In search of tomorrow's cybersecurity protections, building on today's foundations

Eric Osterweil
eoster@gmu.edu

A brief glimpse of our threatscape

DDoS attack that disrupted internet was

BBC Sign in Menu

World US & Canada UK Business Tech Science Stories

anada

er-attack: US and UK blame North
a for WannaCry

nber 2017 Share

[Not-so] Surreptitious MitM BGP attack
--Google '18 (this week!)

oto used against us as ransomware
annaCry '17

MASON UNIVERSITY

We can see the roots of these attacks

- Largest DDoS attacks ← source-address spoofing (lies)
- Data exfiltration ← data at-rest not protected
- Compromised systems ← malcode reactively detected
- IoT ← hemorrhaging vulnerable systems
- Man-in-the-Middle (MitM) attacks ← actively being *sought out*

- We have not been able to operationalize protections/remediations

Good news: Long sought after capabilities *now exist in* the Internet's core, and it is poised to be the foundation cybersecurity needs!

GEORGE MASON UNIVERSITY

Getting started with protections we need

- Key building blocks
  - **Resource Certification**
  - **Object-level encryption**
  - **Provenance**

- Advanced protections can be built on our distributed **core** architecture
  - We can ***operationalize*** secure handshakes, secure objects, etc. and ***measure*** efficacy

The science of cybersecurity:
We must analyze and *learn* what security properties work *operationally* and *at scale*

GEORGE MASON UNIVERSITY

Outline

- Background: cybersecurity in the Internet's core
- Current research challenges
- Future directions

GEORGE MASON UNIVERSITY

Background: the status of the Internet's core

- Internet's core: inter-domain **routing** and **naming** services
  - The Border Gateway Protocol (BGP)
  - The Domain Name System (DNS)
- Like the Internet's Interstate Highway system
  - … and do any of us ever stop to think about what **I-95's** asphalt is made of?

- Historical, no security protections, but *now* they are getting them now
  - BGP: The Resource Public Key Infrastructure (**RPKI**)
  - DNS: The DNS Security Extensions (**DNSSEC**)

Focusing on DNSSEC:
Capable of being a foundational substrate that will enable *other* security needs!

GEORGE MASON UNIVERSITY

DNS Security Extensions (DNSSEC)

- DNSSEC = DNS' namespace
  + Internet-scale crypto

- Over 13 years + **millions** of zones
  [SecSpider.net]

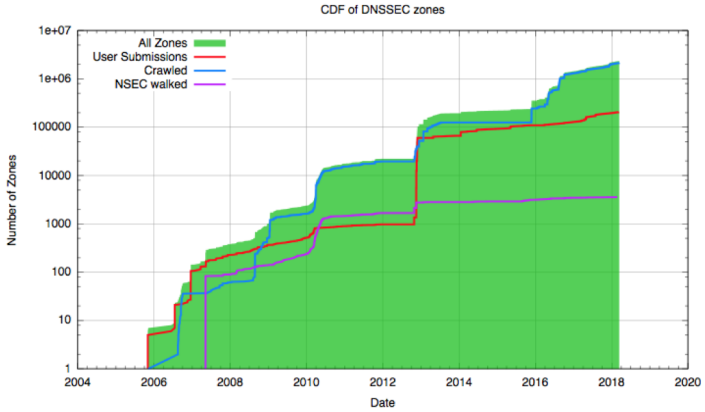

- First **core** Internet protocol to secure with crypto

- Zones sign all RRsets (i.e. **extensible data**) and resolvers use DNSKEYs to verify them

GEORGE MASON UNIVERSITY

DNSSEC in a nutshell

- DNSSEC zones create public/private keys
- Uses a **single root** key (Root KSK)
- A*ll* verification flows recursively down the hierarchy

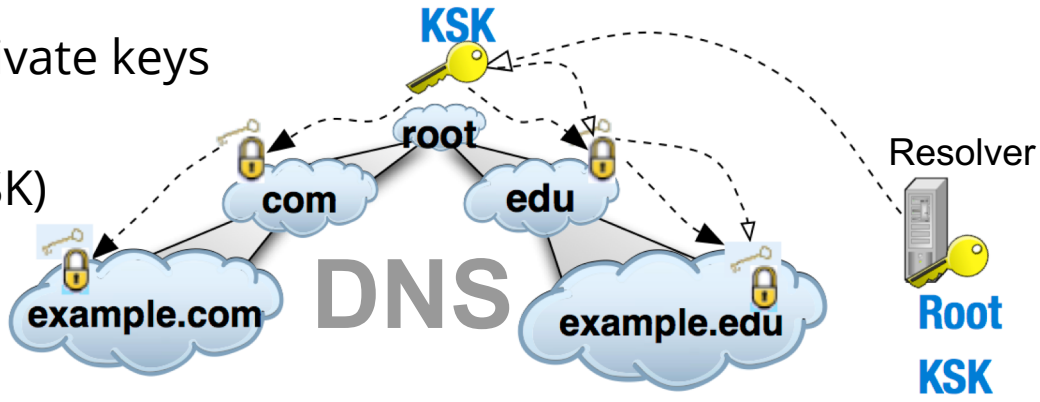

- An operational Public Key Infrastructure (PKI) at scale not previously possible
- Other systems have tried to create a **single root** key, but have failed
  - **PEM, Web PKI**, **RPKI**, etc.

