Nominum

The Internet and Identifiers

Paul V. Mockapetris

Sigcomm 2005

8/23/2005

What are Today's Digital Identifiers?

Nominum.

- Conventions associating one piece of data to another
  - www.nominum.com to see web page
  - "Anna Kournikova" into Google window
  - Shell.nominum.com for SSH
  - 160.192.177.128.in-addr.arpa for email verification
  - pvm@Nominum.com for email
  - pvm@a21.com to log on to Amazon
  - Dial +1-650-381-6100 on a phone
- Anything we type or click on to identify what we want
- The first step in any communication; they are the nouns and pronouns of the language of the Internet
- The ultimate way to get paid per click

One Way to Evaluate Their Significance…

Nominum.

- .COM
  - Verisign has $6.5 billion market capitalization
  - Registrar gets $2+ per name at retail
  - Registry (central database) gets $6 per name
  - Over 30,000,000 names in .com
- Google
  - ~~$46~~ ~~$80~~ $77 billion market cap
- Phone numbers
  - In 2002, US phone companies, desperate for cash, raised over $10 billion by selling phone directory operations

The technology landscape

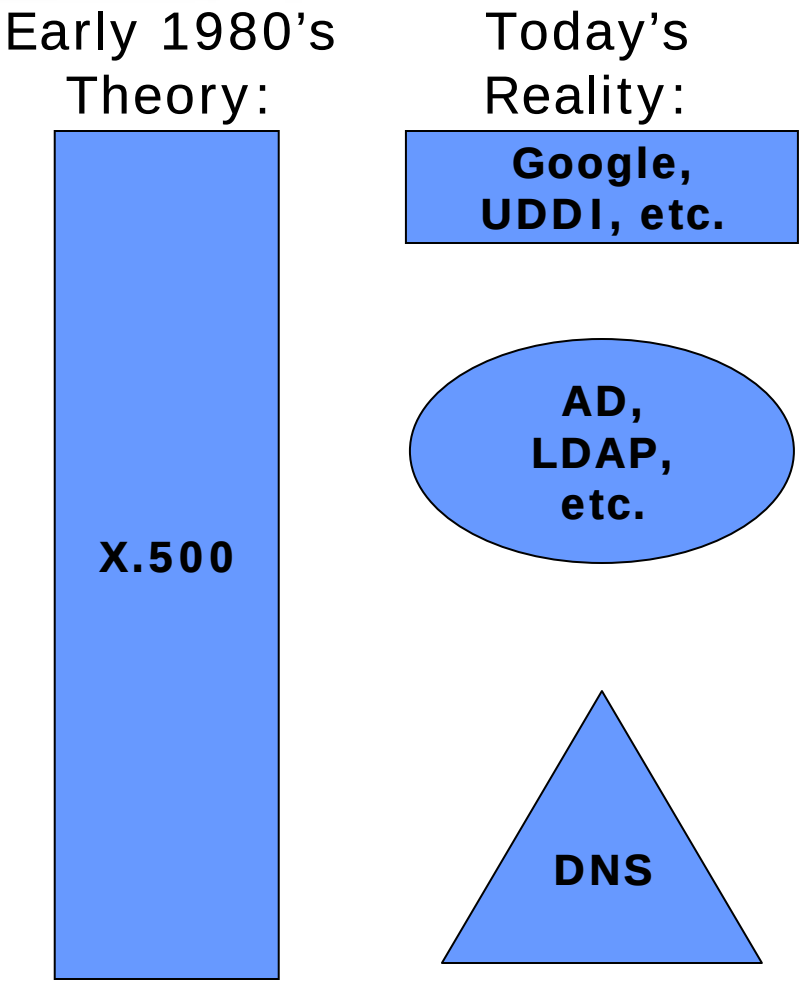

- In the beginning, theory said there would be one monolithic service – X.500
  - Searches
  - Lookups
  - Schema
  - Access Control
- In practice, there are many services & applications, with different properties, at 3 levels:
  - Web
  - Directory
  - DNS

Niches and specialization

Nominum.

	Openness	Speed	Reach	Data Format	Functions
Web Based	**Usually proprietary**	**Seconds**	**Internet subset**	**Varies**	**Any •SEARCH**
Directory	**Mostly open**	**10+ millisecond**	**Single organization**	**Heavily structured**	**•search • Lookup • Update**
DNS	**Open & interoperable**	**Sub millisecond**	**Internet & intranet Universal**	**Slightly structured**	**•Lookup •Update**

Architectures that Create Digital Identifiers

Nominum.

DNS

AD, LDAP, etc.

Google, UDDI, etc.

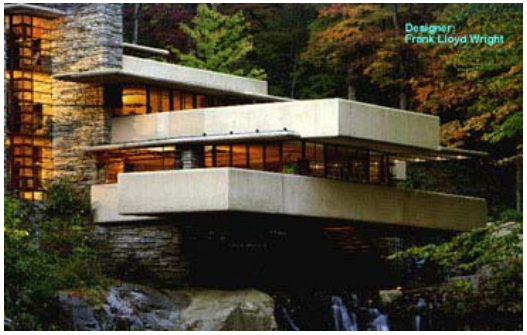

Is this separation natural?

Nominum.

- Should we / will we always have a speedy lower layer that spans the Internet?
- Does Moore's law trump efficiency?
- Does Darwin favor AD over open source LDAP simply because schemas can be enforced?

Conjectures for today

Nominum.

- We need innovation at all levels of these systems.
- We can learn from experience.
- There's no guide for what the Internet should look like, we have to create a vision.
- We can imagine what a DNS replacement might do.

(For the rest of the talk, assume: DNS=today's DNS or its successor)

The Obstacles

Nominum.

- TDOS attacks
  - Technobabble attacks, e.g. can't add generic TLDs because of security and stability concerns, but can add 200 country TLDs
  - The cure: Be objective.
- EDOS
  - Everything changes the Internet; you can't build a useful service that satisfies every bureaucrat in every country and the IETF …
  - The Cure: Build tools that are orthogonal.

What does a DNS system do?

- 3 original (1983) functions:
  - Distribute itself
  - Provide host names
  - Be extensible
- Today
  - Tens of applications and datatypes added
  - VOIP & ENUM & URIs
  - RFID – it's the standard, stupid
    - Unify 6+ numbering schemes
  - モッカペトリス.jp, 모카페트리스.kr, 莫卡派乔斯.cn
  - May have dozens of DNS administrators in an enterprise
- DNS is the distributed database of the Internet

