VoIP – not your grandmother's phone any more

Henning Schulzrinne
Columbia University

Overview

- VoIP as black phone replacement → interactive communications enabler
- Presence as a service enabler
- Peer-to-peer VoIP
- Integrating VoIP with cellular
- Fax-over-IP

Outline

- VoIP maturing: vision vs. reality
  - overview of protocol zoo
  - presence and location-based services
  - user-programmable services
- New VoIP challenges
  - emergency calling
  - peer-to-peer systems
- The state of SIP standardization
  - trouble in standards land
  - interoperability

The three Cs of Internet applications

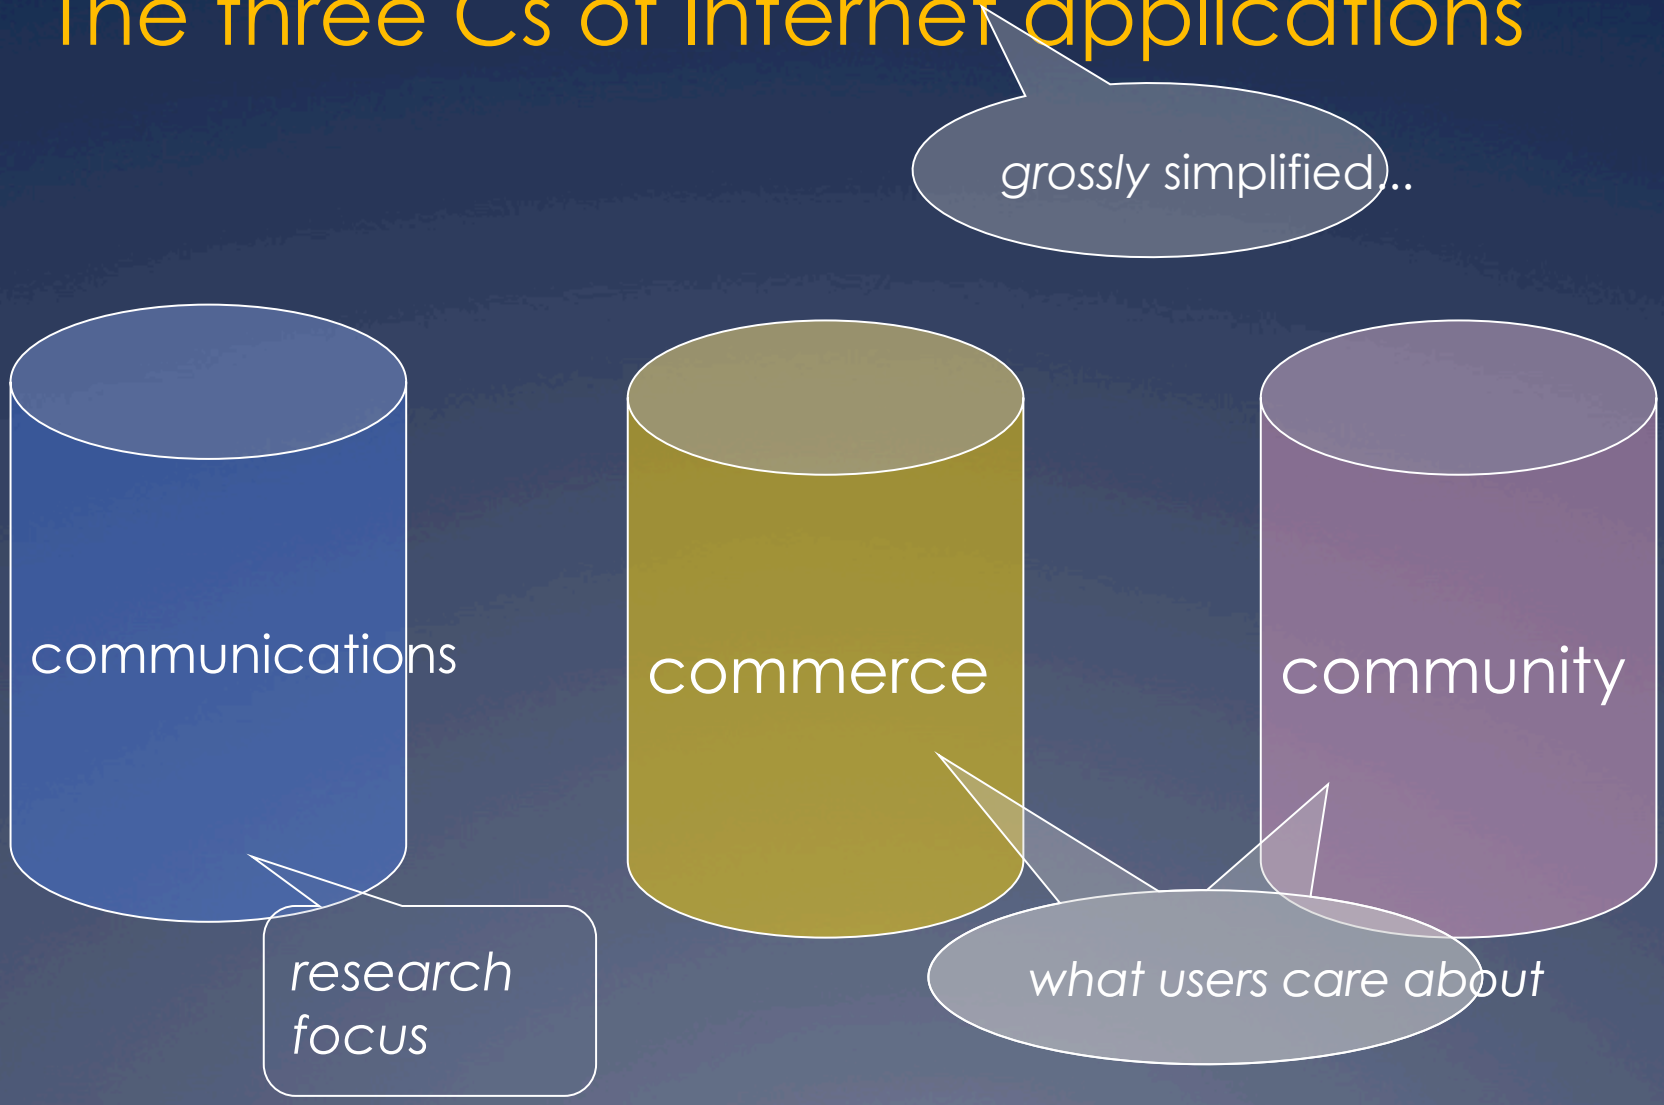

Killer Application

- Carriers looking for killer application
  - justify huge infrastructure investment
  - "video conferencing" (*1950 – †2000)
  - ?
- "There is no killer application"
  - Network television block buster → YouTube hit
  - "Army of one"
  - Users create their own custom applications that are important to them
  - Little historical evidence that carriers (or equipment vendors) will find that application if it exists
- Killer app = application that kills the carrier

Collaboration in transition

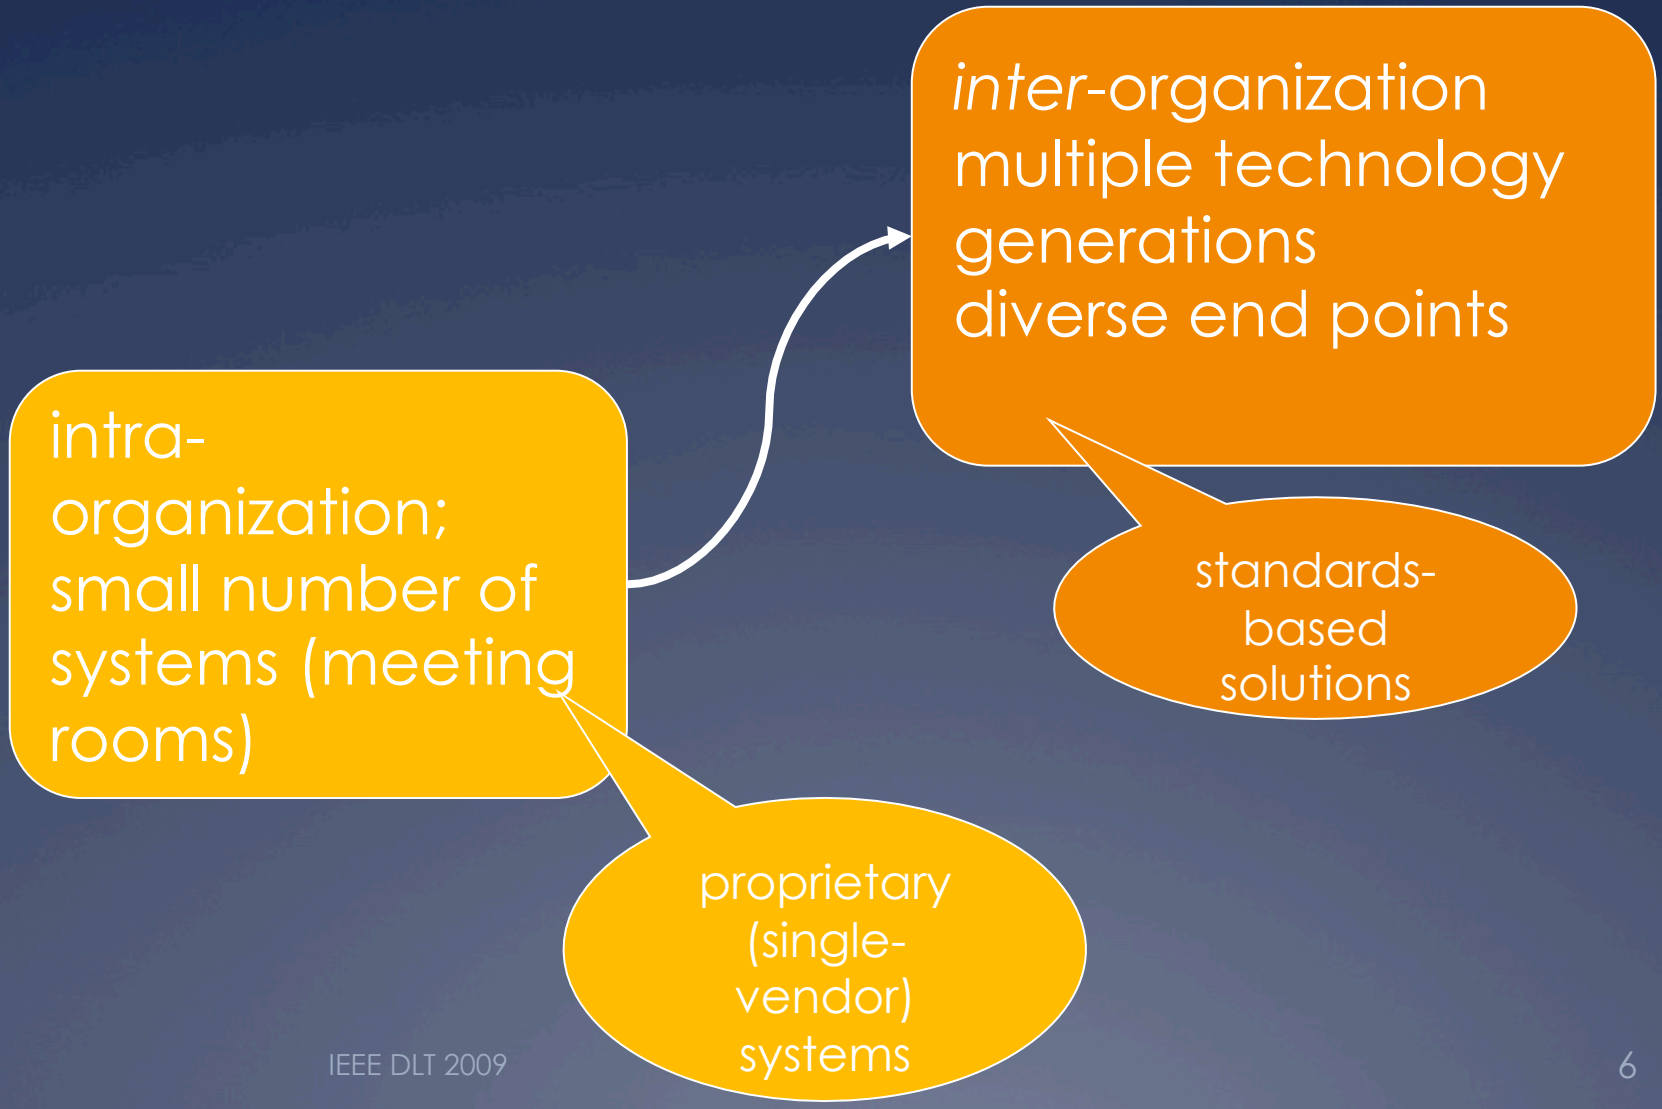

Evolution of VoIP

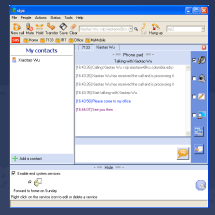

long-distance calling, ca. 1930

"amazing – the phone rings"

"does it do call transfer?"

catching up with the digital PBX

"How can I make it stop ringing?"

going beyond the black phone

"Can it really replace the phone system?"

replacing the global phone system

1996-2000 2000-2003 2004-2005 2006-

IETF VoIP & presence efforts

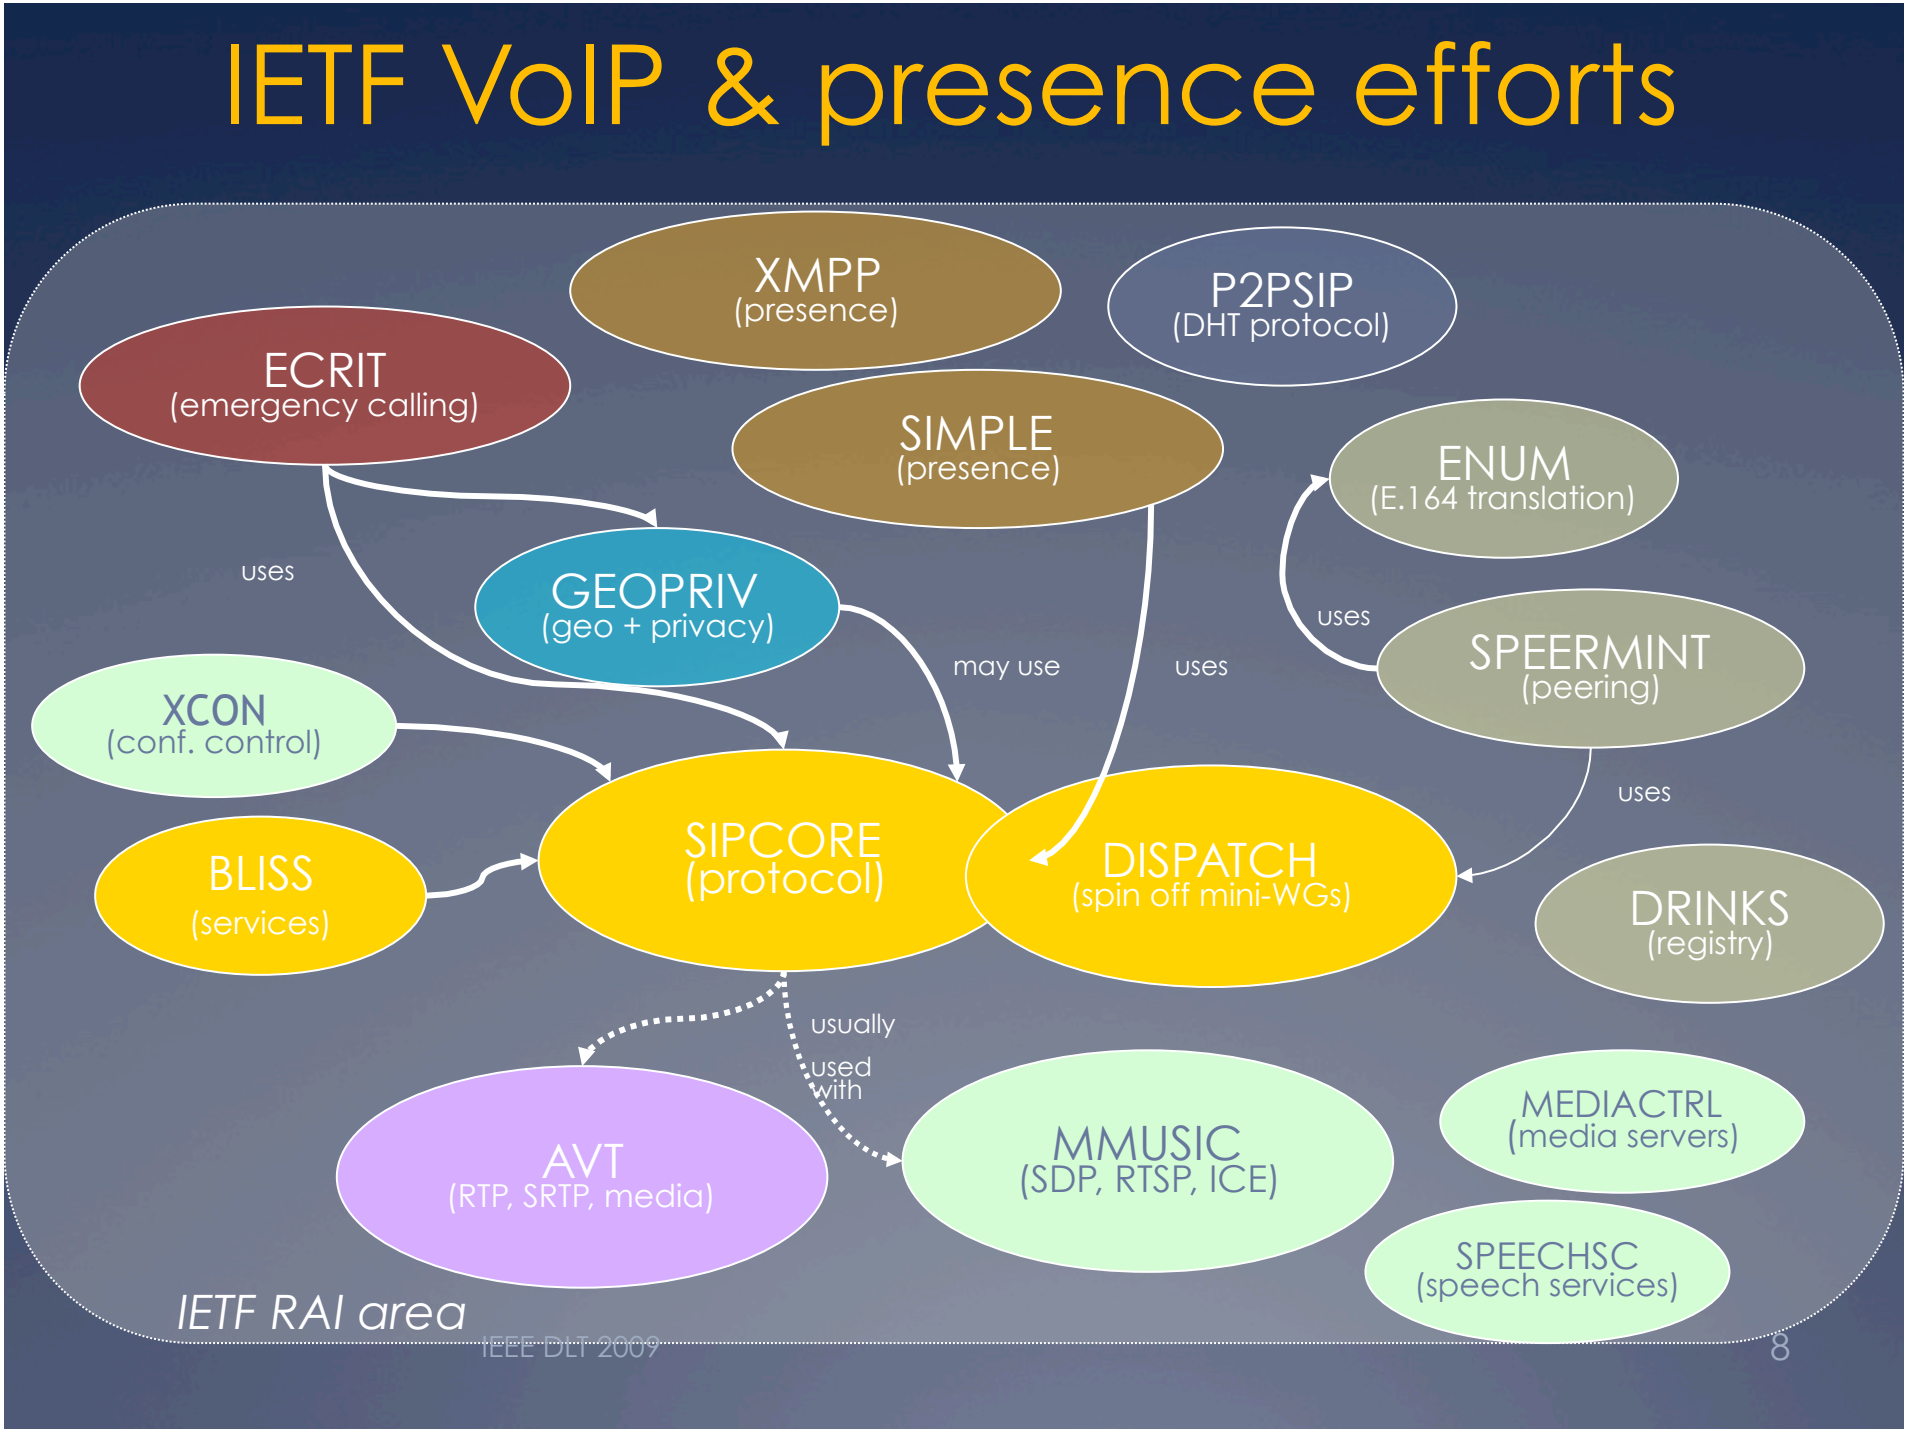

Old vs. new

	old reality	new idea	new reality
service provider	**ILEC**, *CLEC*	email-like, run by enterprise, homes	E.164-driven; MSOs, some ILECs, Skype, European SIP providers, Vonage, ~~SunRocket~~
media	4 kHz audio	wideband audio, video, IM, shared apps, …	4 kHz audio
services	CLASS (CLID, call forwarding, 3-way calling, ...)	user-created services (web model) presence	still CLASS GrandCentral The New Way to Use Your Phones jaiku
user IDs	E.164	email-like	E.164 IM handles

SIP overview

Internet services – the missing entry

Service/ delivery	synchronous	asynchronous
push	instant messaging presence event notification session setup media-on-demand	messaging
pull	data retrieval file download remote procedure call	peer-to-peer file sharing

Filling in the protocol gap

Service/ delivery	synchronous	asynchronous
push	SIP RTSP, RTP	SMTP
pull	HTTP ftp SunRPC, Corba, SOAP	(not yet standardized)

SIP as service enabler

- Rendezvous protocol
  - lets users find each other by only knowing a permanent identifier
- Mobility enabler:
  - personal mobility
    - one person, multiple terminals
  - terminal mobility
    - one terminal, multiple IP addresses
  - session mobility
    - one user, multiple terminals in sequence or in parallel
  - service mobility
    - services move with user

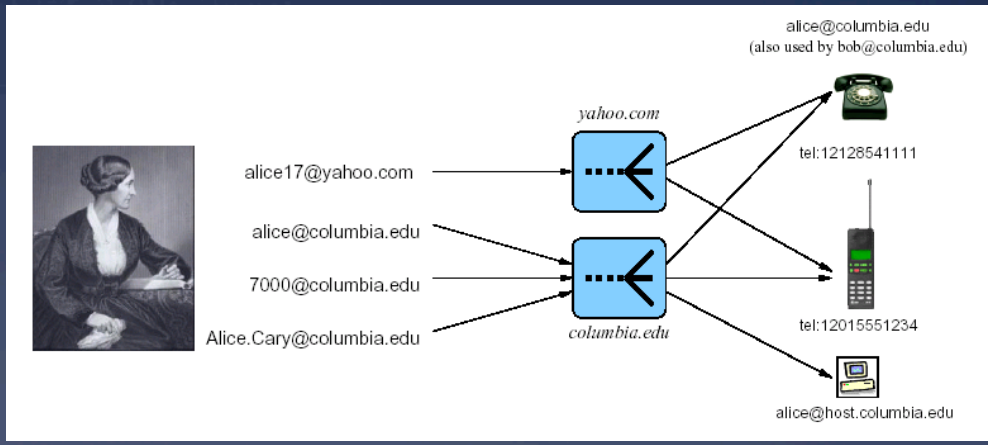

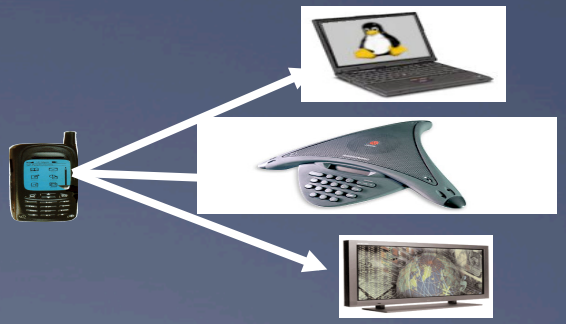

What is SIP?

