

Autonomica Anycast Service

APTLD meeting

Kuala Lumpur, 18-24 May 2008

Nurani Nimpuno



Who is Autonomica / Netnod?

- Neutral, non-profit & independent organisation
 - Owned by the TU foundation
- Operator of i.root-servers.net
 - One of 13 root servers in the world
 - Currently 31 anycast instances globally
- TLD unicast & anycast slave service provider
 - Have provided production anycast services since 2003
- Operator of exchange points in Sweden
 - through Netnod IX



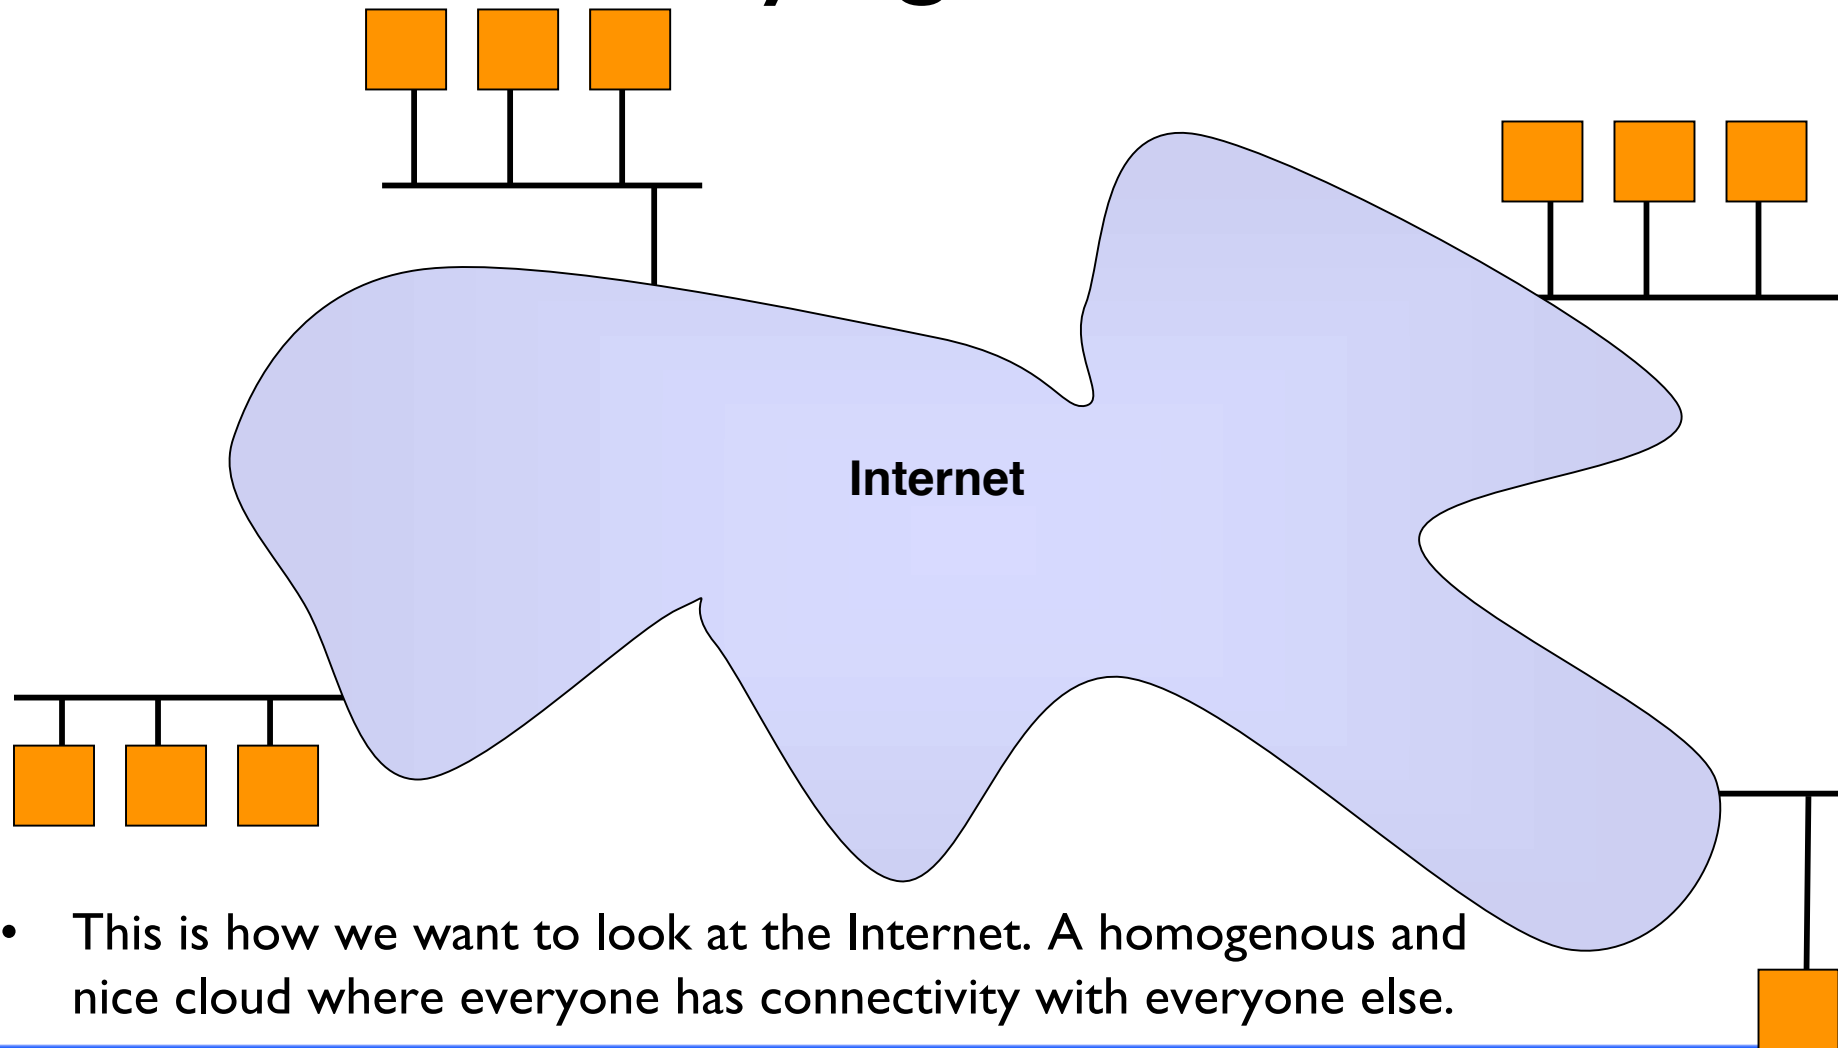
What is anycast?

