Your world. Our domain.

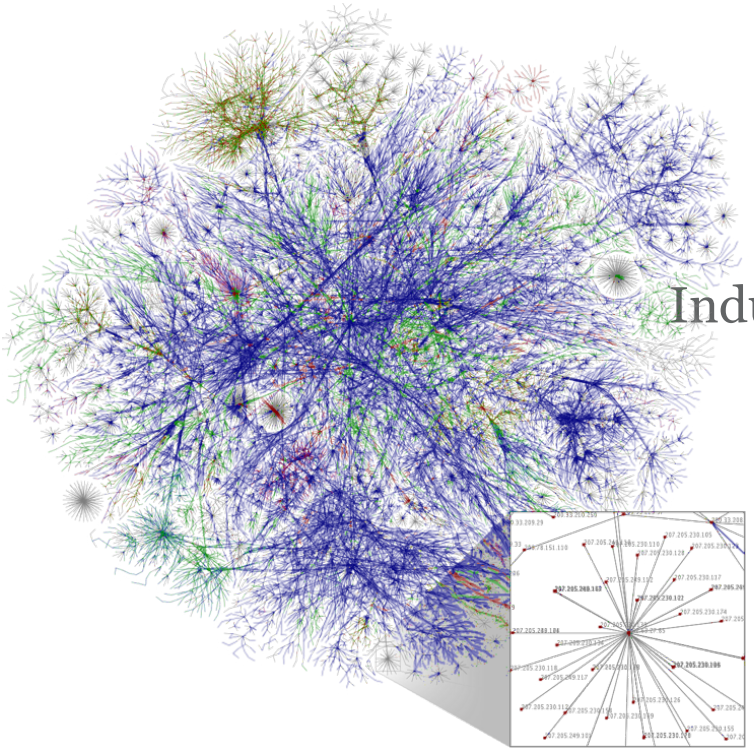

Industry perspective:

Names, numbers, routes

Marco Davids

Extra Lecture for course Security Services for the IoT (SSI)
[virtual session] – April 22th 2020, 11:00 – 12:30

@marcodavids

SIDN Labs

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

SIDN Labs

SIDN Labs contributes to improving the:

- security, privacy,
- stability, resilience

of the .nl ‘ecosystem’ and the broader (and even the future) internet.

By means of applied, measurement-based research and technology-development.

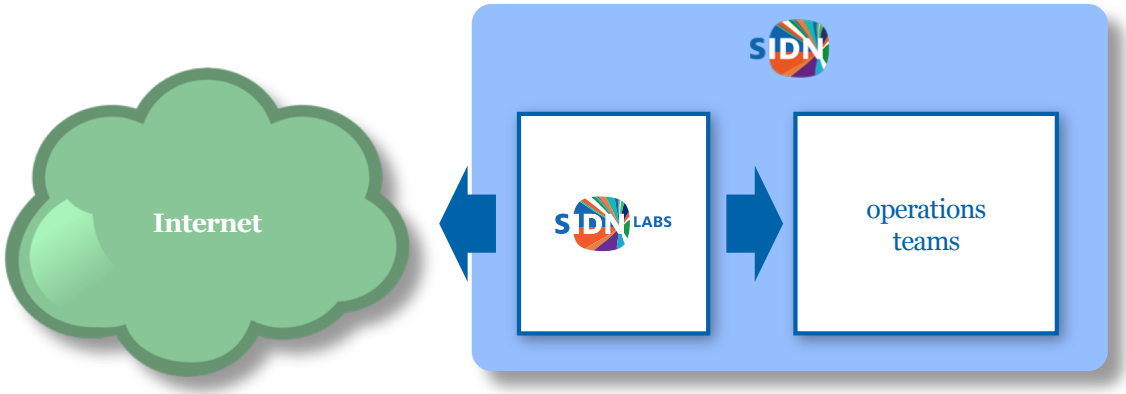

SIDN Labs

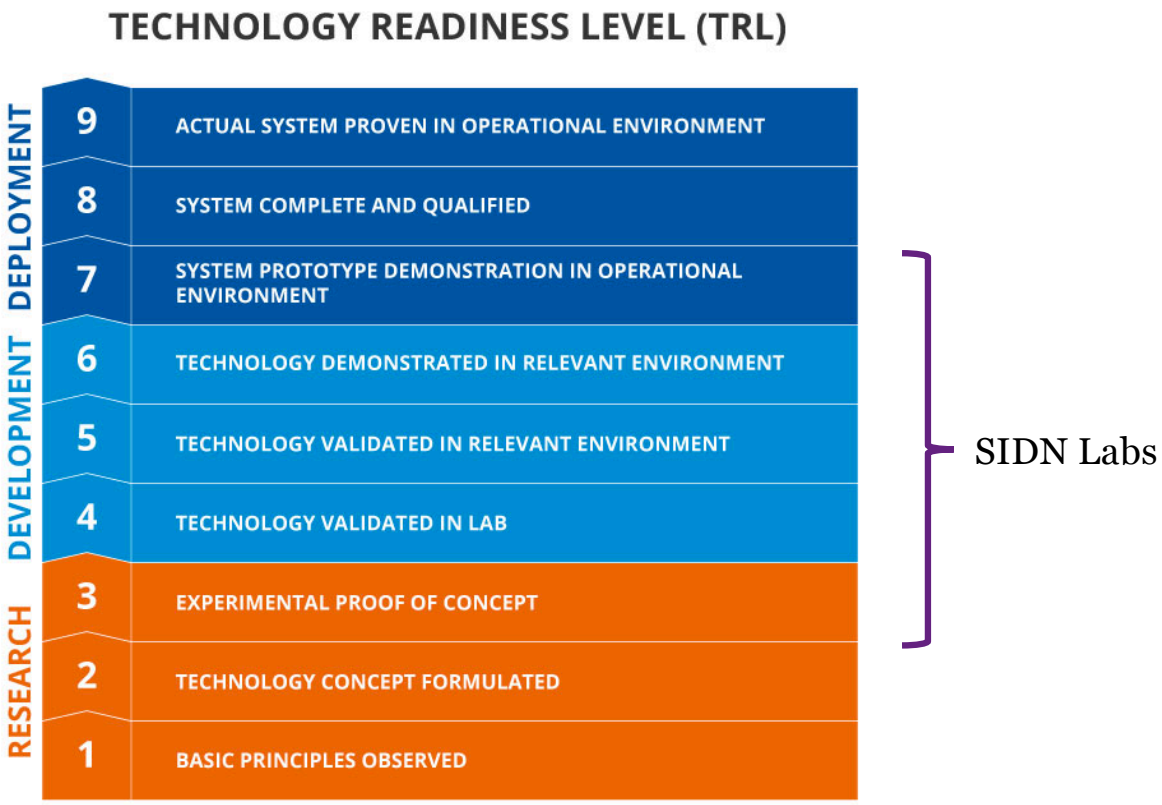

(end of part 0: intro)

How the internet works, in one tweet:

A name (of a resource) indicates what we seek.

An address indicates where it is.

A route indicates how to get there.

– RFC760 and RFC761

John Shoch

Names

Some extensions look quit familiar:

www.sidn.nl

Some maybe not so familiar:

www.sidn.team

Top-level domains

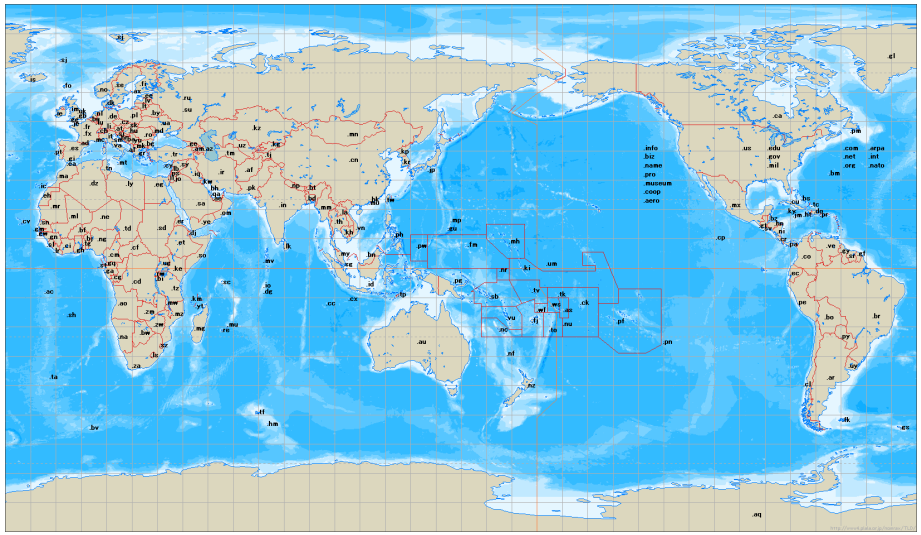

±250

±1300

http://www.marco.panizza.name/dispenseTM/slides/TLD/ccTLD_worldmap.html

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

Domain Name System (DNS)

- Won’t explain it here, you (should) know the drill
- Concept is simple (like chess)
- Reality is not quite that simple (understatement)
- Remember; very crucial component!
- Running a critical DNS infrastructure is a story by itself
- We’ll get to that

Domain Name System (DNS)

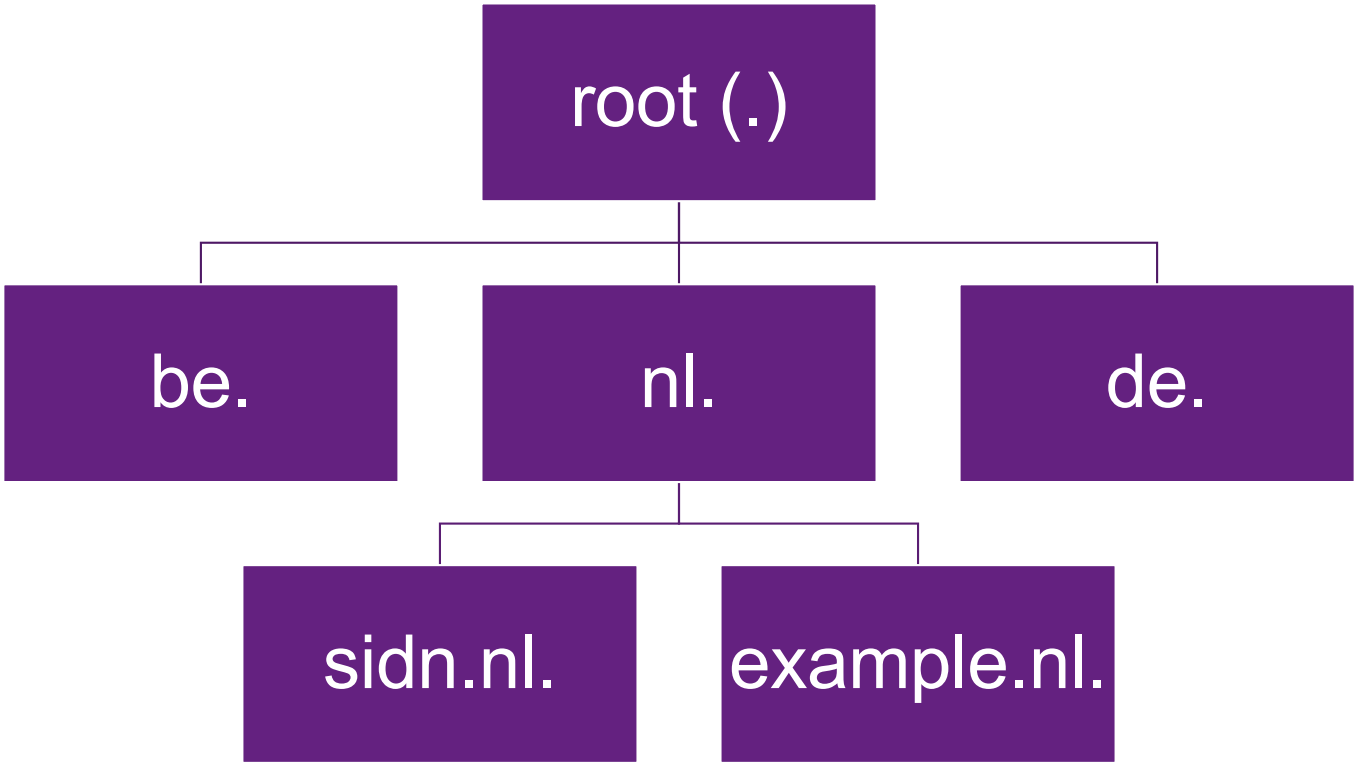

DNS

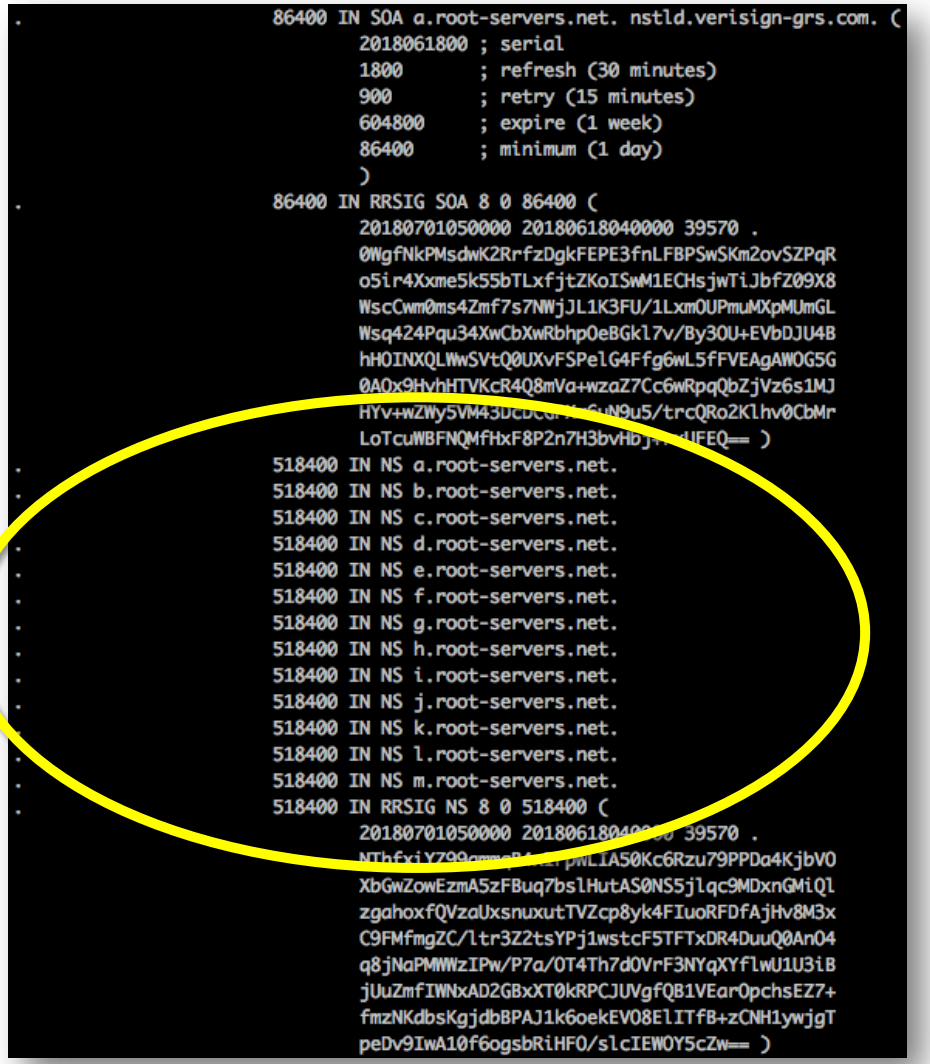

SOA record

NS set

DNS

DNSSEC

DNSSEC

SDN LABS

DNSSEC Key Signing Ceremony

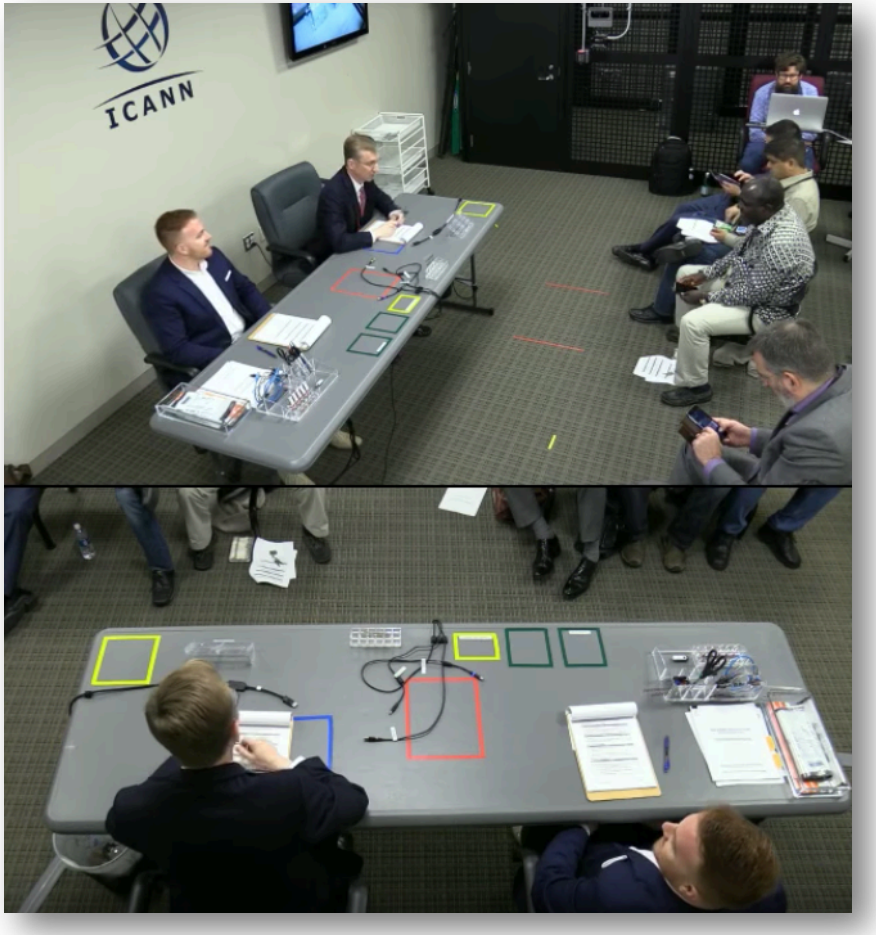

https://www.youtube.com/watch?v=ZTxweLGjZSU

DNSSEC Key Signing Ceremony

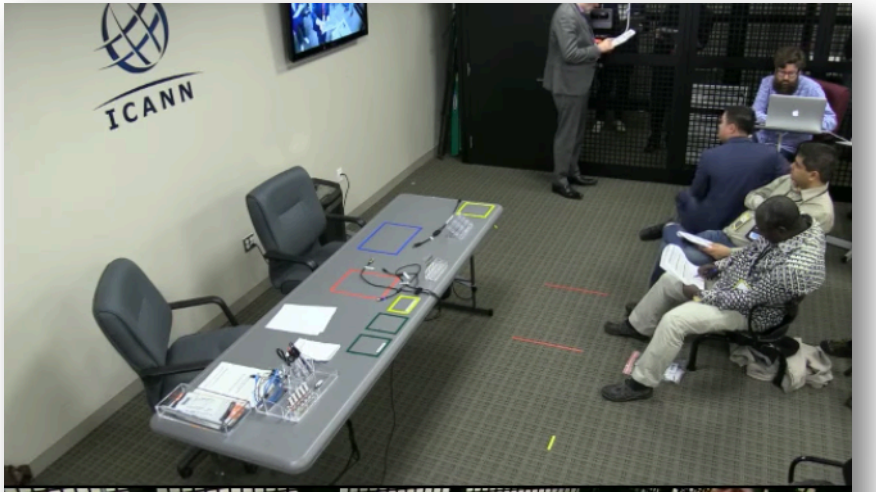

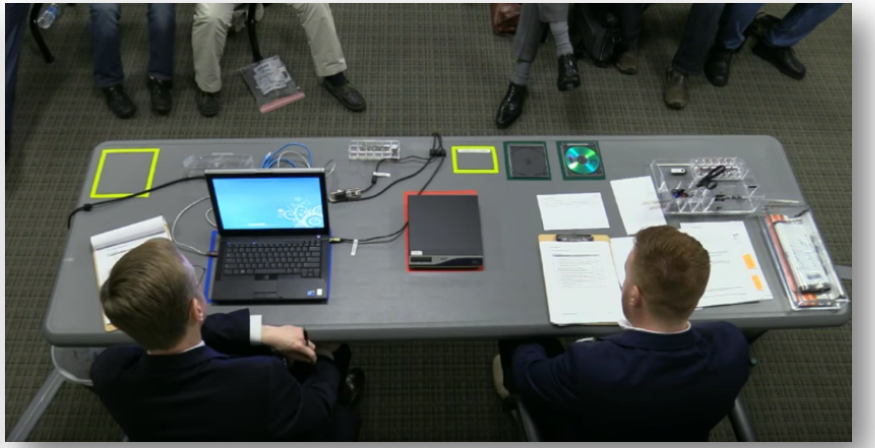

Open the Credential Safe #2

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.		

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

About SIDN

Stichting Internet Domeinregistratie Nederland

- **Registry and designated manager for .nl** ccTLD
  - .nl exists since 1986, SIDN since 1996
  - ~100 FTE (~40% at ICT, 12% at Labs)
- ~ 5.92 million .nl domain names
  - > 55% signed with DNSSEC
- **Registry system + DNS infrastructure**
- RSP for .politie, .amsterdam and .aw
- Located in Arnhem (NL)

SIDNfonds

SIDN LABS

SIDN, the registry for .nl

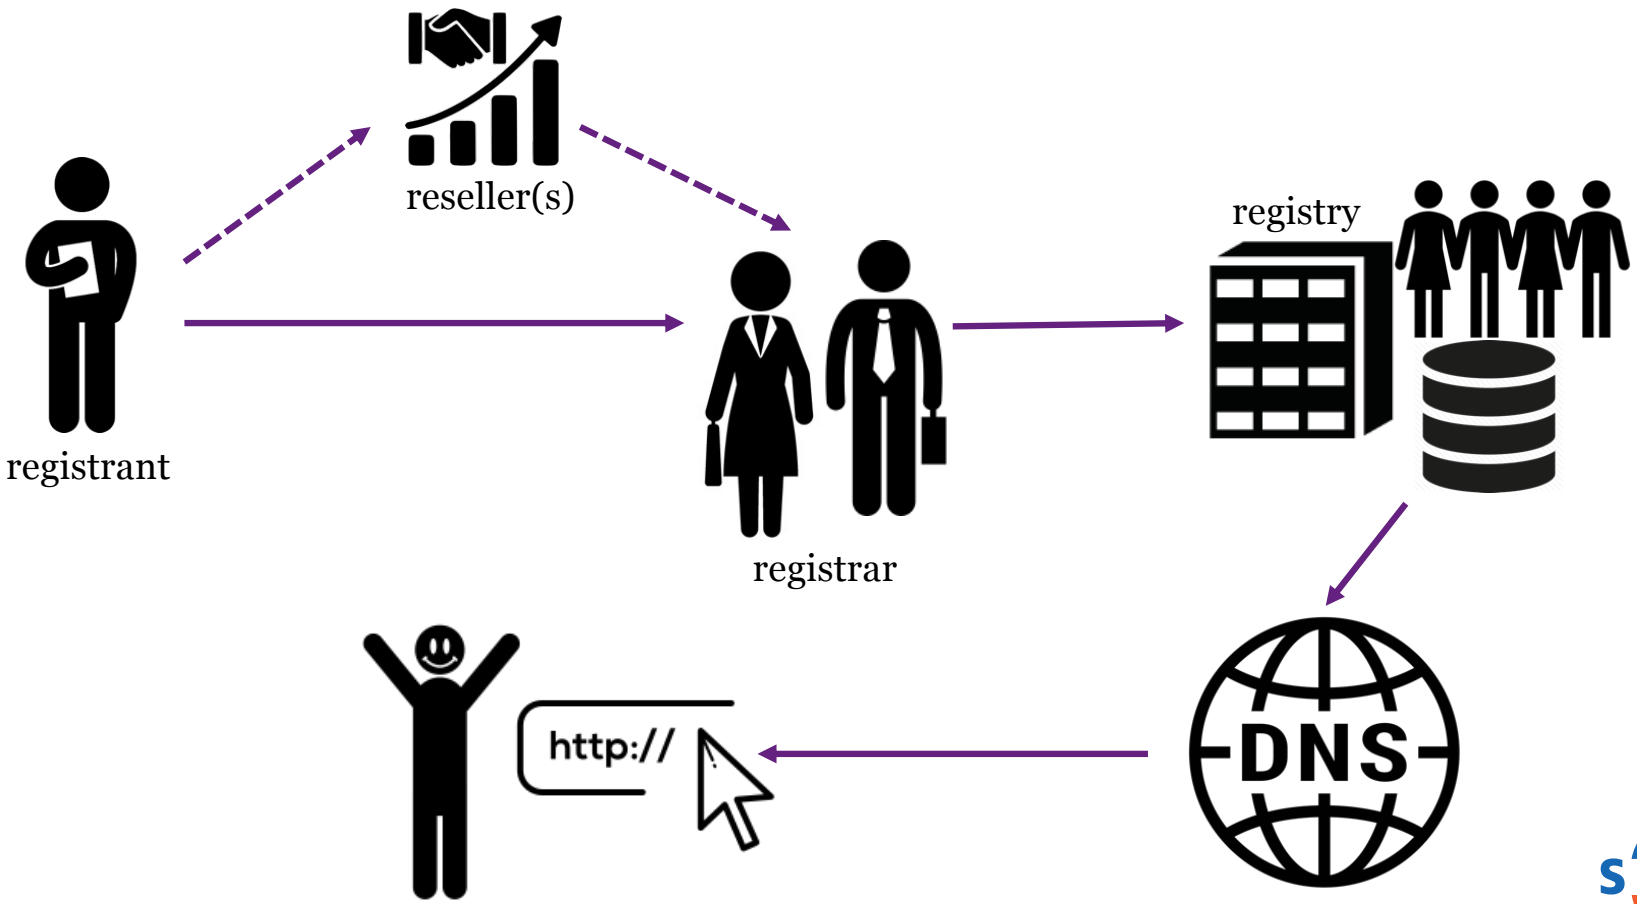

COVID-19 (we noticed something too)

Thousands of related domains registered.

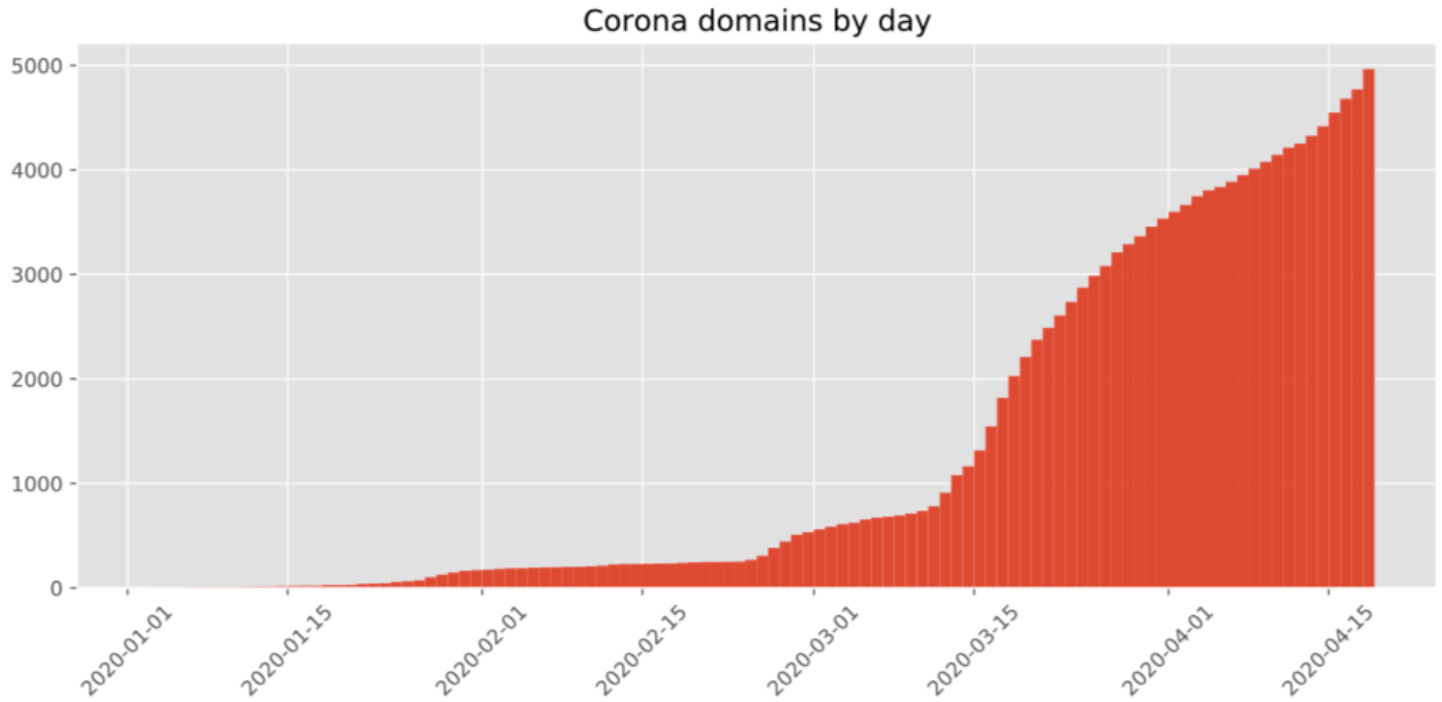

(end of part 1: names)

Numbers (aka addresses)

Legacy scheme: 198.51.100.123

New scheme: 2001:db8::198:51:100:123

Numbers (aka addresses)

(nerdy detail)

IP-address notations are in user friendly format.

