Zorgeloos online

(For confidence online)

Background image: Thijs van den Hout & Midjourney

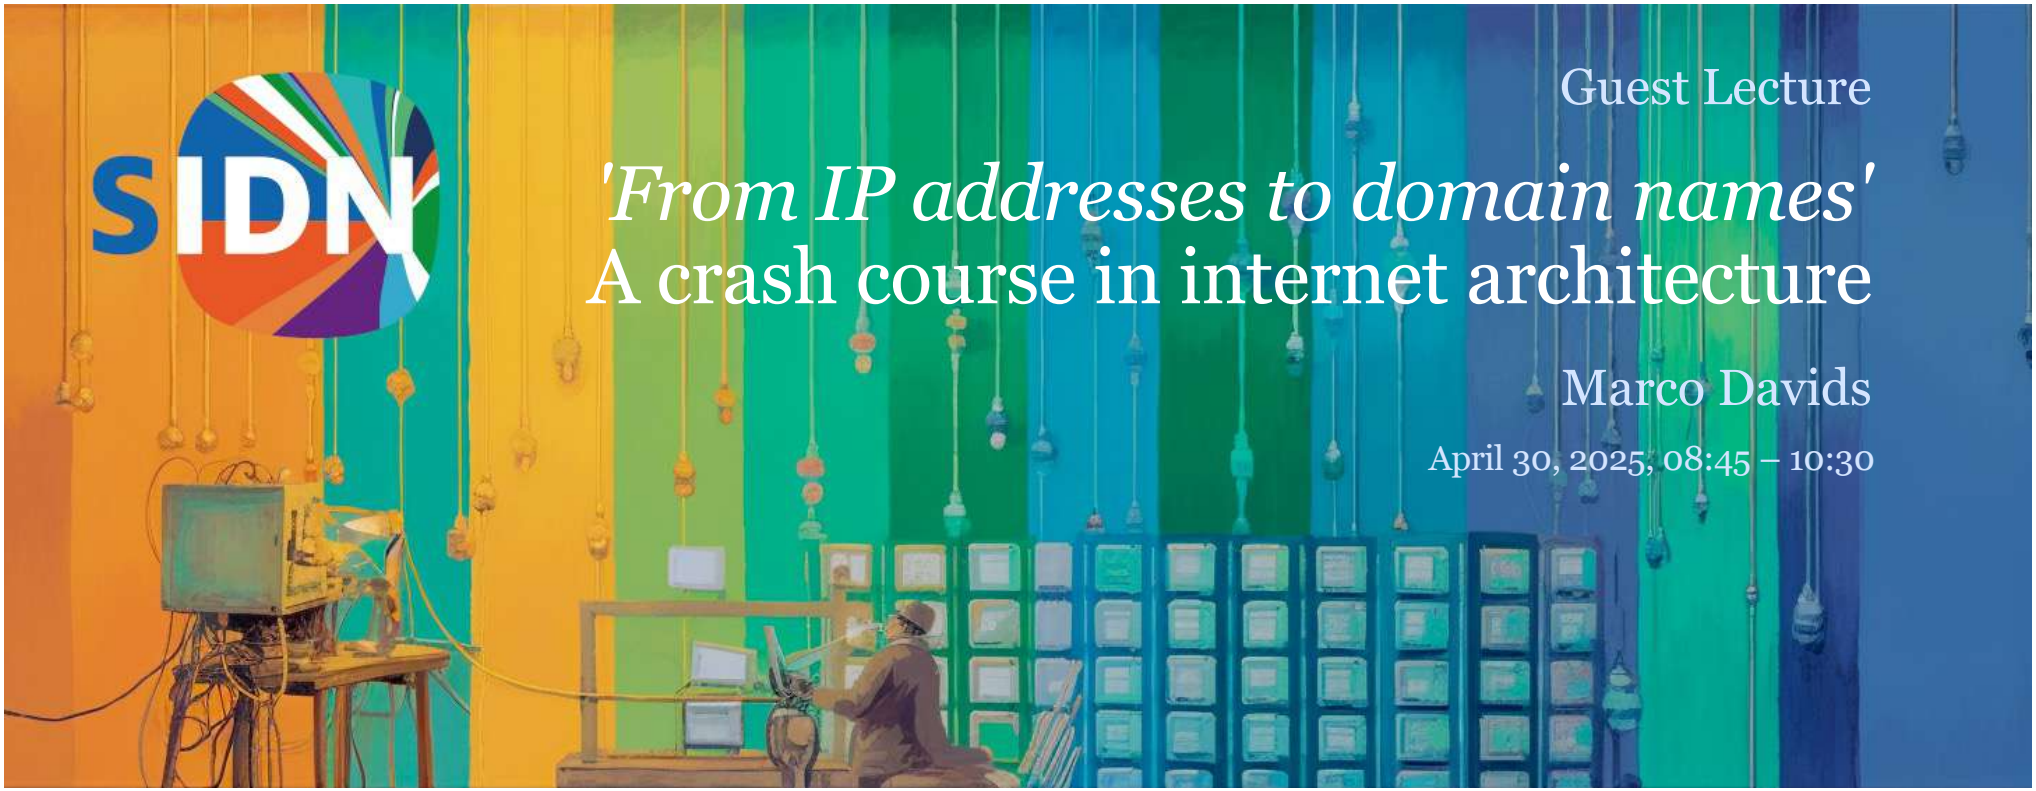

Guest Lecture

'From IP addresses to domain names'
A crash course in internet architecture

Marco Davids

April 30, 2025, 08:45 – 10:30

@marcodavids

SIDN LABS

"A name indicates what we seek.
An address indicates where it is.
A route indicates how to get there."

-- Jon Postel, RFC 791: Internet Protocol (1981)

About SIDN

- *Registry* for the .nl *country code top-level* domain
  - Nowadays also .amsterdam, .politie and .aw (technical management)
  - 6,2 million .nl domain names - 61% DNSSEC
  - Brand monitoring, .nl-Control, portfolio checker, abuse204.nl, etc.
- SIDN Fonds
- SIDN Labs

https://www.sidn.nl/en/about-sidn/what-we-do

Number of .nl domain names: 4th largest ccTLD

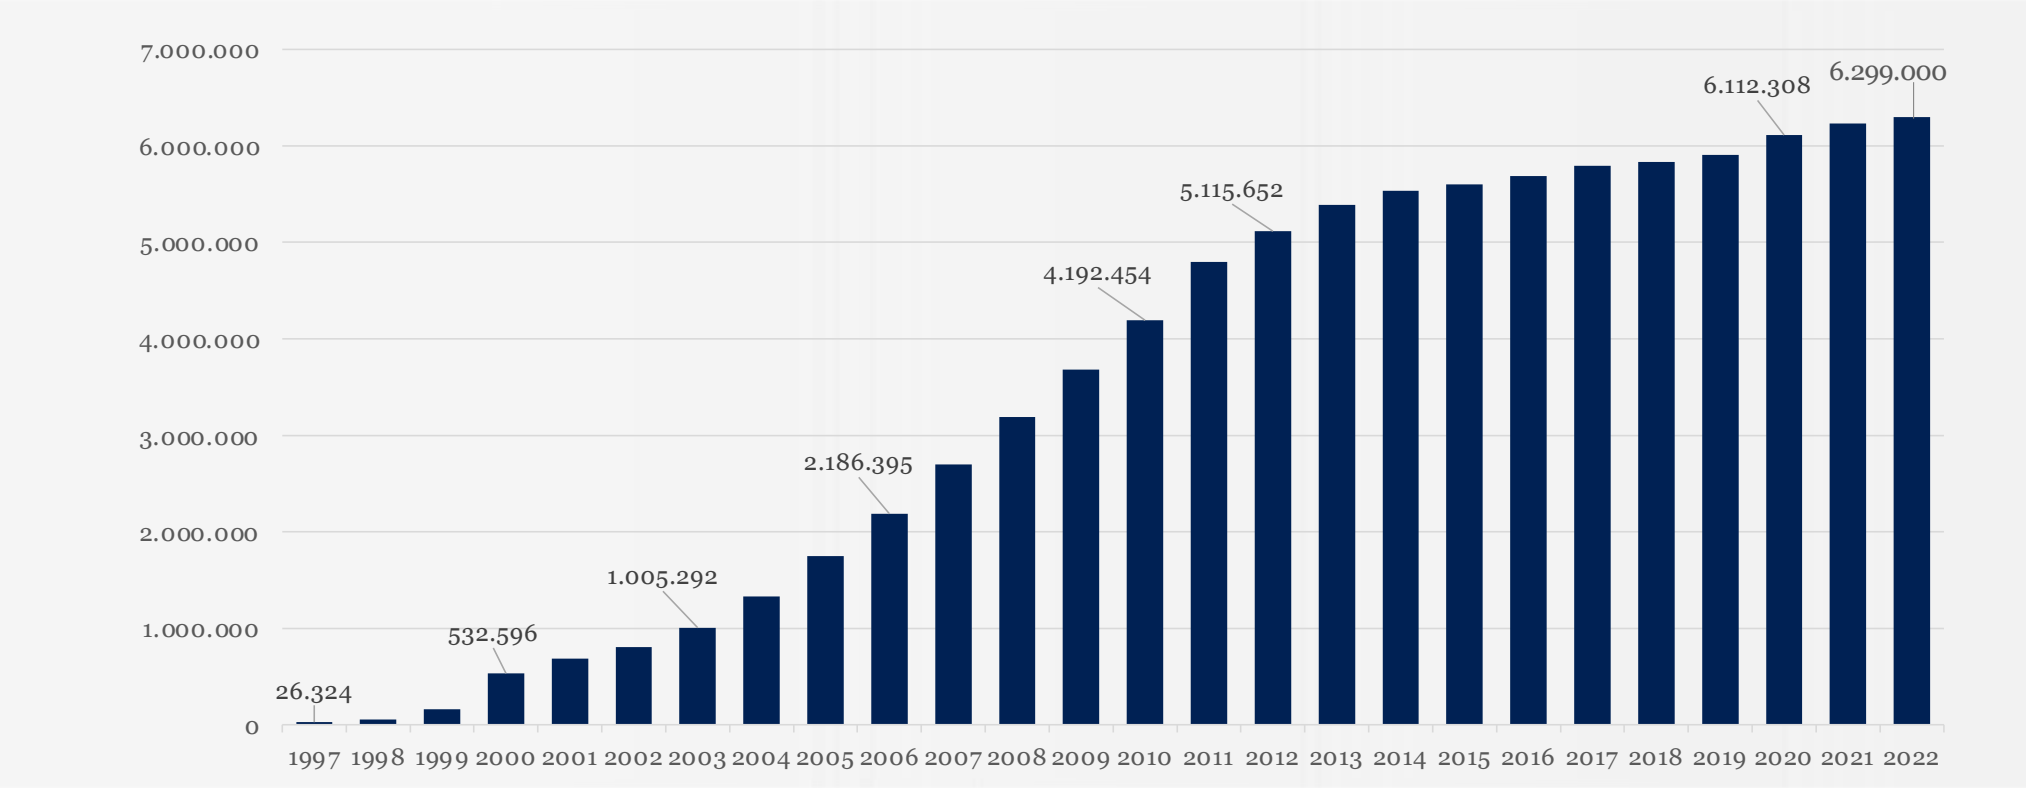

Number of .nl domain names: 4th largest ccTLD

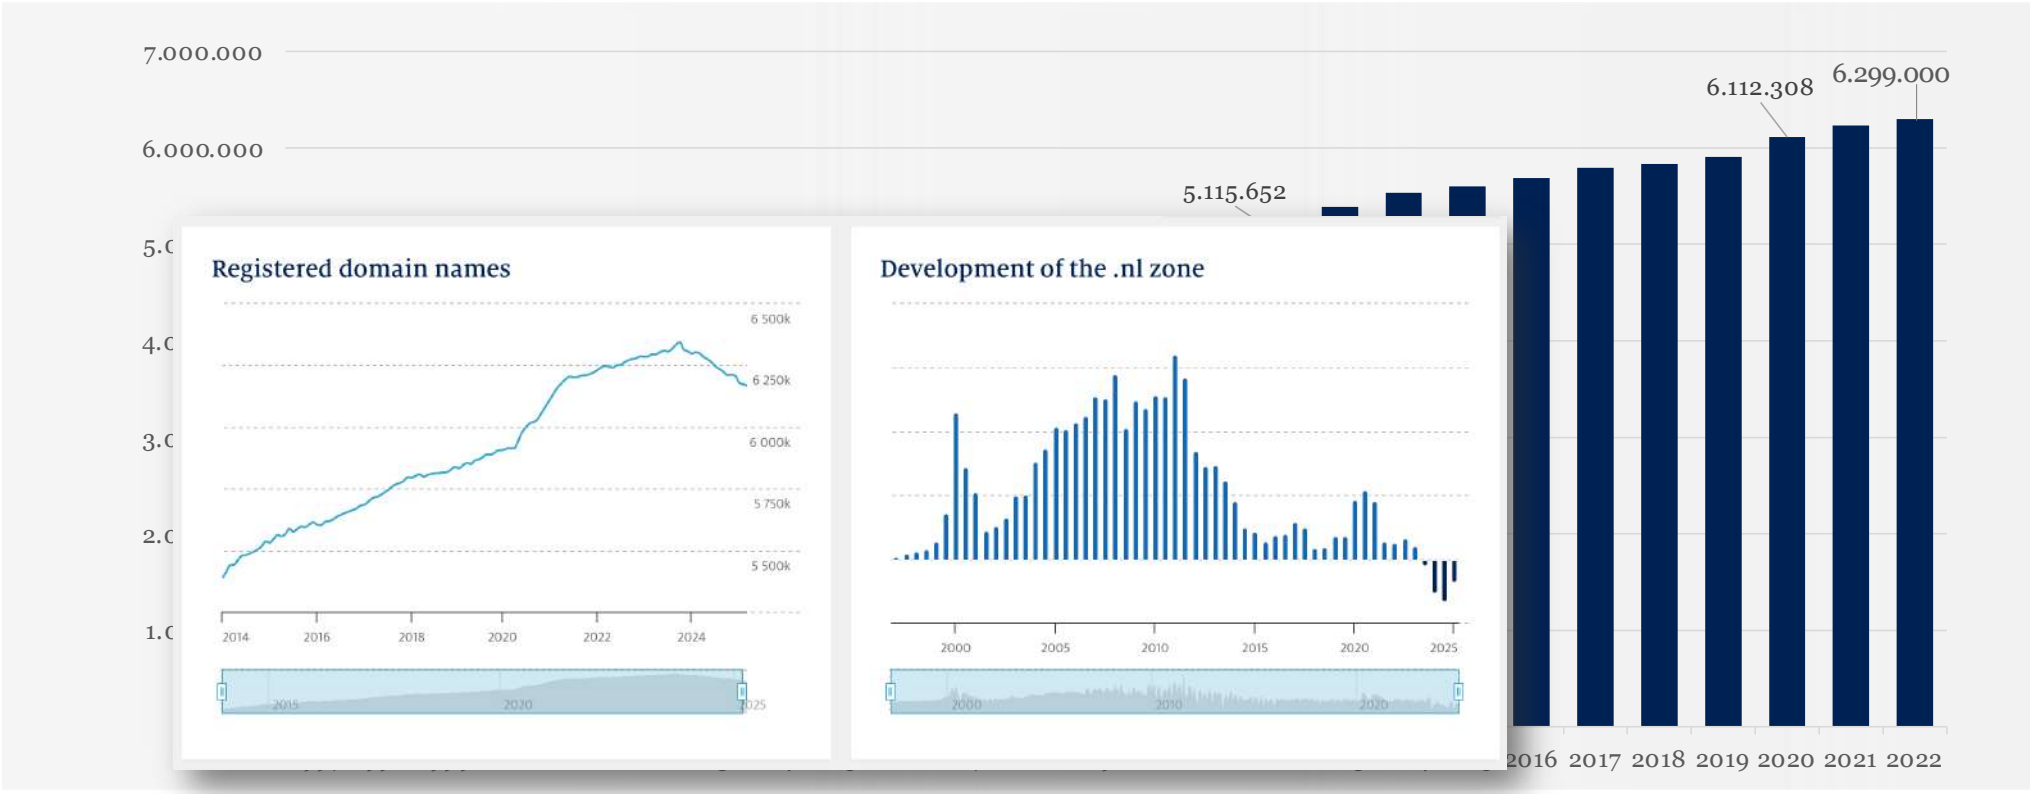

About SIDN Labs

Applied technical research

into the security of internet infrastructure

- Three themes:
  - Domain name security
  - Infrastructure security
  - Emerging internet technologies

https://www.sidnlabs.nl/en/about-sidnlabs

About SIDN Labs - examples

- ENTRADA
- DMAP
- LogoMotive
- RegCheck
- TimeNL
- IETF / RIPE / DNS-OARC / CENTR
- ICANN-studies

https://www.sidnlabs.nl/over-sidnlabs

Terminology

- Registrant wants domain name

Terminology

- Registrant wants domain name
- Goes to Registrar

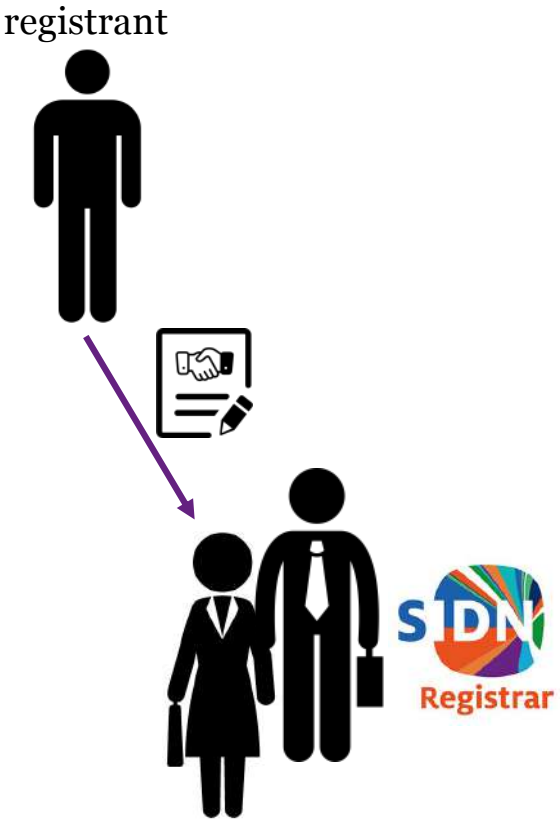

https://www.sidn.nl/nl-domeinnaam/registrar-zoeken

SIDN LABS

Terminology

- Registrant wants domain name
- Goes to Registrar
- (possibly via Reseller)

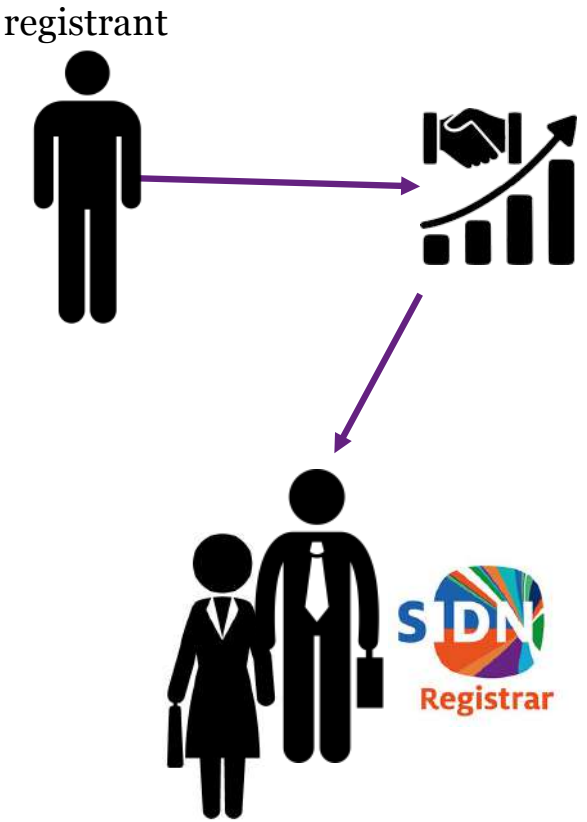

https://www.verenigingvanregistrars.nl/

Terminology

- Registrar is affiliated with Registry
- That's us ☺

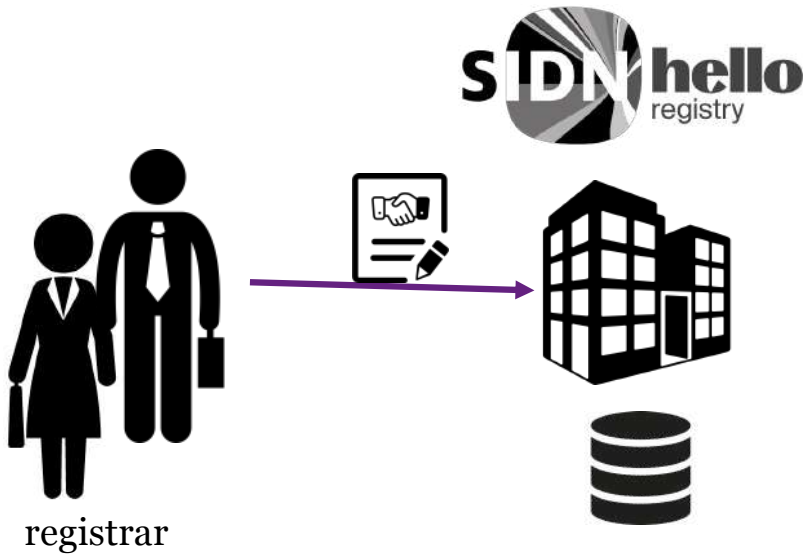

SIDN LABS

Terminology

- The domain name is registered!

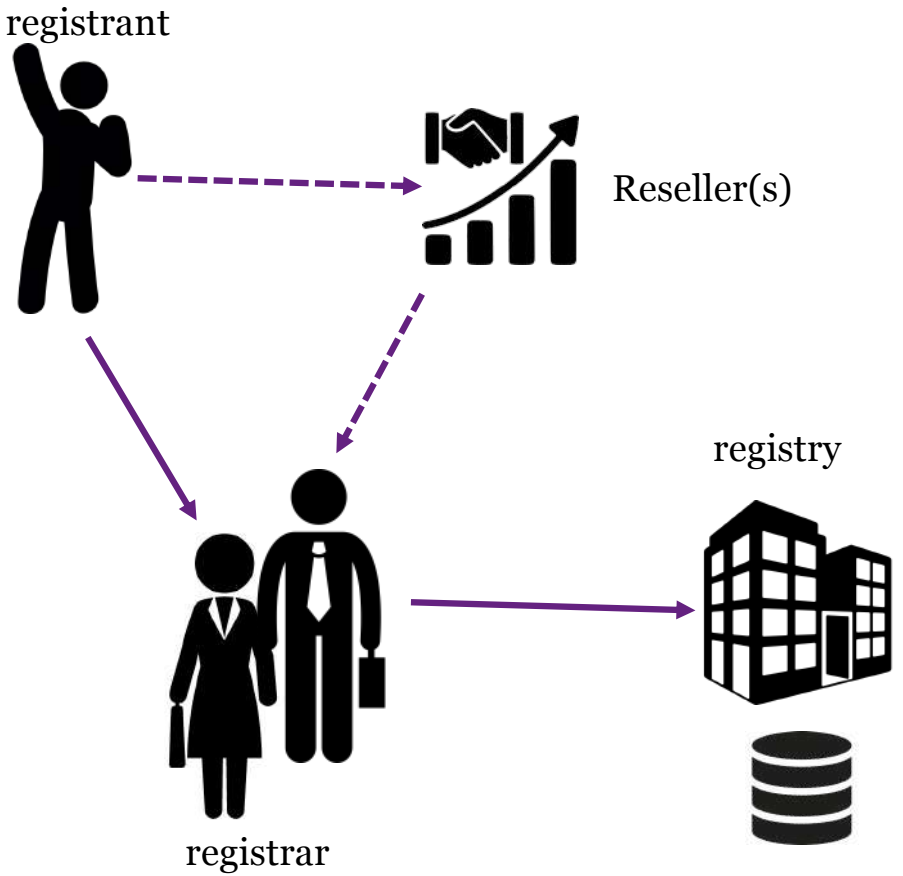

https://whois.nl/

Terminology

- 1 Registry (per TLD)
- ~1100 Registrars
- ??? Resellers
- ??? Unique registrants
- RRR model

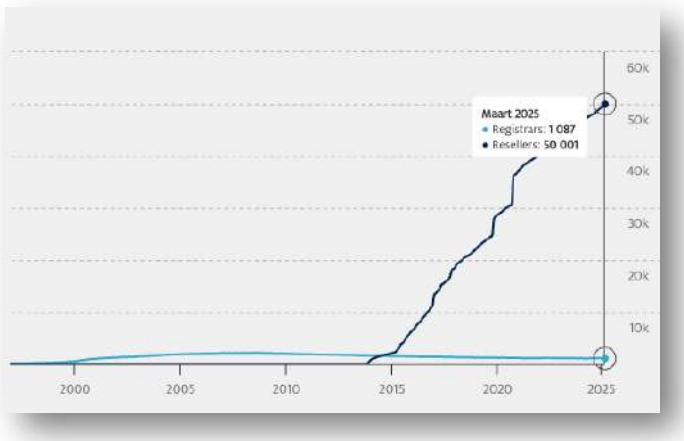

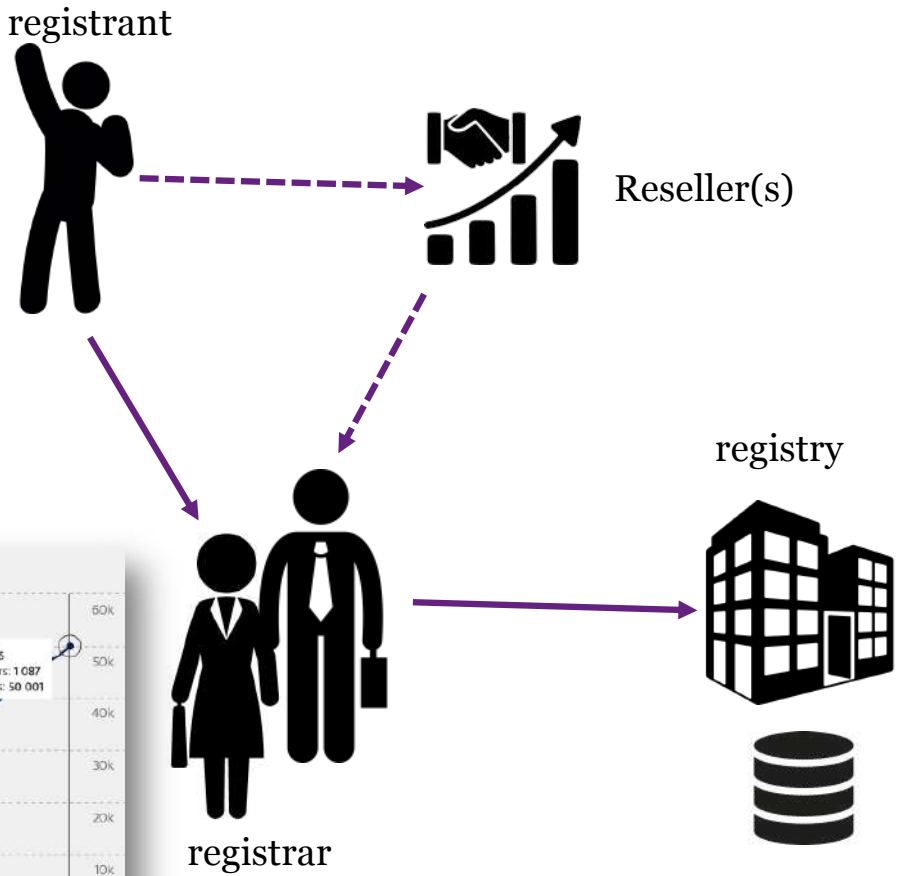

https://whois.nl/

Terminology

- The domain name is registered!

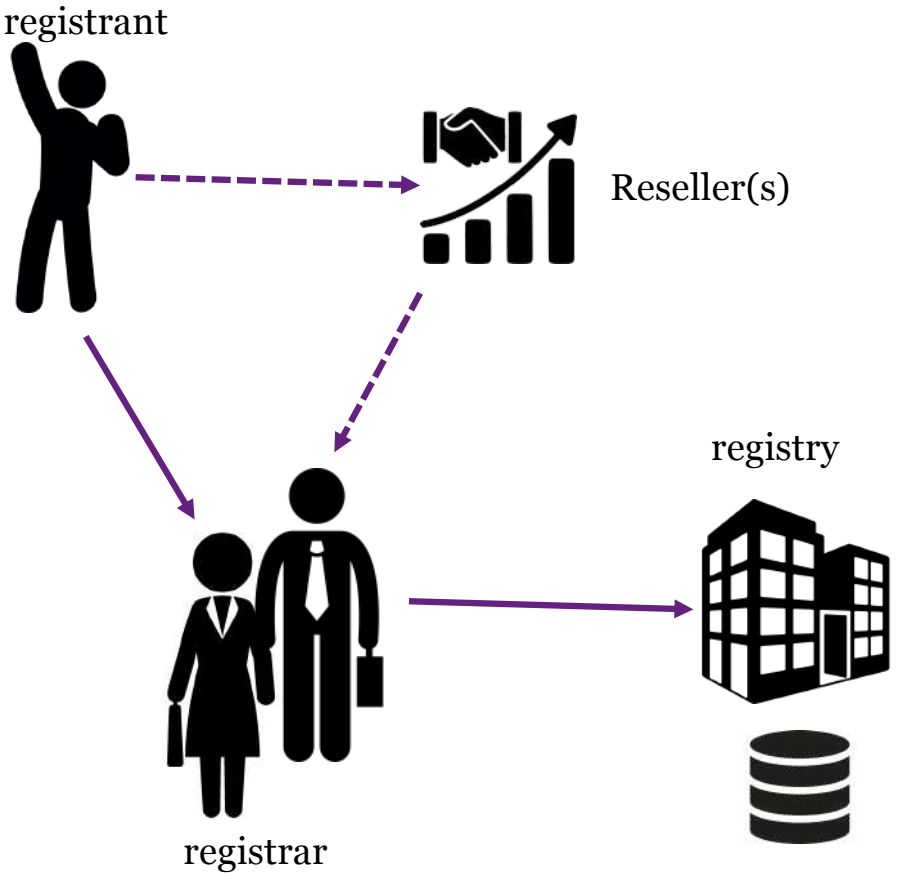

https://whois.nl/

Terminology

- The domain name is registered
- The domain name has been published!
- (only then does it work on the **internet**)

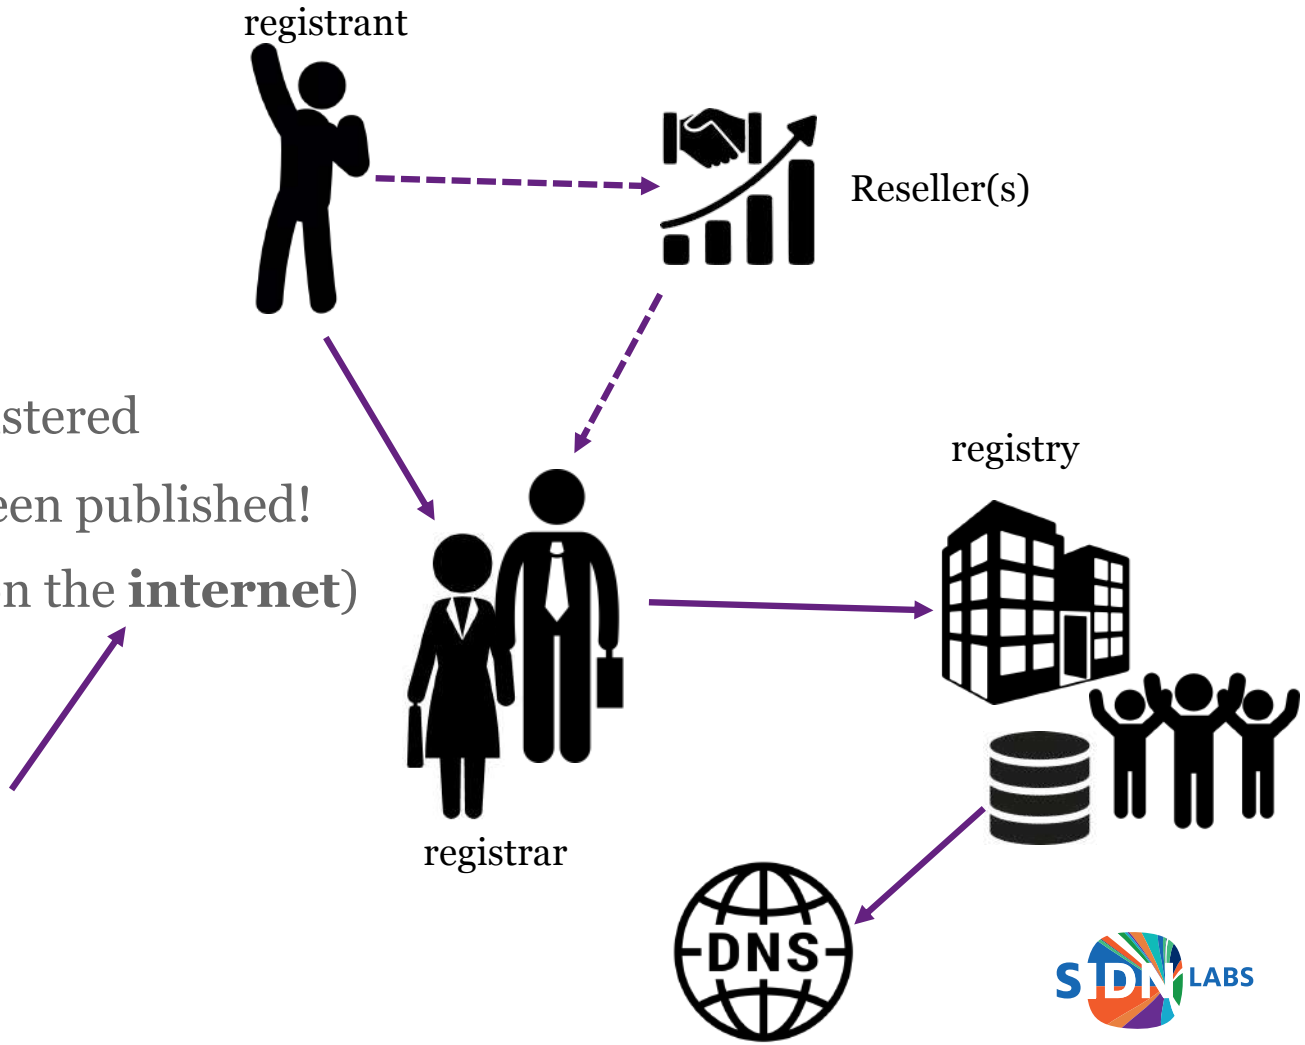

Important task!

The term 'internet'

What do we mean by that?

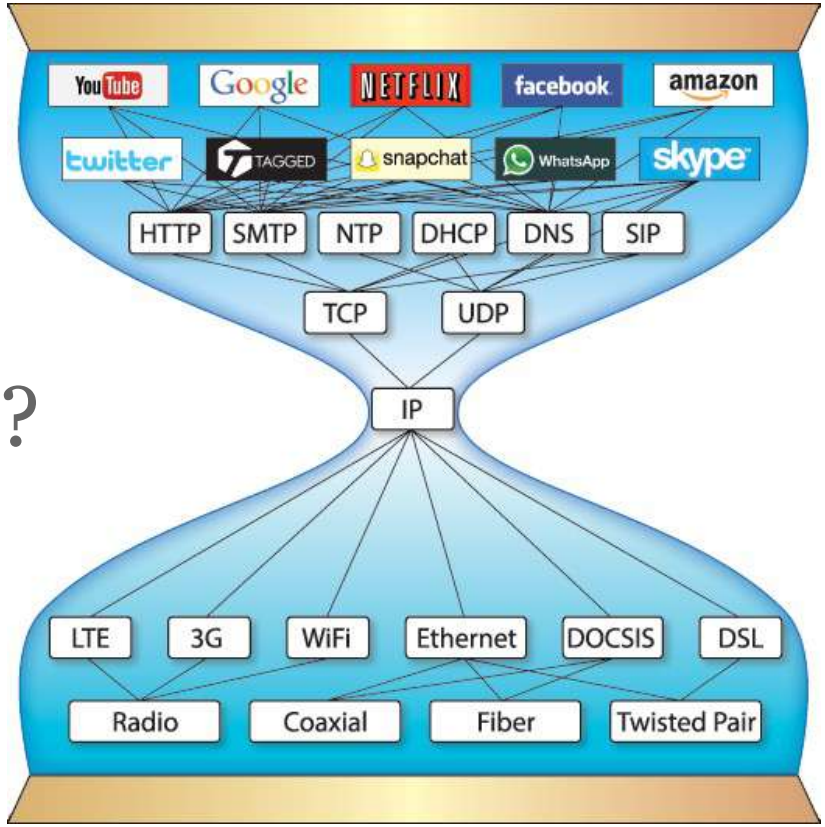

SDN LABS

The internet

Most people

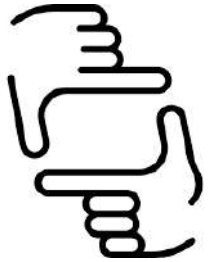

Me

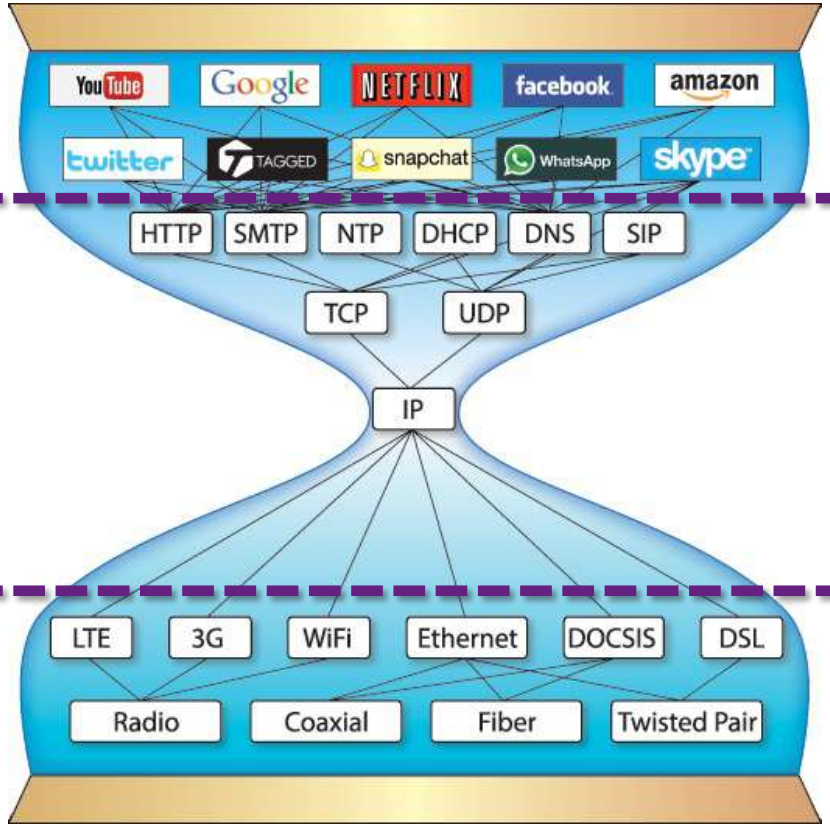

SDN LABS

The internet

Application (Presentation / session)

Transport

Internet (network)

Link (datalink)

(Physical)

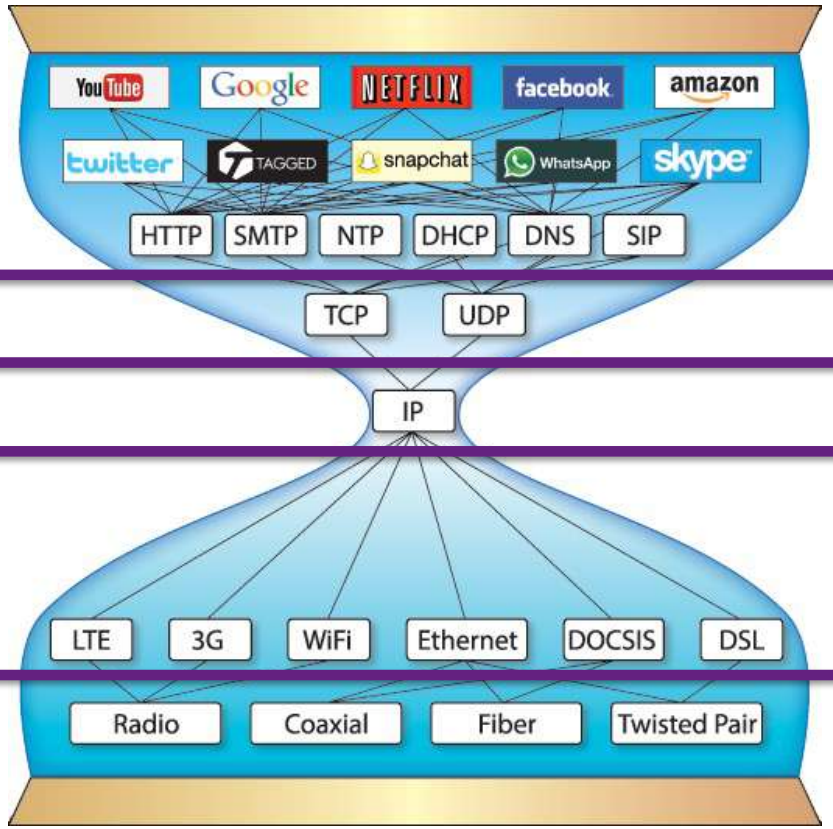

https://en.wikipedia.org/wiki/Internet_protocol_suite#Comparison_of_TCP/IP_and_OSI_layering

The internet

Fast

Slow!

