Root zone update for TLD managers

Mexico City, Mexico
March 2009

Kim Davies
Manager, Root Zone Services

ICANN
Internet Corporation for
Assigned Names & Numbers

A quick census

280

delegated

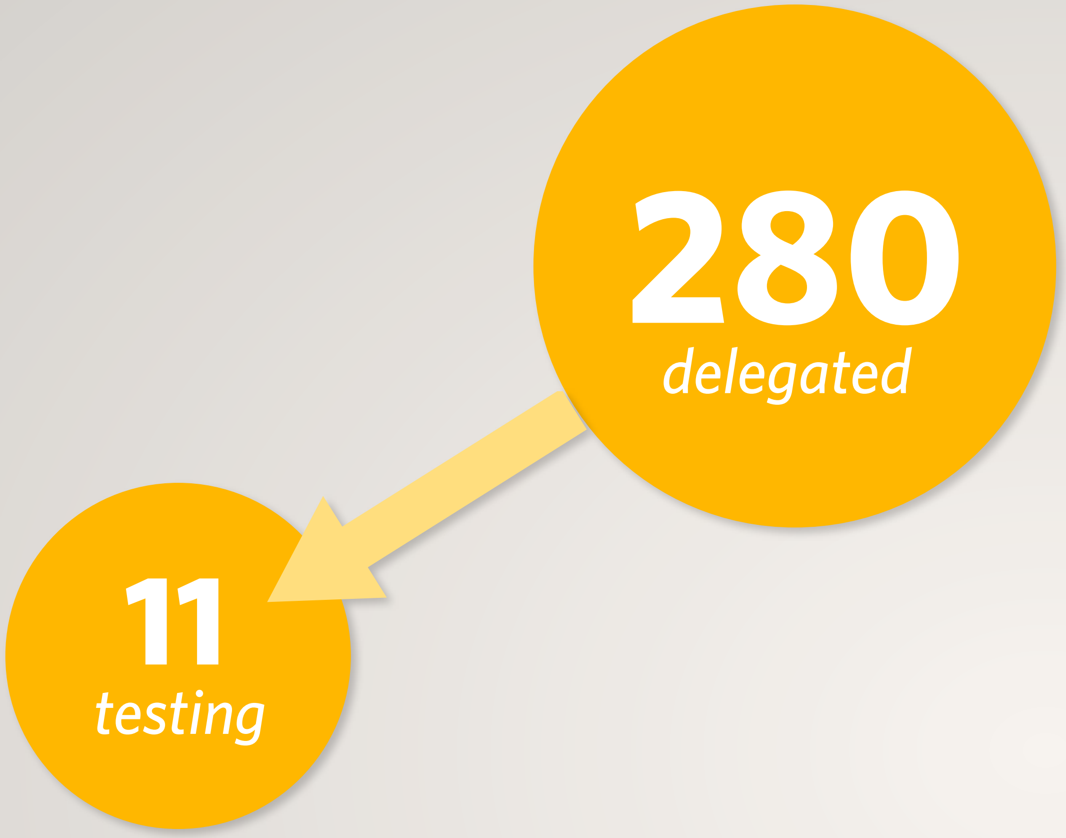

280
delegated

11
testing

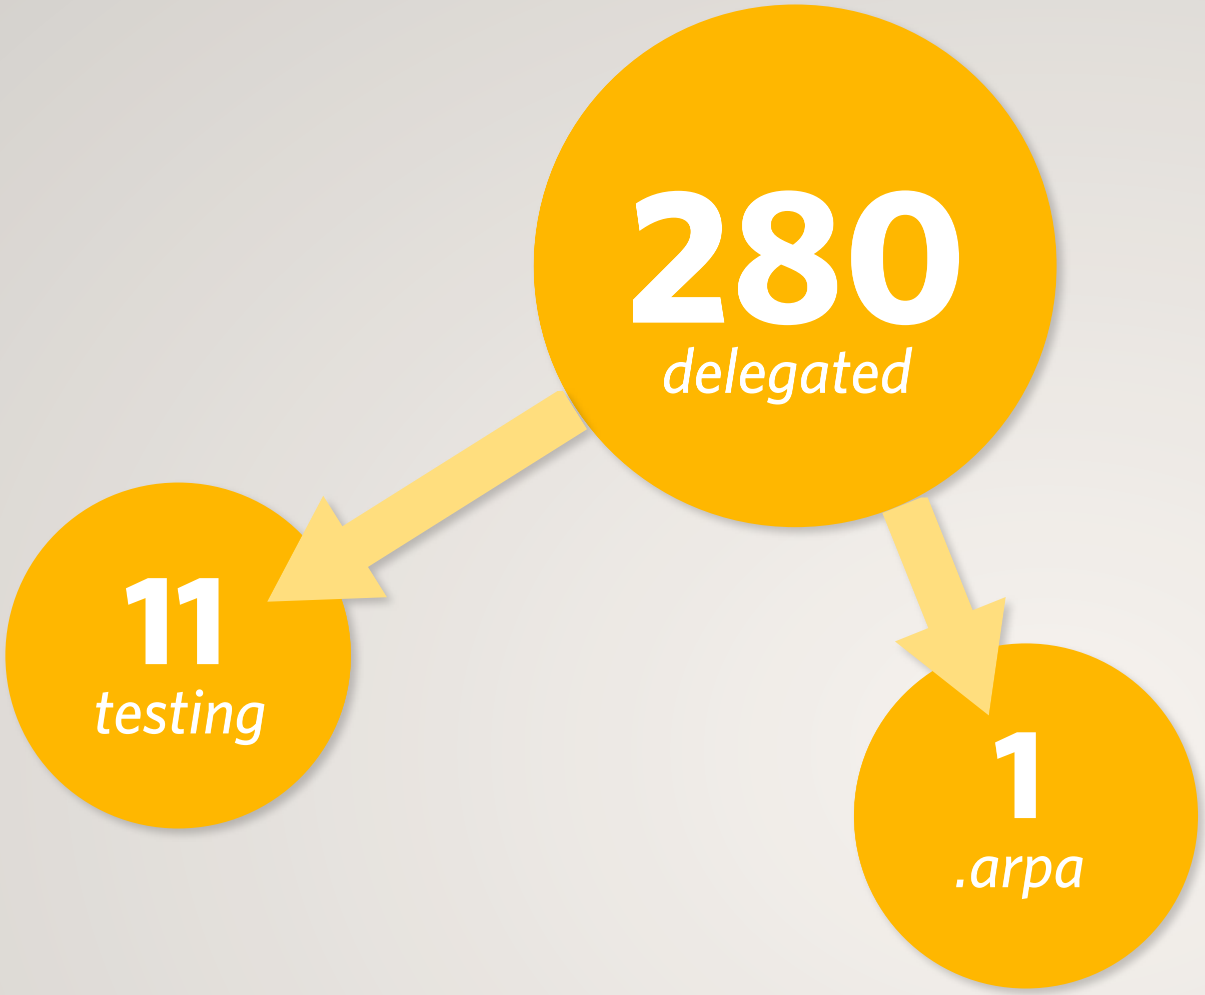

280
delegated

11
testing

1
.arpa

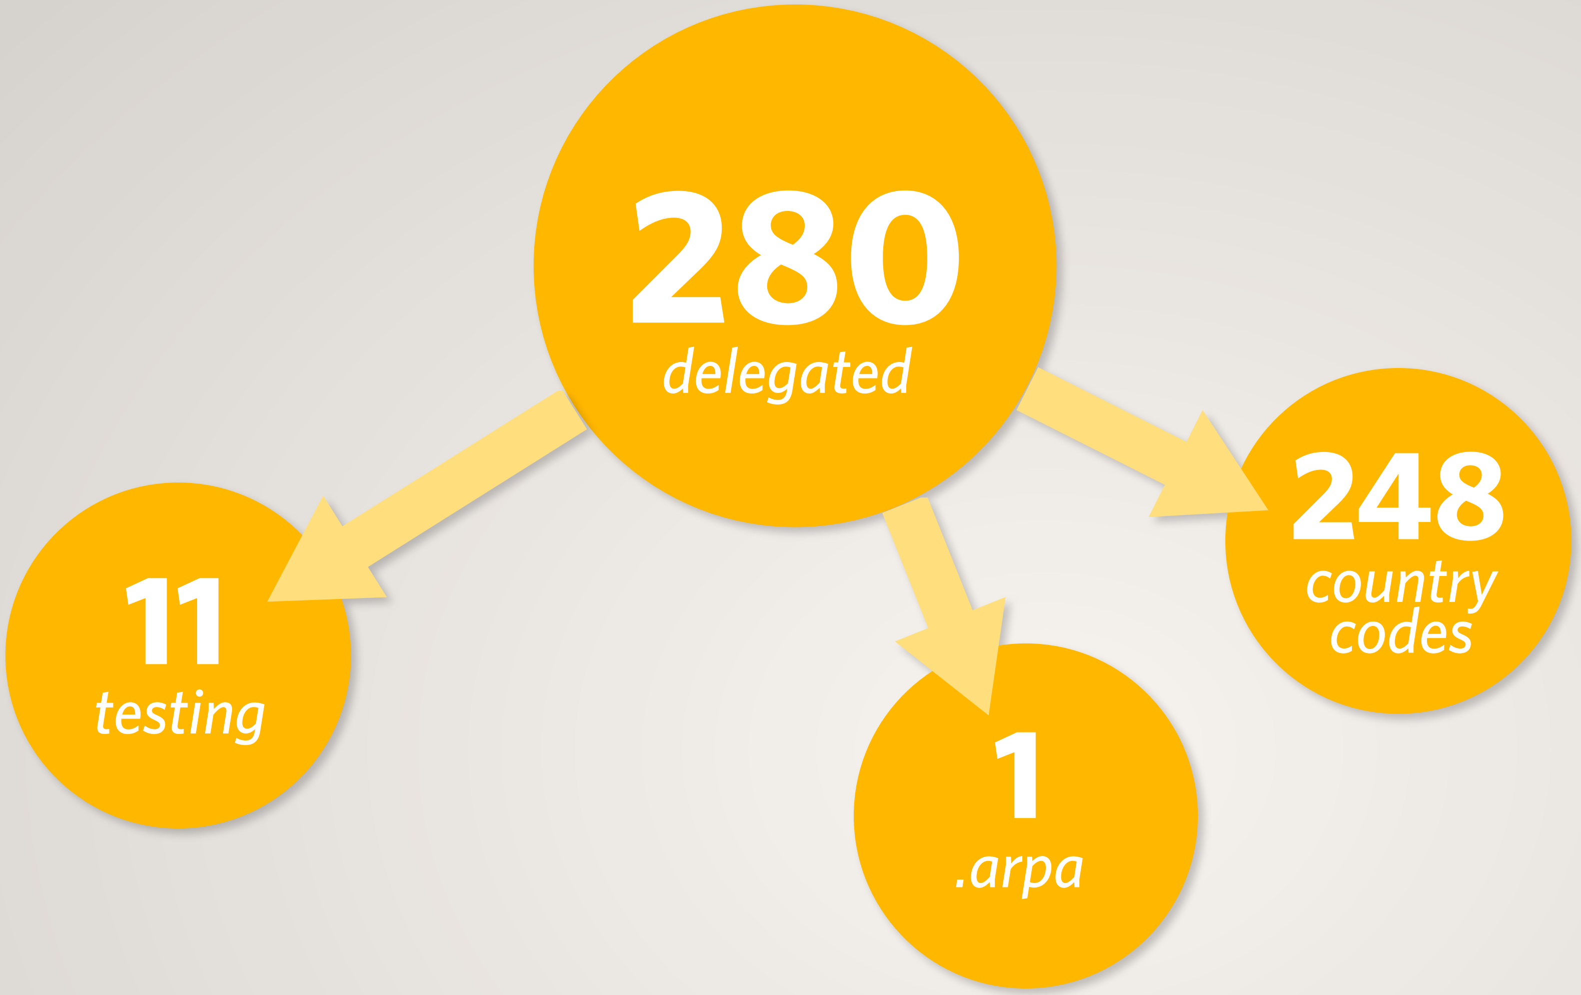

280
delegated

11
testing

1
.arpa

248
country
codes

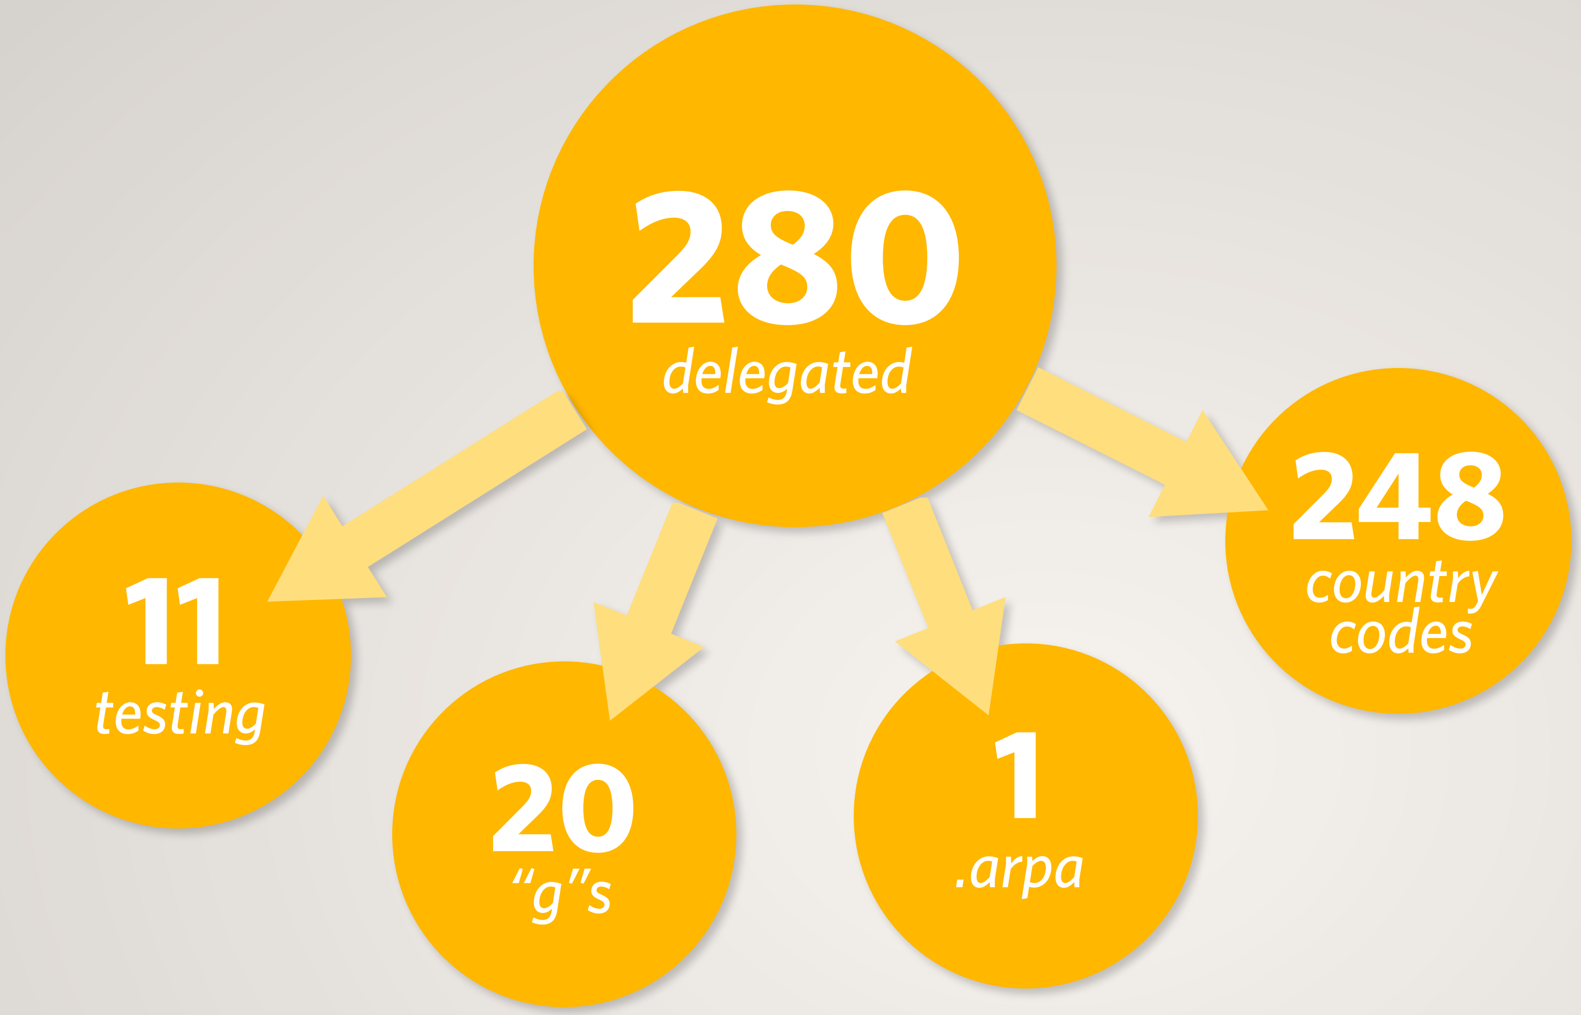

280
delegated

11
testing

20
“g”s

1
.arpa

248
country codes

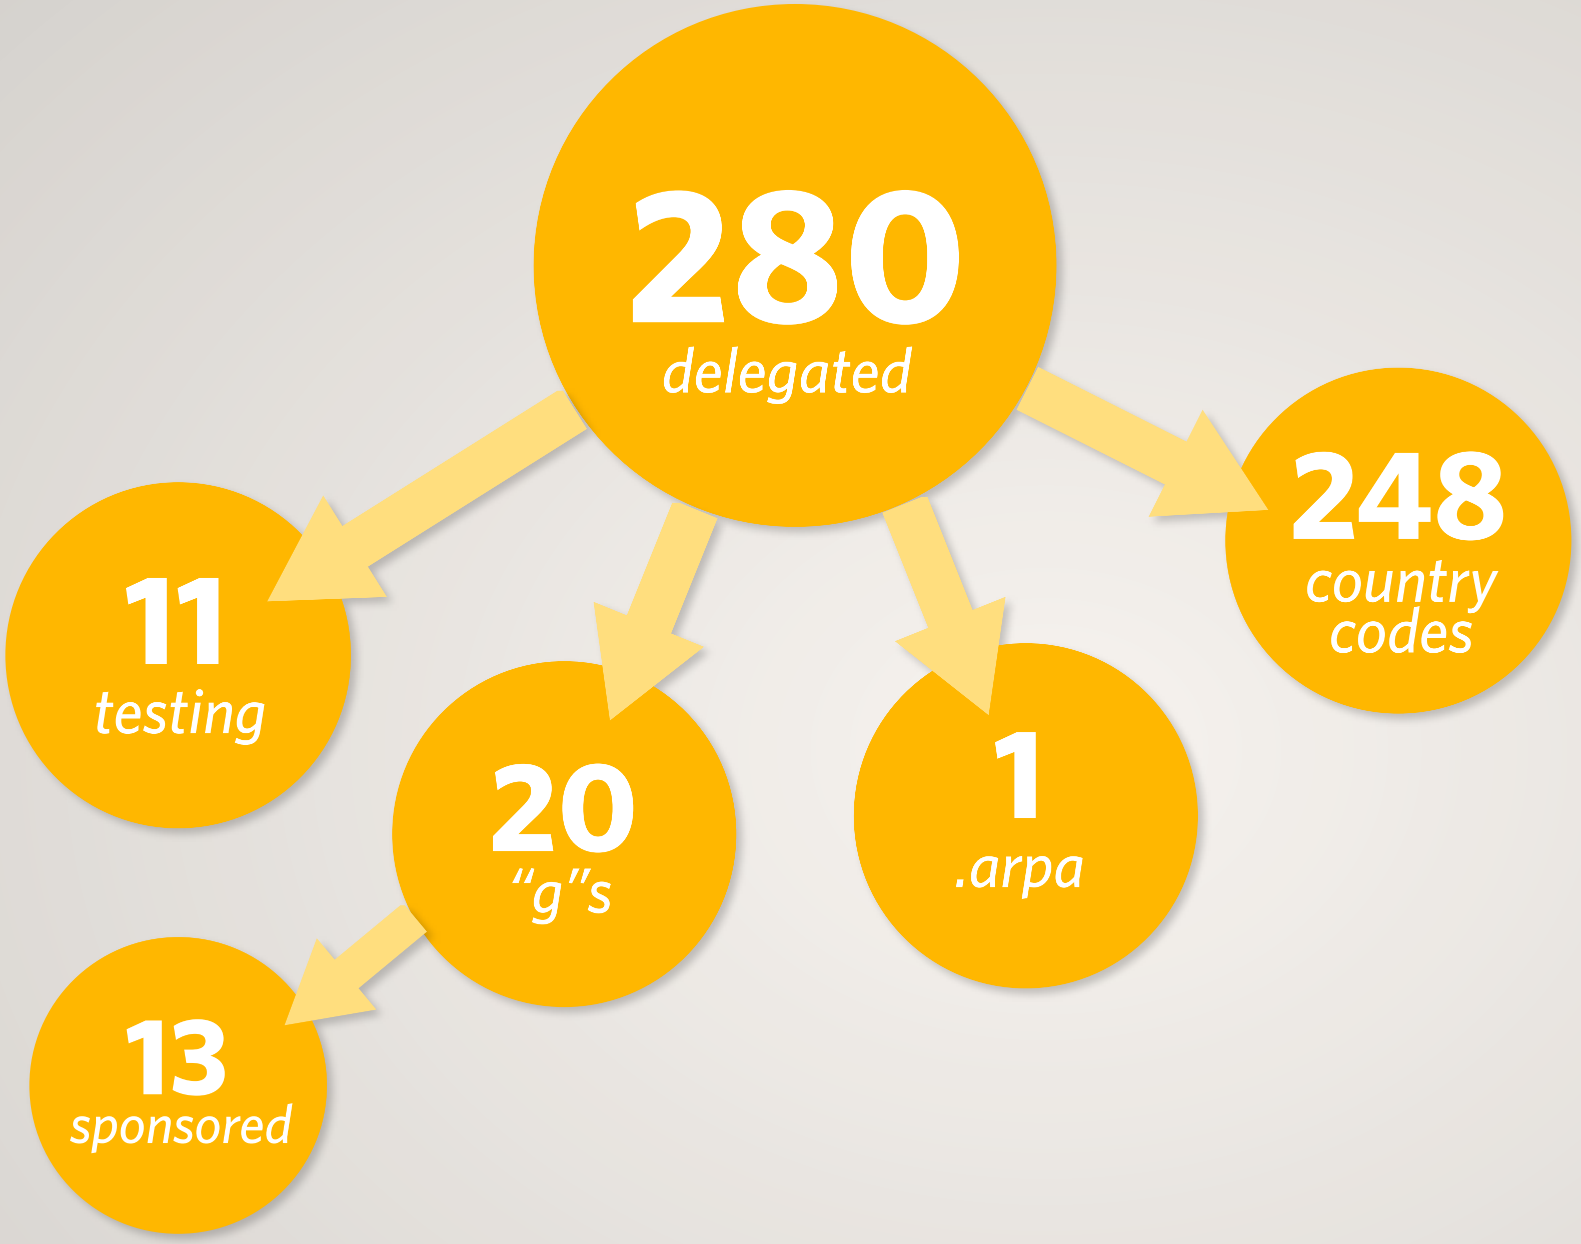

280 delegated

11 testing

20 "g"s

13 sponsored

1 .arpa

248 country codes

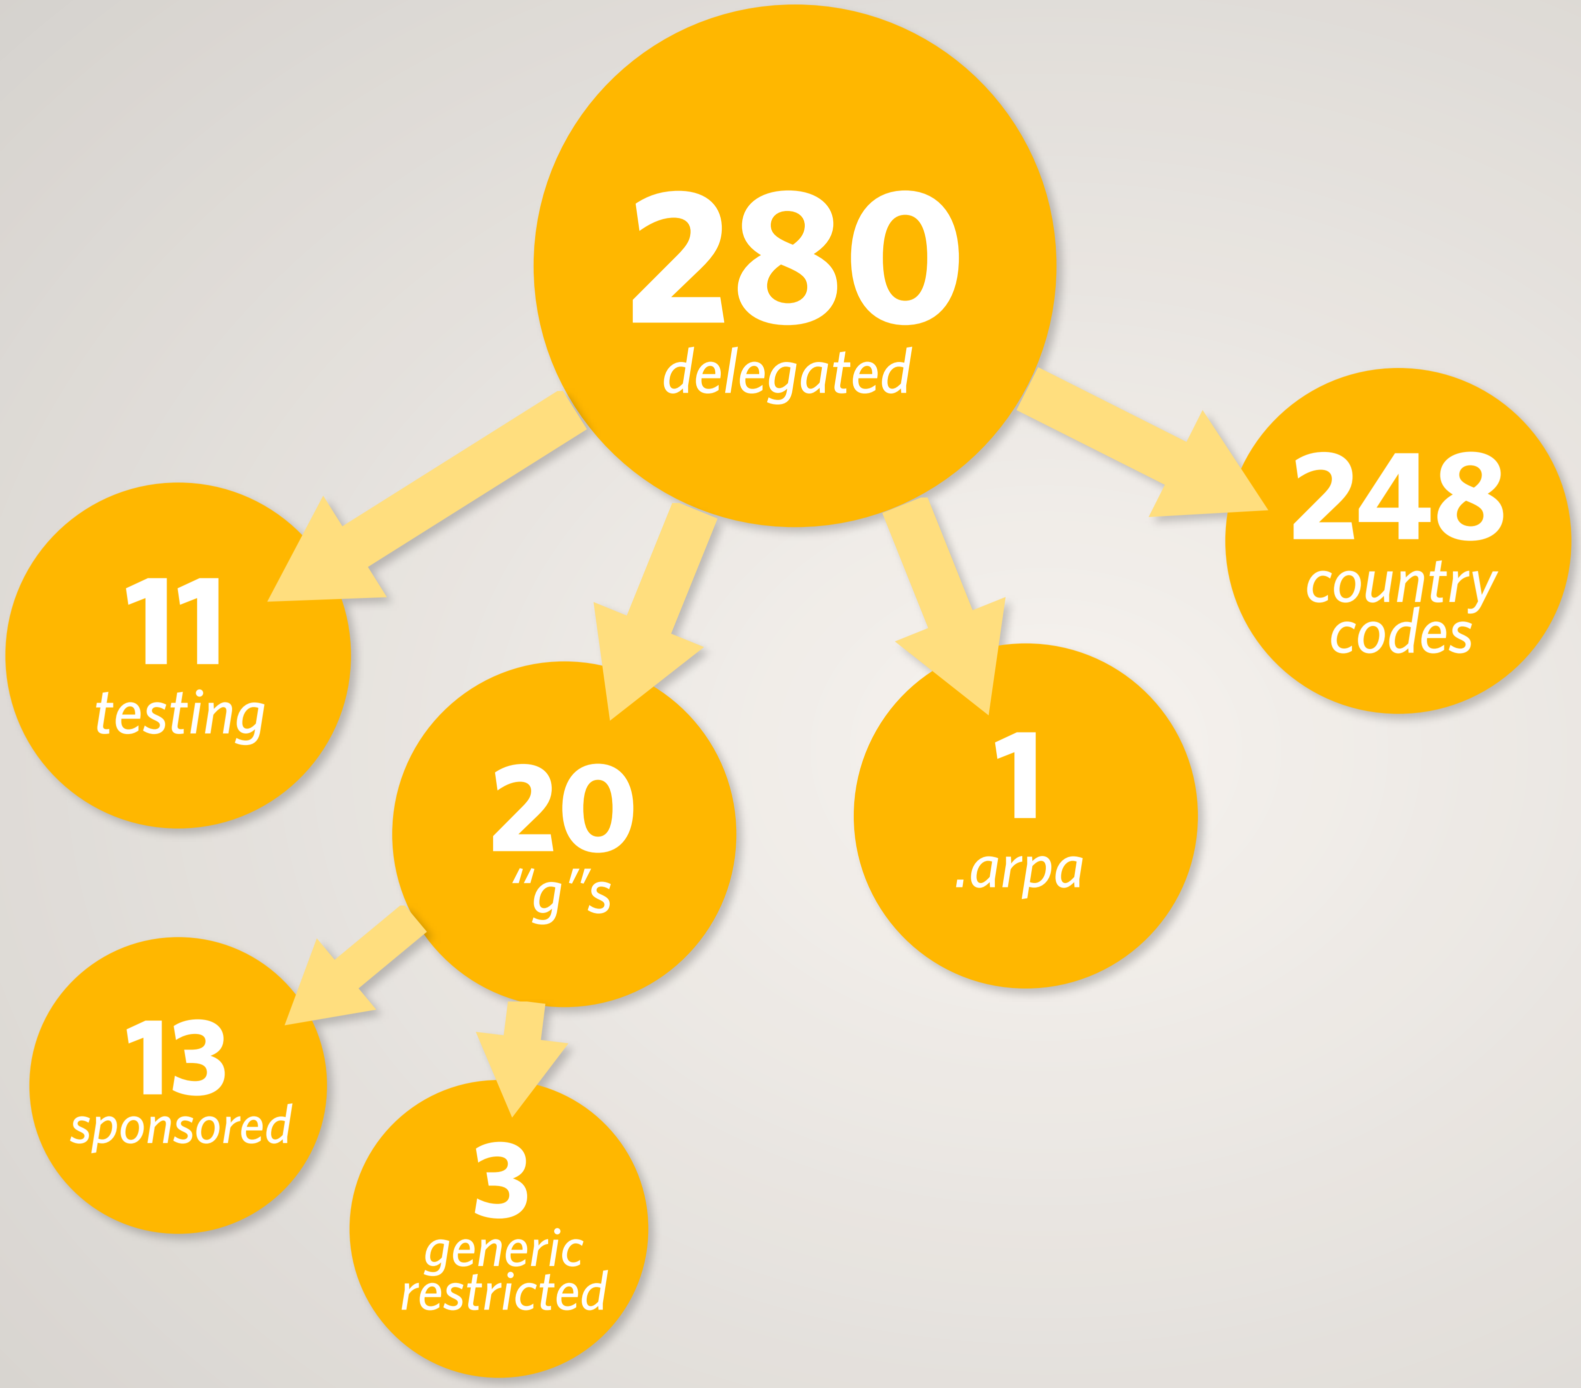

280 *delegated*

11 *testing*

20 *“g”s*

13 *sponsored*

3 *generic restricted*

1 *.arpa*

248 *country codes*

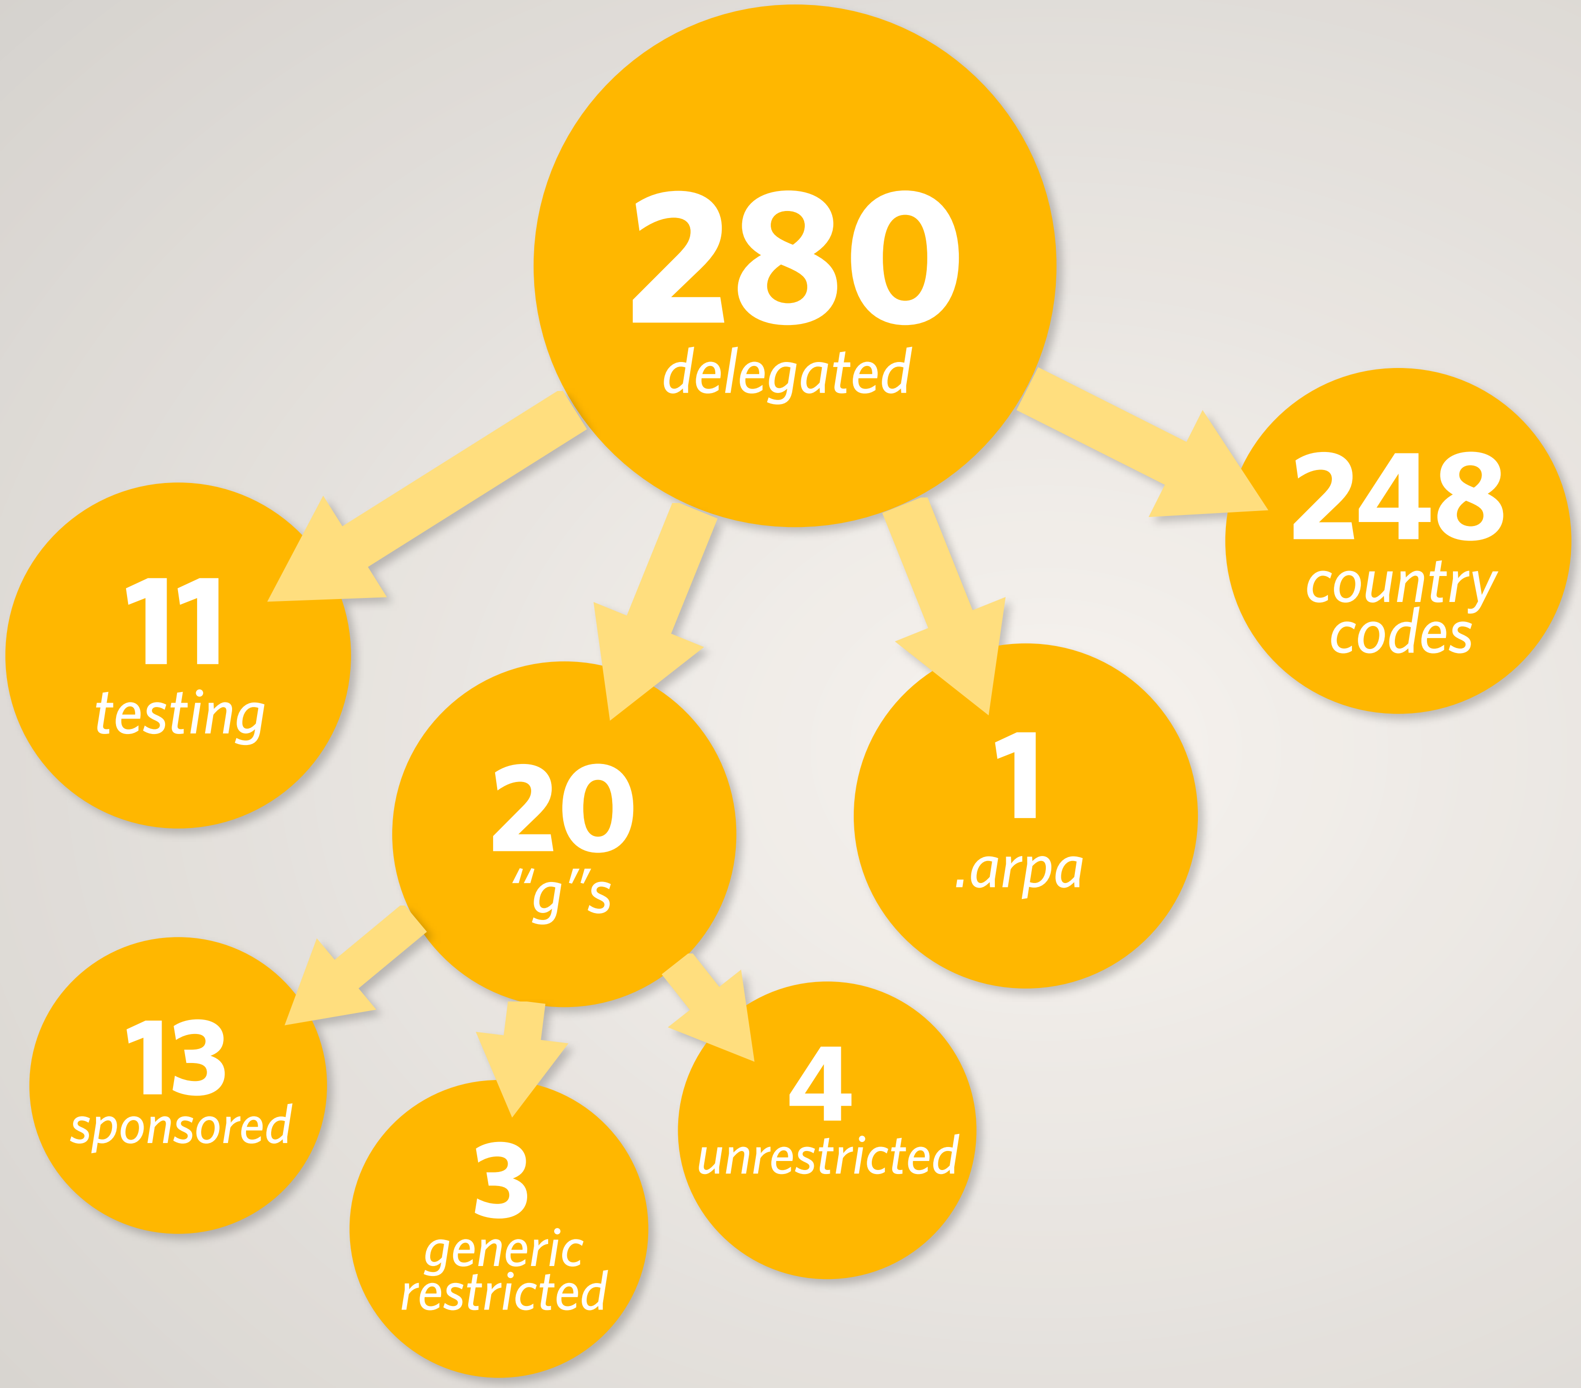

280 delegated

11 testing

20 “g”s

1 .arpa

248 country codes

13 sponsored

3 generic restricted

4 unrestricted

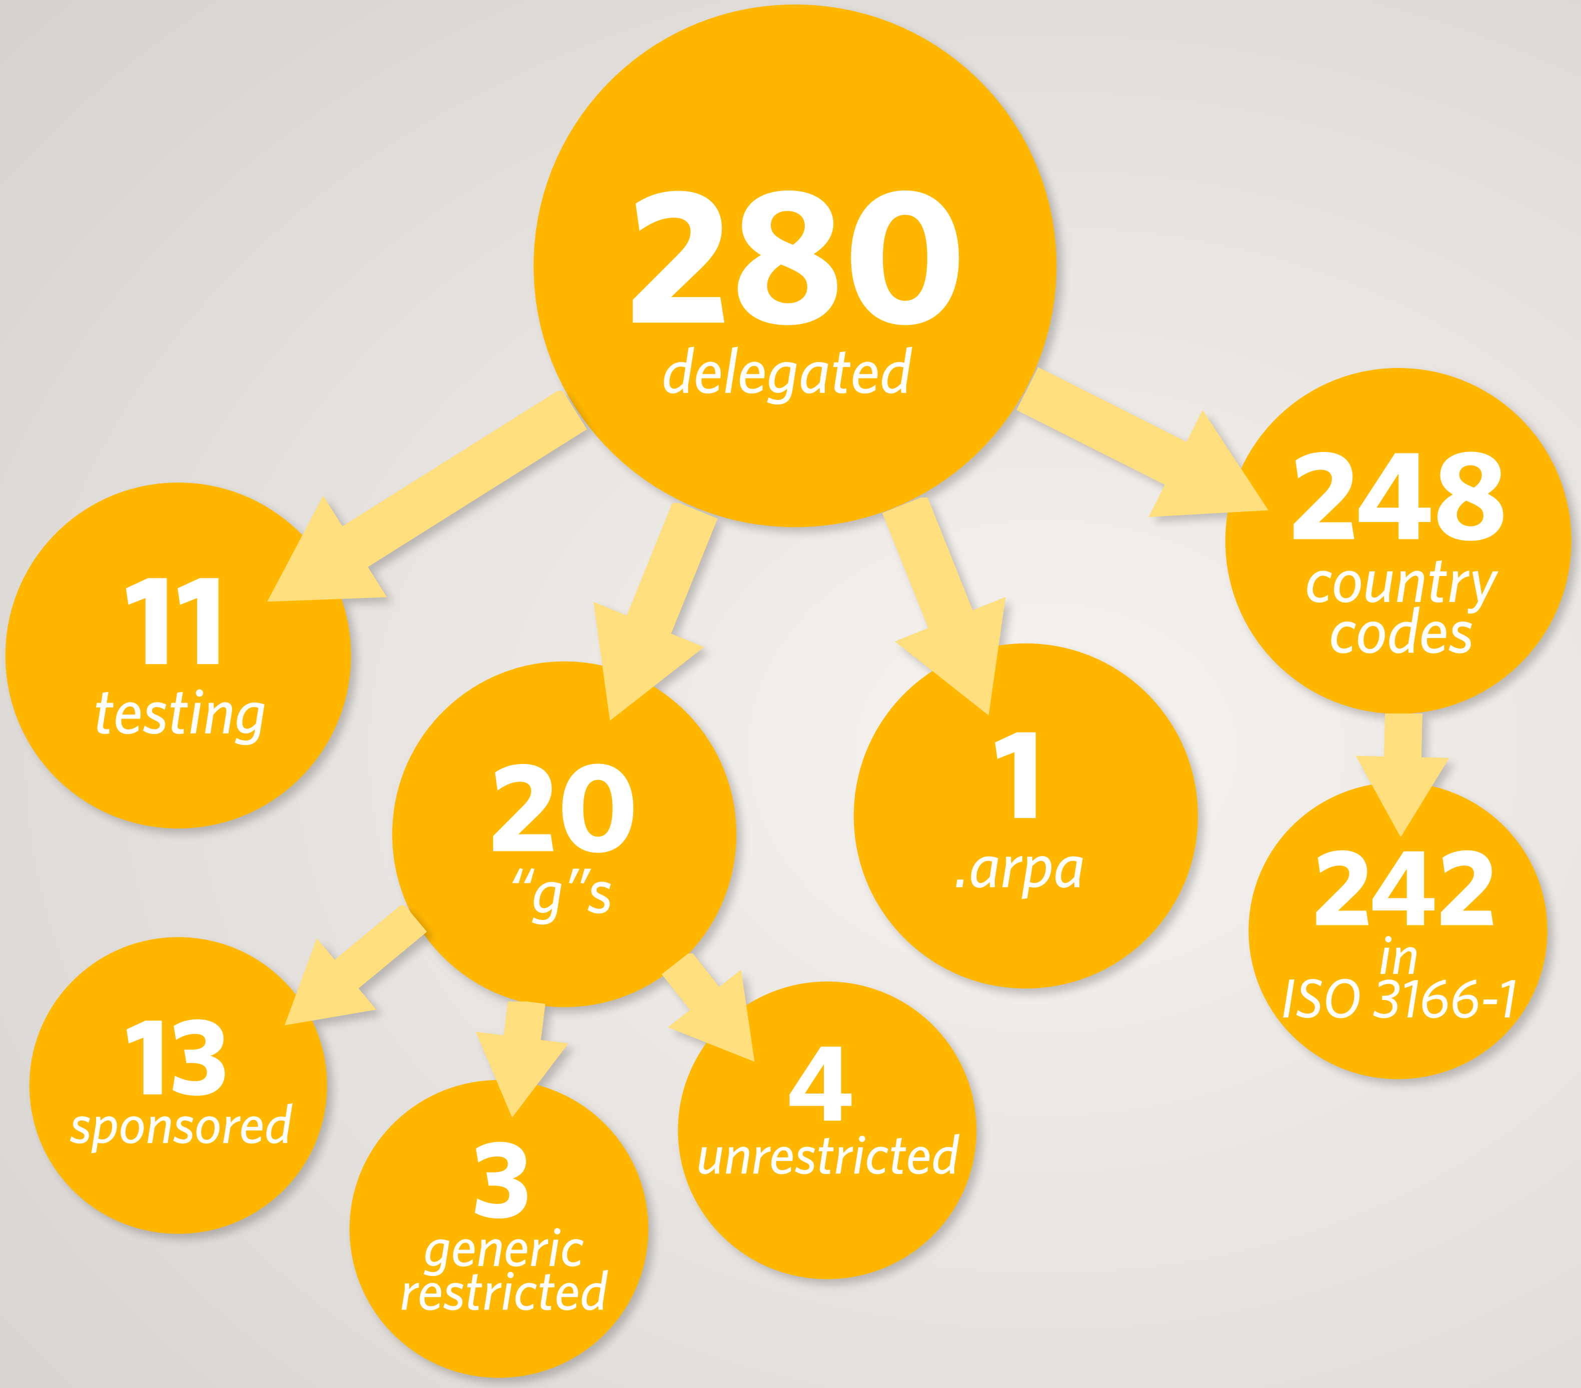

280 delegated

11 testing

20 "g"s

13 sponsored

3 generic restricted

4 unrestricted

1 .arpa

248 country codes

242 in ISO 3166-1

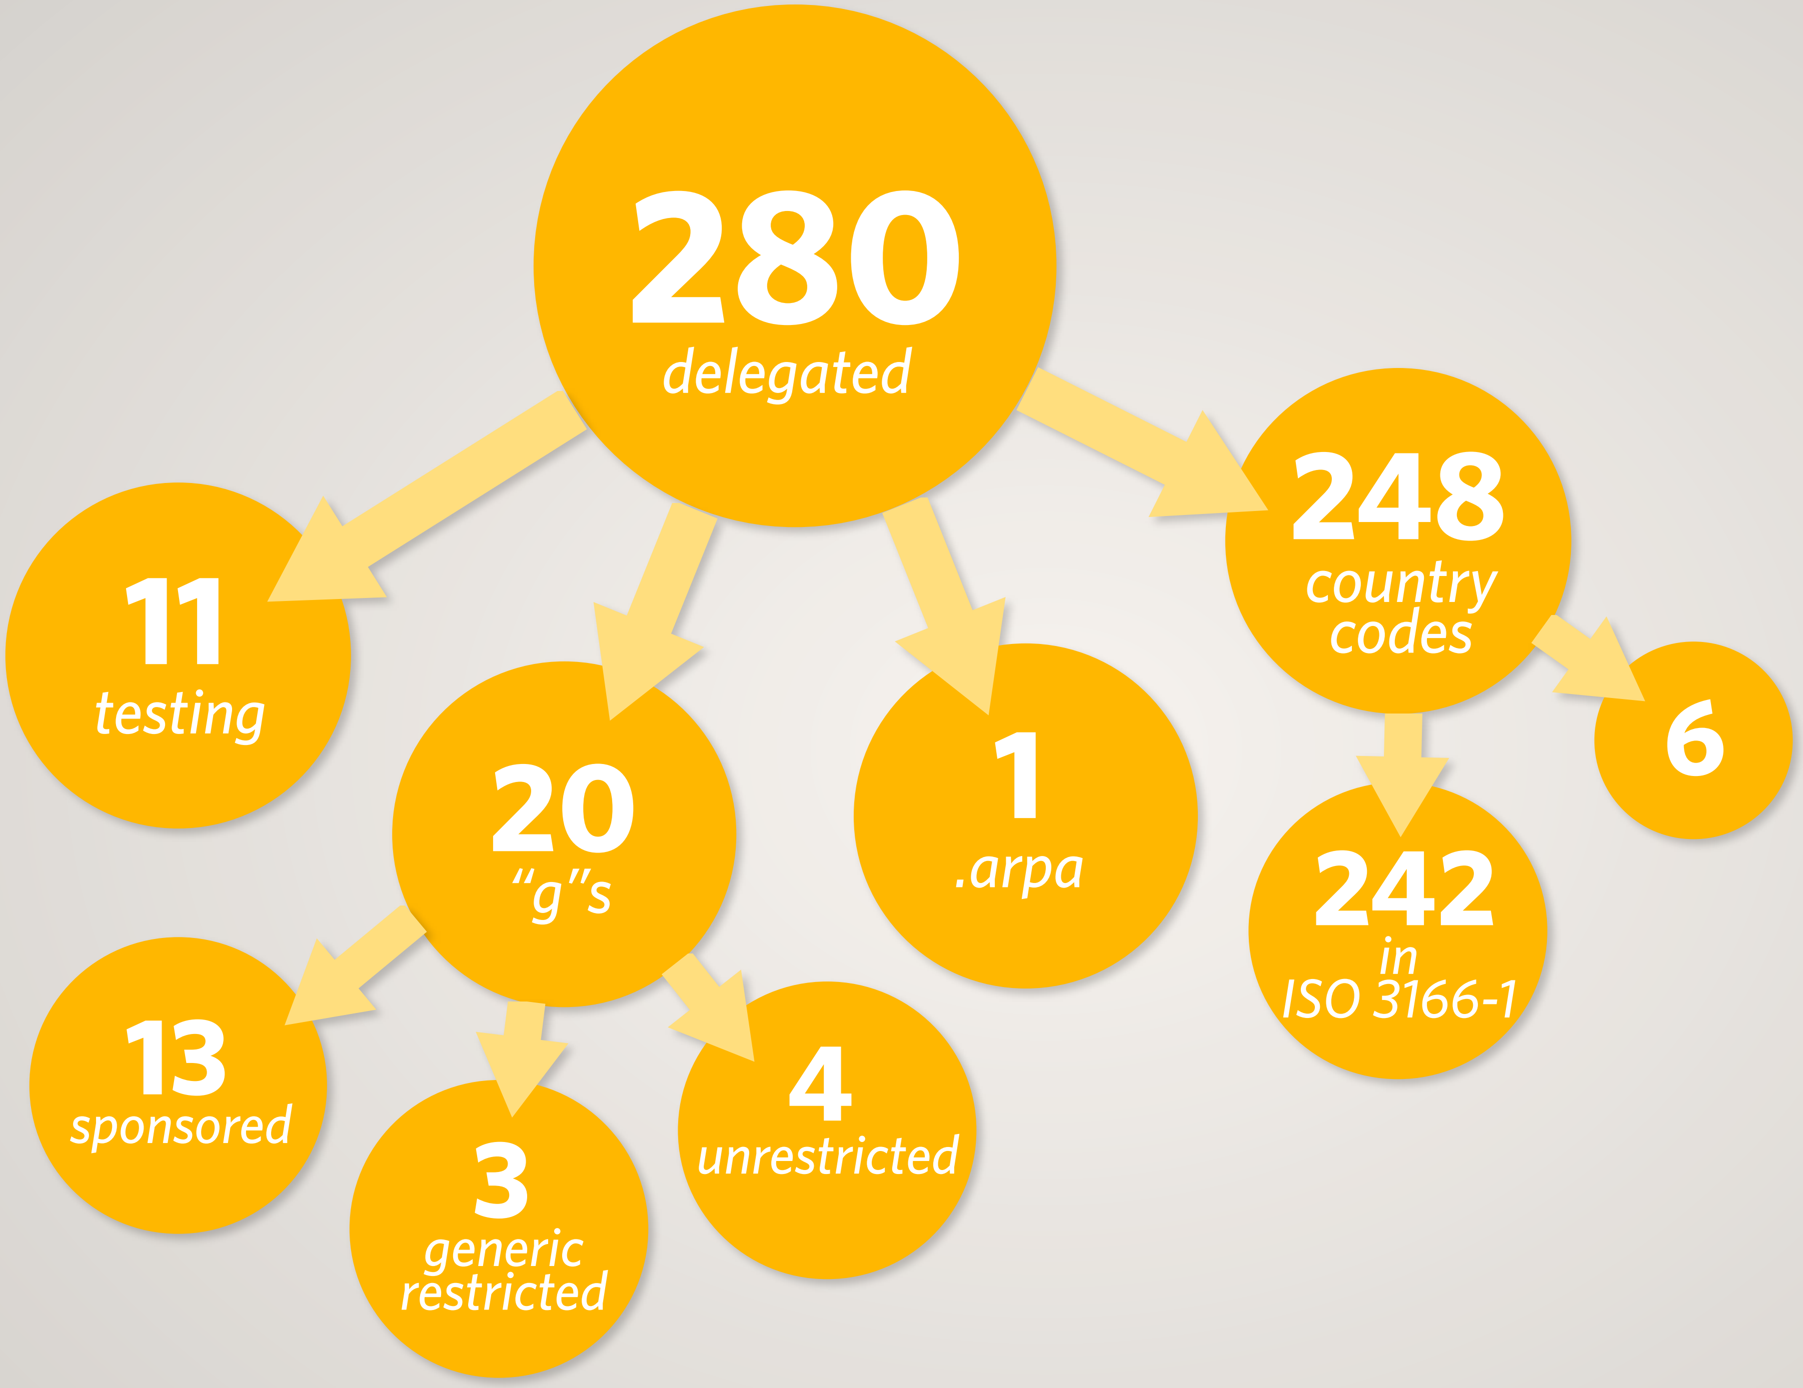

280 delegated

11 testing

20 “g”s

13 sponsored

3 generic restricted

4 unrestricted

1 .arpa

248 country codes

242 in ISO 3166-1

6

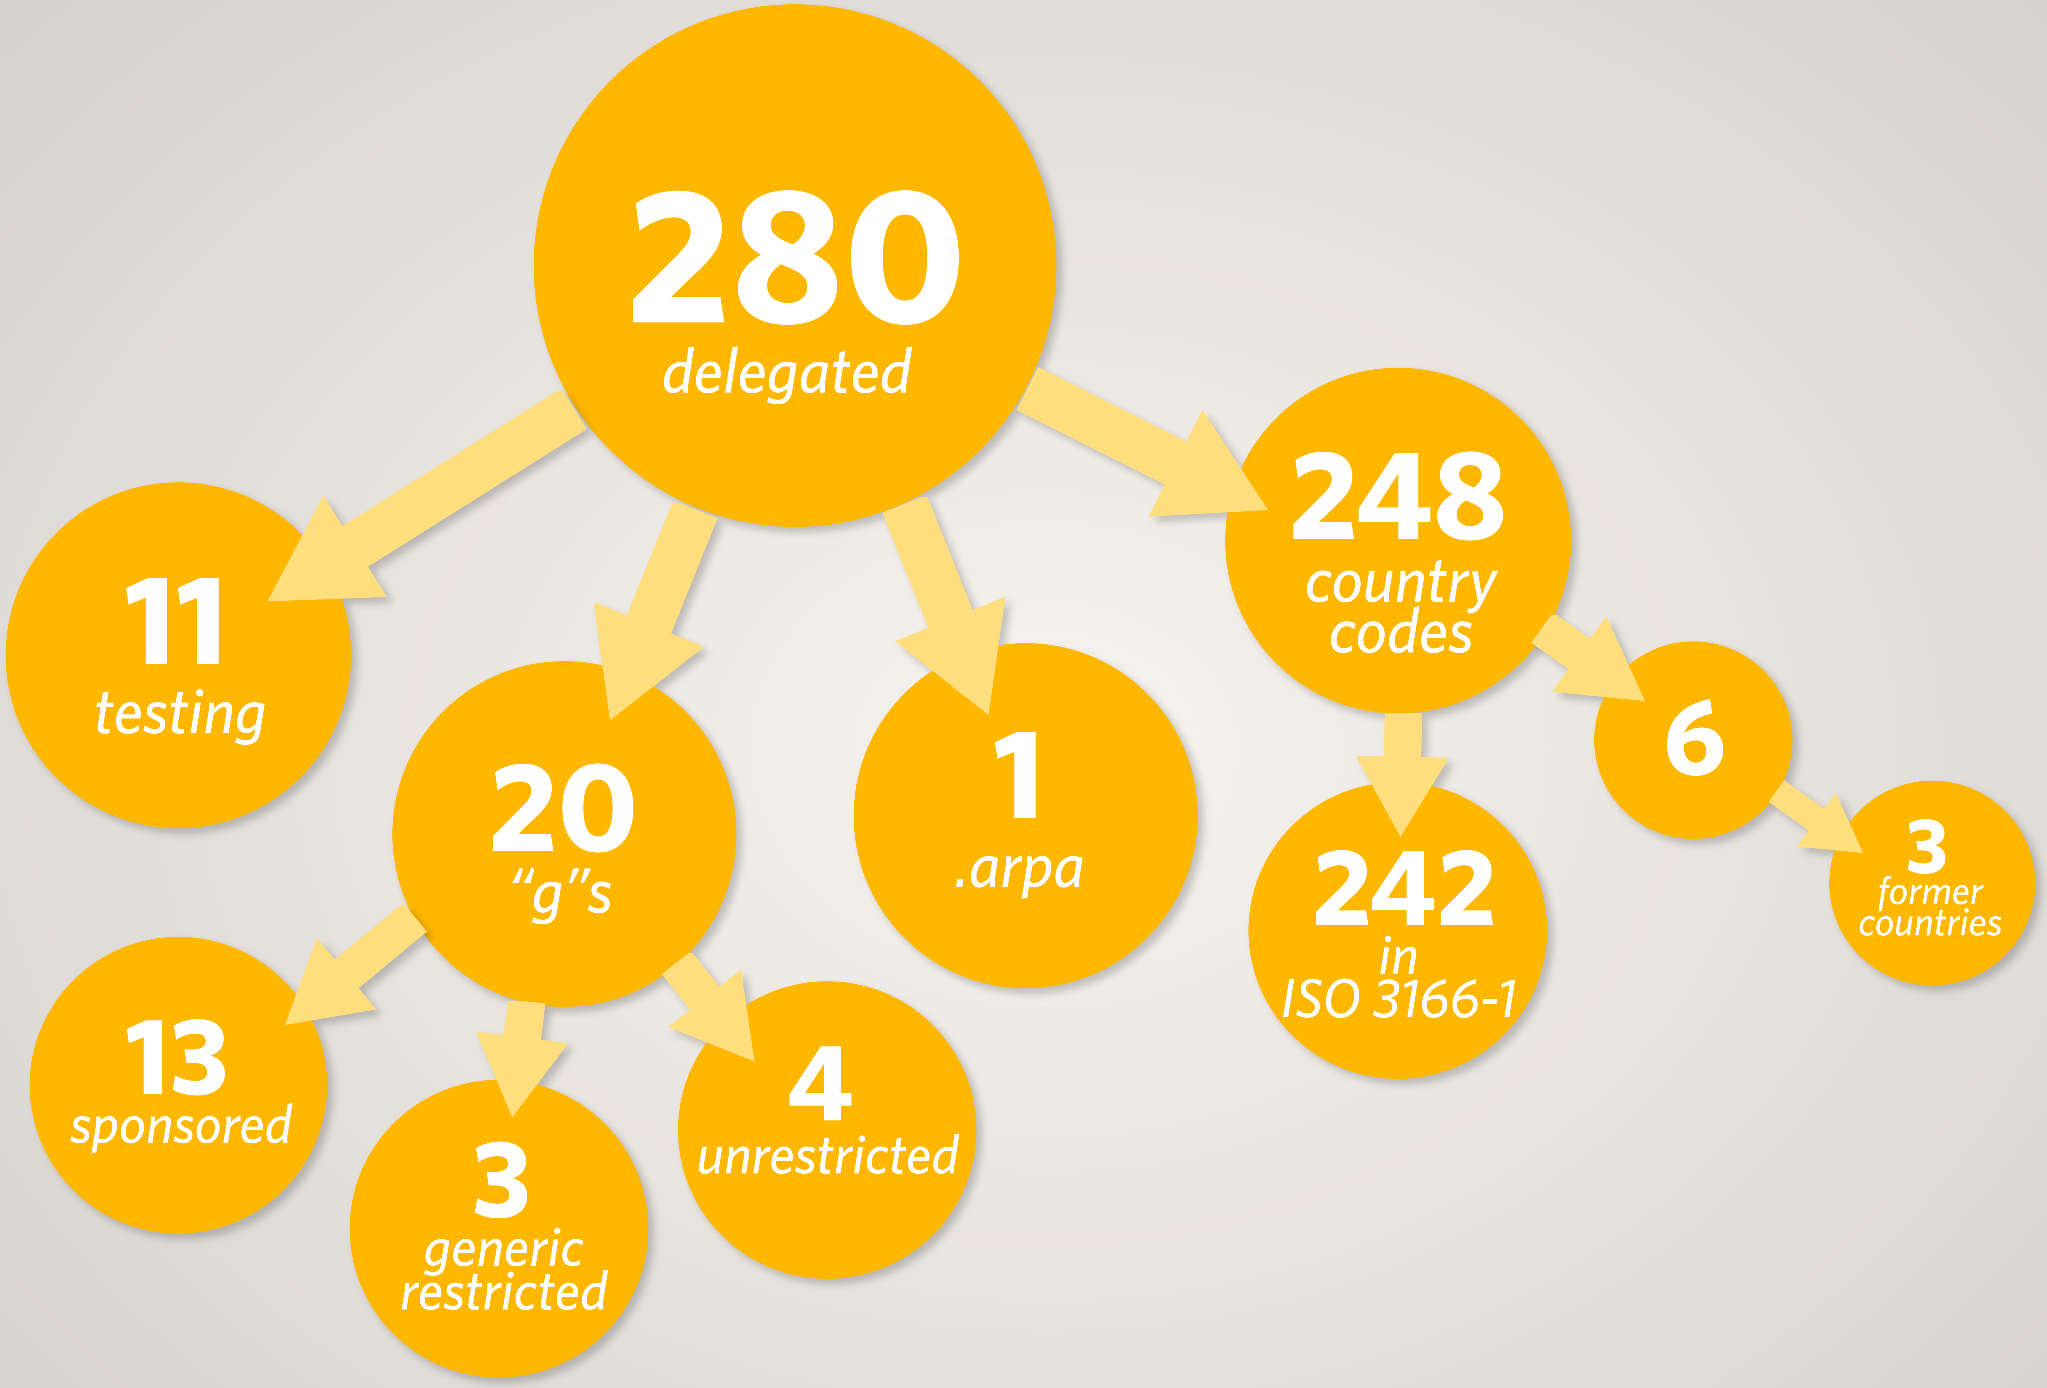

280
delegated

11
testing

20
"g"s

13
sponsored

3
generic
restricted

4
unrestricted

1
.arpa

248
country
codes

242
in
ISO 3166-1

6

3
former
countries

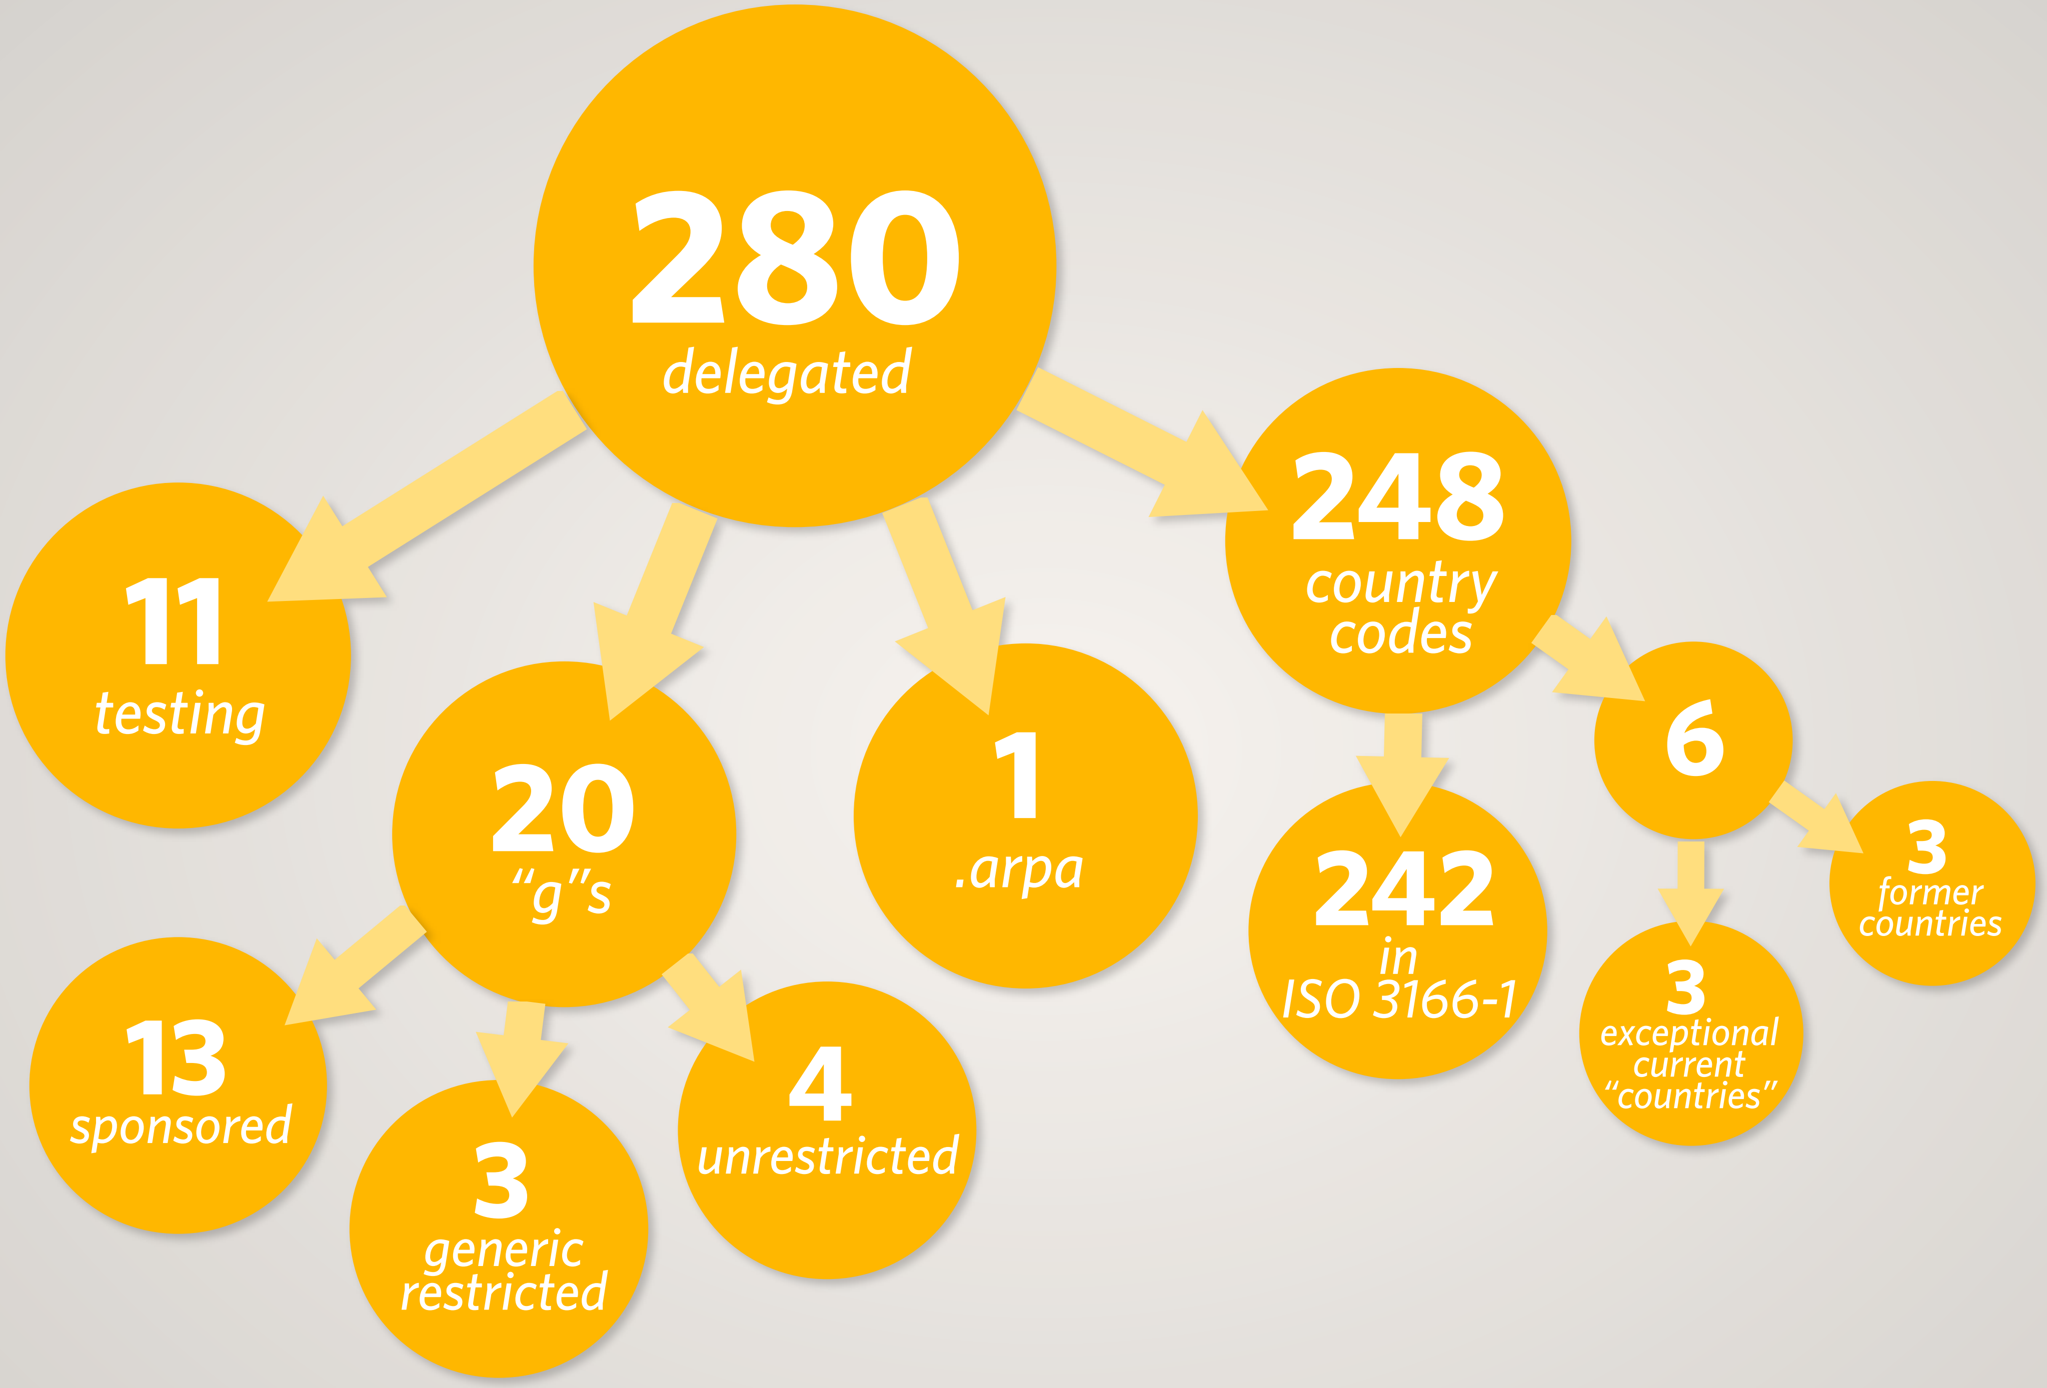

280 delegated

11 testing

20 “g”s

13 sponsored

3 generic restricted

4 unrestricted

1 .arpa

248 country codes

242 in ISO 3166-1

6

3 exceptional current “countries”

3 former countries

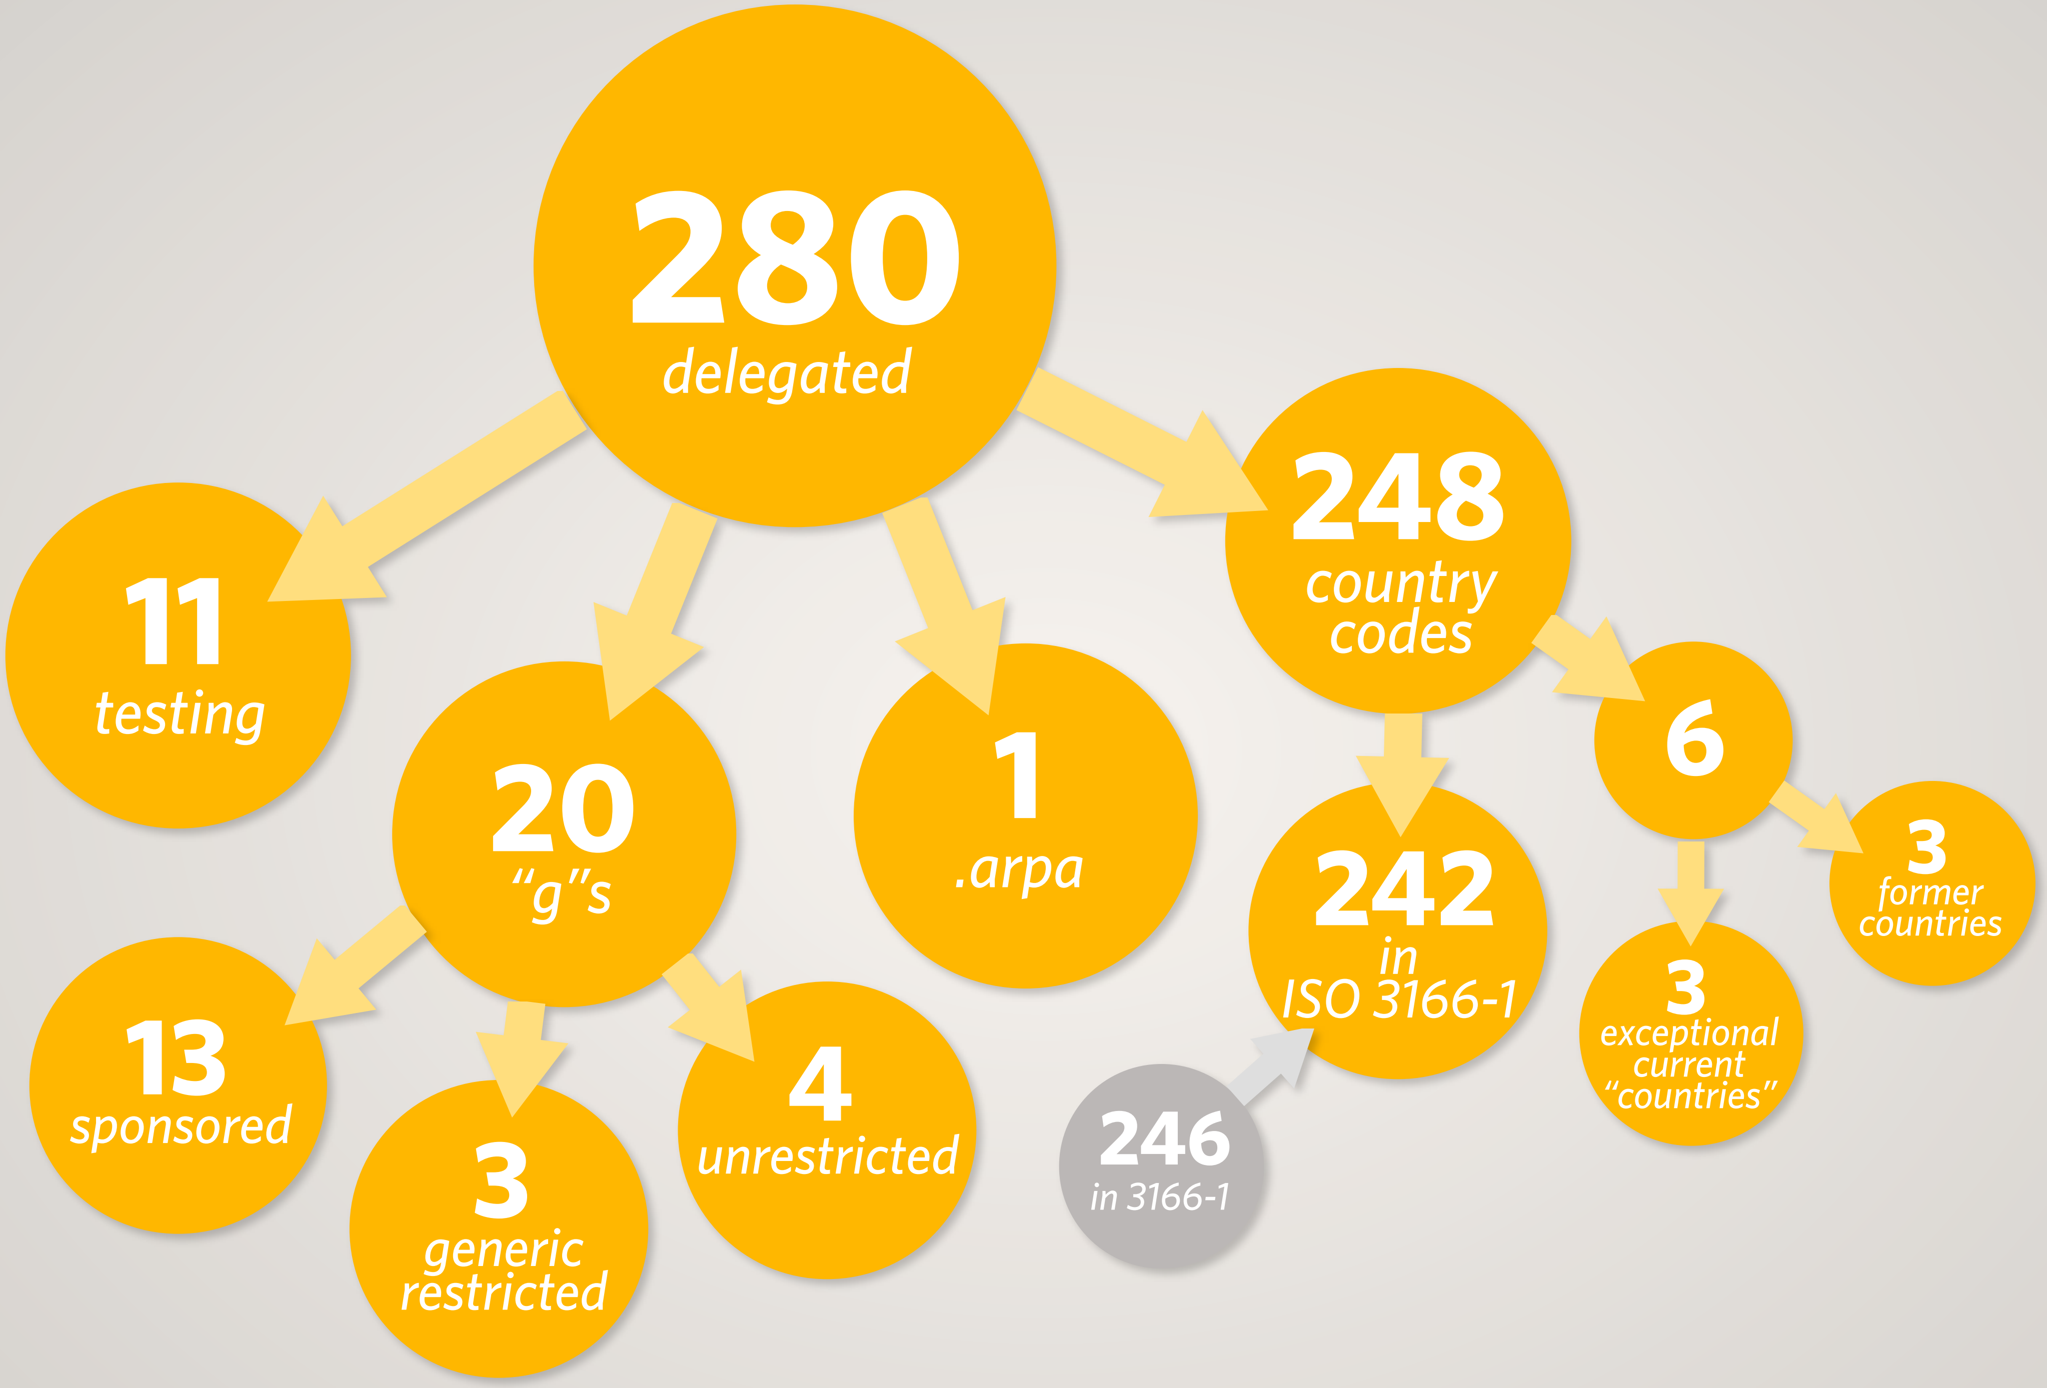

280 delegated

11 testing

20 “g”s

13 sponsored

3 generic restricted

4 unrestricted

1 .arpa

248 country codes

242 in ISO 3166-1

246 in 3166-1

6

3 exceptional current “countries”

3 former countries

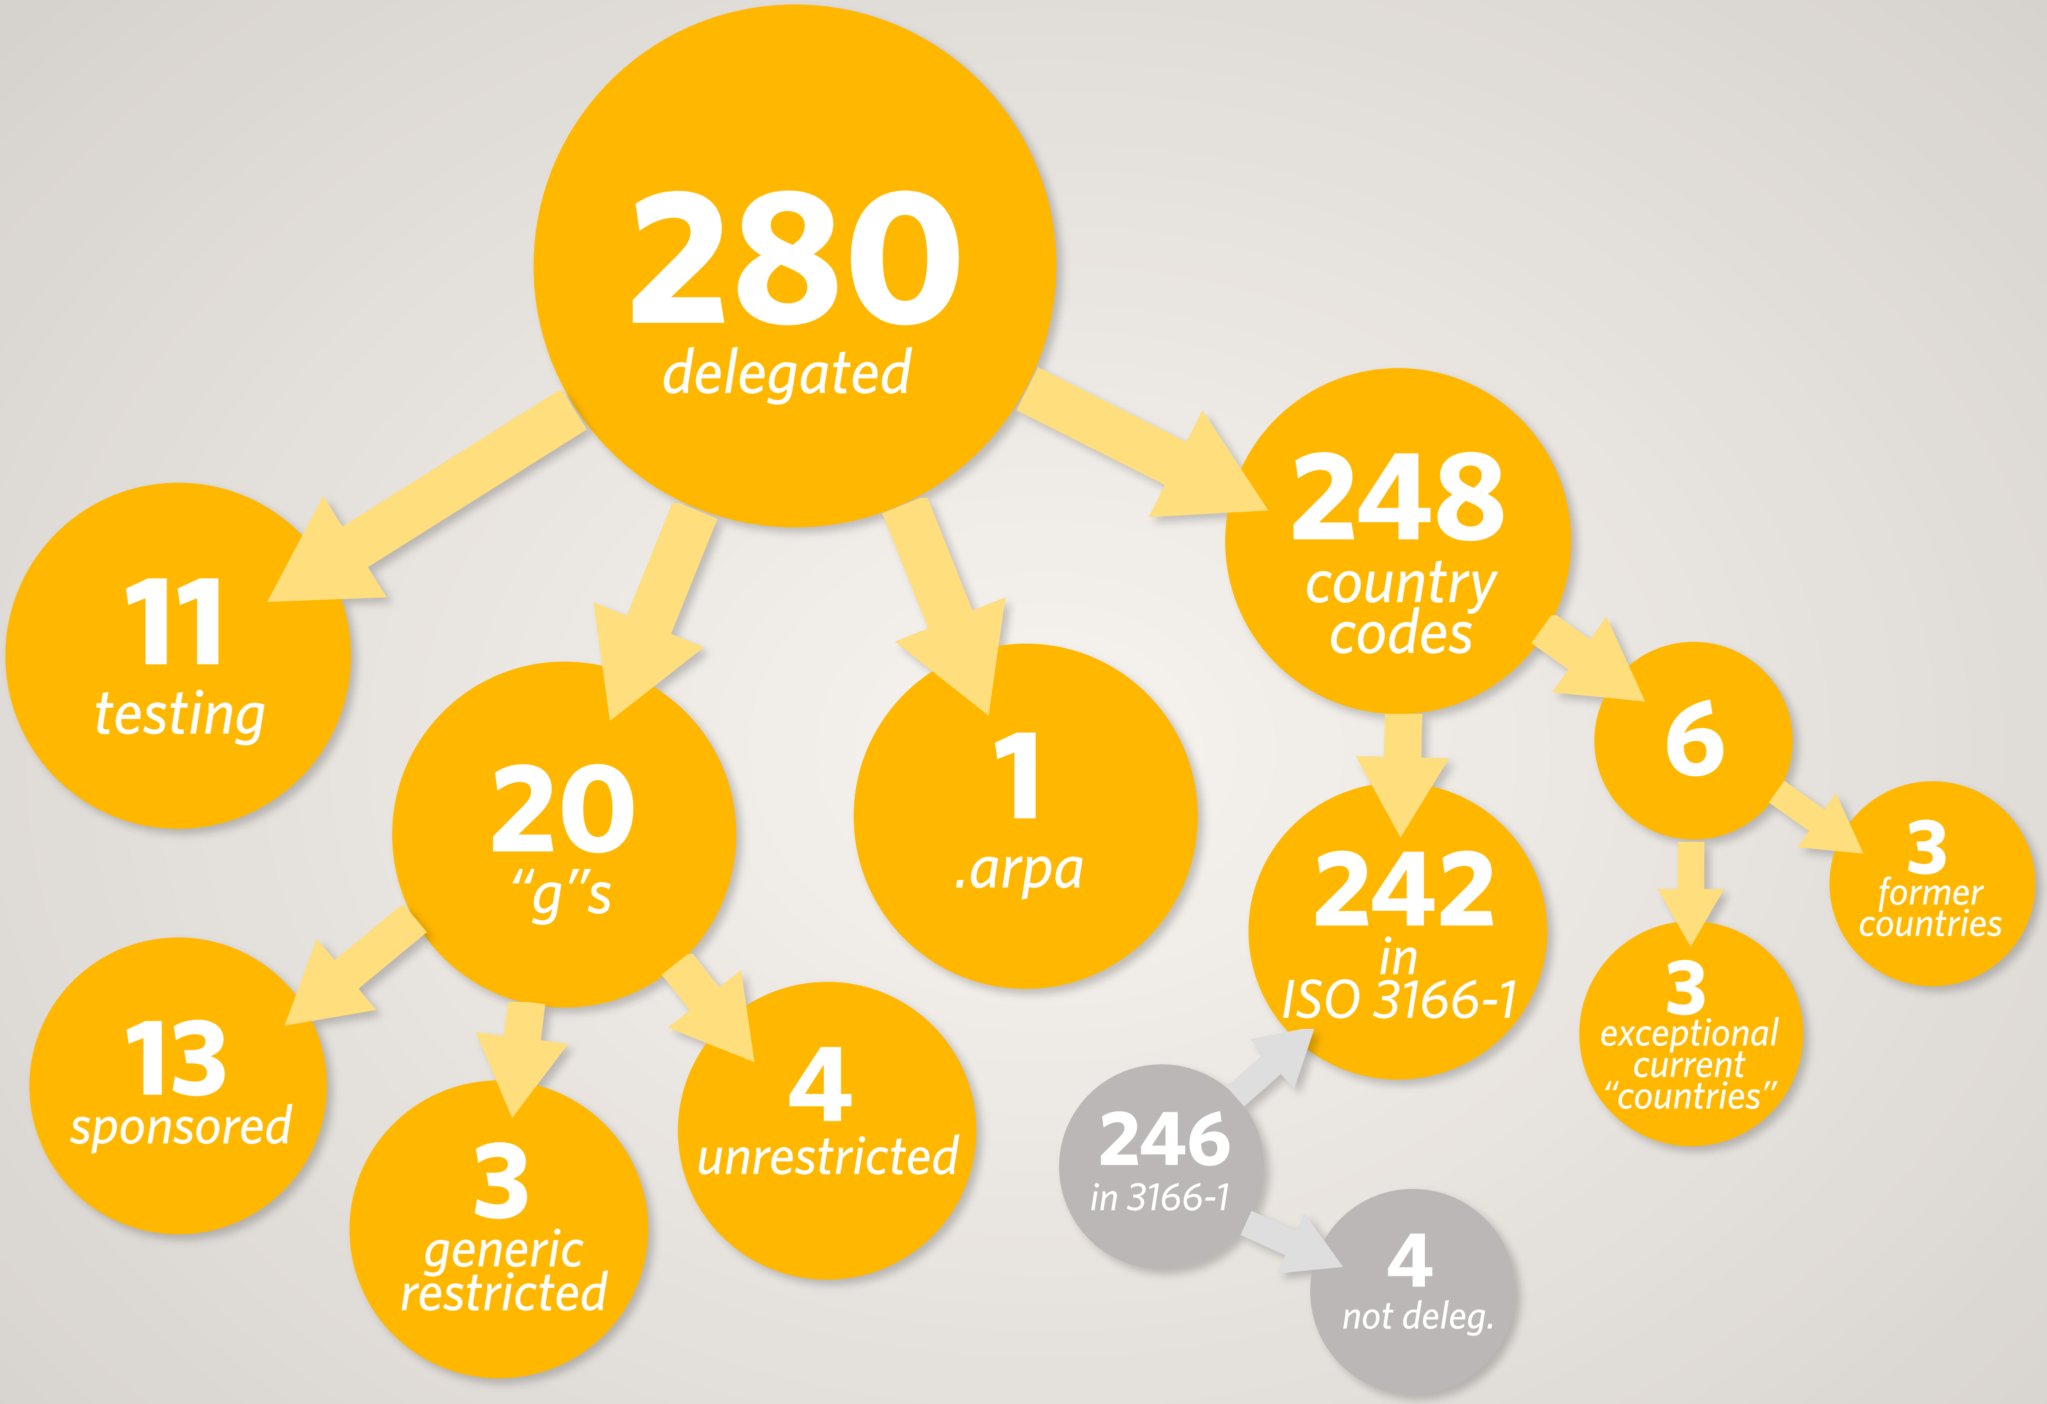

280 delegated

11 testing

20 "g"s

13 sponsored

3 generic restricted

4 unrestricted

1 .arpa

248 country codes

242 in ISO 3166-1

246 in 3166-1

4 not deleg.

6

3 exceptional current "countries"

3 former countries

Technical Conformance

Technical Conformance

- Bring our minimum technical criteria for root zone changes up to date
