LocalRoot++

Making LocalRoot the Default

Disclaimers:

- Joint work with Wes Hardaker, Geoff Huston, Jim Reid, but input from many many others.
- Ongoing discussions in the IETF DNSOP working group.
  - Things are subject to change

Ignaz Semmelweis

Wikimedia Commons/Semmelweis Museum of Medical History

(not) Ignaz Semmelweis

Image created with Gemini AI

A refresher on how resolution works...

How do I reach
www.a-random-domain.net?

Webserver

www.a-random-domain.net

A refresher on how resolution works...

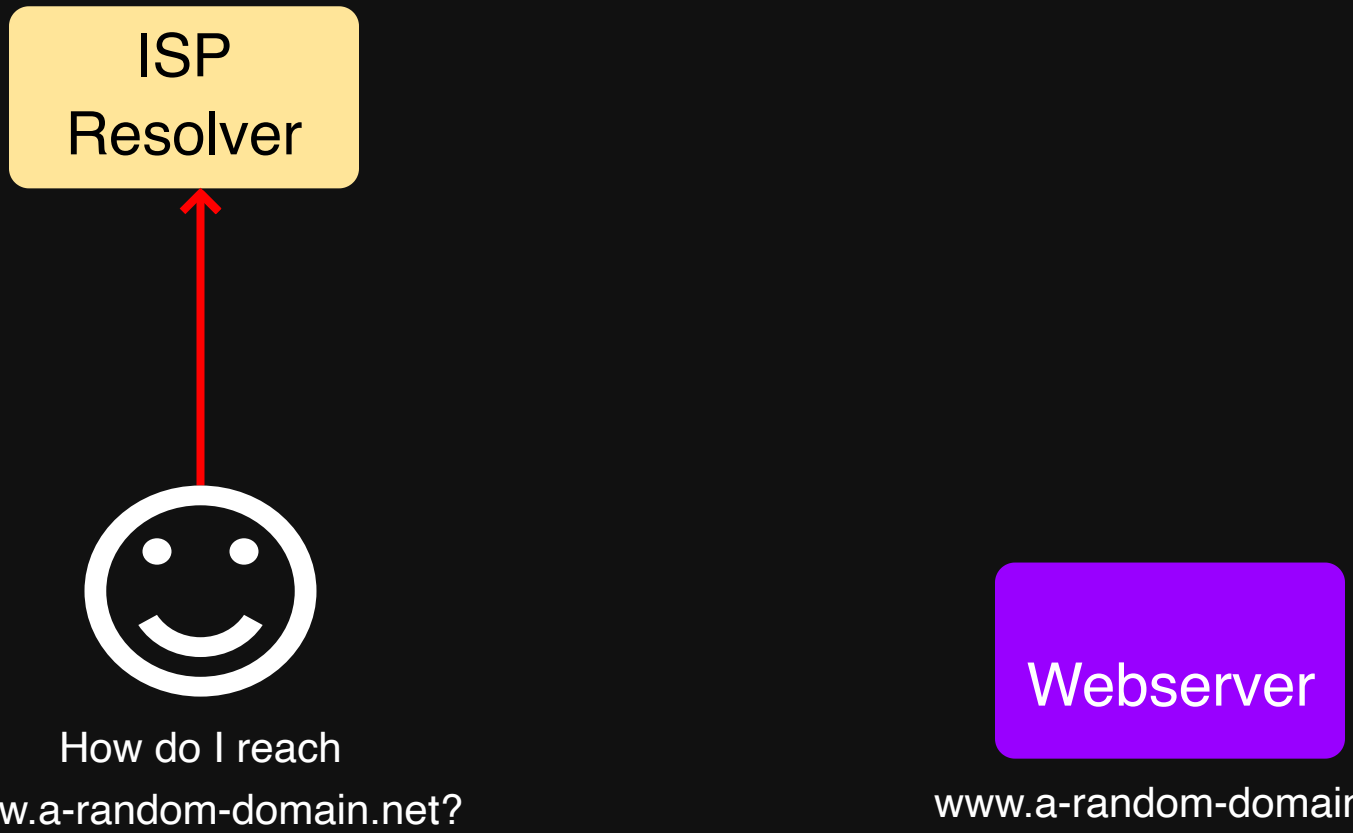

A refresher on how resolution works...

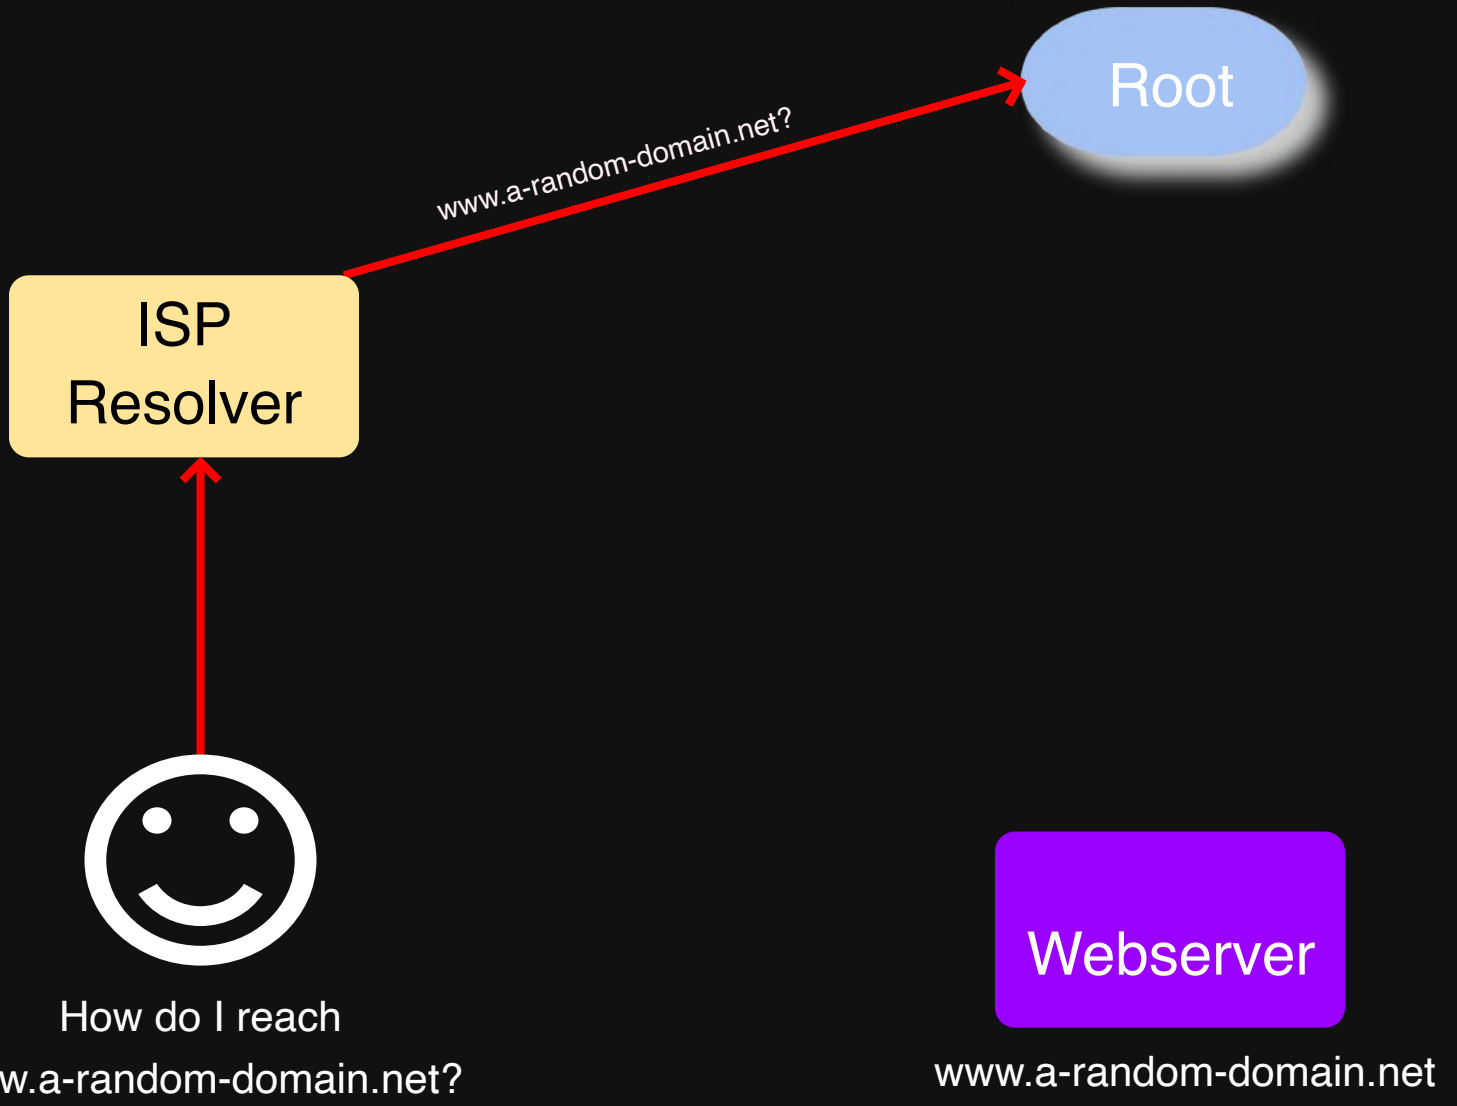

A refresher on how resolution works...

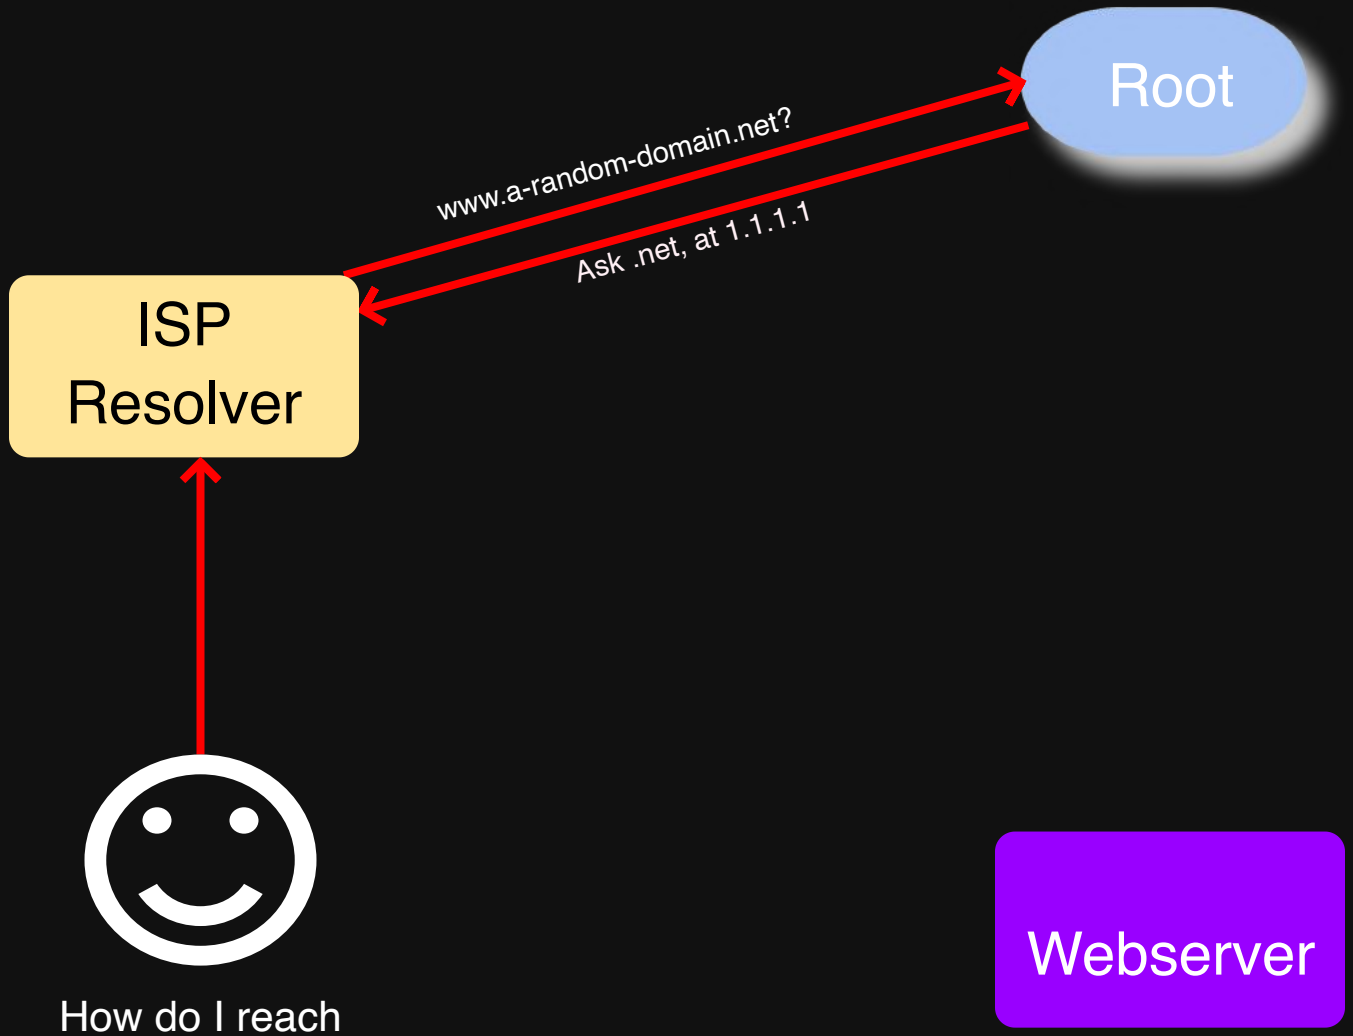

A refresher on how resolution works...

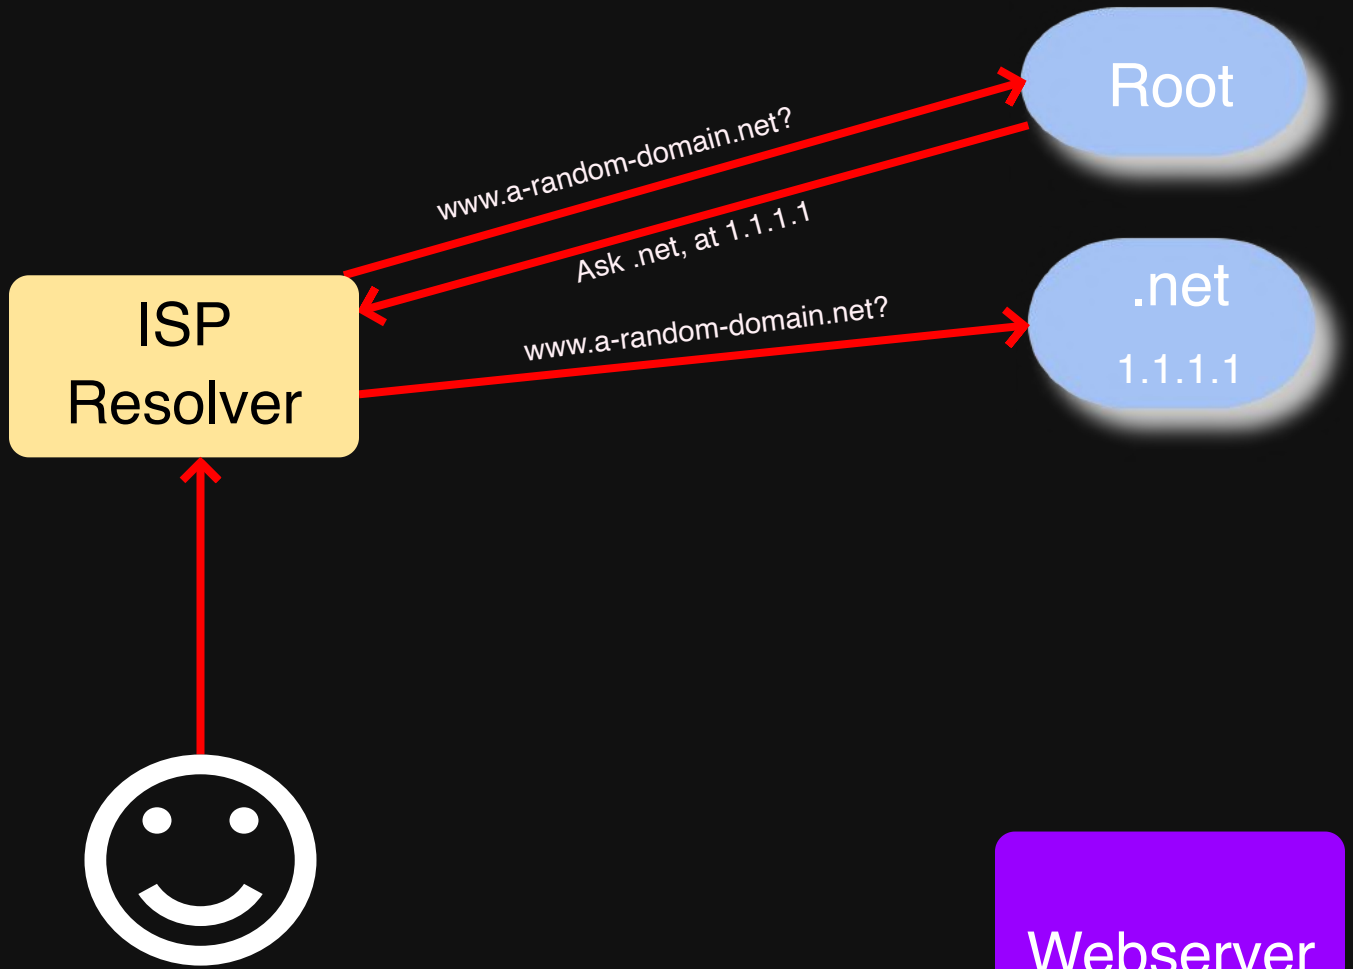

A refresher on how resolution works...

