

YYCIX

an Internet Exchange for Calgary

Theo de Raadt

Who am I?

OpenBSD / OpenSSH project leader / release engineer

One of those outdoorsy programmer types

Calgarian

YYCIX network manager (IX operations)

AlbertaIX board member



What is OpenBSD?

Operating System Security Idea Incubator Organization

pf / tmux / security hardening / privilege-separation /
privilege-drop / stackghost / OpenSSH / strlcpy /
auditing procedures / stack-protector / OpenBGPD / OpenNTPD /
isakmpd / IKEv2 / relayd / new resolver / soft-error ICMP/RST /
randomized malloc / hw.sensors / hundreds more...

We invent/adopt new "security ideas", complete/integrate them into the whole system, then ship a release to push the idea to maturity.

ie. first OS to ship IPSEC, IPv6, crypto, SSH, strong random, stack-protector, other security mitigation methods, ...

17-year success story.

Why OpenBSD matters a lot

These ideas change the world because vendors copy the idea / code



People have no idea how much free code is used in commercial products

Over 20 years, Linux + SSH + BSD + GNU compiler tools
have become **CRITICALLY IMPORTANT INFRASTRUCTURE!**

These kinds of projects must continue...

OK, what does that have to do with networking?

OpenBSD has 100-150 developers all around the world

All dedicated volunteers to the cause -- work is cheap :)

Machines are cheap - people give them to us :)

Electricity costs are covered :)

But?

The local providers stop us from building a good network, GRRR!

Why is a Shaw cable modem not enough?

We need a lot of redundancy, for many IP addresses, and push a lot of traffic outwards

We need the best routing choices to reach our developers



To protect ourselves, OpenBSD has built a Real Network (TM)

It wasn't very nice to show that picture...

In reality, planning against failure is simply *Good Engineering*

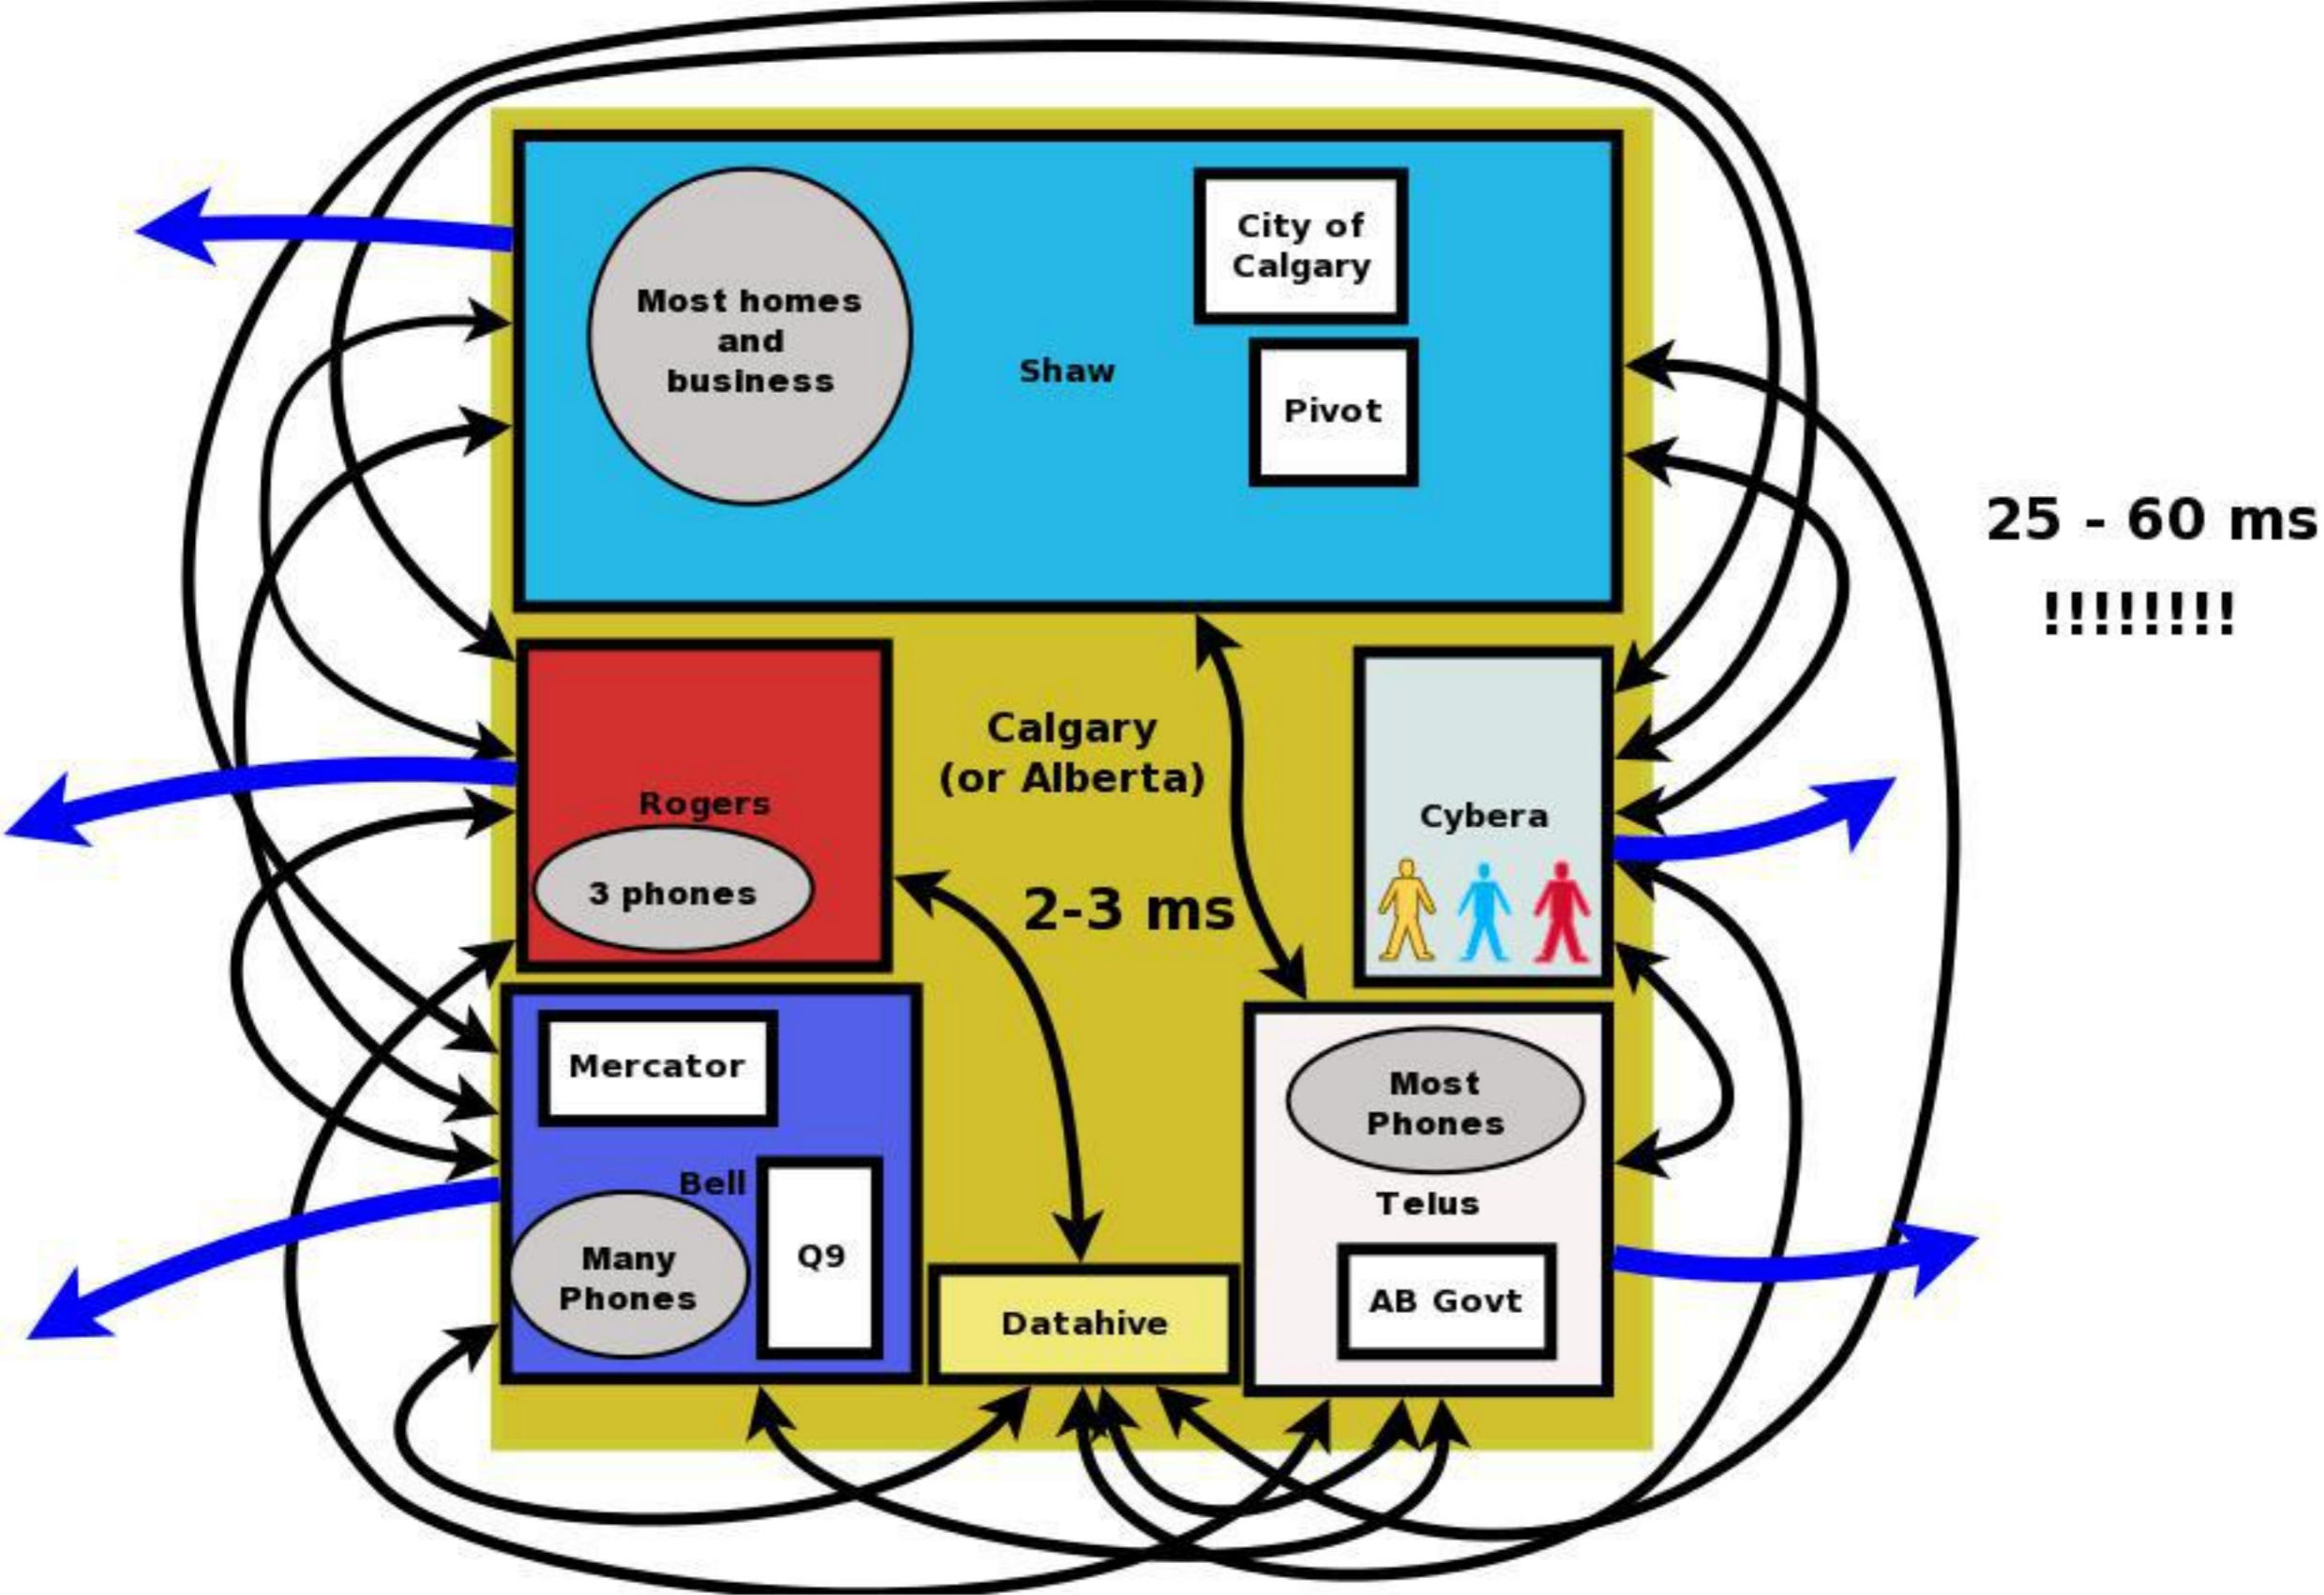
Real Networks (TM) guard against single/multiple points of failure

- They own an AS to define themselves
- They own their own network prefixes
- They use BGP on multiple edge routers
- They advertise their network to multiple providers and peers
- They use multiple physical network links, at multiple locations



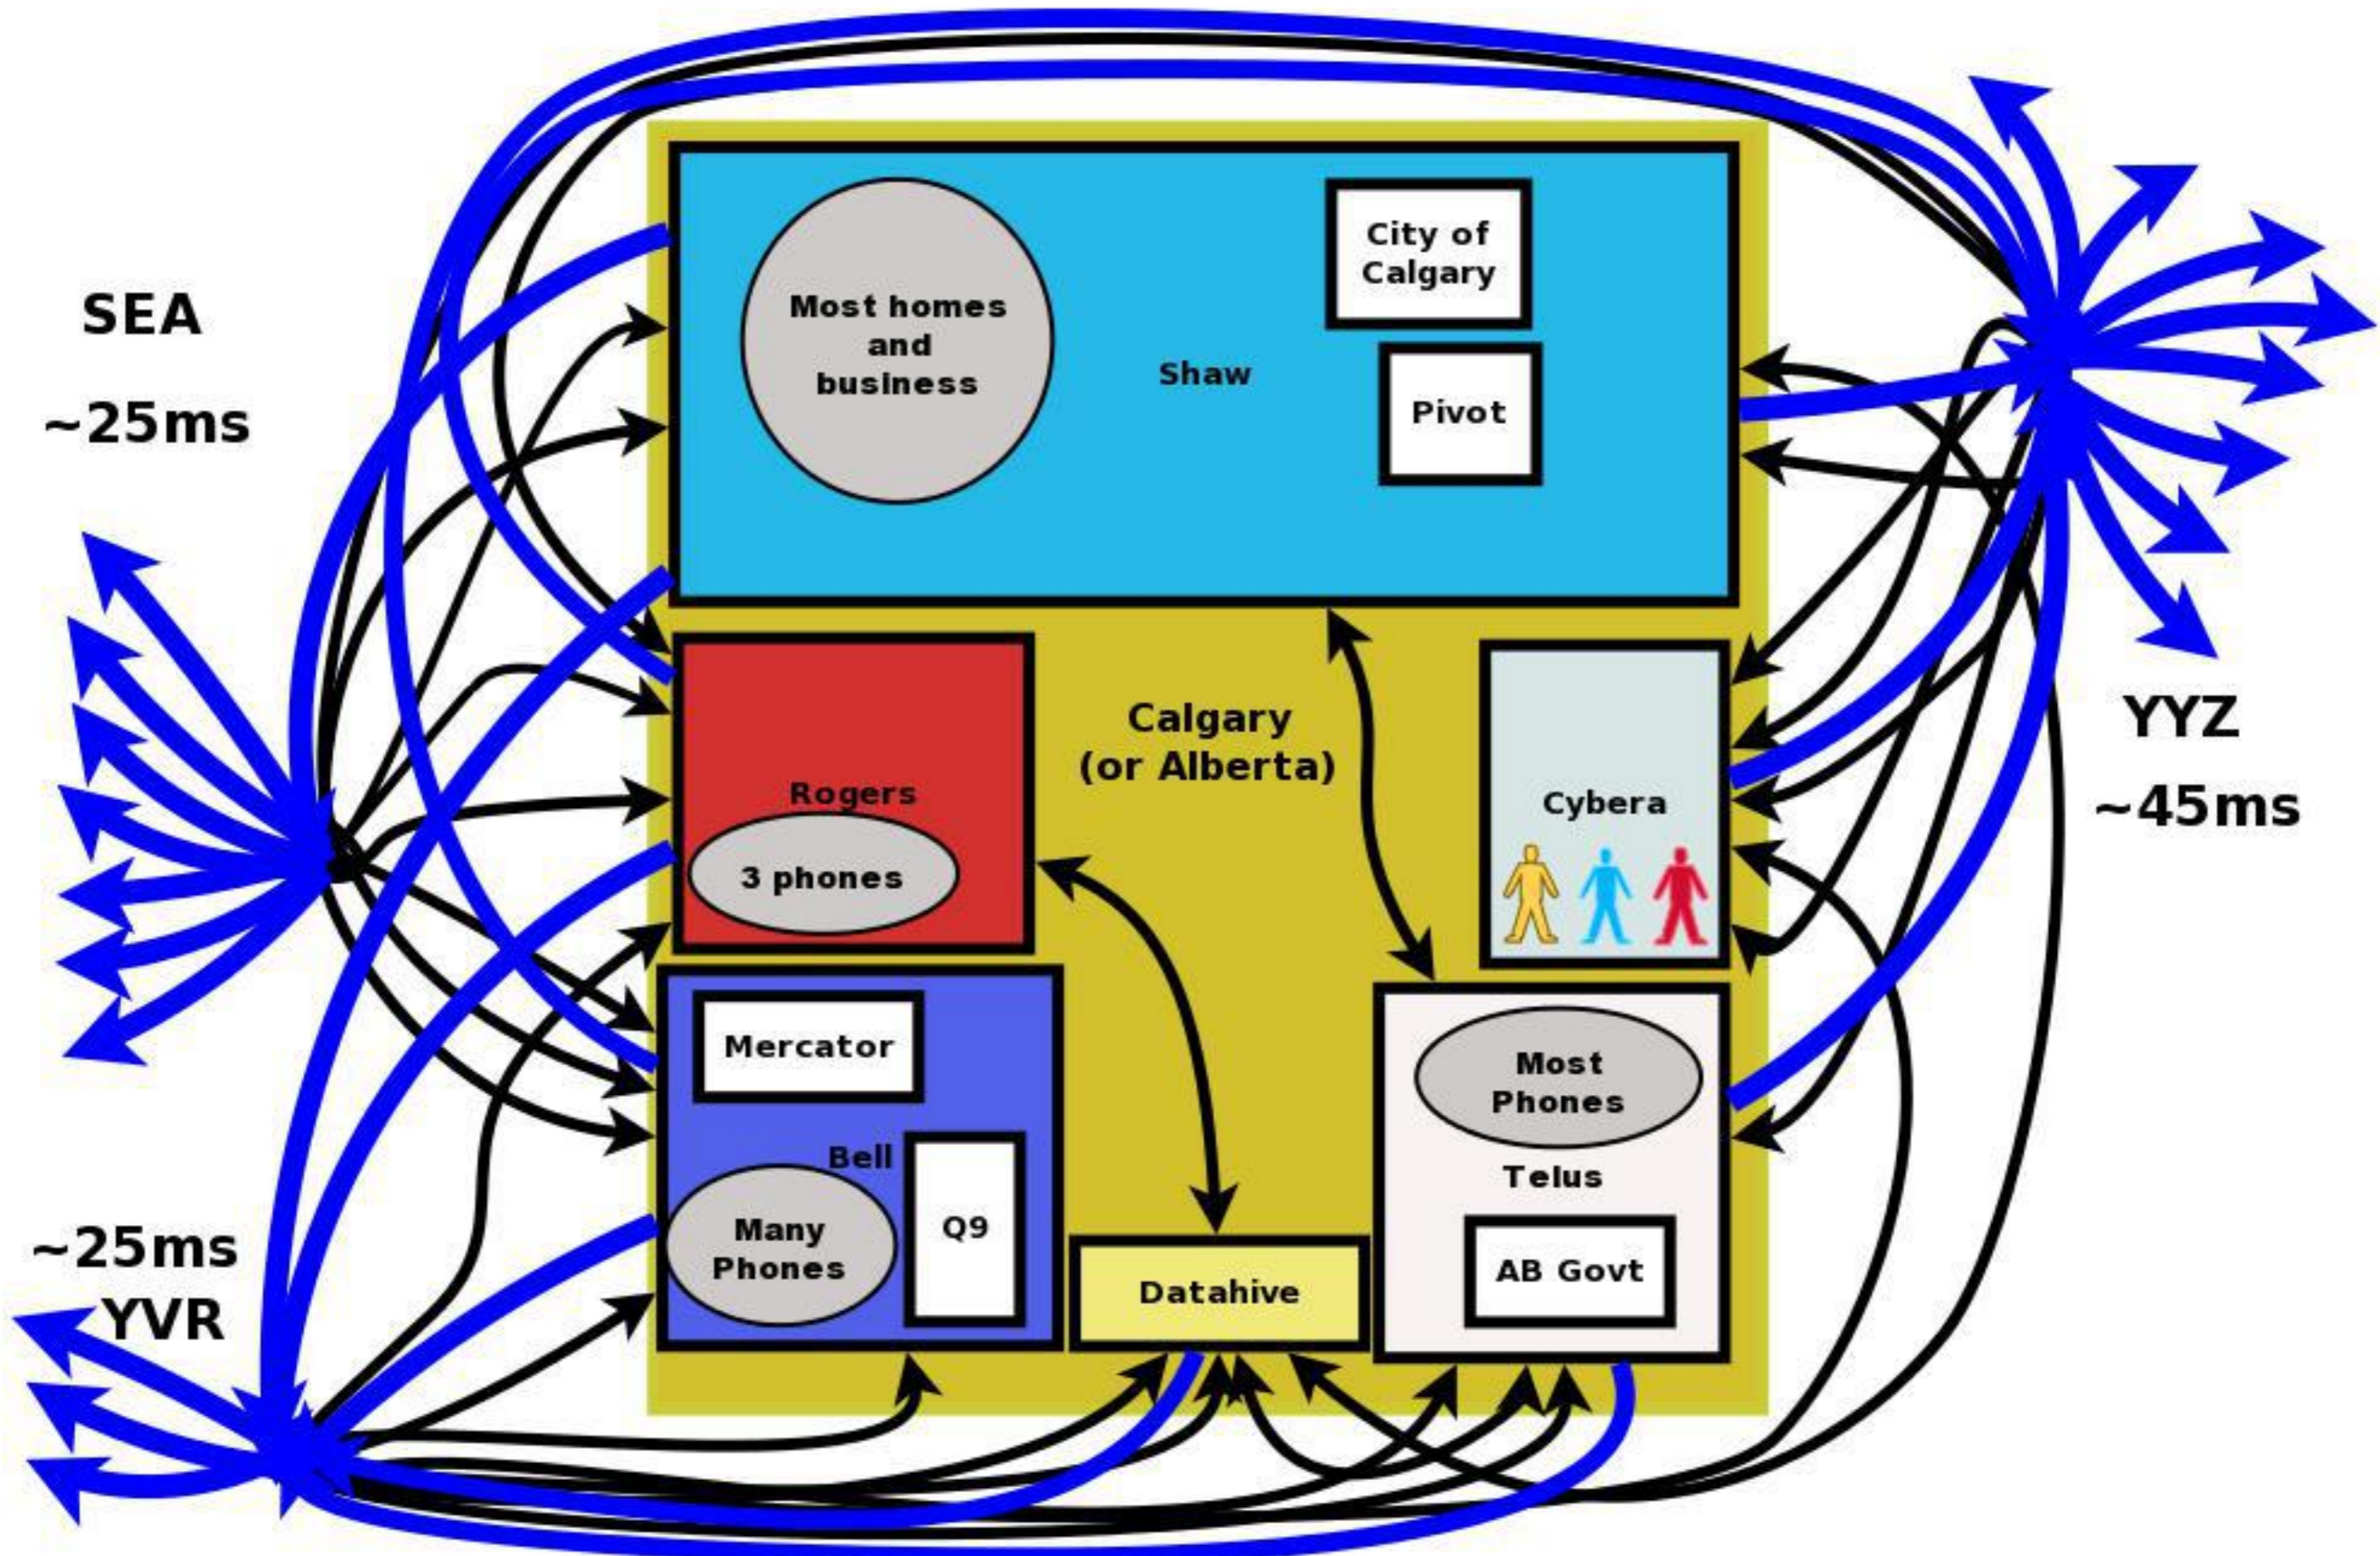
Multi-homing is what makes the Internet resilient.

So let's look at Calgary today



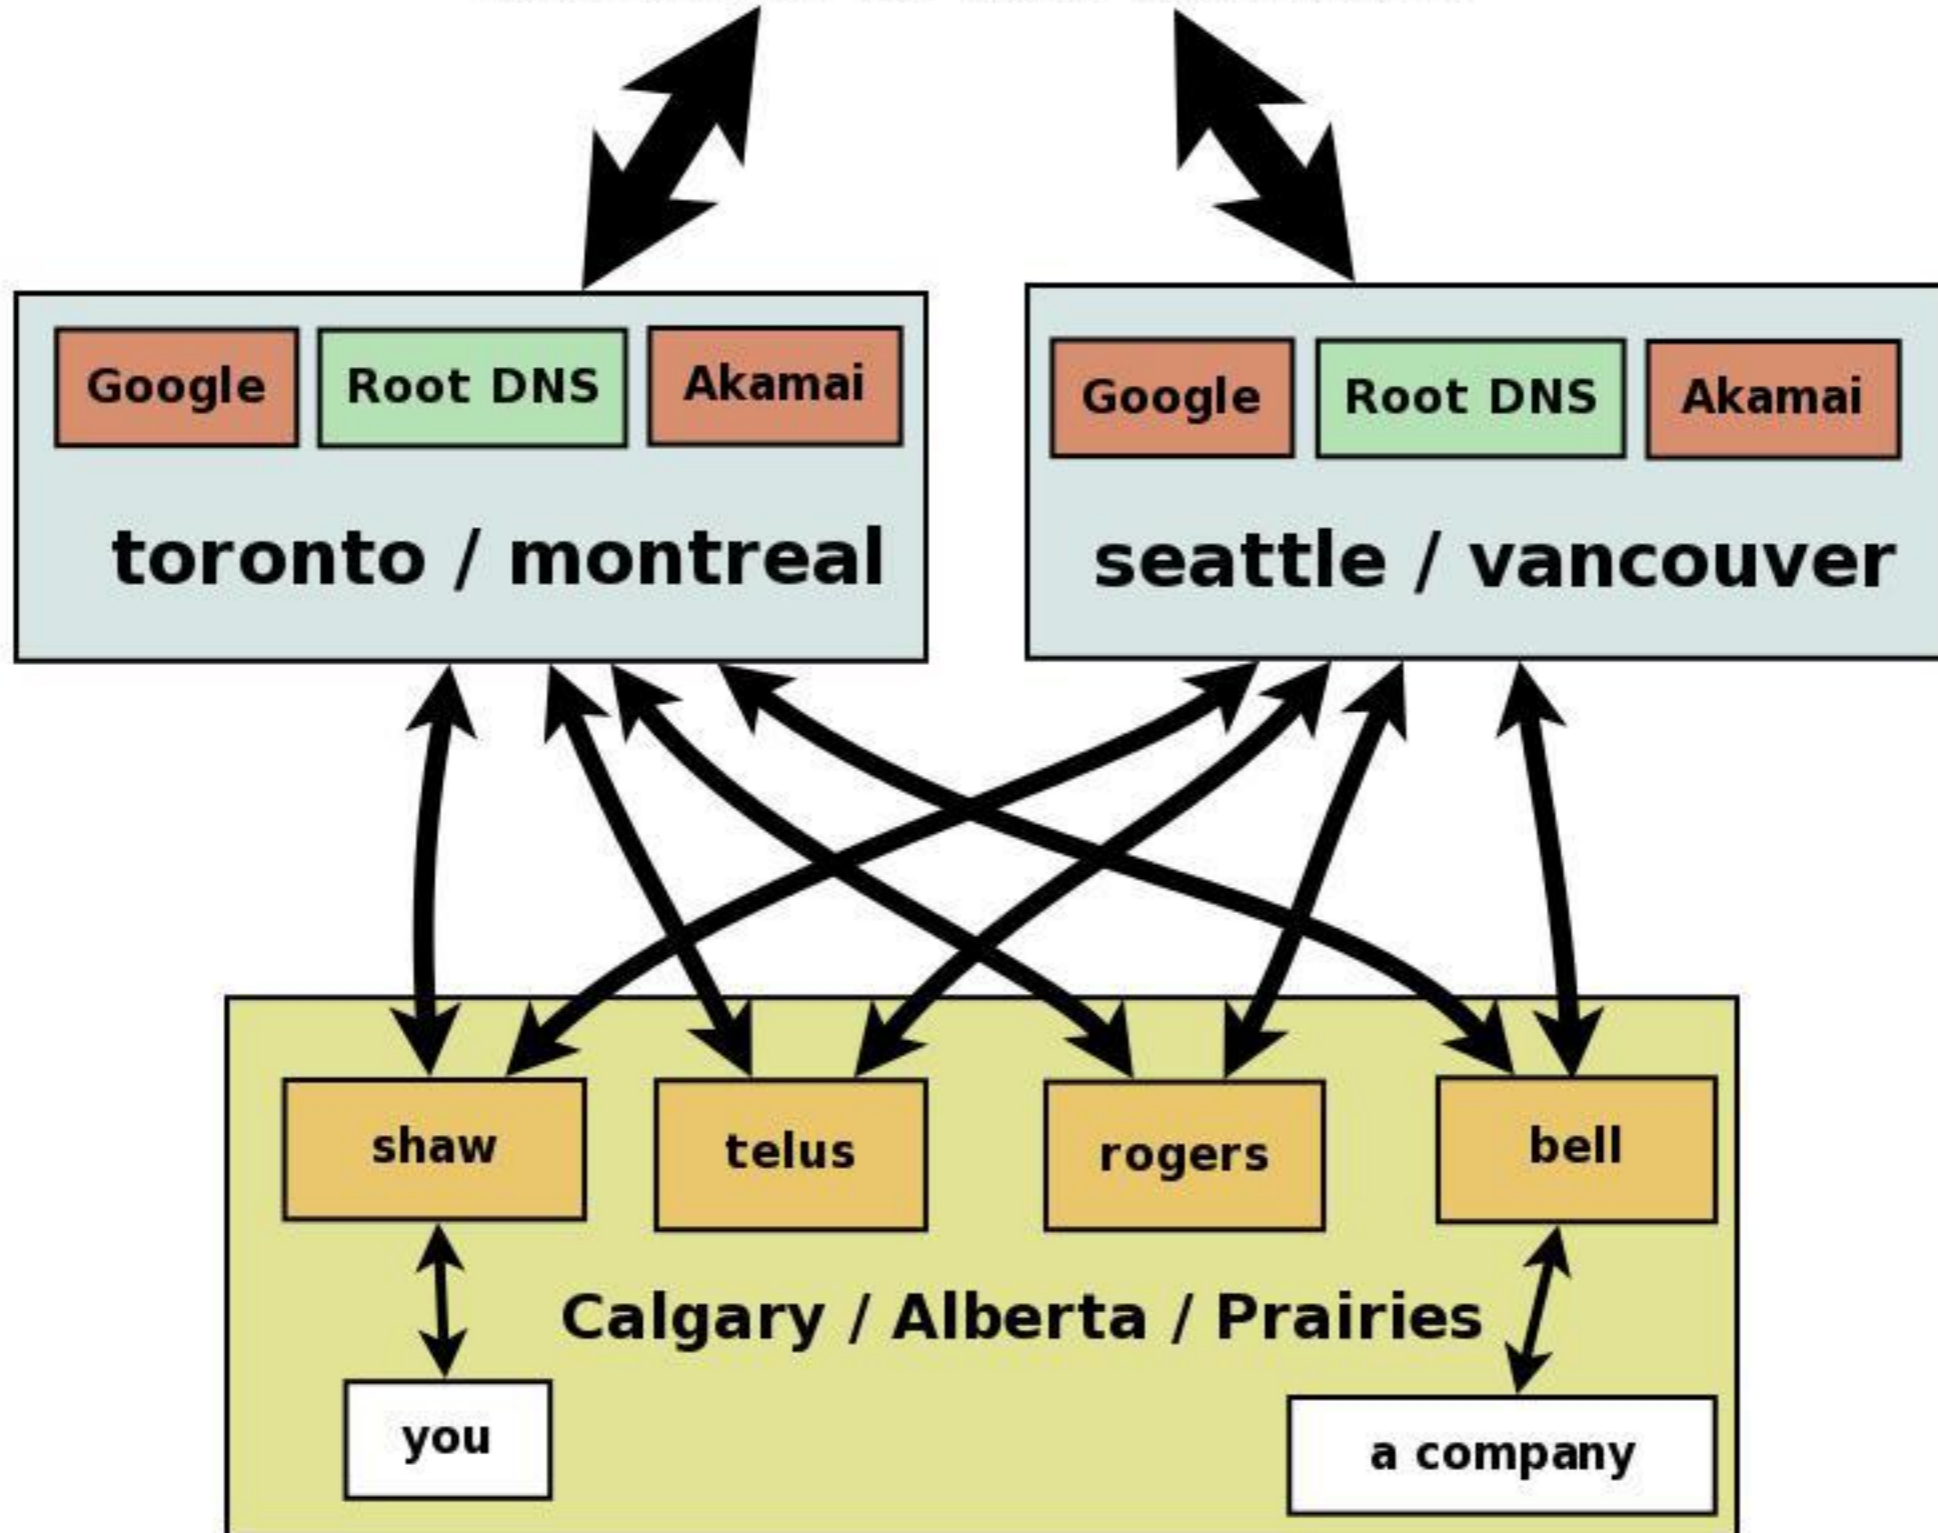
Hurting us locally, but they do have lots of redundant paths... wait

