Anycastning of DNS-servers

Kurt Erik Lindqvist

Netnod / Autonomica

Netnod

Who am I?

- Ålcom (EUnet Finland) -> 1997
- EUnet Sweden 1997-1998
- KPNQwest Sweden 1998-2000
- KPNQwest 2000-2002
- 2002- Consultant (Netnod)

- Other
  - Current chairman of Euro-IX
  - Chair IETF multi6 WG
  - Chair RIPE ”NCC-Services WG”.
  - Chairman Swedish Operators Forum
  - Member of IETF ops-dir and addr-dir

Who is Netnod?

- Owned by the TU-Foundation
- Operates exchanges in four cities in Sweden
  - Stockholm, Gothenburg, Malmö and Sundsvall
- Fully owned daughter company Autonomica
  - Operates the exchanges
  - Operates i.root-servers.net

Netnod Corporate Structure

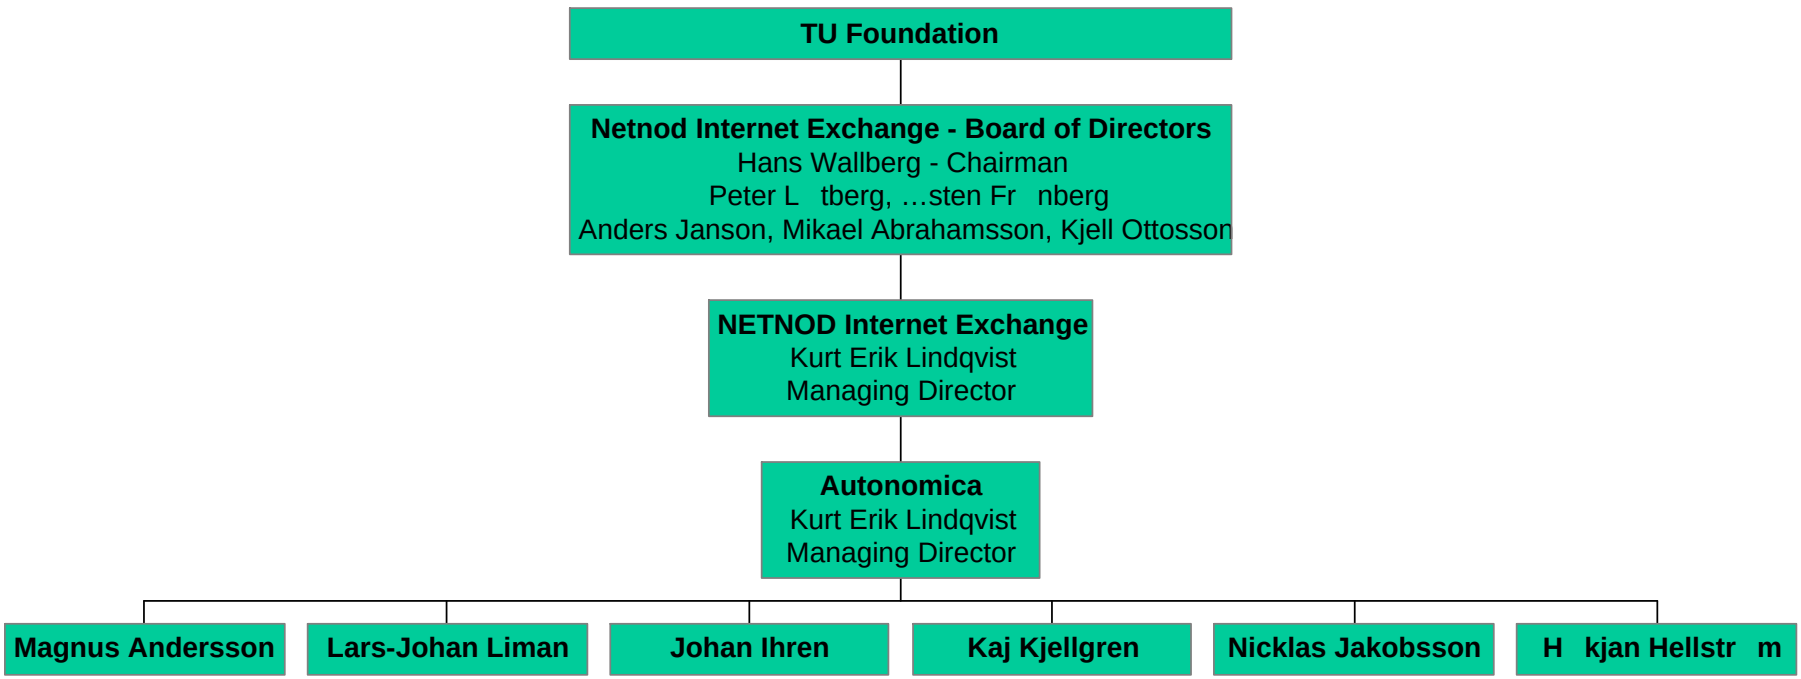

Root-nameservers - what are they?

- The entrance to the DNS-systems database
- Root-nameservers can always point further in the system, or tell that something does not exist

A lot of queries

- Stockholm: around 4.000 qps
- .... Equivalent of 345.600.000 queries/day
- a.root-servers.net is usually at around 14.000 qps

STUPID queries

- 8-10% is from net 10.0.0.0/8 (RFC1918) etc.
  - Can not be replied to. Filtered.
- 3,6 % queries **ABOUT** net 10.in-addr.arpa, 168.192.in-addr.arpa.
  - Even though they have been delegated
  - AS112
- 4% queries on "localhost"
  - Around 13.720.000 queries/day
- 3% queries for ".local"
- 2,3% recursive queries

STUPID queries

- Poorly configured

  _ldap._tcp.Standardname-des-ersten-Standorts._sites.gc._msdcs.USD.local

- ” The Marjasinproblem”

  SC1DREV_TByggesagerIgangv\145rende\032sagerSag\032011.09\032Ny\032receptionsskranke\032afsnit\0323981Rekvisition_J\032Pihl.doc

Explanations?

- No, not really
- Information and education
- Get over it…
- ".local" special
  - Political problem that can be a considered norm, needs to be treated carefully…

Code problems

1. Code diversity
2. Security vulnerabilities in DNS-code?
   - Very close relationship with the developers
3. Security vulnerabilities in operating systems
   - Known systems with open code and many users
4. Router operatingsystem problems?
   - Good relation to the developers

Attacks

- Deliberately “broken” queries trying to exploit security holes
- Cache pollution?
  - SEP
- Distributed Denial of Service Attacks.
  - We don't lik'em
  - But we get them so often that most of them are just noise

Loadbalancing

- Queries are distributed between the several servers in the same installation
- Almost all root-servers do this today

