Stony Brook University

CSE 361: Web Security

DNS and Web Security

Nick Nikiforakis

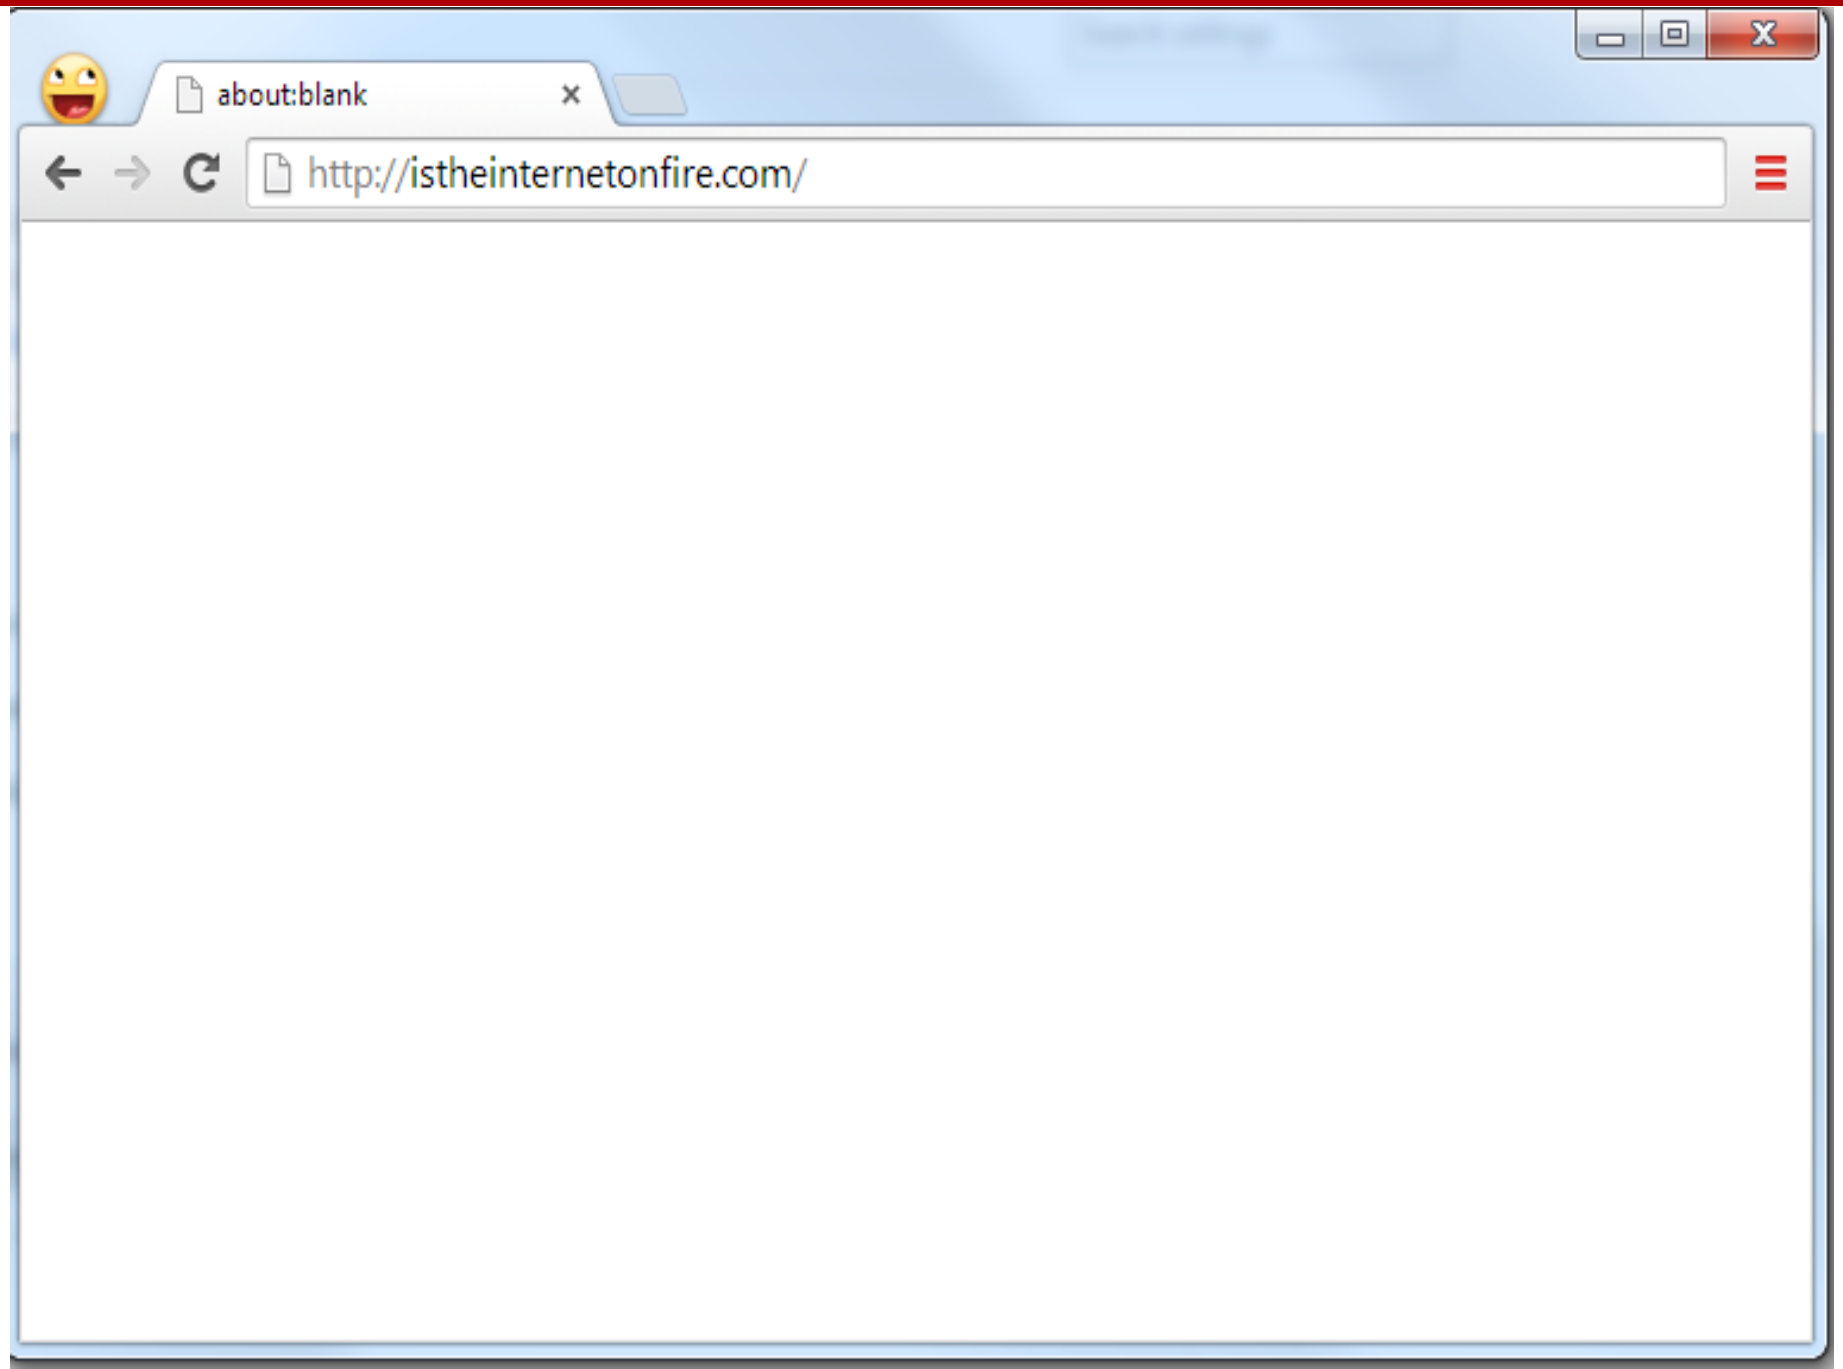
about:blank
http://istheinternetonfire.com/

