Fragmentation Considered Harmful?

Geoff Huston

APNIC

… “considered harmful”

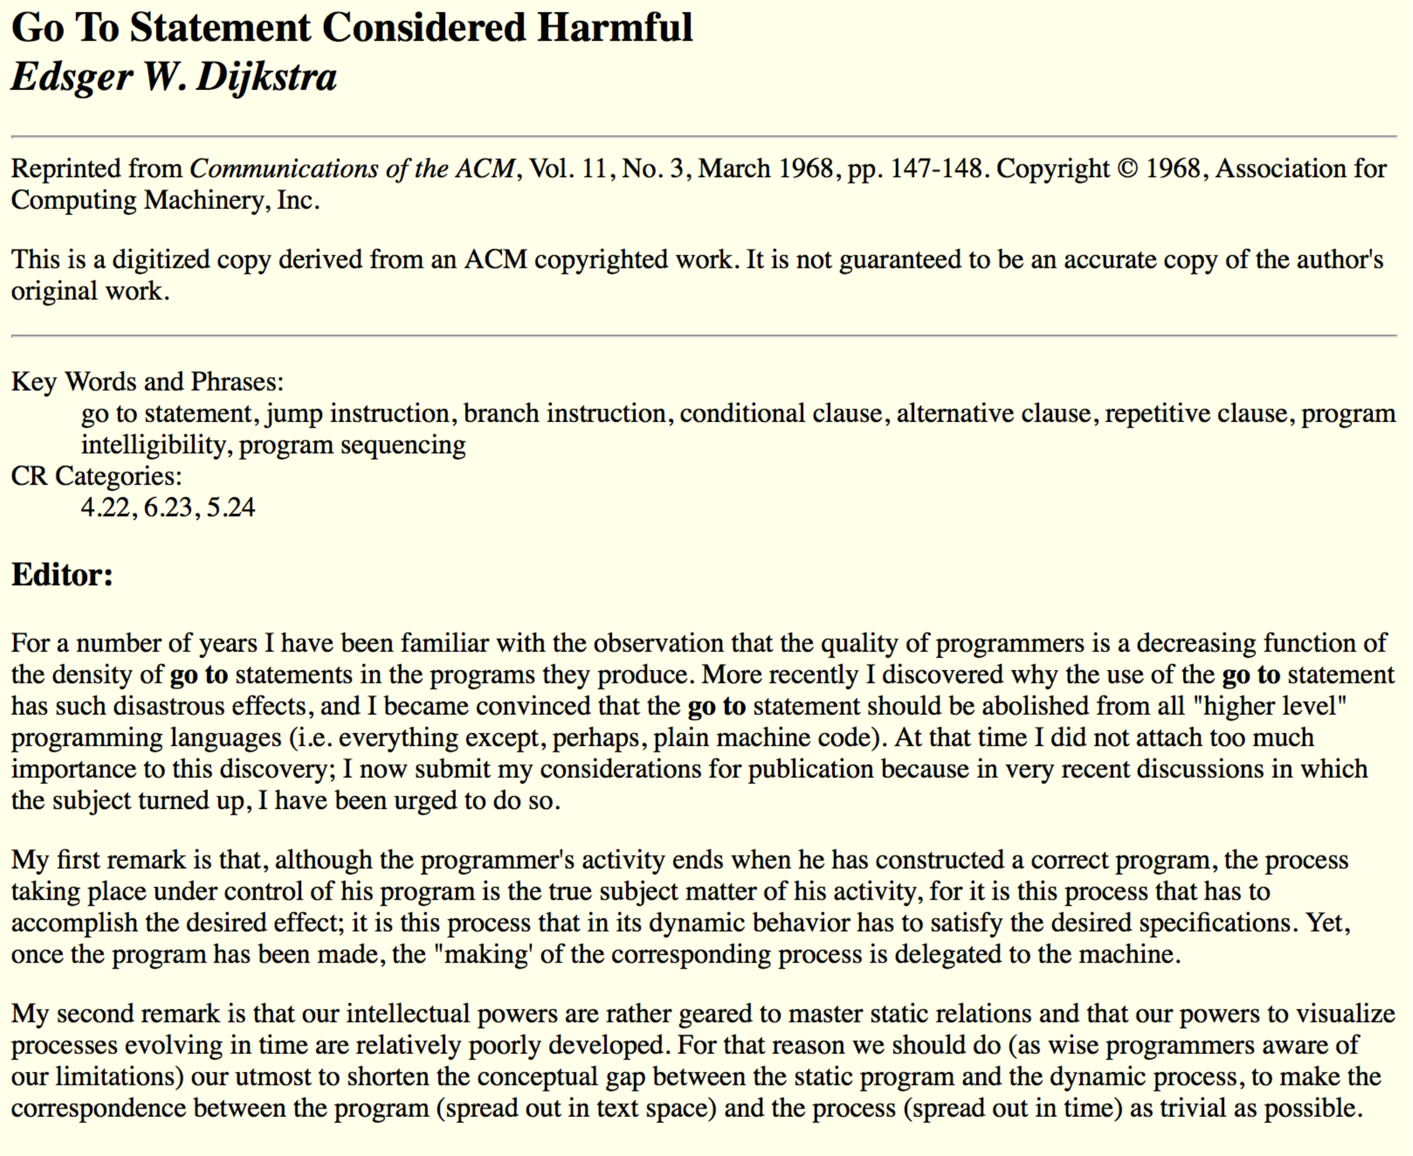

Go To Statement Considered Harmful

Edsger W. Dijkstra

Reprinted from *Communications of the ACM*, Vol. 11, No. 3, March 1968, pp. 147-148. Copyright © 1968, Association for Computing Machinery, Inc.

Key Words and Phrases:
go to statement, jump instruction, branch instruction, conditional clause, alternative clause, repetitive clause, program intelligibility, program sequencing
CR Categories:
4.22, 6.23, 5.24

Editor:

For a number of years I have been familiar with the observation that the quality of programmers is a decreasing function of the density of **go to** statements in the programs they produce. More recently I discovered why the use of the **go to** statement has such disastrous effects, and I became convinced that the **go to** statement should be abolished from all "higher level" programming languages (i.e. everything except, perhaps, plain machine code). At that time I did not attach too much importance to this discovery; I now submit my considerations for publication because in very recent discussions in which the subject turned up, I have been urged to do so.

My first remark is that, although the programmer's activity ends when he has constructed a correct program, the process taking place under control of his program is the true subject matter of his activity, for it is this process that has to accomplish the desired effect; it is this process that in its dynamic behavior has to satisfy the desired specifications. Yet, once the program has been made, the "making' of the corresponding process is delegated to the machine.

My second remark is that our intellectual powers are rather geared to master static relations and that our powers to visualize processes evolving in time are relatively poorly developed. For that reason we should do (as wise programmers aware of our limitations) our utmost to shorten the conceptual gap between the static program and the dynamic process, to make the correspondence between the program (spread out in text space) and the process (spread out in time) as trivial as possible.

What's "Fragmentation" anyway?

Some time ago we used to refer to this as "Network Balkanisation"

- It was a term used when talking about "alternate DNS roots" and other forms of of customised DNS name systems

These days I guess that "Network Fragmentation" is the politically correct term

- And it can refer to a broad range of diversity, in name resolution, packet routing, security realms, application behaviour, character sets, and similar possibilities for differential outcomes
- Although it still retains a flavour of enforcing geo-political boundaries and elimination cross-border dependencies

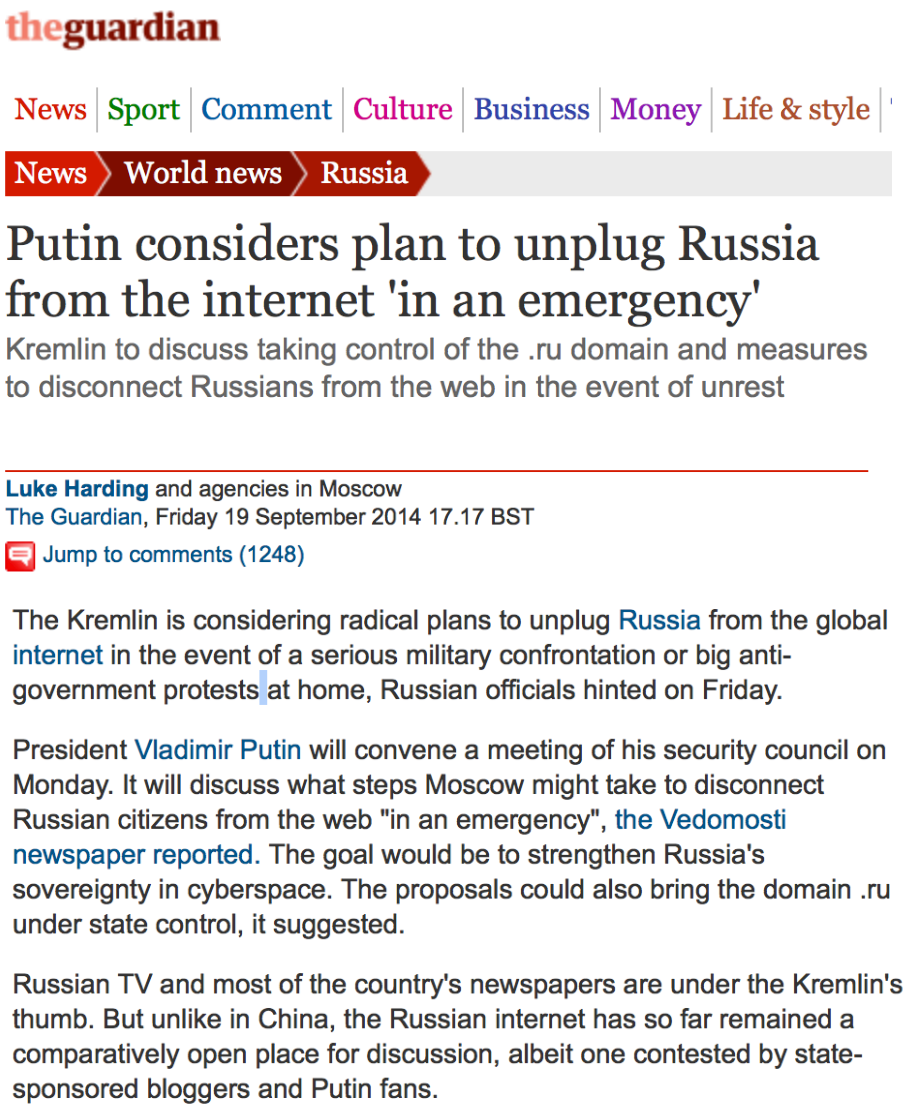

Putin considers plan to unplug Russia from the internet 'in an emergency'

Kremlin to discuss taking control of the .ru domain and measures to disconnect Russians from the web in the event of unrest

Luke Harding and agencies in Moscow
The Guardian, Friday 19 September 2014 17.17 BST

The Kremlin is considering radical plans to unplug Russia from the global internet in the event of a serious military confrontation or big anti-government protests at home, Russian officials hinted on Friday.

President Vladimir Putin will convene a meeting of his security council on Monday. It will discuss what steps Moscow might take to disconnect Russian citizens from the web "in an emergency", the Vedomosti newspaper reported. The goal would be to strengthen Russia's sovereignty in cyberspace. The proposals could also bring the domain .ru under state control, it suggested.

Russian TV and most of the country's newspapers are under the Kremlin's thumb. But unlike in China, the Russian internet has so far remained a comparatively open place for discussion, albeit one contested by state-sponsored bloggers and Putin fans.

Its possible that the overt US control of the root zone of the DNS lies behind much of these reported Russian concerns – if .ru was arbitrarily removed from the DNS root zone then many communications within Russia, and external communications to and from Russia would be disrupted.

Could the Russians “isolate” Russia from the Internet? Probably. Syria, Egypt, and others have done so in the recent past. If you cut the wires then most communications just stop.

Could the Russians force continuity of connectivity in the face of external efforts to isolate Russia? What if .ru was expunged from the DNS? Could this be circumvented? The answer to that hypothetical question is far harder. “Probably not” is about the best answer today.

Why?

- Why does the Internet have such a critical level of interdependence?
- And why is this dependence asymmetric?
- Is it even possible to conceive of an Internet as a loose coalition of national assets bound by selective bilateral arrangements and a loose regulatory binding structure managed through an international treaty structure?
- Could we build an Internet as a collection of national fragments? Or are there some fundamental components in the Internet's architecture that make this inconceivable?

