Your world. Our domain.

Local Anycast at SIDN

Marco Davids

Joint 38th CENTR Technical / 12th CENTR R&D workshop
Moscow (RU) - May 30rd 2018

SIDN

Registry for .nl ccTLD

- And a number of other things (.aw, .politie, .amsterdam)

https://www.sidn.nl/

.nl domain names

DNSSEC .nl domain names

(as per may 9th 2018)

SIDN Labs






Vacancies:

- 1 x Machine Learning Engineer
- 2 x Research Engineers on Emerging Internet Architectures

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

Botnets / DDoS

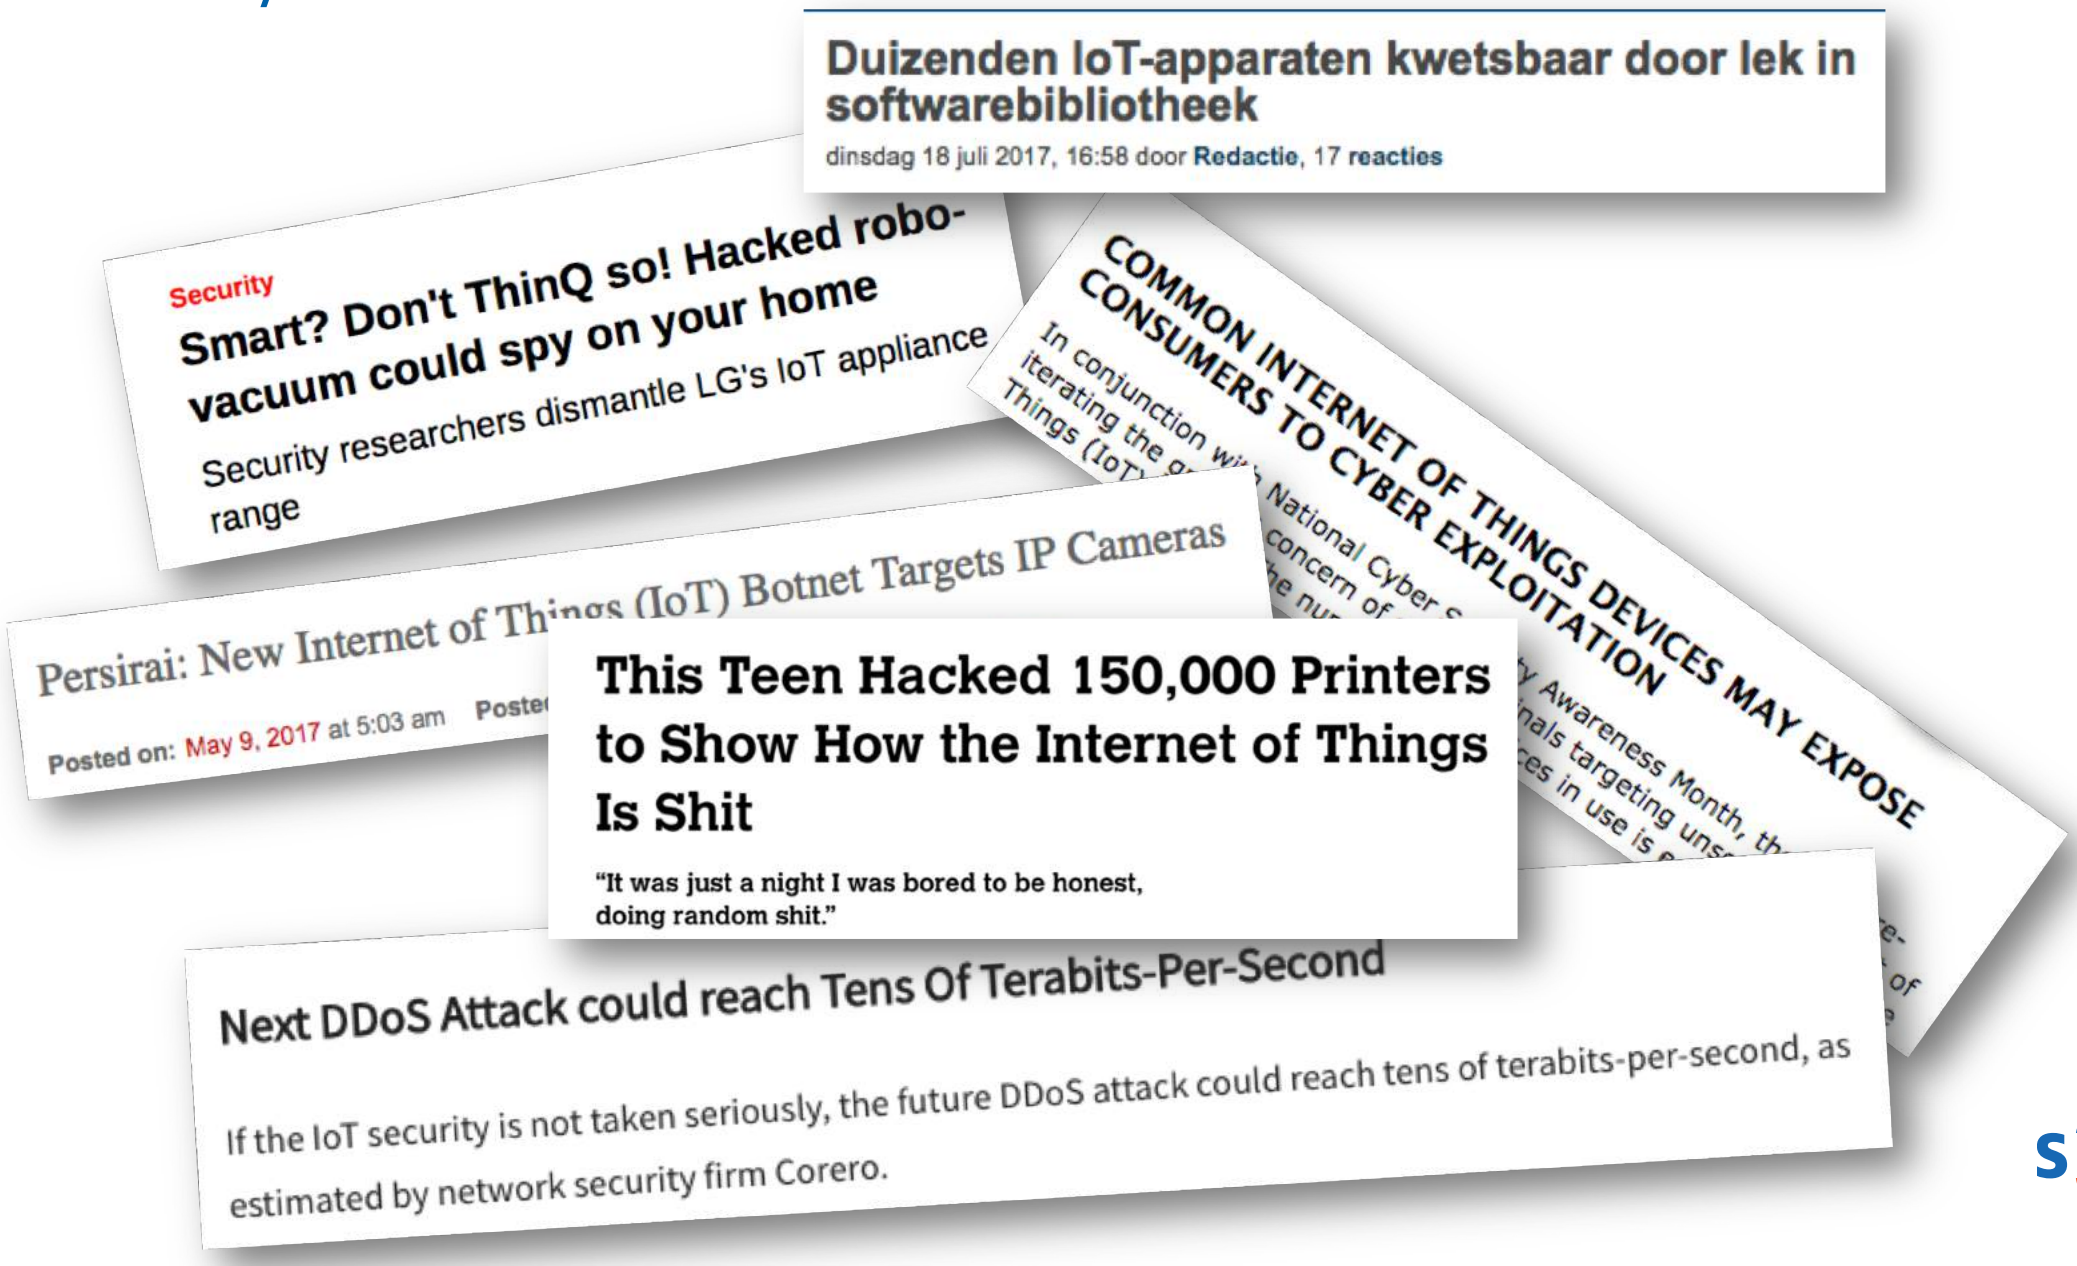

Why we care

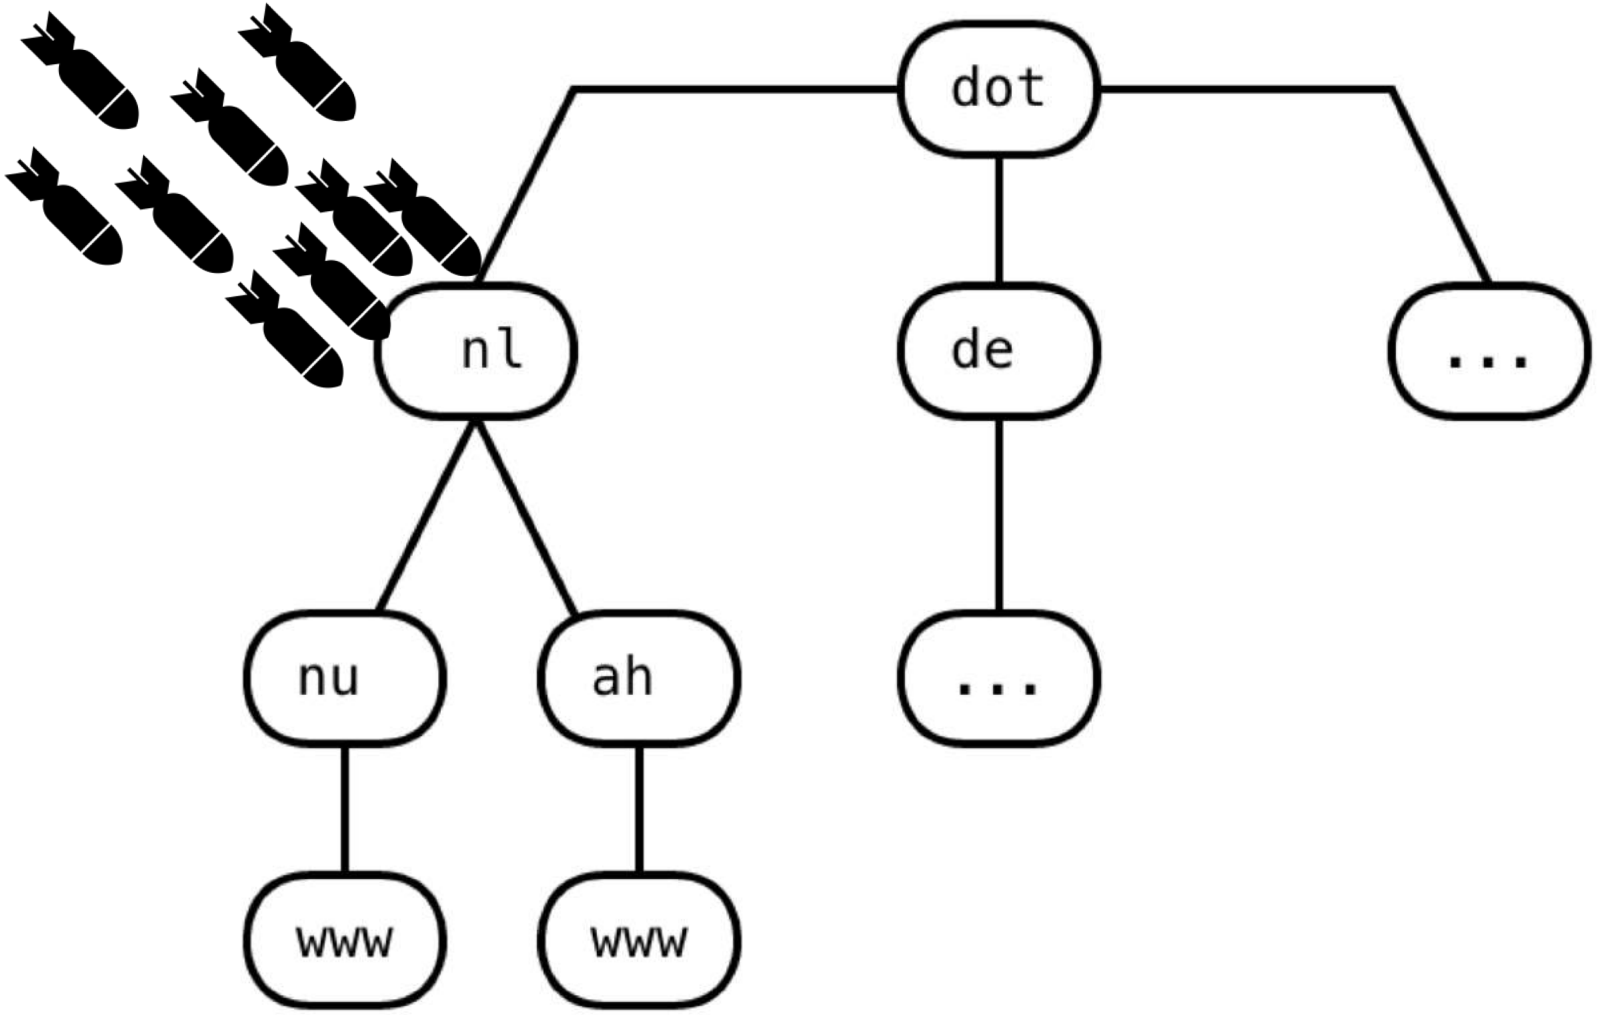

Why we care

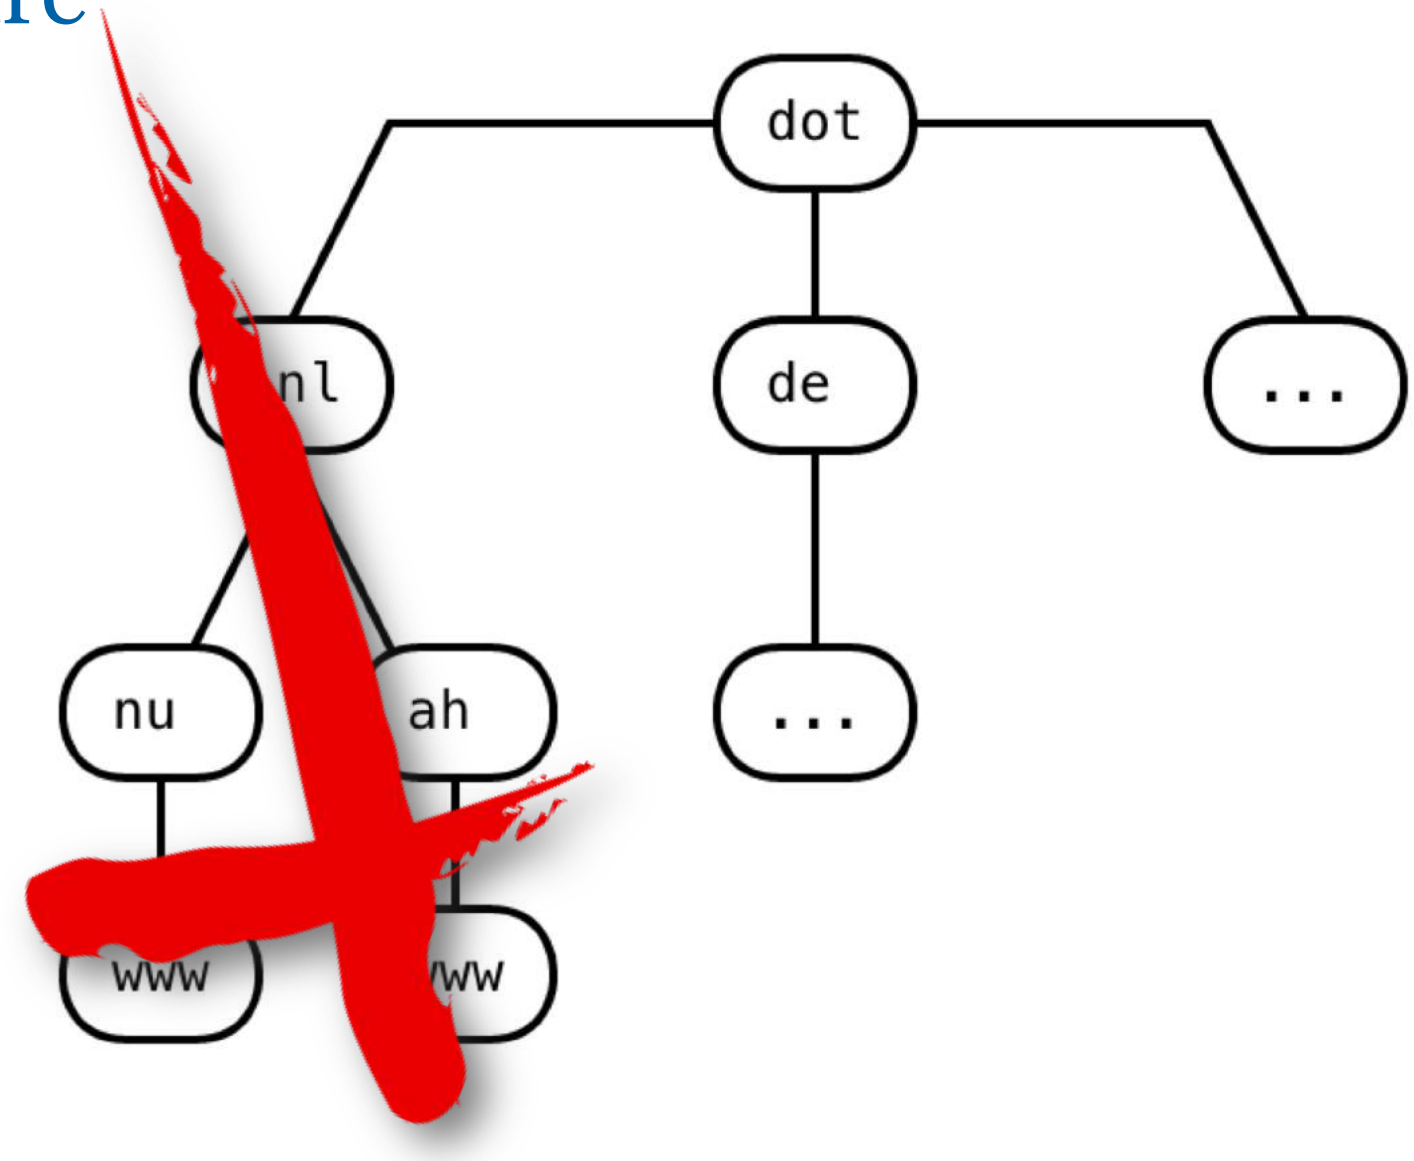

The Solution: DNS global anycast

- 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
  - 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)

SIDN LABS

DNS global anycast (in a nutshell)

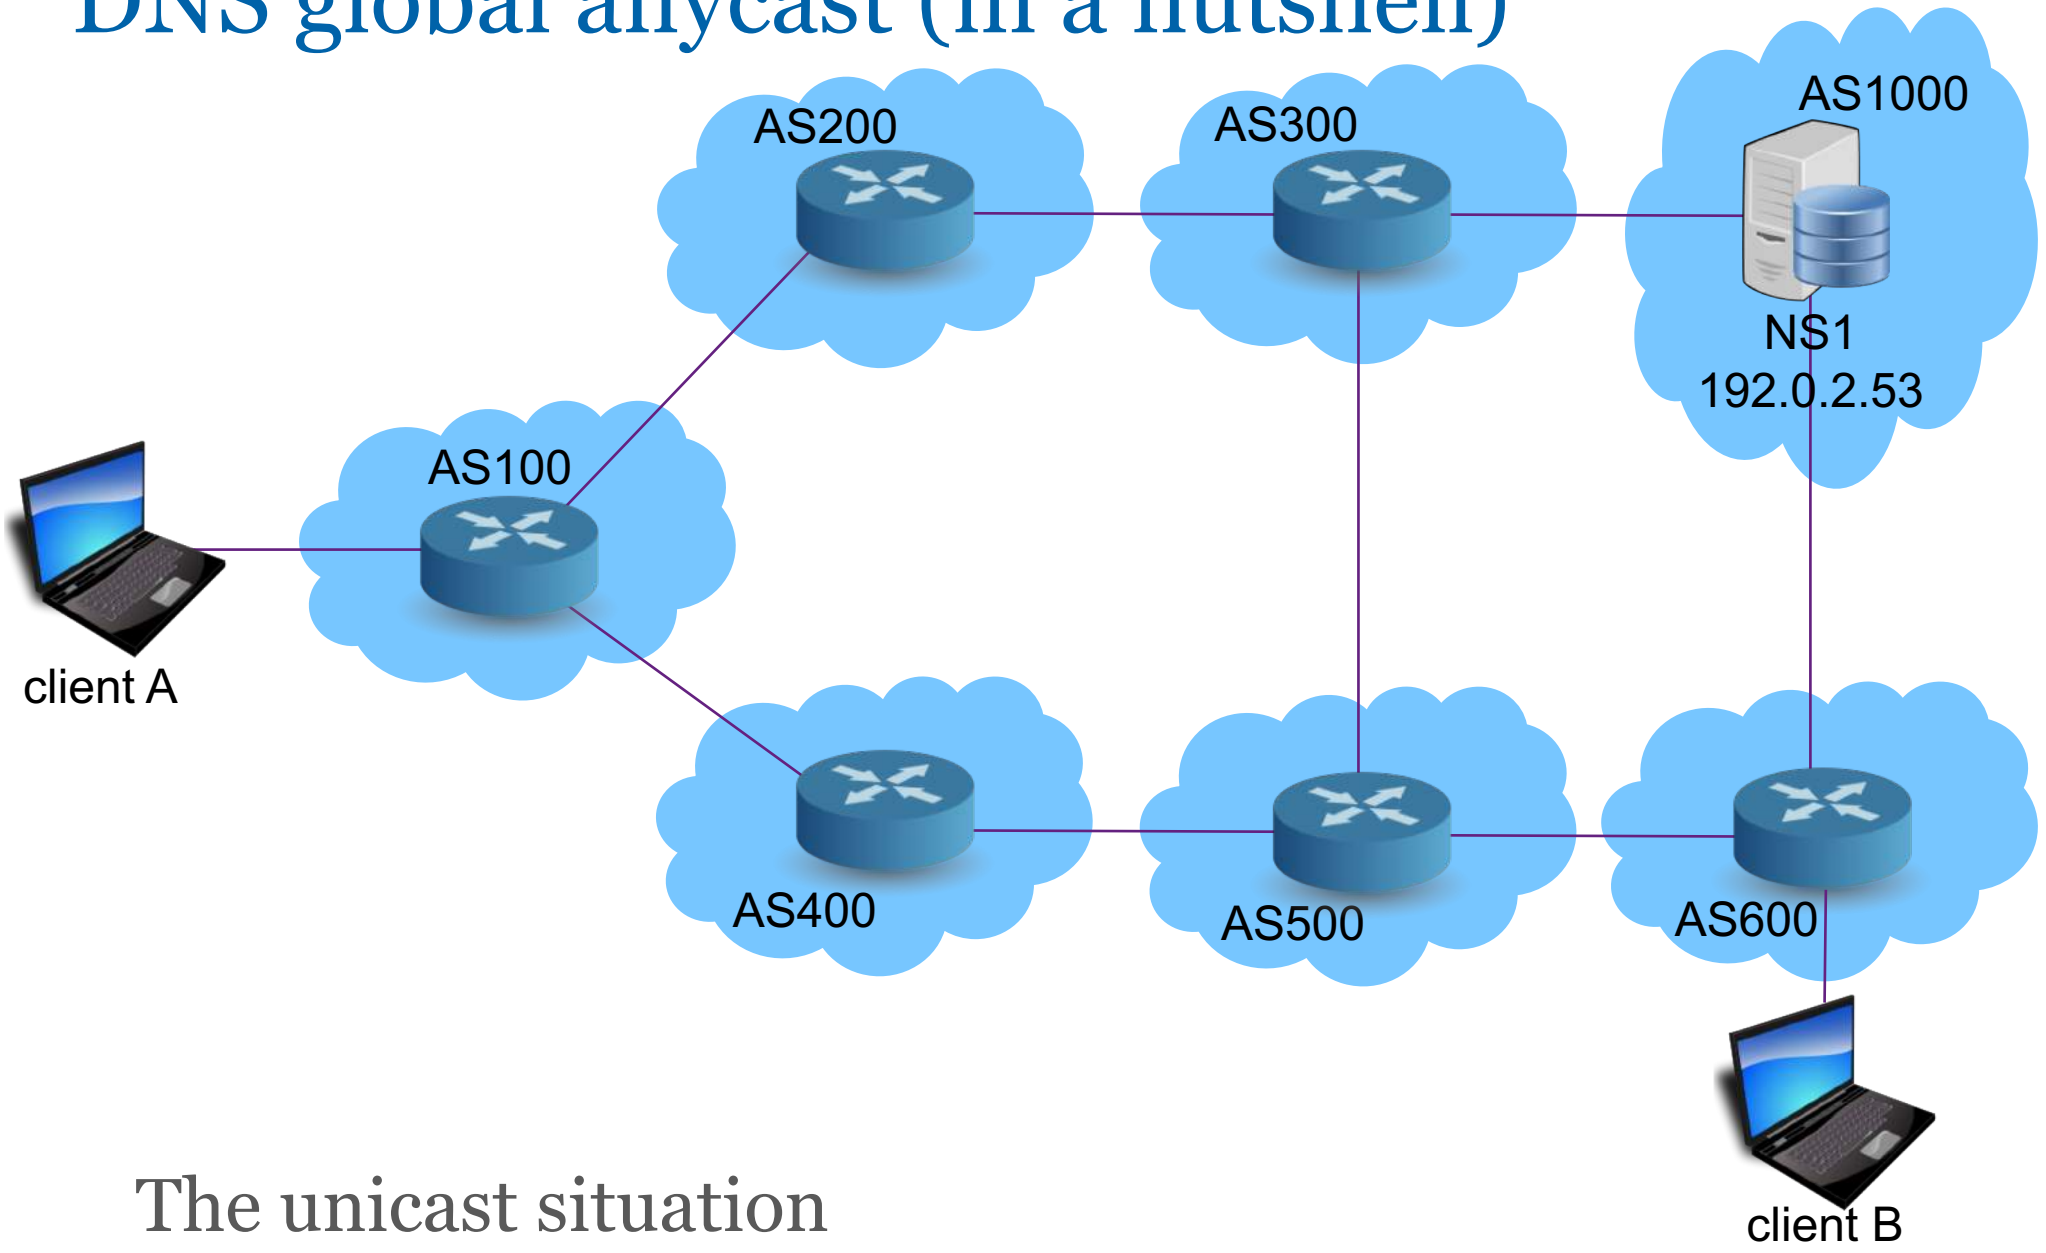

The unicast situation

(simplified, so one server, only IPv4, etc.)