In truth, those loops go through the same places

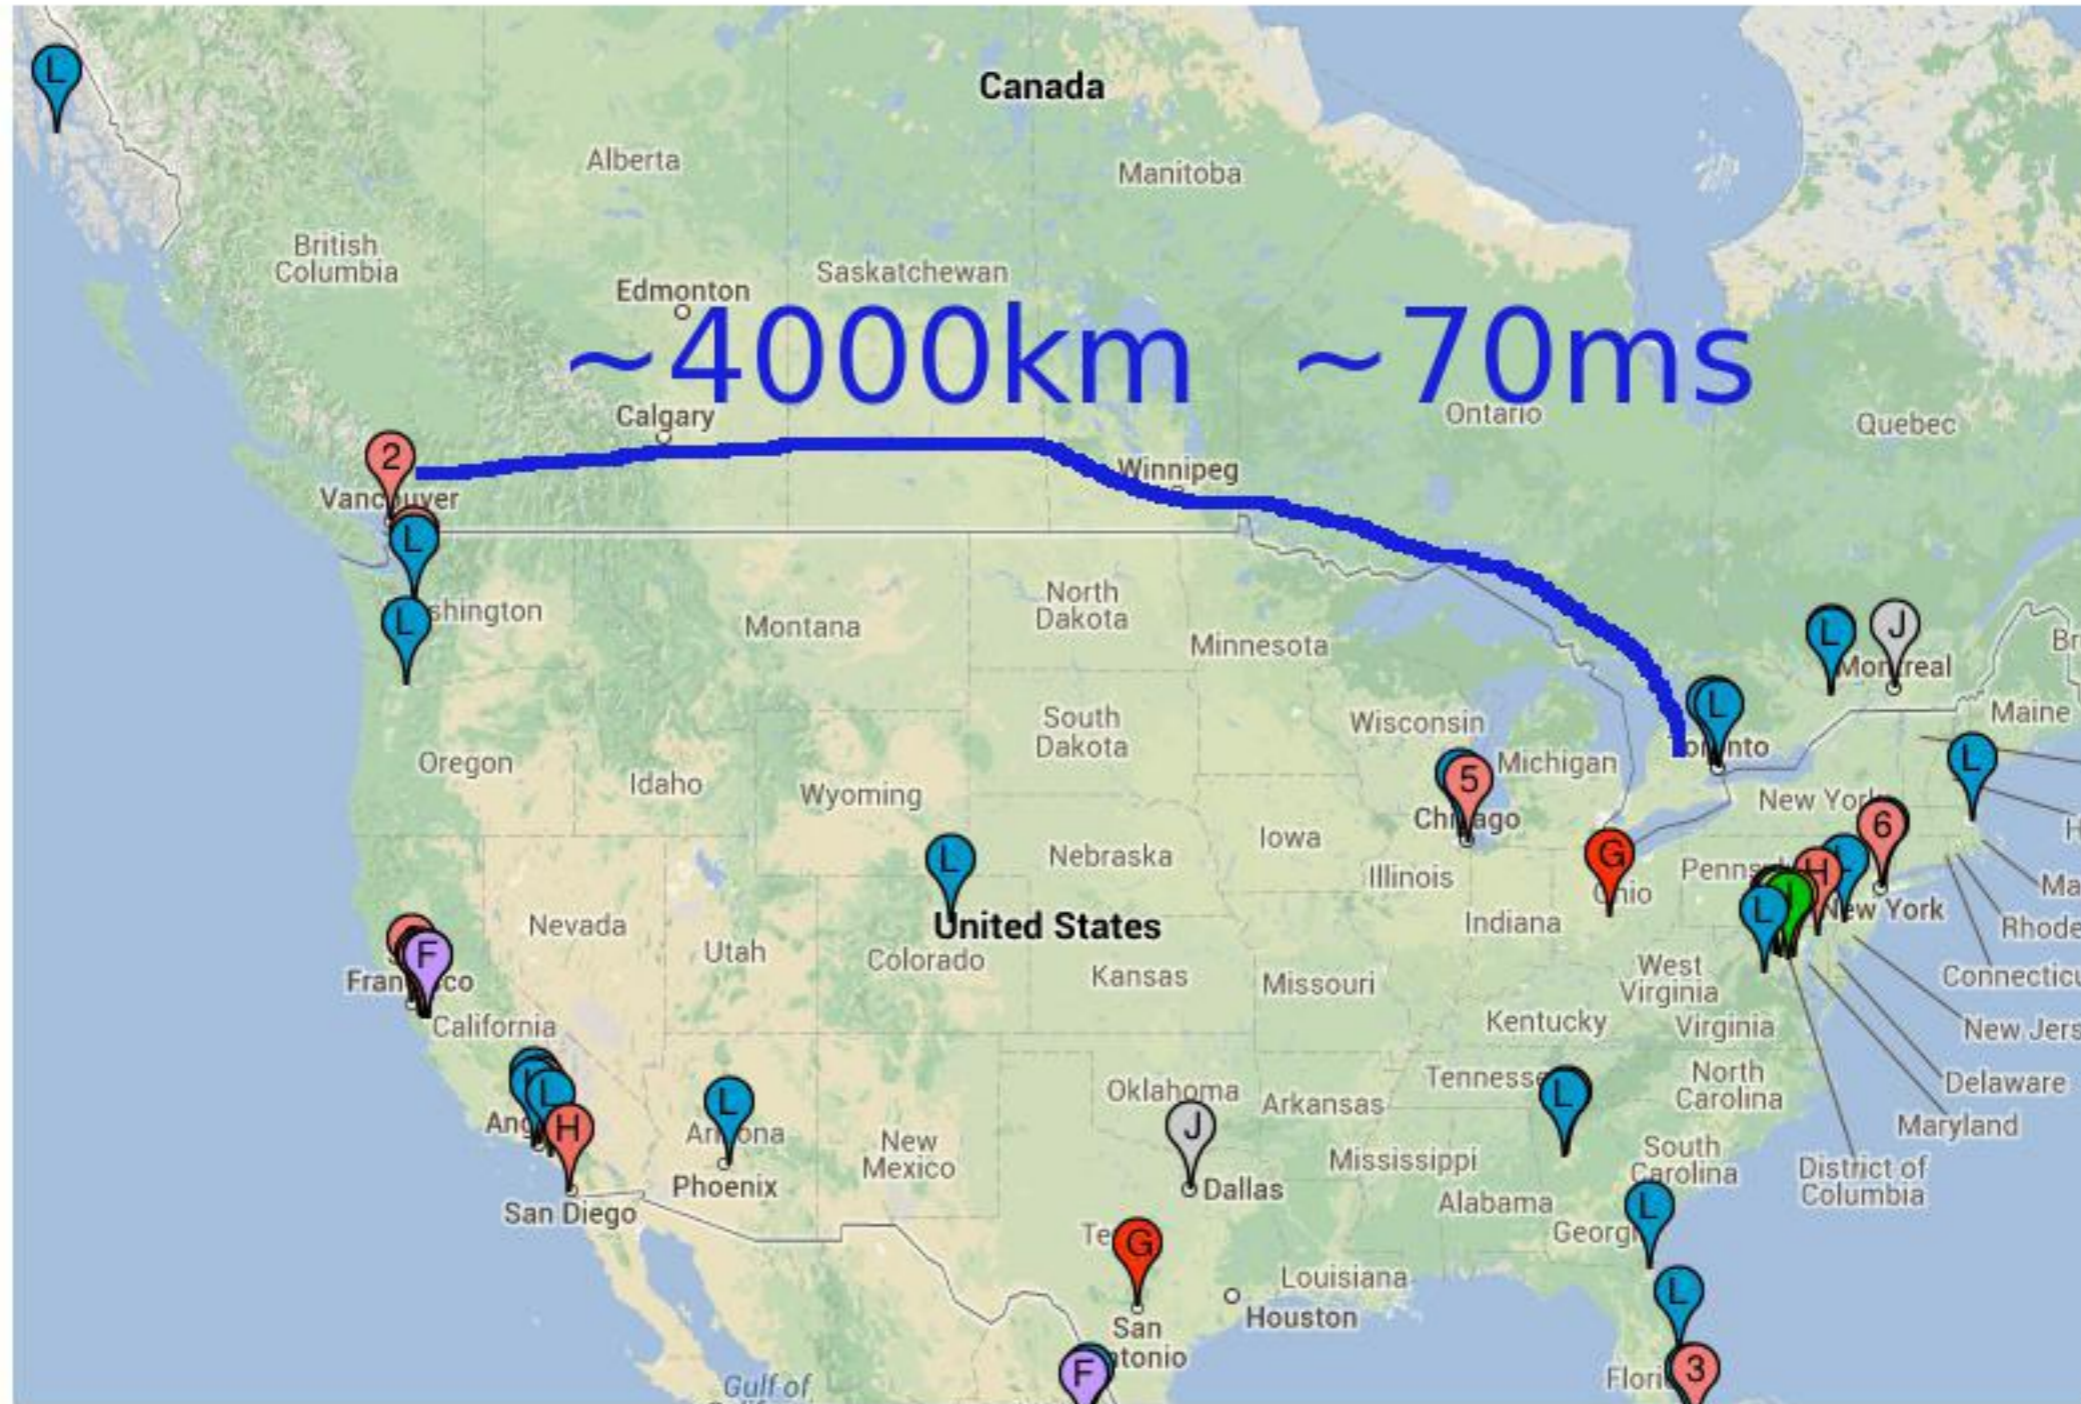


Sending traffic via "carrier hotels" in Seattle or Toronto

The rest of the Internet



North American Internet seen from Space



(root/com/net DNS server locations in North America -- <http://root-servers.org>)

Very few DNS servers in Canada... why?

DNS server operators prefer to deploy at Internet Exchange points

What are these Internet Exchange points you talk of...

Internet Exchanges started popping up to meet growth challenges identified around 1990

This is not 2001. The Internet does not work as a result of the 'invisible hand' of the legacy telco carriers.

Internet Exchanges are critical to the good functioning of the Internet.

They are locations/infrastructure which encourages traffic exchange.


[will talk about the electronic bits much later]

Compare ourselves to the world



(root/com/net DNS server locations in World -- <http://root-servers.org>)

And if we could see the fiber...



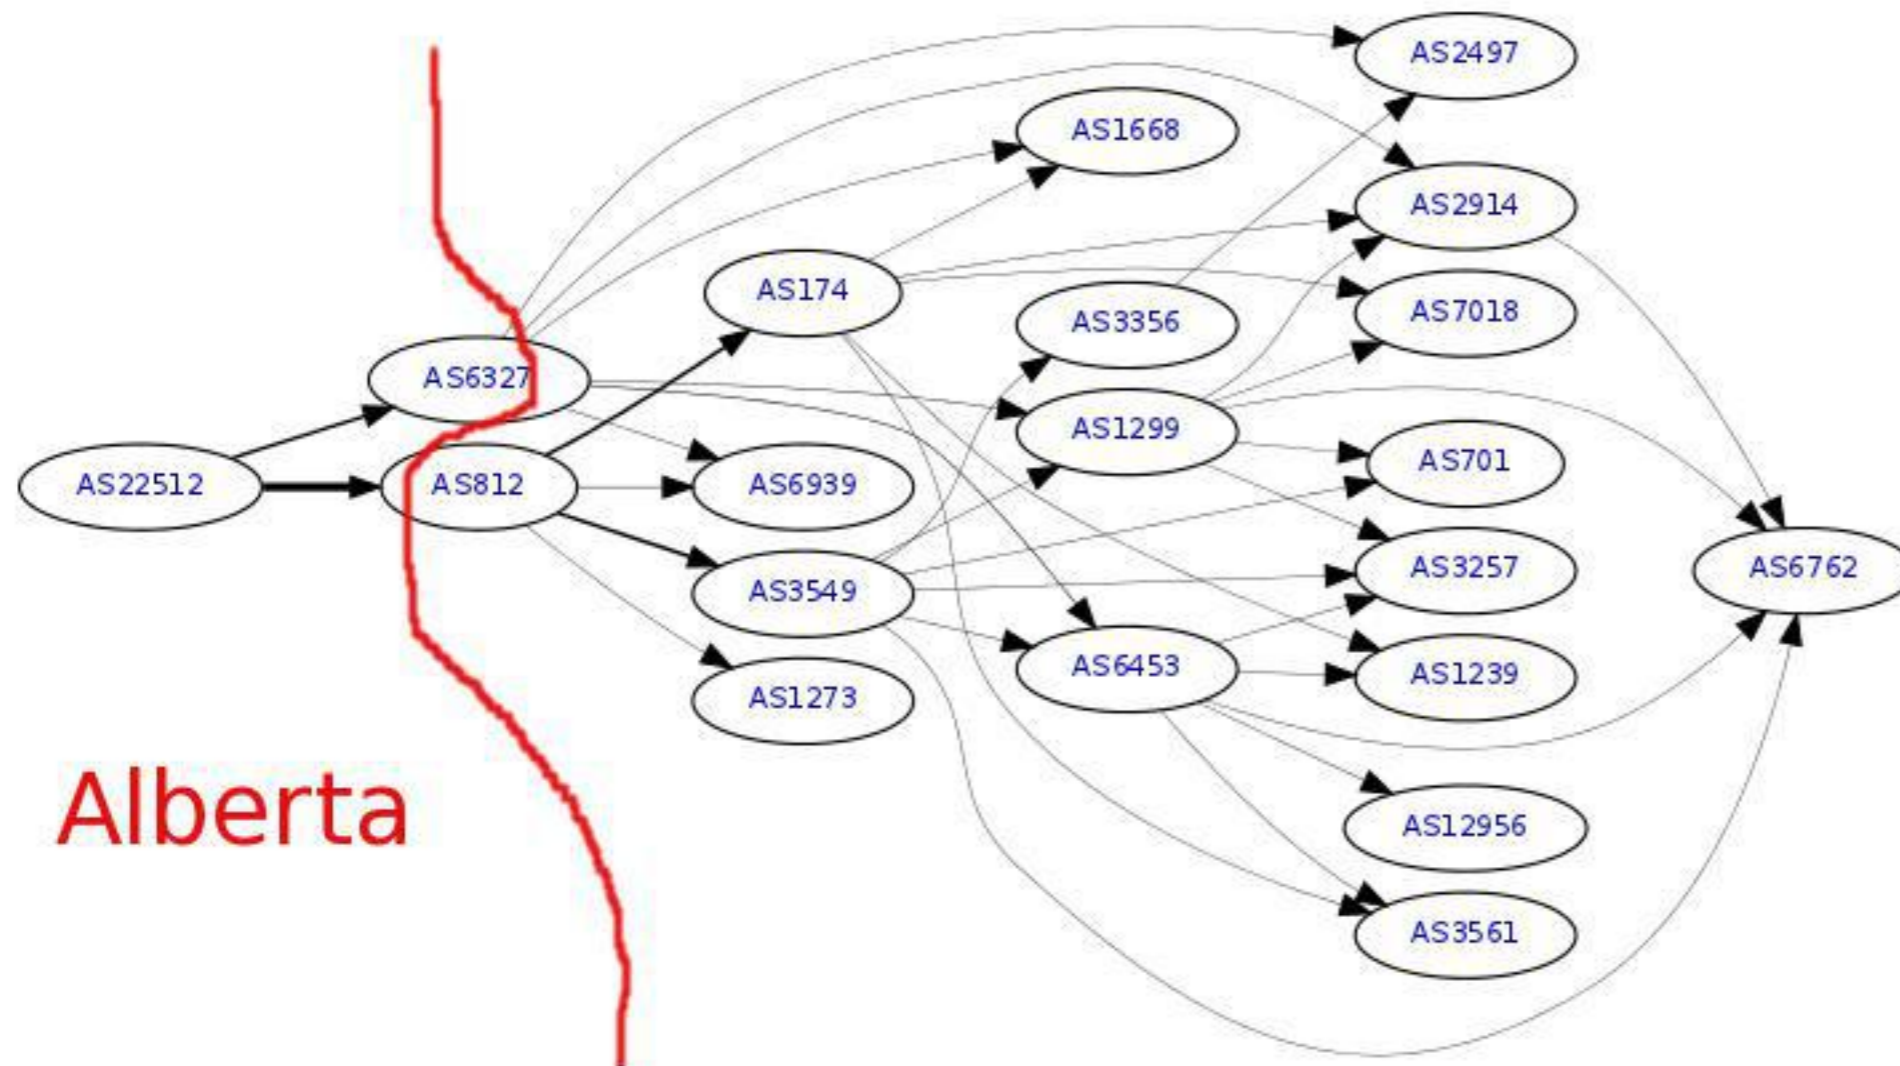
West-Canadian Internet is a line, north of another line

Our providers take all traffic to the exchange locations

Internet is an ad-hoc coupling of independent networks

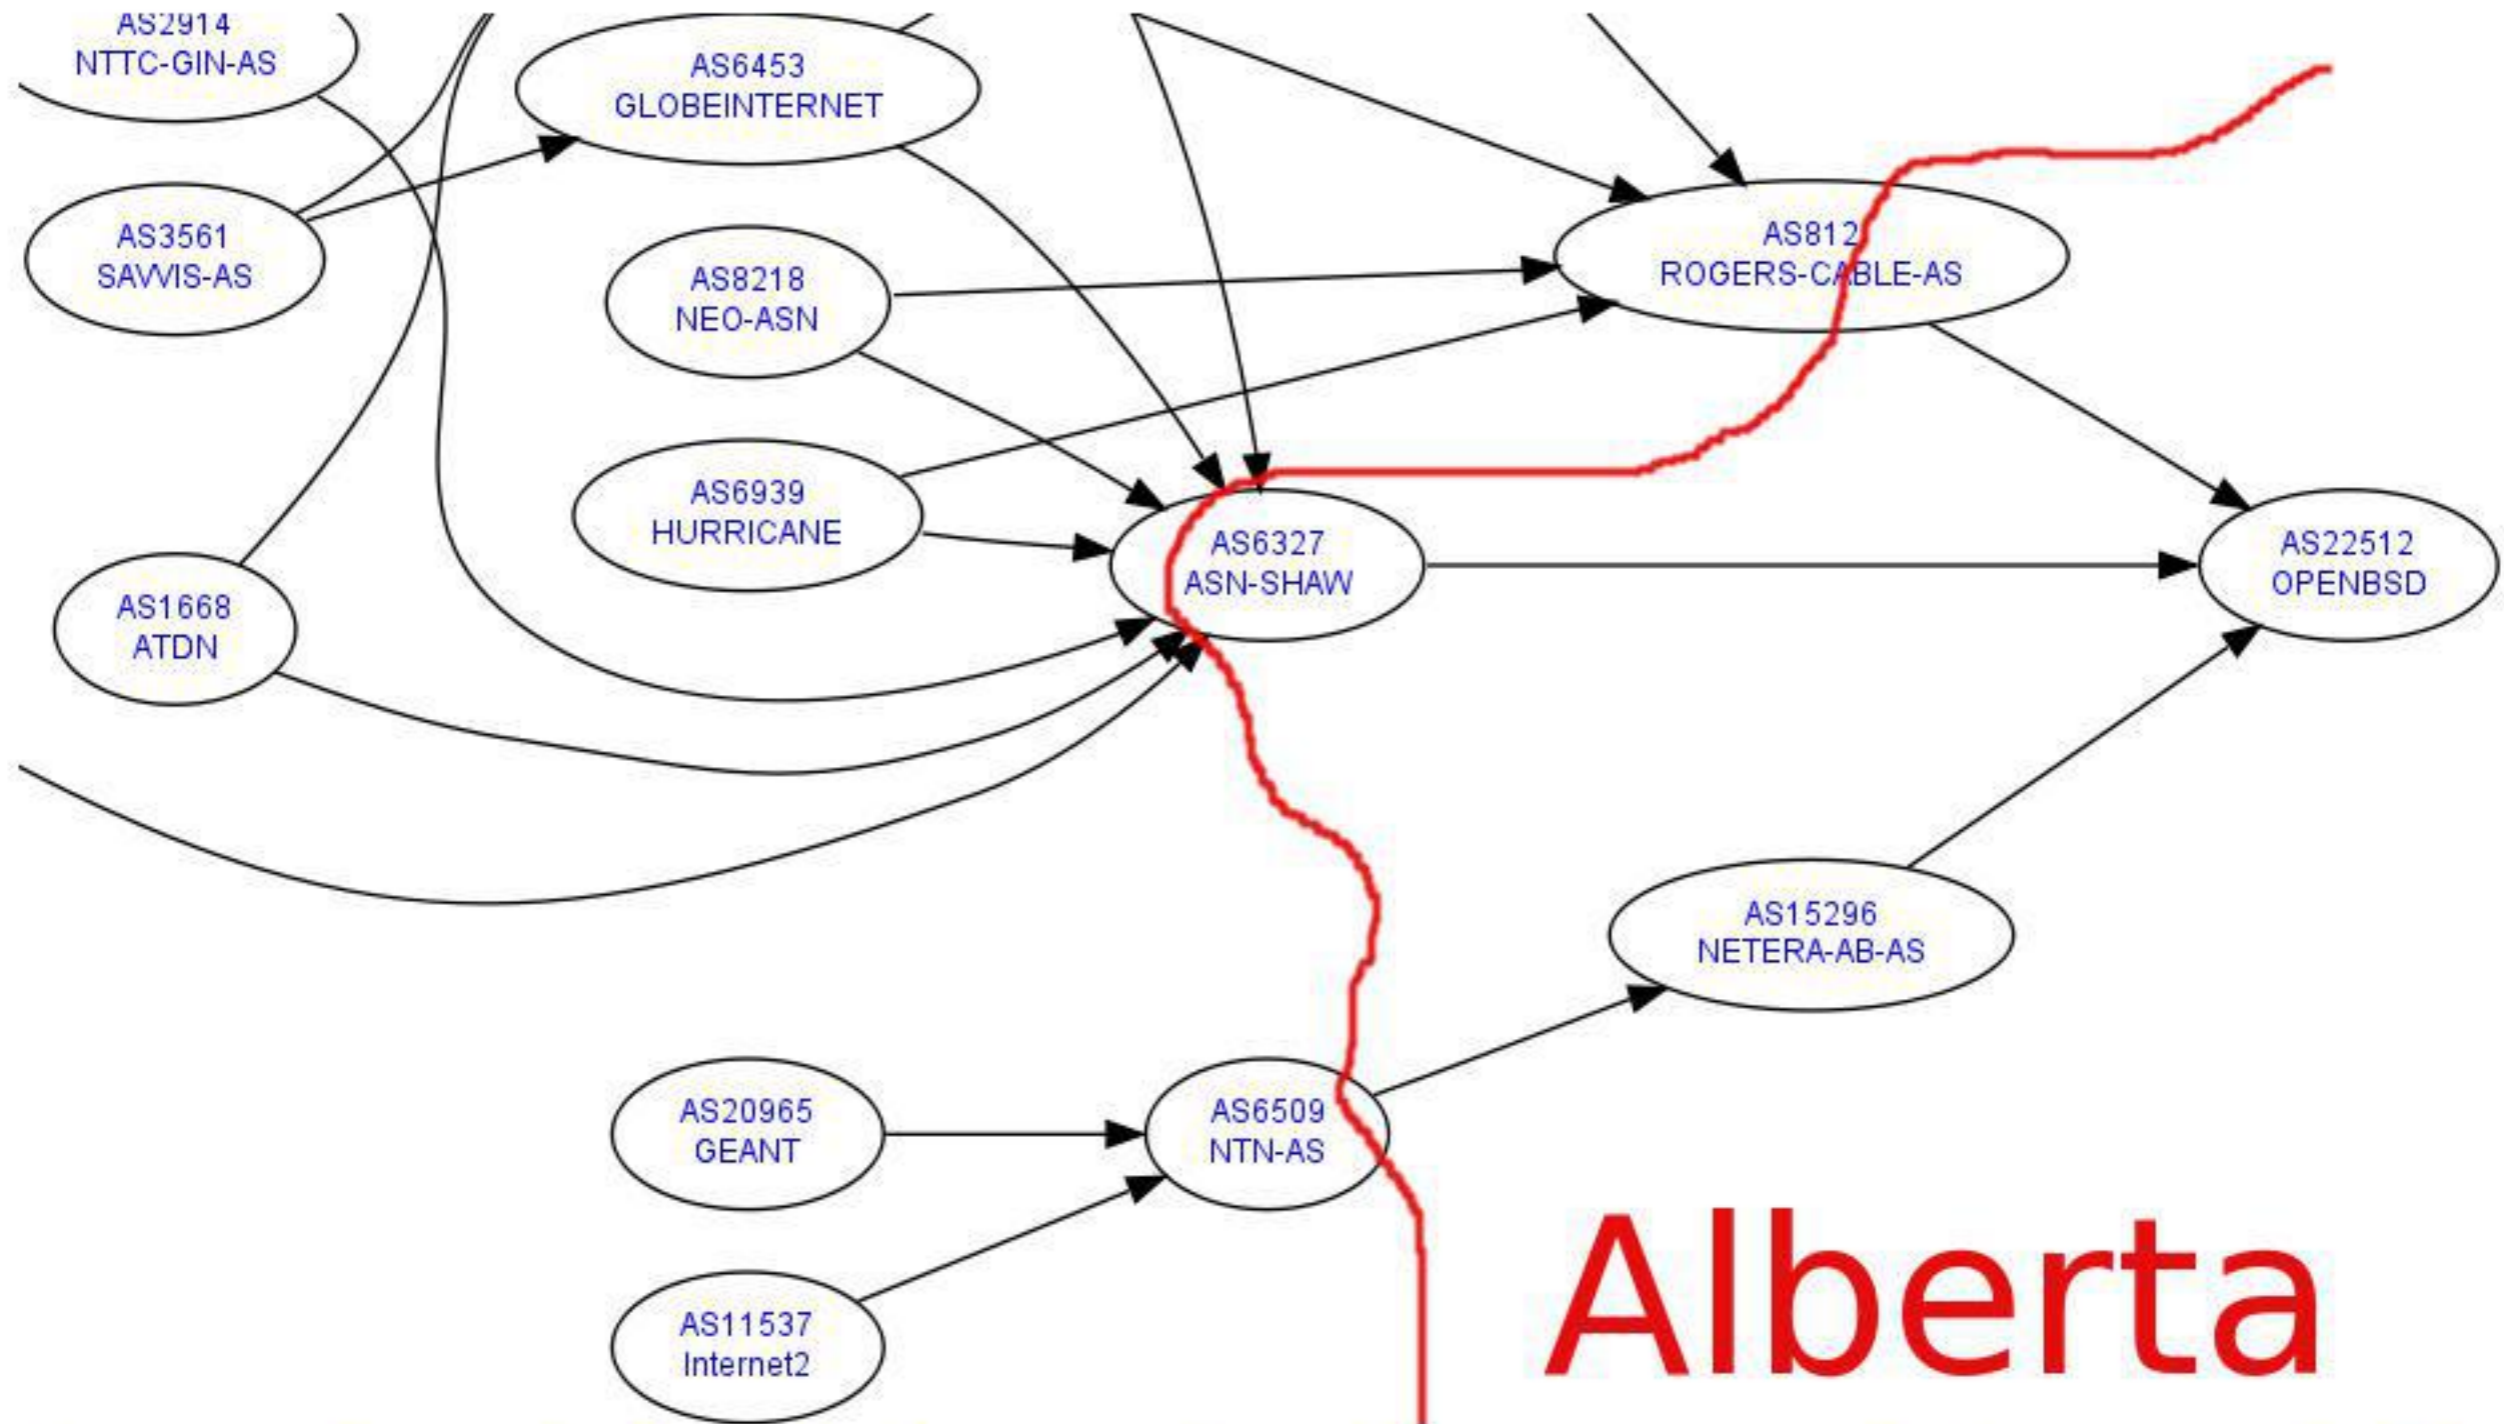
"series of tubes" --> easier to understand as "nodes"

node = a network = "Autonomous System + Network Prefixes"



