In The Line of Fire: Risks of DPI-Triggered Data Collection

Ariana Mirian

CSET 2023, August 7 2023

All research begins with a story

All research begins with a story

We requested a resource on a machine

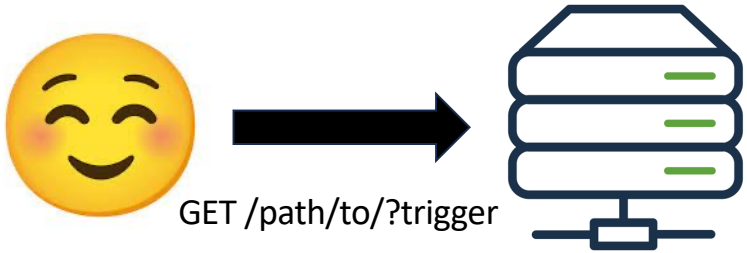

All research begins with a story

We requested a resource on a machine

Started receiving requests for that same exact path

All research begins with a story

We requested a resource on a machine

Started receiving requests for that same exact path

Suspicious because:

1) the server was not configured to allow directory listing
2) entire directory was protected by HTTP basic authentication
3) the server wasn’t advertising to the public

All research begins with a story

We requested a resource on a machine

Started receiving requests for that same exact path

Suspicious because:

1) the server was not configured to allow directory listing
2) entire directory was protected by HTTP basic auth
3) the server wasn't advertising to the public

When we changed the path…same behavior occurred


GET /path/to/?trigger
GET /path/to/?trigger

Oh no, we're pwned

Oh no, we're pwned

Spent next 48 hours coordinating and working with IR team at UCSD

Finally confirmed that this was behavior from a FireEye protection mechanism

FireEye Network Security - NX Series

Effective protection against cyber breaches for midsize to large organizations

NX 2550, NX 3500, NX 5500, NX 10450 (not pictured NX 2500, NX 4500, NX10550)

Not pwned, just “protected” (by FireEye)

Threat Intelligence collection system that fetches resources

Specifically identifies suspicious files by name

We had inadvertently triggered it by naming our file with a specific string

Not pwned, just “protected” (by FireEye)

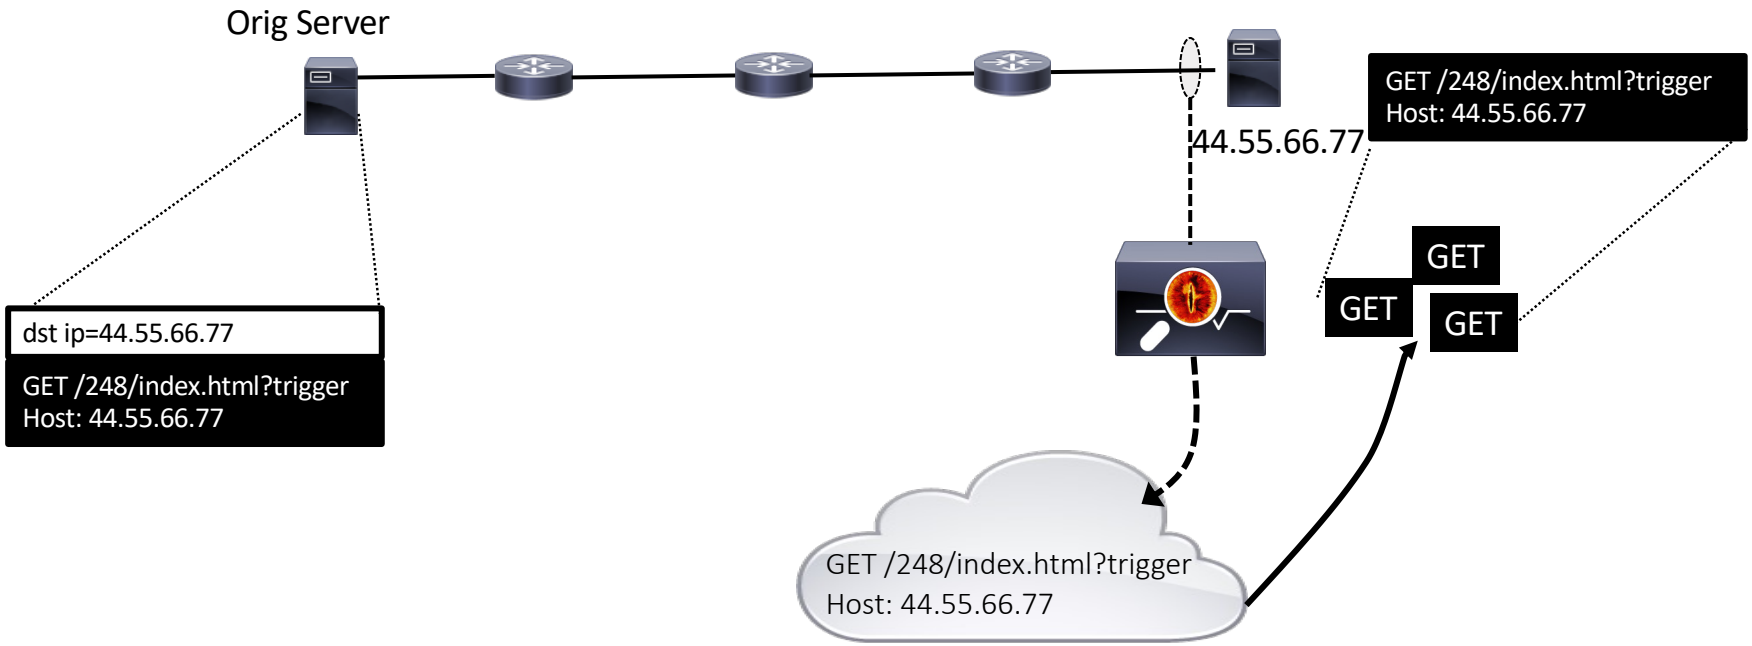

Can we measure this?

