Middleboxes, NAT, and the End of End-to-End

CS249i: The Modern Internet

Stanford University

End-to-End Principle

Application-specific features reside in the communicating end nodes of the network, rather than in intermediary nodes, such as gateways and routers, that exist to establish the network

Originated by Paul Baran in the 1960s, which addressed the requirement of network reliability when the building blocks are inherently unreliable

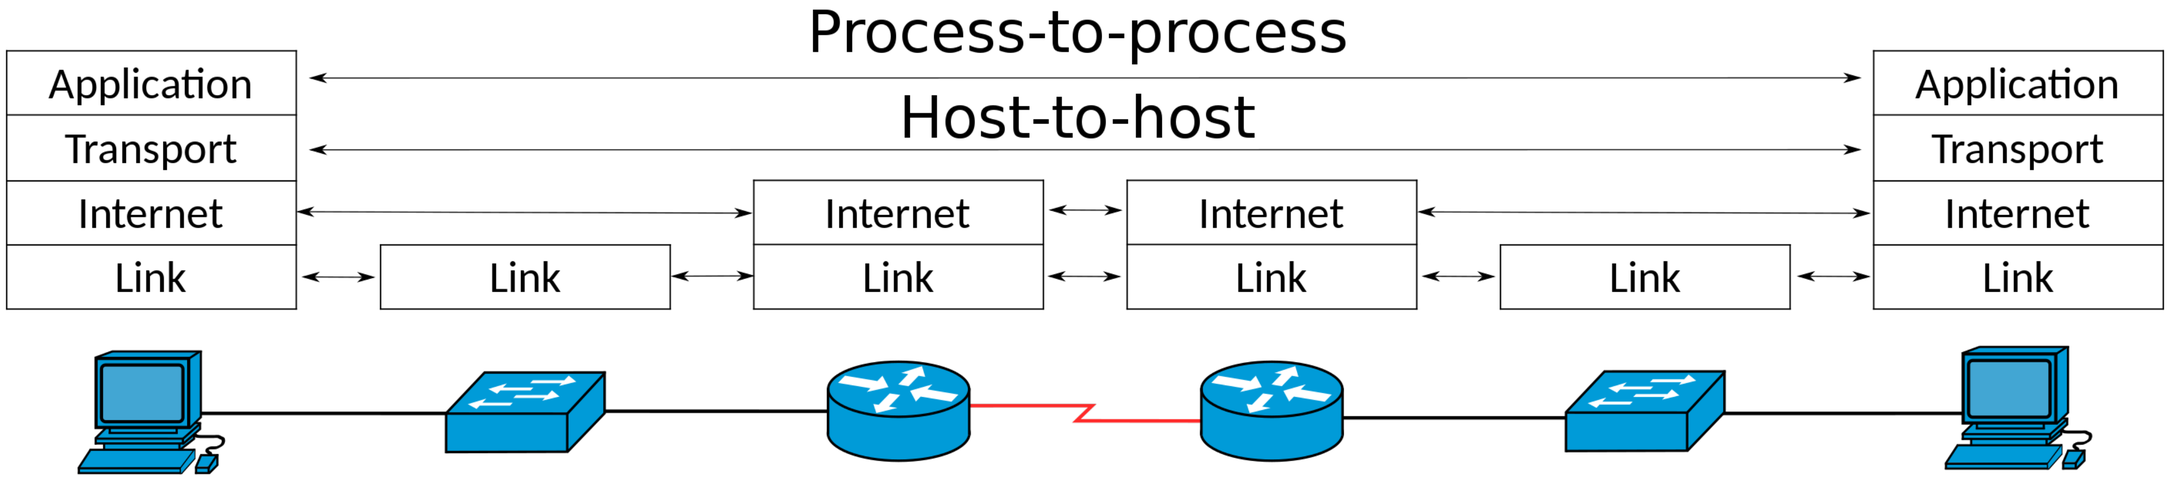

What are Middleboxes?

A middlebox is a networking device that transforms, inspects, filters, and manipulates traffic for purposes other than packet forwarding.

Includes firewalls, network address translators (NATs), load balancers, and deep packet inspection (DPI).

Middleboxes are Everywhere

Type of appliance	*Number*
Firewalls	166
NIDS	127
Media gateways	110
Load balancers	67
Proxies	66
VPN gateways	45
WAN Optimizers	44
Voice gateways	11
Total Middleboxes	***636***
Total routers	***~900***

(Sherry et al, SIGCOMM' 12)

Example Enterprise

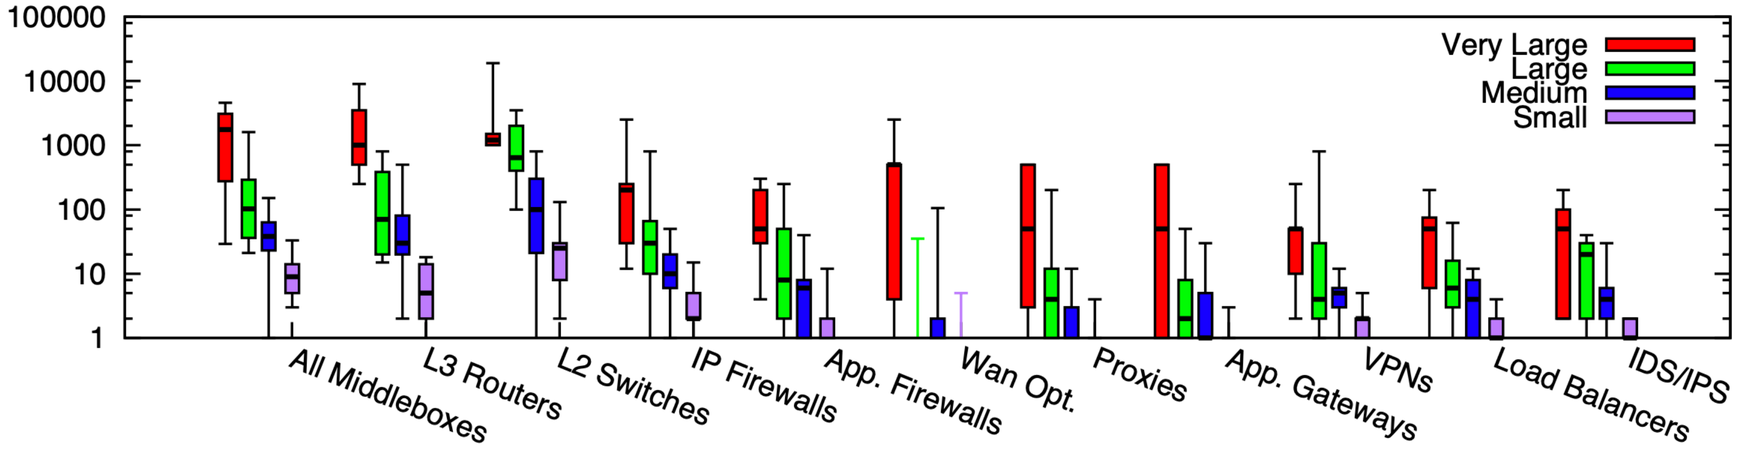

Number of Devices Across ~60 Enterprises

Verizon Advertising Header

How Verizon's Advertising Header Works

Over the past couple of days, there's been an outpouring of concern about Verizon's adver tising practices. Verizon Wireless is injecting a unique identifier into web requests, as data transits the network. On my phone, for example, here's the extra HTTP header.[1]

```
X-UIDH: OTgxNTk2NDk0ADJVquRu5NS5+rSbBANlrp+13QL7CXLGsFHpMi4LsUHw
```