- A technique that allows several (identical) servers on the Internet to share the same IP address
 - BGP directs packets to the topologically closest instance
 - Anycast shifts the redundancy management from the DNS layer to the routing layer
- Mitigates impact of DDos attack
 - By localising the attack with an increased foot print

Redundancy

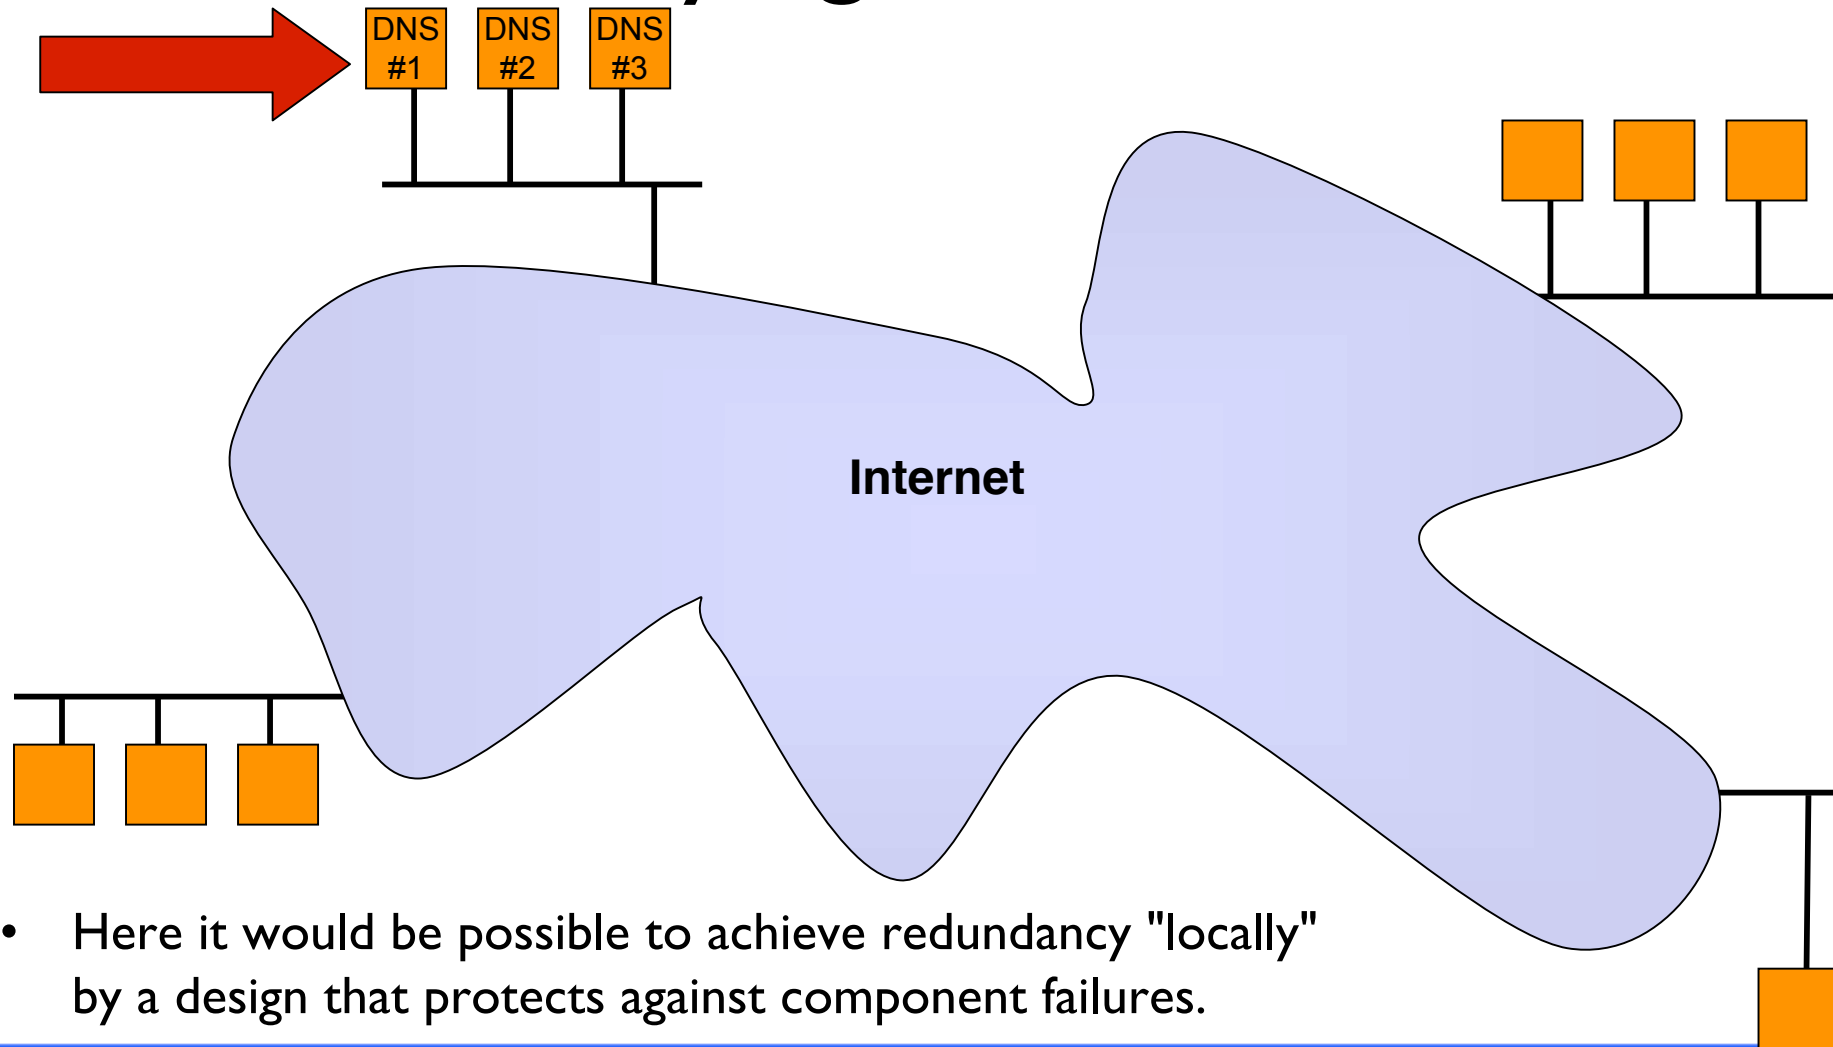
- Any important service should of course be redundant and robust
 - Can lead to very expensive and complicated machines with special "high-availability" design.
 - These things are usually very, very much more expensive than standard components
- In DNS this was part of the thinking from the outset
 - DNS protocol has provision for "redundancy" in the application layer
 - i.e. multiple name servers for the same "zone"
 - Therefore the general opinion for many years was that the redundancy needs of DNS was a solved problem

