The Evolving Architecture of the Web

Nick Sullivan

CLOUDFLARE

CLOUDFLARE

Head of Cryptography

CFSSL
Universal SSL
Keyless SSL
Privacy Pass
Geo Key Manager

Recently
Standards work
TLS 1.3

Competing Goals

make browsing more

HTTP

DNS

UTD

DNS

Hosts

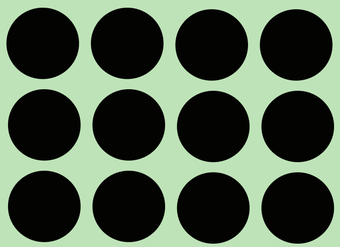

Clients

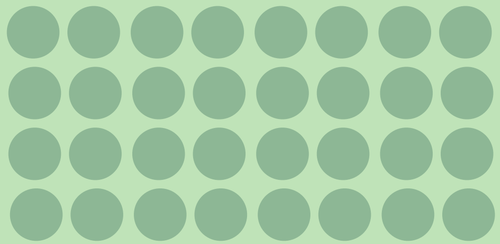

HTTP ⟶

Clients

Hosts

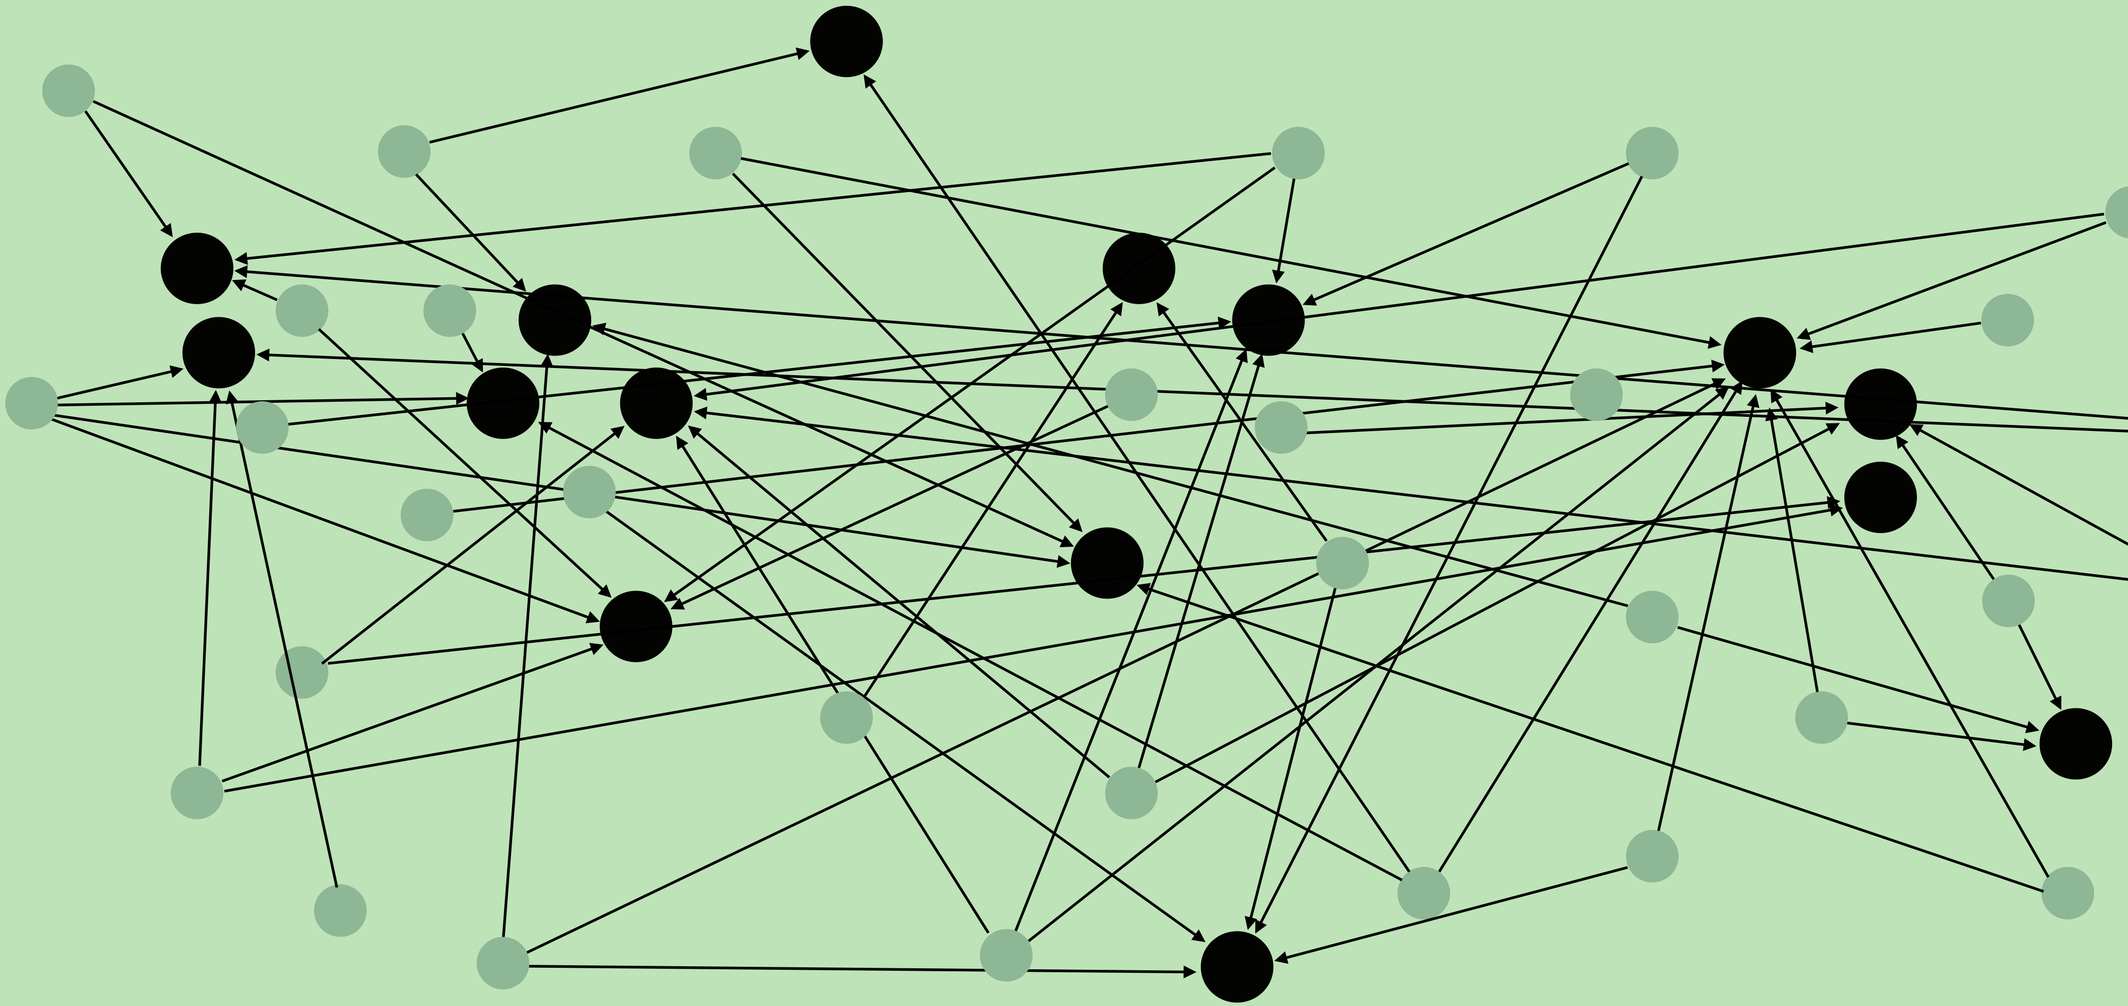

Geographically Centralized

Administratively Diverse

One IP per Hostname

What a network observer can see

HTTP ⟶

Unique Client IP

Unique Server IP

Server URL

Website content

Clients

Hosts

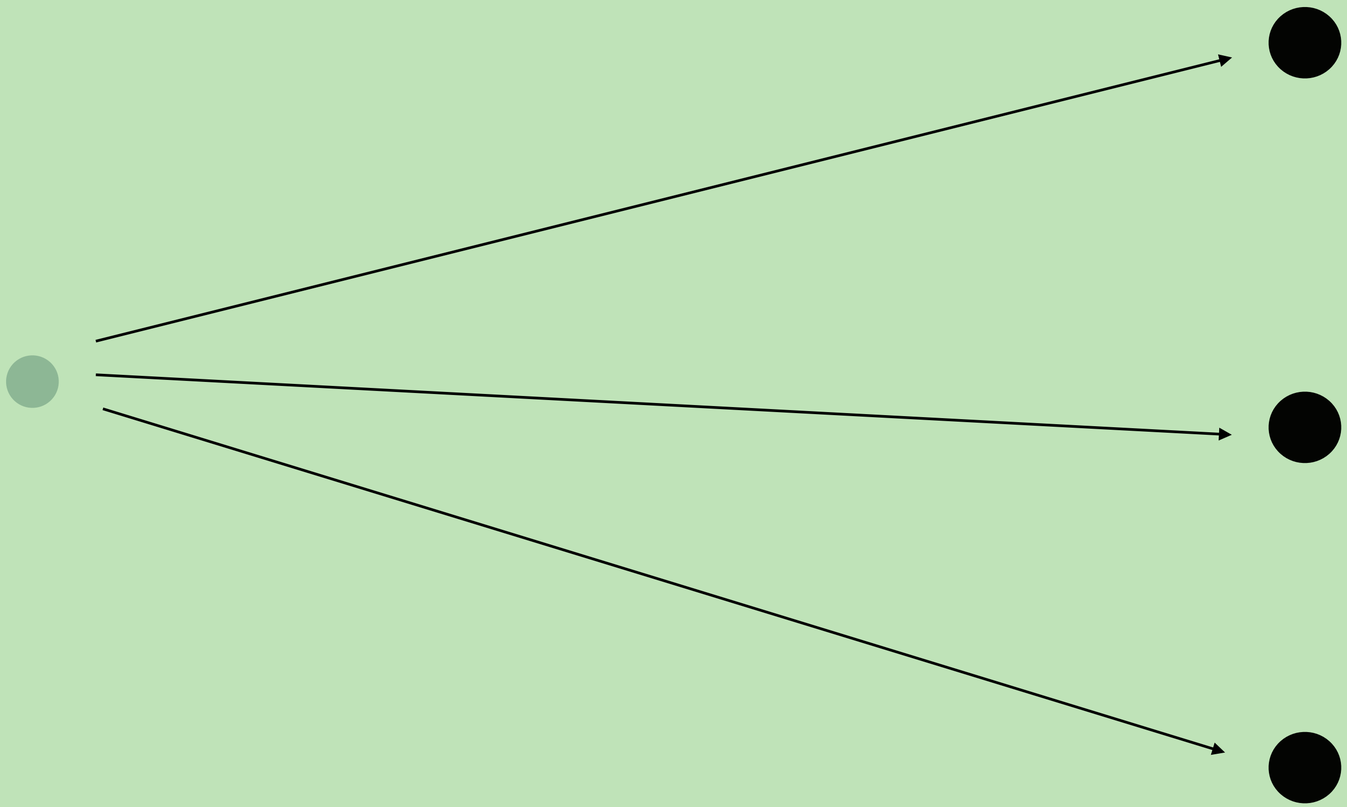

	Anonymity set
Client IP	**1**
Server IP	**1**

IPv4

4.3 Billion Addresses

Not enough for every user

What a network observer can see

HTTP ⟶

Client Proxy IP

Unique Server IP

Server URL

Website content

Clients

Hosts

Proxy

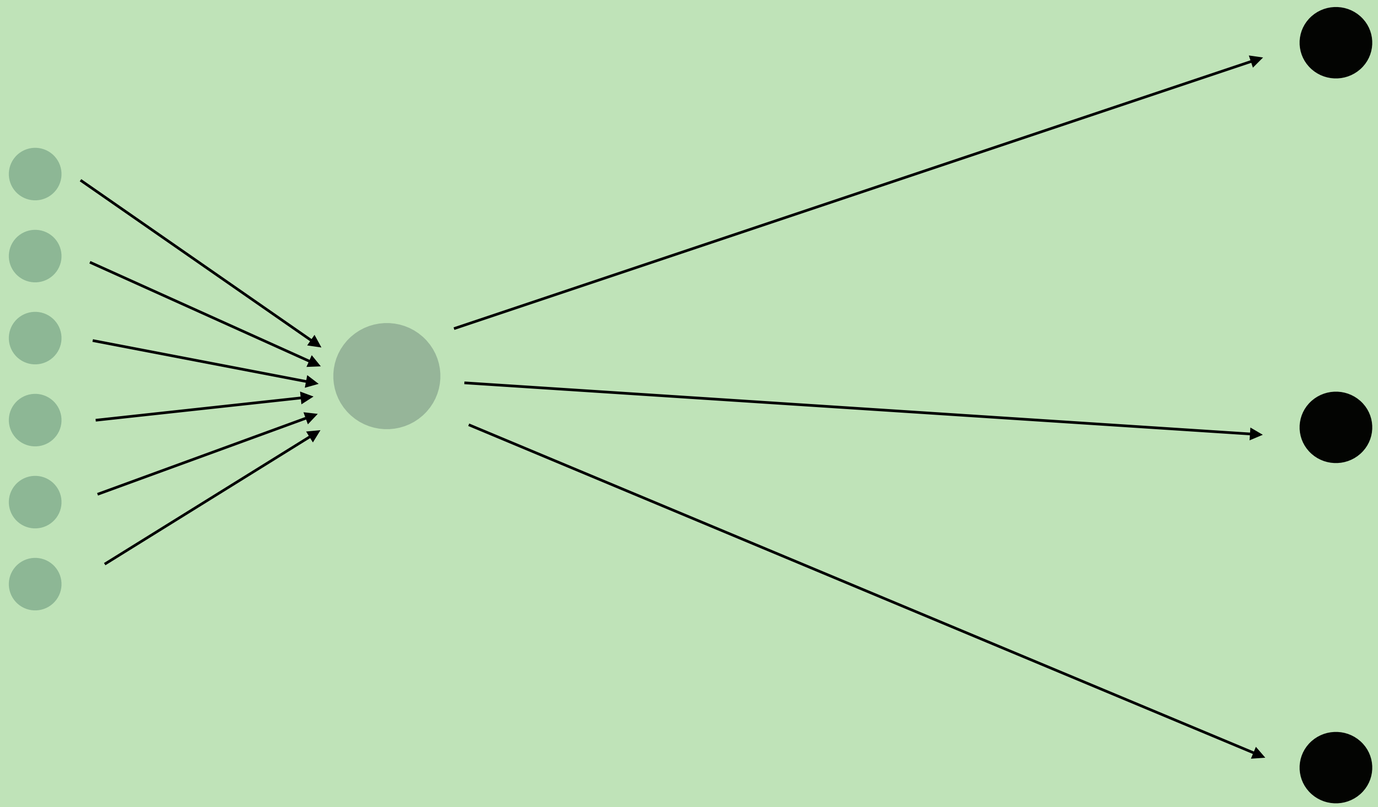

	Latency Cost
Tor	**3 round-the-world**
VPN	**1 round-the-world**
Carrier NAT	**Small**