Can we measure this?

Could we scan and trigger the FireEye protection mechanism to:

1) Understand the global footprint of FireEye?
2) Characterize the proxy (request) network?

Host header is the key to the measurement

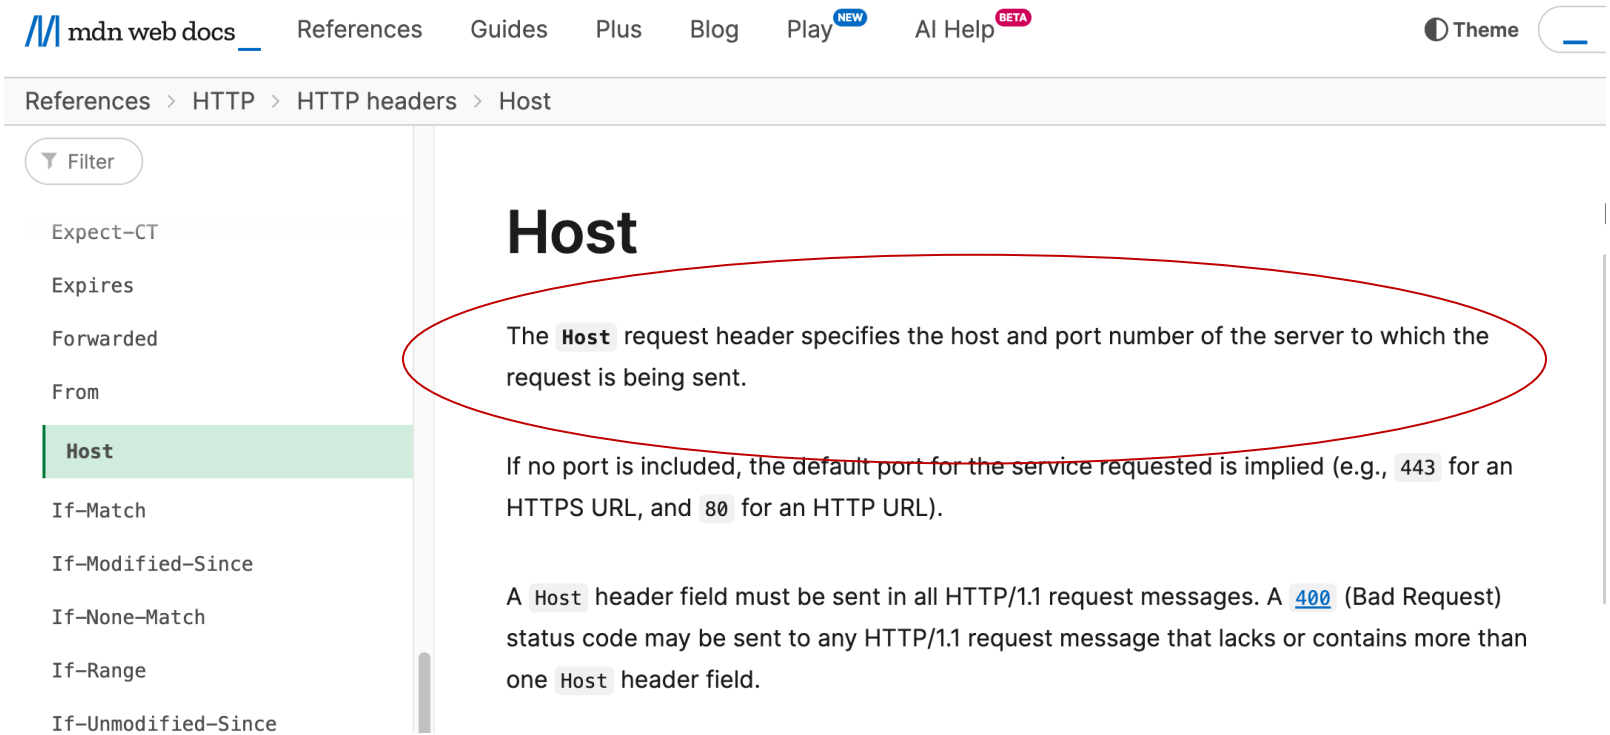

Initial Setup

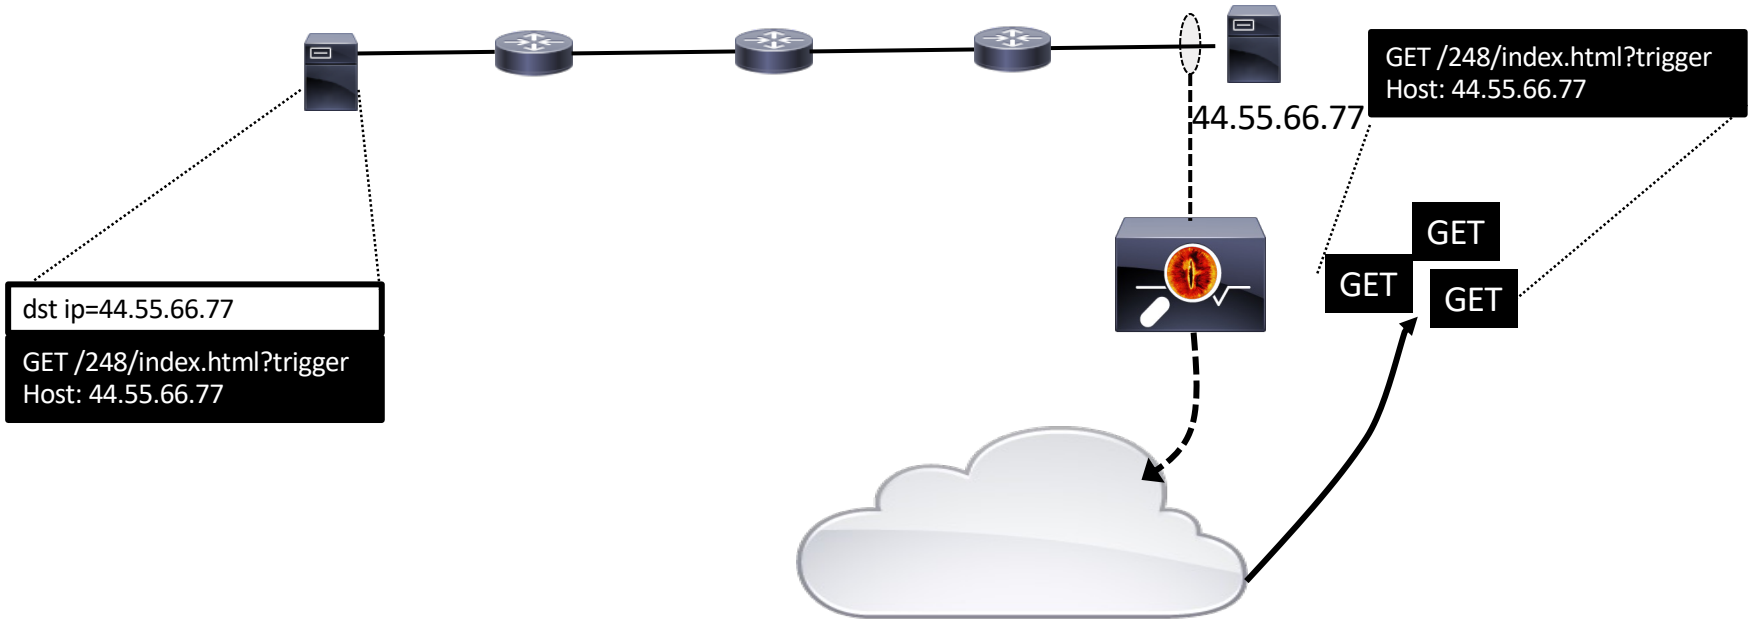

Measuring FireEye coverage

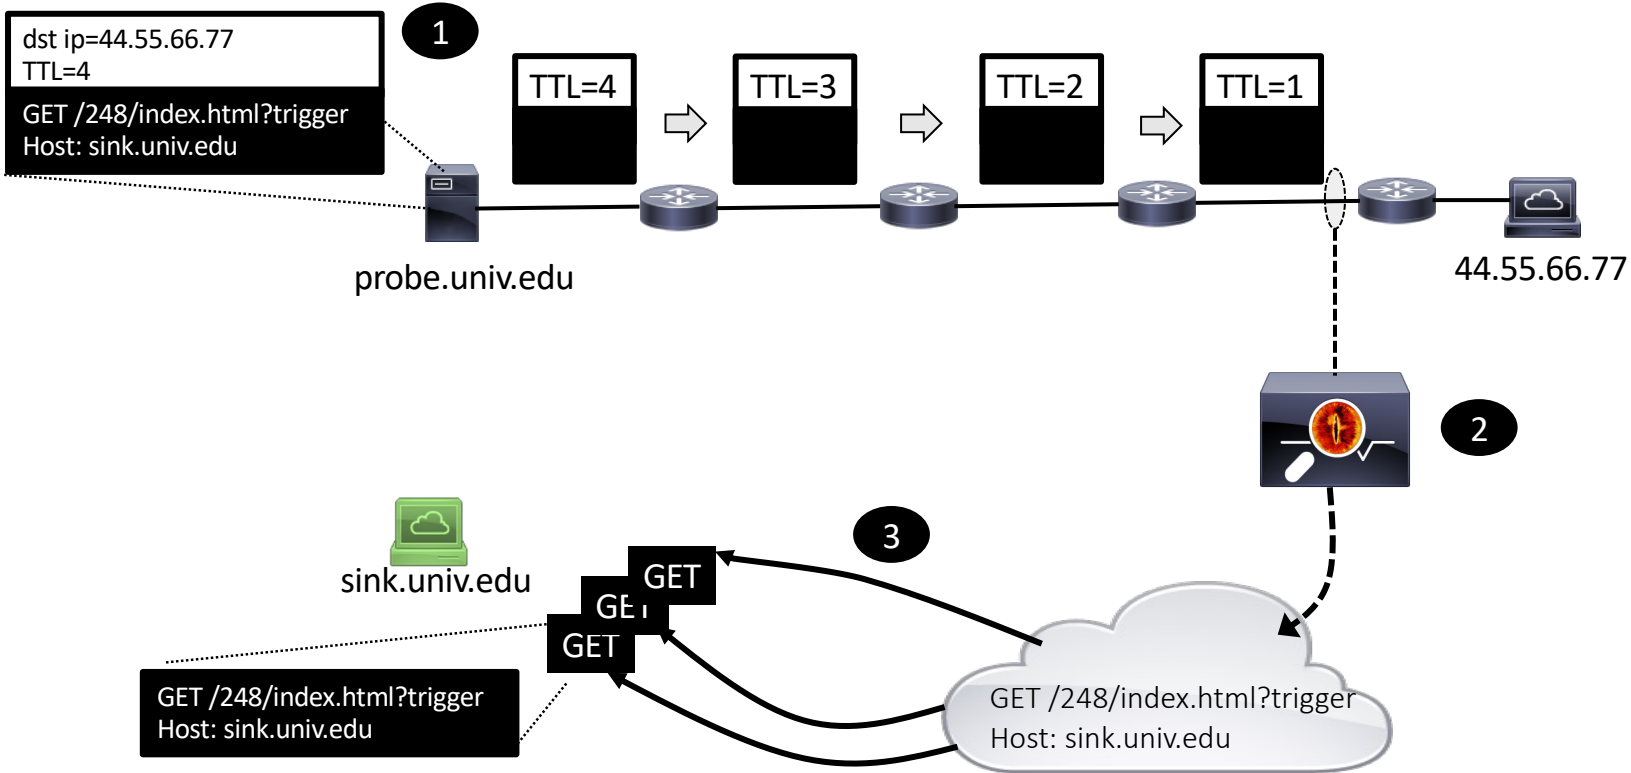

Scan Mechanisms

Traceroute to calculate the forward path hop count, N

Send five probes (GET requests) to target using TTL of N-1

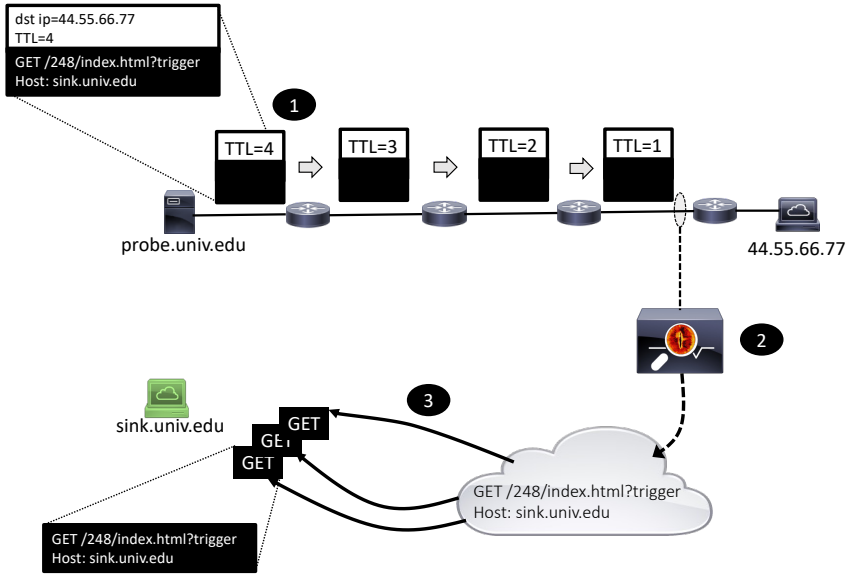

Measuring FireEye coverage

Downloaded 80M IPS that offer service on port 80 from Censys

Subsampled to mitigate our effect on the network --- 80M to 3M

Performed 3 scans on each IP

Of 3M IPs, 50k were responsive

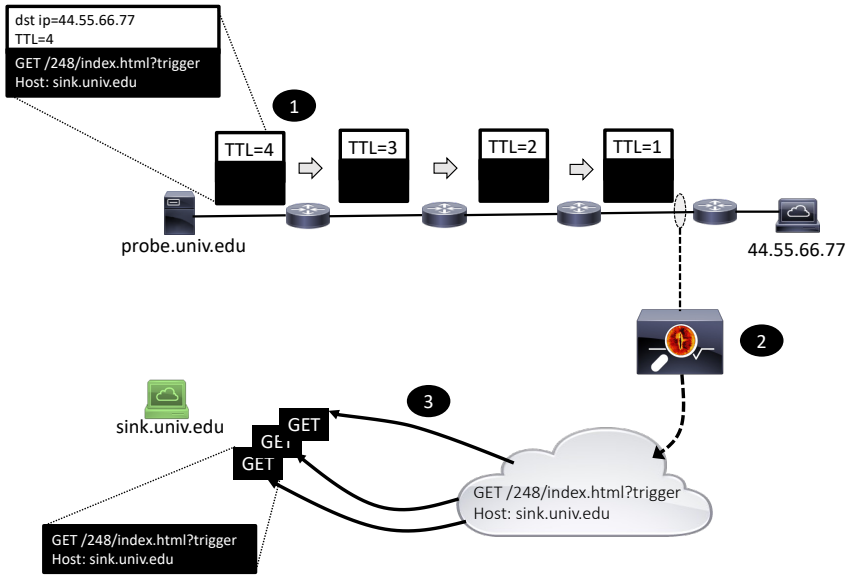