DNS use is growing exponentially

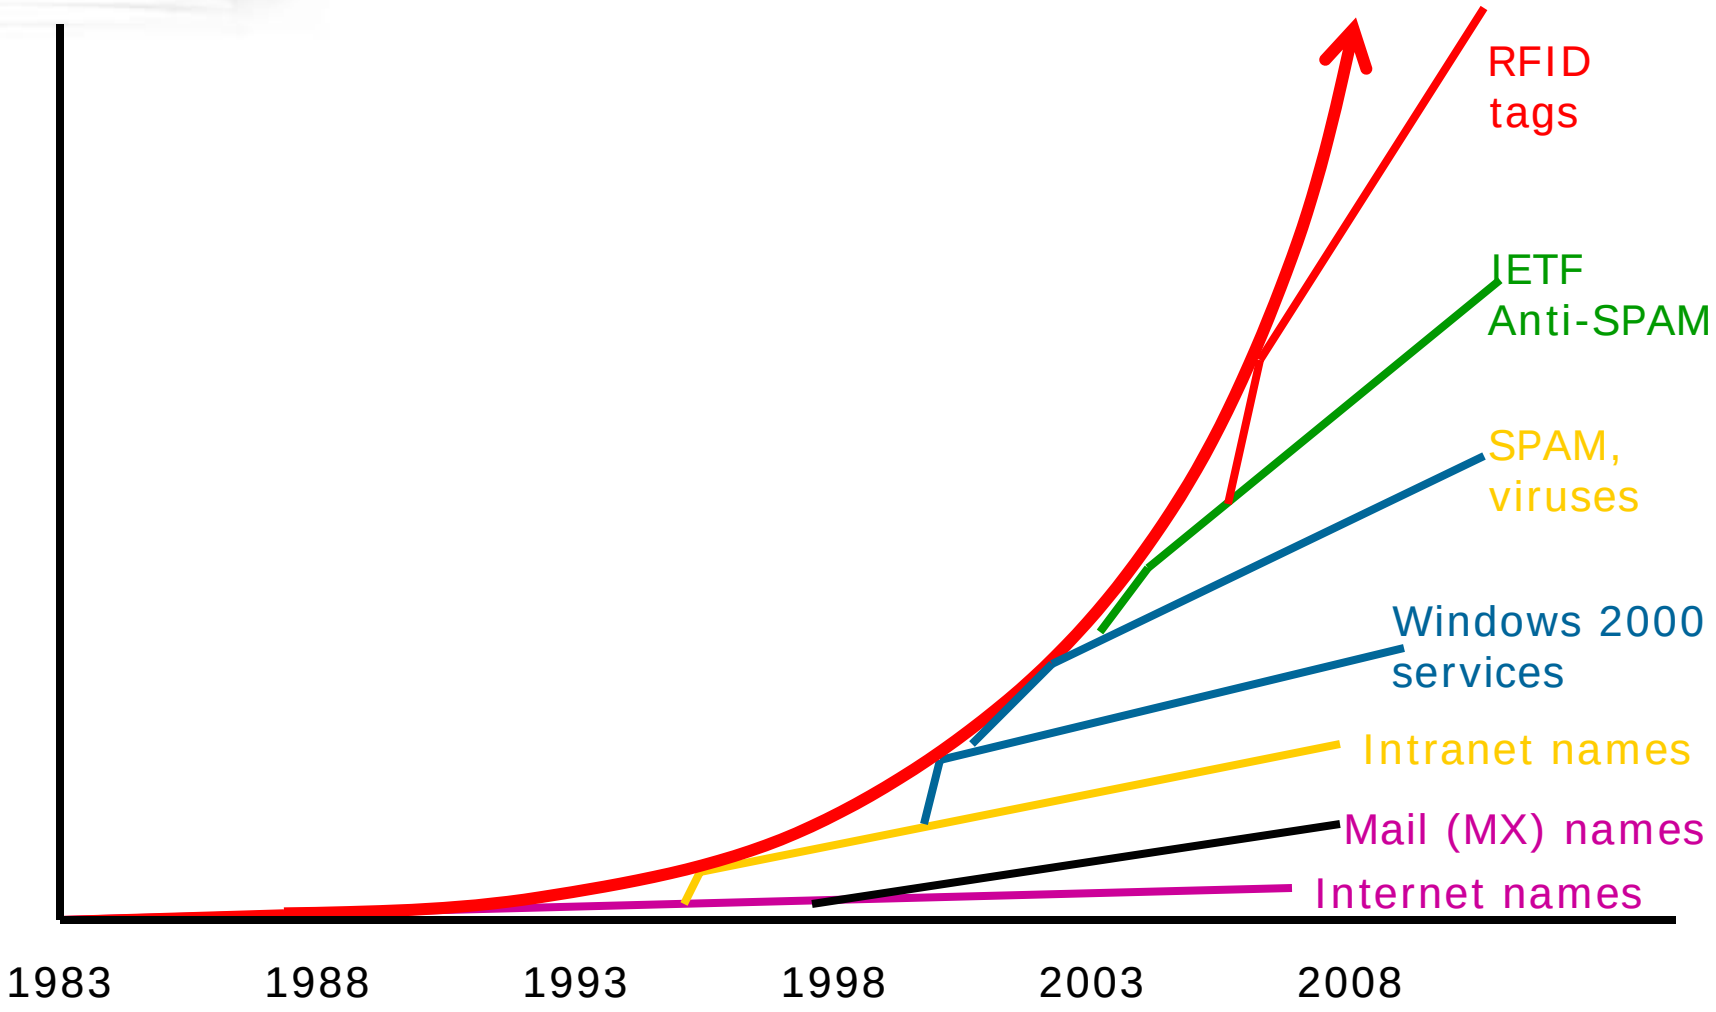

How large is DNS?

Nominum.

The largest distributed database in the world!

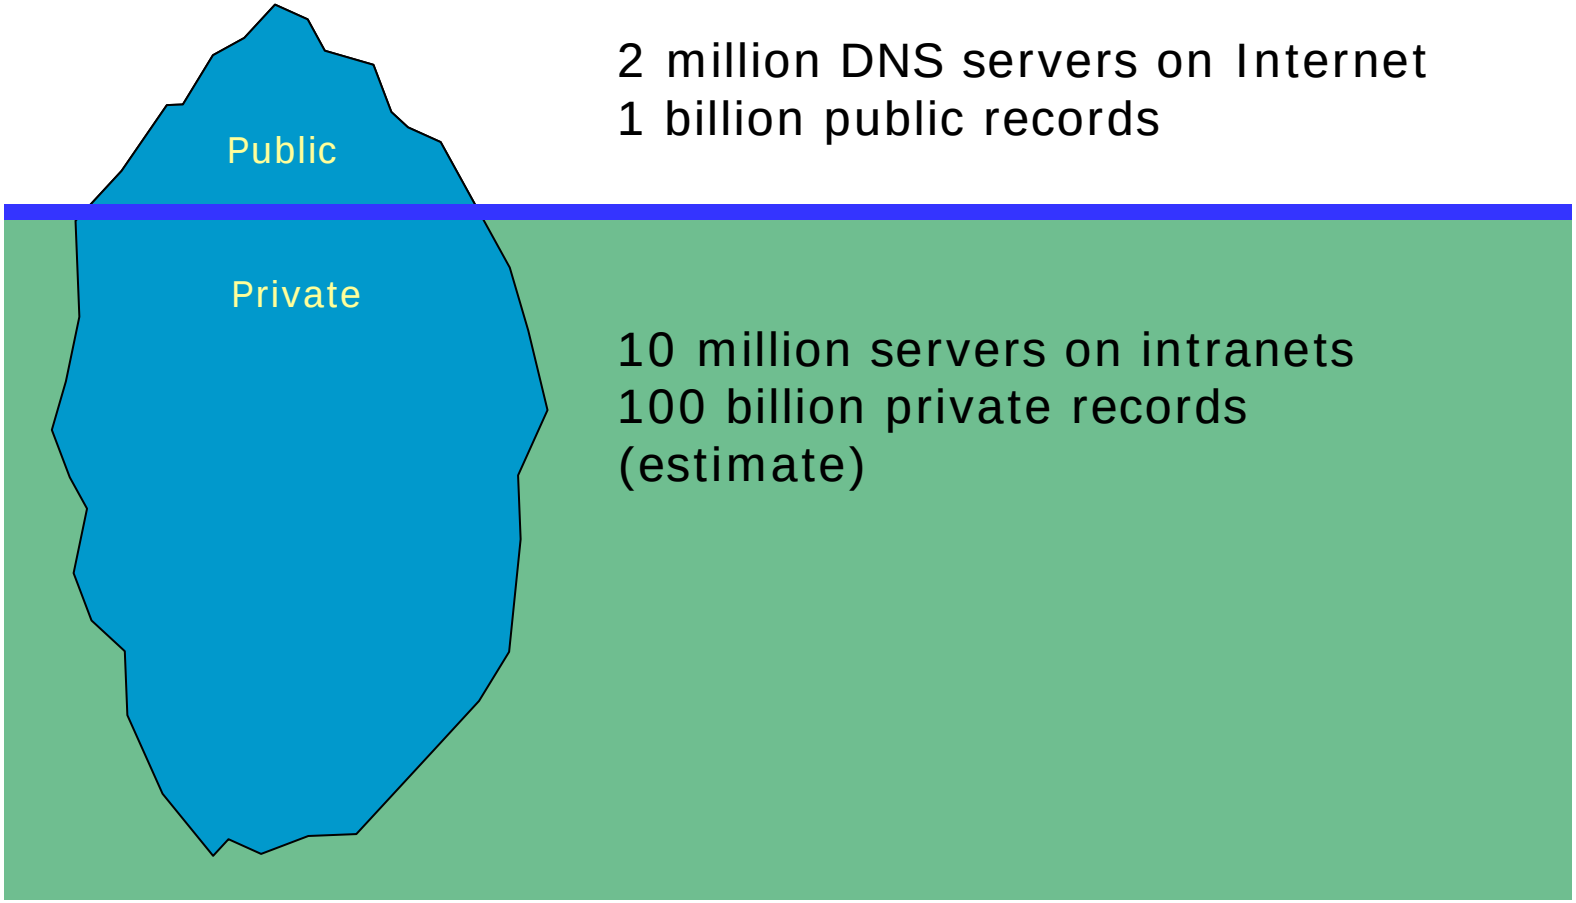

2 million DNS servers on Internet
1 billion public records

10 million servers on intranets
100 billion private records
(estimate)

Evolution of DNS data

MD, MF

More than one type of answer, multiple instances

MX

One type, separate weights, post query selection

SRV

Multiple metrics

NAPTR

Send a “program” as an answer; Compute local custom answer

Learning from experience

Nominum.

