The World* Turned Upside Down[†]

* I.e., the DNS Root Server System

Jeff Osborn
President, ISC (F-Root, BIND9) & Chair, RSSAC

RIPE NCC, Krakow, May 2024

[†] With apologies to Lin-Manuel Miranda, Hamilton the Musical

Problem

- Most policy makers do not understand the Root Server System
- Some policy makers need to understand the Root Server System
    - Not only what it is (theoretically)
    - Also what it means (operationally)

The “usual” way to explain DNS

- Assumes cold start scenario: resolver knows nothing
- Focus on name space organisation logic, not operational mechanics
- Overstates short-term dependency (86,400,000 ms) on RSS
- Understates operational significance of resolvers
- Understates or ignores operational role of IANA/RZM

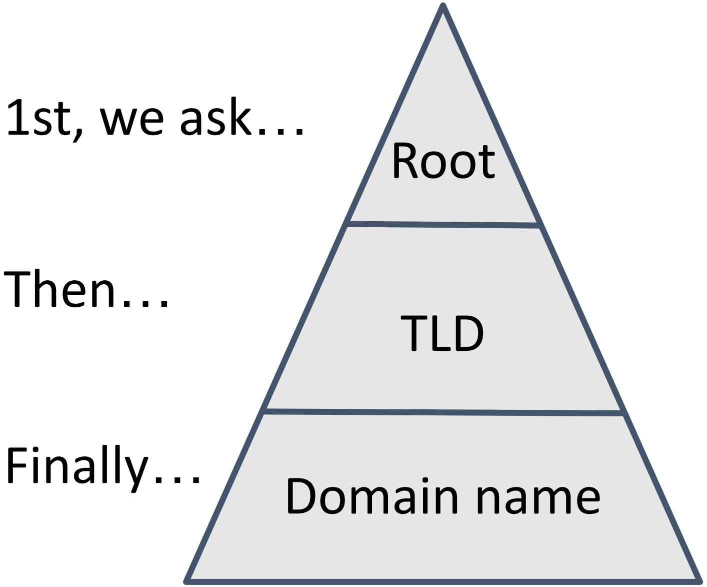

What’s the harm (with the usual approach)?

- Creates the FALSE impression that RSS is a “gatekeeper” to the Internet; RSS as on-ramp (slip road) entry point to the Internet
- Politicians invest too much meaning in some engineering terms like “hierarchy”
- Fails to explain the close-to-zero observable impact if some components of the RSS were to fail briefly
    - The Root Server System as a whole has never failed in 40 years
    - RSS now comprised of 1700+ server instances; with anycast; operators act independently; no technological single point of failure; no institutional single point of failure

Solution… invert to show resolver reality and <u>frequency</u>

1st, we ask…

When we need to, we ask…

When we need to, we ask…

<u>If we really know nothing</u>, we ask…

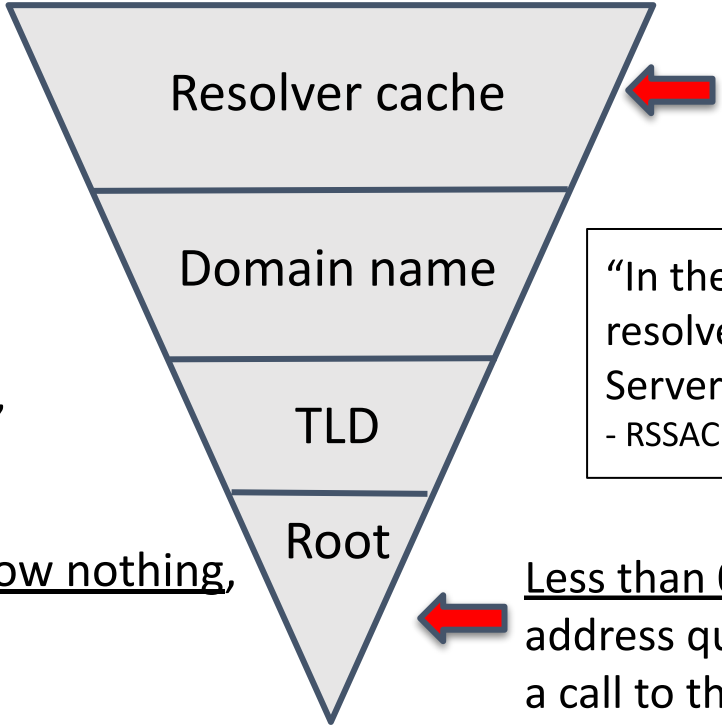

<u>More than 90%</u> of all address queries are resolved here

> “In the millisecond world of a resolver, queries to the Root Server System are rare”
> - RSSAC (forthcoming publication)

<u>Less than 0.02%</u> of all address queries require a call to the RSS

How to present this message?