MASON UNIVERSITY

DNSSEC data signing example

Using a zone's key on a standard RRset (the NS)

Signature (RRSIG) is only verifiable by the DNSKEY if *no* data was modified

GEORGE MASON UNIVERSITY

DNSSEC has already taught lessons

- Simple design has taught **many lessons**
  - Replay vulns [SecSpider-NPSec-07]
  - PMTU issues [SecSpider-IMC-08, SecSpider-ACSAC-09]
  - key rollovers
- Unforeseen **design principles**
  - **Security** from **untrusted** servers [SysDeps-NPSec-14]
- This operational PKI has taught novel lessons

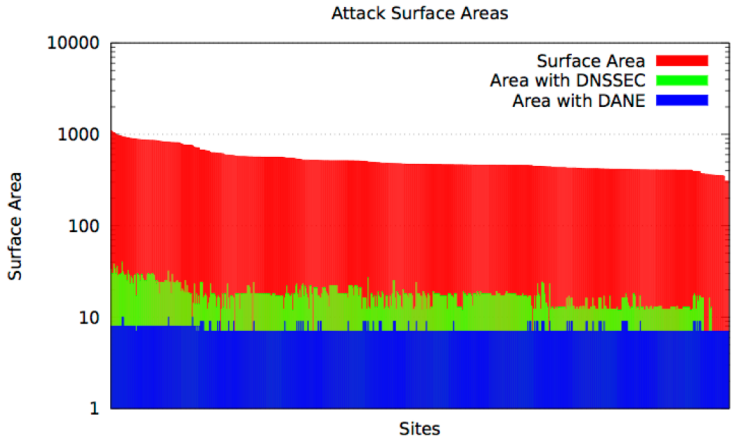

But, what is DNSSEC really?
Not just a PKI, but a globally distributed (now secure) general database

MASON UNIVERSITY

Outline

- Background: cybersecurity in the Internet's core
- Current research challenges
- Future directions

GEORGE MASON UNIVERSITY

Current research: is DNSSEC ready to be built on?

- Have operators learned to manage crypto?
- Is it well suited to **build on**: design, mature operations, and scale?

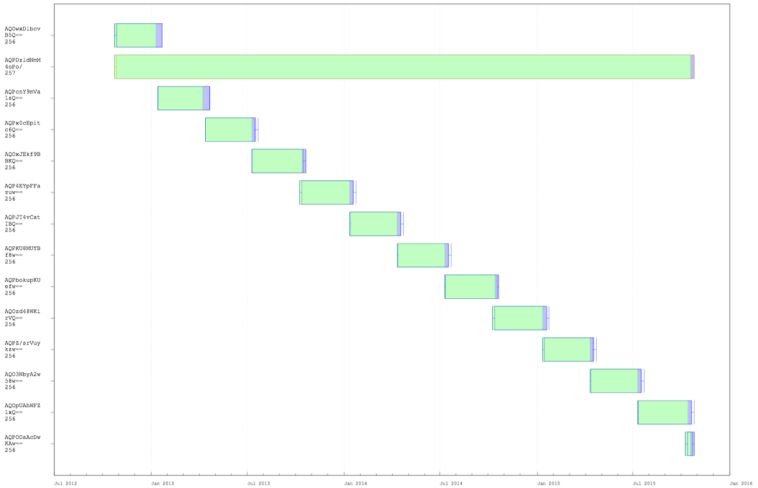

- Reusable new protocols: DNS-based Authentication of Named Entities (DANE)
  - Crypto key learning using domain names (**S/MIME** for email, etc.)
  - **Communities** forming: IETF, NIST/NCCoE, LIGHTest
  - Resource certification, object-security and provenance

GEORGE MASON UNIVERSITY

Example, Internet number resource certification

- Consider 2016's Dyn DDoS: largest ever
  - Very complex MIRAI botnet -> massive spoofed packets
  - No **BCP-38/BCP-84** : but that's an **open research challenge**
- Ops don't have this today!

- Resource Certification: ending spoofing
  - Could've mitigated attack *at origins* by BCP-38/BCP-84 filters
- *Many open questions:*
  - Policies to encode? Enforcement?