The "Original" Internet Model

Network transactions are intentionally consistent:

https://www.mypage.com pulls down precisely the same content, no matter who requests the URL

We all use the same DNS data base

We all use the same resource object semantics

We all use the same protocols

We all use the same security framework

We all are connected to a common network

How was this consistency supported?

One DNS root hierarchy, and one name family

The unsuccessful "experiments" in alternate root zone files in the late 90's were perhaps one of the earliest forms of network fragmentation

They reinforced the criticality of the single root zone as the glue of the Internet

End-to-end application behaviour

The absence of a requirement for active middleware allowed each client to directly interact with the intended server. As long as the server was consistent, then every client was handled consistently and the network played no active role.

So, is this the case today?

Is the Internet still “consistent”?

Geo-Located Content serving

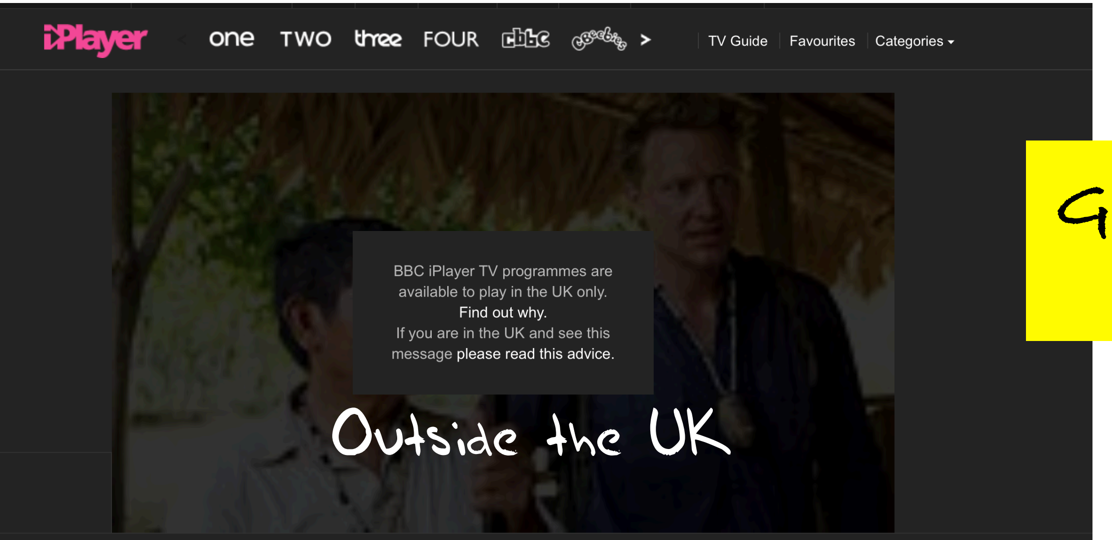

BBC iPlayer TV programmes are available to play in the UK only.
Find out why.
If you are in the UK and see this message please read this advice.

Outside the UK

The Taino people of the Caribbean had a multicultural society complete with drug-infused rituals, strange skulls and amazing navigation.

First shown: 9pm 22 Sep 2014
Available for 20 days
60 mins

Full description | Programme website | Credits

Download
AD Watch with AD
SL Watch with SL
Add to Favour
Share this pag

The assumed location of the end client (by source IP address) determines what content is passed back to the user

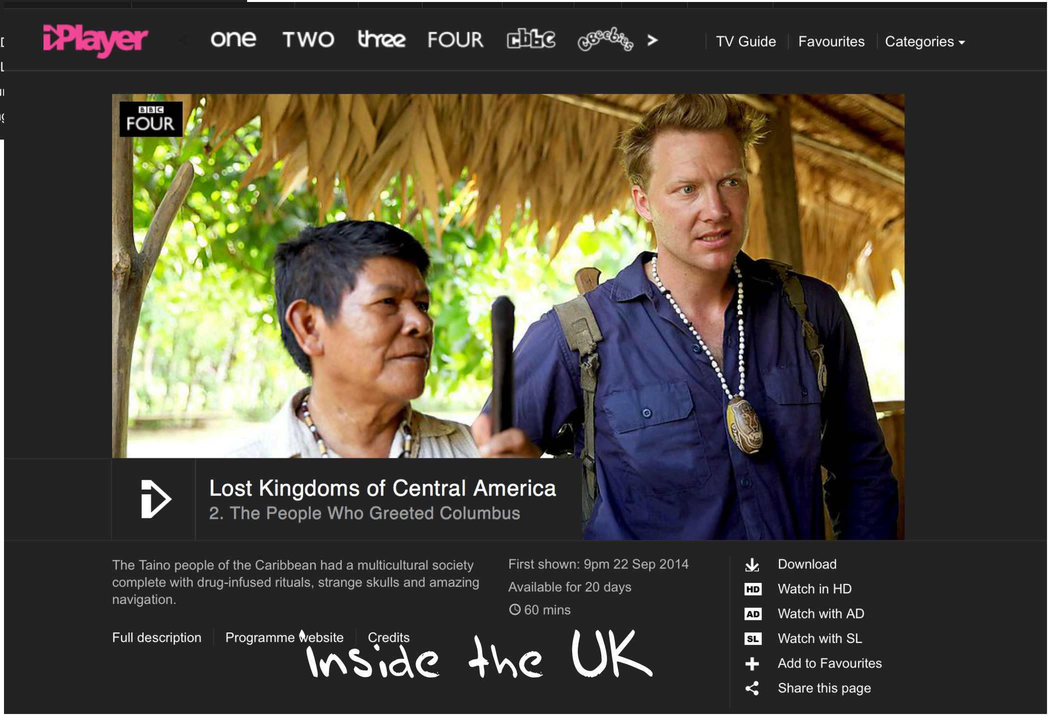

Lost Kingdoms of Central America
2. The People Who Greeted Columbus

The Taino people of the Caribbean had a multicultural society complete with drug-infused rituals, strange skulls and amazing navigation.

First shown: 9pm 22 Sep 2014
Available for 20 days
60 mins

Full description | Programme website | Credits

Download
HD Watch in HD
AD Watch with AD
SL Watch with SL
Add to Favourites
Share this page

Inside the UK

Content De-Fragmentation

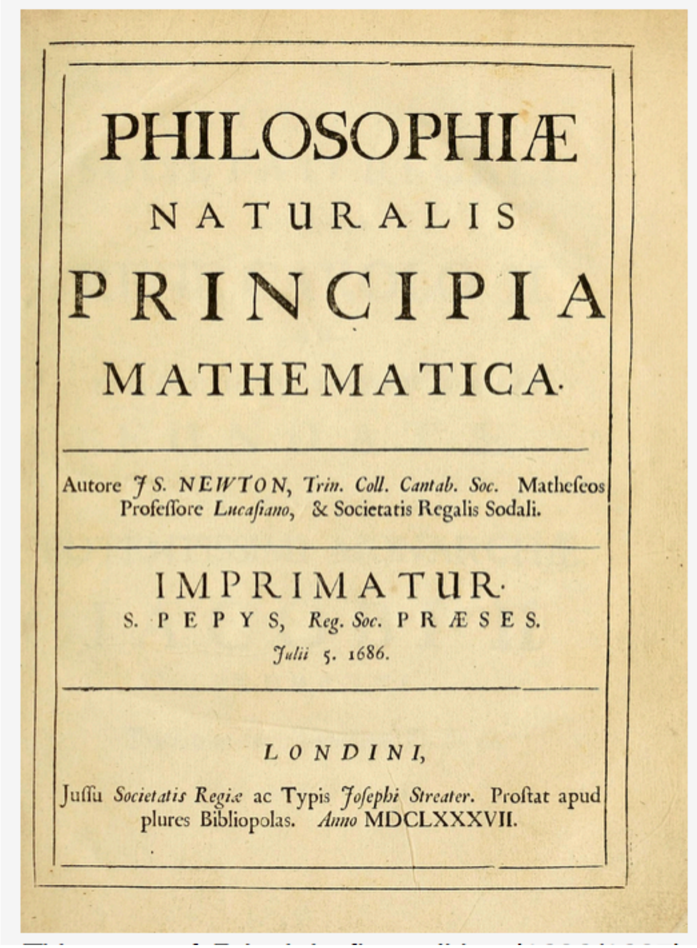

PHILOSOPHIÆ
NATURALIS
PRINCIPIA
MATHEMATICA.

Autore *J S.* NEWTON, *Trin. Coll. Cantab. Soc.* Matheſeos
Profeſſore *Lucaſiano*, & Societatis Regalis Sodali.

IMPRIMATUR.
S. PEPYS, *Reg. Soc.* PRÆSES.
Julii 5. 1686.

LONDINI,
Juſſu *Societatis Regiæ* ac Typis *Joſephi Streater*. Proſtat apud
plures Bibliopolas. *Anno* MDCLXXXVII.

Sometimes the remedies for network fragmentation are regulatory, sometimes through voluntary adoption of common codes of operational practice and common technology standards, and sometimes remedies lie in the deliberate actions of users to circumvent the imposed fragmentation

"Every action has an
equal and opposite
reaction"

The Rise of the Retail VPN Provider

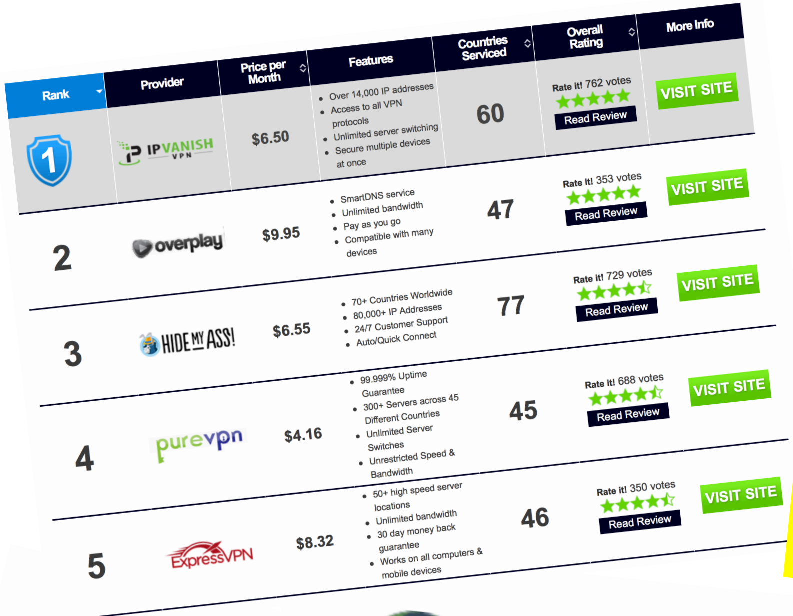

Rank	Provider	Price per Month	Features	Countries Serviced	Overall Rating	More Info
1	IPVANISH VPN	$6.50	• Over 14,000 IP addresses • Access to all VPN protocols • Unlimited server switching • Secure multiple devices at once	60	Rate it! 762 votes ★★★★★ Read Review	VISIT SITE
2	overplay	$9.95	• SmartDNS service • Unlimited bandwidth • Pay as you go • Compatible with many devices	47	Rate it! 353 votes ★★★★★ Read Review	VISIT SITE
3	HIDE MY ASS!	$6.55	• 70+ Countries Worldwide • 80,000+ IP Addresses • 24/7 Customer Support • Auto/Quick Connect	77	Rate it! 729 votes ★★★★☆ Read Review	VISIT SITE
4	purevpn	$4.16	• 99.999% Uptime Guarantee • 300+ Servers across 45 Different Countries • Unlimited Server Switches • Unrestricted Speed & Bandwidth	45	Rate it! 688 votes ★★★★☆ Read Review	VISIT SITE
5	ExpressVPN	$8.32	• 50+ high speed server locations • Unlimited bandwidth • 30 day money back guarantee • Works on all computers & mobile devices	46	Rate it! 350 votes ★★★★☆ Read Review	VISIT SITE

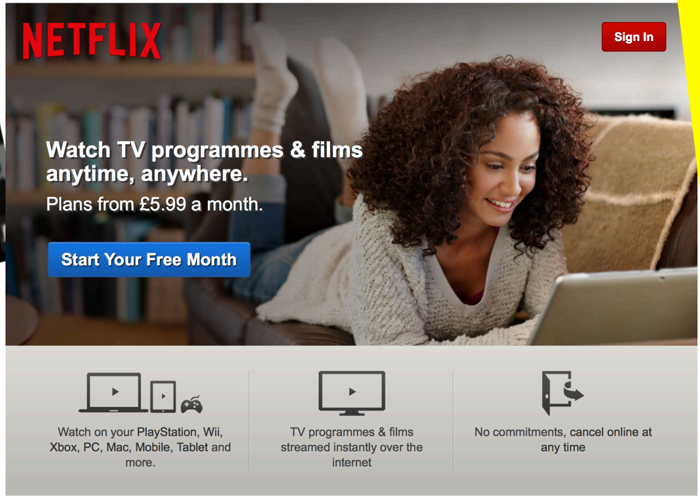

Watch TV programmes & films anytime, anywhere.