DNS

- istheinternetonfire.com does not mean anything to a computer
  - So first your browser needs to find the IP address belonging to that domain name

```
nslookup istheinternetonfire.com
Server:         97.107.133.4
Address:        97.107.133.4#53

Non-authoritative answer:
Name:   istheinternetonfire.com
Address: 166.84.7.99
```

How does that work?

- DNS (Domain Name System) works through distributed hierarchical database of DNS servers

- Your computer has what is called a “stub resolver”.
  - This stub resolver does two things:
    - 1. Ask your recursive resolver (typically provided to you by your ISP) to resolve domains for it
    - 2. Remember (cache) the answer of recent queries

How does that work?

- Given that this is the first time you tried to go to this website, your stub resolver asks your network's recursive resolver the same question
  - If another user asked that question recently, your recursive resolver (like your stub resolver) remembers the answer and provides it immediately
  - If not then the recursive resolver ask the root servers
    - Root server == "Gate keepers of worldwide DNS"
    - 13 Root servers distributed across the world managed by various entities
      - E.g. Verisign operates 2 out of the 13 servers

Where are Verisign's root servers?

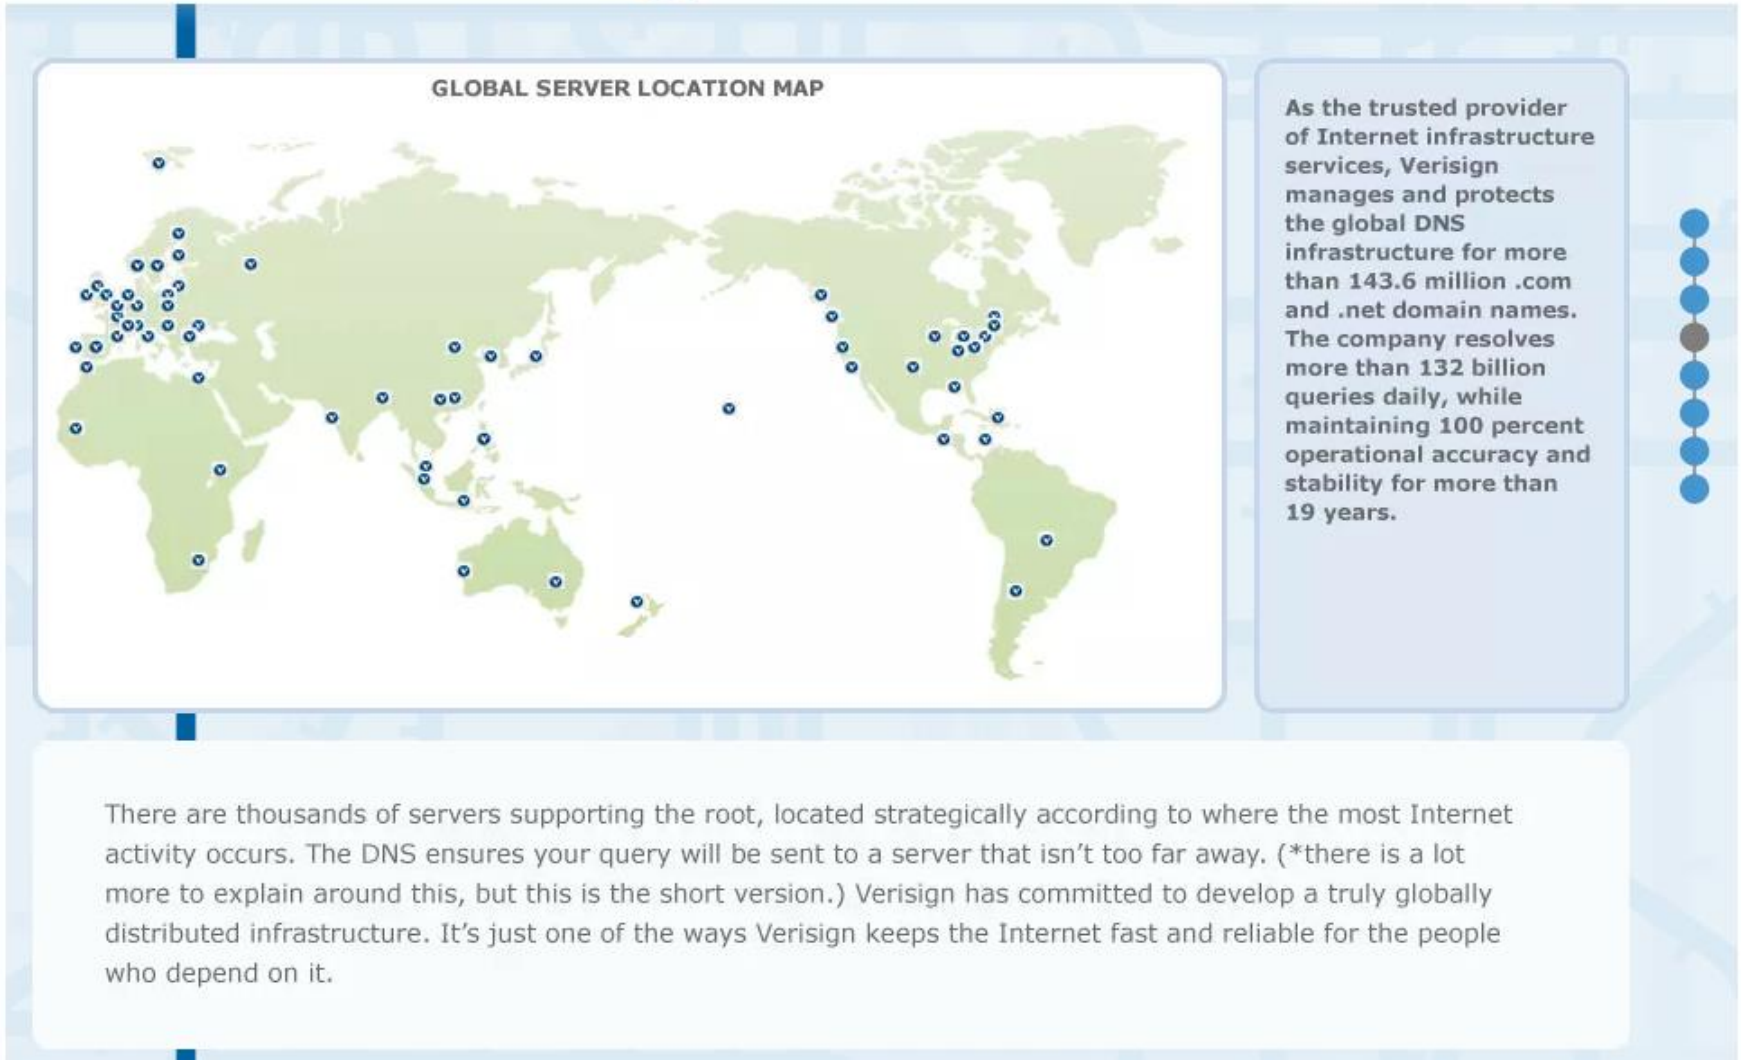

Note: 2 root servers DOES NOT mean two physical machines

Root servers

- The only thing that root servers know, is where the TLD name servers are
  - Servers for .com, .net, .org, etc.
- When your ISP's recursive resolver asks a root server for the address of istheinternetonfire.com the answer is:
  - I don't know, but here is a list of .com nameservers that will probably know

TLD Nameserver

- Q: Hey .com Nameserver, what is the IP address of istheinternetonfire.com ?

- A: I don't know, but go ask the nameservers that are responsible for resolving it, a.dns.gandi.net, b.dns.gandi.net, c.dns.gandi.net
  - Notice that the NS server is located on the .net TLD
  - To save us the trip up to the root and down the .net server, the .com nameserver provides the IP address of the nameserver in its response
    - This is possible because .com and .net are both operated by Verisign

Authoritative Nameserver

- Q: Hey b.dns.gandi.net what is the IP address of istheinternetonfire.com ?