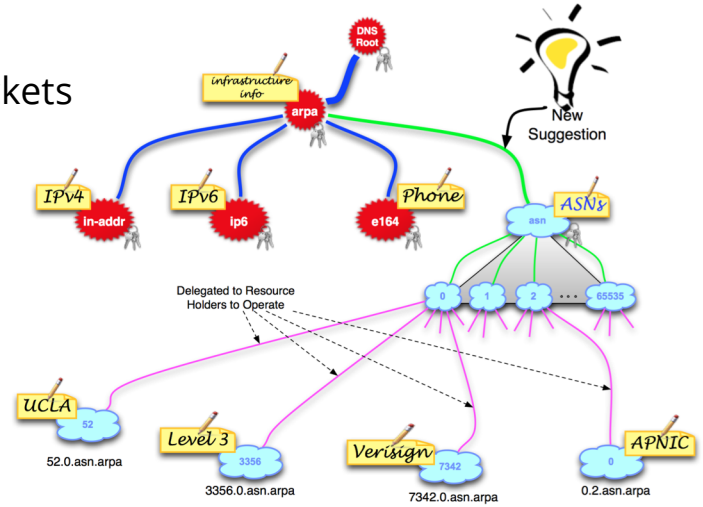

Also consider, object-security for data at rest

- In 2014, Sony's email was **exfiltrated** (unencrypted at rest)
- Why was data unprotected?
  - Existing techniques can't securely learn keys
- DANE securely learns keys (deployable today)
  - Email addrs are domain names + keys from DANE
- Still leaves *many* open challenges:
  - Email key lifecycle management, archival vs. online keys, and likely many more we don't yet know to ask

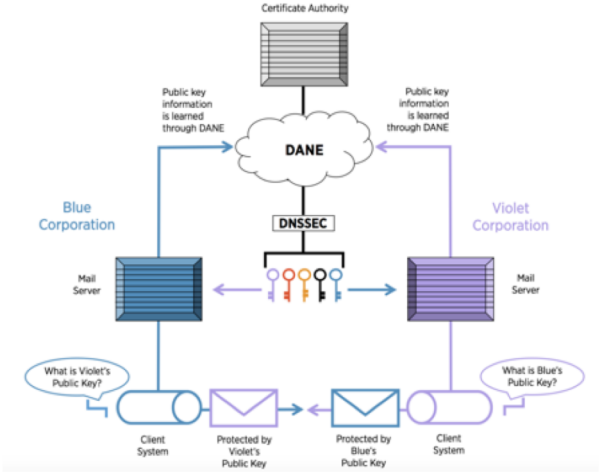

https://nccoe.nist.gov/sites/default/files/library/fact-sheets/dns-secure-email-fact-sheet.pdf

Is DNS' namespace ready? the Devil's in the details

- These solutions raise the stakes, which raises deeper research questions
  - Approaches like DANE do put a heavy onus on using DNSSEC's namespace
- Cybersecurity substrate will need closer examination
  - As long as that namespace is managed properly and used correctly, we gain its benefits
- In DNS, domain names don't collide: each has a single authority
  - Only Google resolves google.com
- However, while the global DNS' namespace is collision-free, other namespaces can lead to **name collisions** within the DNS

GEORGE MASON UNIVERSITY

Name collisions [Oakland-16, US-CERT-16]

- *A name collision occurs whenever a name is resolved in a namespace other than where it was expected to be resolved*

- Example, a corp expects free reign choosing systems' names
- Uses a fictitious Top-Level Domain (TLD) "`.network`"
  - `mail.network`
  - `www.network`
  - `ad.network` (an ActiveDirectory server)

Why should they **not** be surprised when *all* of their online credentials stolen, and all of their online transactions intercepted?

Things used to seem easy…

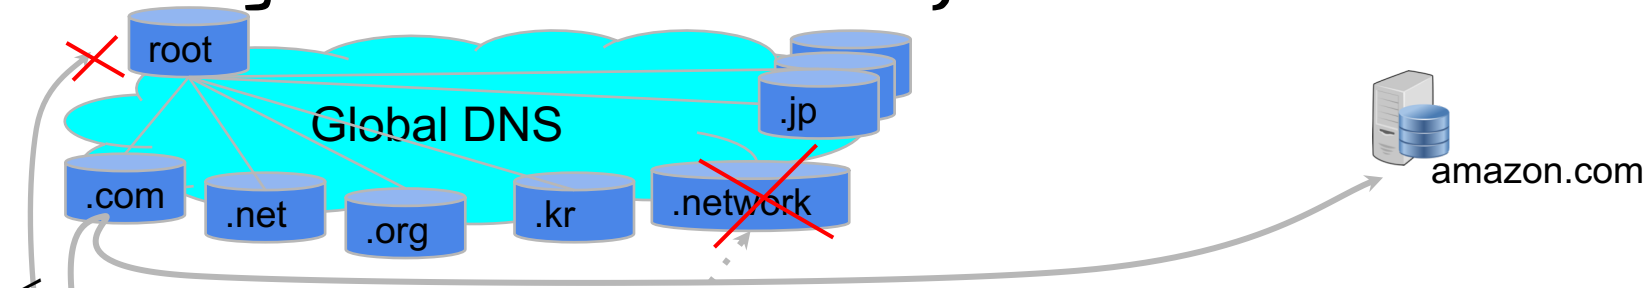

www.network (?)

