Were It So Easy

Using TLS in the Real World

David Adrian
@davidcadrian

Who am I?

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

I am advised by Professor J. Alex Halderman

Computer security, systems, network measurement, research

Where is Zakir?

We just don’t know.

He’s probably on fire.

What am I talking about?

What is TLS and why study it?

What happens when TLS fails?

What can we do to prevent TLS failing in the future?

TLS provides a **secure channel** for communication that other protocols can be built on top of.

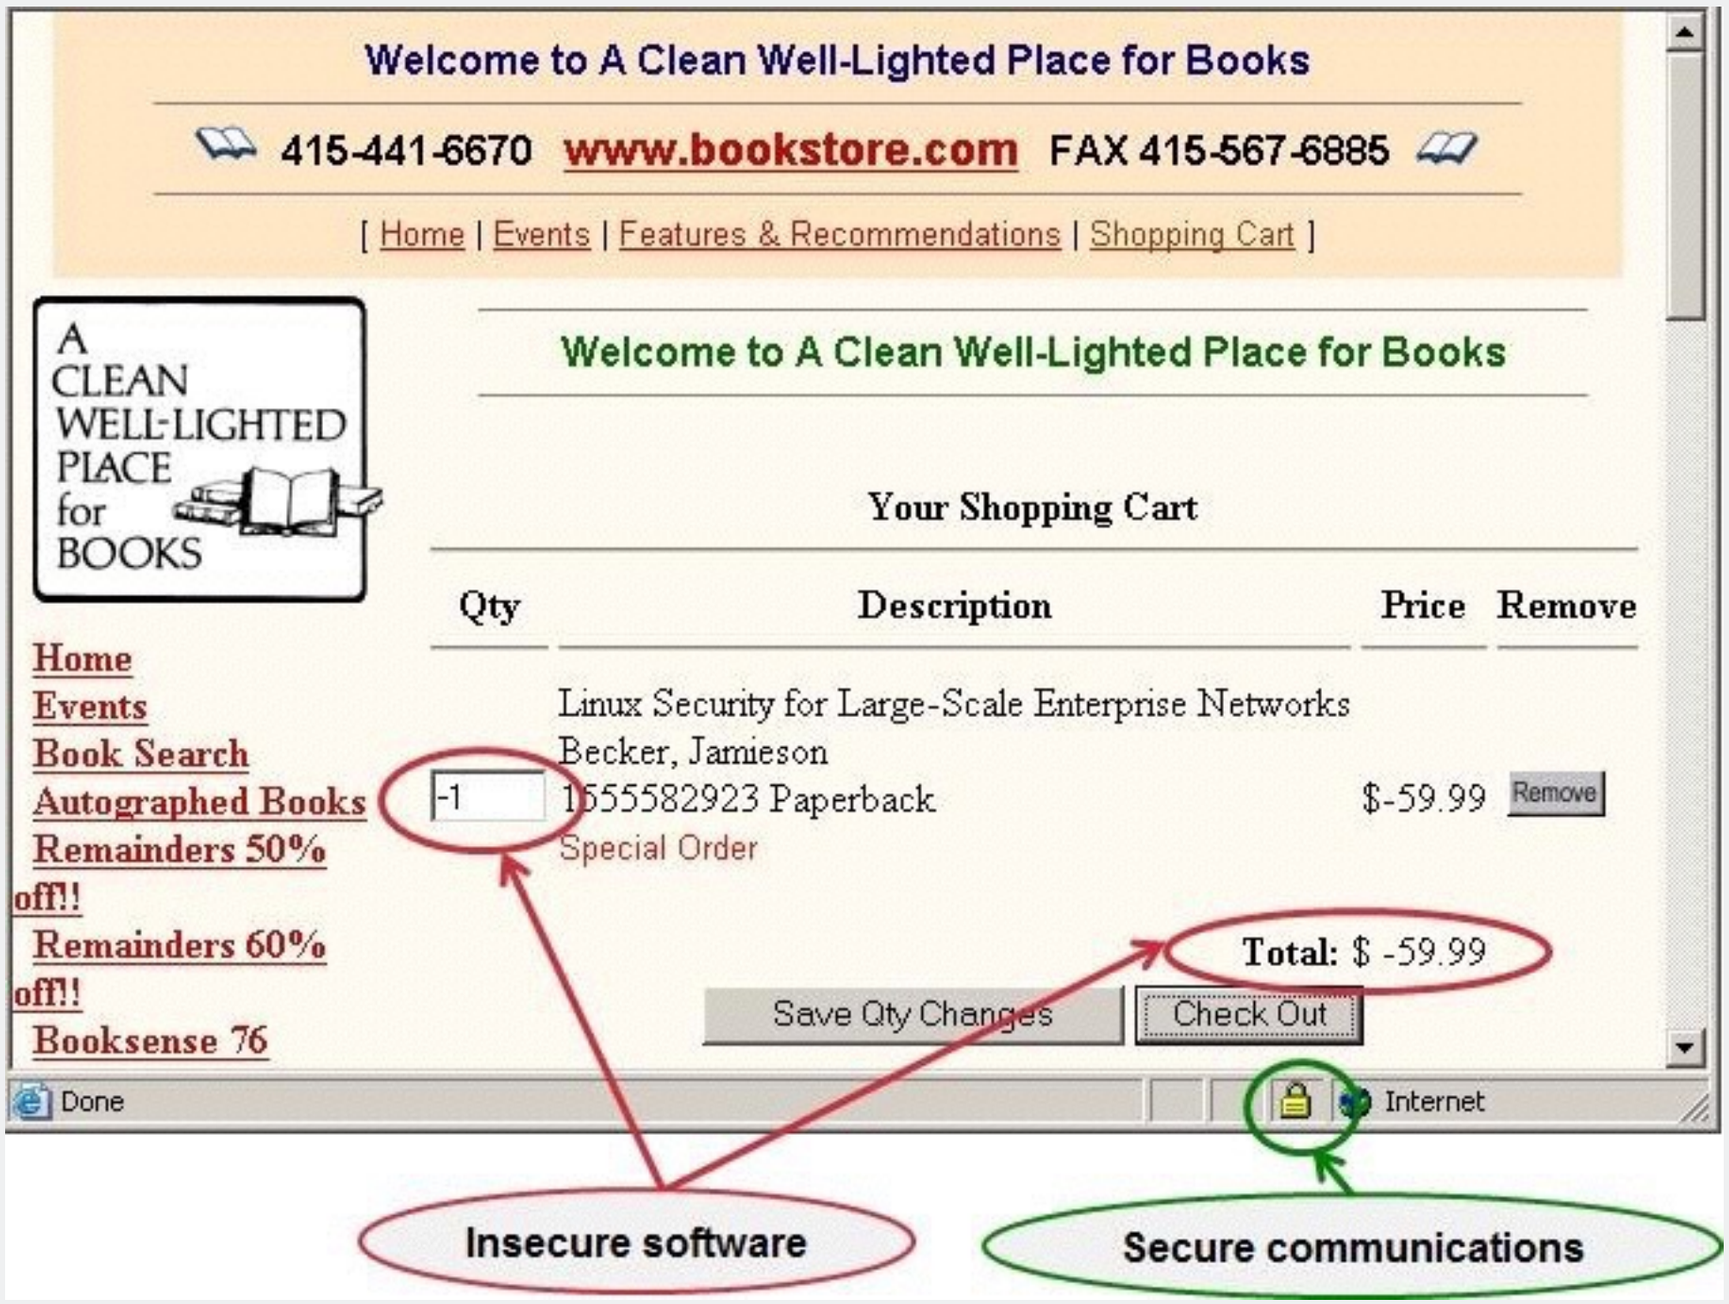

Does not provide application-level security!

https://twitter.com/ericbaize/status/492777221225213952

Confidentiality

Attacker cannot read messages

Integrity

Attacker cannot modify or replay messages

Authentication

Attacker cannot impersonate the recipient

TLS, SSL, HTTPS, oh my…

SSL: Secure Socket Layer

- Originally developed by Netscape in the 90s
- It didn't catch on until the third version (SSLv3)

TLS: Transport Layer Security

- Successor to SSL (TLS v1.0 = SSLv3.1)

HTTPS: HTTP ran inside of TLS

What happens when TLS fails?

Heartbleed

In April 2014, OpenSSL disclosed a bug in their implementation of the TLS Heartbeat Extension

Vulnerability allowed attackers to dump private cryptographic keys, logins, and other private user data

Potentially effected any service that used OpenSSL for TLS—including web, mail, messaging and database servers

An estimated 24-55% of HTTPS websites were initially vulnerable

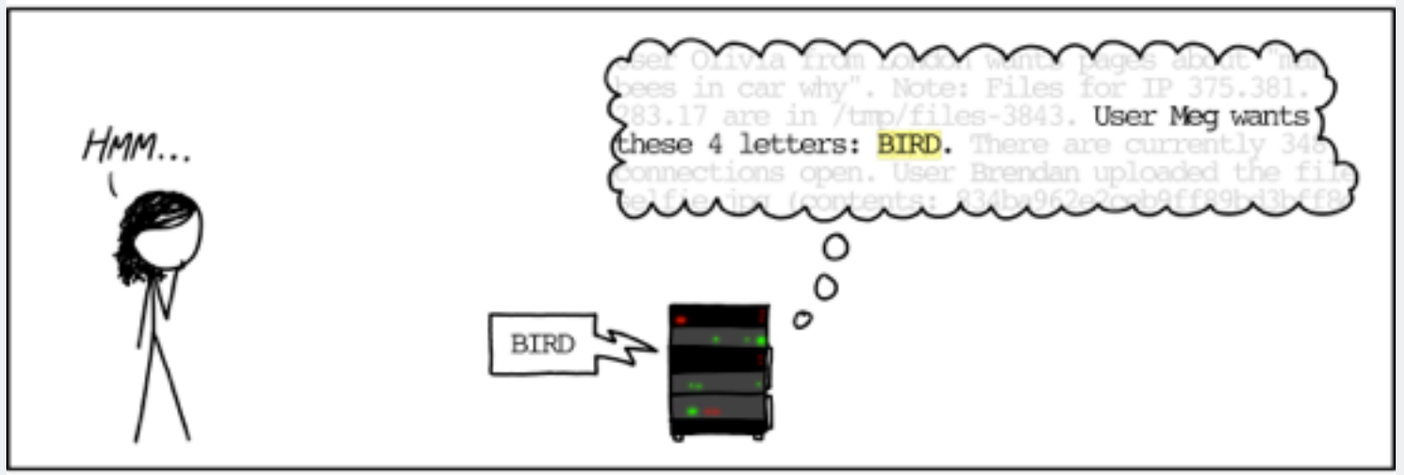

https://xkcd.com/1354/

What can we **learn** about TLS from Heartbleed?

How can we **measure the impact** of Heartbleed on TLS a whole?

ZMap

Network port scanner capable of completing an Internet-wide TCP SYN scan, in **under five minutes** on a 10G uplink, and under forty-five minutes on a 1G uplink, from a single machine.

We use ZMap to identify and connect to **all HTTPS servers** on the Internet.

Data Collection

- Began scanning 48 hours after public disclosure
- Scanned Alexa Top 1 Million and 1% samples of IPv4 every 8 hours

RFC 6520 Compliant

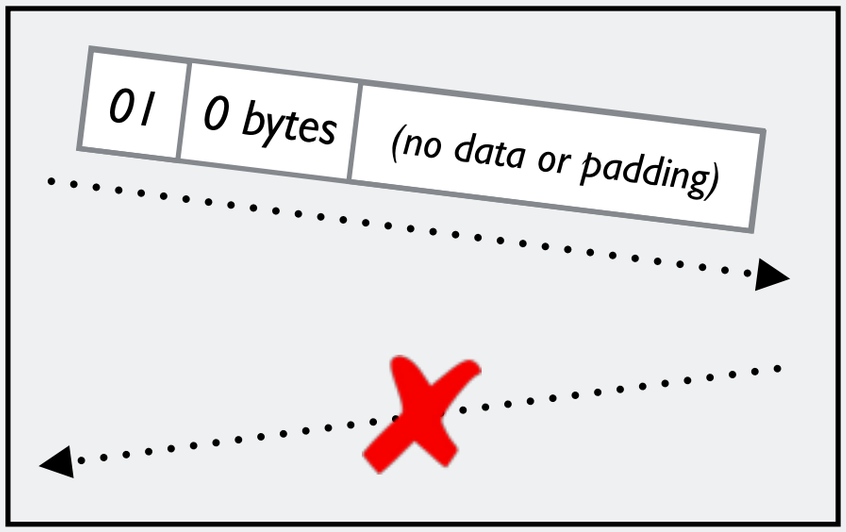

Scanning for Heartbleed

- Instead of exploiting the vulnerability, we checked for non-compliant behavior of vulnerable OpenSSL version

https://zmap.io/heartbleed

Vulnerable OpenSSL

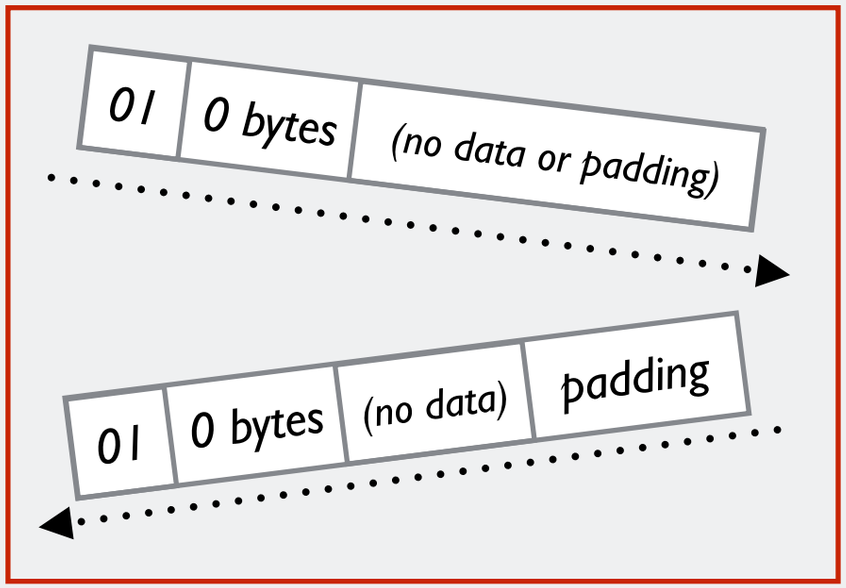

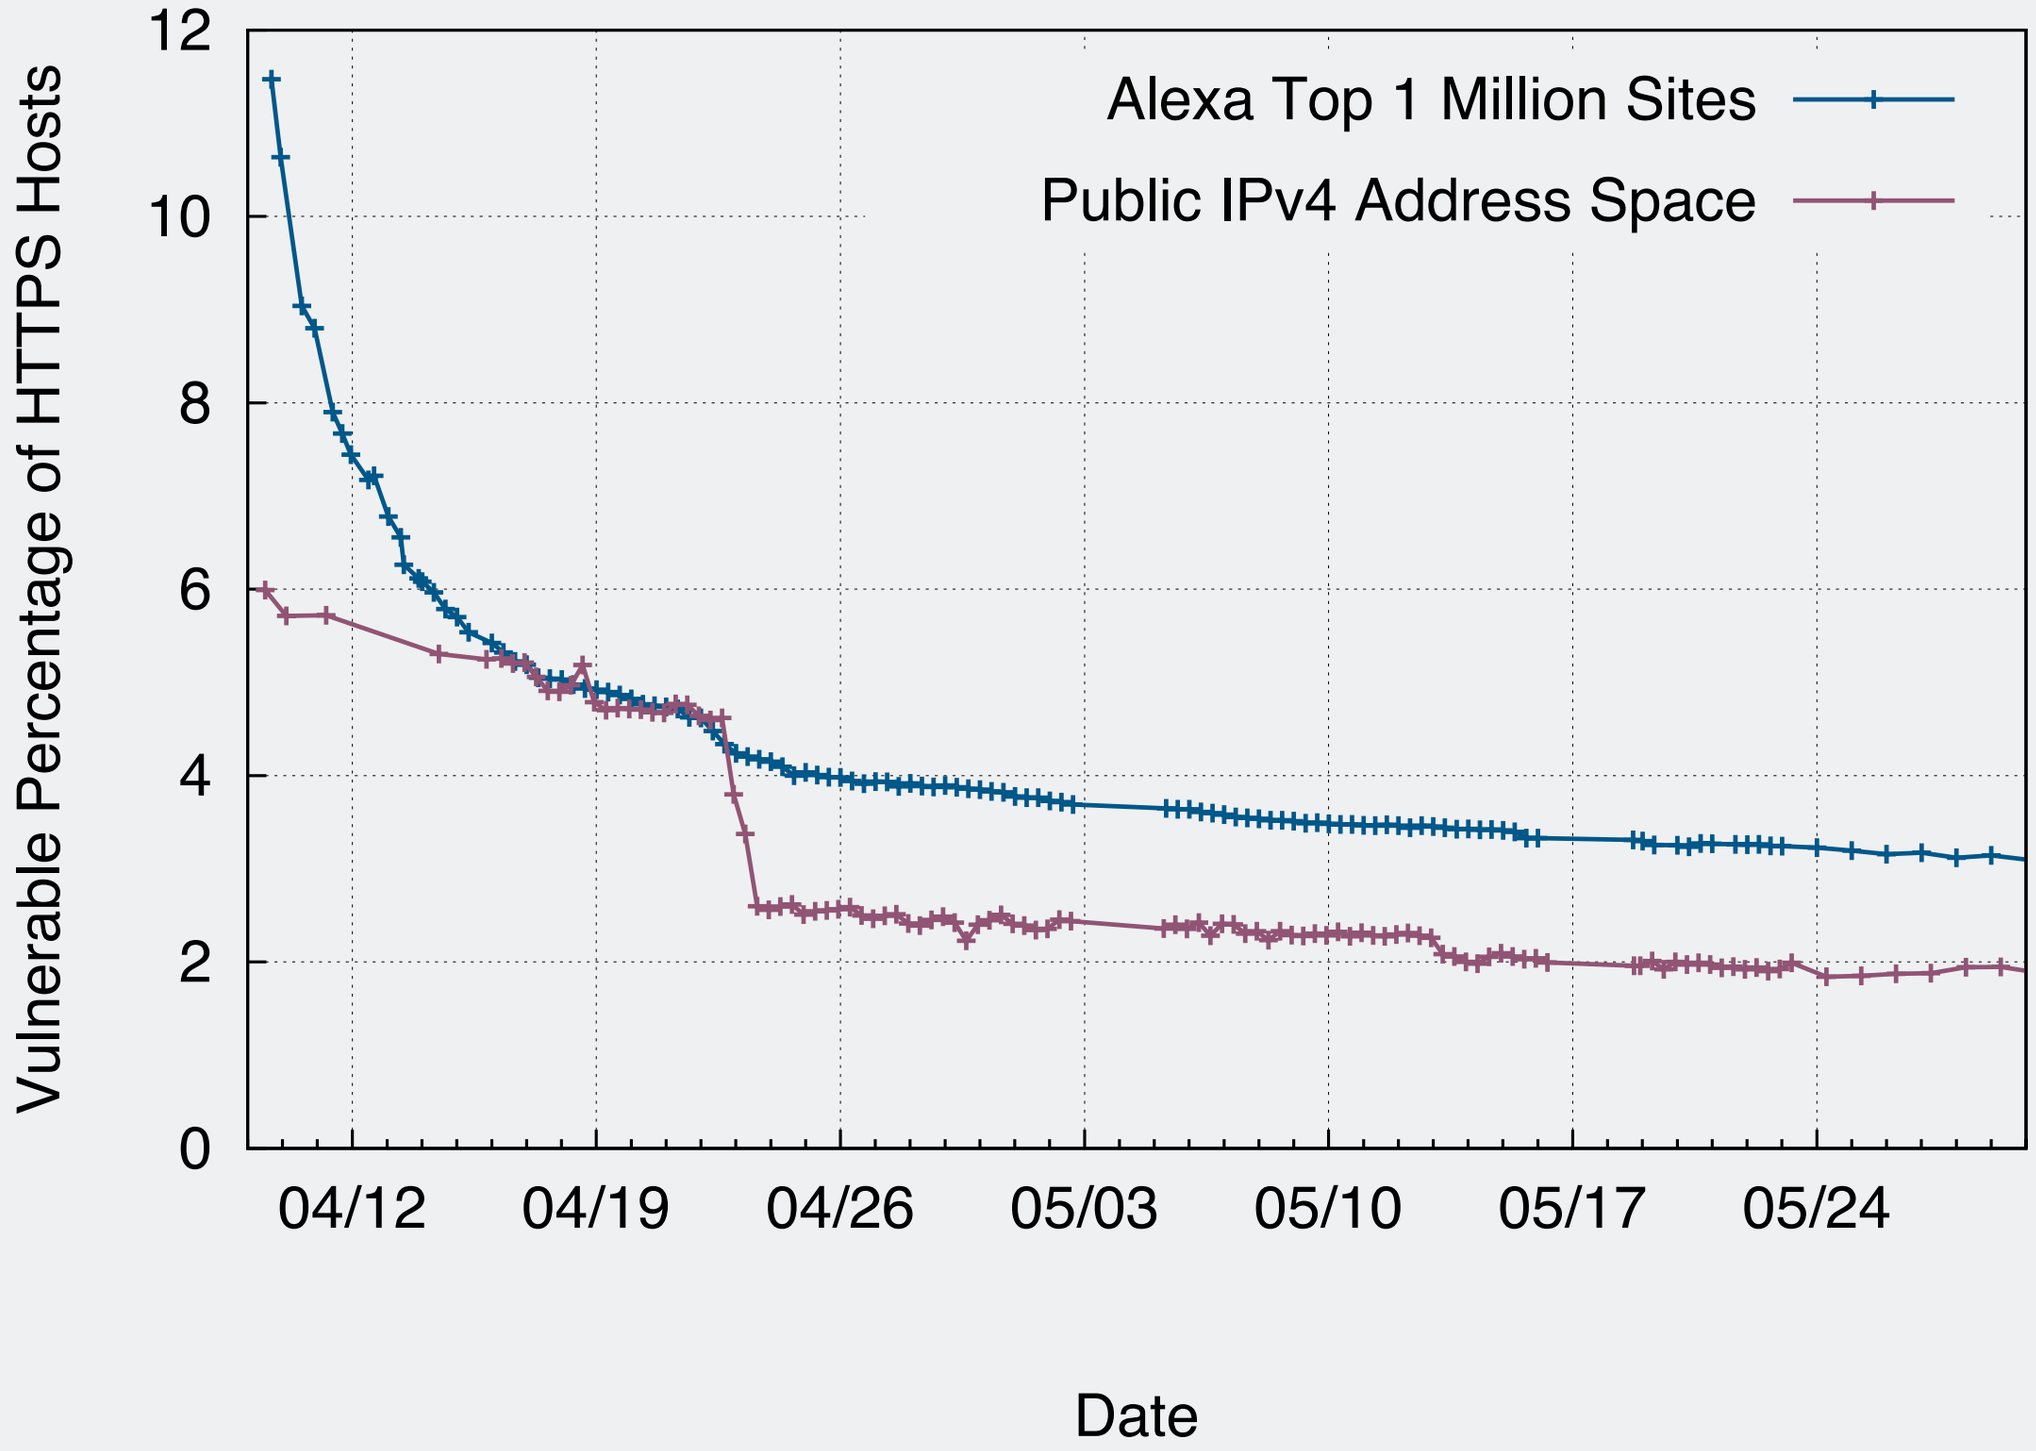
Alexa Top 1 Million Sites
Public IPv4 Address Space
Vulnerable Percentage of HTTPS Hosts
12
10
8
6
4
2
0
04/12
04/19
04/26
05/03
05/10
05/17
05/24
Date

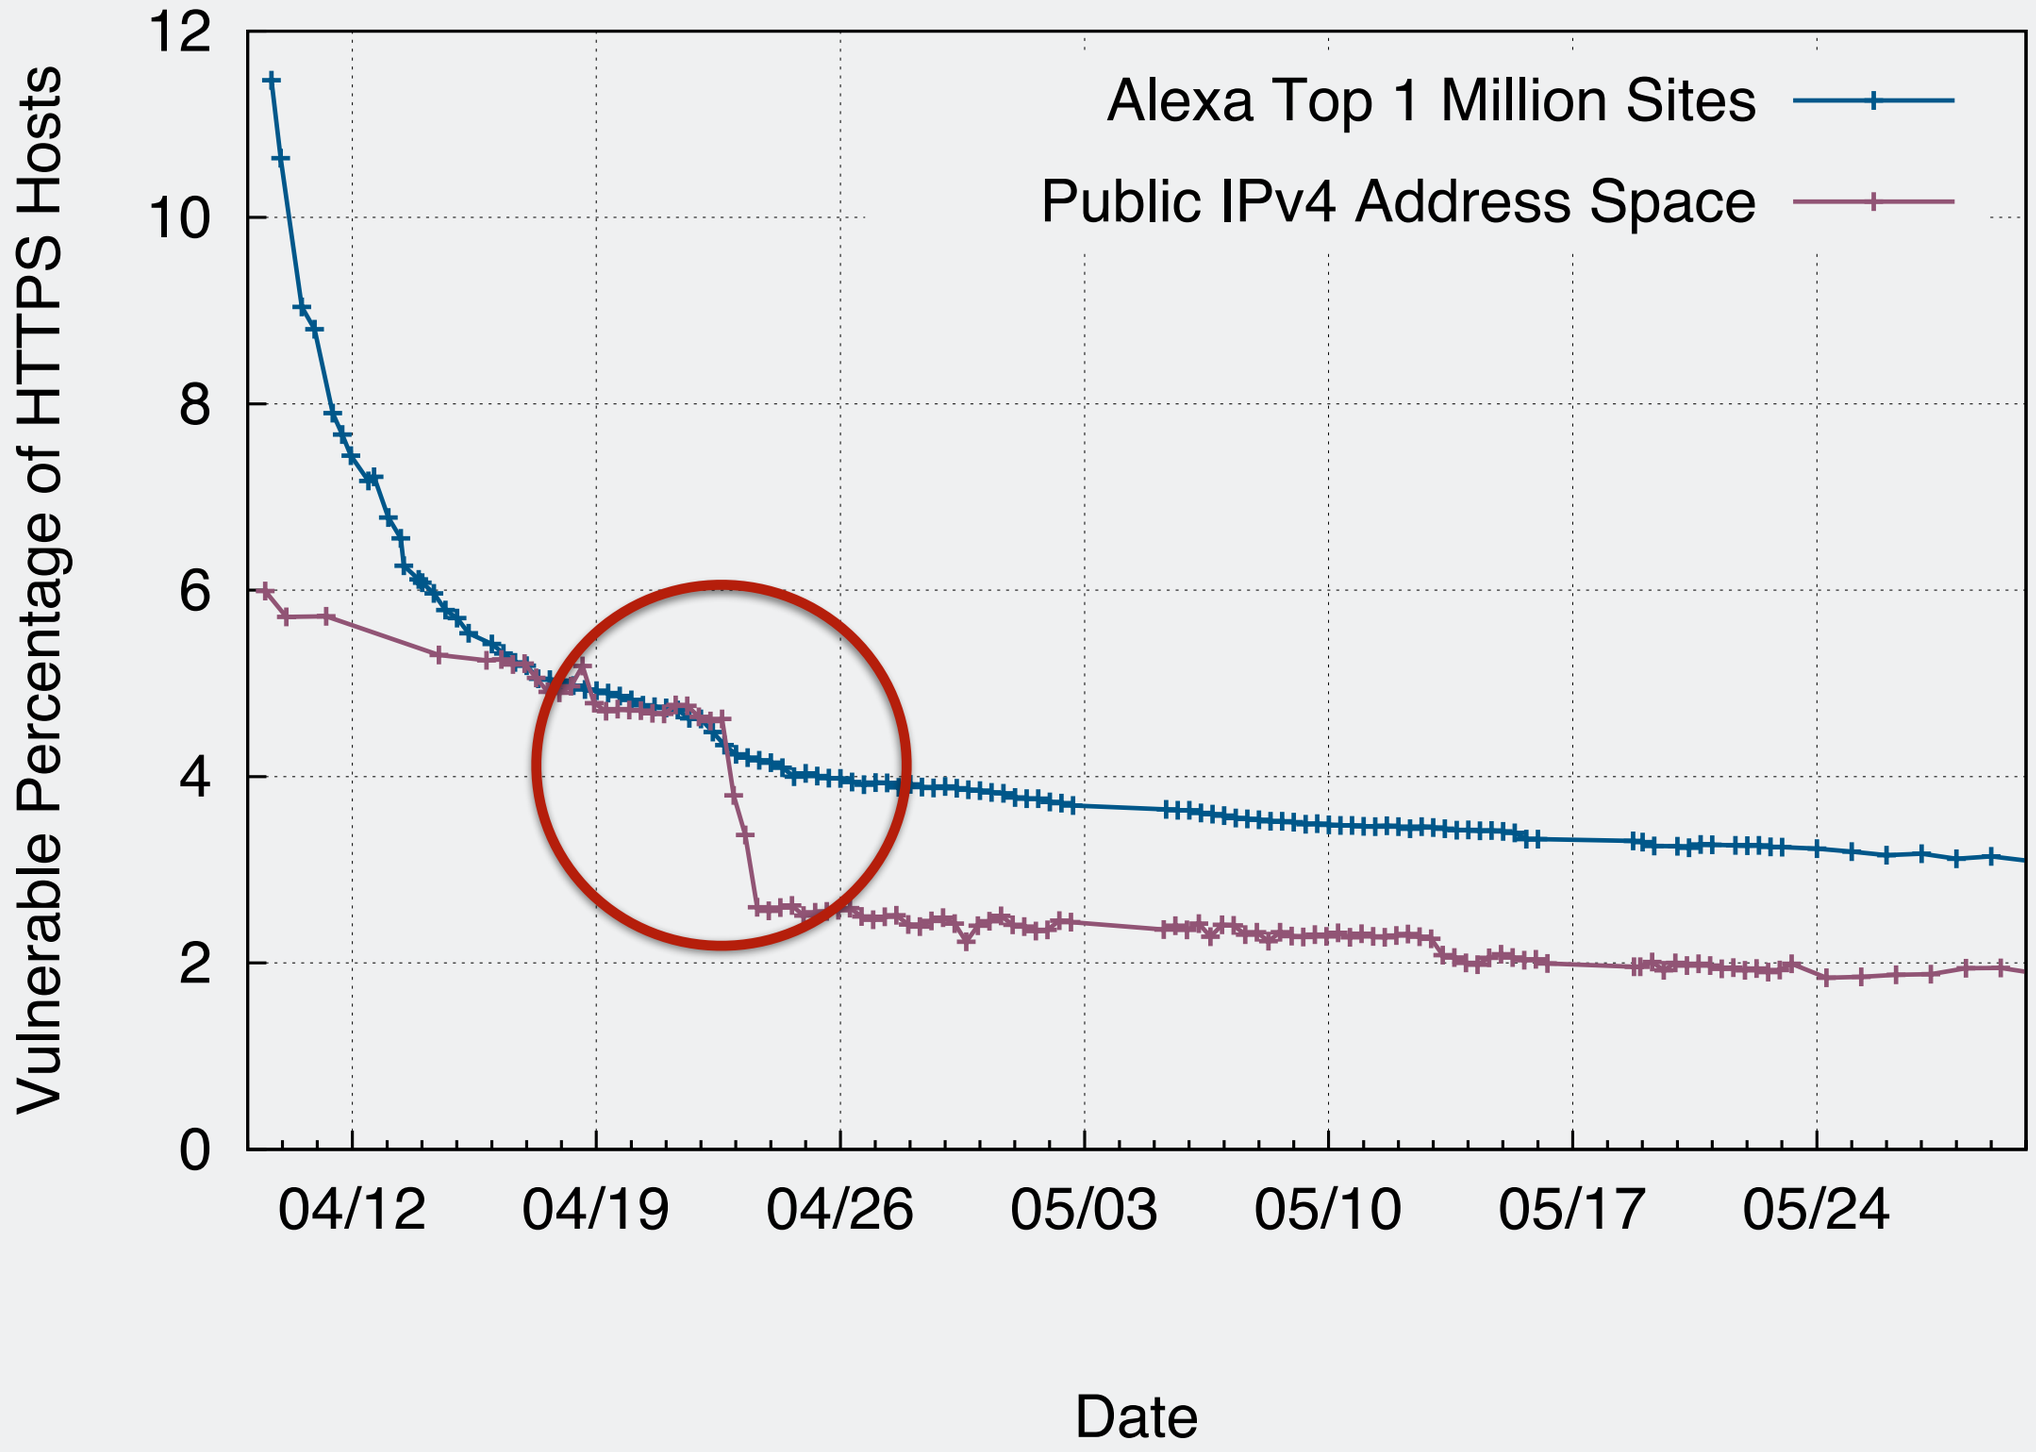
Vulnerable Percentage of HTTPS Hosts
12
10
8
6
4
2
0
04/12
04/19
04/26
05/03
05/10
05/17
05/24
Date
Alexa Top 1 Million Sites
Public IPv4 Address Space