- Session Initiation Protocol → protocol that establishes, manages (multimedia) sessions
  - also used for IM, presence & event notification
  - uses SDP to describe multimedia sessions
- Developed at Columbia U. (with others)
- Standardized by
  - IETF (RFC 3261-3265 et al)
  - 3GPP (for 3G wireless)
  - PacketCable
- About 100 companies produce SIP products
- Microsoft's Windows Messenger (≥4.7) includes SIP

Philosophy

- Session establishment & event notification
- Any session type, from audio to circuit emulation
- Provides application-layer *anycast* service
- Provides terminal and session *mobility*
- Based on HTTP in syntax, but different in protocol operation
- Peer-to-peer system, with optional support by proxies
  - even stateful proxies only keep transaction state, not call (session, dialogue) state
  - transaction: single request + retransmissions
  - proxies *can be* completely stateless

Basic SIP message flow

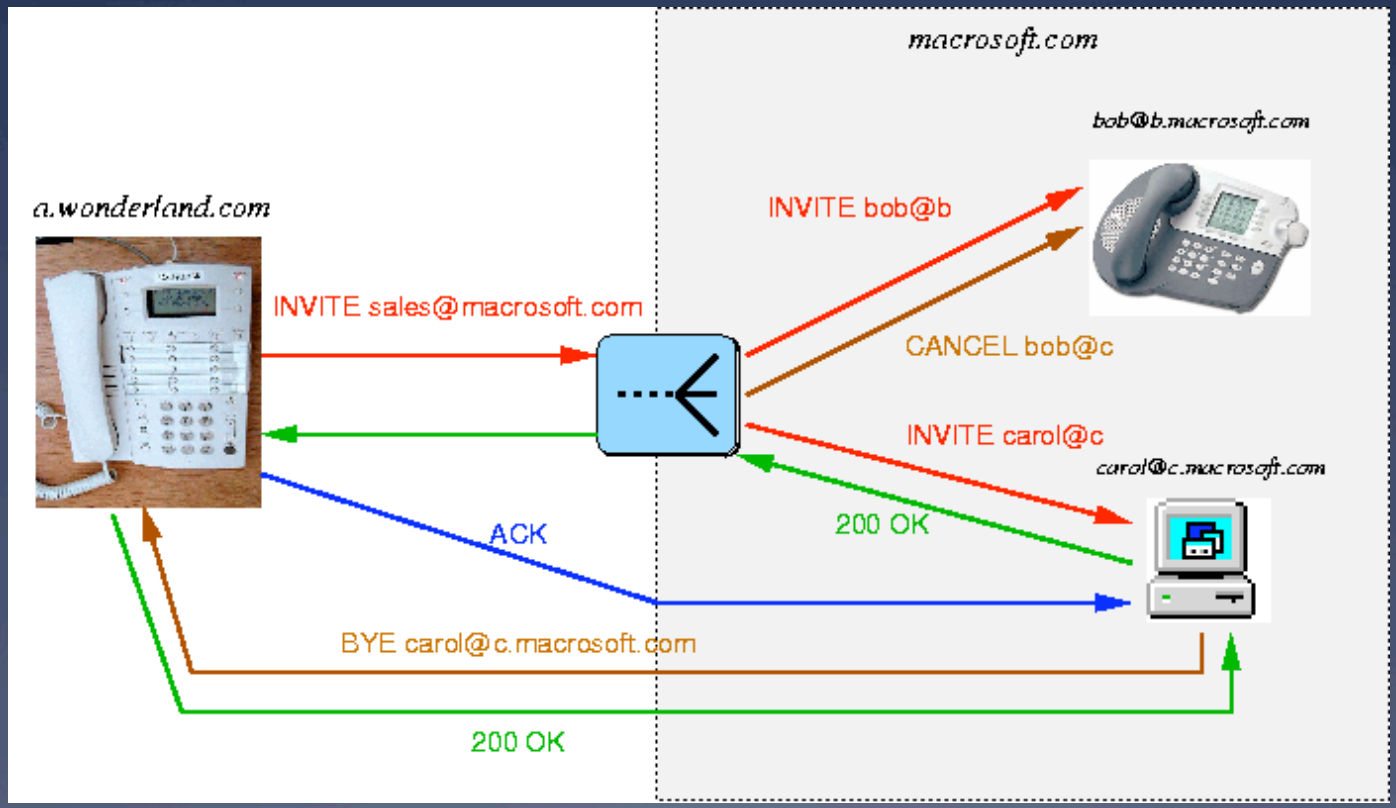

SIP trapezoid

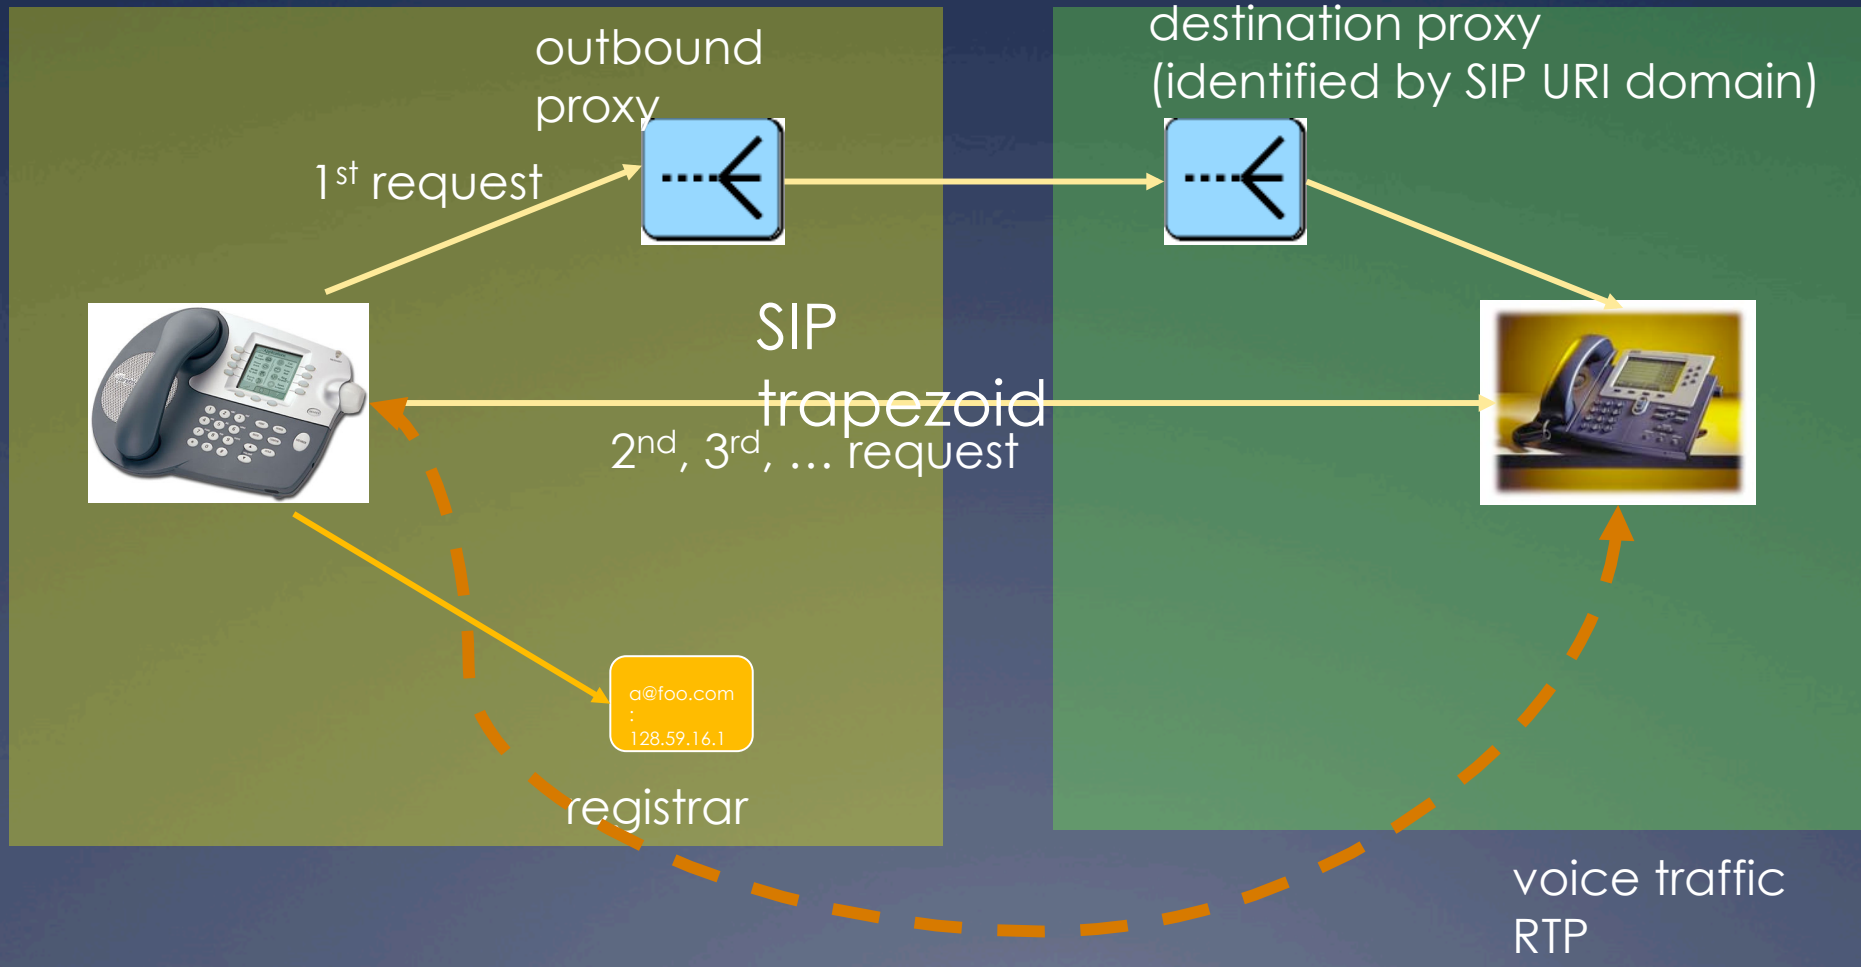

SIP message format

PSTN vs. Internet Telephony

PSTN:

Internet telephony:

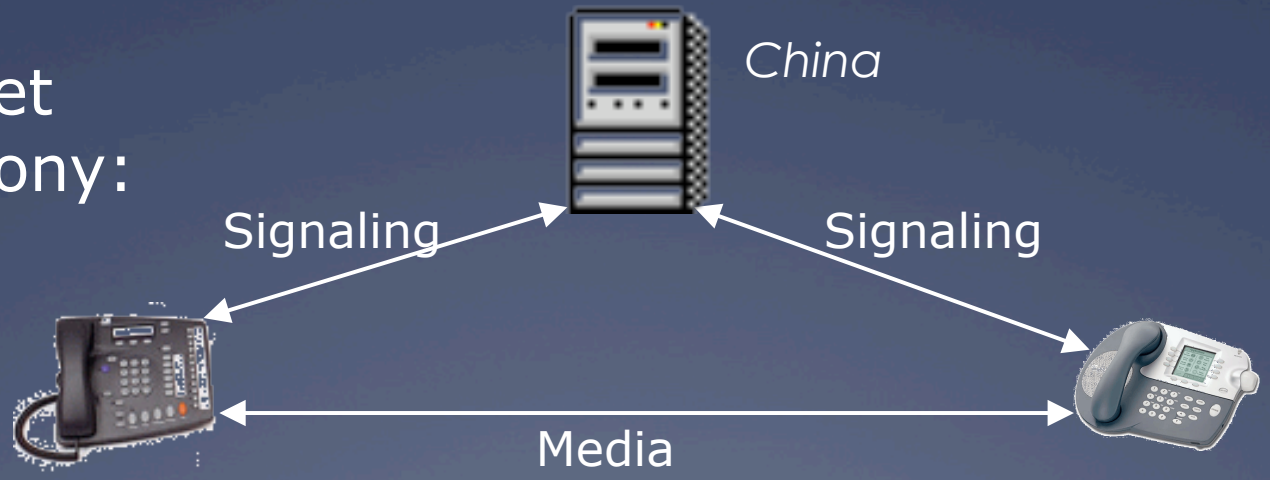

Belgian customer, currently visiting US

Australia

SIP addressing

- Users identified by SIP or tel URIs
  - `sip:alice@example.com`
- tel: URIs describe E.164 number, not dialed digits (RFC 2806bis)
- tel URIs → SIP URIs by outbound proxy
- A person can have any number of SIP URIs
- The same SIP URI can reach many different phones, in different networks
  - sequential & parallel forking
- SIP URIs can be created dynamically:
  - GRUUs
  - conferences
  - device identifiers (sip:foo@128.59.16.15)
- Registration binds SIP URIs (e.g., device addresses) to SIP "address-of-record" (AOR)

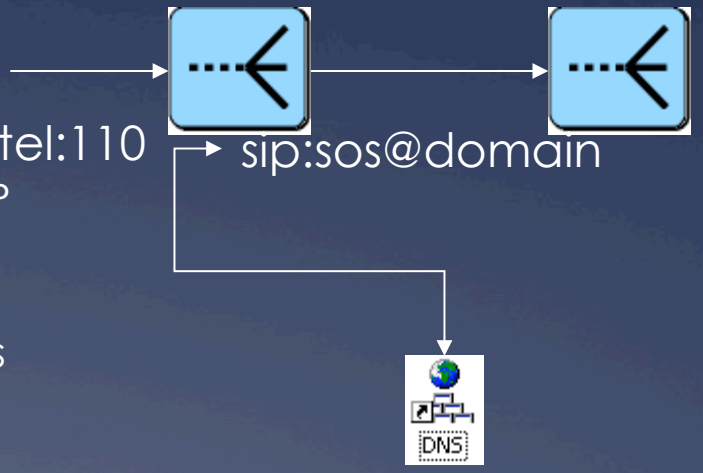

3G Architecture (Registration)

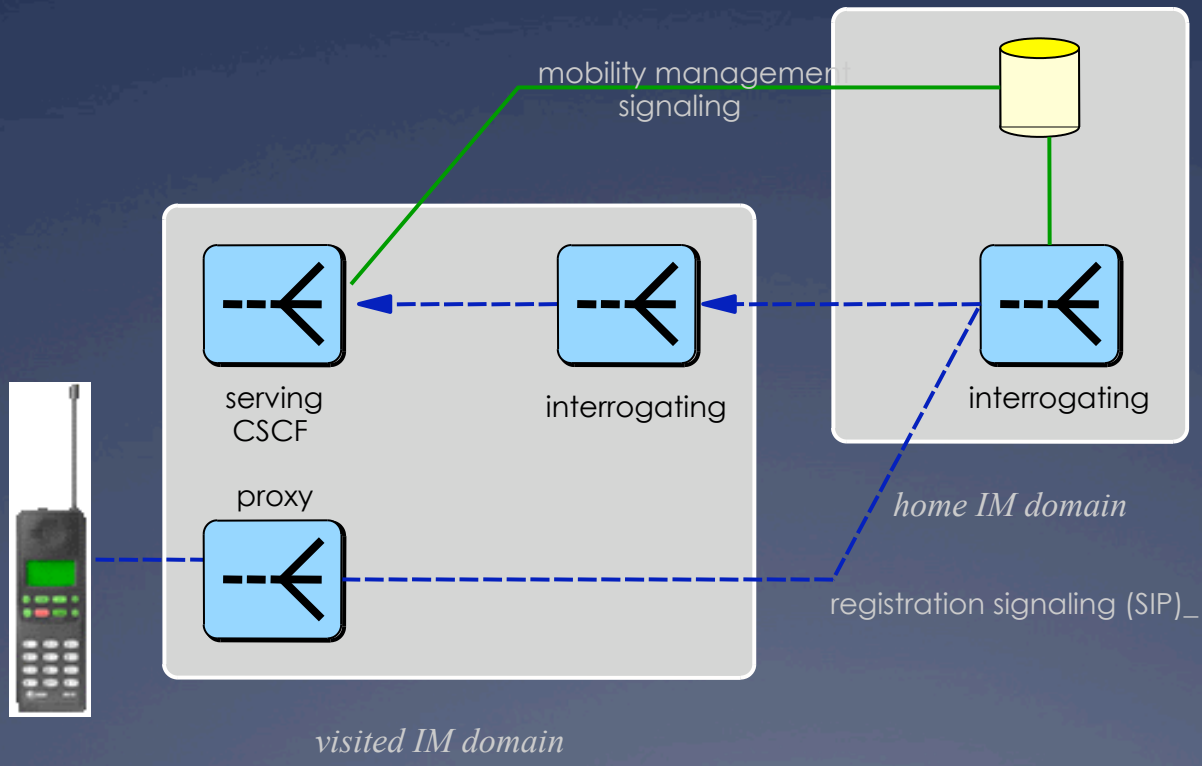

Presence & events

We need glue!

- Lots of devices and services
  - cars
  - household
  - environment
- Generally, stand-alone
  - e.g., GPS can't talk to camera
- Home
  - home control networks have generally failed
    - cost, complexity
- Environment
  - "Internet of things"
  - tag bus stops, buildings, cars, ...

Left to do: event notification

- notify (small) group of users when something of interest happens
  - presence = change of communications state
  - email, voicemail alerts
  - environmental conditions
  - vehicle status
  - emergency alerts
- kludges
  - HTTP with pending response
  - inverse HTTP --> doesn't work with NATs
- Lots of research (e.g., SIENA)
- IETF efforts starting
  - SIP-based
  - XMPP

Context-aware communication

- context = "the interrelated conditions in which something exists or occurs"
- anything known about the participants in the (potential) communication relationship
- both at caller and callee

time	CPL
capabilities	caller preferences
location	location-based call routing location events
activity/availability	presence
sensor data (mood, bio)	privacy issues similar to location data

The role of presence

- Guess-and-ring
  - high probability of failure:
    - "telephone tag"
    - inappropriate time (call during meeting)
    - inappropriate media (audio in public place)
  - current solutions:
    - voice mail → tedious, doesn't scale, hard to search and catalogue, no indication of when call might be returned
    - automated call back → rarely used, too inflexible
  - → most successful calls are now scheduled by email
- Presence-based
  - facilitates *unscheduled* communications
  - provide recipient-specific information
  - only contact in real-time if destination is willing and able
  - appropriately use synchronous vs. asynchronous communication
  - guide media use (text vs. audio)
  - predict availability in the near future (timed presence)

Prediction: almost all (professional) communication will be presence-initiated or pre-scheduled

GEOPRIV and SIMPLE architectures

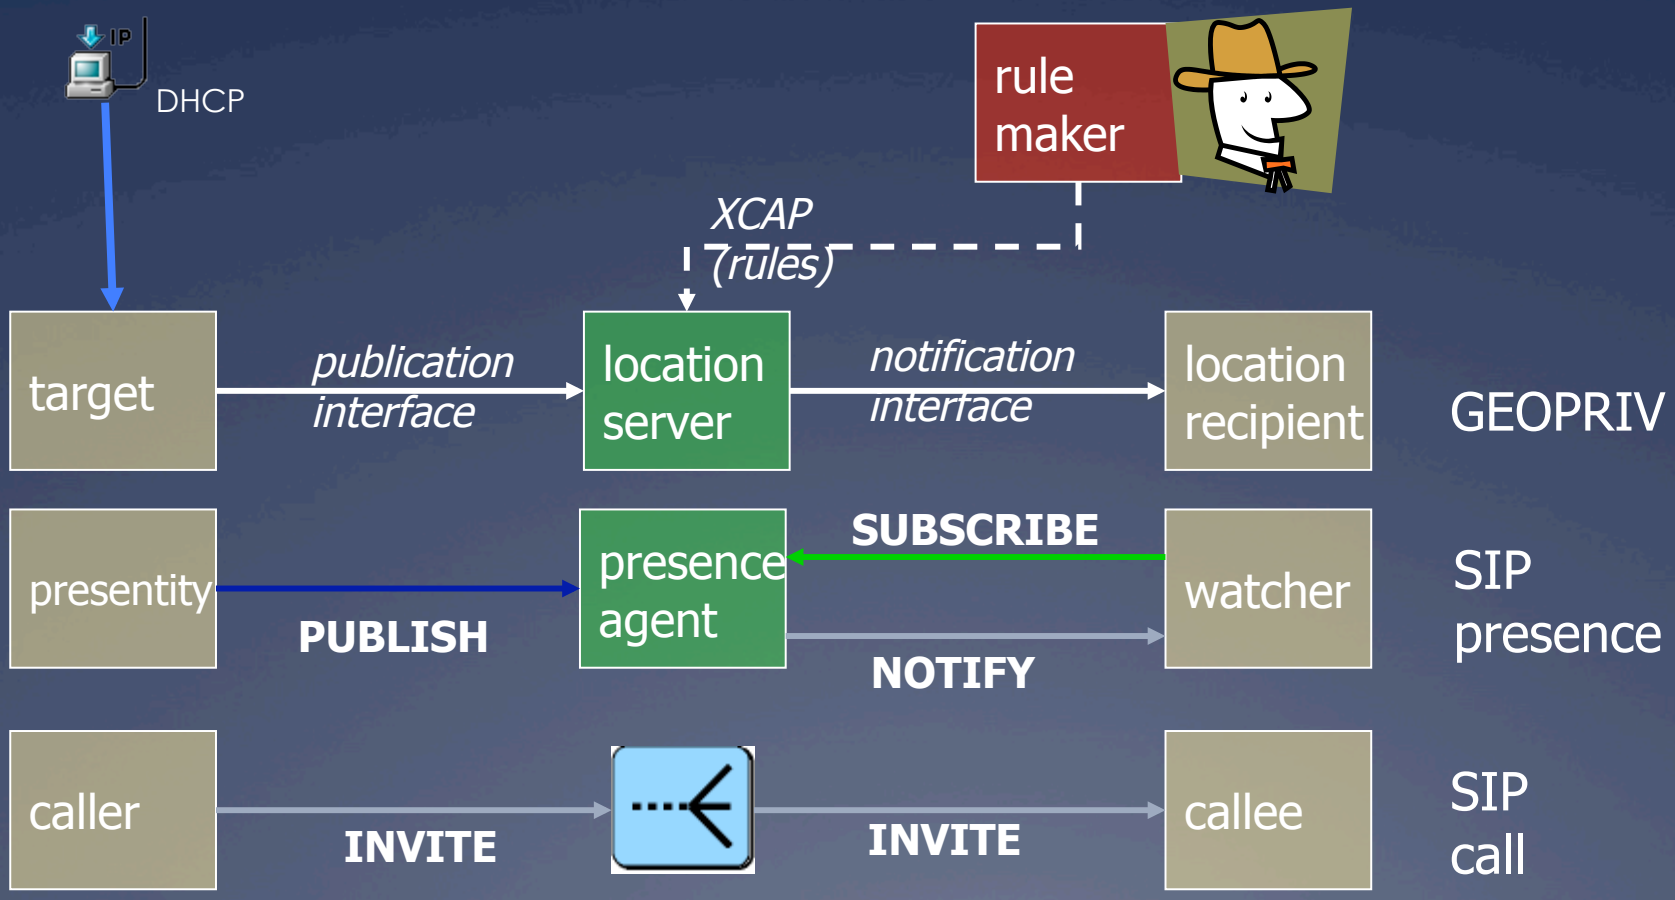

Presentity and Watchers

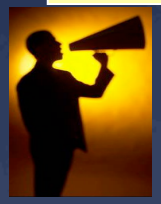

Bob's Presentity

PUBLISH

Bob's status, location

Presence Server (PS)

SUBSCRIBE

NOTIFY

Watchers

Available, Busy, Somewhat available, Invisible

Bob's Filters (Rules), PIDF *)

PUBLISH

R u there ?

