Structure of the Internet

David Conrad
drc@isc.org
Internet Software Consortium

Introduction

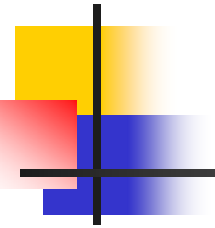

- "The Internet" does not exist!
  - You can't buy it
    - No one owns it
  - You can't connect to it
    - Really.
  - You can't regulate it
    - You can regulate parts
- However, you can become a part of it
  - But what exactly is "it"?
  - How is it structured?

Agenda

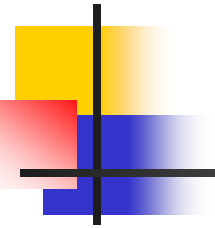

- Organizational Structure
  - How the Internet is put together
  - Internet History
  - Executive TCP/IP Summary
- Administrative Structure
  - Support, development, user, and operations groups

Structure of the Internet

- Today, the Internet is composed of over 61,000 autonomous run interconnected networks
  - The only guarantee is that all networks use the TCP/IP protocols
- Each machine on each network can offer useful services
  - Or not -- entirely dependent on the site
  - Each user of the network can use those services
- First **bi-directional** electronic communications medium that supports:
  - One to one (unicast, e.g. email) One to many (multi-cast, e.g., electronic news)
- Key point: Internet component networks are operated **AUTONOMOUSLY**
  - You don't connect to the Internet, you become part of it

Structure of the Internet (cont'd

- In the Beginning
  - The Internet was a research network with lots of compute geeks jabbering at each other
    - Sometimes even getting interesting work done
  - A very homogeneous community with essentially the same culture/language/values
- Now, the Internet has commercialized
  - People are trying to use the Internet for real work
    - And for the most part, succeeding
  - Some of those people may be nice or not so nice
    - UCE, crackers, con-artists, pornographers, etc.
    - Good Samaritans, hackers, philanthropists, etc.
      - The Internet reflects society

Structure of the Internet (cont'd)

- From a million miles up, the Internet looks hierarchical
  - Big ISPs (backbones) provide services to smaller ISPs
  - Smaller ISPs provide services to yet smaller ISPs
  - Somebody eventually provides services to end users

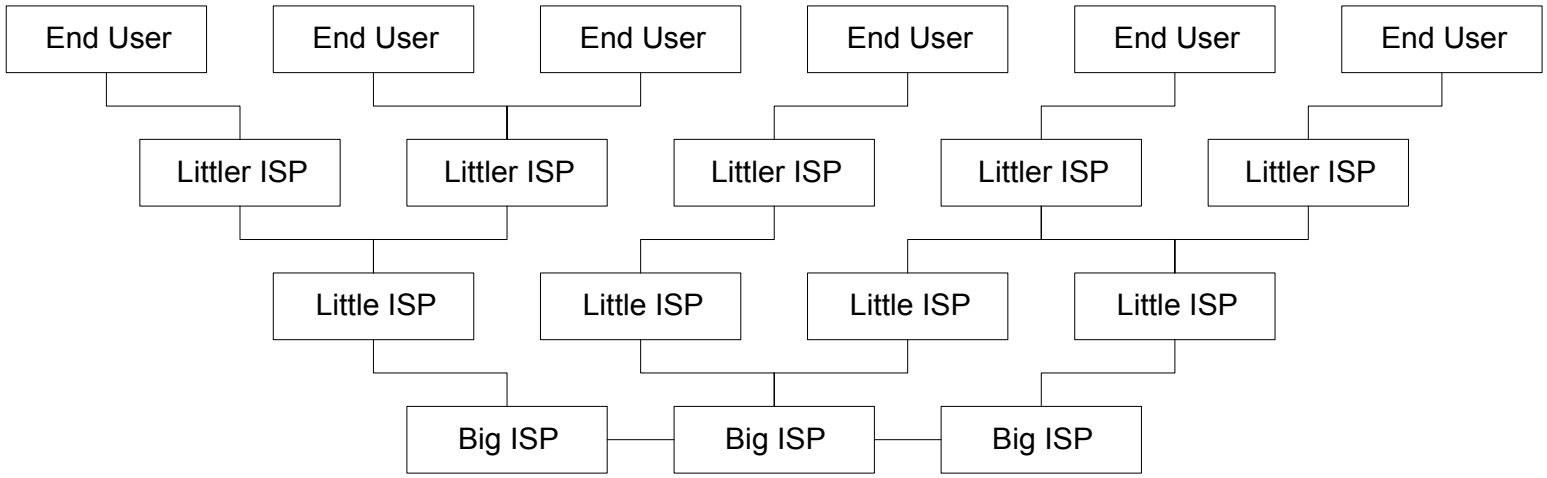

Structure of the Internet (cont'd)

- However, as you get closer in, this nice hierarchy breaks down
  - private deals, performance fixes, etc.
- Everybody provides services to the end users
- Service providers connect to multiple service providers
- End users connect to multiple service providers

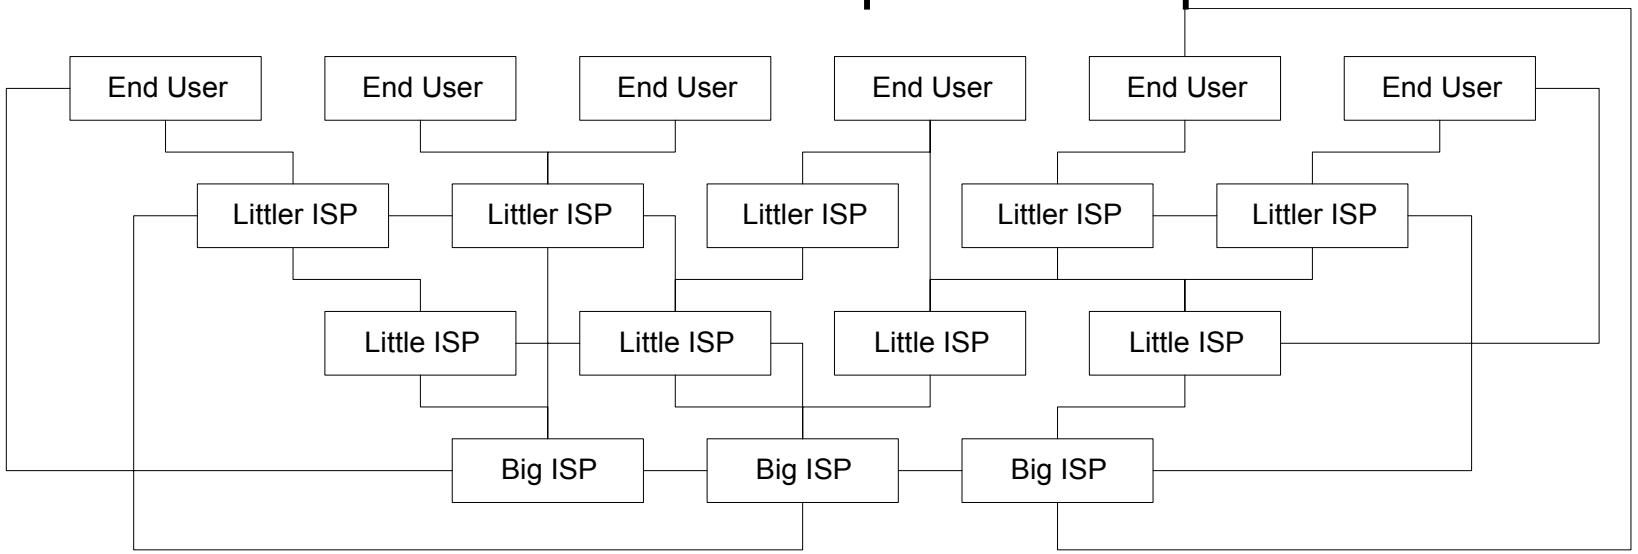

Connectivity

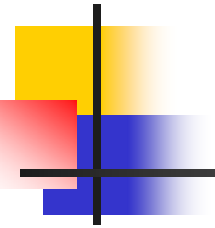

- In order to exchange traffic, an organization must connect with another organization and agree to pass traffic back and forth
  - Many possible conditions and clauses
  - Two types of relationships between the 2 organizations
    - Customer/Provider
    - Peer to peer
- In a given geographical region, organizations tend to want to mee at the same place
  - Easier to make deals with a bunch of people at the same place, than deals multiple places
    - Also saves on infrastructure costs
- Thus: the creation of a connection point
  - Known as *Internet Exchanges* (IXes)
    - NAPs, MAEs, etc.