Global Spread of FireEye

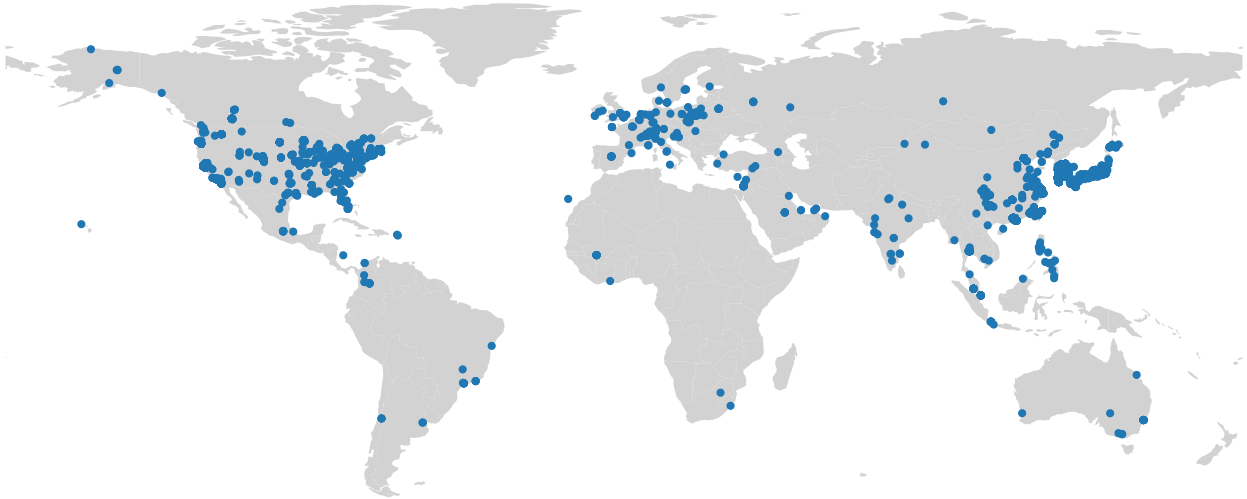

Global Spread of FireEye

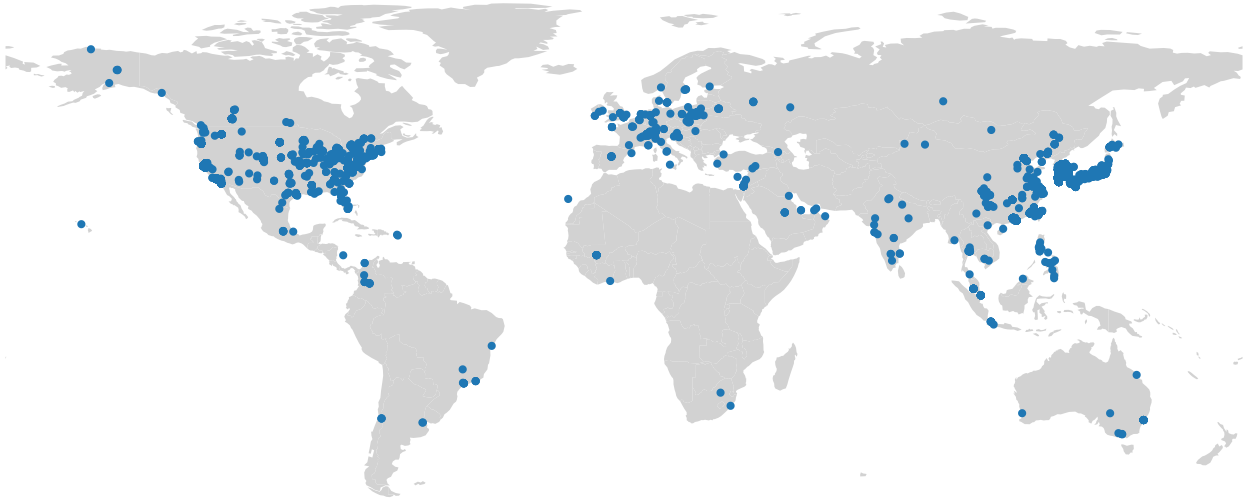

Large spread of FireEye in US, Europe, Asia

Global Spread of FireEye

Probed IP Addresses		Responsive IP Addresses	
ASN Name	% of IPs (#)	ASN Name	% of IPs (#)
COMCAST-7922	4.07% (135810)	SKB-ASSKBroadbandCoLtd	48.99% (24903)
AMAZON-02	2.97% (99152)	KIXS-AS-KRKoreaTelecom	35.90% (18247)
KIXS-AS-KRKoreaTelecom	2.82% (94140)	HWCSNETHuaweiCloudServicedatacenter	1.77% (901)
DTAGInternetserviceprovideroperations	2.38% (79549)	UCSD	0.52% (265)
ATT-INTERNET4	2.25% (75267)	UCLA	0.52% (223)
AMAZON-AES	1.57% (52453)	VA-TECH-AS	0.44% (207)
FranceTelecom-Orange	1.40% (46776)	CHINANET-IDC-BJ-AP	0.37% (187)
OCNNTTCommunicationsCorporation	1.31% (43740)	BIZLAND-SD	0.33% (169)
