



Internet Performance Impacts of Canadian Content Hosting

Jim Cowie, Chief Scientist

DNS-OARC 2015

Montreal, October 2015

**INTERNET
PERFORMANCE.
DELIVERED.**

 dyn.com  [@dyn](https://twitter.com/dyn)

Global Content, Local Hosting

- Content localization is a form of international trade - import/export
- Choices made in each country affect enduser performance and are visible in the operation of their ccTLDs
- Policy folks are looking for ways to measure and understand Internet evolution
- **Zone operators have something important to contribute here...**



Zones Are Maps, Let's Peruse Them

- How does the “domain space” within a given TLD (including ccTLDs) map onto geographic space?
- Using Canada as a specific example, let's see where domains “live,” and what impact that might have on consumers
- We tried two different experiments, one starting with published **Alexa** data, the other with **ccTLD zone** data



EXPERIMENT 1:

FINDING THE ALEXA 100


Where's The Popular Content?

- **Retrieve** Alexa.com's public list of the top 100 sites for various countries
- **Resolve** the apex domains to determine one or more A records for each domain in the zone
- **Geolocate** each service IP and use BGP routing data to examine the origin and transit providers for each domain
- **Calculate** the fraction of each country's popular content that is locally hosted according to this metric



Canada's Alexa 100 Hosting (1/2)

google.ca	domestic	has ipv6	reddit.com	foreign	ipv4 only	tumblr.com	foreign	ipv4 only	xvideos.com	foreign	ipv4 only
google.com	domestic	has ipv6	imgur.com	foreign	ipv4 only	diply.com	foreign	ipv4 only	t.co	foreign	ipv4 only
facebook.com	foreign	has ipv6	ziddu.com	foreign	ipv4 only	ebay.com	foreign	ipv4 only	microsoft.com	foreign	ipv4 only
youtube.com	domestic	has ipv6	msn.com	foreign	ipv4 only	craigslist.ca	foreign	ipv4 only	nhl.com	foreign	ipv4 only
yahoo.com	foreign	ipv4 only	pinterest.com	foreign	ipv4 only	theweathernetwork.com	domestic	ipv4 only	bmo.com	foreign	ipv4 only
wikipedia.org	foreign	has ipv6	netflix.com	foreign	ipv4 only	bing.com	foreign	ipv4 only	stackoverflow.com	foreign	ipv4 only
amazon.ca	foreign	ipv4 only	paypal.com	foreign	ipv4 only	wordpress.com	foreign	ipv4 only	go.com	foreign	ipv4 only
twitter.com	foreign	ipv4 only	royalbank.com	foreign	ipv4 only	apple.com	foreign	ipv4 only	aliexpress.com	foreign	ipv4 only
live.com	foreign	ipv4 only	instagram.com	foreign	ipv4 only	pornhub.com	foreign	ipv4 only	cibc.com	foreign	ipv4 only
amazon.com	foreign	ipv4 only	cbc.ca	foreign	ipv4 only	rbcroyalbank.com	foreign	ipv4 only	theglobeandmail.com	domestic	ipv4 only
kijiji.ca	foreign	ipv4 only	td.com	foreign	ipv4 only	cnn.com	foreign	ipv4 only	scotiabank.com	foreign	ipv4 only
linkedin.com	foreign	ipv4 only	imdb.com	foreign	ipv4 only	kickass.to	domestic	ipv4 only	cra-arc.gc.ca	domestic	ipv4 only




Canada's Alexa 100 Hosting (2/2)

indeed.com	foreign	ipv4 only	livejasmin.com	foreign	ipv4 only	thepiratebay.se	foreign	ipv4 only	weather.gc.ca	domestic	ipv4 only
gfyca.com	foreign	has ipv6	vice.com	foreign	ipv4 only	godaddy.com	foreign	ipv4 only	gouv.qc.ca	domestic	ipv4 only
tdcanadatrust.com	foreign	ipv4 only	redflagdeals.com	foreign	ipv4 only	etsy.com	foreign	ipv4 only	mailchimp.com	foreign	ipv4 only
buzzfeed.com	foreign	ipv4 only	pof.com	domestic	ipv4 only	cyberpresse.ca	domestic	ipv4 only	bestbuy.ca	foreign	ipv4 only
about.com	foreign	ipv4 only	tripadvisor.ca	foreign	ipv4 only	futureshop.ca	foreign	ipv4 only	alibaba.com	foreign	ipv4 only
ask.com	foreign	ipv4 only	walmart.ca	foreign	ipv4 only	flickr.com	foreign	ipv4 only	cnet.com	foreign	ipv4 only
huffingtonpost.com	foreign	ipv4 only	espn.go.com	foreign	ipv4 only	onclickads.net	foreign	ipv4 only	costco.ca	foreign	ipv4 only
dropbox.com	foreign	ipv4 only	realtor.ca	domestic	ipv4 only	googleadservices.com	domestic	ipv4 only	pcfinancial.ca	domestic	ipv4 only
tsn.ca	domestic	ipv4 only	thestar.com	domestic	ipv4 only	nytimes.com	foreign	ipv4 only	bbc.co.uk	foreign	ipv4 only
dailymail.co.uk	foreign	ipv4 only	canadapost.ca	domestic	ipv4 only	craigslist.org	foreign	ipv4 only	hootsuite.com	foreign	ipv4 only
bbc.com	foreign	ipv4 only	desjardins.com	domestic	ipv4 only	github.com	foreign	ipv4 only	twitch.tv	foreign	ipv4 only
adcash.com	foreign	ipv4 only	ctvnews.ca	domestic	ipv4 only	reference.com	foreign	ipv4 only	expedia.ca	foreign	ipv4 only



Canada's Alexa 100 Hosting: Summary



All told, **16.1%** of Canada's most popular domains are hosted (at least partially) within Canada, adding up partial percentages

CDNs **do** push significant site content closer to users – these are just the resolutions of the “front door” domains.

EXPERIMENT 2:

*.CA

Examining .CA (thank you CIRA!!)

- Examined the zone in June and July 2015, and again at the end of September 2015 --- the statistics are v. stable
- NS records provided for **+2.3M** zones, pointing at **+79k** authoritative servers
- We performed authoritative resolutions of each individual apex domain against each provided authority, geolocated them, identified their BGP transit



Top contributors of NS authority

- domaincontrol.com
- netfirms.com
- name-services.com
- domainsatcost.ca
- webnames.ca
- registrar-servers.com
- namespro.ca
- dreamhost.com
- hostpapa.com
- hostgator.com

...and a long tail of nearly 40,000 more

The authoritative nameservers for the zones within .CA certainly tend to have High Canadian Content...