DNS global anycast

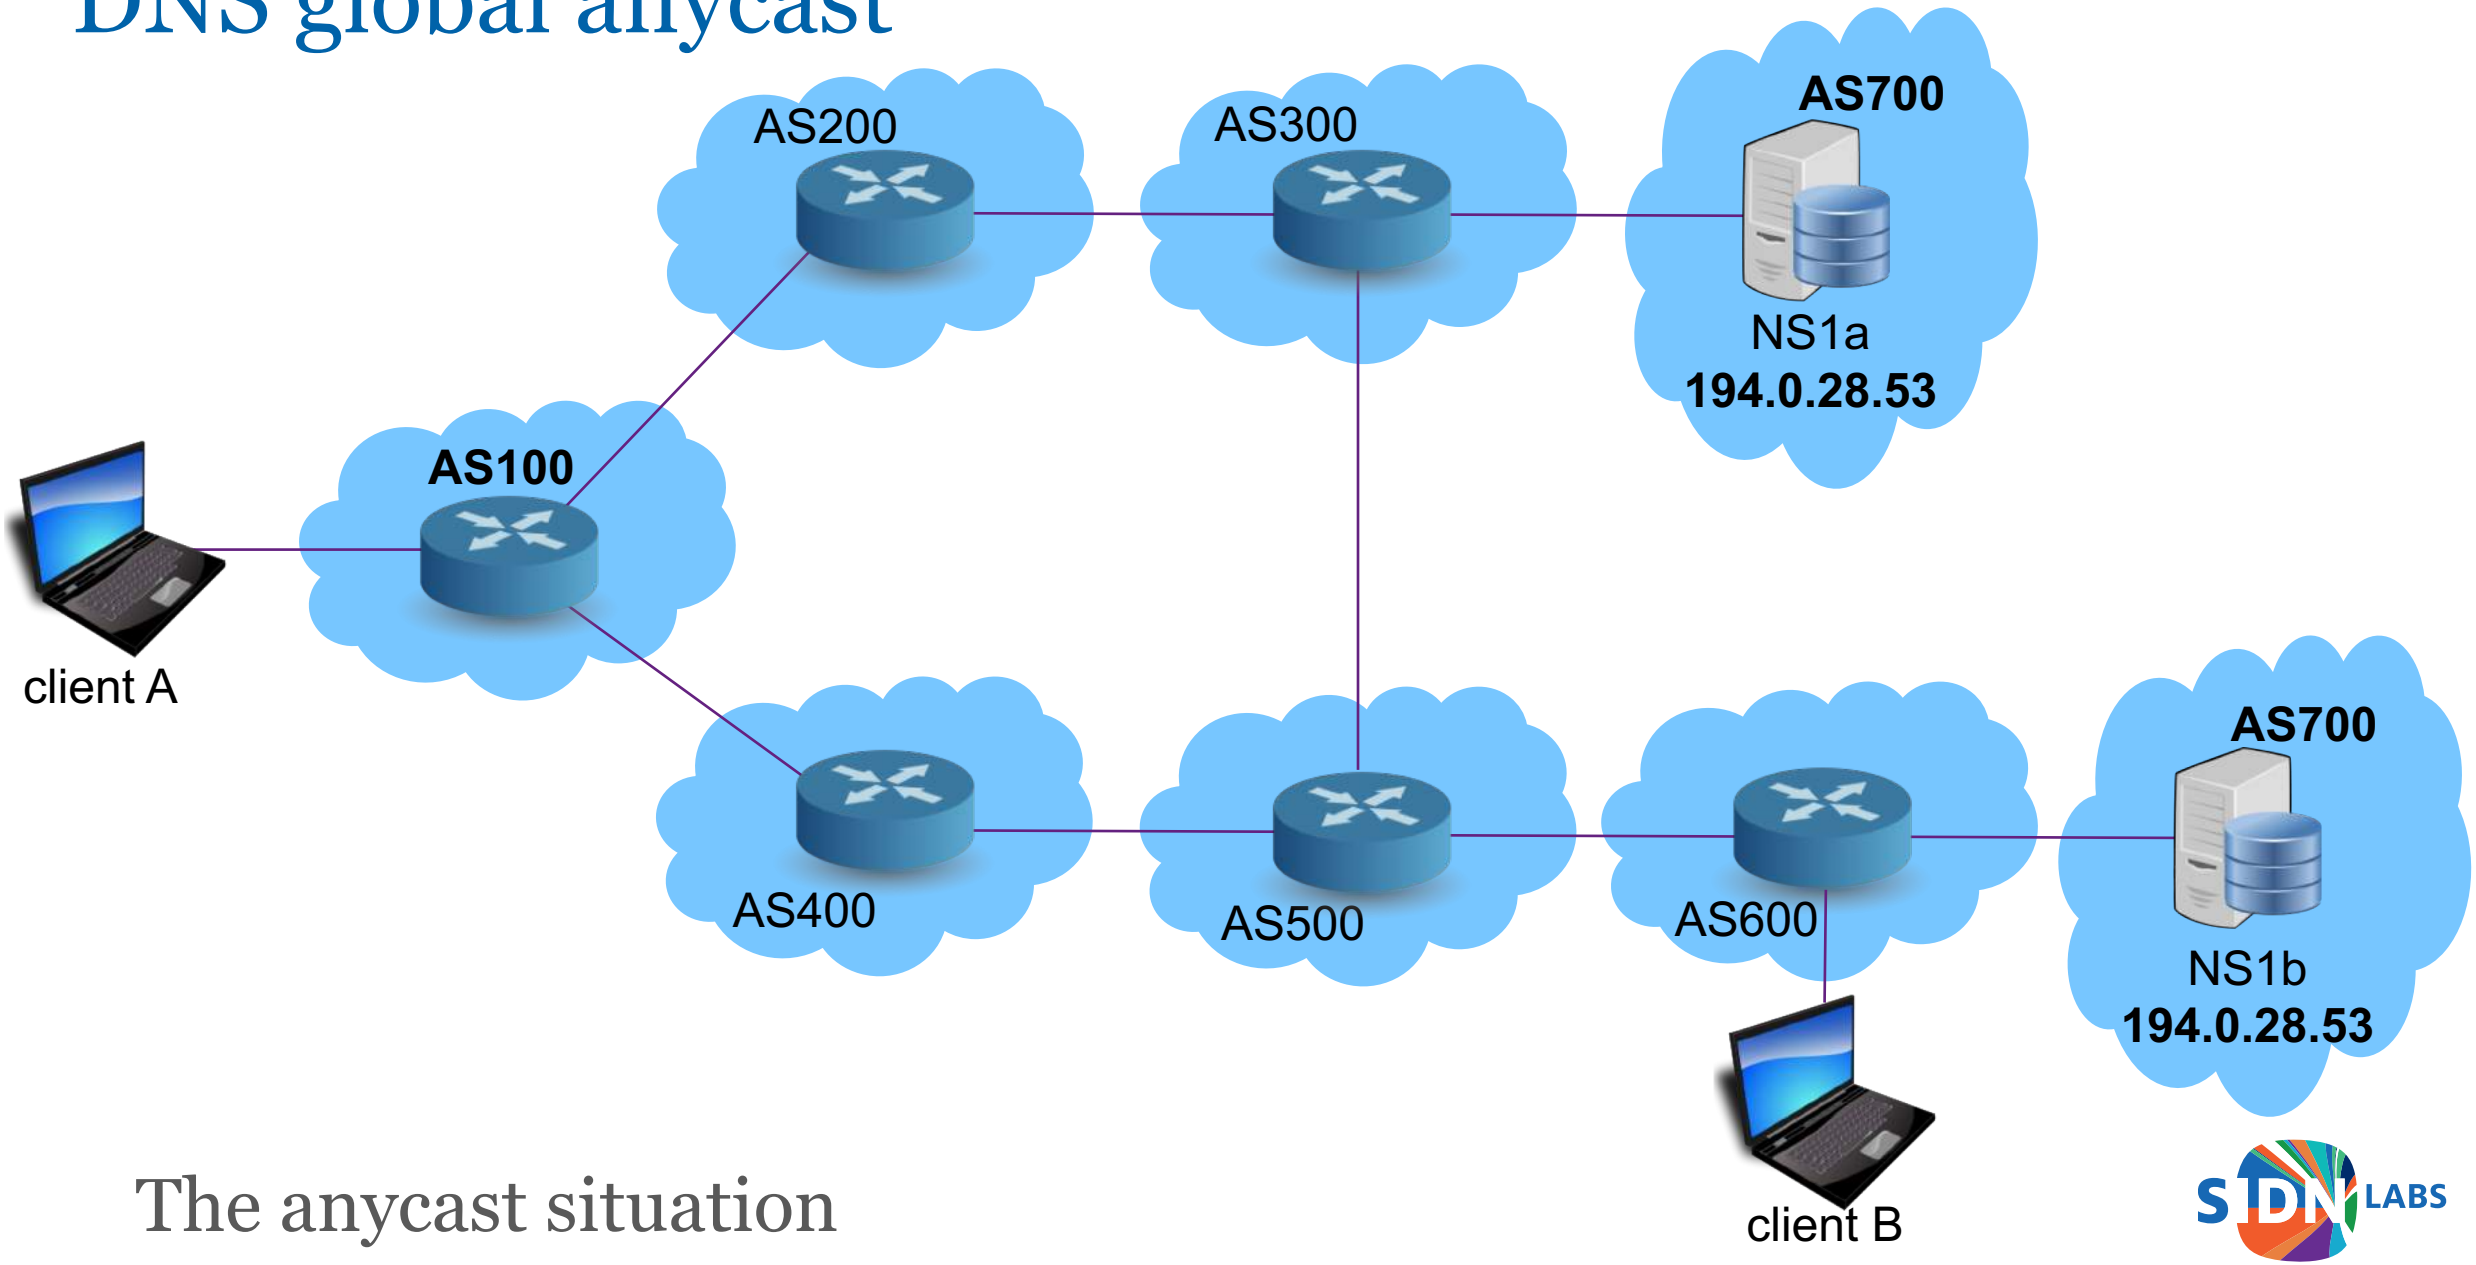

The anycast situation

DNS global anycast

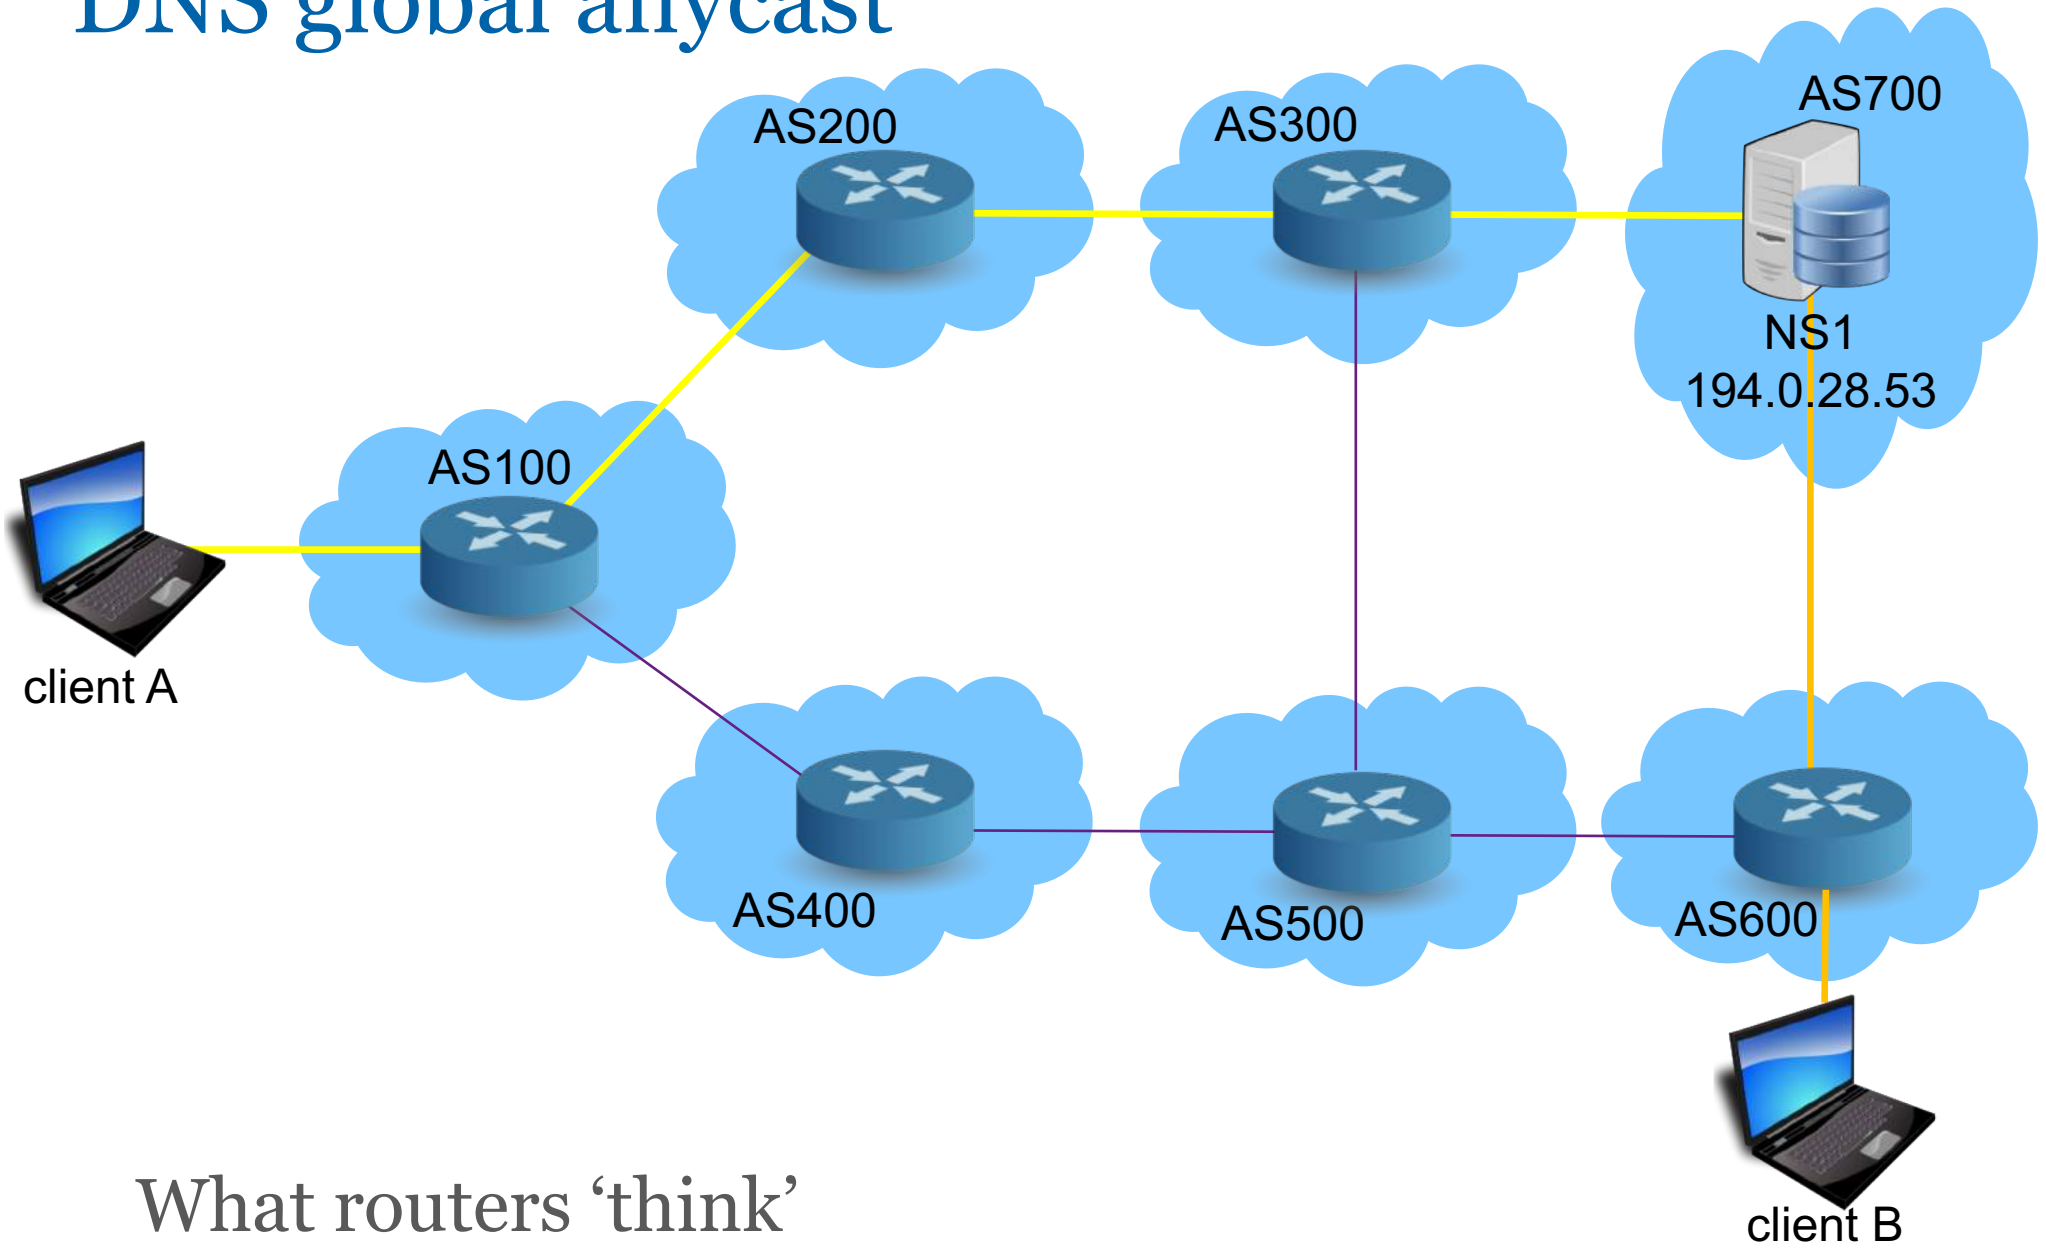

What routers ‘think’