	Anonymity set
Client	**k**
Server	**1**

New Trends

ENCRYPT THE WEB

HTTPS

Client

ISP

Host

TLS 1.2

TLS 1.3: coming soon

HTTP ⟶ HTTPS⟶

What a network observer can see

Unique Client IP

Unique Server IP

~~Server URL~~

~~Website content~~

Clients

Hosts

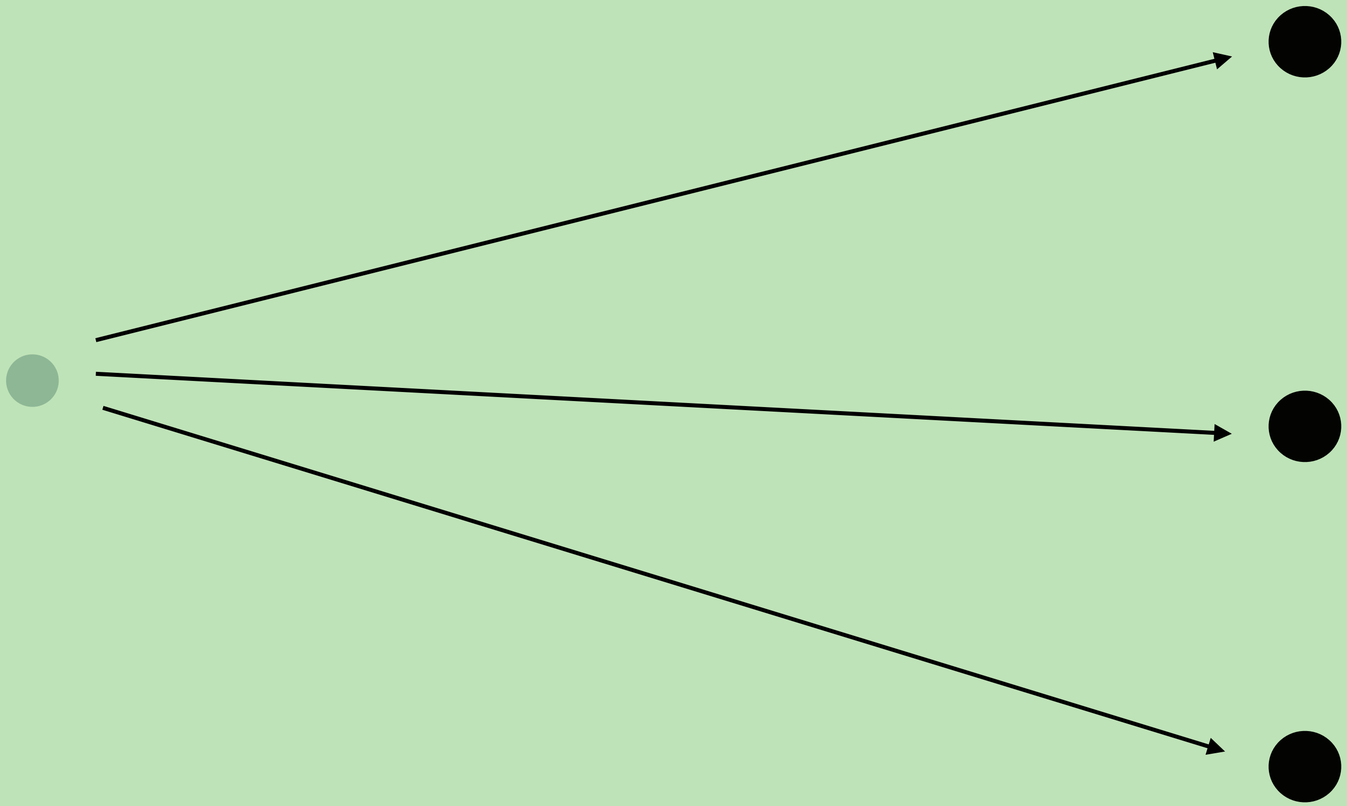

	Anonymity set
Client	**1**
Server	**1**

IPv4

4.3 Billion Addresses

Not enough for every website

HTTP ⟶

Clients

Hosts

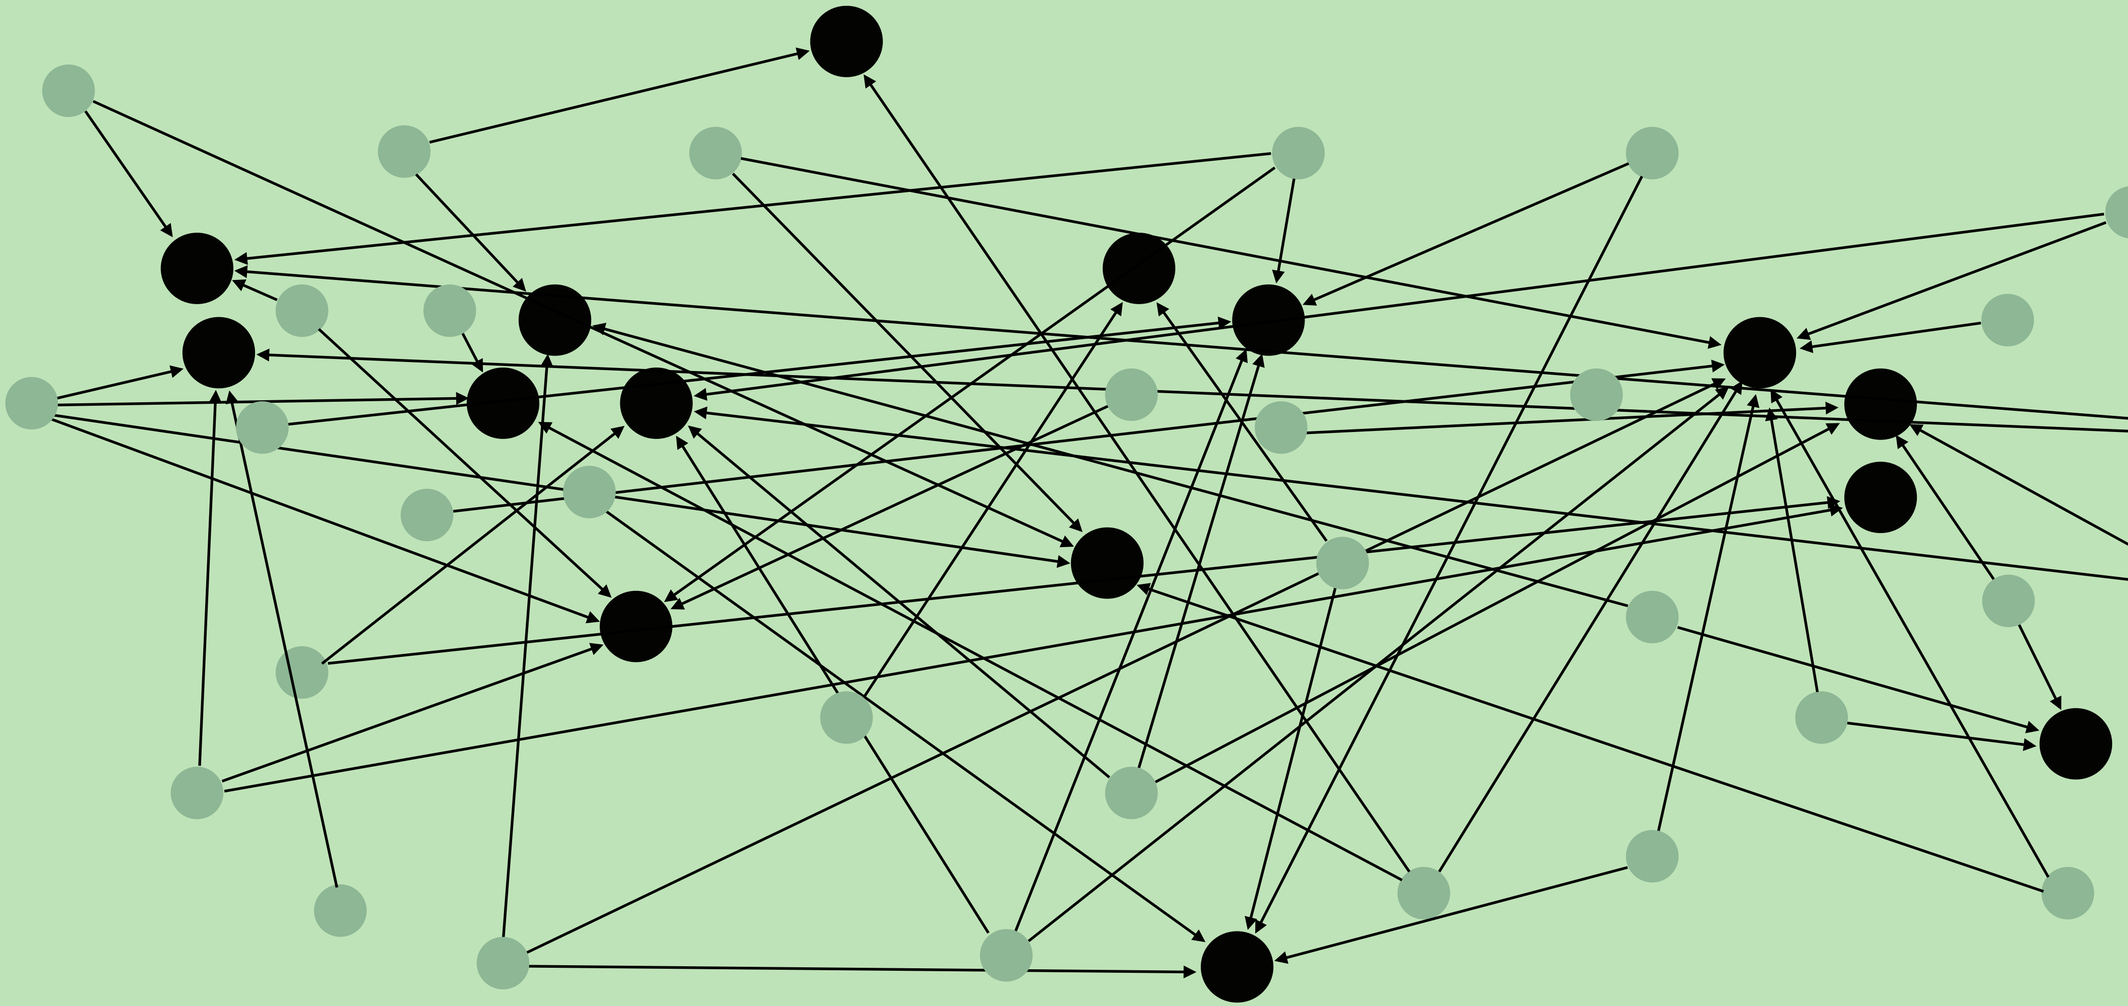

Geographically Centralized

Administratively Diverse

One IP per Hostname

HTTP → HTTPS →

Clients

Hosts

Shared Hosts

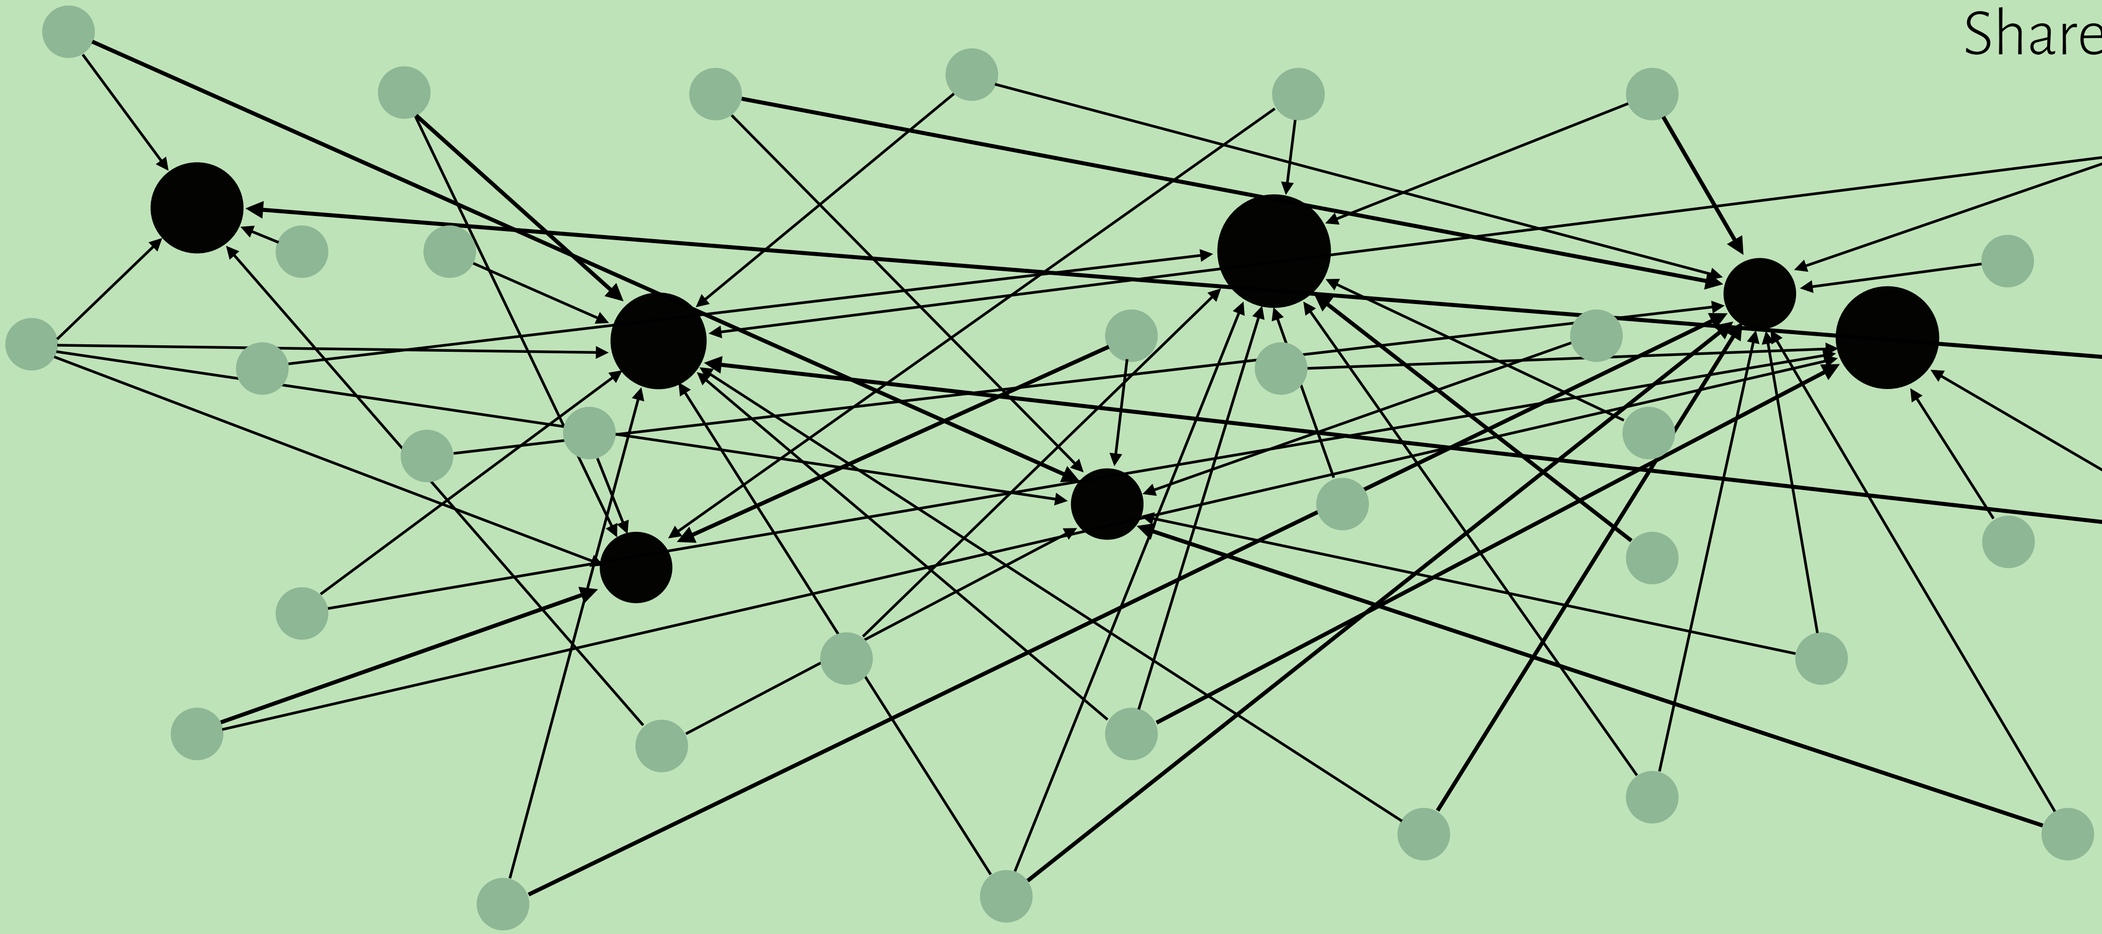

More Geographically Centralized

More Administratively Centralized

Multiple Hostname per IPs

SNI

Virtual Hosting

Send the hostname to the server
so it can choose the certificate

100.0%
95.0%
90.0%
85.0%
80.0%
75.0%
Jan/14
Jul/14
Jan/15
Jul/15
Jan/16
Jul/16
Jan/17
Source: Akamai

HTTP ⟶ HTTPS⟶

What a network observer can see

Client Unique IP

Shared Server IP

Hostname

Clients

Hosts

Shared Hosts

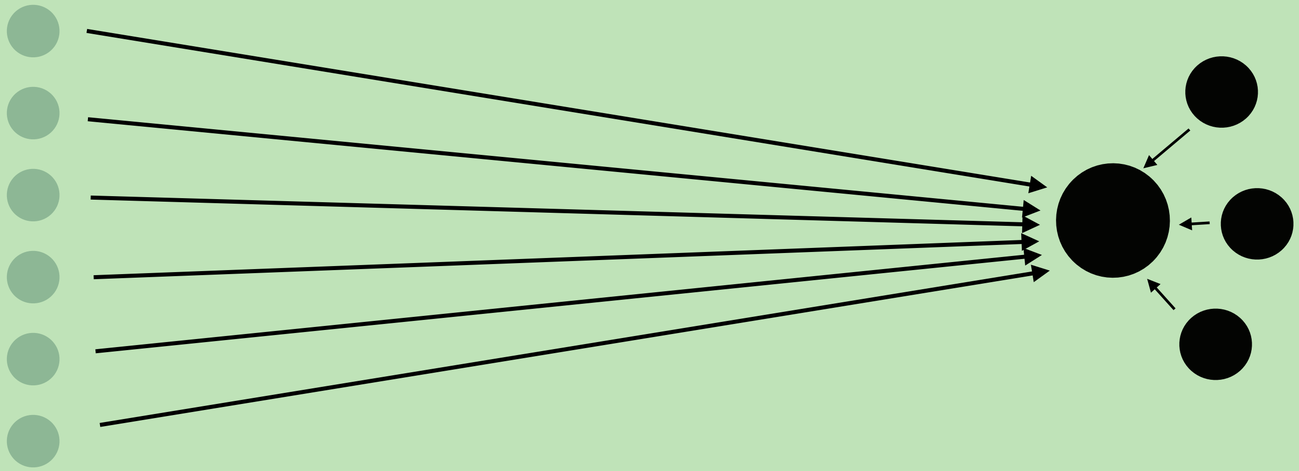

	Anonymity set
Client	**1**
Server (Shared IP+Hostname)	**1**

Internet Scans and IPv6

Privacy Evolves

Certificate Transparency

Wildcard certificates

Edge Services

Edge Services

- Websites and are delegating to globally distributed parties
- Authorized to terminate TLS
- Reduced Latency
- Improved DDoS resilience
- Anycast to reduce number of IPs needed