- A:The IP address of istheinternetonfire.com is 166.84.7.99

Now the recursive resolver caches the result and returns the address to your stub resolver running in your operating system

Visually

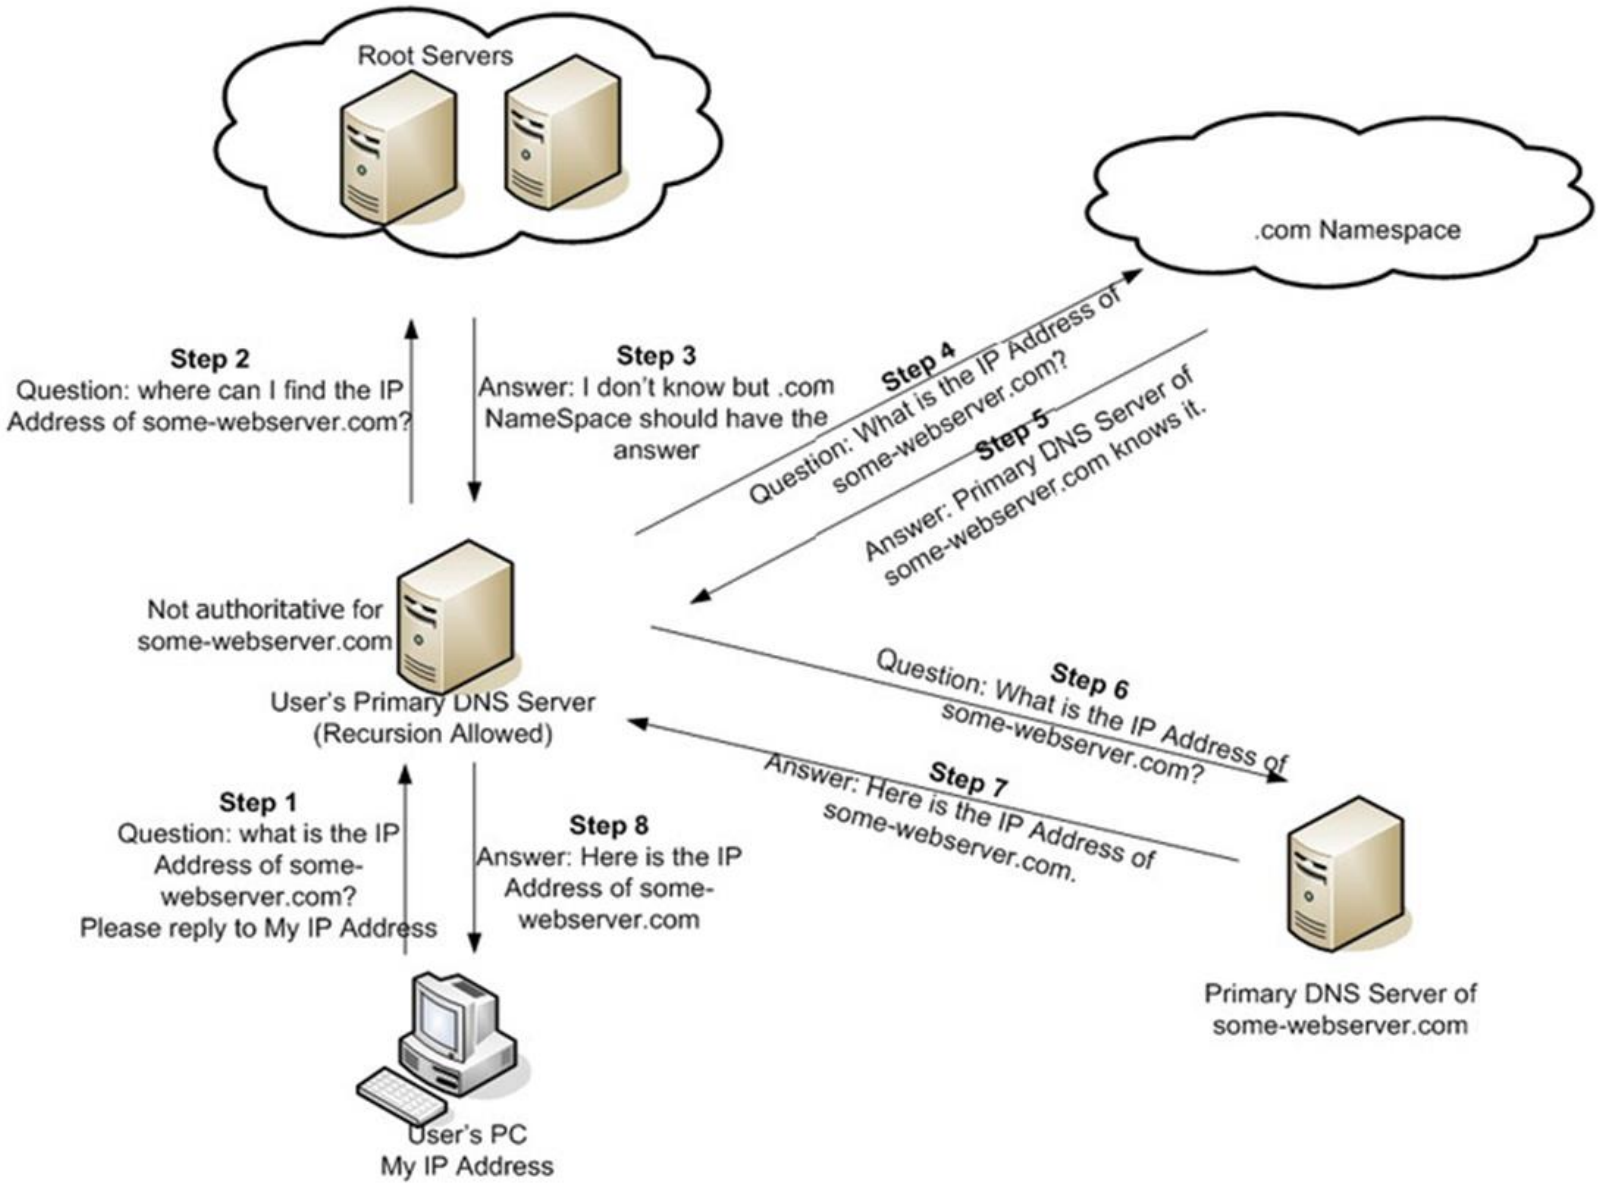
Root Servers
.com Namespace
Step 2
Question: where can I find the IP Address of some-webserver.com?
Step 3
Answer: I don't know but .com NameSpace should have the answer
Step 4
Question: What is the IP Address of some-webserver.com?
Step 5
Answer: Primary DNS Server of some-webserver.com knows it.
Not authoritative for some-webserver.com
User's Primary DNS Server (Recursion Allowed)
Step 6
Question: What is the IP Address of some-webserver.com?
Step 7
Answer: Here is the IP Address of some-webserver.com.
Step 1
Question: what is the IP Address of some-webserver.com? Please reply to My IP Address
Step 8
Answer: Here is the IP Address of some-webserver.com
Primary DNS Server of some-webserver.com
User's PC
My IP Address