DNS global anycast

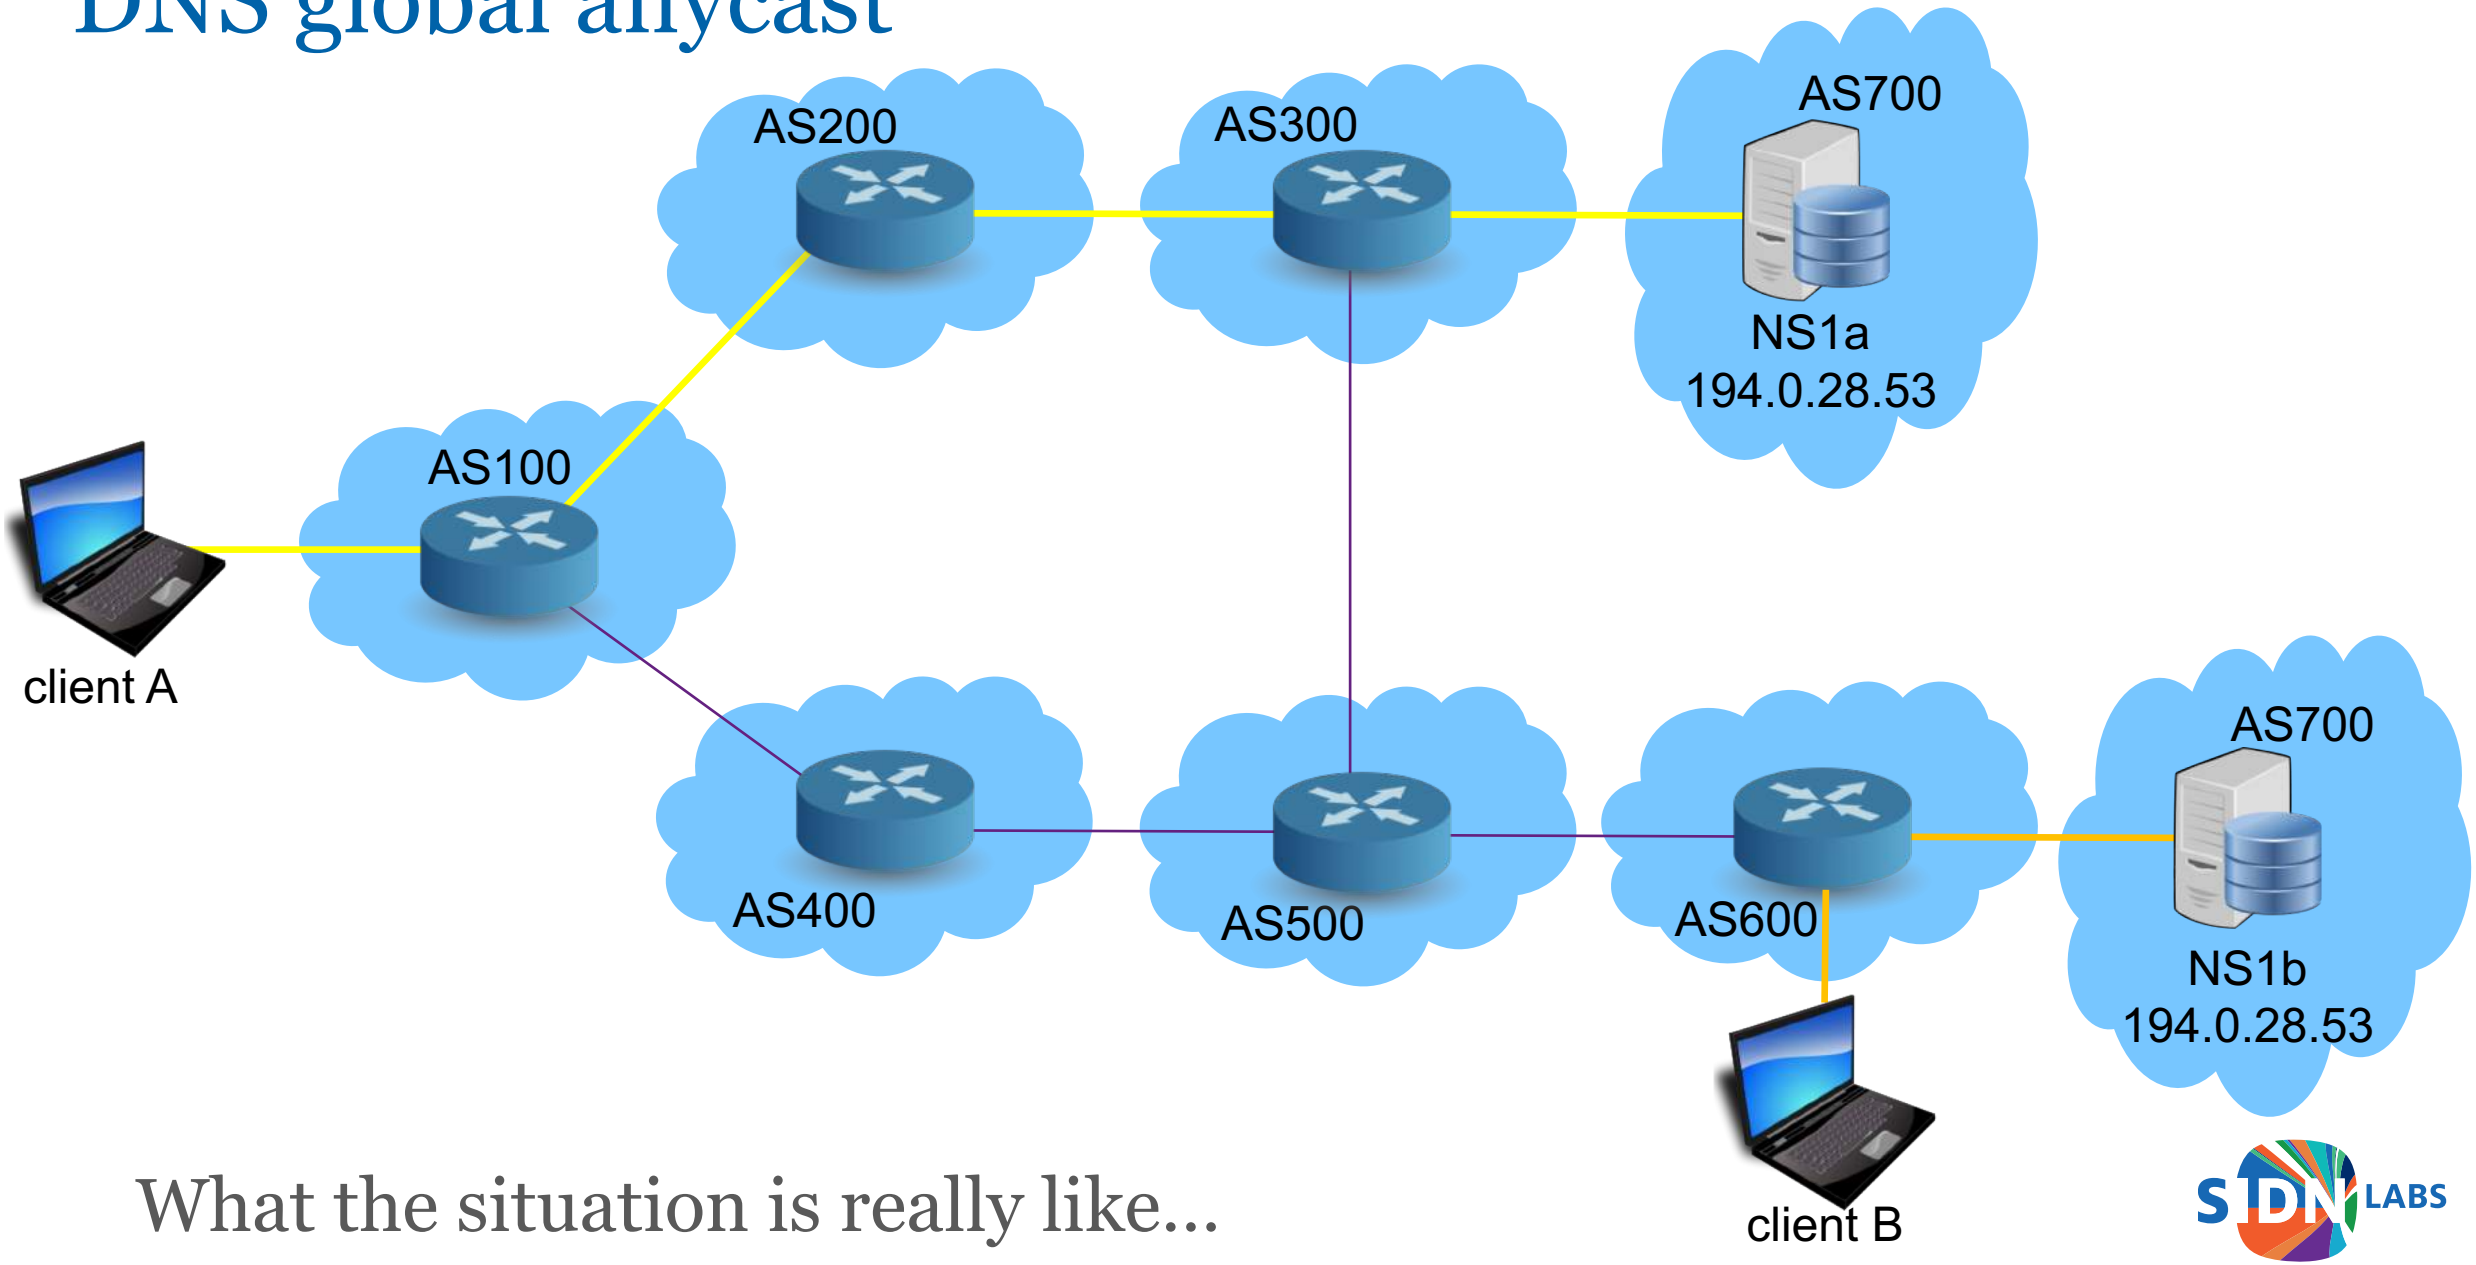

What the situation is really like…

DNS global anycast

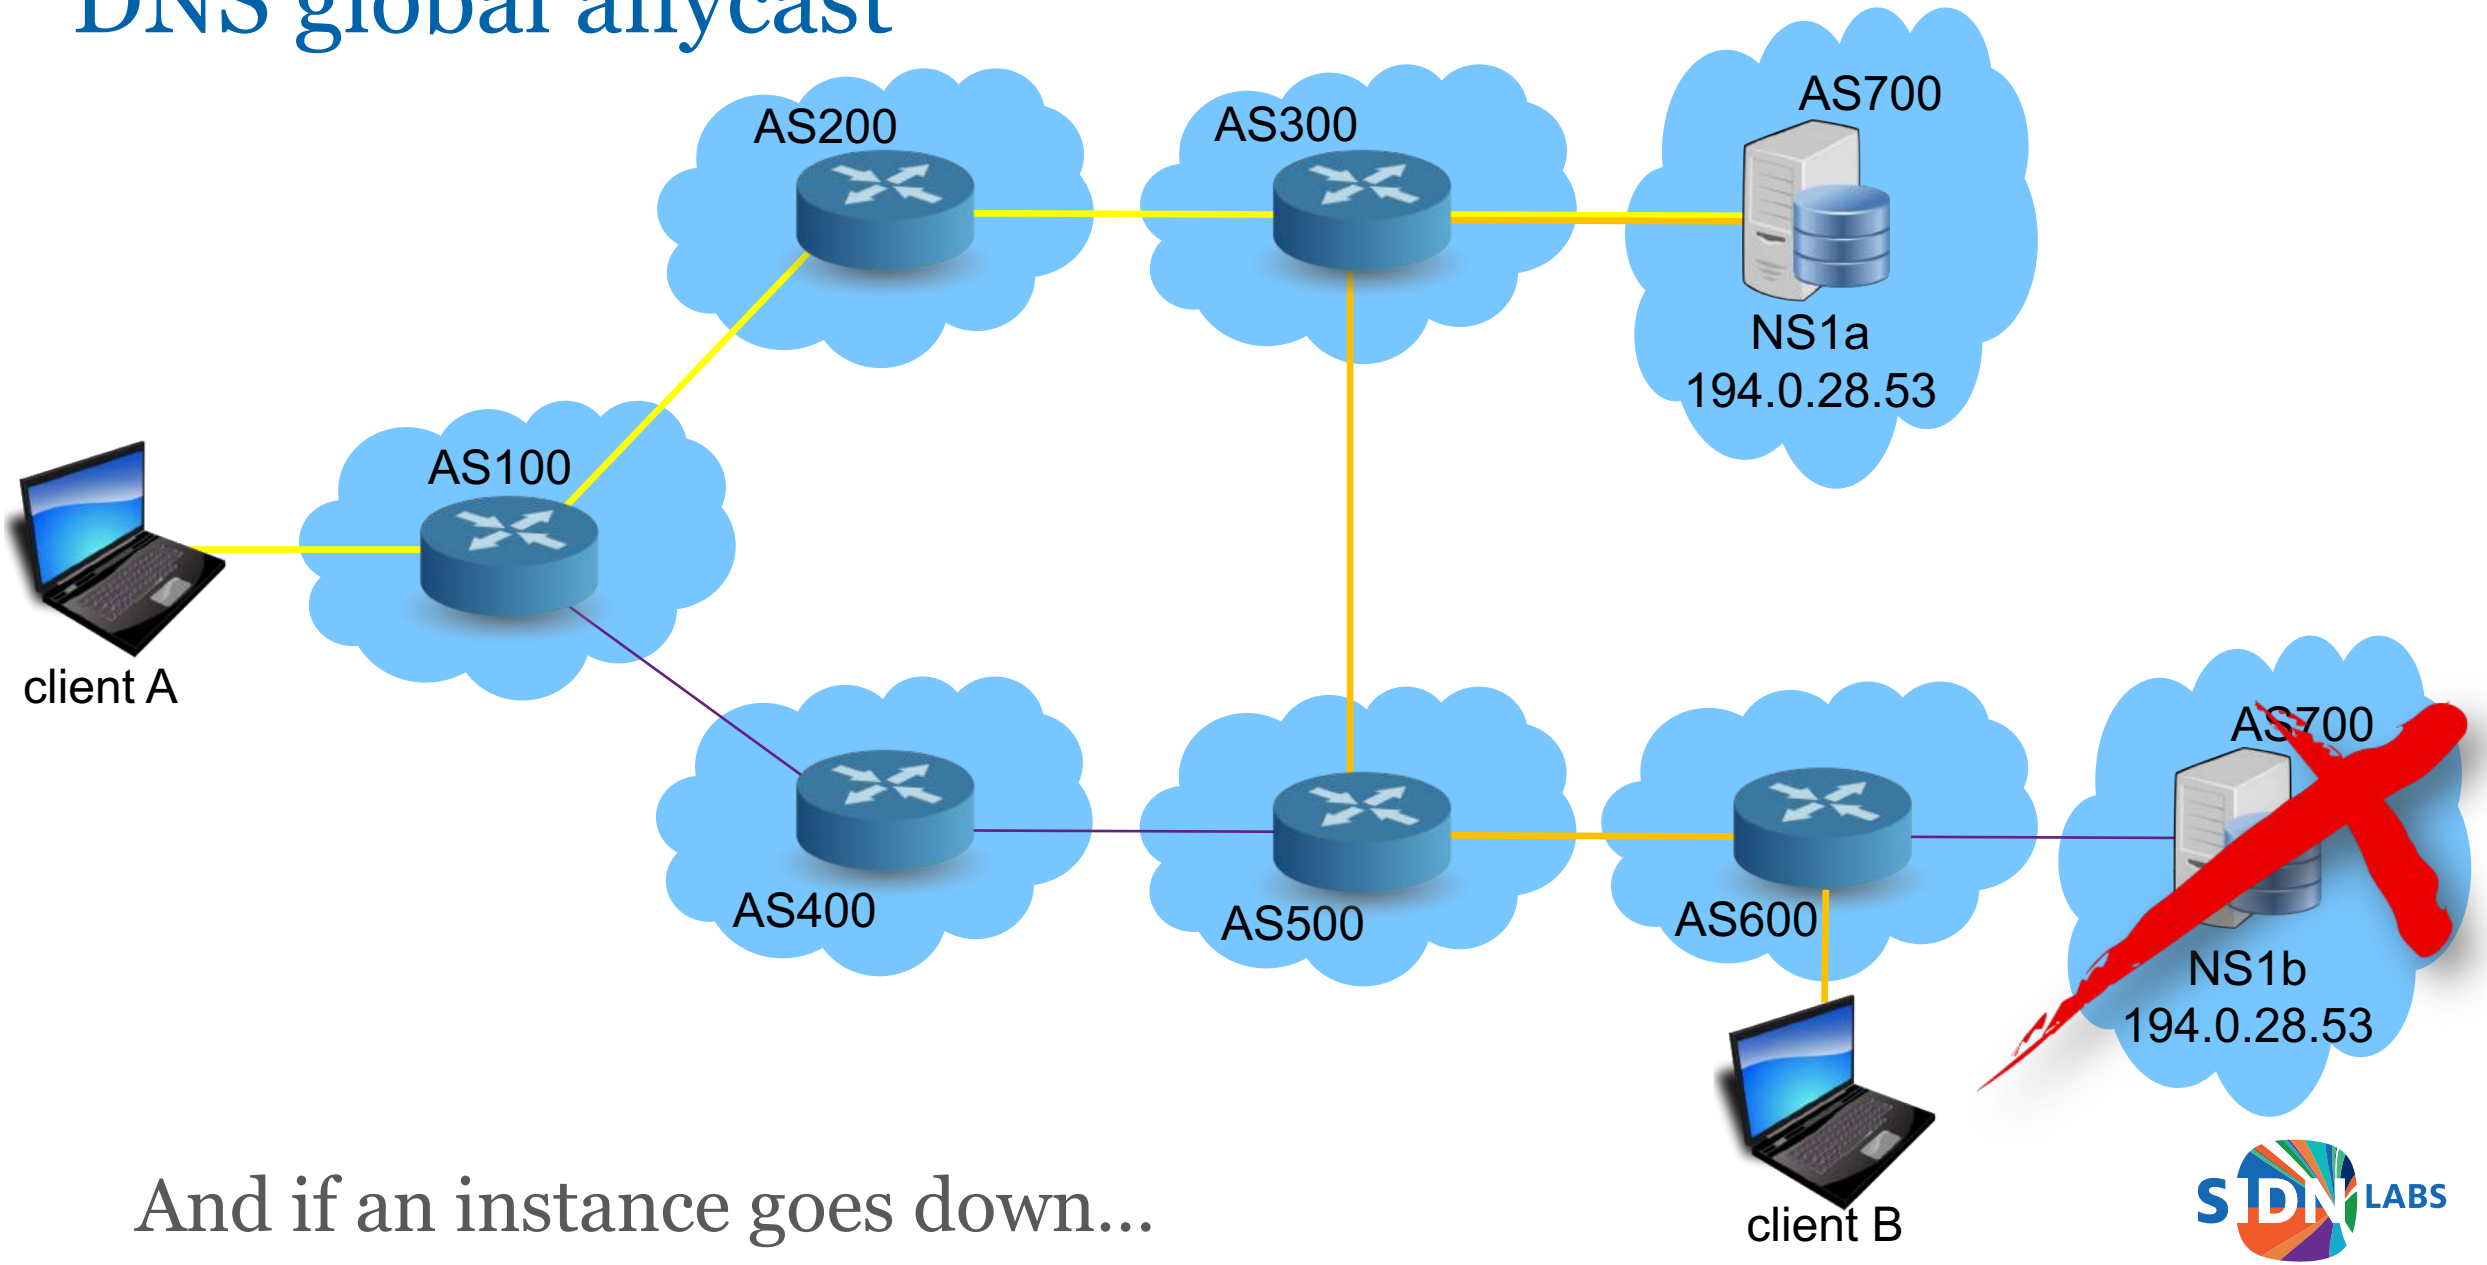

And if an instance goes down…

Problem solved…?

Problem solved, or…?