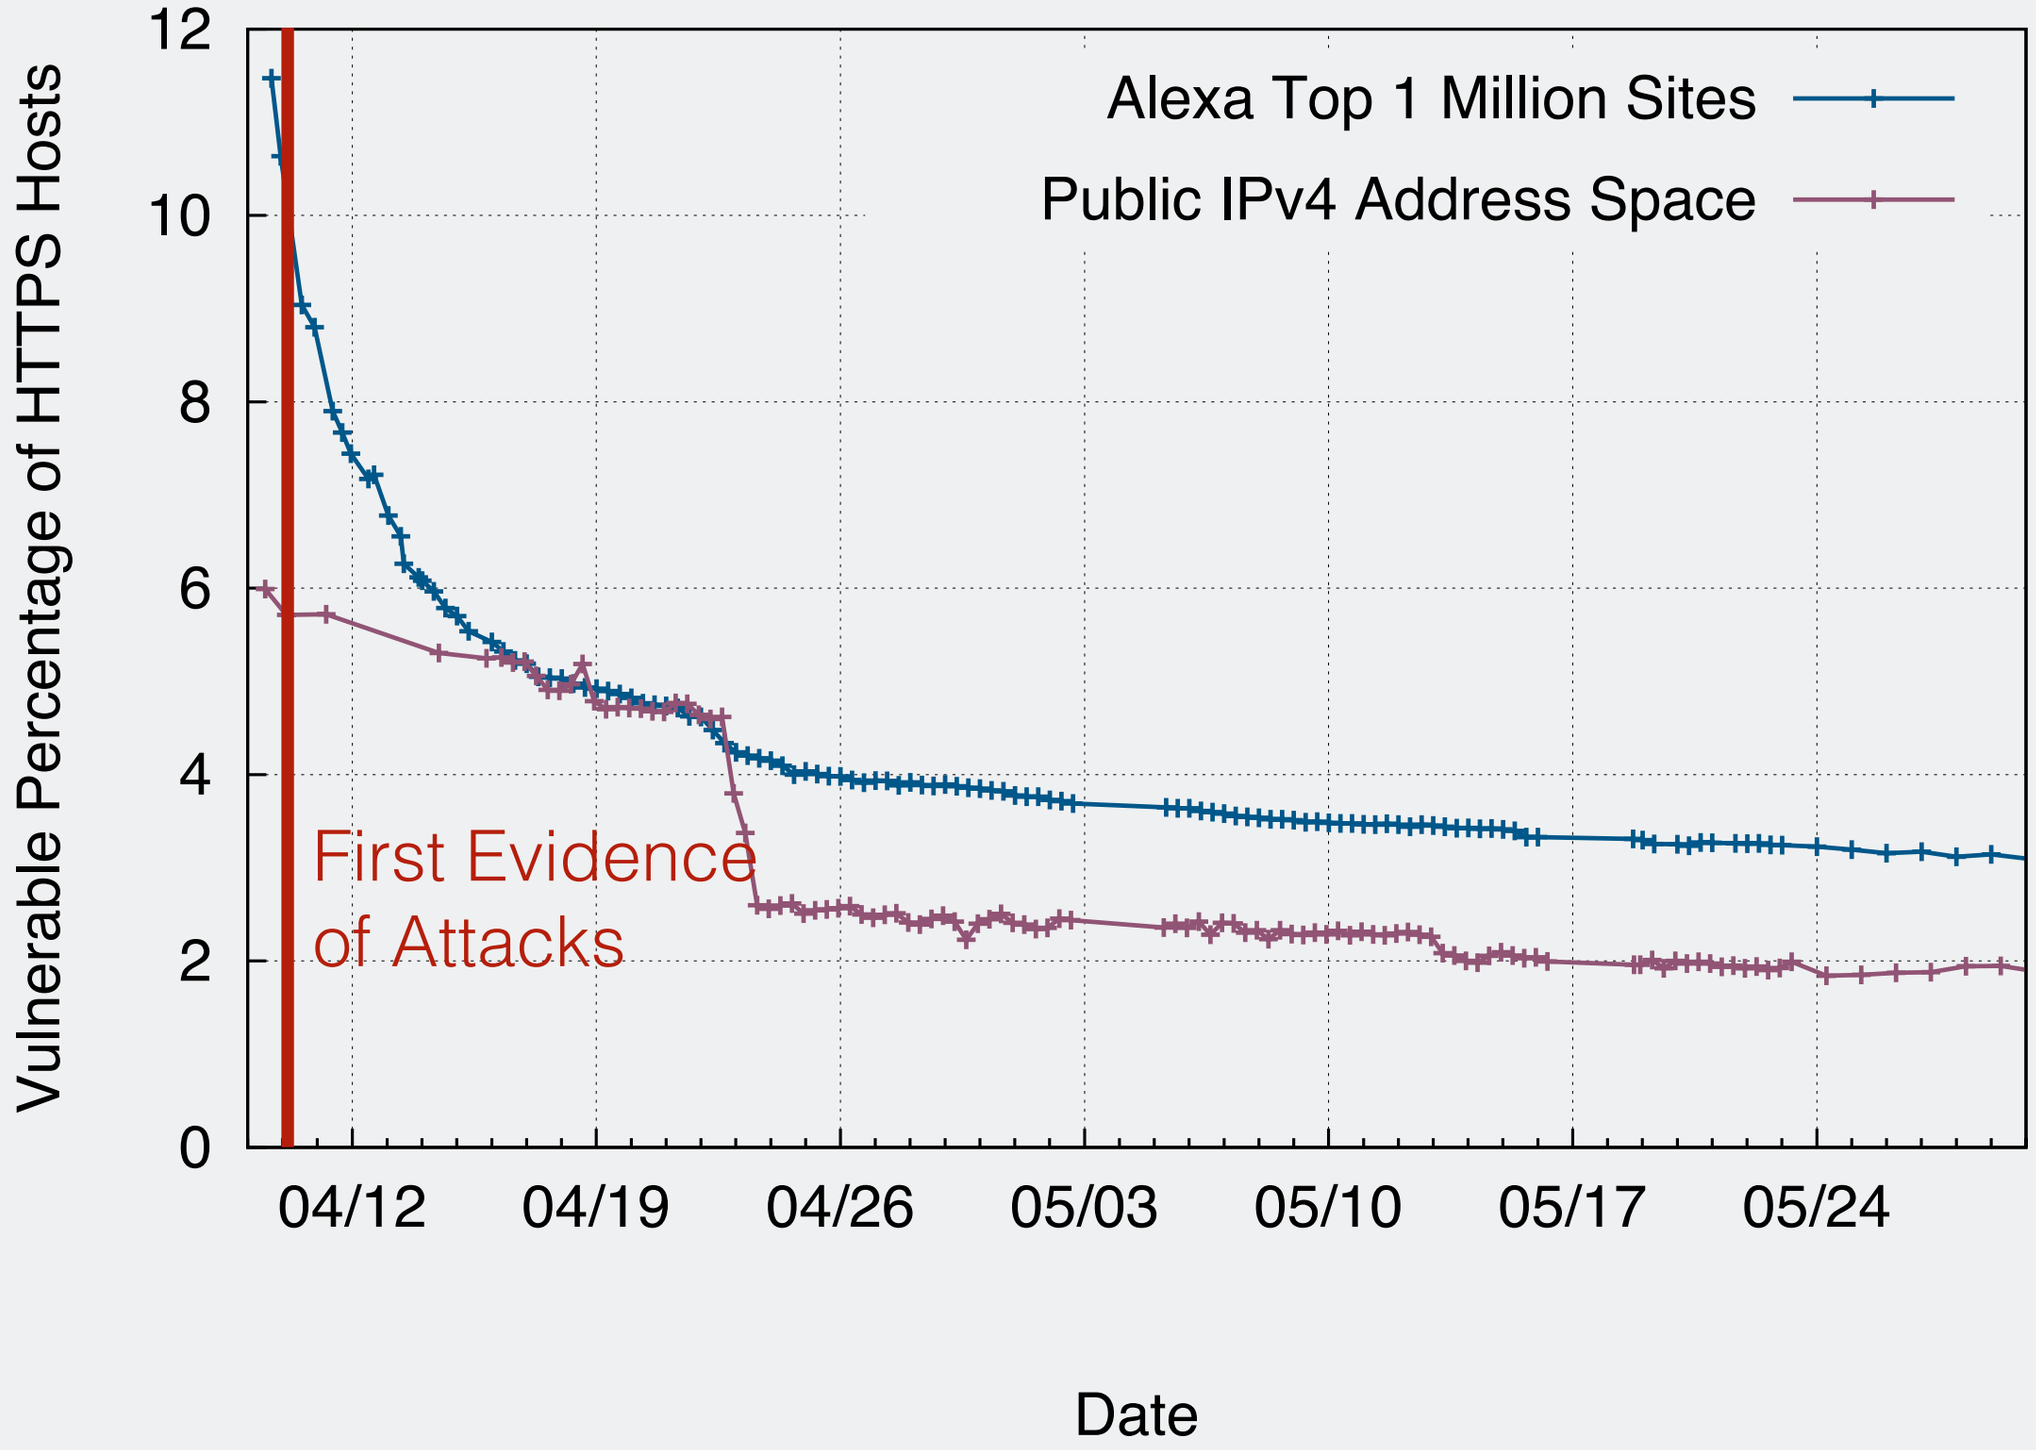
Vulnerable Percentage of HTTPS Hosts
12
10
8
6
4
2
0
04/12
04/19
04/26
05/03
05/10
05/17
05/24
Date
Alexa Top 1 Million Sites
Public IPv4 Address Space
First Evidence
of Attacks

What happens if you email everybody on the Internet who is vulnerable and tell them to patch?

What happens if you email everybody on the Internet who is vulnerable and tell them to patch?

Only two people threaten to sue you!

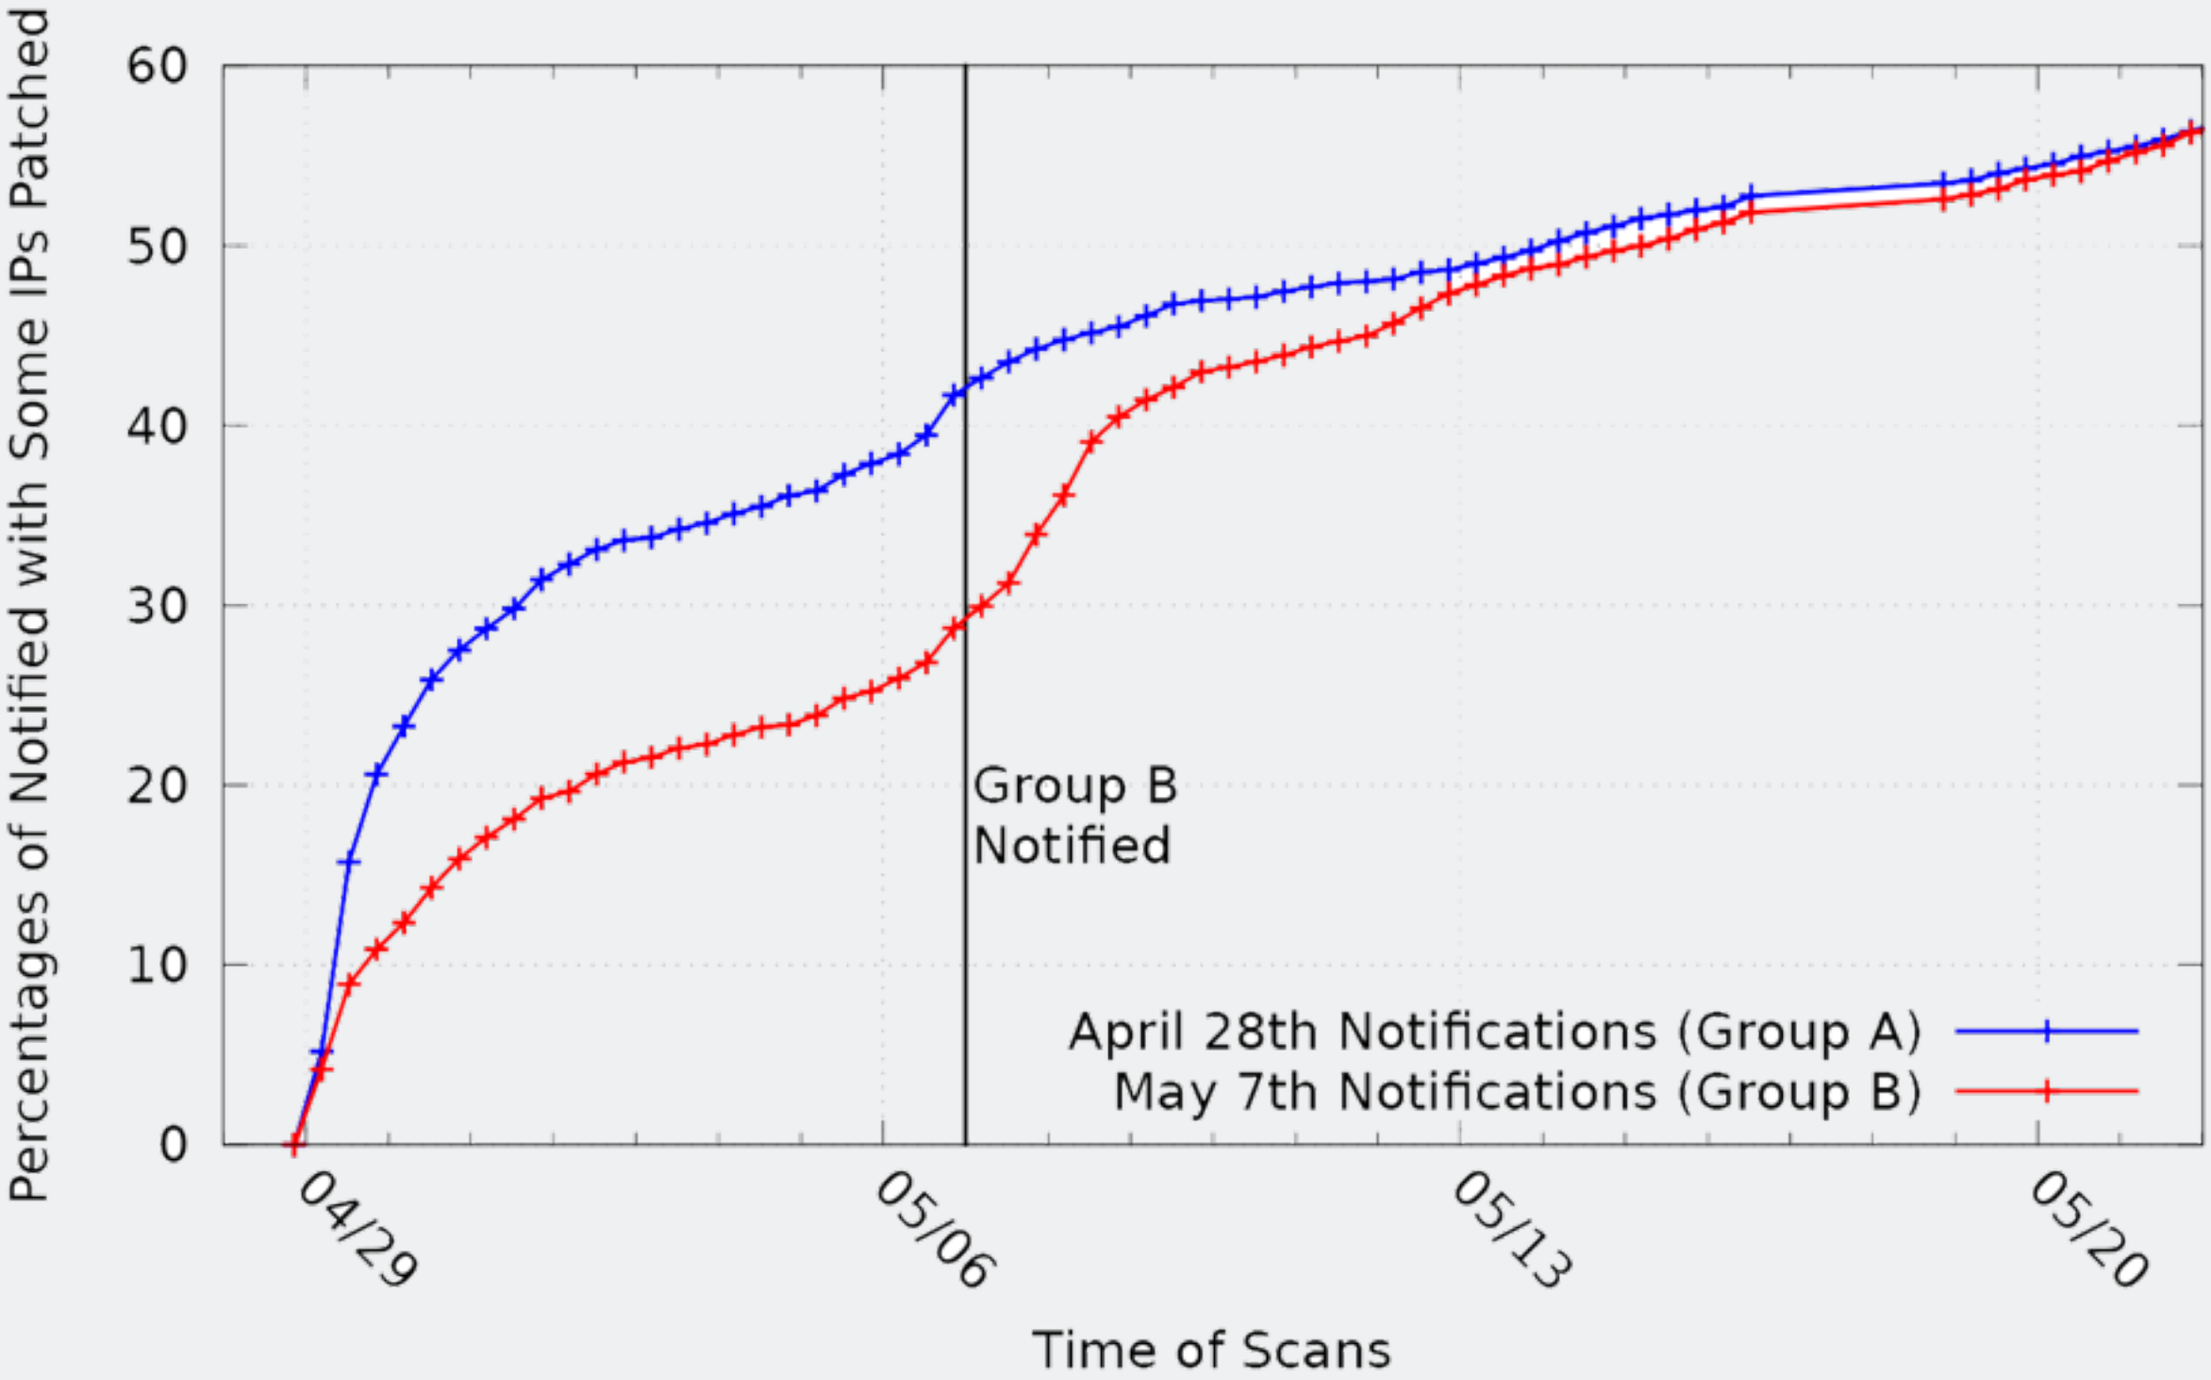
Percentages of Notified with Some IPs Patched
60
50
40
30
20
10
0
04/29
05/06
05/13
05/20
Time of Scans
Group B
Notified
April 28th Notifications (Group A)
May 7th Notifications (Group B)

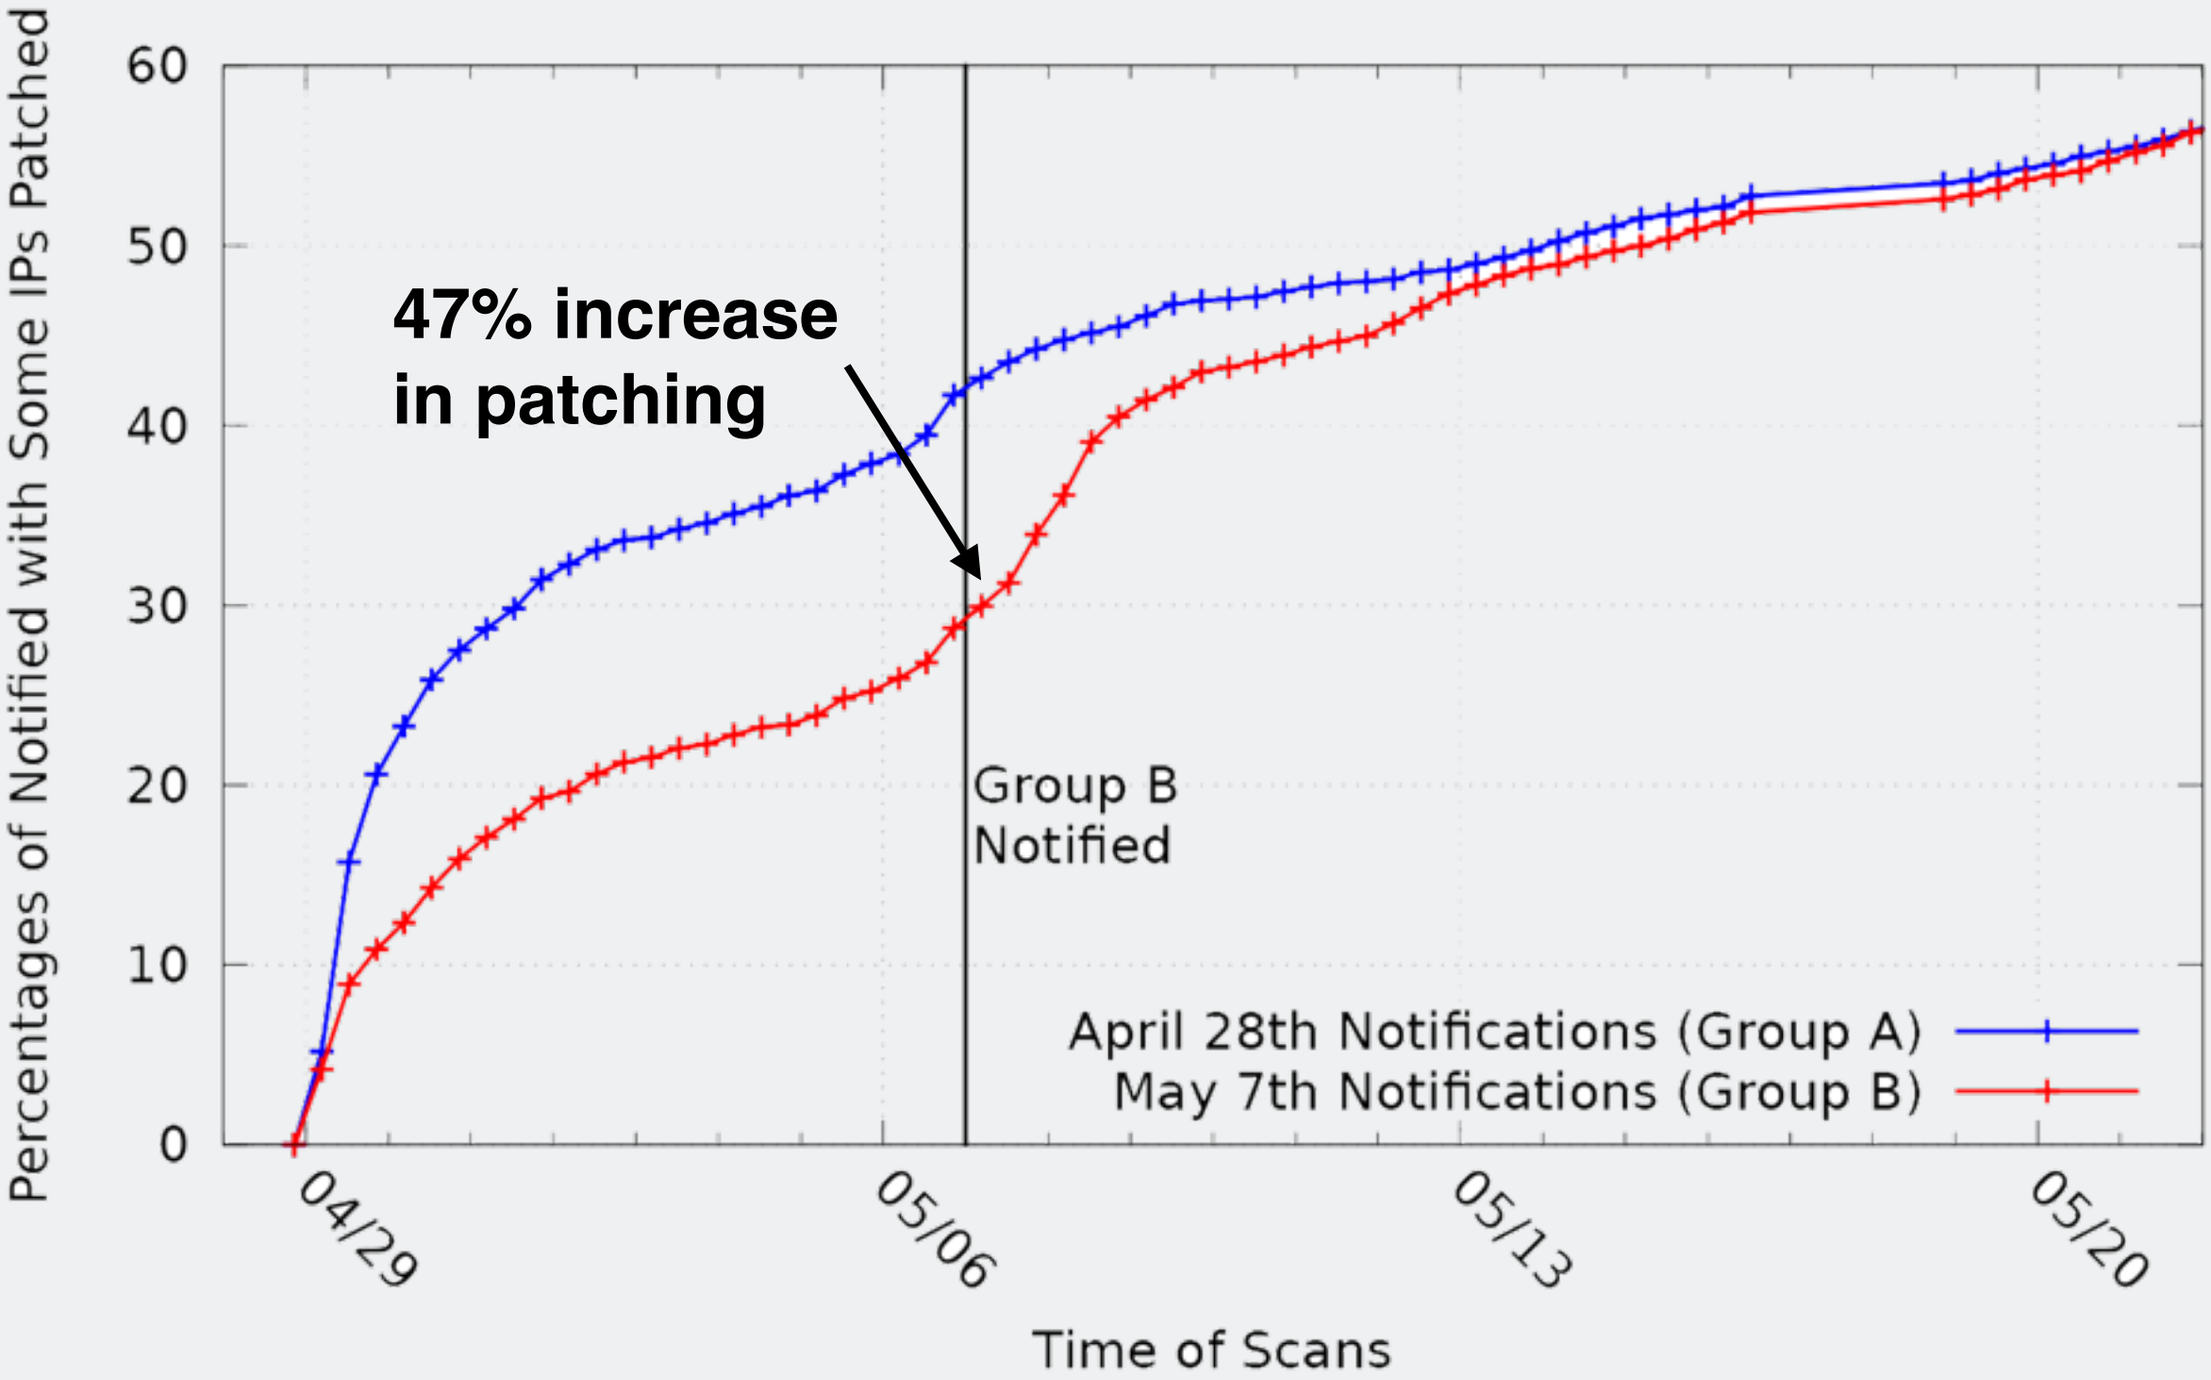
Percentages of Notified with Some IPs Patched
60
50
40
30
20
10
0
47% increase
in patching
Group B
Notified
April 28th Notifications (Group A)
May 7th Notifications (Group B)
04/29
05/06
05/13
05/20
Time of Scans

How do we **better defend** TLS in the future? Can we use ZMap and measurement to proactively identify new ways that TLS can fail in practice?

How do we **better defend** TLS in the future? Can we use ZMap and measurement to proactively identify new ways that TLS can fail in practice?

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 that allows a man-in-the-middle to downgrade connections to "export-grade" Diffie-Hellman. To carry out this attack, we implement the number field sieve discrete log algorithm. After a week-long precomputation for a specified 512-bit group, we can compute arbitrary discrete logs in this group in minutes. We find that 82% of vulnerable servers use a single 512-bit group, allowing us to compromise connections to 7% of Alexa Top Million HTTPS sites. In response, major browsers are being changed to reject short groups.

We go on to consider Diffie-Hellman with 768- and 1024-bit 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 been lost among practitioners deploying cryptosystems. We exploit it to obtain the following results:

Active attacks on export ciphers in TLS. We identify a new attack on TLS, in which a man-in-the-middle attacker can downgrade a connection to export-grade cryptography. This attack is reminiscent of the FREAK attack [6], but applies to the ephemeral Diffie-Hellman ciphersuites and is a TLS protocol flaw rather than an implementation vulnerability. We present measurements that show that this attack applies to 8.4% of Alexa Top Million HTTPS sites and 3.4% of all

Diffie-Hellman

Two parties agree on a secure, secret key over an insecure channel

Only share public data, agree on shared private secret

Eavesdropper cannot determine the secret key

Leverages a “hard” mathematical problem known as discrete log

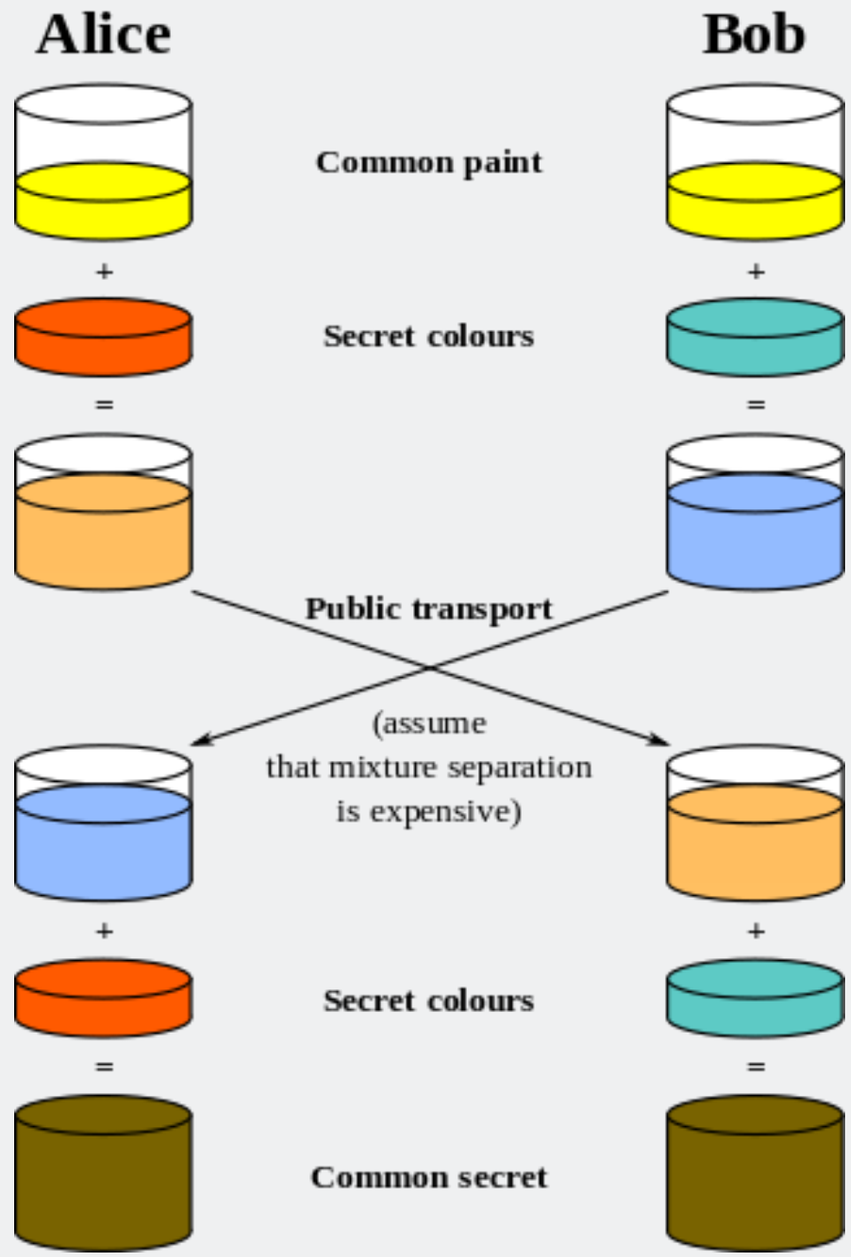

Alice and Bob agree on two numbers, ***p*** and ***g***, where p is prime. Alice picks a secret ***a***, and Bob picks ***b***.

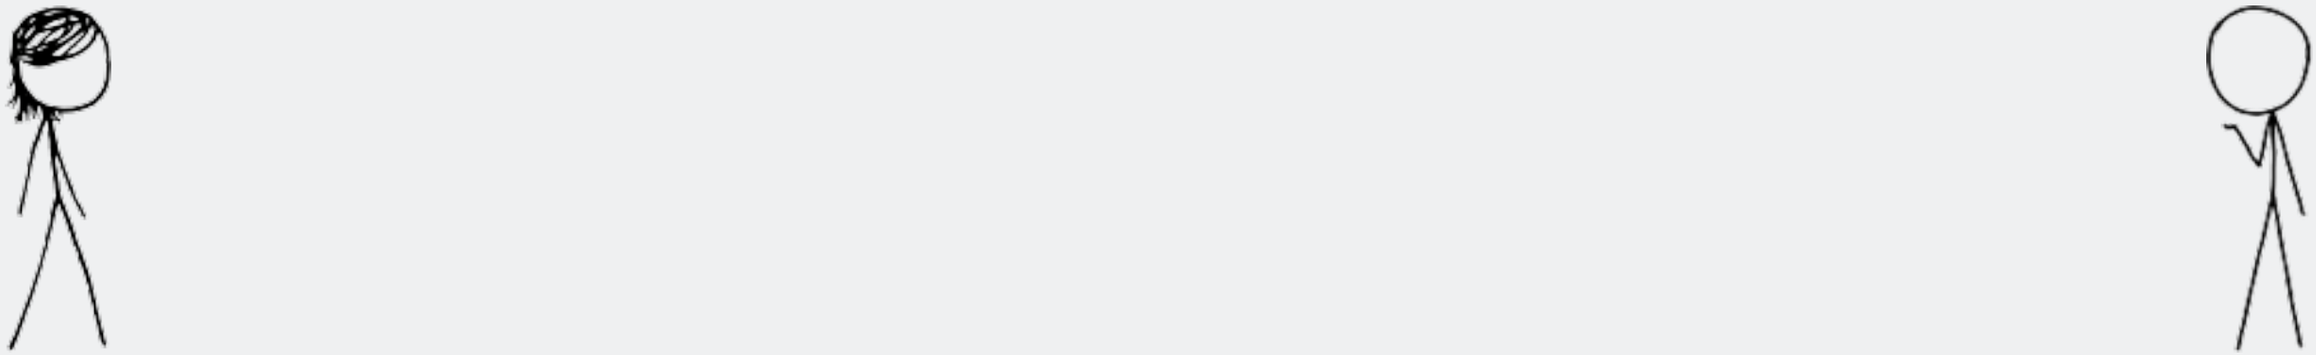

Alice and Bob agree on two numbers, ***p*** and ***g***, where p is prime. Alice picks a secret ***a***, and Bob picks ***b***.

$$A = g^a \bmod p$$

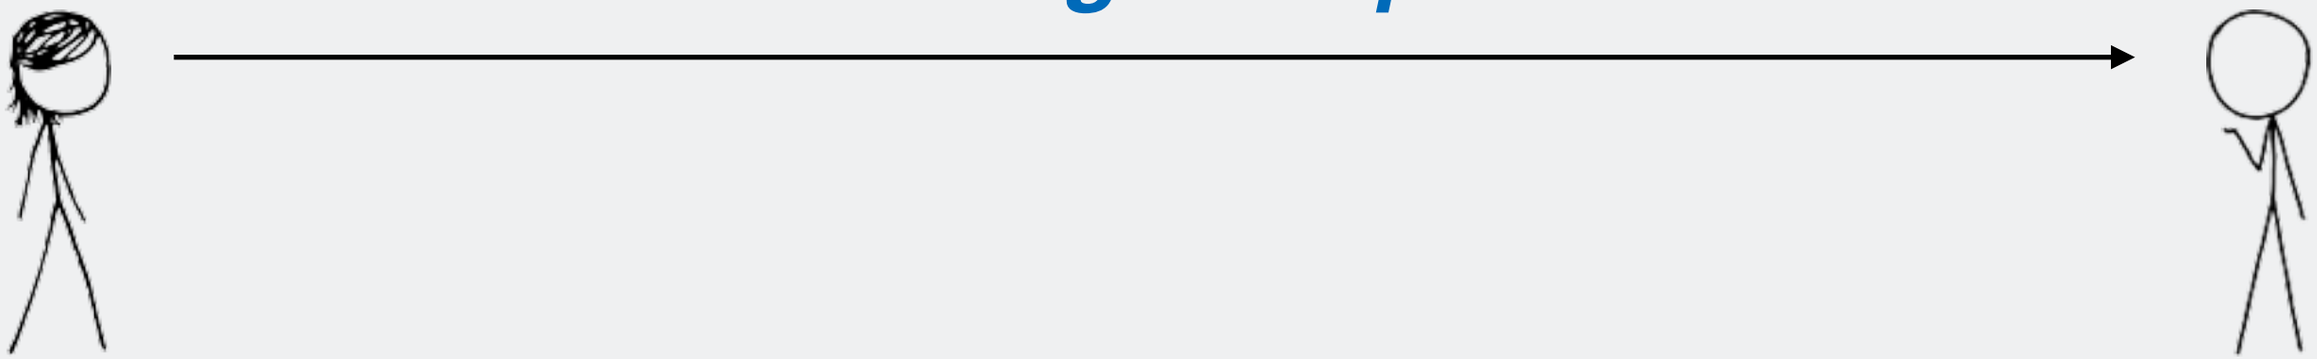

Alice and Bob agree on two numbers, ***p*** and ***g***, where p is prime. Alice picks a secret ***a***, and Bob picks ***b***.

