Rohit Khare

Information & Computer Science
4K Associates

What's in a name?

- ***Any problem in Computer Science can be solved by another layer of indirection***

  — David Wheeler

  (Chief EDSAC Programmer)

- We name objects in order to:
  - *Abstract away details of location, access, user interface*
  - *Interpose another layer of control, to allow relocation, e.g.*
- Naming is a *choice*
  - *To share a common name is to share trust in its meaning*

What's Internet Scale About, Anyway?

- This workshop series is dedicated to the proposition that successful Internet applications require *more than scalable algorithms*:
  - *They must scale across time* — *longevity*
  - *They must scale across space* — *latency*
  - *They must scale across organizations* — *liability*
  - *Economic, Political, and Social criteria are just as critical*
- Internet Scale is about more than large numbers...

1000 m

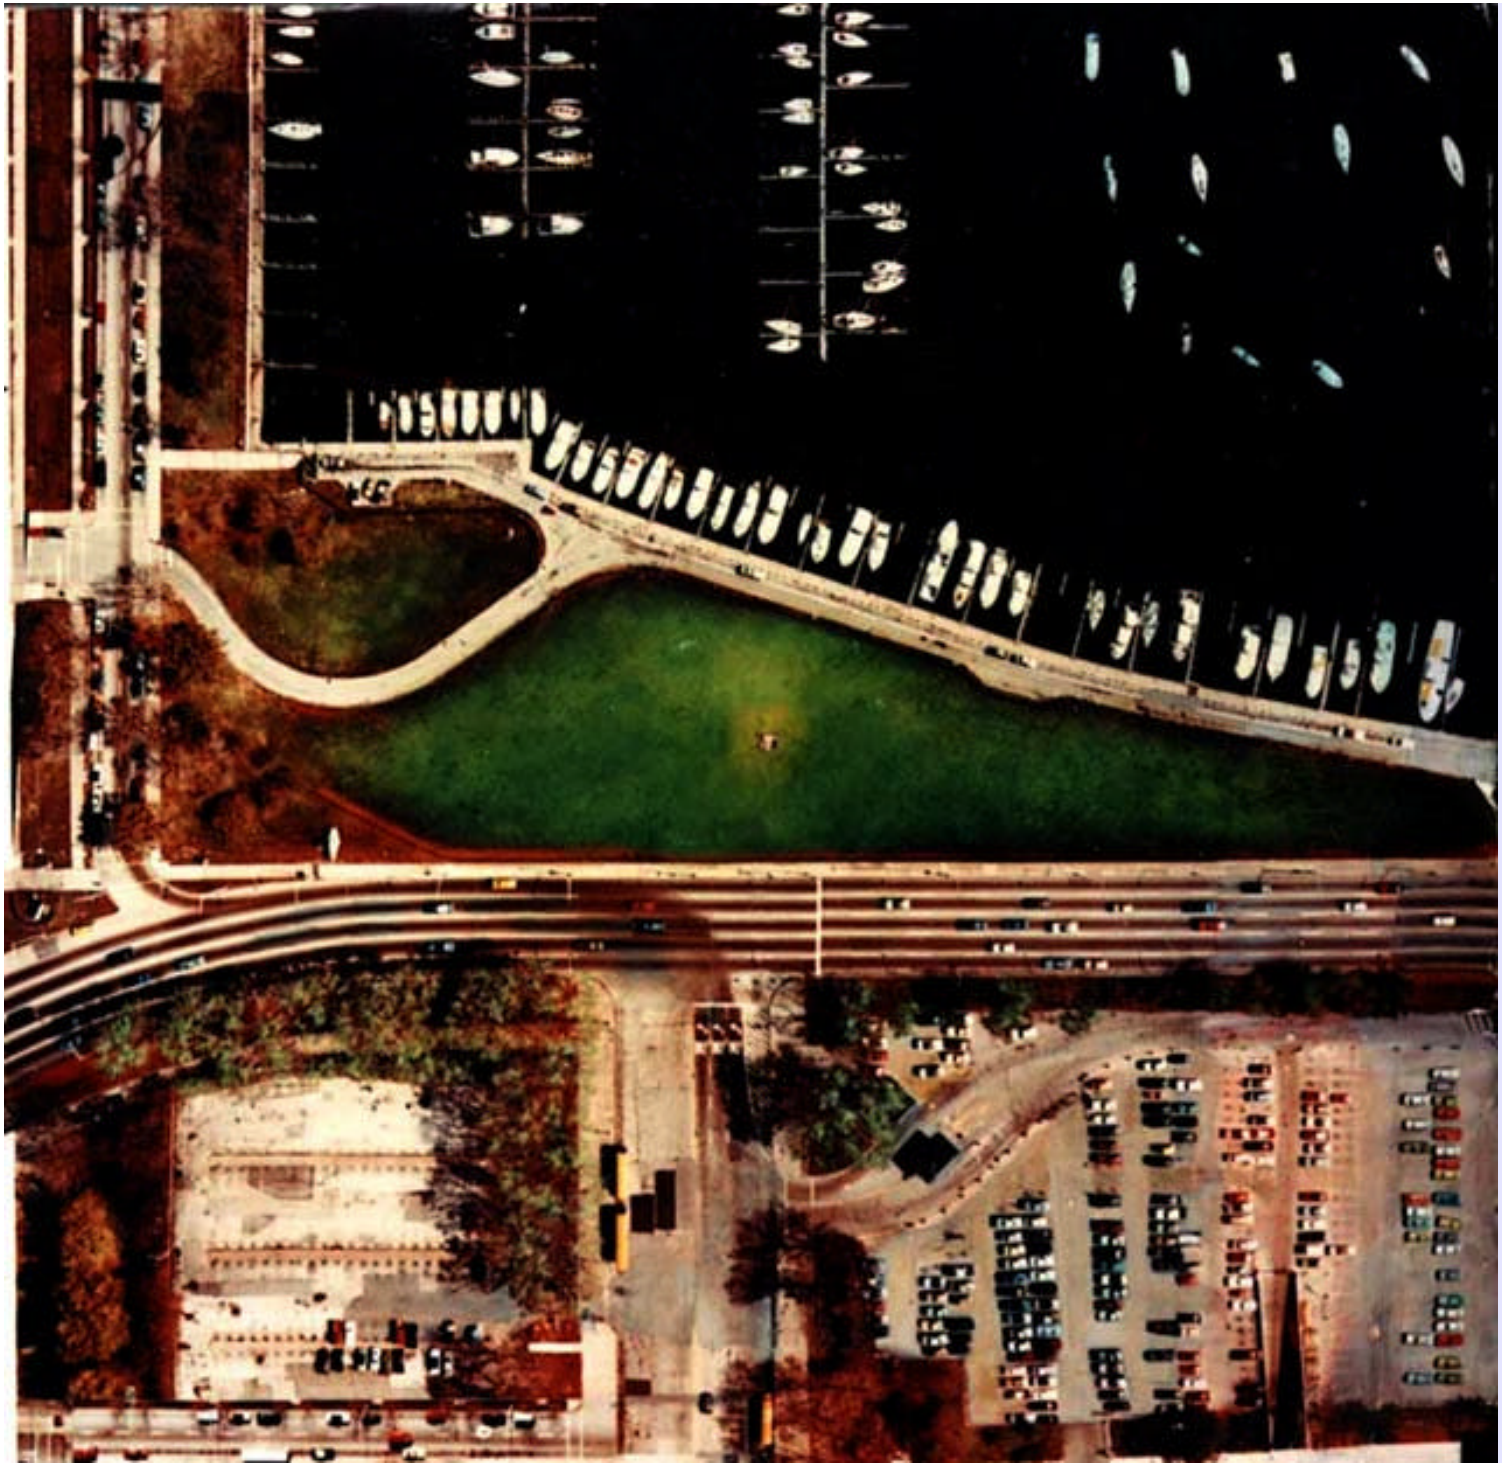

100 m

10 m

1 m

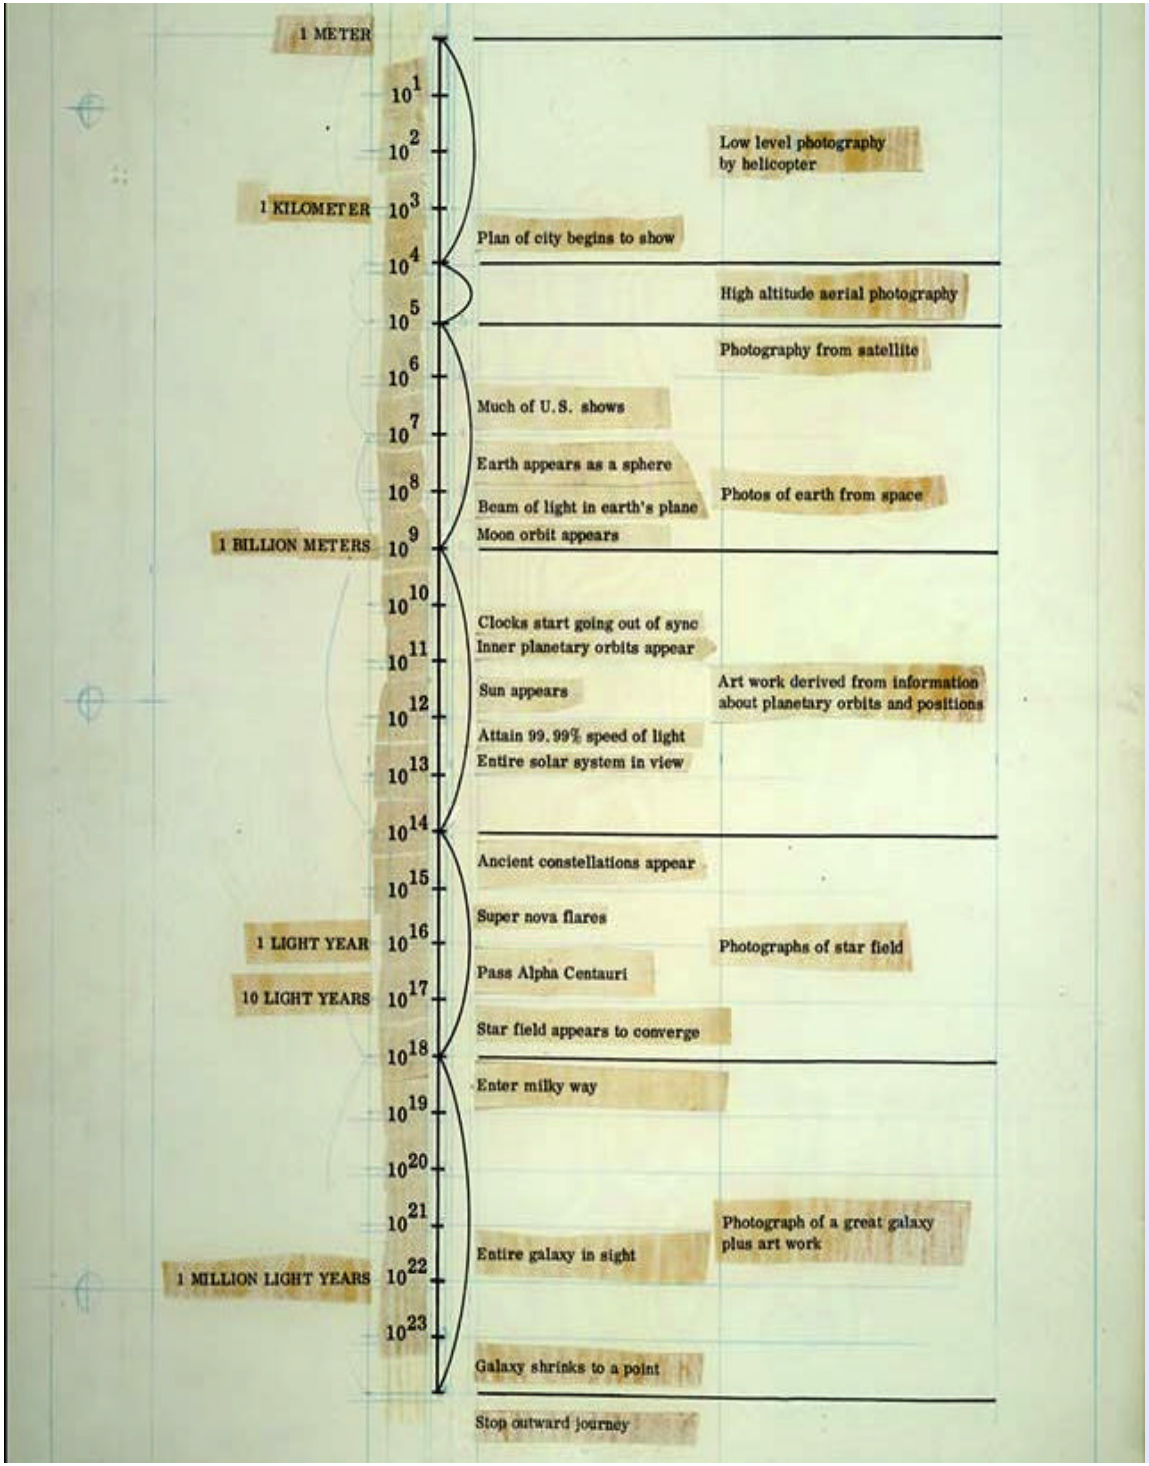

Powers of Ten

- *Powers of Ten* illustrates the different rules governing different scales of existence
  - *Meteorology*, *Biology*, *Chemistry*, *Quantum Mechanics…*
  - *Geology*, *Astronomy*, *Cosmology…*
- Yet the *same* rules apply, too!
  - *Physics is scale-invariant*
- *Let's try 'zooming in' on an Internet-scale name…*
- A film from the office of Charles & Ray Eames, 1977, running time 8:47

http://www.united.com/Itinerary/NQSS5A

- URI — Uniform Resource Identifier
- Resolved by — Web Browser
- Hierarchical — Left-to-Right
- Format by — IETF RFC 1630 (6/94), IETF RFC 2396 (8/98)
- Entries by — Server Administrator
- Internationalization — US-ASCII (UTF-8)
- Number — 10^{10} +
- Lifetime — 10^{1} - 10^{8} sec

- Replaced complex recipes for fetching network information with a single string
- Composed from four namespaces
  - *Scheme, domain, port, path*
- Can also have username, password

