Nominum

What should we learn from 25 years of the Internet:

A DNS case study

Paul V Mockapetris
pvm@Nominum.com

What's the Point?

- Four reasons:

  1. Just record the history
  2. Have a good party and talk about how we were geniuses (cue Bruce Springsteen, "Glory Days")
  3. Learn principles for the future
  4. Make fun of the "clean slate"

- All have issues, let's stick to #3.

Source: Nominum

Why a DNS case study?

Nominum.

- It's my area of expertise
- Experts for other protocols often
  - Disagree about age of Internet
    - 40th birthday if you include ARPAnet
    - ~120th birthday if you think about Hertz/Marconi
    - Internet era may have ended with HTTP / web2.0
  - Get into credit food fights
  - Like the history or party idea better
- DNS has evolved by many hands
- DNS touches most of the rest anyway, so may be a good place to start

Source: Nominum

EarlyTimeline

Nominum.

- Nov 1983 – RFCs 882, 883
- 1985/1986 machines without host tables
- Jan 1986 – MX style mail routing
- Nov 1987 – RFC 1034, 1035
- Aug 1988 – “Development of the Domain Name System”, Sigcomm 88
  - AKA DoDNS

Source: Nominum

Then - 1983

Nominum.

- Previously, the IP/TCP transition meant that every system could be rethought
  - For example, FTP->FTP & separate email
- Many, many things to rethink
  - Important folks rethought what were seen as important issues, for instance
    - Routing
    - Card images in TCP
    - Design of “The Directory”
  - Less important folks did things like
    - DNS
    - Datagrams
  - Some things seemed simple
    - Managing & allocating names

Source: Nominum

Intent of DNS protocol design 1983

Nominum.

- Provide a design that was just lightweight enough to take off
- Provide a design that had orthogonal features that could be combined to produce lots of possibilities
- More of a recipe than an invention
- Core values
  - Simple wins
  - Reliable through replication
  - Must be inherently fast
  - Distribution of authority and control

Source: Nominum

Later Additions

- Dynamic Update
- DNSSEC
- TSIG

- Many false starts

Source: Nominum

Important other issues

- DNS -> DN$
  - Marketing
  - Trademarks
  - ICANN
  - Etc
- Simple numbers
  - e.g. DoDNS
    - Root does 1 query/sec
    - Good queries take 100 msec

Source: Nominum

What would Buffett Say

- “You can get in way more trouble with a good idea than a bad idea”
  - Ben Graham
- …because you forget that the good idea has limits
  - Warren Buffett
- “Life is like a snowball. The important thing is finding really wet snow and a really long hill.”
  - Warren Buffett

Source: Nominum

Scalability & Extensibility

Source: Nominum

Scalability

Nominum.

- Should MTU be in bits or time?
- For example:
  - 1990 ATM cell @ OC-3 = ~350 ns
  - 2008 Ether @ 10G = ~150 ns

- DNSSEC fundamentals suffering from inability to carry large signatures easily
- DNS-only expansion isn't the answer
- TCP isn't the answer

Source: Nominum

It’s the API, stupid

Nominum.

- Ethernet API has survived:
  - Change from passive multidrop to point to point
  - Copper to wireless to optical
  - Frame and address idea survived

- DNS API
  - RRs OK for a decade
  - Needs update now
    - Based on simple concepts
      - Set theory
      - hierarchy
    - Self defining new types

Source: Nominum

Standardizing can be tough

Nominum.

- The affair "_"
  - ISC outlaws the "_"
  - Microsoft makes it required
- The IETF
  - "Don't overload the DNS"
  - We'll tell you what you can use in your DNS
  - Can't be used for data needing security, except that it is.

Source: Nominum

Lessons

Nominum.

- We need a new, larger, datagram.
- We should rethink the conceptual model and clean it up, and extend it, in the process. Define the API.
- We don’t expect the file system to approve content; we shouldn’t do so in the DNS either.

Source: Nominum

Security

Source: Nominum

Today's Discussion

- The DNS is exposed
- Subscribers are under attack
- The "Gap" between future technologies and today
- The Trusted Internet Experience – The TRUE Architecture

Source: Nominum

Rapidly Growing Problem

How to determine the difference between safe and harmful requests in real time?

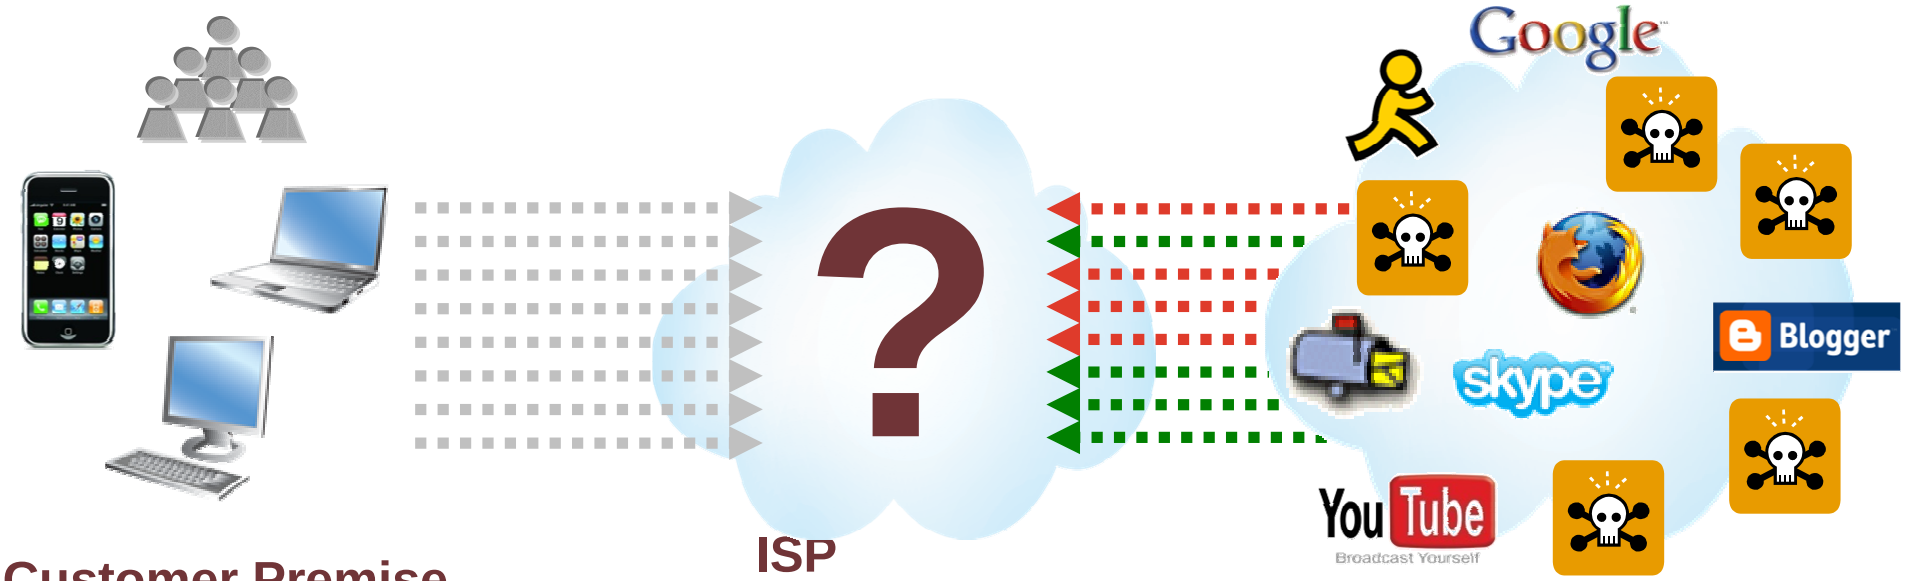

Customer Premise

ISP Broadband Network

Internet
150 *million* malicious hosts and growing…

How Can the Service Provider Help?

Source: Nominum

DNS History (past and future)

Nominum.