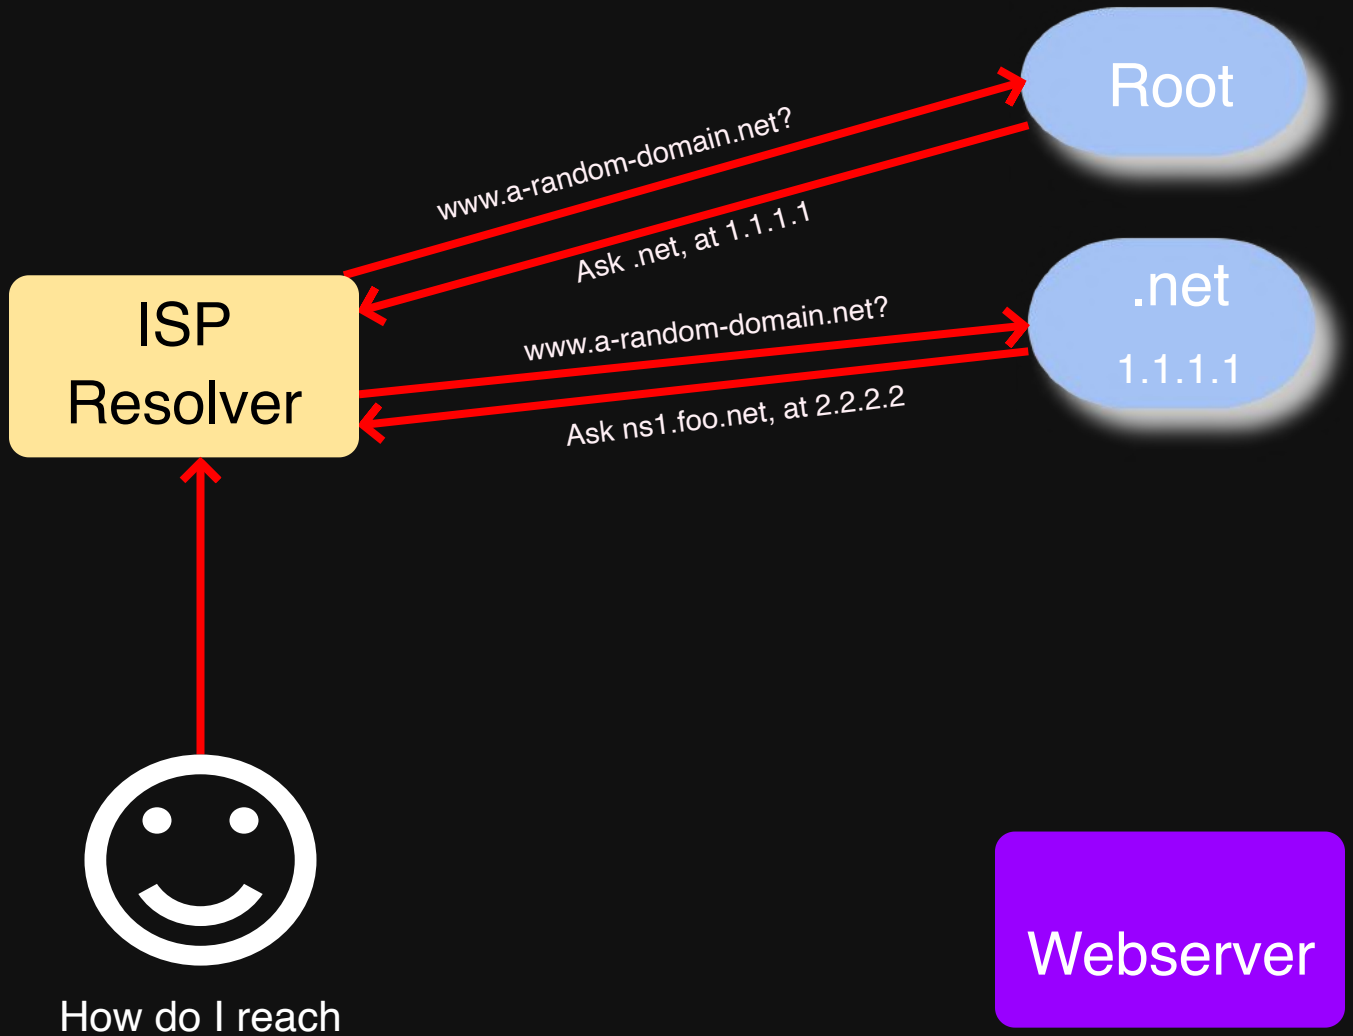

A refresher on how resolution works...

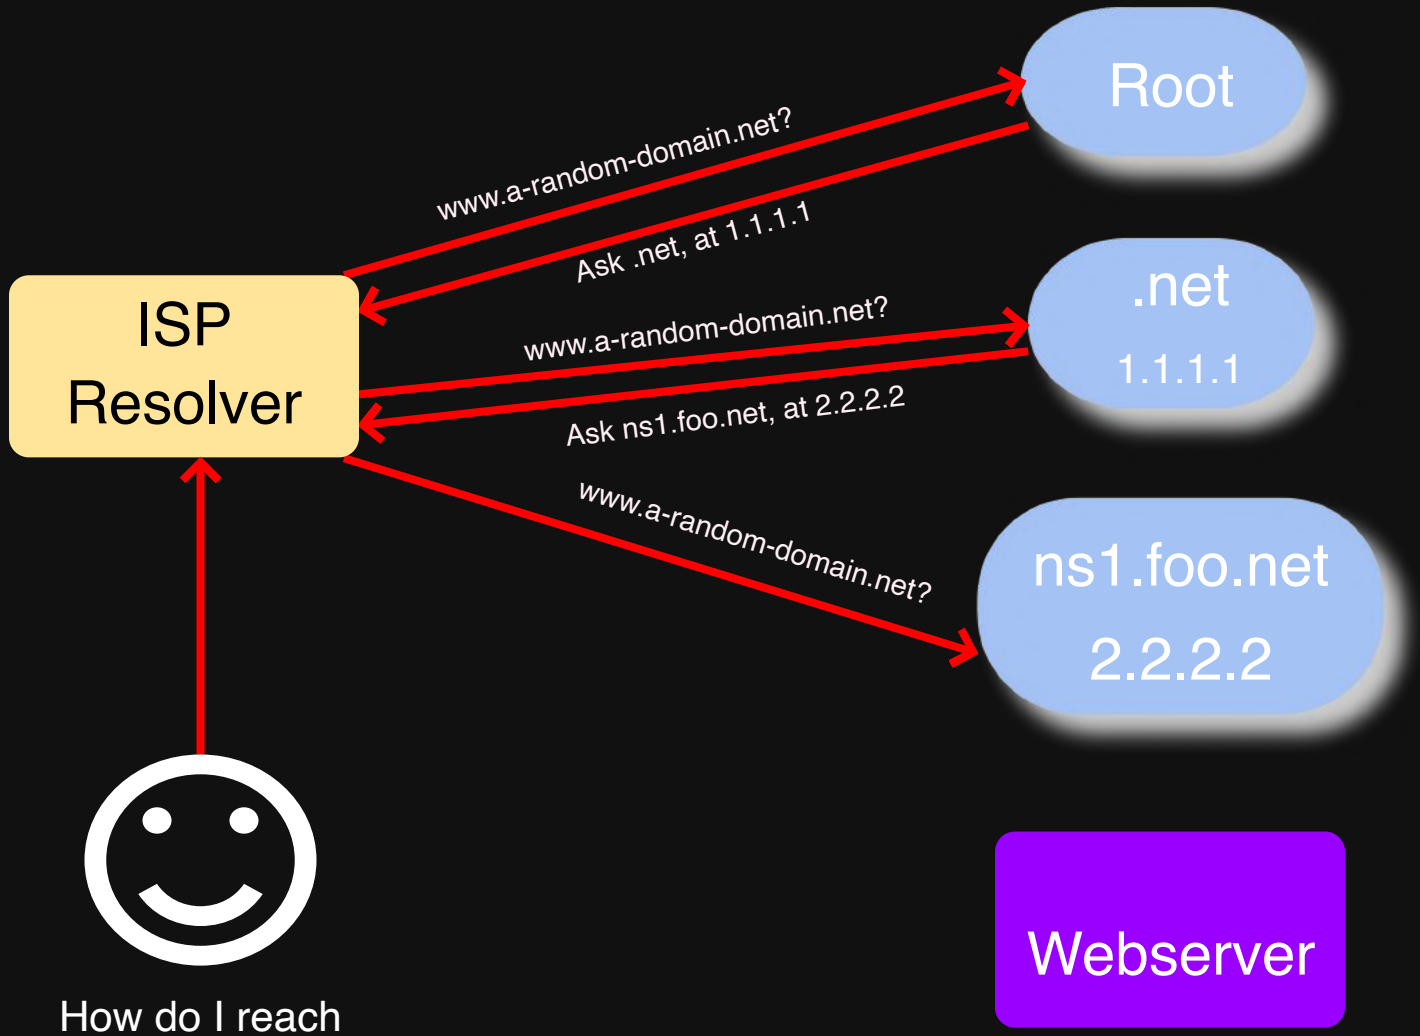

A refresher on how resolution works...

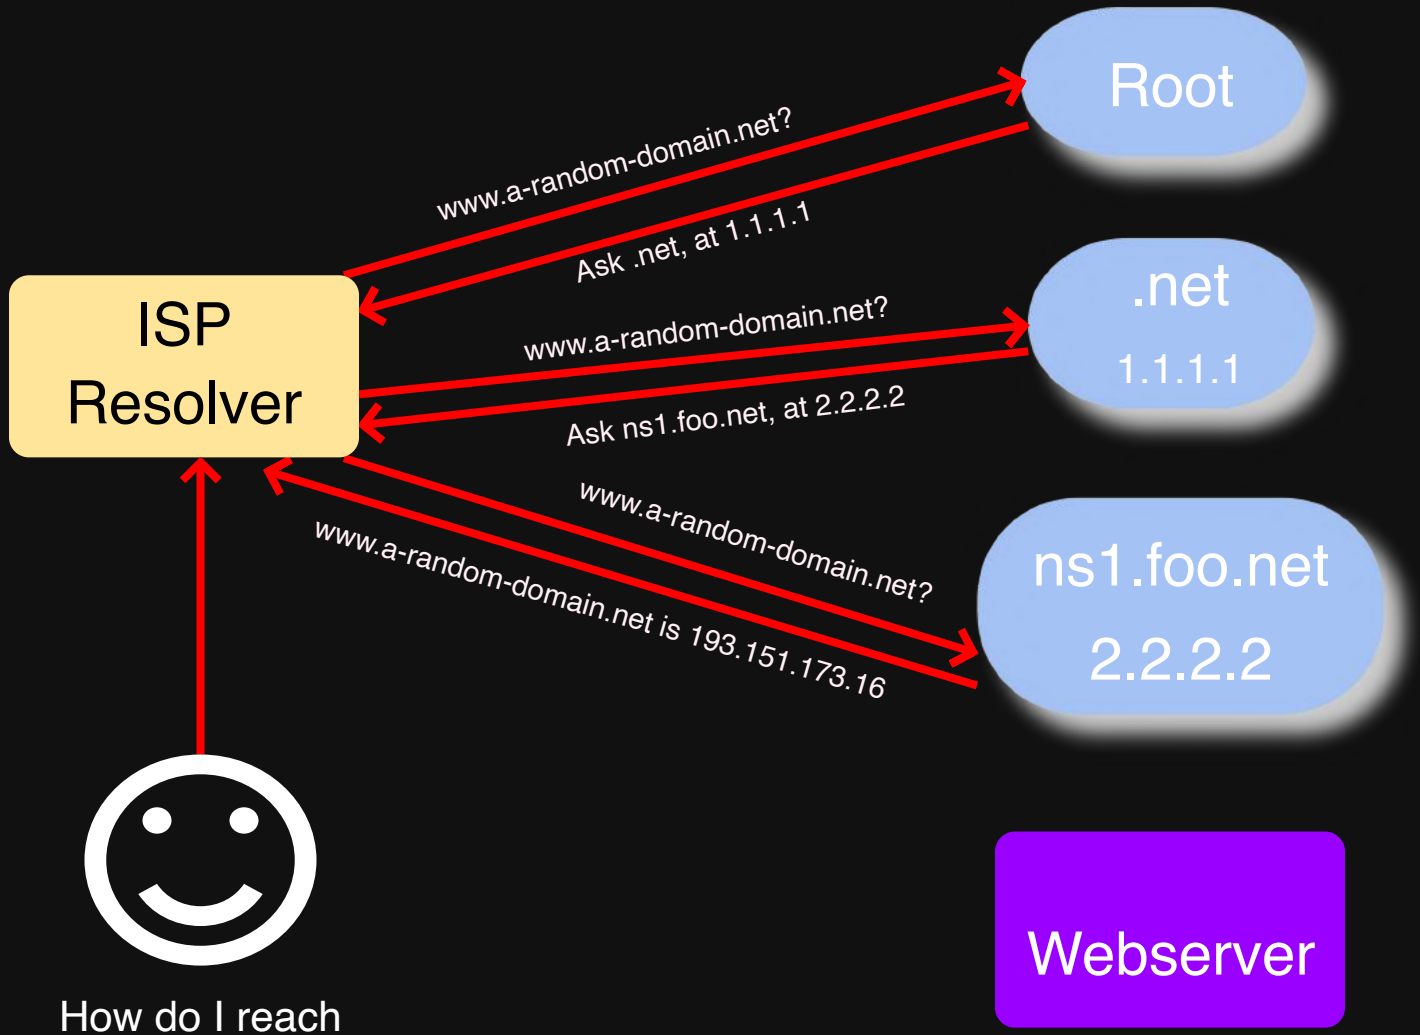

A refresher on how resolution works...

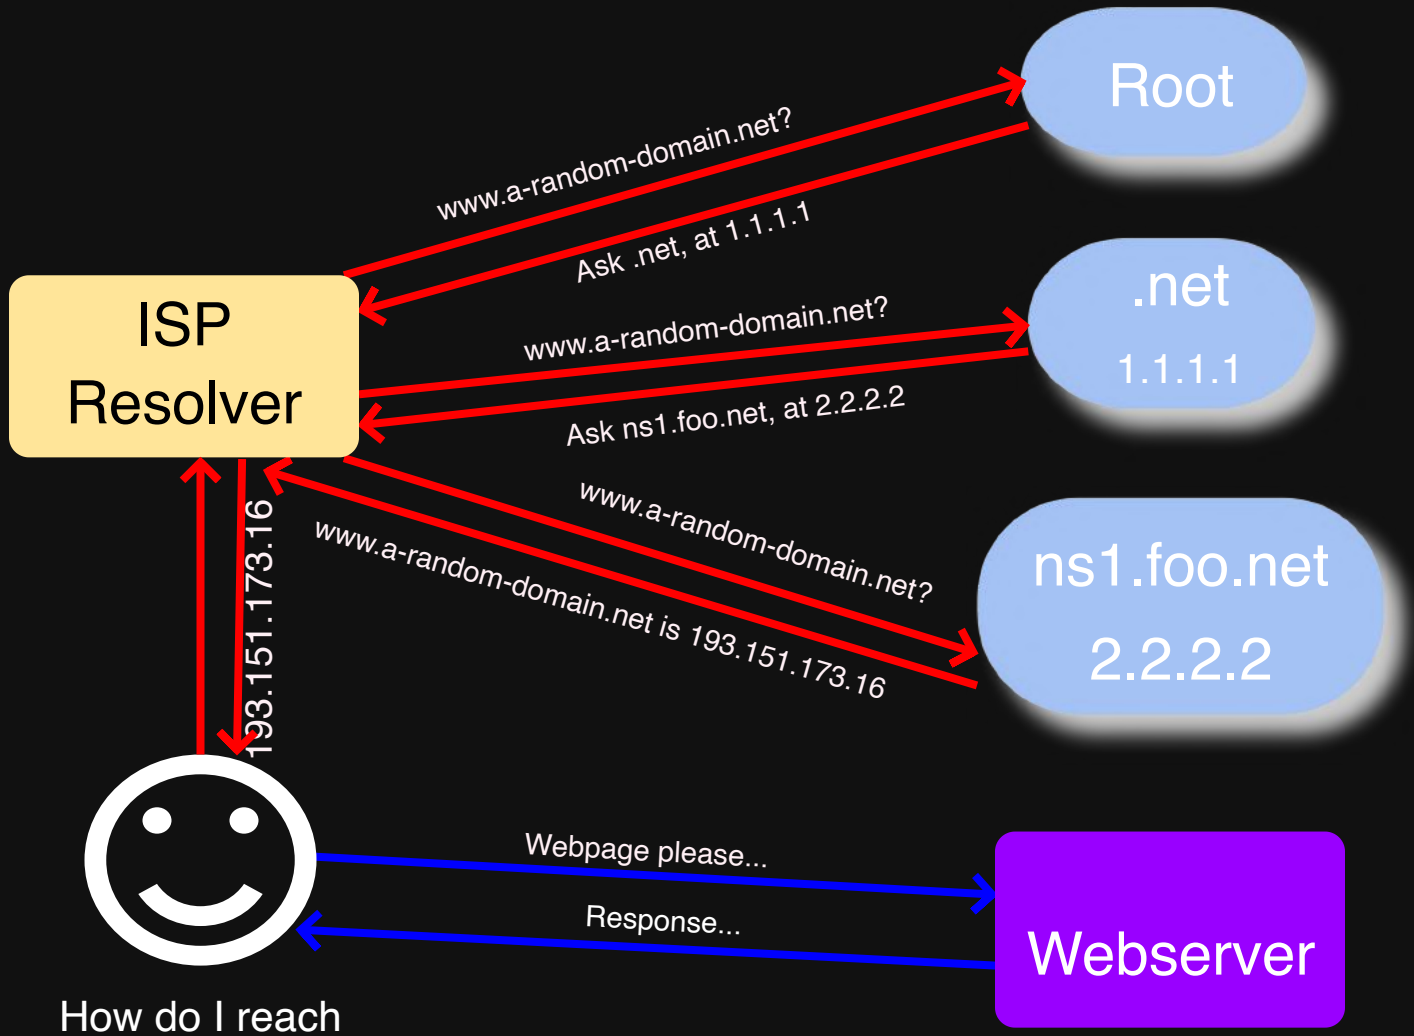

A refresher on how resolution works...