This also works:

```
ping 1590075171
```

Or:

http://1590075171

Numbers (aka addresses)

Make no mistake…

192.168.0.1

192.168.1.1

192.168.100.1

192.168.2.1

Numbers (aka addresses)

Anyway… as you know:

Every device directly connected to the internet needs a unique IP address.*

* except for anycast, but more on that later

Managing the IP address space

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/

Managing the IP address space

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/

Managing the IP address space

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

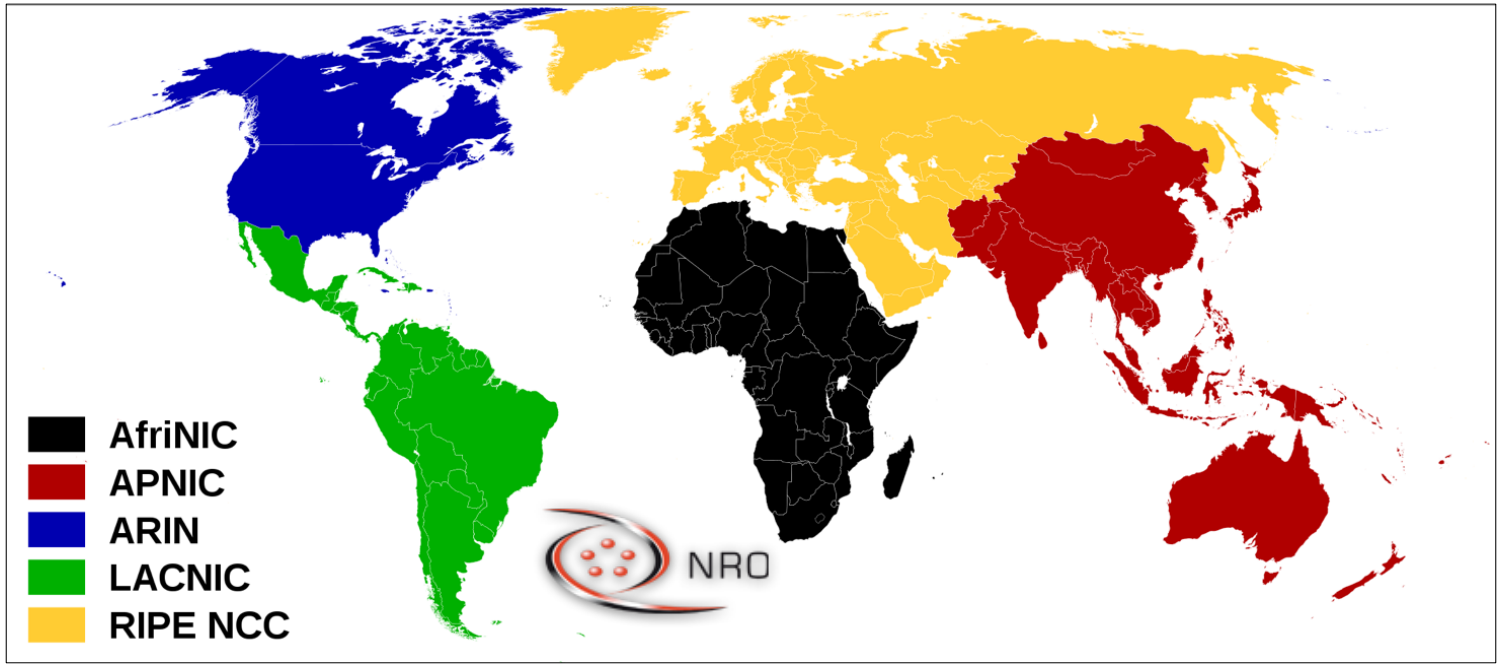

IANA (Internet Assigned Numbers Authority) → RIRs → LIRs

https://www.nro.net/

Managing a whole lot more! (protocol assignments)

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

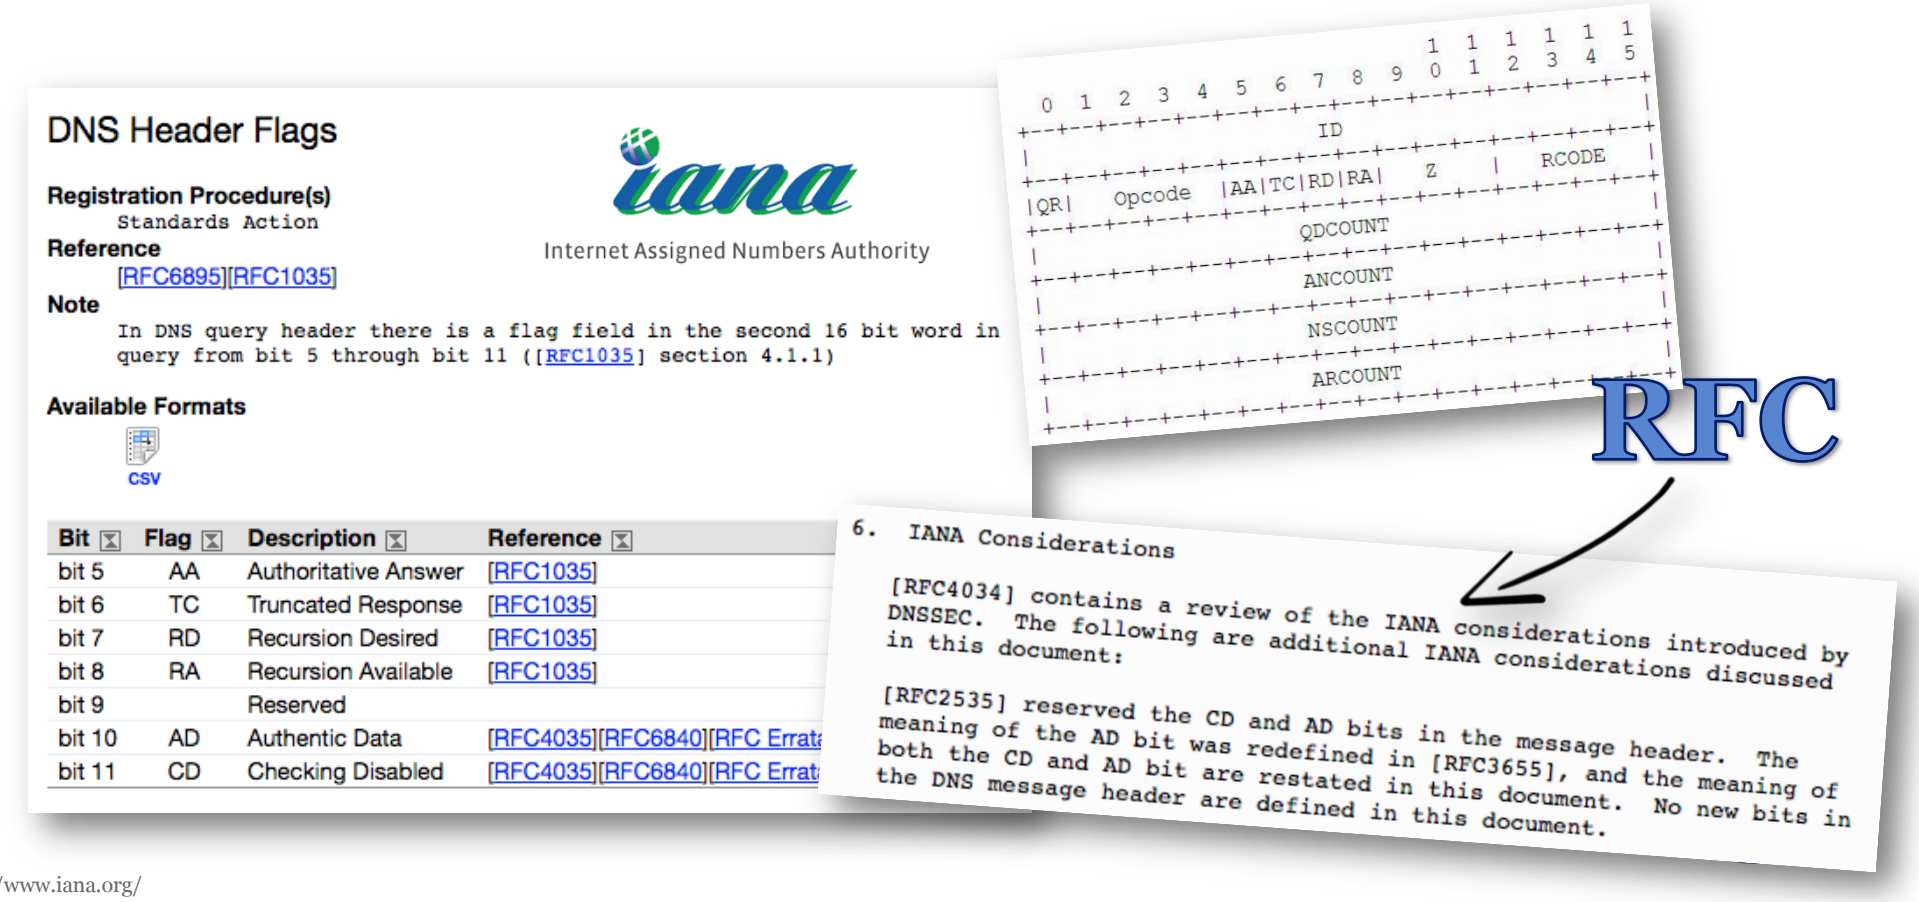

DNS Header Flags

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

```
                                1  1  1  1  1  1
  0  1  2  3  4  5  6  7  8  9  0  1  2  3  4  5
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                      ID                       |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|QR|   Opcode  |AA|TC|RD|RA|   Z    |   RCODE   |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                    QDCOUNT                    |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                    ANCOUNT                    |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                    NSCOUNT                    |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                    ARCOUNT                    |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
```

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/

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, with no formal membership
- Everyone can join in (in person or via mailing lists)
- Under the auspices of the Internet Society (ISOC)
- Large number of working groups and informal discussion groups
- Rough consensus* is the primary basis for decision making.
- Often slow processes!
- But lots of RFC's ! Over 8778 and many more drafts.

* https://tools.ietf.org/html/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.

https://tools.ietf.org/html/rfc2026

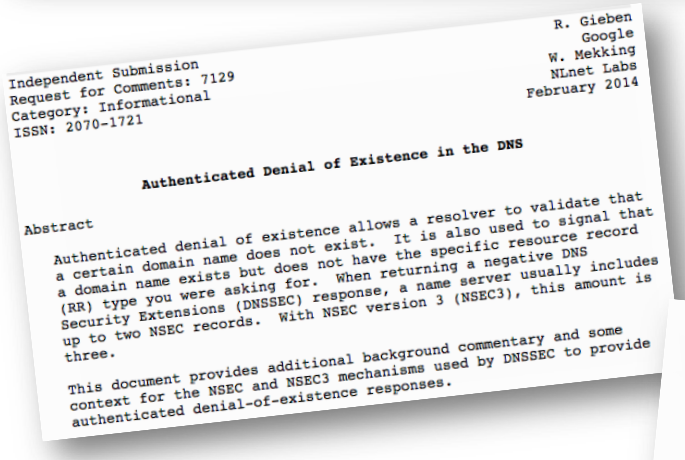

Independent Submission
Request for Comments: 7129
Category: Informational
ISSN: 2070-1721

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 that a certain domain name does not exist. It is also used to signal that 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 is 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.

- Informational
- Experimental
- BCP
- Standards track
- Historic
- Unknown

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 poll queue defined in RFC 5730.

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

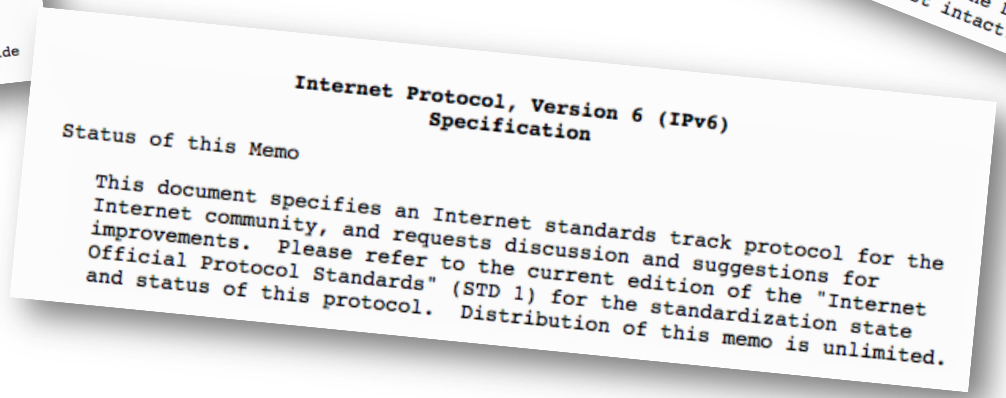

Internet Protocol, Version 6 (IPv6) Specification

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

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

SIDN LABS

A personal favorite: RFC1925 😄

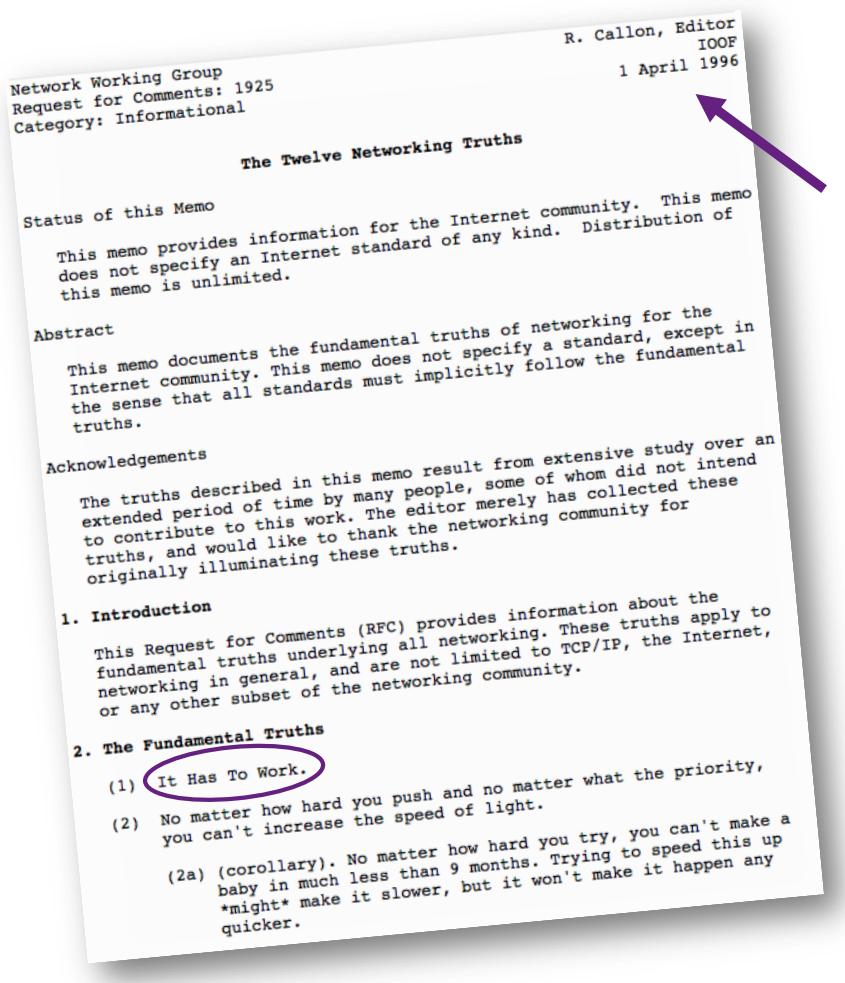

Network Working Group
Request for Comments: 1925
Category: Informational

R. Callon, Editor
IOOF
1 April 1996

The Twelve Networking Truths

Status of this Memo

This memo provides information for the Internet community. This memo does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Abstract

This memo documents the fundamental truths of networking for the Internet community. This memo does not specify a standard, except in the sense that all standards must implicitly follow the fundamental truths.

Acknowledgements

The truths described in this memo result from extensive study over an extended period of time by many people, some of whom did not intend to contribute to this work. The editor merely has collected these truths, and would like to thank the networking community for originally illuminating these truths.

1. Introduction

This Request for Comments (RFC) provides information about the fundamental truths underlying all networking. These truths apply to networking in general, and are not limited to TCP/IP, the Internet, or any other subset of the networking community.

2. The Fundamental Truths

(1) It Has To Work.

(2) No matter how hard you push and no matter what the priority, you can't increase the speed of light.

(2a) (corollary). No matter how hard you try, you can't make a baby in much less than 9 months. Trying to speed this up *might* make it slower, but it won't make it happen any quicker.

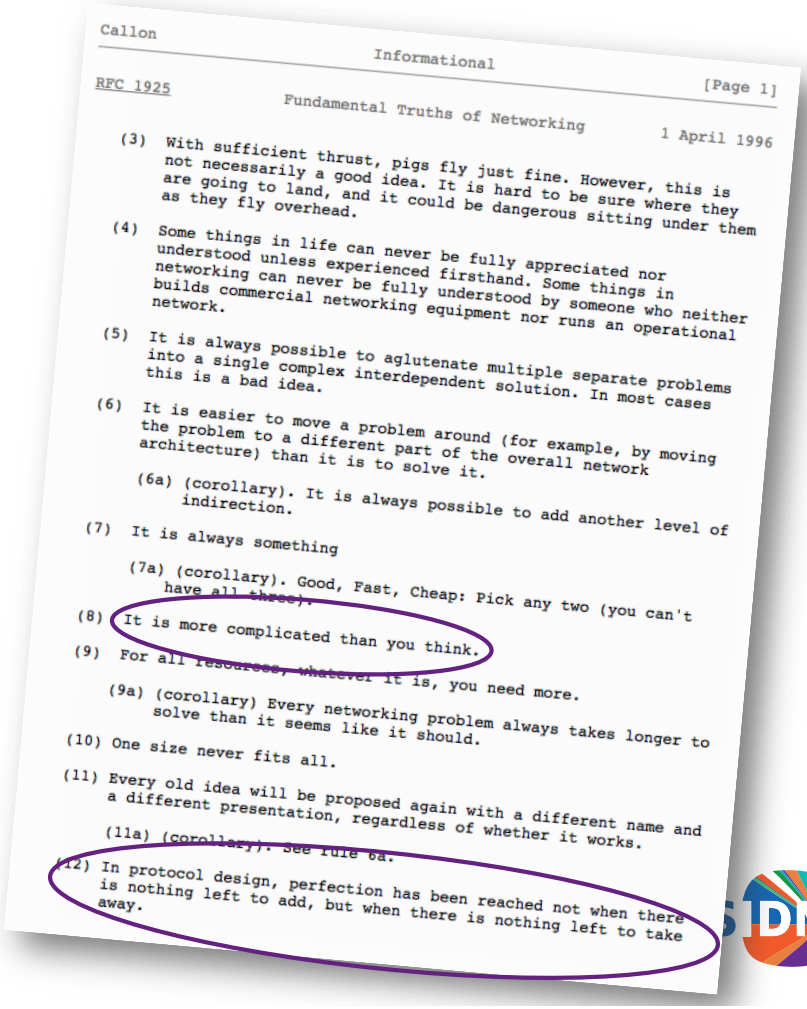

Callon Informational [Page 1]

RFC 1925 Fundamental Truths of Networking 1 April 1996

(3) With sufficient thrust, pigs fly just fine. However, this is not necessarily a good idea. It is hard to be sure where they are going to land, and it could be dangerous sitting under them as they fly overhead.

(4) Some things in life can never be fully appreciated nor understood unless experienced firsthand. Some things in networking can never be fully understood by someone who neither builds commercial networking equipment nor runs an operational network.

(5) It is always possible to aglutenate multiple separate problems into a single complex interdependent solution. In most cases this is a bad idea.

(6) It is easier to move a problem around (for example, by moving the problem to a different part of the overall network architecture) than it is to solve it.

(6a) (corollary). It is always possible to add another level of indirection.

(7) It is always something

(7a) (corollary). Good, Fast, Cheap: Pick any two (you can't have all three).

(8) It is more complicated than you think.

(9) For all resources, whatever it is, you need more.

(9a) (corollary) Every networking problem always takes longer to solve than it seems like it should.

(10) One size never fits all.

(11) Every old idea will be proposed again with a different name and a different presentation, regardless of whether it works.

(11a) (corollary). See rule 6a.

(12) In protocol design, perfection has been reached not when there is nothing left to add, but when there is nothing left to take away.

Playing field of IETF?

The Internet Hourglass

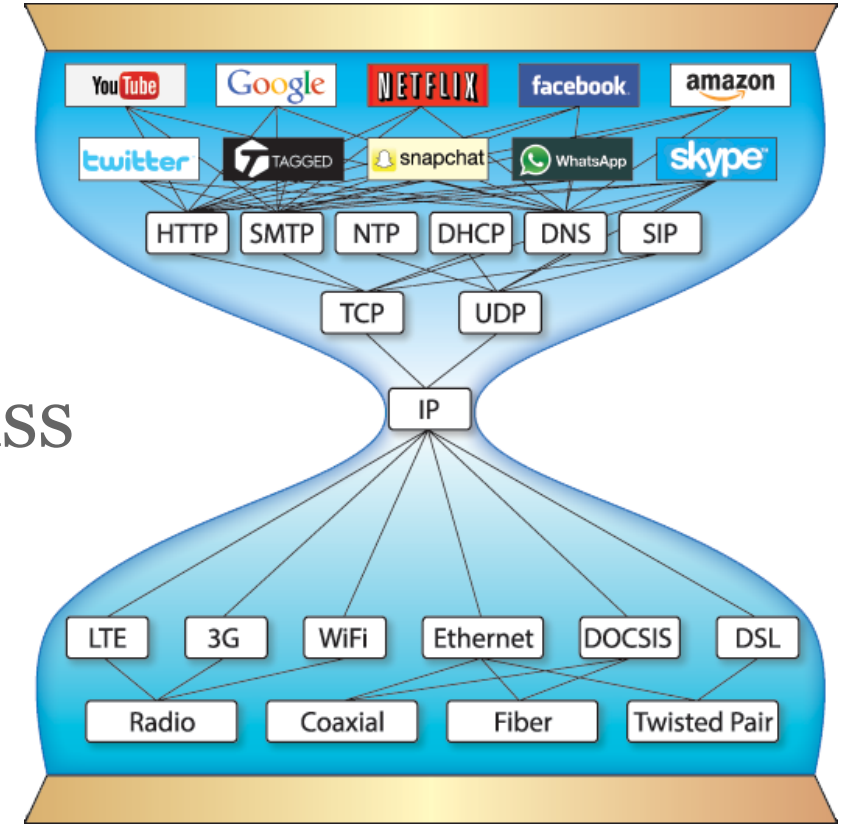

SDN LABS

Abstraction layers always +1

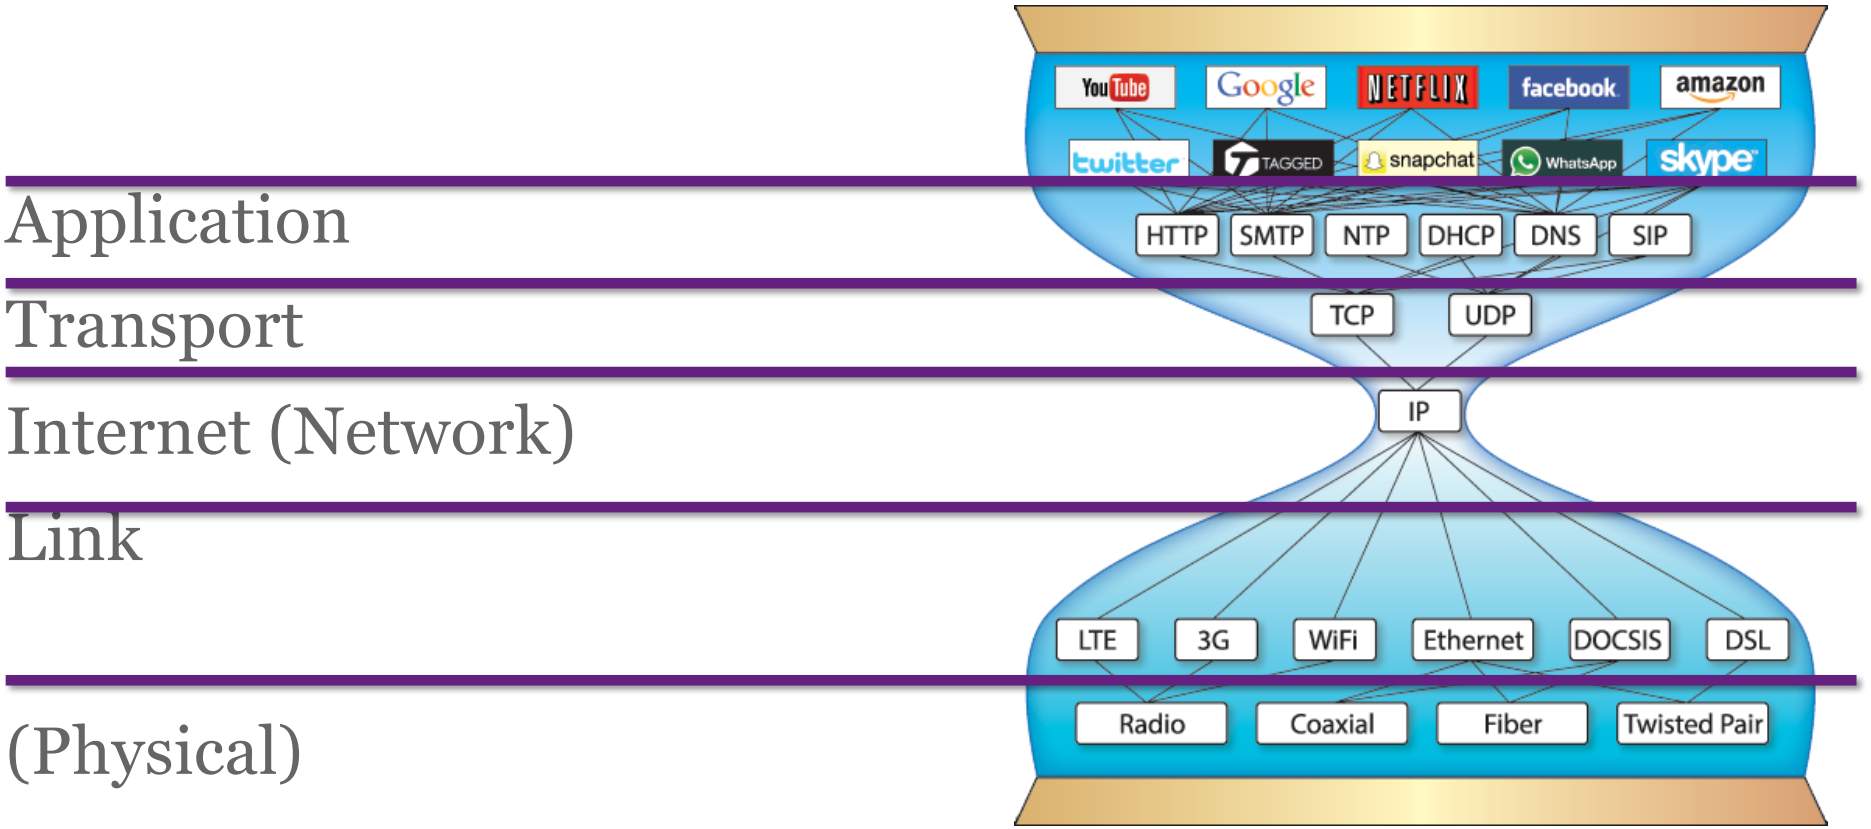

SDN LABS

Also…

Fast

Slow!

Fast

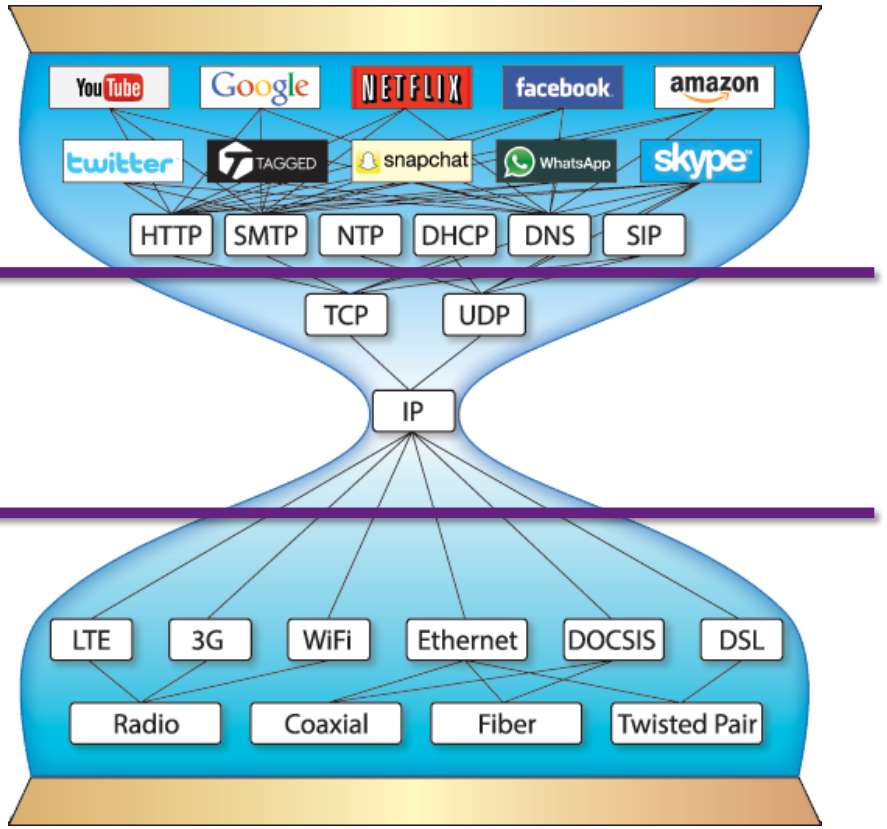

SDN LABS

Sounds familiar?

Most people

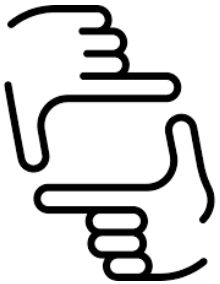

Me (you?)

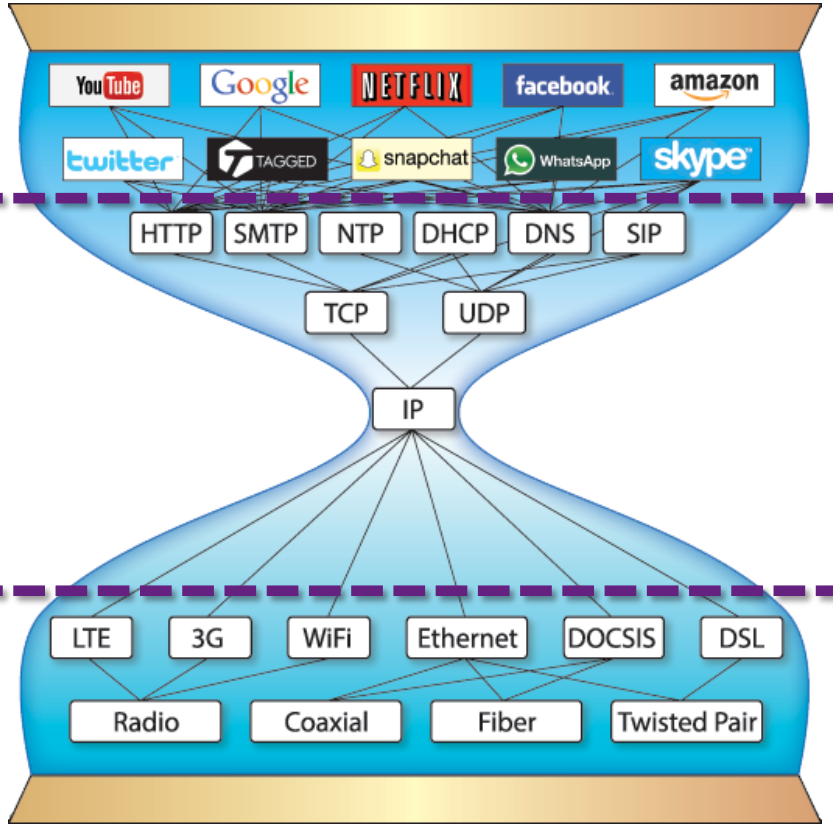

SDN LABS

Top is mostly what the news is about

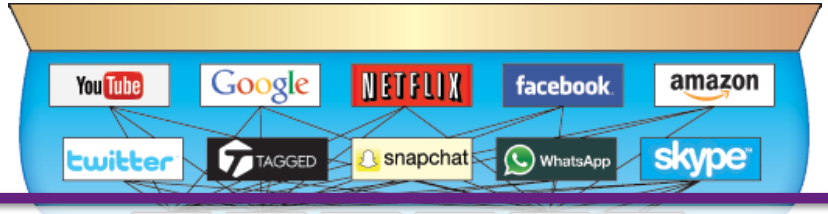

Piracy / DRM
Innovation
Cloud
Darknet
Anti trust
Spam
AI
Ad blocking
Phishing
Commerce
Geo blocking
Making money
Hype
SEO
Net neutrality
Cat pictures
Exploits
Domotica
Shadow ban
Privacy
Bitcoin
ID theft
Smart society
Hacktivism
NSA/Snowden
Echo chamber
Covid-app
Social impact
Digital rights
Addiction
Censorship

SIDN LABS

Bottom is also very interesting, but

http://www.eurofiber.nl

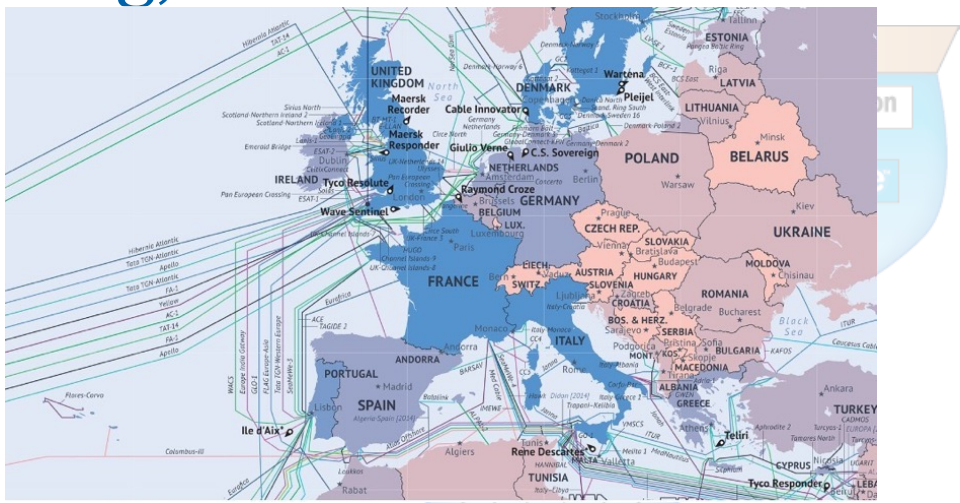

https://www.fiberoptictel.com/submarine-fiber-optic-cables-international-communications/

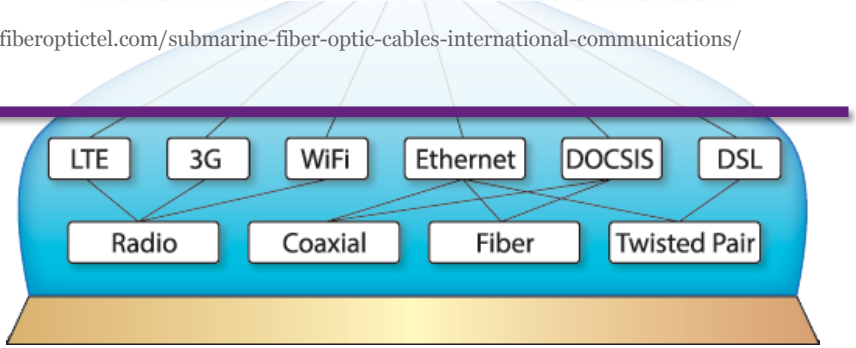

IEEE

(and others, like Bluetooth®)

SDN LABS

Playing field of IETF:

IETF®

(and others, like W3C®)

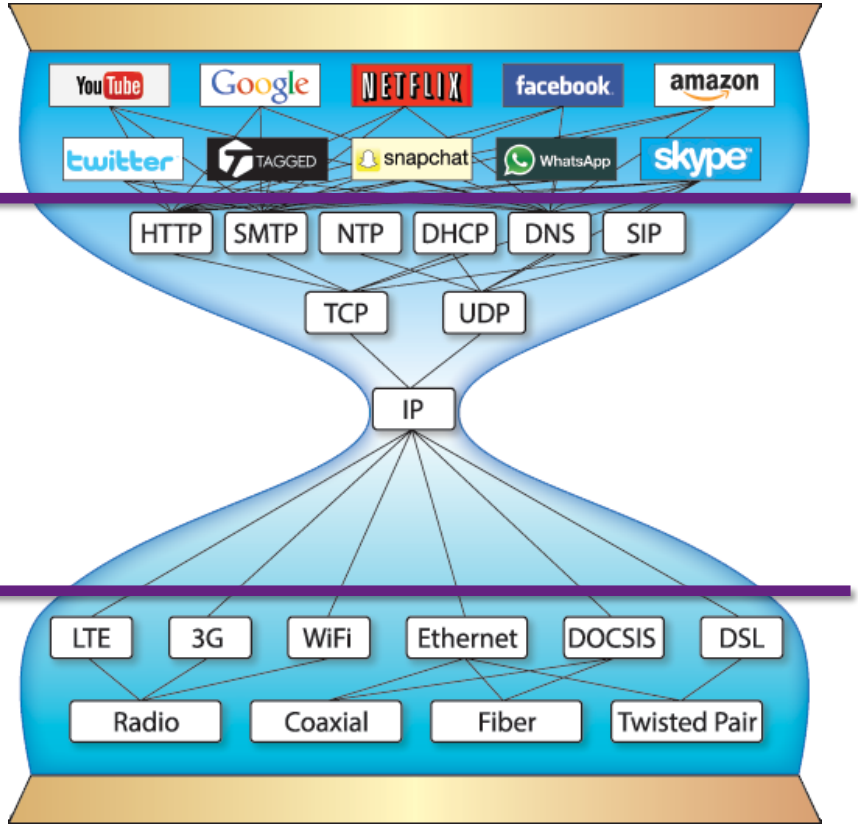

SDN LABS

(end of part 2: numbers, etc.)

