Imperfect Forward Secrecy

David Adrian
@davidcadrian

UNIVERSITY OF
MICHIGAN

Who am I?

I'm **David Adrian**, graduate student at the University of Michigan

An Academic.

"Halfademic"

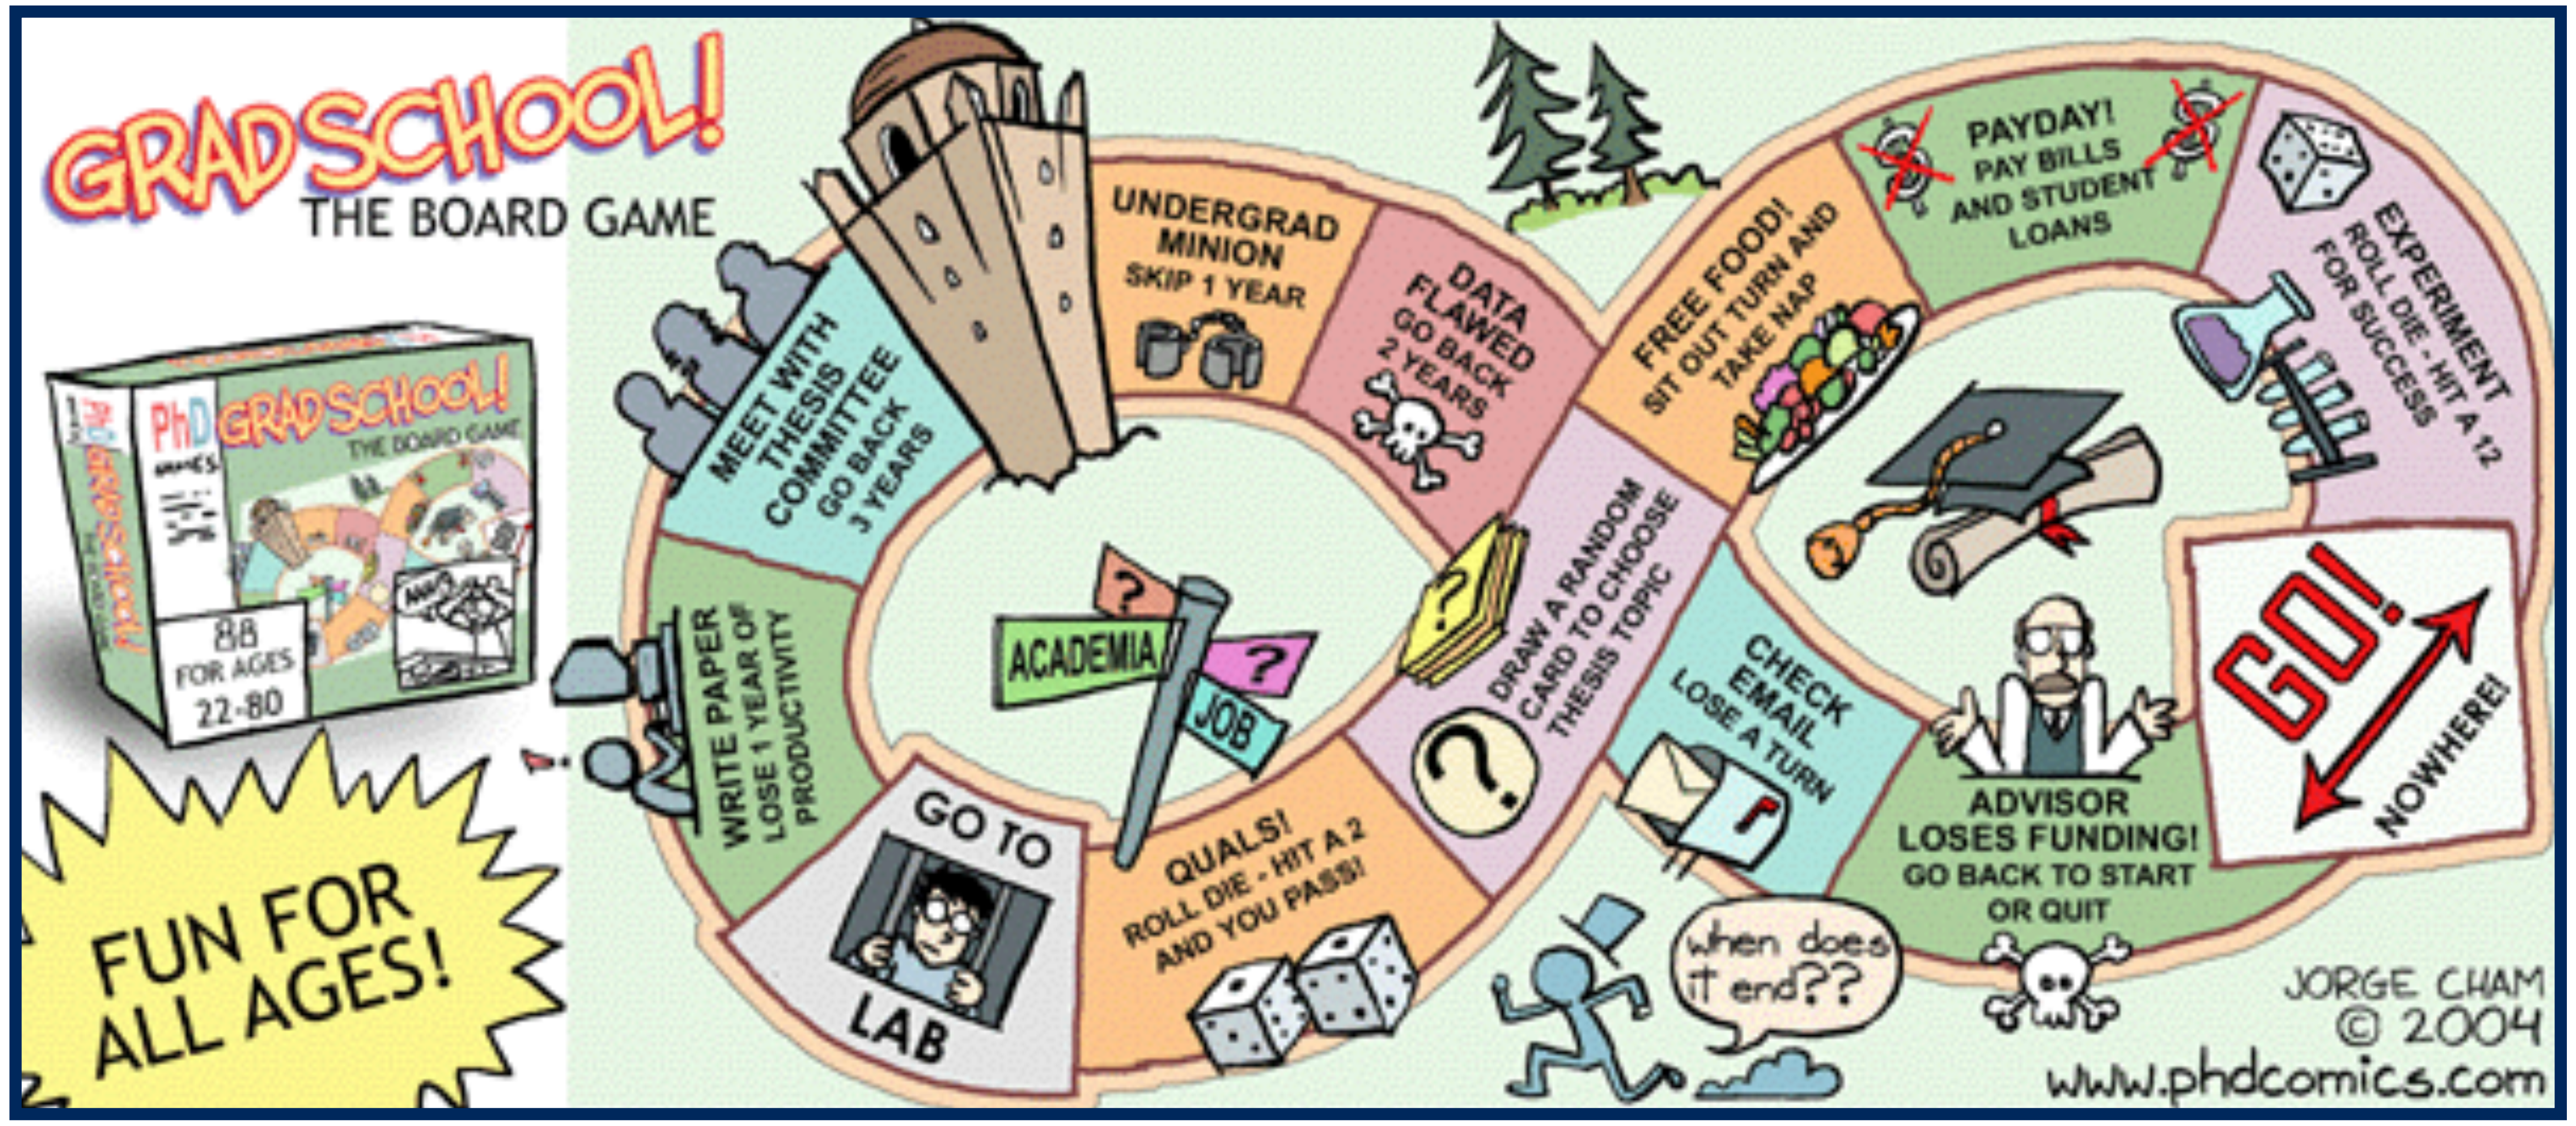

What do I do?

What is this?

Logjam

Internet-scanning?

Weak Diffie-Hellman

Mail security?

Imperfect Forward Secrecy: How Diffie-Hellman Fails in Practice

David Adrian[¶] Karthikeyan Bhargavan[*] Zakir Durumeric[¶] Pierrick Gaudry[†] Matthew Green[§]
J. Alex Halderman[¶] Nadia Heninger[‡] Drew Springall[¶] Emmanuel Thomé[†] Luke Valenta[‡]
Benjamin VanderSloot[¶] Eric Wustrow[¶] Santiago Zanella-Béguelin[‖] Paul Zimmermann[†]

[*] INRIA Paris-Rocquencourt [†] INRIA Nancy-Grand Est, CNRS and Université de Lorraine
[‖] Microsoft Research [‡] University of Pennsylvania [§] Johns Hopkins [¶] University of Michigan

For additional materials and contact information, visit WeakDH.org.

ABSTRACT

We investigate the security of Diffie-Hellman key exchange as used in popular Internet protocols and find it to be less secure than widely believed. First, we present a novel flaw in TLS

logs in that group, amortizing the cost over all targets that share this parameter. The algorithm can be tuned to reduce individual log cost even further. Although this fact is well known among mathematical cryptographers, it seems to have