- 1983 — DNS starts
  - » Intentional omissions include security, dynamic update, etc, etc
- 1986 — DNS liftoff
- 1989 — Cache Poisoning observed
  - » “Don’t cache data just because somebody sends it to you”
- 1989-2008 — Various cache poisoning attacks
  - » Multiplexing technology adapted for security
  - » Other defenses deployed
- 1993 — DNSSEC starts
- ~2000 — Search makes “the missing directory” irrelevant
- 2008 — Kaminsky fast poisoning attack

…

- 201X — Majority of DNS secured with digital signatures

Source: Nominum

Statistical Attacks

Nominum.

Password

- Type login command
- Guess password
- Repeat till success

- Odds/guess:
  - Using “a-z, A-Z, 0-9” ~6 bits/character
  - 2 chars 1 in 3,884
  - 3 chars 1 in 238,328
  - 4 chars 1 in 14,776,336

Kaminsky

- Send query so server listening for answer
- Send guesses while target DNS waits for real answer
- Repeat till success

- Odds/guess
  - 16 bit ID, 1 in 65536

Source: Nominum

How do Computers Navigate the Network?

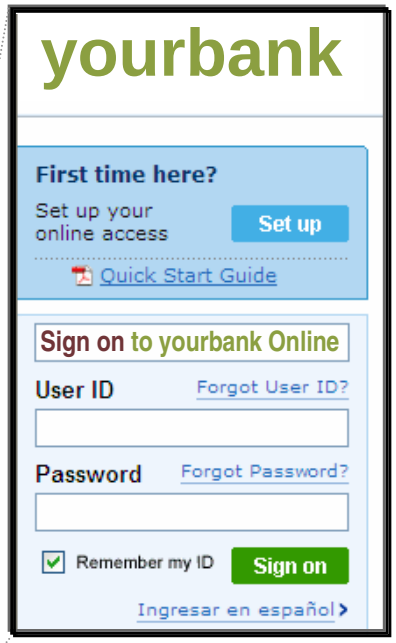

To get to www.yourbank.com, the computer asks its local name server for directions.
For a company, it's the company's DNS server.
For a broadband user, it's the ISP's.

ISP or Enterprise Caching DNS

How do I reach www.yourbank.com ?