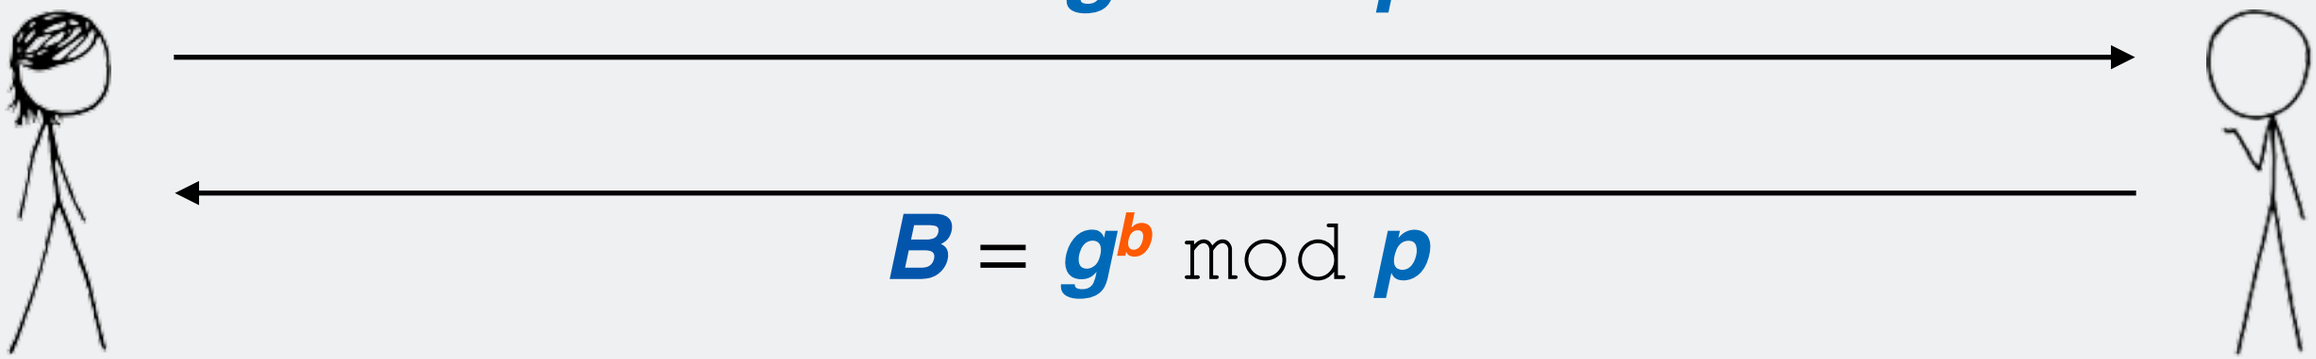

Alice and Bob agree on two numbers, ***p*** and ***g***, where p is prime. Alice picks a secret ***a***, and Bob picks ***b***.

$$A = g^a \bmod p$$

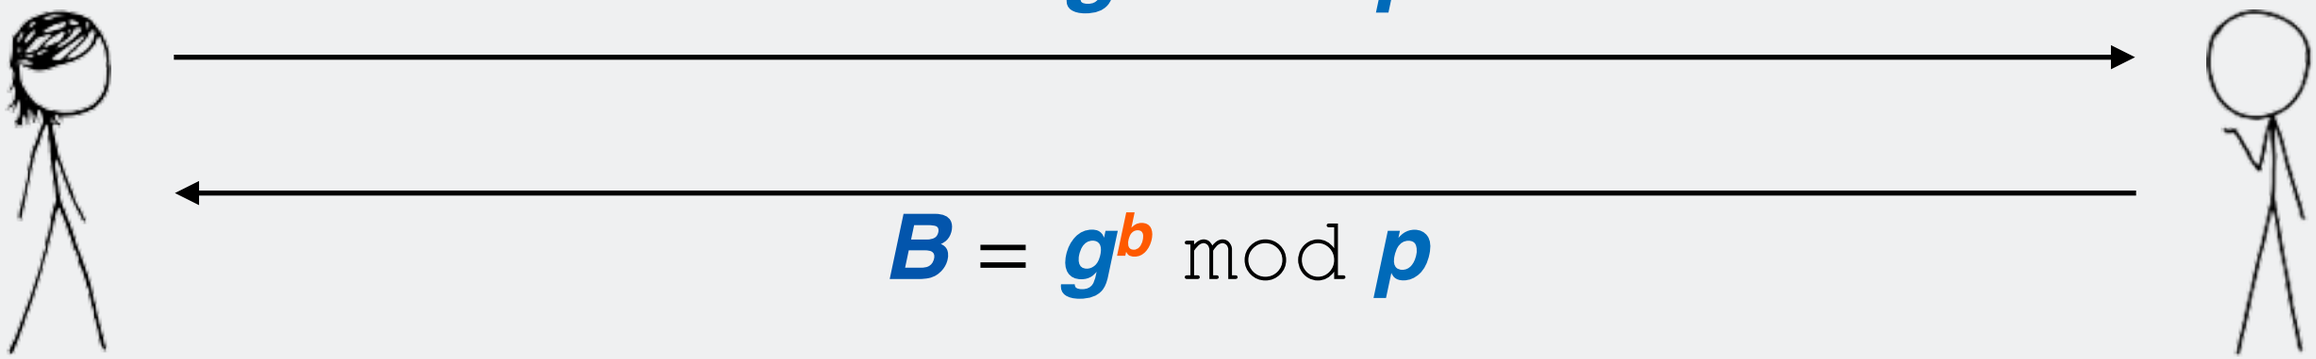

$$B = g^b \bmod p$$

$$s = B^a \bmod p$$

$$s = A^b \bmod p$$

Alice and Bob agree on two numbers, ***p*** and ***g***, where p is prime. Alice picks a secret ***a***, and Bob picks ***b***.

$$A = g^a \bmod p$$

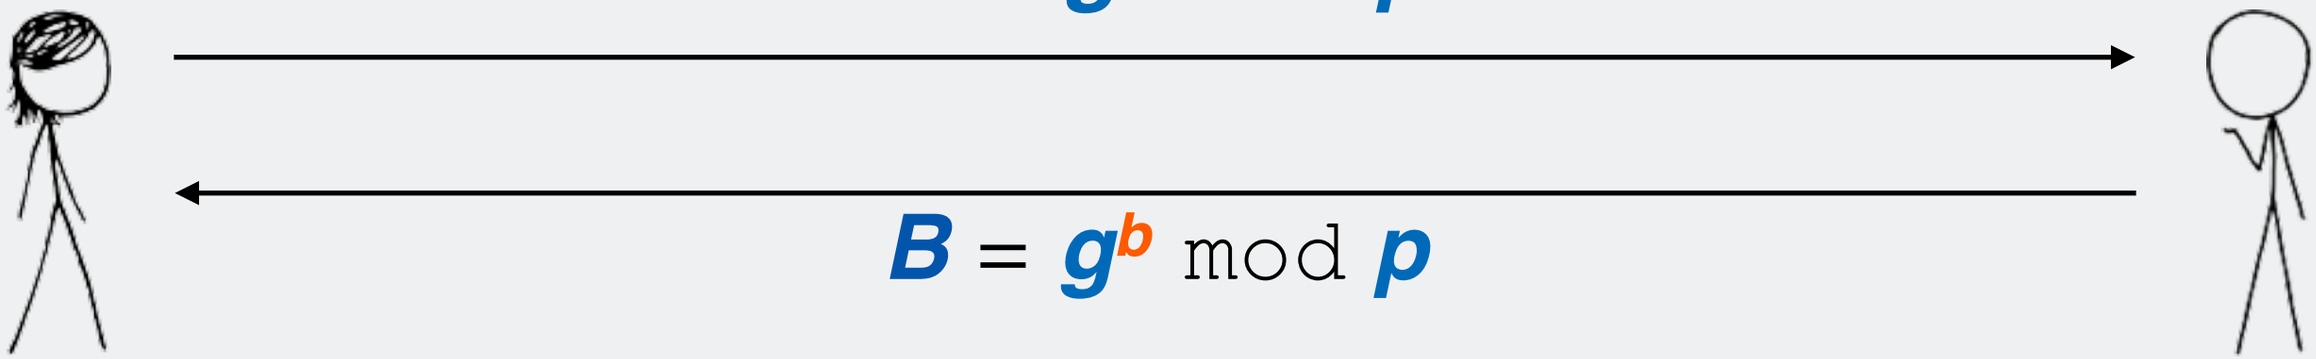

$$B = g^b \bmod p$$

$$s = B^a \bmod p$$

$$g^{ab} \bmod p$$

$$s = A^b \bmod p$$

$$g^{ab} \bmod p$$

Why does it work?

To break Diffie-Hellman, you need to be able to calculate $g^{ab} \bmod p$, given only $g^a \bmod p$ and $g^b \bmod p$

This is the **discrete log** problem. Mathematicians tell us this is hard and infeasible to compute, so long as p is **sufficiently large**.

Recommended to use 2048-bit primes or higher.

You can use the **number-field sieve** algorithm to calculate discrete log and break Diffie-Hellman.

It turns out the algorithm almost entirely depends on ***p***, not ***a*** or ***b***.

Feasible for academics to break **512-bit** primes.

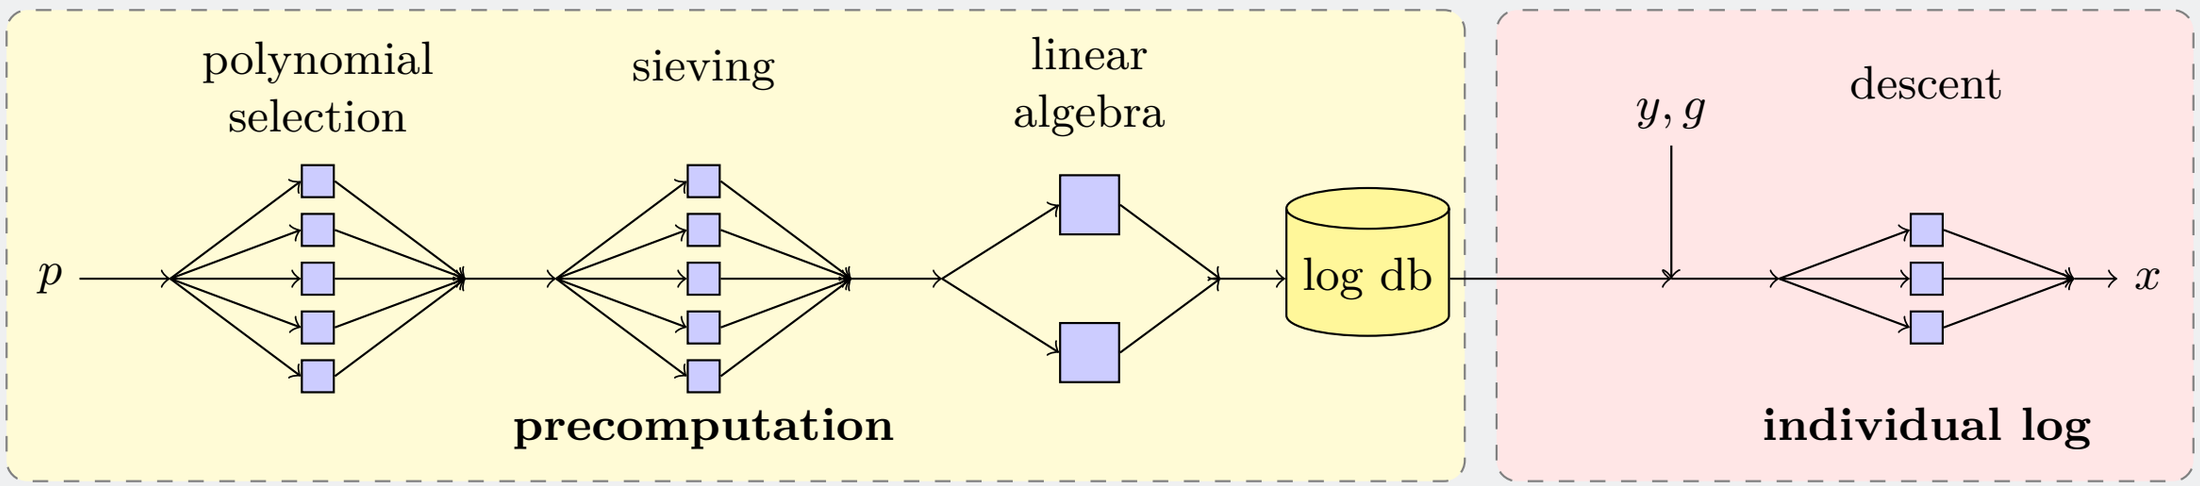

Two-Factor Authentication

DUO SECURITY, INC. [US] https://www.duosecurity.com

DUO SECURITY, INC.

Your connection to this site is private.

Permissions | Connection

The identity of DUO SECURITY, INC. at Ann Arbor, Michigan US has been verified by DigiCert SHA2 Extended Validation Server CA and is publicly auditable.

Certificate Information

Your connection to www.duosecurity.com is encrypted with modern cryptography.

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.

Site information

You have never visited this site before today.

What do these mean?

Pricing Resources Documentation About Blog Sign U

LABLE

naos with the

ed 2FA solution.

Learn More

What is Platform?

Two-Factor Authentication

DUO SECURITY, INC. [US] https://www.duosecurity.com

DUO SECURITY, INC.

Your connection to this site is private.

Permissions | Connection

The identity of DUO SECURITY, INC. at Ann Arbor, Michigan US has been verified by DigiCert SHA2 Extended Validation Server CA and is publicly auditable.

Certificate Information

Your connection to www.duosecurity.com is encrypted with modern cryptography.

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.

Site information
You have never visited this site before today.

What do these mean?

Certificate

Pricing Resources Documentation About Blog Sign U

LABLE

naos with the
ed 2FA solution.

Learn More

What is Platform?

Two-Factor Authentication

DUO SECURITY, INC. [US] https://www.duosecurity.com

DUO SECURITY, INC.

Your connection to this site is private.

Permissions | Connection

The identity of DUO SECURITY, INC. at Ann Arbor, Michigan US has been verified by DigiCert SHA2 Extended Validation Server CA and is publicly auditable.

Certificate Information

Your connection to www.duosecurity.com is encrypted with modern cryptography.

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.

Site information

You have never visited this site before today.

What do these mean?

Certificate

Cipher Suite

Pricing Resources Documentation About Blog Sign U

LABLE

naos with the

ed 2FA solution.

Learn More

What is Platform?

Client

Server

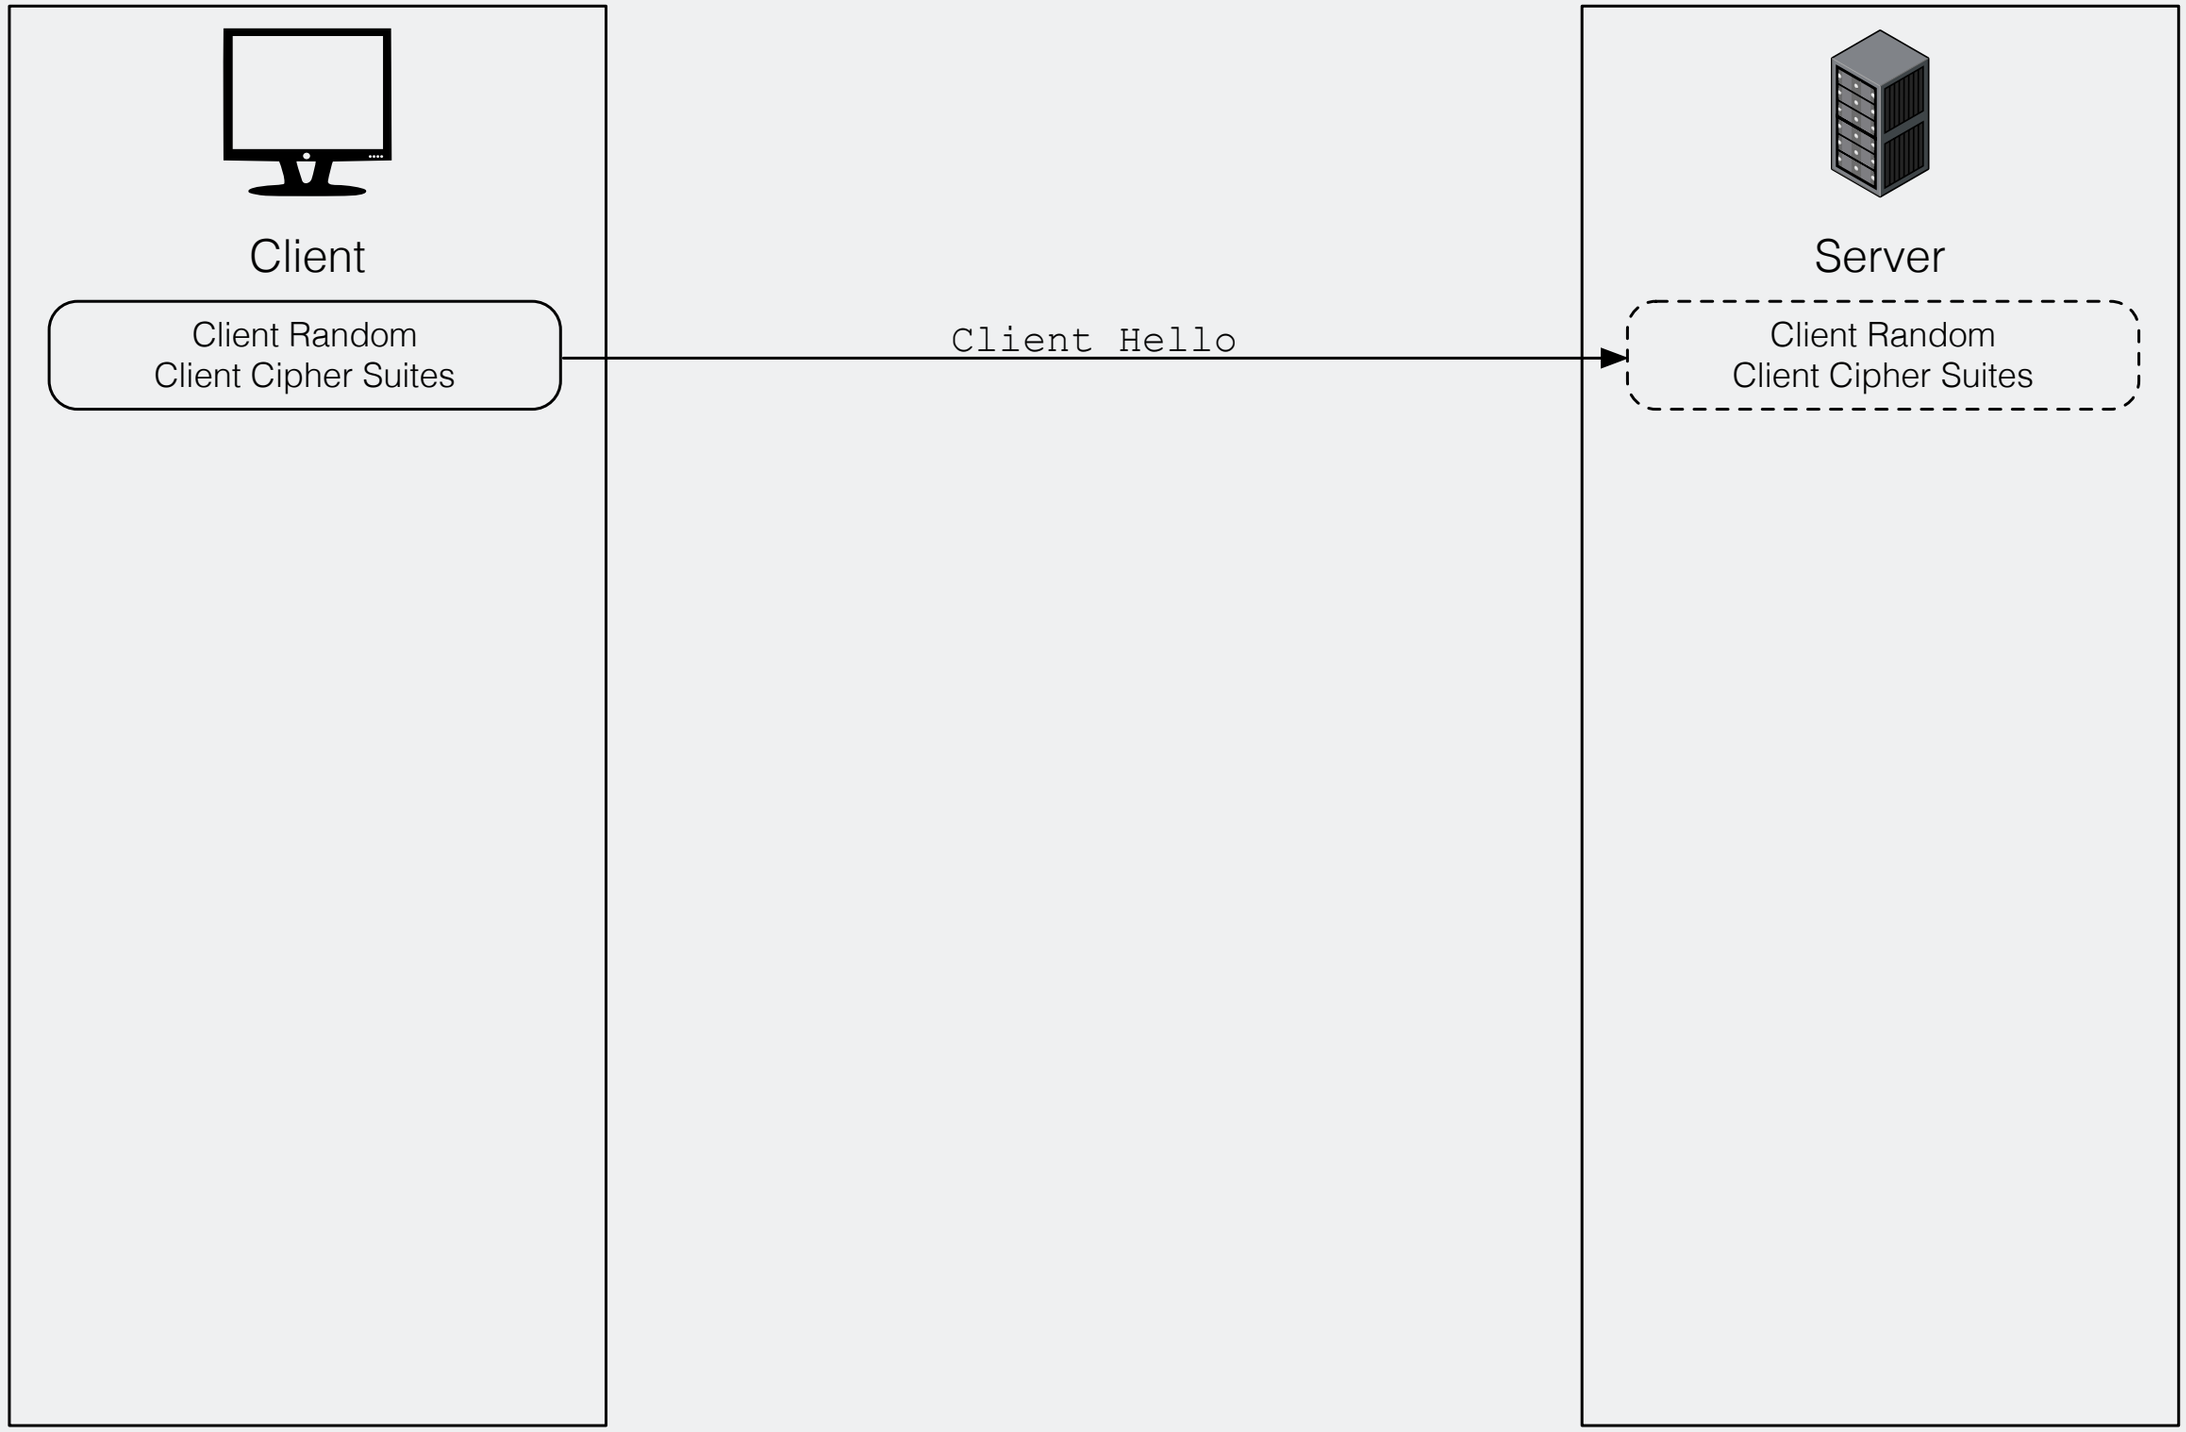
Client
Client Random
Client Cipher Suites
Client Hello
Server
Client Random
Client Cipher Suites

Client

Client Random
Client Cipher Suites

Client Hello

Server

Client Random
Client Cipher Suites

Server Random
Selected Cipher Suite

Server Certificate
Server Public Key

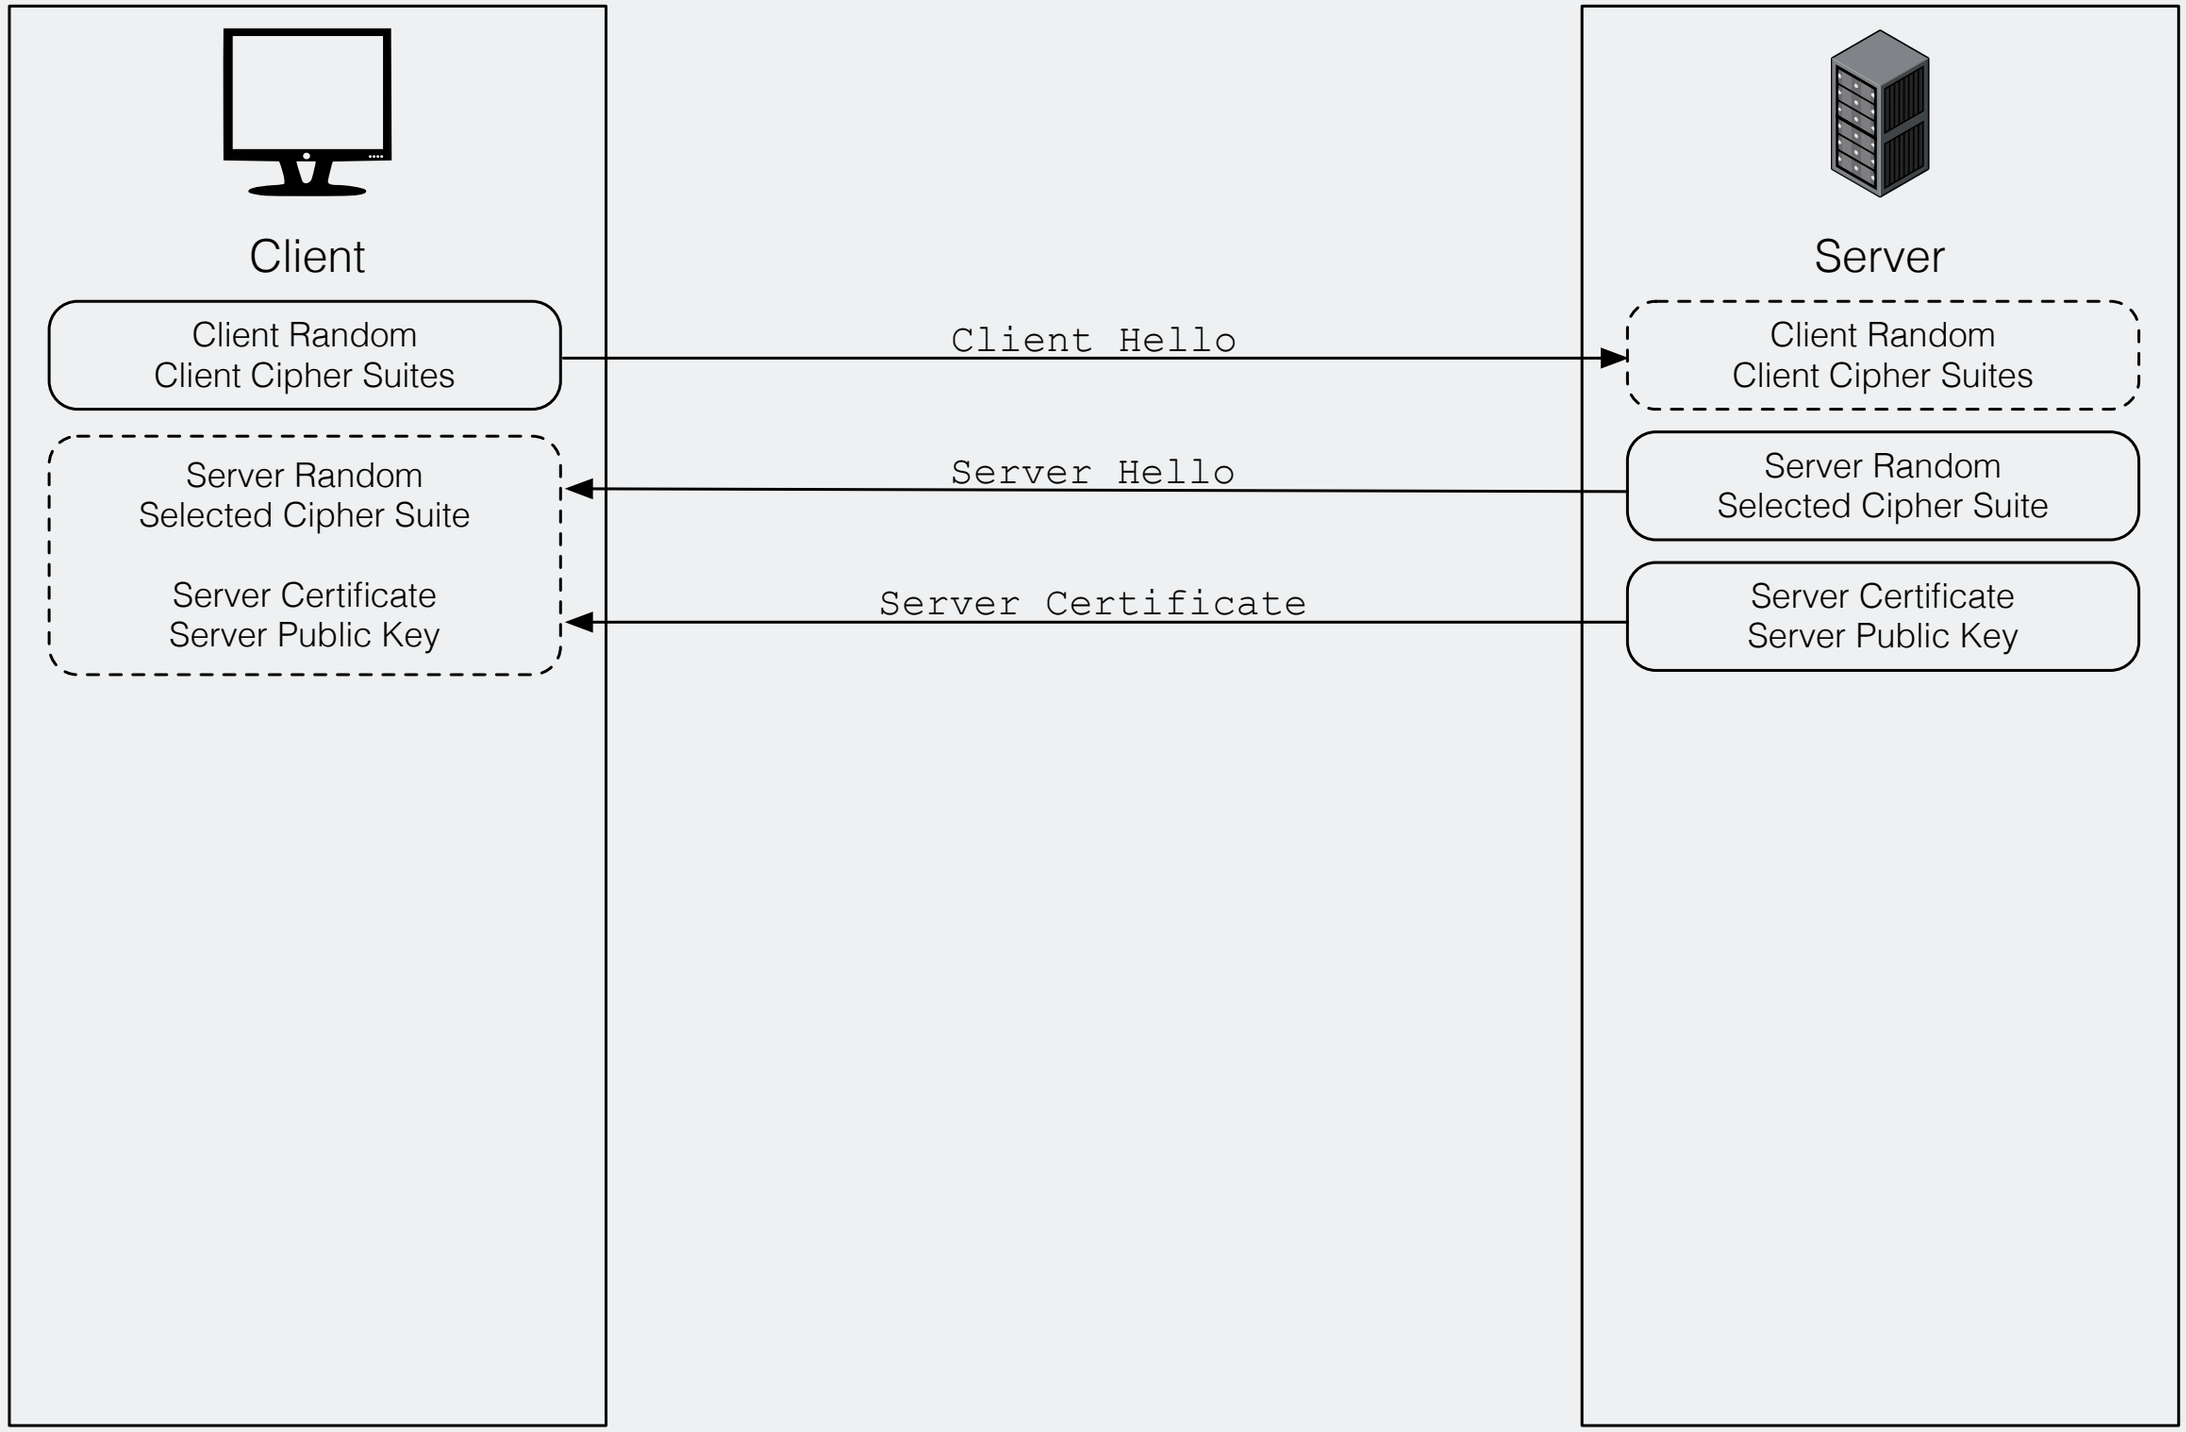
Client
Client Random
Client Cipher Suites
Server Random
Selected Cipher Suite
Server Certificate
Server Public Key
Client Hello
Server Hello
Server Certificate
Server
Client Random
Client Cipher Suites
Server Random
Selected Cipher Suite
Server Certificate
Server Public Key