Underlying structure



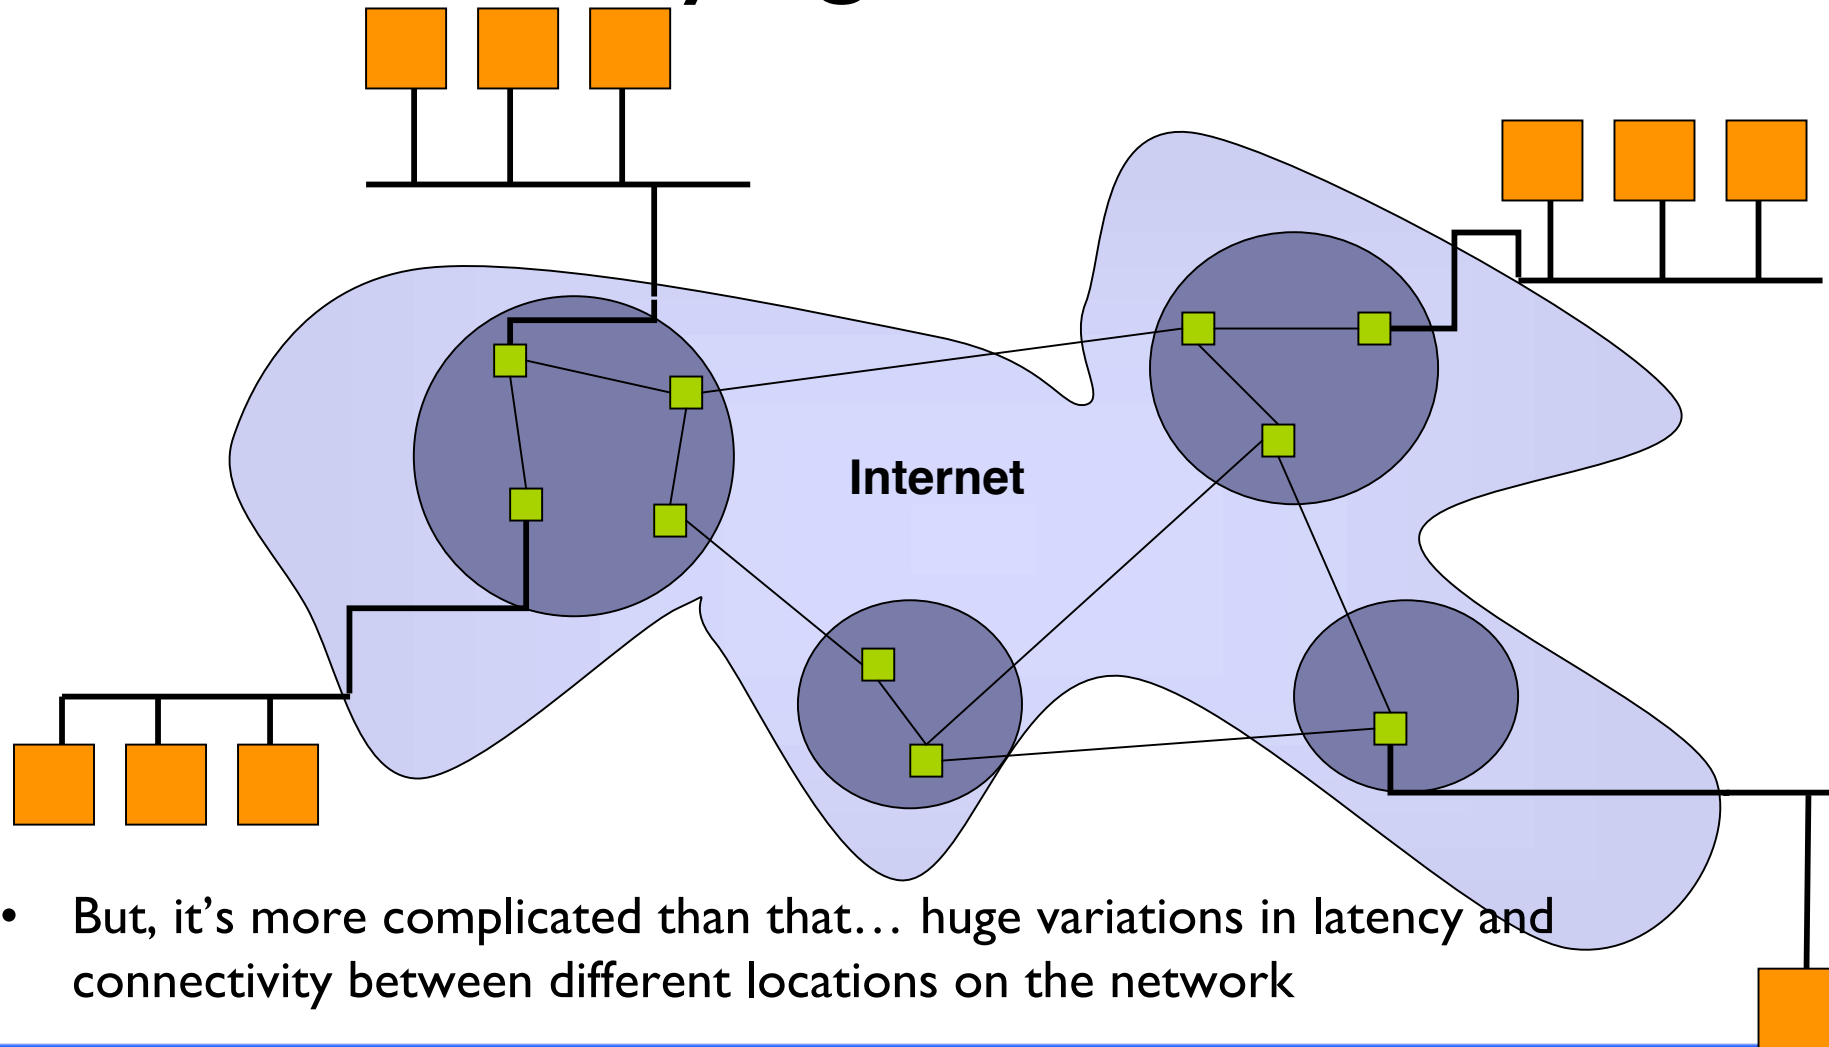
- This is how we want to look at the Internet. A homogenous and nice cloud where everyone has connectivity with everyone else.