Verizon Advertising Header

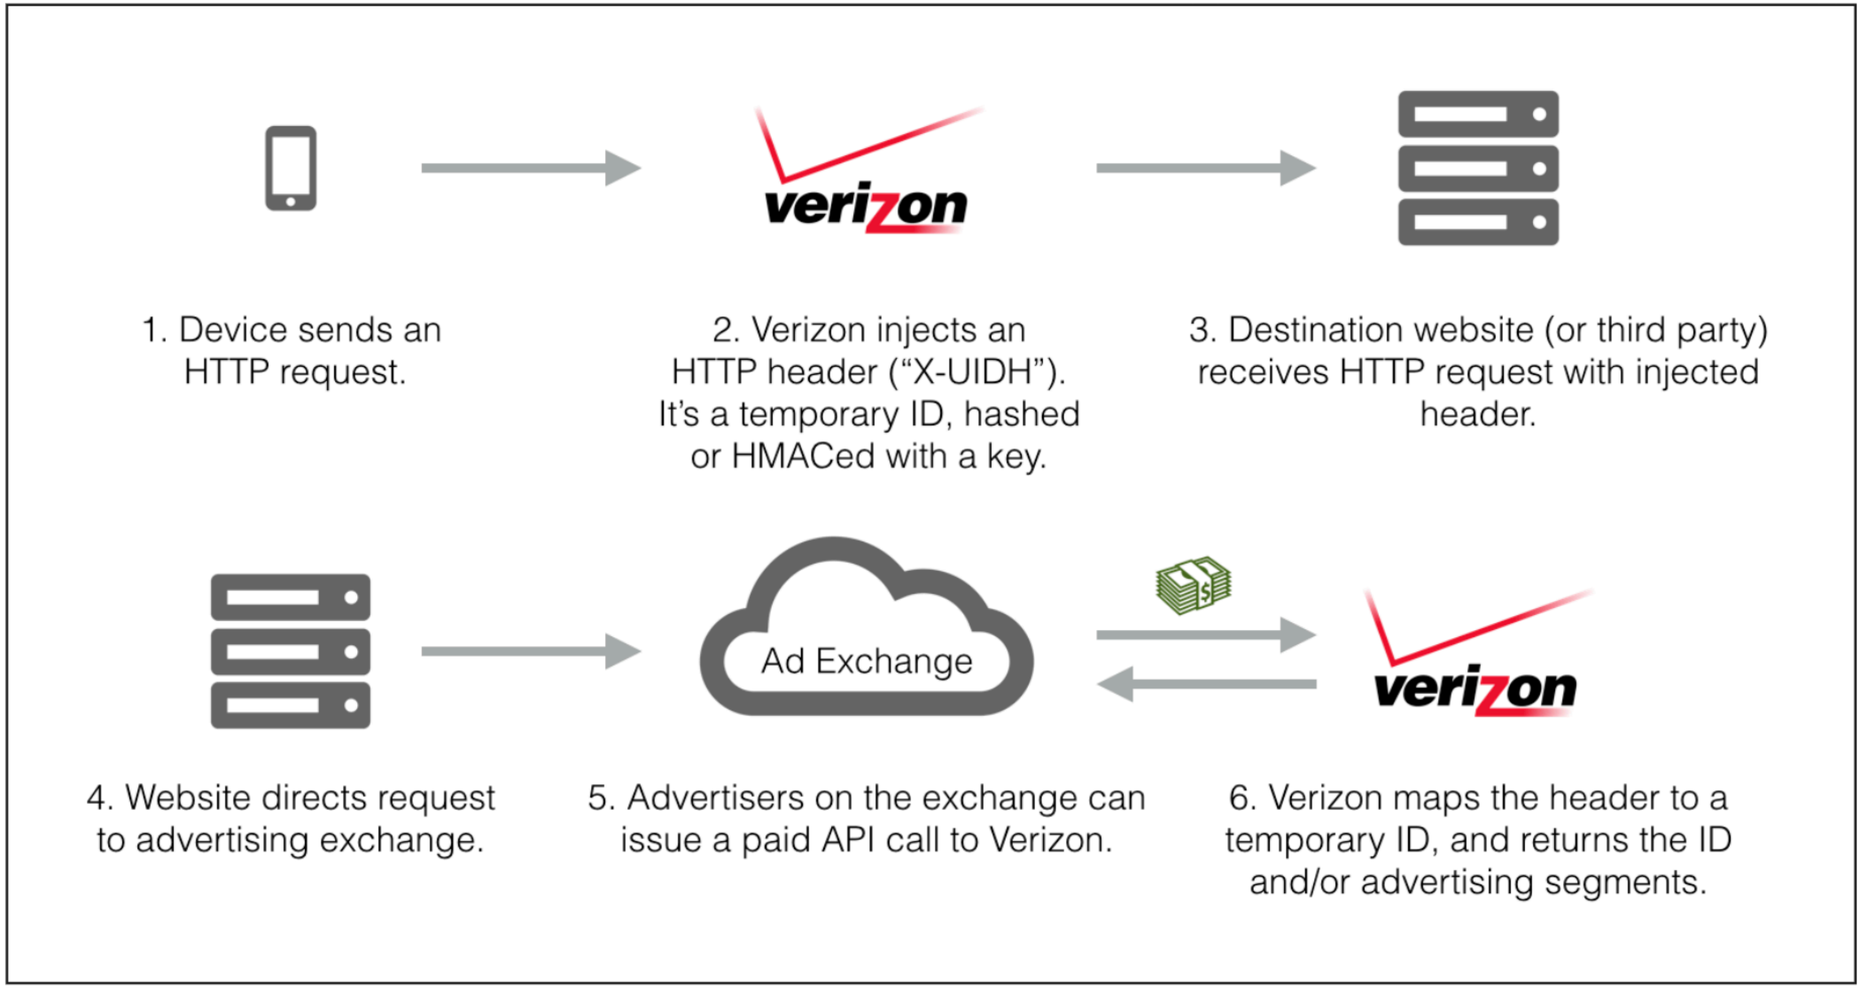

Carrier Grade NAT

NAT44 / Carrier Grade NAT ("CGNAT")

End sites (e.g., homes) are assigned private IP addresses. Middleboxes translate connections from private to public space.

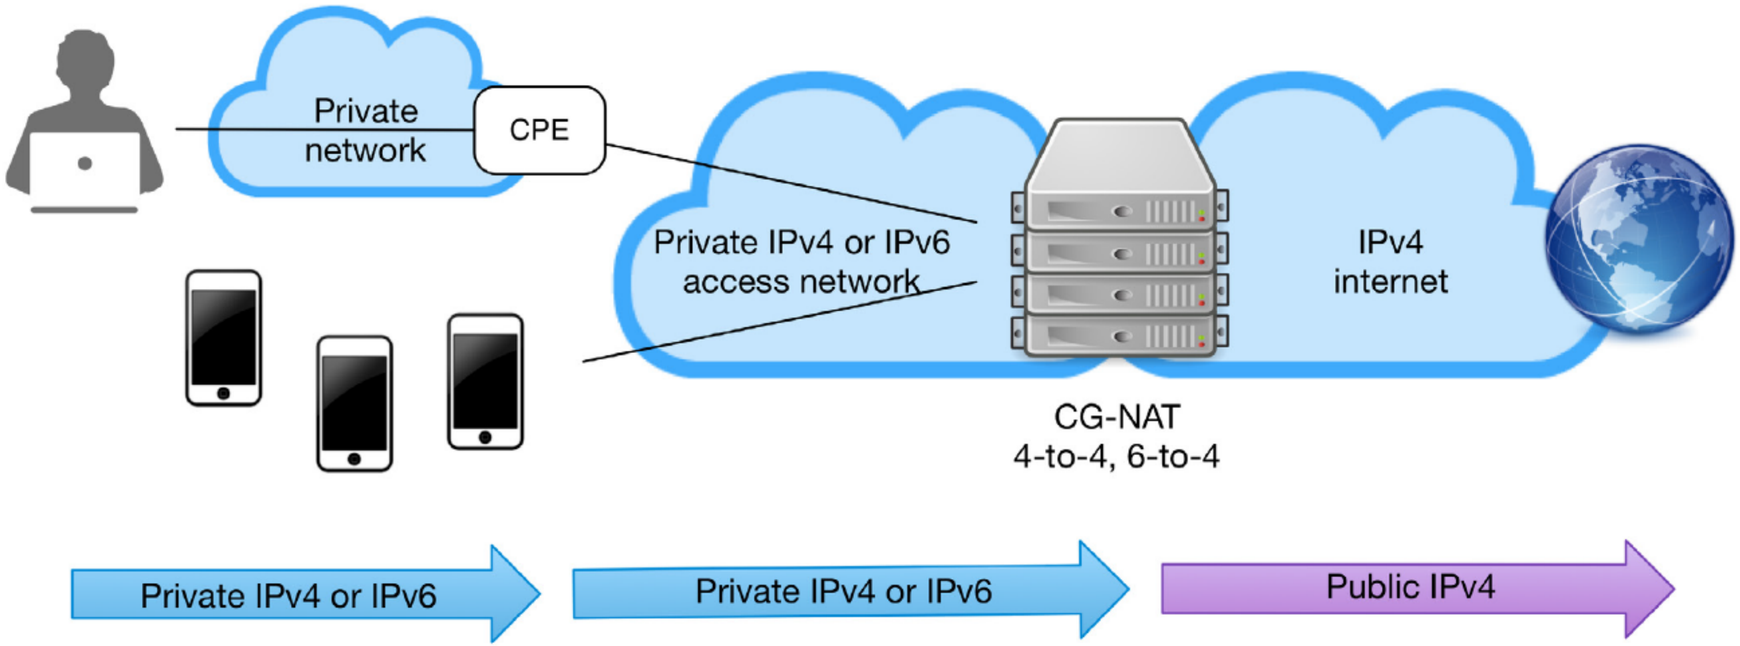

NAT444 / Carrier Grade NAT ("CGNAT")

NAT occurs both insider of home networks and ISP edge

(NAT44 + NAT44 = NAT444)

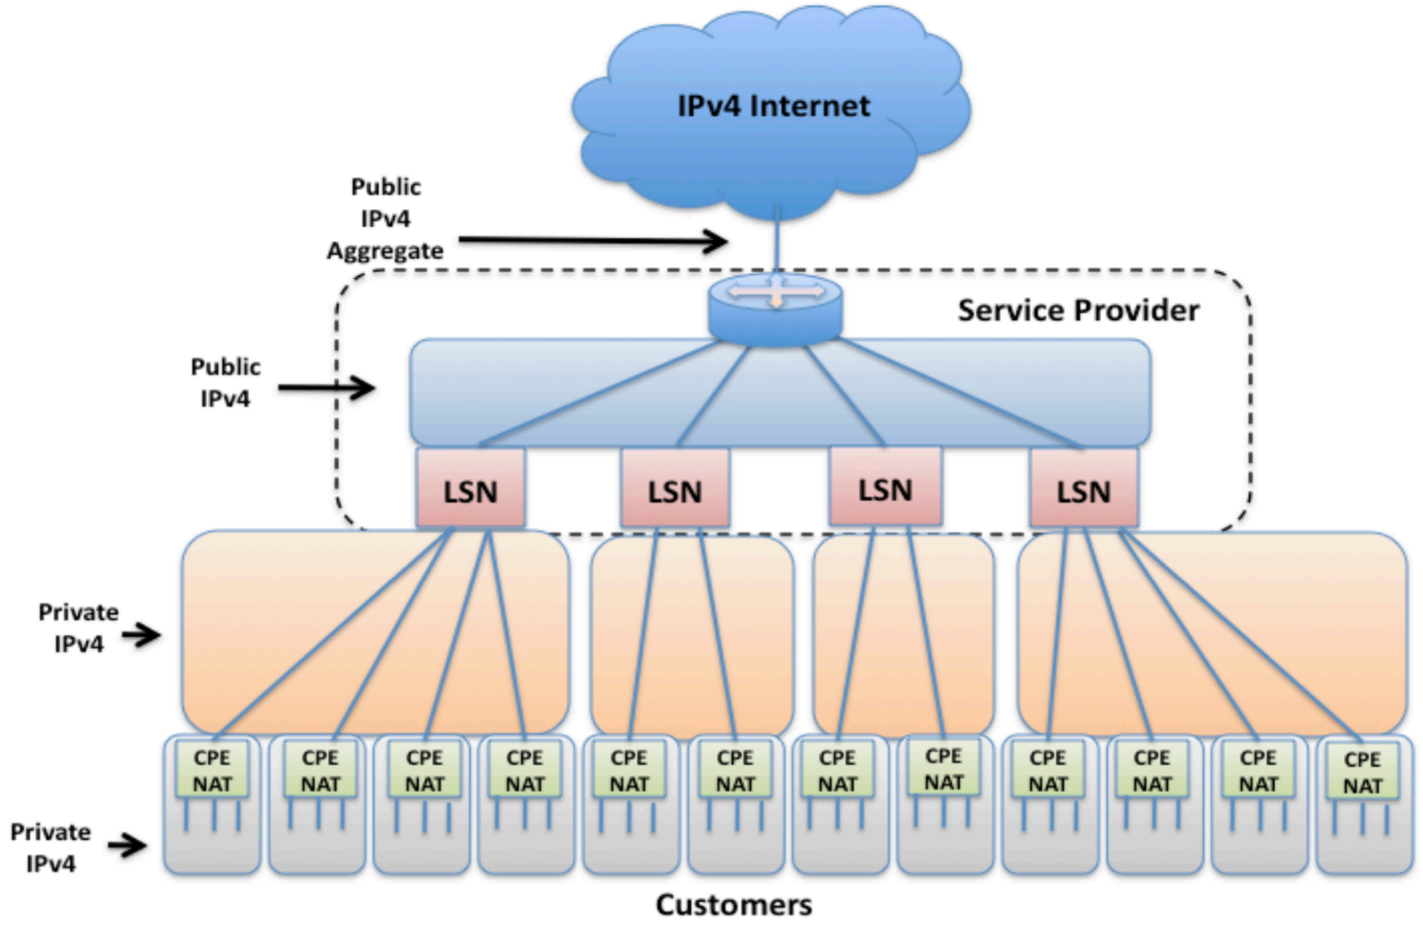

NAT444 / Carrier Grade NAT ("CGNAT")