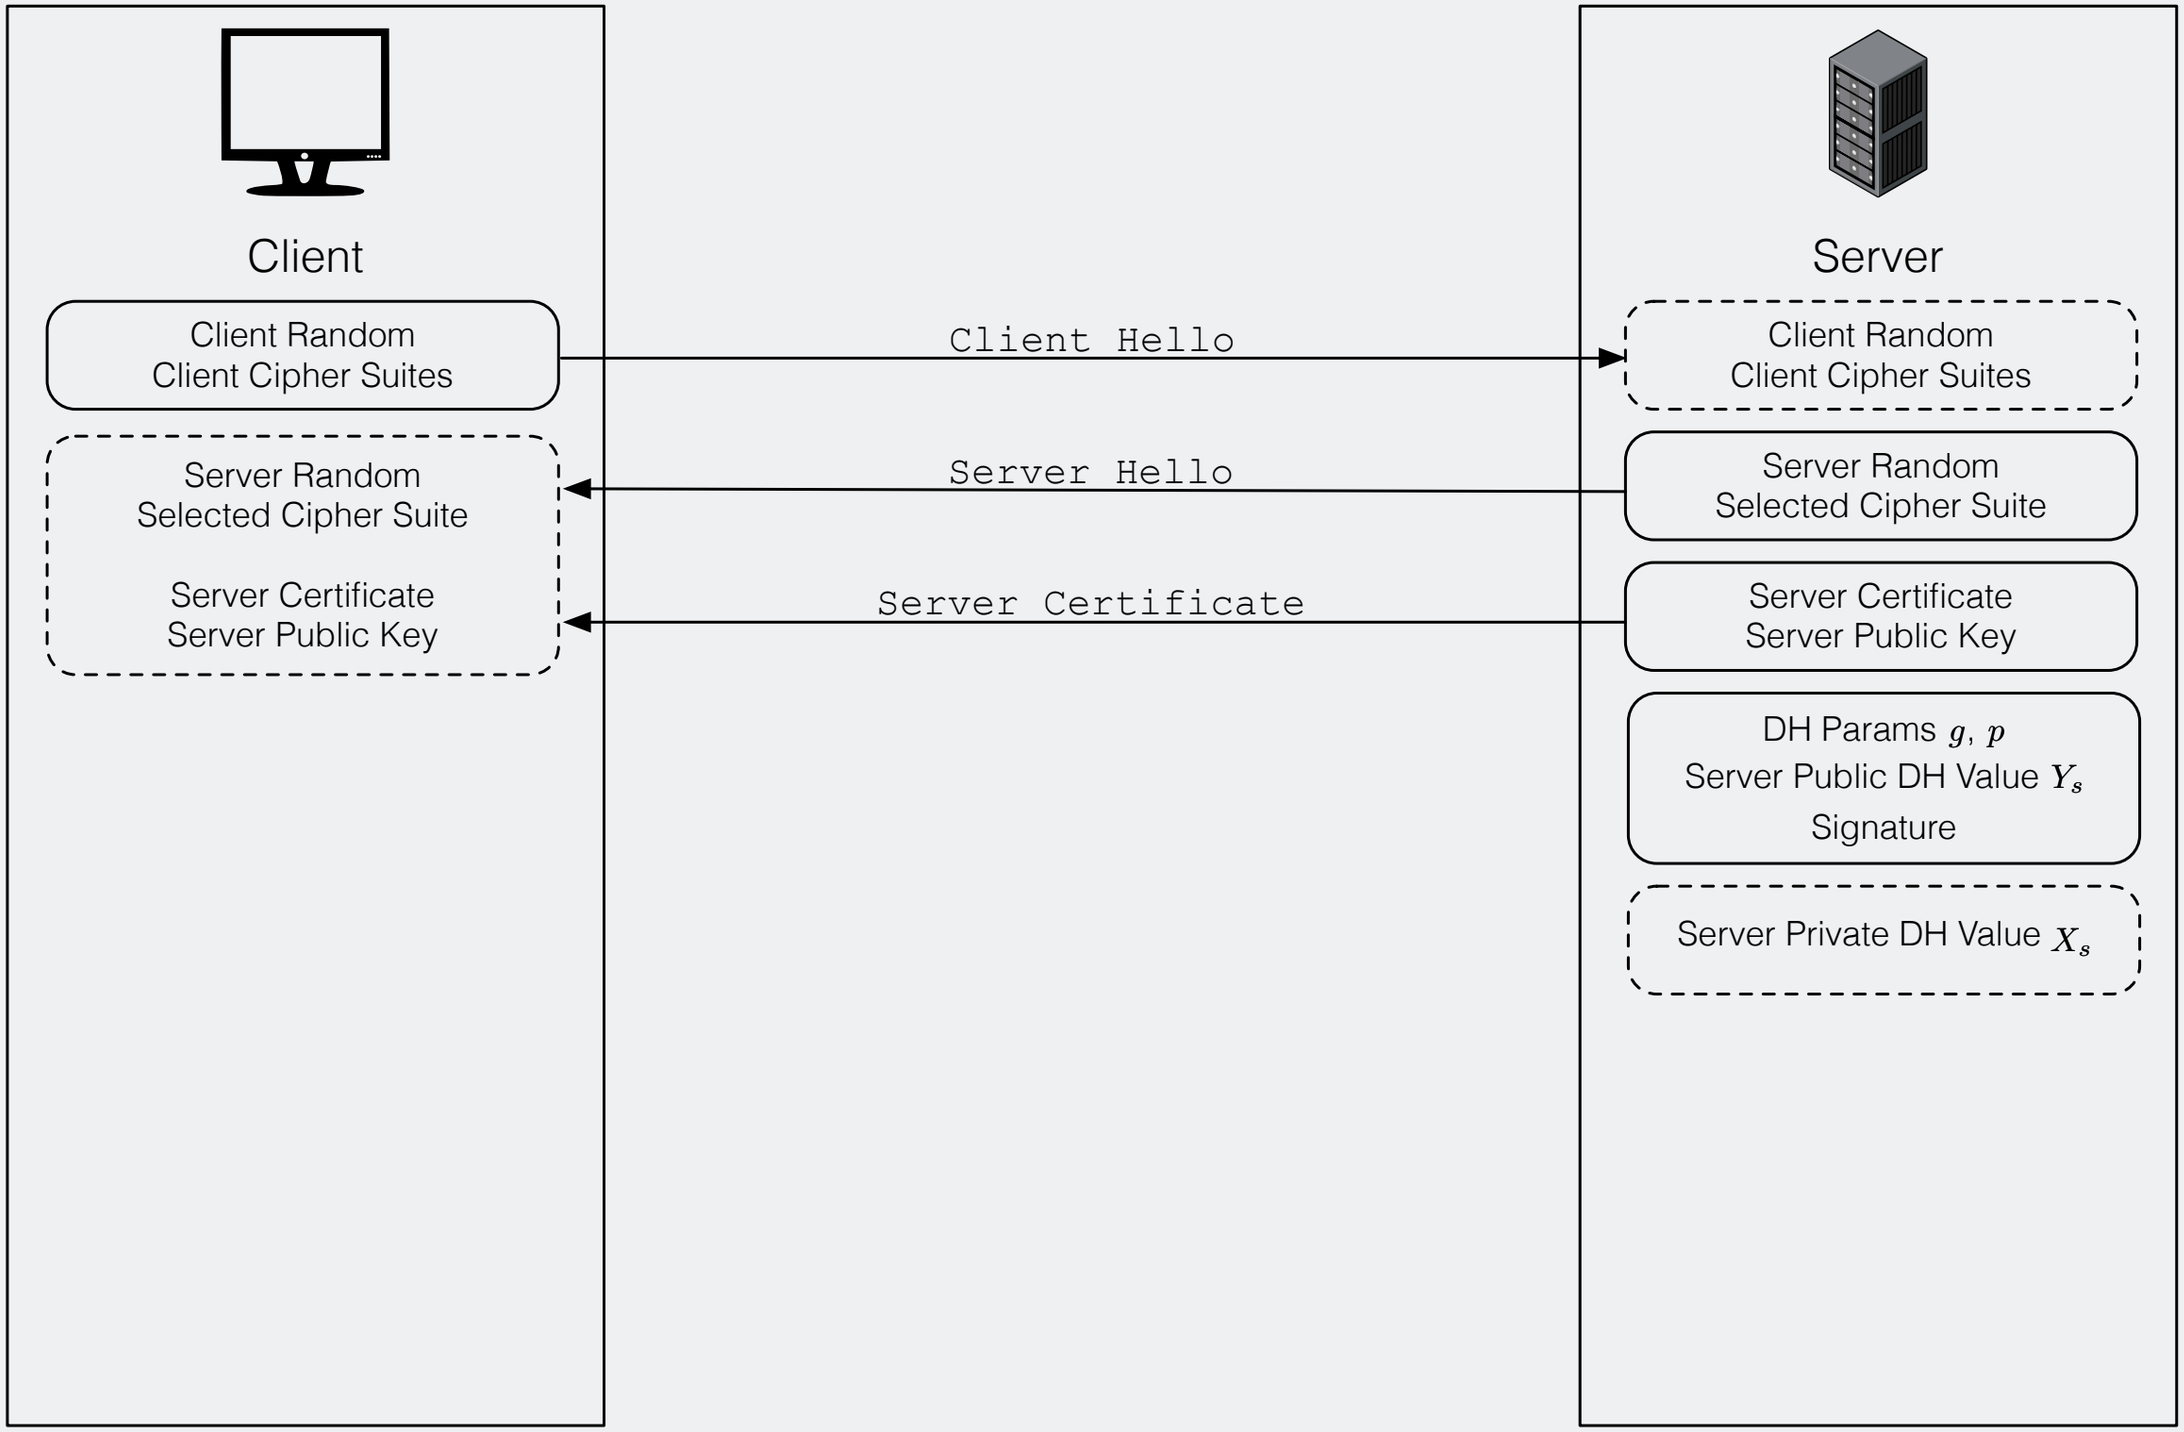
Client
Client Random
Client Cipher Suites
Server Random
Selected Cipher Suite
Server Certificate
Server Public Key
Client Hello
Server Hello
Server Certificate
Server
Client Random
Client Cipher Suites
Server Random
Selected Cipher Suite
Server Certificate
Server Public Key
DH Params g, p
Server Public DH Value Y_s
Signature
Server Private DH Value X_s

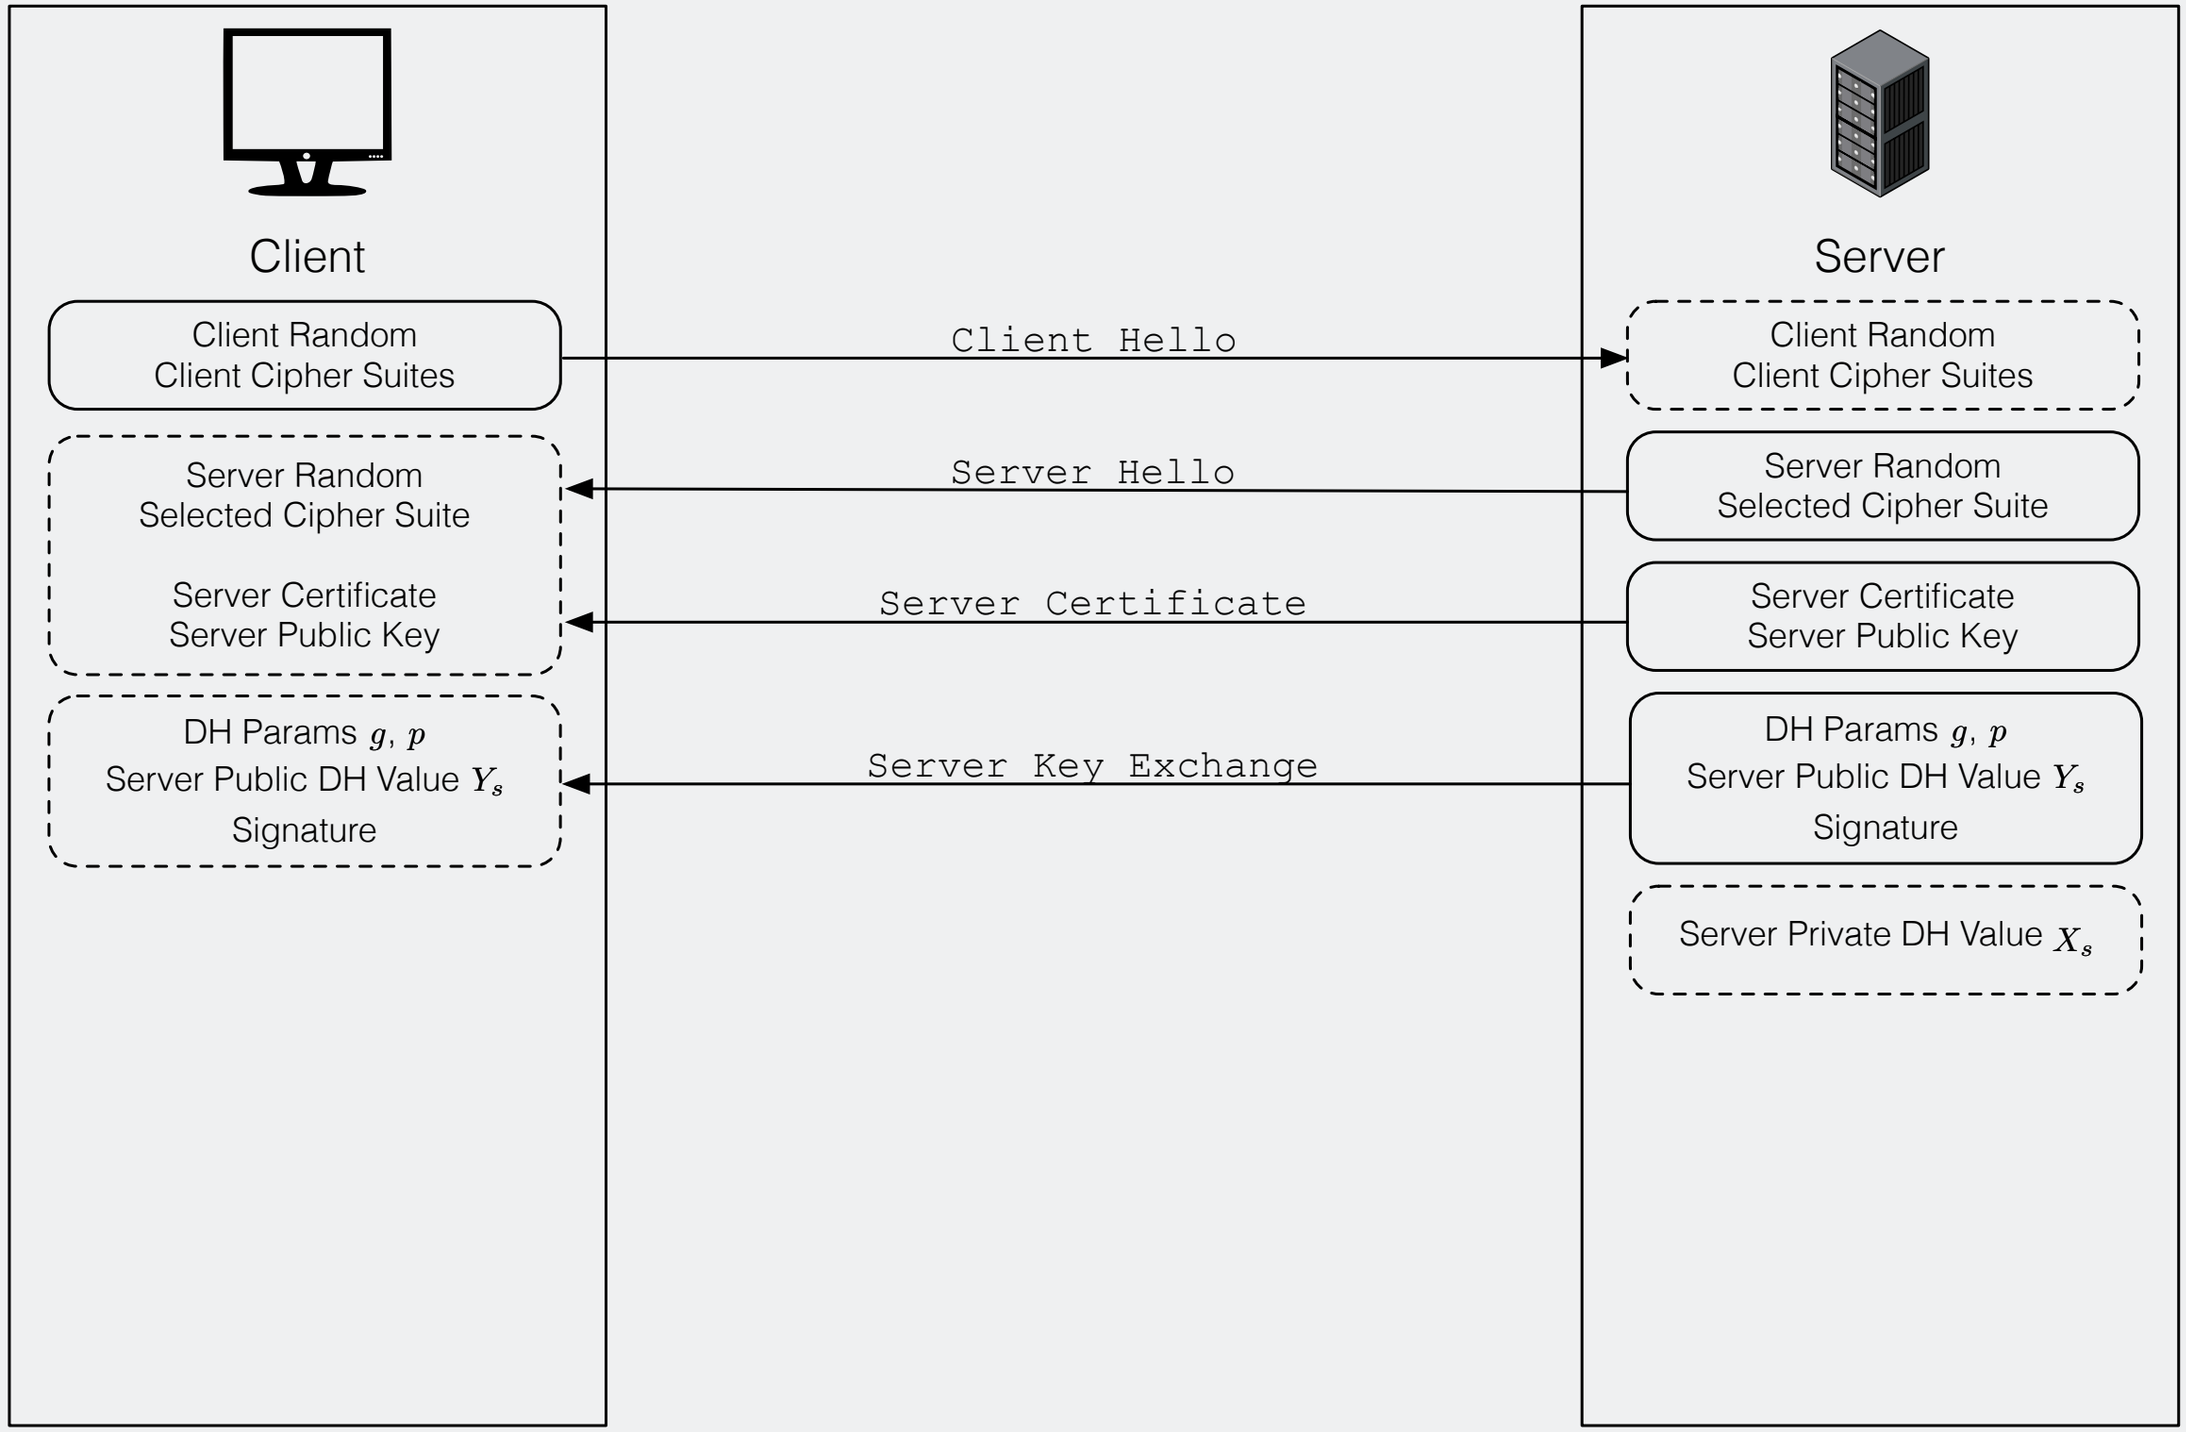
Client
Client Random
Client Cipher Suites
Server Random
Selected Cipher Suite
Server Certificate
Server Public Key
DH Params g, p
Server Public DH Value Y_s
Signature
Client Hello
Server Hello
Server Certificate
Server Key Exchange
Server
Client Random
Client Cipher Suites
Server Random
Selected Cipher Suite
Server Certificate
Server Public Key
DH Params g, p
Server Public DH Value Y_s
Signature
Server Private DH Value X_s

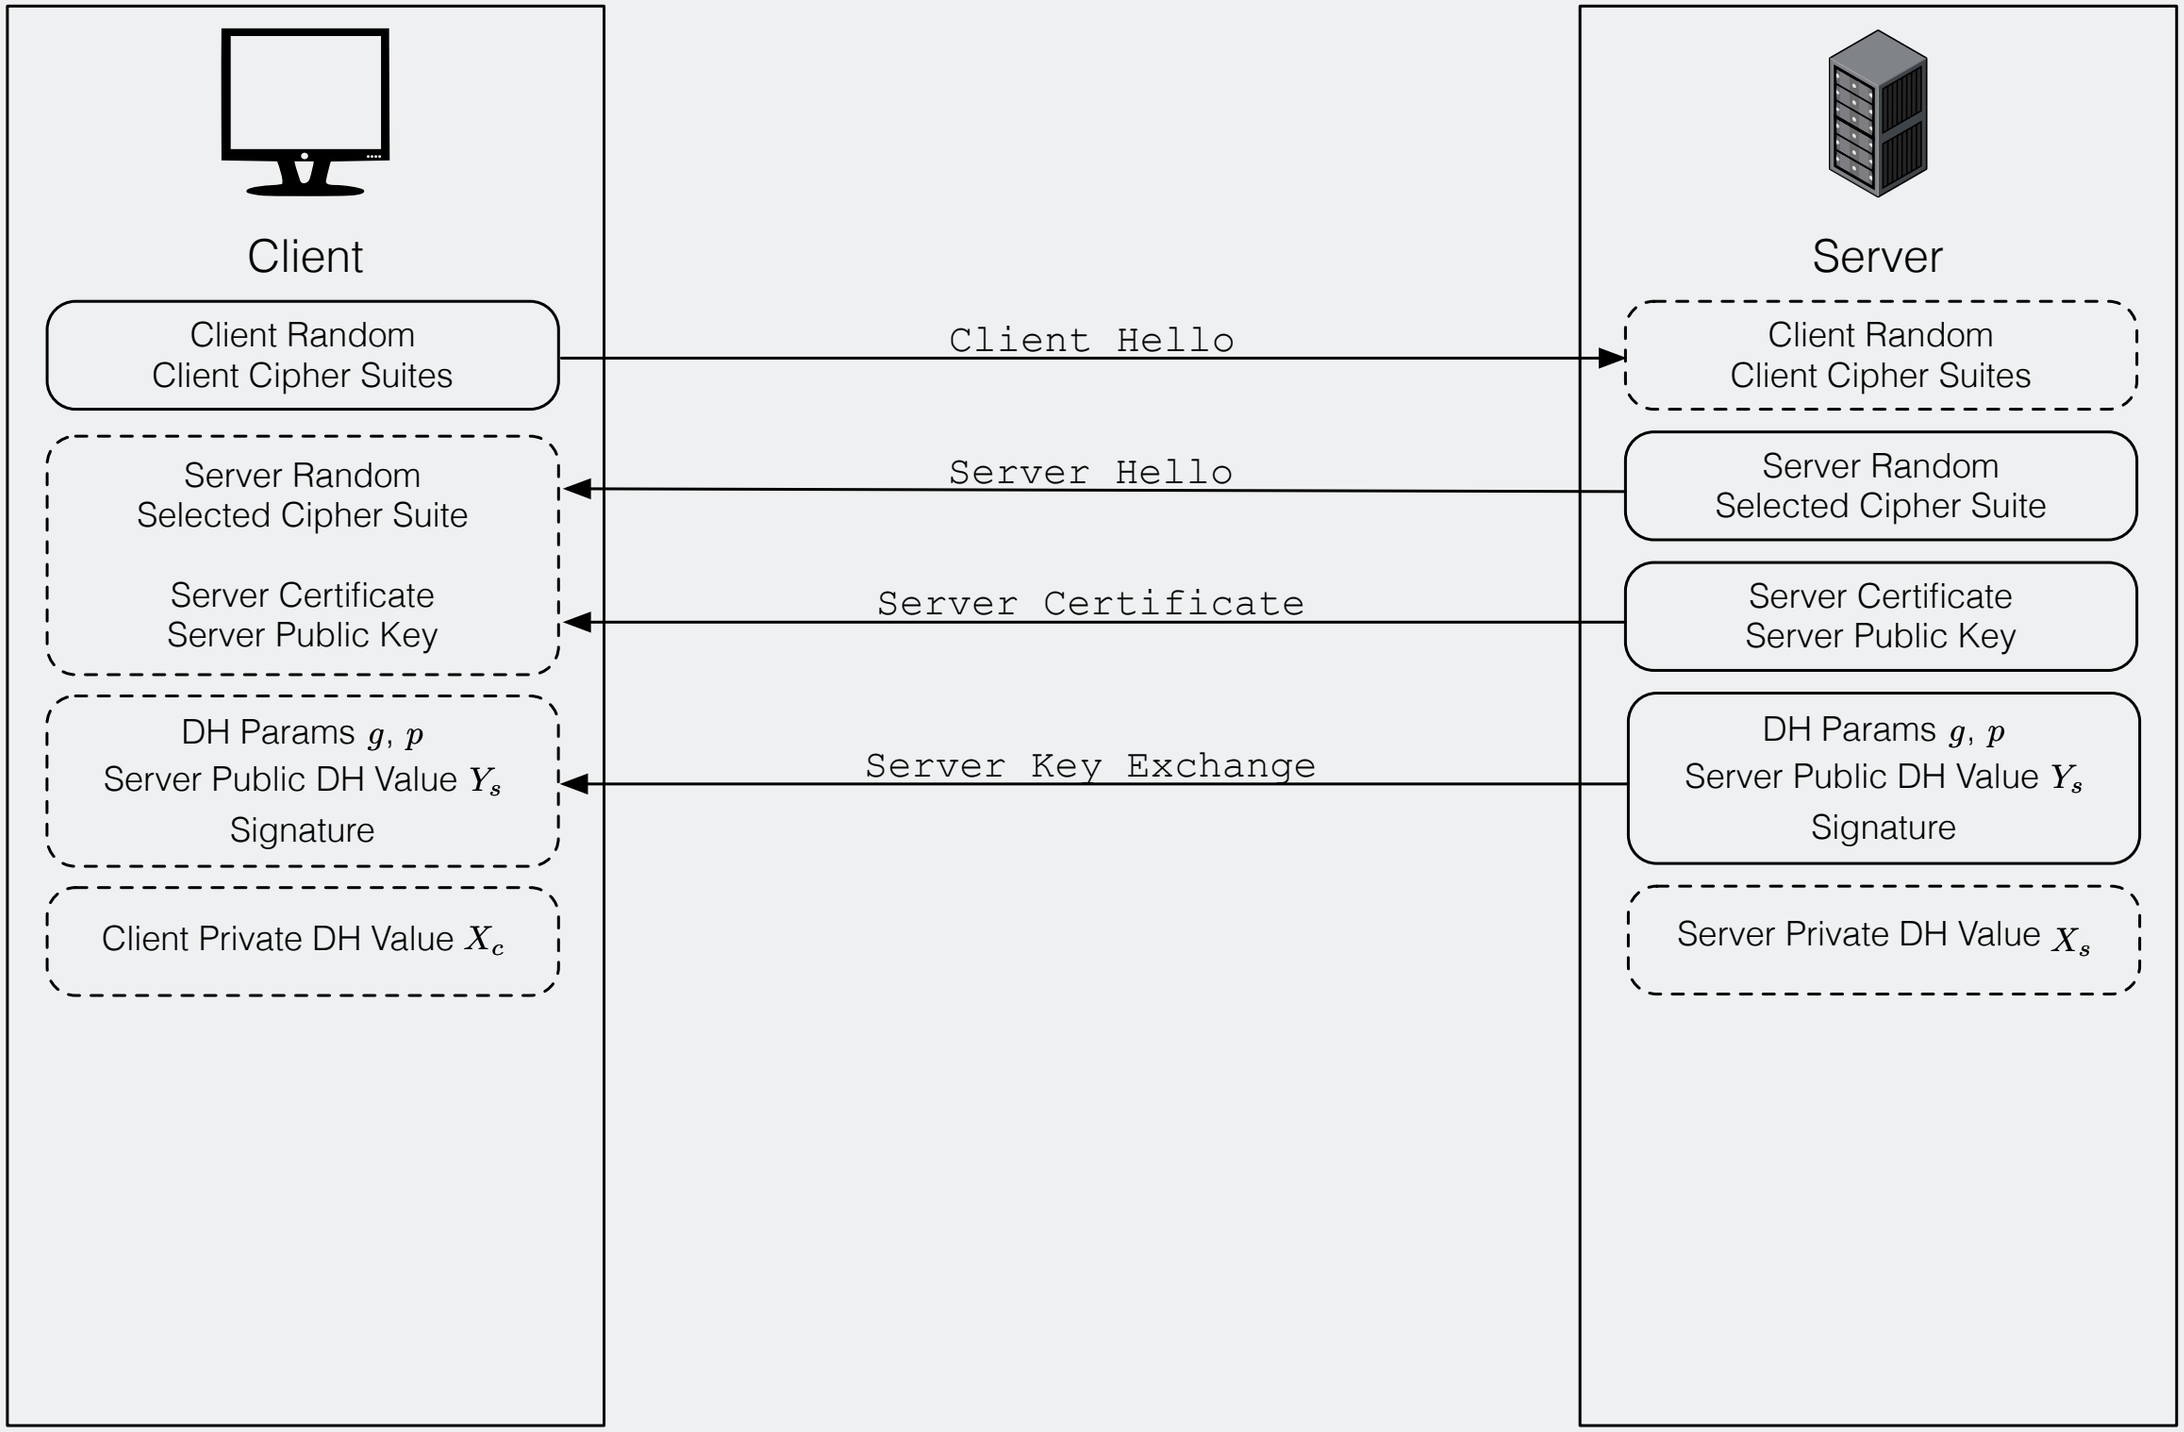
Client
Client Random
Client Cipher Suites
Server Random
Selected Cipher Suite
Server Certificate
Server Public Key
DH Params g, p
Server Public DH Value Y_s
Signature
Client Private DH Value X_c
Client Hello
Server Hello
Server Certificate
Server Key Exchange
Server
Client Random
Client Cipher Suites
Server Random
Selected Cipher Suite
Server Certificate
Server Public Key
DH Params g, p
Server Public DH Value Y_s
Signature
Server Private DH Value X_s