Exchange Points

- Who connects to exchange points?
  - Almost always Internet Service Providers
    - The question is: who is a customer and who is a peer?
      - Customers pay, peers don't
  - Non-ISPs can probably place equipment at an exchange point
    - But ISPs want customers who pay, not peers, so the ISPs probably won't accept the customer's traffic
- Who runs exchange points?
  - Usually a neutral body, unbiased to any organization connecting there.
    - Tending towards commercial organizations

Exchange Points (cont'd)

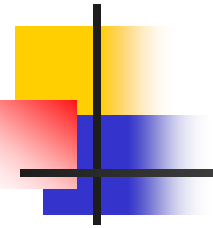

- These exchange points become Internet focal points
  - Major exchanges exist in:
    - The US: MAE-East, MAE-West, the NAPs, etc.
    - Europe: The D-GIX points, LINX, etc.
    - Asia and Pacific Rim: HKIX, NZIX, JPIX, NSPIXP-2, etc.
  - Generally, exchanges spontaneously generate where there is a high concentration of ISPs
  - Provide a way to keep local traffic local
    - Sometimes ISPs may not want this
- Large ISPs may not want to connect
  - IXes encourage smaller ISPs to want to "peer" with large ISP

Current Structure of the Internet

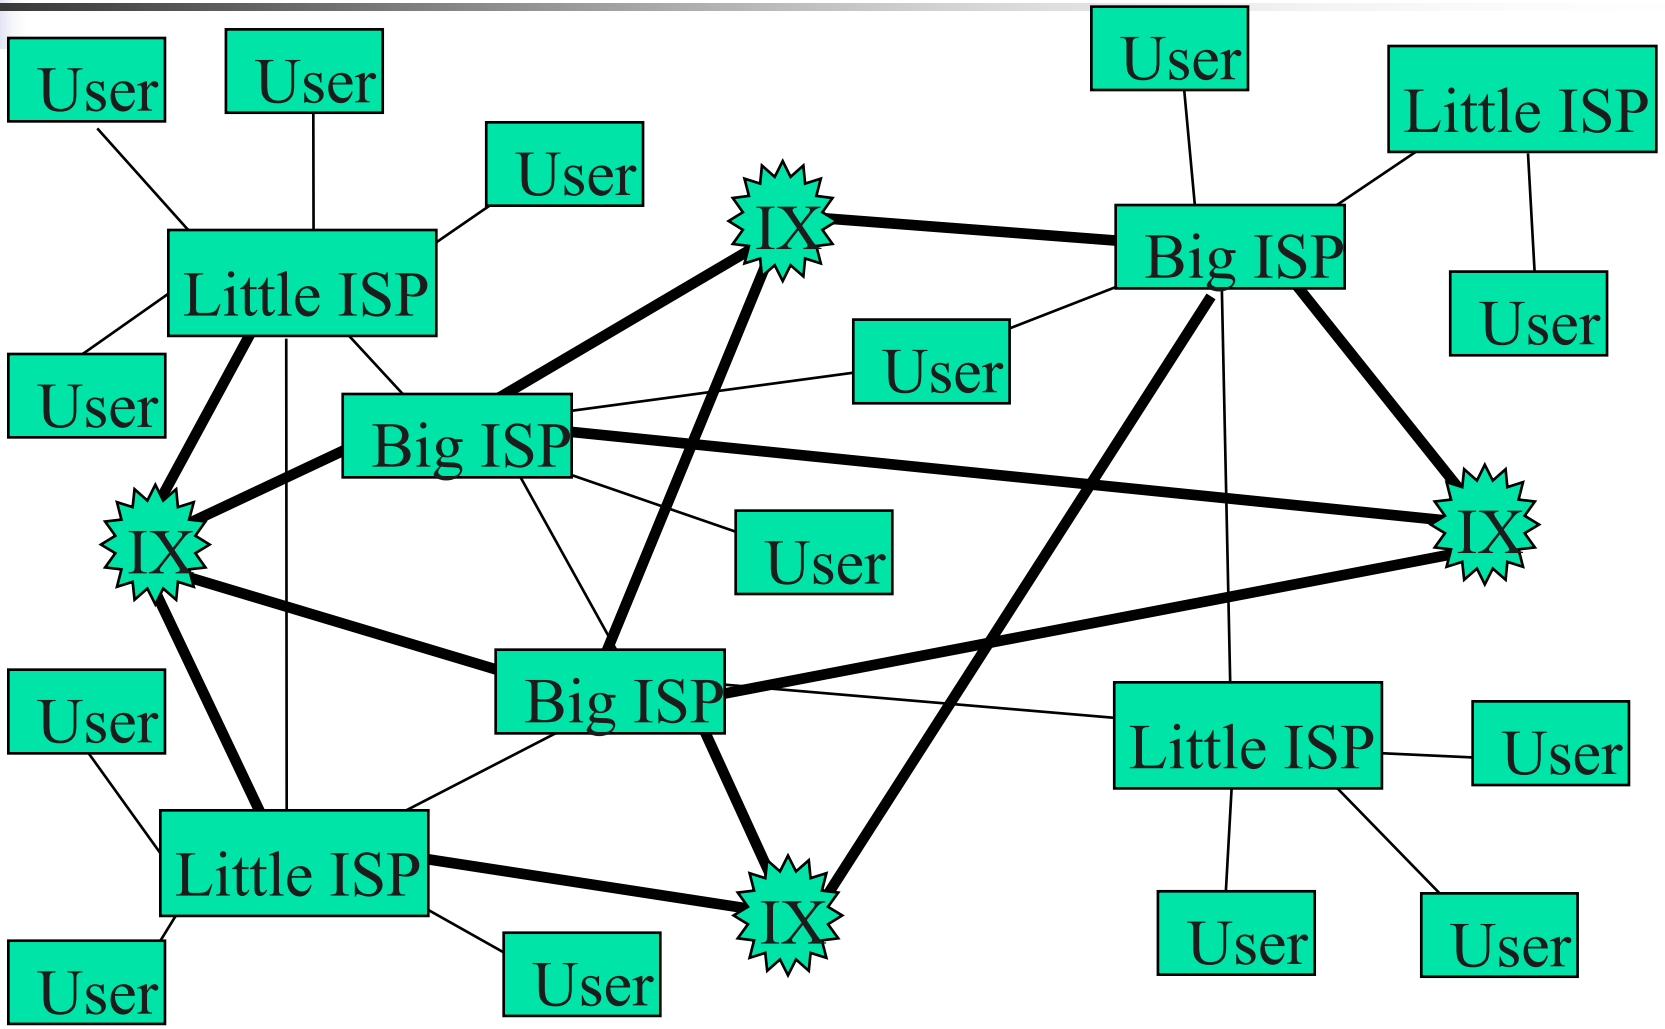

History of the Internet

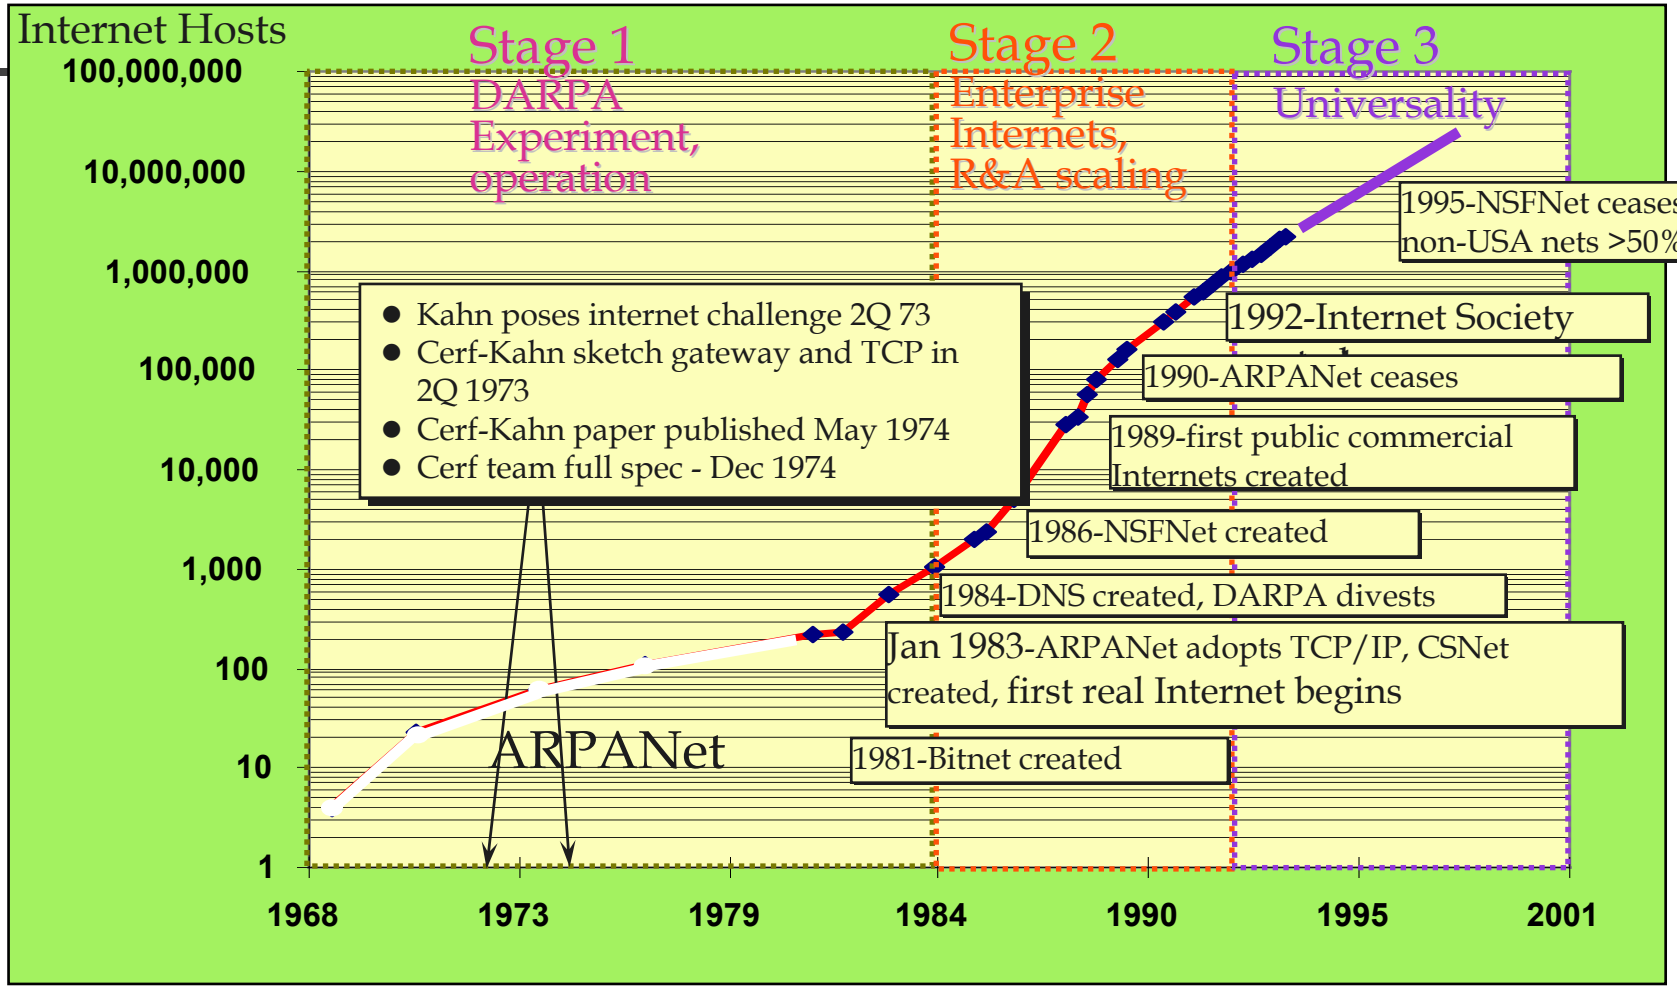

History of the Commercial Internet

- 1989 -- First commercial ISPs (UUNet and PSI) start taking customers
  - NSFNet disallowed commercial use
  - NSFNet operator soon forms commercial side-business (ANS CO+RE)
  - Many NSFNet regionals start commercial services
- European Internet hampered by TCP/IP vs. OSI wars
  - First commercial ISP (EUNet) around 1991
- AP Region Internet also OSIfied
  - First AP region commercial ISP (IIJ) in Japan 1992

Commercial Internet Exchange

- NSFNet provided only AUP compliant transit
  - ANS CO+RE provided commercial transit
    - Wanted to charge other commercials per byte
- Existing commercial ISPs were unhappy
  - Created the CIX
    - Explicitly disallowed usage based charging
- ANS became CIX member
  - CIX filtered transit between ANS and its resellers
    - But filters only applied to ANS
      - Legal wrangling occurred
- CIX forced to apply filters to all non-CIX members
  - No one noticed despite much concern

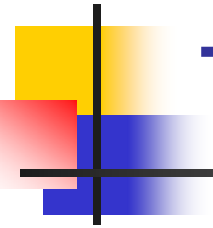

The Internet Today

- A self organizing collection of autonomous components
  - If you think the Internet can be controlled, see Chaos and Complexity Theory
    - However, you can control parts
- Commercial traffic far outweighs non-commercial
  - Electronic commerce will accelerate this trend
  - Non-commercial networks heading off on their own
- Exponential growth in nearly all areas
  - Number of users
  - Amount of traffic