- Deliverable 1: detailed tutorial/explainer
    - Written for non-technical audience
    - Current RSSAC draft 15 pages, potentially finished at ICANN 80 (June 9-14, 2024)
- Deliverable 2: Slide shows based on Deliverable 1
    - Draft versions presented to friendly audiences at ICANN 79 (Feb 3-8, 2024) for feedback
    - Now, it looks like this…

The DNS Root Server System

Introduction for a non-technical audience

(PREVIEW EDITION)

Introducing DNS (the Domain Name System)

- DNS uses human names to find computer addresses
  - Humans know the domain names like: www.amazon.com
  - Computers know IP addresses like: 18.239.62.181
  - DNS looks up “www.amazon.com” and gets “18.239.62.181”
  - For the most part, numbers change, but names don’t
- Most connected devices need DNS to find things
  - Computers & servers
  - Smart phones
- Questions use a domain name; answers use IP addresses

Benefits of DNS

- Human-friendly identifiers
  - **`www.example.com`** is easier to use than **`192.168.45.99`**
- Service portability
  - Resource owners control address mapping in their domain
  - DNS follows you to your new online home
- It’s a huge distributed network that’s easy to use
  - Flexible delegated management of hundreds of millions of directories
  - World’s largest distributed database

Devices get addresses from resolvers

- There are millions of resolvers around the world
- It’s like resolvers can read all the world’s phone books
  - The phone books are authoritative servers
  - The phone book listings are zone data
- What is the number for www.amazon.com?
- The number for www.amazon.com (for now) is 18.239.62.181
  - This happens in milliseconds
  - This happens about 500 trillion times every day

Resolvers get addresses from authoritative servers

- The resolver remembers addresses
  - This is called caching
  - This is where answers come from most of the time
- Once in a while, it needs a new number or to confirm an old one
- Depending how much it needs, it will ask:
  1. A domain name’s authoritative server
  2. A domain name’s authoritative server, and a TLD’s authoritative server
  3. A domain name’s authoritative server, and a TLD’s authoritative server, and a root server

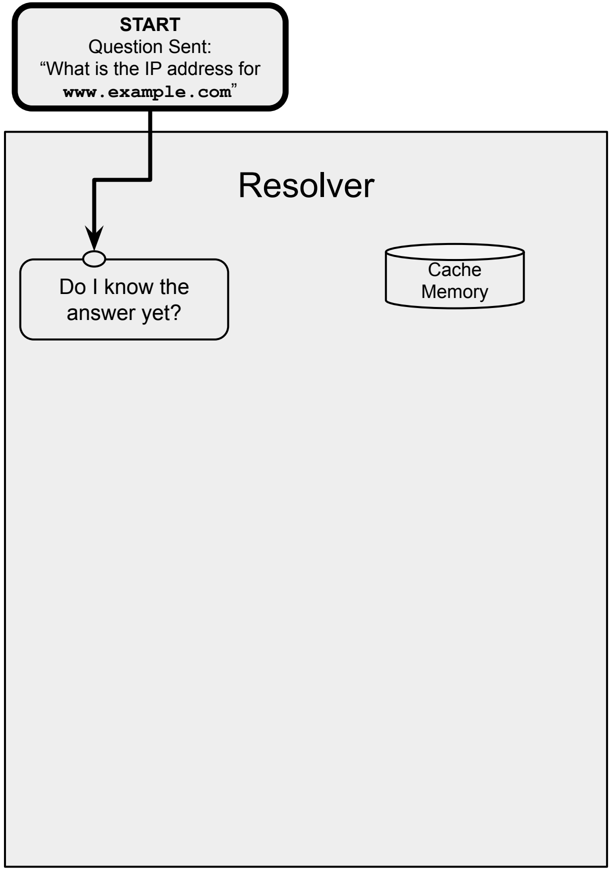
START
Question Sent:
"What is the IP address for **`www.example.com`**"
Resolver
Do I know the answer yet?
Cache Memory