- Phasing in:
  - Prohibition on open recursive name servers
  - More appropriate name server diversity requirement
  - No fragmentation of root zone referrals

❶ Open recursive name servers

- Not good network citizens
  - Open to cache poisoning attacks (Kaminsky, et.al)
  - Open to amplification attacks
- Not required for authoritative service

❶ Open recursive name servers

- Not good network citizens
  - Open to cache poisoning attacks (Kaminsky, et.al)
  - Open to amplification attacks
- Not required for authoritative service

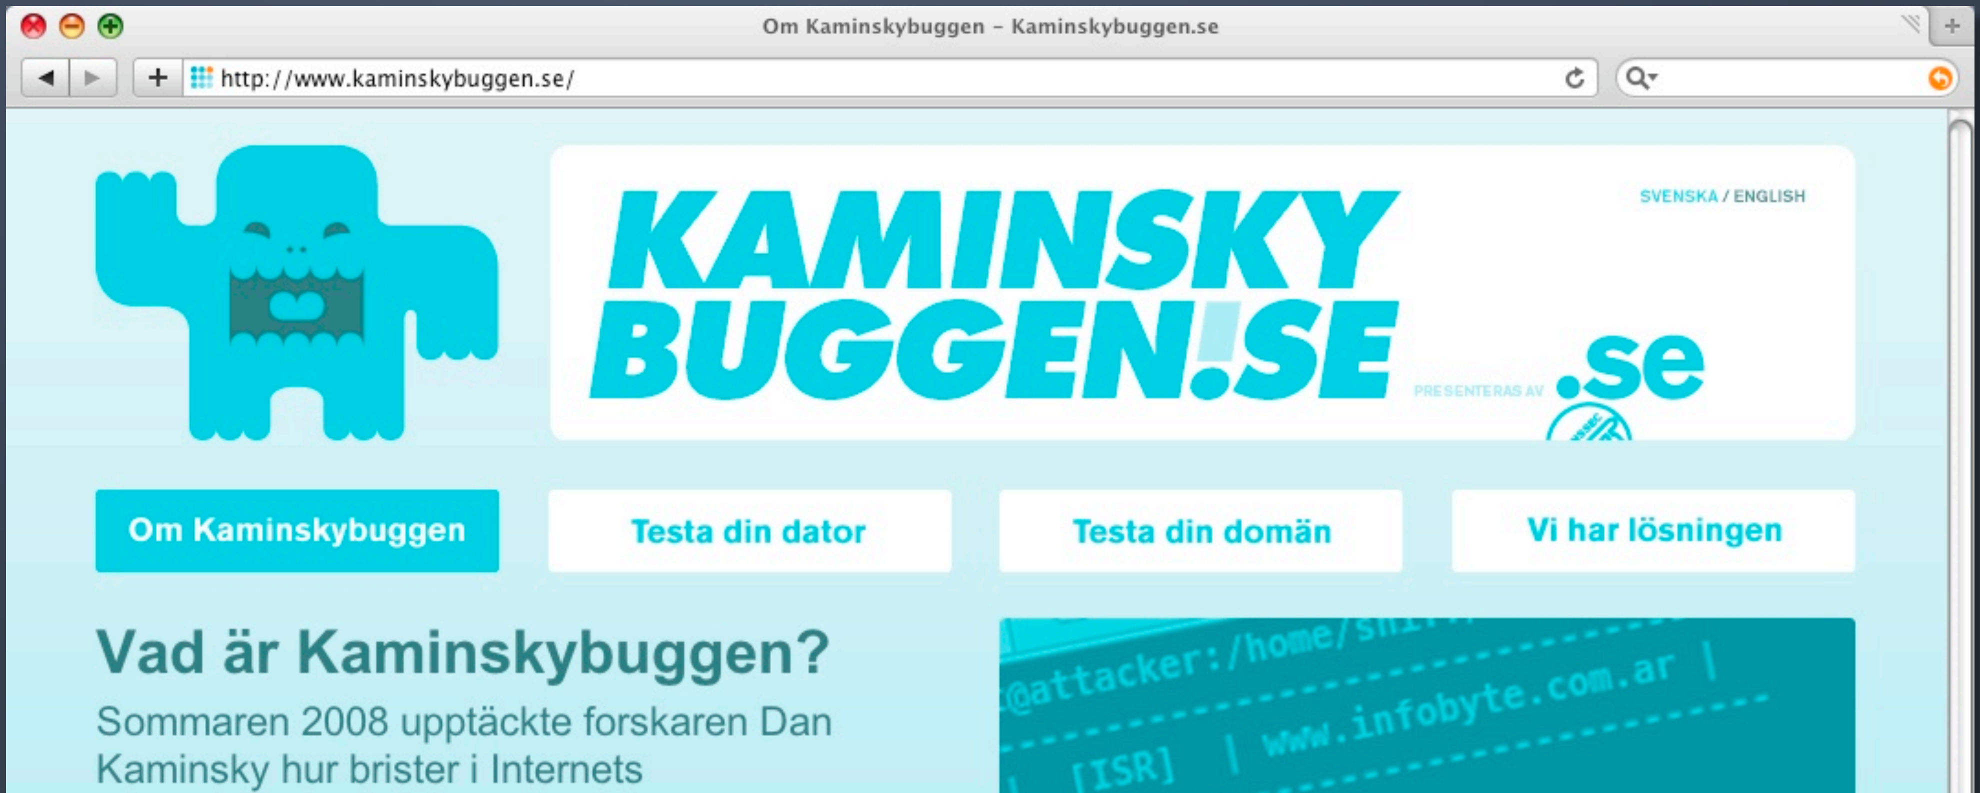

Vad är Kaminskybuggen?

Sommaren 2008 upptäckte forskaren Dan Kaminsky hur brister i Internets

❷ Network diversity for name servers

- Current informal rule is a minimum of two “not in the same /24 subnet”
    - Not very relevant to networks today
- Each IP address on the Internet’s network location is derived through announcements in the “global routing table” using BGP
- Each network is roughly organised into a group called an “autonomous system”
- Require name servers to be announced in at least two different autonomous systems

.CX

ns.cx-nic.org.nz[203.119.12.245]
ns.anycast.nic.cx[204.61.216.16]
cx1.dyntld.net[208.78.70.77]
cx2.dyntld.net[204.13.250.77]
cx3.dyntld.net[208.78.71.77]
cx4.dyntld.net[204.13.251.77]

.CX

ns.cx-nic.org.nz[203.119.12.245]	Hostway Corporation Pty Ltd
ns.anycast.nic.cx[204.61.216.16]	WoodyNet
cx1.dyntld.net[208.78.70.77]	Dynamic Network Services, Inc.
cx2.dyntld.net[204.13.250.77]	Dynamic Network Services, Inc.
cx3.dyntld.net[208.78.71.77]	Dynamic Network Services, Inc.
cx4.dyntld.net[204.13.251.77]	Dynamic Network Services, Inc.

.CX

ns.cx-nic.org.nz[203.119.12.245]	Hostway Corporation Pty Ltd
ns.anycast.nic.cx[204.61.216.16]	WoodyNet
cx1.dyntld.net[208.78.70.77]	Dynamic Network Services, Inc.
cx2.dyntld.net[204.13.250.77]	Dynamic Network Services, Inc.
cx3.dyntld.net[208.78.71.77]	Dynamic Network Services, Inc.
cx4.dyntld.net[204.13.251.77]	Dynamic Network Services, Inc.

3 distinct networks

.CX

ns.cx-nic.org.nz[203.119.12.245]	Hostway Corporation Pty Ltd
ns.anycast.nic.cx[204.61.216.16]	WoodyNet
cx1.dyntld.net[208.78.70.77]	Dynamic Network Services, Inc.
cx2.dyntld.net[204.13.250.77]	Dynamic Network Services, Inc.
cx3.dyntld.net[208.78.71.77]	Dynamic Network Services, Inc.
cx4.dyntld.net[204.13.251.77]	Dynamic Network Services, Inc.

