

# DNS Out There ...

Lars-Johan Liman, M.Sc.  
Sr. Systems Specialist  
Netnod Internet Exchange



©2011 Lars-Johan Liman 201109-ipop-DNS-liman KTH 2011-09-27

www.netnod.se

# Bear in Mind ...



# One DB to Rule Them All ...

- DNS in a database.
- DNS is ***one*** database ...
  - ▶ ... that is distributed over many servers,
  - ▶ ... that can contain a variety of information,
  - ▶ ... that is hierarchical, but
- It is ***ONE*** database.
- The full service resolver (local nameserver) hides the complexity of the server tree from the stub resolver (end client).



# Search vs. Lookup

- A *search* is a step by step process to "home in" on the exact information:
- Example, phonebook:
  - ▶ Persons in Stockholm with last name "Liman"
    - 30 people.
  - ▶ ... that hold an M.Sc. degree ...
    - 2 people
  - ▶ ... whose first name is "Lars" ...
    - **FOUND!**



# Search vs. Lookup

- A *lookup* uses a unique handle as a direct link to the correct information.
- Example, tax authorities:
  - ▶ Address information for 640825-xxxx.
    - FOUND!
- **DNS only does lookups!**
  - ▶ The unique handle is the triple  
Domain name, Class, Record Type  
`www.kth.se. IN A`

# DNS – Fair and Equal

- ***ALL*** domain names are treated equal!
  - ▶ There is ***NO*** structural difference between
    - arpa.
    - lab.csc.kth.se.
    - www.volvo.se.
    - 10.23.19.8.in-addr.arpa.
    - 10.ugly.23.hack.55.to.33.confuse.12.students.se.
  - (or even)
    - .



# Trick Questions

- Can you have an A record for SE.?

www.netnod.se



netnod

©2011 Lars-Johan Liman 201109-ipop-DNS-liman KTH 2011-09-27

7

# Trick Questions

- Can you have an A record for SE.?
  - ▶ YES!

se.            3600            IN            A            10.23.19.2

www.netnod.se



netnod

©2011 Lars-Johan Liman 201109-ipop-DNS-liman KTH 2011-09-27

8

# Trick Questions

- Can you have an A record for SE.?
  - ▶ YES!

se. 3600 IN A 10.23.19.2

- Can you have a PTR record for kth.se?

# Trick Questions

- Can you have an A record for SE.?
  - ▶ YES!

se. 3600 IN A 10.23.19.2

- Can you have a PTR record for kth.se?
  - ▶ YES!

kth.se. 600 IN PTR server14.kth.se.



netnod

# Trick Questions

- Can you have an A record for SE.?
  - ▶ YES!

se. 3600 IN A 10.23.19.2

- Can you have a PTR record for kth.se?
  - ▶ YES!
- Can you have an MX record for "."?

kth.se. 600 IN PTR server14.kth.se.

# Trick Questions

- Can you have an A record for SE.?
  - ▶ YES!

se. 3600 IN A 10.23.19.2

- Can you have a PTR record for kth.se?

- ▶ YES!

kth.se. 600 IN PTR server14.kth.se.

- Can you have an MX record for "."?

- ▶ YES!

. 86400 IN MX 10 smtp.cafax.se.



# Categories of Record Types

- The DNS uses its own records to maintain the structure of the database itself.
- Two categories of records:
  - ▶ **Data carrying records:** Used by applications to reach various resources.  
A, AAAA, MX, PTR, SRV, NAPTR, ...
  - ▶ **Structure carrying records:** Only used by the DNS system itself to locate the sought data records.  
SOA, NS, CNAME, DNAME, ...

# Cacheing is Vital

- The cacheing mechanism in the full service resolver helps off-load traffic from the root name servers.
  - ▶ The higher up, the more traffic (almost ...).
- Without cacheing, every query would have to start at the root.
  - ▶ Not A Good Idea™
  - ▶ The root servers receive enough junk as it is ...



netnod

# Speaking of Root Name Servers

www.netnod.se



# Root Name Servers

- So what's the deal?
  - ▶ What, why, who, where, when?
- 13 root name servers on the Internet:
  - ▶ X.root-servers.net. (where  $X \in \{ a, b, \dots, m \}$ )
- Operated by 12 organisations.
  - ▶ i.root-servers.net is operated by Netnod.
  - ▶ Verisign operates 2 letters: a and j.
- Overview:
  - ▶ <http://www.root-servers.ORG/>

# Root Server System Expansion

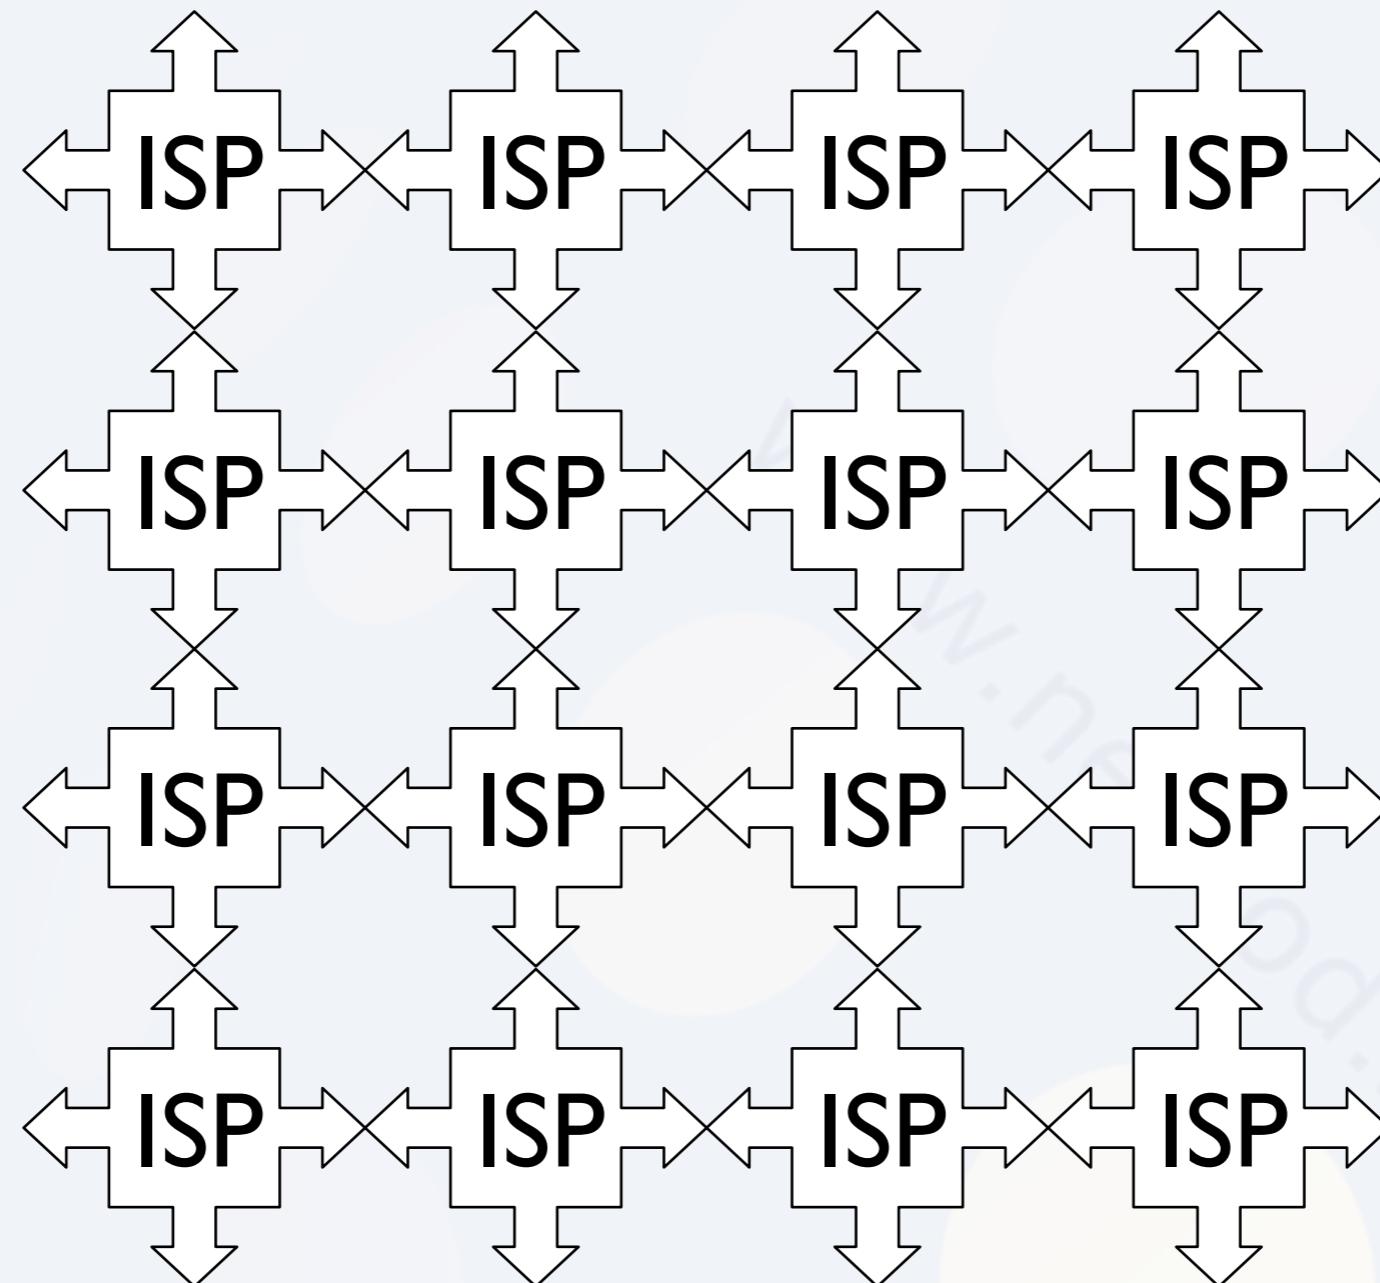
- 10 years ago it was 13 machines.
  - ▶ 10 x USA + Sweden + UK + Japan.
- Big nations started to complain.
  - ▶ China, India, France, Germany, ...
- Technical limitation on the number of servers!
  - ▶ What to do?
- Ugly hack: ANYCAST!
  - ▶ Now 255 sites in total!