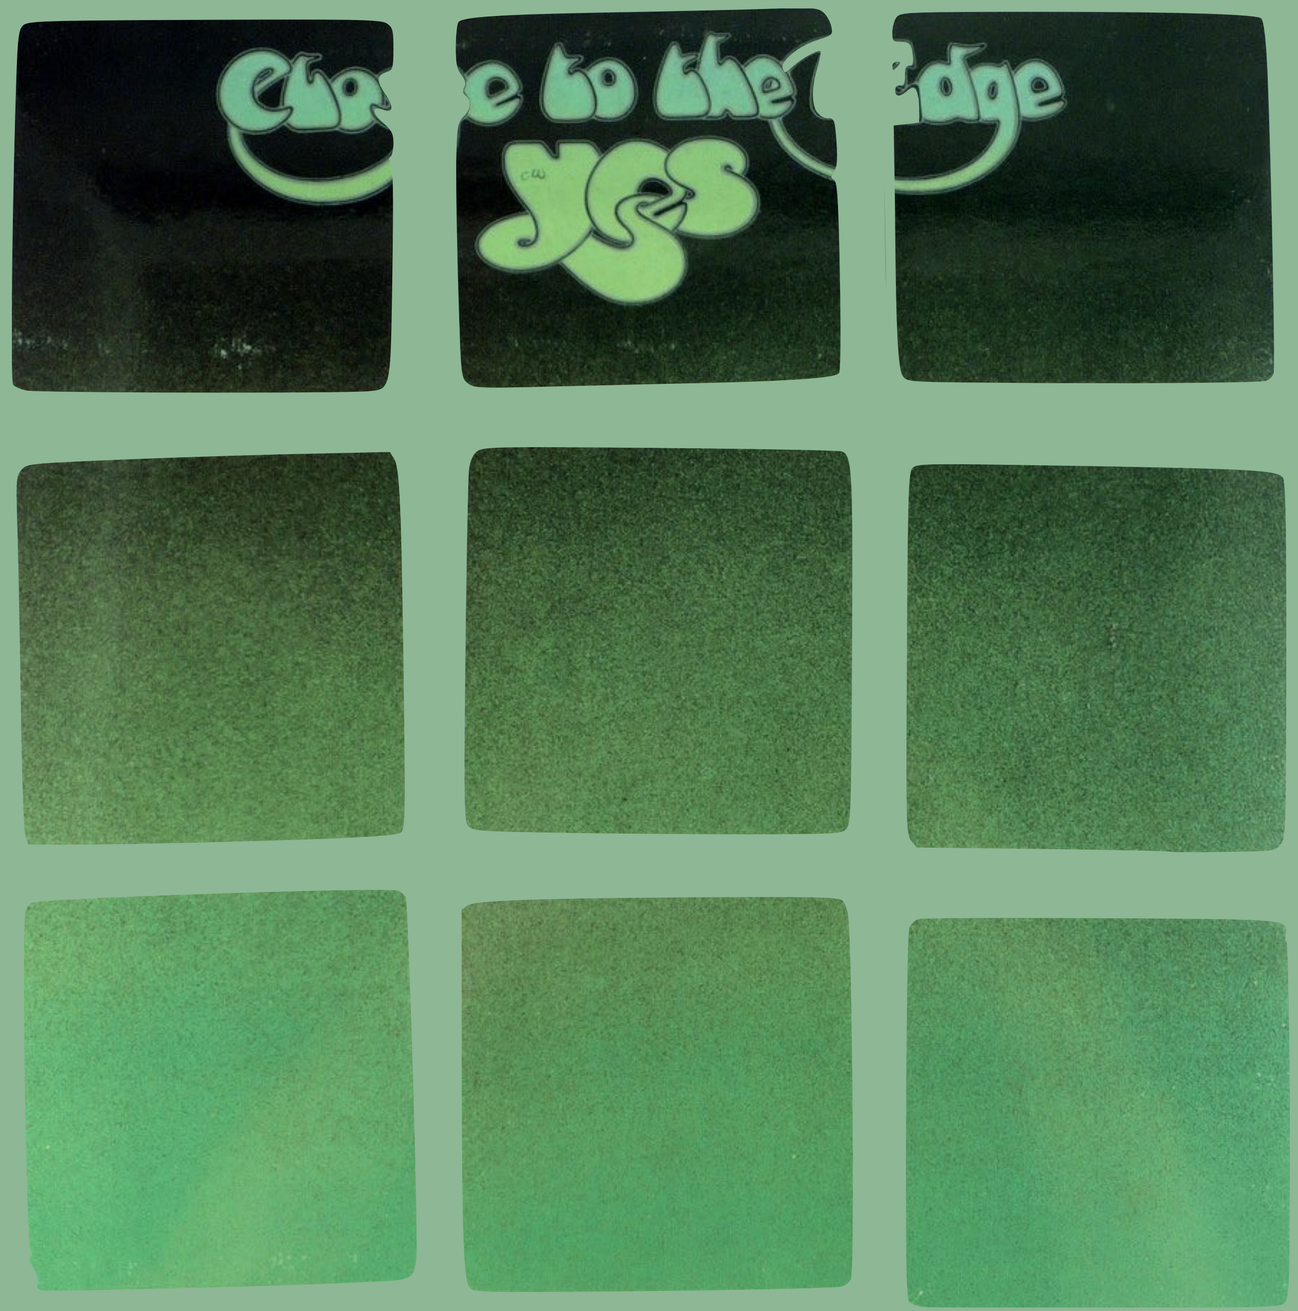

HTTP → HTTPS →

Clients

Hosts

Shared Hosts

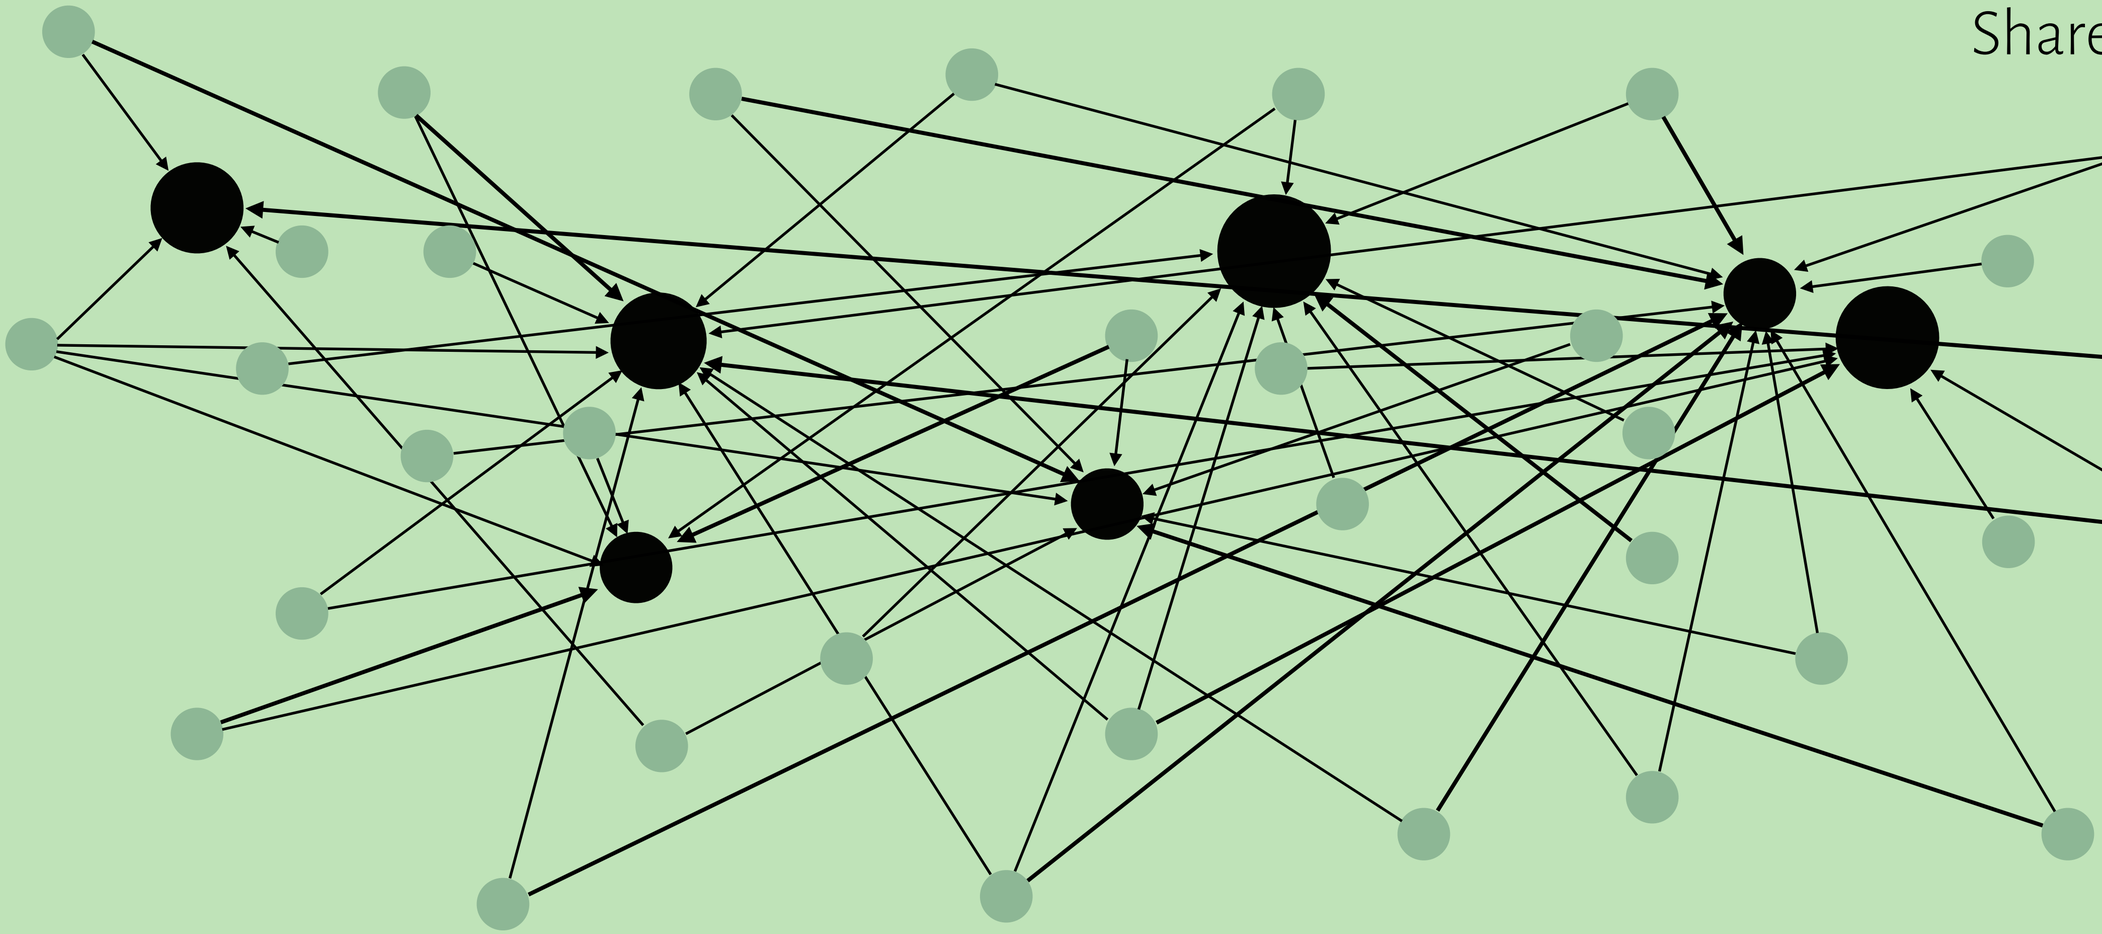

More Geographically Centralized

More Administratively Centralized

Multiple IPs per Hostname

HTTP → HTTPS →

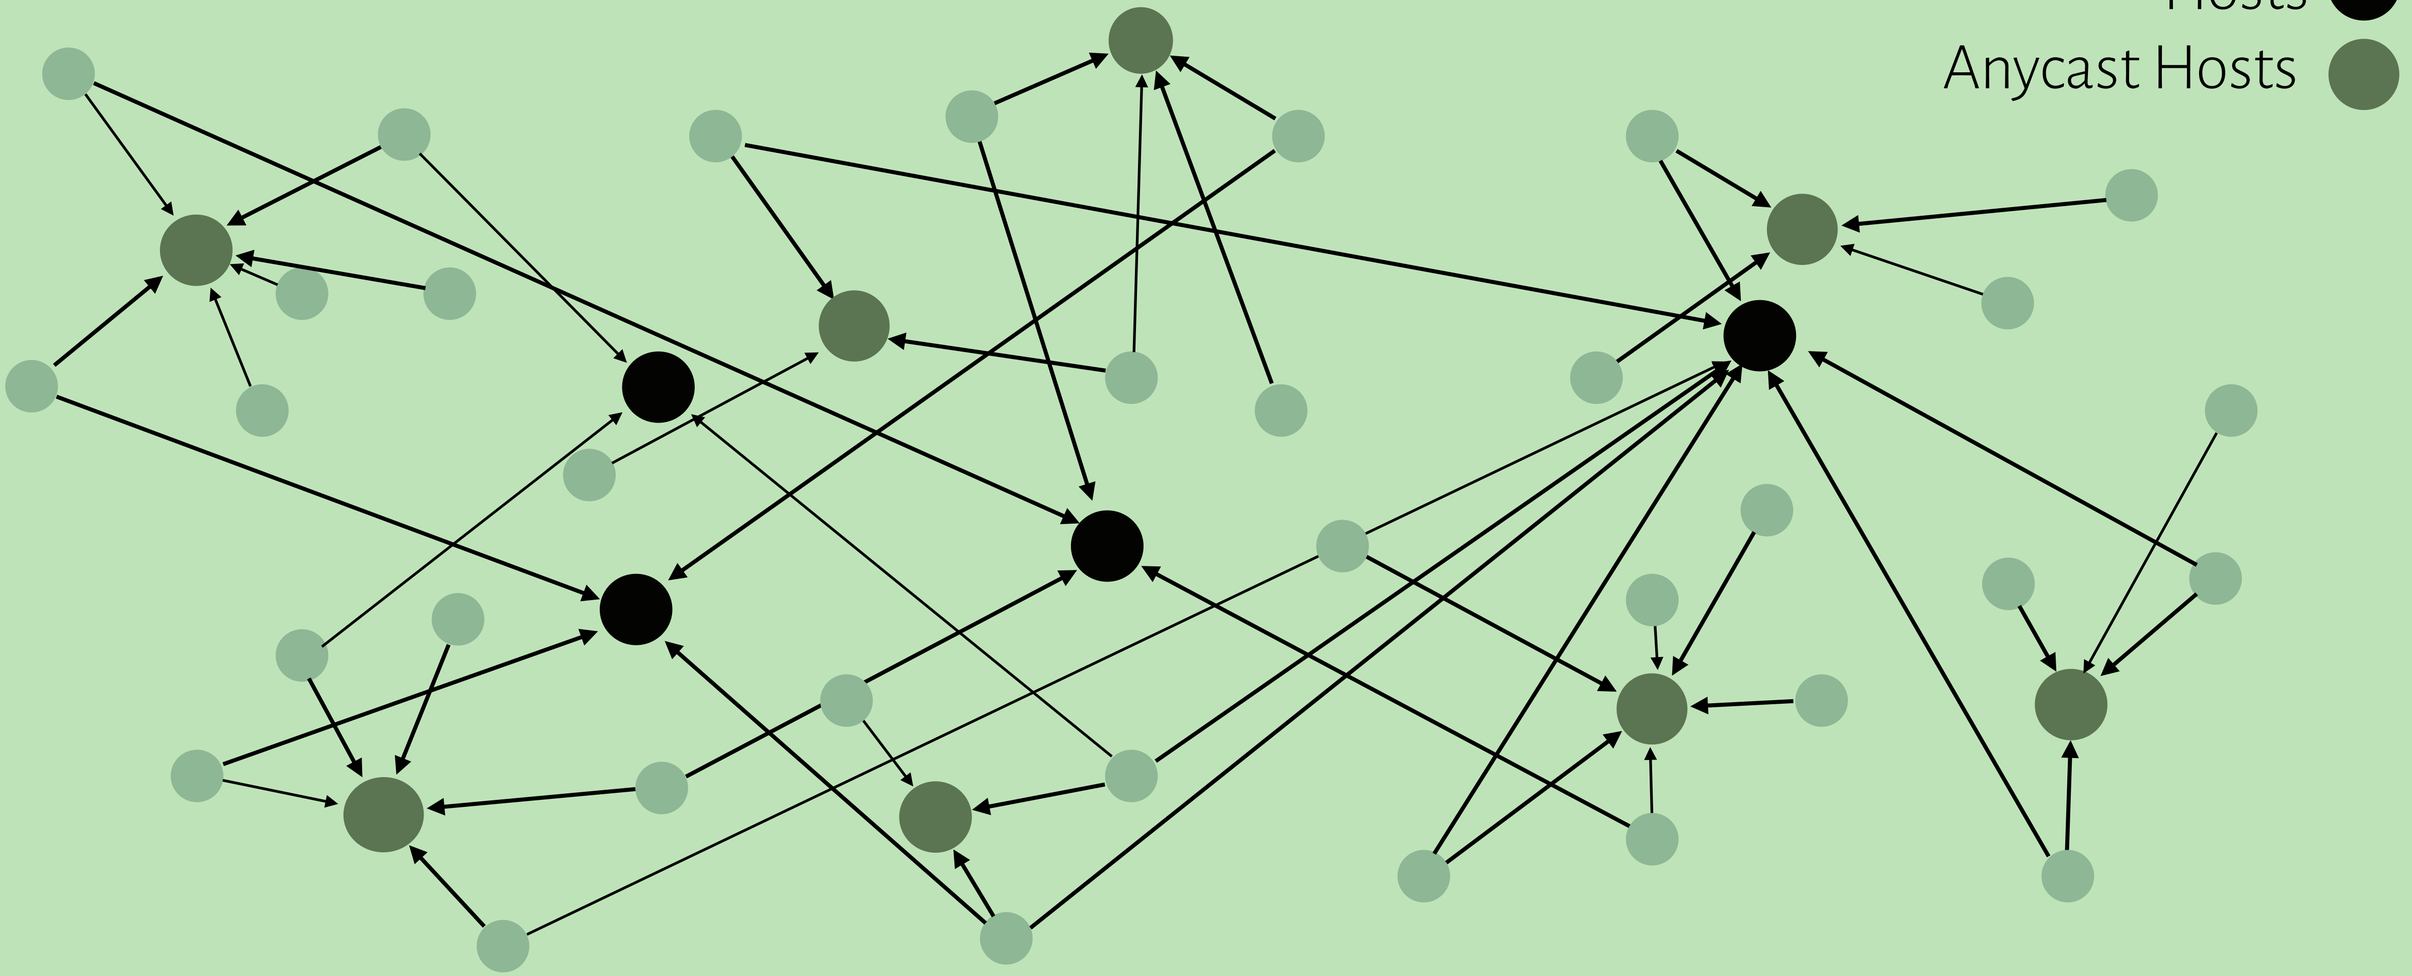