3 distinct networks ✔

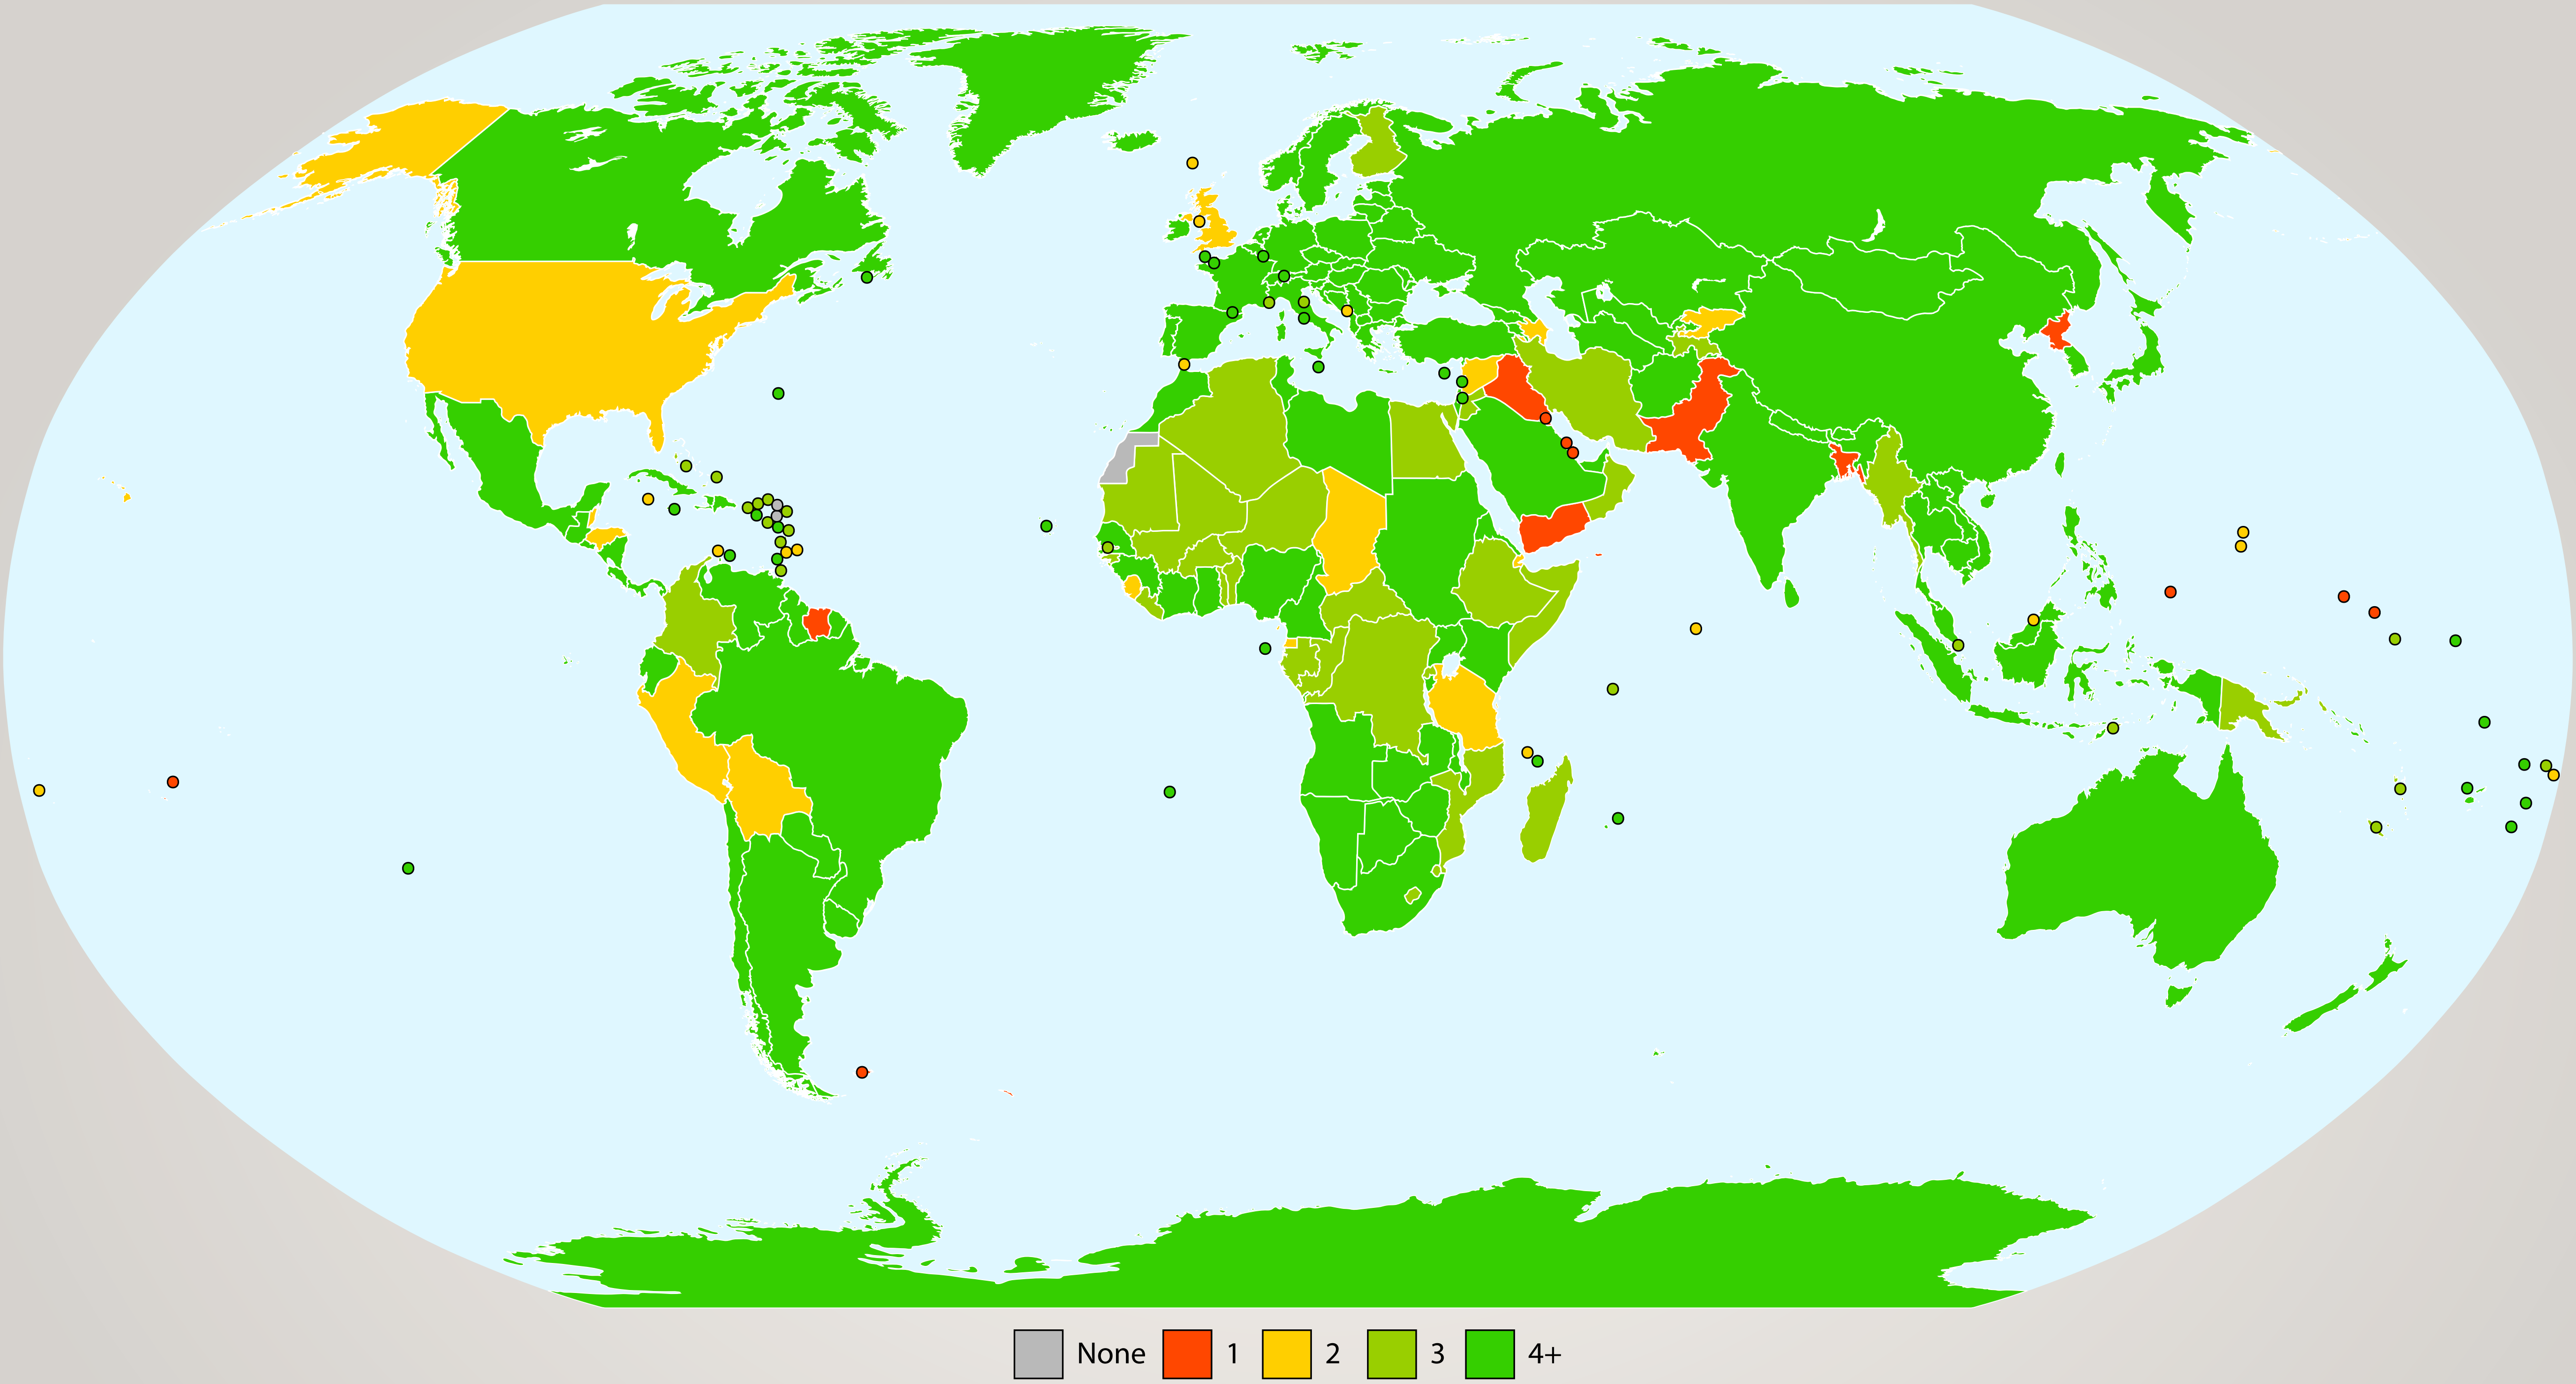

ccTLDs with AS diversity

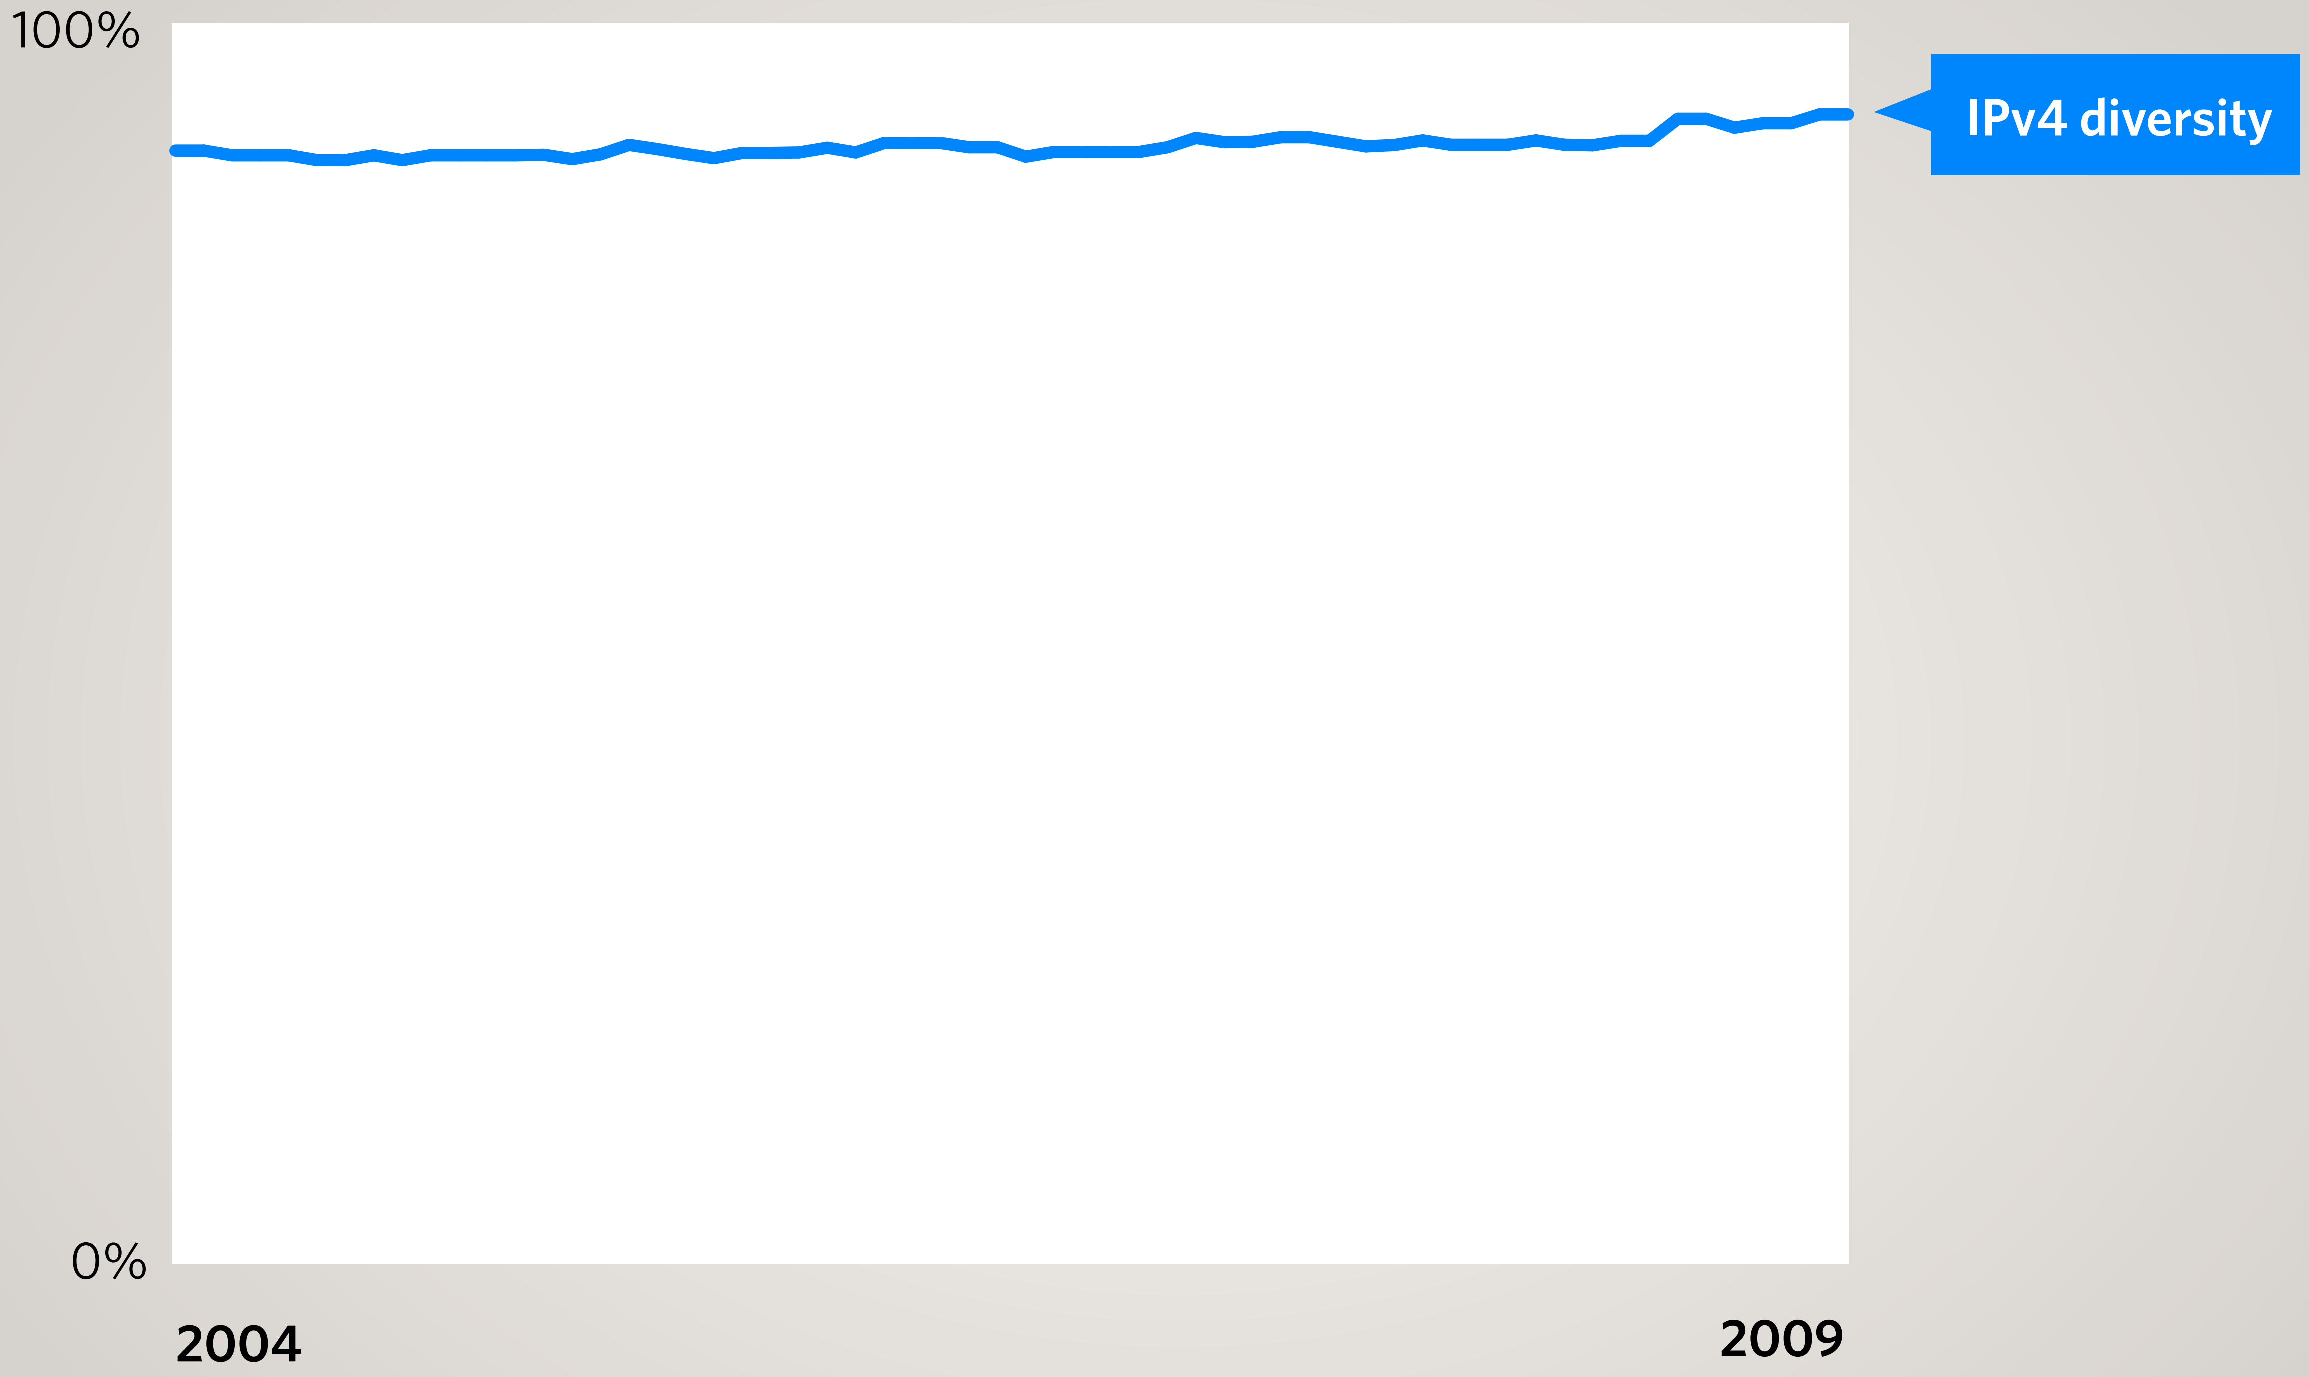
100%
0%
2004
2009
IPv4 diversity

Pushing the envelope...

"IANA currently has a minimum set of technical requirements for IPv4 name service. These include two nameservers separated by geography and by network topology, that each serve a consistent set of data, and are reachable from multiple locations across the globe. The registry will meet this same criterion for IPv6, requiring IPv6 transport to their network."

—Evaluation Criterion #40
Draft gTLD Applicant Guide Book

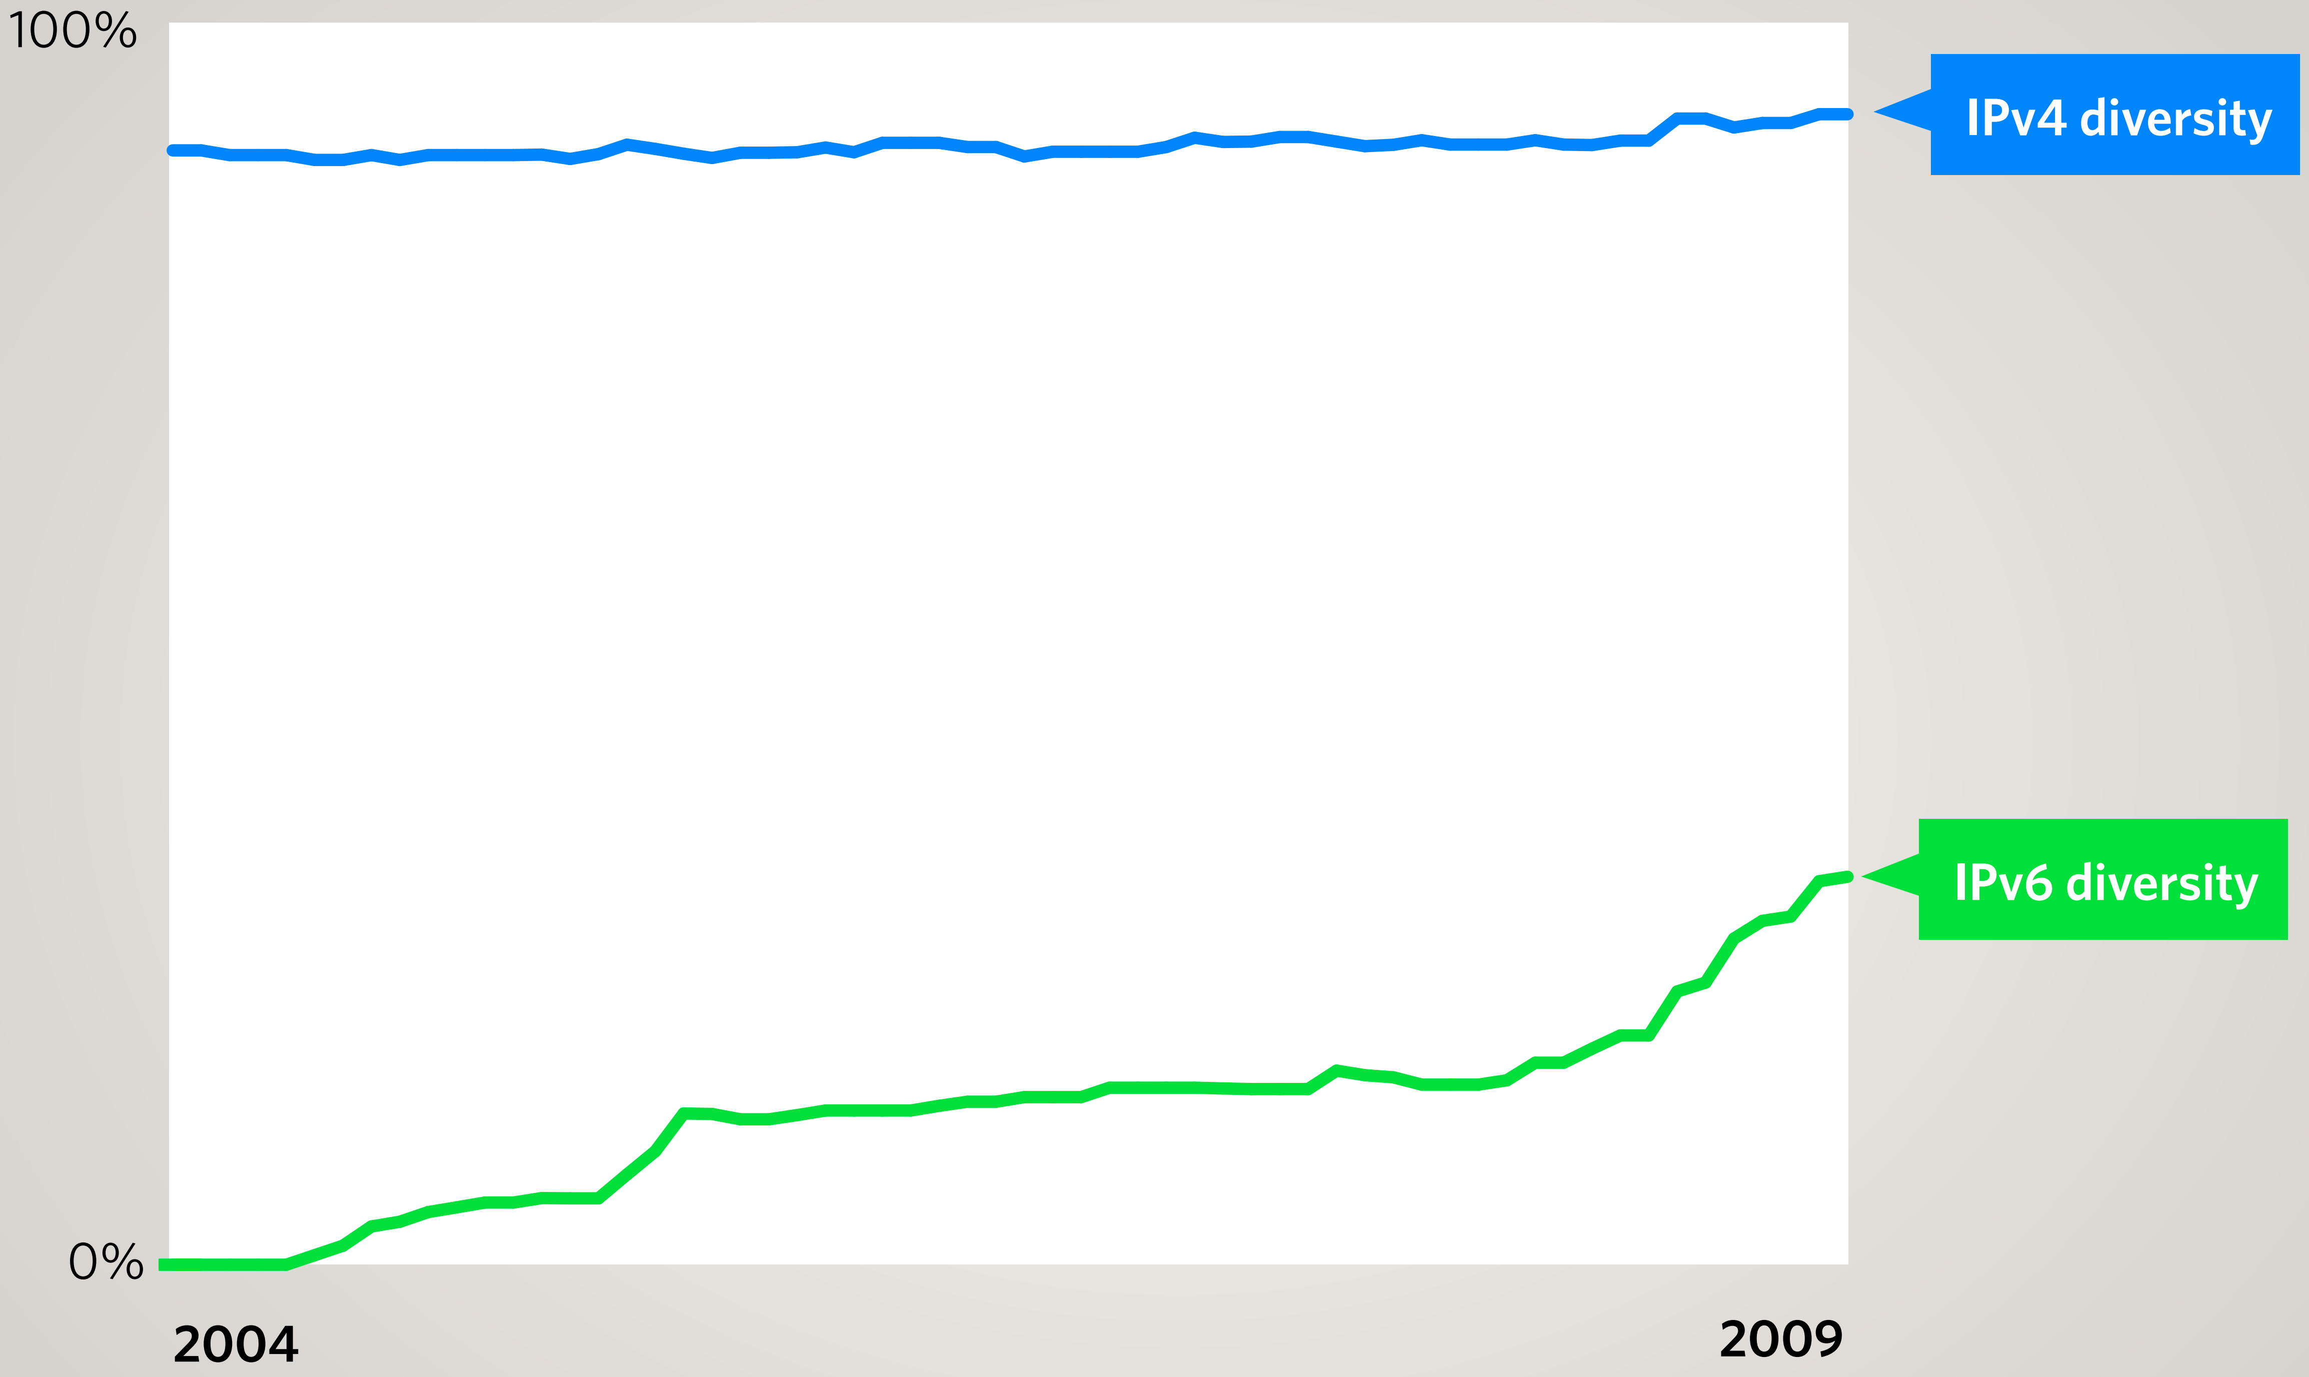
100%
0%
2004
2009
IPv4 diversity
IPv6 diversity

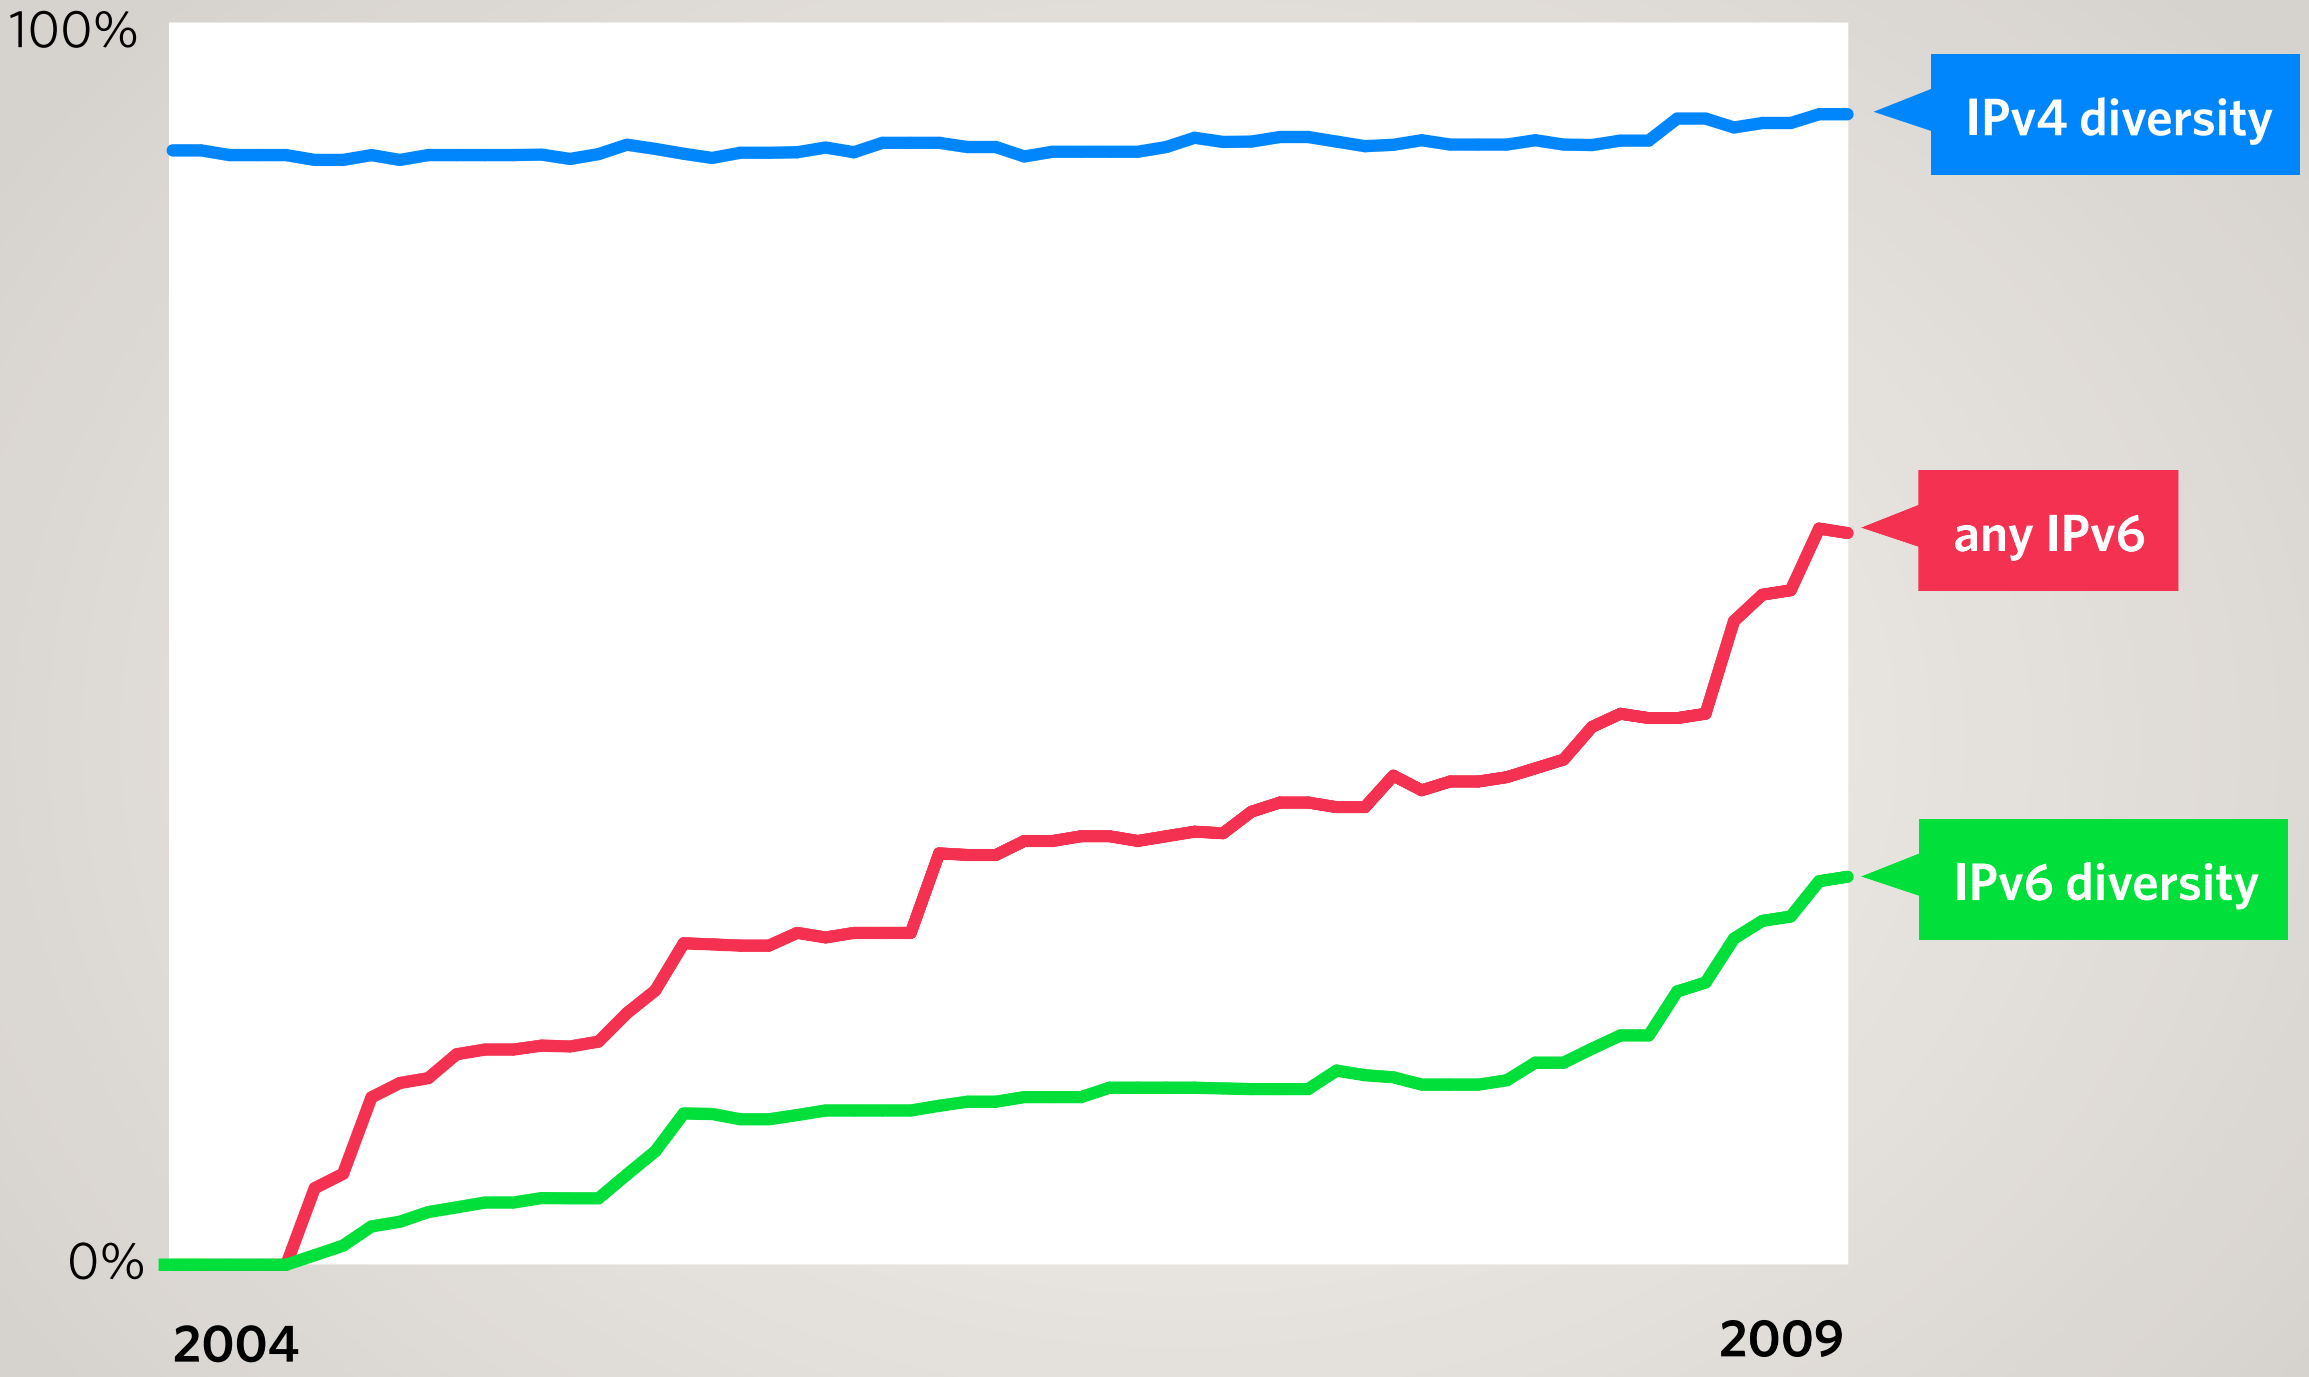
100%
0%
2004
2009
IPv4 diversity
any IPv6
IPv6 diversity

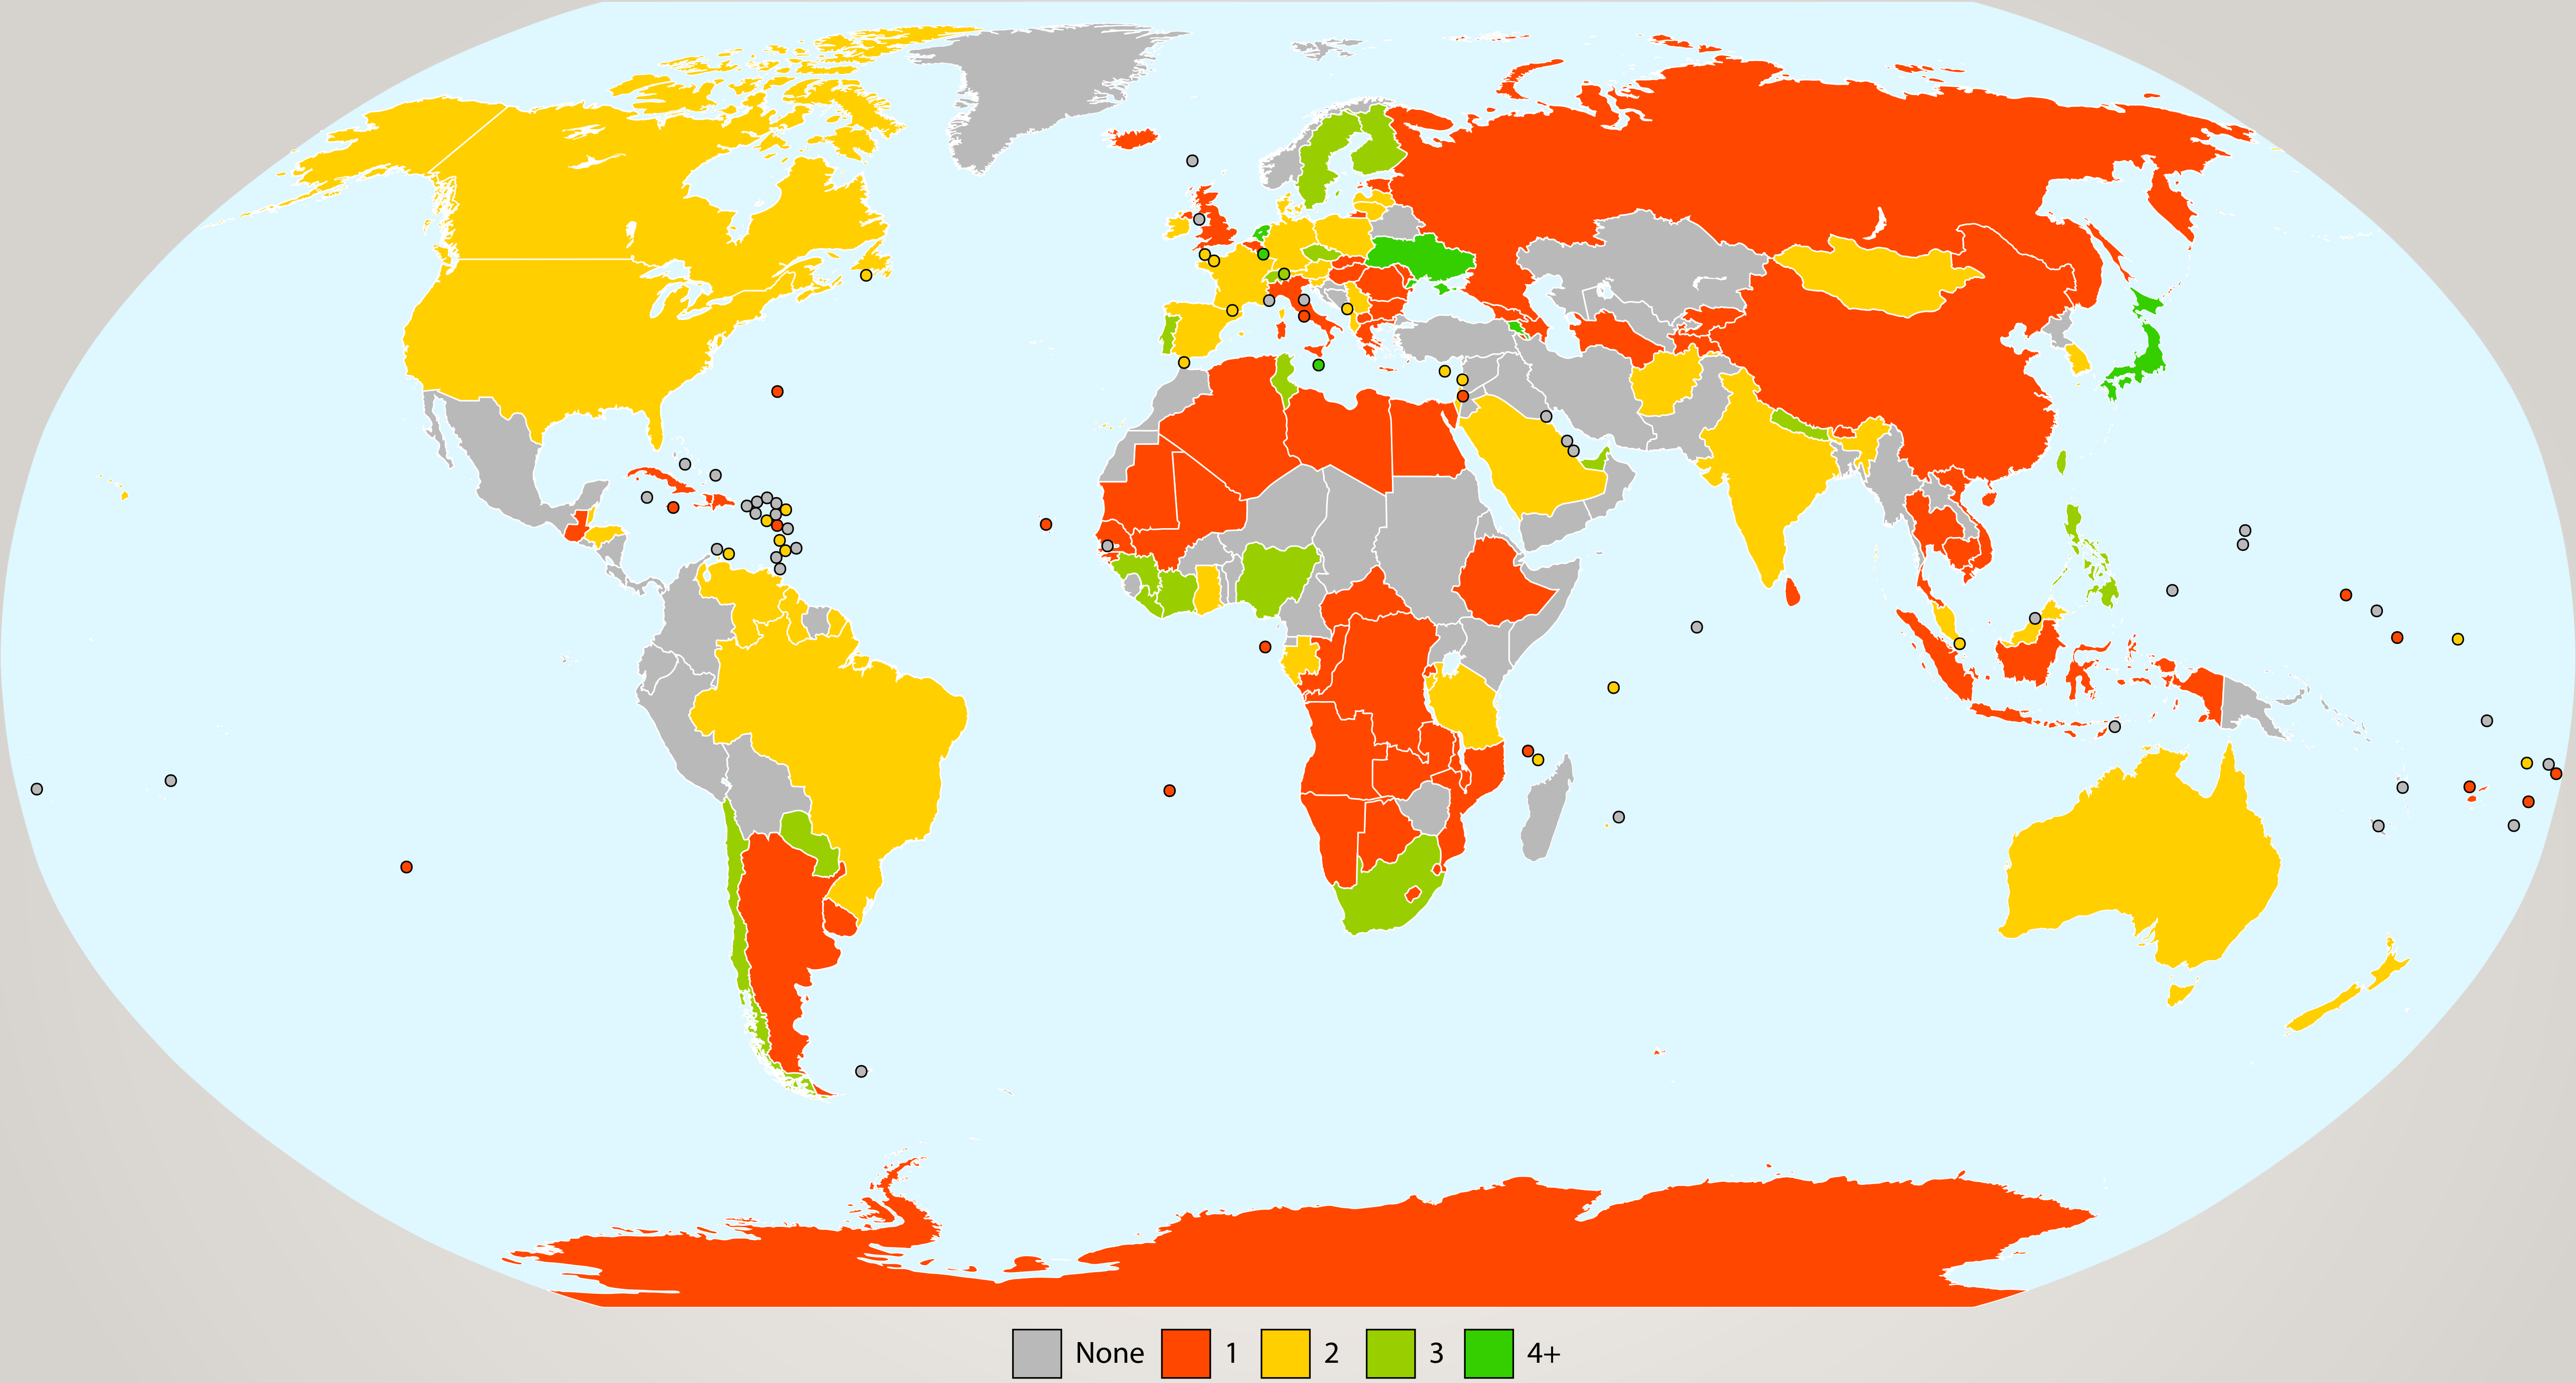

ccTLDs with AS diversity over IPv6

❸ Referrals should not fragment

- A query for a domain name to the root servers results in a referral to a TLD’s authorities

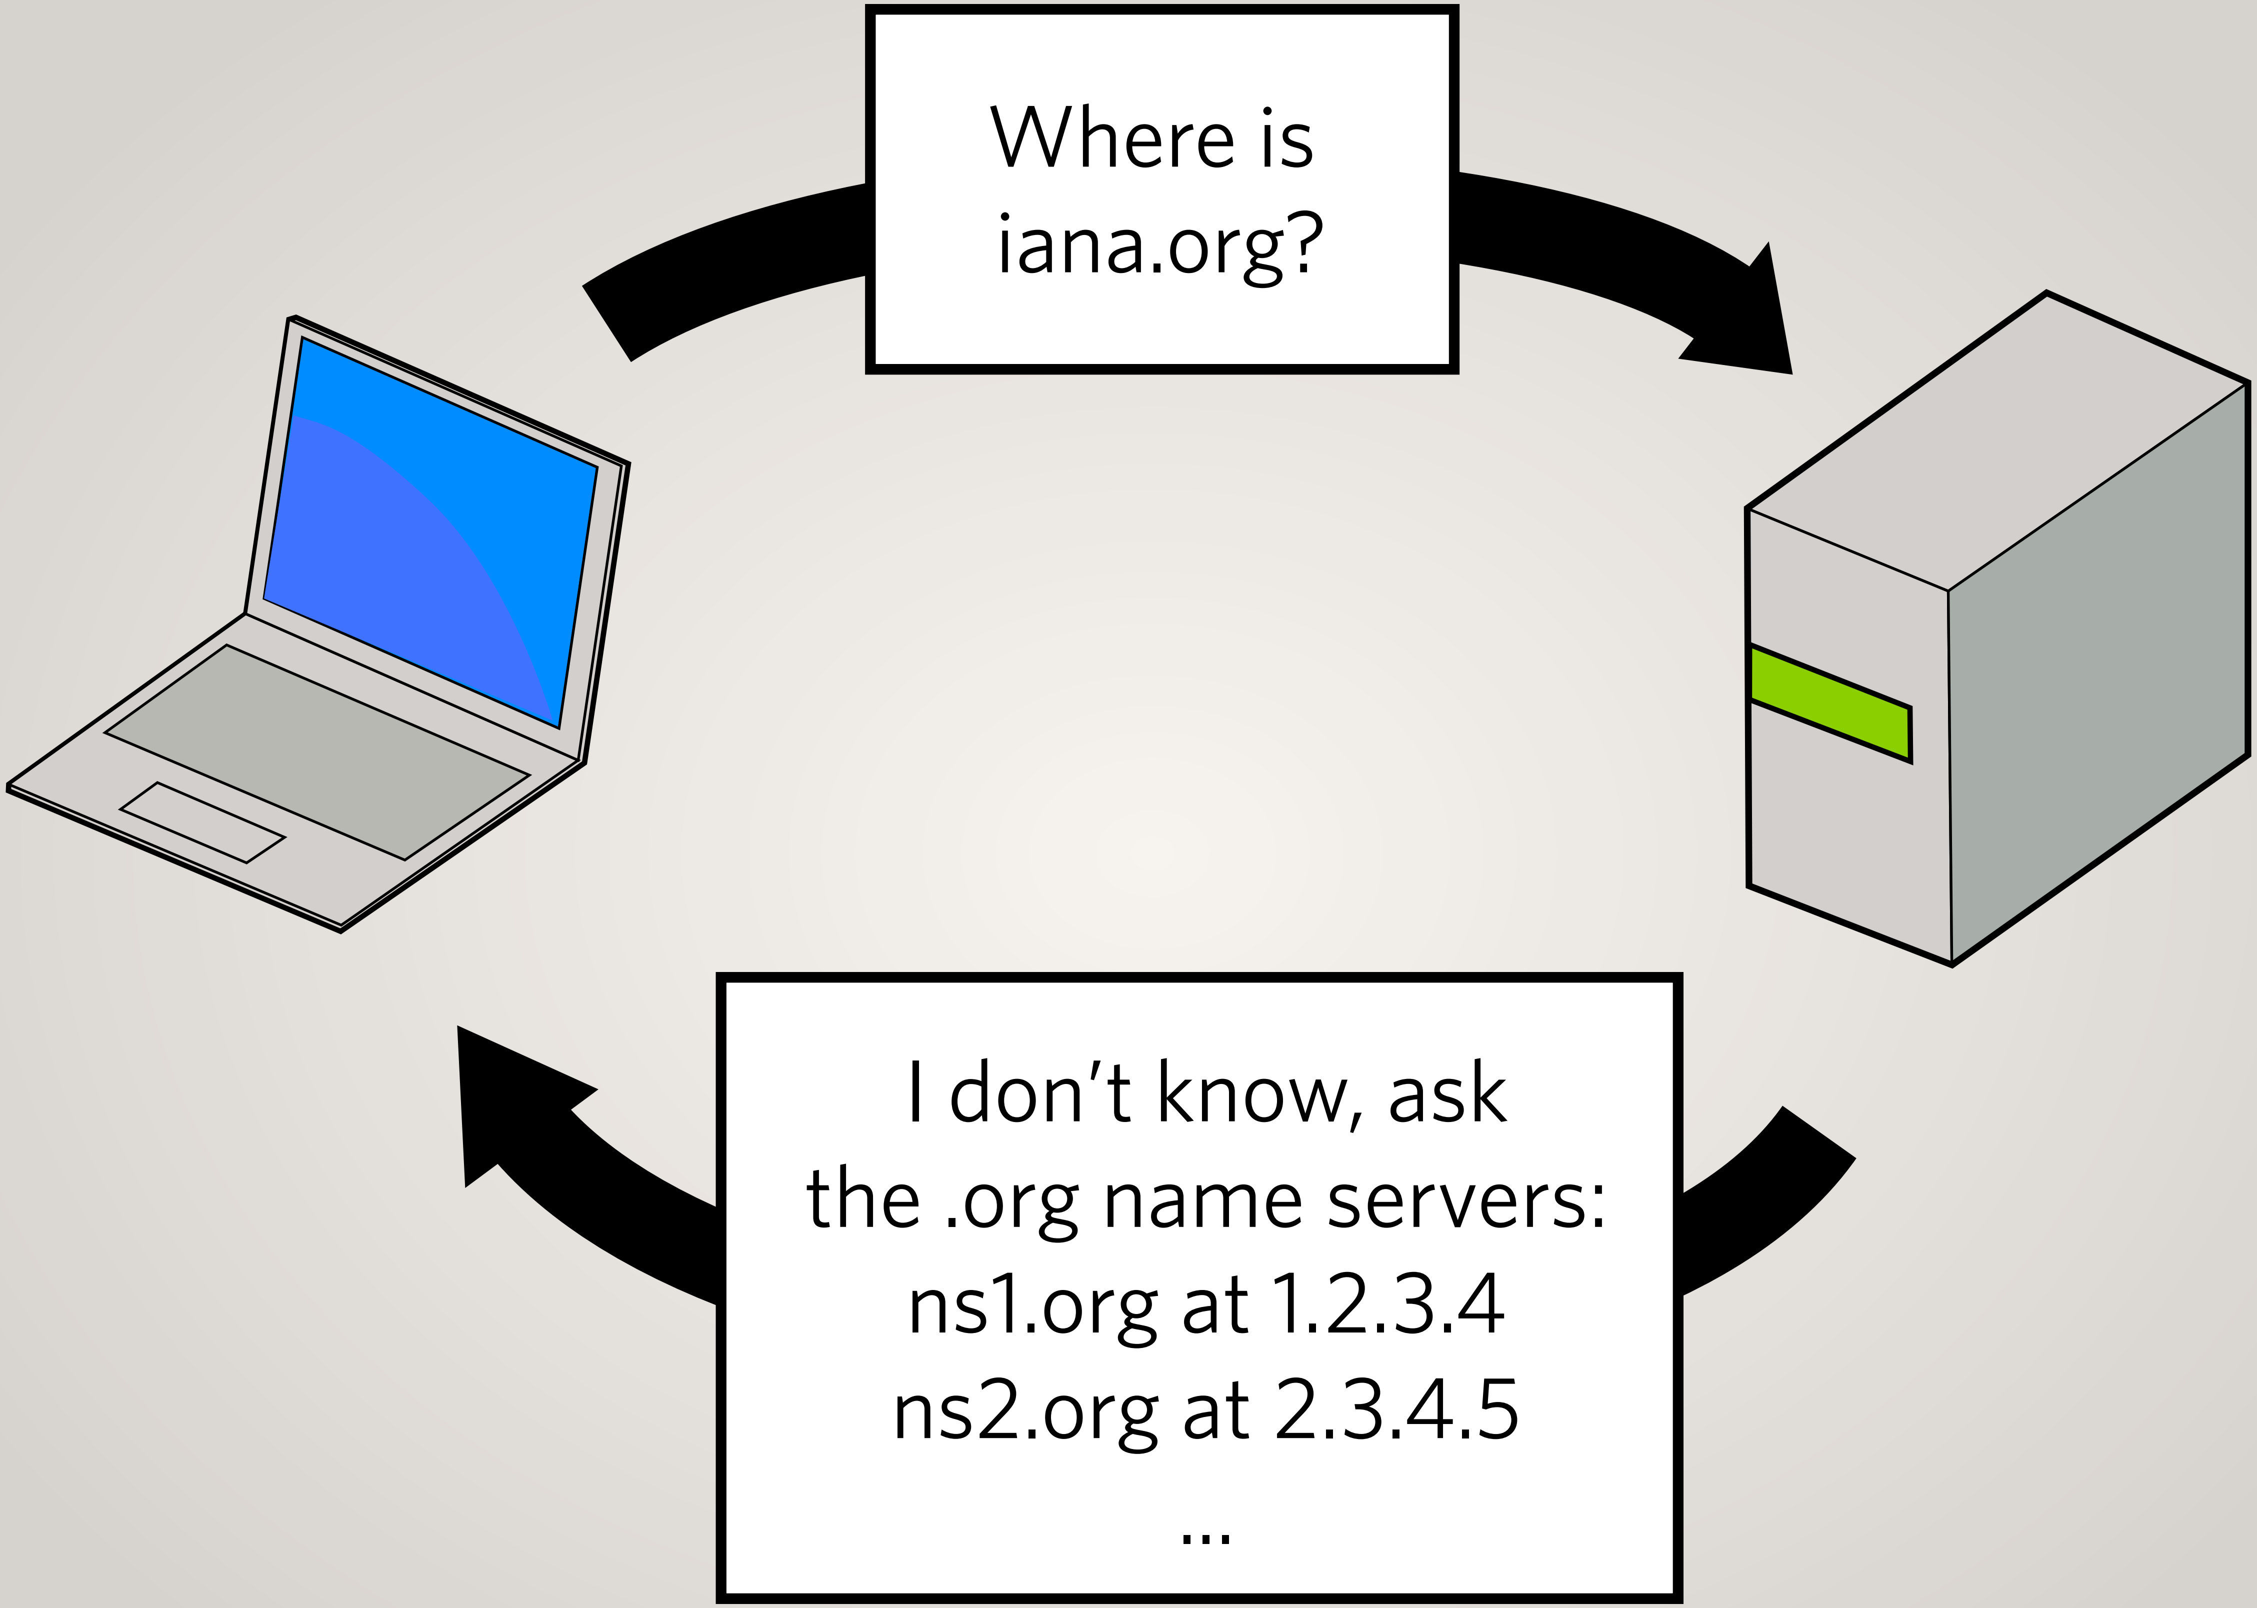

Where is iana.org?