... but the A records they return for the apex domains of those zones are significantly less Canadian, as we'll see.

Digressions for Evangelists



- 0.16% of .CA zones return a DNSKEY record for the apex
- Only 130 zones (0.003%?) have a DS record in .CA – obstacles to adoption at registrars?
- Compare to 1.11% signed, with 0.3% under DS, in a recent .com sample
- Similar proportions of apex domains with AAAA records ☹️

Results: Canada's .CA Apex Hosting



18.7% of the domains within .CA are hosted (at least partially) within Canada

This is a little better than the results for the top 100 (16.1%), but the US still represents the most common hosting location for .ca apex domains



Top Origin Providers: Worldwide .CA

<u>Provider</u>	<u>Global%</u>	<u>Organization</u>
1. AS26496	34%	GoDaddy.com, LLC
2. AS29873	6.10%	The Endurance International Group, Inc
3. AS13768	4.80%	Peer 1 Network (USA) Inc.
4. AS14618	4%	Amazon.com, Inc.
5. AS46606	2.80%	Unified Layer
6. AS32613	2.10%	iWeb Technologies Inc.
7. AS22576	1.80%	Layered Technologies, Inc.
8. AS21844	1.60%	ThePlanet.com Internet Services, Inc.
9. AS20013	1.50%	CyrusOne LLC
10. AS14116	1.50%	InternetNamesForBusiness.com

The US numbers are skewed high by huge numbers of parked domains whose “location” is presumably not critical to enduser experience

...and a very long tail of smaller hosting providers



Top Origin Providers: *Canadian .CA*

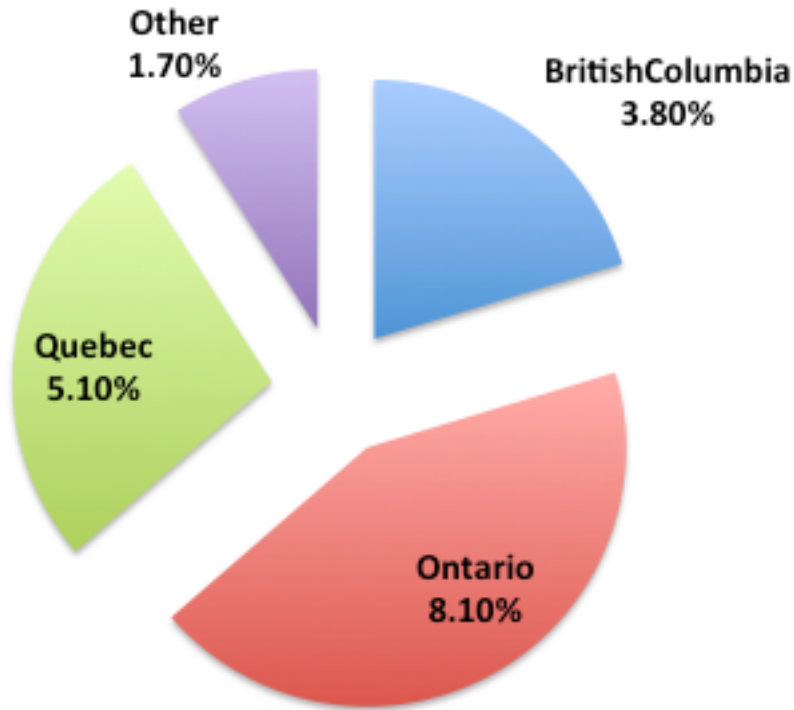
<u>Provider</u>	<u>Canada%</u>	<u>Organization</u>
1. AS13768	22.20%	Peer 1 Network (USA) Inc.
2. AS32613	11.30%	iWeb Technologies Inc.
3. AS15348	6.90%	Tucows.com Co.
4. AS16276	5.80%	OVH SAS
5. AS12188	2.50%	Q9 Networks Inc.
6. AS10929	2.50%	Netelligent Hosting Services Inc.
7. AS6327	2.10%	Shaw Communications
8. AS852	2%	Telus
9. AS14116	1.60%	InternetNamesForBusiness.com
10. AS30176	1.60%	Priority Colo Inc

Substantial domain hosting is focused in three major metro locations: Toronto, Vancouver, and Montreal

...and a very long tail of smaller hosting providers



Provincial .CA Apex Hosting



Of the Canadian 18.7% ..

- Ontario is nearly half
- Quebec next, followed by BC

Top Origin Providers: Ontario .CA

<u>Provider</u>	<u>Ontario%</u>	<u>Organization</u>
1. AS13768	22.70%	Peer 1 Network (USA) Inc.
2. AS15348	16.10%	Tucows.com Co.
3. AS12188	5.90%	Q9 Networks Inc.
4. AS14116	3.80%	InternetNamesForBusiness.com
5. AS30176	3.80%	Priority Colo Inc
6. AS13657	3.50%	E-Gate Communications Inc.
7. AS36351	2.90%	SoftLayer Technologies Inc.
8. AS21949	2.10%	Beanfield Technologies Inc.
9. AS14166	2%	Softcom Technology Consulting Inc.
10. AS15290	1.80%	Allstream Corp.
11. AS577	1.50%	Bell Canada

Top Ontario providers tend to mirror top Canadian providers, since Ontario has the largest provincial share of the zone's origination.



Top Origin Providers: Quebec .CA zones

<u>Provider</u>	<u>Quebec%</u>	<u>Organization</u>
1. AS32613	41.20%	iWeb Technologies Inc.
2. AS16276	19.20%	OVH SAS
3. AS10929	9.10%	Netelligent Hosting Services Inc.
4. AS36666	4%	GloboTech Communications
5. AS5769	3.30%	Videotron Telecom Ltee
6. AS13768	3.20%	Peer 1 Network (USA) Inc.
7. AS22652	1.40%	Fibrevoire Internet Inc.
8. AS16532	1.10%	B2B2C Inc
9. AS4540	0.90%	SOGETEL INC
10. AS36354	0.80%	SherWeb inc.

Top Quebec providers, on the other hand, are quite distinct. Language is presumably a strong driver of content hosting differentiation within the .CA zone.



