Internet 2.0 – Challenges for the Future Internet

Henning Schulzrinne
Dept. of Computer Science
Columbia University
New York, NY

IEEE DLT 2009

Overview

- The Internet as core civilizational infrastructure
- Devices and services
- The Internet is more than web 2.0
- Challenges
  - Network address exhaustion
  - Routing table explosion
  - Network ossification
  - Securing the network infrastructure
  - Usability & towards self-managed networks

IP as a core infrastructure interface

A taxonomy of Internet-connected devices

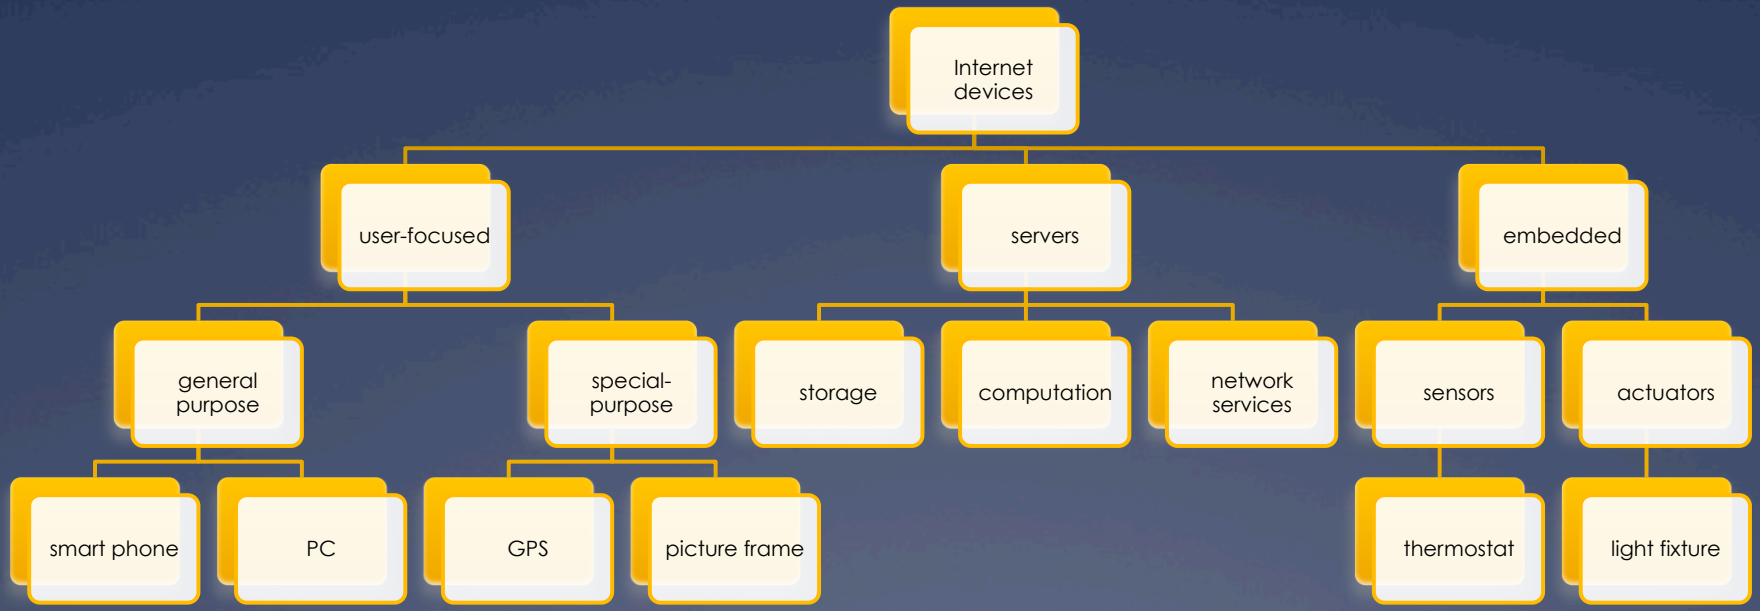

Internet-connected display devices

	Smart phone	Net book	Laptop	PC	TV
					
Screen	2-3"	7-12"	13-17"	19-22"	24-60"
Weight	< 0.5 lbs	2-3 lbs	3-8 lbs		
Sensors	Δ v/t, light, compass, GPS, microphone, camera		microphone, camera		

The great infrastructures

* Technical structures that support a society → "civil infrastructure"
  * Large
  * Constructed over generations
  * Not often replaced as a whole system
  * Continual refurbishment of components
  * Interdependent components **with well-defined interfaces**
  * High initial cost

water

energy

transportation

The Internet as core civil infrastructure

- Involved in all information exchange
  - (in a few years)
- Crucial to
  - commerce
  - governance
  - coordination
  - inter-personal communication
- Assumed to just be there
  - "plumbing", "pipes", …

Interfaces: Energy

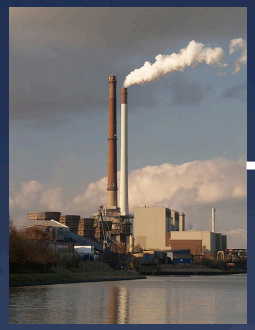

110/220V

1904

- Lots of other (niche) interfaces
- Replaced in a few applications

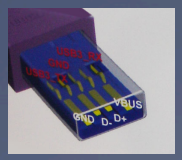

1901

Interfaces: Paper-based information

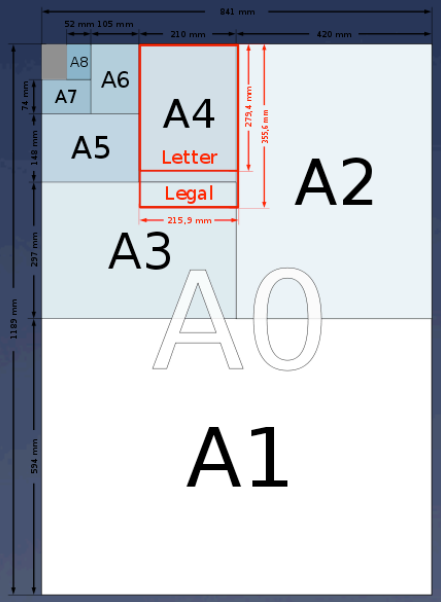

1798, 1922 (DIN)

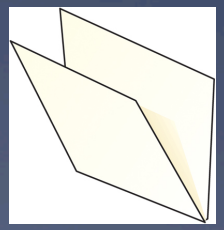

Interfaces: Transportation

1435 mm

1830 (Stephenson)
1846 UK Gauge Act

About 60% of world railroad mileage

12’

What makes interfaces permanent?

- Widely distributed, uncoordinated participants
- Capital-intensive
  - depreciated over 5+ years
  - see Y2K problem
- Allocation of cost vs. savings
  - e.g., ISP saves money, end user pays
- Hard to have multiple at once
  - "natural monopoly"

Extrapolating from history

- IP now "the" data interface
- Unclear that any packet-based system can be
  - ≥ 10 times cheaper
  - ≥ 10 times more functionality
  - ≥ 10 times more secure
- Replacing phone system due to generality, not performance
  - IP offers general channel
- → We're stuck with IPv4/IPv6
  - except for niche applications (car networks, BlueTooth, USB, …)

Technology evolution

- Early technology stages:
  - make it work
  - make it cheap
  - make it fashionable
  - This happened in the auto industry. Early cars barely worked at all, every journey was an adventure. In the 1920s Ford broke the automobile patent and built a car for the common man, a car that did not need the skills of a mechanic to drive. Reliability improved gradually until the 1970s when there was a sudden realization that consumers would pay more for a car that was not designed to rust. Today most cars will go 10,000 miles between services and not need major repairs beyond a clutch plate for 50,000 or even 100,000 miles
- Completion of conversion from analog to digital/ packet media
- Patterson: **S**ecurity, **P**rivacy, **U**sability, **R**eliability
  - phishing attacks, DDOS
  - cost of purchase vs. cost of ownership
  - dependability (crashes & reboots)

What defines the Internet?

Networks beyond the Internet, cont'd

Network model	route stability	motion of data routers
Internet	minutes	unlikely
mobile ad-hoc	$3\ \tau$	disruptive
store-carry-forward	$< 3\ \tau$	helpful

More than just Internet Classic

Network	wireless	mobility	path stability	data units
Internet "classic"	last hop	end systems	> hours	IP datagrams
mesh networks	all links	end systems	> hours	
mobile ad-hoc	all links	all nodes, random	minutes	
opportunistic	typical	single node	≈ minute	
delay-tolerant	all links	some predictable	some predictable	bundles
store-carry-forward	all nodes	all nodes	no path	application data units

IP model

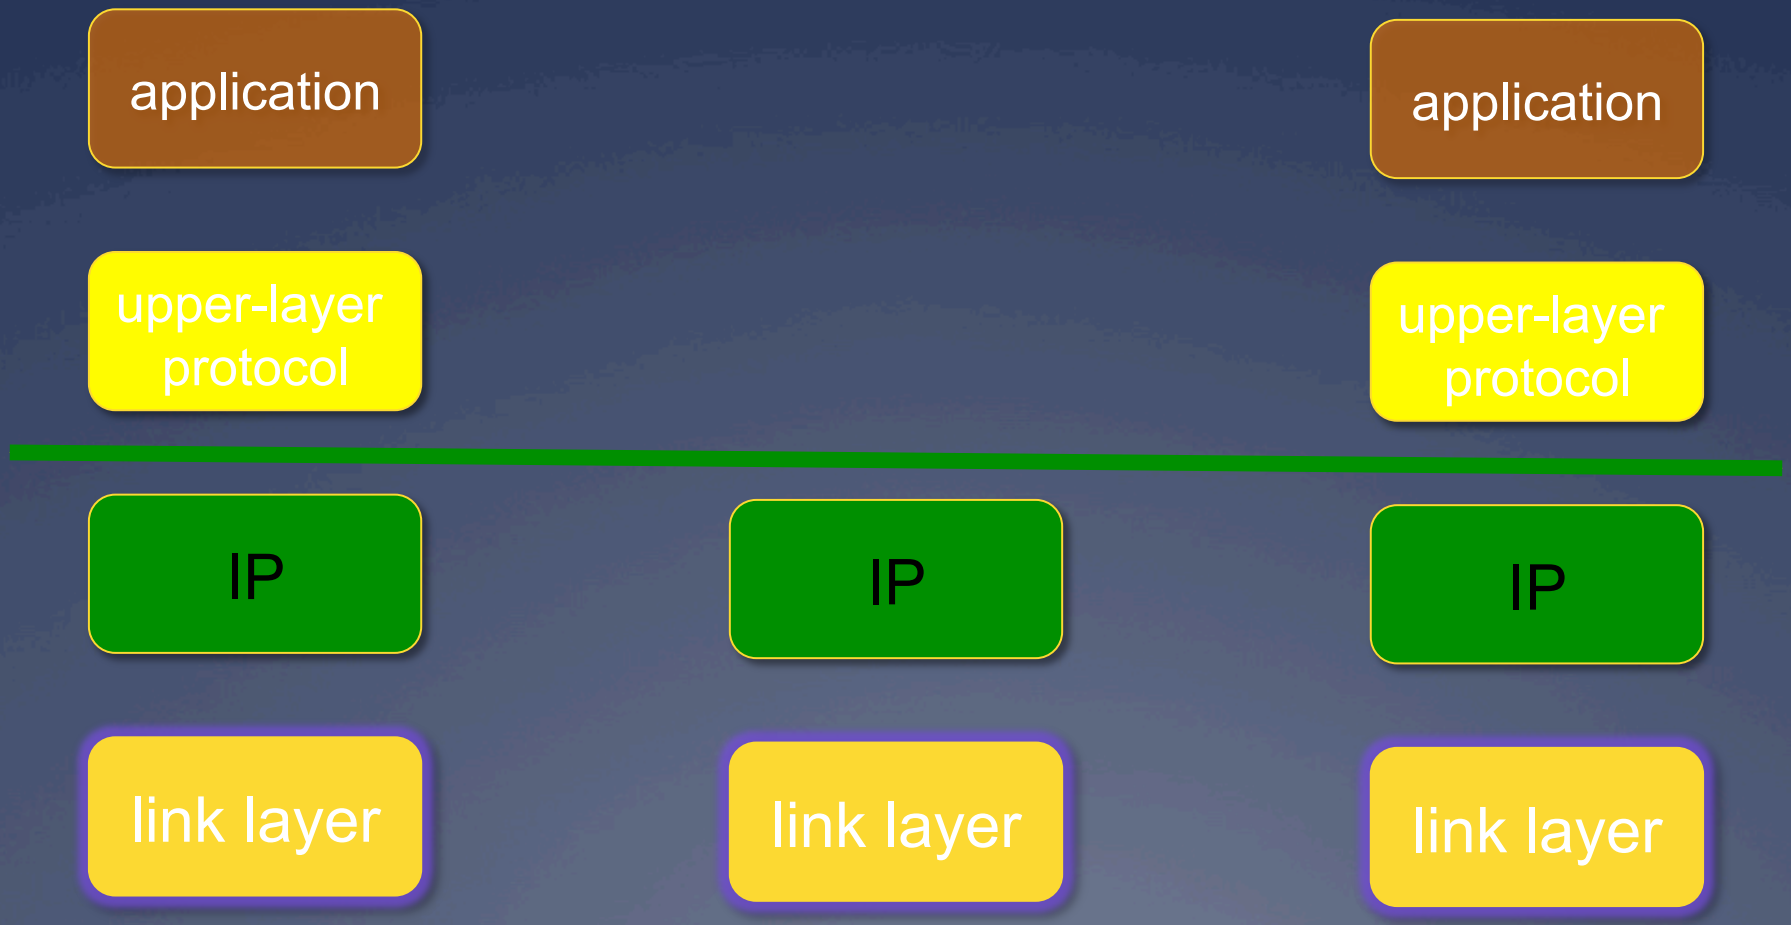
application
application
upper-layer protocol
upper-layer protocol
IP
IP
IP
link layer
link layer
link layer

Basic IP service model

- Unchanged since 1978
- Send without signaling
- Receive at provisioned address, without signaling
  - but: permission-based sending
- Variable-sized packets < ≈ 1,500 bytes
- Packets may be lost, duplicated, re-ordered

Addressing assumptions

- A host has only one address & one interface
  - apps resolve name and use first one returned
  - address used to identify users and machines
  - machine-wide DHCP options
- Failing
  - multi-homing on hosts (WiFi + Ethernet + BlueTooth + 3G)
- Attempts to restore
  - MIP: attachment-independent address
  - HIP: cryptographic host identify

Myth #1: Addresses are global & constant

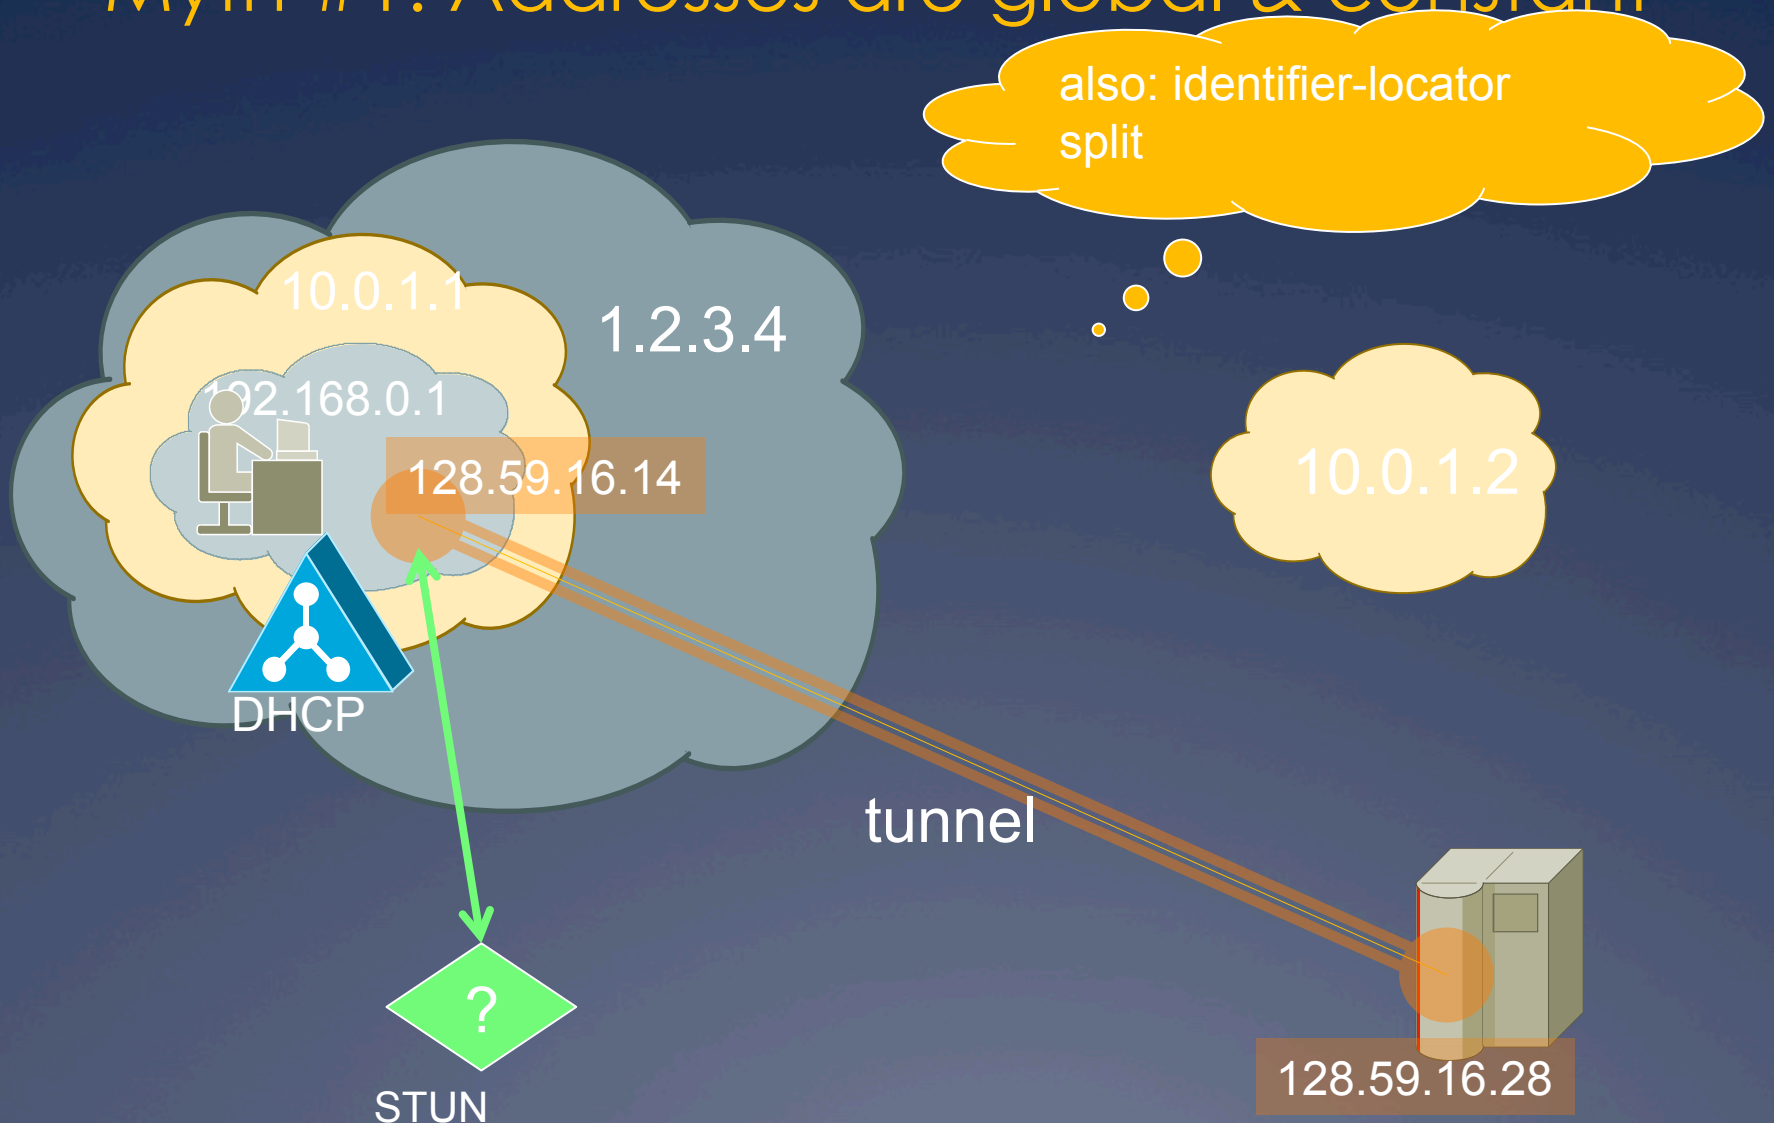

Myth #2: Connectivity commutes, associates

- Referals, call-backs, redirects
- Assumptions:
  - A connects to B → B can connect to A
  - A connects to B, B to C → C can connect to A
- May be time-dependent