If it's not a TLD, then queries to the global root will return non-existent domain (NXDomain)

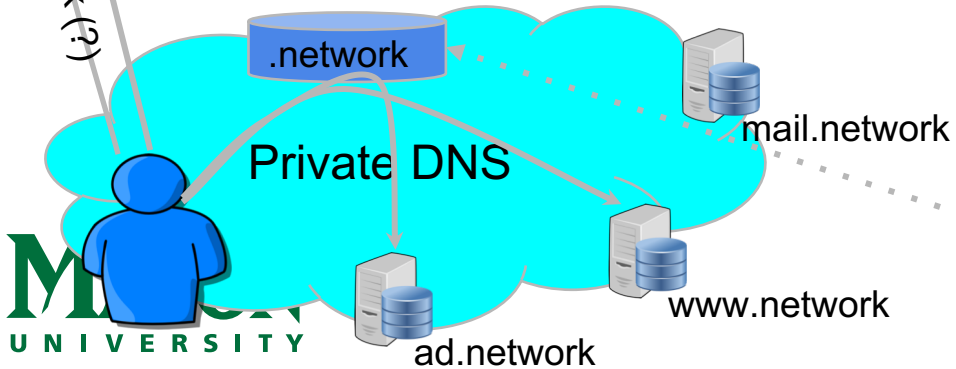

In this case, .network is deployed as an internal Top-Level Domain (iTLD)