TCP/IP Executive Summary

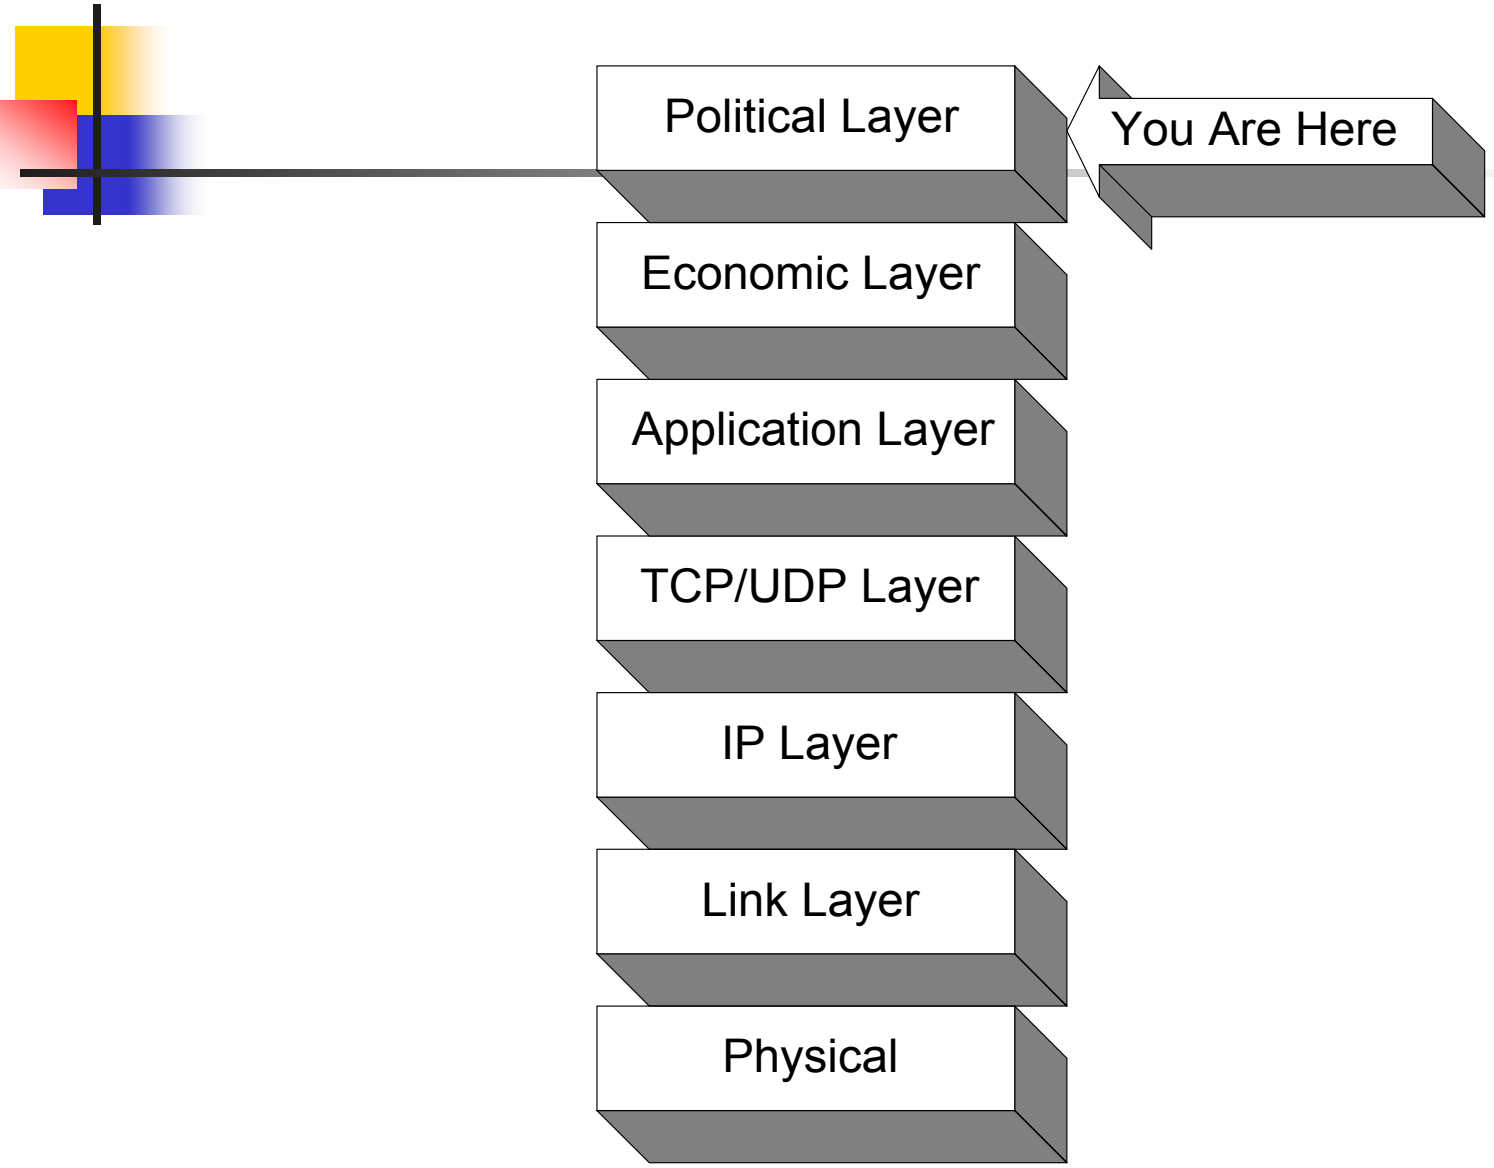

The Physical and Link Layers

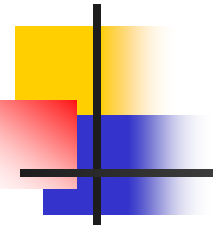

- TCP/IP runs on almost anything
  - Ethernet, FDDI, Cable TV, telephone wires, radio waves, barbed wire, carrier pigeon, etc.
- It runs better on some things than others
  - It is best not to shred the data (read: ATM)
  - It is best not to do error checking under IP (read: X.25)
  - It is best not to asynchronously switch (read: Frame Relay)
- It's all a question of overheads
  - low prices can make even "icky" things appealing
  - or when you don't have a choice...

The Internet Protocol

- IP is an abstraction of the underlying hardware
  - Ethernet, serial, barbed wire all look the same from above
  - The problem of dealing with multiple hardware types hidden
- IP is unreliable
  - throw it to the hardware and hope
- IP handles fragmentation and re assembly
  - But you don't want to fragment
- Want more? See RFC 791

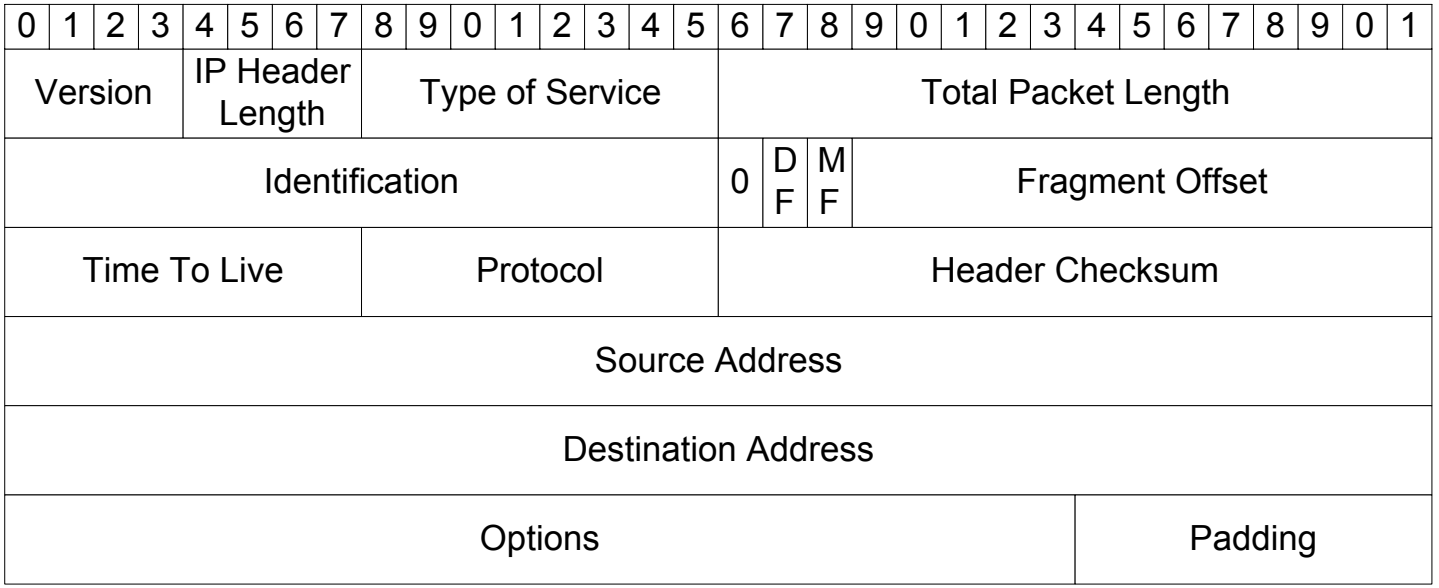

The User Datagram Protocol

- UDP provides an unreliable datagram facility to the user
  - No retransmission, duplicate detection or flow control
  - Optional corruption detection
  - Extends IP to allow multiplexing
- Want more? See RFC 768

- About as minimal as you can get
  - Gives the application protocol developer complete freedom
  - But developer usually has to re-implement TCP features
- Mostly used for DNS and network file system protocols (e.g., Sun's NFS)

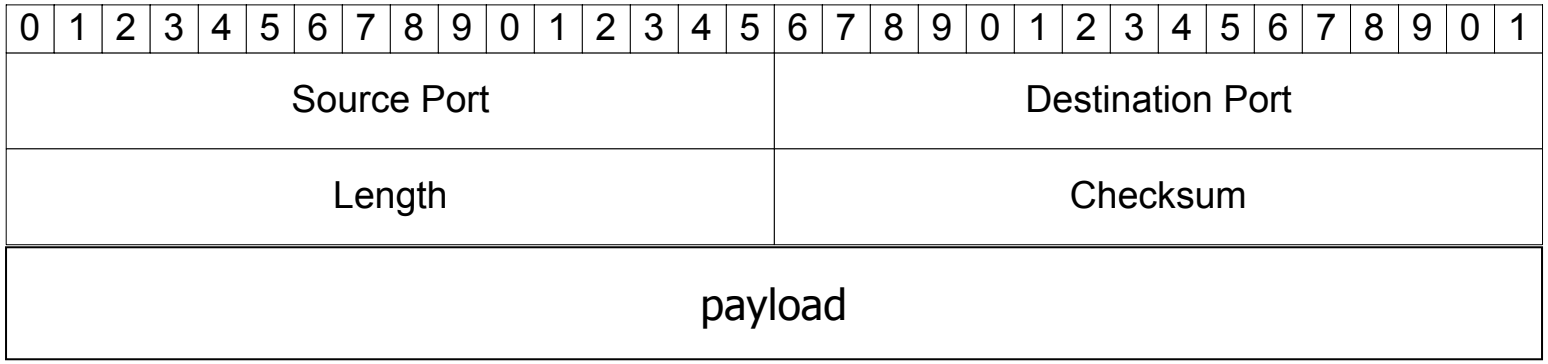

The Transmission Control Protoco

- TCP provides a reliable connection oriented stream of data
  - No framing or concept of blocks of data -- just a stream of bytes
  - Concept of "urgent" data
- Want more? See RFC 793

- TCP handles:
  - connection establishment, termination and reset
  - Detection of data corruption an duplication
  - Retransmissions and flow contr
    - Uses "sliding window"

