CDN and Anycast

With a CDN (Content Delivery Network) does the IP address on each packet use Anycast on TV streaming

Yes, the IP address on TV streaming packets can use **Anycast** via a Content Delivery Network (CDN), though it depends heavily on how the CDN and the streaming provider are architected.

How Anycast Works in TV Streaming

- **Shared IP Address:** With Anycast, multiple geographically distributed CDN edge servers advertise the exact same public IP address to the internet via Border Gateway Protocol (BGP).
- **Topological Routing:** When your TV or streaming device sends a packet to that IP address, internet routers automatically send it to the "closest" edge node based on network path and hop count, rather than a specific physical server with a unique unicast address.
- **Initial Connection & DNS:** CDNs frequently use Anycast for initial handshakes, API requests, manifest/playlist fetching, or DNS resolution so that users are seamlessly handed off to a nearby data center.

Why CDNs Often Mix Anycast and Unicast for Video

While Anycast is great for finding a nearby entry point and absorbing traffic spikes or DDoS attacks, heavy video streaming (like live TV or 4K video) introduces specific challenges:

- **Route Flapping & Stability:** During a long streaming session, if internet routing paths shift mid-stream, your subsequent packets might land on a *different* edge server advertising that same Anycast IP, which can cause connection resets or buffering.
- **Session and Origin Control:** Many large streaming services use **Unicast** (where each edge server has a unique IP address) or hybrid approaches after the initial setup. Once a client requests a video stream, the CDN often redirects the client to a dedicated unicast server or specific node to maintain a stable, long-lived TCP/QUIC connection for heavy chunk/segment delivery.

TCP vs QUIC

QUIC significantly outperforms TCP for TV streaming by eliminating connection delays and preventing head-of-line blocking. While TCP remains a widely supported foundational protocol, modern Content Delivery Networks (CDNs) and major streaming platforms increasingly use QUIC (which underlies HTTP/3) to deliver a smoother, buffer-free viewing experience.

Direct Comparison: TCP vs. QUIC

Feature	TCP (Transmission Control Protocol)	QUIC (Quick UDP Internet Connections)
Underlying Protocol	Base layer protocol (Layer 4)	Built on top of **UDP** (User Datagram Protocol)
Connection Setup	Requires a **3-way handshake** (plus an extra TLS handshake for HTTPS)	Combined **0-RTT or 1-RTT** connection and TLS 1.3 encryption handshake
Head-of-Line Blocking	**Present.** If one packet is lost, the entire stream stalls until it is resent.	**Eliminated.** Multiple streams run independently; a lost packet only delays its specific stream.
Connection Migration	**Poor.** Tightly bound to the device's IP address. Changing networks drops the connection.	**Seamless.** Uses a unique Connection ID. Smoothly transitions from Wi-Fi to mobile data without dropping.
Congestion Control	Built into the operating system kernel (harder to update/tune).	Implemented in user space (rapidly updatable and highly customisable by CDNs).

Key Architectural Differences

1. Connection Setup Time (Latency)

TCP requires a multi-step sequence to establish a secure connection, forcing multiple round-trip times (RTT) before any actual video data is transferred. QUIC combines the transport and cryptographic handshakes into a single step, allowing video chunks to start downloading almost instantly.

2. Head-of-Line (HoL) Blocking

When streaming video over TCP, all video segments travel down a single, strict queue. If one packet is dropped on a congested network, TCP stops processing everything behind it. The viewer's player starves of data, resulting in a **buffering spinner**, even if the subsequent data packets have already arrived at the device.

QUIC solves this by multiplexing. It splits the data into independent streams. If a packet containing video segment A is lost, QUIC requests a retransmission for that specific packet while simultaneously continuing to deliver video segment B to the player.

3. Network Migration (The "Living Room to Mobile" Scenario)

TCP identifies a connection using a 4-tuple: source IP, source port, destination IP, and destination port. If you start watching a stream on your phone via home Wi-Fi and step outside onto a 5G network, your IP address changes. The TCP connection breaks instantly, forcing the streaming app to negotiate a brand-new connection from scratch.

QUIC uses an abstract **Connection ID** independent of your IP address. When your network switches, QUIC simply sends a packet with the same Connection ID from the new IP, and the CDN edge server keeps streaming video chunks without a single frame dropping.

Cloudflare, Akamai, AWS CloudFront Are they the big three in Australia CDNs

Yes, **Cloudflare, Akamai, and AWS CloudFront make up the "Big Three" CDNs in Australia**, but they dominate the market in completely different ways.