ASN-IBSNAZ	1.21% (40474)	ICNDP-AS-KRNamincheonBrodcastingCo.	0.31% (160)
UninetS.A.deC.V.	1.20% (39999)	WSU-AS	0.31% (158)

Global Spread of FireEye

Probed IP Addresses		Responsive IP Addresses	
ASN Name	% of IPs (#)	ASN Name	% of IPs (#)
COMCAST-7922	4.07% (135810)	SKB-ASSKBroadbandCoLtd	48.99% (24903)
AMAZON-02	2.97% (99152)	KIXS-AS-KRKoreaTelecom	35.90% (18247)
KIXS-AS-KRKoreaTelecom	2.82% (94140)	HWCSNETHuaweiCloudServicedatacenter	1.77% (901)
DTAGInternetserviceprovideroperations	2.38% (79549)	UCSD	0.52% (265)
ATT-INTERNET4	2.25% (75267)	UCLA	0.52% (223)
AMAZON-AES	1.57% (52453)	VA-TECH-AS	0.44% (207)
FranceTelecom-Orange	1.40% (46776)	CHINANET-IDC-BJ-AP	0.37% (187)
OCNNTTCommunicationsCorporation	1.31% (43740)	BIZLAND-SD	0.33% (169)
ASN-IBSNAZ	1.21% (40474)	ICNDP-AS-KRNamincheonBrodcastingCo.	0.31% (160)