- *Browsers* resolve *URIs* to *Web Pages*

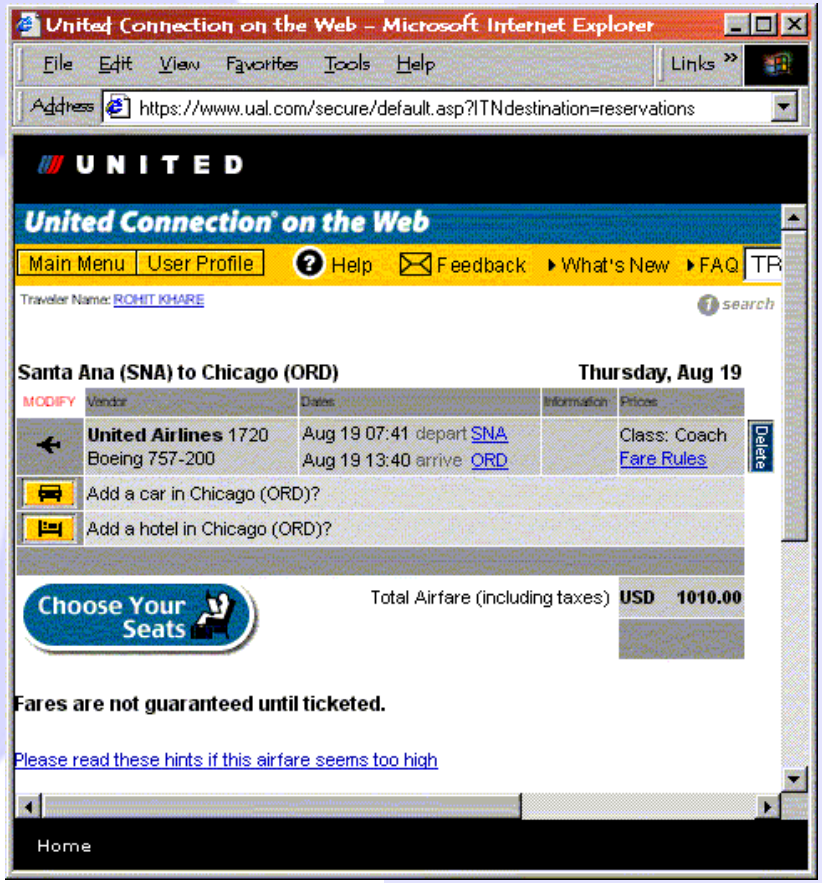

http://www.united.com/Itinerary/NQSS5A

- URI Scheme
- Resolved by — Web Browser
- Atomic — ASCII string
- Format by — IETF RFC 1738
- Entries by — IANA Registry
- Internationalization — none
- Number — 10^{1}+
- Lifetime — 10^{8}- 10^{9} sec

- Quickly identifies information-access system which can resolve the URI path
- Resolves to IANA assigned port numbers
  - *Not injective: HTTP and IPP both at 80*
- Can be an address, too, as with *data:*

- *Web Browsers* resolve *URI Schemes* into *connection protocols and ports*

Scheme	*Protocol*	*RFC*	*Port*
FTP	File Transfer Protocol	1738	21
Telnet	Interactive Sessions	1738	23
Gopher	The Gopher Protocol	1436	70
HTTP	Hypertext Transfer	2616	80
NNTP	Netnews Transfer	977	119
WAIS	Wide Area Inf. Svc	1625	210
Z39.50s	Z39.50 Session	ANSI	210
Mailto	Invoke mailer	821	25

- Https (443), snews(563), ftps(990)
  - *Single-bit security flag*

http://www.united.com/Itinerary/NQSS5A

- DNS — Domain Name
- Resolved by — DNS Protocol
- Hierarchical — Right-to-Left
- Format by — IETF RFC 883 (11/83)
- Entries by — ICANN-delegated registrar
- Internationalization — [A-Z][a-z][0-9]-
- Number — 10^8+ (63/254 char limit)
- Lifetime — 10^7 - 10^8 sec

- Composed of hierarchical namespaces
  - *com (ICANN), united (NSI), www (United)*
- Uniqueness requirement forces political solutions: United Van Lines or Air Lines?
  - *Actually, neither: this domain is disputed*
- Resolved by an 13-rooted planetary tree

- *DNS Resolvers* resolve *Hostnames* into *Internet Addresses*
- National TLDs ISO-3166 two-letter codes
  - *Iceland* — *this.is/keyword*
  - *Monteserrat* — *linux.versus.ms*
- Original TLD intentions:
  - *.gov* — *US Federal Gov't*
  - *.net* — *Network service providers*
  - *.int* — *International treaty orgs*
- 'Localhost' is a reserved name
- Reverse lookups
  - *213.21.195.128.arpa.in-addr*
- Competing global trademark registries
  - *RealNames, WHOIS (RFC 2345)*
- Urgent need to expand number of roots
  - *Allow several to masquerade as one*

http://128.192.21.213/Itinerary/NQSS5A

- IP — Internet Protocol Address
- Resolved by — TCP/IP Stack
- Hierarchical — Right-to-Left
- Format by — IETF RFC 791 (9/81)
- Entries by — IANA-delegated IP registry
- Internationalization — none
- Number — $2^{31} = 10^{10}+$
- Lifetime — 10^{1} - 10^{7} sec

- Composed of subnet and link numbers
  - *Class A, B, C and CIDR net mask prefixes*
  - *Topological consistency of net ranges*
- Also demuxed by 16-bit TCP port number
- Network Address Translators (NATs) fudge injectivity – address collisions poss.