- How do we add an application?
- Marid
- RFID
- ENUM

Add an application to DNS

Nominum.

- Map name space onto DNS name space
- Add data at nodes
- See RFC 1101, TPC.INT

- Invented multiple times.
- Patented multiple times.

MARID *et al.*

Nominum.

- MX mail routing was the first new application added to the original DNS.
- Recently we had about 10 new proposals for ways to stop email when its spam; pretty much all used the DNS to store one form of authentication info or another
- Should have been easy
  - We know how to map mail addresses
  - Just decide on the data formats
- Has not been easy; Cisco's DKIM is the latest

RFID's Origins

Nominum.

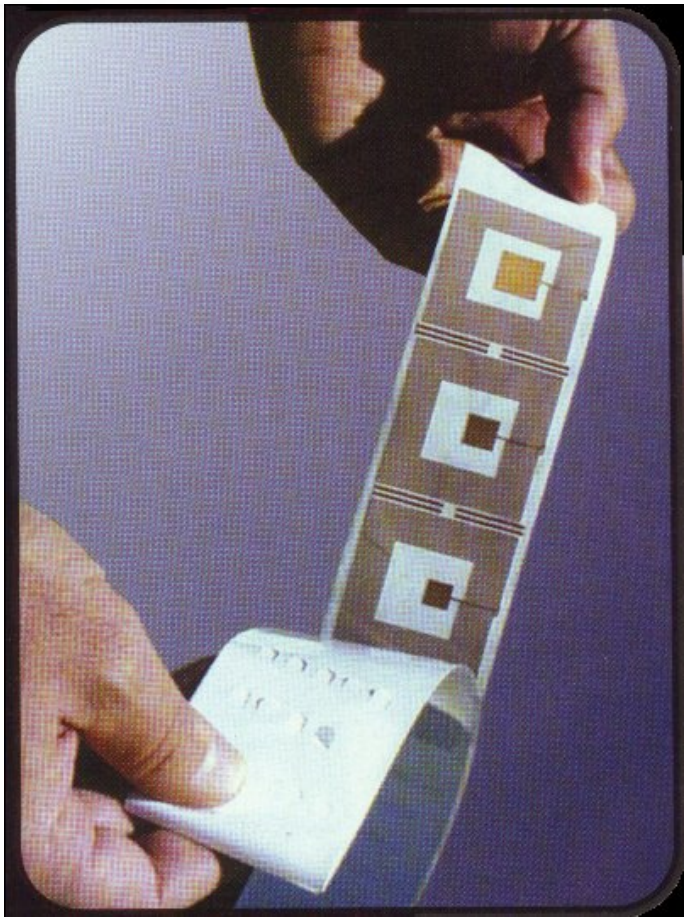

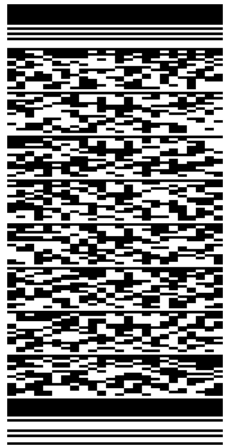

Why RFID is hard

Nominum.

- Legacy
  - Multiple existing name spaces
  - Multiple objectives (e.g. pallets vs. razor blades)
  - Varying Tag intelligence
    - Active (powered)/passive
    - Internal smarts
- Future
  - Privacy concerns
  - Standards body structure
    - Hardware IPR vs. software IPR

How we got to today

Nominum.

- MIT AutoID Center, with industry, defines:
  - Set of physical tag standards
  - Format for the binary string tags return
- Results turned over to EPCGlobal, a standards organization, with bar code experience, et al.

The Curious Devolution of the ONS Standard

Nominum.

- MIT Auto-ID Center defines
  - 96 bits of data per RFID tag
  - Object Naming System (v 0.5)
    - Layer over DNS
    - Variable sequence of fields for encoding all 96 bits a la subnetting inside an IP address
    - Different number trees could use different structures
    - Customize by
      - Computing the query
      - Customize the result

The Curious Devolution of the ONS Standard

Nominum.

- EPC Global “improves” to
  - 96 bits of data per RFID tag
  - Object Naming System (v 1.0)
    - Layer over DNS
    - Fixed 3 levels
      - Header (numbering scheme)
      - General Manager (subowner of name space, e.g. company)
      - Object Class (e.g. SKU)
    - Remaining bits up to other protocol
  - This allows different industry verticals to keep incompatible protocols and numbering formats

The Curious Devolution of the ONS Standard

Nominum.

- "Logic" behind the solution
  - If you can query individual serial numbers, there will be too much network traffic.
  - If there are errors reading tags, you can get the wrong unit data.
  - We need more powerful query technology.

- Bottom line: Database may fragment along industry verticals.
  Will database be like LDAP?
  (powerful but incompatible)

What's today's purpose of ENUM?

Nominum.

well known and standardized telephone number

The data might be:

- URI of a SIP phone
- Mailbox for voicemail

The ENUM data economy

- "Owners" of data
  - Multiple service providers: TDM, VOIP, VM …
  - Individuals
  - Registrars / Outsourcers
- "Slicers and dicers"
  - Verisign, Neustar
  - Private peers
- DNS transit
  - Complete datasets, queries/dips
- Post processing, local updates

What is ENUM?

Nominum.

The best hope for an open-standards-based approach to communications identifiers and signaling for the next decade:

Phone Numbers in the DNS

(but not just phone numbers)

Types of ENUM Deployments

Nominum.

Public ENUM

Publicly available, shared database

Carrier ENUM

Database shared on the basis of bi- or multi-lateral agreements

Private ENUM

Non-public database

Where does DNS appear?

ENUM DNS Service

3. DNS returns NAPTR record containing SIP URL to Calling Party UA

2. Calling party proxy UAC queries DNS for endpoint location

Query
4.3.2.1.5.5.2.0.2.1.e164.arpa?

Response
sip:name@domain.com

"Call Setup"

1. The caller dials the person's telephone number

Dial
+1-202-555-1234

Sip Proxy

Sip
sip:name@domain.com

4. Calling party UA connects the call

Sip Proxy

Why multiple ENUM types?

Nominum.

- Theory One:
  - The Internet wasn't relevant until there were multiple networks.
  - ENUM won't be relevant until we get a critical mass of VOIP implementations that use/need it.
- Theory Two:
  - Its just a matter of preserving ownership/control of something valuable, e.g.
    - Inside an enterprise
    - Between partners
    - Outsourcing while owning
    - Can Internet style ENUM triumph?