And its services are resolved in the internal network

UNIVERSITY

But, since 2013, we have almost 1k new TLDs

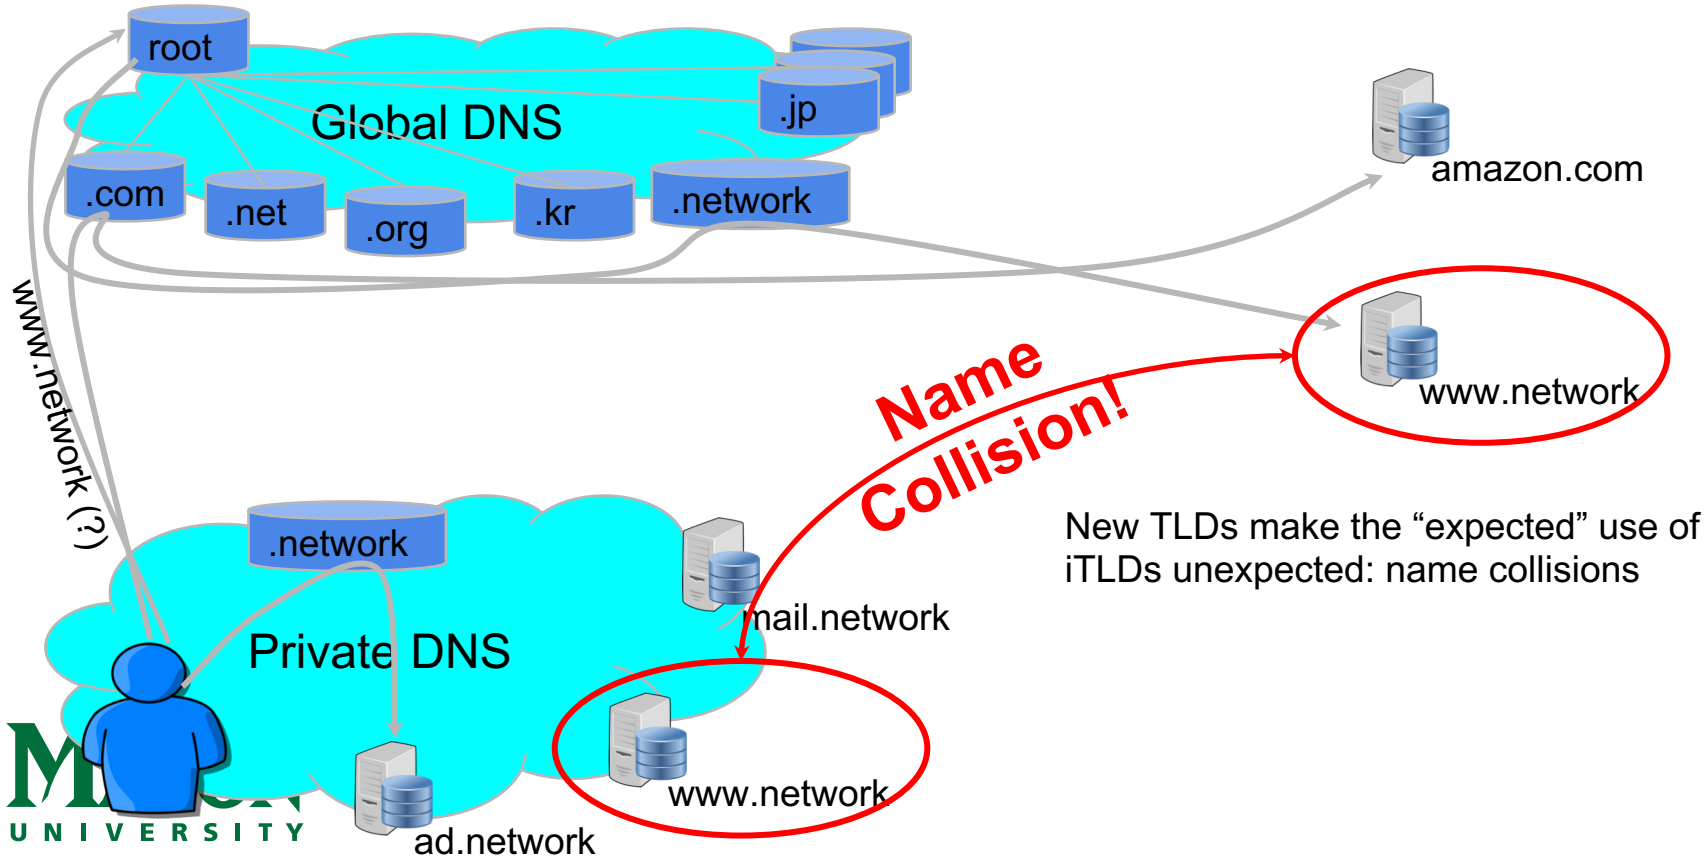

Name collisions can expose users to MitM

- Getting the wrong server is bad, but with DNS-based Service Discovery (**DNS-SD**) protocols, it can be much worse

- WPAD automatically sets up a web proxy for browsers *surreptitiously*
  - Clients look-up "`wpad.<configured-home-domain>`" and **trust** whatever comes back!

- After that first name collision, **all web traffic is owned** from that point on