- *TCP/IP Stacks* resolve *Internet Addresses* to *MAC (physical) Addresses* or next-hop *Internet Addresses*
- Regional IP numbering registries
  - *Europe - RIPE, Asia- APNIC, US- ARIN*
  - *Allocation Policy set by RFC 2050*
- Reserved ranges
  - *This network: 0.x.x.x*
  - *Broadcast: 255.255.x.x (Class B)*
  - *Multicast: 224.0.0.0 and up*
  - *Reserved loopback address: 127.0.0.1*
- Sample Netnumbers circa 1981 (RFC 790)

Internet Address	Name	Network
001.rrr.rrr.rrr	BBN-PR	BBN Packet Radio Network
002.rrr.rrr.rrr	SF-PR-1	SF Packet Radio Network
003.rrr.rrr.rrr	BBN-RCC	BBN RCC Network
004.rrr.rrr.rrr	SATNET	Atlantic Satellite Net
005.rrr.rrr.rrr	SILL-PR	Ft. Sill Packet Radio
007.rrr.rrr.rrr	CHAOS	MIT CHAOS Network

http://128.192.21.213/Itinerary/NQSS5A

90:ca:fe:de:ca:de

- MAC Media/Multiple Access Control
- Resolved by LAN Address Res. Protocol
- Hierarchical Org. Unique ID + device ID
- Format by IEEE 802.3 &c
- Entries by IEEE Registration Authority
- Internationalization none
- Number $2^{48} = 10^{14}$
- Lifetime 10^{8} - 10^{10} sec

- Maps onto individual link endpoints (network stations)
- Absolutely must be link-unique
  - *Analogous MACs for ATM, Token Ring*
- ARP uses a simple lookup table

- *ARP & RARP* resolve *Ethernet Addresses* to/from *Internet Addresses*
- Blocks of 4,096 are sold to Ethernet adapter manufacturers for $500, after a $1,250 “initiation” fee
- What portions of this space are reserved?

- Address Resolution Protocol, RFC 826, November 1982
  - *Announce own IP, request target IP’s MAC*
- Reverse ARP, RFC 903, June 1984
  - *Broadcast a request to get an IP address*

1-800-296-3892

- PPPAddress — Phone number
- Resolved by — Point-to-Point Prot. + modem
- Hierarchical — Left-to-right
- Format by — ITU E.164 (Bell, 1947)
- Entries by — N. Am. Numbering Plan
- Internationalization — country codes
- Number — 10^{10}
- Lifetime — 10^5 - 10^9 sec

- PPP Link driver itself operates over a phone circuit
- Phone numbers hierarchically assigned
  - *Nation, Area, Exchange, Subscriber*
- Absolutely must be world-unique
- Indexed by Yellow and White Pages

- *Telephones* resolve *phone numbers* into *circuits*
- Networks can be countries, too: Iridium satellite phone subscribers get +8816
- Phone numbers are represented in many common forms
  - *ITU form: +1-(626)-806-7574*
  - *DNS form: 4.7.5.7.6.0.8.6.2.6.1.tpc.int*
  - *URI form: phone://16268067574/*
    - Tel: , fax: and modem: proposals, too
  - *Geocoded: MAdison 437 (archaic)*
- Reserved portions
  - *555 testing & information services*
  - *800-855-xxx Teletype toll-free info (TDD)*
- Caller-ID (ANI) reveals source address
- Reverse lookup possible too

http://www.united.com/Itinerary/NQSS5A

- URL Pathname
- Resolved by — Web Server
- Hierarchical — Left-to-right
- Format by — IETF RFC 2396
- Entries by — Server Administrator
- Internationalization — US-ASCII (UTF-8)
- Number — 10^{10} +
- Lifetime — 10^{1} - 10^{8} sec

- Maps onto individual resources
  - *But representation on the wire may still depend on content language, media type, authorization, &c.*
- Must be server-unique; may be aliased
- No versioning by default; can rot

- *Web Servers* resolve *Pathnames* into *HTTP Representations* (replies)
- Becomes BASE for resolving relative URLs
  - *This resource identifier resolves to the HTML outline of a page that is completed with several subsidiary resources (graphics, sounds, style sheets)*
- Can be a collection resource (DAV)
  - *Supports enumeration, searching of directories*
- Can have properties (DAV)
  - *Such as Author, Words, Cost…*
    - Which come from yet other property namespaces...

http://www.united.com/Itinerary/NQSS5A
/usr/local/www/db/reservations.msql

- Filename
- Resolved by — Web Server
- Hierarchical — Left-to-right
- Format by — Operating System
- Entries by — Content Administrator
- Internationalization — ad-hoc
- Number — $1 - 10^6$
- Lifetime — $10^1 - 10^8$ sec

- Maps onto individual files *or processes*
- Server typically rewrites the URL by substituting root, user directory, extension
- Security and accounting controlled by OS, not necc. the web server's control

- *Web Servers* resolve *path components* into *filenames*
- *Operating Systems* resolve *filenames* into *inodes*
- *Disk Drivers* resolve *inode* into *track+sector addresses*
- *Disk Controllers* resolve *track+sector addresses* into *data blocks*

http://www.united.com/Itinerary/NQSS5A

- PNR — Passenger Name Record
- Resolved by — Airline Distribution System
- Atomic — Alphanumeric picture string
- Format by — Length, pattern vary by GDS
- Entries by — Airline
- Internationalization — none
- Number — 10^8 +
- Lifetime — 10^1 - 10^7 sec

- Maps onto individual reservations
  - *Every booking and confirmation is kept until flight time*
- Resolves to an IATA+Airline ticket number
  - *Permanent identifier lasts for years*
- Must be unique over itsdesign lifetime

- A *Reservation Database Process* resolves *PNR keys* into *Reservation records*
- Talking to a process, not a bag of bits.
- Interoperability standards are crucial for interline ticketing, but still fragmented by each GDS (Sabre, Apollo, Amadeus, etc).