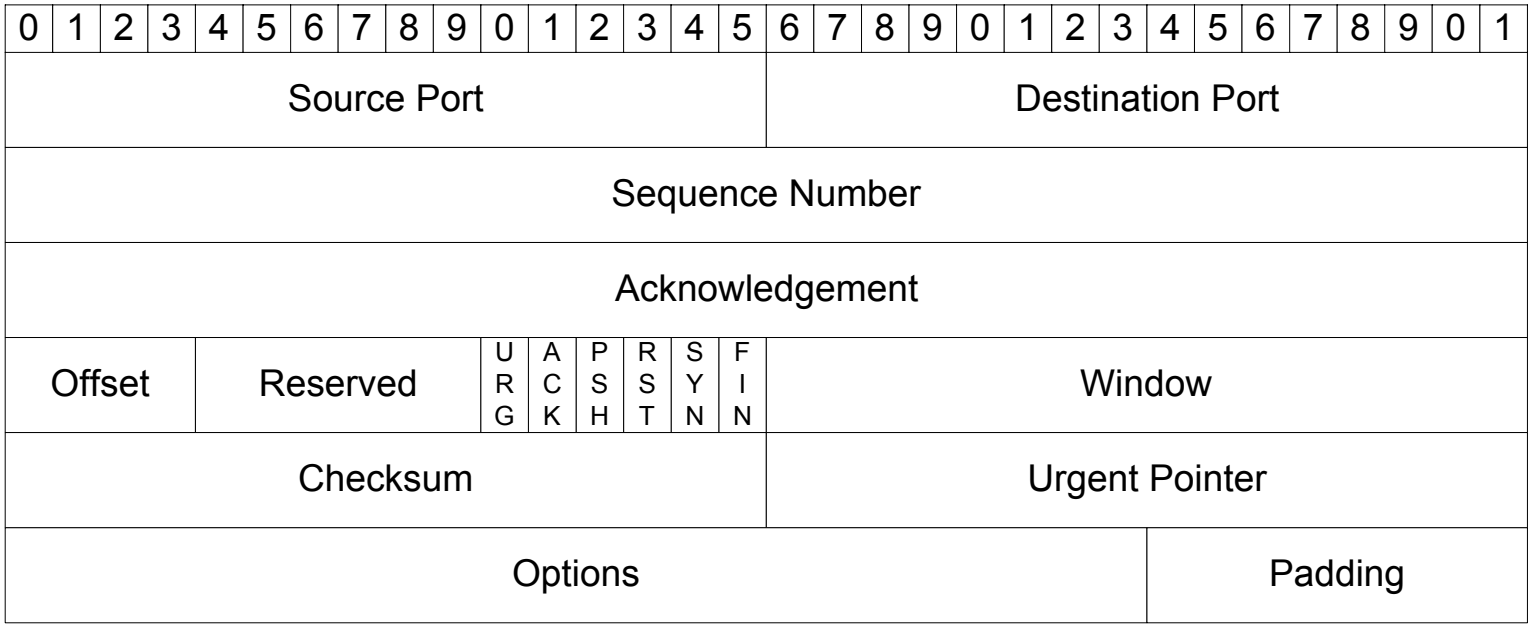

TCP vs. UDP

- TCP
  - Gives you exactly what you started with
  - Makes efficient use of the network, but has overhead
- UDP
  - Might (or might not) give you what you started with
  - Minimal possible overhead

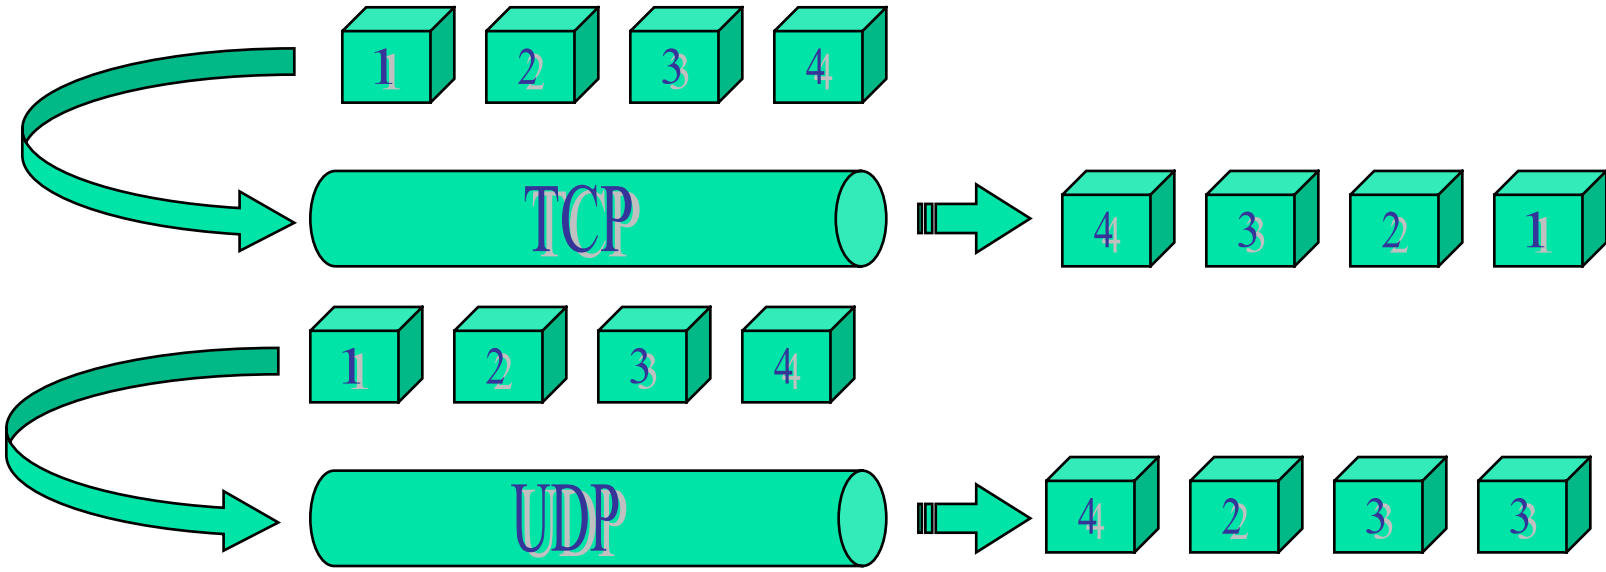

An Application Sampler

- TCP Applications
  - HTTP (WWW)
  - SMTP (Email)
  - NNTP (Electronic News)
  - Telnet (remote login)
  - FTP (remote file access)
  - Network File System
  - Whois
  - Rlogin/Rsh/Rdump, etc.
- UDP Applications
  - Domain Name System
  - Network File System
  - RADIUS (remote authentication)
  - Network Time Protocol
  - Trivial File Transfer Protocol

Administrative Support Groups

- Internet Hierarchy
- The IANA
- Numbers Support
  - Regional Registries
  - Local Registries
- Names Support
  - Top Level DNS Administrators
    - ISO 3166 TLD Administrators
    - Generic TLD Administrators
- Other support groups
  - Protocol development

Internet Hierarchy (Historical View)

