September 2024

Let's Measure the Australian Internet Ecosystem

Internet Society

Aftab Siddiqui siddiqui@isoc.org

- Launched December 2020.
- We curate Internet measurement data from trusted sources to help everyone gain deeper, data-driven insight into the Internet.

Trusted data from multiple sources:

- Benefit: Helps to assess whether efforts to ensure that the Internet remains open, globally connected, secure, and trustworthy are working.
- Benefit: Allows policymakers, researchers, journalists, network operators, civil society groups, and others to better understand the health, availability, and evolution of the Internet.





But.

- I'll be sharing insights that are not available through the Pulse platform (at the moment), as they derive from dataset we're examining as part of our research team.
- Data analysis is a fascinating process—it can range from uncovering trivial details to discovering invaluable information.
- Through our analysis, we aim to transform these complex datasets into meaningful information that can benefit all of us.



Presentation vs Interpretation Bias

- Bias exists in all types of data analysis and is difficult to eliminate.
- Bias can occur at any stage of the process, and it can impact the validity and reliability of your findings, leading to "misinterpretation" of data.
- Interpretation Bias: If the analyst is making a conclusion that are not directly supported by the dataset.
- Presentation Bias: If dataset are presented without interpretation, there is a high likelihood that readers may draw misleading conclusions.
 - Let's address some presentation biases.







Presentation vs Interpretation Bias

 # curl -s 'https://ftp.ripe.net/pub/stats/ripencc/nro-stats/latest/nrodelegated-stats' | grep "apnic|AU|asn" | awk -F '|' '{print \$5}' | awk '{sum += \$1} END {print sum}'

→ 2901

- # curl -s 'https://ftp.ripe.net/pub/stats/ripencc/nro-stats/latest/nrodelegated-stats' | grep "|AU|asn" | awk -F '|' '{print \$5}' | awk '{sum += \$1} END {print sum}'
- → 2940
- \rightarrow AU delegations from all RIRs



Multiple transit providers ensure continuous service and reduce the risk of outages (Not exactly rocket science, Sherlock!)



Although it may appear that the 'relaxed' regulatory environment in Australia for ISPs allows for a variety of transit provider options, we can pick n choose as per the business requirement and there are 00s of networks providing transit services, so is it working?



Note: some networks do appear in Peeringdb as members of IXPs in AU but can't be verified through routeviews or RIPE Stat



Who are they relying upon for transit?





900+ Networks with Single Transit



140



What does that all show?





Upstream/Transit Diversity – ROV Status

Who is protecting you?

御







Top 35 networks based on transit ranking

IPv6 Adoption

IPv6 adoption varies significantly around the world and across different types of networks. Some regions and countries have significantly higher adoption rates. Major content delivery networks have been at the forefront.



IPv6 Adoption Statistics in Australia



https://www.akamai.com/internet-station/cyber-attacks/state-of-the-internet-report/ipv6-adoption-visualization



Region Map for Australia and New Zealand (053)

APNIC Labs https://stats.labs.apnic.net/ipv6/AU



IPv6 Adoption Statistics in Australia

IPv4 vs. IPv6 in Australia



Google IPv6

https://www.google.com/intl/en/ipv6/statistics.html#tab=per-country-ipv6-adoption

Cloudflare Radar

https://radar.cloudflare.com/adoption-and-usage/au?dateStart=2024-08-01&dateEnd=2024-08-21



IPv6 Adoption Statistics in Australia

Akamai: **26.9%** APNIC Labs: **43.06%** Cloudflare: **26.7%** Google IPv6: **31.22%**

Which number is correct?









Of the 2,940 ASNs allocated in Australia, over 1,200 are currently inactive. Review of the allocation dates reveals a widespread distribution, with a significant portion assigned between 2015 and 2021.



6

Active ASN with no IPv6 Announcements









BROEYOURSEF DNS ISSUES ARE COMING



makeameme.org

- DNS is fundamental to Internet functionality; nothing more on that!
- Proximity and Providers:
 - Critical to assess how far away the DNS service is located.
 - Important to identify who is providing the DNS service—whether it's a local service provider or a public DNS provider such as Google Public DNS, Cloudflare 1.1.1.1, Quad9, or OpenDNS, etc.
- Usage Patterns: Understanding what DNS services people are using is important.
- Contextual Relevance:
 - In open and democratic environments like Australia, the choice of DNS might not seem critical but still presents an informative landscape.
 - In other regions, the choice of DNS service can provide significant insights into local internet governance and access restrictions.



origi cont parts the l distr	1	CHINAM OBILE-CN China Mobile Communications Group Co., Ltd., CN
	2	ALIBABA-CN-NET Hangzhou Alibaba Advertising Co.,Ltd., CN
	3	AS-CHOOPA, US
	4	CHINANET-BACKBONE No.31, Jin-rong Street, CN
	5	CDNEXT, GB
	6	AKAMAI-A SN1, NL
	10	WOODYNET-1, US
	15	CISCO-UMBRELLA, US
	30	GOOGLE, US
	50	CDN77 _, GB
		CLOUDFLARENET, US
00 400)	

Out of 10k DNS queries, only 1,675 resolvers originated outside of the ISP network. This contrasts sharply with trends observed in other parts of the world, there is a correlation between the level of restrictions in a country and its resolver distribution

1000

600

800

1200

1400



25

1546

1600





321 Unique resolvers out of 10,550 queries. We verified the RPKI validation status for these 321 IP addresses.



Lack of a valid ROA can expose DNS resolvers to BGP route hijacking.