Plans from £5.99 a month.

Watch on your PlayStation, Wii, Xbox, PC, Mac, Mobile, Tablet and more.

TV programmes & films streamed instantly over the internet

No commitments, cancel online at any time

The VPN approach can be seen as a response to imposed content fragmentation, where the user's assumed identity is translated to one That can access the content

These VPNs conceal both the payload to the network, but also conceals the parties to a communication from the network and potentially from each other.

This is a case of the law of unintended outcomes in action, where encrypted tunnels have passed over from limited commercial use into popular retail products. This has its own Consequences not only for dis-intermediation in content distribution, but for LEA and national security functions

So, is this the case today?

Is the Internet still “consistent”?

China, late 2013

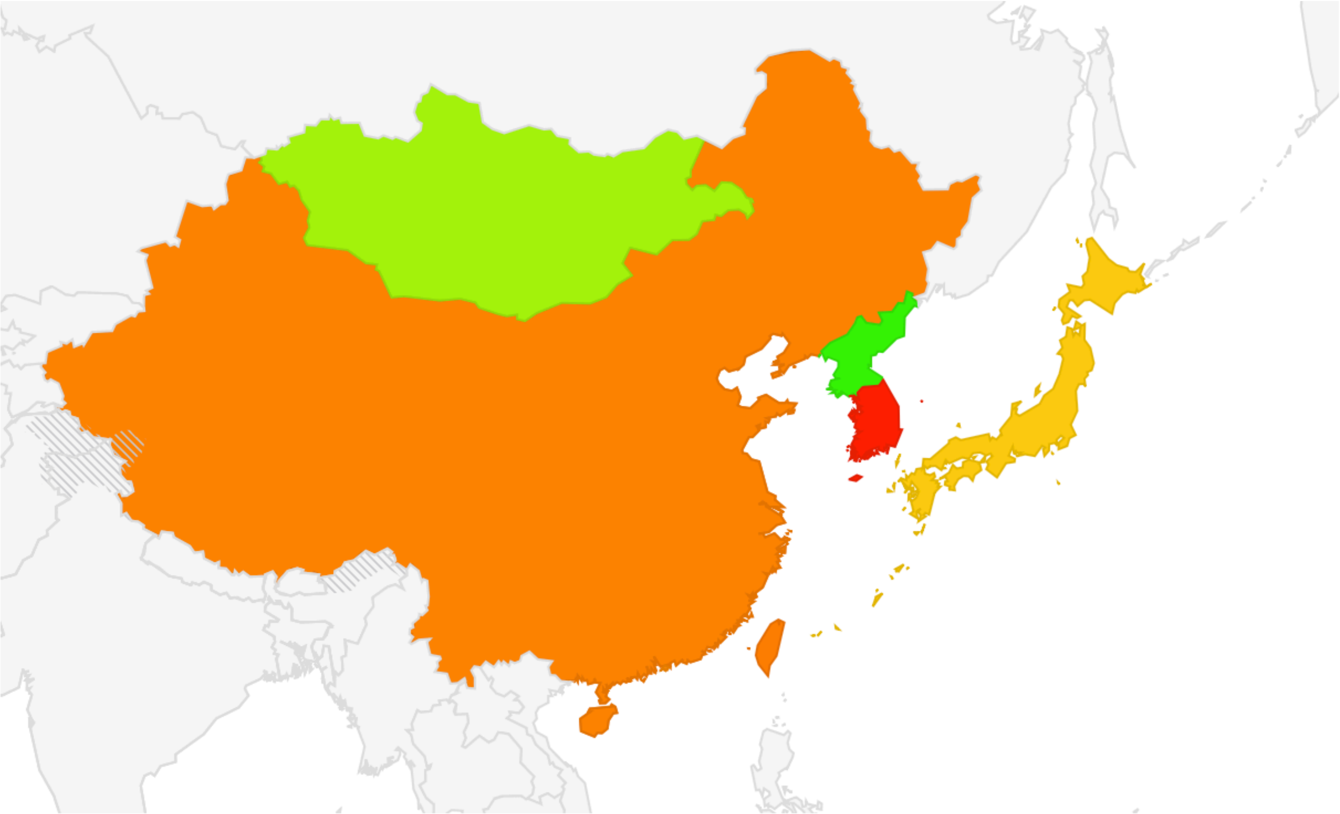

Outside China

```
$ dig www.facebook.com
69.171.229.25
$ dig AAAA www.facebook.com
2a03:2880:10:6f08:face:b00c:0:1
```

Inside China

```
$ dig www.facebook.com
1.1.1.1
$ dig AAAA www.facebook.com
2001:da8:112::21ae
```

Outside China

```
$ dig www.facebook.com
69.171.229.25
$ dig AAAA www.facebook.com
2a03:2880:10:6f08:face:b00c:0:1
```

Inside China

```
$ dig www.facebook.com
1.1.1.1
$ dig AAAA www.facebook.com
2001:da8:112::21ae
```

In this case the DNS providing different answers depending on where I was. Inside China I often received the address 1.1.1.1 in response to my DNS query for the name "www.facebook.com"

When I used a VPN to jump outside of China, I received the "real" DNS answers

Outside

```
NetRange:       69.171.224.0 - 69.171.255.255
CIDR:           69.171.224.0/19
OriginAS:       AS32934
NetName:        TFBNET3
NetHandle:      NET-69-171-224-0-1
Parent:         NET-69-0-0-0-0
NetType:        Direct Assignment
RegDate:        2010-08-05
Updated:        2012-02-24
Ref:            http://whois.arin.net/rest/net/NET-69-171-224-0-1

OrgName:        Facebook, Inc.
OrgId:          THEFA-3
Address:        1601 Willow Rd.
City:           Menlo Park
StateProv:      CA
PostalCode:     94025
Country:        US
RegDate:        2004-08-11
Updated:        2012-04-17
Ref:            http://whois.arin.net/rest/org/THEFA-3

OrgTechHandle: OPERA82-ARIN
OrgTechName:   Operations
OrgTechPhone:  +1-650-543-4800
OrgTechEmail:  noc@fb.com
OrgTechRef:    http://whois.arin.net/rest/poc/OPERA82-ARIN

OrgAbuseHandle: OPERA82-ARIN
OrgAbuseName:   Operations
OrgAbusePhone:  +1-650-543-4800
OrgAbuseEmail:  noc@fb.com
OrgAbuseRef:    http://whois.arin.net/rest/poc/OPERA82-ARIN


#
# ARIN WHOIS data and services are subject to the Terms of Use
# available at: https://www.arin.net/whois_tou.html
#
```

Inside

```
inetnum:        1.1.1.0 - 1.1.1.255
netname:        Debogon-prefix
descr:          APNIC Debogon Project
descr:          APNIC Pty Ltd
country:        AU
admin-c:        AR302-AP
tech-c:         AR302-AP
mnt-by:         APNIC-HM
mnt-routes:     MAINT-AU-APNIC-GM85-AP
mnt-irt:        IRT-APNICRANDNET-AU
status:         ASSIGNED PORTABLE
changed:        hm-changed@apnic.net 20110922
source:         APNIC

irt:            IRT-APNICRANDNET-AU
address:        PO Box 3646
address:        South Brisbane, QLD 4101
address:        Australia
e-mail:         abuse@apnic.net
abuse-mailbox:  abuse@apnic.net
admin-c:        AR302-AP
tech-c:         AR302-AP
mnt-by:         MAINT-AU-APNIC-GM85-AP
changed:        hm-changed@apnic.net 20110922
source:         APNIC

role:           APNIC RESEARCH
```

Outside

Facebook, INC

```
NetRange:       69.171.224.0 - 69.171.255.255
CIDR:           69.171.224.0/19
OriginAS:       AS32934
NetName:        TFBNET3
NetHandle:      NET-69-171-224-0-1
Pa
Ne
Re
Up
Ref:            http://whois.arin.net/rest/net/NET-69-171-224-0-1

OrgName:        Facebook, Inc.
OrgId:          THEFA-3
Address:        1601 Willow Rd.
City:           Menlo Park
StateProv:      CA
PostalCode:     94025
Country:        US
RegDate:        2004-08-11
Updated:        2012-04-17
Ref:            http://whois.arin.net/rest/org/THEFA-3

OrgTechHandle: OPERA82-ARIN
OrgTechName:   Operations
OrgTechPhone:  +1-650-543-4800
OrgTechEmail:  noc@fb.com
OrgTechRef:    http://whois.arin.net/rest/poc/OPERA82-ARIN

OrgAbuseHandle: OPERA82-ARIN
OrgAbuseName:   Operations
OrgAbusePhone:  +1-650-543-4800
OrgAbuseEmail:  noc@fb.com
OrgAbuseRef:    http://whois.arin.net/rest/poc/OPERA82-ARIN


#
# ARIN WHOIS data and services are subject to the Terms of Use
# available at: https://www.arin.net/whois_tou.html
#
```

Inside

Not Facebook!

```
inetnum:        1.1.1.0 - 1.1.1.255
netname:        Debogon-prefix
descr:          APNIC Debogon Project
descr:          APNIC Pty Ltd
country:        AU
adm
tec
mnt
mnt
mnt-irt:        IRT-APNICRANDNET-AU
status:         ASSIGNED PORTABLE
changed:        hm-changed@apnic.net 20110922
source:         APNIC

irt:            IRT-APNICRANDNET-AU
address:        PO Box 3646
address:        South Brisbane, QLD 4101
address:        Australia
e-mail:         abuse@apnic.net
abuse-mailbox:  abuse@apnic.net
admin-c:        AR302-AP
tech-c:         AR302-AP
mnt-by:         MAINT-AU-APNIC-GM85-AP
changed:        hm-changed@apnic.net 20110922
source:         APNIC

role:           APNIC RESEARCH
```

~~Outside~~ Inside

So let’s focus on that inside address: 1.1.1.1

Is this the “real thing” or just a local route to some local black hole?

Let’s resume the Inside/Outside examination, but focus just on the address 1.1.1.1

Outside

Inside

```
$ traceroute 1.1.1.1
traceroute to 1.1.1.1 (1.1.1.1), 64 hops max, 52 byte packets
 1  202.158.221.221 (202.158.221.221)  0.266 ms  0.202 ms
0.194 ms
 2  ge-4-0-0.bb1.b.cbr.aarnet.net.au (202.158.208.81)  0.491 ms
0.438 ms  0.424 ms
 3  so-0-1-0.bb1.b.mel.aarnet.net.au (202.158.194.42)  8.001 ms
8.000 ms  7.992 ms
 4  xe-0-0-0.pe1.a.mel.aarnet.net.au (202.158.200.12)  8.110 ms
8.088 ms  8.078 ms
 5  gw1.xe-0-0-2.pe1.a.mel.aarnet.net.au (202.158.210.41)
12.165 ms  12.193 ms  12.159 ms
 6  66.249.95.232 (66.249.95.232)  12.823 ms  13.254 ms
    66.249.95.234 (66.249.95.234)  20.282 ms
 7  209.85.249.52 (209.85.249.52)  134.637 ms  109.393 ms
109.426 ms
 8  64
    64
 9  20
    20
10  64
211.694 ms
11  72.14
264.518 m
12  216.2
    72.14
    216.2
13  72.14
    66.24
14  72.14
    72.14.236.147 (72.14.236.147)  278.543 ms
    72.14.236.99 (72.14.236.99)  412.531 ms
15  209.85.252.47 (209.85.252.47)  256.836 ms
    209.85.252.81 (209.85.252.81)  253.370 ms
    209.85.252.47 (209.85.252.47)  254.984 ms
16  65.210.126.78 (65.210.126.78)  255.713 ms  253.913 ms
253.865 ms
```

```
$ traceroute 1.1.1.1
traceroute to 1.1.1.1 (1.1.1.1), 64 hops max, 52 byte packets
 1  254 (220.247.145.254)  6.620 ms  1.774 ms  2.561 ms
 2  * * *
 3  192.168.9.5 (192.168.9.5)  6.718 ms  3.322 ms  5.324 ms
 4  159.226.253.189 (159.226.253.189)  25.557 ms  26.145 ms
27.191 ms
 5  8.130 (159.226.253.57)  26.059 ms  27.225 ms  25.889 ms
 6  8.198 (159.226.253.50)  27.060 ms  29.788 ms  29.476 ms
 7  8.192 (159.226.254.254)  64.767 ms  66.801 ms  66.768 ms
 8  72.14.221.138 (72.14.221.138)  100.753 ms  105.117 ms
99.613 ms
 9  209.85.241.56 (209.85.241.56)  101.914 ms
    209.85.241.58 (209.85.241.58)  105.561 ms  101.748 ms
10  66.249.94.31 (66.249.94.31)  175.902 ms
    209.85.255.58 (209.85.255.58)  154.349 ms
                                             .374 ms
                                                 ms
                                             13 ms
                                             13 ms
17  72.14.236.147 (72.14.236.147)  313.601 ms
    72.14.236.149 (72.14.236.149)  261.617 ms  305.524 ms
18  209.85.252.47 (209.85.252.47)  306.003 ms
    209.85.252.81 (209.85.252.81)  276.541 ms  363.910 ms
19  65.210.126.78 (65.210.126.78)  398.681 ms  361.306 ms
366.265 ms
```