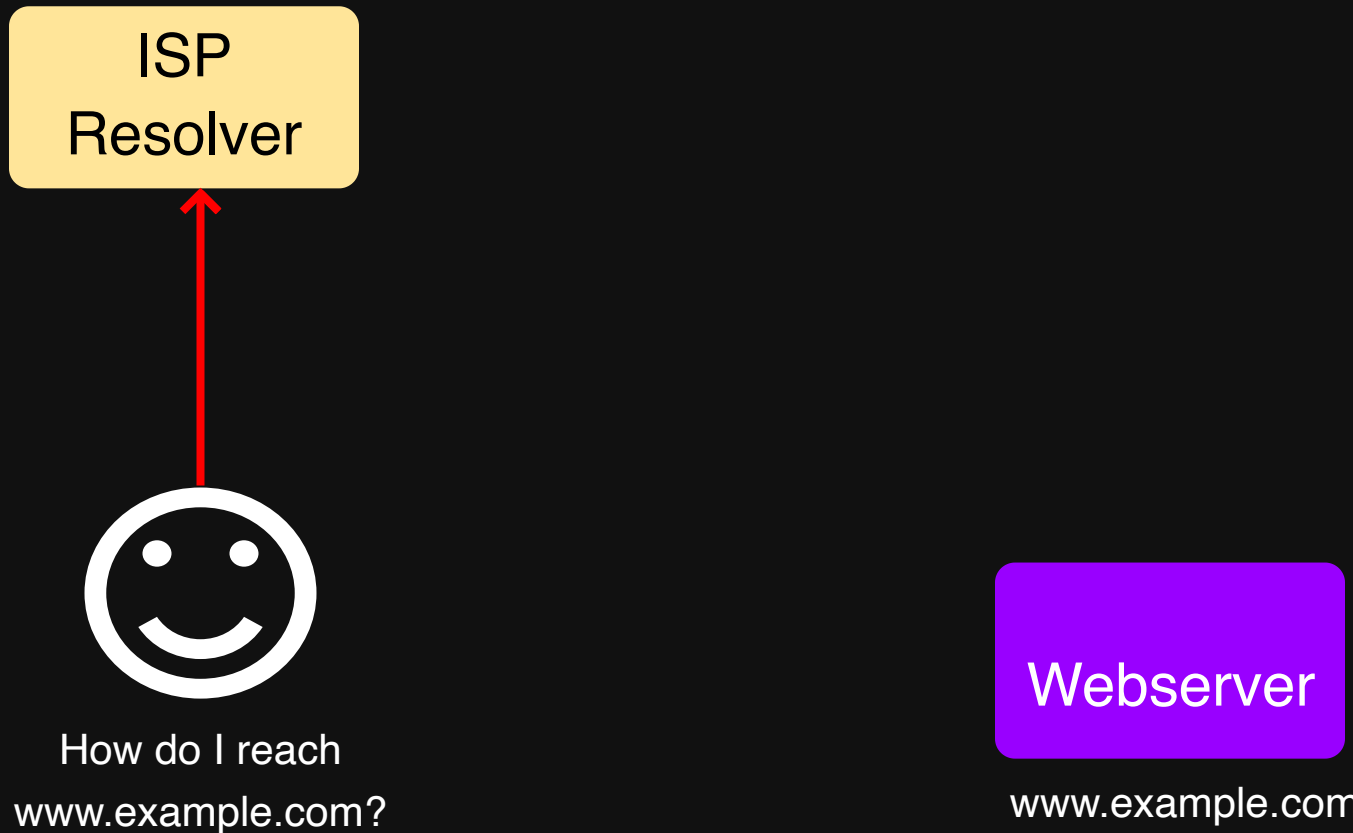

A refresher on how resolution works...

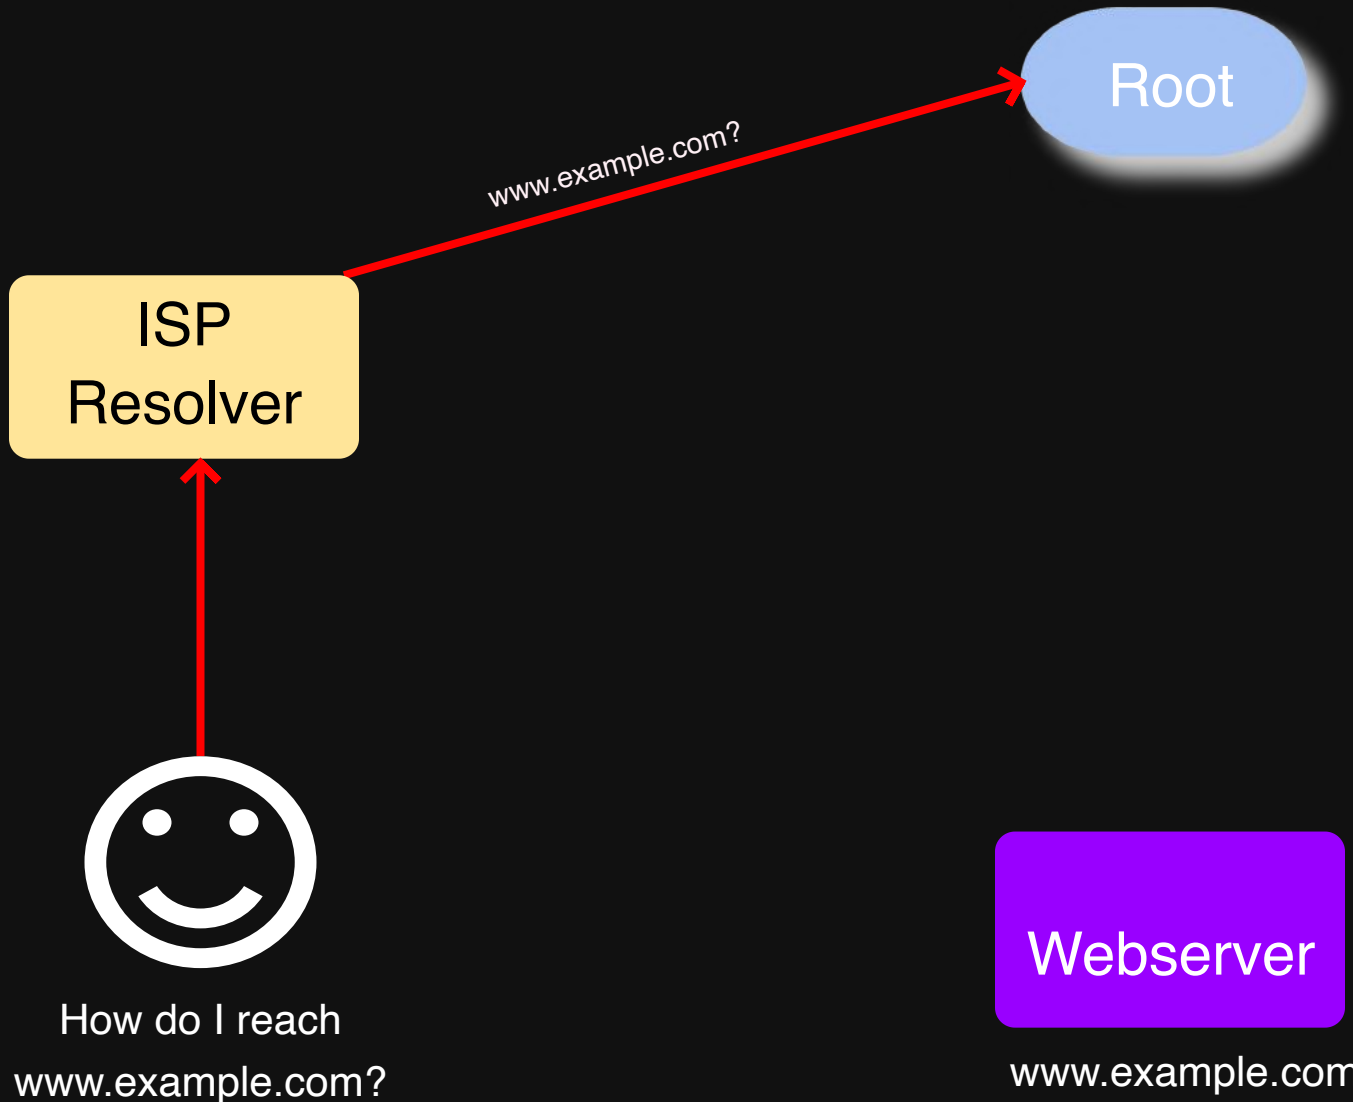

A refresher on how resolution works...

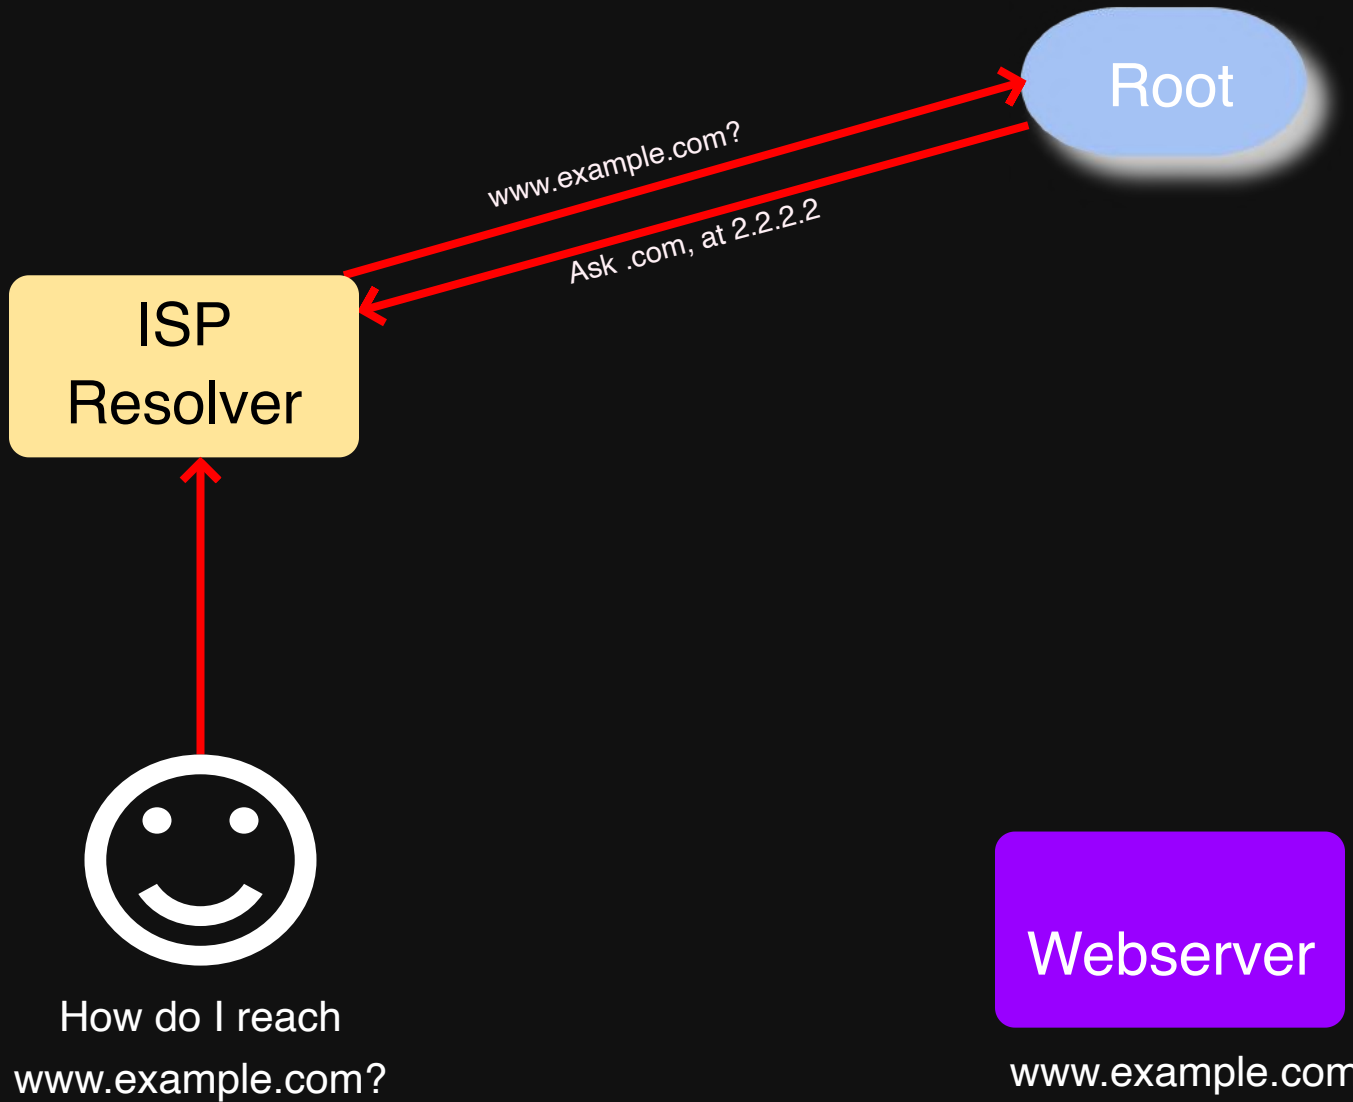

A refresher on how resolution works...

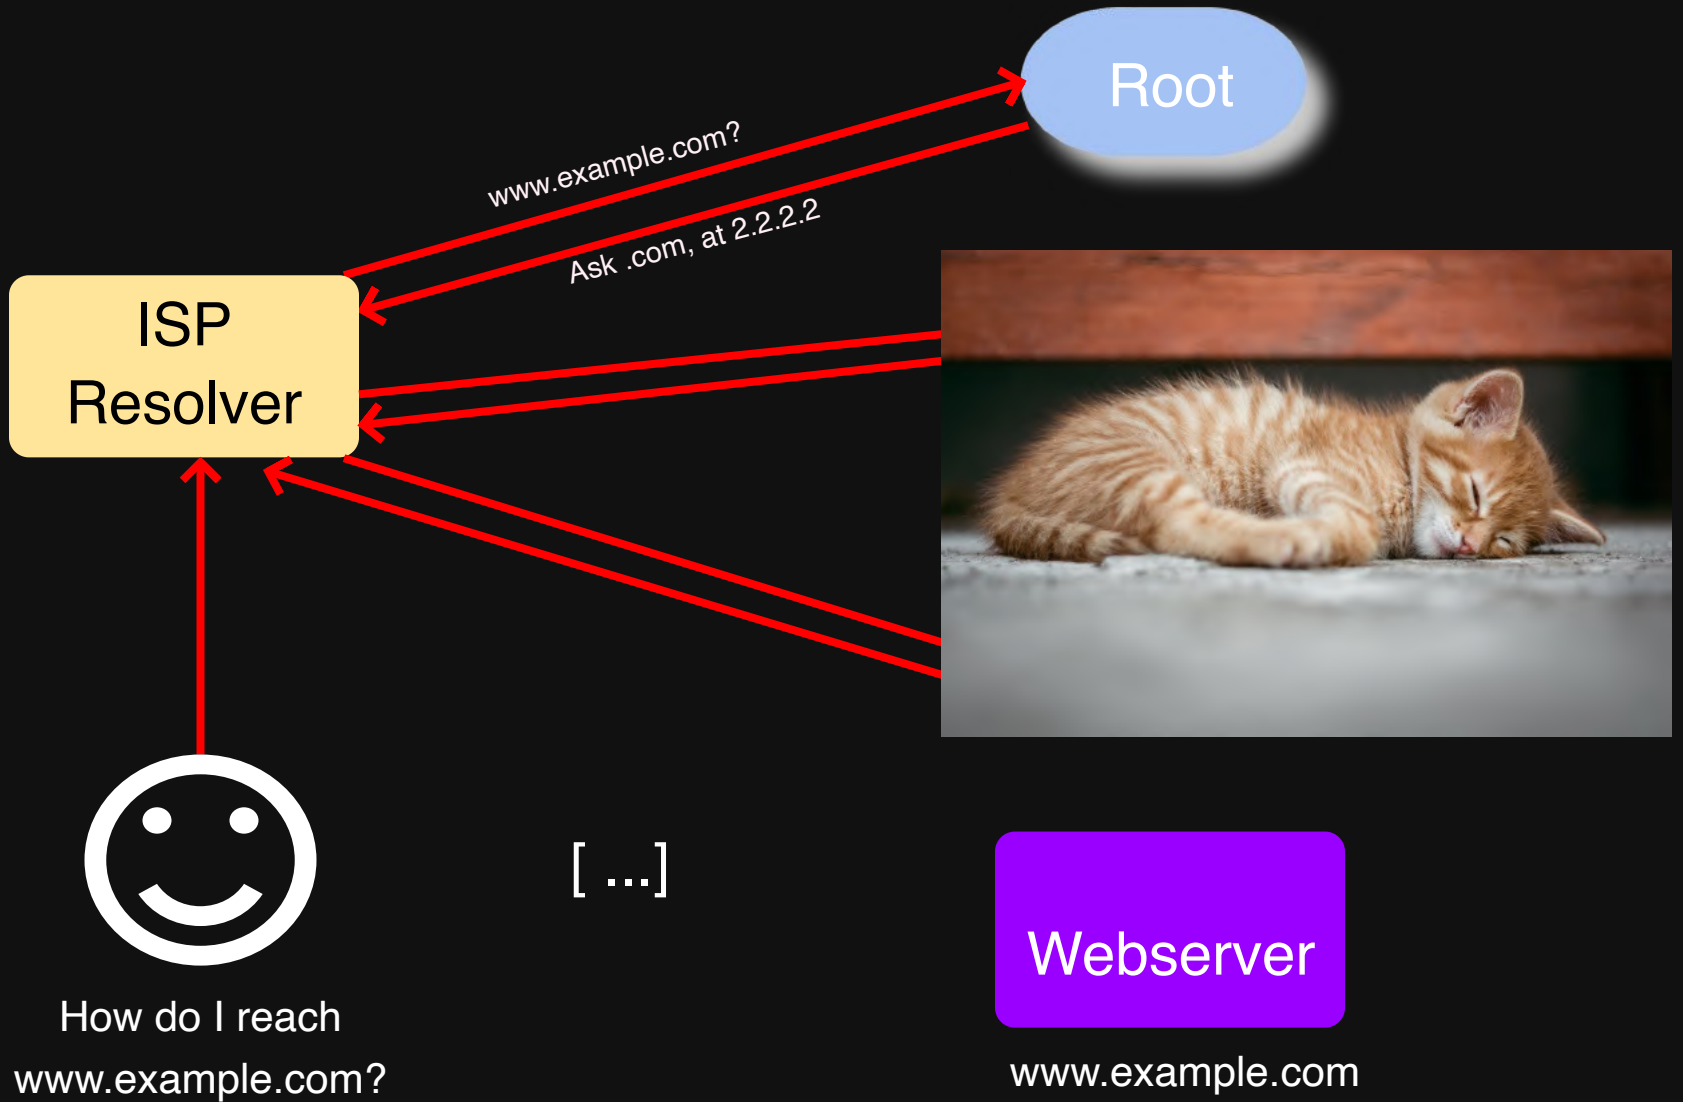

A refresher on how resolution works...