UninetS.A.deC.V.	1.20% (39999)	WSU-AS	0.31% (158)

Skewed concentration in two large Korea Telecom ASes

Organizational Categorization

ASN Category	% of ASes (#)
Computer and Information Technology	43.63% (315)
Education and Research	18.98% (137)
Government and Public Administration	5.96% (43)
Finance and Insurance	5.96% (43)
Service	5.96% (43)
Community Groups and Nonprofits	3.74% (27)
Retail Stores, Wholesale, and E-commerce	3.60% (26)
Manufacturing	2.77% (20)
Media, Publishing, and Broadcasting	2.22% (16)
Construction and Real Estate	1.39% (10)

Large focus on Computer and Information Technology

PTR Record Domains

Open INTEL

PTR records are DNS records that map IPs to their DNS names

OpenIntel has historical PTR records

Of the 50K IPs, we identified 229K historical PTR records for 8.5K IPs

These 8.5K IPs map to 860 registered domains

>50 large educational institutions, >40 US Govt agencies, >20 commercial

Can we measure this?

Could we scan and trigger the FireEye protection mechanism to:

1) Understand the global footprint of FireEye?
2) Characterize the proxy (request) network?

Can we measure this?

Could we scan and trigger the FireEye protection mechanism to:

1) Understand the global footprint of FireEye?
    1) Yes!
    2) Can map responsive IPs to their ASN name, location, and organization
2) Characterize the proxy (request) network?

Can we measure this?

Could we scan and trigger the FireEye protection mechanism to:

1) Understand the global footprint of FireEye?
    1) Yes!
    2) Can map responsive IPs to their ASN name, location, and organization
2) Characterize the proxy (request) network?