Business decisions cause exchange to happen at \$cheap exchange points

44,811 networks make up the Internet

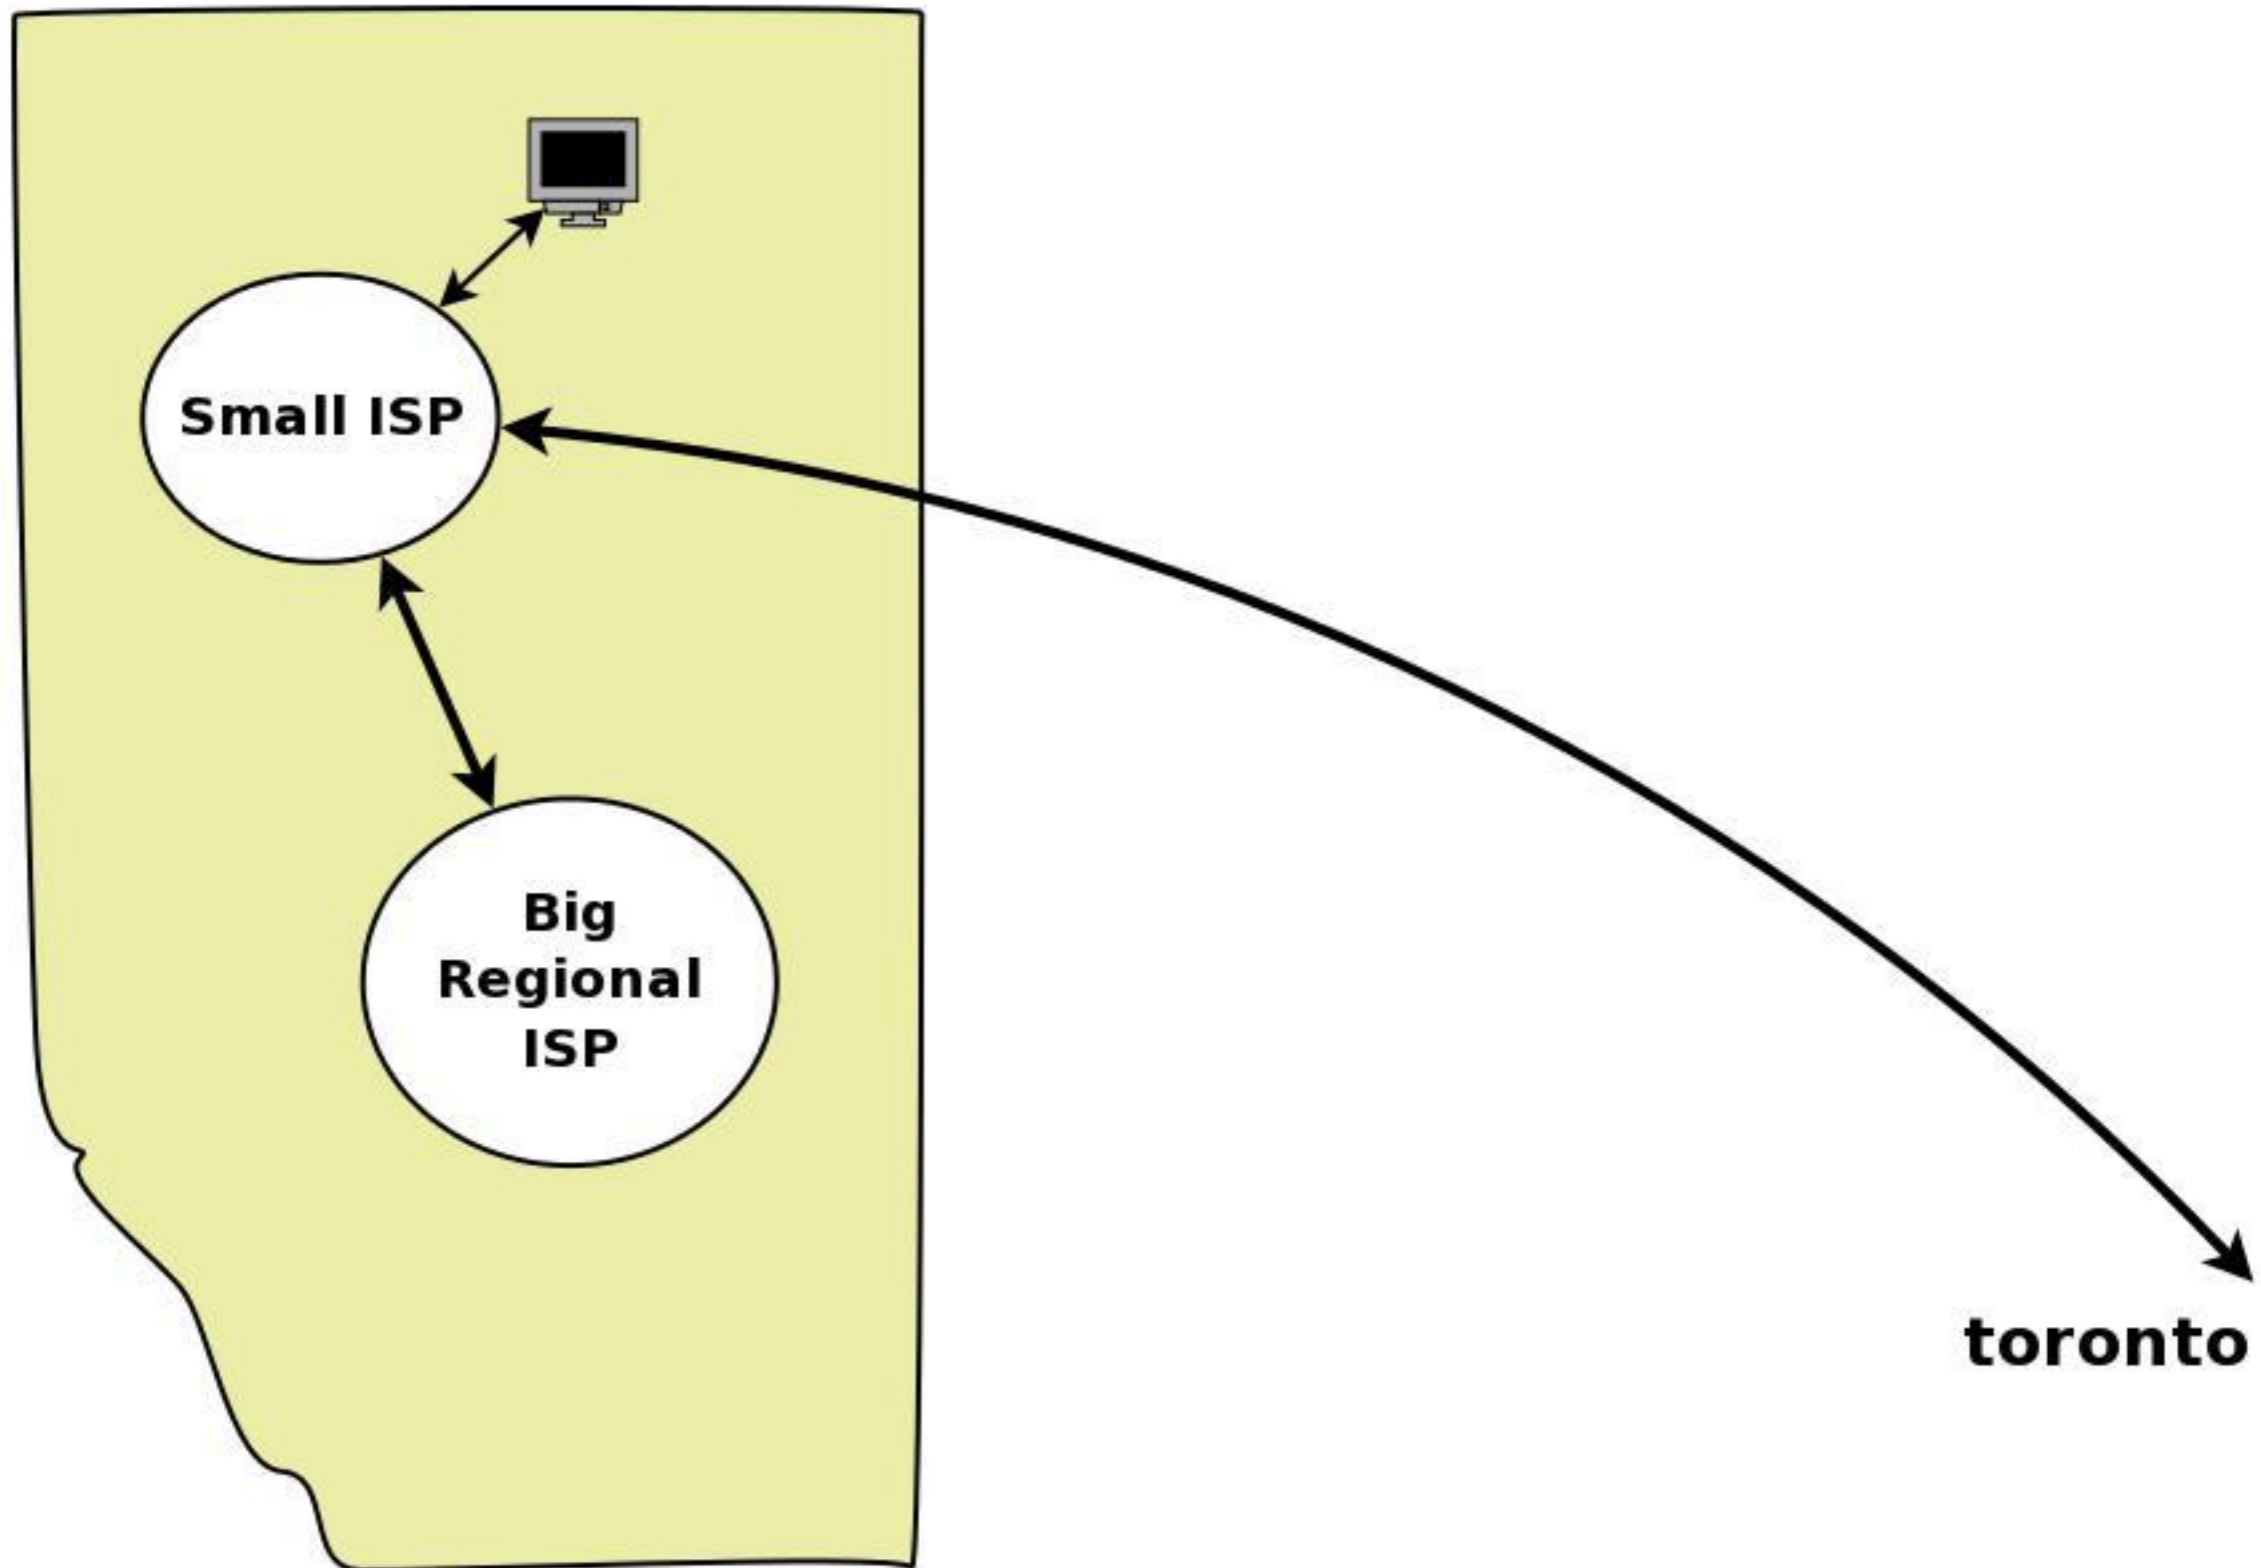


Alberta

My providers don't talk here. All exchange in YYZ/YUL or YVR/SEA

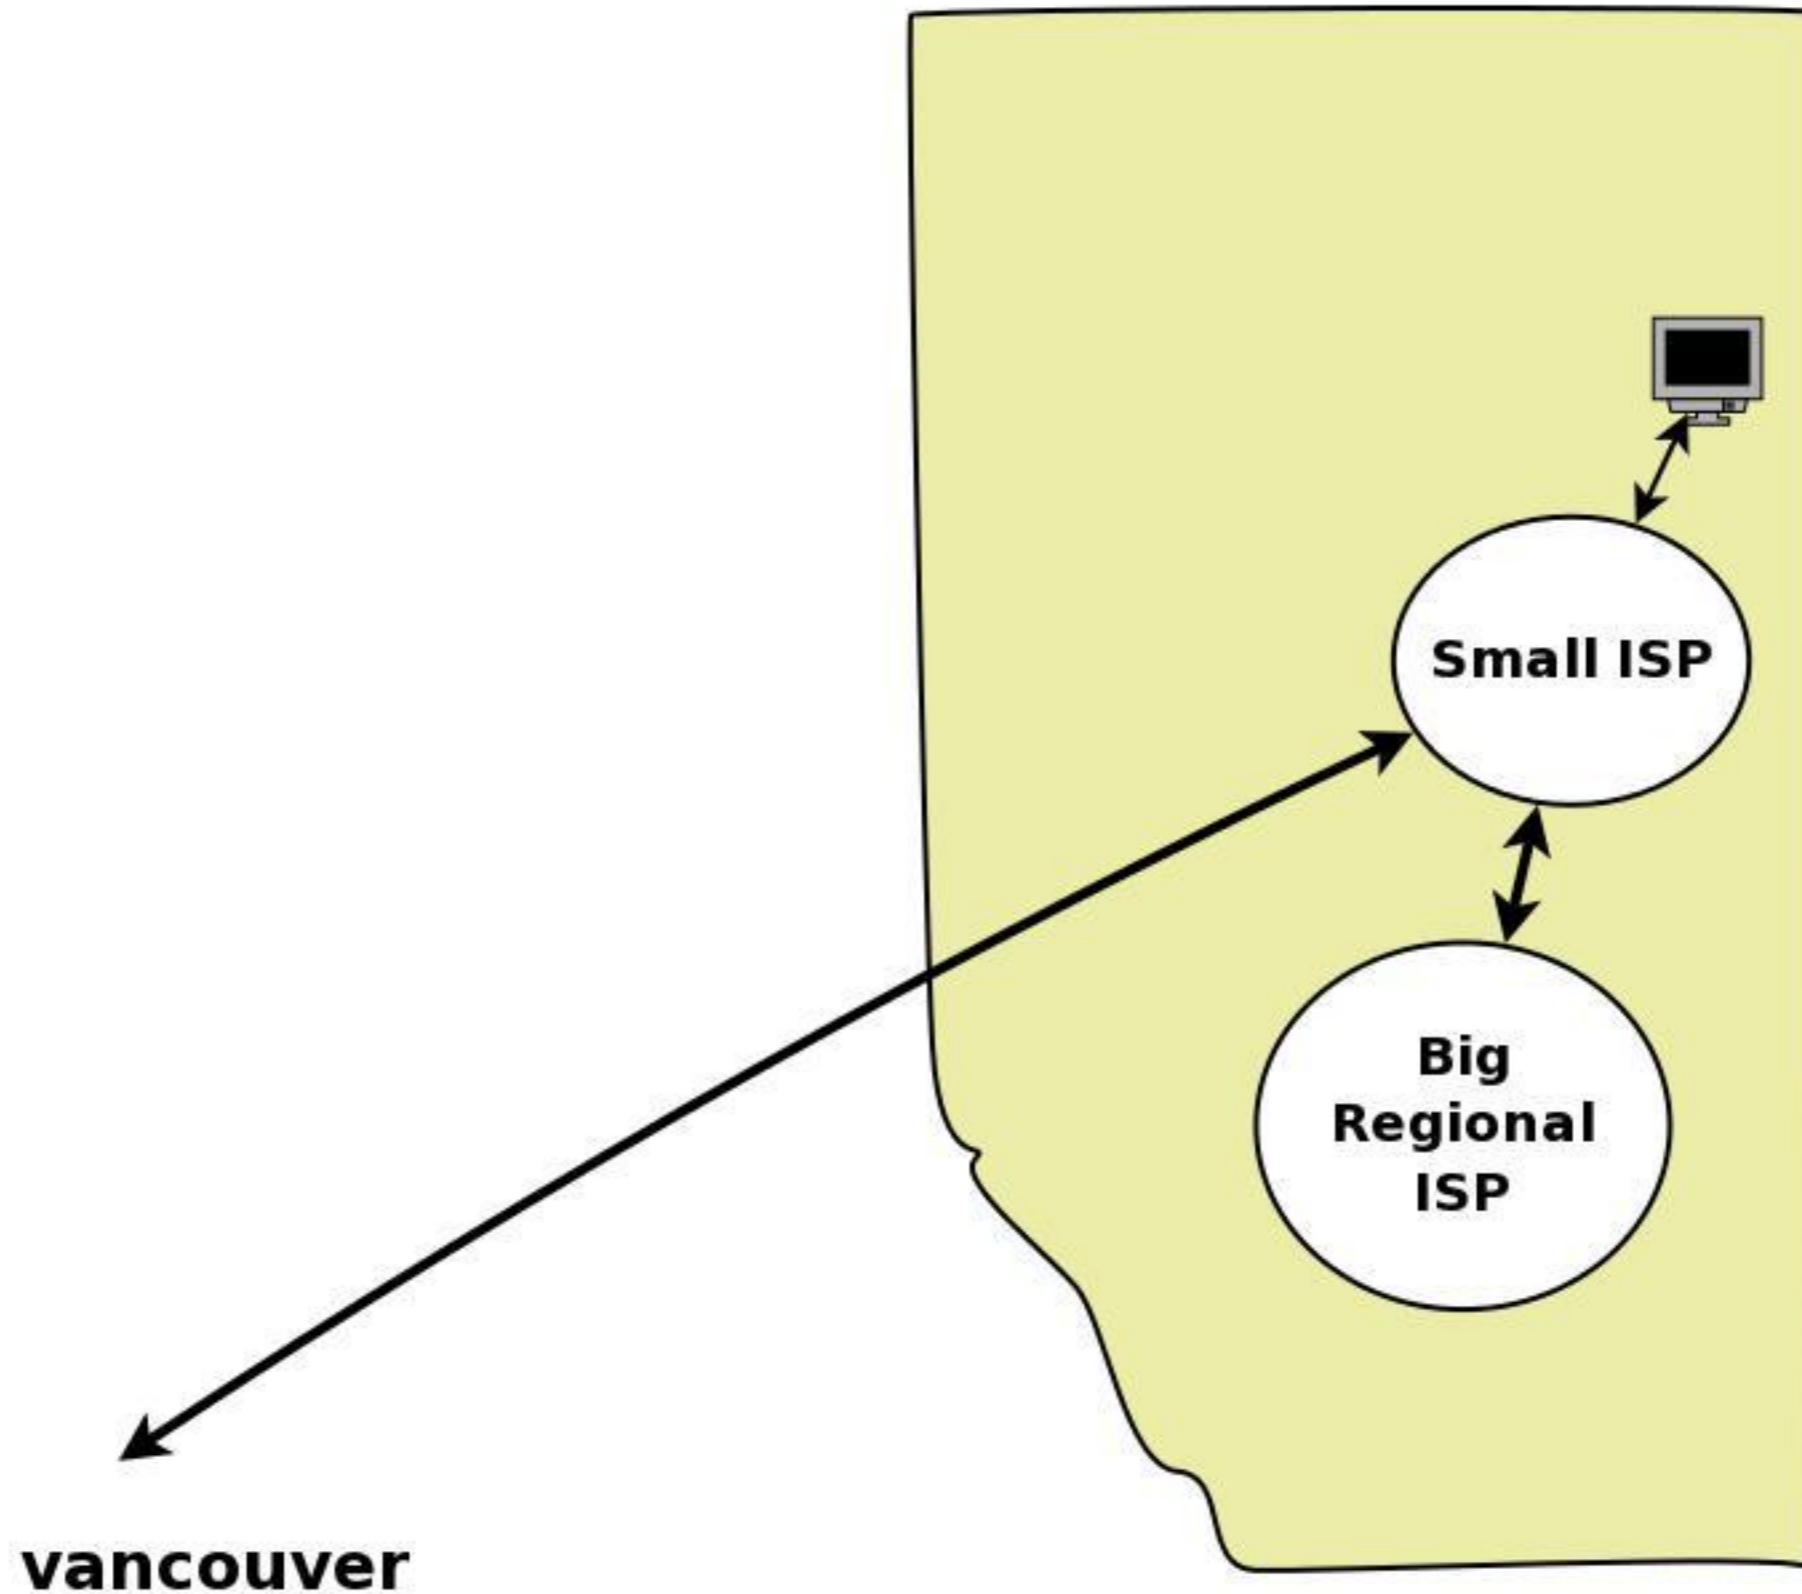
Choice even more limited when on a down-stream ISP

"My customers want Shaw and the world"



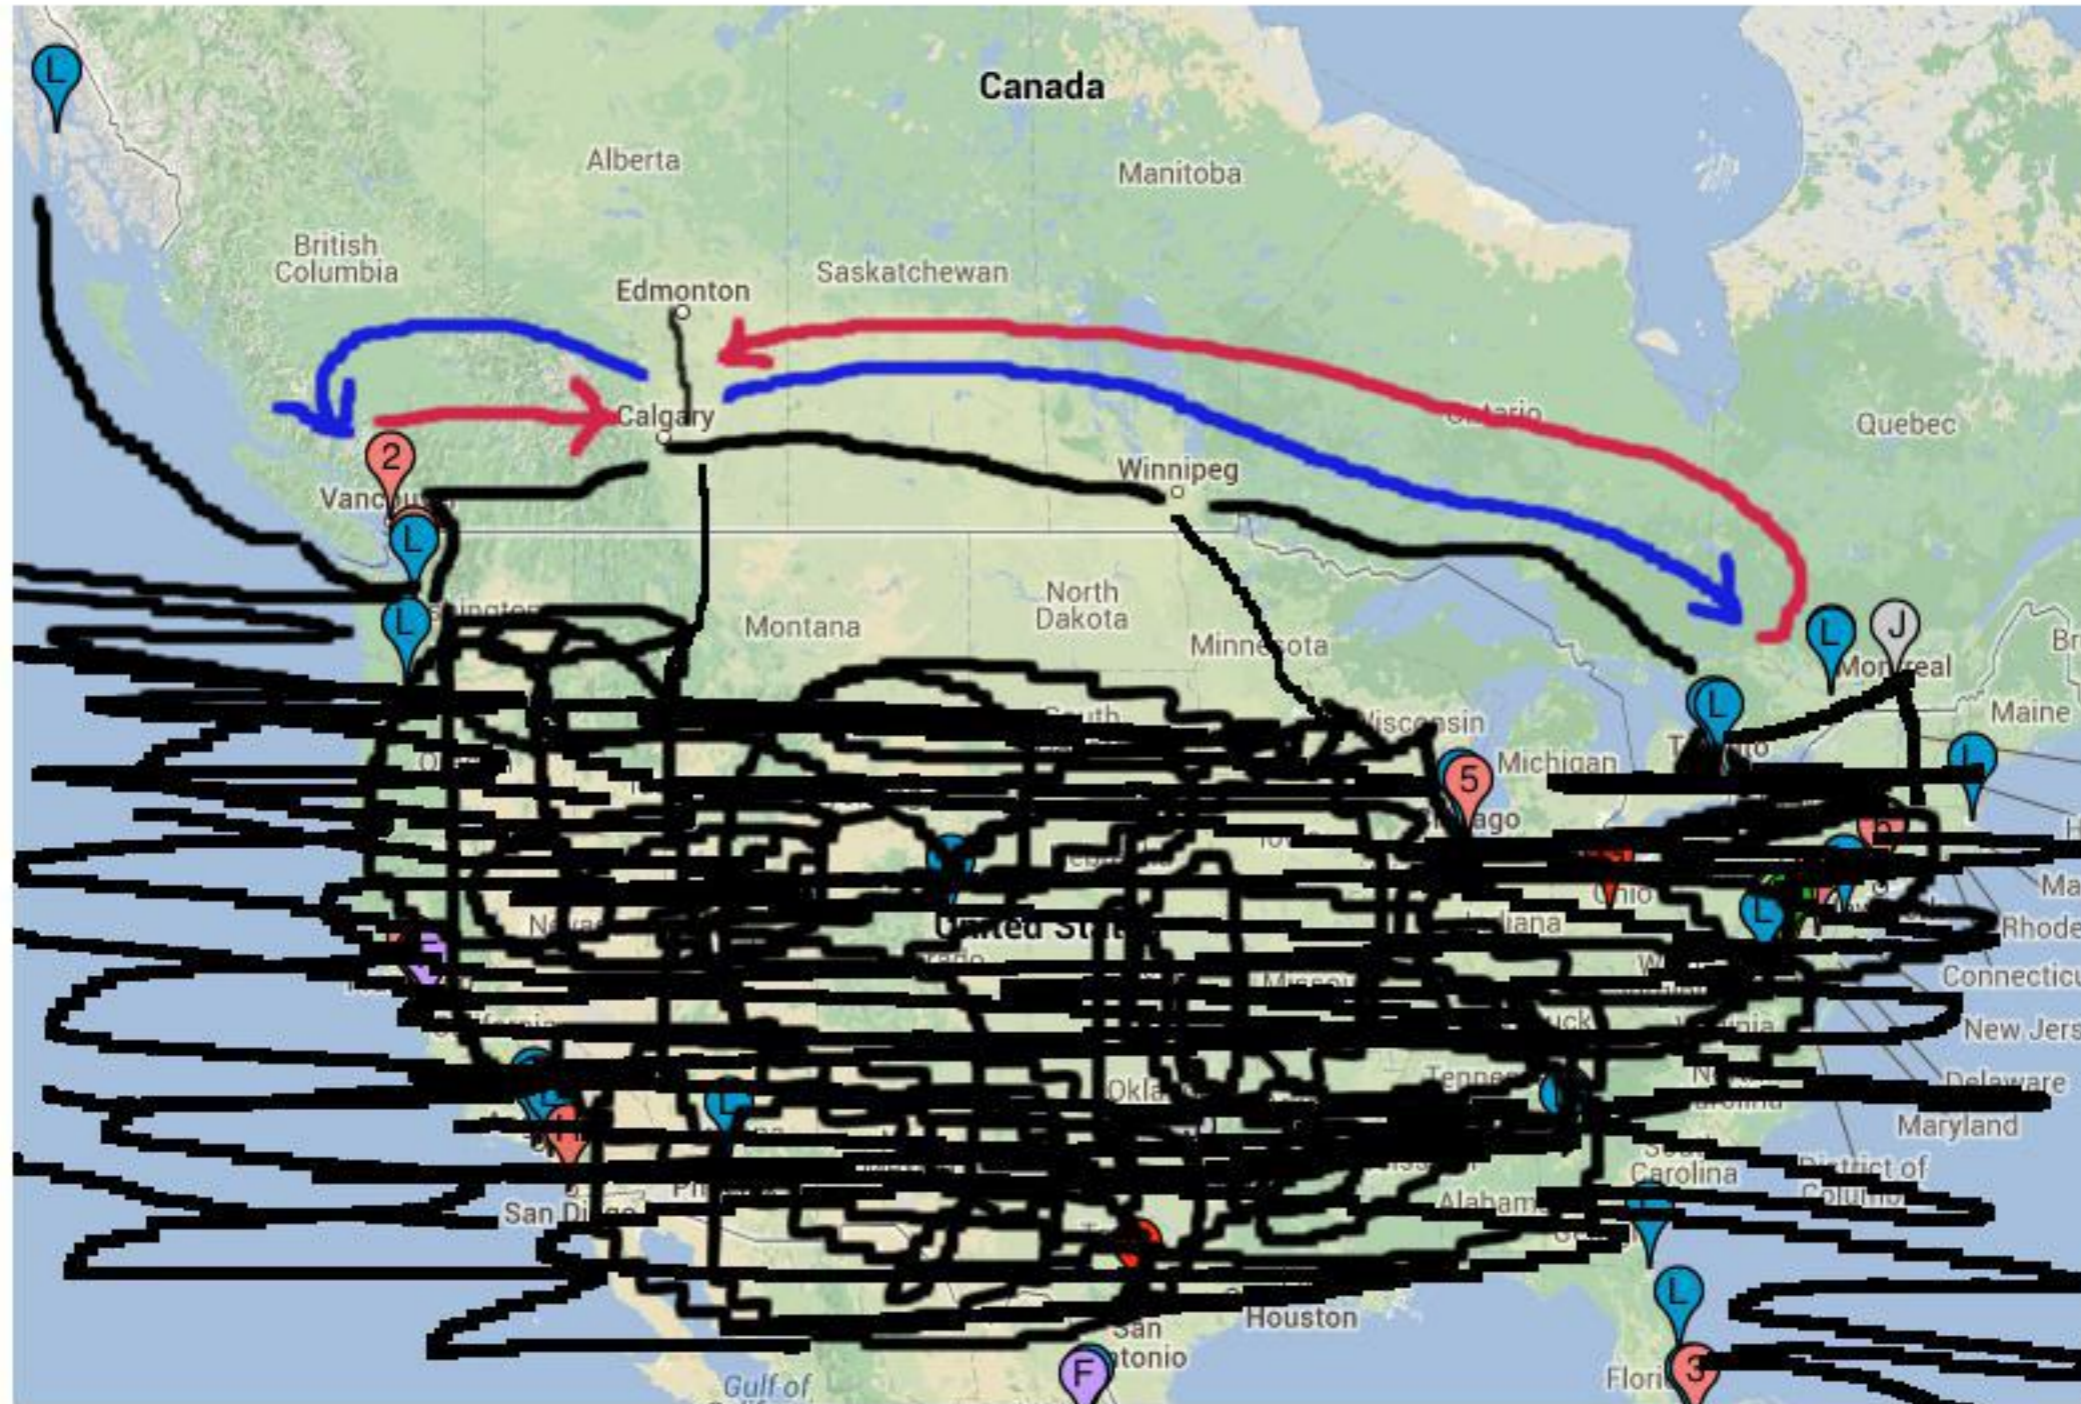
As a result, half of Alberta mid-to-small ISPs have done this

"My customers want Bell/Telus and the world"