SDN LABS

Main cause: (insecure) IoT devices

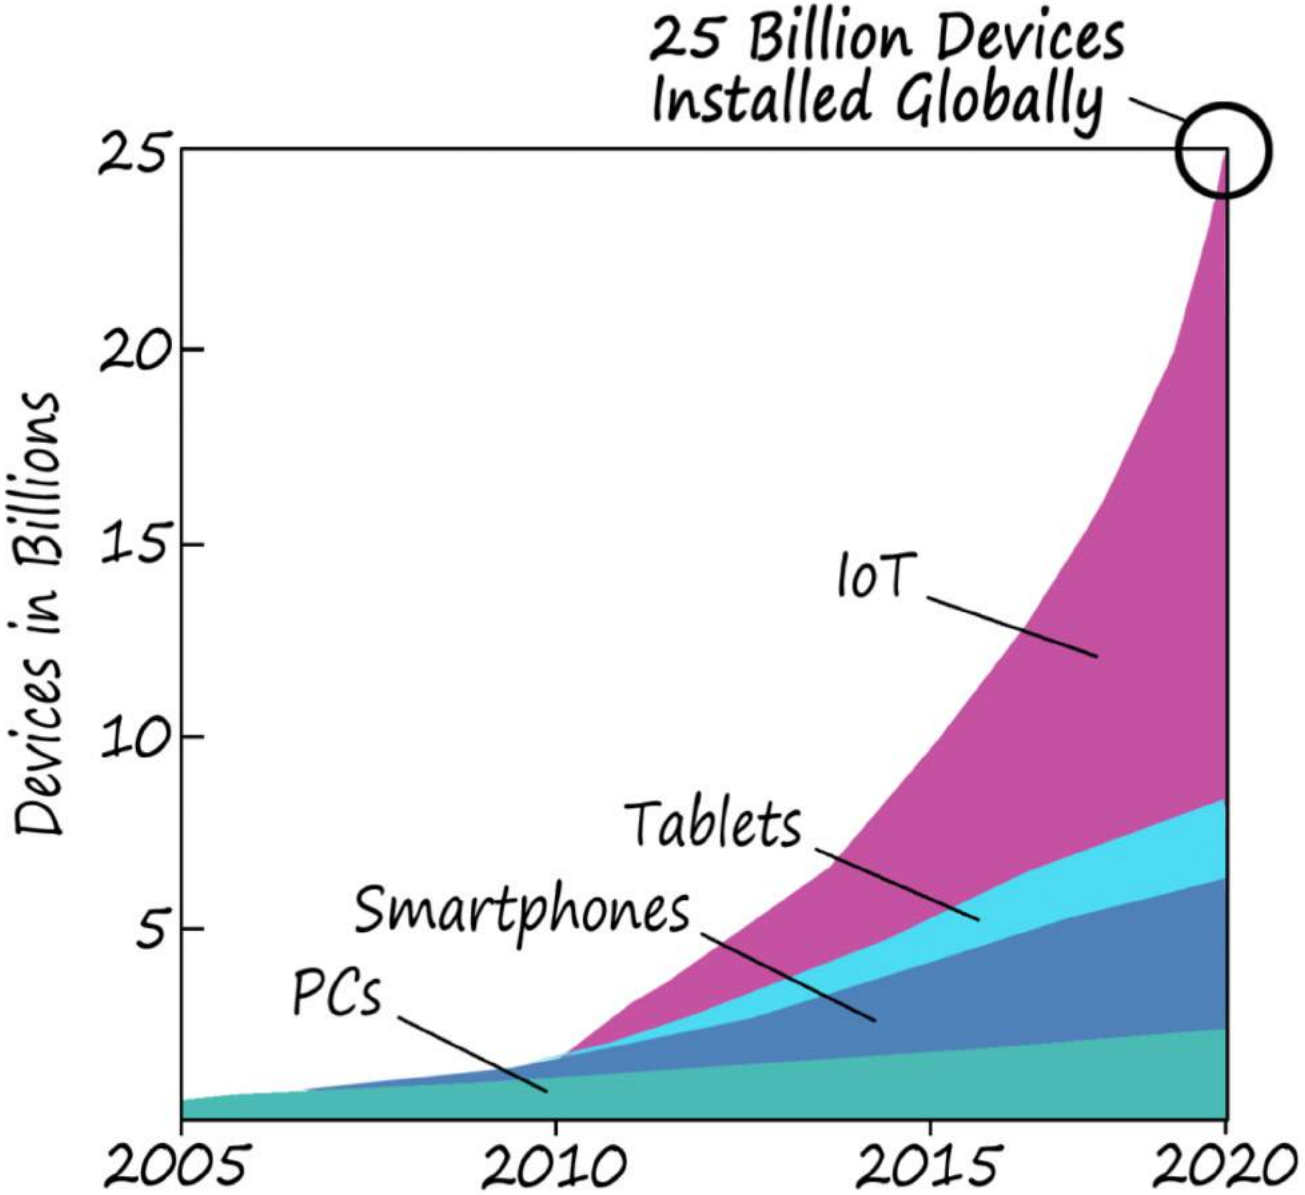

Main cause: (insecure) IoT devices

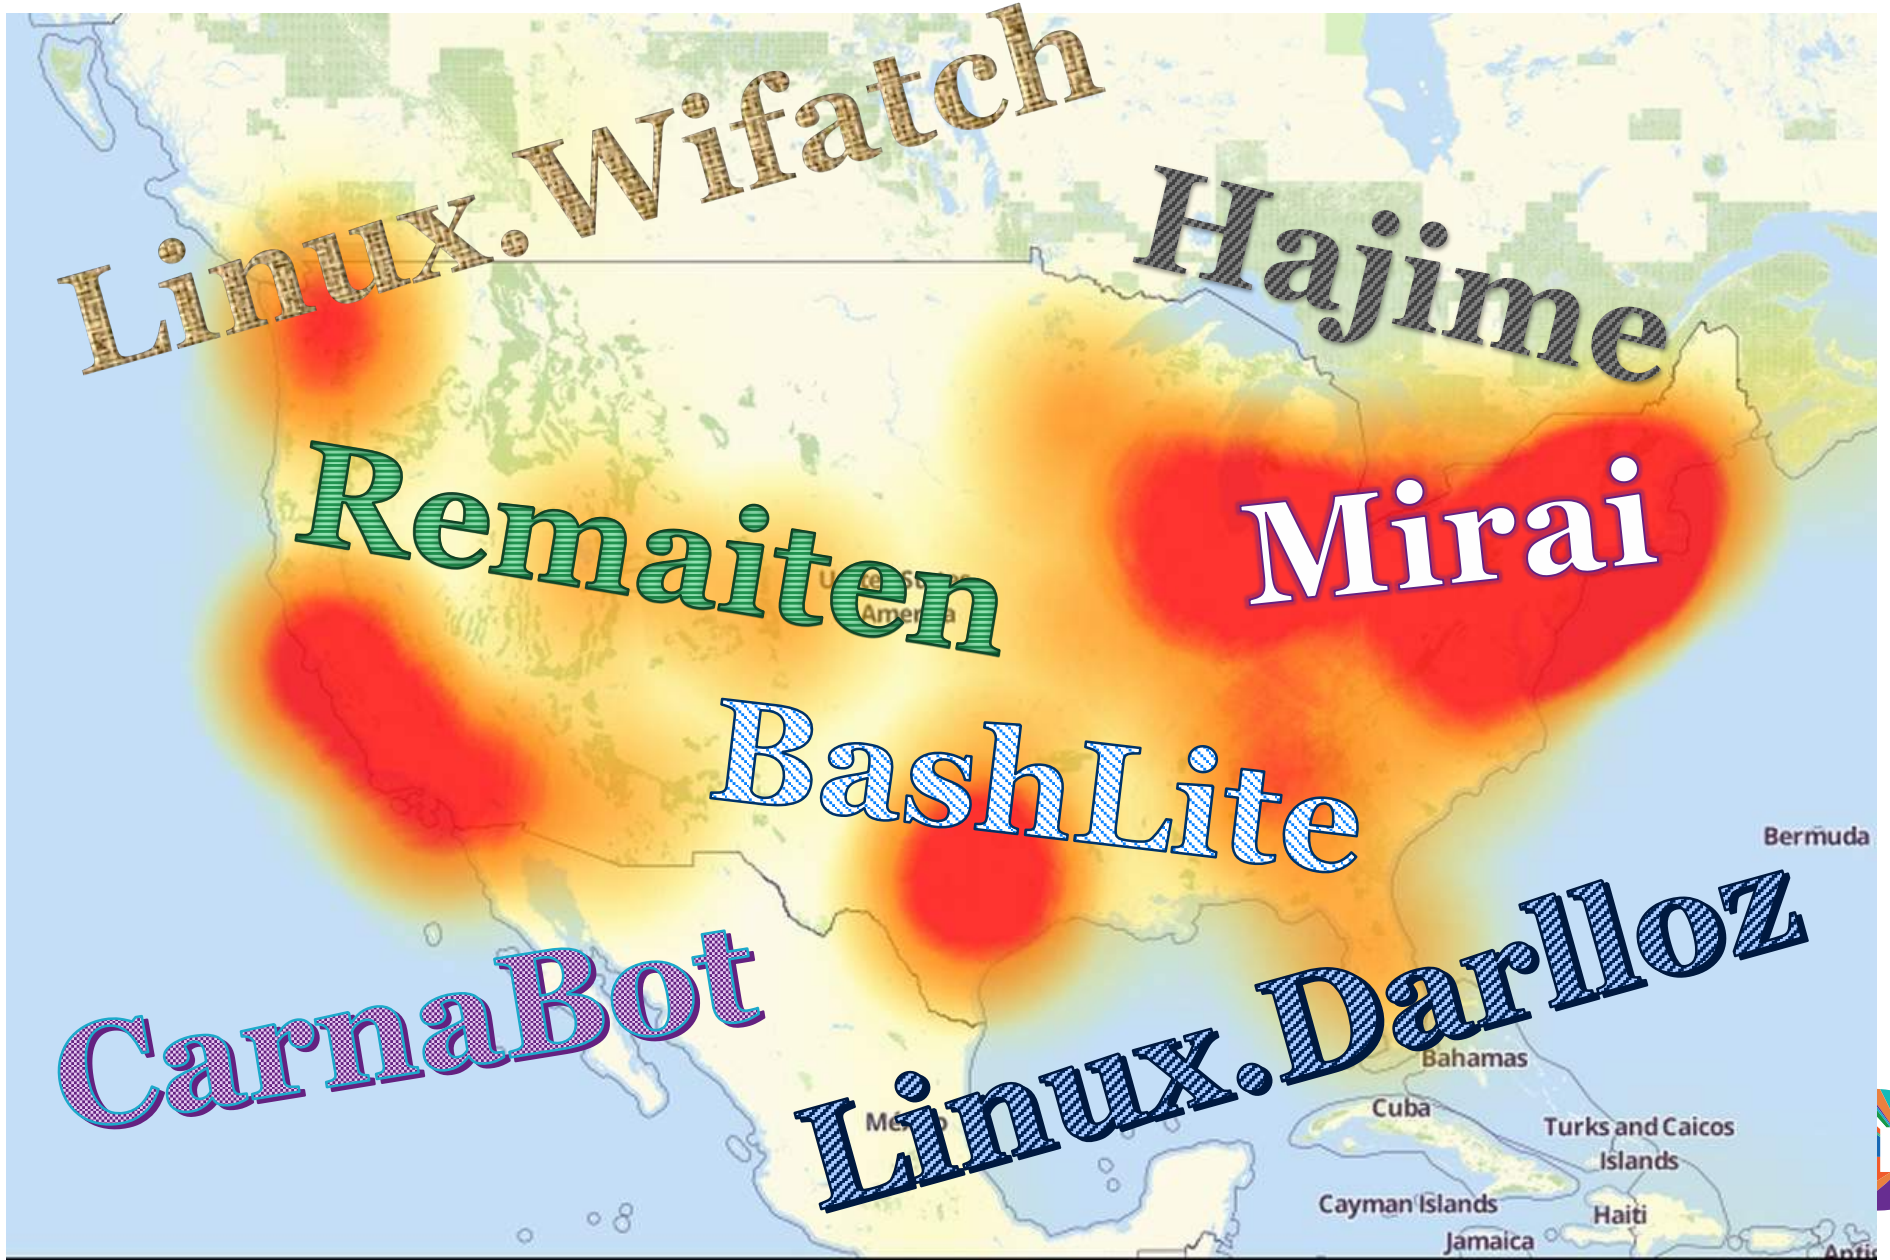

LABS

Main cause: Record breaking DDoS attacks

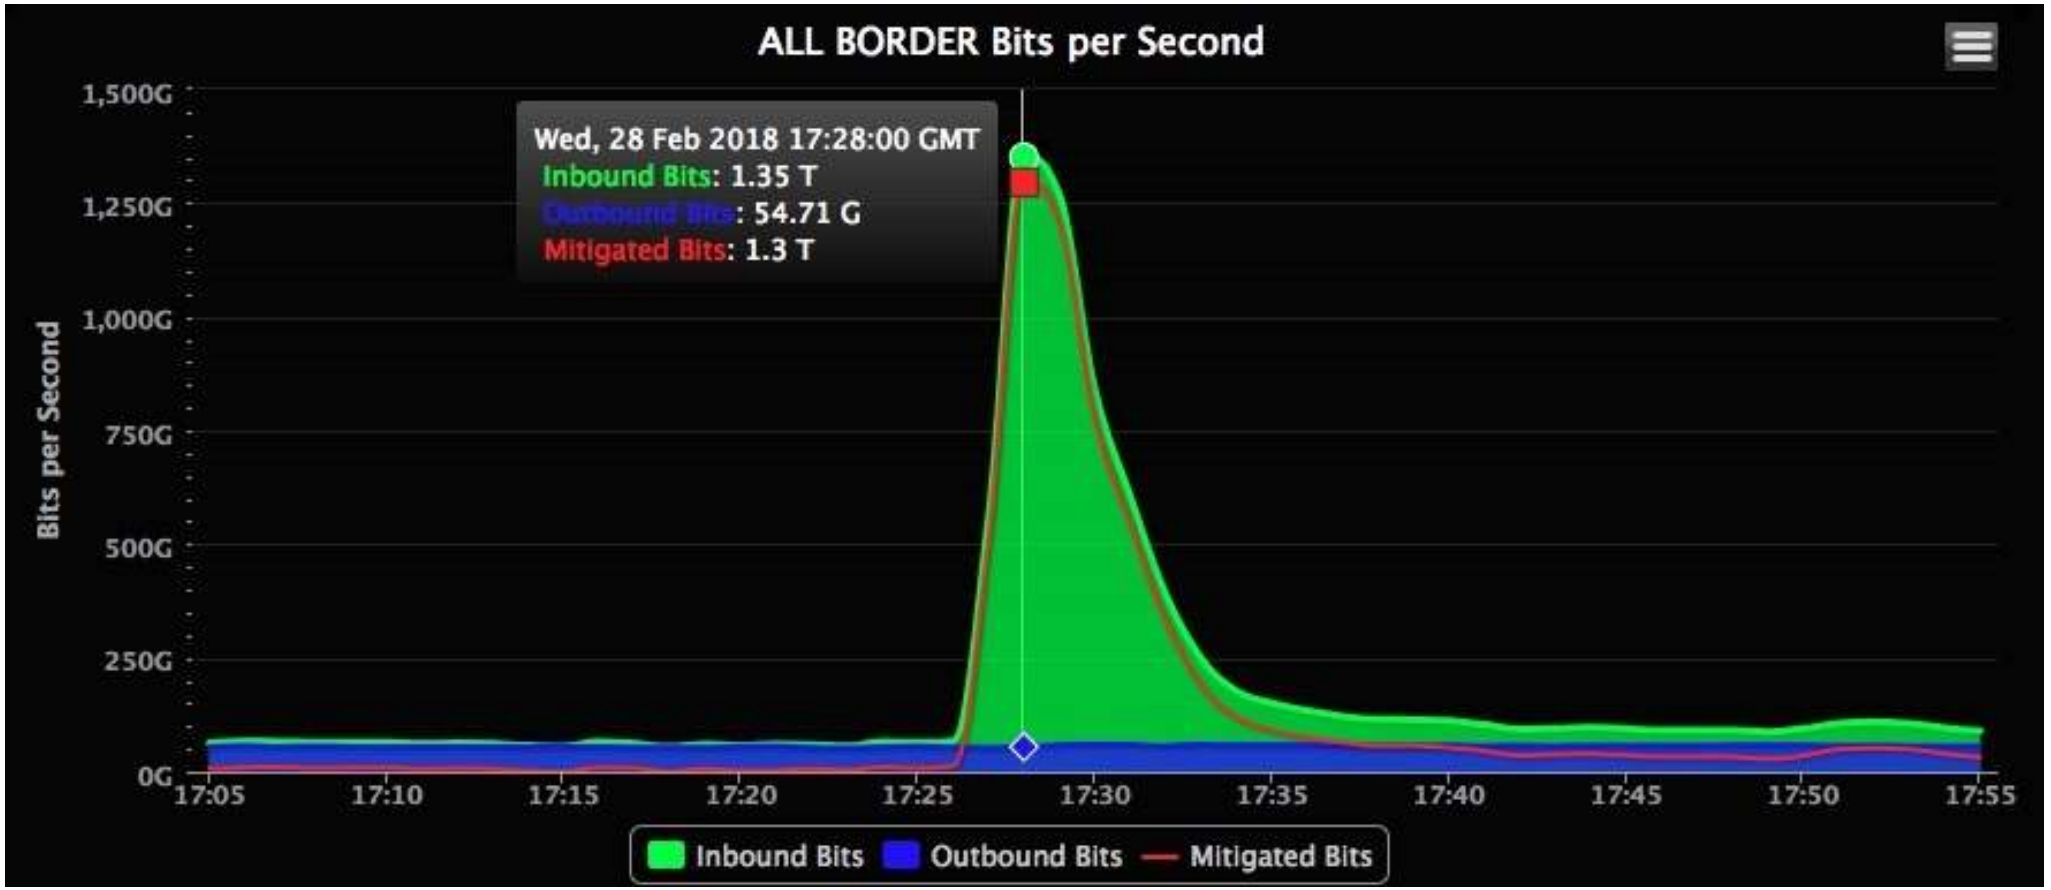

A rat race we can't win.

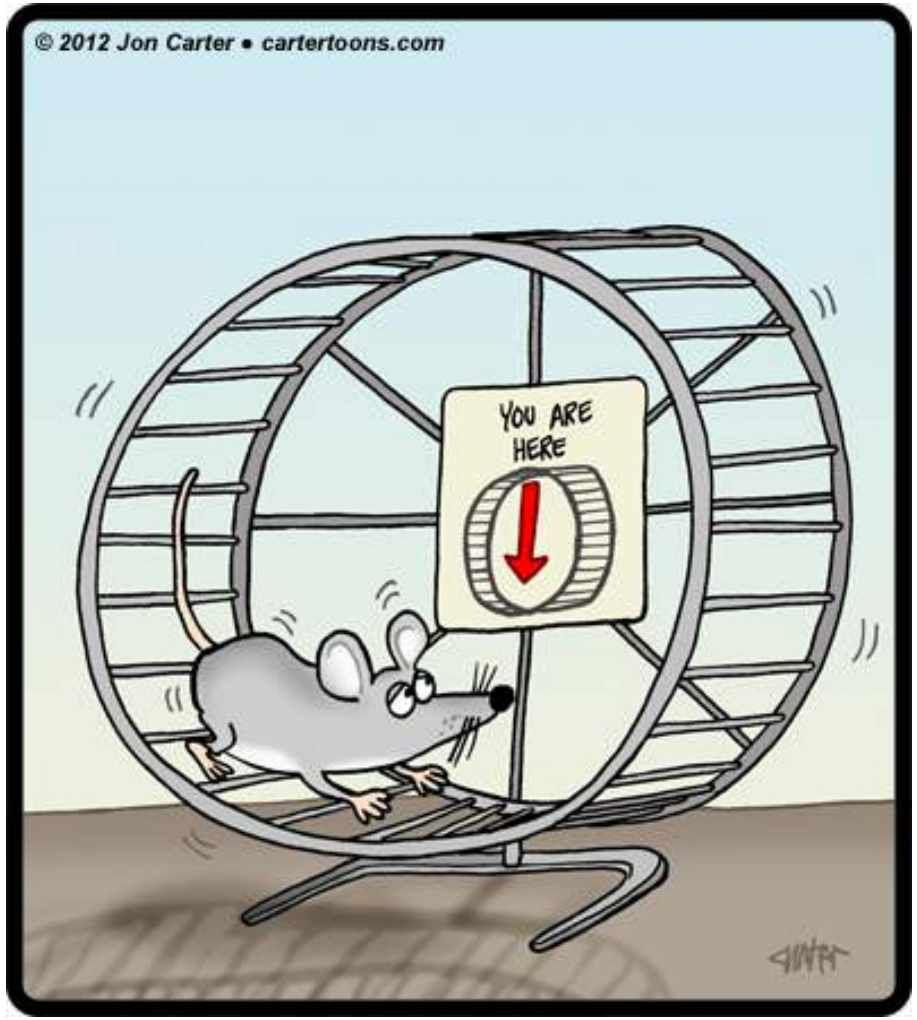

Paradigm shift

It is better to remain available for some than to be not available at all.

(All resolvers are equal, but some resolvers are more equal than others)