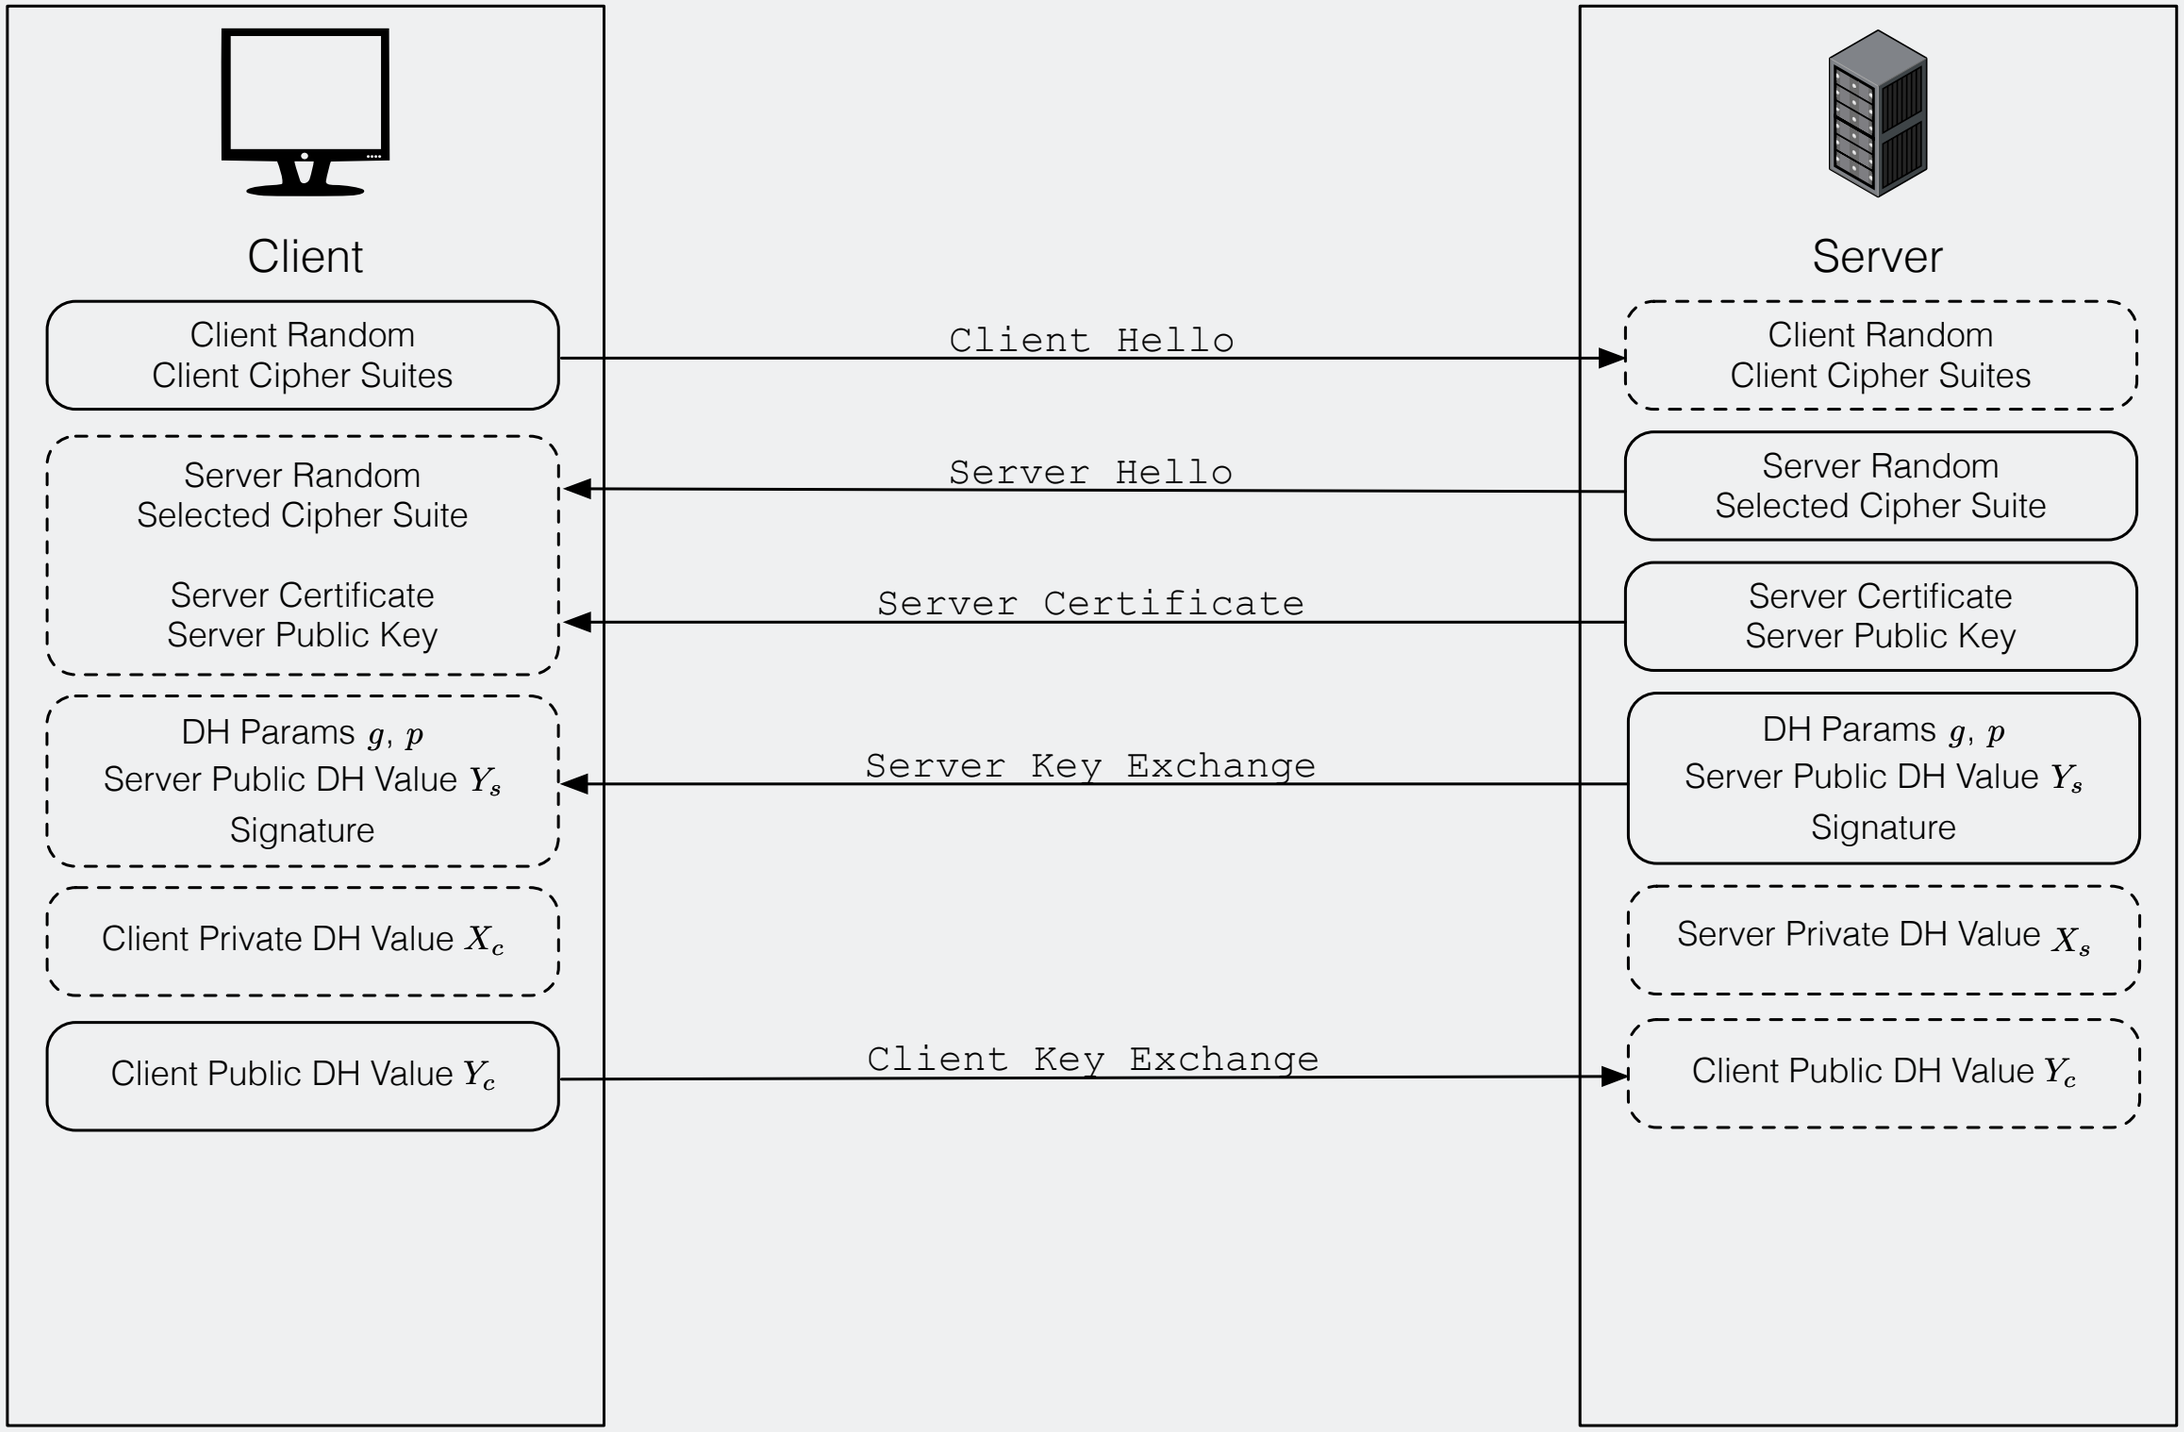
Client
Server
Client Random
Client Cipher Suites
Client Hello
Client Random
Client Cipher Suites
Server Random
Selected Cipher Suite
Server Hello
Server Random
Selected Cipher Suite
Server Certificate
Server Public Key
Server Certificate
Server Certificate
Server Public Key
DH Params g, p
Server Public DH Value Y_s
Signature
Server Key Exchange
DH Params g, p
Server Public DH Value Y_s
Signature
Client Private DH Value X_c
Server Private DH Value X_s
Client Public DH Value Y_c
Client Key Exchange
Client Public DH Value Y_c

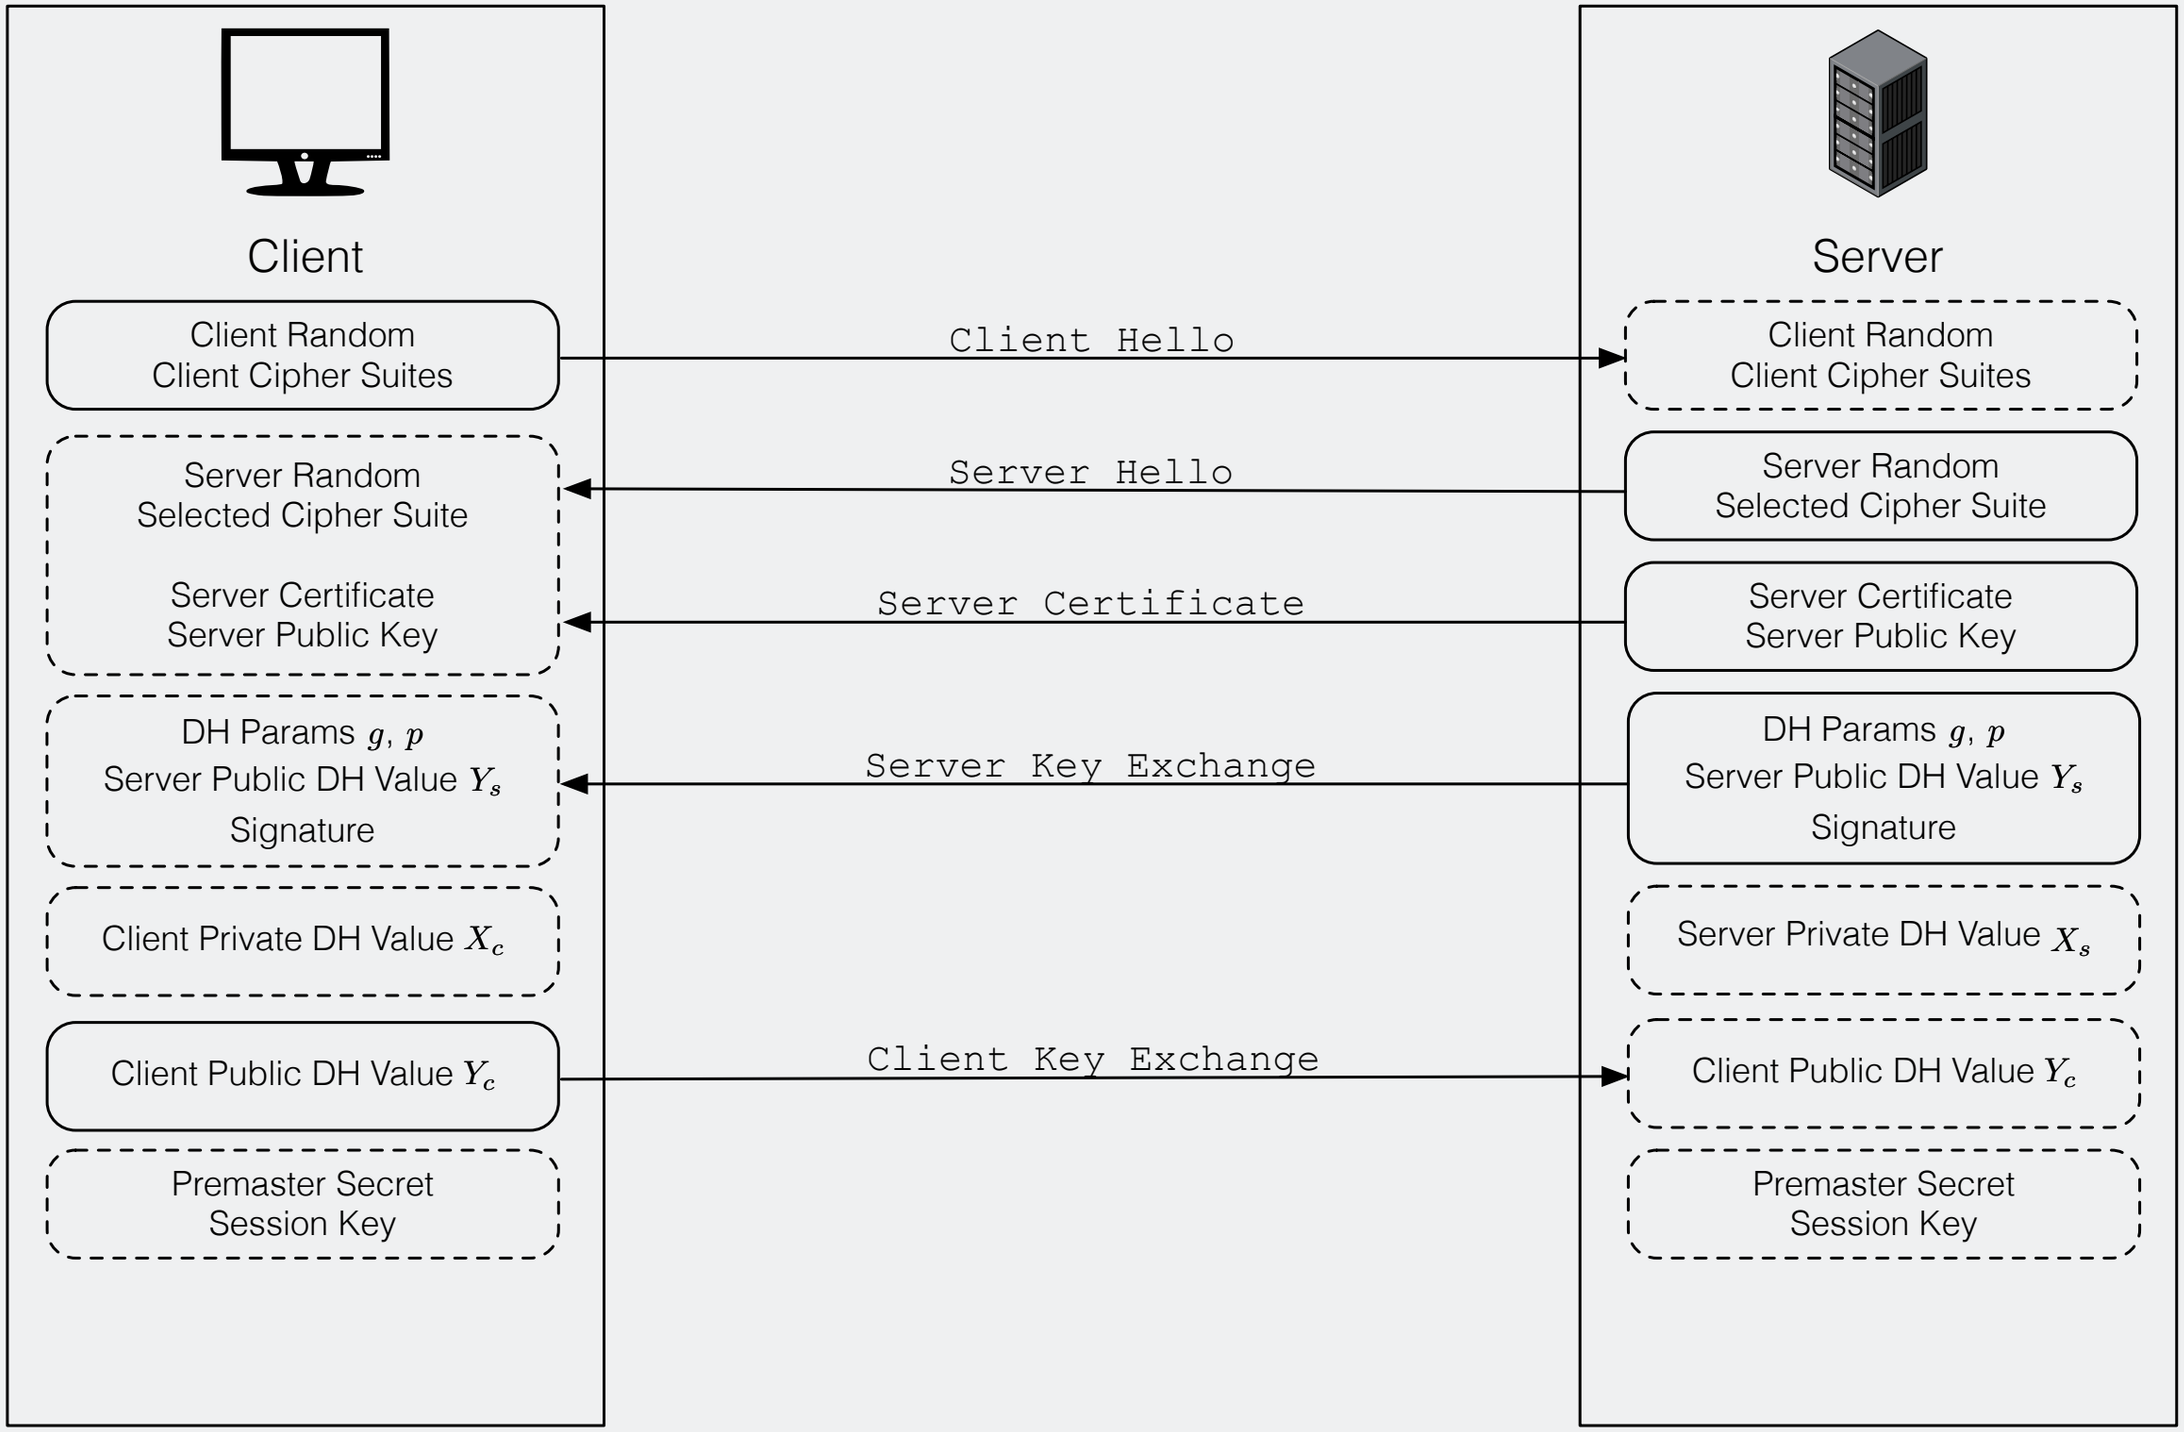
Client
Client Random
Client Cipher Suites
Server Random
Selected Cipher Suite
Server Certificate
Server Public Key
DH Params g, p
Server Public DH Value Y_s
Signature
Client Private DH Value X_c
Client Public DH Value Y_c
Premaster Secret
Session Key
Client Hello
Server Hello
Server Certificate
Server Key Exchange
Client Key Exchange
Server
Client Random
Client Cipher Suites
Server Random
Selected Cipher Suite
Server Certificate
Server Public Key
DH Params g, p
Server Public DH Value Y_s
Signature
Server Private DH Value X_s
Client Public DH Value Y_c
Premaster Secret
Session Key

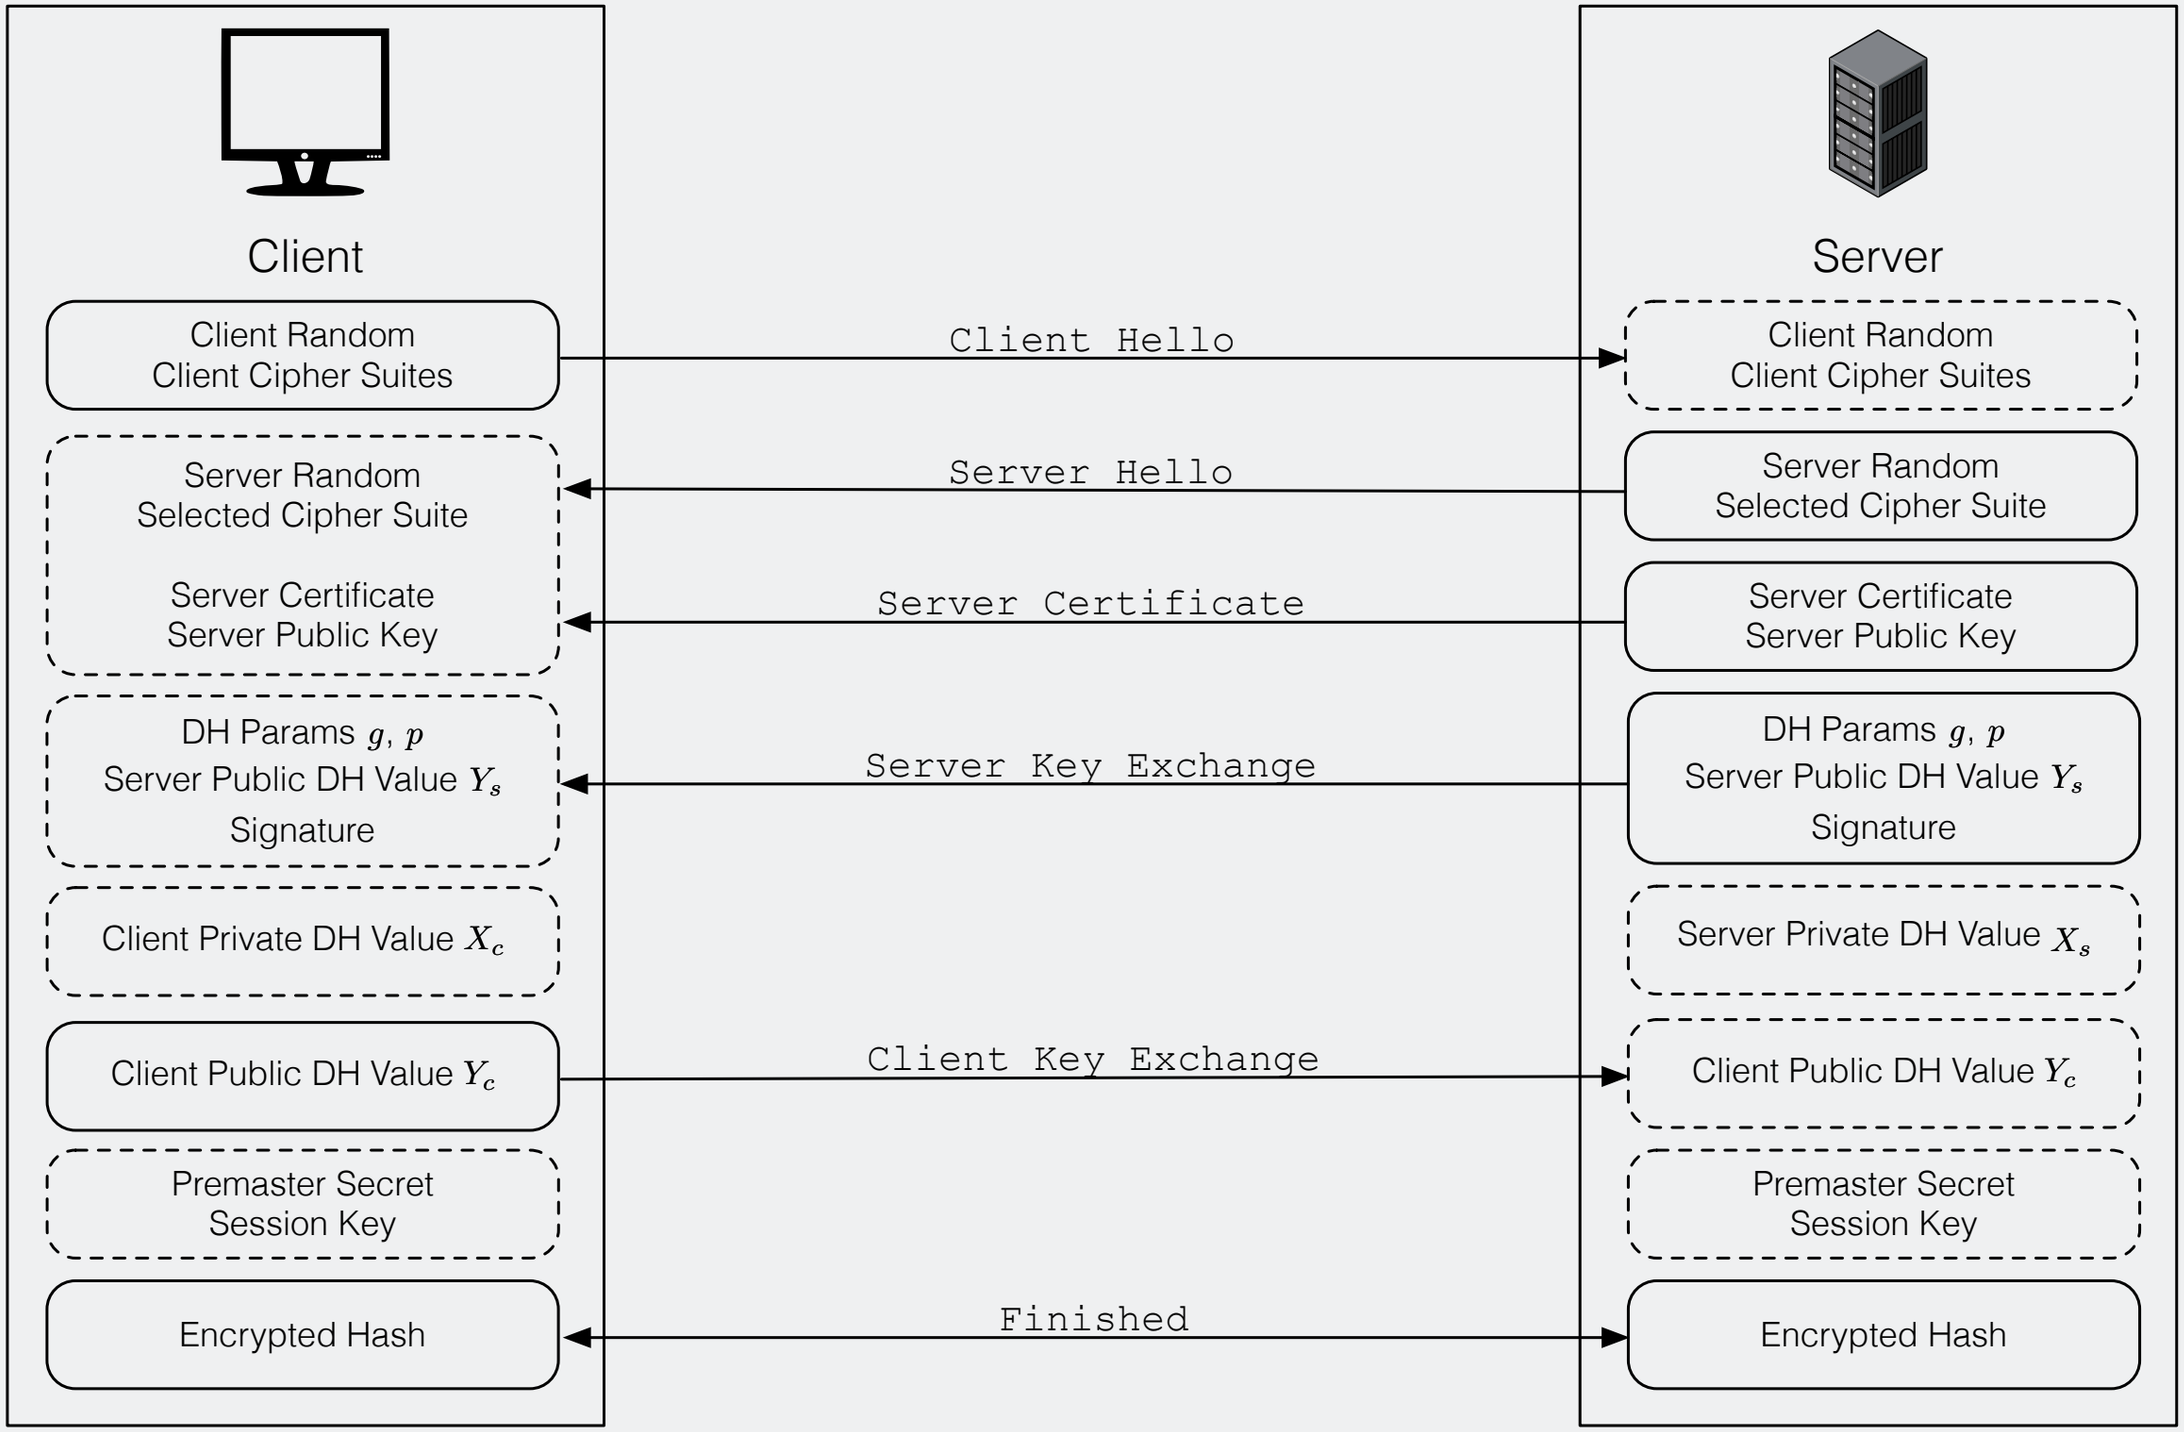
Client
Server
Client Random
Client Cipher Suites
Client Hello
Client Random
Client Cipher Suites
Server Random
Selected Cipher Suite
Server Hello
Server Random
Selected Cipher Suite
Server Certificate
Server Public Key
Server Certificate
Server Certificate
Server Public Key
DH Params g, p
Server Public DH Value Y_s
Signature
Server Key Exchange
DH Params g, p
Server Public DH Value Y_s
Signature
Client Private DH Value X_c
Server Private DH Value X_s
Client Public DH Value Y_c
Client Key Exchange
Client Public DH Value Y_c
Premaster Secret
Session Key
Premaster Secret
Session Key
Encrypted Hash
Finished
Encrypted Hash

“Surely, no one uses 512-bit primes?”

Export Ciphers

Remnant of the 90s “crypto wars”

It used to be illegal to export “strong crypto” outside of the United States

9th Circuit Court overturned the law 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

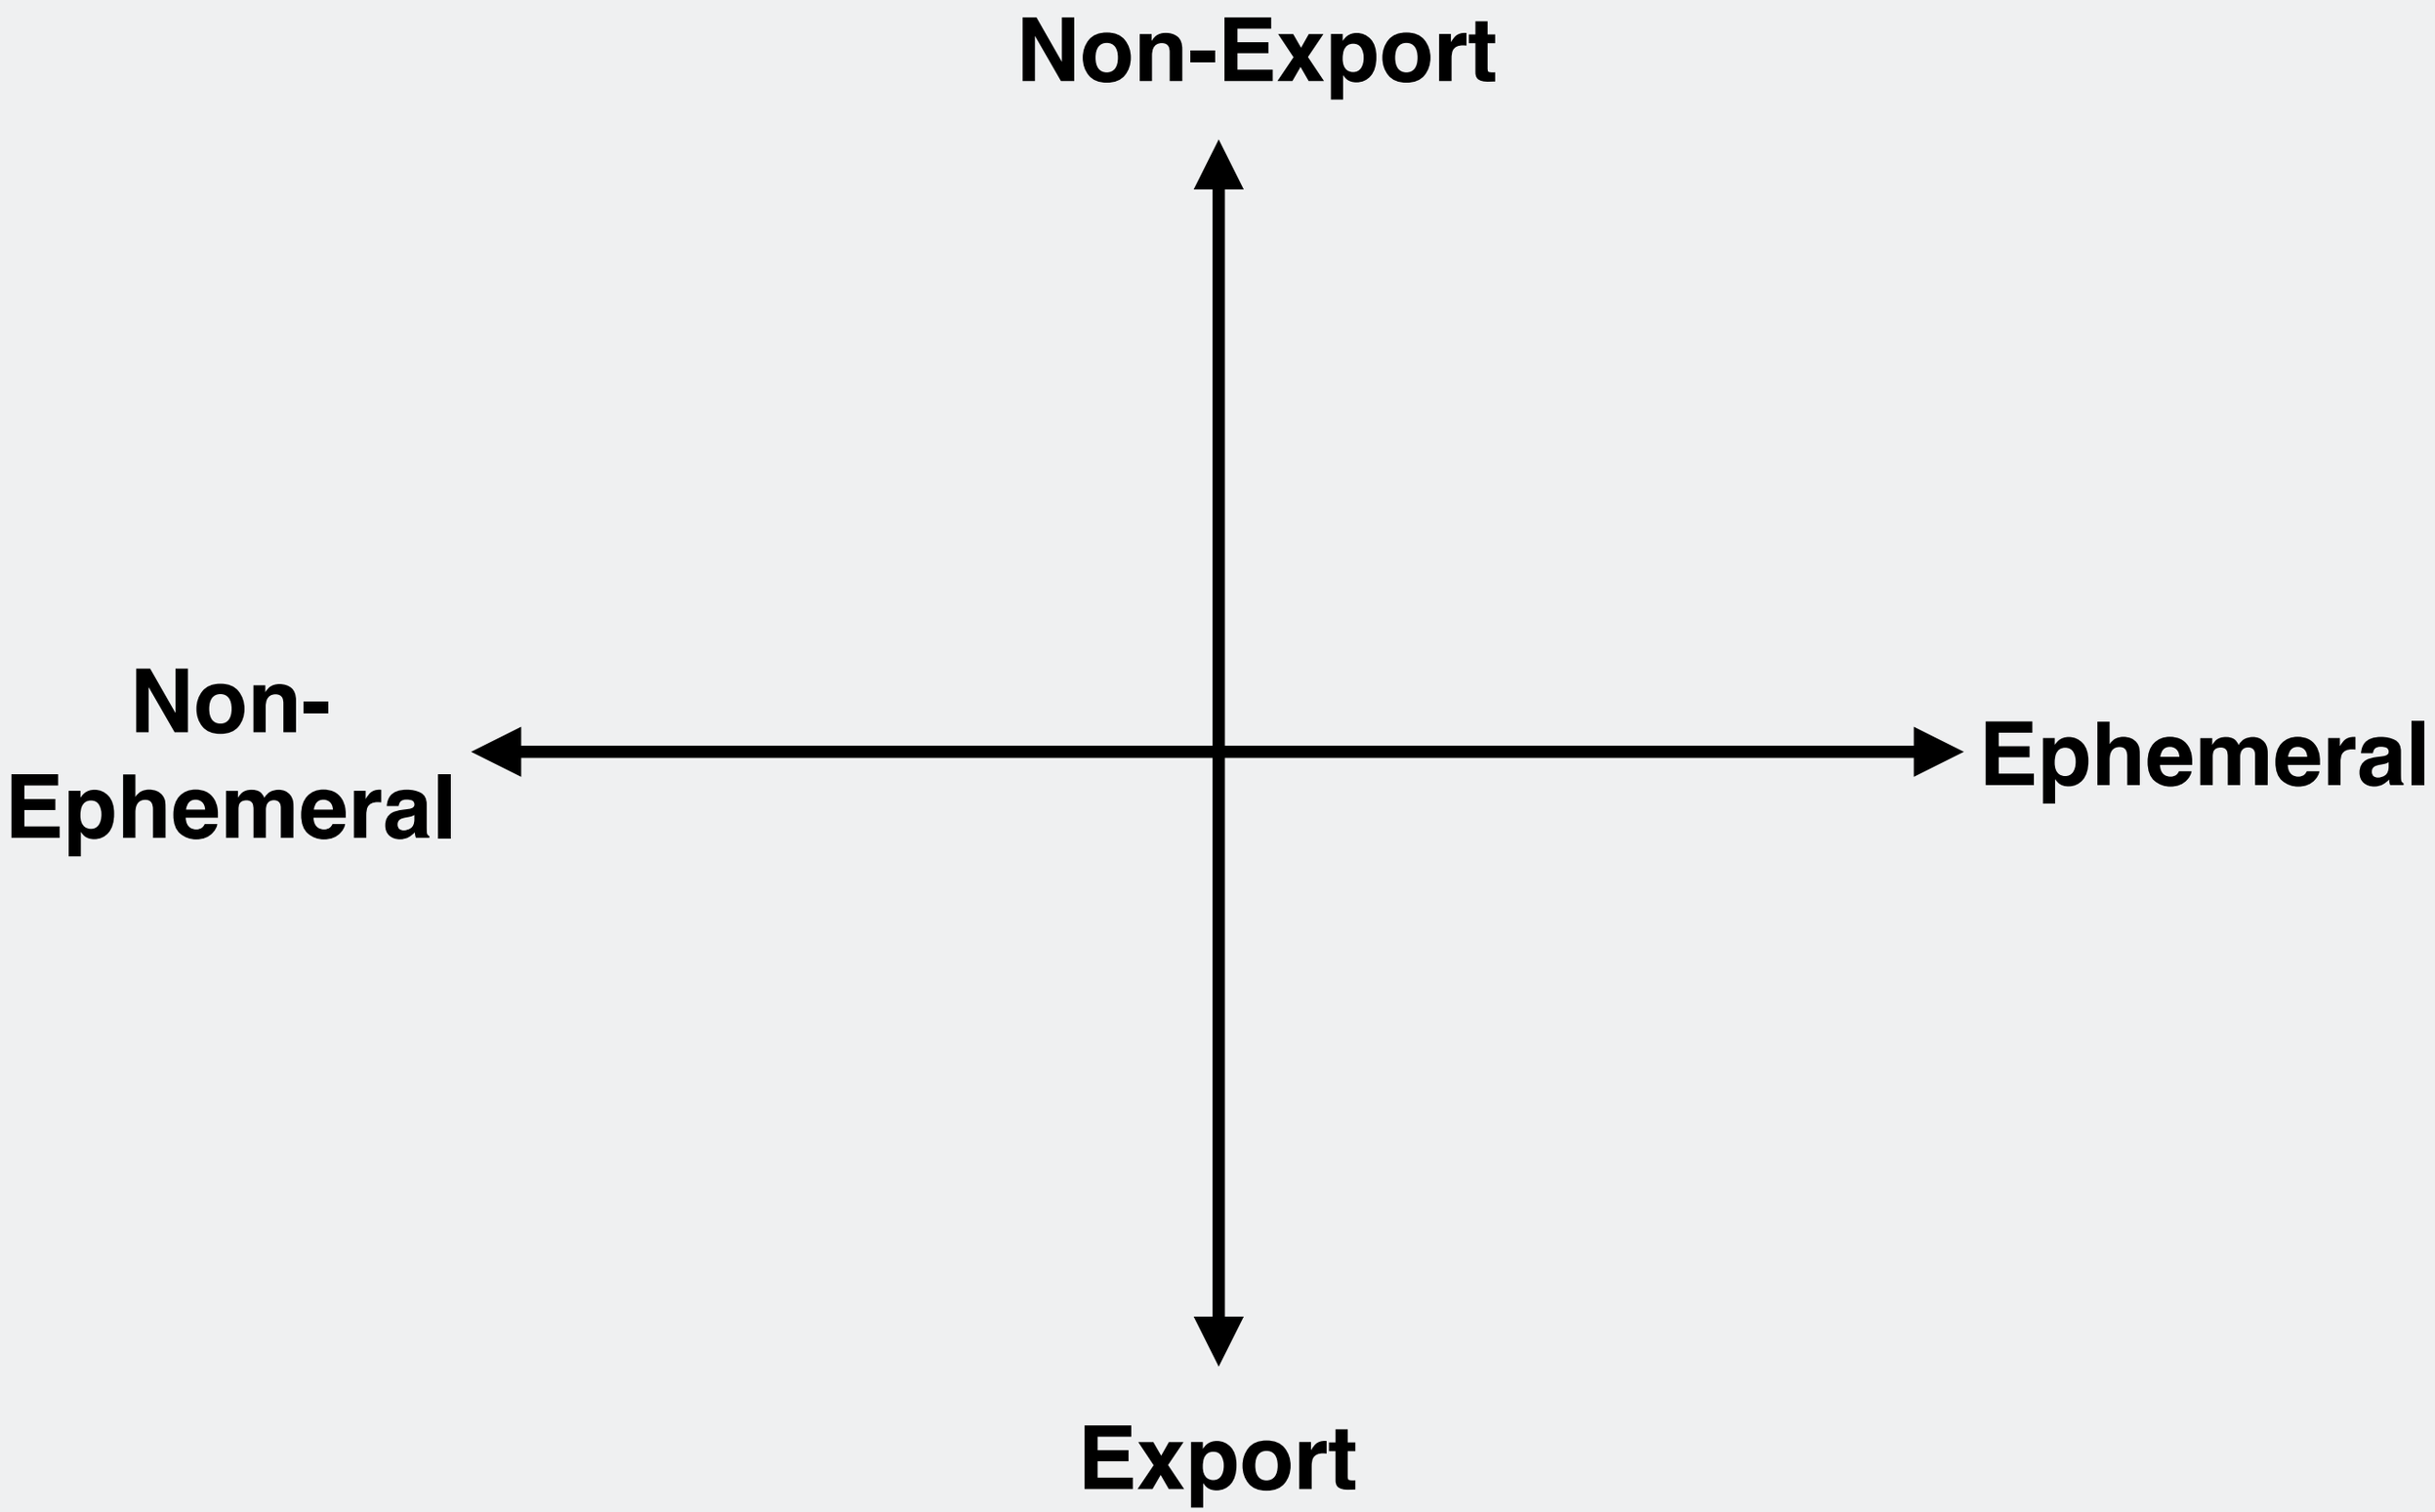
Non-Export
Non-Ephemeral
Ephemeral
Ephemeral
Export