BUZZ

Cell Phone PC-IM Client Bob's play station

Bob's Presence User Agents (PUA)

*) - PIDF = Presence Information Data Format

wife

son

friend

externo world

Basic presence

- Role of presence
  - initially: "can I send an instant message and expect a response?"
  - now: "should I use voice or IM? is my call going to interrupt a meeting? is the callee awake?"
- Yahoo, MSN, Skype presence services:
  - on-line & off-line
  - useful in modem days – but many people are (technically) on-line 24x7
  - thus, need to provide more context
  - + simple status ("not at my desk")
- entered manually → rarely correct
- does not provide enough context for directing interactive communications

Presence data architecture

presence sources

PUBLISH

create view (compose)

raw presence document

privacy filtering

XCAP

select best source
resolve contradictions

depends on watcher

XCAP

composition policy

privacy policy

(not defined yet)

draft-ietf-simple-presence-data-model

Presence data architecture

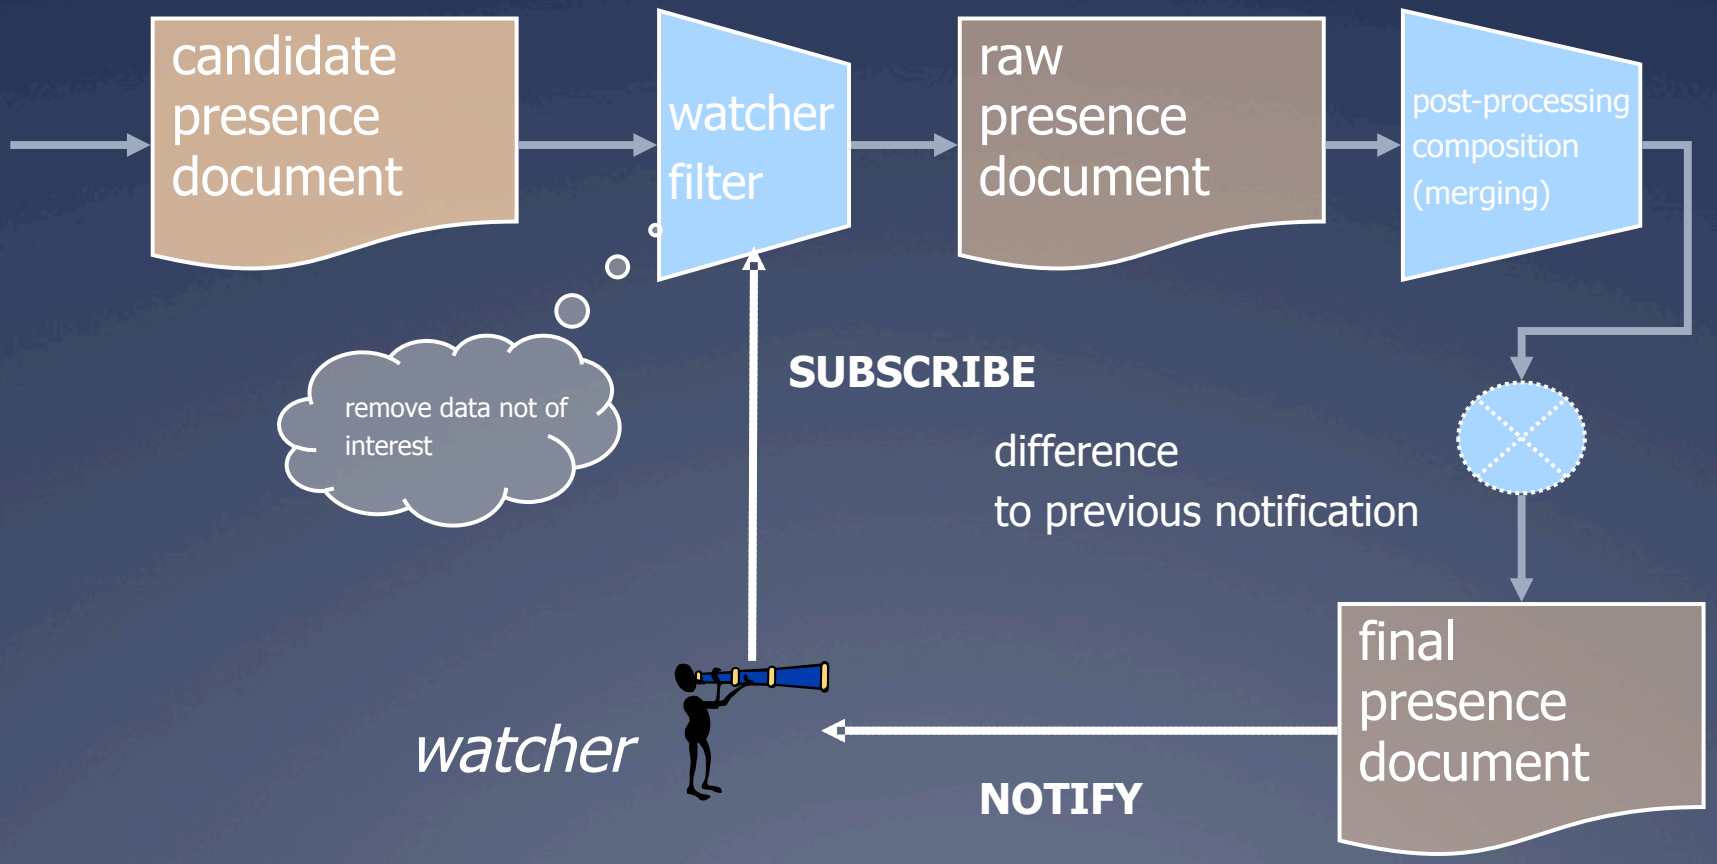

Rich presence

- Provide watchers with better information about the what, where, how of presentities
- facilitate appropriate communications:
  - "wait until end of meeting"
  - "use text messaging instead of phone call"
  - "make quick call before flight takes off"
- designed to be derivable from calendar information
  - or provided by sensors in the environment
- allow filtering by "sphere" – the parts of our life
  - don't show recreation details to colleagues

Rich presence

- automatically derived from
  - sensors: physical presence, movement
  - electronic activity: calendars
- Contains:
  - multiple contacts per presentity
    - device (cell, PDA, phone, …)
    - service (“audio”)
  - activities, current and planned
  - surroundings (noise, privacy, vehicle, …)
  - contact information
  - composing (typing, recording audio/video IM, …)

The role of presence for call routing

- Two modes:
  - watcher uses presence information to select suitable contacts
    - advisory – caller may not adhere to suggestions and still call when you're in a meeting
  - user call routing policy informed by presence
    - likely less flexible – machine intelligence
    - "if activities indicate meeting, route to tuple indicating assistant"
    - "try most-recently-active contact first" (seq. forking)

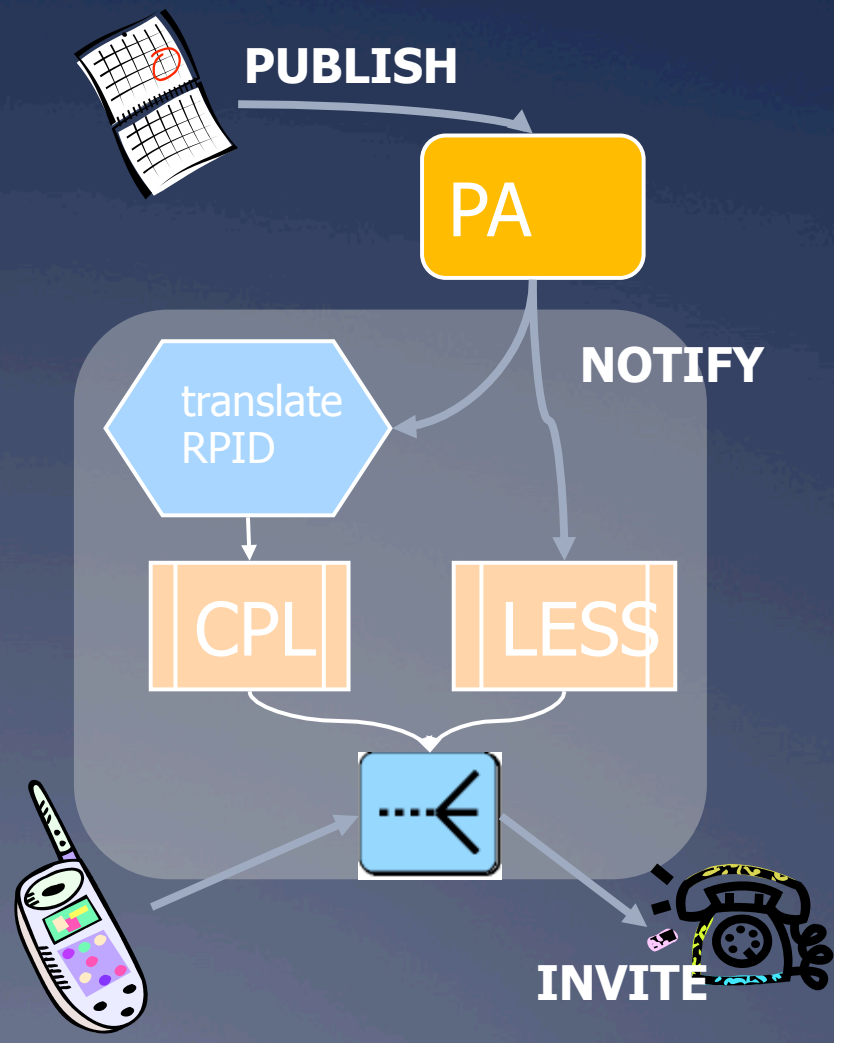

Presence and privacy

- All presence data, particularly location, is highly sensitive
- Basic location object (PIDF-LO) describes
  - distribution (binary)
  - retention duration
- Policy rules for more detailed access control
  - who can subscribe to my presence
  - who can see what when

```
<tuple id="sg89ae">
  <status>
    <gp:geopriv>
      <gp:location-info>
        <gml:location>
          <gml:Point gml:id="point1"
          srsName="epsg:4326">
              <gml:coordinates>37:46:30N 122:25:10W
            </gml:coordinates>
             </gml:Point>
        </gml:location>
      </gp:location-info>
     <gp:usage-rules>
       <gp:retransmission-allowed>no
       </gp:retransmission-allowed>
       <gp:retention-expiry>2003-06-23T04:57:29Z
       </gp:retention-expiry>
     </gp:usage-rules>
    </gp:geopriv>
  </status>
  <timestamp>2003-06-22T20:57:29Z</timestamp>
</tuple>
```

Privacy rules

- Conditions
  - identity, sphere
  - time of day
  - current location
  - identity as <uri> or <domain> + <except>
- Actions
  - watcher confirmation
- Transformations
  - include information
  - reduced accuracy
- User gets maximum of permissions across all matching rules
  - privacy-safe composition: removal of a rule can only reduce privileges
- Extendable to new presence data
  - rich presence
  - biological sensors
  - mood sensors

Example rules document

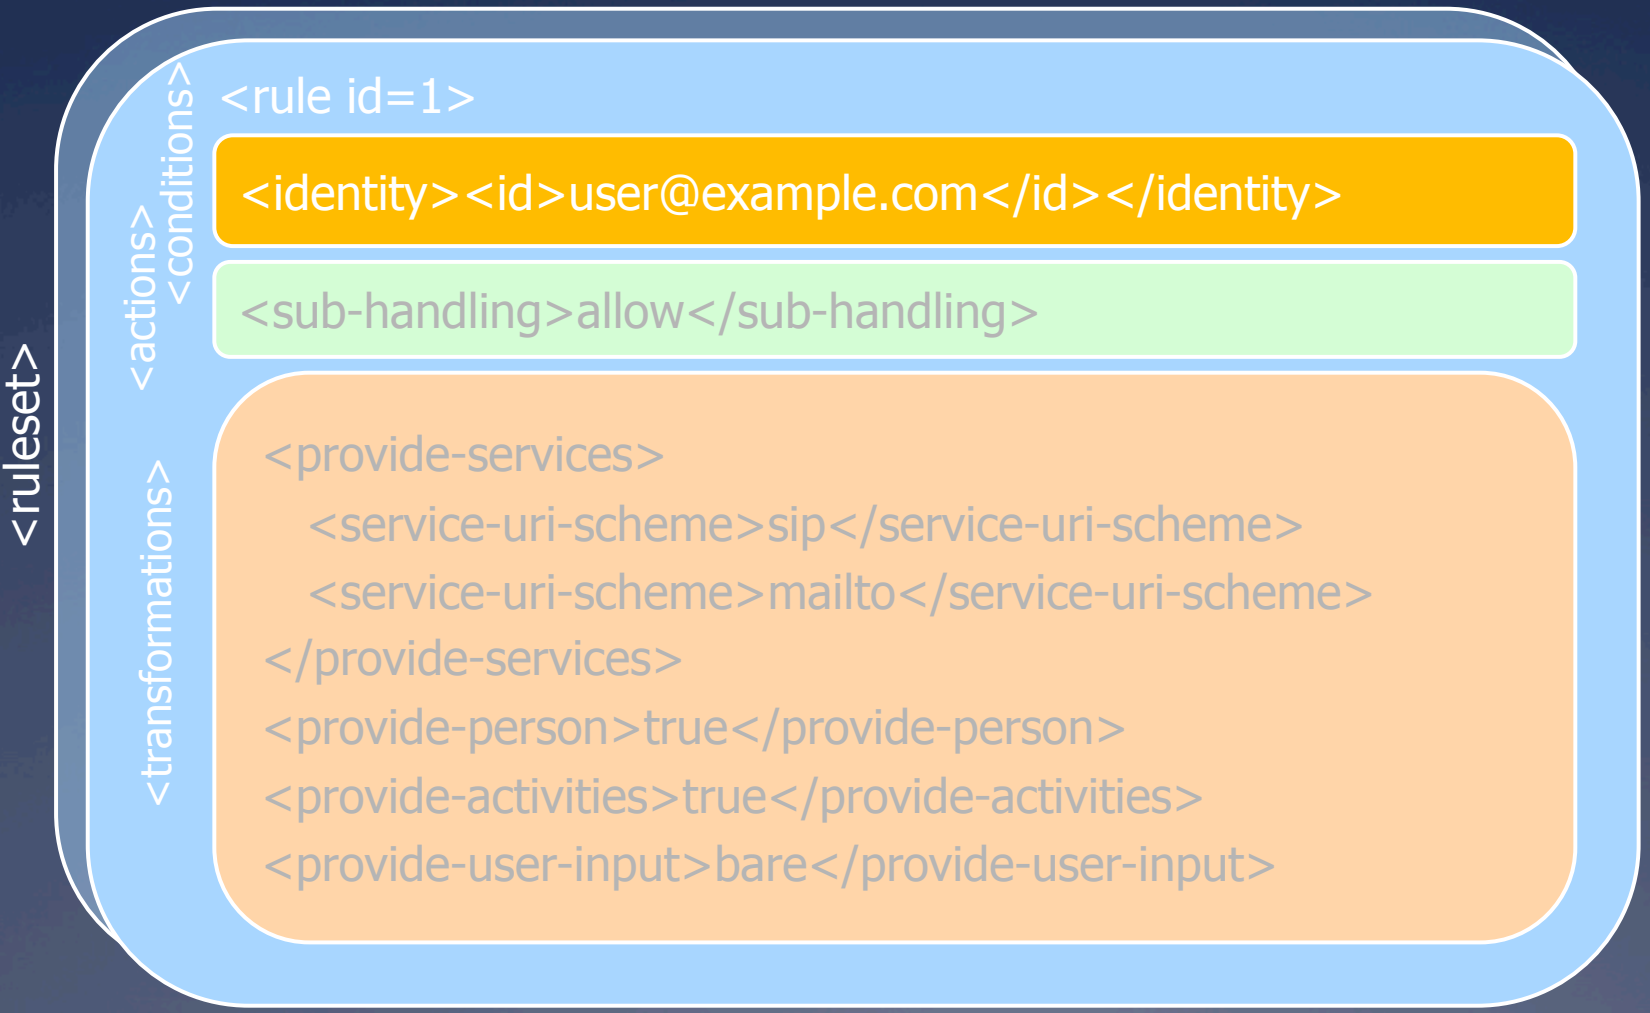

Location-based services

Location-based services

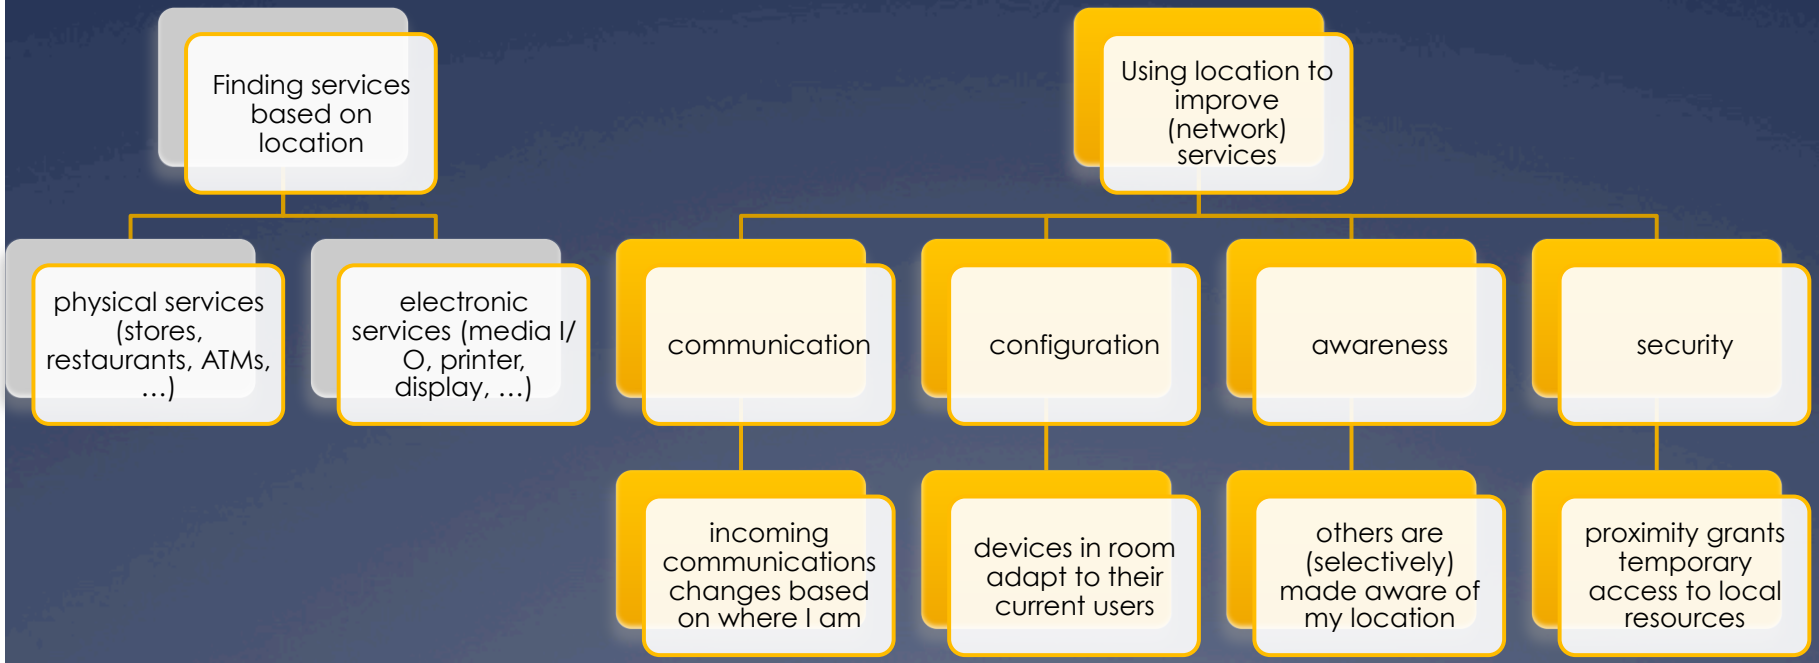

Location-based SIP services

- Location-aware inbound routing
  - do not forward call if time at callee location is [11 pm, 8 am]
  - only forward time-for-lunch if destination is on campus
  - do not ring phone if I'm in a theater
- outbound call routing
  - contact nearest emergency call center
  - send delivery@pizza.com to nearest branch
- location-based events
  - subscribe to locations, not people
  - Alice has entered the meeting room
  - subscriber may be device in room → our lab stereo changes CDs for each person that enters the room

Location delivery

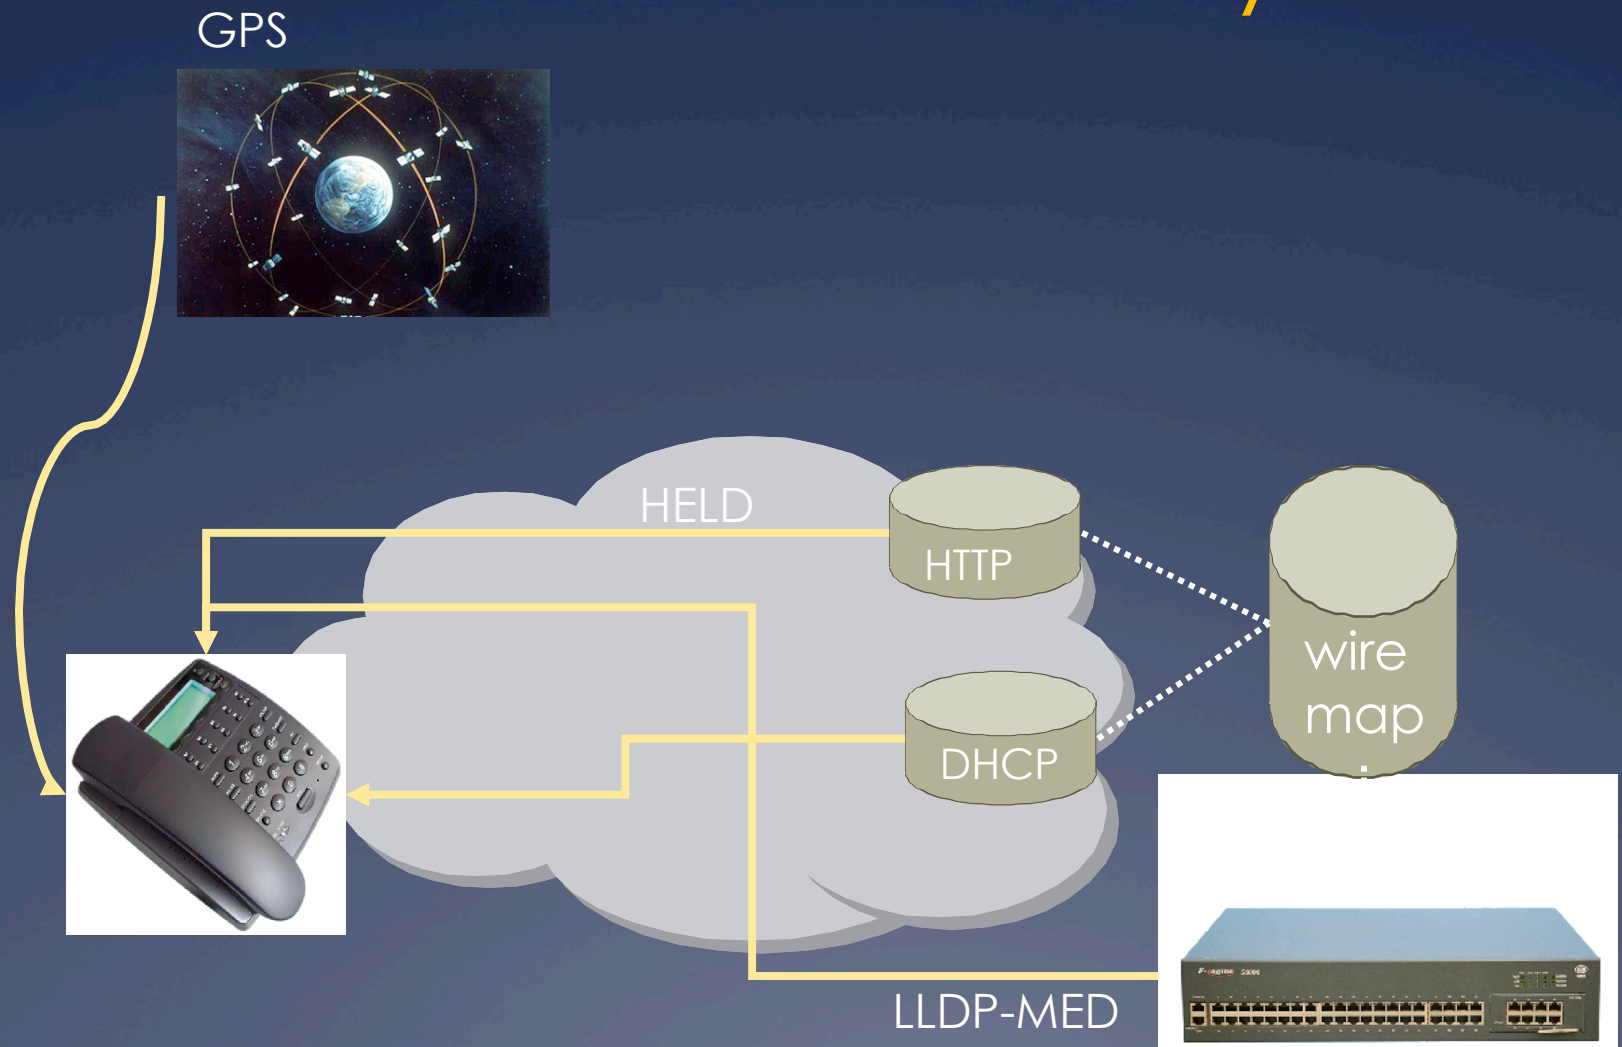

Location determination options

Method	CDP or LLDP-MED	DHCP	HELD	GPS	manual entry
Layer	L2	L3	L7 (HTTP)	-	user
advantages	• simple to implement • built into switch • direct port/room mapping	• simple to implement • network locality	• traverses NATs • can be operated by L2 provider	• accurate • mobile devices • no carrier cooperation	• no infrastructure changes • no carrier cooperation
problems	may be hard to automate for large enterprises	mapping MAC address to location?	mapping IP address to switch port?	• indoor coverage • acquisition time	• fails for mobile devices • unreliable for nomadic
Use	Ethernet LANs	Enterprise LANs Some ISPs	DSL, cable	mobile devices	fall back

