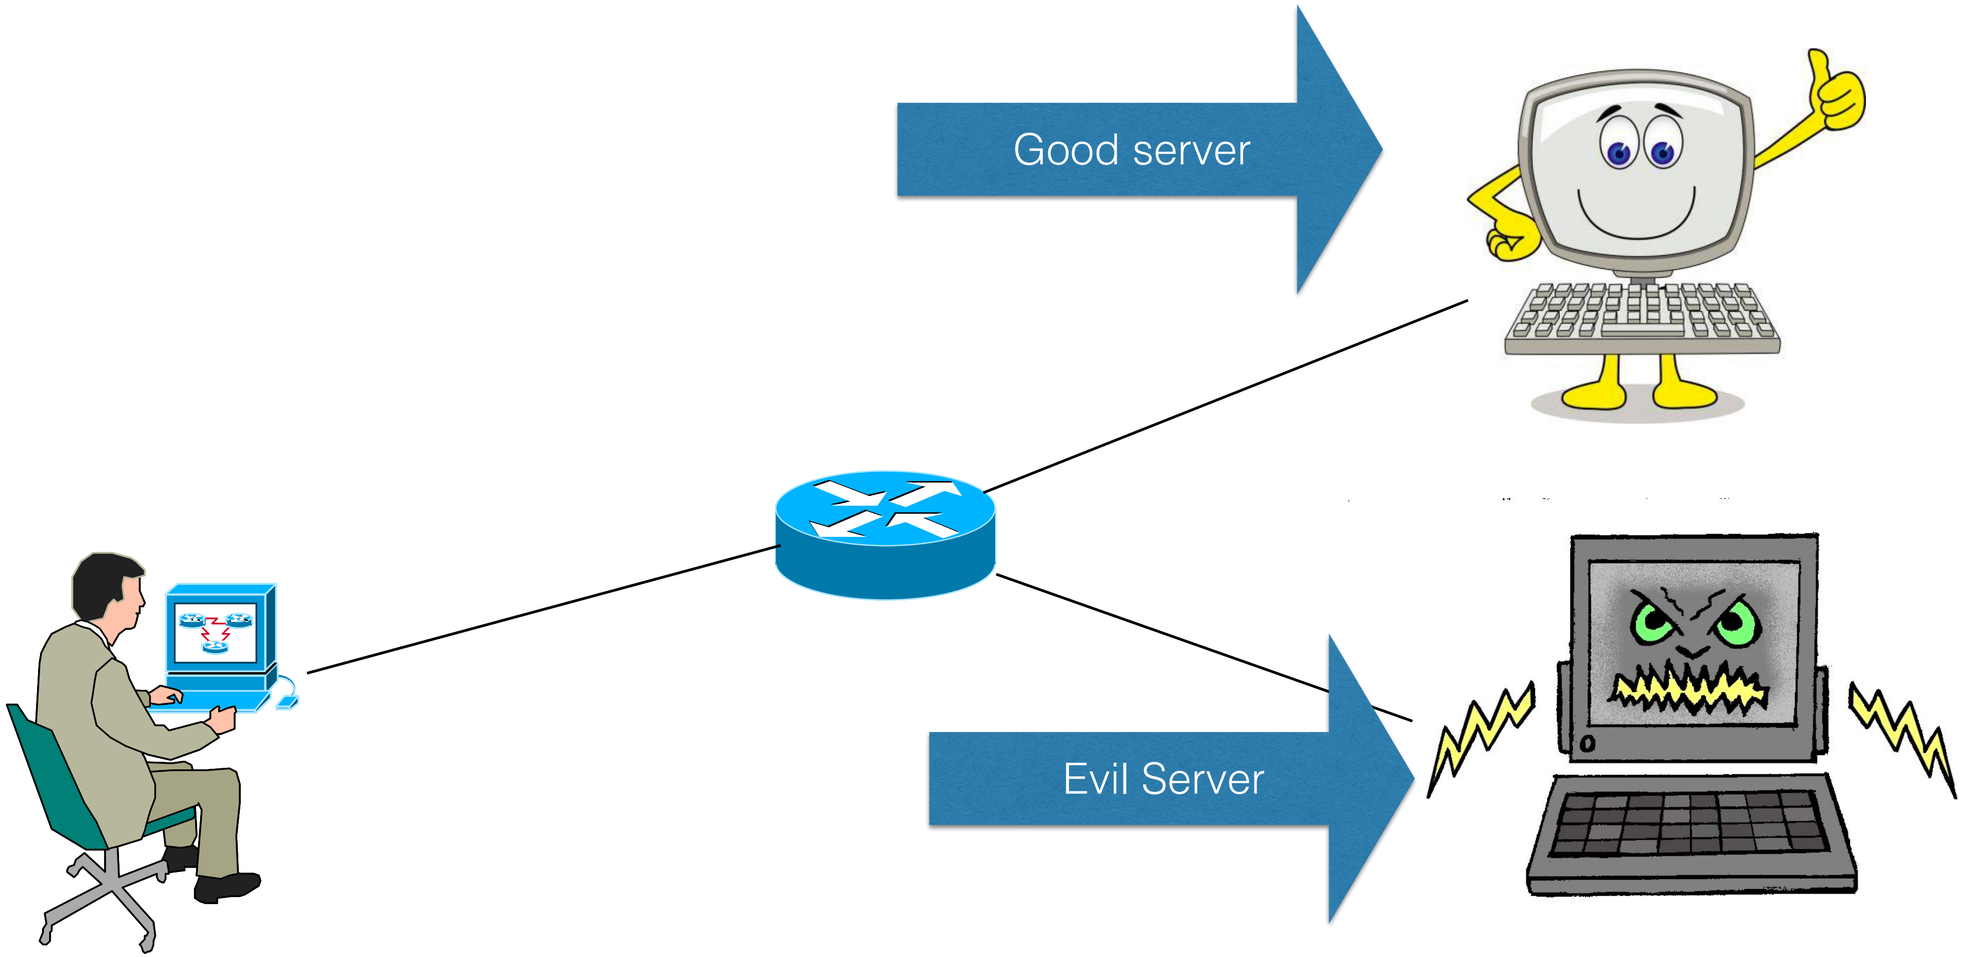

Firewalls and Intrusion Detection

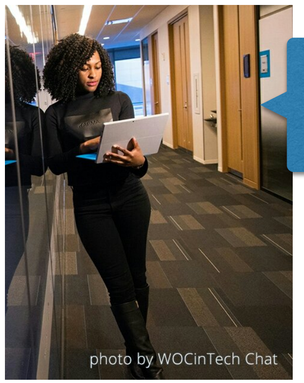

I need to protect my users!

Web traffic, email

IRC, strange port numbers

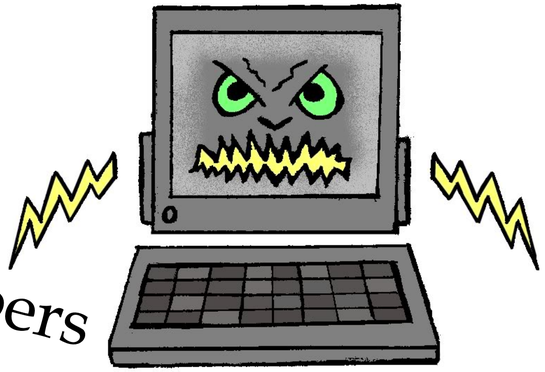

Firewalls and Intrusion Detection

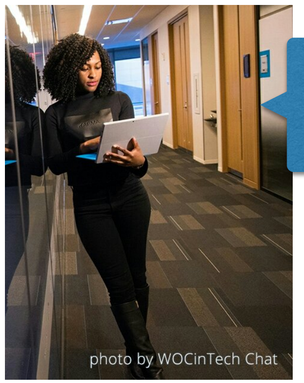

I need to protect my users!

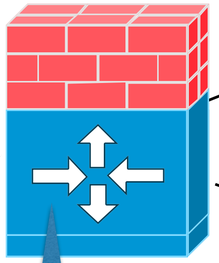

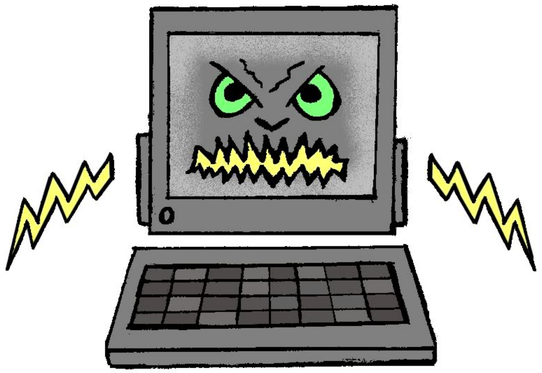

Only allow web and email!

Firewalls and Intrusion Detection

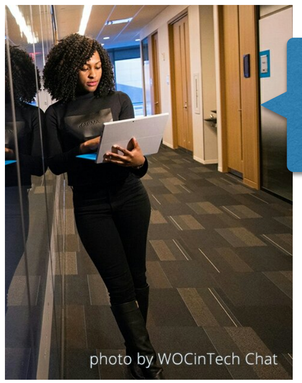

I need to protect my users!

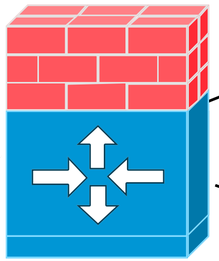

But what if I have a cool new app?

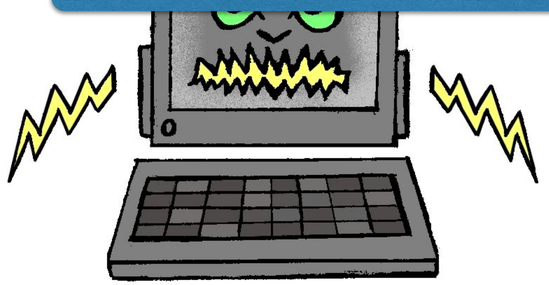

End to End Argument: Recap

- Basic argument: If you can implement functionality correctly and completely at endpoints, do it there and not at a lower layer.
  - It saves on redundant work in the system, and avoids confusion later. Exceptions okay for performance optimizations.
- Strong argument: Avoid putting unneeded functionality at lower layers of your system altogether because it's harmful!
  - Extra functionality at low layers constrains how applications are designed at higher layers.

Now you know

- The 5 layers define the architecture of the Internet today.
- You already know about two of them!
- Starting today we'll dig into the Network Layer with Prof Steenkiste. Next week we'll move into transport. And then — since you've built your Liso Server at the Application Layer — you'll know how the Internet works :-)
- In general, we try to push as much functionality as we can into higher layers, rather than lower ones, because of the End to End Argument.