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

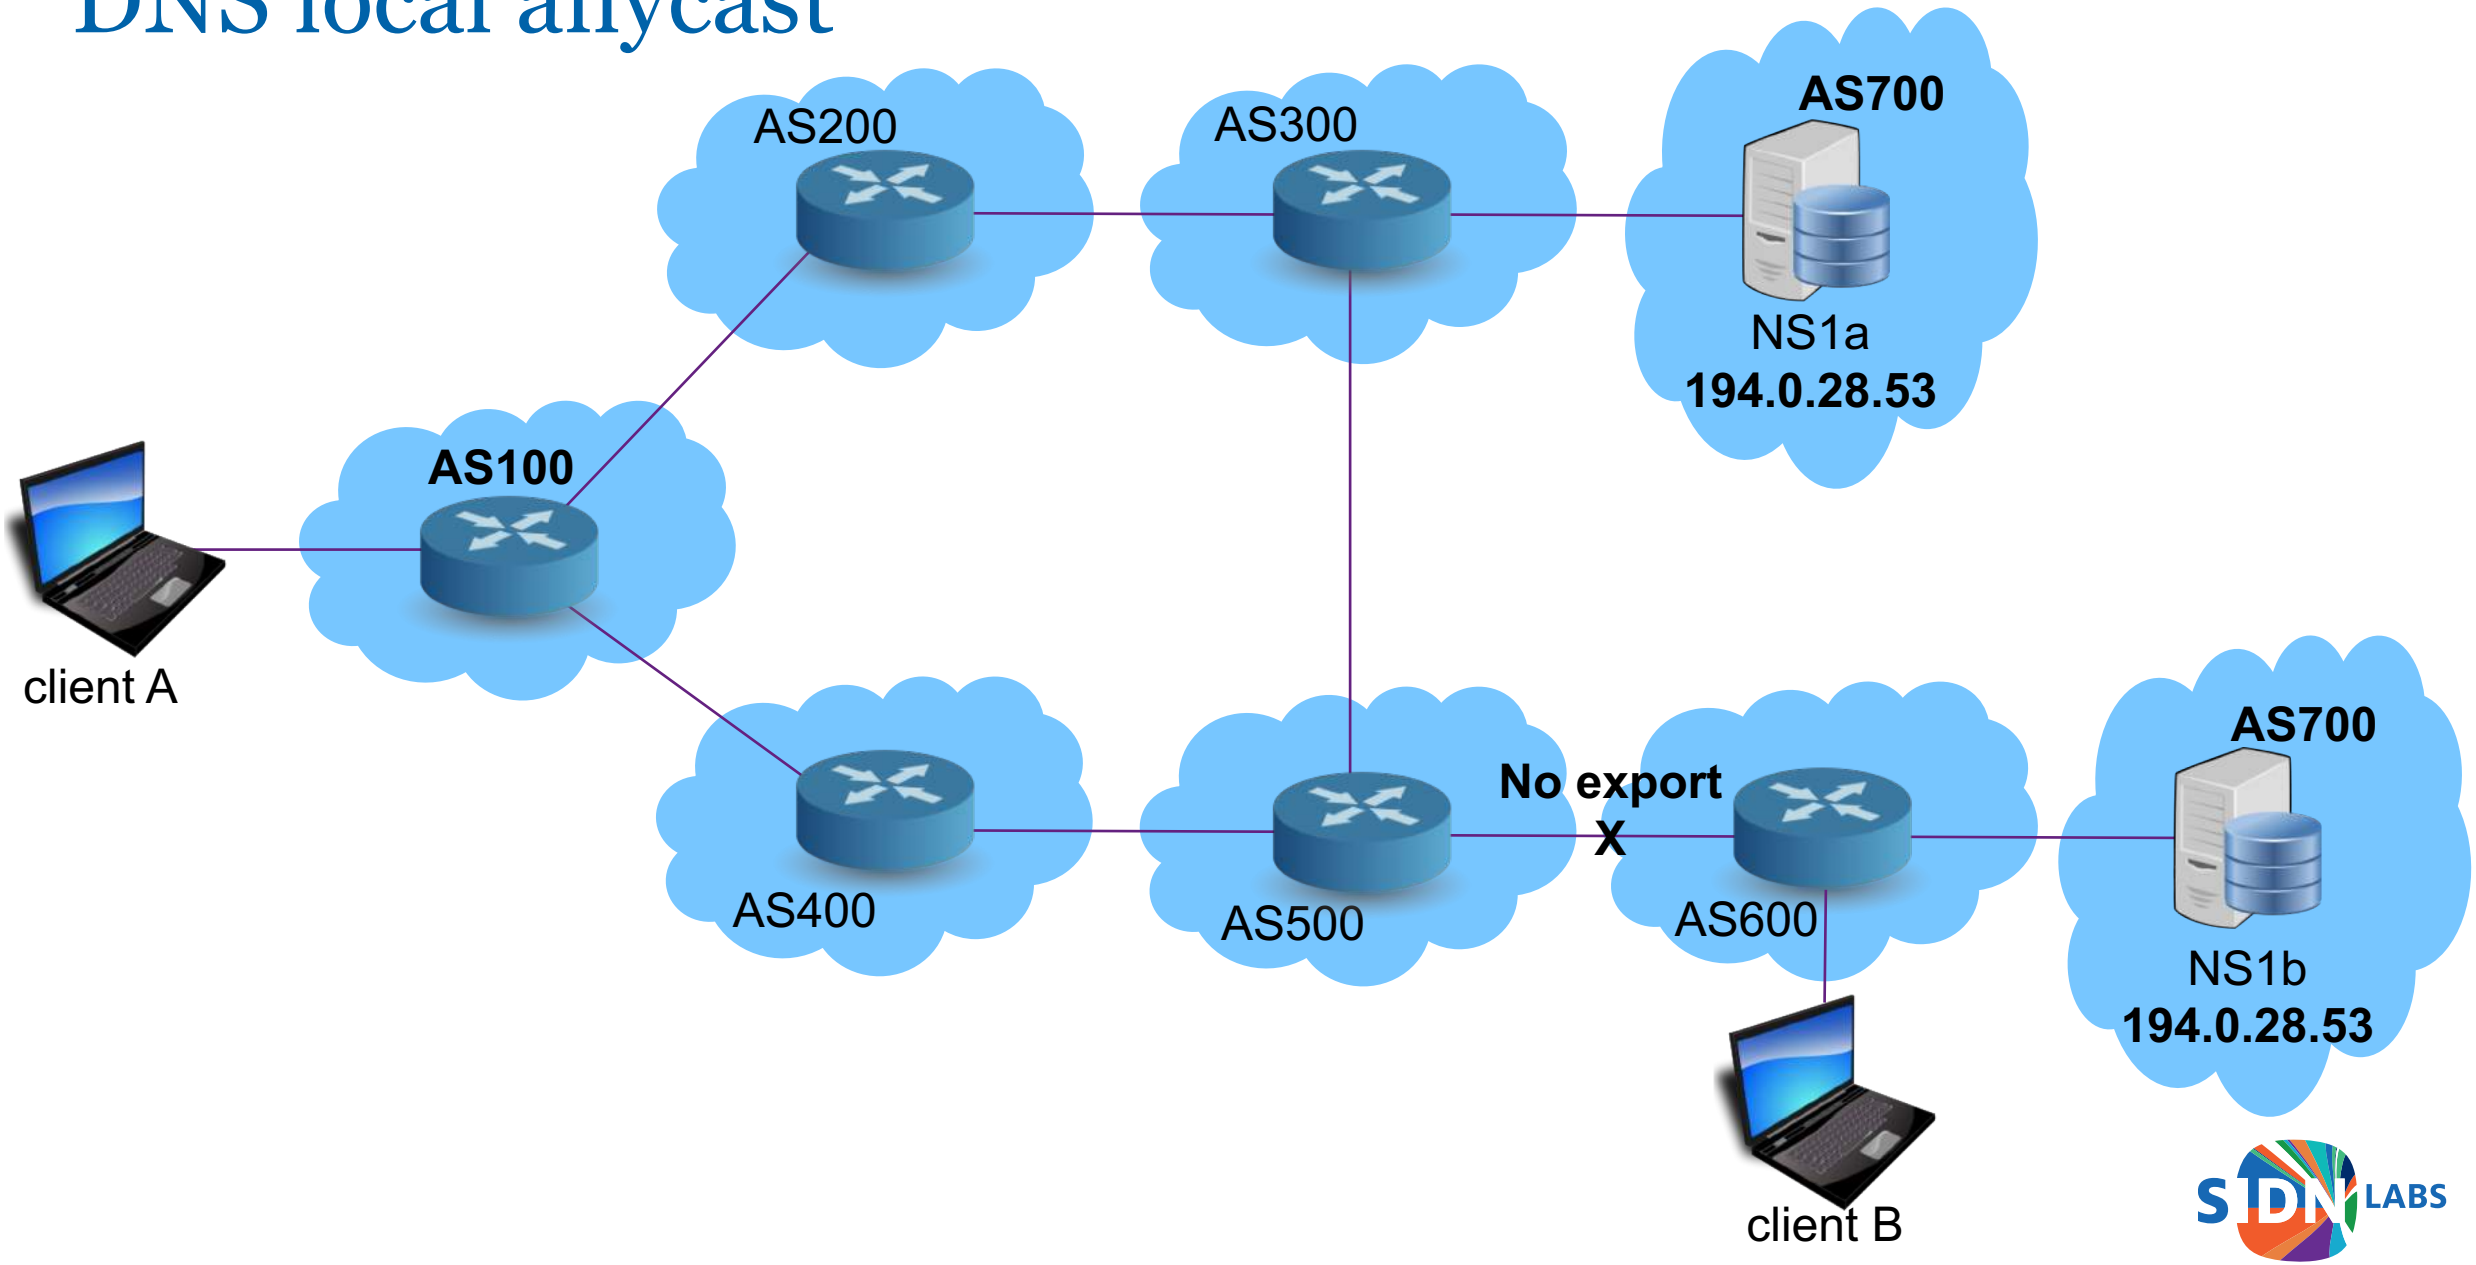

DNS local anycast

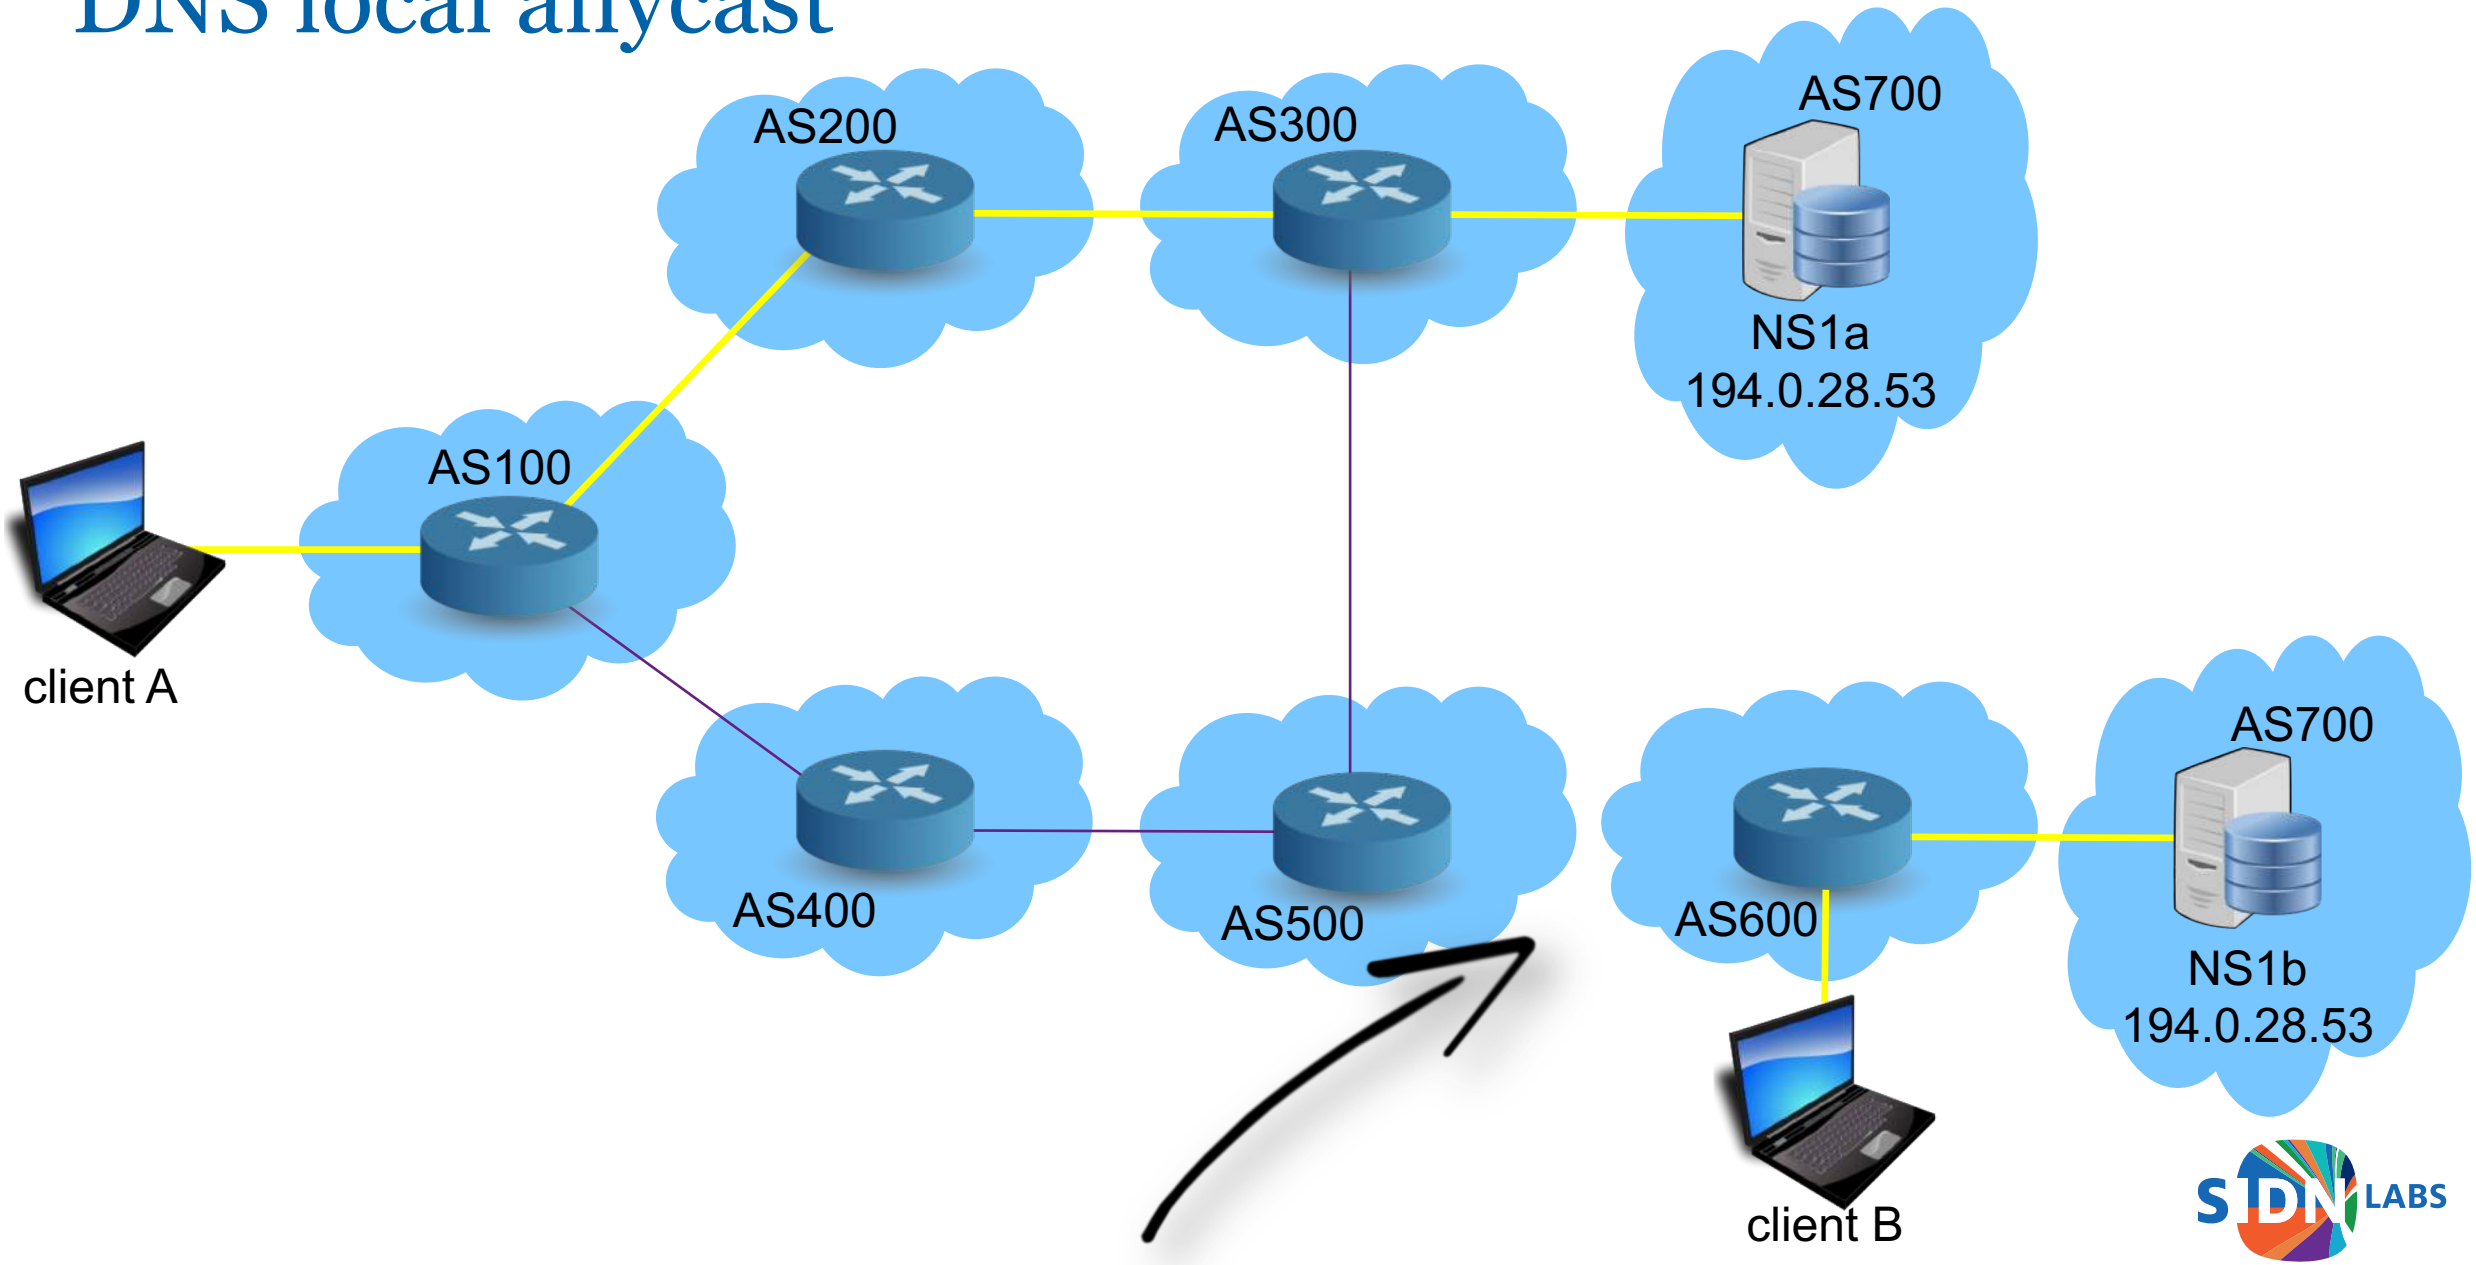

DNS local anycast

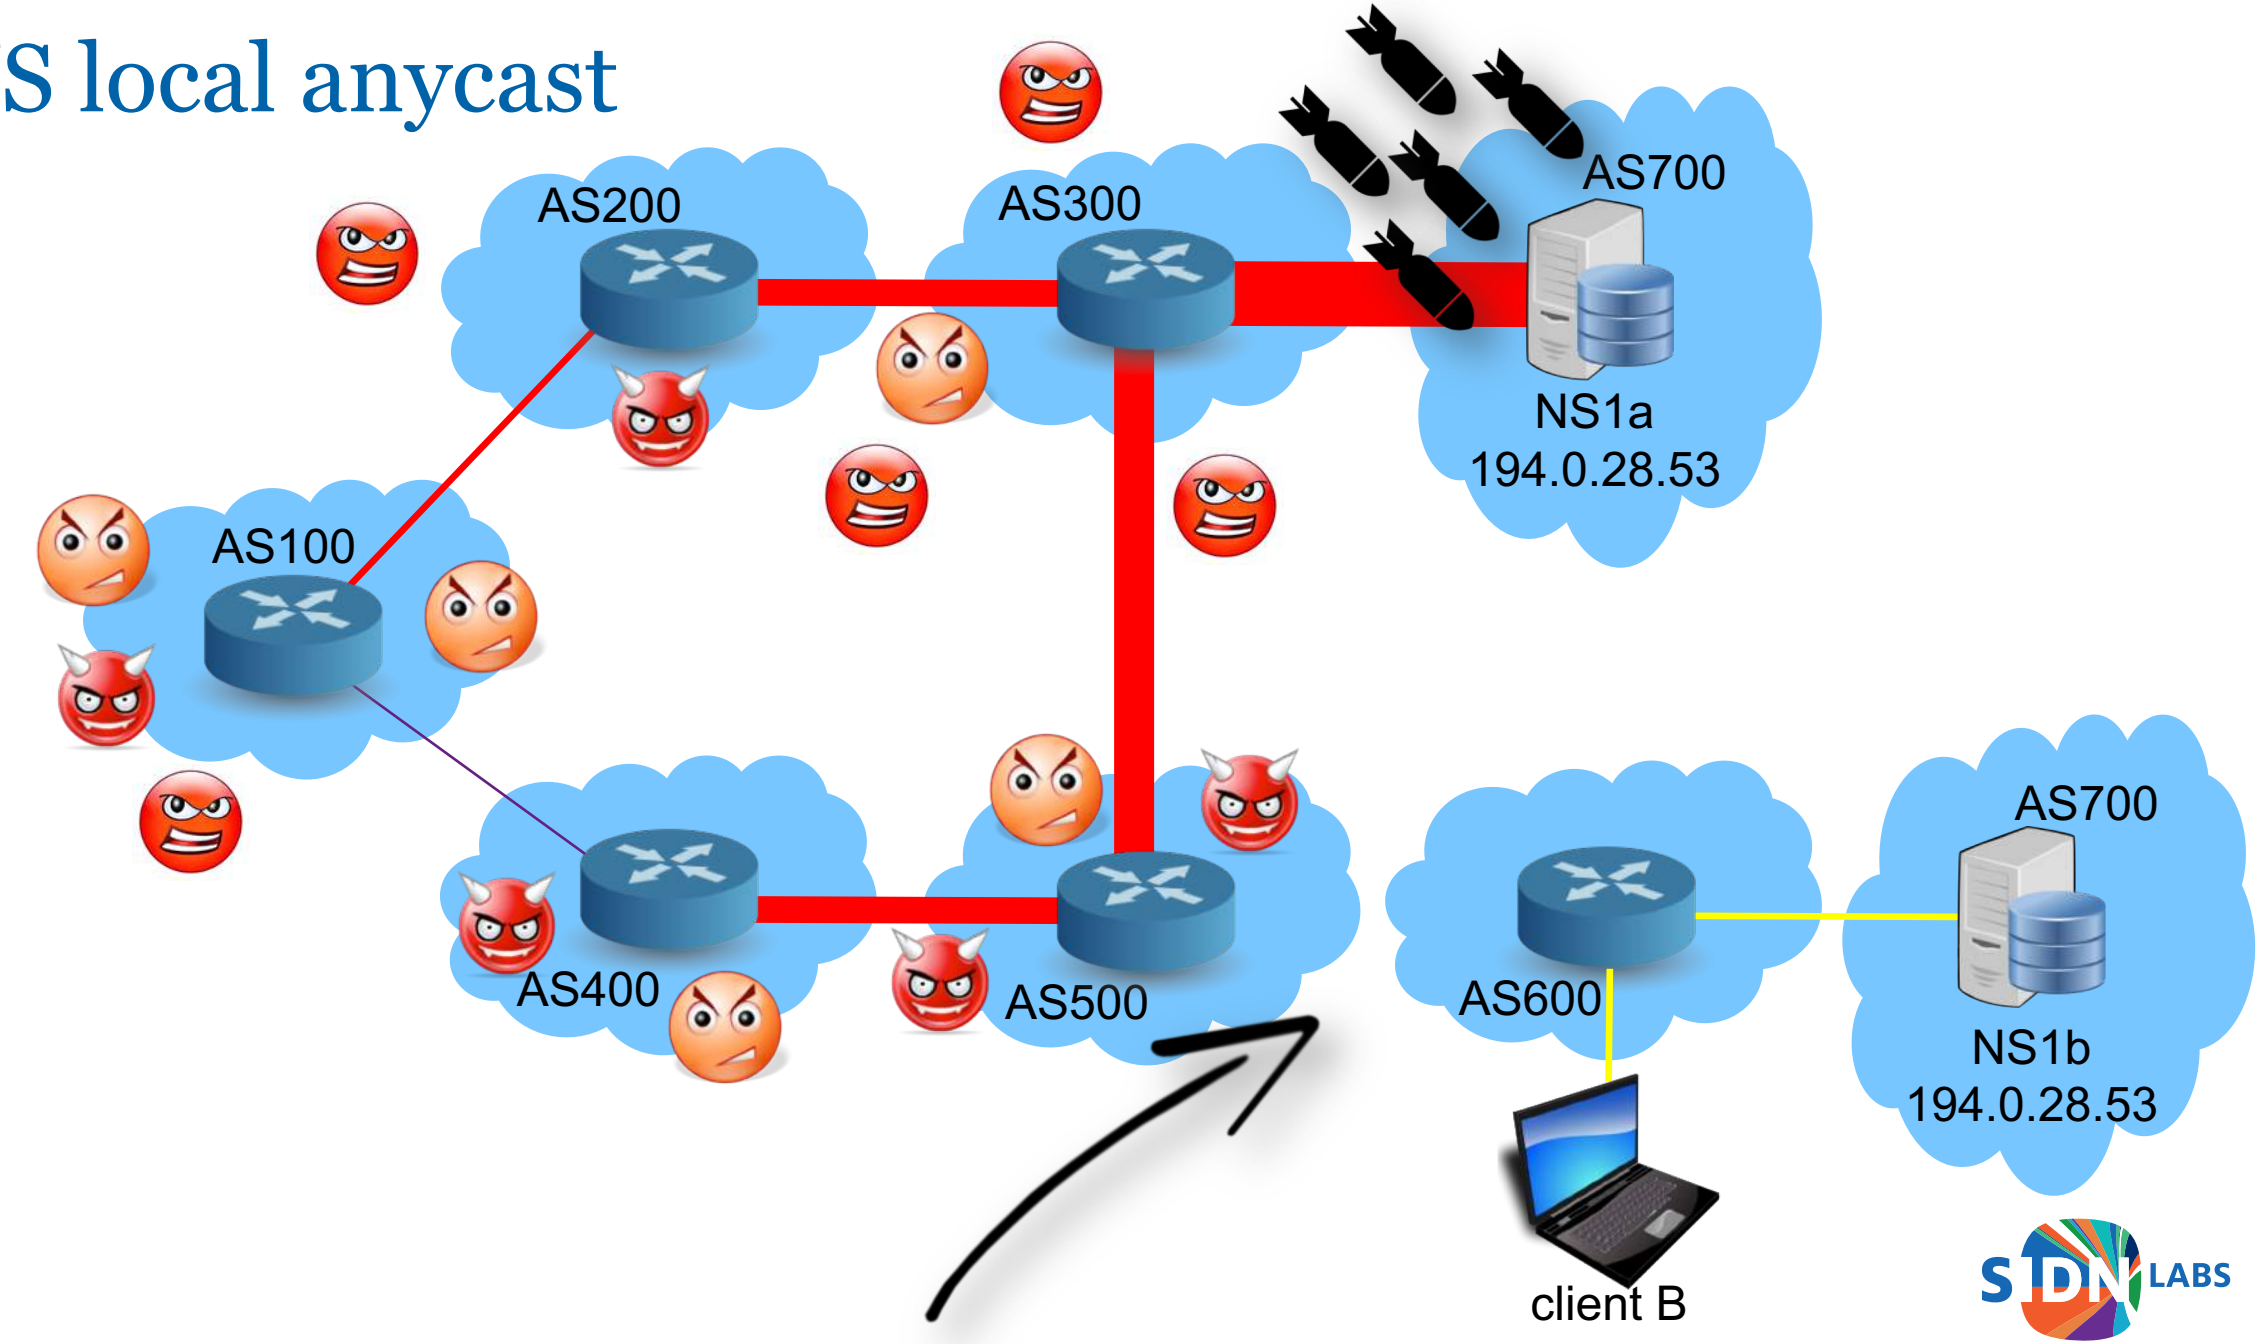

DNS local anycast

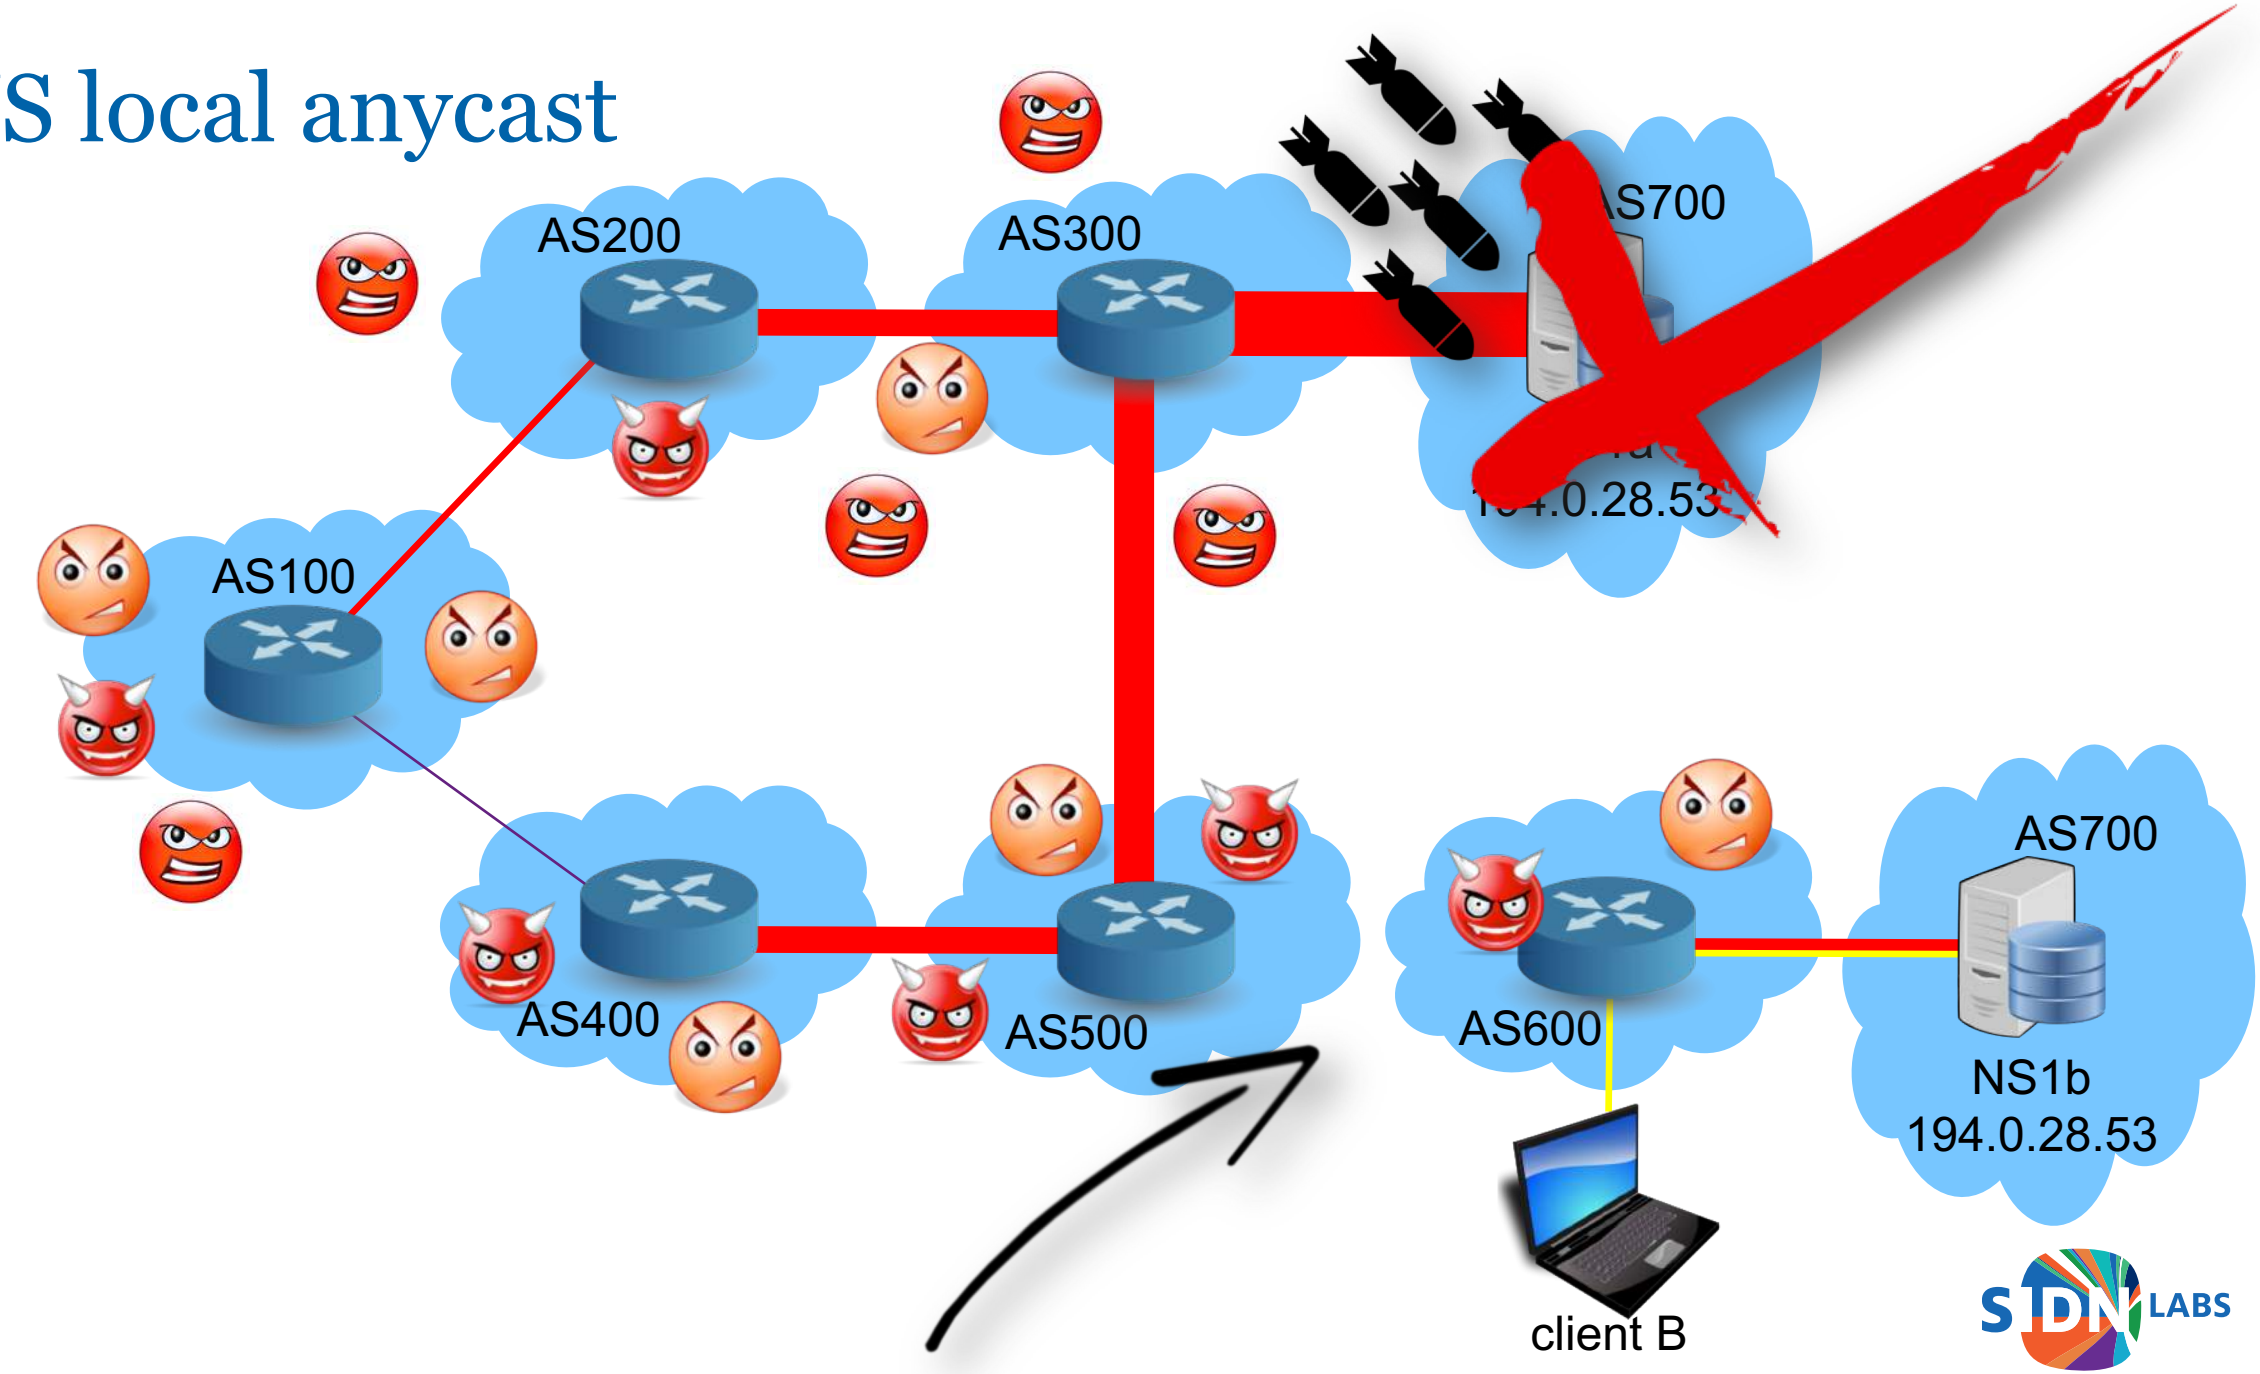

DNS local anycast

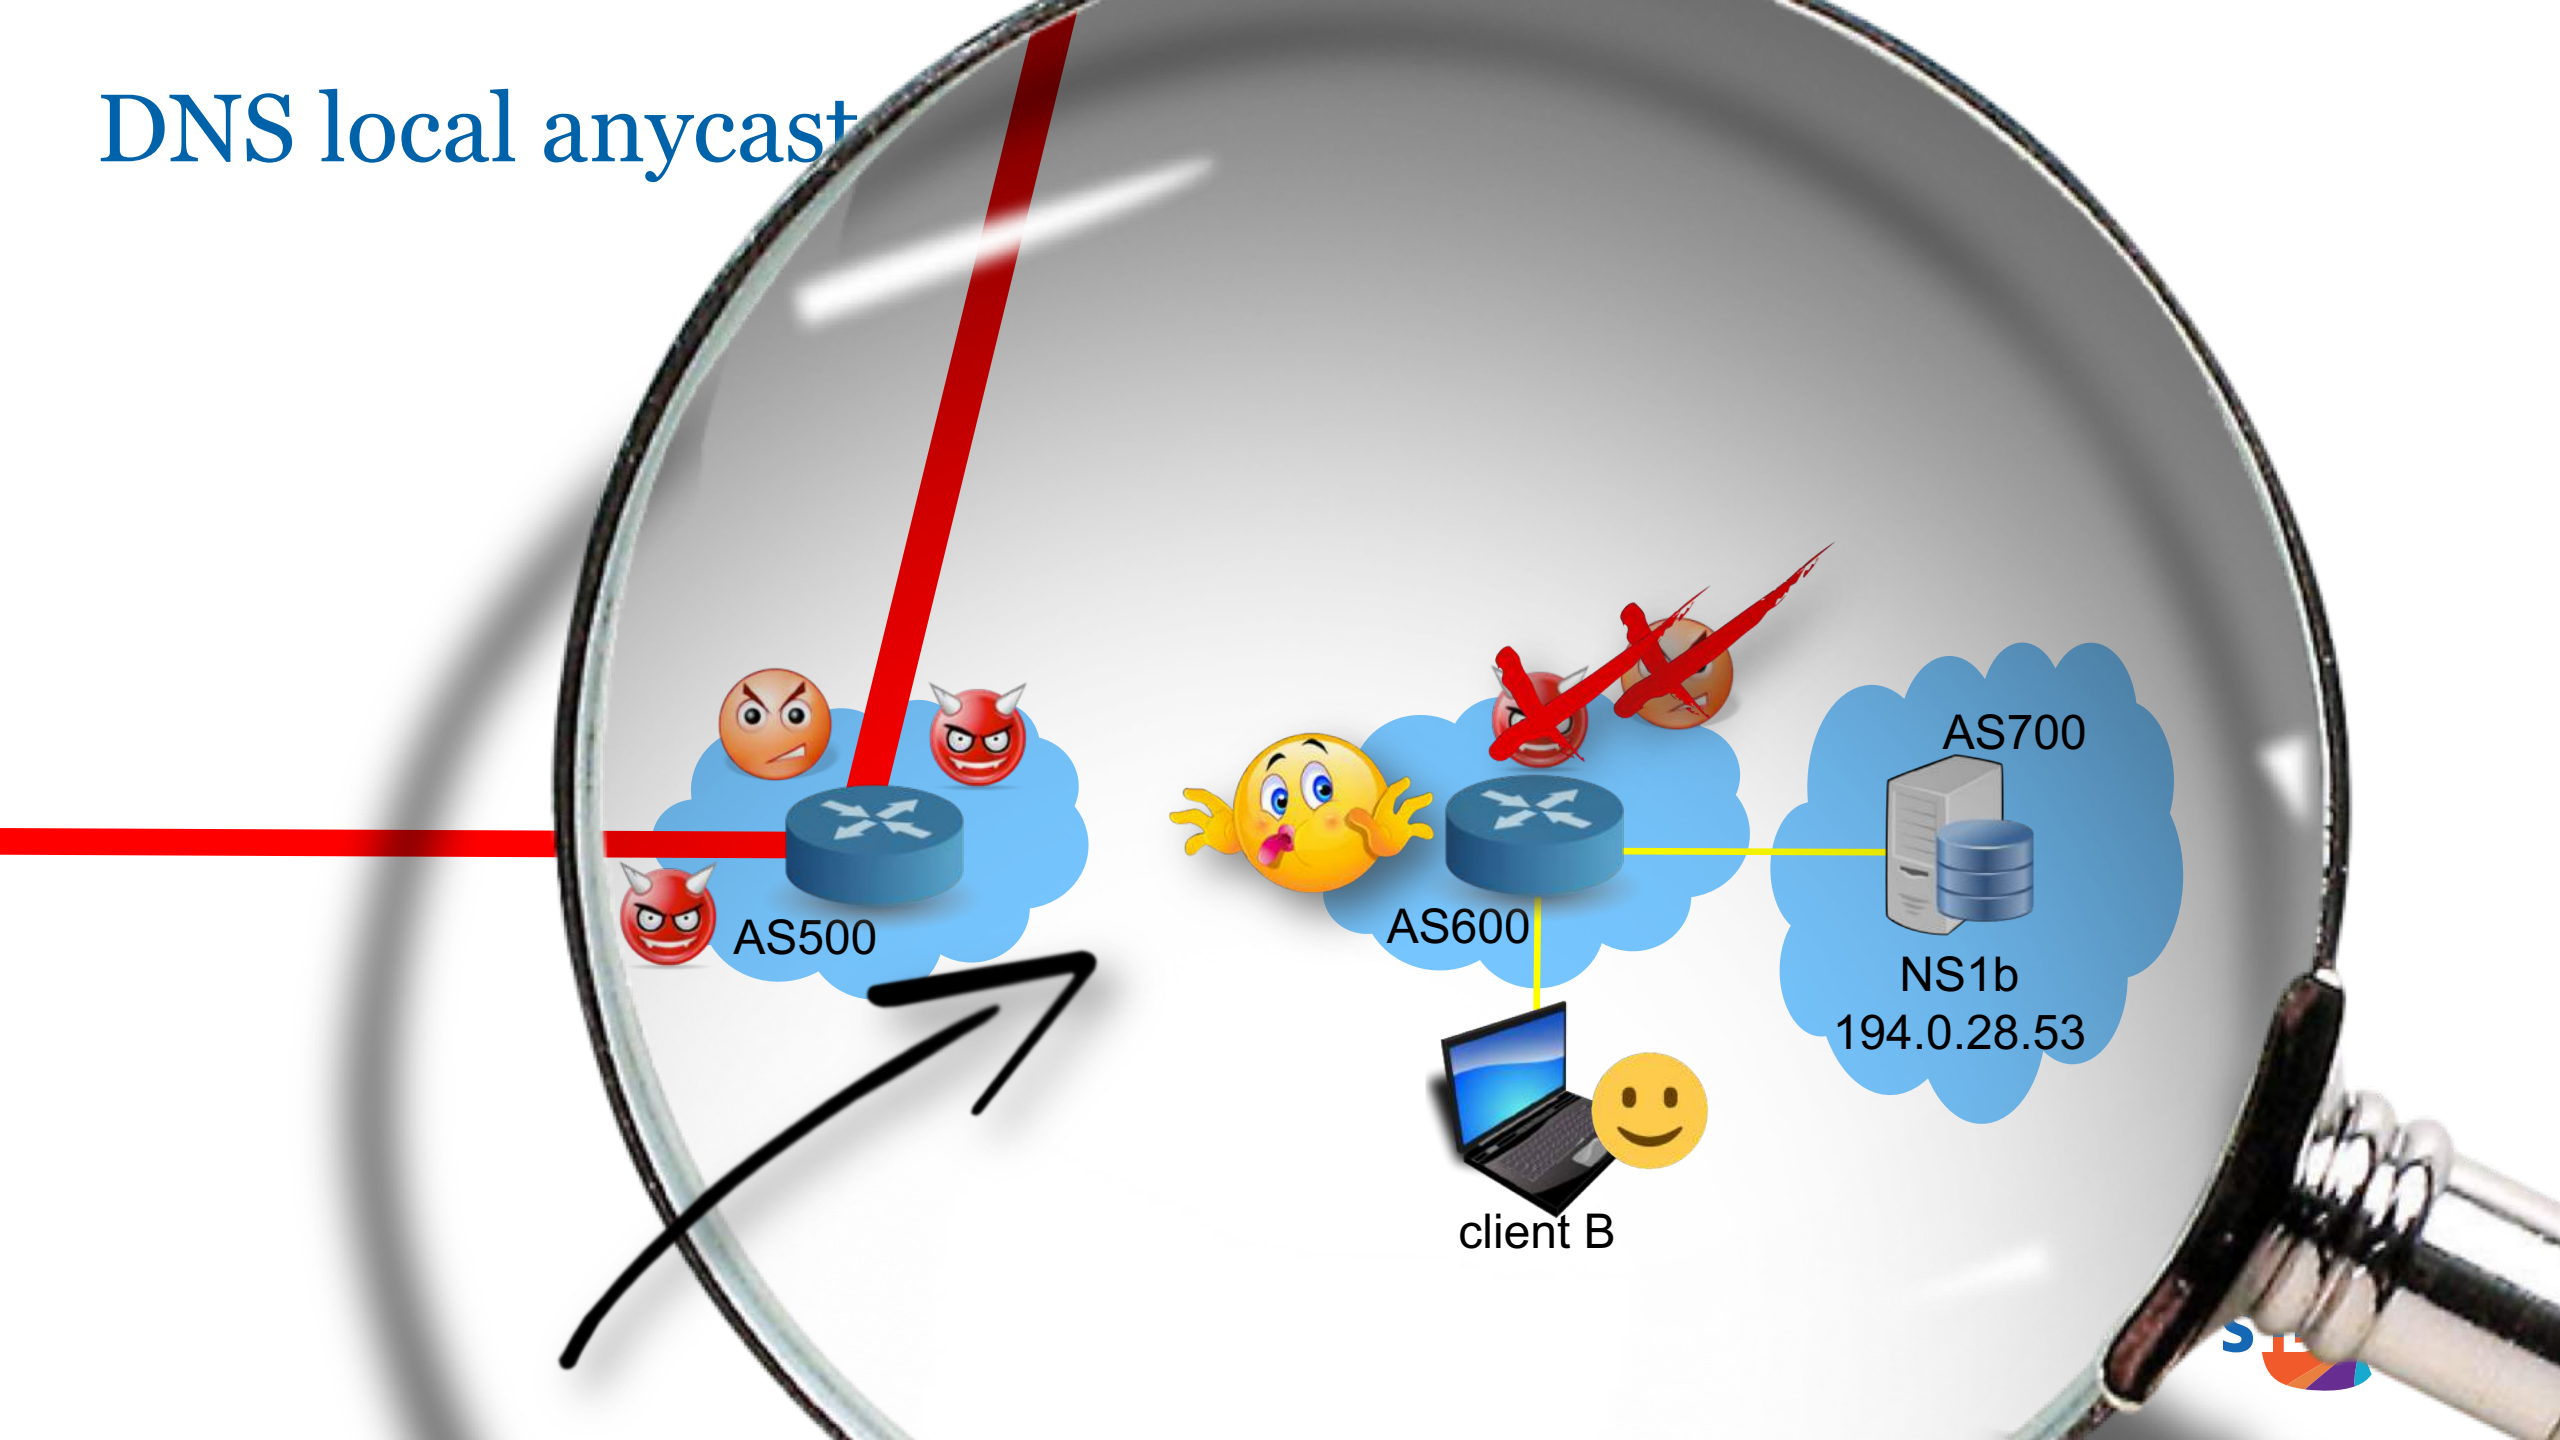

DNS local anycast – incident at AMS-IX

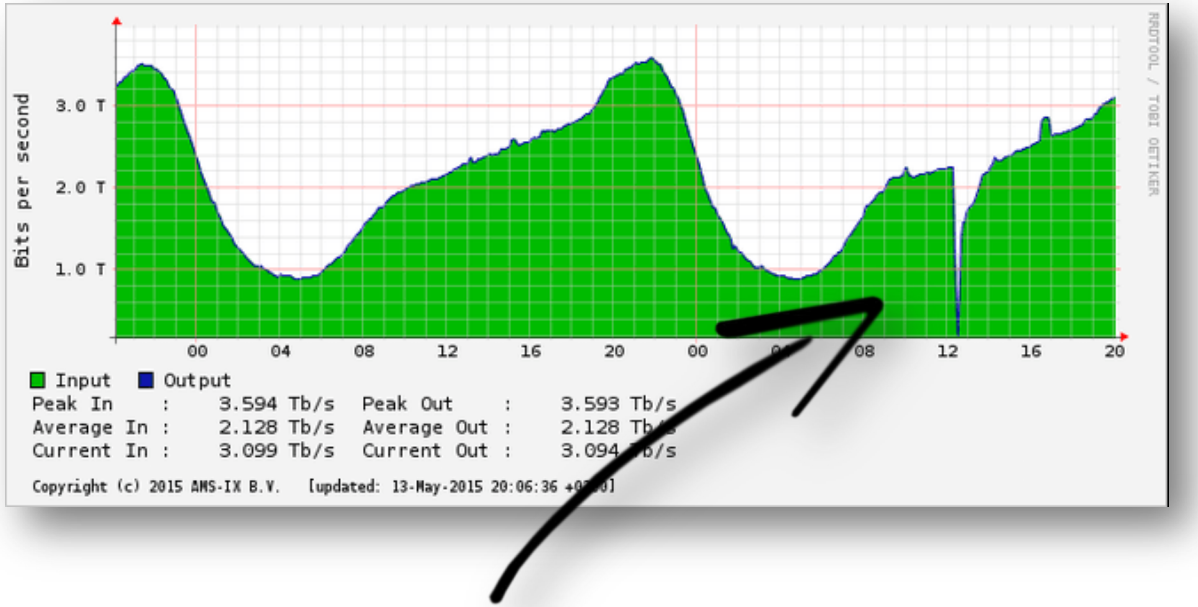

DNS local anycast – incident at AMS-IX