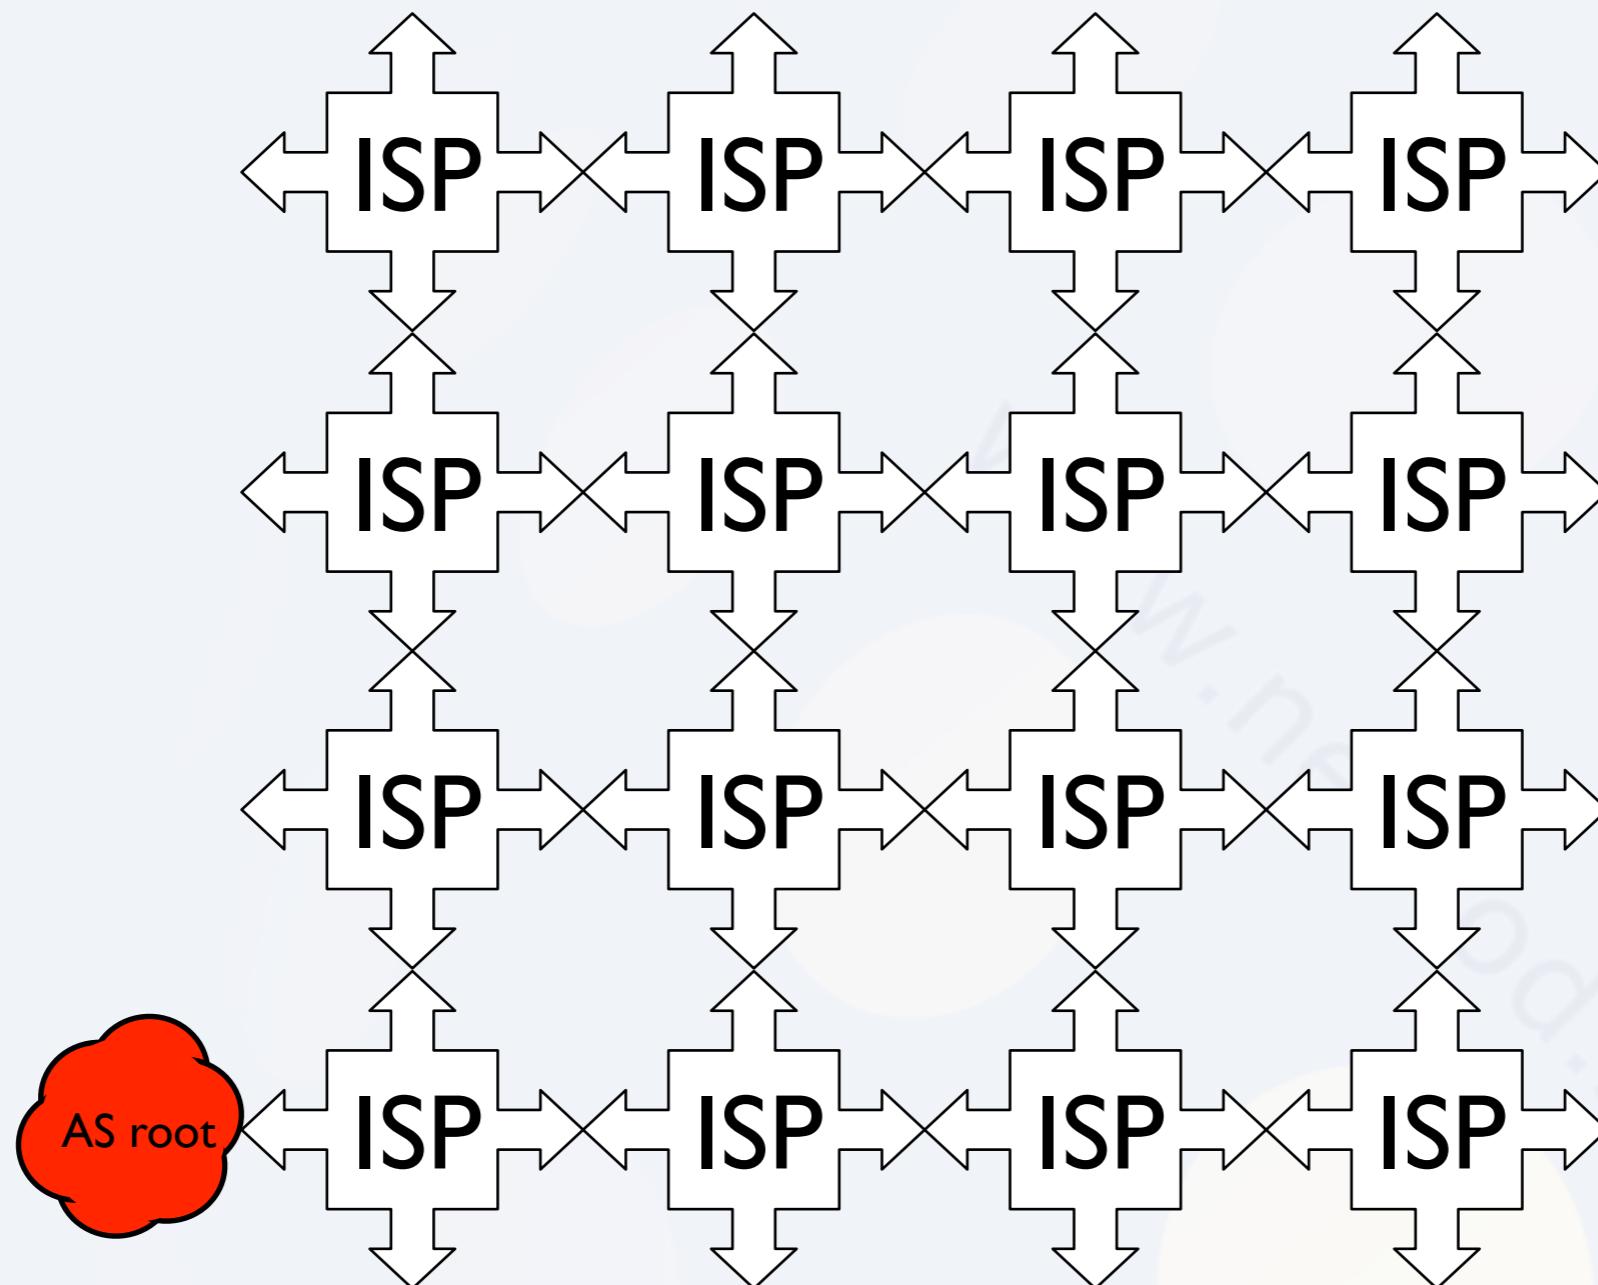


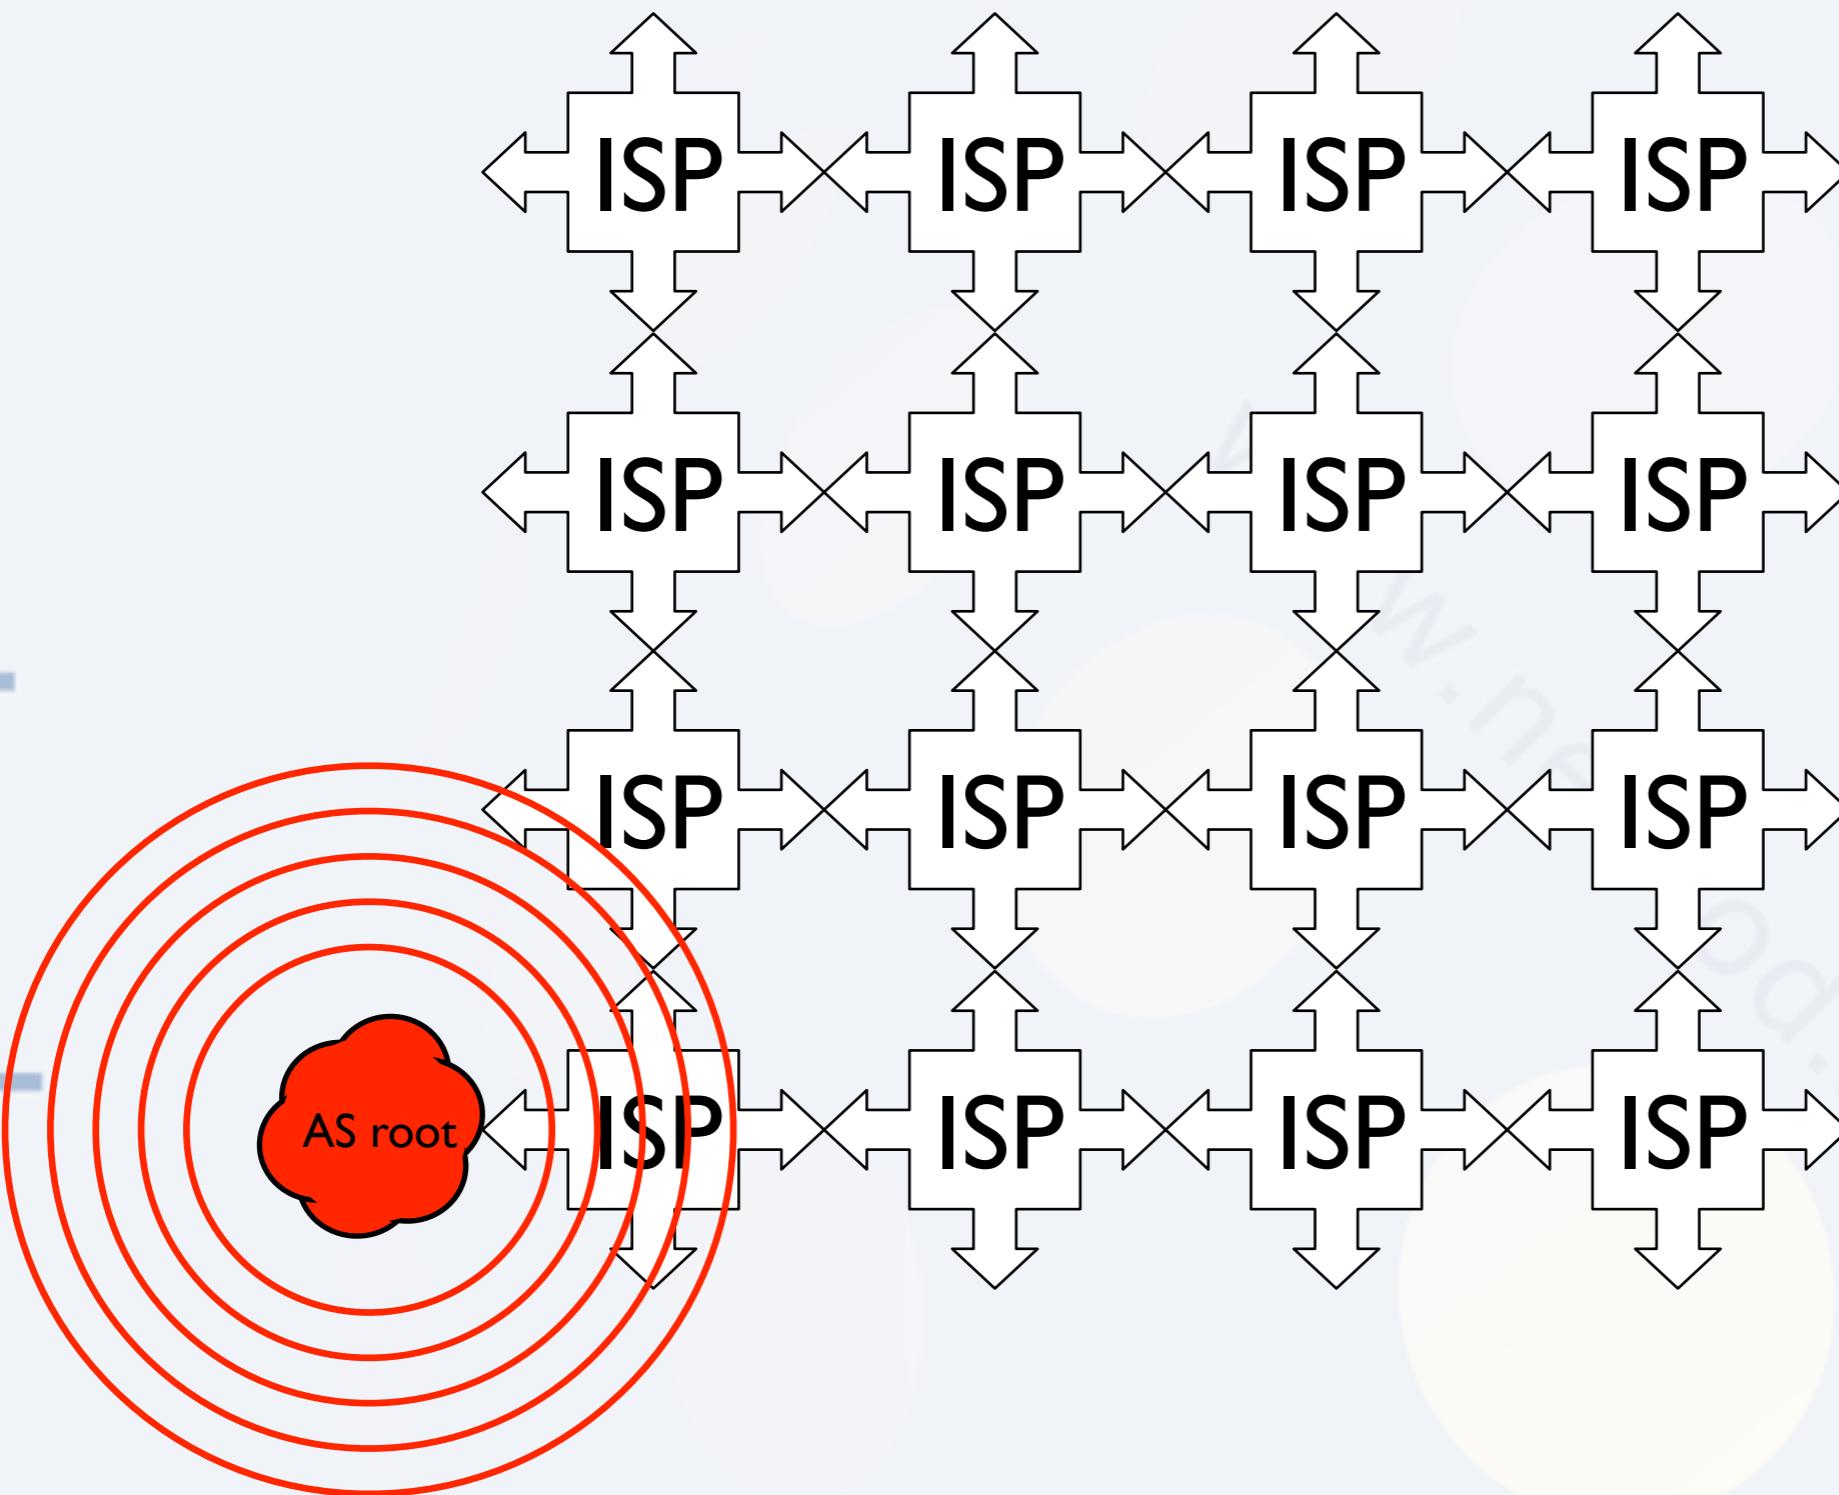
www.netnod.se

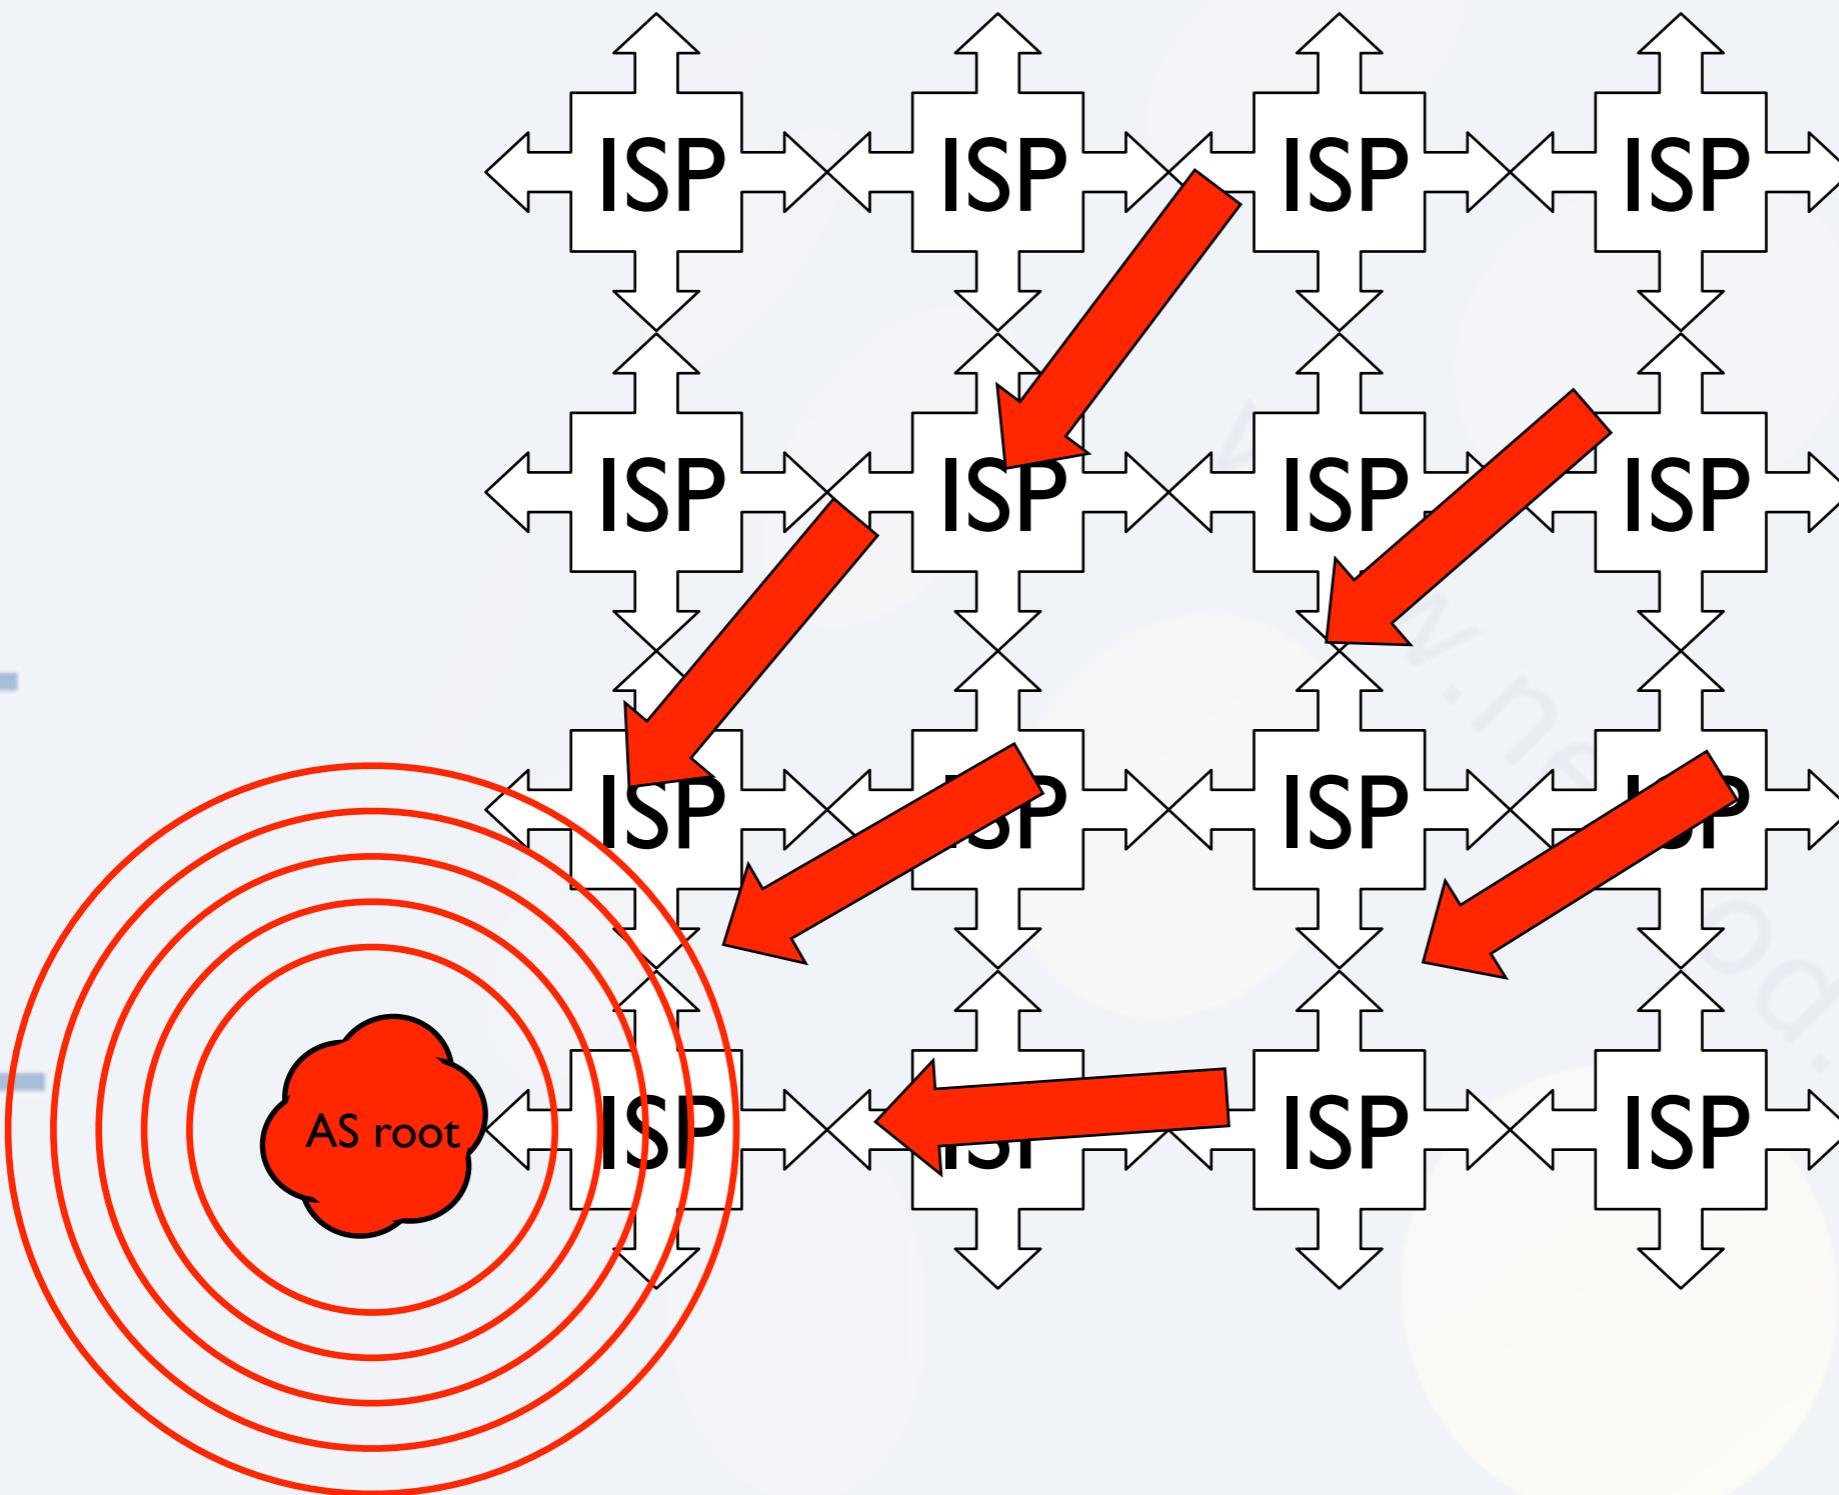
# Anycast

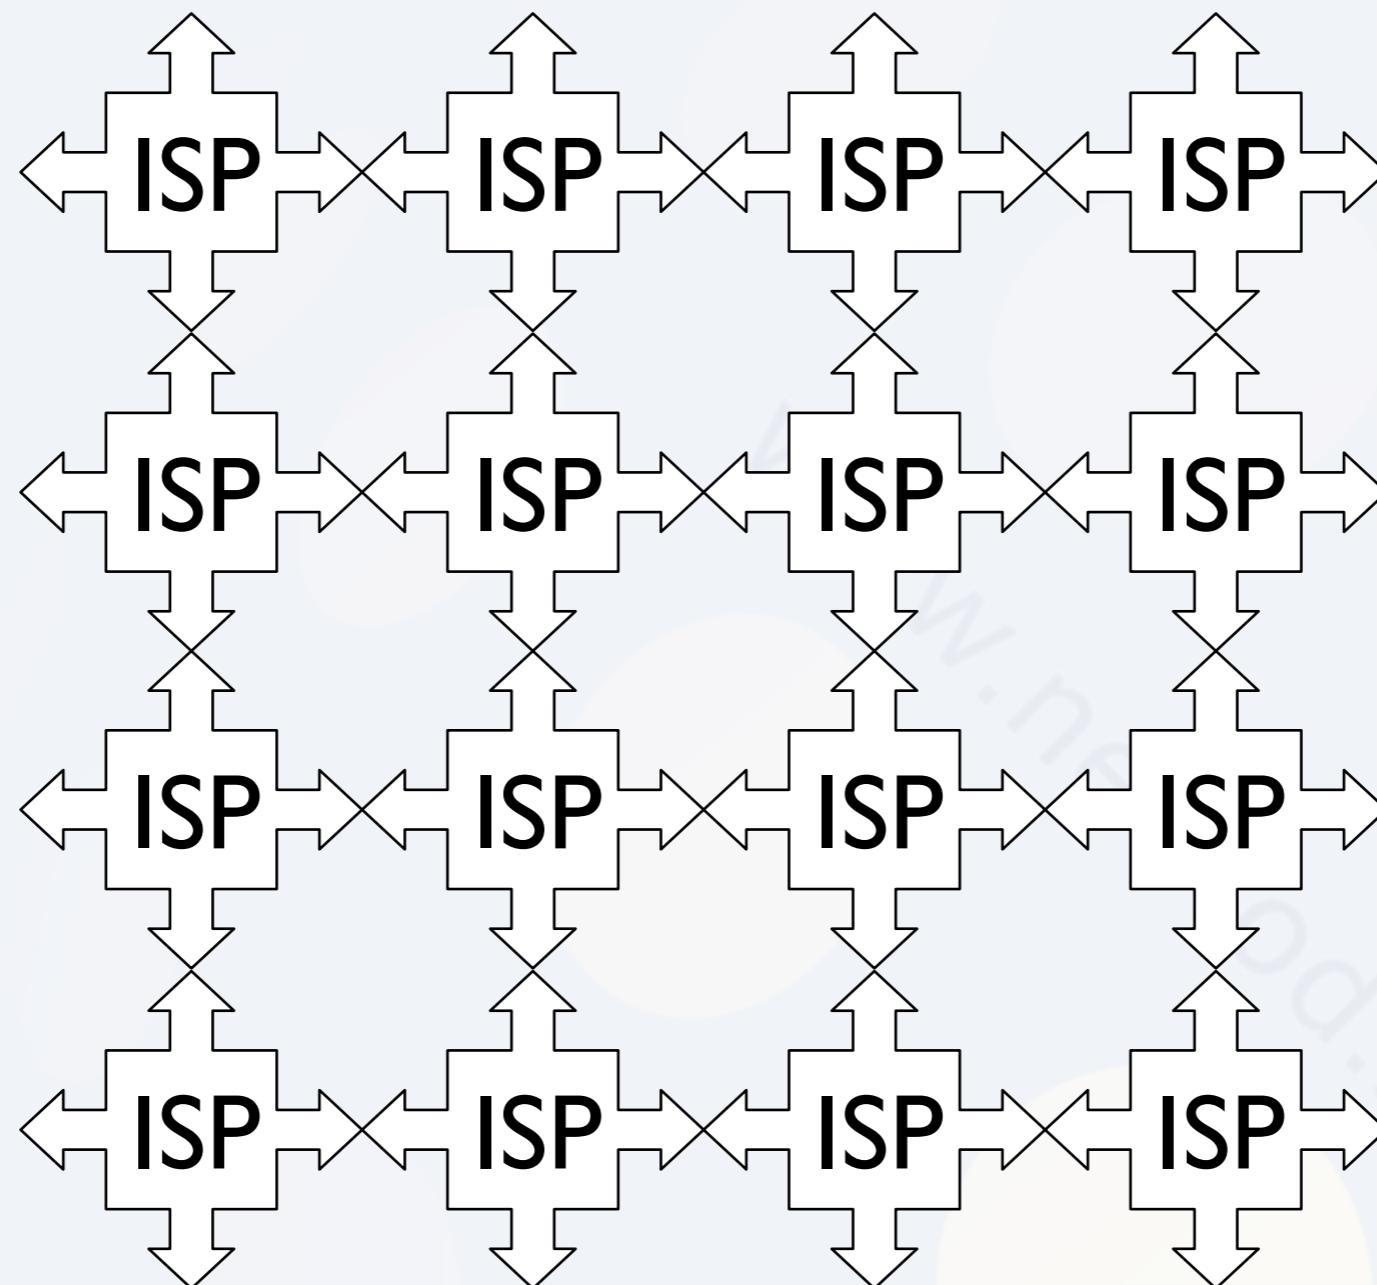