Fast

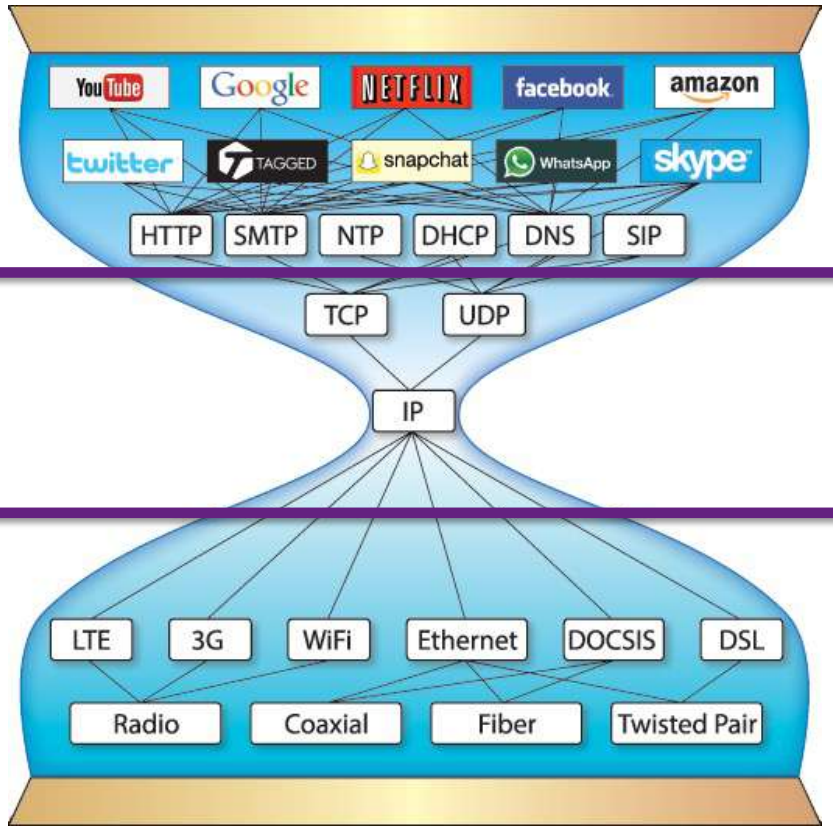

SDN LABS

The internet

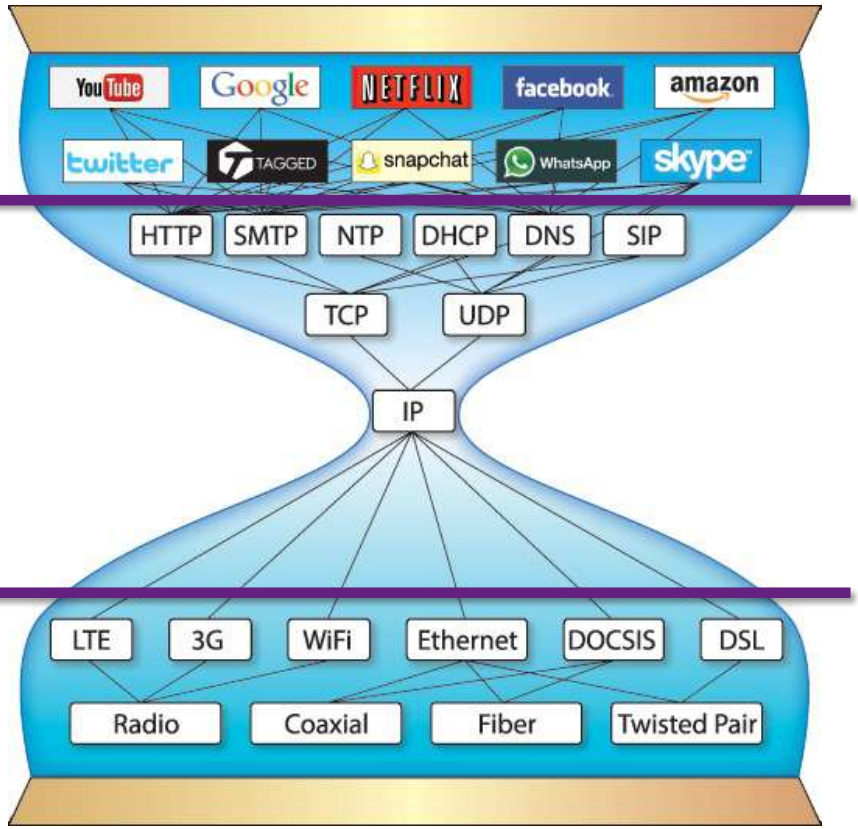

IETF®

(and others, such as W3C®)

IEEE®

(and others, such as Bluetooth®)

SDN LABS

Internet Standards

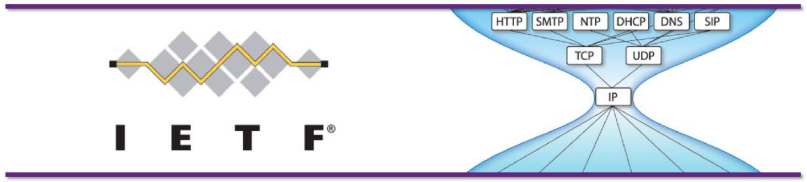

"We reject kings, presidents and voting.
We believe in rough consensus and running code"
-- David Clark

IETF, Internet Engineering Task Force:

- Open standards organization, without formal membership
- Everyone can participate (in person or via mailing lists)
- Under the auspices of the Internet Society (ISOC)
- Many working groups and informal discussions
- Rough consensus* is the primary basis for decision-making
- Often slow processes!
- But many RFCs ! About 9.800 and a multitude of drafts

* https://www.rfc-editor.org/info/rfc7282

IETF: bottom-up standards development

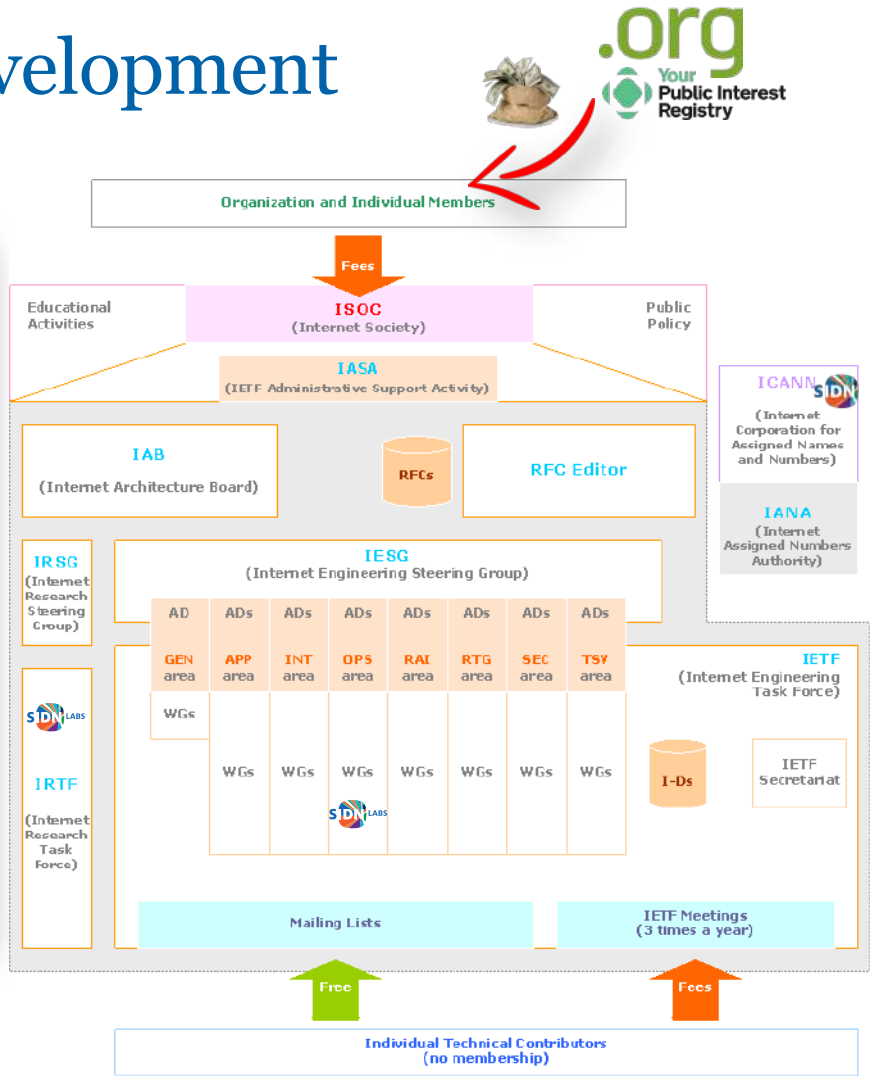

IETF: many RFC's

The mission of the IETF is to make the Internet work better by producing high quality, relevant technical documents that influence the way people design, use, and manage the Internet.

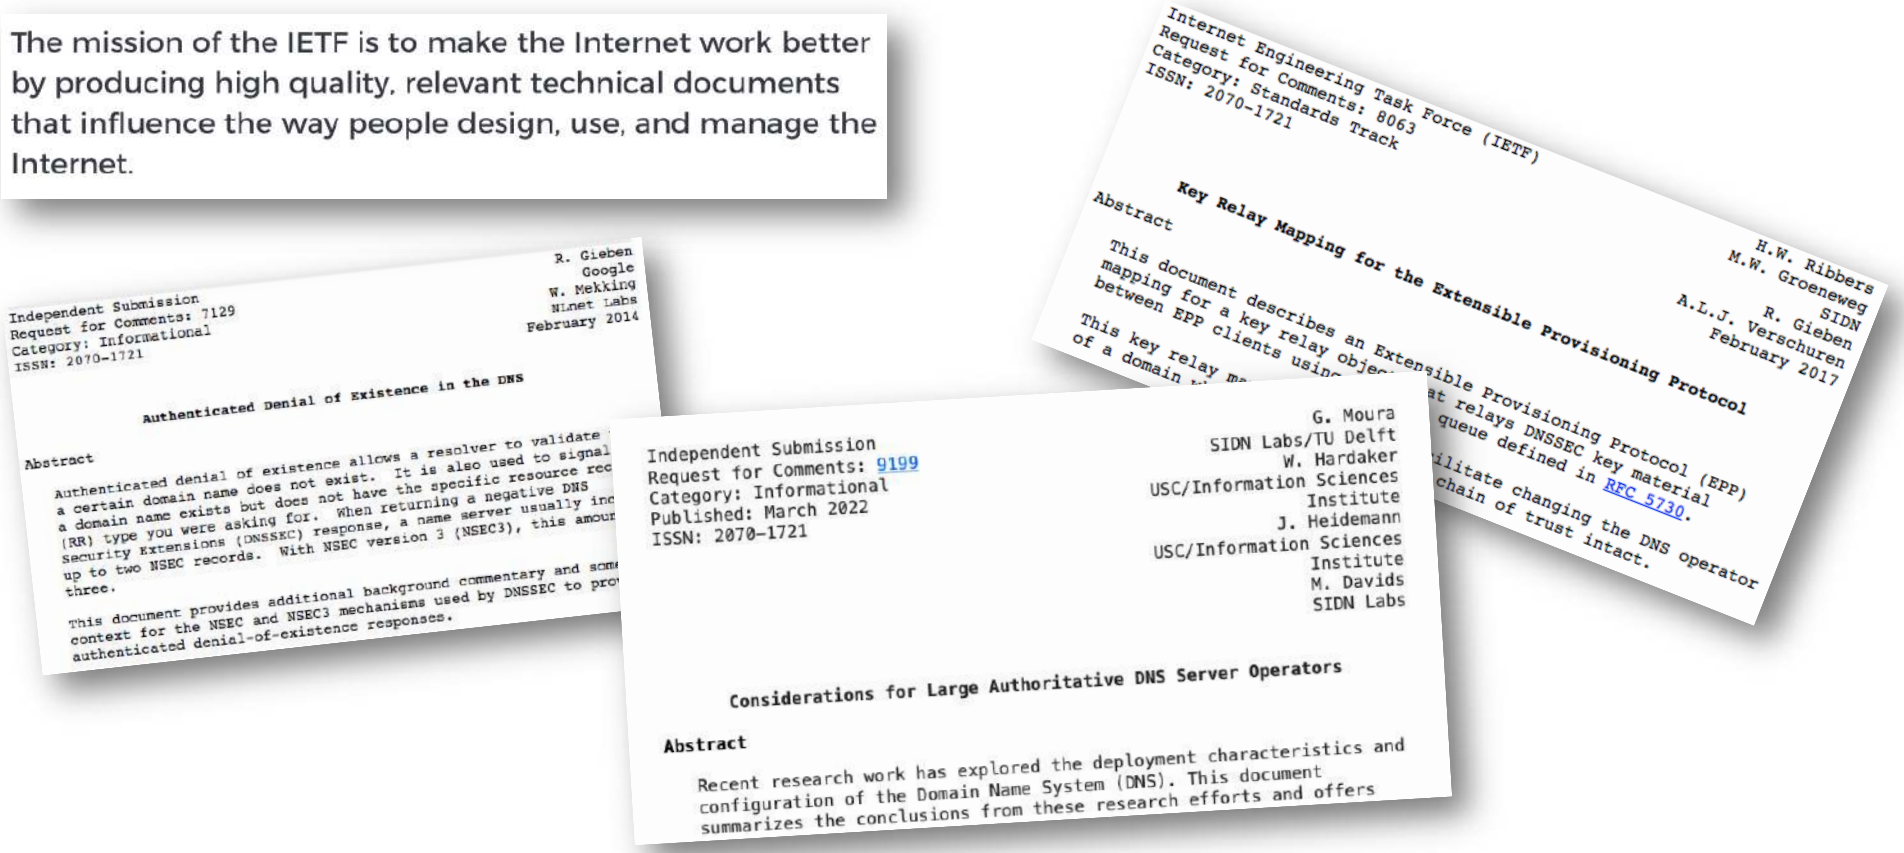

Independent Submission
Request for Comments: 7129
Category: Informational
ISSN: 2070-1721

R. Gieben
Google
W. Mekking
NLnet Labs
February 2014

Authenticated Denial of Existence in the DNS

Abstract

Authenticated denial of existence allows a resolver to validate a certain domain name does not exist. It is also used to signal a domain name exists but does not have the specific resource record (RR) type you were asking for. When returning a negative DNS Security Extensions (DNSSEC) response, a name server usually includes up to two NSEC records. With NSEC version 3 (NSEC3), this amount three.

This document provides additional background commentary and some context for the NSEC and NSEC3 mechanisms used by DNSSEC to provide authenticated denial-of-existence responses.

Internet Engineering Task Force (IETF)
Request for Comments: 8063
Category: Standards Track
ISSN: 2070-1721

H.W. Ribbers
M.W. Groeneweg
SIDN
R. Gieben
A.L.J. Verschuren
February 2017

Key Relay Mapping for the Extensible Provisioning Protocol

Abstract

This document describes an Extensible Provisioning Protocol (EPP) mapping for a key relay object that relays DNSSEC key material between EPP clients using the queue defined in RFC 5730.

This key relay mapping facilitate changing the DNS operator of a domain while keeping the chain of trust intact.

Independent Submission
Request for Comments: 9199
Category: Informational
Published: March 2022
ISSN: 2070-1721

G. Moura
SIDN Labs/TU Delft
W. Hardaker
USC/Information Sciences Institute
J. Heidemann
USC/Information Sciences Institute
M. Davids
SIDN Labs

Considerations for Large Authoritative DNS Server Operators

Abstract

Recent research work has explored the deployment characteristics and configuration of the Domain Name System (DNS). This document summarizes the conclusions from these research efforts and offers

🇳🇱 👍🏽 http://www.arkko.com/tools/allstats/thenetherlands.html

SIDN LABS

Really a lot, maybe even too much…

https://emaillab.jp/wp/wp-content/uploads/2017/11/RFC-DNS.pdf

246 RFC's
4156 pages
~245.000 lines
~2.206.900 words

This is ~4 times "The C++ Programming Language" (4th ed.)

Some good words on this are in RFC 8324

https://datatracker.ietf.org/meeting/101/materials/slides-101-dnsop-sessa-the-dns-camel-01

DNS(SEC)

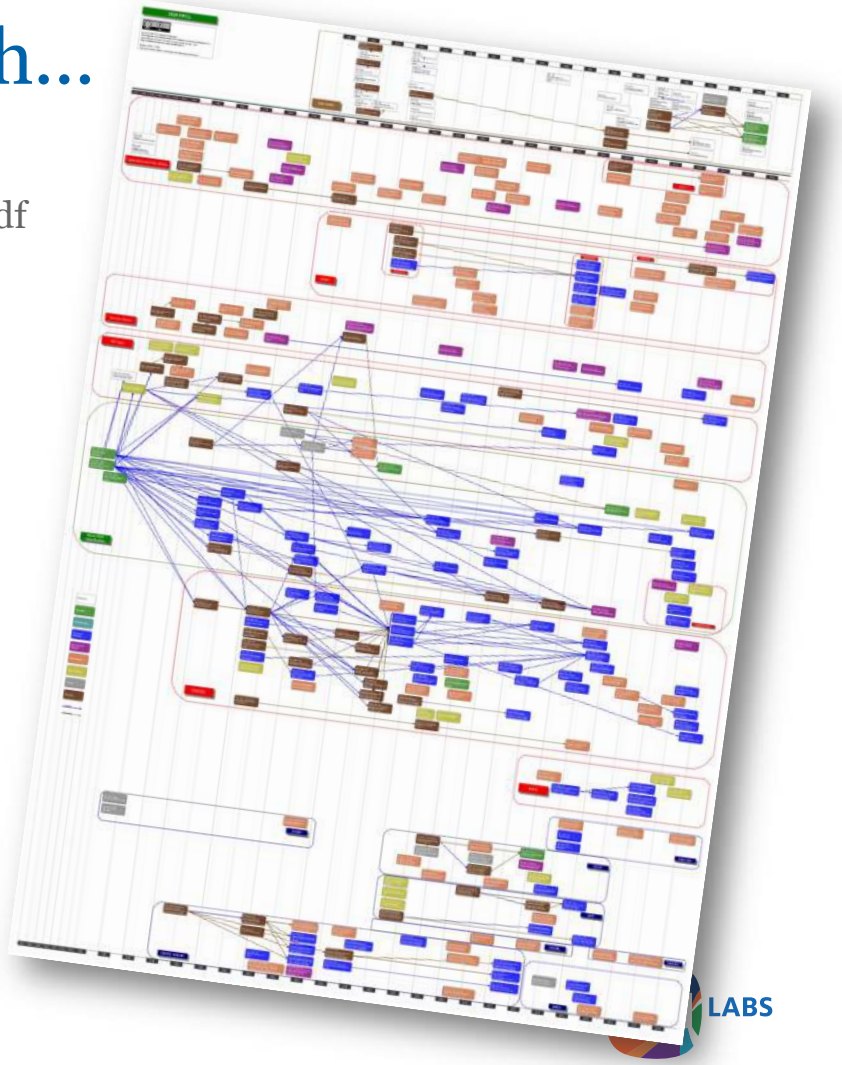

See also: https://powerdns.org/dns-camel/

Names and numbers

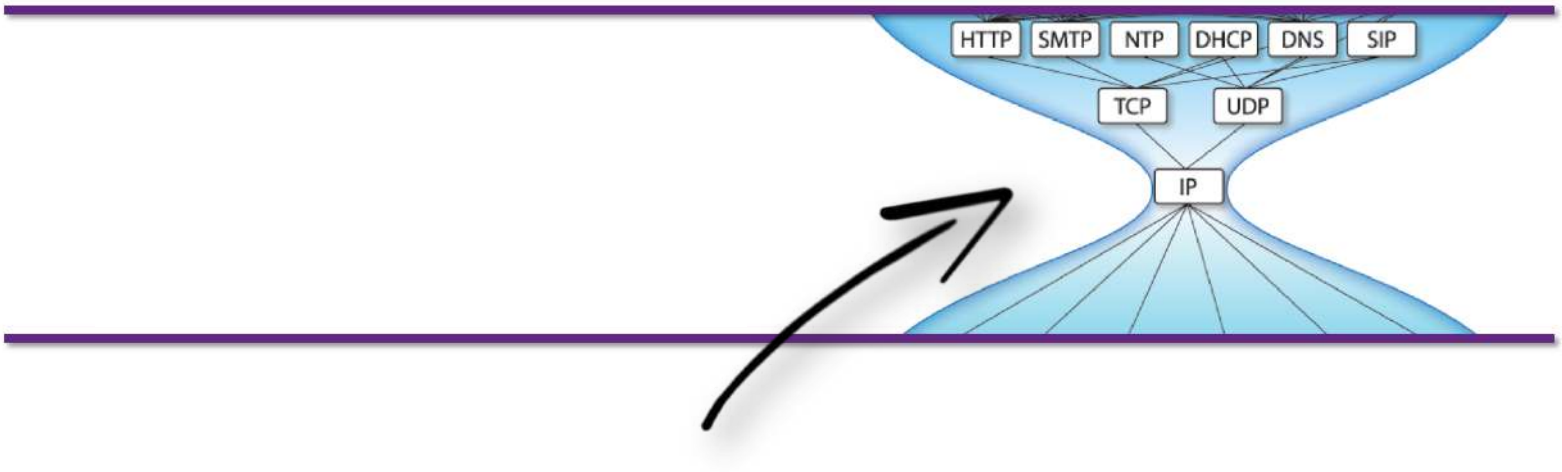

Names and numbers

Do you recognize these?

192.168.0.1

192.168.1.1

192.168.100.1

192.168.2.1

https://www.rfc-editor.org/info/rfc1918

Names and number

198.51.100.123

A 'global' IP address

But this one?

2001:db8::198:51:100:123

Names and numbers

Good to know:

*Every device connected **directly** to the Internet requires a unique* IP address.*

* exception: anycast

Control of IP address space issuance

The mission of ICANN is to ensure the stable and secure operation of the Internet's unique identifier systems

ICANN (the Internet Corporation for Assigned Names and Numbers)

https://www.icann.org/resources/pages/governance/bylaws-en/

Control of IP address space issuance (and more)

DNS space

AS numbers

IP space

The mission of ICANN is to ensure the stable and secure operation of the Internet's unique identifier systems

'Global number registries'

ICANN (the Internet Corporation for Assigned Names and Numbers)

https://www.icann.org/

Issuance of IP address space

https://www.iana.org/assignments/ipv4-address-space/
https://www.iana.org/assignments/ipv6-address-space/

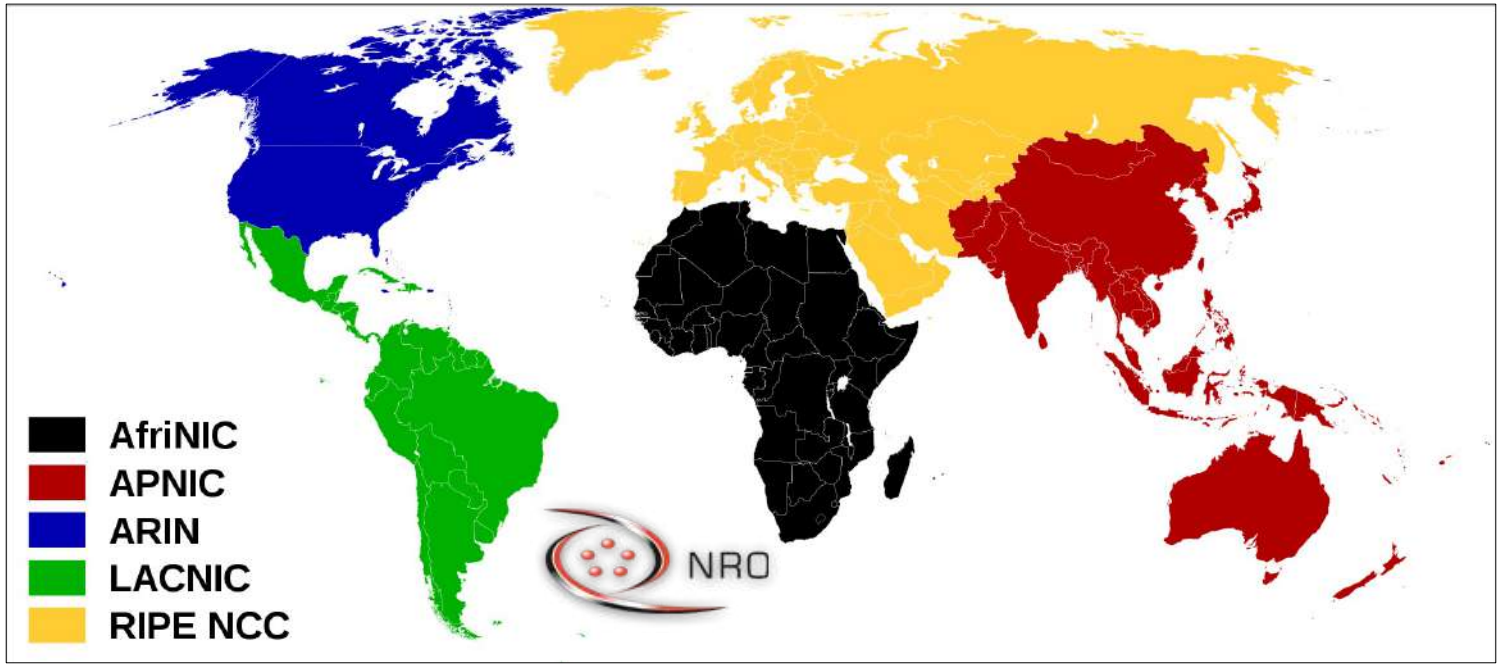

IANA (Internet Assigned Numbers Authority) → RIRs → LIRs

https://www.nro.net/

But ICANN does more (top level domains)

https://www.iana.org/domains/root/db

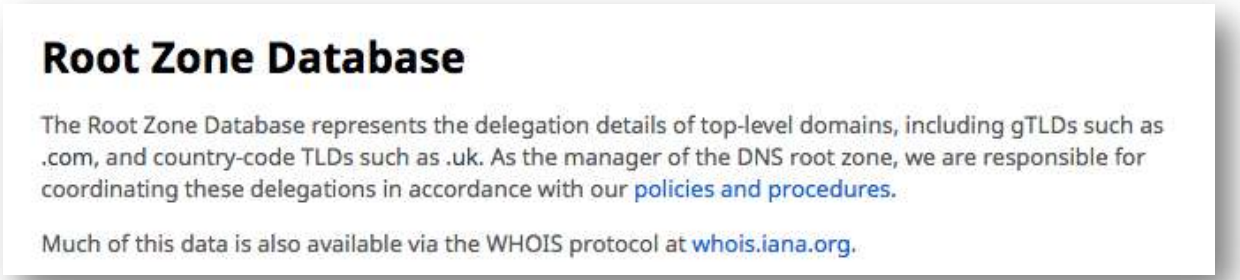

Root Zone Database

The Root Zone Database represents the delegation details of top-level domains, including gTLDs such as .com, and country-code TLDs such as .uk. As the manager of the DNS root zone, we are responsible for coordinating these delegations in accordance with our policies and procedures.

Much of this data is also available via the WHOIS protocol at whois.iana.org.

iana

Internet Assigned Numbers Authority

https://www.iana.org/

And ICANN does even more (protocol assignments)

https://www.iana.org/assignments/dns-parameters/dns-parameters.xhtml#dns-parameters-12

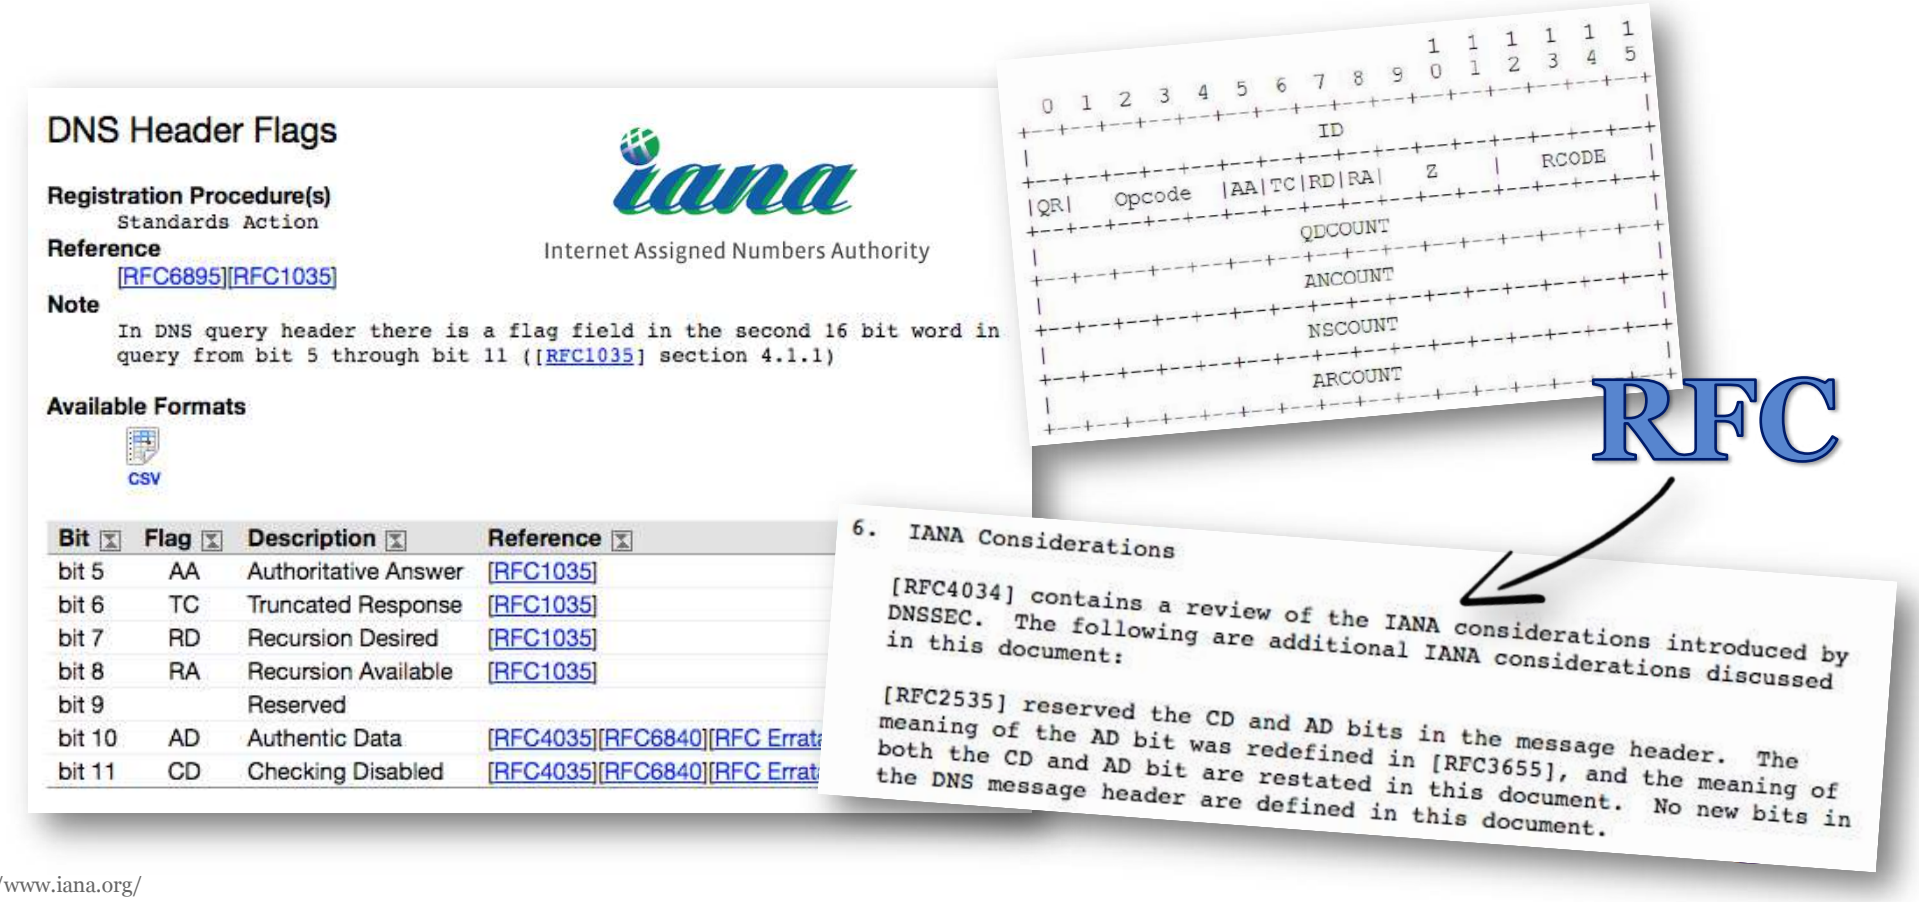

DNS Header Flags

Registration Procedure(s)
Standards Action

Reference
[RFC6895][RFC1035]

Note
In DNS query header there is a flag field in the second 16 bit word in query from bit 5 through bit 11 ([RFC1035] section 4.1.1)

Internet Assigned Numbers Authority

Available Formats

CSV

Bit	Flag	Description	Reference
bit 5	AA	Authoritative Answer	[RFC1035]
bit 6	TC	Truncated Response	[RFC1035]
bit 7	RD	Recursion Desired	[RFC1035]
bit 8	RA	Recursion Available	[RFC1035]
bit 9		Reserved	
bit 10	AD	Authentic Data	[RFC4035][RFC6840][RFC Errata
bit 11	CD	Checking Disabled	[RFC4035][RFC6840][RFC Errata

RFC

6. IANA Considerations

[RFC4034] contains a review of the IANA considerations introduced by DNSSEC. The following are additional IANA considerations discussed in this document:

[RFC2535] reserved the CD and AD bits in the message header. The meaning of the AD bit was redefined in [RFC3655], and the meaning of both the CD and AD bit are restated in this document. No new bits in the DNS message header are defined in this document.

https://www.iana.org/

Overview ICANN / IANA

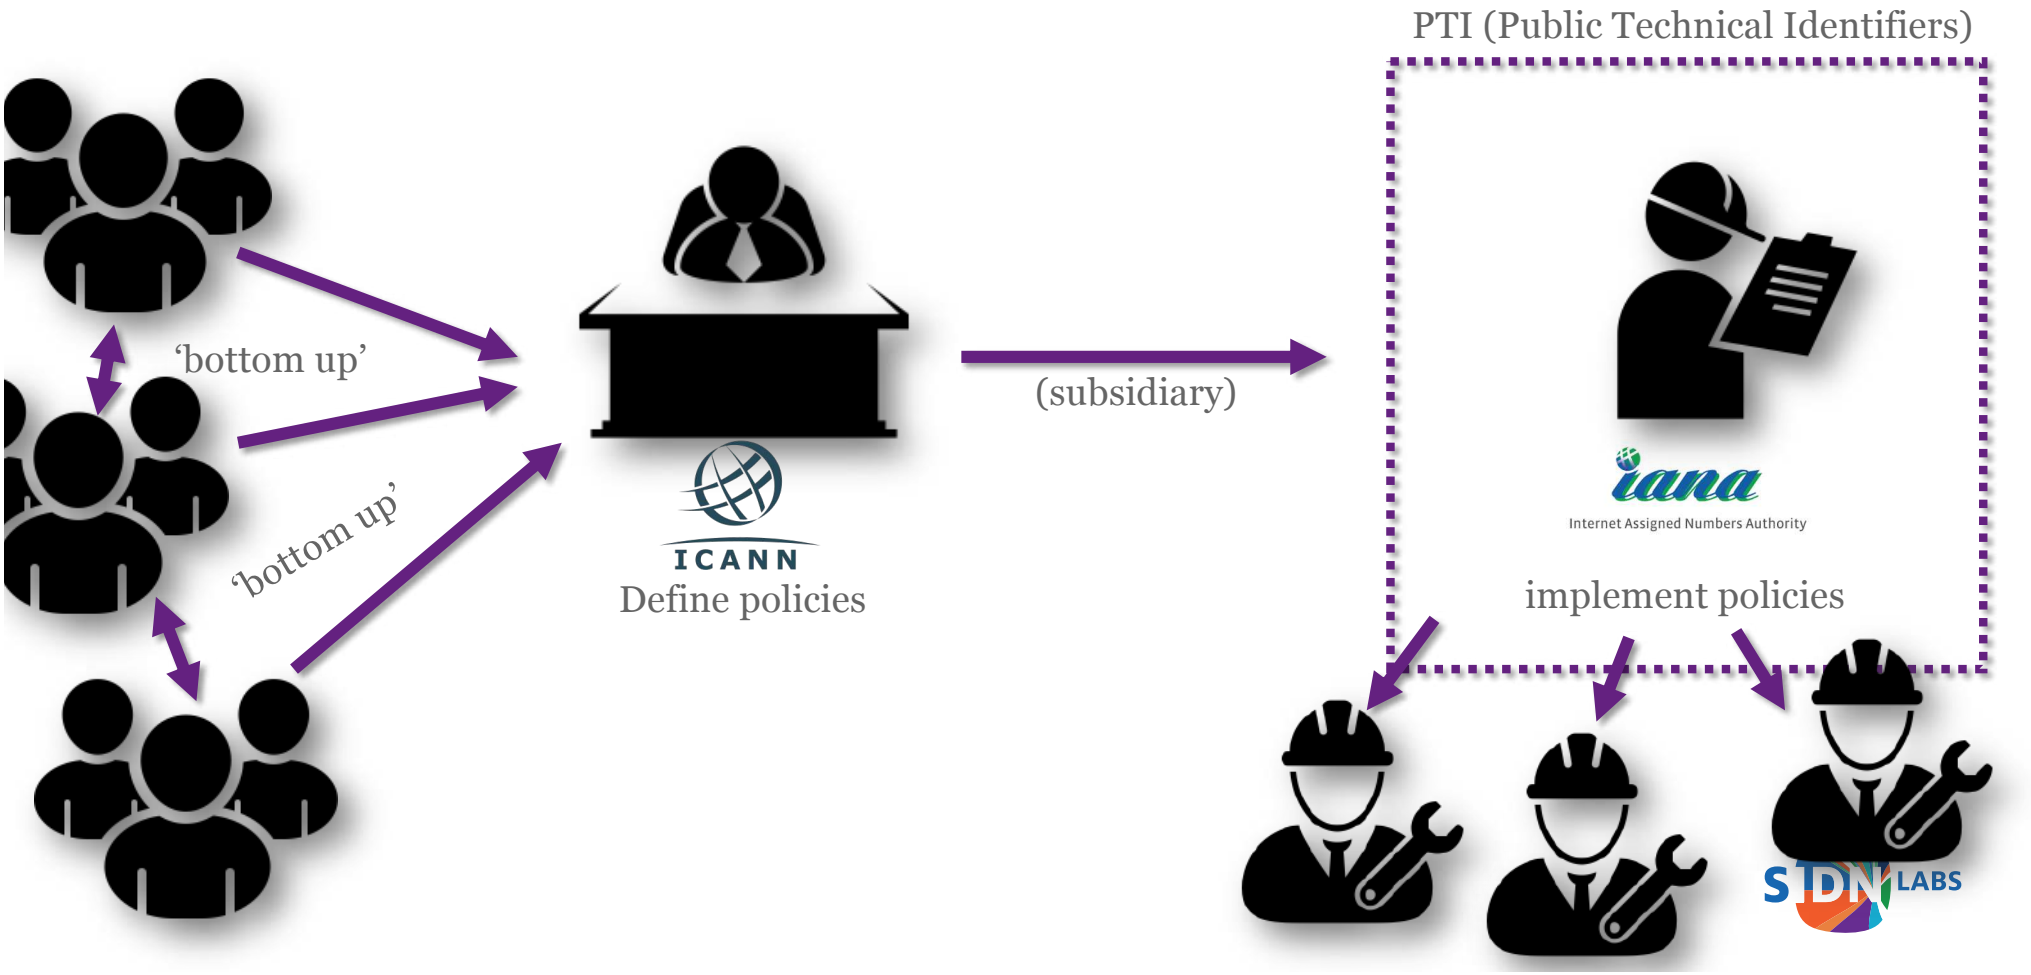

(back to) Names and numbers

www.example.nl

This will be familiar.