Underlying structure #2



- Here it would be possible to achieve redundancy "locally" by a design that protects against component failures.

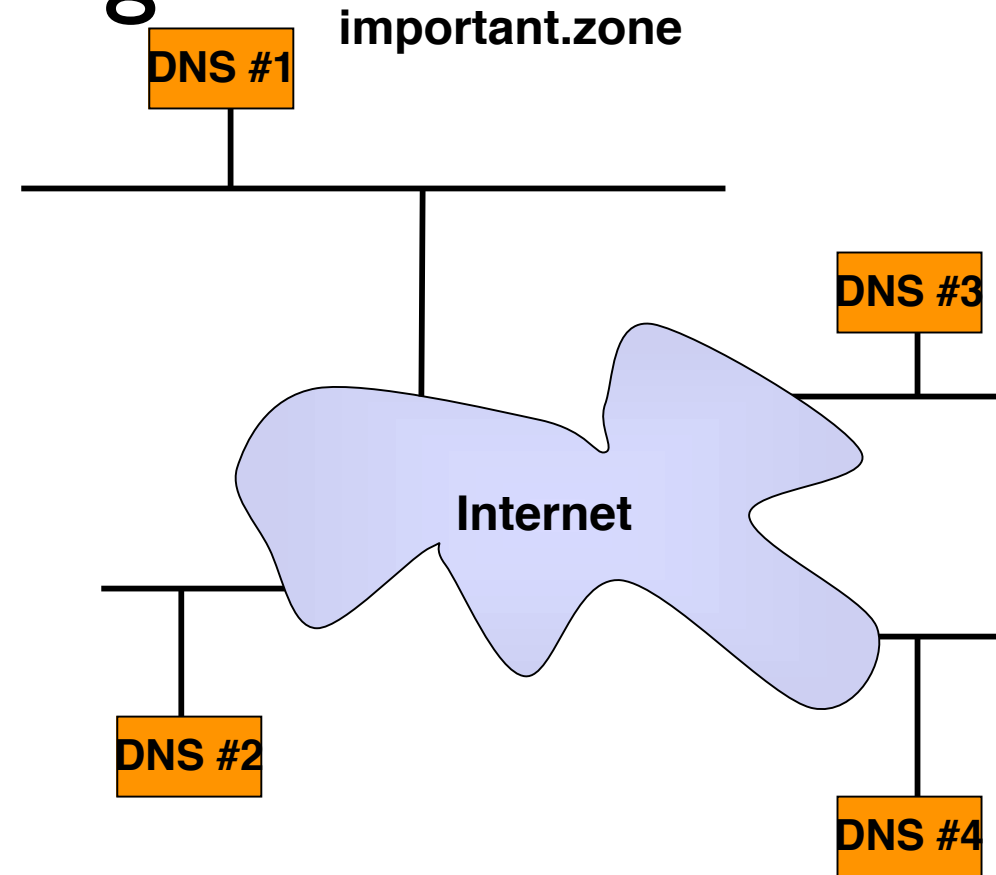
Underlying structure #3



- But, it's more complicated than that... huge variations in latency and connectivity between different locations on the network