GEORGE MASON UNIVERSITY

MitM from name collisions: WPAD

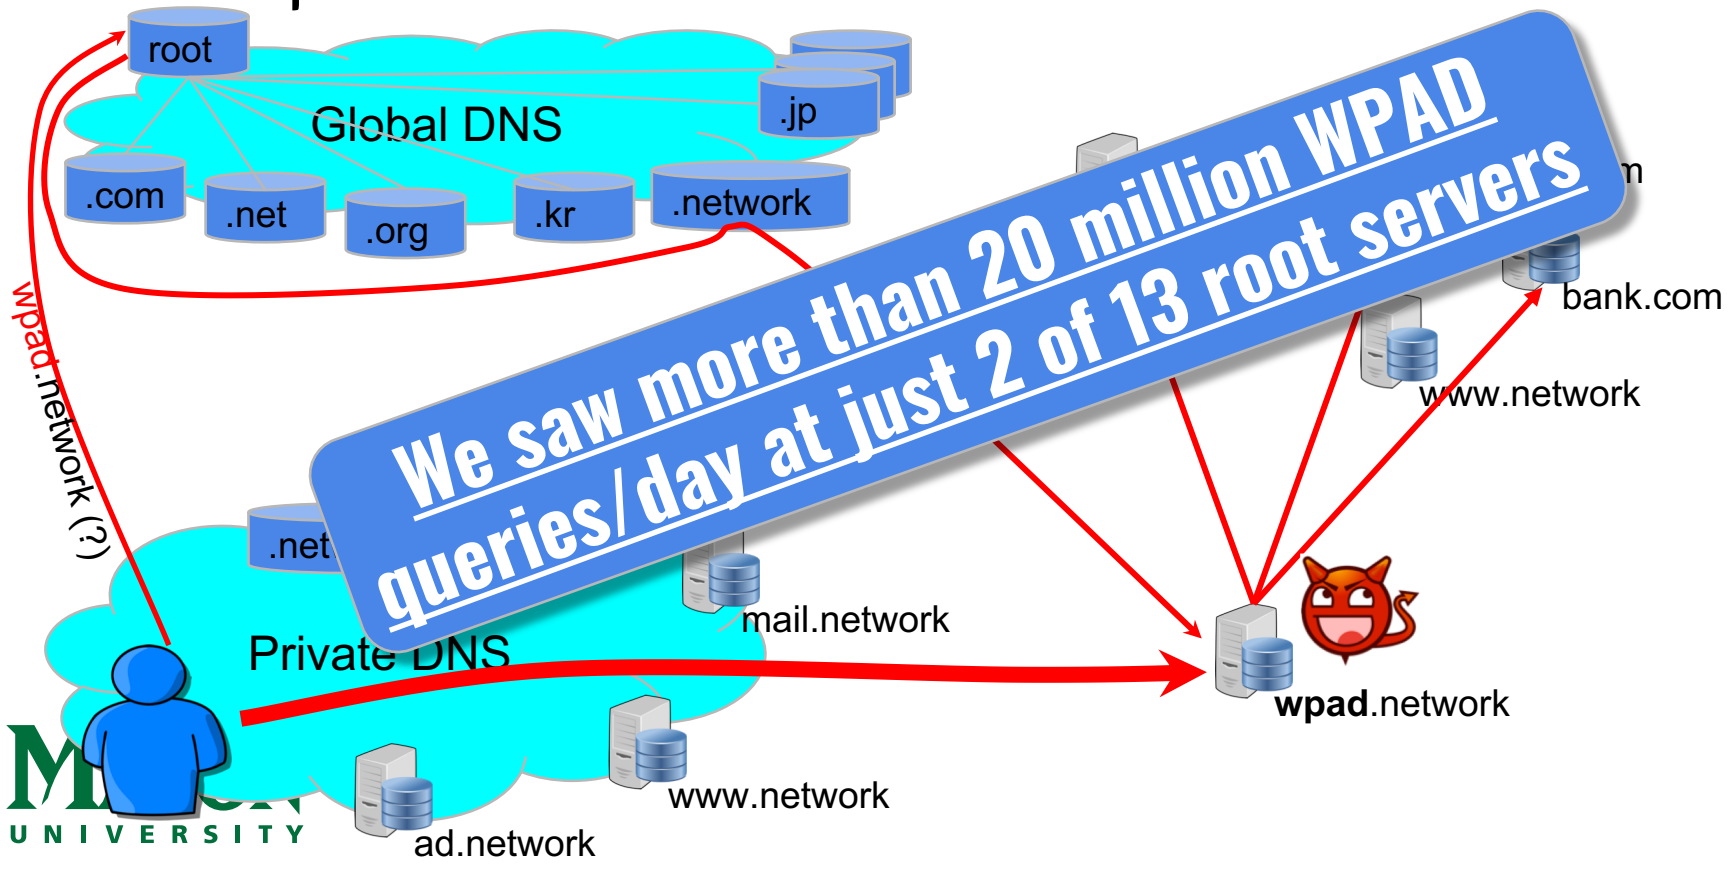

Why are these collisions happening?

- To find out **why**, we started with "**who**" (i.e. who is "vulnerable")?
- Looked networks (Autonomous Systems, ASes) leaking vulnerable queries
- Found 12 ASes source 85% of vulnerable queries
  - *All* home access or open DNS resolver networks
- **Names** indicated corporations

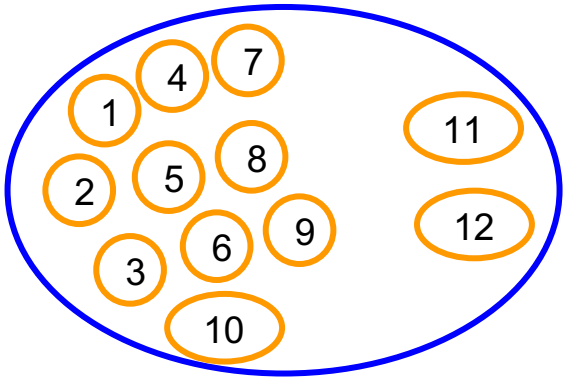

GEORGE MASON UNIVERSITY

Tracking the cause

- Hypothesis: vulnerable queries from corp. **laptops** when on **remote** networks
  - Corollary: should be most evident in networks that aggregate **diverse** remote clients

- *Evaluated*: entropy metric, to measure **diversity** of vulnerable queries:

$$-\sum_{suf \in sld.tld} p_{suf} \ln p_{suf}$$

- *Results*: 10 of the top 12 most vulnerable ASes were also 10 of the top 11 high-entropy sources!