Top *Transit* Providers: Quebec .CA zones

<u>Provider</u>	<u>Quebec%</u>	<u>Organization</u>
1.AS6453	94.80%	Tata Communications
2.AS3356	92.20%	Level 3
3.AS3257	70.70%	GTT
4.AS174	50.40%	Cogent
5.AS1299	48%	TeliaSonera International Carrier

Transit providers (wholesalers who are BGP upstreams of the hosting providers) are visibly different in Quebec and Ontario, due to the historic role of Teleglobe (now Tata Communications).


<u>Provider</u>	<u>Ontario%</u>	<u>Organization</u>
1.AS1299	90.90%	TeliaSonera International Carrier
2.AS6453	72.10%	Tata Communications
3.AS3356	69.10%	Level 3
4.AS3257	65.60%	GTT
5.AS2914	57.60%	NTT



WHY LOCAL CONTENT MATTERS:

IMPACTS ON ENDUSER EXPERIENCE

US-based .CA Apex Hosting

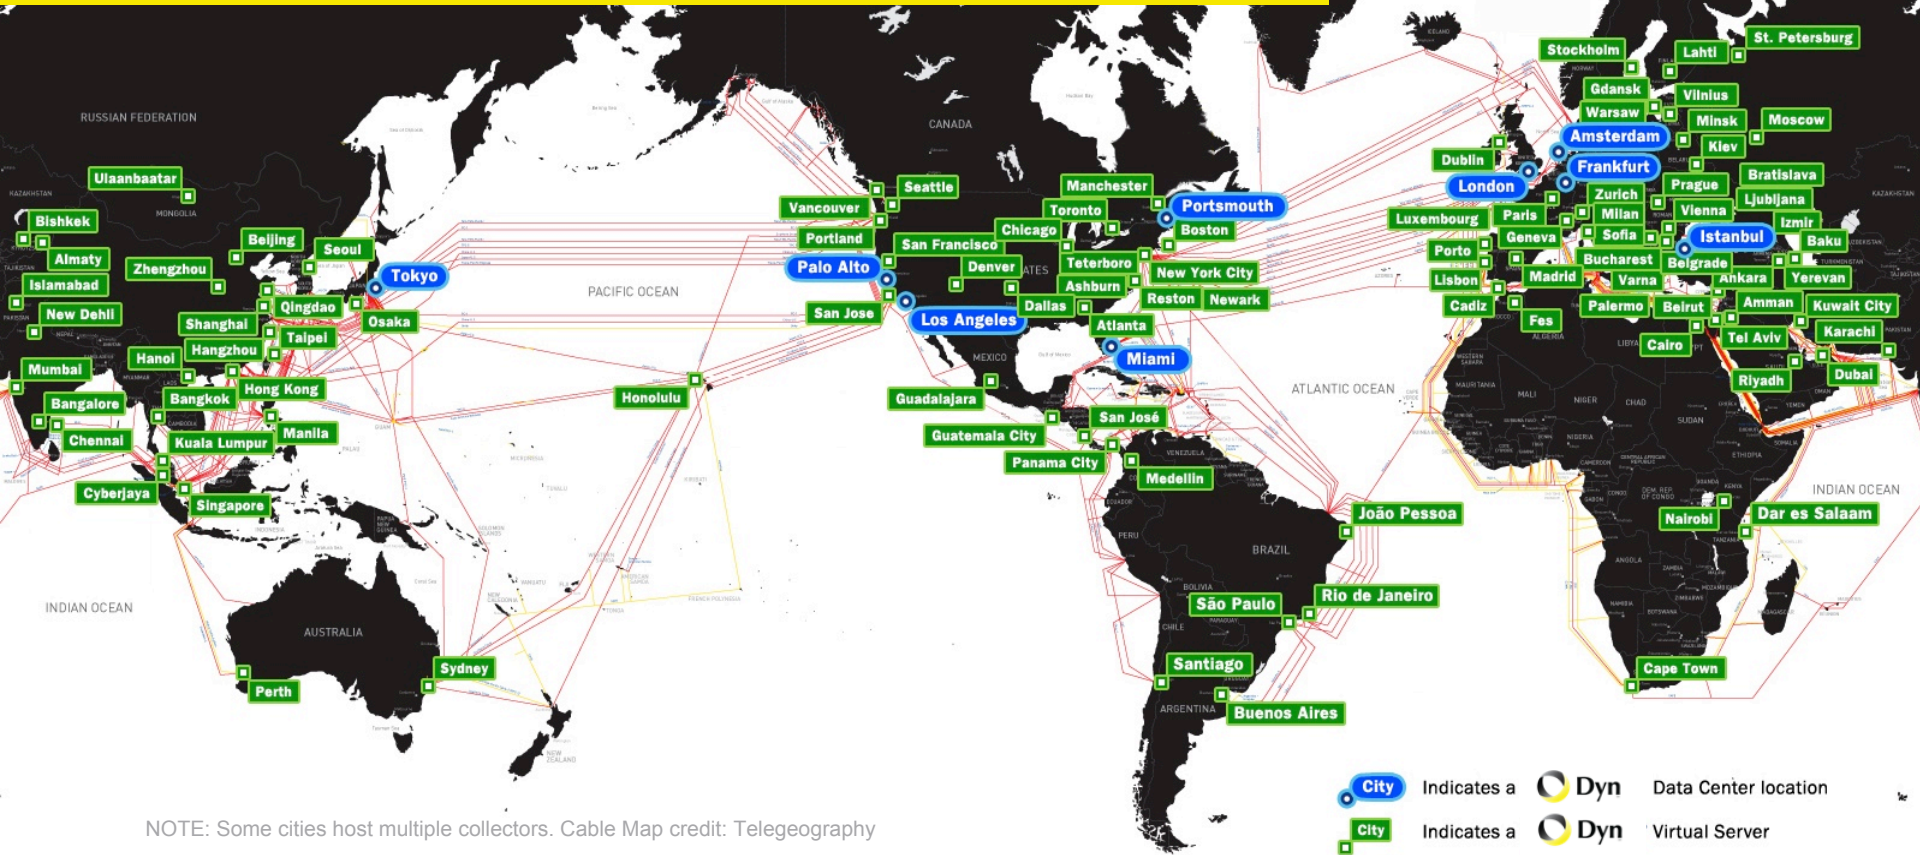


78.1% of the domains within .CA are hosted (at least partially) within the USA, though a high percentage are parked