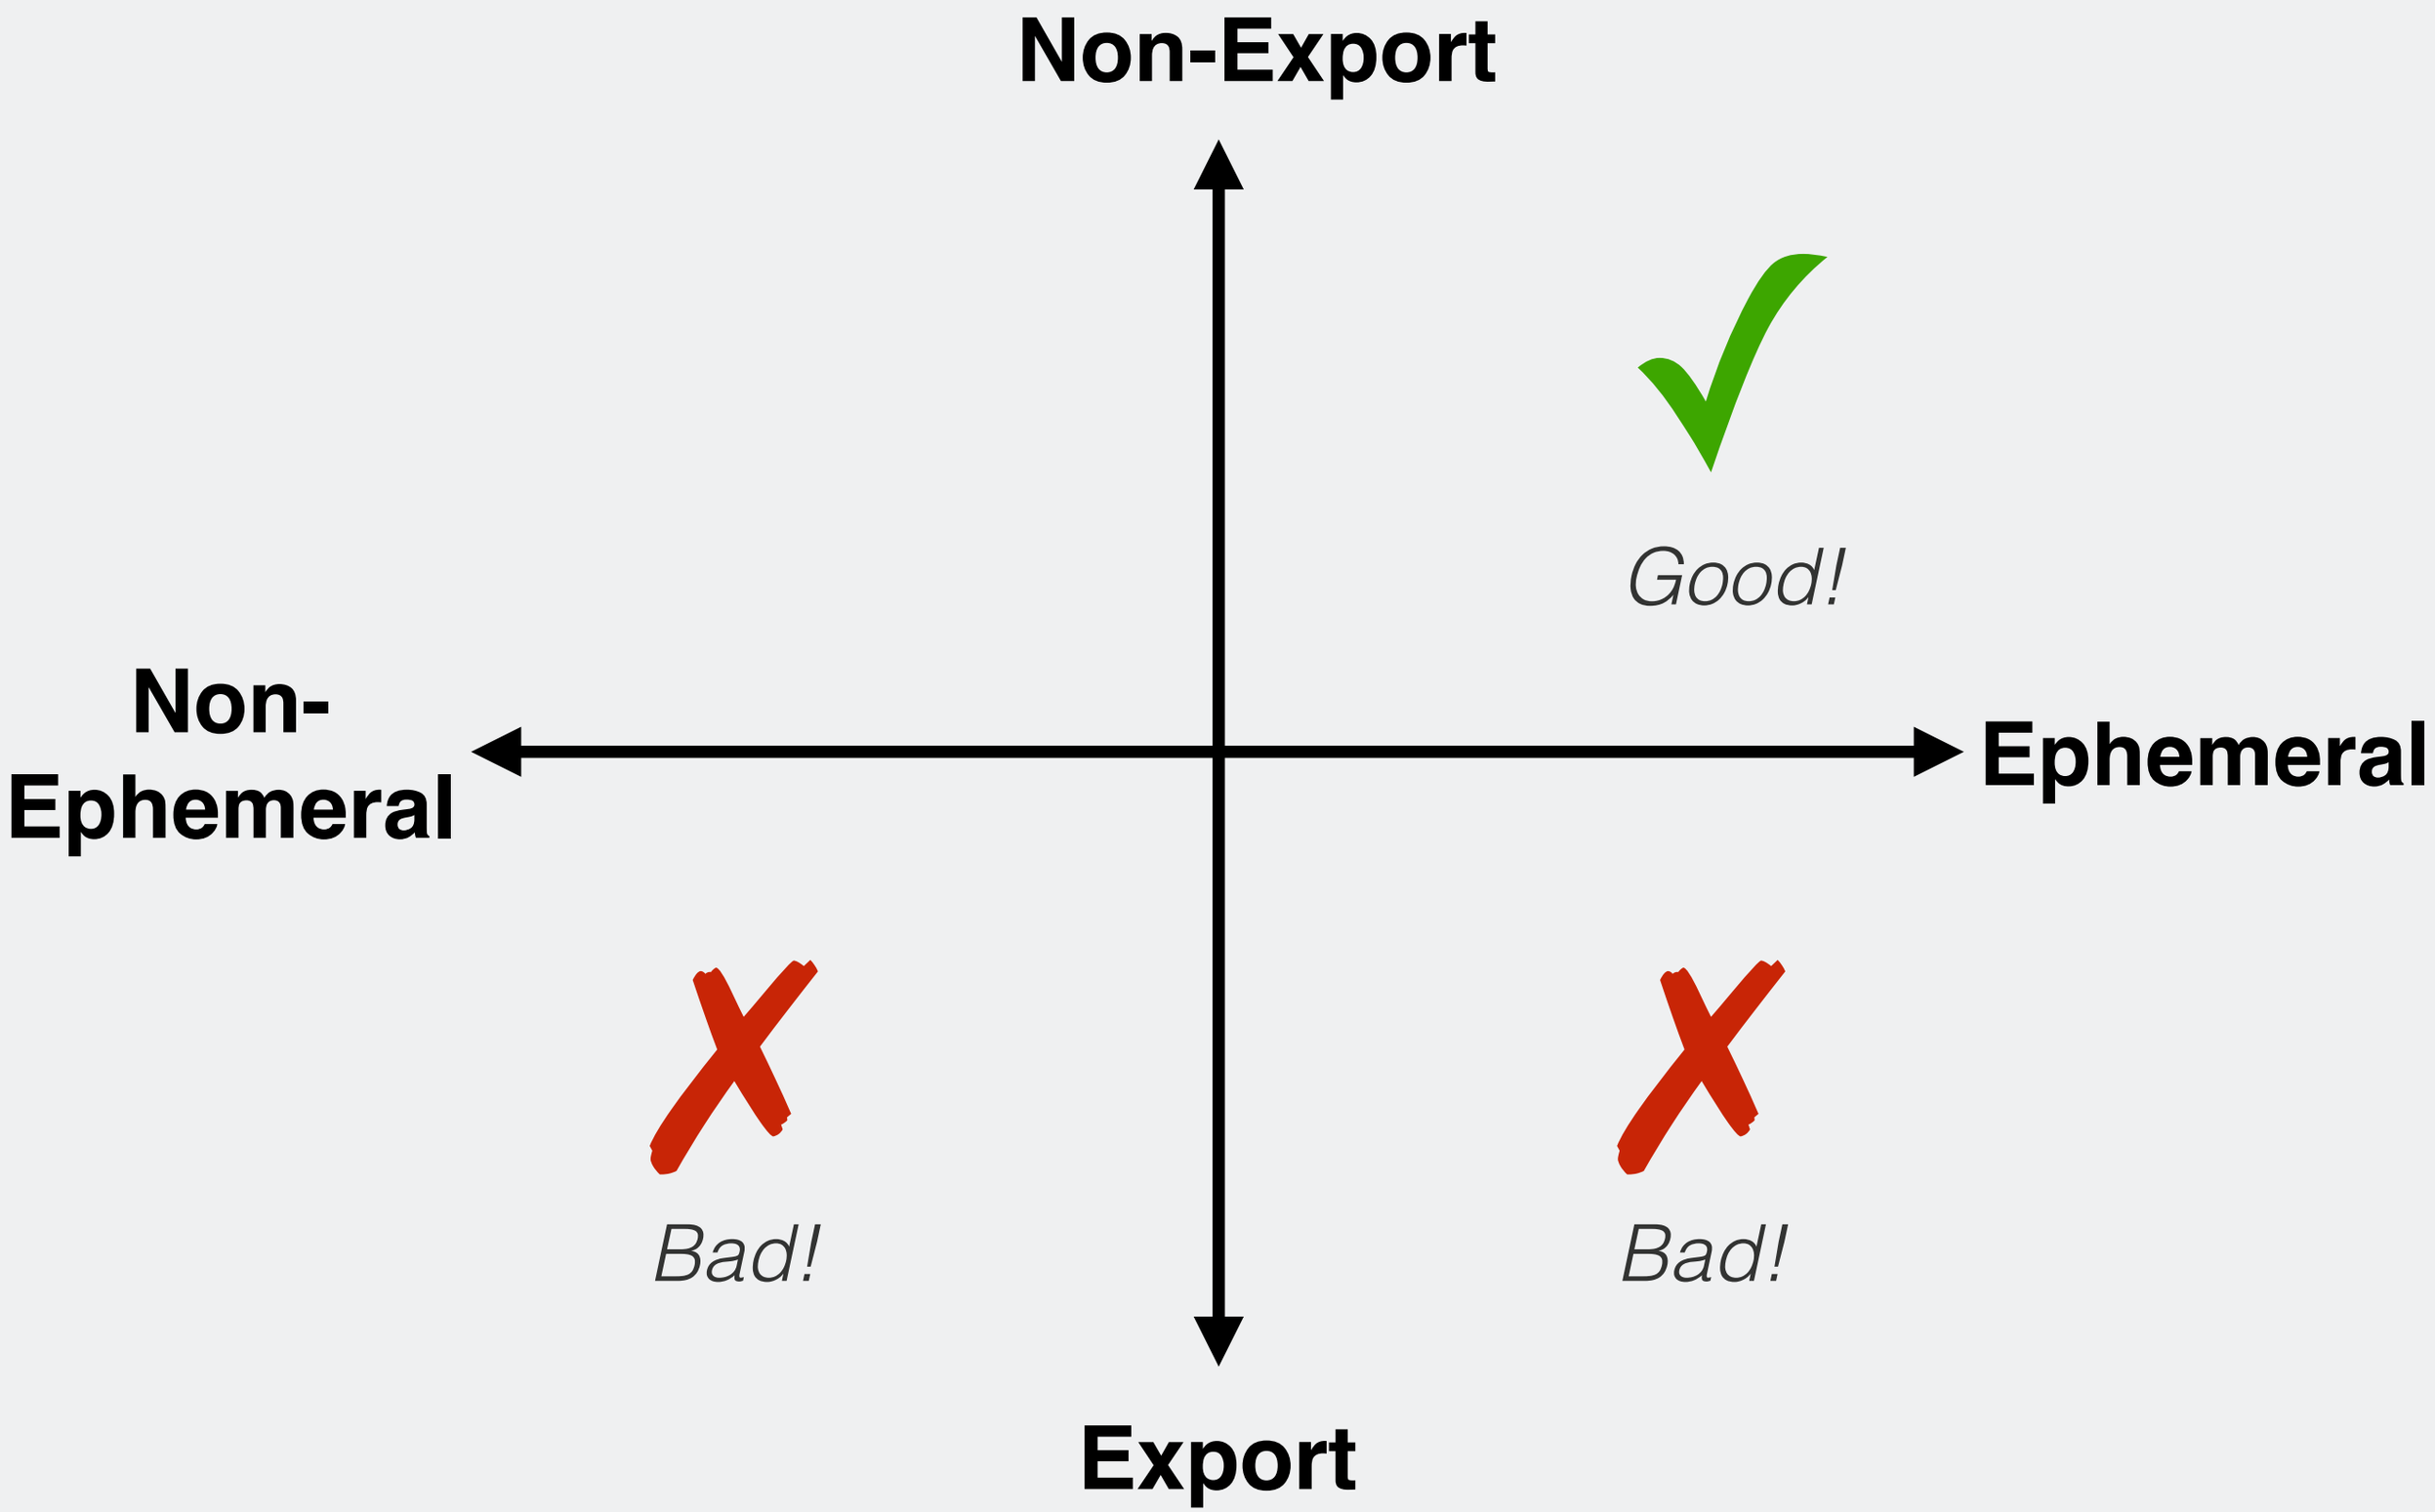
Non-Export
Non-Ephemeral
Ephemeral
Ephemeral
Good!
Bad!
Bad!
Export

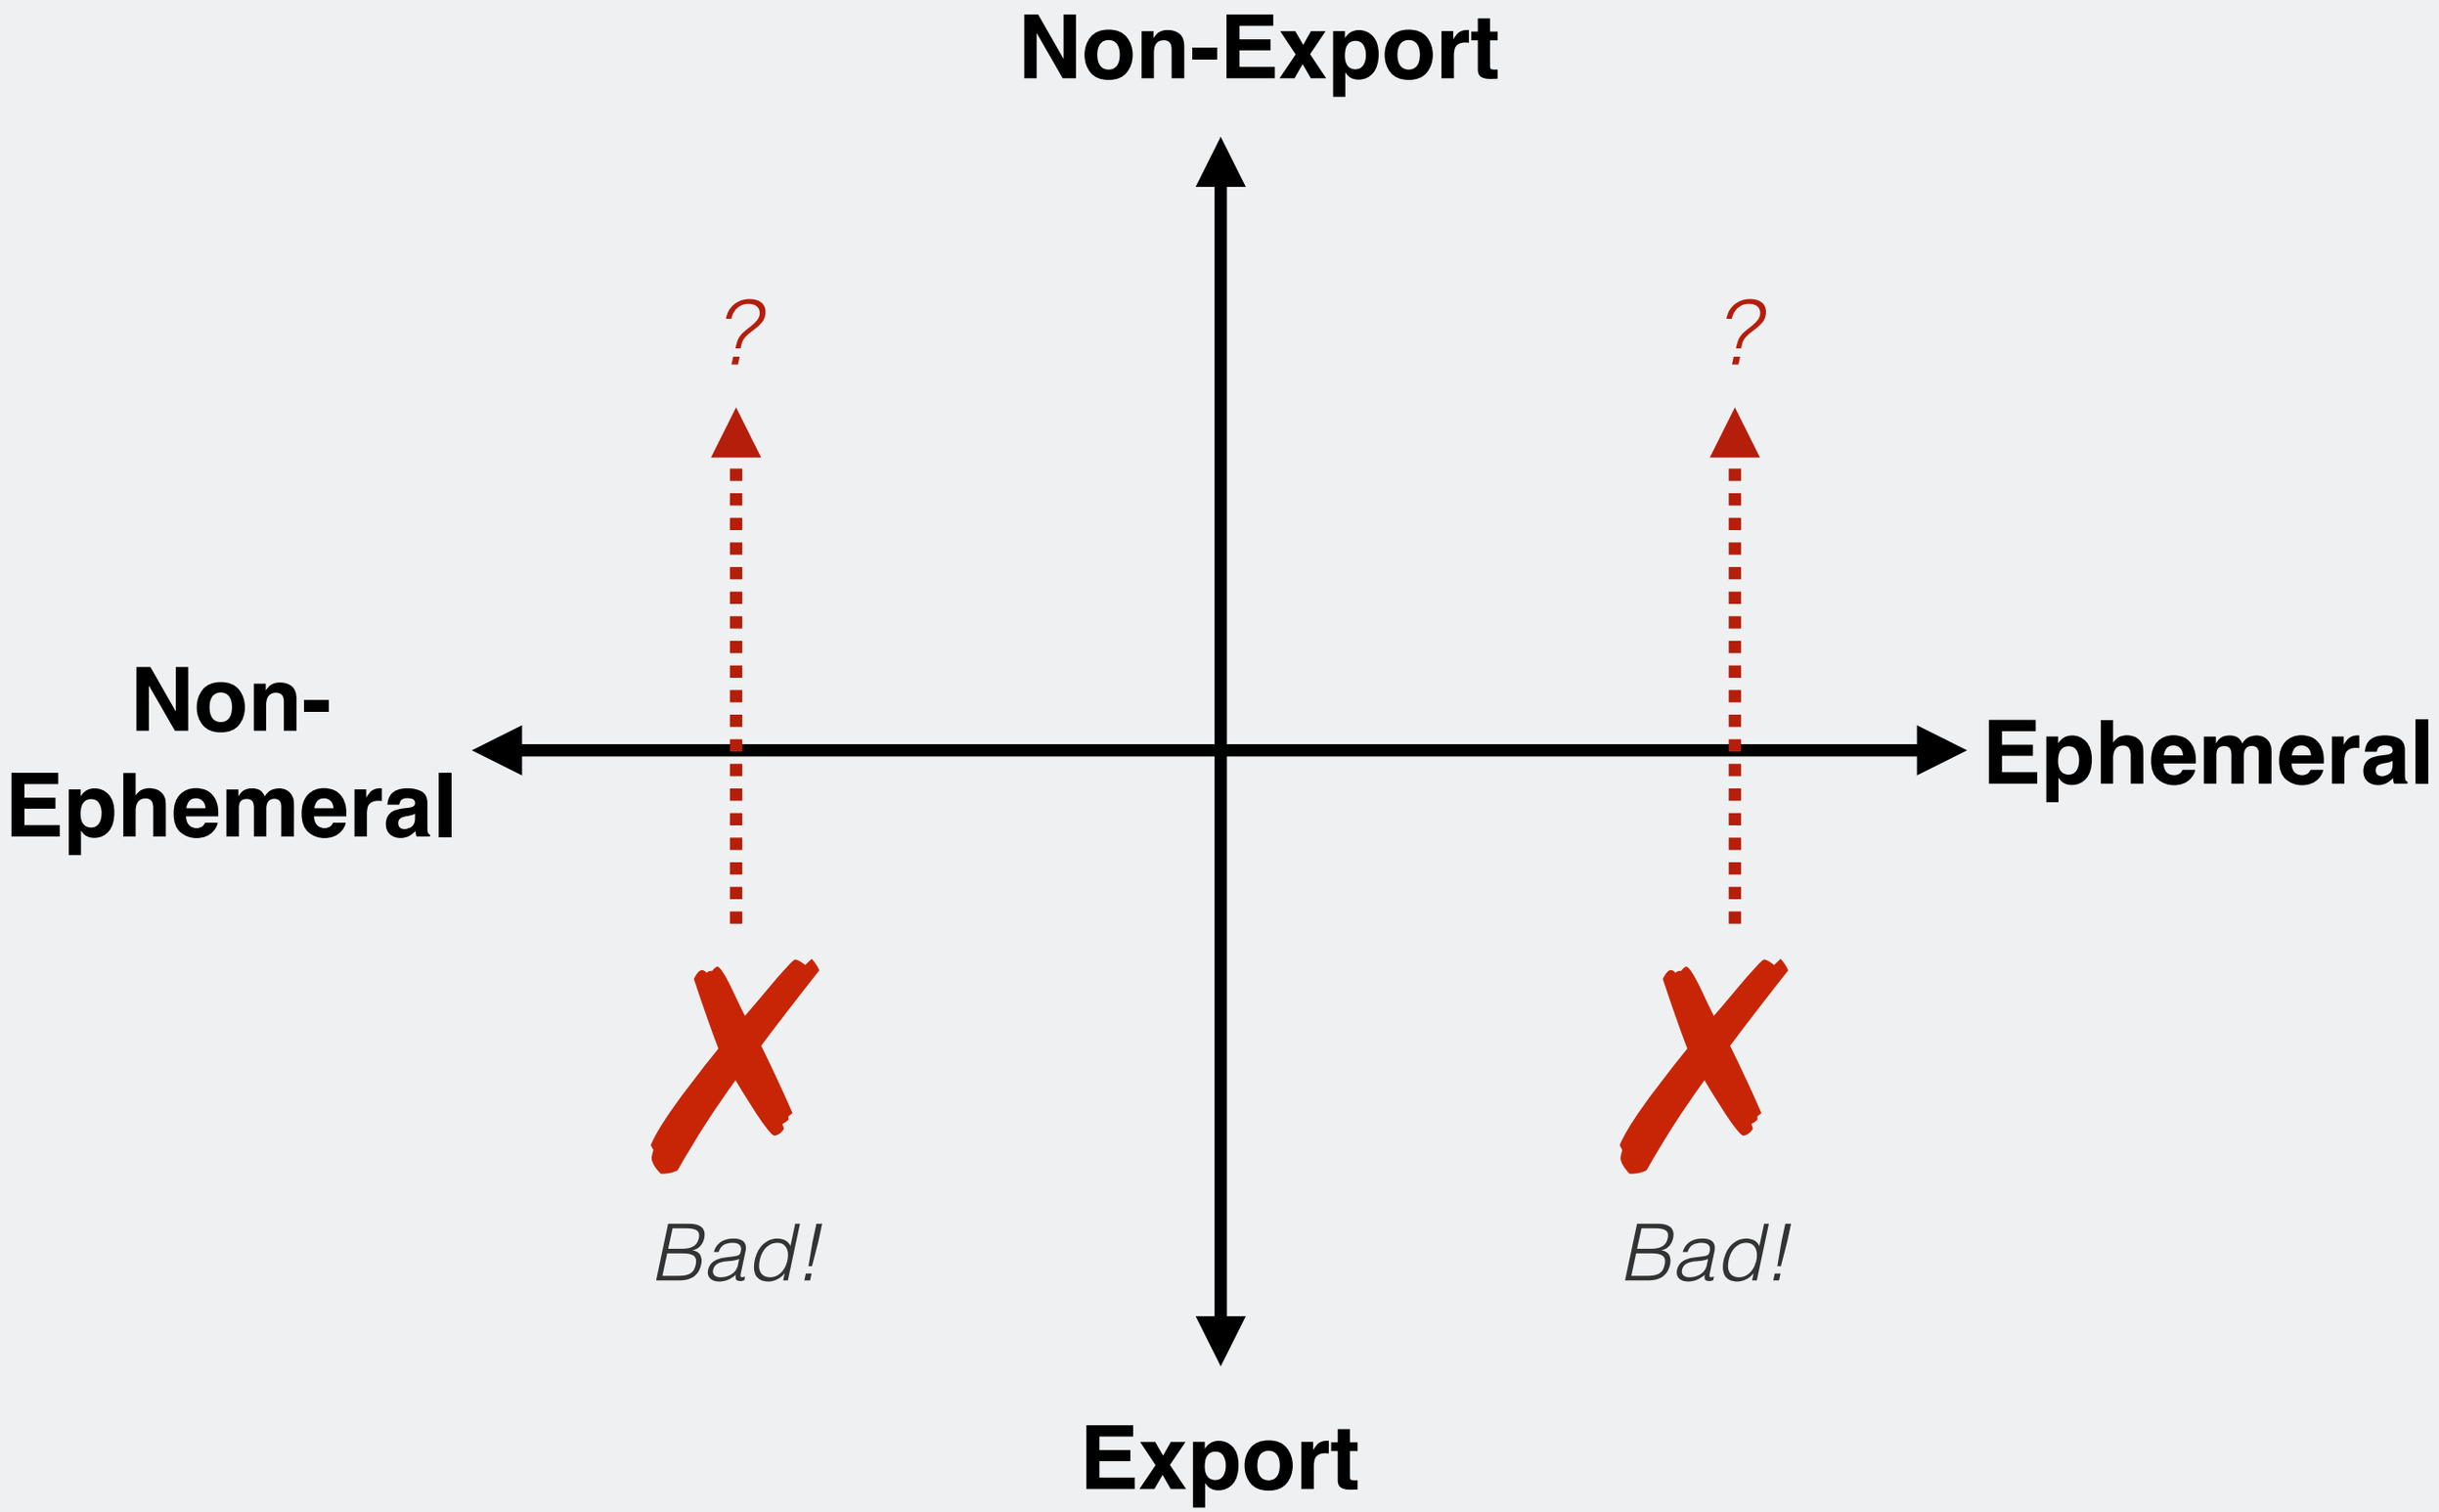
Non-Export
?
?
Non-
Ephemeral
Ephemeral
Bad!
Bad!
Export

Logjam 🌲🍯

Logjam is a **downgrade attack against TLS** that enables a man-in-the-middle to read and modify data passed over the connection.

Logjam affects any server that supports **DHE_EXPORT** ciphers.

Client

MitM

Server

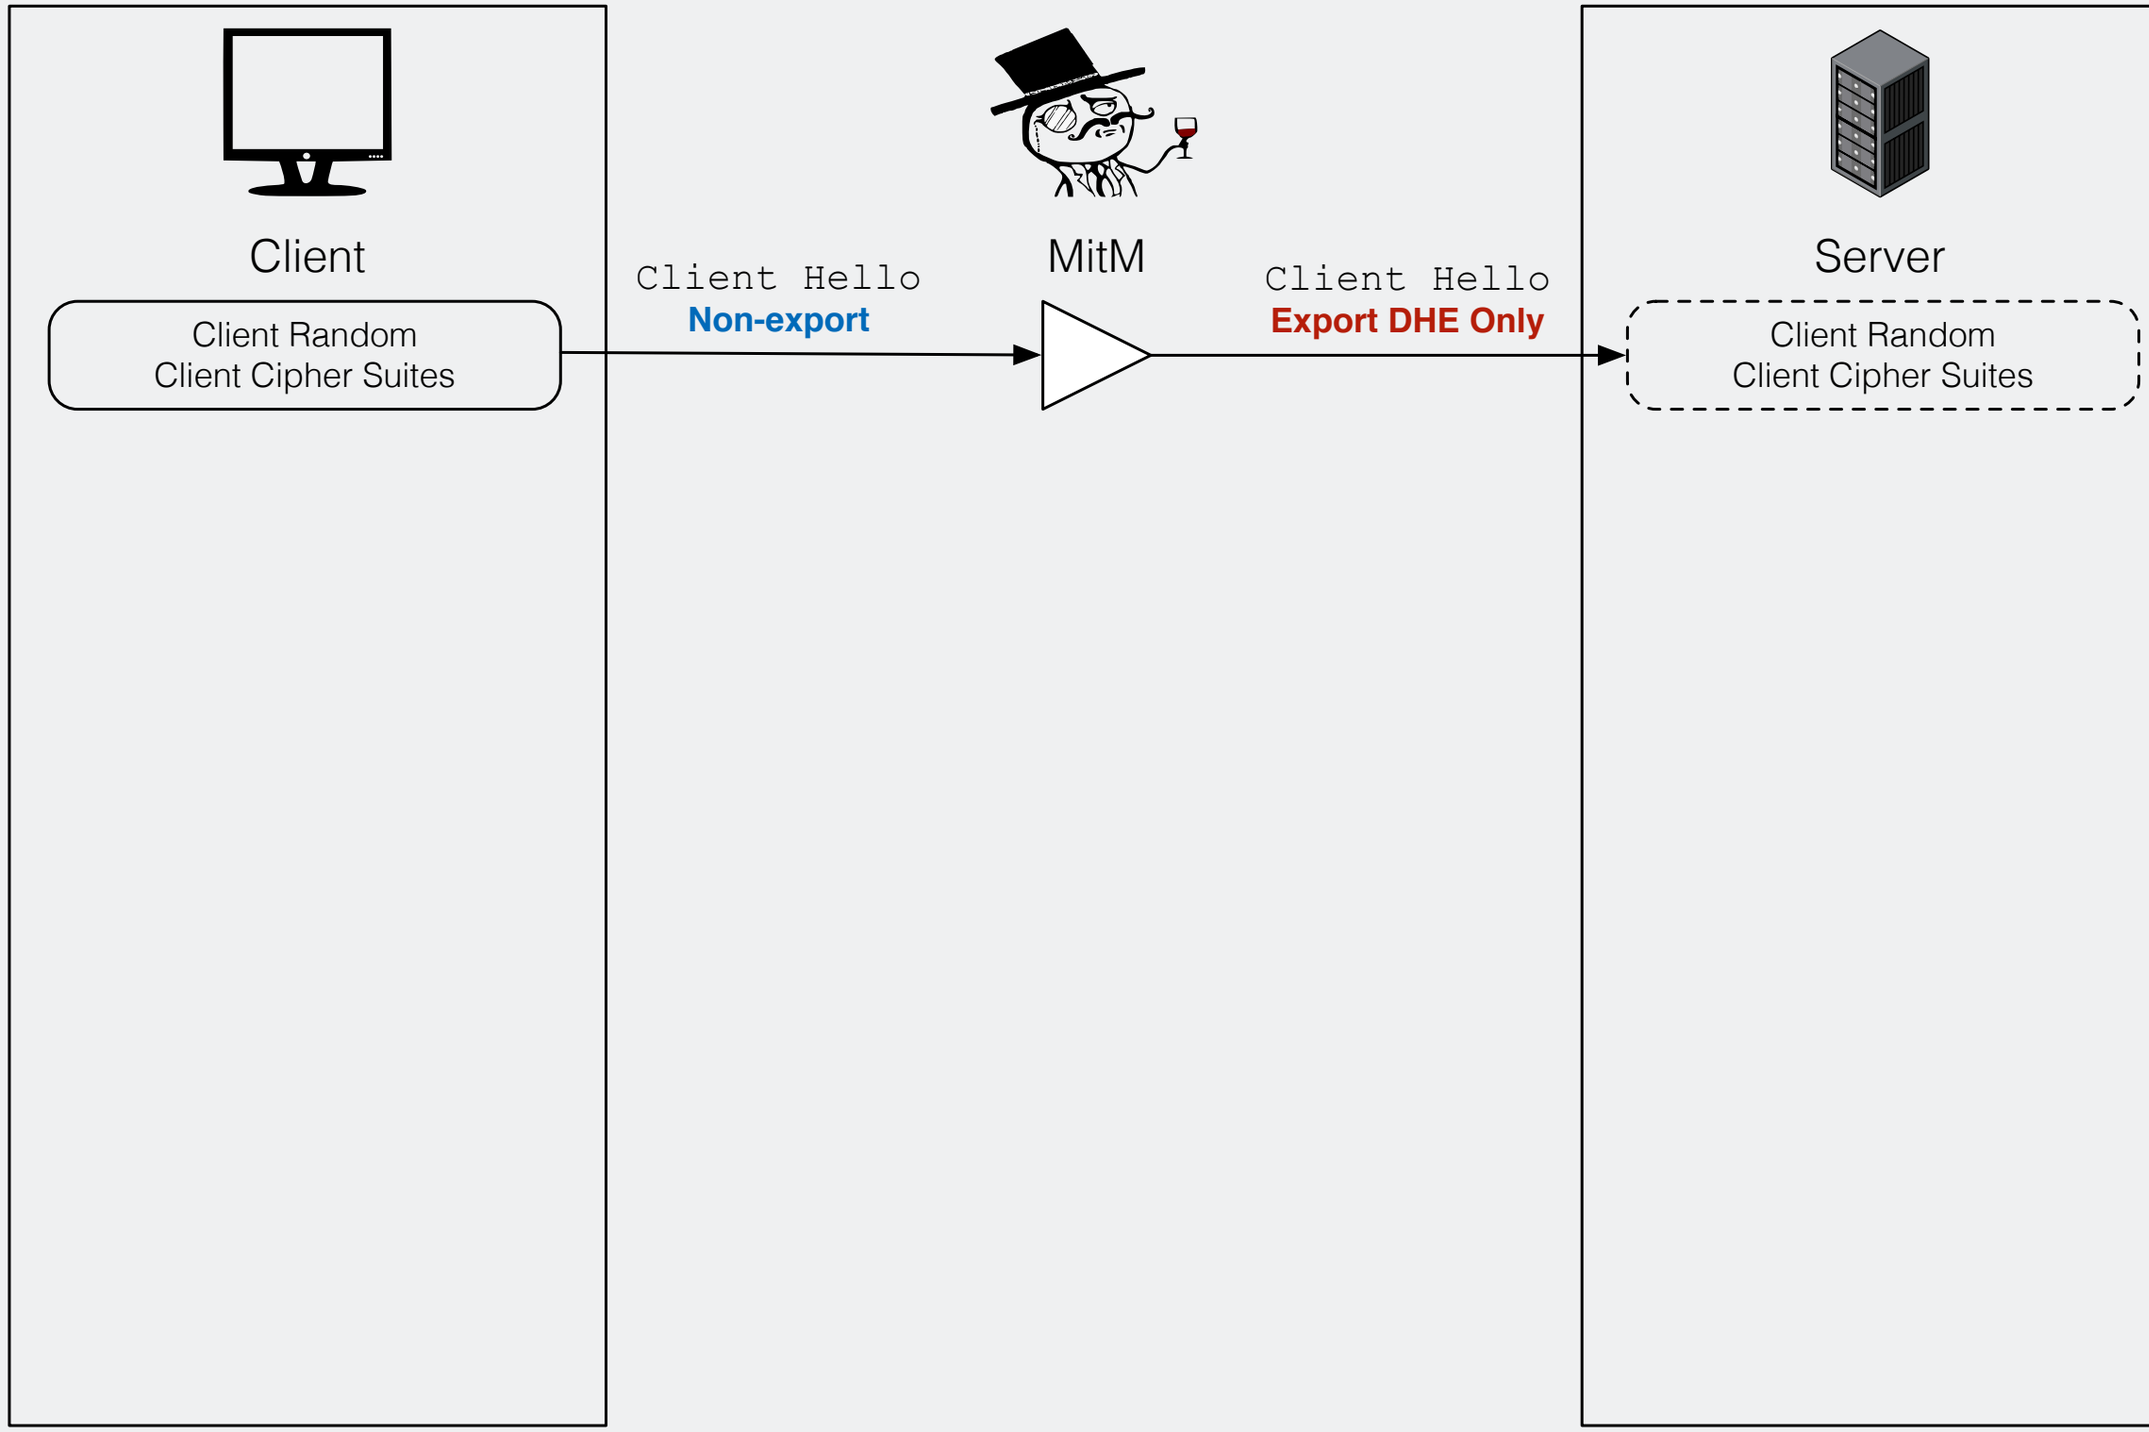
Client
Client Random
Client Cipher Suites
Client Hello
Non-export
MitM
Client Hello
Export DHE Only
Server
Client Random
Client Cipher Suites