Diffie-Hellman Key Exchange

Diffie-Hellman Key Exchange

First published key-exchange algorithm

Two parties agree on a **shared secret key** over an unsecured channel

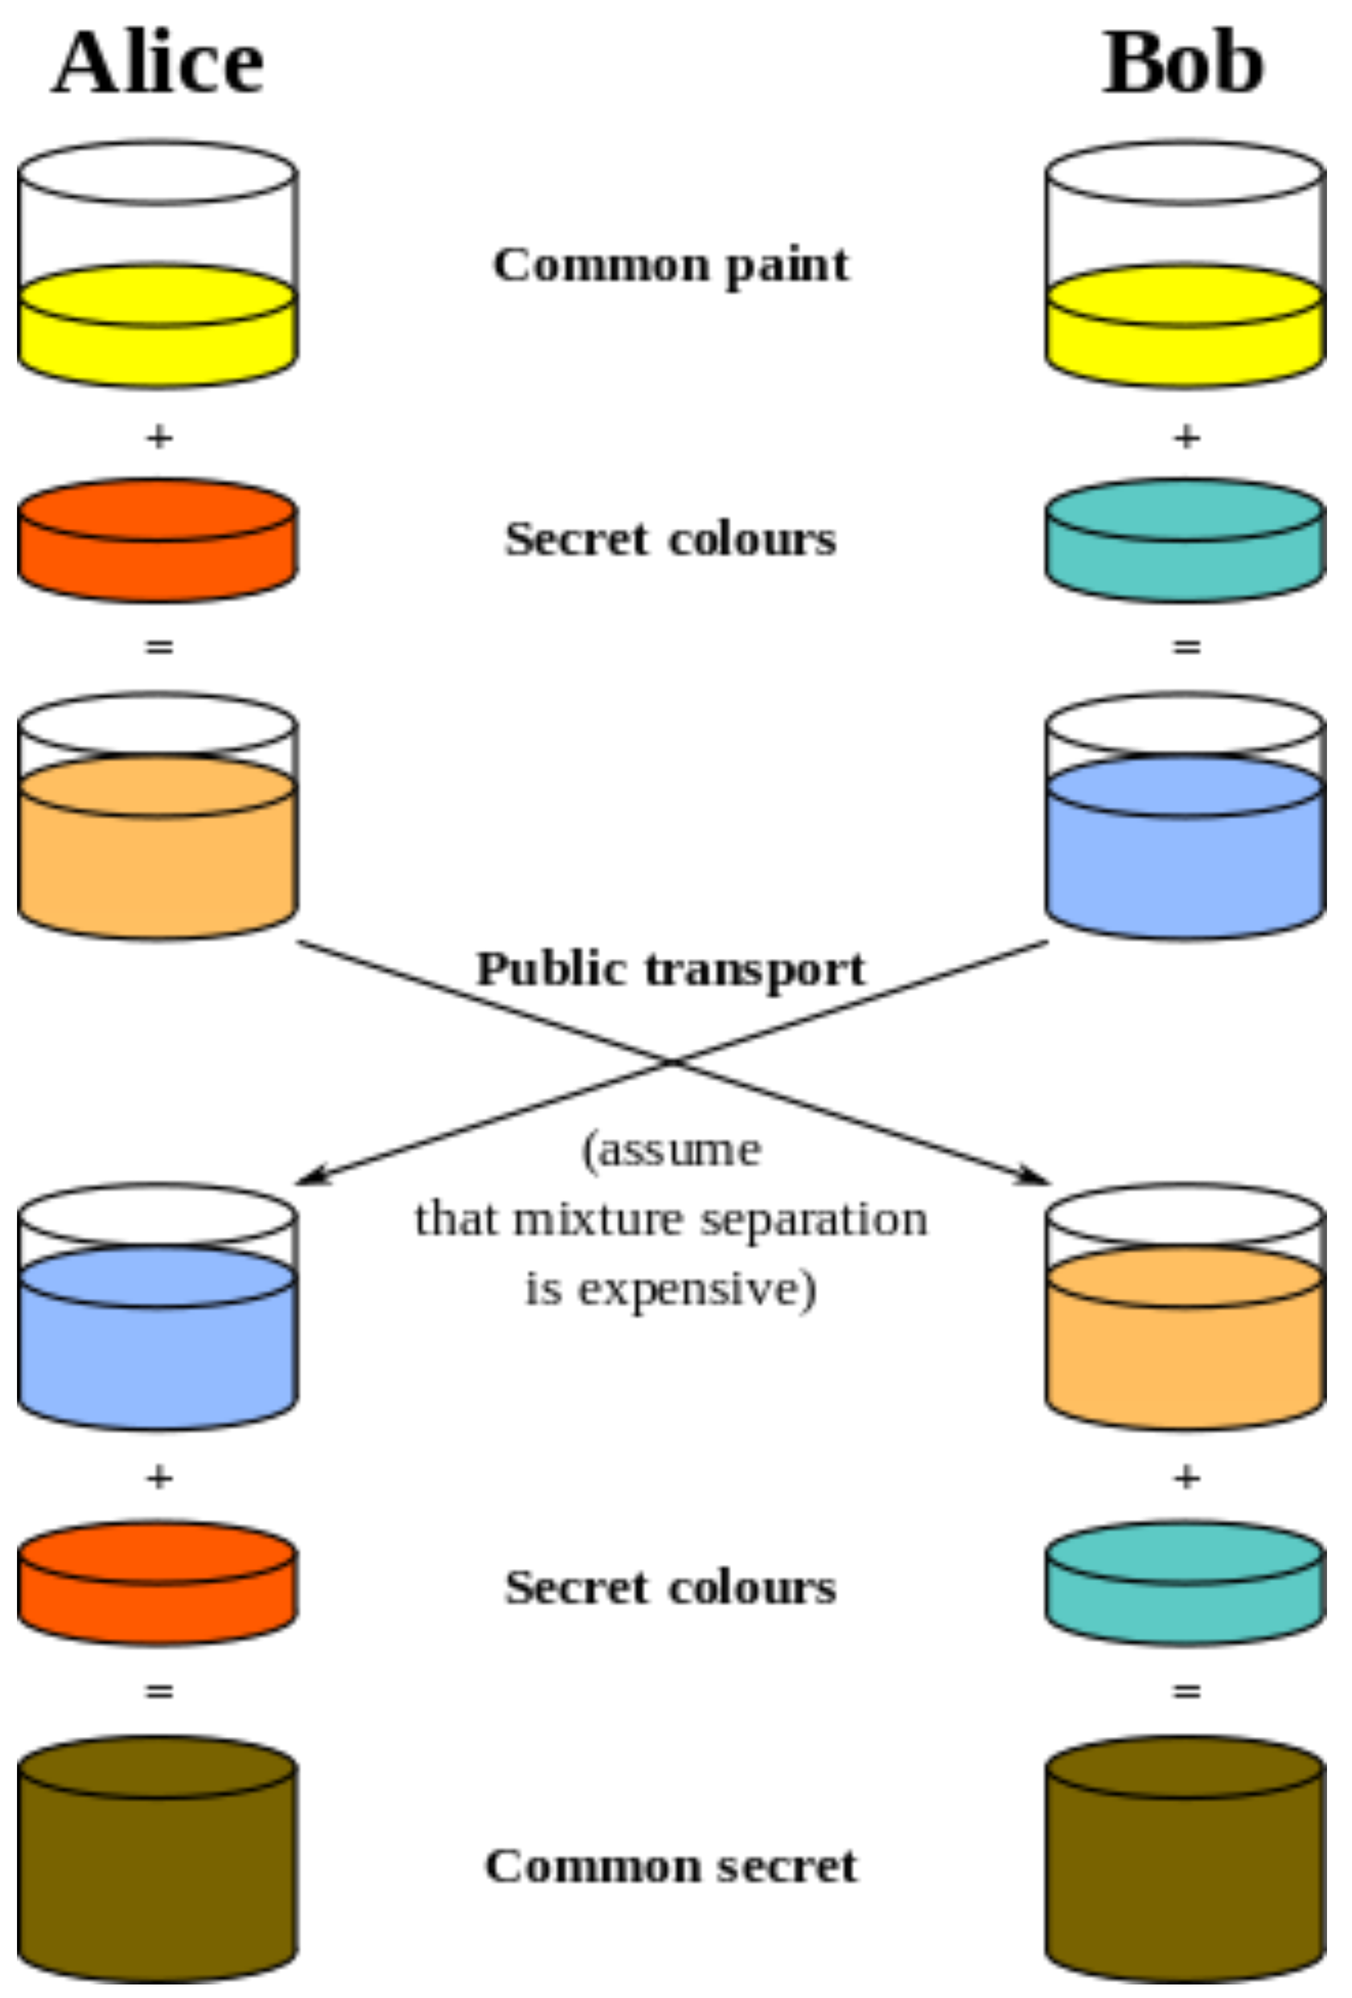

Public parameters

- p, a large prime
- g, a generator for a group modulo p

$g^a \bmod p$ →

← $g^b \bmod p$

$$\boxed{g^{ba} \bmod p == g^{ab} \bmod p}$$

Shortcomings

Unauthenticated *Fix by signing with a long-term key (certificate)*

How to pick *g* and *p*? *Standardize in protocol, or allow server to choose*

Perfect Forward Secrecy

When a new DIffie-Hellman key exchange is completed at the start of every connection, you gain **perfect forward secrecy**.

Breaking one-connection or the long-term key of a server, does not allow an adversary to decrypt past connections.

“Sites that use perfect forward secrecy can provide better security to users in cases where the encrypted data is being monitored and recorded by a third party.”

“With Perfect Forward Secrecy, anyone possessing the private key and a wiretap of Internet activity can decrypt nothing.”

Breaking Diffie-Hellman

Requires finding the solution to a “hard” mathematical problem called **discrete log**

Given $g^x \equiv y \bmod p$, compute x

Breaking Diffie-Hellman

Requires finding the solution to a “hard” mathematical problem called **discrete log**