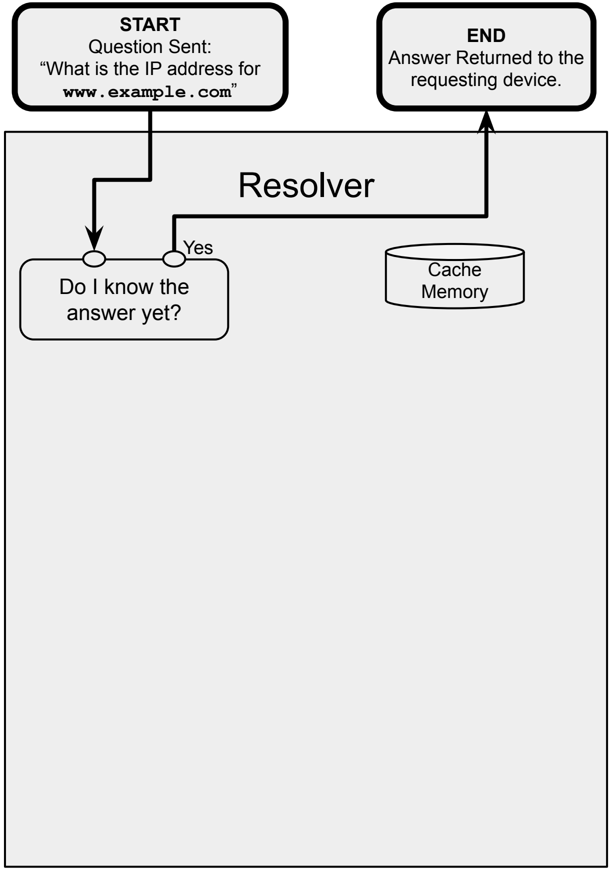

Frequency (estimates):
On average, how often do Resolvers consult at this level to answer a question?

Routine:
> 90% of answers are returned needing cache memory only

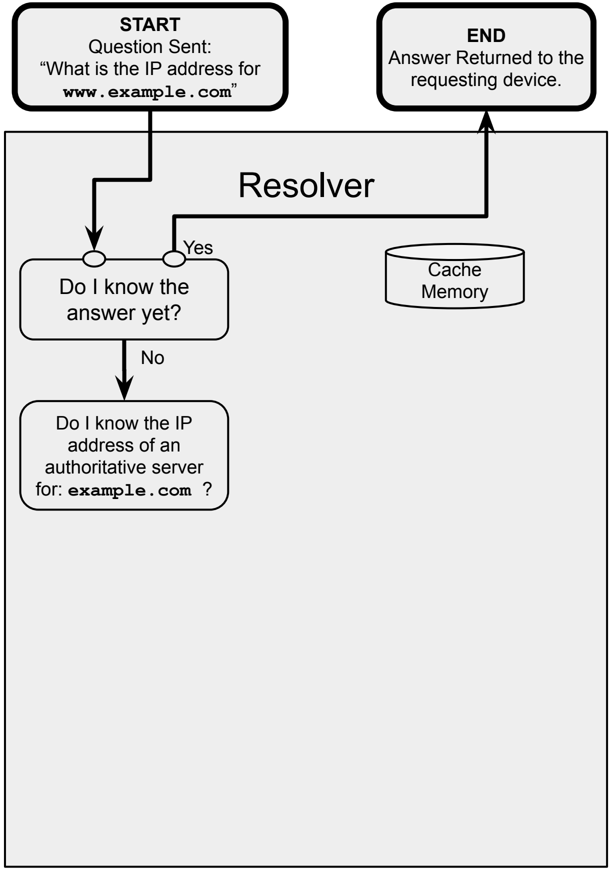

Frequency (estimates):
On average, how often do Resolvers consult at this level to answer a question?