I don’t know, ask the .org name servers:
ns1.org at 1.2.3.4
ns2.org at 2.3.4.5
…

❸ Referrals should not fragment

- A query for a domain name to the root servers should result in a referral to the TLD’s authorities
- Classical limit for response size is 512 bytes
- If the root server needs to send back more than 512 bytes of in a response, it will need to establish a much more complicated TCP connection, rather than use the simpler UDP protocol.
- This is not good for load and reliability

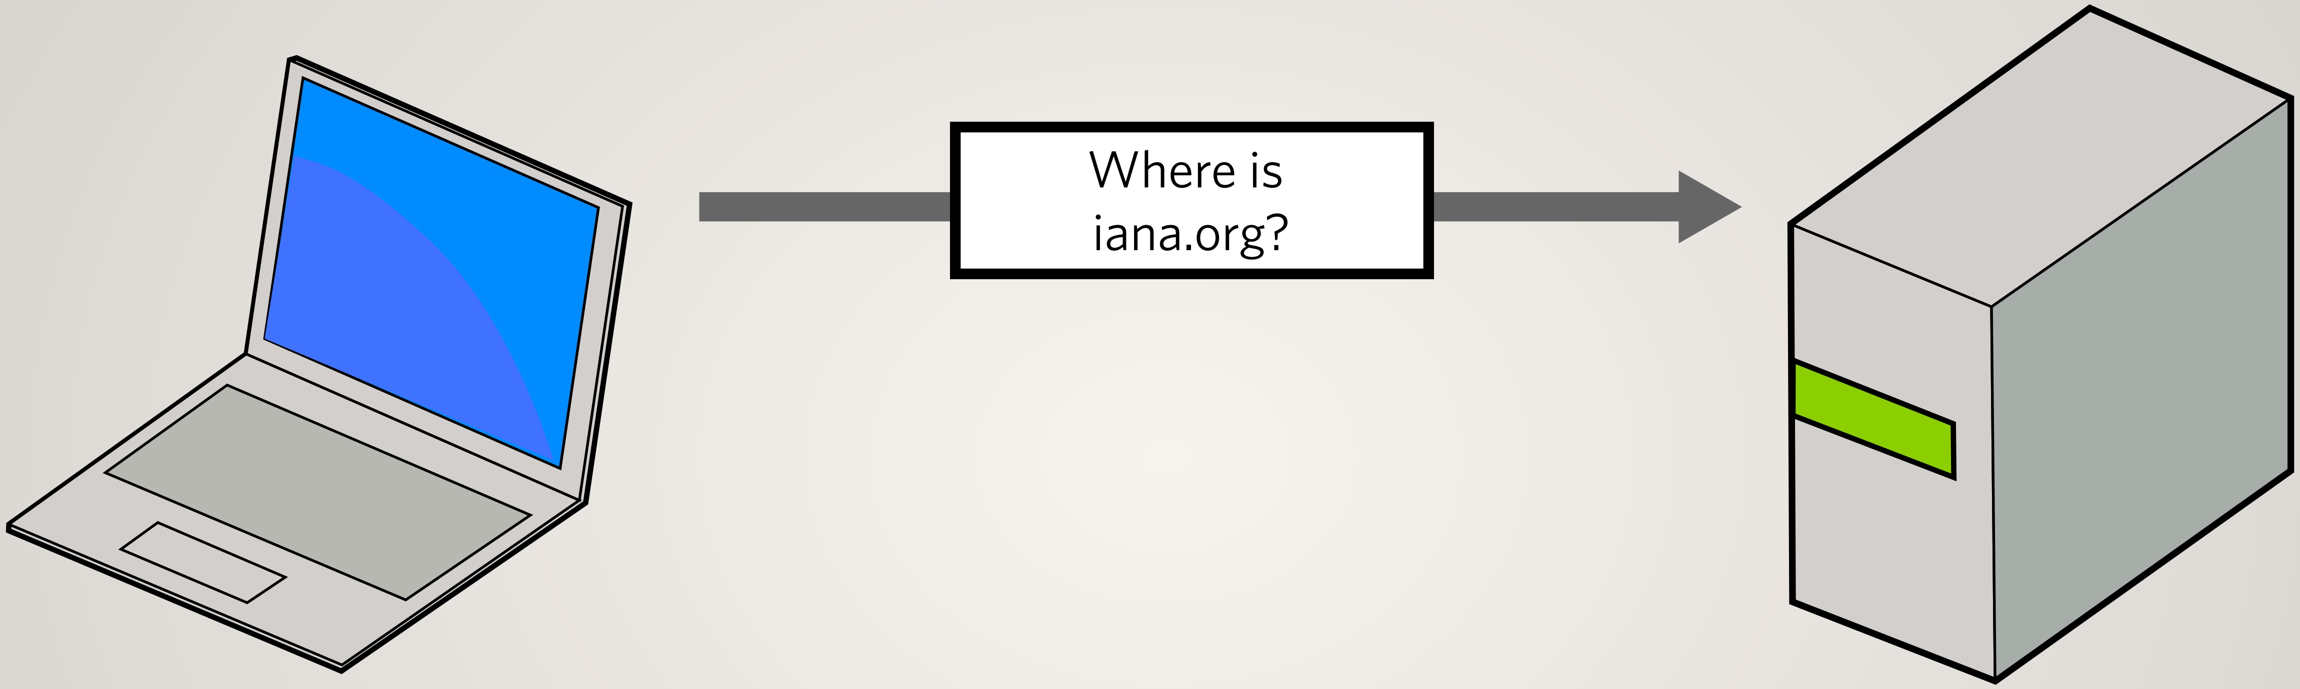

Where is iana.org?

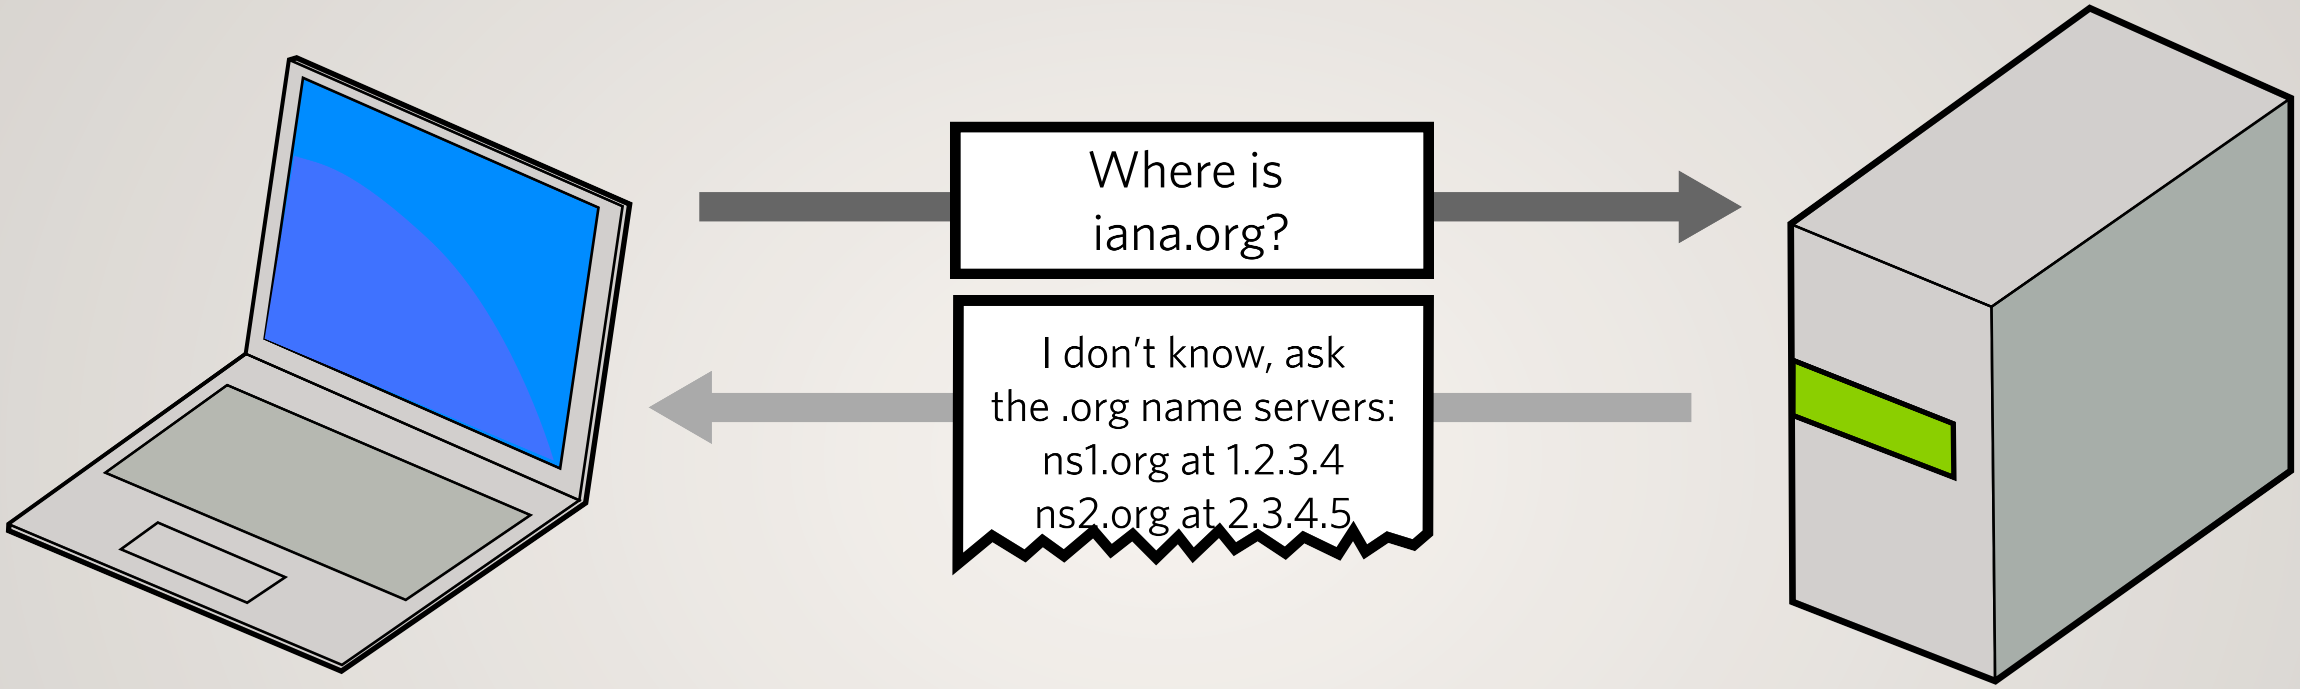

Where is iana.org?

I don't know, ask the .org name servers:
ns1.org at 1.2.3.4
ns2.org at 2.3.4.5

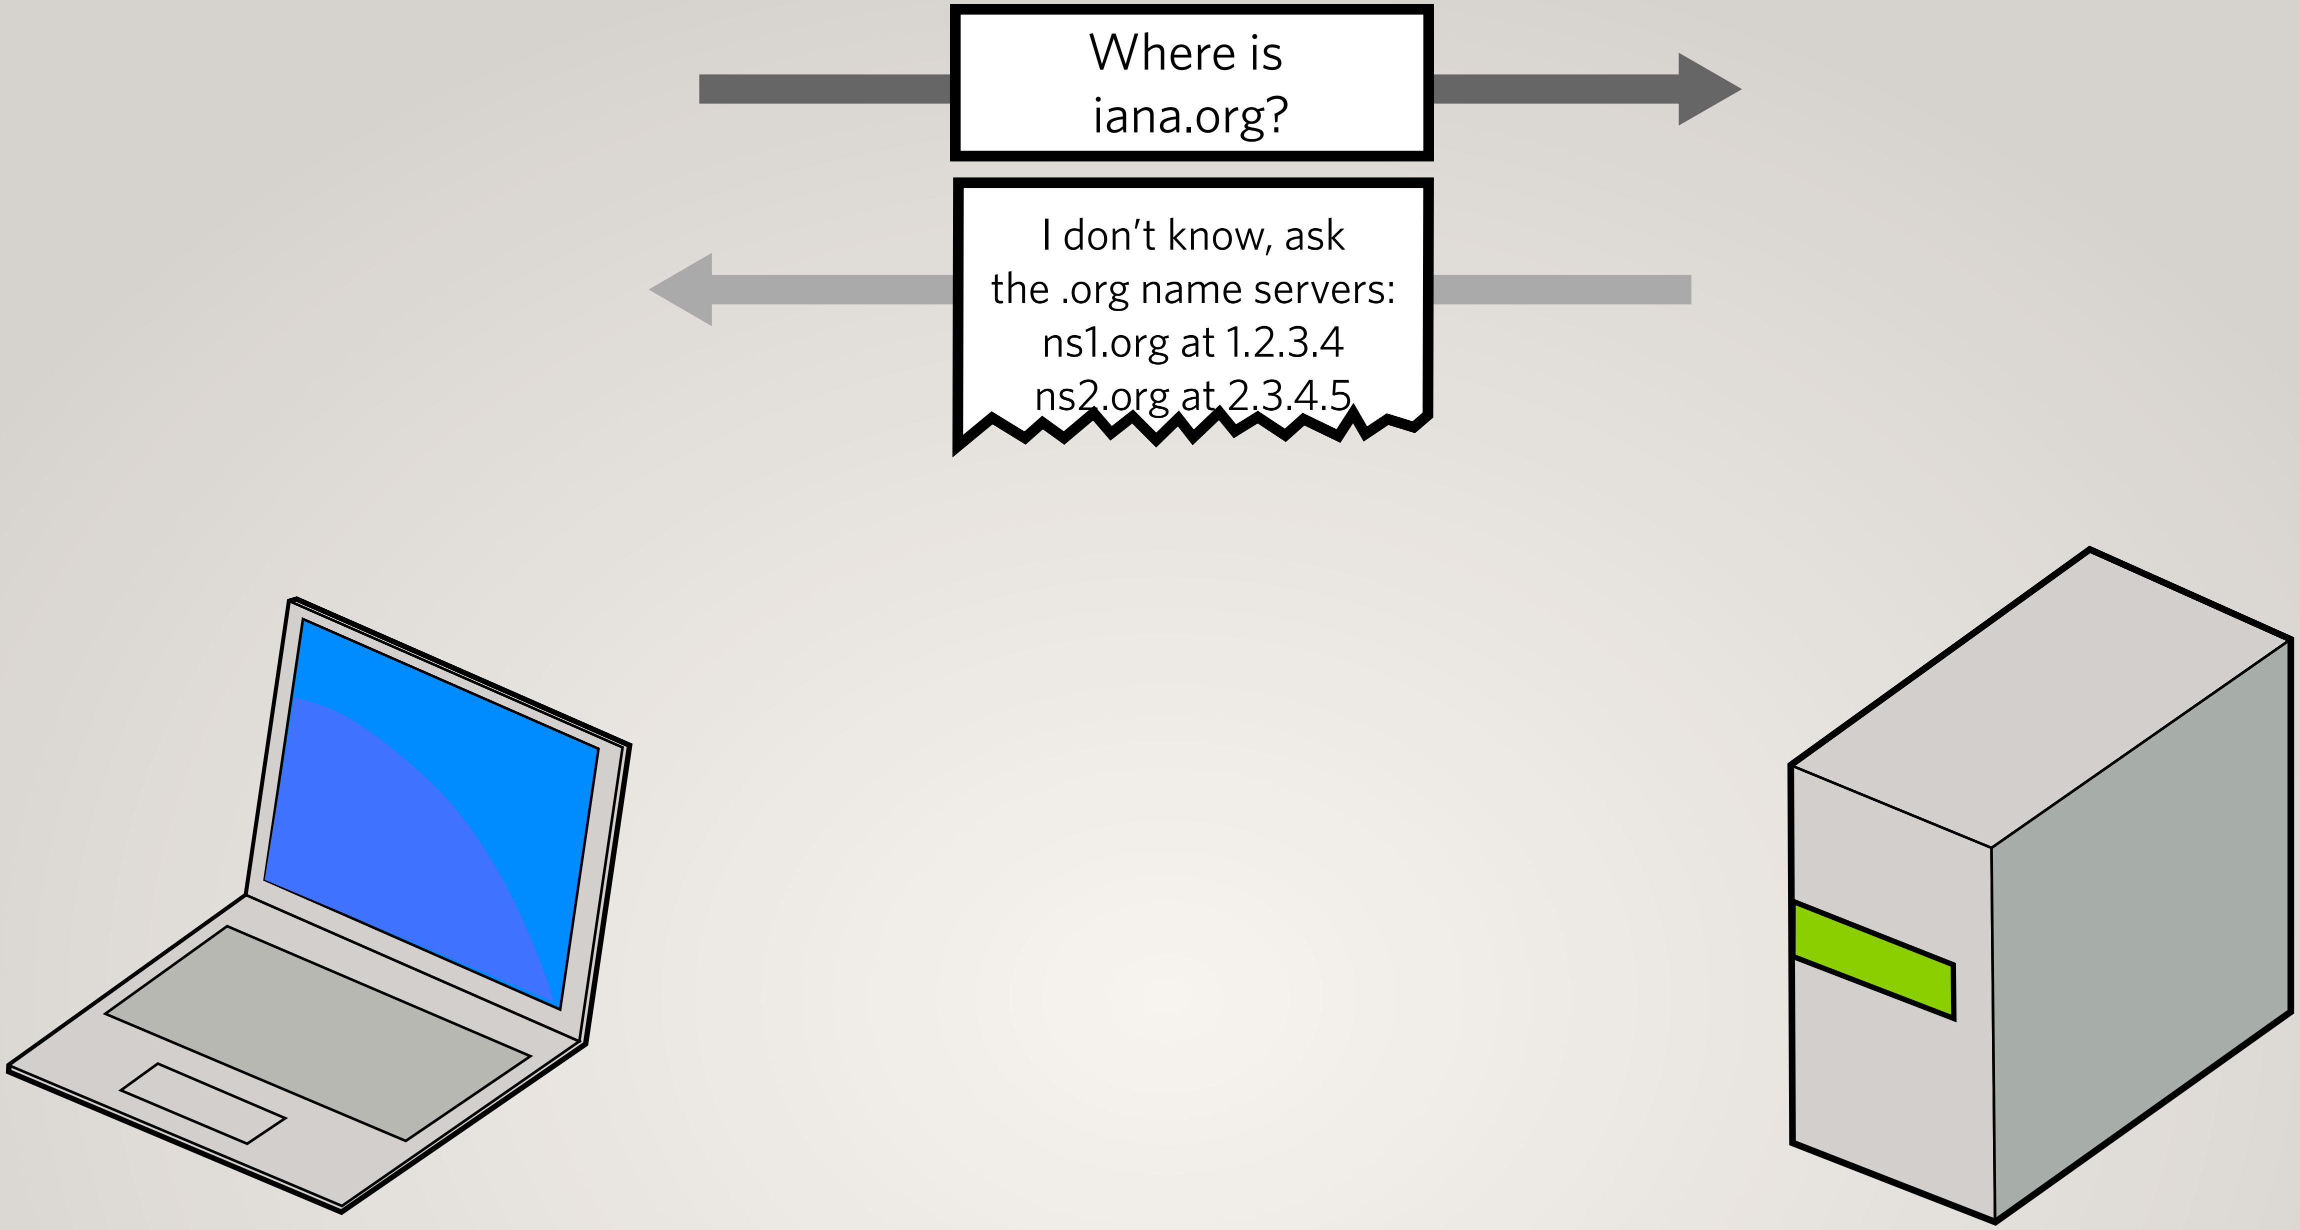

Where is
iana.org?

I don't know, ask
the .org name servers:
ns1.org at 1.2.3.4
ns2.org at 2.3.4.5

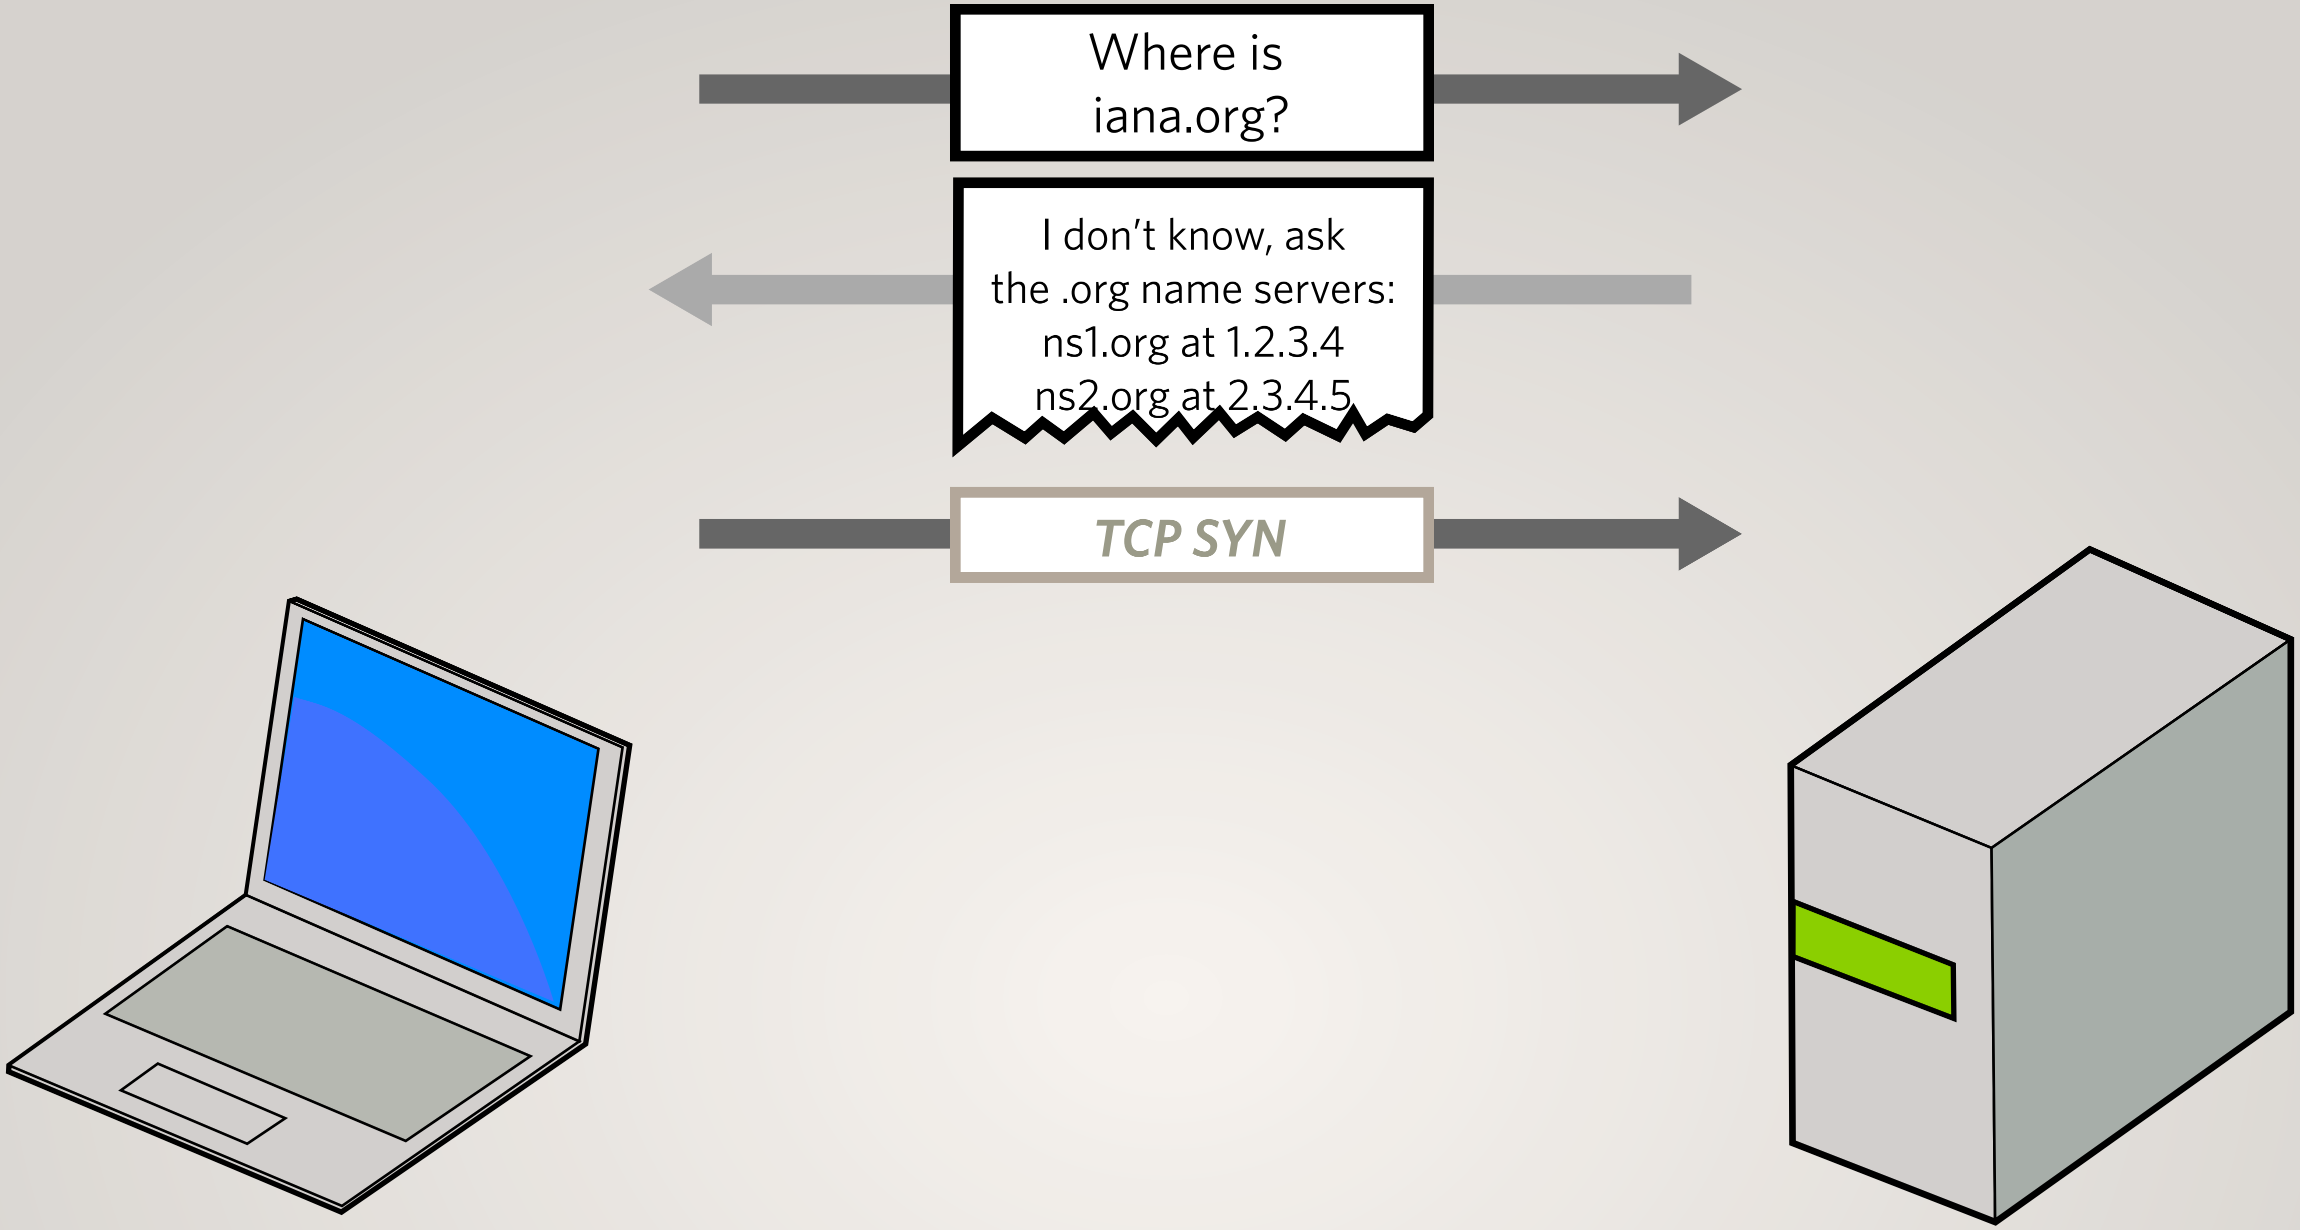

Where is iana.org?

I don’t know, ask the .org name servers:
ns1.org at 1.2.3.4
ns2.org at 2.3.4.5

TCP SYN

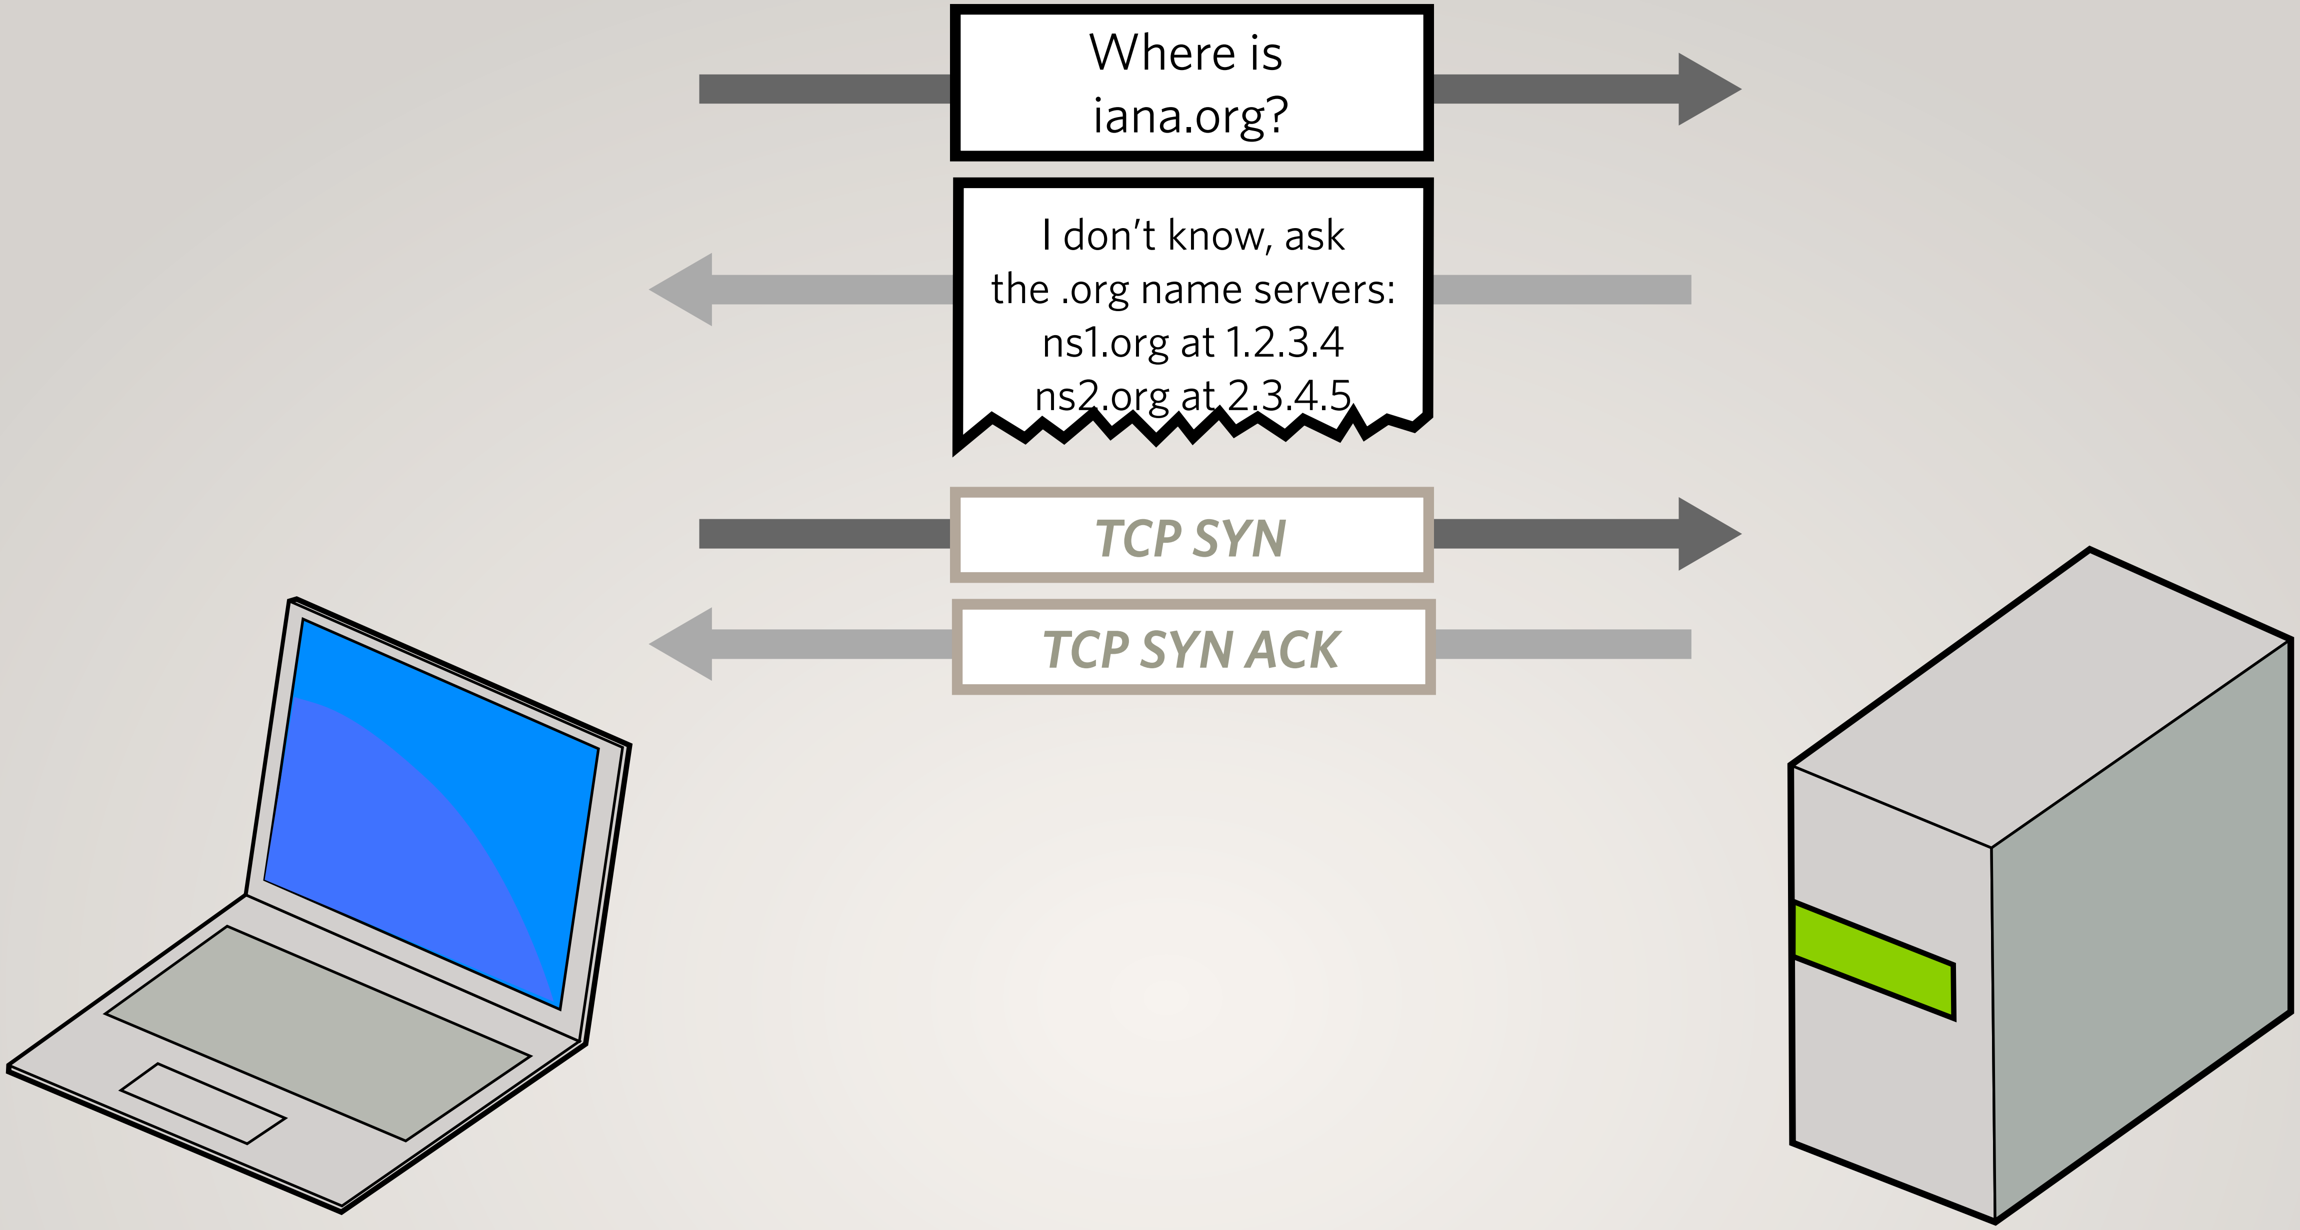

Where is iana.org?

I don't know, ask the .org name servers:
ns1.org at 1.2.3.4
ns2.org at 2.3.4.5

TCP SYN

TCP SYN ACK

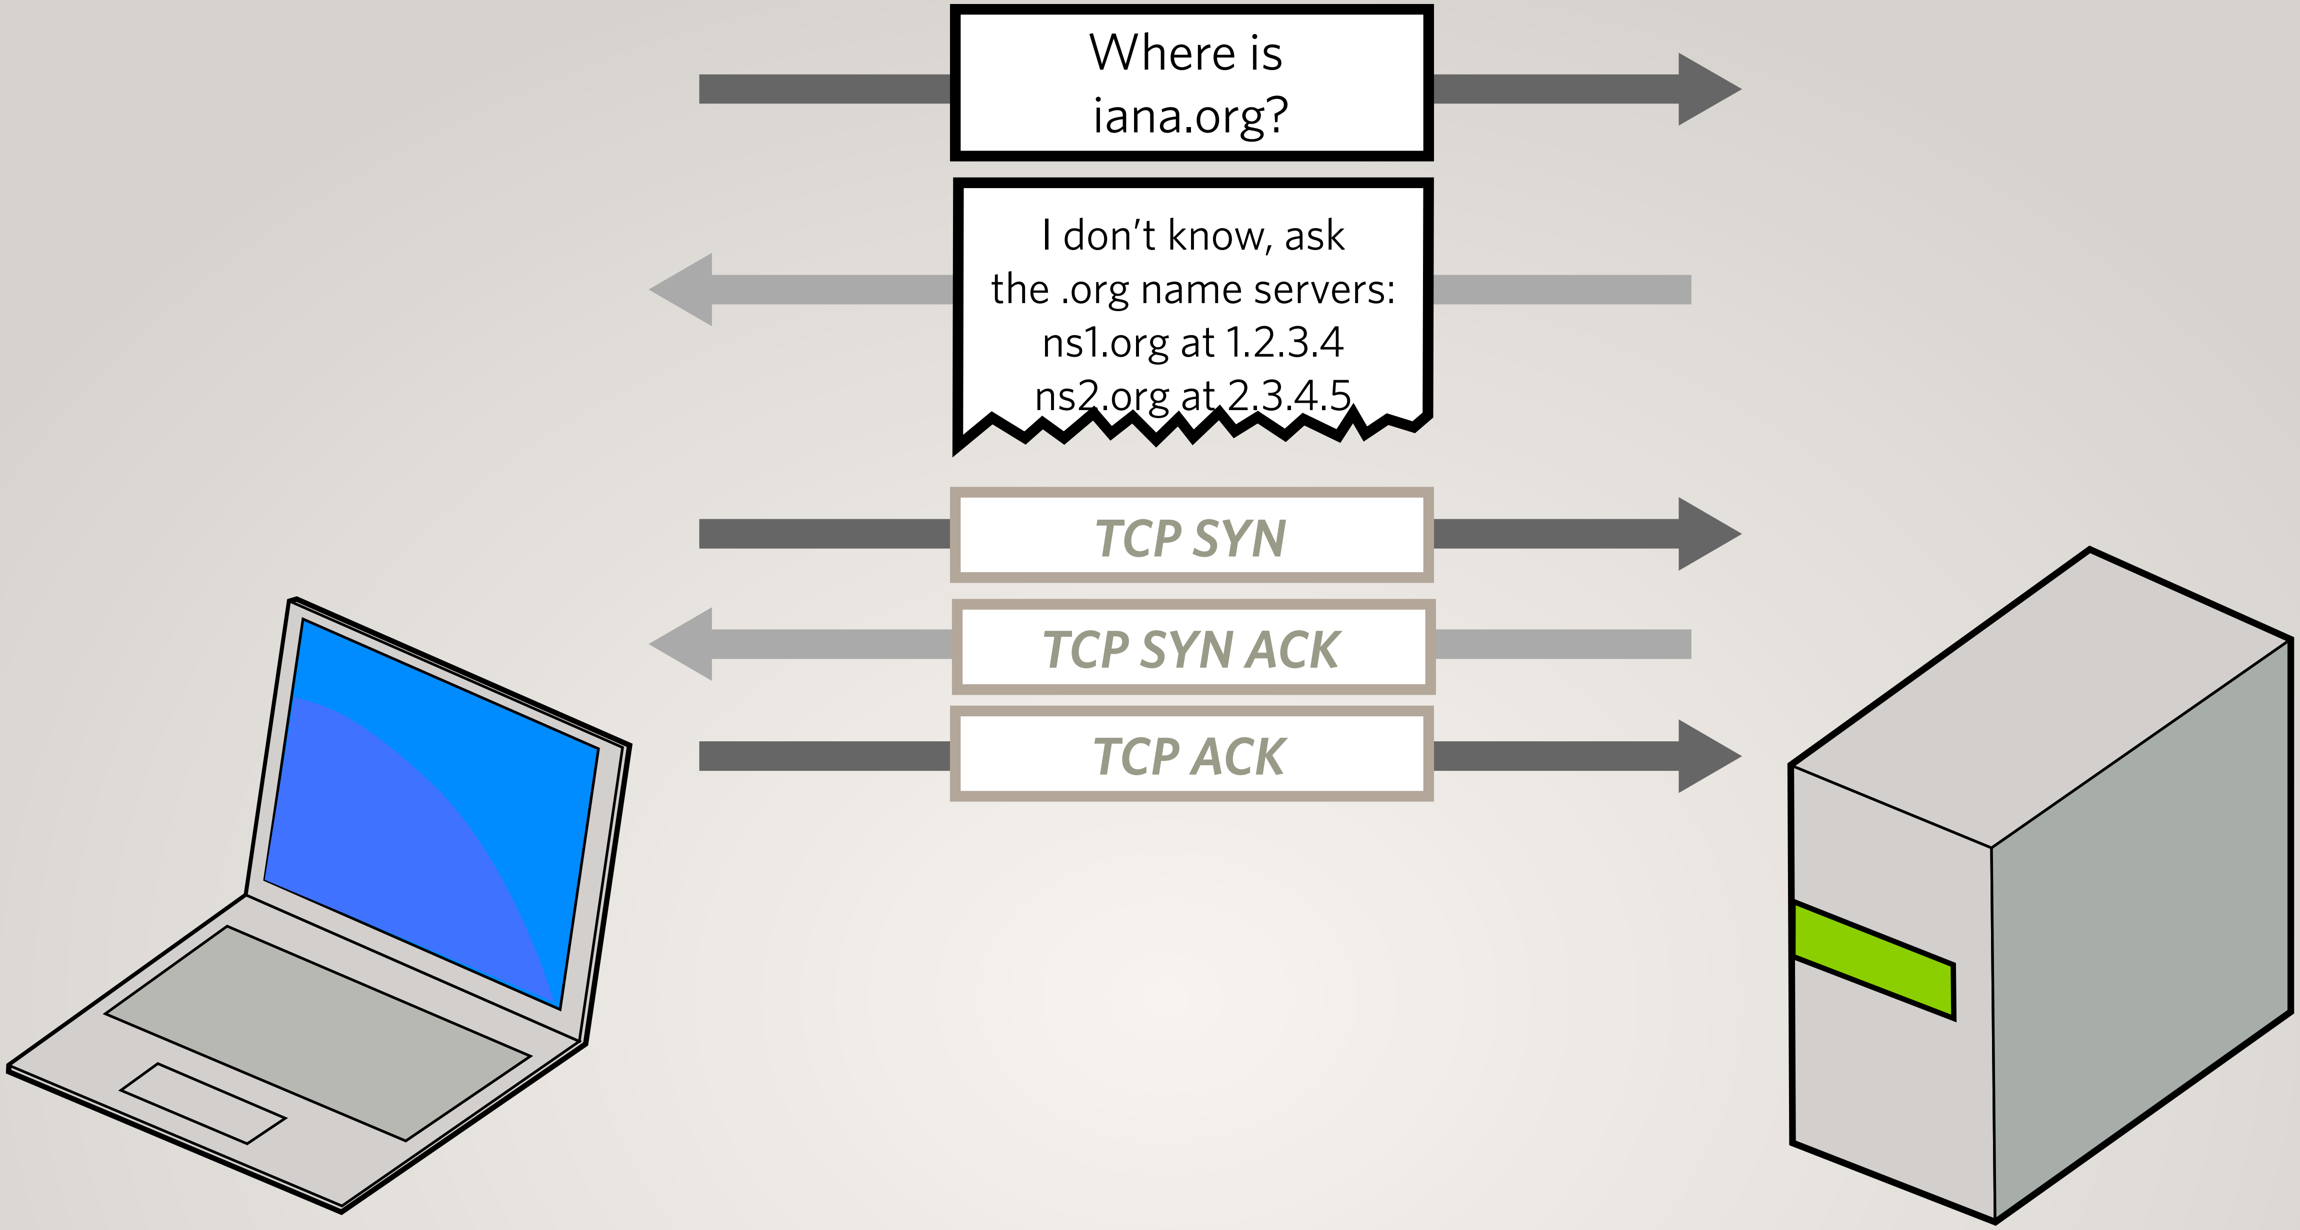

Where is iana.org?

I don't know, ask the .org name servers:
ns1.org at 1.2.3.4
ns2.org at 2.3.4.5

TCP SYN

TCP SYN ACK

TCP ACK

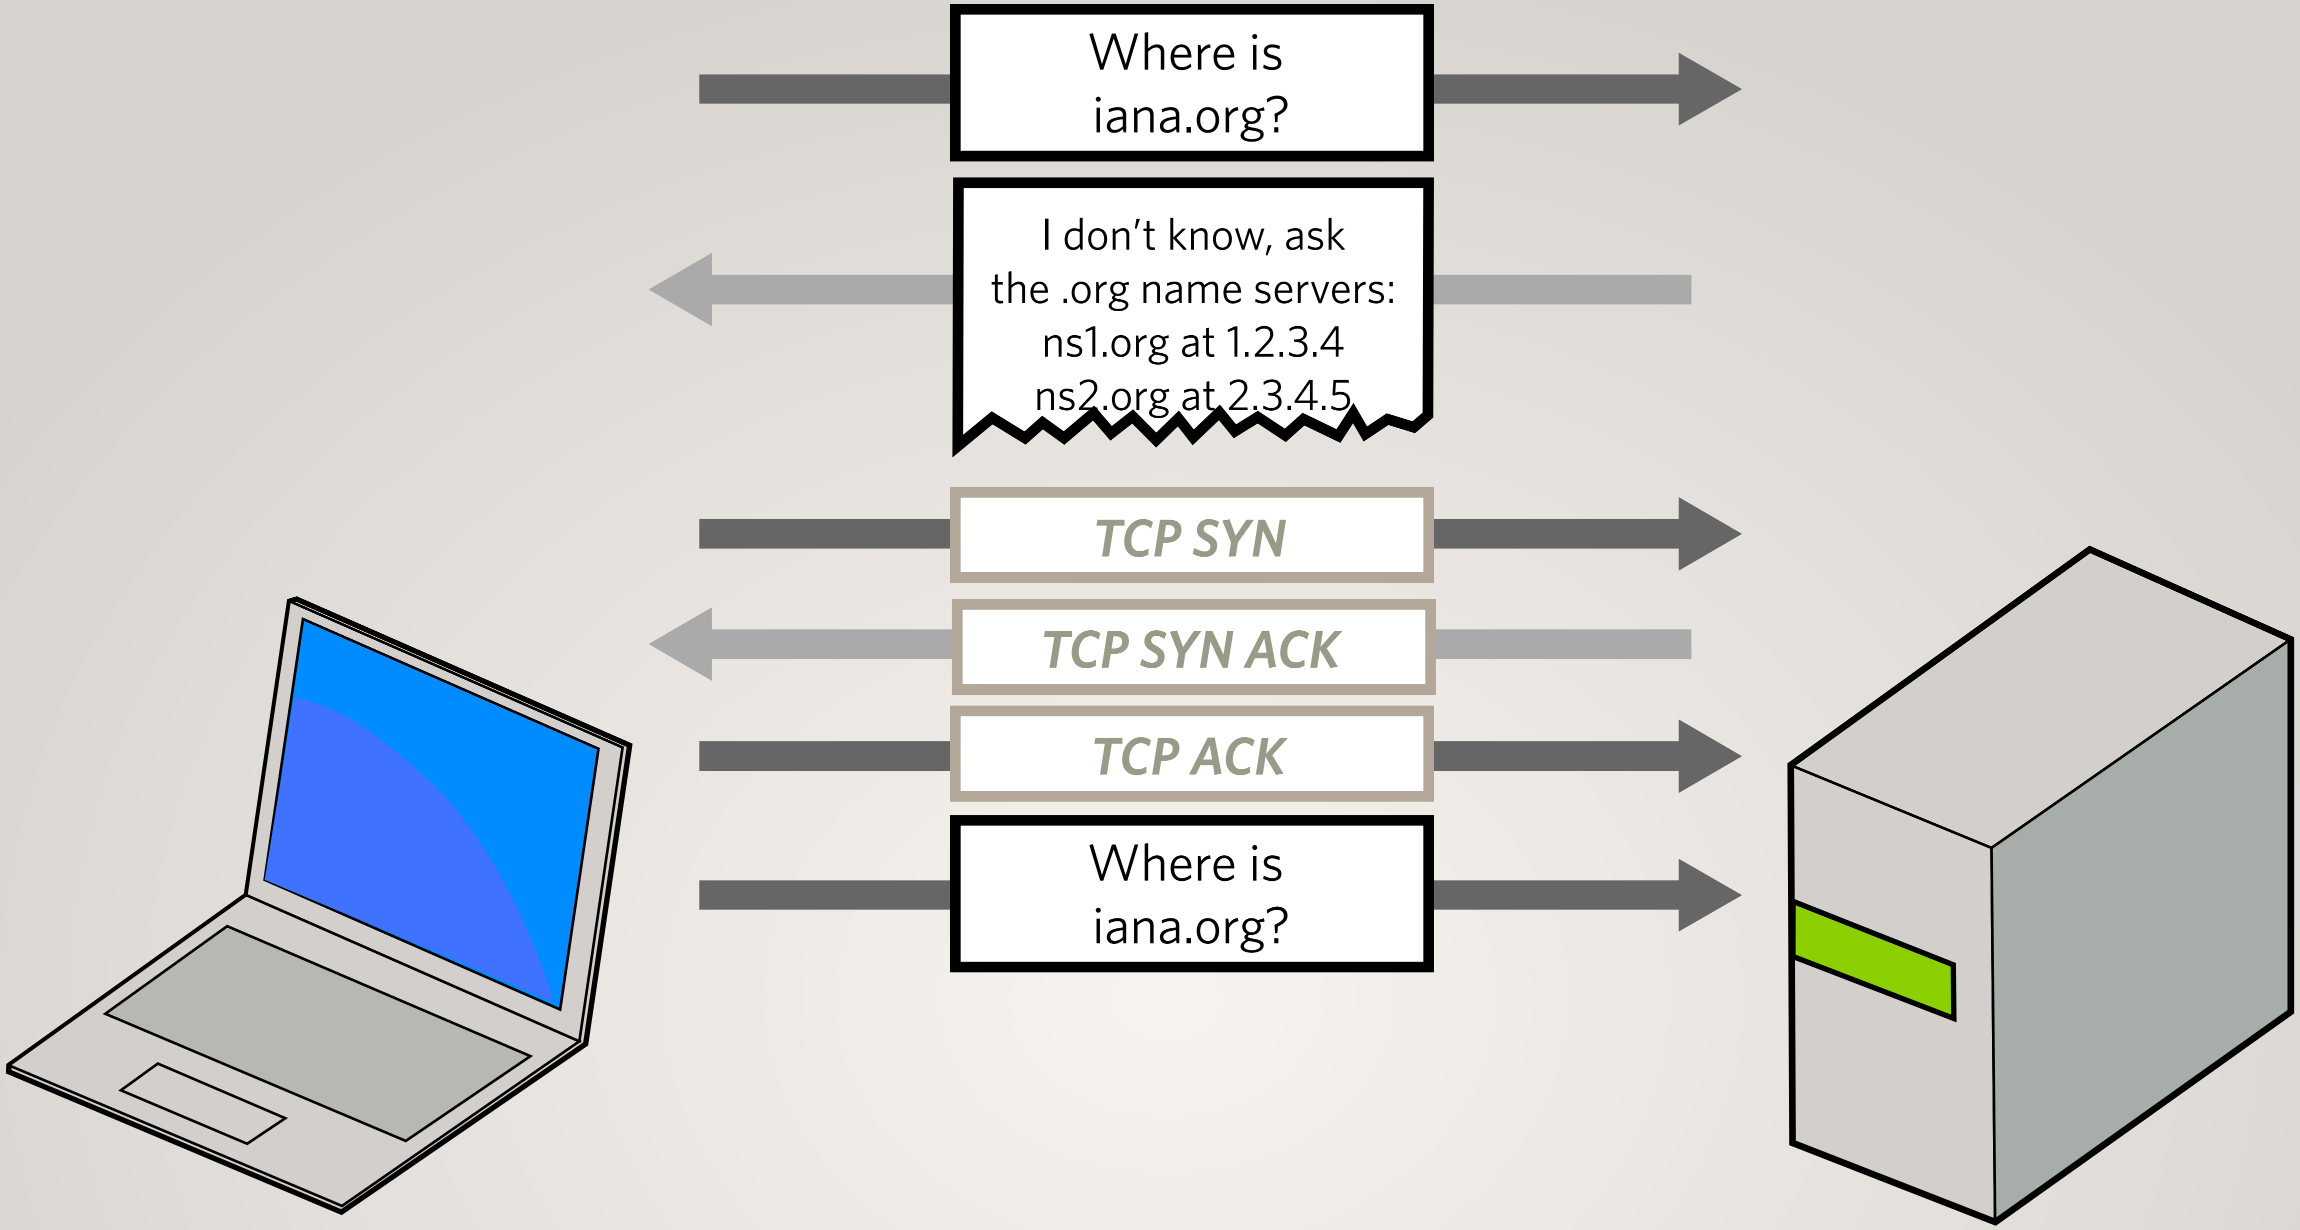

Where is iana.org?

I don’t know, ask the .org name servers: ns1.org at 1.2.3.4 ns2.org at 2.3.4.5

TCP SYN

TCP SYN ACK

TCP ACK

Where is iana.org?