The Situation: Islands of VoIP Connected through the PSTN

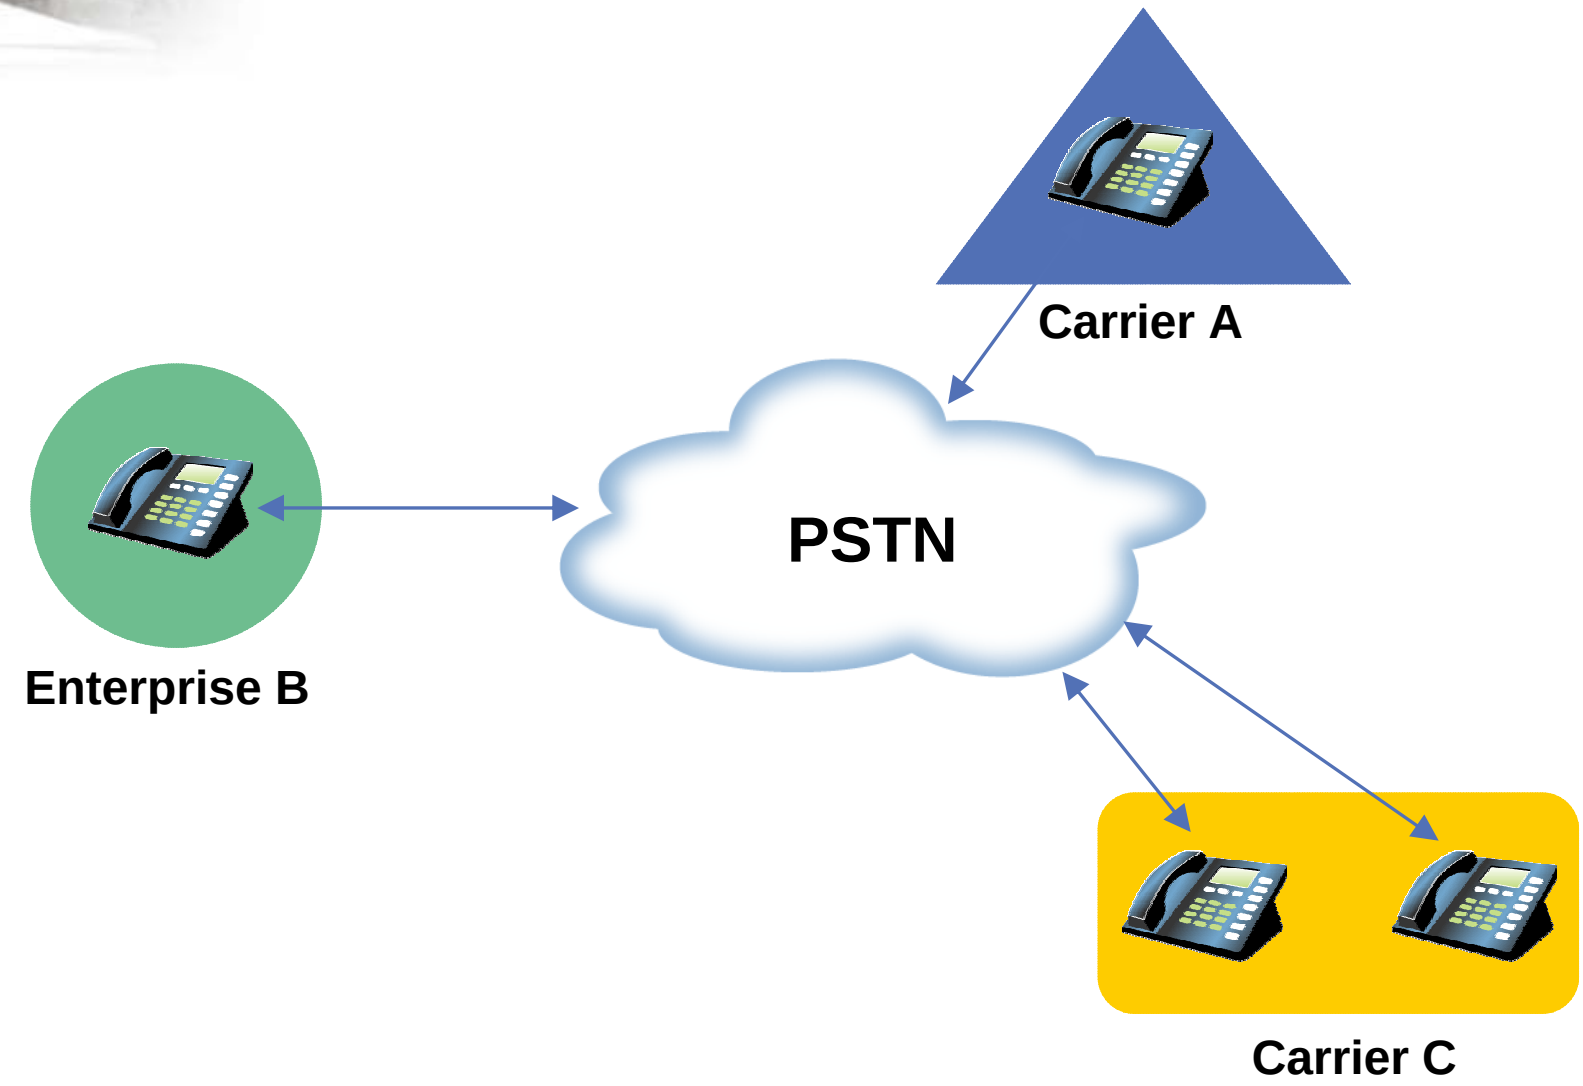

Why Internet ENUM?
Efficient Communications

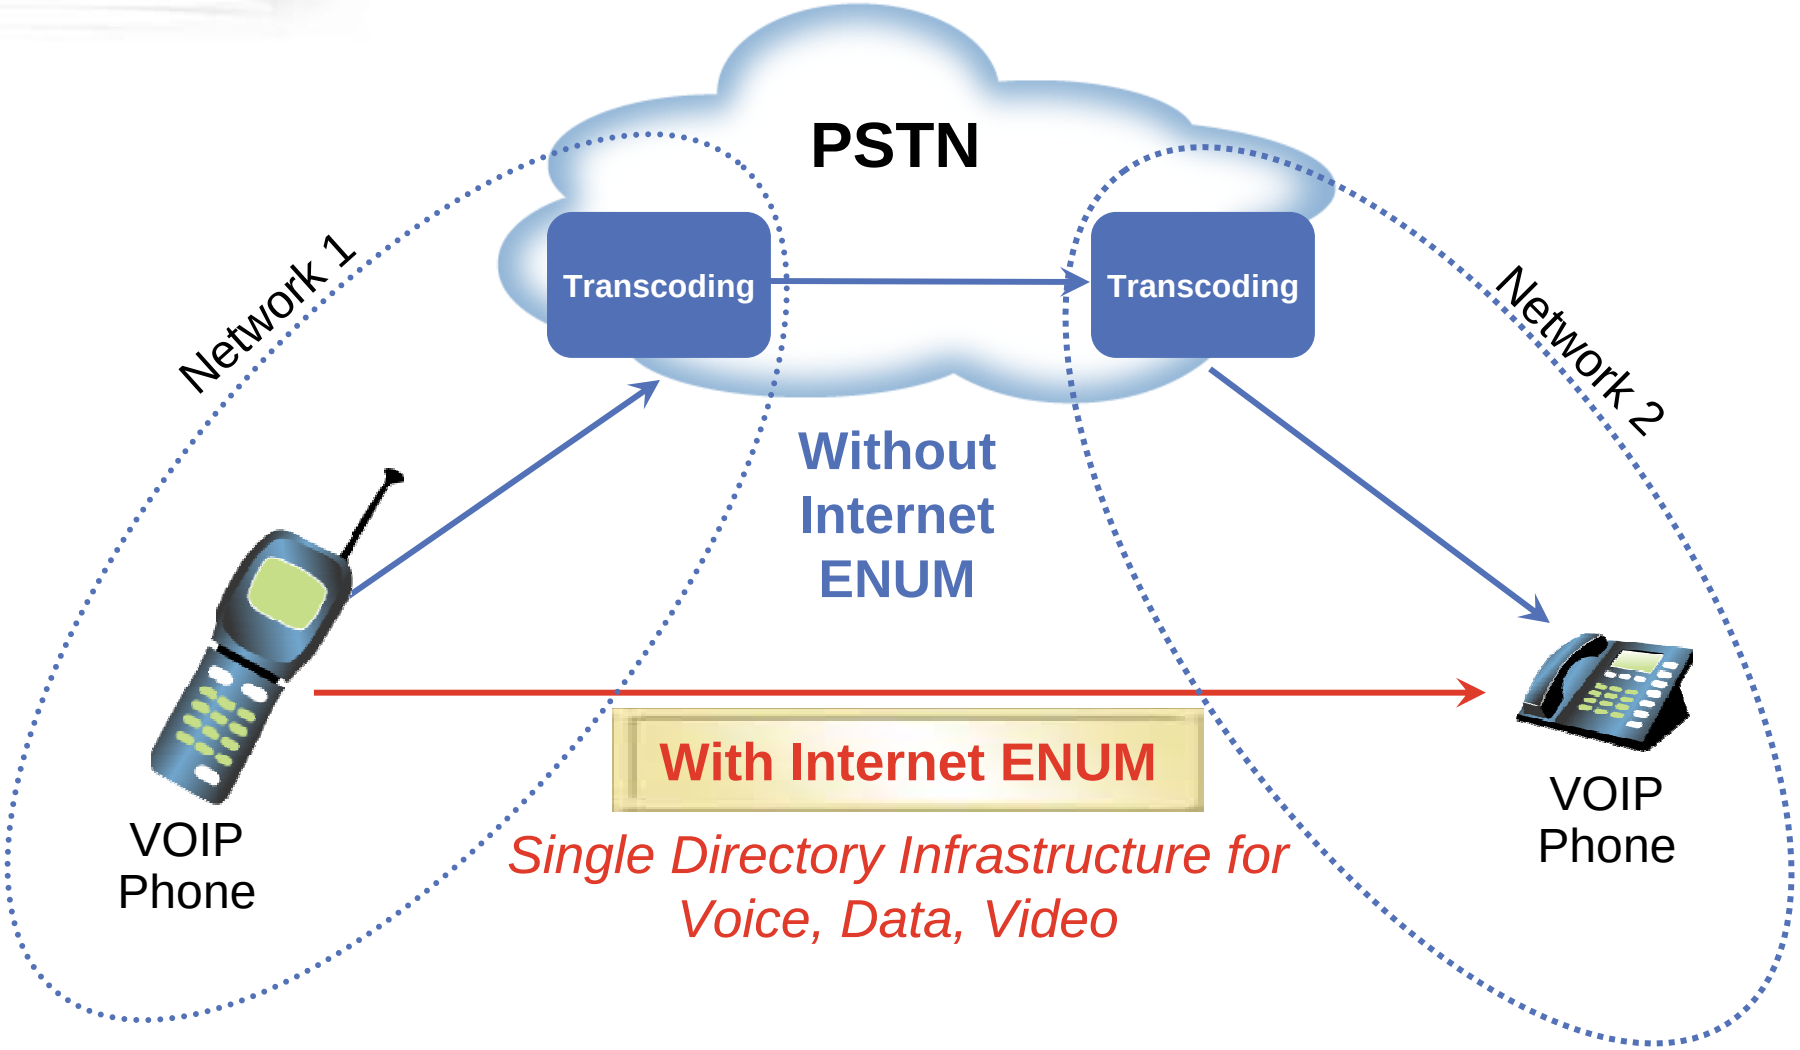

The “Wholesale” level model