Want

Given $g^x \equiv y \bmod p$, compute x

Have

Conceptually easy ✓, computationally hard ✗

Breaking Diffie-Hellman

State-of-the-art is the **number-field sieve** algorithm

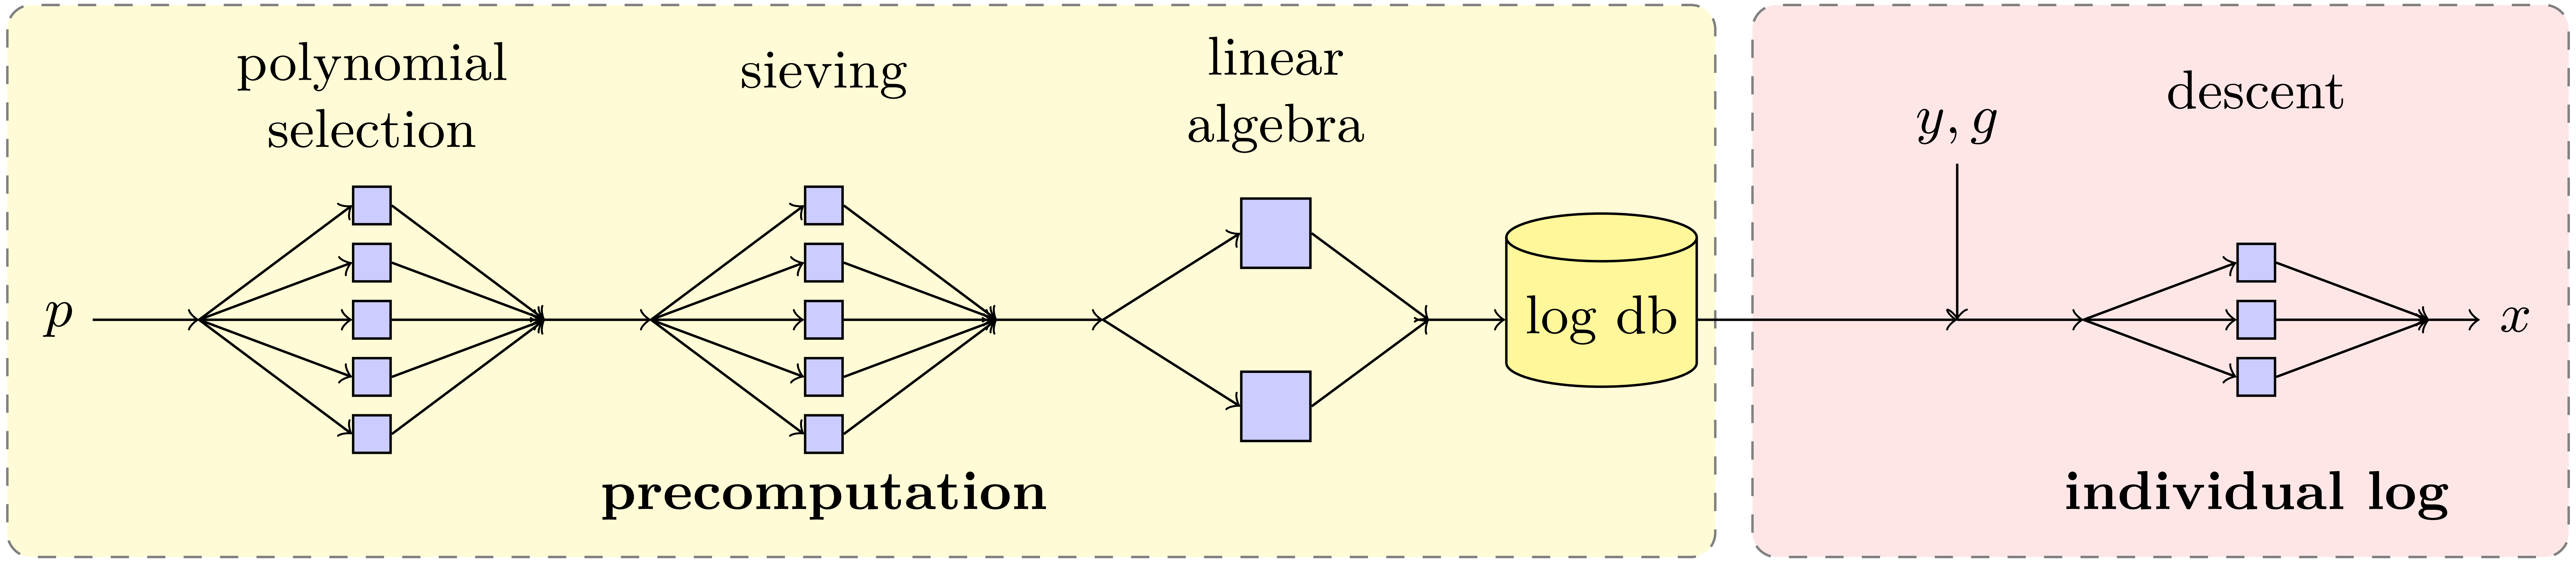

Number-Field Sieve

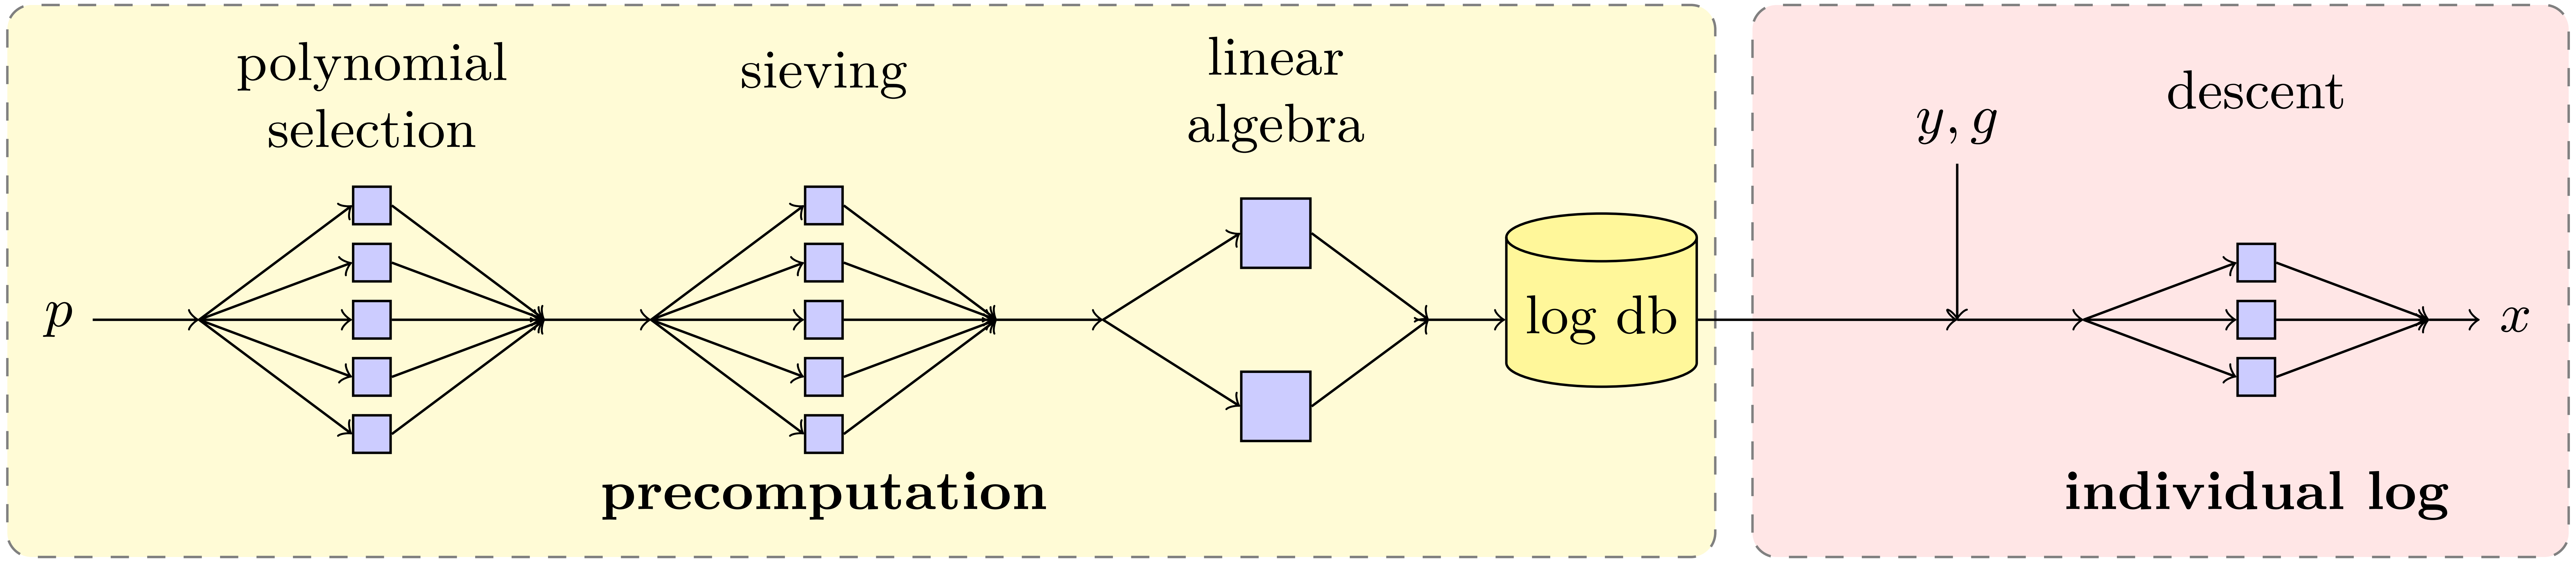

	Sieving	Linear Algebra	Descent
DH-512	2.5 core years	7.7 core years	10 core min.

Precomputation depends solely on *p*!

Logjam Attack on TLS


Google
https://www.google.com/?gws_rd=ssl
www.google.com
Your connection to this site is private.
Permissions
Connection
The identity of this website has been verified by Google Internet Authority G2. Valid Certificate Transparency information was supplied by the server.
Certificate Information
Your connection to www.google.com is encrypted using a modern cipher suite.
The connection uses TLS 1.2.
The connection is encrypted and authenticated using AES_128_GCM and uses ECDHE_RSA as the key exchange mechanism.
What do these mean?
Certificate
Cipher
David
Gmail
Images
Google Search
I'm Feeling Lucky