And the other half of Alberta mid-to-small ISPs have done this

Wasting our time and our money



Our providers take all traffic to the exchange locations

... even traffic which should have stayed in Calgary

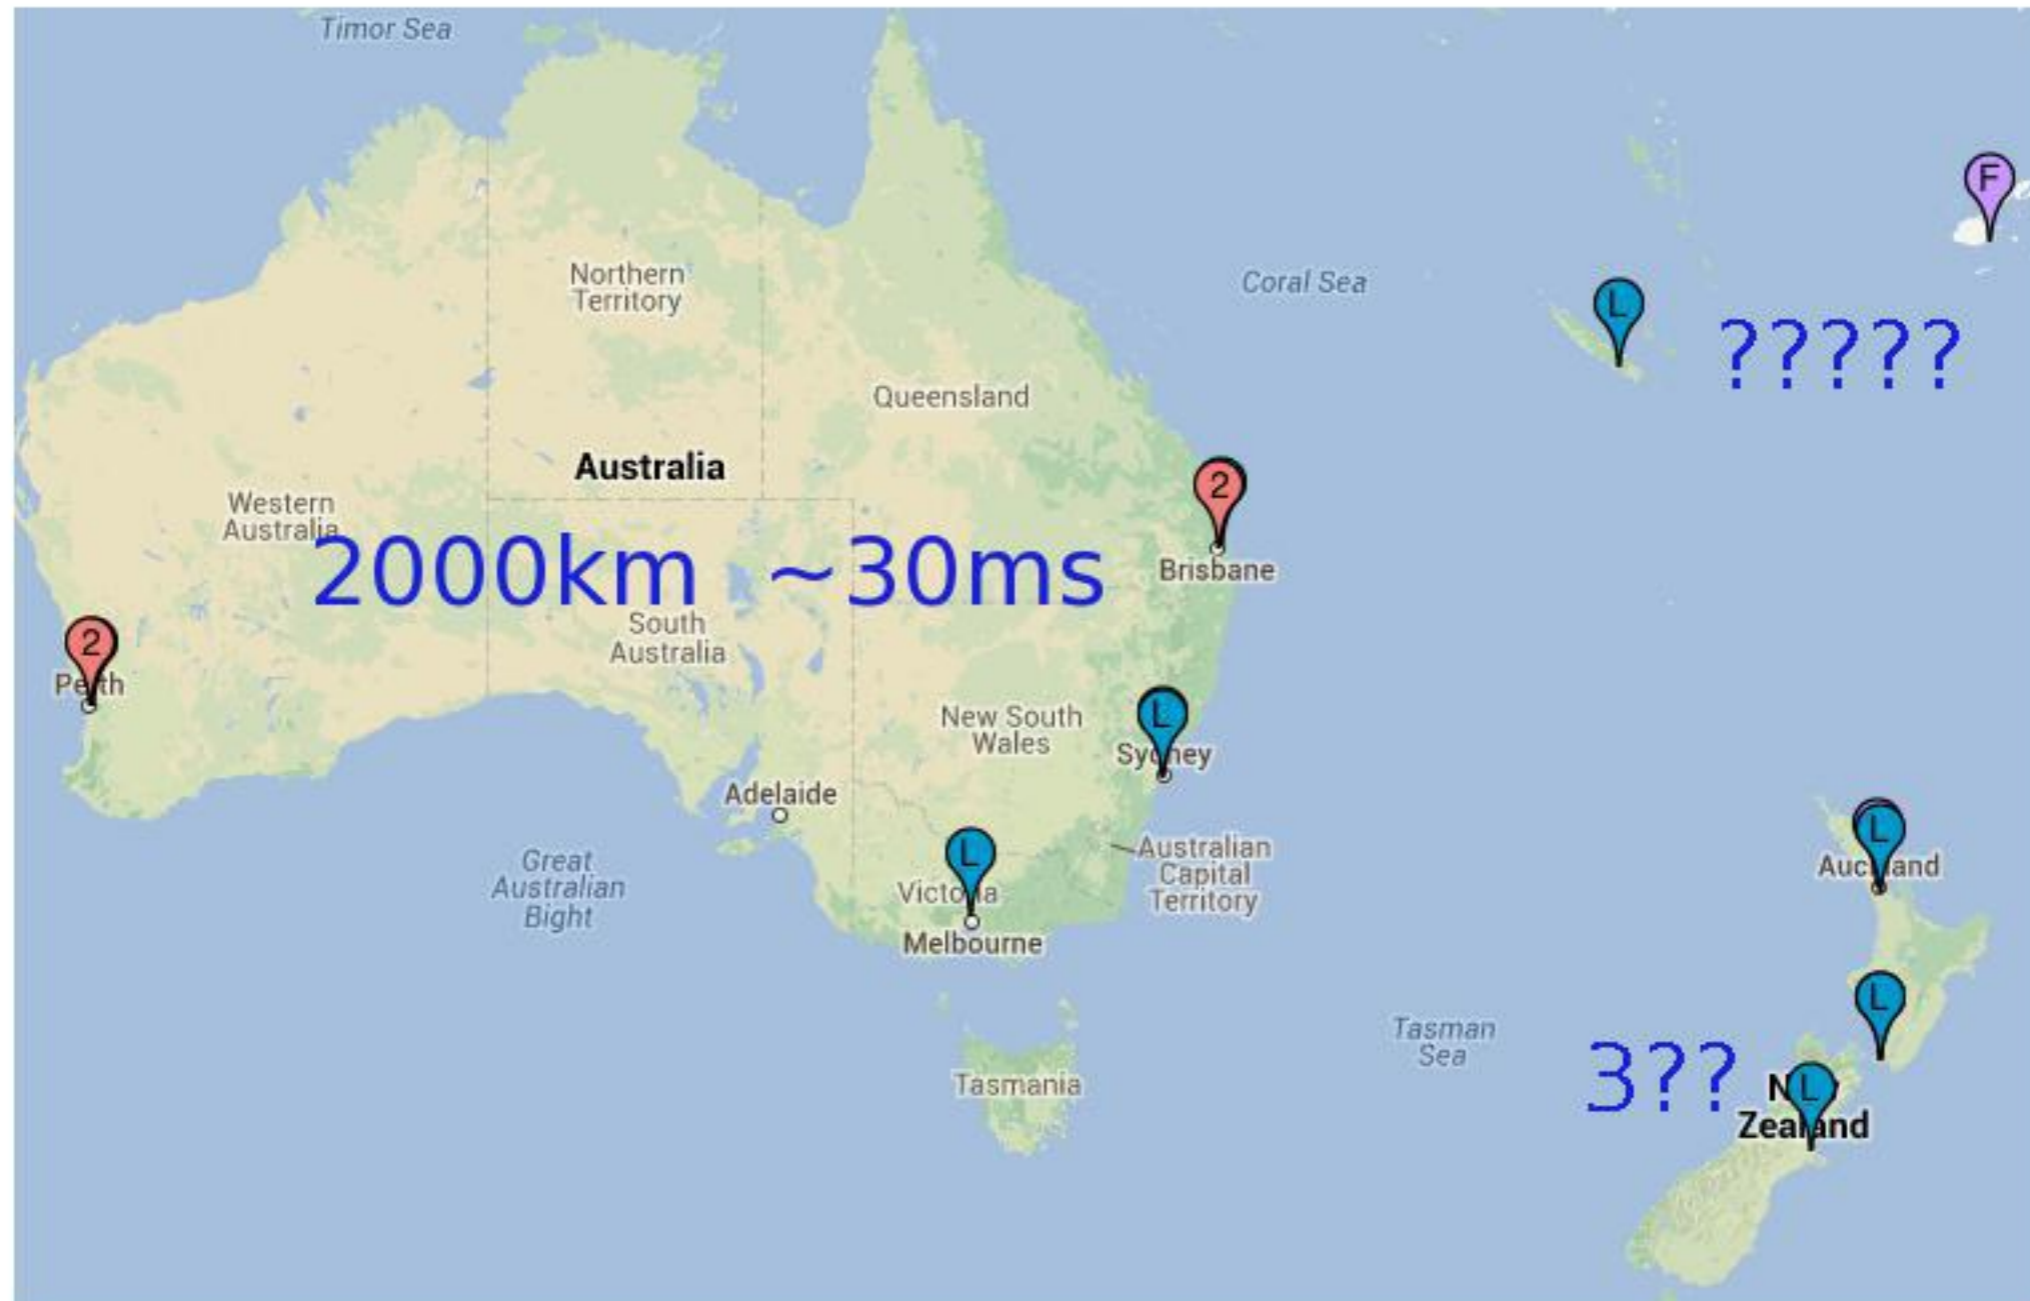
1/3 of Canadians



Under-served

Stay on our network or you will suffer

"but Canada has Geographical Challenges"

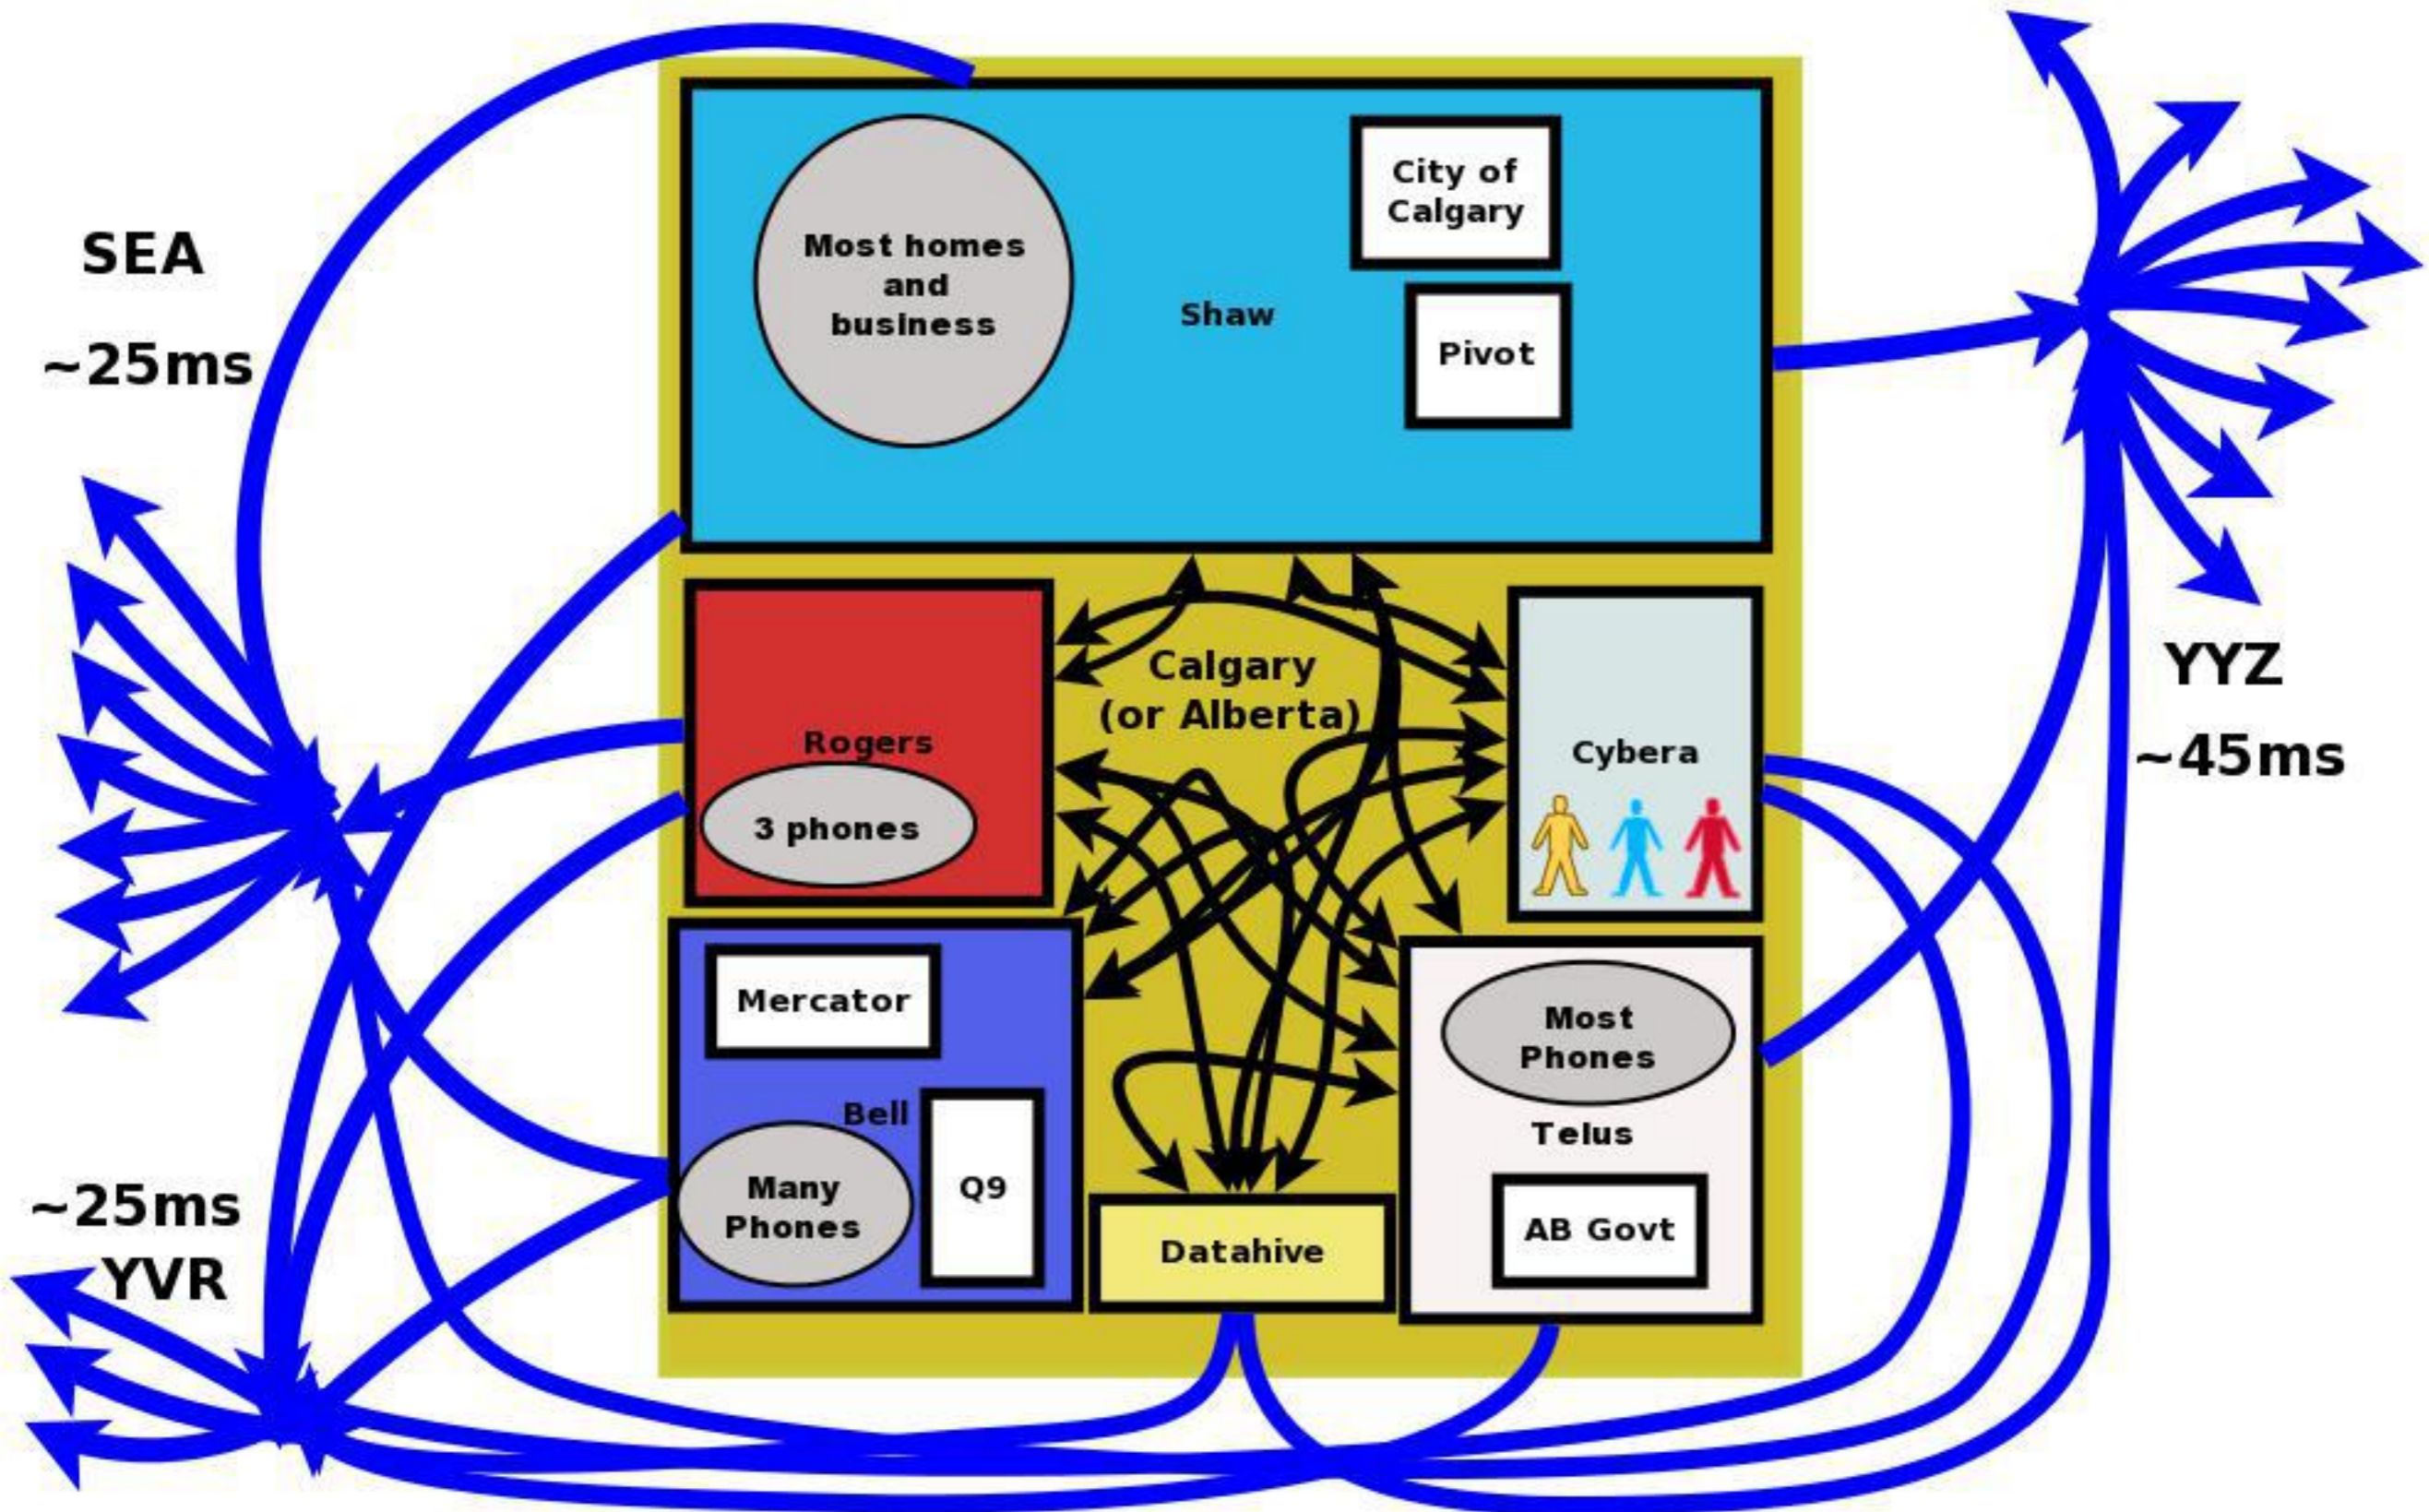


Kangaroos and wallabys in Alice Springs have better service than Calgarians

Australia + New Zealand: smaller population, more local exchange

New Zealand ISPs do not exchange traffic in Australia

Calgary: a model for the future



Providers are at war, so getting this won't happen fast

Some of that jumble can be improved using a switch

ADVICE ON INTERNET SAYS TO START SMALL

- A (very good) ethernet switch
- Some ability to get at fiber + copper parts
- 2-3 management machines

- A datacenter willing to play along (DataHive)
- Some peers who are eager from from the start

- Some know-how (... with a little help from our friends...)
- A community of techies, network operators, and supporters
- Communications with the global network community (to get CDNs)

- Willingness to grow/adapt as necessary

Not communism. Most IX are co-ops run by network geeks

How do the electronic bits get used?

An Internet Exchange is intended to be the very fast but inexpensive network path

Peers run BGP and have other providers -- so they can deal with failure!

Small network interruptions are OK

Promises or SLA's don't matter; BGP fails over invisibly

Not rocket science

Volunteers can do this, and it is very common; if it grows, change the model

Don't take my word -- go read about SIX (Seattle) or hundreds of exchanges around the world!

But how do we make this happen?

How does an Internet Exchange blossom?

Content servers come if there are enough eyeball networks

- then serve up content for free to....

The eyeball networks which come because content is cheap & fast

- bring their eyeballs to their advertising/content ...

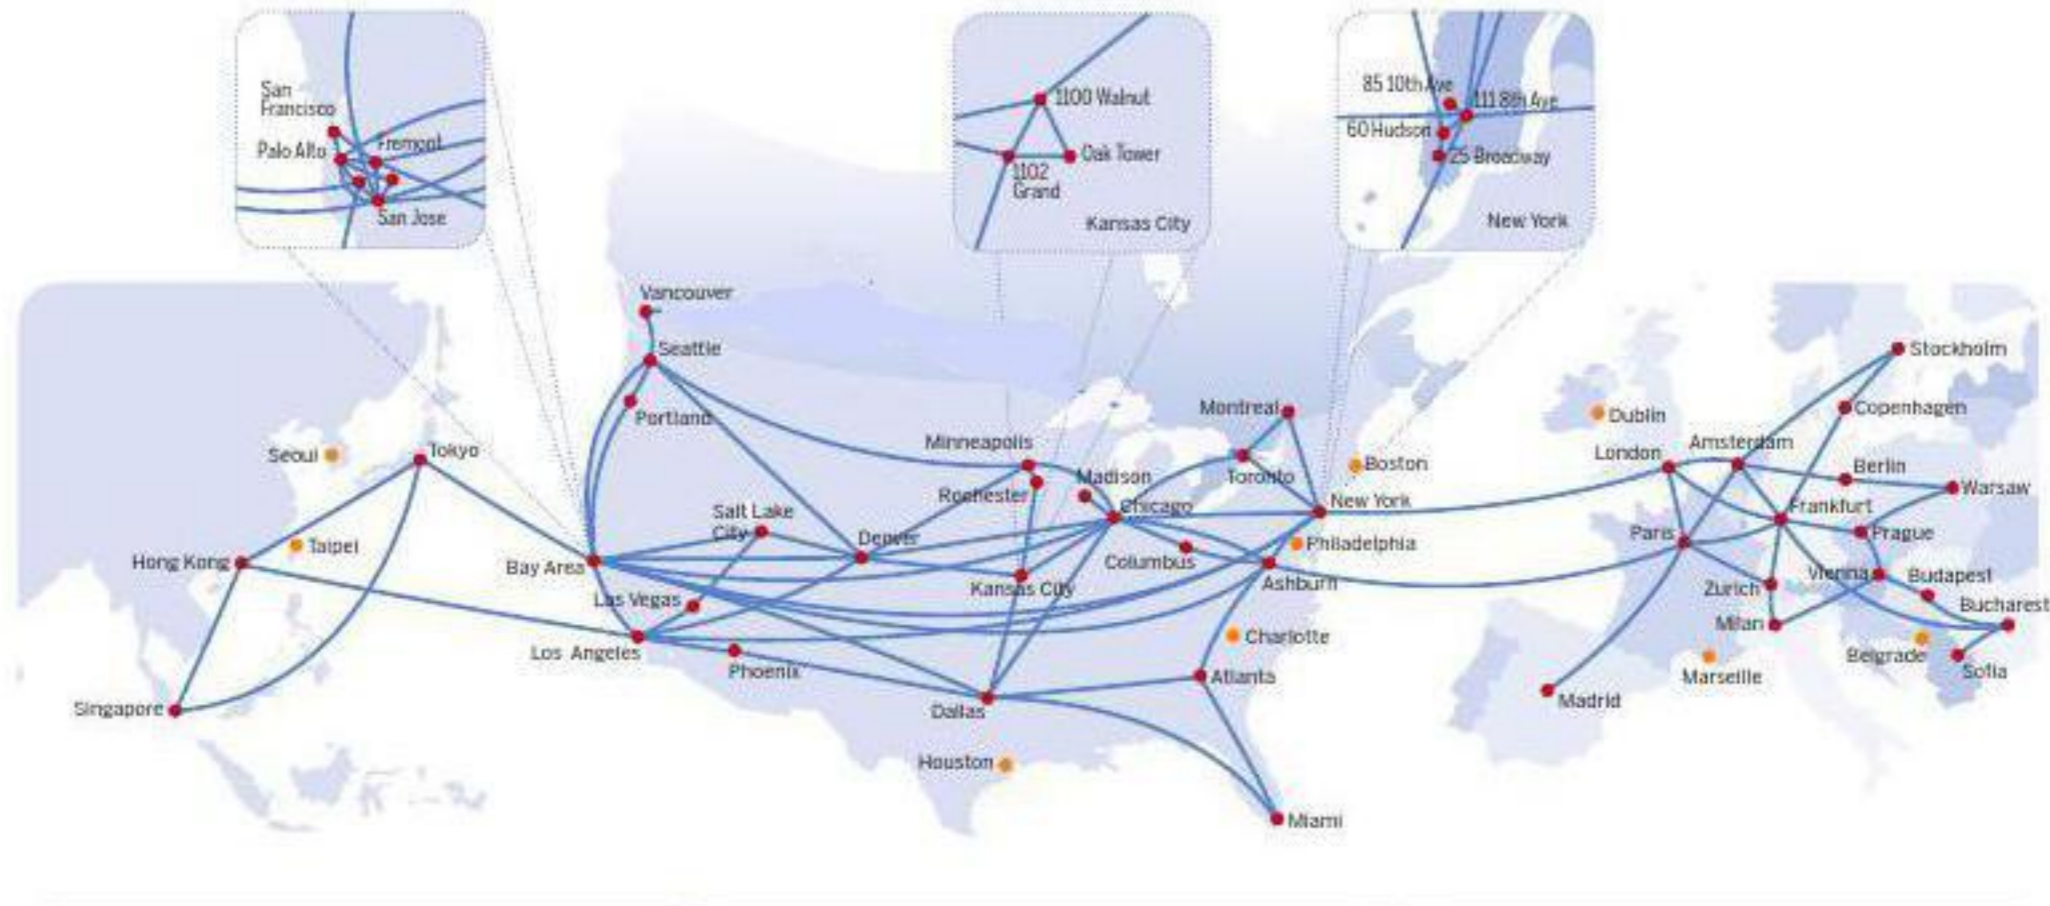
Put many networks in the same location, what happens?

- cheap carriers arrive to sell at that location
- everyone wants cheap network, right? No..



But if everyone waits for everyone else, nothing happens!

I hate when nothing happens.



I spent weeks talking to global carriers, especially Hurricane Electric

I believe in having options and opportunity, and not being held hostage

We need a carrier to help us put a pin on the map!

Other people hate when nothing happens, too.

Other concerned people saw a problem with a 3-month delay

Their perception was there was stalling / positioning going on

They decided to form a board independent of AlbertaIX

YYCIX is formed: Bob Beck, Bob Kitella, Mark Cordingley

Asked me to install a peering switch at DataHive right away

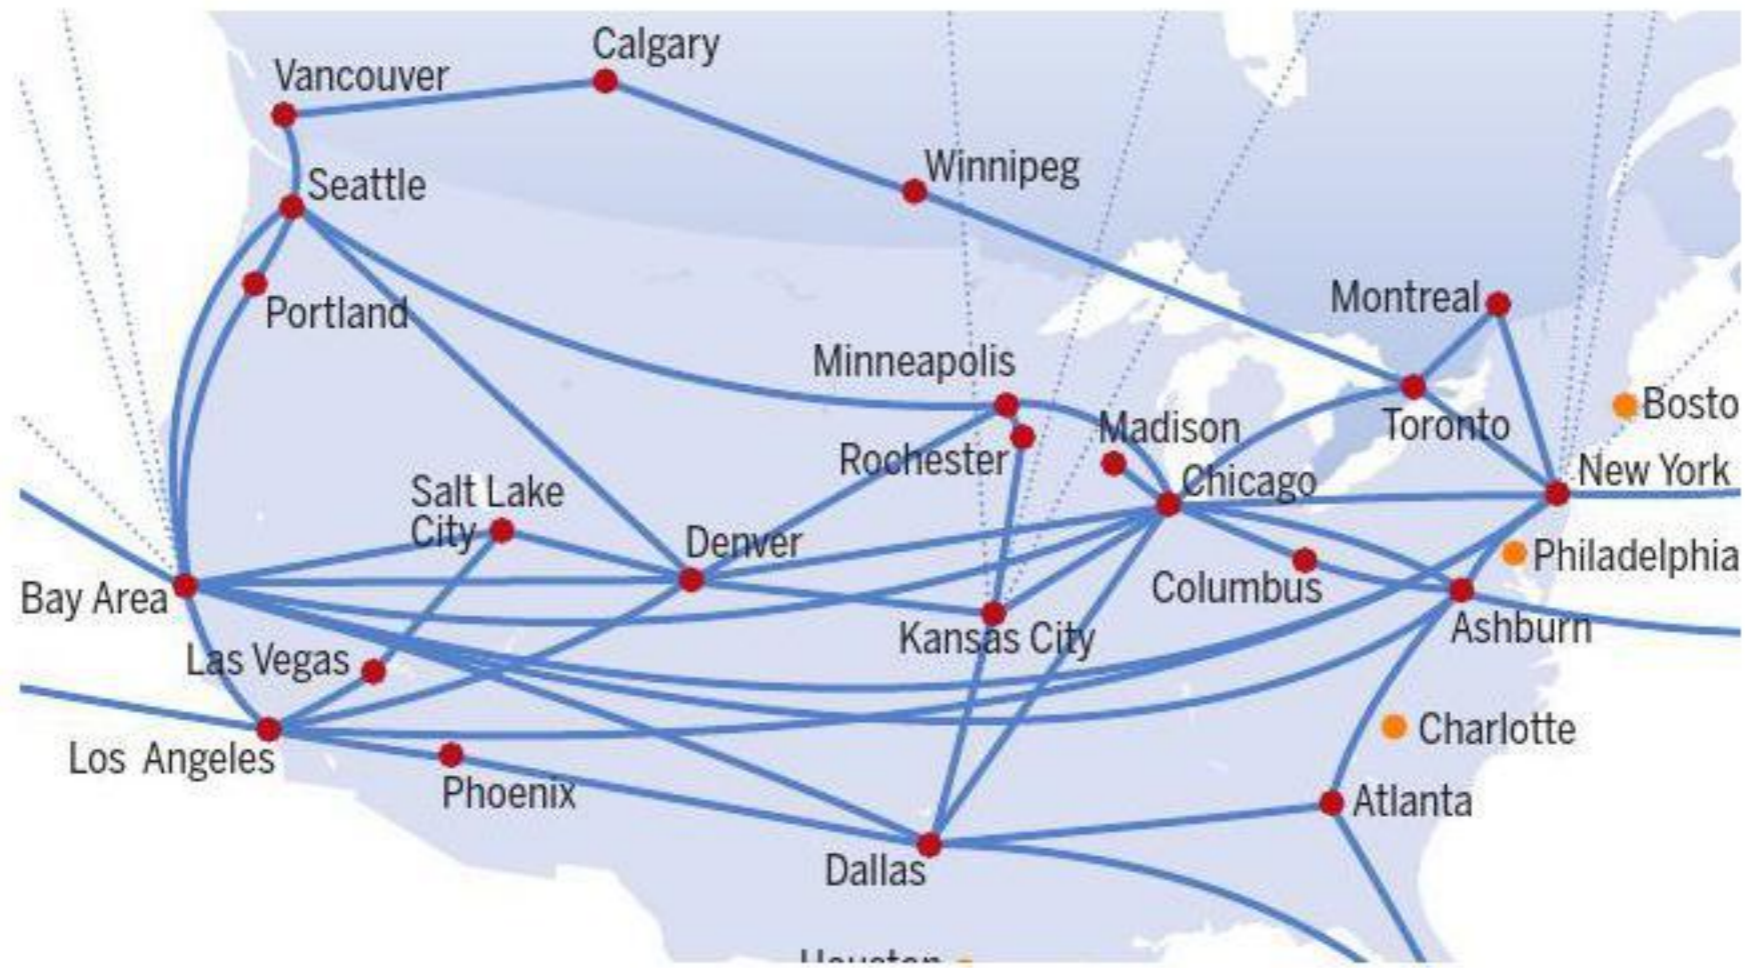
Easily done.

... some conflict developed, but then:

..... Wait for it



!!!!!!!!!!!! Pin on the map !!!!!!!!!!!!!