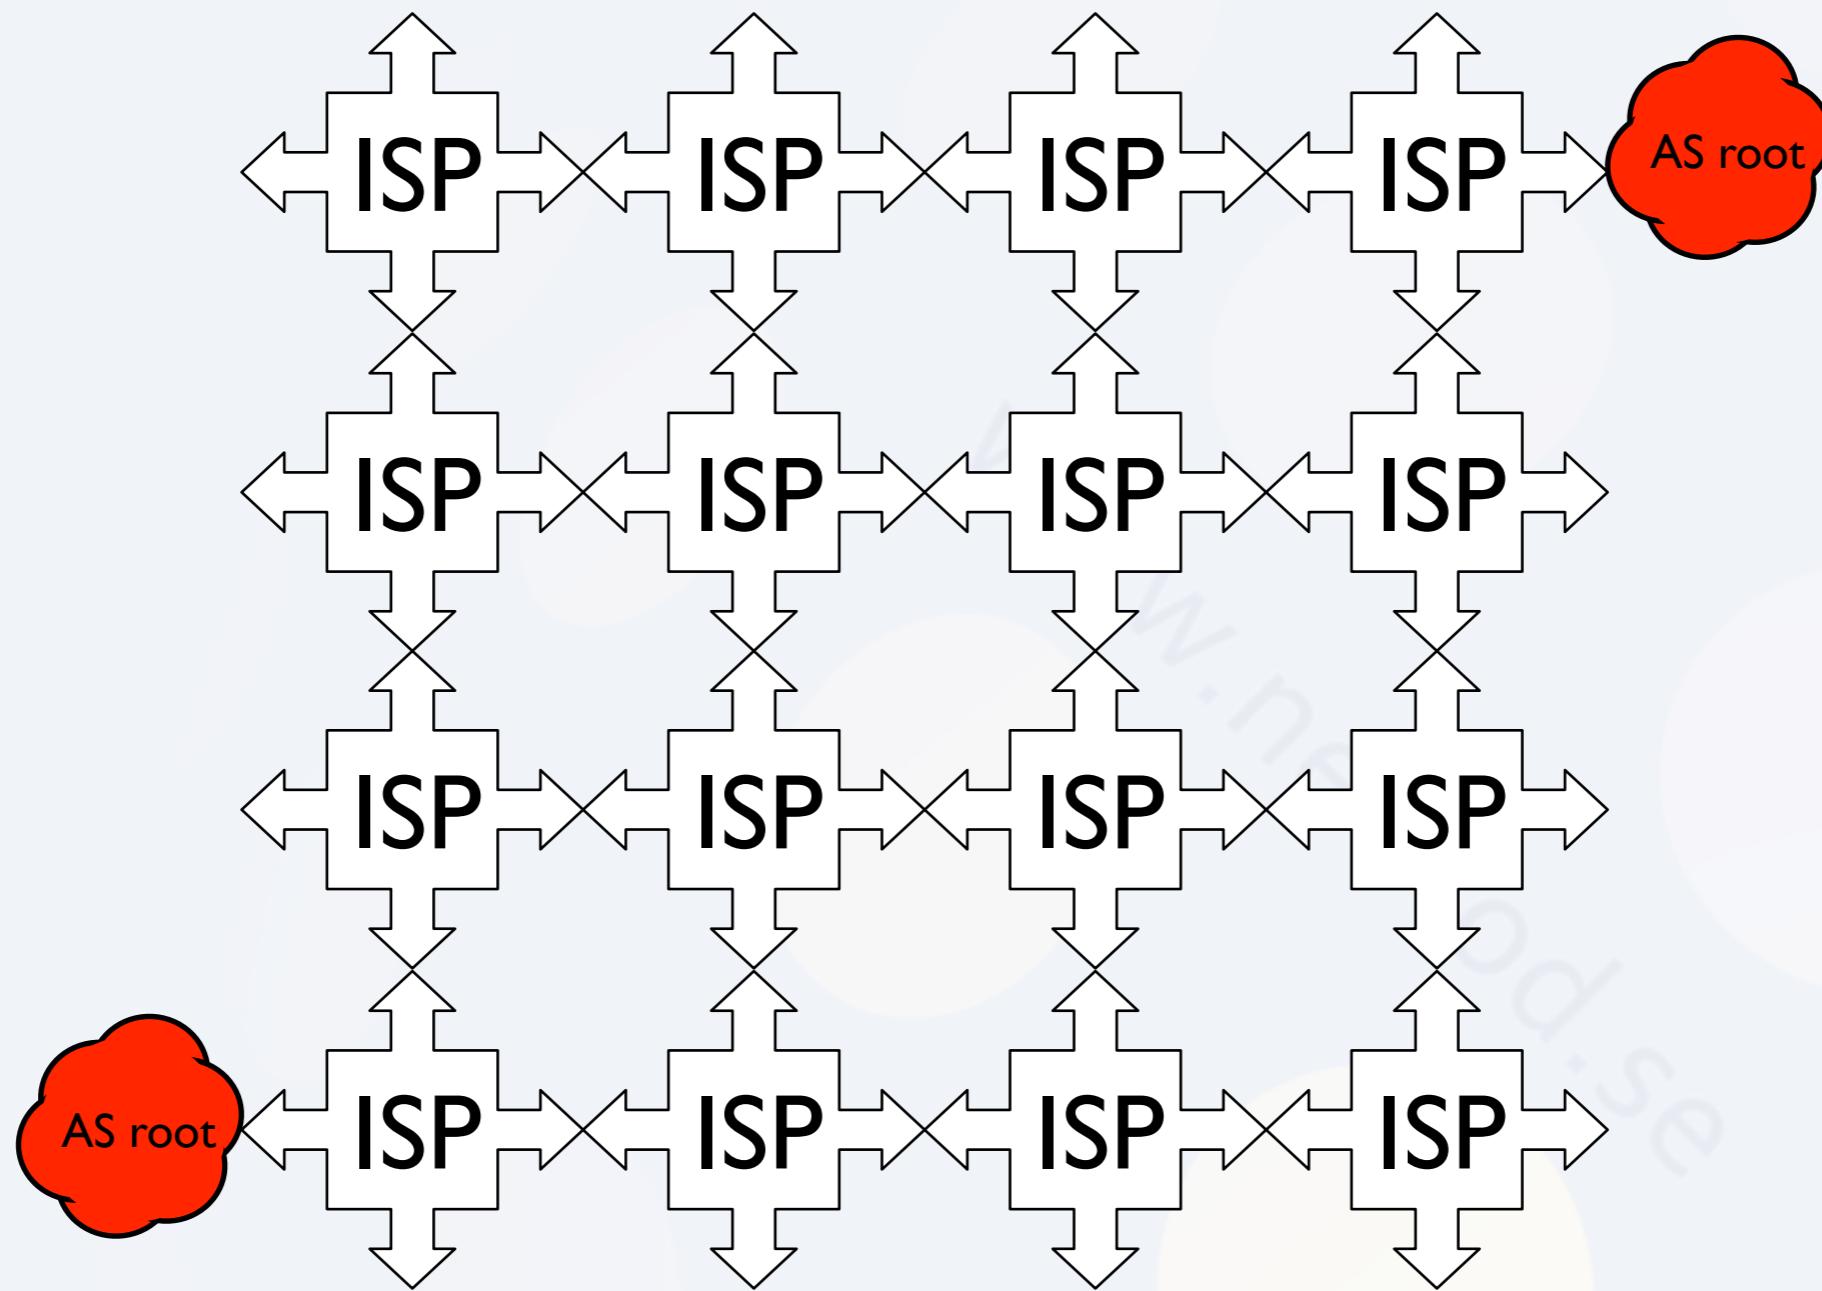


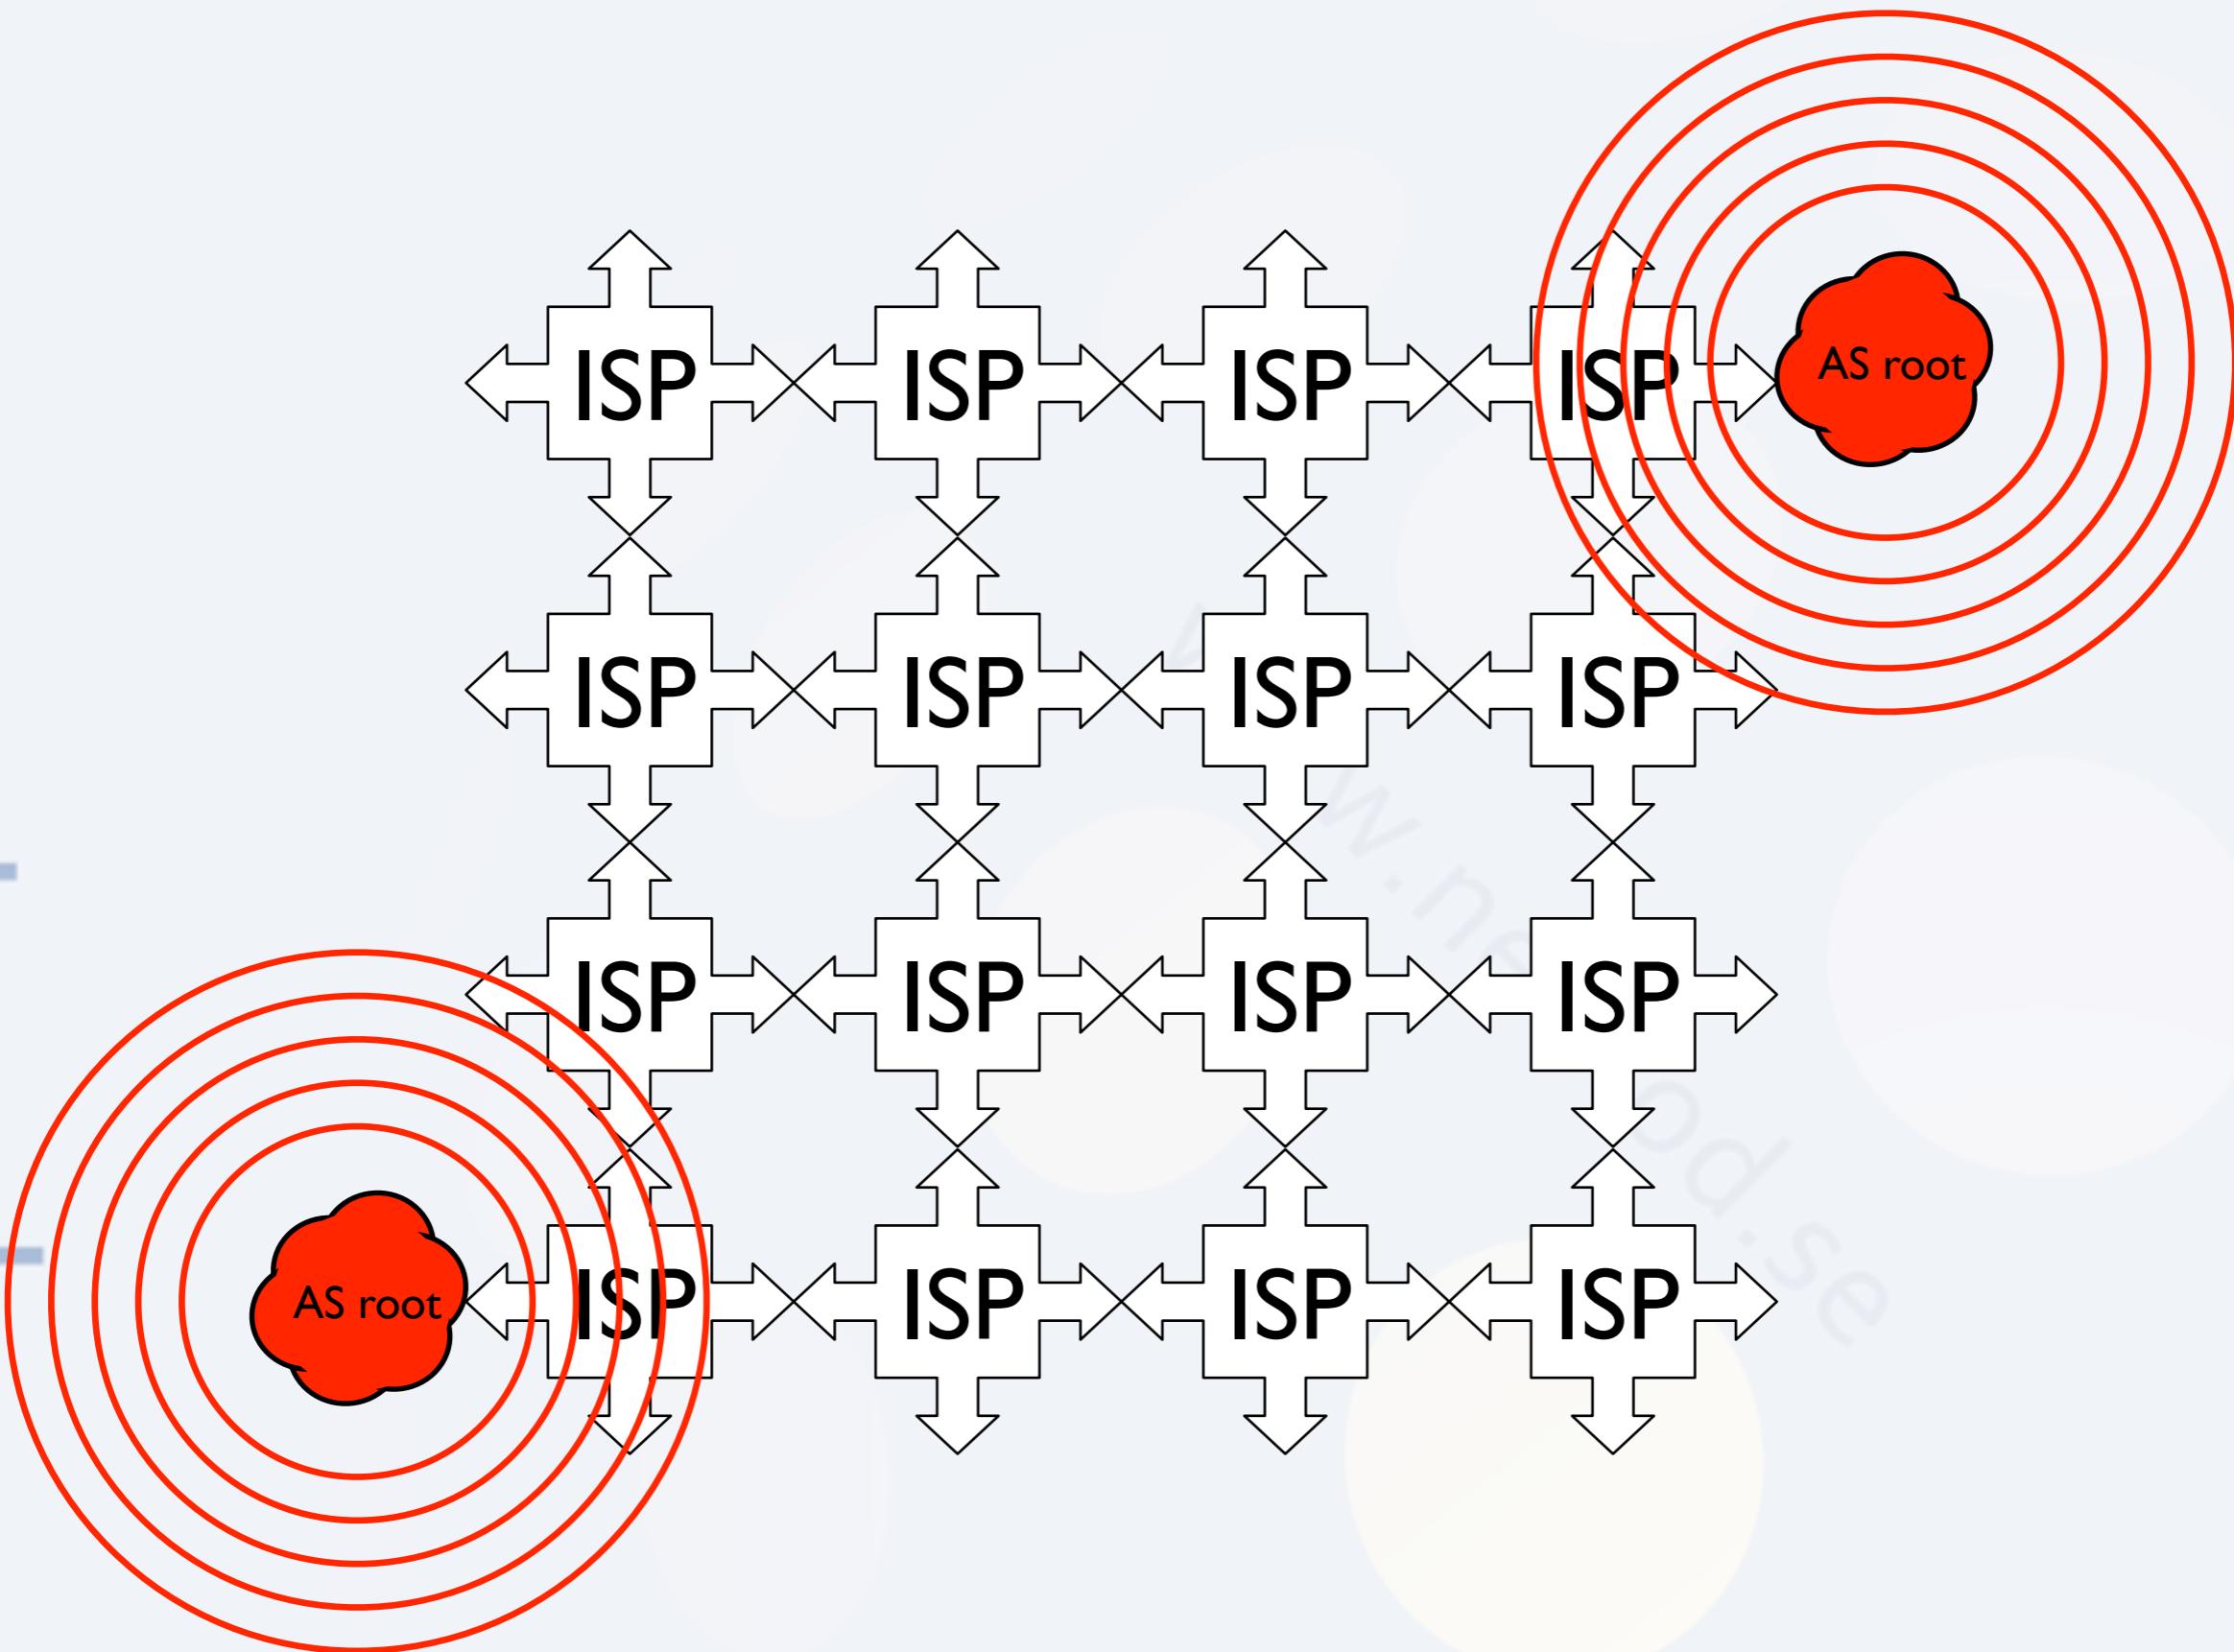


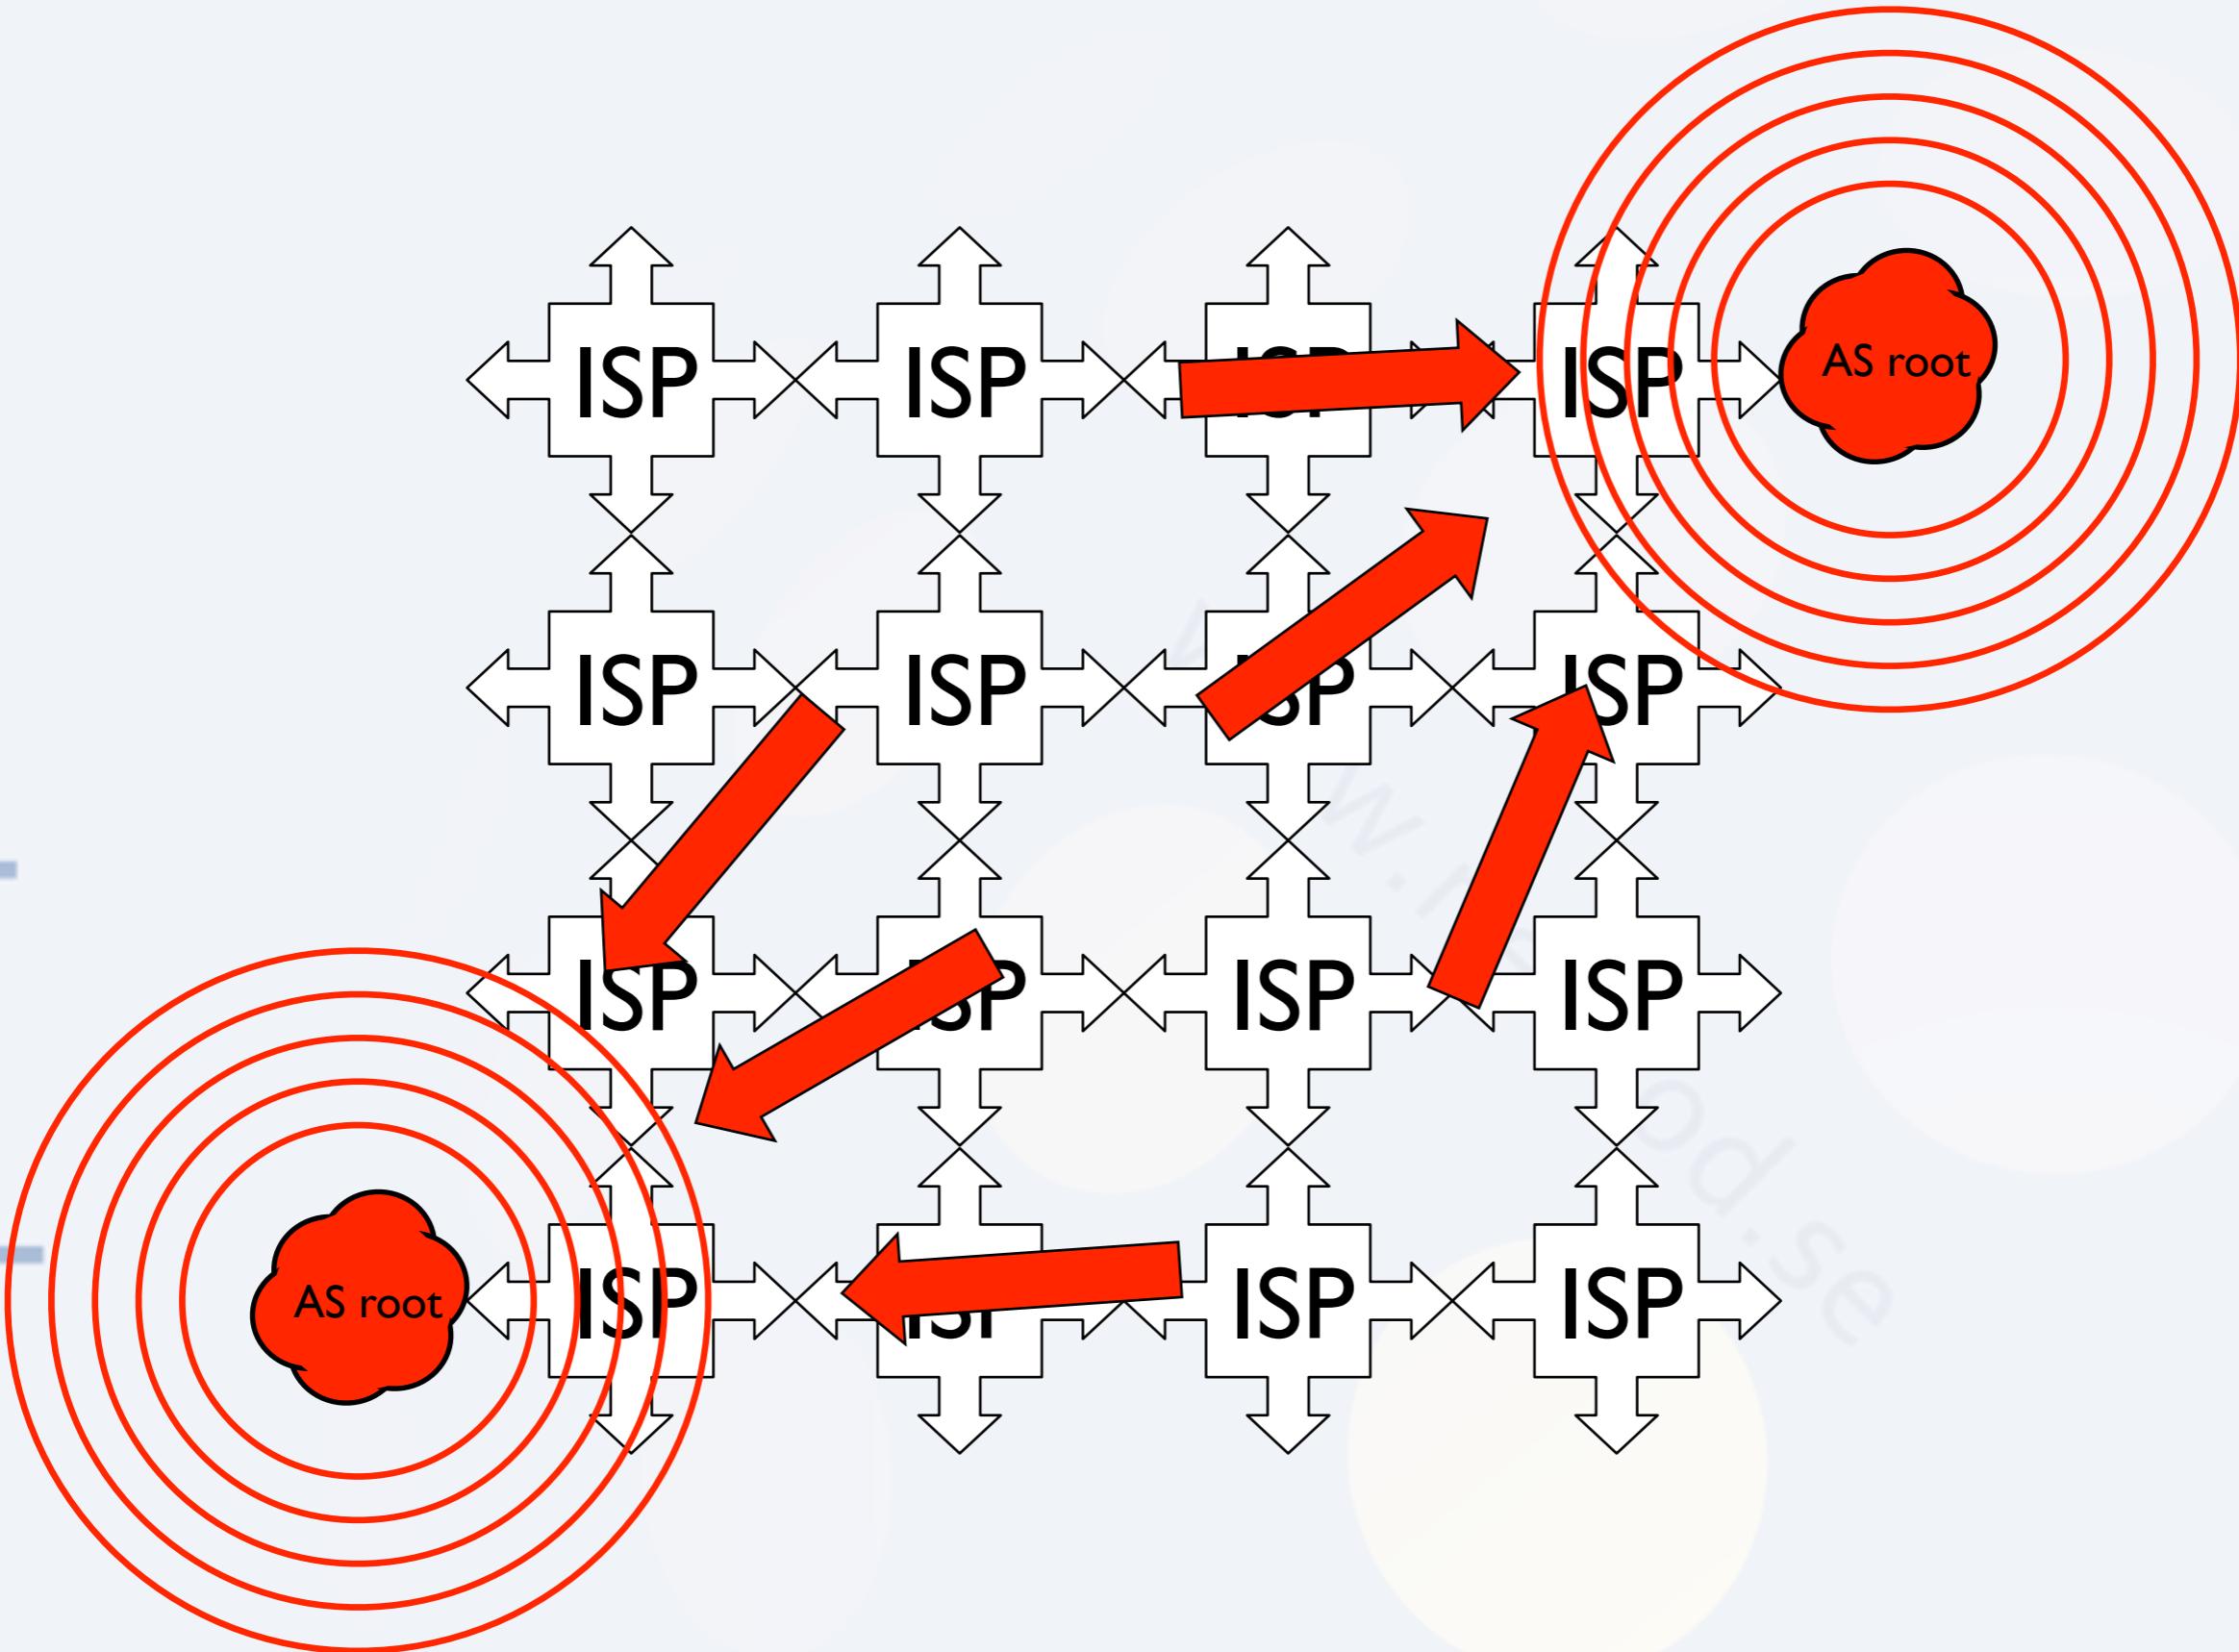


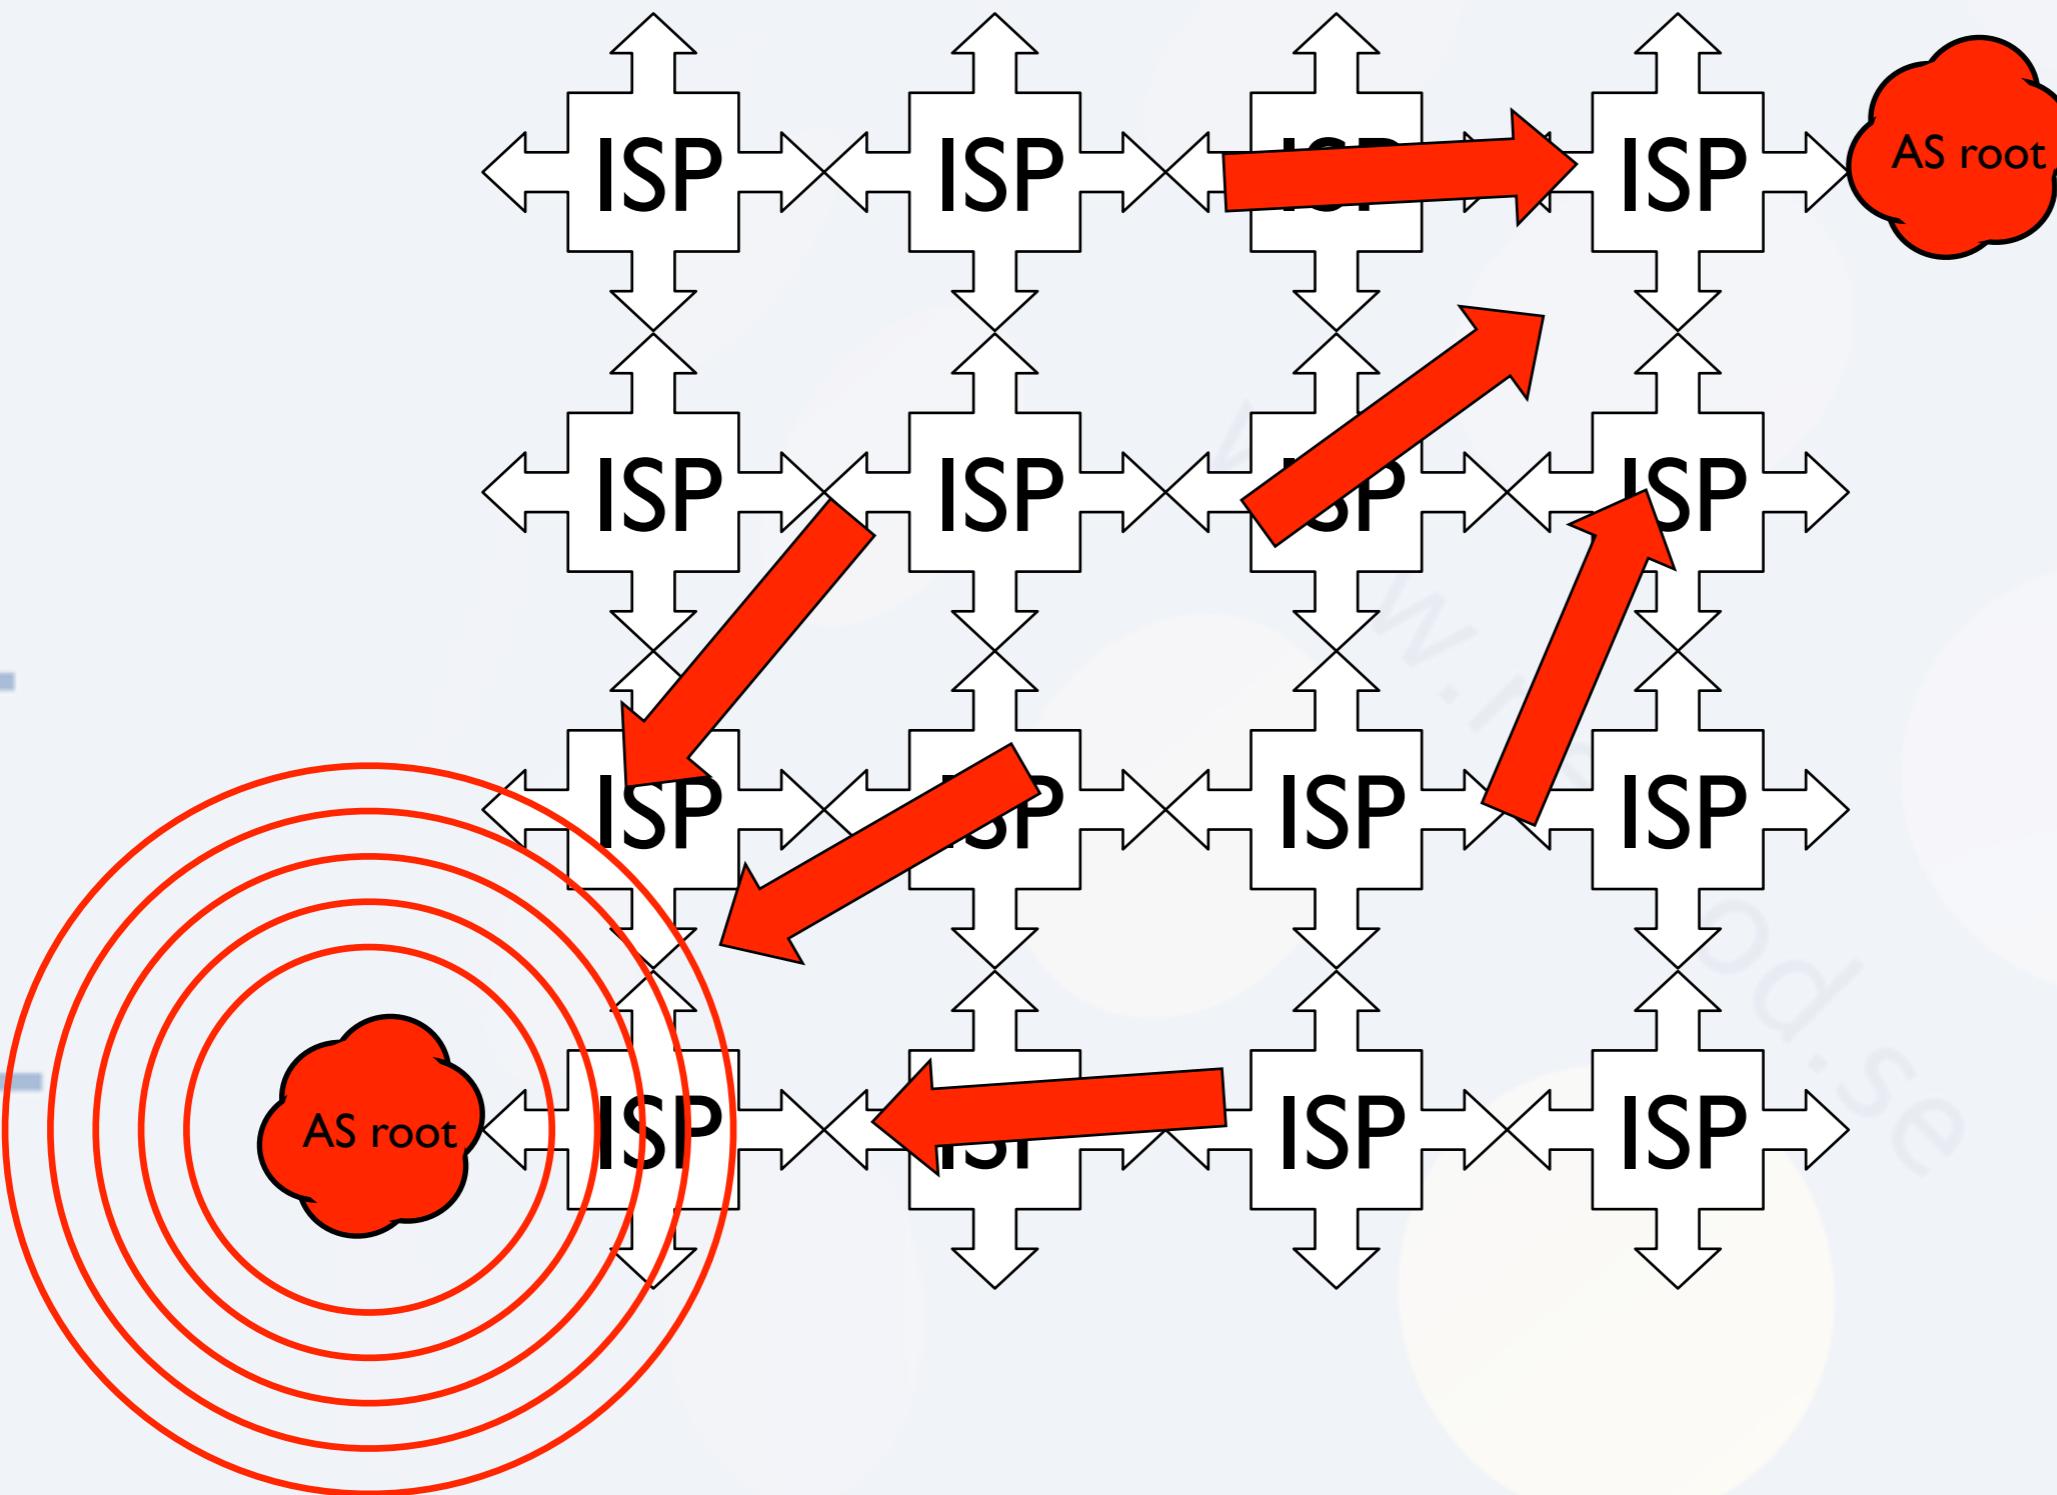


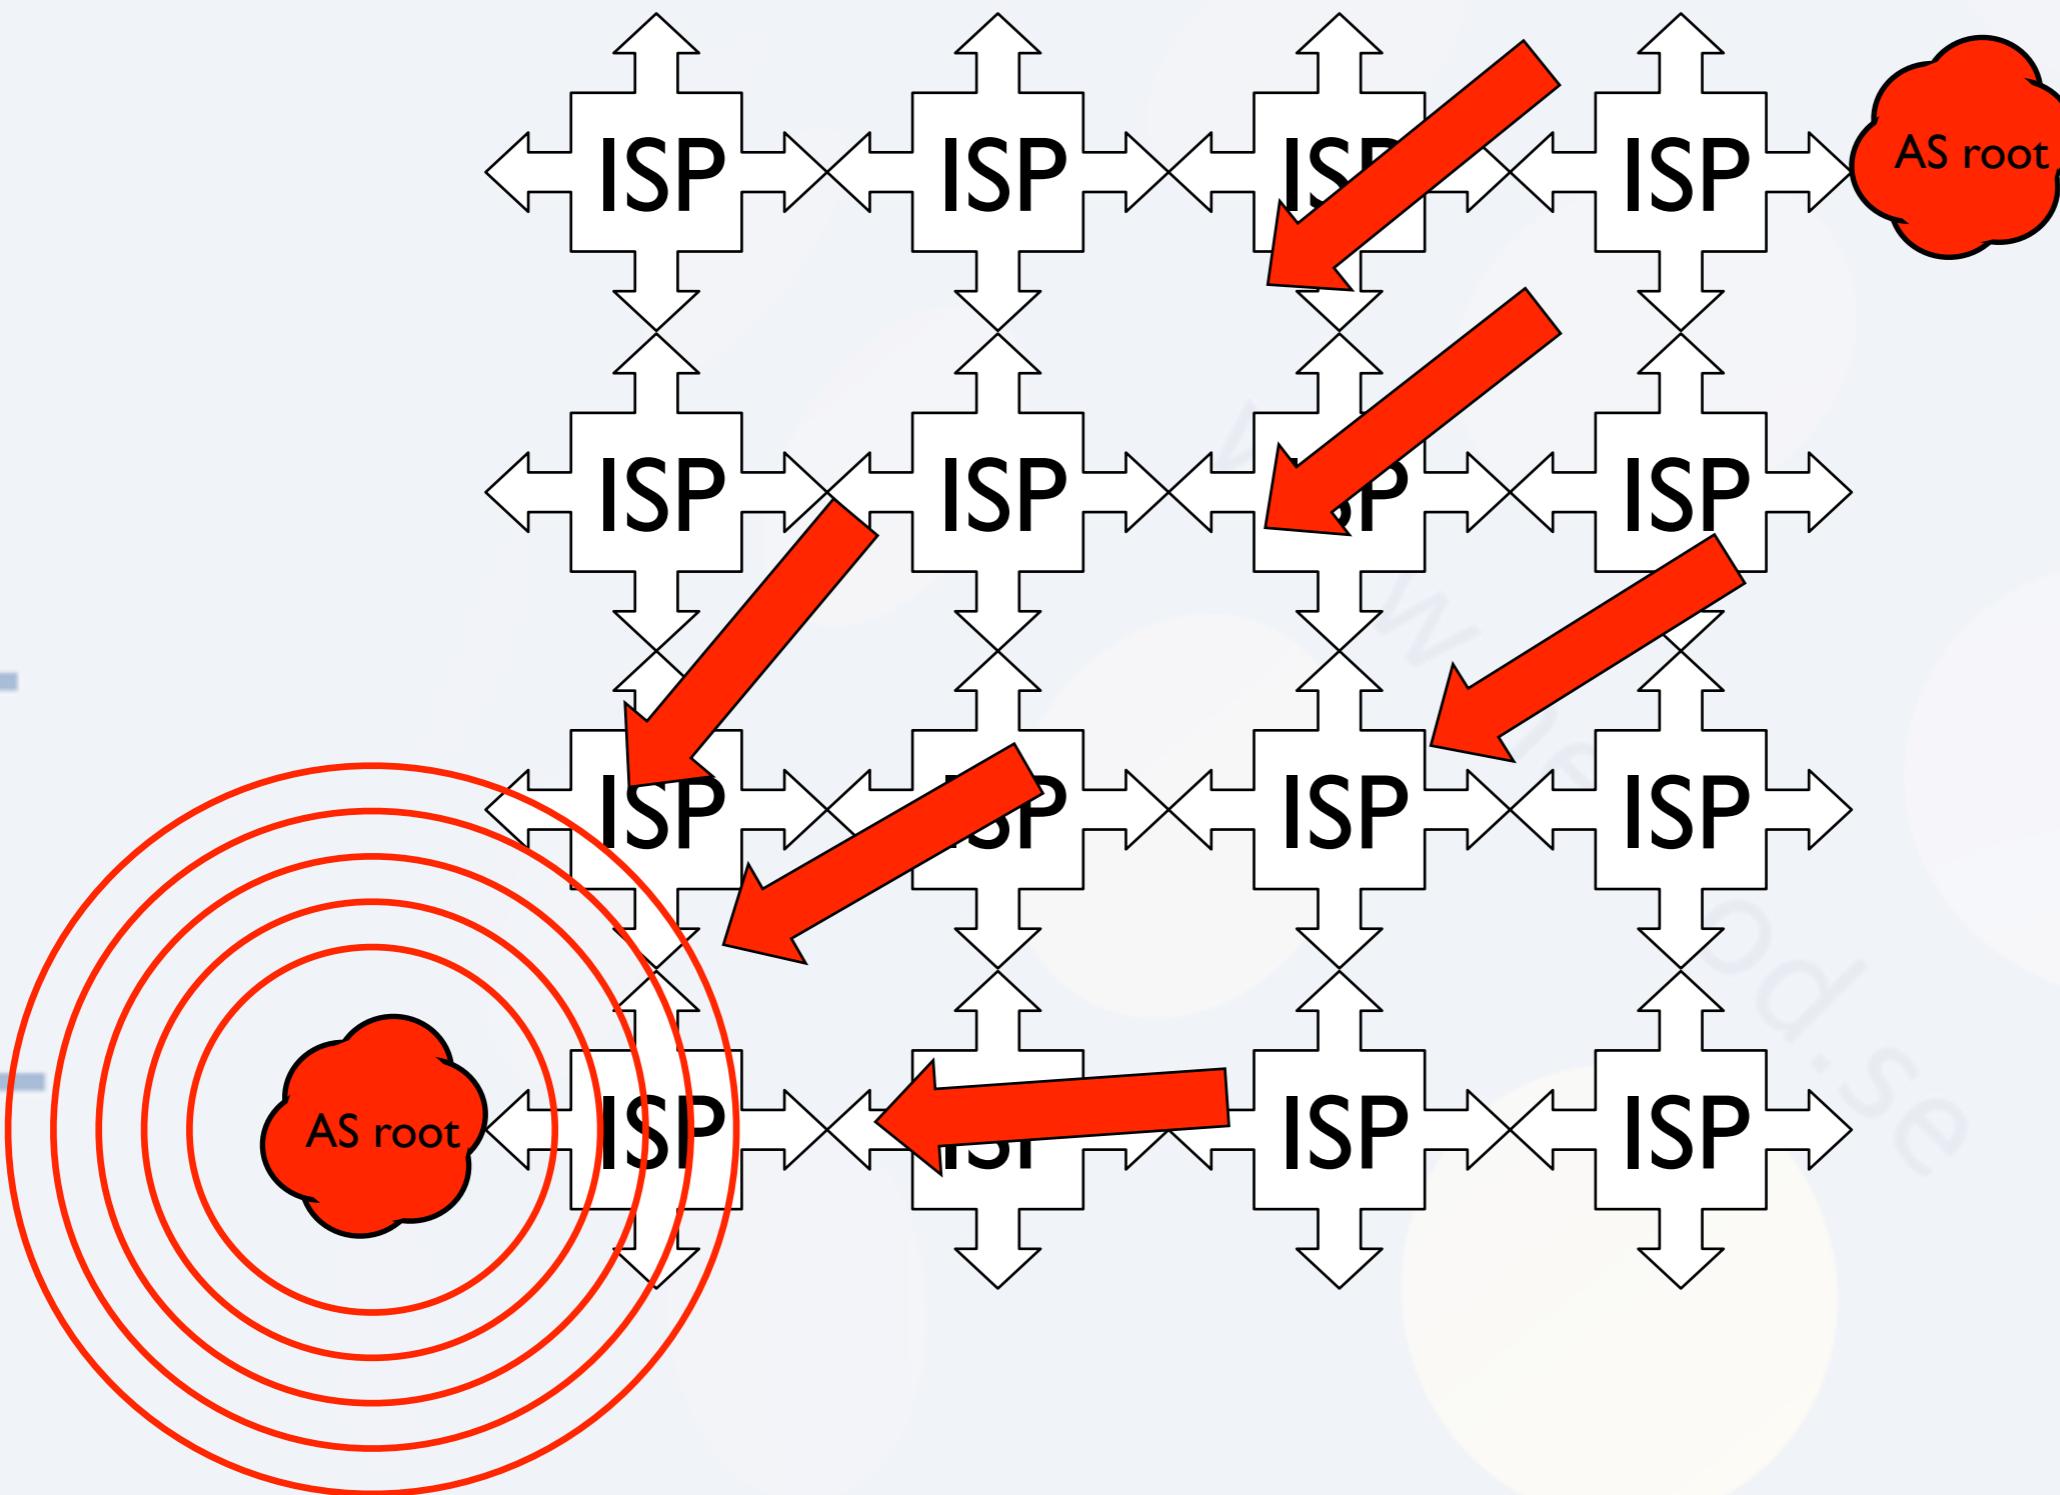












# Current Server Map



netnod

©2011 Lars-Johan Liman 201109-ipop-DNS-liman KTH 2011-09-27

29

www.netnod.se

# Root Zone Administration



# The Good Old Days™

- Used to be Jon Postel (RIP), head of the IANA (Internet Assigned Numbers Authority).
  - ▶ ... and that was that.
  - ▶ ... and then he died!
- Jon initiated the creation of an organisation for the purpose.
  - ▶ This has developed into ICANN (Internet Corporation for Assigned Names and Numbers).
  - ▶ The IANA is now a subdivision of ICANN.



# ICANN

- ICANN was designed as a multi-stakeholder, bottom-up organisation.
  - ▶ Tried to let everyone have a say – from single citizens to business, and governments.
  - ▶ This works only so-so.
- ICANN is a policy organisation, and it sets policies for new TLDs, and decides which to establish.