Myth #2a: Bidirectional connectivity

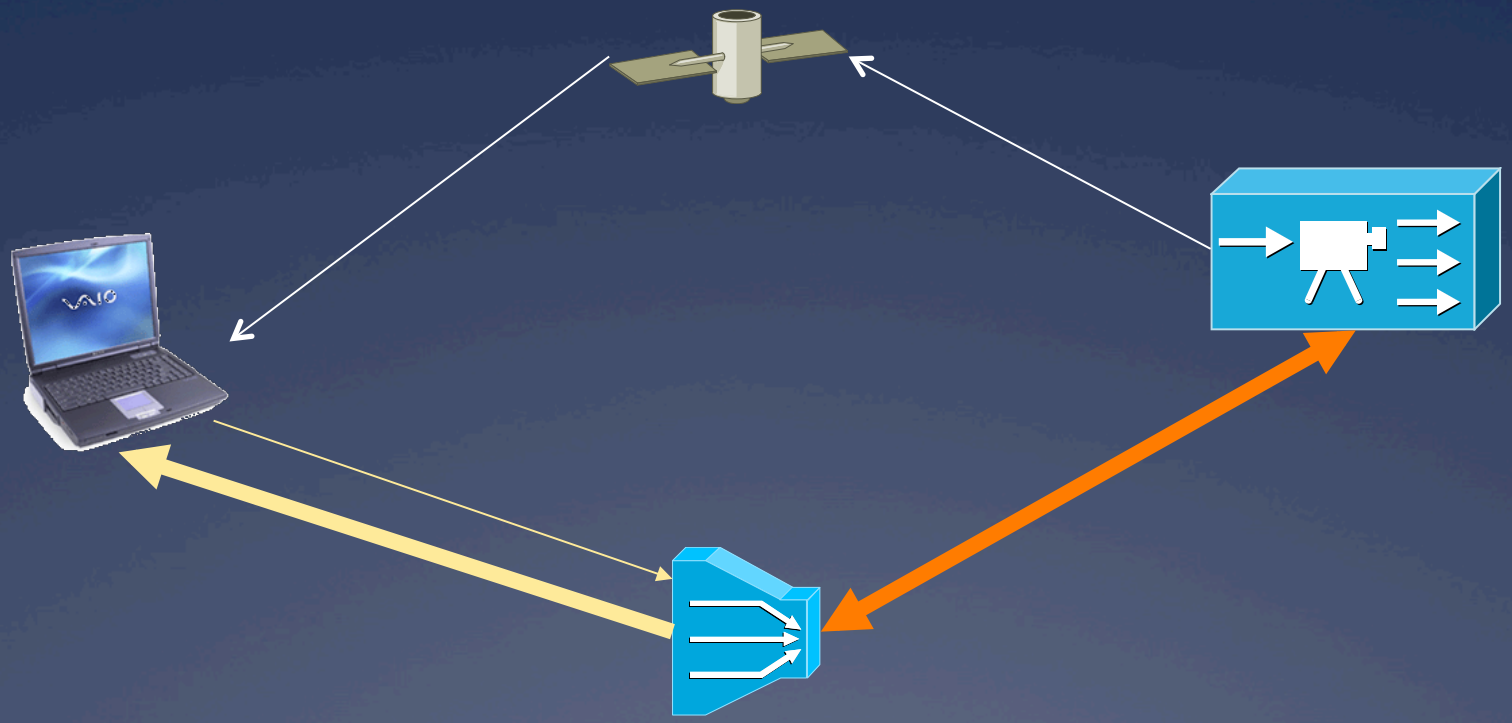

Myth #3: End-to-end delay of 1st packet typical

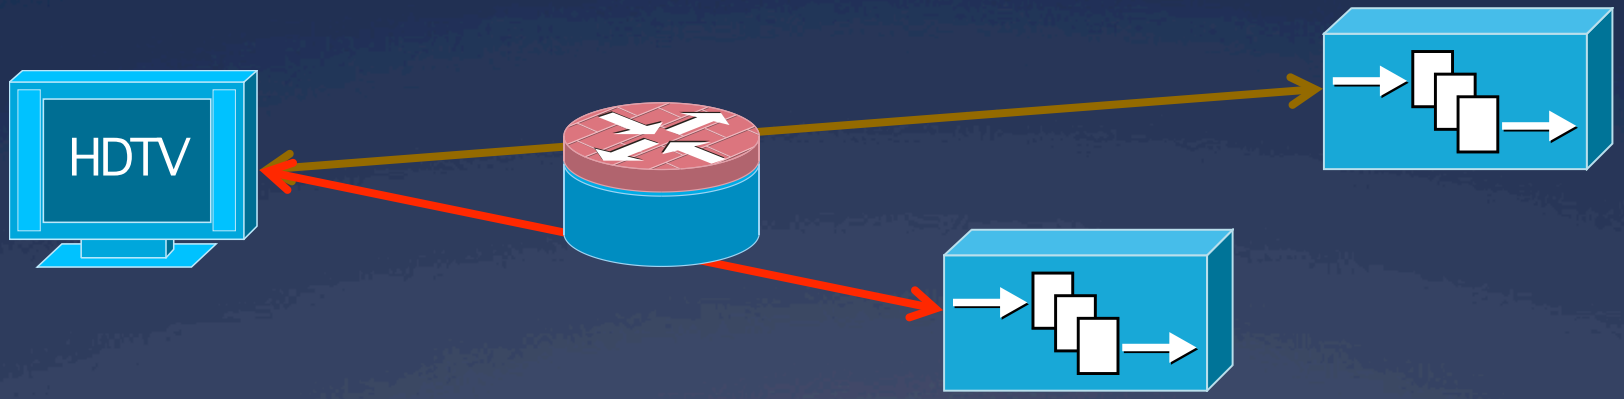

- 1st packet may have additional latency
  - ARP, flow-based routers
- MIPv6, PIM-SM, MSDP: fixed path during initial data burst
- → Choice of server may be suboptimal
  - higher delay, lower throughput, inefficient network usage

Challenges

A^7: Anytime Anywhere Affordable Access to Anything by Anyone Authorized

Jeanette Wing, NSF, Assistant Director for CISE

- Anytime and anywhere
  - From chip-level and biological networks to global scale
- Anything
  - Digital artifacts to services
- Anyone
  - "young and old, rich and poor, abled and disabled, literate and illiterate"
- Access
  - "Only authorized users will have the relevant access rights."
- Affordable
- Authorized

http://www.cra.org/CRN/articles/nov08/Wing-A7.html

User challenges vs. research challenges

* Are we addressing real user needs?
  * Engineering vs. sports
* My guesses

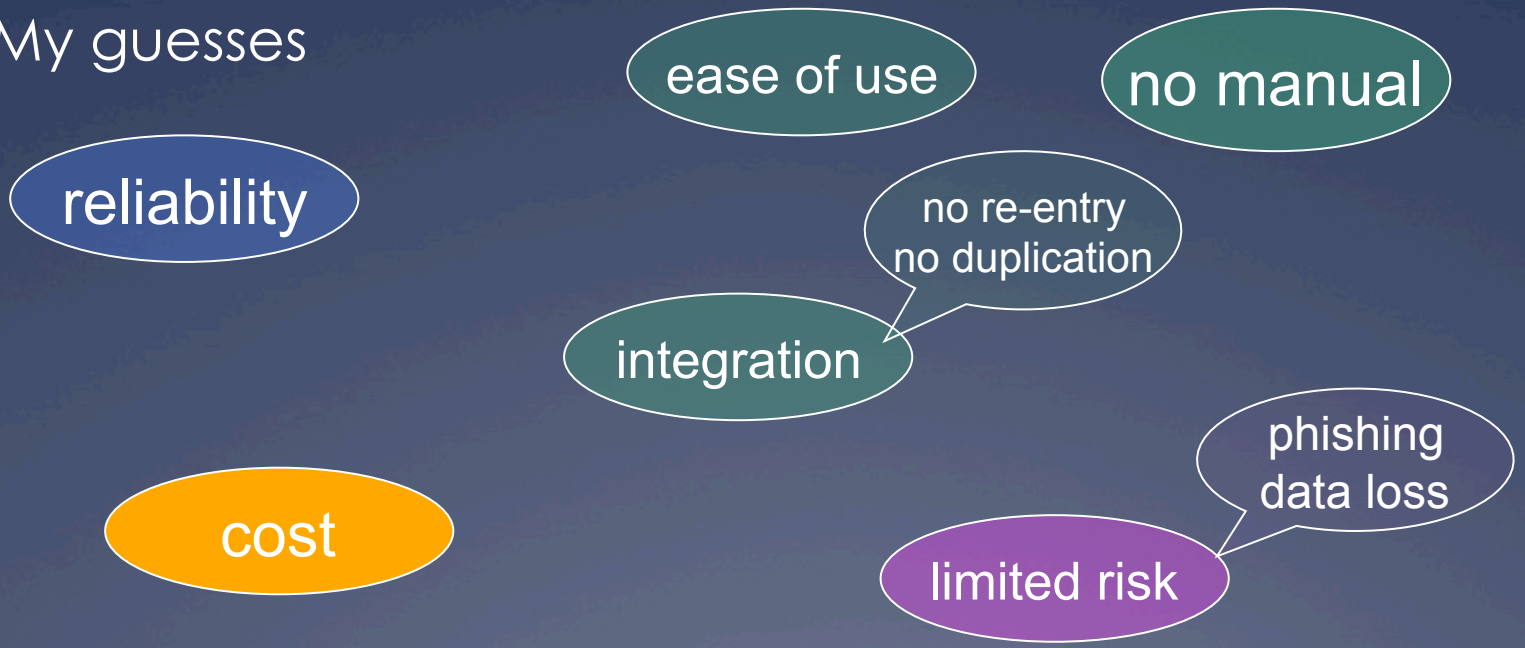

Cause of death for the next big thing

	QoS	multi-cast	mobile IP	active networks	IPsec	IPv6
not manageable across competing domains	✞	✞	✞	✞		
not configurable by normal users (or apps writers)	✞			✞	✞	
no business model for ISPs	✞	✞	✞	✞	✞	✞
no initial gain	✞	✞	✞	✞		✞
80% solution in existing system	✞	✞	✞	✞	✞	✞ (NAT)
increase system vulnerability	✞	✞	✞	✞		

Network ossification

Challenges

Why is the Internet ossifying?

- Lack of network transparency
  - NATs
    - ➔ only UDP + TCP
    - ➔ only client-server
  - Firewalls
    - only HTTP
- Standardization delays
  - No major new application-layer protocol since 1998
  - Protocols routinely take 5+ years
- Deployed base
  - Major OS upgrade every 7-8 years
  - But: automatic software updates
    - encourages proprietary application protocols

Which Internet are you connected to?

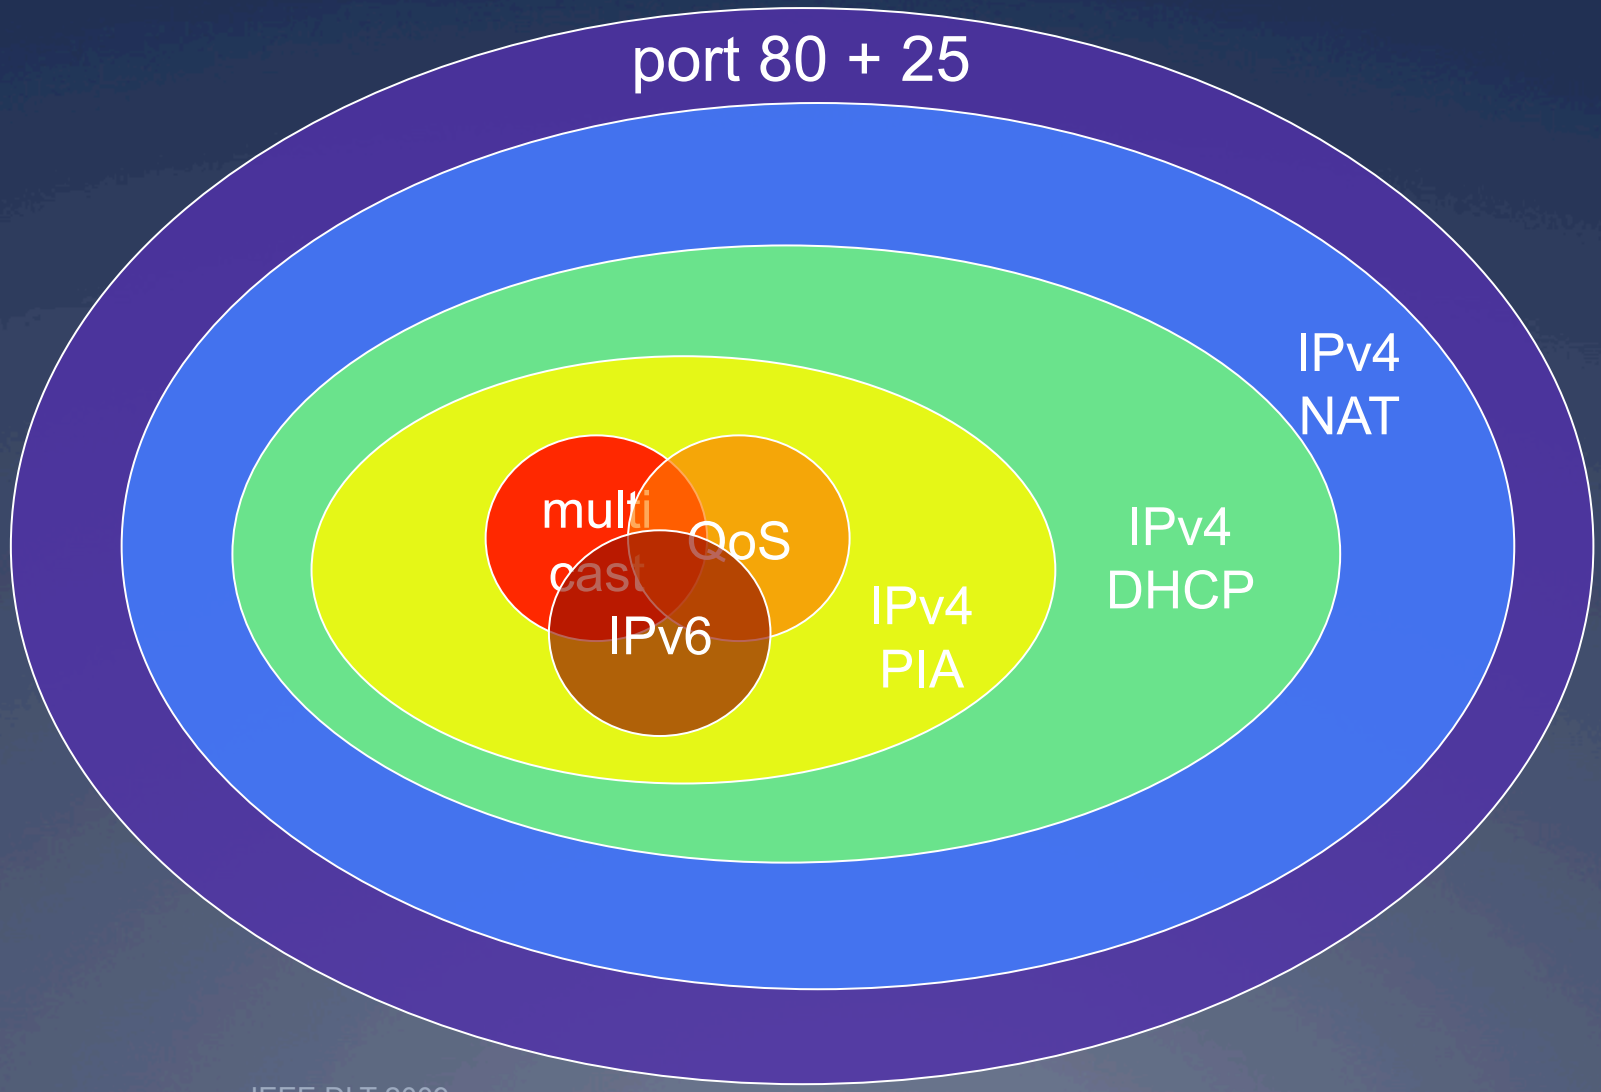

The two-port Internet

- Many public access systems only allow port 80 (HTTP) and maybe 25 (SMTP)
  - e.g., public libraries
- Everything tunneled over HTTP
  - Web-based email
  - Flash video delivery (e.g., YouTube)
  - HTTP CONNECT for remote login

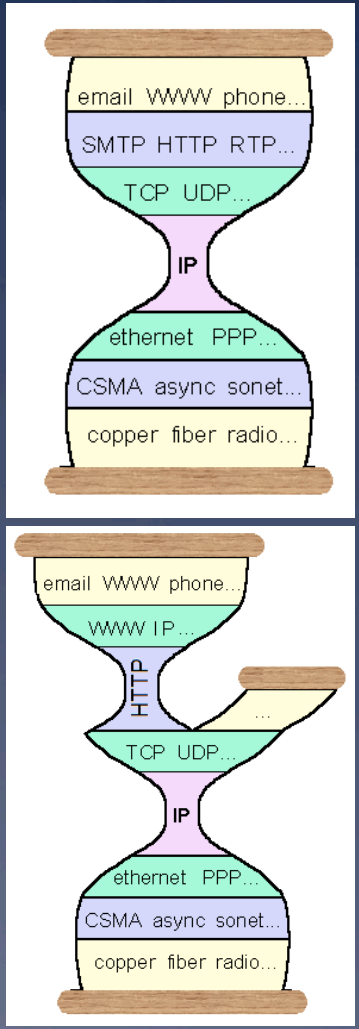

Dave Thaler

Causes

- Link-layer technologies
  - satellite, DSL
  - NBMA
- Network-layer technologies
  - security: broken by design vs. broken by accident?
  - NATs
  - Ill-defined meaning of IP addresses and names
    - theoretically, single network interface
    - practically, often more than that
      - virtualization
      - multi-homing
      - fail-over

Network challenges

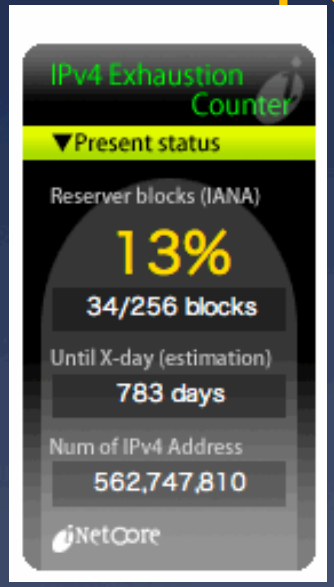

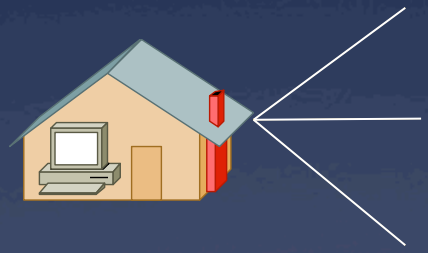

multi-homing

routing table explosion

+2 years

+5 years

+8 years

99.9 → 99.999%

zero configuration

The end of IP(v4) as we know it

Challenges

Where do IP addresses come from?