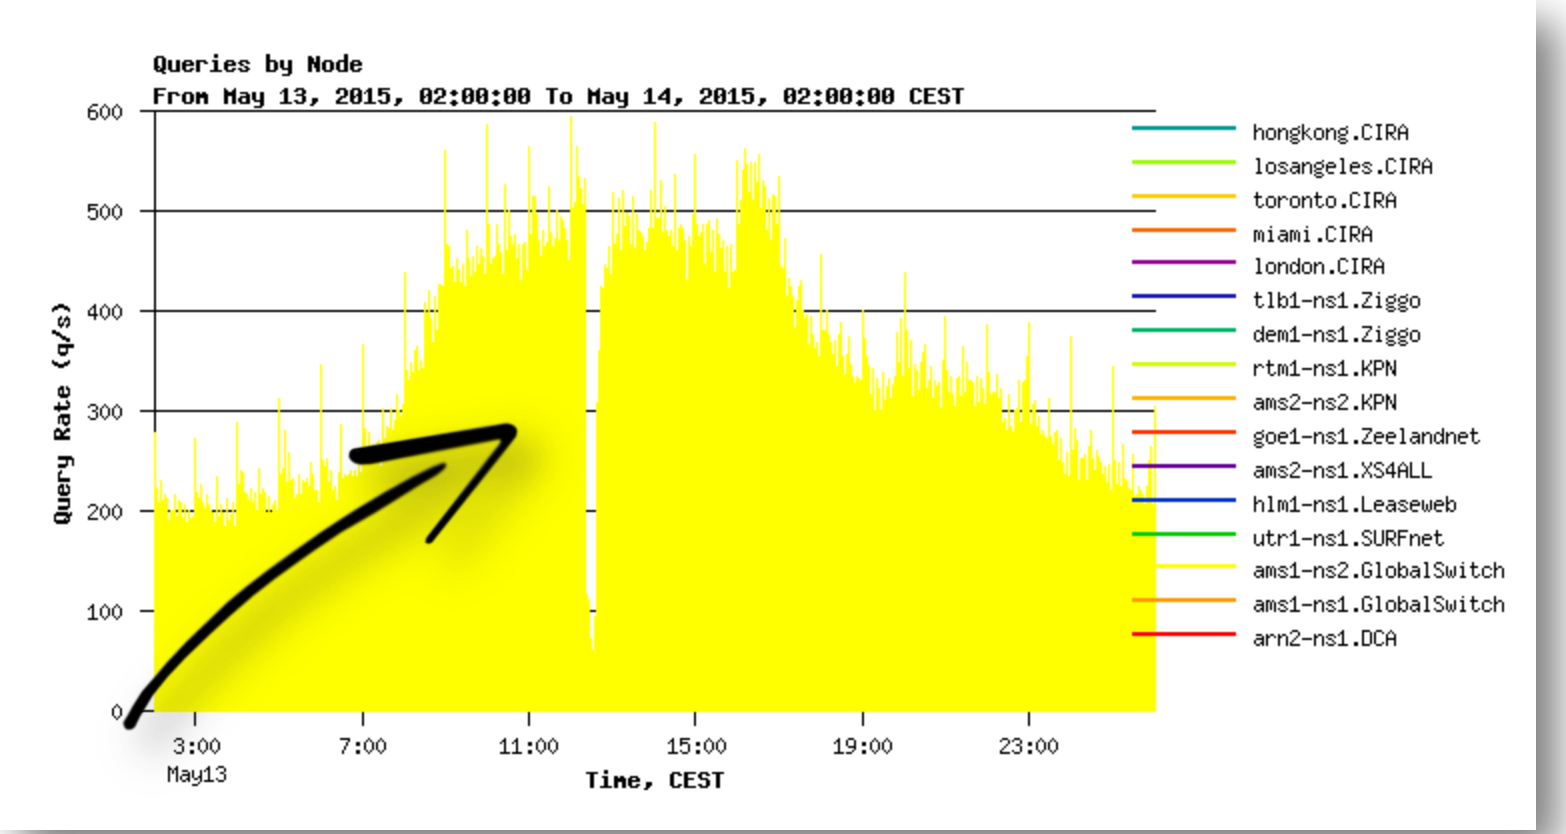

DNS local anycast – incident at AMS-IX

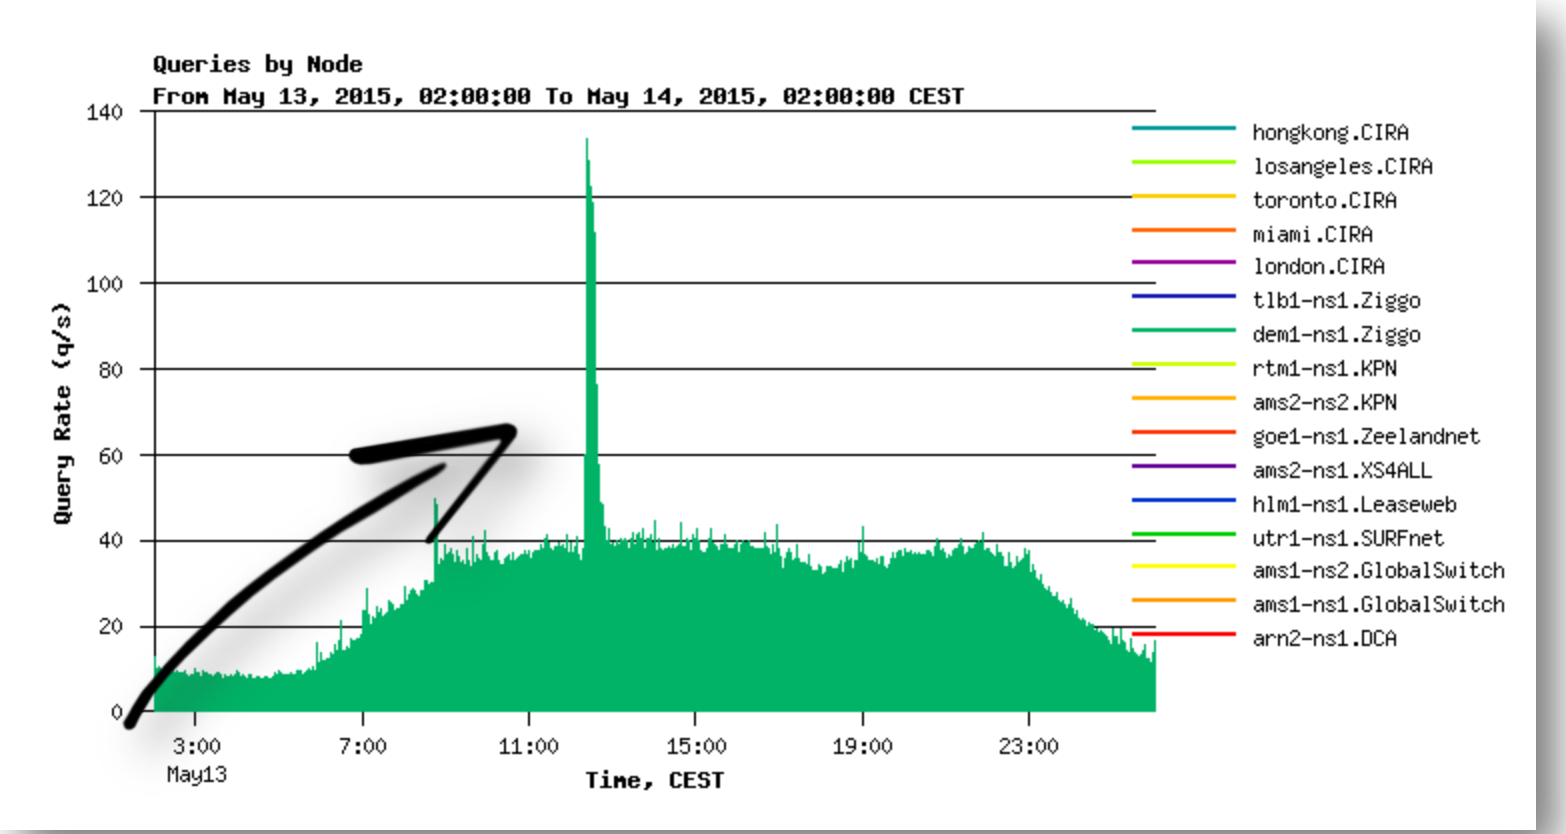

DNS local anycast – incident at AMS-IX

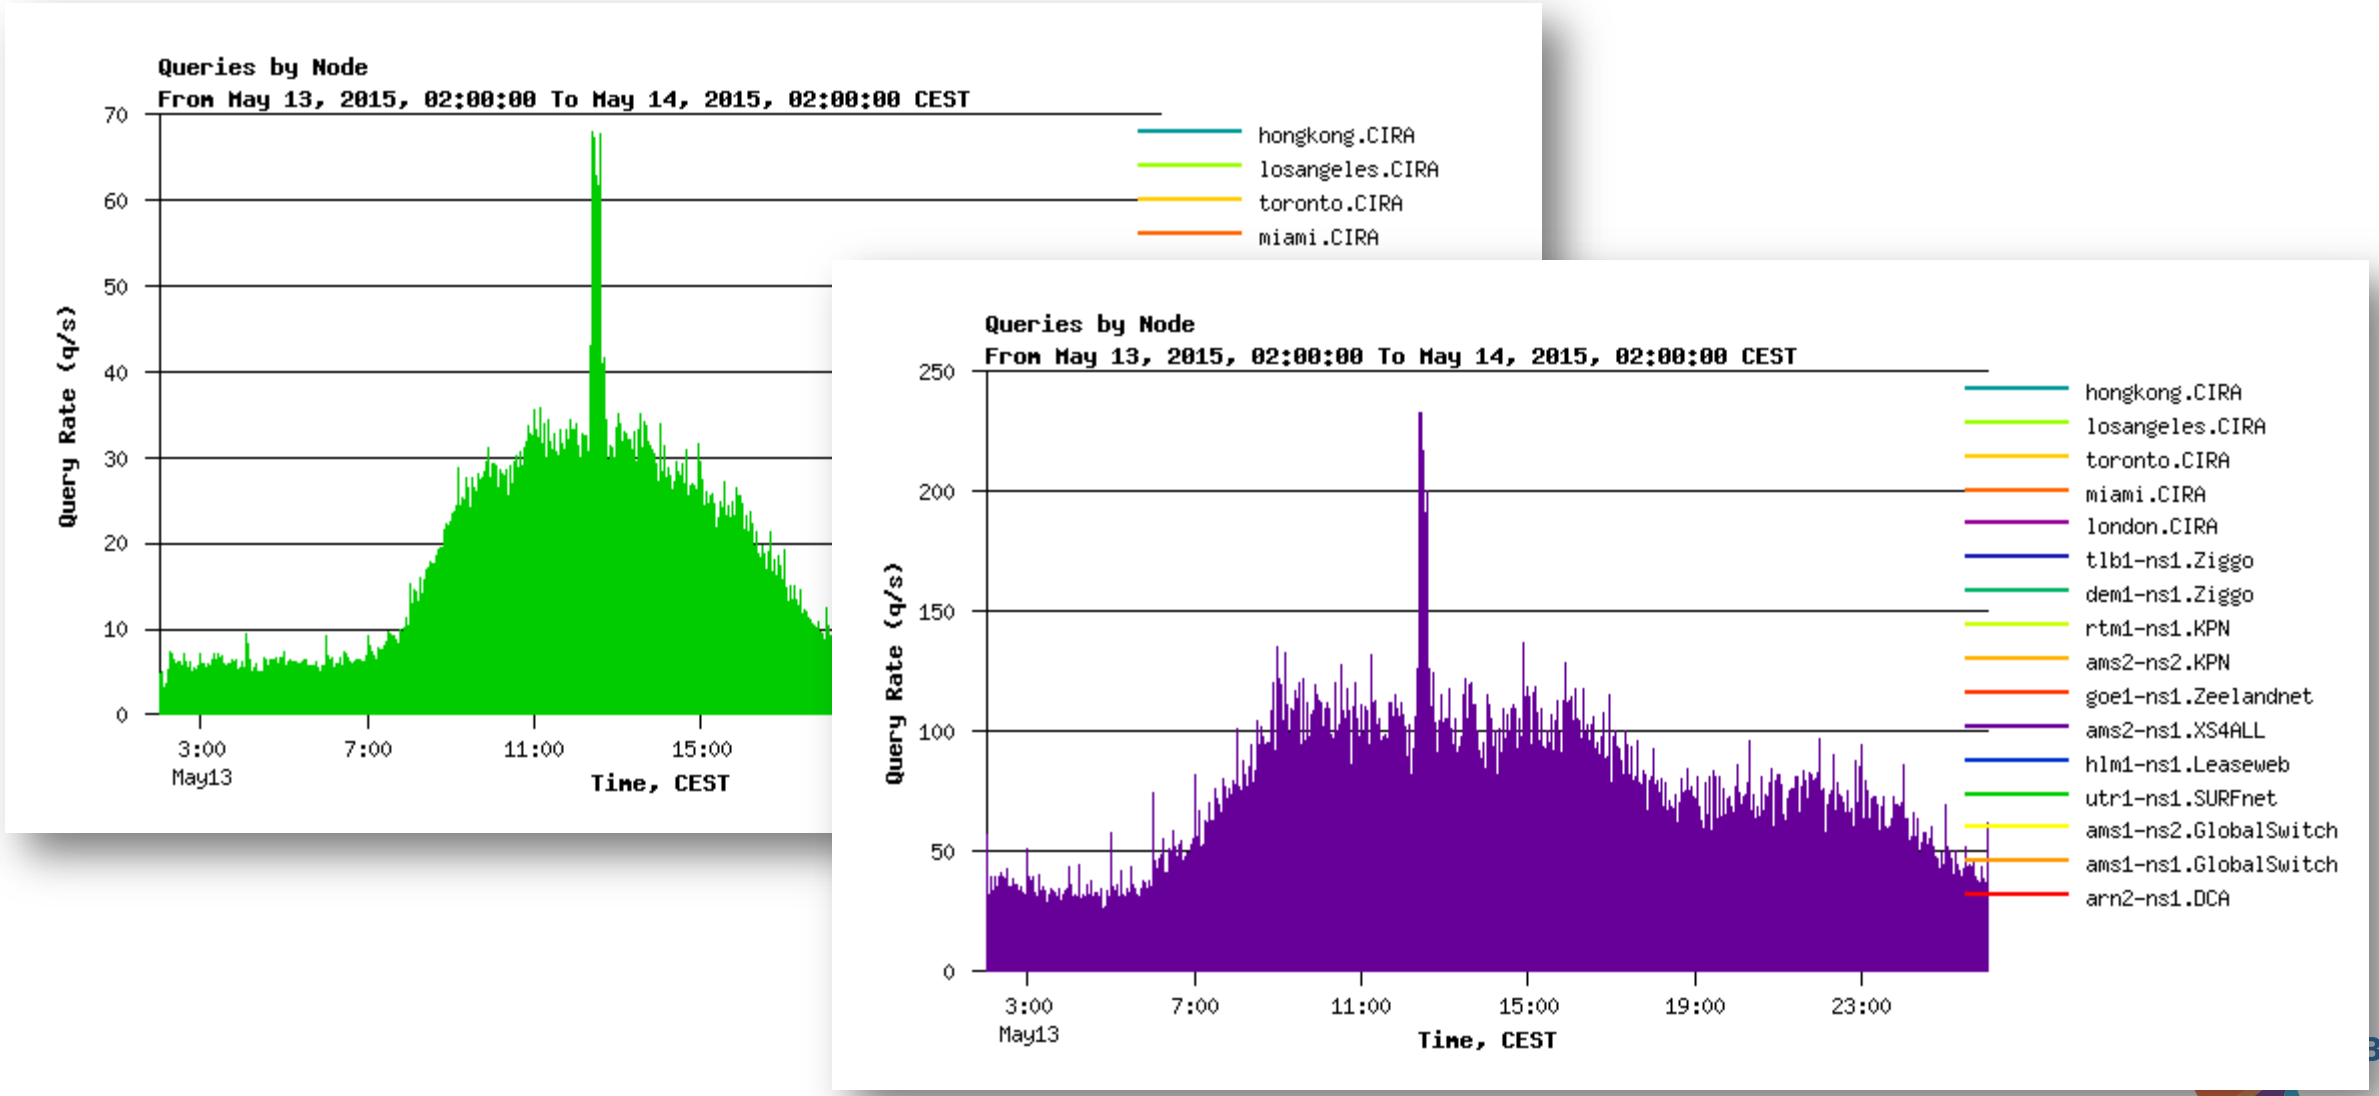

DNS local anycast – conclusion

It actually works!

- For multiple TLD's
  - .amsterdam, .aw, .politie

DNS local anycast – ‘business model’

For mutual benefit:

- ISP / datacentre provides bandwidth, rack space, power and sometimes ‘remote hands’