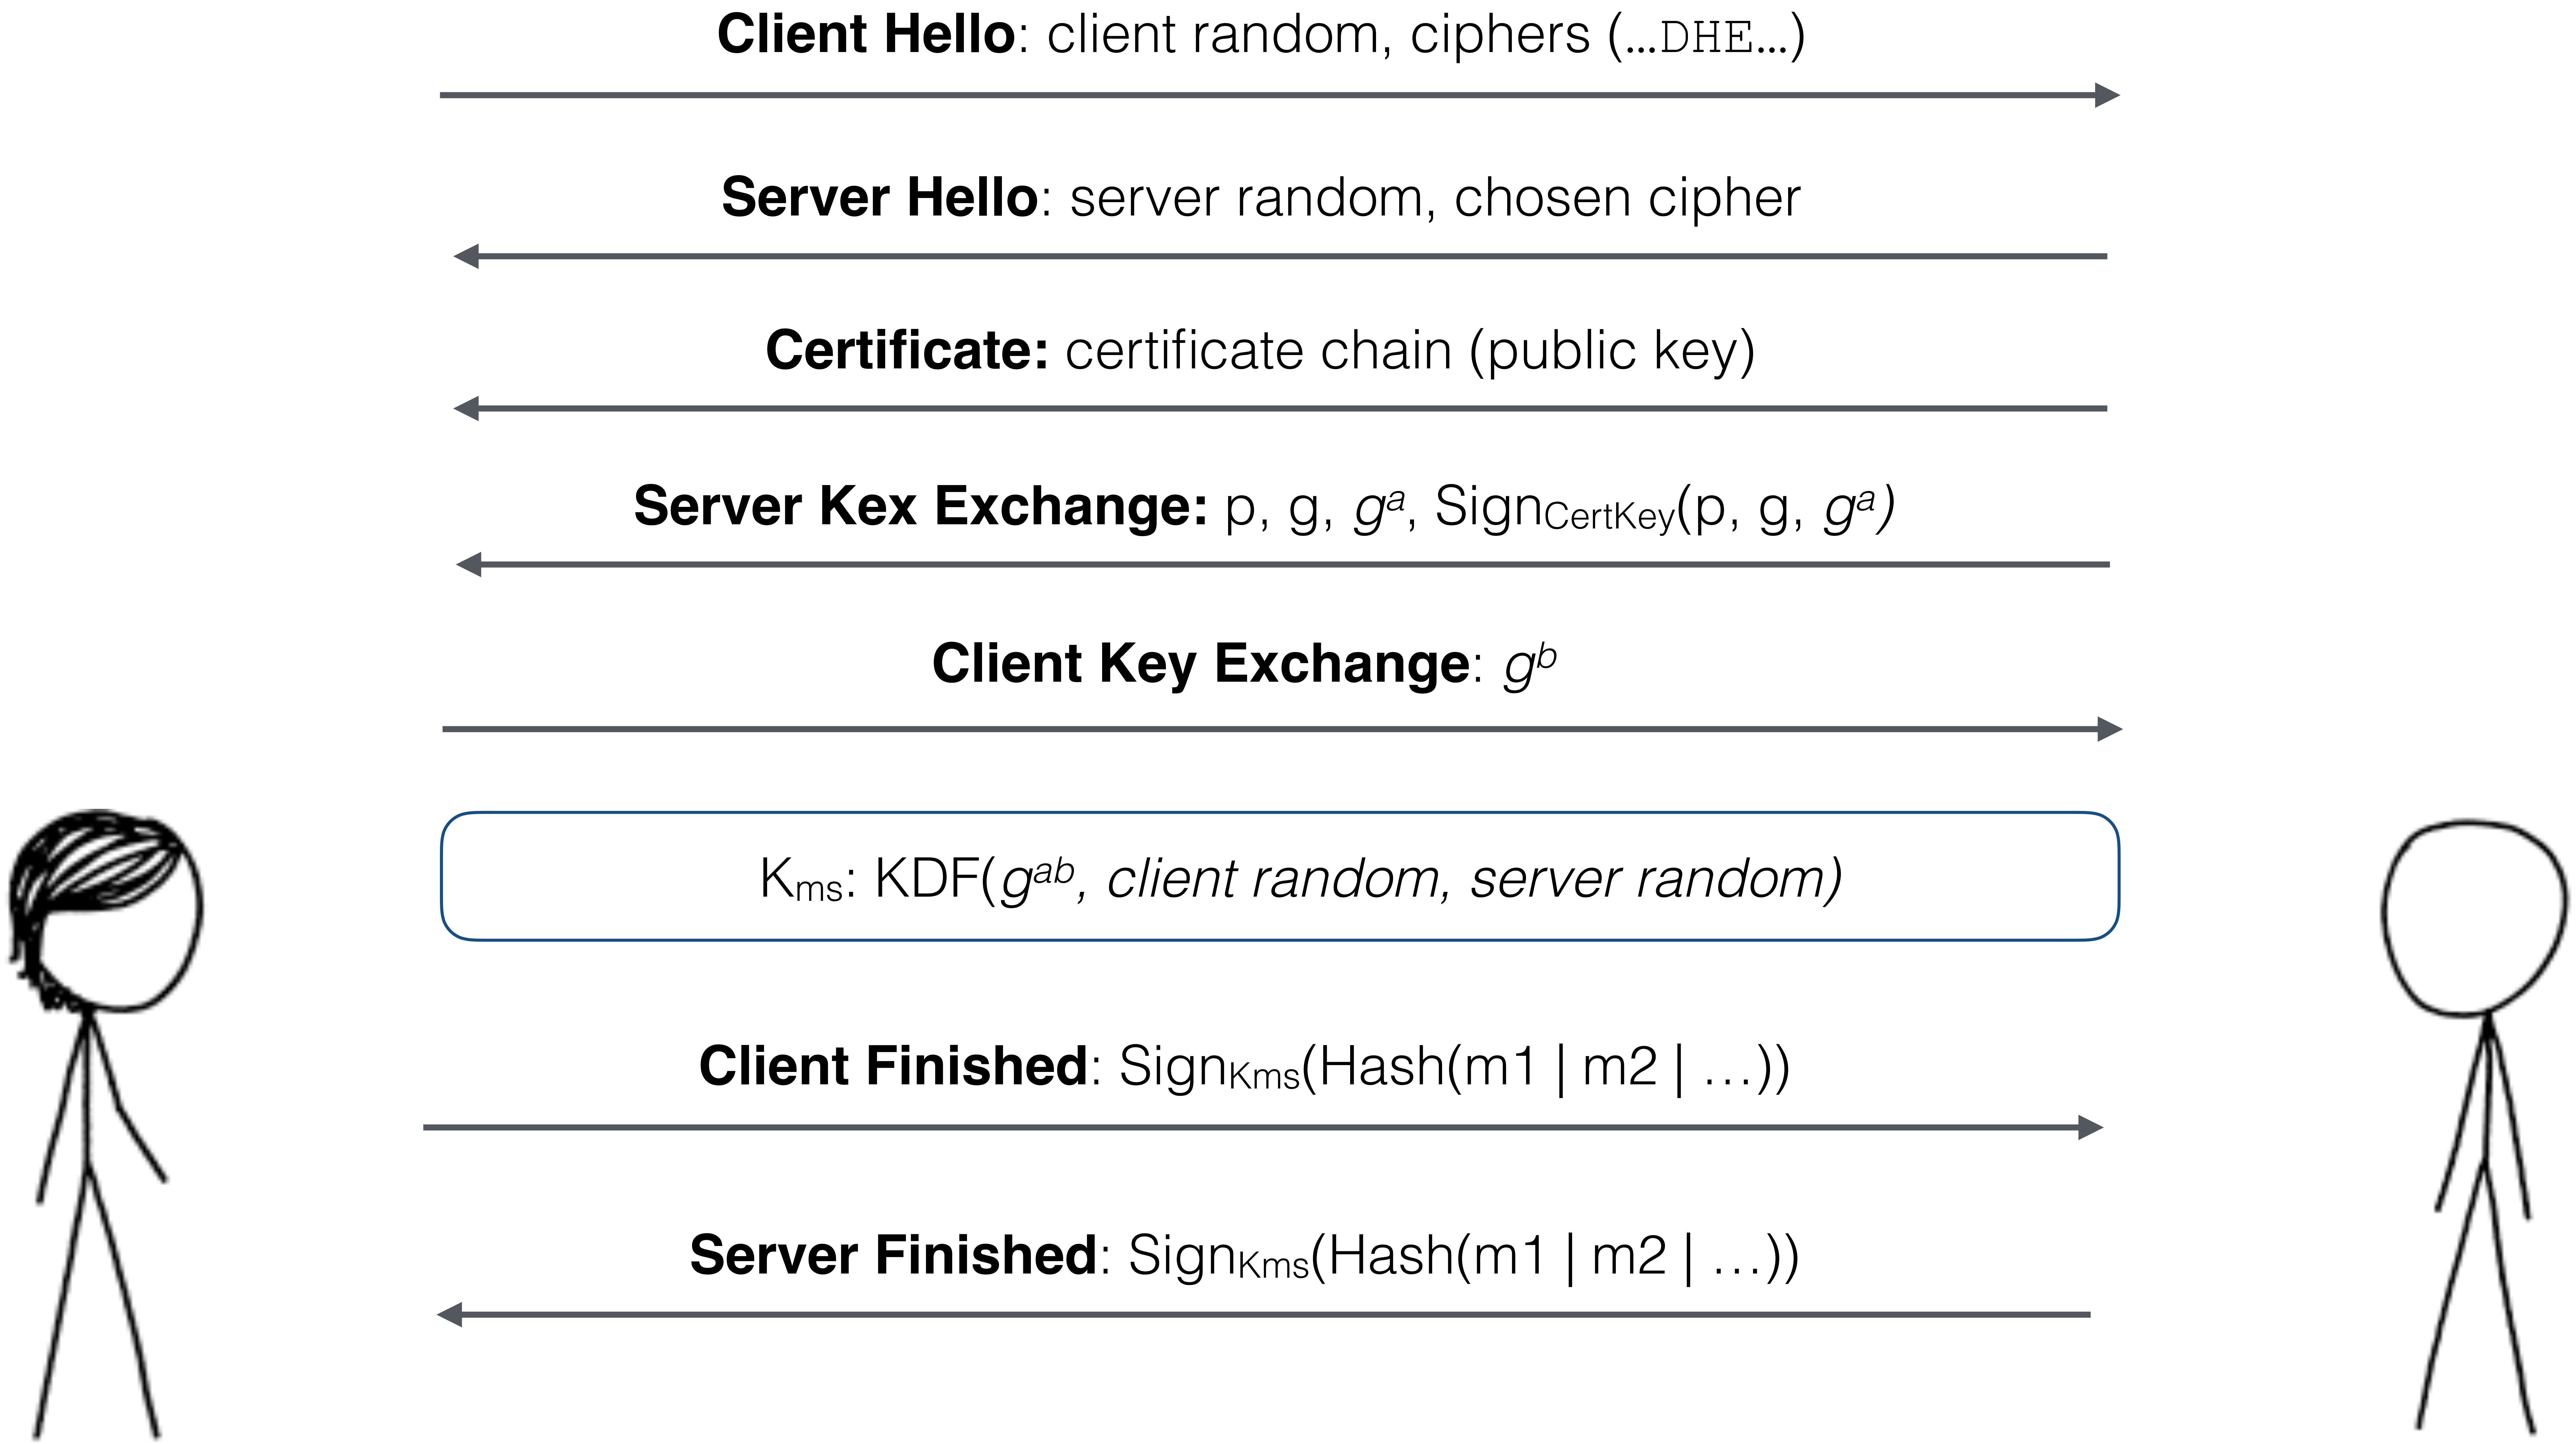

Client Hello: client random, ciphers (...DHE...)
Server Hello: server random, chosen cipher
Certificate: certificate chain (public key)
Server Kex Exchange: p, g, g^a, $\mathrm{Sign}_{\mathrm{CertKey}}(p, g, g^a)$
Client Key Exchange: g^b
K_{ms}: KDF(g^{ab}, *client random, server random)*
Client Finished: $\mathrm{Sign}_{K_{ms}}$(Hash(m1 | m2 | ...))
Server Finished: $\mathrm{Sign}_{K_{ms}}$(Hash(m1 | m2 | ...))

Export Ciphers in TLS

Remnant of the 90s “crypto wars”

It used to be illegal to export “strong crypto” outside of the United States, law was overturned in *Bernstein vs. United States of America*

TLS was designed before the law was overturned

Included weak (short-key) “export ciphers” for use outside of the United States, e.g. `DHE_EXPORT`

`DHE_EXPORT` uses 512-bit primes!

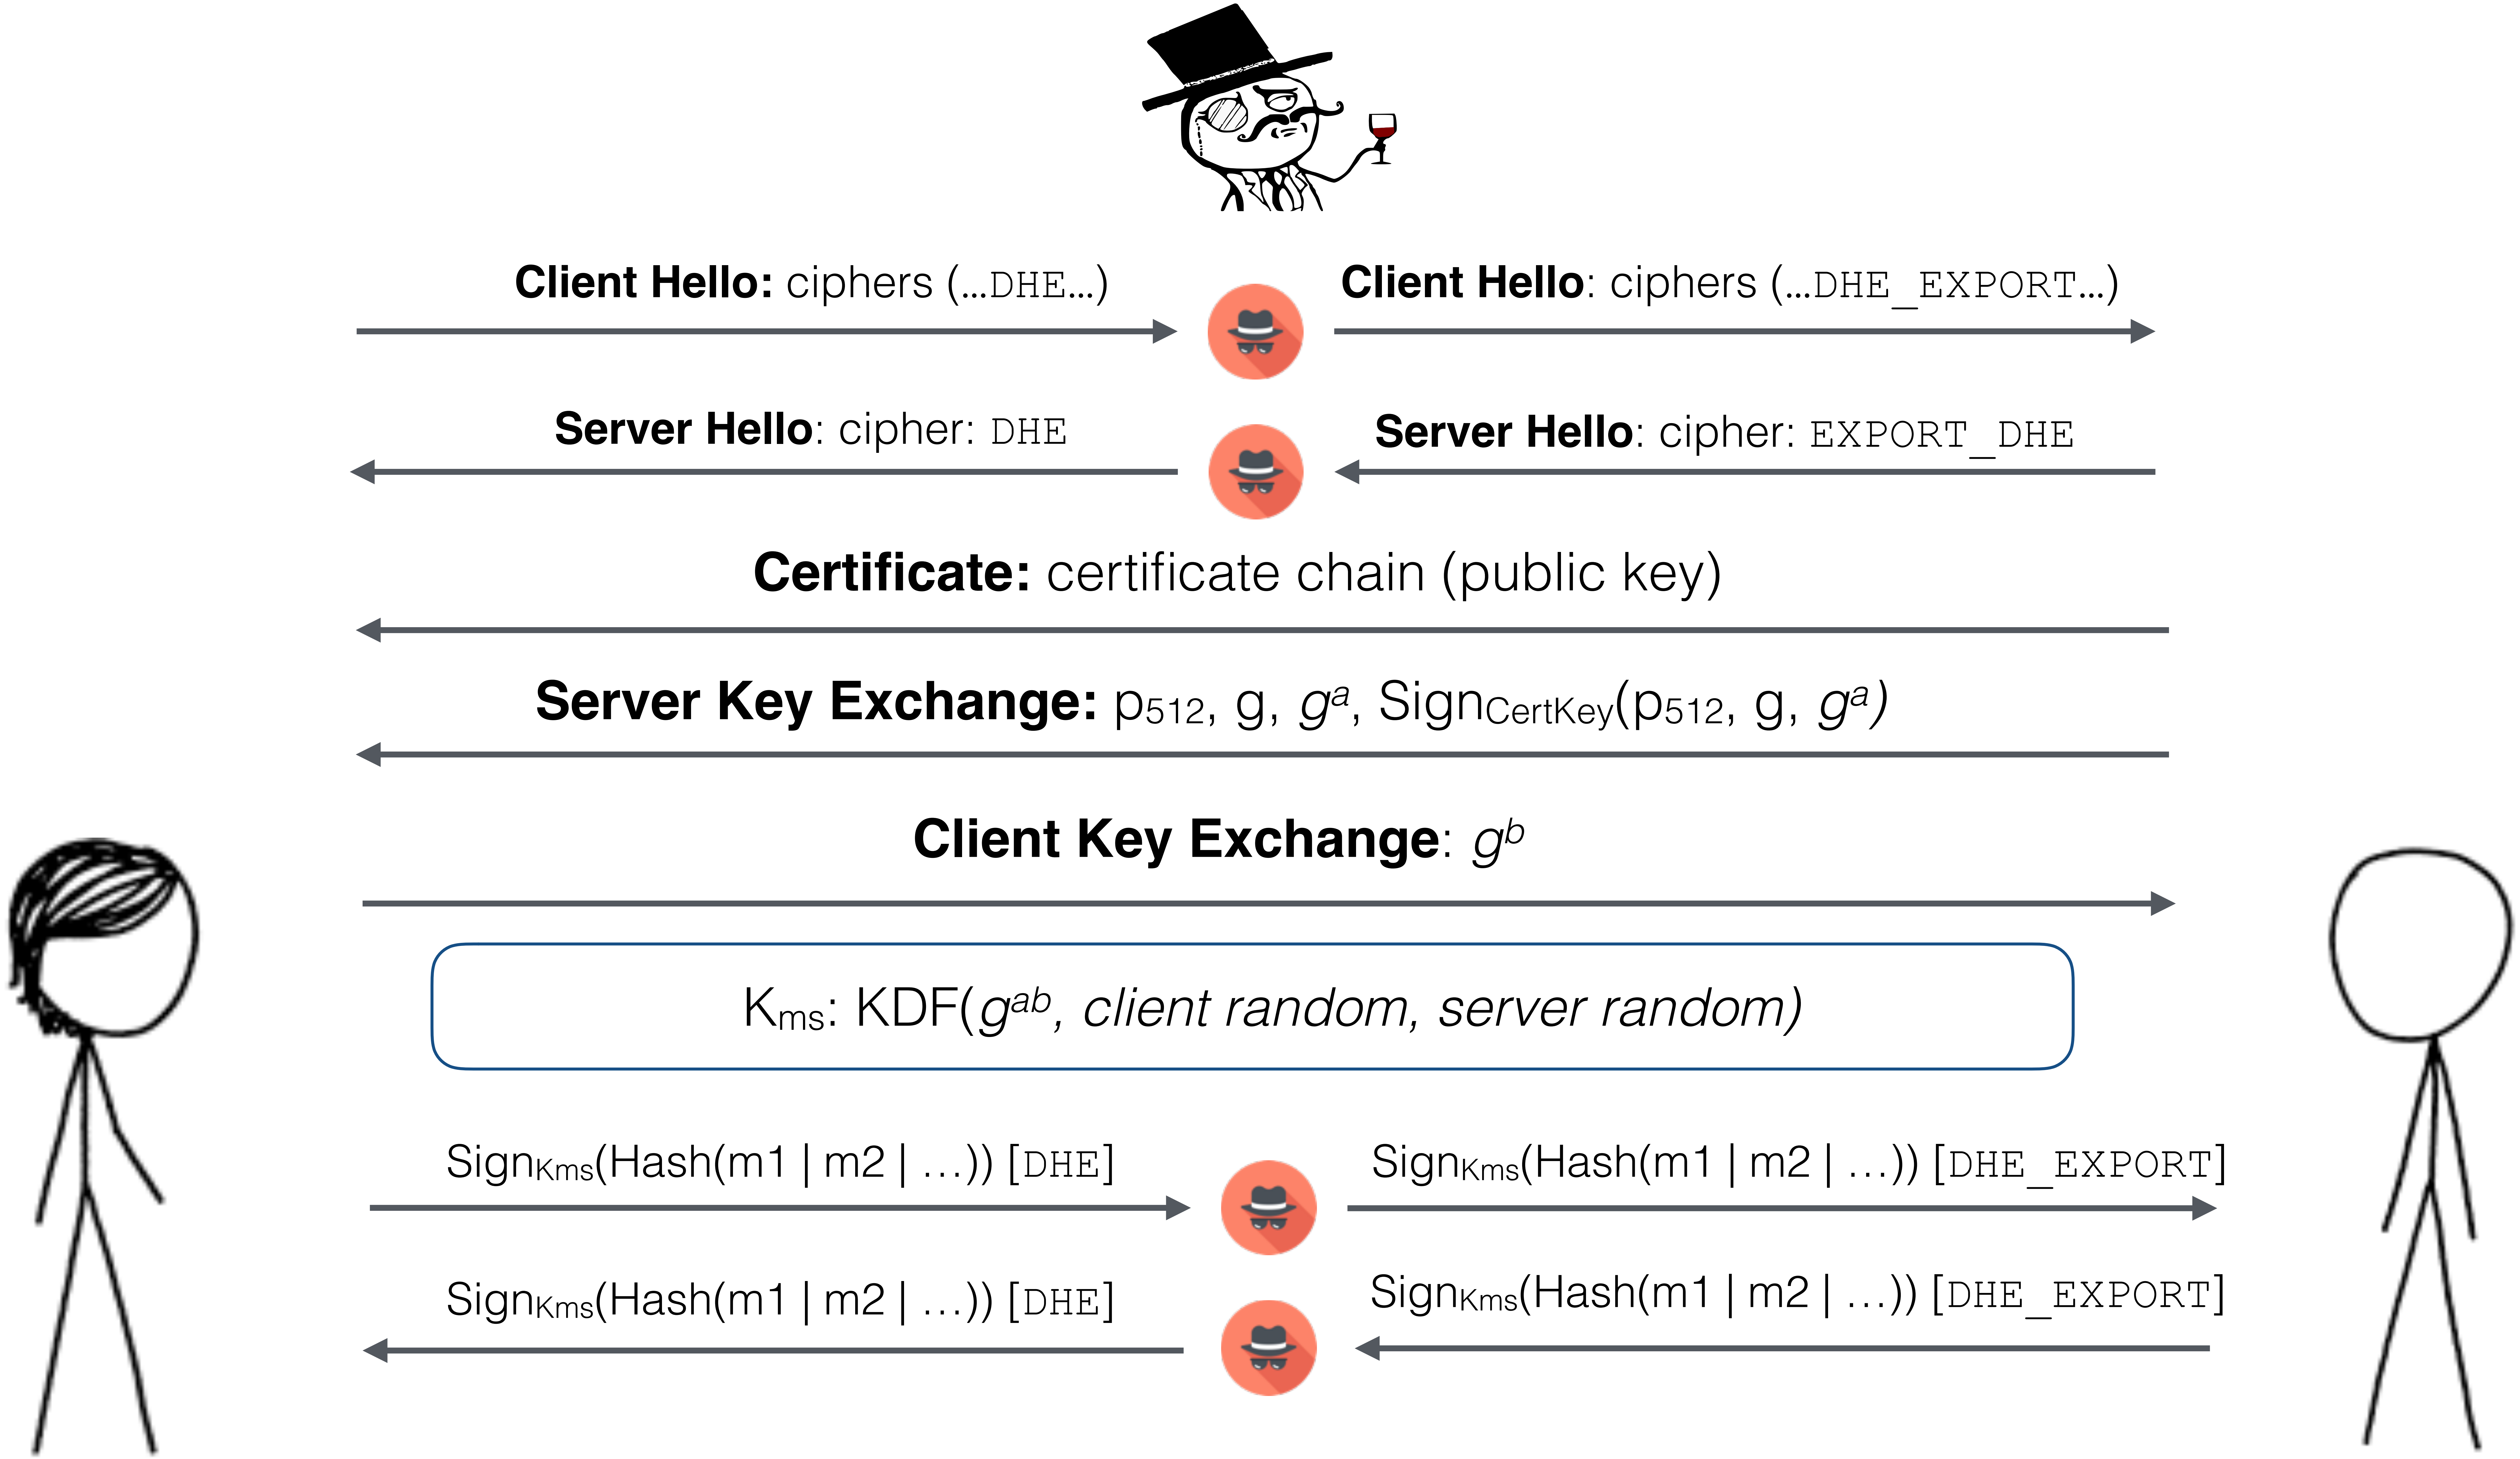
Client Hello: ciphers (...DHE...)
Client Hello: ciphers (...DHE_EXPORT...)
Server Hello: cipher: DHE
Server Hello: cipher: EXPORT_DHE
Certificate: certificate chain (public key)
Server Key Exchange: p_{512}, g, g^a, $\text{Sign}_{\text{CertKey}}(p_{512}, g, g^a)$
Client Key Exchange: g^b
K_{ms}: KDF(g^{ab}, client random, server random)
Sign_{Kms}(Hash(m1 | m2 | ...)) [DHE]
Sign_{Kms}(Hash(m1 | m2 | ...)) [DHE_EXPORT]
Sign_{Kms}(Hash(m1 | m2 | ...)) [DHE]
Sign_{Kms}(Hash(m1 | m2 | ...)) [DHE_EXPORT]