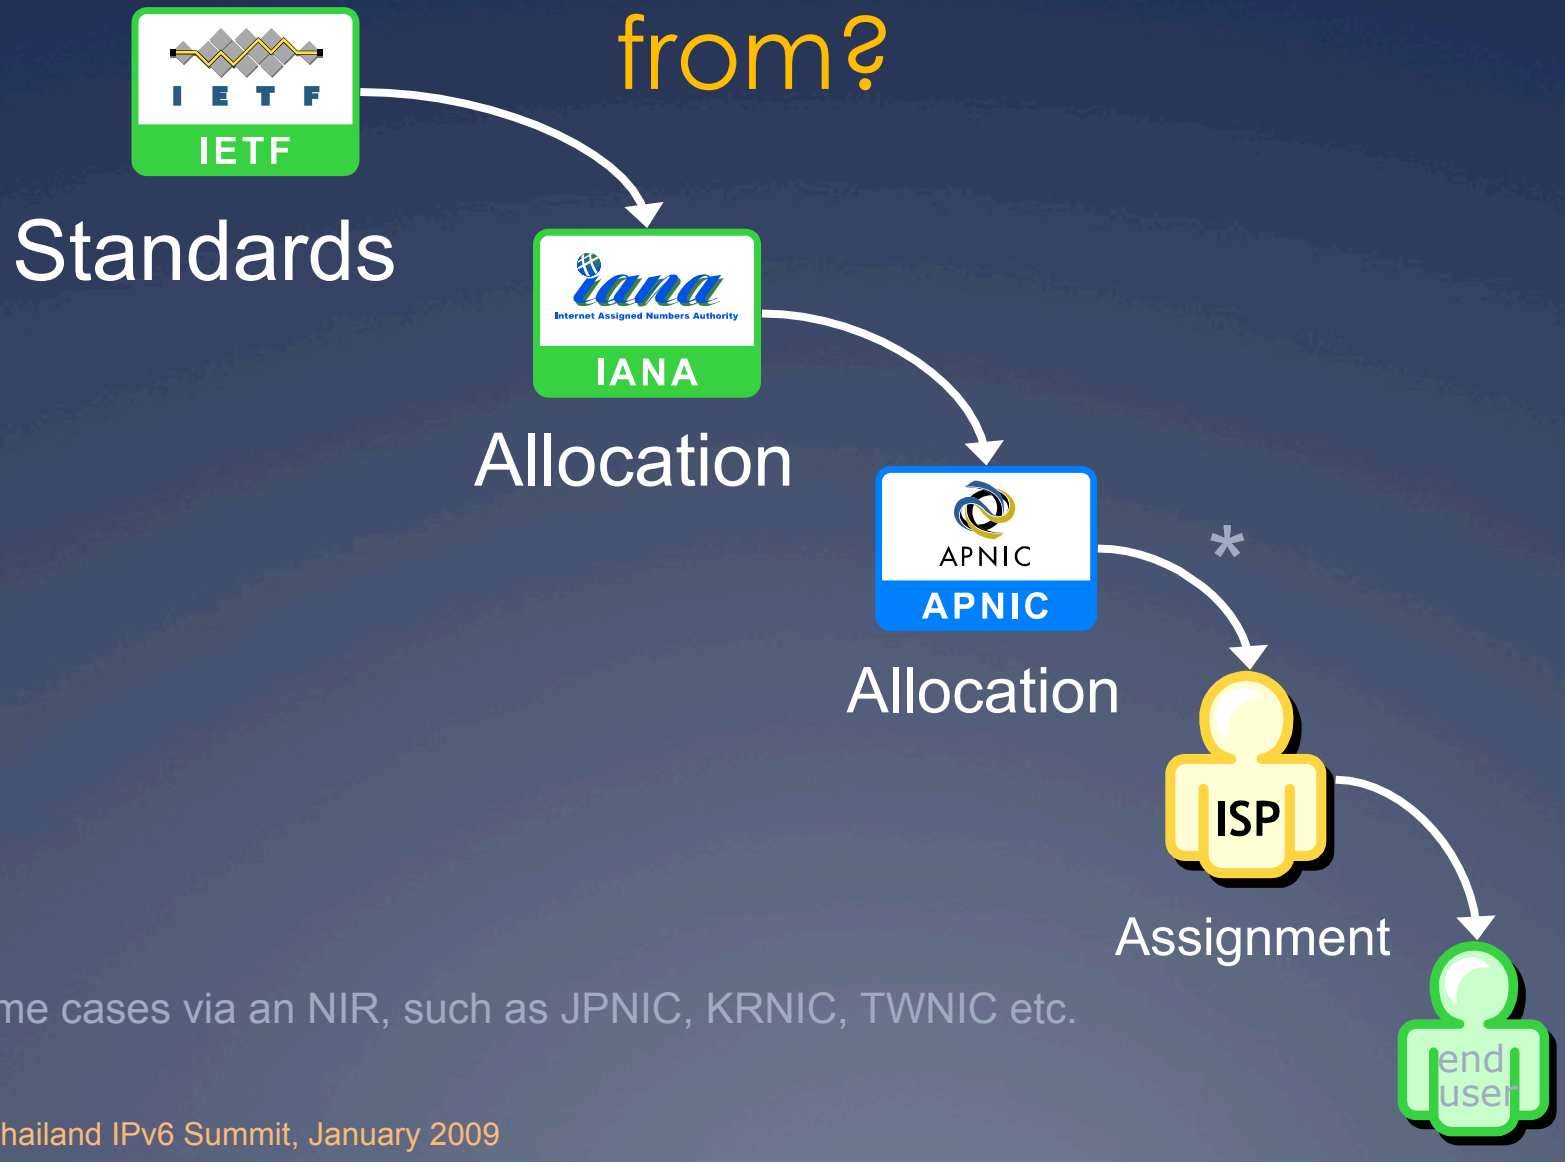

* In some cases via an NIR, such as JPNIC, KRNIC, TWNIC etc.

Regional Internet Registries

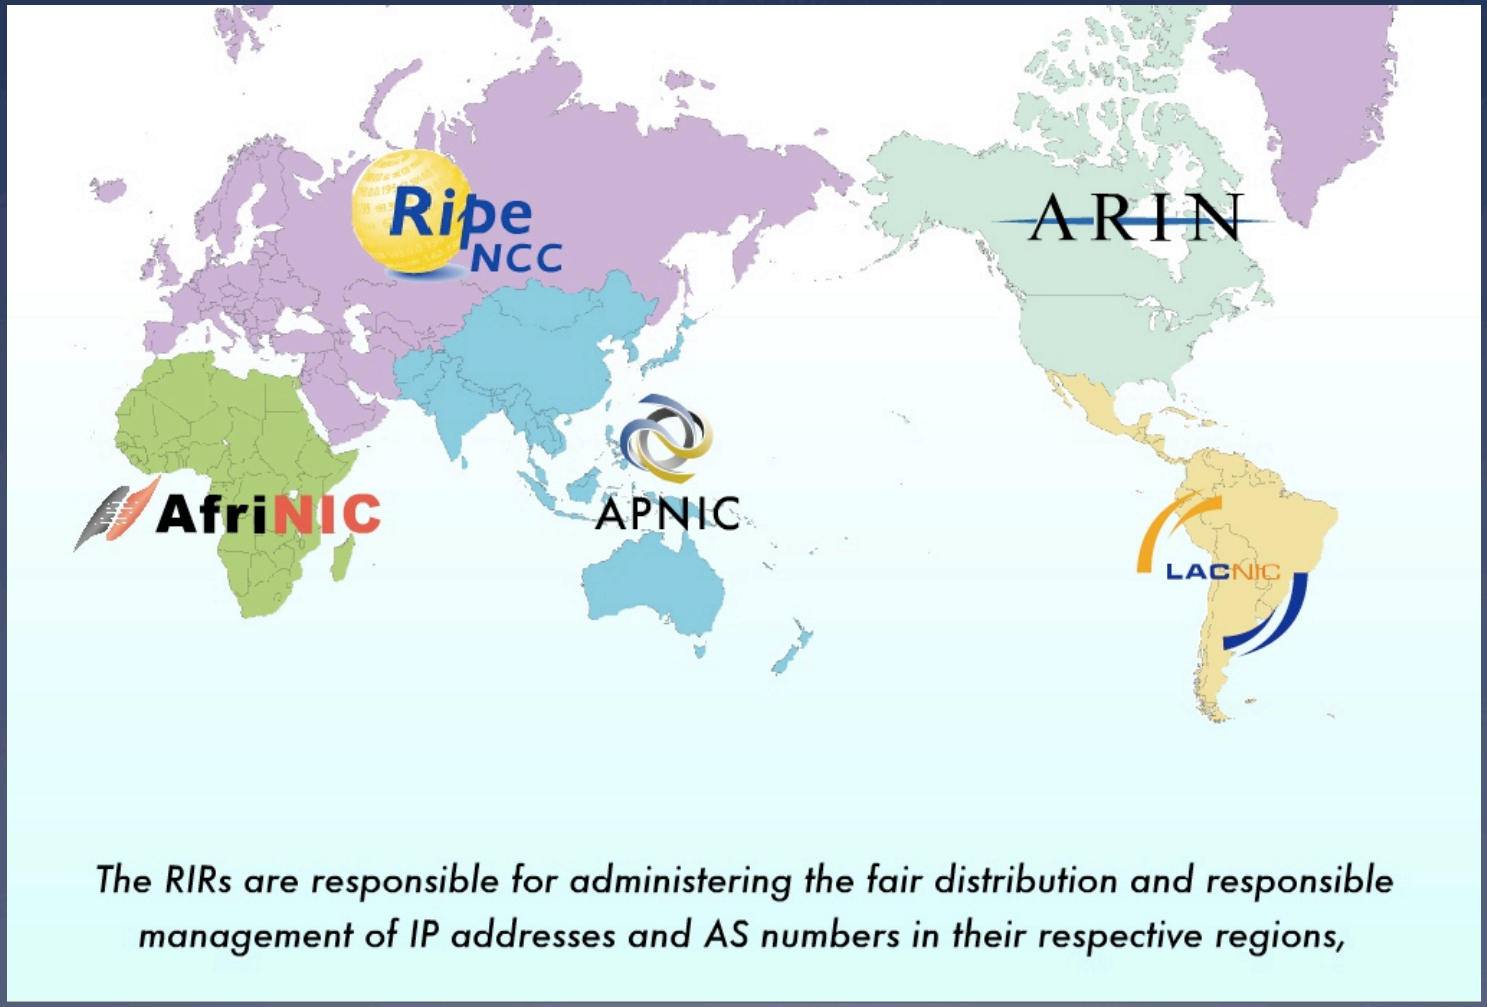

The RIRs are responsible for administering the fair distribution and responsible management of IP addresses and AS numbers in their respective regions,

IPv4 consumption – Projection

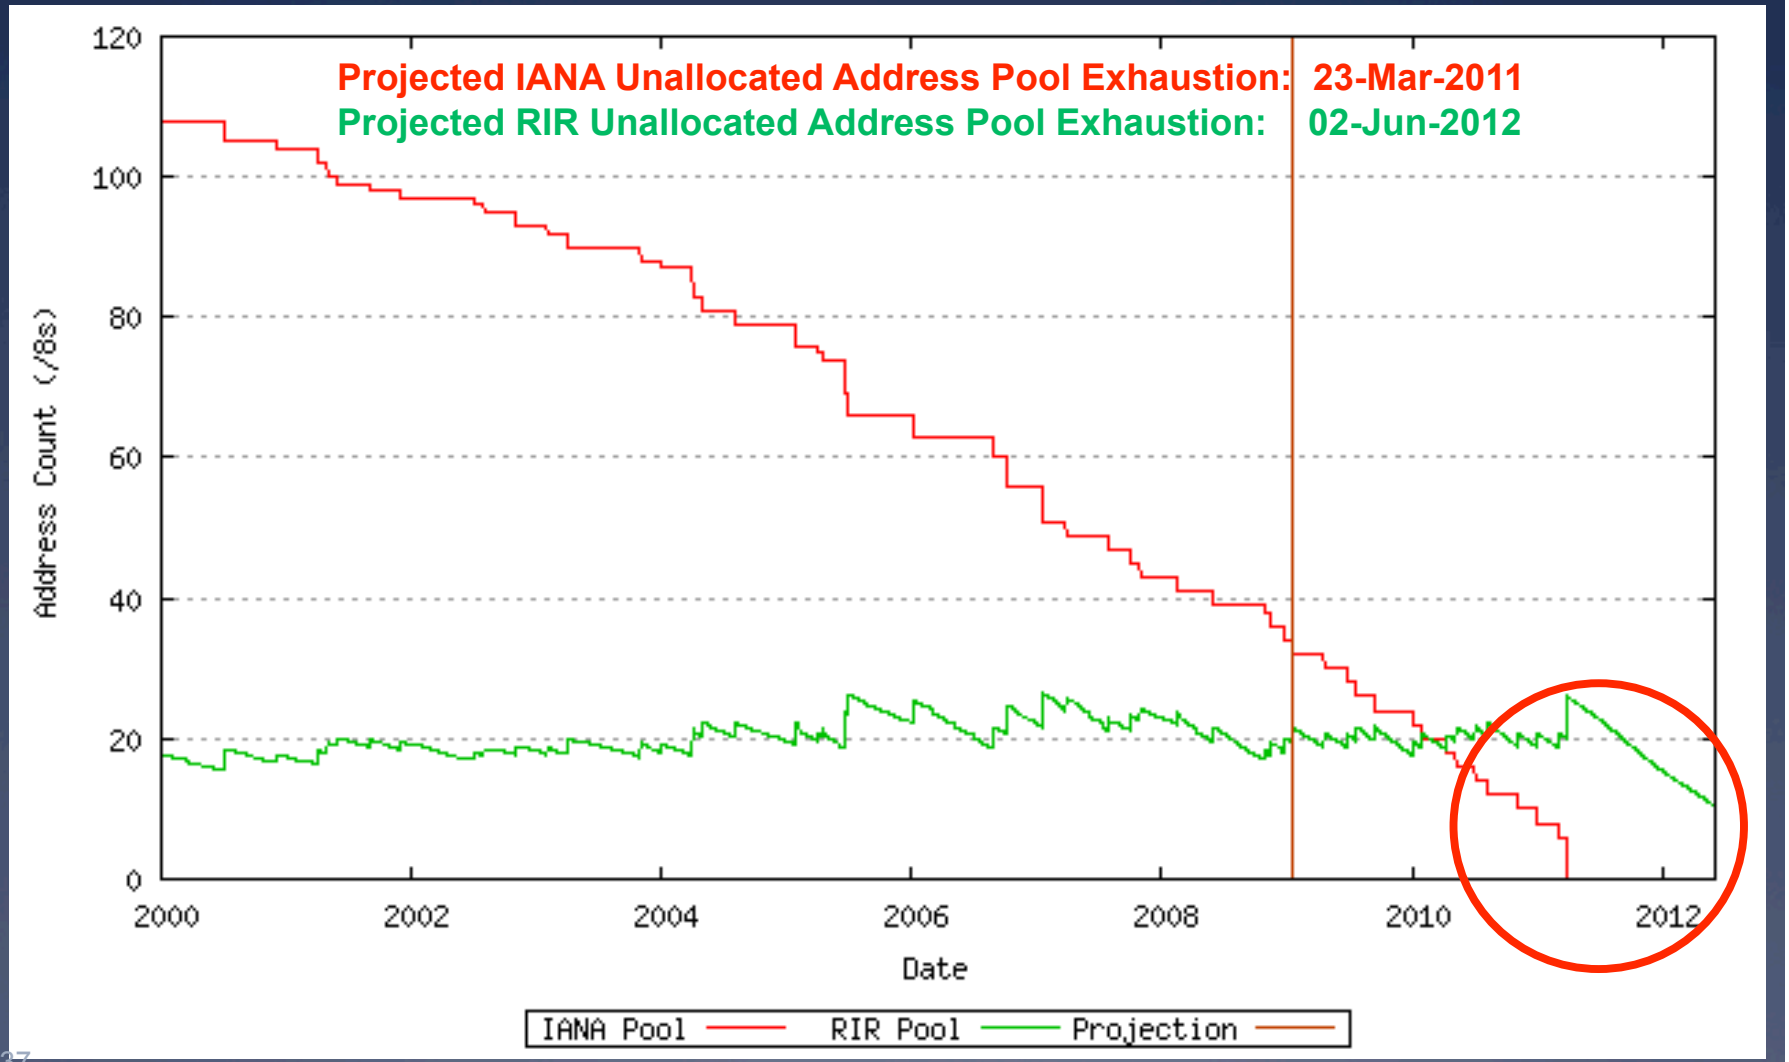

The transition to IPv6

- IPv4 needed for at least a decade
- Dual stack transition
  - but IPv6 server + non-IPv6 network + dual-stack server fail annoyingly
- NAT IPv4 ↔ IPv6
  - longer term, RFC 1918 (192.168.*.*) + global IPv6 address
- Decreasing IPv4 address demand
  - multi-layer ("carrier-grade") NATs →
    - limited effectiveness (hundreds of ports for BitTorrent or web page)
    - reliability problems
- Increasing IPv4 address supply
  - recycle unused /8s → few months supply
  - address auctions → router table size

