Failure,
挫折 (setback)

Blair Harrison
Kiwicon 2013

About

- (mostly) Linux Sysadmin
- Done lots of work for ISPs and Hosting providers
- Seems to break a lot of stuff
- On Twitter - @trogs

Summary

- A bit about some faily stuff (Not entirely my fault)
- Some stuff I failed at (mostly vendors fault)
- Some pictures (all mine)
- Some Advice (That I learned the hard way)

What is failure?

- When stuff happens that we didn't expect or anticipate
- Not always bad or catastrophic outcomes
- Can we learn from it?
- You guys learn a lot when you find bugs, vendors learn a lot too
- Hopefully everything gets better

System Failure

- Disk related failure
- /tmp filled up. Now you can't install anything.
- /var is full. No more logs for you today.
- / is full. Everything is fucked. Can't login.
- How long have you been running on 1 power supply?
- How do we prevent this?

Big file systems

- If you have more than about 4TB of disk space, you're going to run into problems with hrStorageTable (1.3.6.1.2.1.25.2.3) or dskTable (1.3.6.1.4.1.2021.9)
- What does 200TB look like if the counter wraps?
- Use extended SNMP attributes to monitor % of space, and maybe how many GB you have used
- in snmpd.conf
- extend diskpercent /scripts/diskpercent.sh
- diskpercent.sh is just
- df -h /mnt/diskwhatever | awk {′print $5′} | tail -n 1 | sed s/\%//g
- this ends up being 1.3.6.1.4.1.8072.1.3.2.3.1.1.″diskpercent″′
- snmpwalk -v2c -c yourcommunity hostname 1.3.6.1.4.1.8072.1.3.2.3.1.1.\”diskpercent\″

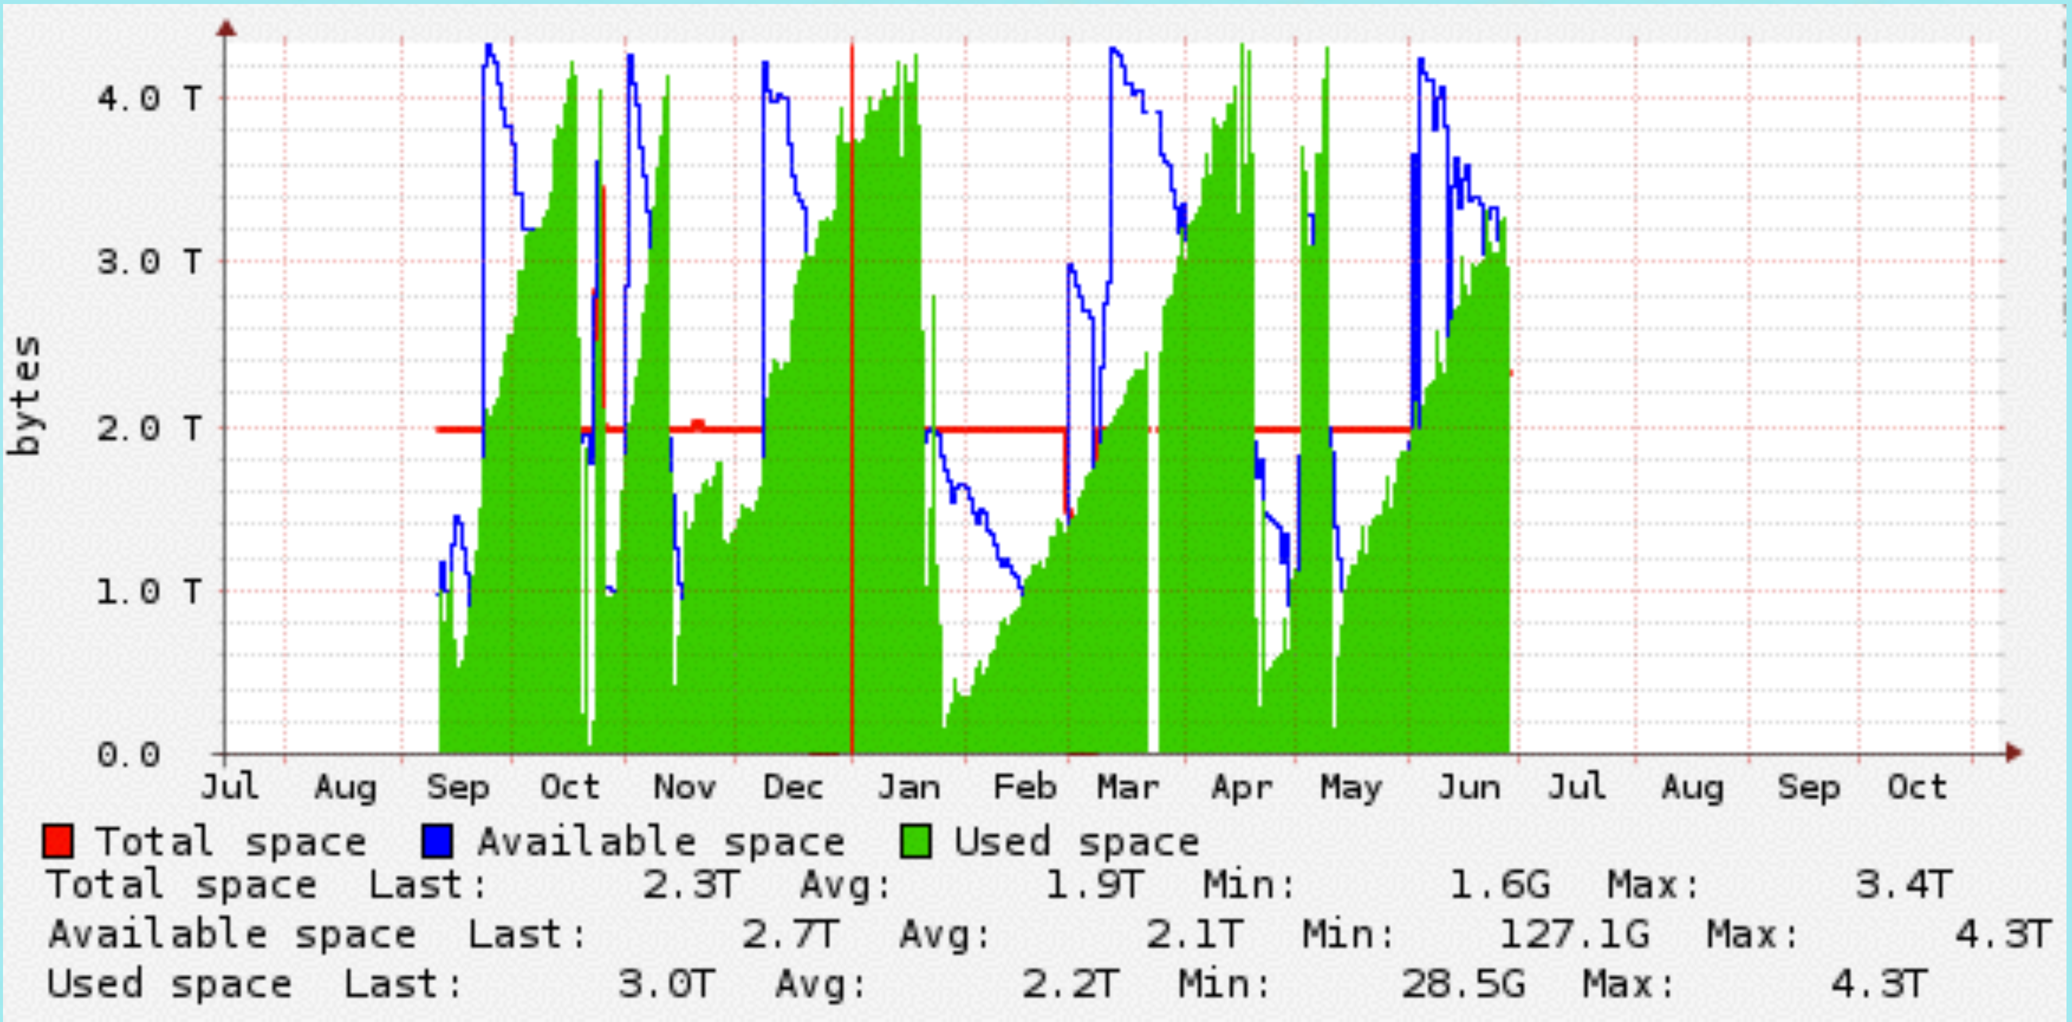

🟥 Total space 🟦 Available space 🟩 Used space

```
Total space  Last:      2.3T  Avg:      1.9T  Min:      1.6G  Max:      3.4T
Available space  Last:      2.7T  Avg:      2.1T  Min:    127.1G  Max:      4.3T
Used space  Last:      3.0T  Avg:      2.2T  Min:     28.5G  Max:      4.3T
```

Monitor your stuff

- Hopefully you have Nagios etc
- Get a usb cellular modem - use that to send text alerts. If your internet goes down, that swanky Twilio API aint gonna be worth shit
- Put ILOs on all your servers. Or use vmware, xen, etc. Easier to reboot from bed rather than hiking to the DC in the middle of the night
- This years lesson is - don't put the ILOs directly on the internet, put them behind VPN on separate mgmt network. Don't use the same switch as for data if possible

Who monitors the monitor?

- Monitor your monitoring box
- Put another box somewhere else, or borrow someone else’s monitoring system
- If your Nagios box dies in the night, you aren’t going to want your customer to be the one to tell you your webserver is busted in the morning

Indescriminate deletion

- rm -rf /<tab> no complete, accidental enter
- fuck
- ILO and restore time!

Backups

- Back up your servers. rsnapshot is Good Enough
- Test your backups!
- Don’t put your backup server in the same place as the servers you’re backing up
- Disk is pretty cheap. 2TB disks are flaky as hell, though.
- WD has a better RMA procedure than Seagate

Bitflipping

- bitflipping / bitsquatting (see last years talk)
- exploit bitflip memory failure in dns lookups
- Last year I did some trademe variants
- Failed at instrumenting the DNS lookups and apache setup entirely
- Did it somewhat better this year, with some other domains, setup a new vm, dedicated DNS server, all logging on etc
- NO FUCKING HITS on the domains I picked
- (except googlebot)
- FAIL

bitflipping fail

- Kind of hard to figure out what is actually going to get you a lot of hits with regard to bitflipping anyway
- best thing seems to be just register lots of domains
- interesting bitflip is . -> n
- eg www.amazon.com -> wwwnamazon.com

jquery

- who uses jquery?
- jque2y.com
- ajaxngoogleapis.com
- already had some hits
- all your jquery are belong to me

So you like SNMP?