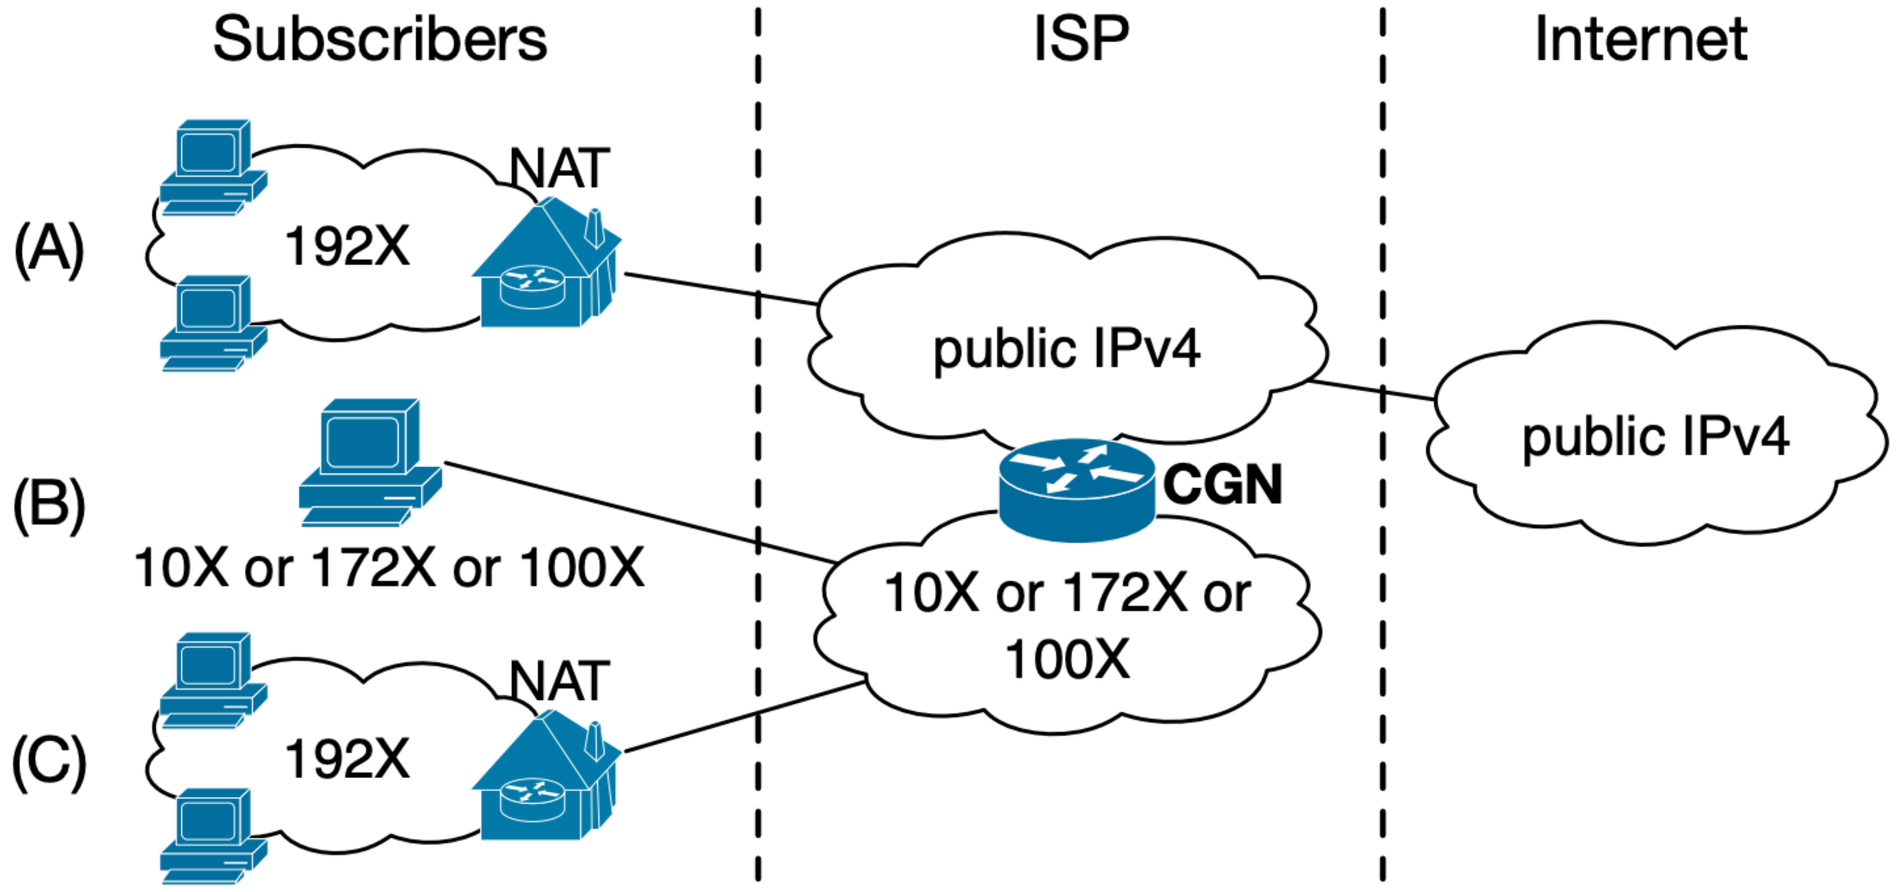

Survey: How Common is CGNAT?

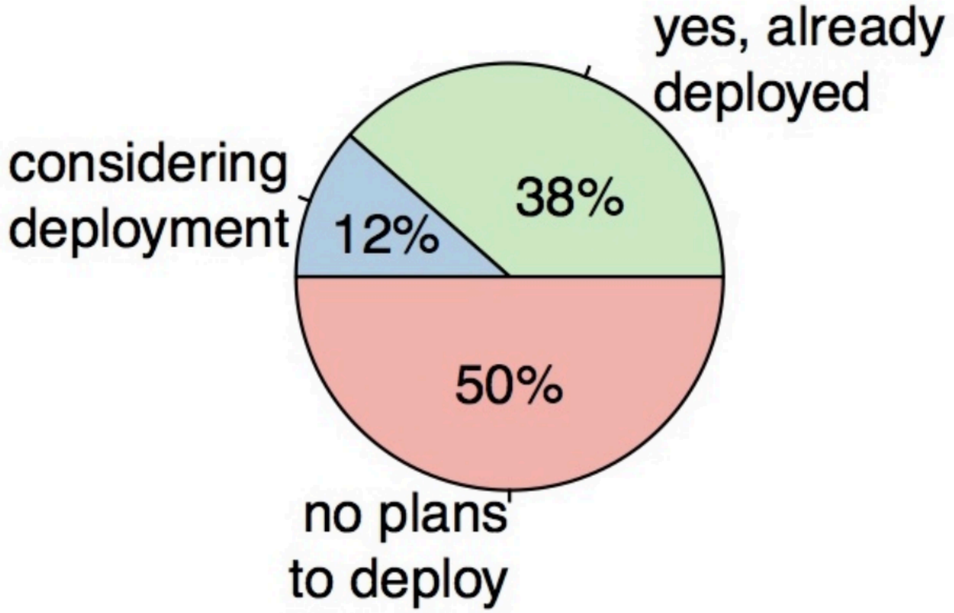

(a) Carrier-Grade NAT.

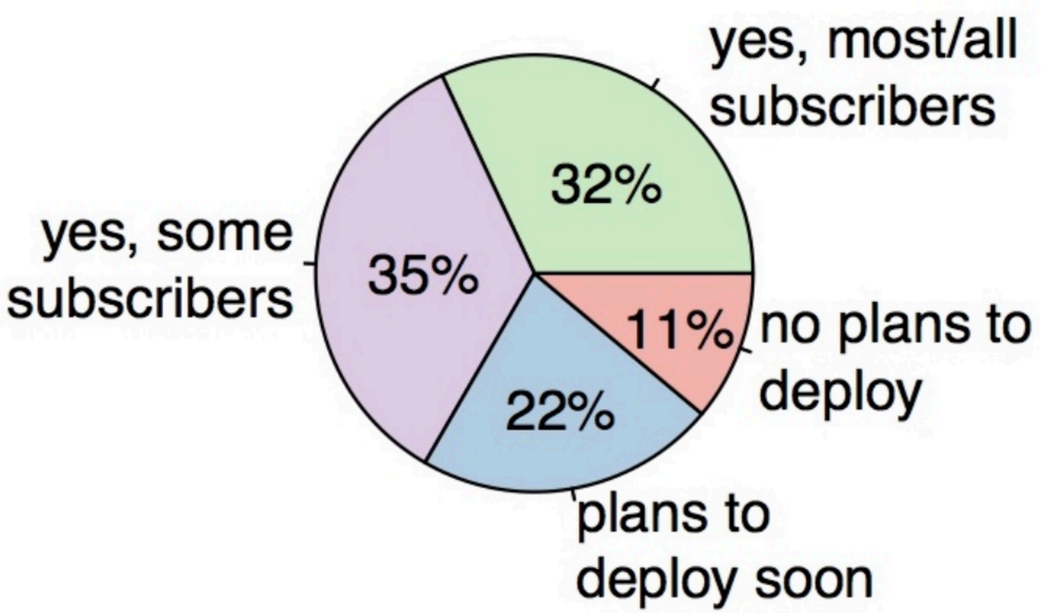

(b) IPv6.

Empirical Results

	routed ASes (N=52K)		eyeball ASes, PBL (N=2.9K)		eyeball ASes, APNIC (N=3.1K)	
	covered	CGN-positive	covered	CGN-positive	covered	CGN-positive
BitTorrent	2,724 (5.2%)	254 (9.40%)	1,673 (57.7%)	180 (10.8%)	1,824 (59.6%)	204 (11.2%)
Netalyzr non-cellular	1,367 (2.6%)	195 (14.3%)	866 (29.8%)	151 (17.4%)	929 (30.4%)	174 (18.7%)
BitTorrent ∪ Netalyzr	3,166 (6.0%)	421 (13.3%)	1,791 (61.7%)	306 (17.1%)	1,946 (63.6%)	350 (18.0%)
Netalyzr cellular	218 (0.4%)	205 (94.0%)	175 (6.0%)	162 (92.6%)	171 (5.6%)	161 (94.2%)

Table 5: Coverage and detection rates of our methods as fraction of all routed ASes, as well as Eyeball ASes, primarily connecting end users, as derived from PBL and APNIC.