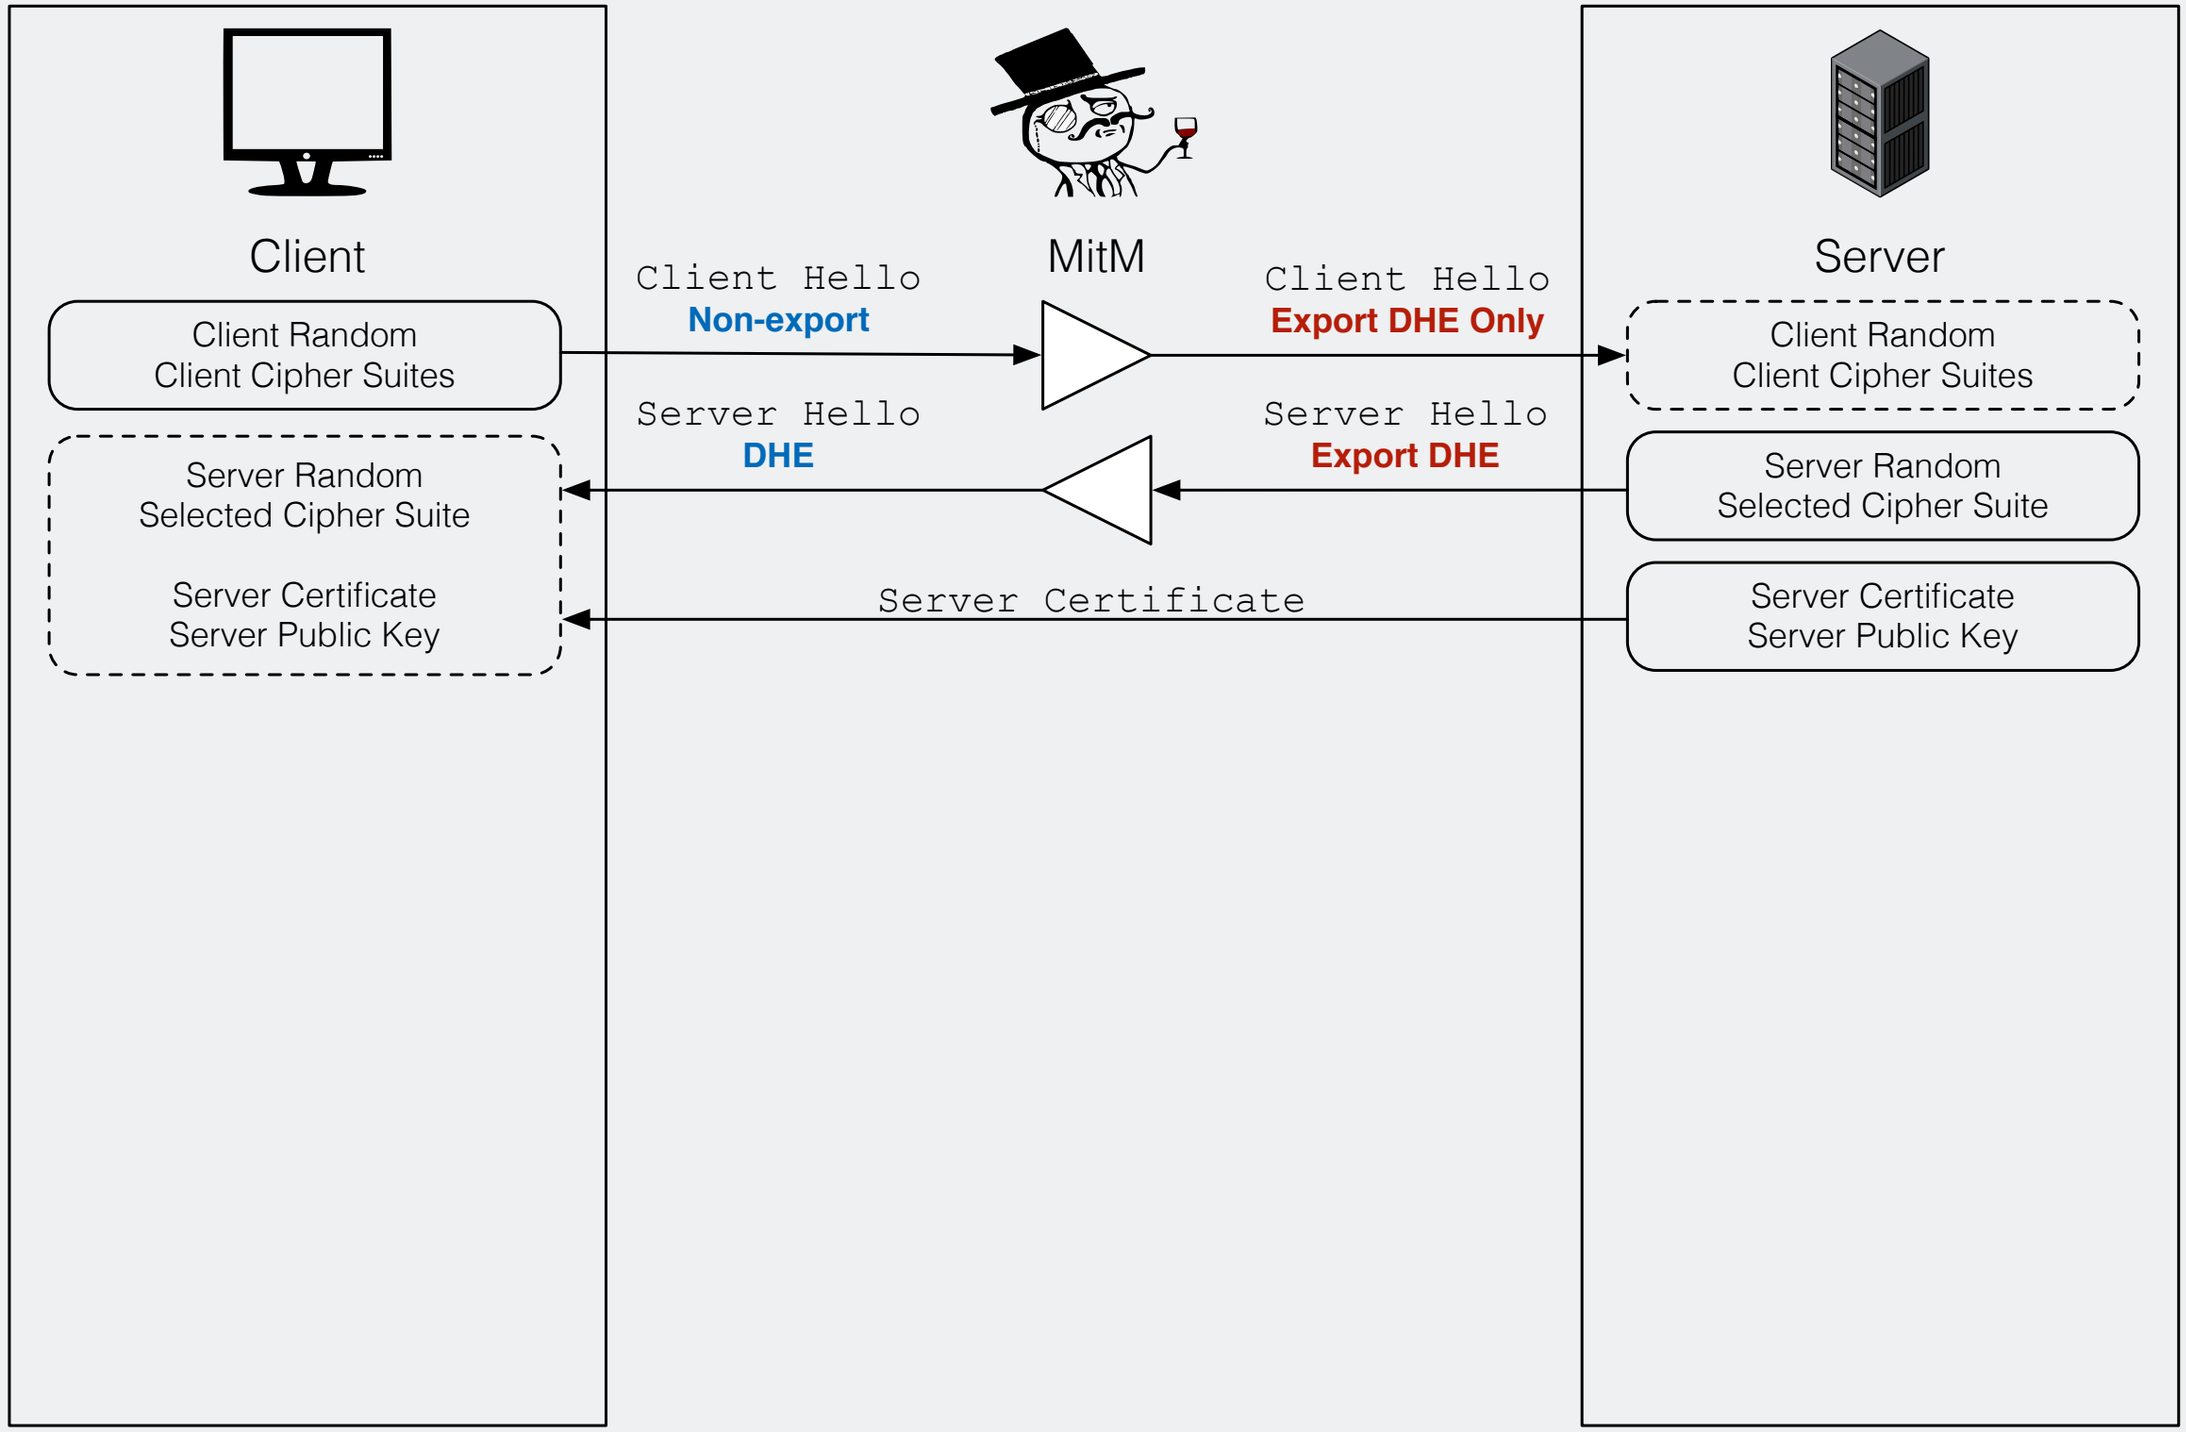
Client
MitM
Server
Client Hello
Non-export
Client Hello
Export DHE Only
Client Random
Client Cipher Suites
Client Random
Client Cipher Suites
Server Hello
DHE
Server Hello
Export DHE
Server Random
Selected Cipher Suite
Server Random
Selected Cipher Suite
Server Certificate
Server Public Key
Server Certificate
Server Certificate
Server Public Key

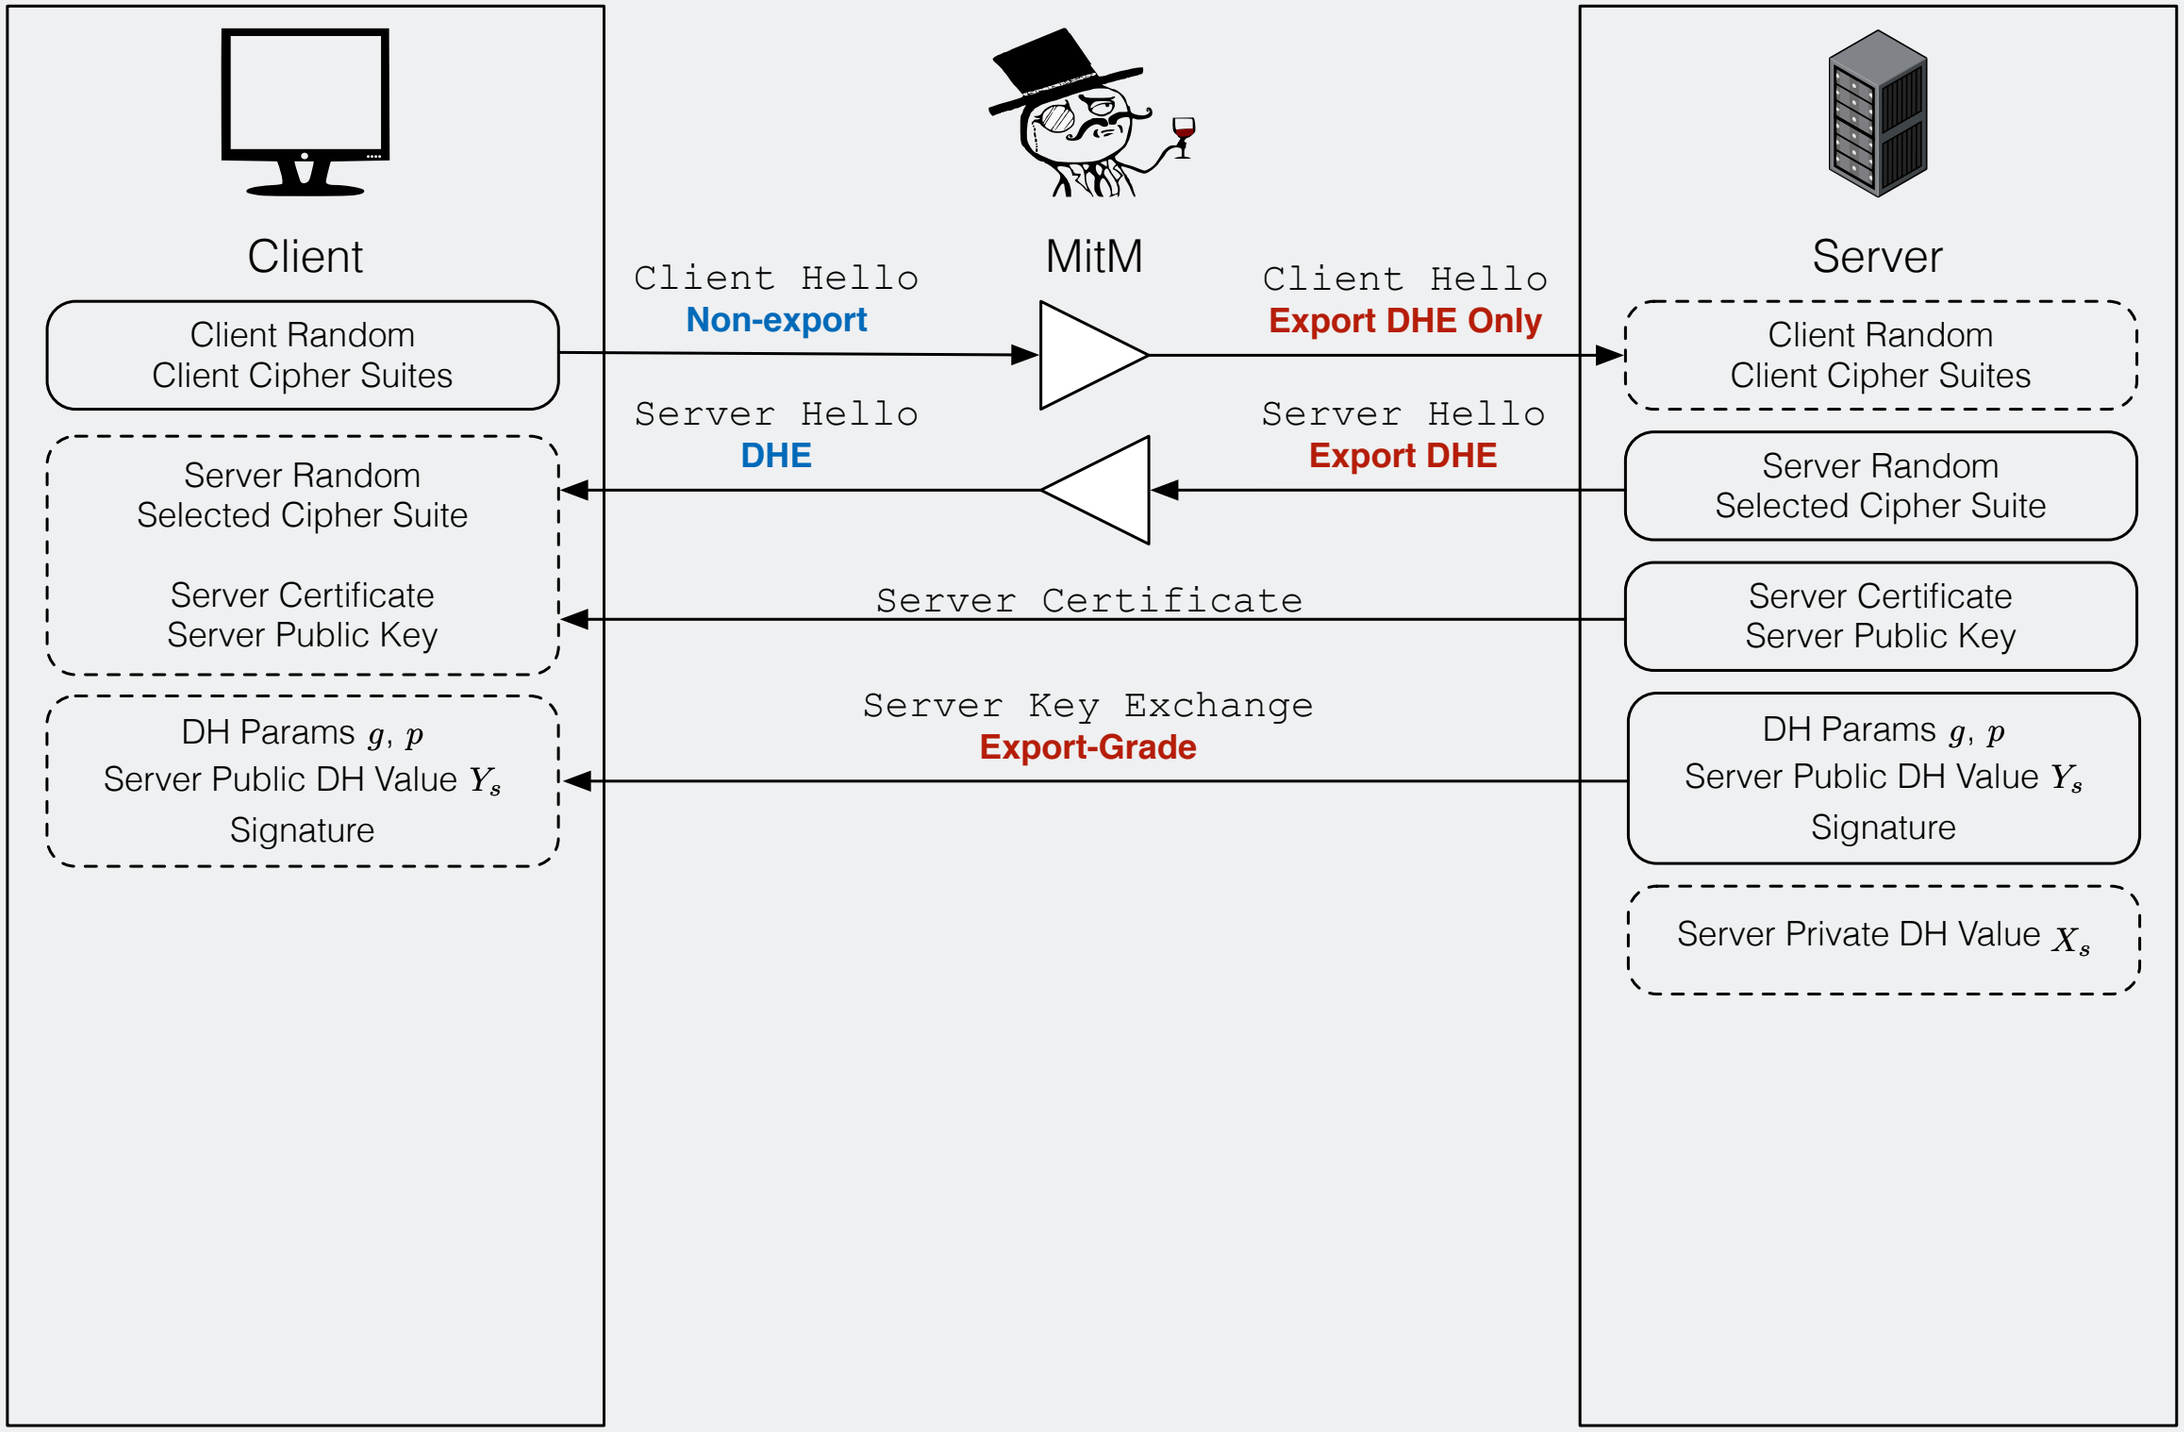
Client
MitM
Server
Client Random
Client Cipher Suites
Client Hello
Non-export
Client Hello
Export DHE Only
Client Random
Client Cipher Suites
Server Hello
DHE
Server Hello
Export DHE
Server Random
Selected Cipher Suite
Server Random
Selected Cipher Suite
Server Certificate
Server Public Key
Server Certificate
Server Certificate
Server Public Key
DH Params g, p
Server Public DH Value Y_s
Signature
Server Key Exchange
Export-Grade
DH Params g, p
Server Public DH Value Y_s
Signature
Server Private DH Value X_s

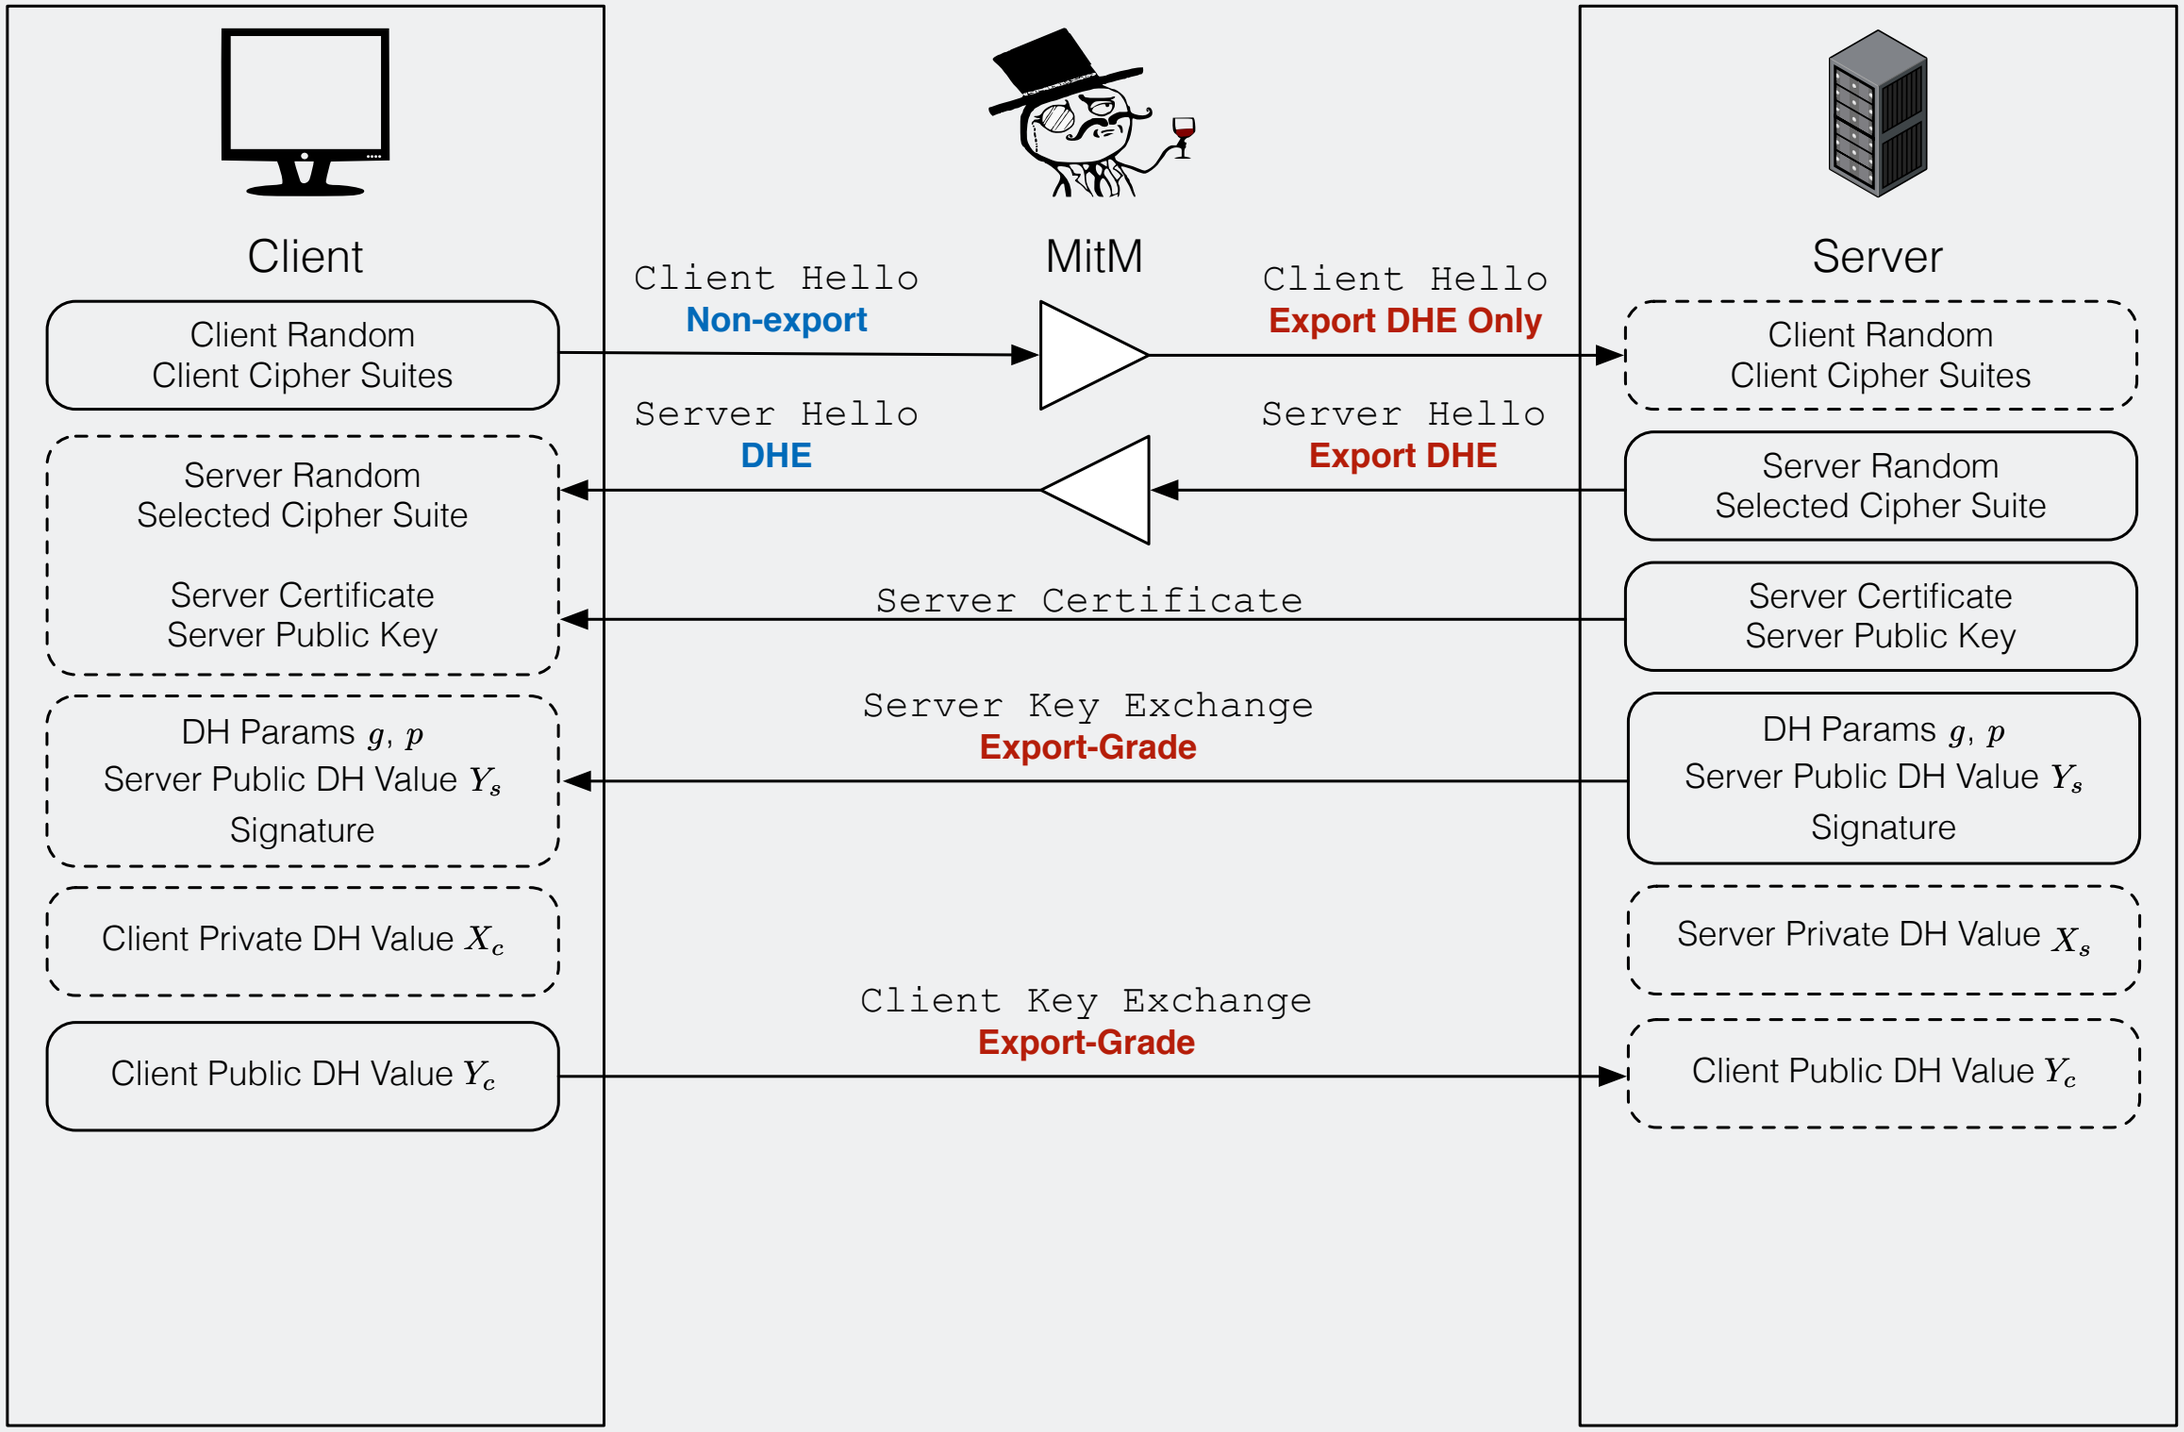
Client
MitM
Server
Client Hello
Non-export
Client Hello
Export DHE Only
Client Random
Client Cipher Suites
Client Random
Client Cipher Suites
Server Hello
DHE
Server Hello
Export DHE
Server Random
Selected Cipher Suite
Server Random
Selected Cipher Suite
Server Certificate
Server Public Key
Server Certificate
Server Certificate
Server Public Key
Server Key Exchange
Export-Grade
DH Params g, p
Server Public DH Value Y_s
Signature
DH Params g, p
Server Public DH Value Y_s
Signature
Client Private DH Value X_c
Server Private DH Value X_s
Client Key Exchange
Export-Grade
Client Public DH Value Y_c
Client Public DH Value Y_c