- Simple Network Management Protocol
- SNMP is a standard, right?
- Should be pretty easy to use standard tools to look at all devices with no issues at all, in 2013??
- ha ha ha

BGP and SNMP

- How many BGP peers are up? how many are down? who is affected by this outage RIGHT NOW?
- Don’t want to mess around sshing to routers running weird commands
- write a script - bgpstatus.sh (see me later if you want it)
- Poll a bunch of OIDs to get BGP status for all your peers
- Nicely formatted display
- Totally useless without extra data because Cisco doesn’t include peer descriptions in SNMP - FAIL

RANCID to the rescue

- Use RANCID to back up your routers
- Can mine this data with bgpstatus.sh
- Now we have a useful BGP status script
- Only good if you don’t have overlapping IPs in VRFs (BGP MIB has no support for this, not really Cisco’s fault, but annoying)
- Also no IPv6 Support < ios 15.2 (cbgpPeer2Table) FAIL

bgpstatus output:

```
Total Peers Configured for xxxxxx-rt1 : 87
IP               Advertised     Denied   Accepted      Limit State   Description
192.168.xxx.xxx       72280          0        495   No Limit Up      xxxxx-RT1
192.168.xxx.xxx       72280          0         97   No Limit Up      xxxxx-RT2
192.168.xxx.xxx       72280          0       4056   No Limit Up      xxxxx-RT3
192.168.xxx.xxx       72280          0     482394   No Limit Up      xxxxx-RT4
192.168.xxx.xxx       72280          0     474874   No Limit Up      xxxxx-RT5
192.168.xxx.xxx       72280          0        858   No Limit Up      xxxxx-RT6
10.100.xxx.xxx            1          0          1        100 Up      CustomerA
10.101.xxx.xxx            0          0          0        100 Active  CustomerB
10.102.xxx.xxx            1         24          1        100 Up      CustomerC
10.103.xxx.xxx            0          0          0        100 Down    CustomerD
10.104.xxx.xxx            0          0          0        100 Down    CustomerE
10.105.xxx.xxx        65035         19         15        100 Up      CustomerF
```

Cisco sh bgp summary output:

```
Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.10.xxx.xxx   4        xxxxx   38723   43490 103858955    0    0 6d17h           0
10.20.xxx.xxx   4        xxxxx  506534  489968 103858955    0    0 11w2d           1
10.30.xxx.xxx   4        xxxxx  474942  567092 103858955    0    0 11w2d           1
10.100.xxx.xxx  4        xxxxx       0       0         1    0    0 never        Idle
```

IOS-XR VRFs

- You can create SNMP contexts in IOS-XR 4.x
- https://supportforums.cisco.com/thread/2219379
- bug referenced is private, though.
- not sure on 5.x, probably still works

Alcatel 7210

- If you configure epipe services on ALU 7210, these aren't in the normal ifTable
- VENDOR FAIL, this makes it annoying to monitor, I assume they hate their customers
- No support by default in a few nms tools I tried
- sapBaseStatsIngressForwardedPackets 1.3.6.1.4.1.6527.6.2.2.2.8.1.1.1.3.random.random.random
- 1.3.6.1.4.1.6527.3.1.2.4.3.2.1.5 - description table
- has speechmarks in the regular ifTable, so don't use that either, annoying to parse output
- 1.3.6.1.4.1.6527.3.1.2.2.4.2.1.5.1 - slightly better formatted interface table

Bluecoat PacketShaper

- HA HA HA You thought you could get 64 bit interface counters
- This is important in high bandwidth applications, because 32 bit counters will wrap within your 5 minute polling period
- Instead of one 64 bit counter, has two 32 bit OIDs, one with the high 32 bits, one with the low 32 bits
- high * 2**32 + low and you get your 64 bit counter
- no direct support for that in most standard nms tools
- we had to write custom target module to support this type of vendor fail

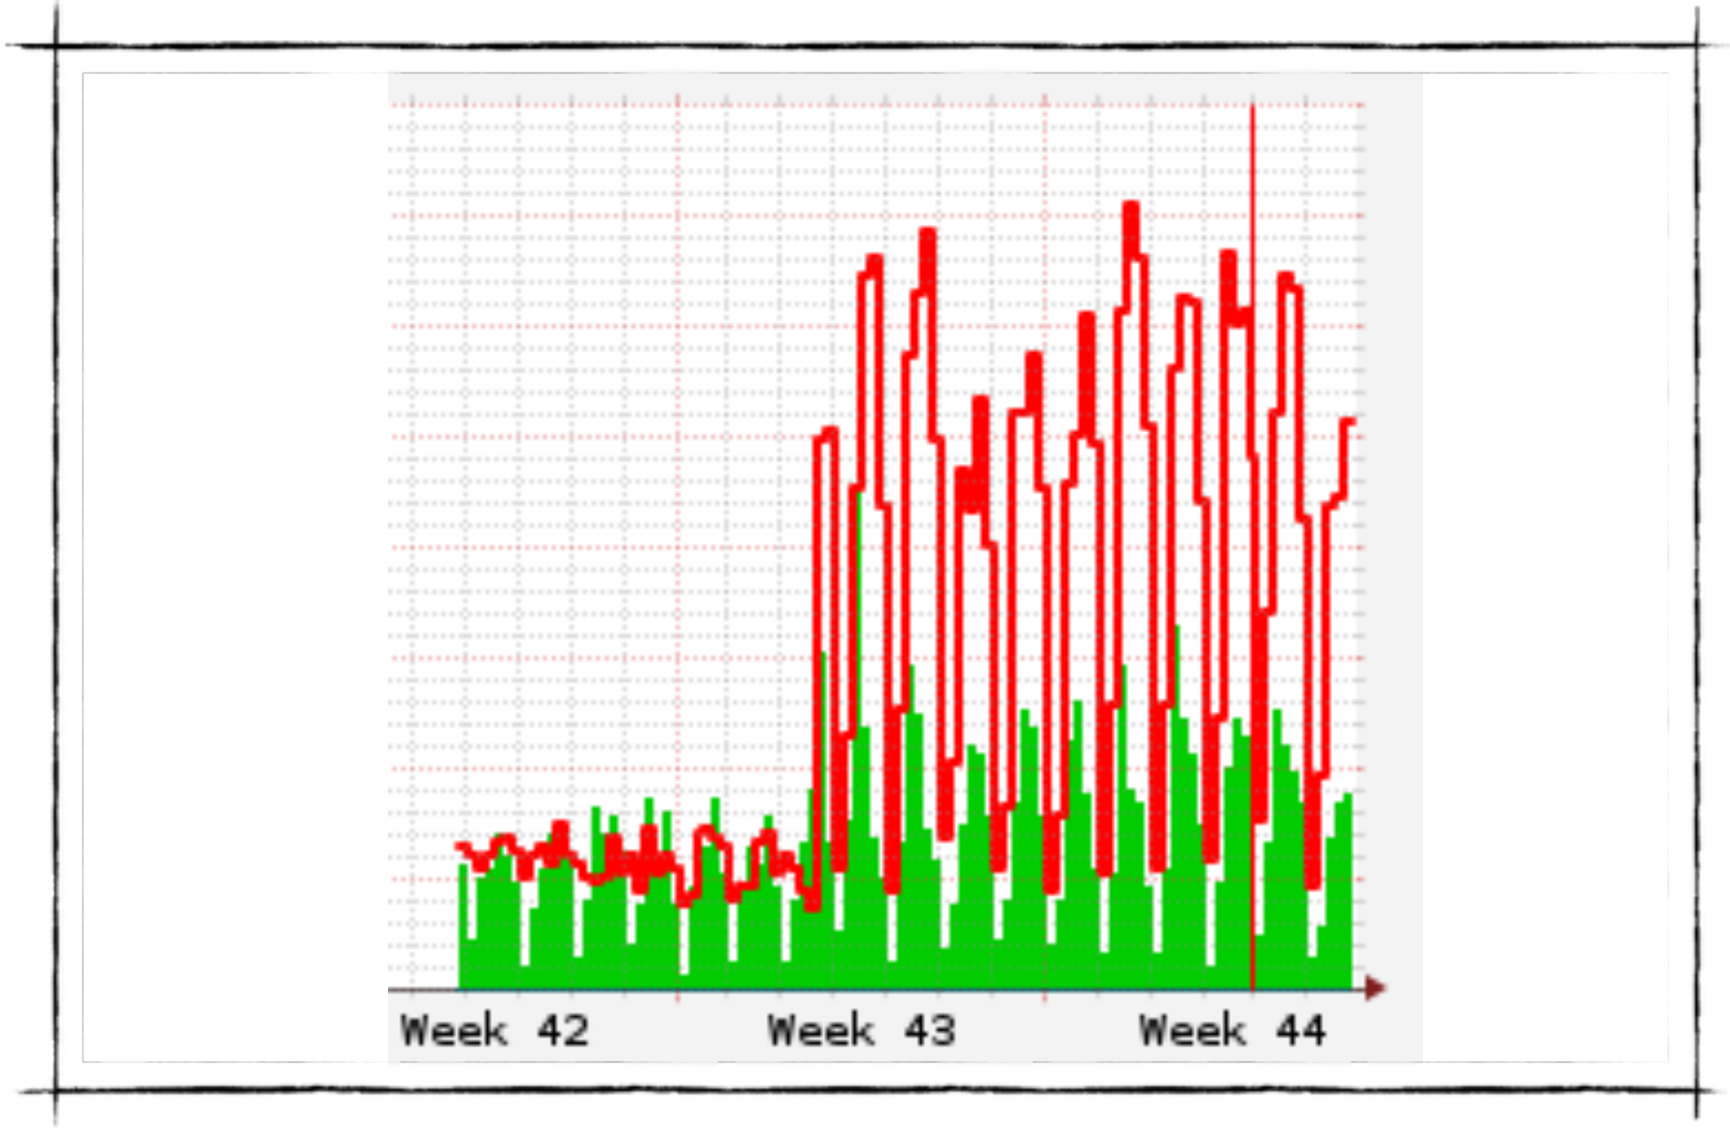

Graph fail + fix

Internet mapping

- Been trying to make a map of the internet for a long time
- A few people have done this already with varying successes
- http://www.caida.org/research/topology/
- I think the maps they made are all boring and fail at being cool
- Surely this can't be too hard, I'll just do it myself!