# Still Going Strong ...

1974!

| HOST NAMES   |                    |                                 |   |
|--------------|--------------------|---------------------------------|---|
| HOSTNAME     | HOST ADDR<br>(Dec) | LIAISON                         | STATUS                                    |
| AFWL-TIP     | 176                | D Hyde (505)247-1711 x3803      | TIP, Up 3-74                              |
| ALOHA-TIP    | 164                | R Binder (808)948-7066          | TIP                                       |
| AMES-11      | 208                | J Hart (415)965-5935            | USER, up 12-73                            |
| AMES-67      | 16                 | W Hathaway (415)965-6033        | SERVER                                    |
| AMES-TIP     | 144                | W Hathaway (415)965-6033        | TIP                                       |
| ANL          | ?                  | L Amiot (312)739-7711 x1309     | SERVER, up 2-74                           |
| ARPA-DMS     | 28                 | S Crocker (202)694-5037         | USER, Agency use only                     |
| ARPA-TIP     | 156                | S Crocker (202)694-5037         | TIP                                       |
| BBN-IIA      | 5                  | R Thomas (617)491-1850 x483     | Peripheral processor<br>for #69, up 12-73 |
| BBN-ID       | 232                | A McKenzie (617)491-1850 x441   | USER                                      |
| BBN-NCC      | 40                 | A McKenzie (617)491-1850 x441   | USER                                      |
| BBN-TENEX    | 69                 | R Thomas (617)491-1850 x483     | SERVER                                    |
| BBN-TENEXB   | 133                | R Thomas (617)491-1850 x483     | SERVER, Limited                           |
| BBN-TESTIP   | 158                | A McKenzie (617)491-1850 x441   | TIP (magtape)                             |
| BELVOIR      | 27                 | W Andrews (703)664-5511         | USER, up 6-74                             |
| BRL          | 29                 | M Romanelli (301)278-4574       | USER                                      |
| CASE-10      | 13                 | J Calvin (216)368-2984          | SERVER                                    |
| CCA-TENEX    | 31                 | R Winter (617)491-3670          | SERVER                                    |
| CCA-TIP      | 159                | R Winter (617)491-3670          | TIP                                       |
| CMU-10A      | 78                 | H Van Zoeren (412)621-2600 x160 | SERVER                                    |
| CMU-10B      | 14                 | H Van Zoeren (412)621-2600,x160 | SERVER                                    |
| CMU-11       | 142                | C Pierson (412)621-2600 x130    | USER, up Spring 74                        |
| CMU-CC       | 206                | D King (412)621-2600 x2683      | USER, up Spring 74                        |
| DOCB-TIP     | 153                | S Stevenson (303)499-1000 x3138 | TIP                                       |
| EGLIN        | ?                  | E Blackwell (904)882-3734       | Up 3/74                                   |
| ETAC         | 20                 | G Petregal (202)433-3911        | USER, up Spring 74                        |
| ETAC-TIP     | 148                | G Petregal (202)433-3911        | TIP (magtape)                             |
| FNWC         | 33                 | M Reese (408)646-2817           | USER, up 2-74                             |
| FNWC-TIP     | 161                | M Reese (408)646-2817           | TIP                                       |
| GWC-TIP      | 152                | A Wells (402)294-2968           | TIP (magtape)                             |
| HARV-1       | 73                 | B Reussow (617)495-4147         | USER                                      |
| HARV-10      | 9                  | B Reussow (617)495-4147         | SERVER                                    |
| HARV-11      | 137                | S Bradner (617)495-3864         | USER                                      |
| HASKINS      | 197                | F Cooper (203)865-6163          | USER (VDH),<br>up Spring 74               |
| HAWAII-500   | 100                | J Davidson (808)944-7455        | SERVER, up 1-74                           |
| HAWAII-ALOHA | 36                 | R Binder (808)948-7066          | USER, Up 12-73                            |
| Ih-TENEX     | 15                 | J McConnell (408)735-0635       | SERVER                                    |
| Ih-TENEXA    | 79                 | J McConnell (408)735-0635       | Peripheral processor<br>for #15           |
| ILL-CAC      | 12                 | T Milke (217)333-8469           | USER                                      |
| ILL-NTS      | 76                 | T Milke (217)333-8469           | USER                                      |
| ISI-DEVTENEX | 150                | R Hoffman (213)822-1511 x190    | USER, up 1-74                             |
| ISI-SPEECH11 | 22                 | R Hoffman (213)822-1511 x190    | USER, up 1-74                             |

55

2011:  
Current  
Chairman of  
ICANN

# The IANA

- The IANA is an administrative function that is actually governed and funded by the US Dept of Commerce\* under very old ARPA regulations.
  - ▶ The IANA functions are performed by ICANN under a contract with the US Gov't.
- The IANA handles DNS related stuff, and performs redelegations of existing TLDs.
- (... and a lot of other stuff that is non-DNS-related.)

\*) Specifically its National Telecommunications and Information Administration (NTIA)



# The DoC (NTIA) and Verisign

- NTIA must approve any changes
  - ▶ ... before they can be implemented by ...
- Verisign, who
  - ▶ Generate DNSSEC signature, and
  - ▶ Change the actual root zone file.
- The root server operators ...
  - ▶ ... pick up the root zone file from Verisign.
  - ▶ The root zone is publicly available at  
<ftp://ftp.internic.net/domain/root.zone>



# Root Server Operators

- 12 separate organisations
  - ▶ Great variety: from military, via universities, to commercial.
- No formal relationship with
  - ▶ DoC (except for Verisign)
  - ▶ IANA (except for Verisign)
  - ▶ ICANN
  - ▶ each other
  - ▶ any parent organisation (none exists!)
- Technically very well coordinated!



# Root Server Operators

- Go by old gentlemens' agreements
  - ▶ ... with Jon Postel (RIP), former head of IANA.
  - ▶ There is no process to change the set of root ops!
- Vast experience in DNS, routing and system administration.
- 100,00 % uptime\* on the system as a whole for more than 10 years.
- Meet 3 times/year for coordination.

\*) Knock on wood ...

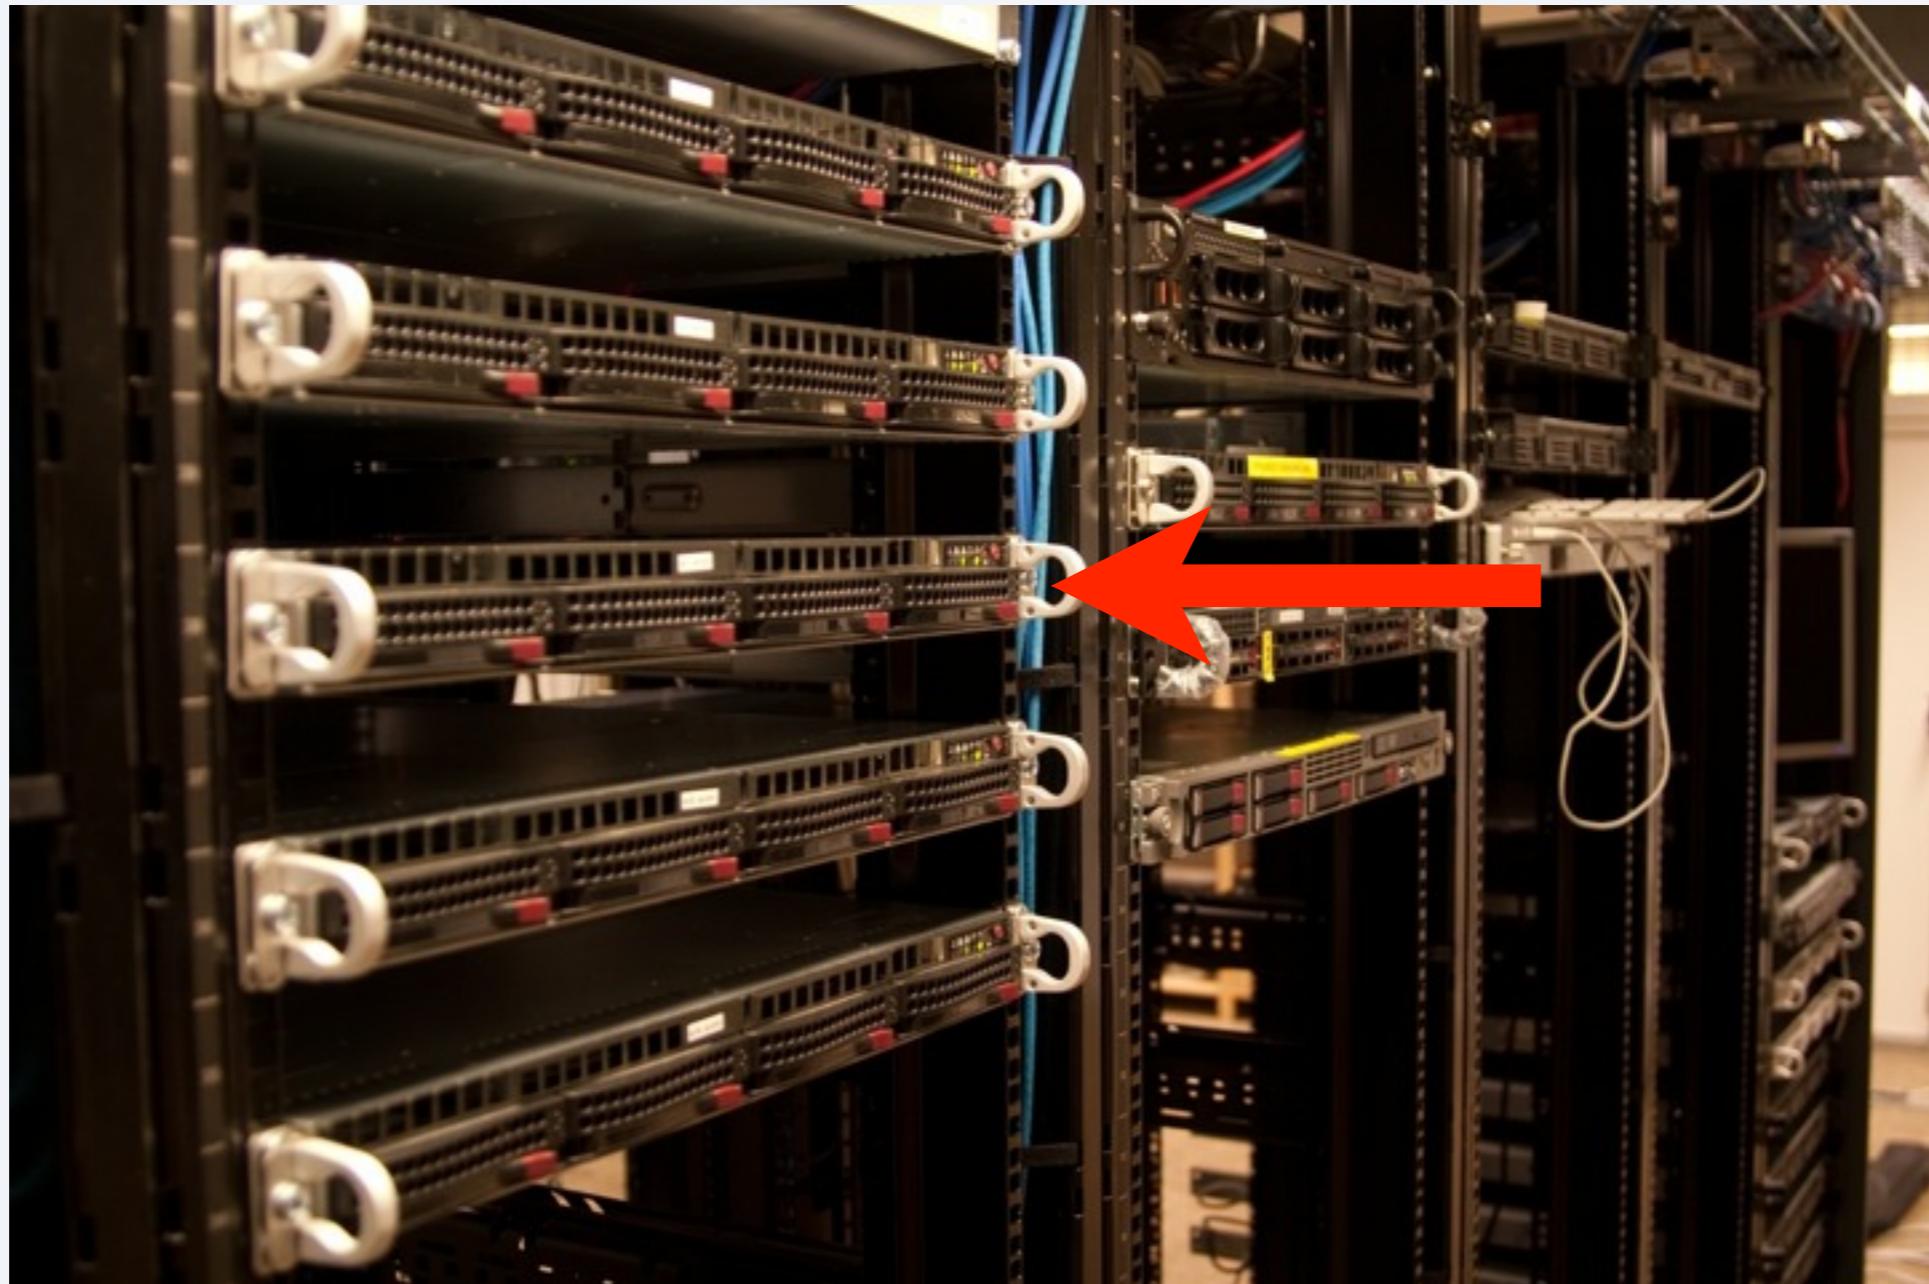


# i.root-servers.net

- Ordinary DNS servers.
  - ▶ Serves the "." zone.
  - ▶ Stockholm: located in communications bunkers provided by PTS.
- Anycast
  - ▶ 44 places across the globe.
  - ▶ From Stockholm to Johannesburg, from Wellington to San Francisco ...
  - ▶ ... including odd places like Colombo, Katmandu, Yerevan, Thimphu, ...
- Serves ~1.8 billion queries/day.



# s1.sth



# Which Node?

- How do I know which node I'm using?

```
$ dig @i.root-servers.net. \
hostname.bind. TXT CHAOS \
+short
```

“s1.sth”



# Contact Information

Lars-Johan Liman  
Netnod Internet Exchange  
Box 30194  
104 25 Stockholm  
Sweden  
Tel: +46-8-562 860 12  
E-mail: liman@netnod.se

Slides available at:  
[http://www.netnod.se/liman/presentations/kth/  
201109-ipop-DNS-liman.pdf](http://www.netnod.se/liman/presentations/kth/201109-ipop-DNS-liman.pdf)

