



# PCH ANYCAST AT .HU

Csillag Tamás - DNS Services engineer

tom(a)pch.net

HUNOG 2023

## Outline

- whoami?
- about PCH?
- anycast technology
- plans to become a secondary for .HU zone
- plans about establishing a PoP at BIX

## **whoami**

- worked at a university
- at one of the 5 biggest banks
- nic.hu / ISZT
- now: PCH - DNS Services Engineer



## About PCH

- Packet Clearing House (PCH) is the global intergovernmental treaty organization responsible for providing operational support and security to critical Internet infrastructure, including Internet exchange points and the core of the DNS, since 1994.



## About PCH

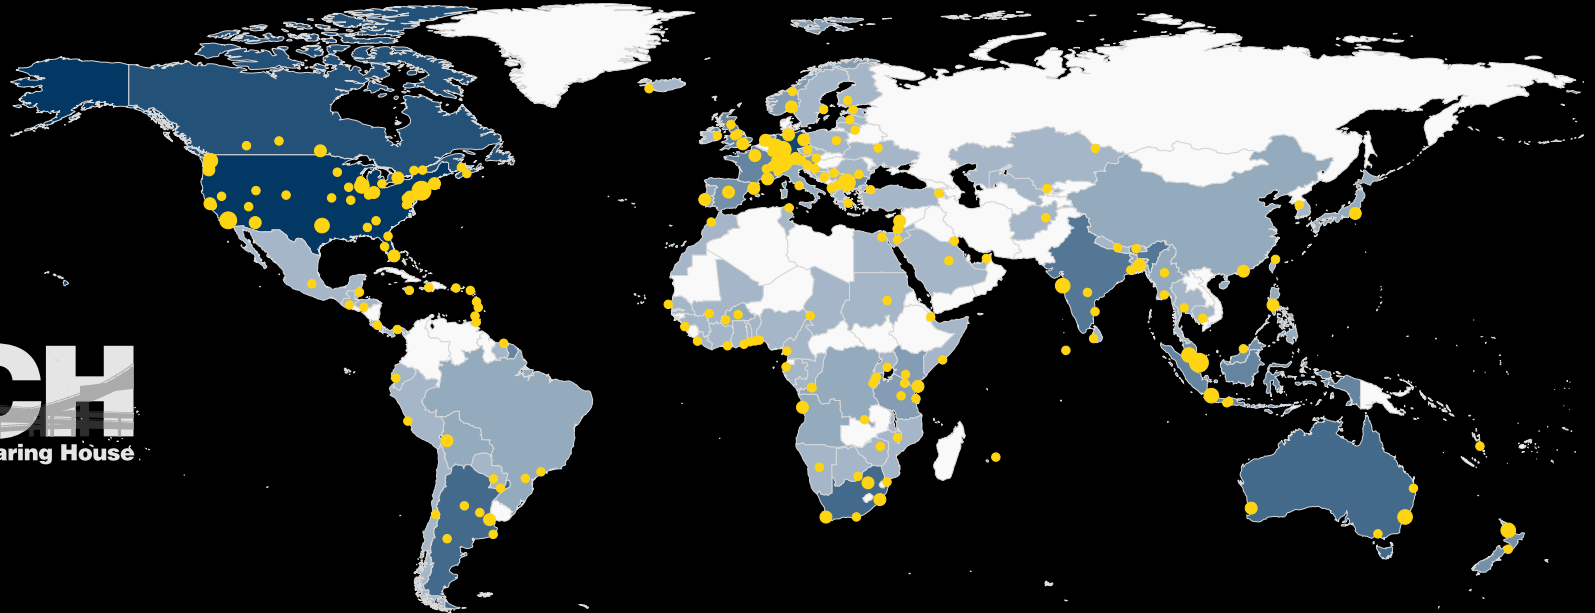
- Providing DNS infrastructure for 2 root server administrators
  - instance D - University of Maryland
  - instance E - NASA
- about 400 TLDs including 130 ccTLDs
- the Quad9 recursive resolver
- ipv4: 9.9.9.9, ipv6: 2620:fe::fe

## **Anycast technology**

- Anycast is a *special network configuration* that is used to have the same IP address online at various locations, at the same time.
- Via BGP, the Internet's global routing system routes users to the anycast node that is closest (topographically) to them.
- Anycast is a popular especially with DNS.