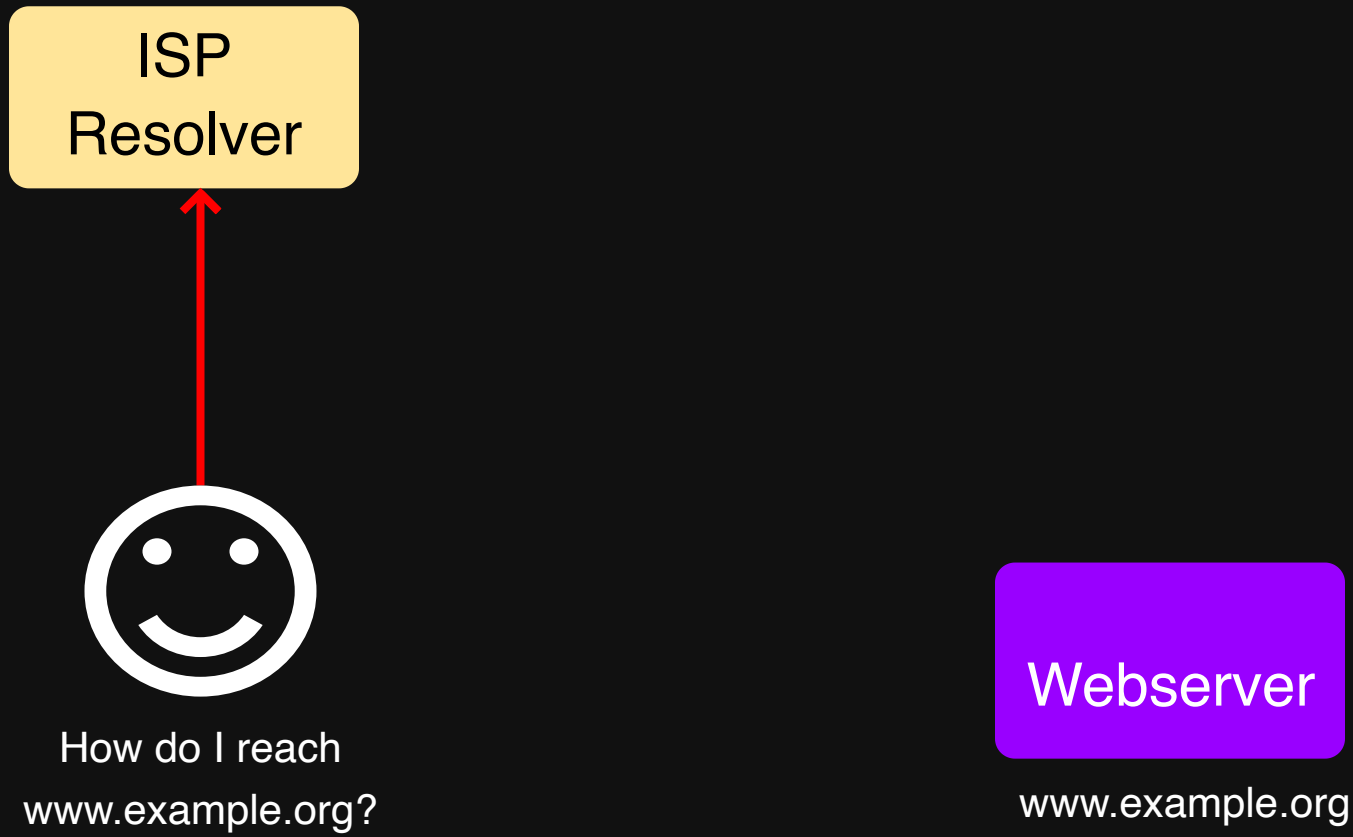

A refresher on how resolution works...

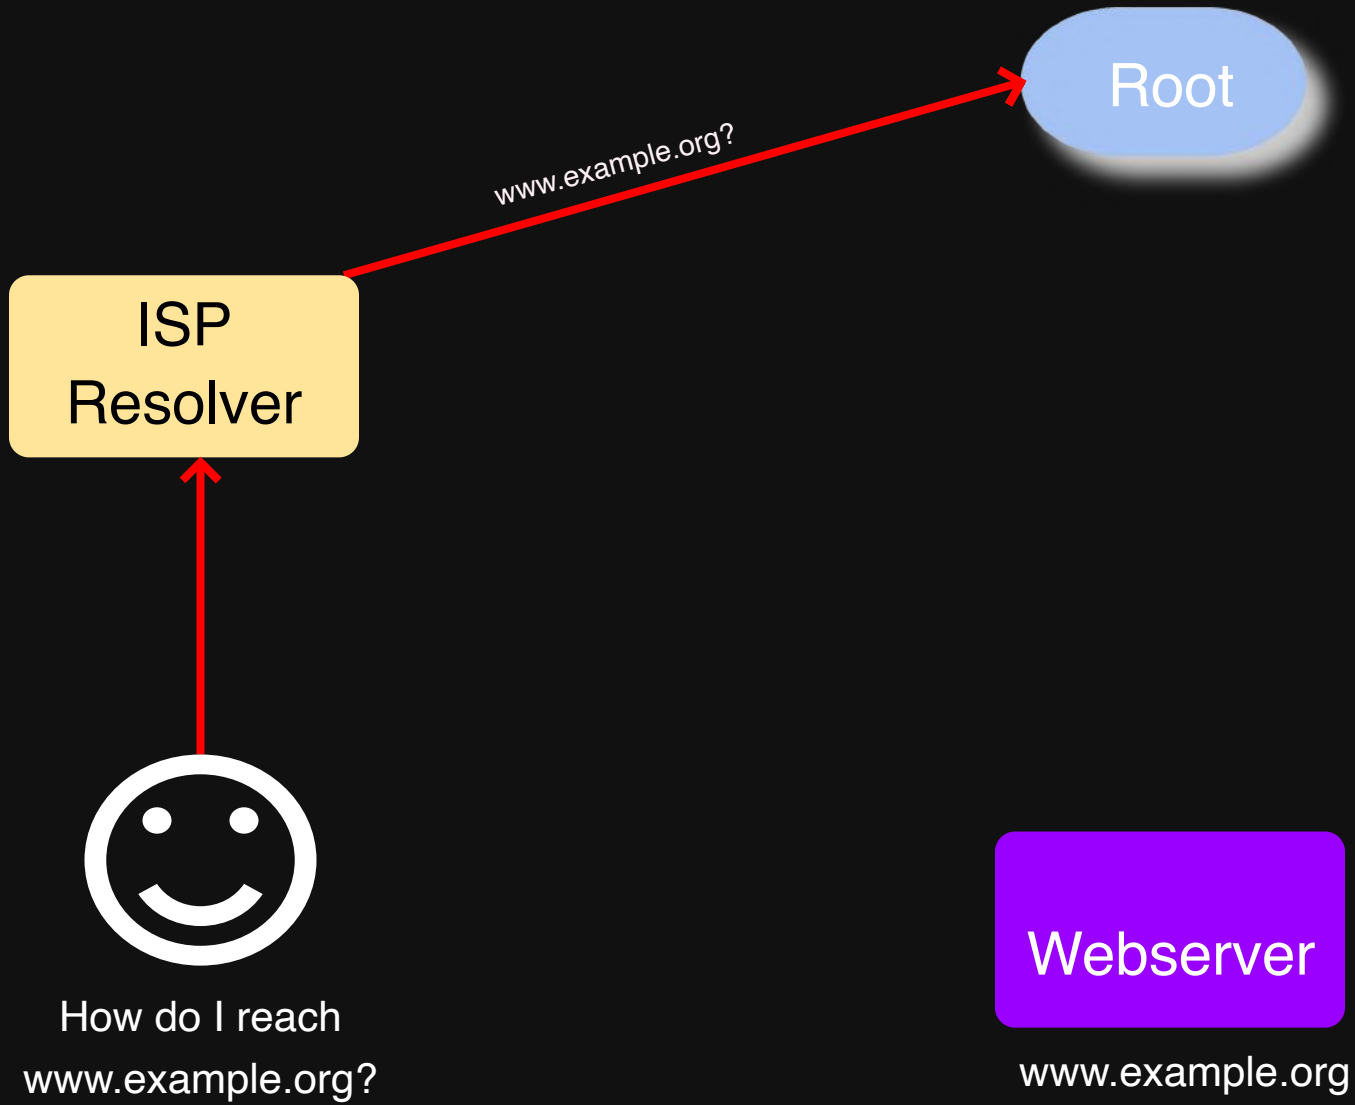

A refresher on how resolution works...

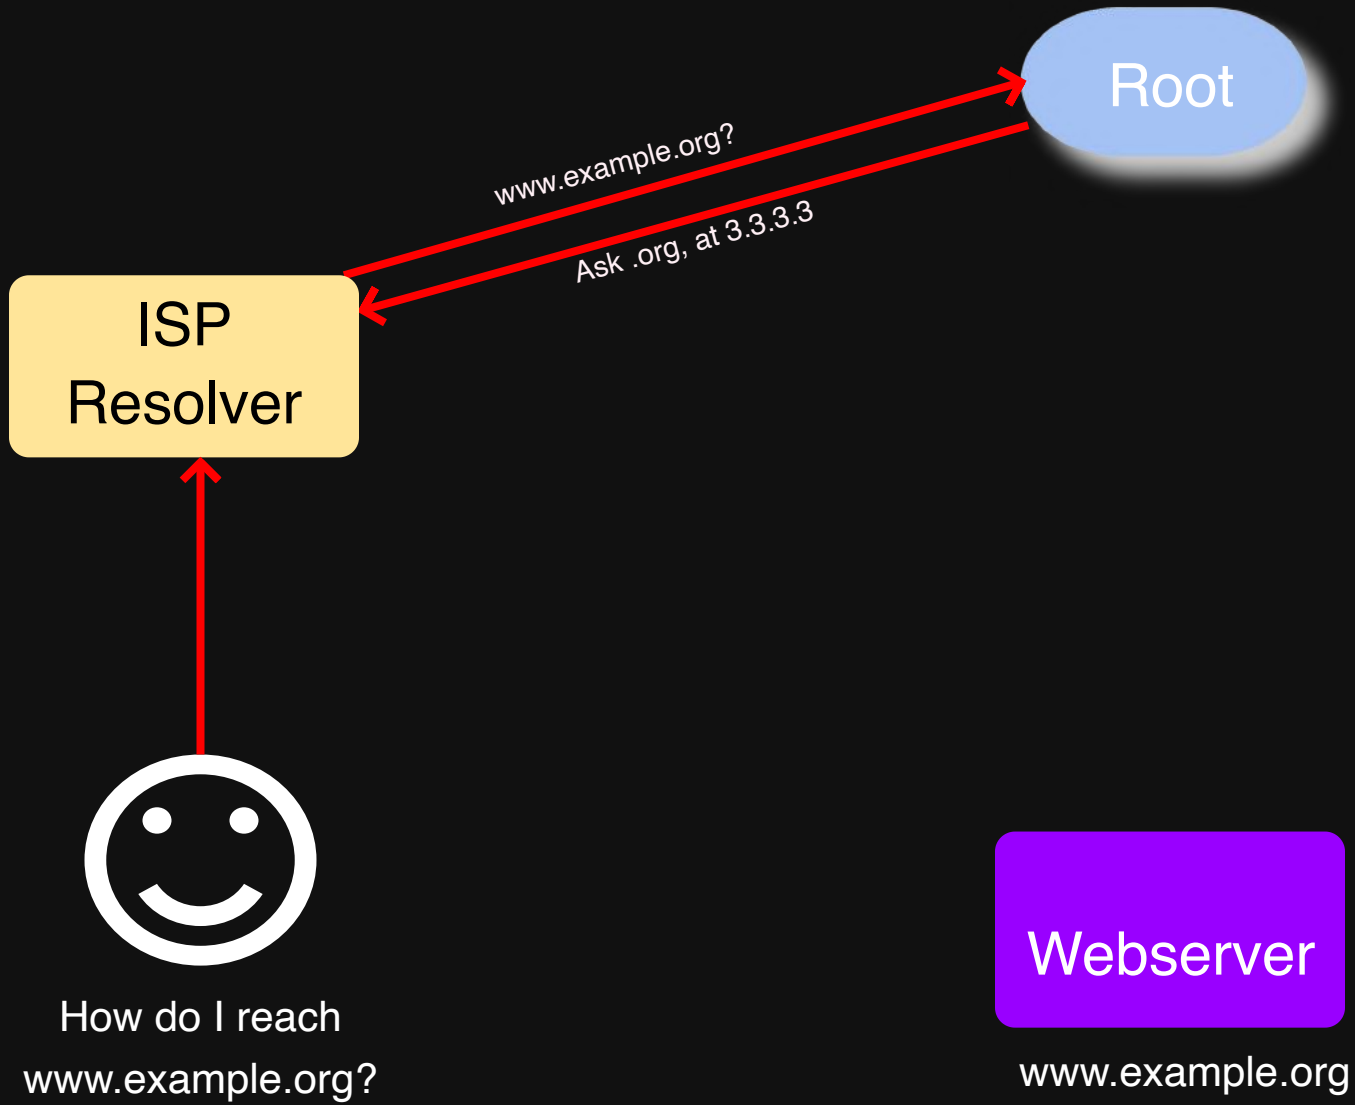

A refresher on how resolution works...