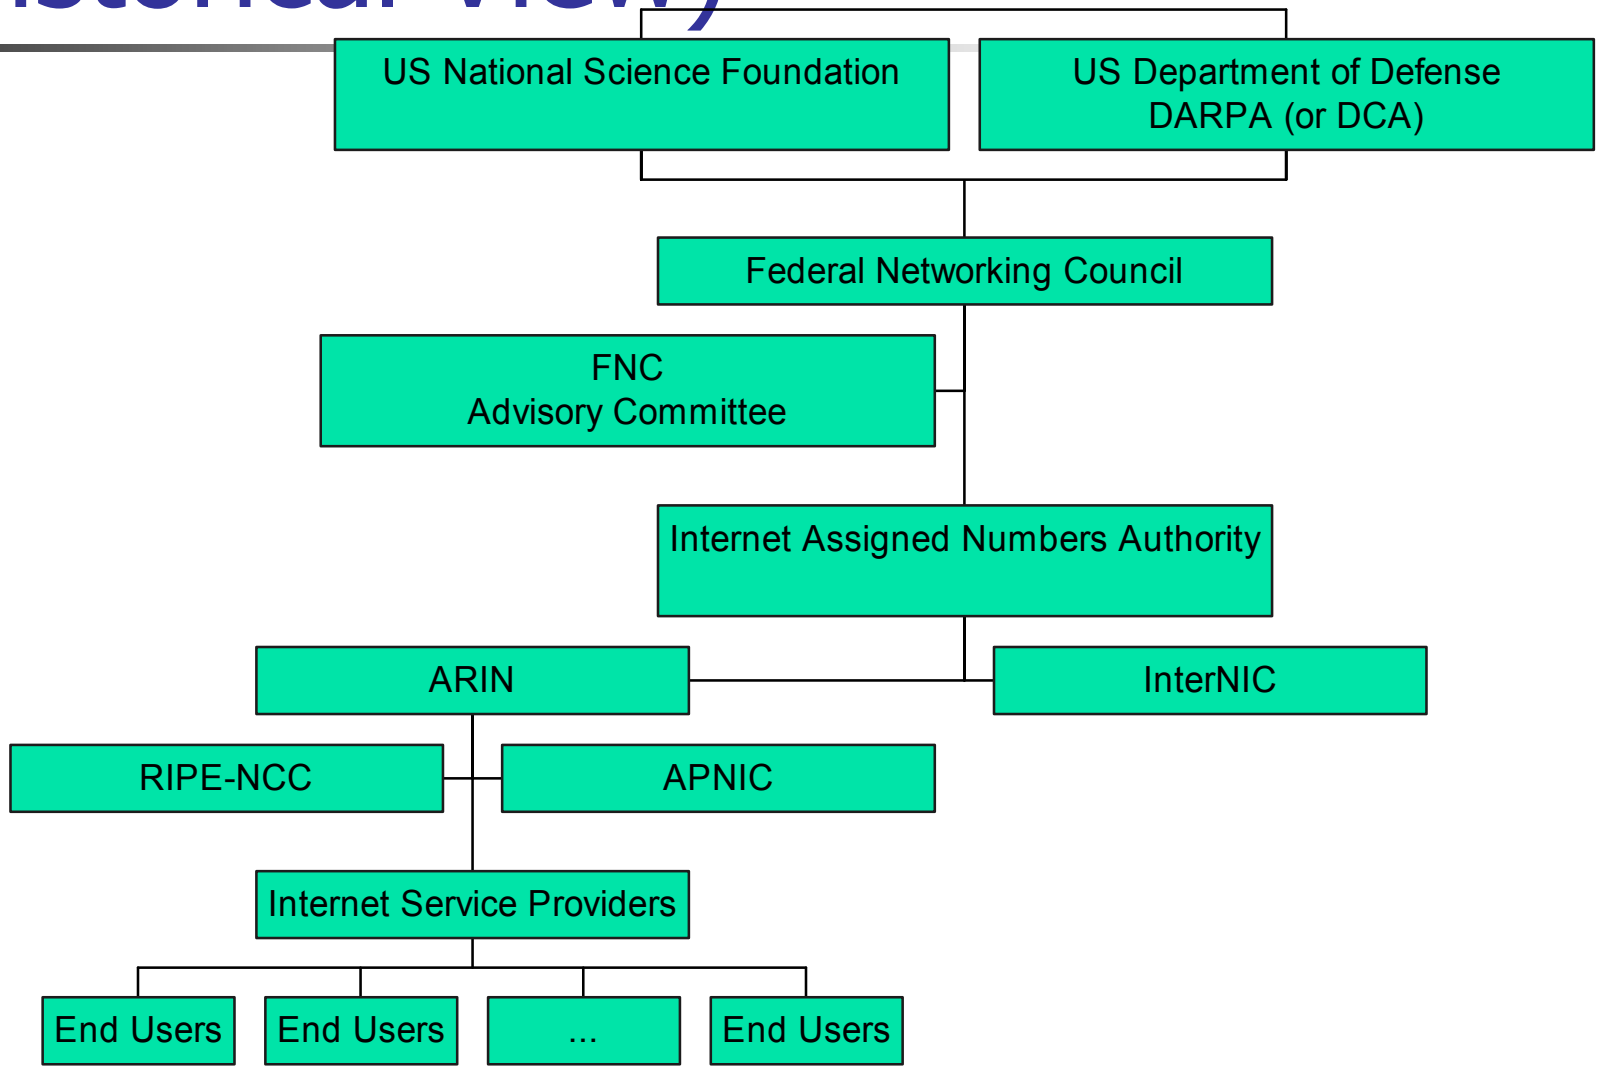

Internet Hierarchy (Politically Correct History)

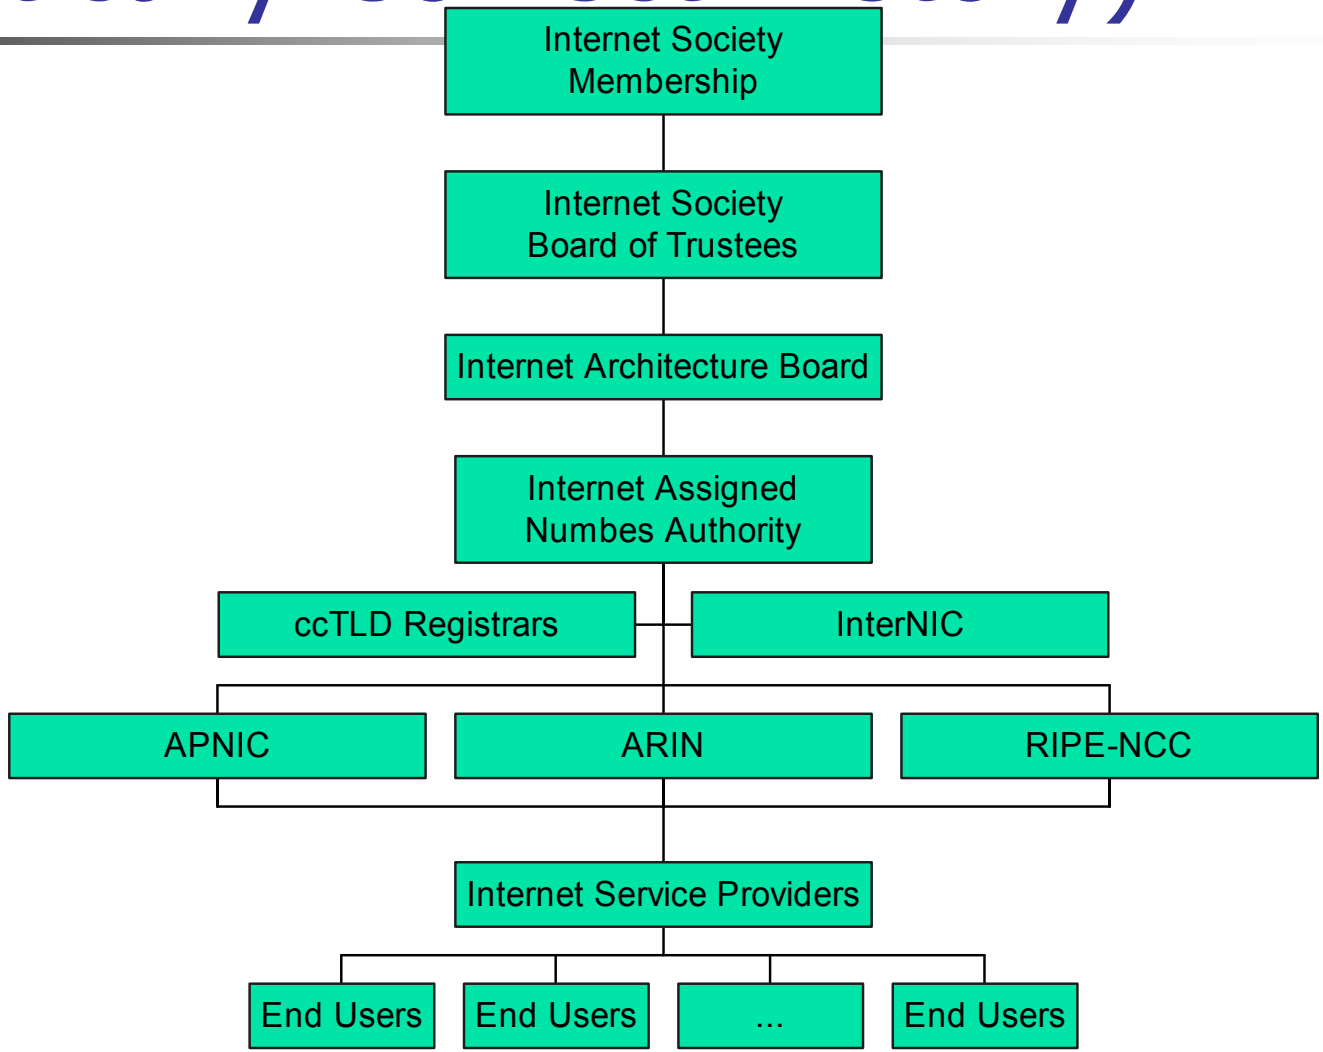

Internet Hierarchy (Bottom Up View)

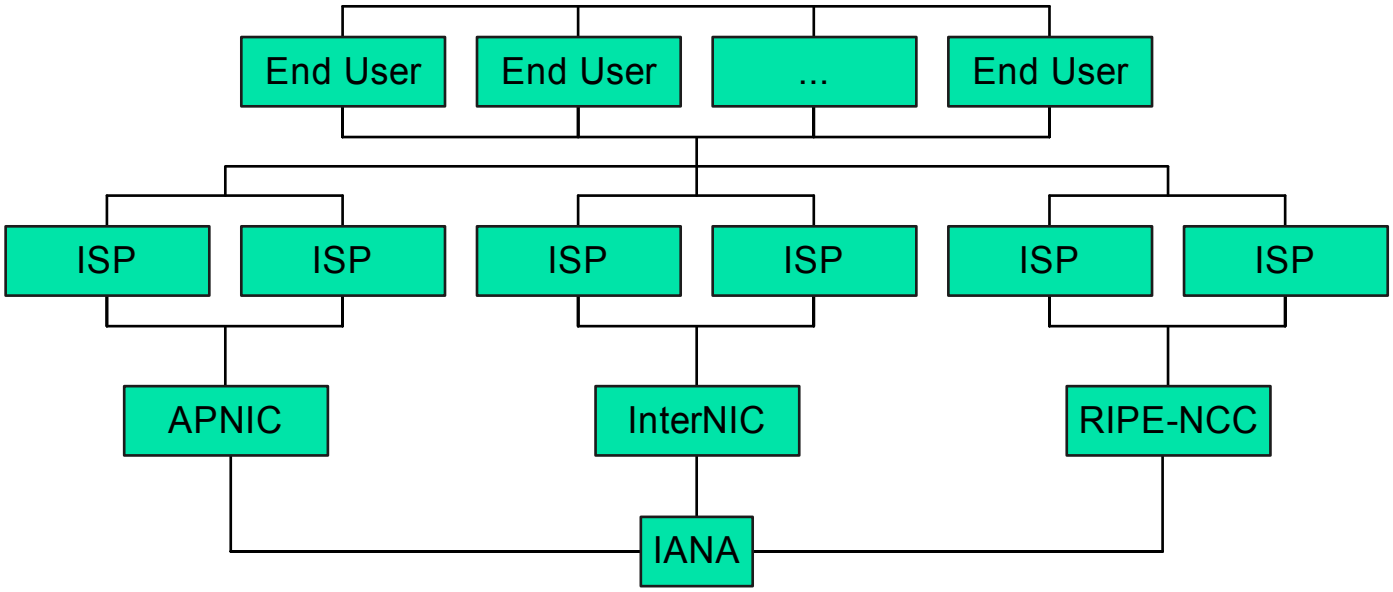

The Internet Assigned Numbers Authority

- The IANA was the parent of all regional registries and top level domain name administrators
  - In some context at least, the IANA could be said to "own" all administrative resources on the Internet
  - Handed out all globally unique numbers (IP addresses, protocol numbers, port numbers, object Ids, etc.)
- The IANA consisted of
  - Four (part-time) people
    - Located at USC ISI in Marina Del Rey, CA, USA
  - Had a yearly budget of roughly US $250,000
  - Delegated as much responsibility as possible

The ICANN phase

- US Government wants to ‘devolve’ the Internet’s administrative function to self-regulatory structures
- A bunch of DNS crazies step into the room
- ICANN attempts to ‘work out’ the DNS mess
- Address infrastructure management actually left to the RIRs
- US Government increasingly nervous about the whole deal
- Watch this space!