Canada

Equinix Toronto TR1

151 Front St, Suite 706, Toronto, ON M5J 2N1

Cologix Vancouver

555 West Hastings St, Vancouver, BC V6B 4N4

Cologix Montreal

1250 Rene Levesque West, Montreal, QC H3B 4W8

DataHive

840 7th Avenue SW, Calgary, AB T2P 3G2

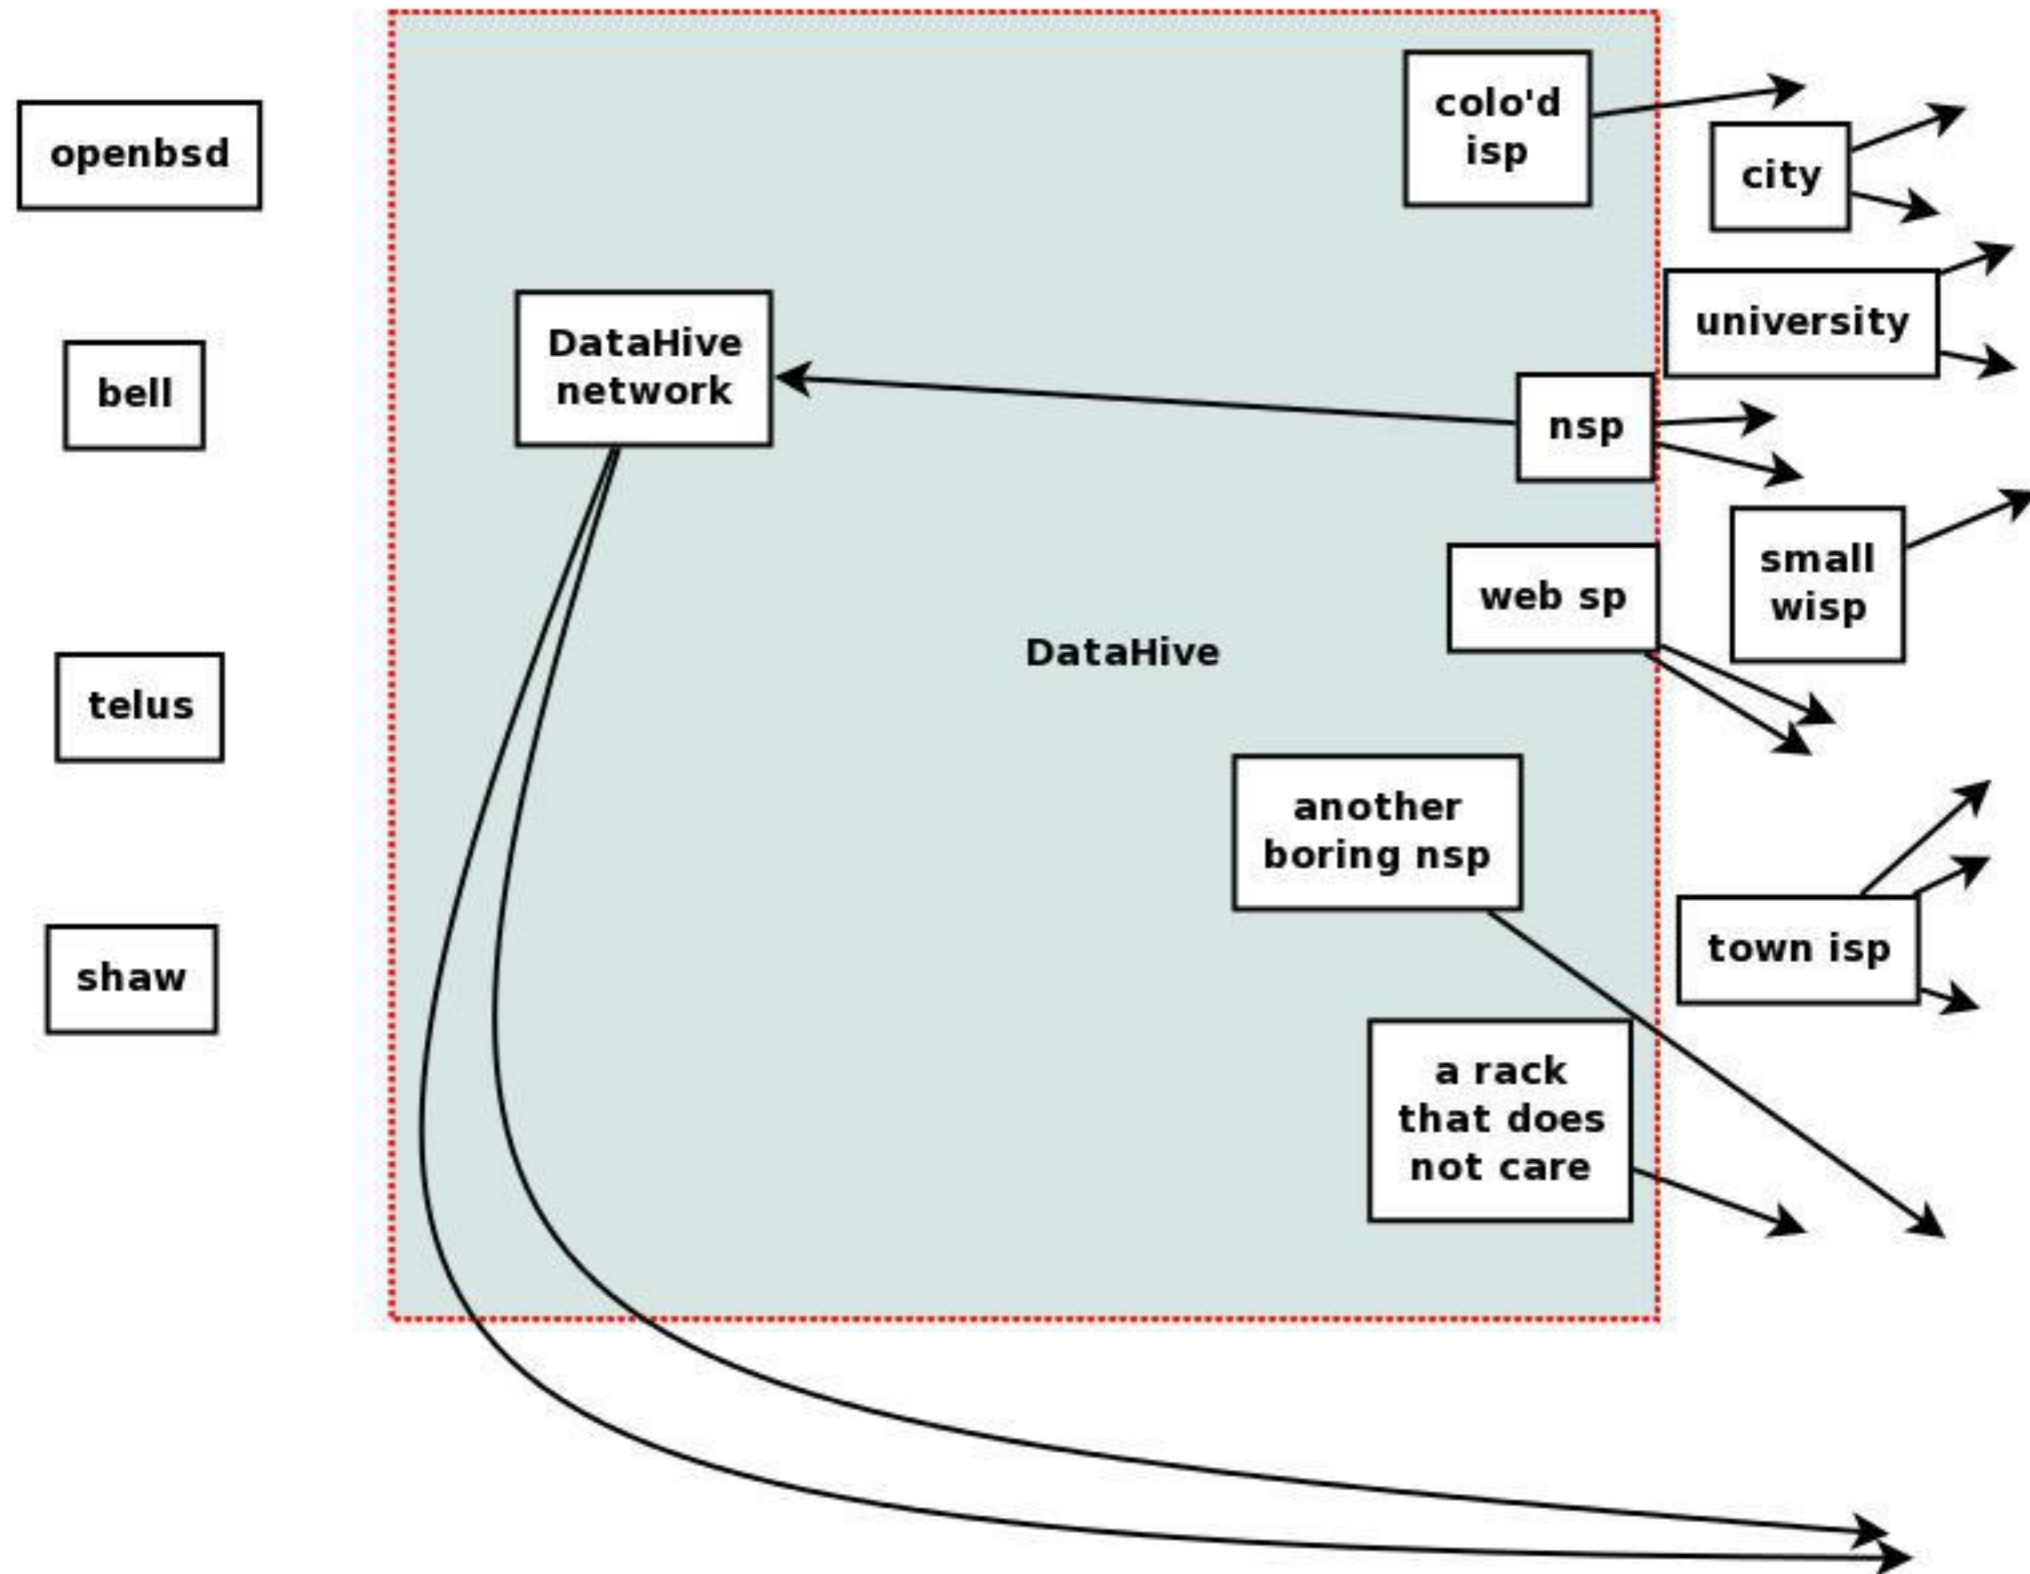
Global Server Centre

167 Lombard Ave, Winnipeg, MB R3B 0V3



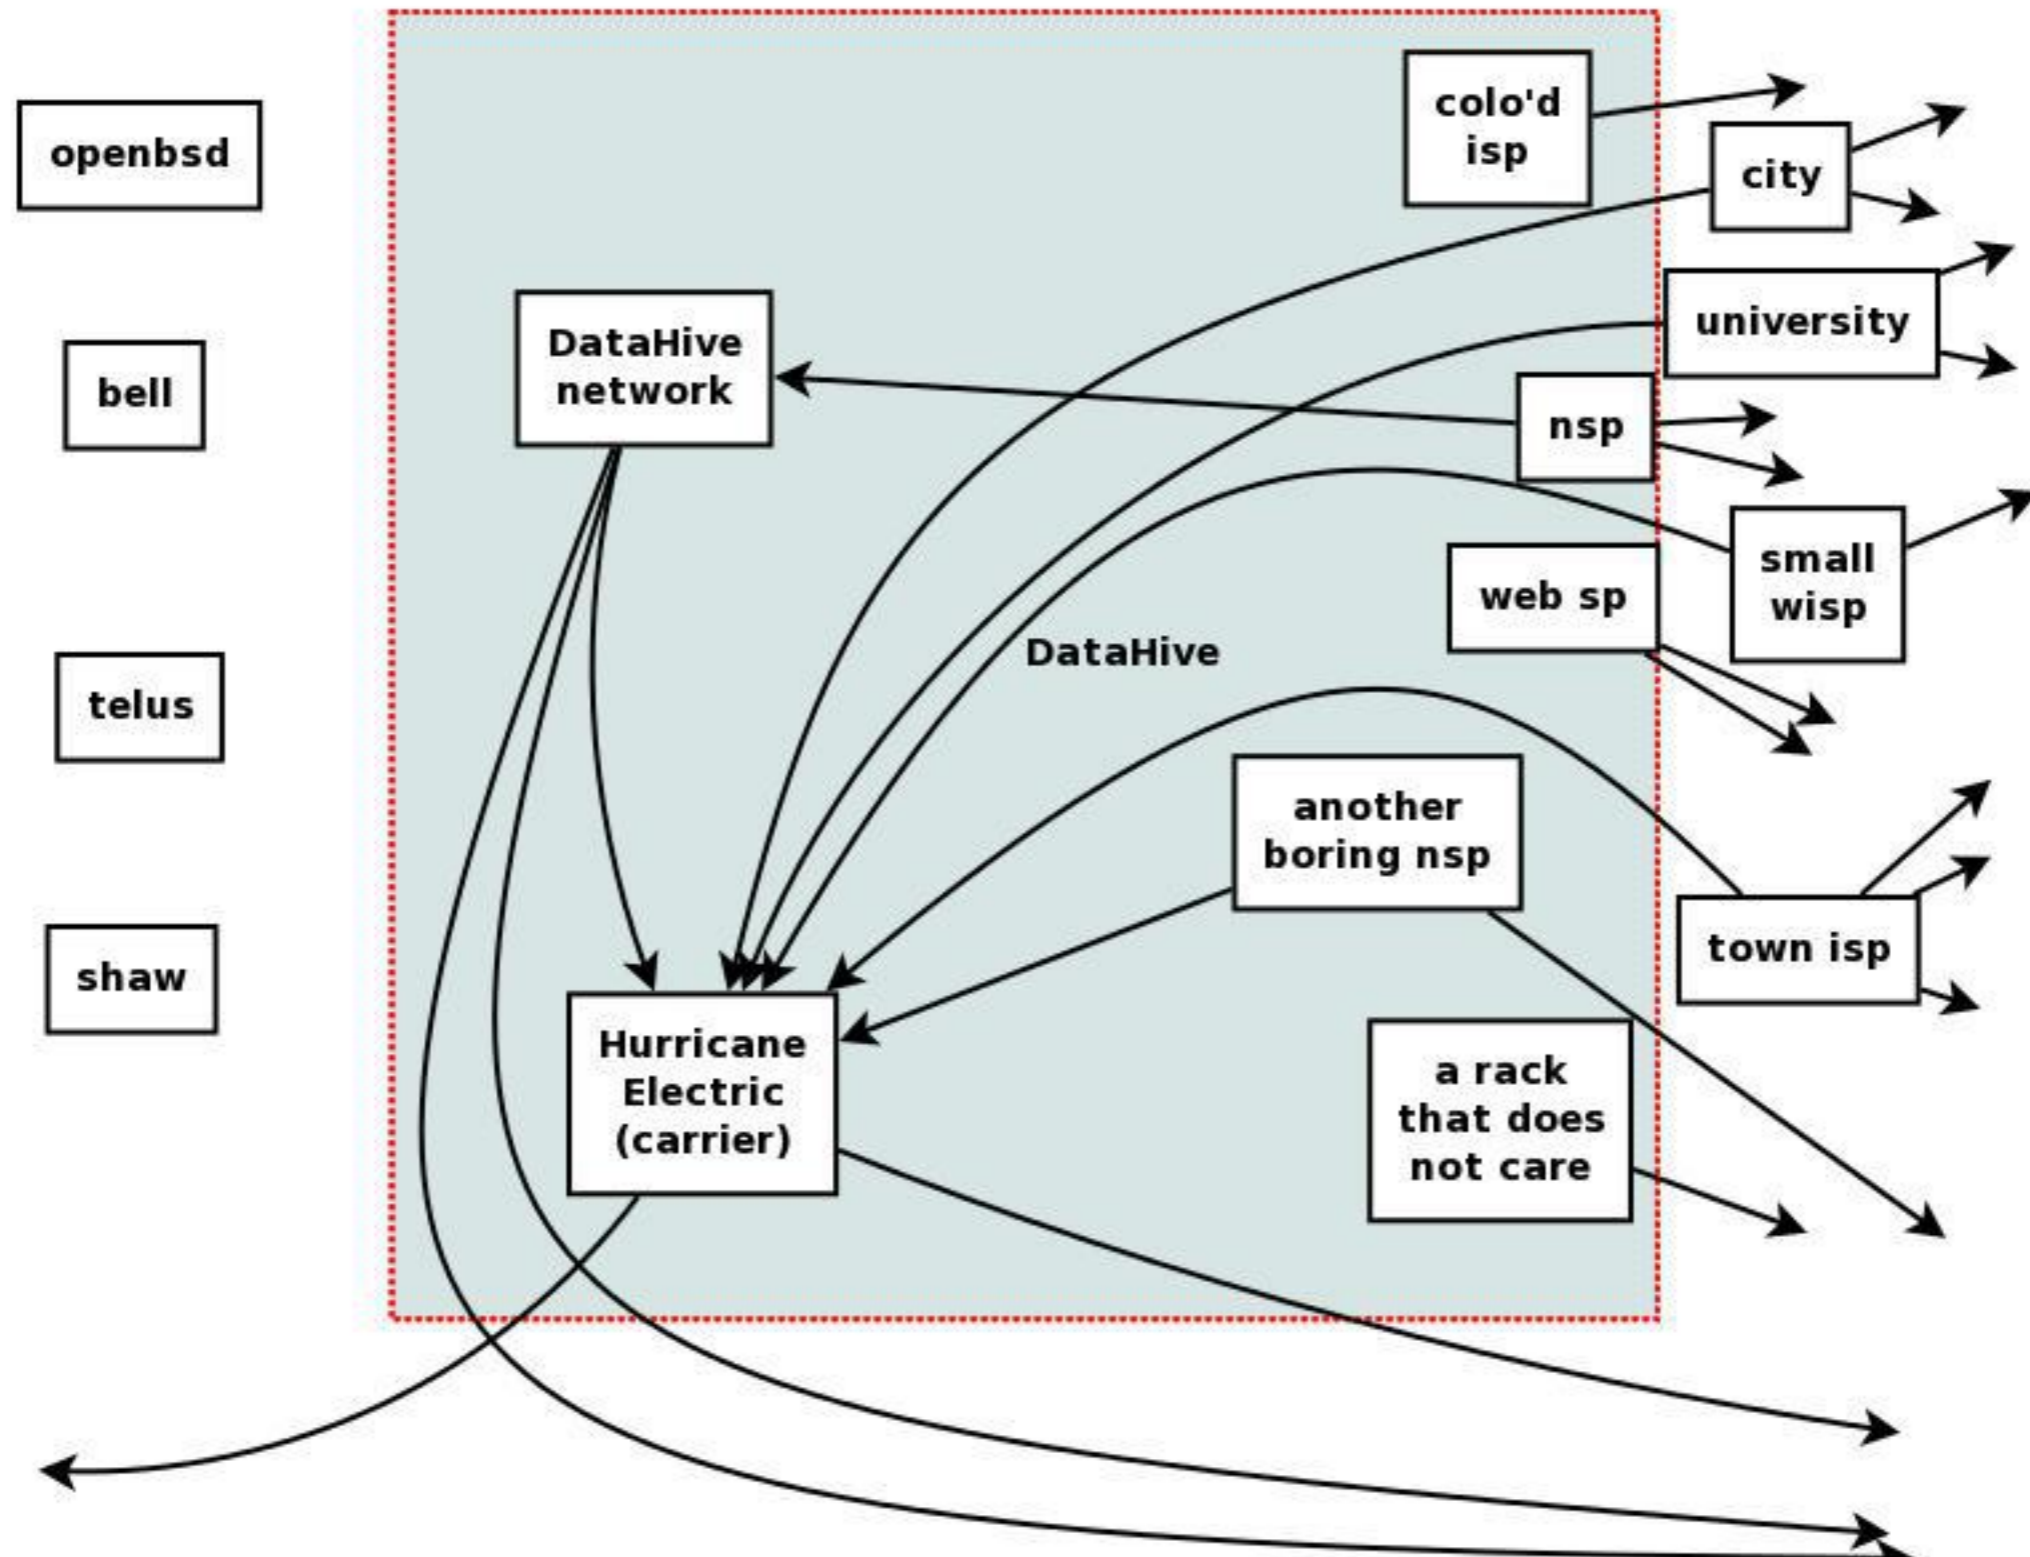
19-87 DRIVING THE LAST SPIKE NOV. 7-1885

DataHive Today



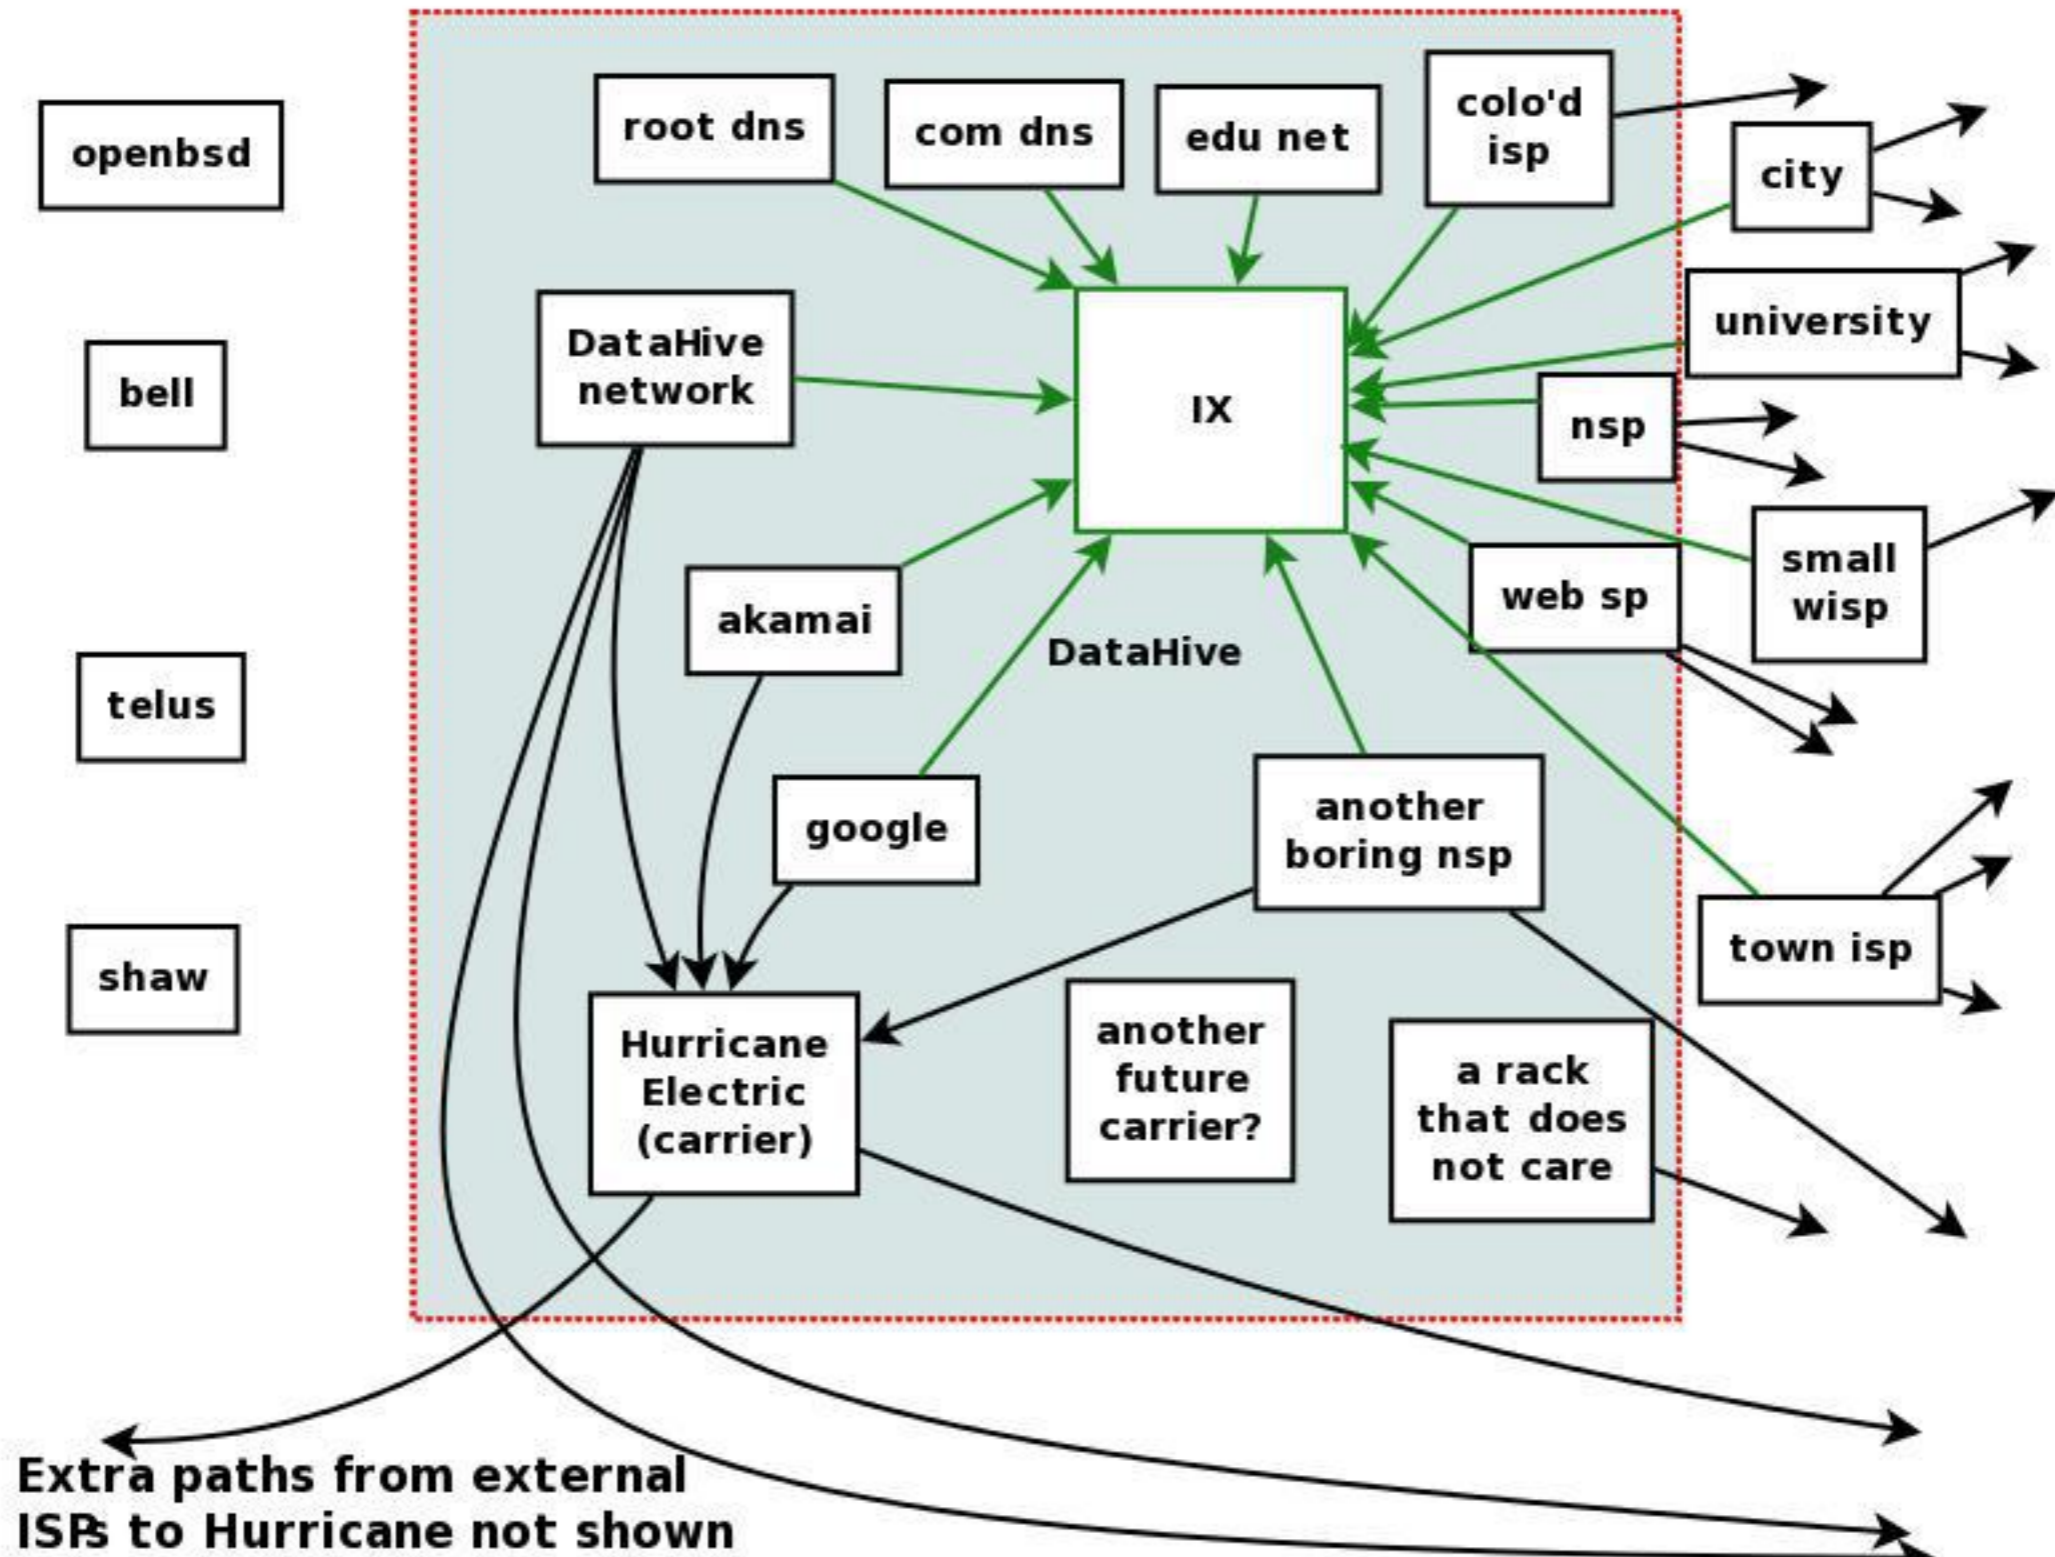
DataHive is a provider-dependent but neutral datacenter.

Hurricane Electric is coming to town...