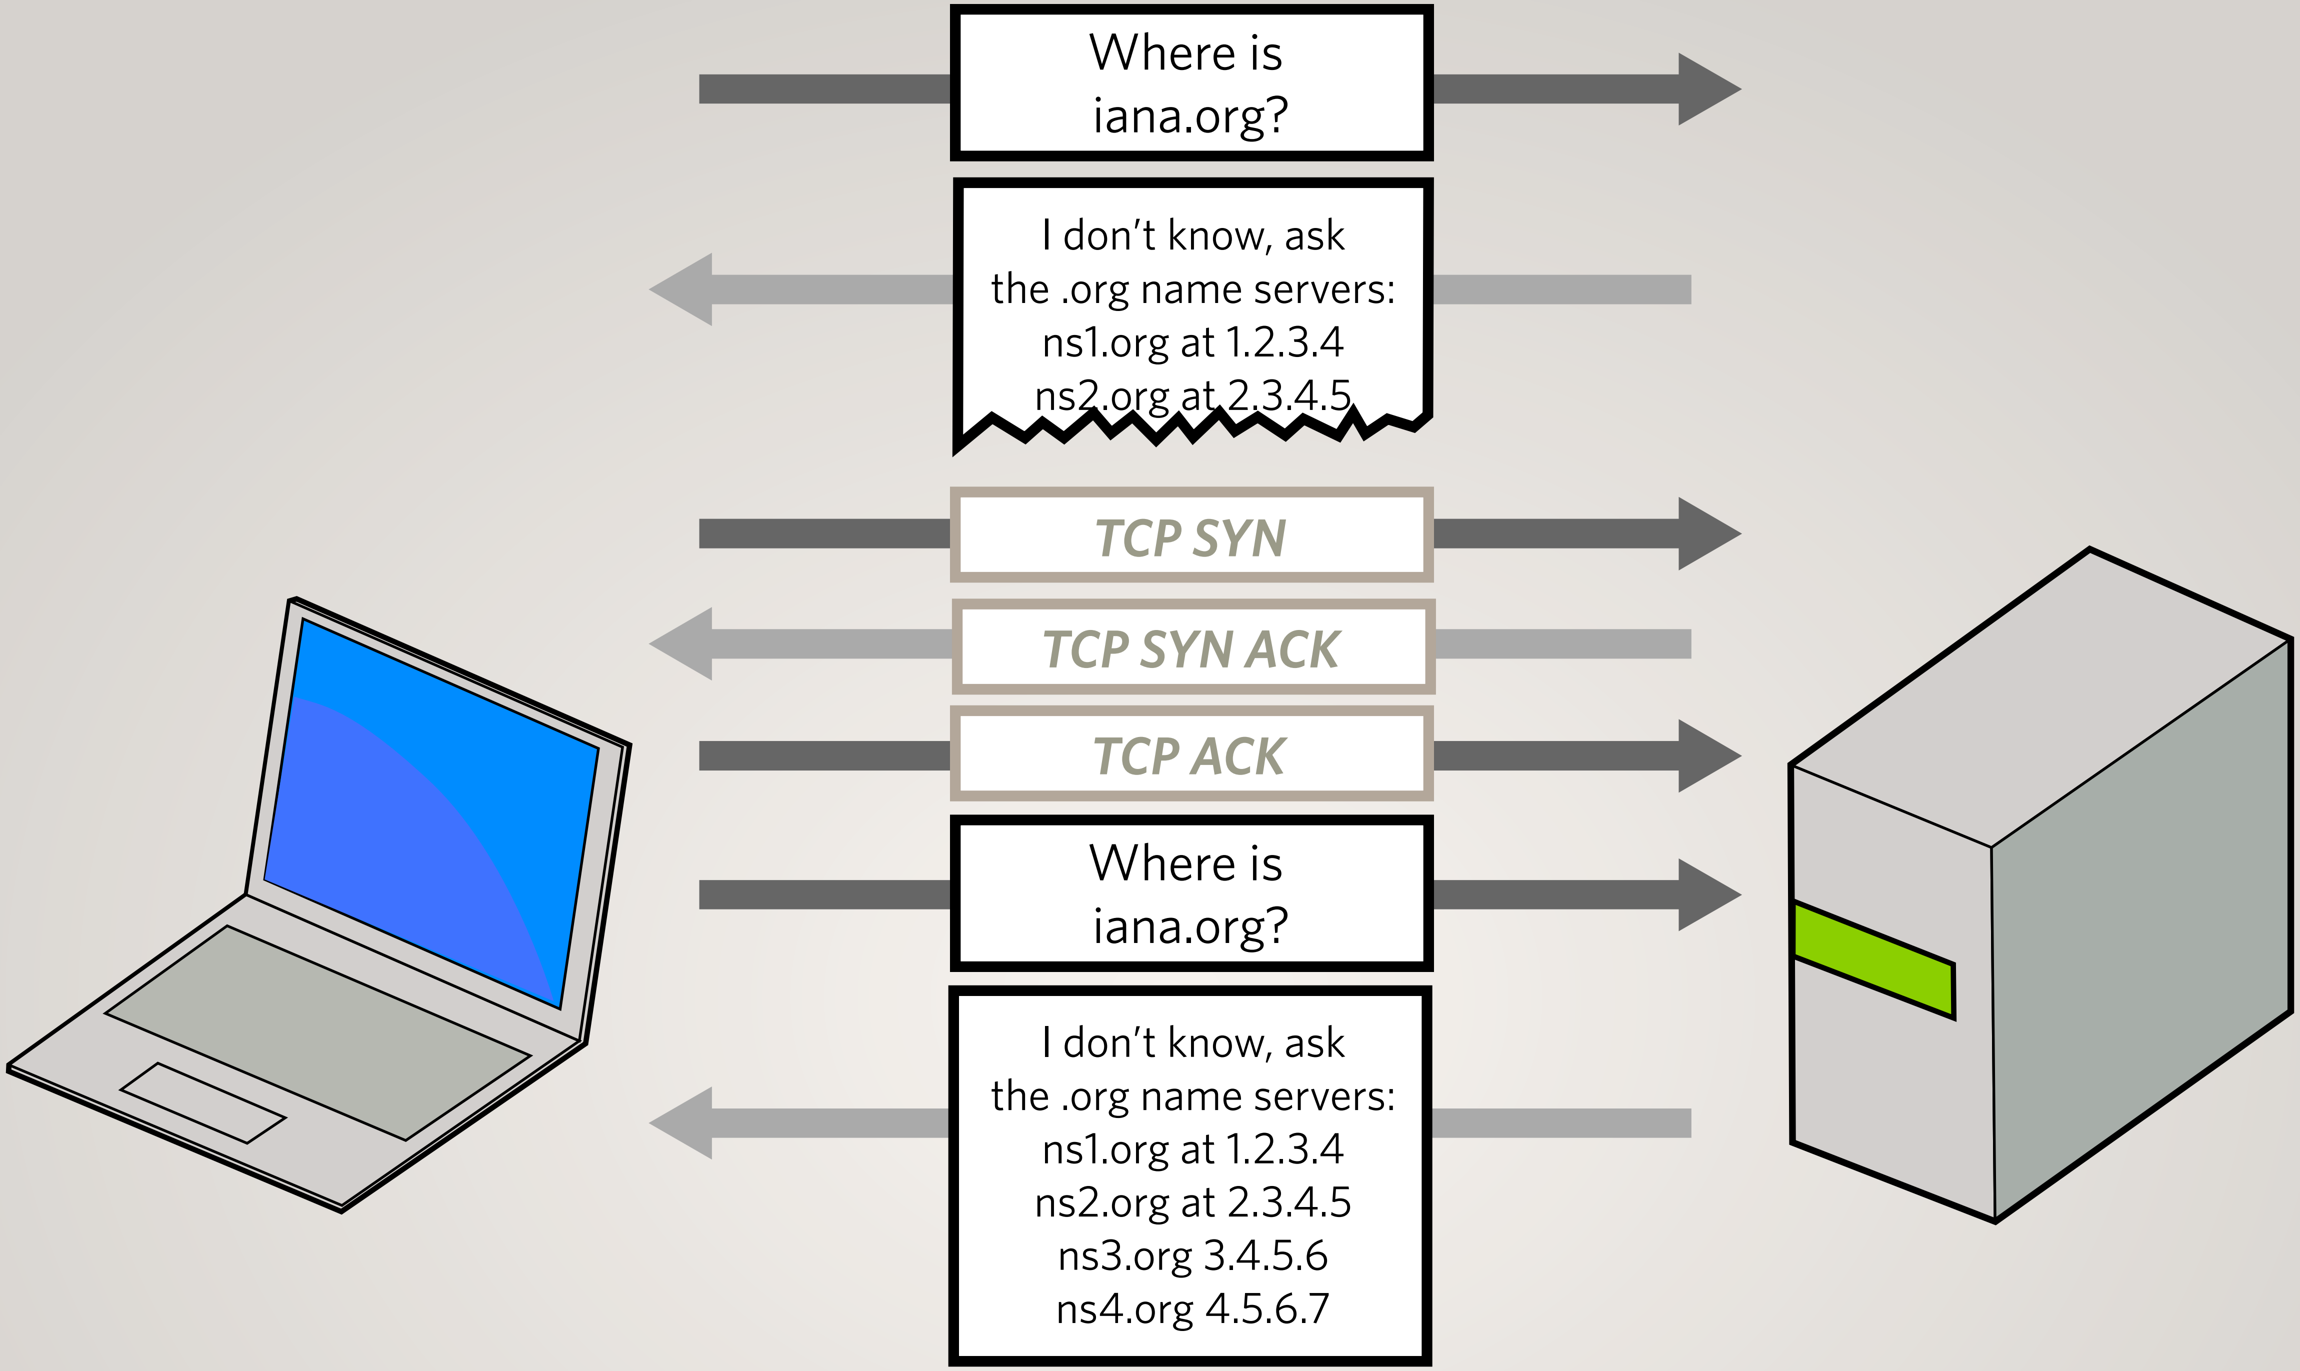

Where is iana.org?

I don’t know, ask the .org name servers:
ns1.org at 1.2.3.4
ns2.org at 2.3.4.5

TCP SYN

TCP SYN ACK

TCP ACK

Where is iana.org?

I don’t know, ask the .org name servers:
ns1.org at 1.2.3.4
ns2.org at 2.3.4.5
ns3.org 3.4.5.6
ns4.org 4.5.6.7

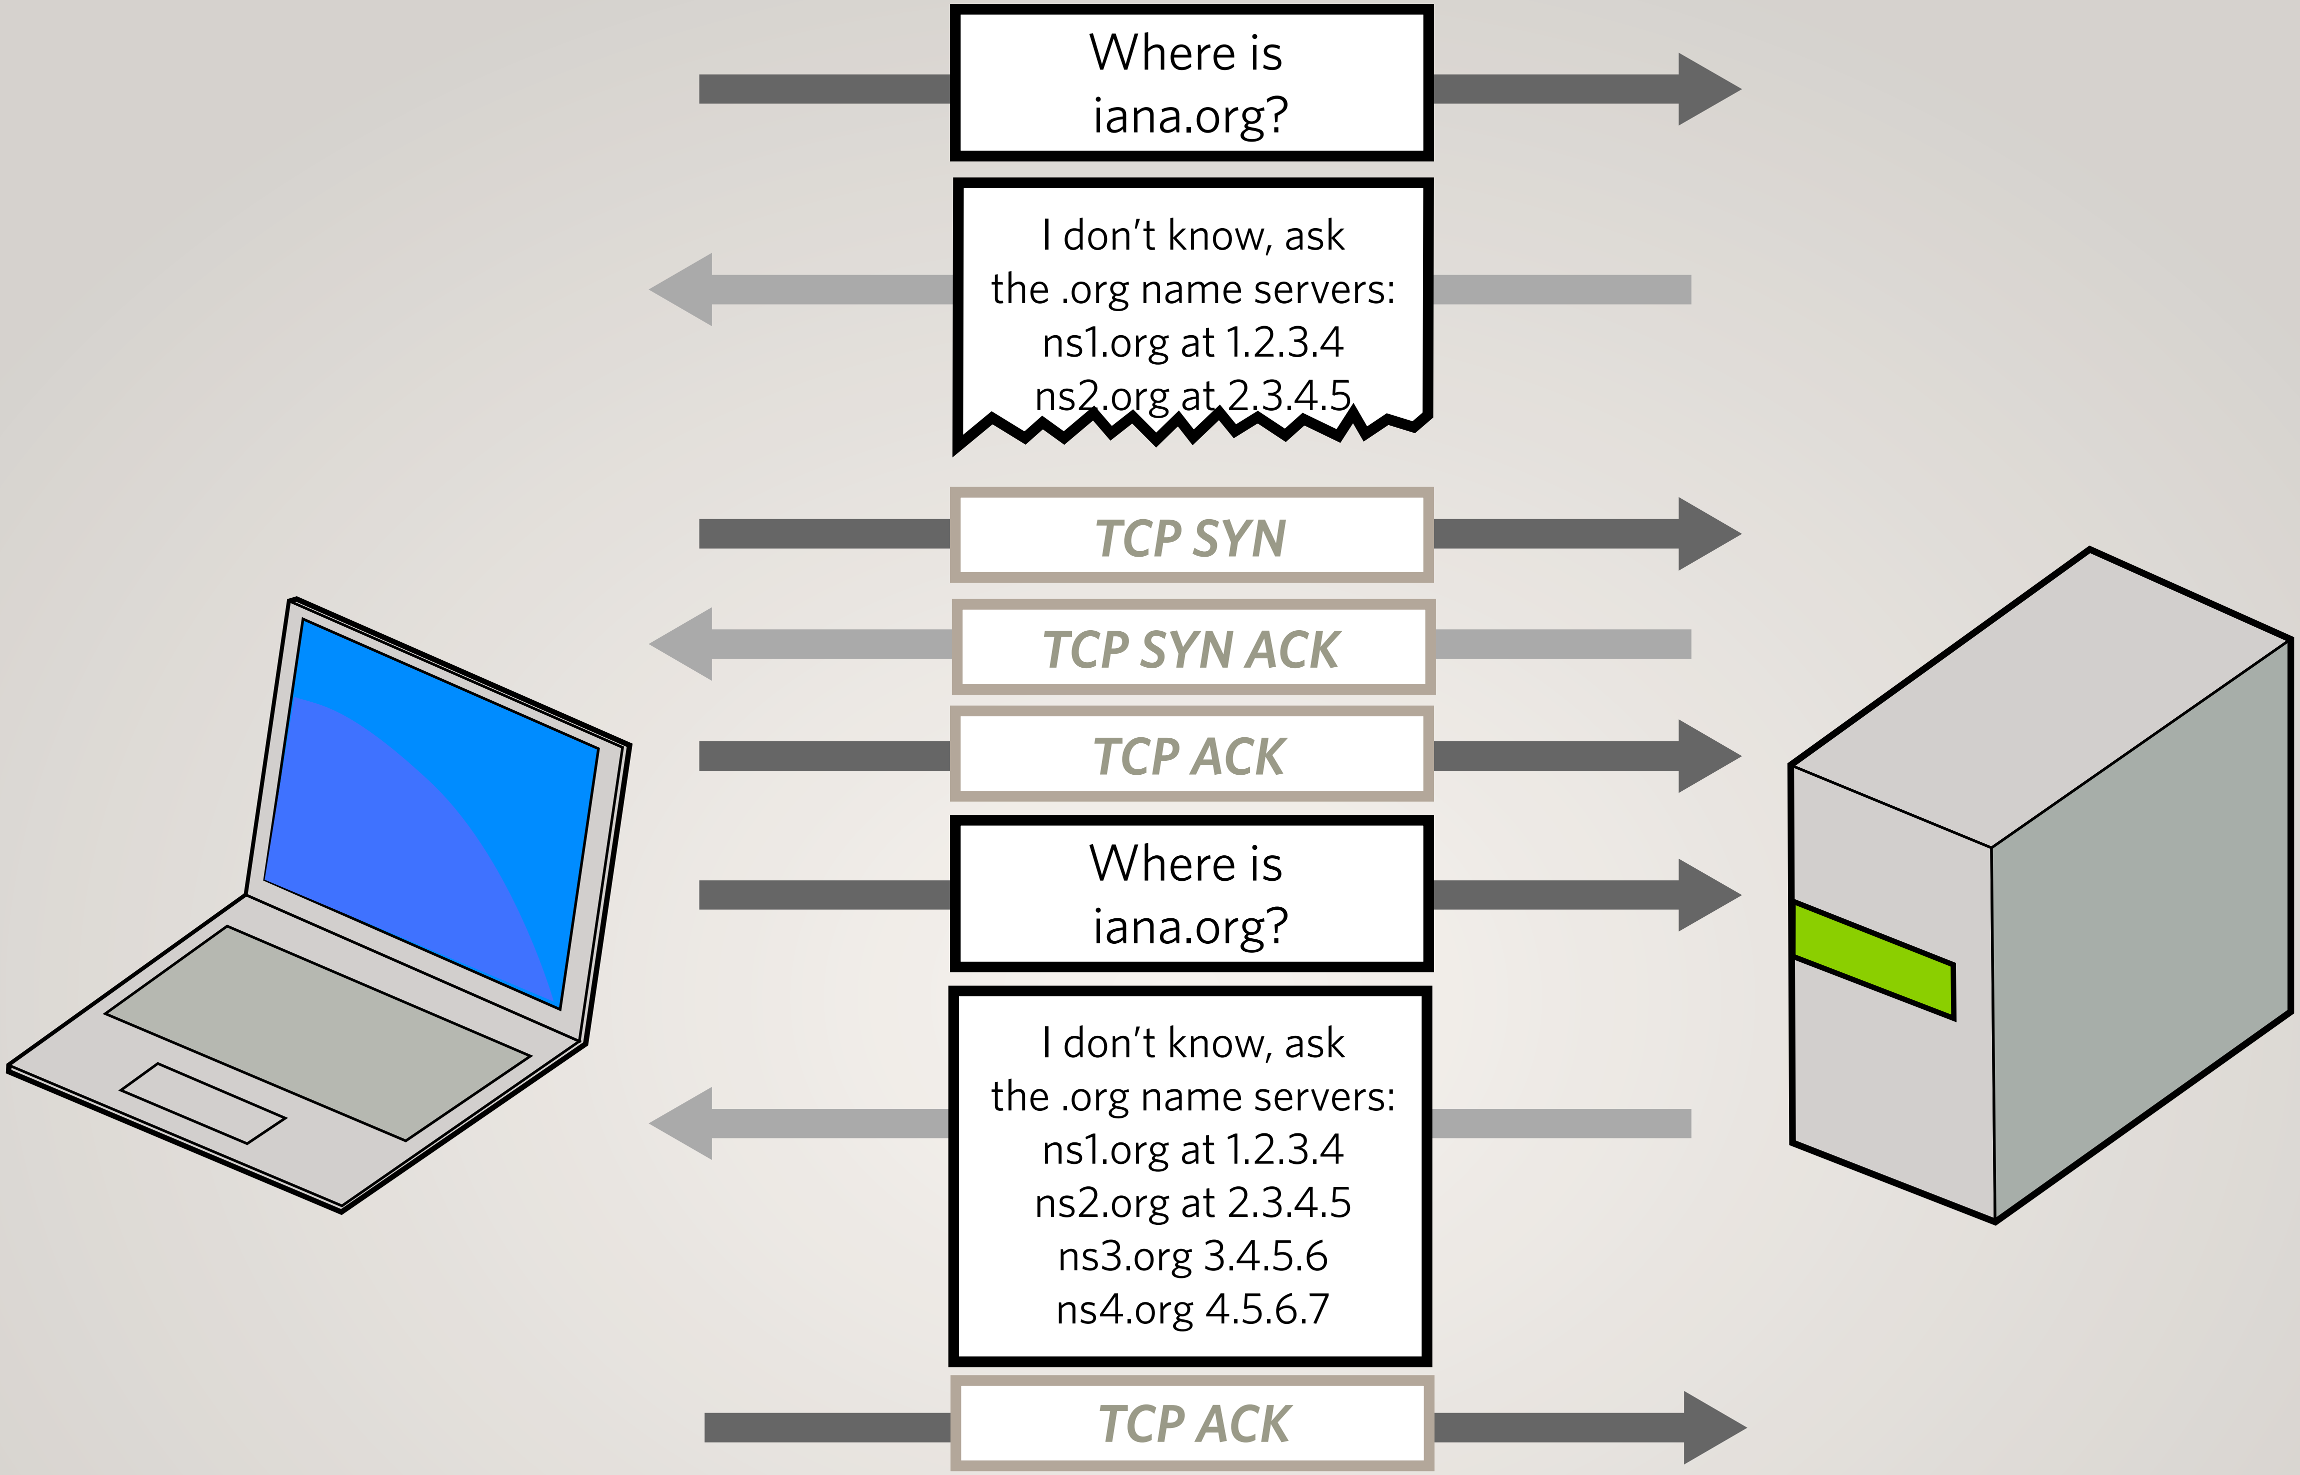

Where is iana.org?

I don’t know, ask the .org name servers:
ns1.org at 1.2.3.4
ns2.org at 2.3.4.5

TCP SYN

TCP SYN ACK

TCP ACK

Where is iana.org?

I don’t know, ask the .org name servers:
ns1.org at 1.2.3.4
ns2.org at 2.3.4.5
ns3.org 3.4.5.6
ns4.org 4.5.6.7

TCP ACK

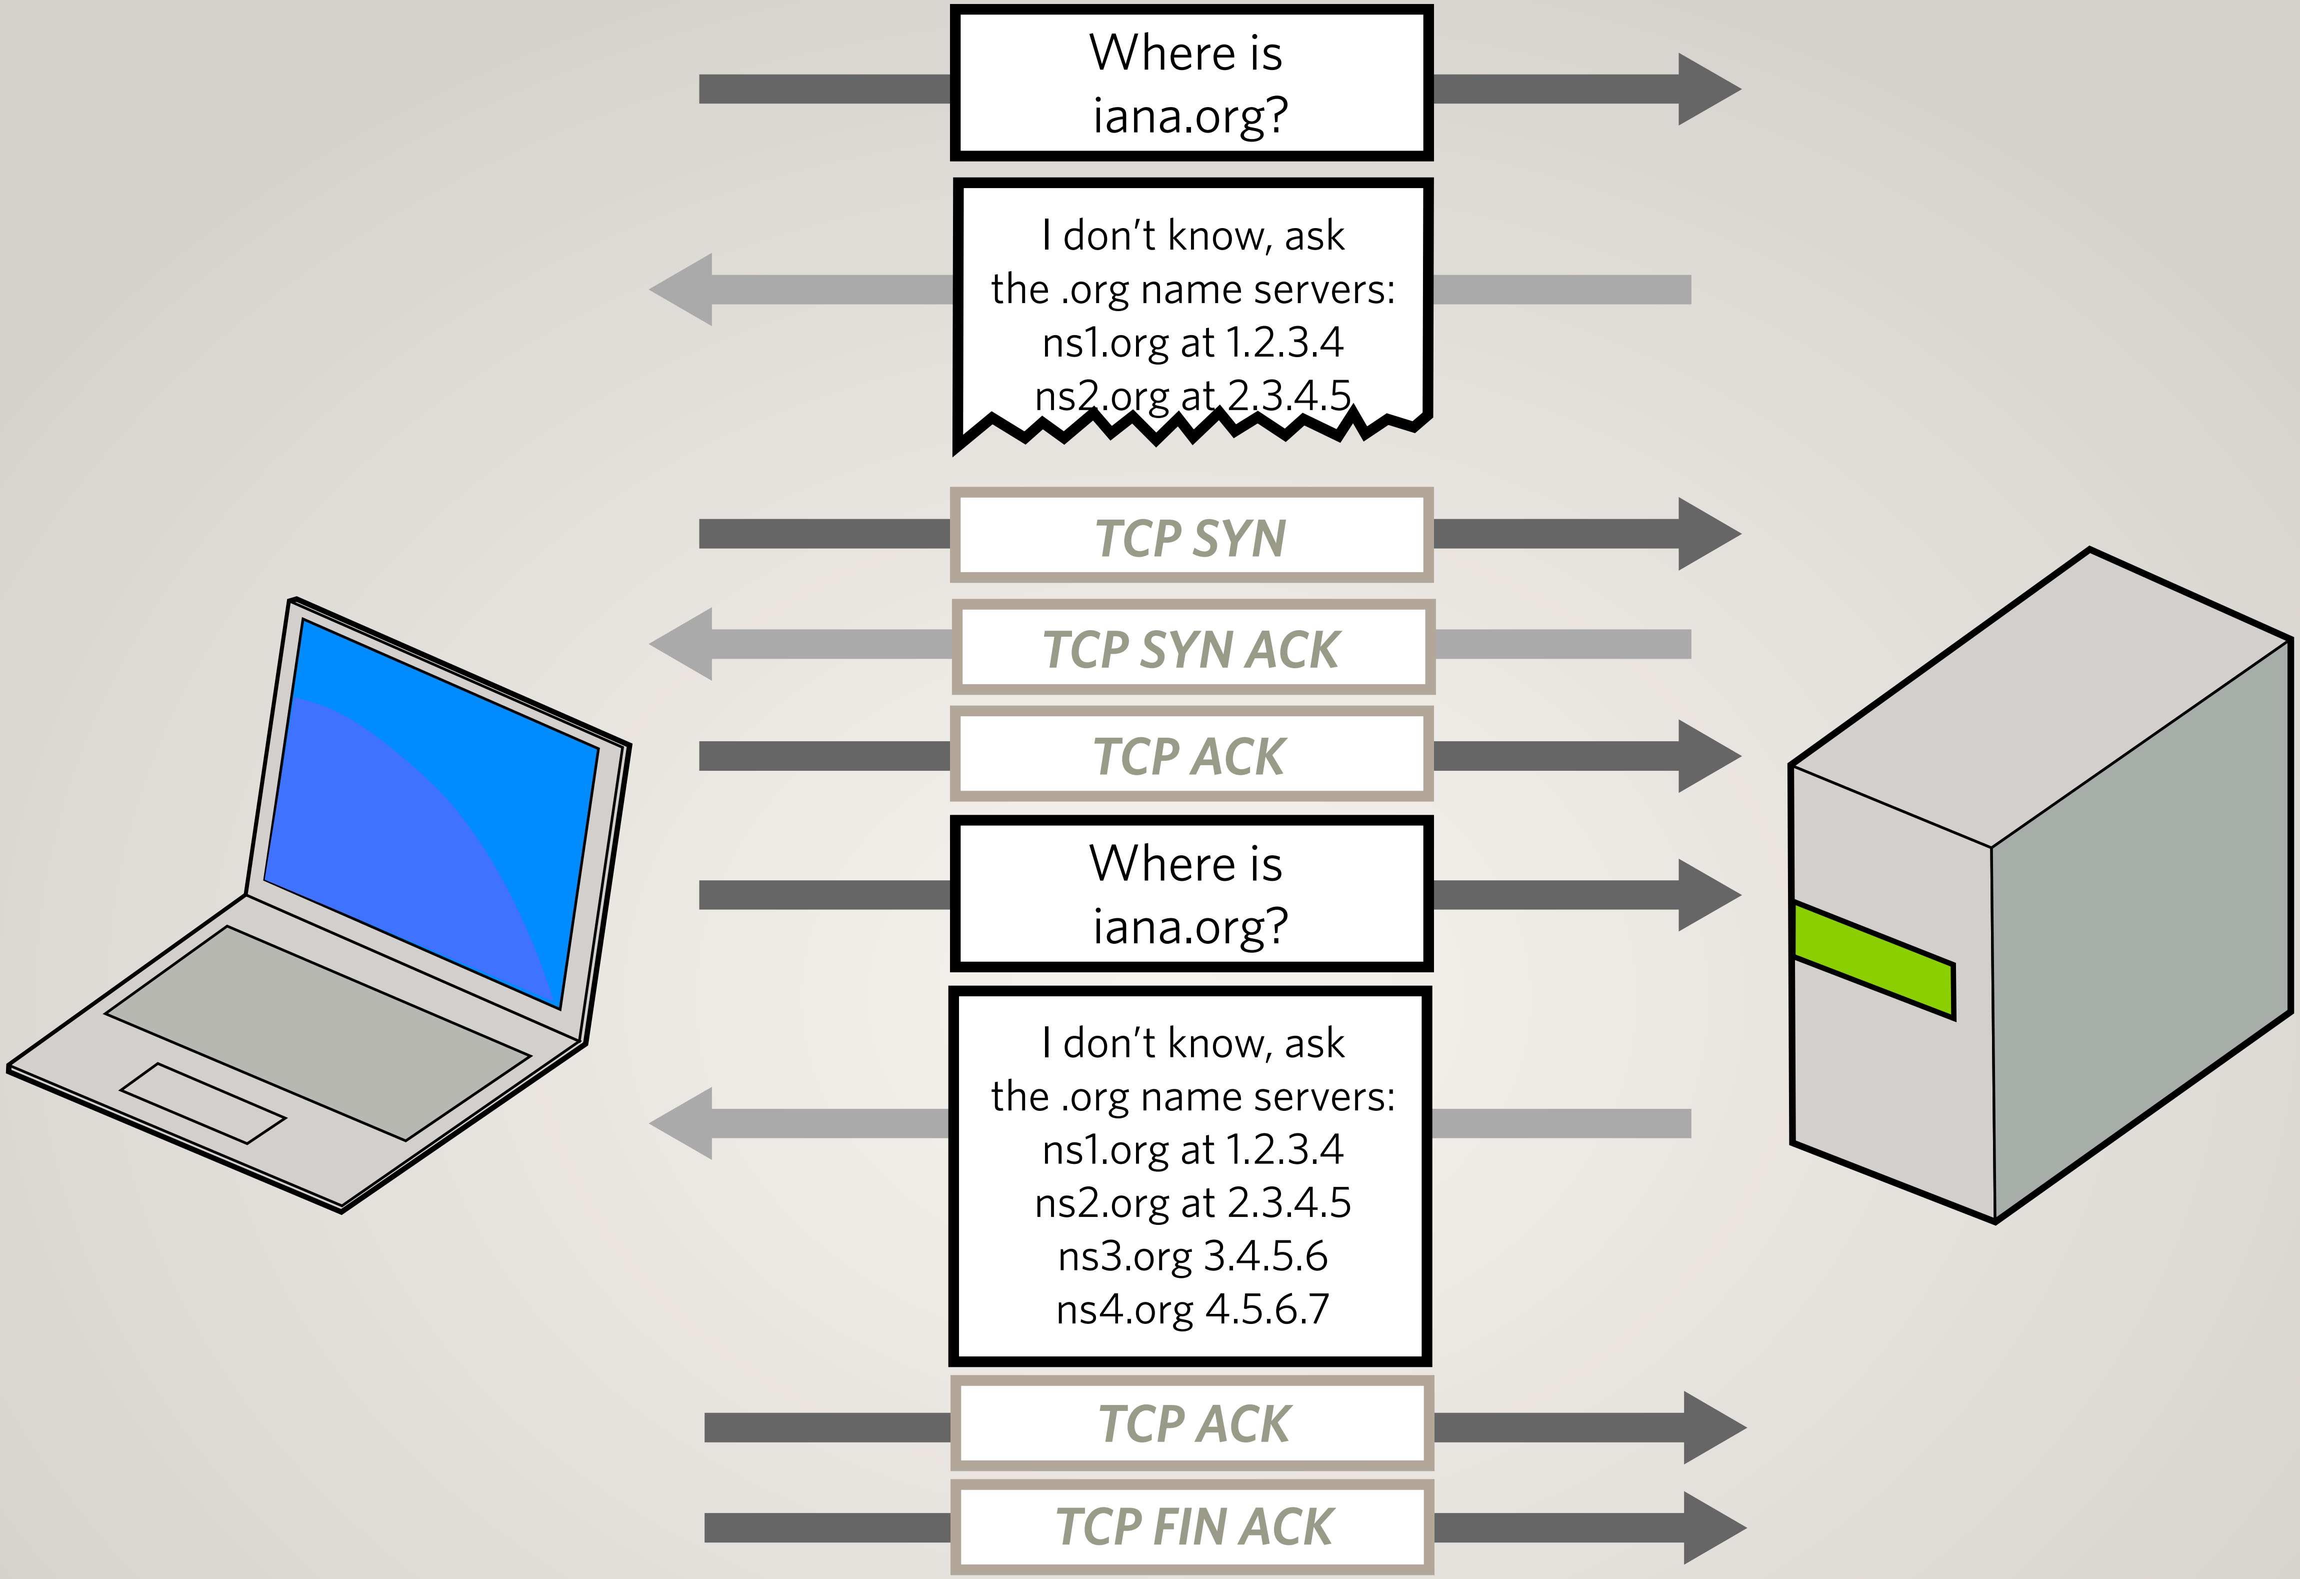

Where is iana.org?

I don't know, ask the .org name servers:
ns1.org at 1.2.3.4
ns2.org at 2.3.4.5

TCP SYN

TCP SYN ACK

TCP ACK

Where is iana.org?

I don't know, ask the .org name servers:
ns1.org at 1.2.3.4
ns2.org at 2.3.4.5
ns3.org 3.4.5.6
ns4.org 4.5.6.7

TCP ACK

TCP FIN ACK

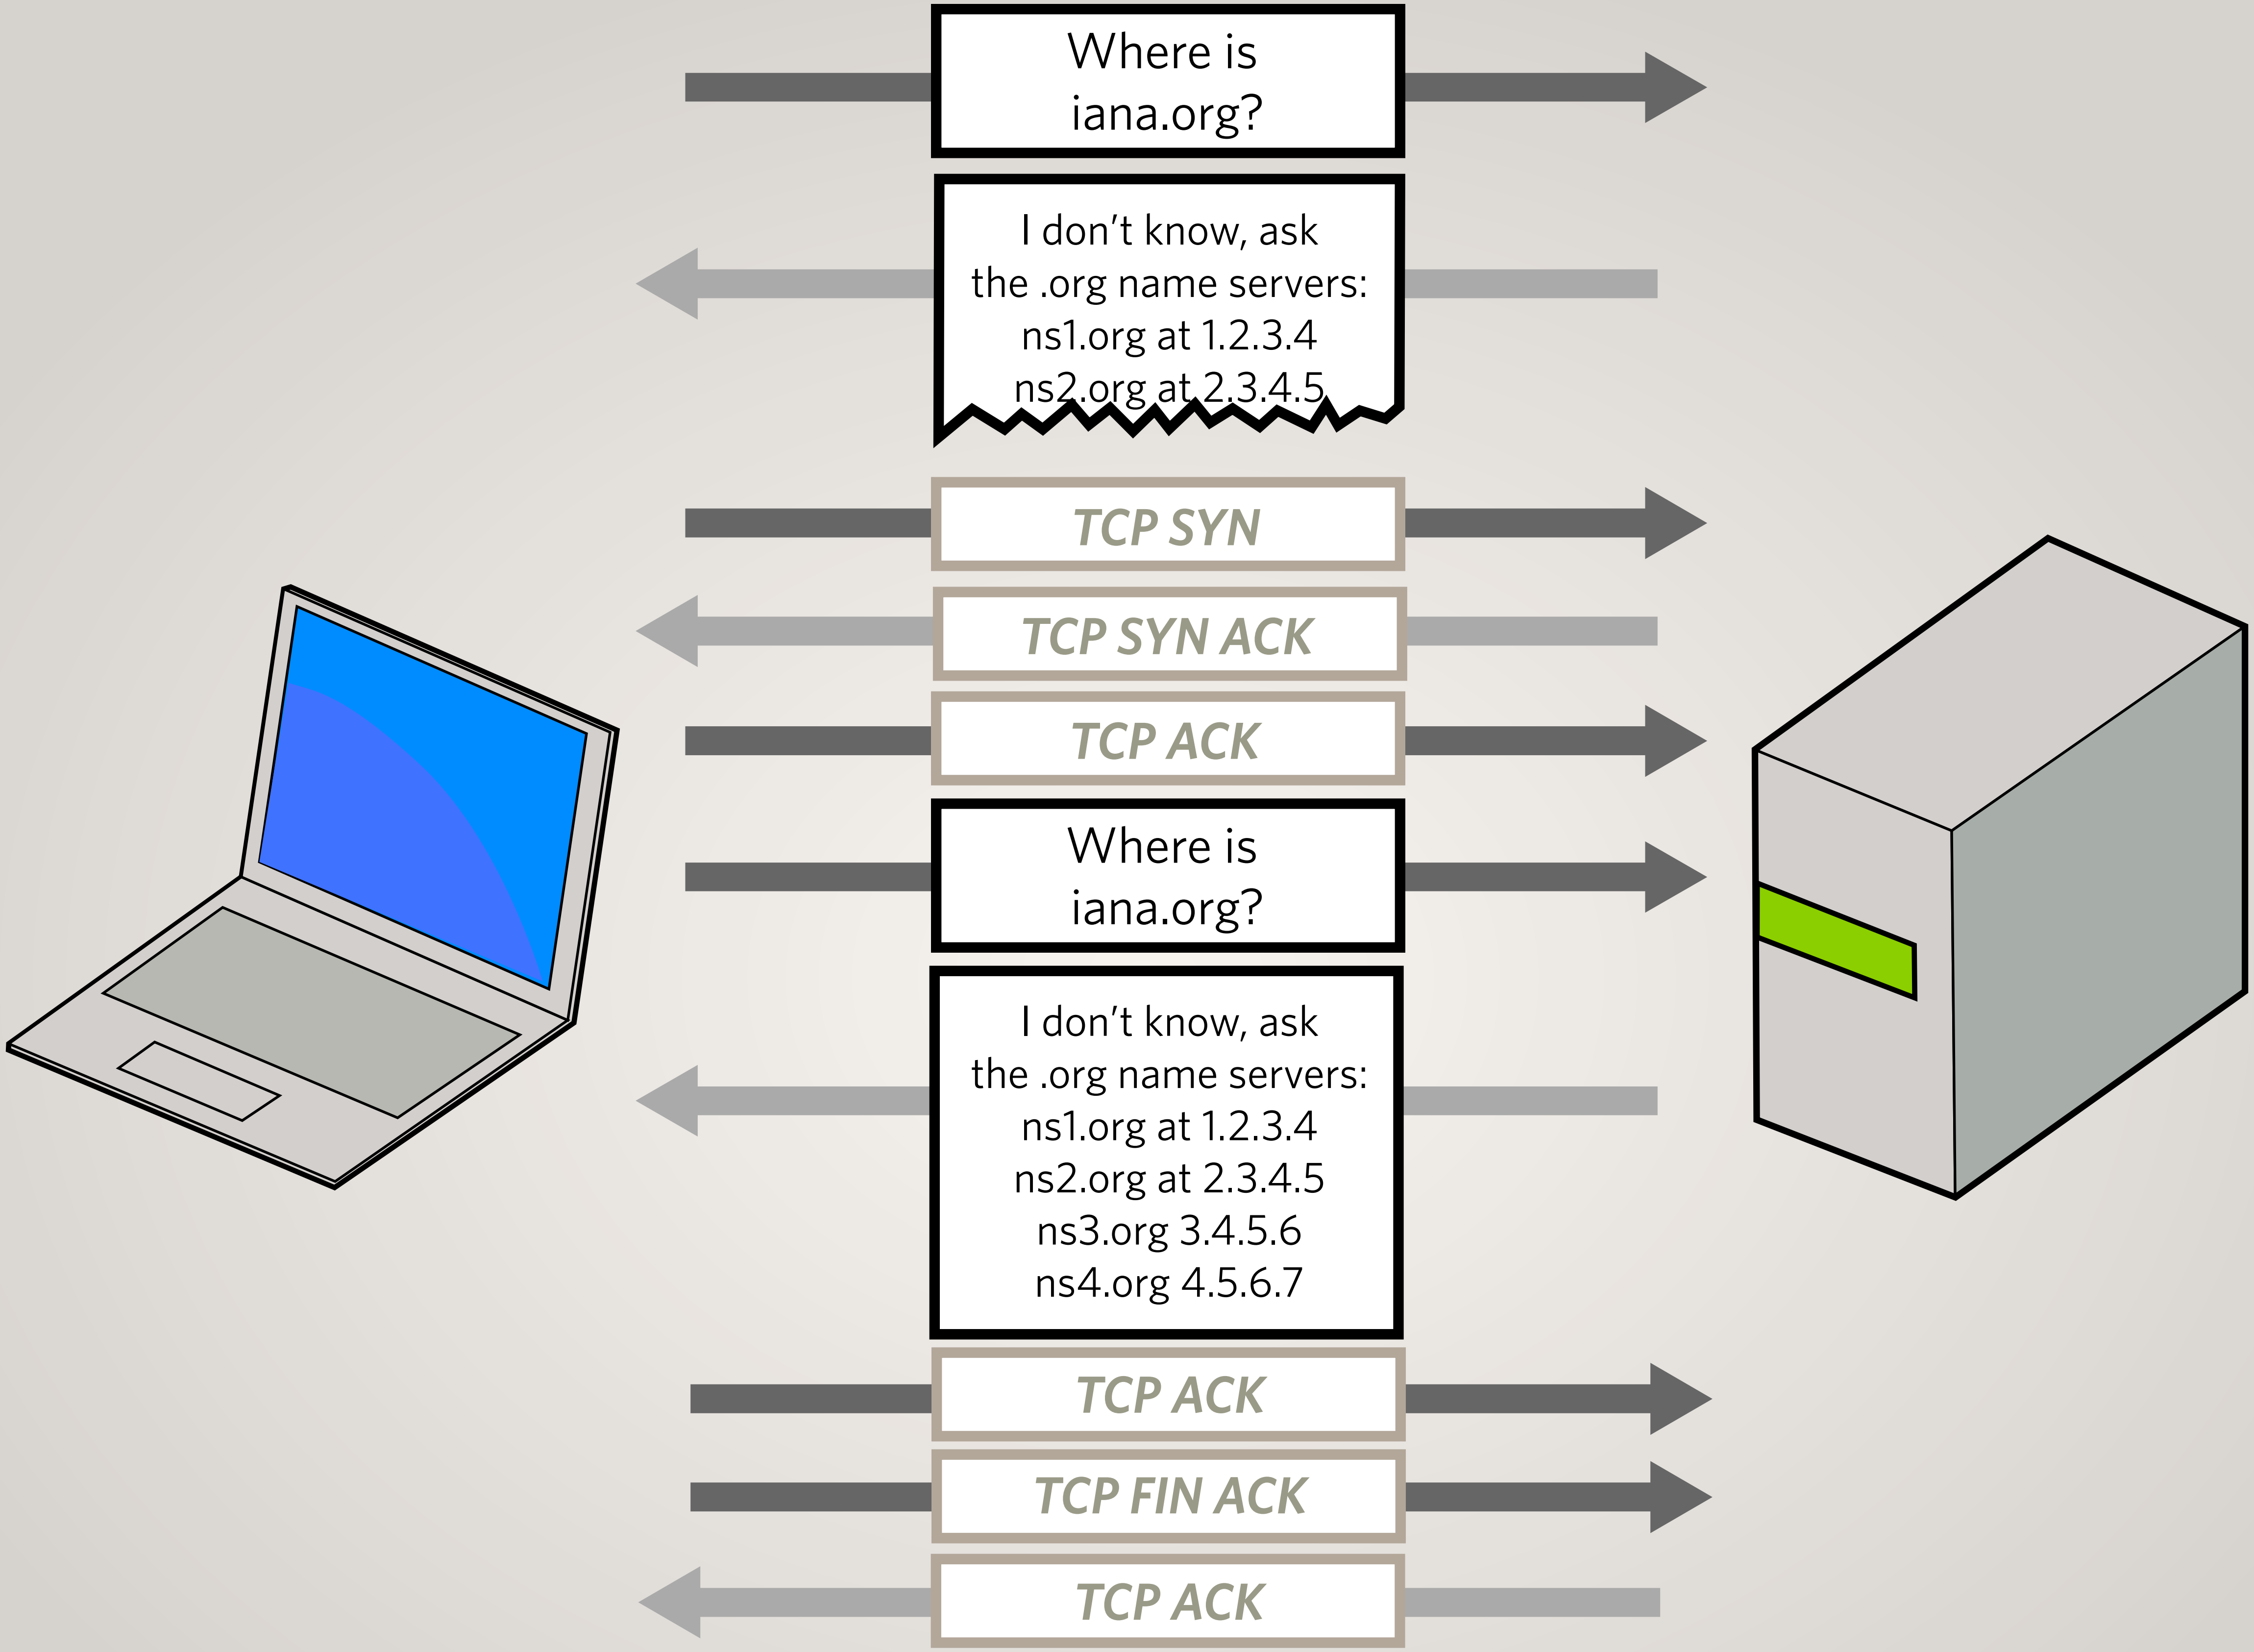

Where is iana.org?

I don't know, ask the .org name servers:
ns1.org at 1.2.3.4
ns2.org at 2.3.4.5

TCP SYN

TCP SYN ACK

TCP ACK

Where is iana.org?

I don't know, ask the .org name servers:
ns1.org at 1.2.3.4
ns2.org at 2.3.4.5
ns3.org 3.4.5.6
ns4.org 4.5.6.7

TCP ACK

TCP FIN ACK

TCP ACK

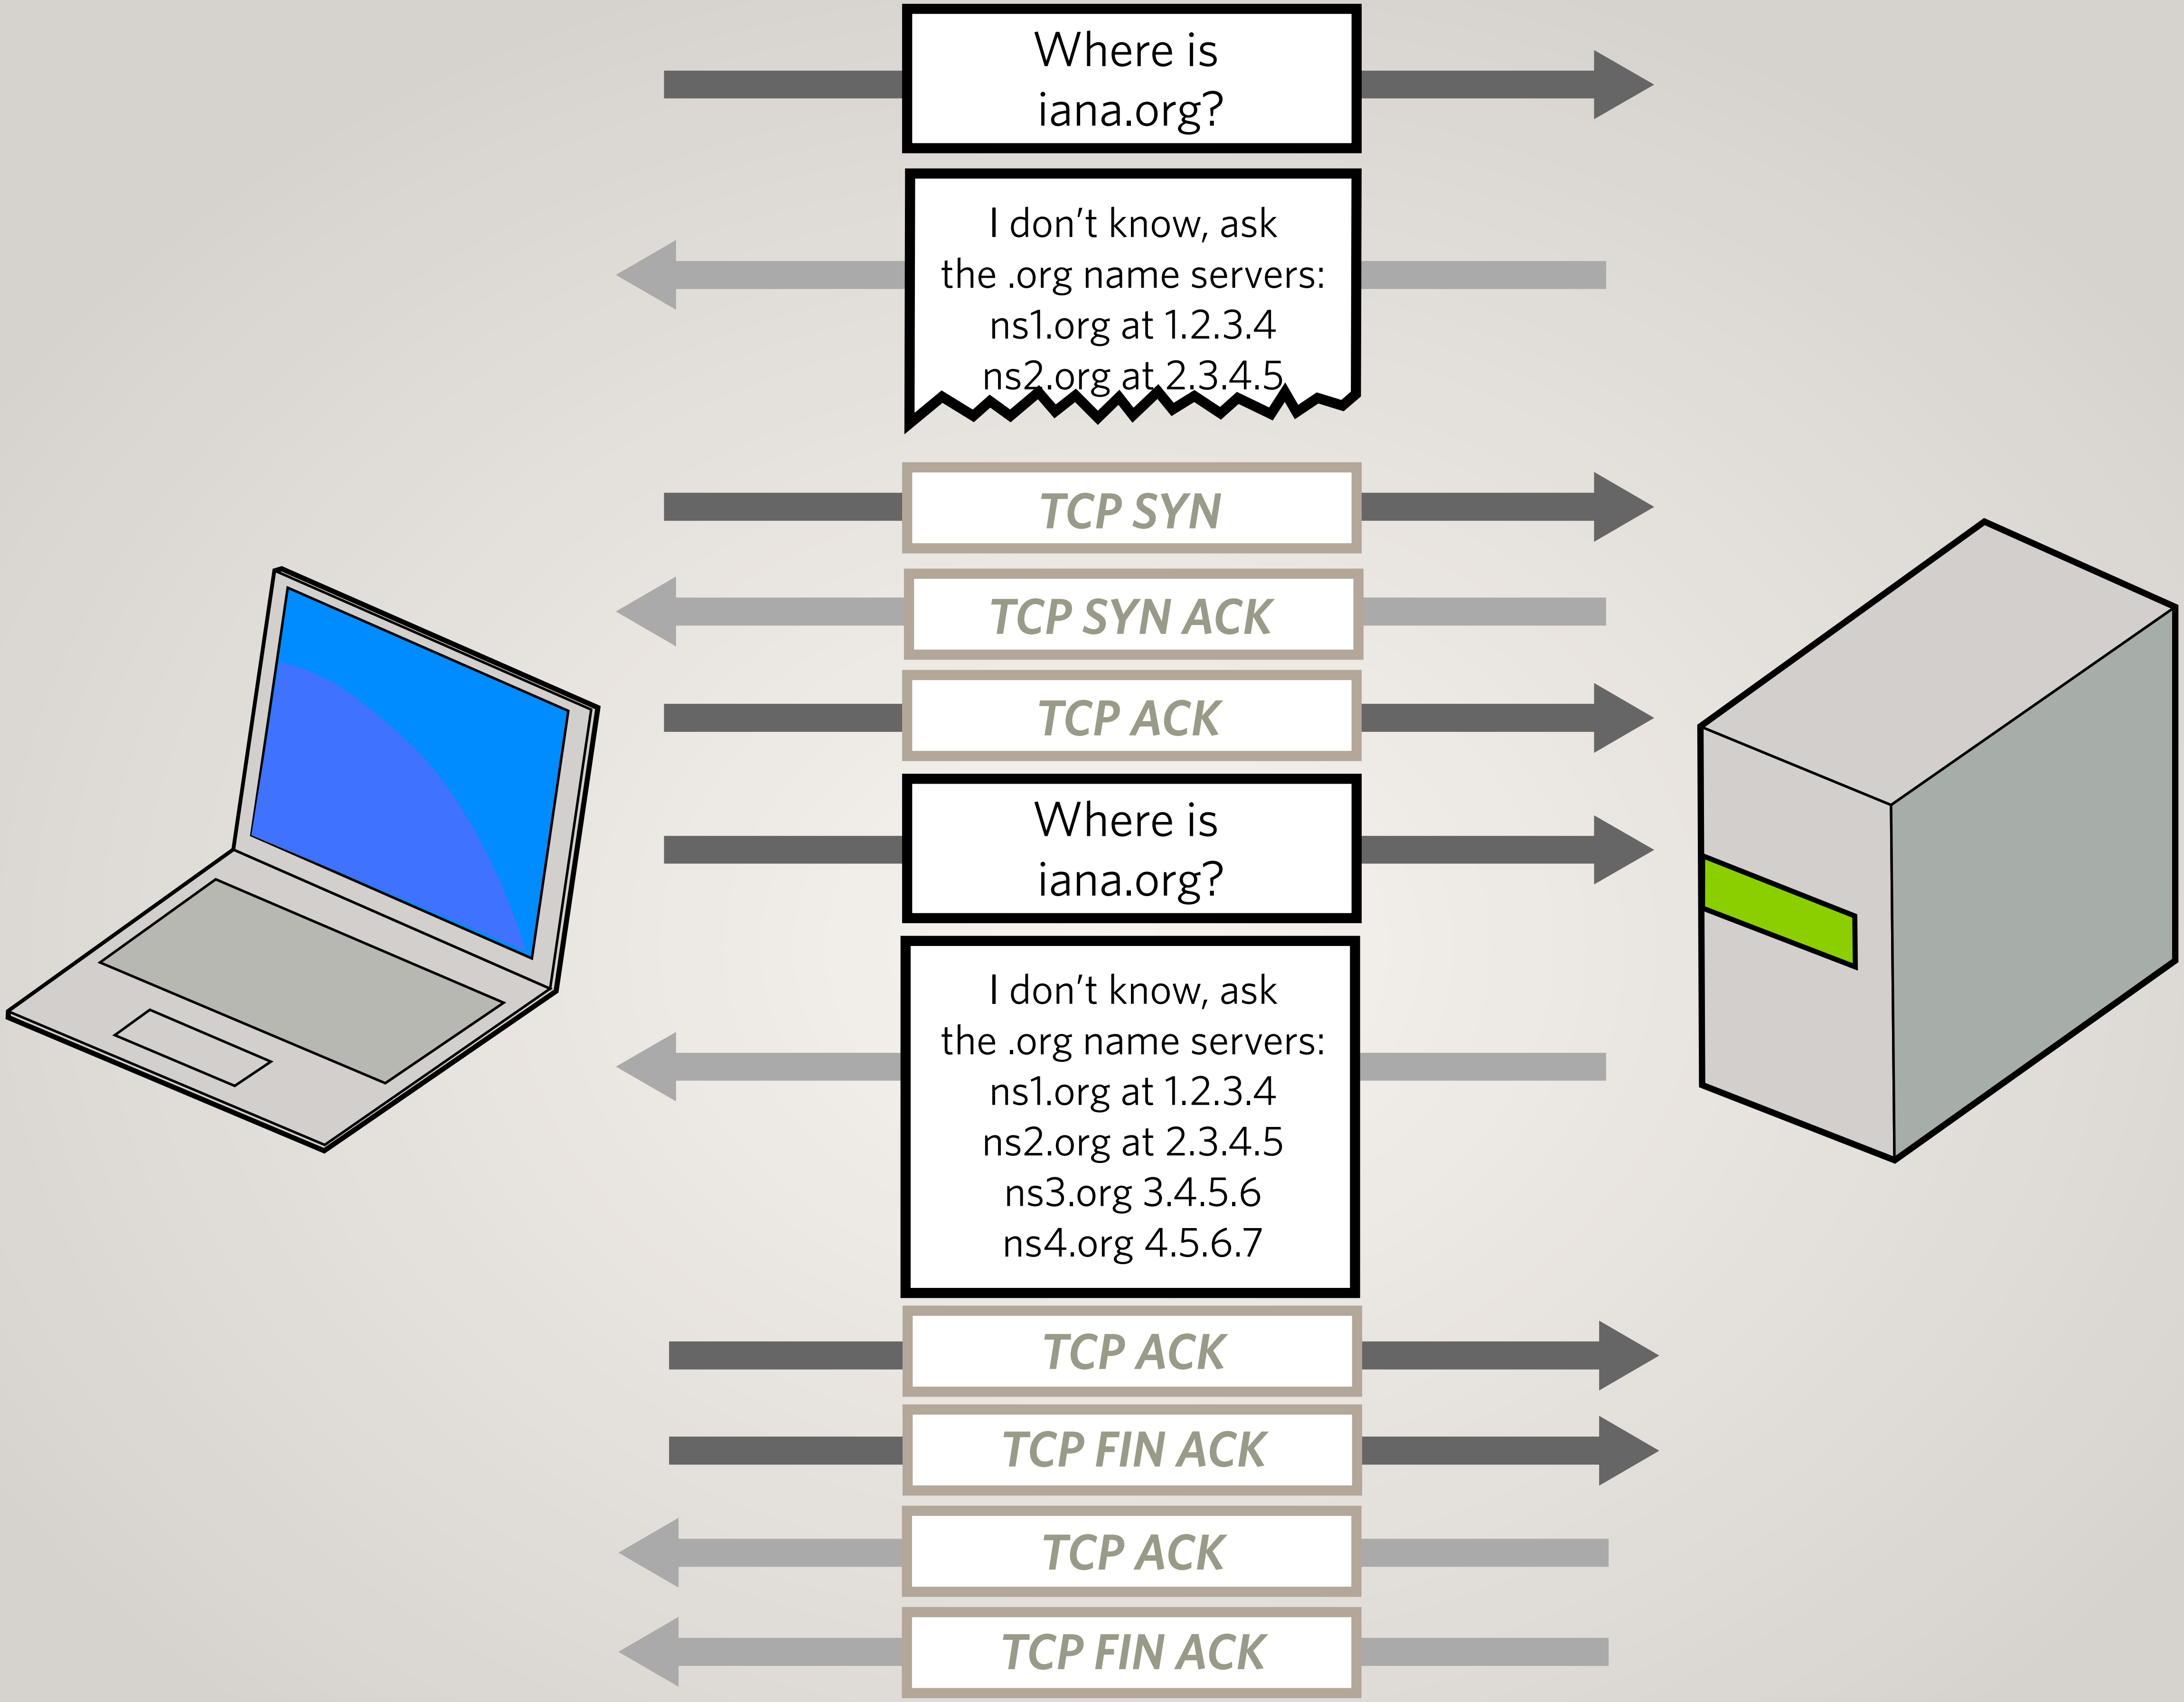

Where is iana.org?