Data Source

- Need data sources
- ape/wix looking glass (fail)
- www.routeviews.org has historical route table archives
- Lets just take the “sh ip bgp” type output
- mrt files are probably useful, but I’ve ignored them for now

BGP Looking Glasses

- So you want some info on the BGP table, but you don’t have a router
- Lets look at the APE/WIX and try and get a full list of routes
- type regexp ^ into http://nzix.net/cgi-bin/lg.cgi
- Select rs1.wix
- If it worked nice you should see at the bottom a total
- Total number of prefixes 1224

BGP Looking Glasses

- Try again with rs1.ape or rs2.ape
- Hmm, nope, don’t get that summary at the end. maybe it was just a temporary blip. Try again.
- Nooooooope. still busted.
- Looking Glass FAIL

ASNs

- ASN = Autonomous system number
- Links between ASNs make up the internet
- Graph those links, and you'll graph the internet (from the perspective of the BGP route collection boxes)
- Use AWK to parse

```
grep "\ [i,e]$" oix-full-snapshot-2012-12-06-0600 | grep -v “{" |

awk '{

if (NF-1 > 7)
for (i = 7; i <= NF-2; i++)

if (($i != $(i+1)) && ($i != 0) && ($(i+1) != 0)) print $i,”,",$(i+1)
}’ | sort | uniq > 2012-12.txt
```

Maxmind

- Maxmind has a free dataset, you may be familiar with it
- Tie that together with aforementioned ASN paths to try and guess at geolocation
- Use some graphing software to make a map of the world

