The Internet: Philosophy & Technology

Scott Bradner
Harvard University
Feb 4 2002

The Network That Was There

- **the** Phone Net from The Phone Company (TPC)
- circuit-based
  - assumed simple & predictable interconnections between hosts
  - assumed requirement for QoS
  - assumption of being carrier-provided
  - voice-oriented

Traditional Phone Network

- circuits & "smart network"
- connection-oriented
- hard state in network devices
- fragile
- central resource control
- socialist? "for the good of all"
- applications in network
  - e.g., phone switch
  - end-to-end touch-tone signaling was a mistake
- predictable development path
  - extended development cycle

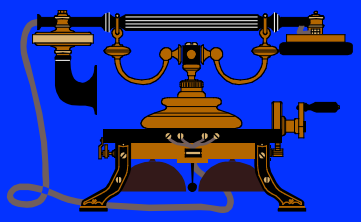

What Was Wrong With That?

- nothing, if you just wanted to talk
- nothing, if you just wanted to talk to Joe
- nothing, if you just wanted one service
- nothing, if you thought innovation had stopped
- nothing, if you thought that AT&T innovated
- nothing, if you wanted your data service provided to the wall by a carrier

  (ISDN is the answer, what was your question?)

So, Lets Make (Not Build) our own

- multiple unrelated efforts (early to mid 1960's)
  - packet switching theory: (Kleinrock) 1961
  - day dreaming: (Licklider's Galactic Network) 1962
  - make use of remote expensive computers: (Roberts) 1964
  - survivable infrastructure for voice and data: (Baron) 1964
- ARPANET (late 1960's)
  - Roberts ARPANET paper 1967
  - RFP for "Interface Message Processor" won by BBN 1968
  - four ARPANET hosts by 1969
  - public demo and email in 1972

Fundamental Goal of Internet Protocols

- multiplexed utilization of **existing** networks
  - different administrative boundaries
  - multiplexing via packets
  - networks interconnected with packet switches
    - called gateways (now called routers)
  - note: international in scope
- did not want to build a new global network
  - too expensive
  - too limiting

Internet Protocols Design Philosophy

- ordered set of 2nd-level goals
  - 1/ **survivability** in the face of failure
  - 2/ support **multiple types** of communications service
  - 3/ accommodate a **variety** of network types
  - 4/ permit **distributed management** of resources
  - 5/ **cost effective**
  - 6/ **low effort** to attach a host
  - 7/ **account** for use of resources
- note: no performance (QoS) or security goals
- not all goals have been met
  - management & accounting functions are limited

Packets!

- basic decision: use packets not circuits
  - Kleinrock's work showed packet switching to be a more efficient switching method
- packet (a.k.a. datagram)

Dest Addr	Src Addr	payload

  - self contained
  - handled independently of preceding or following packets
  - contains destination and source **internetwork** address
  - **may** contain processing hints (e.g. QoS tag)
  - **no delivery guarantees**
    - net may drop, duplicate, or deliver out of order
    - reliability (where needed) is done at higher levels
  - **no authentication of packet header**

Routing

- sub parts of the network are connected together by computers that forward packets toward destination
  - these computers are called "**routers**"
- routers use destination address in packet to make forwarding decision
- routers exchange reachability information with other routers to build tables of "next hops" toward specific local networks
  - exchange of reachability information done with "**routing protocol**"

A Quote

"*the lesson of the Internet is that **efficiency is not the primary consideration**. Ability to grow and adapt to changing requirements is the primary consideration. This makes simplicity and uniformity very precious indeed.*"

Bob Braden

End-to-End Argument

- 1981 paper by Saltzer, Reed & Clark
- "smart networks" do not help
  - adding functions into network can be redundant since actual function is **end-to-end**
    - e.g. encryption, data reliability
  - also harder to change to support new technology
    - also see Lampson ***Hints for Computer System Design***
- e2e argument projected to mean
  - no per-session knowledge or state in the network
    - but some "soft-state" (auto refreshed) may be OK
  - network should be transparent to end-to-end applications

Internet

- packets & e2e
- soft state in network devices
- resilient
- competitive resource control
- capitalist? "individual initiative"
  - but too much selfishness hurts all
  - must play by the same rules - but no enforcement
    - **the tragedy of the commons**
- applications in hosts at edges (end-to-end)
  - and in 3rd party servers anywhere on the net
- hard to predict developments
  - chaos at the rate of "Internet time"

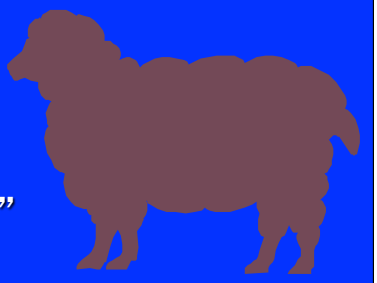

Smart vs. Stupid Networks

- phone network technology: self-named "Intelligent Network" (IN)
  - many network-based services
    - admission control, number translation, accounting, ...
- Isenberg's ***Rise of the Stupid Network*** compared phone network's "Intelligent Network" to Internet
  - Isenberg's basic messages:
    - network (i.e. carrier) -based services slow to change
    - voice is not all there is
    - carrier gets in the way
    - just "deliver the bits" works

But!!