- SIDN provides equipment, operations and the service

DNS local anycast – current situation

- Local presence at 8 sites at ISP's
- One shared node (will explain later)
- ~ >80% of Dutch consumers "covered"

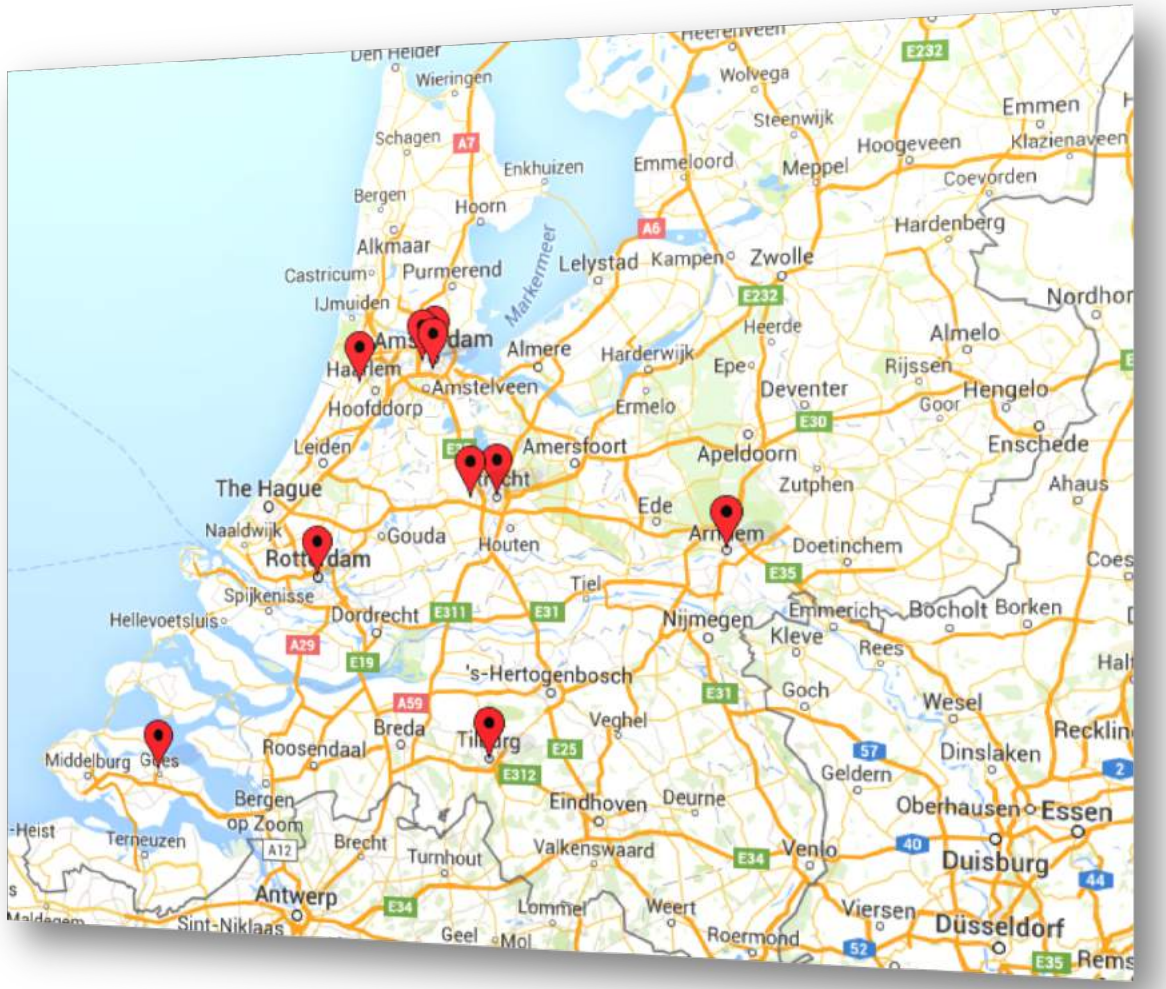

DNS local anycast - management

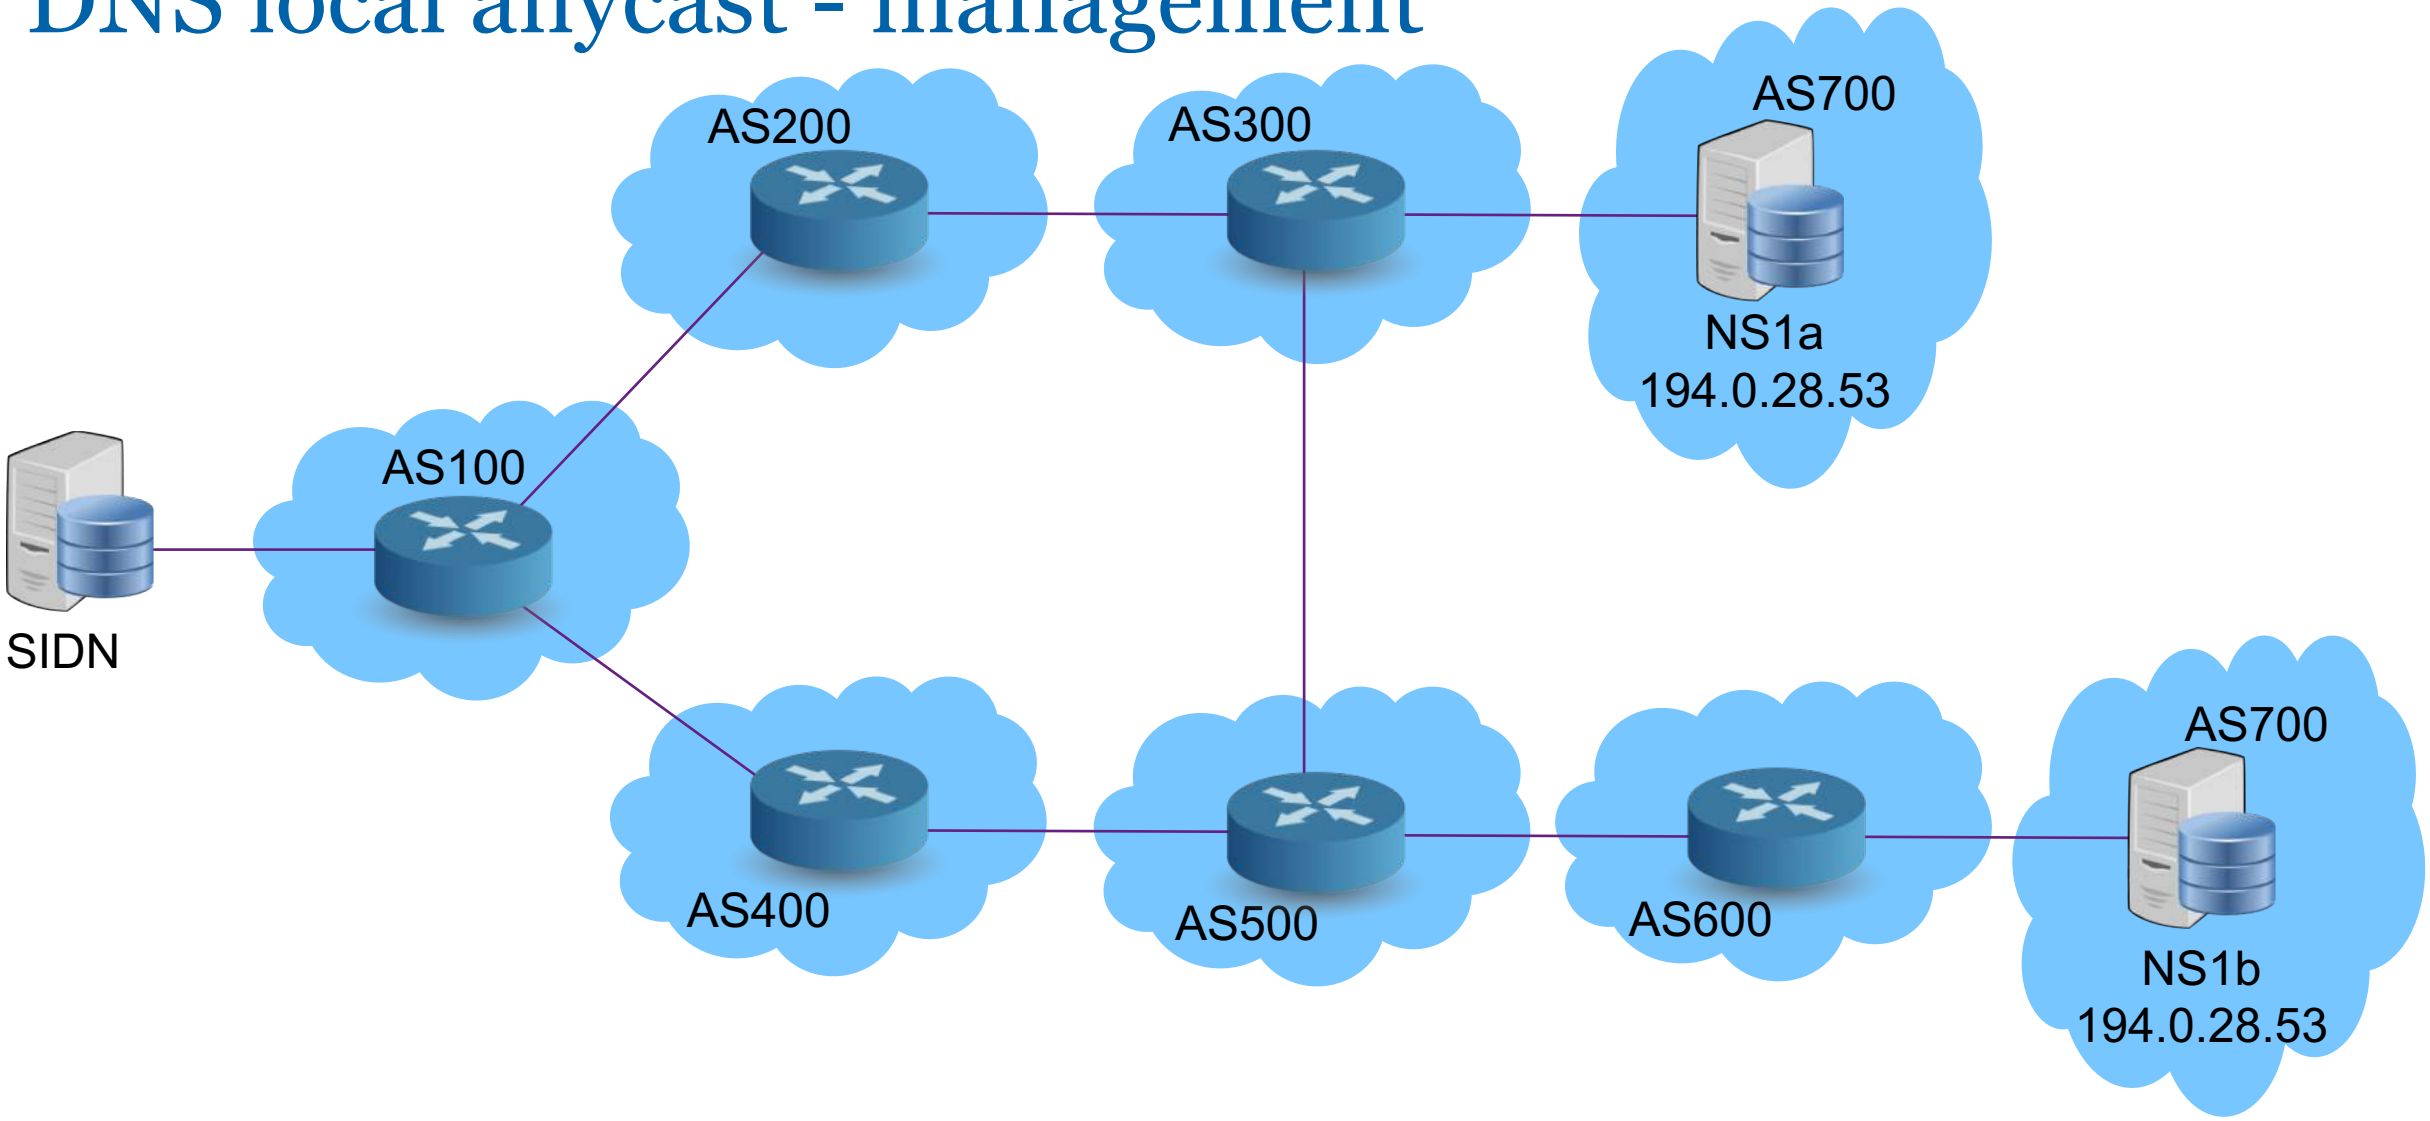

DNS local anycast - management

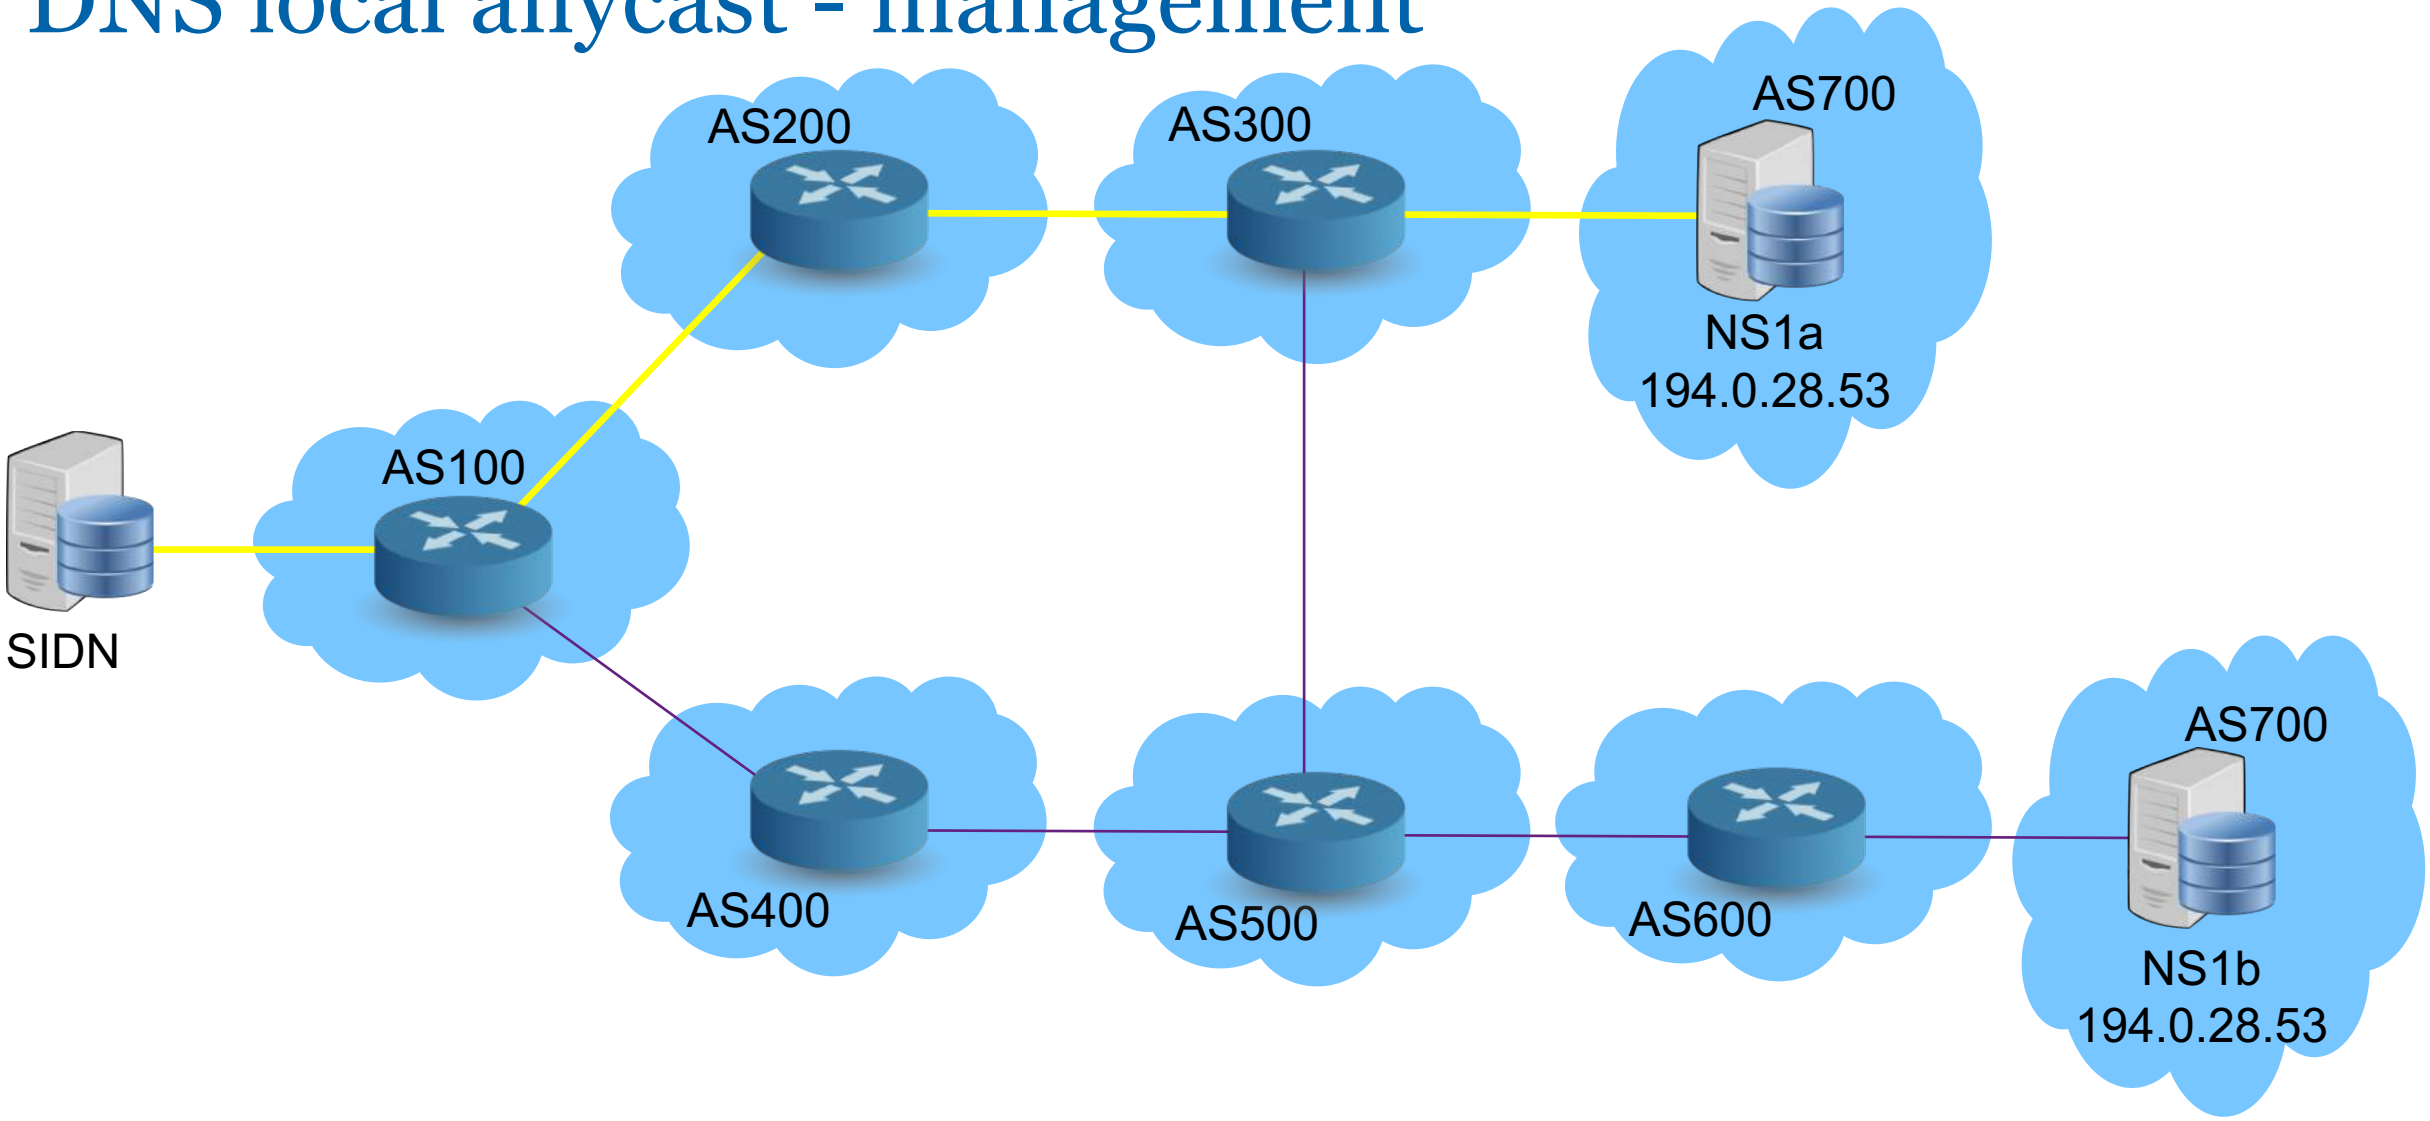

DNS local anycast – management (via VPN's)