Hang on… BOTH inside and outside ROUTE the packets addressed to 1.1.1.1 to the same endpoint: **65.210.126.78**

What we are seeing here is that the content blocking system is focused on the name component, rather than on attempting to block the traffic flow

Within China this form of DNS filtering is commonplace

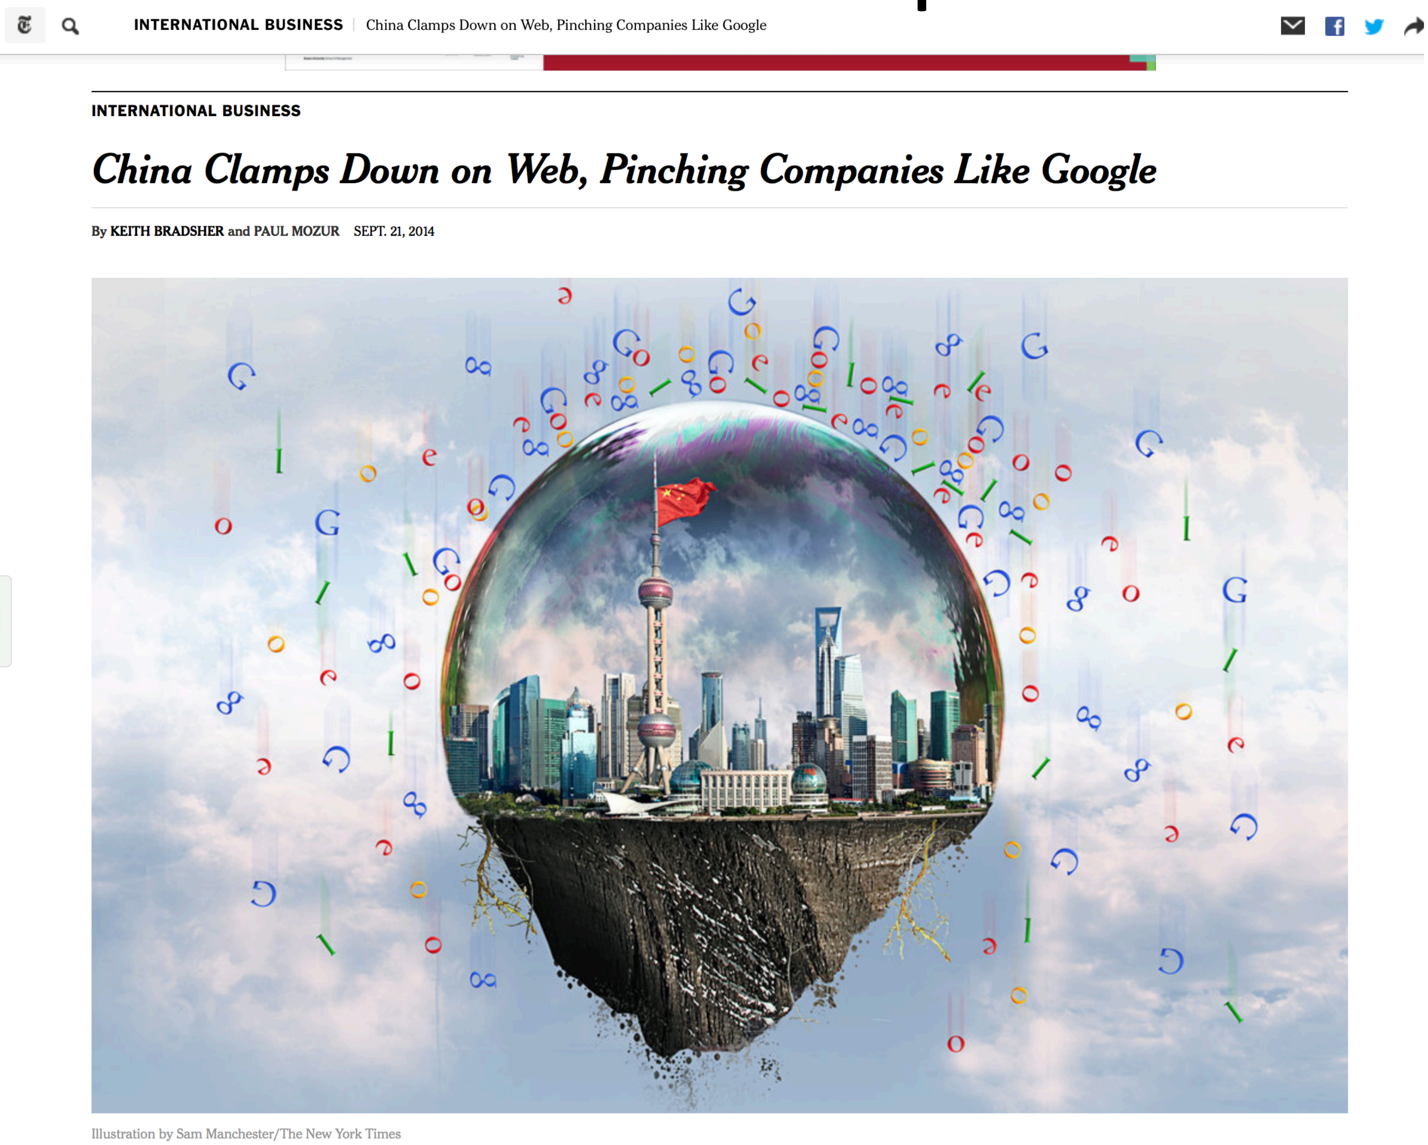

INTERNATIONAL BUSINESS | China Clamps Down on Web, Pinching Companies Like Google

INTERNATIONAL BUSINESS

China Clamps Down on Web, Pinching Companies Like Google

By KEITH BRADSHER and PAUL MOZUR SEPT. 21, 2014

Illustration by Sam Manchester/The New York Times

Within China this form of DNS filtering is commonplace

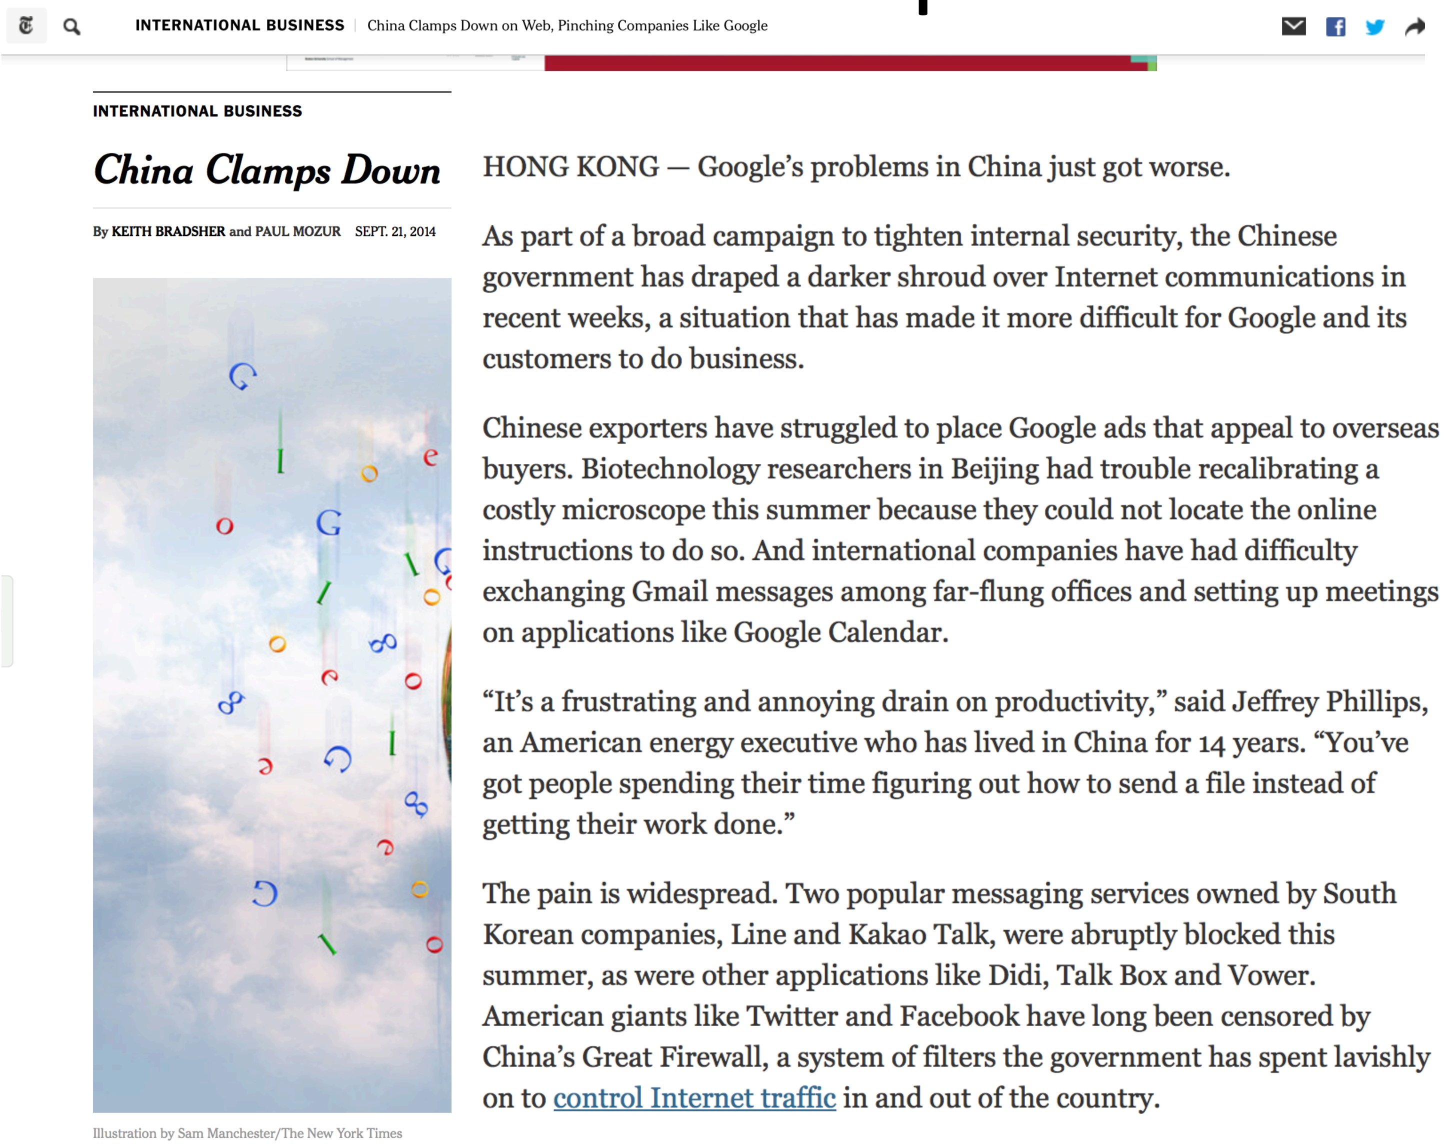

INTERNATIONAL BUSINESS | China Clamps Down on Web, Pinching Companies Like Google

INTERNATIONAL BUSINESS

China Clamps Down

By KEITH BRADSHER and PAUL MOZUR SEPT. 21, 2014

Illustration by Sam Manchester/The New York Times

HONG KONG — Google's problems in China just got worse.

As part of a broad campaign to tighten internal security, the Chinese government has draped a darker shroud over Internet communications in recent weeks, a situation that has made it more difficult for Google and its customers to do business.