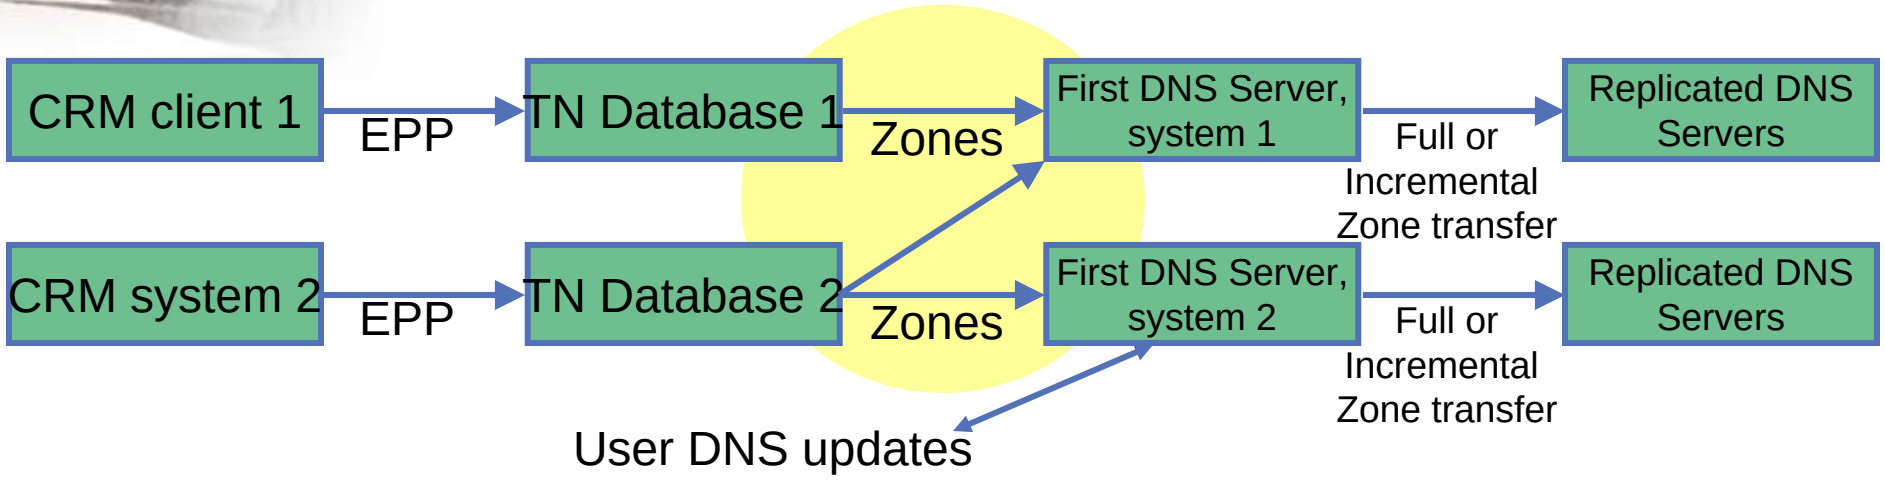

- The first DNS function occurs when the TN databases output zones to a first level DNS.
- Typically can be done in a secure manner using a variety of tunneling techniques