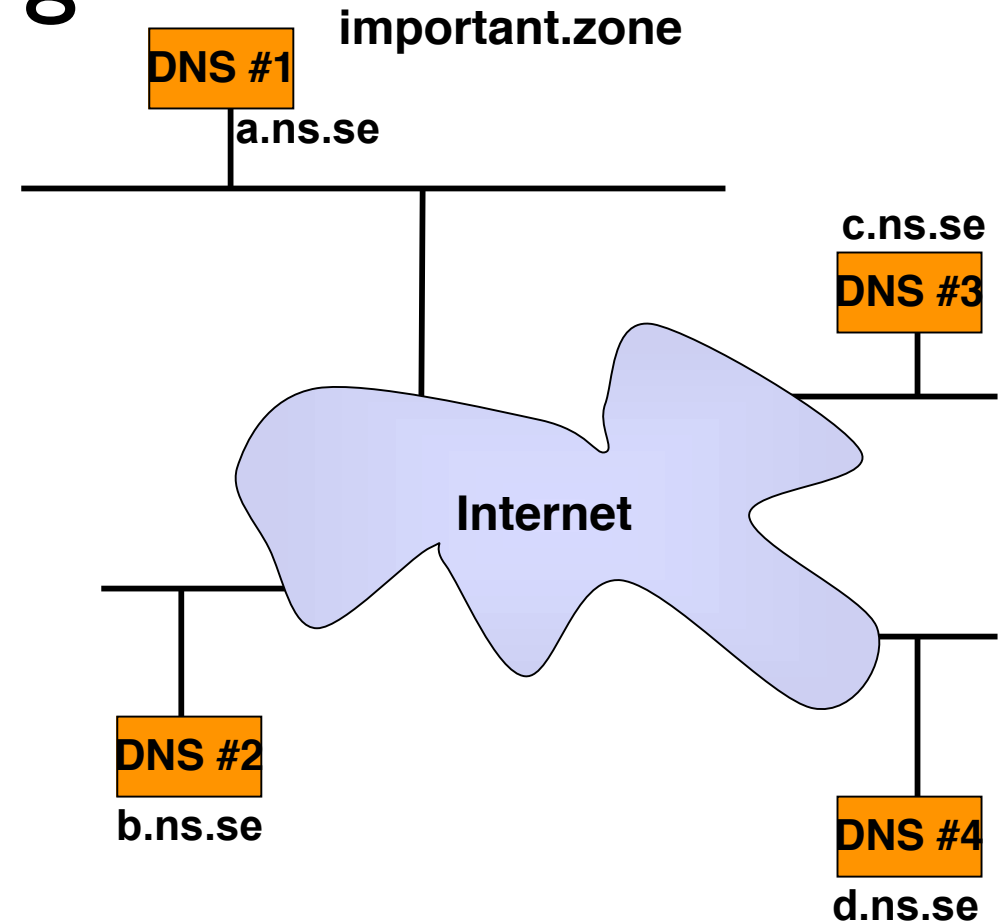
Design

- Obvious solution to redundancy problem is to separate the authoritative servers as much as possible
 - different locations
 - different transit providers
 - different IP prefixes
 - different service providers
 - (different platforms)
 - etc



Design #2

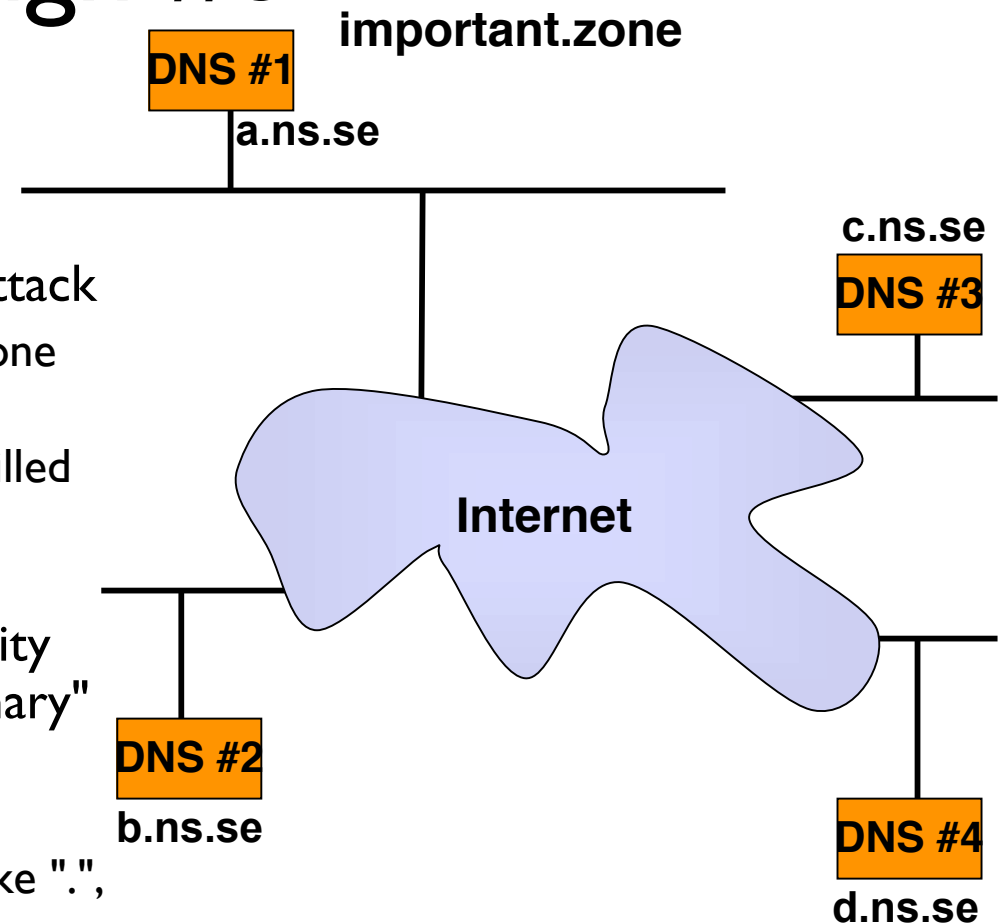
- Obvious solution to redundancy problem is to separate the authoritative servers as much as possible
 - different locations
 - different transit providers
 - different IP prefixes
 - different service providers
 - (different platforms)
 - etc



- All this is obviously possible to achieve with "standard" DNS.


Design #3

- When is this redundancy not sufficient?
 - High perceived risk of DDOS attack
 - usually each "server" has only one connection to the network
 - this connection is quite easily filled up during an attack
 - Difficult to achieve sufficient geographic and topologic diversity with a smaller number of "ordinary" (unicast) servers
 - probably mostly a concern for zones with truly global usage like ".", "com", "in-addr.arpa", etc.