Support for Export Ciphers

8.5% of the Alexa Top 1M support `DHE_EXPORT`

Prime	Popularity
Apache mod_ssl	82%
nginx	10%
Other (463 primes)	8%

Breaking 512-bit

We did the precomputation for the two most popular 512-bit primes.

	polysel	sieving	linalg	descent
	2000-3000 cores		288 cores	36 cores
DH-512	3 hours	15 hours	120 hours	70 seconds

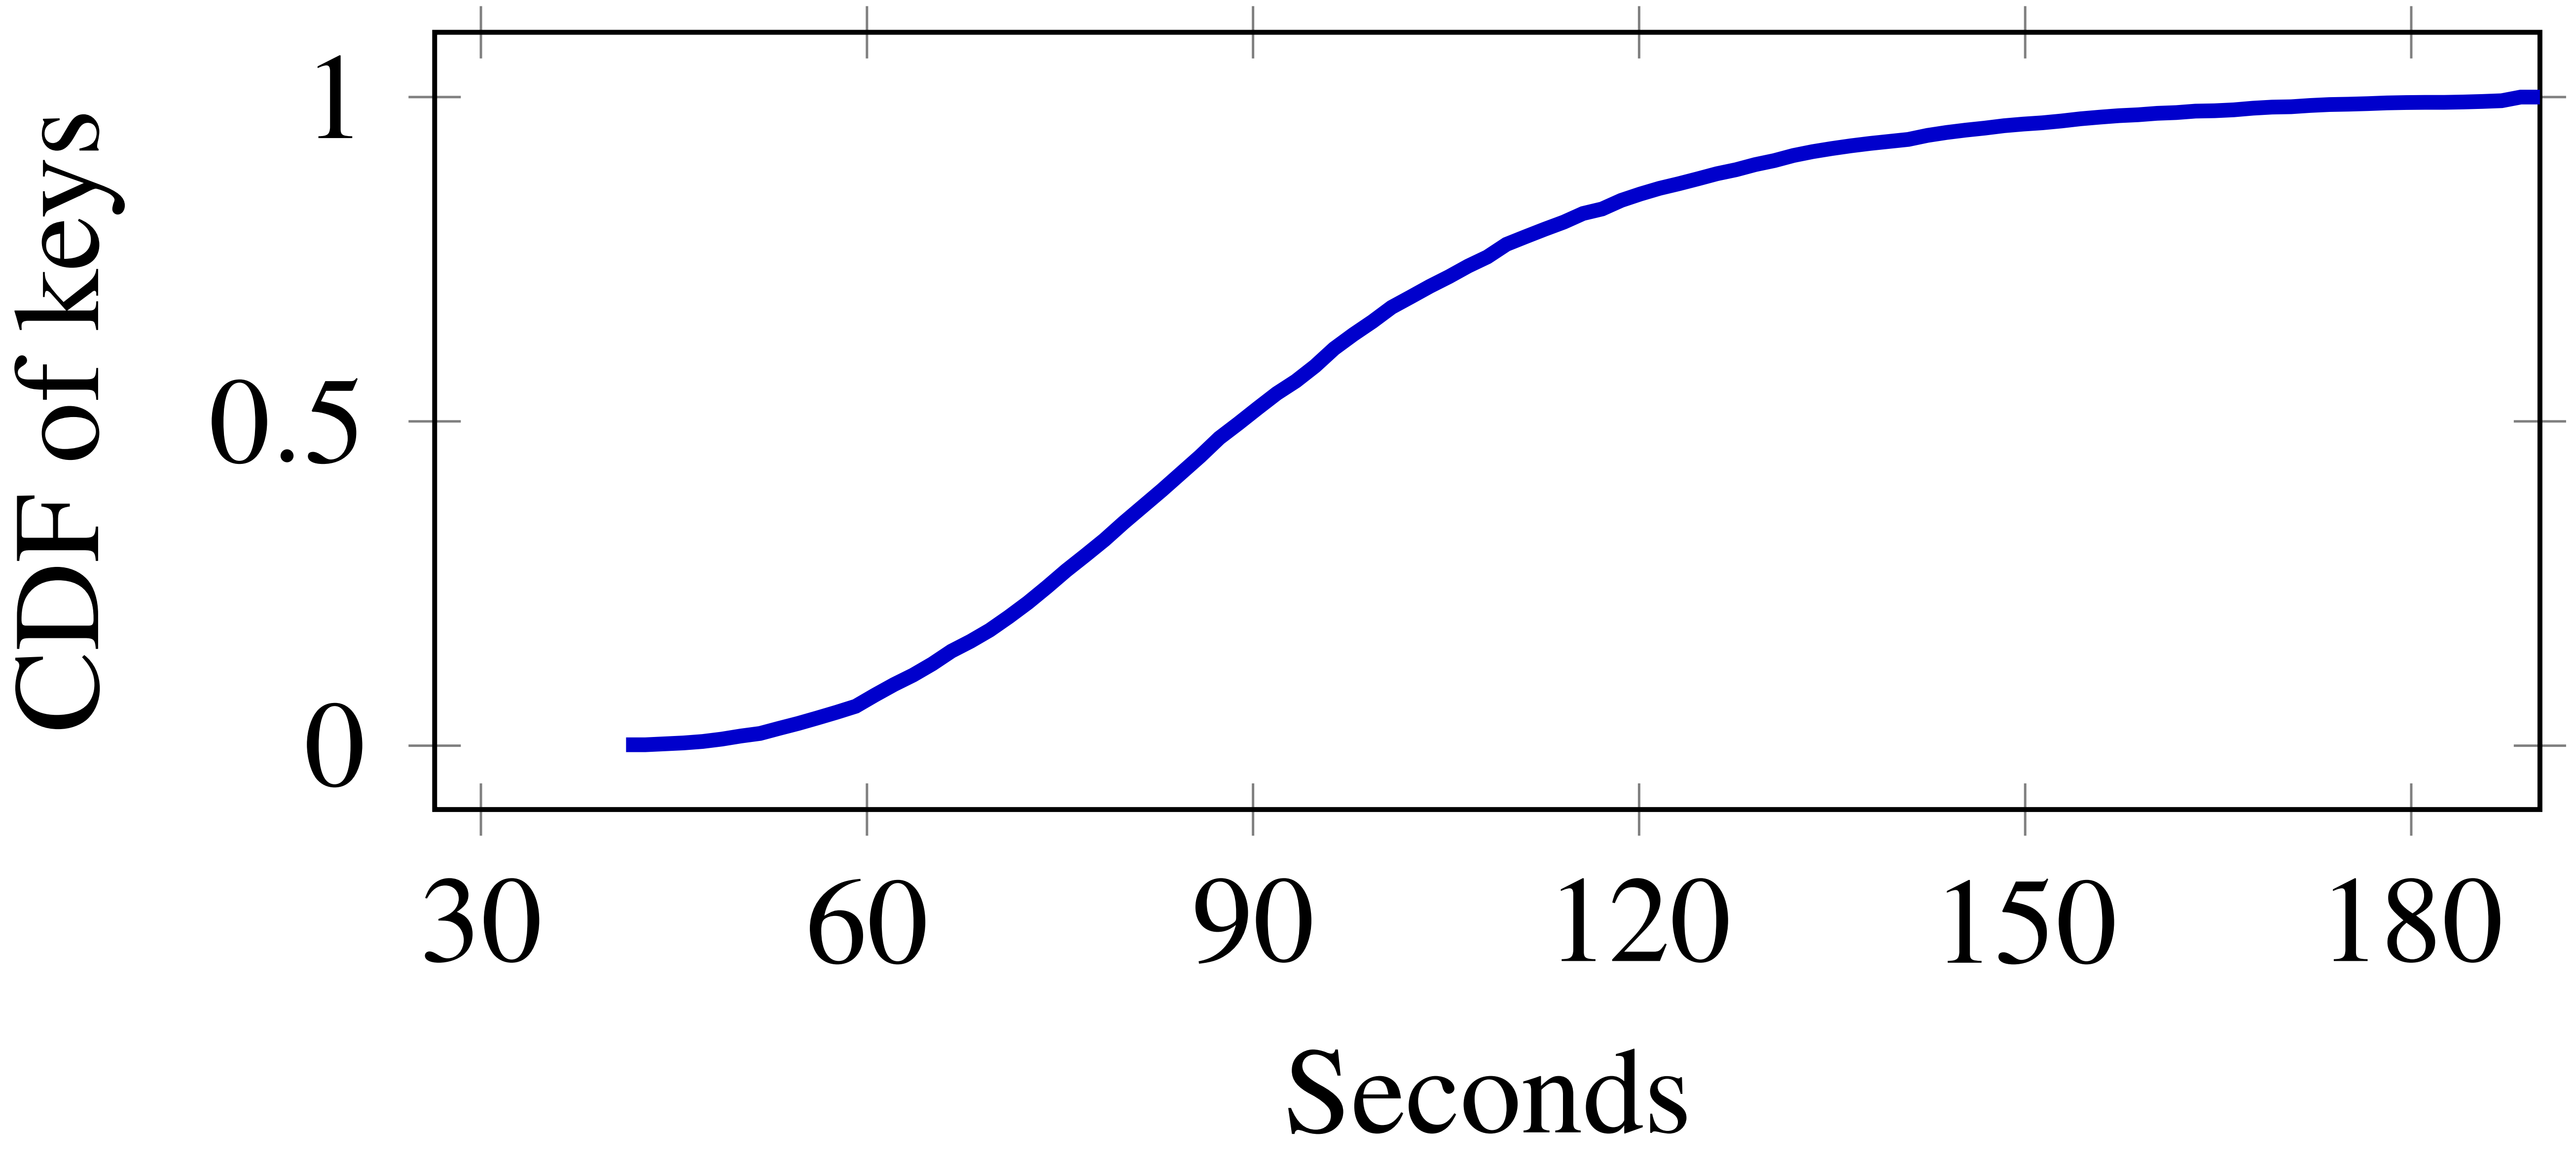

Mitigations

Browsers

- No longer support 512-bit
- Will be sunsetting 768-bit and 1024-bit

Server Operators

- Disable DHE_EXPORT
- Move to 2048-bit or elliptic curve variant

What about 1024-bit?

Cost of NFS

Rough estimations based on asymptotic complexity

	Sieving core-years	Linear Algebra core-years	Descent core-time
RSA-512	0.5	0.33	
DH-512	2.5	7.7	10 mins
RSA-768	800	100	
DH-768	8,000	28,500	2 days
RSA-1024	1,000,000	120,000	
DH-1024	10,000,000	35,000,000	30 days

Custom Hardware

If you were actually attempting this, would use custom hardware.

Prior work suggests **~80x speedup** from equivalent cost in custom hardware

Money

Prime Length	Broken By...	Precomputation Time
512-bit	Academics	1 week
768-bit	Academics	1 month
1024-bit	Nation State Large Organization	1 year ~$100-300M

Impact of a 1024-bit break

Precomputing on **one** 1024-bit prime (Oakley Group 2) would allow passively decrypting connections with:

- 66% of IPSEC VPN servers
- 26% of SSH servers

The second most common prime (Apache):

- 18% of top 1 million websites
- 6.6% of all browser trusted websites

Server Support

	Top Prime	Top 10
HTTPS Top 1M	205K (37.1%)	309K (56.1%)
HTTPS All	1.8M (12.8%)	3.4M (23.8%)
SSH	3.6M (25.7%)	3.6M (25.7%)
IKE (VPN)	1.7M (66.1%)	1.7M (66.1%)

Is NSA Breaking 1024-bit?