And what do we see here?

www.example.team

(back to) Names and numbers

www.example.испытание

What about this one?

www.日本レジストリサービス.jp

(back to) Names and numbers

www.bücher.de

Allowed or…?

www.café.nl

www.café.be

Internationalized domain name (IDN)

www.bücher.de

www.xn--bcher-kva.de

IDNA / punycode / ACE-string

www.café.be

www.xn--caf-dma.be

Reading tip: https://en.wikipedia.org/wiki/IDN_homograph_attack

apple vs apple

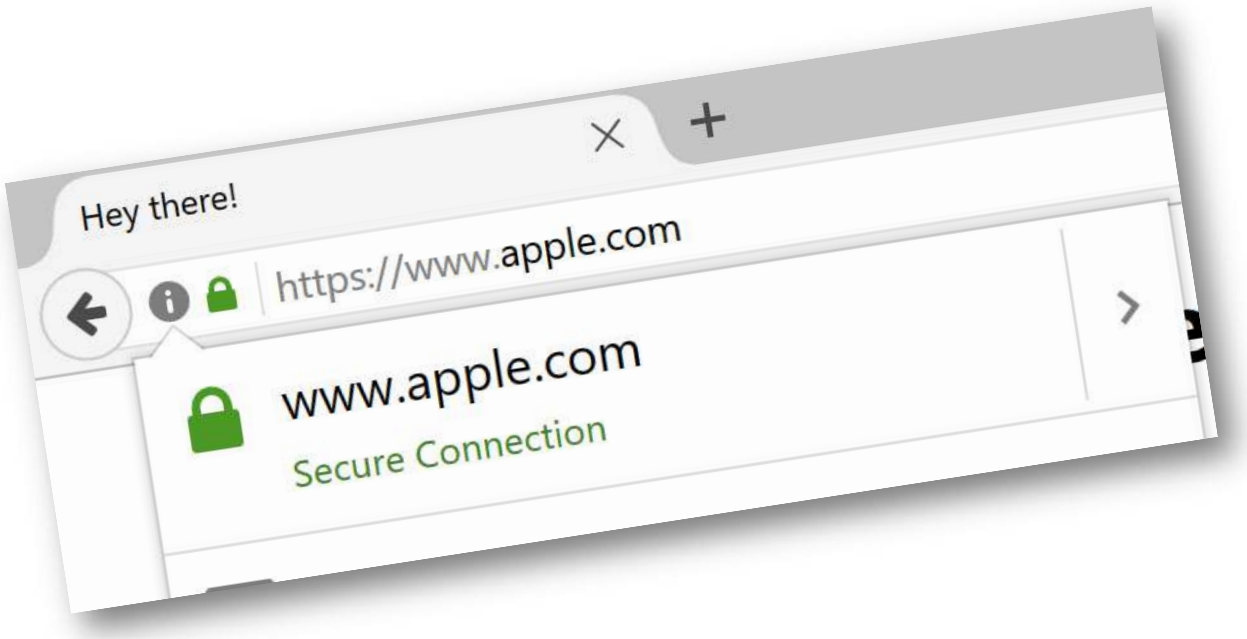

Reading tip: https://www.xudongz.com/blog/2017/idn-phishing/

apple vs apple

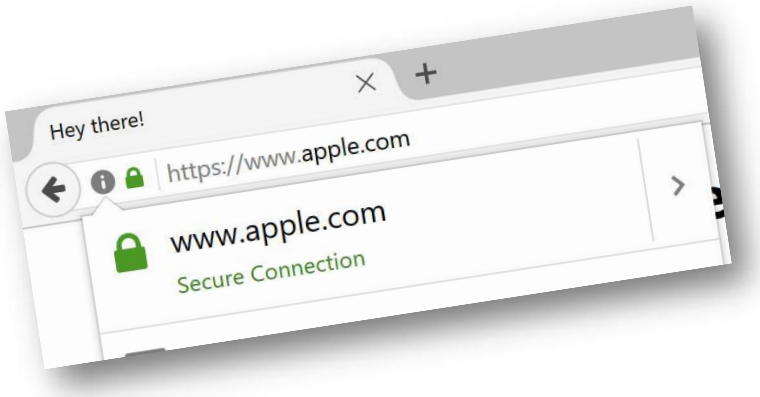


```
package main

import "fmt"

func main() {
        for _, value := range "apple" {
                fmt.Println(string(value), value)
        }
        fmt.Println()
        for _, value := range "appIe" {
                fmt.Println(string(value), value)
        }
}
```

Question:
How many characters of apple and apple are different?

Reading tip: https://www.xudongz.com/blog/2017/idn-phishing/

apple vs apple

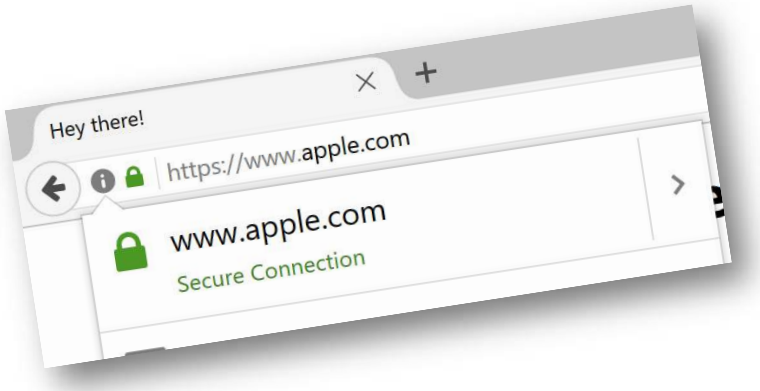


```
package main

import "fmt"

func main() {
        for _, value := range "apple" {
                fmt.Println(string(value), value)
        }
        fmt.Println()
        for _, value := range "appIe" {
                fmt.Println(string(value), value)
        }
}
```

```
a 97
p 112
p 112
l 108
e 101
```

```
a 1072
p 1088
p 1088
I 1231
e 1077
```

Reading tip: https://www.xudongz.com/blog/2017/idn-phishing/

apple vs apple

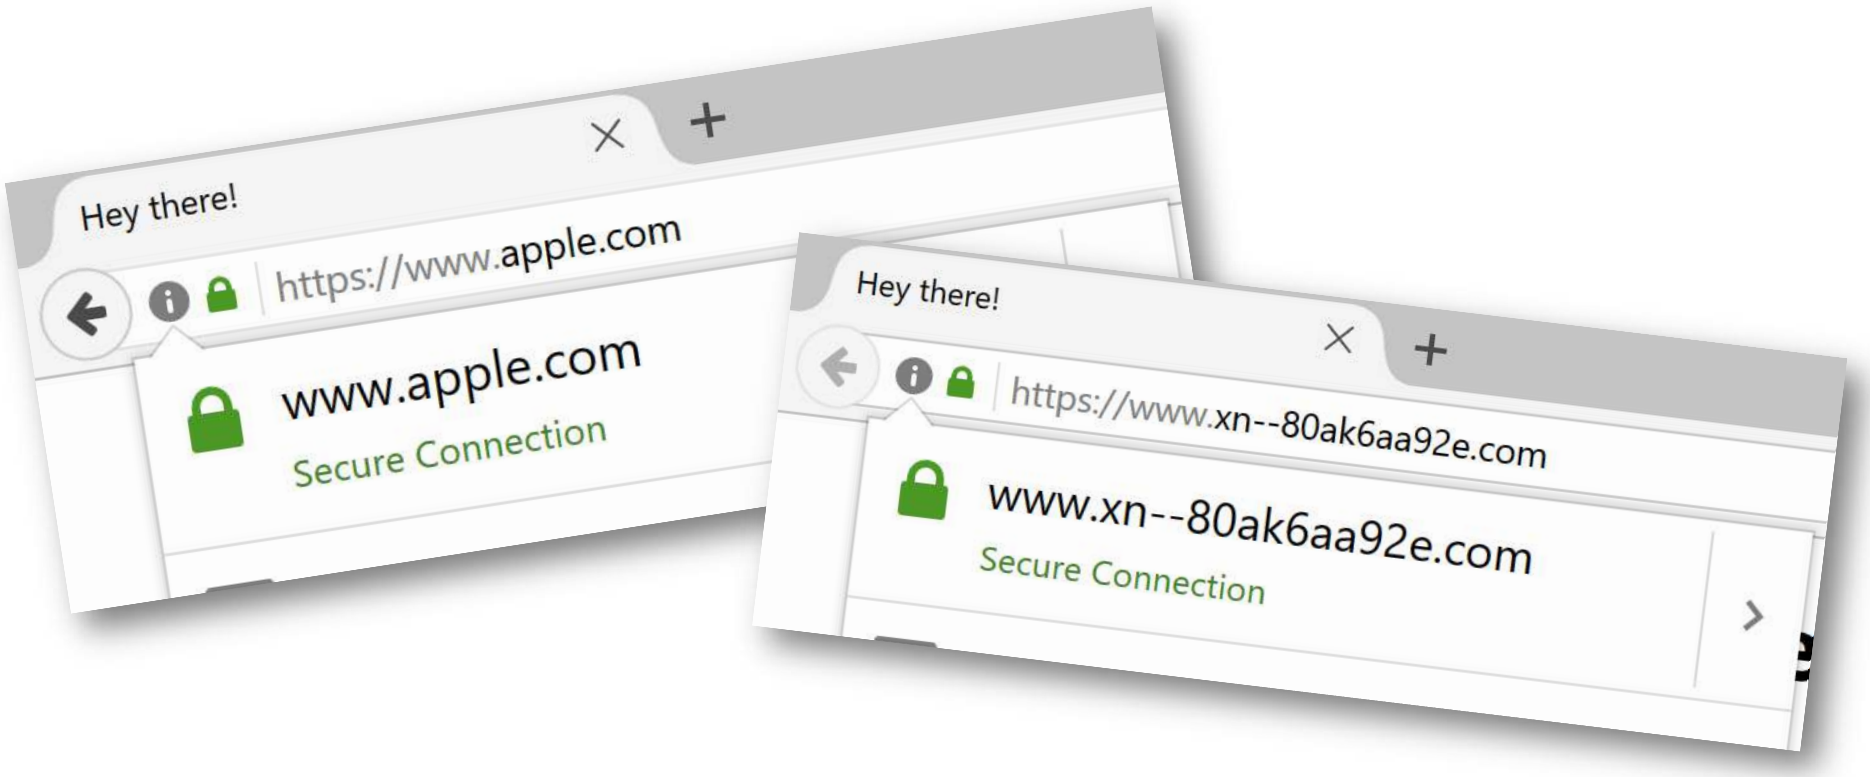

Reading tip: https://www.xudongz.com/blog/2017/idn-phishing/

apple vs apple

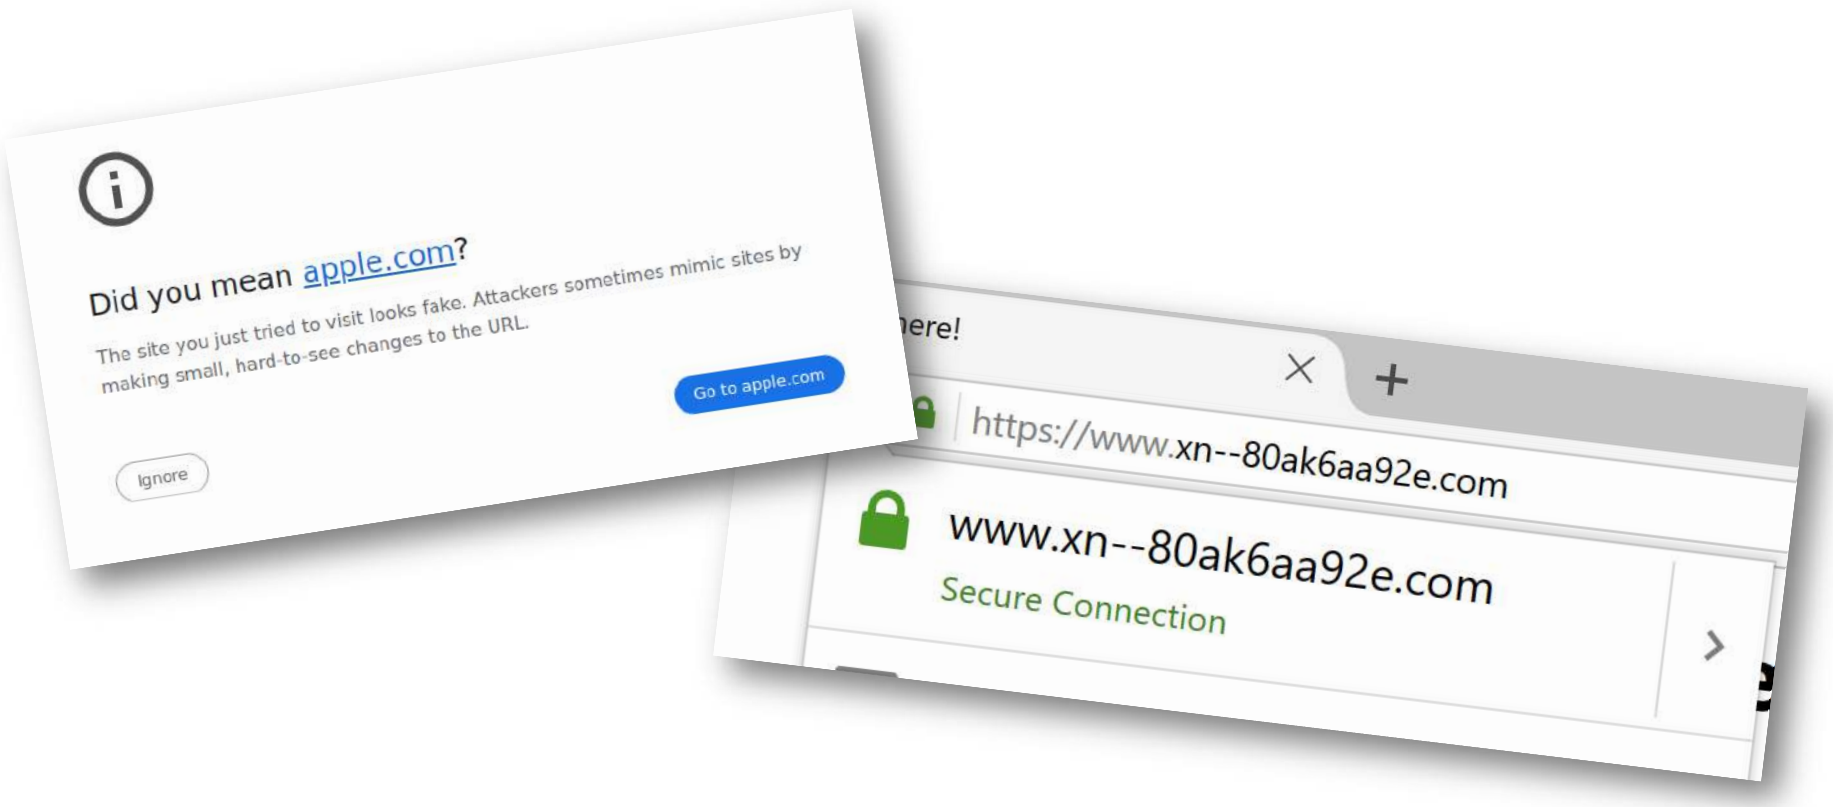

Reading tip: https://www.xudongz.com/blog/2017/idn-phishing/

Names ~~and numbers~~

And have you ever seen these ones?

marco-gw.home.arpa

darknetsite.onion

e164.arpa

eat.kiwi

Top-level domains

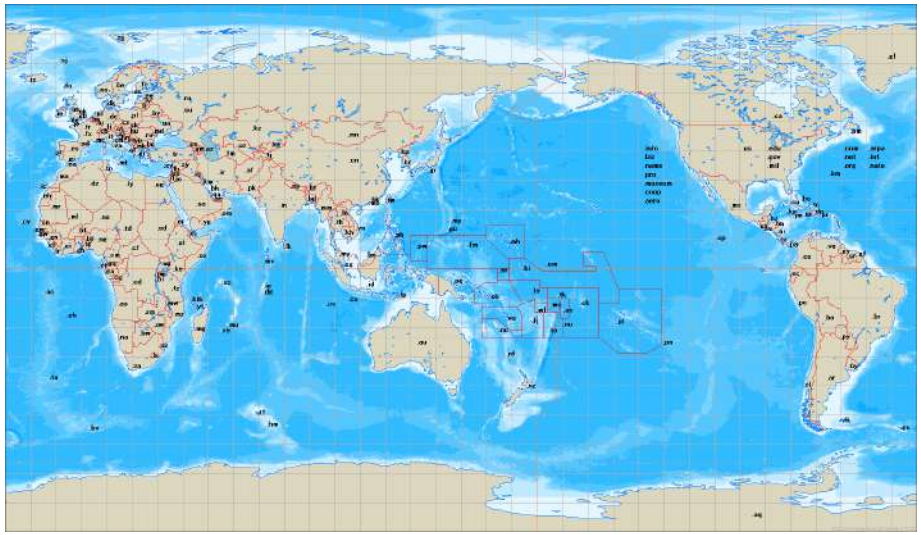

±308

±1241

http://www.marco.panizza.name/dispenseTM/slides/TLD/ccTLD_worldmap.html
https://en.wikipedia.org/wiki/Country_code_top-level_domain

https://newgtlds.icann.org/en/program-status/statistics

Domain Name System (DNS)

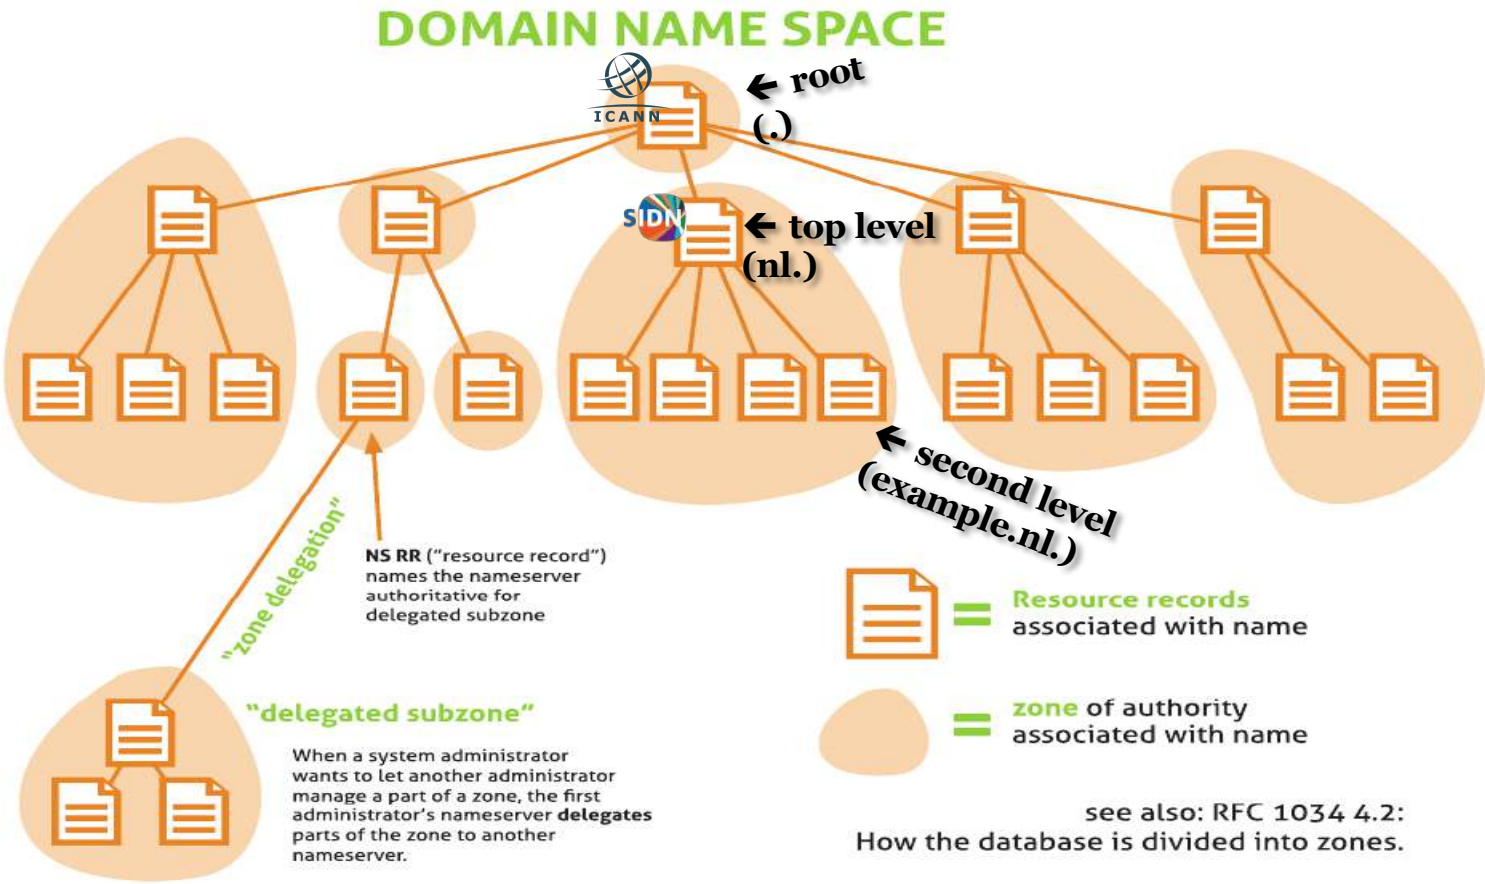

LABS

About DNS

- DNS is a kind of ‘internet signage’
- Some compare it with a telephone directory
- You know the name (sidn.nl), your browser can't do much with it
- Your browser wants the IP address (2600:1901:0:7947::0)
- DNS makes this possible

How do you get from blank page to website?

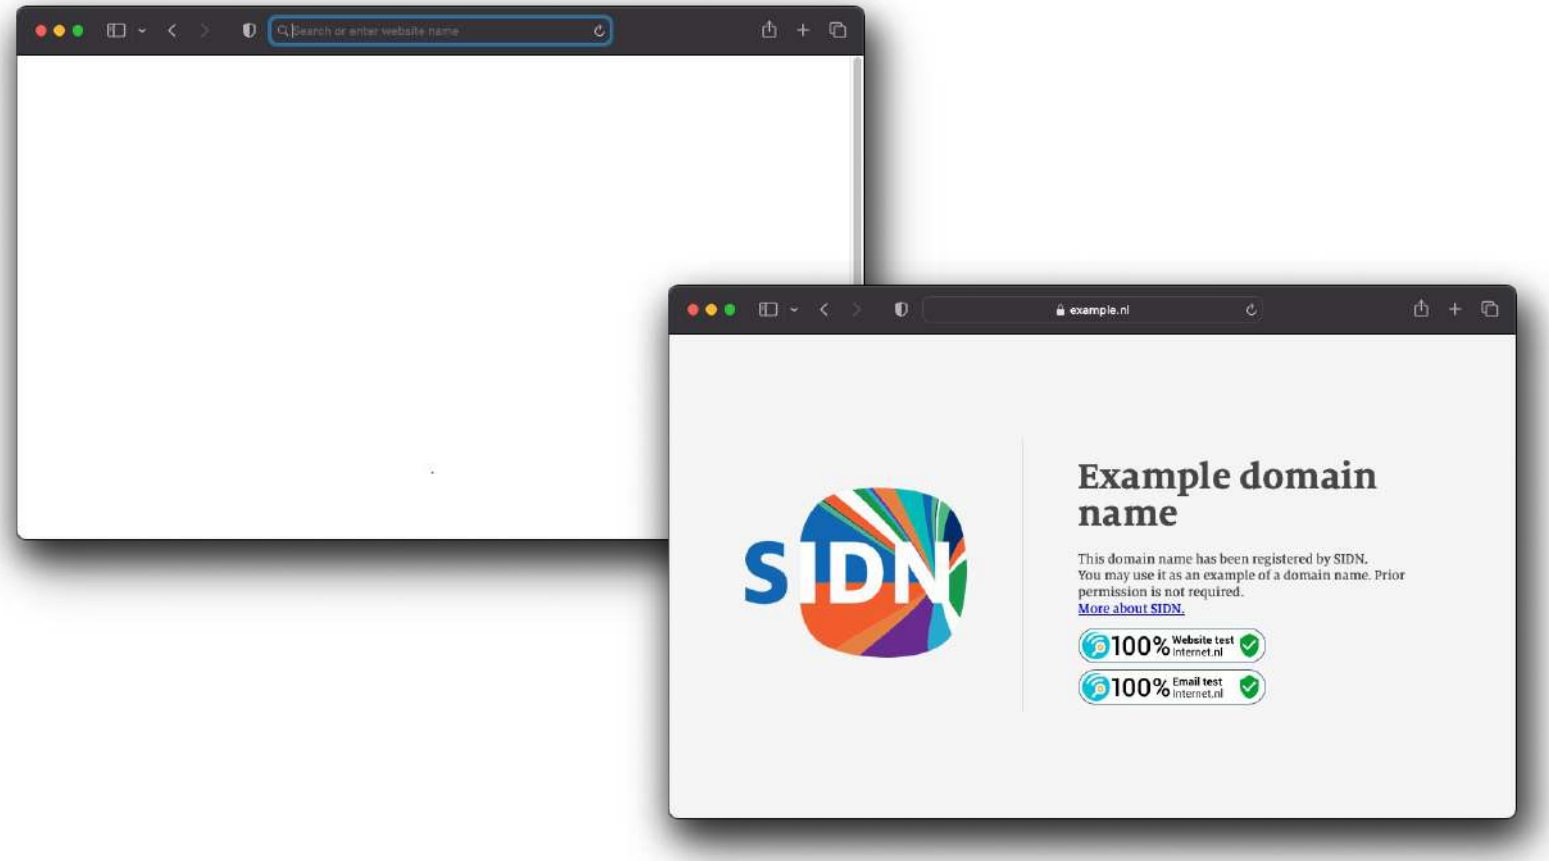

DNS: What happens when you surf to http://example.nl?

Local user

Resolver

Autoritatieve name servers

DNS: The 'stub-resolver' asks a question to the 'resolver'...

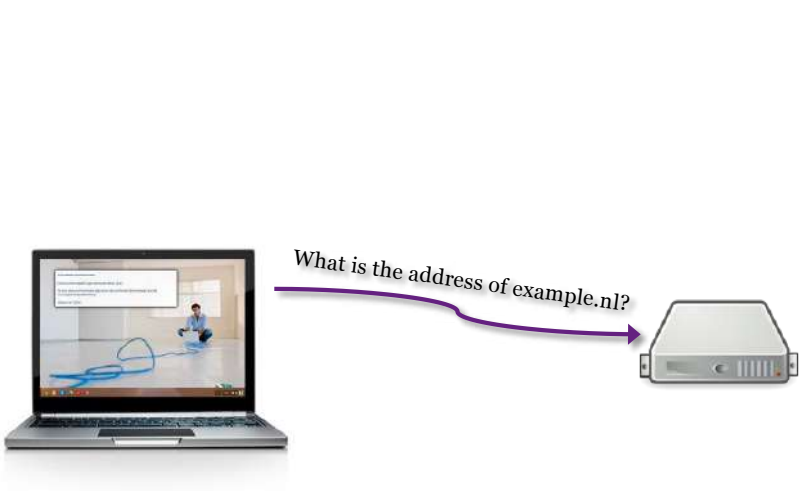

DNS: The resolver starts looking for the answer…

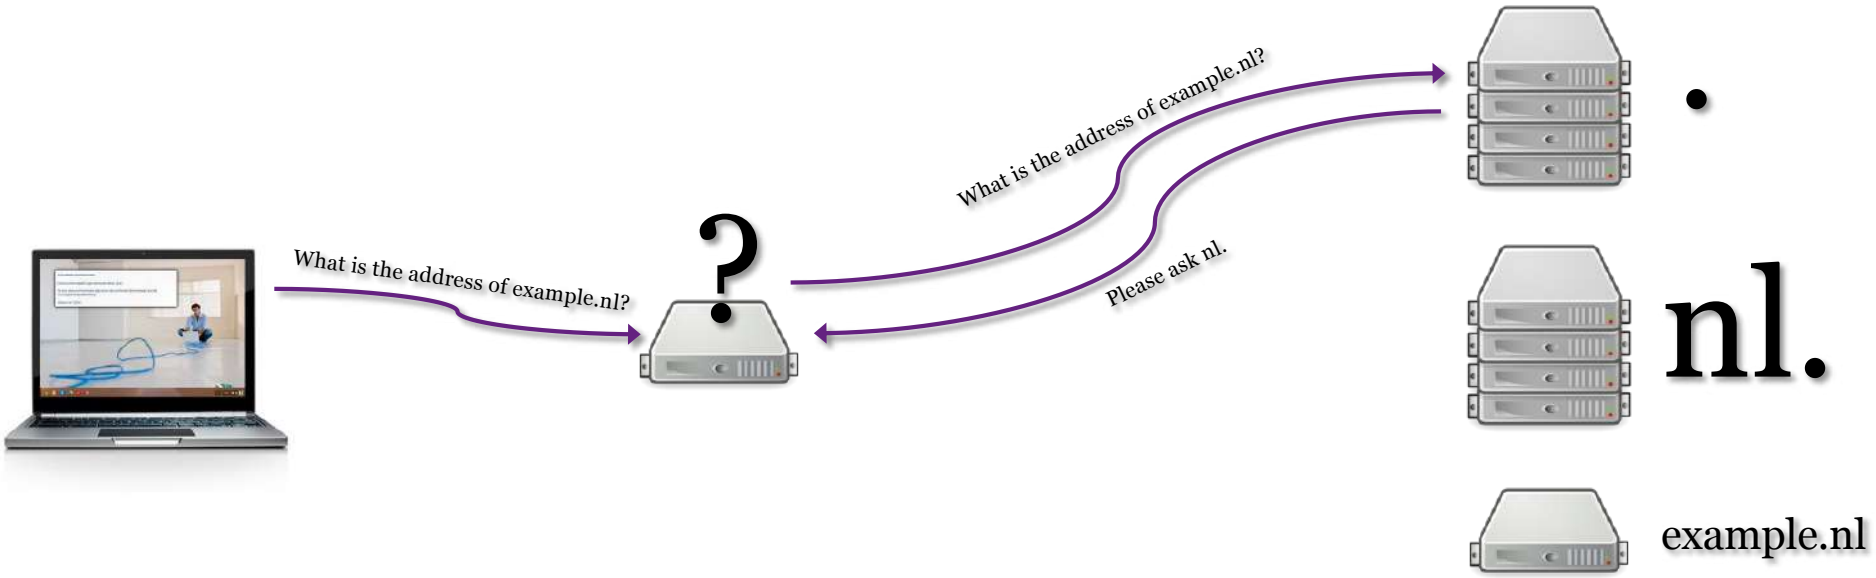

DNS: The resolver starts looking for the answer...

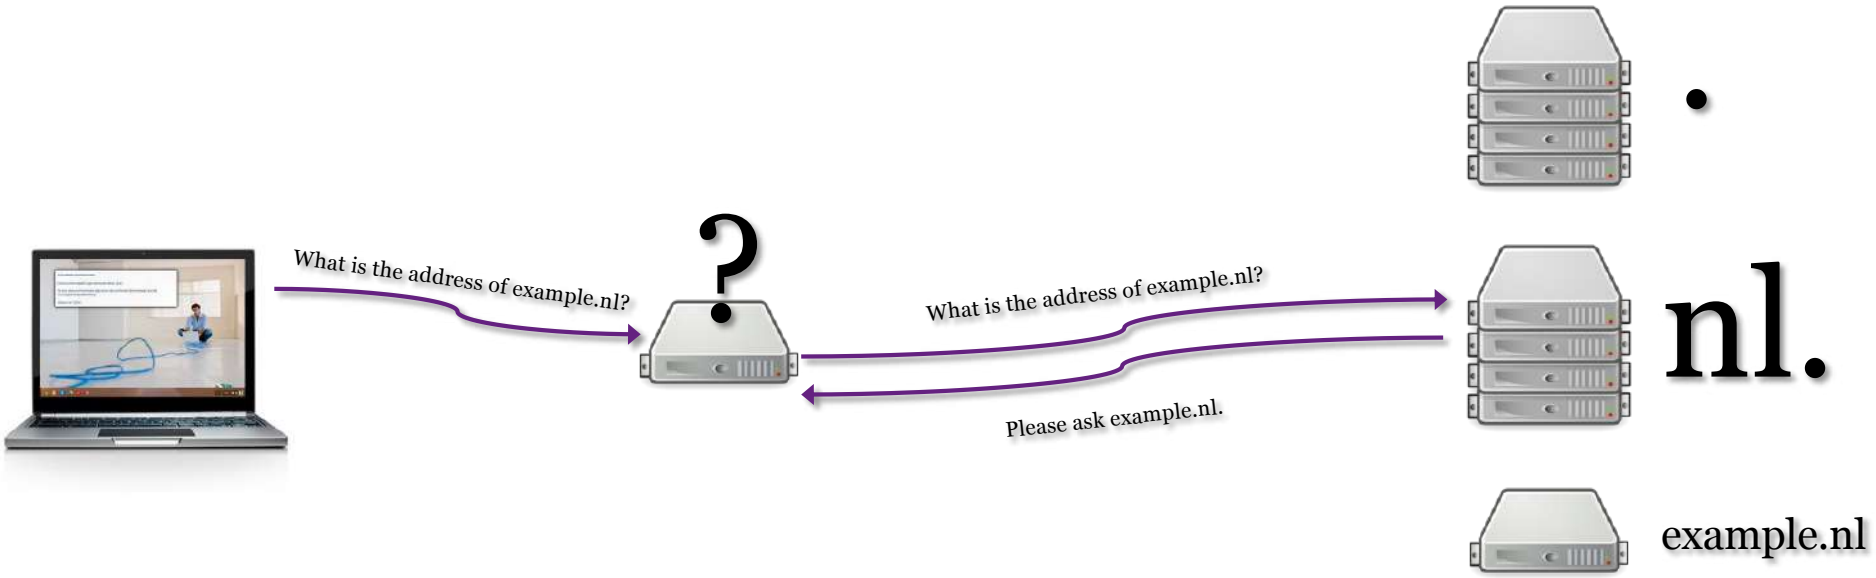

DNS: The resolver starts looking for the answer...

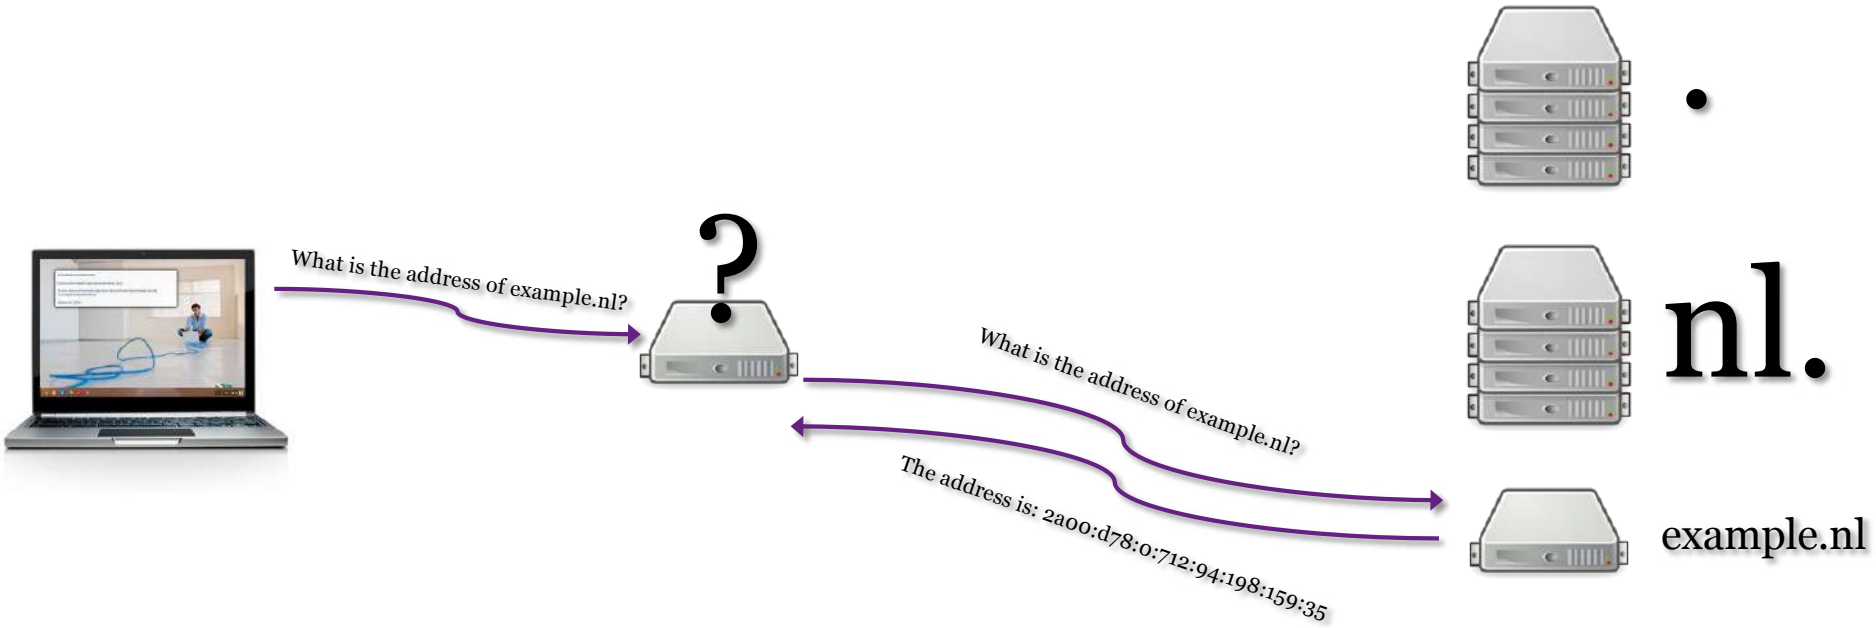

DNS: The resolver starts looking for the answer…

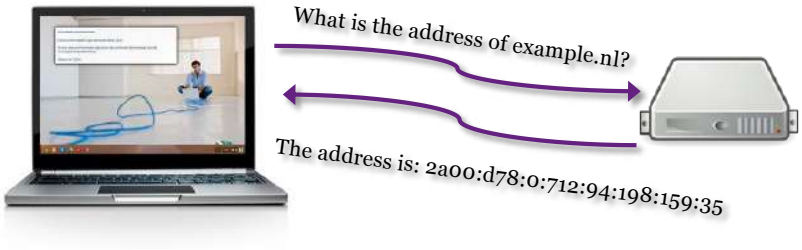

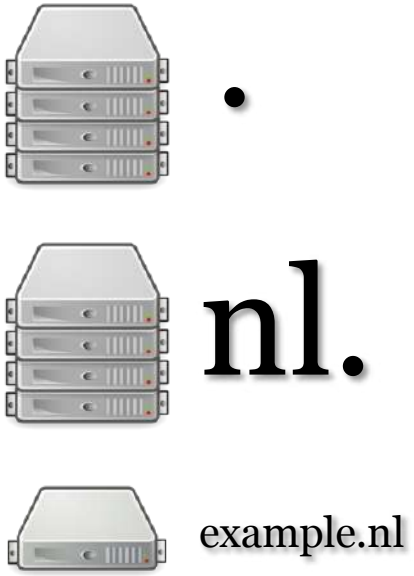

DNS: The resolver starts looking for the answer...