While they all maintain heavy infrastructure footprints across major Australian internet exchange points (primarily in Sydney, Melbourne, Brisbane, Perth, and Adelaide), they cater to distinct tiers of the market:

How the "Big Three" Split the Australian Market

CDN Provider	Market Dominance Style	Typical Australian Customer Base	Strengths in the Local Market
Cloudflare	**The Volume Leader** (by sheer website count)	Small-to-medium businesses (SMBs), tech startups, and major digital platforms.	Massive network density across regional Australia, exceptional DDoS mitigation, and a developer-friendly serverless edge (Cloudflare Workers).
Akamai	**The Enterprise & Revenue Leader**	Major Australian banks, federal/state government agencies, and massive TV streaming networks (e.g., live sports broadcasting).	Deeply entrenched ISP partnerships inside Australia (caching servers placed directly *inside* Telstra, Optus, and TPG networks), unmatched scale for heavy media streaming.
AWS CloudFront	**The Cloud Integrator**	Any business running its core infrastructure inside the AWS Sydney (`ap-southeast-2`) or Melbourne (`ap-southeast-4`) regions.	Seamless integration with AWS services (S3, EC2, ALB). Zero data egress fees out of AWS when caching via CloudFront, making it highly cost-effective for AWS-heavy architectures.

The Challenger: Fastly

While not part of the traditional "Big Three", **Fastly** is a major player in Australia. They are highly favoured by **premium media publishers and dynamic tech companies** (like Reddit or Kayo/Canva-style architectures) because of their instant cache invalidation (purging old content globally in under 150 milliseconds) and highly customisable Varnish Configuration Language (VCL).

Why the Distinction Matters for TV Streaming

Tying this back to **Anycast vs. Unicast** and **QUIC vs. TCP**:

- **Cloudflare** heavily champions an Anycast-first architecture and has been a massive driver of HTTP/3 and QUIC adoption across the web.

- **Akamai** frequently uses a hybrid or multi-unicast approach for major live Australian sporting events (like the AFL, NRL, or cricket streams) to map users to deterministic, high-throughput edge blocks directly inside an ISP's network to guarantee zero stream degradation.
- **CloudFront** utilizes AWS's massive global private network, meaning once an Australian user's packet hits a local CloudFront edge location, it travels over Amazon's private fiber rather than the public internet.

Timeline of Anycast in Australia

The deployment of Anycast routing in Australia evolved alongside the growth of the internet, transforming from an experimental tool used to protect global DNS infrastructure into the foundation of modern content delivery.

Phase 1: The Root Server Pioneers (2002–2006)

Before Anycast, if an Australian device wanted to talk to an internet Root DNS Server, the packet had to travel across submarine cables to California.

- **2002–2003: The Wake-Up Call:** Following massive DDoS attacks on the global internet root server system, operators rapidly pivoted to Anycast.
- **2004: First Local Roots:** The F-Root (operated by ISC) and I-Root (operated by Netnod) deployed some of Australia's first Anycast instances in Sydney. Instead of packets crossing the Pacific, Australian ISPs could route root DNS queries to a shared IP address answered locally inside Australian borders.

Phase 2: Protecting the Nation's Core (2006–2010)

As the local internet footprint matured, Australia's own infrastructure required resilience.

- **2007–2009: The .au Namespace Moves to Anycast:** **auDA** (the .au Domain Administration) migrated the country's country-code top-level domain (ccTLD) nameservers to Anycast infrastructure. Rather than relying on a couple of unicast servers in Melbourne or Sydney, identical `.au` registry instances were mirrored across Australian capital cities. If an exchange point in Brisbane went down, the BGP routing table instantly swung traffic to Sydney without breaking domestic web traffic.

Phase 3: The Enterprise Edge & CDN Explosion (2010–2016)

Anycast shifted from a DNS-only protection mechanism to a tool for rapid content and website delivery.

- **2012: The Rise of Cloudflare:** **Cloudflare** aggressively built its global network on an Anycast-only architecture. They established major presences in Sydney and Melbourne, allowing millions of global websites to instantly route Australian traffic to local data centers using the exact same public IP addresses.
- **2014–2015: Internet Exchange Evolution:** Local internet exchanges like **IX Australia** and Megaport grew dramatically. CDNs and Anycast operators began "peering" directly with Australian retail ISPs (like Telstra, Optus, and TPG). Instead of a packet hopping through multiple networks, an Anycast IP could be reached in just one or two network hops from a user's home.

Phase 4: Video Streaming & Hybrid Maturity (2016–Present)

As heavy video streaming (Netflix, Stan, Kayo, and live sports broadcasting) took over Australian networks, Anycast evolved into its current state.