The IPv6 choke points

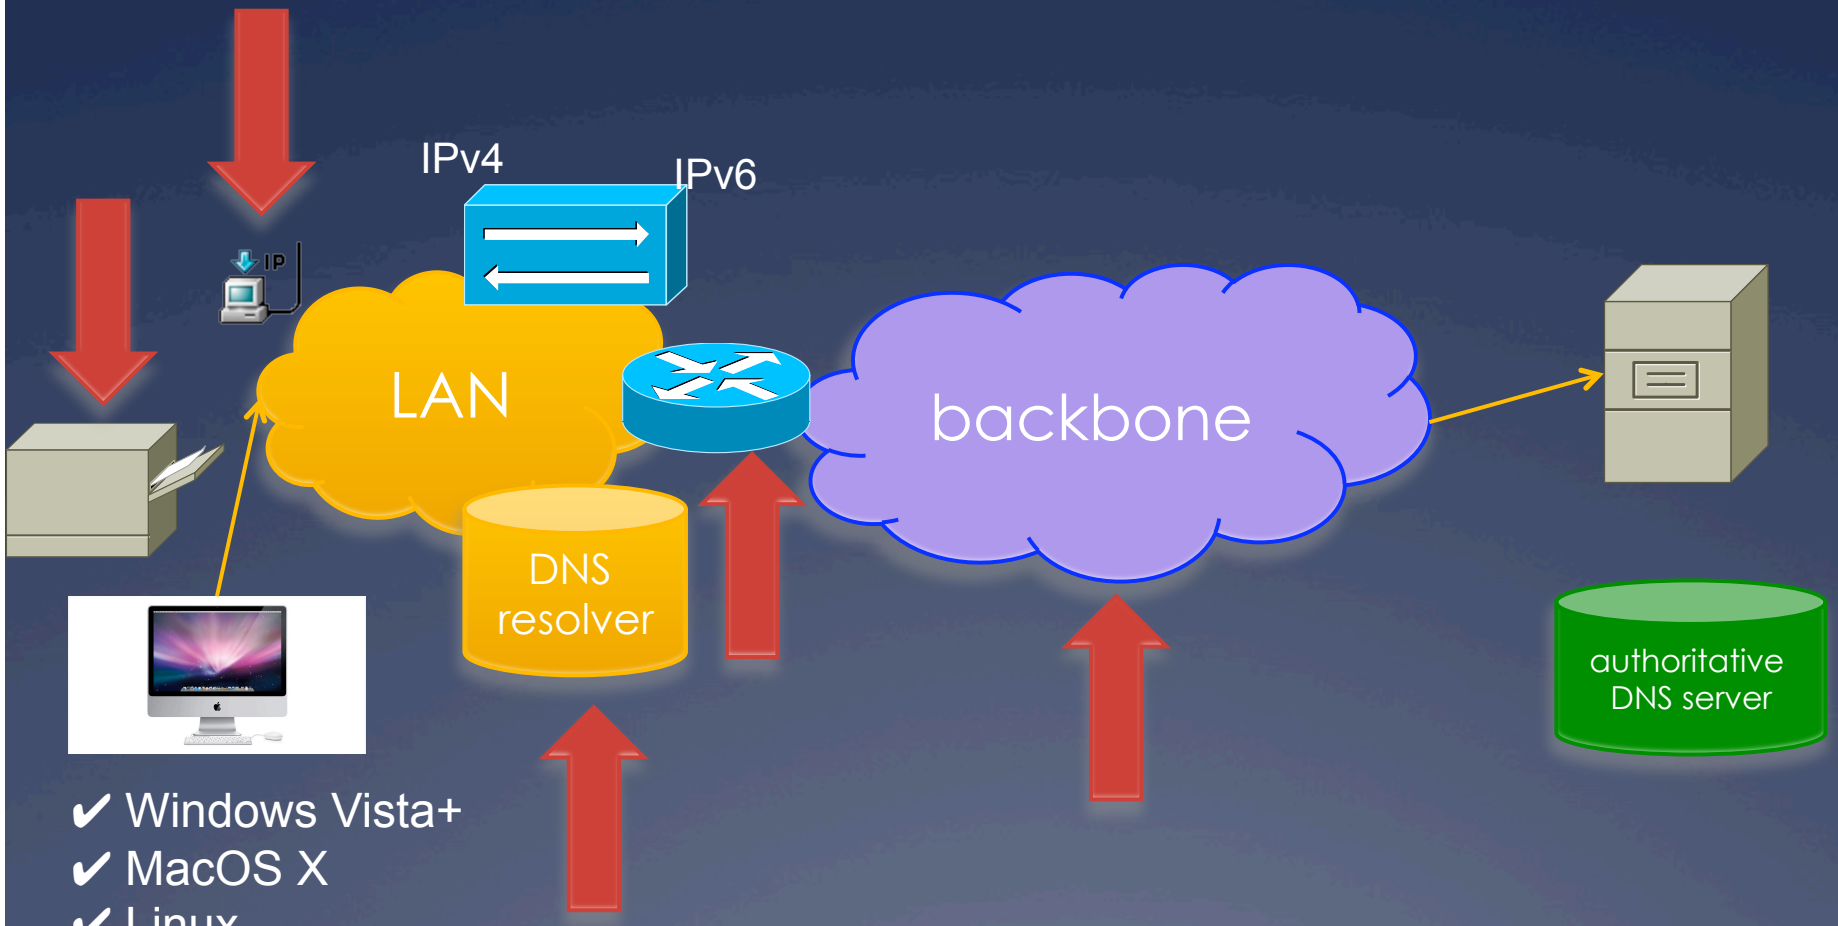

- ✔ Windows Vista+
- ✔ MacOS X
- ✔ Linux

IPv4 vs IPv6 – 2004 to present

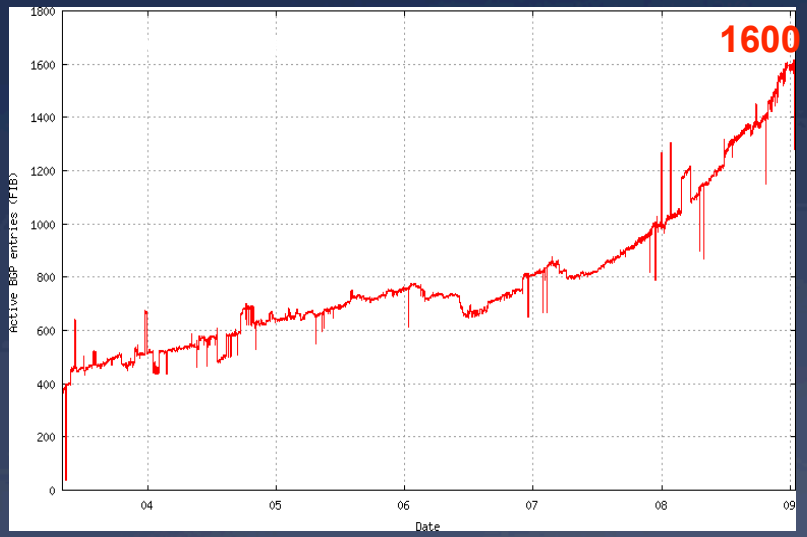

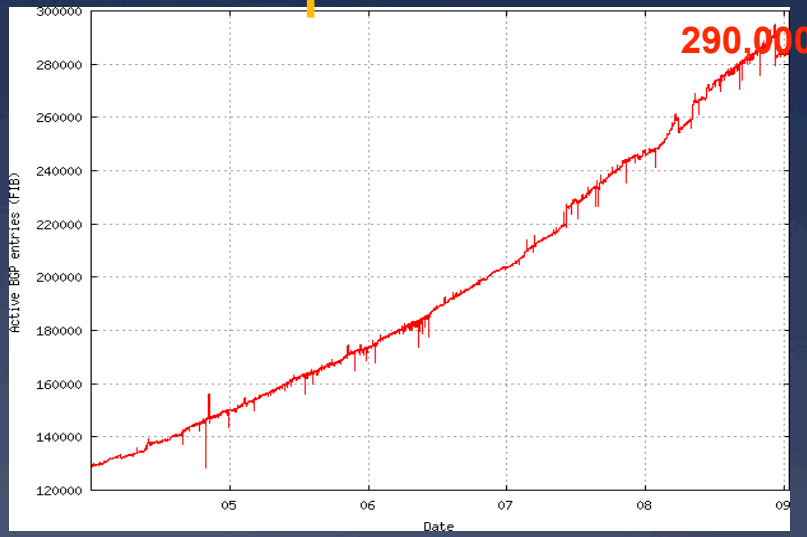

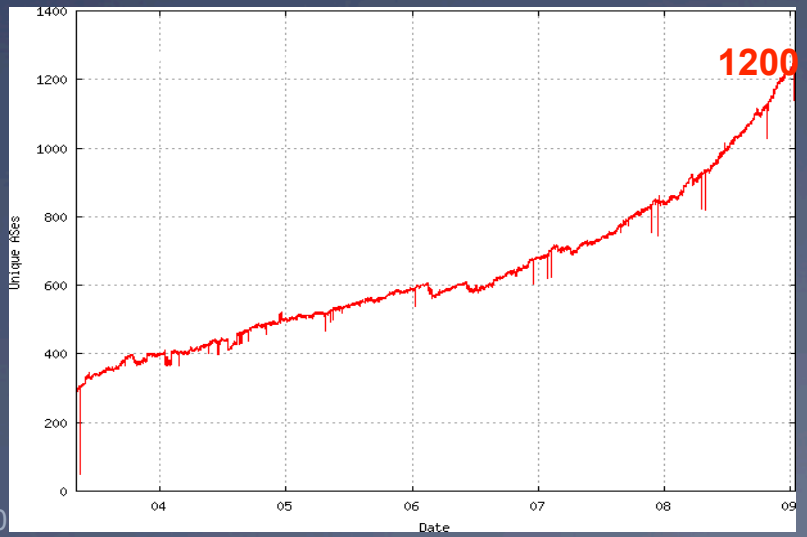

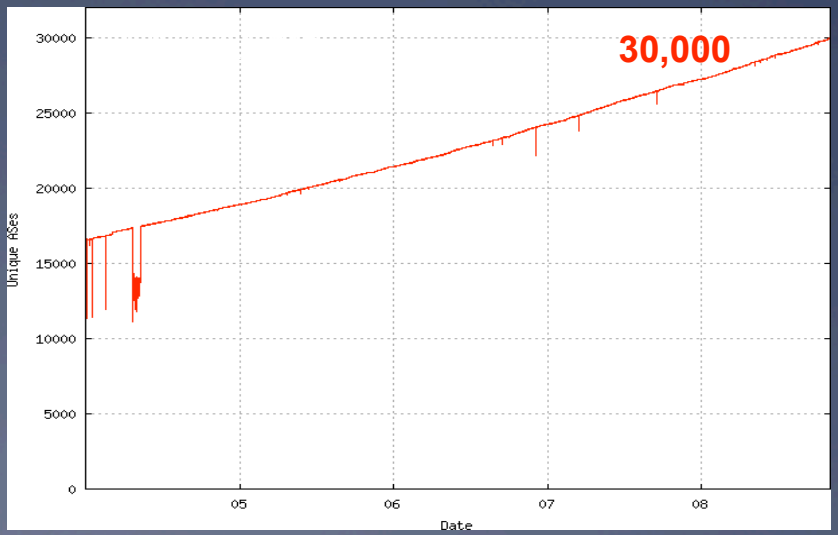

http://bgp.potaroo.net/ as of 15/01/2009

Pervasive multihoming

Challenges

Network of the (near) future

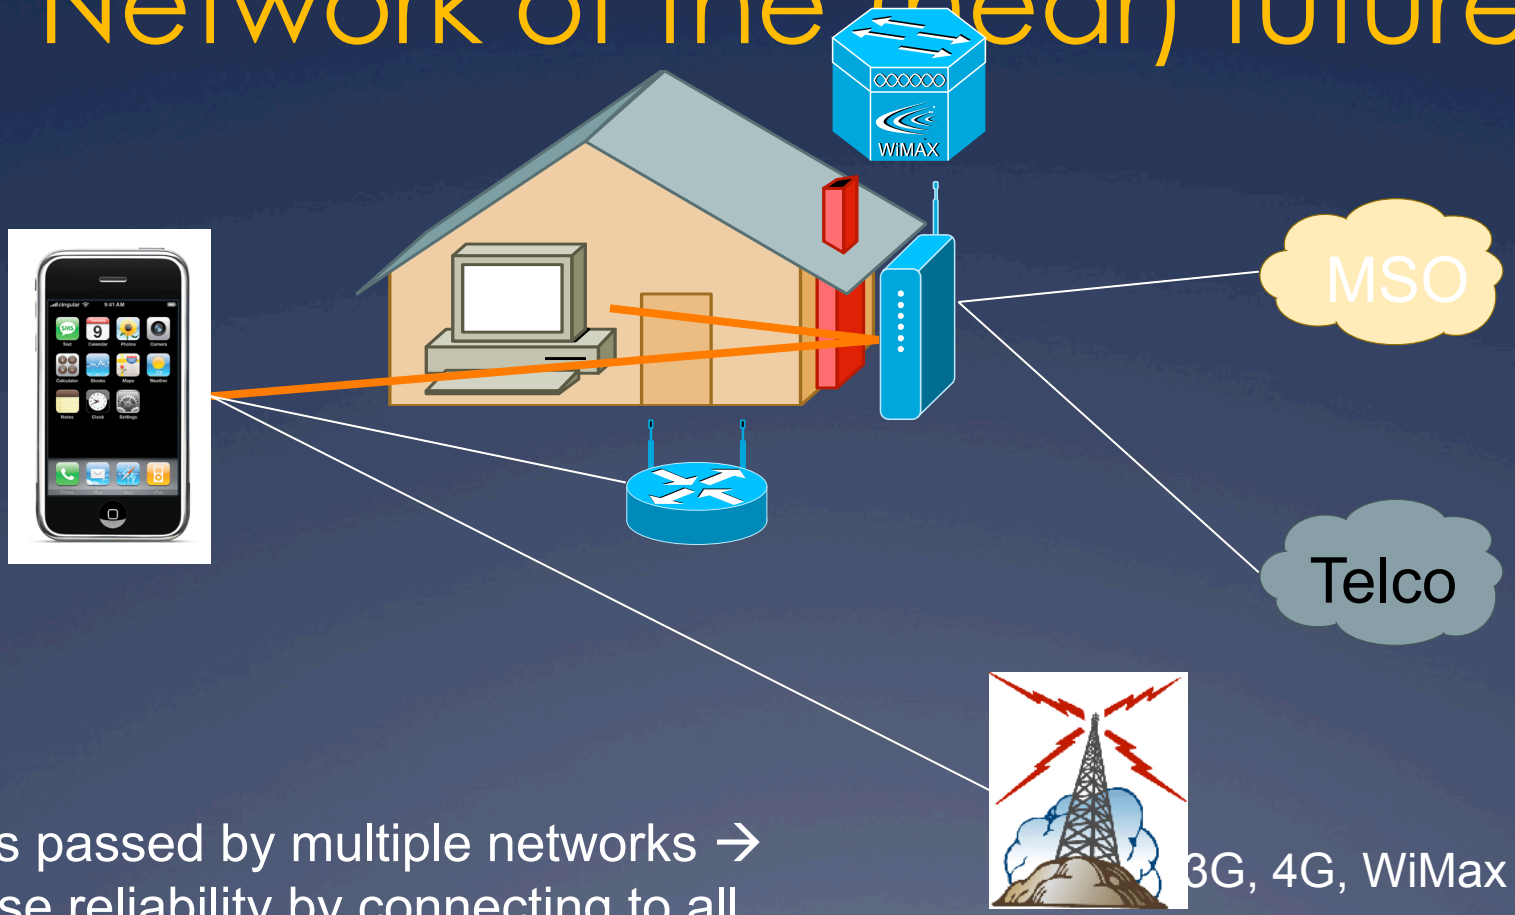

Homes passed by multiple networks → increase reliability by connecting to all ("reliable system out of unreliable components")

Multihoming (& mobility)

- Current IPv4 address ➔
  - identifier = unique host or **interface**
  - locator = network that serves host (provider)
- One system, multiple addresses:
  - multihoming: at the same time
  - mobility: sequentially
- Multihoming:
  - connections need to be aware of network path
  - socket interface makes it hard to program
- Solutions:
  - HIP: cryptographic host identifier
  - SHIM6
  - LISP: two network addresses
  - DNS: SRV, NAPTR