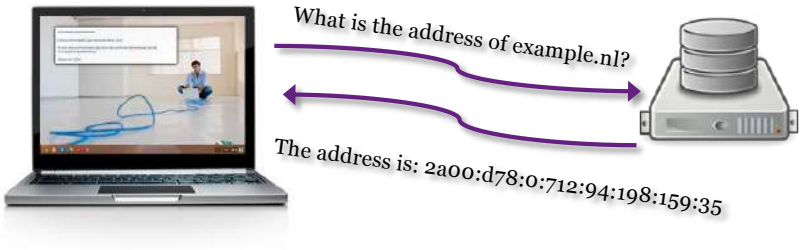

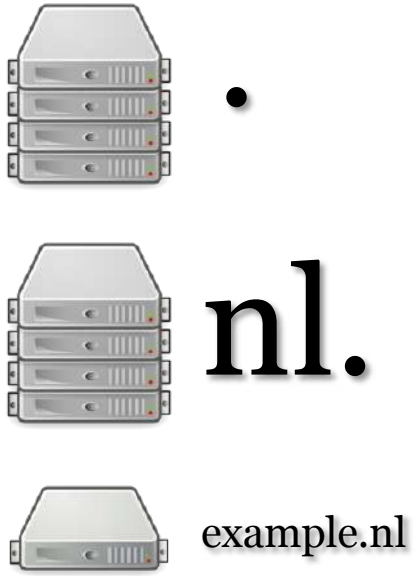

DNS: The resolver starts looking for the answer…

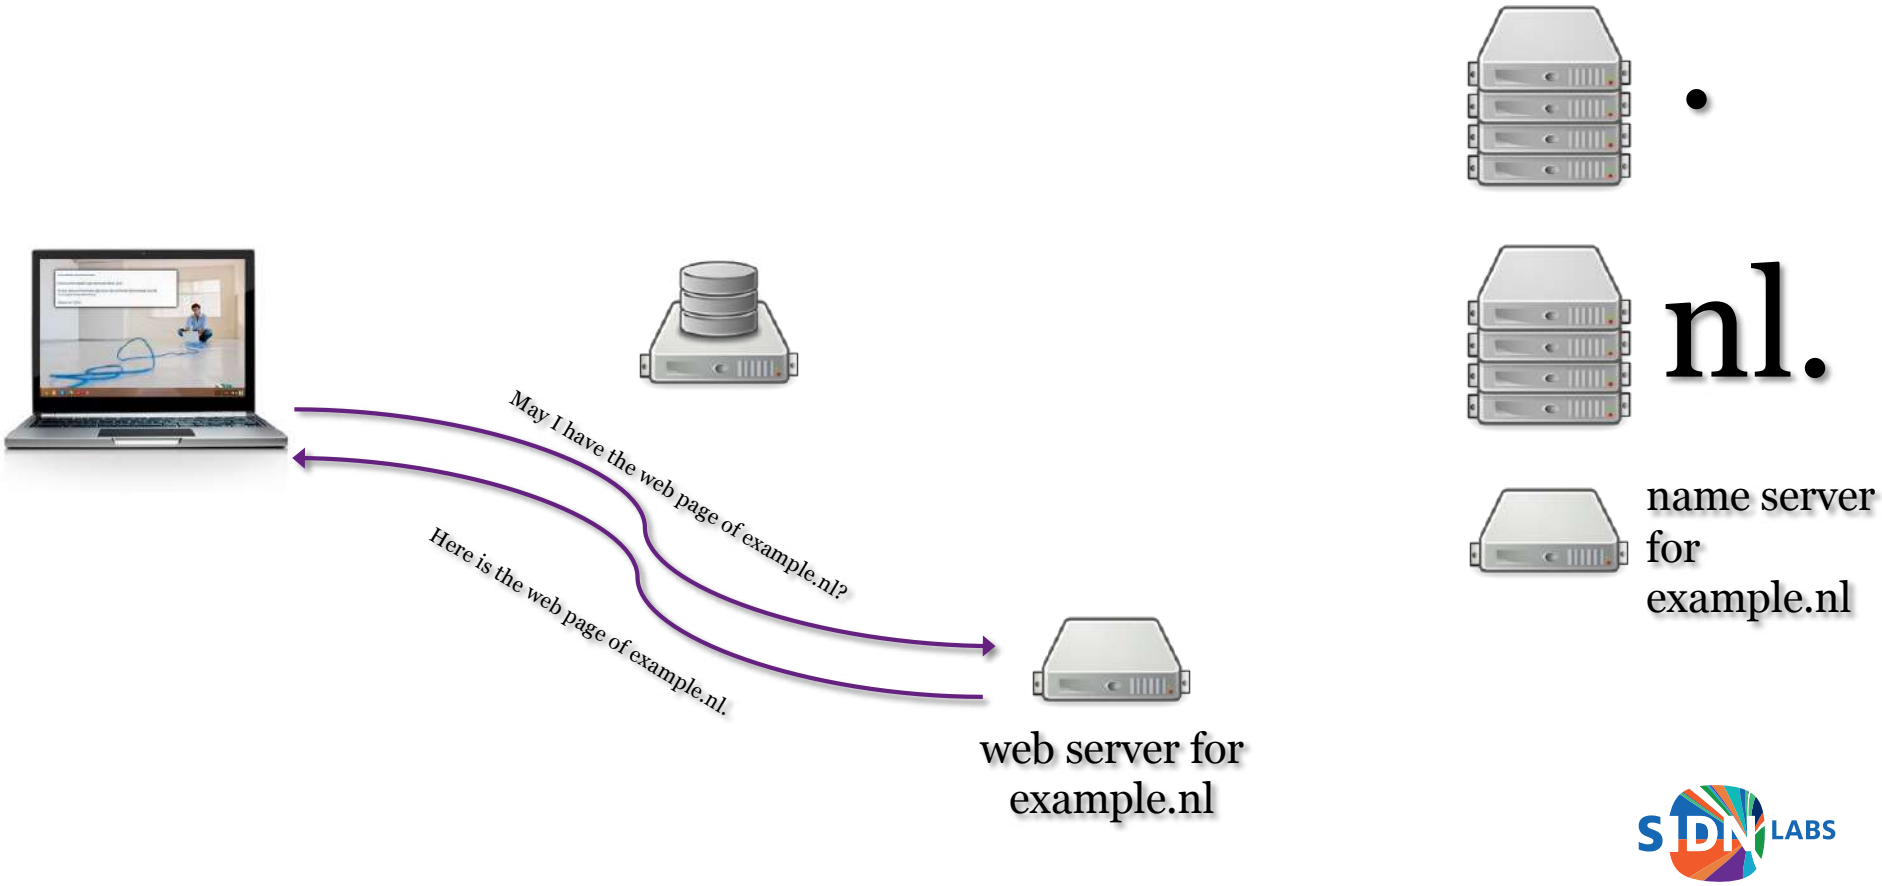

DNSSEC

About DNSSEC

Workings:

- DNS responses are digitally signed
- These digital signatures are checked by the resolver
- If the answer has been tampered with, the signature is incorrect

What could go wrong?

Local user

Resolver

Malicious hacker

Autoritatieve name servers

The resolver looks for the answer…

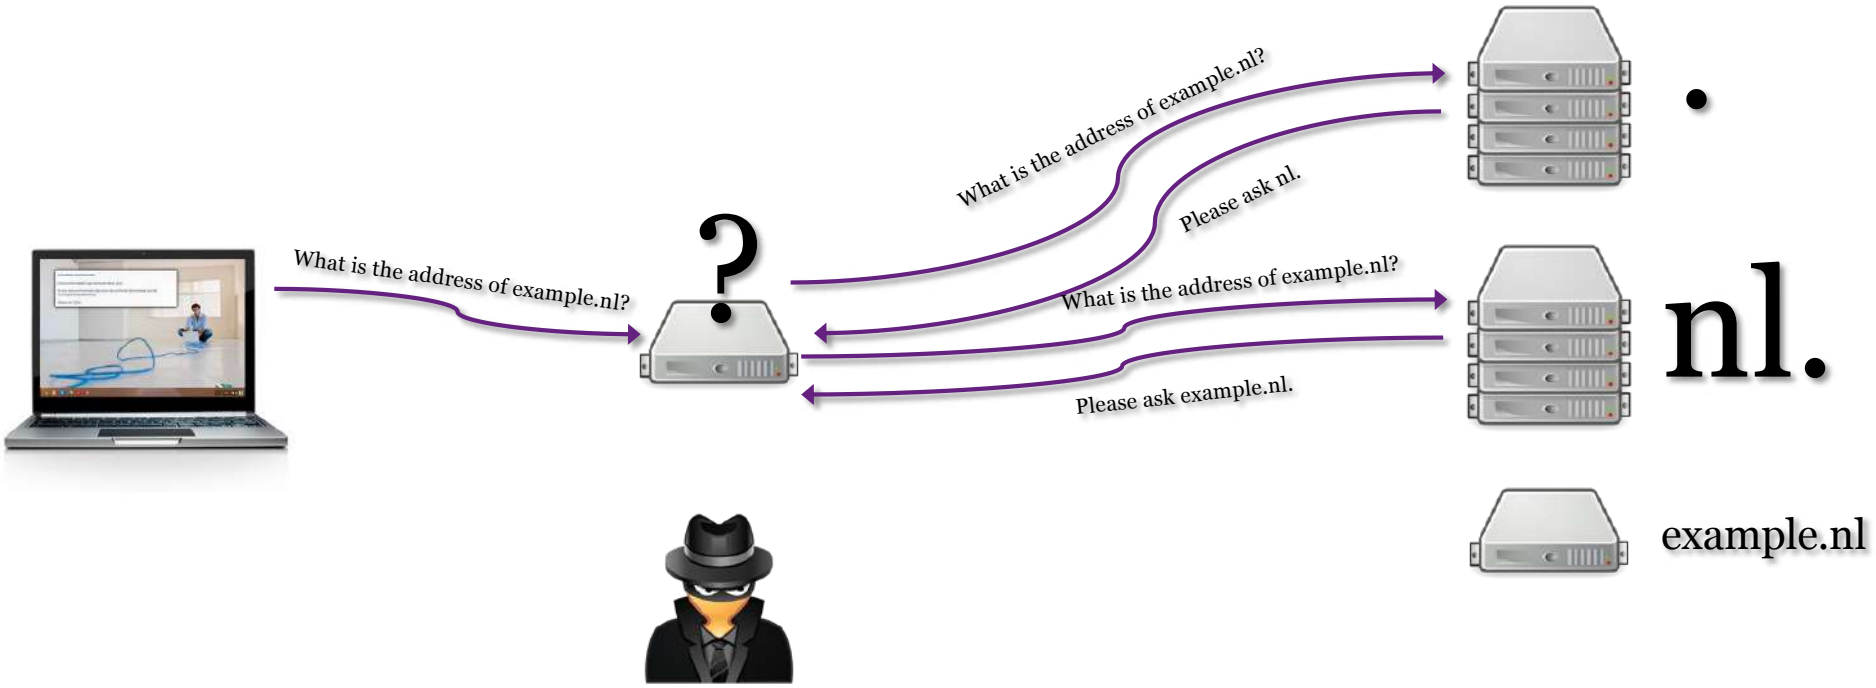

Lots of fake answers!

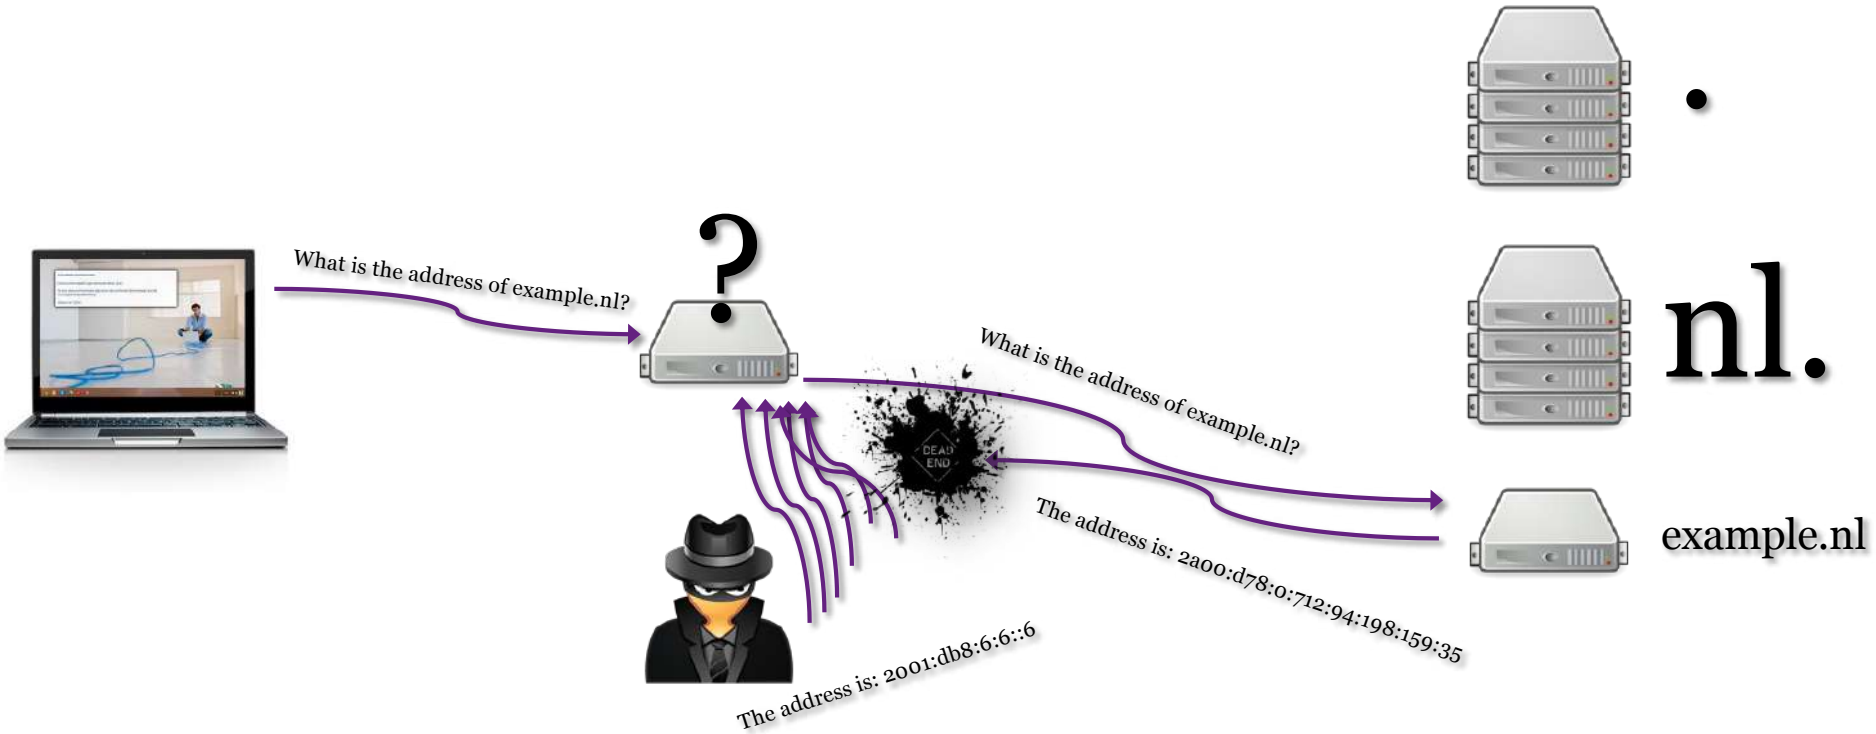

The resolver has fallen for a false answer...

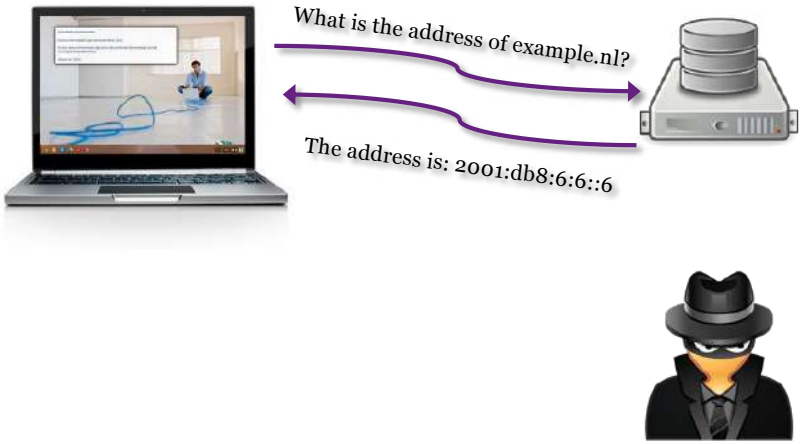

.

nl.

example.nl

SIDN LABS

The user has been redirected...

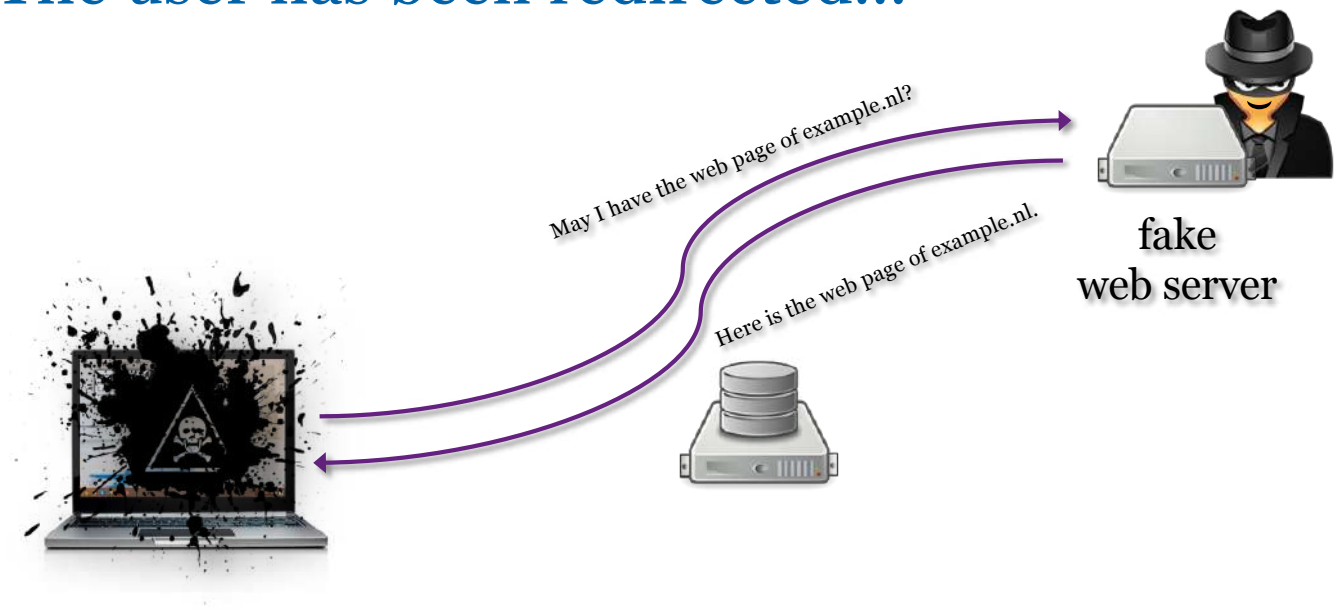

The user has been redirected…

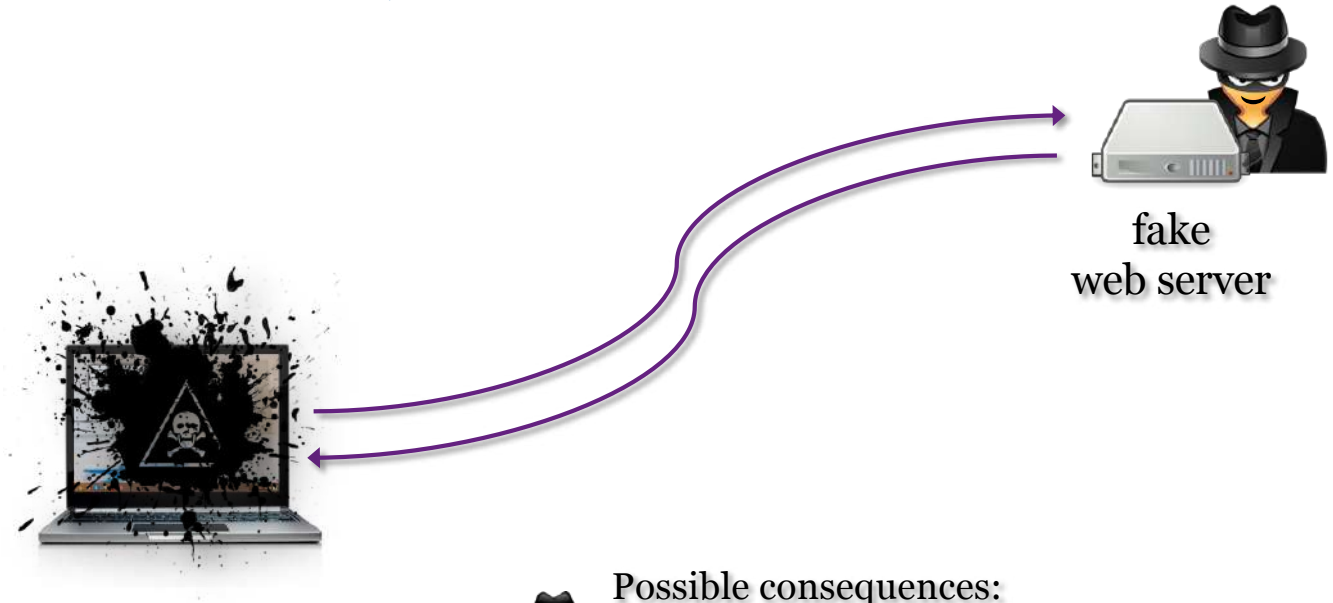

Possible consequences:

- Interception (such as emails or passwords)
- Presenting fake content (including malware)
- All the consequences as a result
  - (financial damage, reputational damage, etc.)

About DNSSEC

```
;; ANSWER SECTION:
example.nl.             3600 IN AAAA 2a00:d78::712:94:198:159:35
example.nl.             3600 IN RRSIG AAAA 5 2 3600 20160514113814 (
                                20160414113121 15516 example.nl.
                                gFgoC1jh7AMNbxDmCfP2kxQ7FJt7rEllAUshps1YIXLN
                                CA2T2z80xZMYUyAT9fxOY0jVIbL6NVFiHAuQ3bz4xSsw
                                +uweGvkIgkRQSQQavlmBrelXE45pdARmkFy0fC7eCX4D
                                4vyvk8QogdpyGxYqZdU0atrZ3lsFmsH9KSTTBYQ= )
```

With DNSSEC,
every DNS response is provided with a digital signature,
so that the content can be checked for authenticity.

SIDN LABS

DNSSEC Key Signing Ceremony

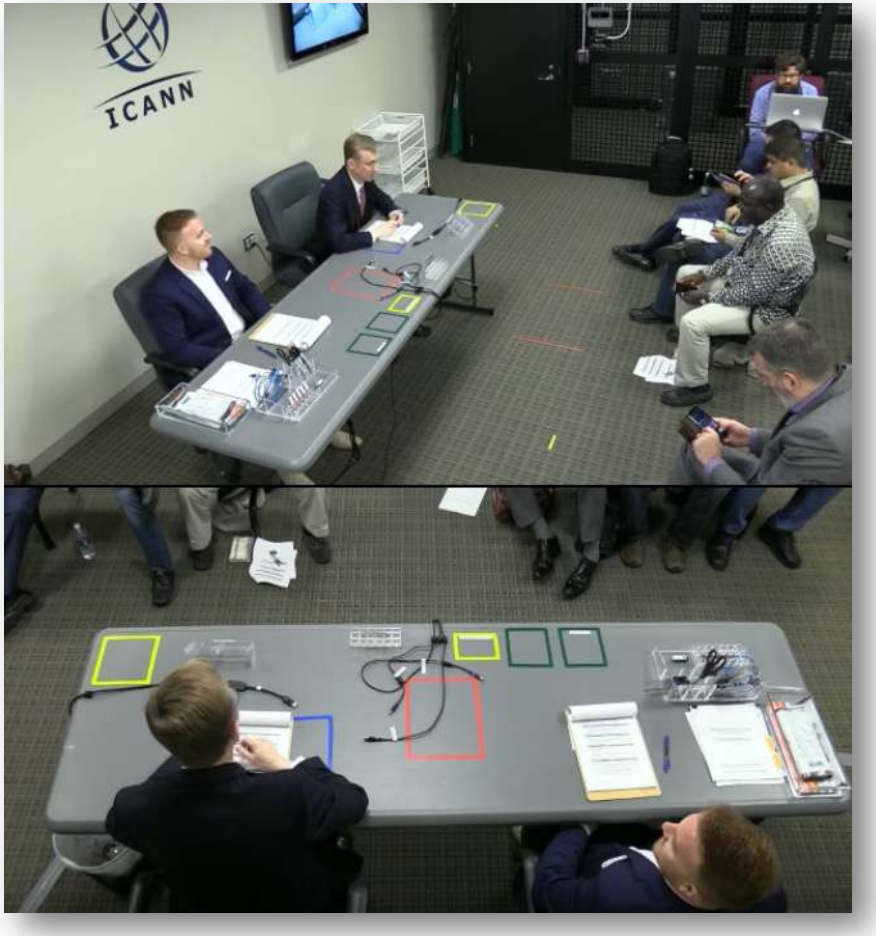

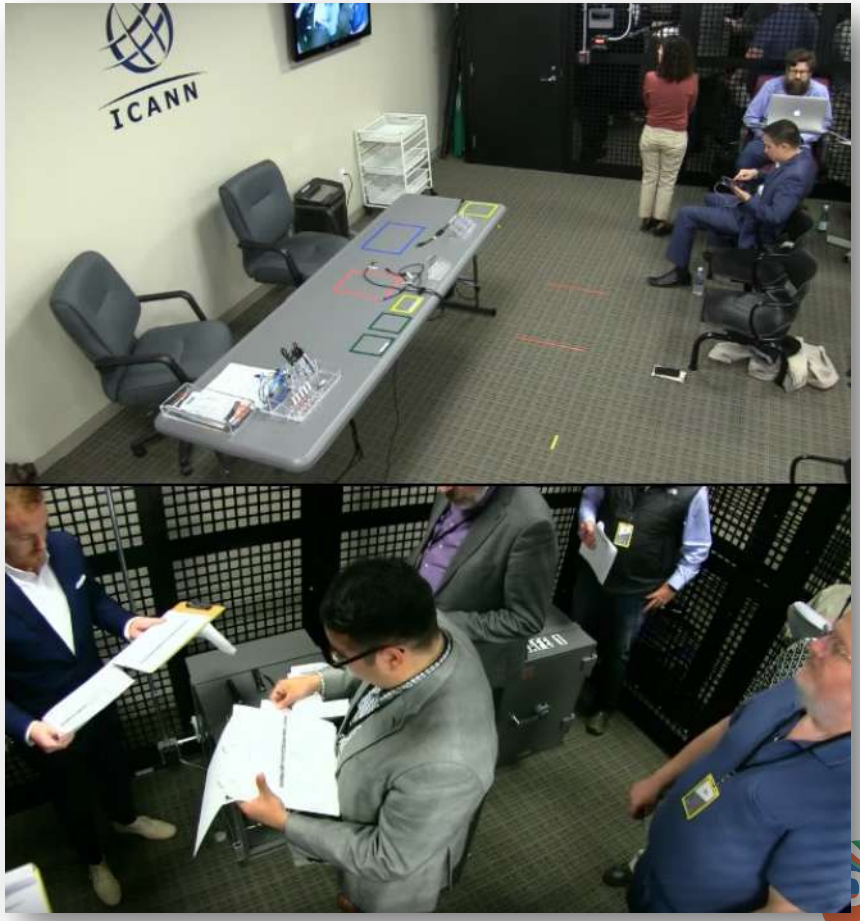

https://www.youtube.com/watch?v=ZTxweLGjZSU

DNSSEC Key Signing Ceremony

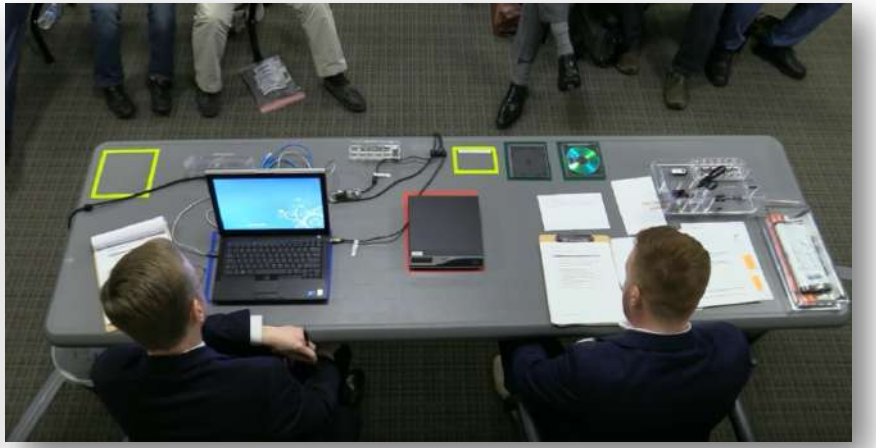

Open the Credential Safe #2

Step	Activity	Initials	Time
7.	CA and IW1 brings a flashlight then escorts SSC2, COs into the safe room.		
8.	SSC2 opens Safe #2 while shielding the combination from the camera.		
9.	SSC2 removes the existing safe log and shows the most recent page to the audit camera. SSC2 obtains the pre-printed safe log from IW1, then writes the date/time and signature on the safe log where "Open Safe" is indicated. IW1 verifies this entry, then initials it.		