www.yourbank.com Is at 1.2.3.4

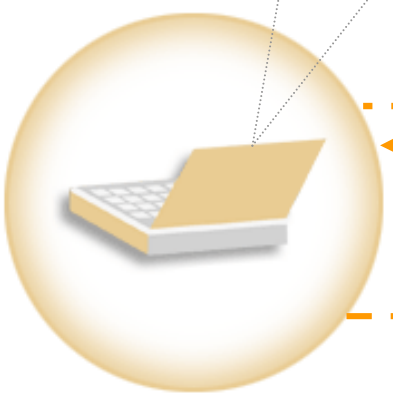

User browsing

TCP Me to: 1.2.3.4

Internet

Your Bank

www.yourbank.com
1.2.3.4

Source: Nominum

Cache Poisoning Attack

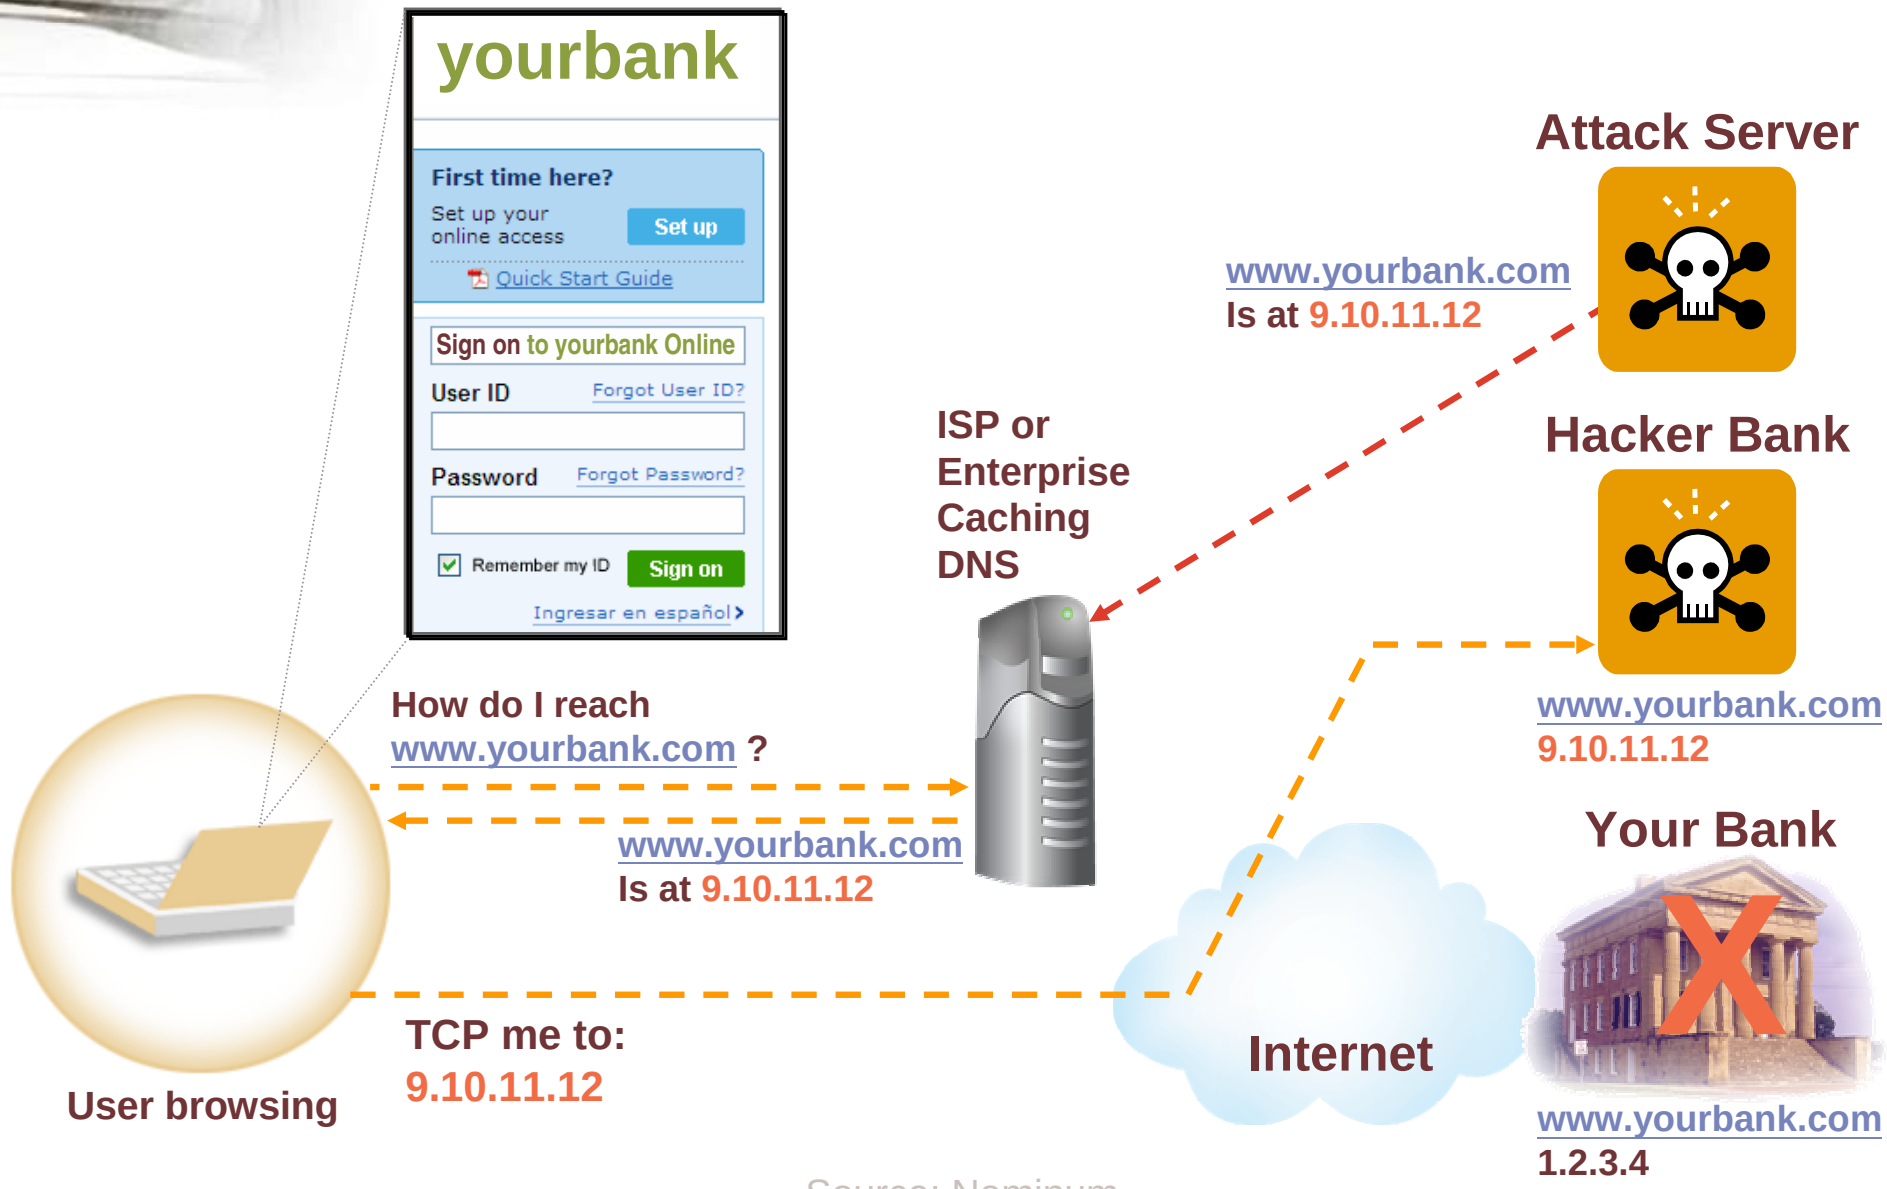

Source: Nominum

The Eye of the Hurricane

Nominum.