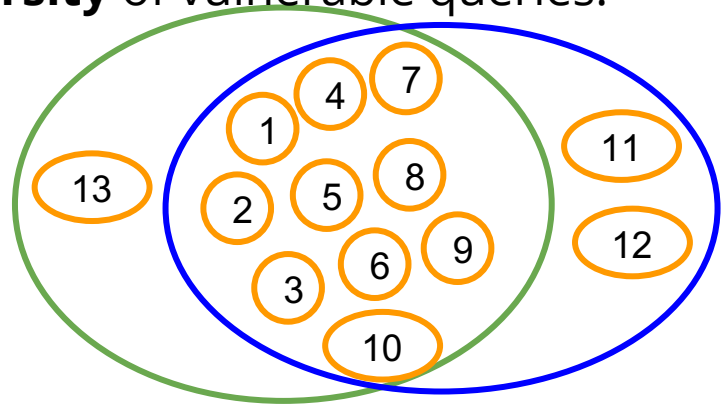

GEORGE MASON UNIVERSITY

Name collisions: leaks vs. entropy

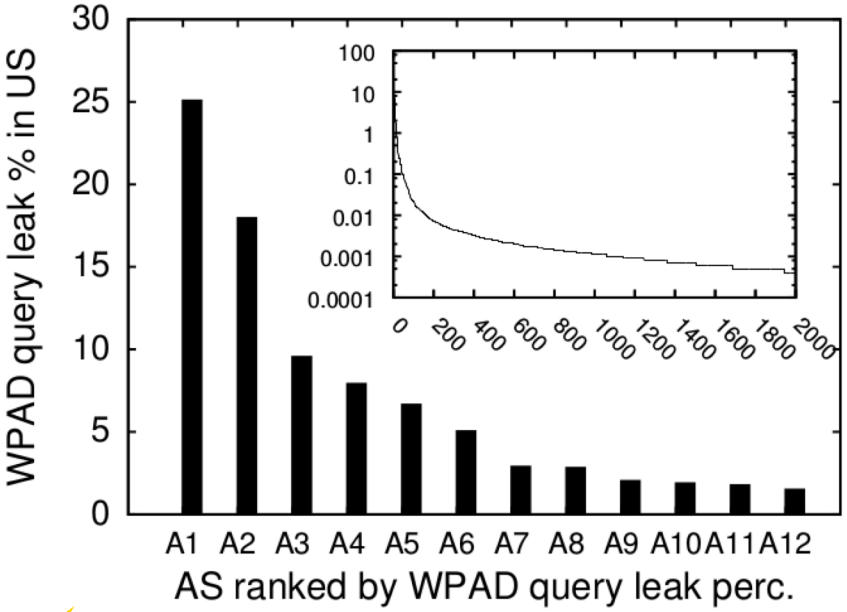

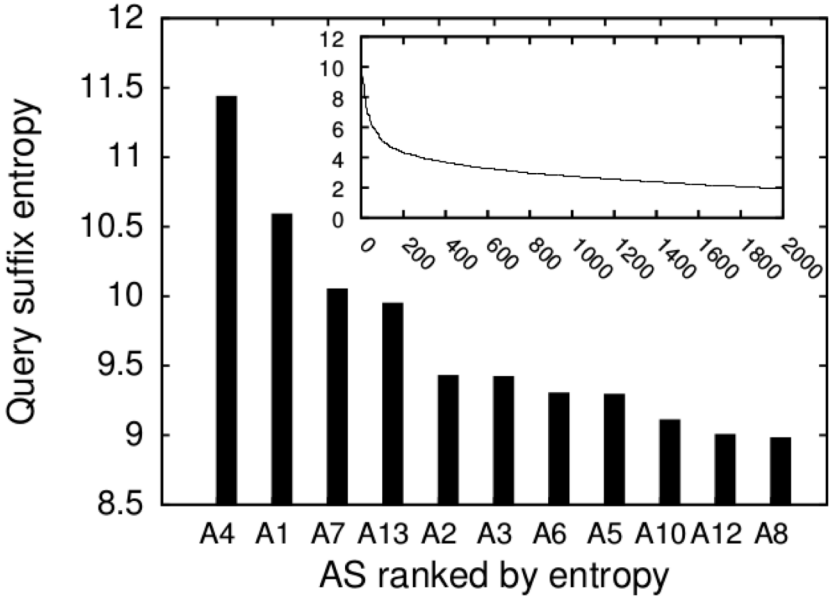

GEORGE MASON UNIVERSITY

Quantifying the attack surface

- Attack surface metric = Highly Vulnerable Domains (HVDs)
  - High volume + persistent: victims an adversary would recoup investment on attacking
- **Evaluated** rate new gTLDs may be being attacked
  - Who is most vulnerable and are they being targeted at registration?
- Among top 10 ASes, HVD metric made up 96.7% of vulnerable queries
  - Attack window opening, but not abnormally
- Found registration rate was same for HVDs and non-HVDs
  - Indicated no malicious intent (at that time)

GE MASON UNIVERSITY

Remediating the name collision vulnerability

- Alerted the community
  - Alerted NTT-ME about Chiba, Japan SmartCity, 2013
  - US-CERT Technical Alert TA16-144A, 2016 [US-CERT-16]
- Advice: name collision triage
  - At the registry: employ registration block-lists for HVDs
  - At the network perimeters: drop query/responses/payload for WPAD
  - At the client: investigate **namespace disambiguation**

- Forward: dependent protocols also need **defense in depth** [CCS-17]

Operationalizing and studying our substrate has taught us a lot, but…
What cybersecurity defenses can it enable?

MASON UNIVERSITY

Outline

- Background: cybersecurity in the Internet's core
- Current research challenges
- Future directions

GEORGE MASON UNIVERSITY

Basic research from the operational Internet

- Operationalizing cybersecurity with principled investigations keeps the Internet's foundation solid

- Some exciting directions
  - Studying and enhancing Internet-scale secure crypto key learning
  - Studying and combating global intern-domain Man-in-the-Middle (MitM) attacks
  - Analyzing the use of privacy techniques in the Internet and evaluating their effects on cybersecurity

For example, some ongoing investigations...

- DNSSEC just rolled its root key over (*on October 11th, 2018*)
  - Measurement suggested concern because validators could not learn the new root key
  - How *should* we securely, stably, and resiliently manage and bootstrap crypto?

- Google just got MitM'ed by China telecom and a Russian ISP, *this Monday*
  - Can we proactively detect weaknesses in deployed inter-domain dependencies?
  - How can we quantify and monitor cross-modal attack surfaces for complex systems?