Scanning

L4 Responsiveness on the Internet is a Lie

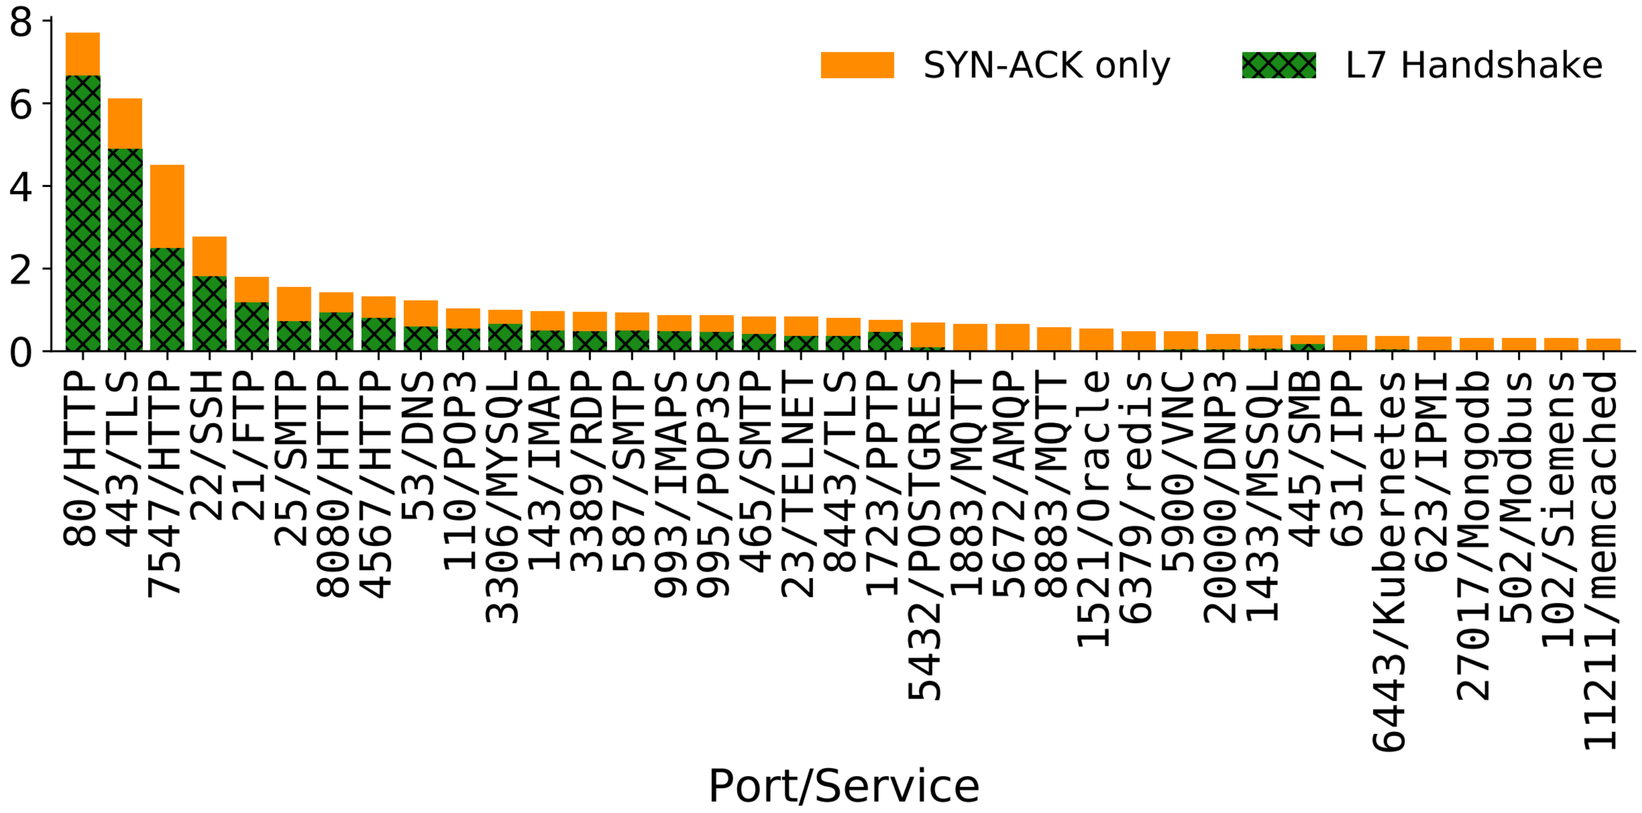

LZR: Identifying Unexpected Internet Services

Running Protocol Per RFC Isn't Sufficient

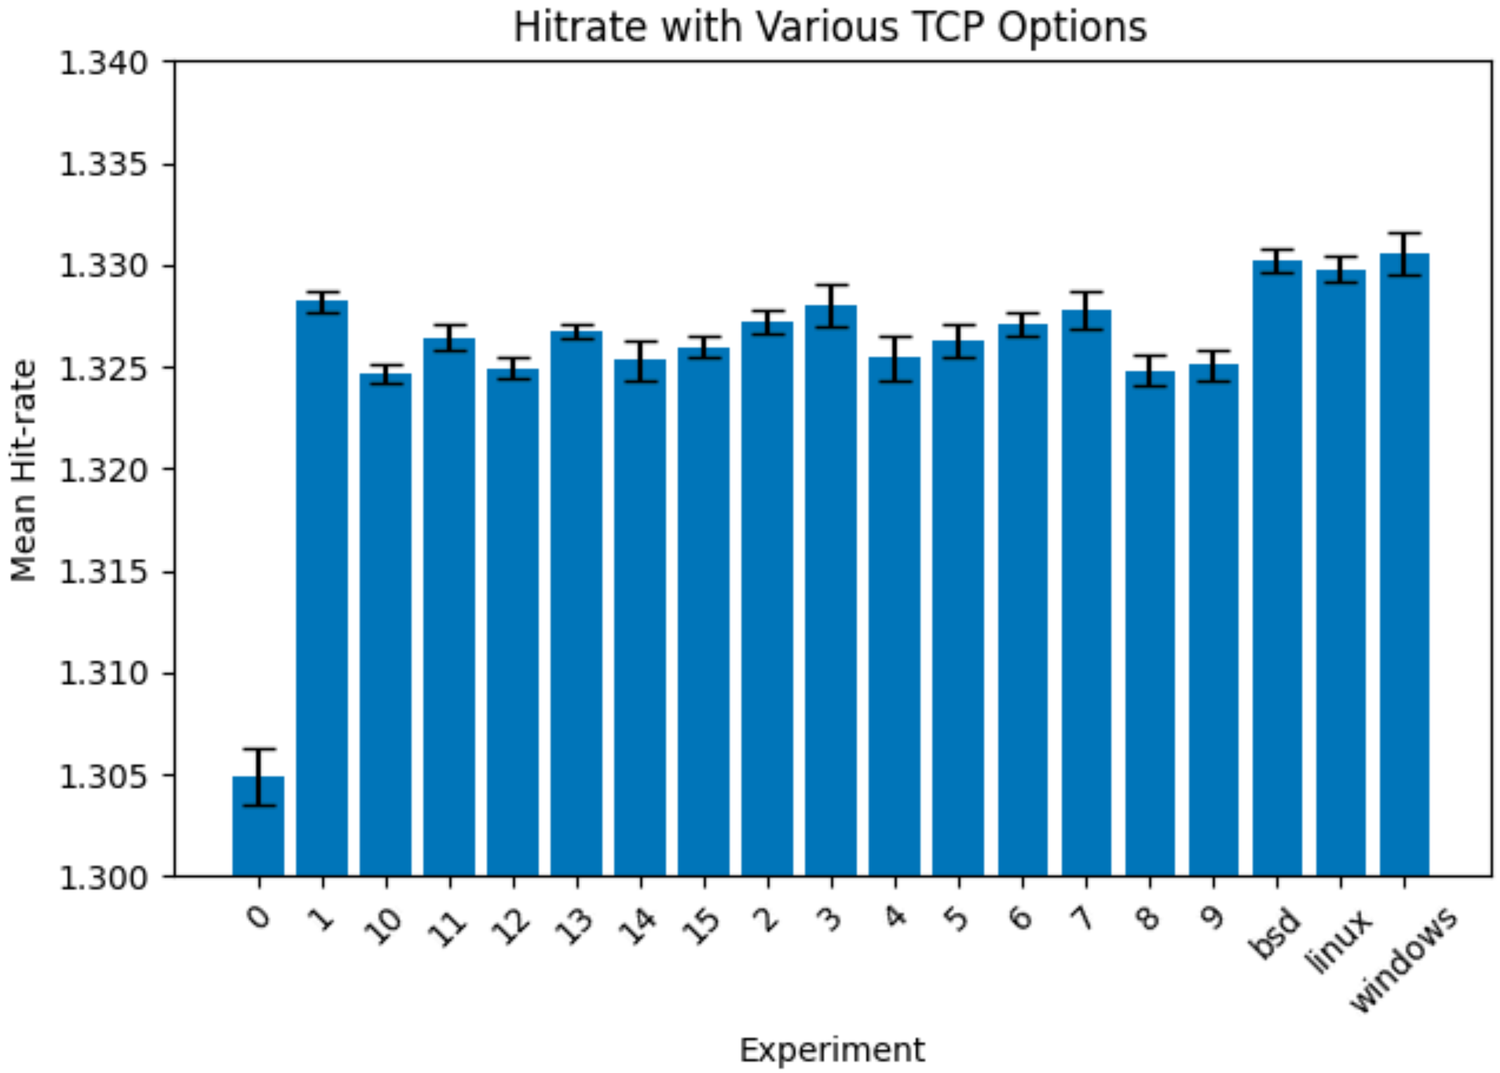

Phillip

TLS Interception

HTTPS Interception

Middleboxes and security software are increasingly intercepting HTTPS connections in order to inspect encrypted content.