- *GDS = Global Distribution System*

Anatomy of a URI

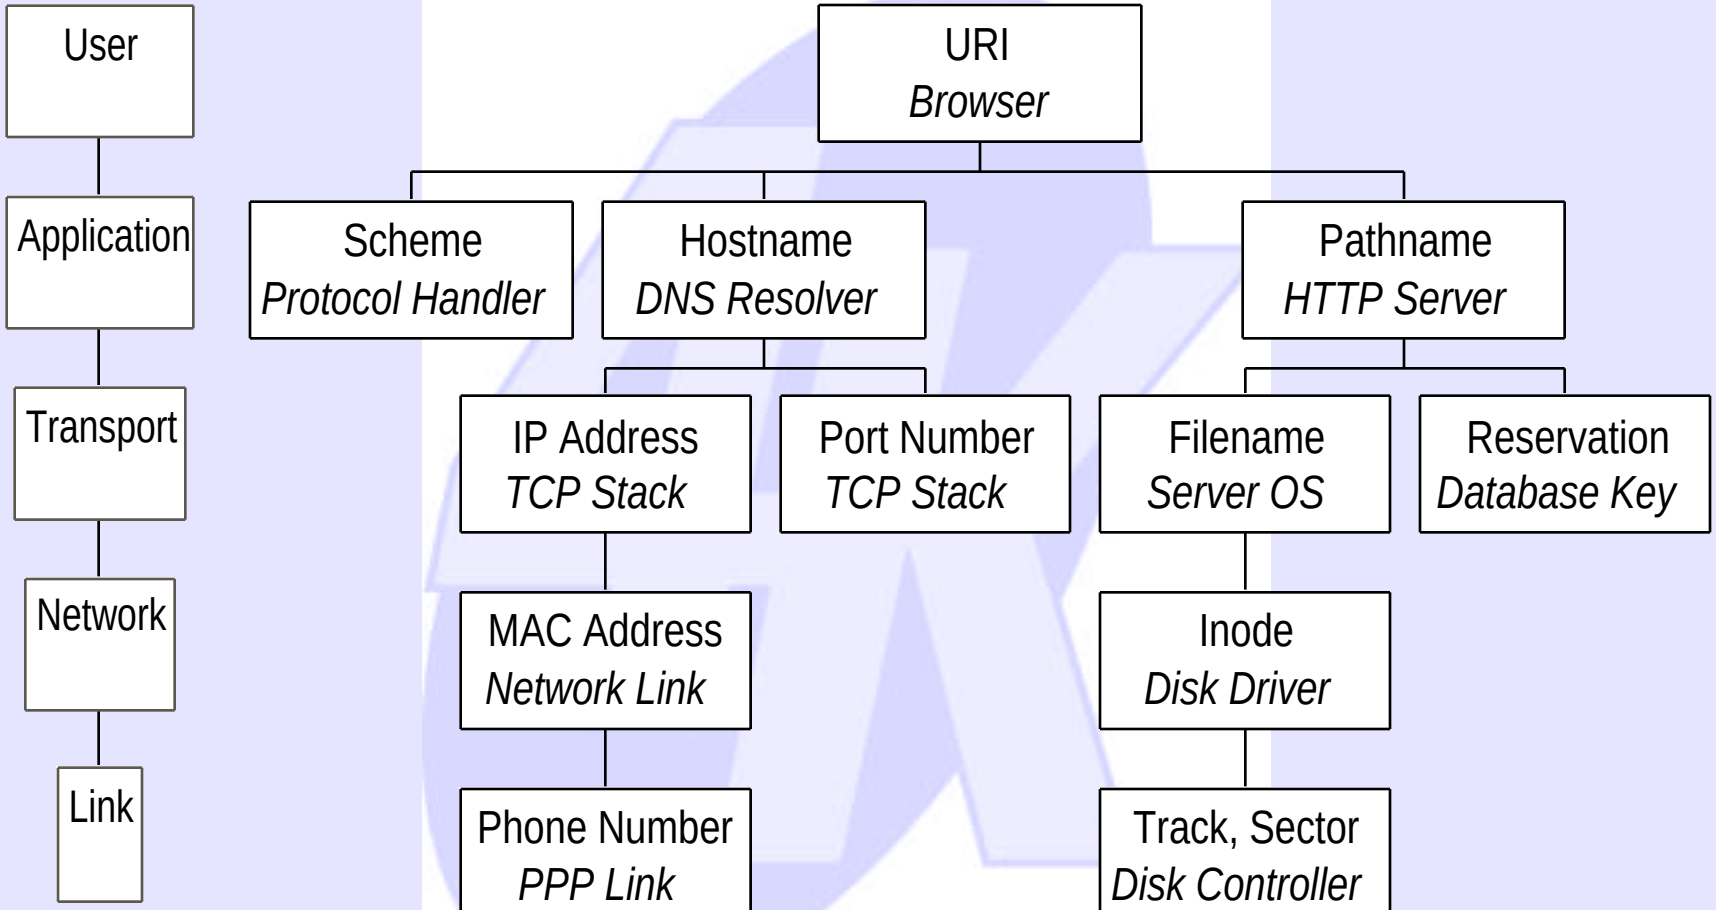

A URI resolves to an HTTP Message

- Zooming further in, an HTTP response message uses several more namespaces

- Method — Standards-track RFC
- Reply Code — RFC; newIANA Registry
- Content-Type — IANA Media-Type
- Content-Language — ISO language codes
- Character-set — IANA ref to ISO Charset
- ETag — Uniquely identify the resource
- PICS label — URI pointing to schema
  - *Here, 'Good Clean Fun' specifies its own suds/density/color ratings namespace*
- Digital Signature — Hash of resource
  - *Algorithm identifiers are URIs, too*
  - *But signing principals are another scale...*

- GET /PICS/DSig/Overview HTTP/1.1
  Host: www.w3.org
- HTTP/1.1 200 OK
  Date: Wed, 18 Aug 1999 21:22:41 GMT
  Server: Apache/1.3.6 (Unix) PHP/3.0.11
  Content-Location: Overview.html
  Vary: negotiate
  Last-Modified: Mon, 06 Apr 1998 20:24:44 GMT
  ETag: "2def30-a2e-35293a0c;35293a2f"
  Accept-Ranges: bytes
  Content-Length: 2606
  Content-Type: text/html; charset=iso-8859-1
- <!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 4.0 Transitional//EN"
  "http://www.w3.org/TR/REC-html40/loose.dtd">
  … <META http-equiv="PICS-Label" content=' (PICS-1.1 "http://www.gcf.org/v2.5"
  by "John Doe" labels for "http://www.w3.org/PICS/DSig/Overview"
  extension (optional "http://www.w3.org/TR/1998/REC-DSig-label/resinfo-1_0"
  ("http://www.w3.org/TR/1998/REC-DSig-label/MD5-1_0" "cdc43463463="
  "1997-02-05T08:15-0500"))
  extension (optional "http://www.w3.org/TR/1998/REC-DSig-label/sigblock-1_0"
  ("AttribInfo" ("http://www.w3.org/PICS/DSig/X509-1_0" "efe64685685=")
  ("http://www.w3.org/PICS/DSig/X509-1_0"
  "http://SomeCA/Certs/ByDN/CN=PeterLipp,O=TU-Graz,OU=IAIK")
  ("http://www.w3.org/PICS/DSig/pgpcert-1_0" "ghg86807807=")
  ("http://www.w3.org/PICS/DSig/pgpcert-1_0"
  "http://pgp.com/certstore/plipp@iaik.tu-graz.ac.at"))
  ("Signature" "http://www.w3.org/TR/1998/REC-DSig-label/RSA-MD5-1_0"
  ("byKey" (("N" "aba212412412=") ("E" "3jdg93fj")))
  ("on" "1996-12-02T22:20-0000") ("SigCrypto" "3j9fsaJ30SD=")))
  on "1994.11.05T08:15-0500"
  ratings (suds 0.5 density 0 color 1))'>

Principals resolve to People & Organizations