Chinese exporters have struggled to place Google ads that appeal to overseas buyers. Biotechnology researchers in Beijing had trouble recalibrating a costly microscope this summer because they could not locate the online instructions to do so. And international companies have had difficulty exchanging Gmail messages among far-flung offices and setting up meetings on applications like Google Calendar.

"It's a frustrating and annoying drain on productivity," said Jeffrey Phillips, an American energy executive who has lived in China for 14 years. "You've got people spending their time figuring out how to send a file instead of getting their work done."

The pain is widespread. Two popular messaging services owned by South Korean companies, Line and Kakao Talk, were abruptly blocked this summer, as were other applications like Didi, Talk Box and Vower. American giants like Twitter and Facebook have long been censored by China's Great Firewall, a system of filters the government has spent lavishly on to control Internet traffic in and out of the country.

But it’s not just China

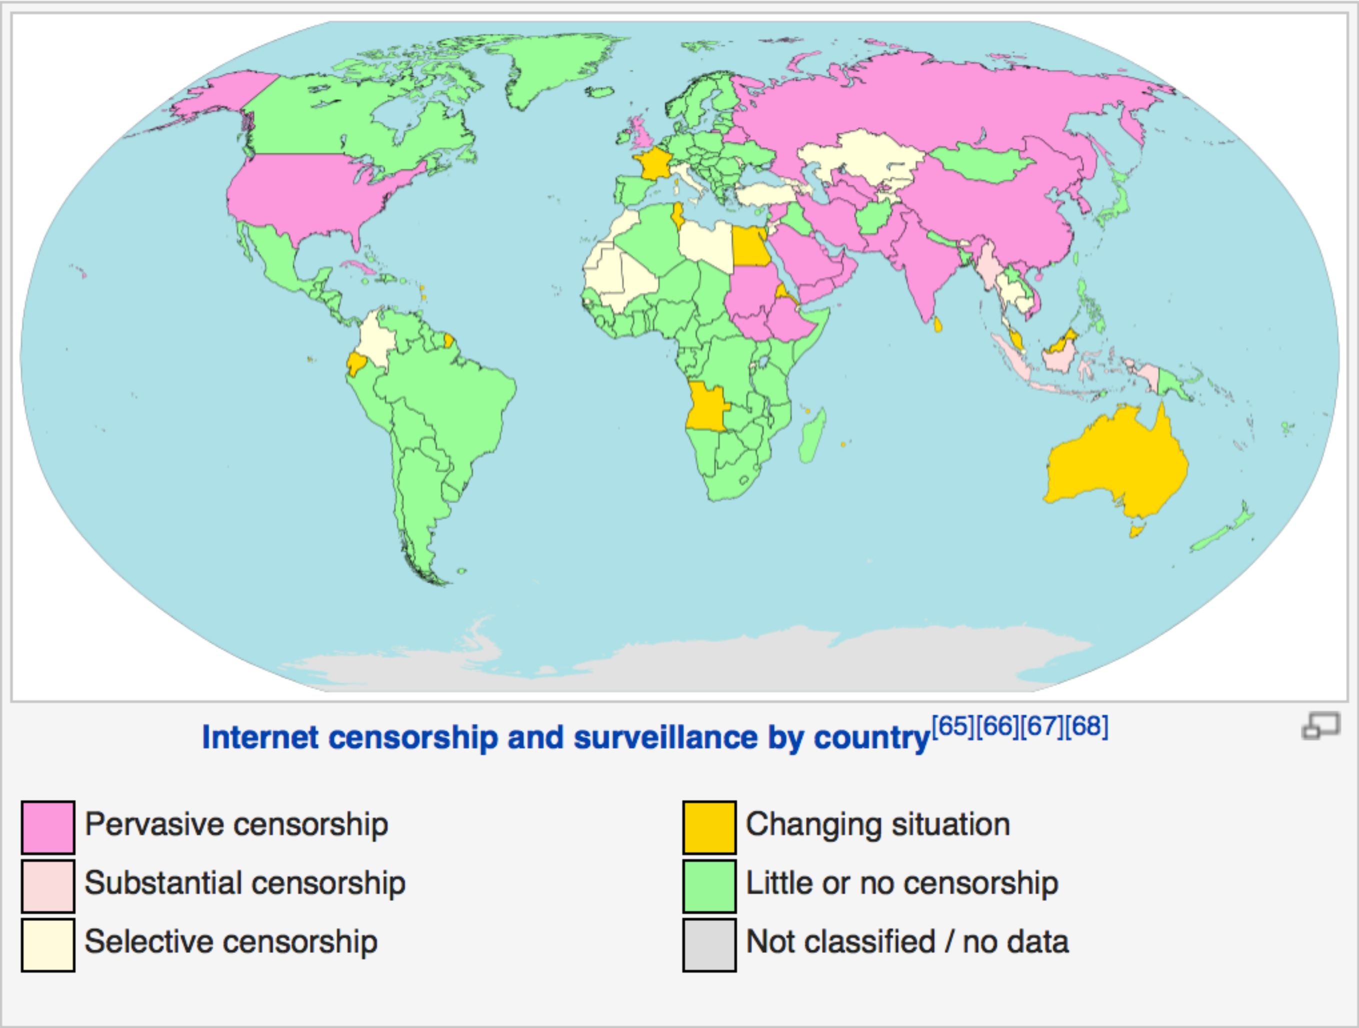

Internet censorship and surveillance by country[65][66][67][68]

- ■ Pervasive censorship
- ■ Substantial censorship
- ■ Selective censorship
- ■ Changing situation
- ■ Little or no censorship
- ■ Not classified / no data

http://en.wikipedia.org/wiki/Internet_censorship

Selective DNS “fragmentation” is widespread

- The DNS is easy to intercept
- Its easy to intercede and synthesize false answers to incoming DNS queries
  - It’s often easier to intercept the DNS than it is to intercept traffic on the wire
- End user applications are too credulous
  - They will believe anything the DNS appears to tell them!

The business of content filtering through DNS manipulation

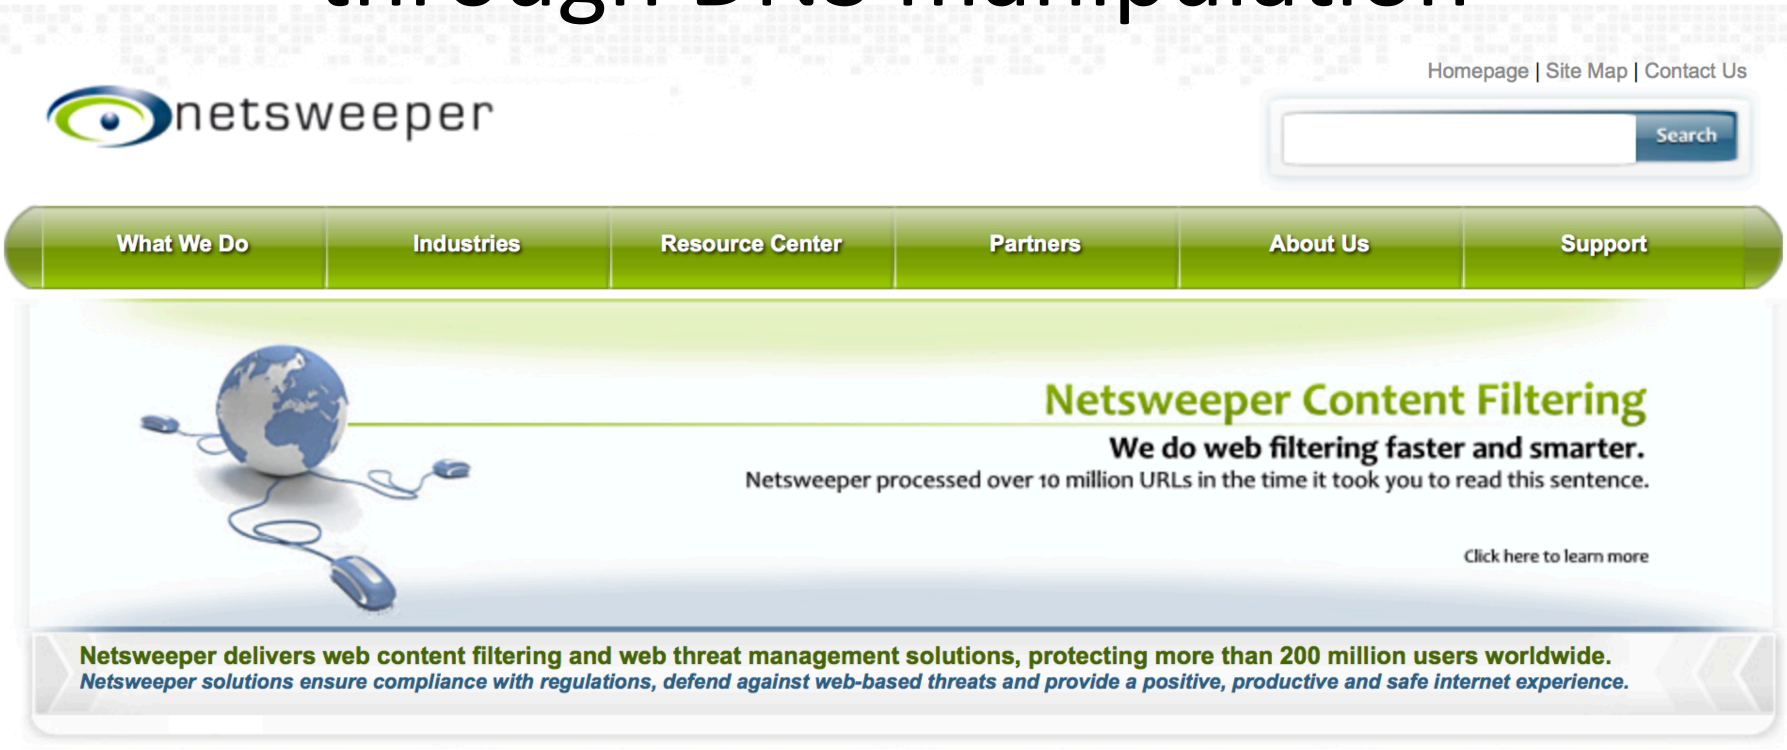

Netsweeper Content Filtering

We do web filtering faster and smarter.

Netsweeper processed over 10 million URLs in the time it took you to read this sentence.

Netsweeper delivers web content filtering and web threat management solutions, protecting more than 200 million users worldwide.

Netsweeper solutions ensure compliance with regulations, defend against web-based threats and provide a positive, productive and safe internet experience.

"Among the most popular filtering software programs is SmartFilter by Secure Computing in California, which was bought by McAfee in 2008. SmartFilter has been used by Tunisia, Saudi Arabia, Sudan, the UAE, Kuwait, Bahrain, Iran, and Oman, as well as the United States and the UK. Myanmar and Yemen have used filtering software from Websense. The Canadian-made commercial filter Netsweeper is used in Qatar, the UAE, and Yemen."

http://en.wikipedia.org/wiki/Internet_censorship

So this is “bad” – right?

“bad” in the sense that it doesn’t really pose a barrier that can’t be solved by any user equipped with a decent search engine and a few minutes to install a workaround, so its “bad” in the sense that it’s a relatively meaningless imposition of an inefficiency within the network

So this is "bad" – right?

- Well, not always.
- What if your business is all about speed?