How HTTPS Interception Works

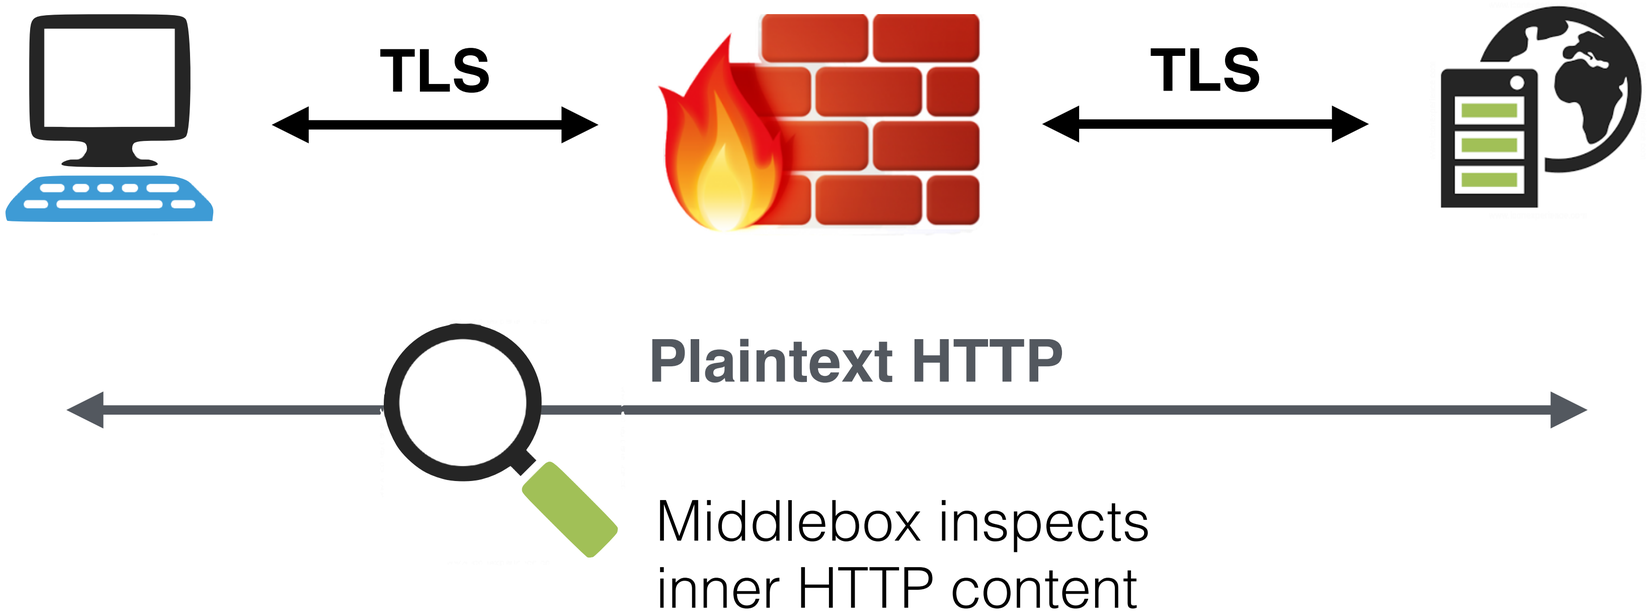

How HTTPS Interception Works

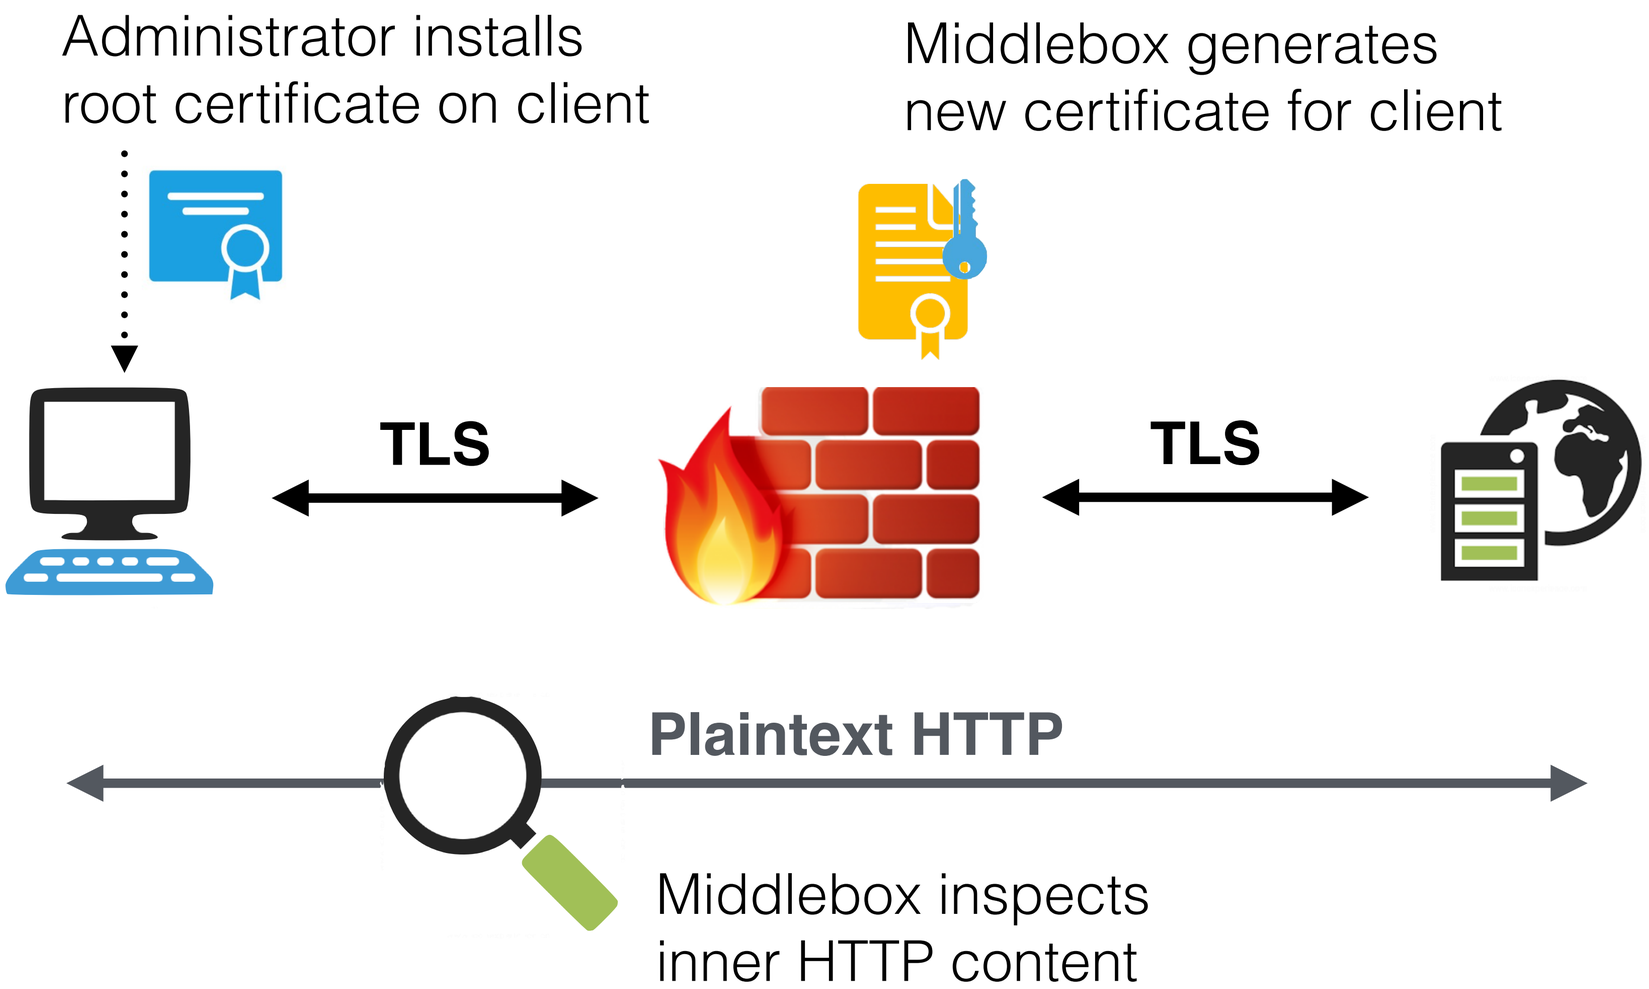

How do we measure the total amount of interception?

Change in TLS Library

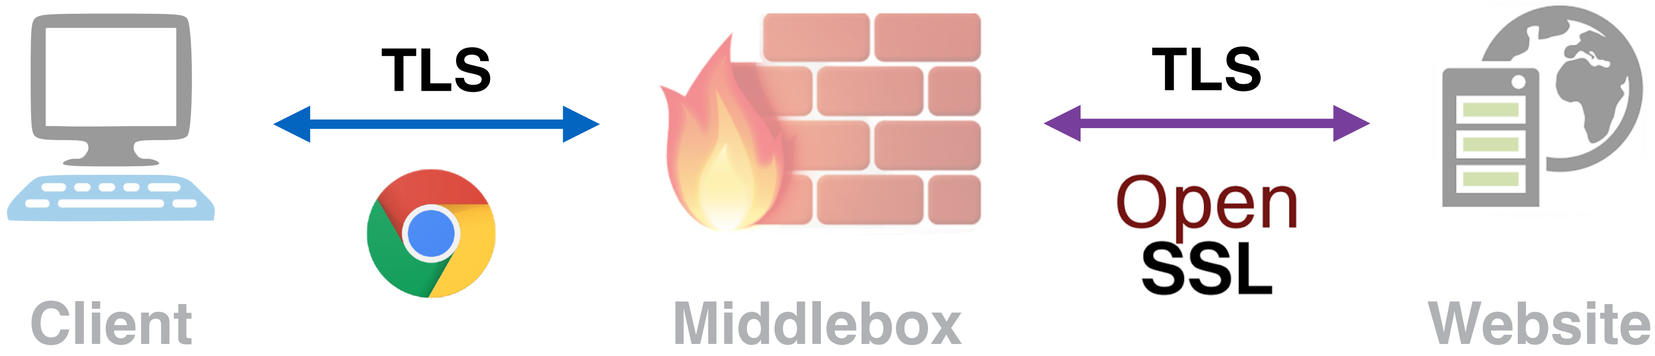

Plaintext HTTP

HTTP User Agent: Chrome

Measuring Interception

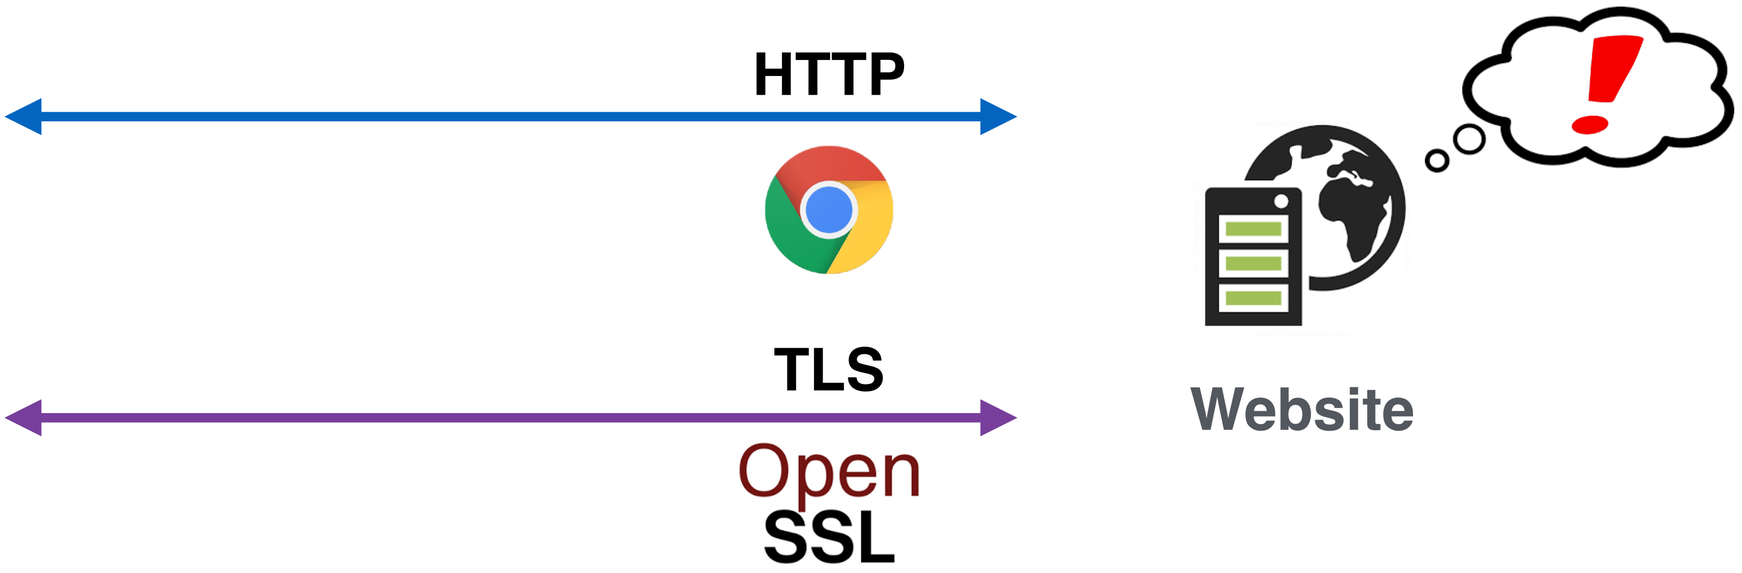

Websites can potentially detect interception by identifying a *mismatch* between network layers