- a "stupid network" is a commodity service
  - the price of a commodity service is driven by the stupidest vendor
- hard to make money delivering commodity services
- new network infrastructure is very expensive
  - fiber optic cables (with installation) & hardware
- access rights can also be very expensive
  - e.g. wireless spectrum licenses
- carriers need something else to make money
  - common dream is that services or content will save the day
    - may be a false dream (other than porno)

But!! (2)

- packets w/o circuits cause problems
  - can not do guaranteed QoS
    - can not control path packets take
    - can not reserve capacity for application
  - security control harder
    - do not have logical "wire" back to source
  - management harder
    - can not see data patterns on the network
    - finding non-catastrophic failures harder
  - service provider interconnections harder
    - no clean interface for problems
- lack of useful formal tools to describe performance

!QoS

Conceptualization Problem

- fundamental disconnect between "Internet" and "phone" people "bell-heads vs. net-heads"
- by their definition the Internet can not work

  and must be fixed - they will rescue us

"You can not build corporate network out of TCP/IP."

IBM circa 1992

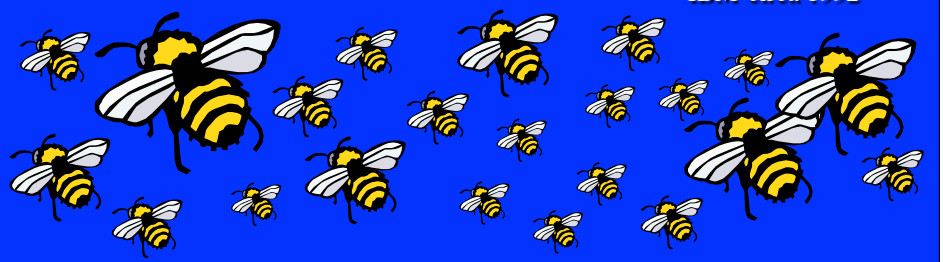

More Conceptualization Problems

- service provided by 3rd parties - not only by carriers

  different from phone world

- a quote from an IETF telephony mailing list

  Hi Roy,

  I still don't understand why it is a "users" choice where the "services" are executed - I would have thought that this would be networks choice

IP as a Common Bearer Service

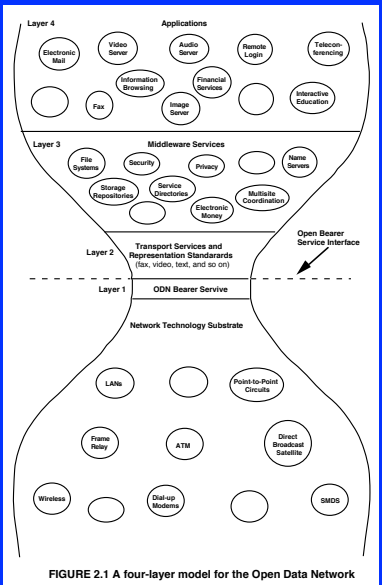

FIGURE 2.1 A four-layer model for the Open Data Network

From: Realizing the Information Future

Trust-Free Environment

- original Internet architecture assumed a trustworthy environment
- no longer the case
  - mistrust net itself (eavesdropping, reliability etc)
  - mistrust that you are talking to the right end point
    - e.g., proxy, redirect, spoofing (MAC & IP address)
  - unsolicited correspondence (spam)
  - anonymity hard to get
  - mistrust own hardware and software
  - 3rd parties insist on being in the middle
    - filters, wiretapping, …