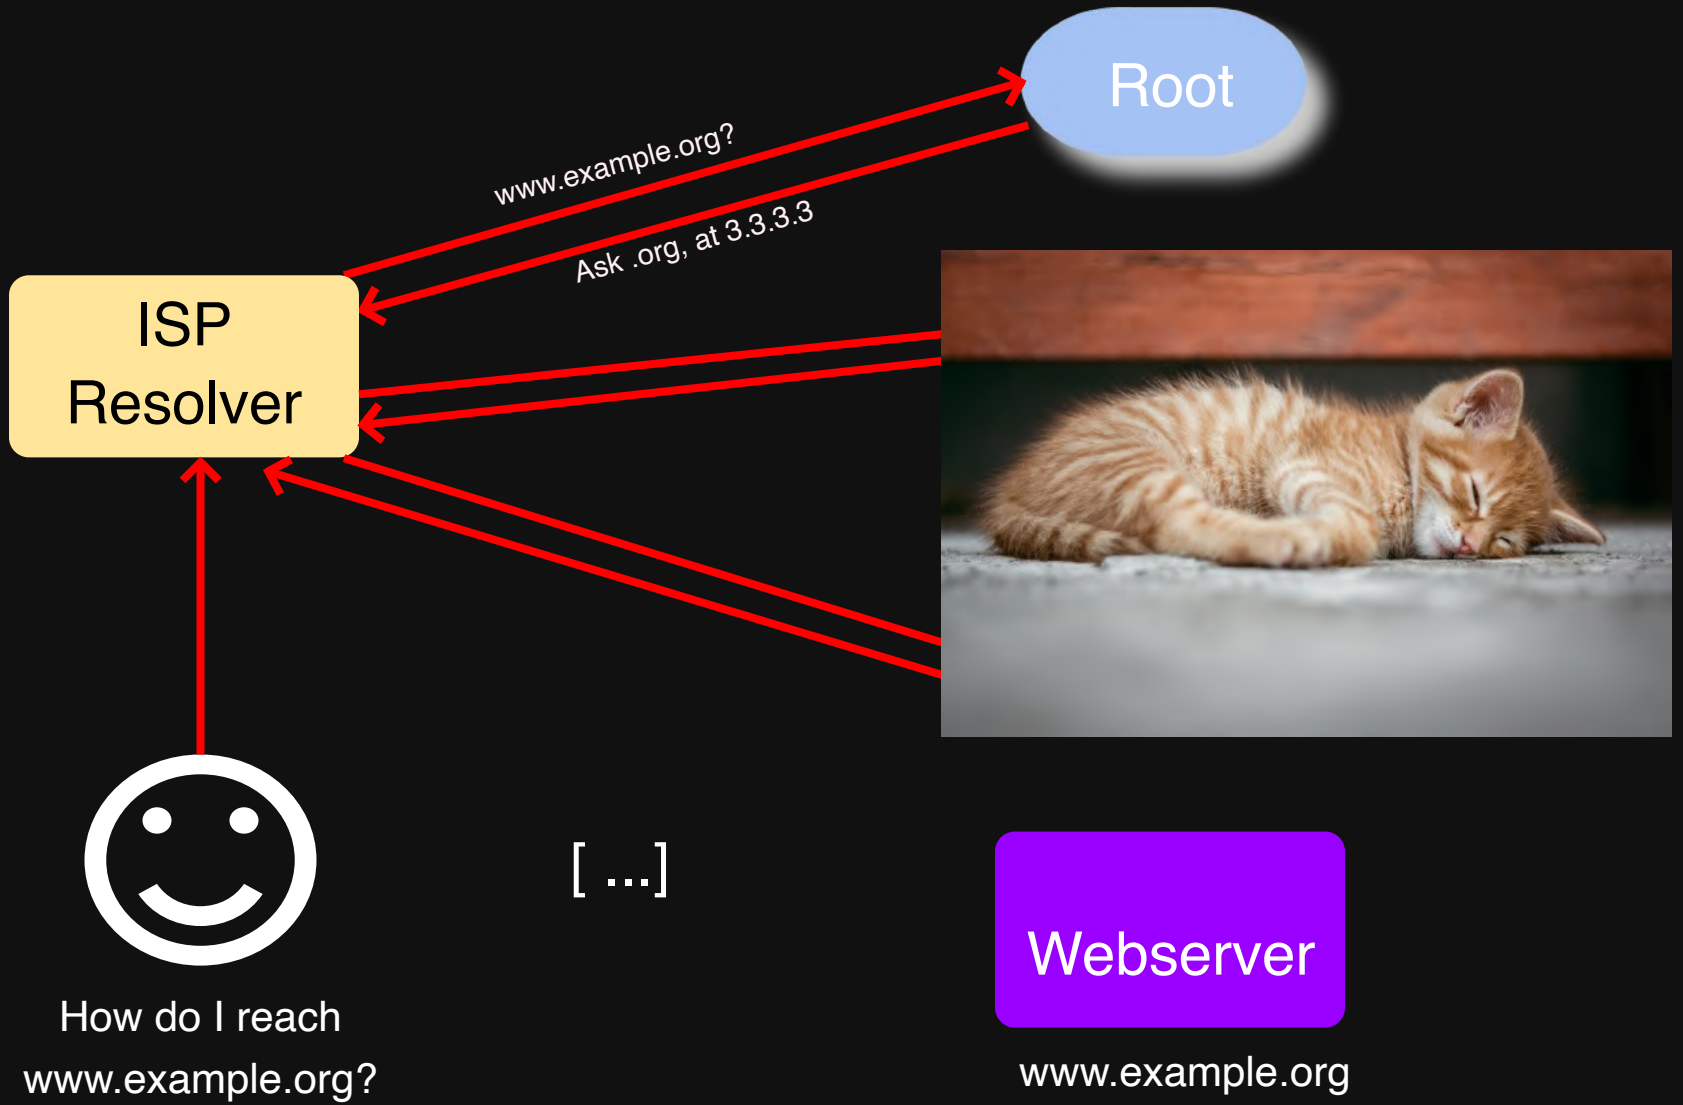

A refresher on how resolution works...

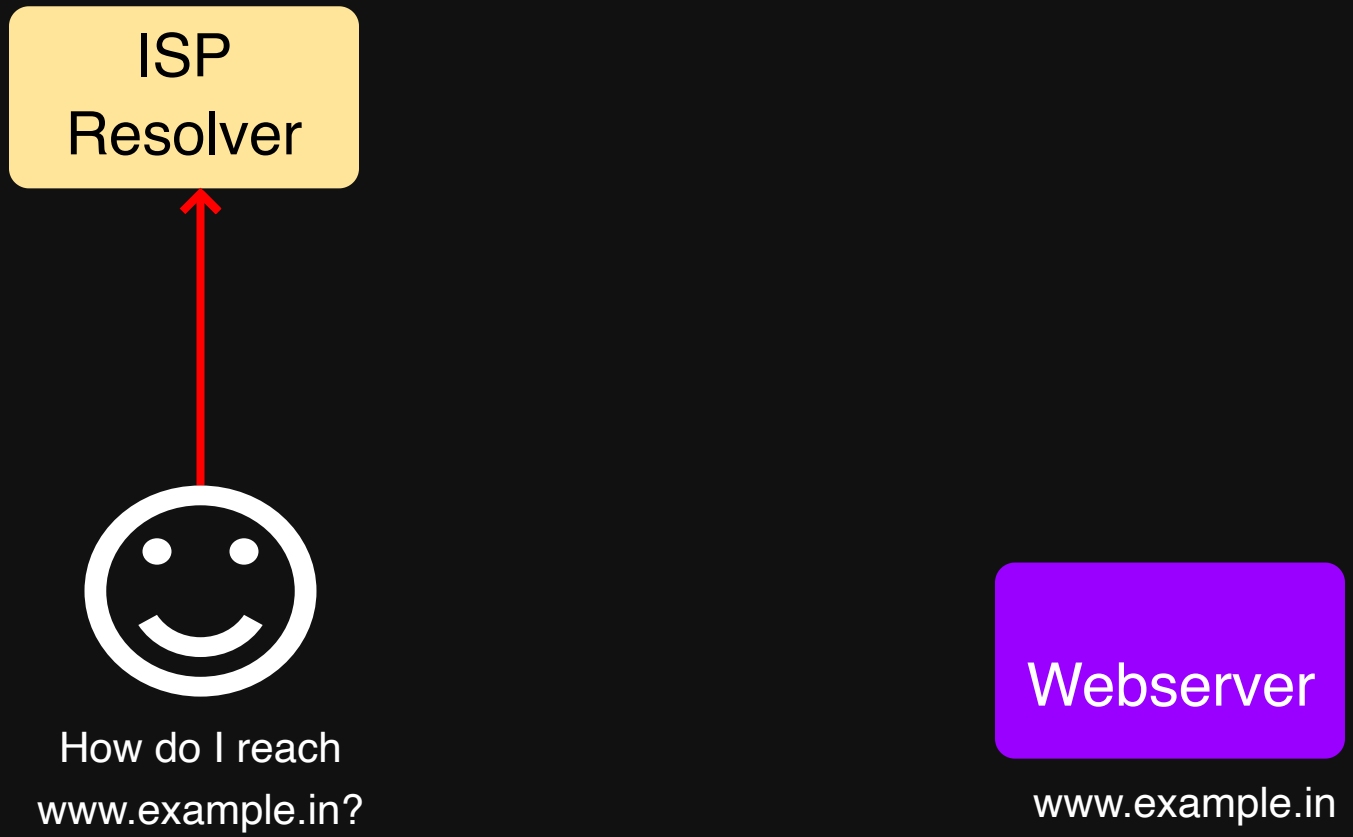

A refresher on how resolution works...

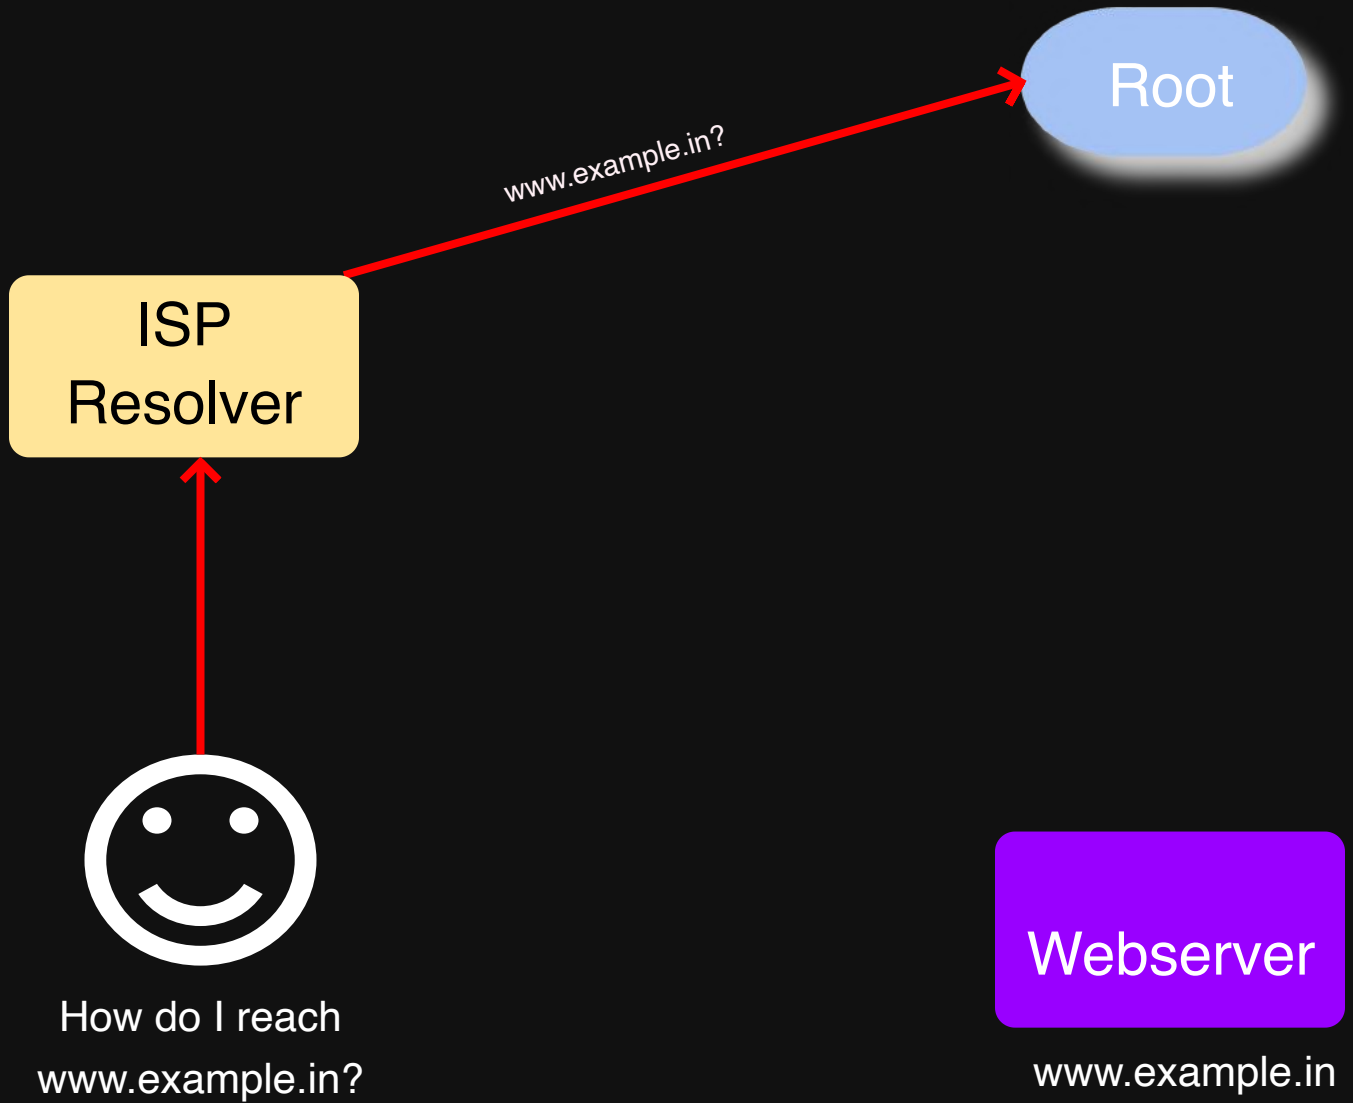

A refresher on how resolution works...

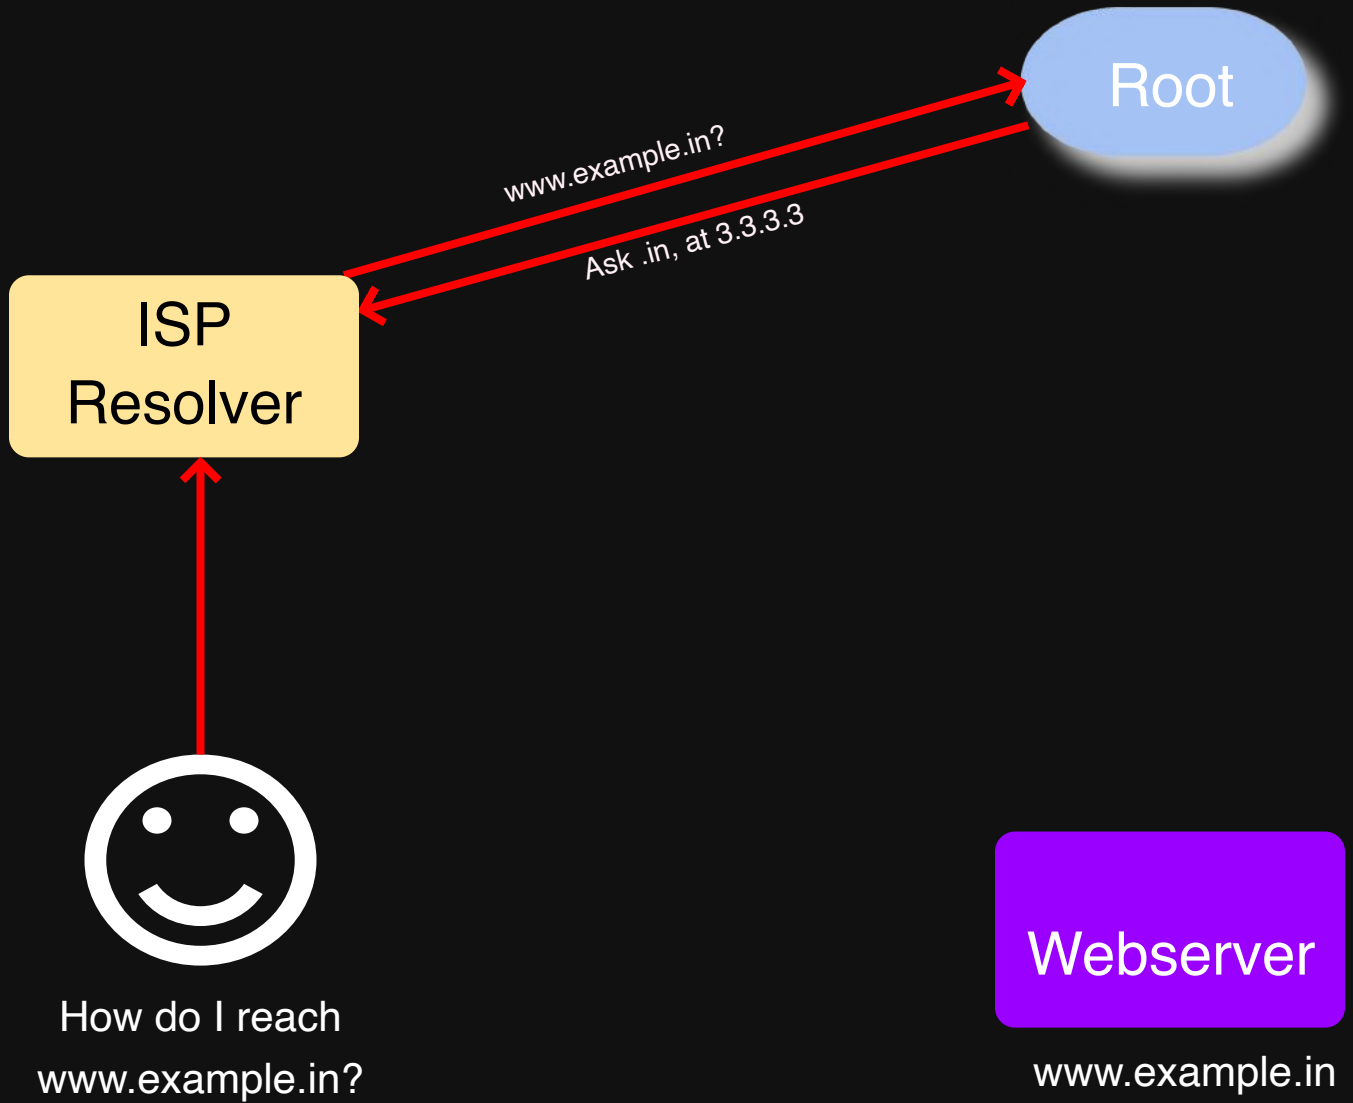

A refresher on how resolution works...