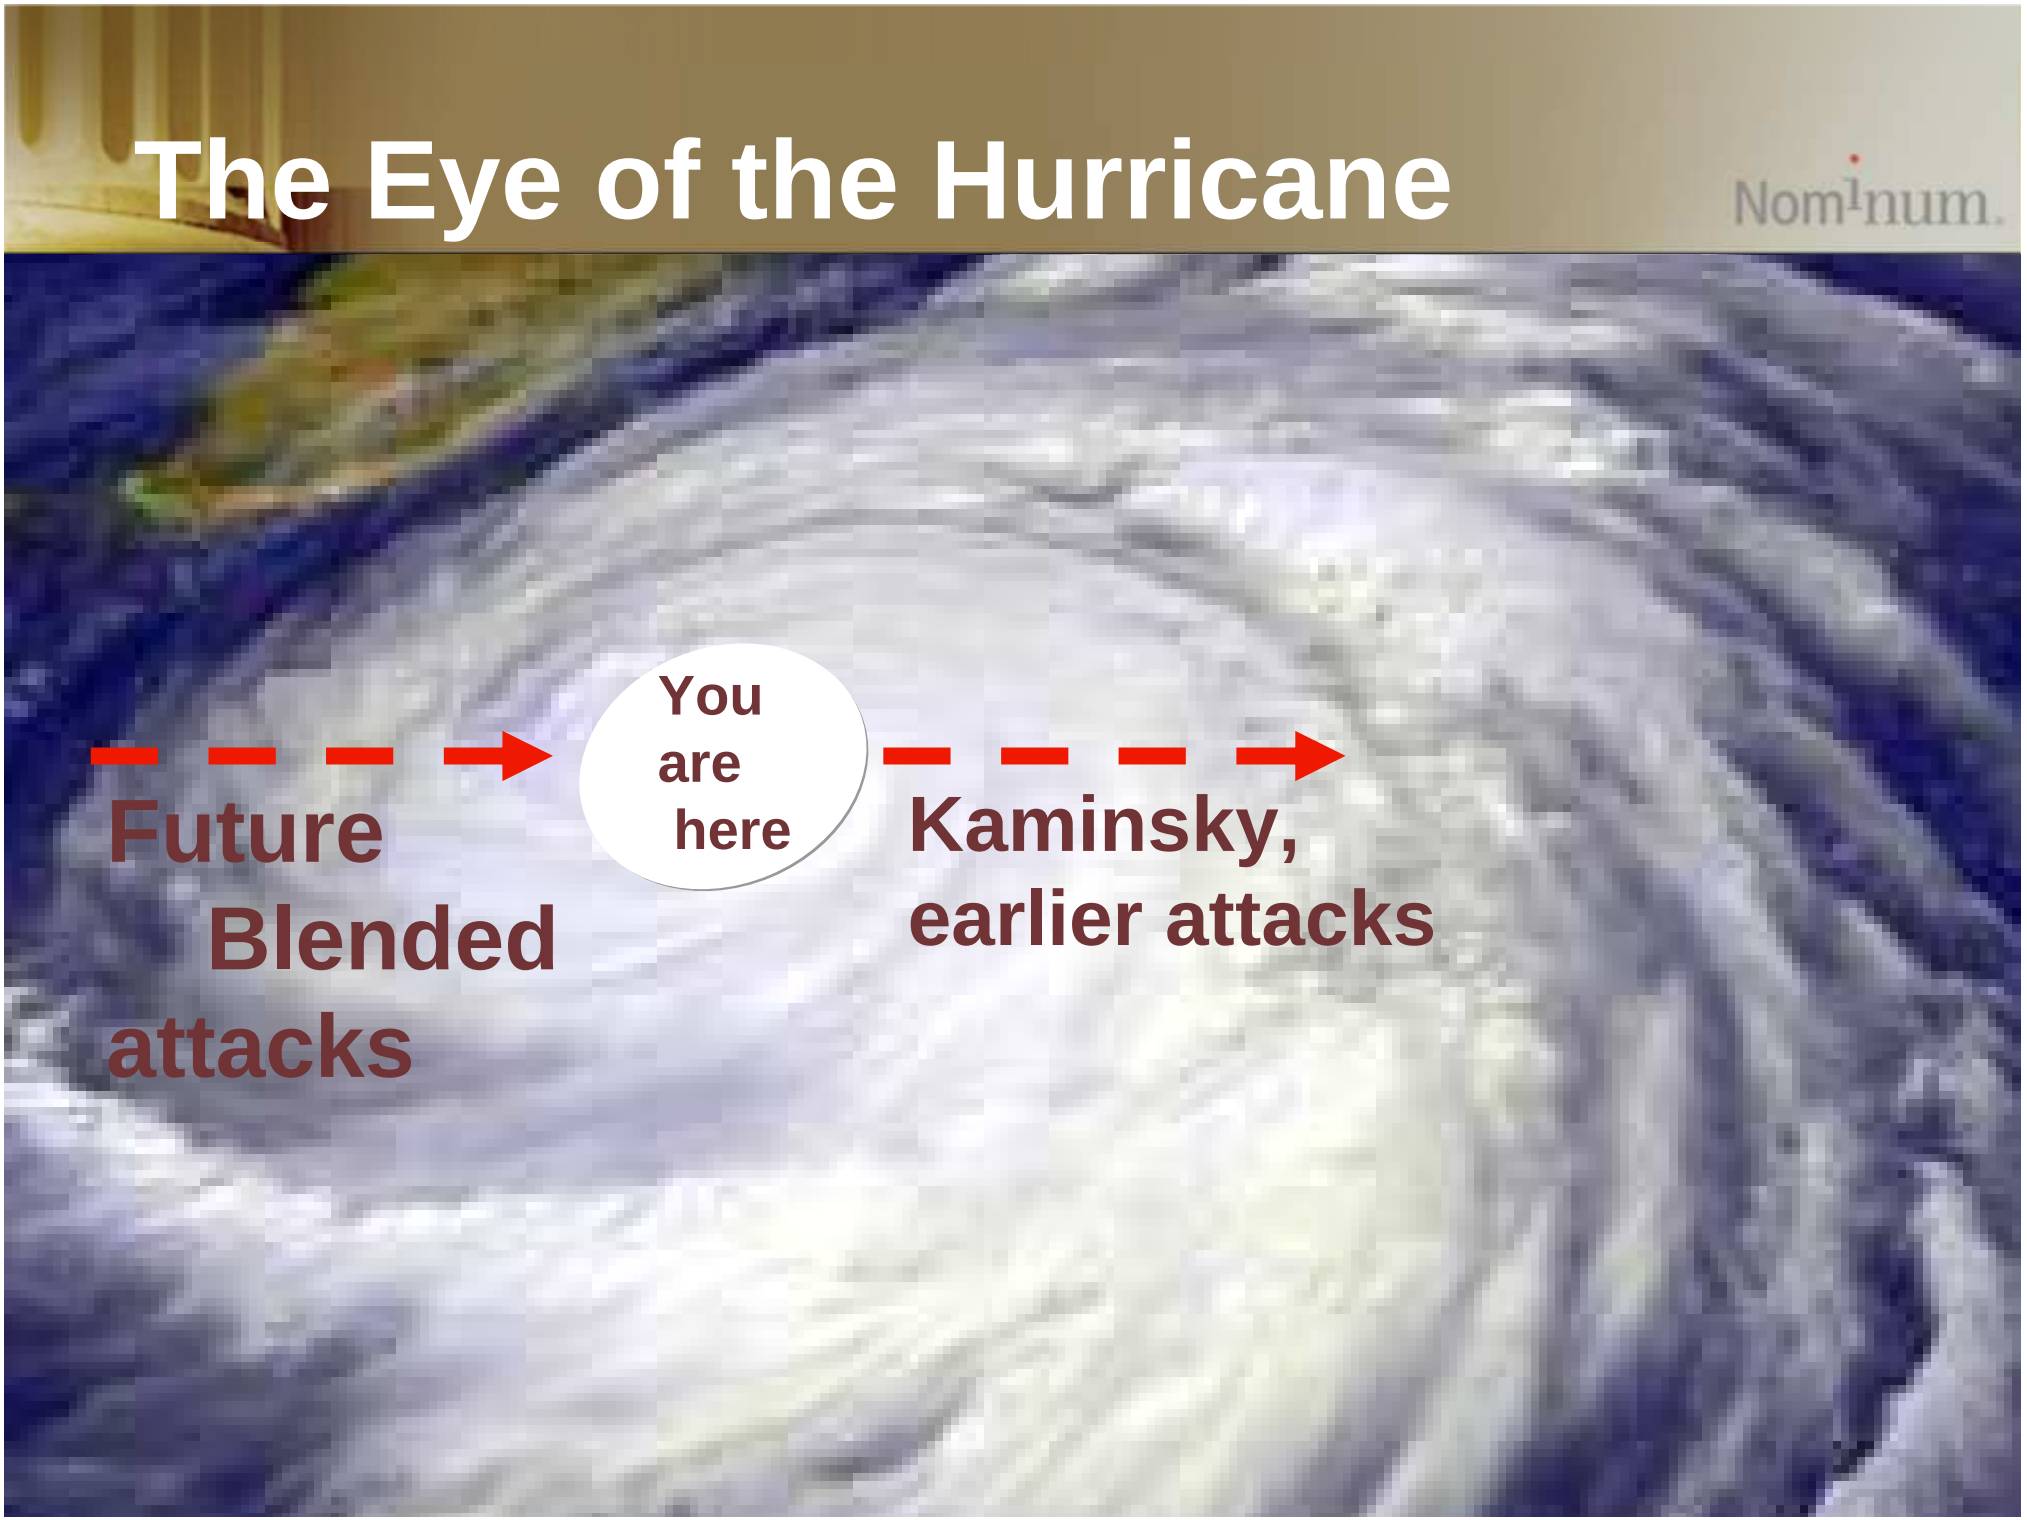

Mail Attack

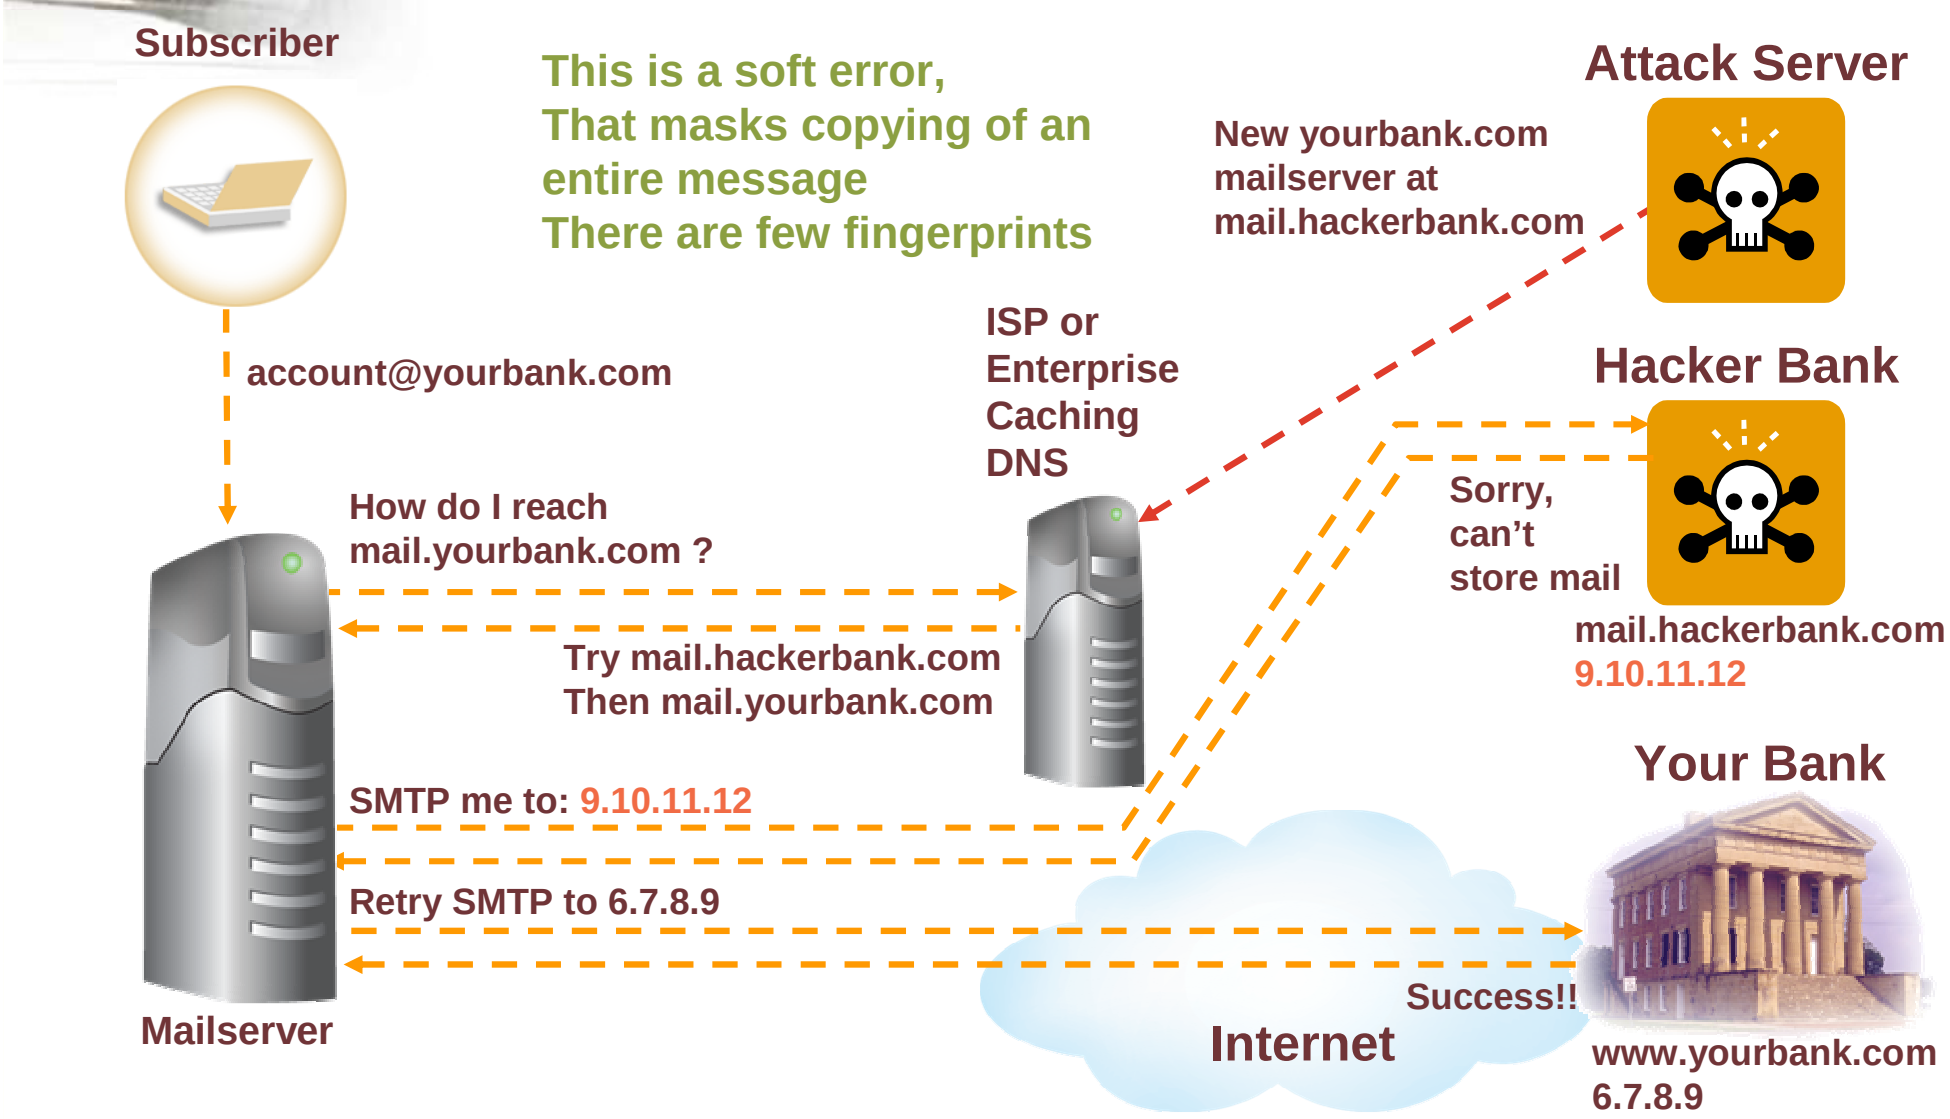

Source: Nominum

Two Messages

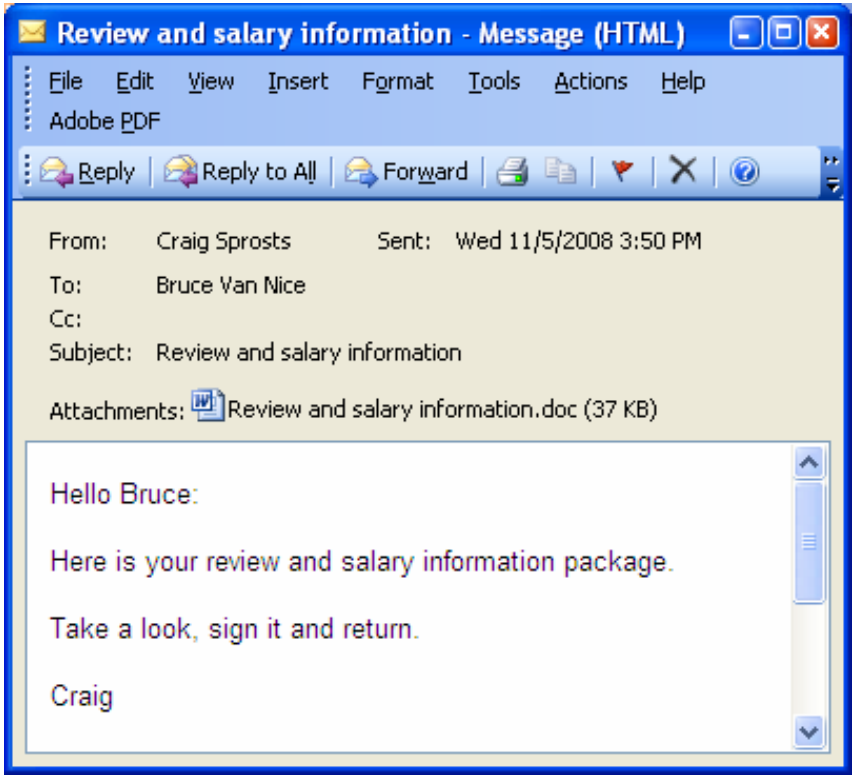

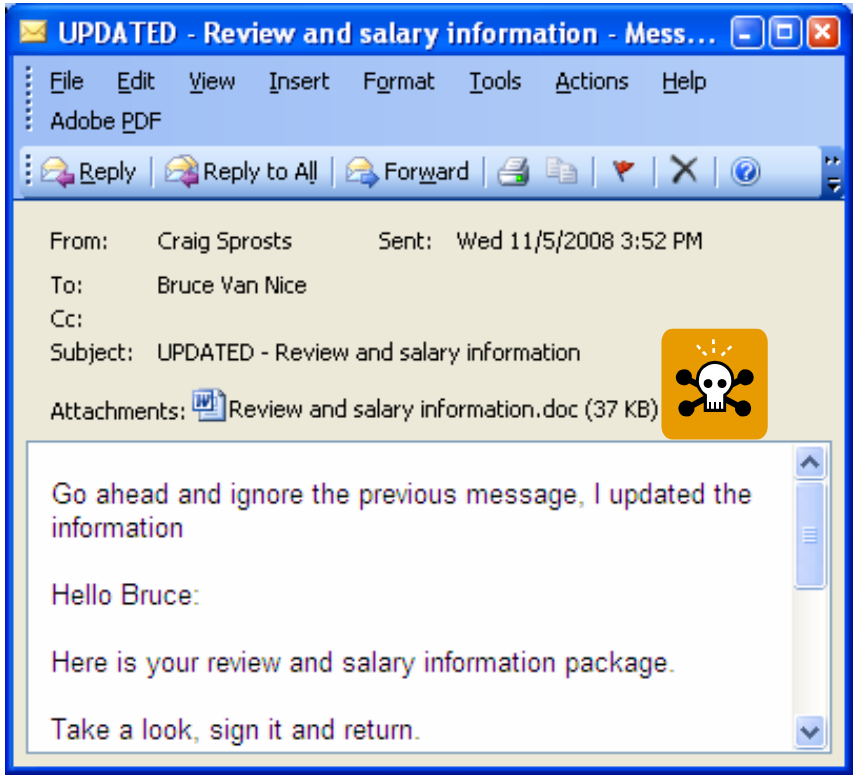

Source: Nominum

IETF USPR response: Augment IDs with ports

Nominum.

- Old ID-only: 1 chance in 65,536
- ID + port: 1 chance in 4,294,967,296

- But
  - Doesn’t work with load balancers
    - Back to 1 chance in 65,536
  - Slows servers

Source: Nominum

Hacker response to USPR: increase attack rate

Nominum.

- One experiment showed that an attack over a gigabit network defeated USPR in 10 hours using 2 machines.
- That attack was unlucky; attack works faster on average
- Coordinated attacks via botnets
- Attack .COM or .JP and own all names below

USPR isn't enough.