Routing

Here's an average network

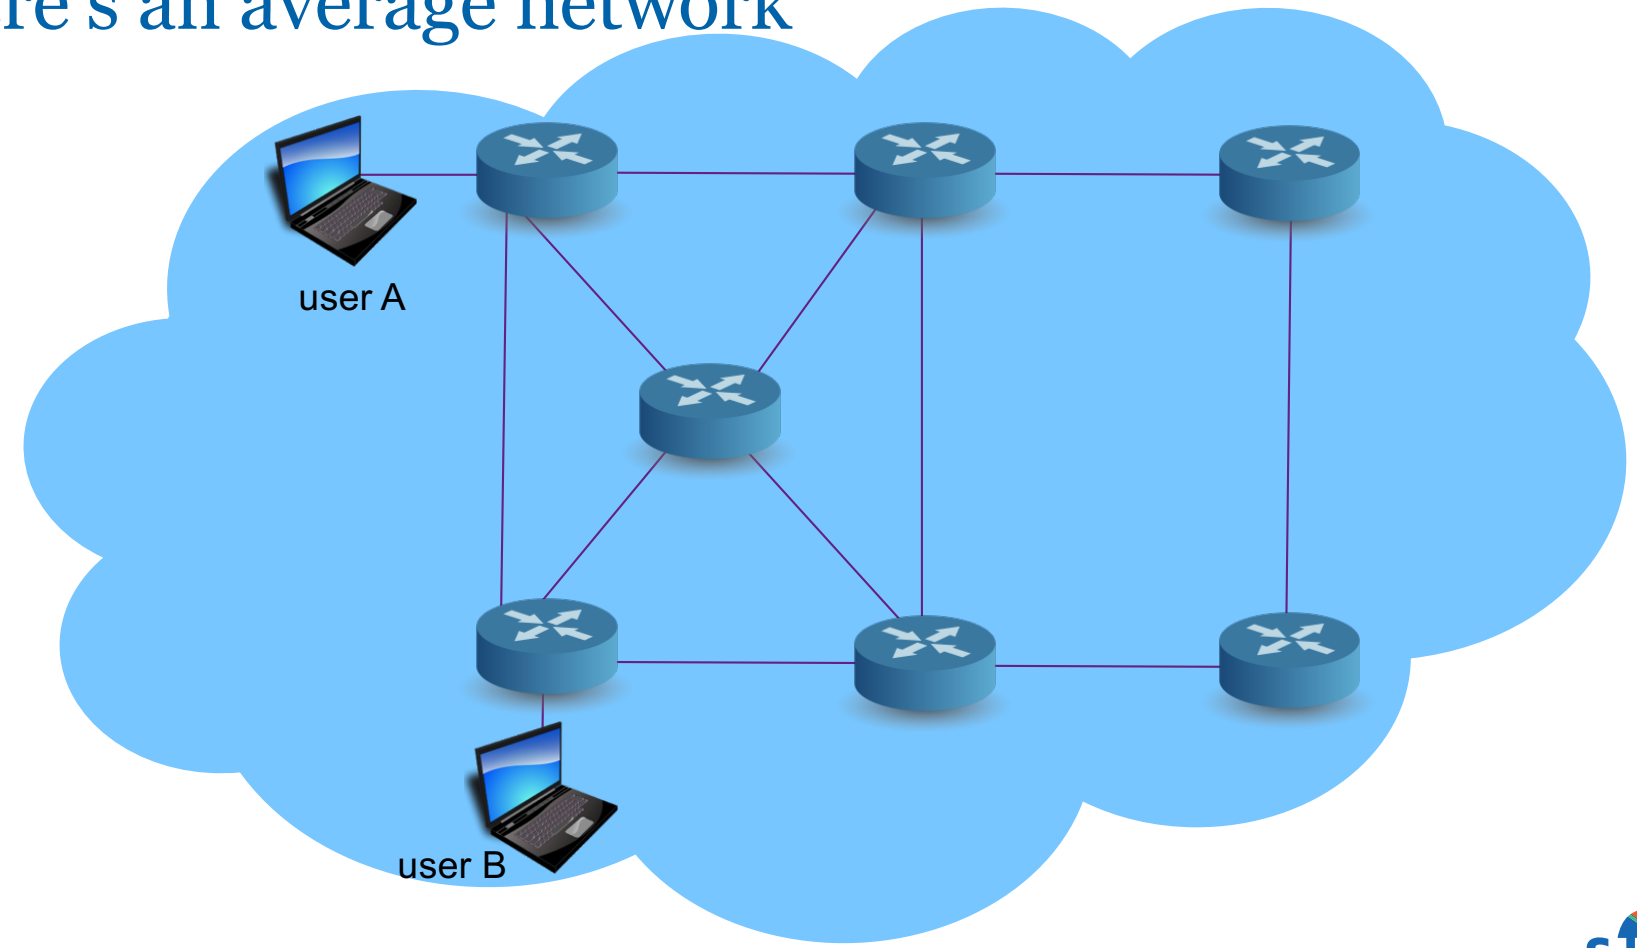

Here's an average network

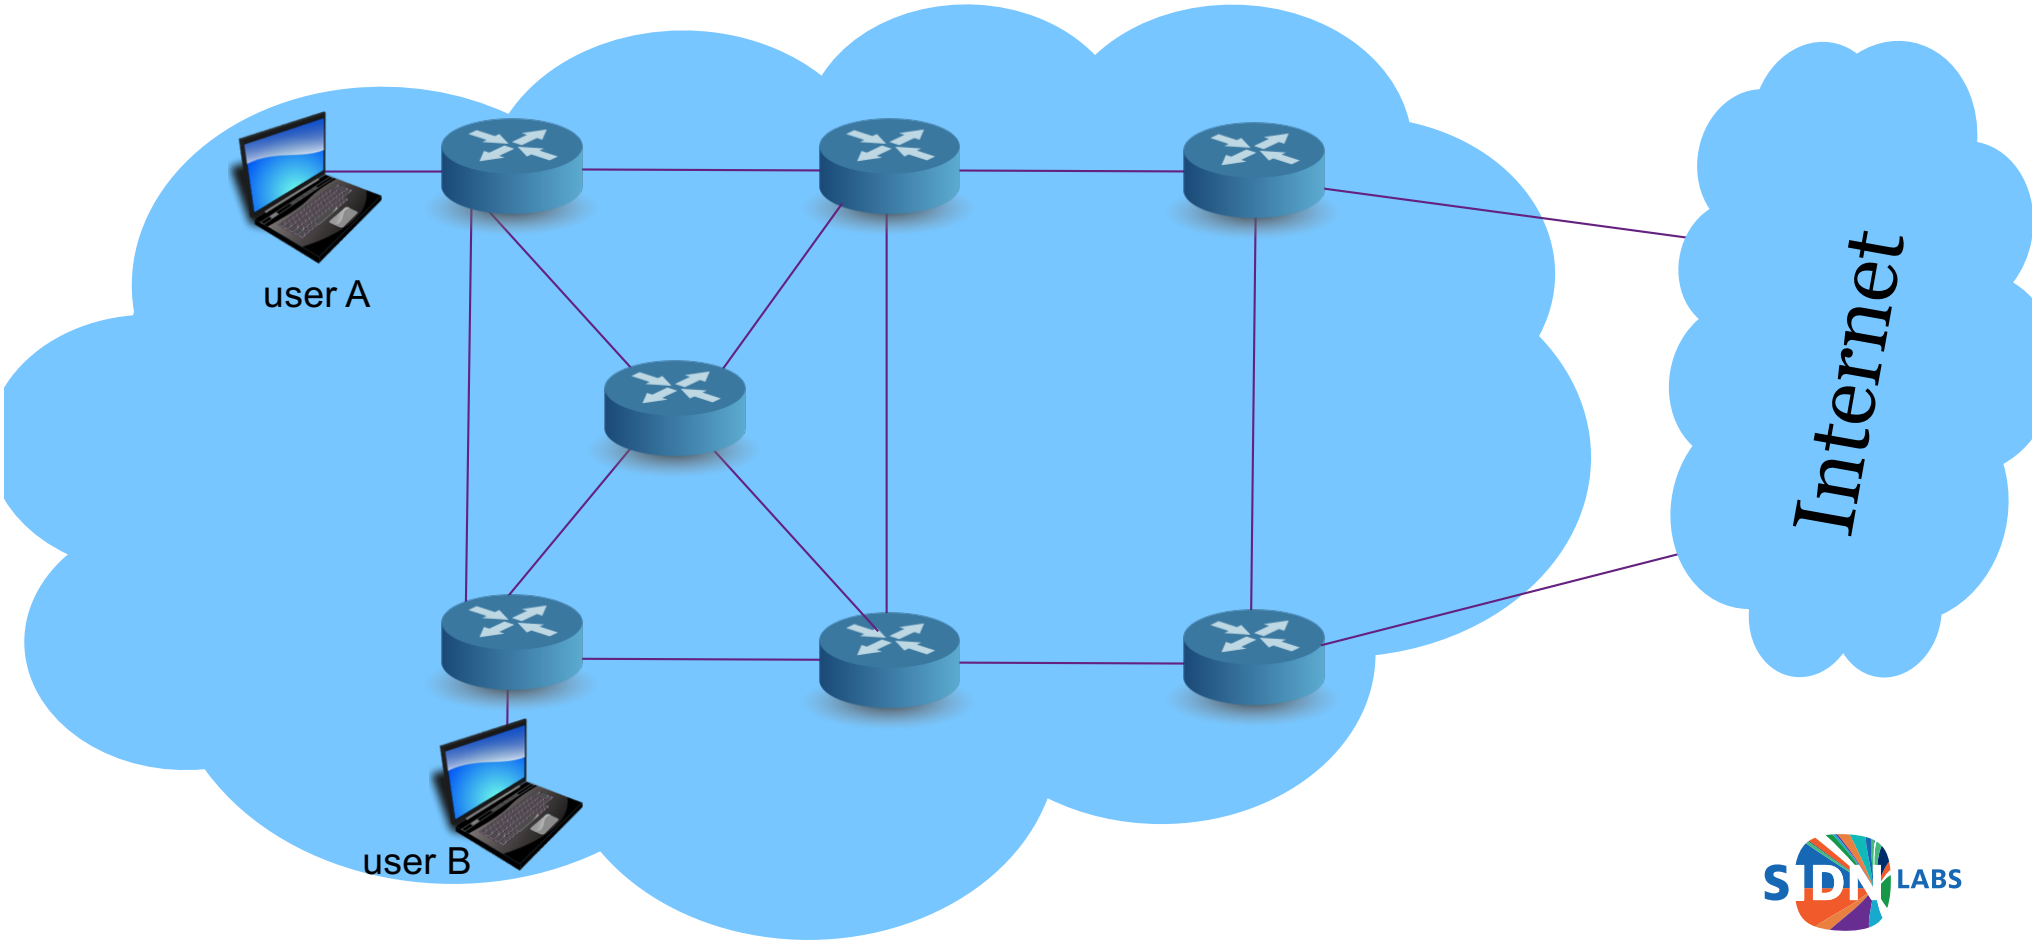

Here's an average network

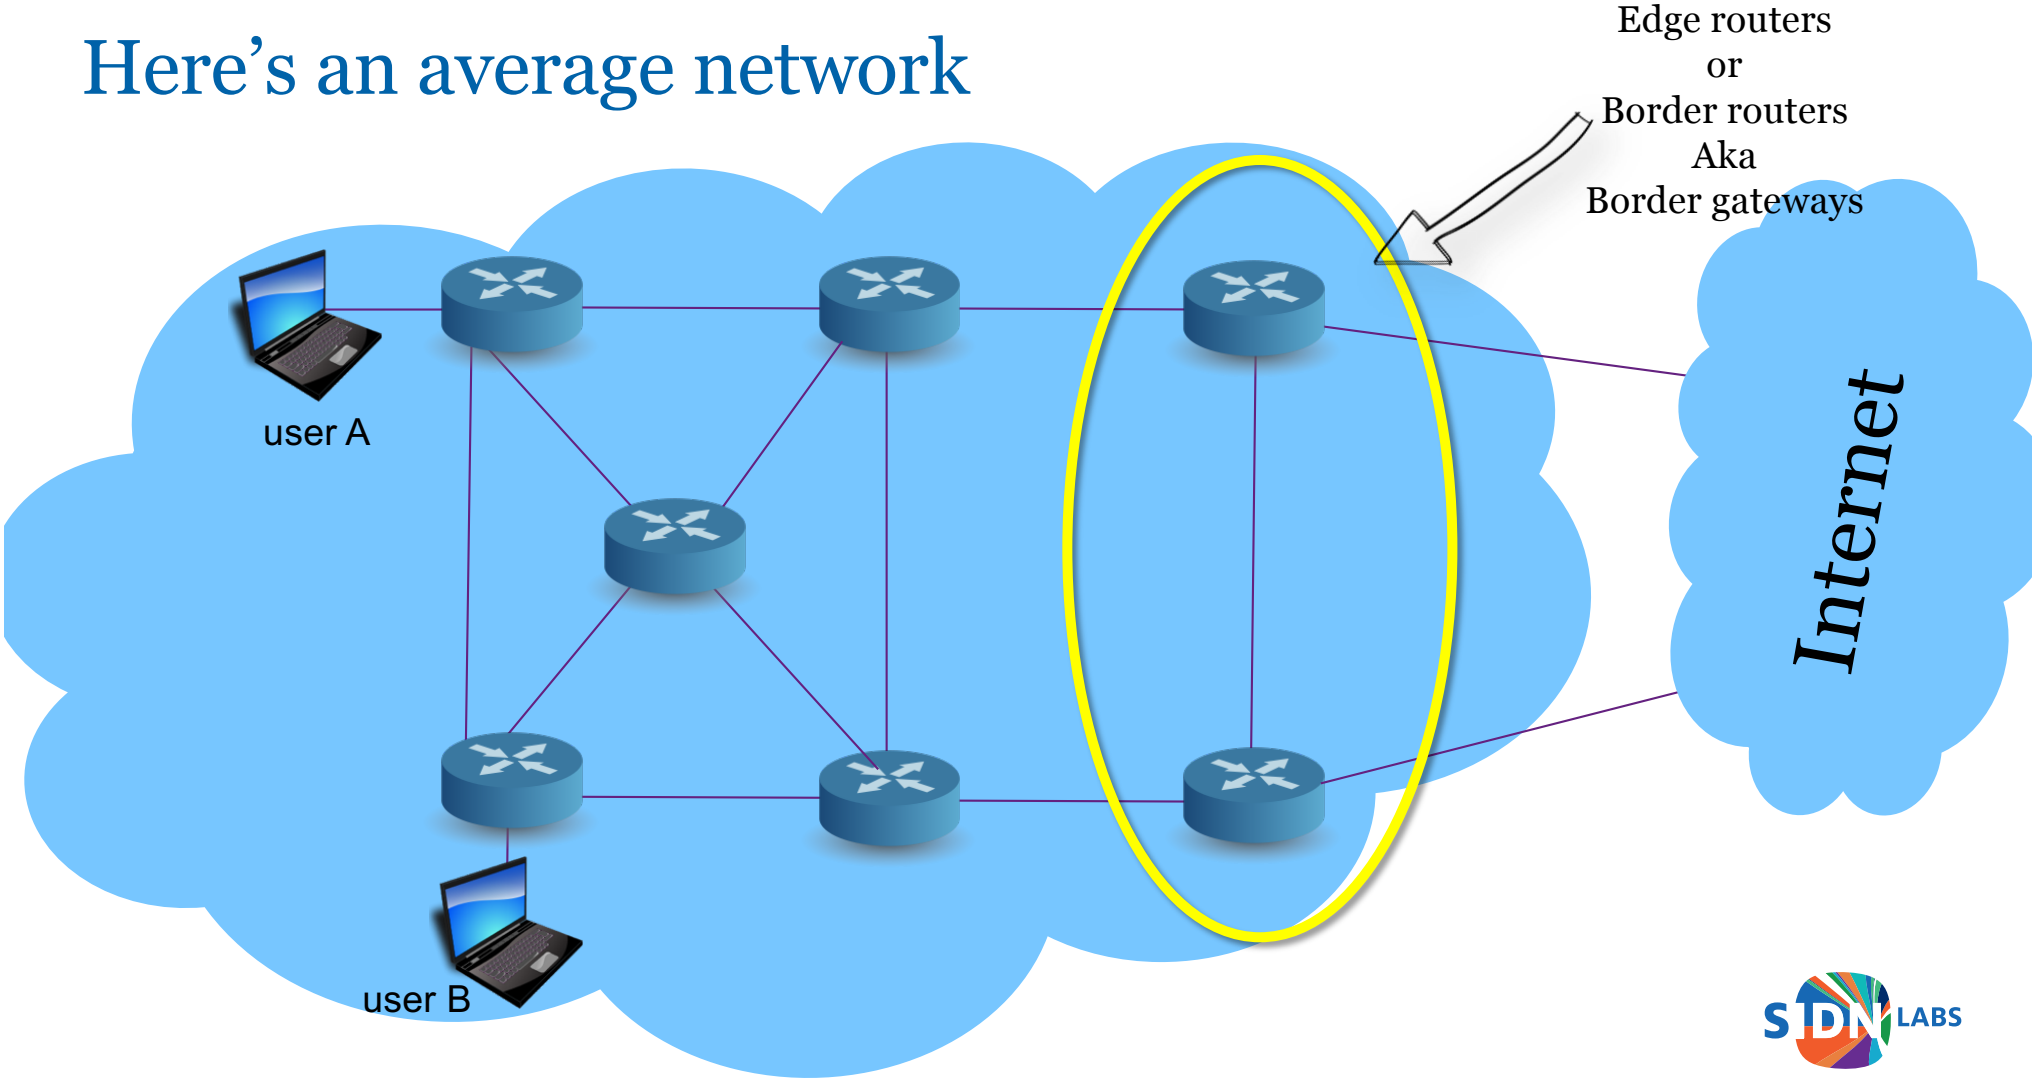

BGP (in a nutshell)

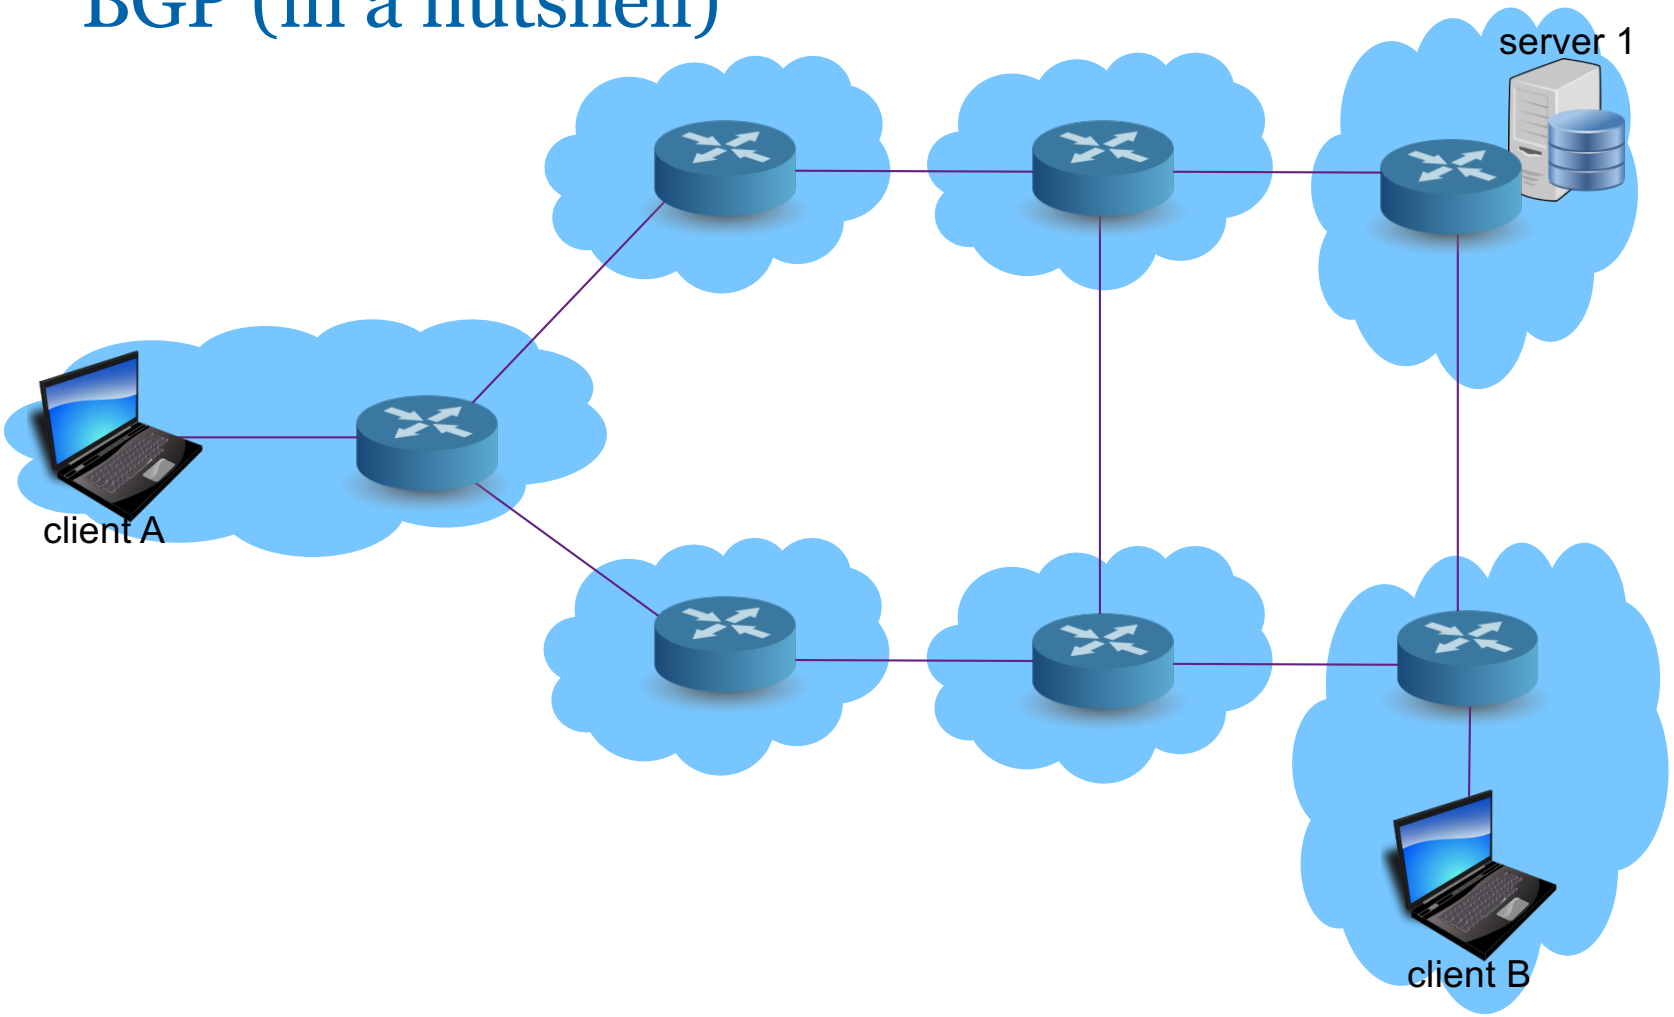

BGP (in a nutshell)

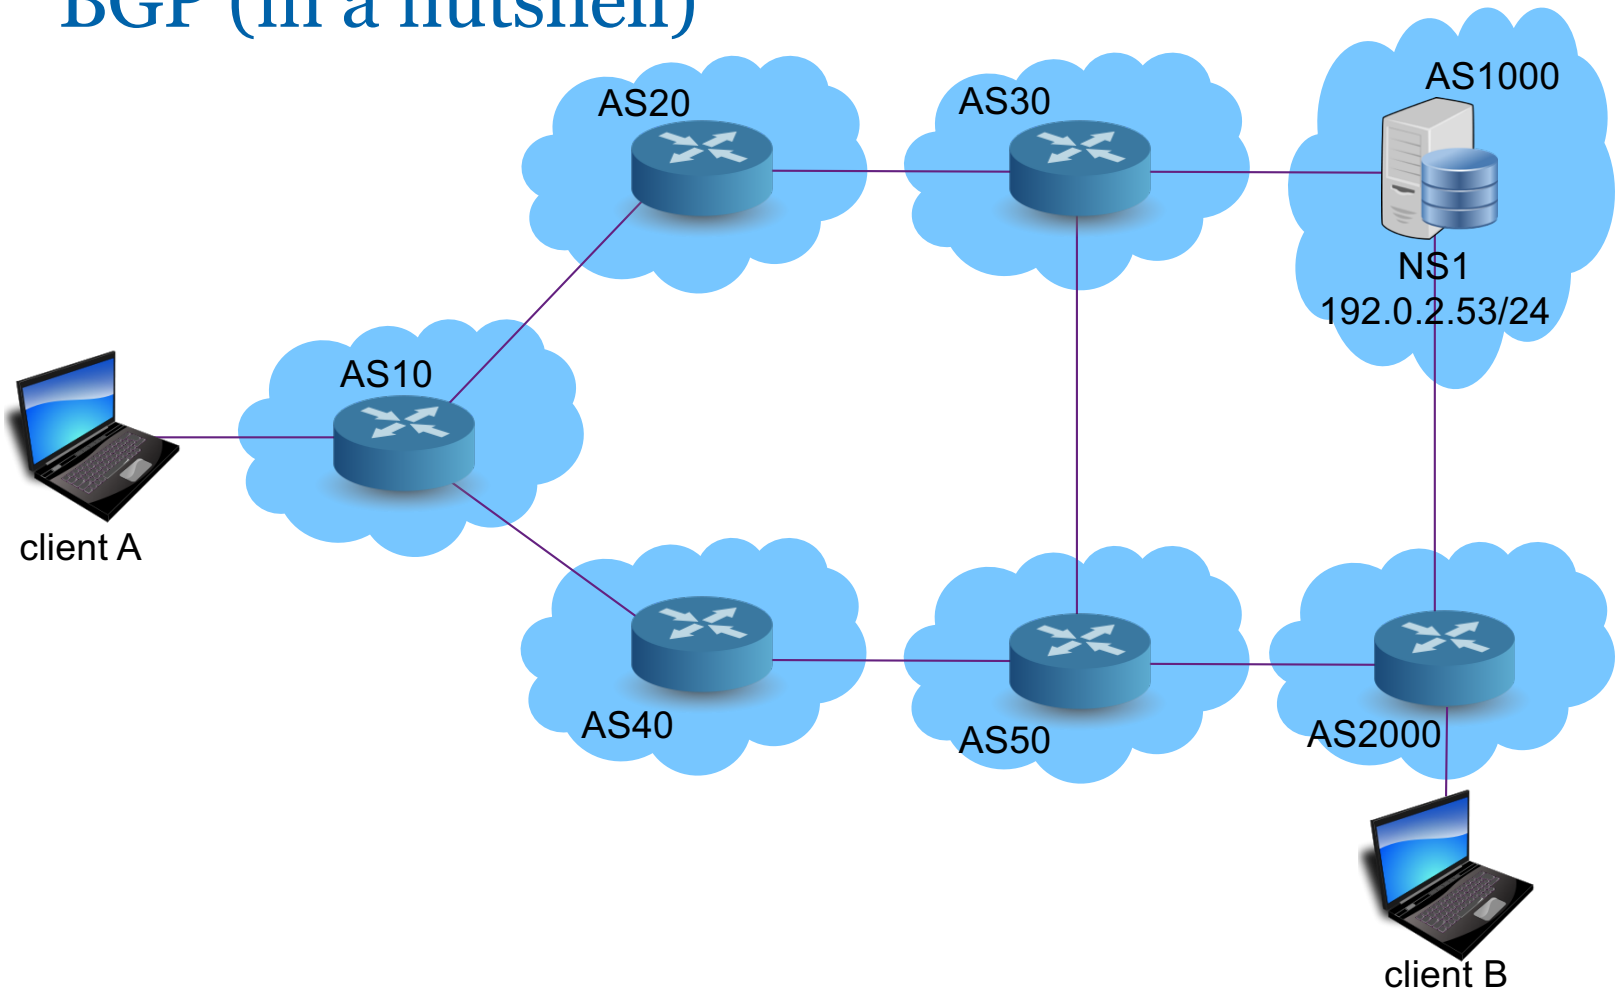

BGP (in a nutshell)