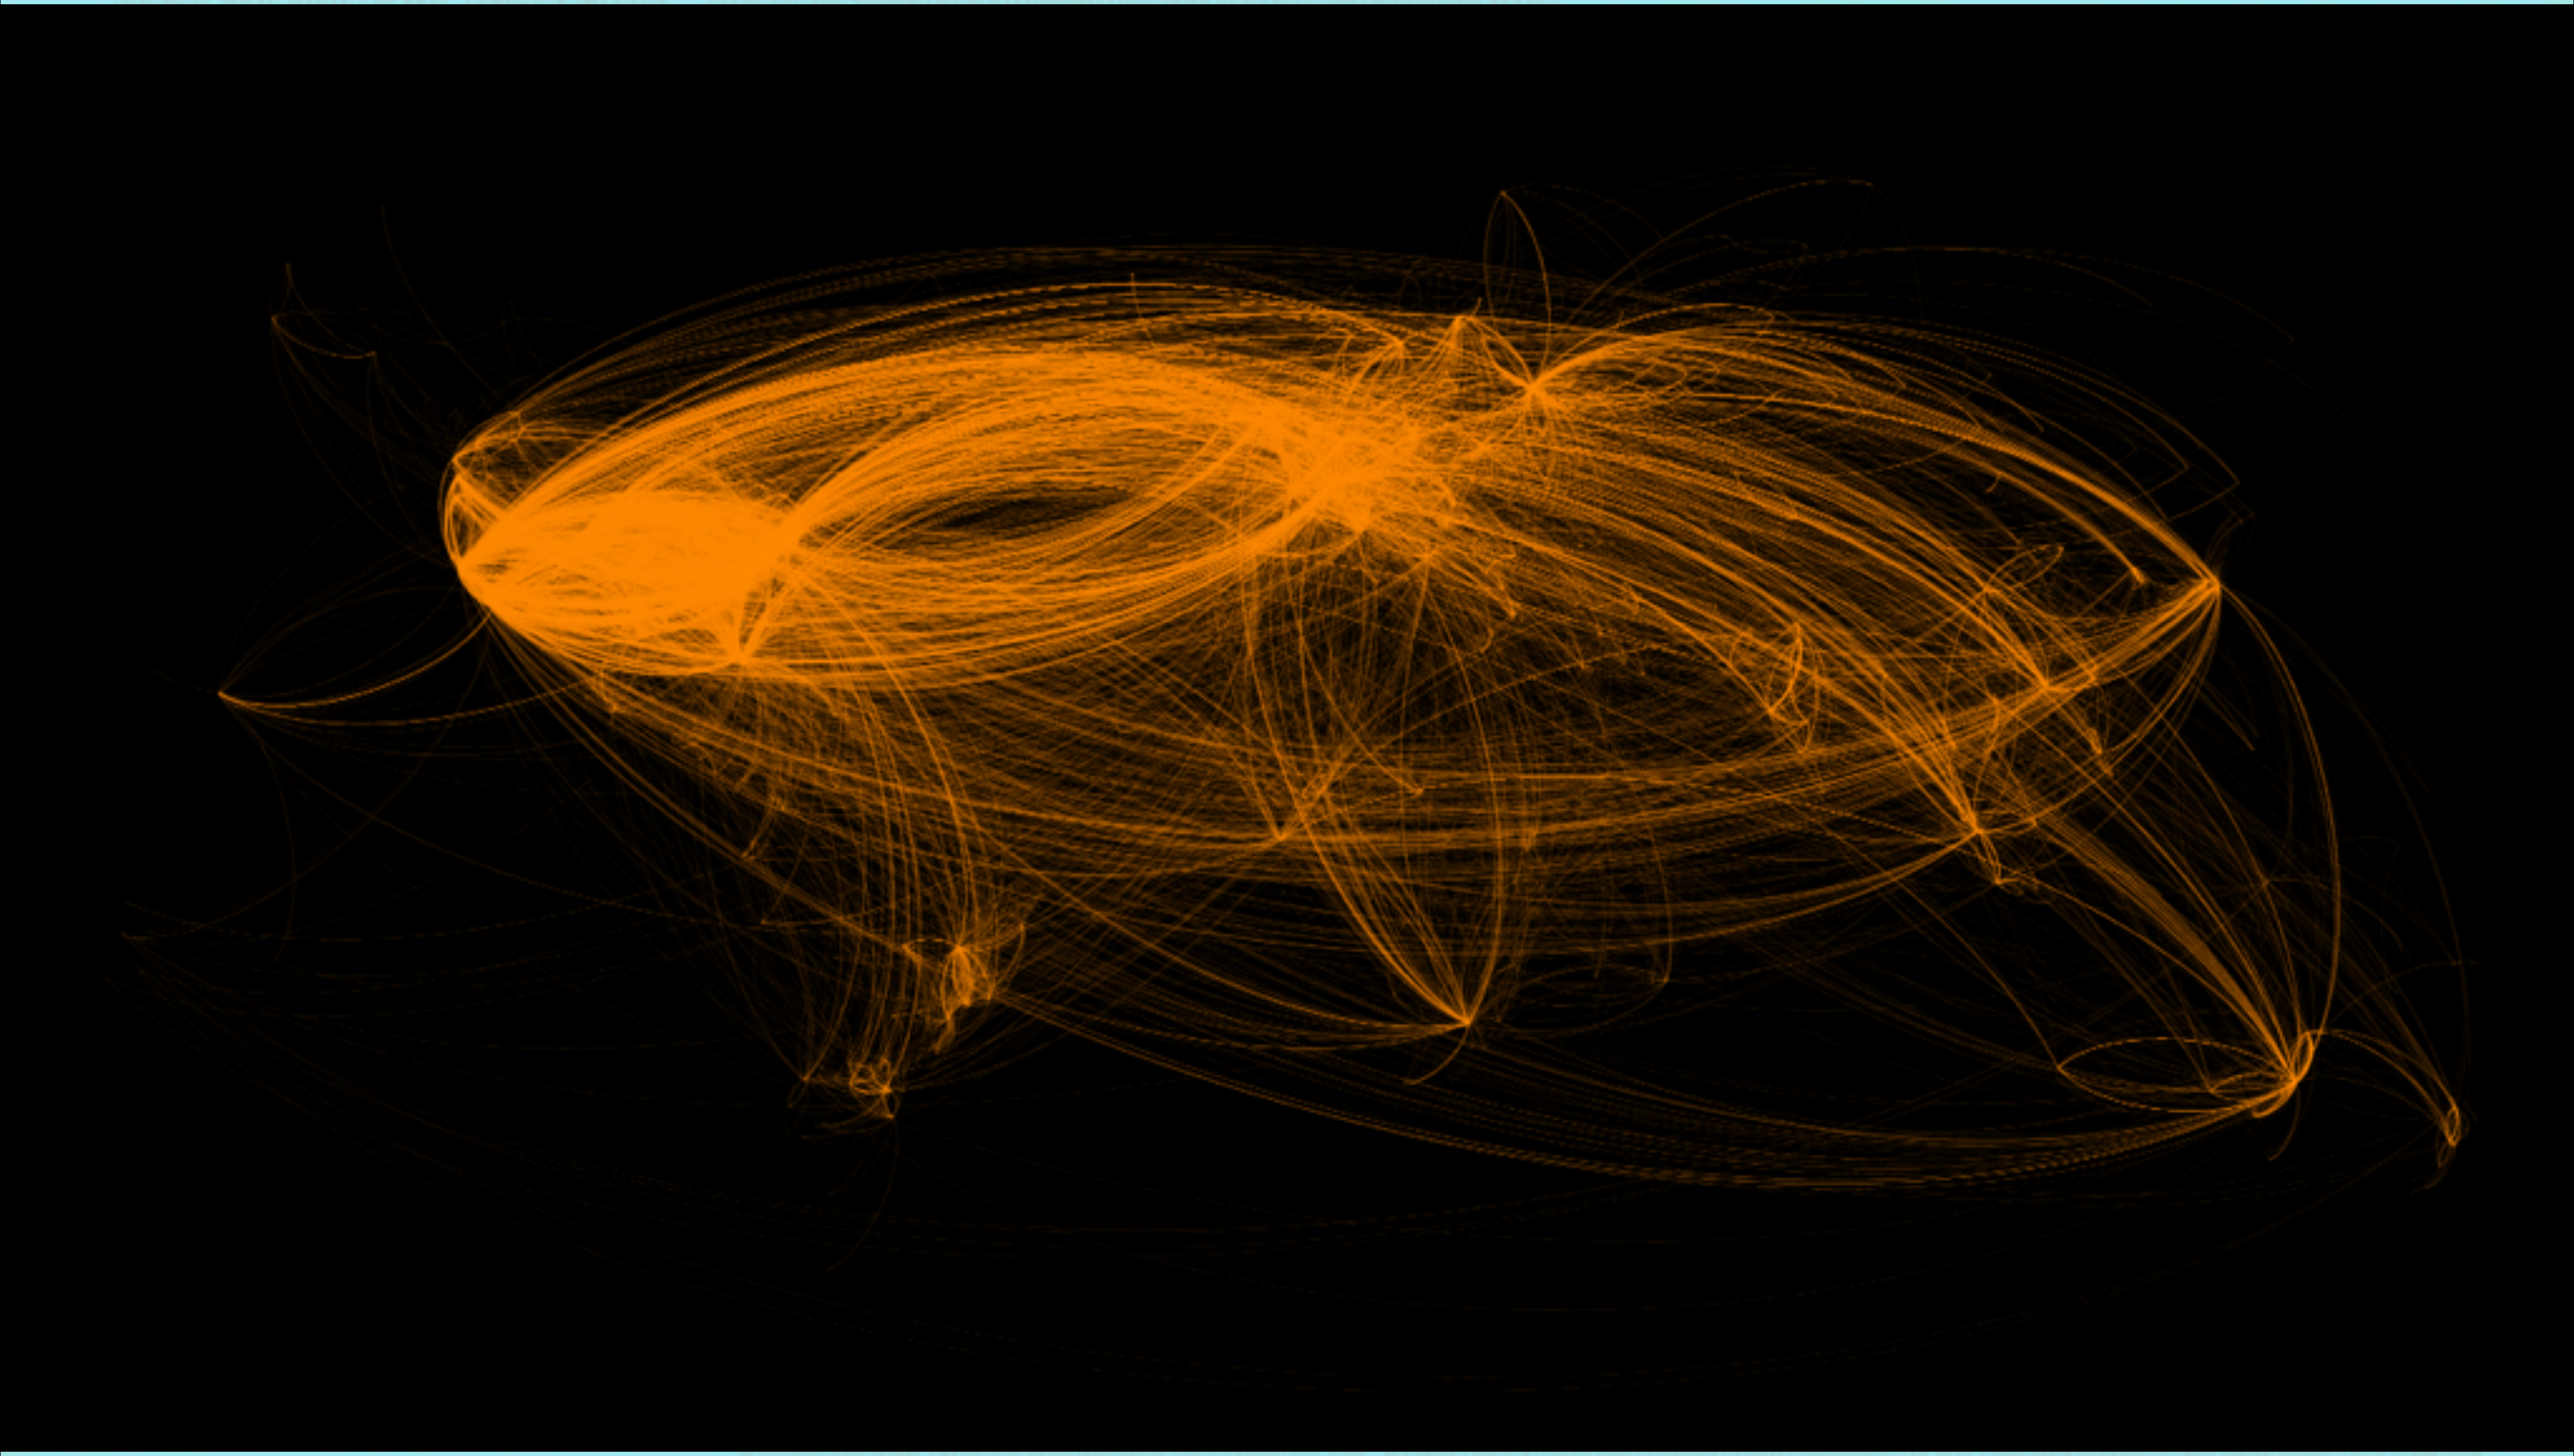

Internet Map

- Well that was kind of boring, actually
- What if I put the AS Names on it?
- Might need to be able to zoom in on it
- end up with stupidly huge PNG files, and you still can't read it properly
- FAIL