The “Wholesale” level model

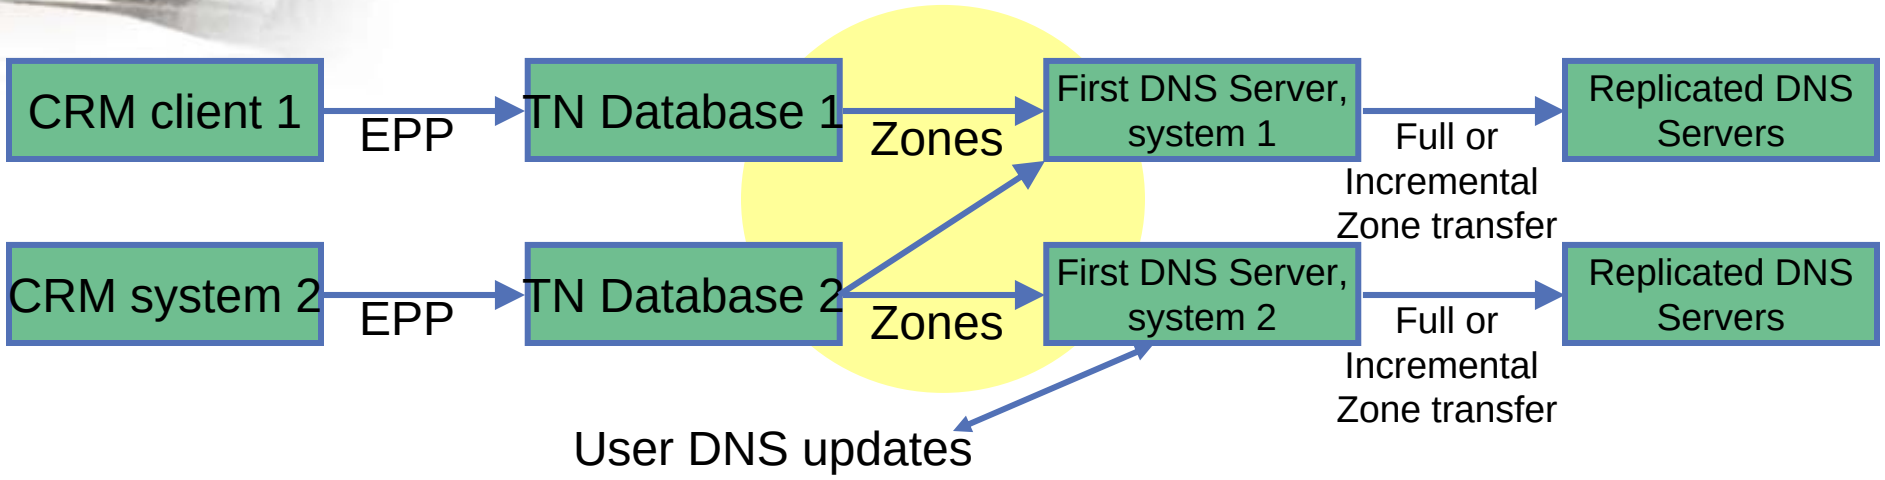

- DNS supports “views” which are basically different zone content for different customers, e.g. an “internal” view *vs* an “external” view.
- Can be used to serve different info to different carriers, subscribers, locations, *etc*.

The “Wholesale” level model

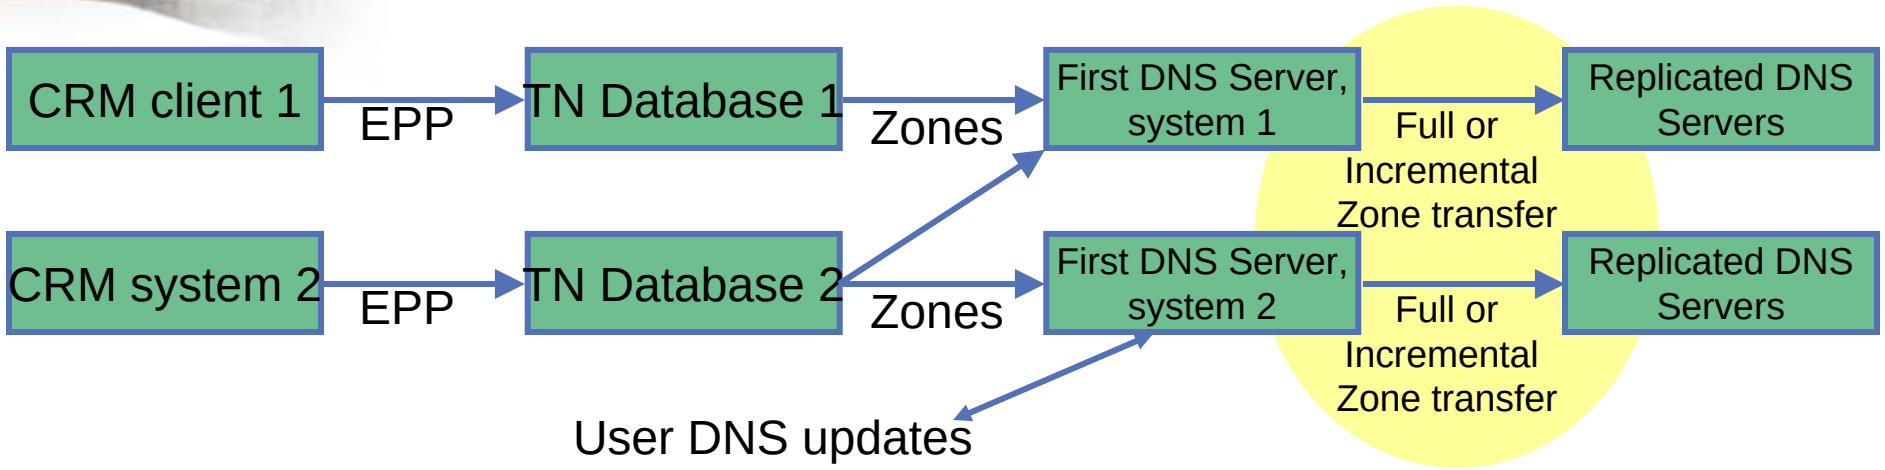

- Replication between DNS servers can be done either on the basis of a full zone transfer, or as incremental changes.
- TSIG to authenticate and prevent replays, but symmetric keying can be problematic.

The “Wholesale” level model

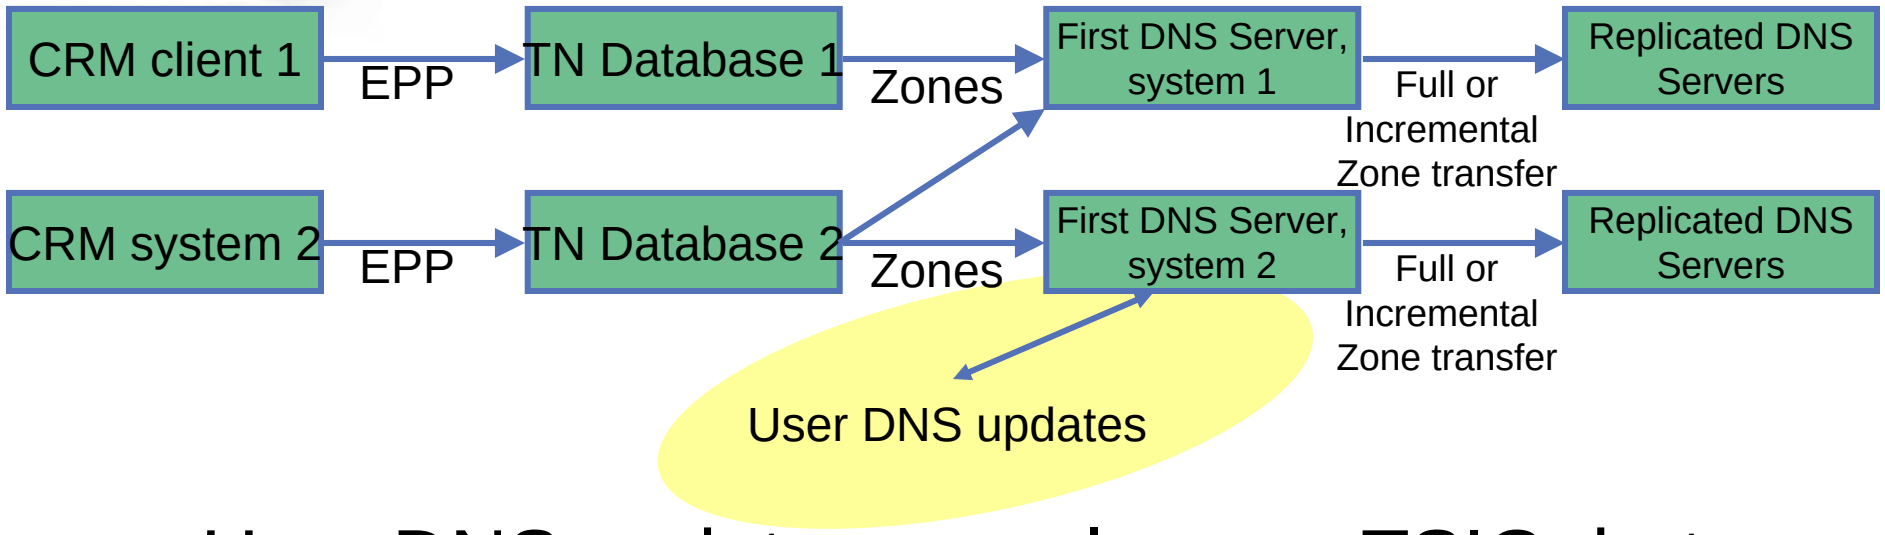

- User DNS updates can also use TSIG, but more of a keying problem.