Source: Nominum

A Changing World

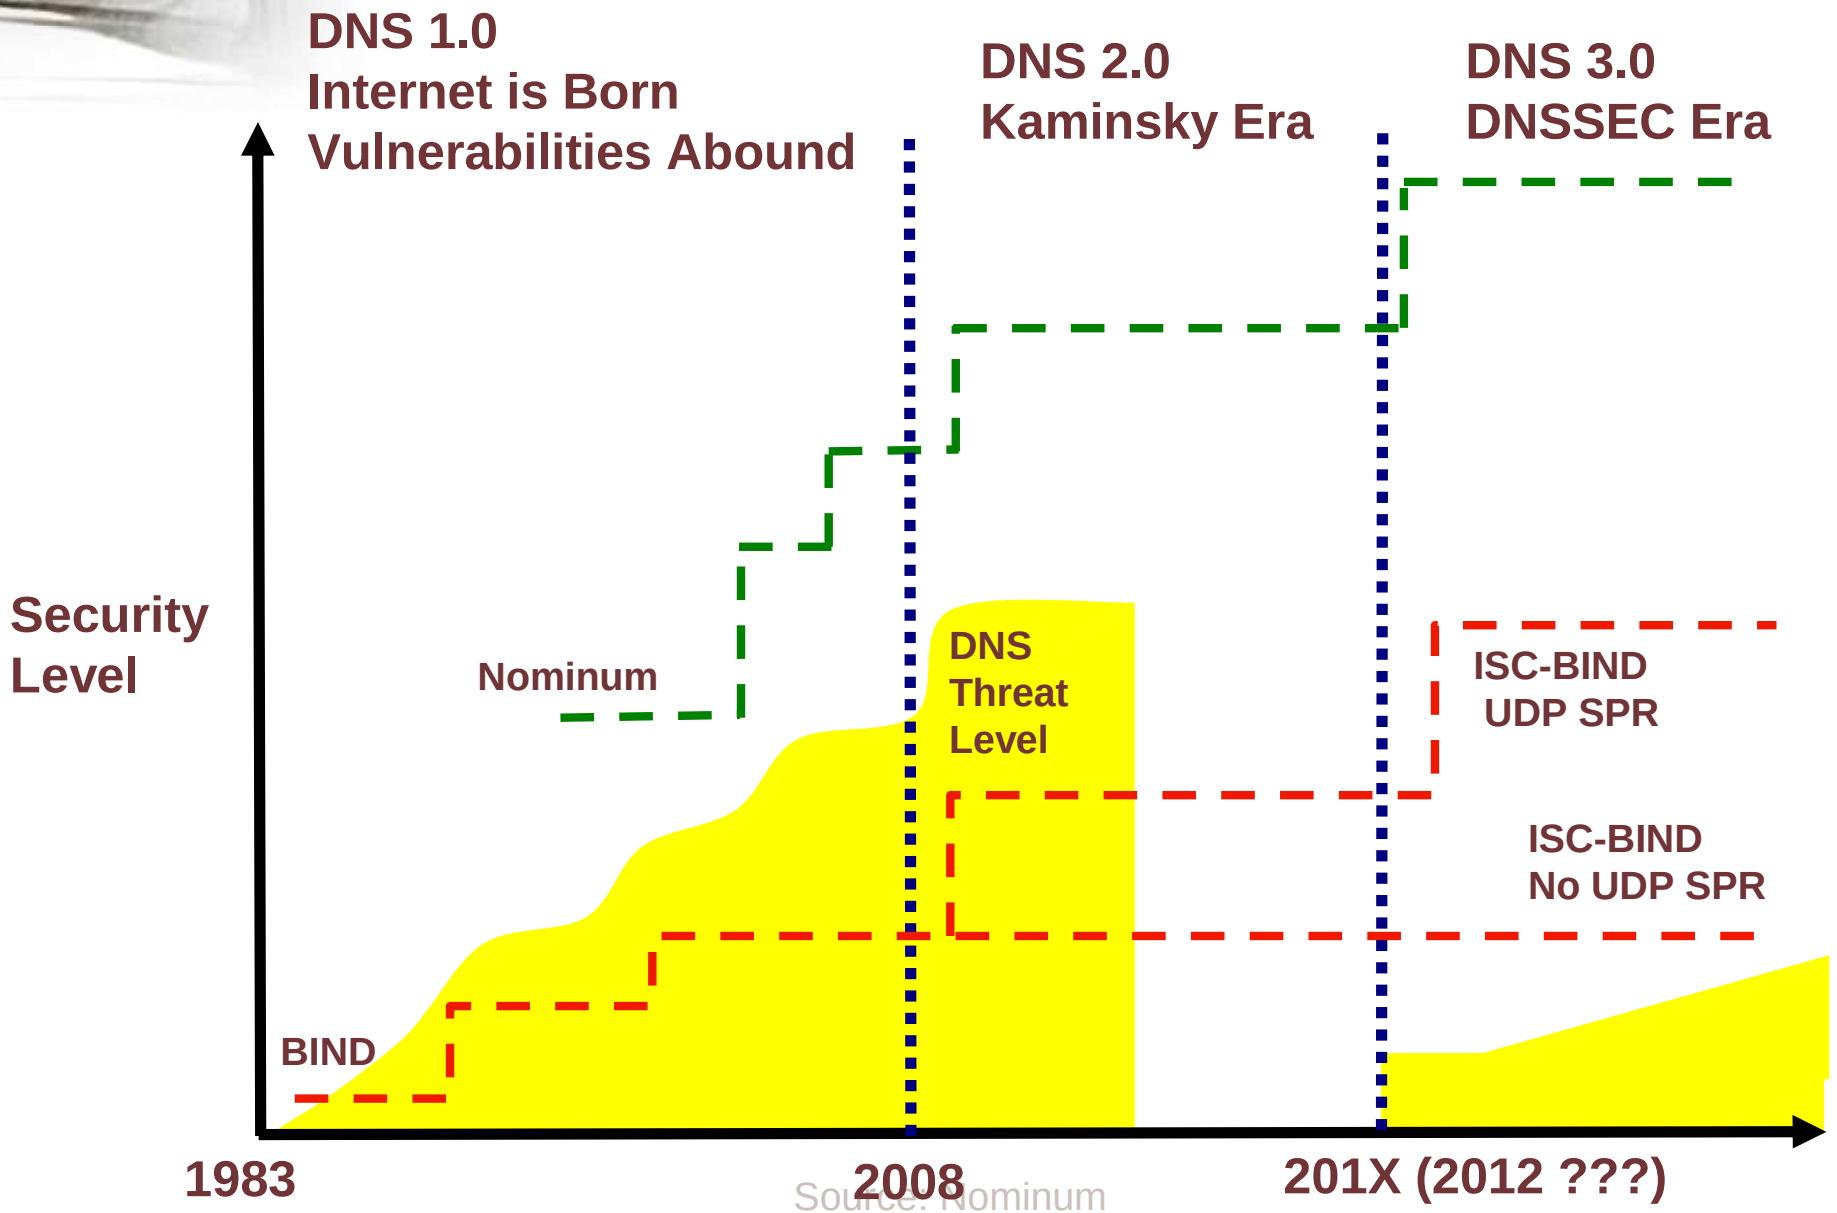

Source: Nominum

How safe is the Internet?