Program location-based services

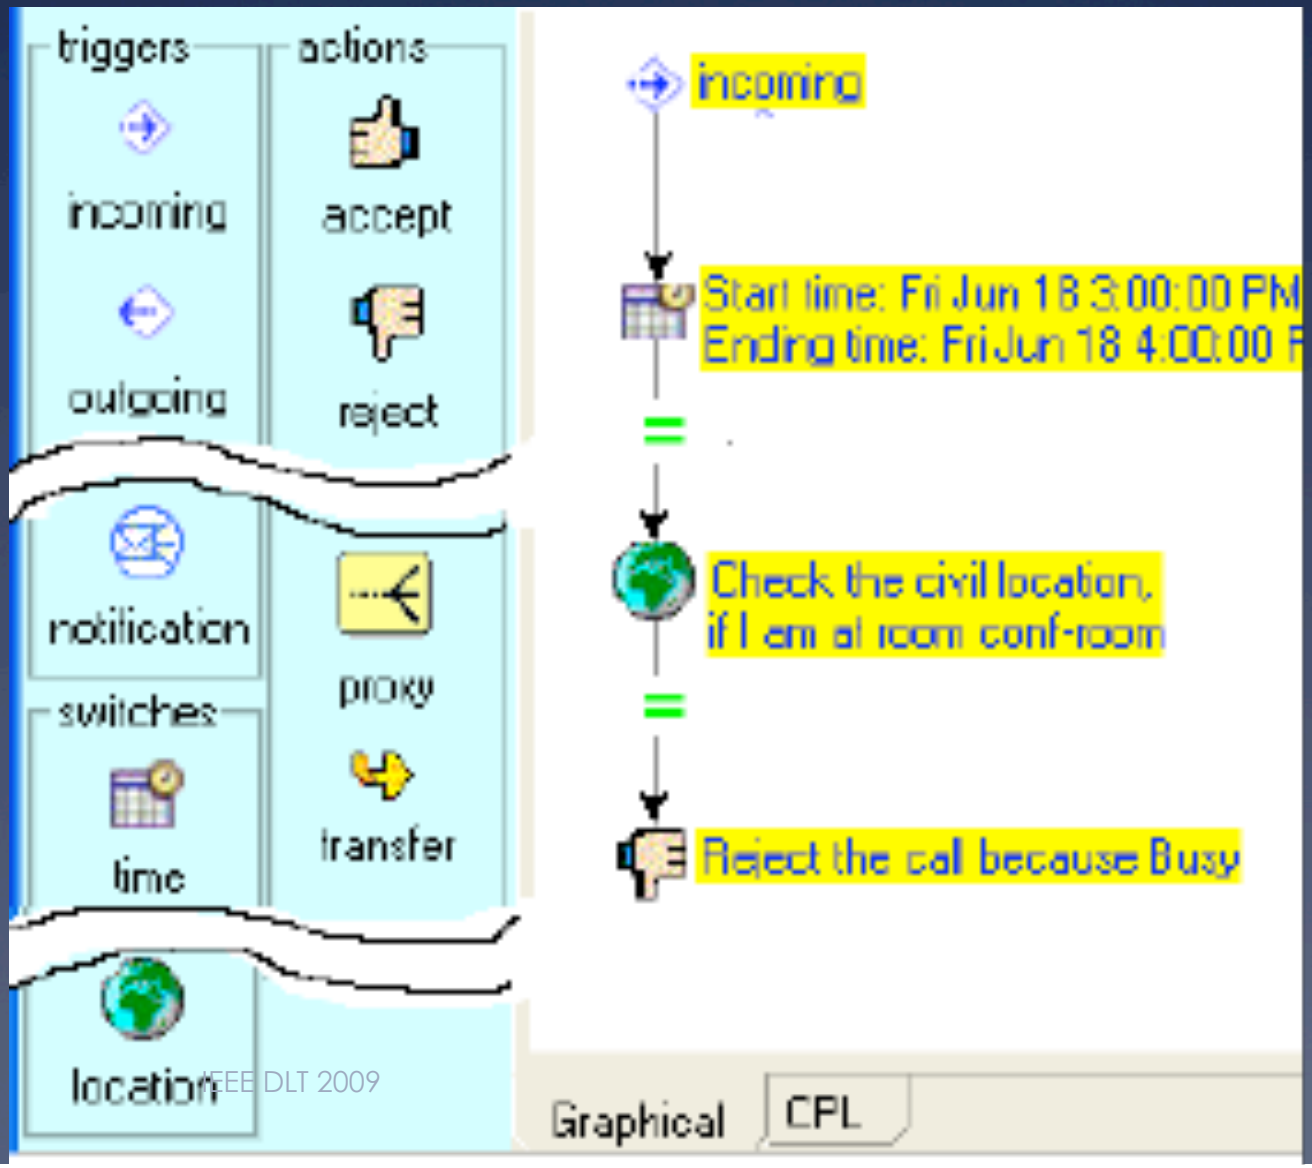

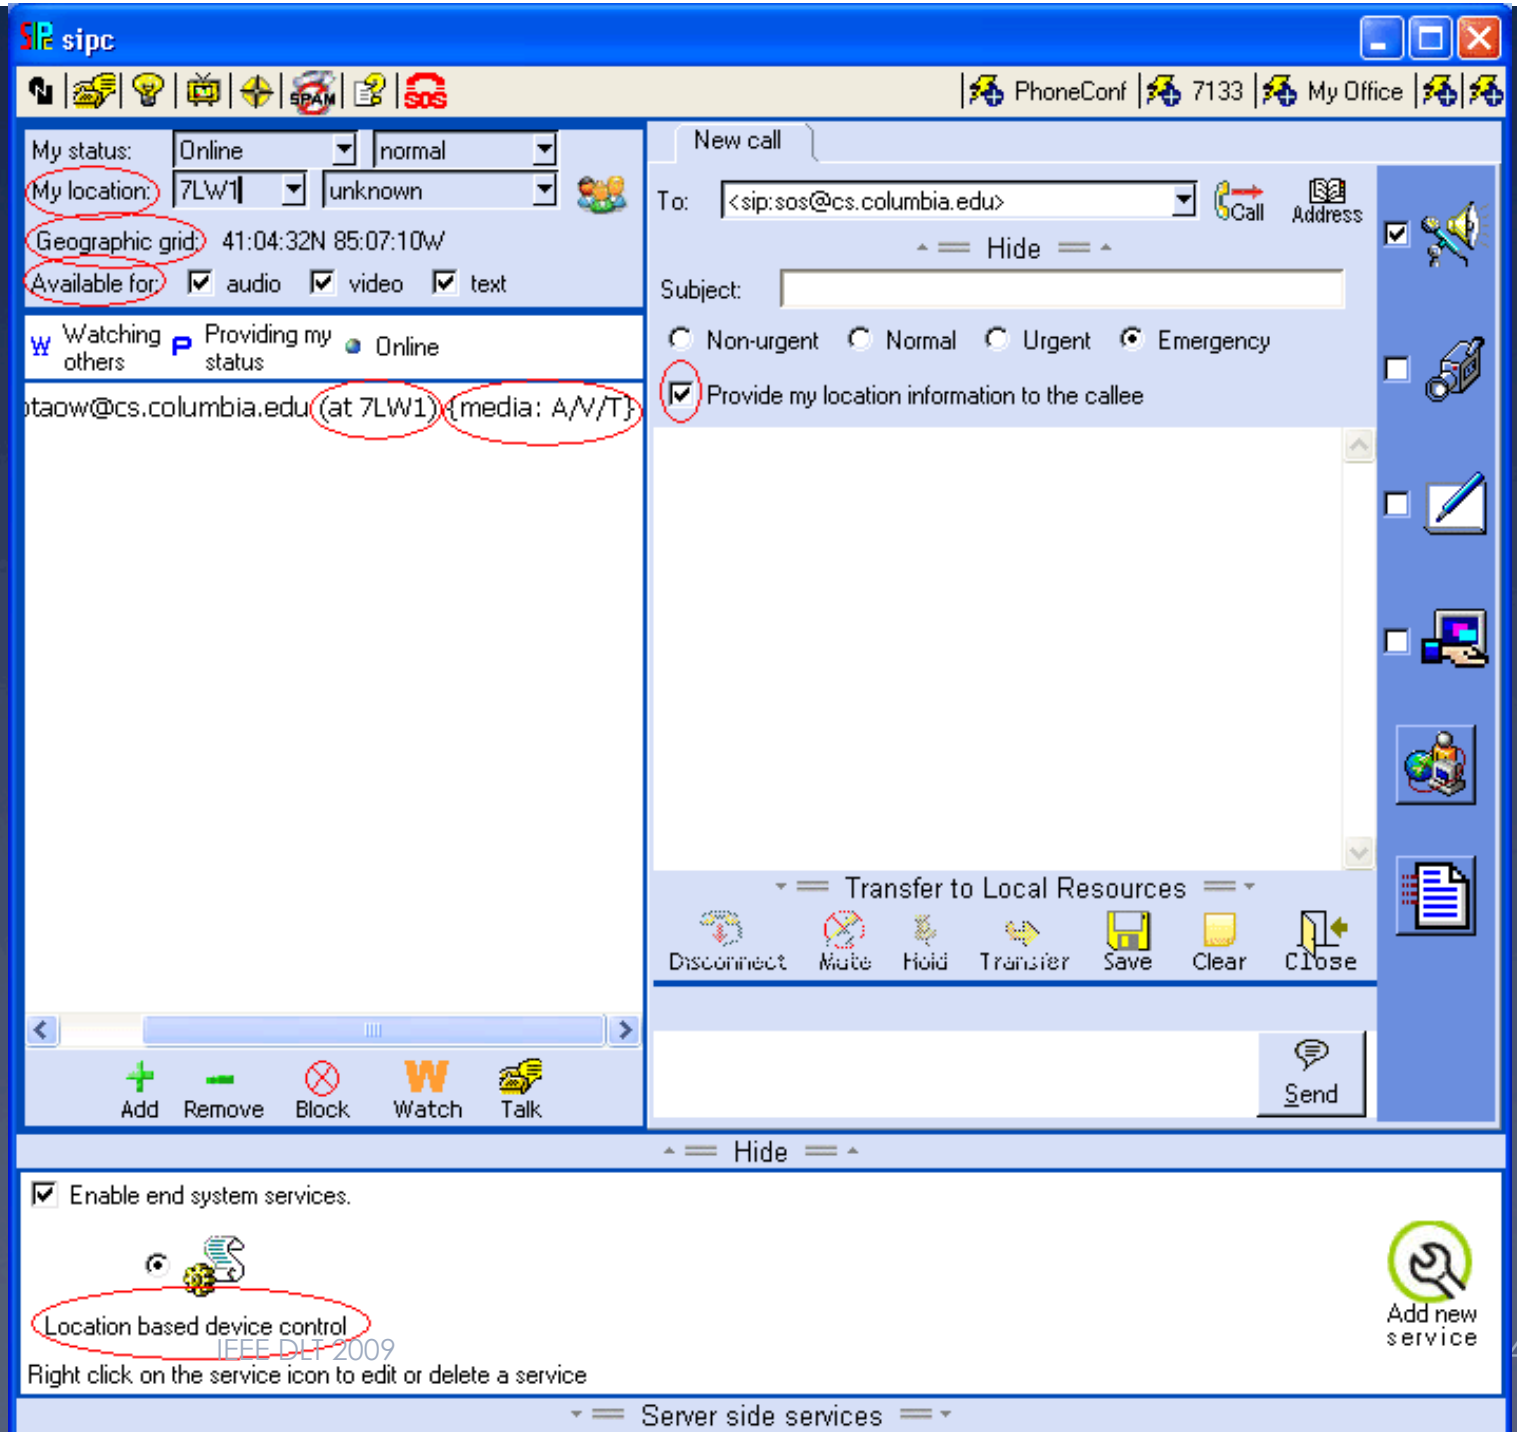
sipc
PhoneConf
7133
My Office
My status:
Online
normal
My location:
7LW1
unknown
Geographic grid:
41:04:32N 85:07:10W
Available for:
audio
video
text
Watching others
Providing my status
Online
taow@cs.columbia.edu
(at 7LW1)
{media: A/V/T}
Add
Remove
Block
Watch
Talk
New call
To:
<sip:sos@cs.columbia.edu>
Call
Address
Hide
Subject:
Non-urgent
Normal
Urgent
Emergency
Provide my location information to the callee
Transfer to Local Resources
Disconnect
Mute
Hold
Transfer
Save
Clear
Close
Send
Hide
Enable end system services.
Location based device control
Right click on the service icon to edit or delete a service
Add new service
Server side services


Emergency calling

Modes of emergency communications

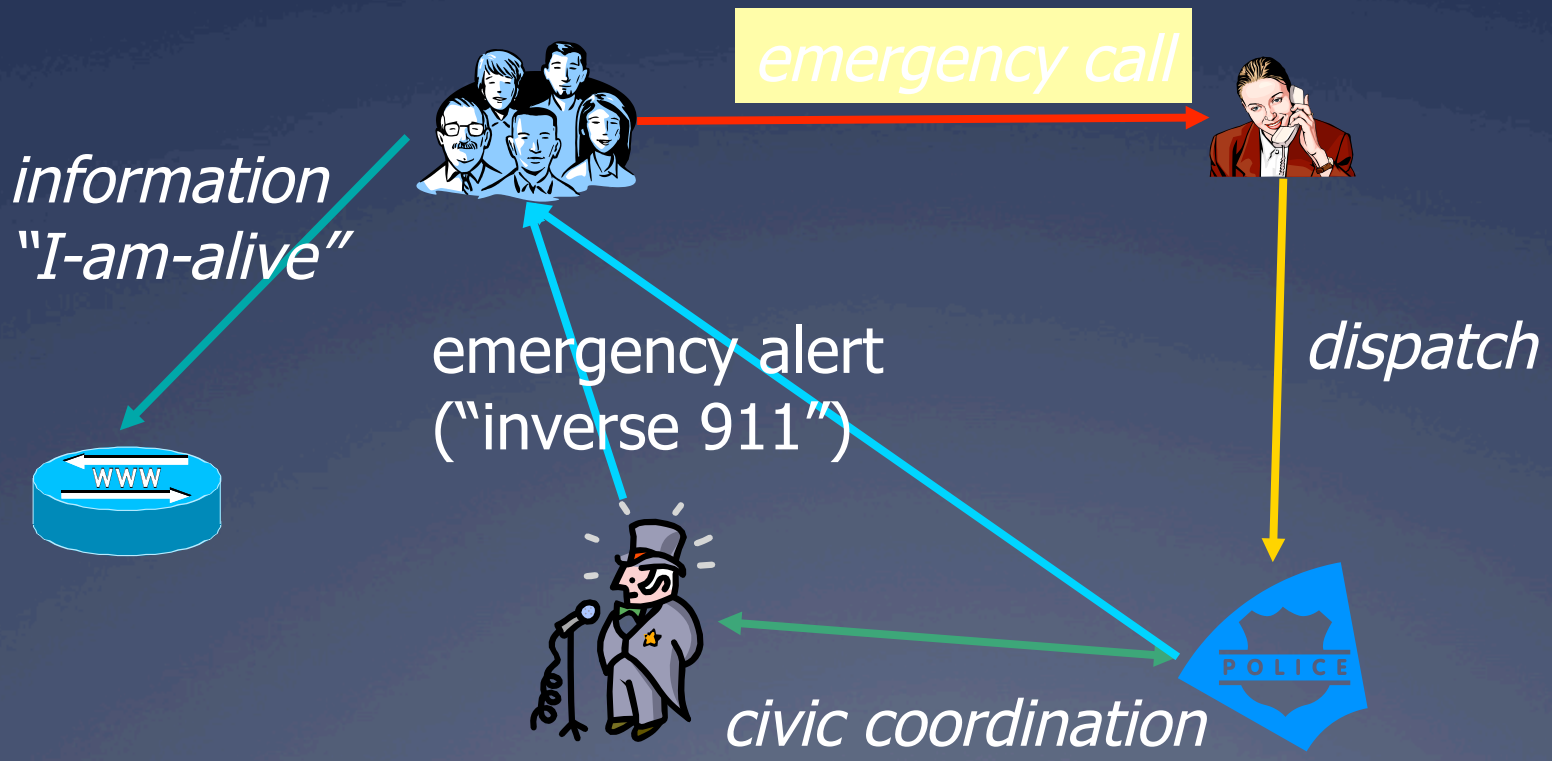

Background on 9-1-1

- Established in Feb. 1968
  - 1970s: selective call routing
  - late 1990s: 93% of population/96% of area covered by 9-1-1
  - 95% of 9-1-1 is Enhanced 9-1-1
  - US and Canada
- Roughly 200 mio. calls a year (6 calls/second)
  - 1/3 wireless
- 6146 PSAPs in 3135 counties
  - most are small (2-6 call takers)
  - 83.1% of population have some Phase II (April 2007)
- "12-15 million households will be using VoIP as either primary or secondary line by end of 2008" (NENA)

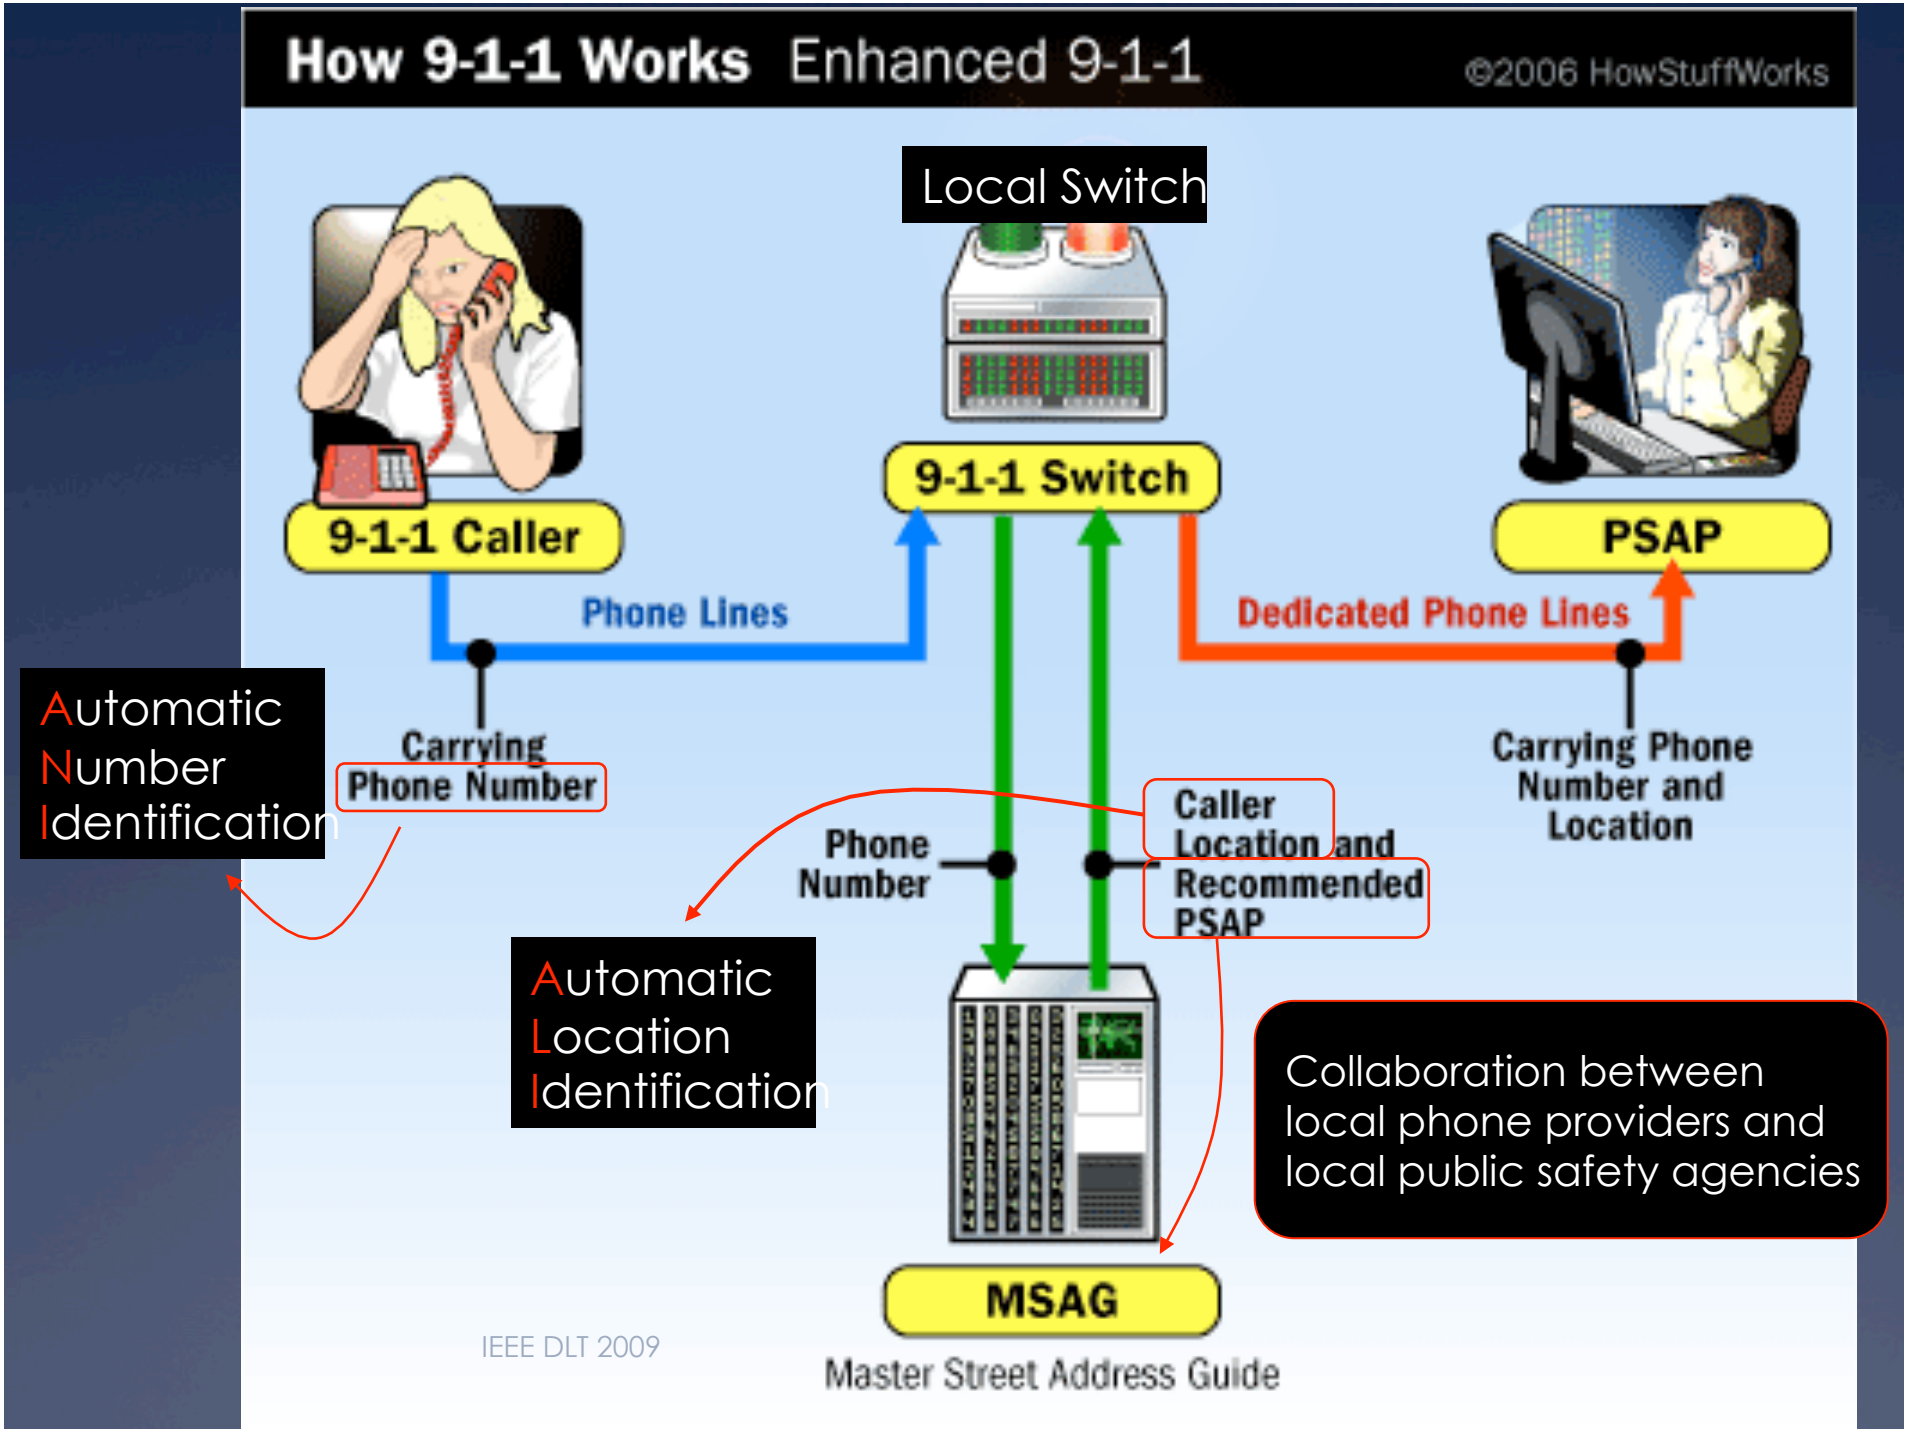
How 9-1-1 Works Enhanced 9-1-1
©2006 HowStuffWorks
Local Switch
9-1-1 Switch
9-1-1 Caller
PSAP
Phone Lines
Dedicated Phone Lines
Carrying Phone Number
Carrying Phone Number and Location
Automatic Number Identification
Phone Number
Caller Location and Recommended PSAP
Automatic Location Identification
Collaboration between local phone providers and local public safety agencies
MSAG
Master Street Address Guide

What makes VoIP 112/911 hard?