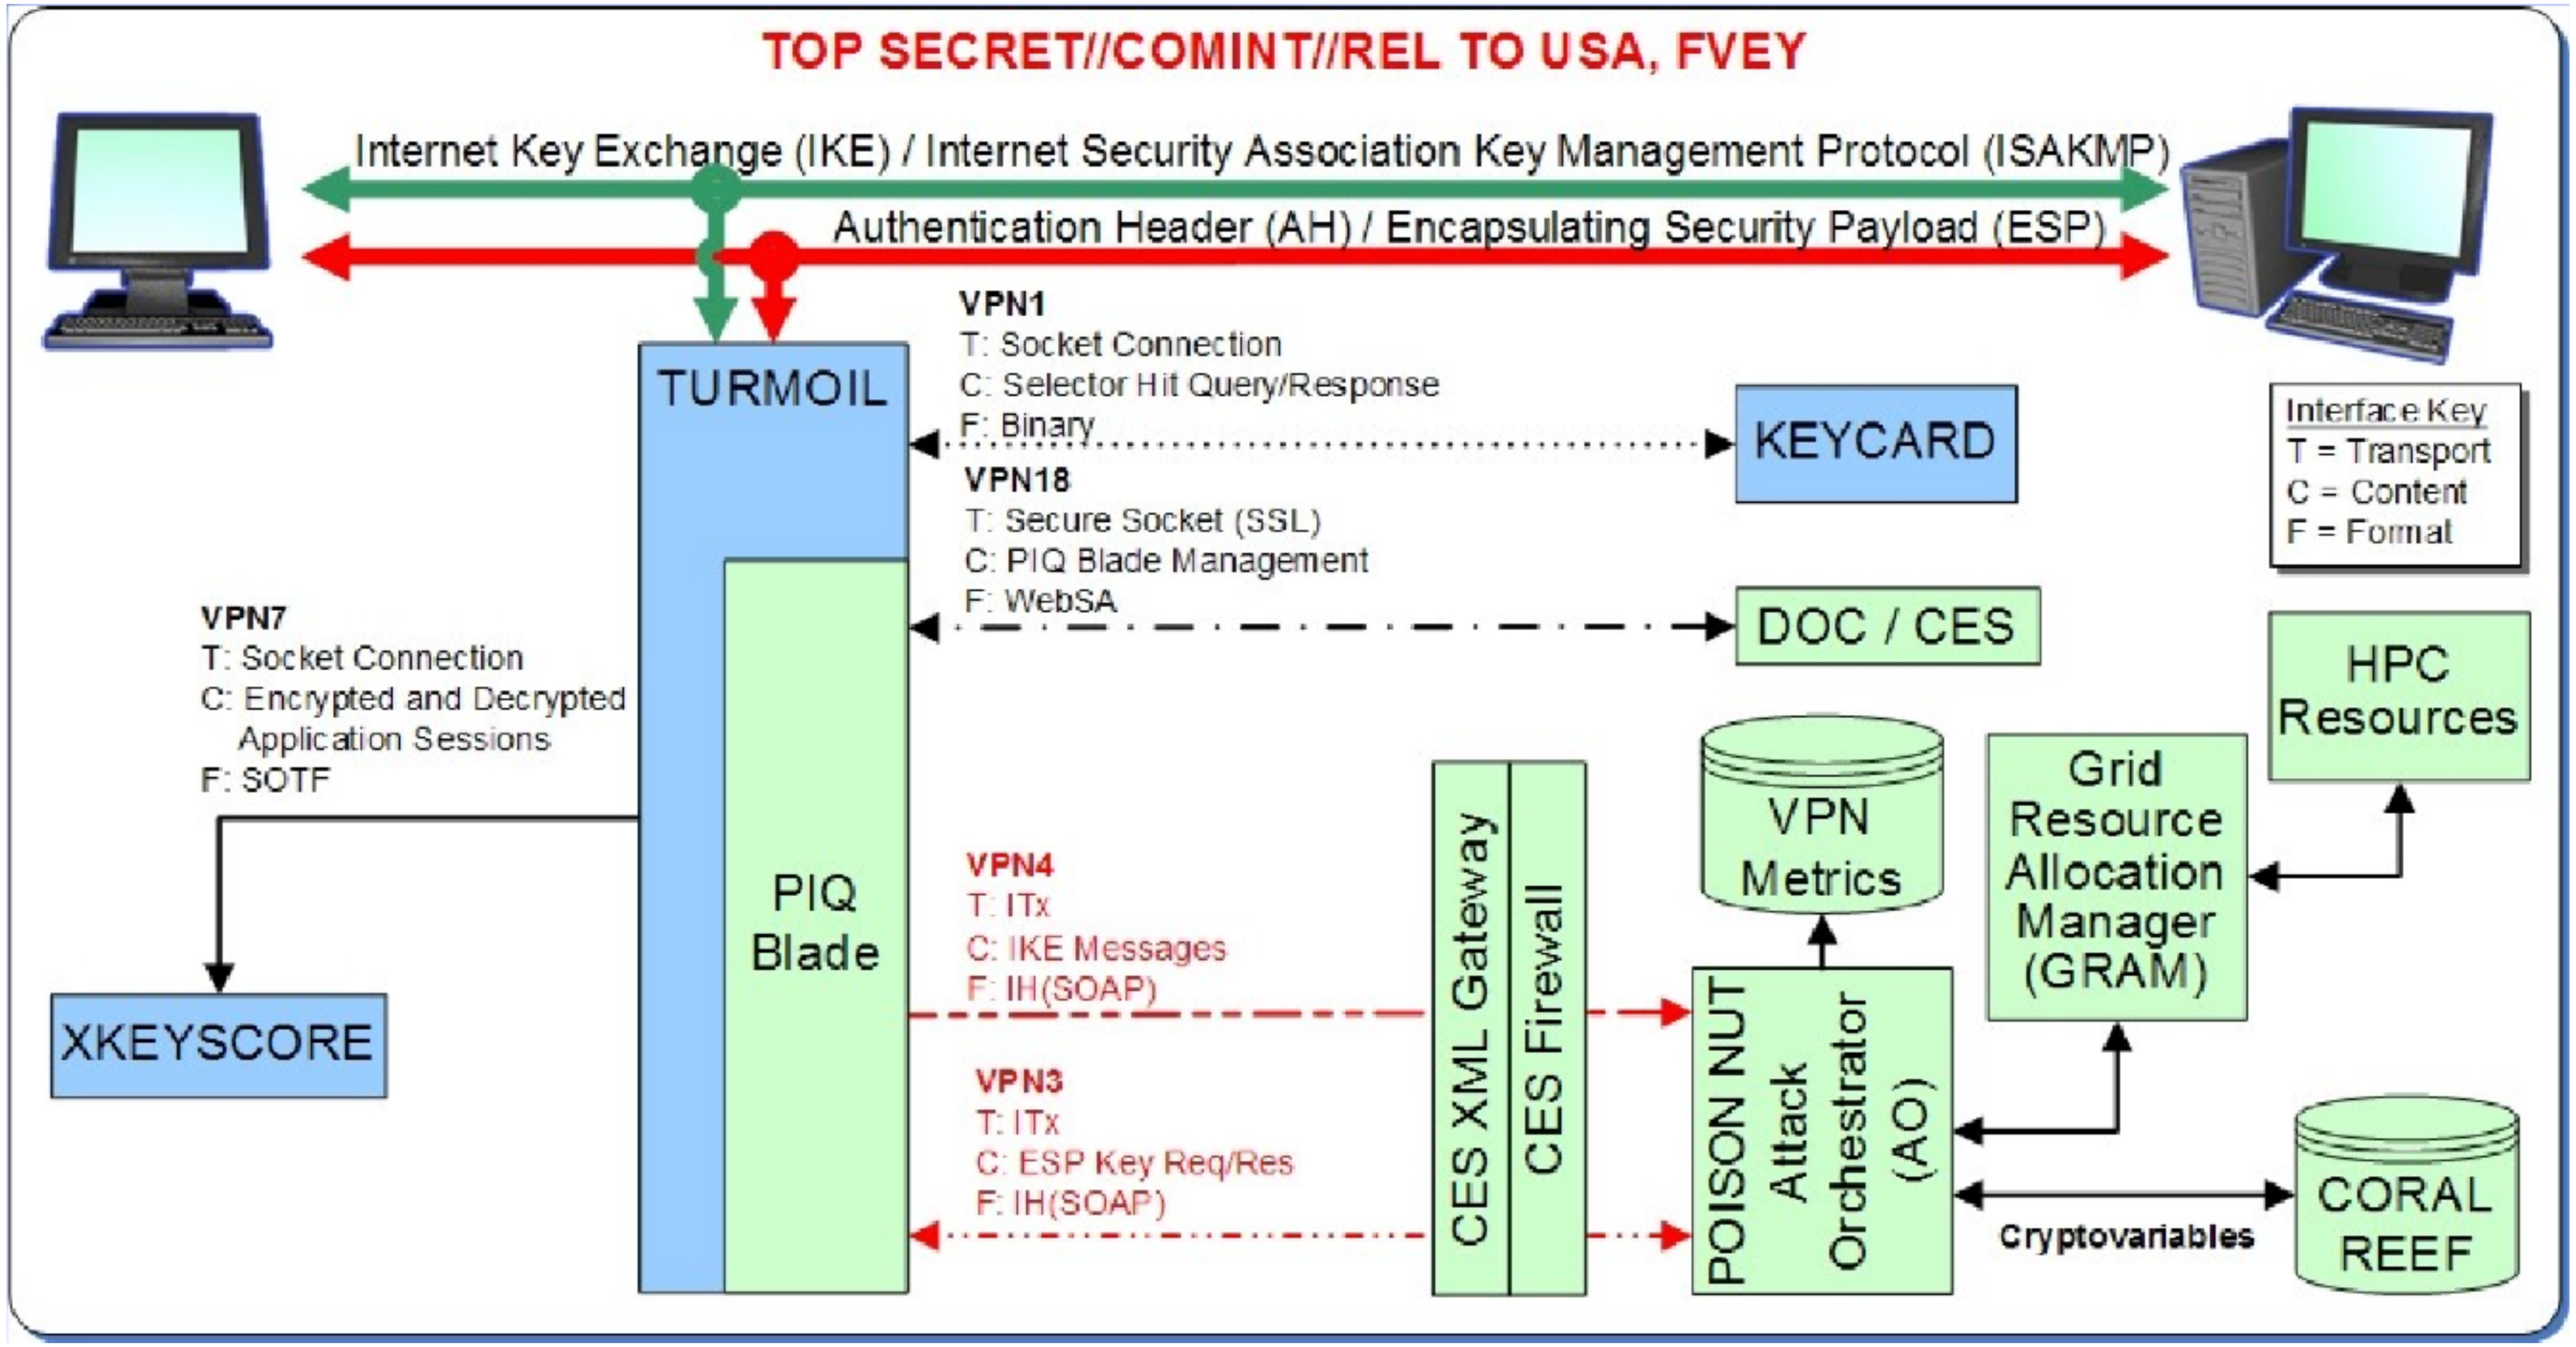

4. Communicate Results

Can we decrypt the VPN traffic?

- If the answer is “No” then explain how to turn it into a “YES!”
- If the answer is “YES!” then…

Happy Dance!!

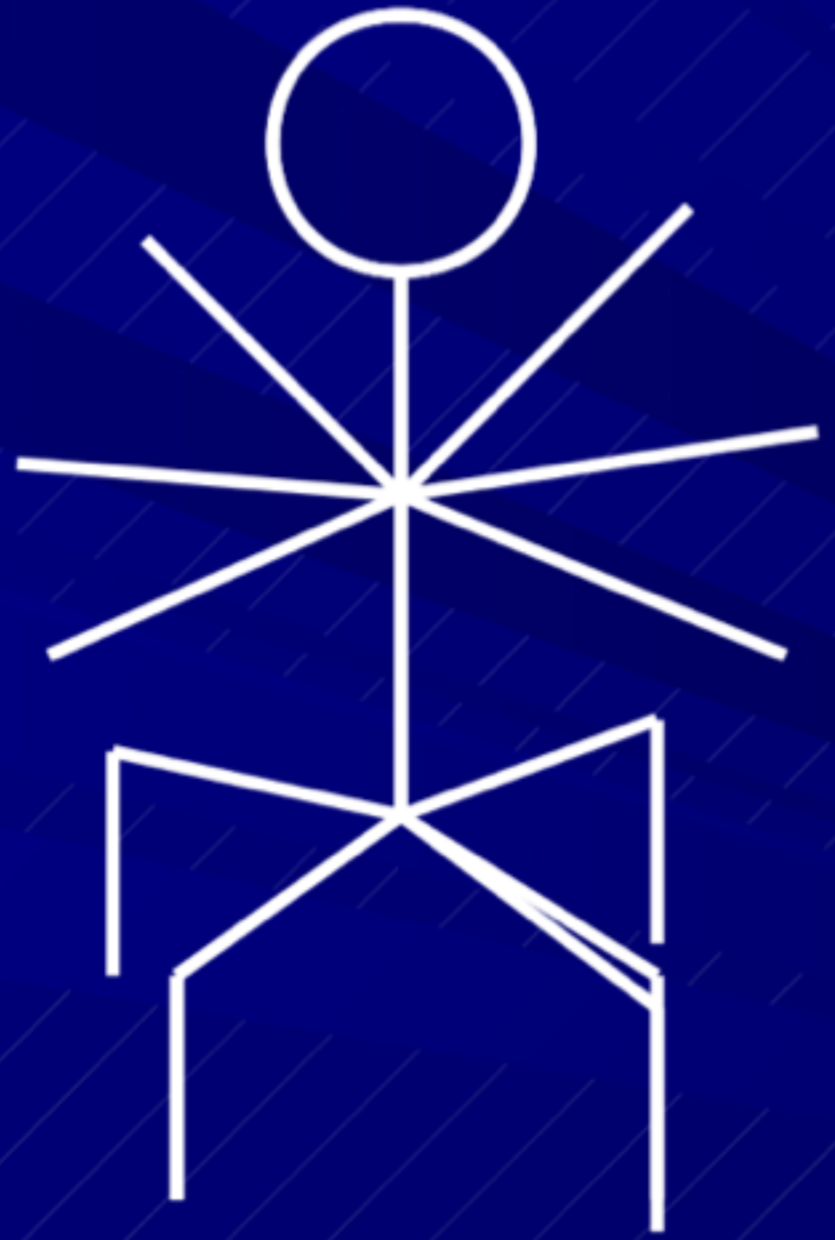

Turn that Frown Upside Down! From "No" to "YES!"

- Depends on why we couldn't decrypt it
- Find Pre-Shared Key
- Locate complete paired collect
- Locate both IKE and ESP traffic
- Have collection sites do surveys for the IP's
- Find better quality collect with rich metadata

Where did all this data come from?