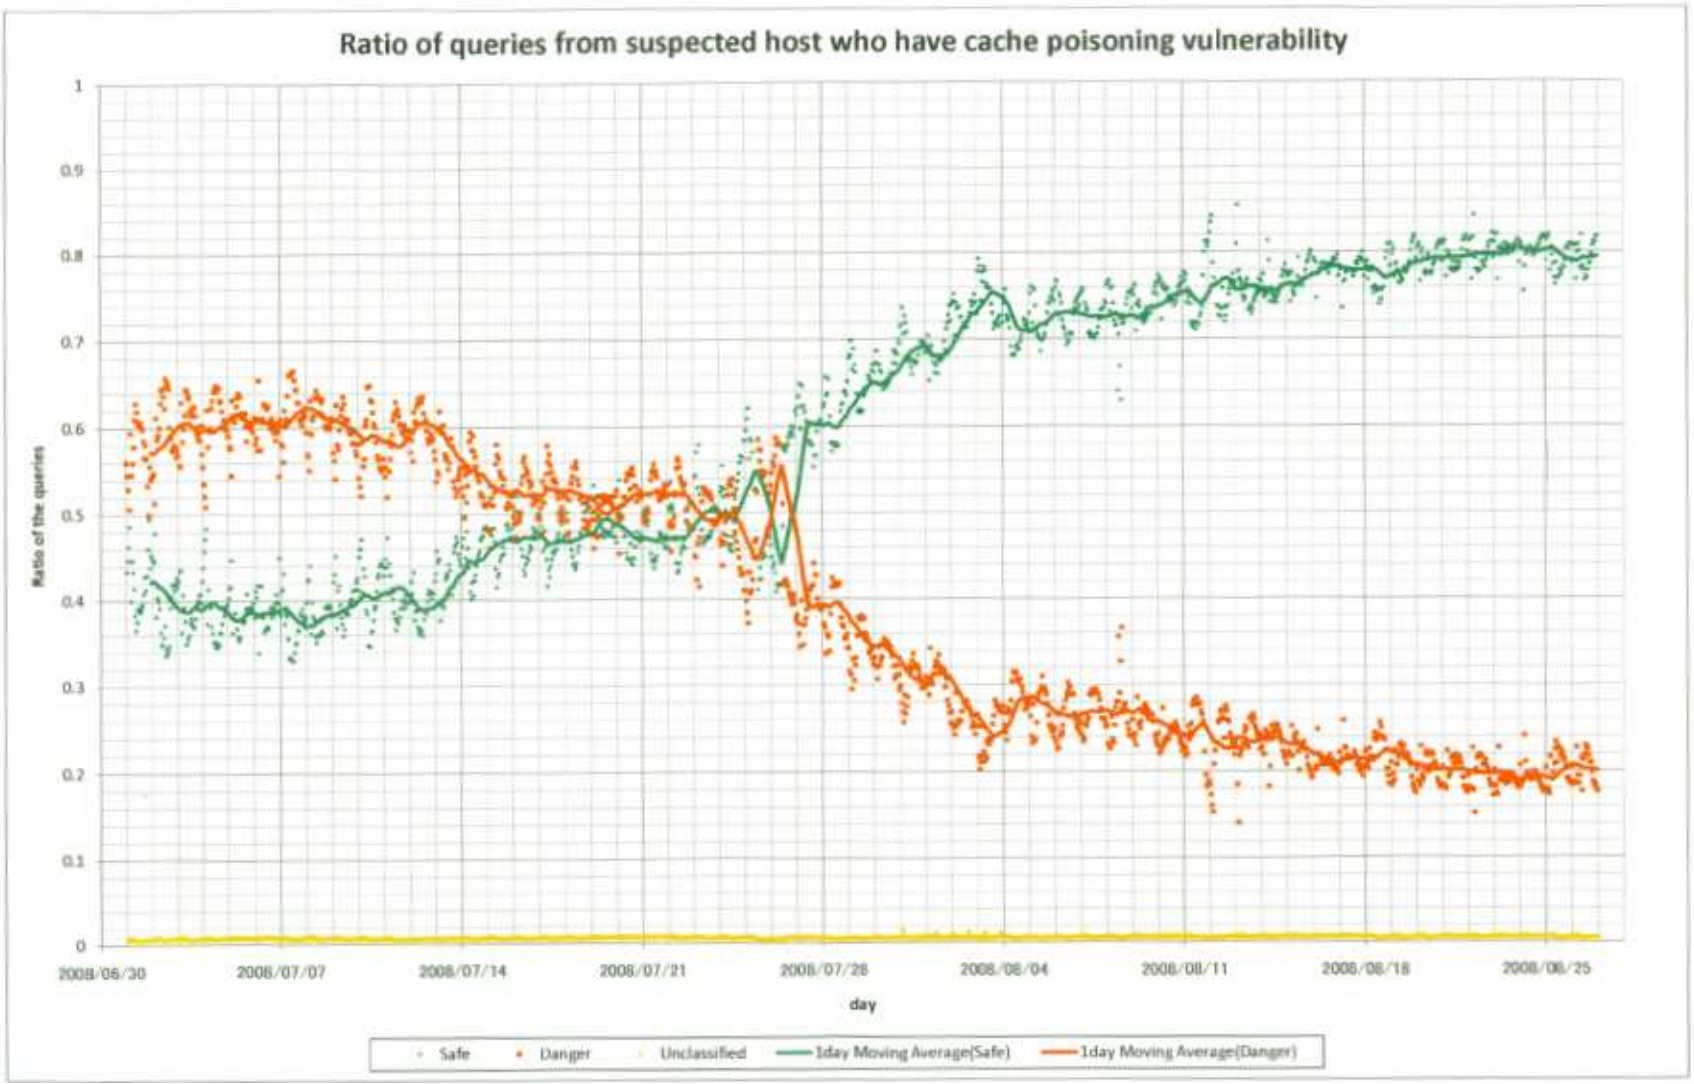

Source: Nominum

Lessons

- We need strategies to improve DNS security
  - Near term which can be deployed now
  - Long term enhancements (DNSSEC?)
- Speed kills (faster nets are more vulnerable)
  - Enterprise at risk from infected machines
  - Secure your DNS with a 10Mbit connection?
- DNS servers embedded in appliances, etc may not be easily upgradable