Routine:
> 90% of answers are returned needing cache memory only

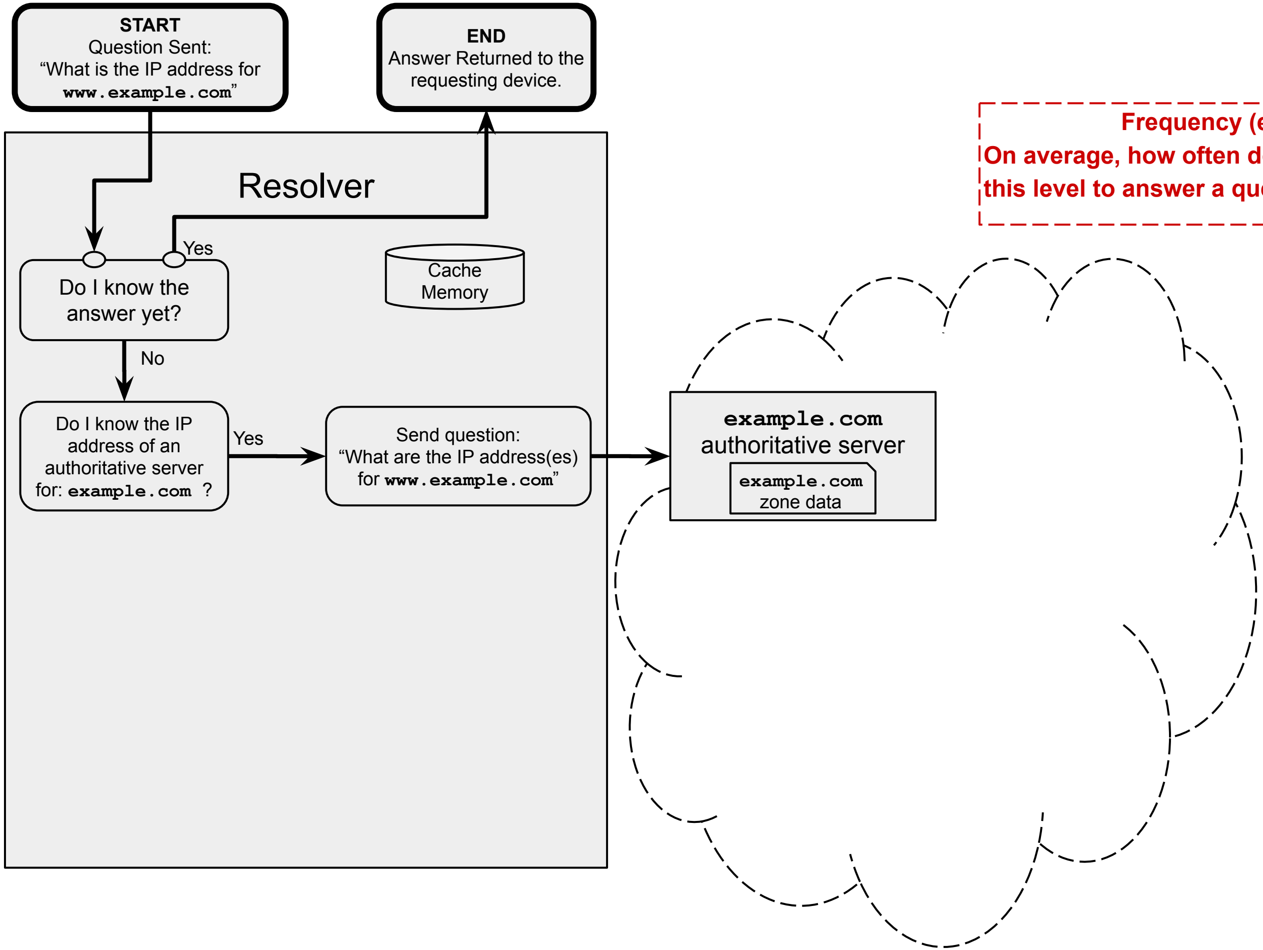

Frequency (estimates):

On average, how often do Resolvers consult at this level to answer a question?

Routine:

> 90% of answers are returned needing cache memory only

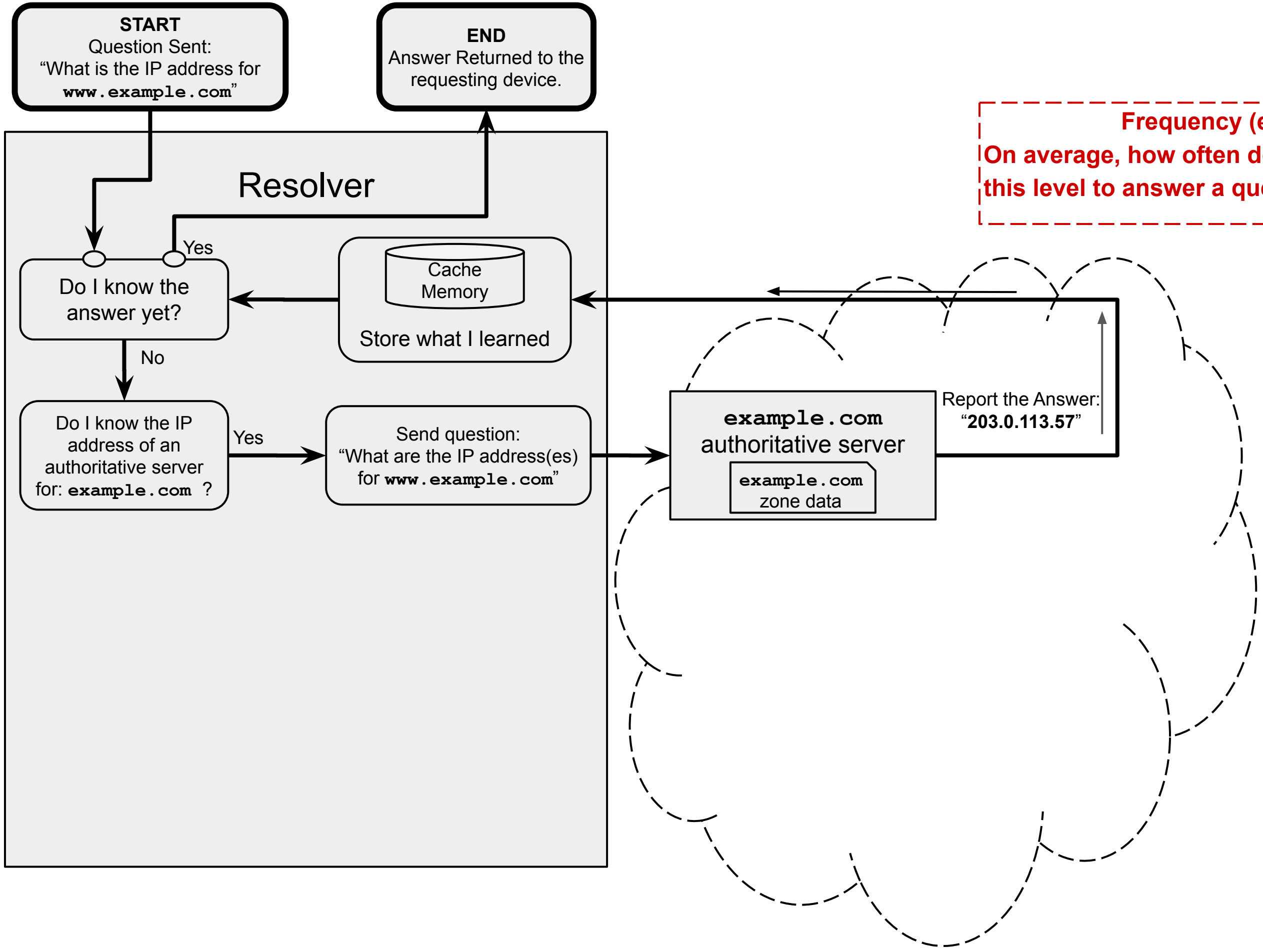

Frequency (estimates):
On average, how often do Resolvers consult at this level to answer a question?

Routine:
> 90% of answers are returned needing cache memory only

Ocasional:
~ 5% of answers require a question to the domain name’s authoritative server

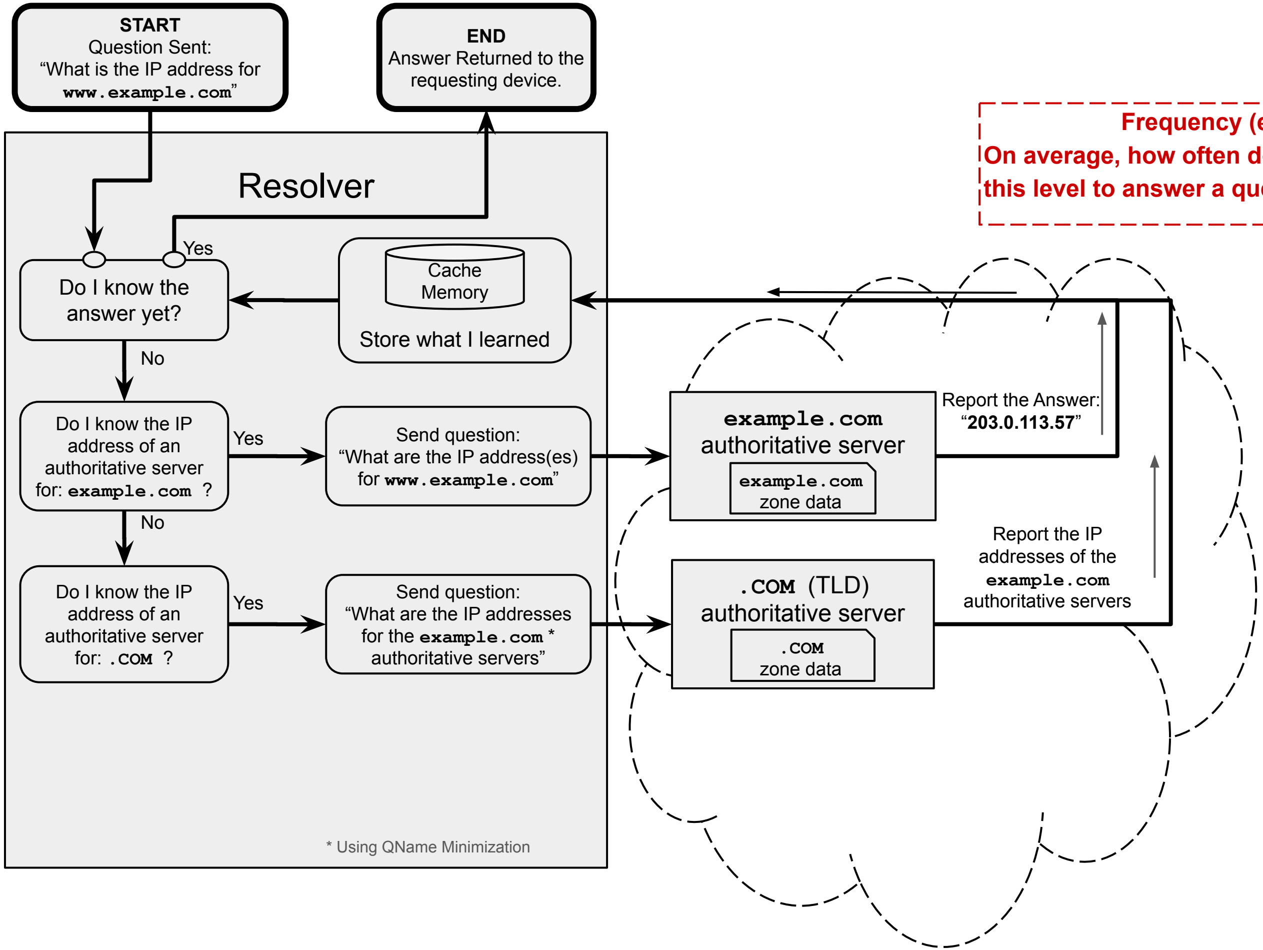

START
Question Sent:
"What is the IP address for www.example.com"

END
Answer Returned to the requesting device.