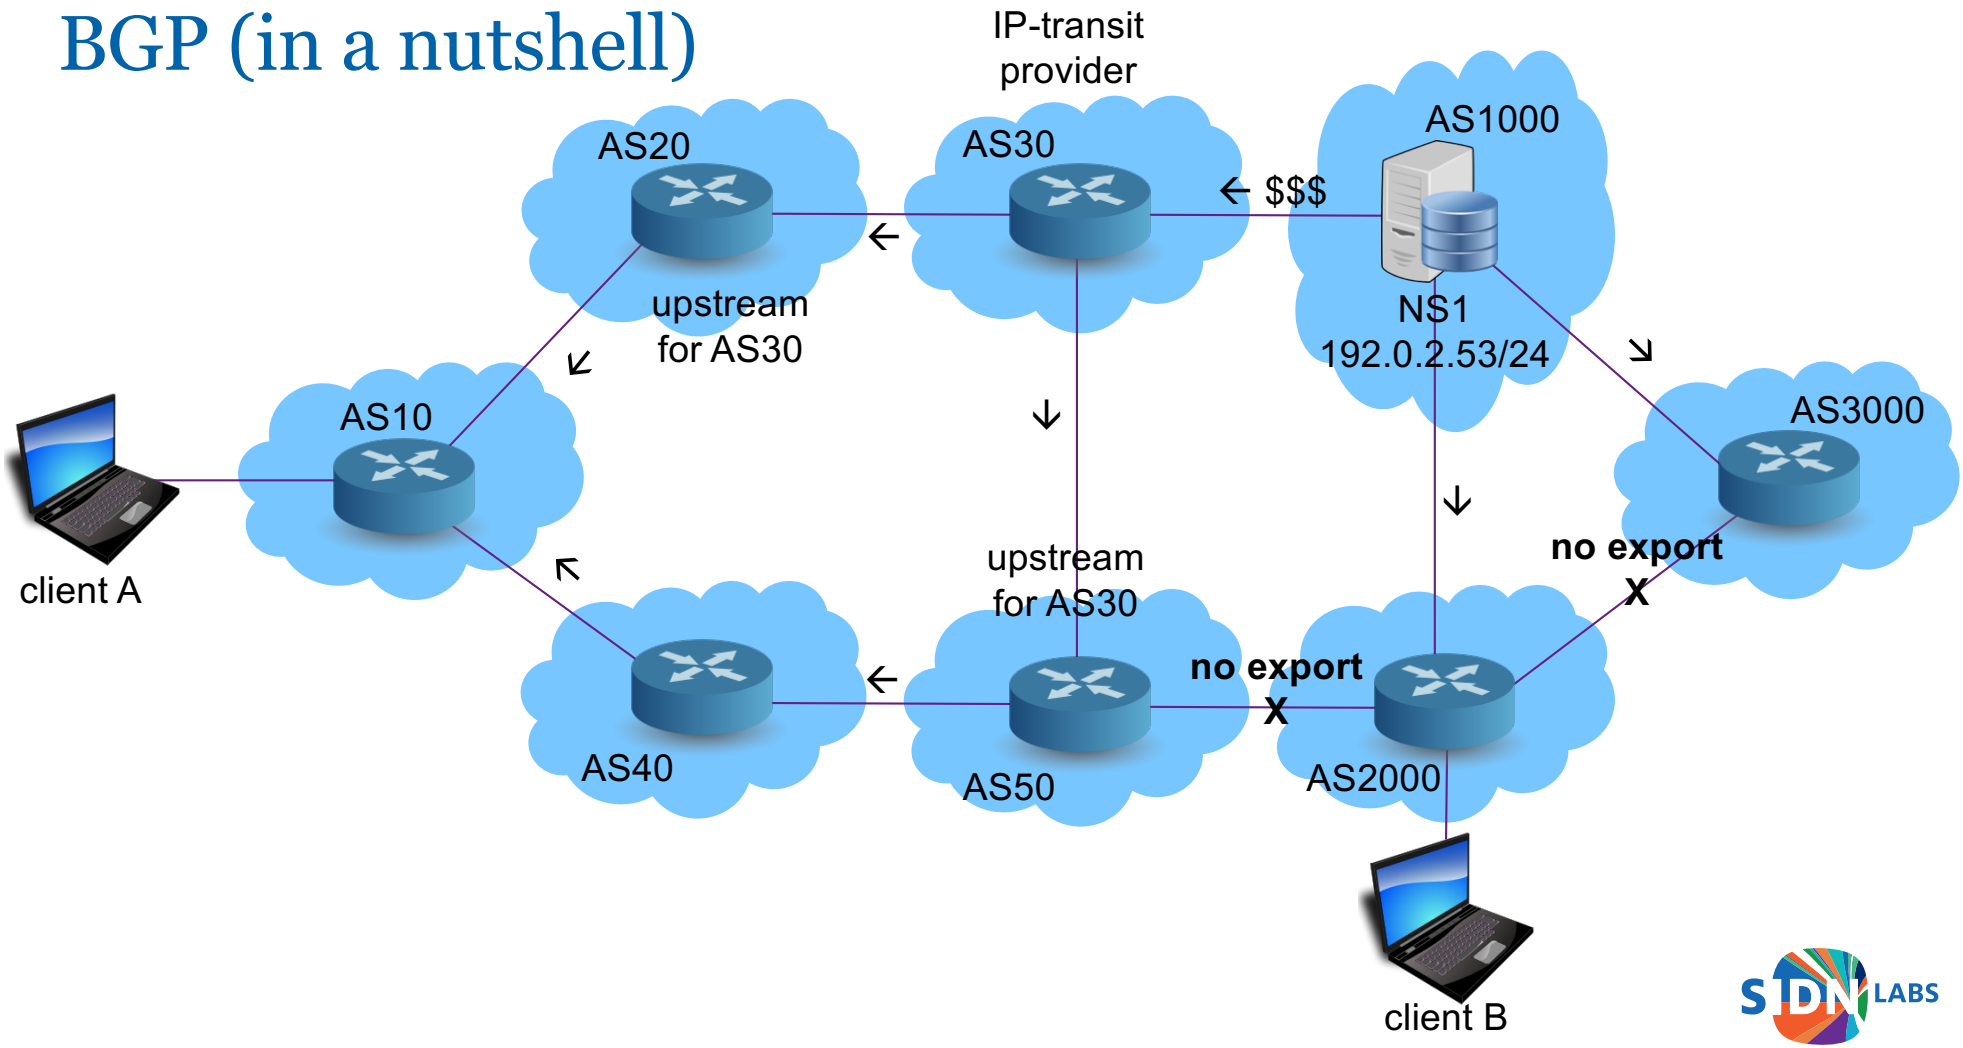

Some terminlogy

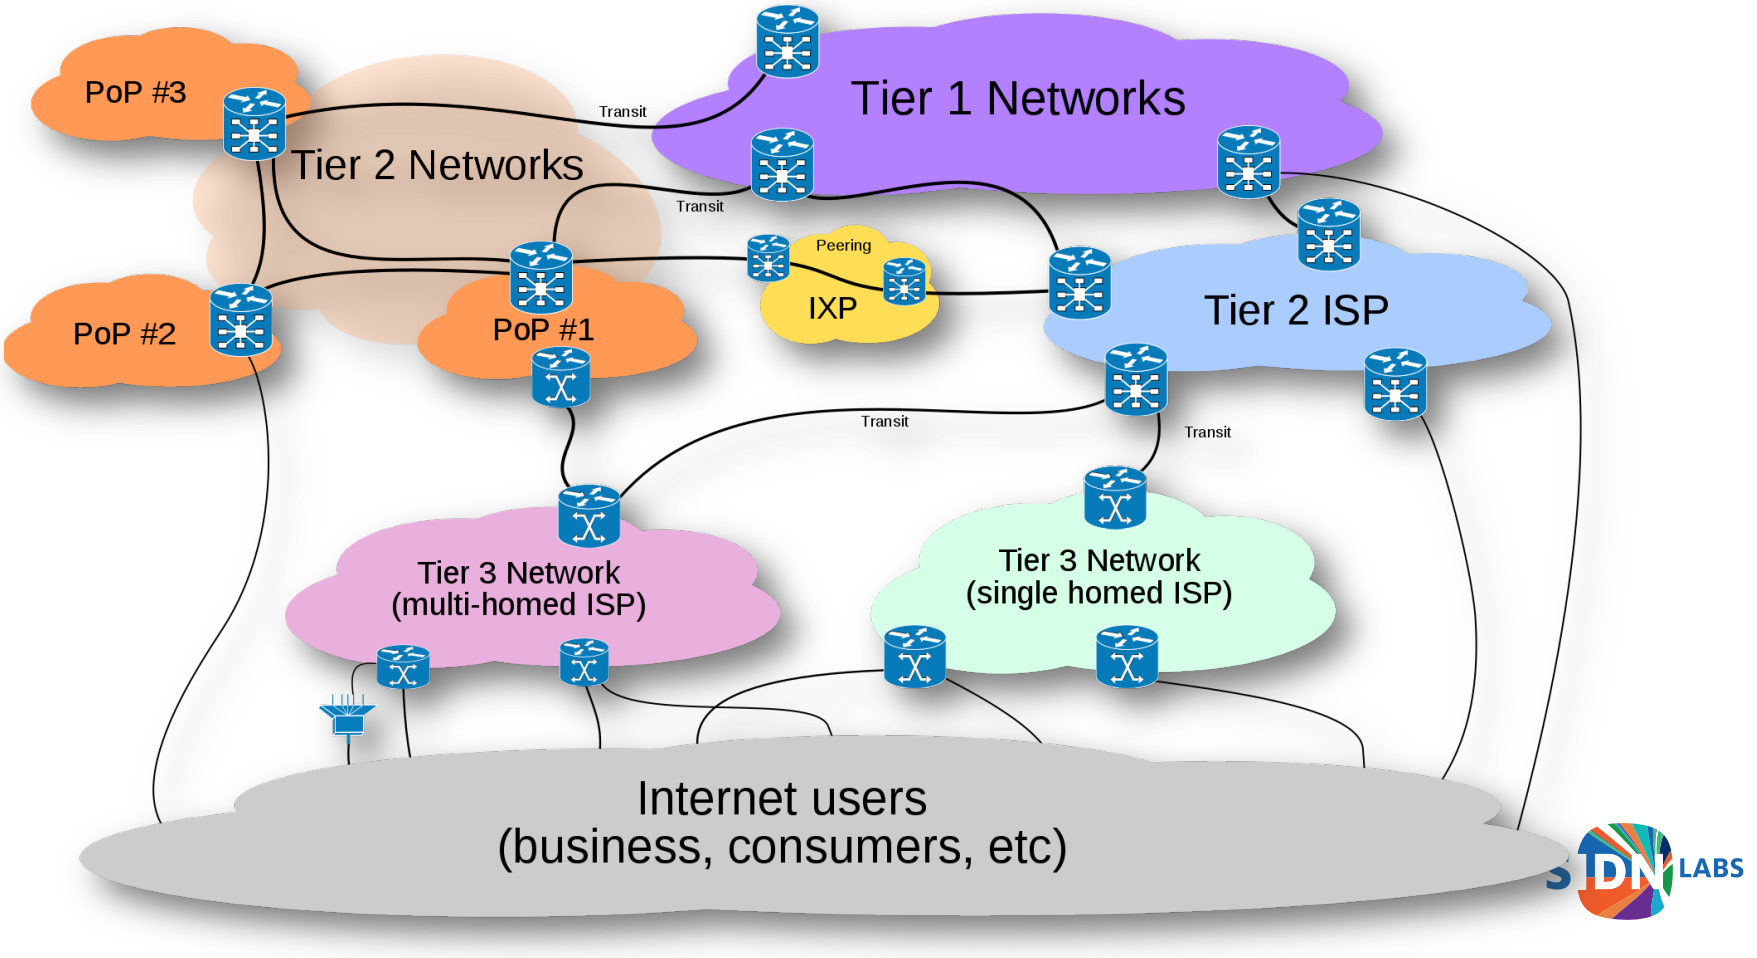

BGP (in a nutshell)

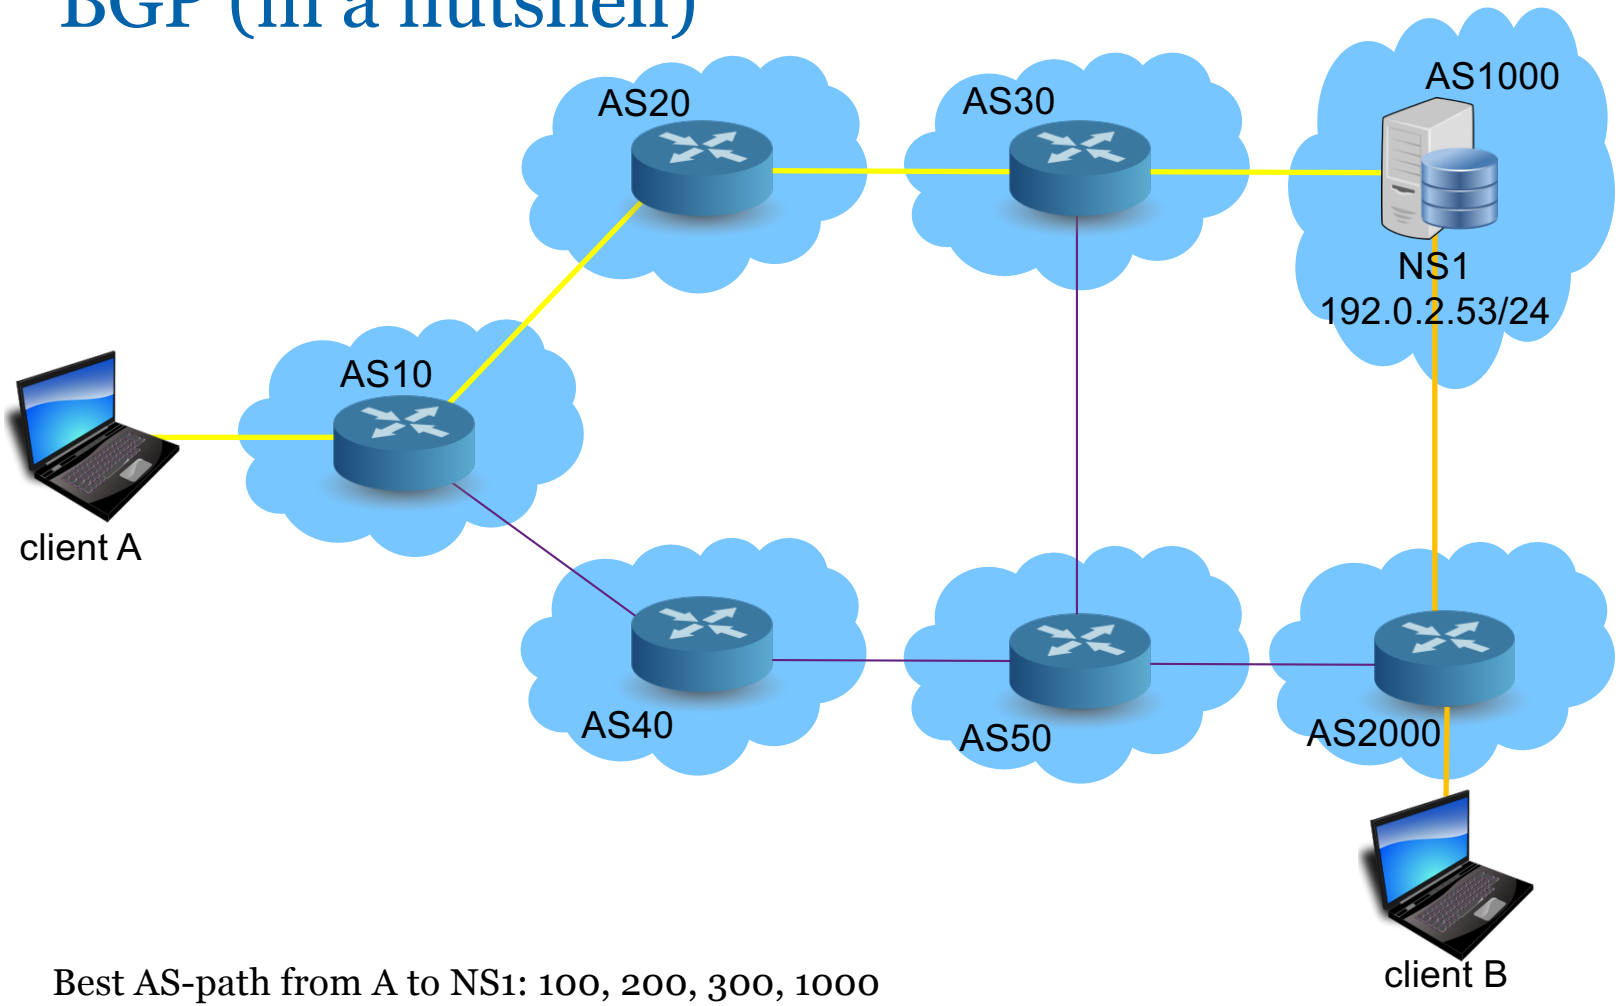

Best AS-path from A to NS1: 100, 200, 300, 1000

BGP (in a nutshell)

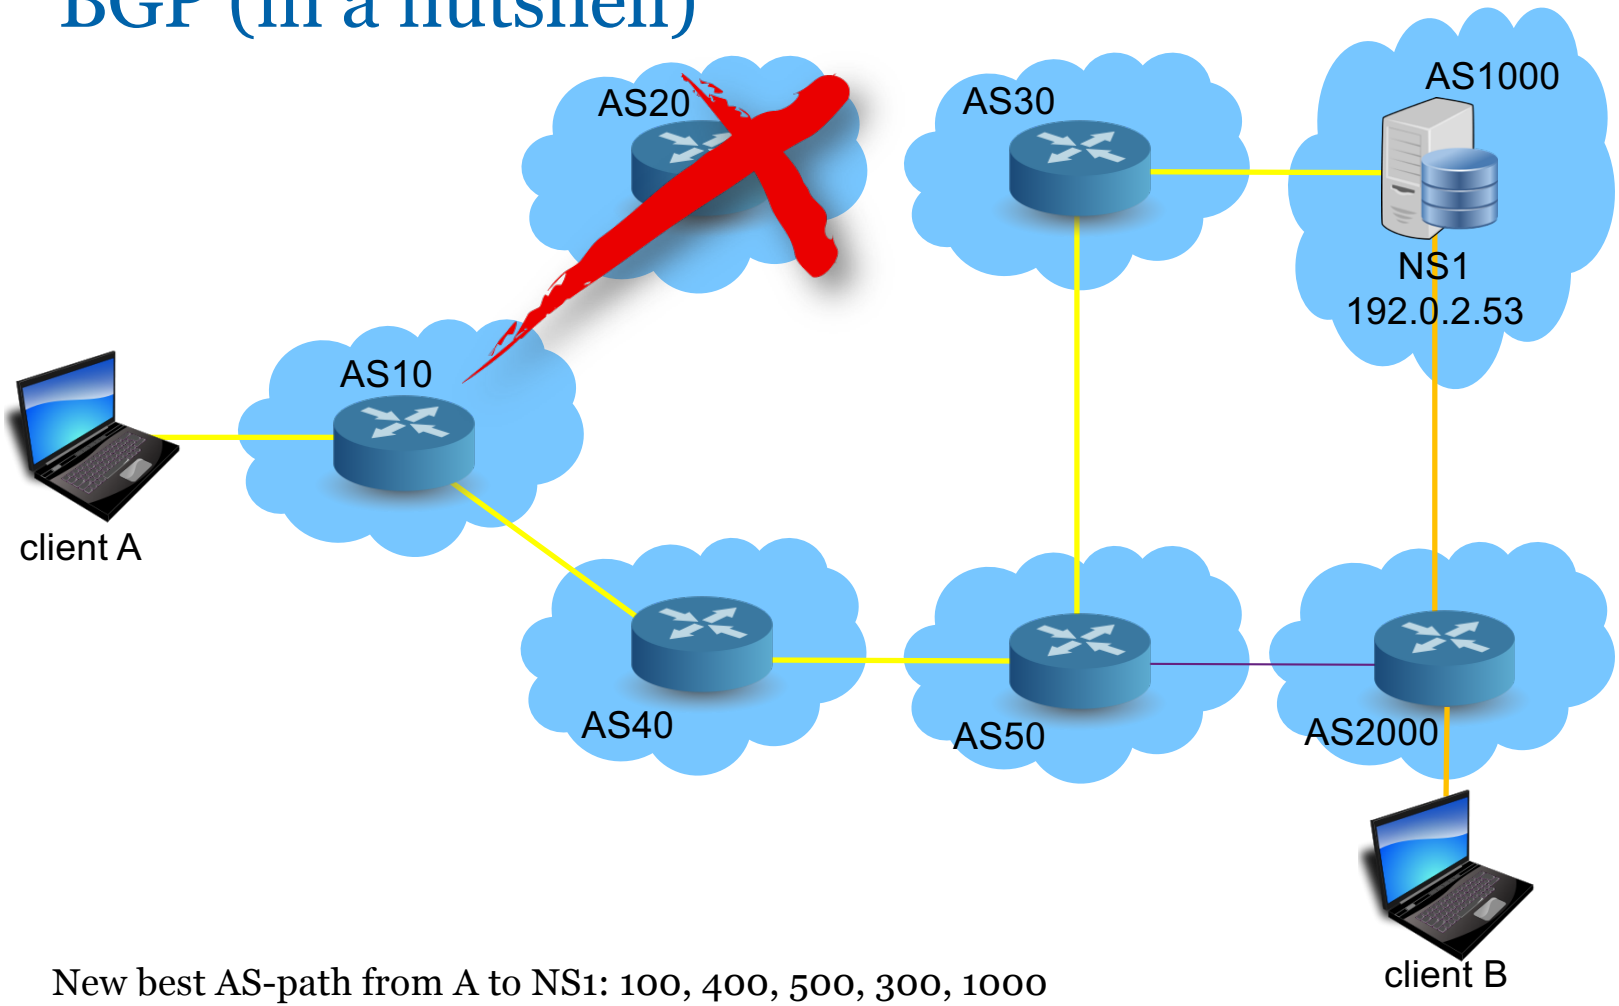

New best AS-path from A to NS1: 100, 400, 500, 300, 1000

BGP traffic engineering: AS-path prepending

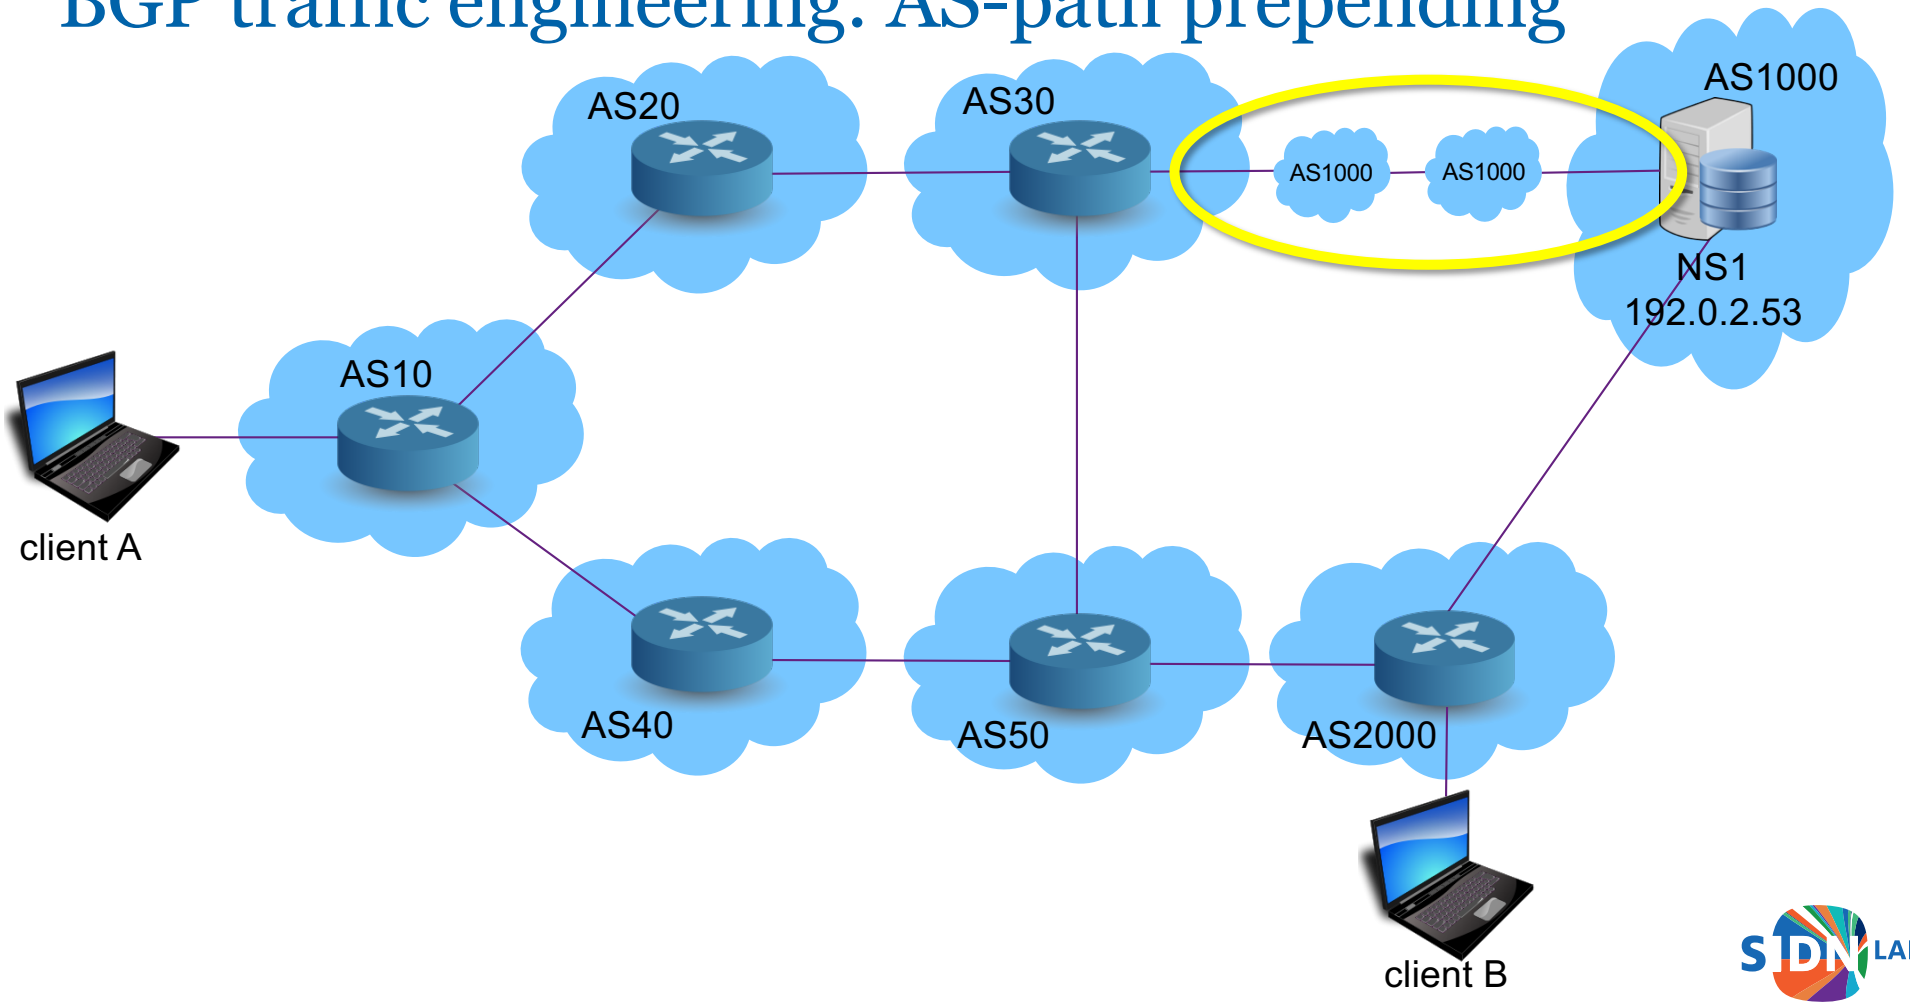

BGP traffic engineering: AS-path prepending

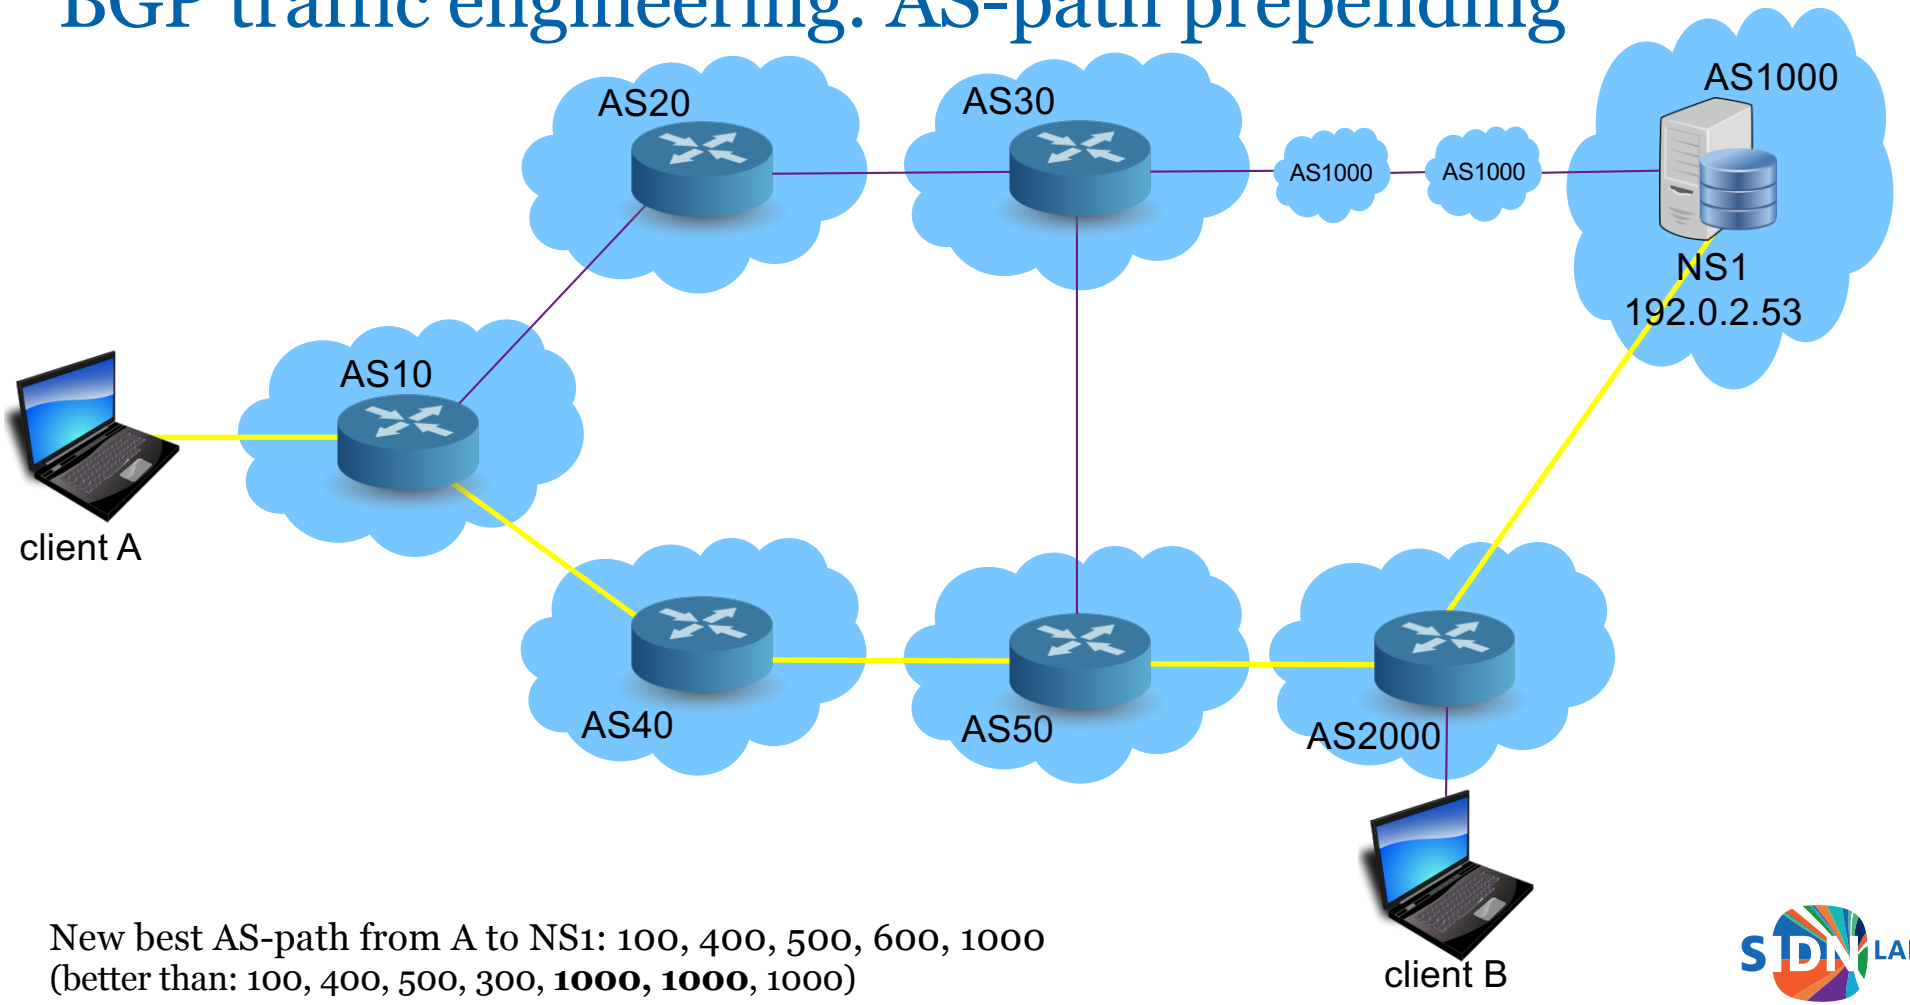

New best AS-path from A to NS1: 100, 400, 500, 600, 1000
(better than: 100, 400, 500, 300, **1000, 1000**, 1000)