Example: BGP growth

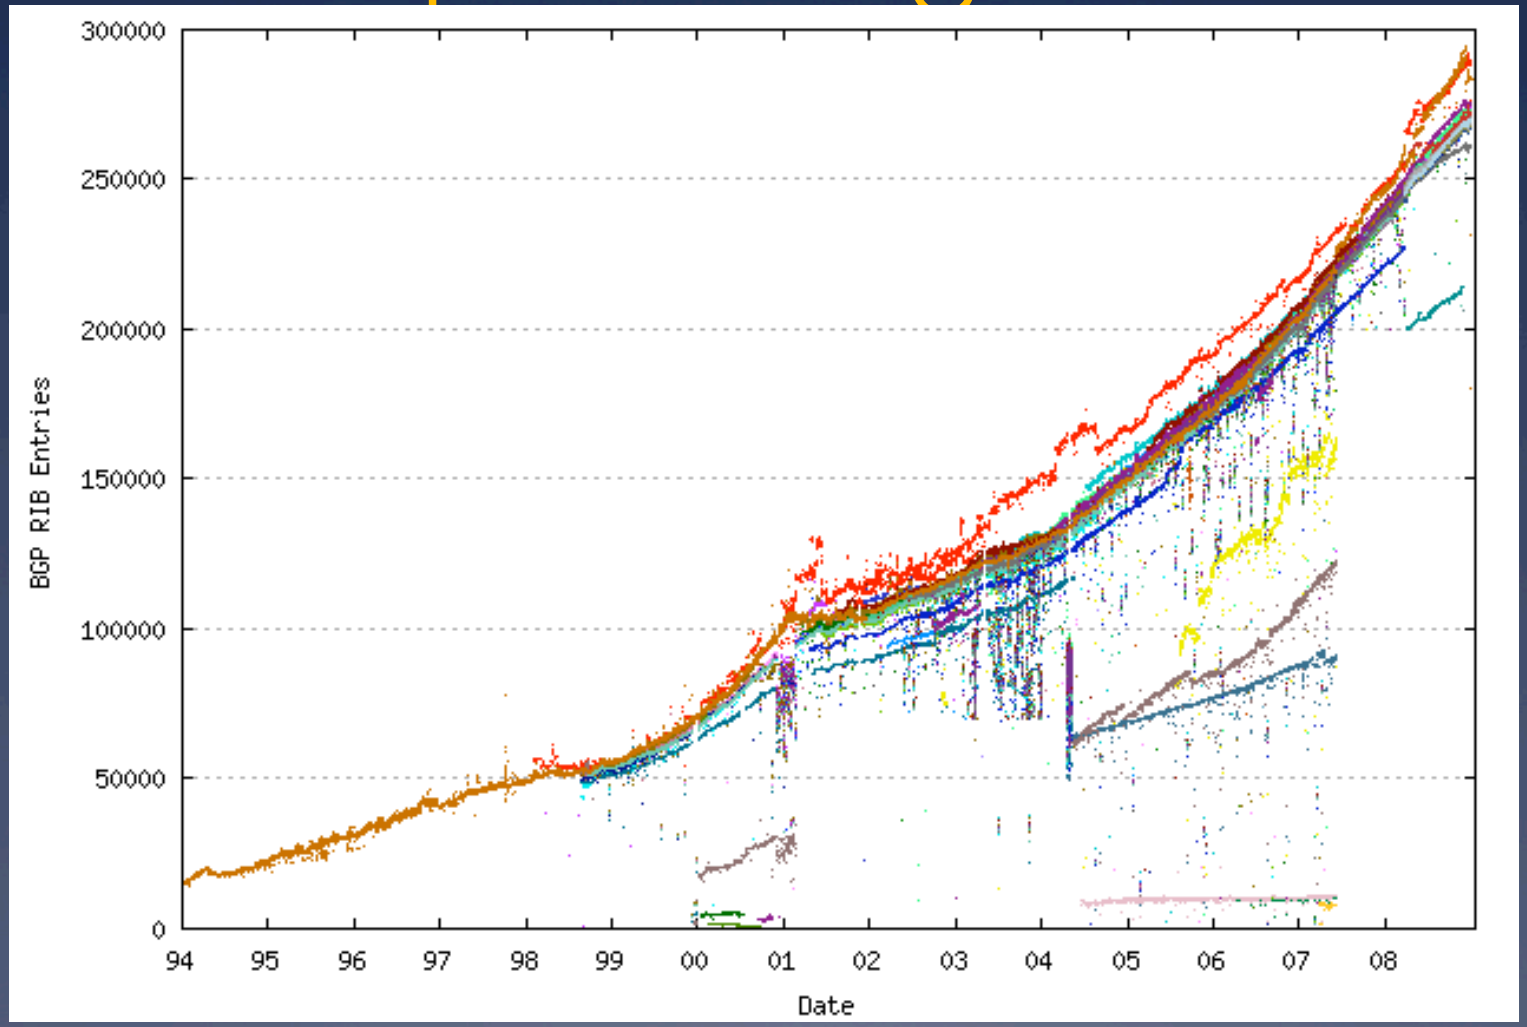

http://bgp.potaroo.net/

Security

Challenges

Network security issues

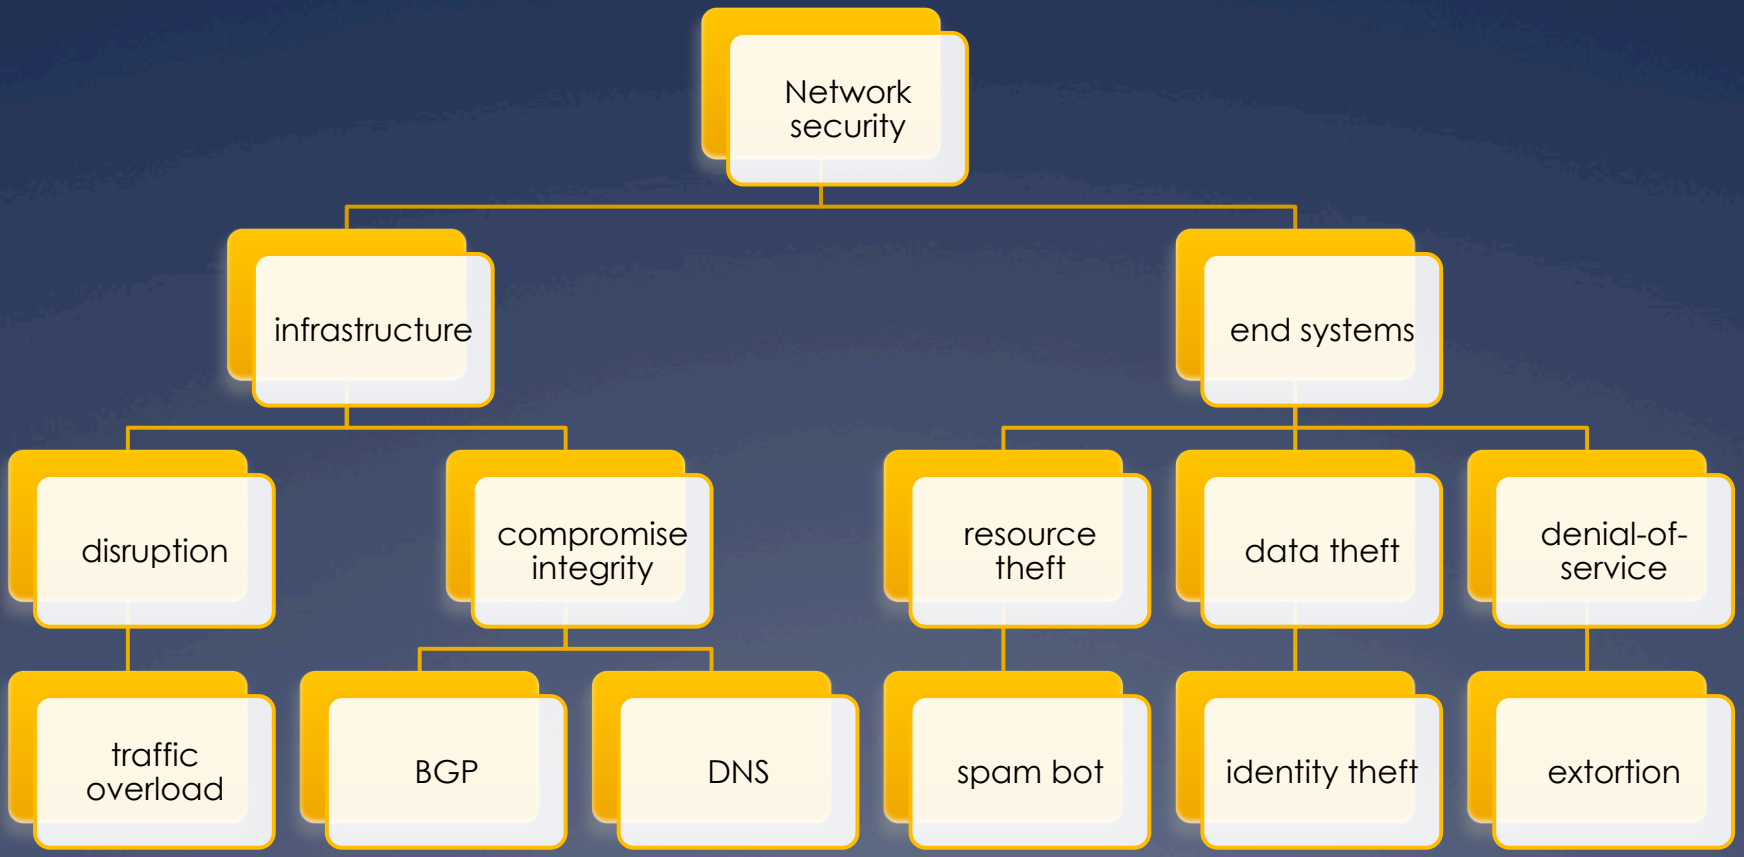

What about security?

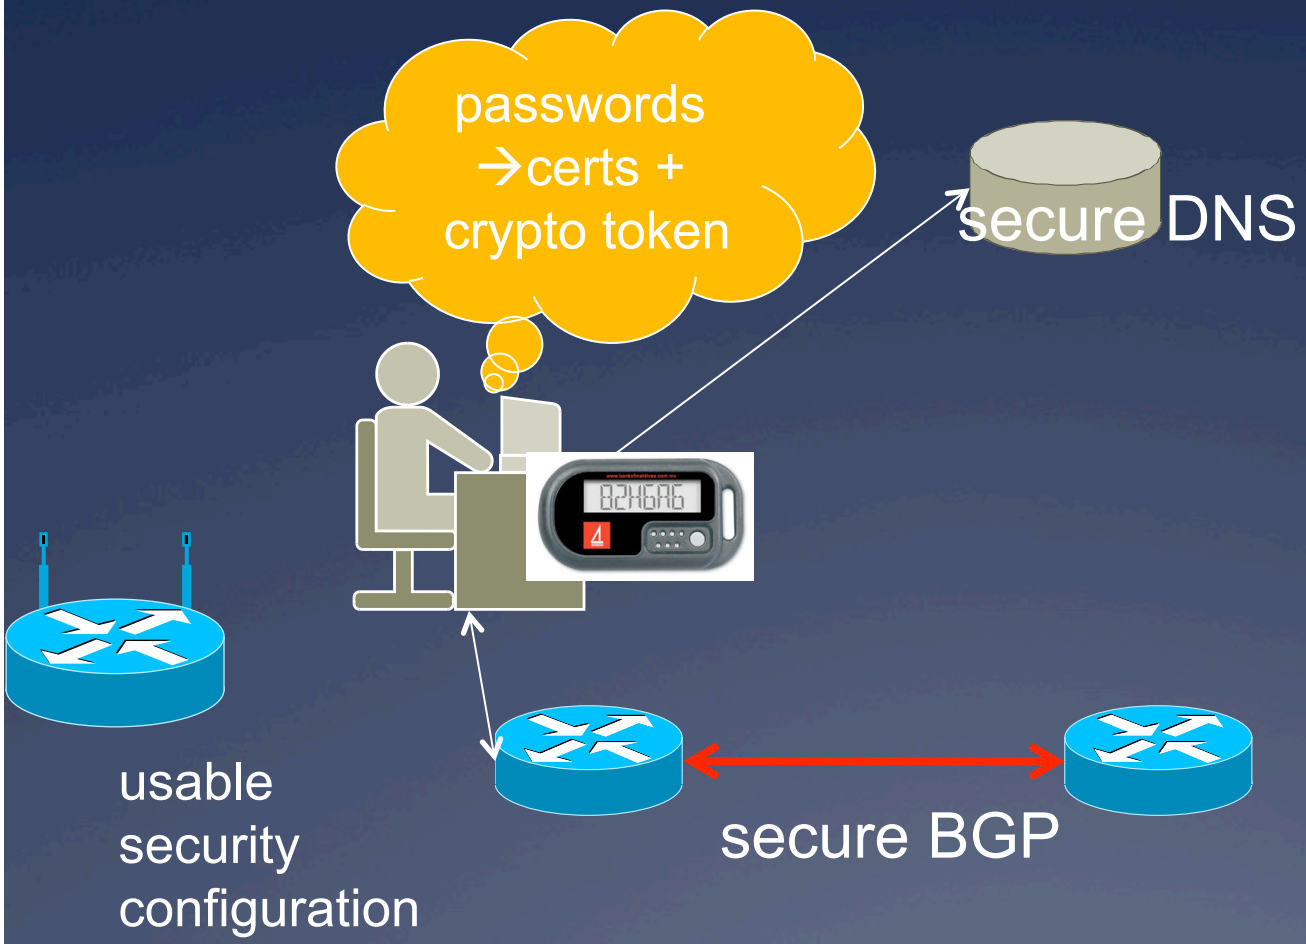

9: Political
8: Financial
Application
Presentation
Session
Transport
Network
Link
Physical

Technologies (mostly) available, but use & deployment hard

What about security?

- "The future Internet must be secure"
- Most security-related problems are **not** network problems
  - spam: identity and access, not SMTP
  - web: (mostly) not TLS, but distinguishing real bank from fake one
  - web: cross-domain scripting, code injection
  - browser vulnerabilities & keyboard sniffers
- Restrict generality
- Black list → white list
  - virus checker → app store
- Automated tools
  - better languages, taint tracking, automated input checking, stack protection, memory randomization, …
- Probably need more trust mediation

Usability

Challenges

Usability: Email configuration

- Application configuration for (mobile) devices painful
- SMTP port 25 vs. 587
- IMAP vs. POP
- TLS vs. SSL vs. "secure authentication"
- Worse for SIP...

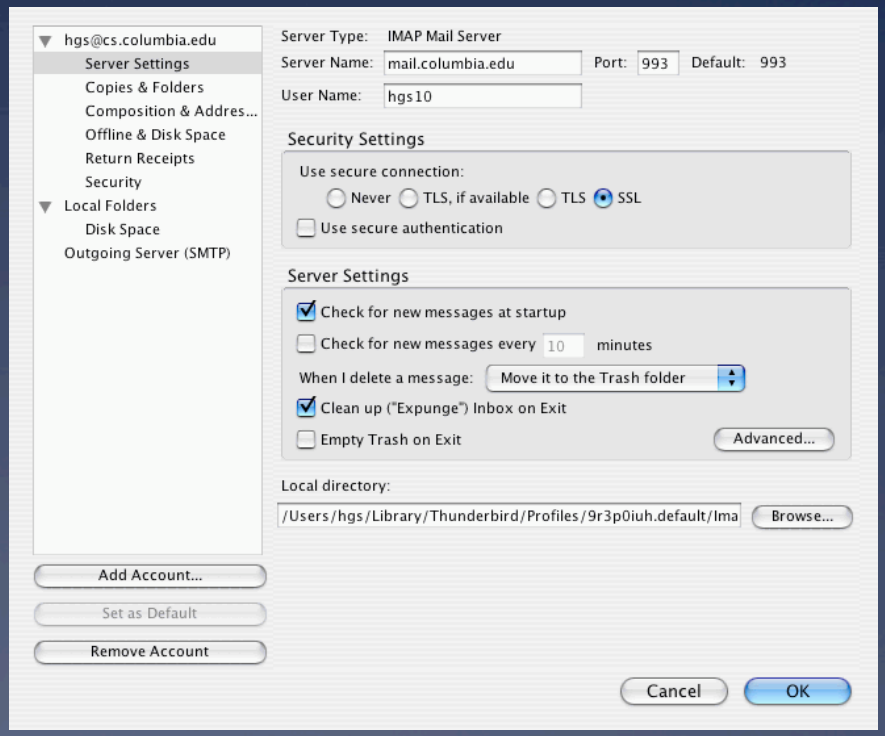

Usability: SIP configuration

partially explains

- highly technical parameters, with differing names
- inconsistent conventions for user and realm
- made worse by limited end systems (configure by multi-tap)
- usually fails with some cryptic error message and no indication which parameter
- out-of-box experience not good

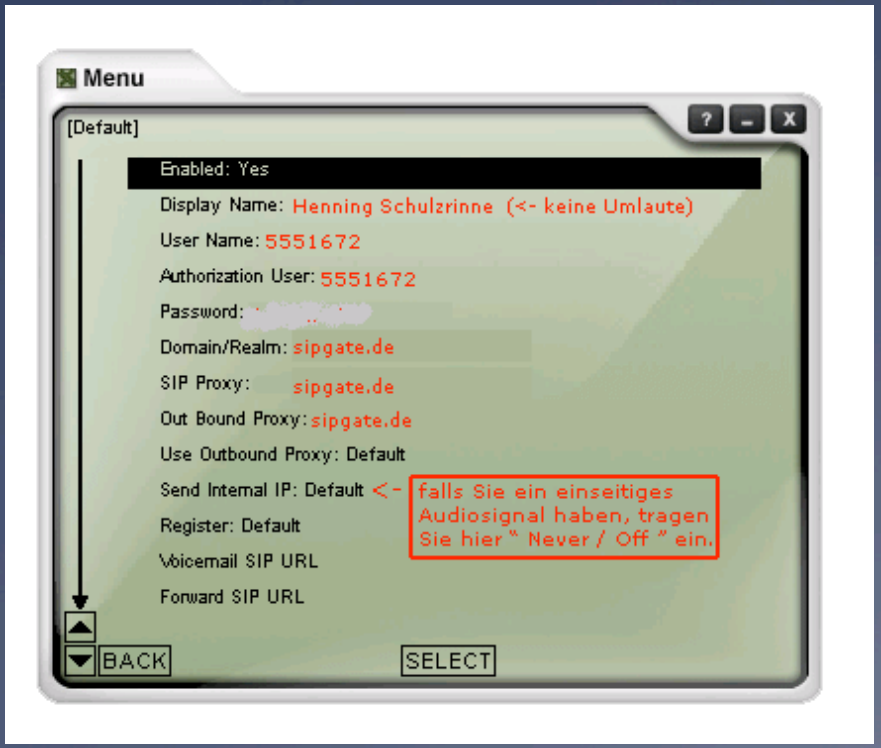

Usability: Interconnected devices

generates TAN

incoming call

time, location

updates location

address book

any weather service
school closings

acoustic alerts

alert, events

Mobile why's

- Not research, but examples of real annoyances
- Why does each mobile device need its own power supply?
- Why do I have to adjust the clock on my camera each time I travel?
- Why do I have to know what my IMAP server is and whether it uses TLS or SSL?
- Why do I have to type in my address book?
- Why do I have to "synchronize" my PDA?
- Why do I have to manually update software?
- Why is connecting a laptop to a projector a gamble?
- Why do we use USB memory sticks when all laptops have 802.11b?