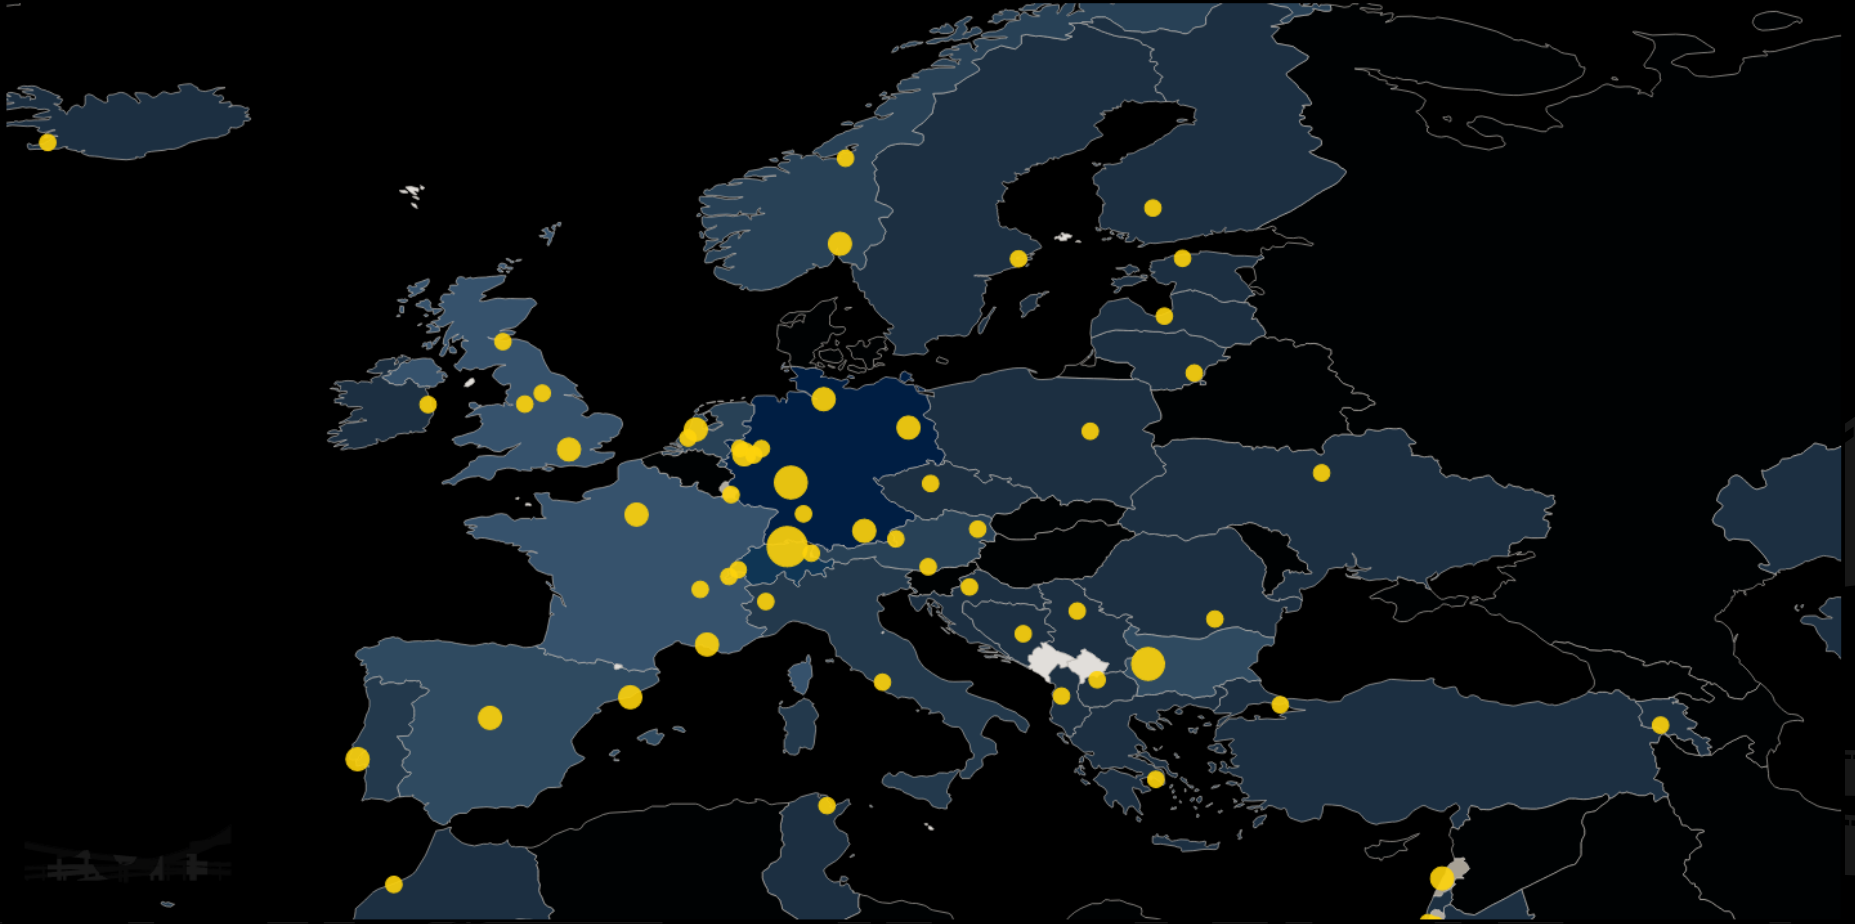
# Locations

- As of September 2023, PCH has Points of Presence (PoPs) in 283 Internet Exchange Points in 125 countries.



**PCH**

# Locations EU





# RTT, traceroute for PCH anycast

[atlas.ripe.net/measurements/59890236/](https://atlas.ripe.net/measurements/59890236/)

atlas.ripe.net/measurements/59890236/#probes

Settings & Status | Latest Results | Map | Tracemon | IPMap | Downloads

Probe	ASN (IPv4)	ASN (IPv6)		Time (UTC)	
1005586	63949	63949	🇺🇸 🟢	2023-09-12 09:37	0
1004284	27928		🇫🇷 🟢	2023-09-12 09:37	0
1005920	208871		🇫🇷 🟢	2023-09-12 09:37	0
1006019	34927		🇩🇪 🟢	2023-09-12 09:37	1
11717	35833		🇩🇪 🟢	2023-09-12 09:37	1
1004280	265867		🇫🇷 🟢	2023-09-12 09:37	1
62784	33885		🇸🇪 🟢	2023-09-12 09:37	1
1004498	10004		🇯🇵 🟢	2023-09-12 09:37	1
1006102	263781		🇮🇪 🟢	2023-09-12 09:37	1
33572	49036		🇩🇪 🟢	2023-09-12 09:37	1
13593	2108		🇫🇷 🟢	2023-09-12 09:37	1
55321	25180		🇬🇧 🟢	2023-09-12 09:37	2
53700	15435		🇯🇵 🟢	2023-09-12 09:37	2
11463	553		🇯🇵 🟢	2023-09-12 09:37	2
1000816	701		🇺🇸 🟢	2023-09-12 09:37	4,123
10048	60377		🇬🇧 🟢	2023-09-12 09:37	4