DNS Hierarchy (it's a tree!)

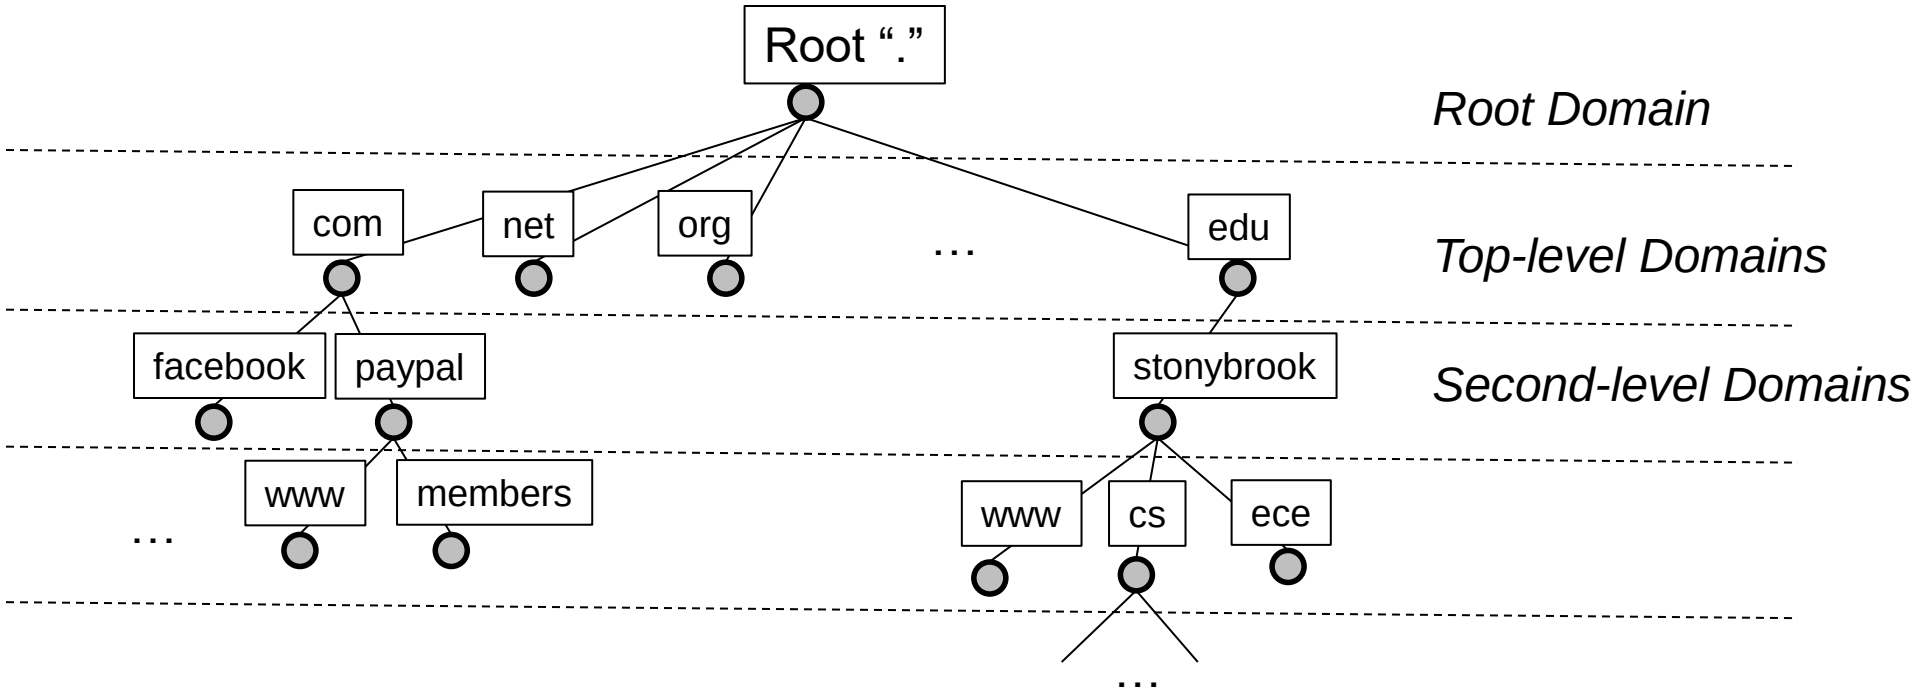

Domains and security

- Domain names are a critical part of web security

- We use domains to:
  - Reference resources on remote servers
    - Scripts, images, stylesheets, objects
  - Make access control decisions
    - Same-Origin Policy (<protocol, host, port>)
  - Configure security mechanisms
    - Allowed domains in CSP
  - Separate different parts of our web application (subdomains)
    - mail.google.com
    - calendar.google.com