Fingerprinting Network Layers

HTTP

Parse HTTP User Agent Header:

Mozilla/5.0 (Macintosh; Intel Mac OS X 10_12_2) AppleWebKit/ 537.36 (KHTML, like Gecko) Chrome/55.0.2883.95 Safari/537.36

TLS

No identifying field. Instead, we built a set heuristics that identify whether a TLS handshake is consistent with a browser.

Typical TLS Handshake

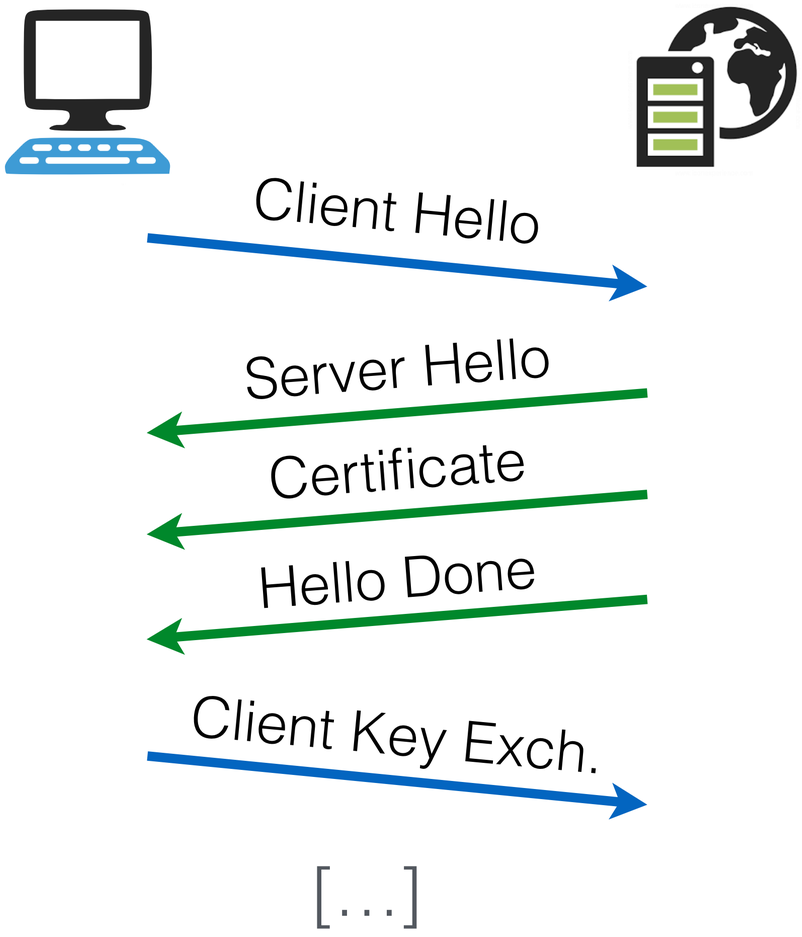


```
Secure Sockets Layer
▼ TLSv1.2 Record Layer: Handshake Protocol: Client Hello
    Content Type: Handshake (22)
    Version: TLS 1.0 (0x0301)
    Length: 217
  ▼ Handshake Protocol: Client Hello
      Handshake Type: Client Hello (1)
      Length: 213
      Version: TLS 1.2 (0x0303)
    ▶ Random
      Session ID Length: 0
      Cipher Suites Length: 36
    ▶ Cipher Suites (18 suites)
      Compression Methods Length: 1
    ▶ Compression Methods (1 method)
      Extensions Length: 136
    ▶ Extension: Unknown 35466
    ▶ Extension: renegotiation_info
    ▶ Extension: server_name
    ▶ Extension: Extended Master Secret
    ▶ Extension: SessionTicket TLS
    ▶ Extension: signature_algorithms
    ▶ Extension: status_request
    ▶ Extension: signed_certificate_timestamp
    ▶ Extension: Application Layer Protocol Negotiation
    ▶ Extension: channel_id
    ▼ Extension: ec_point_formats
        Type: ec_point_formats (0x000b)
        Length: 2
        EC point formats Length: 1
      ▼ Elliptic curves point formats (1)
          EC point format: uncompressed (0)
    ▶ Extension: elliptic_curves
    ▶ Extension: Unknown 43690
```

(Client Hello)

Firefox vs. GnuTLS Client Hellos

Extensions
Server Name (SNI)
Extended Master Secret
Renegotiation Info
Elliptic Curves
[…]

Ciphers
ECDHE_ECDSA_AES128_GCM_SHA256
ECDHE_RSA_AES128_GCM_SHA256
ECDHE_RSA_CHACHA20_SHA2156
ECDHE_ECDSA_AES256_GCM_SHA384
[…]

Curves
secp256r1
secp384r1
secp521r1

Extensions
Extended Master Secret
Encrypt then MAC
OCSP Status Request
Server Name (SNI)
[…]

Ciphers
ECDHE_ECDSA_AES128_GCM_SHA256
ECDHE_ECDSA_AES128_GCM_SHA386
ECDSA_CAMELLIA_128_GCM_SHA256
ECDSA_CAMELLIA_128_GCM_SHA384
[…]

Curves
secp256r1
secp384r1
secp521r1
secp224r1
secp192r1

Offhand, seems easy to detect

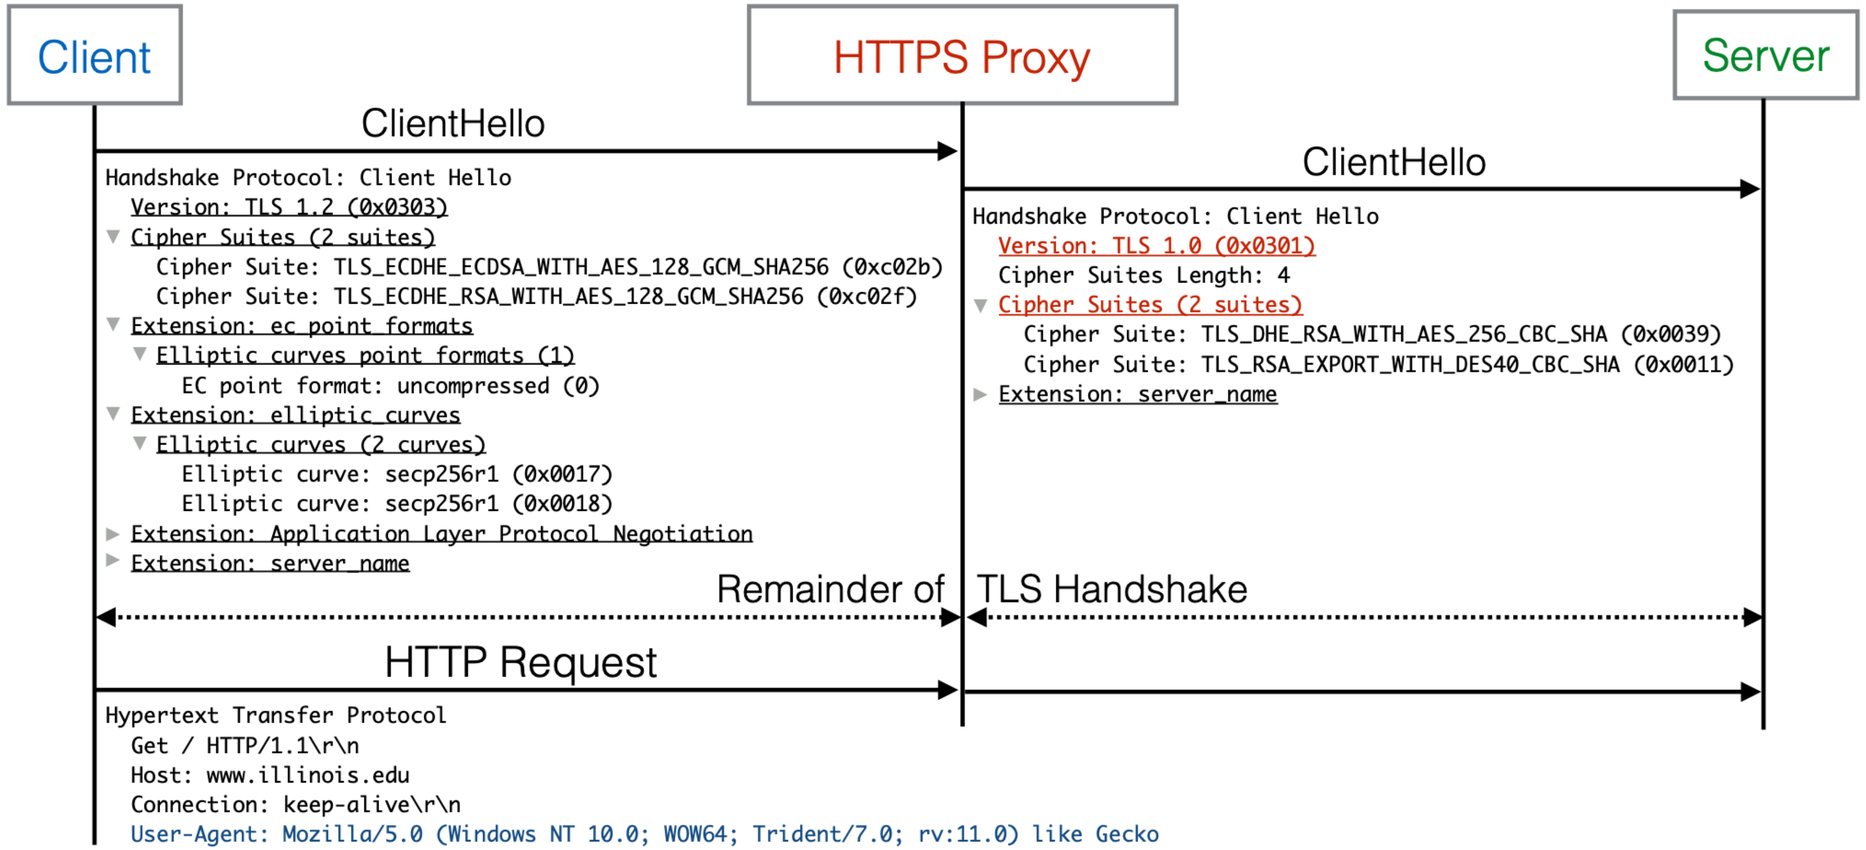

Investigating Common Products

We analyzed the TLS Client Hello messages from popular browsers browsers, middle boxes, client security software, and malware

Every product we investigated produced a unique TLS Client Hello message

Not always possible to identify product based on the handshake, but possible to detect whether a handshake is incompatible with a given browser

Browser Fingerprintability

Firefox — highly consistent (static) across platforms and OSes. NSS very different from OpenSSL/Boring/…

Chrome — optimization logic changes handshakes based on HW and OS. BoringSSL looks a bit like OpenSSL.