BGP traffic engineering: local pref

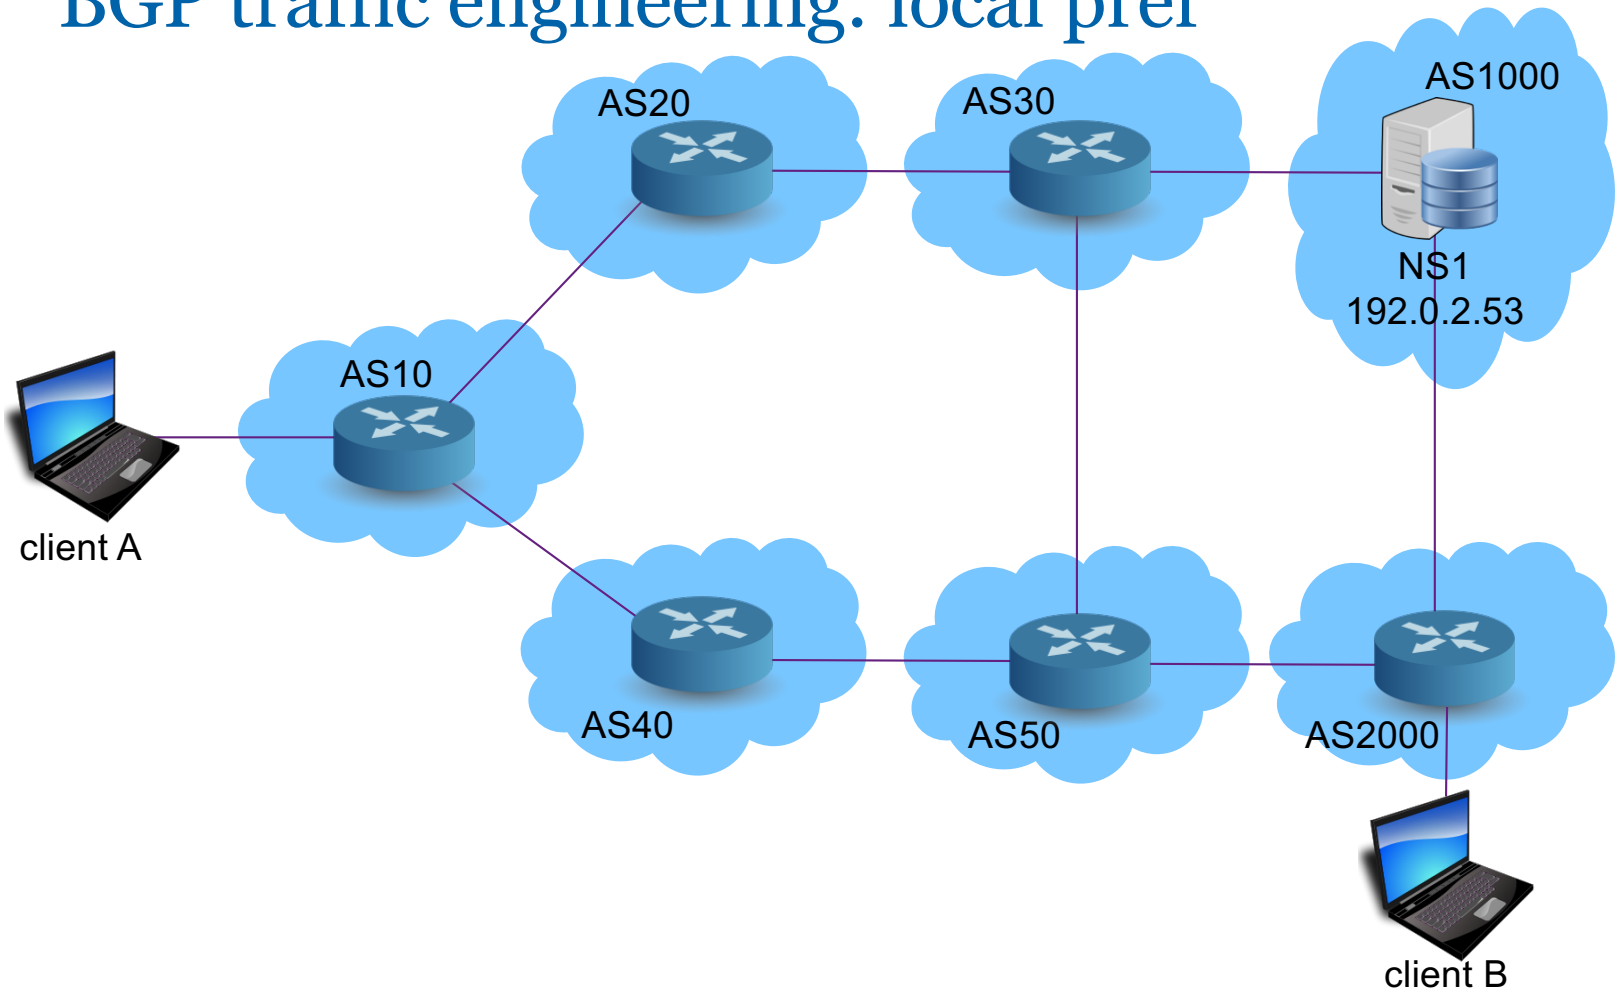

BGP traffic engineering: local preference

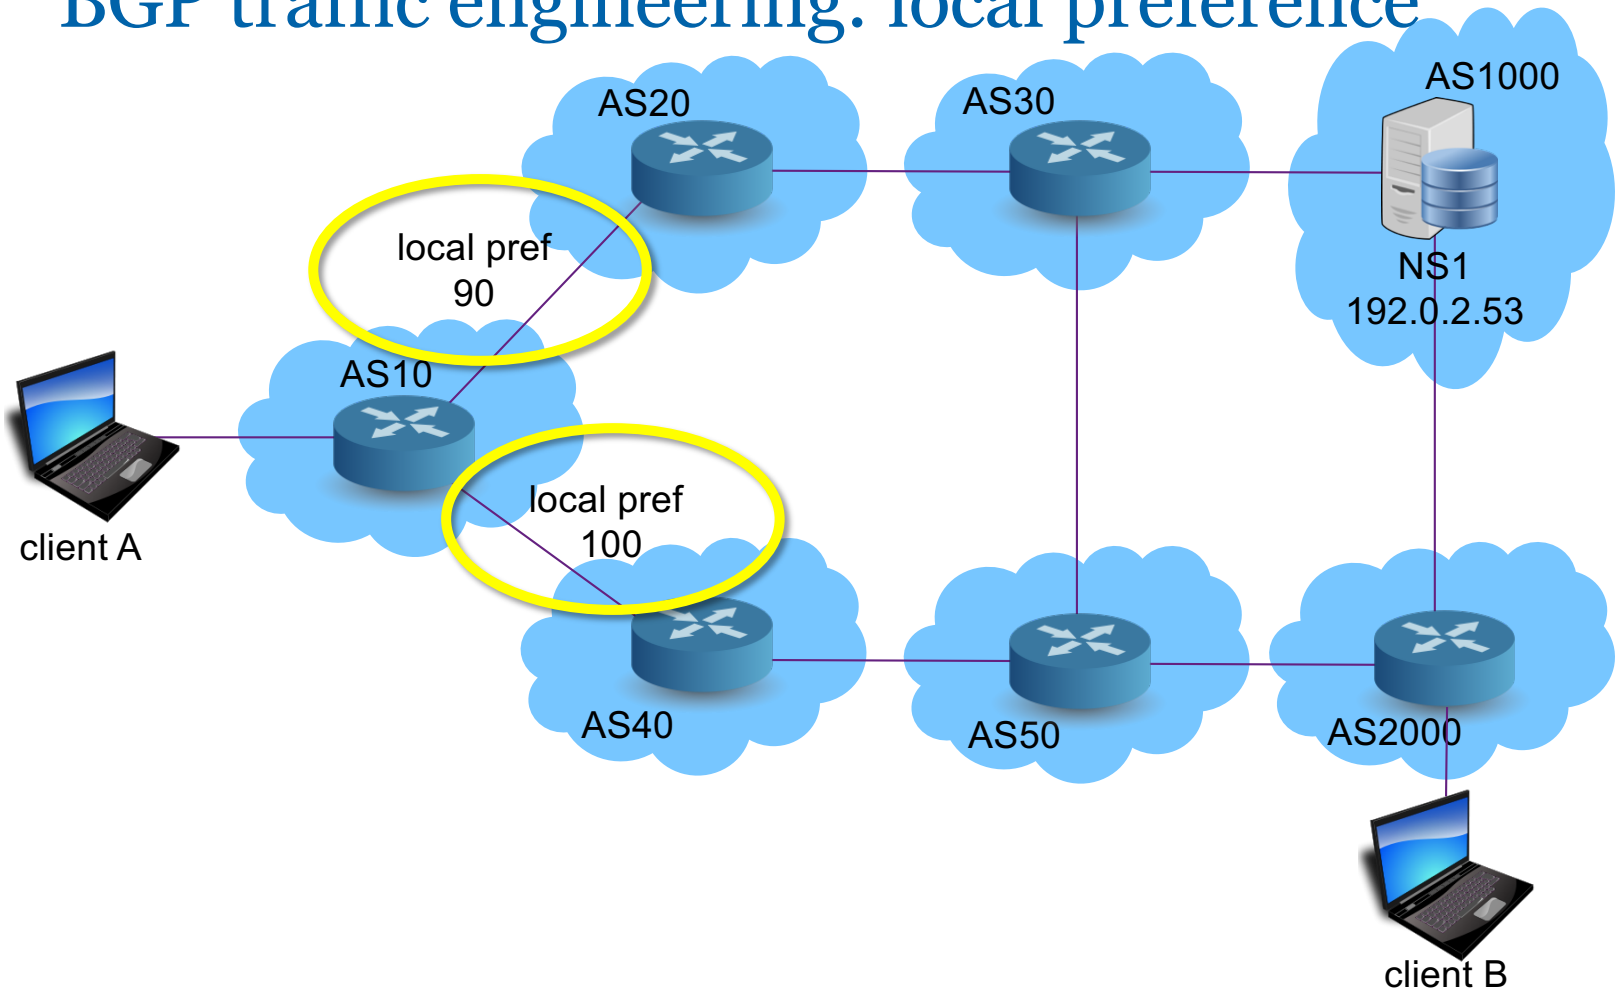

BGP traffic engineering: local preference

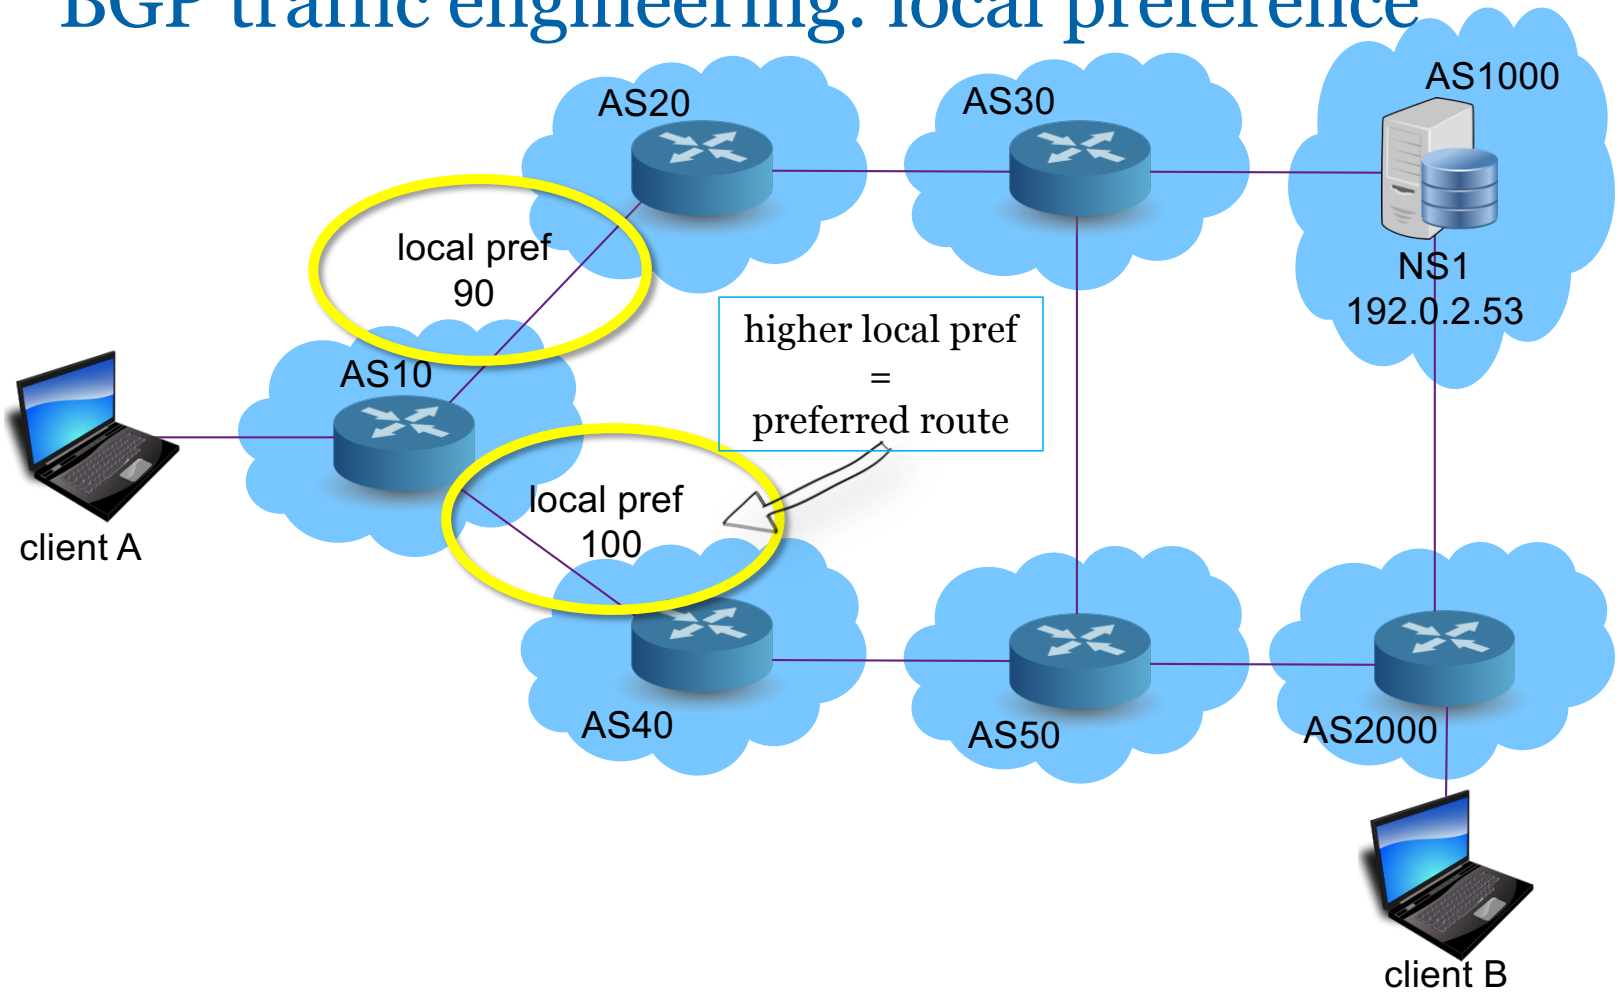

BGP traffic engineering: local preference

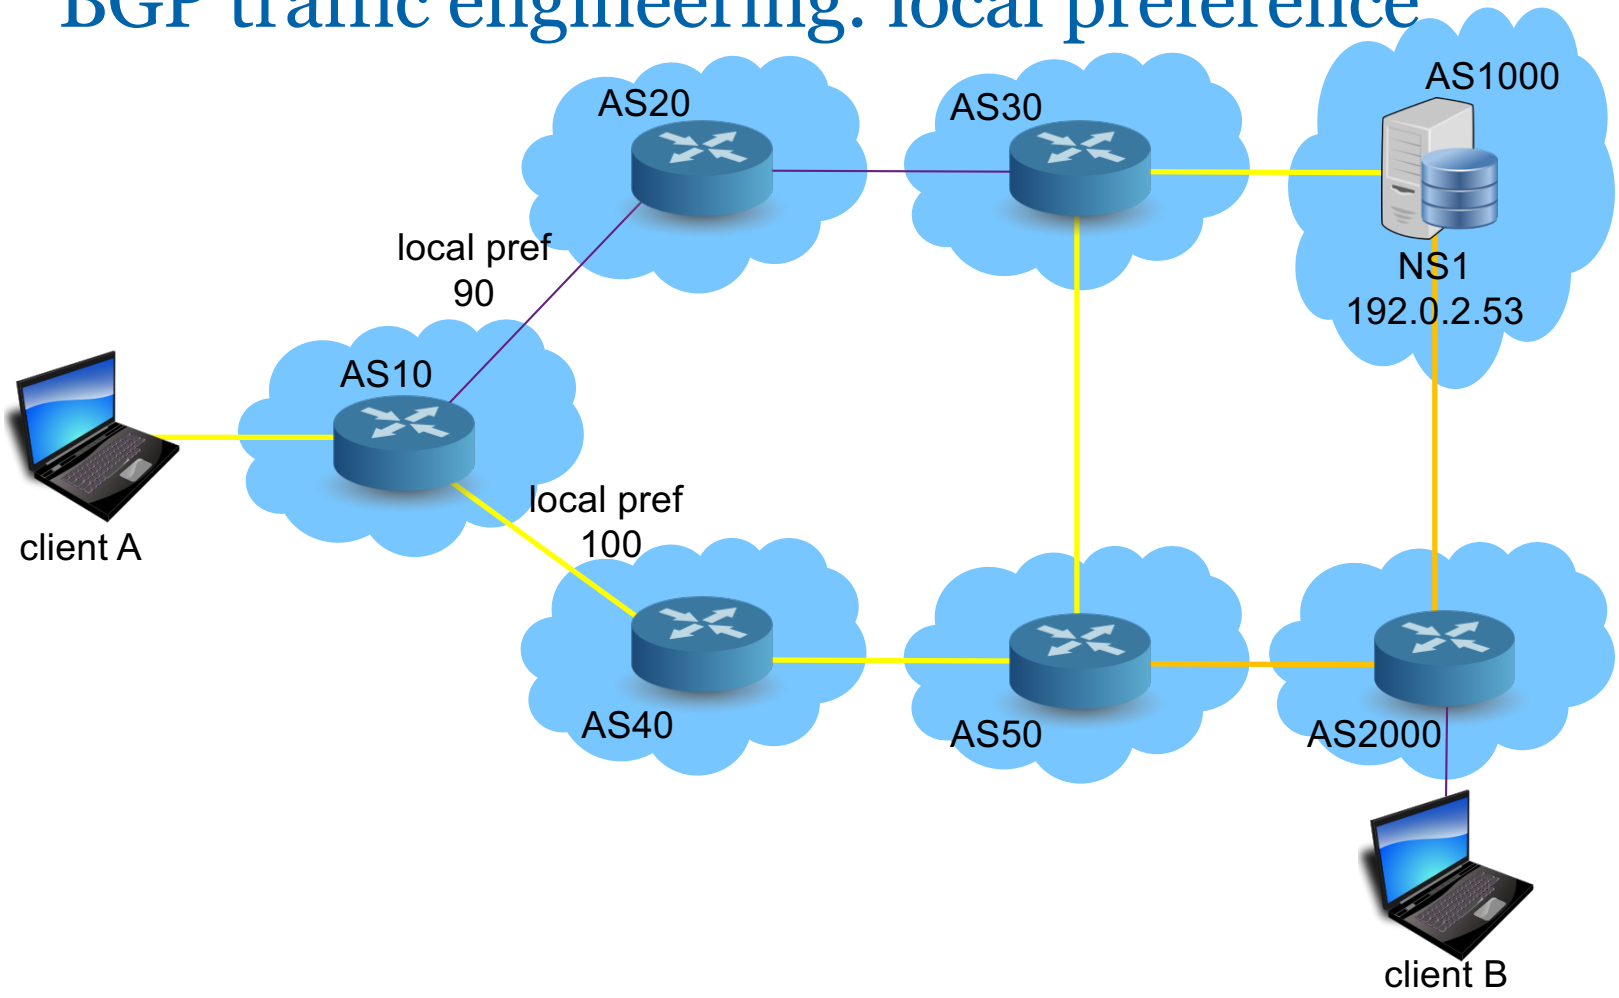

BGP traffic engineering: local preference

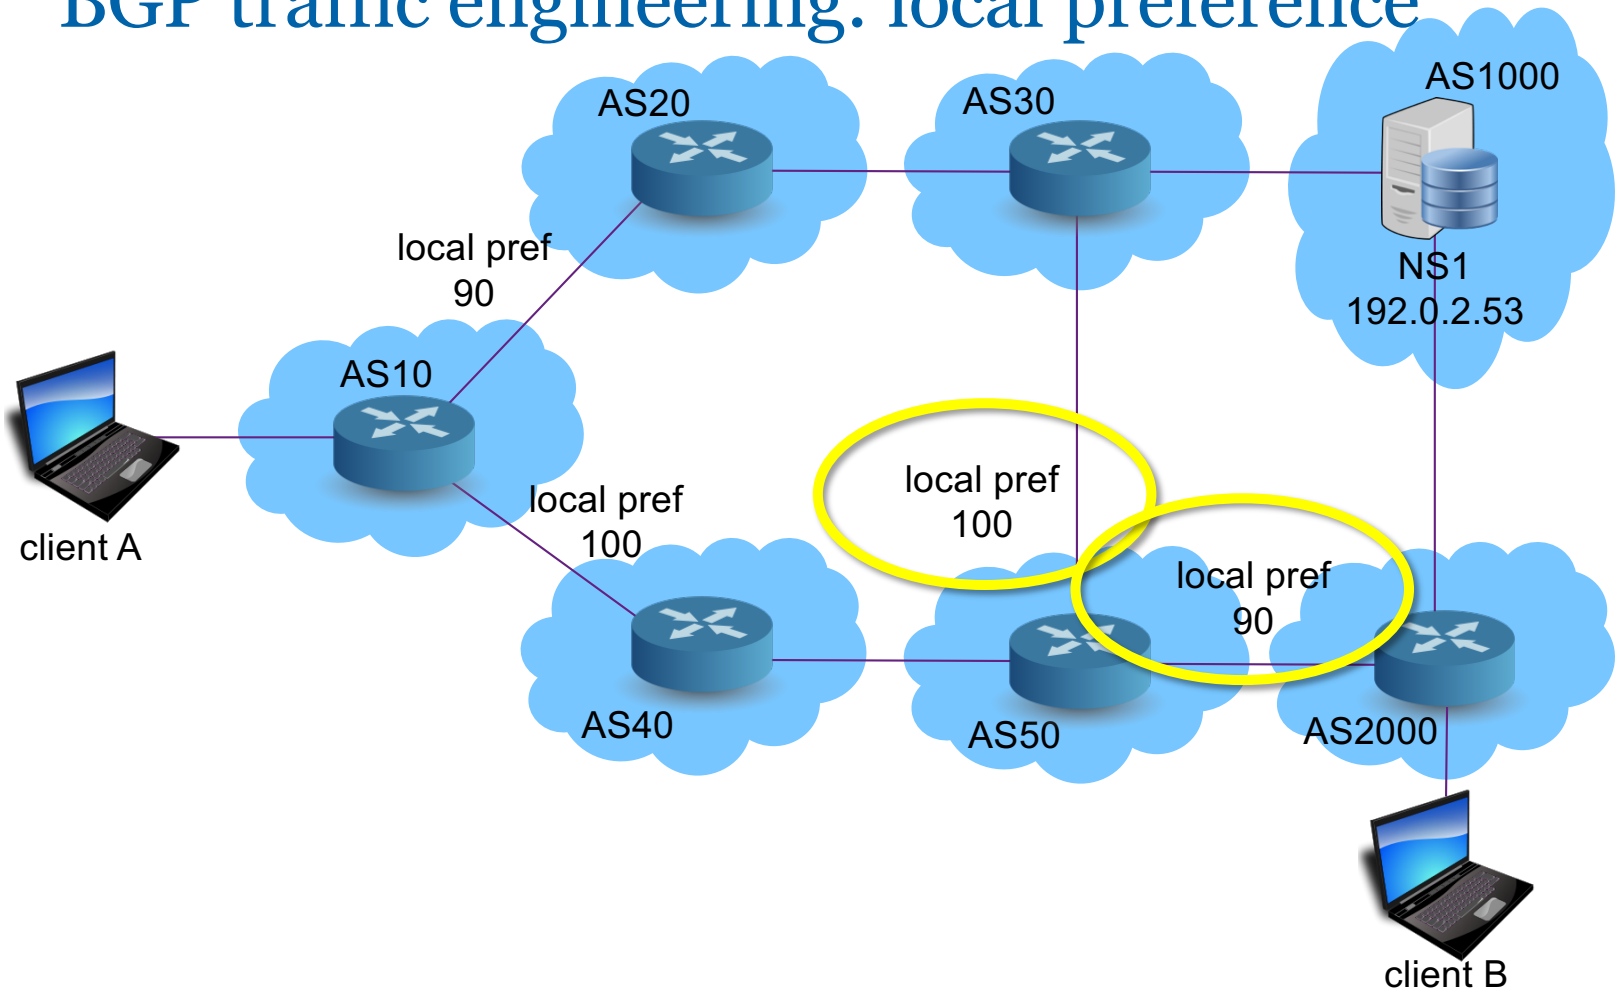

BGP traffic engineering: local preference

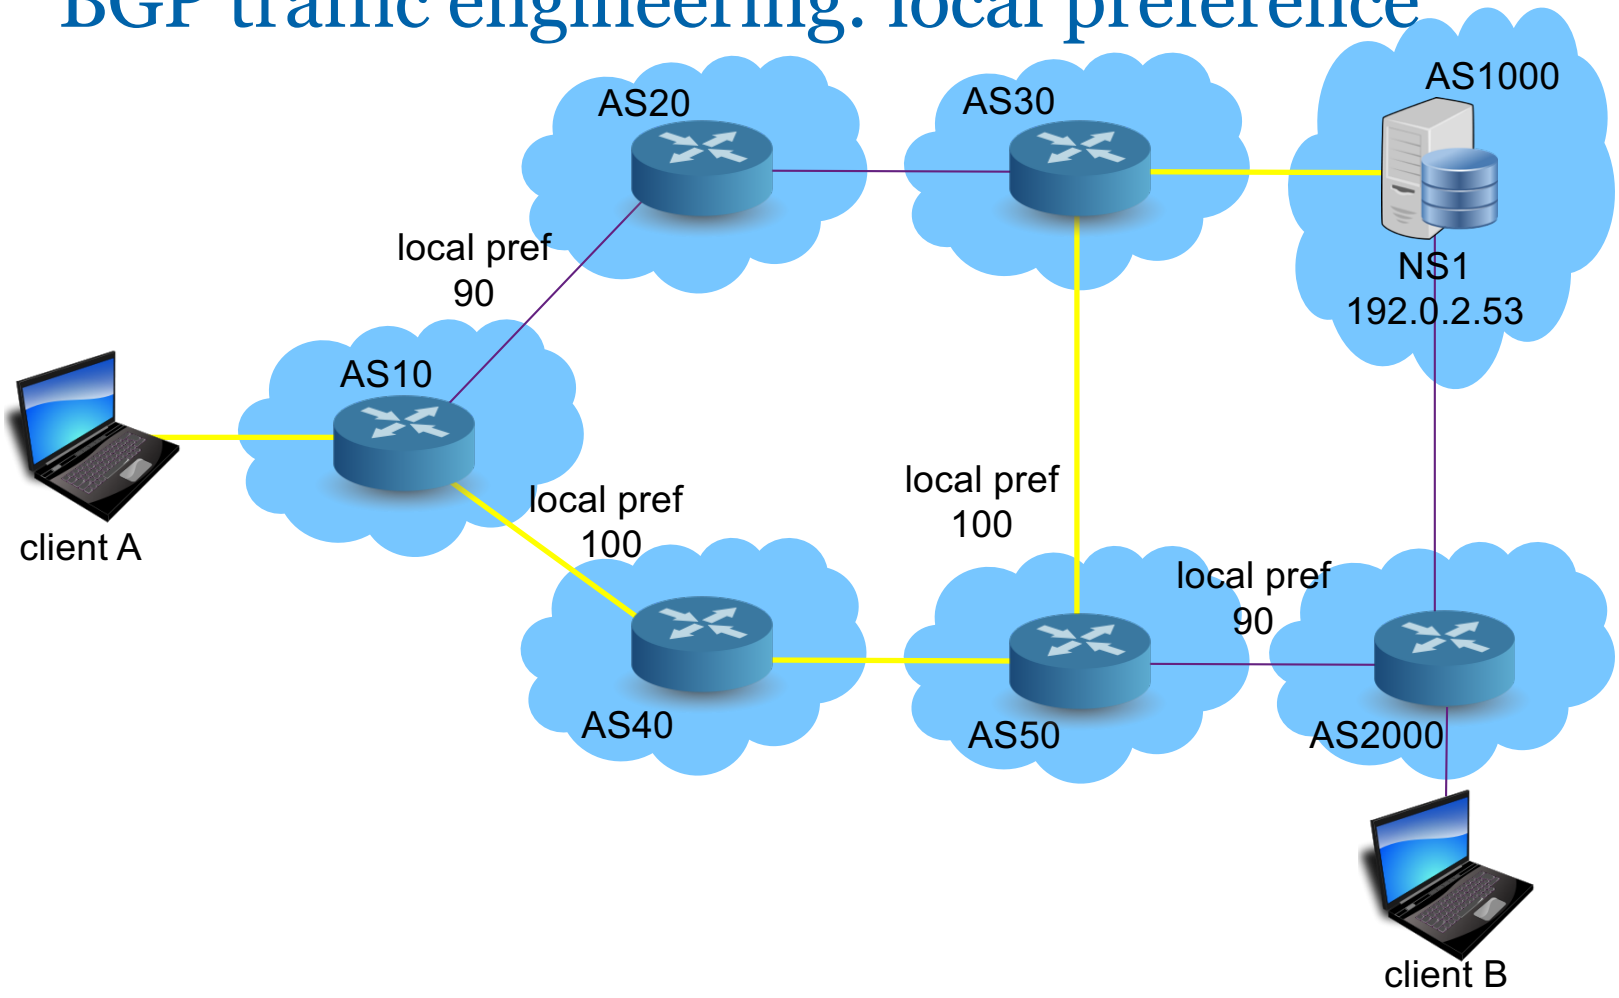

Traffic engineering with BGP communities

- Transitive attribute tags that can be applied on incoming or outgoing prefixes to achieve a certain goal.
- For example: local pref adjustments, geographic restictions, AS-path prepending or blackholing.
- No universal definitions, except some well-known ones

```
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: Resource Public Key Infrastructure

- A public key infrastructure framework designed to secure BGP
- Resource certification of IP-prefixes / ASN combination
- Prevents (to some extend) route-hijacking

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

RPKI (in a nutshell)