Resolver

Do I know the answer yet?

Yes

No

Cache Memory

Store what I learned

Do I know the IP address of an authoritative server for: example.com ?

Yes

No

Send question:
"What are the IP address(es) for www.example.com"

Do I know the IP address of an authoritative server for: .COM ?

Yes

Send question:
"What are the IP addresses for the example.com * authoritative servers"

* Using QName Minimization

example.com authoritative server

example.com zone data

Report the Answer:
"**203.0.113.57**"

.COM (TLD) authoritative server

.COM zone data

Report the IP addresses of the example.com authoritative servers

Frequency (estimates):
On average, how often do Resolvers consult at this level to answer a question?

Routine:
> 90% of answers are returned needing cache memory only

Ocasional:
~ 5% of answers require a question to the domain name's authoritative server

Uncommon:
~ 2% of answers require a question to the TLD authoritative server

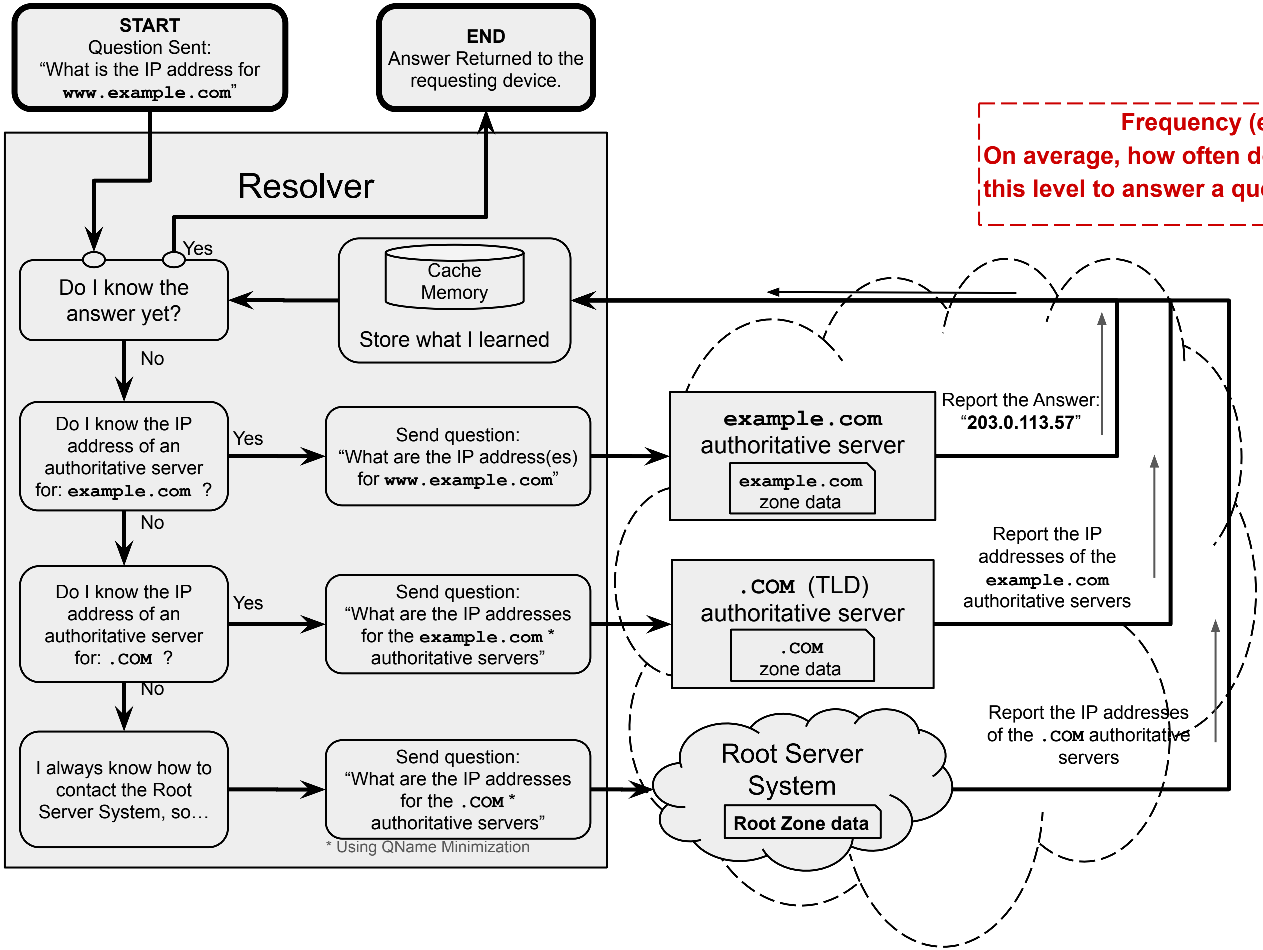

START
Question Sent:
"What is the IP address for `www.example.com`"

END
Answer Returned to the requesting device.