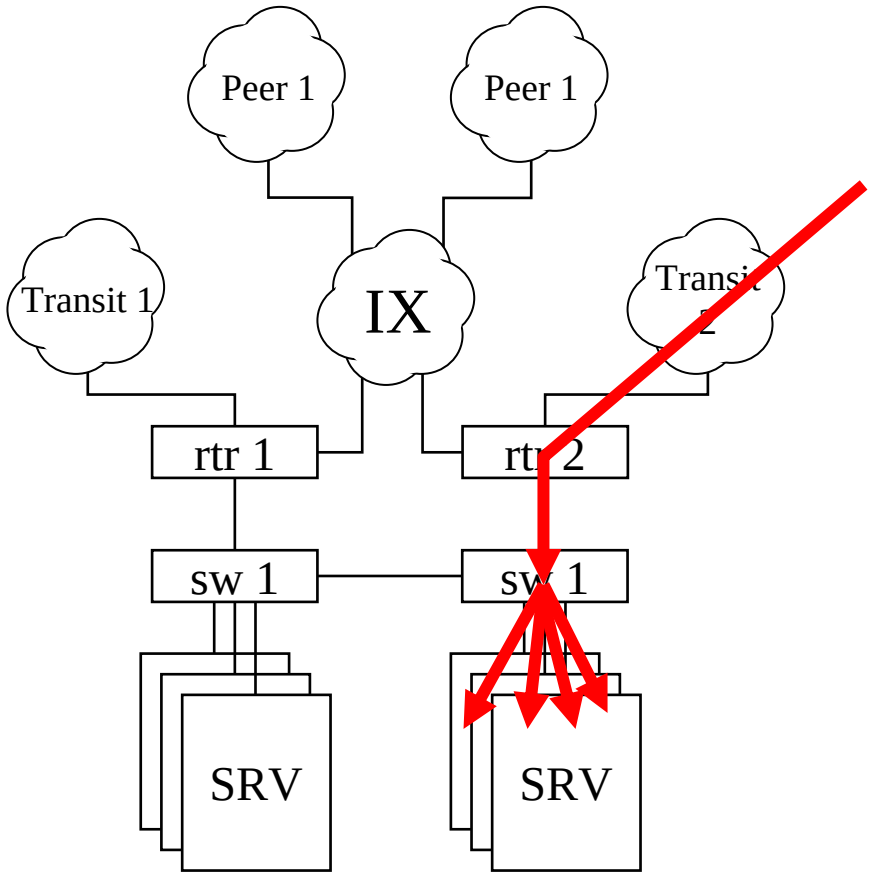

Loadbalancing

- Good idea
  - Increases query capacity linear
- But ...
- The edge of the network will always be larger than any given server cluster

Anycast

What is anycast?

- A way to install multiple copies at multiple locations

Why anycast?

- Better service to more users
  - Noone could decide on where to locate new roots. Root-ops do it themselves
- Cancels the effects of the DDoS-attacks

How does anycast work?

- Servers locates around the world
- ***THE SAME*** network information
- **THE SAME** data
- ***DIFFERENT*** servers
- The routing system will decide where the query is sent