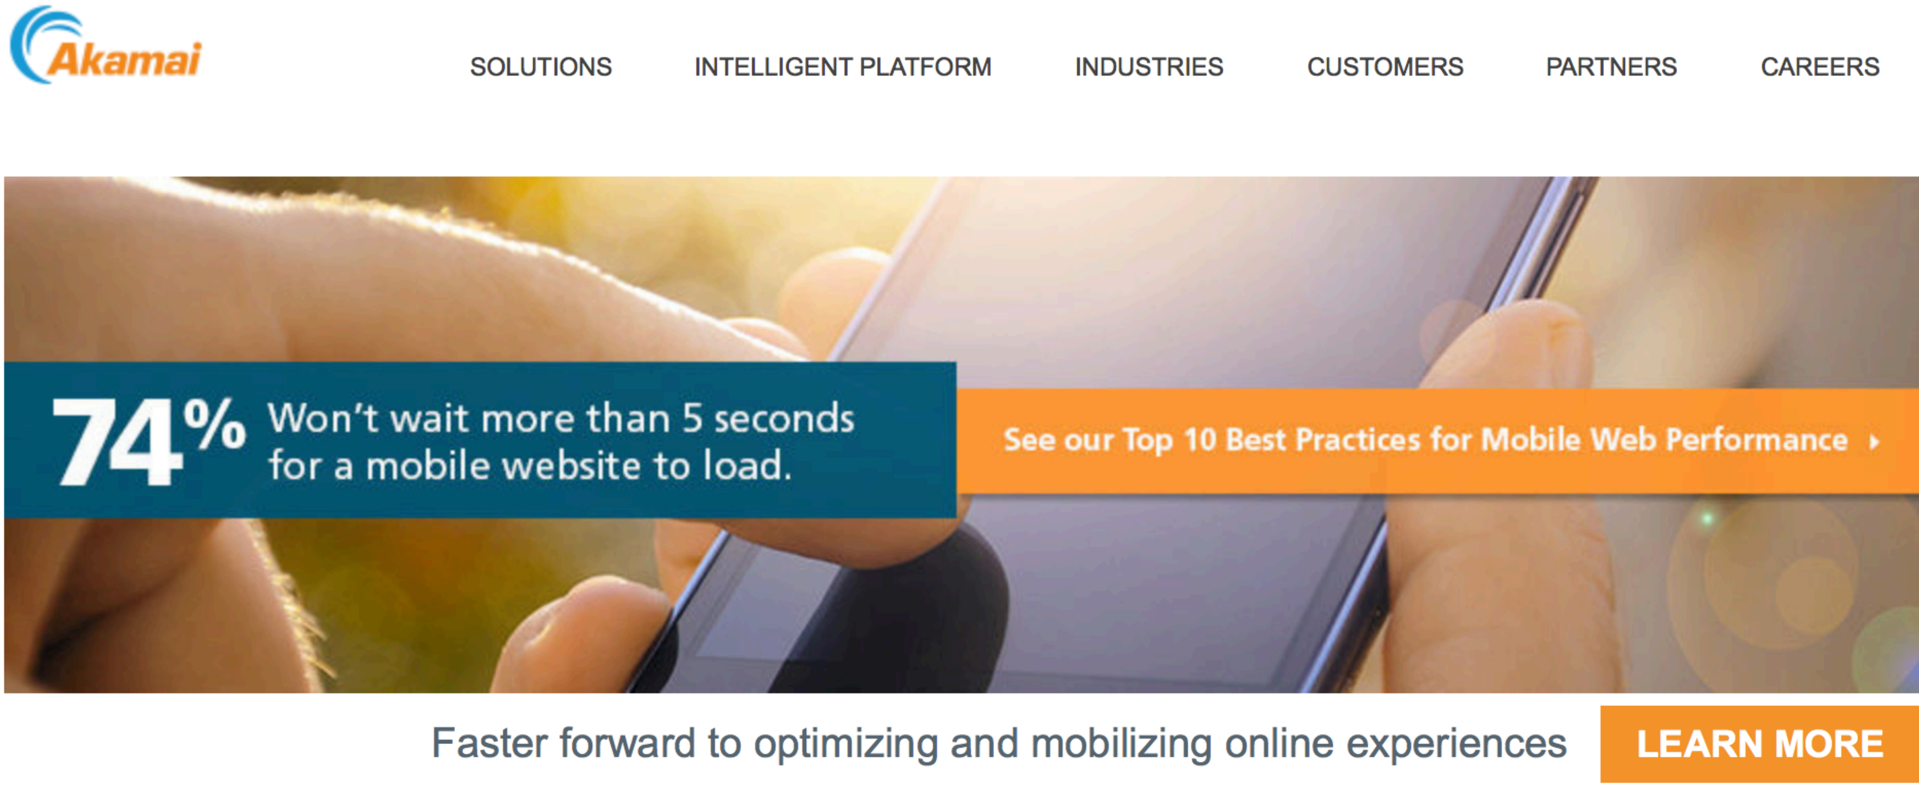

74% Won’t wait more than 5 seconds for a mobile website to load.

See our Top 10 Best Practices for Mobile Web Performance ›

Faster forward to optimizing and mobilizing online experiences

The workings of a CDN

- The DNS provides a different answer depending on where the CDN thinks that the user is located
- In other words different users will get different DNS responses from the same query
- Is this “fragmentation”?

  Well, yes. Different users see different DNS outcomes
- Is this aiding users and content publishers?

  Yes! Bringing content closer to users aids significantly in network efficiency

The workings of a CDN-DNS

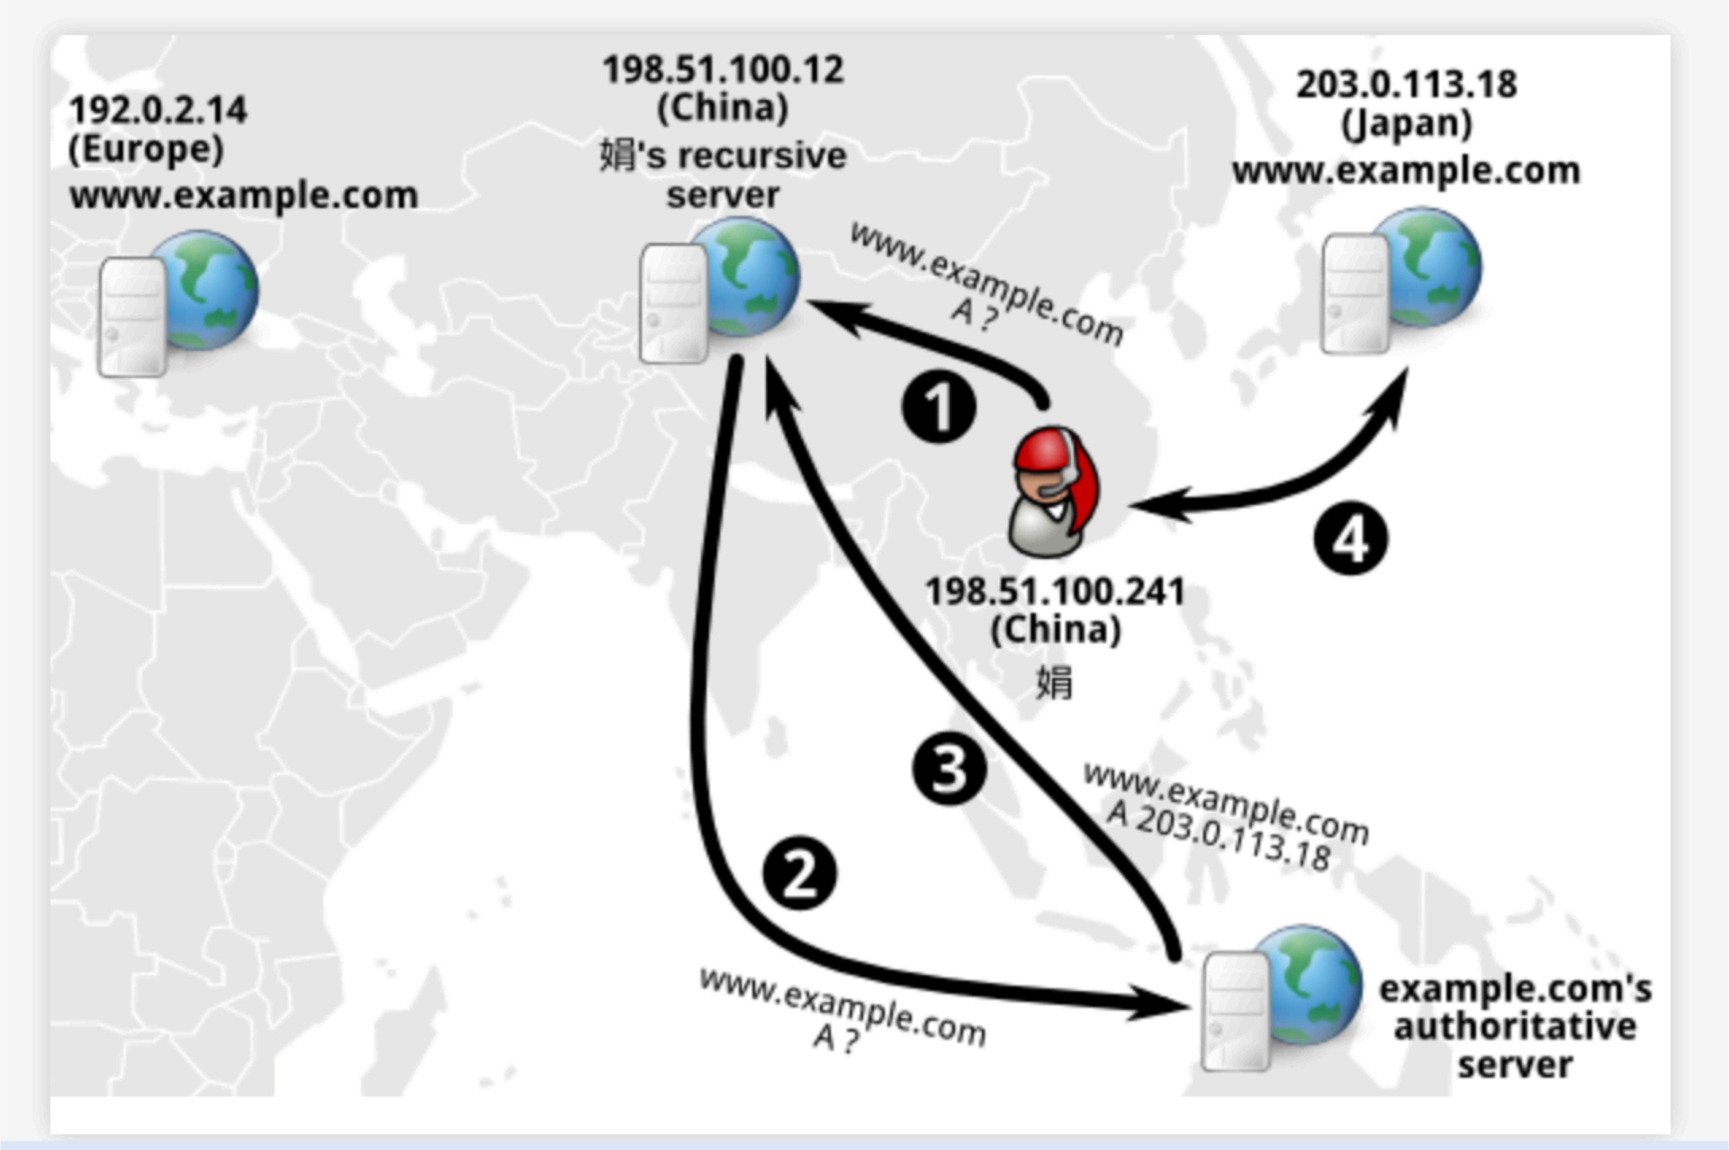

So this is "good" – right?

- Well, not always.
- Because lying in the DNS is still a lie. And accepting lies introduces a new set of security vulnerabilities

User A

```
$ dig +short www.mybank.example.com
203.10.60.10
```

User B

```
$ dig +short www.mybank.example.com
23.10.60.10
```

Which one is genuine?

Telling Good from Bad

What we use today to disambiguate “good” from “bad” when we see such variance is domain name certificates

The conventional response is that user can weed out “bad” DNS responses in a secure (TLS) context by validating the offered key

- Essentially an “out of band” way of saying that a certain named service is secured with a crypto key value which is only known to the domain name holder

And if the name is NOT used in a TLS context?

- You lose! You can’t tell which is the genuine response.

Security to the rescue?

Will our current name security tools allow users to identify attempts to misdirect the user through synthetic name responses?

Is this distributed system of trust adequately secure?

DigiNotar®

A VASCO COMPANY

DigiNotar®
A VASCO COMPANY

Enough said?

No?