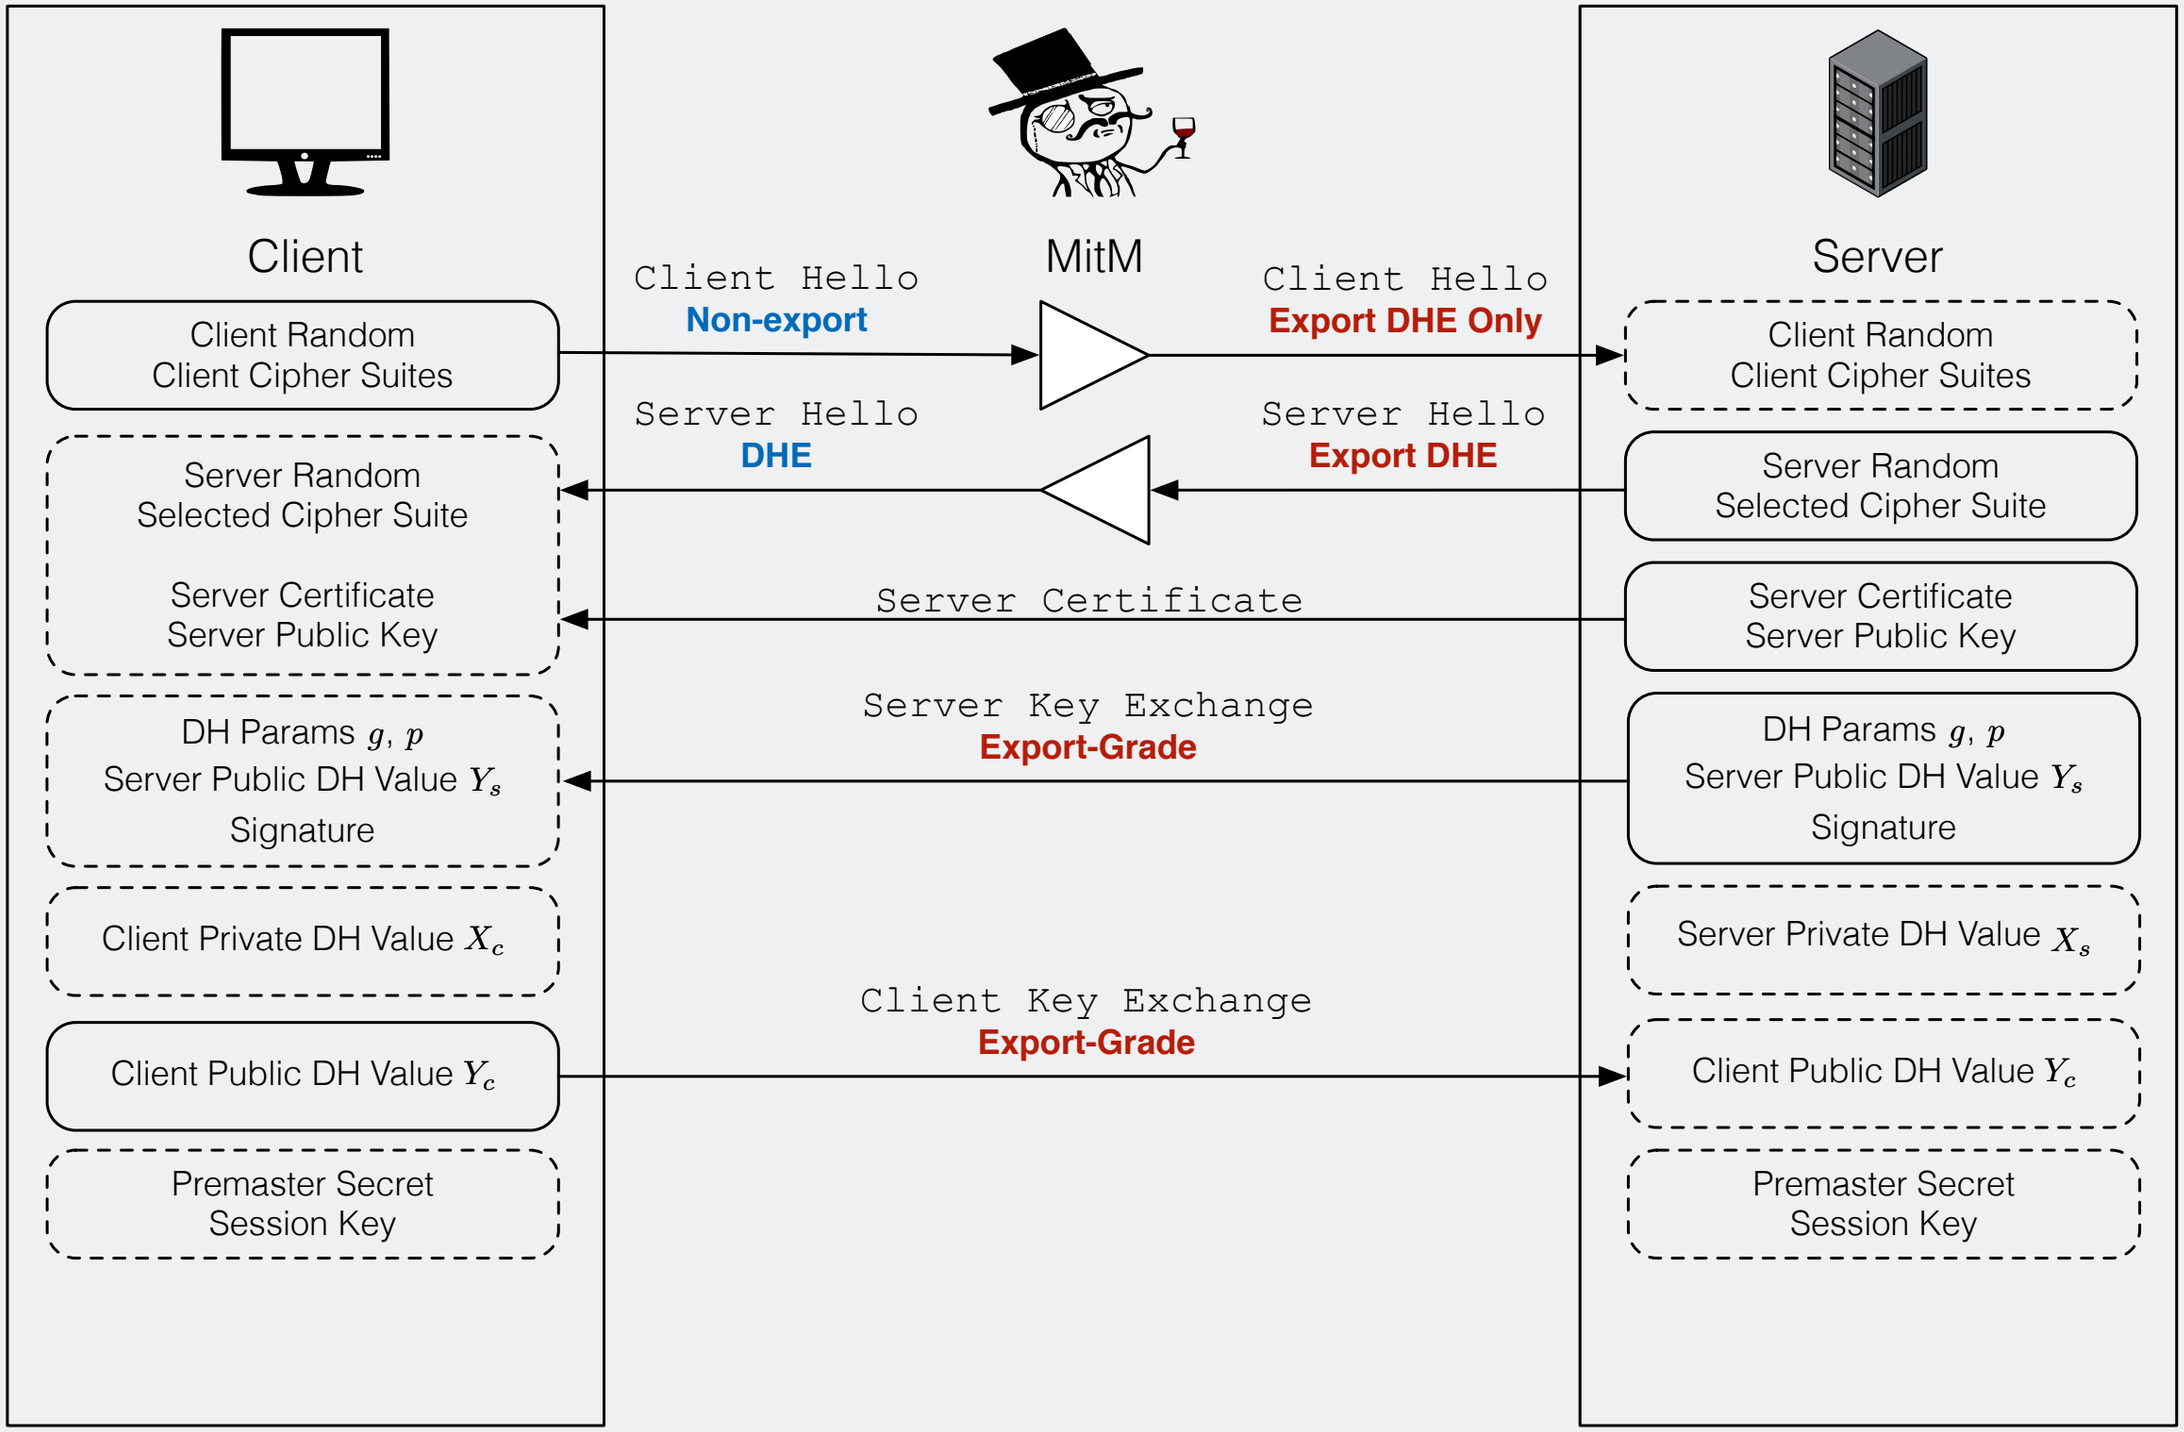
Client
MitM
Server
Client Hello
Non-export
Client Hello
Export DHE Only
Client Random
Client Cipher Suites
Client Random
Client Cipher Suites
Server Hello
DHE
Server Hello
Export DHE
Server Random
Selected Cipher Suite
Server Random
Selected Cipher Suite
Server Certificate
Server Public Key
Server Certificate
Server Certificate
Server Public Key
Server Key Exchange
Export-Grade
DH Params g, p
Server Public DH Value Y_s
Signature
DH Params g, p
Server Public DH Value Y_s
Signature
Client Private DH Value X_c
Server Private DH Value X_s
Client Key Exchange
Export-Grade
Client Public DH Value Y_c
Client Public DH Value Y_c
Premaster Secret
Session Key
Premaster Secret
Session Key

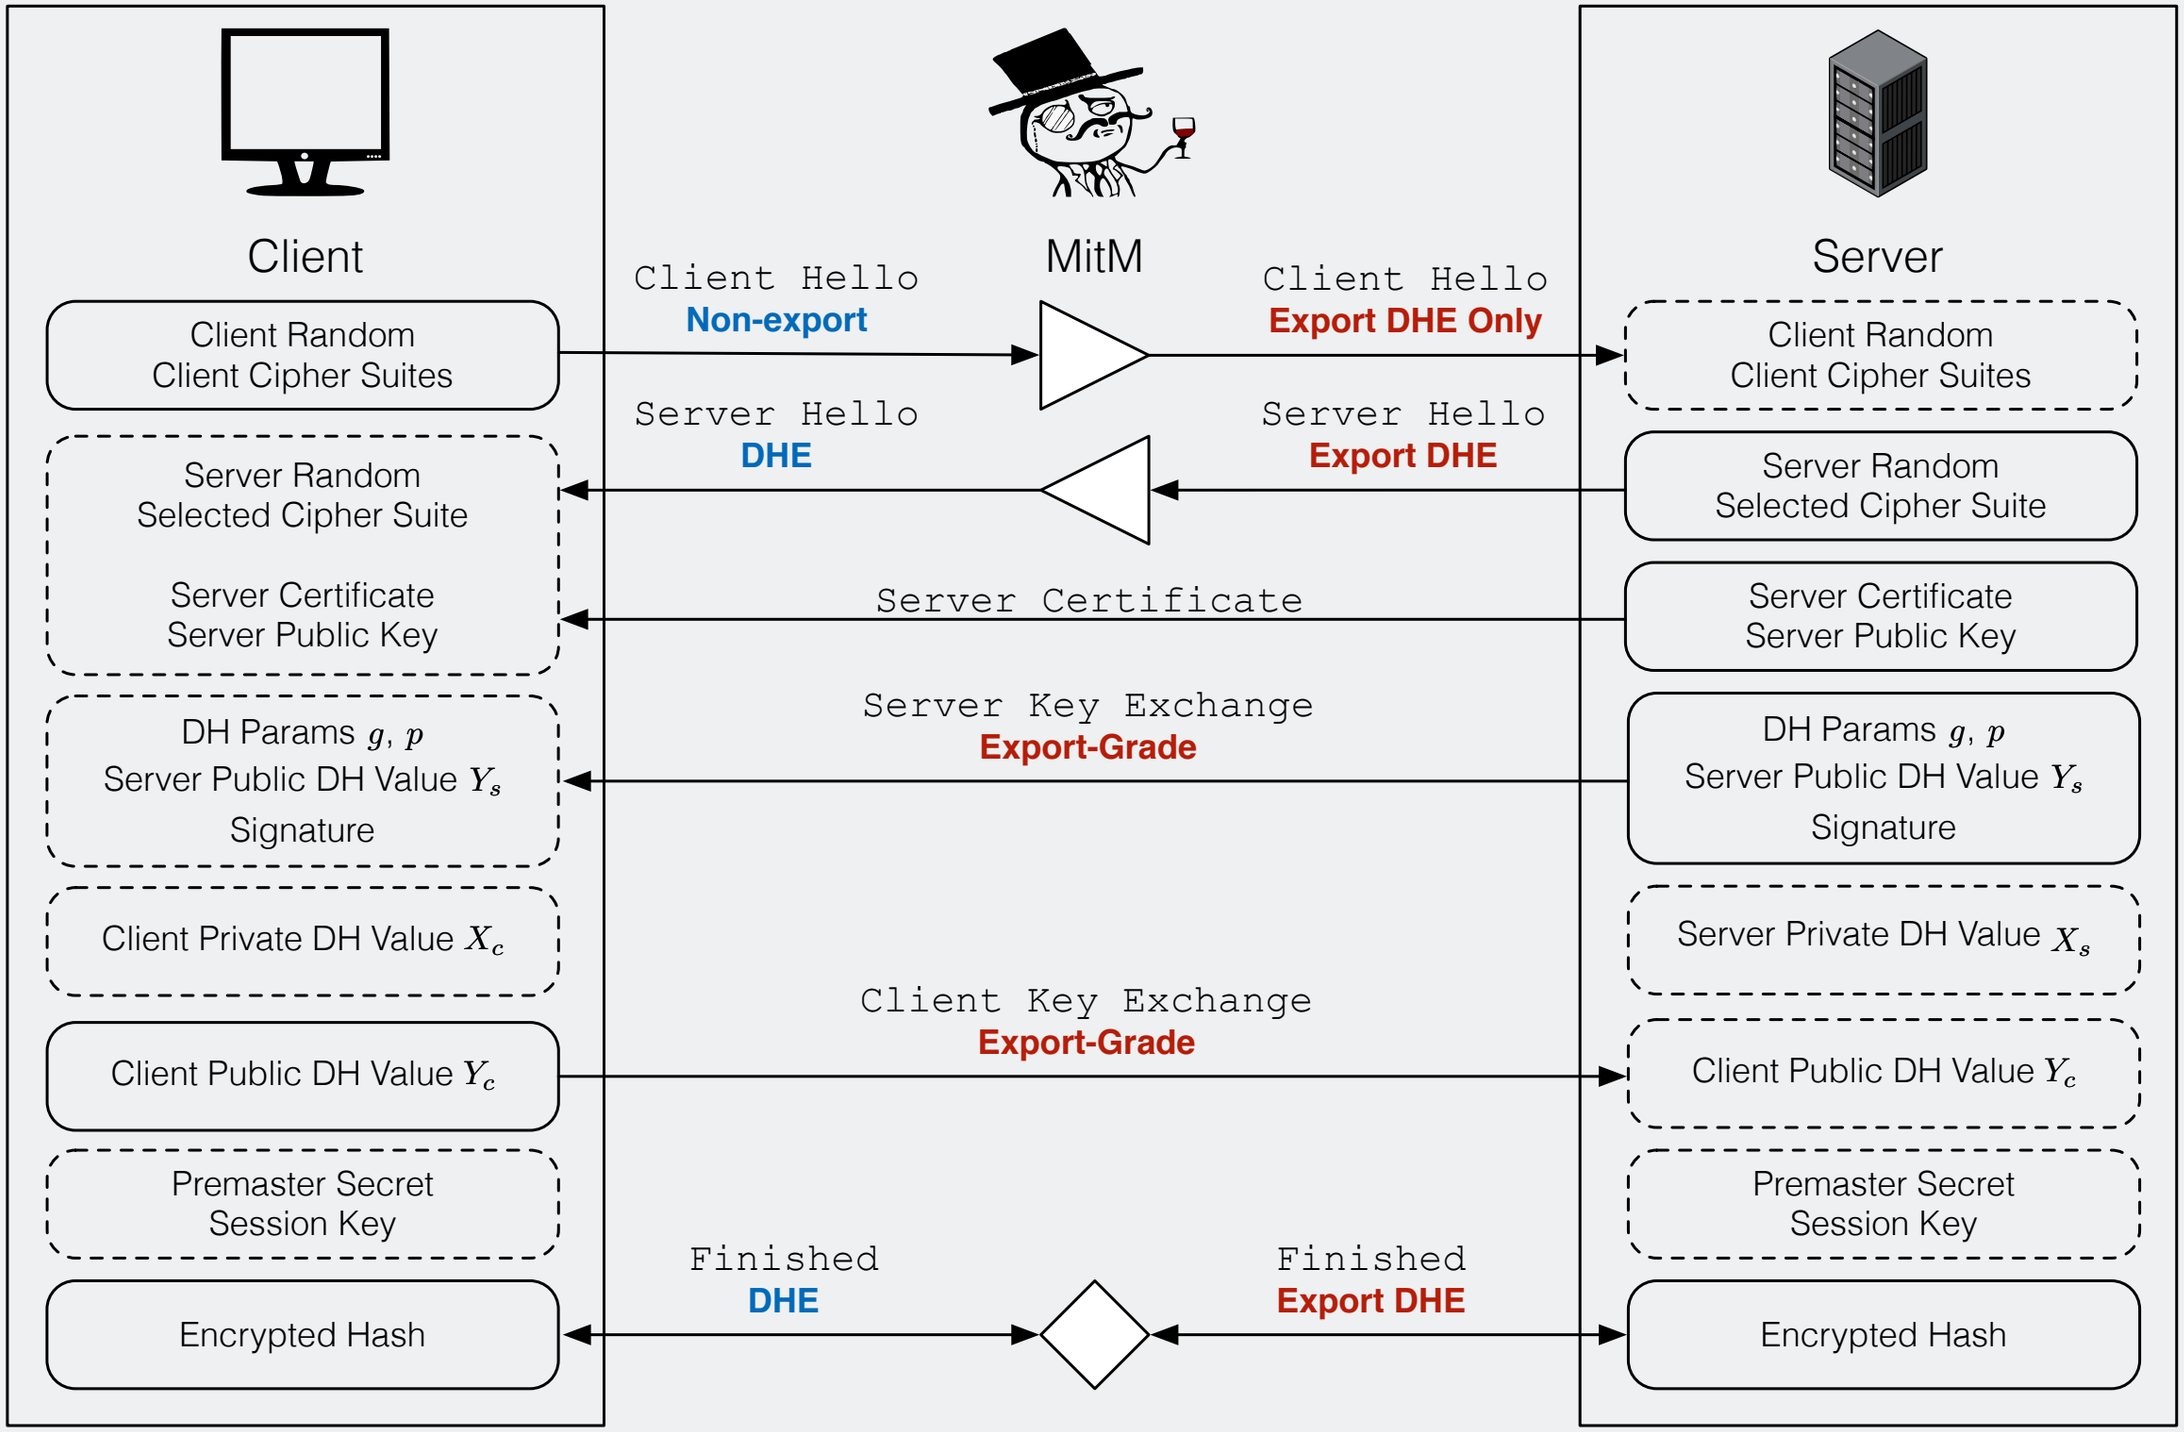
Client
MitM
Server
Client Hello
Non-export
Client Hello
Export DHE Only
Client Random
Client Cipher Suites
Client Random
Client Cipher Suites
Server Hello
DHE
Server Hello
Export DHE
Server Random
Selected Cipher Suite
Server Random
Selected Cipher Suite
Server Certificate
Server Public Key
Server Certificate
Server Certificate
Server Public Key
Server Key Exchange
Export-Grade
DH Params g, p
Server Public DH Value Y_s
Signature
DH Params g, p
Server Public DH Value Y_s
Signature
Client Private DH Value X_c
Server Private DH Value X_s
Client Key Exchange
Export-Grade
Client Public DH Value Y_c
Client Public DH Value Y_c
Premaster Secret
Session Key
Premaster Secret
Session Key
Finished
DHE
Finished
Export DHE
Encrypted Hash
Encrypted Hash

~~“Surely, no one uses 512-bit primes?”~~

“Surely, no one uses export ciphers?”

In April 2015, **8.4%** of the Top 1 Million HTTPS domains support DHE_EXPORT.

- **82%** used the most common prime, and **10%** used the second most common

We carried out the precomputation on these primes in **7 days** each

- Break single connections in **under two minutes**.

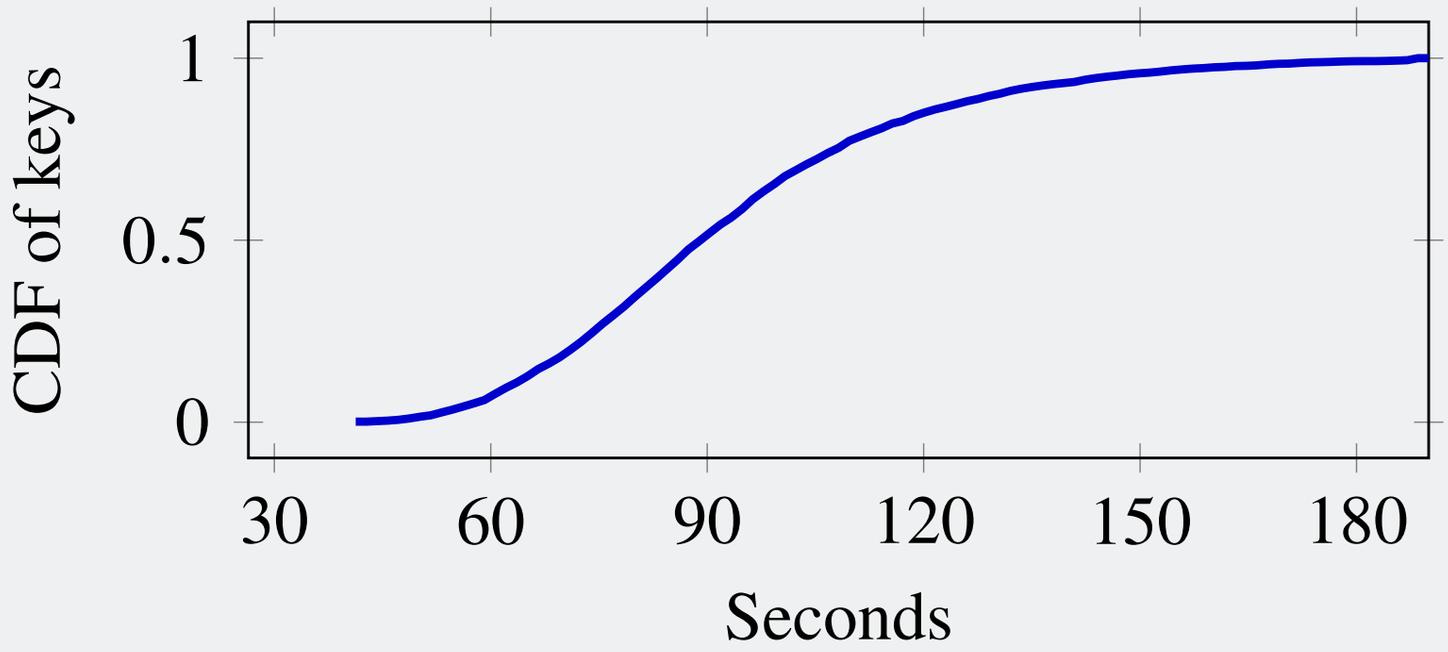

~~“Surely, no one uses 512-bit primes?”~~

~~“Surely, no one uses export ciphers?”~~

“Surely, no one is actually exploiting this?”

What if we could break **1024-bit** Diffie-Hellman?

Regular, non-export connections already uses 1024-bit primes. No need to downgrade, just passively decrypt!

Prime Length	Could Be Broken By	Precomputation Time
512 bits	**Academics**	7 days
768 bits	**Academics**	*~1 month*
1024 bits	**Nation State** **Large Organization**	*~1 year* *~$100-300 million*

Prime Length	Could Be Broken By	Precomputation Time
512 bits	**Academics**	7 days
768 bits	**Academics**	*~1 month*
1024 bits	**Nation State** **Large Organization**	*~1 year* *~$100-300 million*

~~“Surely, no one uses 512-bit primes?”~~

~~“Surely, no one uses export ciphers?”~~

~~“Surely, no one is actually exploiting this?”~~

“Surely, people are using more than one prime?”

	Top 1024-bit Prime	Top Ten 1024-bit Primes
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%)

	Top 1024-bit Prime	Top Ten 1024-bit Primes
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%)

	Top 1024-bit Prime	Top Ten 1024-bit Primes
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%)

"Also, we are investing in groundbreaking cryptanalytic capabilities to defeat adversarial cryptography and exploit Internet traffic"

–NSA, 2013

"Also, we are investing in groundbreaking cryptanalytic capabilities to defeat adversarial cryptography and exploit Internet traffic"

–NSA, 2013

Also, the have a $10 billion budget just to break crypto.

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

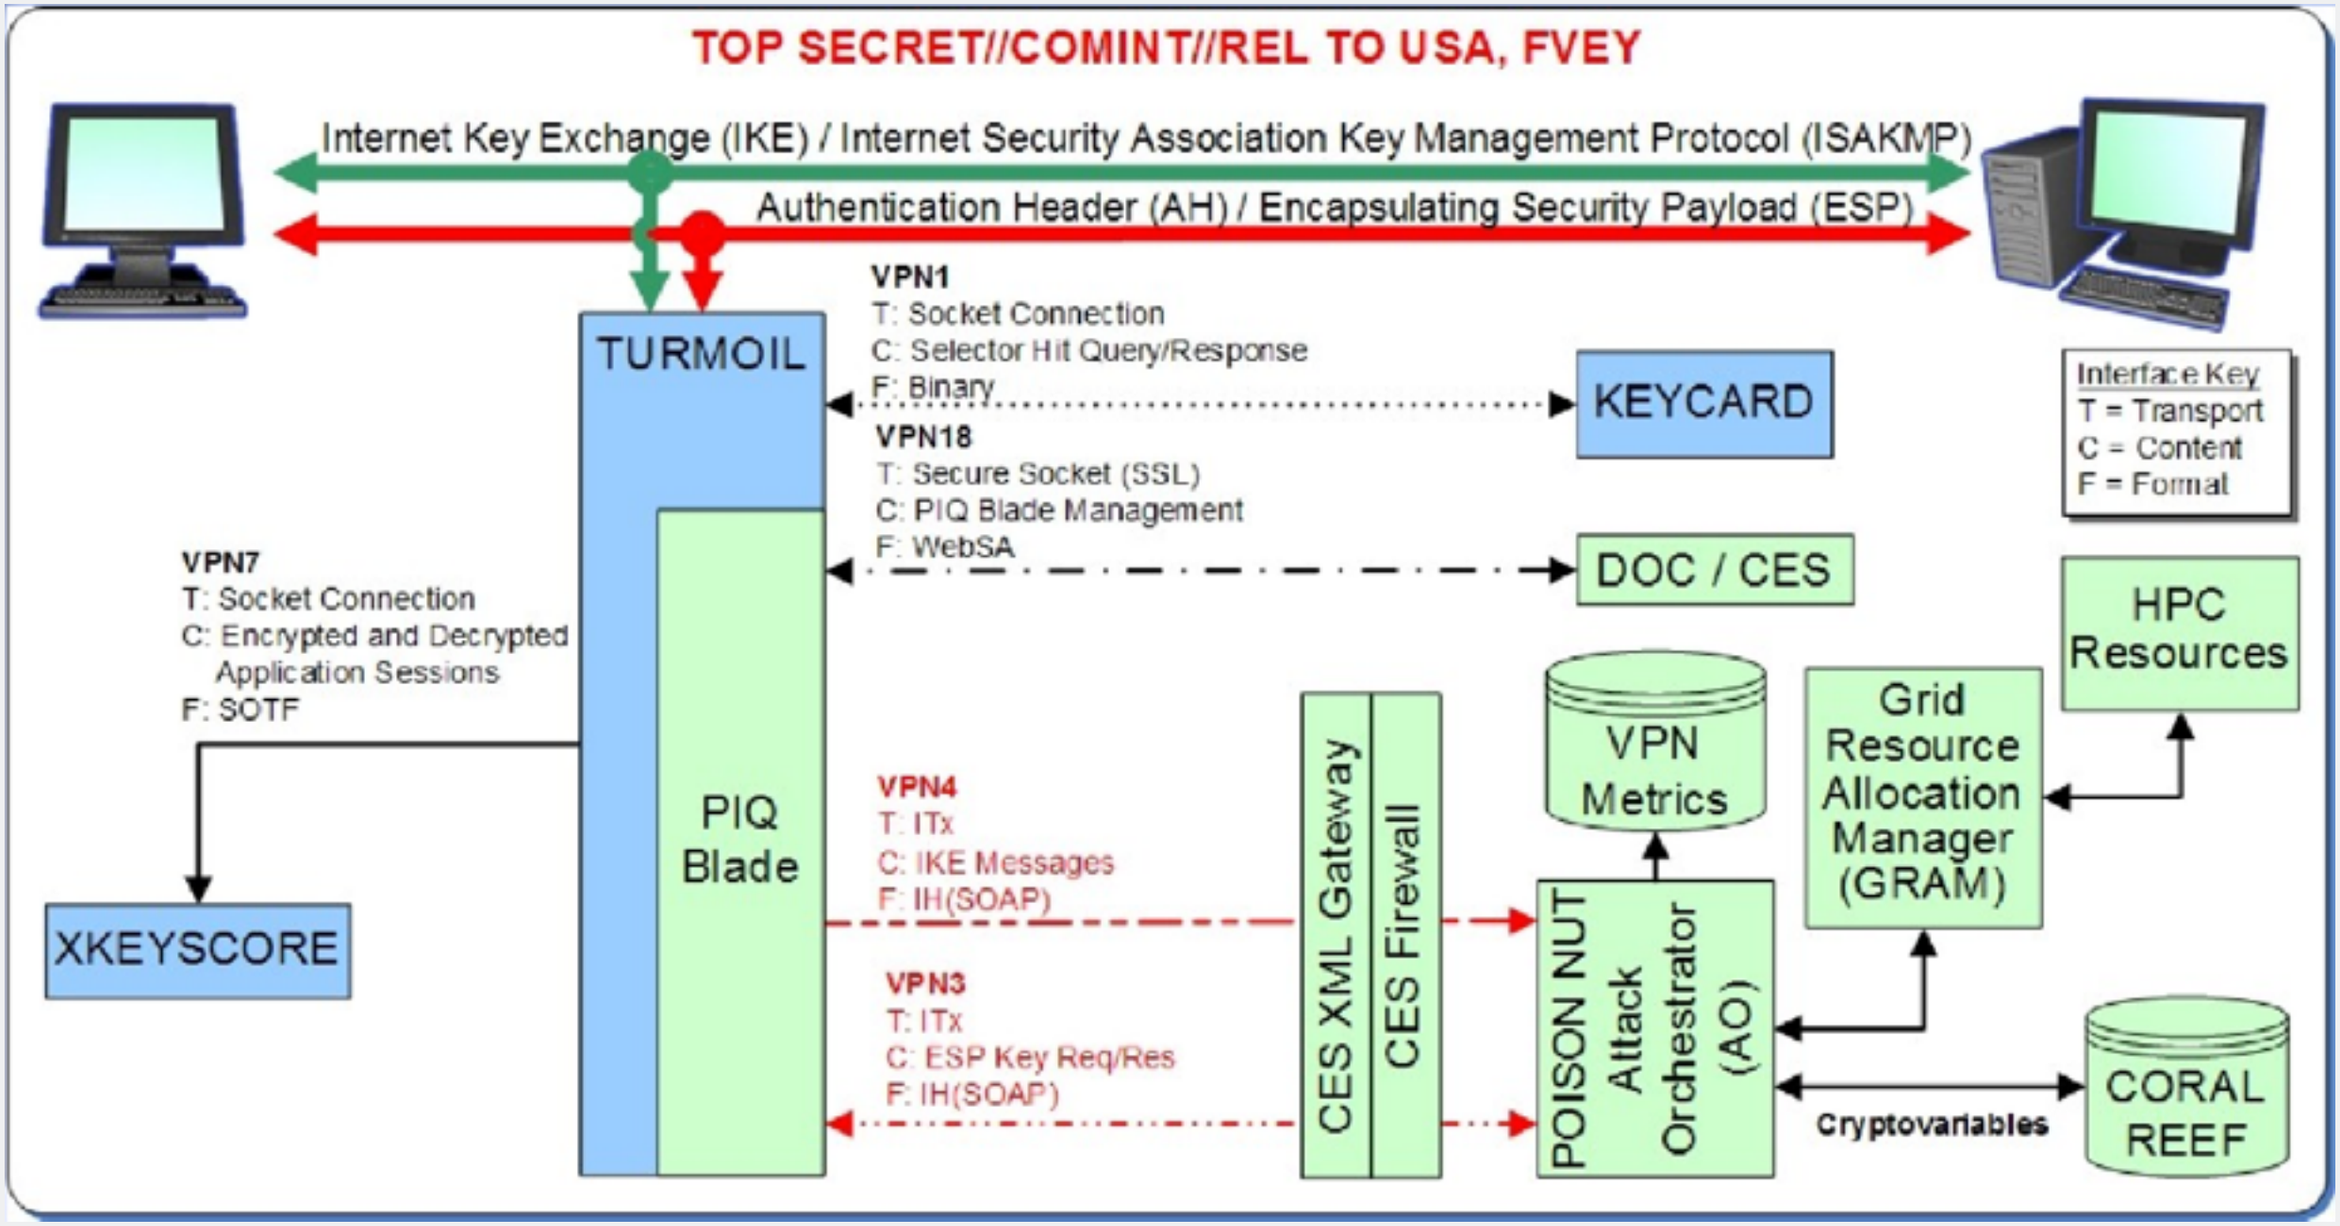

Passive decryption of VPN connections using a broken 1024-bit prime is consistent with Snowden documents

Mitigations and Lessons

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

Browser will reject 512-bit groups, and are sunsetting 768-bit and 1024-bit

Turn export ciphers off!

Questions?

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

@davidcadrian

https://weakdh.org

https://zmap.io

https://davidadrian.org

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-Beguelin, and Paul Zimmermann.