version 2.0 Page 4 of 27

https://www.youtube.com/watch?v=ZTxweLGjZSU

Part 2: Routing

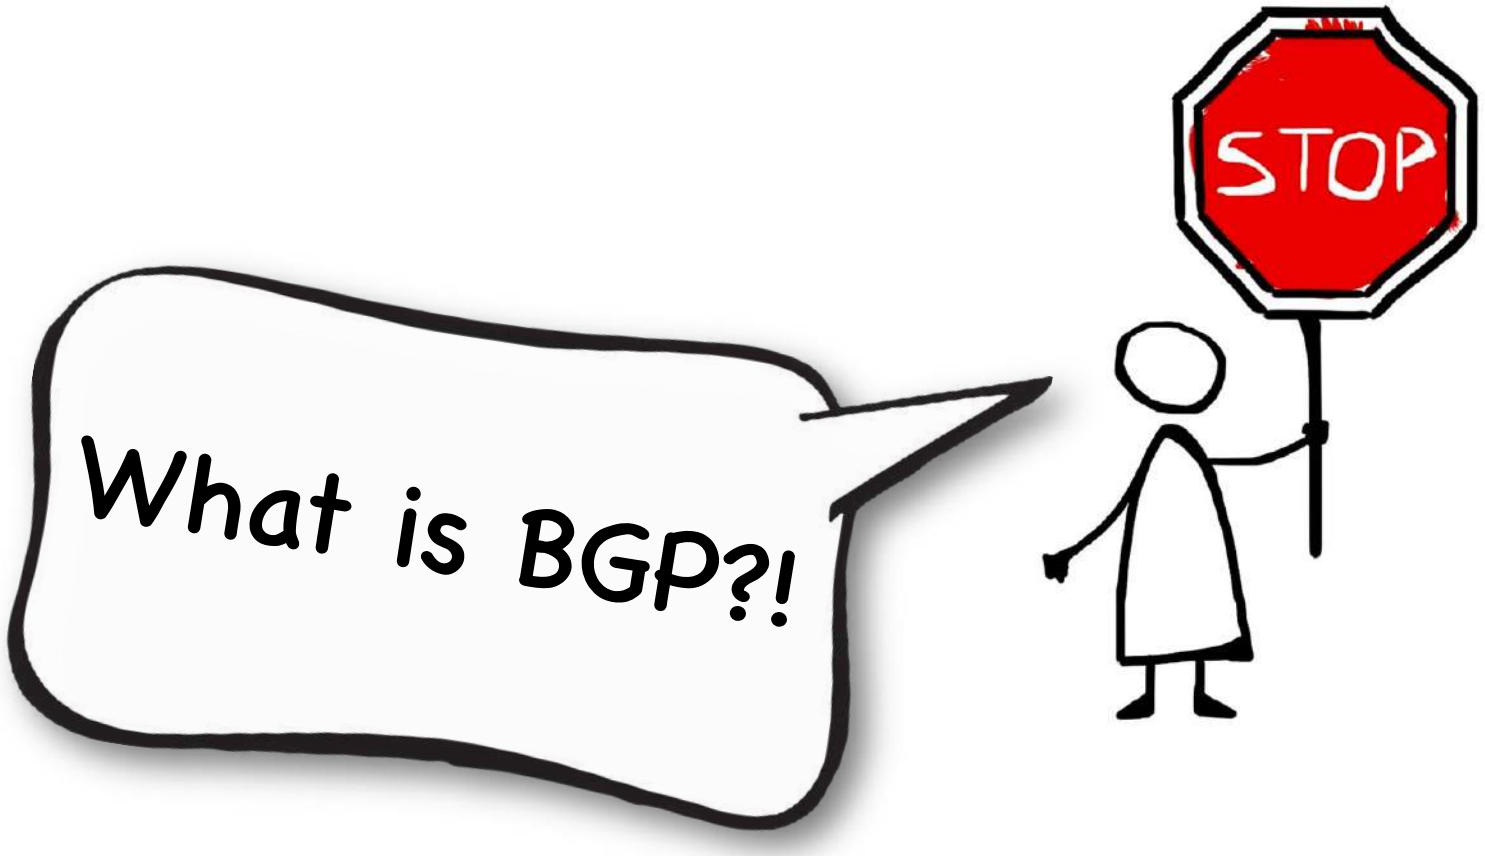

A typical network

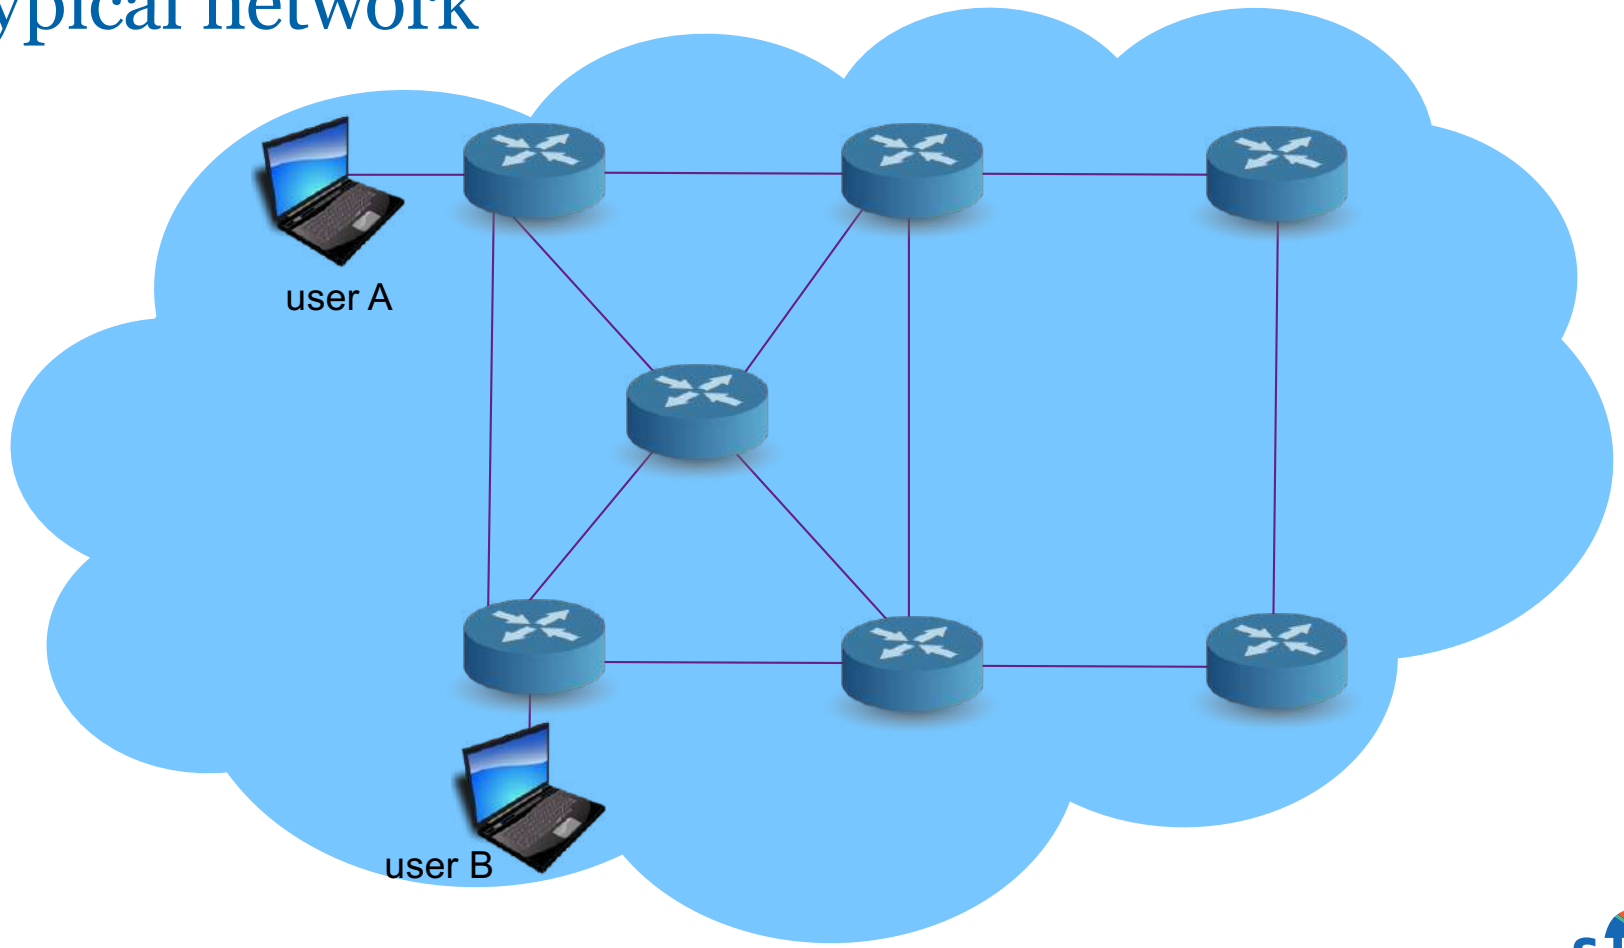

A typical network

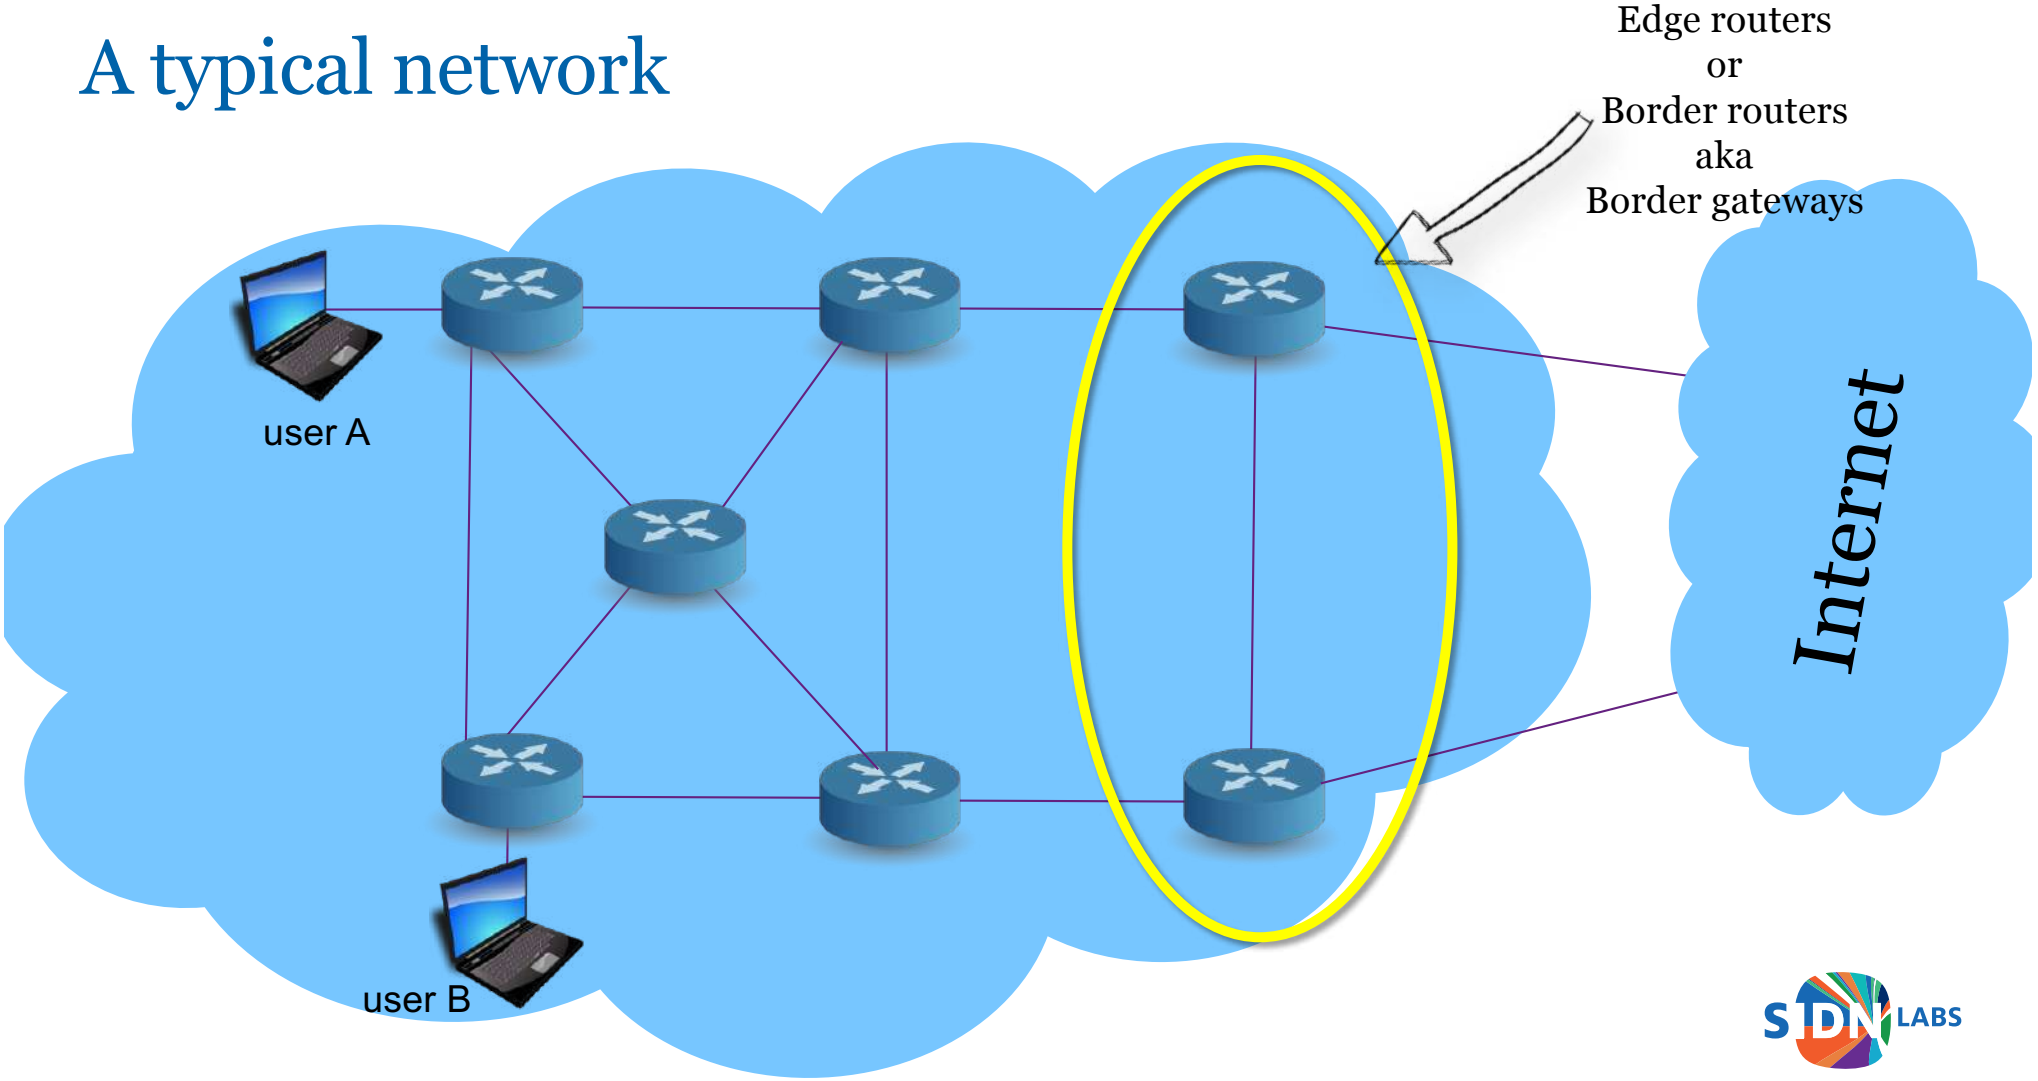

Average network

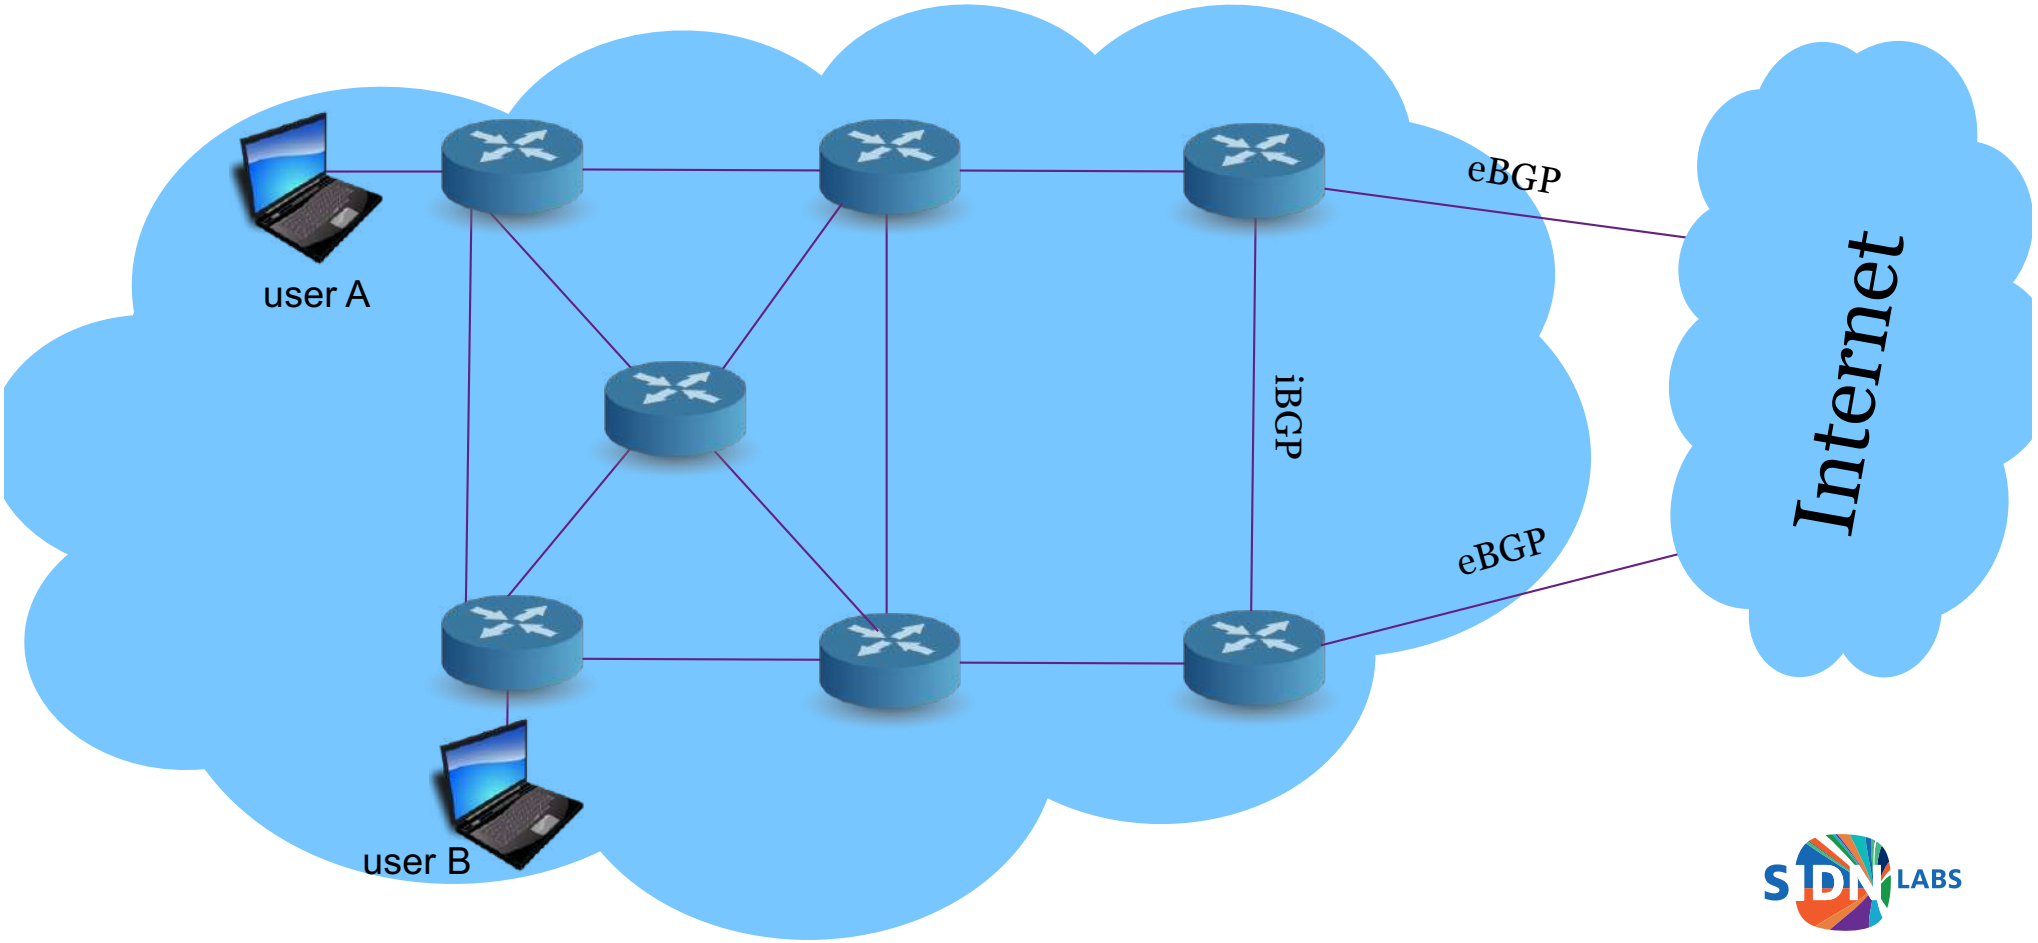

BGP (in a nutshell)

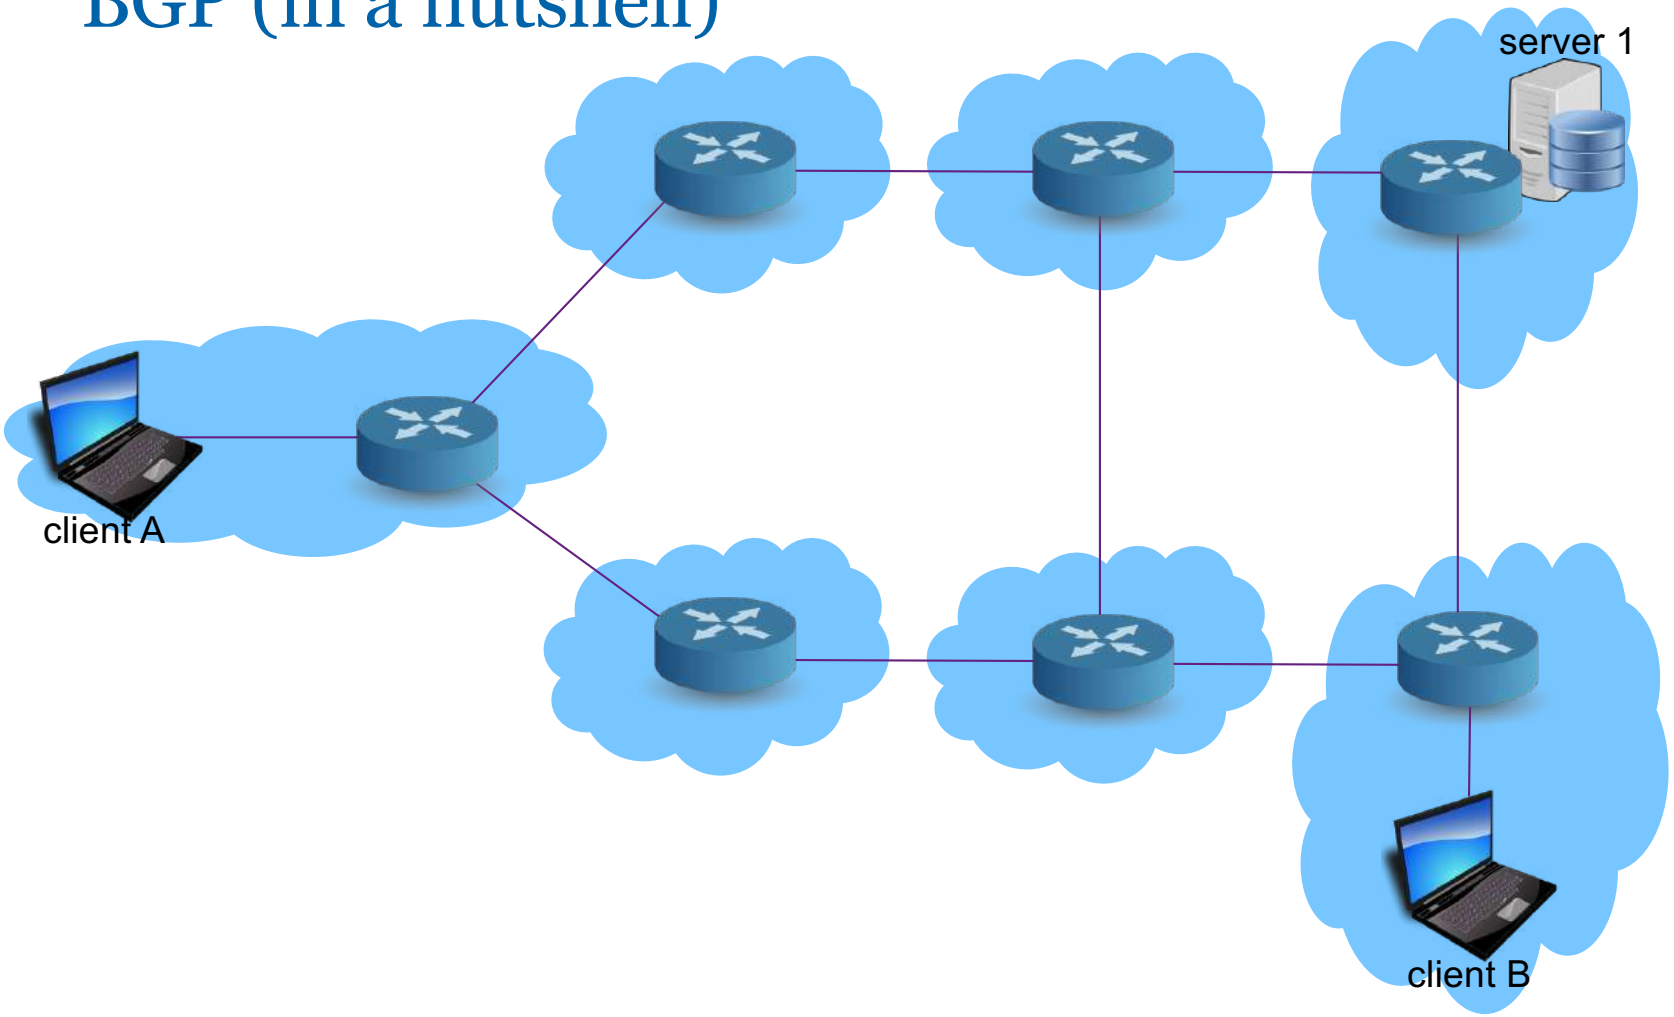

BGP (in a nutshell)

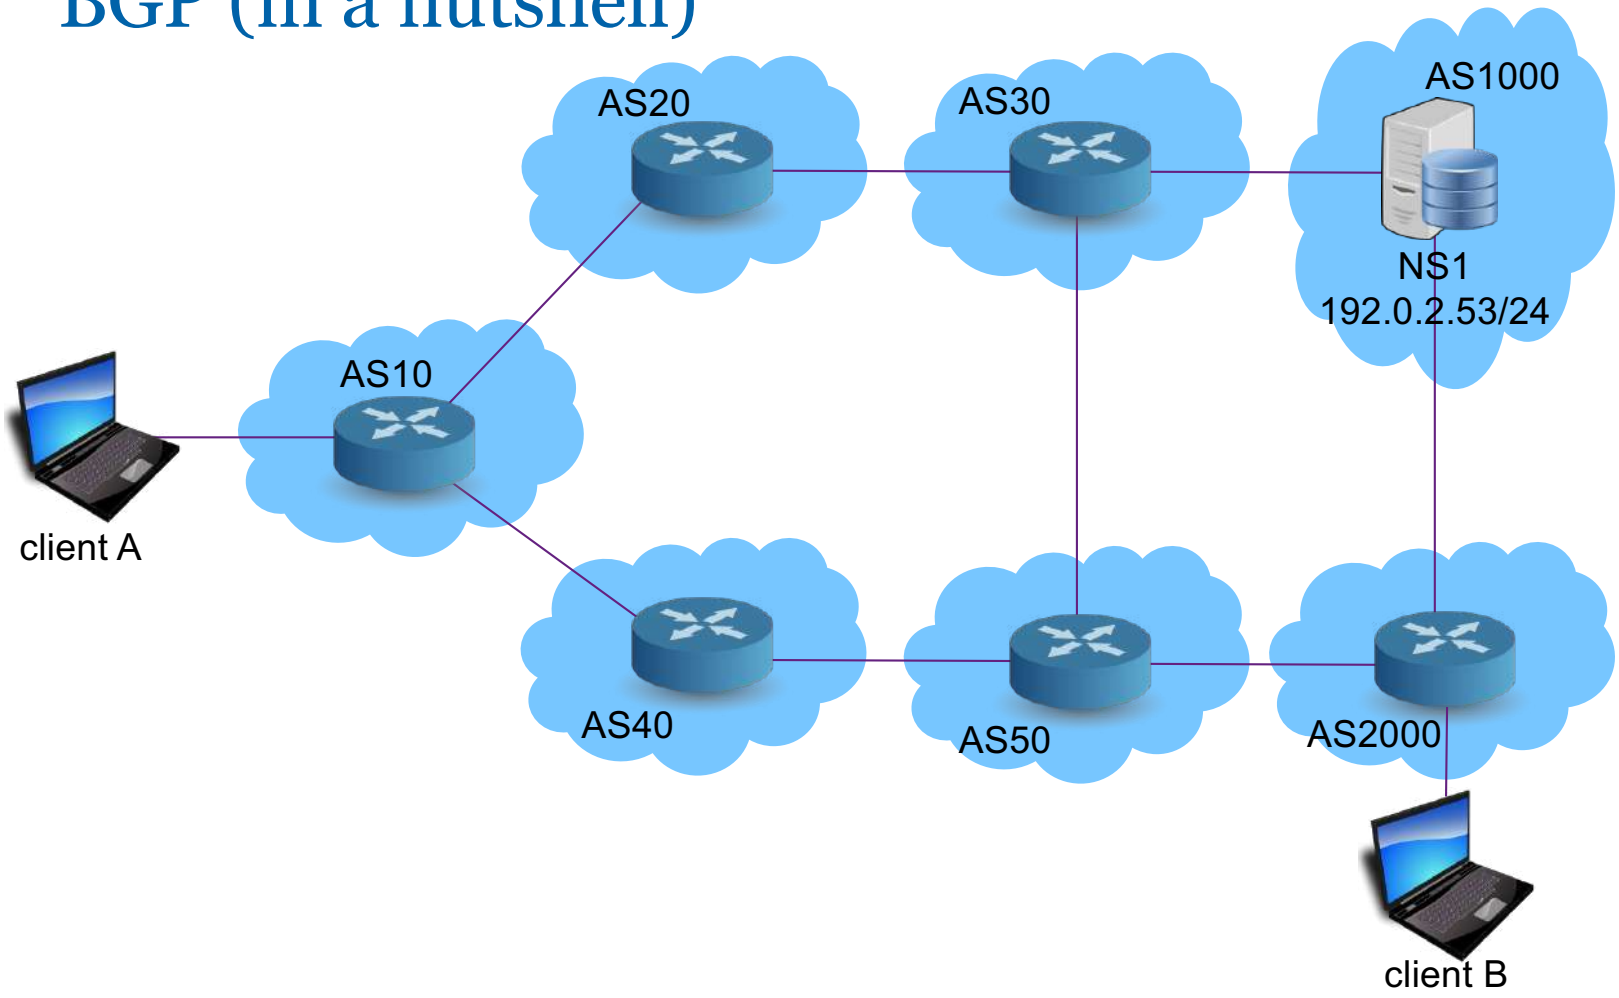

BGP (in a nutshell)

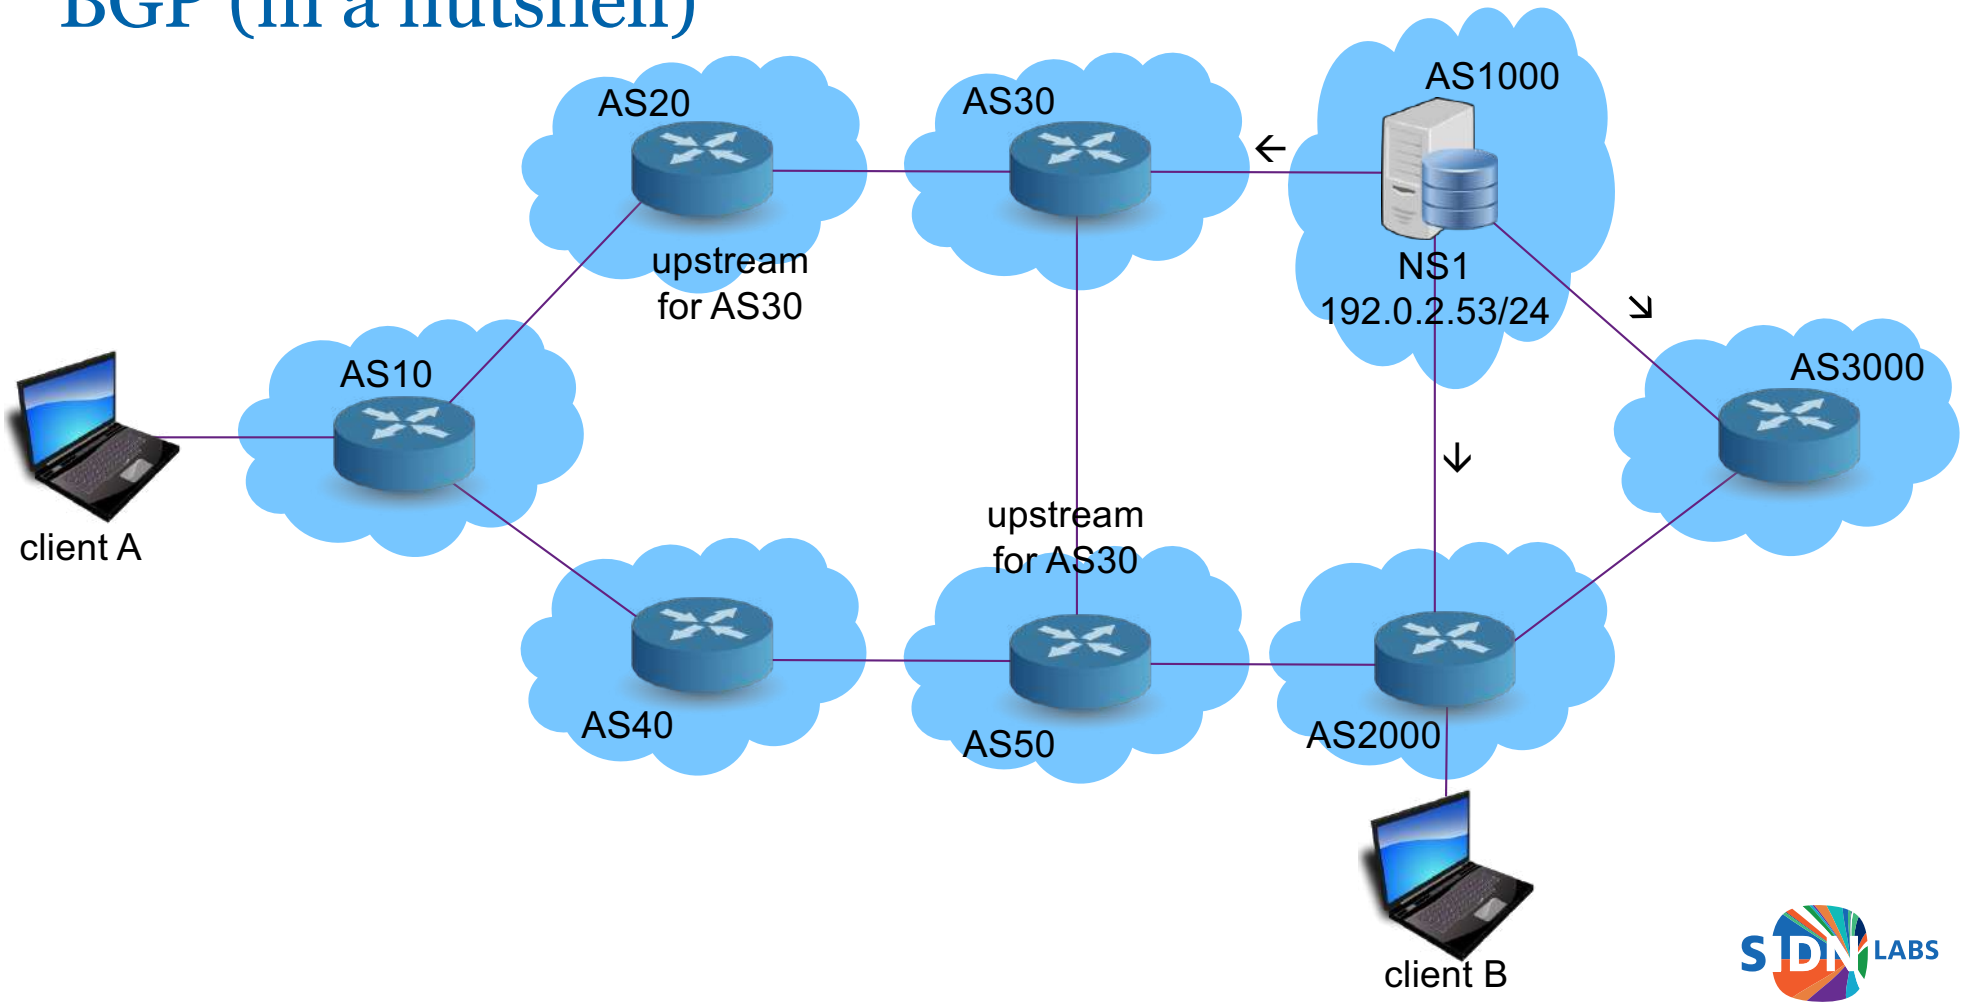

BGP (in a nutshell)

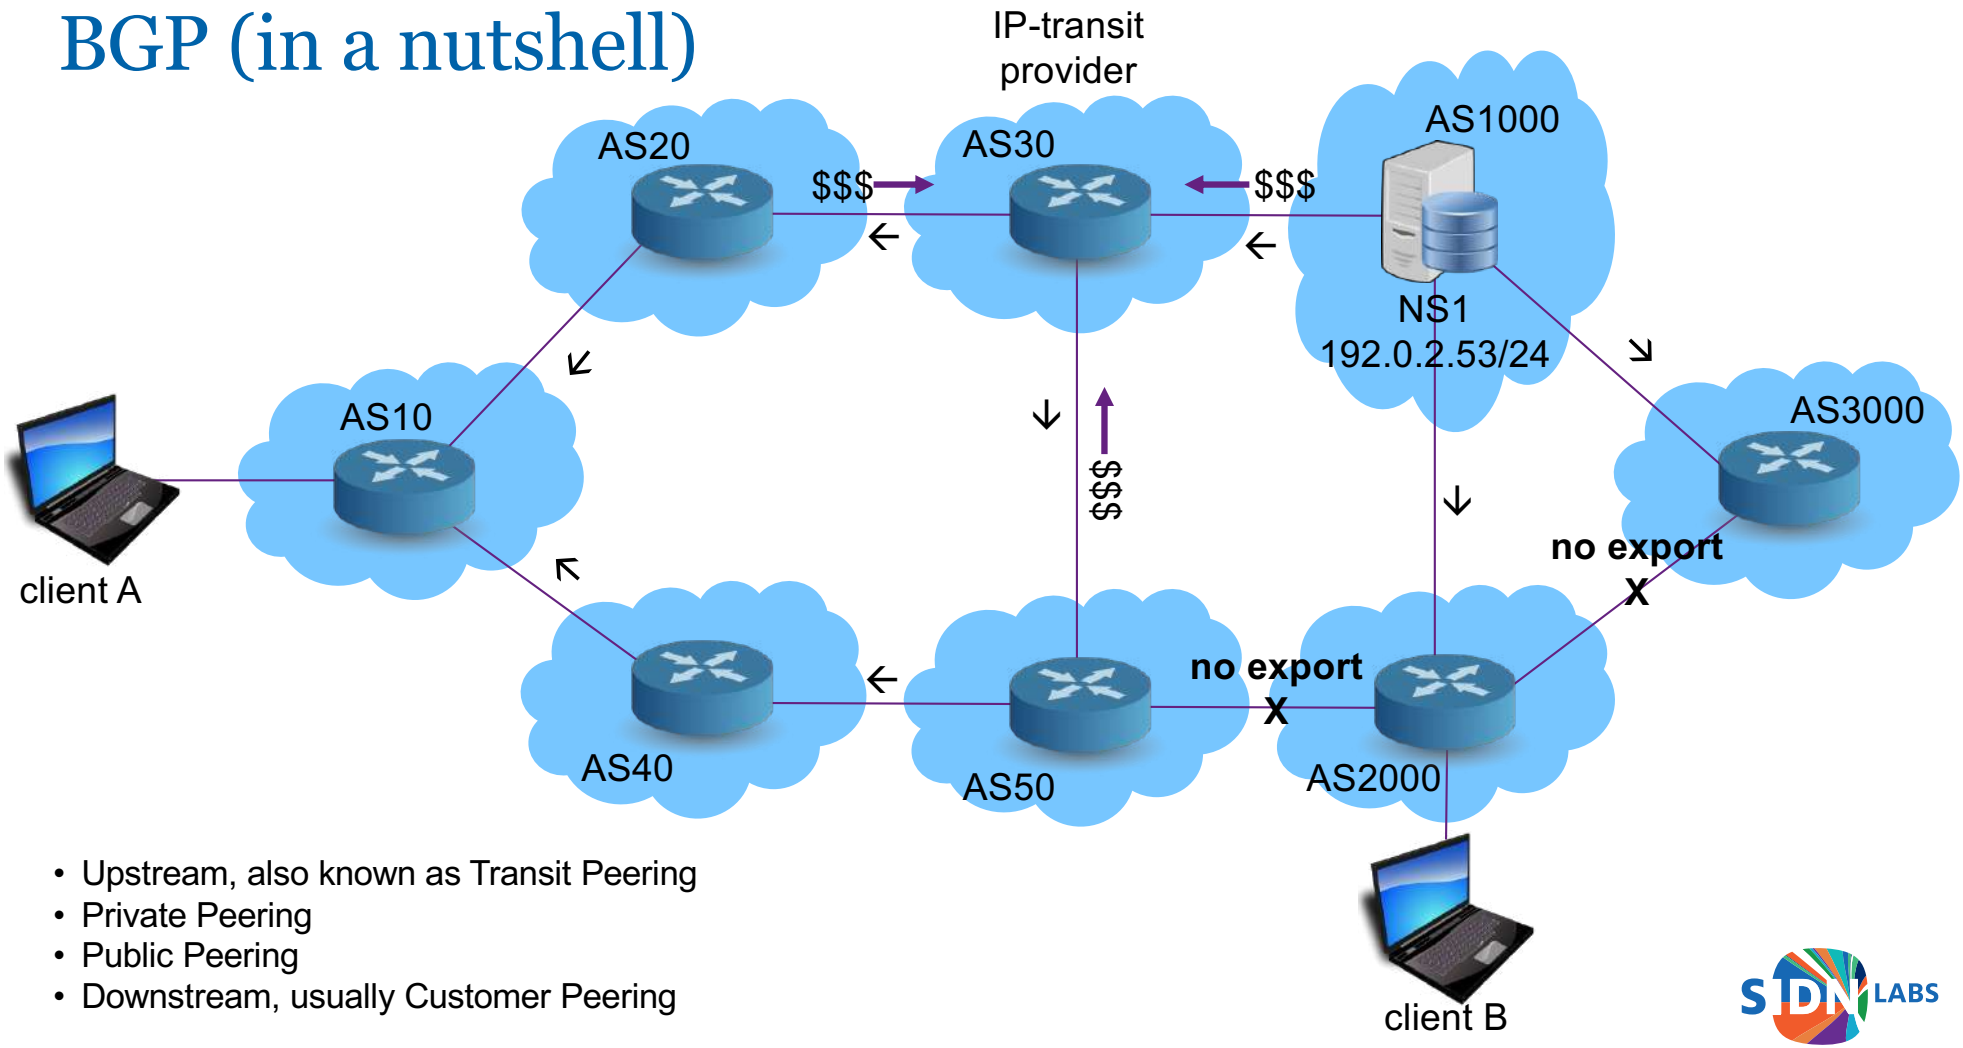

- Upstream, also known as Transit Peering
- Private Peering
- Public Peering
- Downstream, usually Customer Peering

BGP (in a nutshell)

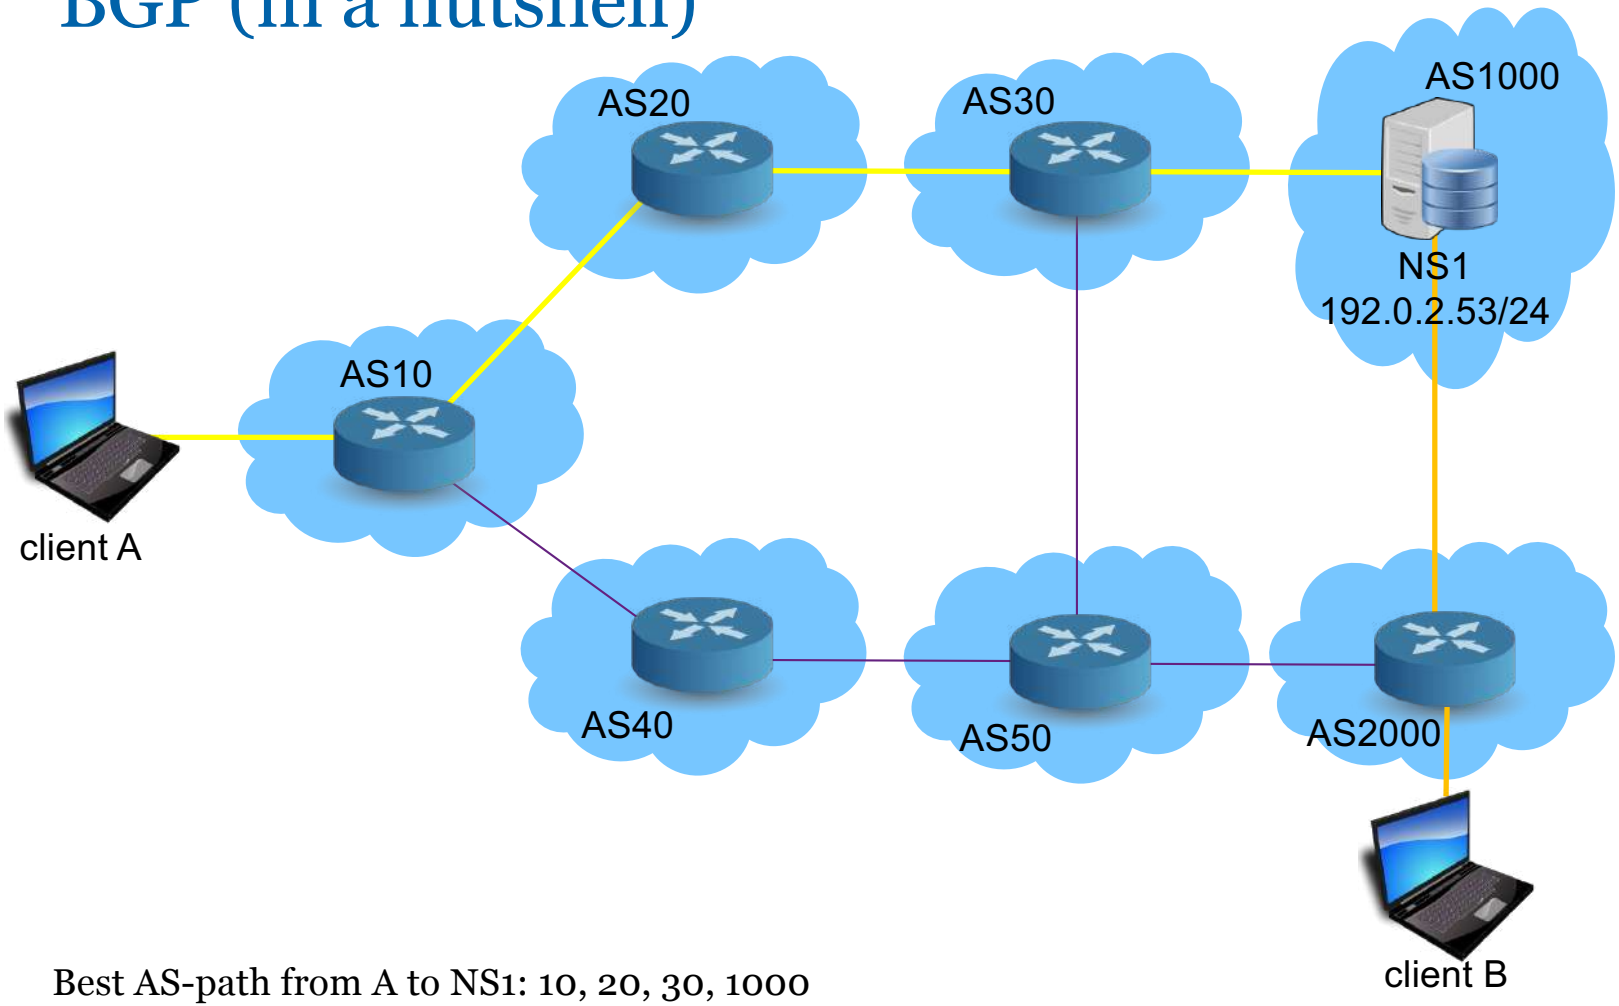

Best AS-path from A to NS1: 10, 20, 30, 1000

BGP (in a nutshell)