Resolver

Do I know the answer yet?

Yes

No

Cache Memory

Store what I learned

Do I know the IP address of an authoritative server for: `example.com` ?

Yes

Send question:
"What are the IP address(es) for `www.example.com`"

No

Do I know the IP address of an authoritative server for: `.COM` ?

Yes

Send question:
"What are the IP addresses for the `example.com`* authoritative servers"

No

I always know how to contact the Root Server System, so…

Send question:
"What are the IP addresses for the `.COM`* authoritative servers"

* Using QName Minimization

`example.com` authoritative server

`example.com` zone data

Report the Answer:
"**203.0.113.57**"

`.COM` (TLD) authoritative server

`.COM` zone data

Report the IP addresses of the `example.com` authoritative servers

Root Server System

Root Zone data

Report the IP addresses of the `.COM` authoritative servers

Frequency (estimates):
On average, how often do Resolvers consult at this level to answer a question?

Routine:
> 90% of answers are returned needing cache memory only

Ocasional:
~ 5% of answers require a question to the domain name's authoritative server

Uncommon:
~ 2% of answers require a question to the TLD authoritative server

Rare:
~ 0.02% of answers require a question to the RSS

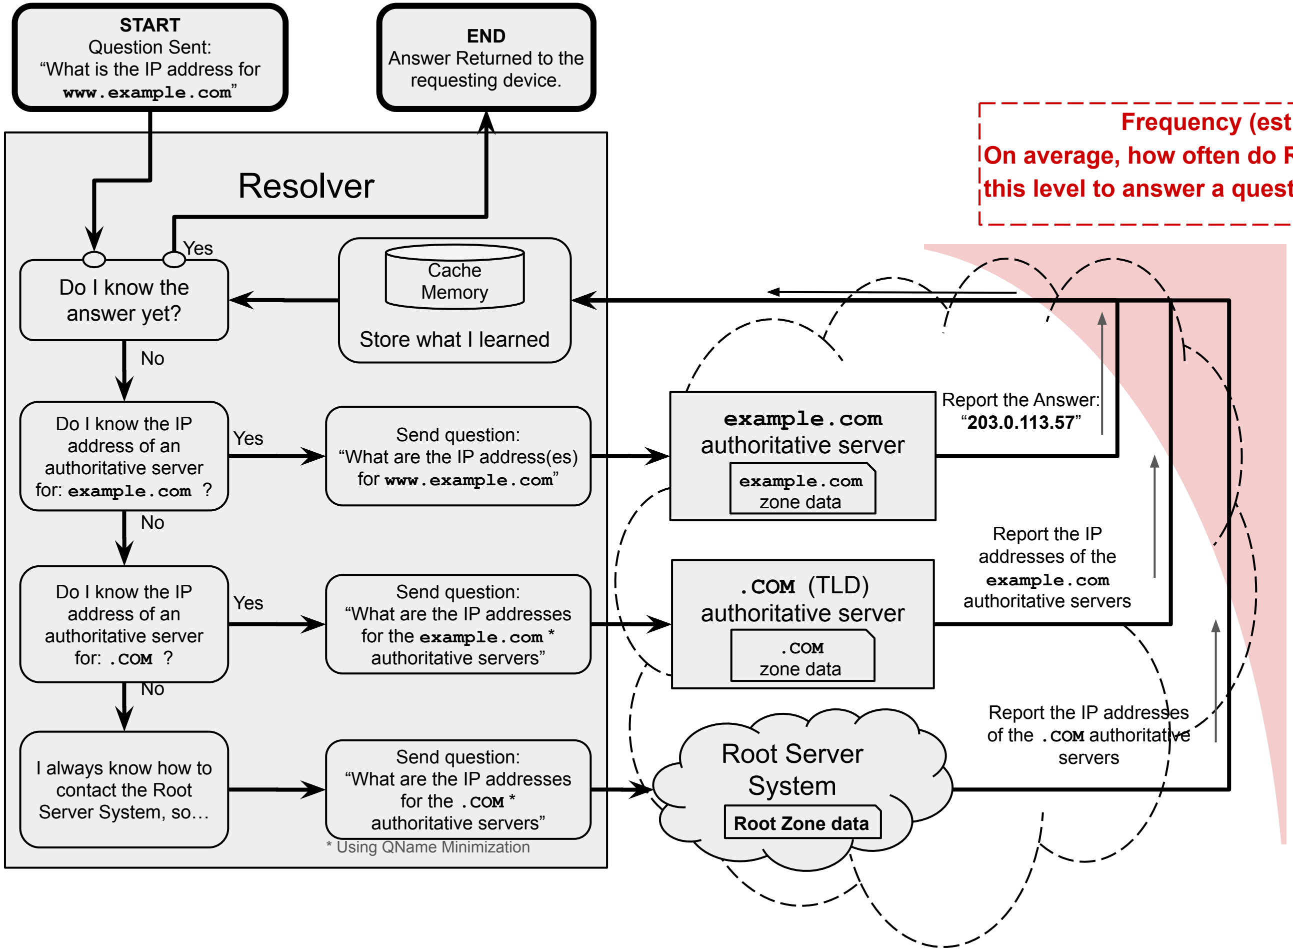

START
Question Sent:
“What is the IP address for **www.example.com**”

END
Answer Returned to the requesting device.

Resolver

Do I know the answer yet?

Yes

No

Cache Memory

Store what I learned

Do I know the IP address of an authoritative server for: **example.com** ?

Yes

Send question:
“What are the IP address(es) for **www.example.com**”

No

Do I know the IP address of an authoritative server for: **.COM** ?

Yes

Send question:
“What are the IP addresses for the **example.com** * authoritative servers”

No

I always know how to contact the Root Server System, so…

Send question:
“What are the IP addresses for the **.COM** * authoritative servers”

* Using QName Minimization

example.com authoritative server

example.com zone data

Report the Answer:
“**203.0.113.57**”

.COM (TLD) authoritative server

.COM zone data

Report the IP addresses of the **example.com** authoritative servers

Root Server System

Root Zone data

Report the IP addresses of the **.COM** authoritative servers

Frequency (estimates):
On average, how often do Resolvers consult at this level to answer a question?

Routine:
> 90% of answers are returned needing cache memory only

Ocasional:
~ 5% of answers require a question to the domain name’s authoritative server

Uncommon:
~ 2% of answers require a question to the TLD authoritative server

Rare:
~ 0.02% of answers require a question to the RSS

The Root Zone holds addresses for less than 0.00005% of the world’s addressable resources