Proxy Network

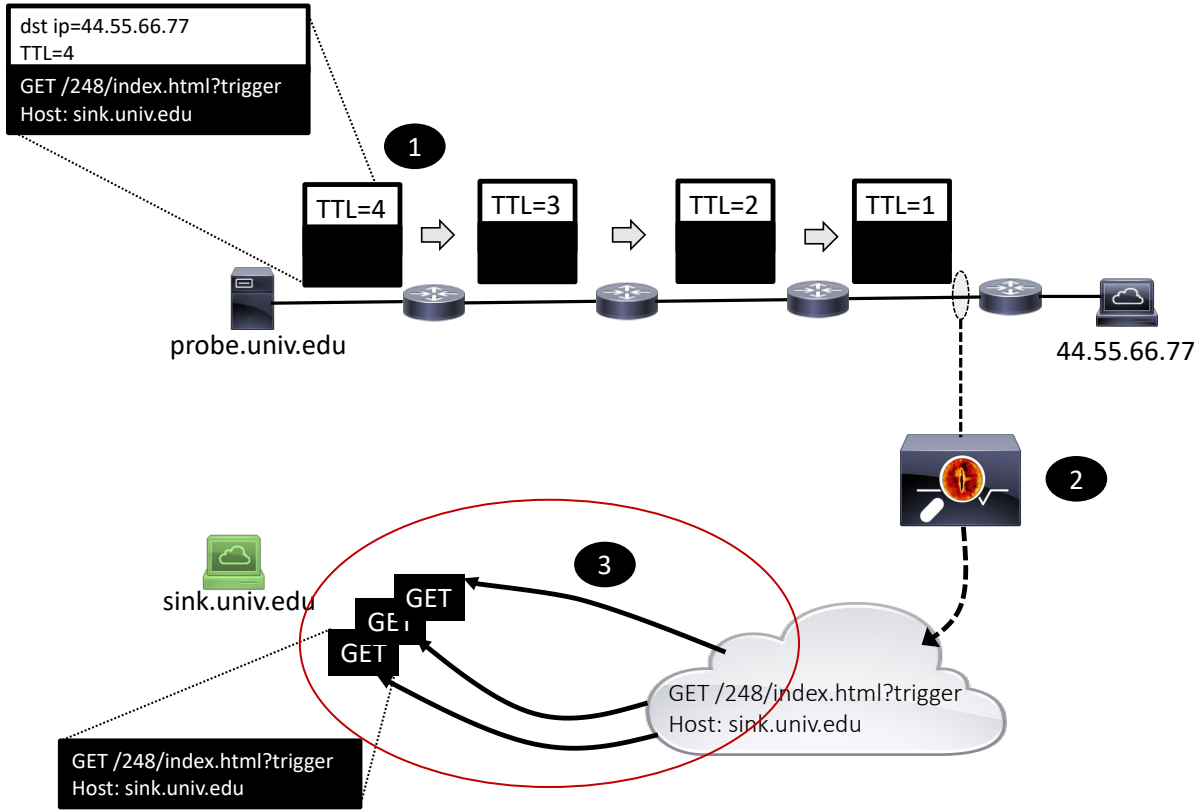

Proxy Network

We observed 568 source proxies

Collectively issued 234K requests
to our sink server

Proxy Network

We observed 568 source proxies

Collectively issued 234K requests to our sink server

Many requests issued promptly

Median time is 14 minutes

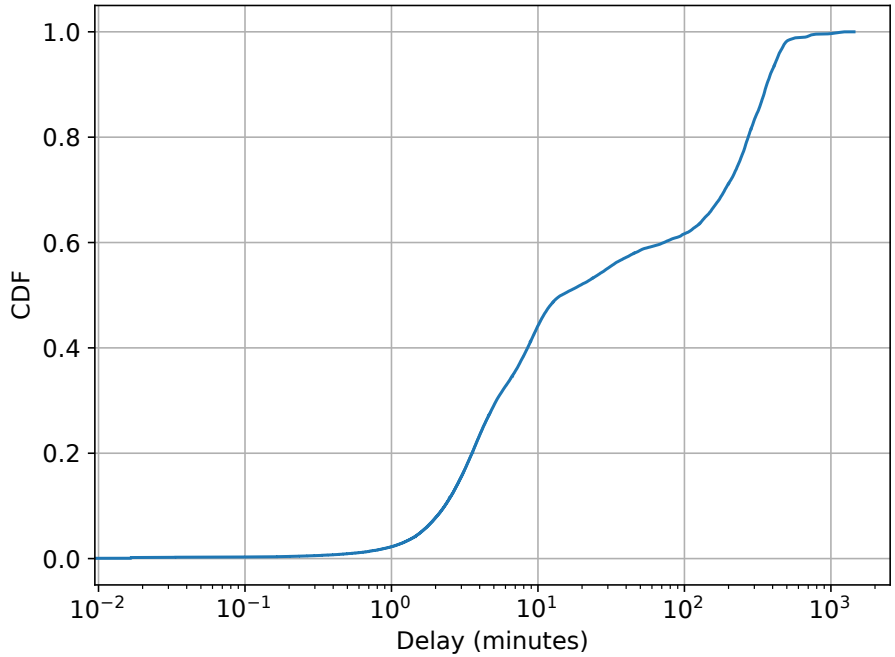

Proxy Network

Multiple requests is the norm

Median of four requests

In most extreme, 29 unique proxies

Proxy Network

Multiple requests is the norm

Median of four requests

In most extreme, 29 unique proxies

Large spread of subsequent proxies

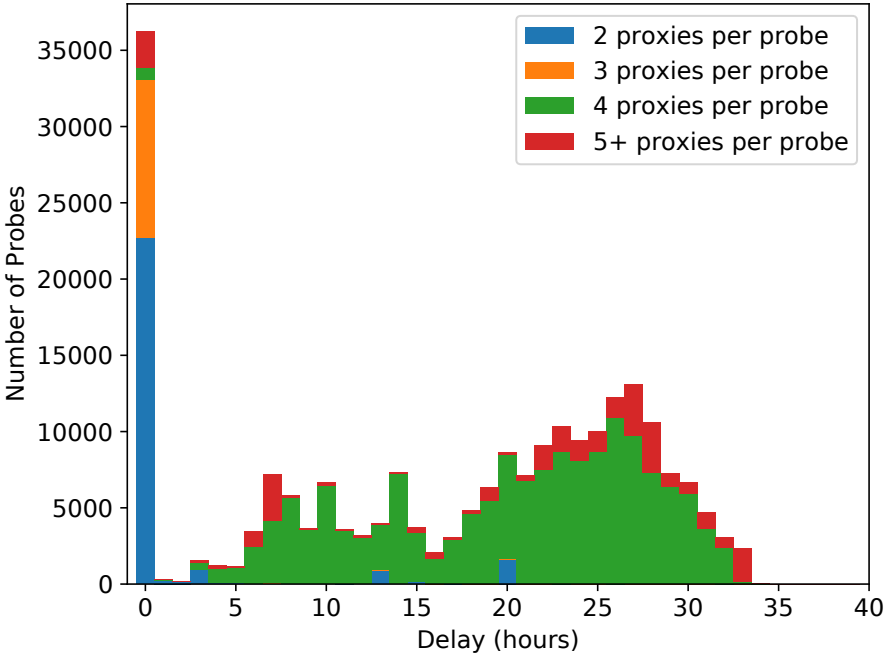

DDOS Potential

Can triggered FireEye cause a DDOS for a client?

Efforts indicate no, but still able to drive over 100 probes per second

Can also name large object files with trigger keyword

Inconclusive decreasing TTL scans

Wanted to find exact link where FireEye resides