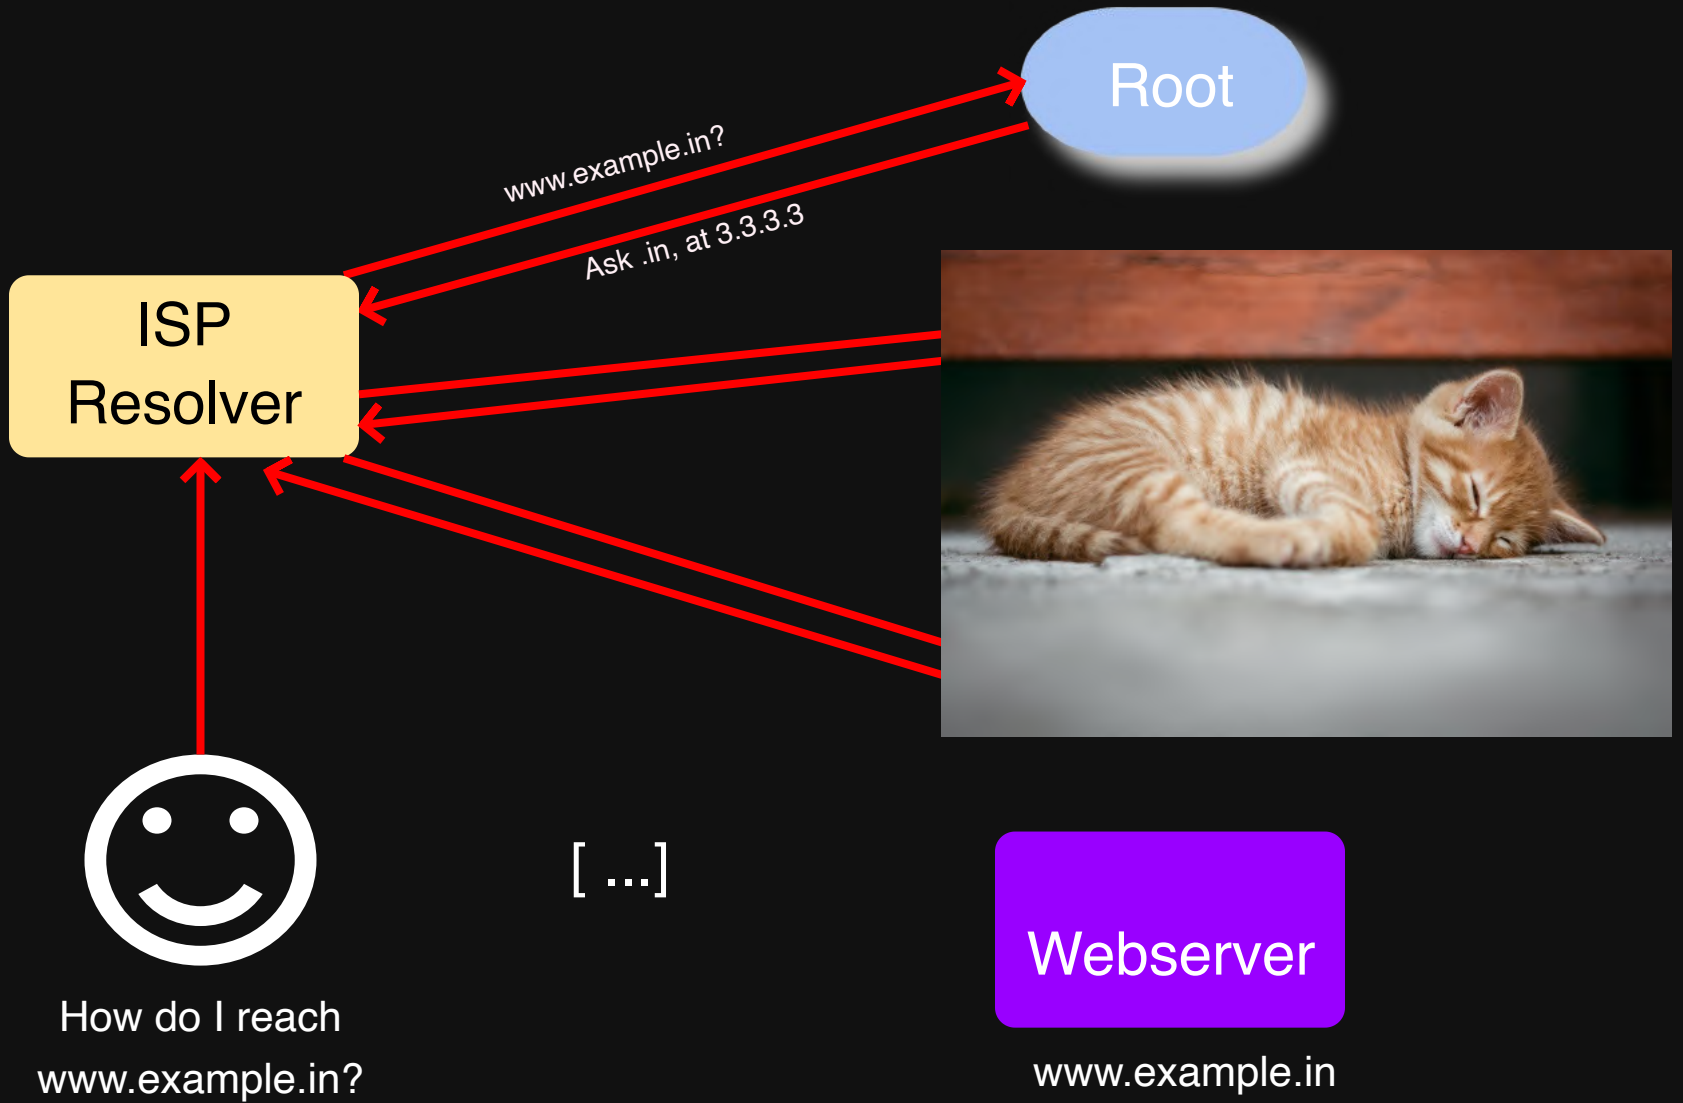

Wouldn't it be nice...

RFC8806 - "Running a Root Server Local to a Resolver"

RFC8806 - "Running a Root Server Local to a Resolver"

- Published in 2020, updates RFC7706 from 2015.

RFC8806 - "Running a Root Server Local to a Resolver"

- Published in 2020, updates RFC7706 from 2015.
- Basically: Fetch the root zone and cache it.

Benefits

Benefits

- Increased performance
  - If you don't need to send a packet, there is no RTT

Benefits

- Increased performance
  - If you don't need to send a packet, there is no RTT
- Increased reliability
  - If you don't need to reach the root, you don't need to be **able** to reach the root

Benefits

- Increased performance
  - If you don't need to send a packet, there is no RTT
- Increased reliability
  - If you don't need to reach the root, you don't need to be **able** to reach the root
- Improved privacy
  - www.alcoholics-anonymous.org
  - www.alcoholics-anonymous.orf

Benefits

- Increased performance
  - If you don't need to send a packet, there is no RTT
- Increased reliability
  - If you don't need to reach the root, you don't need to be **able** to reach the root
- Improved privacy
  - www.alcoholics-anonymous.org
  - www.alcoholics-anonymous.orf
- DoS Mitigation
  - Enumeration attacks

Yes, I'm aware of RFC7816 - "DNS Query Name Minimisation to Improve Privacy" and RFC8198 - "Aggressive Use of DNSSEC-Validated Cache" (I'm an author)

How?

Getting the root zone

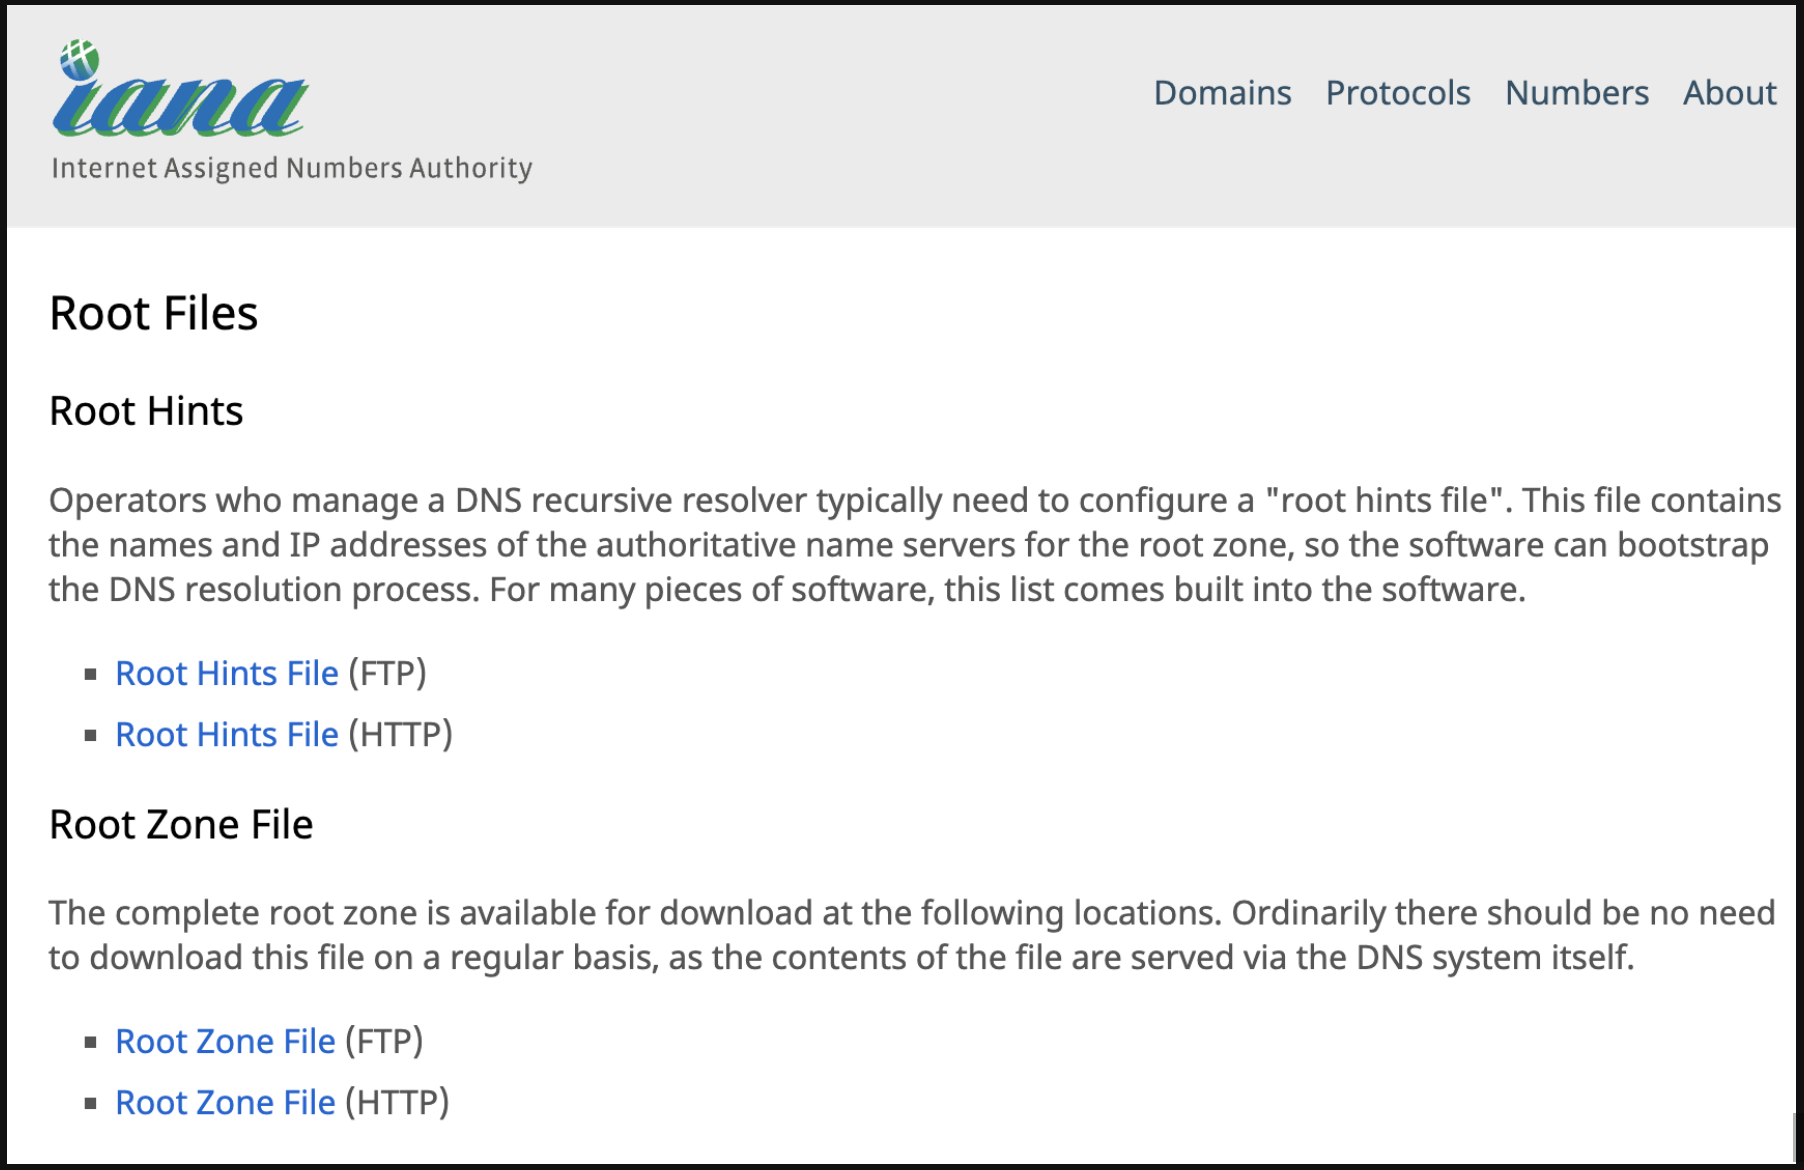

Getting the root zone

iana
Internet Assigned Numbers Authority

Domains Protocols Numbers About

Root Files

Root Hints

Operators who manage a DNS recursive resolver typically need to configure a "root hints file". This file contains the names and IP addresses of the authoritative name servers for the root zone, so the software can bootstrap the DNS resolution process. For many pieces of software, this list comes built into the software.

- Root Hints File (FTP)
- Root Hints File (HTTP)

Root Zone File

https://www.internic.net/domain/root.zone

The complete root zone is available for download at the following locations. Ordinarily there should be no need to download this file on a regular basis, as the contents of the file are served via the DNS system itself.

- Root Zone File (FTP)
- Root Zone File (HTTP)

Fetching using curl / wget

```
$ wget https://www.internic.net/domain/root.zone

--2025-10-17 11:37:51--  https://www.internic.net/domain/root.zone
Resolving www.internic.net (www.internic.net)... 192.0.46.9
Connecting to www.internic.net (www.internic.net)|192.0.46.9|:443... connected.
HTTP request sent, awaiting response... 200 OK
Length: 2246592 (2.1M) [text/plain]
Saving to: ‘root.zone’

root.zone                          100%[==============>]   2.14M  2.28MB/s    in 0.9s

2025-10-17 11:37:53 (2.28 MB/s) - ‘root.zone’ saved [2246592/2246592]
```

Fetching using AXFR

```
$ dig axfr . @b.root-servers.net

; <<>> DiG 9.10.6 <<>> axfr . @b.root-servers.net
;; global options: +cmd
.           86400   IN  SOA a.root-servers.net. nstld.verisign-grs.com. 2025101700 1800 900 60
.           518400  IN  NS  a.root-servers.net.
.           518400  IN  NS  b.root-servers.net.
.           518400  IN  NS  c.root-servers.net.
.           518400  IN  NS  d.root-servers.net.
.           518400  IN  NS  e.root-servers.net.
.           518400  IN  NS  f.root-servers.net.
...
...
ns1zim.telone.co.zw.    172800  IN  A   41.220.30.81
ns1zim.telone.co.zw.    172800  IN  AAAA    2c0f:f758:0:a::81
ns2zim.telone.co.zw.    172800  IN  A   41.220.30.82
ns2zim.telone.co.zw.    172800  IN  AAAA    2c0f:f758:0:a::82
.           86400   IN  SOA a.root-servers.net. nstld.verisign-grs.com. 2025101700 1800 900 60
;; Query time: 172 msec
;; SERVER: 170.247.170.2#53(170.247.170.2)
;; WHEN: Fri Oct 17 11:30:52 EDT 2025
;; XFR size: 24904 records (messages 86, bytes 1422842)
```