Domains and security

- Domain names can
  - Expire
  - Be sold to third parties
  - Be compromised and transfer control to attackers
- What happens to our existing links when all of the above happens?
- Nothing…
  - Our web applications will happily keep resolving domain names and contacting the appropriate servers


PRODUCT
EXPIRED

Expiration of domain names

- Each day, 100K+ domains expire and are returned to the pool of available domains
  - Failed businesses
  - Merging
  - Bad speculating
  - Accidentally

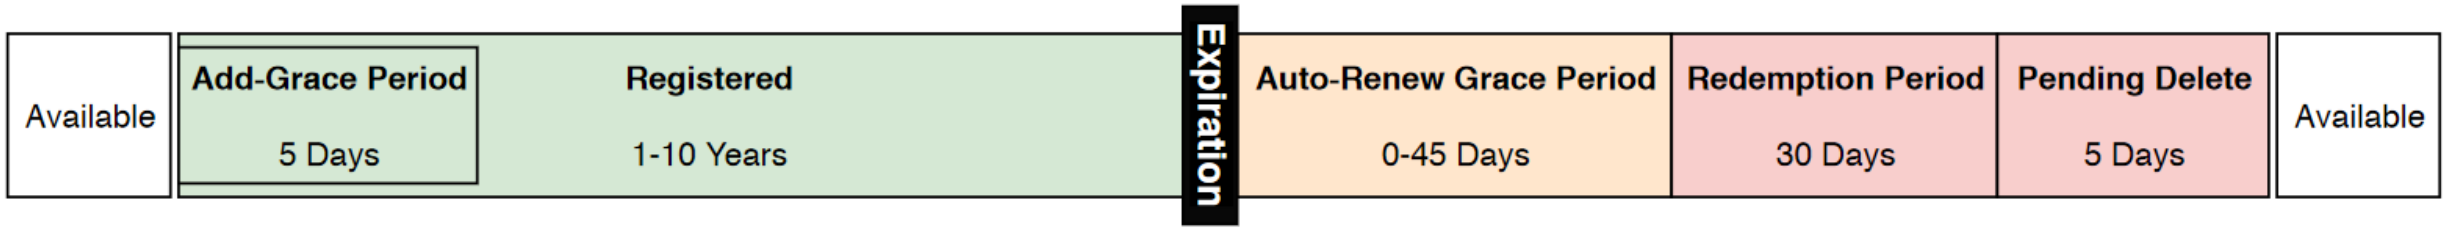

Barron et al. "Now You See It, Now You Don't: A Large-scale Analysis of Early Domain Deletions", RAID 2019

Who buys expiring domains

- Dropcatchers
  - An entire business revolving around identifying attractive domains and re-registering them as fast as possible
    - bikes.com is more valuable than speedy-bikes2019.xyz
    - A site that used to be part of Alexa top 100K is more valuable than one that was never in the top 1M
  - Domains are either then resold or are developed
    - Most players have opportunistic but benign intentions