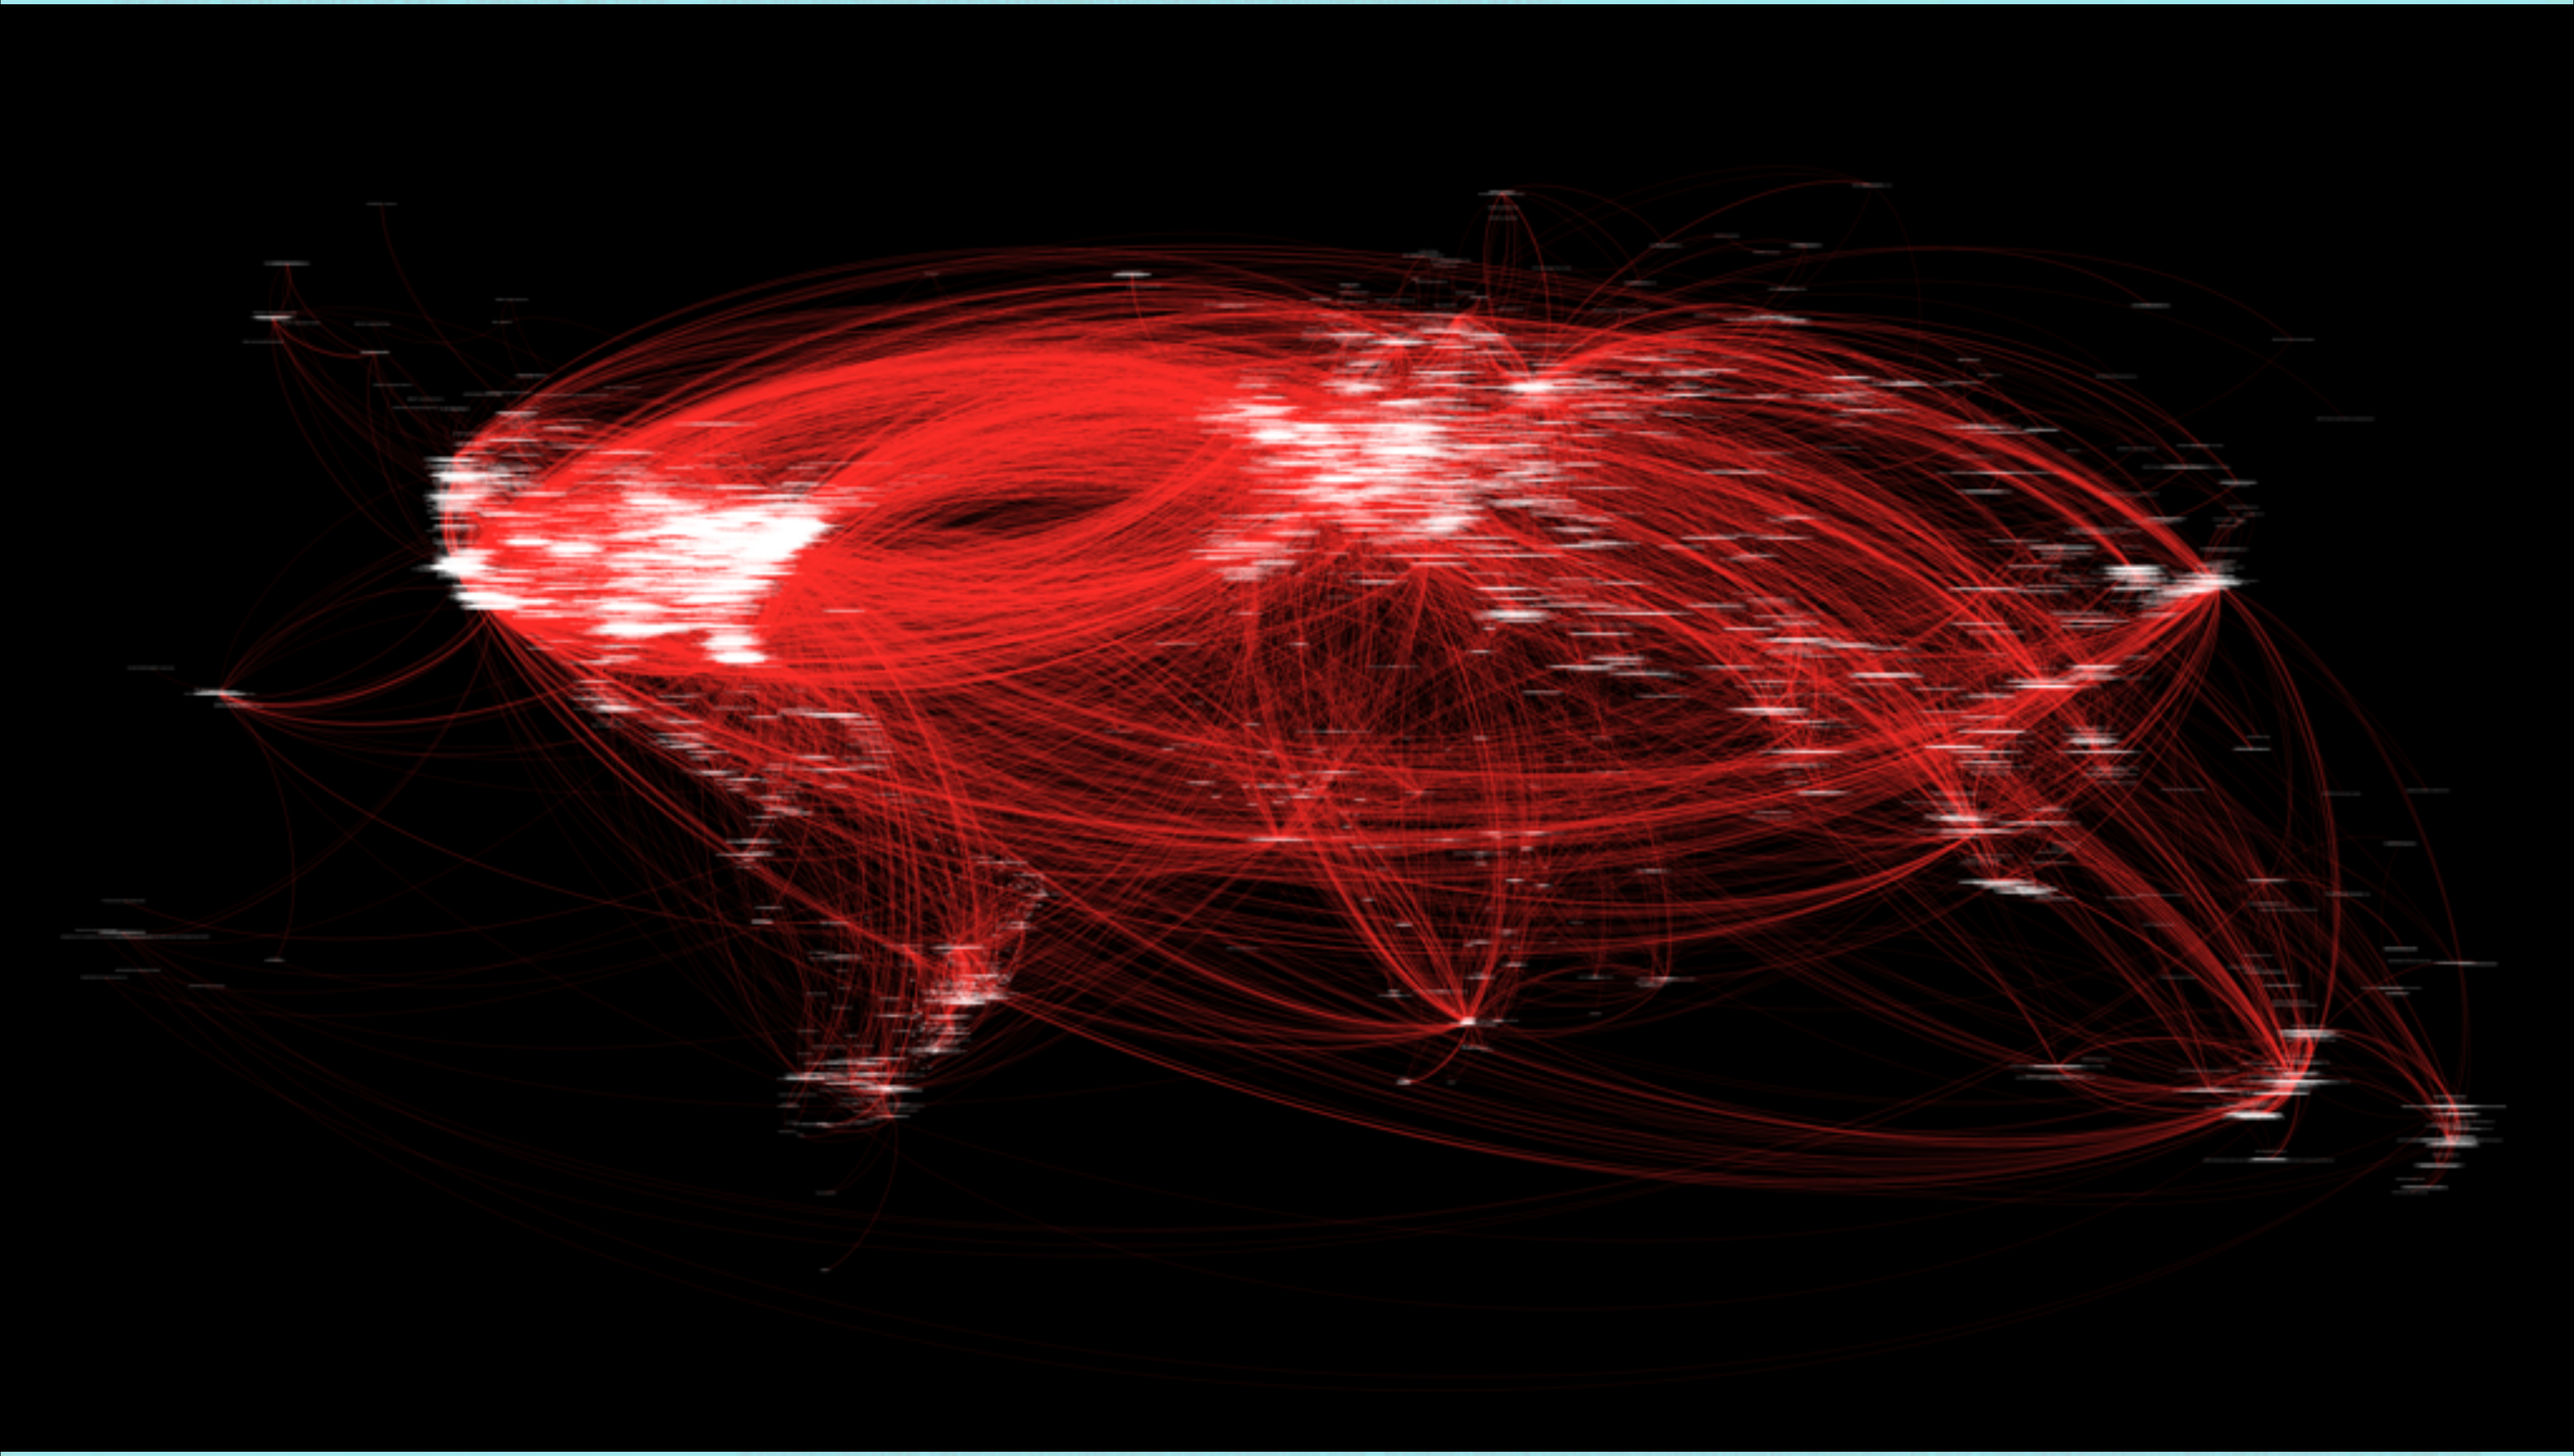

SVG? KML?

- Browsers suck-ass at this kind of map represented as an SVG file
- Maybe try a KML?
- Google earth will load it but it's just a mass of white everywhere
- About ready to give up?