Geographically Distributed
Administratively Centralized

Multiple IPs per Hostname

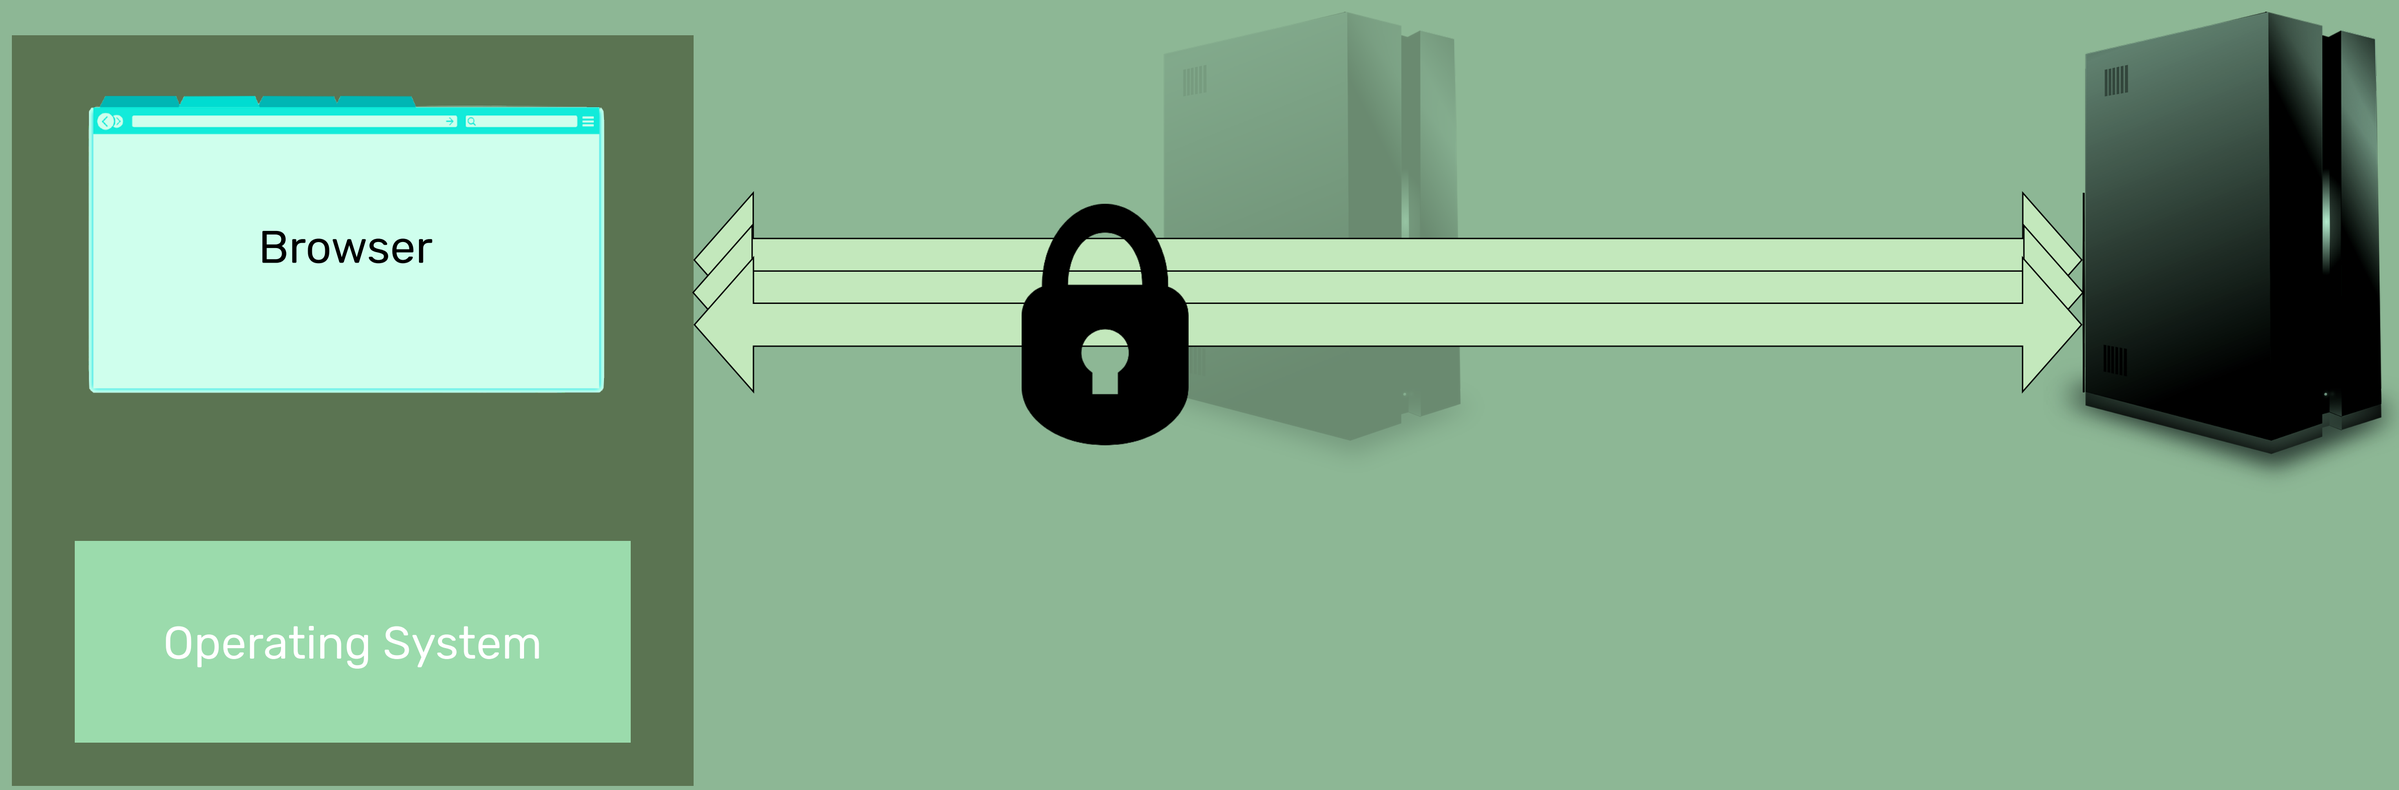

HTTPS

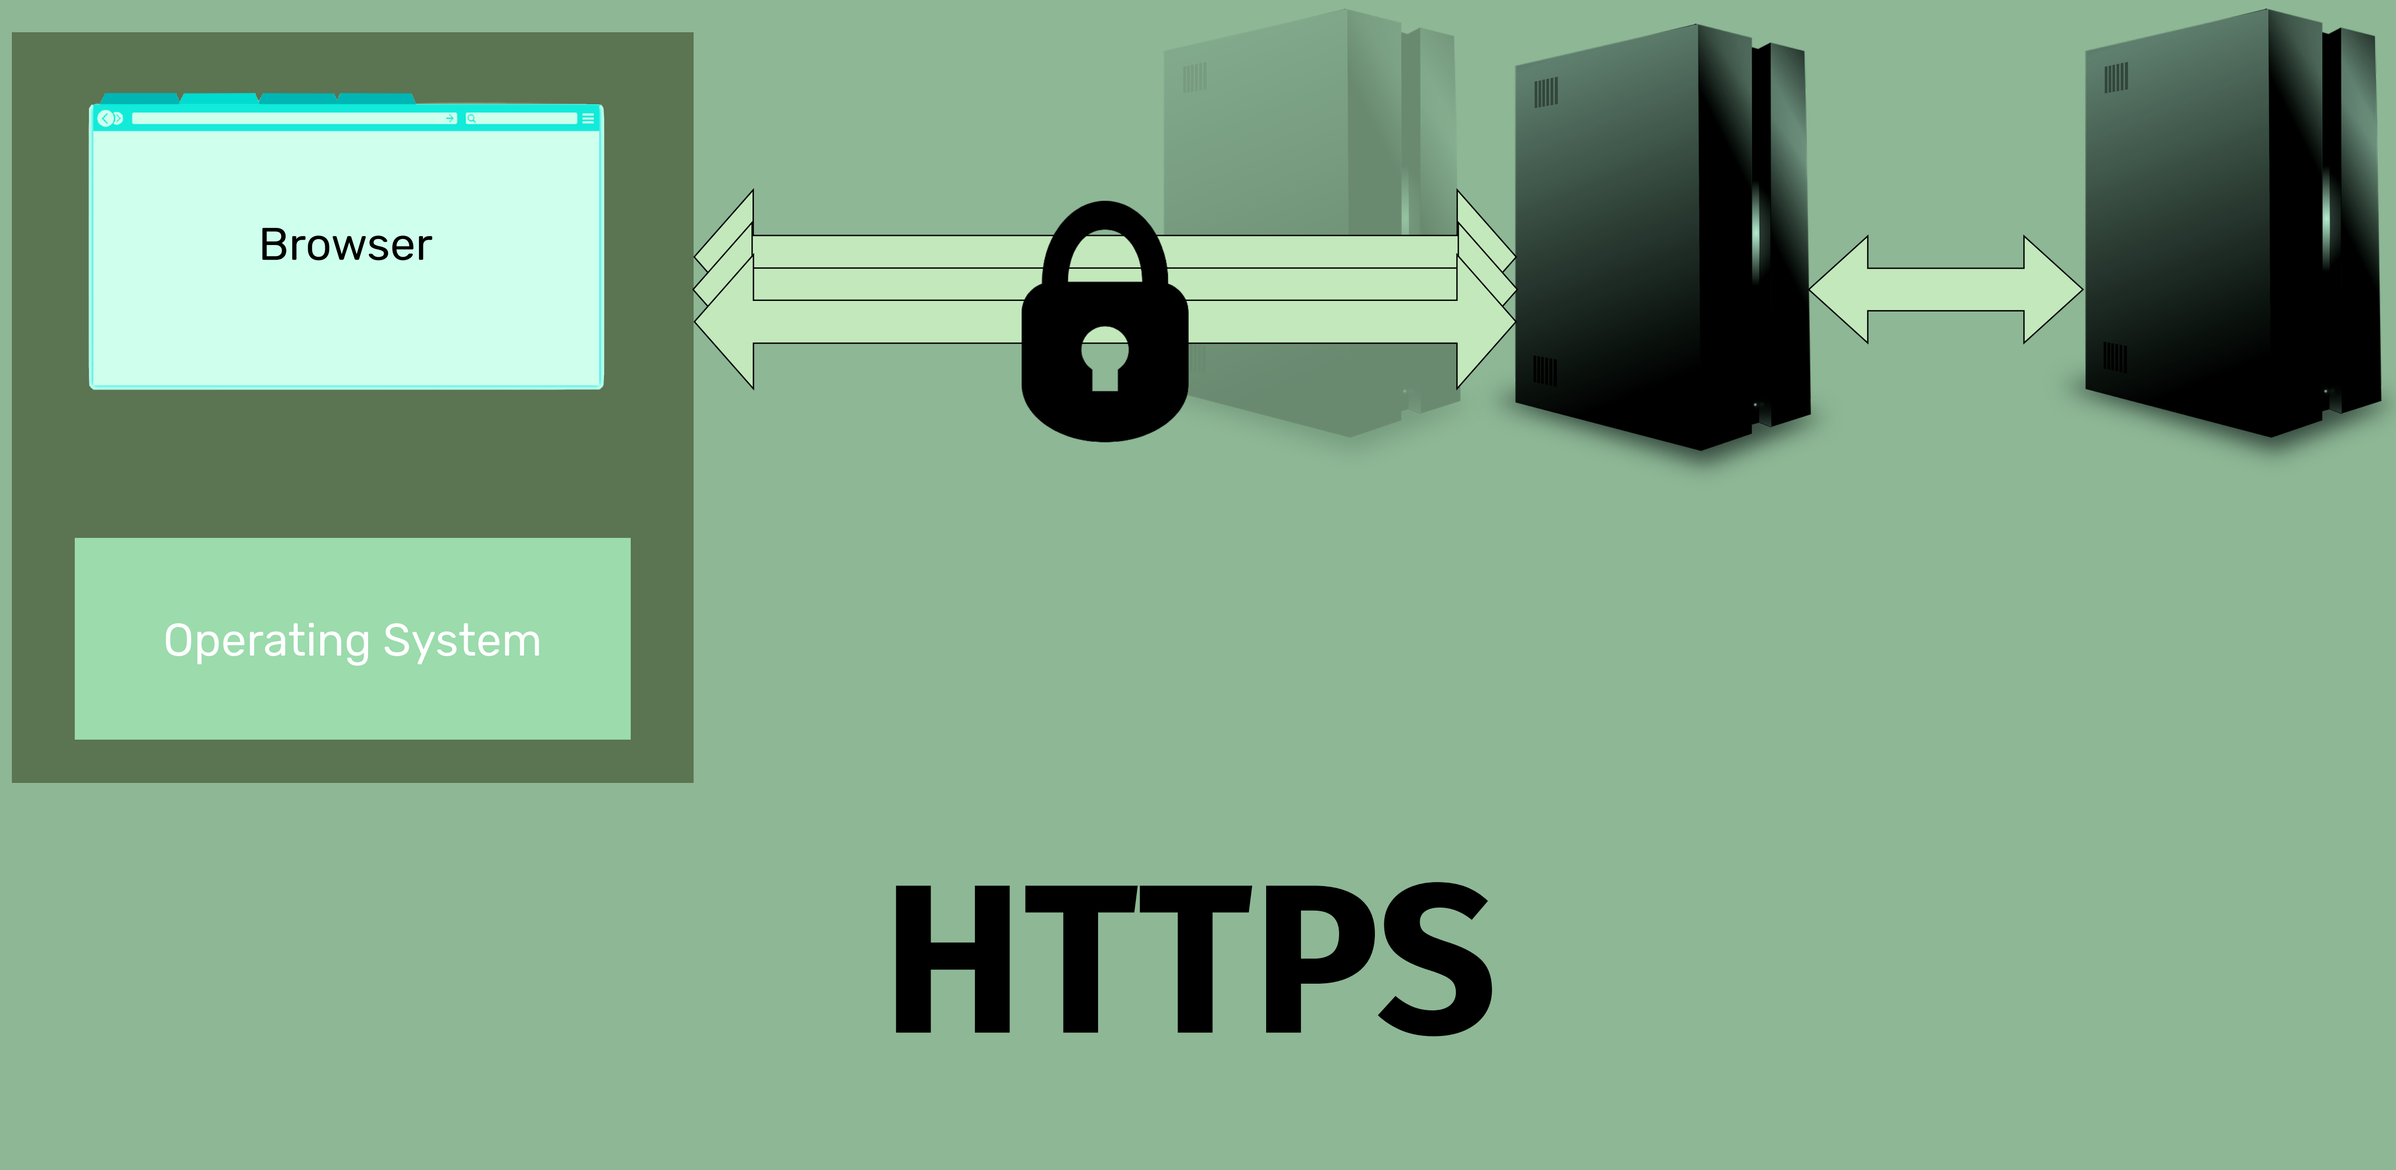

HTTPS

Questions

Can we improve ***privacy***?
Can we improve ***latency***?
Can we improve ***both***???