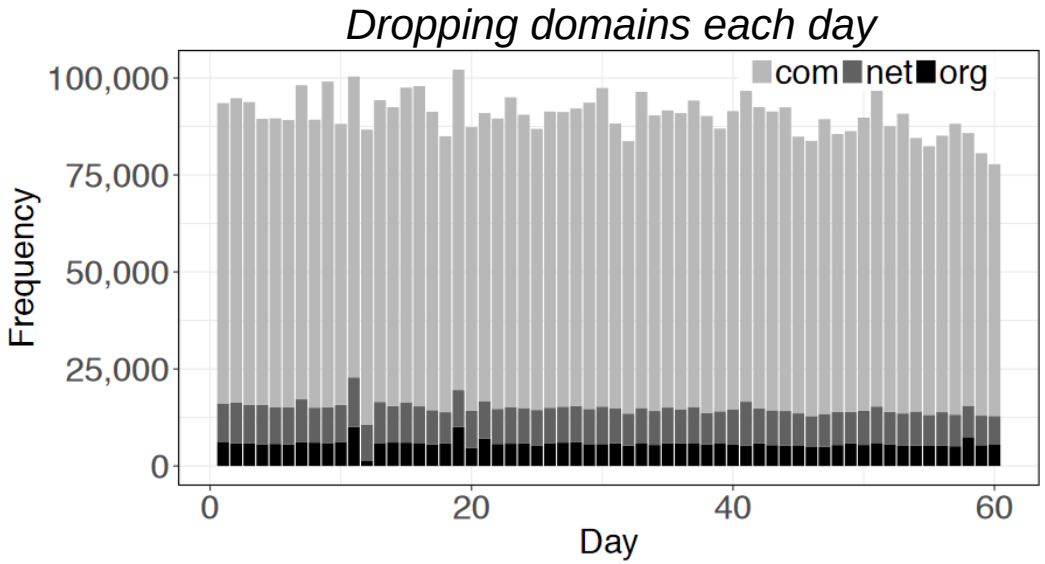

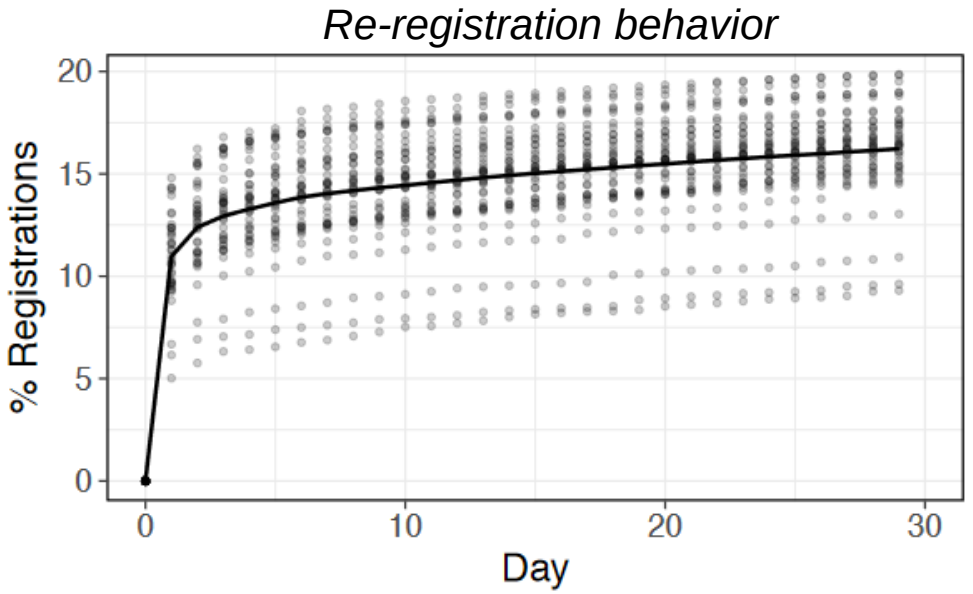

Miramirkhani et al. "Panning for gold.com: Understanding the dynamics of domain dropcatching" WWW 2017

Residual trust

- Potentially sensitive domains are in the hands of
  - New owners who know nothing about their past use
  - New potentially malicious owners who registered these domains
  - No one, just waiting to be rediscovered

- This is called "residual trust" and is straightforwardly abusable by attackers

Residual trust - JavaScript

- In 2012, Nikiforakis et al. discovered that popular websites requested JS from expired domains
  - 56 domains used in 47 sites in the top 10K most popular websites of the Internet
- Attack:
  - Just re-register the domains, and serve scripts where the existing requests expect them to be

	blogtools.us	hbotapadmin.com
Visits	80,466	4,615
Including domains	24	4
Including pages	84	41

Table 5: Results from our experiment on expired remotely-included domains

Intended domain	Actual domain
googlesyndication.com	googlesyndicatio.com
purdue.edu	purude.edu
worldofwarcraft.com	worldofwaircraft.com
lesechos.fr	lessechos.fr
onegrp.com	onegrp.nl

Table 6: Examples of mistyped domains found in remote JavaScript inclusion tags

Nikiforakis et al. "You Are What You Include: Large-scale Evaluation of Remote JavaScript Inclusions" CCS 2012

Residual trust – malicious infrastructure

- In 2016, Lever et al. studied the overlap between malicious operations and expired domains
  - 8.7% overlap between domain blocklists and lists of expired domains
    - Attackers weaponizing known-good domains
- Presented examples of residual trust in
  - Browser extensions
  - Name servers
  - Email servers

Lever et al. "Domain-z: 28 Registrations Later: Measuring the exploitation of residual trust in domains", IEEE S&P 2016

Residual Trust - CSP

- In 2020, Roth et al. investigated the evolution of CSP policies over the years

- One of the experiments was regarding trusted domain names in CSP policies
  - 41 cases of domains that could be abused due to residual trust, typos, and local resolution

Category	Vulnerable domains	Duration	Impacted domains
Expired	16		15
Example	*sushissl.com*	39 days	*zomato.com*
Typo	11		11
Example	*optmster.com*	7 months	*experian.com*
Local address	15		26
Example	*marketo.net*	3 months	*dropbox.com*
Total	41		50

TABLE II: *Vulnerable whitelisted domains and the number of sites that allowed these domains in their whitelists. One example for each category with a high-profile site that included it and duration of attack opportunity.*

Roth et al. "Complex Security Policy? A Longitudinal Analysis of Deployed Content Security Policies", NDSS 2020

No shortage of real-world examples

200k+ Parked/Expired Domains Used to Distribute Malicious Ads

JUNE 29, 2016 DOUGLAS SANTOS

The .io Error – Taking Control of All .io Domains With a Targeted Registration

Related attacks – Dangling resolutions

- Keeping control of the domain but losing control of the resolving IP address

- Scenario
  - Create subdomain for your site (testdev.example.com)
  - Spin up VM in public cloud and assign testdev.example.com to IP address of VM
  - At a later point, abandon project, delete VM, while forgetting to delete the NS record
  - Attackers discover the "danging" DNS record and keep asking for VMs until they get the one that's yours

Related attacks – Dangling resolutions

- In 2018, Borgolte et al. measured the vulnerability of domains to this attack
  - Analyzed 130M domains resolving to public clouds
  - Identified 700K dangling cases (i.e. pointing to IP addresses that were "free")
  - Calculated that attackers need about $1 to cycle through IP addresses until they find the right one