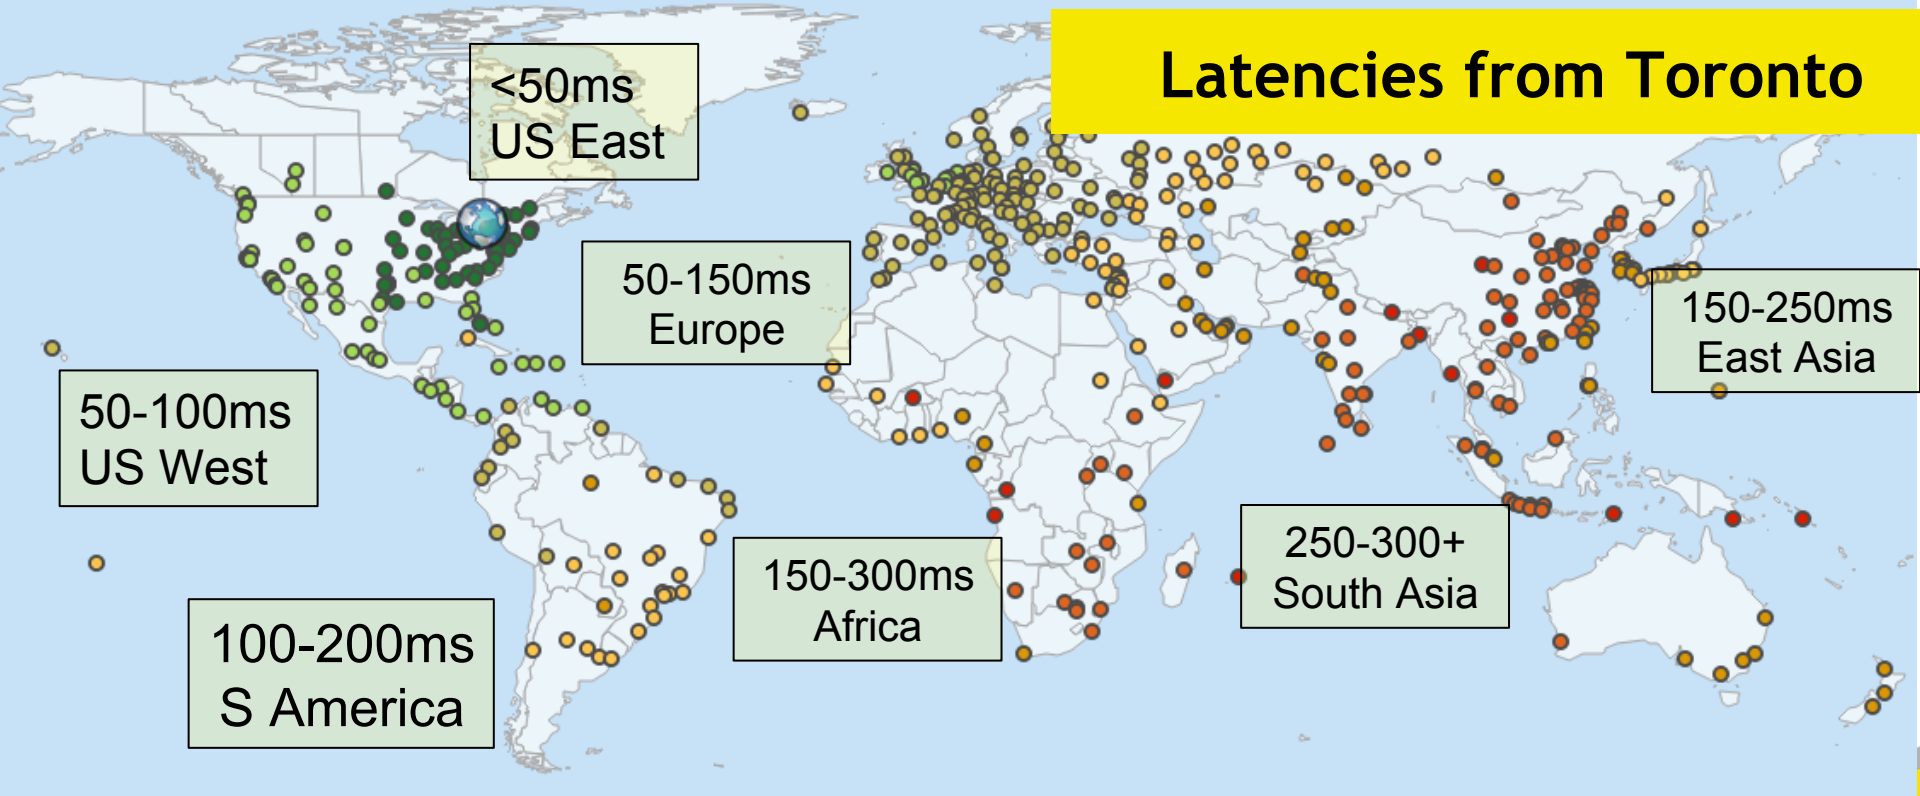
- **Phoenix** (34%, GoDaddy)
- **Boston** (7%, Endurance Intl)
- **DC Area** (5%, Amazon)
- **Utah** (3%, Unified Layer)
- Other markets have longer tails of hosting providers