Regional Registry History

- First NIC at Stanford Research Institute (SRI-NIC)
  - Located in California (near Stanford University)
  - Funded by DOD DARPA
- SRI replaced by GSI in Washington DC area
  - Lowest bidder
    - Unpleasant transition
  - DOD DCA provided funding
- NSF issued InterNIC 5 year Cooperative Agreement
  - Cooperative agreement issued in 1992
  - AT&T, General Atomics, and Network Solutions, Inc. each awarded part of InterNIC

InterNIC History

- InterNIC consisted of 3 parts
  - Registration Services operated by NSI
  - Database and Directory Services operated by AT&T
  - Information services operated by General Atomics
- In 1995, NSF held a mid-term review
  - NSI commended, but told to stop relying on US Government for funding
  - AT&T told to do more
  - GA fired
- In 1997, InterNIC spins off Address allocator (ARIN)
  - NSI focuses on (now hugely profitable) domains

Address Registry Hierarchy

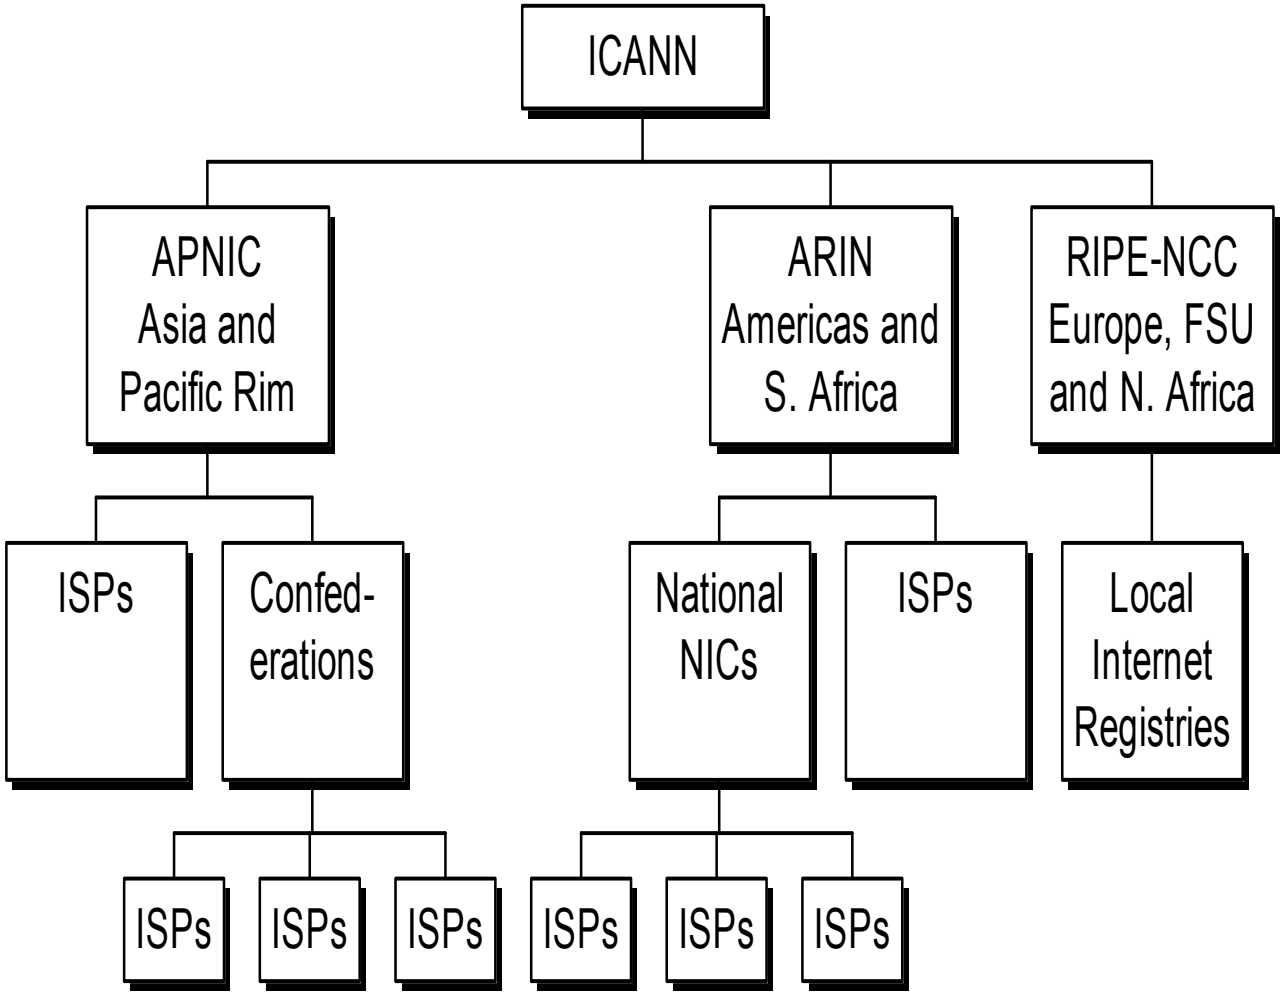

Regional Address Registries

- Registries allocate numbers
  - Internet addresses
    - (plus in-addr.arpa domains)
  - Autonomous System Numbers
- Currently three regional registries exist
  - APNIC, InterNIC, RIPE-NCC
    - All but InterNIC are self-funded
  - The IANA may create others as needs arise
    - AfriNIC and ALyCNIC have been discussed recently

Regional Registries (cont'd)

- Regional Registries are **NOT** regulatory bodies
  - They do not "license" ISPs
    - This is a national governmental issue
  - They are not the authority for who can or cannot connect to the Internet
    - Anyone can who is permitted by law in their country
  - They cannot control any organization
    - So complaining to them is pretty pointless
- Regional registries also don't make the rules
  - The Internet community does
    - Regional Registries simply implement policy, they don't invent it

Regional Registry Funding

- The Internet community specified the creation of regional registries
  - But didn't indicate how they would be funded
  - APNIC, ARIN and RIPE-NCC derive funding from their 'membership'

APNIC

- Started as an APCCIRN/APEPG Pilot Project in Sept., 1993, received address space from IANA in April, 1994, Incorporated in April 1996
- Uses a tiered membership fee with self-determined tiers
  - US $10,000/$5,000/$2,500 for Large/Medium/Small
    - Has about 400 members
  - Also has a non-member fee structure
- Located in Brisbane, Australia
- More info: see http://www.apnic.net

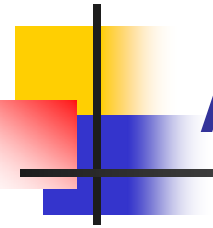

ARIN

- Created out of InterNIC in 1997
- Costs recovered via membership and allocation fees
  - Membership fee: US $1000/year
  - Allocation fees depend on the amount of space allocated the previous year ($20,000, $10,000, $5,000, or $2,500 depending on last year's allocations)
- Has a staff of about 20
- Based in Reston, Virginia US
- More info: see http://www.arin.net

RIPE-NCC

- Created in 1990 as the IP networking special interest group of RARE, a EU funded group working to deploy OSI networks in Europe
- Membership based organization using a self-determined tiered fee structure
  - Membership fees originally (about) US $10,000/ $5,000/$2,500, but changed to US $5000/$3500/$3000
    - Has about 600 members
- Has a staff of about 30
- Based in Amsterdam, The Netherlands
- More info: see http://www.ripe.net

Local Internet Registries

- Regional Registries delegate authority to "Local Internet Registries" to allocate resources
  - Usually Internet Service providers
  - Sometime confederations of service providers (e.g., national NICs)
- Local Internet Registries sub-delegate to customers
  - Either end users or other ISPs
- Each Local Internet Registry may have its own rules, but all must follow the rules of their parent registry