Current Internet Architecture

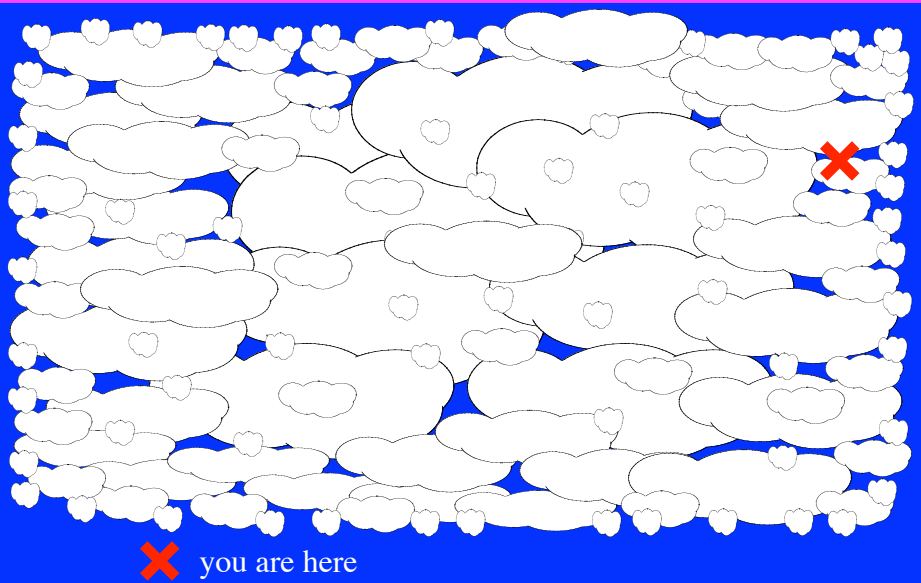

Numbers and Names

- nodes on IP networks have addresses
  - currently addresses are 32-bit values (IPv4)
    - total possible addresses: 4,294,967,295
    - written as 4 short numbers separated by periods
      - e.g., 128.103.60.212
  - IPv6 uses 128-bit addresses
    - total possible addresses:
    - 340,282,366,920,938,463,463,374,607,431,768,211,456
- half of IPv4 addresses have been assigned
  - address assignments are conservative these days
  - IPv6 developed to deal with shortage

Uniqueness of Addresses

- addresses have to be unique within scope
  - scope = connected network - e.g., the Internet
  - since address used to direct packet to destination
- can have address translators (NAT) if not unique
  - but NATs hurt end-to-end model
- blocks of addresses assigned by regional IP address registries
  - each with a unique geographic scope
  - competition is not appropriate when trying to conserve a scarce resource

Names for Addresses

- addresses change and are hard to remember
  - addresses change when networks are reconfigured
- service can be provided from more than one computer
  - for load distribution and/or reliability
- started with centrally maintained table that people downloaded
  - but that quickly became too big to stay accurate
- Domain Name System (DNS) developed to allow distributed database for mapping

Domain Names

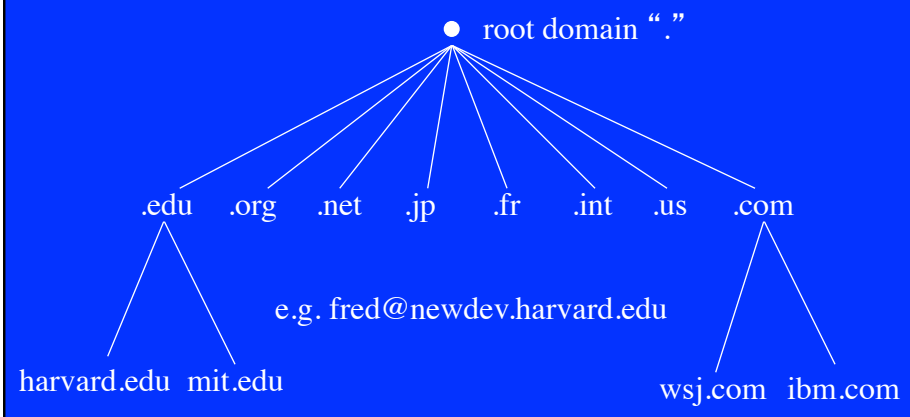

name servers for each domain
with database of next lower level entries

Uniqueness of Names

- single DNS tree required to ensure consistency
  - if >1 root then if you & I look something up we may get different responses if using different roots
- some proposals for >1 root
  - motivated by desire to not have single control point
  - but no technical way to ensure consistency

Standards

- a common (standard) transport is needed for interoperability
  - IP is the common bearer service for the Internet
- a common (standard) congestion control mechanism is needed to keep the net from collapsing
  - TCP & SCTP are the IETF congestion control protocols
- common application technology needed within each application for interoperability
  - e.g., email, www
  - counterproductive to prohibit alternates: innovation is good

Coordination

- uniqueness is a requirement in a number of things
  - addresses
  - names
  - protocol parameters
- unique things have to be coordinated
  - i.e., one authoritative database
- ICANN coordinates some Internet things: "IANA"
  - continuing work of Jon Postel
  - addresses & dns top-level domains
  - protocol parameters from the IETF

Limited Standardization

- IETF (and others) create "standards" for the Internet
- but use of the Internet not restricted to these "standards"
- can be an issue when a company refuses to open technology (or to support a standard)
  - e.g., instant messaging