- Data privacy is a very large topic, but it competes with threat intelligence
  - The General Data Protection Regulation (GDPR) has hamstrung Abuse Response
  - How can we **evaluate** tradeoffs between privacy and security?

Building toward future needs

- We've known how to build end-to-end security for email for years (S/MIME, PGP, etc.)
  - But why can't we make it *operational* and ubiquitous today?
  - Because **previously** we couldn't securely learn crypto keys!

- Investigating the next challenges
  - Patching **connected vehicles**
  - Security substrate for **IoT**
  - Securing and upgrading **SmartCities**
  - **Medical** devices (mHealth)
  - **Fuse telemetry** for cybersecurity information sharing

Summary and future directions

- We are entering a golden-age of cybersecurity
  - Protections that we've sought for years are now in reach

- As we operationalize key building blocks we meet threats head on
  - Tomorrow's cybersecurity needs us to be rigorous and treat security like a *science*

True "defense in depth" begins in the Internet's foundation.
The lessons that we can learn will predict the success of our defenses.

GEORGE MASON UNIVERSITY

Bibliography

- [CCS-17] "Client-side Name Collision Vulnerability in the New gTLD Era: A Systematic Study" Qi Alfred Chen, Matthew Thomas, Eric Osterweil, Yulong Cao, Jie You, Z. Morley Mao, In *ACM Conference on Computer and Communications Security (CCS '17)*, November 2017
- [US-CERT-16] "WPAD Name Collision Vulnerability" United States Computer Emergency Readiness Team (US-CERT) Alert (TA16-144A), May 2016
- [Oakland-16] "MitM Attack by Name Collision: Cause Analysis and Vulnerability Assessment in the New gTLD Era" Qi Alfred Chen, Eric Osterweil, Matthew Thomas, Z. Morley Mao, In *IEEE Symposium on Security and Privacy (S&P), 2016*, pp. 675-690. IEEE, 2016.
- [SysDeps-NPSec-14] "The Shape and Size of Threats: Defining a Networked System's Attack Surface" Eric Osterweil, Danny McPherson, and Lixia Zhang, In *9th IEEE Workshop on Secure Network Protocols (NPSec), 2014* , pp. 636-641. *Best Paper Award*
- [pd-tpds-14] "Verifying Keys through Publicity and Communities of Trust: Quantifying Off-Axis Corroboration" Eric Osterweil, Dan Massey, Danny McPherson, Lixia Zhang, *IEEE Transactions on Parallel and Distributed Systems* 25, no. 2 (2014): 283-291, *Supplemental text here*
- [ROVER] Gersch, Joseph, and Dan Massey. "Rover: Route origin verification using dns." In *Computer Communications and Networks (ICCCN), 2013 22nd International Conference on*, pp. 1-9. IEEE, 2013.
- [TASRS] "TASRS: Towards a Secure Routing System Through Internet Number Resource Certification" Eric Osterweil, Danny McPherson, Verisign Labs Technical Report #1130009, February 2013
- [ResCert-Hotnets-11] "The Great IPv4 Land Grab: Resource Certification for the IPv4 Grey Market" Eric Osterweil, Shane Amante, Danny McPherson, Dan Massey, In *10th ACM Workshop on Hot Topics in Networks (HotNets)*, p. 12, 2011.

GE MASON UNIVERSITY

Bibliography (2)

- [trust-fist-09] "Managing Trusted Keys in Internet-Scale Systems" Eric Osterweil, Dan Massey, Lixia Zhang, In *Ninth Annual International Symposium on Applications and the Internet (SAINT), 2009*. pp. 153-156. IEEE, 2009.
- [SecSpider-ACSAC-09] "Deploying and Monitoring DNS Security (DNSSEC)" Eric Osterweil, Dan Massey, Lixia Zhang, In *25th Annual Computer Security Applications Conference (ACSAC), 2009* pp. 429-438. IEEE, 2009.
- [SecSpider-IMC-08] "Quantifying the Operational Status of the DNSSEC Deployment" Eric Osterweil, Michael Ryan, Dan Massey, Lixia Zhang, In *8th ACM SIGCOMM conference on Internet measurement (IMC), 2008*, pp. 231-242.
- [SecSpider-NPSec-07] "Observations from the DNSSEC Deployment" Eric Osterweil, Dan Massey, Lixia Zhang, In *3rd IEEE Workshop on Secure Network Protocols (NPSec), 2007*. pp. 1-6.
- [pski-hotsec-06] "Security Through Publicity" Eric Osterweil, Dan Massey, Batsukh Tsendjav, Beichuan Zhang, Lixia Zhang, In *1st USENIX Workshop on Hot Topics in Security (HotSec)*, 2006. pp. 3-3.
- [SecSpider] SecSpider: Global DNSSEC Deployment Tracking, Eric Osterweil http://secspider.verisignlabs.com/
- [libsmaug] libsmaug A C++ library for DANE protocols, Eric Osterweil https://github.com/verisign/smaug

GEORGE MASON UNIVERSITY

Thank you!
Questions?