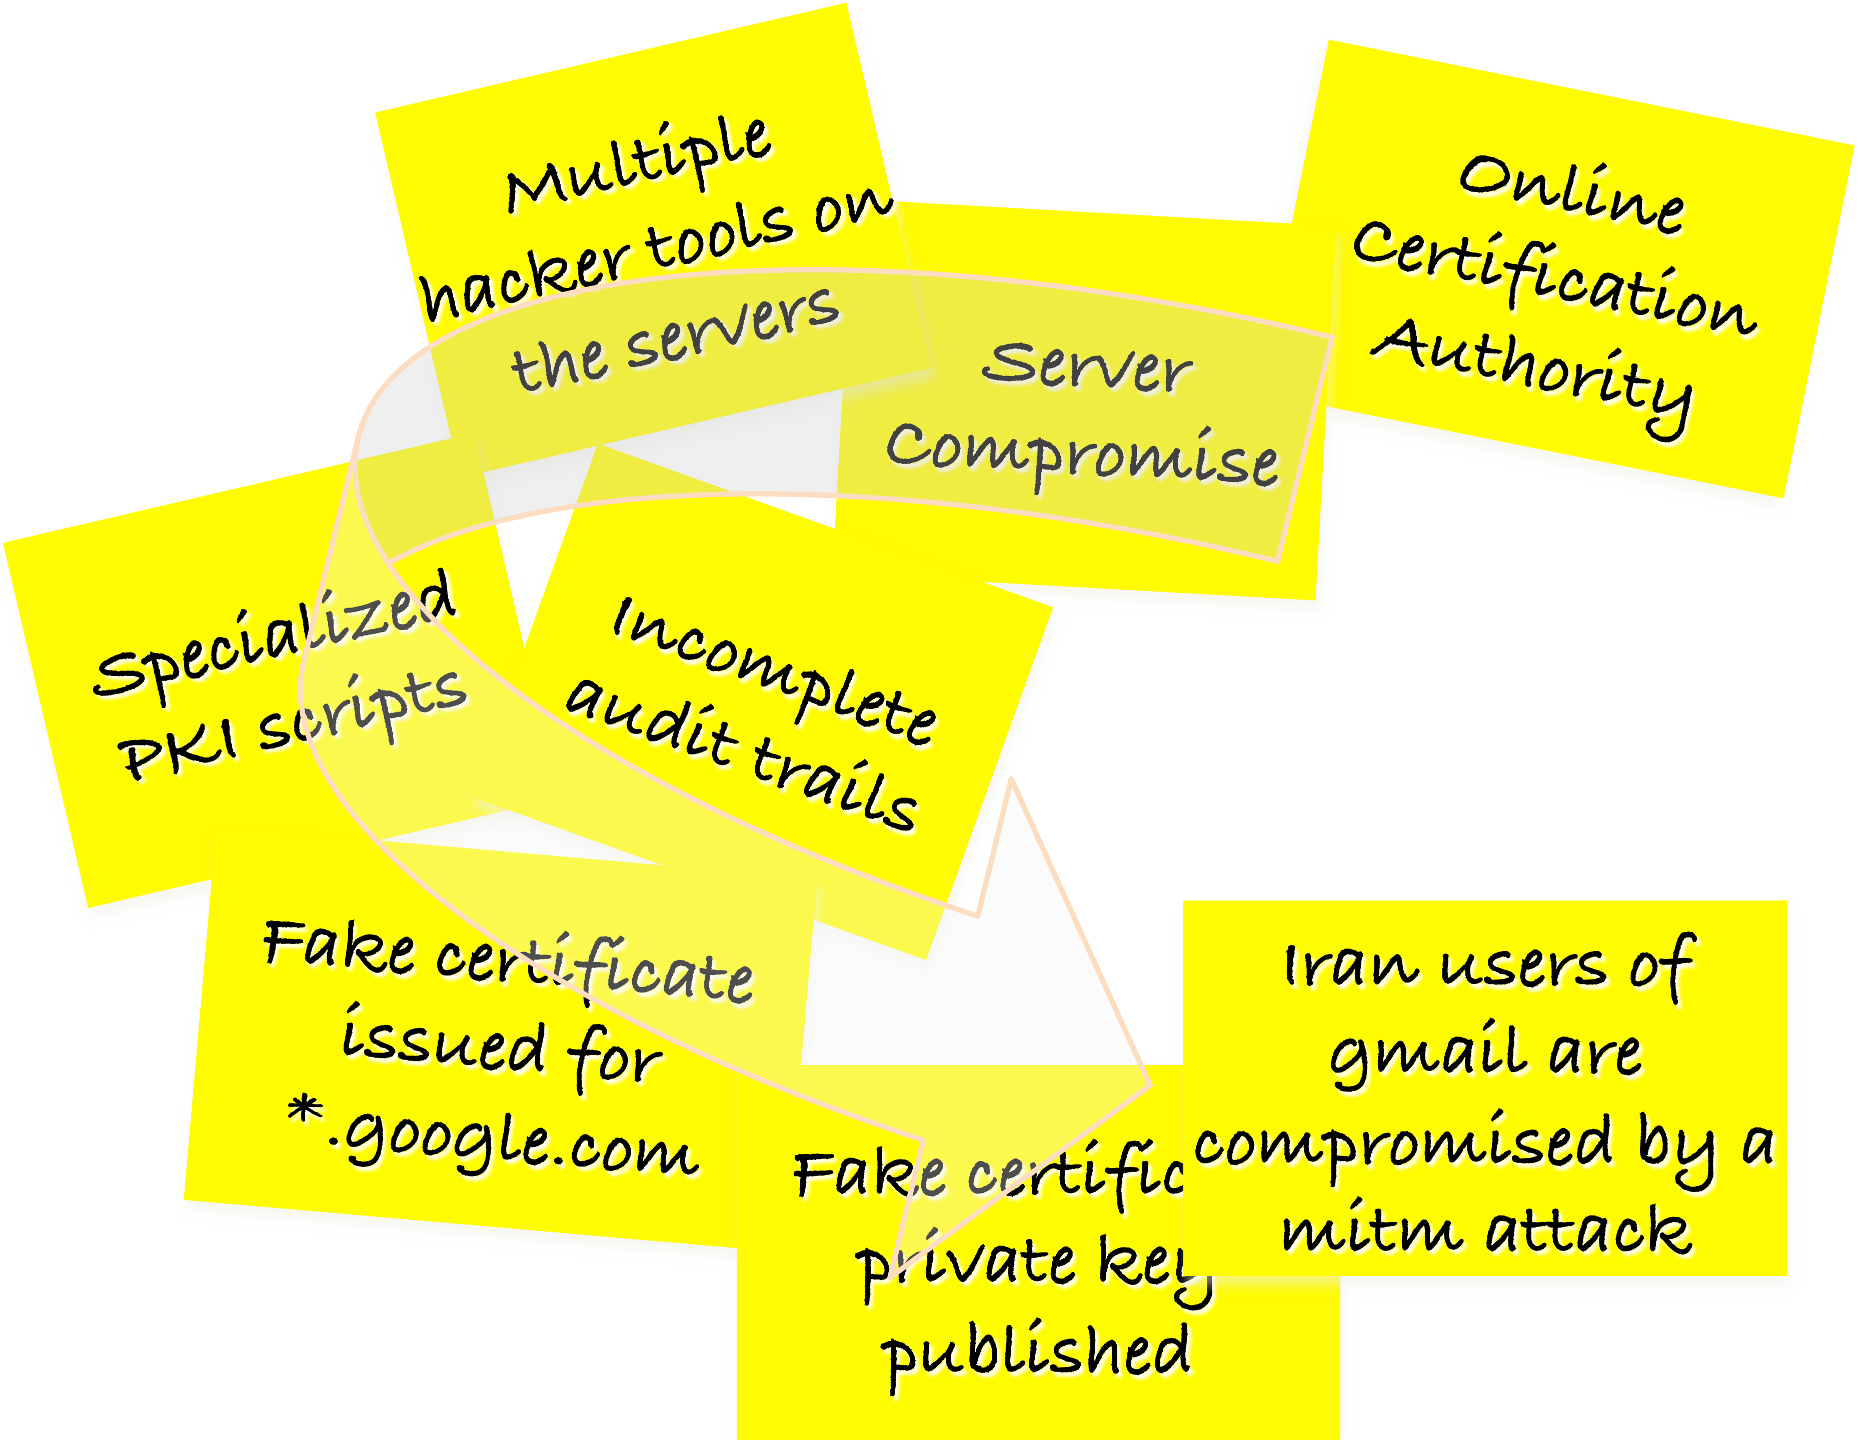

Online Certification Authority

Multiple hacker tools on the servers

Server Compromise

Specialized PKI scripts

Incomplete audit trails

Fake certificate issued for *.google.com

Fake certific private key published

Iran users of gmail are compromised by a mitm attack

Fake certificate issued for *.google.com

+

Fake certificate private key published

=

Any attacker-in-the-middle can intercept a connection request for mail.google.com, and initiate a “secure” connection using the fake certificate, and your browser would be fooled into believing that this was the genuine server!

Two problems:

1. I may not have landed up where I wanted to be:
   - DNS cache poisoning
   - DNS resolver compromise
   - Local host compromise
   - Routing compromise

2. The domain name certificate may be fake

The combination of the two implies that I, and the browser I use, may not even notice that we have been mislead. This is bad.

This is broken economics!

Domain Name certification should use trust and integrity of operation as a differentiator

- If you pay more money you would expect to use a service that operates with greater levels of care and data protection of your data and users of your service would be “more secure” – right?

But a compromised CA can issue a domain name certificate for ANY domain name

- If you trust this compromised CA then you are going to trust its products

The entire Domain Name CA operation is only as good as the worst CA!

- It does not matter what CA service you use, because any compromised CA can compromise users of your service

So what can we do about it?

User A

```
$ dig +dnssec www.mybank.example.com
203.10.60.10
```

✓

User B

```
$ dig +dnssec www.mybank.example.com
SERVFAIL
```

✗

One possible response to inconsistency in DNS responses is to place digital signatures into the DNS – by using DNSSEC

Synthetic DNS responses can be rejected because of the lack of a valid digital signature path

So DNSSEC helps here?

- Yes
  - If we are looking for ways we can identify if the name space is being tampered with by third parties, then DNSSEC can help in identifying when the DNS response is synthetic
- But
  - It does not help with the other purpose of a domain Name Certificate, namely to convey a public key to be used as part of a TLS (secure transport) session
  - For that we also need “DANE”

DANE

DNS-Based Authentication of Named Entities

How to represent and authenticate "named entities" in the DNS, using DNSSEC

Web Sites

Email addres

Jabber IDs

Valid CERTs and/or CAs are stored in the the DNS

If DANE provides the CA's identity, then DANE offers the protection that you are looking at a valid Certificate issued by the CA that performed the EV validation checks in the first place

CA compromise then has limited liability to those certificates issued by the compromised CA

i.e. your service is compromised only if your chosen CA is compromised!

How are we going with DNSSEC?..

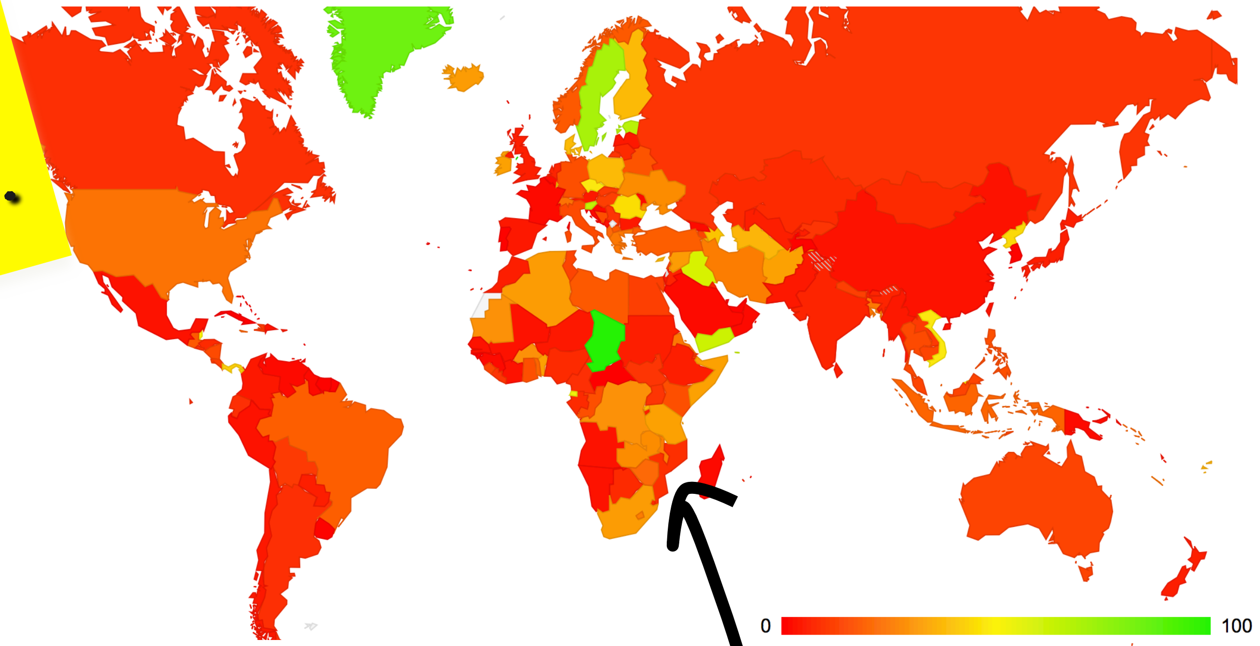

High levels of DNSSEC use seen in Africa, Eastern and Northern Europe

Use of DNSSEC Validation for World (XA)