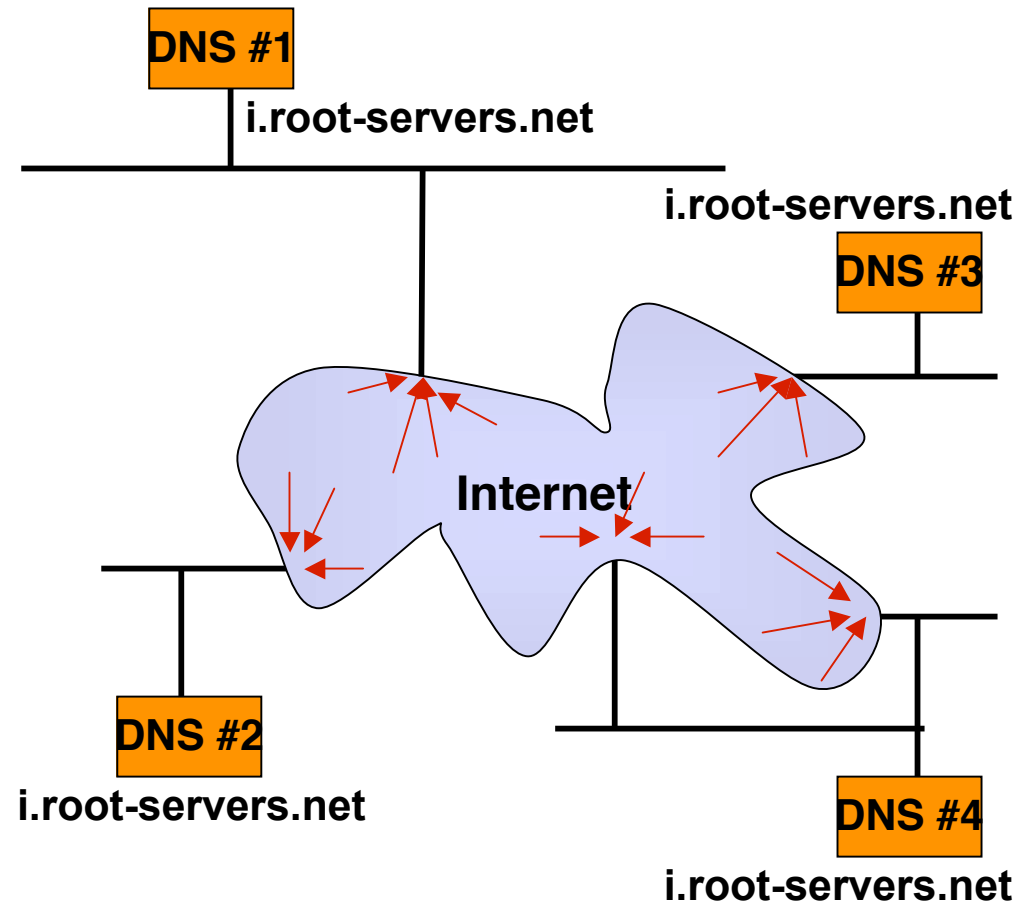
Anycast #1

- An advantage with anycast is that :
 - the number of servers and available bandwidth can be increased in a way that is invisible to the "DNS layer"
 - i.e. from a DNS perspective this is just one single server regardless of the number of physical machines and the aggregate amount of bandwidth and server capacity



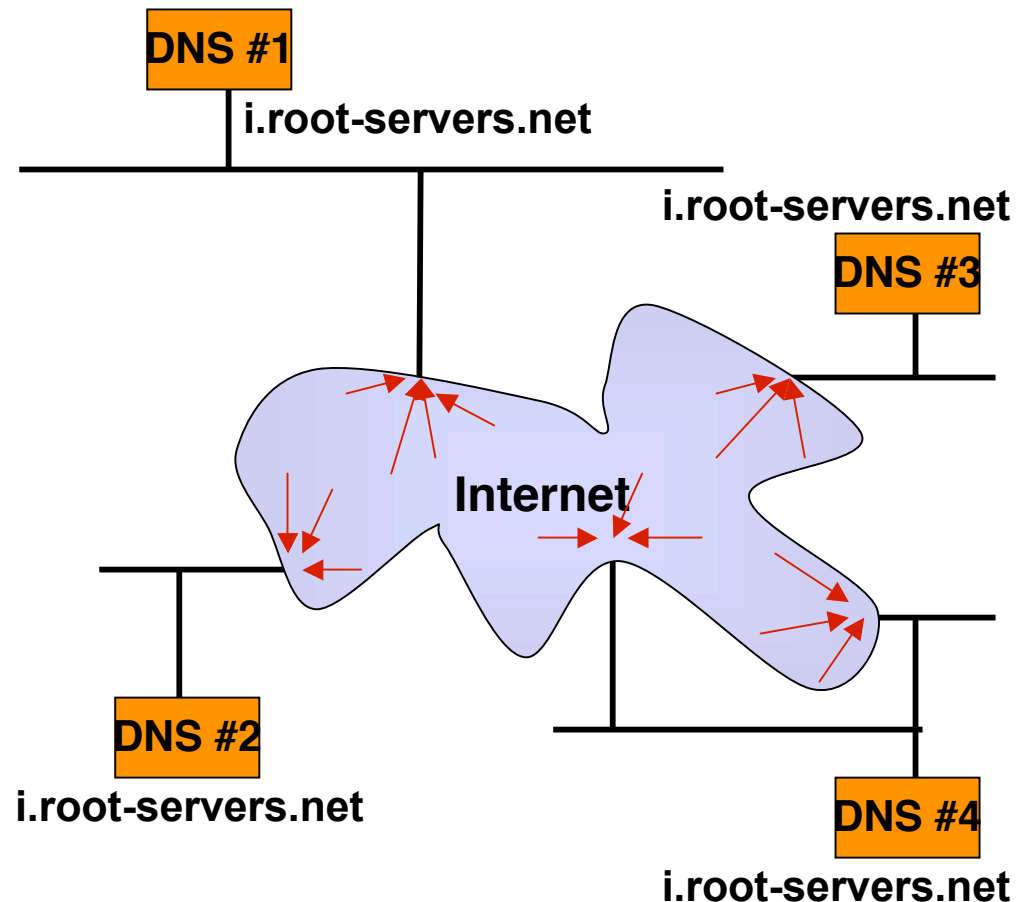
Anycast #2

- Another advantage with anycast is:
 - traffic is automatically "localized" to the closest instance of the service
 - "selection of optimal server" is moved from "application" (DNS) to "transport" (routing) layer
 - raises the barrier for a global attack significantly, since "DDOS armies" rarely are evenly distributed across the entire Internet



