17 March 2015

Disruption and the DNS:

(Caring for Senior Citizen Protocols)

Paul Mockapetris

Paul.Mockapetris@ICANN.ORG
Paul Mockapetris <Paul-Vincent.Mockapetris@npa.lip6.fr>

Thesis: We need more Disruption!

Today’s Agenda

- Philosophy
- How did we first Disrupt?
- Planned Disruption ends
- Today
- Future Directions

Philosophy

All Distributed Systems have 3 Parts Today:

Hardware

Software

Configuration

Why is it always so messy?

- Because we always build systems that challenge:
  - the competition
  - the complexity we can handle
- So priority one is reducing complexity

Simple ideas win, but may take time

Look what Zog do!

- 1909 Deutsche Luftschiffahrt Aktien Gesellschaft (DELAG) – First commercial airline (using zepplins)
- 1914 St. Petersburg-Tampa Air Line (flying boats)
- 1949 Comet - First commercial jet airliner

- 1970 Bernard Sadlow adds wheels to luggage (lying flat)
- 1989 Robert Plath invents the wheelie bag (2 wheels and handle we have today)

How did we first disrupt?

My Original Marching Orders from Jon Postel

- Find something better than hosts.txt
- Look at 5 or so proposals, find a compromise
- But very clear that we needed something that scaled differently…

Intent of DNS protocol design 1983

- Provide a design that was just lightweight enough to take off – some things left out
- 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
  - Prepare for evolution

- Left Out
  - Security
  - Clever Replication
  - Access control
  - Class definition
  - Other data types

What happened?

RFC 882/883

1. Little “DNA” from the original proposals
2. UDP and Server Redundancy recipe is novel
3. RFC 882 & 883 (1983) lead to small changes and 1034 & 1035 (1987)