Active Measurement Infrastructure



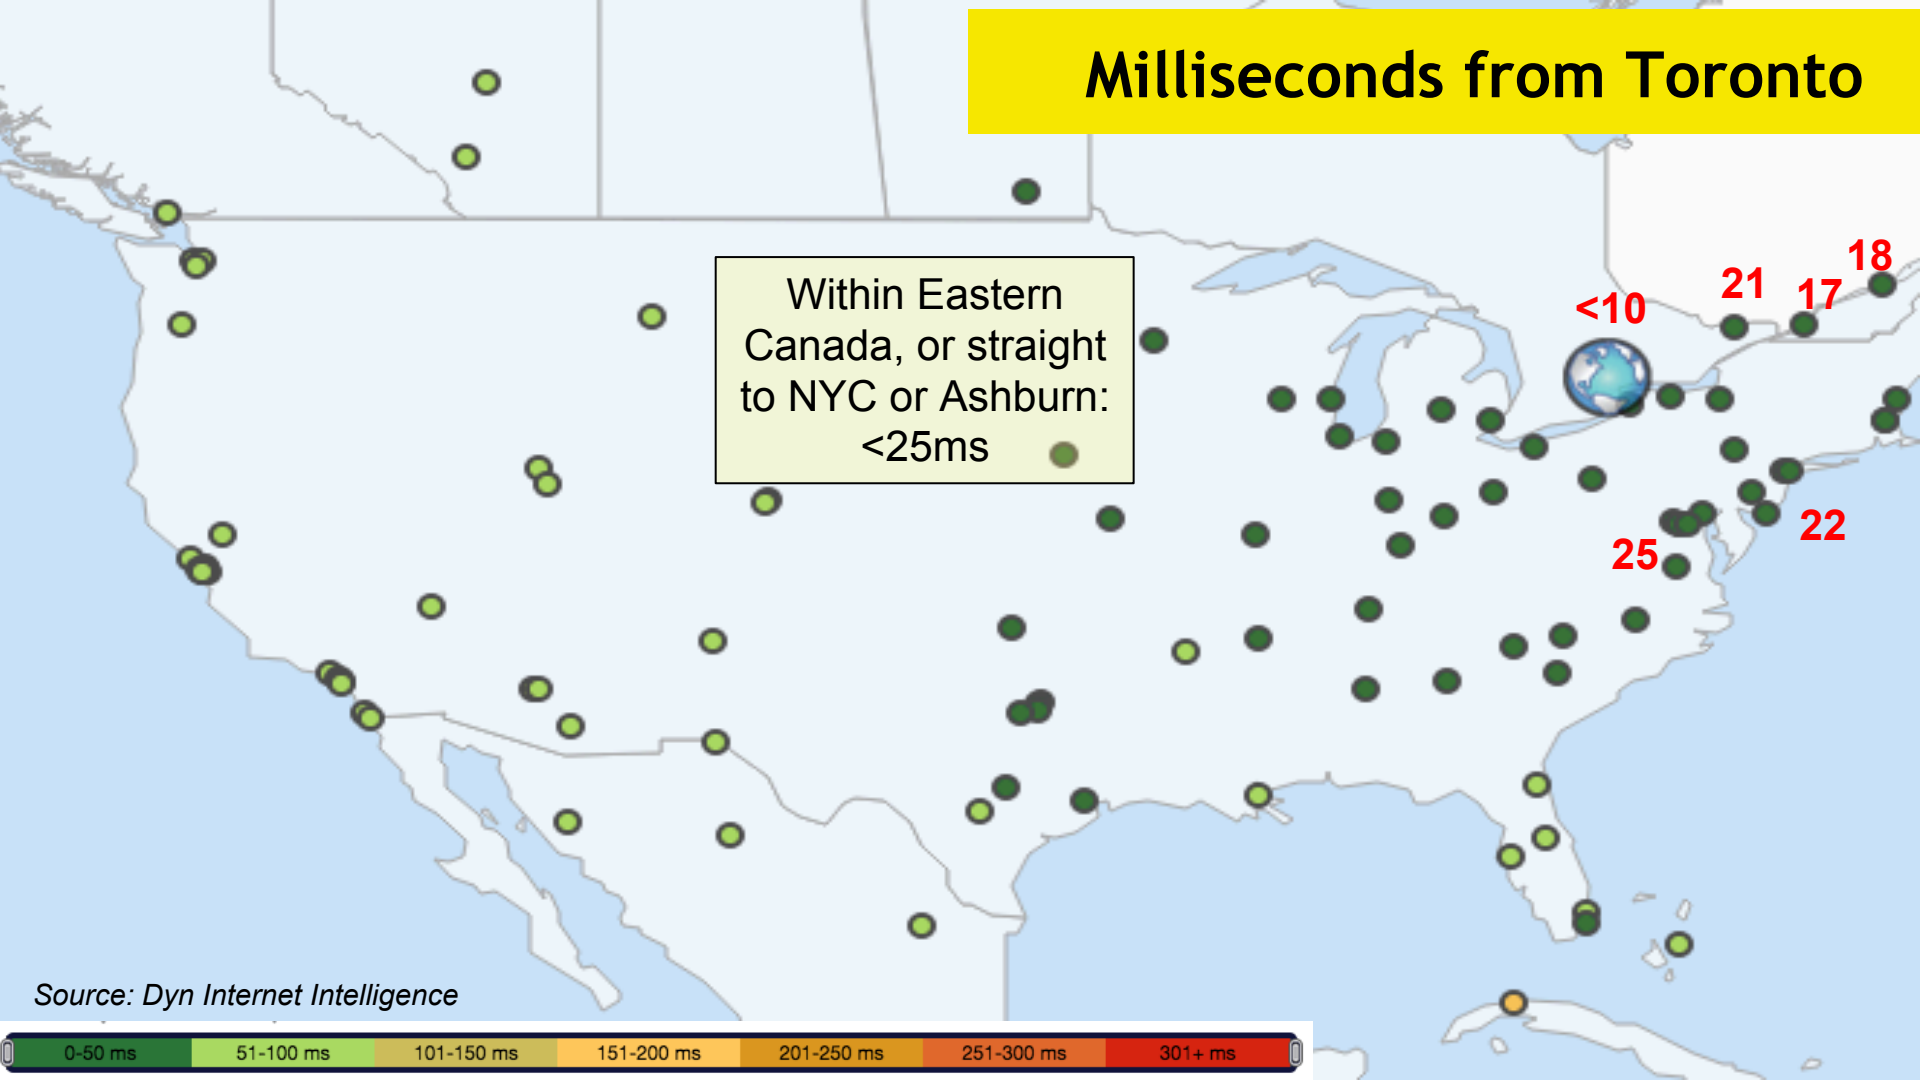
Latencies from Toronto



Source: Dyn Internet Intelligence

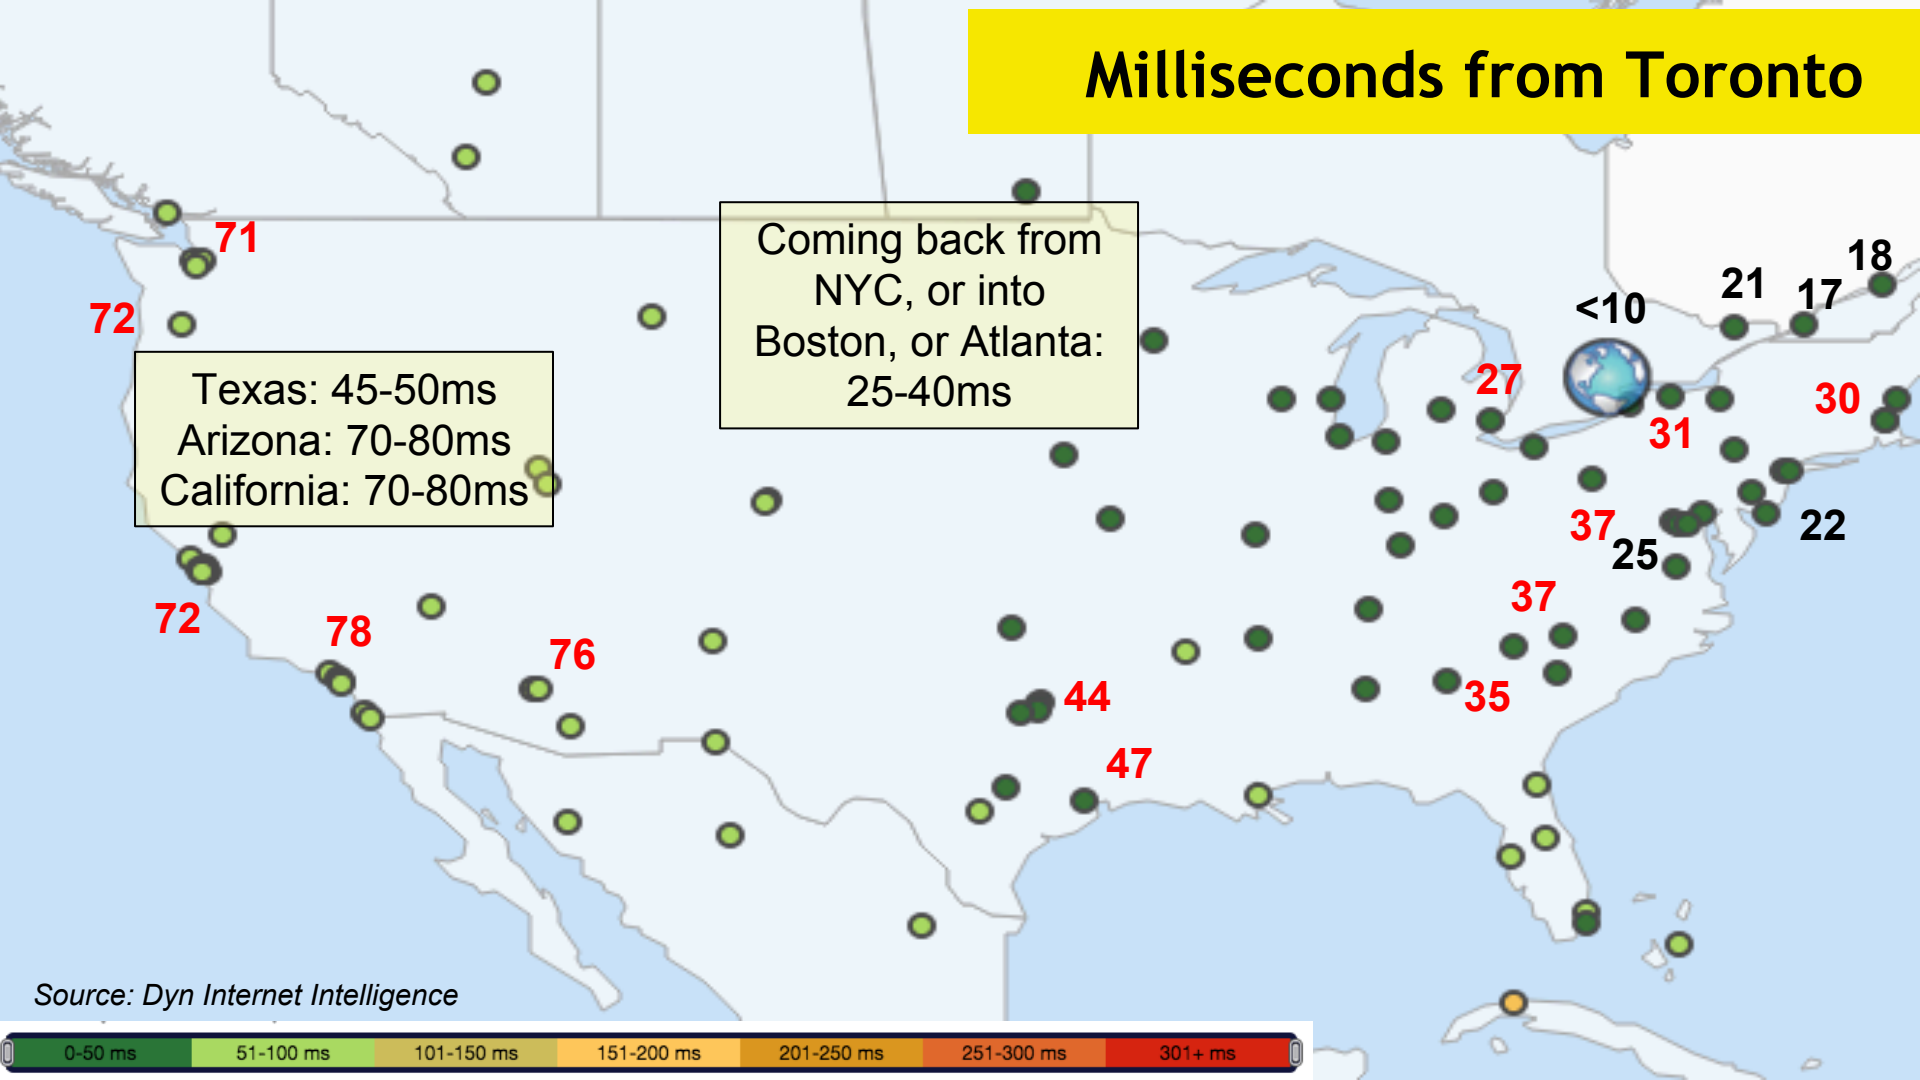


Milliseconds from Toronto




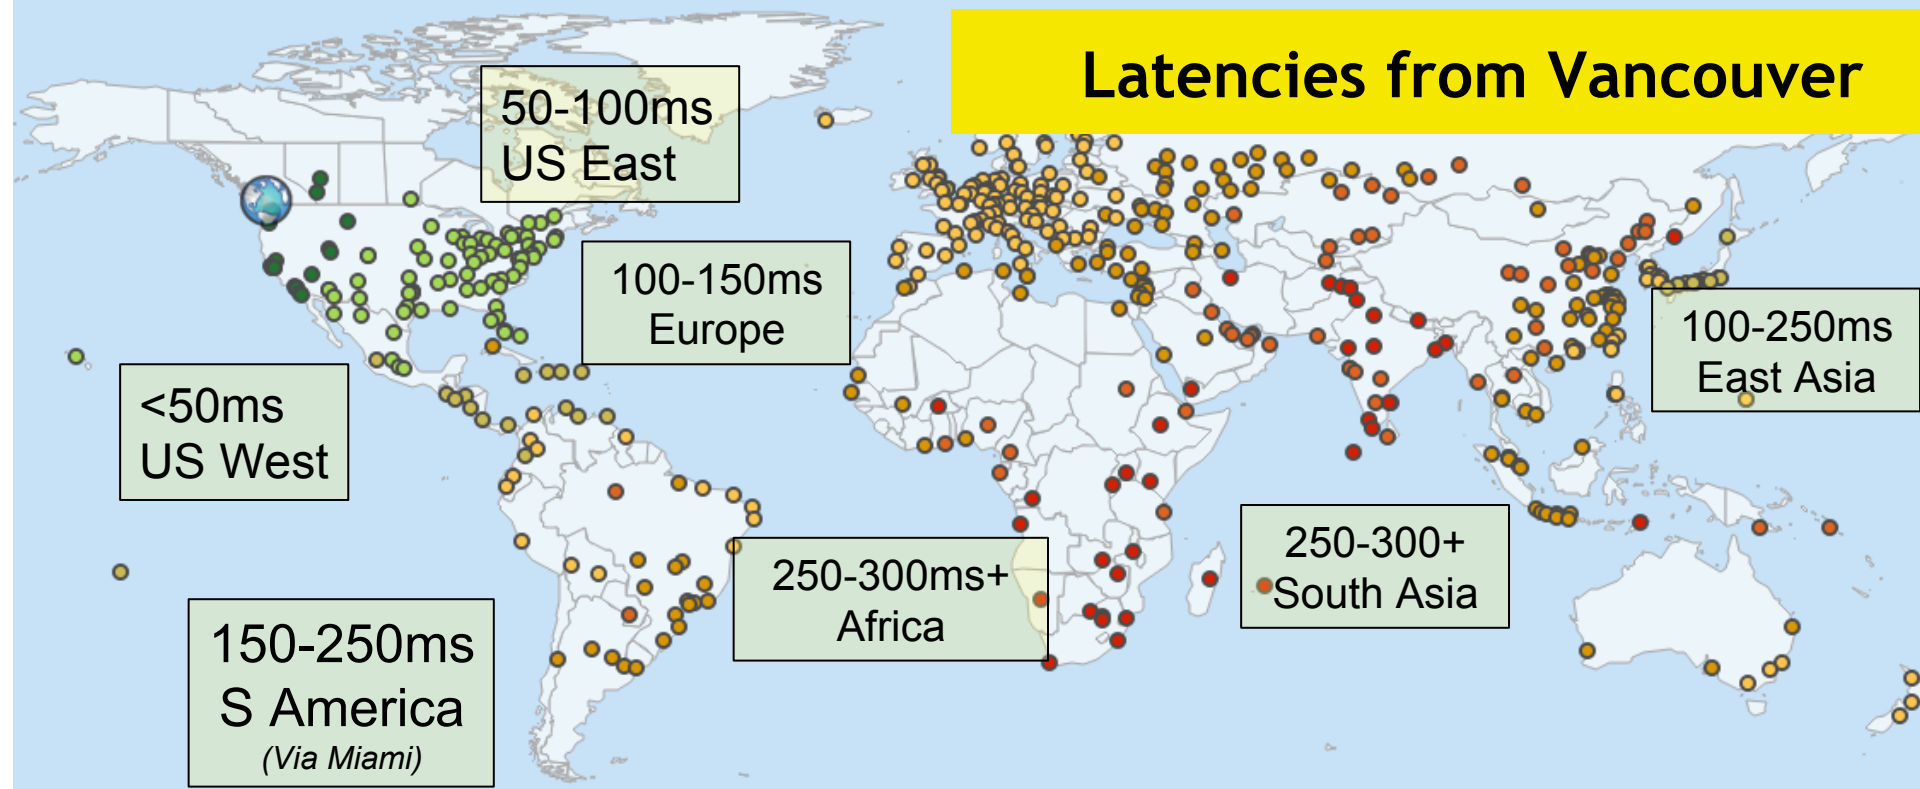
Source: Dyn Internet Intelligence

Milliseconds from Toronto



Source: Dyn Internet Intelligence

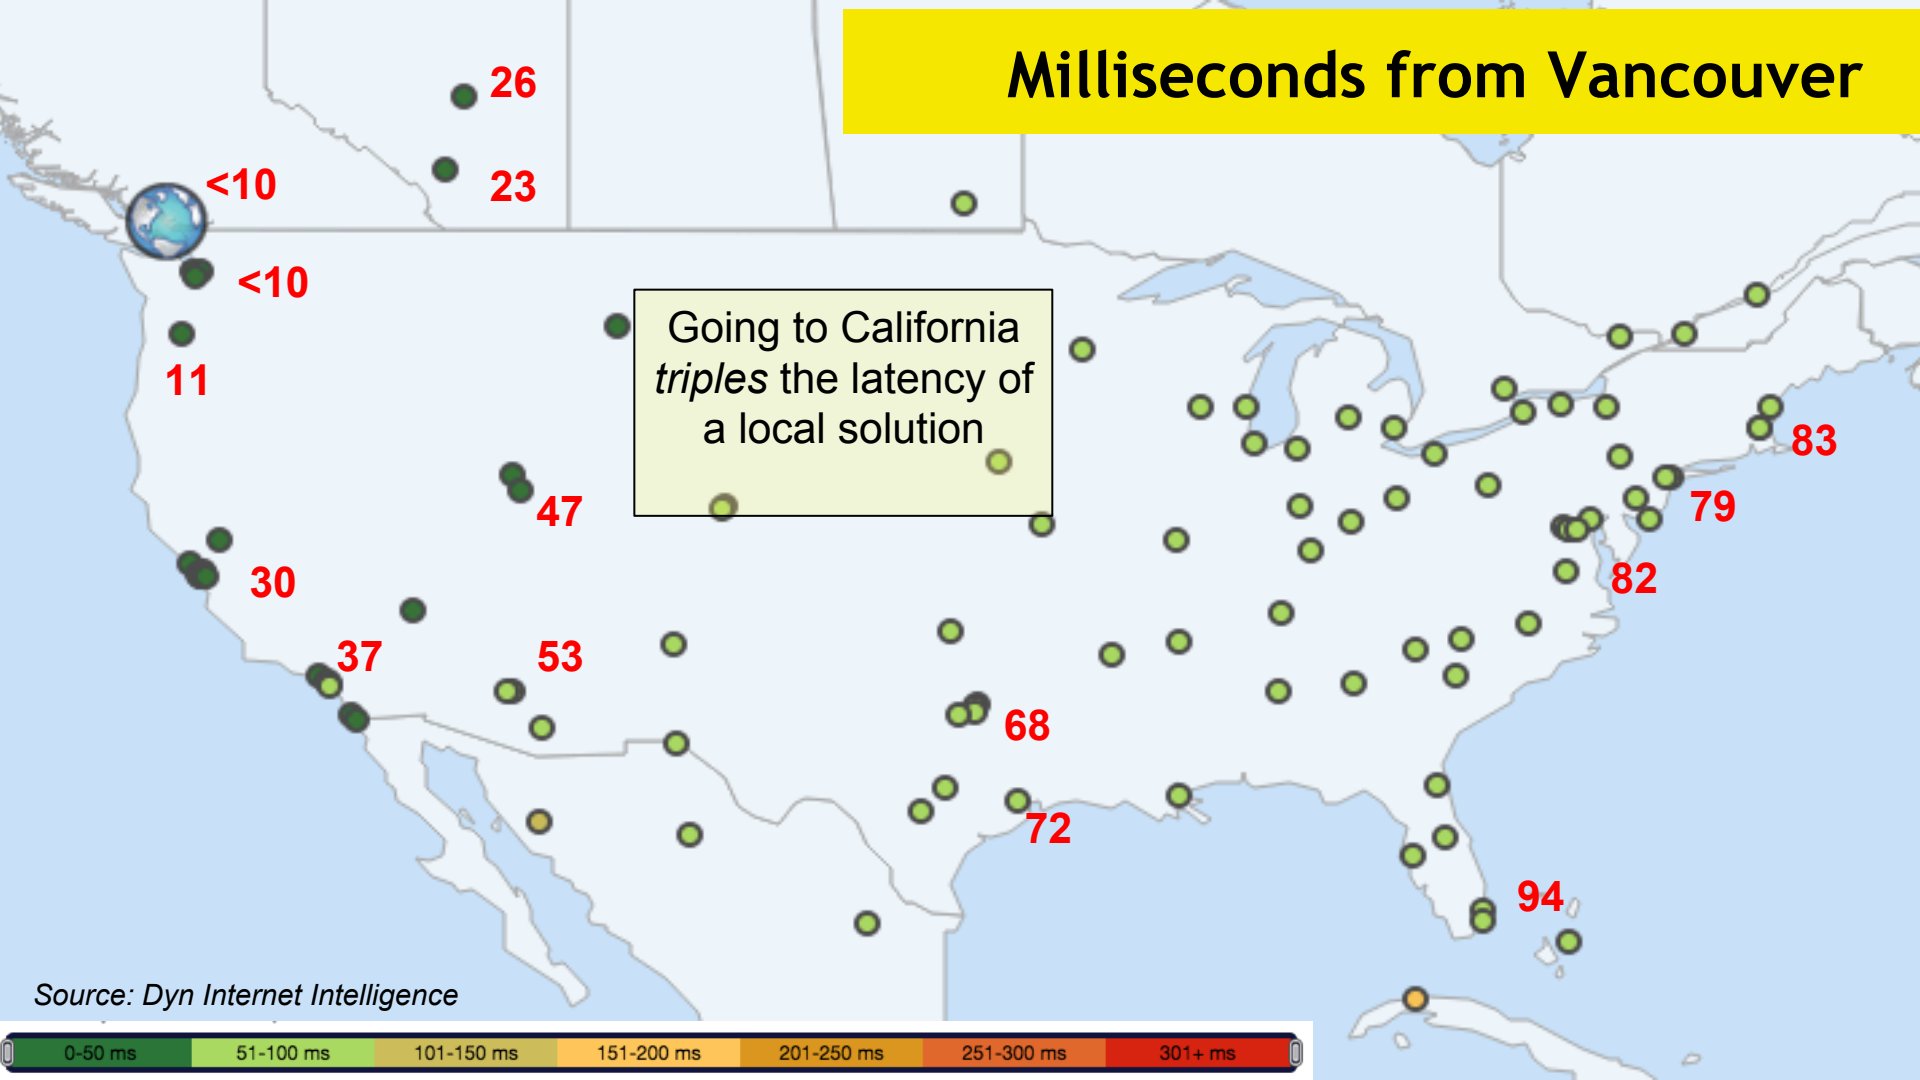
Latencies from Vancouver




Source: Dyn Internet Intelligence



Milliseconds from Vancouver



...but where do queries come from?




Resolver geography is counterintuitive.