ZMap

2013 A **1200x performance improvement** over Nmap for an Internet-wide single port TCP scan

2014 Scan the Internet in **under 5 minutes**.

2015 Popular in industry and academia, used by over **104** academic studies

ZMap Vision

Goals	Reality
Enable new and exciting research	Not all researchers can run ZMap
Decrease the barriers to entry for Internet-wide surveys	Negotiate with network administrators for bandwidth and address space
Anyone can scan the entire Internet using a single host	Maintain an opt-out list and respond to complaints

Search engine that allows researchers to **ask questions** about the *devices* and *networks* that compose the Internet censys

443.https.dhe_export.support: true

Search

Example

What hosts still support DHE_EXPORT?

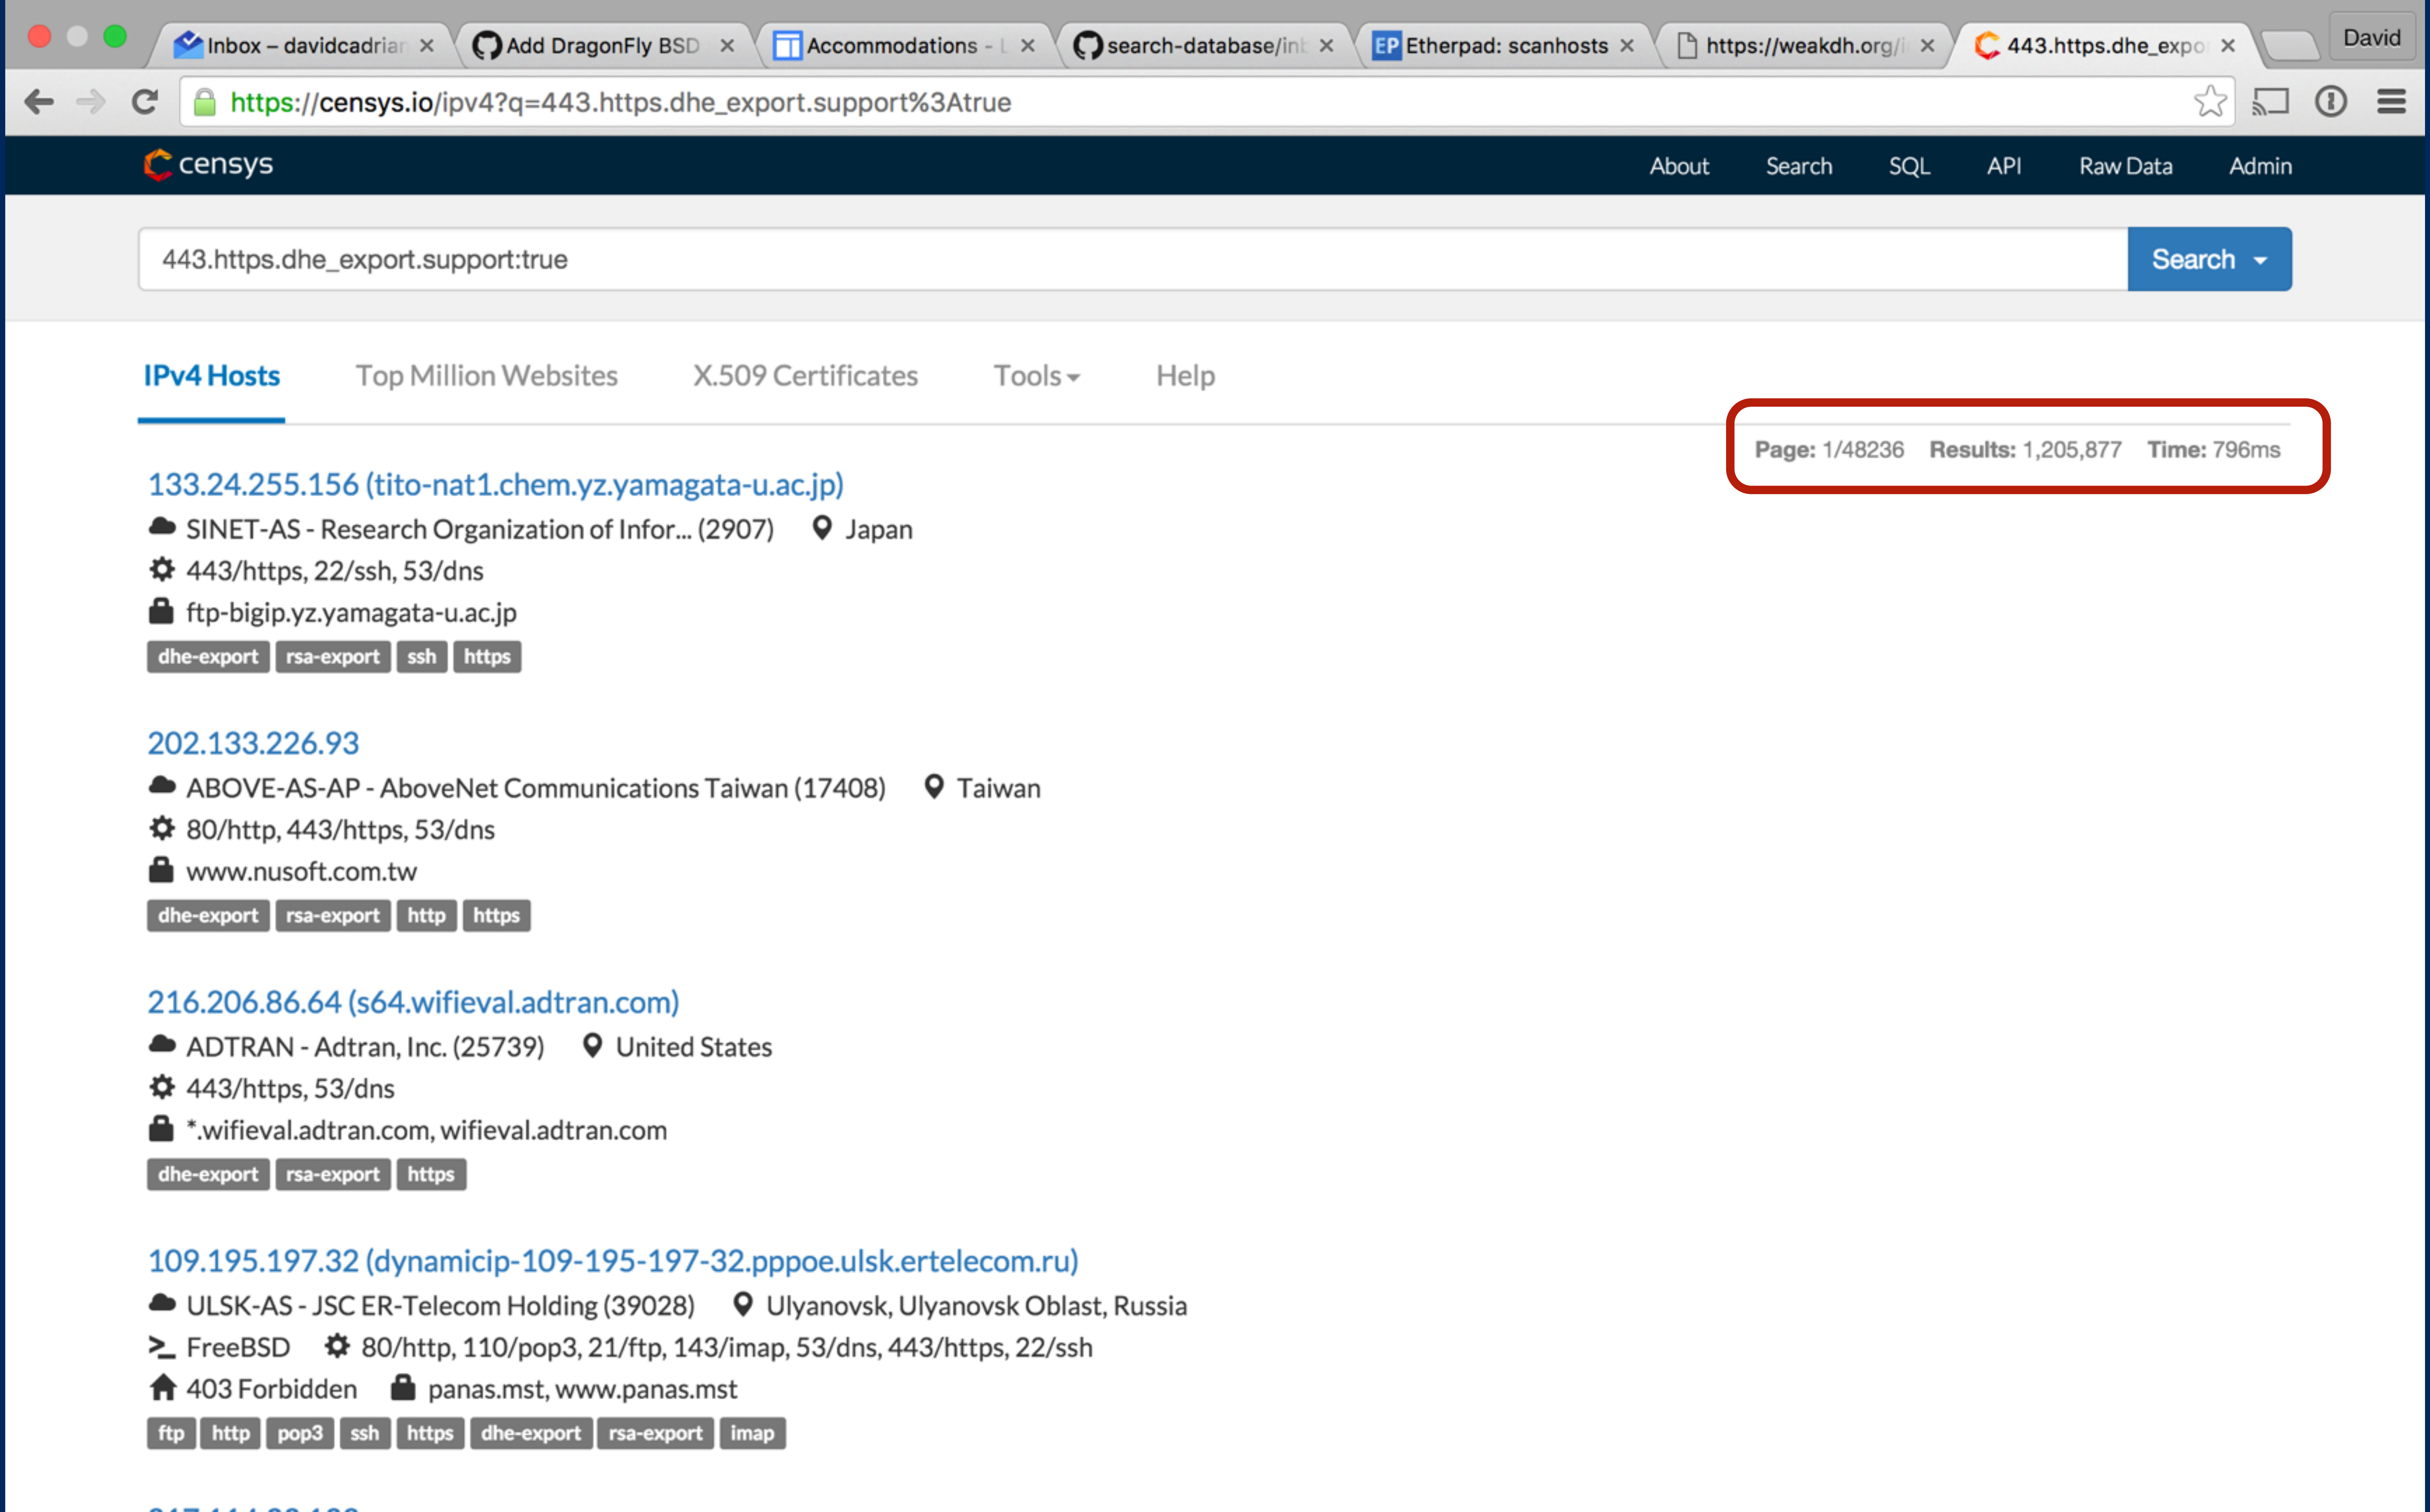

https://censys.io/ipv4?q=443.https.dhe_export.support%3Atrue
censys
About
Search
SQL
API
Raw Data
Admin
443.https.dhe_export.support:true
Search
IPv4 Hosts
Top Million Websites
X.509 Certificates
Tools
Help
Page: 1/48236 Results: 1,205,877 Time: 796ms
133.24.255.156 (tito-nat1.chem.yz.yamagata-u.ac.jp)
SINET-AS - Research Organization of Infor... (2907) Japan
443/https, 22/ssh, 53/dns
ftp-bigip.yz.yamagata-u.ac.jp
dhe-export rsa-export ssh https
202.133.226.93
ABOVE-AS-AP - AboveNet Communications Taiwan (17408) Taiwan
80/http, 443/https, 53/dns
www.nusoft.com.tw
dhe-export rsa-export http https
216.206.86.64 (s64.wifieval.adtran.com)
ADTRAN - Adtran, Inc. (25739) United States
443/https, 53/dns
*.wifieval.adtran.com, wifieval.adtran.com
dhe-export rsa-export https
109.195.197.32 (dynamicip-109-195-197-32.pppoe.ulsk.ertelecom.ru)
ULSK-AS - JSC ER-Telecom Holding (39028) Ulyanovsk, Ulyanovsk Oblast, Russia
FreeBSD 80/http, 110/pop3, 21/ftp, 143/imap, 53/dns, 443/https, 22/ssh
403 Forbidden panas.mst, www.panas.mst
ftp http pop3 ssh https dhe-export rsa-export imap

censys

443.https.tls.validation.browser_trusted:true

Search

Example

What hosts still support DHE_EXPORT?