Client

HTTP 1.1

Resolver

Edge

Browser

Operating System

SNI: burrito.com

SNI: beans.com

Q: burrito.com A: 1.2.3.4

Q: beans.com A: 1.2.3.5

burrito.com

beans.com

Safety in Numbers

Meek

Client

Edge

Origin

Browser

Operating System

SNI: burrito.com

burrito.com
Host

beans.com
Host

Resolver

Q: burrito.com A: 1.2.3.4

Meek

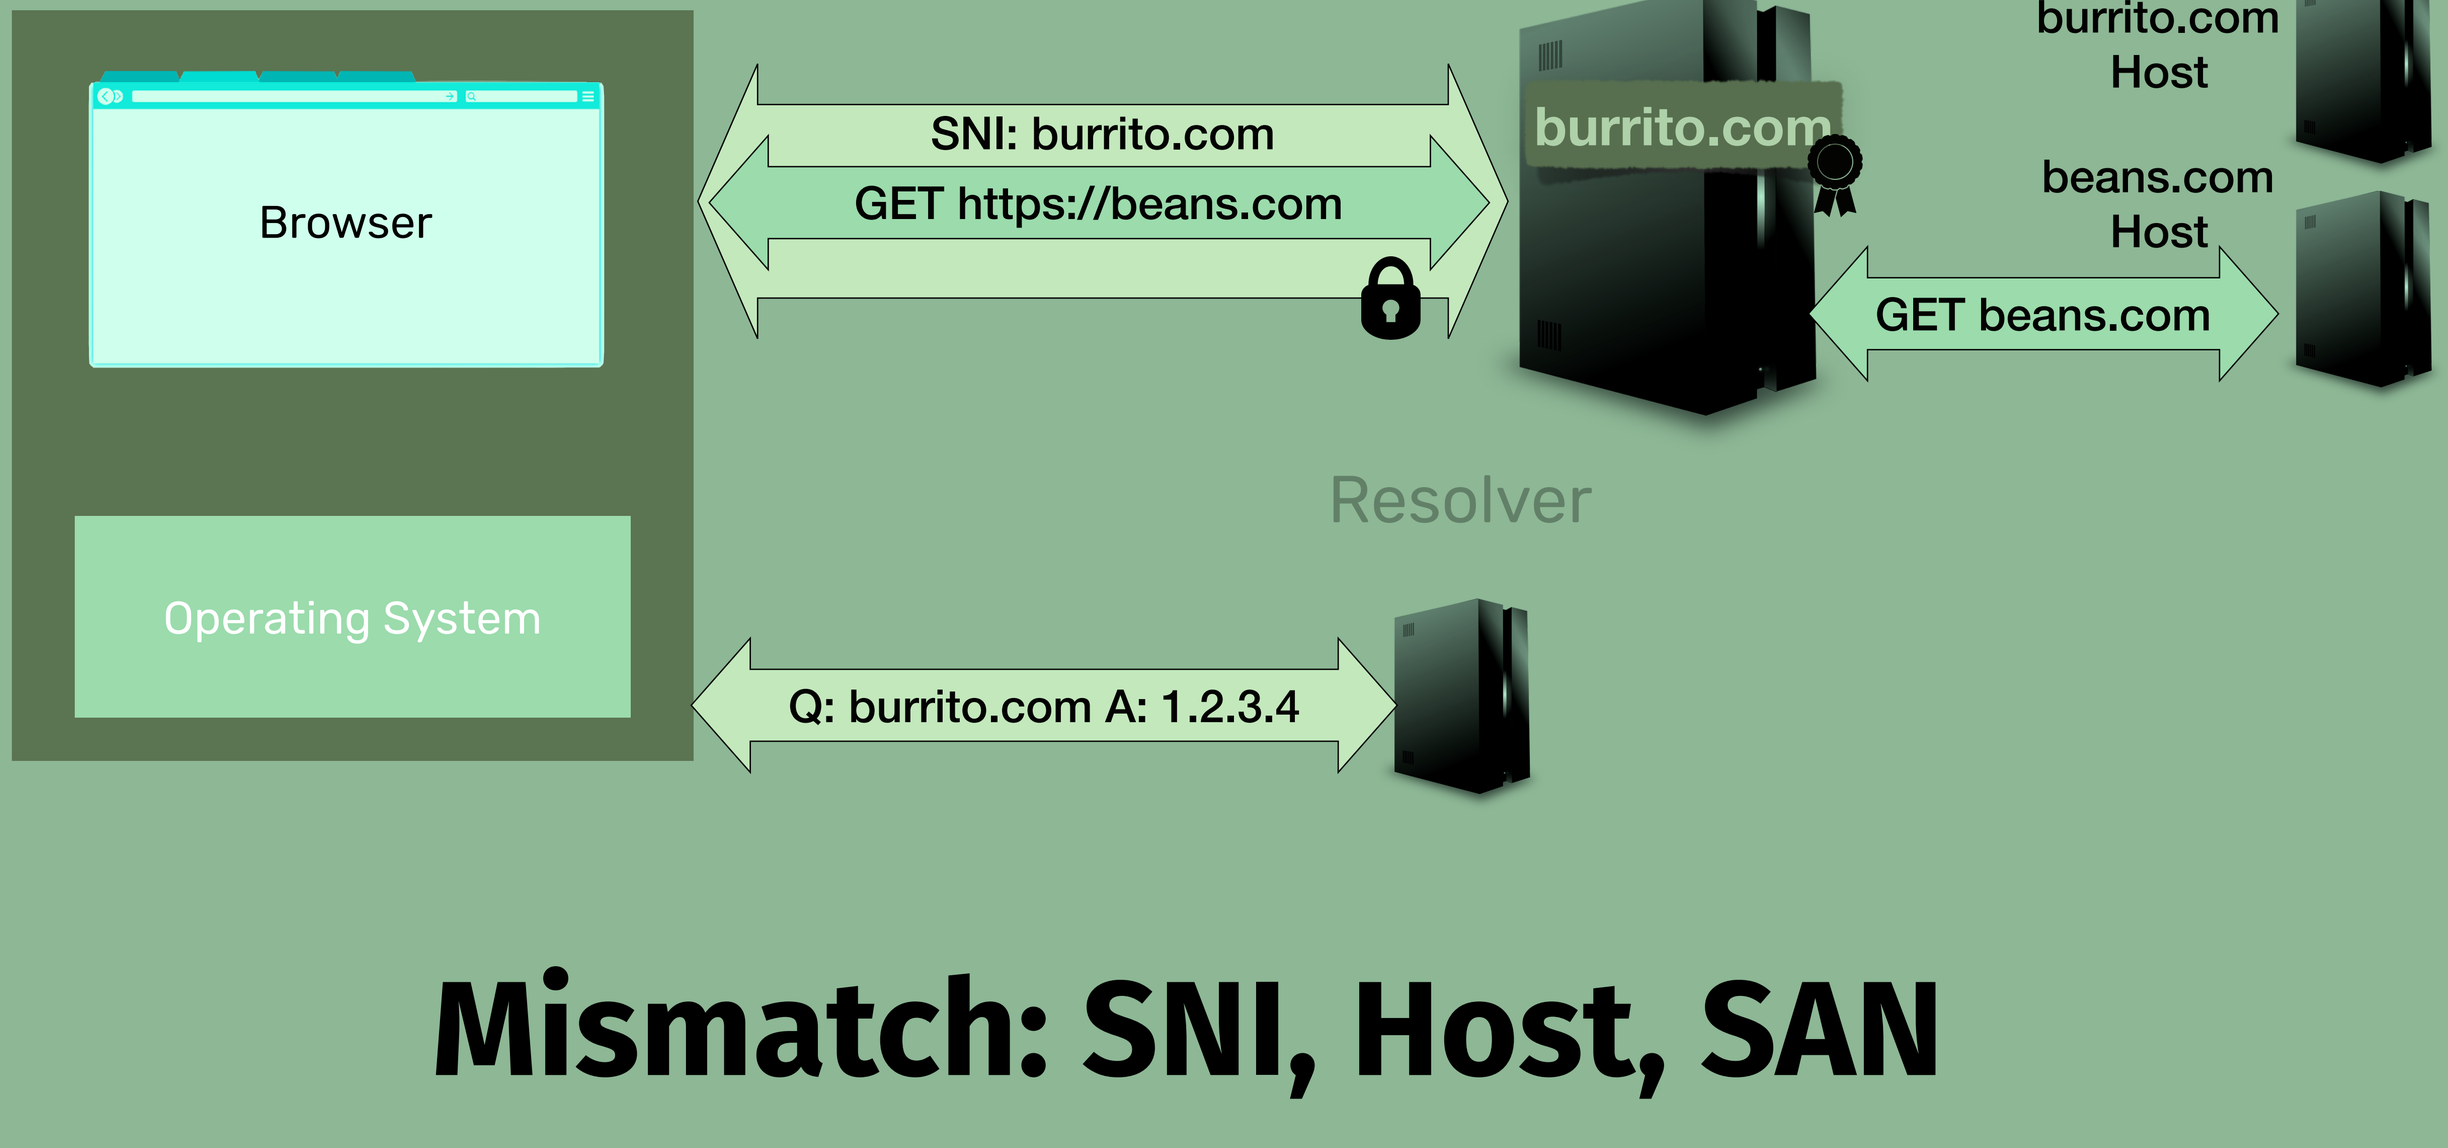

Mismatch: SNI, Host, SAN

Client

HTTP/2

Resolver

Edge

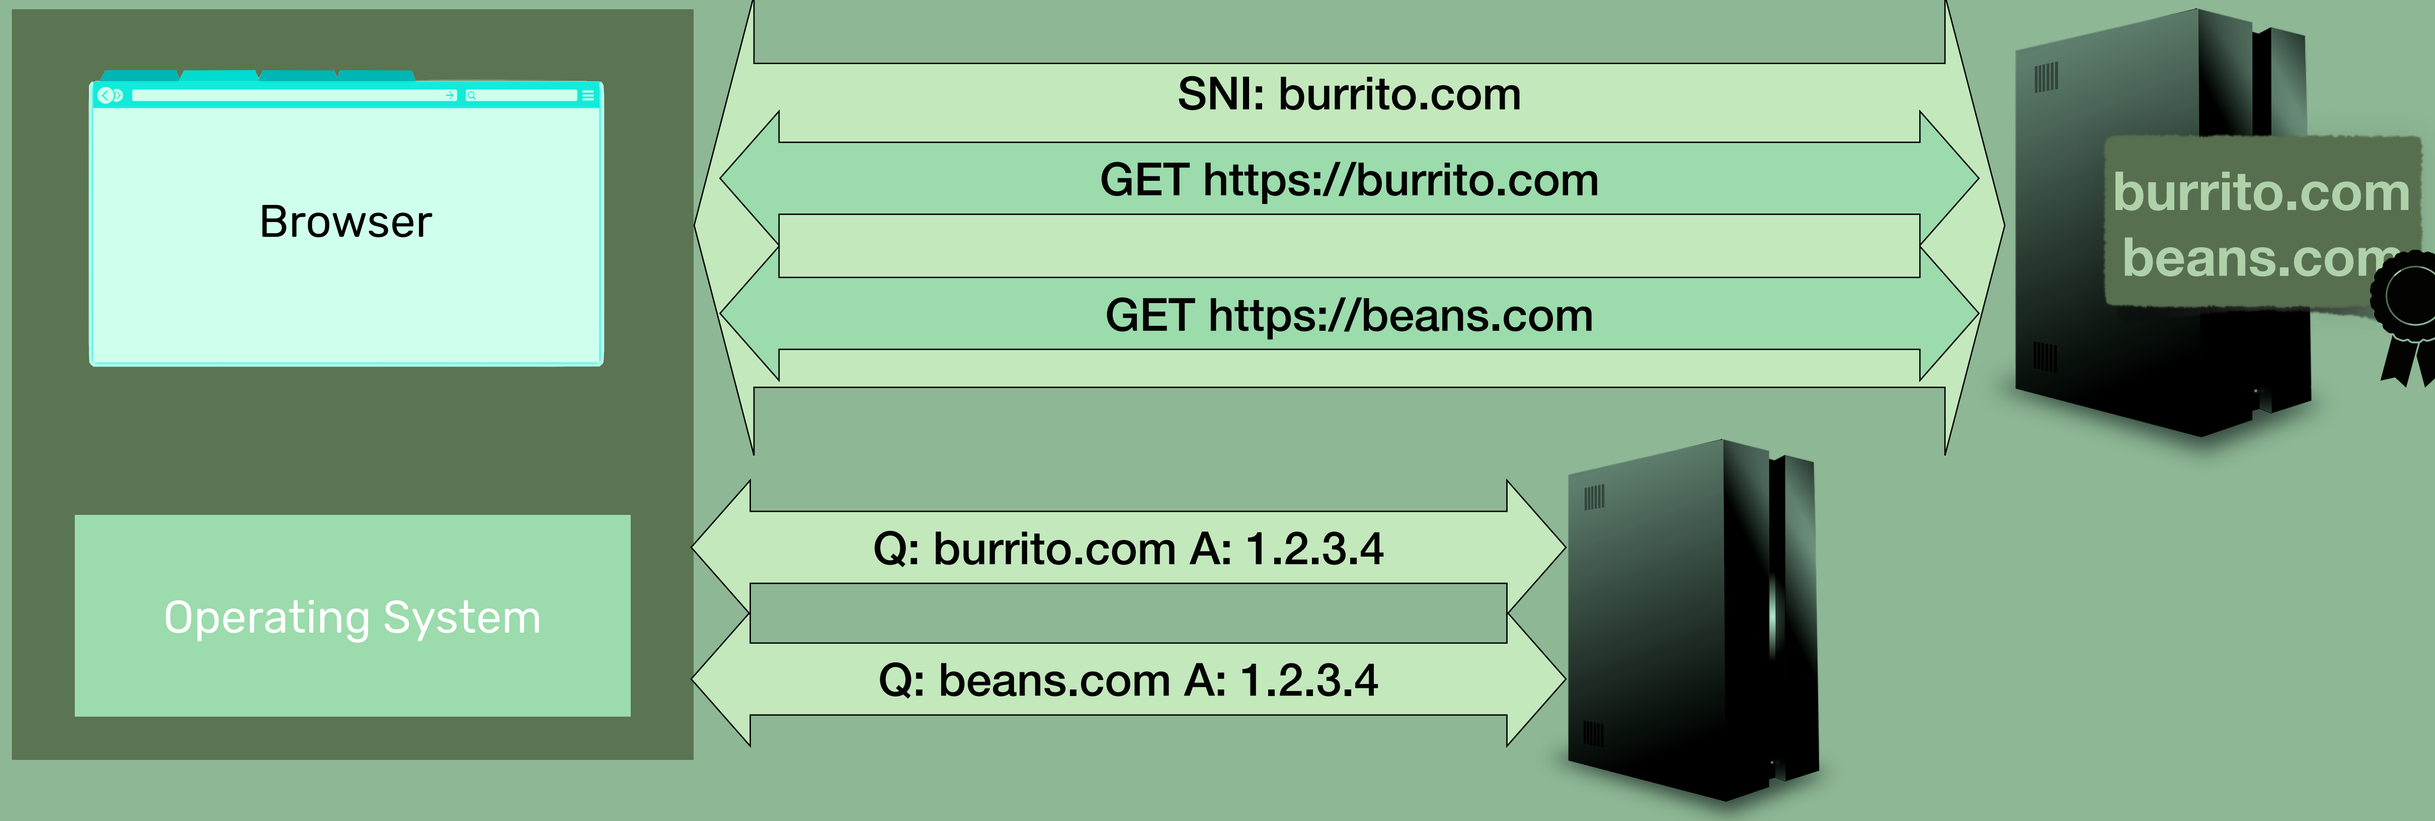

Connection Coalescing

Client

HTTP/2

Resolver

Edge

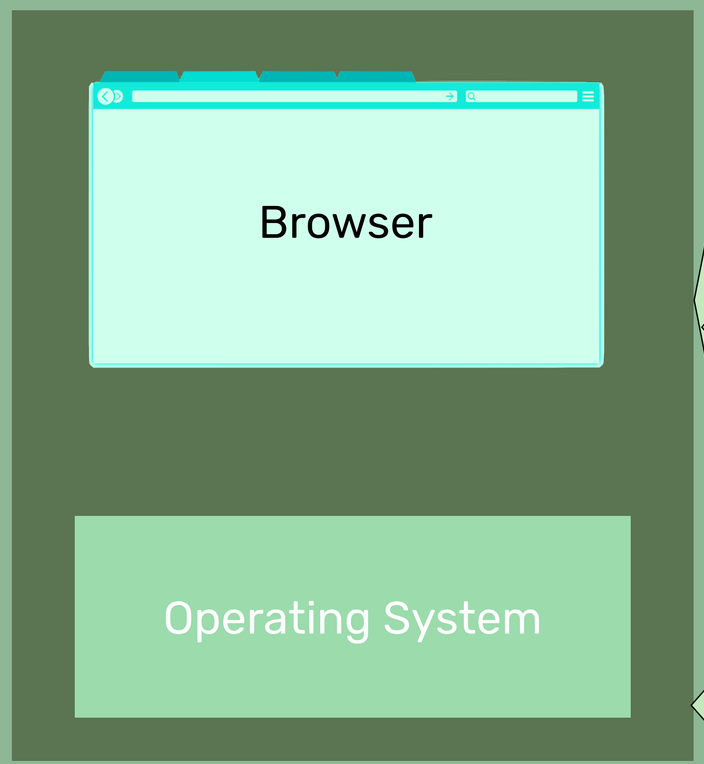

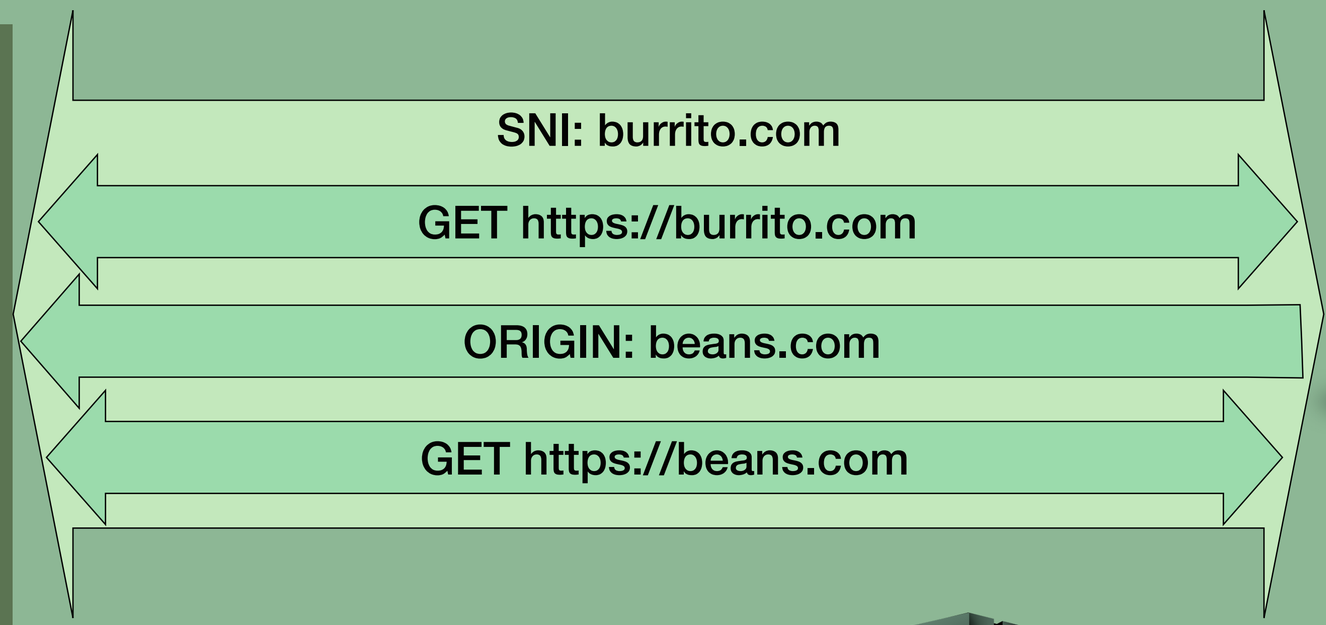

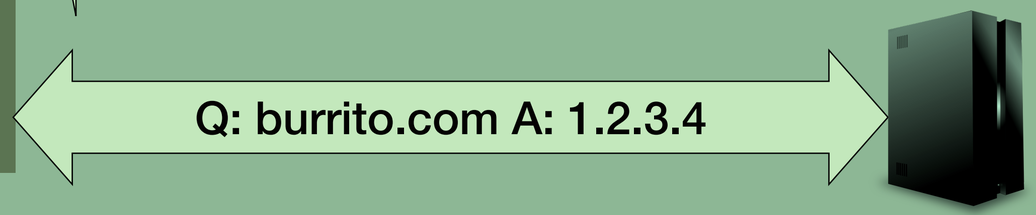

ORIGIN Frame

What a network observer can see

HTTP → HTTPS →

Clients
Hosts
Anycast Hosts

Client Unique IP

Shared Server IP

First Hostname

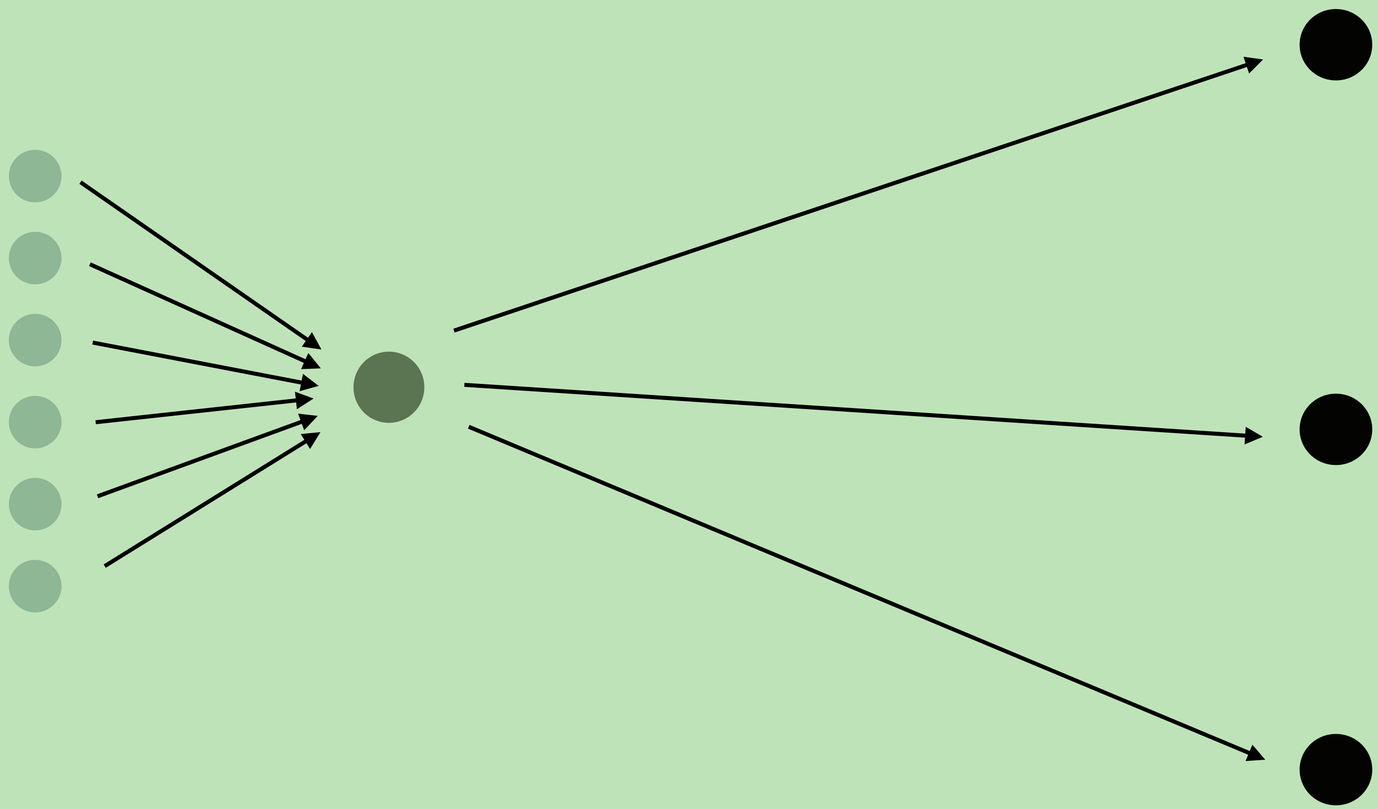

	Anonymity set
Client	**1**
Server (Shared IP+Certificate)	**~20**

Client