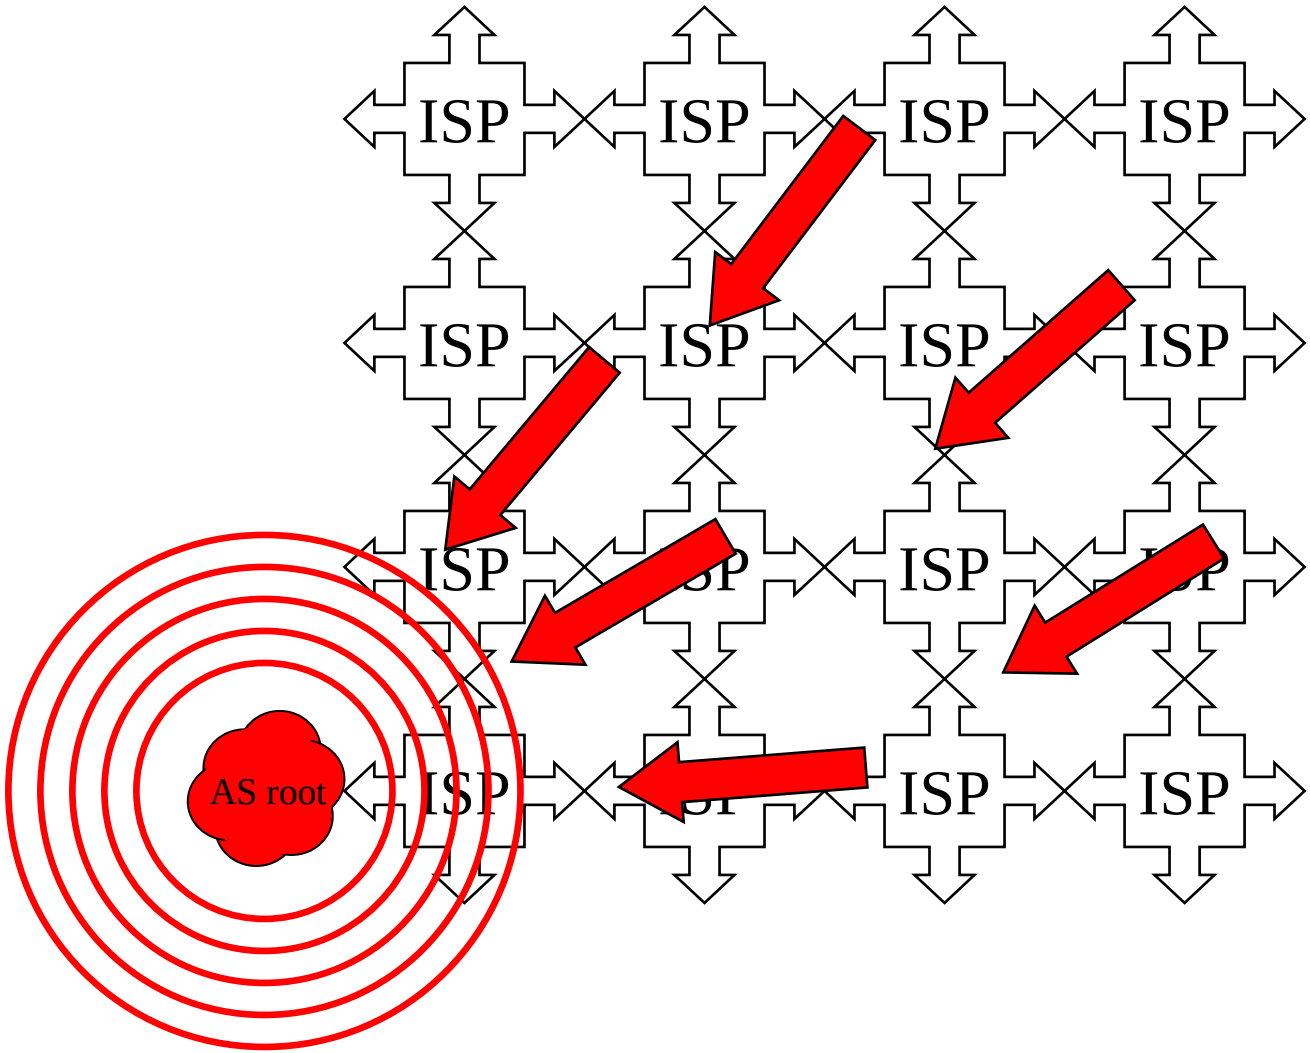

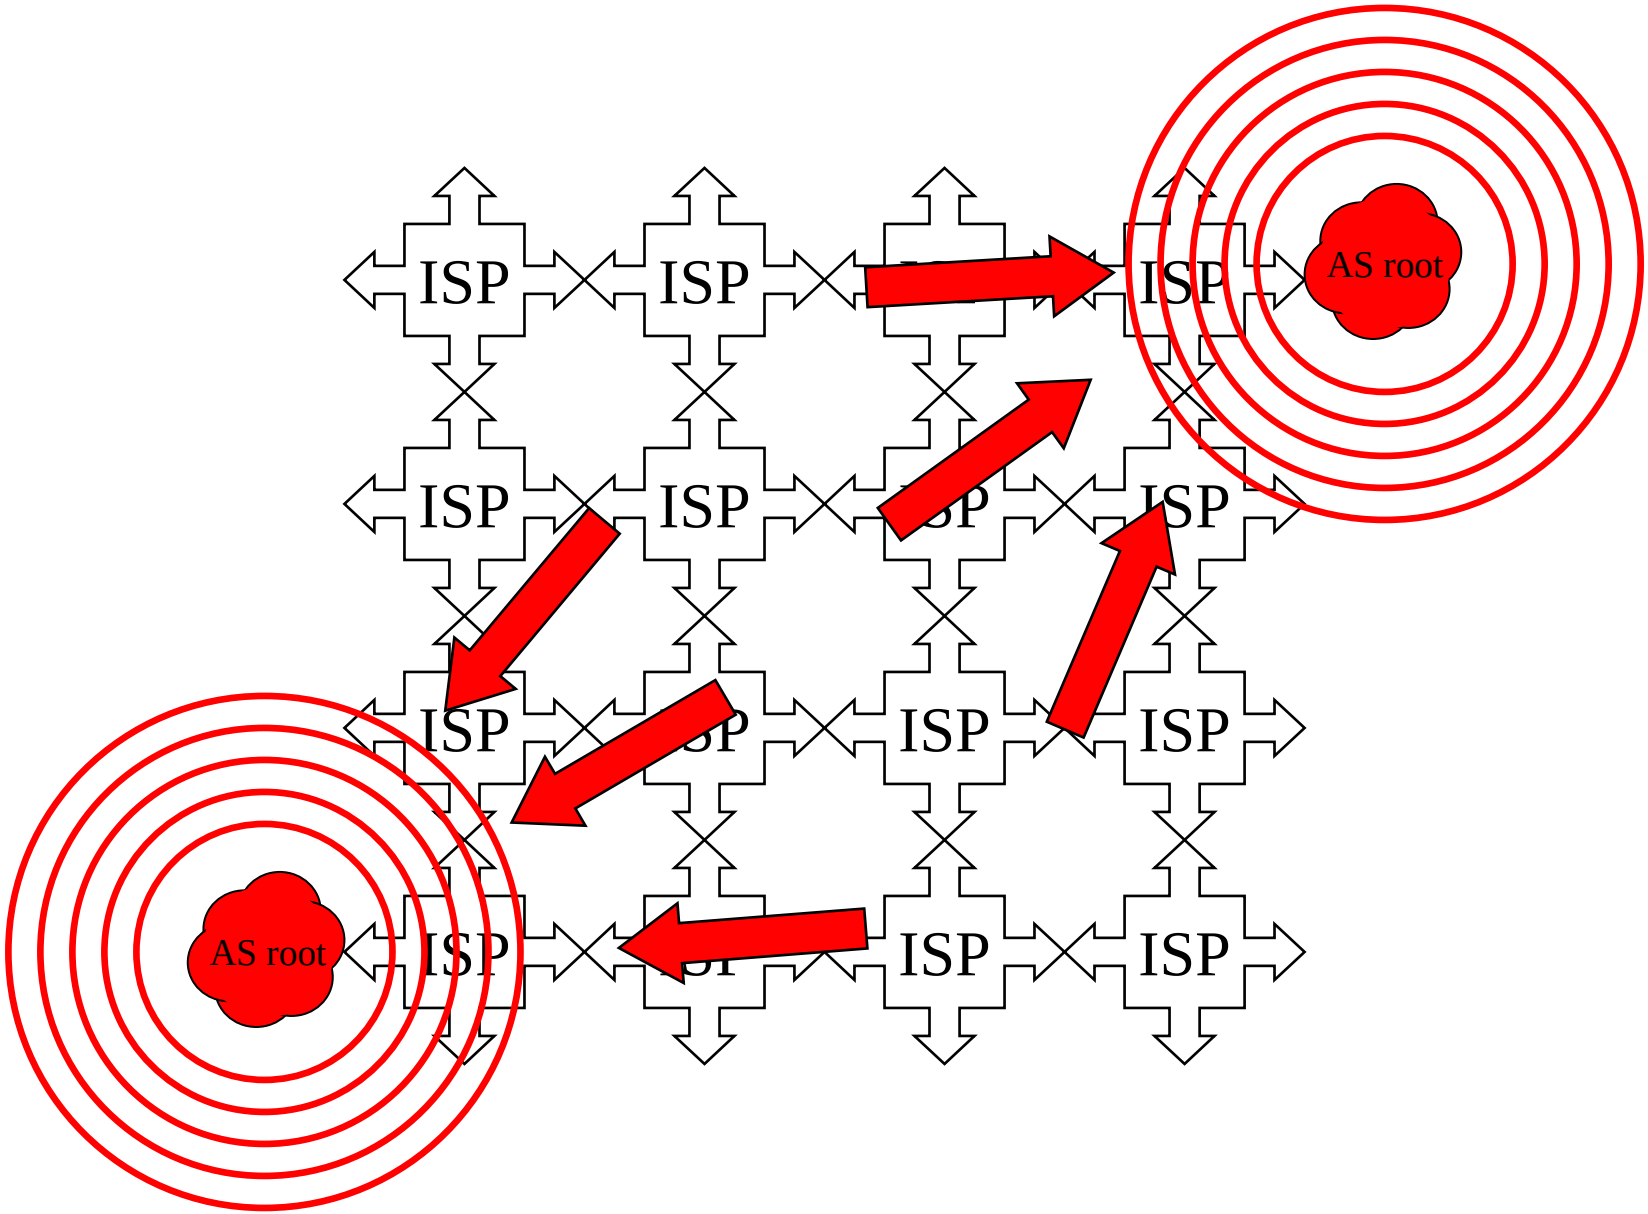

Advantages with anycast

- The service is closer to the users
- Automatic loadbalancing
- Automatic fail-over.
- Localization of attacks

Global vs Local nodes

- Global anycast nodes
  - Announce the prefix with no limitations
  - Will need to be transited
  - Have potential to service all of the Internet
- Local nodes
  - Announced with some form of limitation (no-export, specific communities etc)
  - Will only service ”local” part of the net

Drawbacks / Advatages with local/global

- Local nodes does not risk taking out other nodes due to routing problems
  - Failed RPF checks
  - Asymetric routing
  - BGP Dampening
- Global nodes have the advantage of providing fall-back
  - We have done fall-back to London to work in Stockholm

