

# **Global Anycast Hosting Platform for ccTLD, in-addr & Root DNS**

**Version 2.2**

**June, 2008**

**Bill Woodcock**

**Packet Clearing House**

## **PCH Background**

Not-for-profit private-sector agency supporting operation of critical Internet infrastructure, including the core of the DNS and Internet exchange points, since 1993.

Funded by grants and service-provision fees from the Internet operations industry.

Eleven engineering and three administrative staff in offices in London, Wellington, San Francisco, Minneapolis, and Kathmandu.

# Anycast Background

PCH and its precursors have run **production anycast services since 1989.**

Bill Woodcock, of PCH, and Mark Kosters, of Verisign, first proposed the idea of anycasting authoritative root and TLD DNS at the Montreal IEPG in 1995.

PCH has been operating production anycast for ccTLDs and in-critical infrastructure in-addr's since 1997, with **100% up-time over more than eleven years.**

PCH first hosted the production anycasting of a root nameserver in 2002.

# Anycast Technology

An **anycast cloud** is a **distributed cluster** of identical **instances** of a server, each typically containing identical data, and capable of servicing requests identically.

Each instance has a regular unique globally routable IP address for management purposes, but...

Each instance also **shares an IP address in common** with all the others.

The Internet's normal global routing system routes every query to the **instance** of the anycast cloud that's **closest** to the user who originated the query.

# PCH Anycast Servers

PCH operates the world's largest, oldest, and most continuously-available anycast server cloud.

PCH has built or designed many of the other large anycast clouds on contract.

Our servers are used to provide ccTLD, root DNS, and in-addr DNS slave service.

There are no fees for the use of this service by small ccTLDs. We negotiate pricing for large ccTLDs on a case-by-case basis. **We have never failed to negotiate a mutually acceptable cost.**

# PCH 6th-Generation Architecture

Routing vendor redundancy: Cisco and Quagga.

Clusters of Sun X2200 servers, currently eight-core, 16GB RAM, upgrading to 64GB RAM as the price of 4GB SIMMs drops.

Each site can answer 300,000 queries per second.

VMware ESX clusters, supporting any X86 32-bit or 64-bit OS.

Hosted servers fully integrated with BGP routing architecture.

OS redundancy: Solaris and CentOS.

DNS redundancy: BIND and NSD.

Long-term strategic relationships with all involved vendors: Cisco, AMD, Sun, VMware, ISC, and NLNet Labs.

# PCH Server Footprint



46 Locations as of June, 2008

# How to Use the Anycast Service

Contact [noc@pch.net](mailto:noc@pch.net) to receive an anycast IP address.

Define a shared TSIG key, if desired.

Permit IXFR and AXFR between your hidden masters and ours.

Observe transfer and availability statistics during a trial period.

Tell us what kind of contractual agreement you want.

Register the new IP address with the IANA.

Find and use other anycast clouds... they're not mutually exclusive!



# Thanks, and Questions?

Copies of this presentation can be  
found in PDF and QuickTime formats at:

**[http:// www.pch.net / resources / papers / anycast-services](http://www.pch.net/resources/papers/anycast-services)**

Bill Woodcock  
Research Director  
Packet Clearing House  
**woody@pch.net**



Produced with the generous support of the  
ICANN Training and Outreach Program