Increasing reliability and usability through end system diagnostics

with Kyung-Hwa Kim, Vishal Singh and Kai Miao

IEEE DLT 2009

Circle of blame

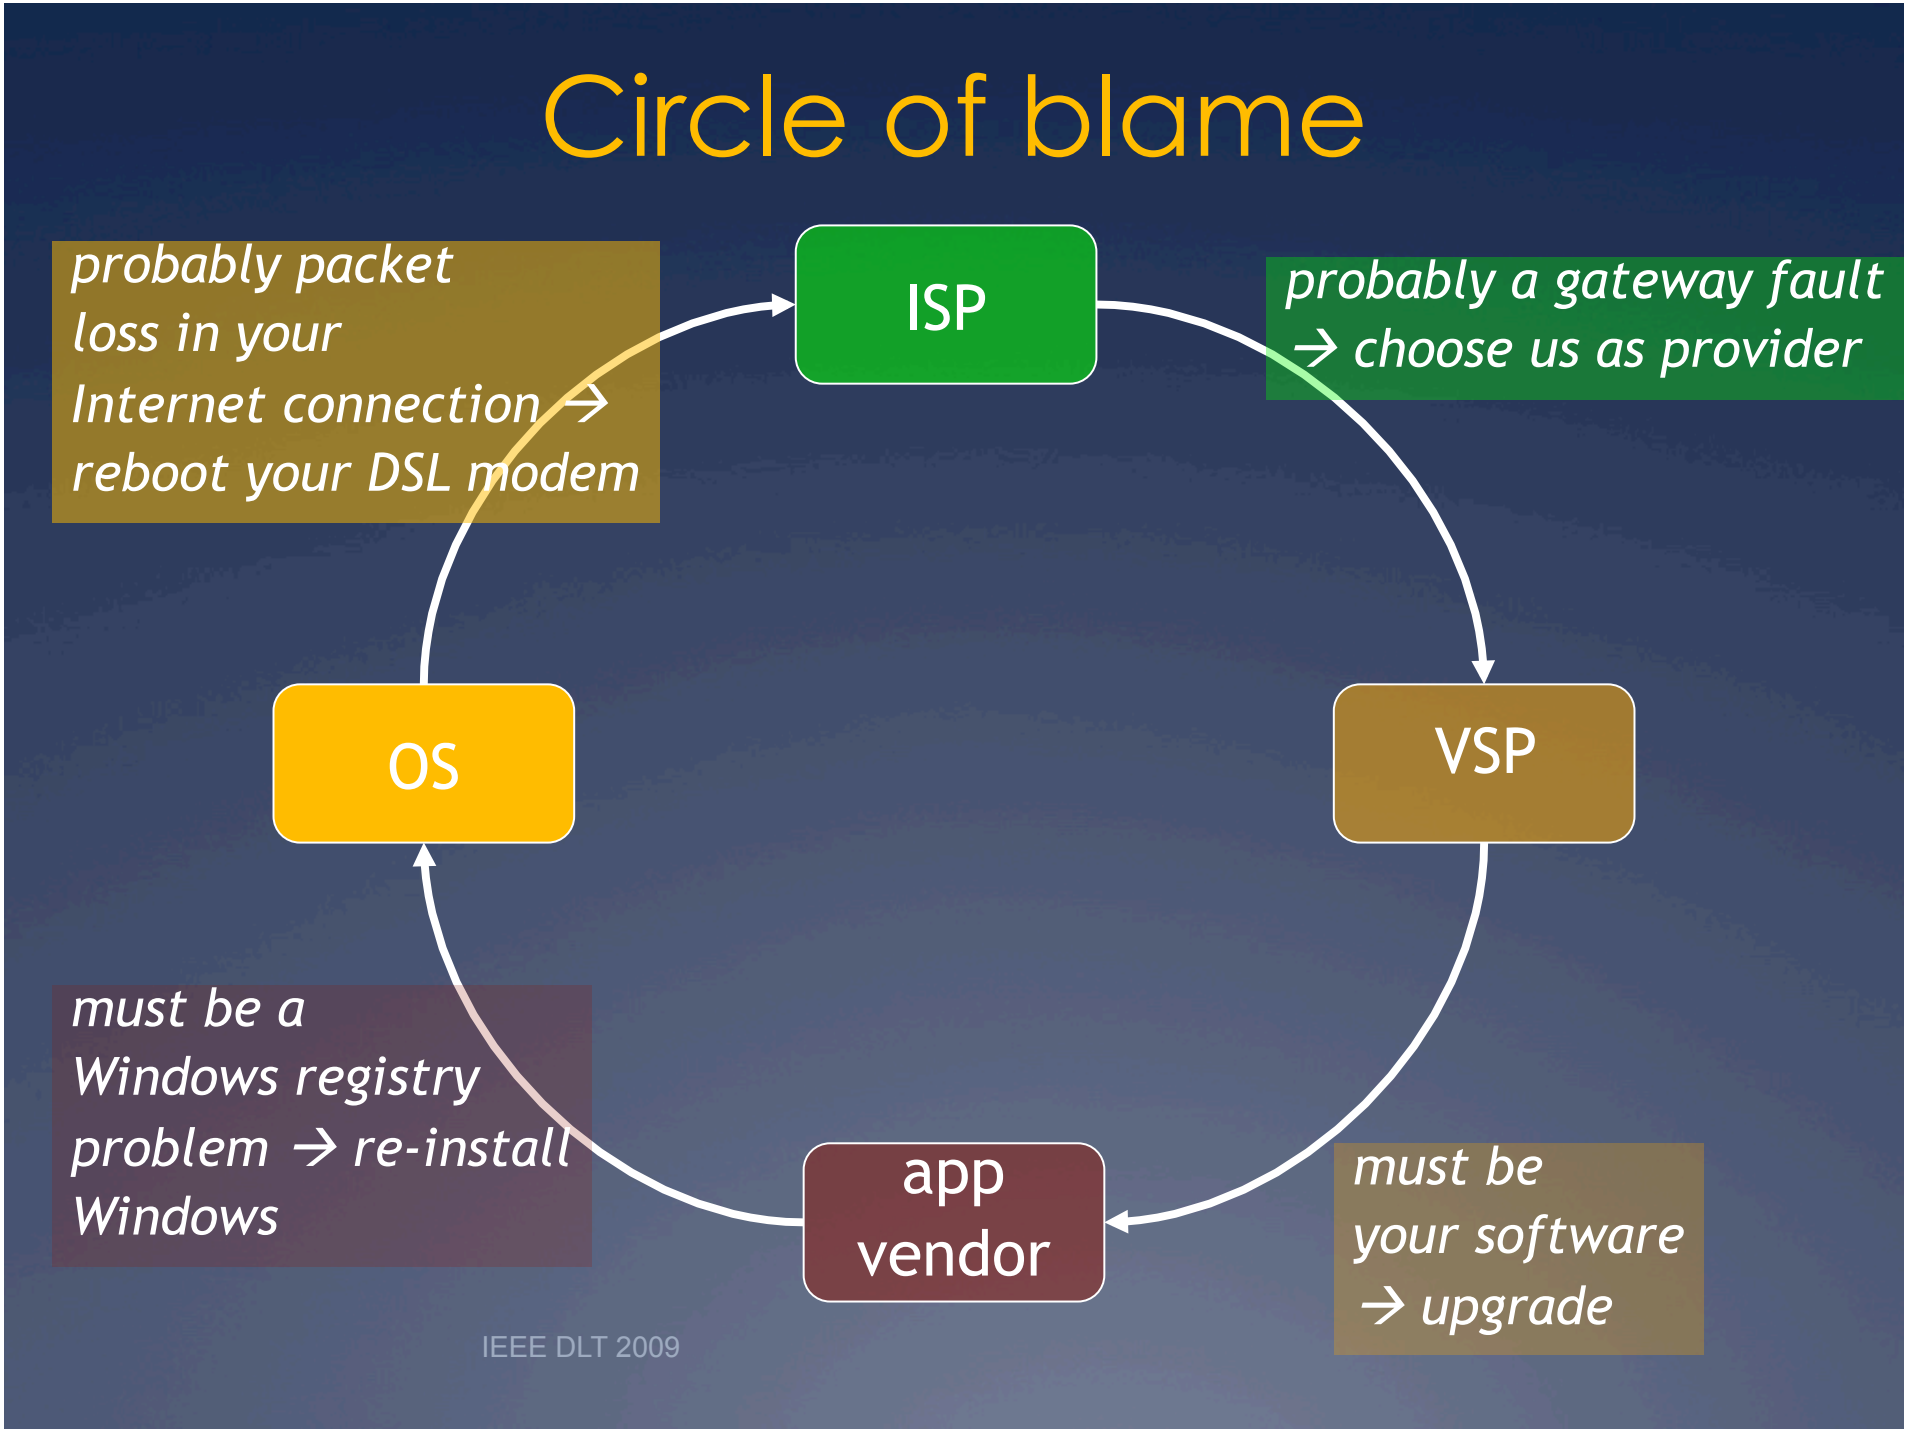

Traditional network management model

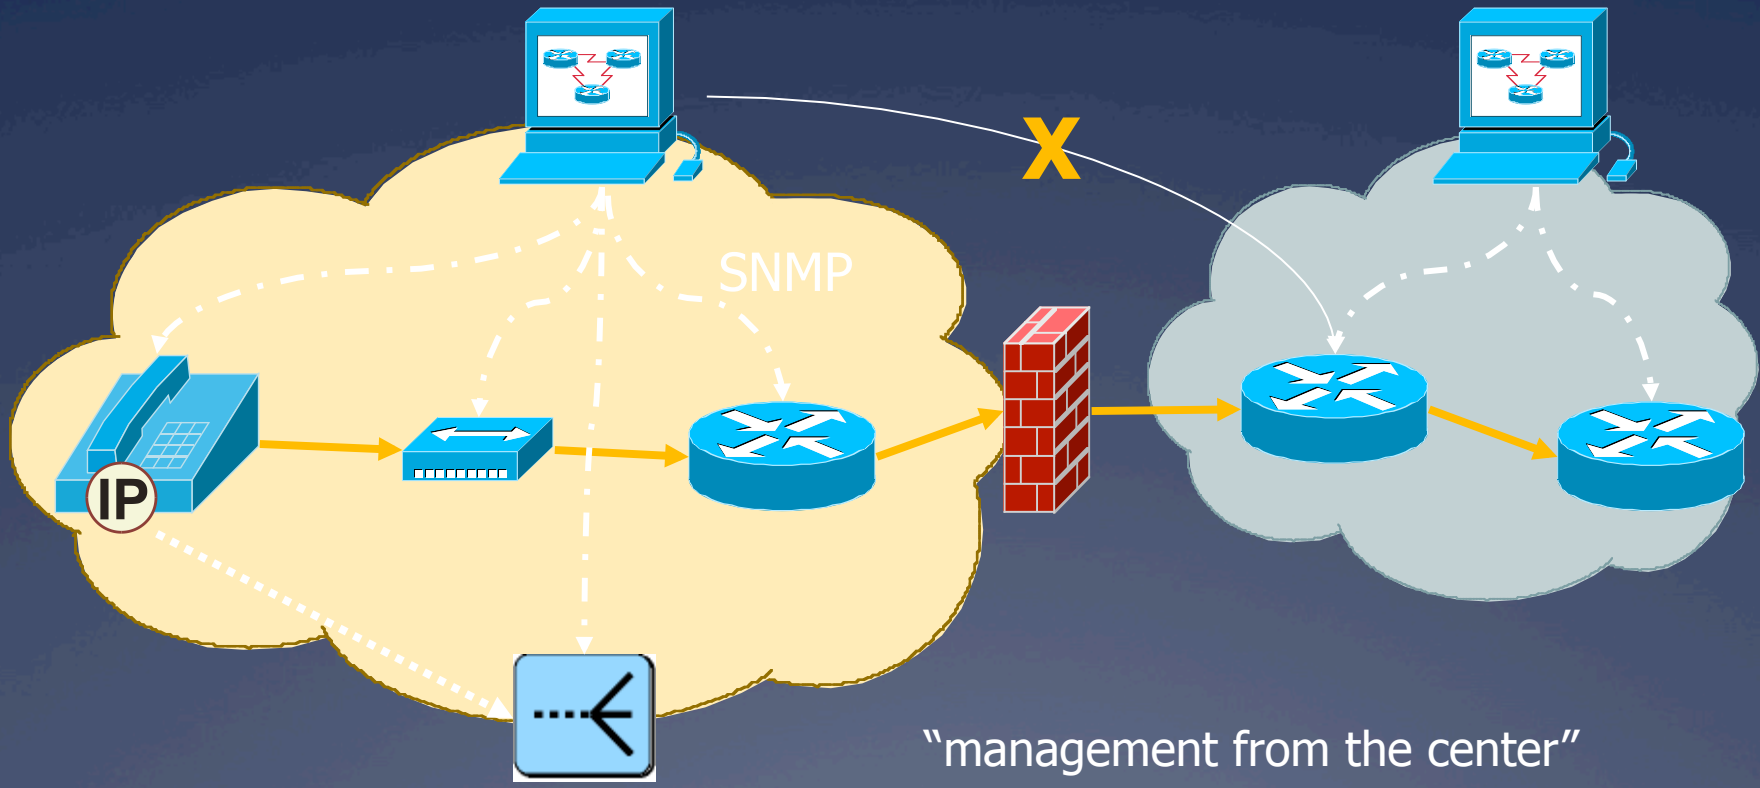

"management from the center"

Old assumptions, now wrong

- Single provider (enterprise, carrier)
  - has access to most path elements
  - professionally managed
- Problems are hard failures & elements operate correctly
  - element failures ("link dead")
  - substantial packet loss
- Mostly L2 and L3 elements
  - switches, routers
  - rarely 802.11 APs
- Problems are specific to a protocol
  - "IP is not working"
- Indirect detection
  - MIB variable vs. actual protocol performance
- End systems don't need management
  - DMI & SNMP never succeeded
  - each application does its own updates

Managing the protocol stack

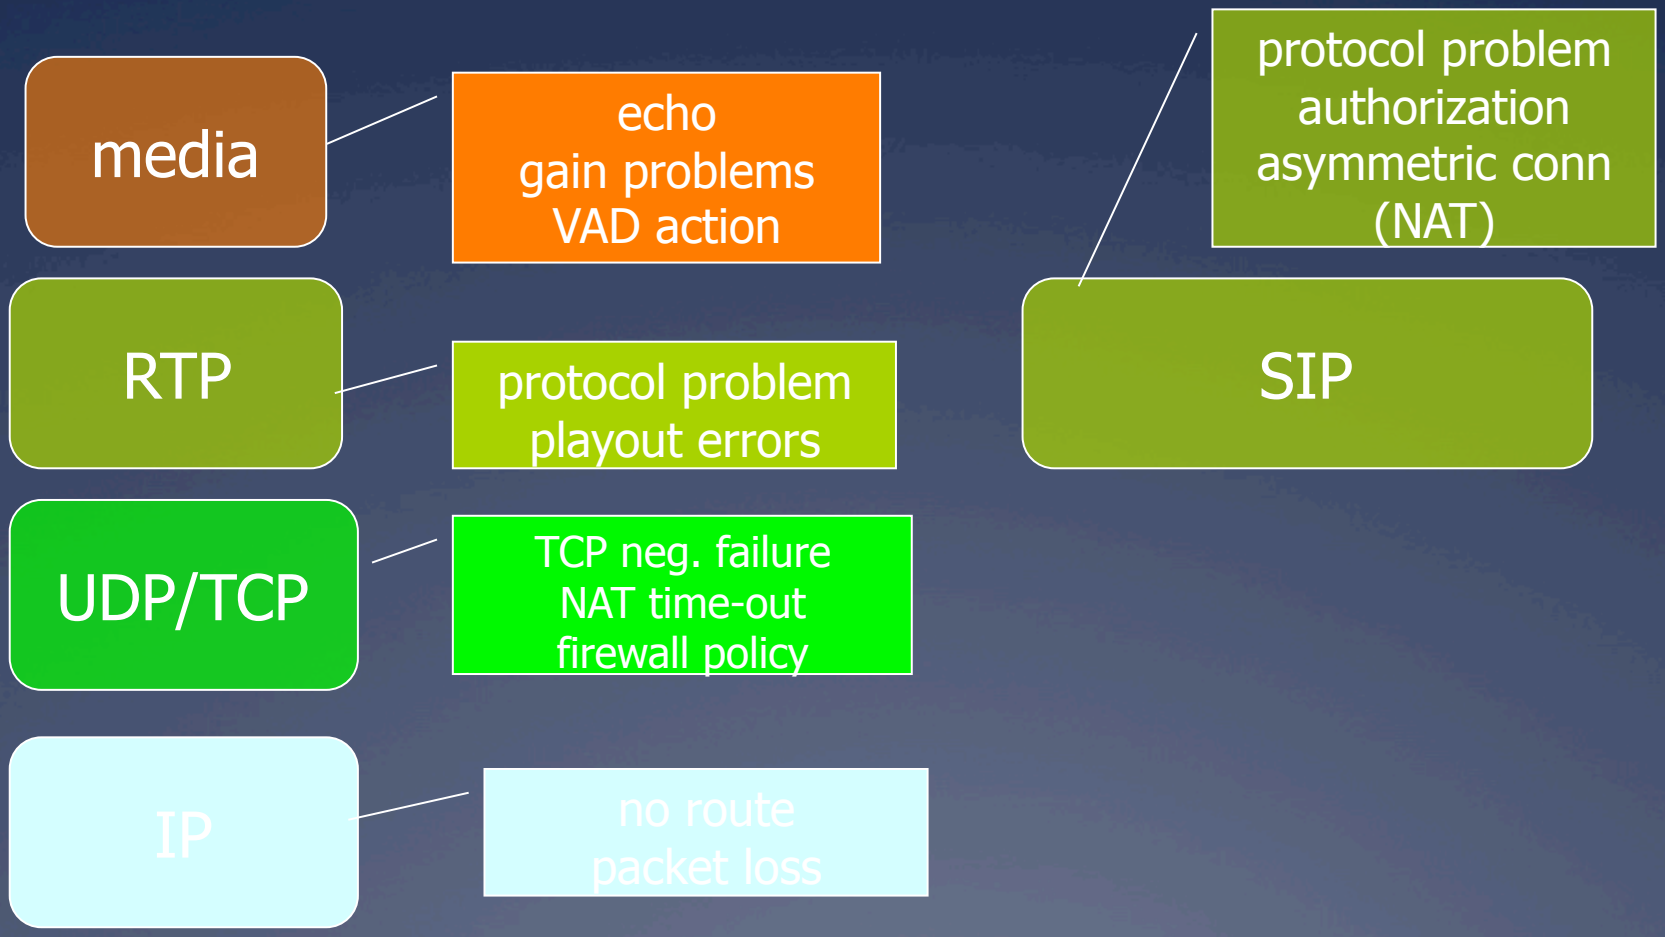

Types of failures

- Hard failures
  - connection attempt fails
  - no media connection
  - NAT time-out
- Soft failures (degradation)
  - packet loss (bursts)
    - access network? backbone? remote access?
  - delay (bursts)
    - OS? access networks?
  - acoustic problems (microphone gain, echo)
  - a software bug (poor voice quality)
    - protocol stack? Codec? Software framework?