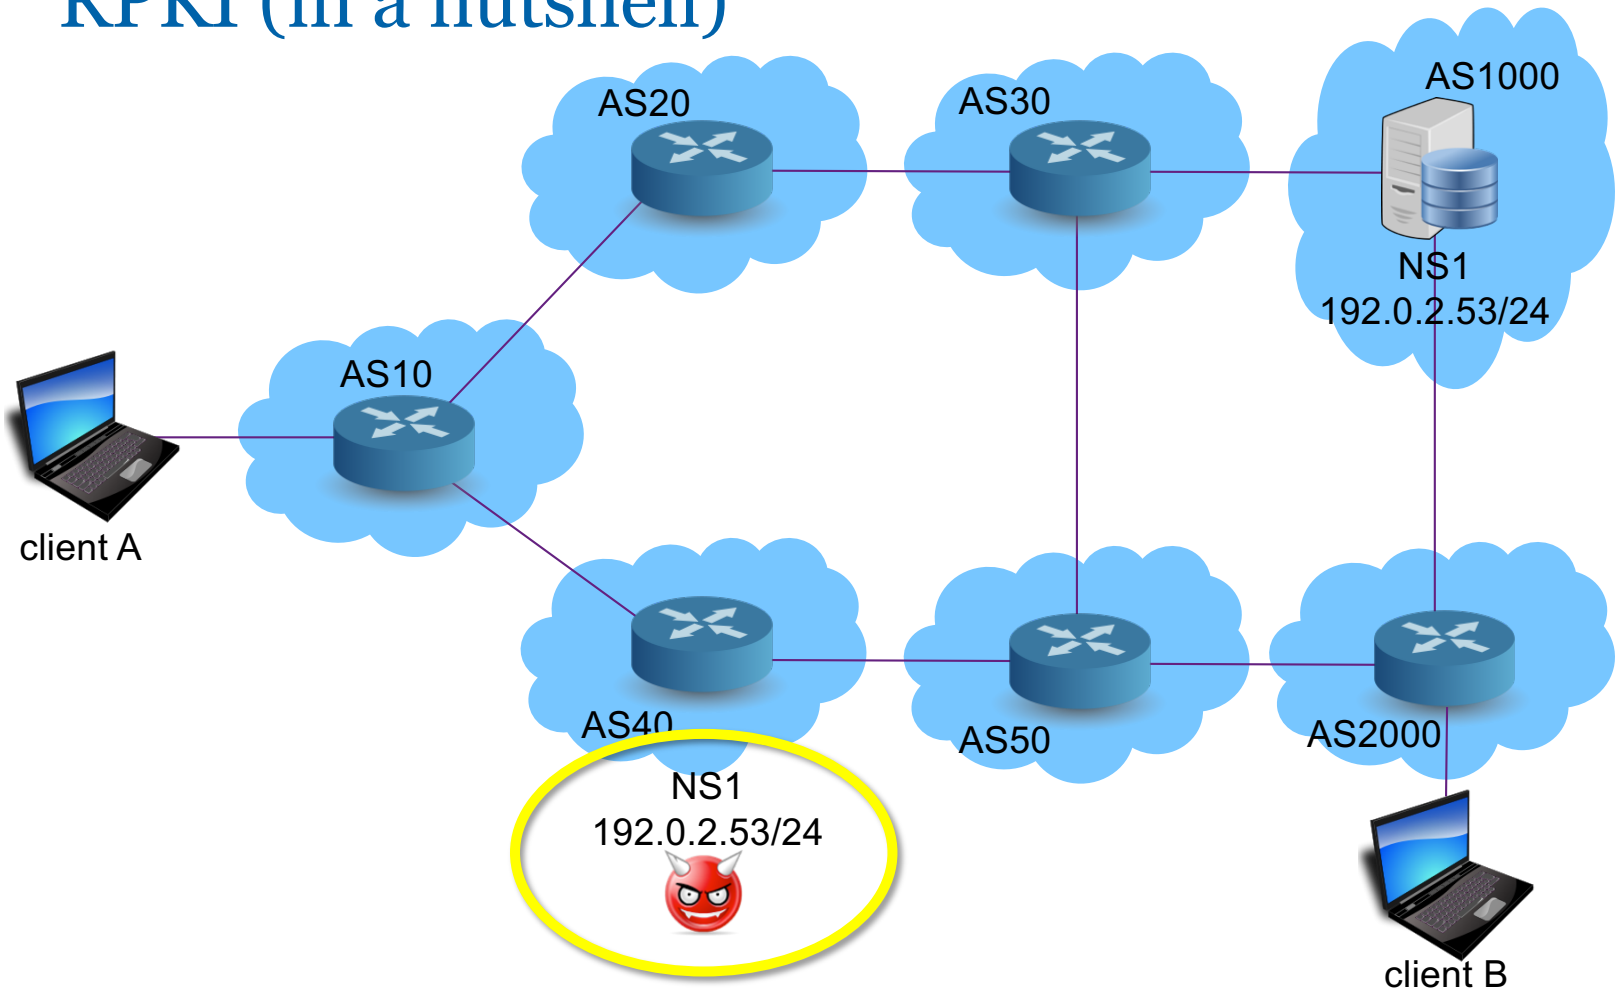

RPKI (in a nutshell)

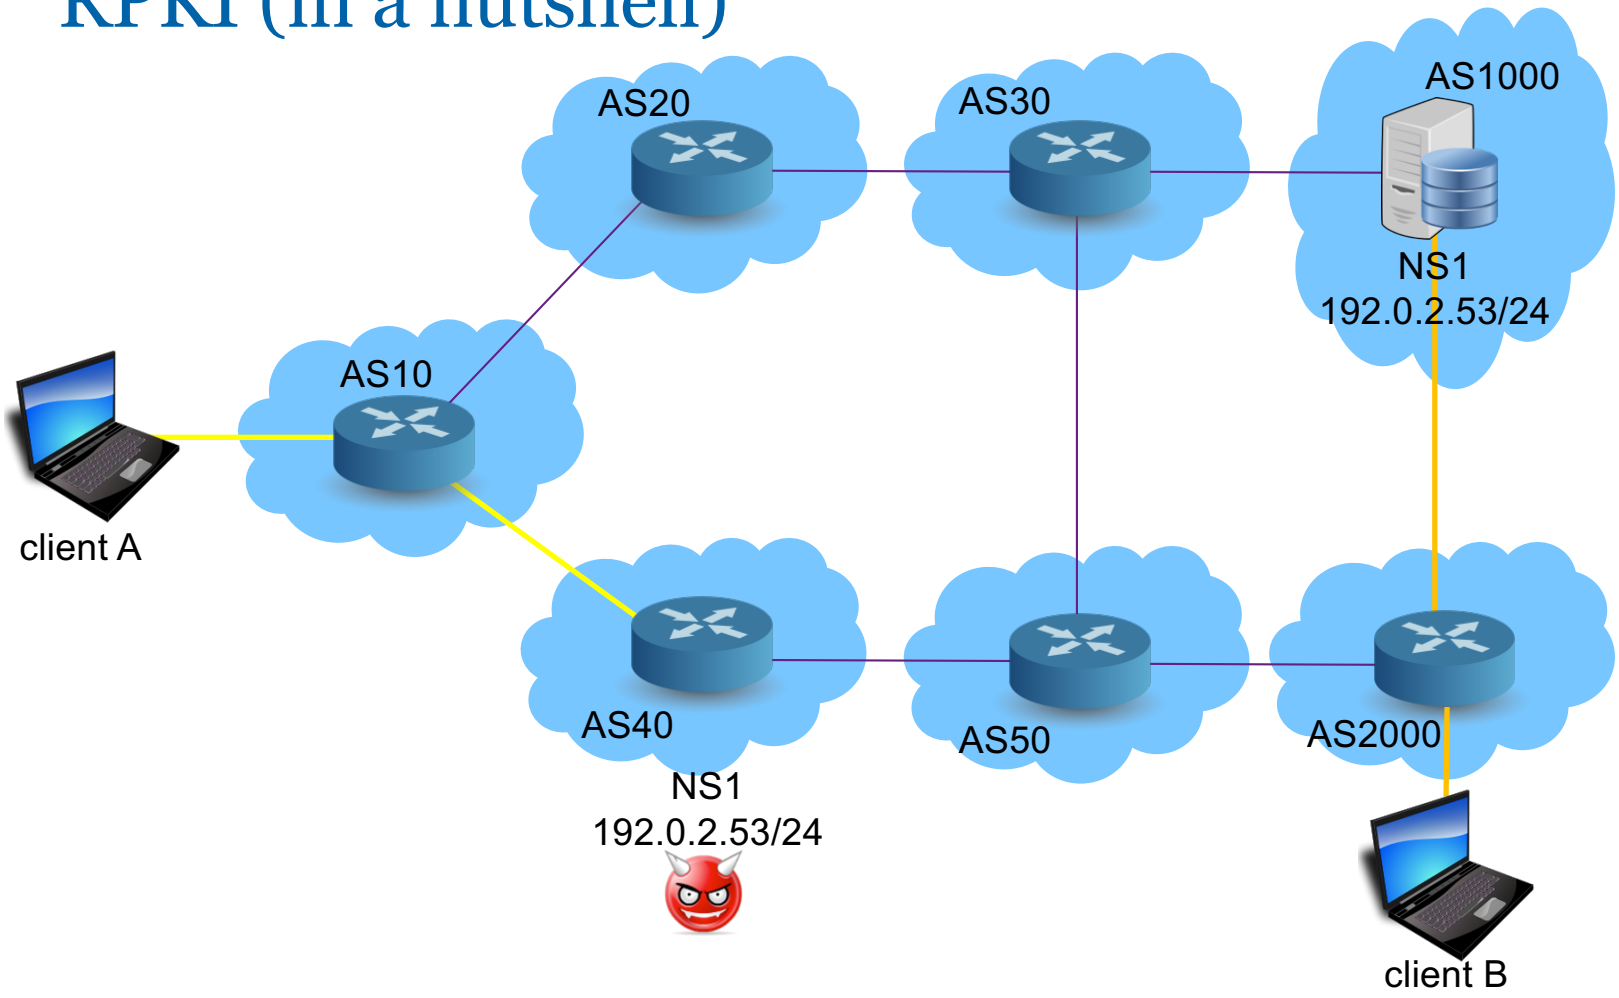

Challenges: RTT

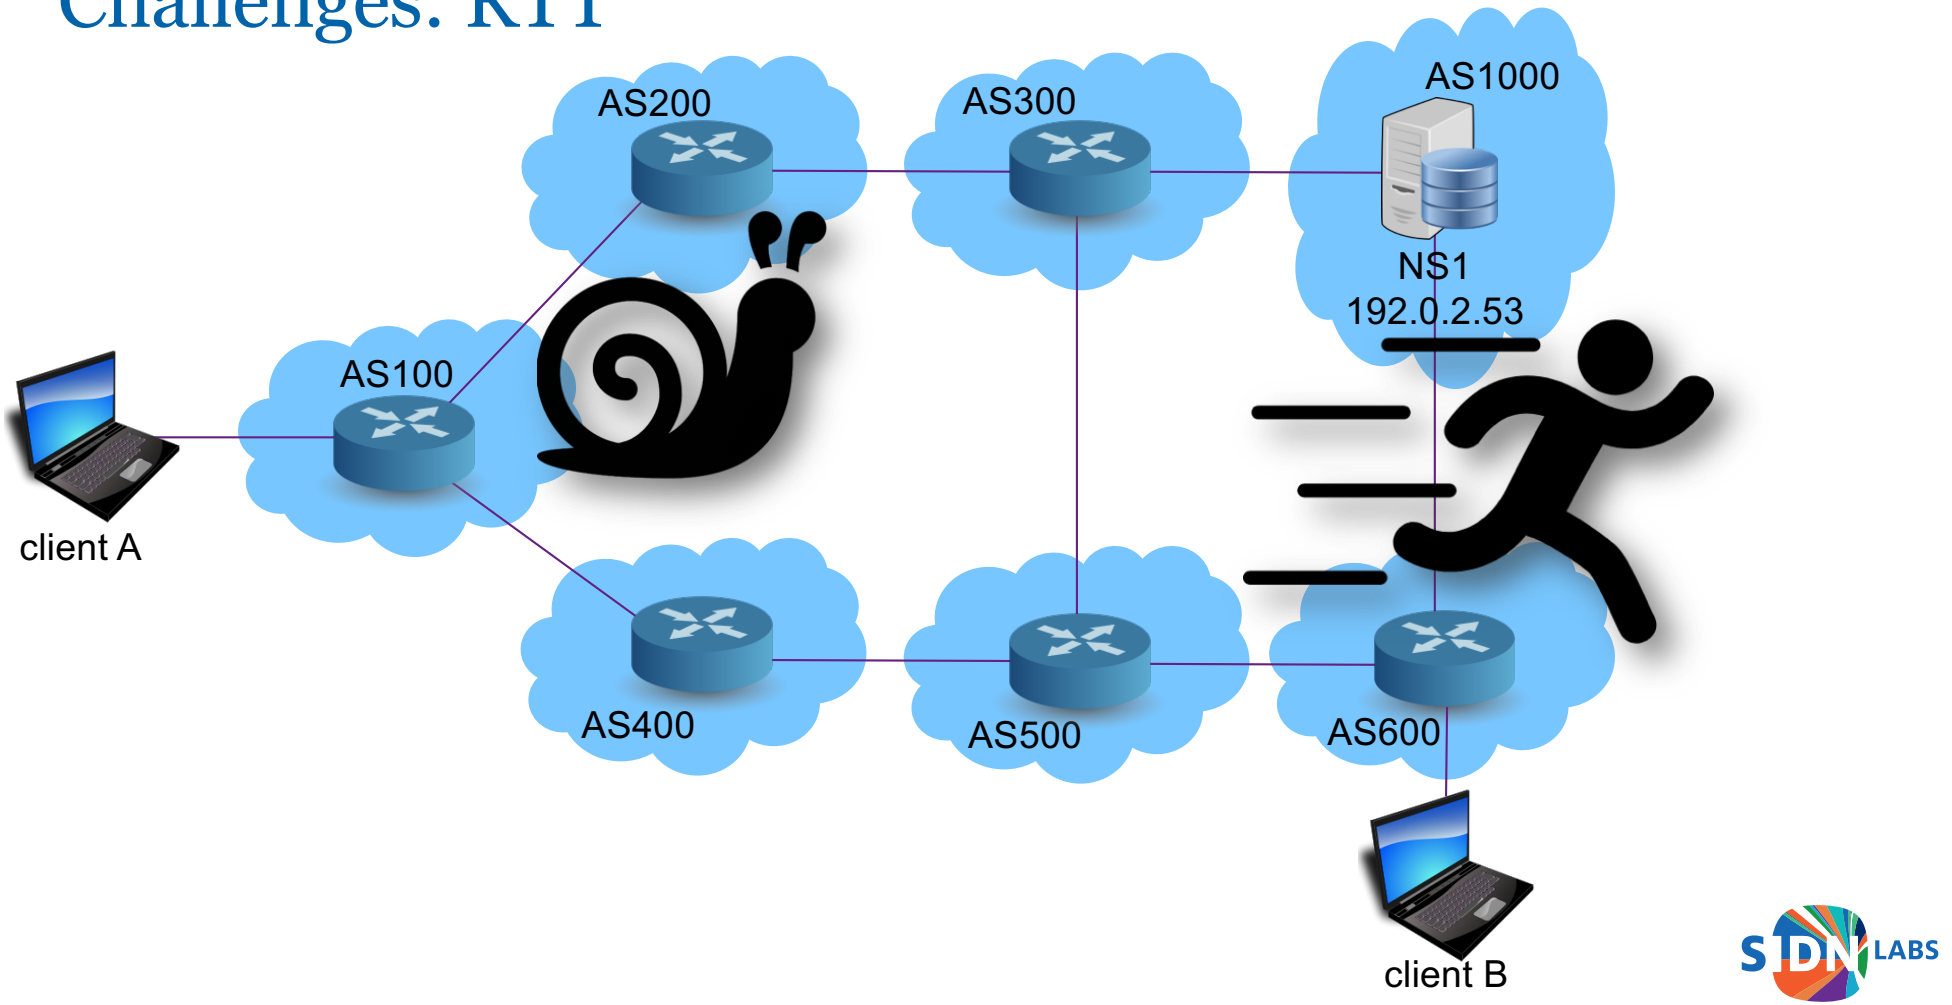

Challenges: DDoS (record breaking sometimes)

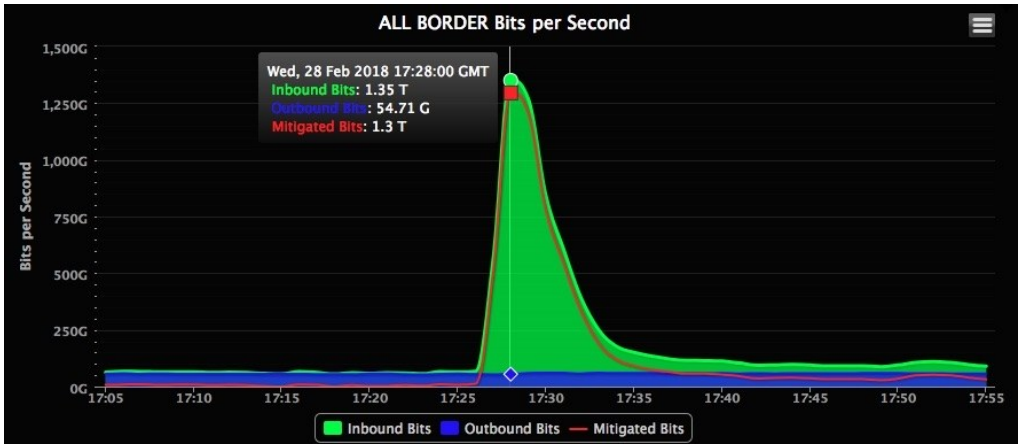

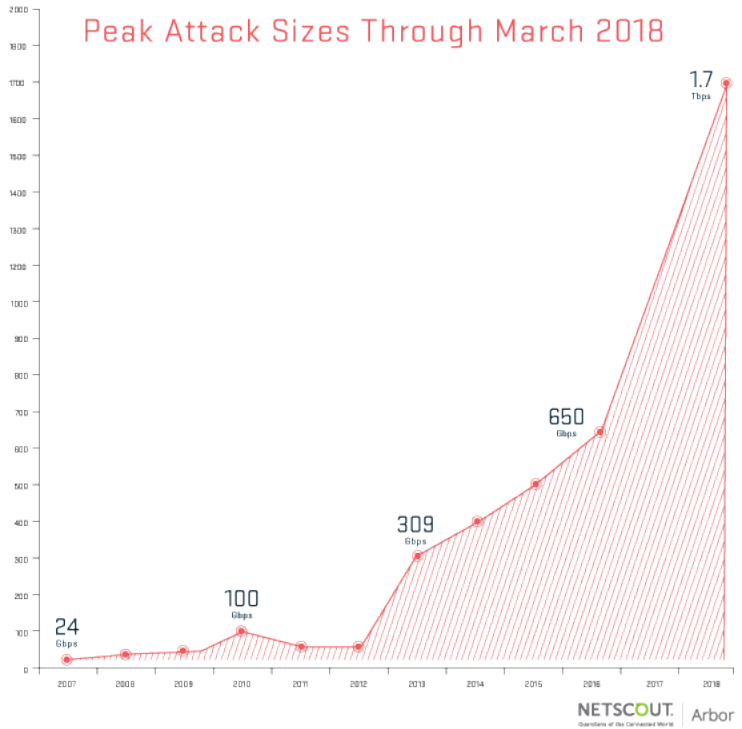

Main reasons: IoT devices

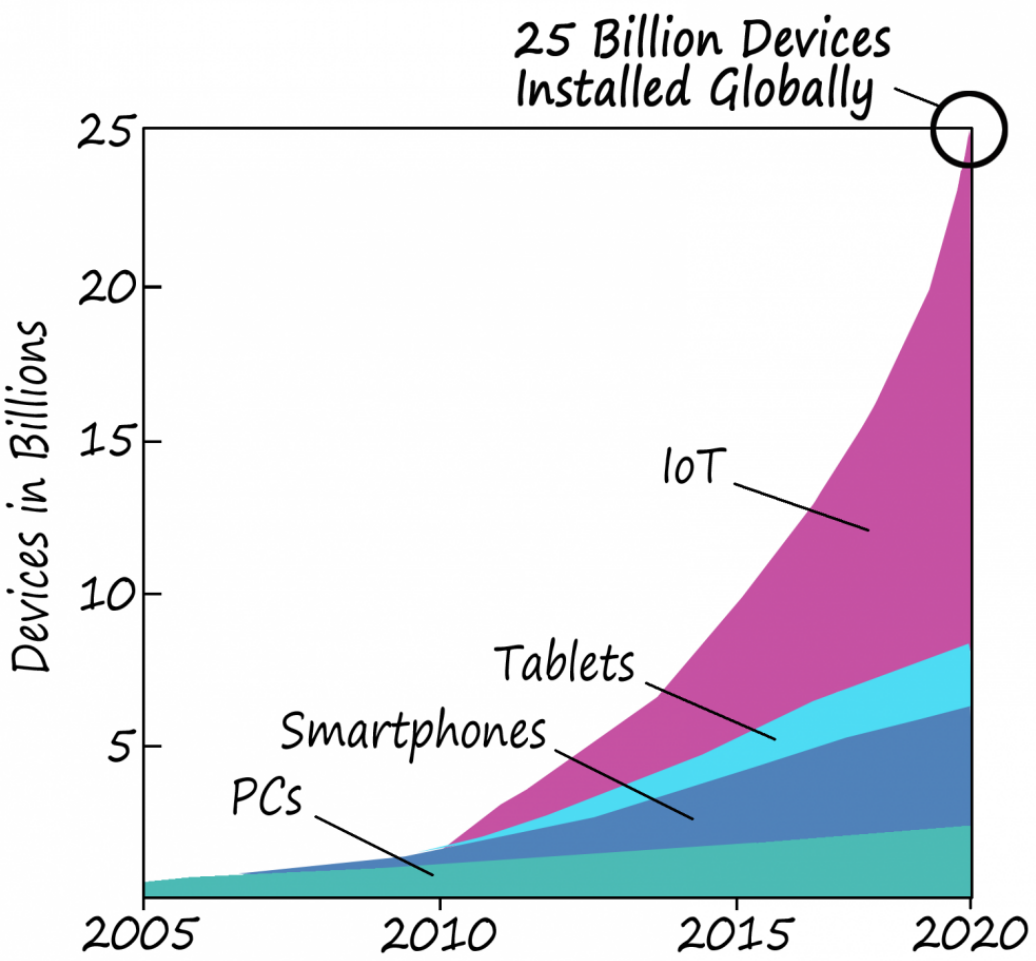

IoT powered botnets

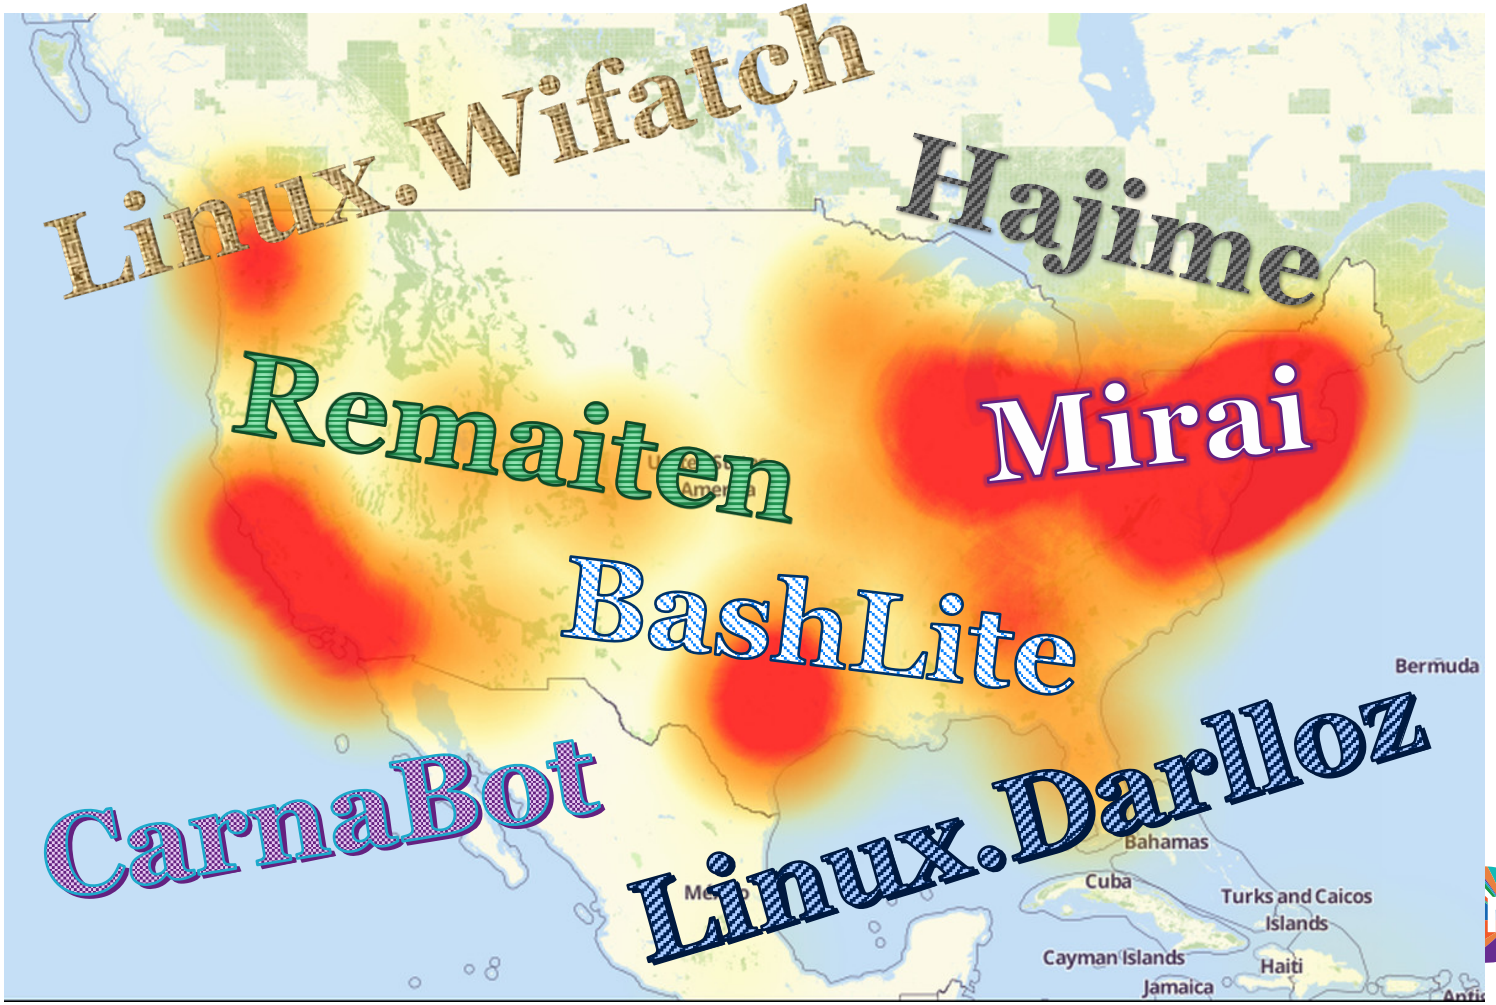

The solution to both challenges: DNS **global** anycast

- Just a clever ‘network hack’ to provide (a lot of) resilience.
  - And better performance (shorter RTT’s)
- Works with BGP
- Well understood solution, deployed in many places
  - The DNS root servers (for many years)
  - 1.1.1.1, 8.8.8.8, 9.9.9.9, 64.6.64.6, OpenDNS and more
- Originally only in UDP environments
  - But proven in TCP environments as well (i.e. CloudFlare)

DNS global anycast

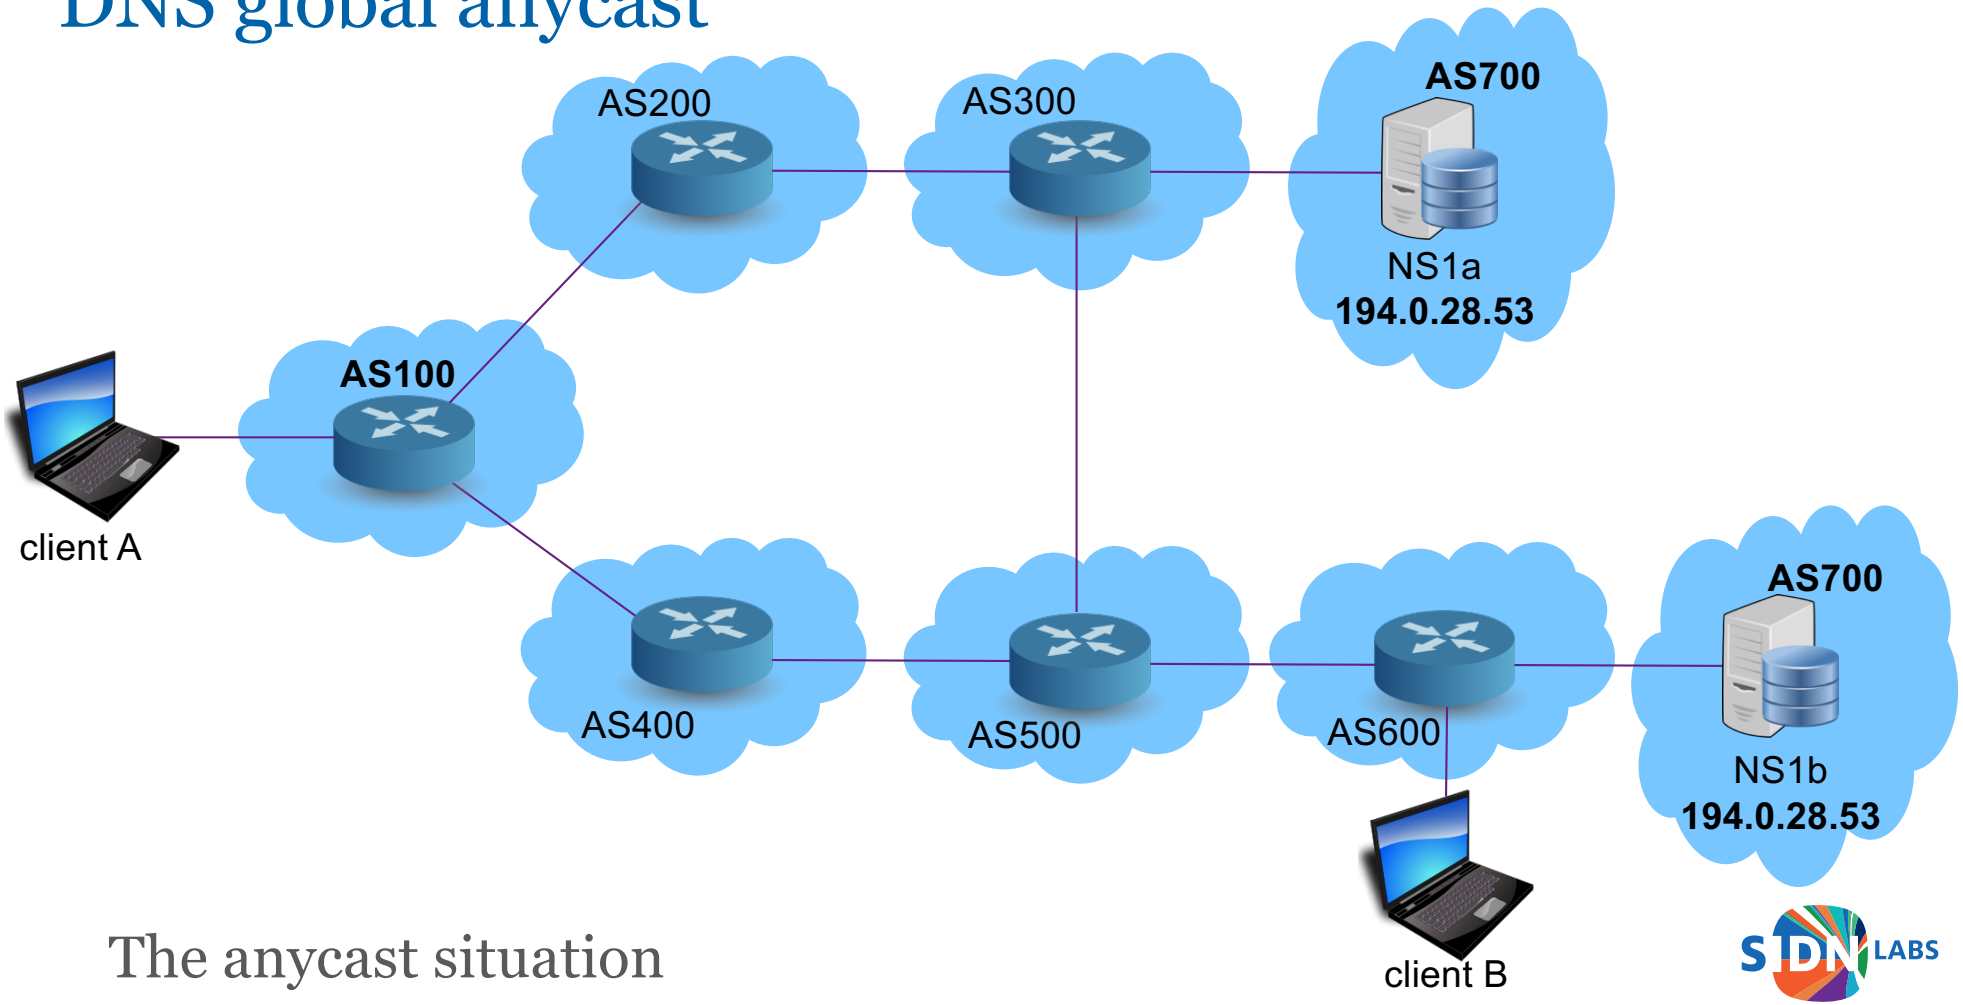

The anycast situation

DNS global anycast

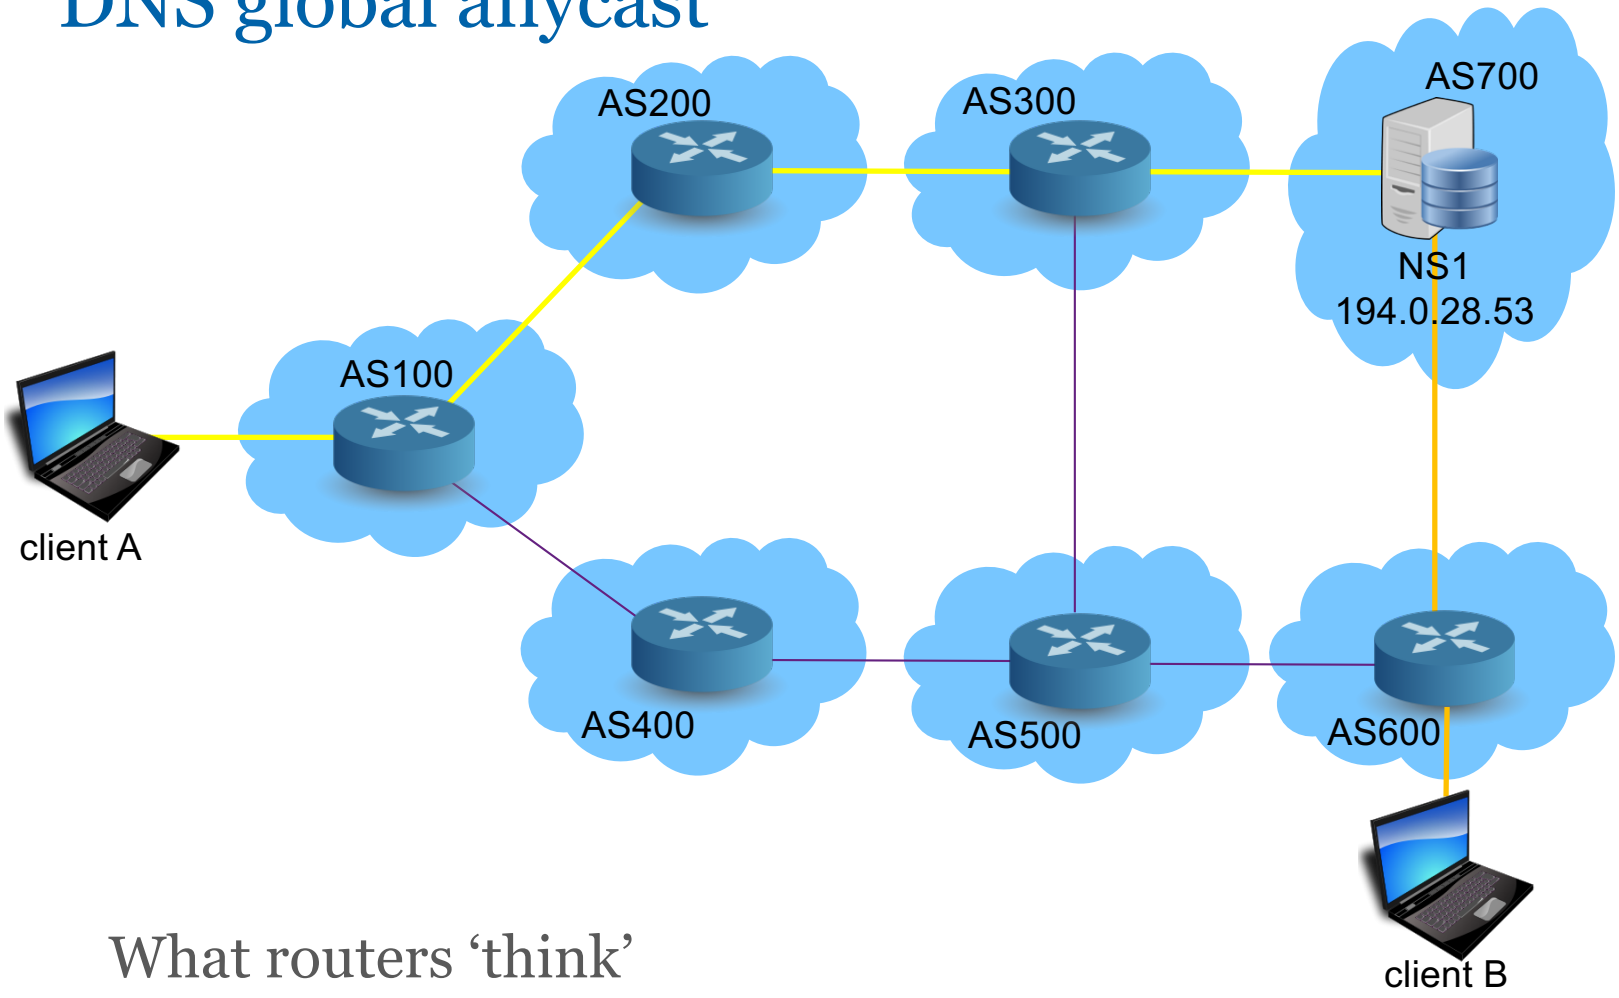

What routers ‘think’

DNS global anycast

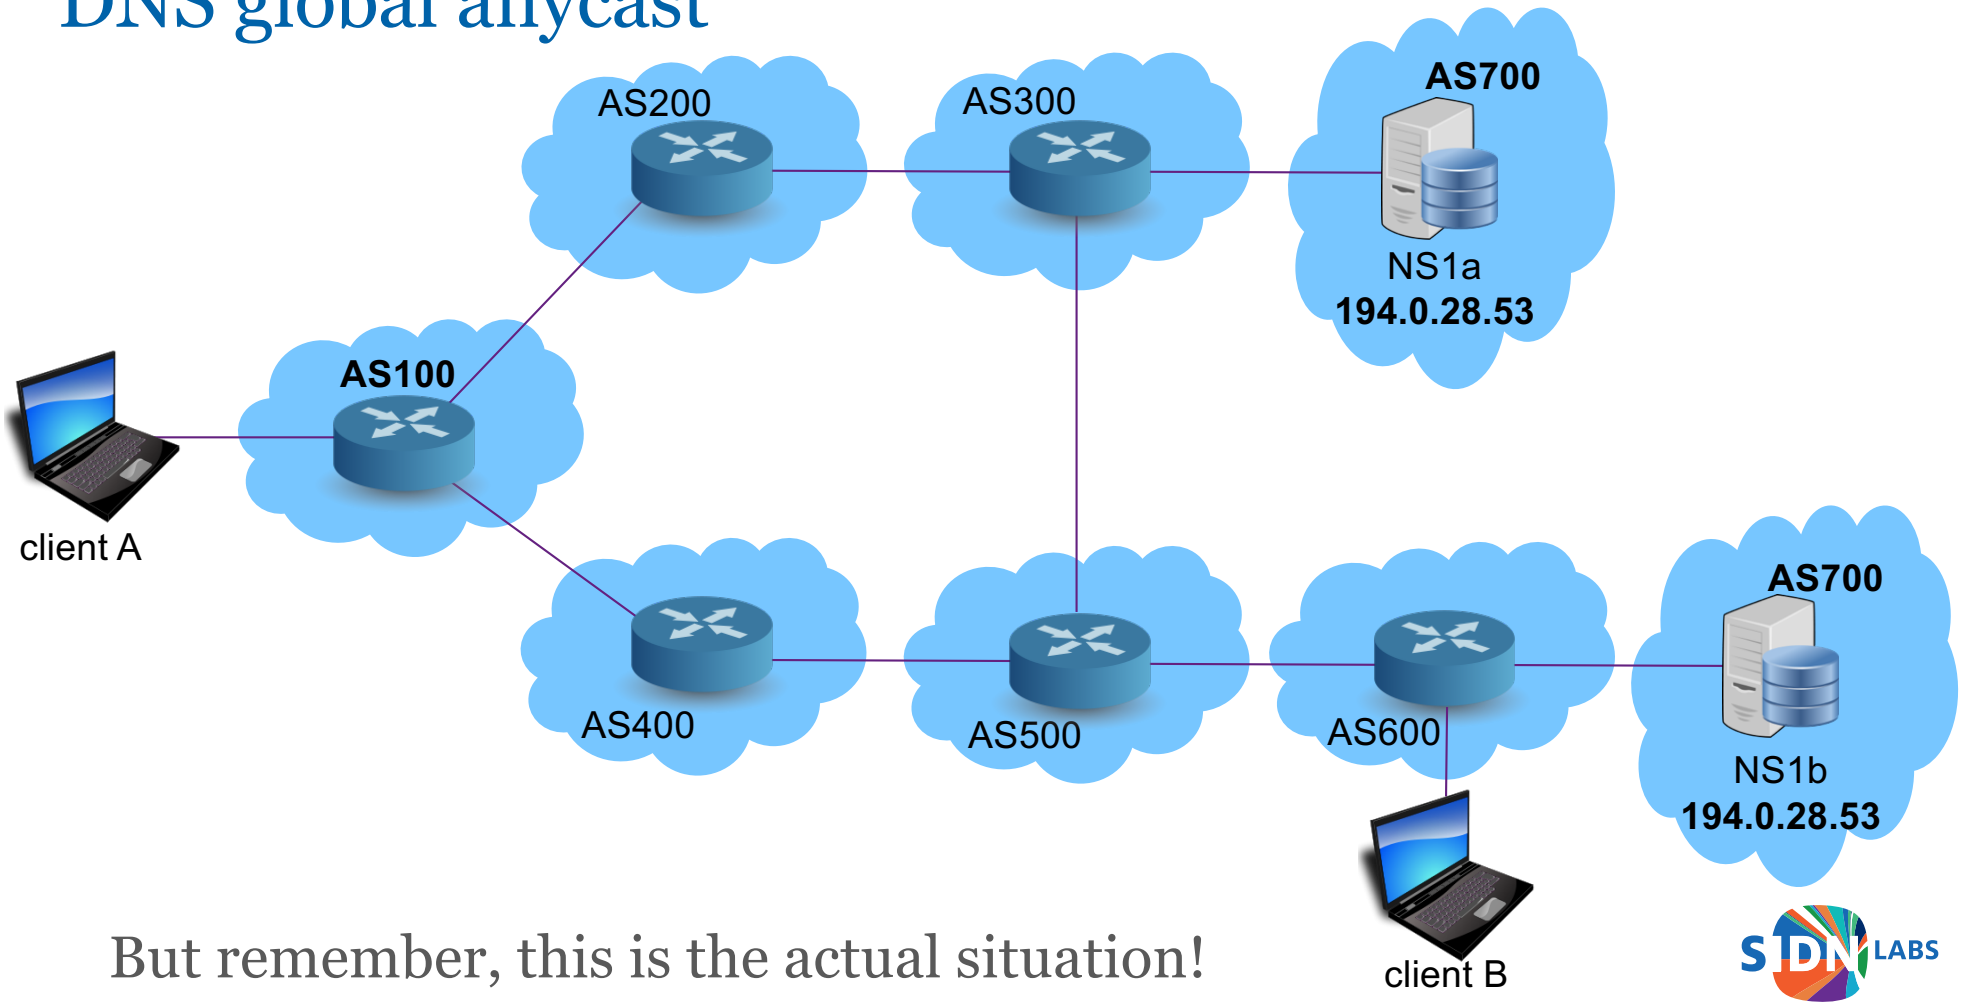

But remember, this is the actual situation!