I don't know, ask the .org name servers:
ns1.org at 1.2.3.4
ns2.org at 2.3.4.5

TCP SYN

TCP SYN ACK

TCP ACK

Where is iana.org?

I don't know, ask the .org name servers:
ns1.org at 1.2.3.4
ns2.org at 2.3.4.5
ns3.org 3.4.5.6
ns4.org 4.5.6.7

TCP ACK

TCP FIN ACK

TCP ACK

TCP FIN ACK

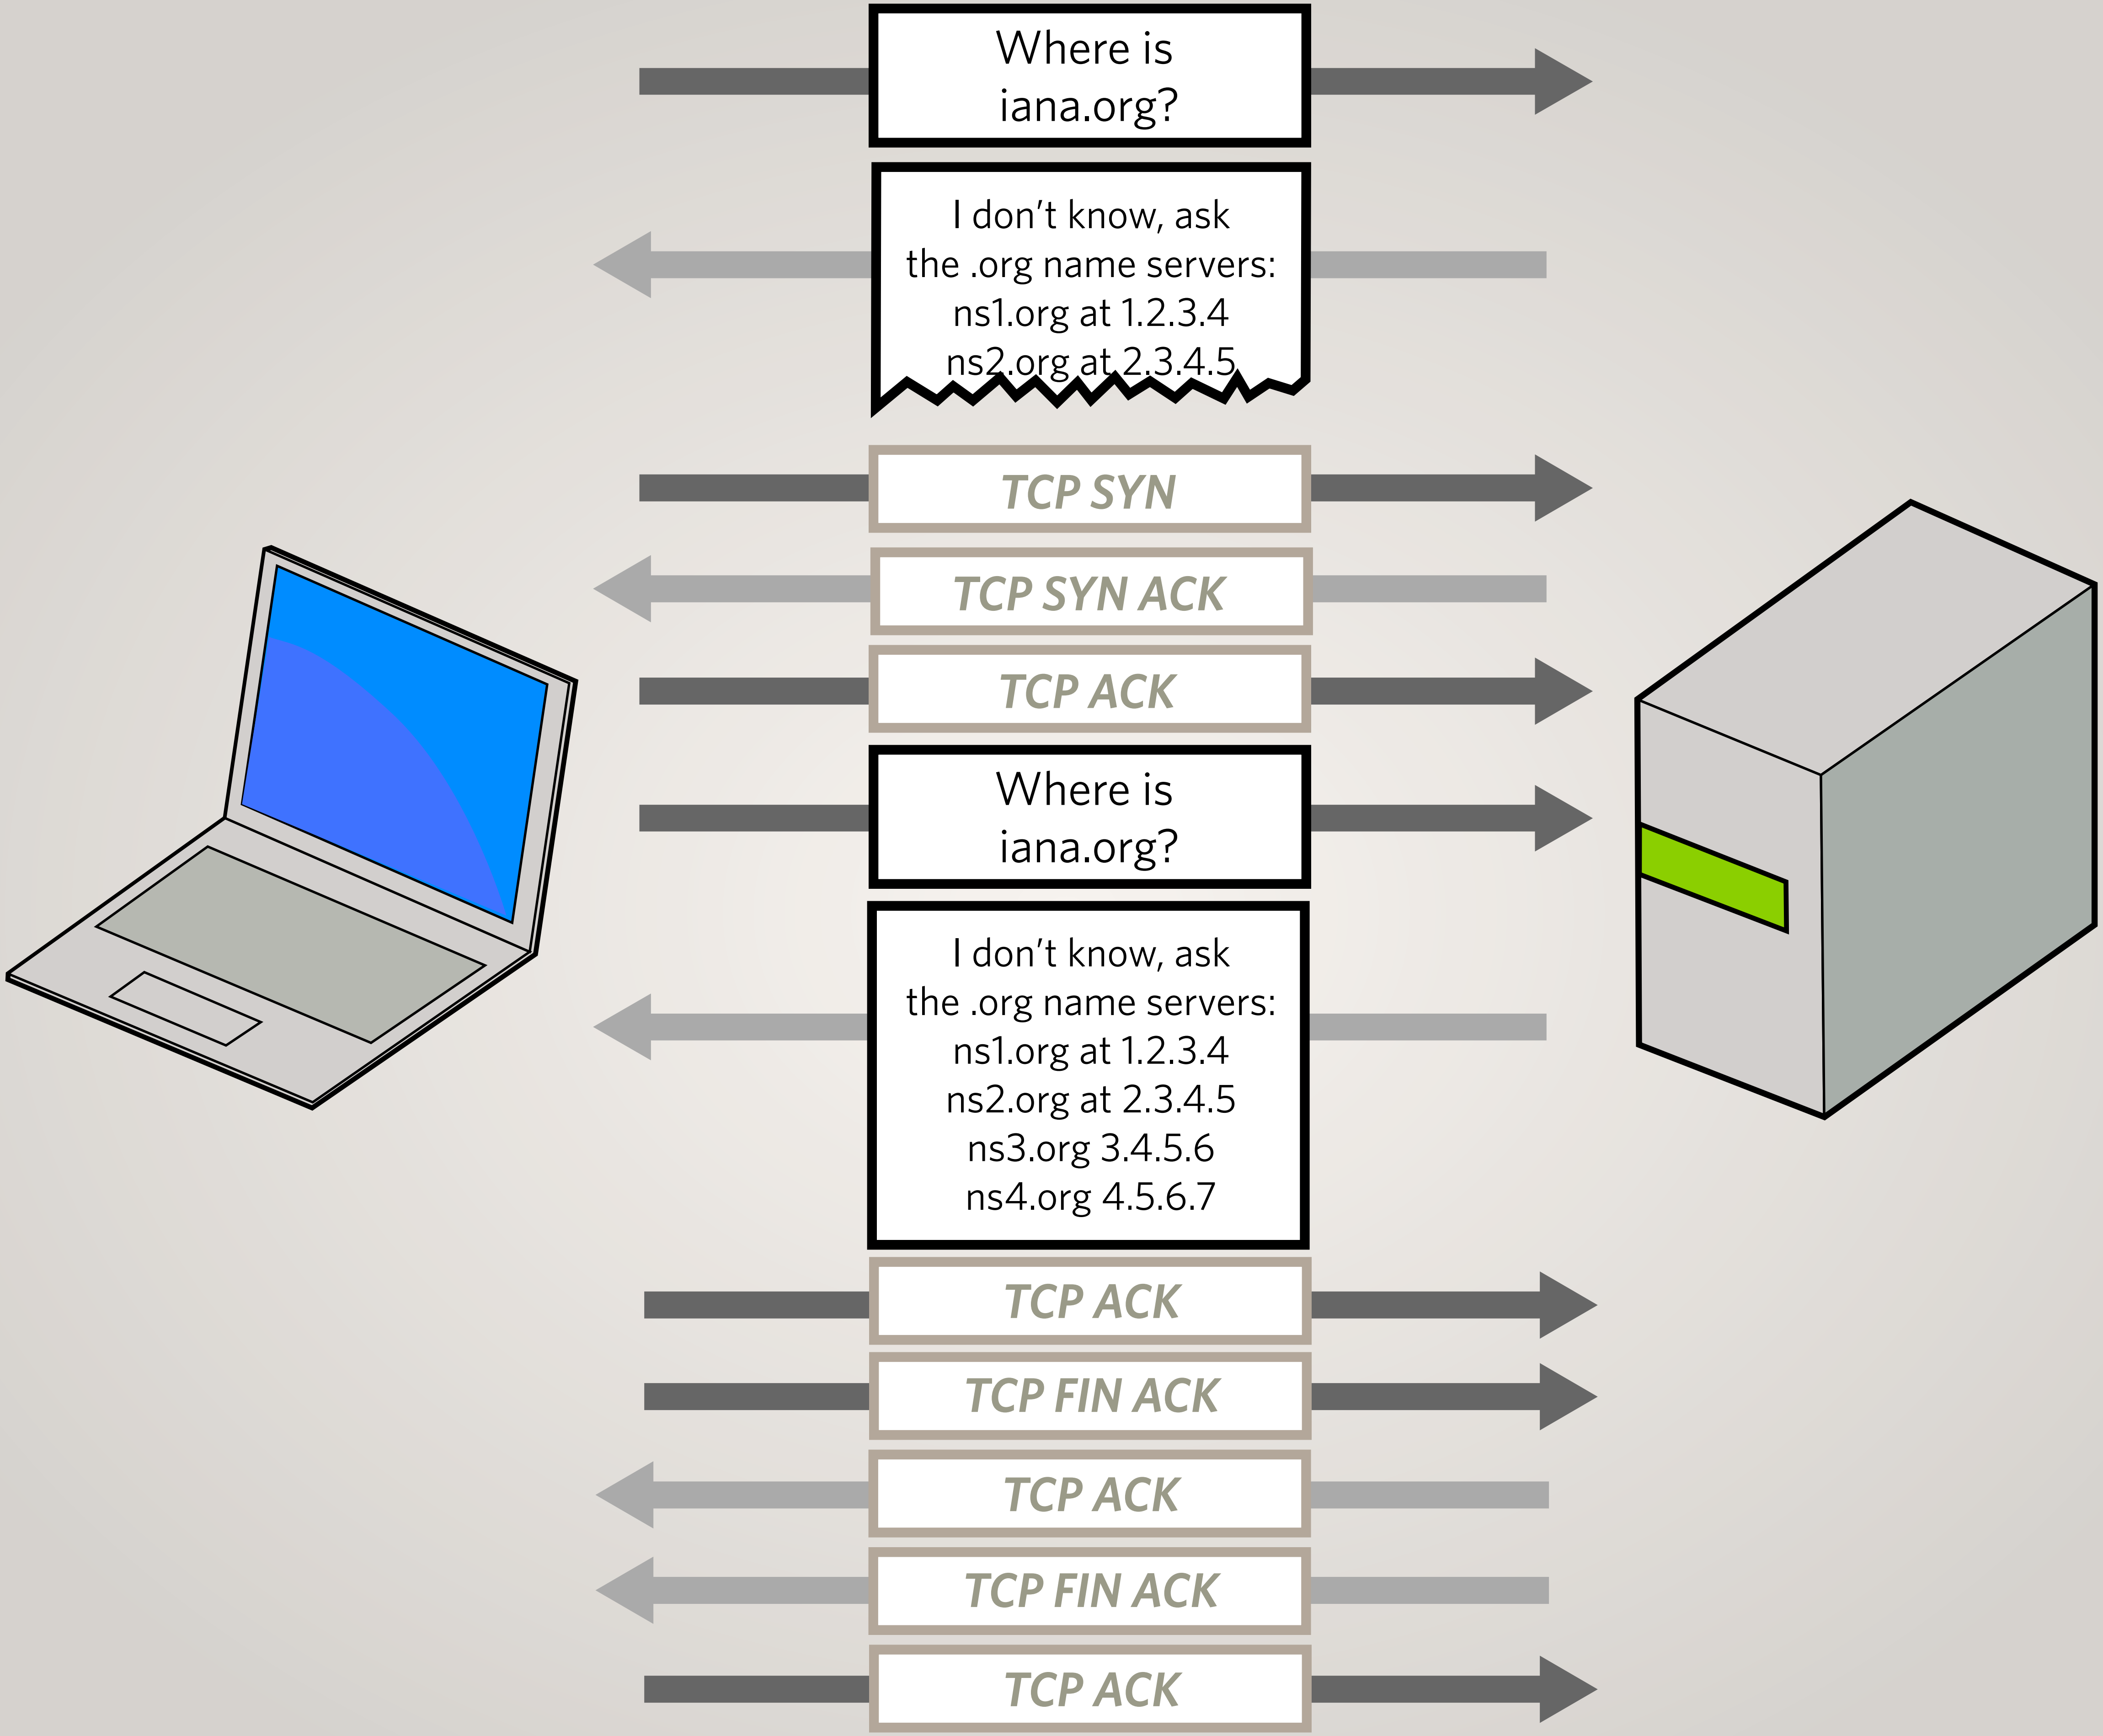

Where is iana.org?

I don’t know, ask the .org name servers:
ns1.org at 1.2.3.4
ns2.org at 2.3.4.5

TCP SYN

TCP SYN ACK

TCP ACK

Where is iana.org?

I don’t know, ask the .org name servers:
ns1.org at 1.2.3.4
ns2.org at 2.3.4.5
ns3.org 3.4.5.6
ns4.org 4.5.6.7

TCP ACK

TCP FIN ACK

TCP ACK

TCP FIN ACK

TCP ACK

Limiting referral size

- Reduce the number of name servers
- Take advantage of name compression

ns1.iana.org and ns2.iana.com

3	n	s	1	4	i	a	n	a	3	o	r	g	0
3	n	s	2	4	i	a	n	a	3	c	o	m	0

Bytes used for names = 28

ns1.iana.org and ns2.iana.com

3	n	s	1	4	i	a	n	a	3	o	r	g	0
3	n	s	2	4	i	a	n	a	3	c	o	m	0

Bytes used for names = 28

ns1.iana.org and ns2.iana.org

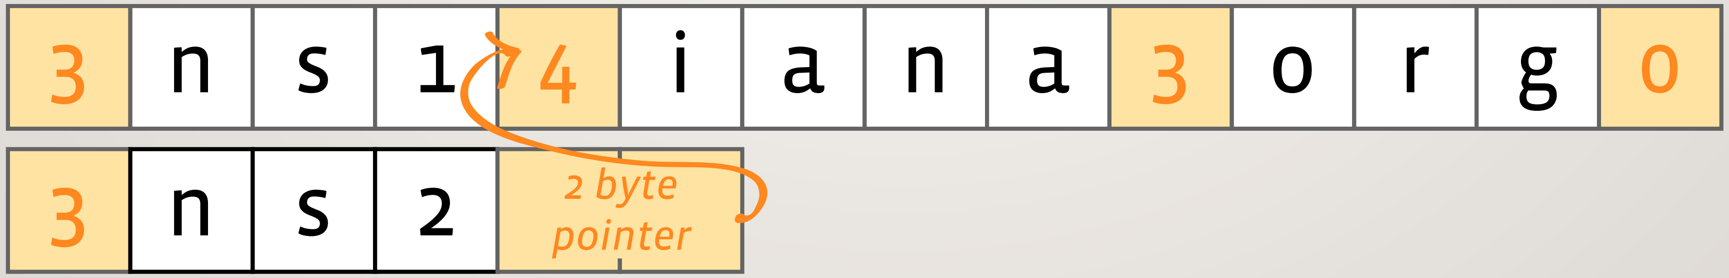

3	n	s	1	4	i	a	n	a	3	o	r	g	0
3	n	s	2	2 byte pointer									

Bytes used for names = 20

8 bytes saved

Limiting referral size

- Reduce the number of name servers
- Take advantage of name compression
- The more domains are shared for authorities, the better the compression outcome
- Tradeoff — you are now more reliant on certain domains

The bottom line

The bottom line

TLDs with open recursive name servers	9.6%

The bottom line

TLDs with open recursive name servers	9.6%
TLDs without diverse IPv4 connectivity	7.2%

The bottom line

TLDs with open recursive name servers	9.6%
TLDs without diverse IPv4 connectivity	7.2%
TLDs without diverse IPv6 connectivity	68.7%

The bottom line

TLDs with open recursive name servers	9.6%
TLDs without diverse IPv4 connectivity	7.2%
TLDs without diverse IPv6 connectivity	68.7%
… without *any* IPv6	41.0%

The bottom line

TLDs with open recursive name servers	9.6%
TLDs without diverse IPv4 connectivity	7.2%
TLDs without diverse IPv6 connectivity	68.7%
… without *any* IPv6	41.0%
TLDs with referrals that can fragment	4.3%

How IDN ccTLD applications will be processed

(in theory)

Signing the Root Zone

Signing the root zone?

- ICANN’s strategic plan is to be “operationally ready”
    - Signed root test bed operating for over a year
    - System is built with advice from current DNSSEC operators, and many other experts in both DNS and cryptography
    - ICANN already signs 11 top-level domains operationally, and incrementally signing the last remaining zones under our control

Signing the root zone?

- ICANN developed a proposal to sign the root zone which was submitted to US Government
- VeriSign followed up with a different proposal to sign the root zone
- The US Government has issued a "Notice of Inquiry" to seek views relating to signing the DNS root zone, which was open to comments until November 24.
  - http://tinyurl.com/3v8akt

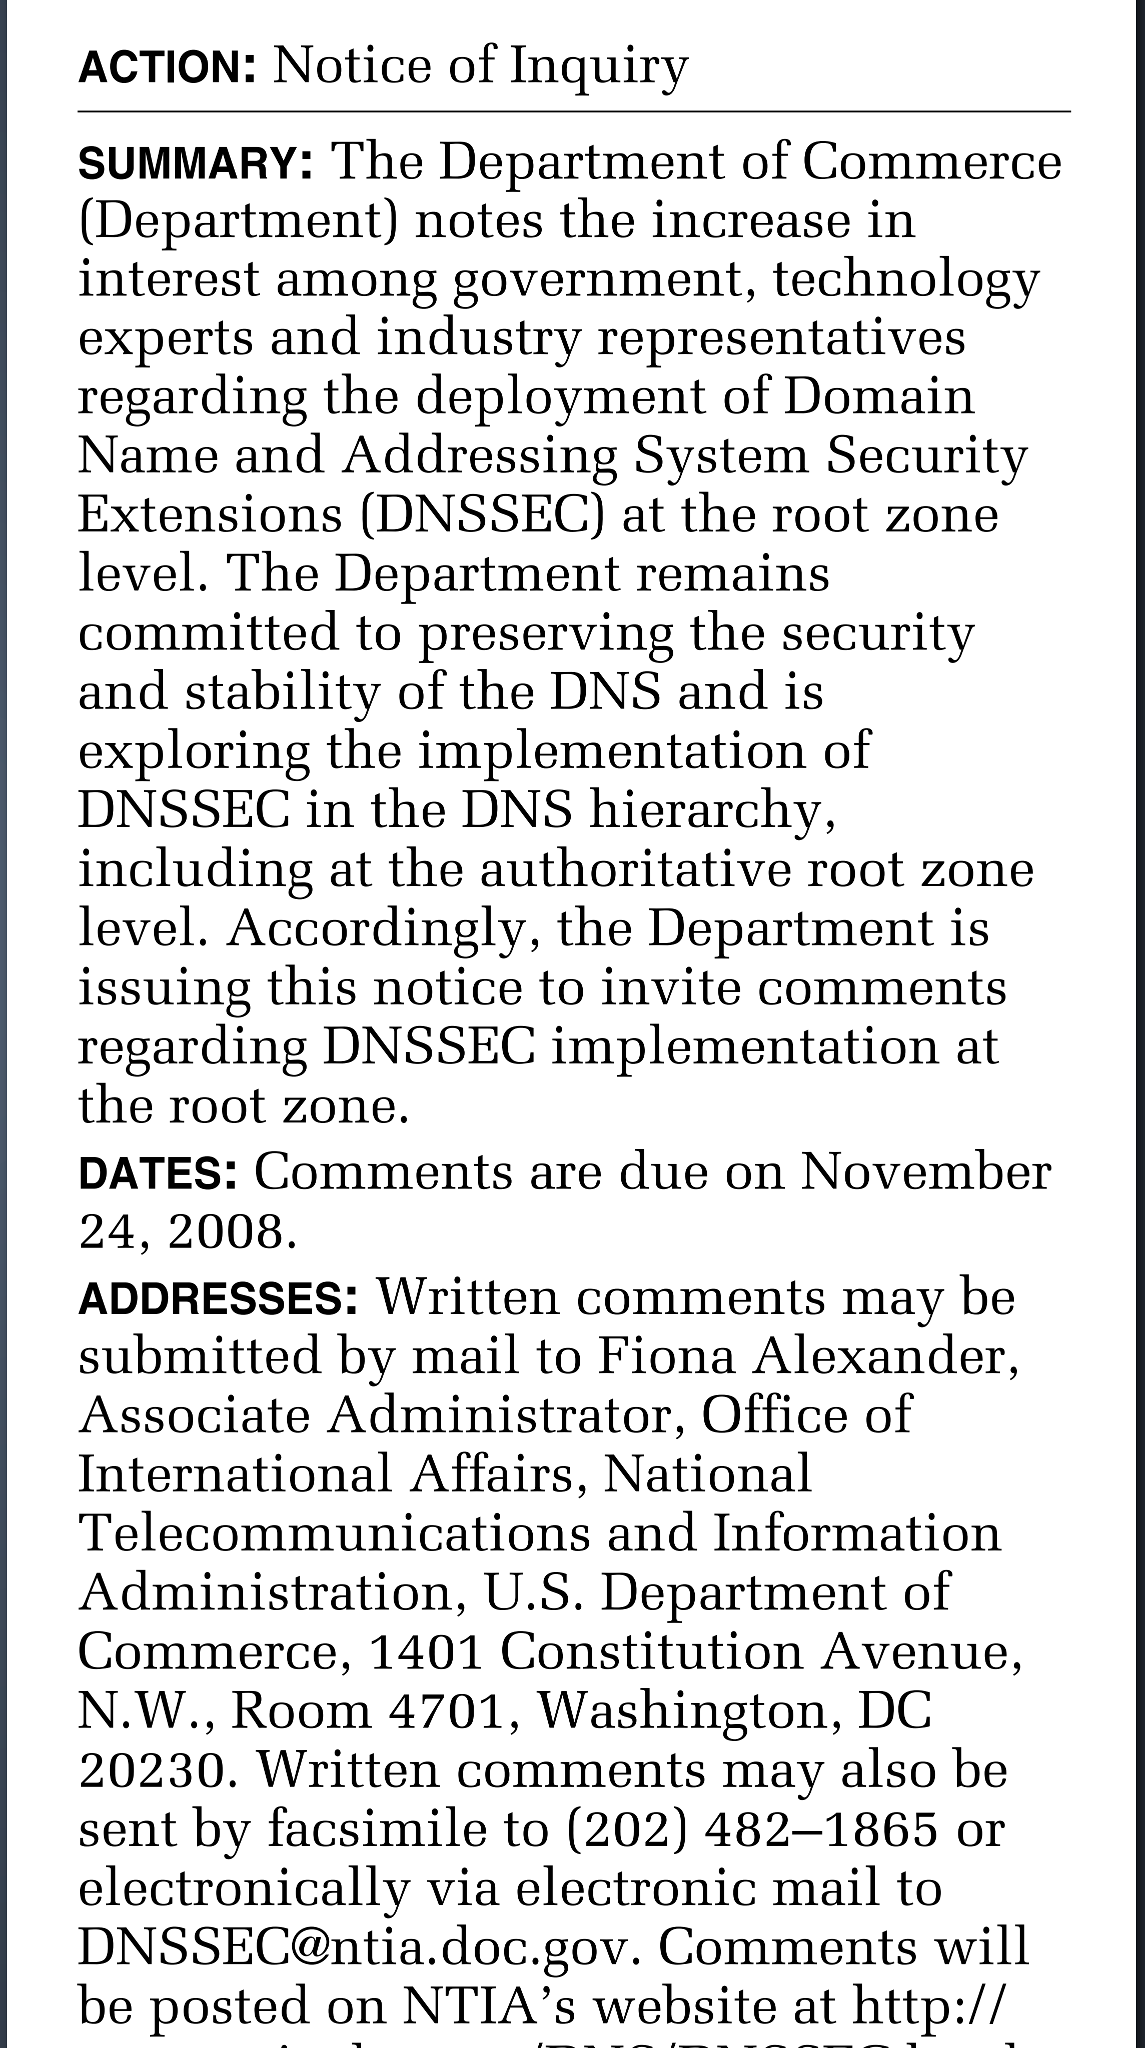

ACTION: Notice of Inquiry

SUMMARY: The Department of Commerce (Department) notes the increase in interest among government, technology experts and industry representatives regarding the deployment of Domain Name and Addressing System Security Extensions (DNSSEC) at the root zone level. The Department remains committed to preserving the security and stability of the DNS and is exploring the implementation of DNSSEC in the DNS hierarchy, including at the authoritative root zone level. Accordingly, the Department is issuing this notice to invite comments regarding DNSSEC implementation at the root zone.

DATES: Comments are due on November 24, 2008.

ADDRESSES: Written comments may be submitted by mail to Fiona Alexander, Associate Administrator, Office of International Affairs, National Telecommunications and Information Administration, U.S. Department of Commerce, 1401 Constitution Avenue, N.W., Room 4701, Washington, DC 20230. Written comments may also be sent by facsimile to (202) 482–1865 or electronically via electronic mail to DNSSEC@ntia.doc.gov. Comments will be posted on NTIA's website at http://

Interim Trust Anchor Repository

Interim Trust Anchor Repository

- A mechanism to publish keys of top-level domains that currently implement DNSSEC
- If the root zone is DNSSEC signed, such a repository is unnecessary
  - Therefore this is a stopgap measure
  - Should be decommissioned when the root is signed

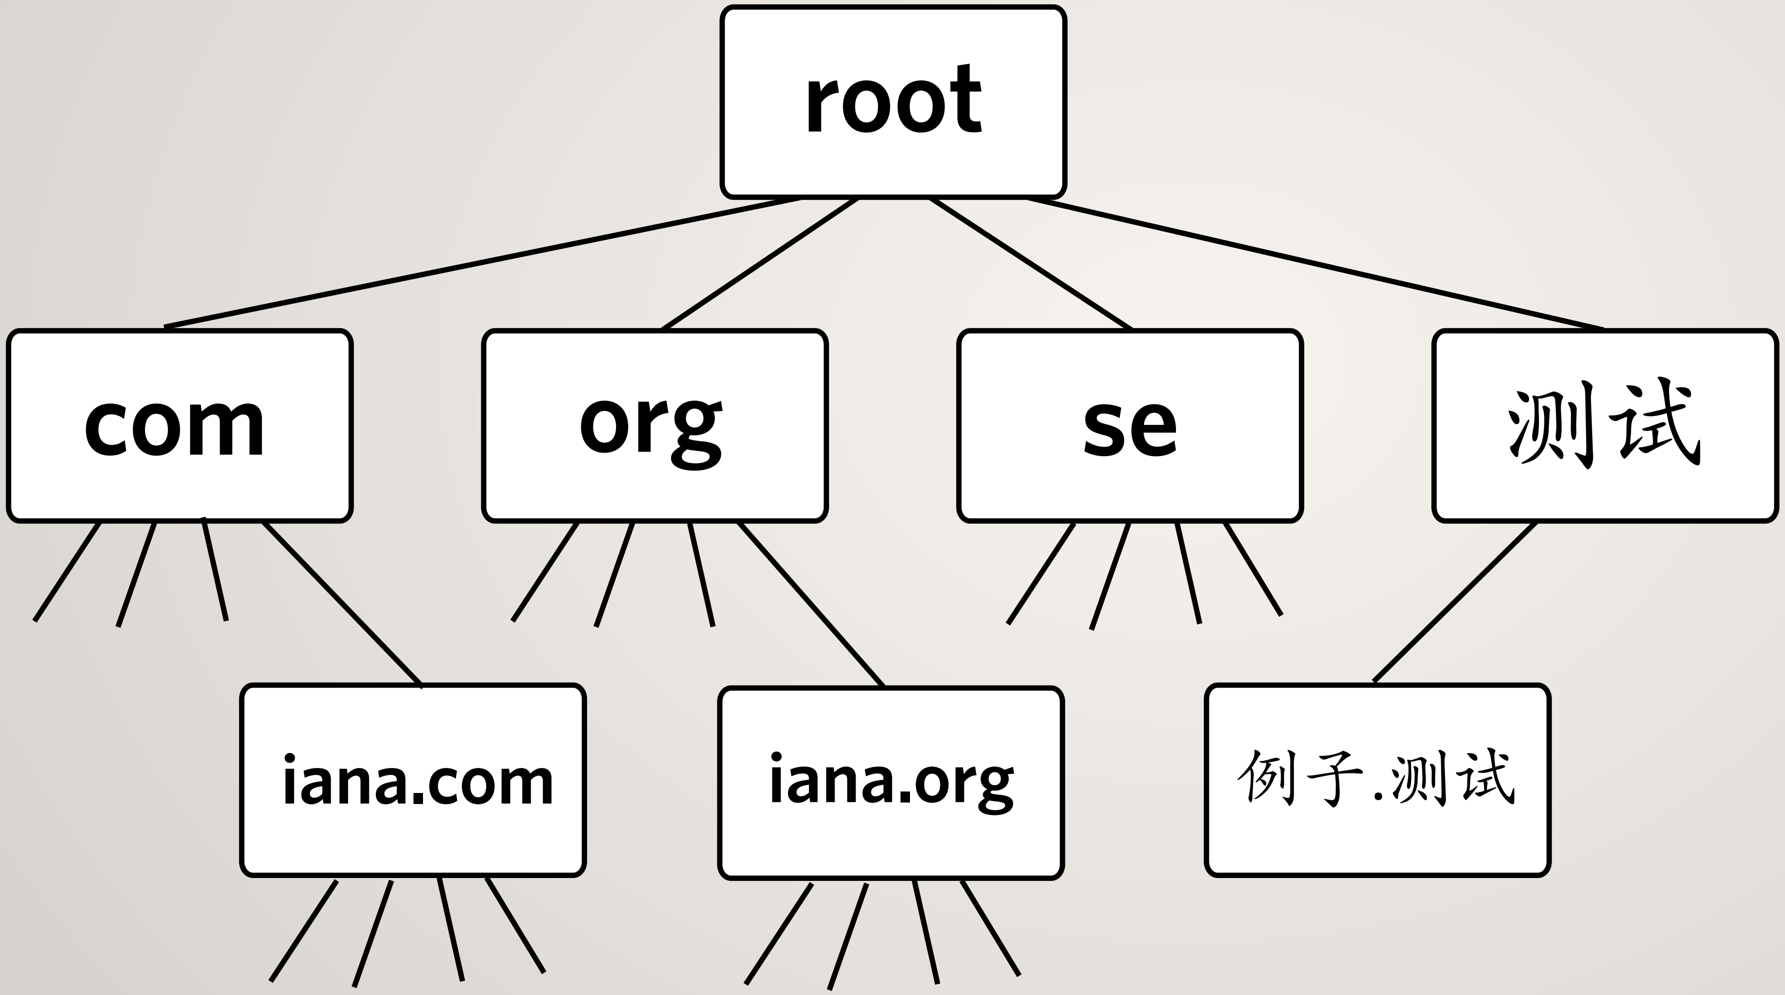
root
com
org
se
测试
iana.com
iana.org
例子.测试

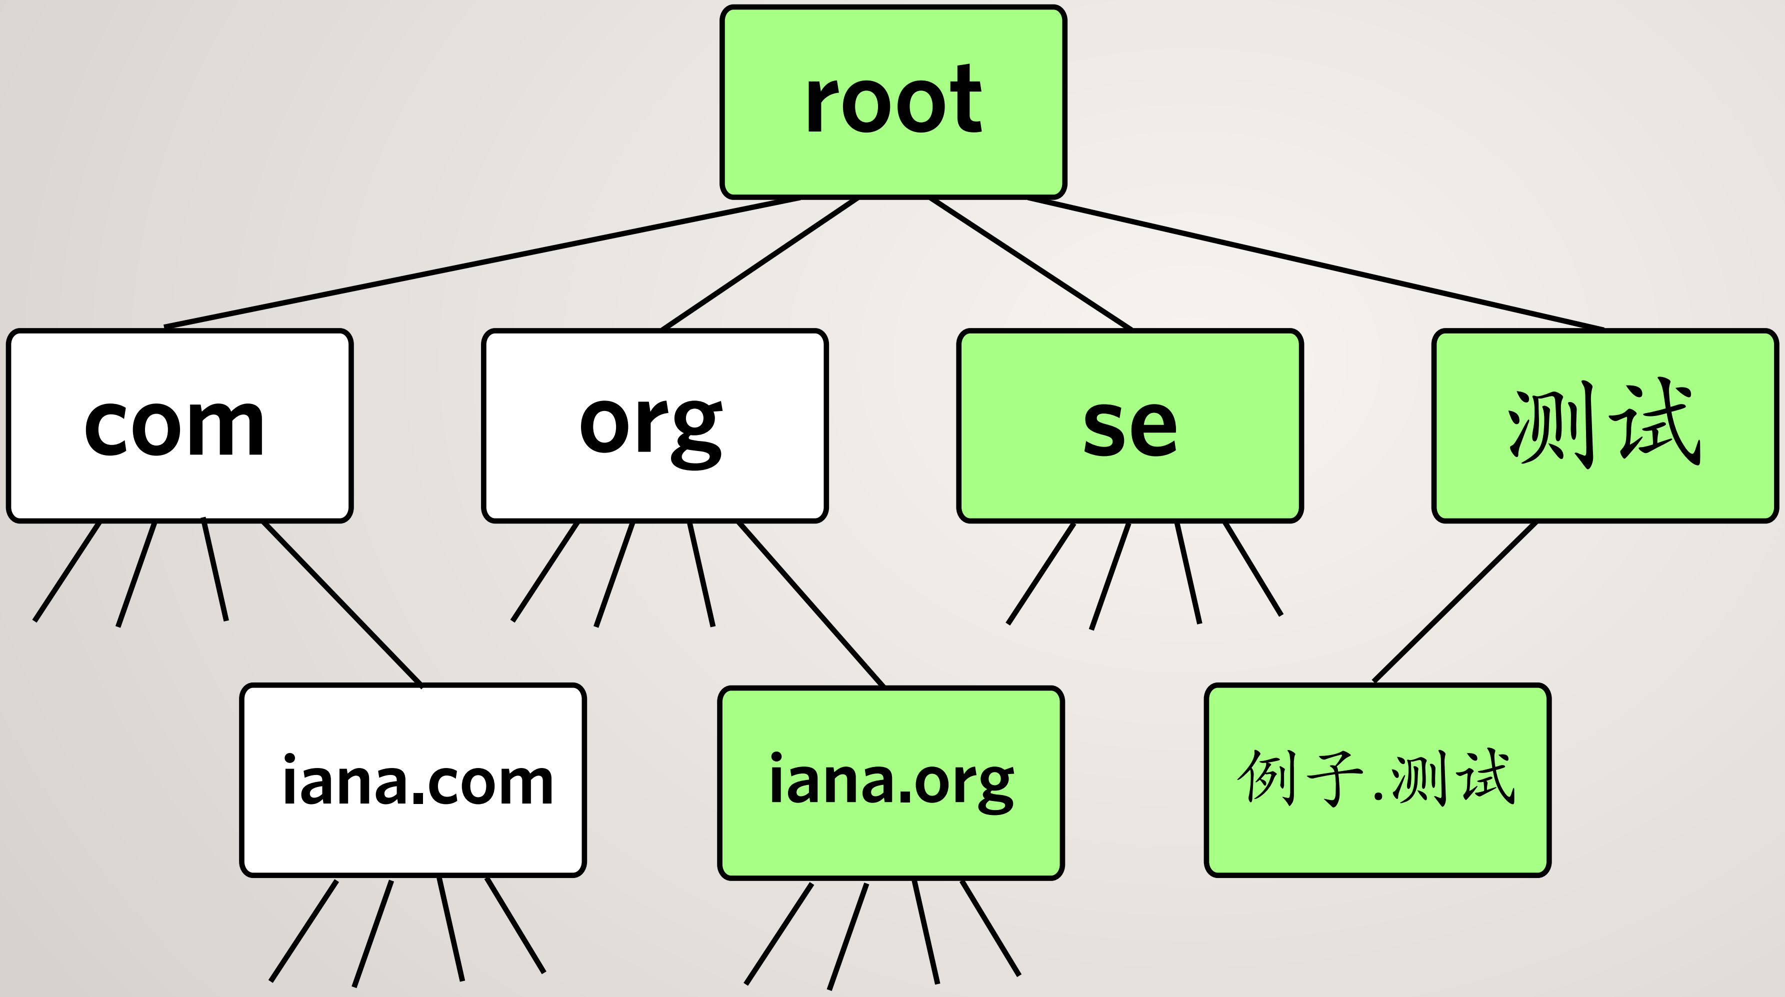
root
com
org
se
测试
iana.com
iana.org
例子.测试

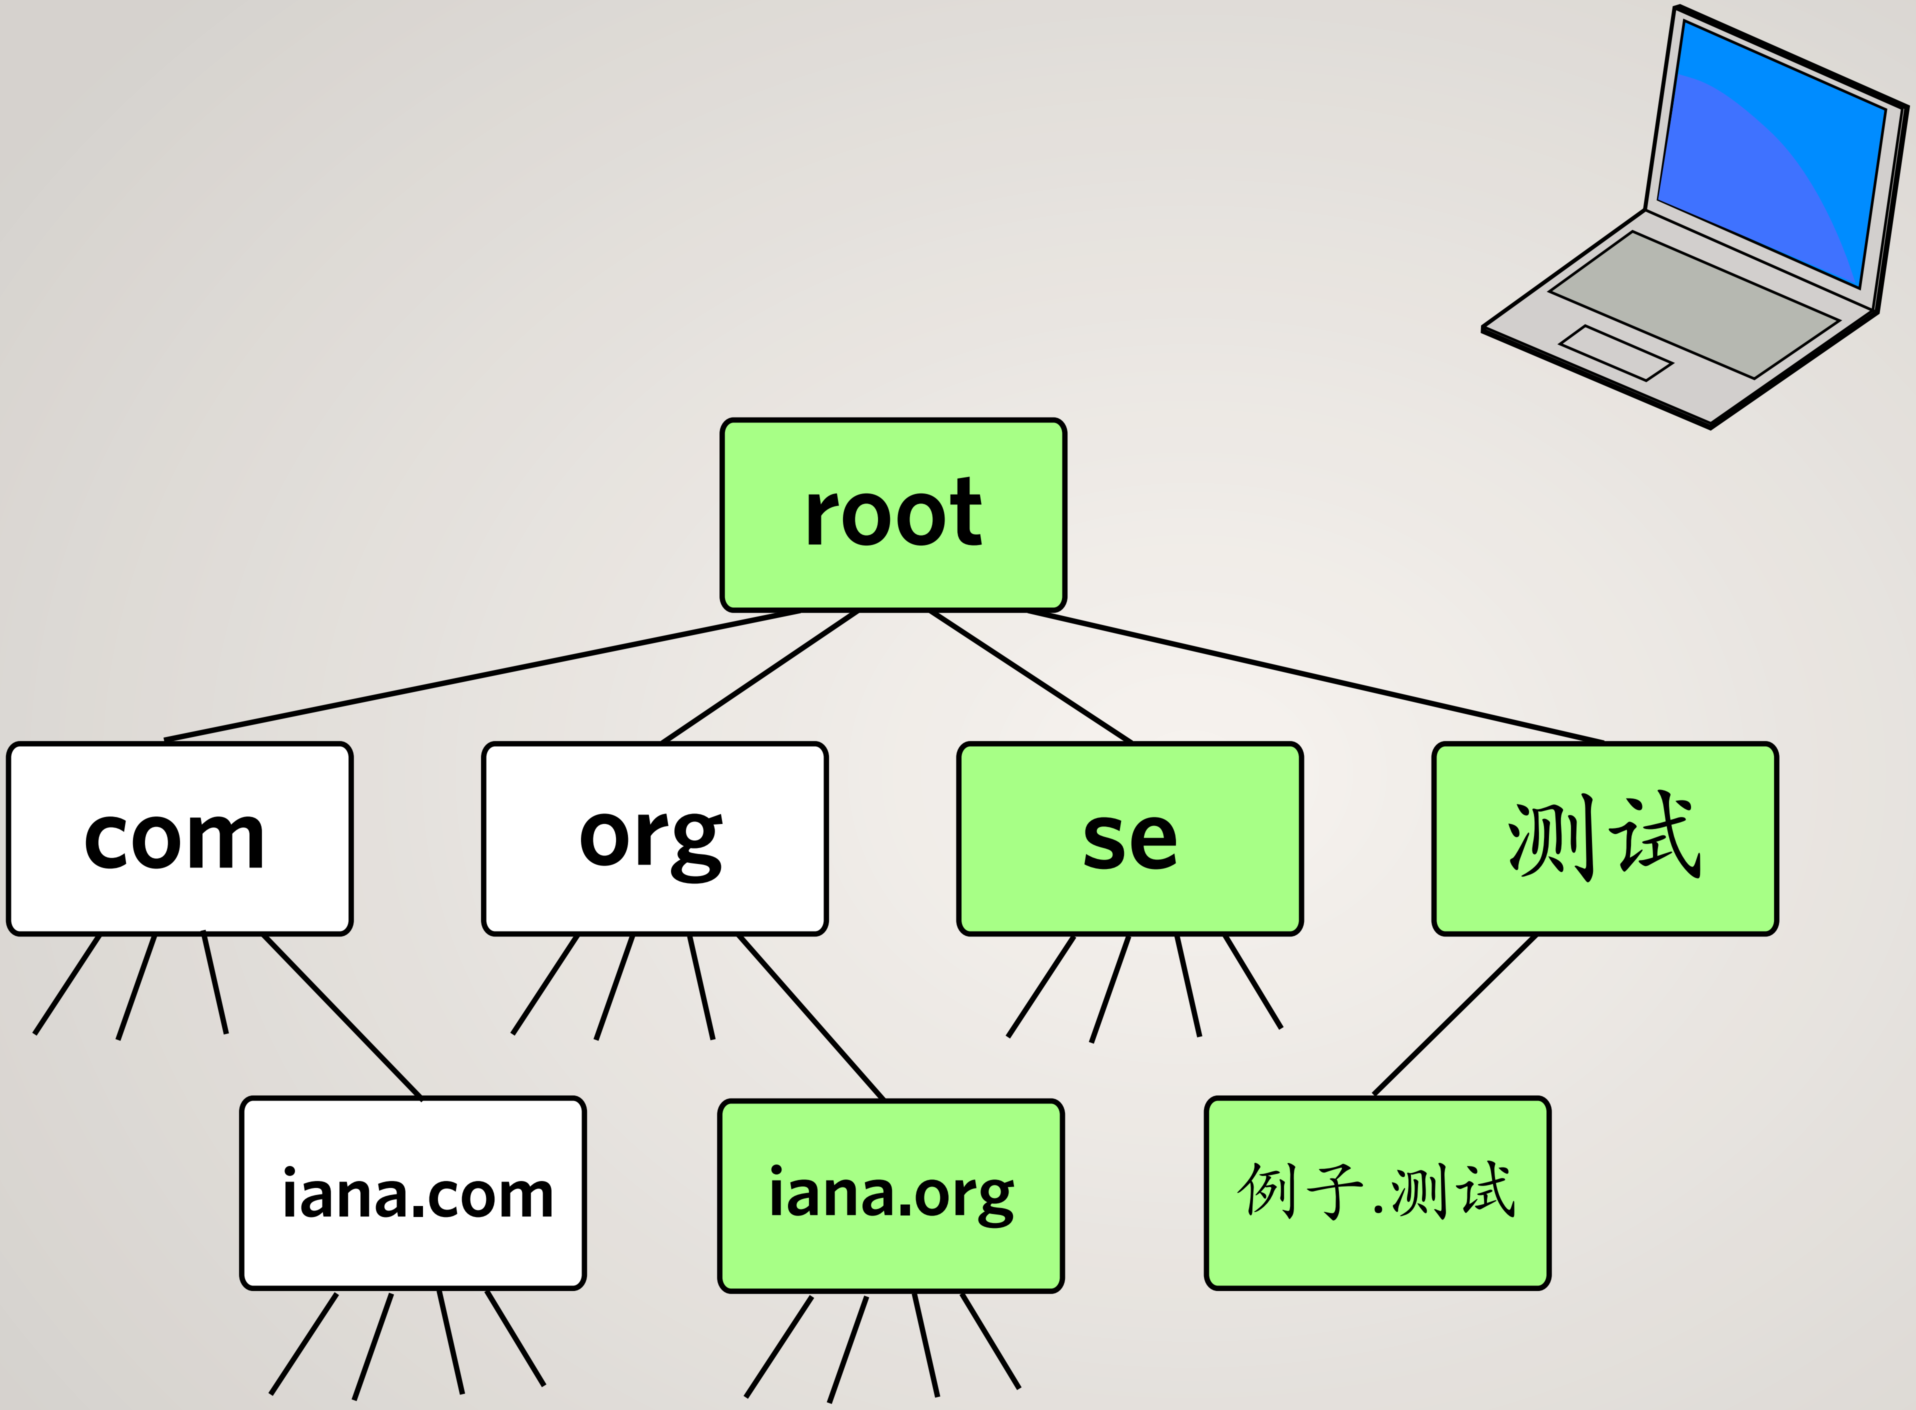
root
com
org
se
测试
iana.com
iana.org
例子.测试

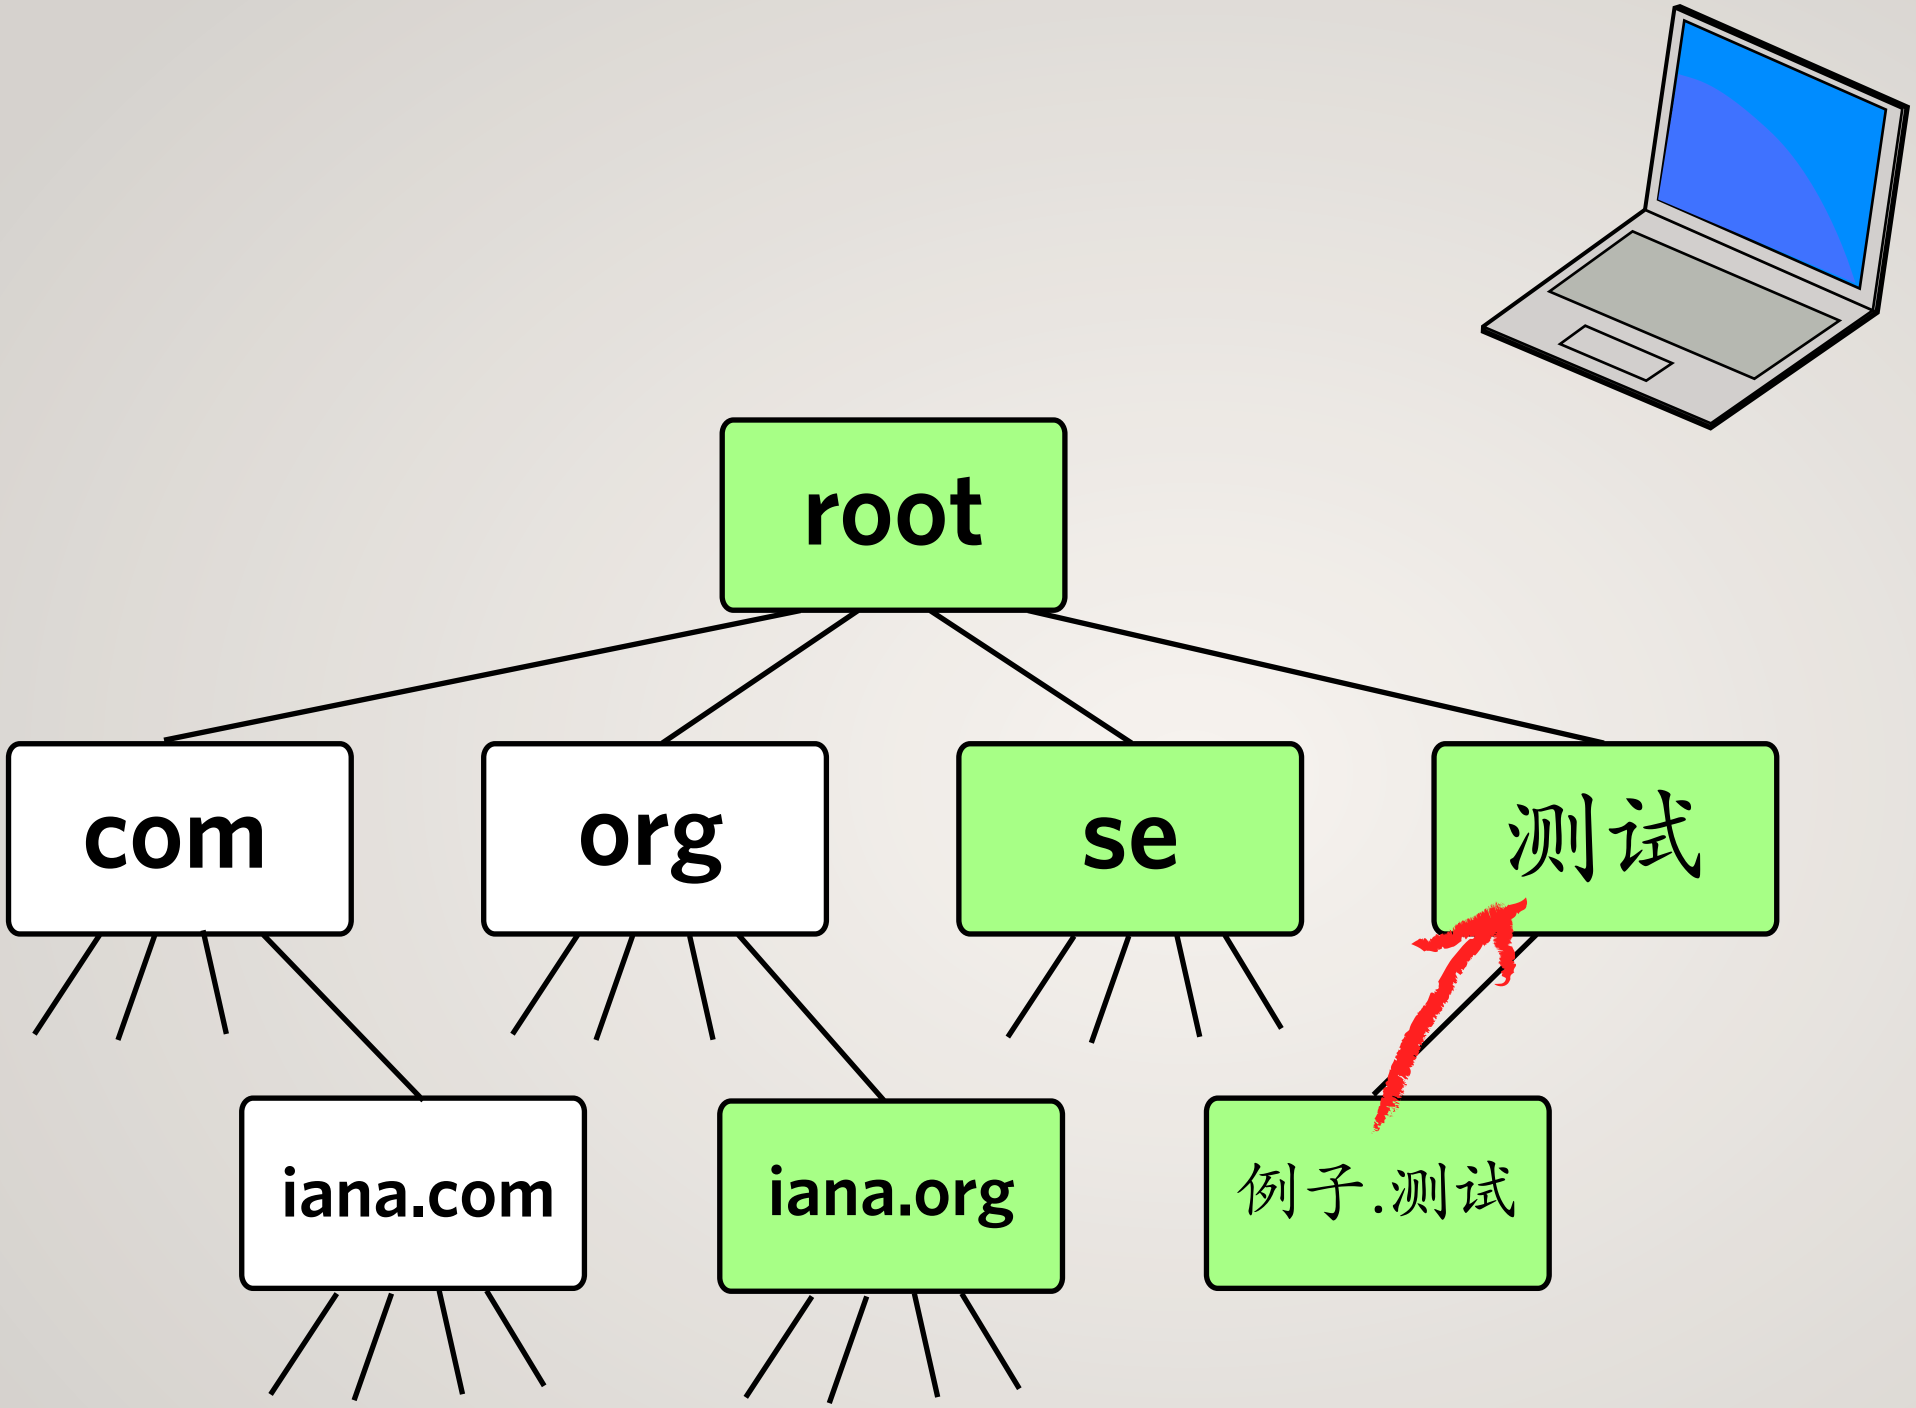
root
com
org
se
测试
iana.com
iana.org
例子.测试