HTTP/2

Resolver

Edge

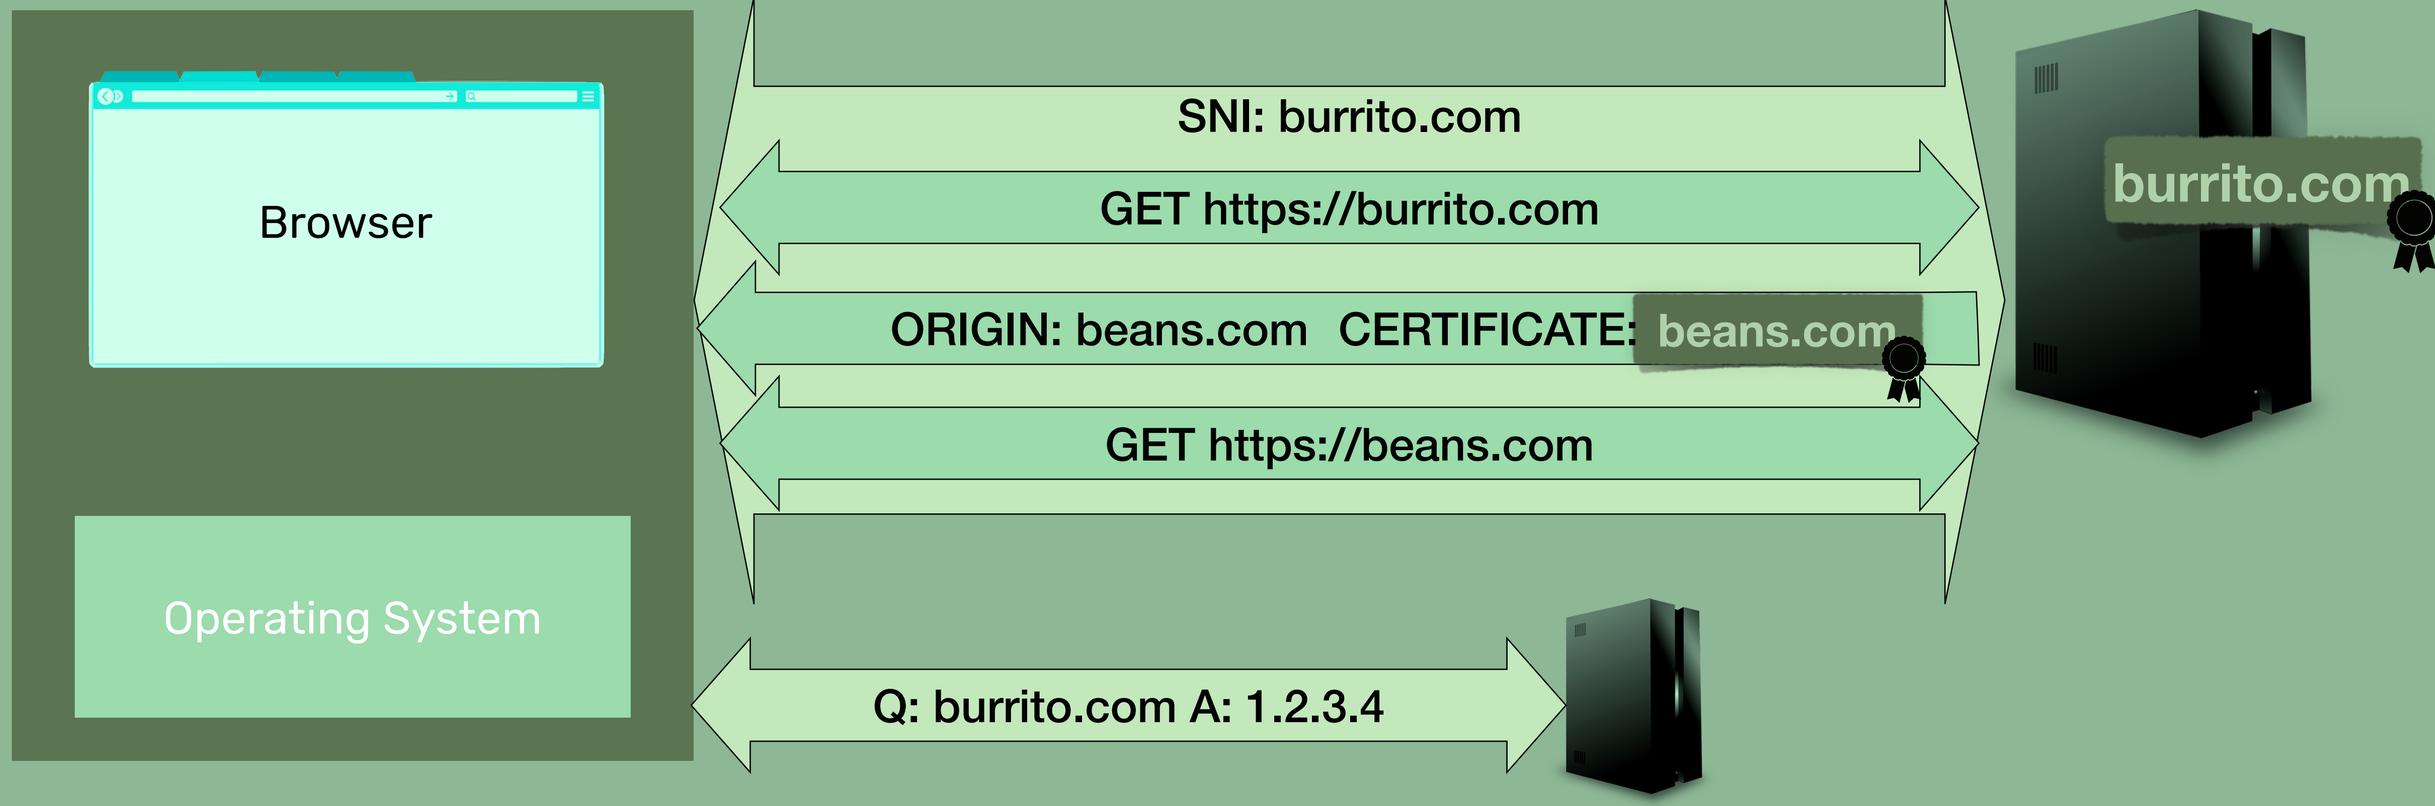

CERTIFICATE Frame

Client

Resolver

Edge

Browser

Operating System

SNI: burrito.com

burrito.com

Q: burrito.com A: 1.2.3.4

What this changes

Having a certificate gives you routing authority

	Anonymity set
Client	**1**
Server (Shared IP+First Hostname)	**k**

k is the set of domains on certificates that can be obtained through “First Hostname”

Meek-like circumvention protection

Only send the CERTIFICATE frame on certain resources

Client

Resolver

Authoritative Server

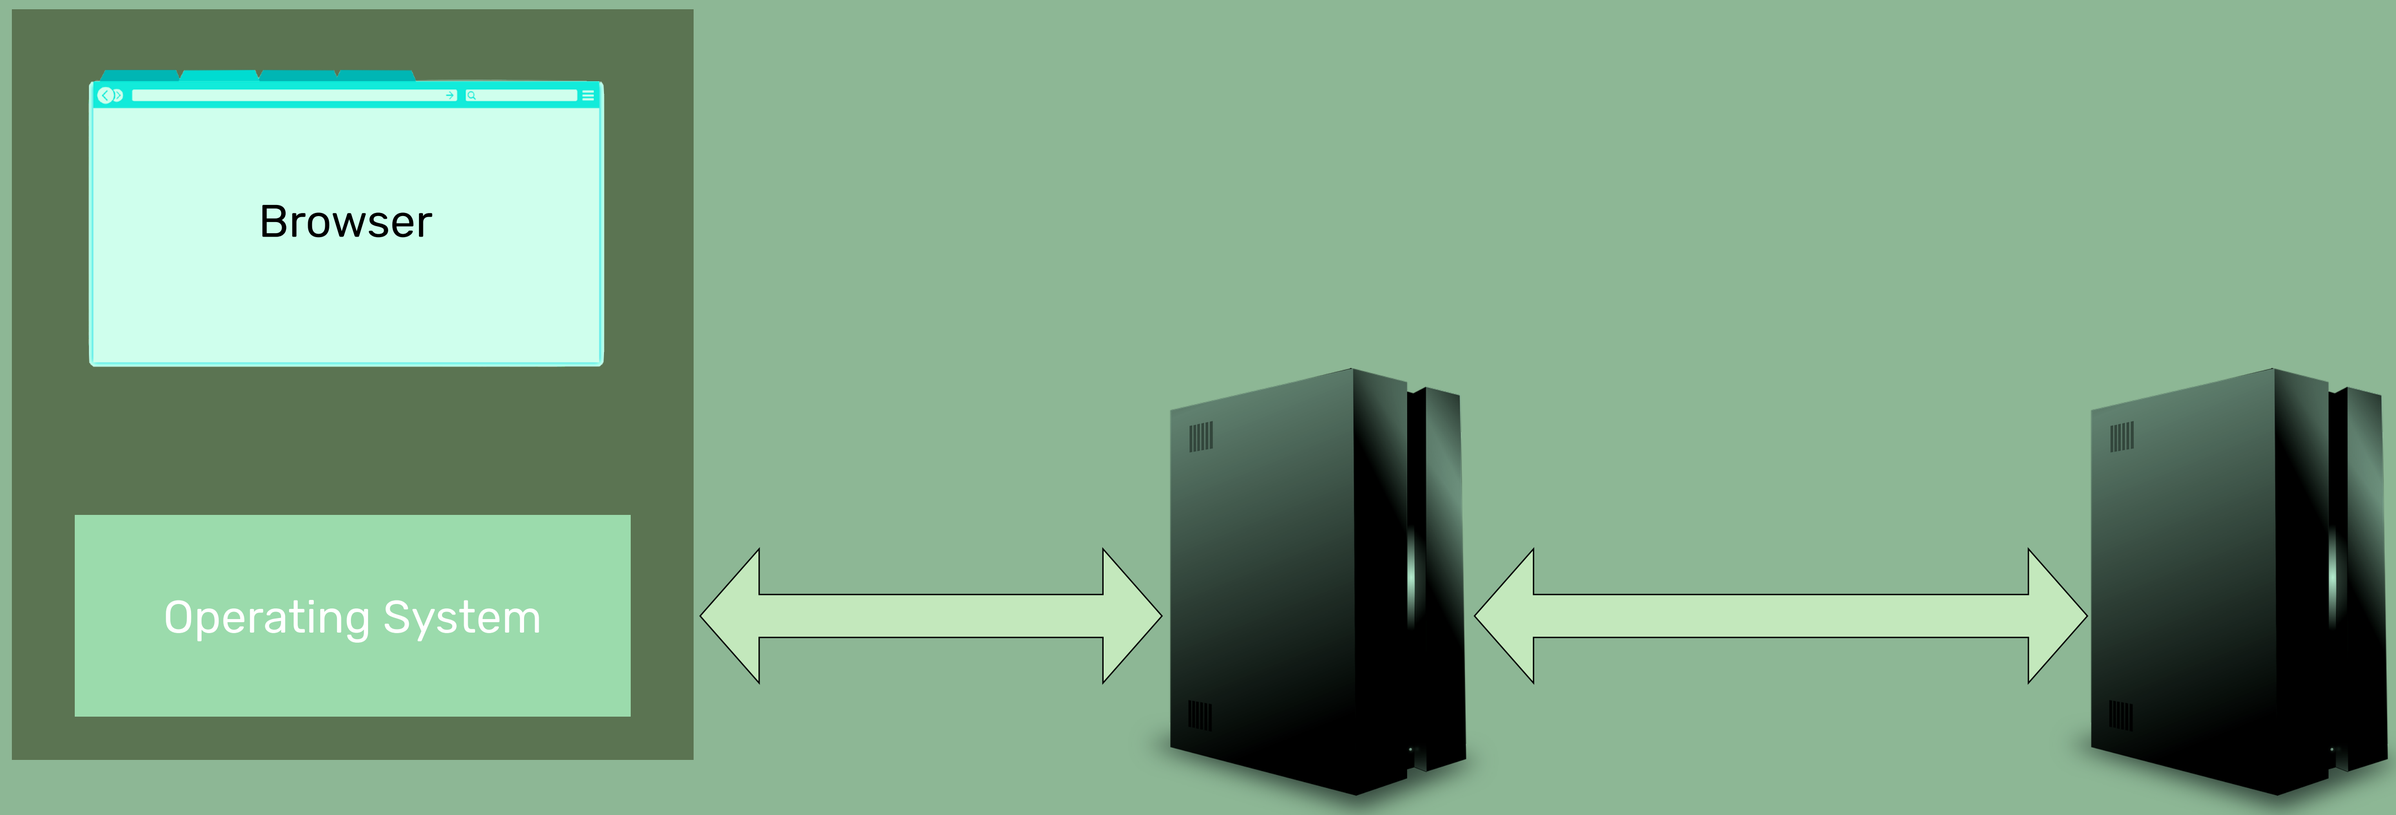

DNS

Client

Resolver

Root Server

Browser

Operating System

me.we.com

TLD Server

me.we.com

Q: e.we.com

Authoritative

me.we.com
Client Subnet

Cache Miss

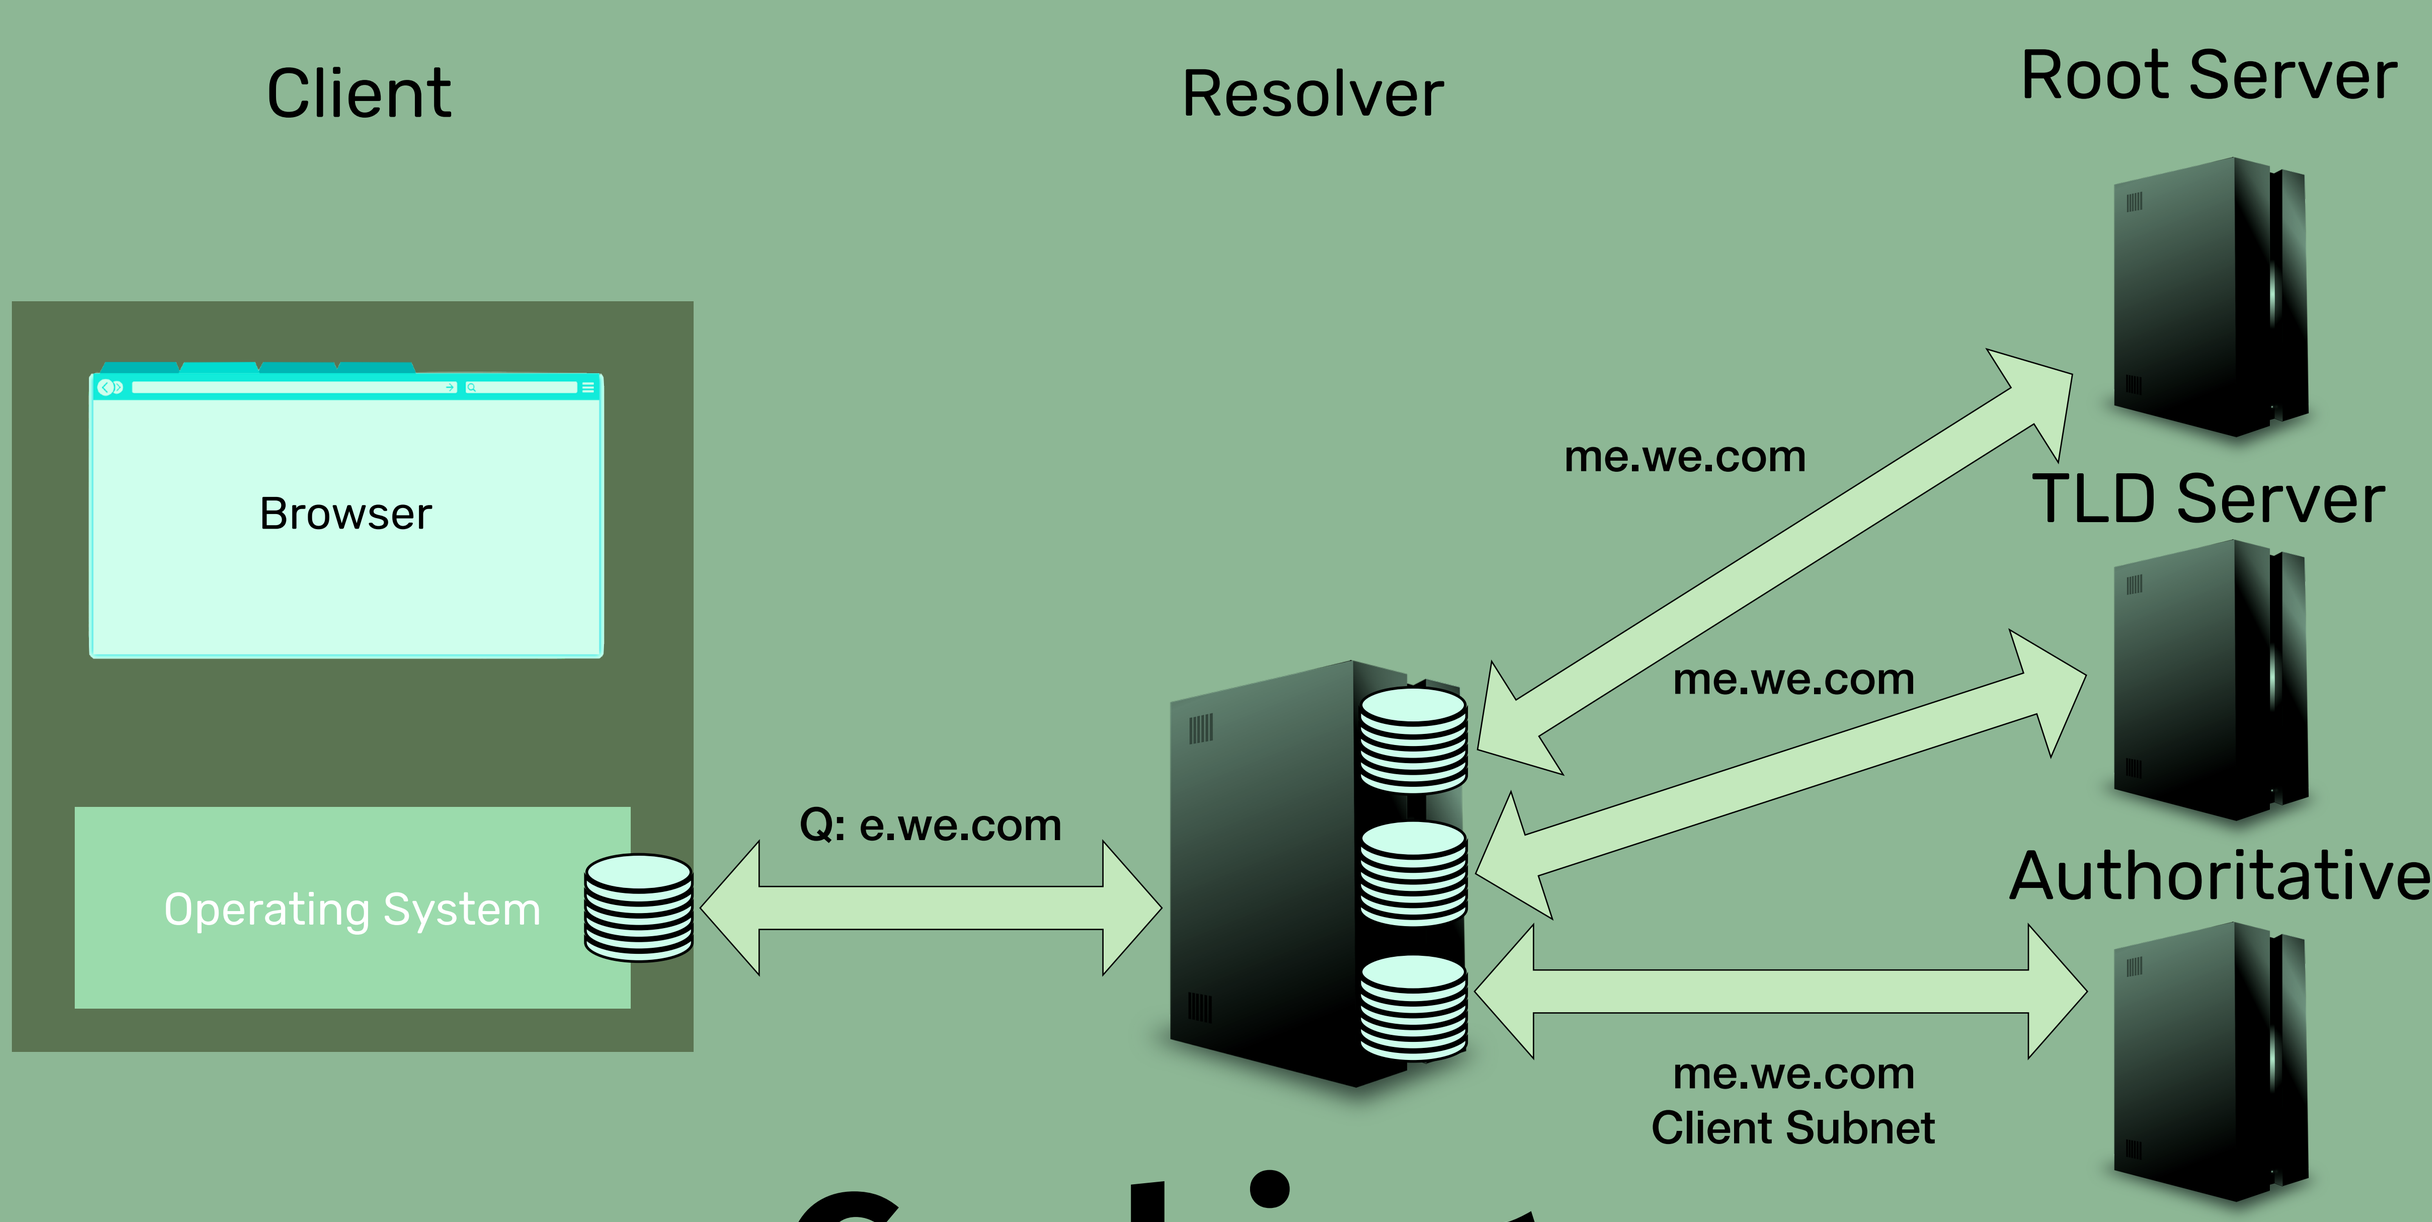
Client
Resolver
Root Server
Browser
Operating System
me.we.com
TLD Server
me.we.com
Q: e.we.com
Authoritative
me.we.com
Client Subnet