- **2016–2018: AWS and Akamai Native Integration:** **AWS CloudFront** heavily integrated Anycast into its global edge routing, meaning an Australian endpoint hitting an AWS API or video manifest file was captured at the nearest Australian edge route.
- **2020–2026: The Dynamic / Hybrid Pivot:** Modern video streaming architectures reached a consensus. Pure Anycast is used across Australia to handle the **DNS resolution, API authentication, and initial video manifests** (the `.m3u8` or `.mpd` files that dictate streaming chunks). However, once the video player knows what files it needs, the CDN dynamically hands the connection off to a localized, unicast IP address or specific cached server inside the ISP's network to prevent route flapping during heavy data streaming.

If you are interested in a deeper look, I can explain:

- How **Border Gateway Protocol (BGP)** determines which Australian capital city your Anycast packet lands in.
- The differences in how **Telstra vs. Optus** peer with Anycast networks.

Let me know what you would like to explore next!

Both

To understand how Anycast packets move across Australia, we have to look at the intersection of global routing protocols and the highly consolidated nature of the Australian telecommunications market.

Here is exactly how Border Gateway Protocol (BGP) dictates where your packets land and how Australia's two largest telcos—Telstra and Optus—handle that traffic differently.

Part 1: How BGP Decides Which Australian Capital City Your Packet Lands In

Border Gateway Protocol (BGP) is the postal system of the internet. It does not look at physical distance (kilometres); it looks at **network topology (Autonomous System hops)** and **routing policies**.

When a CDN uses Anycast, it advertises the exact same IP address from its routers in Sydney, Melbourne, Brisbane, Perth, and Adelaide. When you press play on a video stream, your router checks its BGP routing table to find the "shortest" path.

```
[ Your Home Router ]
         |
         v
  [ Your ISP Network ]
         |
         +---> Option A: CDN Router (Sydney)    -- BGP Path: 1 hop  <-- WINNER
```

```
        |
        +---> Option B: CDN Router (Brisbane)  -- BGP Path: 2 hops
```

BGP decides the winner using a strict hierarchy:

1. **AS-Path Length (The Shortest Route):** BGP prefers the route that crosses the fewest number of distinct networks (Autonomous Systems, or ASNs). If your ISP connects directly to the CDN in Sydney (1 hop) but has to pass through a transit provider to reach the CDN in Brisbane (2 hops), your traffic goes to Sydney—**even if you are physically sitting in Brisbane.**
2. **BGP Local Preference:** ISPs manually configure "Local Pref" values. They will always prefer a network path that costs them nothing (free public peering) over a path that costs them money (paid IP transit), regardless of geographic location.
3. **MED (Multi-Exit Discriminator):** If an ISP and a CDN connect to each other in multiple cities (e.g., Sydney, Melbourne, and Perth), the CDN can use MED attributes to tell the ISP's network, *"If you are in Western Australia, use our Perth exit; if you are in Victoria, use our Melbourne exit."*

Part 2: How Telstra vs. Optus Peer with Anycast Networks

This is where the theoretical design of Anycast collides with Australian commercial realities. Telstra and Optus have fundamentally different strategies for handling content delivery networks.

1. Telstra (AS4608): The Protective Incumbent

Telstra owns the largest domestic fiber footprint in Australia. Because they control the "last mile" to millions of Australian homes, they treat their network access as a premium asset.

- **Restrictive Peering Policy:** Telstra notoriously does not participate in open public peering at major Australian internet exchanges (like IX Australia or Megaport). They refuse to let CDNs hand over traffic for free at a public switch.
- **The "Paid Transit" or On-Net Model:** To ensure low latency and prevent Anycast packets from routing overseas, major CDNs (like Akamai or Cloudflare) must either buy paid IP transit from Telstra or enter into commercial agreements to place CDN caching servers **directly inside Telstra's private data centres** ("On-Net" nodes).
- **The Routing Result:** If a CDN has an On-Net arrangement with Telstra, an Anycast packet from a Telstra home internet user stays highly localized (e.g., a Brisbane user hits a Brisbane node). If the CDN does *not* pay Telstra, Telstra may route that Anycast packet through a circuitous path—sometimes sending a Queensland or Victorian user's packet all the way to Sydney or even Los Angeles/Singapore if that is where their cheapest external transit link sits.

2. Optus (AS4804): The Open Pragmatist

Optus historically takes a much more open, peer-friendly approach to lower its own network backhaul and transit costs.