Examples of additional problems

- `ping` and `traceroute` no longer works reliably
  - WinXP SP 2 turns off ICMP
  - some networks filter all ICMP messages
- Early NAT binding time-out
  - initial packet exchange succeeds, but then TCP binding is removed ("web-only Internet")
- policy intent vs. failure
  - "broken by design"
  - *"we don't allow port 25"* vs. *"SMTP server temporarily unreachable"*

Fault localization

- Fault classification – local vs. global
  - Does it affect only me or does it affect others also?
- Global failures
  - Server failure
    - e.g., SIP proxy, DNS failure, database failures
  - Network failures
- Local failures
  - Specific source failure
    - node A cannot make call to anyone
  - Specific destination or participant failure
    - no one can make call to node B
  - Locally observed, but global failures
    - DNS service failed, but only B observed it

Do You See What I See?

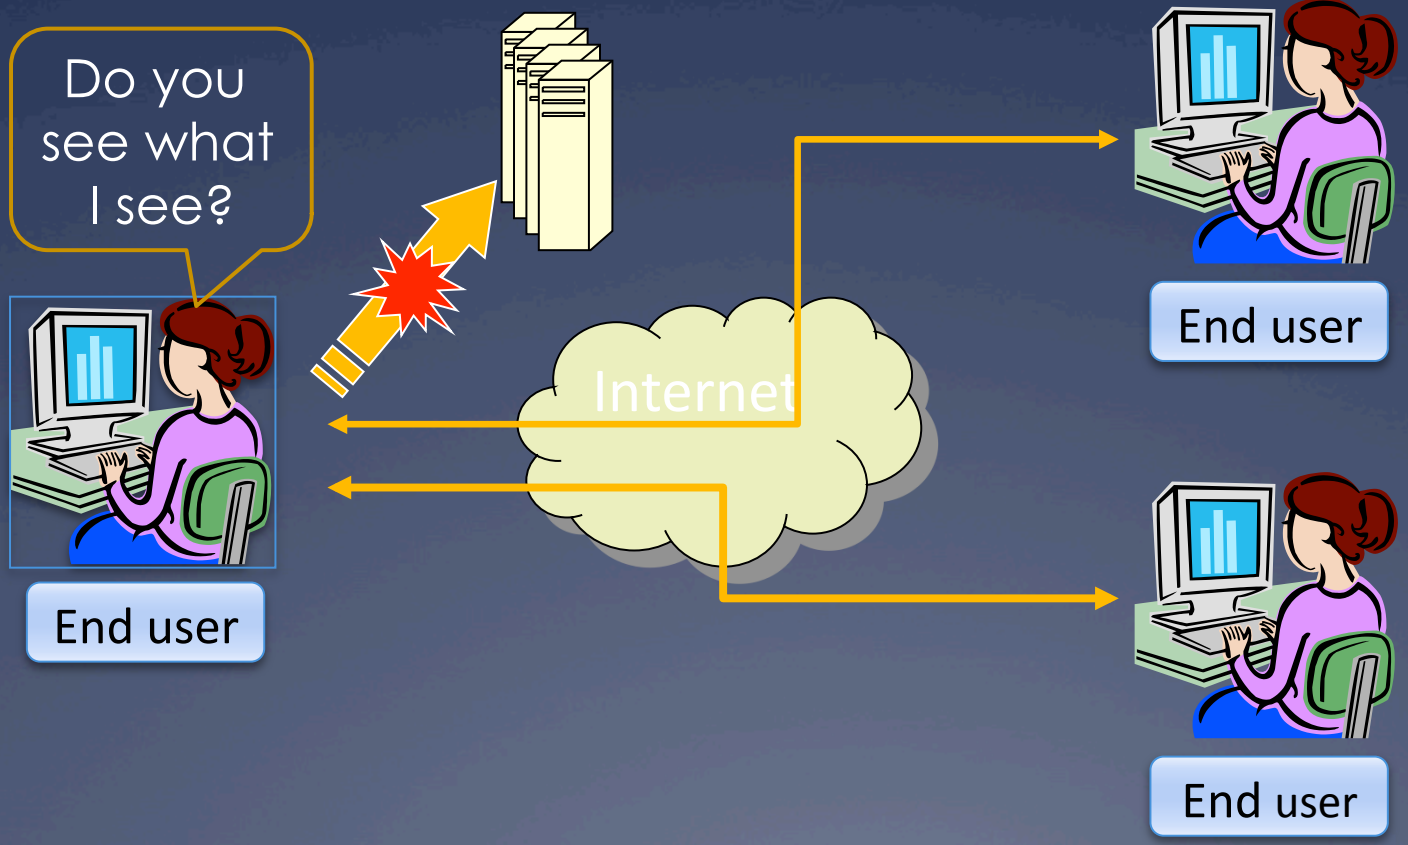

Project: "Do You See What I See?"

DYSWIS

- Each node has a set of active and passive measurement tools
- Use intercept (NDIS, pcap)
  - to detect problems automatically
    - e.g., no response to SIP, HTTP or DNS request
    - deviation from normal protocol exchange behavior
  - gather performance statistics (packet jitter)
  - capture RTCP and similar measurement packets
- Nodes can ask others for their view
  - possibly also dedicated "weather stations"
- Iterative process, leading to:
  - user indication of cause of failure
  - in some cases, work-around (application-layer routing) → TURN server, use remote DNS servers
- Nodes collect statistical information on failures and their likely causes