DNS global anycast

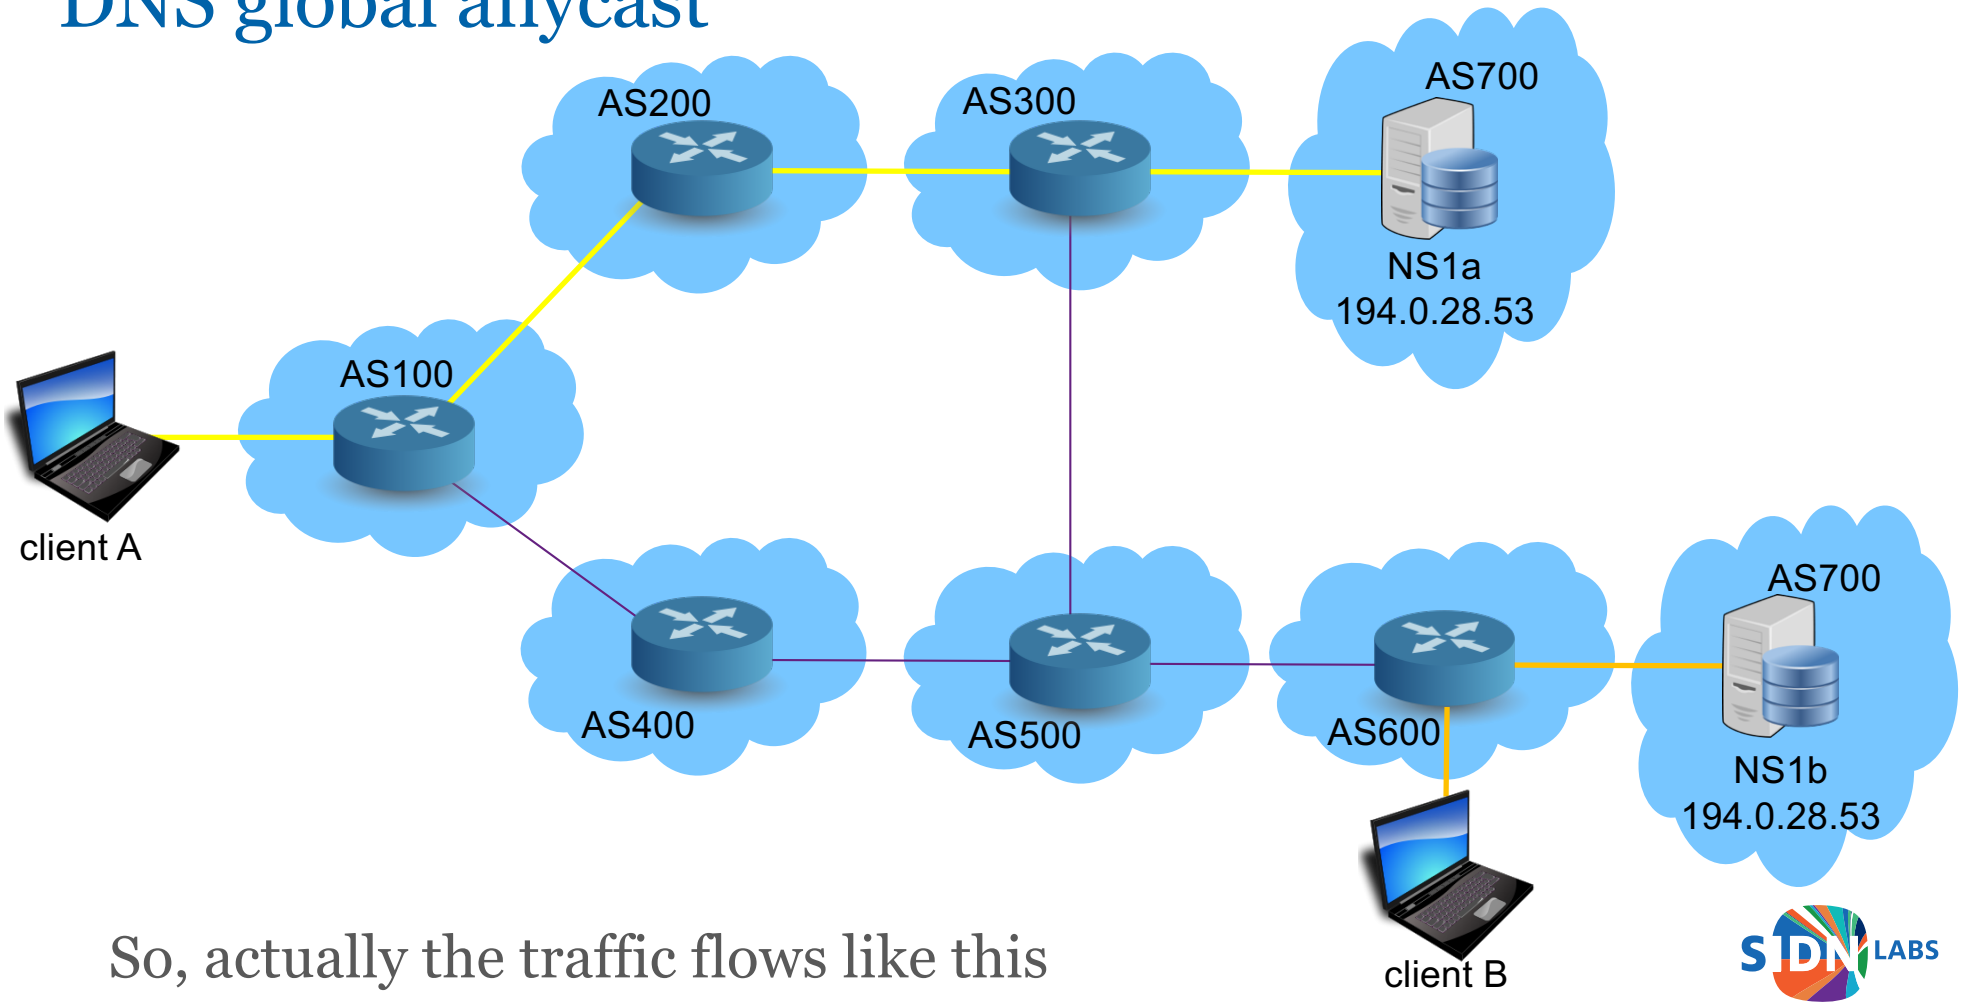

So, actually the traffic flows like this

DNS global anycast

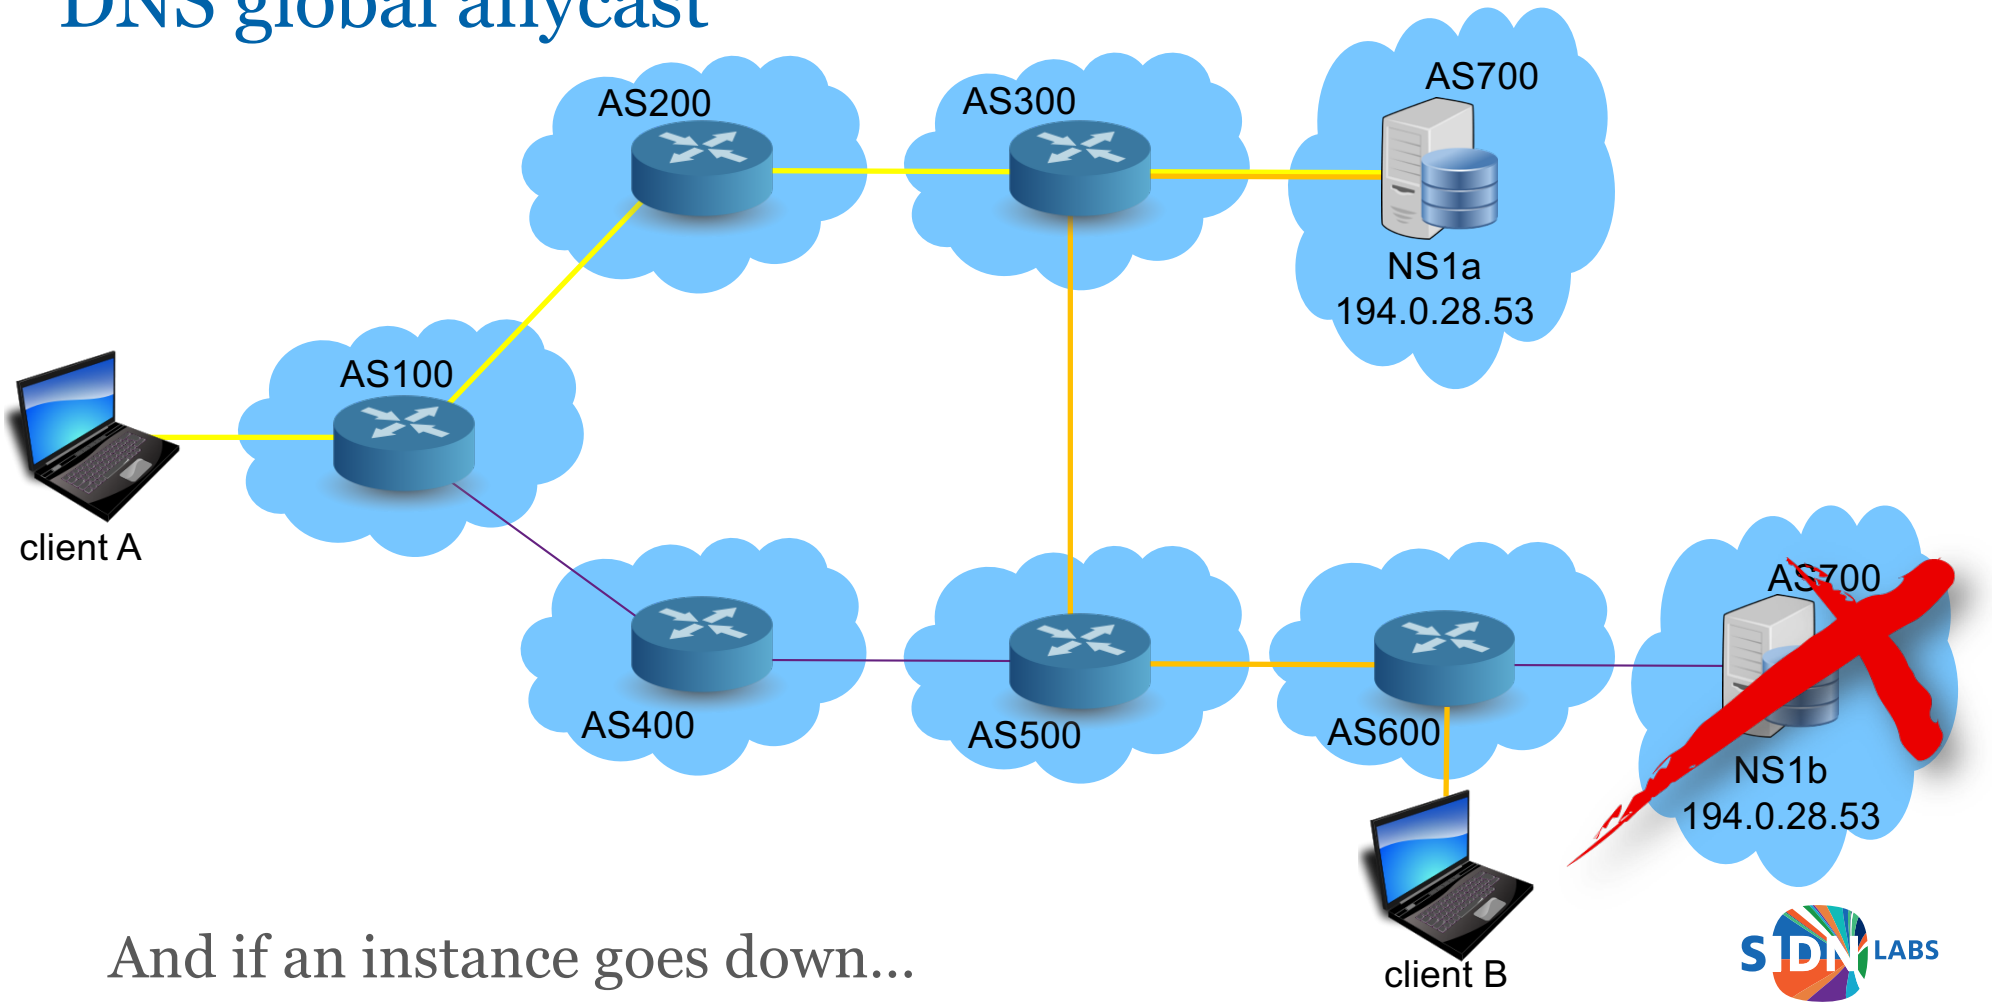

And if an instance goes down...

DNS global anycast (for .)

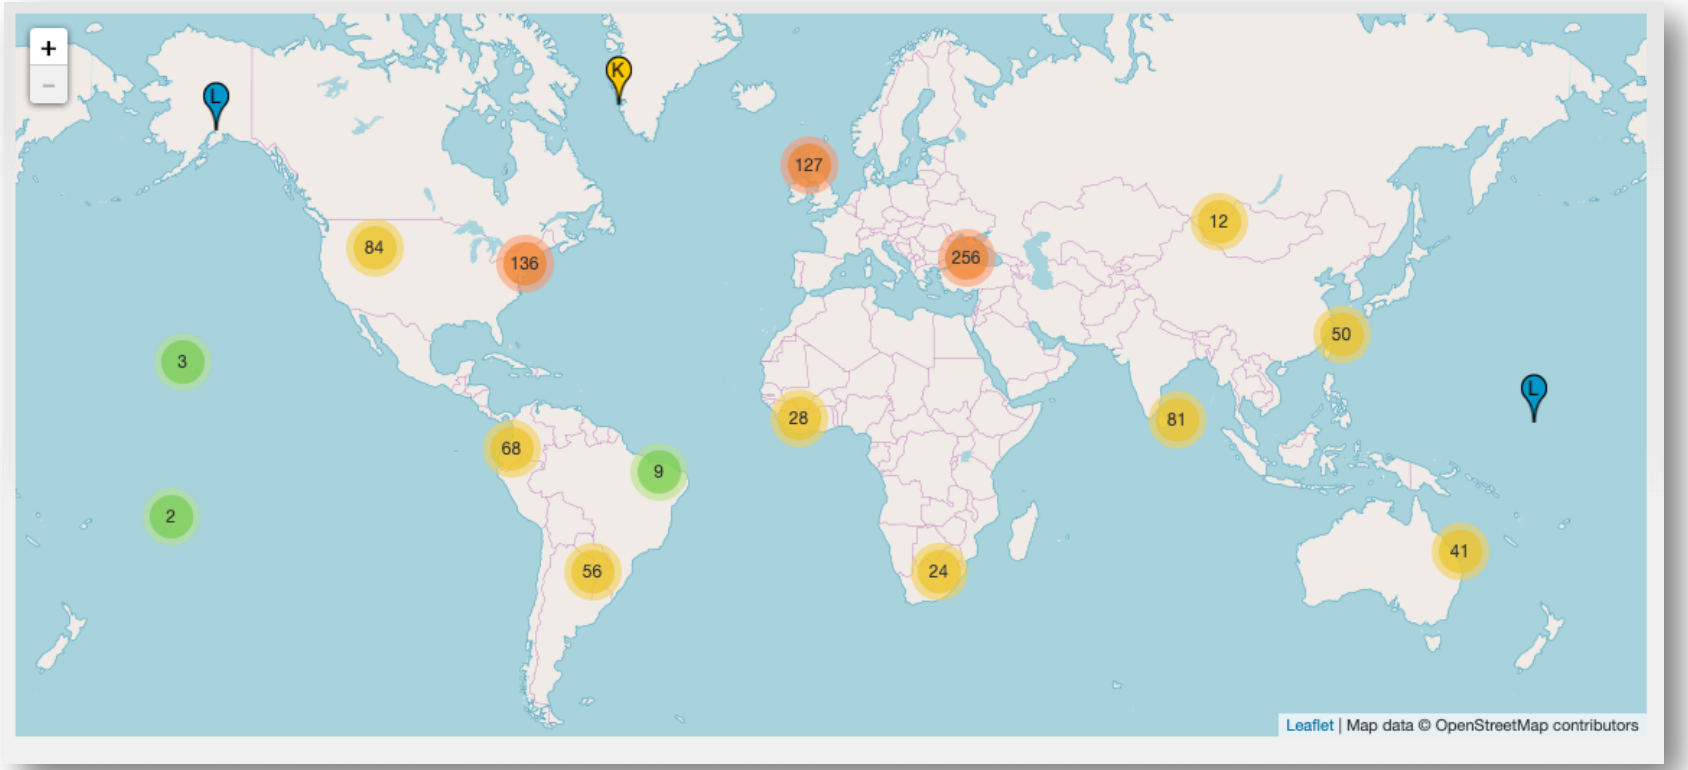

Guess how many?

DNS global anycast (for .)

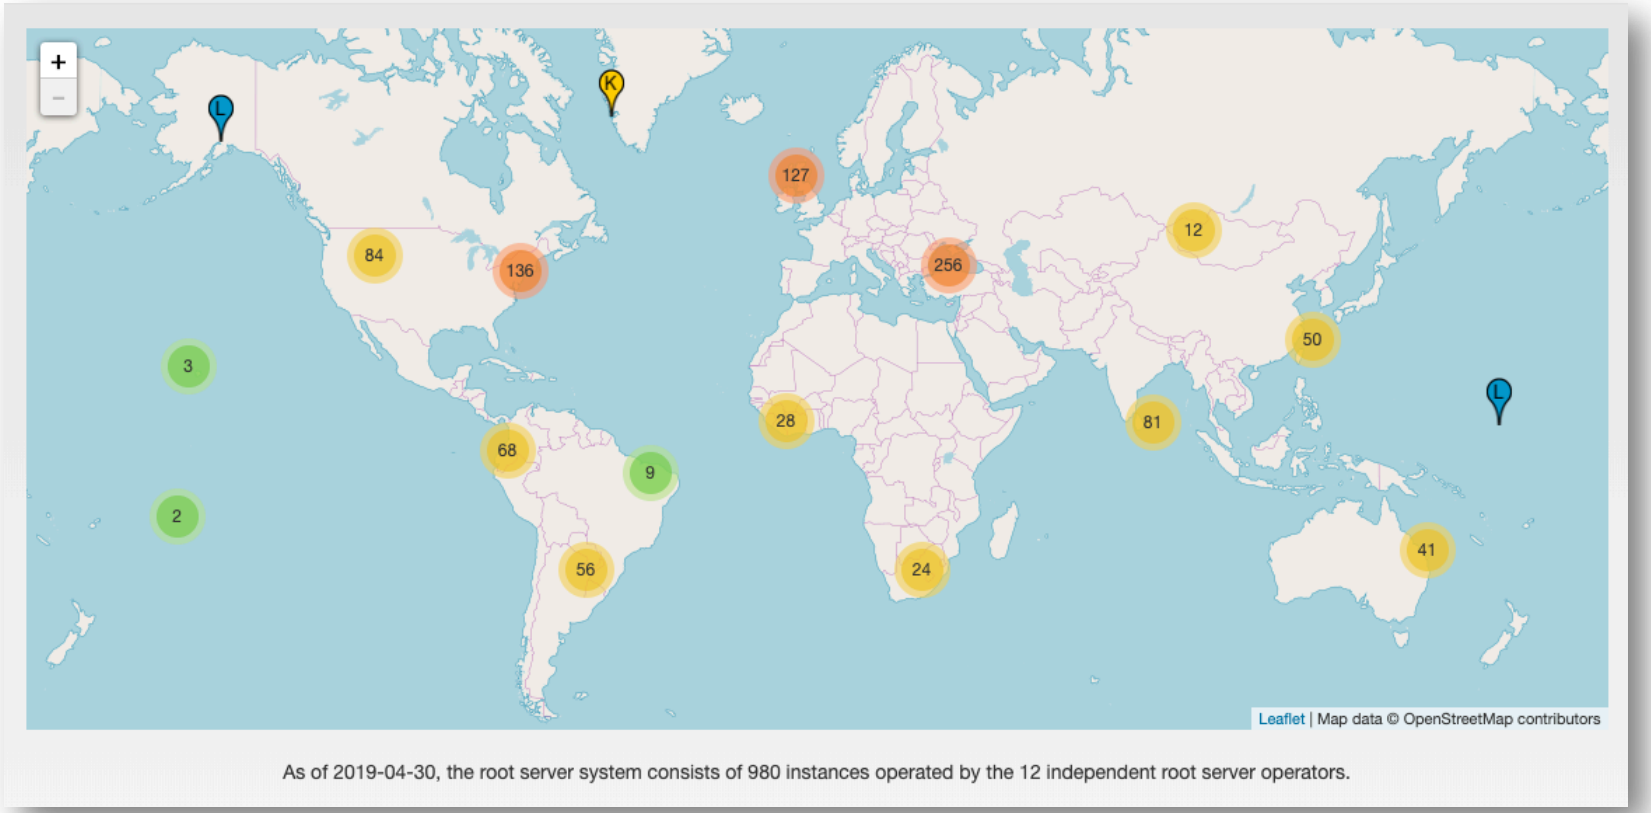

1088 servers!

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

SDN LABS

DNS global anycast (for .nl)

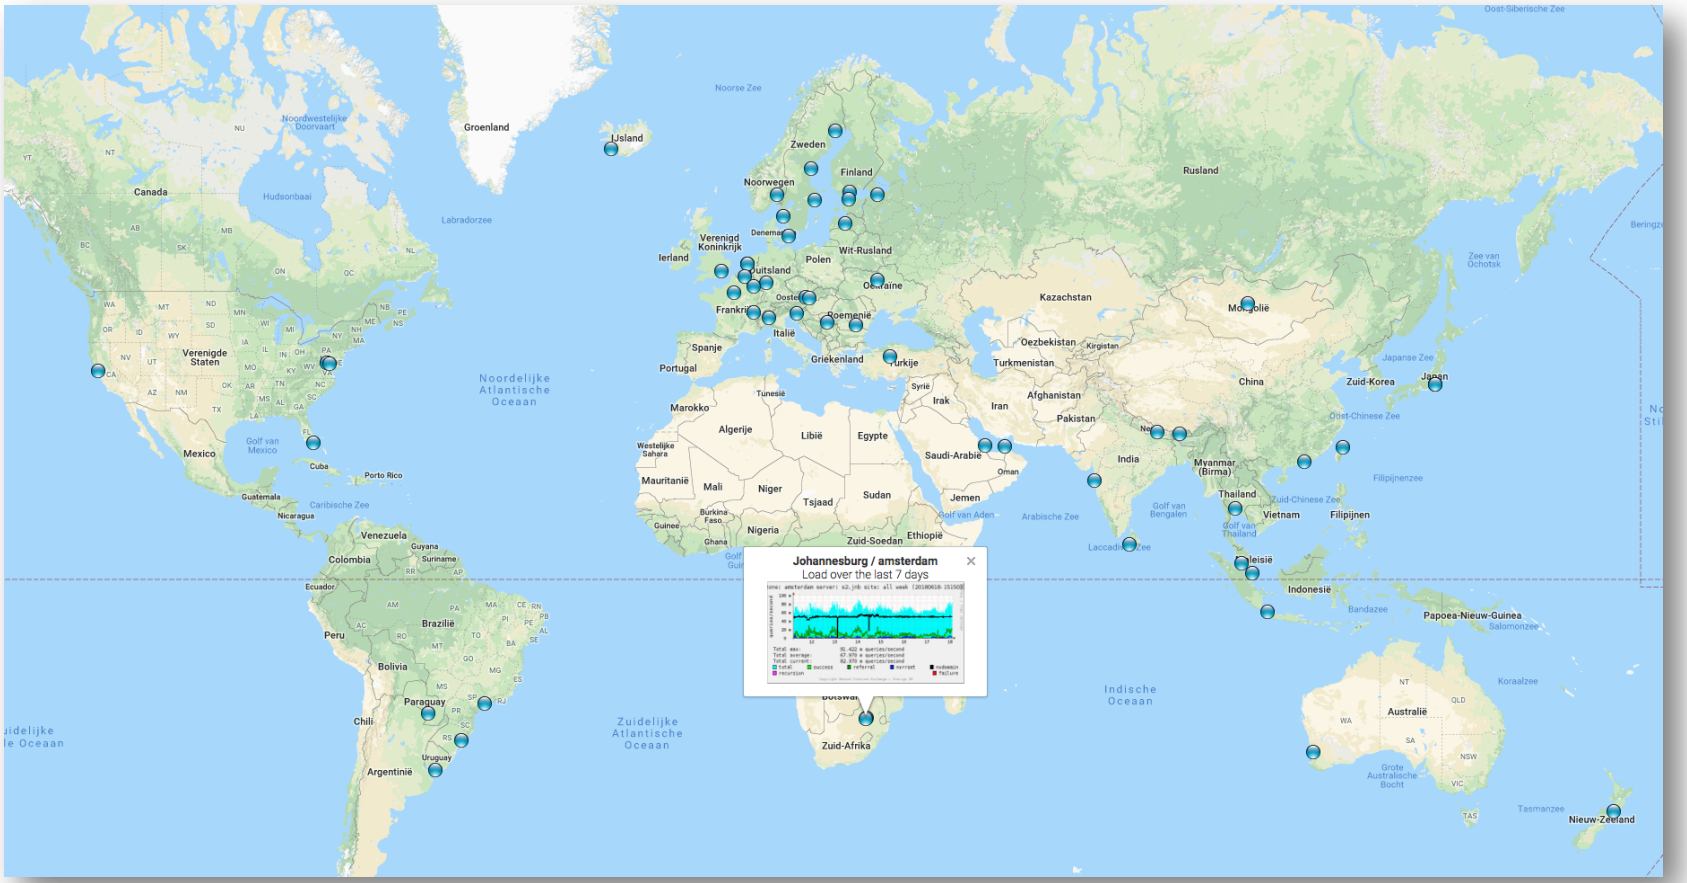

Netnod instances for .nl

Additional approach: DNS ***local*** anycast

- In essence the same principle as global anycast
- But with a deliberately restricted catchment.
- Dedicated instances for exclusive use by (big) ISP's
  - Focus on Netherlands
  - Must have reasonable abuse response capacities
  - Must comply to certain requirements (like BCP38 and IPv6)
- Nothing more, nothing less (basically)

Goals	Non Goals
Resilience (win the rat race)	Latency (in contrast to global anycast)
Availability (at least for our most important users)	Bandwidth (DNS doesn't consume that much, yet)

SIDN LABS

DNS local anycast