Thank you ARPA for supporting ISI and UCB and …

But the fire was lit – DNS RFC family tree

Lately the Marketeers and Politicians have been Disrupting Things

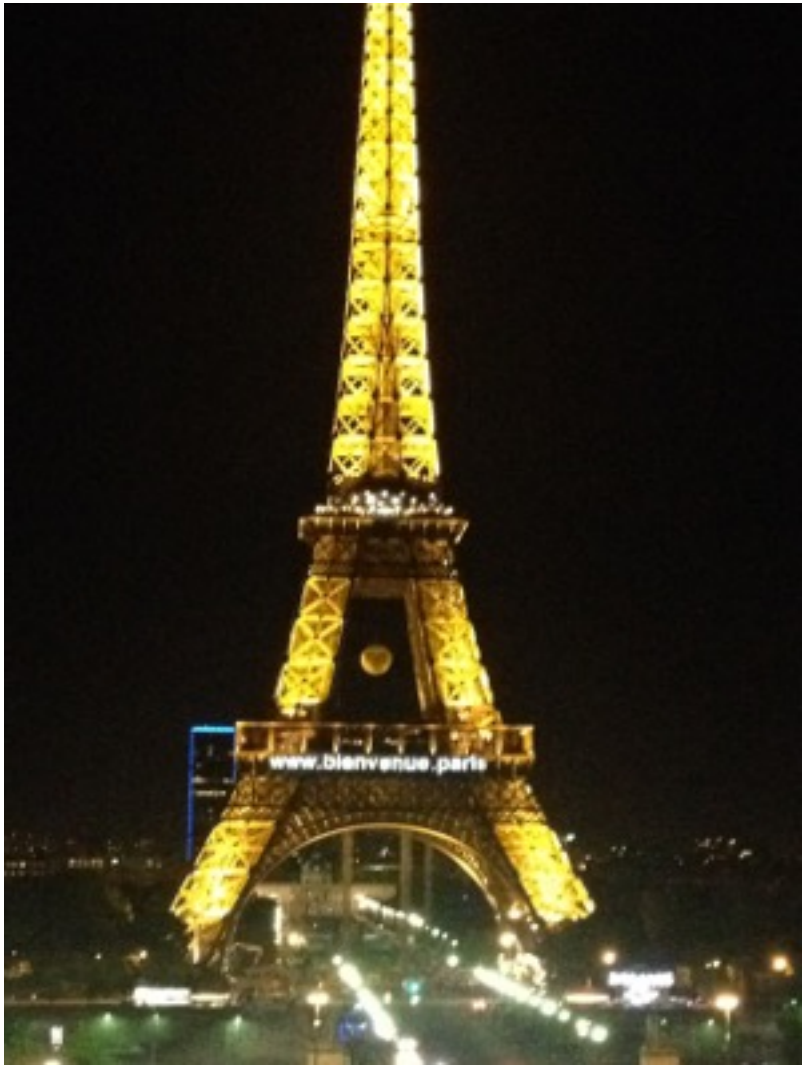

Progress:

Over $300,000,000 in

Application fees

.kosher live Feb 2014

.vin not yet

DNS->DN$?