DYSWIS overview

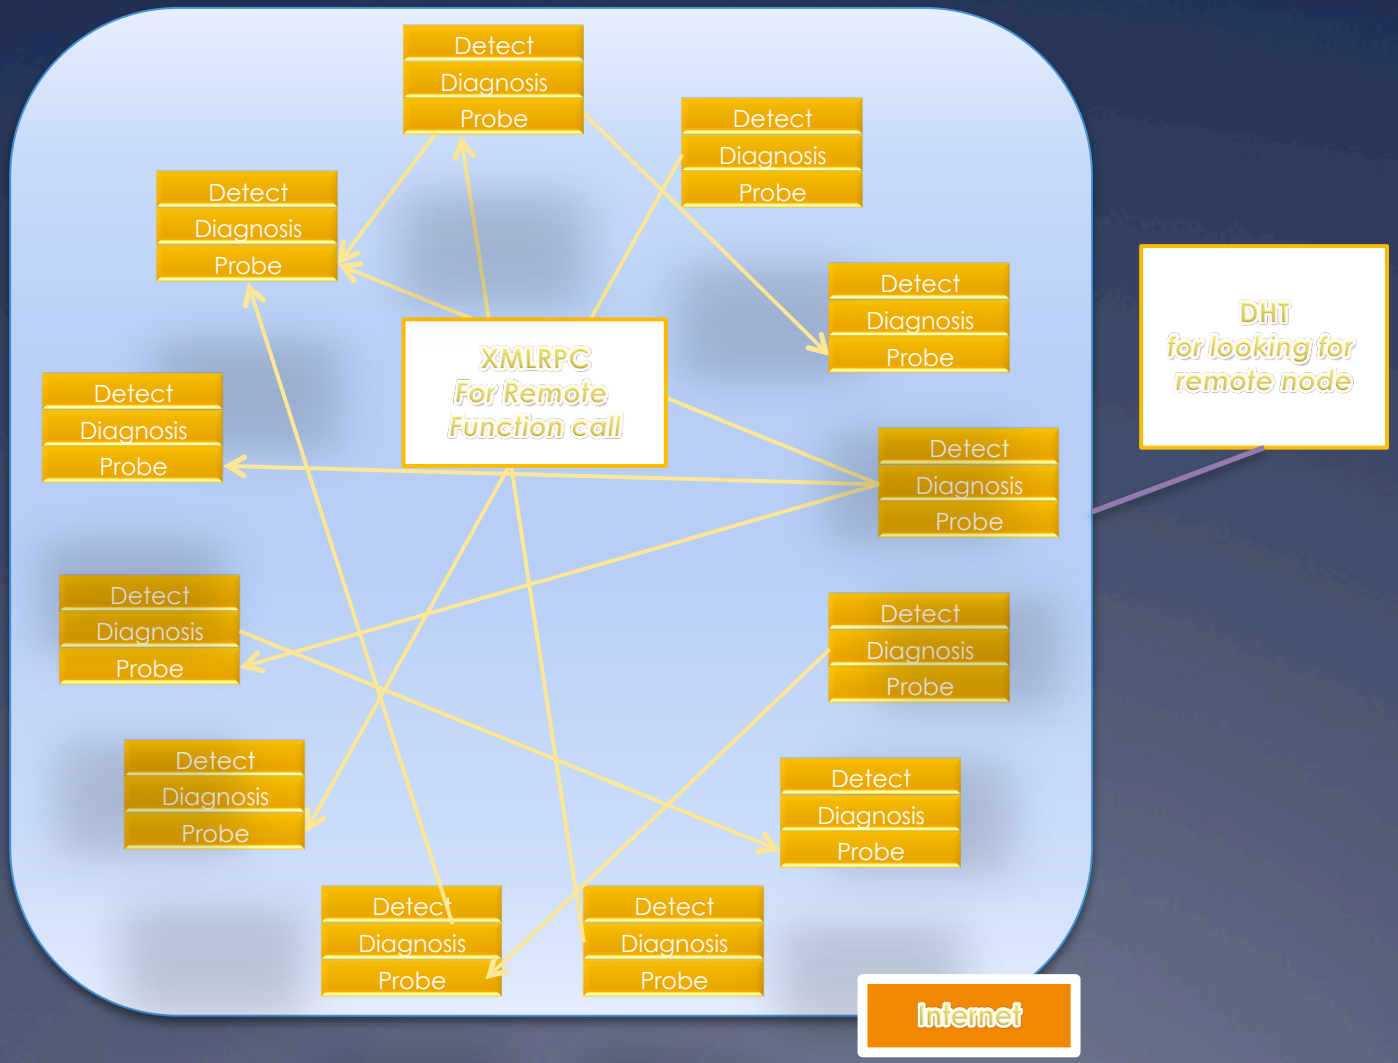

Architecture

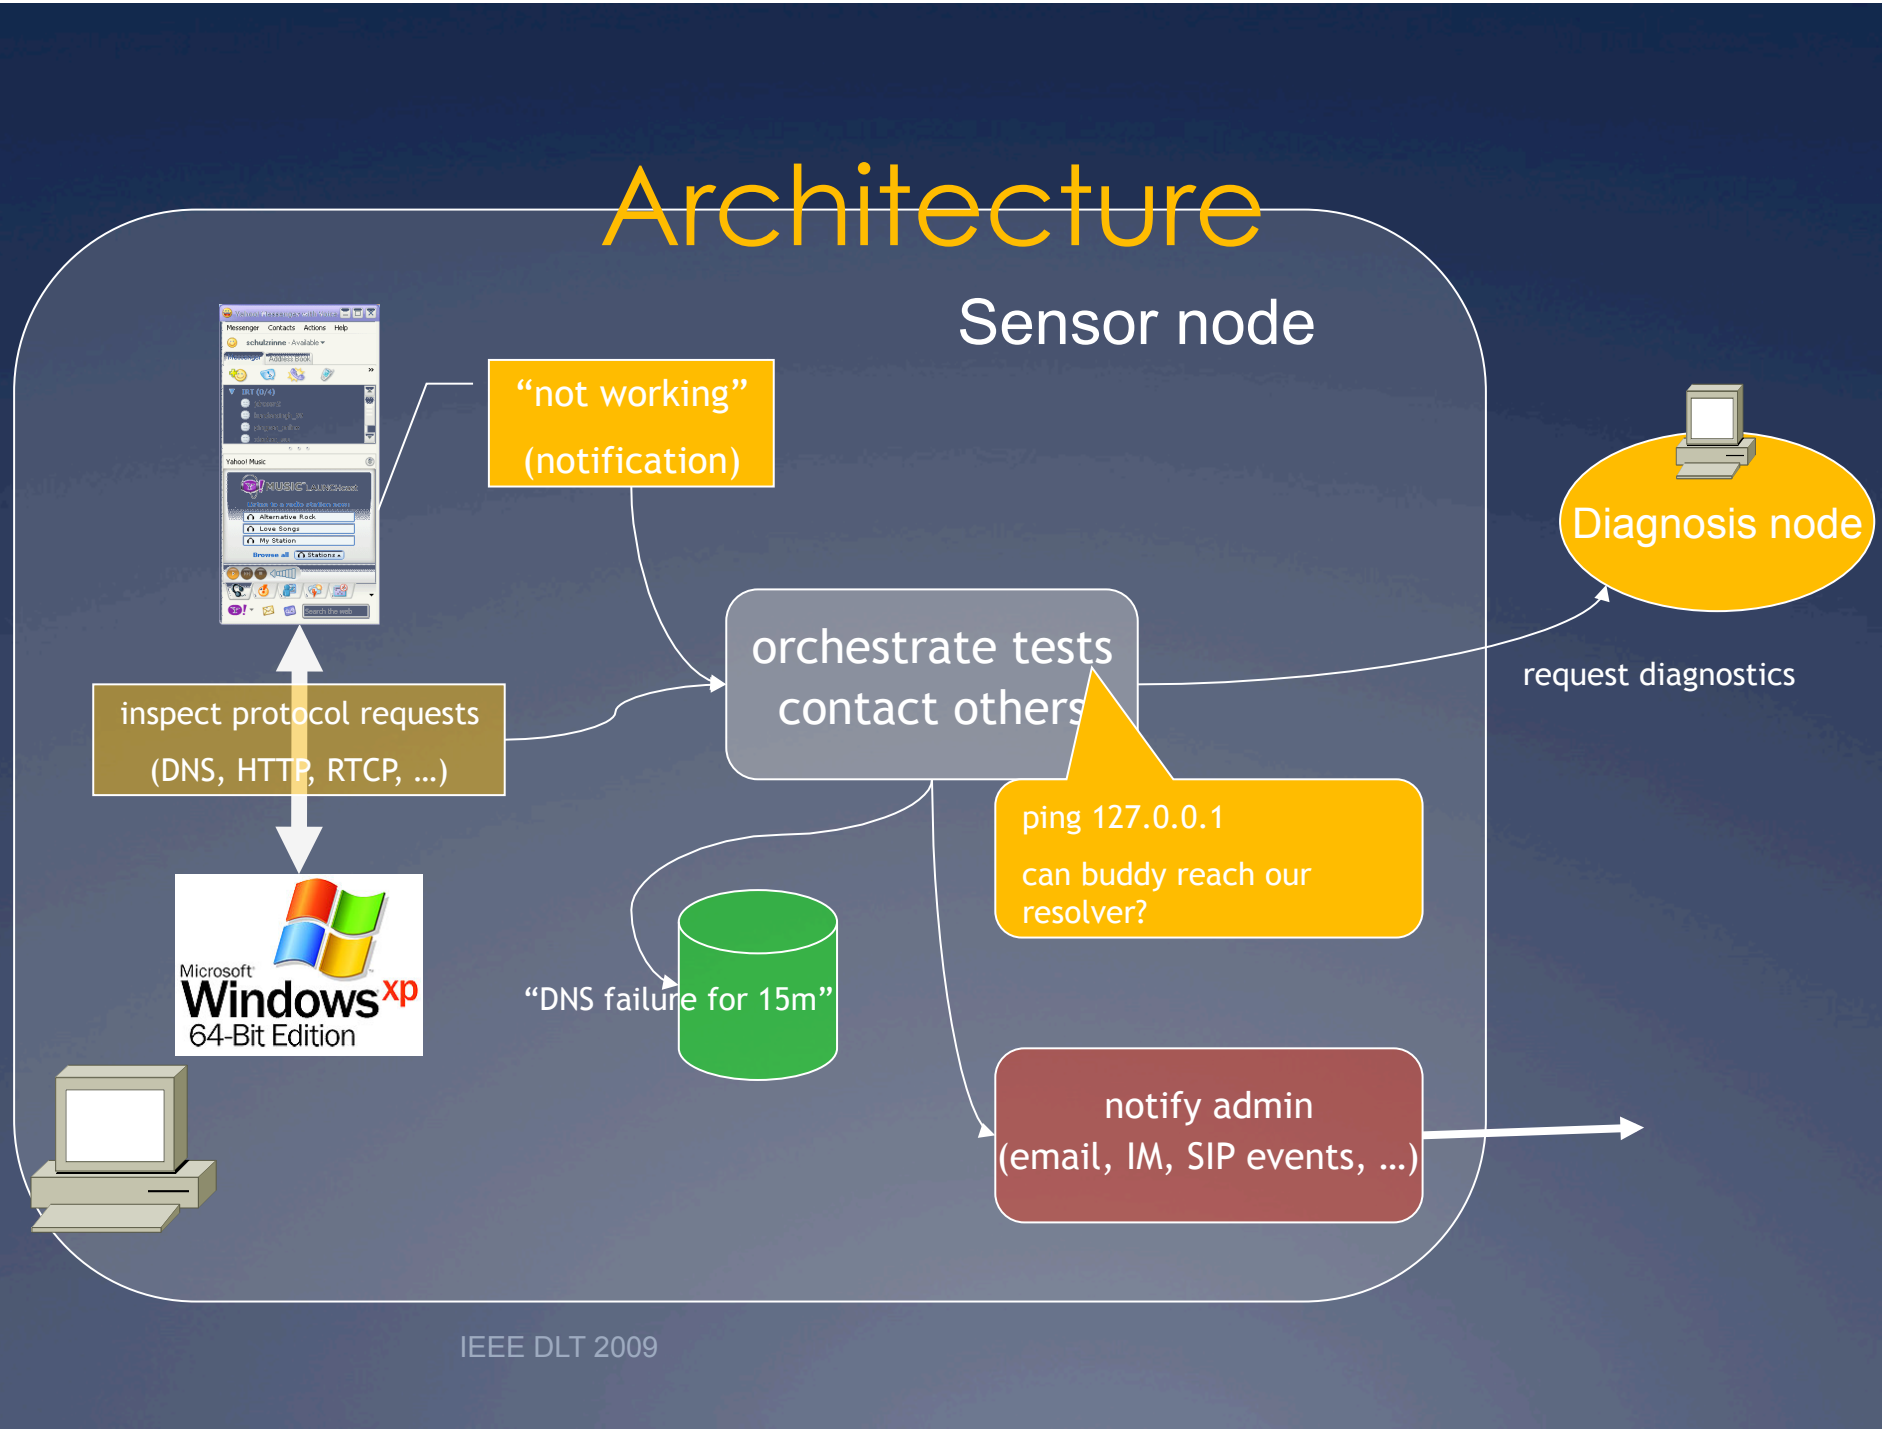

7DS and opportunistic networks: exploring networks beyond the Internet

with Suman Srinivasan, Arezu Moghadam

Contacts are
- opportunistic
- intermittent

802.11 ad-hoc mode
BlueTooth

Web Delivery Model

* 7DS core functionality: Emulation of web content access and e-mail delivery

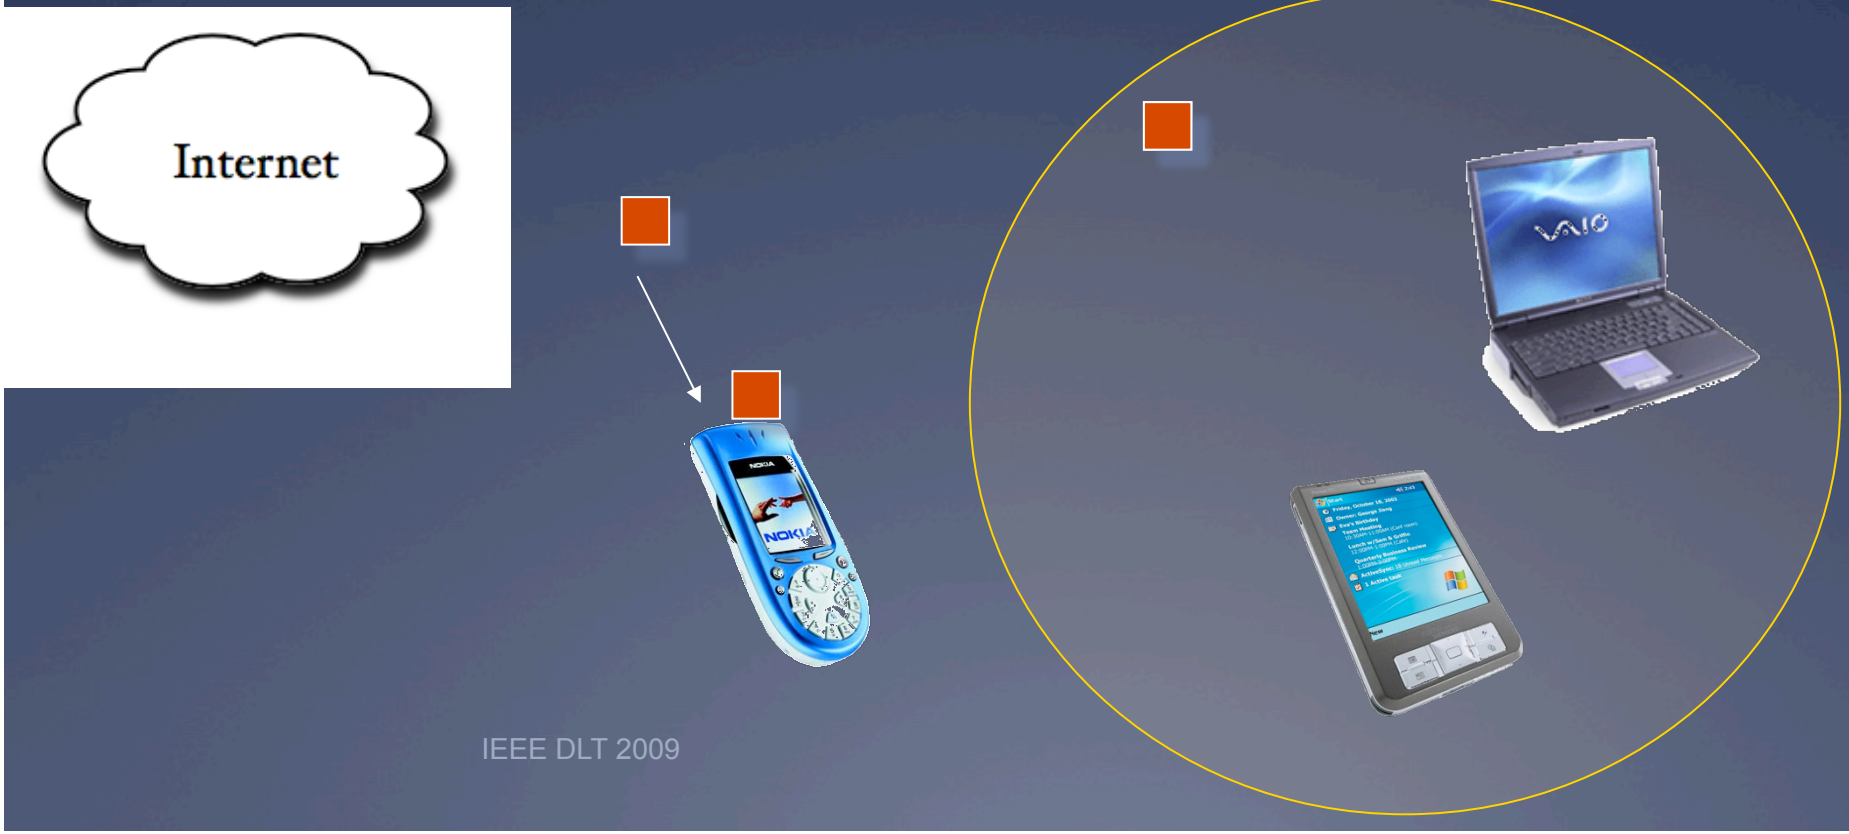

Search Engine

- Provides ability to query self for results
- Searches the cache index using Swish-e library
- Presents results in any of three formats: HTML, XML and plain text
- Similar in concept to Google Desktop

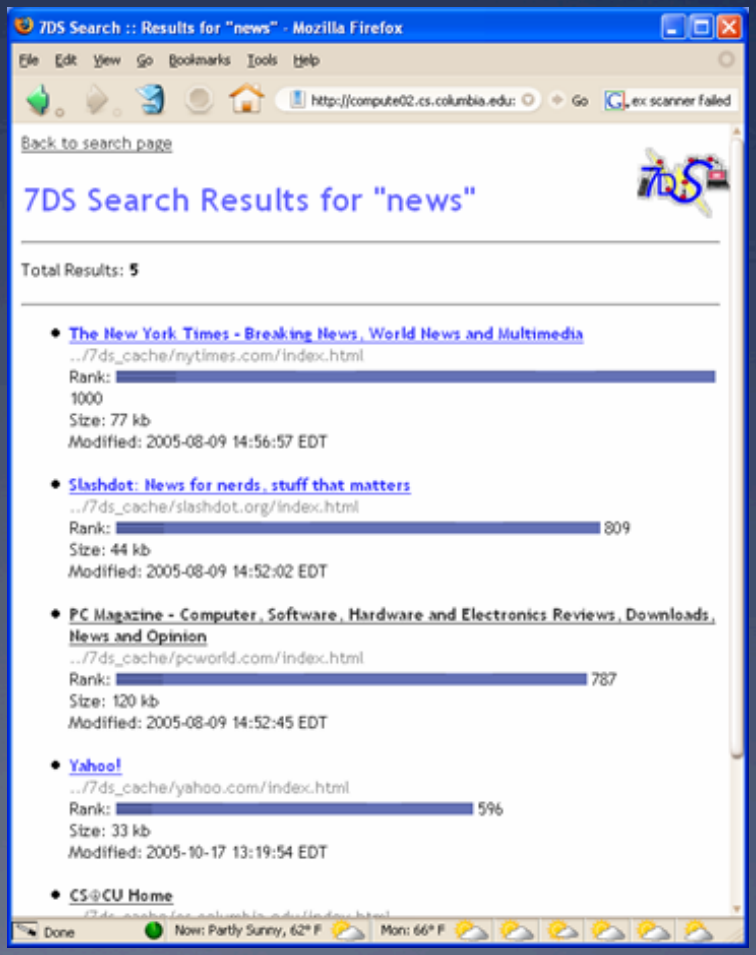

Email exchange

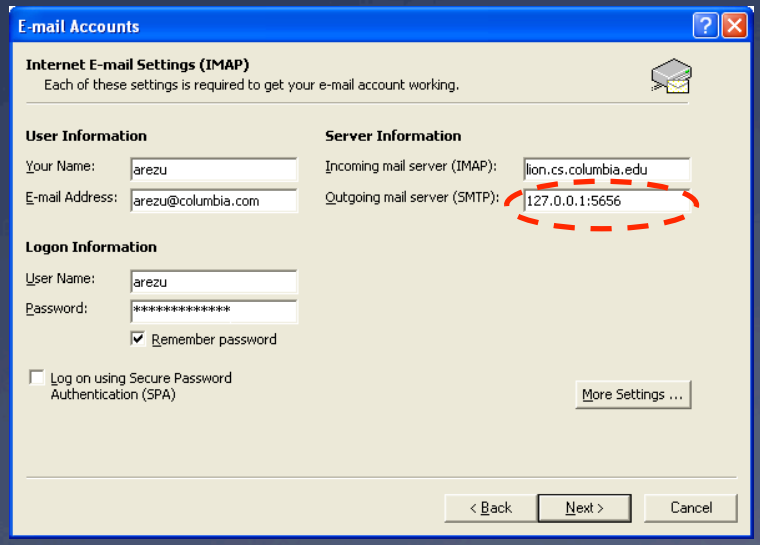

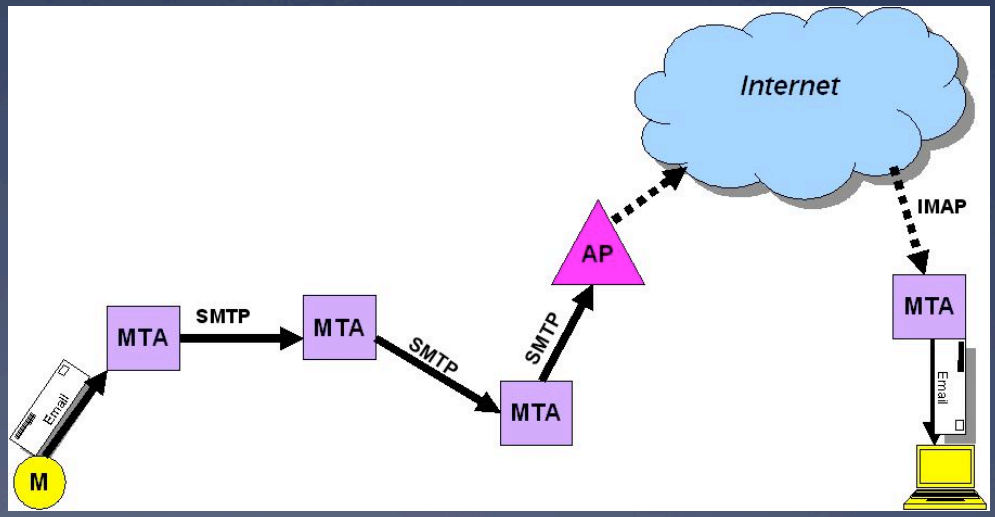

BonAHA framework

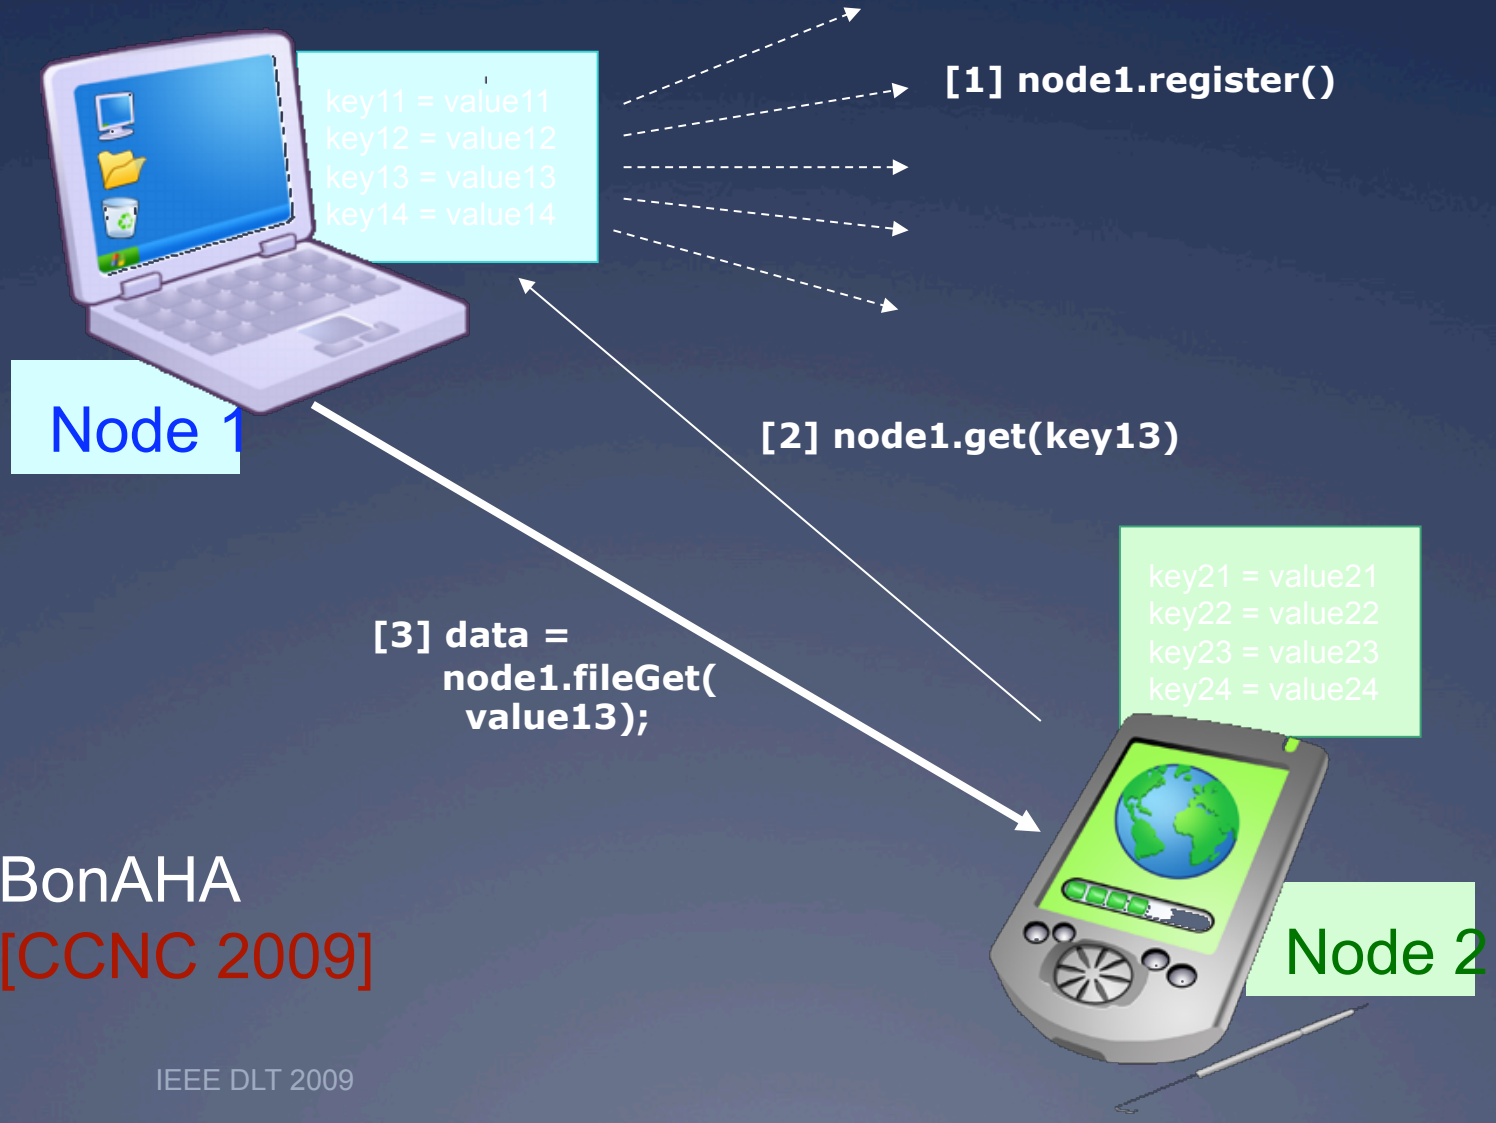

BonAHA
[CCNC 2009]

Bulletin Board System

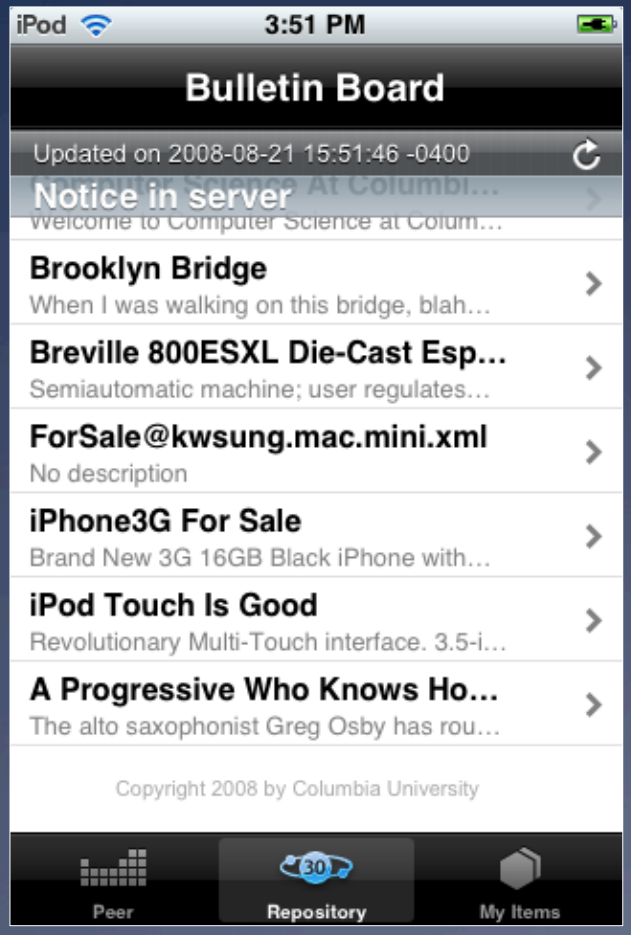

Written in Objective-C, for iPod Touch

Conclusion

- Abandon notion of a clean-slate next-generation Internet
  - that magically fixes all of our problems
- Need for good engineering solutions
  - with user needs, not (just) vendor needs
- Research driven by real, not imagined, problems
  - factor 10 problems: reliability & OpEx
  - more reliability and usability, less sensor networks
- Build a 5-nines network out of unreliable components
- Make network disruptions less visible
- Transition to "self-service" networks
  - support non-technical users, not just NOCs running HP OpenView or Tivoli