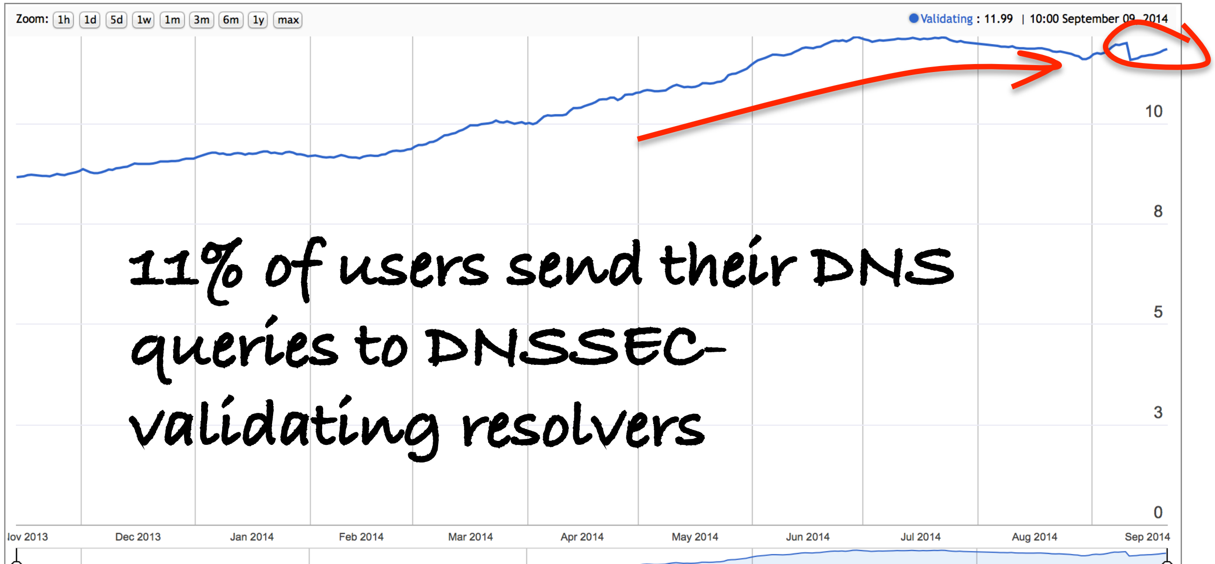

11% of users send their DNS queries to DNSSEC-validating resolvers

What needs to happen?

- The local name management infrastructure should support the use of DNSSEC in all aspects of name management
- ISPs should add DNSSEC validation to their forwarding resolvers
- We should be measuring and reporting on the level of support for DNSSEC validating in both DNS resolvers and DNS names

Name “Fragmentation”

- When a DNS name is resolved into different results it becomes a challenge to distinguish between what’s genuine and what’s a malicious fake
- Adding more trust points to the mix makes the issue worse, not better
- One approach is to use a hierarchical security framework to create an interlocking chain of dependency - DNSSEC
  - On the upside it locks out third party attempts to synthesize fake responses at the edge – the DNS lie cannot be prevented, but it can be exposed
  - On the downside it makes the apex of the name system even more critical, and imposes a common crypto model across the entire Internet

What about Addresses?

If we see continuing pressure for fragmentation and differential behaviours in the Internet's name space, then why don't we see the same fragmentation pressures in the address system?

Addresses are already losing coherence!

In some ways we've already lost that struggle for coherence in addresses – most of the client side of the Internet is addressed behind NATs, and increasing numbers of services are shared on a common platform. Individual IP addresses no longer have a strong association with individual network endpoints.

Addresses these days are just ephemeral conversation tokens without longer term significance

(Aside: Given that we've pushed the Internet into this space where addresses are just ephemeral transaction tokens and the name is the critical point, then will we ever roll back this step with IPv6?)

Where is this heading?

- Are we heading back to the fragmented environment of the 1980s, which a set of islands of networks bridged by arbitrary application level gateways?

I don't believe so – most of the talk of fragmentation remains in the realm of political sabre rattling, and the efforts to restrict, block and redirect end up creating and sustaining novel markets for services that counter such measures, such as retail VPN services, for example

Where is this heading?

- Or is this a more subtle form of fragmentation that operates at the level of discrimination of content delivery system

By leveraging differential responses in the DNS, there is now a specialised market for replicated data delivery services:

- *Akamai, Limelight, Cloudflare, Amazon, Netflix, Google, Apple, Microsoft, DYN, …*

Where is this heading?

Is “Fragmentation” an expression of frustration with the current framework of control of critical infrastructure for the Internet?

Rather than fragmentation, are we witnessing a much grander and far more widespread form of aggregation within the Internet space?

Google Announcement

|< < Prev Random Next > >|

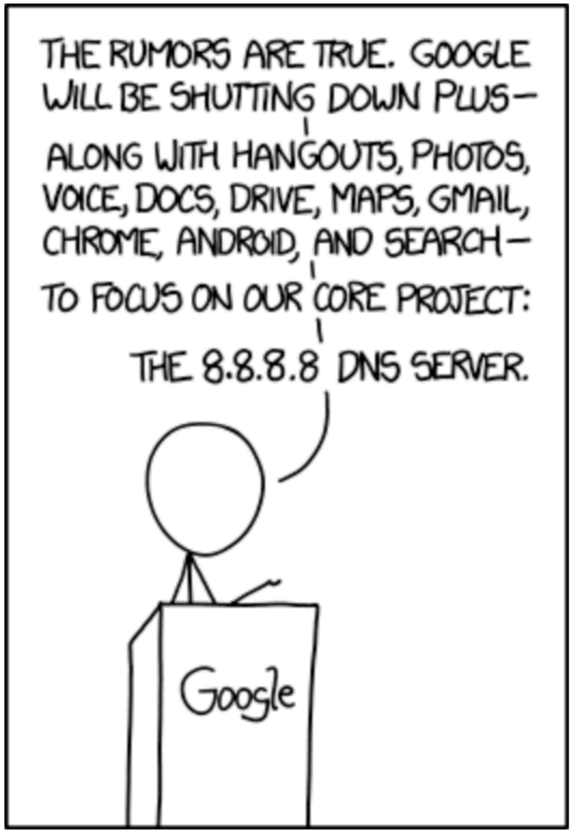
THE RUMORS ARE TRUE. GOOGLE WILL BE SHUTTING DOWN PLUS—
ALONG WITH HANGOUTS, PHOTOS, VOICE, DOCS, DRIVE, MAPS, GMAIL, CHROME, ANDROID, AND SEARCH—
TO FOCUS ON OUR CORE PROJECT:
THE 8.8.8.8 DNS SERVER.
Google

|< < Prev Random Next > >|

Permanent link to this comic: http://xkcd.com/1361/
Image URL (for hotlinking/embedding): http://imgs.xkcd.com/comics/google_announcement.png

Internet-wide Market Share for Google's Public DNS Service

Use of DNSSEC Validation for World

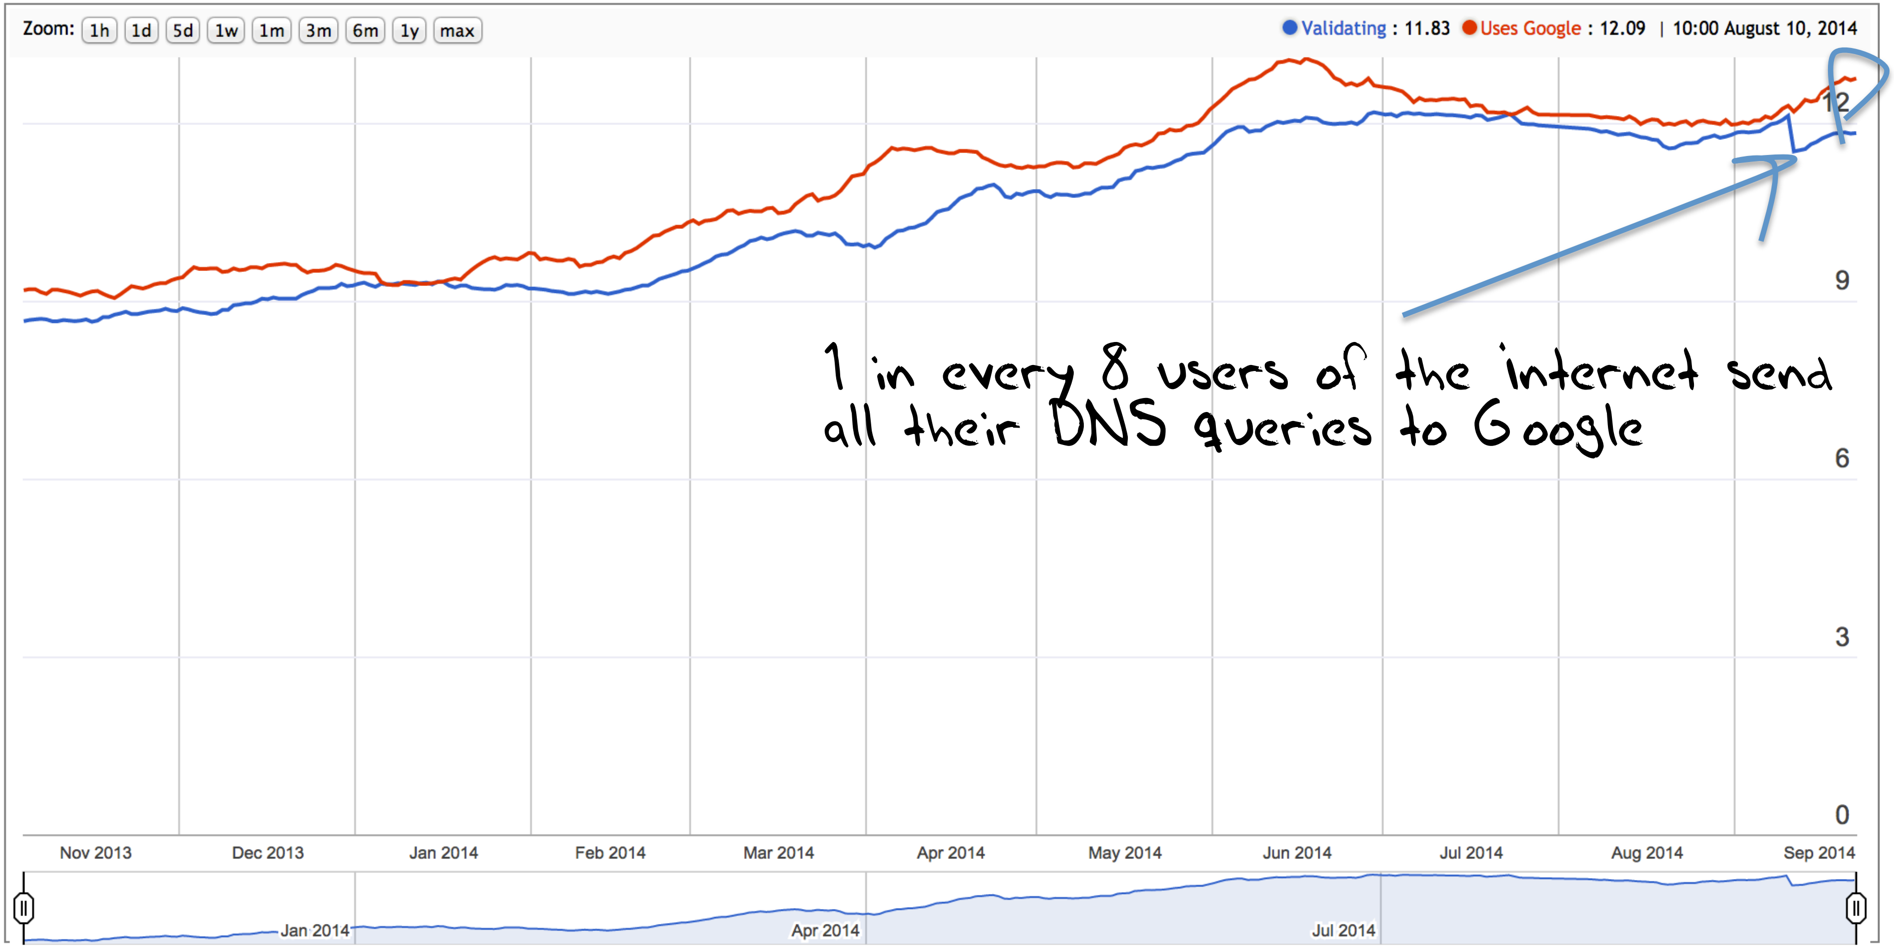

1 in every 8 users of the internet send all their DNS queries to Google

That's it!