Where is the content?



Content Locality

Having content hosted locally is important, especially for a country like Australia. As a vast island nation, having data hosted within the country is more efficient than relying on servers across multiple continents. Not every island is as lucky as we are!

We use Tranco List (lists of popular domains), which is a combination of Google CrUX, Cloudflare Radar, Farsight, Majestic and Cisco Umbrealla.

For this work we are only focusing on .au cctld. We picked 35,000 domains from the Full Tranco List (1 million domains).



Content Locality

副



https://github.com/ISOC-Research/findcdn

Content Locality – CDN Hosting

501

328

incapsula

139

netity

65

rackspace

57

sectionio

30

level3

15

conTI

13

bunnycon

12

wordpress

bolt

51

3110303



Interestingly, the majority of content within the .au ccTLD is hosted by Cloudflare and Amazon, with a substantial portion also managed by Microsoft, Google, and Akamai. The rest is distributed among various smaller players. It's important to note that these observations pertain solely to the .au country-code top-level domain.

telefonica

22101

2

yottaa

conetworks

Content Locality – Non-CDN Hosting



Content Locality – gov.au

- AS45 671-NET-AU Wholesale Services Provider, AU 5
- MARKETCREATION S-AS-AP MARKET CREATION S, AU 5
 - CITEC-AU-AP QLD Government Business IT, AU 5
 - SUCURI-SEC, US 5
 - NETTAS-AS-AP Networking Tasmania, AU 🚦 5
 - GOHOSTING-AS-AP GoHosting, AU 5
 - DIGITALOCEAN-ASN, US 🛛 6
- VOCUS-BACKBONE-AS Vocus Connect International Backbone, AU
- AARNET-AS-AP Australian Academic and Research Network AARNet, AU
 - WA-GOVERNMENT-AS-AP WA Government project, AU
 - VOCUS-RETAIL-AU Vocus Retail, AU 🧧 6
 - OPTU SCOM-AS01-AU SingTel Optus Pty Ltd, AU 7
 - CLOUDFLARES PECTR UM Cloudflare, Inc., US 7
 - ASN-TELSTRA Telstra Limited, AU 🔳 10
 - AAPT AAPT Limited, AU 🔳 10

26

31

34

45

50

65

76

76

100

150

200

250

300

- MCT-SYDNEY Macquarie Technology Operations Pty Limited, AU 🗾 11
 - VZB-AU-AS Verizon Australia PTY Limited, AU 🛛 🔤 13
 - CENTRELINK Department of Human Services, AU 20
 - SECTIONIO-AS-AP SQUIXA PTY LIMITED, AU 23
 - AKAMAI-ASN1, NL 24
- DEWR-AS-AP Department of Employment and Workplace Relations, AU
 - FASTLY, US
 - MTG-MG Macquarie Technology Operations Pty Limited, AU
 - SQUIZ-AS-AP Squiz Pty Ltd, AU

INCAPSULA, US

0

- CLOUDFLARENET, US
- MICROSOFT-CORP-MSN-AS-BLOCK, US
 - AMAZON-02, US

Most .gov.au websites are hosted by major CDN providers, while the remainder are supported by a diverse array of local service providers.

343

400

Content Locality – gov.au – RPKI Validation Status

6

10

10

15

20

20



Over 150 .gov.au websites are hosted on IP addresses lacking valid Route Origin Authorizations (ROAs), with an RPKI status of "not-found." These websites are predominantly hosted by 48 providers, most of which are local networks.



Limitations



Limitations

- The data is pulled from external public sources, not always up-to-date.
- Without in-country measurements, it's difficult to validate the data.
 - RIPE Atlas and OONI are doing great work in this area, but more is needed.
- Some of the data undergoes processing, normalization, and weighing, we use a methodology that is reproducible.
 - You can see raw numbers via API. Email us for access pulse@isoc.org
- The idea is to help decision makers recognize gaps and weaknesses to conduct further study into validating these and work towards addressing them.



Take aways

- •It's crucial to thoroughly understand the measurements to assess whether they convey the real picture or relevant or accurate picture
- •Every data representation carries certain biases, which are manageable as long as they are acknowledged and addressed.
- •Final word: We are not doing good in IPv6 as a technical community!



Subscribe, Review, Contribute

Subscribe to the Pulse newsletter



Contribute to Pulse pulse@isoc.org

Review the Pulse IRI methodology









Aftab Siddiqui siddiqui@isoc.org