Safari — behavior based on OS lib not browser version

Internet Explorer — administrators can enable new ciphers, disable default ciphers, and arbitrarily reorder cipher suites through Windows Group Policy. SChannel behaves differently depending on both Windows updates and browser version

Deploying Heuristics

We deployed our heuristics for one week at three large service providers:

- Mozilla Firefox Update Servers
- Cloudflare CDN (0.5% sample all traffic)
- Popular E-commerce Site and Cloud Provider

Overall Interception Rates

We find a varying amount of interception between vantage points:

	No Interception	Likely Interception	Confirmed Interception
Cloudflare	88.6%	0.5%	10.9%
Firefox	96.0%	0.0%	4.0%
E-Commerce	92.9%	0.9%	6.2%

Inconsistencies Used

Detection Method	Firefox	E-commerce	Cloudflare
Invalid Extensions	16.8%	85.6%	89.0%
Invalid Ciphers	98.1%	54.2%	68.7%
Invalid Version	–	2.0%	–
Invalid Curves	–	5.5%	9.4%
Invalid Extension Order	87.7%	33.9%	40.4%
Invalid Cipher Order	98.8%	21.2%	21.1%
Missing Required Ext.	97.9%	91.1%	50.9%
Injected HTTP Header	–	14.0%	–

82% of all intercepted connections indicated support for the heartbeat extension

98% of intercepted Firefox connections were found based on the inclusion of ciphers never implemented in NSS

Bias By OS (E-Commerce)

OS	All Traffic	Intercepted	Of Intercepted
Windows 7	23.3%	9.6%	56.6%
Windows 10	22.5%	9.3%	14.3%
iOS	17.3%	0.1%	1.1%
Mac OS	15.8%	2.1%	6.5%
Android	9.4%	1.0%	0.5%
Windows 8.1	6.9%	8.3%	15.8%

Bias By Browser (E-Commerce)

Browser	All Traffic	Intercepted	Of Intercepted
Chrome	40.3%	8.6%	56.2%
Explorer	16.8%	7.4%	19.6%
Firefox	13.5%	8.4%	18.2%
Safari	10.2%	2.1%	3.4%
Chromium	7.6%	0.1%	0.1%
Mobile Safari	7.6%	0.9%	1.1%
Other	4.0%	4.0%	2.4%

Firefox Interception

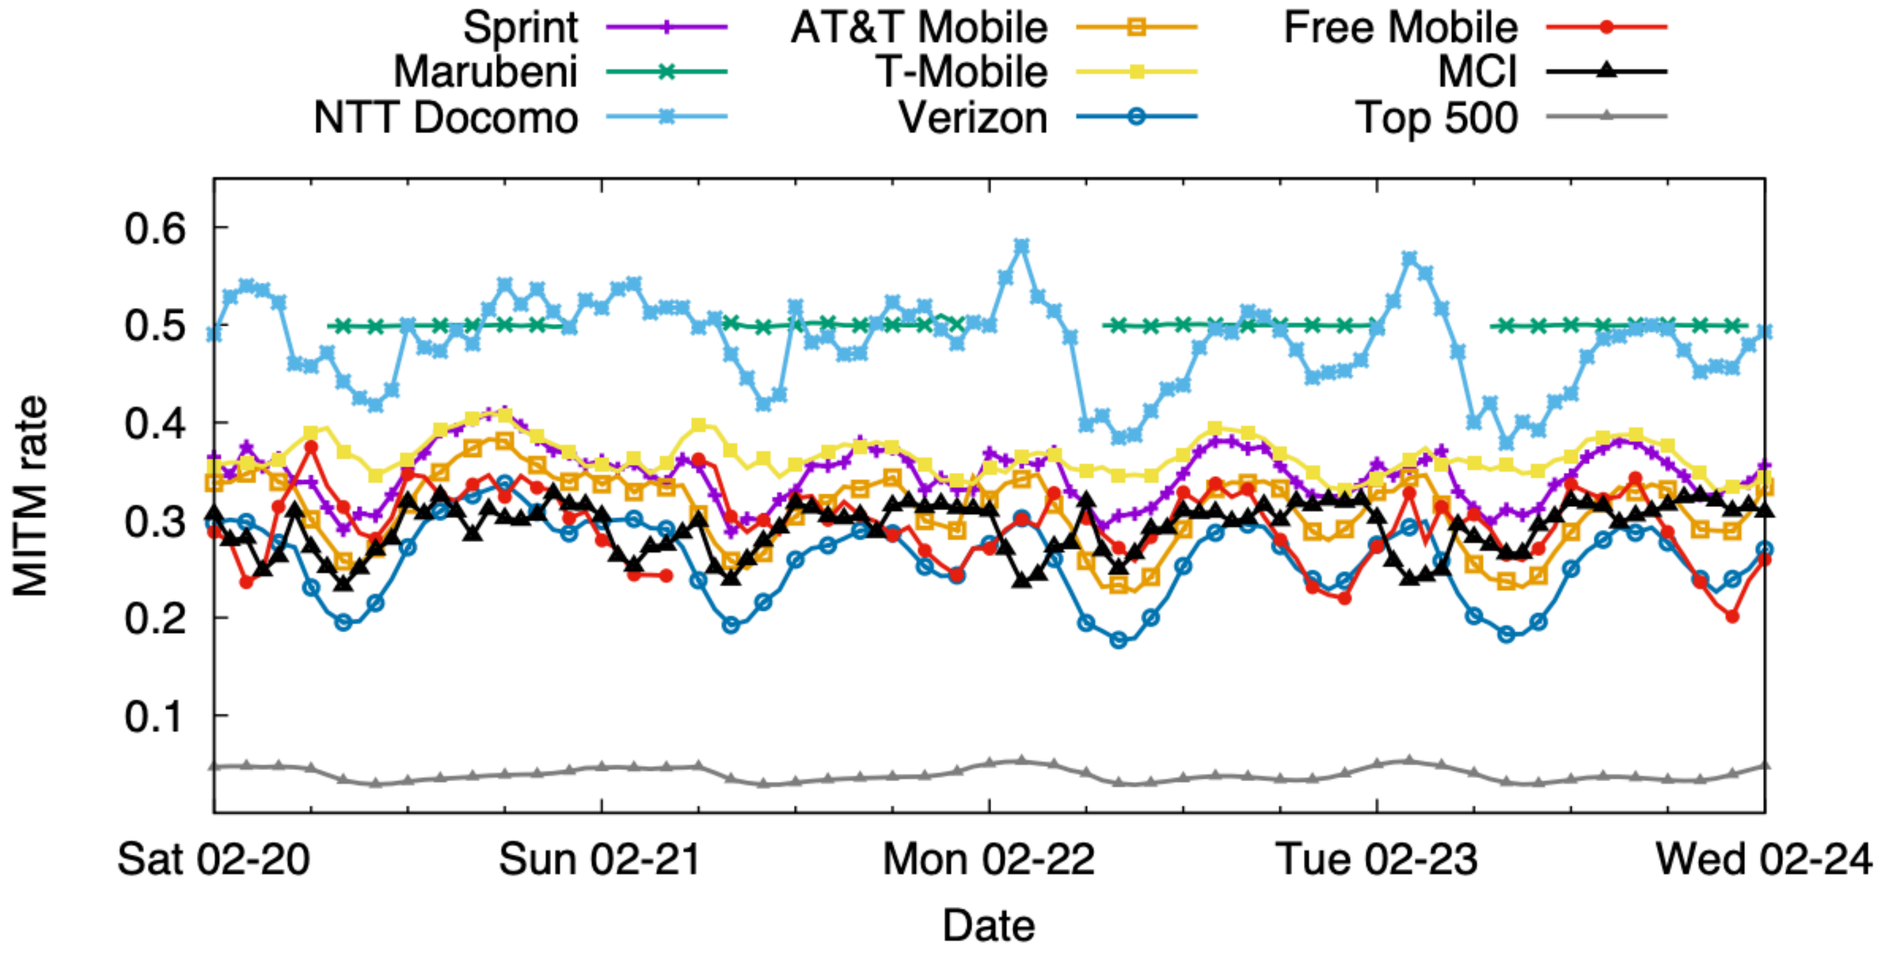

Countries with Firefox Interception

Country	MITM %	Country	MITM %
Guatemala	15.0%	Kiribati	8.2%
Greenland	9.9%	Iran	8.1%
South Korea	8.8%	Tanzania	7.3%
Kuwait	8.5%	Bahrain	7.3%
Qatar	8.4%	Afghanistan	6.7%

Fingerprints of Interceptors

	Fingerprint Description	% Total
E-commerce	*Unknown*	17.1%
	Avast Antivirus	10.8%
	Unknown	9.4%
	Blue Coat	9.1%
	Unknown	8.3%
Cloudflare	Avast Antivirus	9.1%
	AVG Antivirus	7.0%
	Unknown. Likely AV; mainly Windows 10/Chrome 47	6.5%
	Kaspersky Antivirus	5.0%
	BitDefender Antivirus	3.1%
Firefox	Bouncy Castle (Android 5)	26.3%
	Bouncy Castle (Android 4)	21.6%
	Unknown. Predominantly India	5.0%
	ESET Antivirus	2.8%
	Dr. Web Antivirus	2.6%

Overall Interception Rates

We find a varying amount of interception between

We estimate that 5-10% of all HTTPS connections are intercepted.