Can fetch from: [b, c, d, f, g].root-servers.net

Fetching using AXFR

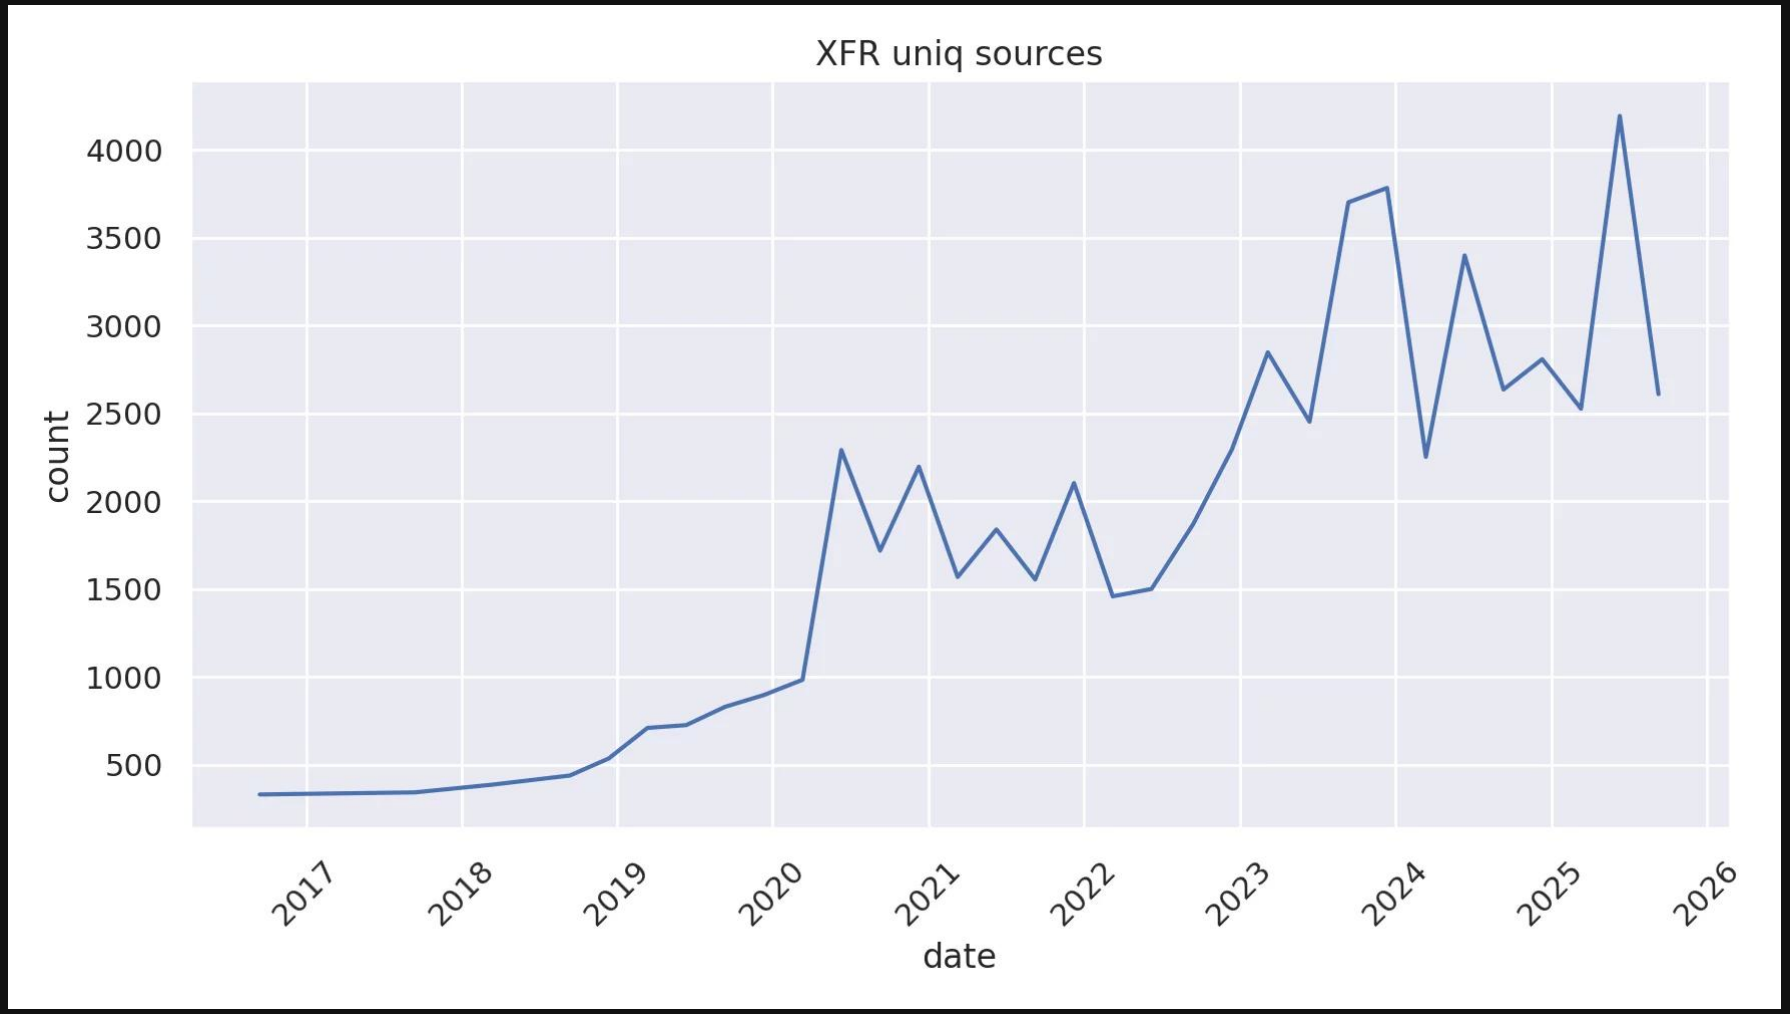

Stats from b.root-servers.net

ISC BIND 9.14 (and above)

Example configuration using a BIND "mirror" zone:

```
zone "." {
      type mirror;
  };
```

Source: BIND documentation for mirror zones

ISC BIND 9.14 (and above)

Example configuration using a BIND "mirror" zone:

Source: BIND documentation for mirror zones

ISC BIND 9.14 (and above)

Example configuration using a BIND "mirror" zone:

Source: BIND documentation for mirror zones

Knot Resolver

Example configuration to prefill cache with root zone using HTTPS:

```
modules.load('prefill')
 prefill.config({
       ['.'] = {
               url = 'https://www.internic.net/domain/root.zone',
               interval = 86400  -- seconds
               ca_file = '/etc/pki/tls/certs/ca-bundle.crt', -- optional
       }
 })
```

Source: Knot Resolver Cache prefilling

Knot Resolver

Example configuration to prefill cache with root zone using HTTPS:

Source: Knot Resolver Cache prefilling

Knot Resolver

Example configuration to prefill cache with root zone using HTTPS:

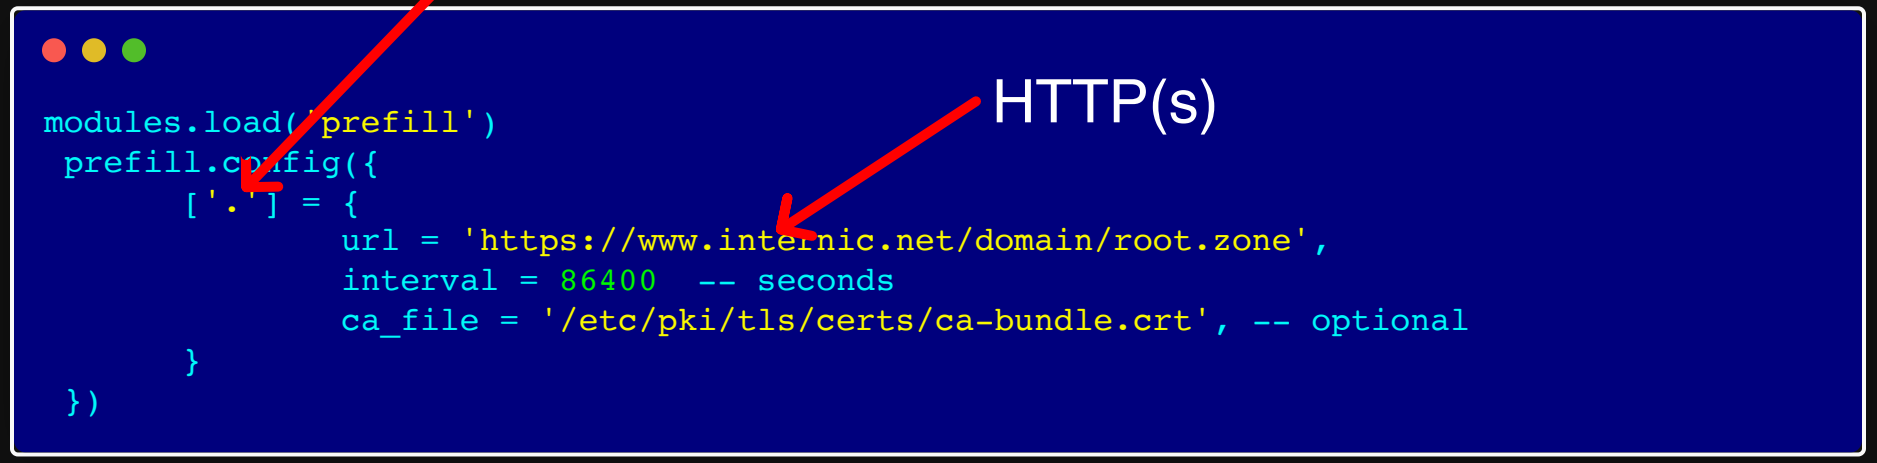

Source: Knot Resolver Cache prefilling

Knot Resolver

Example configuration to prefill cache with root zone using HTTPS:

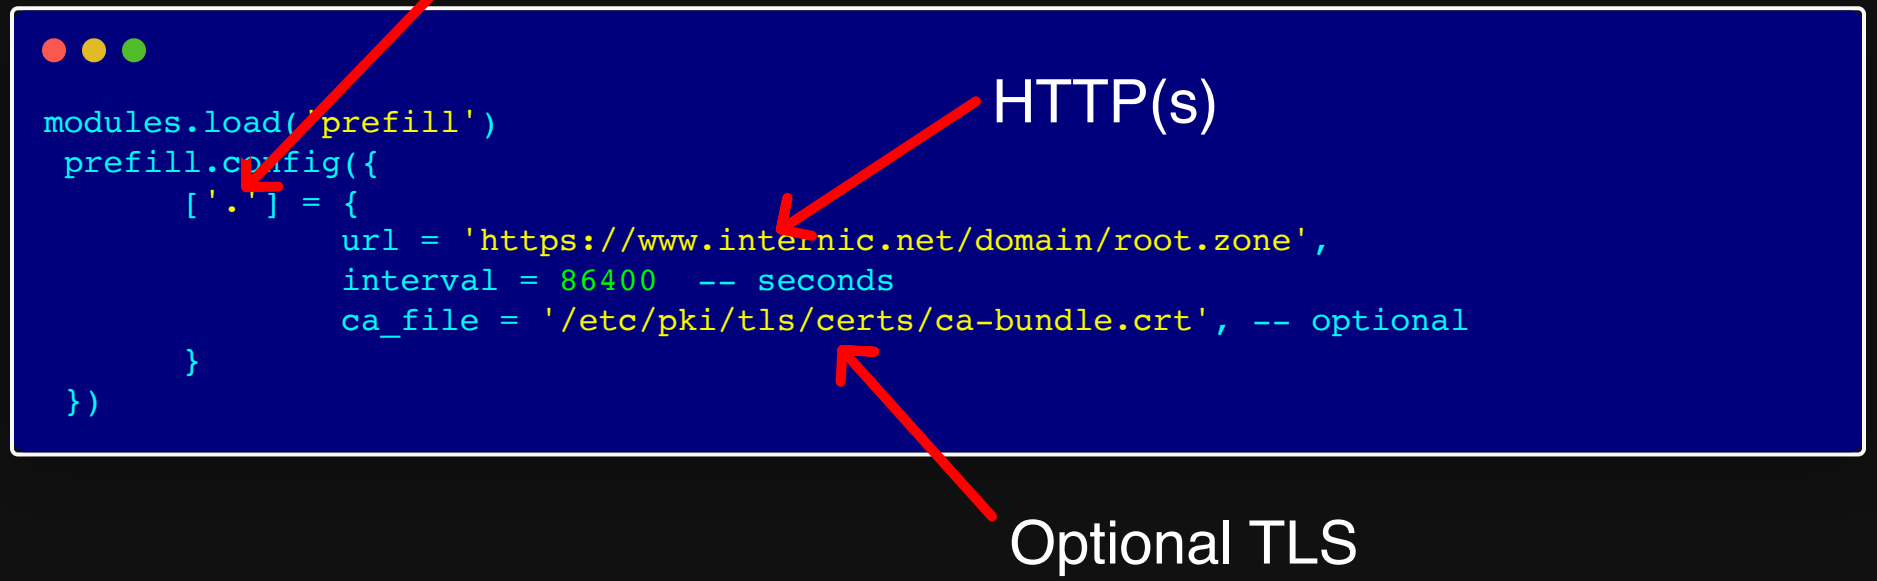

Source: Knot Resolver Cache prefilling

Unbound 1.9 (and above)

Unbound Auth Zones example:

```
auth-zone:
     name: "."
     url: "https://www.internic.net/domain/root.zone"
     zonefile: "root.zone"
          fallback-enabled: yes
       for-downstream: no
       for-upstream: yes
       zonefile: "root.zone"
     prefetch: yes
```

Source: Unbound Authority Zone Options

Unbound 1.9 (and above)

Unbound Auth Zones example:

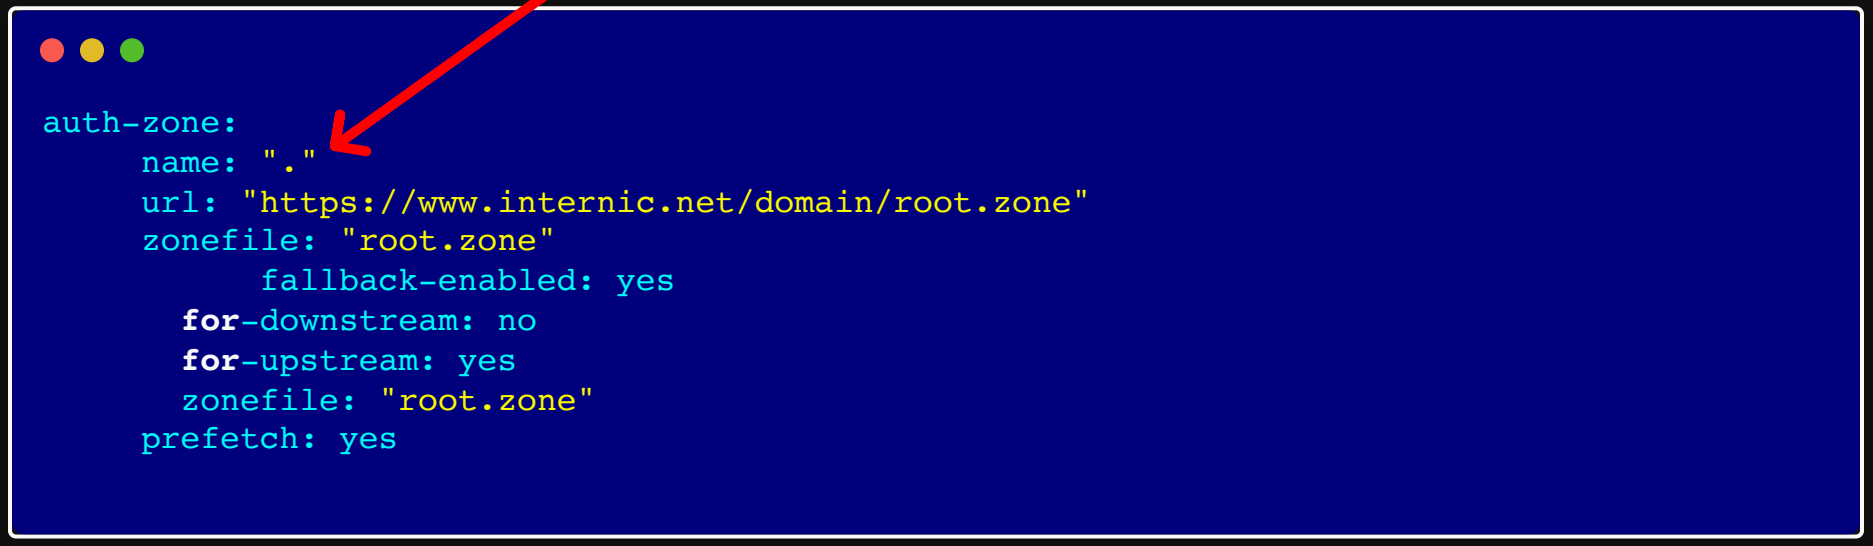

Source: Unbound Authority Zone Options

Unbound 1.9 (and above)

Unbound Auth Zones example:

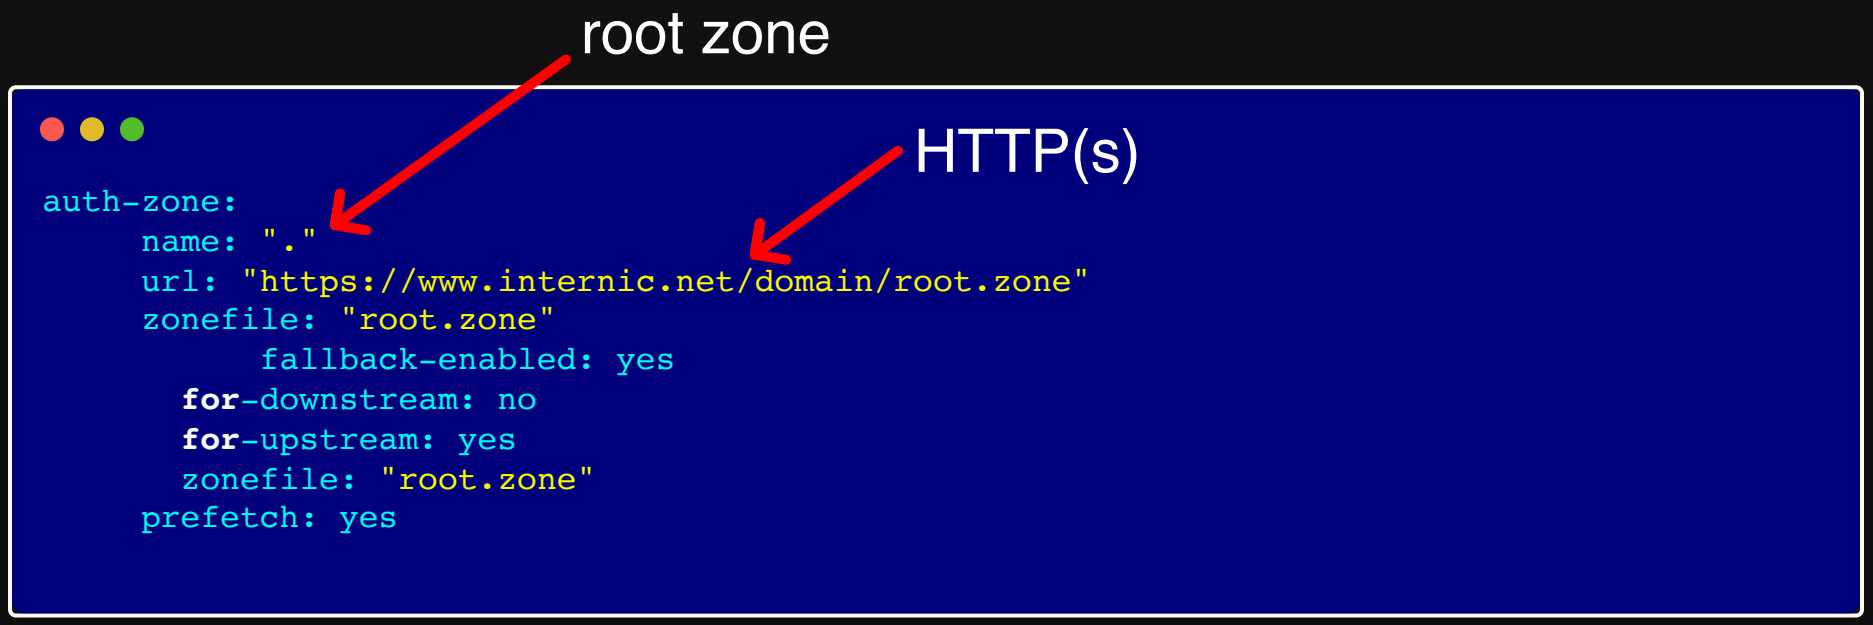

Source: Unbound Authority Zone Options

So, how much data is this?!