RPF example

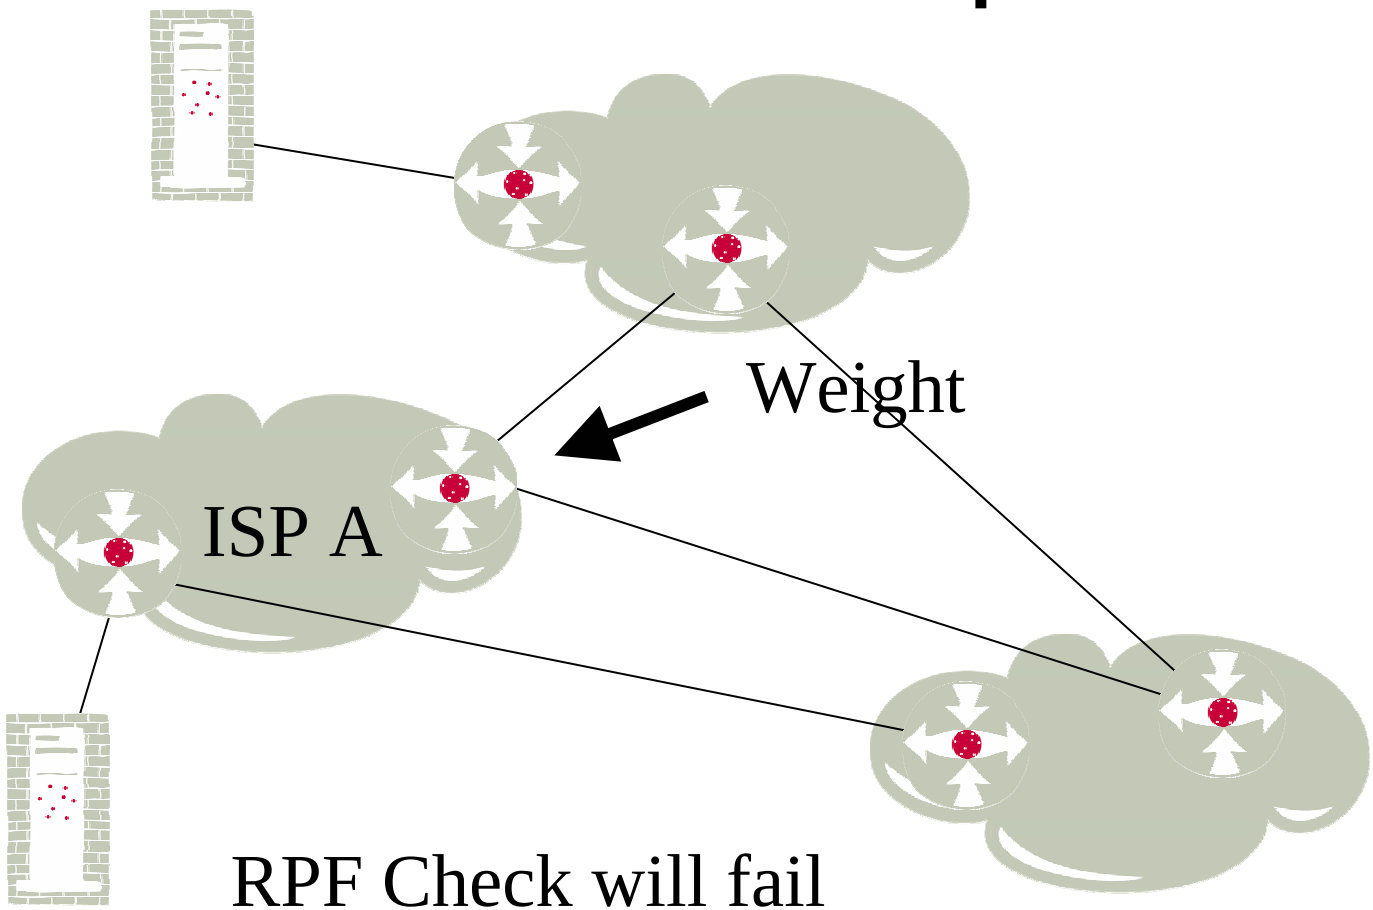

RPF Check will fail

Monitoring

- Ping doesn’t really work :)
- Monitor
  - Routing
  - RTT / Packet drop etc
  - Service
- Routing is tracked by a number of people
  - RIPE RIS
  - Route-views
- RTT / Packet drop
  - Own meassurement probes
  - RIPE DNSMON
  - But which site are we querying?

Monitoring

- Node identification

```
laptop2:~$ dig @i.root-servers.net hostname.bind txt ch i.root-servers.net

; <<>> DiG 9.2.2 <<>> @i.root-servers.net hostname.bind txt ch i.root-servers.net
;; global options:  printcmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 11990
;; flags: qr aa rd; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 0

;; QUESTION SECTION:
;hostname.bind.                 CH      TXT

;; ANSWER SECTION:
hostname.bind.          0       CH      TXT     "s1.sth"
```

Monitoring

- Identification
  - Used for debugging
  - For statistics to be generated per host

Drawbacks with anycast?

- Changes the balance
- Complex
  - Violates all known priciples (KISS, PLS, 1-to-1).
- Hard to administer
  - Monitoring
  - Reachability
  - Data transfer
- Hard to debug

How do you do it?

- Needs a ”service interface” on the host