GTLD Progress - Halloween 2014

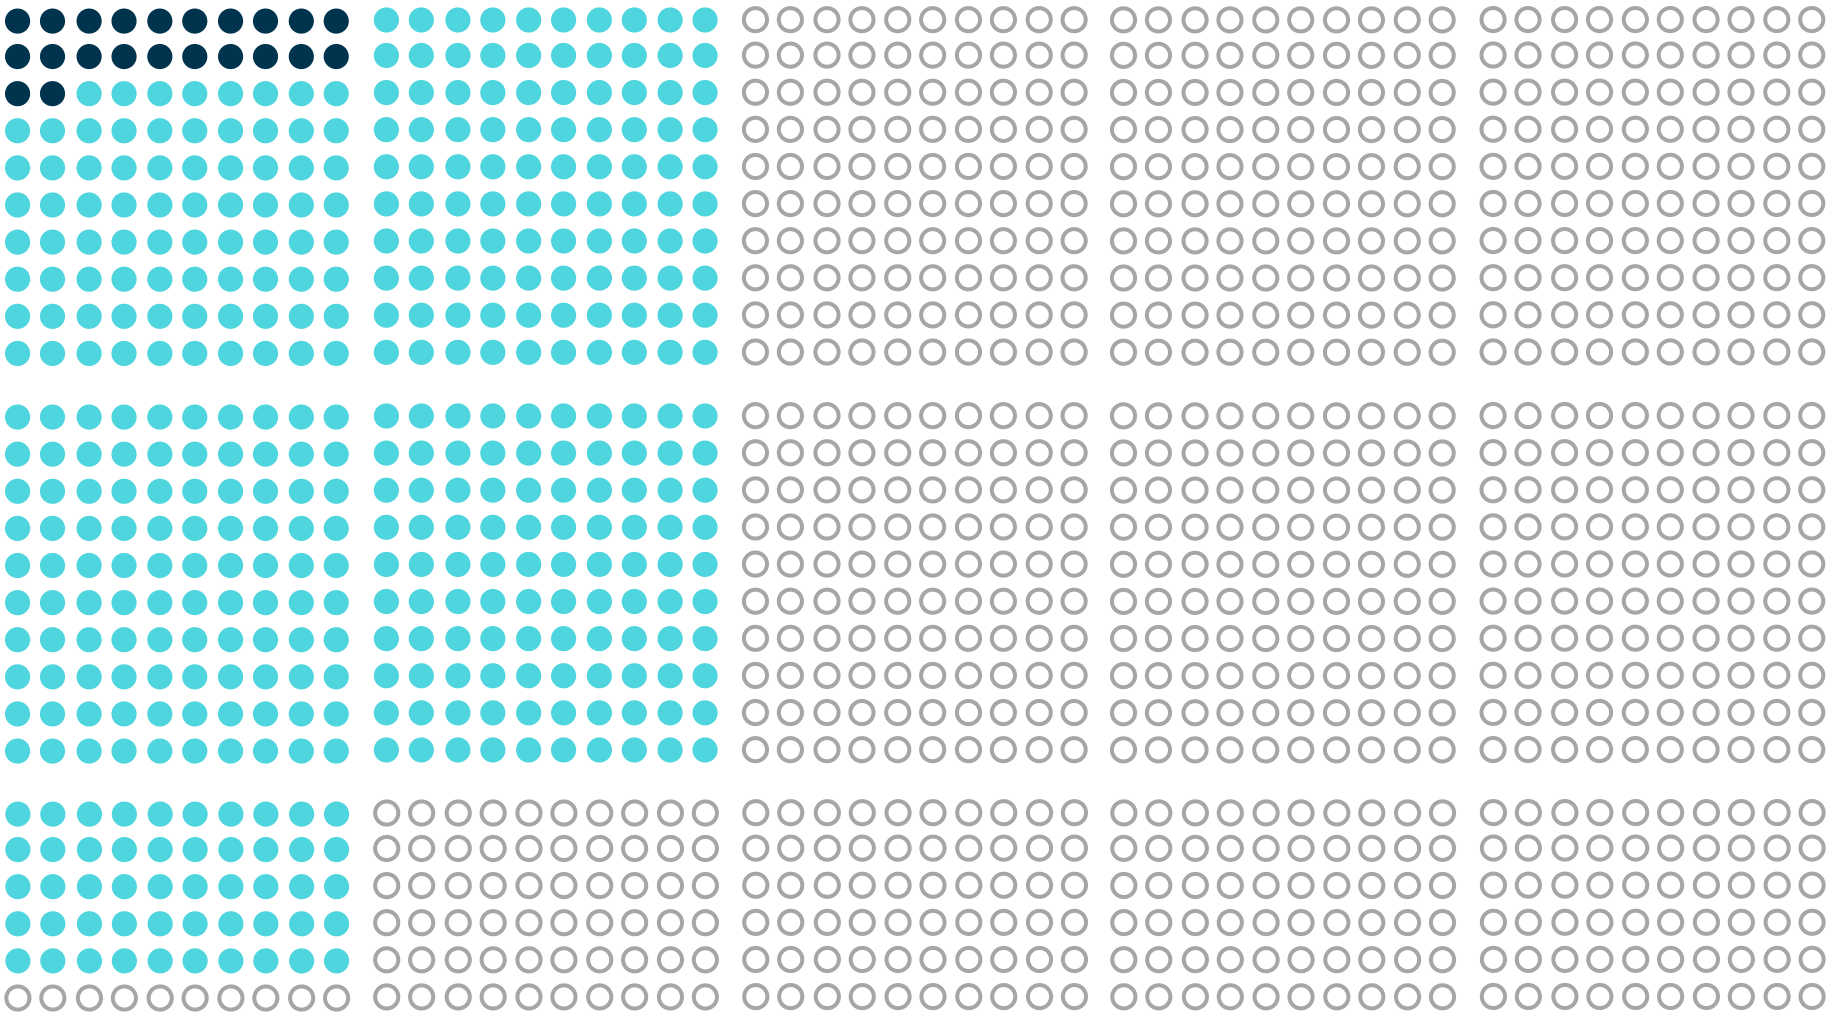

Key

Prior to October 2013

Current New gTLDs

Potential New gTLDs

Planned Disruption Ends

It’s 1989 - NSF, Want to improve DNS?

- Propose:
  - Fix bind
  - Address
    - Incremental update
    - Security
    - Crawl the Internet and build a distributed index stored in the DNS
    - Abuse (accidental DDOS)