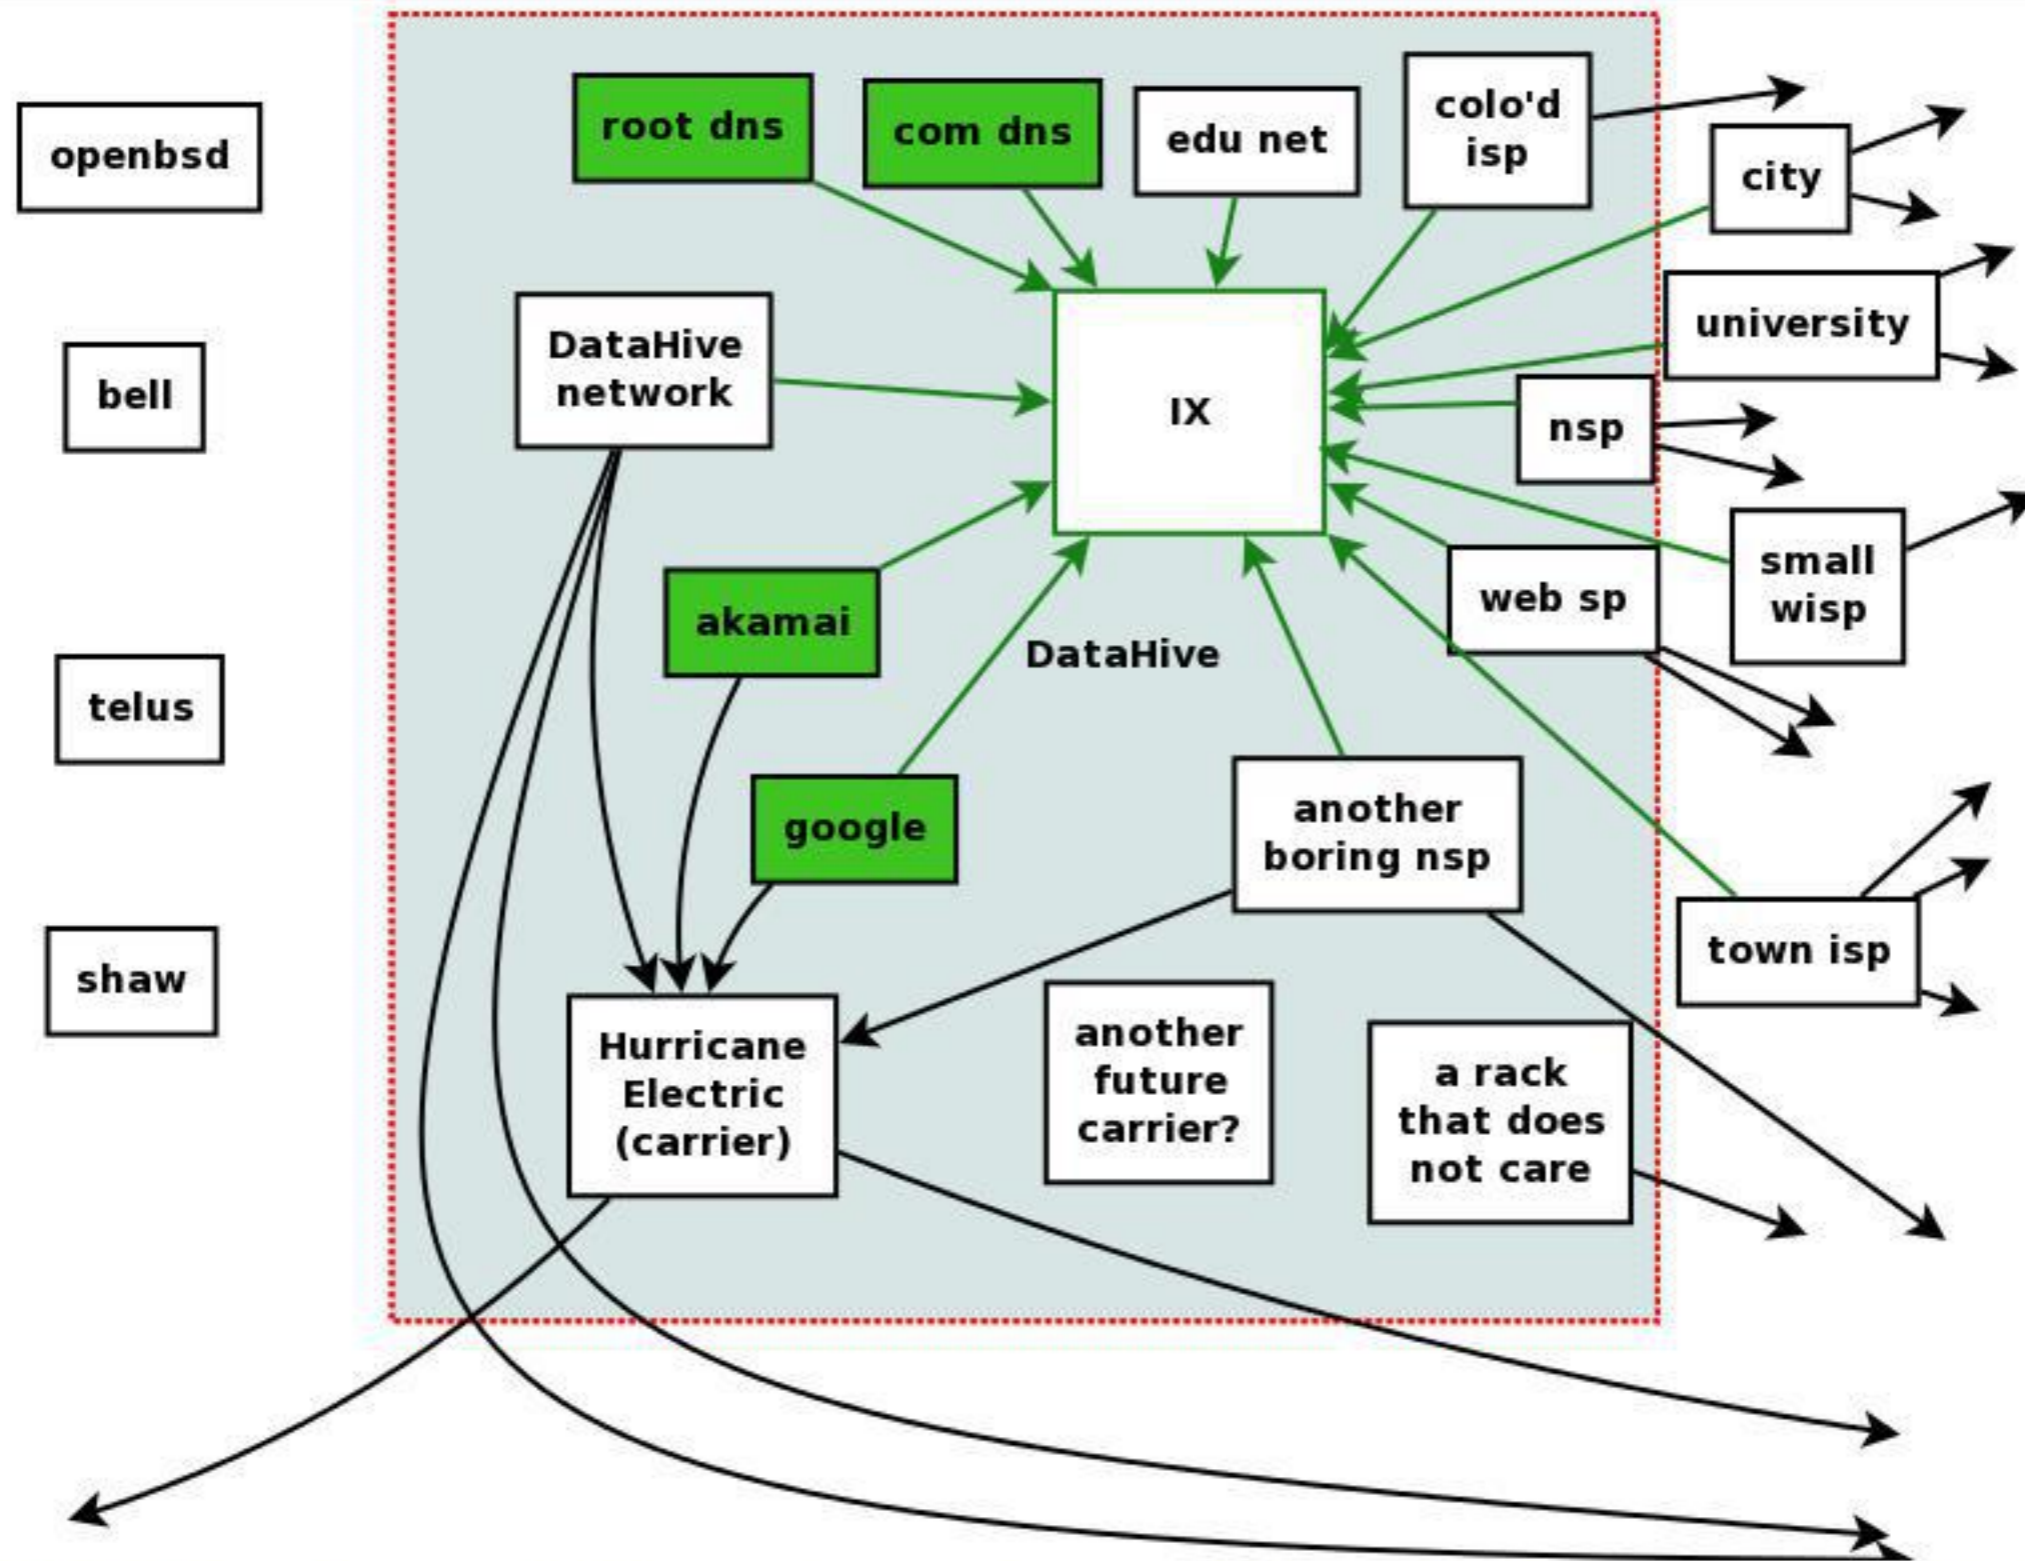
\$20-80/Mbps current prices ----> \$0.50/Mbps in DataHive
(1Gbps and 10Gbps ports, includes native IPv6)

Let's setup an Internet Exchange!



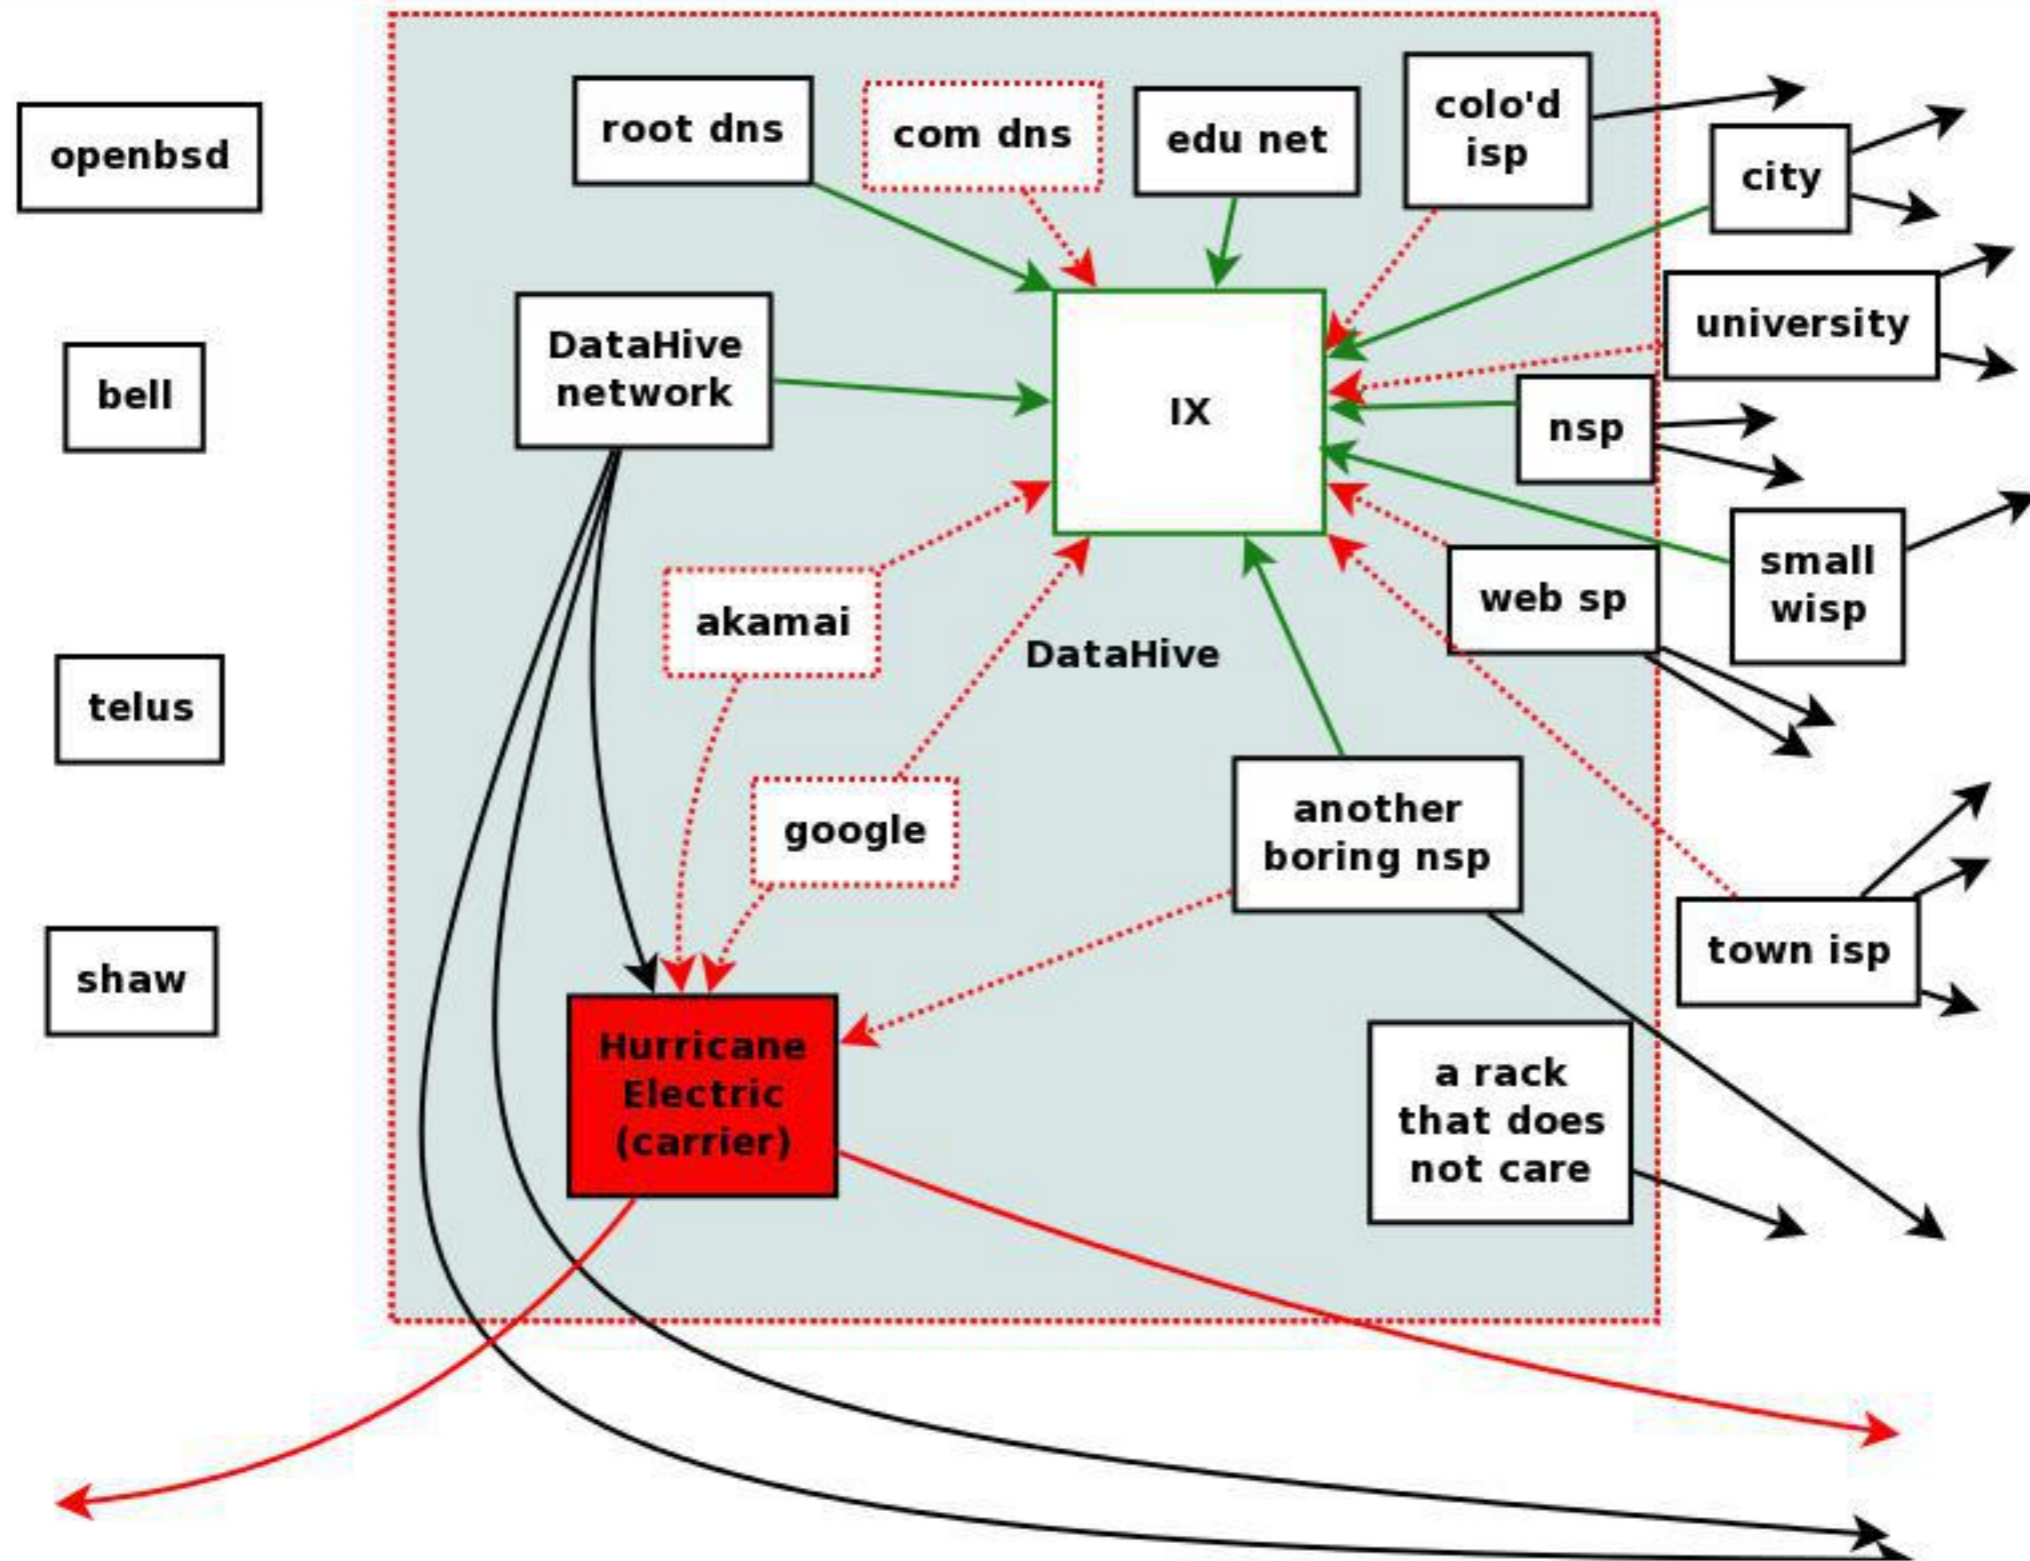
DataHive now a good location to attempt building an Internet Exchange

The secret is in the secret sauce: CDN arrival



DNS, Google and Akamai services will arrive if we do it right!

We could not afford the secret sauce before



Hurricane Electric (or any inexpensive carrier) makes it possible for the CDN caches to feed themselves

Content Delivery Networks

Google, Akamai, etc

Content providers **STRONGLY PREFER** to install at Internet Exchanges (therefore, nowhere in Western Canada thus far)

(they sometimes show up on private networks under agreement)

They typically bring in 1-5 racks of equipment and serve any network who is on the Internet Exchange

Why place a rack in Los Angeles when you want to serve eyeballs in Calgary?

They need to feed their caches (therefore, want cheap network at the location)

DNS and why it matters

DNS matters too!

Via my network paths, the closest DNS root servers are in Toronto or Seattle/Vancouver

Those are also authoritative for .com, .net, and in-addr.arpa

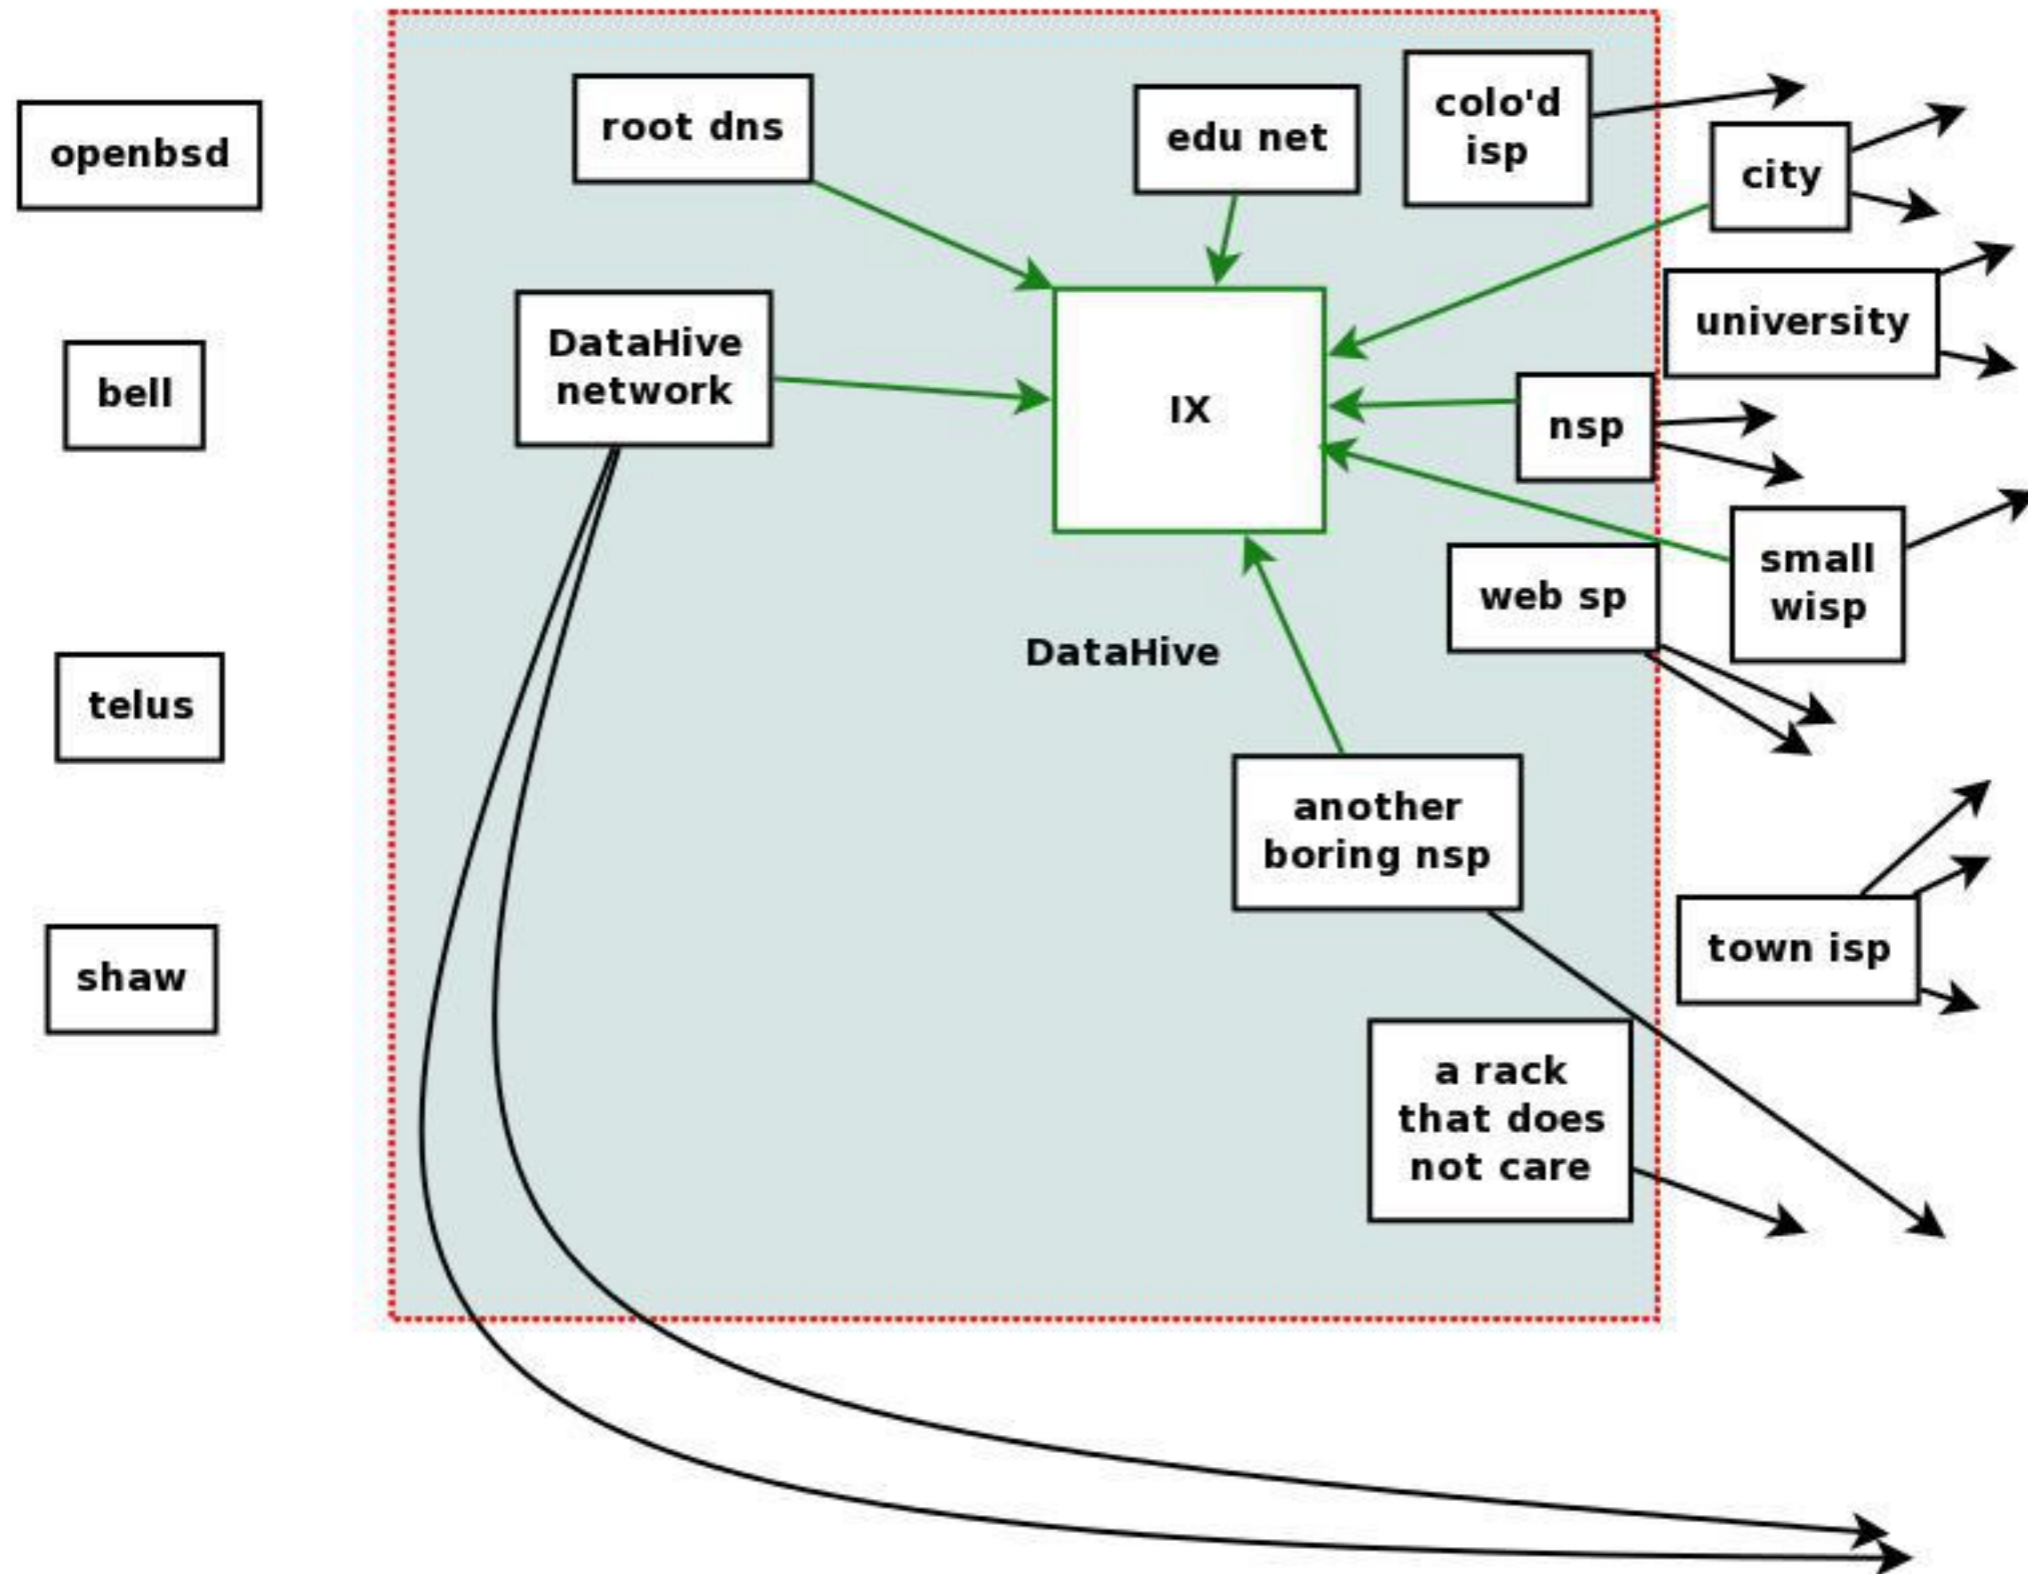
Some providers are doing a lot of DNS caching to circumvent the problem, but it will only get worse and worse over time