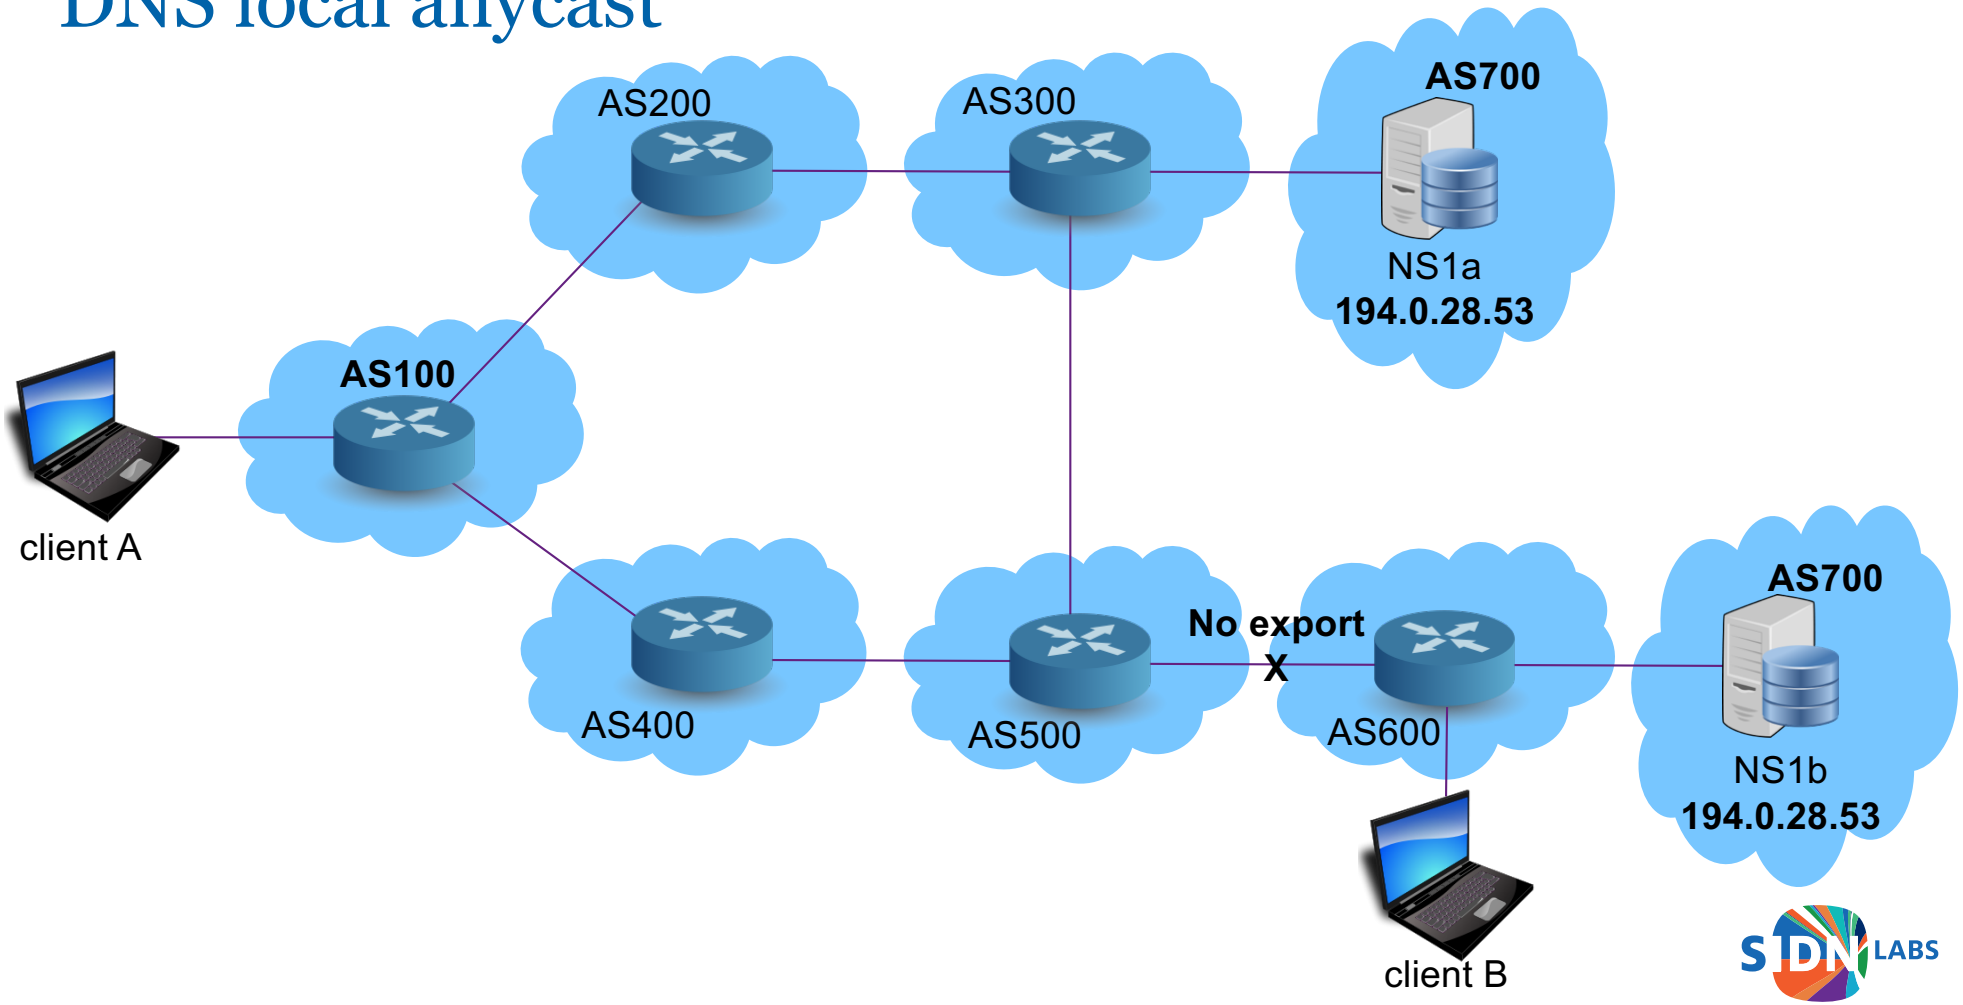

DNS local anycast

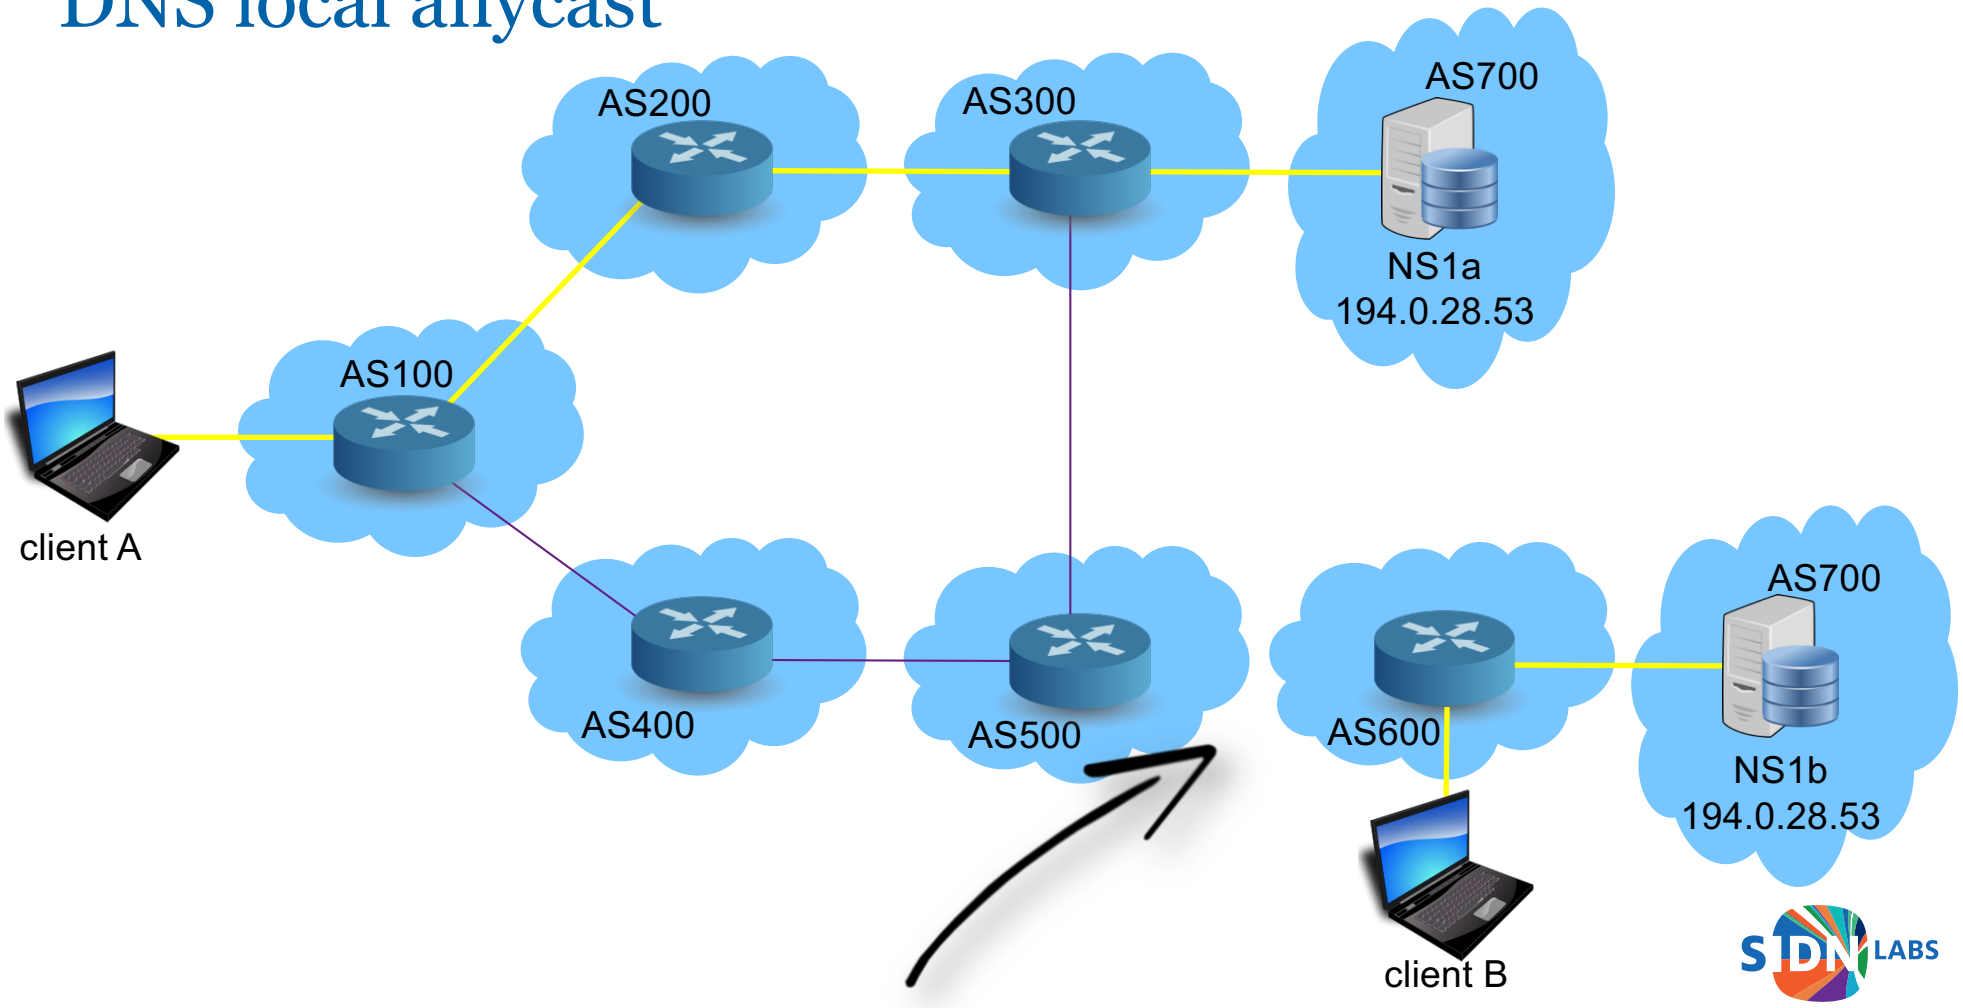

DNS local anycast

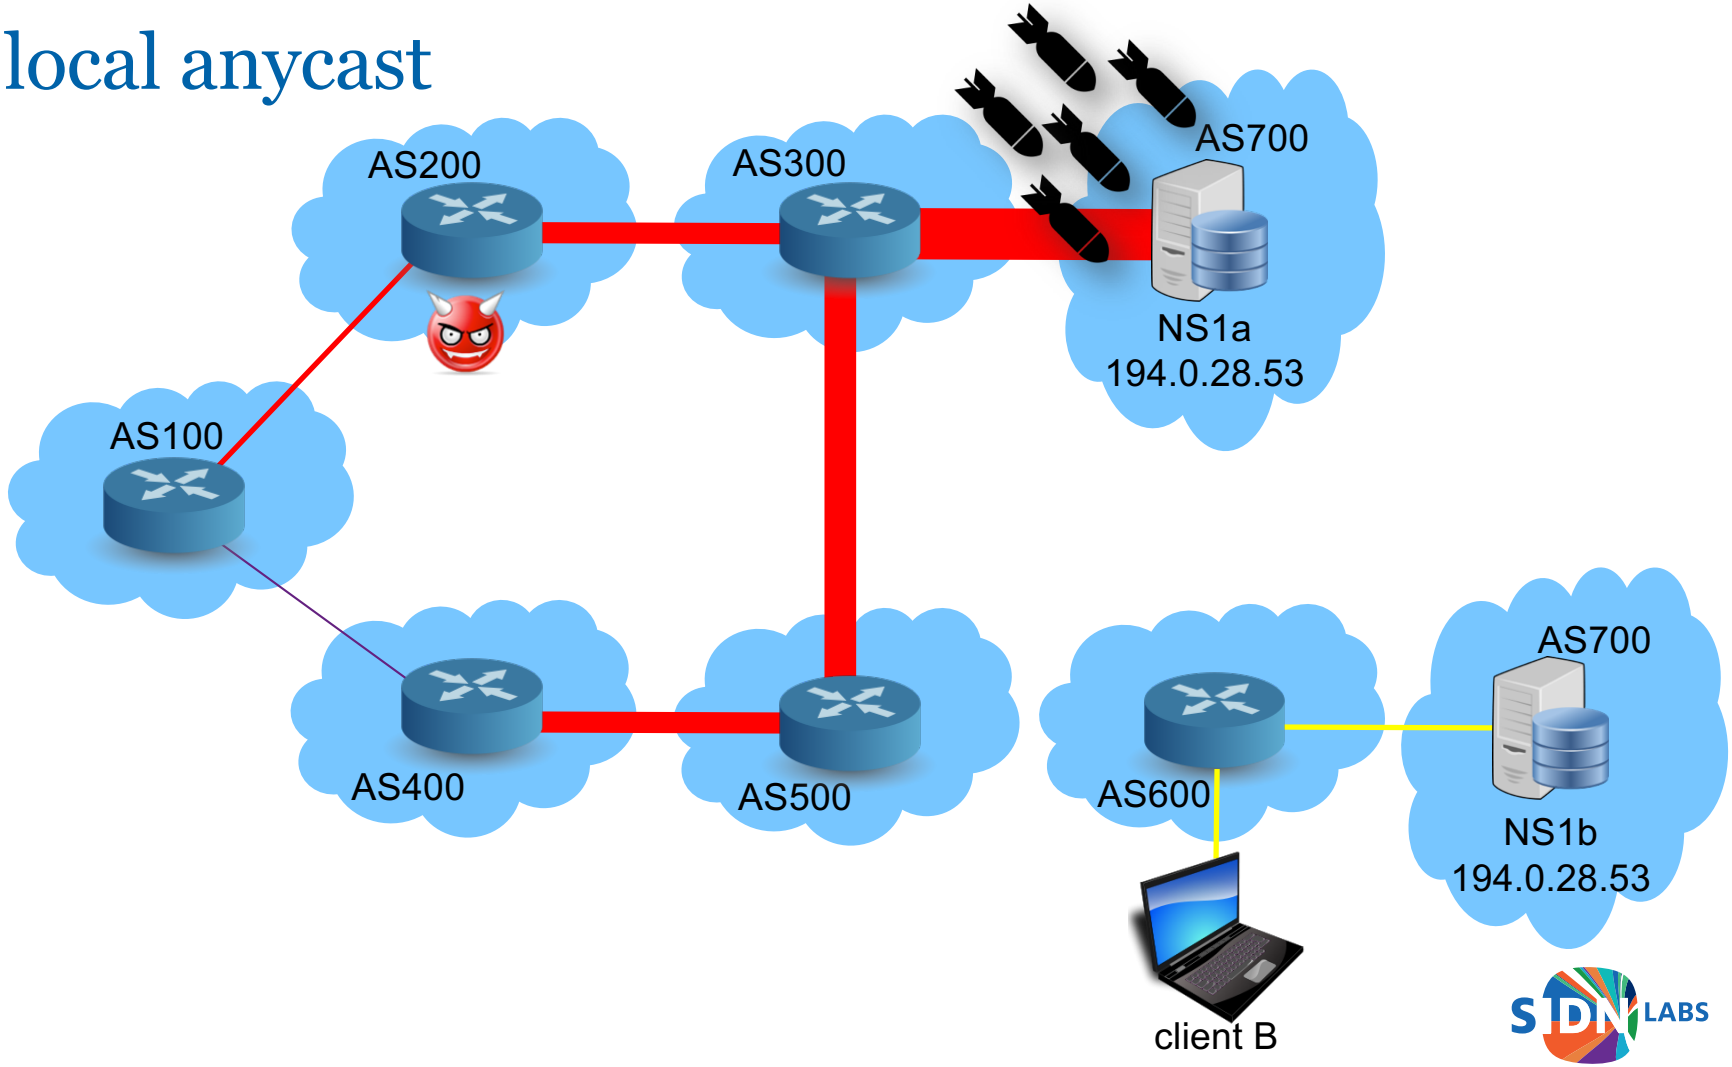

DNS local anycast

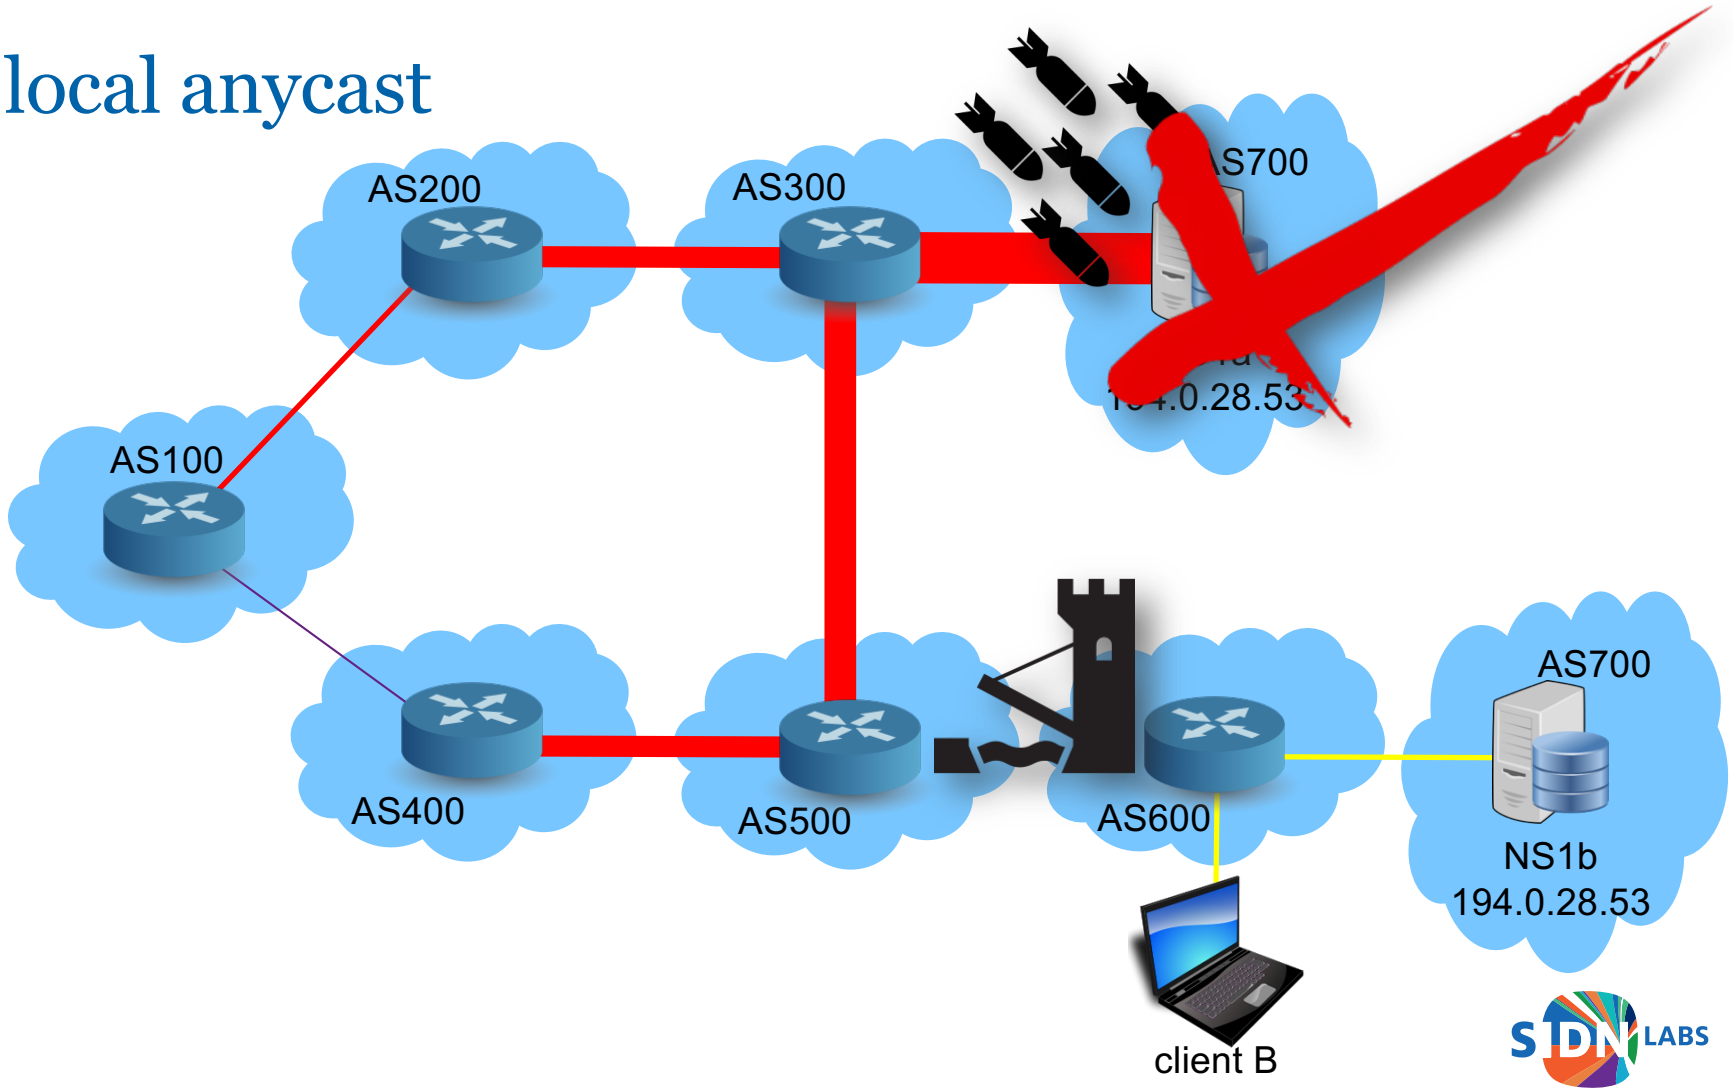

DNS local anycast – current situation

- Local presence at 8 sites at ISP's
- One shared node (will explain later)
- ~ >80% of Dutch consumers "covered"
- Try it: ns1.dns.nl

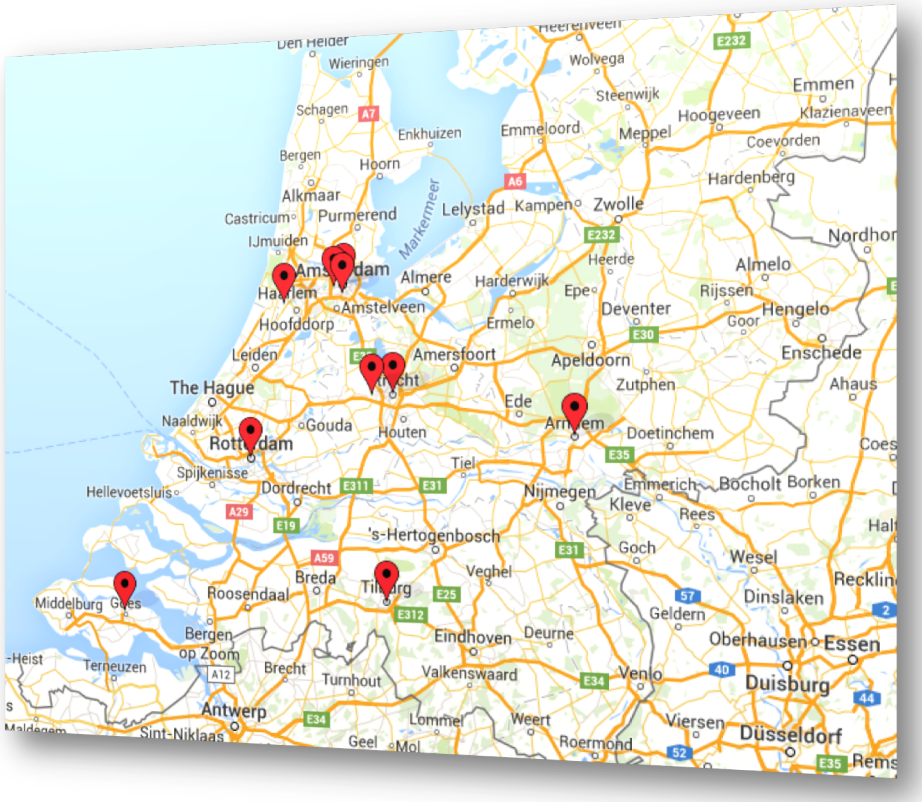

Concluding

- We learned (a bit) about names, numbers, routes
  - and about IETF, ICANN, SIDN
- Running the core of the internet is not a trivial task
  - many people, quite a lot of organizations and stake holders
- Many challenges have been overcome, a lot more to go
  - abuse, politics, legislation, dependency
  - resilience (anycast)
  - addressing, scaling (keep IoT in mind)
- The internet needs constant maintenance and innovation
- Together we can make that happen 😁

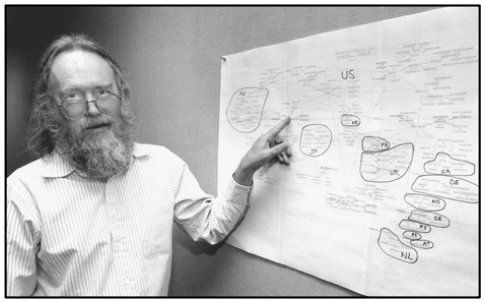

"The Internet works,
because a lot of people **cooperate**
to do things together"

– Jon Postel
(1943-1998)

Questions, discussion?

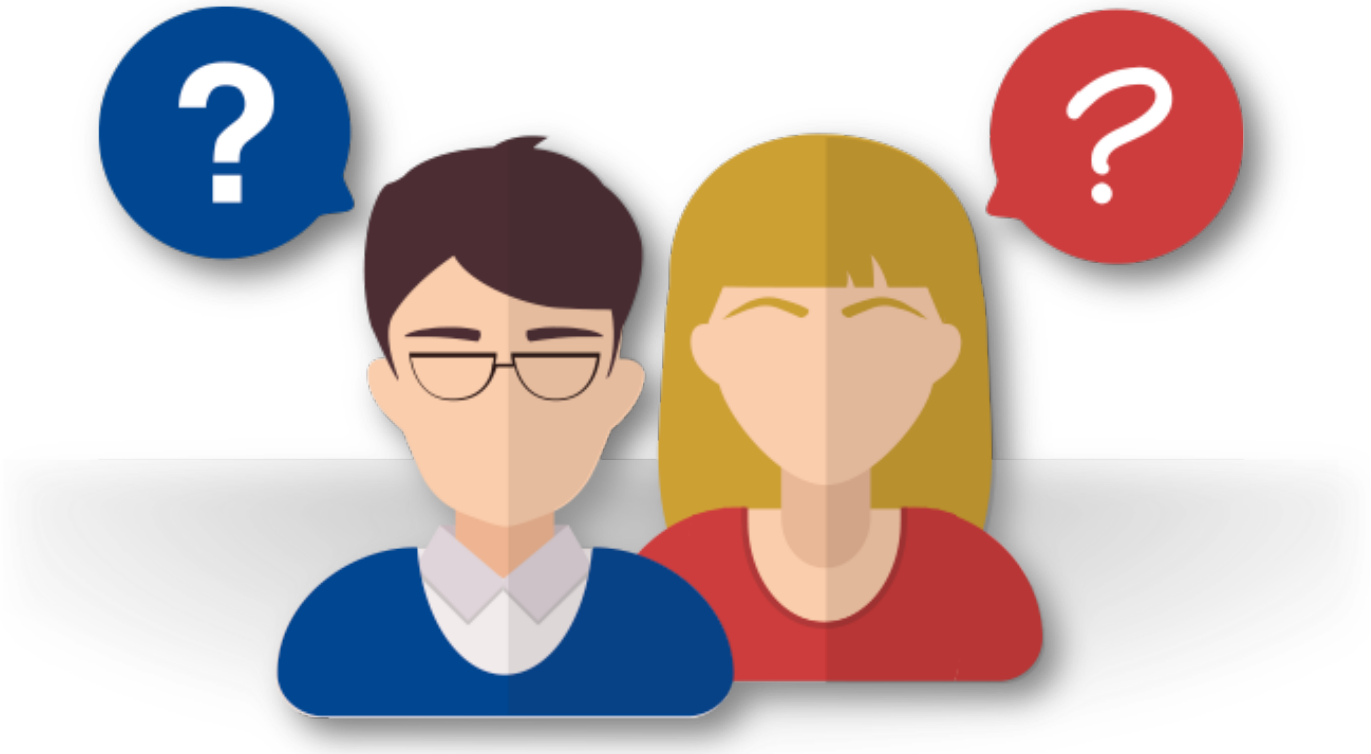

Thank You!