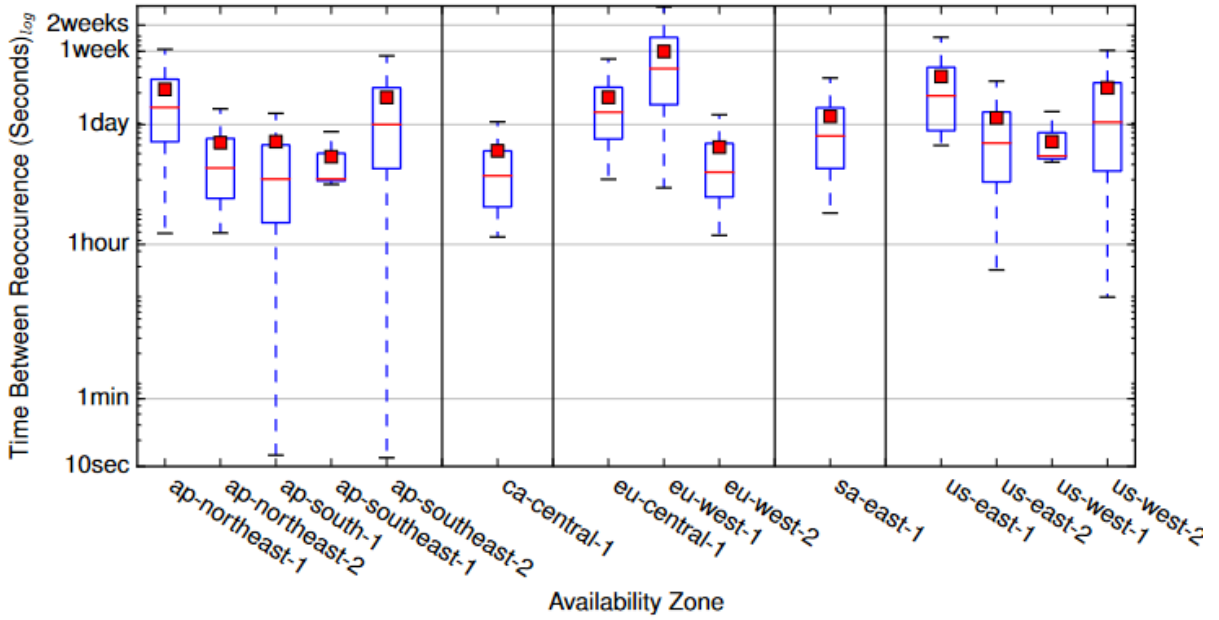

Borgolte et al. "Cloud Strife: Mitigating the Security Risks of Domain-Validated Certificates", NDSS 2018

Defenses

- Asset management
  - Cataloguing all the external dependencies of a web application
  - Searching for outages and anomalies
    - Most domains will stop resolving long before they switch hands
    - Monitor the resolution failures of your infrastructure
- Integrity verification
  - SRI for static JS resources
    - Not always possible but valuable when it is

Summary

DNS Hierarchy (it's a tree!)

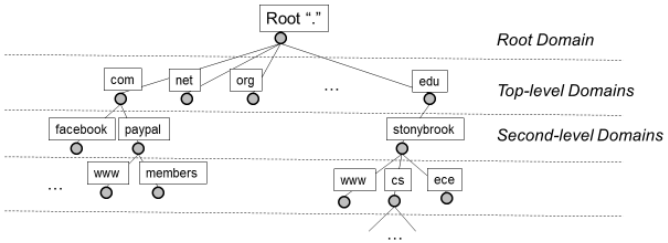

Residual trust - JavaScript

- In 2012, Nikiforakis et al. discovered that popular websites requested JS from expired domains
  - 56 domains used in 47 sites in the top 10K most popular websites of the Internet
- Attack:
  - Just re-register the domains, and serve scripts where the existing requests expect them to be

	blogtools.us	hbotapadmin.com
Visits	80,466	4,615
Including domains	24	4
Including pages	84	41

Table 5: Results from our experiment on expired remotely-included domains

Intended domain	Actual domain
googlesyndication.com	googlesyndicatio.com
purdue.edu	purude.edu
worldofwarcraft.com	worldofwaircraft.com
lesechos.fr	legsechos.fr
onegrp.com	onegrp.nl

Table 6: Examples of mistyped domains found in remote JavaScript inclusion tags

Nikiforakis et al. "You Are What You Include: Large-scale Evaluation of Remote JavaScript Inclusions" CCS 2012

Expiration of domain names

- Each day, 100K+ domains expire and are returned to the pool of available domains
  - Failed businesses
  - Merging
  - Bad speculating
  - Accidentally

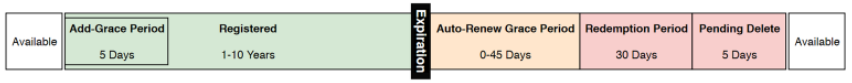

Barron et al. "Now You See It, Now You Don't: A Large-scale Analysis of Early Domain Deletions", RAID 2019

Related attacks – Dangling resolutions

- In 2018, Borgolte et al. measured the vulnerability of domains to this attack
  - Analyzed 130M domains resolving to public clouds
  - Identified 700K dangling cases (i.e. pointing to IP addresses that were "free")
  - Calculated that attackers need about $1 to cycle through IP addresses until they find the right one

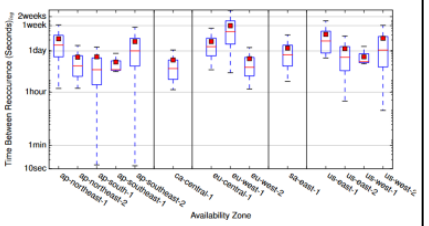

Borgolte et al. "Cloud Strife: Mitigating the Security Risks of Domain-Validated Certificates", NDSS 2018

Credits

- Original slide deck by Nick Nikiforakis