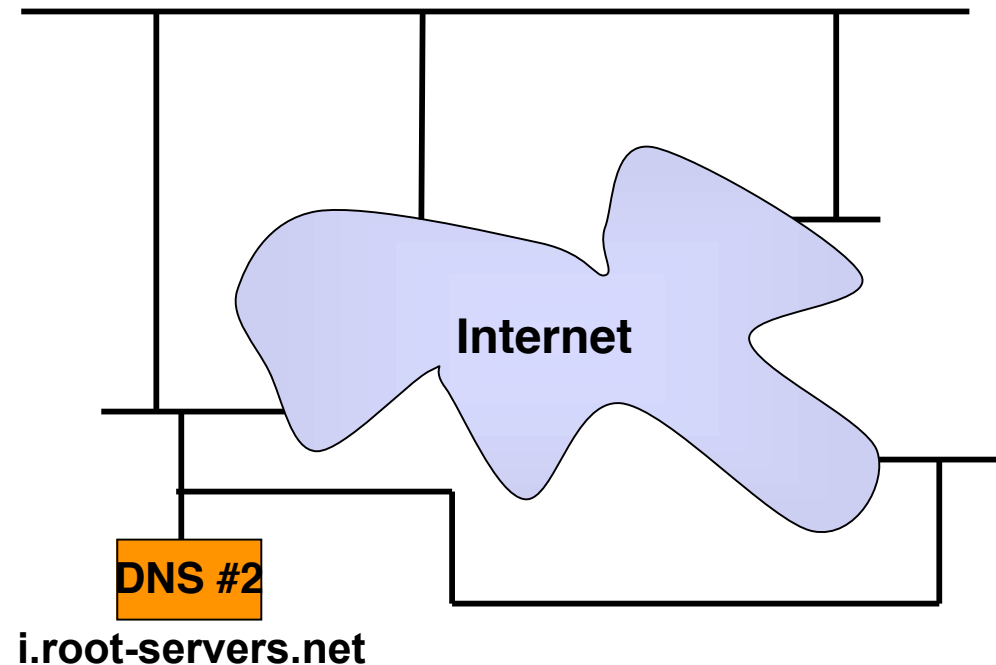
Anycast #3

- A third advantage with anycast is that the "service" is improved:
 - lower latency (i.e. shorter distance and therefore shorter time delay from server to user)
 - automatic fail-over on errors (i.e. if one "site" goes away for whatever reason the traffic will automatically be directed to the other sites)
 - automatic load sharing



Anycast

- From "the outside" a constellation of anycast servers look exactly like a single server that is "multi-homed", i.e. connected to the Internet via multiple connections.
 - Multi-homing is "known technology"
 - usually ISPs are connected to each other in multiple places
- From "the inside" the main difference is that there usually is no internal connectivity between the different points of connection.
 - i.e. the servers are individual "embassies" rather than part of the connected "country"



Anycast Pros

- Pros
 - Better redundancy
 - Automatic fail-over
 - Automatic load sharing
 - Higher resistance to DDOS attacks
 - Lower latency
 - Short distance -> shorter delays
 - Lower jitter
 - Higher availability
 - Better odds in times of network partitioning
 - e.g. Taiwan earth quake

Anycast Cons

- Higher system complexity
 - (Usually a bad idea)
 - Need to weigh needs and cost
- Troubleshooting much more complex
 - Which server is causing problems?
 - How do we reach it?
- System and site design radically different for hardware on other side of planet compared to server room down the hall

Anycast of the DNS root servers

- The redundancy issue considered by the root server operators before the WTC attack 9/11 2001
 - Concerns that "anycast" would be considered "irresponsible"
- DDOS attack on DNS root name server system 2002 highlighted need for action
 - Anycast chosen as best available alternative
 - Agreed that initially only 'F' to experiment
 - in case of unforeseen consequences
 - Concluded to be responsible and well-functioning
 - several letters now actively deploying anycast

i.root sites



Autonomica TLD Anycast service

- Additional servers deployed at the anycast sites for TLD anycast service
 - To allow TLDs to take advantage of anycast
 - To help TLDs take advantage of the experience Autonomica has built up with anycast through i.root-servers.net
 - A way to do cost recovery for the installed hardware
- TLD anycast delivered of same high quality of service as for i.root-servers.net
- Autonomica co-wrote BCP for operation of Anycast Services (RFC 4786)