POTS	PSTN-emulation VoIP	end-to-end VoIP
(landline) phone number limited to limited area	landline phone number anywhere in US (cf. German 180)	no phone number or phone number anywhere around the world
regional carrier	national or continent-wide carrier	enterprise "carrier" or anybody with a peer-to-peer device
voice provider = line provider (~ business relationship)	voice provider ≠ ISP	voice provider ≠ ISP
national protocols and call routing	probably North America + EU	international protocols and routing
location = line location	mostly residential or small business	stationary, nomadic, wireless

Emergency numbers

- Each country and region has their own
  - subject to change
- Want to enable
  - traveler to use familiar home number
  - good samaritan to pick up cell phone
- Some 3/4-digit numbers are used for non-emergency purposes (e.g., directory assistance)

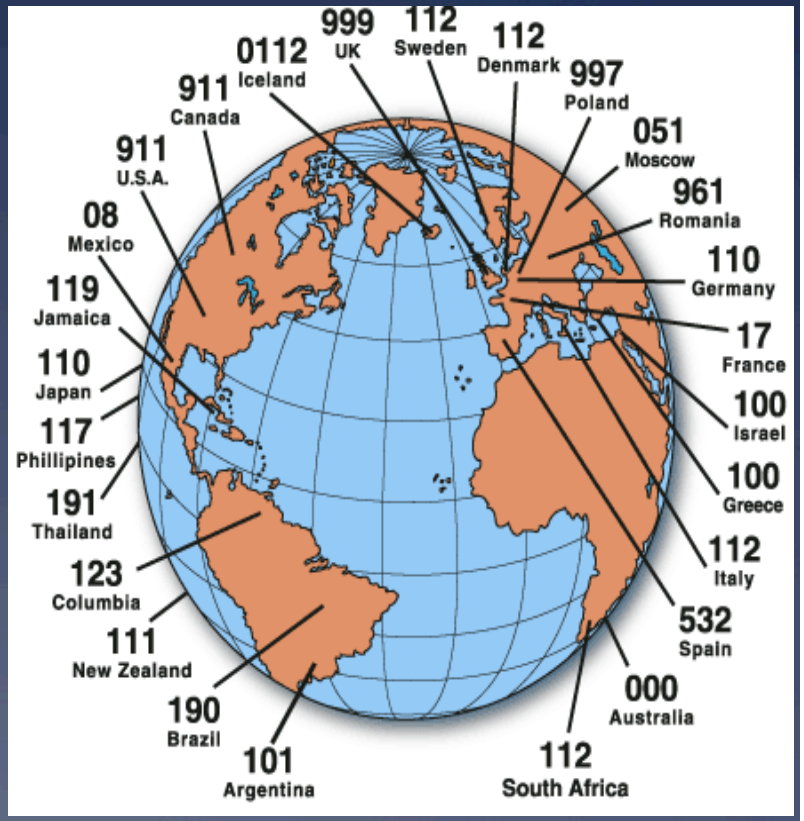

Emergency number

Service URN

- Idea: Identifiers to denote emergency calls
  - and other generic (communication) services
- Described in IETF ECRIT RFC 5031
- Emergency service identifiers:

sos	General emergency services
sos.animal-control	Animal control
sos.fire	Fire service
sos.gas	Gas leaks and gas emergencies
sos.marine	Maritime search and rescue
sos.mountain	Mountain rescue
sos.physician	Physician referral service
sos.poison	Poison control center
sos.police	Police, law enforcement

LoST: Location-to-URL Mapping

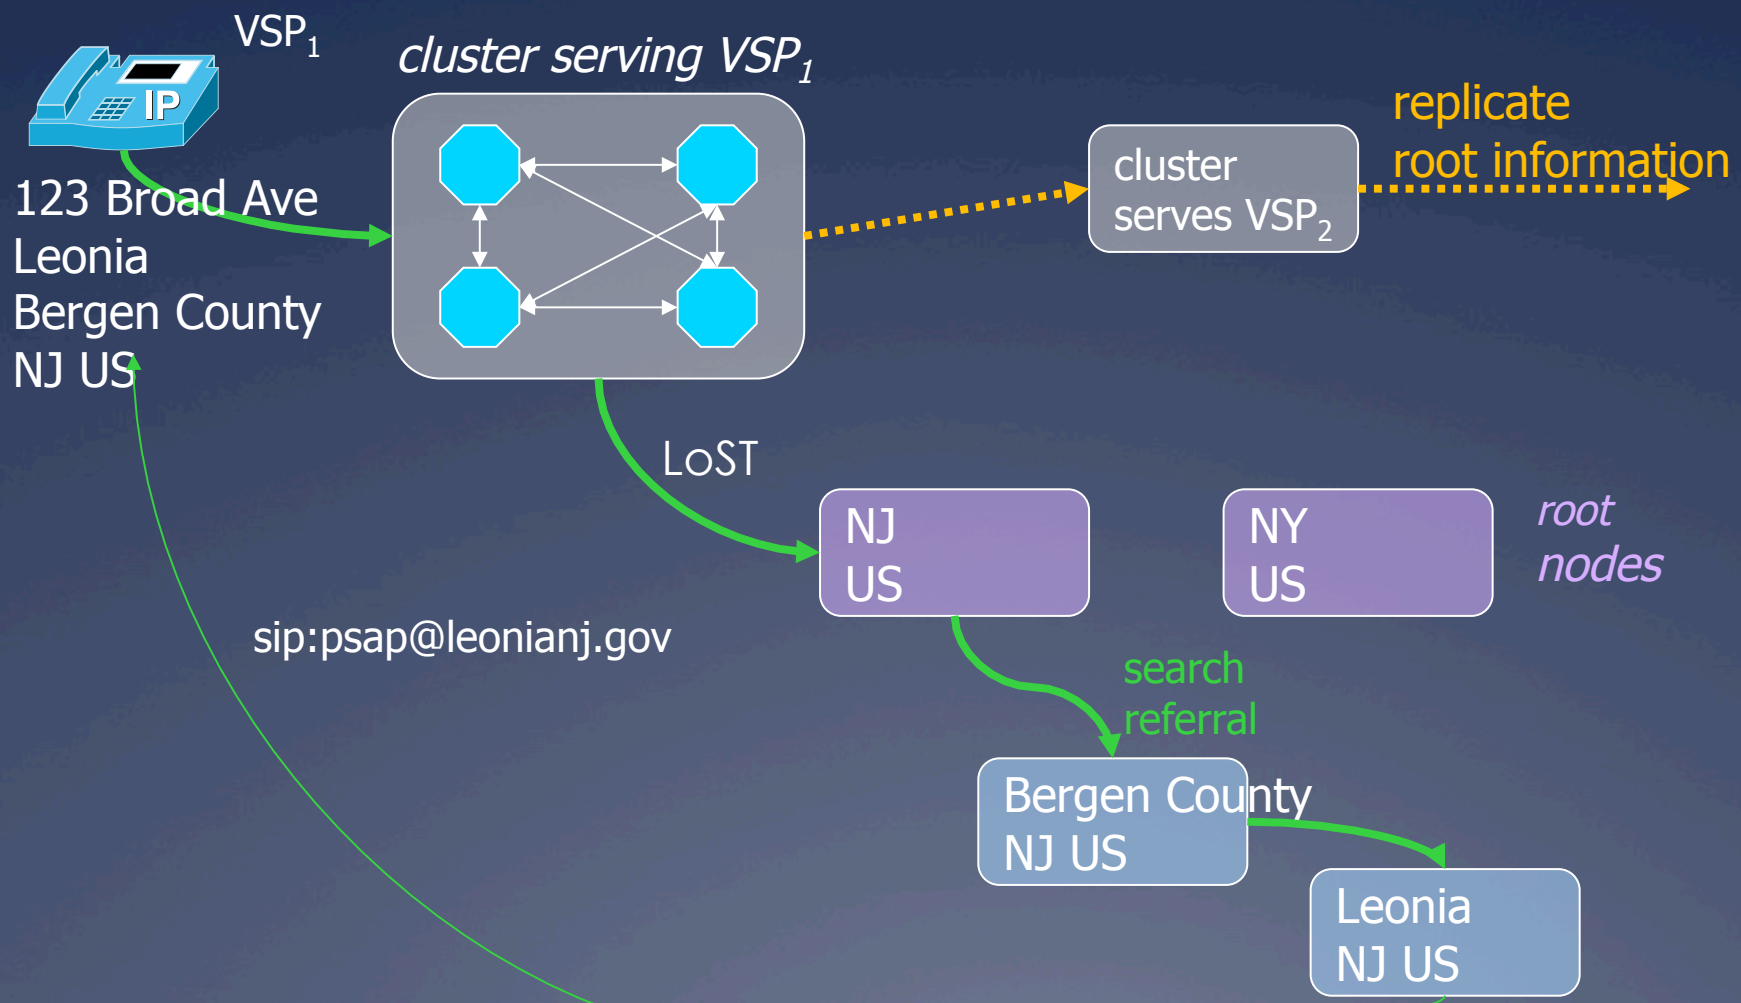

Cellular

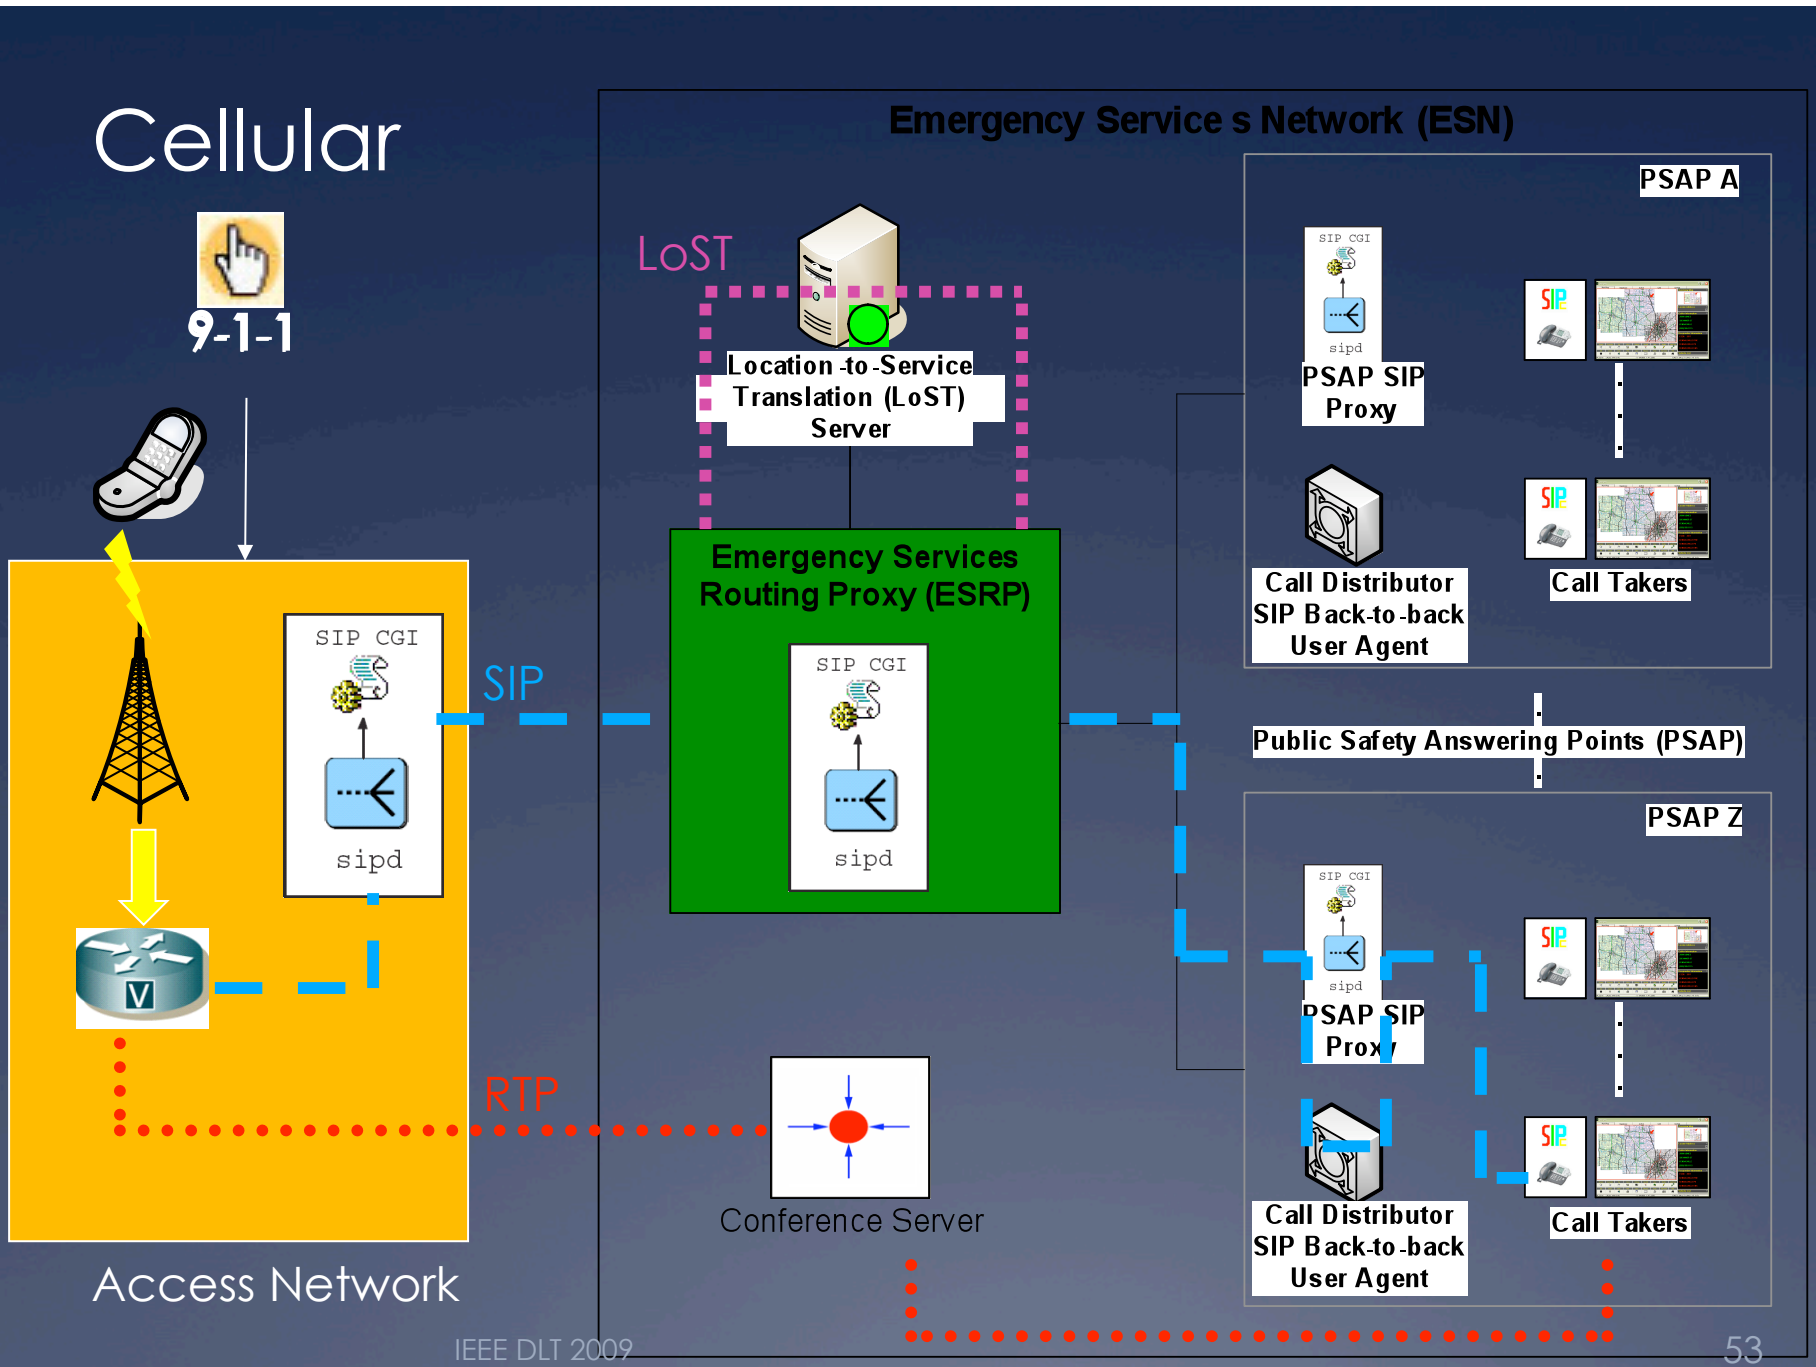
9-1-1
Emergency Service s Network (ESN)
LoST
Location -to-Service Translation (LoST) Server
Emergency Services Routing Proxy (ESRP)
SIP CGI
sipd
SIP
RTP
PSAP A
PSAP SIP Proxy
Call Distributor SIP Back-to-back User Agent
Call Takers
Public Safety Answering Points (PSAP)
PSAP Z
Conference Server
Access Network

The POC system is deployed in 5 real PSAPs and 3 labs across the USA.
PSAP: Public Safety Answering Point (=Emergency call center)

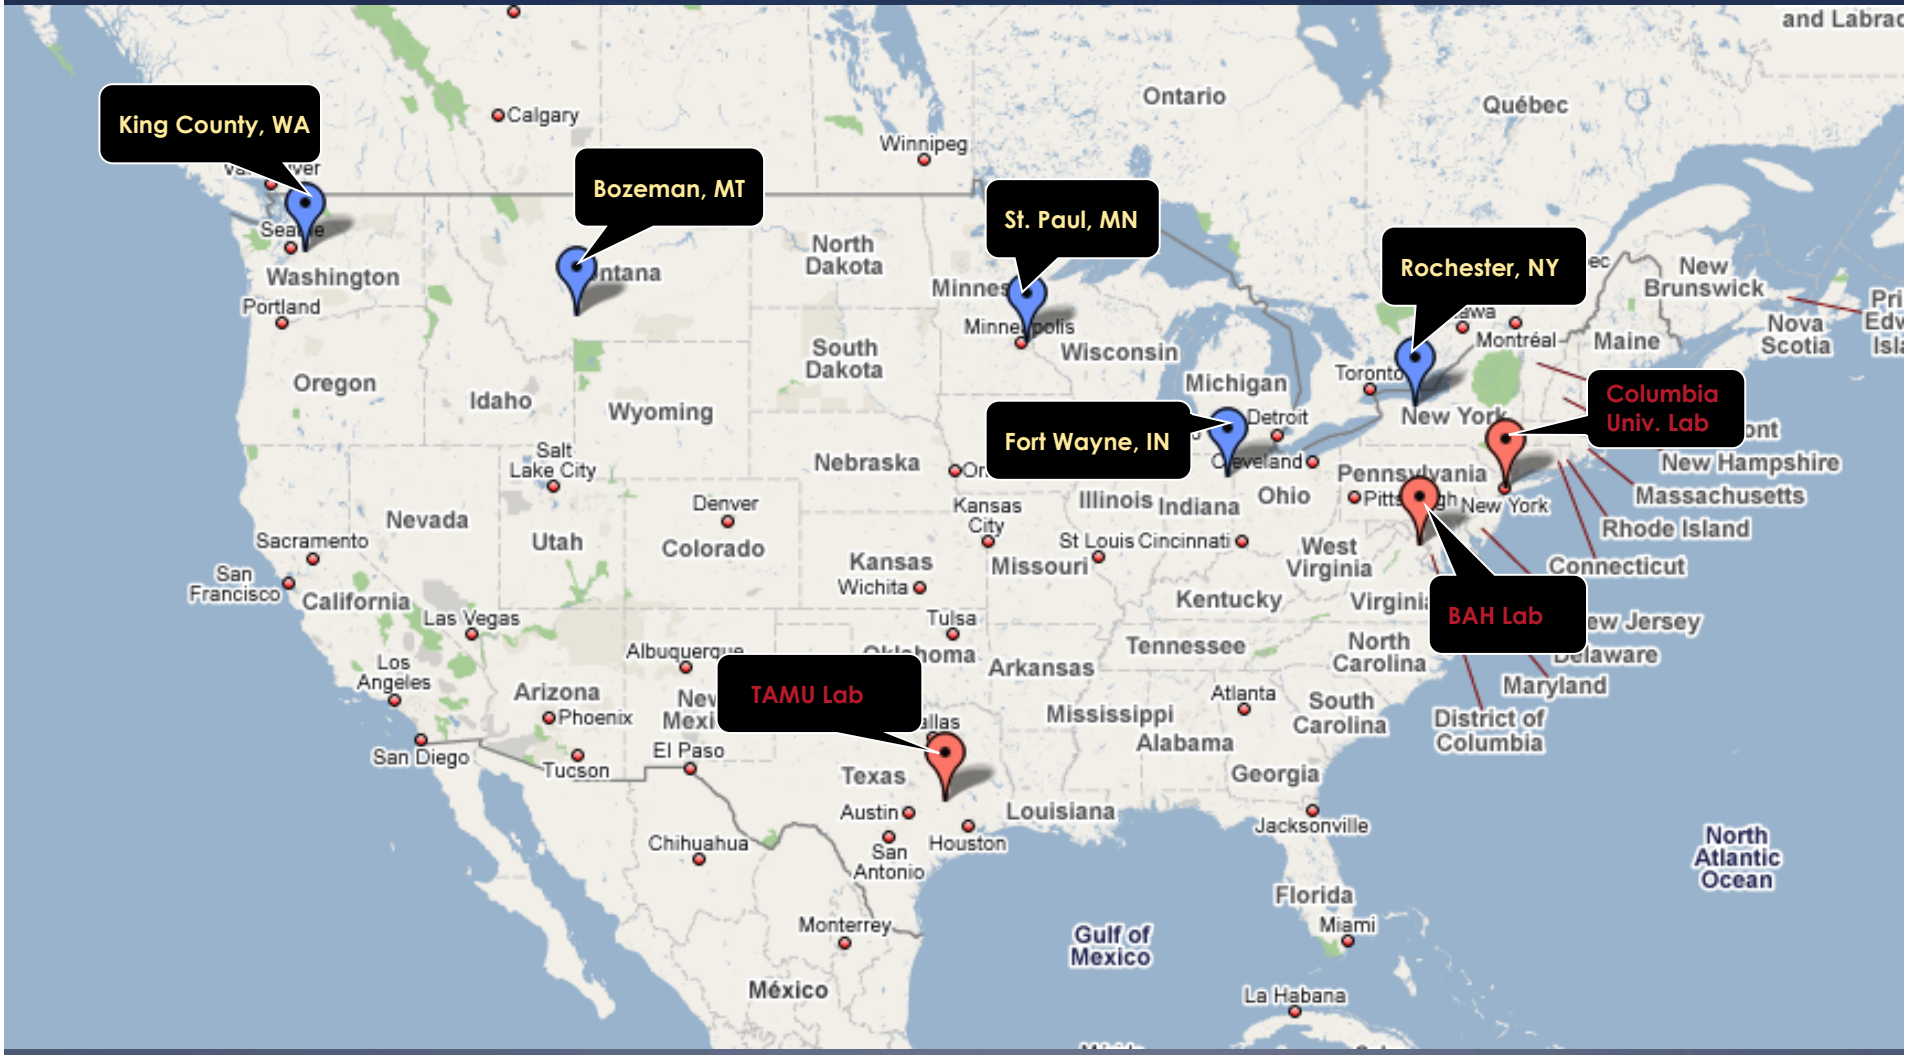

P2P

IEEE DLT 2009

Defining peer-to-peer systems

Each peer must act as both a client and a server.