The “Retail” level model

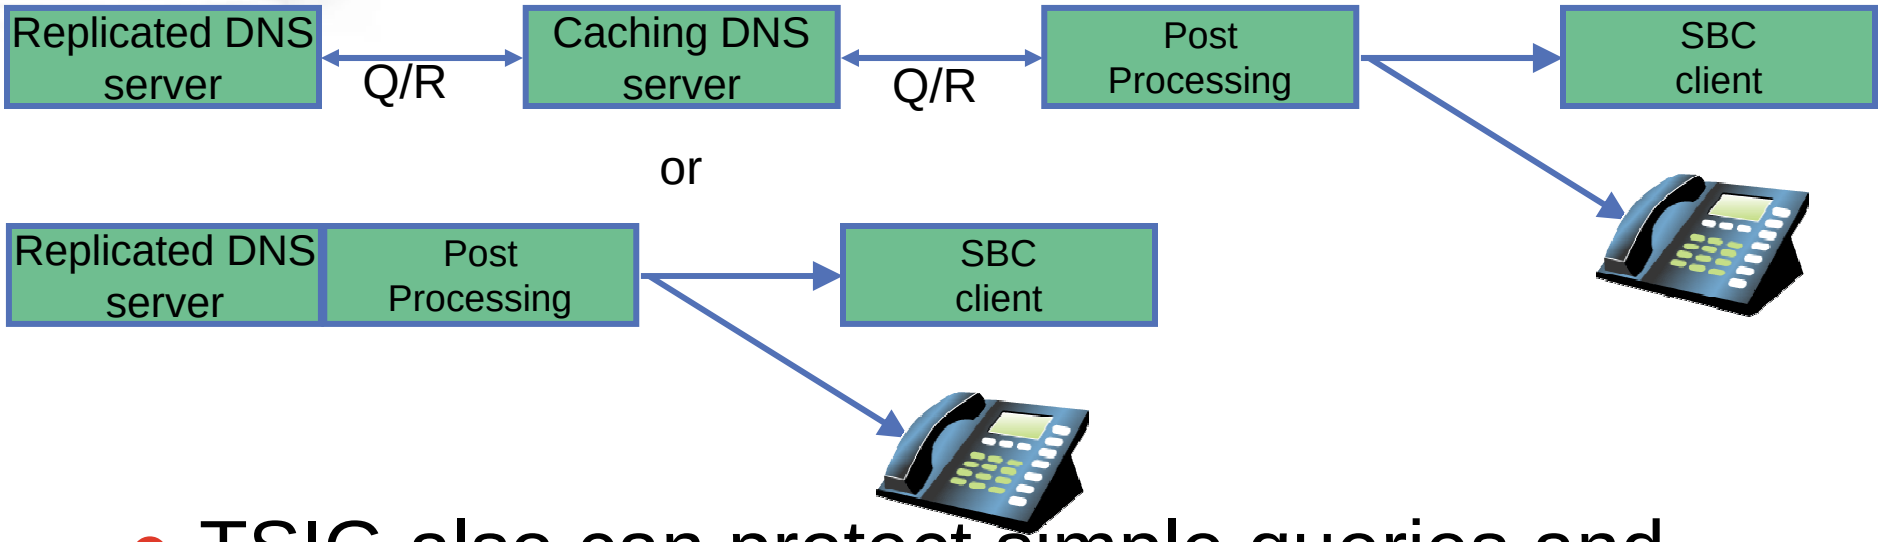

- TSIG also can protect simple queries and responses, although keying is severe problem if clients are numerous.
- May justify switching to DNSSec
- Where should post processing go if needed?

Evolution of ENUM ?