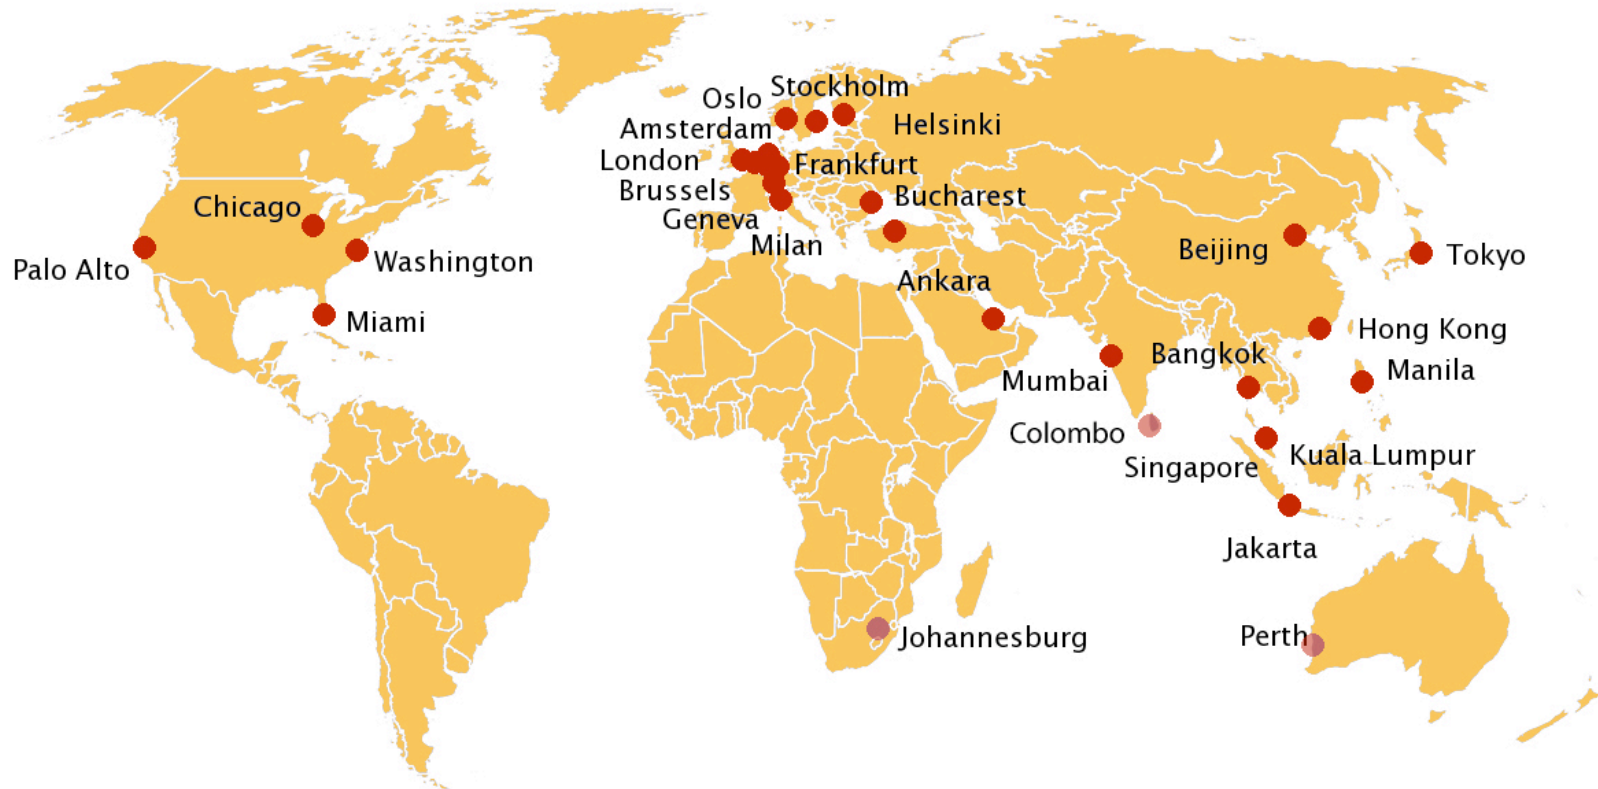


Fee model for TLD anycast

- Netnod / Autonomica is a non-profit organisation
 - Any surplus poured back into organisation
 - Research, testing
 - Improvement of services etc
 - Community participation, standards development
 - TLD anycast income helps fund the operation of i.root-servers.net
 - This way, those who mostly benefit the i.root-servers.net service contributes to the operation of the it
 - Fees not set to make a profit
 - Fees reduced over the board in 2008
 - As more TLDs come on-board, fees will continue to decrease

Current TLD anycast footprint

24 (27 shortly)



New sites underway in Perth, Colombo and Johannesburg

Autonomica anycast cloud

- The management infrastructure behind a commercial grade anycast cloud *will* be complicated (20+)
 - Sites need to have a good spread at a decent distance from each other
- Autonomica operates a mix of sites with great geographical and topological spread
 - Both at edges of Internet and a number of select core locations



Where are we now?

- A mature stable service
 - 100% uptime guarantee over 3 regions (EMEA 3, Asia 2, Americas 2)
 - ~17 TLDs
 - From small TLDs (<50k delegations) to several large TLDs (IM delegations)
 - With and without DNSSEC (.se)
- Growing requirements for more detailed stats
 - Stats per site & traffic analysis
 - Autonomica working on fine tuning stats tools to meet customer needs
- Growing privacy concerns
 - From TLDs & EU privacy legislation
 - Autonomica ensures the data isn't leaked
- Growing interest in DNSSEC
 - Autonomica has vast experience in DNSSEC



DNSSEC

- Many false starts
 - Think we finally got it right
- Autonomica has provided DNSSEC production service since 2005
 - .SE was the first TLD to deploy DNSSEC
 - Provided more than 100.000.000.000 responses from signed zones
- Implementations tend to need even more false starts than protocols
 - We believe in cooperation and active contribution to the technical community & standards development

Do you need anycast?

- Evaluate your current risks, costs and benefits
 - DDos attacks
 - Low value, high visibility
 - High value, high visibility
 - High value, low visibility
- It's hard to compare anycast services
 - Talk to your fellow ccTLDs about their experience
 - Look at footprint, operational experience, stability, additional services, cost, SLA etc.

Summary

- Global presence, Industrial scale
 - 24 sites (and growing)
 - 40+ M Resource Records
 - Authoritative DNS service for more than 100 zones (anycast + unicast)
- Extensive experience
 - Anycast production since 2003
 - TLD anycast provided at same QoS as i.root
 - Production DNSSEC since 2005
 - Active participation in technical community
- Non profit
 - TLD revenue contributes to the operations of i.root
 - Commitment to further lower prices as customer base grows



Terimah kasih

nurani@autonomica.se