Caching

Client

Resolver

Root Server

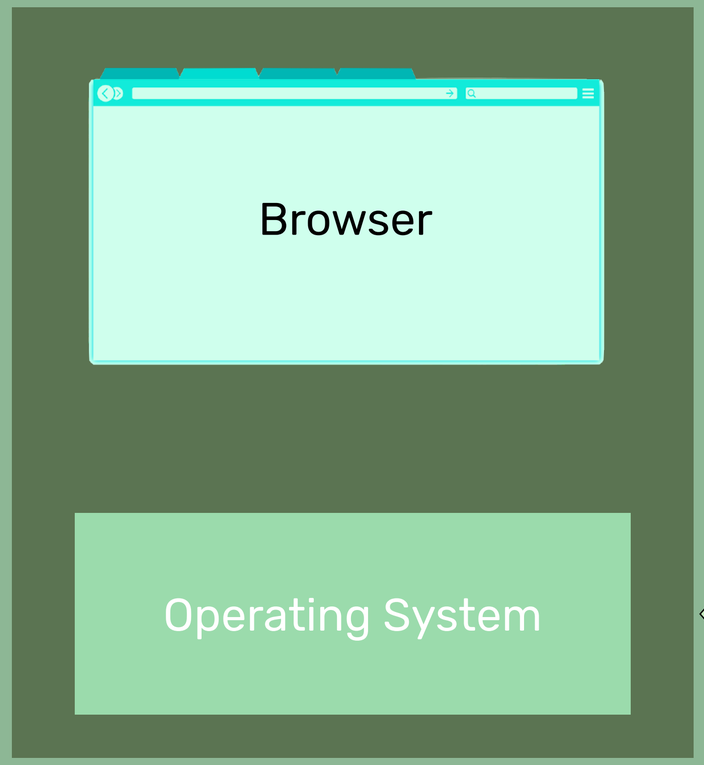

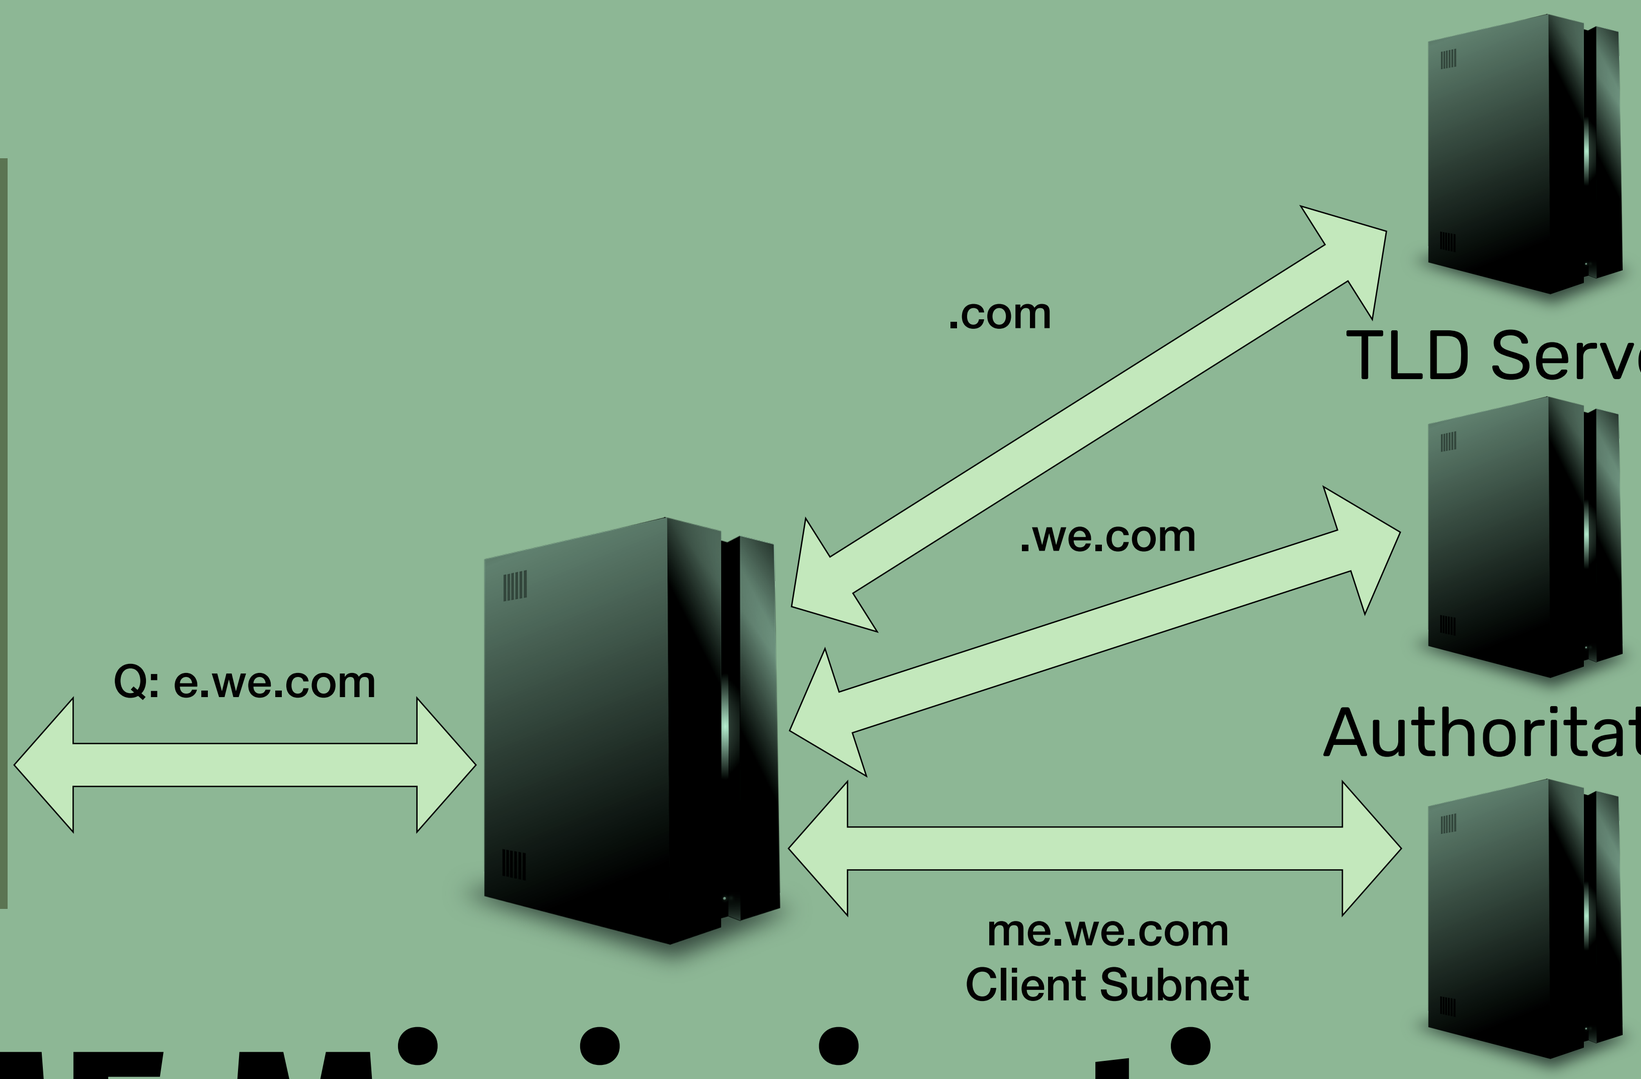

TLD Server

Authoritative

QNAME Minimization

DOH

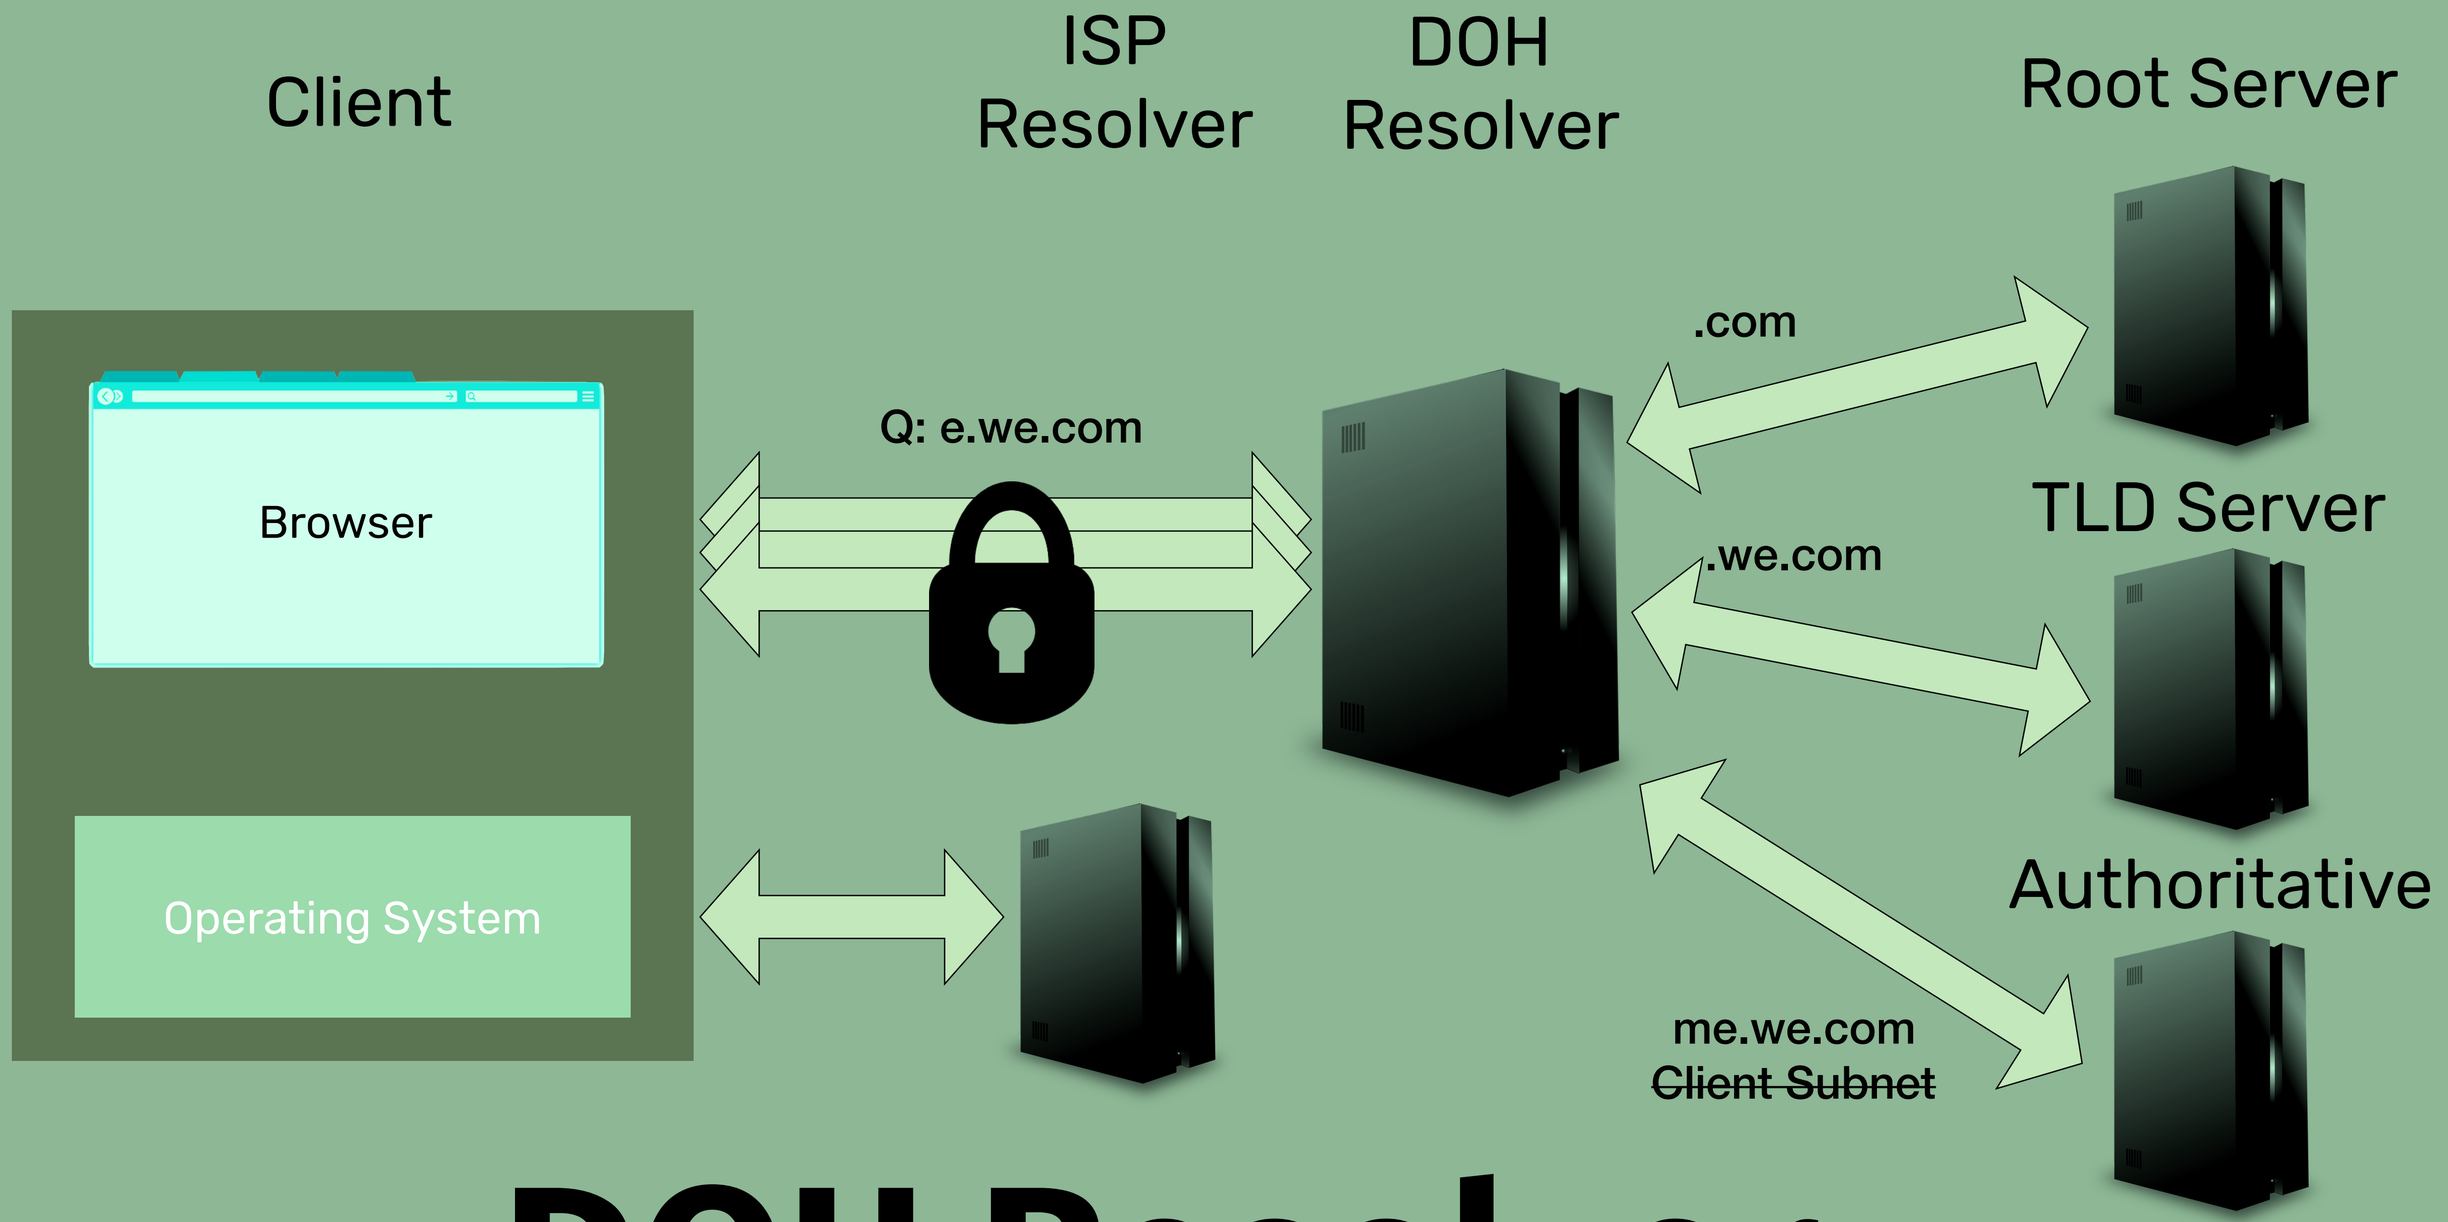

DOH Resolver

Latency

ISPs

Edge DoH

Closer to user
Smaller cache
UDP

Globally Distributed
TLS 1.3 0RTT

Challenges in the Enterprise

HTTP/2

Client

DoH Resolver

Edge

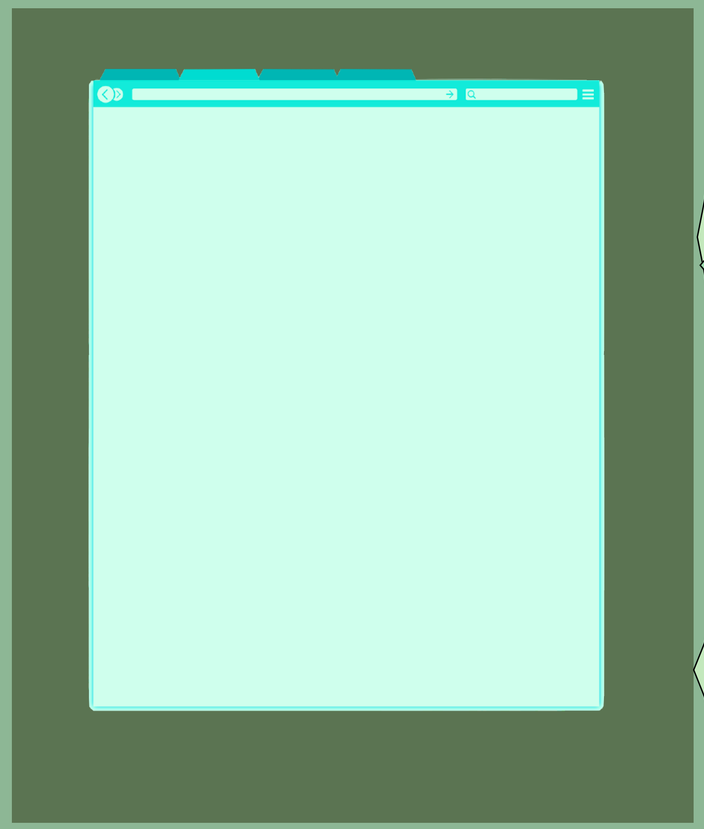

SNI: burrito.com

GET https://burrito.com

ORIGIN: beans.com CERTIFICATE: beans.com

GET https://beans.com

SNI: resolver.com

Q: burrito.com A: 1.2.3.4

ORIGIN/CERT + DoH

Client

HTTP/2

DoH Resolver

Edge

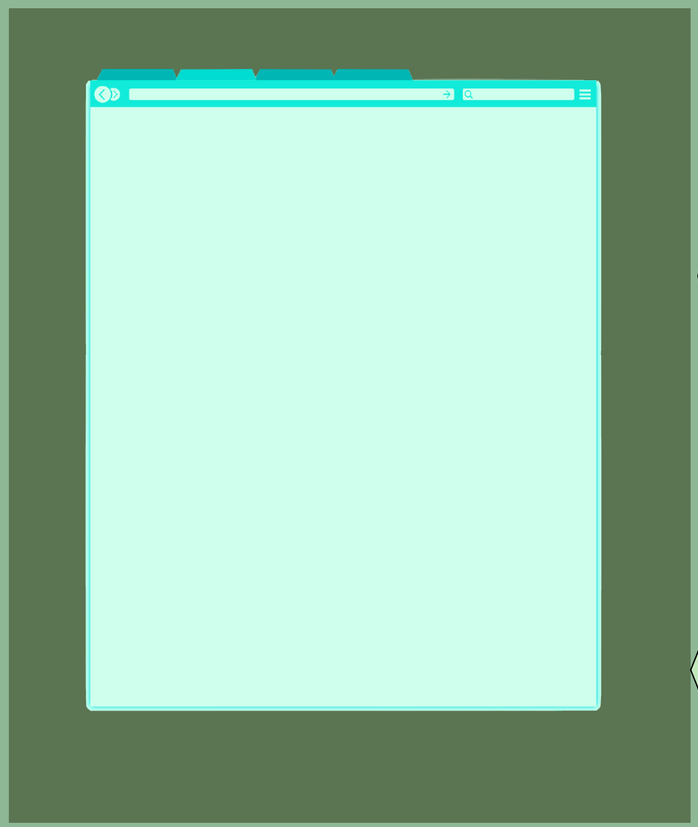

SNI: burrito.com

SNI: resolver.com

ORIGIN/CERT + DoH

SNI

Encryption

Encrypt SNI with client ephemeral key + server public key from DNS

Client

TLS 1.3

DoH Resolver

Edge

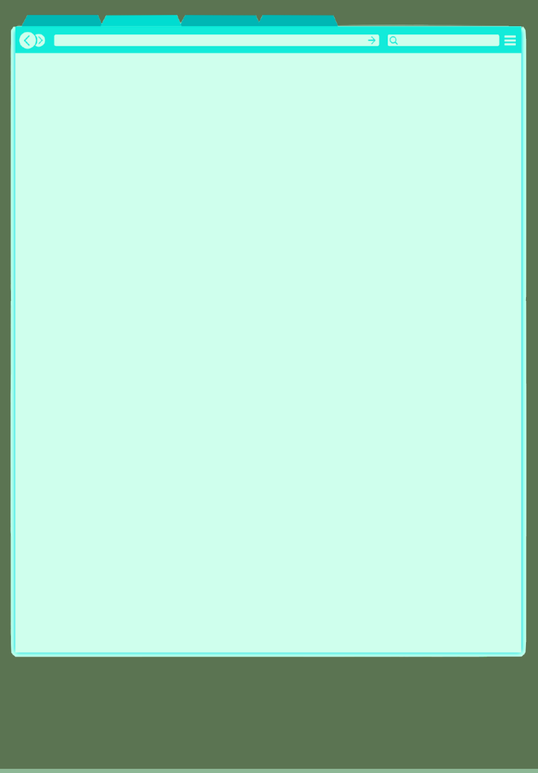

eSNI: E(burrito.com)

SNI: resolver.com

Q: burrito.com A: 1.2.3.4, **PubKey**

O/C + DoH + eSNI

What a network observer can see

HTTP → HTTPS →

Clients
Hosts
Anycast Hosts

Client Unique IP

Shared Server IP

~~First Hostname (SNI)~~

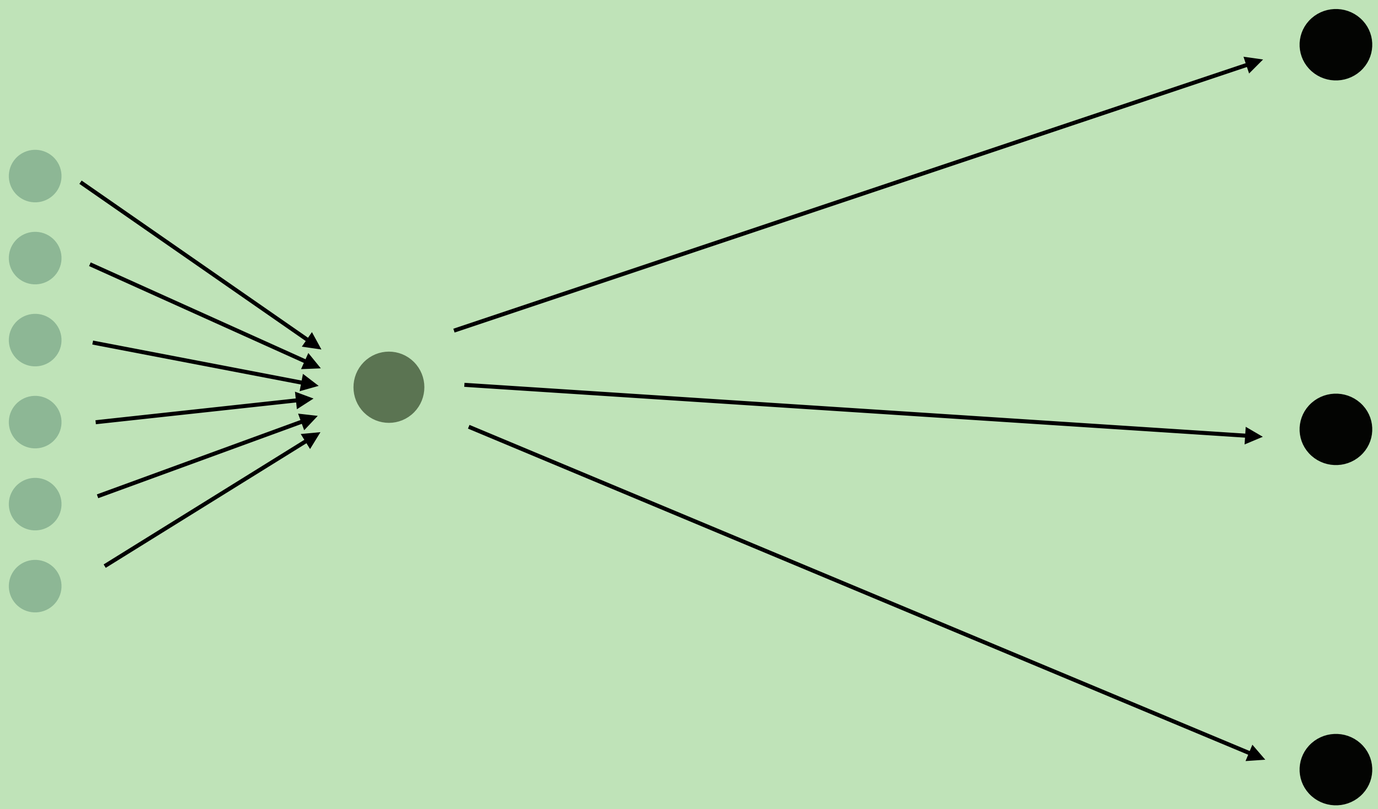

	Anonymity set
Client	**1**
Server	**K**

K is the set of domains that can be served on the IP

Caveat: If Server IP is static, then this give a hint about first hostname.

Client

HTTP/2

Resolver

Edge