Peers provide computational or storage resources for ***other*** peers.

Self-organizing and scaling.

1 & 2 are not sufficient:

- DNS resolvers provide services to others
- Web proxies are both clients and servers
- SIP B2BUAs are both clients and servers

P2P systems are …

vs.

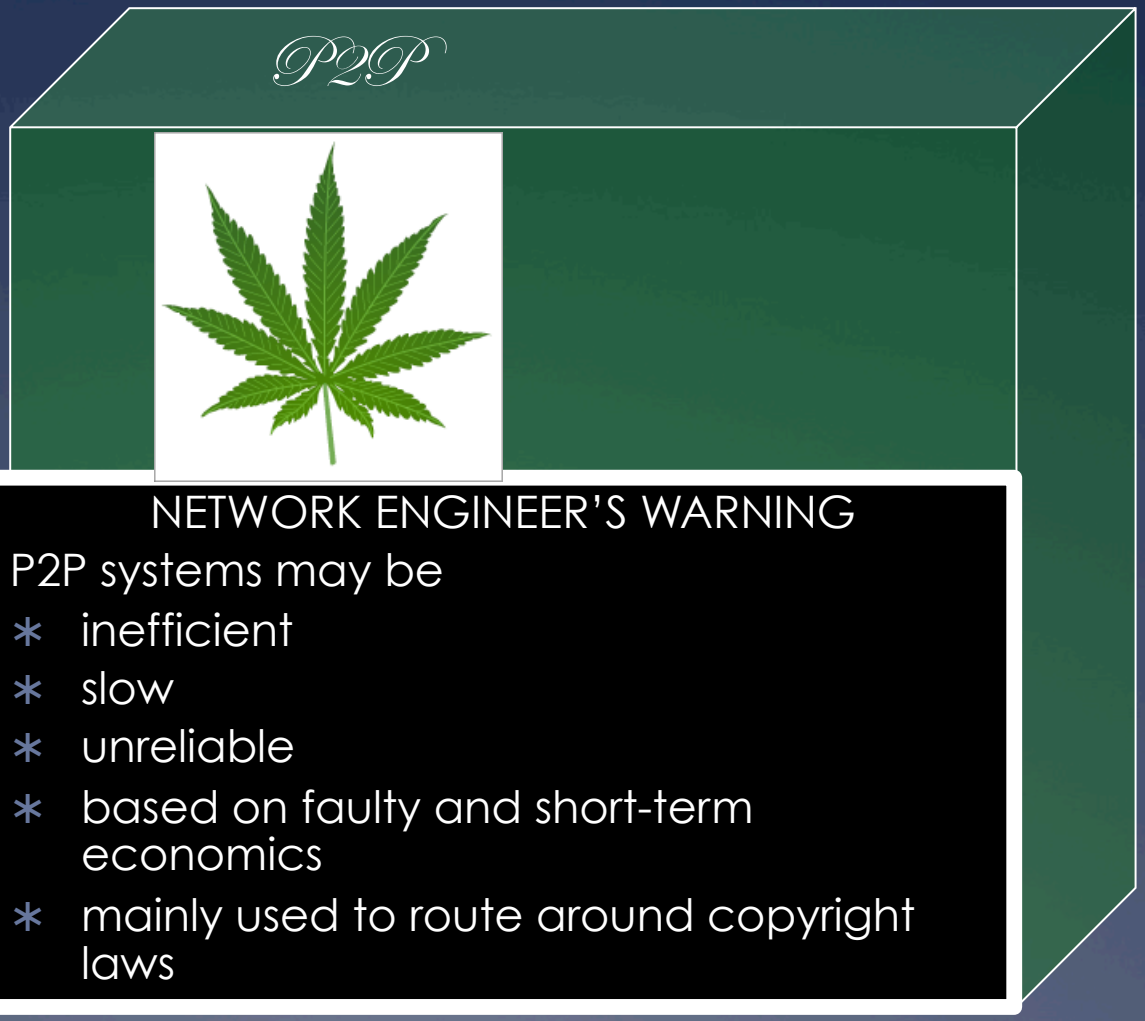

NETWORK ENGINEER'S WARNING

P2P systems may be

* inefficient
* slow
* unreliable
* based on faulty and short-term economics
* mainly used to route around copyright laws

Motivation for peer-to-peer systems

- Saves money for those offering services
  - addresses market failures
- Scales up automatically with service demand
- More reliable than client-server (no single point of failure)
- No central point of control
  - mostly plausible deniability
- Networks without infrastructure (or system manager)
- New services that can't be deployed in the ossified Internet
  - e.g., RON, ALM

P2P traffic is not devouring the Internet…

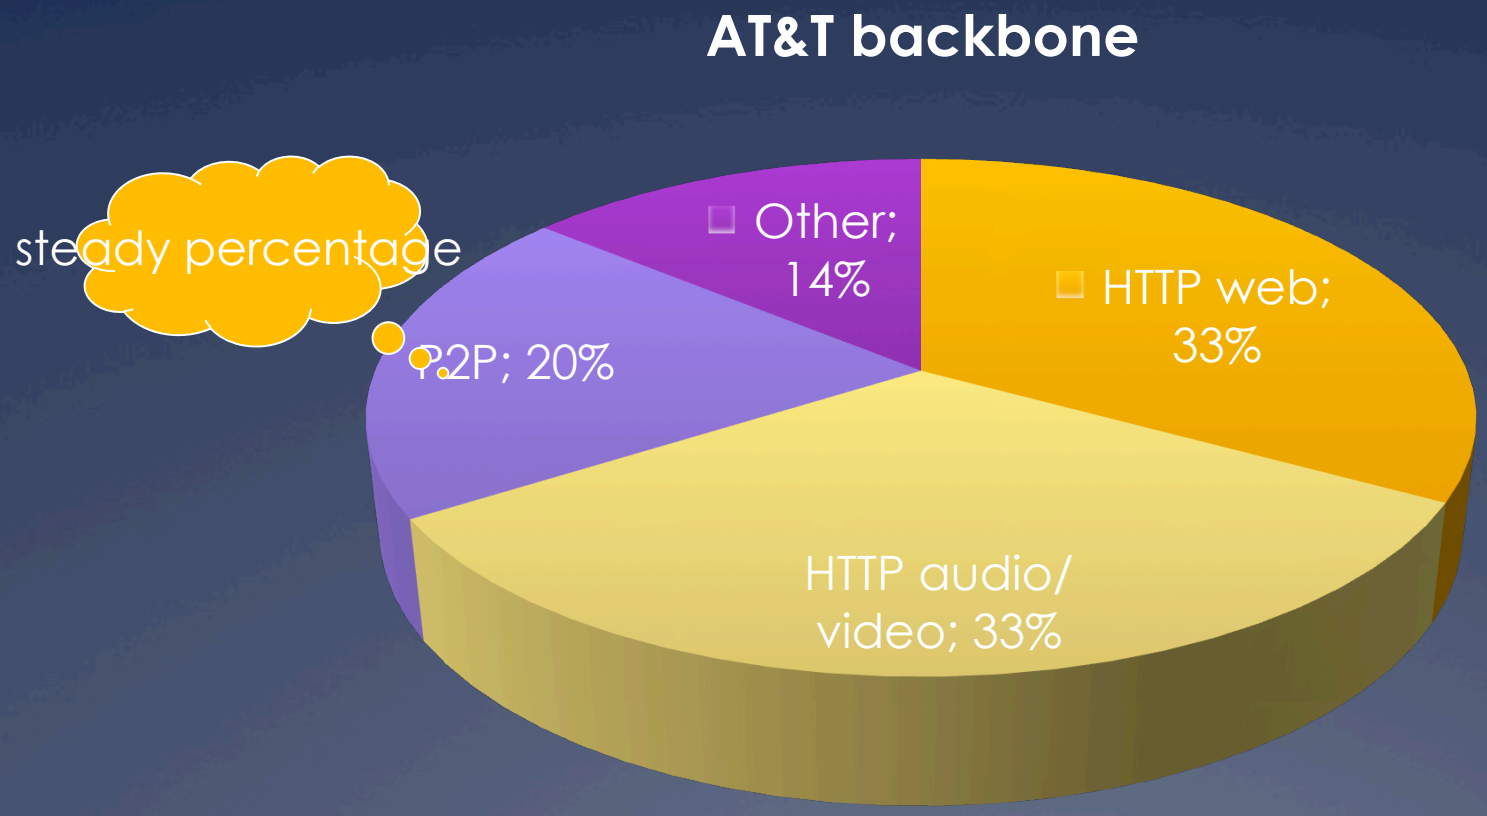

Energy consumption

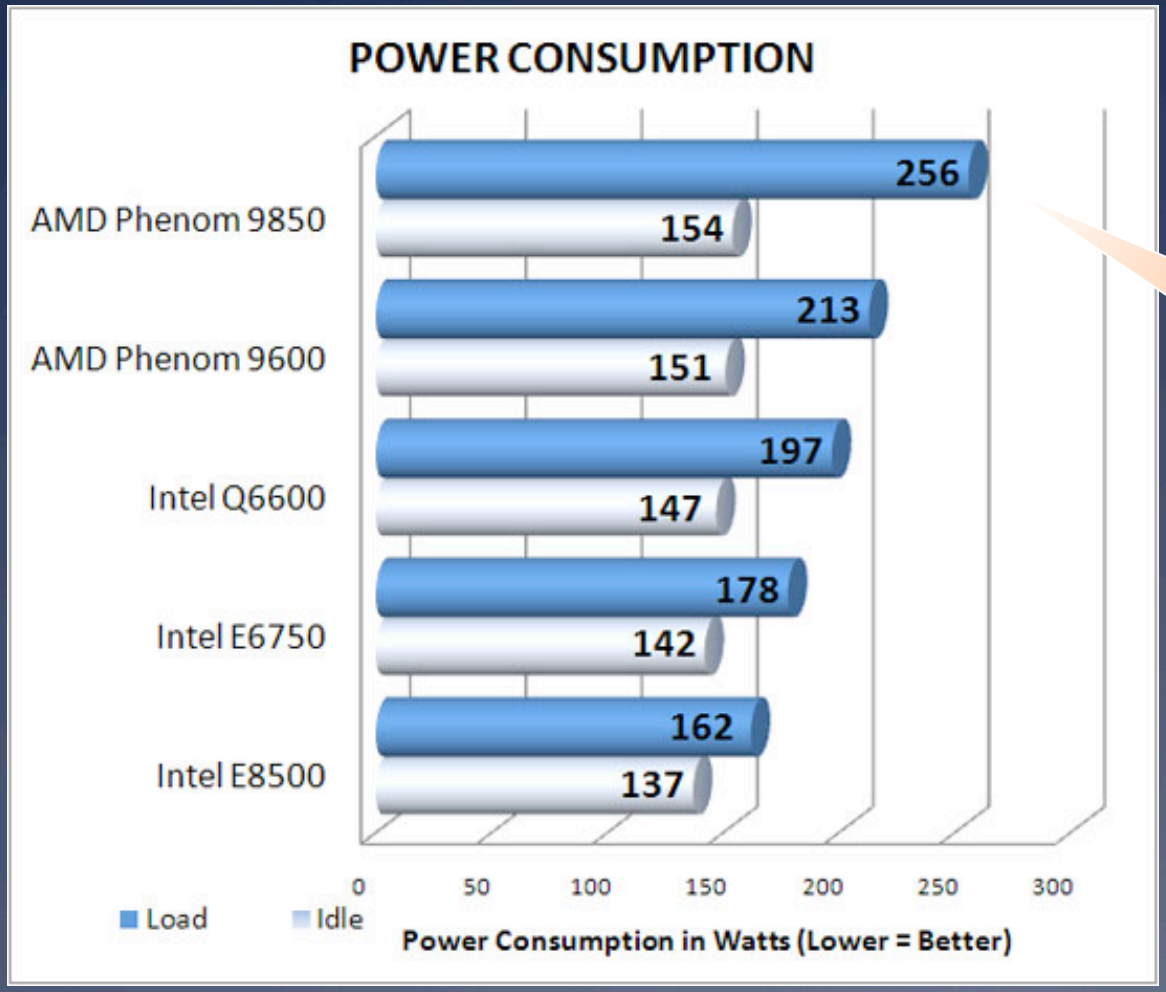

Monthly cost = $37 @ $0.20/kWh

http://www.legitreviews.com/article/682/

Bandwidth costs

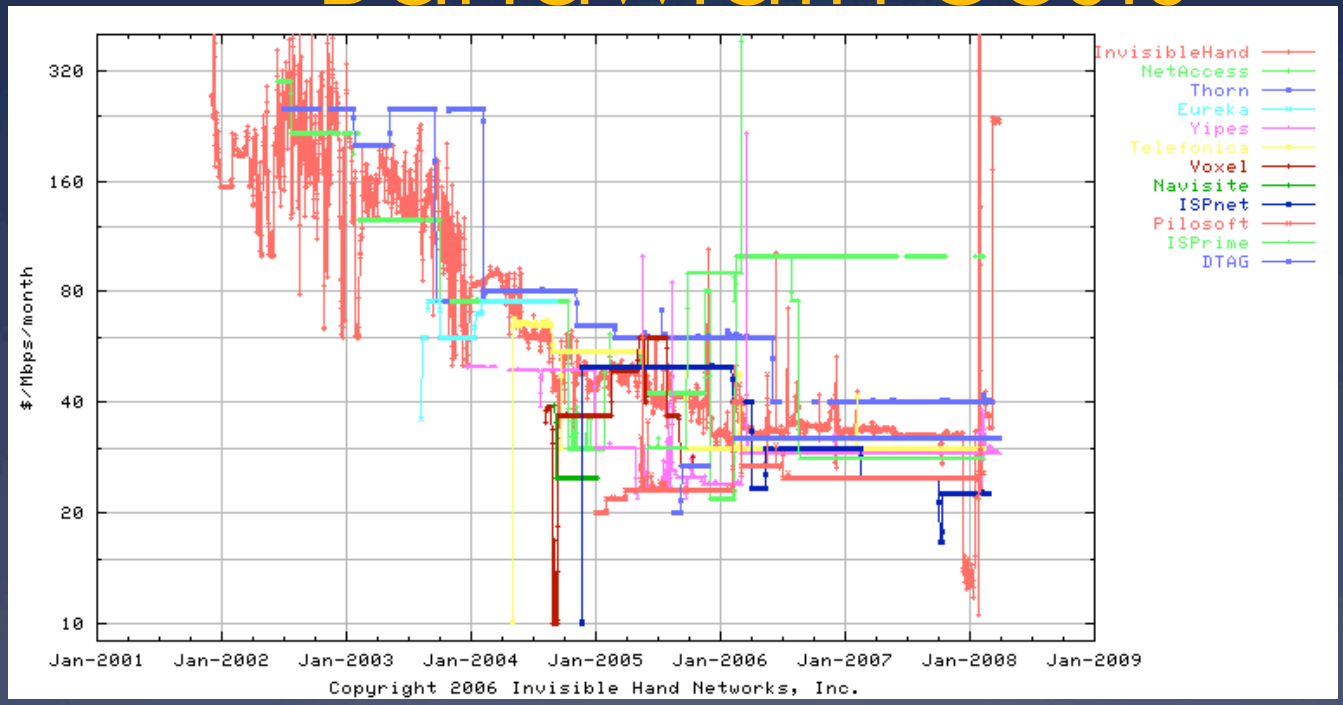

* Transit bandwidth: $40 Mb/s/month ~ $0.125/GB
* US colocation providers charge $0.30 to $1.75/GB
  * e.g., Amazon EC2 $0.17/GB (outbound)
  * CDNs: $0.08 to $0.19/GB

Economics of P2P

- Service provider view
  - save $150/month for single rented server in colo, with 2 TB bandwidth
  - but can handle 100,000 VoIP users
- But ignores externalities
  - home PCs can't hibernate → energy usage
    - about $37/month
  - less efficient network usage
  - bandwidth caps and charges for consumers
    - common in the UK
    - Australia: US$3.20/GB
- Home PCs may become rare
  - see Japan & Korea

Which is greener – P2P vs. server?

- Typically, P2P hosts only lightly used
  - energy efficiency/computation highest at full load
  - → dynamic server pool most efficient
  - better for distributed computation (SETI@home)
- But:
  - CPU heat in home may lower heating bill in winter
    - but much less efficient than natural gas (< 60%)
  - Data center CPUs always consume cooling energy
    - AC energy ≈ server electricity consumption
- Thus,
  - deploy P2P systems in Scandinavia and Alaska

Reliability

- CW: "P2P systems are more reliable"
- Catastrophic failure vs. partial failure
  - single data item vs. whole system
  - assumption of uncorrelated failures wrong
- Node reliability
  - correlated failures of servers (power, access, DOS)
  - lots of very unreliable servers (95%?)
- Natural vs. induced replication of data items

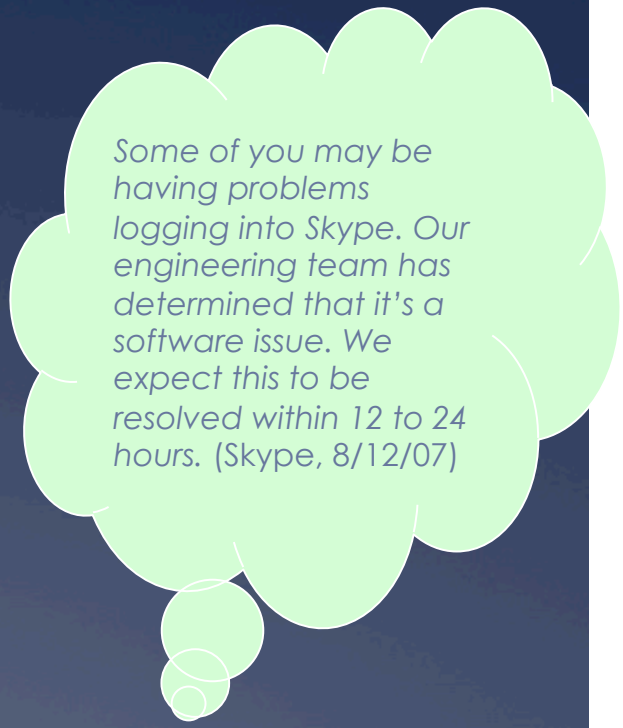

Security & privacy

- Security much harder
  - user authentication and credentialing
    - usually now centralized
  - sybil attacks
  - byzantine failures
- Privacy
  - storing user data on somebody else's machine
- Distributed nature doesn't help much
  - same software → one attack likely to work everywhere
- CALEA?

OA&M

- P2P systems are hard to debug
- No real peer-to-peer management systems
  - system loading (CPU, bandwidth)
    - automatic splitting of hot spots
  - user experience (signaling delay, data path)
  - call failures
- Later: P2PP & RELOAD add mechanisms to query nodes for characteristics
- Who gathers and evaluates the overall system health?

P2P for VoIP

The role of SIP proxies

tel:1-212-555-1234

REGISTER

sip:line1@128.59.16.1

sip:alice@example.com

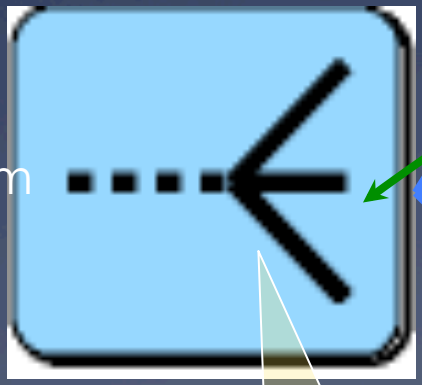

Translation may depend on caller, time of day, busy status, …

sip:6461234567@mobile.com

P2P SIP

- Why?
  - no infrastructure available: emergency coordination
  - don't want to set up infrastructure: small companies
  - Skype envy :-)
- P2P technology for
  - user location
    - only modest impact on expenses
    - but makes signaling encryption cheap
  - NAT traversal
    - matters for relaying
  - services (conferencing, transcoding, …)
    - how prevalent?
- New IETF working group formed
  - multiple DHTs
  - common control and look-up protocol?

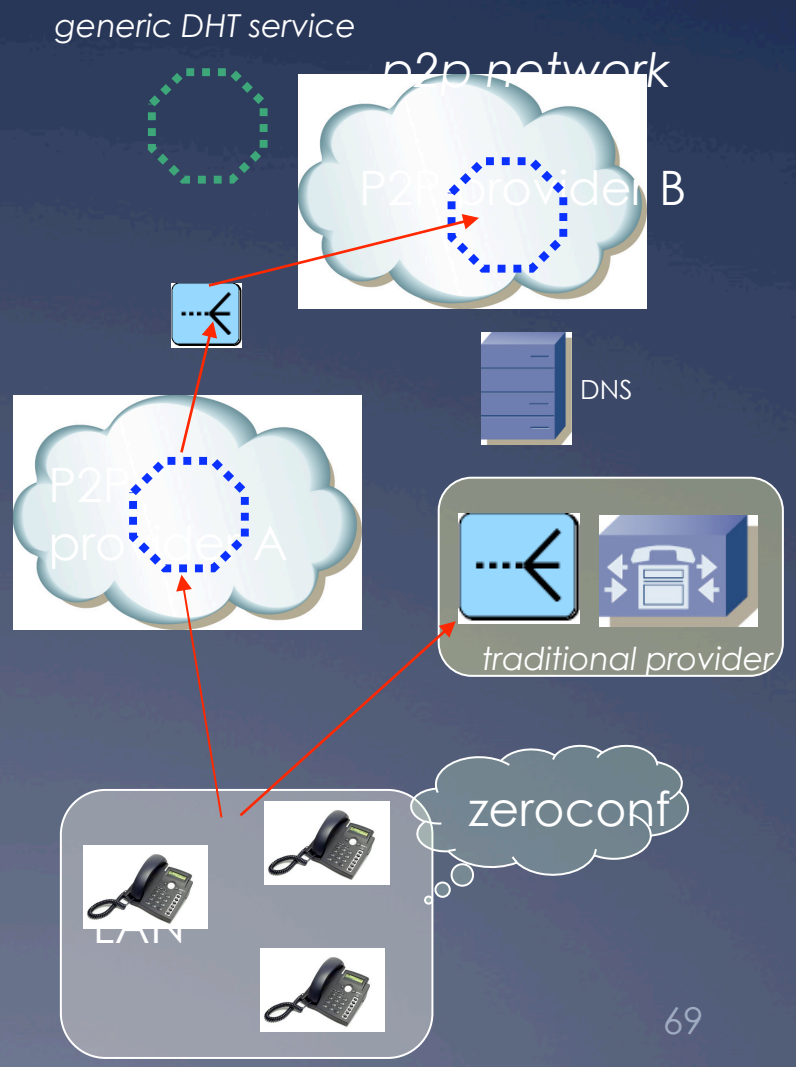

More than a DHT algorithm

Finger table

Tree

Routing-table stabilization

Periodic recovery

Lookup correctness

Prefix-match

Modulo addition

Routing-table size

Parallel requests

Recursive routing

Bootstrapping

Updating routing-table from lookup requests

Leaf-set

XOR

Proximity neighbor selection

Lookup performance

Reactive recovery

Successor

Hybrid

Proximity route selection

Strict vs. surrogate routing

Routing-table exploration

P2P SIP -- components

- Multicast-DNS (zeroconf) SIP enhancements for LAN
  - announce UAs and their capabilities
- Client-P2P protocol
  - GET, PUT mappings
  - mapping: proxy or UA
- P2P protocol
  - get routing table, join, leave, …
  - independent of DHT
  - replaces DNS for SIP and basic proxy