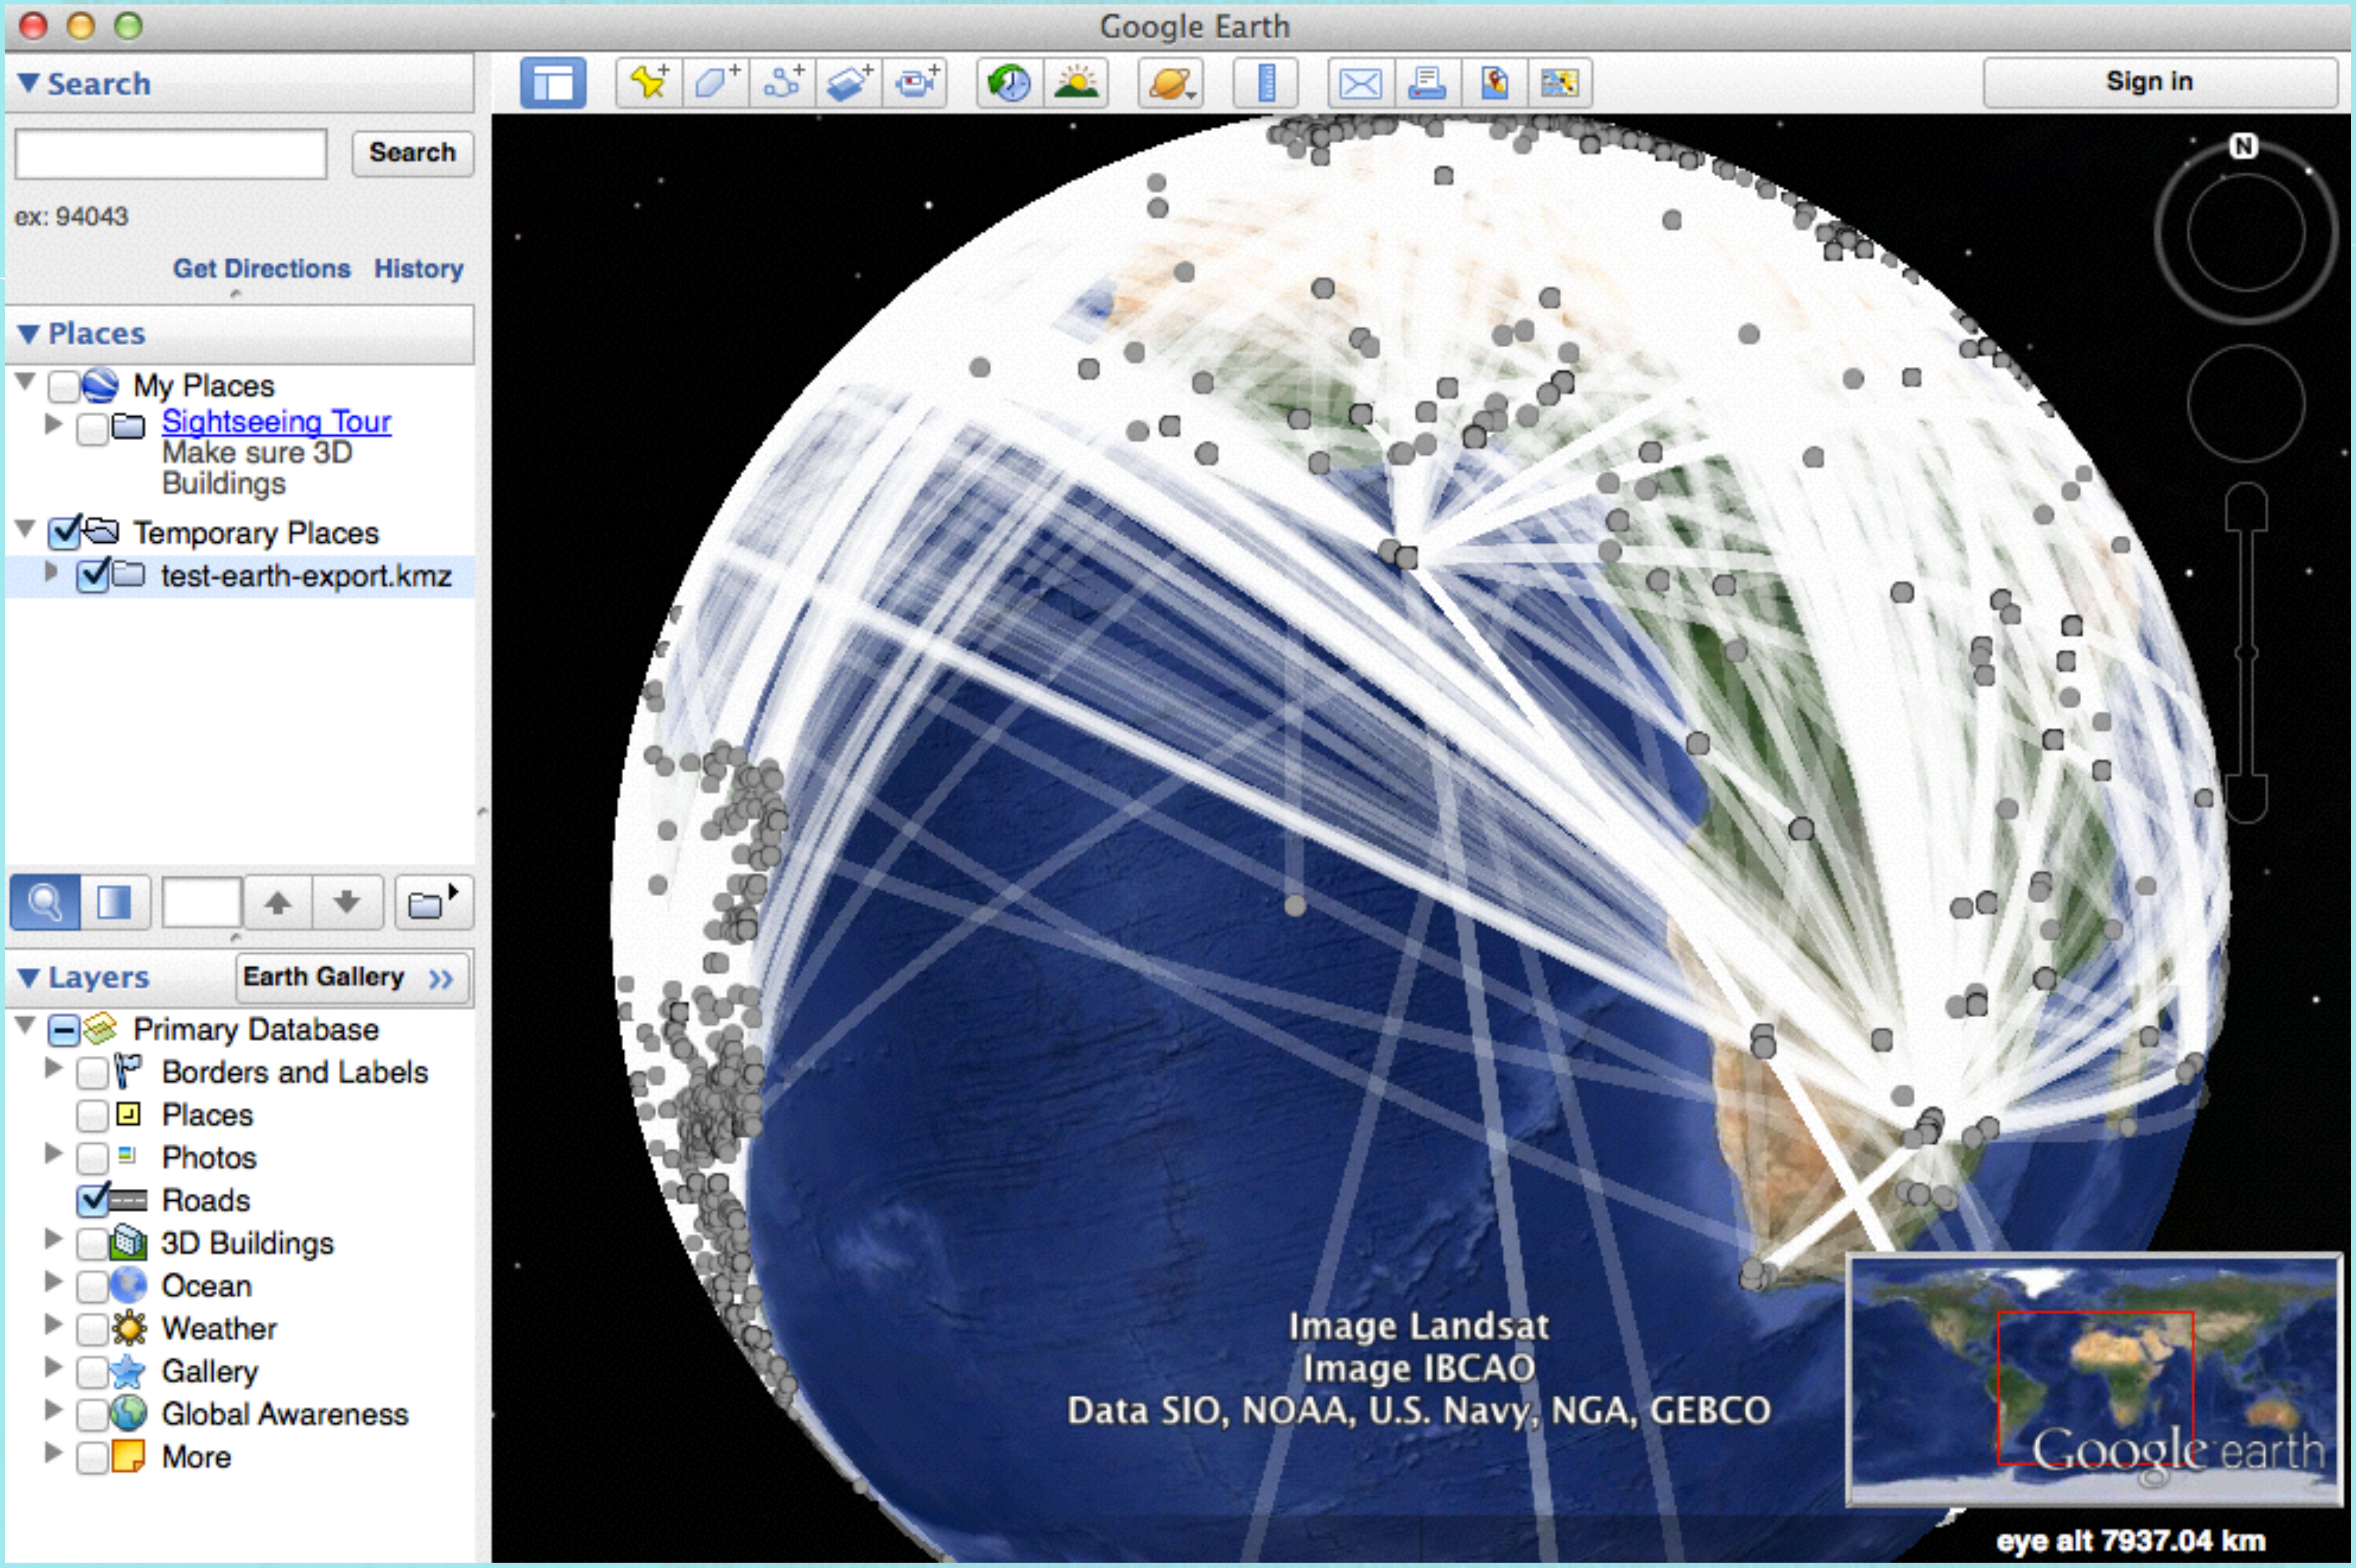

☐ My Places
☐ Sightseeing Tour
Make sure 3D Buildings
☑ Temporary Places
☑ test-earth-export.kmz

☐ Primary Database
☐ Borders and Labels
☐ Places
☐ Photos
☑ Roads
☐ 3D Buildings
☐ Ocean
☐ Weather
☐ Gallery
☐ Global Awareness
☐ More