A

B

C

Internet ENUM

1. Carrier ENUM
2. Public Enum
3. Private Enum

ENUM Lookups Resolution

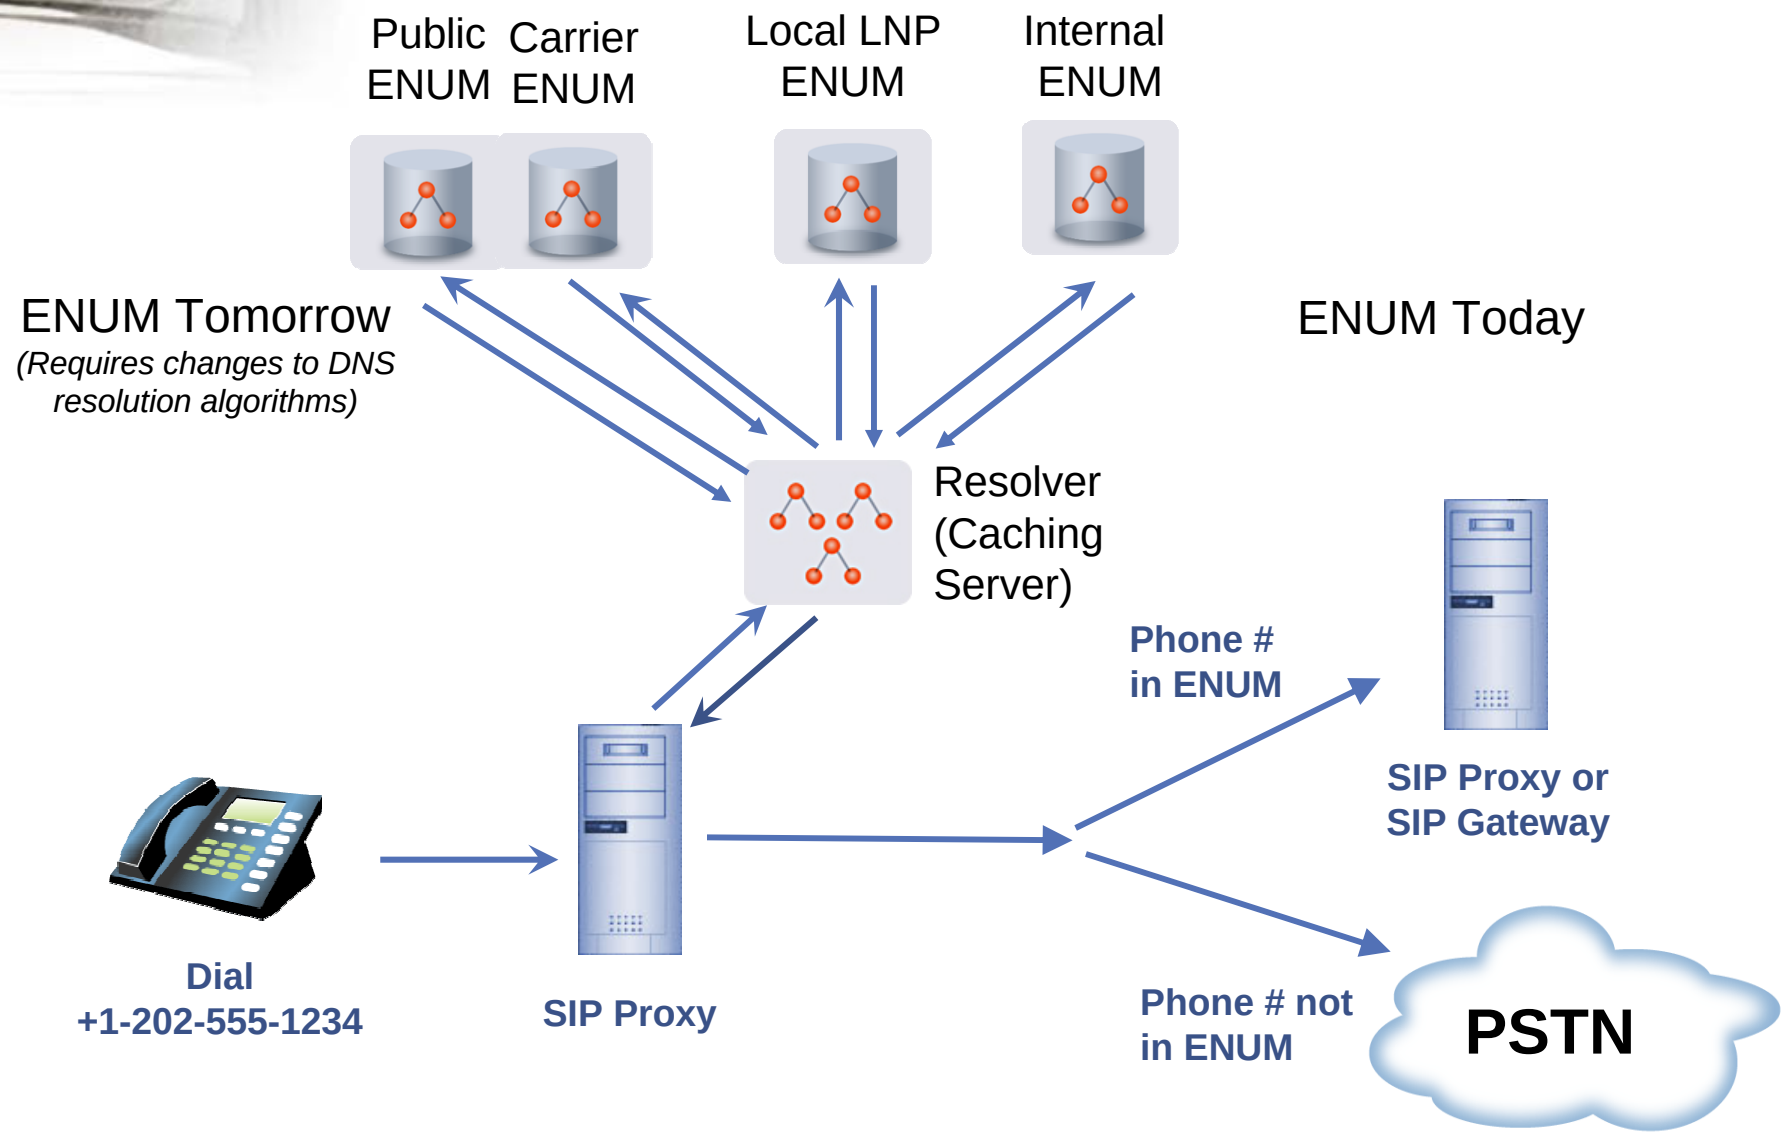

Nominum

Making sense of all this and moving forward..

- What's changed?
- What might we do about it?

Security and Metcalfe

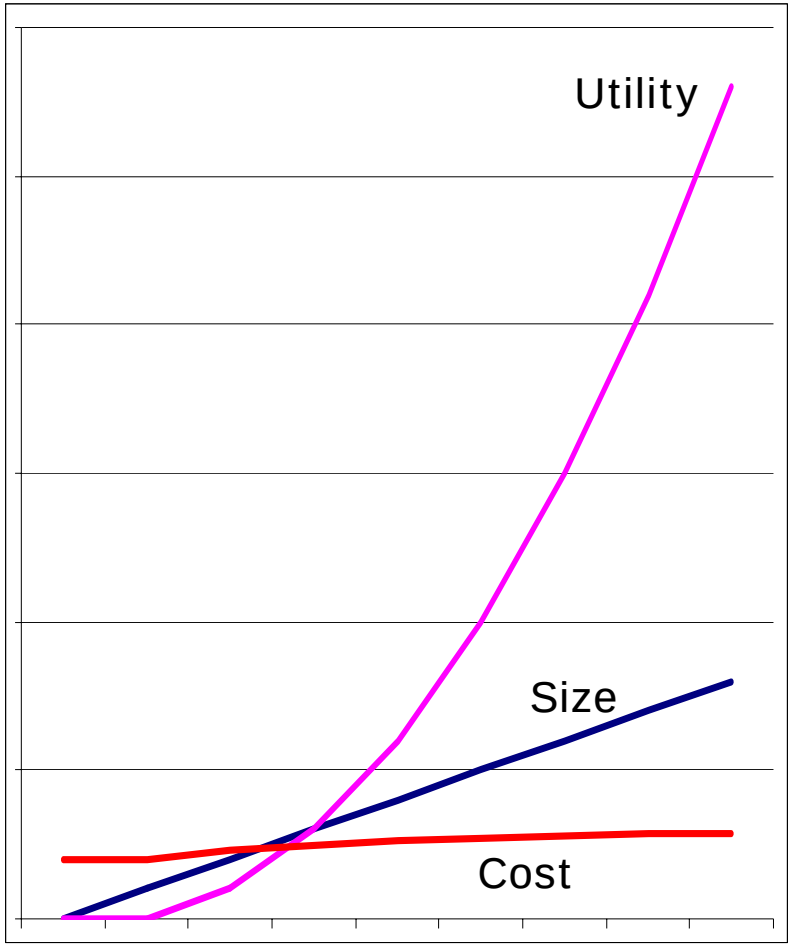

- Metcalfe's law says utility of network is proportional to square of number of members.
- Or utility proportional to number of potential connections.
- Challenge has been to make sure cost grows (much) less than utility and less than size if possible
- With the commercialization of the Internet, law breaks down (e.g. spam, $)

The changing metrics

Nominum.

- Yesterday $U=n^2$

  n = number of parties to the network

- Today $U=g^2-b^2$

  g = number of good guys you can talk to

  b = number of bad guys you can talk to

Implications

Nominum.

- Keeping undesirables out is the new job for the directories
- How does my wireless USB camera talk only to my wireless USB computer and wireless USB hard drive?
- 2005 – How do I store the ENUM for VOIP?
  2008 – How do I disconnect SPIT?

- ***Security needs to become an enabler of new applications, rather than a delayer***

Structuring the data

Nominum.

- We need a way to standardize and deploy new data types
  - DNSSec signed schemas?
- We need to be able to imbed data processing in the data distribution path
  - Data flow in the Internet?
- Multiple name spaces are the rule not the exception
  - No, I haven't met an alternate root I like.
  - Yes, its time to think about what it means.

Making it scale

Nominum.

Semiconductors		DNS	
1947	Transistor	1983	Domain Names, RRs
1958	Integrated Circuit	1993	Dynamic update, DHCP integration
1965	Moore's Law	2005	?

How have DNS concerns changed?

Nominum.

1983

- How do we get researchers to adopt DNS technology?

2005

- How can users get dependable DNS service?
- Managing the data
- Defending against risks
- Reducing costs
- Designing new functions

How have DNS systems changed?

Nominum.

1983

- Where do I get the code for DNS to compile and install?

2005

- I need a system that can do moves, adds, and changes without restarting
- I need to manage 100 servers as a unit
- I need to manage 20 system administrators
- I need integrated DNS and DHCP

Conclusions

- Today, DNS holds roughly a billion names; will double every year for at least next 5 years
- Old management practices will not work as DNS disappears into the infrastructure and becomes mission-critical for all Internet users, even those who don’t know they are using it (e.g., IP telephony)
- Integration between directory levels is the next opportunity we face

Technical & Political Needs

Nominum.

- Ownership/control of:
  - 1,000,000,000,000 identifiers
  - By 10,000,000,000 owners
  - 100,000,000,000 transactions/day
- New security models
  - Faster than 10 years/standard (DNSSec)
  - Easier to use than X.509
- Cooperation model for
  - Standards bodies
  - Companies
  - Governments
  - Lawyers

A word of caution

Nominum.

- “Paul, you are putting too much function into the DNS, these ideas will be too difficult to implement and control, and there are better tools coming that will properly handle this problem.”

  Internet Architecture Board (IAB) 1982 (use x.500 instead)

- “The DNS doesn’t need new features and data primitives”

  National Science Foundation (NSF) 1988 (DNS growth is over)

- “Are you crazy?”

  Coworkers 1978-present

Nominum

Q&A