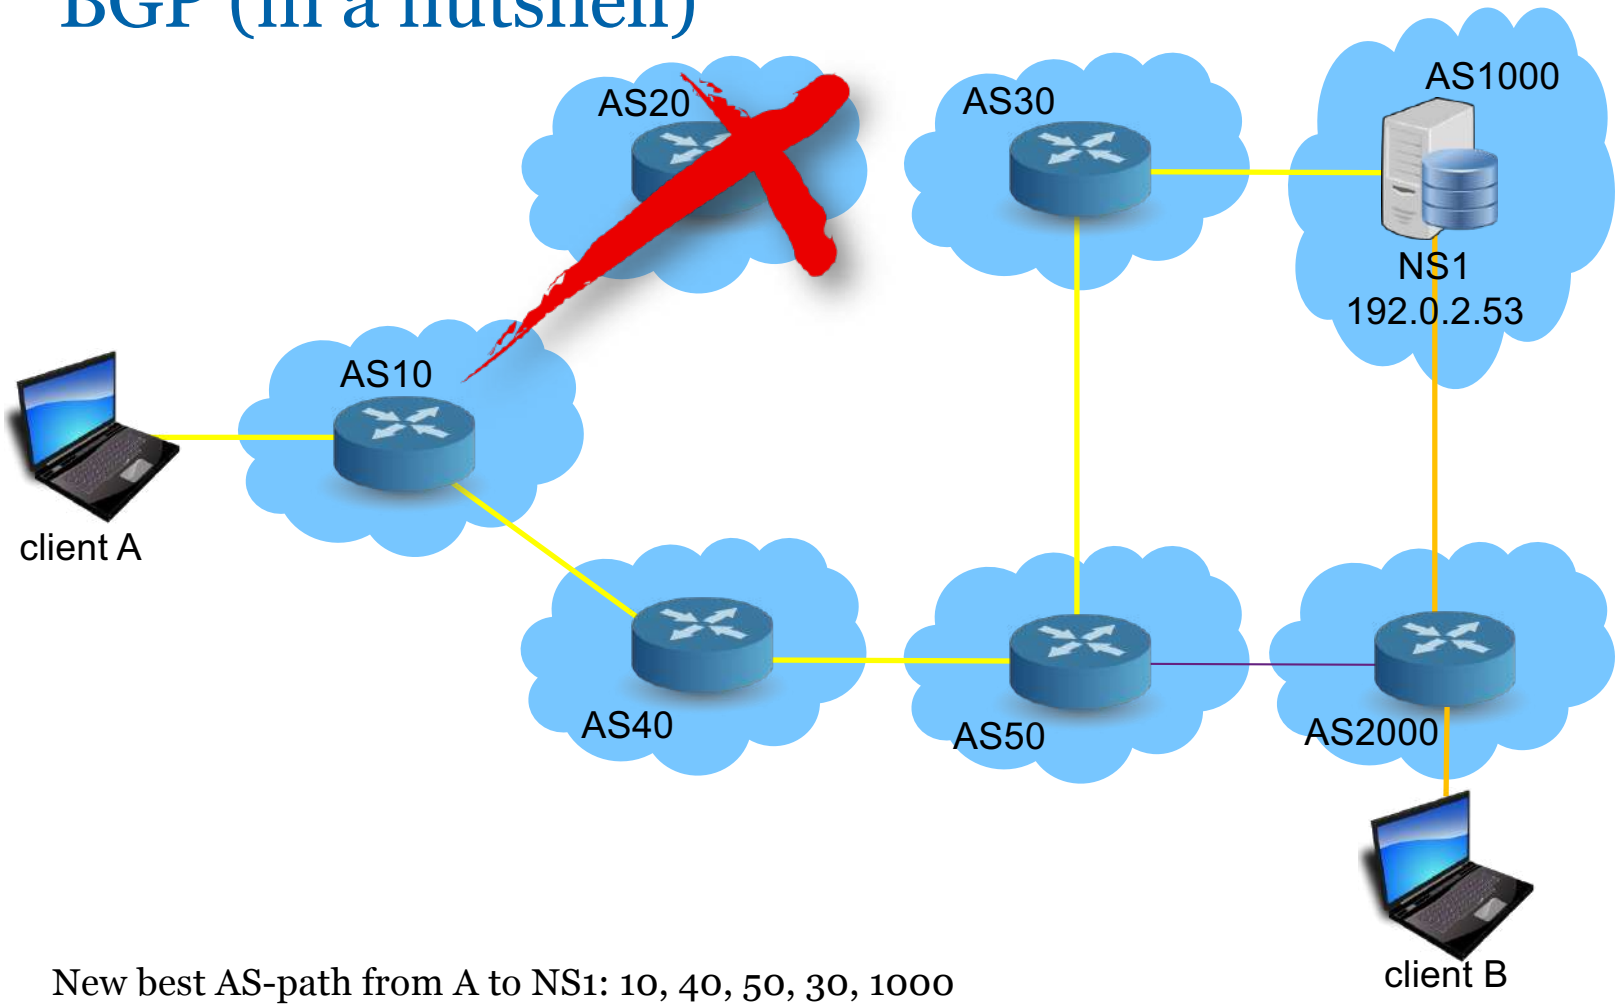

New best AS-path from A to NS1: 10, 40, 50, 30, 1000

Traffic engineering

BGP traffic engineering: local preference

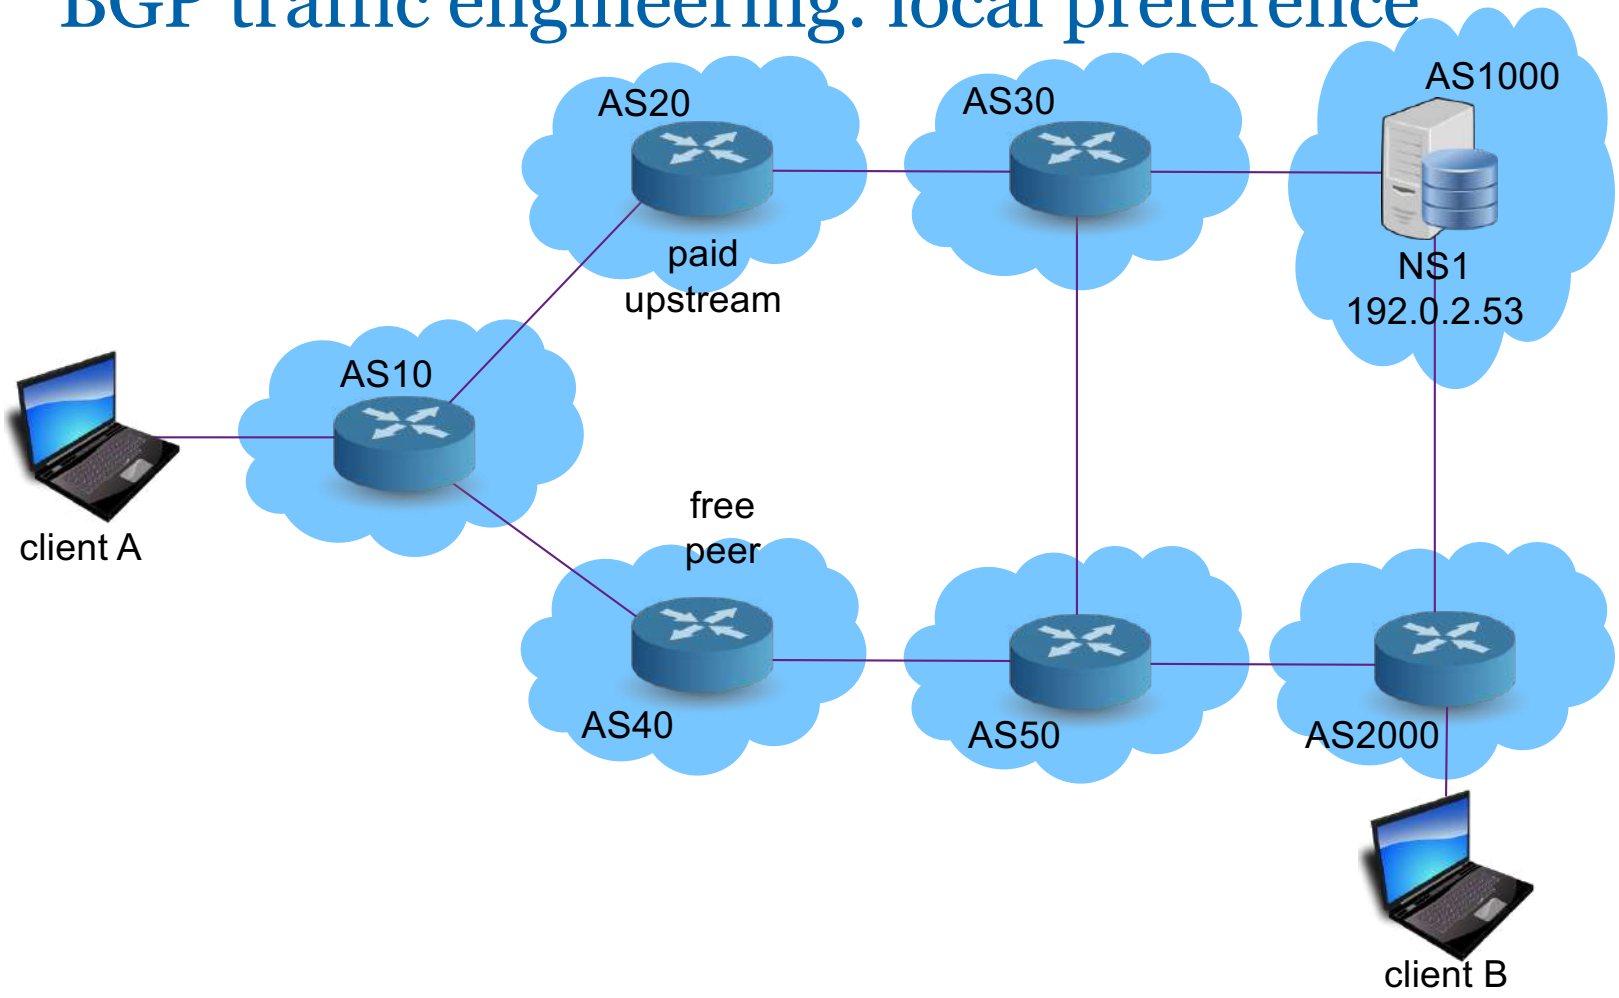

BGP traffic engineering: local preference

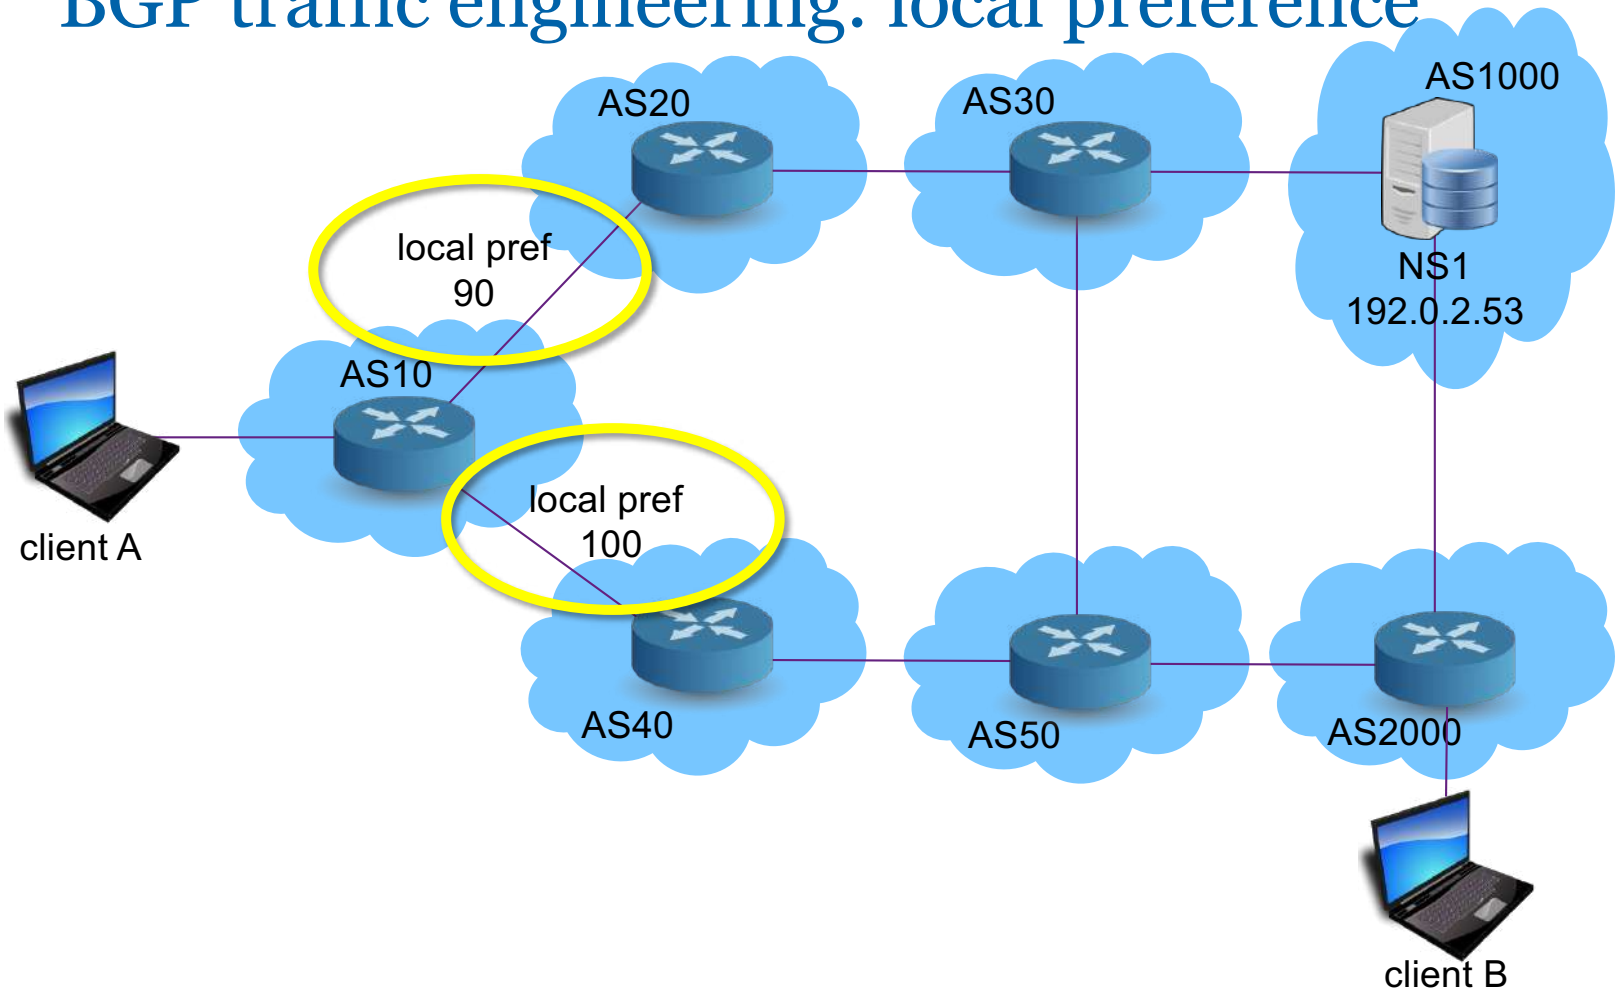

BGP traffic engineering: local preference

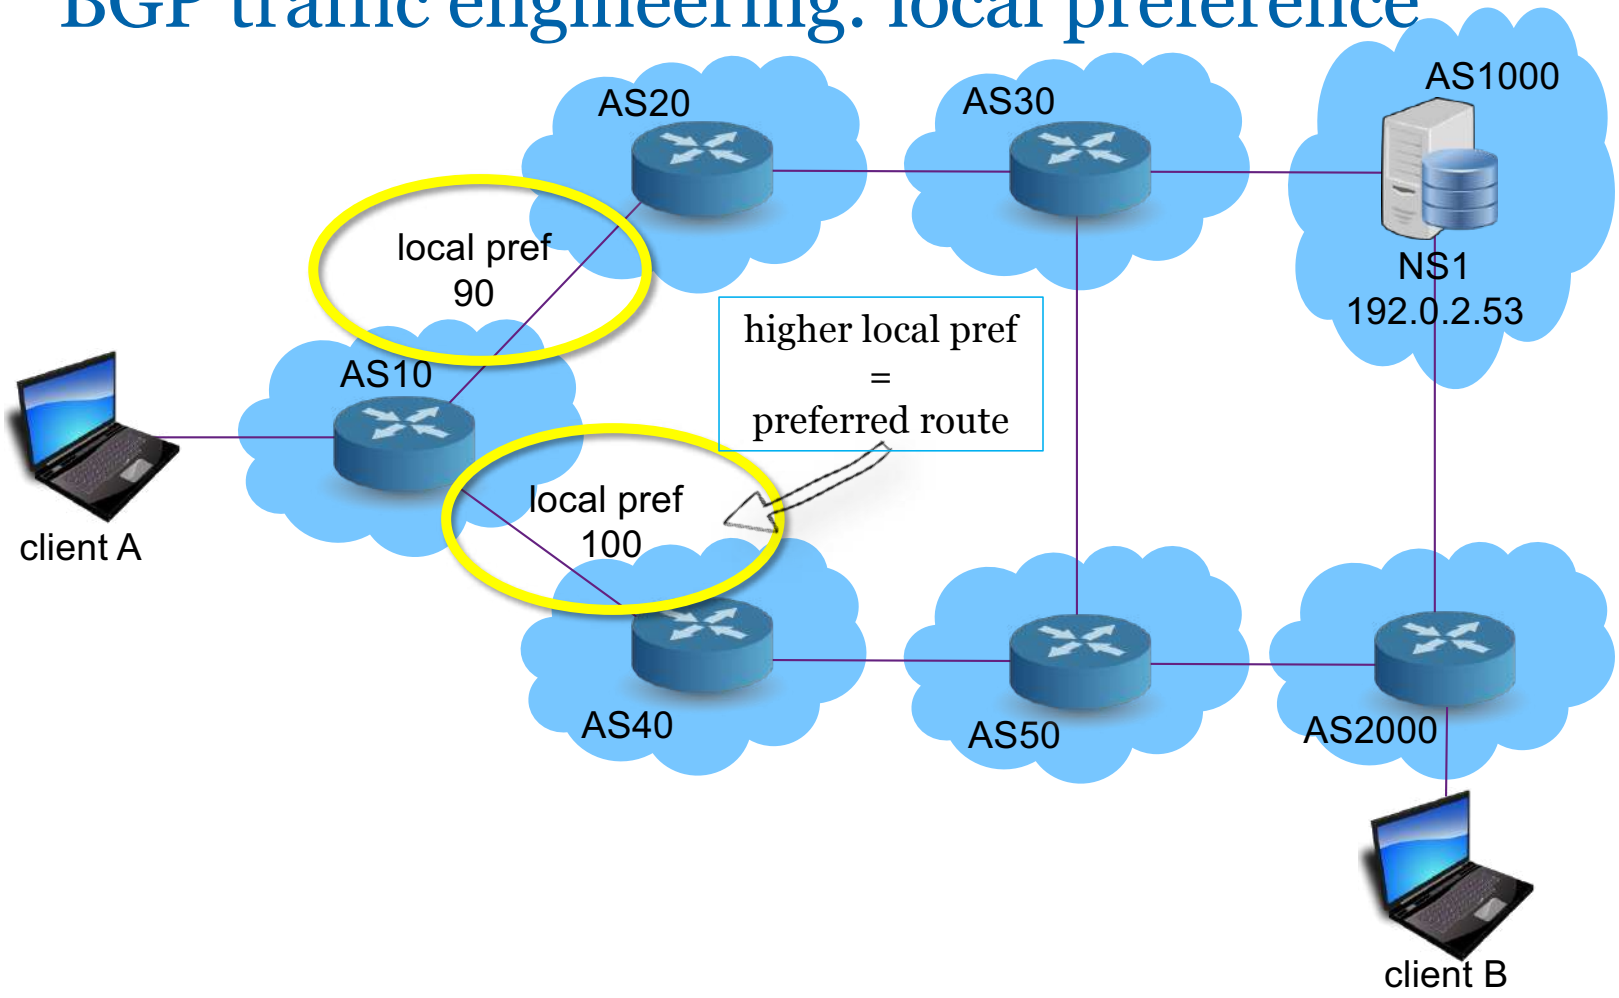

BGP traffic engineering: local preference

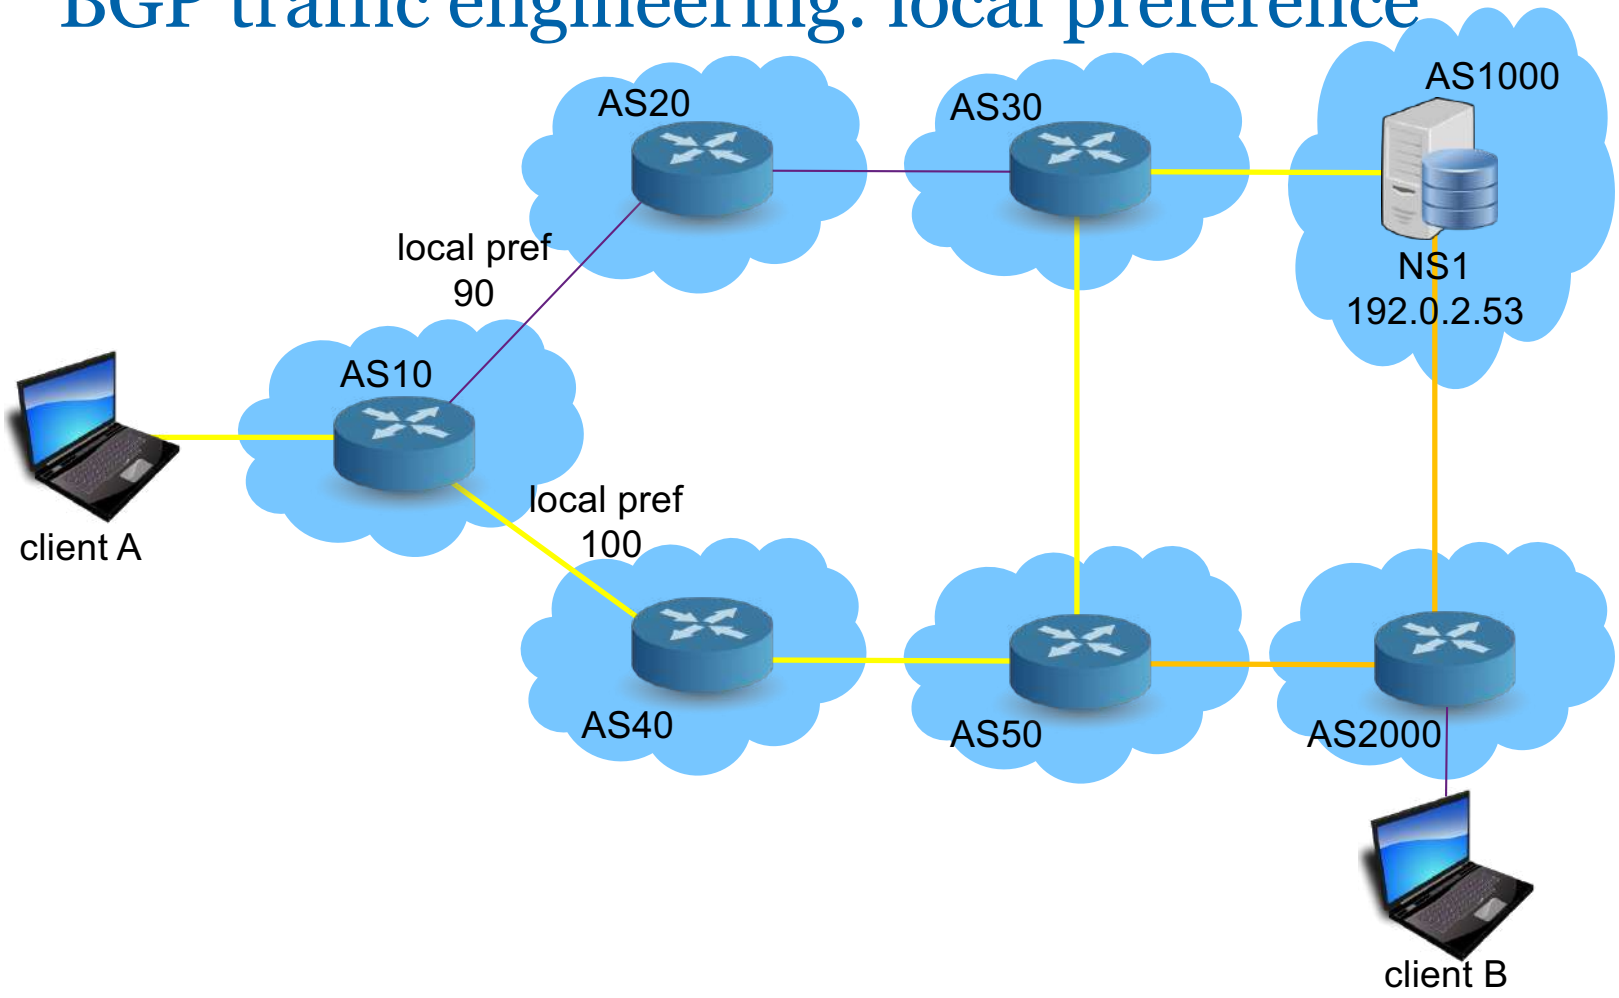

BGP traffic engineering: AS-path prepending

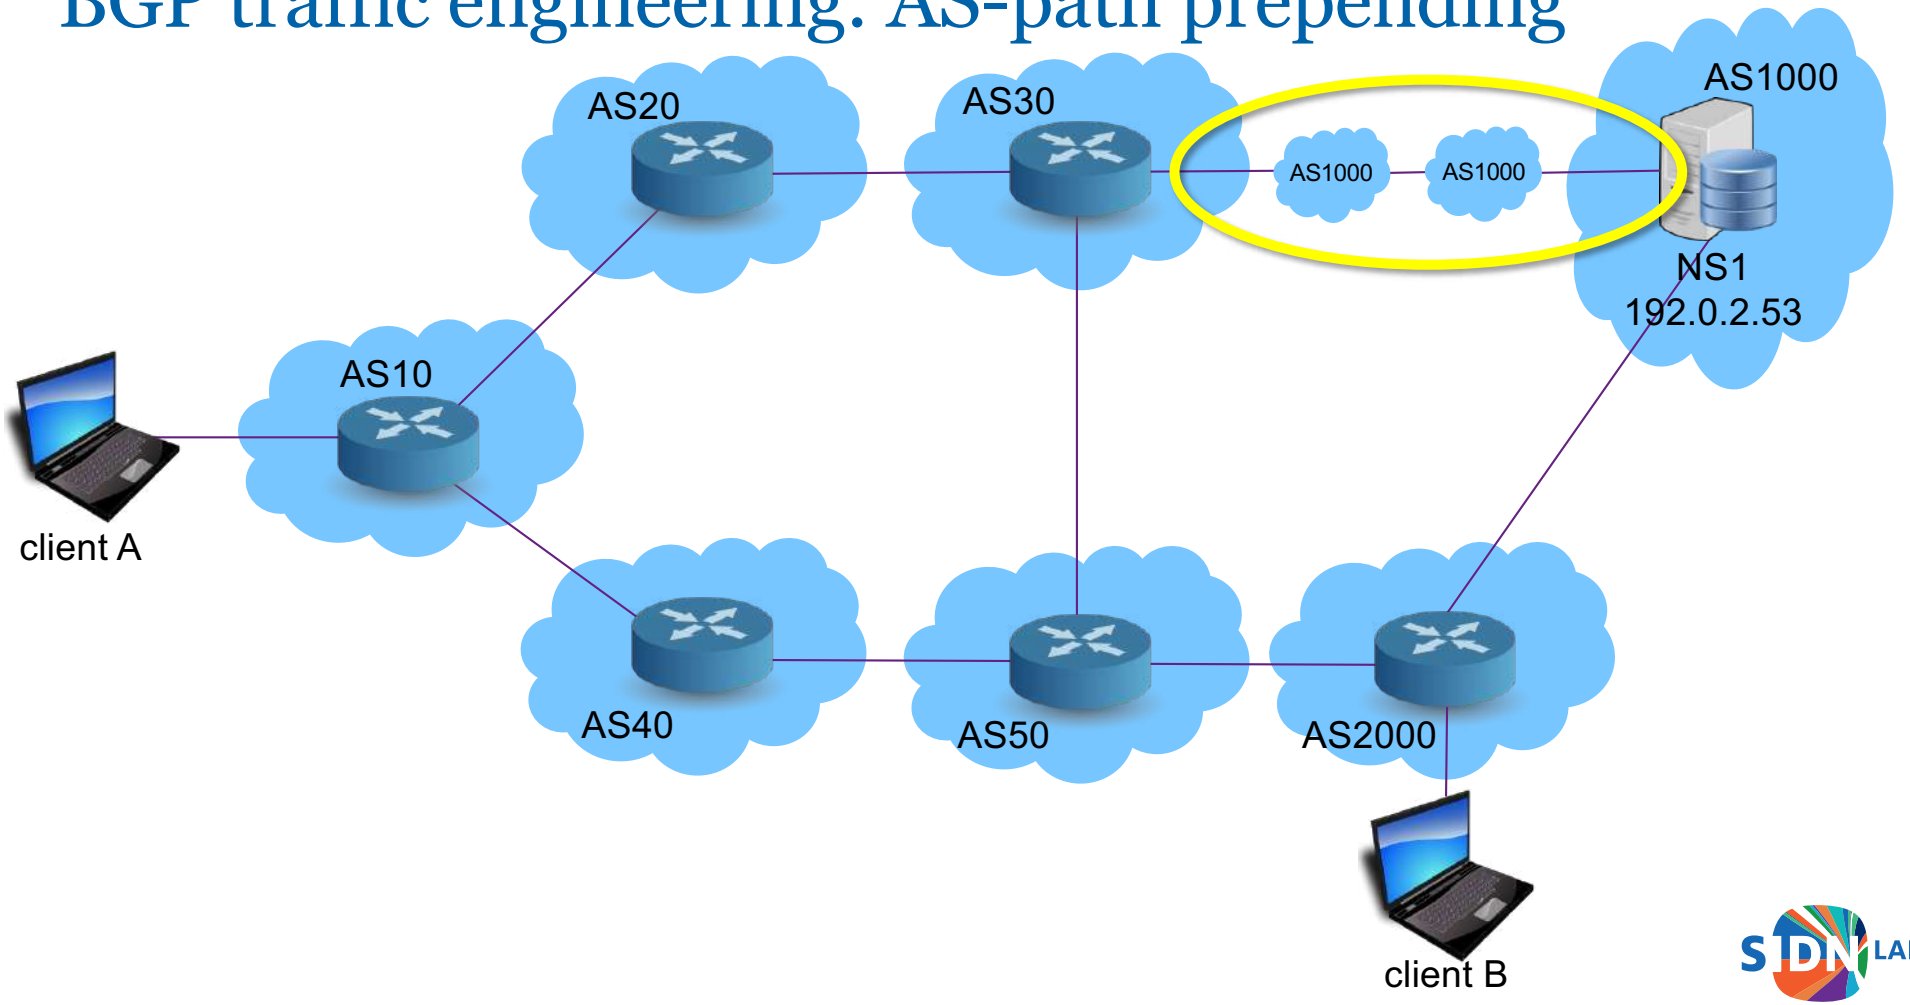

BGP traffic engineering: AS-path prepending

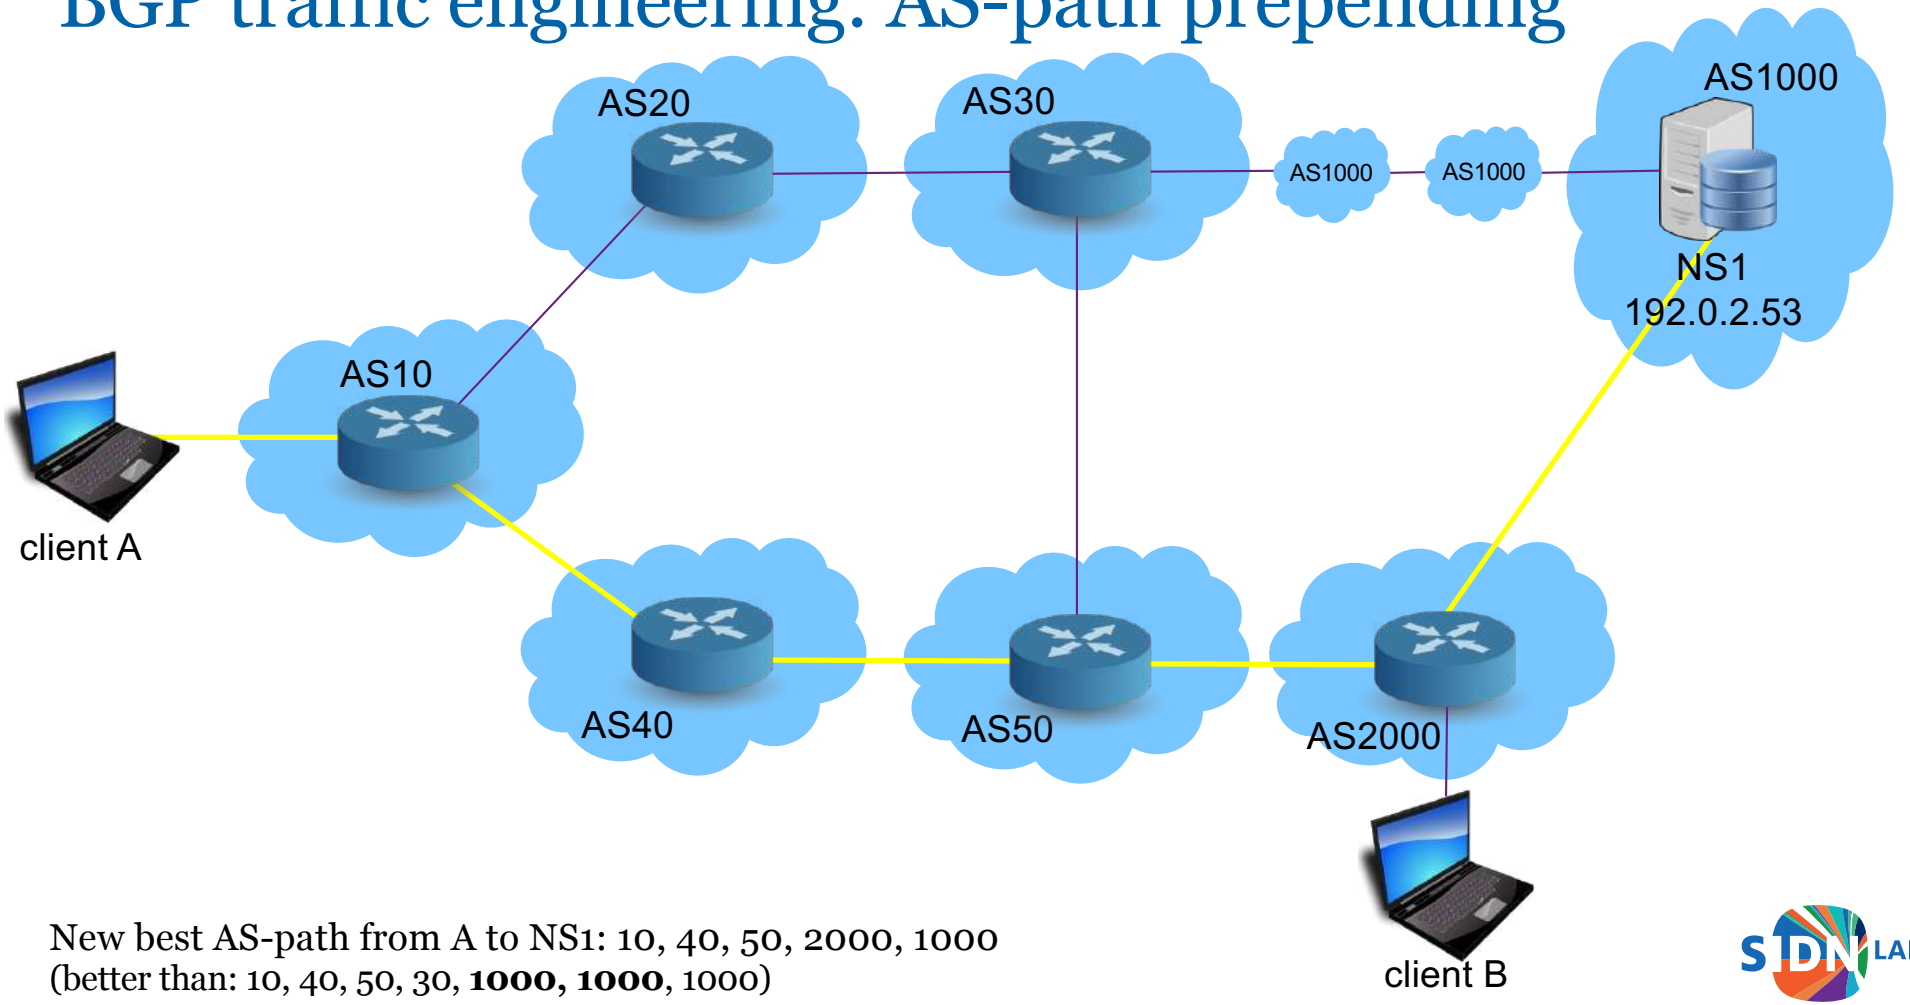

New best AS-path from A to NS1: 10, 40, 50, 2000, 1000
(better than: 10, 40, 50, 30, **1000, 1000**, 1000)

BGP Best Path Selection Algorithm

BGP – Routing Algorithm*

*According to RFC4271 – Implementations are vendor-specific

1. Check if *next hop* is reachable
2. ➔ **Choose route with the highest *Local Preference***
3. ➔ **Prefer the route with the shortest *AS path***
4. Prefer the route with the lowest *origin attribute*
5. ➔ **Prefer the route with the lowest *MED* value**
6. Prefer routes received from *eBGP* over *iBGP*
7. Prefer the nearest *exit* from your network (in terms of your internal routing protocol)
8. ➔ **Implementation dependent: Prefer *older* (= more stable) routes**
9. Prefer routes learned from the router with lower *router ID*
10. Prefer routes learned from the router with lower *IP address*

➔ = most important rules

This is where you prefer peering over upstream

Next hop reachable?	continue if "yes"
Local Preference	higher wins
AS path	shorter wins
Origin Type	IGP over EGP over incomplete
MED	lower wins
eBGP, iBGP	eBGP wins
Network exit	nearest wins
Age of route	older wins
Router ID	lower wins
Neighbor IP	lower wins

Version 1.0

Source: https://www.de-cix.net/en/resources/bgp-basics

SDN LABS

Traffic engineering with BGP communities

- *Transitive attribute tags* that can be applied to inbound or outbound prefixes to achieve a particular goal.
- For example: local pref adjustments, geographical adjustments, AS-path prepending or blackholing.
- No universal definitions, except the so-called *well-known*

```
route-server> show ip bgp 194.0.5.0/24
BGP routing table entry for 194.0.5.0/24
Paths: (23 available, best #18, table Default-IP-Routing-Table)
  Not advertised to any peer
  20473 210004
    206.53.202.75 from 216.218.252.190 (216.218.252.167)
      Origin IGP, metric 0, localpref 100, valid, internal
      Large Community: 6695:1000:1 20473:0:3021840115 210004:3000:1004
      Originator: 216.218.252.167, Cluster list: 216.218.252.190
      Last update: Wed Apr 15 16:06:36 2020
```

RPKI (in a nutshell)

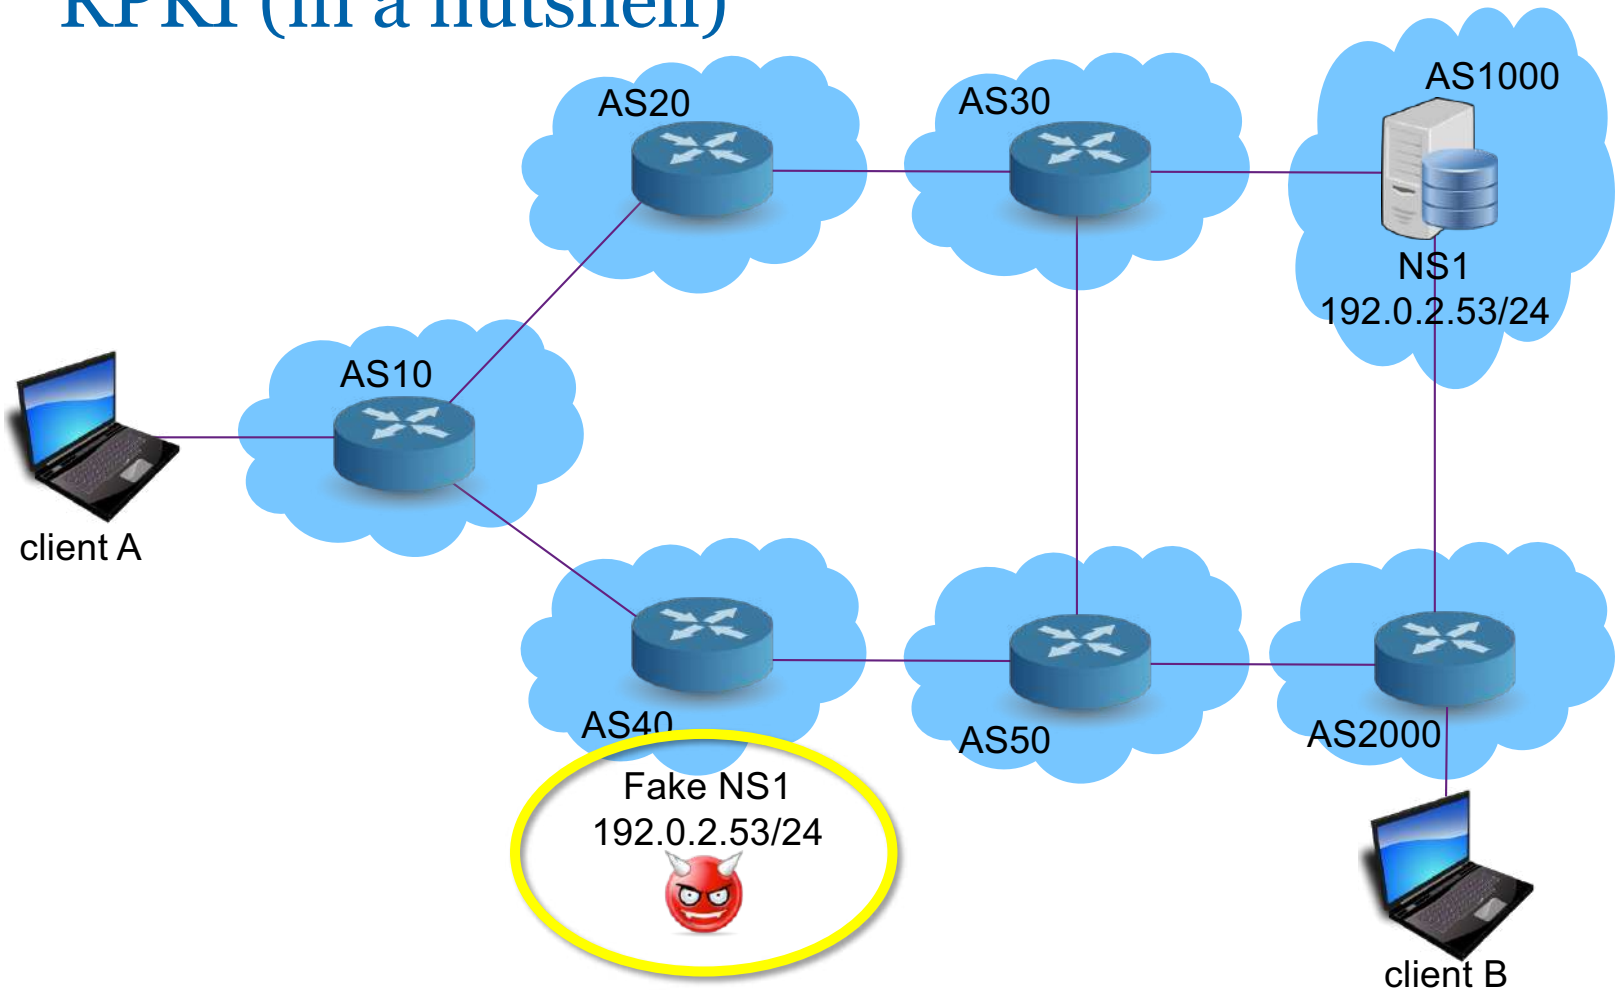

RPKI (in a nutshell)