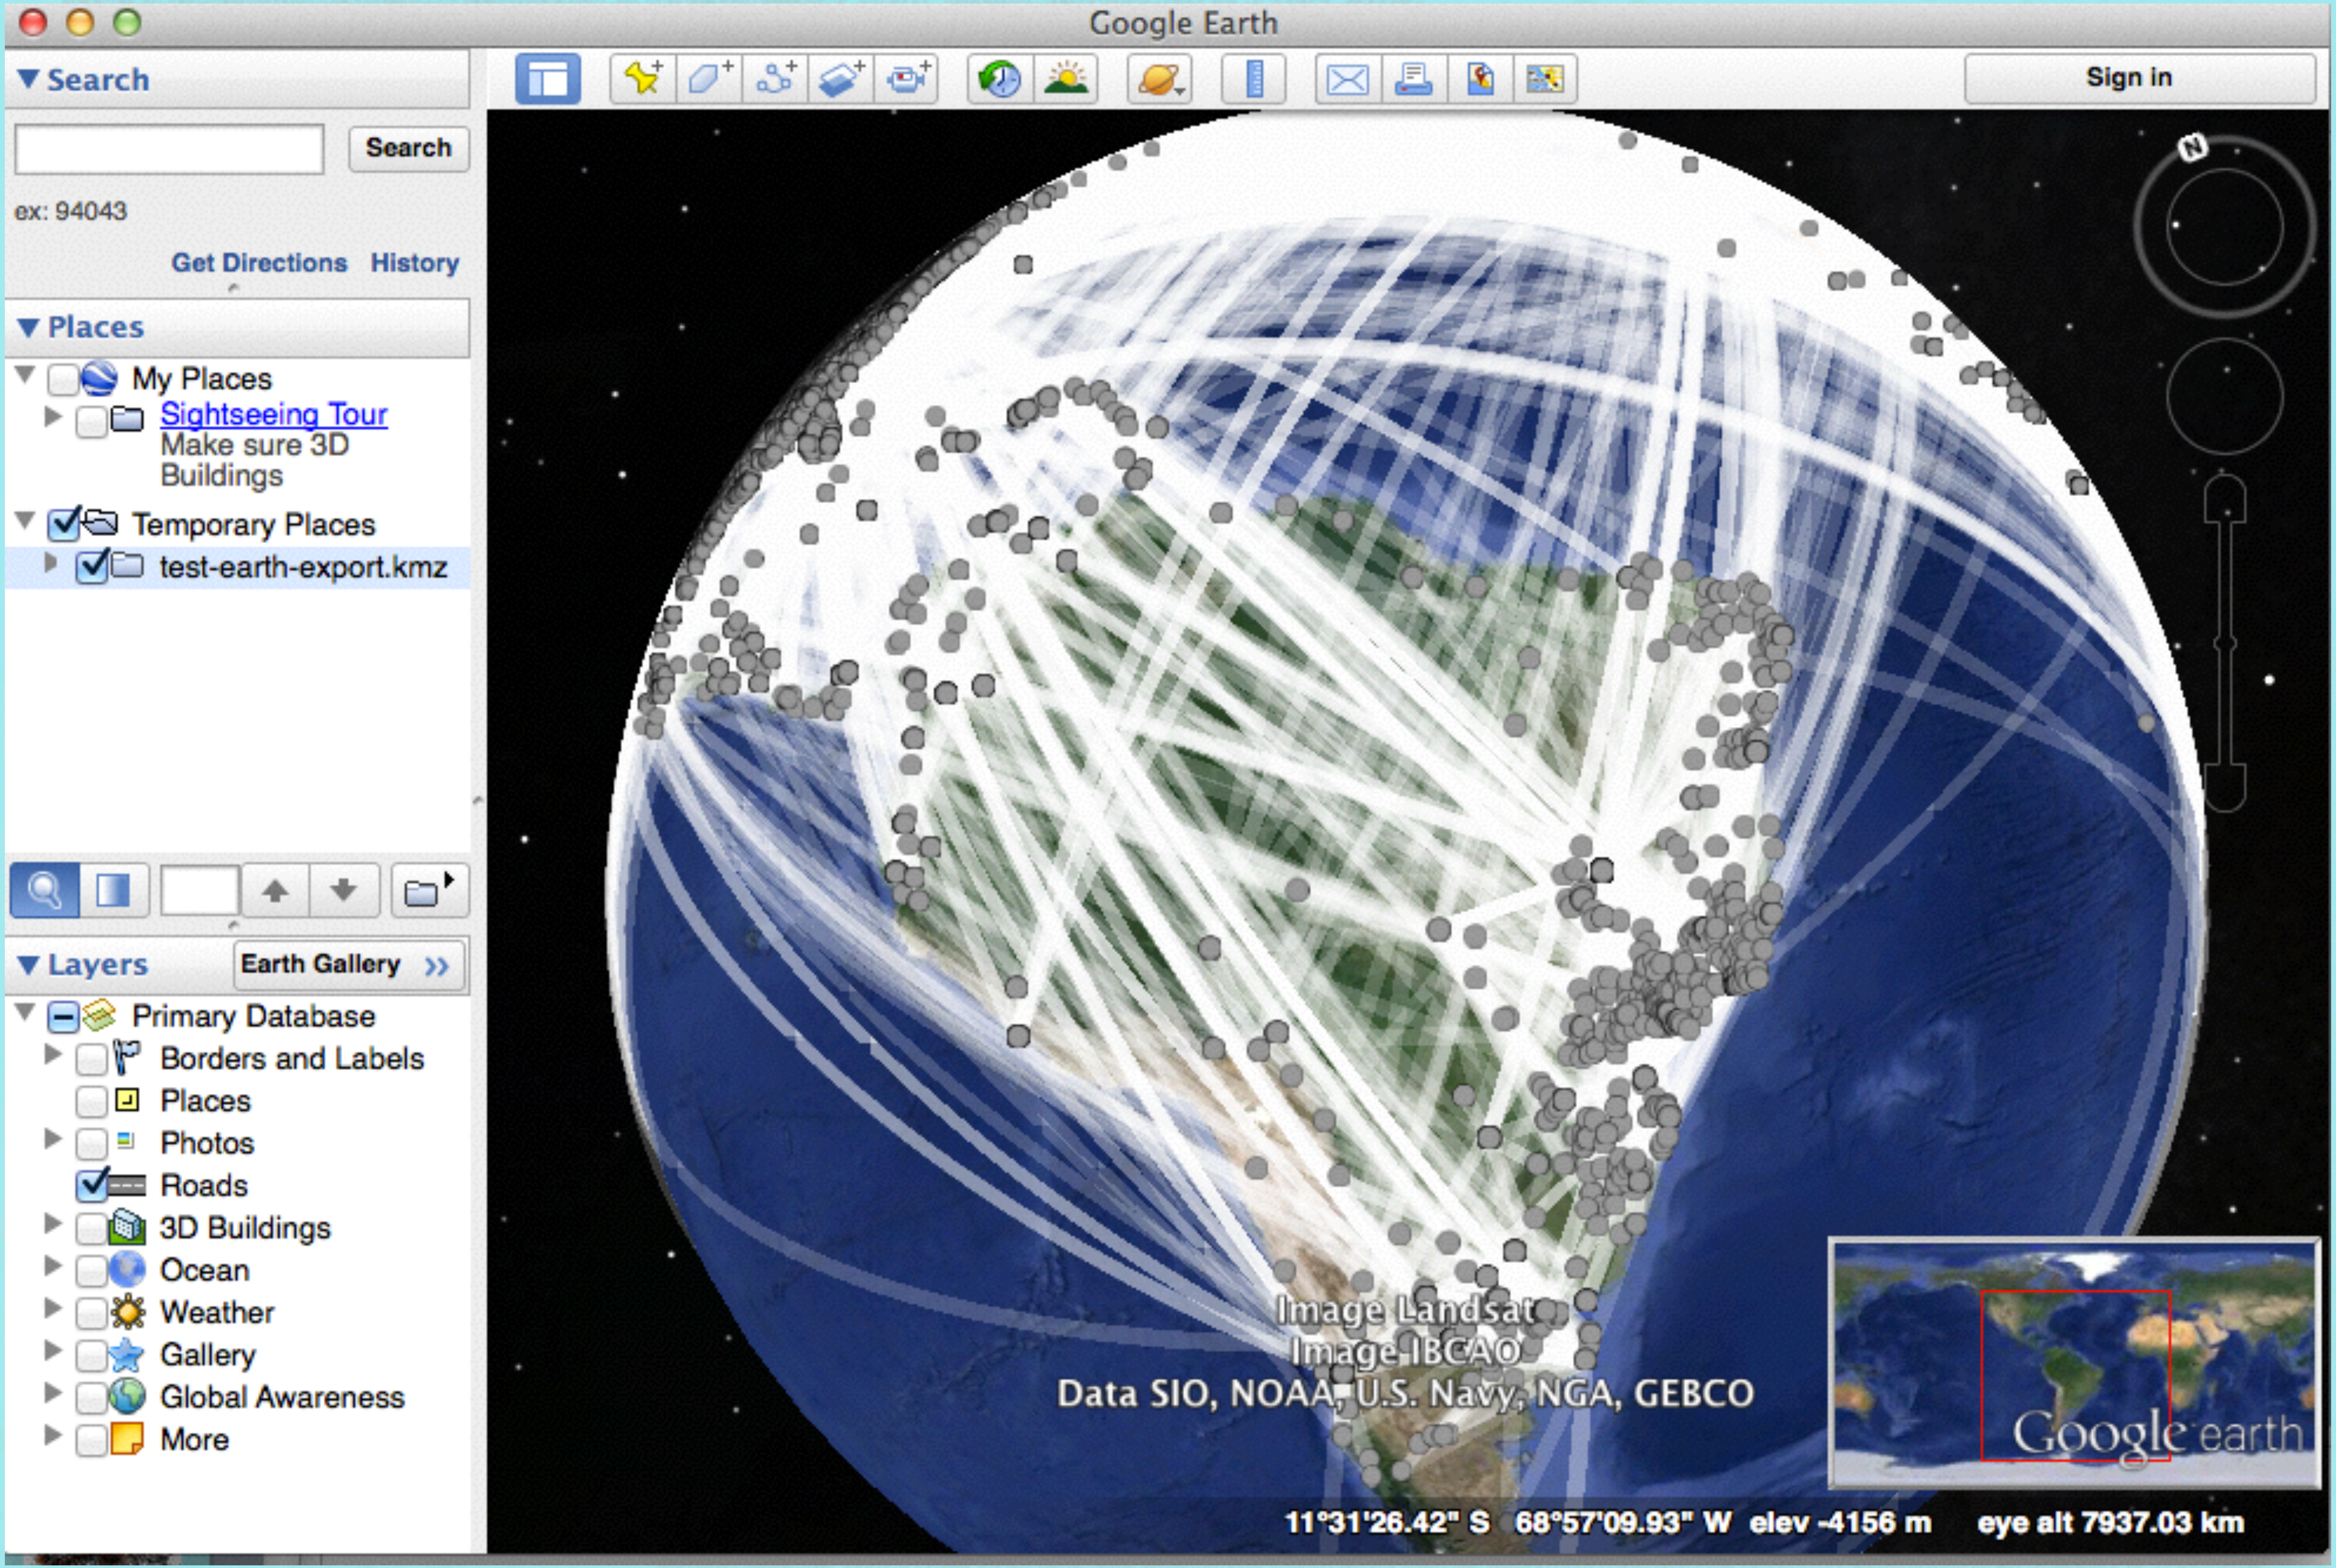

☐ My Places
☐ Sightseeing Tour
☑ Temporary Places
☑ test-earth-export.kmz
⊟ Primary Database
☐ Borders and Labels
☐ Places
☐ Photos
☑ Roads
☐ 3D Buildings
☐ Ocean
☐ Weather
☐ Gallery
☐ Global Awareness
☐ More