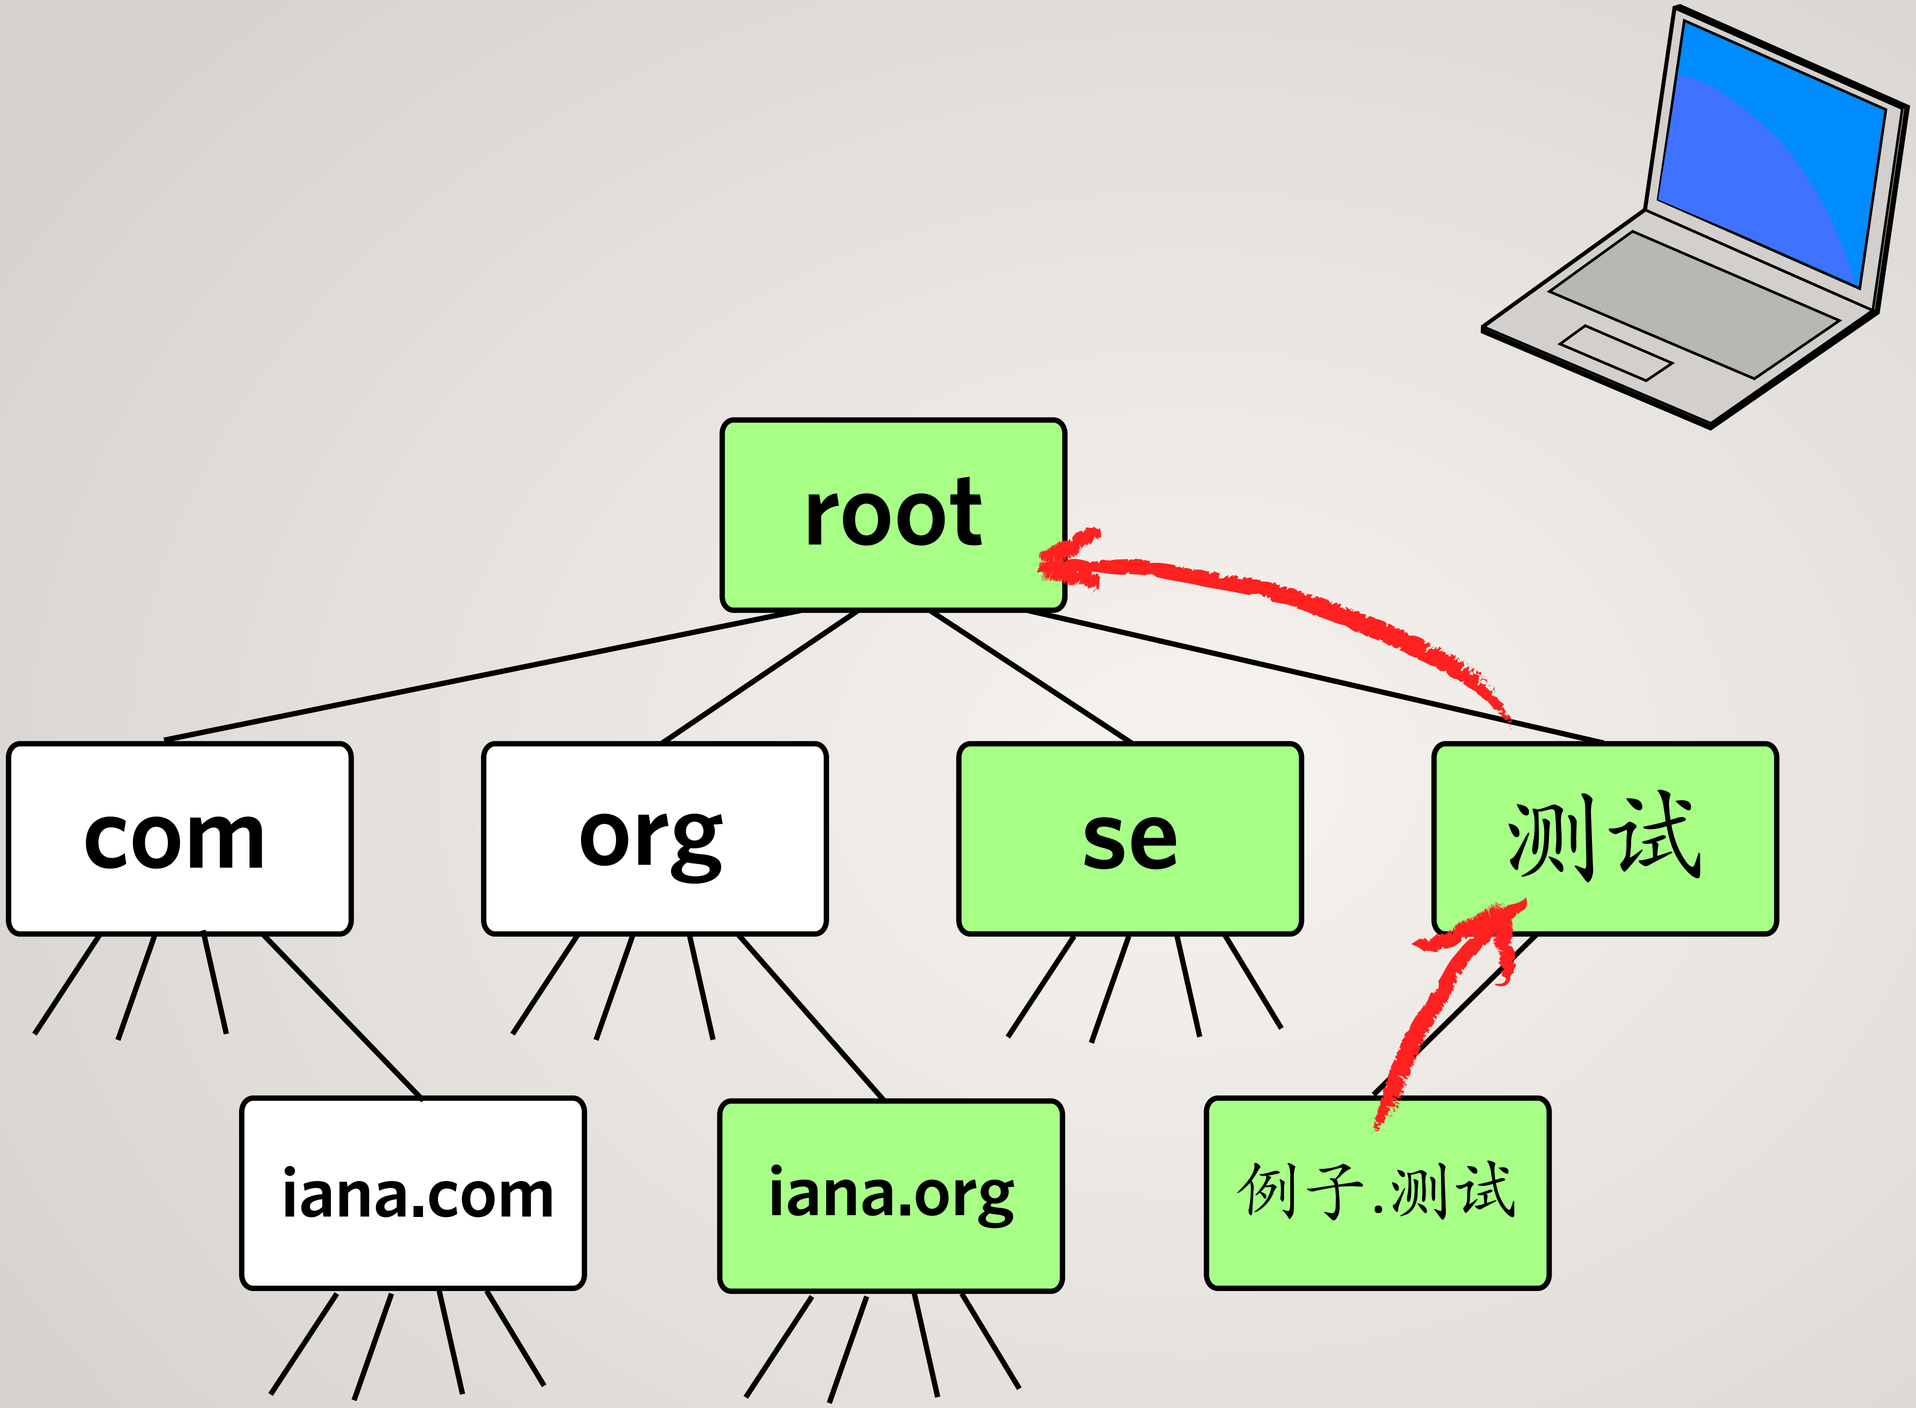

root
com
org
se
测试
iana.com
iana.org
例子.测试

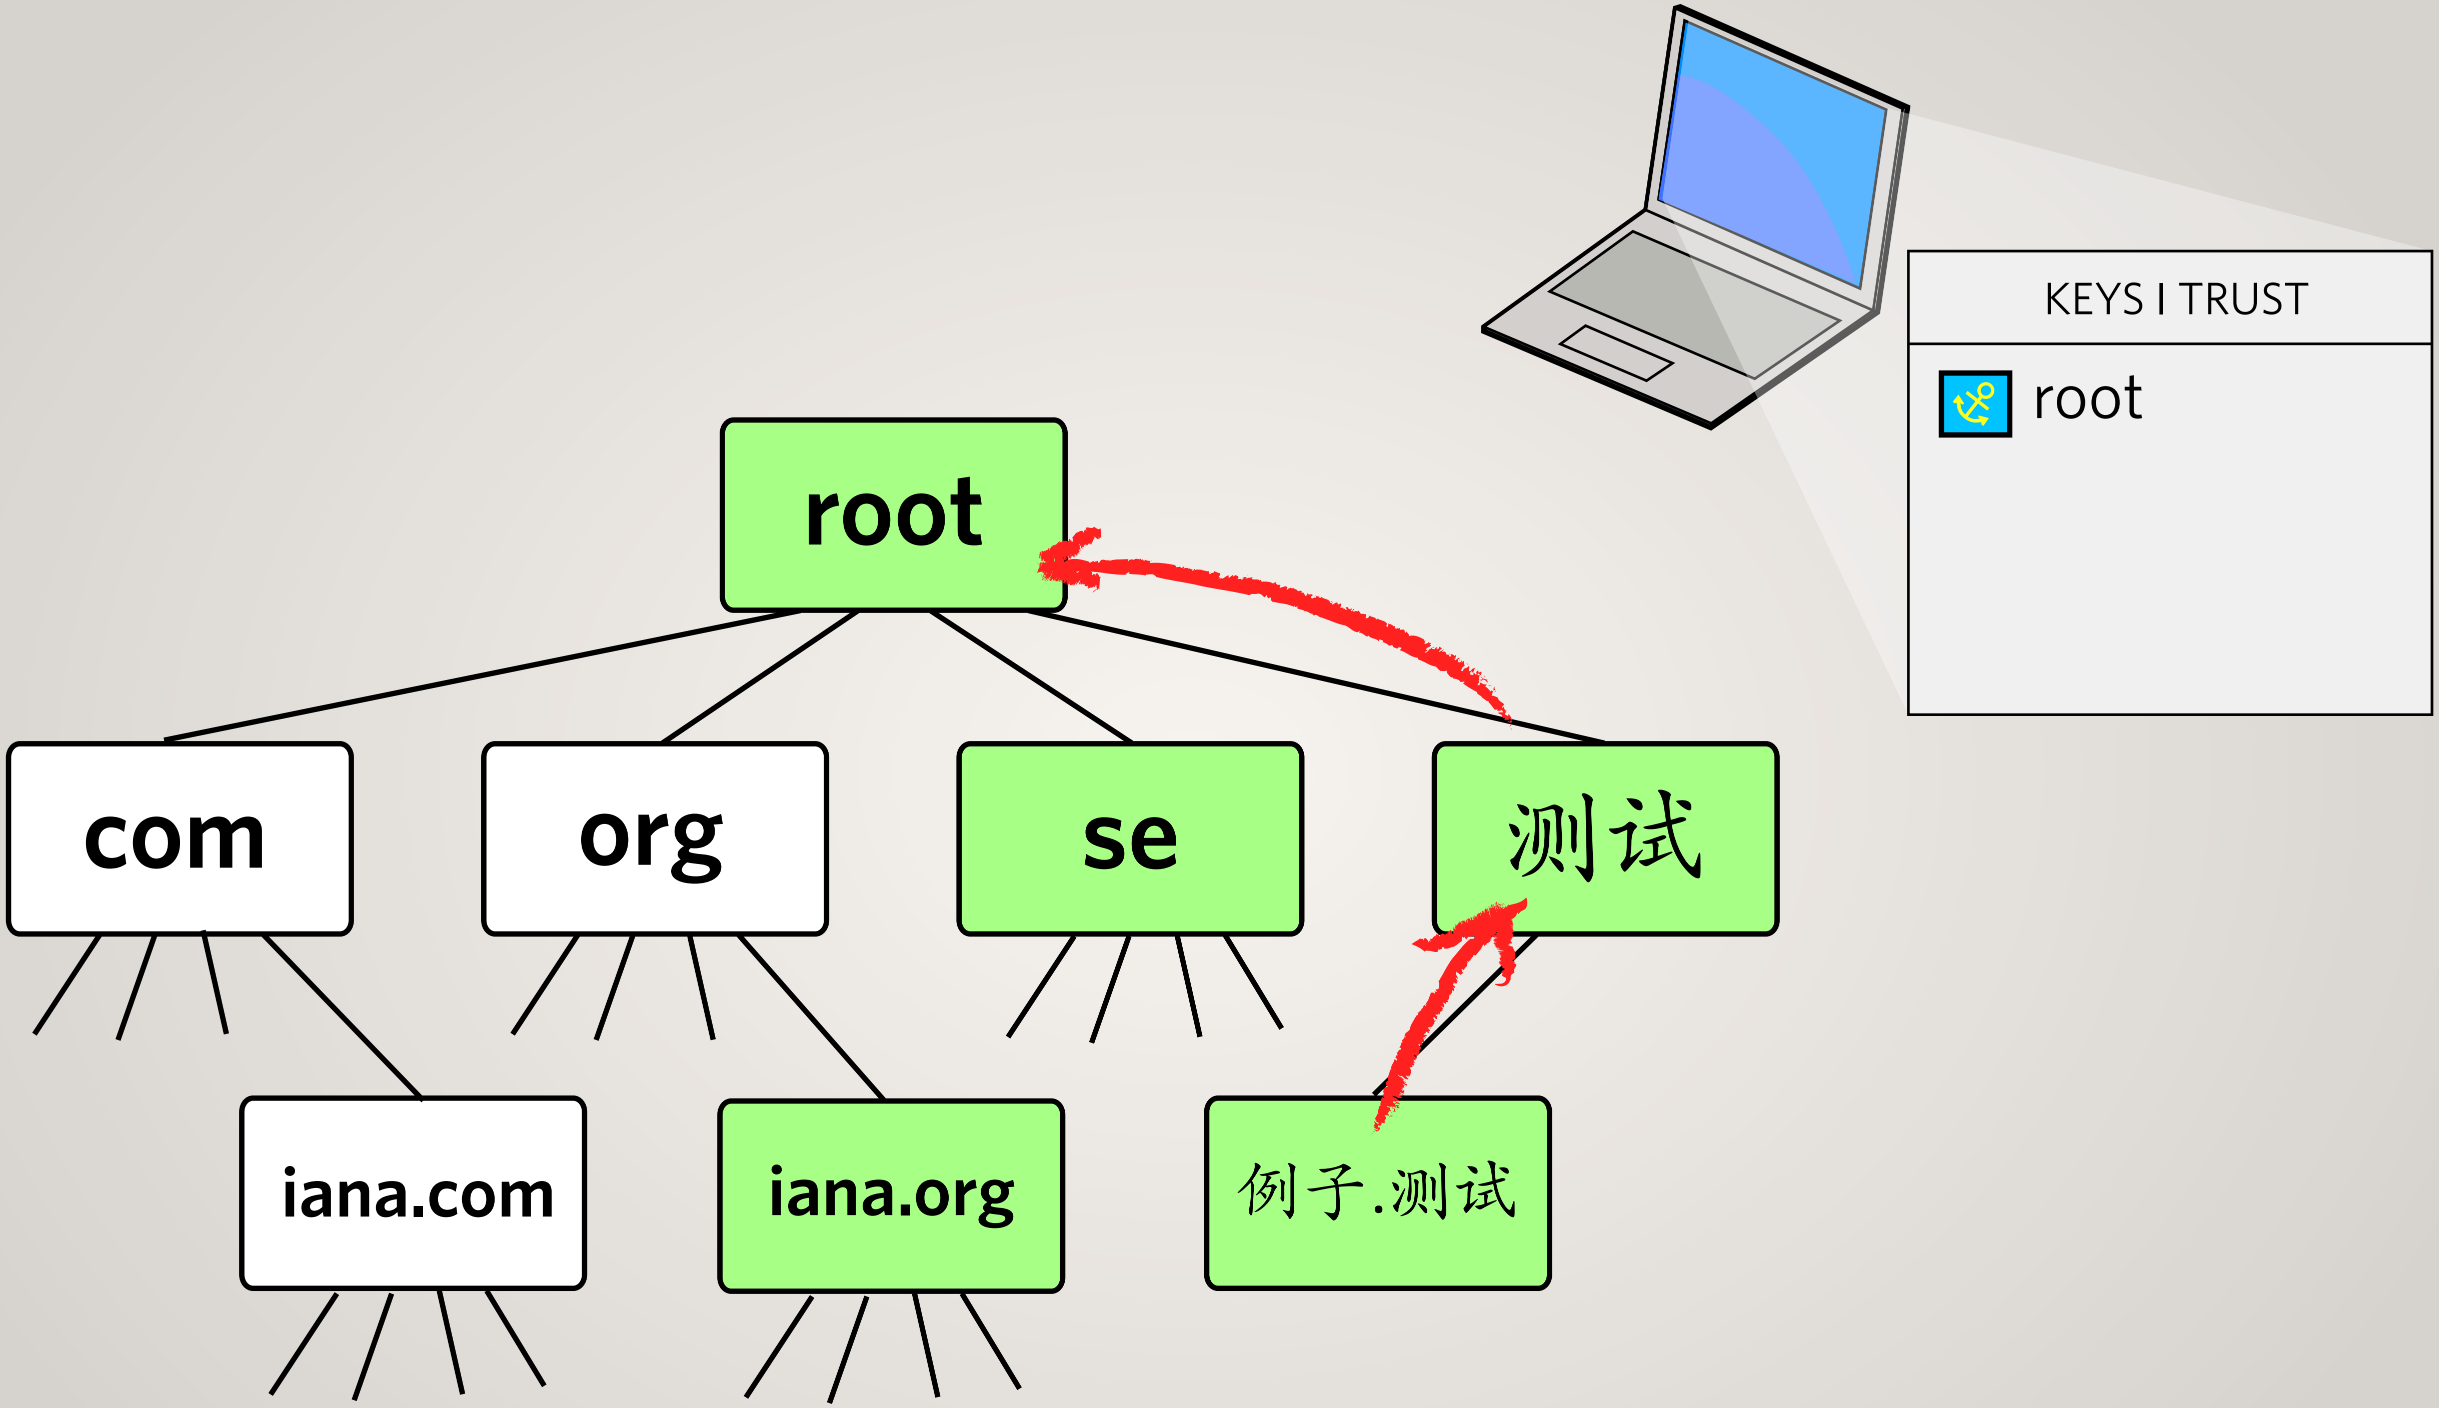
KEYS I TRUST
root
root
com
org
se
测试
iana.com
iana.org
例子.测试

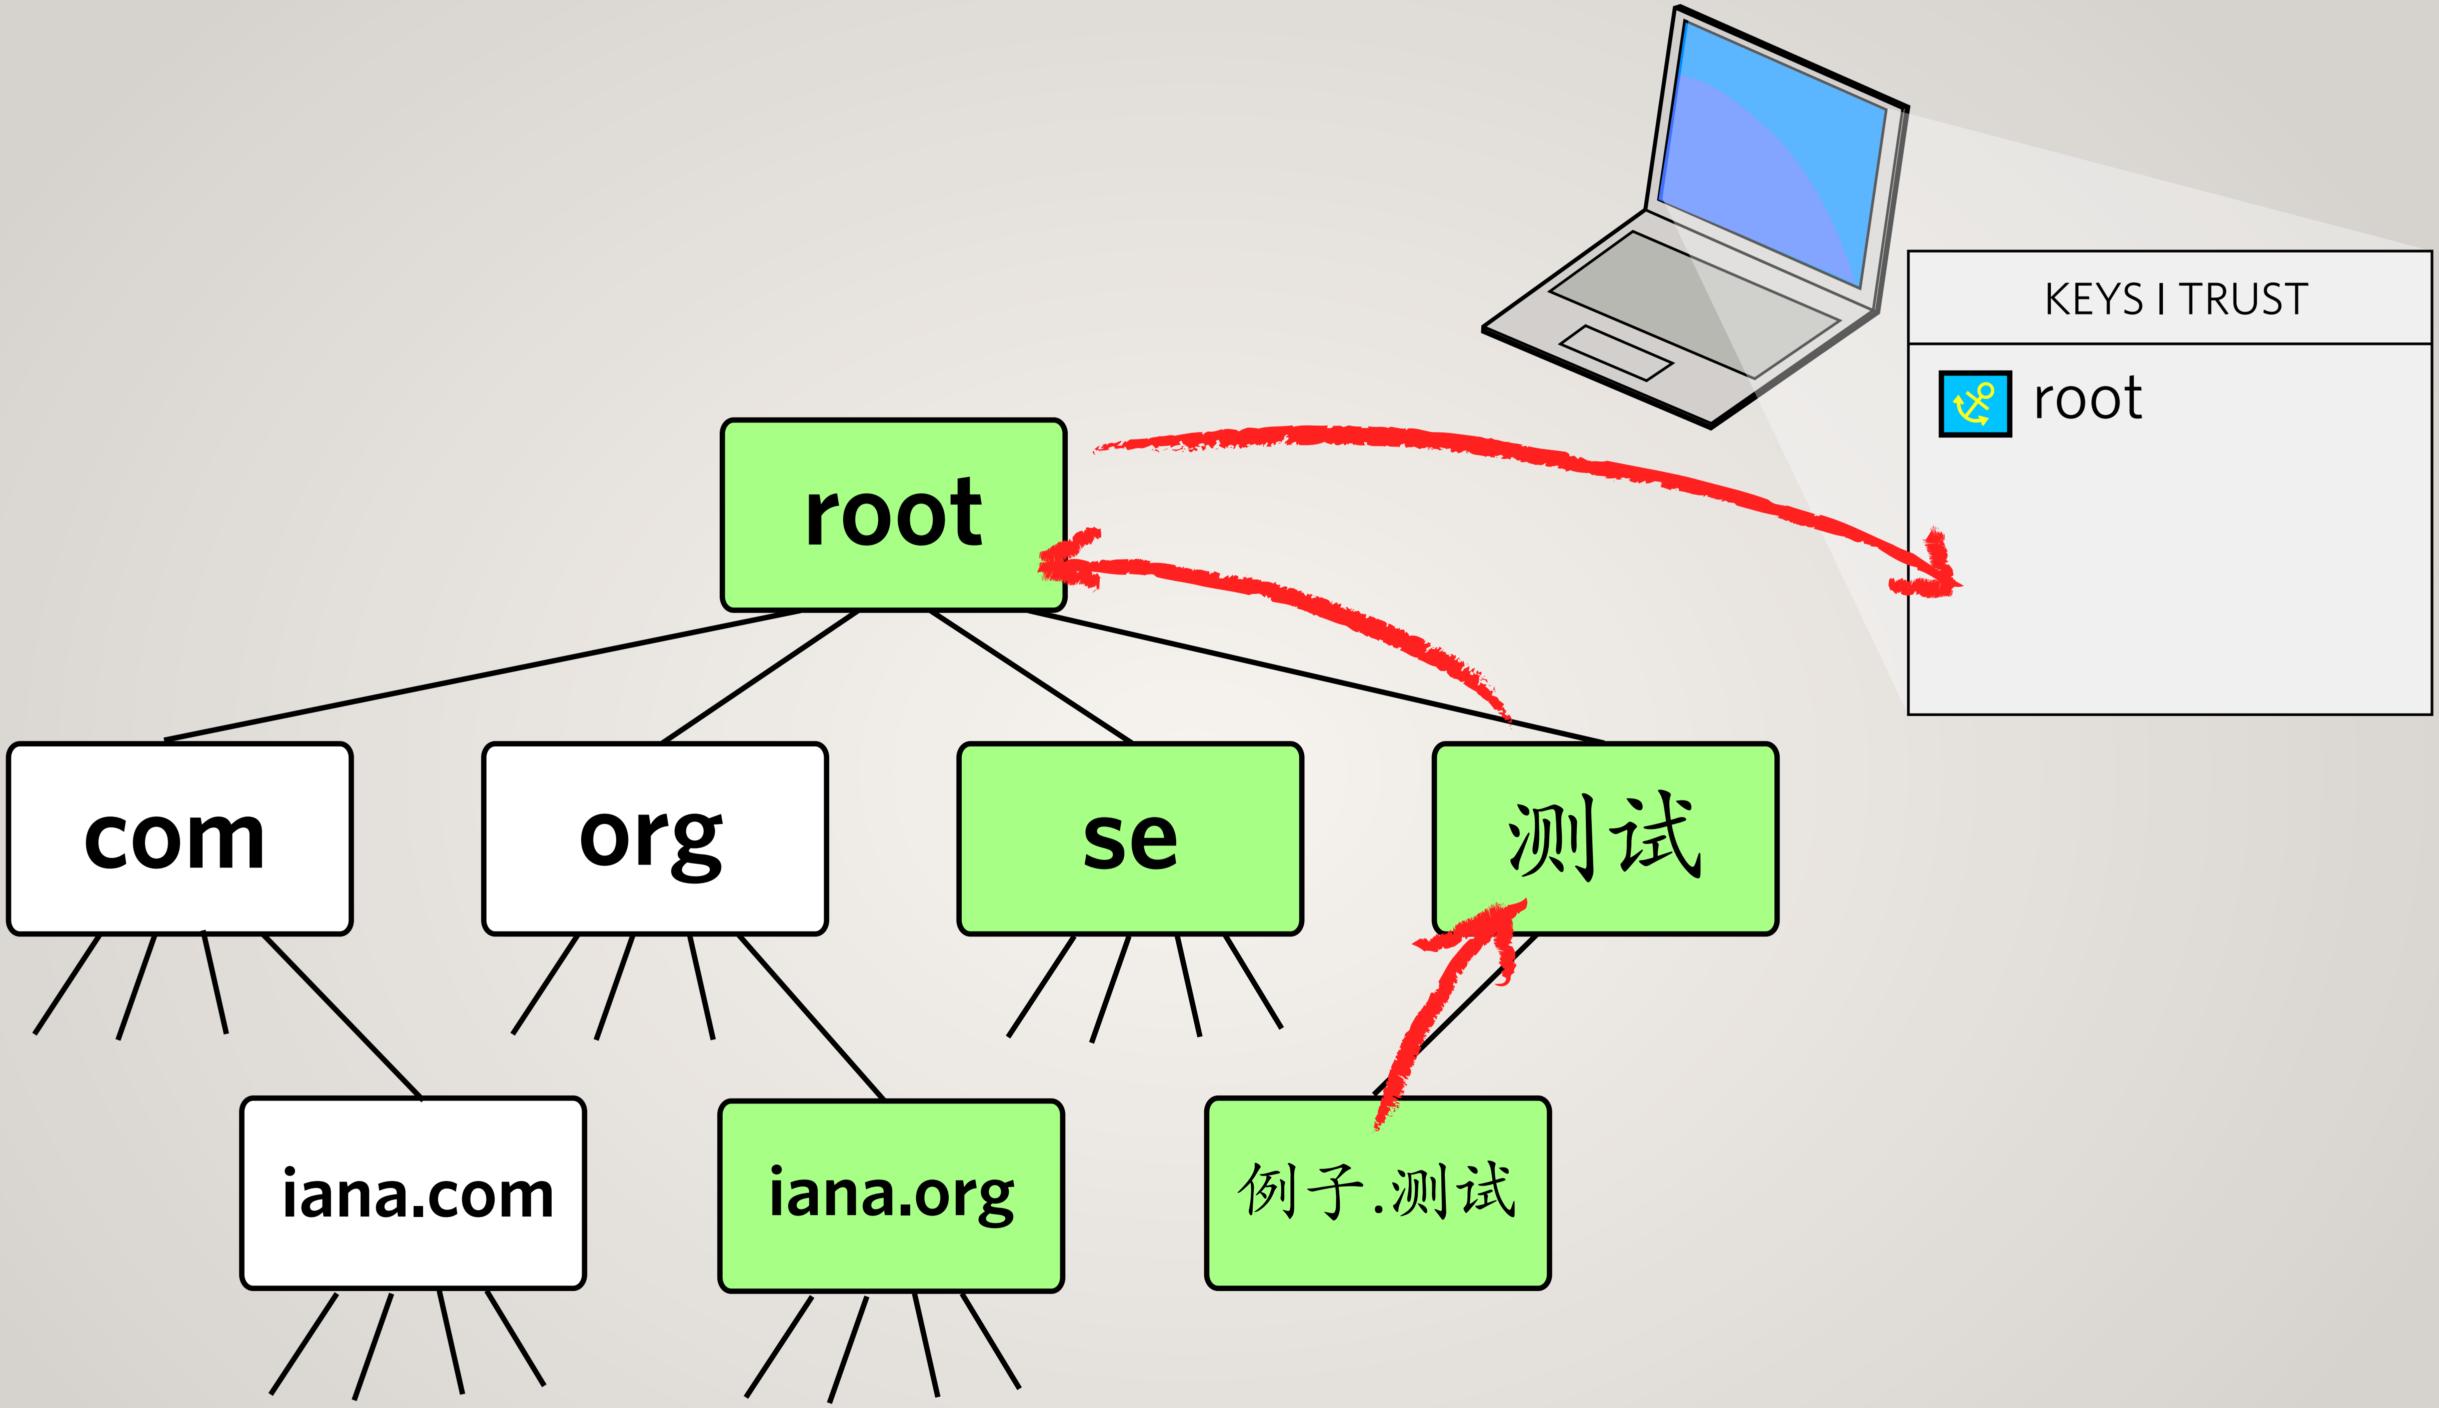
root
com
org
se
测试
iana.com
iana.org
例子.测试
KEYS I TRUST
root

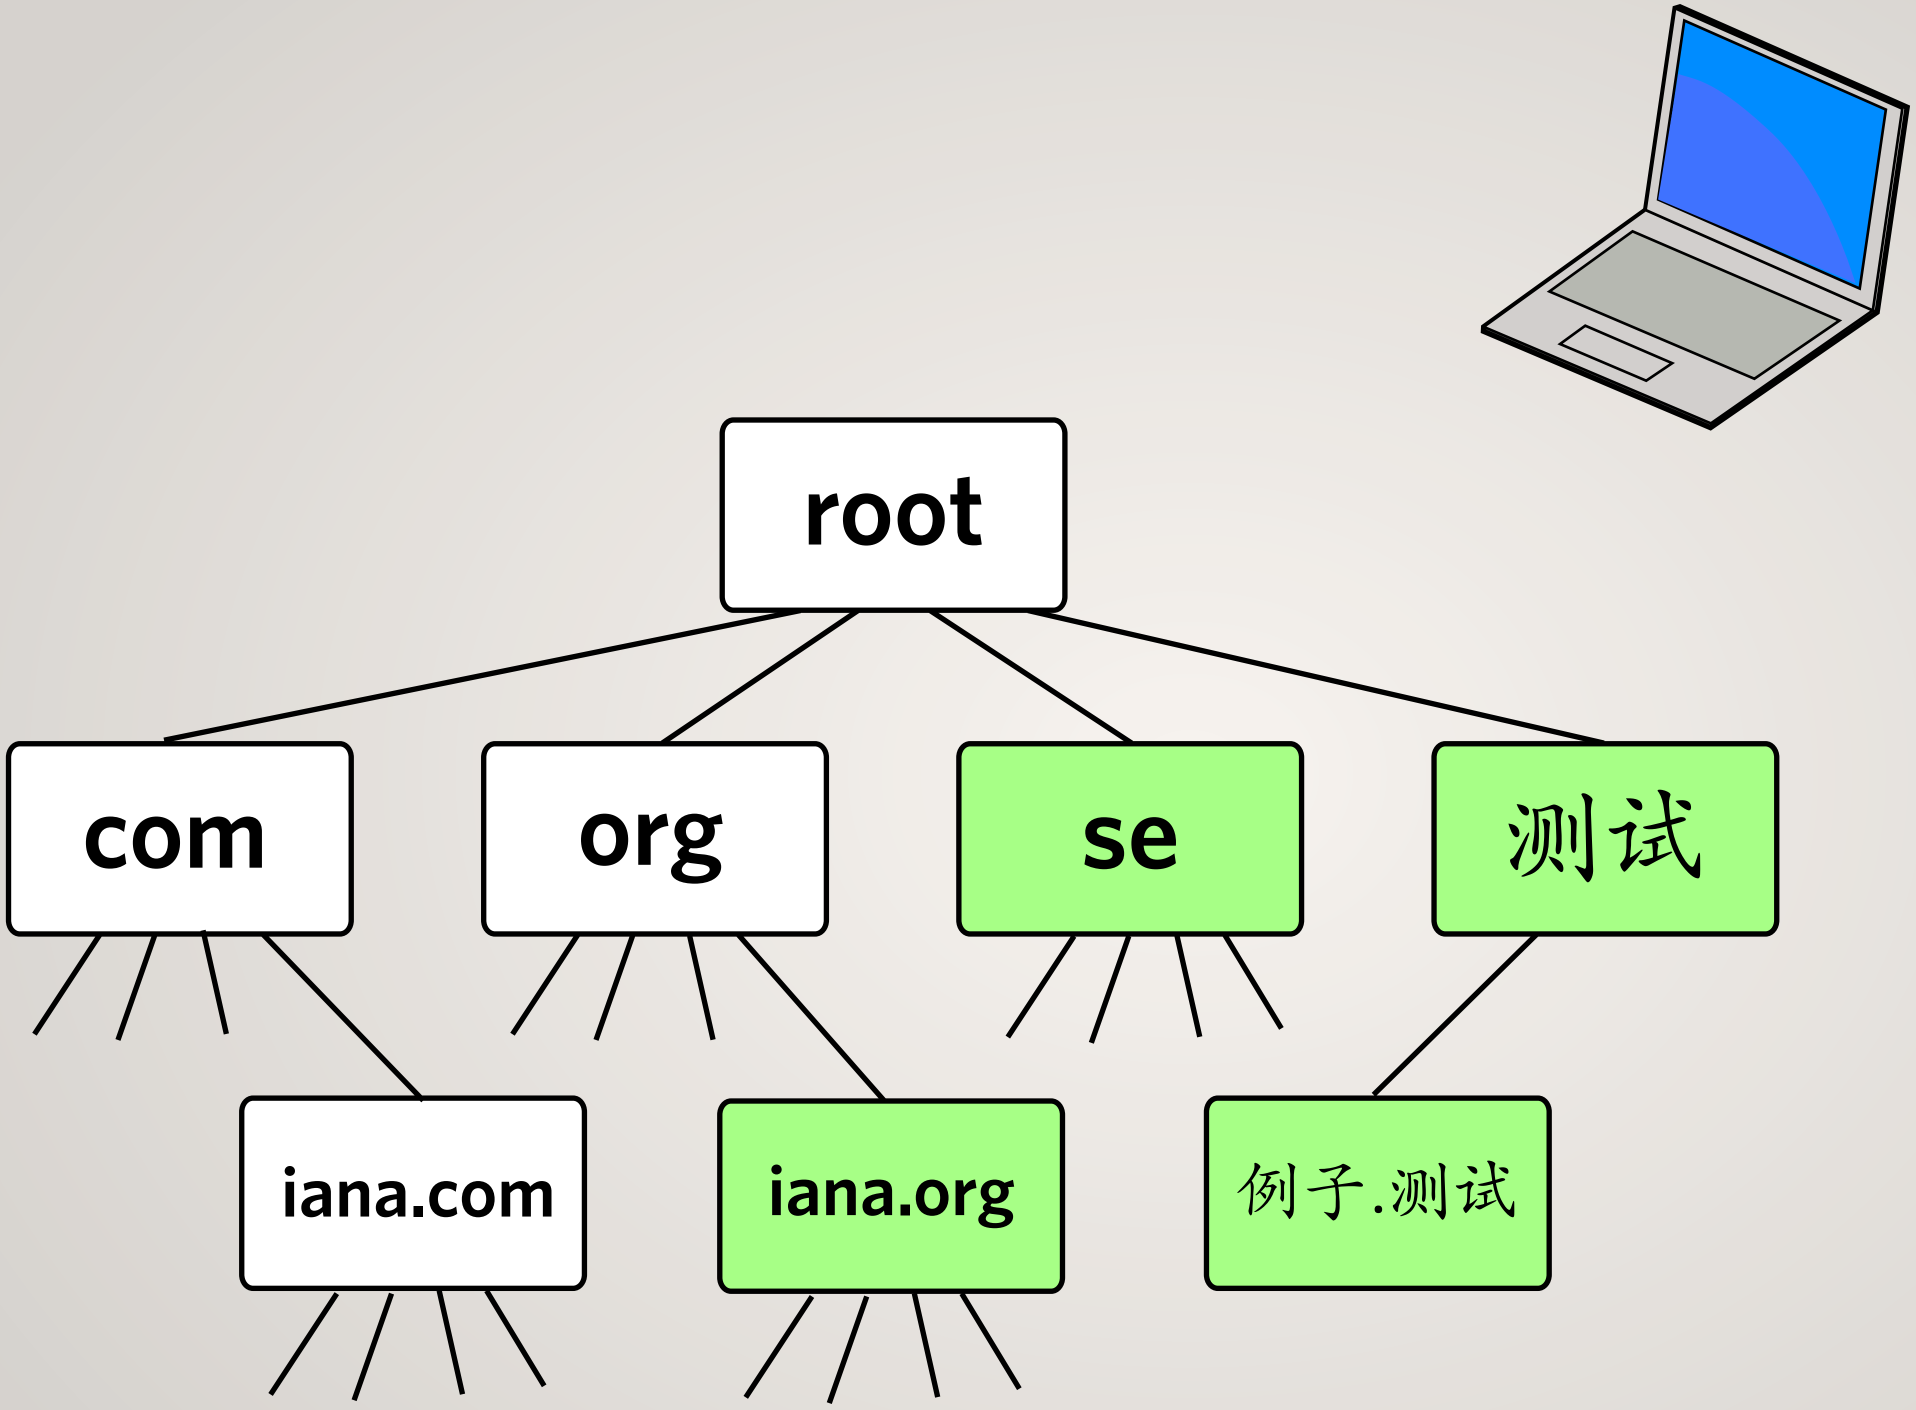

root

com

org

se

测试

iana.com

iana.org

例子.测试

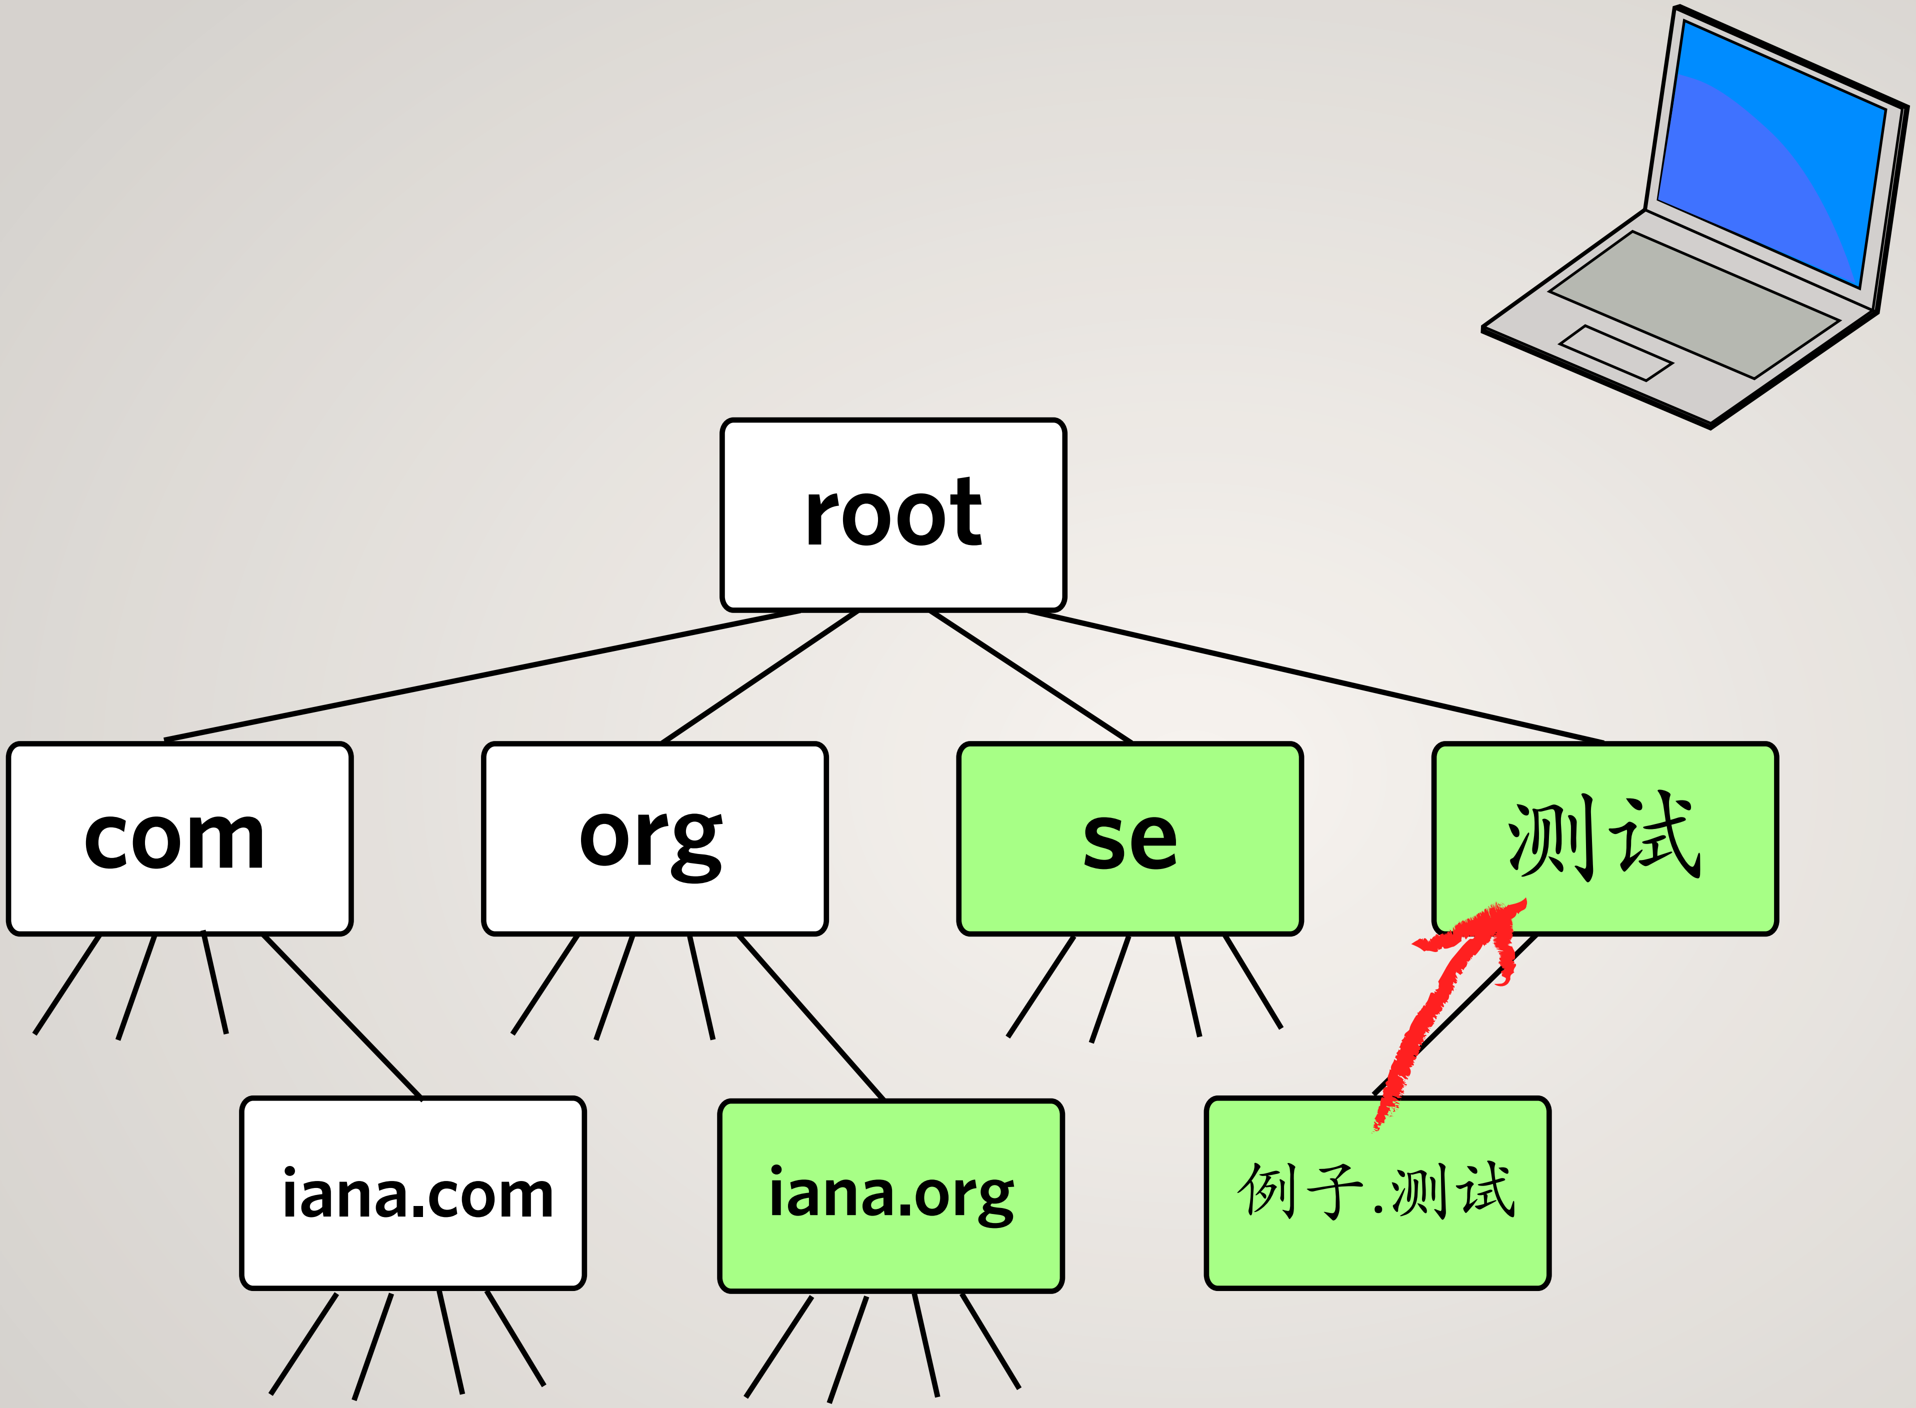

root

com

org

se

测试

iana.com

iana.org

例子.测试

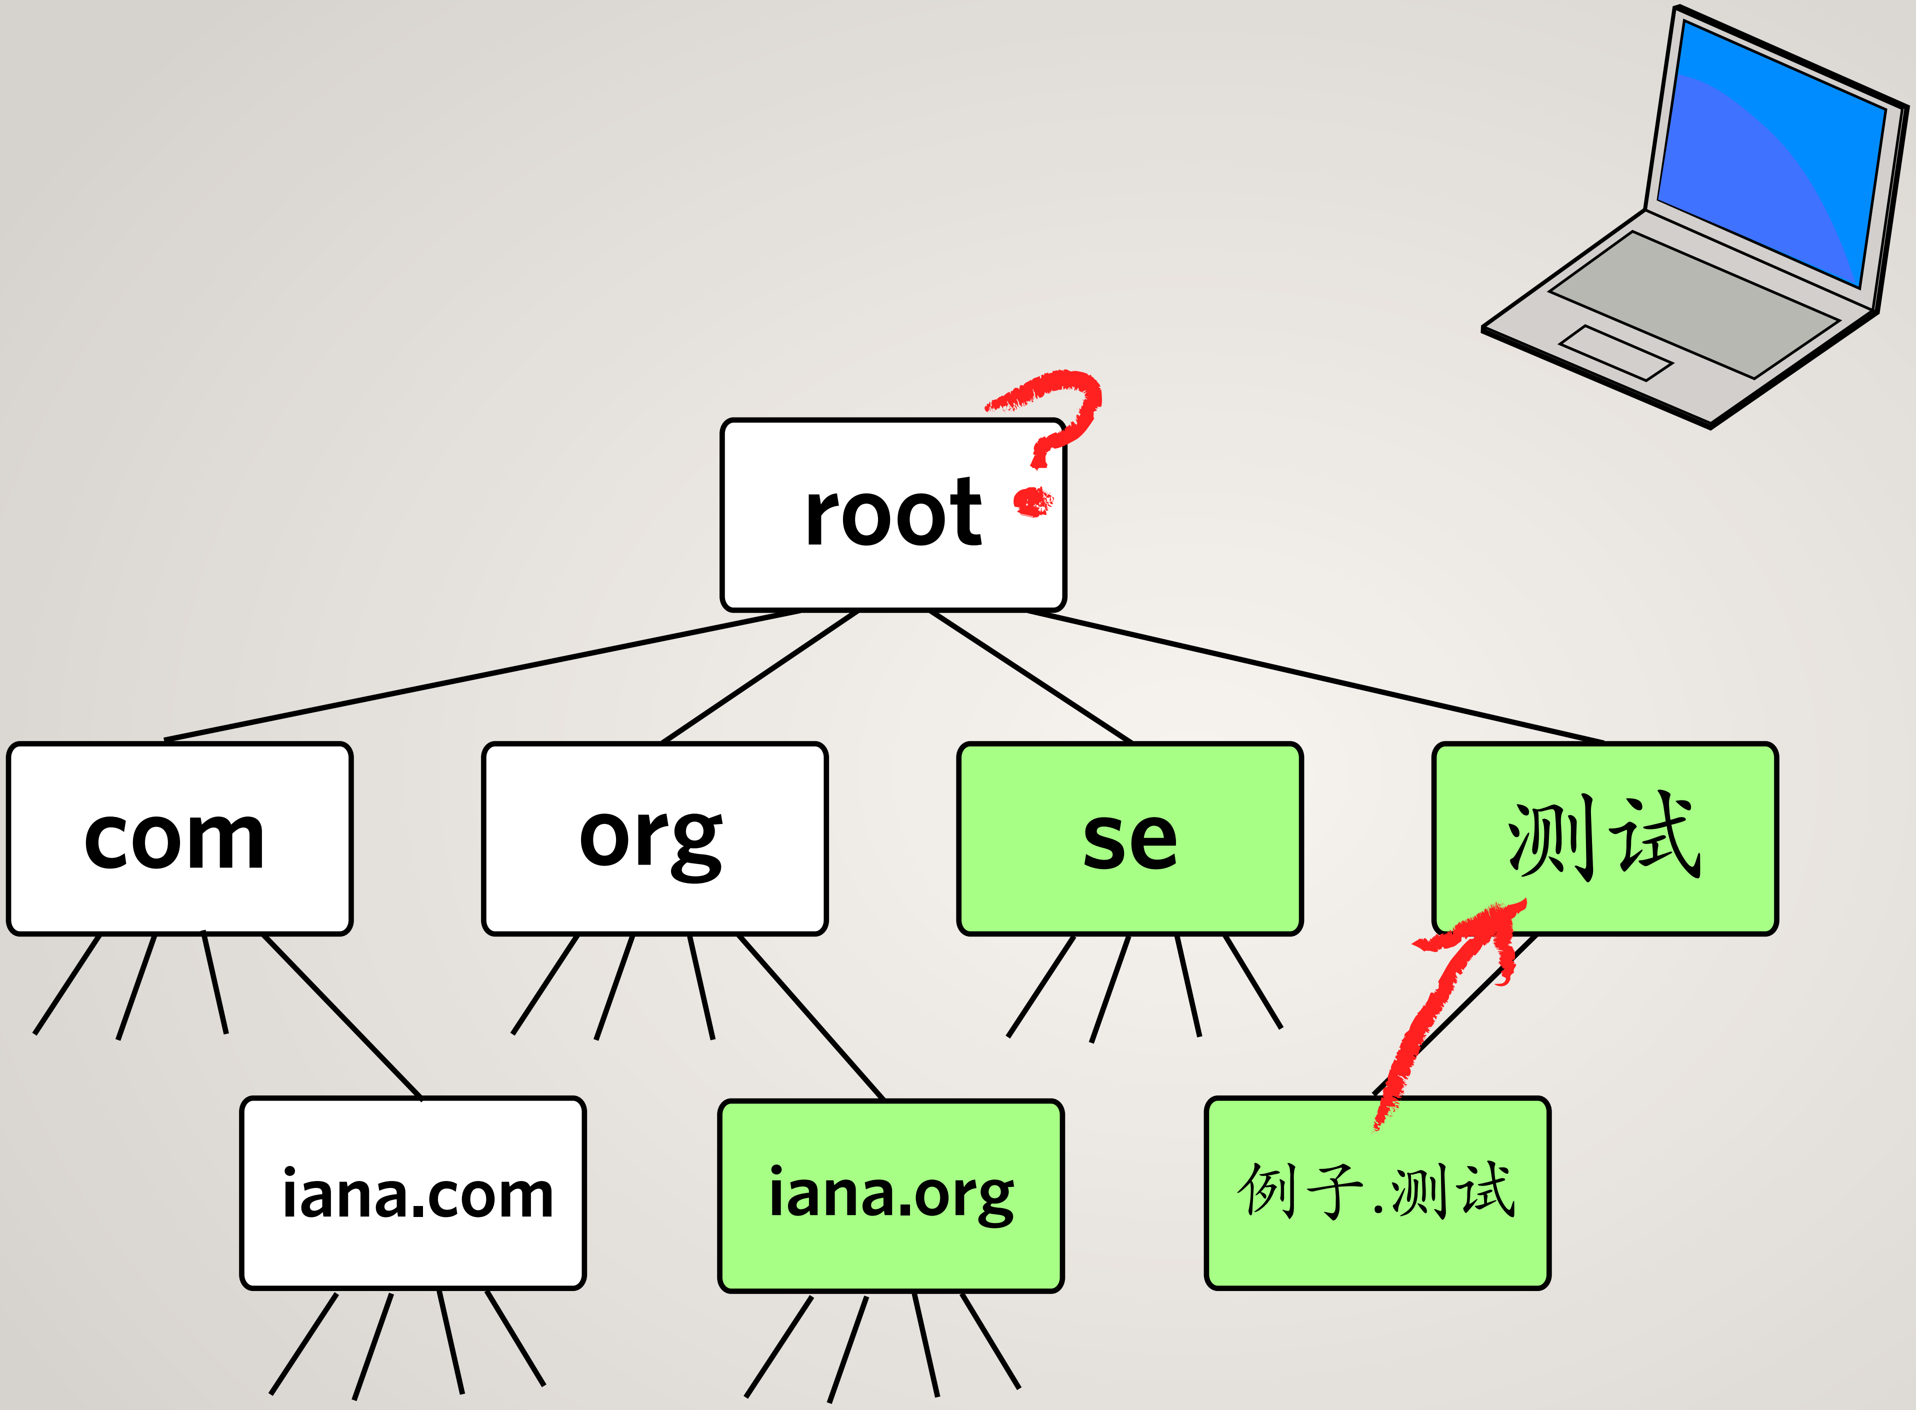
root
com
org
se
测试
iana.com
iana.org
例子.测试