- **Open Peering at Exchanges:** Optus peers extensively at major domestic internet exchange points. They happily interconnect with Cloudflare, AWS, Fastly, and Akamai over public fabrics.
- **Offloading Traffic Early:** Optus's primary goal is to get heavy video data off their long-distance internal network as fast as possible. By peering aggressively with Anycast CDNs in every capital city, they ensure that if an Optus customer in Perth requests a stream, the packet is handed over to the CDN's Perth Anycast node immediately, saving Optus from carrying that data across the Nullarbor Plain from Sydney.
- **The Routing Result:** Anycast tends to behave much more "geographically predictably" on Optus. Packets usually hit the closest physical CDN edge node within your state, provided the CDN is present at the local capital city exchange.

Summary of the Streaming Impact

Because of these dynamics, video streaming providers rarely rely on a single CDN in Australia. If Akamai has an exclusive internal cache alignment with Telstra, but AWS CloudFront performs better over Optus, a streaming platform's architecture will dynamically shift your connection to the specific CDN that holds the most optimal, lowest-hop BGP path for your exact ISP and location.

How non-capital cities (like Cairns, Townsville, Rockhampton, Caloundra, Southport, Newcastle, or Geelong) get served by these CDNs

Serving regional and mid-sized Australian cities introduces a logistical challenge: it is incredibly expensive to build full-scale data centres in every regional hub, yet long-distance transit costs back to Sydney can heavily degrade video performance.

To overcome this, CDNs utilize a multi-tiered approach that separates the cities into two distinct groups based on how close they are to a state capital.

Category 1: The 'Backhaul' Cities (Cairns, Townsville, Rockhampton)

For deep-regional cities in North and Central Queensland, there are rarely major standalone public CDN data centres. Instead, traffic is handled through an integrated pipeline heavily reliant on Australia's **National Broadband Network (NBN)** architecture and **ISP deep-edge caching**:

- **The NBN Handoff (POIs):** The NBN terminates consumer connections at 121 physical Points of Interconnect (POIs) across Australia (e.g., the local NBN facility in Townsville or Cairns). Retail ISPs like Telstra or Optus take the traffic from that local POI and carry it back over their private long-haul fibre networks to the nearest capital city.
- **Akamai & Netflix 'Deep Edge' Boxes:** Because carrying 4K video streams 1,400 kilometres from Brisbane to Cairns is highly inefficient, providers like **Akamai** work directly with Telstra and Optus to place physical, high-storage caching appliances *directly inside* the regional exchange or aggregator nodes. If a popular TV show drops, it is downloaded to a local Telstra rack in Townsville **once**. When 5,000 people in Townsville watch it, they fetch it directly from that internal hardware layer, bypassing the long journey to Brisbane.
- **The Anycast Layer:** If you look up a website on **Cloudflare** from Cairns, your initial BGP route will almost certainly direct you to **Brisbane** (or occasionally Sydney, depending on the ISP's backhaul loop). Brisbane functions as the Anycast gatekeeper for the state.

Category 2: The 'Satellites' (Geelong, Newcastle, Southport, Caloundra)

For cities sitting on the periphery of major capital zones (like Geelong outside Melbourne, Newcastle outside Sydney, or Southport/Caloundra on the Gold and Sunshine Coasts), the network dynamic is entirely different:

- **Geographic Proximity Dominance:** Because a fibre optic signal takes only a few milliseconds to travel from Geelong to Melbourne, or from Southport to Brisbane, CDNs see no functional benefit in deploying local server racks in those mid-sized cities.
- **The Capital City Drain:** When an Anycast packet is fired from a smart TV in Southport, the BGP path length to **Brisbane** is identical to a path length inside Brisbane itself. The packet rides directly into the capital's massive open internet exchanges (like IX Australia or Megaport) and hits Cloudflare, AWS, or Akamai's major state footprints.
- **Why It Doesn't Buffer:** Because the physical distances are short and the backhaul pipe sizes between these satellite cities and capitals are massive, there is virtually zero latency penalty. Your TV in Newcastle will fetch video segments from Sydney just as quickly as a TV located in downtown Sydney.

On 1st July 2026 the variable, shared-bandwidth charge (**CVC**) that caused peak-hour congestion was phased out entirely, moving the entire network to a flat-rate **AVC-only model**.