Caches tend to run into problems

Low-latency DNS matters, and it is a shame we don't have it here

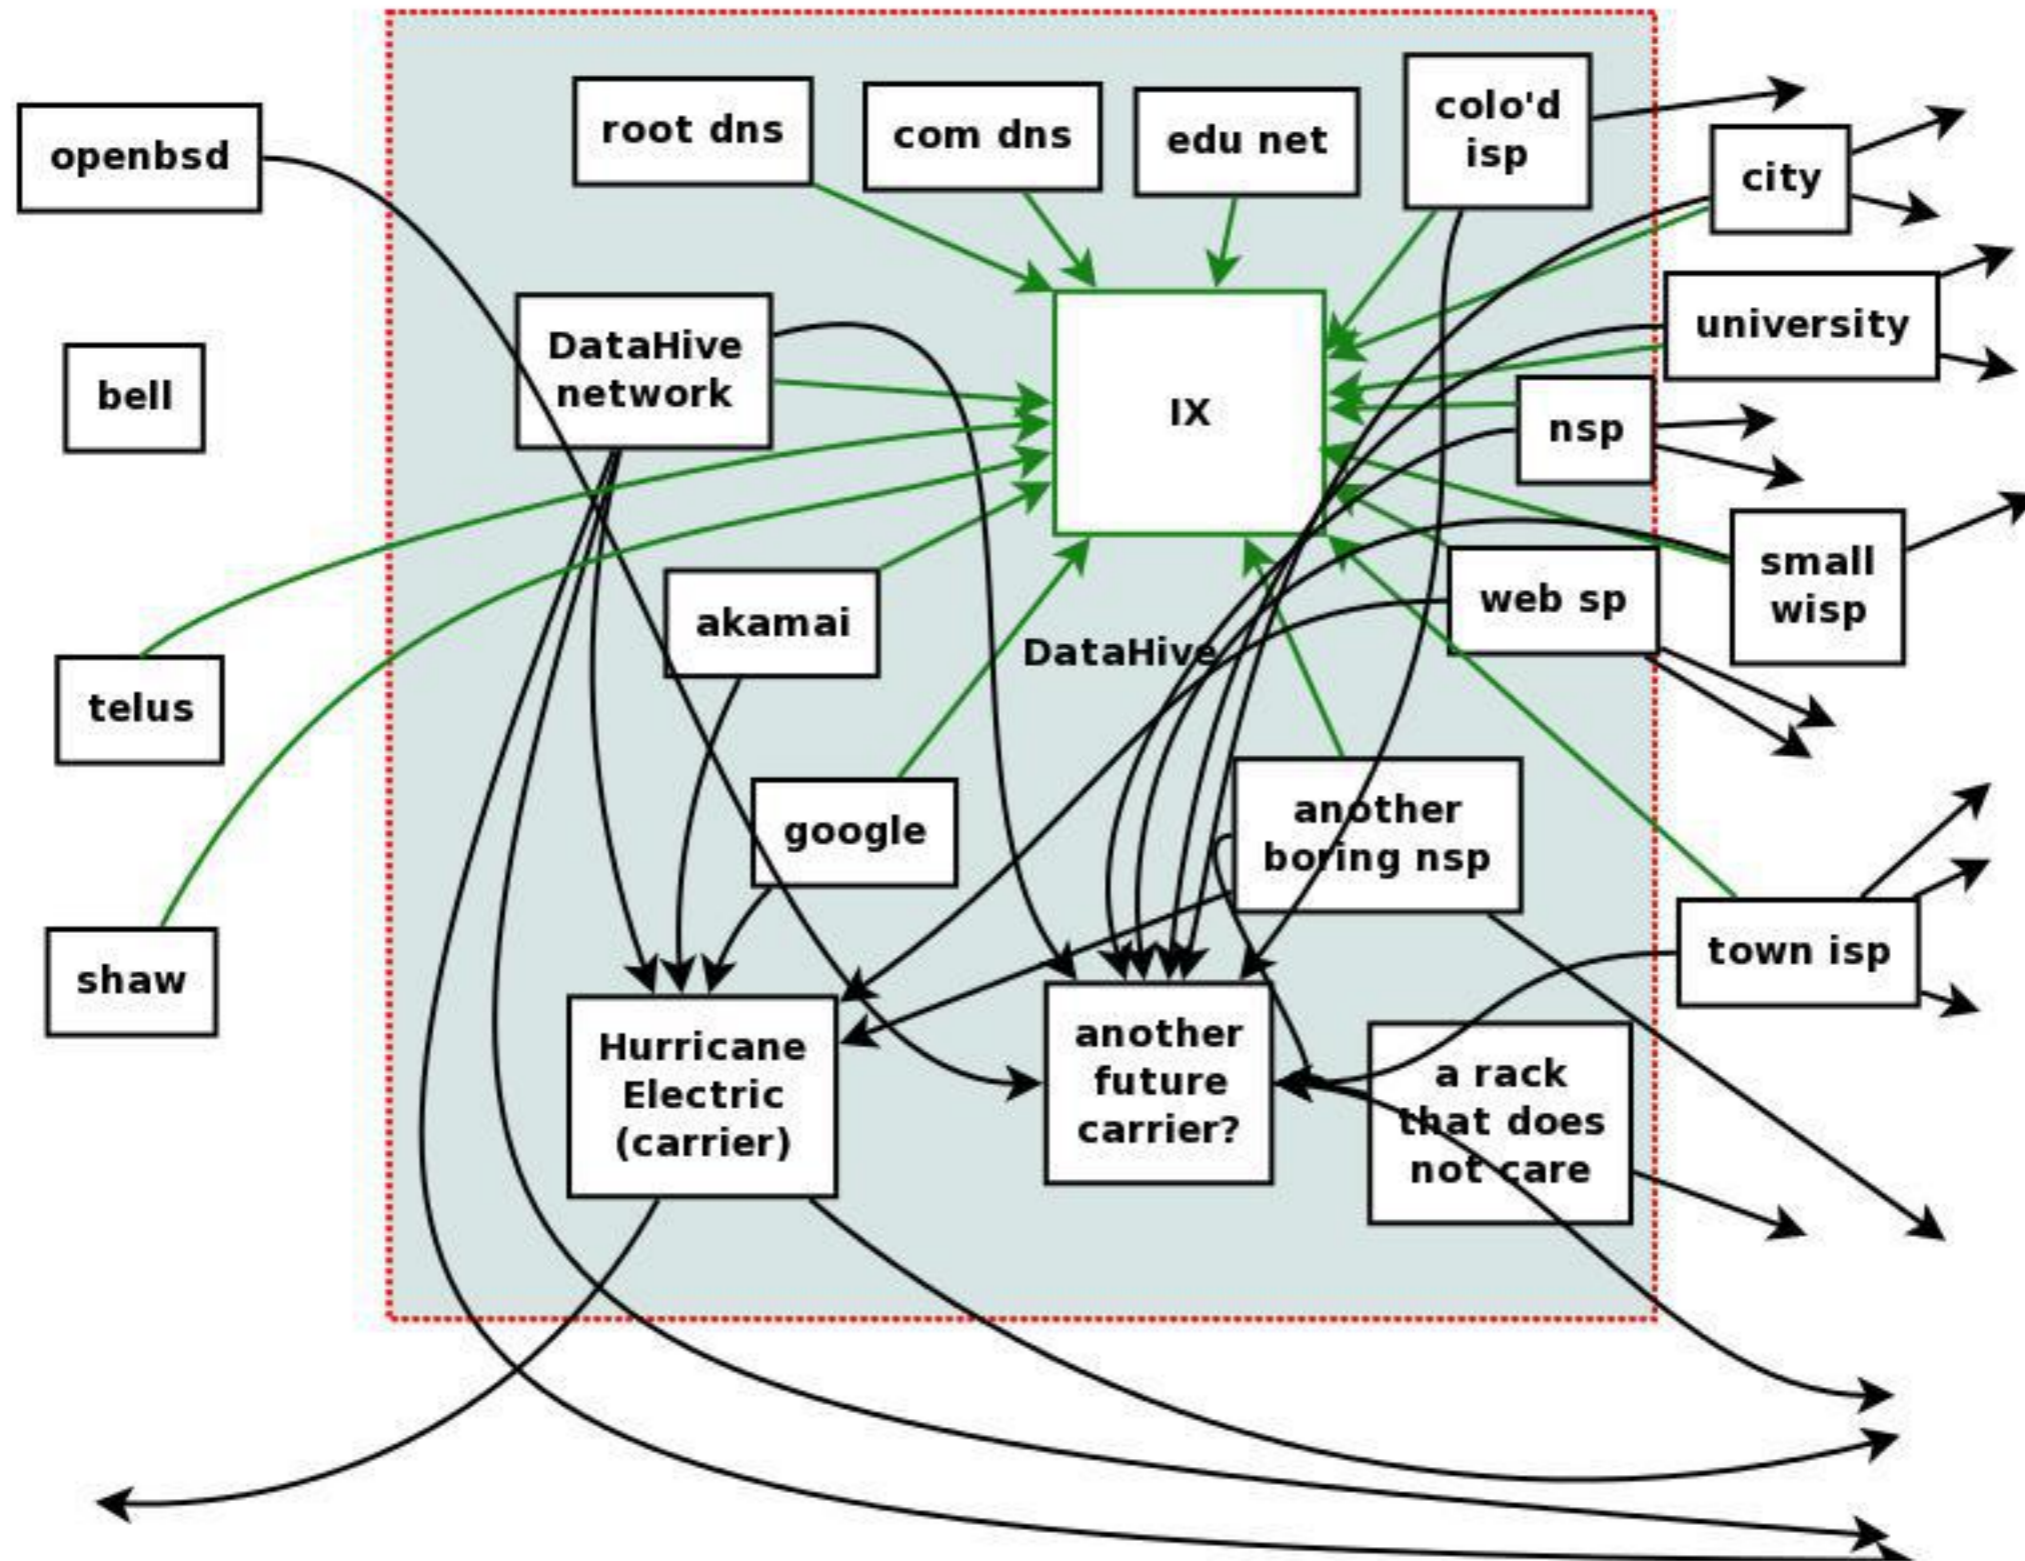
At an Internet Exchange, a pretty big setup is about 1+ racks

an Internet Exchange is pretty dry without the secret sauce



Looks really hard to do without a carrier...

But what if it goes really well?



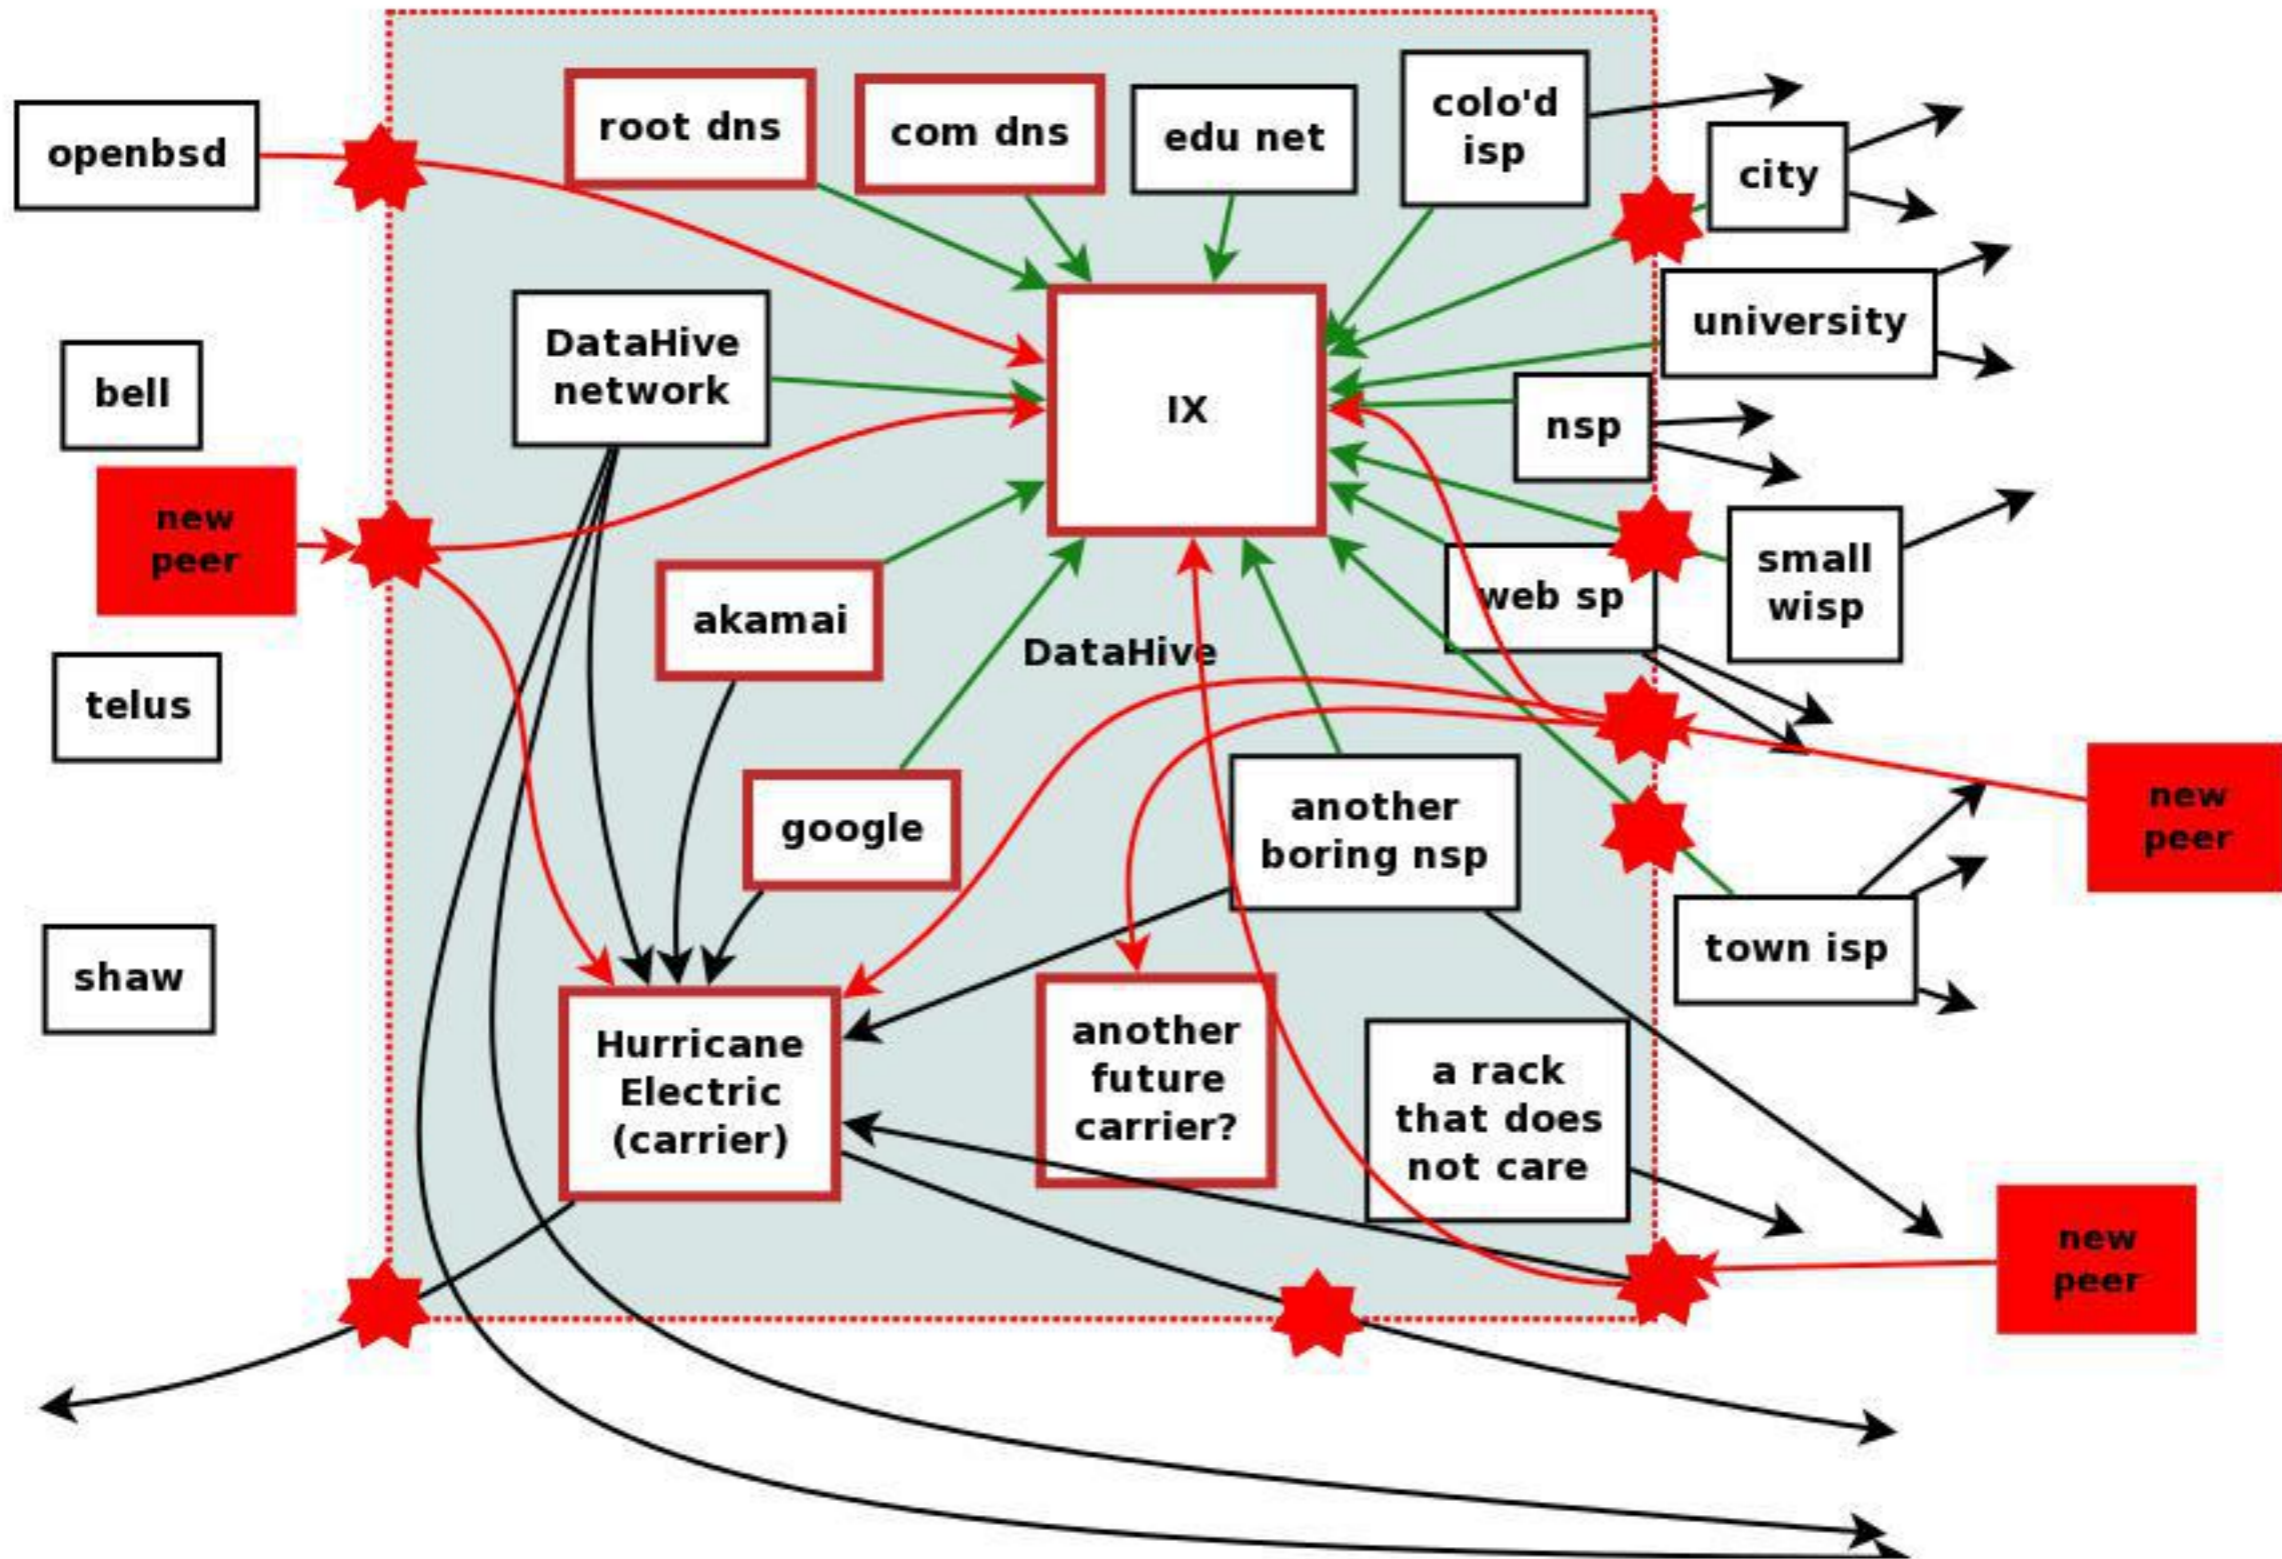
Not without precedent. When IX booms, it booms.
(especially when the alternatives are a rip-off)

All high-density wealthy populations have an exchange



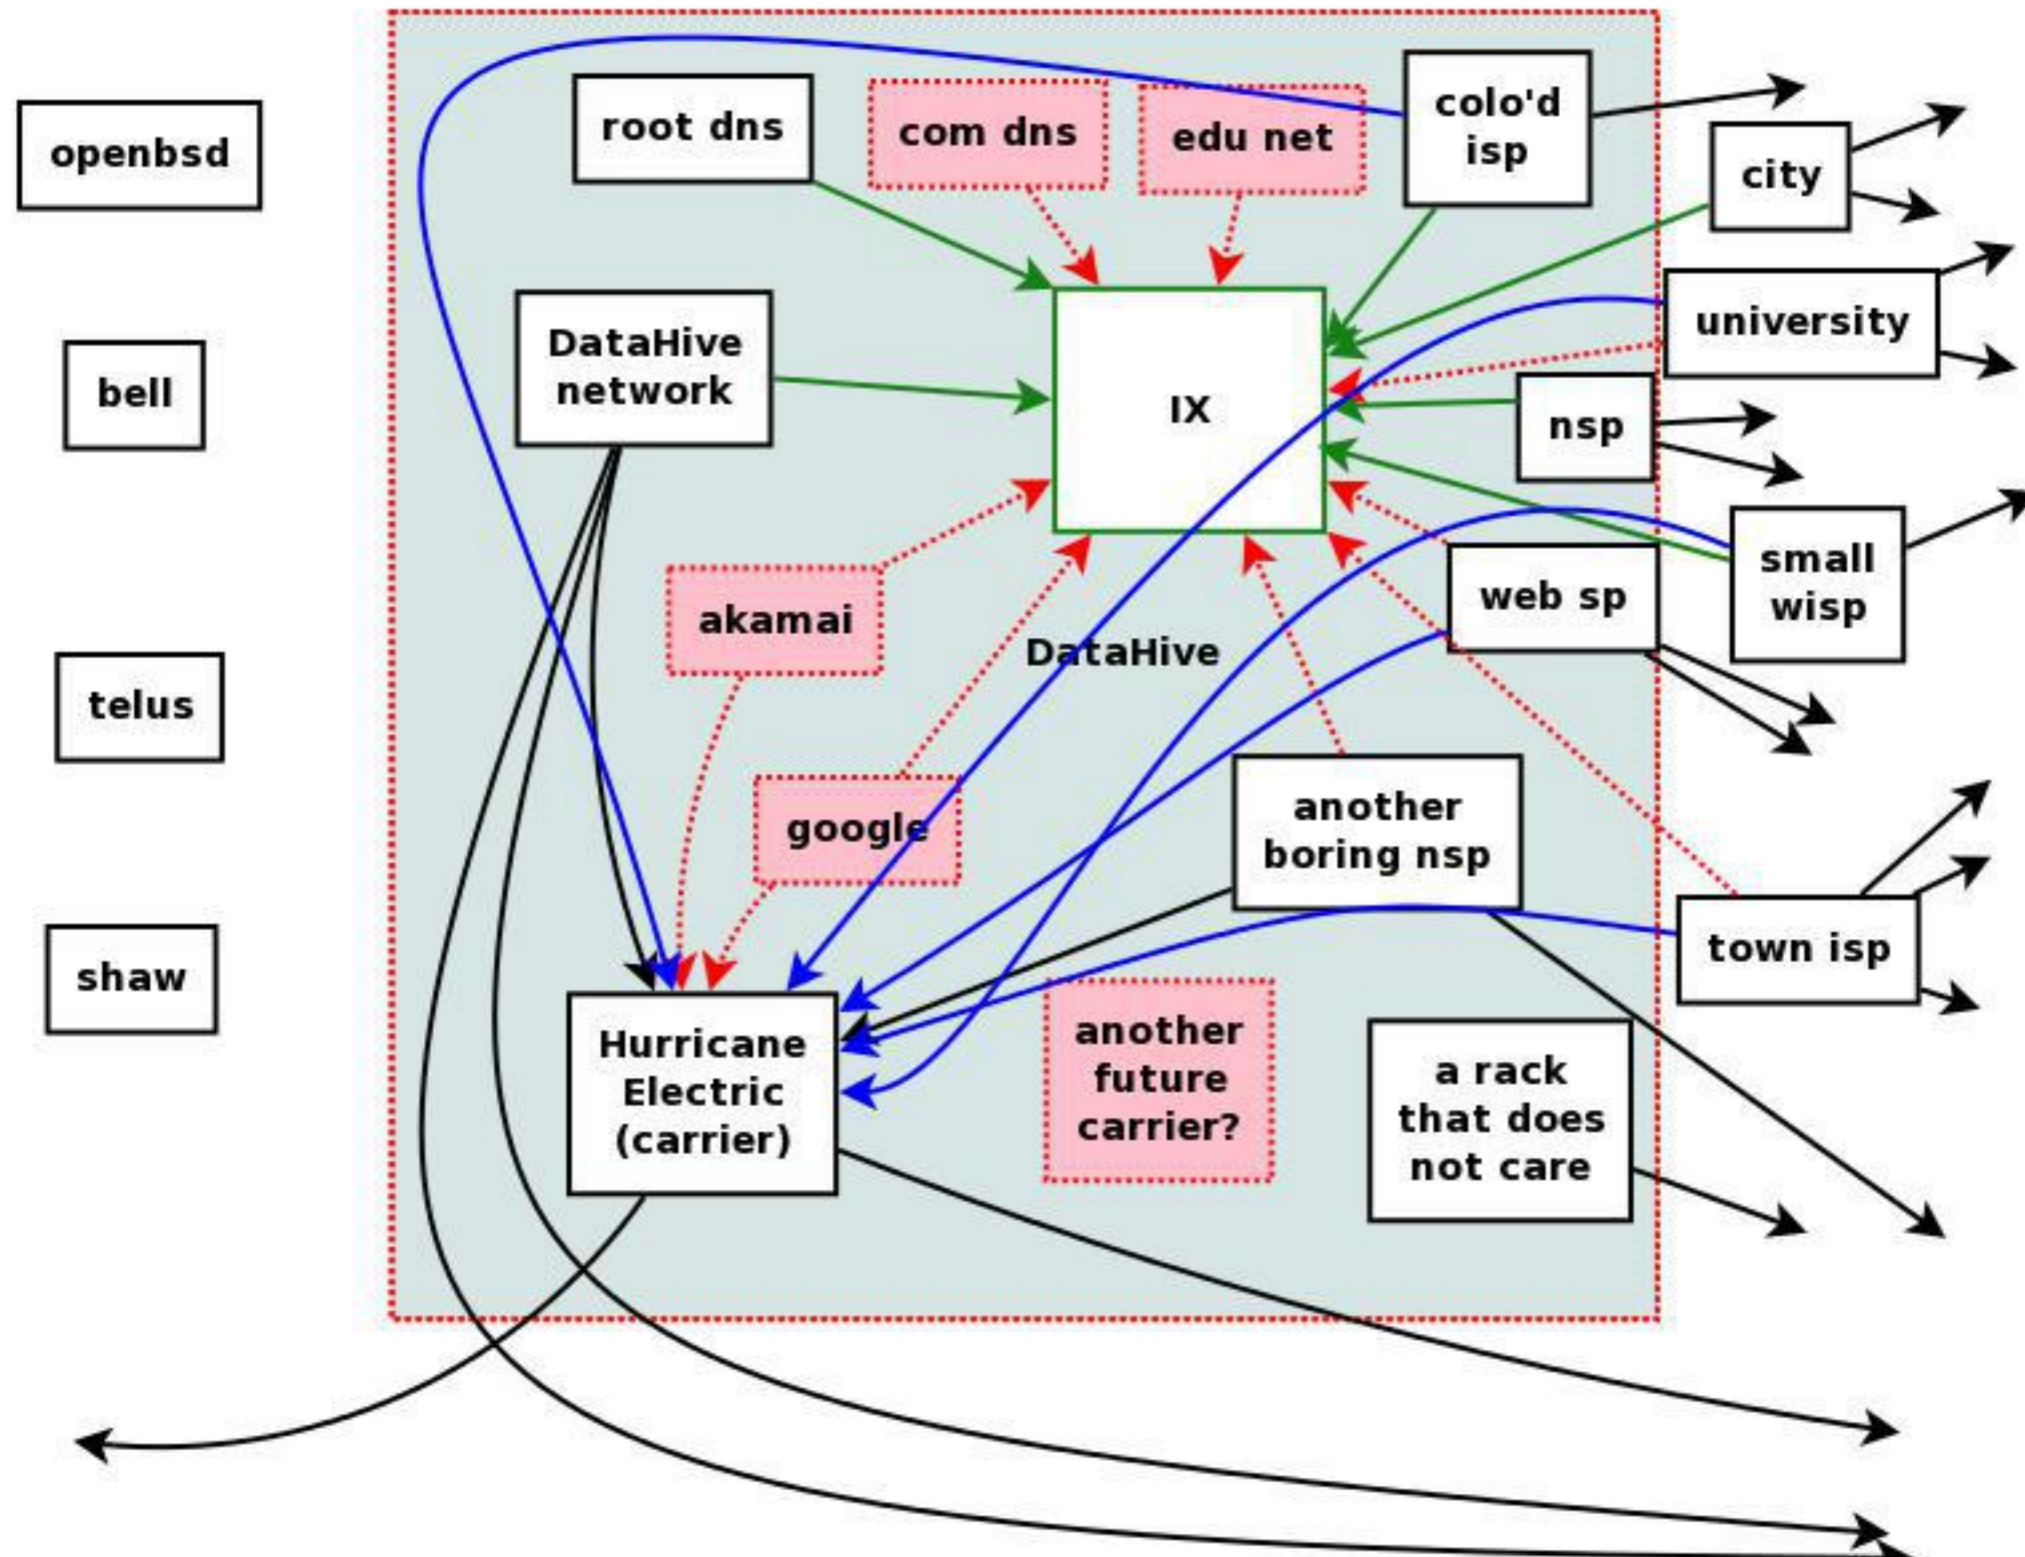
If they don't, their experience is diminished

The dominant local providers might try to stop it.



Actually, they are already trying to cause pain
Probably giving favor\$ to those who stay away from DataHive..

If enough pain, progress could be slow



Leading to DataHive just becoming a "cheap-net pickup point"