Signing Stuff

digital signatures

Encryption

- symmetric (shared secret) system
  - same key used to encrypt and decrypt

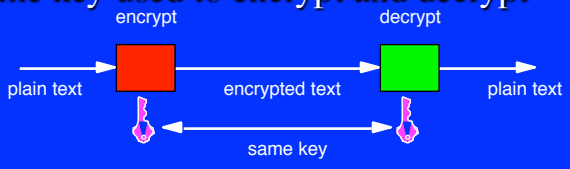

- asymmetric (public key) system
  - separate keys for encryption and decryption
  - data encrypted by one key can only be decrypted by the other

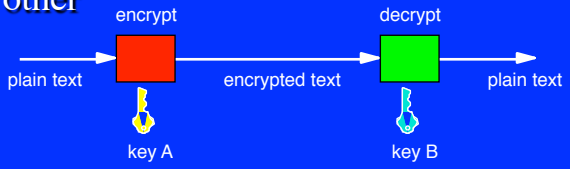

Digital Signature

- need method to be sure that message came from *A* and was not changed
  - use *Digital Signature*
  - appended to message before sending
- procedure for using a digital signature
  - ***A*** computes a one-way hash function of the contents of the message
  - ***A*** encrypts hash code with its private key the result is appended to the message
  - when it gets the message ***B*** computes the same hash function on the body of the message
  - ***B*** then decrypts the received hash code using ***A***'s public key
  - if the hash codes match, the message came from ***A*** and the contents were not altered in transmission

Digital Signature contd.

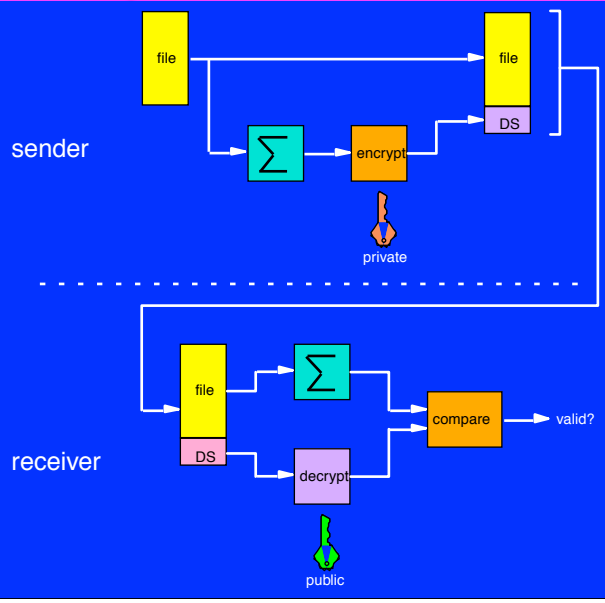

Digital Signatures

- data integrity
  - ensure that the data did not change since DS created
- data origin authentication
  - only person with knowledge of private key can create DS so I can be sure you created it
- non-repudiation of origin
  - different way to say data origin authentication
  - I can show that it must have been you who created DS
    - unless you can show that your private key was compromised

Public Keys

- I need to find out your public key to send you a secure message
- you need to find out my public key to authenticate a message from me
- need to get key in a secure, non-forgeable way

Certificates

- public key with digital signature(s)
  - can have more than one digital signatures
- "signed" by someone or some organization you trust - personal knowledge vs. certificate authority
- X.509 is ISO standard for certificates
  - x.509 v3 adds DNS name & loosens hierarchy requirements

Certificate Issues

- revoking certificates
  - certificate revocation list (CRL)
- expiration date & renewal process
- is a signed document "legal"?

Web-of-Trust Vs. PKI

- web-of-trust
  - you send me certificate signed by someone I know
- PKI
  - hierarchical infrastructure of certificate authorities
  - chain of trust

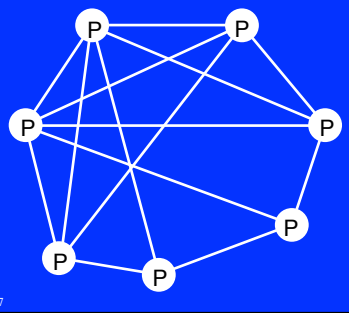

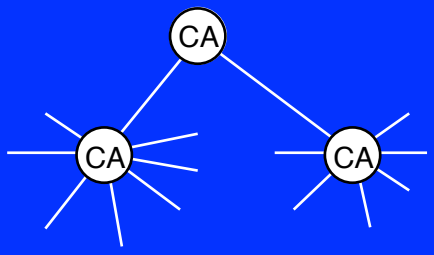

PKI Issues

- a PKI would be good except
  - need system that covers all relevent users
    - corporate-wide for corporate applications
    - world-wide for general Internet commerce
  - liability issues: what could CA be liable for?
  - privacy issues: identity assurance - how about anonymity?
  - jurisdictional relationships: what laws to follow?
  - local CA procedures: what identity assurance was used?
- will not happen soon