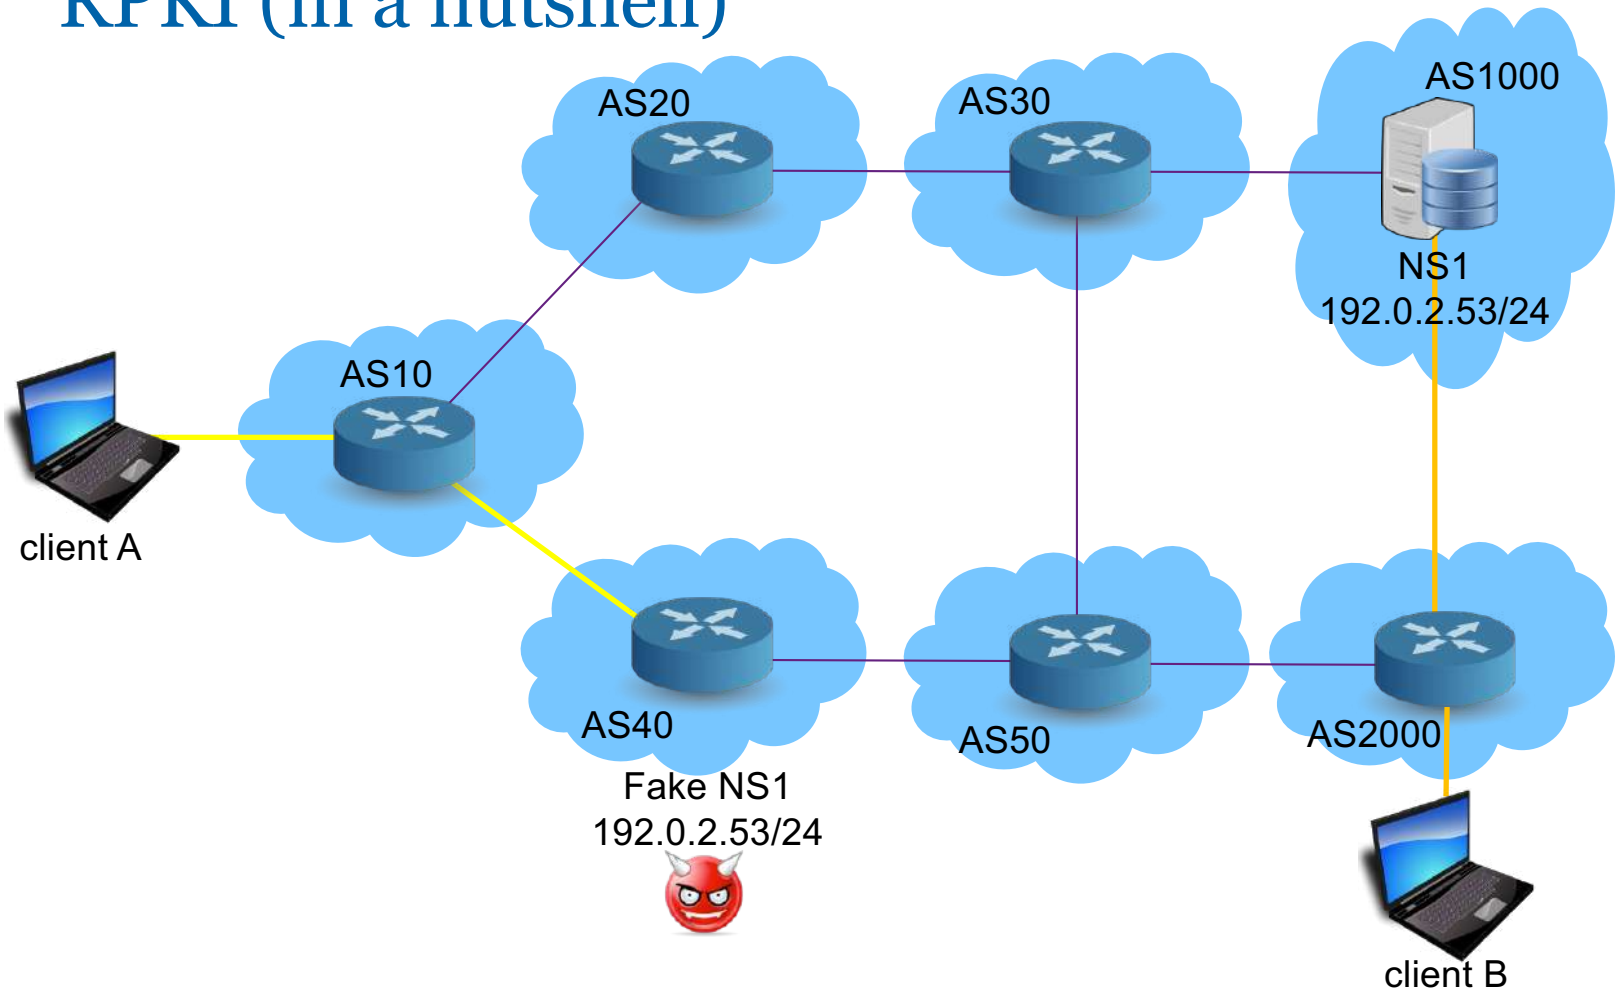

RPKI: Resource Public Key Infrastructure

- A public key infrastructure to secure BGP
- Resource certification of IP prefixes / ASN combination
- Prevents (to some extent) route hijacking
- There are two sides: publishing ROAs and validating them
- Origin validation, **not** path validation (that's BGPSEC, still in the works)

Try your own ISP: https://isbgpsafeyet.com/

SIDN LABS

Anycast and why it's a good idea

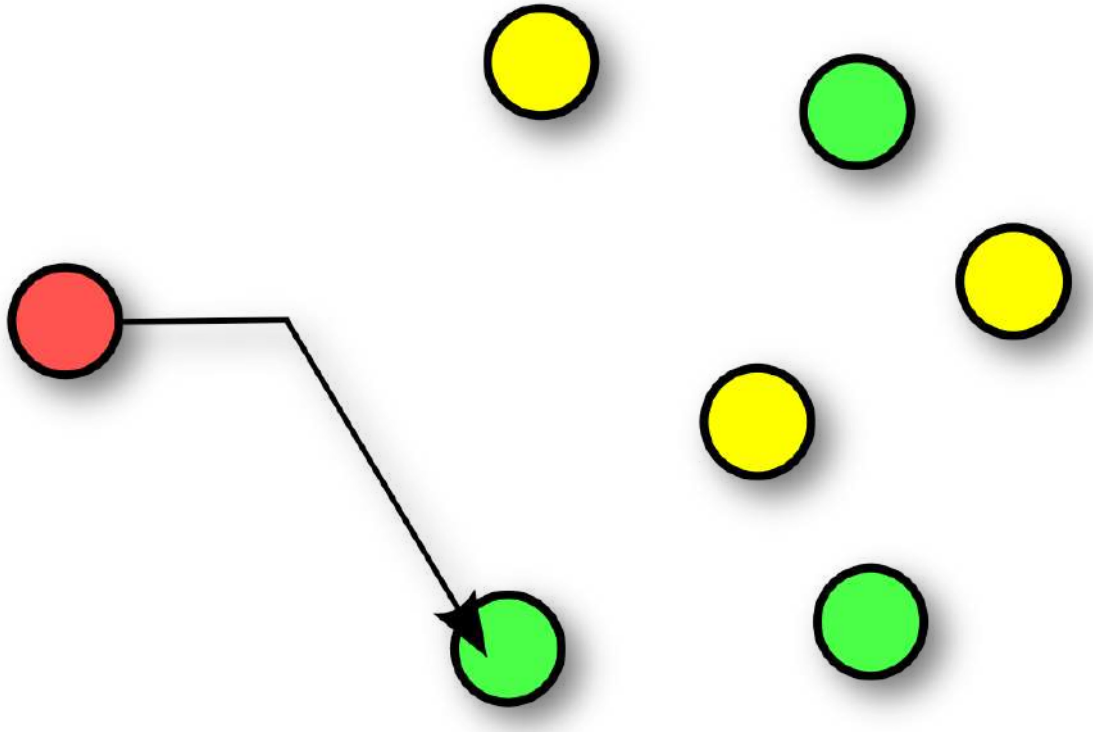

SDN LABS

Challenge 1: RTT

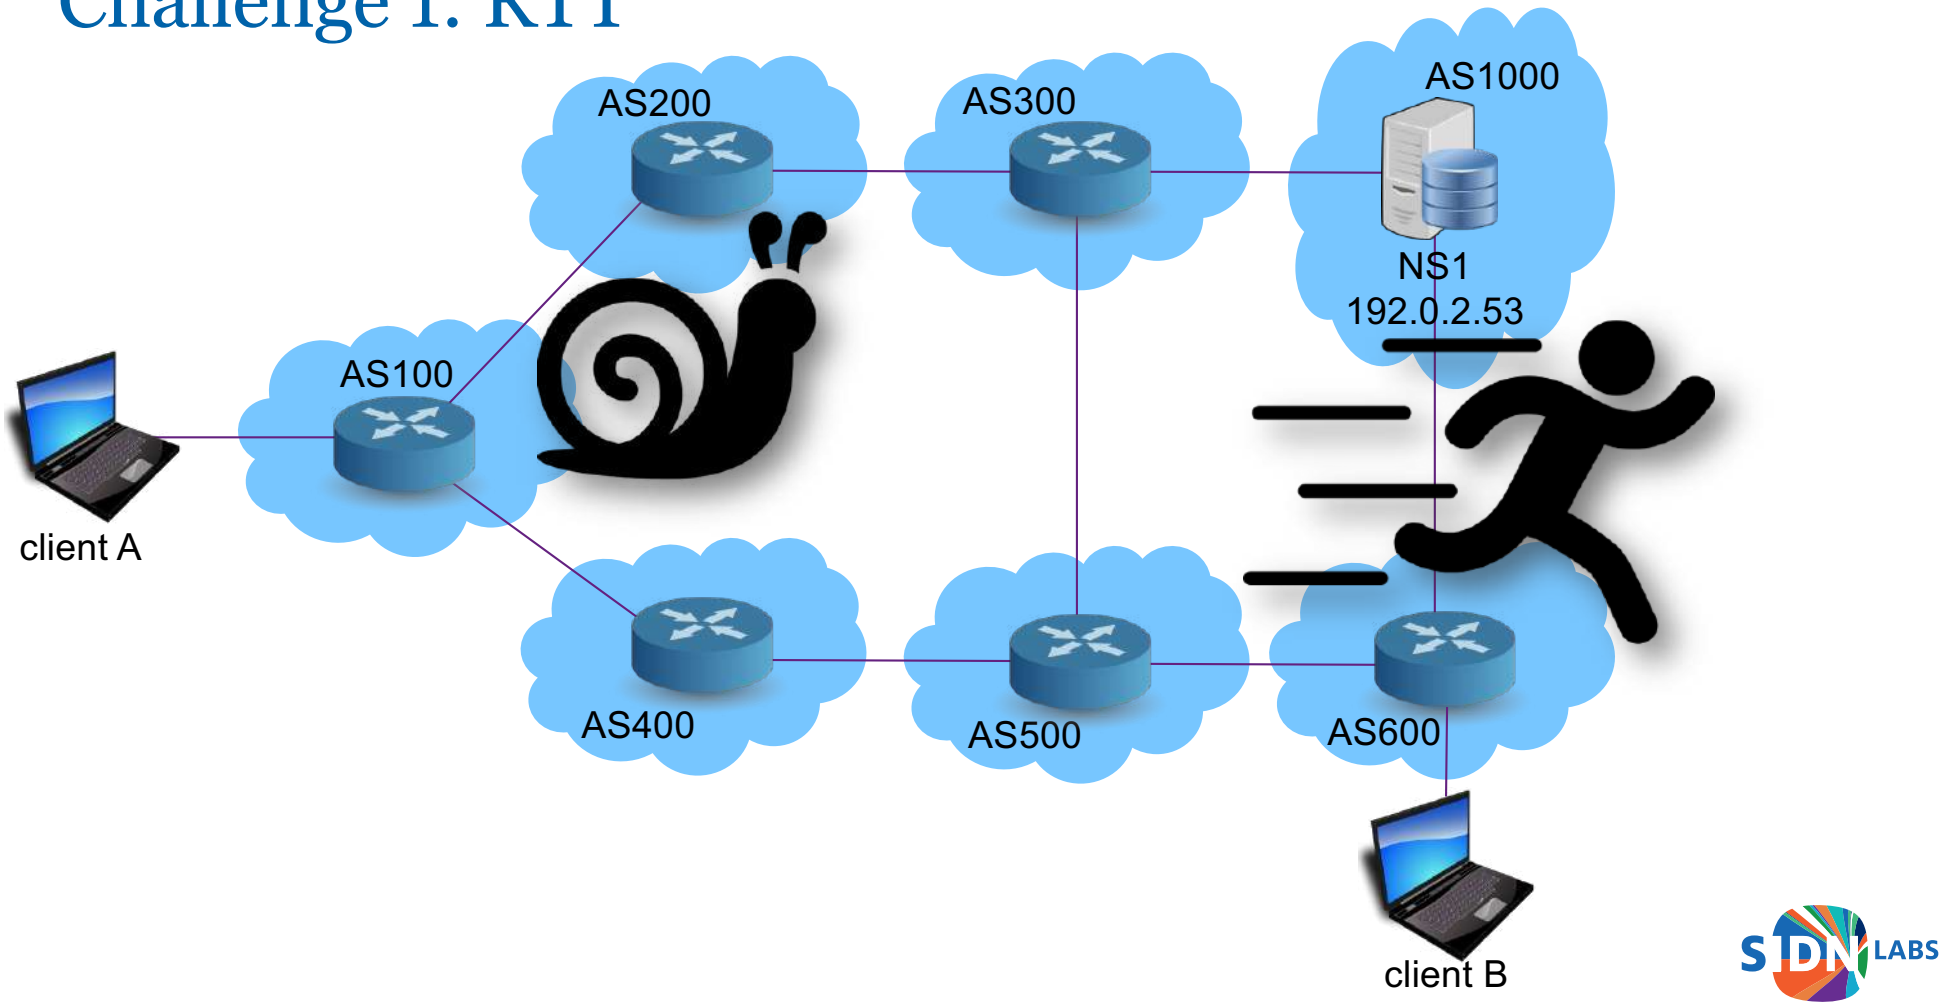

Challenge 2: DDoS

https://stats.ams-ix.net/nbip.html

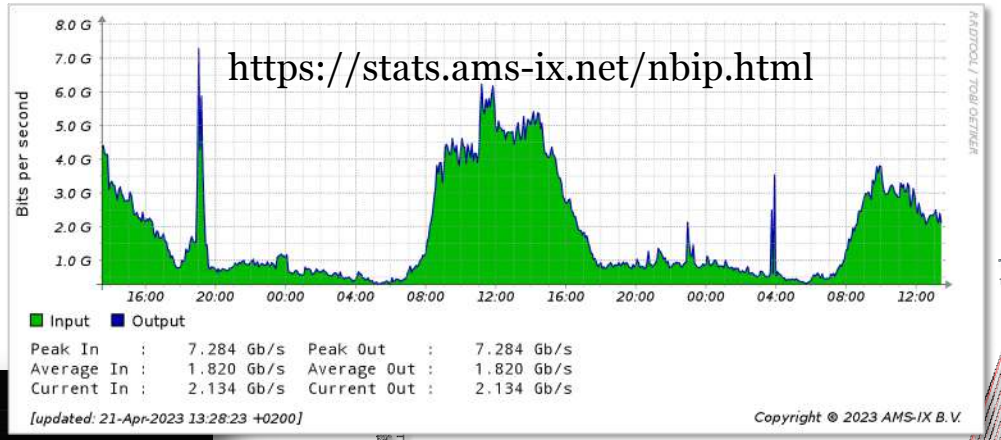

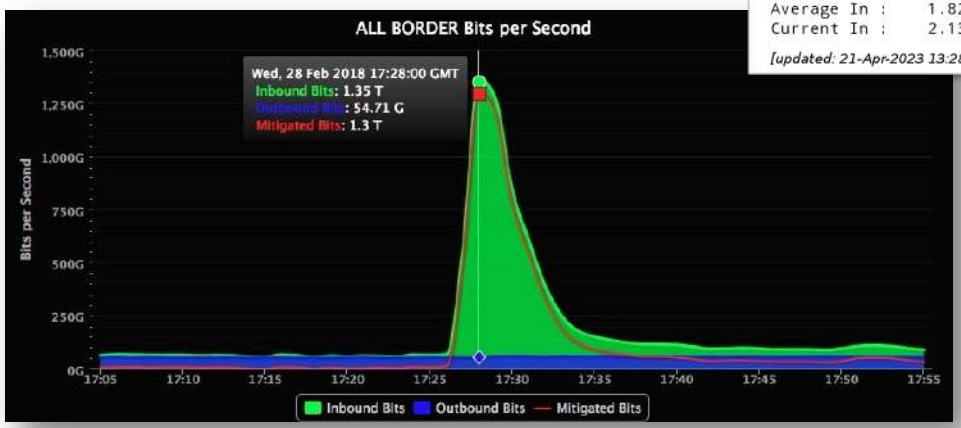

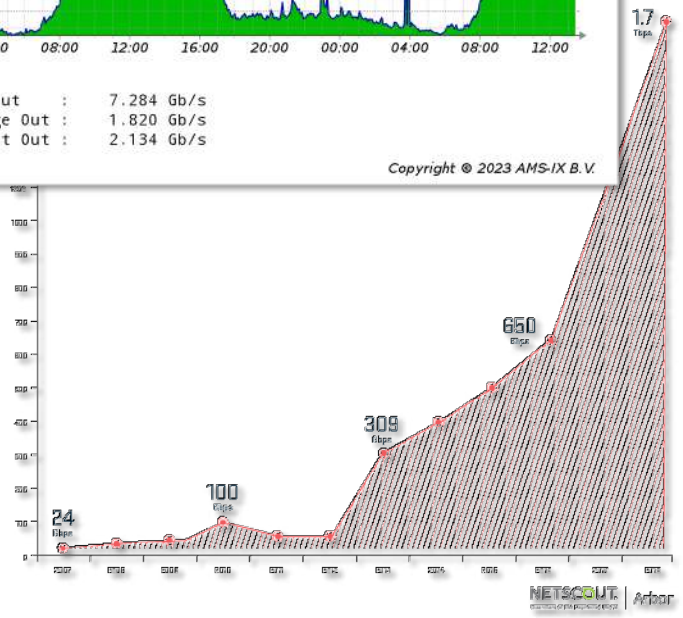

Main reason: IoT devices

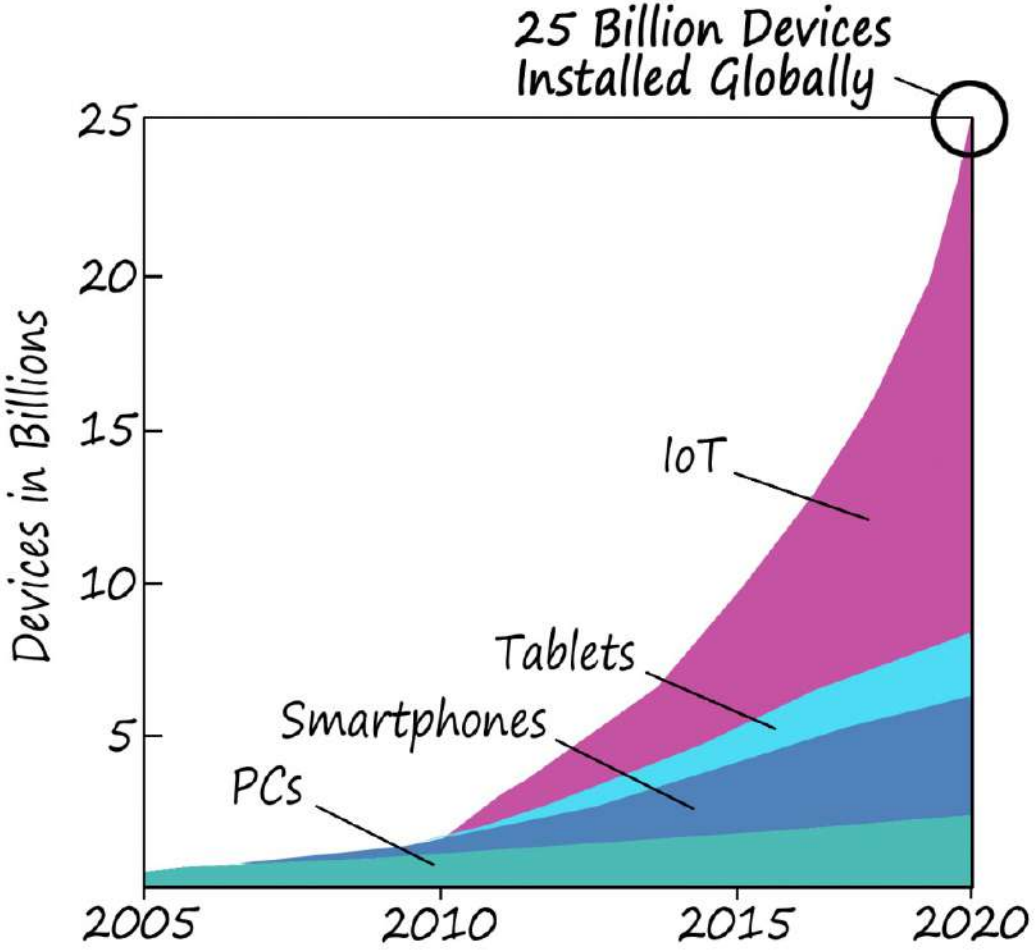

IoT botnets

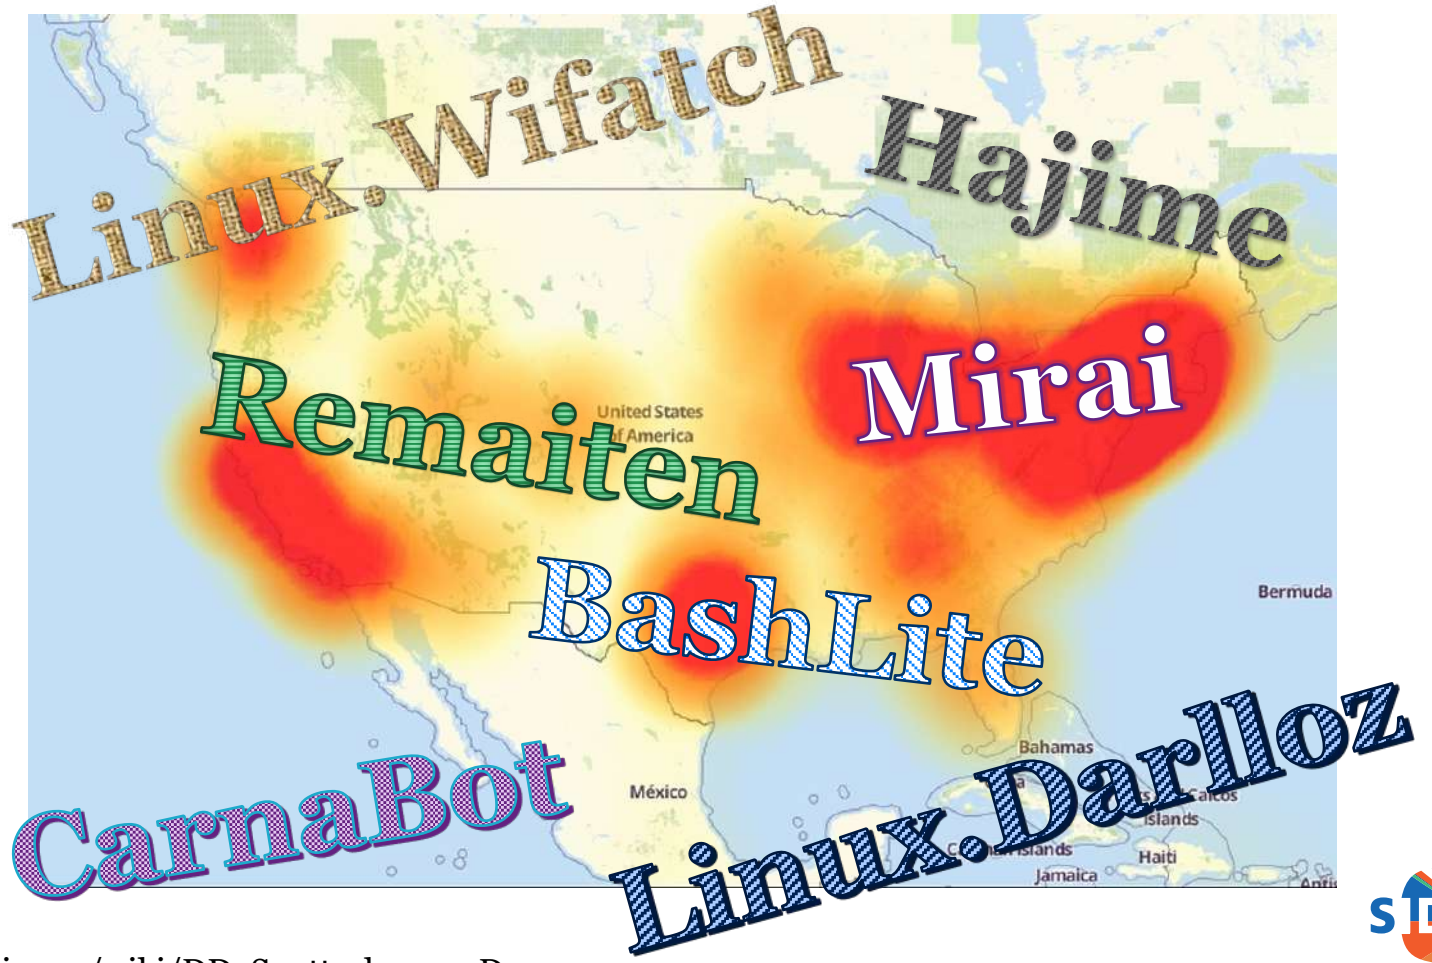

Bron: https://en.wikipedia.org/wiki/DDoS_attacks_on_Dyn

DNS global anycast

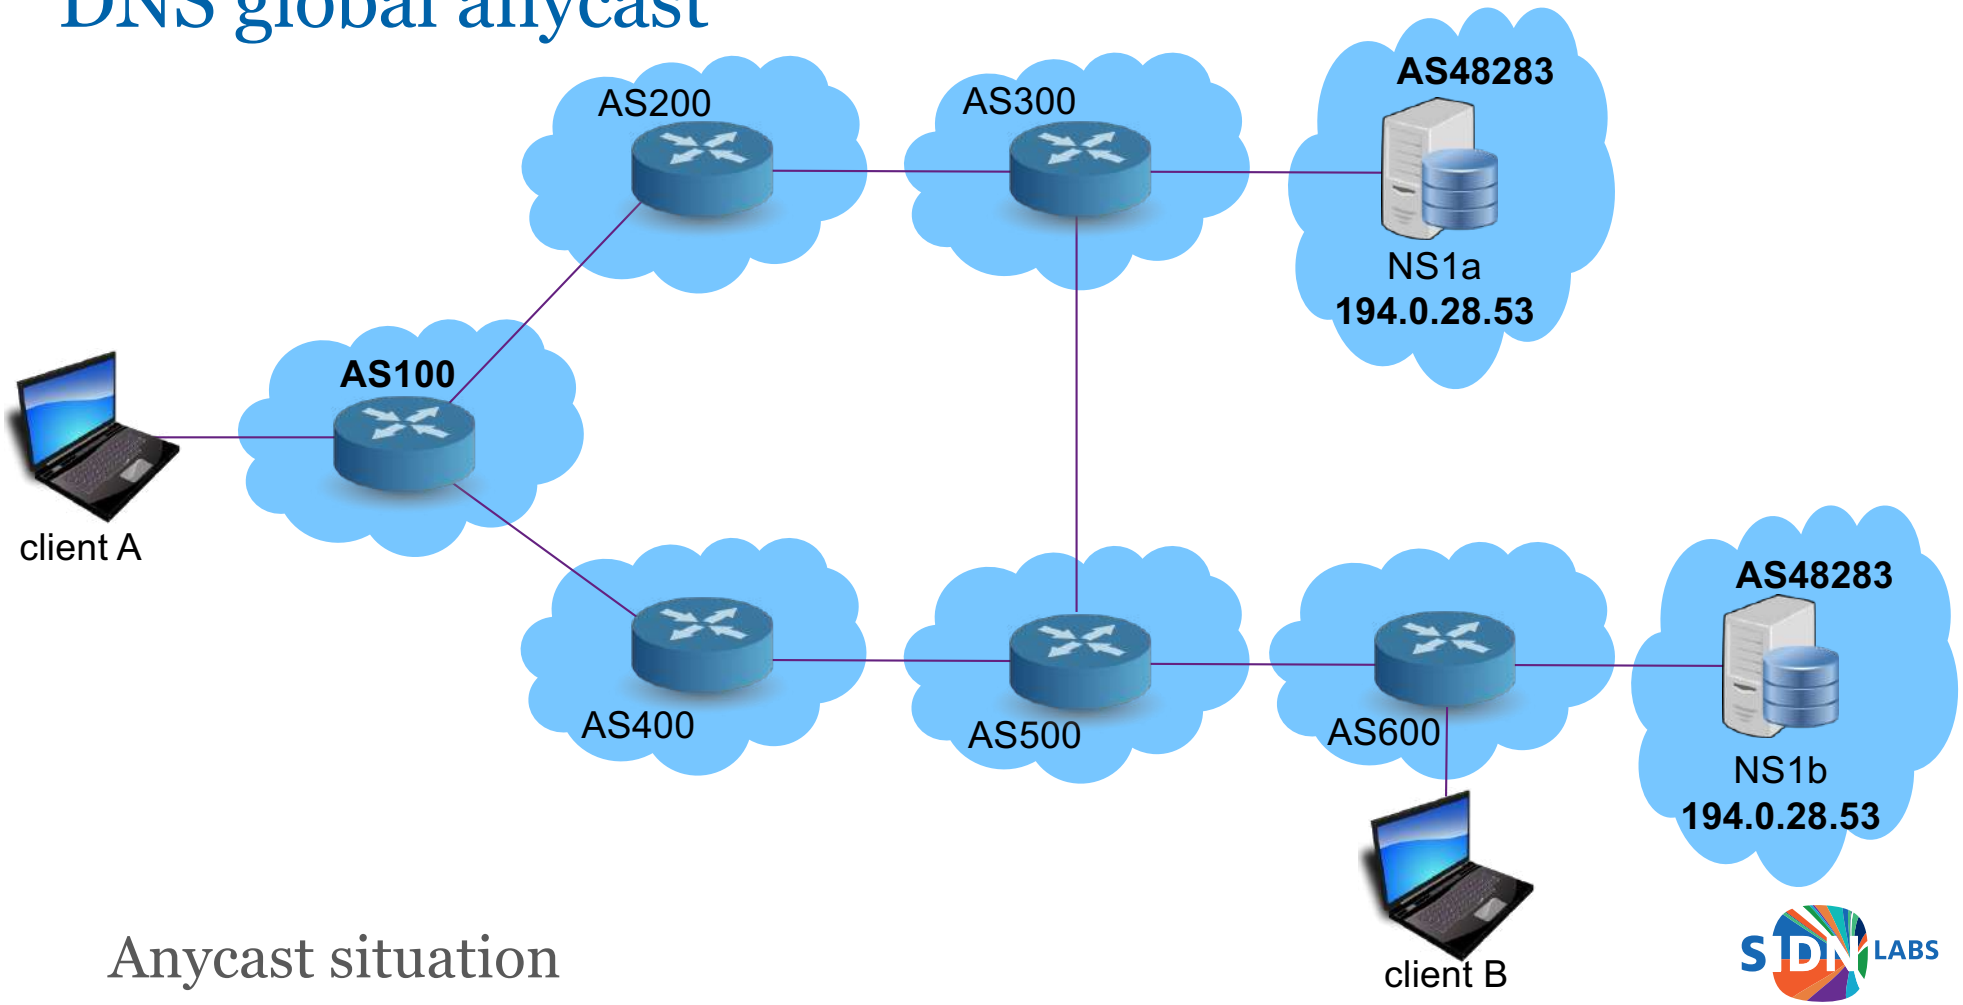

Anycast situation

DNS global anycast

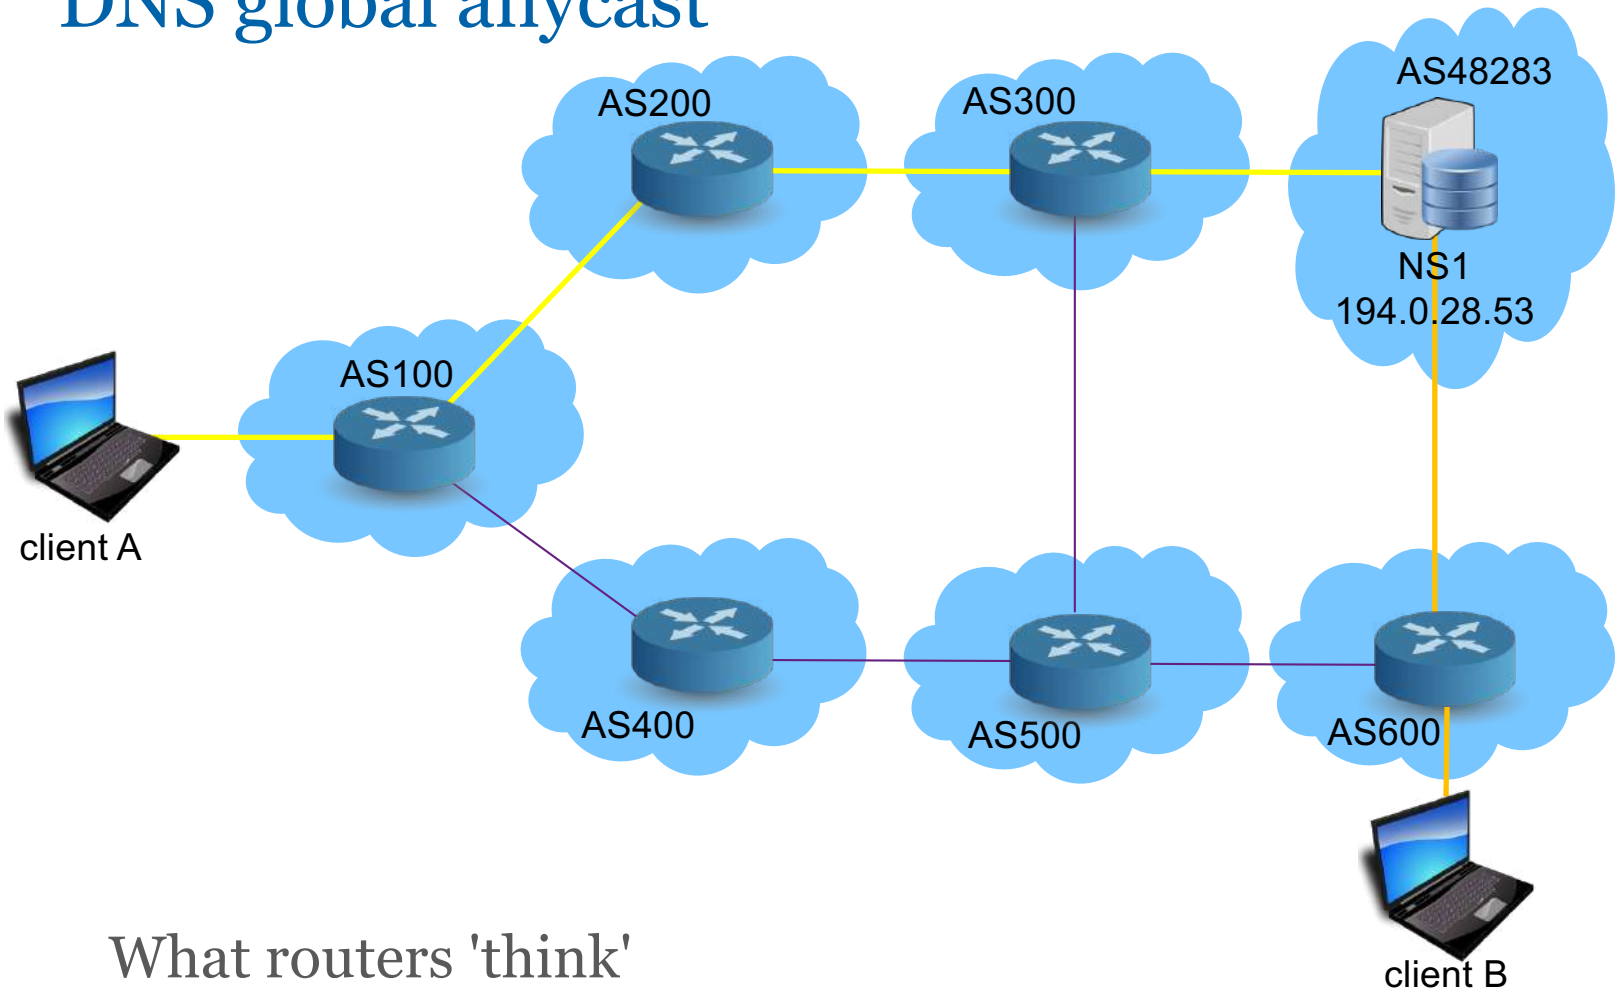

What routers 'think'

DNS global anycast

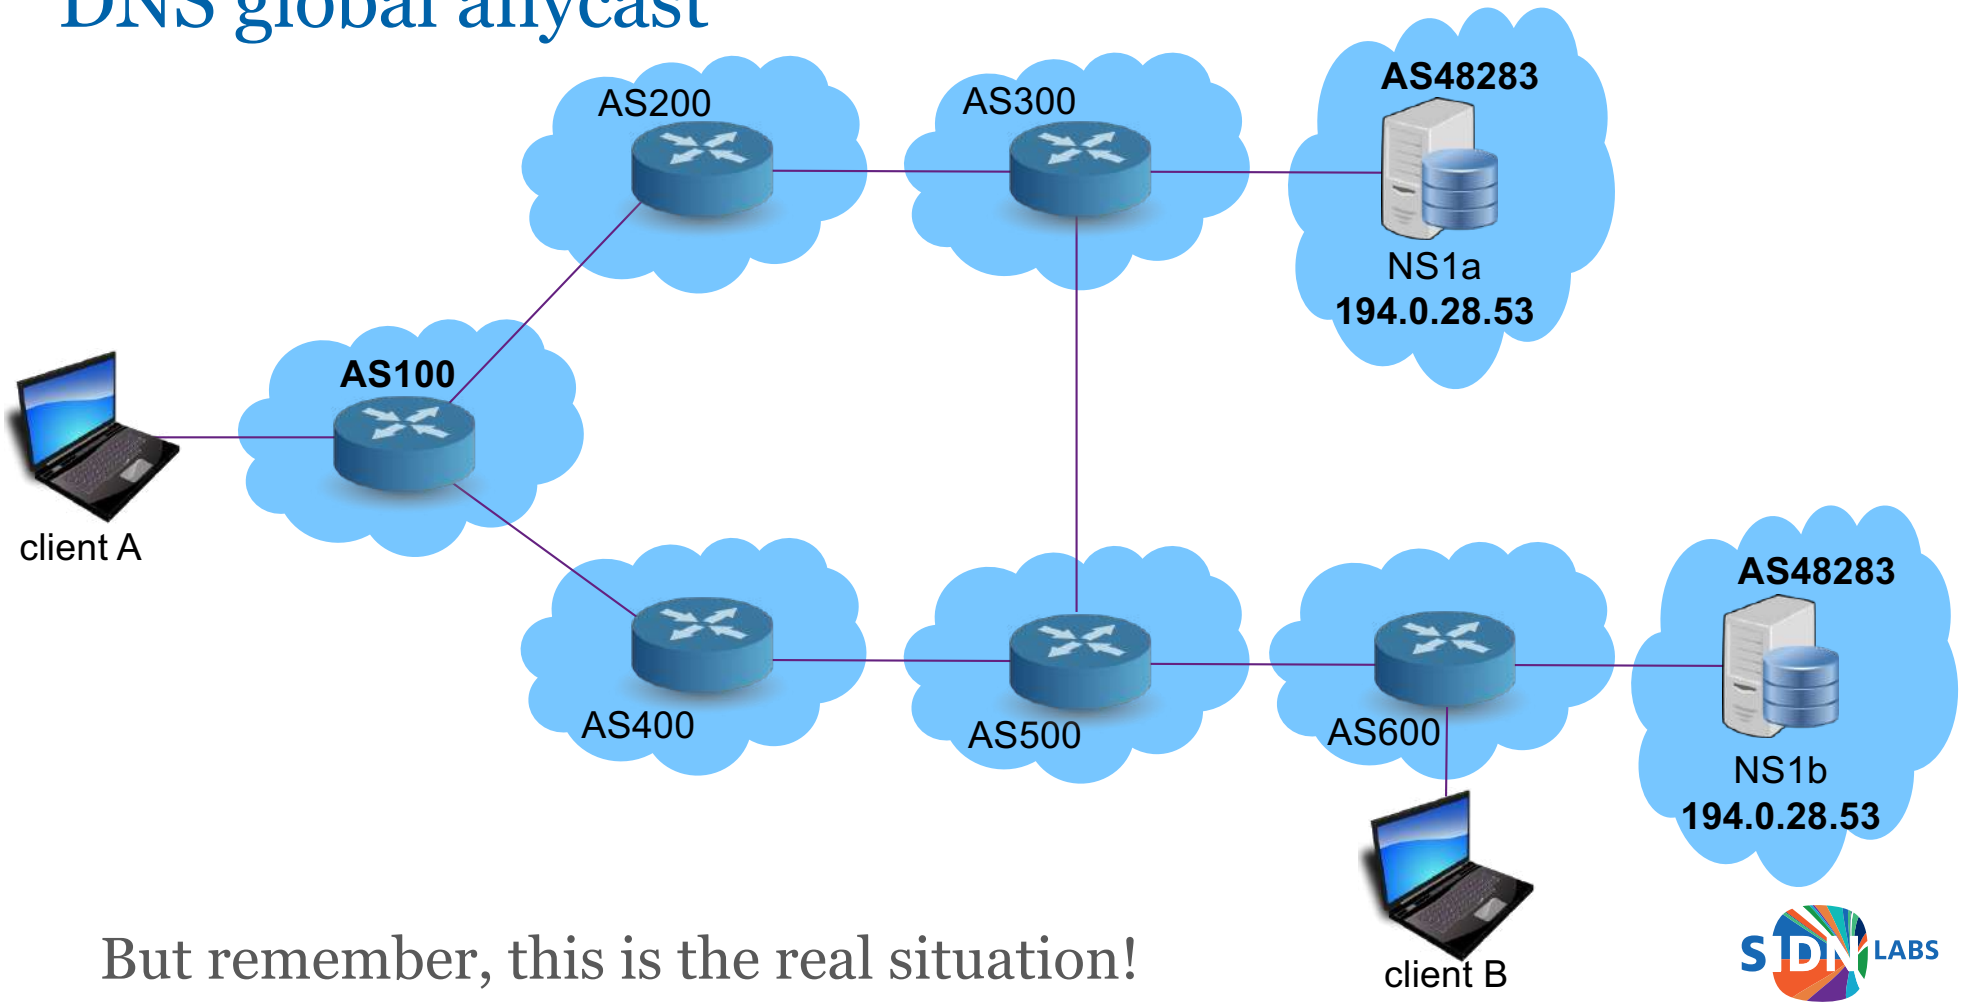

But remember, this is the real situation!