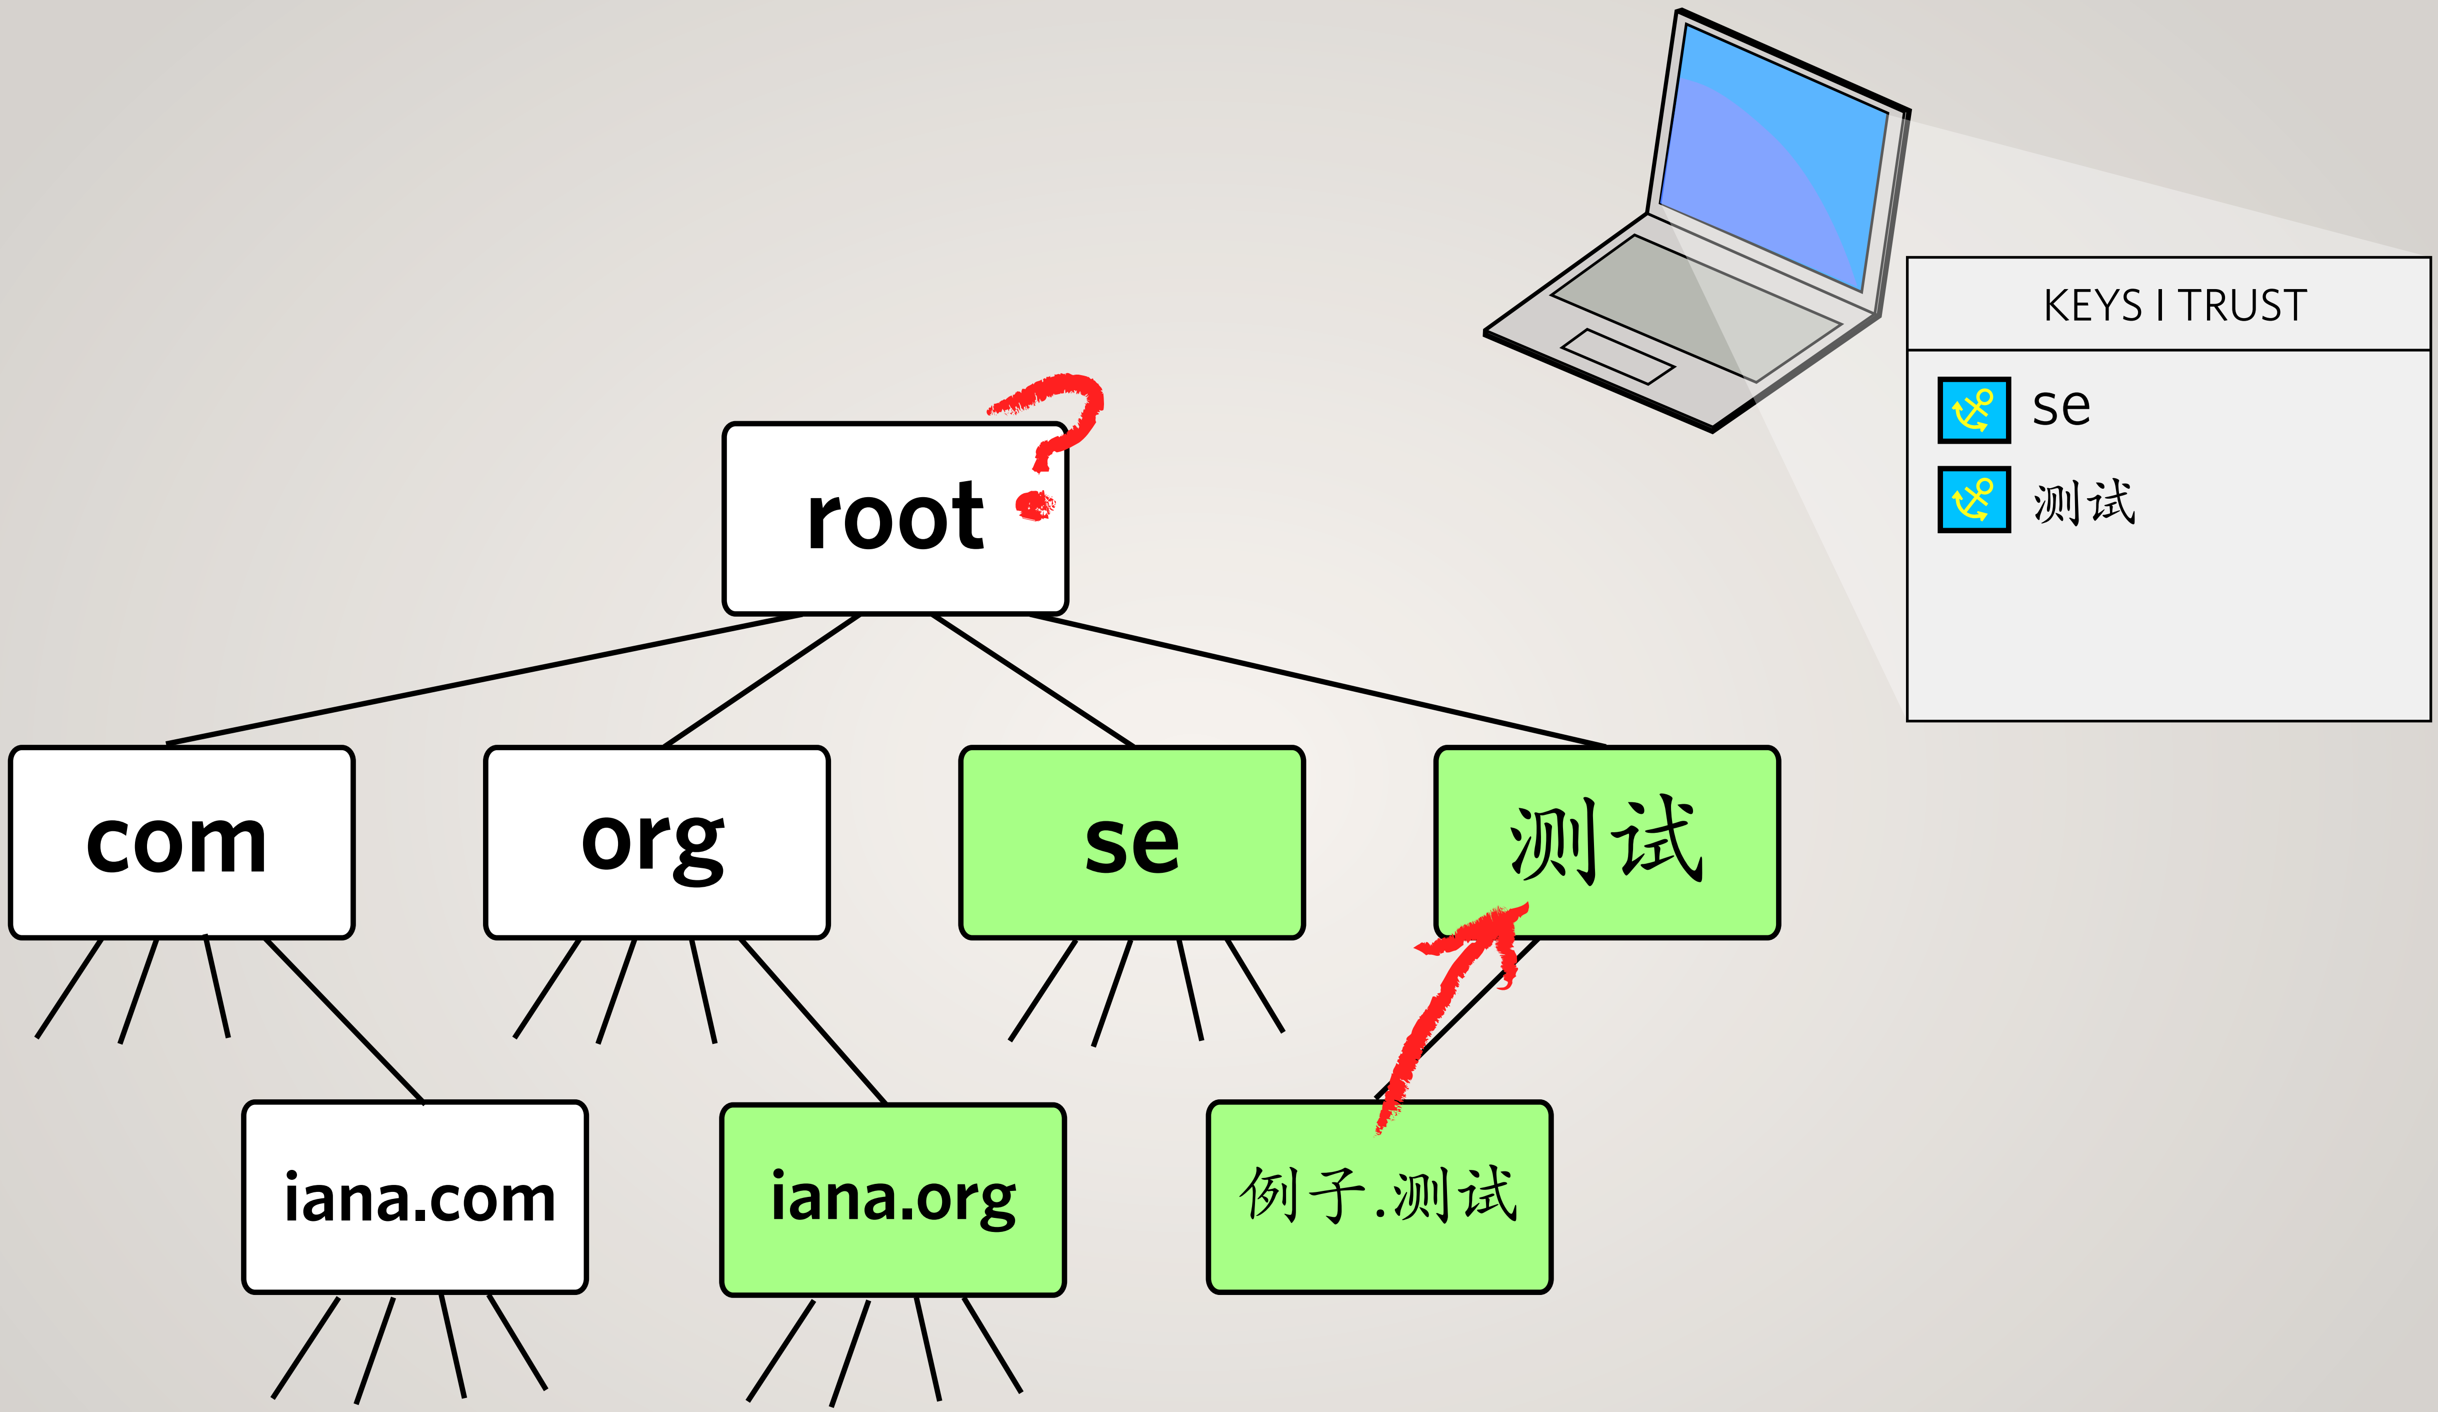

root

com

org

se

测试

iana.com

iana.org

例子.测试

KEYS I TRUST

⚓ se

⚓ 测试

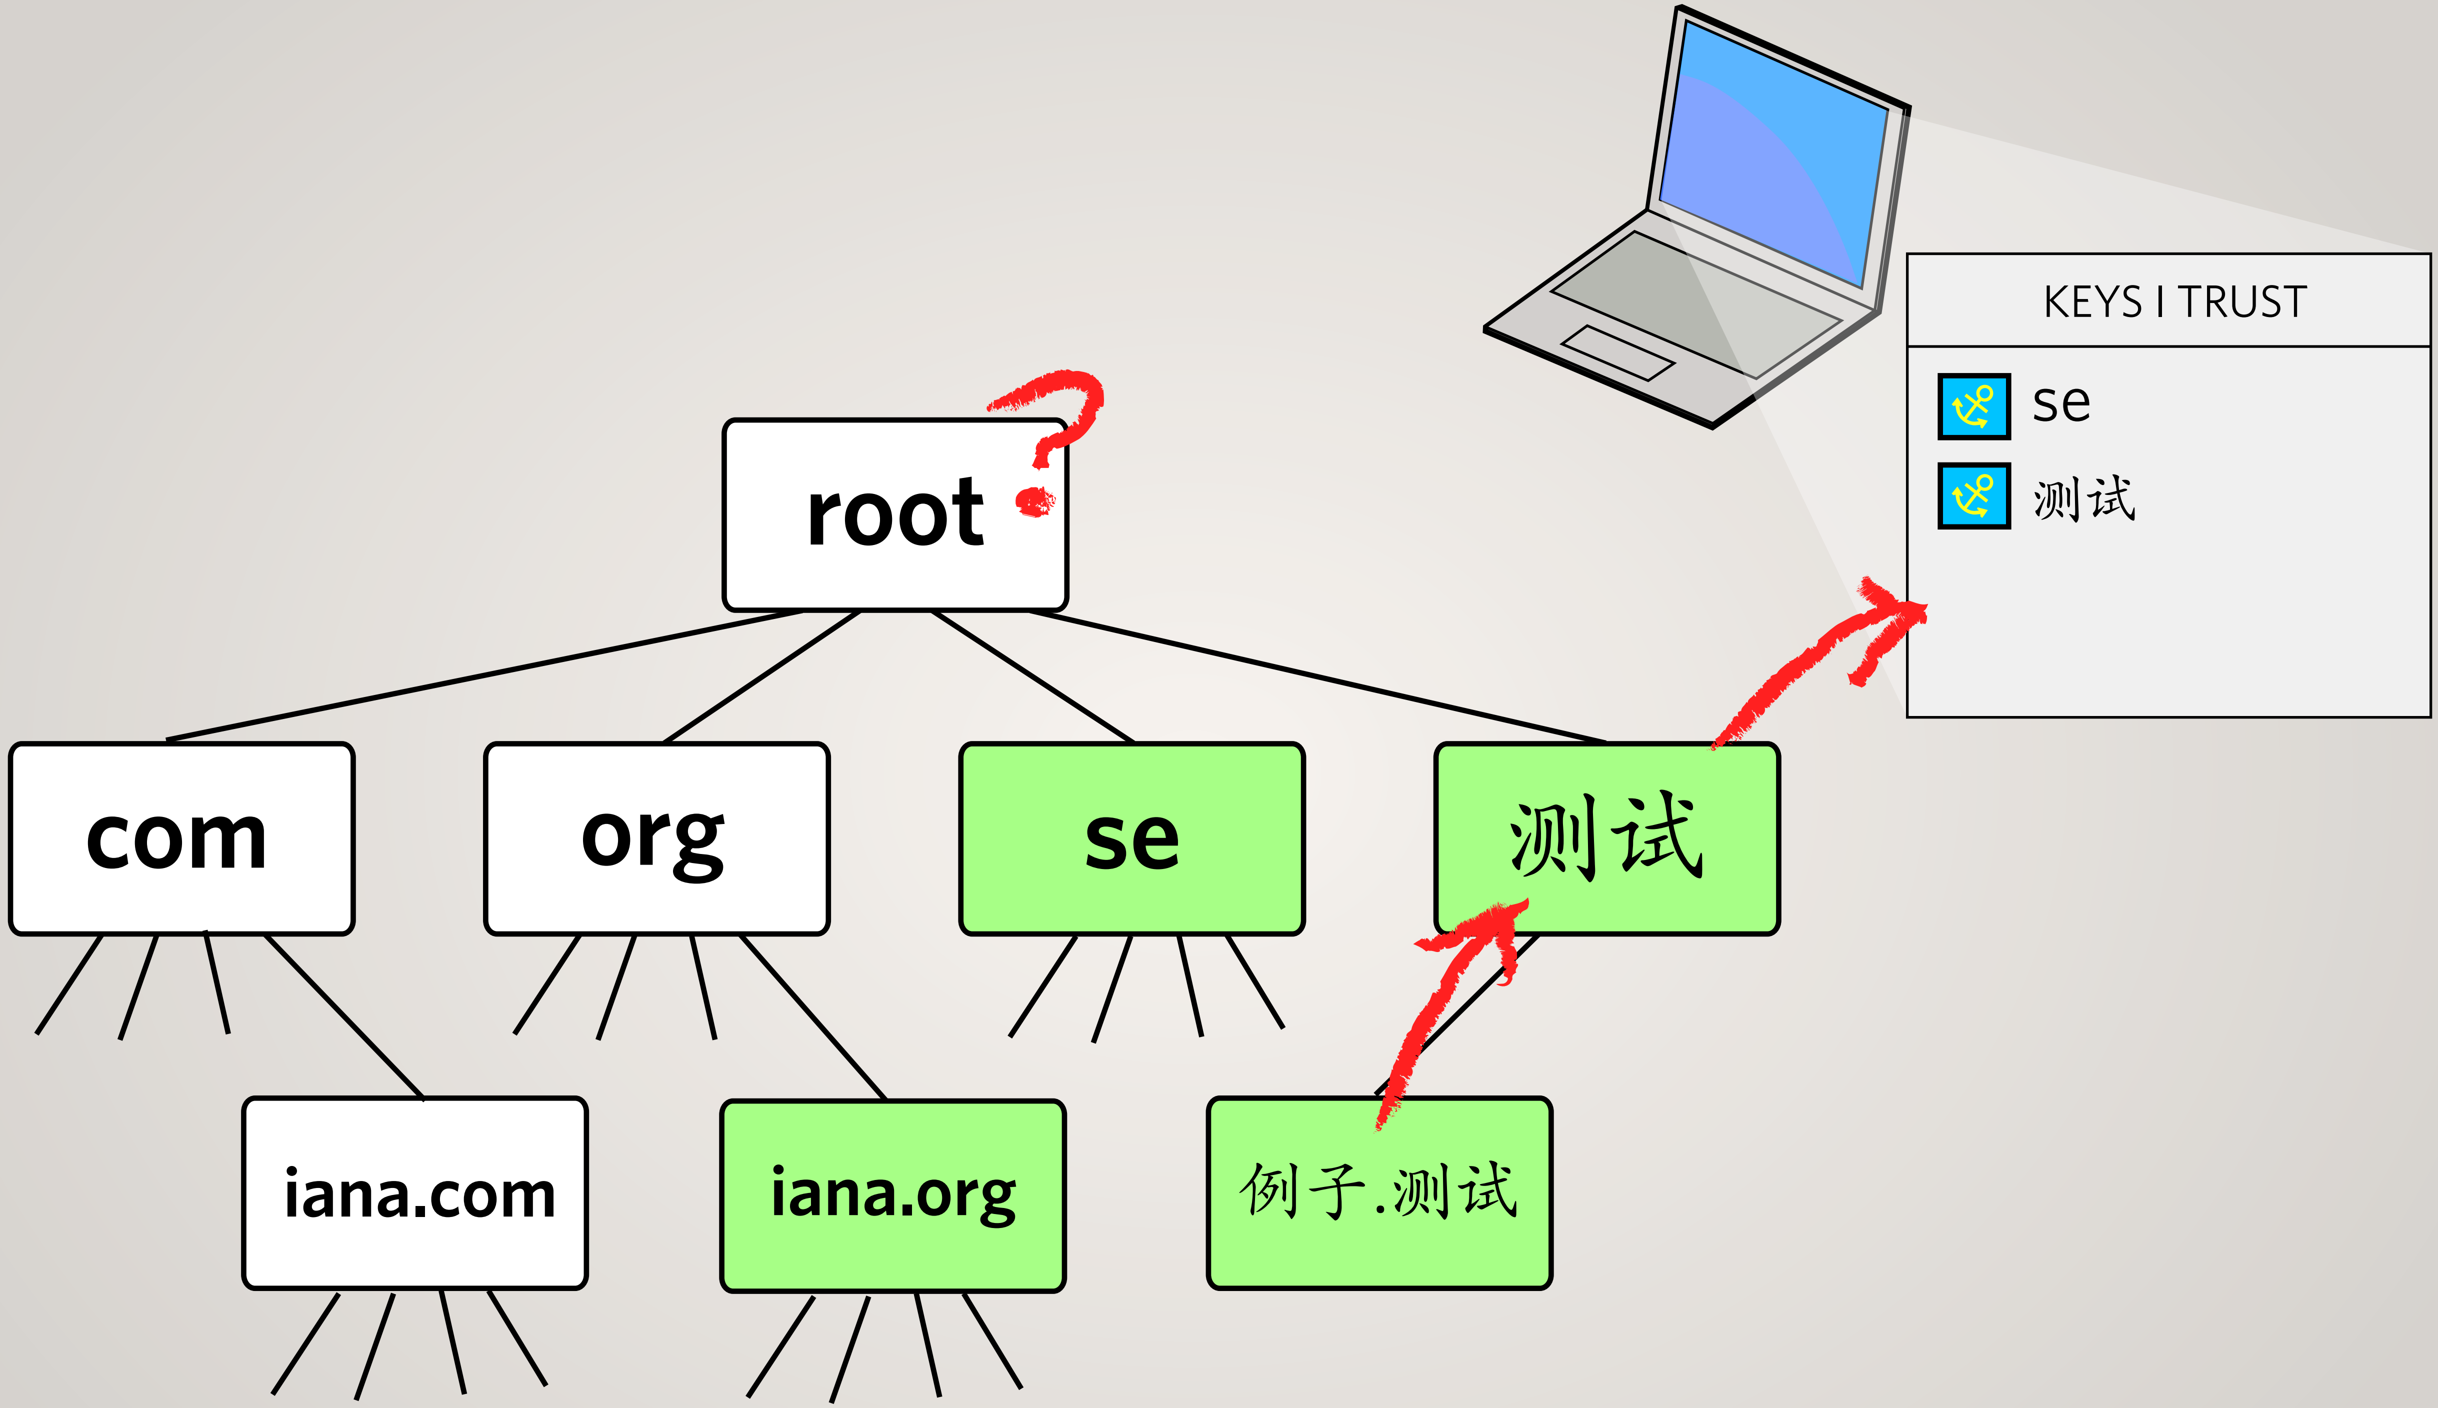

root

com

org

se

测试

iana.com

iana.org

例子.测试

KEYS I TRUST

se

测试

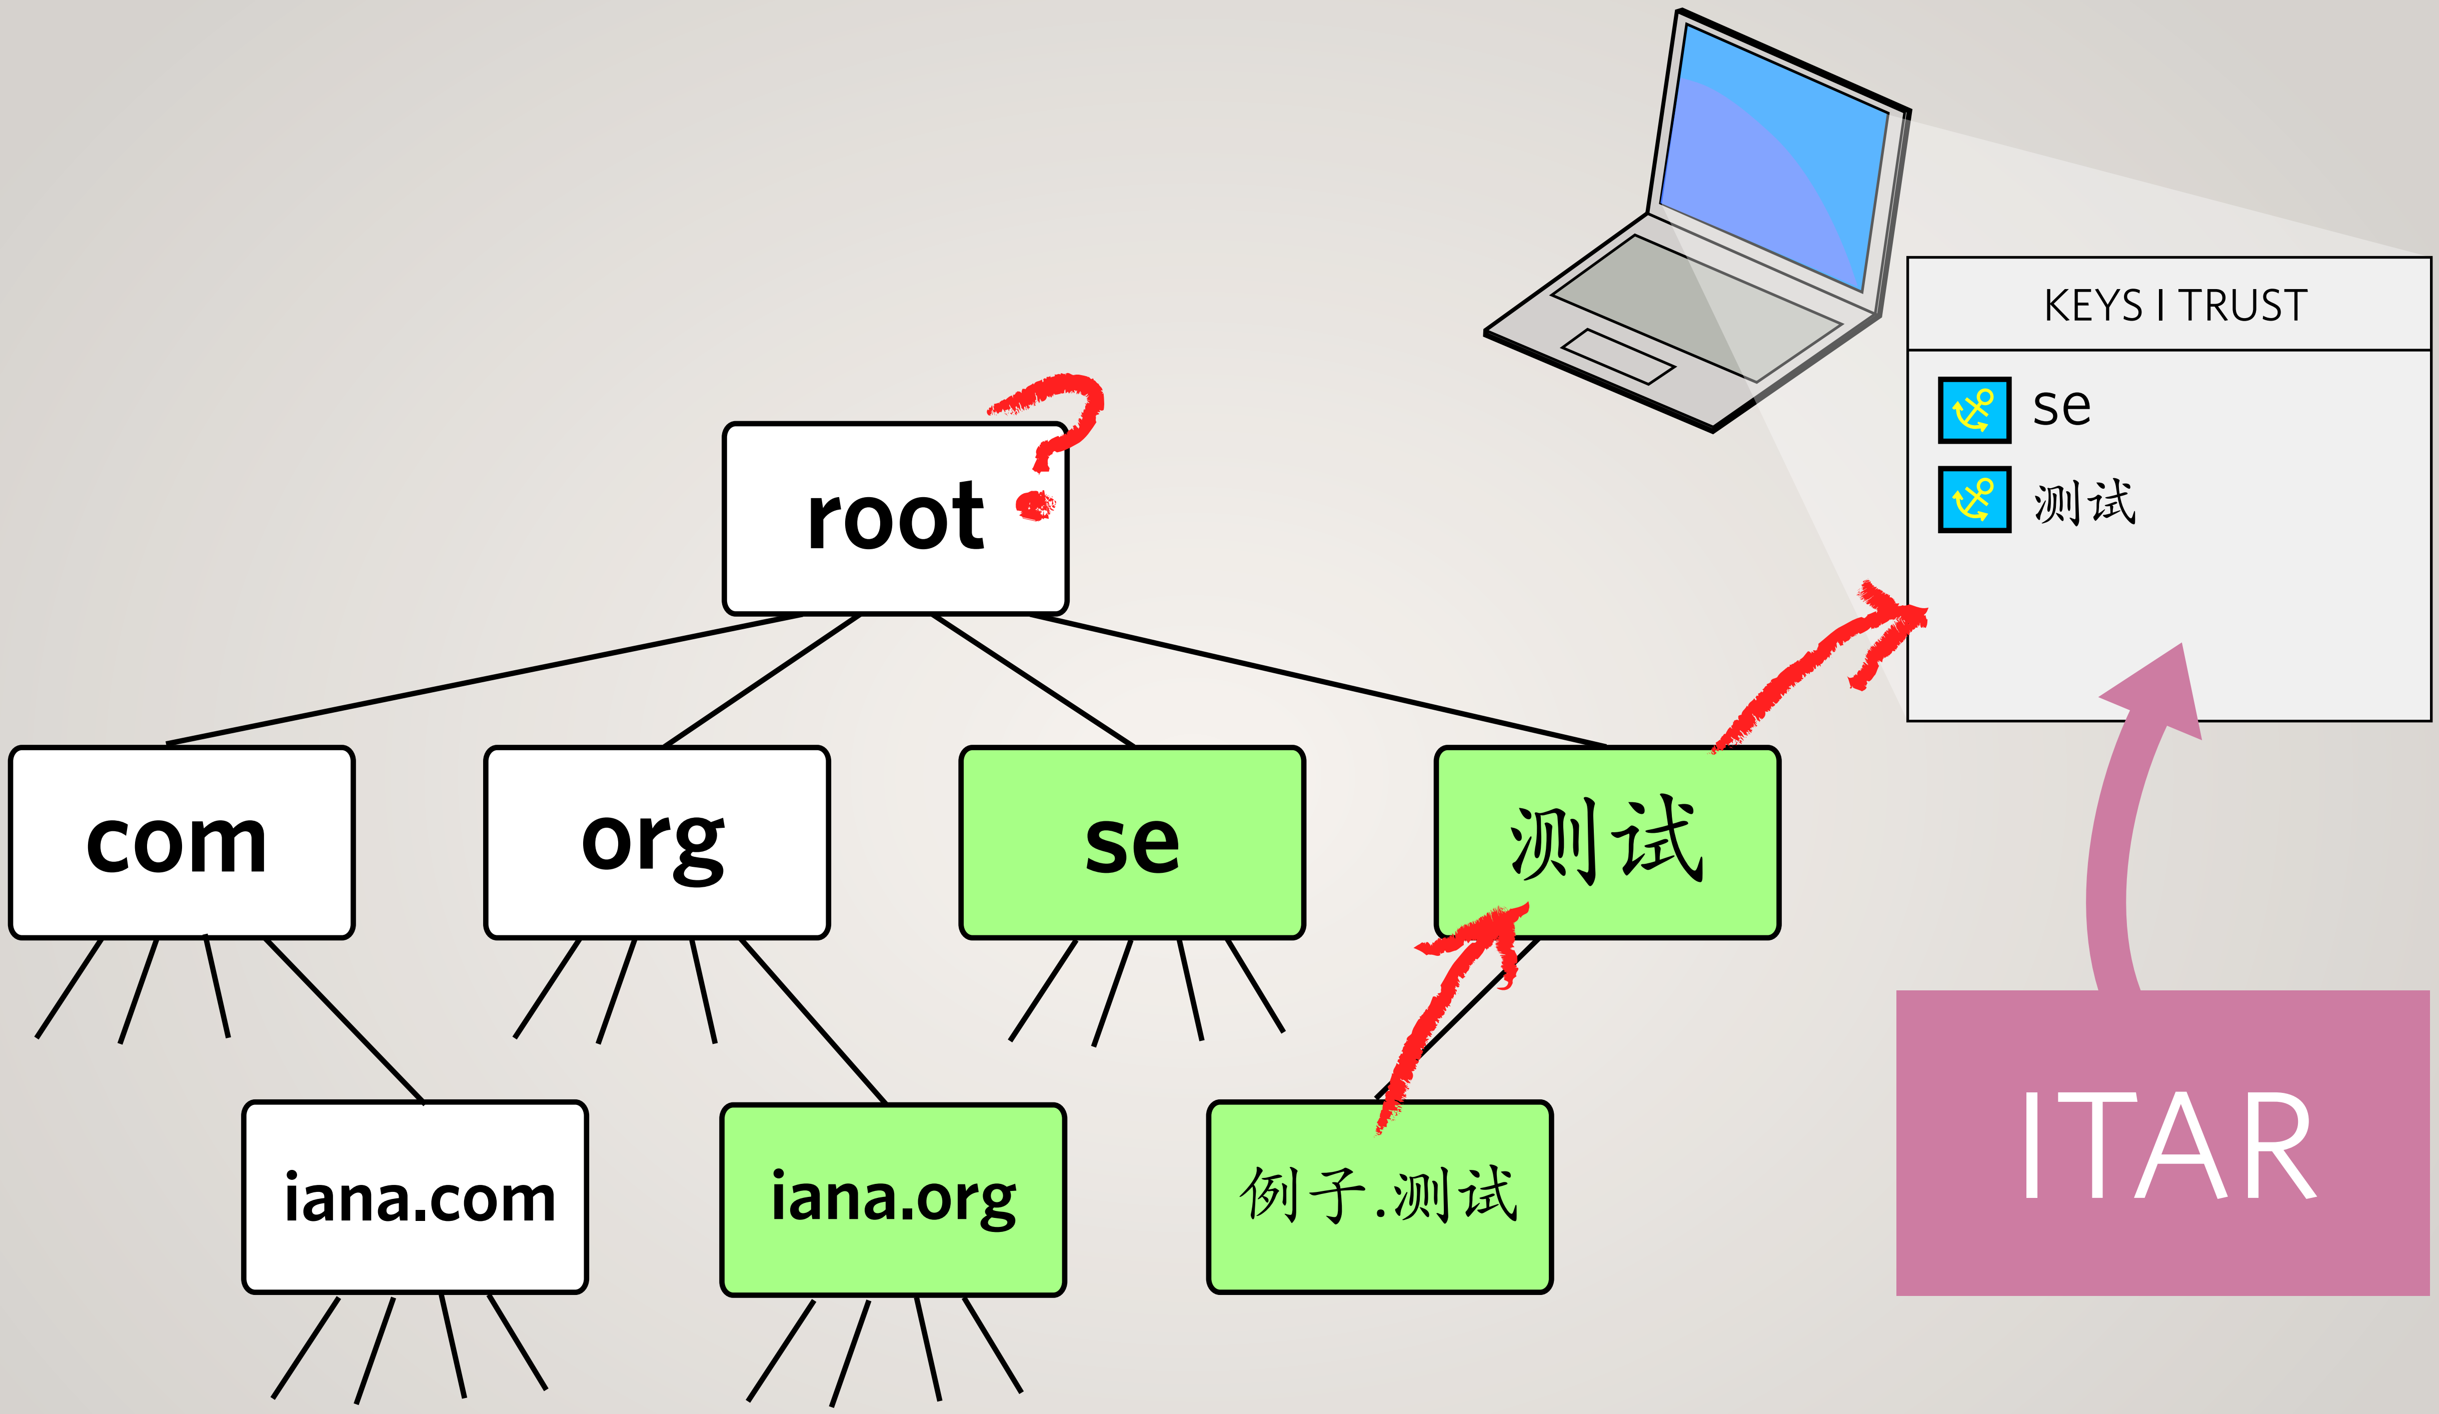
root
com
org
se
测试
iana.com
iana.org
例子.测试
KEYS I TRUST
se
测试
ITAR

Benefits

- Fully meets a set of recommendations provided by RIPE
- Simple to use for both top-level domain operators, and end users.
- Works with different DNS software, different protocols, etc. Non proprietary.
- Almost fully automated
- Helps DNSSEC deployment

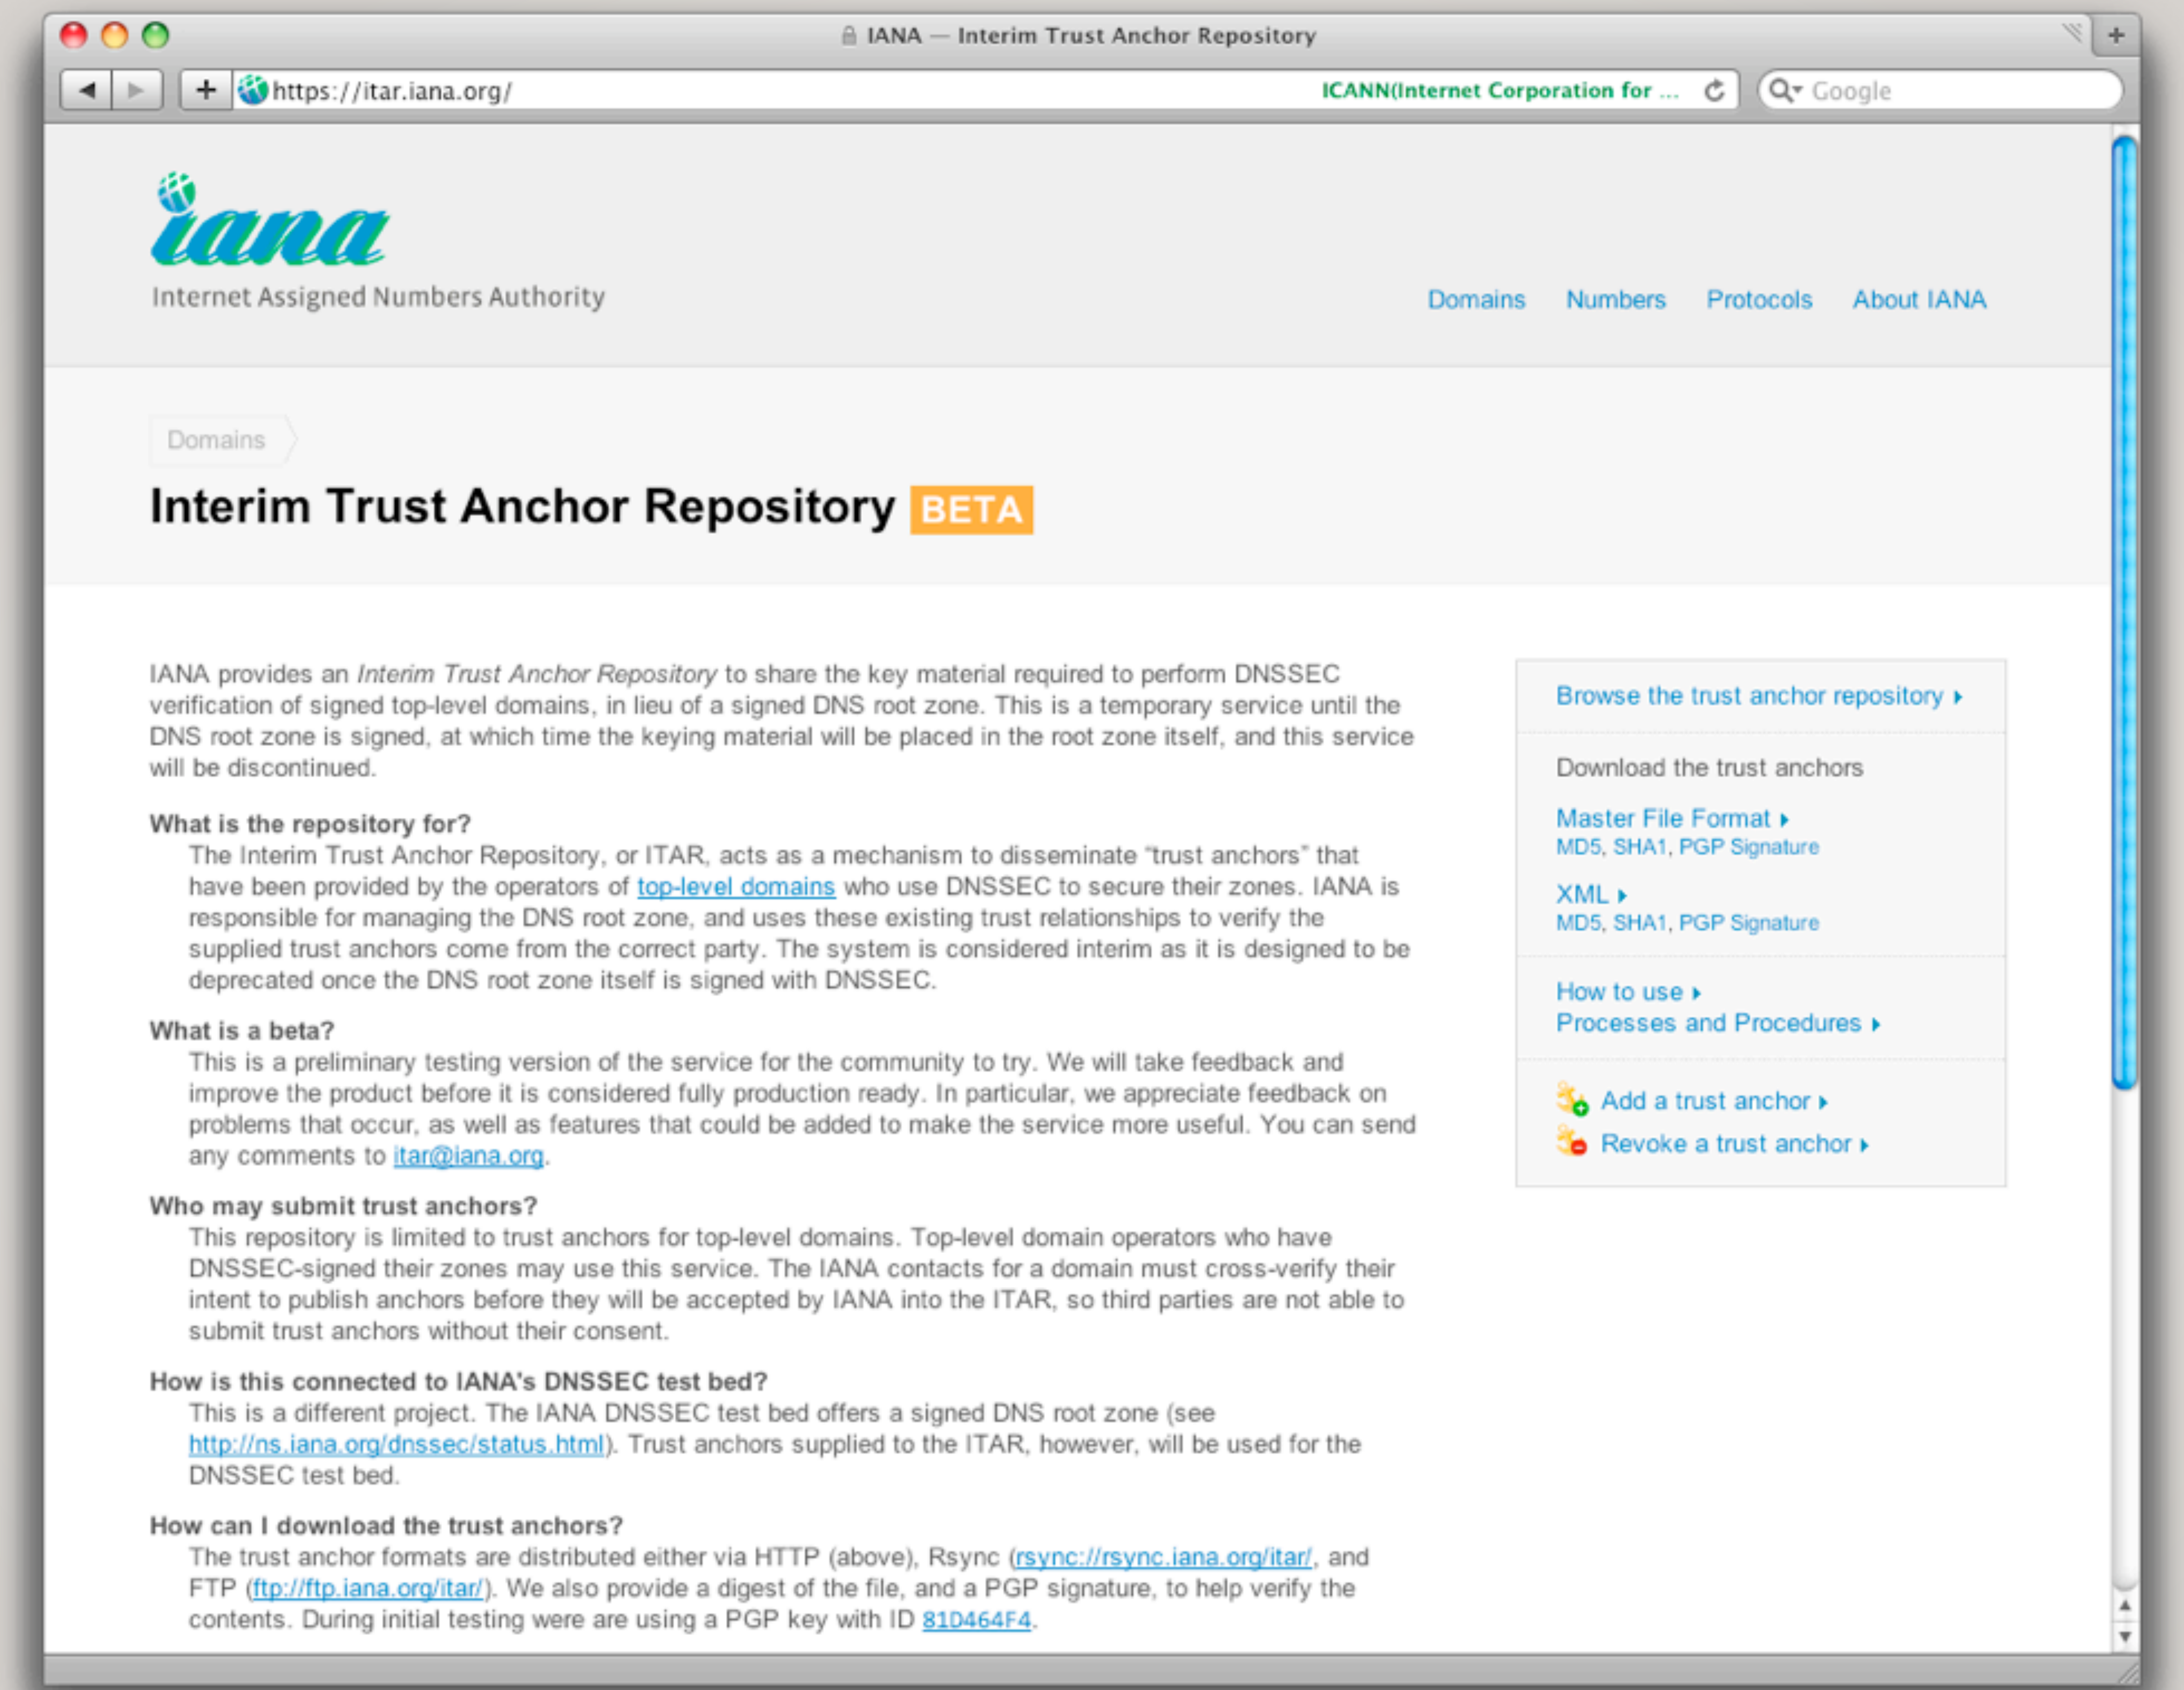

Interim Trust Anchor Repository BETA

IANA provides an *Interim Trust Anchor Repository* to share the key material required to perform DNSSEC verification of signed top-level domains, in lieu of a signed DNS root zone. This is a temporary service until the DNS root zone is signed, at which time the keying material will be placed in the root zone itself, and this service will be discontinued.

What is the repository for?

The Interim Trust Anchor Repository, or ITAR, acts as a mechanism to disseminate "trust anchors" that have been provided by the operators of top-level domains who use DNSSEC to secure their zones. IANA is responsible for managing the DNS root zone, and uses these existing trust relationships to verify the supplied trust anchors come from the correct party. The system is considered interim as it is designed to be deprecated once the DNS root zone itself is signed with DNSSEC.

What is a beta?

This is a preliminary testing version of the service for the community to try. We will take feedback and improve the product before it is considered fully production ready. In particular, we appreciate feedback on problems that occur, as well as features that could be added to make the service more useful. You can send any comments to itar@iana.org.

Who may submit trust anchors?

This repository is limited to trust anchors for top-level domains. Top-level domain operators who have DNSSEC-signed their zones may use this service. The IANA contacts for a domain must cross-verify their intent to publish anchors before they will be accepted by IANA into the ITAR, so third parties are not able to submit trust anchors without their consent.

How is this connected to IANA's DNSSEC test bed?

This is a different project. The IANA DNSSEC test bed offers a signed DNS root zone (see http://ns.iana.org/dnssec/status.html). Trust anchors supplied to the ITAR, however, will be used for the DNSSEC test bed.

How can I download the trust anchors?

The trust anchor formats are distributed either via HTTP (above), Rsync (rsync://rsync.iana.org/itar/, and FTP (ftp://ftp.iana.org/itar/). We also provide a digest of the file, and a PGP signature, to help verify the contents. During initial testing were are using a PGP key with ID 81D464F4.

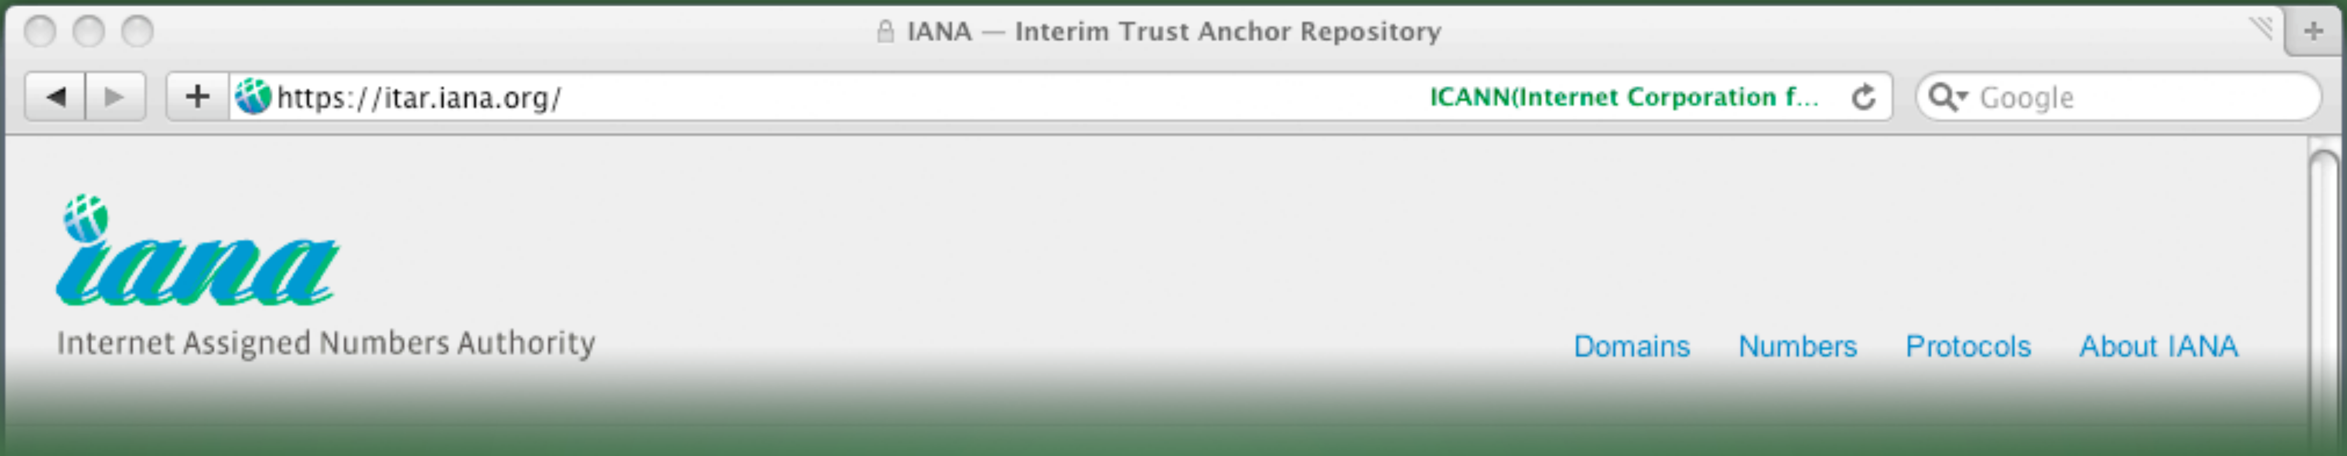

IANA — Interim Trust Anchor Repository

https://itar.iana.org/

ICANN(Internet Corporation f...

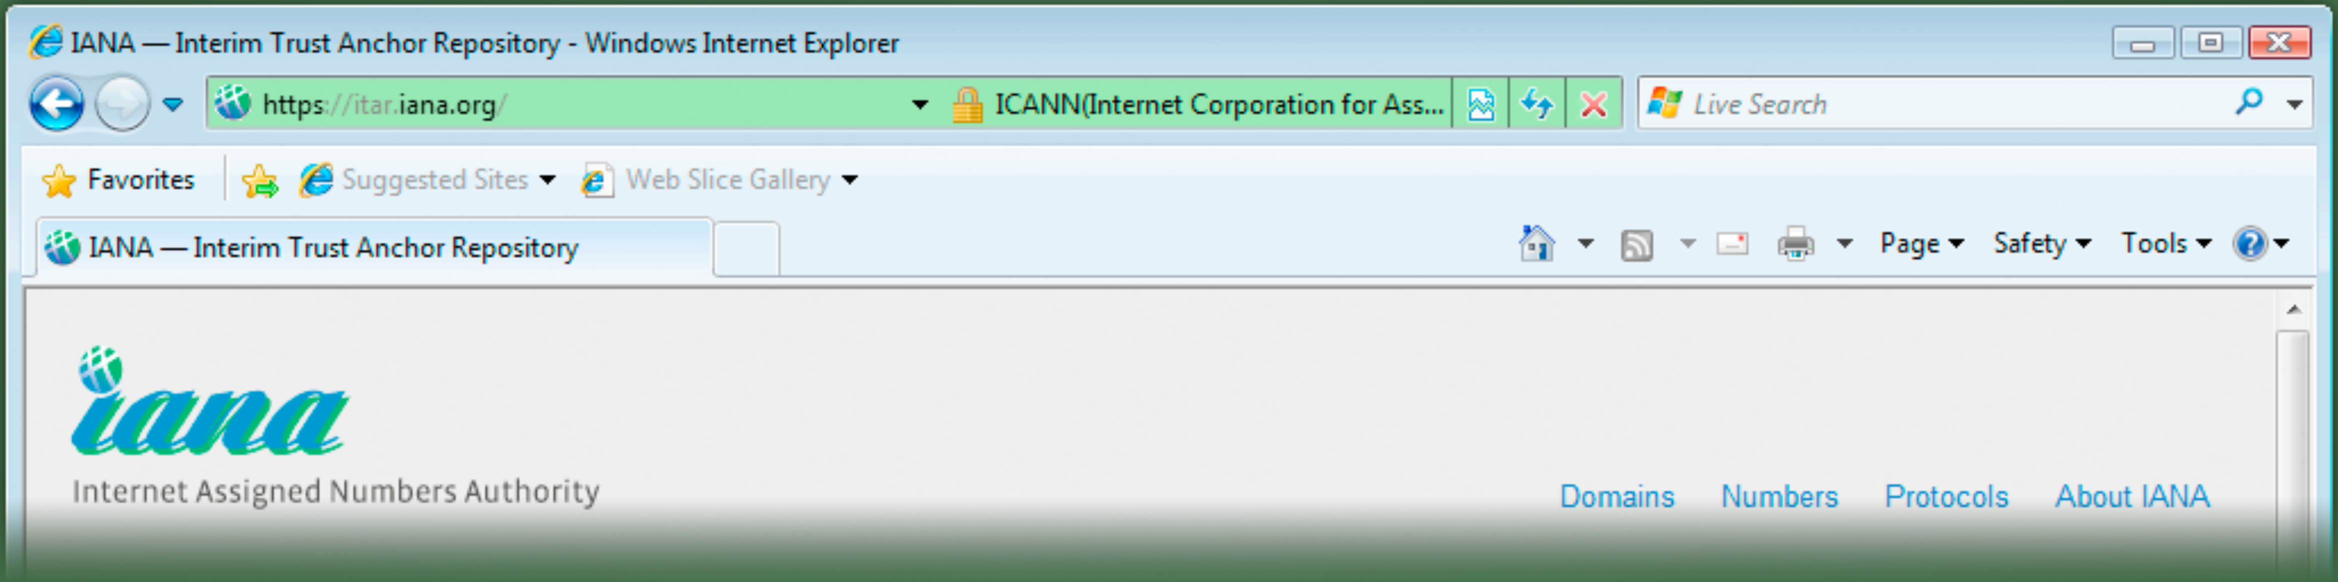

IANA — Interim Trust Anchor Repository - Windows Internet Explorer

https://itar.iana.org/

ICANN(Internet Corporation for Ass...

itar.iana.org

Thanks!

kim.davies@icann.org