- The host needs to communicate routing
  - IGP
  - BGP

Root-servers today?

- 7 of 12 root-server operators are running it
  - **C:**Does it differently. Only inside their own network. Herndon VA; Los Angeles; New York City; Chicago
  - **F:** Ottawa; Palo Alto; San Jose CA; New York City; San Francisco; Madrid; Hong Kong; Los Angeles; Rome; Auckland; Sao Paulo; Beijing; Seoul; Moscow; Taipei; Dubai; Paris; Singapore; Brisbane; Toronto; Monterrey
  - **I:** Stockholm, Helsingfors, Milano, London, Geneva, Amsterdam, Bangkok, Hongkong
  - **J:** Dulles VA; Mountain View CA; Sterling VA (2 locations); Seattle WA; Amsterdam; Atlanta GA; Los Angeles CA; Miami; Stockholm; London
  - **K:** London; Amsterdam; Frankfurt
  - **M:** servers in Tokyo and Osaka since 1998.

i.root-servers.net plans

- Build around 20 sites within a year
  - Distributed across the globe
- We are so far paying for all the hardware
- We are looking for major IXPs
- Also for large ISPs

?

Contact

Netnod Internet Exchange i Sverige AB

Bellmansgatan 30*I*
SE-118 47 Stockholm
Sweden

Office address: Bellmansgatan 30*I*

Telephone: +46-8-615 85 70
Telefax: +46-8-442 09 67
E-mail: kurtis@netnod.se
URL: http//www.netnod.se/