DNS global anycast

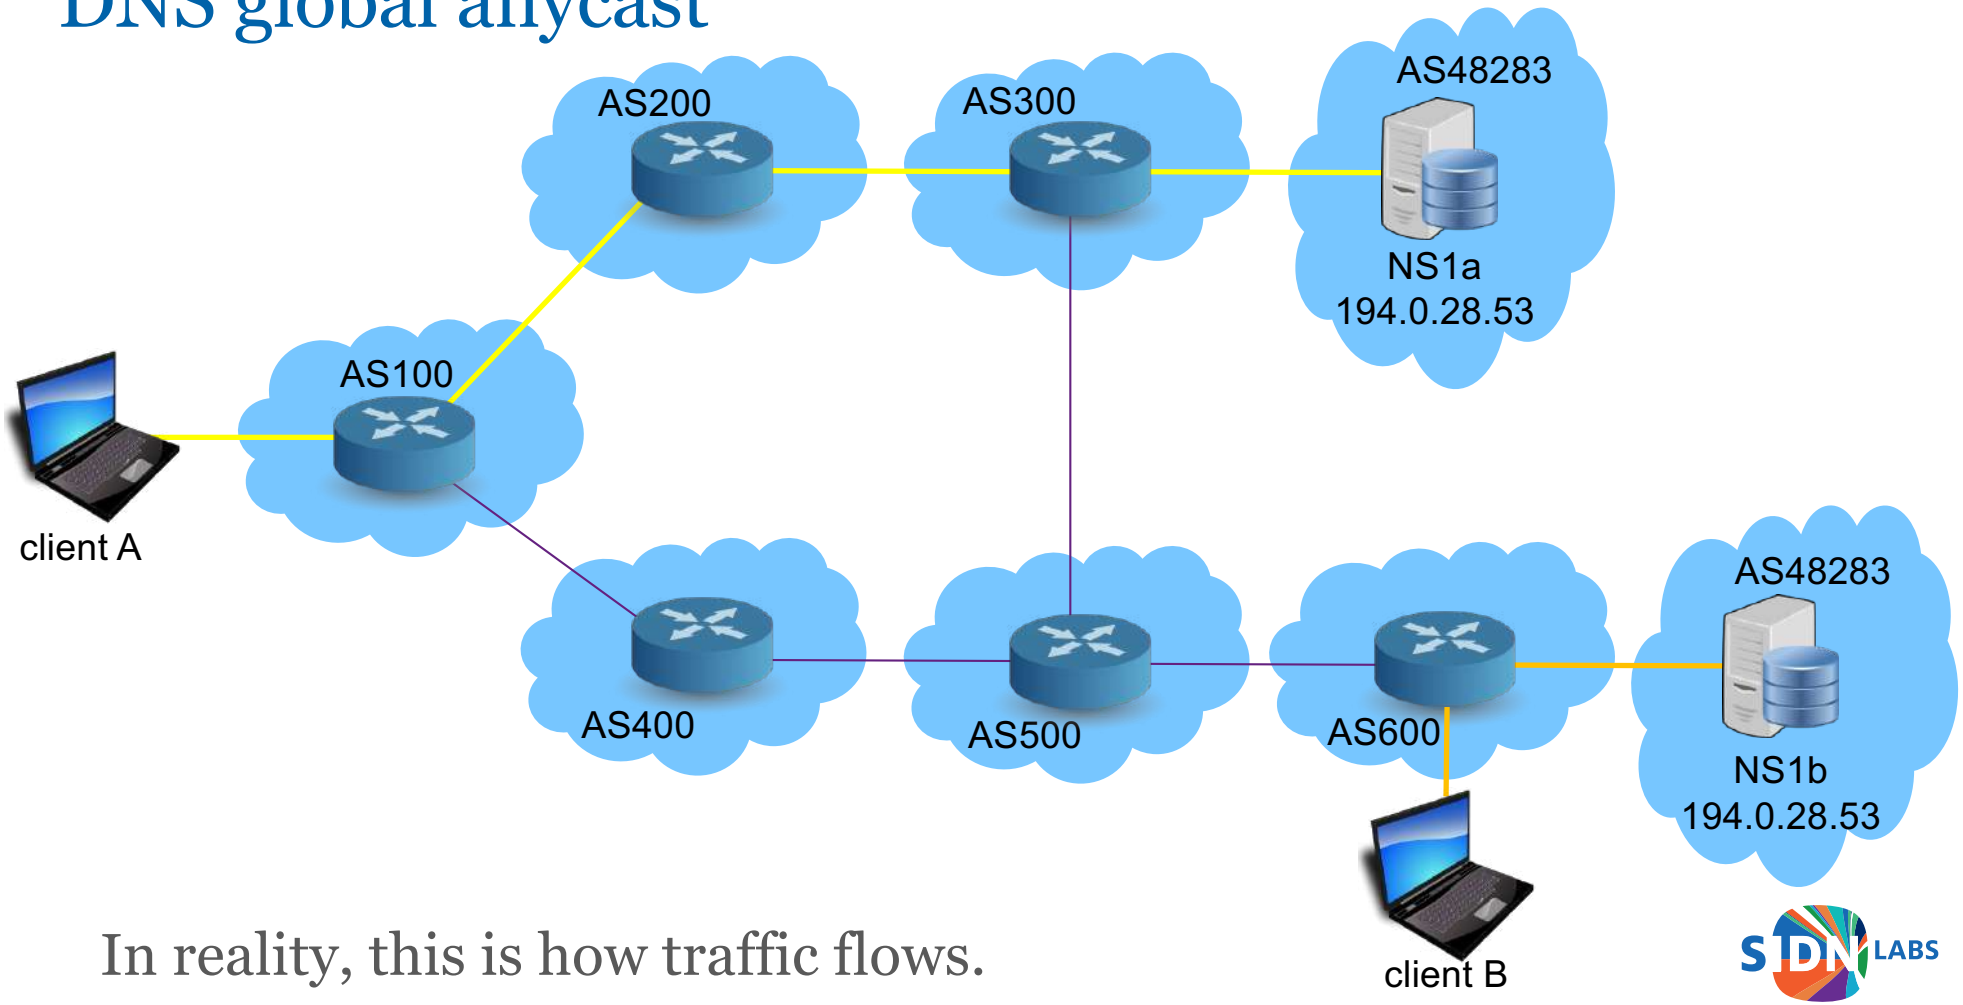

In reality, this is how traffic flows.

DNS global anycast

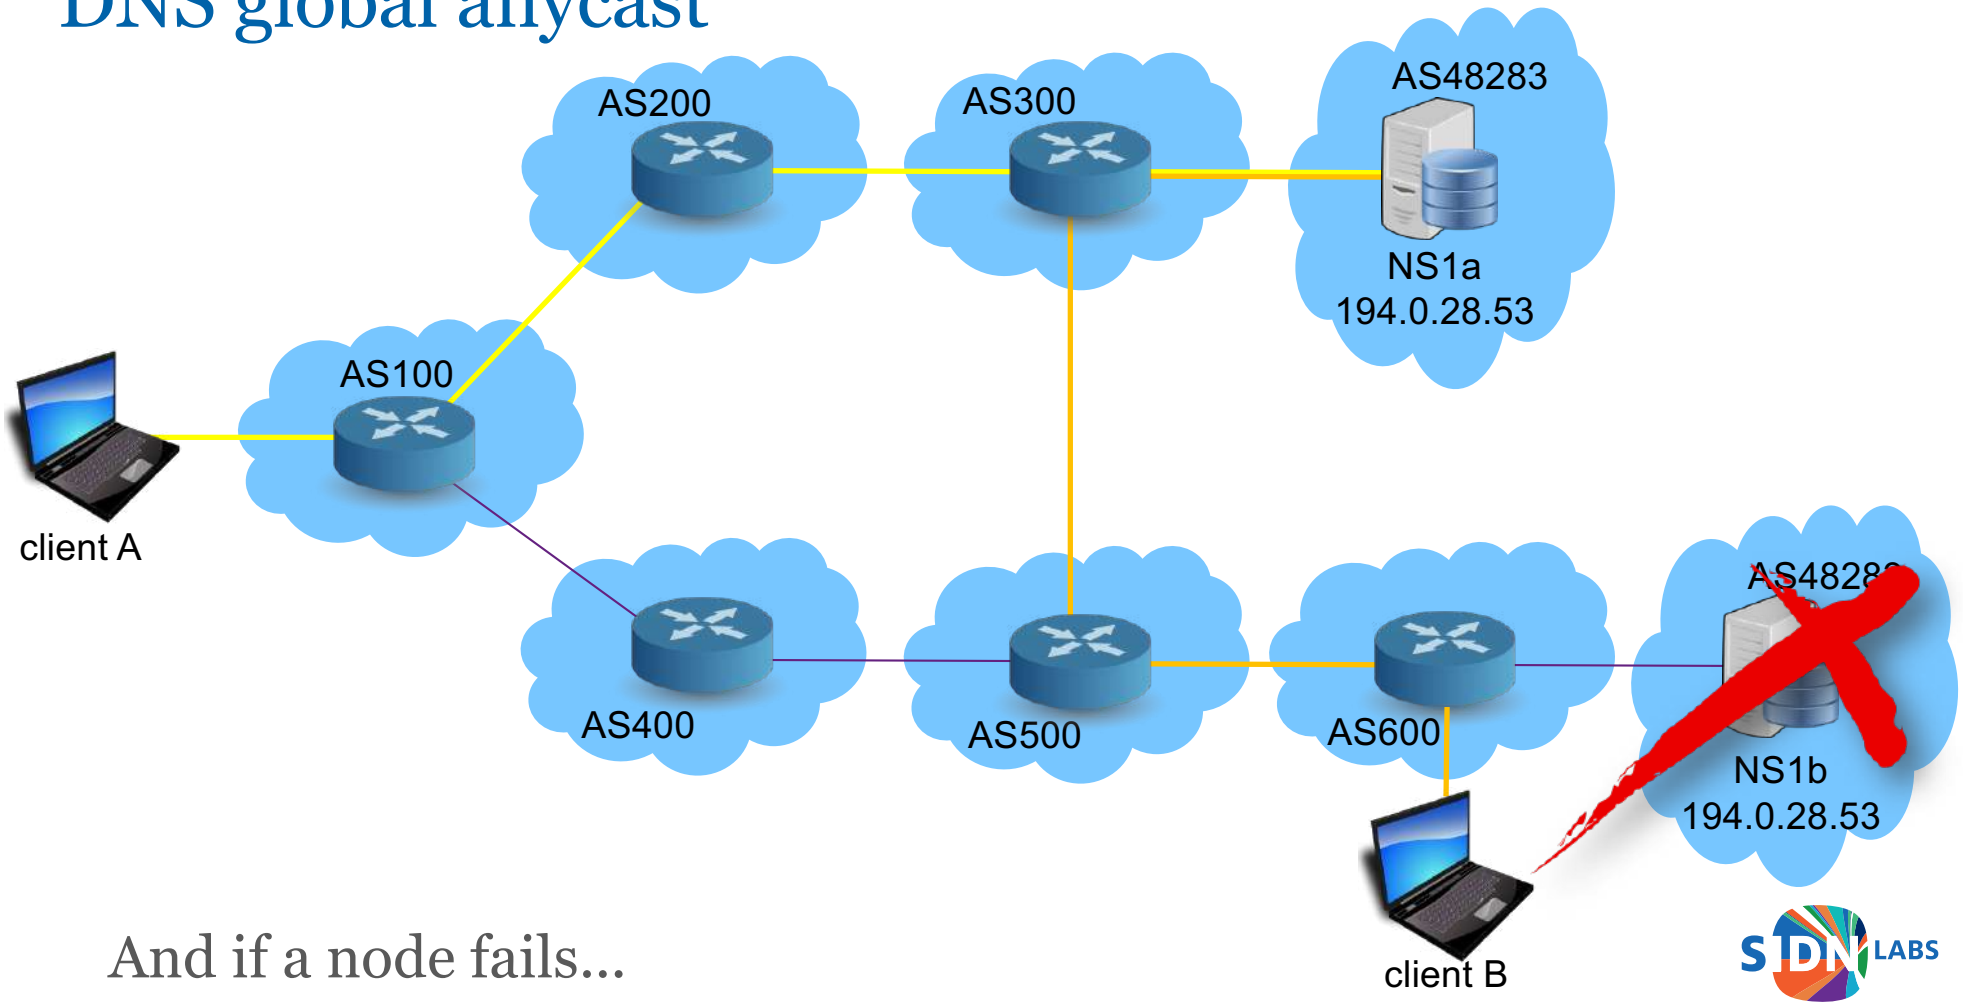

And if a node fails...

DNS global anycast (for .)

1915 servers!

http://www.root-servers.org/

SDN LABS

BGP catchment

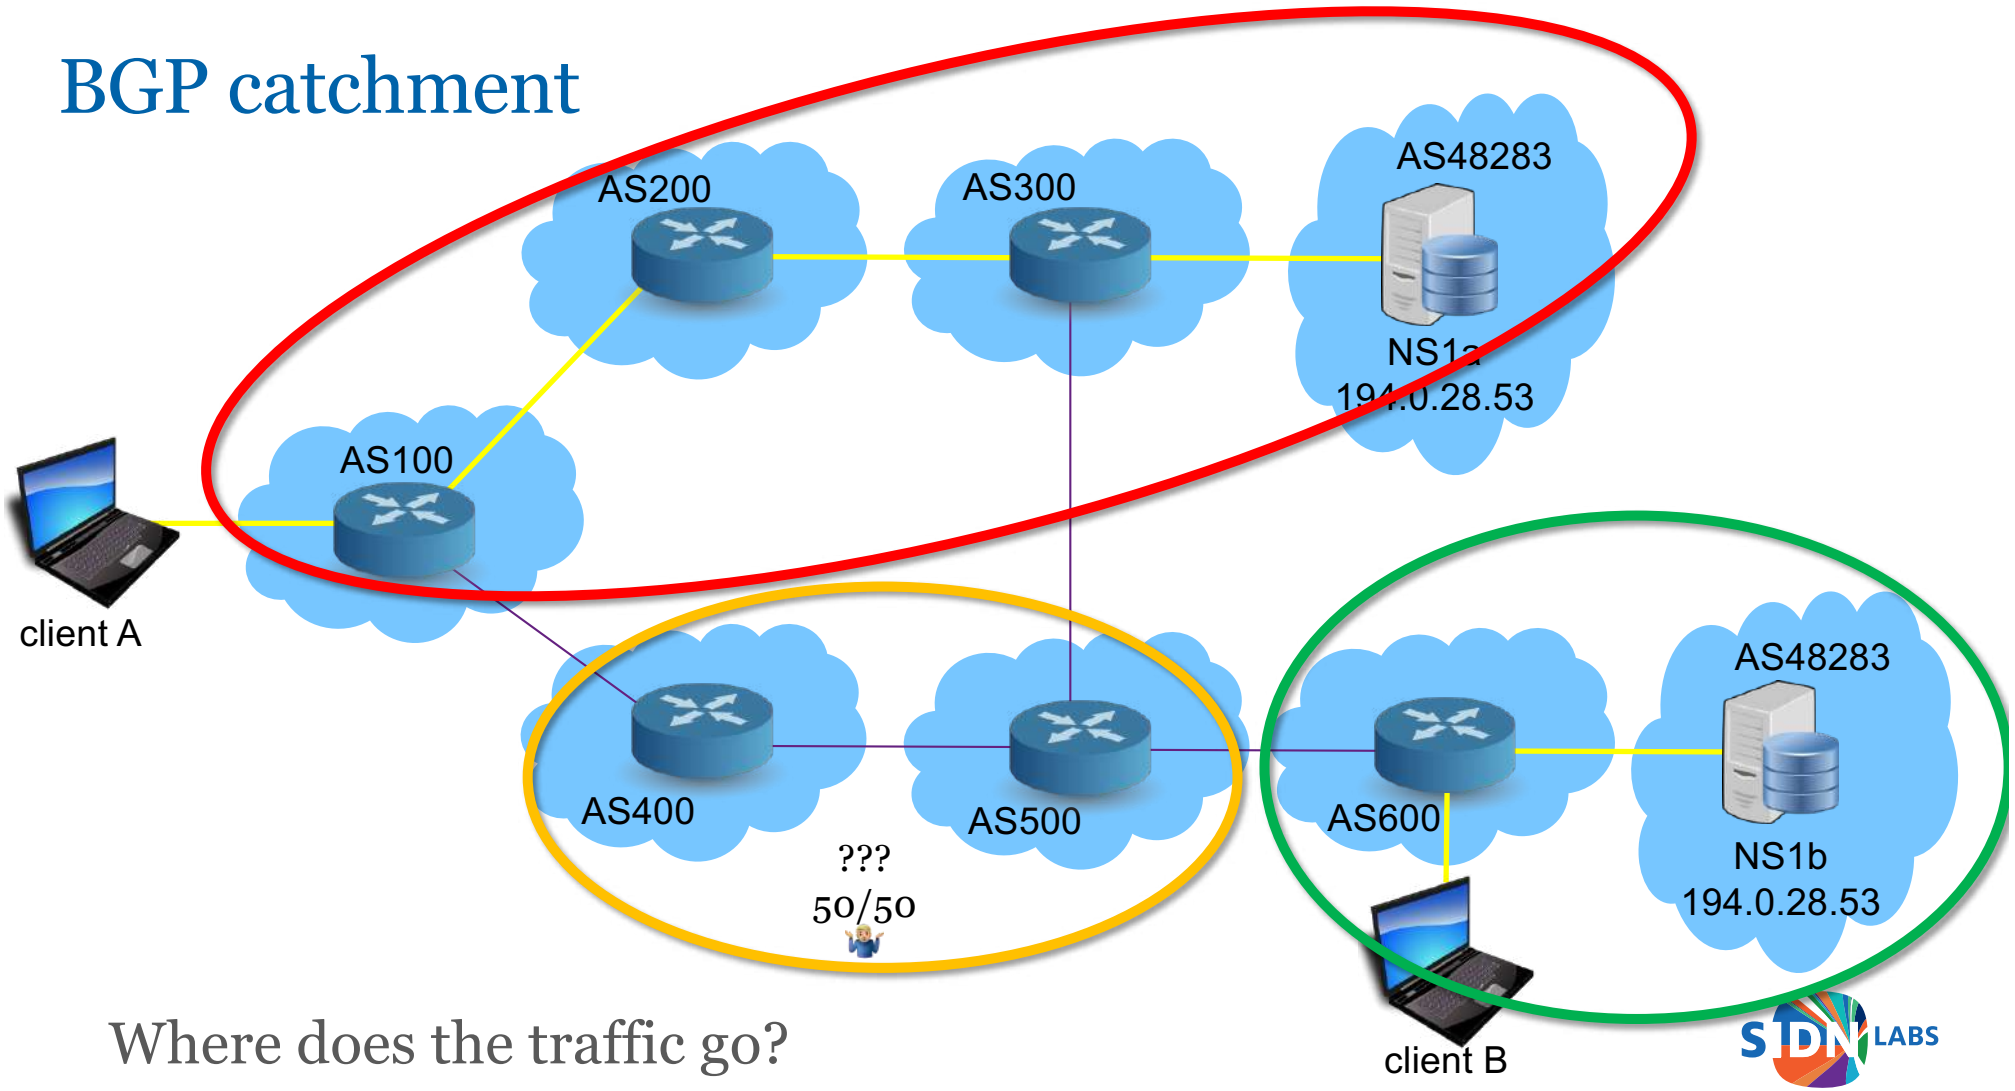

Where does the traffic go?

BGP catchment Santiago ✅

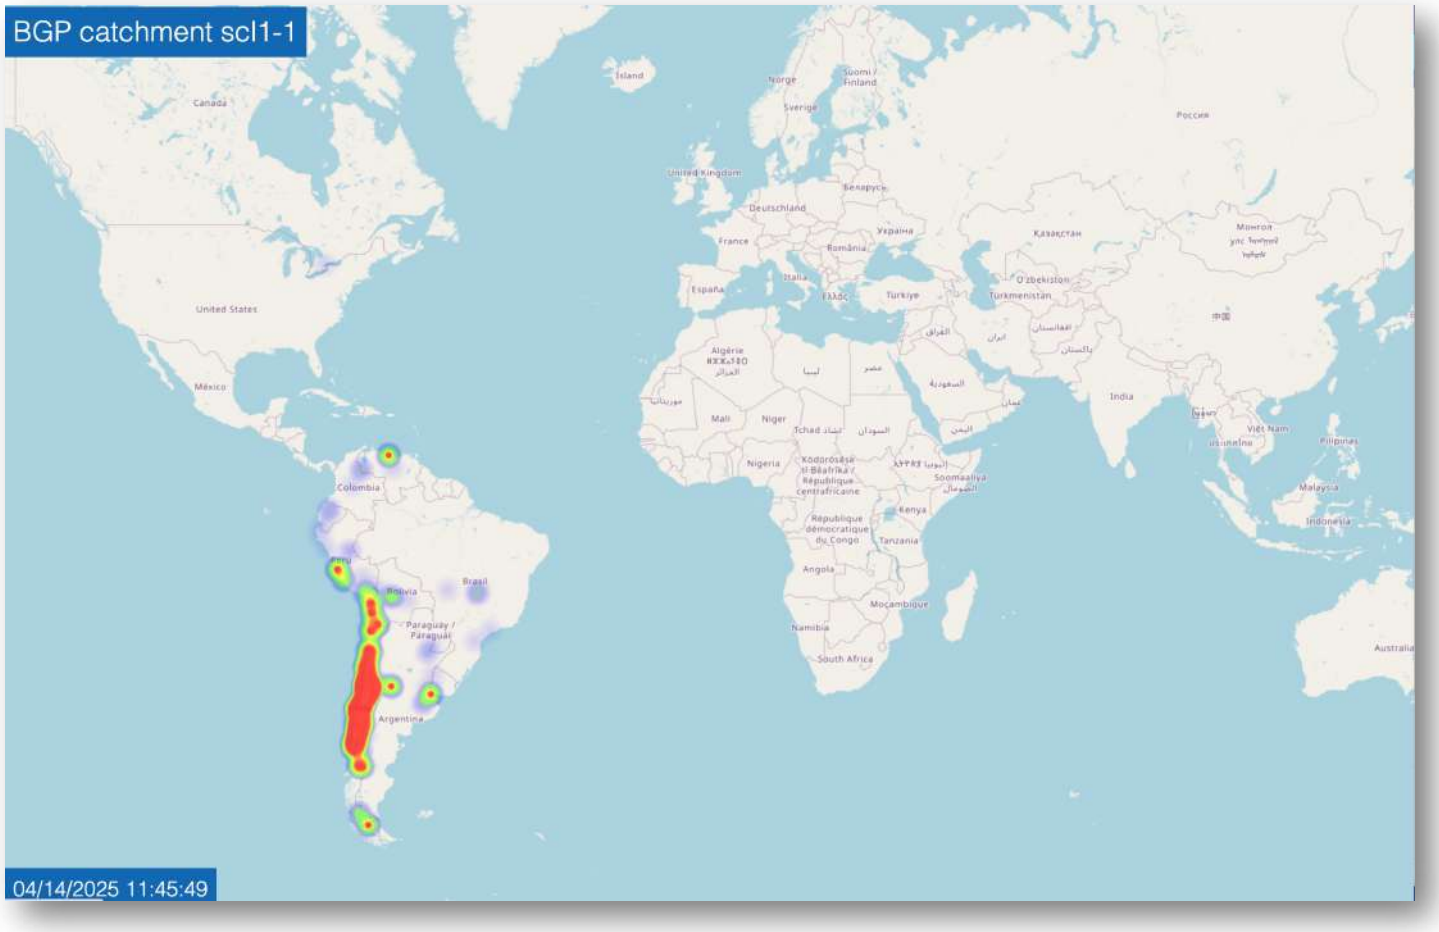

BGP catchment Frankfurt ⚠️

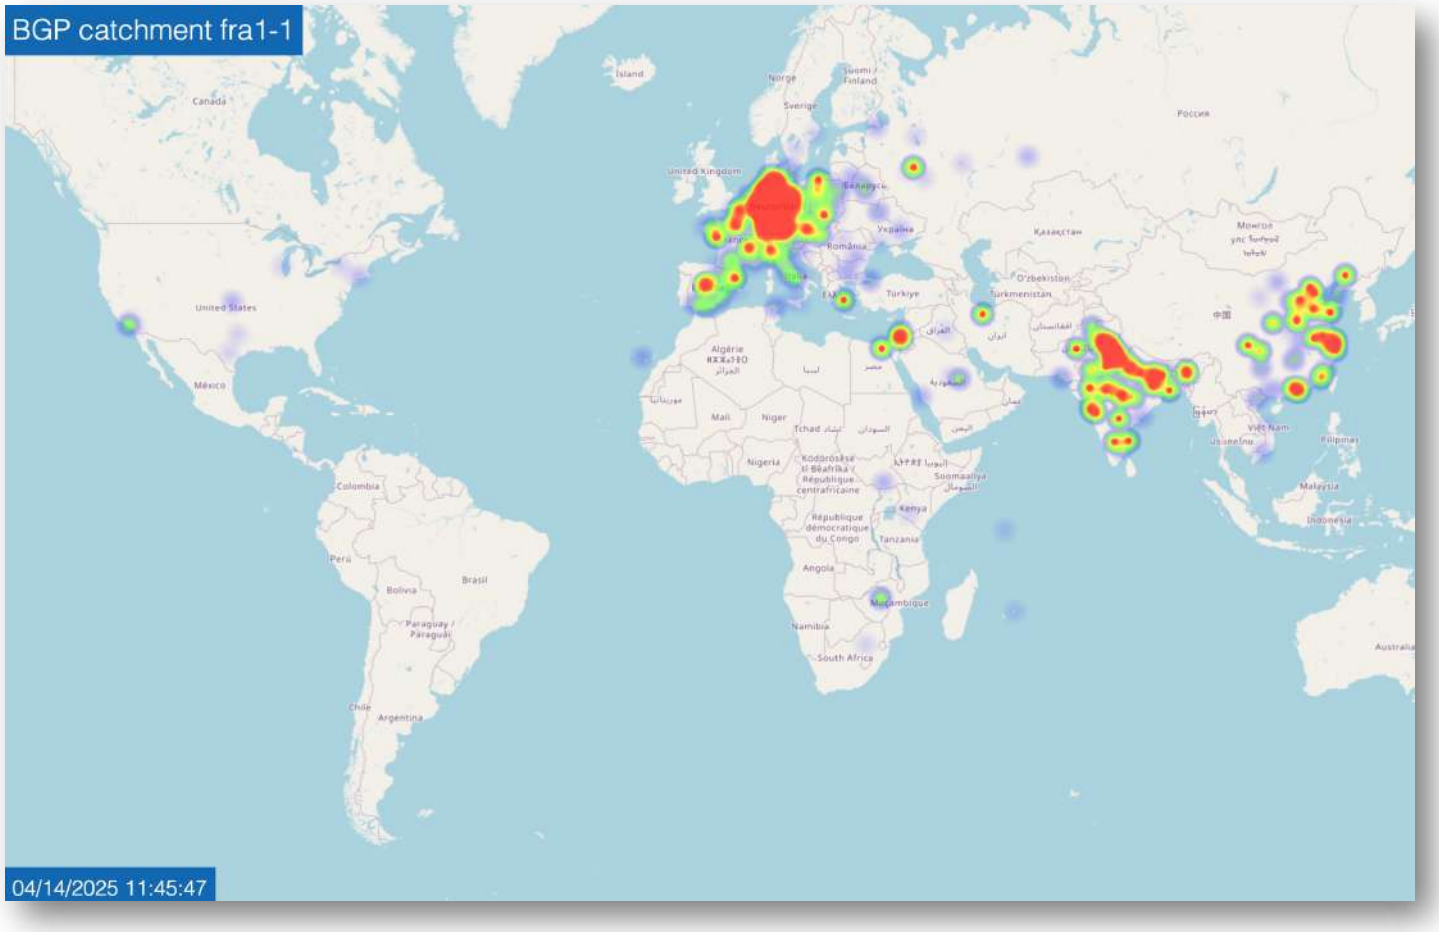

‘Scenic’ routing

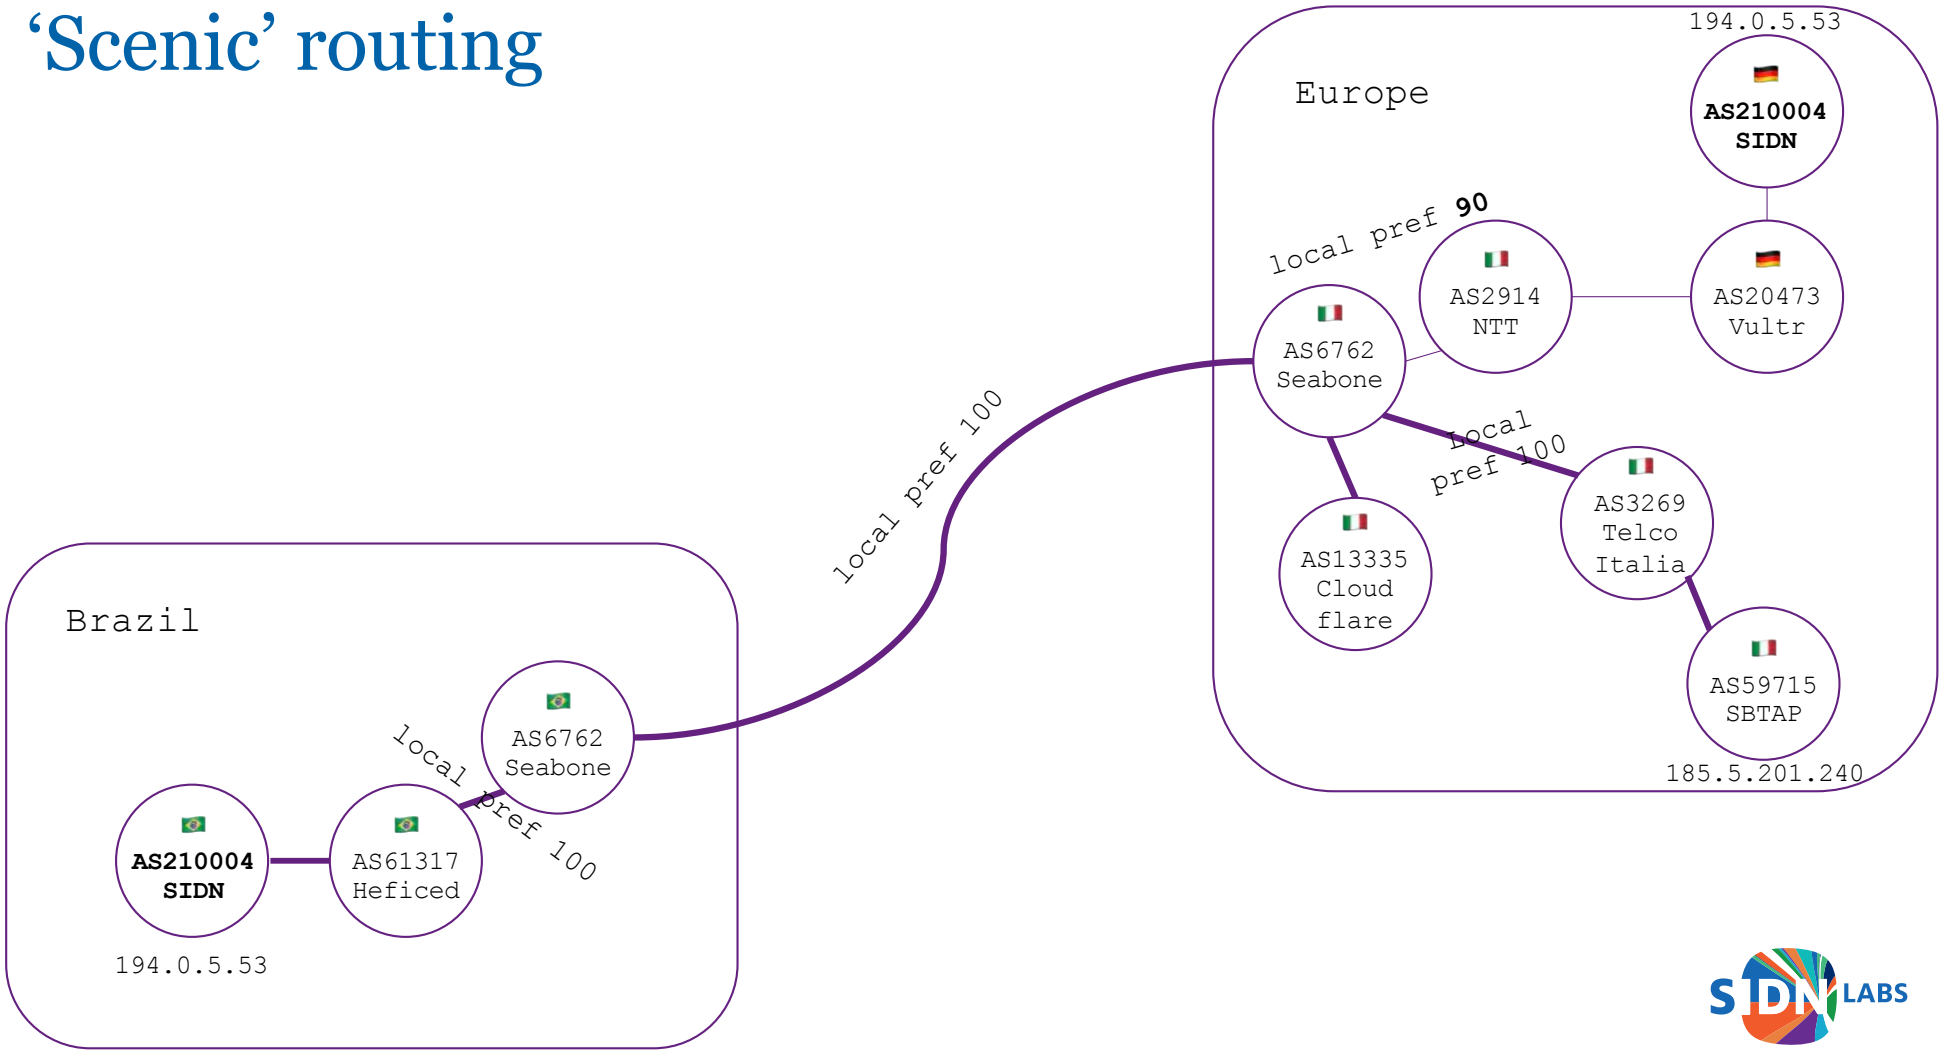

Was Jon Postel indeed ‘the God of the internet’ ?

Postel:

“Of course, there isn’t any 'God of the Internet’ !

The Internet works, because a lot of people cooperate to do things together.”

https://www.linkedin.com/pulse/thanks-jon-john-gods-knights-mario-duhanic/

SIDN LABS

Questions, discussion

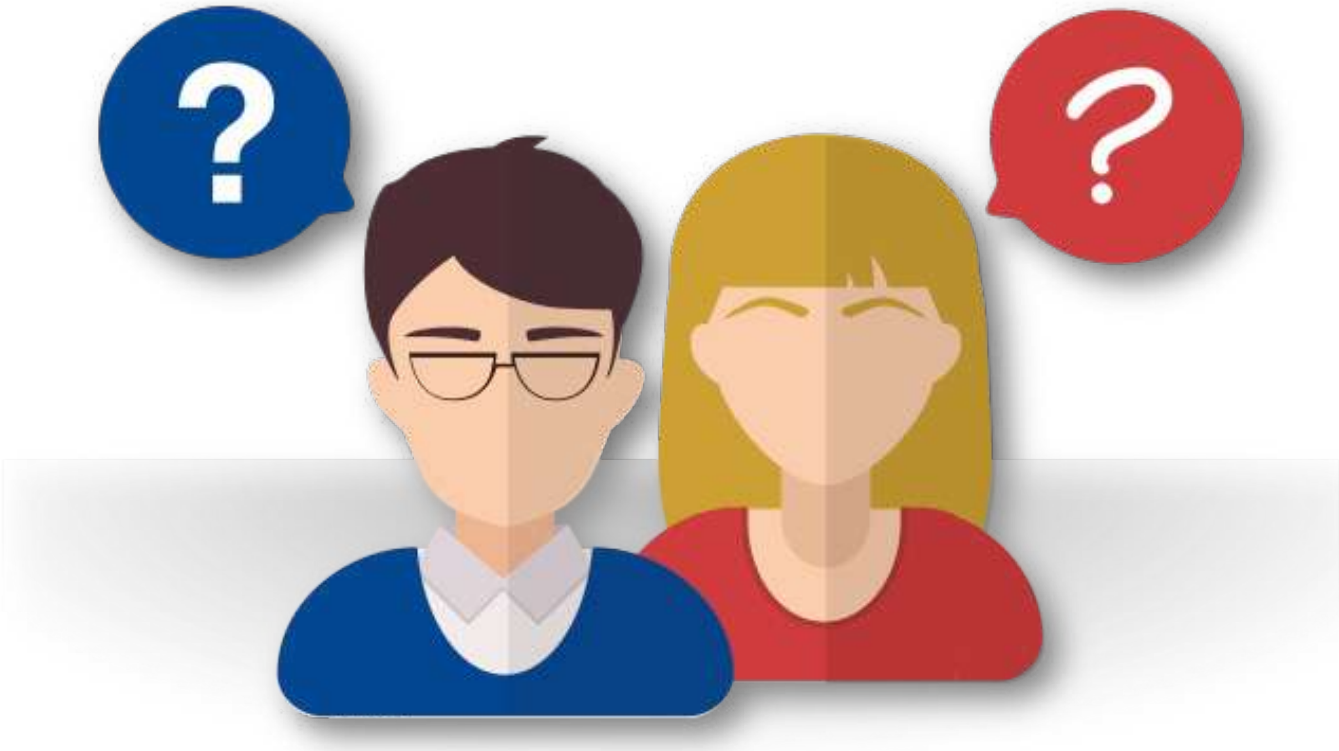

Thank you!