Performed decreasing TTL scans to try to find patterns at scale

Inconclusive, but possible for targeted actors

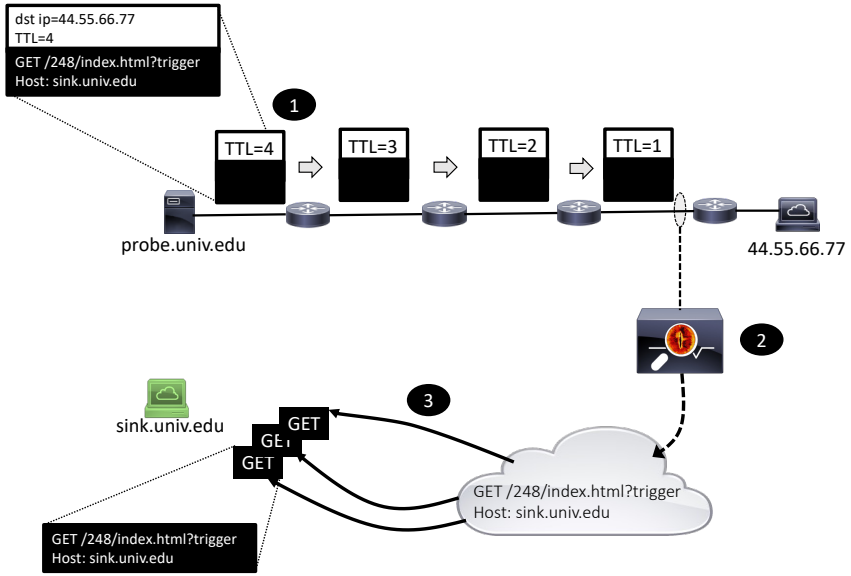

Ethical Considerations

Disclosure to FireEye

TTL N-1 scans to reduce load on end host

Limited probes and scanning

Close each opened connection with RST packet

Takeaways

Threat Intelligence gathering systems can enable reconnaissance

When customers are the data, can create inadvertent confusion

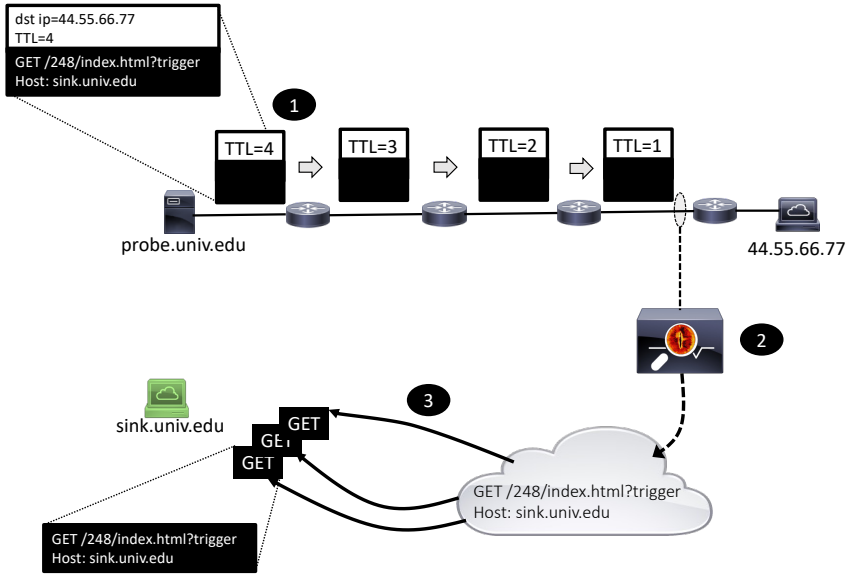

Thank you!

Alisha Ukani, Ian Foster, Gautam Akiwate, Taner Halicioglu, Cynthia T. Moore, Alex C. Snoeren, Geoffrey M. Voelker, Stefan Savage

UCSD IT and SDSC IT staff and incident response teams

Questions?

arianamirian.com

arianamirian28@gmail.com

@arimirian

@amirian@infosec.exchange

Extra Slides

Limitations

One moment in time

We didn’t have other strings to test against

Load balancing/changing network topology