Does the .au registry use Anycast

Yes, the .au top-level registry utilizes Anycast DNS infrastructure to guarantee global availability and speed.

The .au Domain Administration (auDA) and its contracted registry operator, Identity Digital Australia, manage the core authoritative nameservers for `.au` and its second-level domains (such as `.com.au` and `.net.au`). They distribute their nameserver footprint across several distinct Anycast clouds to minimize latency and withstand DDoS attacks.

How the .au Registry's Anycast Nodes are Distributed

Official transparency logs from auDA Registry Monthly Stats indicate that the registry measures its uptime and performance across distinct Anycast networks, which include:

- **a.au:** Managed directly by auDA.
- **c.au and d.au:** Handled via Cloudflare's massive global Anycast network.
- **m.au:** Maintained through GoDaddy Registry infrastructure.

By assigning the authoritative nameservers to these global Anycast providers, a query for a `.au` website from anywhere in the world will automatically route to the topologically nearest physical data centre, eliminating the need for traffic to travel back and forth across oceans just to resolve a domain name.

Timeline of BGP Anycast

BGP Anycast evolved from an experimental load-balancing concept in **1989** into the foundational routing architecture powering modern DNS, Content Delivery Networks (CDNs), and cloud edge infrastructure. Rather than requiring a dedicated protocol, it works by announcing the same IP prefix from multiple global locations via the Border Gateway Protocol (BGP).

Origins & Early Standardization (1989–1997)

The foundational rules for BGP and IP Anycast emerged during the rapid commercialization of the early internet.

- **1989 (First Documented Use & BGP Debut).** Engineers documented the earliest use of anycast routing for topological load-balancing of internet services. In the same year, BGP was introduced to replace the older Exterior Gateway Protocol (EGP).
- **1993 (Formal IETF Definition).** Craig Partridge, Trevor Mendez, and Walter Milliken published RFC 1546 (Host Anycasting Service), formally defining the anycast service model. While the RFC originally envisioned a dedicated anycast address class, operators discovered that announcing standard IP prefixes via BGP achieved the same goal without protocol changes.
- **1997 (First Root DNS Experimentation).** The Internet Systems Consortium (ISC) began early anycast deployments for the F-root DNS nameserver, testing how BGP routing could distribute DNS query loads.

Critical Infrastructure Adoption (2001–2008)

Following major Distributed Denial-of-Service (DDoS) attacks against the DNS root system in October 2002, network operators rapidly adopted BGP Anycast to improve global resilience.

- **2001–2003 (Root Nameserver Rollouts).** Critical infrastructure began deploying anycast at scale, highlighted by Netnod launching multi-instance anycast for the I-root nameserver in August 2003.
- **2006 (Operational Best Practices & IPv6 Updates).** The IETF published RFC 4786 (Operation of Anycast Services), establishing standard guidelines for route advertisements, service health signaling, and security. That same year, RFC 4291 updated how IPv6 addresses reserve subnet anycast identifiers.
- **2006–2008 (Widespread Root DNS Coverage).** Multiple additional root operators—including C, F, J, K, and M—completed anycast rollouts, proving that stateless UDP services could scale across continents using a single IP address.

Expansion to Stateful Traffic & CDNs (2009–2019)

Early network engineers worried that BGP route fluctuations would break long-lived TCP connections mid-session. Over this decade, providers solved those routing stability challenges and brought anycast to HTTP/HTTPS web traffic.

- **2009–2010 (Public Anycast DNS Resolvers).** Major tech companies launched global anycast recursive resolvers, including Google Public DNS (`8.8.8.8`), bringing low-latency anycast routing directly to everyday consumers.
- **2010–2014 (TCP Anycast & Modern CDNs).** Cloud and security providers like Cloudflare proved that BGP Anycast could reliably serve stateful TCP and HTTP traffic across hundreds of data centres. In 2014, the IETF published RFC 7094 (Architectural Considerations of IP Anycast), documenting best practices for both stateless and stateful anycast applications.
- **2017–2019 (Domain Infrastructure Consolidation).** Internet measurement studies showed rapid anycast adoption across Top-Level Domains (TLDs) and Second-Level Domains (SLDs), driven by major cloud DNS providers.

Modern Edge & Cloud Era (2020–Present)

Today, BGP Anycast serves as the default traffic-steering layer for global cloud platforms, DDoS mitigation networks, and mobile edge computing.

- **2021 (Majority DNS Adoption).** Research from the APNIC Blog revealed that over 97% of Top-Level Domains and more than 62% of responsive Second-Level Domains relied on anycast nameservers.
- **2024–Present (Massive Global Scale).** The 13 logical DNS root server addresses now map to well over **1,750 to 1,950 physical server instances** worldwide. Modern architectures also combine BGP Anycast with Equal-Cost Multi-Path (ECMP) routing directly on server hosts and 5G mobile edge networks to achieve millisecond-level failover and ultra-low latency.

** End of Page