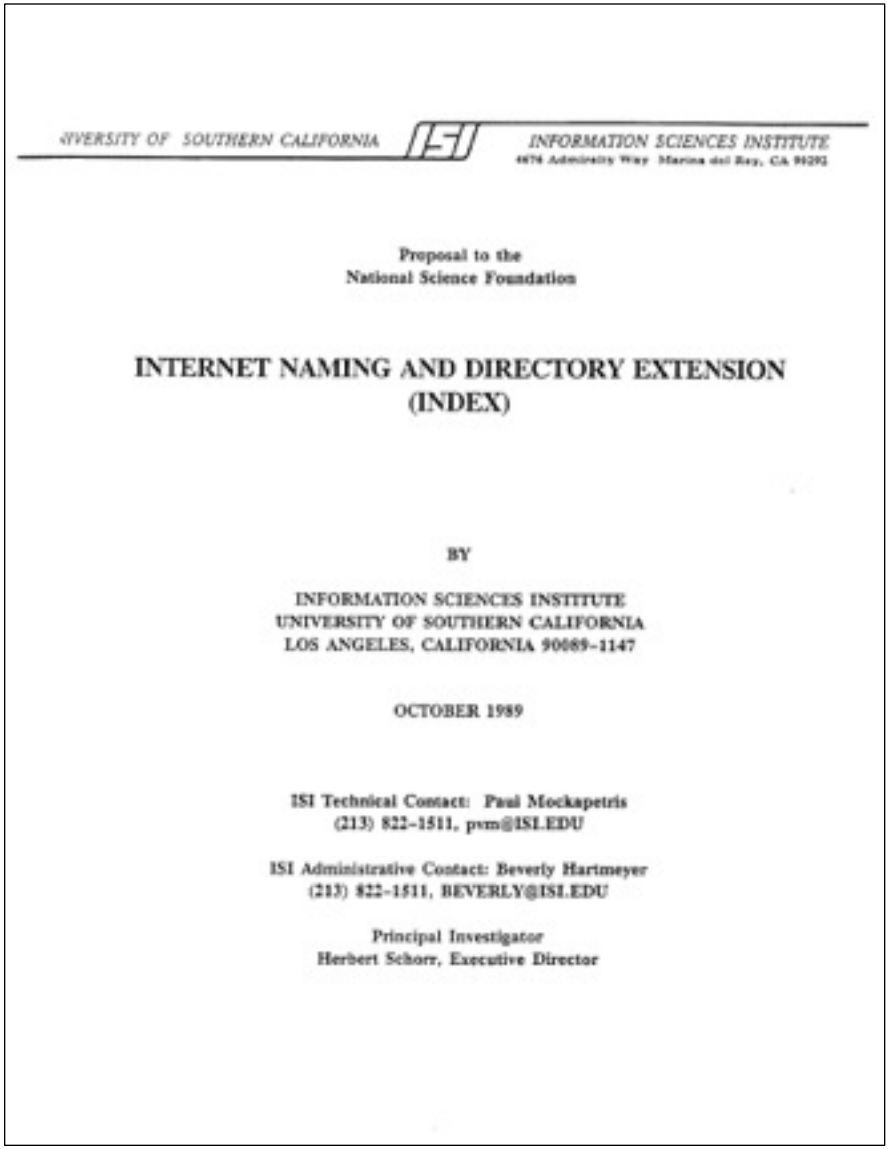

...IVERSITY OF SOUTHERN CALIFORNIA ISI INFORMATION SCIENCES INSTITUTE
4676 Admiralty Way Marina del Rey, CA 90292

Proposal to the
National Science Foundation

INTERNET NAMING AND DIRECTORY EXTENSION
(INDEX)

BY

INFORMATION SCIENCES INSTITUTE
UNIVERSITY OF SOUTHERN CALIFORNIA
LOS ANGELES, CALIFORNIA 90089-1147

OCTOBER 1989

ISI Technical Contact: Paul Mockapetris
(213) 822-1511, pvm@ISI.EDU

ISI Administrative Contact: Beverly Hartmeyer
(213) 822-1511, BEVERLY@ISI.EDU

Principal Investigator
Herbert Schorr, Executive Director

NSF feedback

- Reviewer 1: Excellent
- Reviewer 2: Very Good (critical, but not research)
- Reviewer 3: Very Good (please just fix bind)