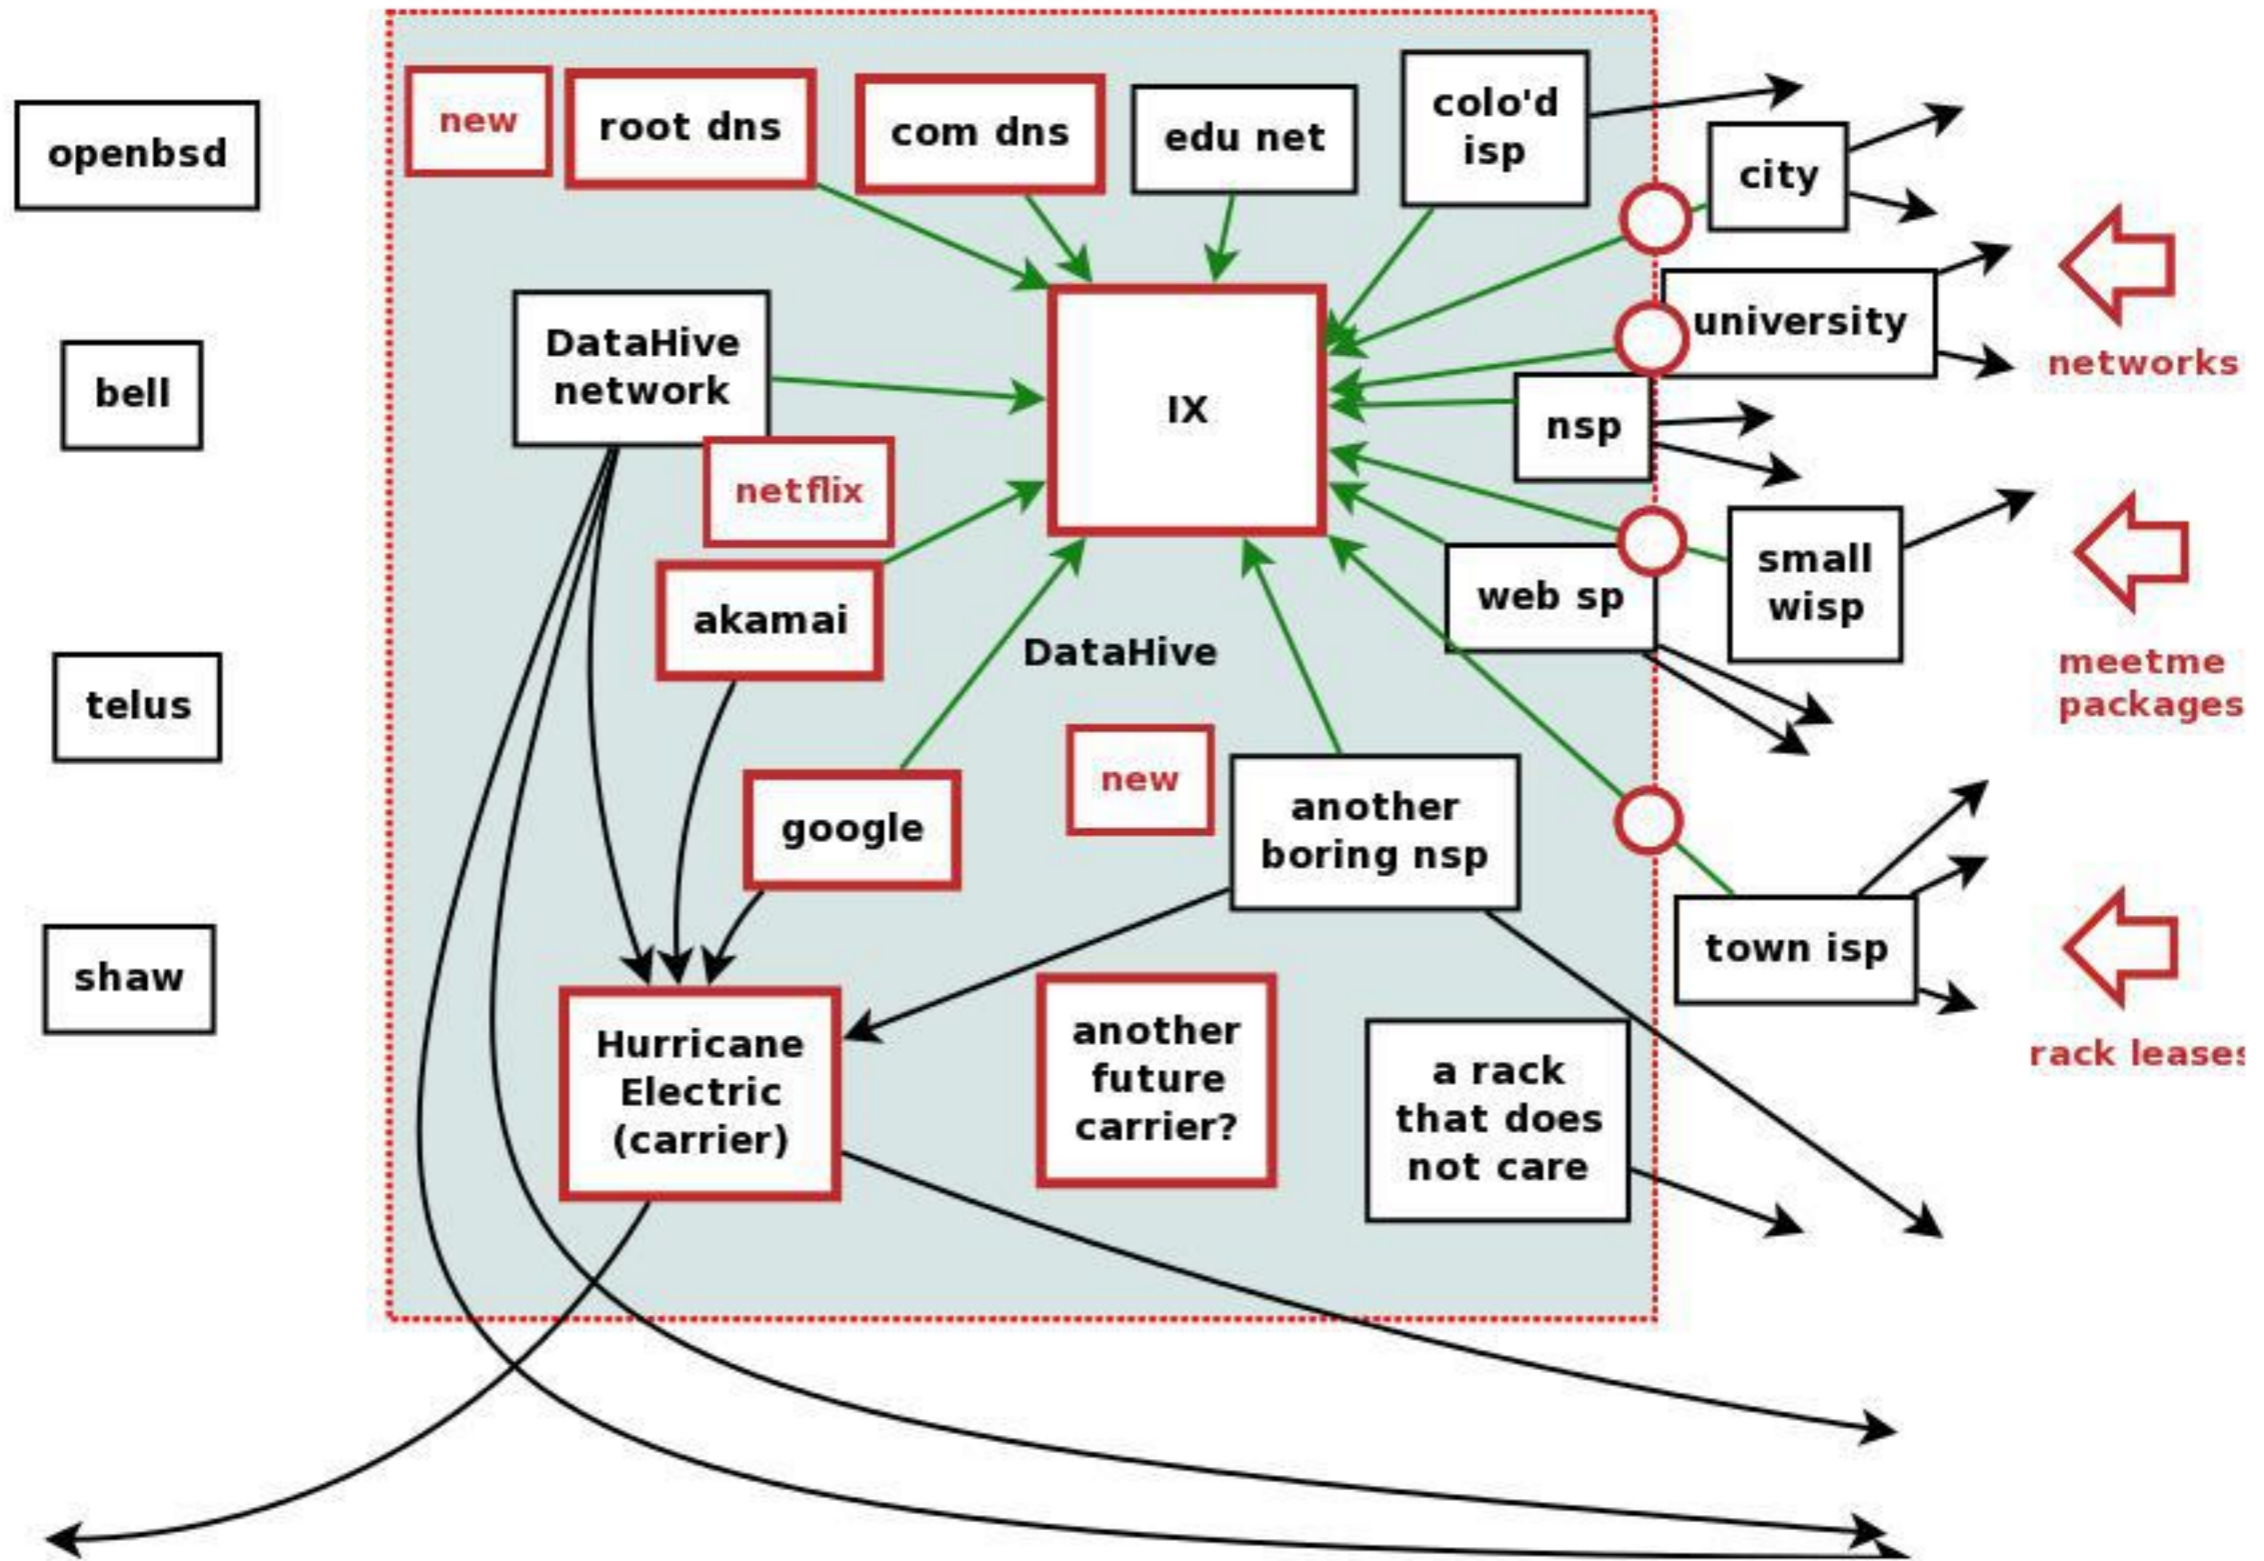
Is DataHive taking a risk here?

Naaaaaw...

I do not speak for DataHive!

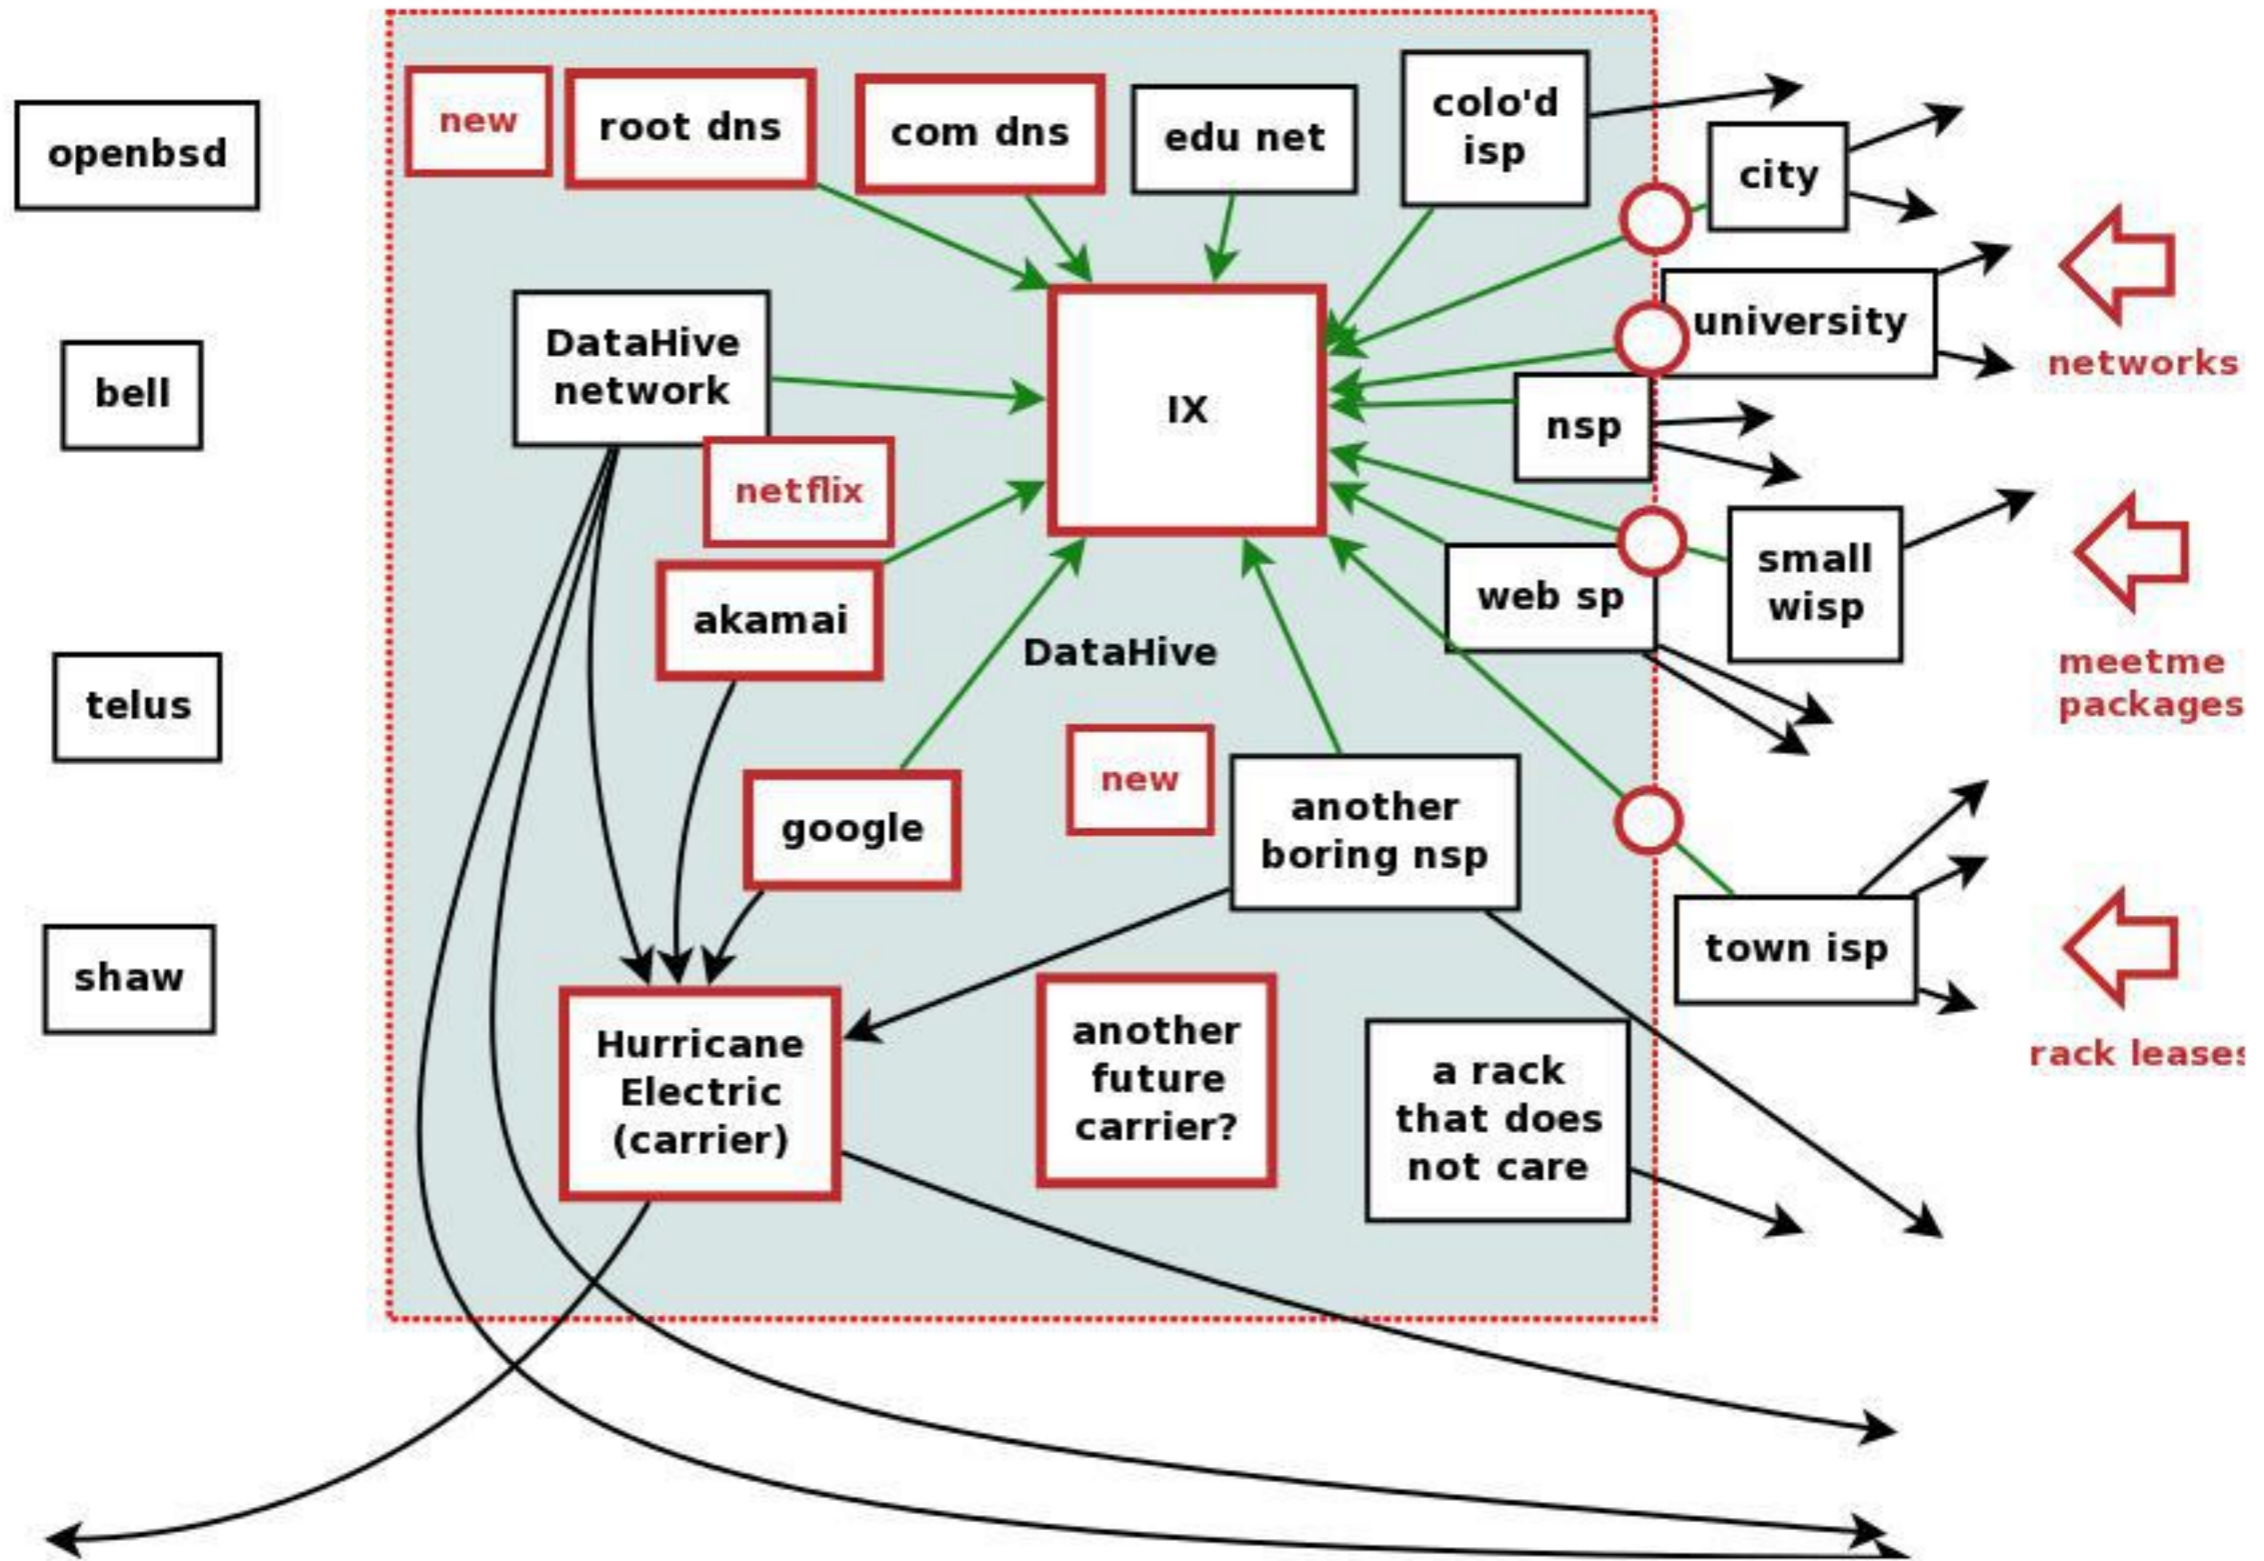
So just theories.

Internet Exchange locations do very good business



DataHive would prefer an IX not be elsewhere!

Internet Exchange locations do very good business



Some peers require proven neutrality and clean governance

But DataHive is not a carrier hotel

Typically exchanges happen where "cabling" pre-existed

When many top-level providers arrive, called a "carrier hotel"

In Calgary, we don't have that, so what is happening at DataHive is an attempt to create it from scratch

Sometimes 166 networks are in the same building: Toronto

.. or 225 networks in buildings in a metro area: Seattle


.. or 595 networks at 12 connected buildings in one country: AMS-IX

DataHive so far? About 20.

Have our providers let us down?

Our providers don't let us build resilient networks, then wire their networks through the same three buildings

One bad actor or event can upset the entire basket for us all!



The "151 Front Street, Toronto" basket

If a remote internet hub has a disaster, does it affect Albertans?

"Downtown explosion knocks out 911, internet service for thousands across Calgary", National Post, July 11/2012 (<http://tinyurl.com/mkonlh9>)

Has our government let us down?

Well, the government hasn't disappointed the stockholders of Shaw, Telus, or Bell

Provincial EDU providers may use CANARIE (CA research network) to carry school traffic to Google at US exchange points

Sprint Canada and UUnet Canada got "acquired" after building more choice for consumers

No other carrier has come near here in a long long time

No place in the world has fiber internet prices as high

All non-Prairie governments have their own AS

Blissful ignorance while our economy de-diversifies further

Internet Exchange Models

Seattle - totally volunteer, small fees

ToRIX - volunteer, slightly higher fees

AMS-IX, DECIX - becoming a big big business, but still cheap

Commercial models: Equinix, Hong Kong?

Not-for-profit models

Big exchanges get greedy --> new small ones pop up

ADVICE ON INTERNET SAYS TO START SMALL

RULES FOR PEERS

Exchange is run by volunteers

Don't want to get burned out by stupidity or politics

Therefore, peers must agree to play nice

If they play nice, they get the benefits

If they don't, they get the boot

This model is working everywhere else

(don't take my word for it -- go do research)

Harder to startup than to operate?

Internet Exchanges are a fascinating mix of good engineering, business, and fighting against monopoly/oligopoly/cartels

Keep local traffic local --> all sectors benefit

Don't take my word for any of this

Do your own research if you are interested

Calgary is probably not a unique snowflake

Challenges

We've faced some big challenges already

There will be more ahead

Some internet exchanges have had trouble starting up...

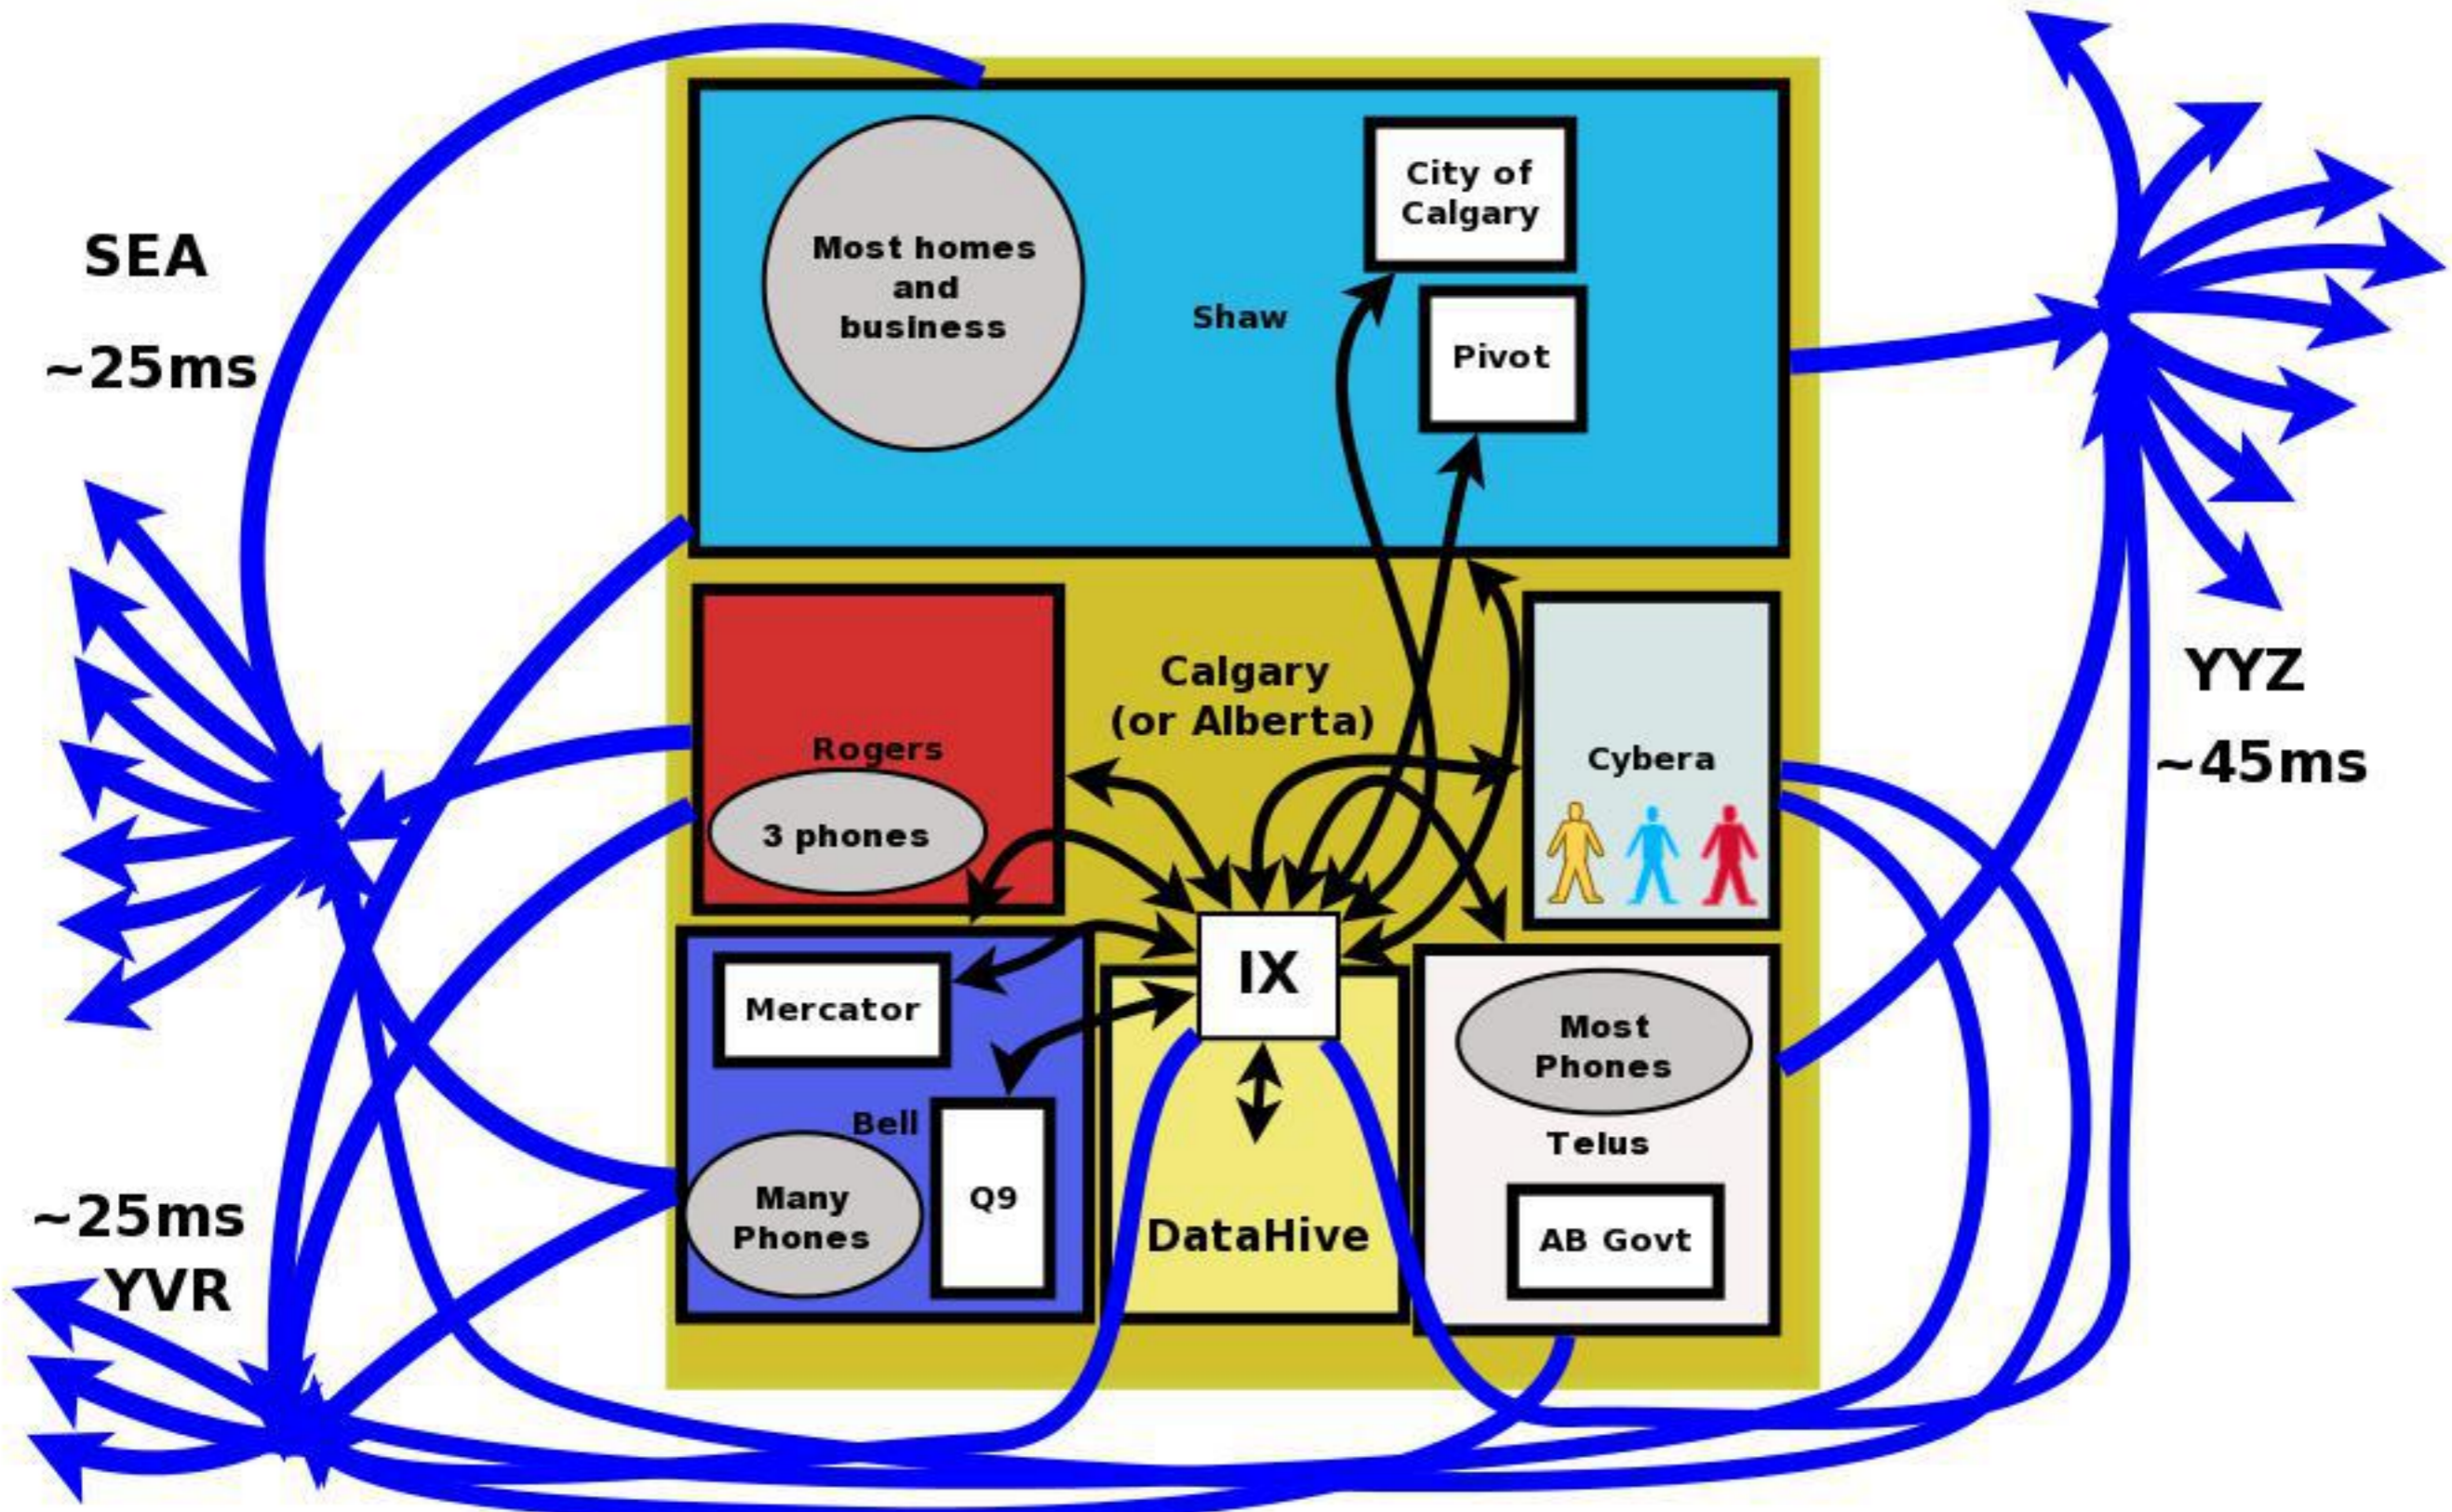
We need concerned people involved to make this happen

Must find a way to keep all the relationships clean

Hopefully the regional network operators will get the message that CDNs don't come unless they unify behind one vision

This is a great opportunity

Calgary: a model for the future



Build cheaper, better Internet for local people

Summary

In general, internet exchanges provide

- the least-pricy network path
- the highest-performance path
- additional redundancy
- very high bandwidth, to local partners
- downside: at a certain pin-point on the map (...or pin-points)
- downside: local-loop connection to there

Benefits:

- Decrease latency
- Improve access to DNS services
- Improve access to Content
- Improve carrier choice
- and of course, drive prices down