We tend to assume that .CA is consumed by Canadians, but in fact .CA domains are resolved by people behind resolvers in all regions of the world.



...but where do queries come from?



Top countries by %pct resolvers

US	27%
Brazil	8%
Germany	7%
Canada	7%
Russia	4%
GreatBritain	4%
France	3%
Japan	3%
Italy	3%
India	2%
Other	33%



TL;DR

Canadian zones with content hosted in the USA increase latency for Canadian users by a significant percentage (and it better be coastally localized)

Global consumers of .CA would be just as happy if the content were in Canada (small/no penalty)

Host in Canada! 



NEXT STEPS:

HOW YOU CAN HELP

Two Kinds of Questions We'd Like Zones To Answer

Trending Questions

- Is local hosting increasing/decreasing? DNSSEC adoption? IPv6?
- Are consumers experiencing more or less latency/congestion as they try to access important content?


Correlation Questions

- As IXPs open, cables land, providers enter and exit markets, what changes in these metrics do we observe in each country?
- What are the regional effects, and how can hosting and transit trends in one country affect its neighbors?



Final Observations

- Interesting dynamics occur at the growing edge of big zones, so periodic measurement seems important
- Locality trends in the resolver audience for authoritative data are equally interesting
- We appreciate the help of [cc]TLD operators who have shared their data as part of this project!





Internet Performance Impacts of Canadian Content Hosting

Jim Cowie, Chief Scientist

DNS-OARC 2015

Montreal, October 2015

**INTERNET
PERFORMANCE.
DELIVERED.**

 dyn.com  [@dyn](https://twitter.com/dyn)