- NSF Result: Can’t decide

- So much for planned evolution…

Today

Google Search Results on PhD Theses

"Domain Name System PhD Thesis"

2,110,000

"Transmission Control Protocol PhD Thesis"

167,000

"Internet Protocol PhD Thesis"

205,000

DNS Trends

- 1983 DNS starts, Paul receives ISO advice “We will bury you.”
- 1986 DNS liftoff – some machines have no host table

- 1989 Cache Poisoning observed “Don’t cache data just because somebody sends it to you”
- 1993 DNSSEC starts

- 2008 Kaminsky fast poisoning attack (We fight AGAINST Moore’s law)
- 2013 Snowden and reactions – opportunistic caching bad?
- 2014 DDOS
- 2015 Internet Governance – where will IANA be?

- 201X Majority of DNS secured with digital signatures

Threats vs. Defense

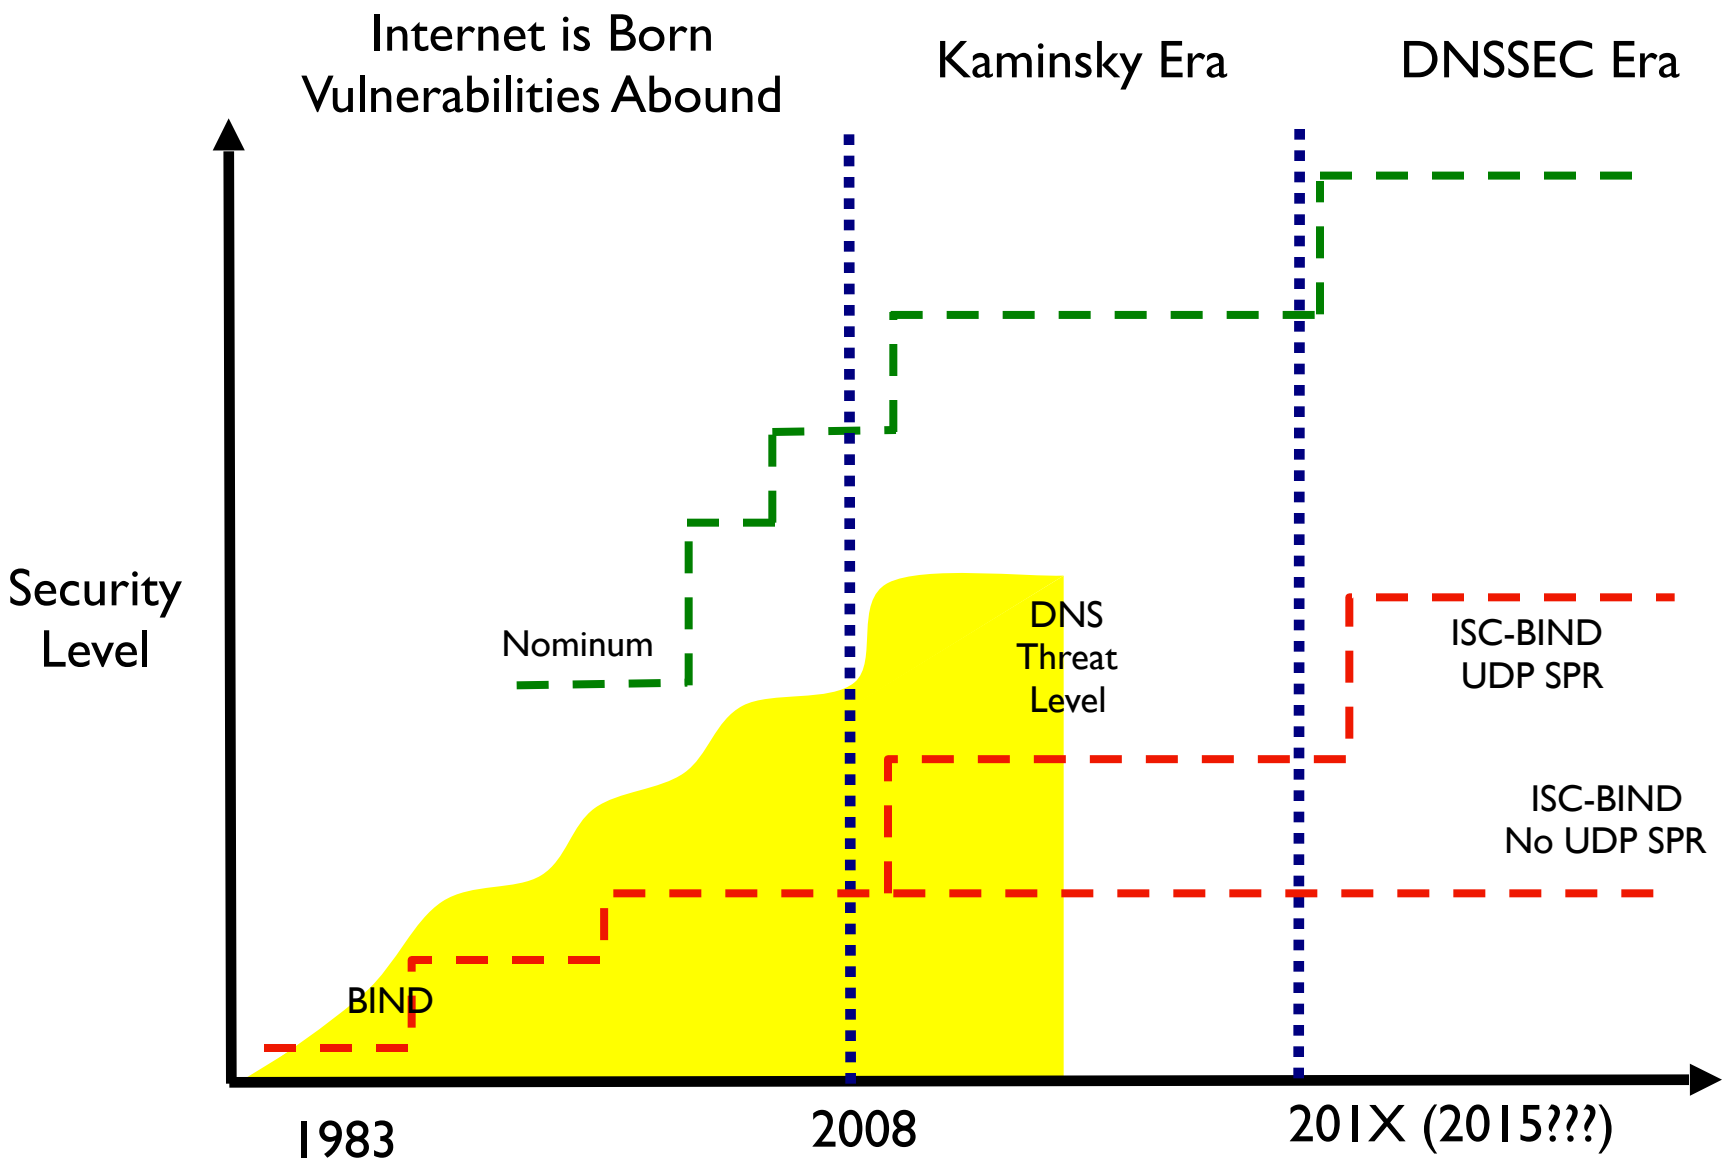

Some of the Internet is always broken

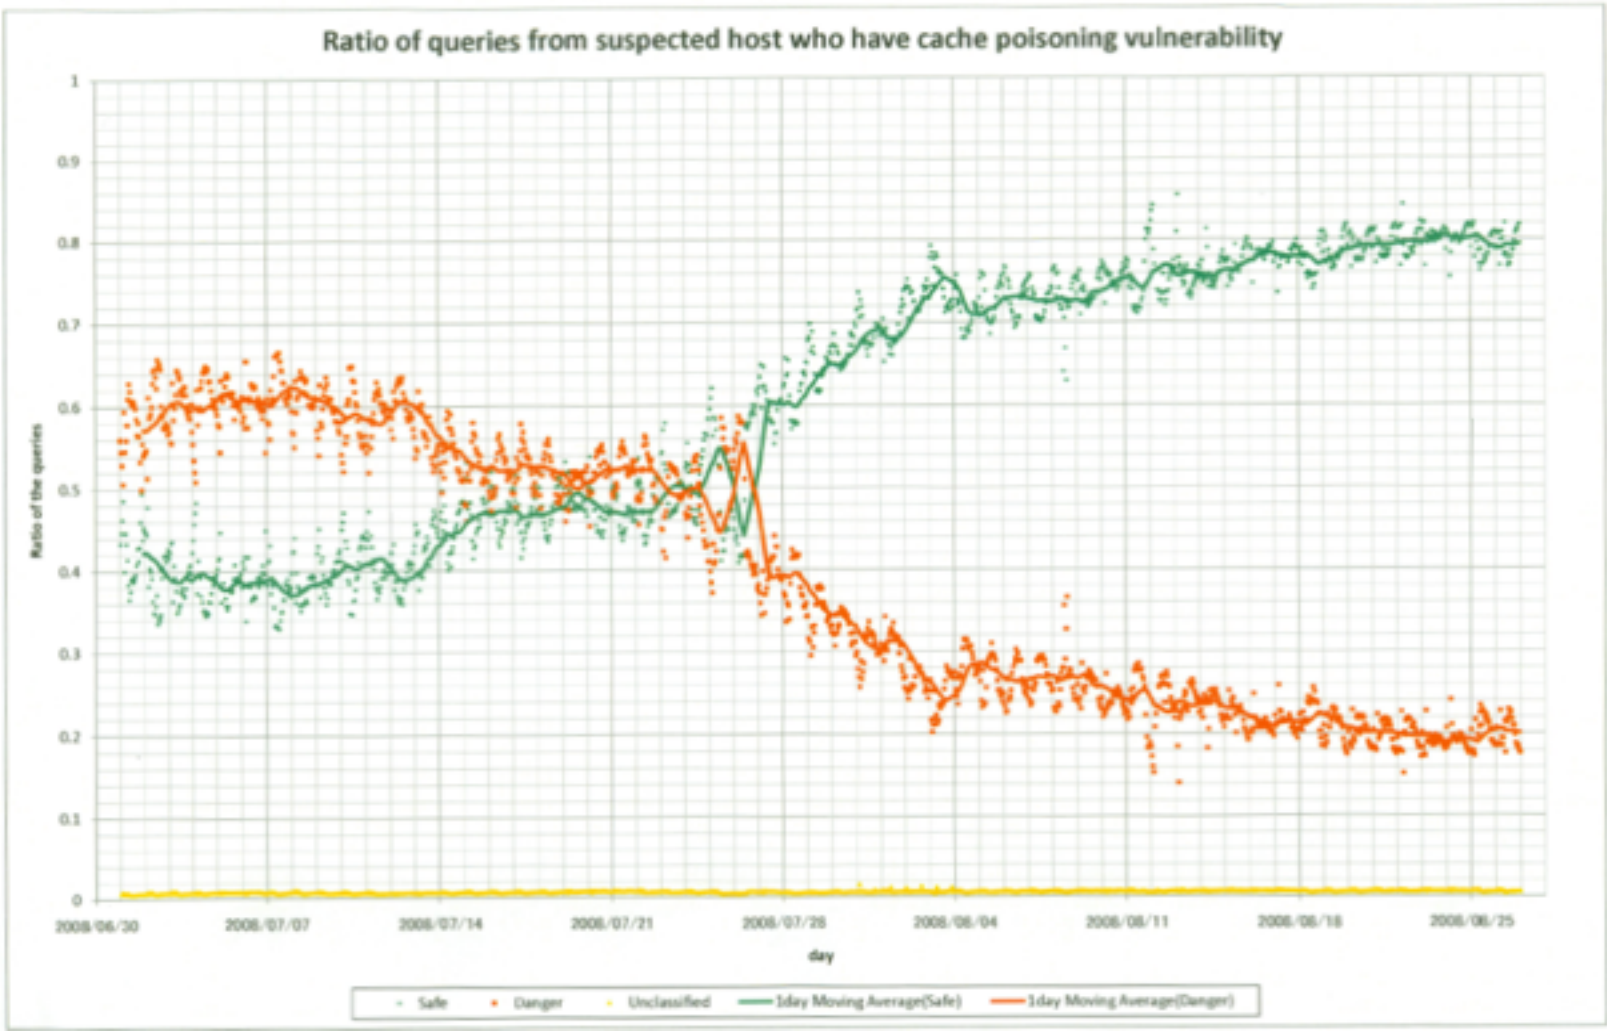

Different Rules for Yesterday and Tomorrow

- Datagrams are fast
- Opportunistic Caching