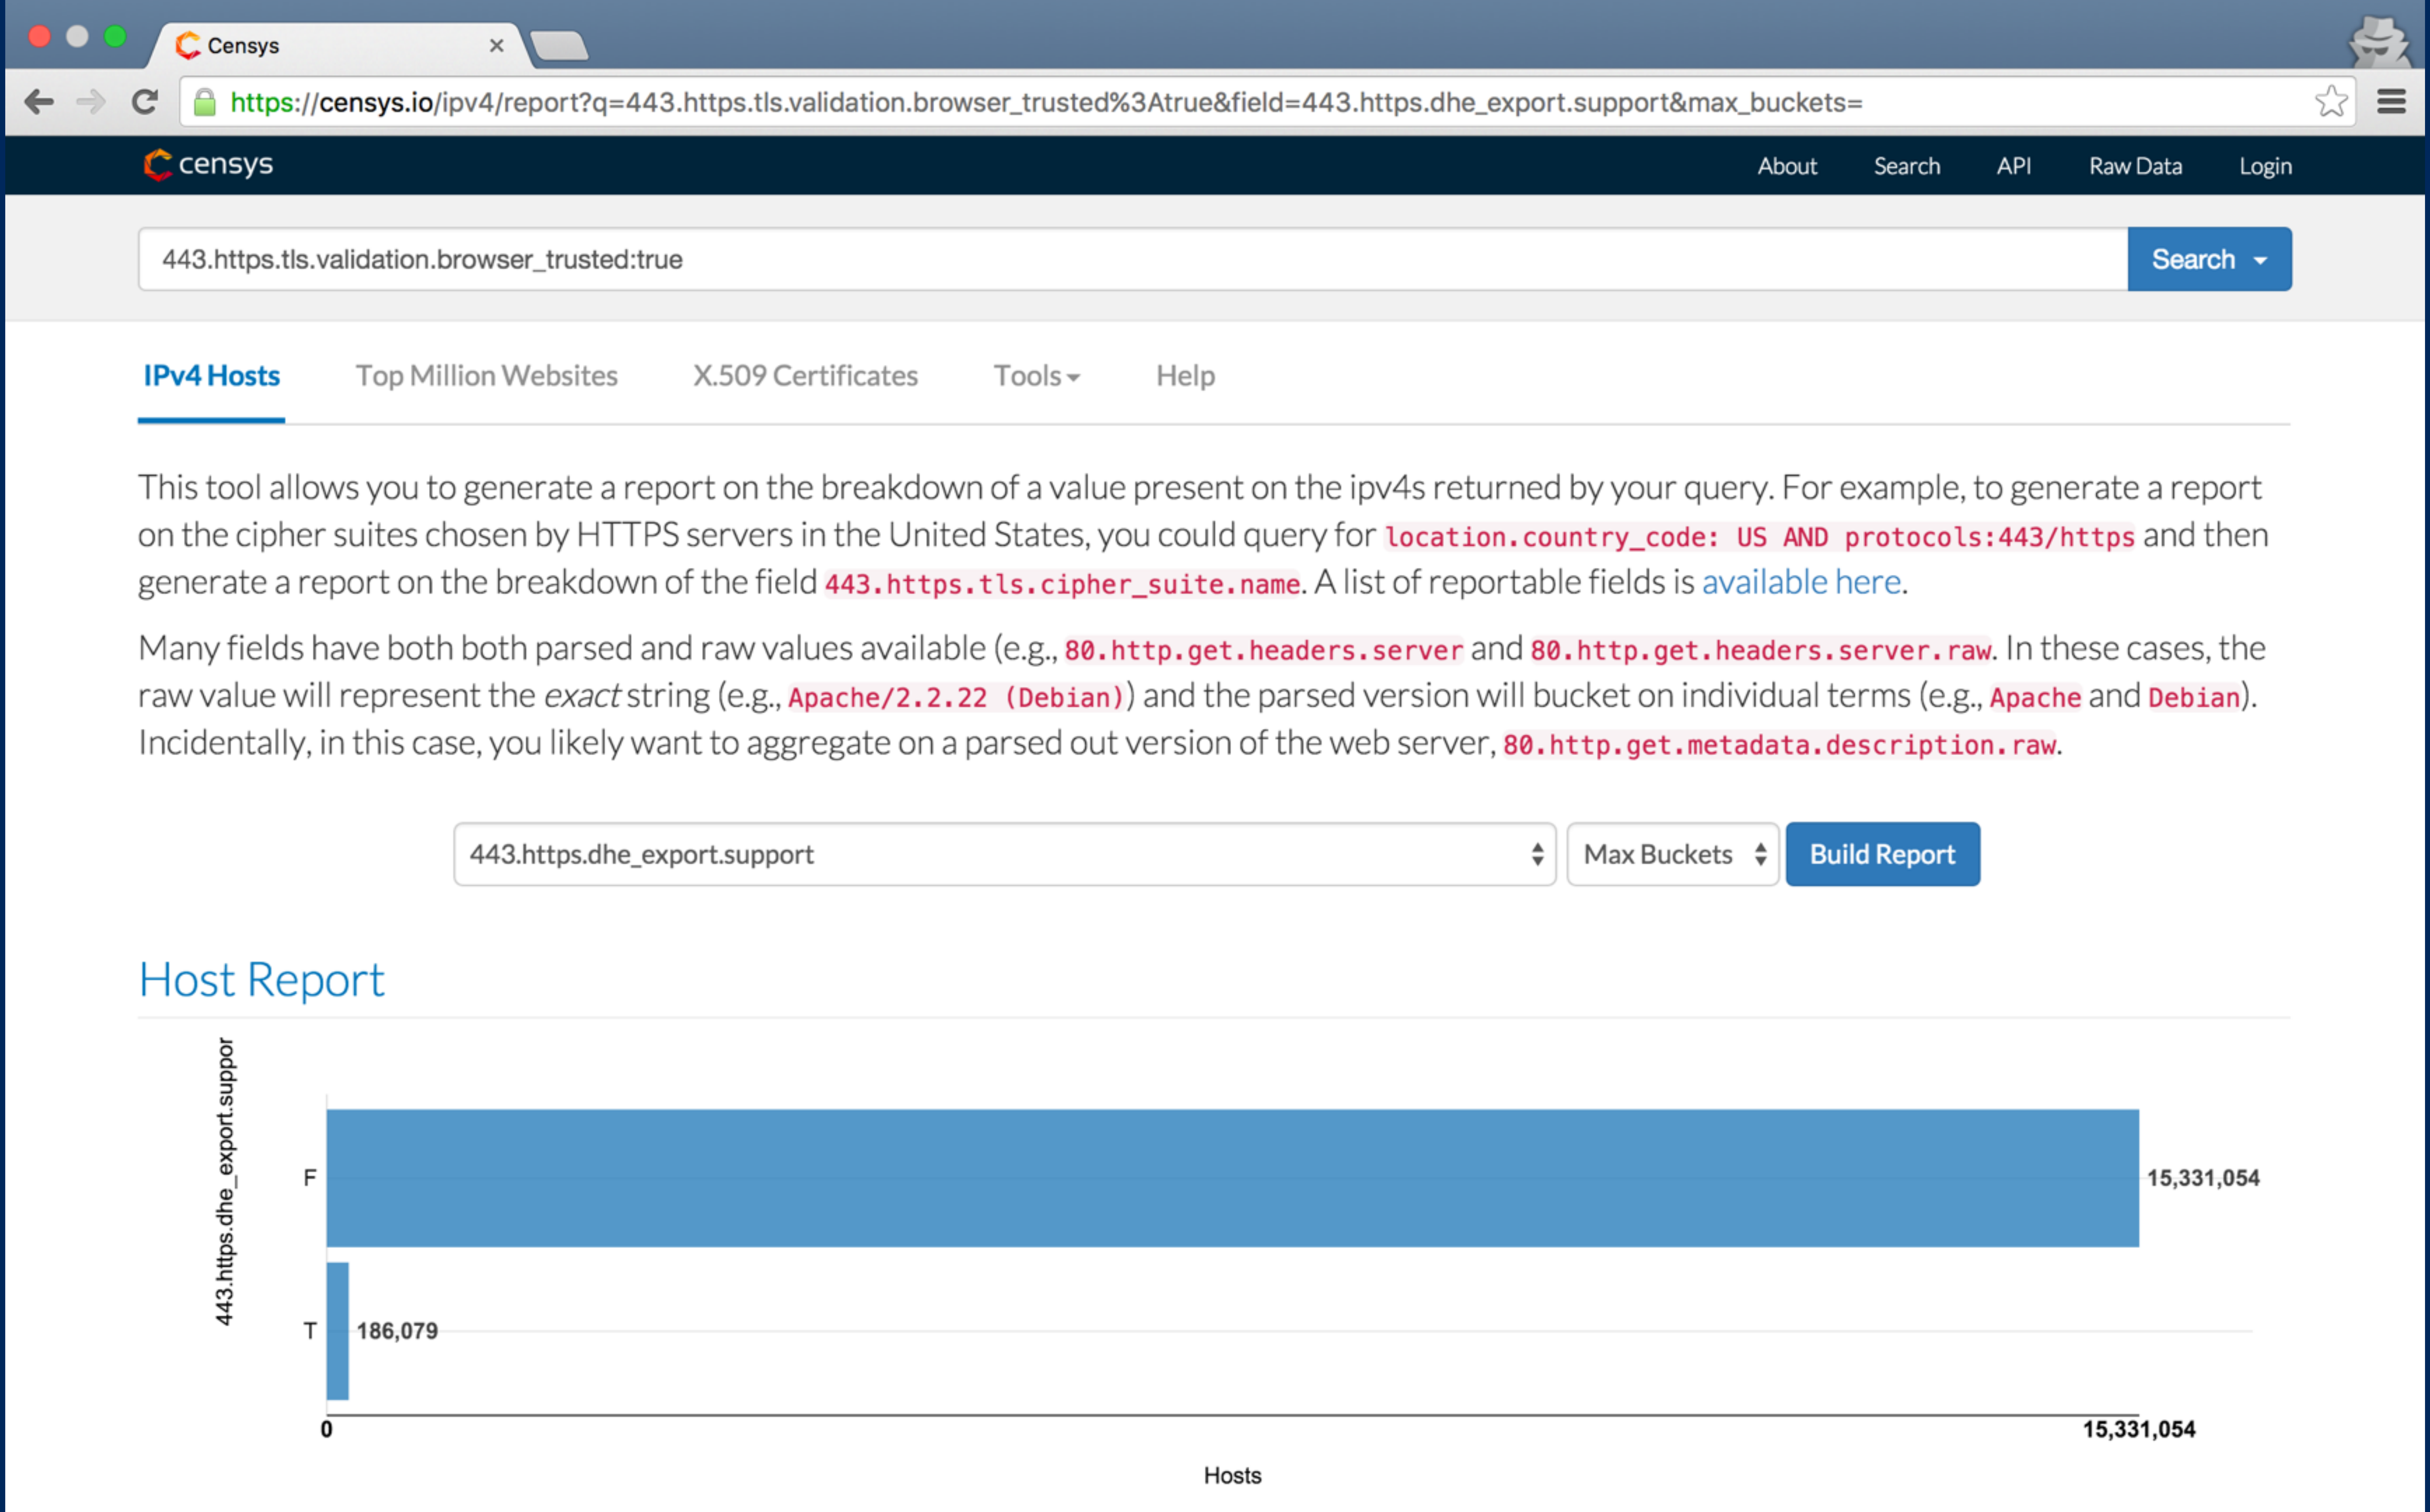
Censys
https://censys.io/ipv4/report?q=443.https.tls.validation.browser_trusted%3Atrue&field=443.https.dhe_export.support&max_buckets=
censys
About
Search
API
Raw Data
Login
443.https.tls.validation.browser_trusted:true
Search
IPv4 Hosts
Top Million Websites
X.509 Certificates
Tools
Help
This tool allows you to generate a report on the breakdown of a value present on the ipv4s returned by your query. For example, to generate a report on the cipher suites chosen by HTTPS servers in the United States, you could query for location.country_code: US AND protocols:443/https and then generate a report on the breakdown of the field 443.https.tls.cipher_suite.name. A list of reportable fields is available here.
Many fields have both both parsed and raw values available (e.g., 80.http.get.headers.server and 80.http.get.headers.server.raw. In these cases, the raw value will represent the exact string (e.g., Apache/2.2.22 (Debian)) and the parsed version will bucket on individual terms (e.g., Apache and Debian). Incidentally, in this case, you likely want to aggregate on a parsed out version of the web server, 80.http.get.metadata.description.raw.
443.https.dhe_export.support
Max Buckets
Build Report
Host Report
443.https.dhe_export.suppor
F
15,331,054
T
186,079
0
15,331,054
Hosts

censys

Full-text search

Current and historical data

SQL

API

https://censys.io

Contributing

Are you extending ZMap, ZGrab, or another scanner with a new protocol?

Do you have annotations to add to our framework?

We'll work with researchers to add new scan modules to Censys

https://github.com/zmap/zmap

https://github.com/zmap/zgrab

https://github.com/zmap/ztag

Finishing Up

Diffie-Hellman Recommendations

Transition to elliptic curve cryptography (ECC)

If ECC isn’t an option, use 2048-bit primes or larger

If 2048-bit isn’t an option, use a fresh 1024-bit prime

All major desktop browsers now reject 512-bit groups, and are sunsetting 768-bit and 1024-bit

Turn export ciphers off!

Censys

Censys strives to be **research enabling more research**

Contribute back scanners and annotations — we do the heavy lifting

Bring **measurement-driven security** to a wider audience

Questions?

David Adrian
@davidcadrian

https://weakdh.org

https://censys.io

https://zmap.io

UNIVERSITY OF MICHIGAN