DNS Layer	Number of unique zones	Typical number of resource addresses	Maintained by
Domain name zone data	350,000,000	Varies Each [`www.`__], [`mail.`__], etc.	The domain name registrant
TLD zones	1,700	1,000 - 10,000,000 domains	The TLD registry
Root Zone	1 (one)	1,700 TLDs	IANA/RZM

In review

- A root server holds a copy of “Root Zone” data
  The Root Zone holds addresses for TLD’s like:
  - .com
  - .nl
  - .jobs (and on and on)
- A TLD’s authoritative server knows the address for the next step
  - All names that end in .com, like amazon.com or tiktok.com
  - All names that end in .nl, like google.nl or amsterdam.nl
  - All names that end in .jobs, like tech.jobs or highpay.jobs
- A domain name’s authoritative server knows
  - The answer to the question about www.amazon.com or mail.amazon.com or info.amazon.com
- The resolver finds and returns the answer

In the millisecond world of a resolver, queries to the Root Server System are rare.

Root Server System Operation

- Massively redundant 1700+ globally distributed server instances
  - Each server instance holds 100% of the Root Zone content
  - Diverse hardware platforms
  - Diverse operating systems
  - Diverse DNS applications
  - Diverse data routing

- Result: No single point of technological failure

Root Server System Operation

- Co-operated by 12 autonomous Root Server Operators (RSO)
  - Each RSO is independent of the others
  - The RSOs collaborate continuously with one another
  - Force majeure event suffered by one (court injunction, etc) has no operational impact on the others

- Result: No single point of institutional failure

Root Server Operators do not choose the content of Root Zone data

- Where does zone data come from?
  - Registrants maintain the zone data for their own domain
  - Registrants provide their authoritative server addresses to TLD registries, via registrars
  - TLD registries provide their authoritative server addresses to IANA for inclusion in the root zone
  - IANA authenticates and sends root zone data changes to the Root Zone Maintainer (RZM)
  - The RZM generates encrypted signatures and makes the root zone data available in the RSS by transmitting it to the RSOs
- The RSOs serve up what IANA sends

40 years of stability, security, and resilience

- The Root Server System has operated since the 1980’s
- It has never suffered a service outage.
    - DDoS attackers have tried; they failed, by design

Summary

- The root server system is an important, if infrequent, component of address resolution
  - Most DNS queries are answered from cache memory
  - Most remaining DNS queries go straight to domain name authoritative servers
- Root server operators do not decide the content of the Root Zone
- The root server system
  - Is massively redundant
  - Is technologically diverse
  - Is institutionally resilient
- The root server system works