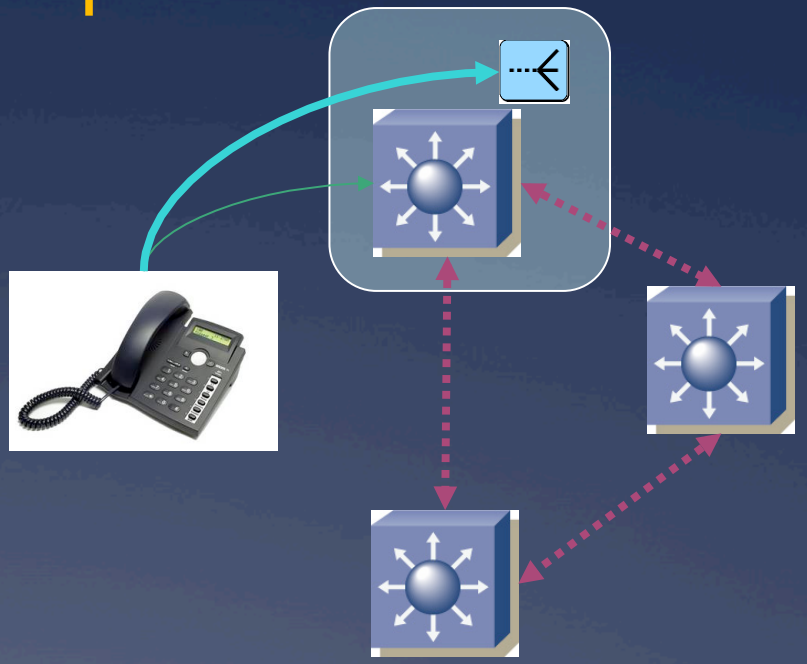

P2PSIP architecture

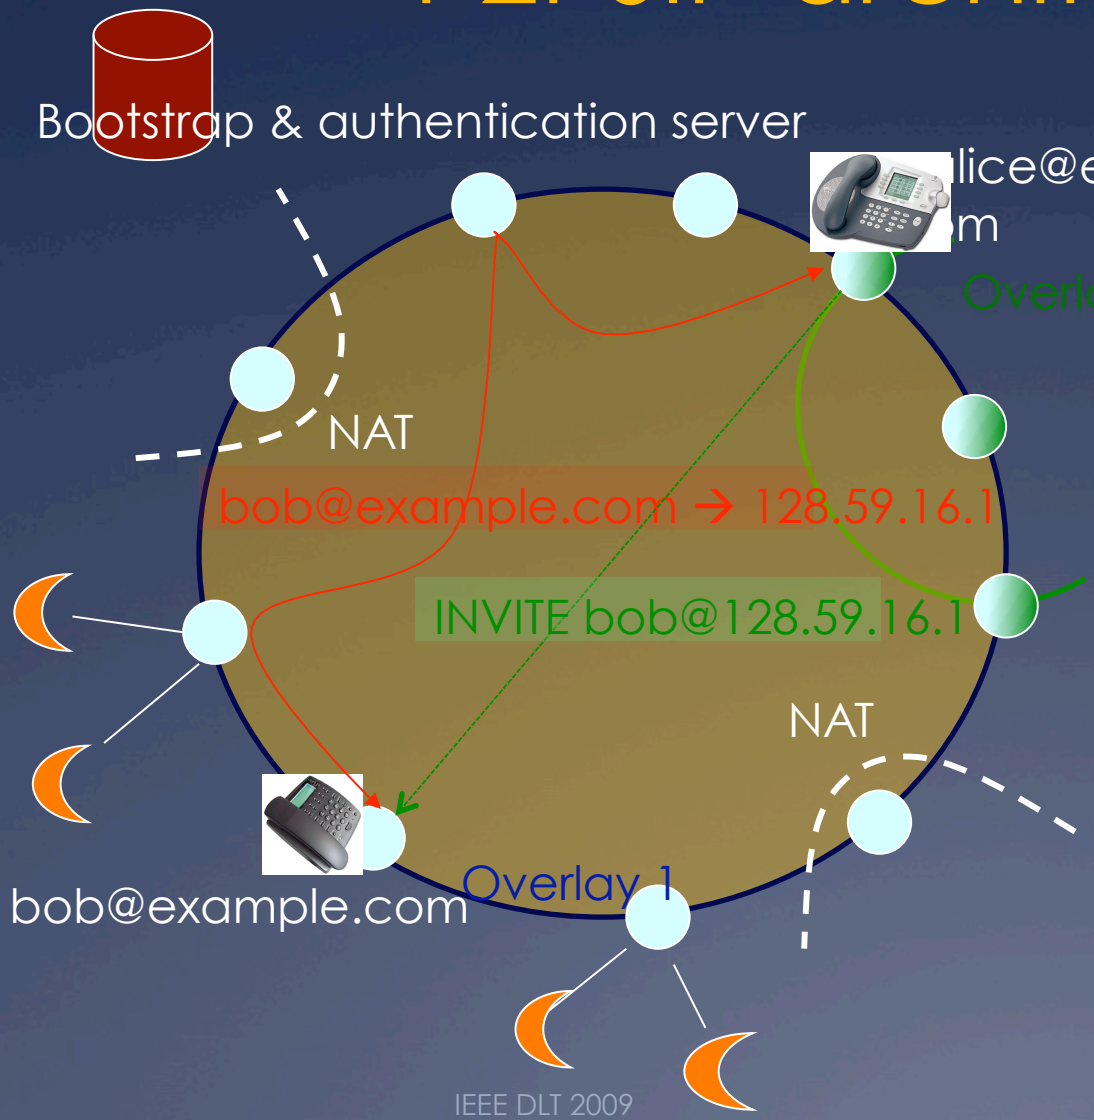

peer in P2PSIP

client

IETF peer-to-peer efforts

- Originally, effort to perform SIP lookups in p2p network
- Initial proposals based on SIP itself
  - use SIP messages to query and update entries
  - required minor header additions
- P2PSIP working group formed
  - now SIP just one usage
- Several protocol proposals (ASP, RELOAD, P2PP) merged
  - still in "squishy" stage – most details can change

RELOAD

- Generic overlay lookup (store & fetch) mechanism
  - any DHT + unstructured
- Routed based on node identifiers, not IP addresses
- Multiple instances of one DHT, identified by DNS name
- Multiple overlays on one node
- Structured data in each node
  - without prior definition of data types
  - PHP-like: scalar, array, dictionary
  - protected by creator public key
  - with policy limits (size, count, privileges)
- Maybe: tunneling other protocol messages

Typical residential access

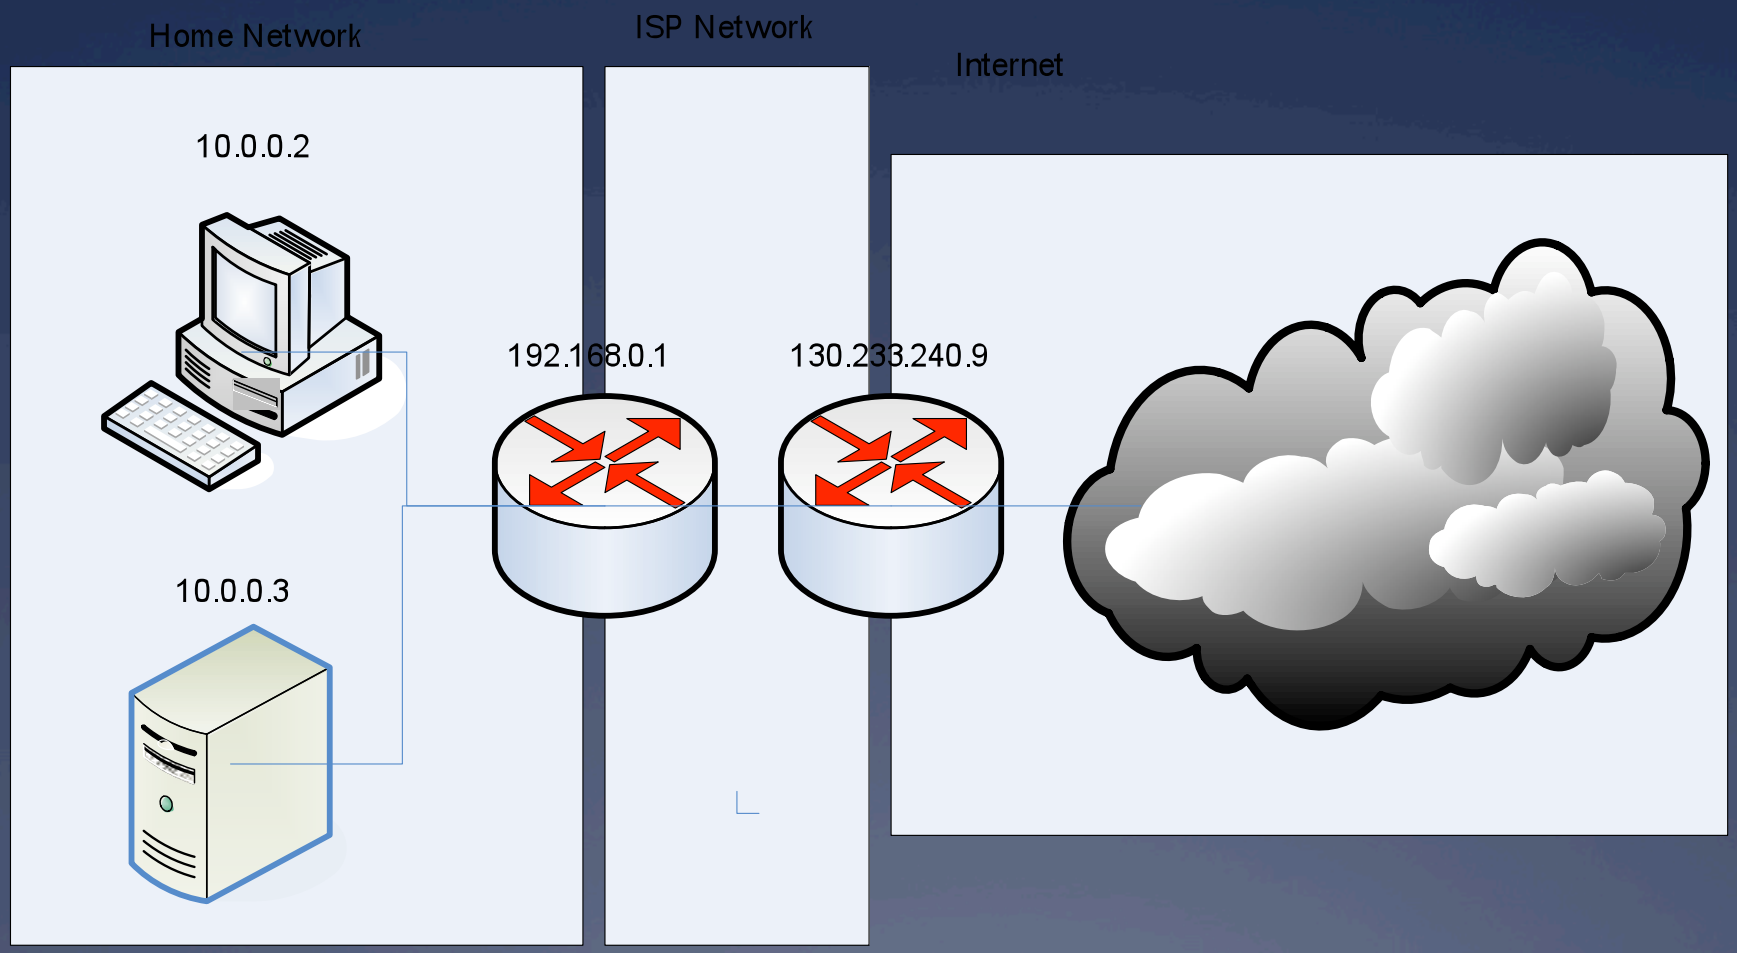

NAT traversal

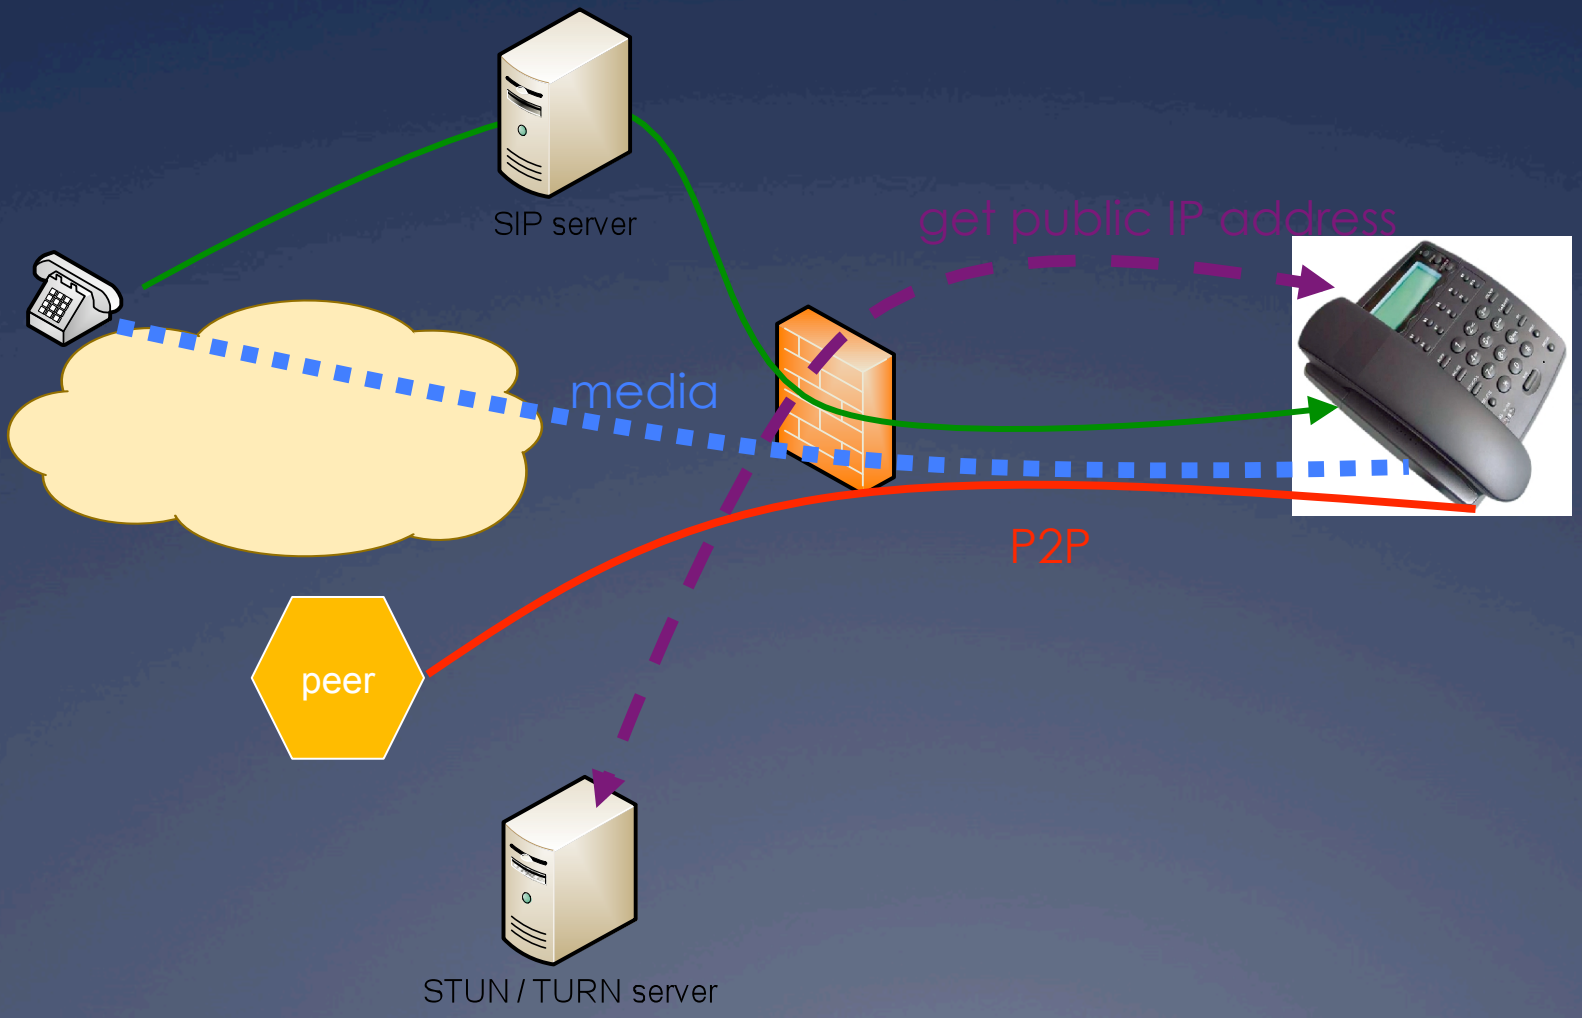

ICE (Interactive Connectivity Establishment)

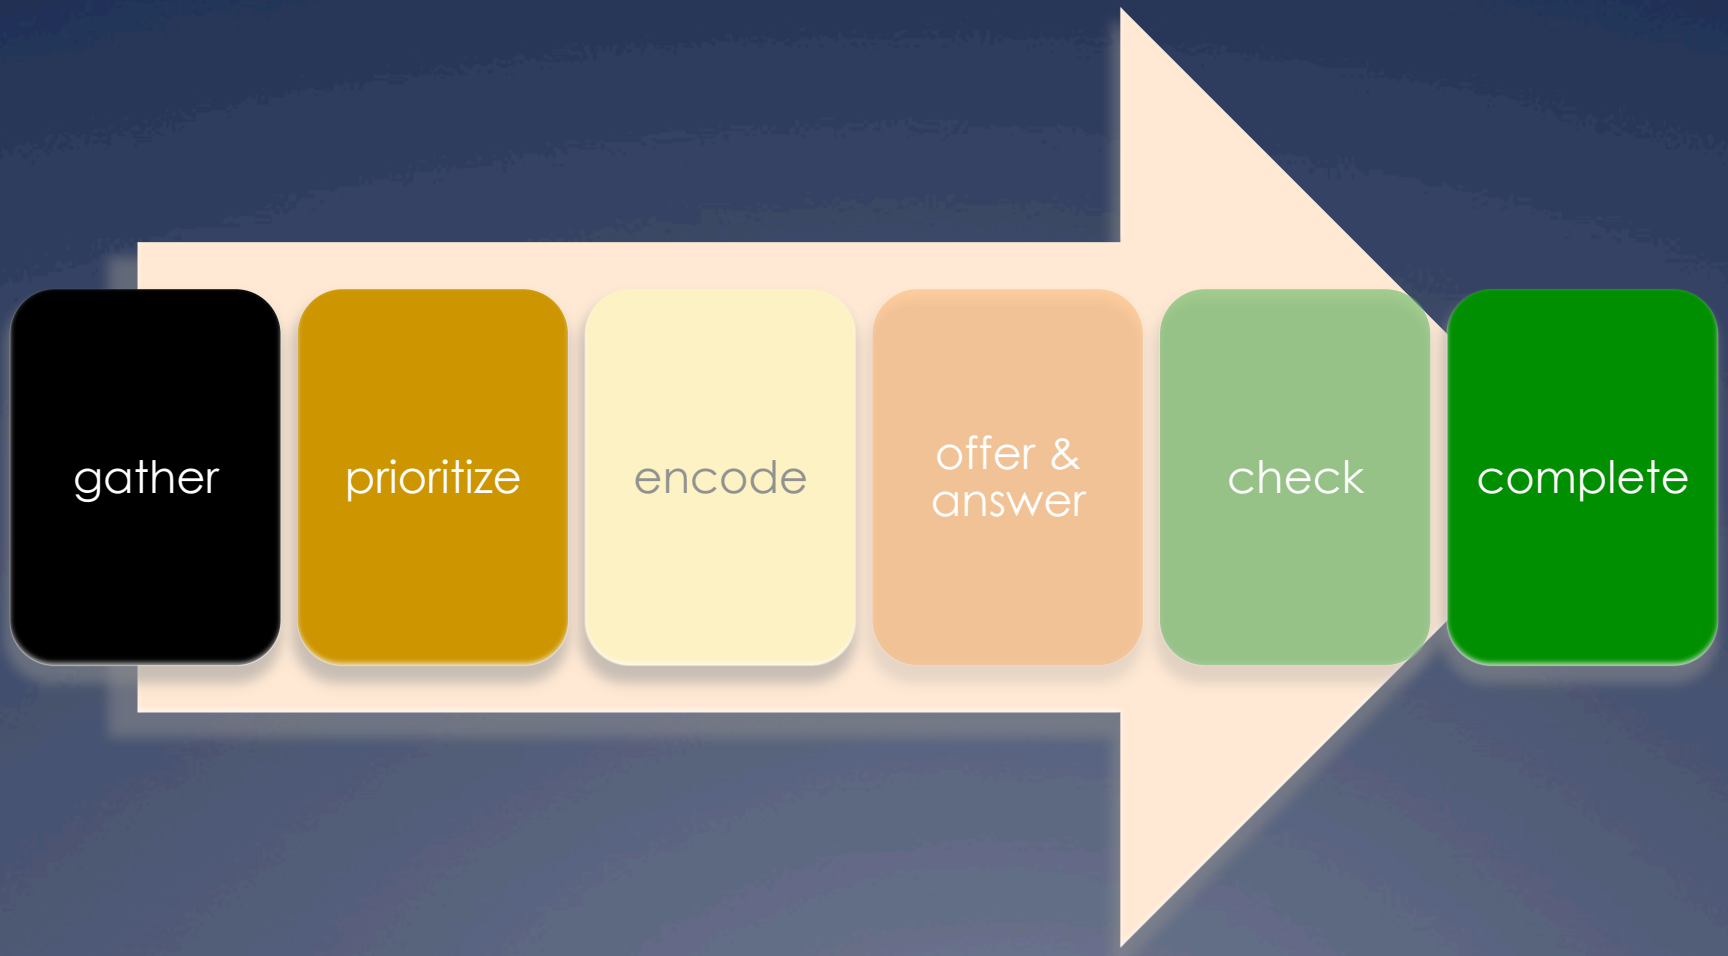

OpenVoIP snapshots

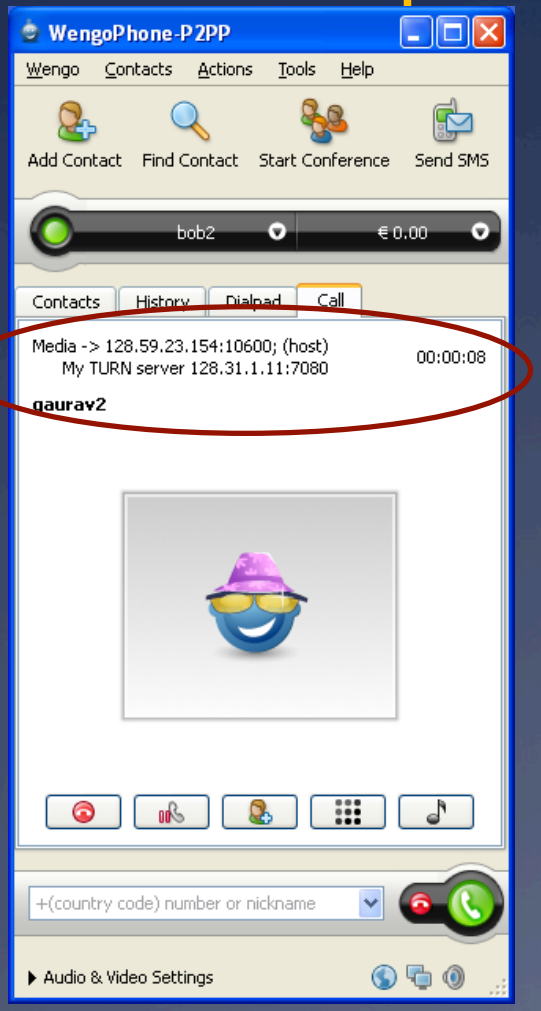

direct

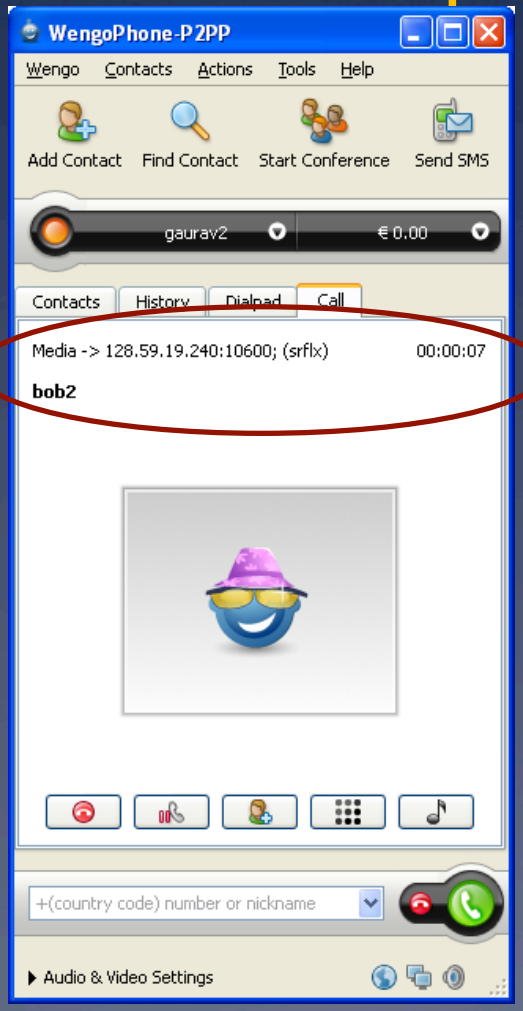

call through a NAT

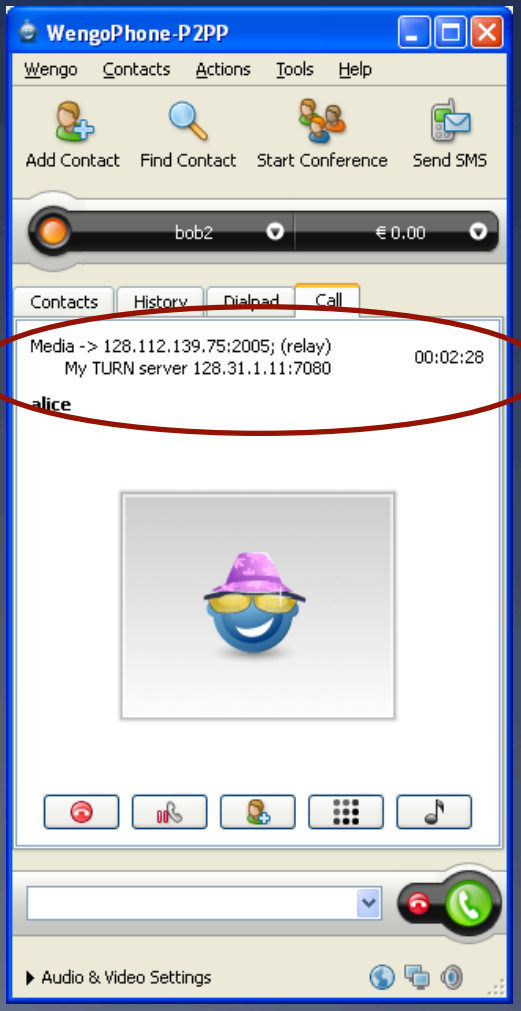

call through a relay

OpenVoIP snapshots

OpenVoIP snapshots

* Tracing lookup request on Google Maps

Operations | Hotmetrics | **Results**

Result:

SIPLookup: unhashed-id: test_call
To: 160.193.163.102:10080
302 response: Next hop: 128.112.139.75:7080
302 response: Next hop: 143.107.111.194:10080
302 response: Next hop: 169.229.50.14:9080
OnSIPLookup: 200
sip:test_call@128.59.19.152:5060

Query time taken: 474.631 ms

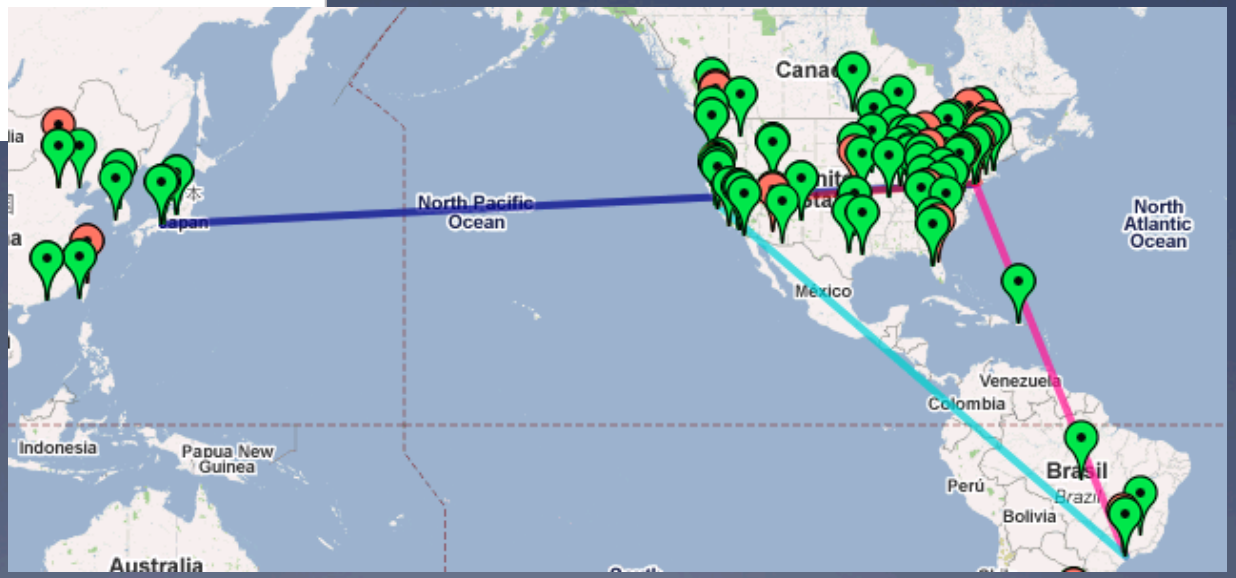

Integrating cellular and 802.11

Integrating VoIP and Cellular

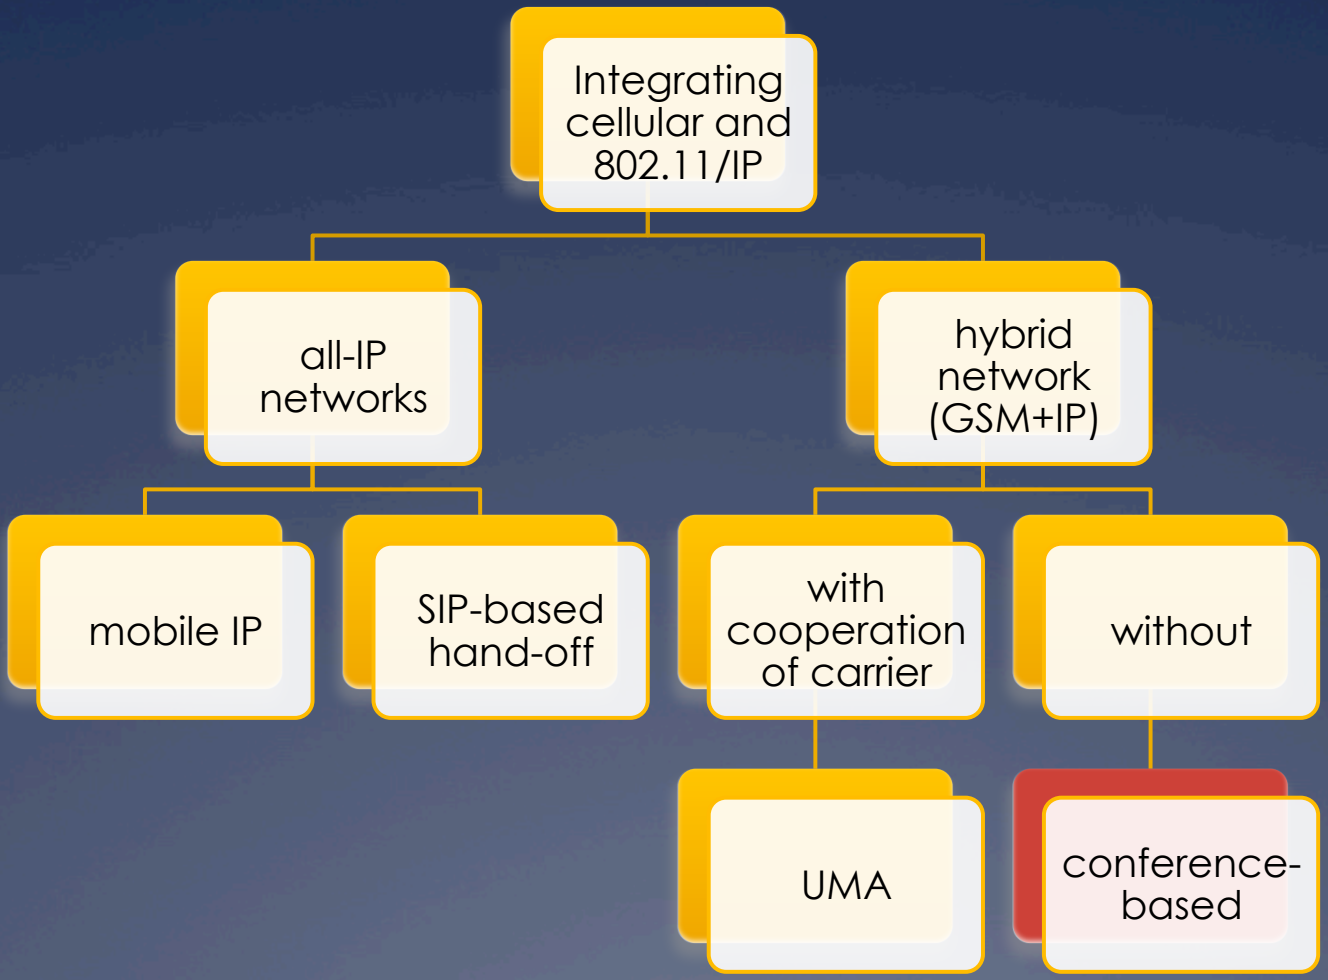

VoIP+GSM testbed

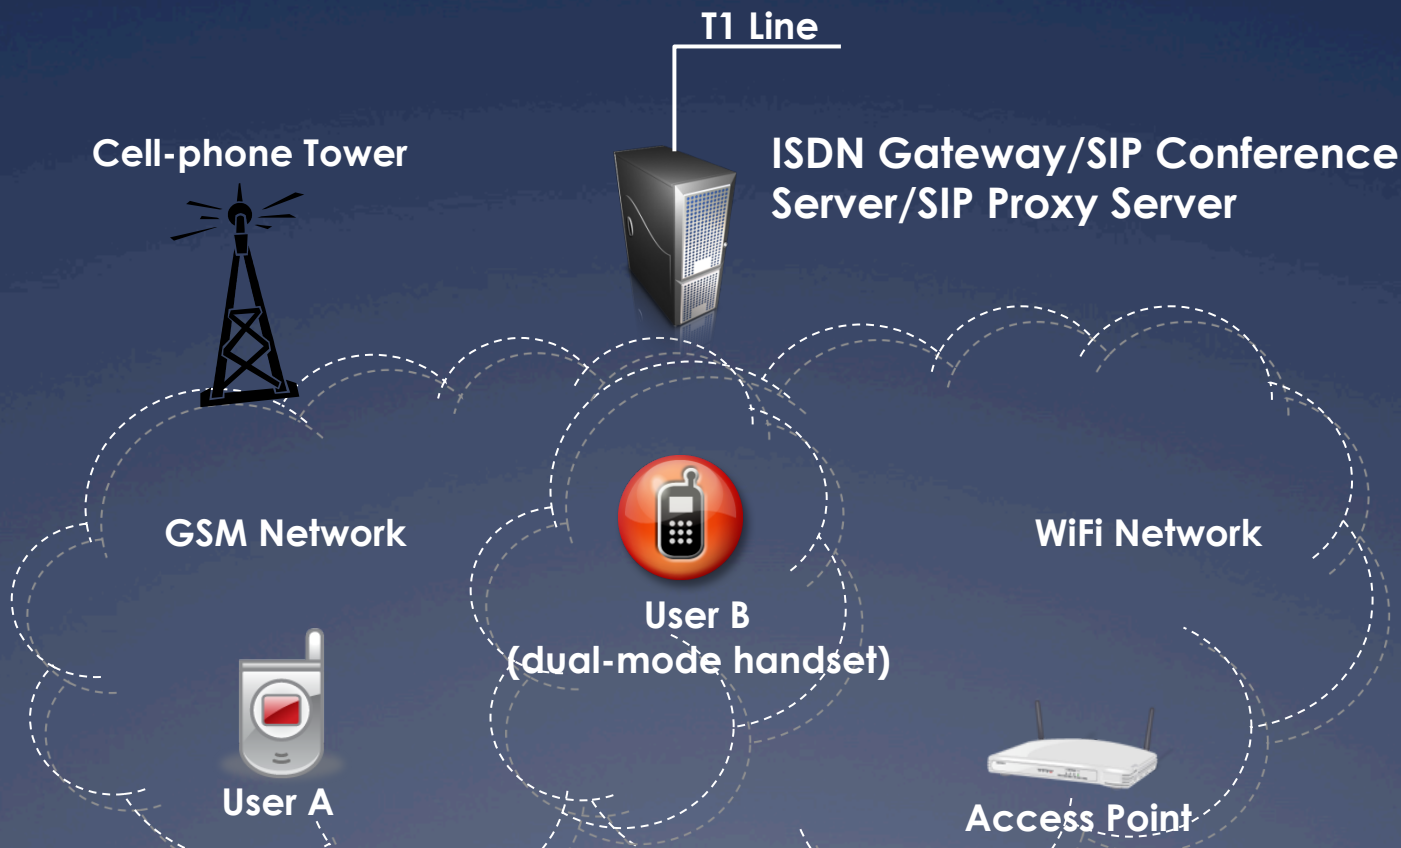

- Dual-mode handset
  - IP interface: X-Lite client
  - GSM interface: Nokia cellphone

Experiments

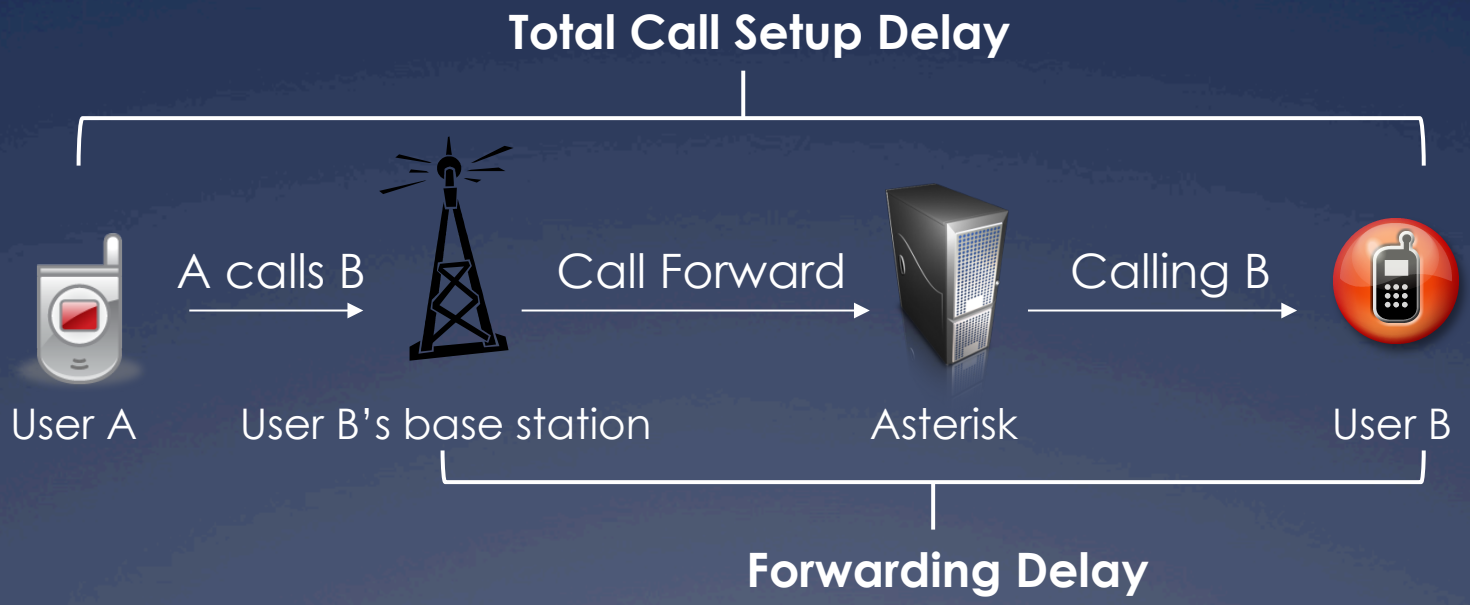

Type of call (A → B)	Forwarding delay	Call-setup delay
Cell-to-cell *	6.7 s	9.6 s
Cell-to-IP **	3.1 s	6.2 s

Fax-over-IP

Fax-over-IP

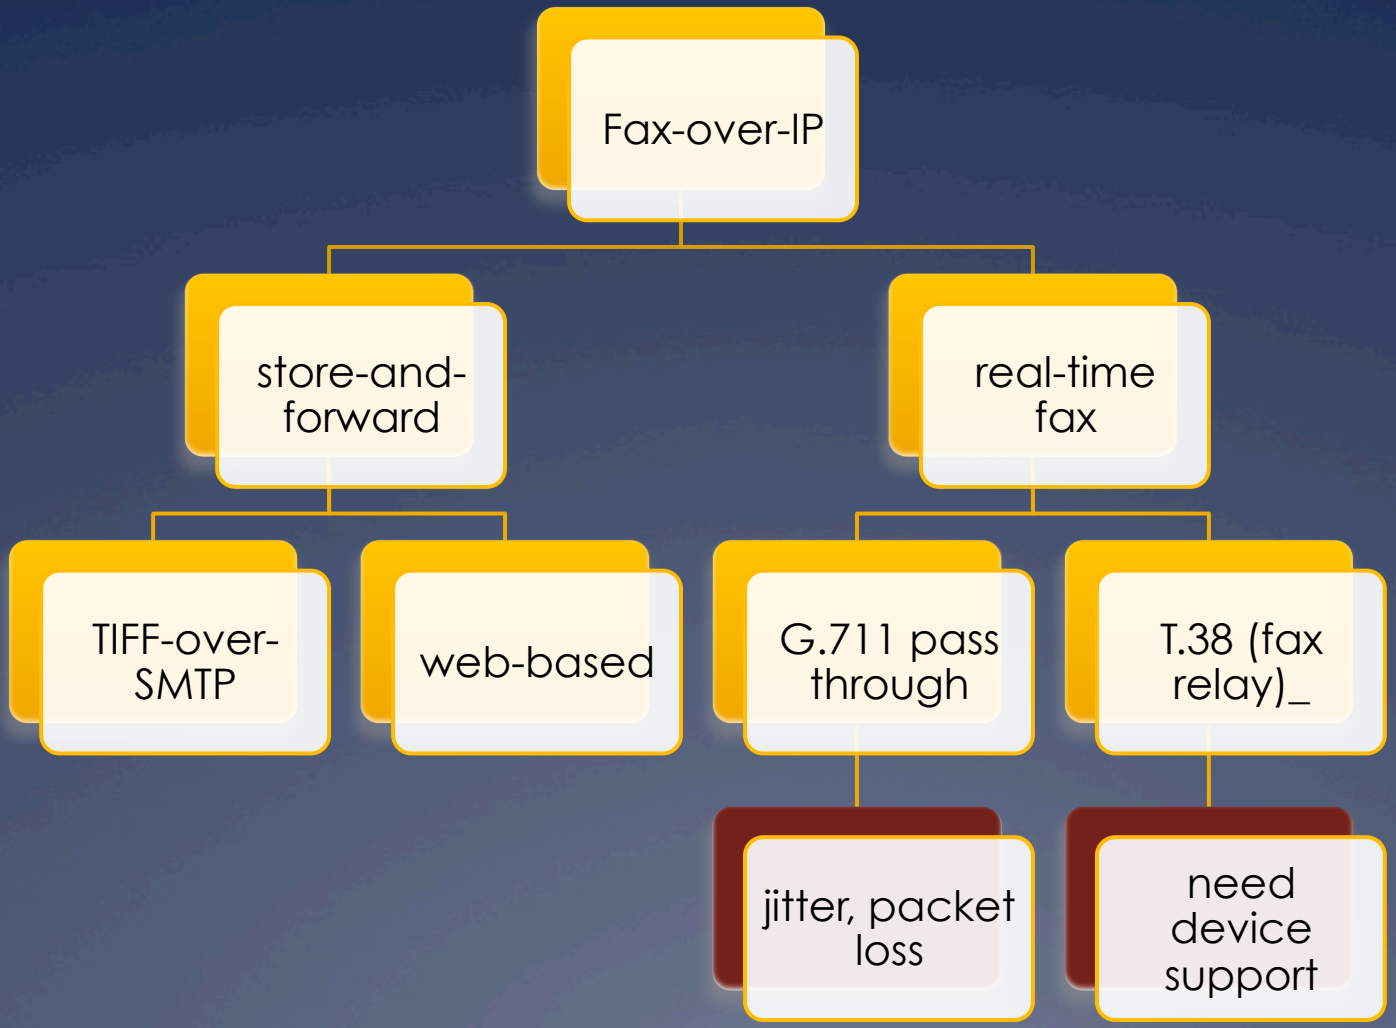
Fax-over-IP
store-and-forward
real-time fax
TIFF-over-SMTP
web-based
G.711 pass through
T.38 (fax relay)_
jitter, packet loss
need device support

Fax pass-through

- Uses G.711 over RTP
  - fax signaling events (RFC 3665)
  - other codecs may not reproduce modem tones
- May be sensitive to packet-specific distortions
  - bit errors → packet loss bursts
  - jitter → delay adaptation gaps
- Fixes:
  - PLC in terminal adapter
  - FEC in RTP stream
  - T.38 in gateway?

Standards & interoperability

Interoperability

- Generally no interoperability problems for basic SIP functionality
  - basic call, digest registration (mostly...), call transfer, voice mail
- Weaker in advanced scenarios and backward compatibility
  - handling TCP, TLS
  - NAT support (symmetric RTP, ICE, STUN, ...)
  - multipart bodies
  - SIP torture tests
  - call transfer, call pick-up
  - video and voice codec interoperability (H.264, anything beyond G.711)

- SIPit useful, but no equivalent of WiFi certification
  - most implementations still single-vendor (enterprise, carrier) or vendor-supplied (VSP)
  - SFTF (test framework) still limited
- Need profiles to guide implementers

- A role for public shaming?

Trouble in Standards Land

- Proliferation of transition standards: 2.5G, 2.6G, 3.5G, …
  - true even for emergency calling…
- Splintering of standardization efforts across SDOs
  - primary:
    - IEEE, IETF, W3C, OASIS, ISO
  - architectural:
    - PacketCable, ETSI, 3GPP, 3GPP2, OMA, UMA, ATIS, …
  - specialized:
    - NENA
  - operational, marketing:
    - SIP Forum, IPCC, …

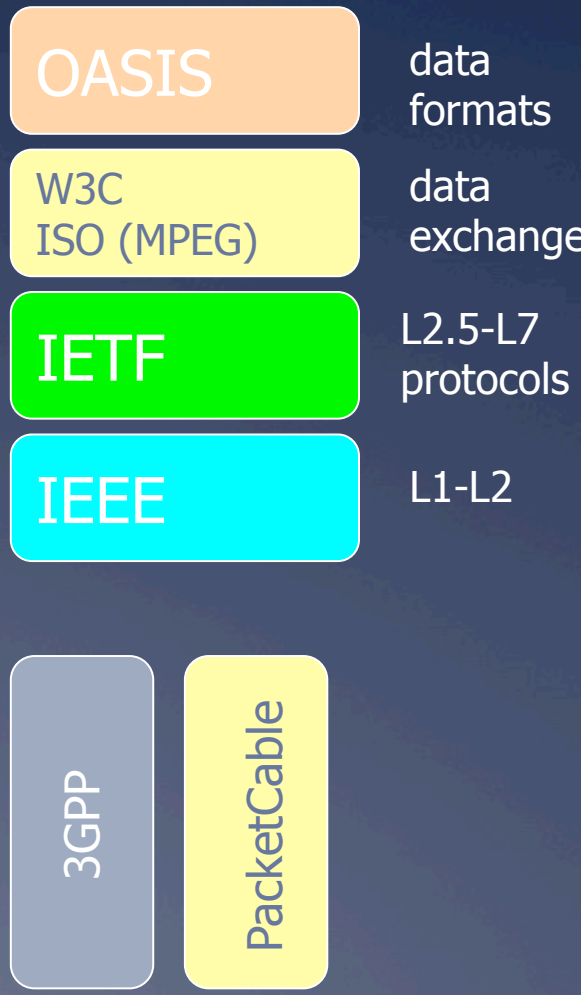

IETF issues

- SIP WGs: small number (dozen?) of core authors (80/20)
  - some now becoming managers…
  - or moving to other topics
- IETF: research → engineering → maintenance
  - many groups are essentially maintaining standards written a decade (or two) ago
    - DNS, IPv4, IPv6, BGP, DHCP; RTP, SIP, RTSP
    - constrained by design choices made long ago
  - often dealing with transition to hostile & "random" network
  - network ossification
- Stale IETF leadership
  - often from core equipment vendors, not software vendors or carriers
- fair amount of not-invented-here syndrome
  - late to recognize wide usage of XML and web standards
  - late to deal with NATs
  - security tends to be per-protocol (silo)
    - some efforts such as SAML and SASL
- tendency to re-invent the wheel in each group

Conclusion

- Even after 10+ years, VoIP mostly still "cheaper calls"
- New services and models:
  - (rich) presence
  - location-based services
  - user-programmable services
  - P2P SIP
- Scaling to carrier-scale and under duress
- Current standardization processes slow and complexity-inducing