- Signing principals must use larger-scale names:
  - *Lifetime of name validity is >> duration of Web transaction*
  - *Social scope of name is >> than just immediate parties*
  - *Typically also used across multiple applications*
- Resolving any identifier onto the range of people and incorporated entities raises non-technical questions
  - *Privacy – is the resolver function known to all? Breakable?*
  - *Trust – such identities are invariably intended to bind legally*
  - *Economic – injectivity creates scarcity (one-to-one map)*
  - *Politics –surjection could be compelled (universal IDs)*

X.500 Directory Hierarchy

- Distinguished Name DN
  - *Common Name* *cn*
  - *Address* *street*
  - *Locality / Region* *l*
  - *State / Province* *st*
  - *Organizational Unit* *ou*
  - *Organization* *o*
  - *Country* *c*
- cn=Rohit Khare, o=4K Associates, c=US
- cn=Rohit Khare, ou=Information and Computer Science, l=Irvine o=University of California, st=CA, c=US
- Took 10 years, CCITT vs ISO friction, too
- *X.509 Certificates & Revocation Lists* resolve *DNs* into *public keys*
- Each component of a DN can be a Certification Authority (CA)
- Yields a pyramid-shaped trust structure, with increasingly liable, larger-scope organizations delegating central authority
- E.g. all https servers must buy certificates from a small number of roots, such as Verisign
- Role/authorization relegated to ‘extended attribute’ fields

PGP Identity Network

- Pretty Good Privacy, by contrast, allows any ASCII string to represent a keyholder
  - *Typically, eMail Address(es)*
- In the beginning, there is the self-signed key
- After verifying key 'fingerprints' offline, correspondents can also sign your key
- Names are imported into your keyring *only when signed by 'trusted' correspondents*
- SDSI works similarly
- Great, if everyone is known to each other… spontaneous messaging requires a bootstrap
- Brian LaMacchia's *PGP Keyserver* is a centralized cache of people's signed public keys
- PGP tools can interactively query it, attempting to construct a chain of trusted 'introducers'
- Thus, we have a radically decentralized namespace –'Rohit' is in the eye of the beholder – but implemented centrally...