Source: Nominum

New Applications

Source: Nominum

DNS use is growing exponentially

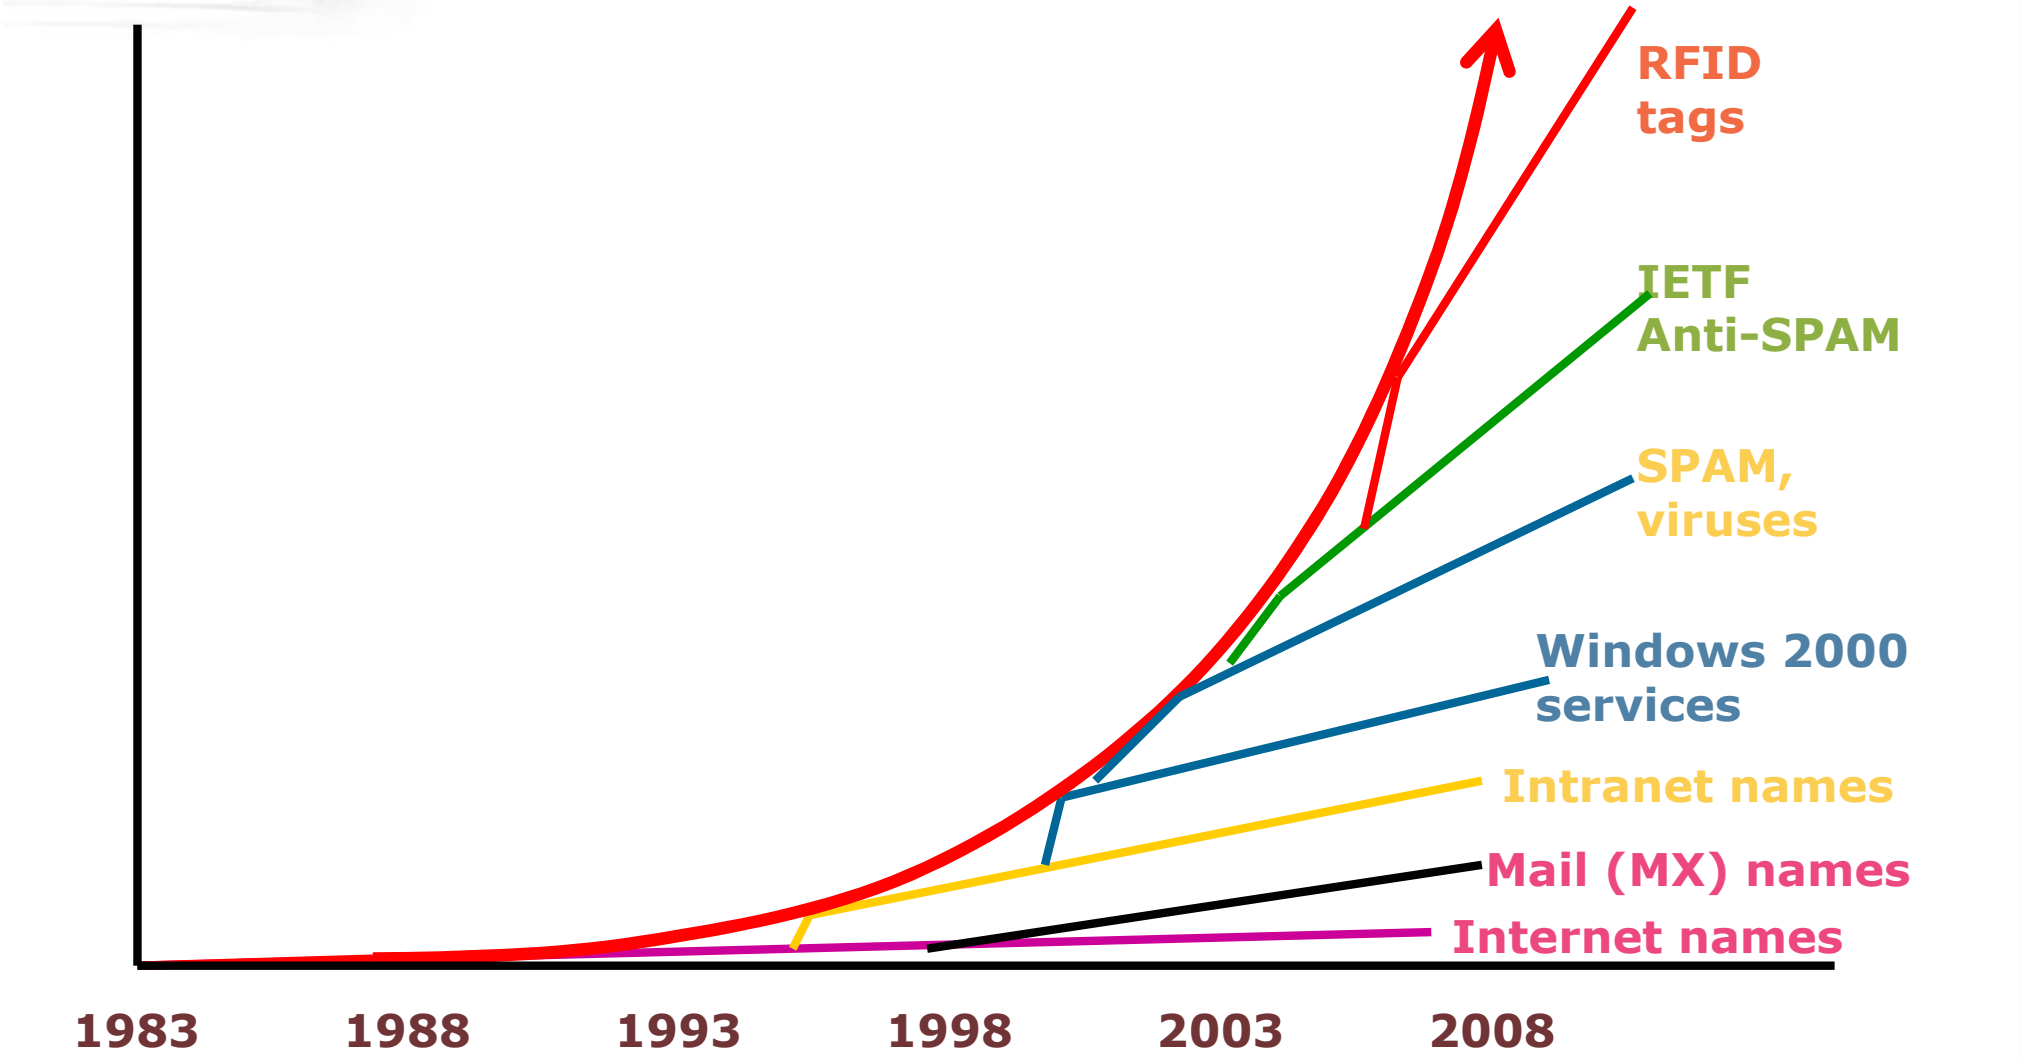

Source: Nominum

RFID's Origins

Nominum.

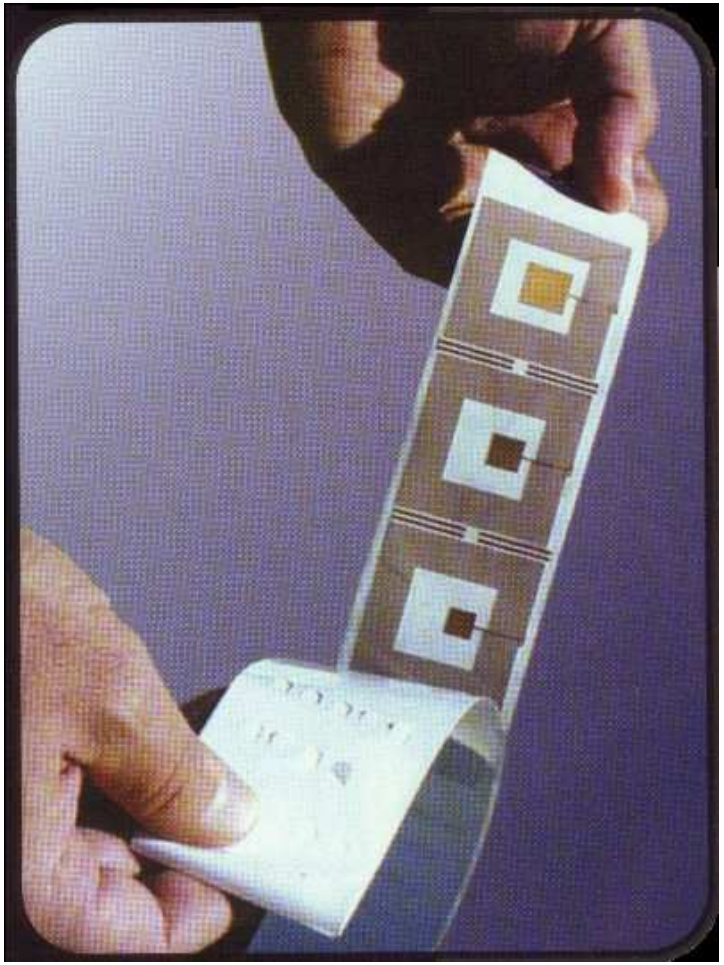

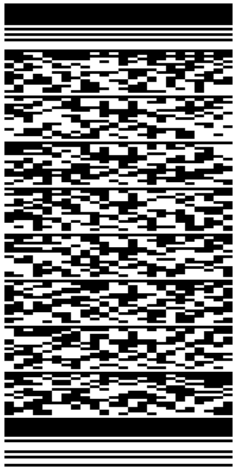

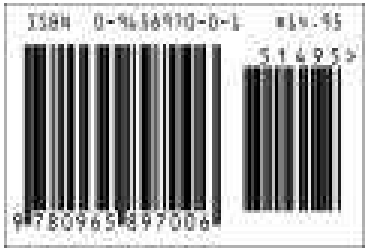

Source: Nominum

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

Source: Nominum

History

- 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.

Source: Nominum

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

- 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

Source: Nominum

ENUM

- Idea: Let’s have a standard that uses the DNS to route phone calls (and other new media)
- Problem: ENUM uses only destination number to route, real world uses more fields than that.
- Problem: Equipment manufacturers want intelligence, i.e. value, in their product.
- Problem: Legacy data owners really don’t want to change ownership scheme.
- Problem: Security is used as issue.

Source: Nominum

Lessons

- Displacing a legacy model is more than technology
- Catalysts for new developments
  - Security
  - Self defining data types
- The next new applications
  - Threat feed and configuration data to all enforcement devices, e.g. firewalls, mail servers, …

Source: Nominum

Final Thoughts

Source: Nominum

Facts to face

- ICANN isn't "too political"
  - ICANN *is* politics
  - Apply the usual political safeguards, checks, and balances
- We shouldn't worry about overloading the DNS
  - We should worry about perfect standards that take decades
  - More evolution, less intelligent design
  - Even if extinction is the next step

Source: Nominum

The future

- Continuing struggle between two factors
  - "The Internet changes everything!"
  - "For every action, there is an equal and opposite reaction."
- The real world pushes back, excesses provoke reform, …

Source: Nominum

Replacing/Extending DNS

- Process:
  I. Assemble set of key problems
  II. Generalize
  III. Prune
  IV. Postulate a solution
  V. Test

Source: Nominum

Worthy Problems

Nominum.

- IPv4 address space exhaustion and LISP
  - Layer of indirection for IPv4 addresses
  - Double size of tracked address space
  - Merge route flap and quasi-static multi-homed assignmants
- AS numbers going to 4 bytes
  - Hard to type
  - Can we distribute mnemonics

Source: Nominum

Nominum®

Q & A