- One key to rule them all

- Datagrams for DDOS
- Privacy of queries and responses
- Multiple trust anchors

Future Directions

1. DNS Basic Algorithms

- Initial algorithms were purposely minimal – We can afford more now!
  - Don’t just go to the top and then down
- Is there a way to kill backward compatibility?
- Is there a way to get people to integrate authoritative and caching servers?

2. Information Centric Networks

- In some ways a better DNS
- Can we:
  - Merge the best ICN ideas into DNS?
  - Kill off DNS, replace with ICN?

- But ICN has its own set of issues:
  - Replacing infrastructure means a IPv6-like timeline, so just layer and get over it
  - More research on name structures, less on hardware
  - Which ICN?

3. Algorithmic Contracts – a personal favorite

- Do away with central management entirely, a la Bitcoin, etc

- Zone management becomes:
  - An accepted set of rules
  - Non-repudiable logs per delegation
  - No jurisdictional locus
  - One or more zone generators

- Extend to other applications
  - Number Portability
  - Contact Sharing
  - ...

Goals

- Create distributed algorithms, sometimes using trusted third parties, sometimes not, that can implement contract workflows, and interface with enforcement, payment, etc functions.

- Today seems to work in practice, e.g. bitcoin, namecoin, but not accepted in theory.

Sample Problems

- Registration
  - Internet TLDs and their management
    - Also addresses, ASNs, …
  - Portable Phone Numbers
  - “Do Not Call” registries
- Connection
  - Require security: car, airplane, smartphone busses
  - Require privacy: IOT tag call home, bluetooth, WiFi tracking
- Peering?
  - End to end QOS?
  - E2E virtual circuits

A brief Introduction to the DNS root

- A database of TLD data which is growing to ~2K entries, some TLDs are countries (ccTLD) e.g. .ES, some generic (gTLD) e.g. .COM. Or .ORG
- New varieties created recently e.g. .BANK
- Each TLD configured by a few records (5-10)
- Example records
  - Nameserver and nameserver addresses
  - Digital signatures

The DNS root (ccTLDs)

- Today:

1. TLD submits change to ICANN / Verisign on even/odd days
2. ICANN vets, Y/N
3. ICANN submits to USG
4. USG vets, Y/N
5. ICANN generates a candidate root zone twice a day, sends to Verisign
6. Verisign vets, Y/N
7. Verisign signs, sends to root operators
8. Root operators distribute

- One possible tomorrow:

1. TLD writes change to its own non-repudiable journal.
2. Other TLDs, ICANN can register requests for reconsideration
3. If TLD doesn't retract, independent zone builder collects from all TLD journals.
4. Sign it somehow (TBD)

Proposed tools

- Workflow description language
  - Perhaps as transition network, e.g. Petri
  - Public transparency vs. privacy
- Primitives
  - Voting
  - Auction
  - Timeout
  - Journals
  - Signature standards
- APIs
  - Payment
  - Notification

Thank You!