XML Namespaces

- Suppose we zoom further into our itinerary web page:
  - *<B> Total: <FARE currency='usd' basis='R'>$6010</FARE>*
- Their XML element for distinguishing fare amounts is an addition to the HTML tag namespace
- XML Namespaces essentially turns tags into URIs:
  - *<HEAD xmlns:u='http://united.com/schemas/fares'>…*
    <u:FARE u:currency='usd' u:basis='R'> $6010 </FARE>
- *But how to compare United's fares to another's?*
  - *XML namespaces are a nifty Internet-scale solution*

IScale Properties of XML Namespaces

- Binding an ontology (vocabulary) to a URI allows communities of different scales to share semantics:
  - *Over time, it could be ratified to http://iata.int/fareschema*
  - *In restricted beta-testing of advanced features, it could be delegated to http://dev.united.com/rel3/fares.v1*
- Versioning is a red-herring: new namespace, new URI
- HTTP content negotiation leaves schema format open
- … but disagreements are still accurately flagged
  - *Disambiguates Air, Hotel, and Auto definitions of <DAY>*

… and many more IScale namespaces:

- Dublin Core
- Library of Congress classifications
- Yahoo! Categories
- ISBN / ISSN numbers
  - *http://isbn.nu/<isbn> - try it!*
- UPC product bar codes
- GPS coordinates (?)
- RFCs & Internet-Drafts
- User & Group profiles
- Printer Descriptions (PPDs)
- Video Codecs
- Fonts
- Colorspaces
- Java class files
- GUIDs (globally unique IDs)
- Social Security Numbers
- DUNS business ID number

Recap: Key Namespace Features

- Name of the Namespace
- Resolver system accepting such addresses
- Authority governing form of names
- Authority governing entries in namespace
- Internal structure of names, if any
  - *Directionality, if hierarchical*
- Lifetime of name (domain)
- Lifetime of address (range)
- Density – current size / potential size
- User Interface implications – internationalization
- Give three example entries in the namespace
- What subspaces are reserved, and for which purposes?
- Formally, is the resolver function a bijection (i.e. injective and surjective – having unique addresses, and names for all addresses?)
- What other namespaces map to it?
  - *E.g. phone numbers are also represented in the .tpc.int domain*
- Context-sensitivity – any additional parameters to the resolver function?
- *What's the resolution algorithm?*

Part II: Identifying IScale Issues

- Recall the three requirements we set forth:
  - *Names must scale across time* — *longevity*
    - Human- and machine-readability
    - Security and reliability
  - *Names must scale across space* — *latency*
    - Scalable, nomadic, decentralized algorithms
    - Geography and other context-dependencies
  - *Names must scale across organizations* — *liability*
    - Names reflect trust decisions
    - Accommodating *anonymity*

IScale: Across Time

- Longevity requires readability
  - *Fixed format standards preserve machine-readability*
  - *Human-readable names for recoverability and usability*
    - Internationalized, graphical, and audio “names” exist, too
- Longevity requires security and reliability
  - *Formats, protocols, and policies must be stable standards*
  - *Resolution services must be audited and bullet-proof*
  - *Reliable on-line access can increase fidelity (up-to-date)*
  - *Mobility, by contrast, calls for agility; rapid updates*

IScale: Across Space

- Physical scale's most salient constraint is *latency*
  - *Far beyond a LAN's RTT of 30 ms, past Internet's 300 ms, all the way to nomadic disconnection for days at a time*
  - *Calls for new mobile, decentralized resolution strategies*
- Physical scale is also an opportunity
  - *'Geospatial hypertext' shows the way to content that resolves specifically for a reader's location*
    - E.g. having "united.com" return the nearest ticket office
  - *Conversely, planetary reach mocks global namespaces*
    - E.g. "tollroad.com", which resolves to a few miles of Hwy 73 at UCI...

IScale: Across Organizations

- Organizational boundaries *are* trust boundaries
  - *Thus, multilaterality is a key IScale issue:*
    - Explicit delegation of naming authority can reduce contention
    - Explicit levels of commitment: private, experimental, public, and so on
  - *Paranoia also follows from strong trust boundaries:*
    - So decentralization is even more of an IScale issue than distribution
- Liability accrues at those boundaries
  - *Drives need to explicitly articulate the namespaces used*
- Anonymity and pseudonymity are also solutions!

A Vision: Postmodern Naming

- *How do human societies handle naming, anyway?*
  - People are not uniquely named
  - Not all people are even uniquely addressable
  - No person or organization can enumerate all people
  - People arguably manage self-organizing namespaces
- Everyone has their own personal namespace, yet we're all only a few degrees of separation apart
- *What will the meaning of a name be when computers have to play 'six degrees of separation', too?*

Take-Home Points

- There are many, many kinds of IScale Namespaces
- There are genuinely Internet Scale issues
  - *Decentralized Algorithms: Protocols, Standards*
  - *Decentralized Policies: Politics, Trust, Economics*
- There are genuinely Internet Scale solution patterns engineers need documented