**ATLANTAIX**

**MITENE.AD.JP**

Latest Traceroute Result for Measurement #59890236

2023-09-12 09:37 UTC

Traceroute to 204.61.216.4 (204.61.216.4), 48 byte packets

1	172.105.134.1	gw-li1940.linode.com	AS63949	0.301ms	0.324ms	0.211ms
2	10.204.64.8			0.438ms	1.419ms	0.397ms
3	10.204.32.12			28.364ms	25.23ms	3.514ms
4	74.207.239.106	lo0-0.gw4.at1.us.linode.com	AS63949	0.585ms	0.571ms	2.914ms
5	198.32.132.37	atlantaix-fe01.woodynet.net		2.49ms	1.343ms	2.038ms
6	204.61.216.4	anyns.pch.net	AS42	0.54ms	0.468ms	0.562ms

Latest Traceroute Result for Measurement #59890236

2023-09-12 09:37 UTC

Traceroute to 204.61.216.4 (204.61.216.4), 48 byte packets


1	157.101.128.185	Sb9.N80659D.rppp.jp	AS10004	0.775ms	0.704ms	0.739ms
2	210.185.159.194	185-159-194.cs1501.b801.rppp.jp	AS10004	1.071ms	1.018ms	1.776ms
3	211.133.159.133		AS10004	0.649ms	0.532ms	0.69ms
4	211.133.145.154	T9a.V9185d3.phoenix-c.or.jp	AS10004	0.852ms	0.89ms	0.883ms
5	218.219.160.197	AS17961		1.027ms	0.967ms	1.067ms
6	202.238.56.100	v90.c76mibbn.mitene.ad.jp	AS17961	1.253ms	0.914ms	1.312ms
7	***					
8	204.61.216.4	anyns.pch.net	AS42	1.493ms	1.42ms	1.505ms

# DNSmon HU testing

## DNSMON

DNS responses for

Protocol:  Servers:



## Root servers RTT from HU

There are 13: a, b, c ... m  
in milliseconds

e.root-servers.net	:	3.62	3.67	3.80
f.root-servers.net	:	3.79	4.05	3.56
c.root-servers.net	:	5.85	5.96	6.13

^ these are at BIX

## Root servers RTT from HU

l.root-servers.net	:	11.8	13.1	37.3
h.root-servers.net	:	18.1	18.2	18.7
i.root-servers.net	:	18.7	20.5	21.1
k.root-servers.net	:	20.6	22.0	31.6

...

^ K is at BIX, routing issue?

## Root servers availability from HU

- common locations...
  - in the US
  - Frankfurt
  - Amsterdam
  - France - multiple locations
  - London

## **In progress #1**

- finalizing a contract with the registry
- to serve .HU zone on the anycast

## In progress #2

- with help from the registry
- ...to deploy a node at BIX (Dataplex)
- which would get all the zones on the anycast locally available
- 2 root DNS servers
- Quad9 resolver

## the end

- questions?
- give feedback
  - to me - tom at pch.net
  - to the registry - HU nyilvántartó
  - [domain.hu/impreszum](http://domain.hu/impreszum)



extra slides



## Deployments

- small, medium and full cluster installations
- routing vendor redundancy: Cisco, FRR and Quagga
- Cisco servers with hardware specs based on deployment types
- VMware ESXi clusters, supporting any x86\_64 OS
- hosted servers fully integrated with BGP routing architecture

## Deployments

- os redundancy: CentOS and Debian (new installs)
- name server redundancy: nsd, bind9, dnsmdist
- long-term strategic relationships with all involved vendors
- Cisco, AMD, VMware, CZ.NIC, ISC, NLnet Labs, PowerDNS

## Monitoring

- multiple layers of monitoring to proactively detect issues that could be leading to a degradation of the service
- hardware layer: CPU, temperature, RAM, Disks
- interconnection layer: ports and traffic levels
- routing layer: AS-PATH and prefix announcements
- service layer: queries per second, replies per second

# Monitoring

- passive monitoring tools
- Cacti/Nagios/Prometheus/LibreNMS with custom plugins for DNS and DNSSEC
- netflow monitoring traffic levels
- active monitoring of global performance using RIPE Atlas and RIPE DNSMon measurements on a regular basis