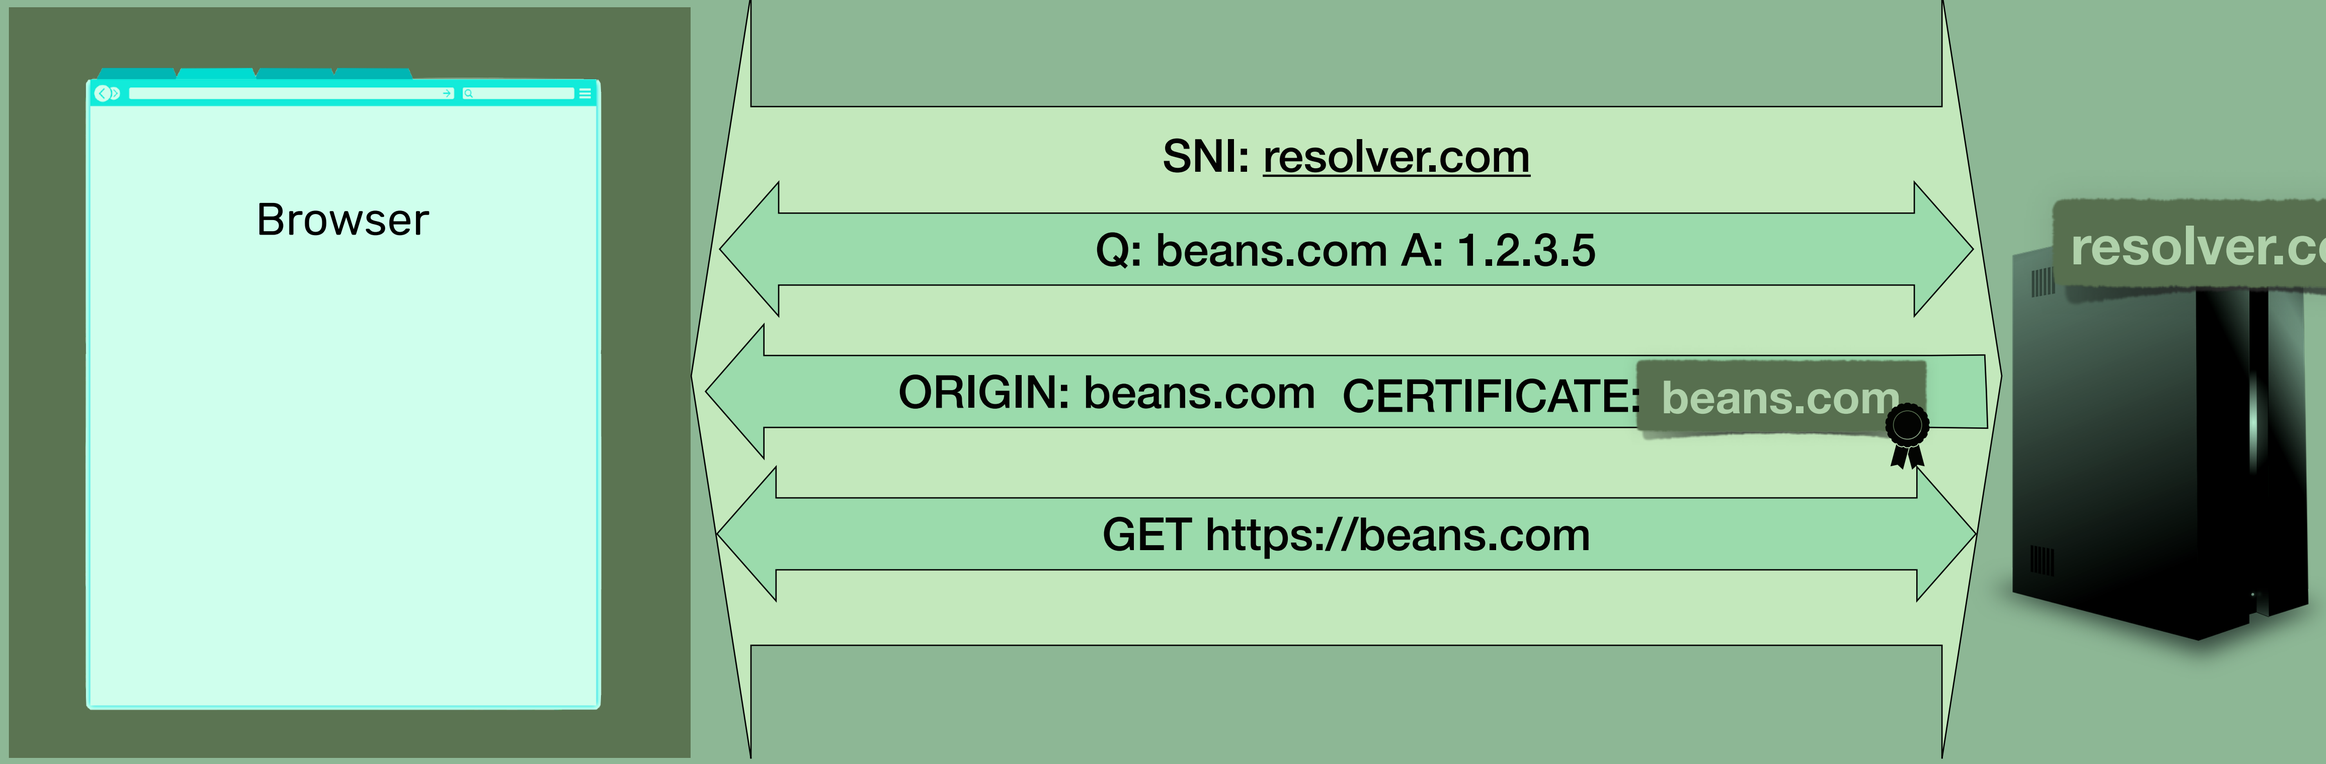

DOH “VPN”

Client

HTTP/2

Resolver

Edge

Browser

SNI: resolver.com

resolver.com

DOH “VPN”

	Anonymity set
Client IP	**1**
Server IP	**K**

K is the set of domains that can be served on the IP

No dynamic IP requirement

Where are we now?

ORIGIN implemented in Firefox

CERTIFICATE being standardized by IETF

DOH supported by Google DNS, 1.1.1.1

eSNI about to be submitted to IETF

ORIGIN

Privacy improvement limited by shared certs

Latency skip both DNS and HTTPS

Security certificate compromise risk

CERTIFICATE

Privacy hide any bean in any burrito

Latency extends origin benefits to any cert

Security exchange DNS for CT or OCSP stapling

DOH

Privacy first hop improvement

Latency depends on provider, TLS 1.3

Security security against attacks, allows passive DNS

eSNI

Privacy first domain privacy given dynamic IPs

Latency depends on DoH for reliability

Security risk of more MiTM

Open Questions

How much privacy does this actually give people?

Does this incentivize further consolidation?

Does increased performance and privacy outweigh the legitimate need for external visibility?

Website Fingerprinting

Removing explicit signals does not protect you from passive ones

Consolidation

Better performance when using a popular provider

Is visibility necessary?

Safety vs. Security

The Evolving Architecture of the Web

Nick Sullivan

CLOUDFLARE