The transition completely alters how regional areas get served and why the billing model changed:

The Old Model vs. The New Model

Feature	**The Old Model (AVC + CVC Mix)**	**The New Model (AVC-Only Flat Rate)**
AVC (Access Virtual Circuit)	A fixed monthly fee the ISP paid NBN Co for your individual connection line.	A single, flat-rate wholesale monthly fee that **now completely bundles your connection line and your data usage**.
CVC (Connectivity Virtual Circuit)	A variable capacity "pipe" the ISP had to buy and share among your neighbourhood. **More usage = higher costs for the ISP.**	**Abolished entirely for all non-satellite fixed-line and wireless tiers.**

The Timeline of the Shift

The transition rolled out in two major phases following the ACCC's approval of NBN Co's new regulatory framework:

1. **December 2023:** CVC charges were completely removed for all high-speed tiers (**100 Mbps and above**).
2. **1 July 2026:** CVC charges were **fully abolished for the remaining entry-level tiers** (12 Mbps, 25 Mbps, and 50 Mbps plans).

How This Changes Regional Streaming

Under the old model, if thousands of households in Cairns or Townsville sat down to stream a live footy match at 8:30 PM, their data usage would spike. If their ISP hadn't purchased enough expensive CVC capacity for that local POI, the shared pipe would max out, forcing the stream to drop from 4K down to blurry standard definition or buffer.

Now that **CVC has been abolished**, ISPs no longer face a financial penalty when regional users consume massive amounts of video data at peak times. Because the wholesale price is a flat-rate AVC fee, the artificial capacity bottlenecks at regional POIs have been removed, meaning video performance in regional cities is dictated entirely by physical infrastructure and the ISP's global routing, rather than budget penny-pinching.

Now that the wholesale model has fully shifted to flat-rate AVC pricing, would you like to see **how retail plan prices have adjusted** for the major tiers, or look closer at **how ISPs advertise their typical evening speeds** under this new framework?

Both

The abolition of CVC fees has fundamentally reshaped how retail internet is priced and advertised in Australia. Because internet service providers (ISPs) no longer face variable capacity penalties, the business model has completely shifted toward a flat-rate framework.

1. How Retail Plan Prices Have Adjusted

The elimination of variable bandwidth charges removed the single biggest variable cost for ISPs. However, under the ACCC-approved framework, **NBN Co** adjusts its wholesale flat-rate AVC pricing annually based on an inflation-linked index. Following the final CVC removal for lower tiers, retail prices have settled into distinctive brackets:

- **Slower/Entry Tiers Increased:** Under the old model, entry-level plans (NBN 12 and NBN 25) were subsidized by low data usage. Because NBN Co increased the base flat-rate wholesale cost (e.g., NBN 25 wholesale price up by **$1.08/month** and NBN 50 up by **$2.34/month**), major telcos like Telstra, Optus, and Belong have bumped their base retail plans up by **$2 to $5 per month**.
- **Premium Tiers Compressed:** High-speed tiers (NBN 100, NBN 250, and NBN 1000) have become significantly cheaper on a relative basis. Because there are no data usage penalties, providers like Superloop, Buddy Telco, and Tangerine are aggressively offering ultra-fast plans at heavily discounted introductory rates (e.g., NBN 1000 plans sitting around **$85–$99/month** for initial promotional periods).

2. How ISPs Advertise 'Typical Evening Speeds'

The removal of the shared CVC bottleneck has caused advertised "Typical Evening Speeds" (measured during the peak usage window of **7:00 PM to 11:00 PM**) to skyrocket.

Historically, ISPs had to specify a performance drop-off during peak times because they were skimping on CVC bandwidth. Today, because capacity is fully unconstrained at the wholesale level, **almost all major ISPs now advertise the maximum theoretical speed of the tier as their typical evening speed** for fixed-line connections (FTTP and HFC).

The **ACCC** actively enforces advertising guidance based on specific "Evening Speed" labels:

NBN Plan Tier	Old Typical Evening Speed (CVC Era)	Modern Advertised Peak Speed (AVC-Only Era)	ACCC Standard Label
NBN 25	20–22 Mbps	**25 Mbps** (100% of line speed)	Standard Evening Speed
NBN 50	42–45 Mbps	**50 Mbps** (100% of line speed)	Standard Plus Evening Speed
NBN 100	80–90 Mbps	**100 Mbps** (100% of line speed)	Premium Evening Speed
NBN 1000	250–500 Mbps	**700–900+ Mbps**	Ultrafast / Hyperfast

The New Bottleneck: Physical Infrastructure

Because the ACCC's broadband monitoring shows that fixed-line NBN networks routinely deliver **over 99% of their maximum plan speeds during busy evening hours**, ISPs are legally safe to advertise flawless performance.

If an Australian connection slows down during peak hours, it is no longer because the ISP is running out of bandwidth. Instead, any degradation is down to **domestic Wi-Fi interference** inside the home, or technical physical limits of old infrastructure like long copper loops on **Fibre-to-the-Node (FTTN)** lines.

** End of Page