Firefox	96.0%	0.0%	4.0%
E-Commerce	92.9%	0.9%	6.2%

Measuring Security Impact

If interception products are performing high quality handshakes, there isn't an inherent security risk

We measured the security impact of interception by grading the security features advertised by the intercepted connection and the original browser

Quantifying Security Impact

We defined a security grading scale base on parameters advertised in Client Hello

Applied to original browsers and the connections we observed in the wild

Grading Scale	
A	**Optimal.** Equivalent to a modern web browser (e.g., Chrome)
B	**Suboptimal.** Non-ideal but not vulnerable to attacks
C	**Known Attack.** Vulnerable to known attack (e.g., RC4)
F	**Severely Broken.** An attacker could easily intercept connection

Security Grade Example

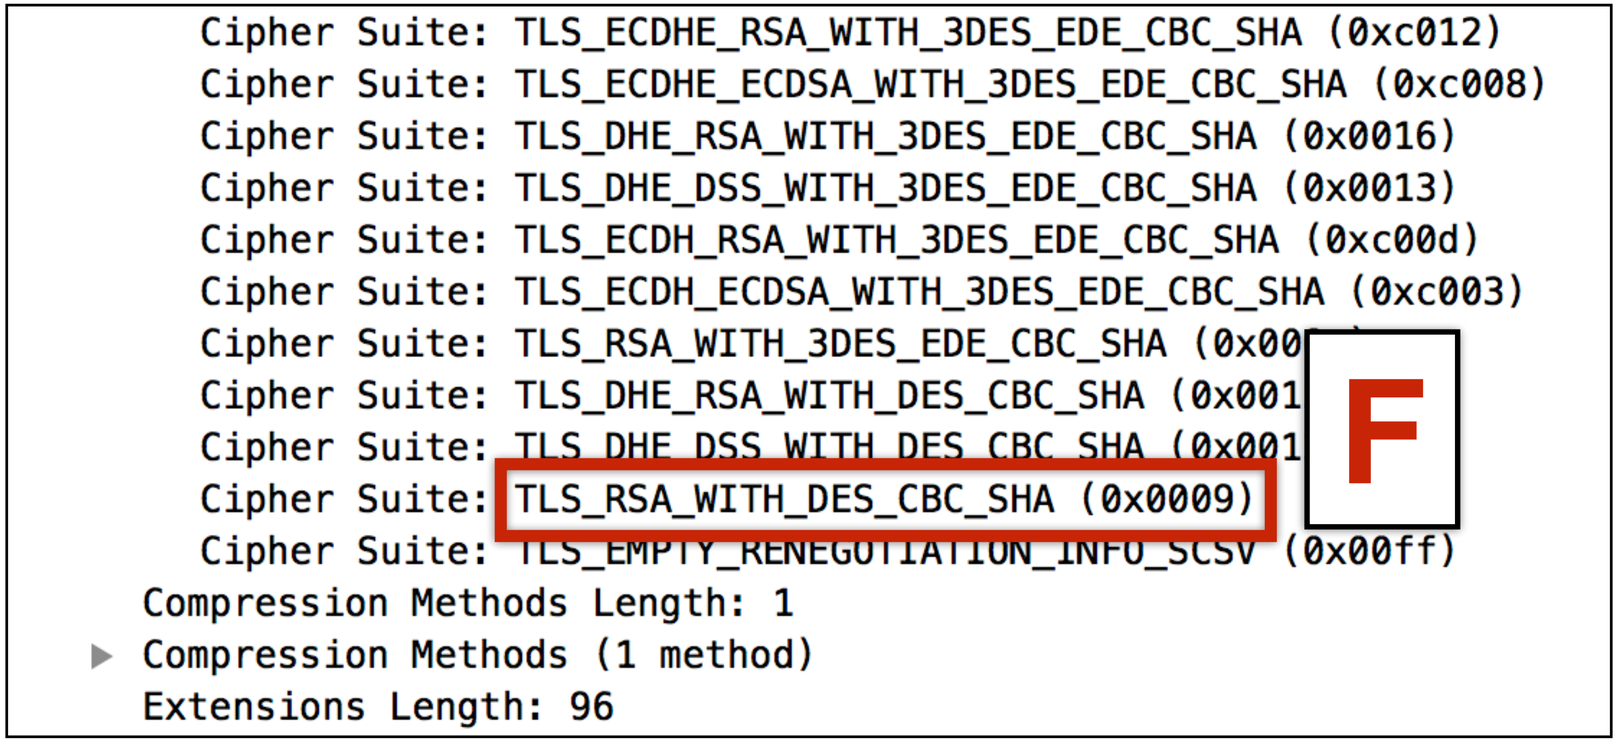

Security Impact of Interception

	Increased Security	Decreased Security	Severely Broken
E-Commerce	4%	27%	18%
Cloudflare	14%	45%	16%
Firefox Updates	0%	66%	37%

Security Impact of Interception

	Increased Security	Decreased Security	Severely Broken
E-Commerce	4%	27%	**18%**
Cloudflare	14%	45%	**16%**
Firefox Updates	0%	66%	**37%**

Middlebox Security

Network middleboxes have a worse security profile than client-side software

62% of connections are less secure

58% are severely broken

```
x-forwarded-for:
192.168.15.56

x-bluecoat-via:
abce6cd5a6733123
```

Why is security suffering?

Investigating Products

We investigated the default configurations of popular interception products:

- Popular middleboxes (e.g., A10, Bluecoat, Cisco)
- Antivirus software (e.g., Avast, AVG, Kaspersky)

We ran a series of automated tests against products. This was ***far from*** in-depth testing (guarantee they all have additional vulns. if you look. Tavis started to…)

Security Profile of Interception Products

	Increased Security	Same Security	Decreased Security	Severely Broken
Client Security Products	0/20	2/20	18/20	10/20
Middleboxes	0/12	1/12	6/12	5/12

Middlebox Security

Product	Grade	Validates Certificates	Modern Ciphers	Advertises RC4	TLS Version	Grading Notes
A10 vThunder SSL Insight	F	✓	✓	Yes	1.2	Advertises export ciphers
Blue Coat ProxySG 6642	A*	✓	✓	No	1.2	Mirrors client ciphers
Barracuda 610Vx Web Filter	C	✓	✗	Yes	1.0	Vulnerable to Logjam attack
Checkpoint Threat Prevention	F	✓	✗	Yes	1.0	Allows expired certificates
Cisco IronPort Web Security	F	✓	✓	Yes	1.2	Advertises export ciphers
Forcepoint TRITON AP-WEB Cloud	C	✓	✓	No	1.2	Accepts RC4 ciphers
Fortinet FortiGate 5.4.0	C	✓	✓	No	1.2	Vulnerable to Logjam attack
Juniper SRX Forward SSL Proxy	C	✓	✗	Yes	1.2	Advertises RC4 ciphers
Microsoft Threat Mgmt. Gateway	F	✗	✗	Yes	SSLv2	No certificate validation
Sophos SSL Inspection	C	✓	✓	Yes	1.2	Advertises RC4 ciphers
Untangle NG Firewall	C	✓	✗	Yes	1.2	Advertises RC4 ciphers
WebTitan Gateway	F	✗	✓	Yes	1.2	Broken certificate validation

Product	OS	Browser MITM				Grade	Validates Certificates	Modern Ciphers	TLS Version	Grading Notes
		IE	Chrome	Firefox	Safari					
Avast ...										
AV 11	Win	●	○	○		A*	✓	✓	1.2	Mirrors client ciphers
AV 11.7	Mac		●	●	●	F	✓	✓	1.2	Advertises DES
AVG ...										
Internet Security 2015–6	Win	●	●	○		C	✓	✓	1.2	Advertises RC4
Bitdefender ...										
Internet Security 2016	Win	●	●	●		C	✓	○	1.2	RC4, 768-bit D-H
Total Security Plus 2016	Win	●	●	●		C	✓	○	1.2	RC4, 768-bit D-H
AV Plus 2015–16	Win	●	●	●		C	✓	○	1.2	RC4, 768-bit D-H
Bullguard ...										
Internet Security 16	Win	●	●	●		A*	✓	✓	1.2	Mirrors client ciphers
Internet Security 15	Win	●	●	●		F	✓	✗	1.0	Advertises DES
CYBERsitter ...										
CYBERsitter 11	Win	●	●	●		F	✗	✗	1.2	No cert. validation, DES
Dr. Web ...										
Security Space 11	Win	●	●	●		C	✓	○	1.2	RC4, FREAK
Dr. Web 11 for OS X	Mac		●	●	●	F	✓	✗	1.0	Export ciphers, DES, RC2
ESET ...										
NOD32 AV 9	Win	●	●	●		F	○	○	1.2	Broken cert. validation
Kaspersky ...										
Internet Security 16	Win	●	●	●		C	✓	✓	1.2	CRIME vulnerability
Total Security 16	Win	●	●	●		C	✓	✓	1.2	CRIME vulnerability
Internet Security 16	Mac		●	●	●	C	✓	✓	1.2	768-bit D-H
KinderGate ...										
Parental Control 3	Win	●	●	●		F	○	✗	1.0	Broken cert. validation
Net Nanny ...										
Net Nanny 7	Win	●	●	●		F	✓	✓	1.2	Anonymous ciphers
Net Nanny 7	Mac		●	●	●	F	✓	✓	1.2	Anonymous ciphers
PC Pandora ...										
PC Pandora 7	Win	●	◐	◐		F	✗	✗	1.0	No certificate validation
Qustodio ...										
Parental Control 2015	Mac		●	●	●	F	✓	✓	1.2	Advertises DES

u643384 Reporter
Description • 2 years ago

User Agent: Mozilla/5.0 (X11; Fedora; Linux x86_64; rv:68.0) Gecko/20100101 Firefox/68.0

Steps to reproduce:

Since today all Internet providers in Kazakhstan started MITM on all encrypted HTTPS traffic.

They asked end-users to install government-issued certificate authority on all devices in every browser: http://qca.kz/

Actual results:

MITM attack: https://i.imgur.com/rFEjXKw.jpg

Message from Internet provider, requires to install this CA: https://i.imgur.com/WyKjOug.jpg

Proofs: https://atlas.ripe.net/measurements/22372655/#!probes

Official site with root CA: http://qca.kz/

Links to certificates:
http://qca.kz/qazca.cer
http://qca.kz/qazca.pem
http://qca.kz/qazca.der

Expected results:

I think this CA should be blacklisted by Mozilla and Firefox should not accept it at all even user installed it manually.

This will save privacy of all Internet users in Kazakhstan.

Starting on July 17, 2019, Kazakhstan launched an HTTPS interception man-in-the-middle (MitM) attack, after instructing citizens to install a government-issued root certificate on all devices and in every browser for “security” purposes.

Our findings show that only a fraction of the Internet traffic inside the country was subject to interception (around 7–24% of the 6,736 TLS hosts measured were affected), and that the path to all of the servers affected by the interception passed through two sets of specific hops in AS9198 (Kazakhtelecom). Of the Alexa Top 10,000 domains, 37 triggered interception.