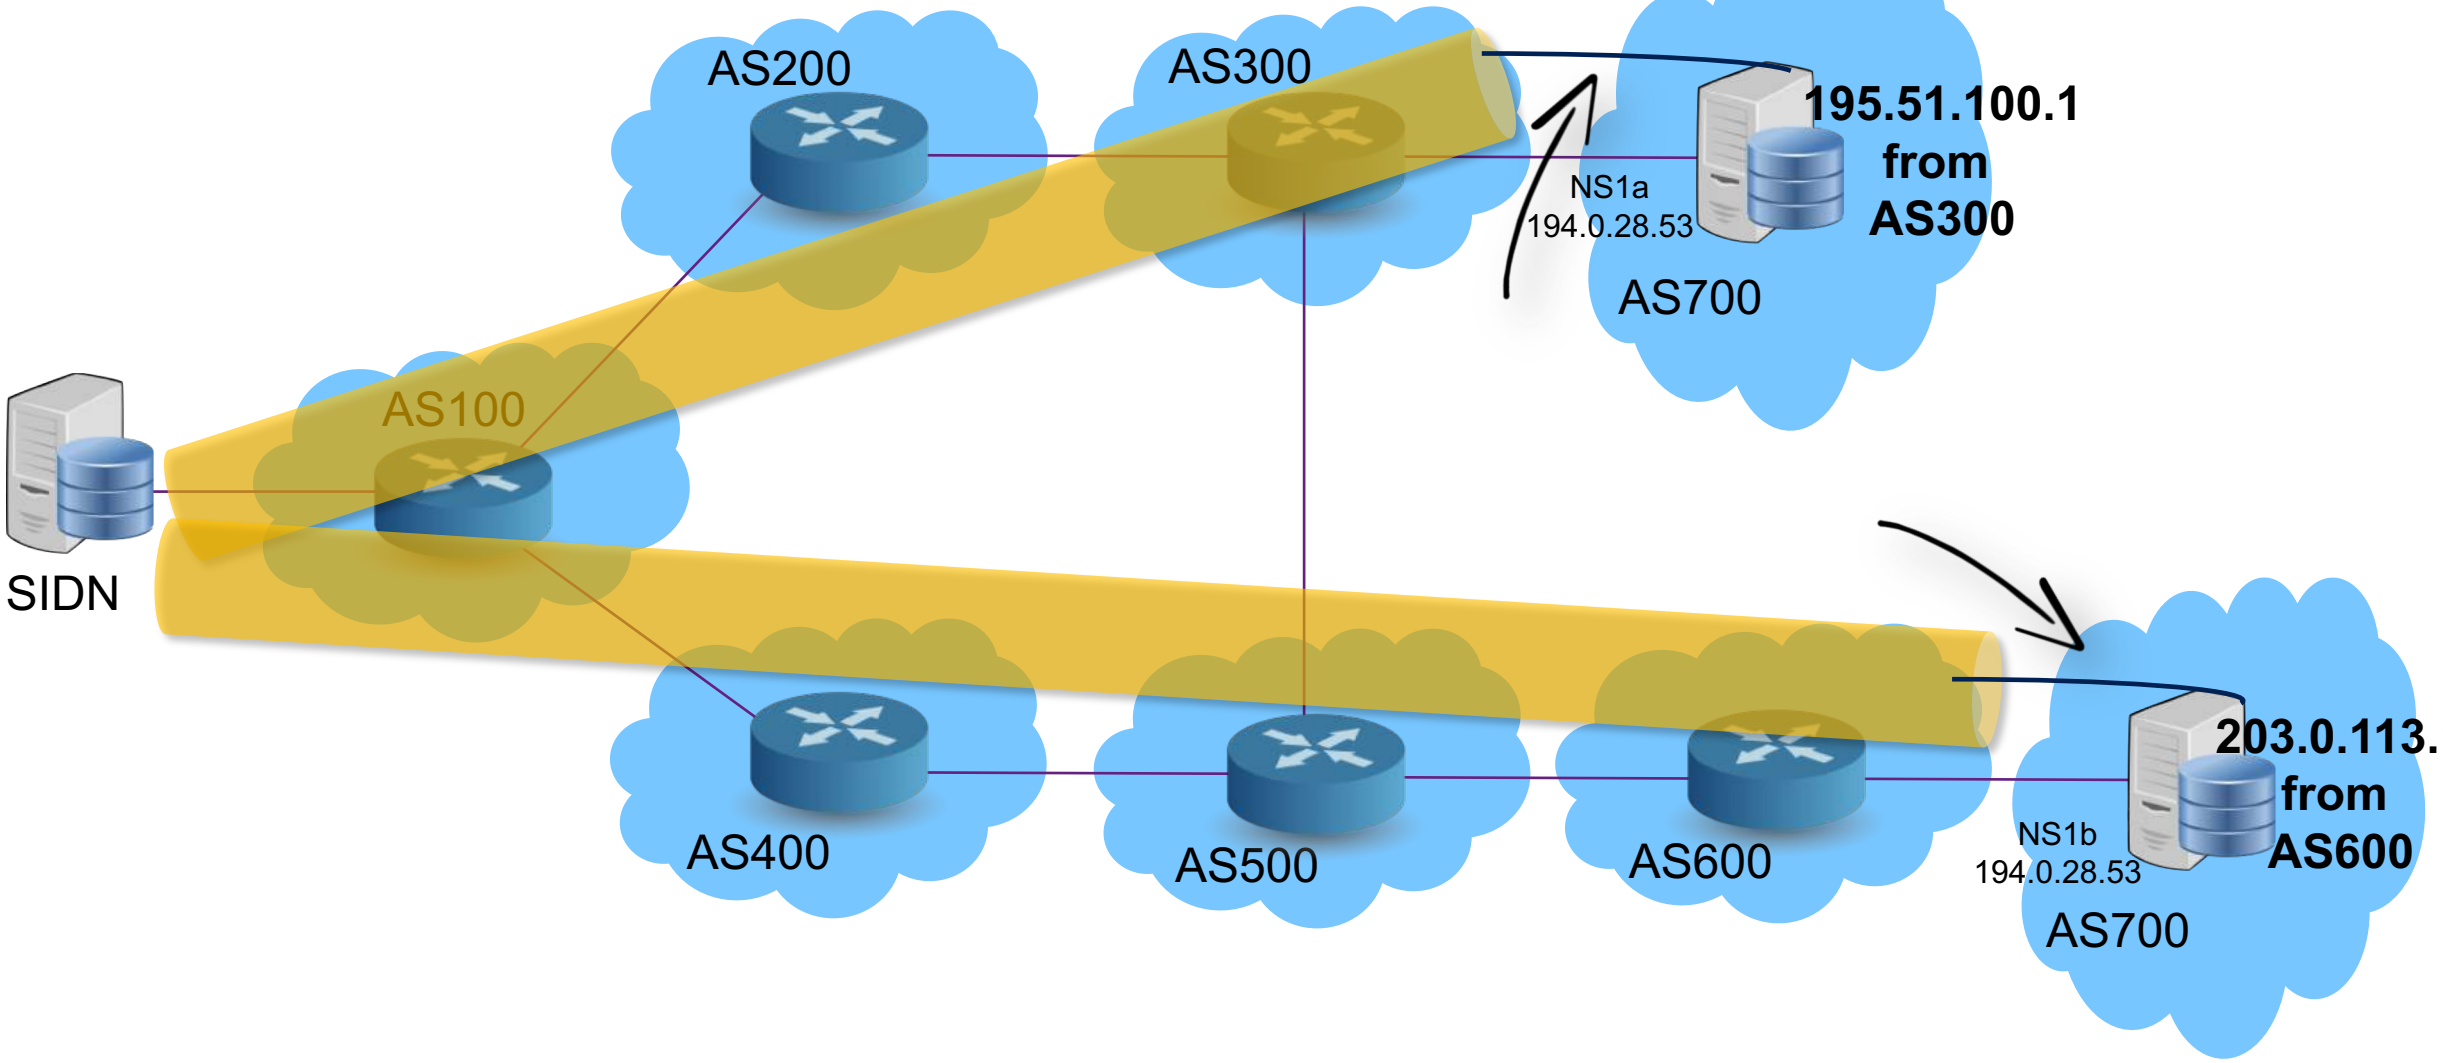

DNS local anycast – management (via VPN's)

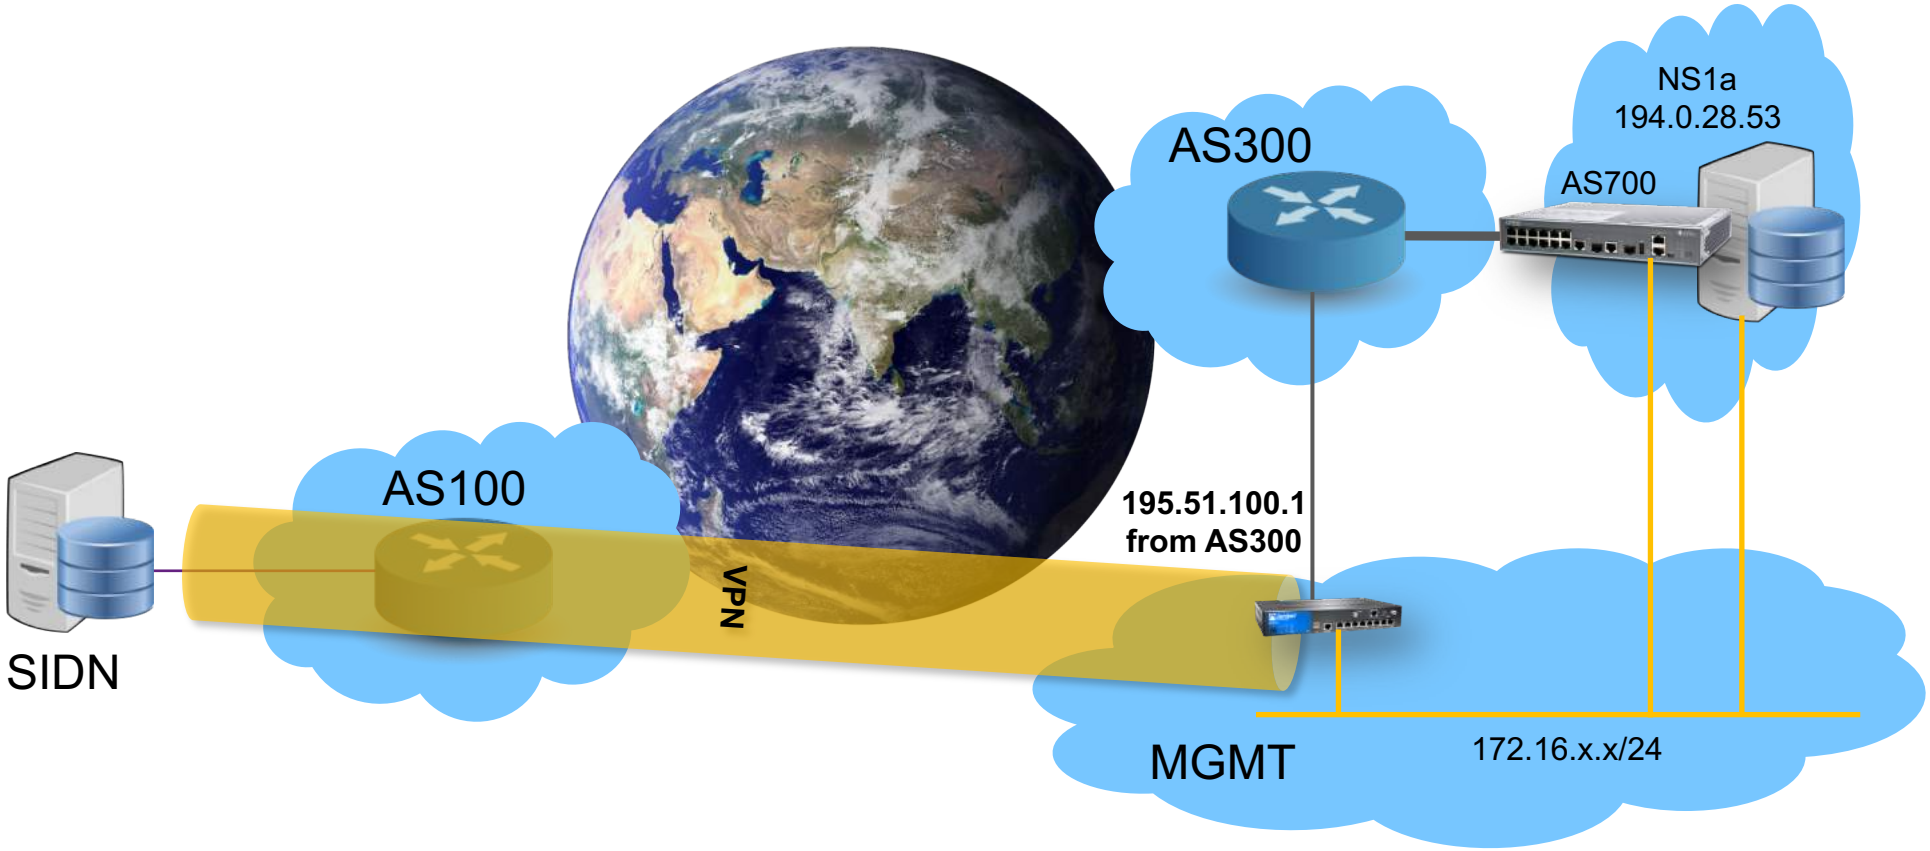

DNS local anycast – lessons learned

Setup is overdone.

- Dell server, 32 Gig RAM, 1U
- Fancy Juniper EX switch for BGP, 1U
- Separate Juniper SRX switch for VPN, 1U
- A bit too much for only 50 qps…

DNS local anycast – lessons learned

Also…

- 'Legal challenges'
- 'Persuasion challenges', or getting in touch with the right people
- 'Not-in-scope challenges' (they want us, we don't really want them)
- It's quite a bit of work to setup and maintain
- Monitoring requires special attention
- So does tuning and tweaking
  - Like making sure partners keep it local and don't export the route

SIDN LABS

DNS local anycast – lessons learned

Also…

- 'Legal' challenges
- Persuasion challenges, or getting in touch with the right people
- Not in scope challenges (they want us, we don't really want them)

So we made a 'shared' local anycast node

SIDN LABS

DNS local anycast – Shared node

- Not with an ISP, but located at an IX
- No exclusive use by one party,
- but used by several, carefully selected peers
  - We may cut them off if they cause too much problems for others
- Works well for smaller parties, or for the other mentioned challenges

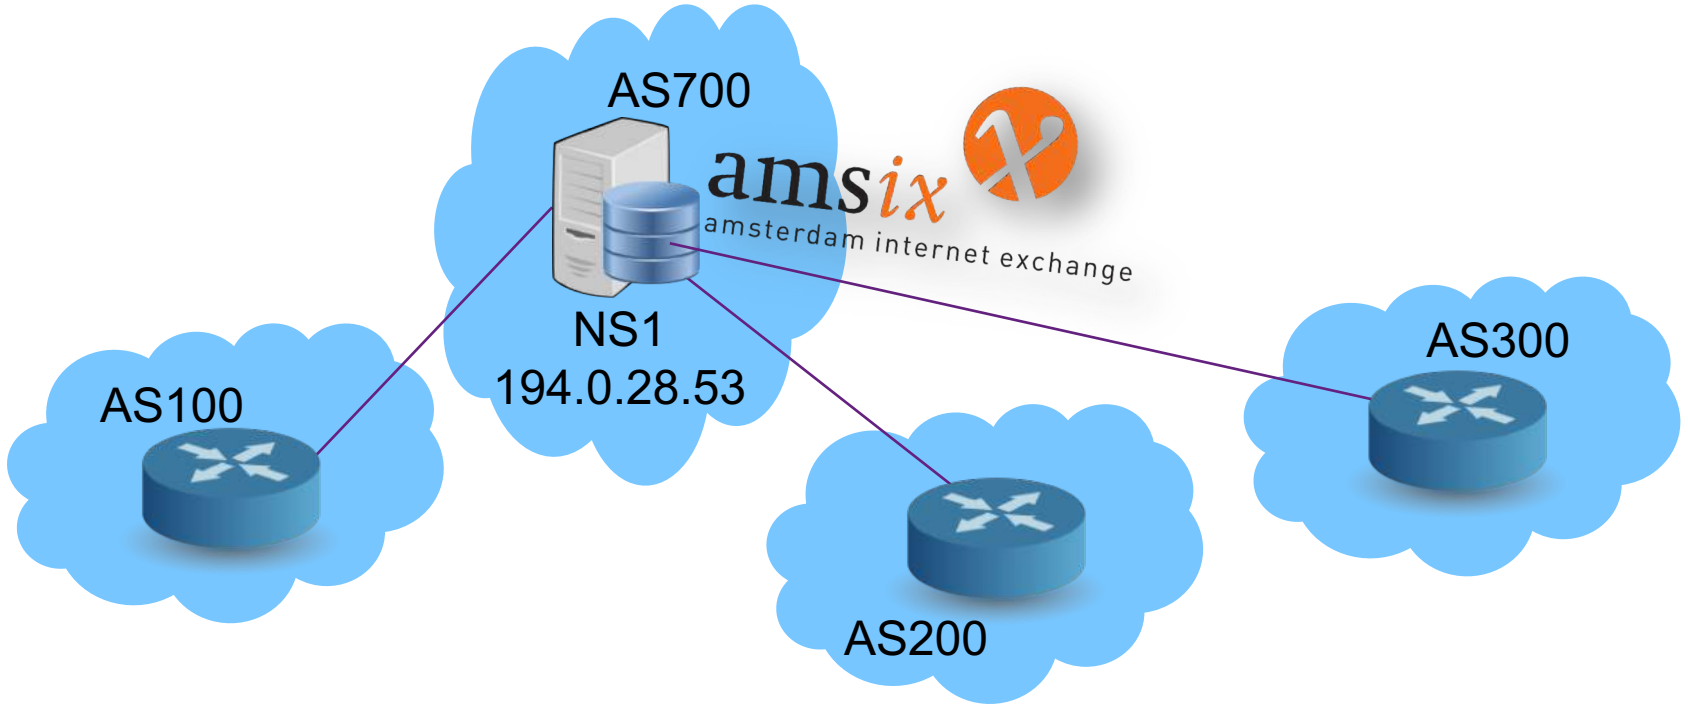

DNS local anycast – Future work

- Maybe a simple front-end and (hidden) big back-end?
  - DNSdist or CoreDNS forward/proxy and cache plugin maybe?

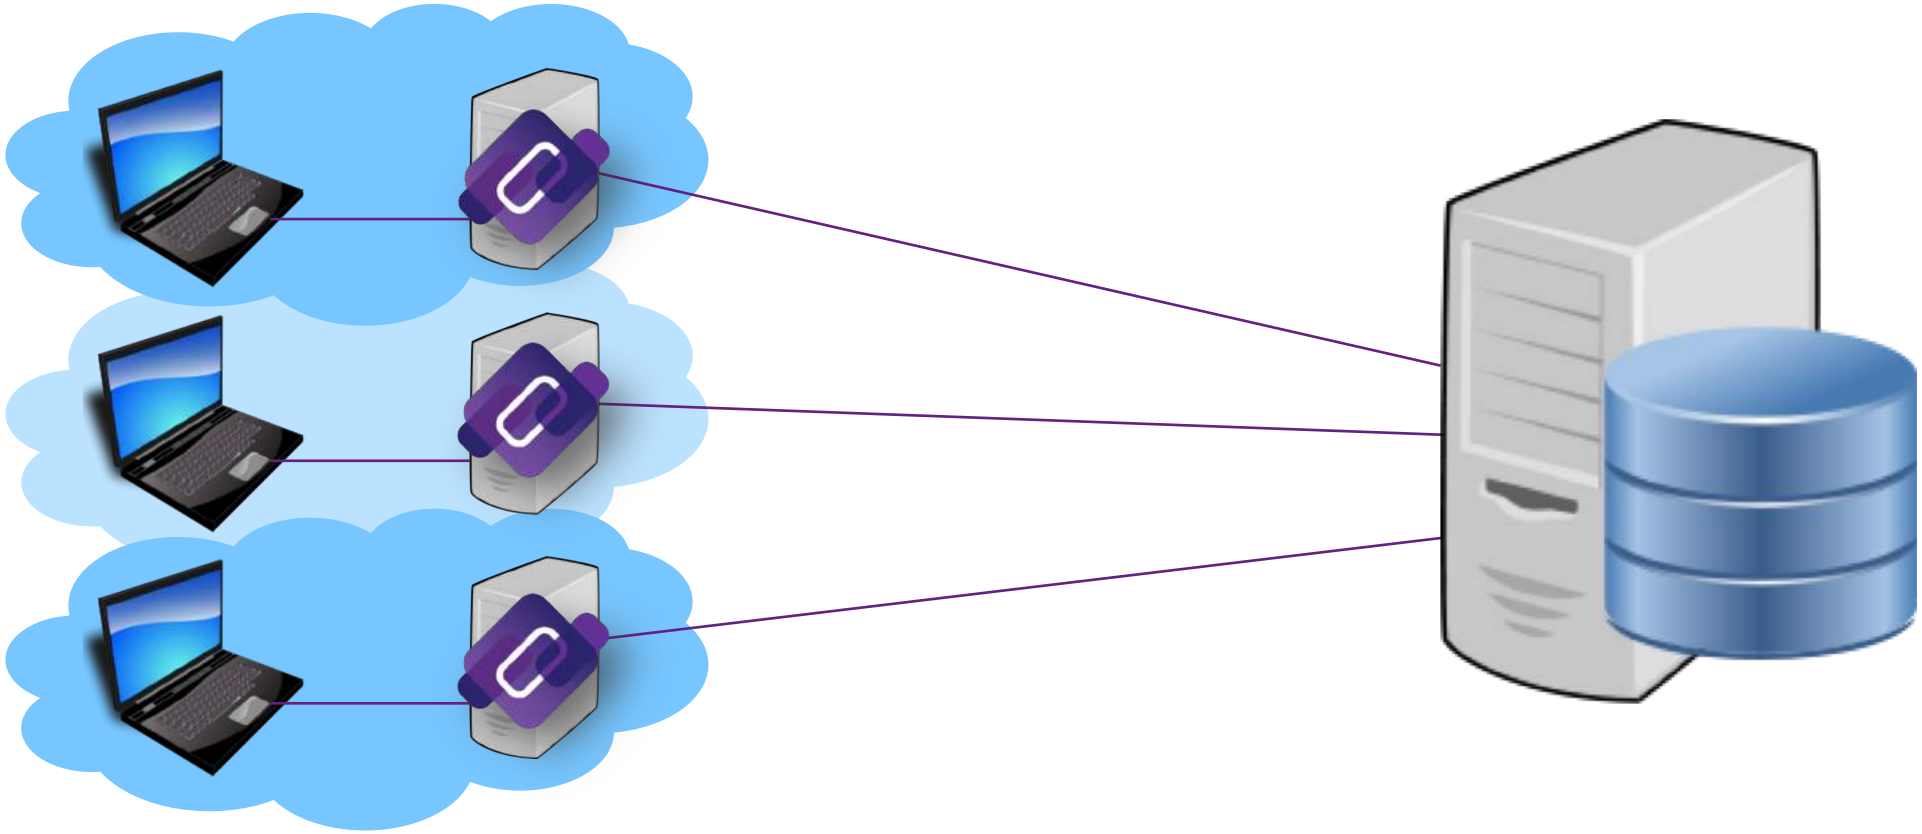

DNS local anycast – Future work

- **Anycast-in-a-box**
  - Single server with BGP (BIRD), DNS (i.e. BIND), VPN (i.e. FreeS/WAN)
  - Can be virtualized (including a Juniper vMX for instance)

SIDN Labs website

https://www.sidnlabs.nl/

Search for ‘anycast’:

Thank You!

SDN LABS