Top Level Domains

? delegates all TLDs

Policies for delegation described in RFC 1591

InterNIC currently administers (and charges for) gTLDs and .ARPA

ISO 3166 TLDs delegated to organizations within the countries

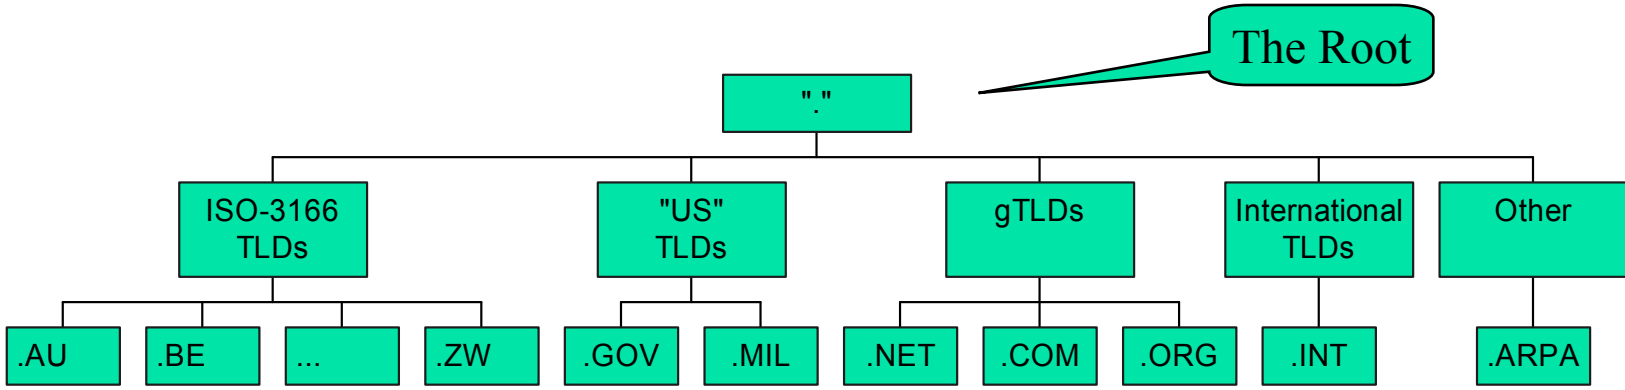

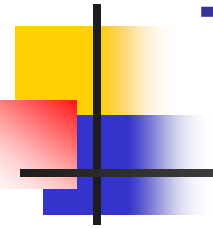

TLD Administrators

- Top Level Domain Administrators are assigned by ?
  - See RFC 1591 for history
- Each TLD has its own rules and restrictions on what names can be delegated

Development Groups

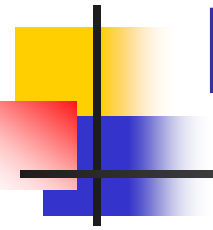

- Internet Engineering Task Force (IETF)
  - Protocol development and standardization
  - Also dabbles in operational and user support issues
- Has evolved from a meeting of 15 people back in the late ‘70s to over 1000 people meeting three times a year.
- Meetings held outside the US every third meeting (theoretically)
  - Next meeting: Pittsburgh, US, August
- More info: http://www.ietf.org

IETF Hierarchy

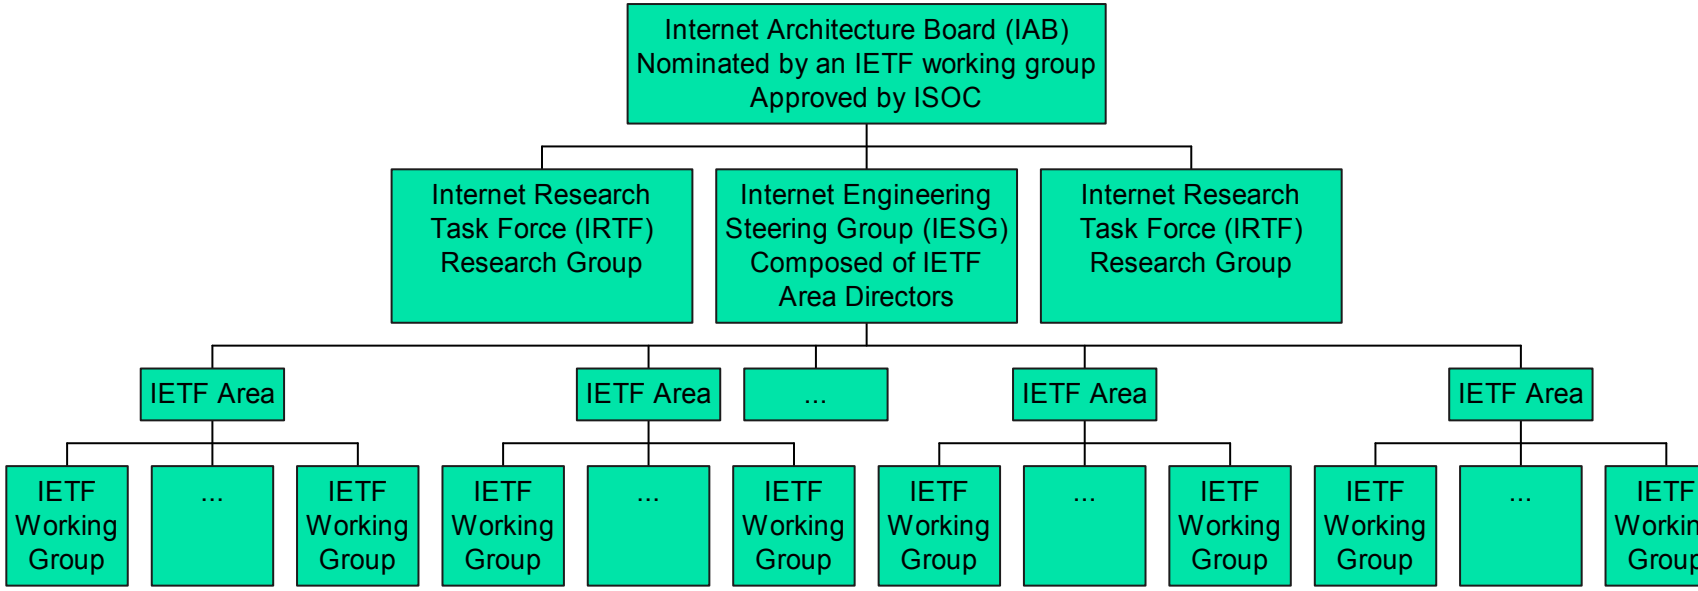

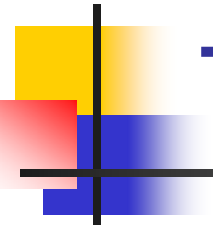

The Internet Society

- Around 1991, it was realized the IETF had no legal protection and as a standards making body, the IAB, IESG, and working group chairs could be sued by anyone unhappy with developing standards
- In addition, there was no organized group which could promote the Internet, act as a clearinghouse for information, etc.
- In 1992, the Internet Society was formed

Operations Groups

- North American Network Operations Group (NANOG)
  - Evolved out of an NSFNet technical working group composed of technical staff of NSFNet regional networks
  - Currently, open membership but targeted at the North American network operators
  - Partially funded by US NSF, coordinated by Merit (the organization that is running the US Routing Arbiter project)
- More information:
  - http://www.nanog.org

Operations Groups (cont'd)

- European Operators Forum (EOF)
  - Created out of the RIPE Community
  - Open membership, but focusing on the European Internet operators community
  - Currently unfunded, operating as a working group of RIPE
- More information:
  - contact ncc@ripe.net

Operations Groups (cont'd)

- Asia Pacific Operators Forum (APOPS)
  - Formed at an Asia Pacific Rim Internet Conference on Operational Technologies (APRICOT)
  - Hosted by APNIC
    - Only exists as an email list right now
      - To subscribe: apops-request@apnic.net
  - Open membership, targeted at the AP region
- More info: send mail to hostmaster@apnic.net

- AFNOG – underway for Africa

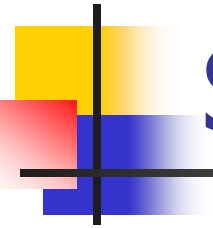

Summary

- The Internet is a collection of privately operated networks that
  - Use TCP/IP as a common protocol
  - agree to cooperate to exchange traffic
    - "Enlightened self-interest" allows the Internet to function
- Various bodies around the world help coordinate Internet
  - But none are essential for its operation

Where to get more information

- TCP/IP Internals
  - The RFCs (available from an archive near you)
  - "Internetworking with TCP/IP, vols. 1, 2, 3" by Douglas E. Comer, published by Prentice Hall, ISBN 0-13-468505-9
- History of the Internet:
  - www.isoc.org web pages

Where to get more information

- Internet Structure
  - NANOG, EOF, APRICOT, AFNOG meetings (the Internet structure changes so rapidly, any documented form will be out of date almost immediately)
- Internet Organizations
  - RIPE-NCC (www.ripe.net), APNIC (www.apnic.net), ARIN (www.arin.net)
  - Internet Society (www.isoc.org)