Individual queries

```
sudo ipset destroy root_servers
sudo ipset -N root_servers iphash

sudo ipset -A root_servers 198.41.0.4
sudo ipset -A root_servers 170.247.170.2
...

sudo iptables -A INPUT -m set --match-set root_servers src -j ACCEPT

sudo unbound-control flush_zone .
sudo iptables -L -Z -v

[ Query each TLD]

sudo iptables -L -vv
sudo iptables -L -vvv
ip filter INPUT 8
  [ match name set rev 4 ]
  [ counter pkts 1537 bytes 1127567 ]
  [ immediate reg 0 accept ]

Chain INPUT (policy ACCEPT 21512 packets, 5564K bytes)
 pkts bytes target     prot opt in     out     source               destination
 1537 1128K ACCEPT     all  --  any    any     anywhere             anywhere             mat
    0     0 ACCEPT     all  --  any    any     anywhere             anywhere             mat
```

Individual queries

```
sudo ipset destroy root_servers
sudo ipset -N root_servers iphash

sudo ipset -A root_servers 198.41.0.4
sudo ipset -A root_servers 170.247.170.2
...

sudo iptables -A INPUT -m set --match-set root_servers src -j ACCEPT

sudo unbound-control flush_zone .
sudo iptables -L -Z -v

[ Query each TLD]

sudo iptables -L -vv
sudo iptables -L -vvv
ip filter INPUT 8
  [ match name set rev 4 ]
  [ counter pkts 1537 bytes 1127567 ]
  [ immediate reg 0 accept ]

Chain INPUT (policy ACCEPT 21512 packets, 5564K bytes)
 pkts bytes target     prot opt in     out     source               destination
 1537 1128K ACCEPT     all  --  any    any     anywhere             anywhere             mat
    0     0 ACCEPT     all  --  any    any     anywhere             anywhere             mat
```

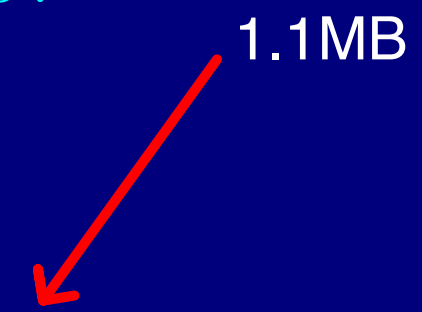

AXFR

```
dig axfr . @b.root-servers.net
```

Bytes: 1,602,476 (2719 packets)

HTTP

HTTPS

```
wget https://www.internic.net/domain/root.zone
```

Bytes: 2,379,902

HTTP

HTTPS

```
wget https://www.internic.net/domain/root.zone
```

Bytes: 2,379,902

HTTPS with nginx - gzip

```
wget https://www.owl-stretching-time.com/root.zone
```

Bytes: 990,949

CDNs - Content Delivery Network

CDNs - Content Delivery Network

- Specifically designed to distribute web objects

CDNs - Content Delivery Network

- Specifically designed to distribute web objects
- Basically every webpage and video is served from a CDN

CDNs - Content Delivery Network

- Specifically designed to distribute web objects
- Basically every webpage and video is served from a CDN
- Designed for high scalability, geographic distribution

Capacity

Recusive Resolvers	1,000,000
Zone Size (MB)	2.5
Zone Size (GZIP MB)	0.9
Updates per Day	3
Uncompressed	
Total Data (Day)	7 TB
Total Data (Month)	222 TB
GZIP Compressed	
Total Data (Day)	3 TB
Total Data (Month)	80 TB
Rate (Transfers / sec)	34.72

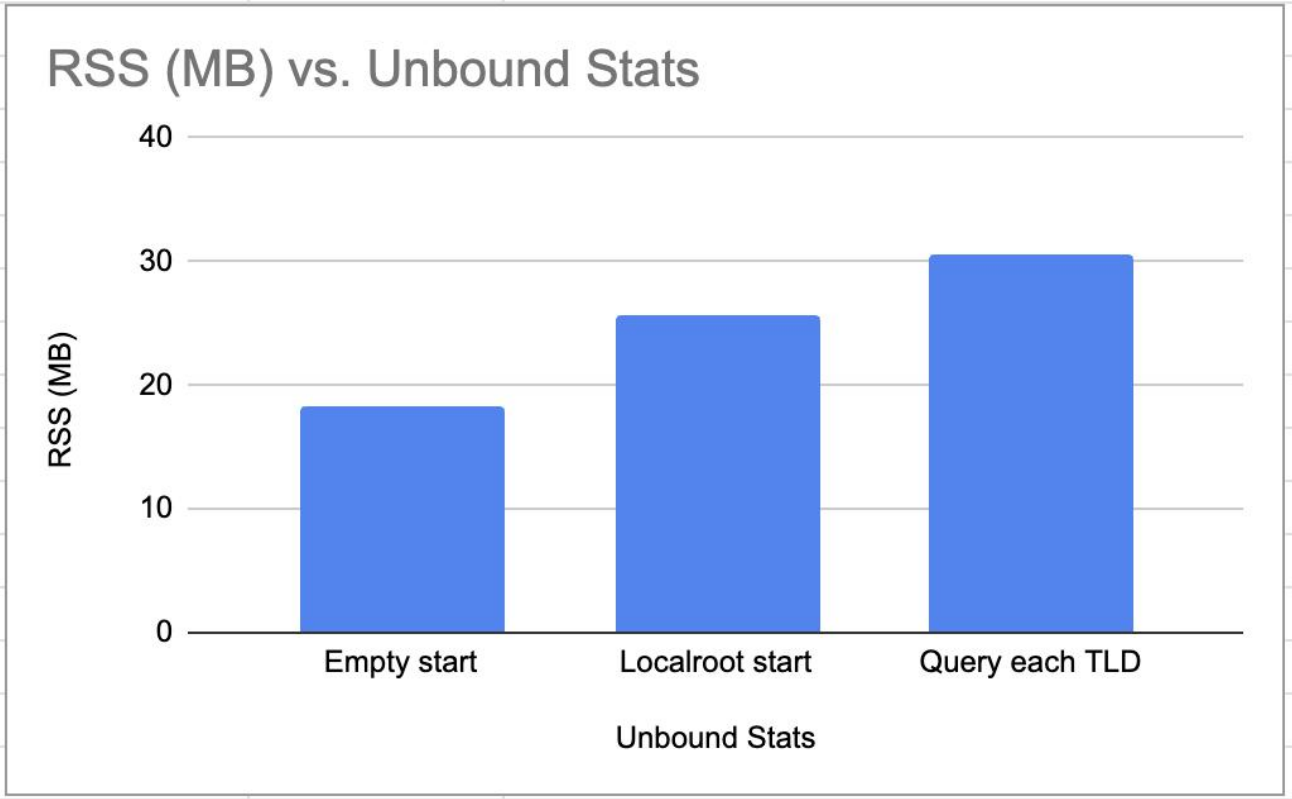

Trust...

RFC 8976 - "Message Digest for DNS Zones"

- "... an RR type that provides a cryptographic message digest of the data in a zone."
- Also known as "ZONEMD"
- DNSSEC signature over the entire zone file

RFC 8976 - "Message Digest for DNS Zones"

- "... an RR type that provides a cryptographic message digest of the data in a zone."
- Also known as "ZONEMD"
- DNSSEC signature over the entire zone file

```
$ dig zonemd .

; <<>> DiG 9.18.39-0ubuntu0.24.04.2-Ubuntu <<>> zonemd .
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 9730
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 1

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 65494
;; QUESTION SECTION:
;.                      IN   ZONEMD

;; ANSWER SECTION:
.                86400   IN   ZONEMD   2025102400 1 1 9679D824471ED710F3C6332217D1A70A96757F564D

;; Query time: 22 msec
```

draft-wkumari-dnsop-localroot-bcp

New document(s) provide:

- A list of places to fetch from
  - And algorithm (AXFR, HTTPS, ...)
- Procedures for updating this list
- Relaxing the implementation requirements
  - What to do (have the root zone locally), not how to do it (removed the "Thou must run an authoritative server...")
- Clarifications on fallback, resilience, etc.

draft-wkumari-dnsop-localroot-bcp

1. Identify locations from where root zone data can be obtained (Section 3.1).

2. Downloading and refreshing the root zone data from one of the publication points (Section 3.2).

3. Integrating and serving the data while performing DNS resolutions (Section 3.3).

- Has new text on retries, fallback and fetching by IP, where and how to get the zone, etc.

Benefits

Benefits

- Increased performance
  - If you don't need to send a packet, there is no RTT

Benefits

- Increased performance
  - If you don't need to send a packet, there is no RTT
- Increased reliability
  - If you don't need to reach the root, you don't need to be **able** to reach the root

Benefits

- Increased performance
  - If you don't need to send a packet, there is no RTT
- Increased reliability
  - If you don't need to reach the root, you don't need to be **able** to reach the root
- Improved privacy
  - www.alcoholics-anonymous.org
  - www.alcoholics-anonymous.orf

Benefits

- Increased performance
  - If you don't need to send a packet, there is no RTT
- Increased reliability
  - If you don't need to reach the root, you don't need to be **able** to reach the root
- Improved privacy
  - www.alcoholics-anonymous.org
  - www.alcoholics-anonymous.orf
- DoS Mitigation
  - Enumeration attacks

Yes, I'm aware of RFC7816 - "DNS Query Name Minimisation to Improve Privacy" and RFC8198 - "Aggressive Use of DNSSEC-Validated Cache" (I'm an author)

Benefits

- Increased performance
  - If you don't need to send a packet, there is no RTT
- Increased reliability
  - If you don't need to reach the root, you don't need to be **able** to reach the root
- Improved privacy
  - www.alcoholics-anonymous.org
  - www.alcoholics-anonymous.orf
- DoS Mitigation
  - Enumeration attacks
- Reduce criticality of the Root Server System

Yes, I'm aware of RFC7816 - "DNS Query Name Minimisation to Improve Privacy" and RFC8198 - "Aggressive Use of DNSSEC-Validated Cache" (I'm an author)

Questions?

Backup slides

Capacity

- Comparison to current root-server system
  - **One** CDN states 405Tbps
  - Current RSS: 1999 instances (2025-10-19)
  - That means each instance would need 202Gbps to match

Google

Cloudflare

Capacity

Through massive distribution, full automation, and a focus on efficiency, Akamai provides:

Better performance

Greater scalability

Much higher reliability

Lower cost

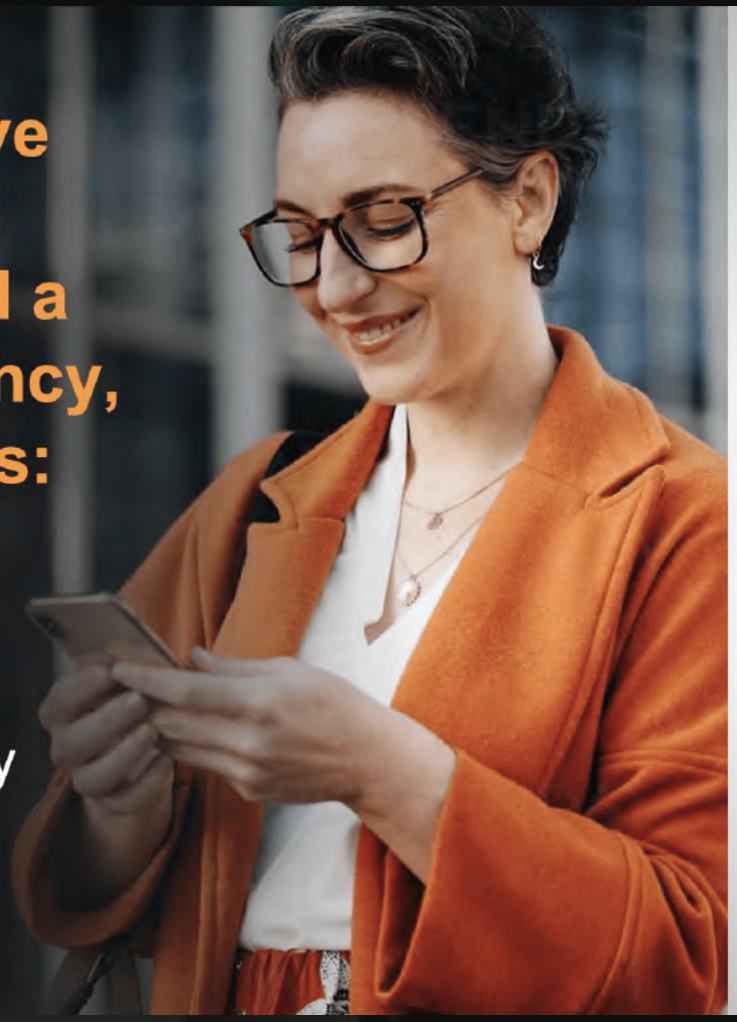

1,000
TBPS OF CAPACITY

4,100+
EDGE PoPs

1,200+
NETWORKS

750+
CITIES

130
COUNTRIES

Source: Akamai's Vision and Strategy Dr. Tom Leighton, Akamai CEO

DDoS Capacity

Aisuru

Source: KrebsOnSecurity

DDoS Capacity

Aisuru

- 30 Tbps / 1998 instances =~ 15Gpbs / instance

Source: KrebsOnSecurity

Root Query

Root Query Load

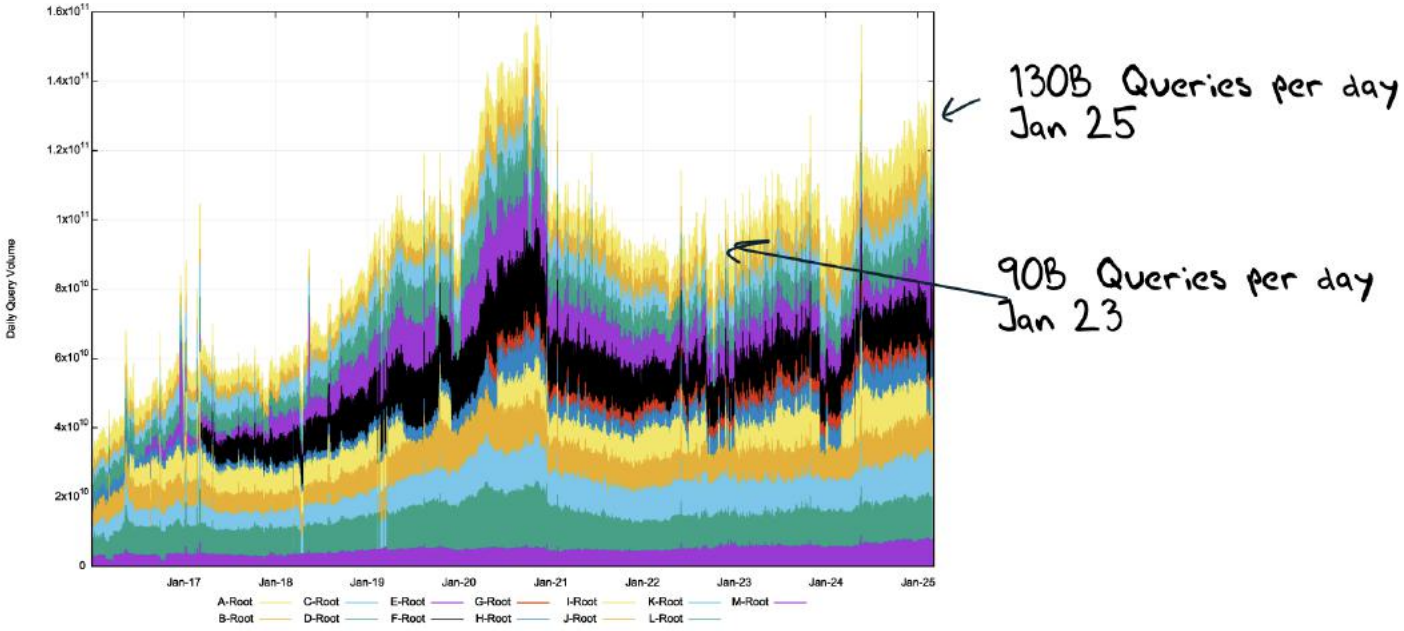

From https://github.com/rssac-caucus/RSSAC002-data

Source: Some Thoughts on the Root of the DNS (Geoff Huston)