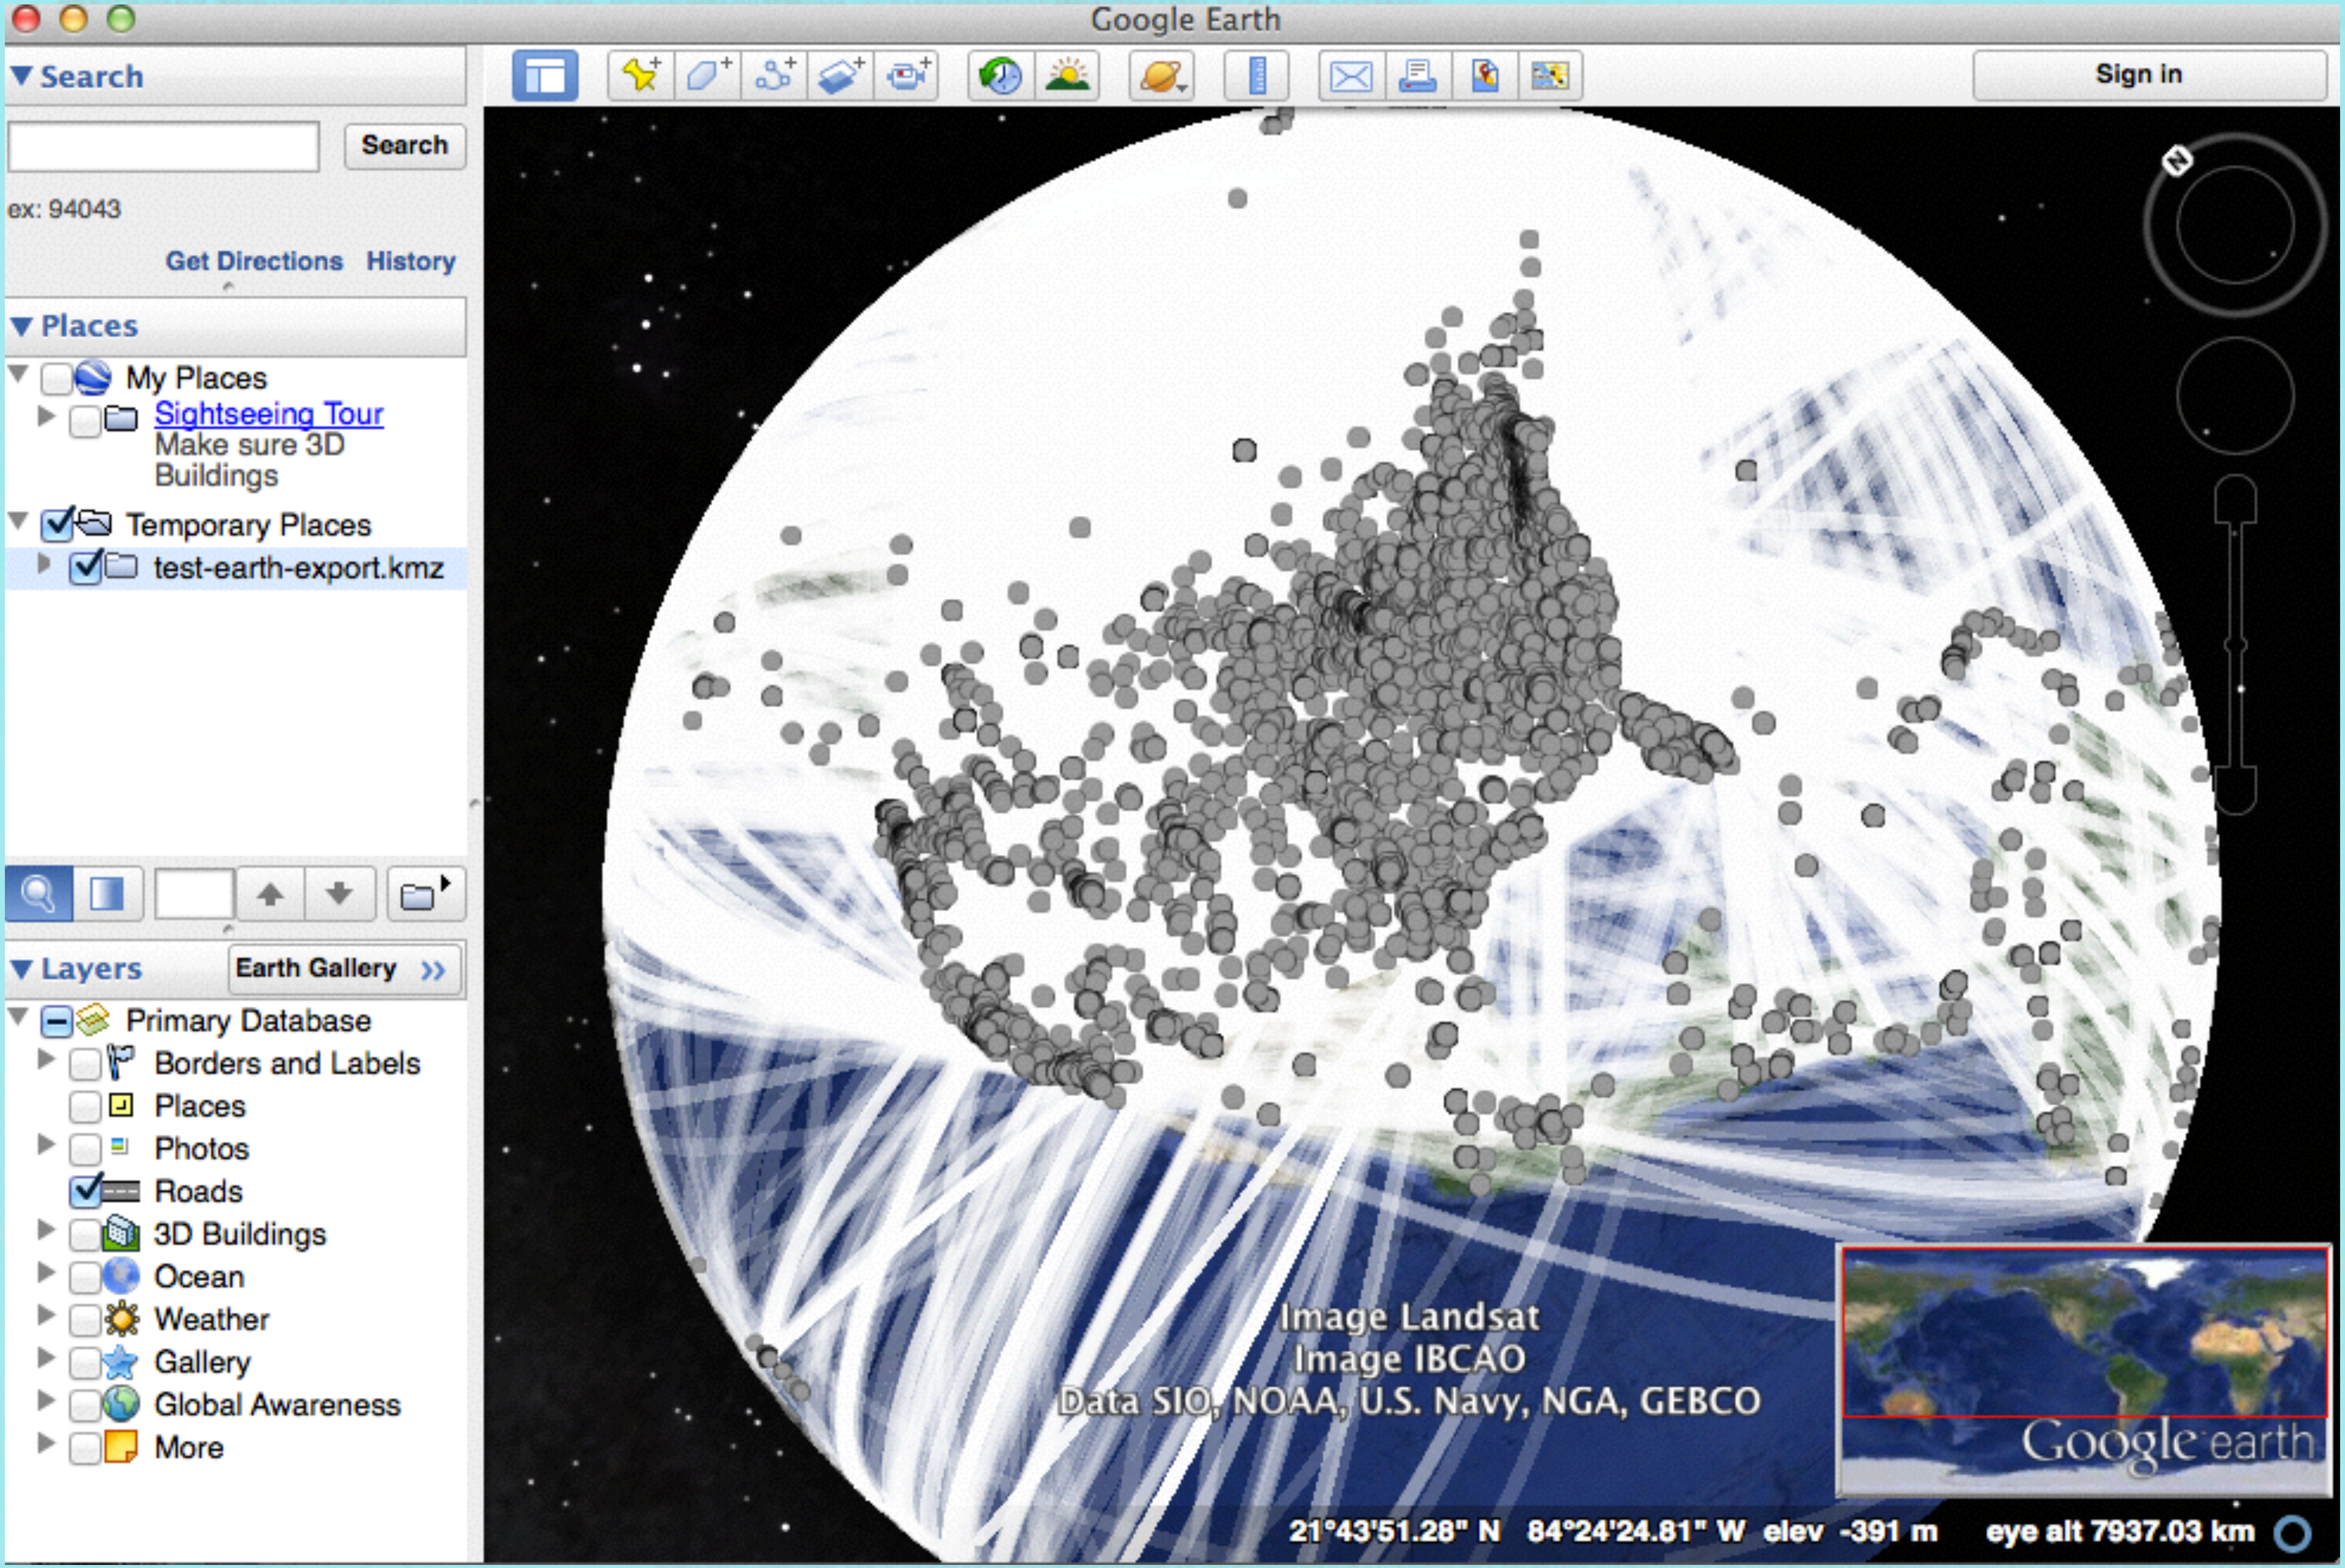

- □ My Places
  - □ Sightseeing Tour
    Make sure 3D Buildings
- ☑ Temporary Places
  - ☑ test-earth-export.kmz

- ⊟ Primary Database
  - □ Borders and Labels
  - □ Places
  - □ Photos
  - ☑ Roads
  - □ 3D Buildings
  - □ Ocean
  - □ Weather
  - □ Gallery
  - □ Global Awareness
  - □ More

21°43'51.28" N 84°24'24.81" W elev -391 m eye alt 7937.03 km

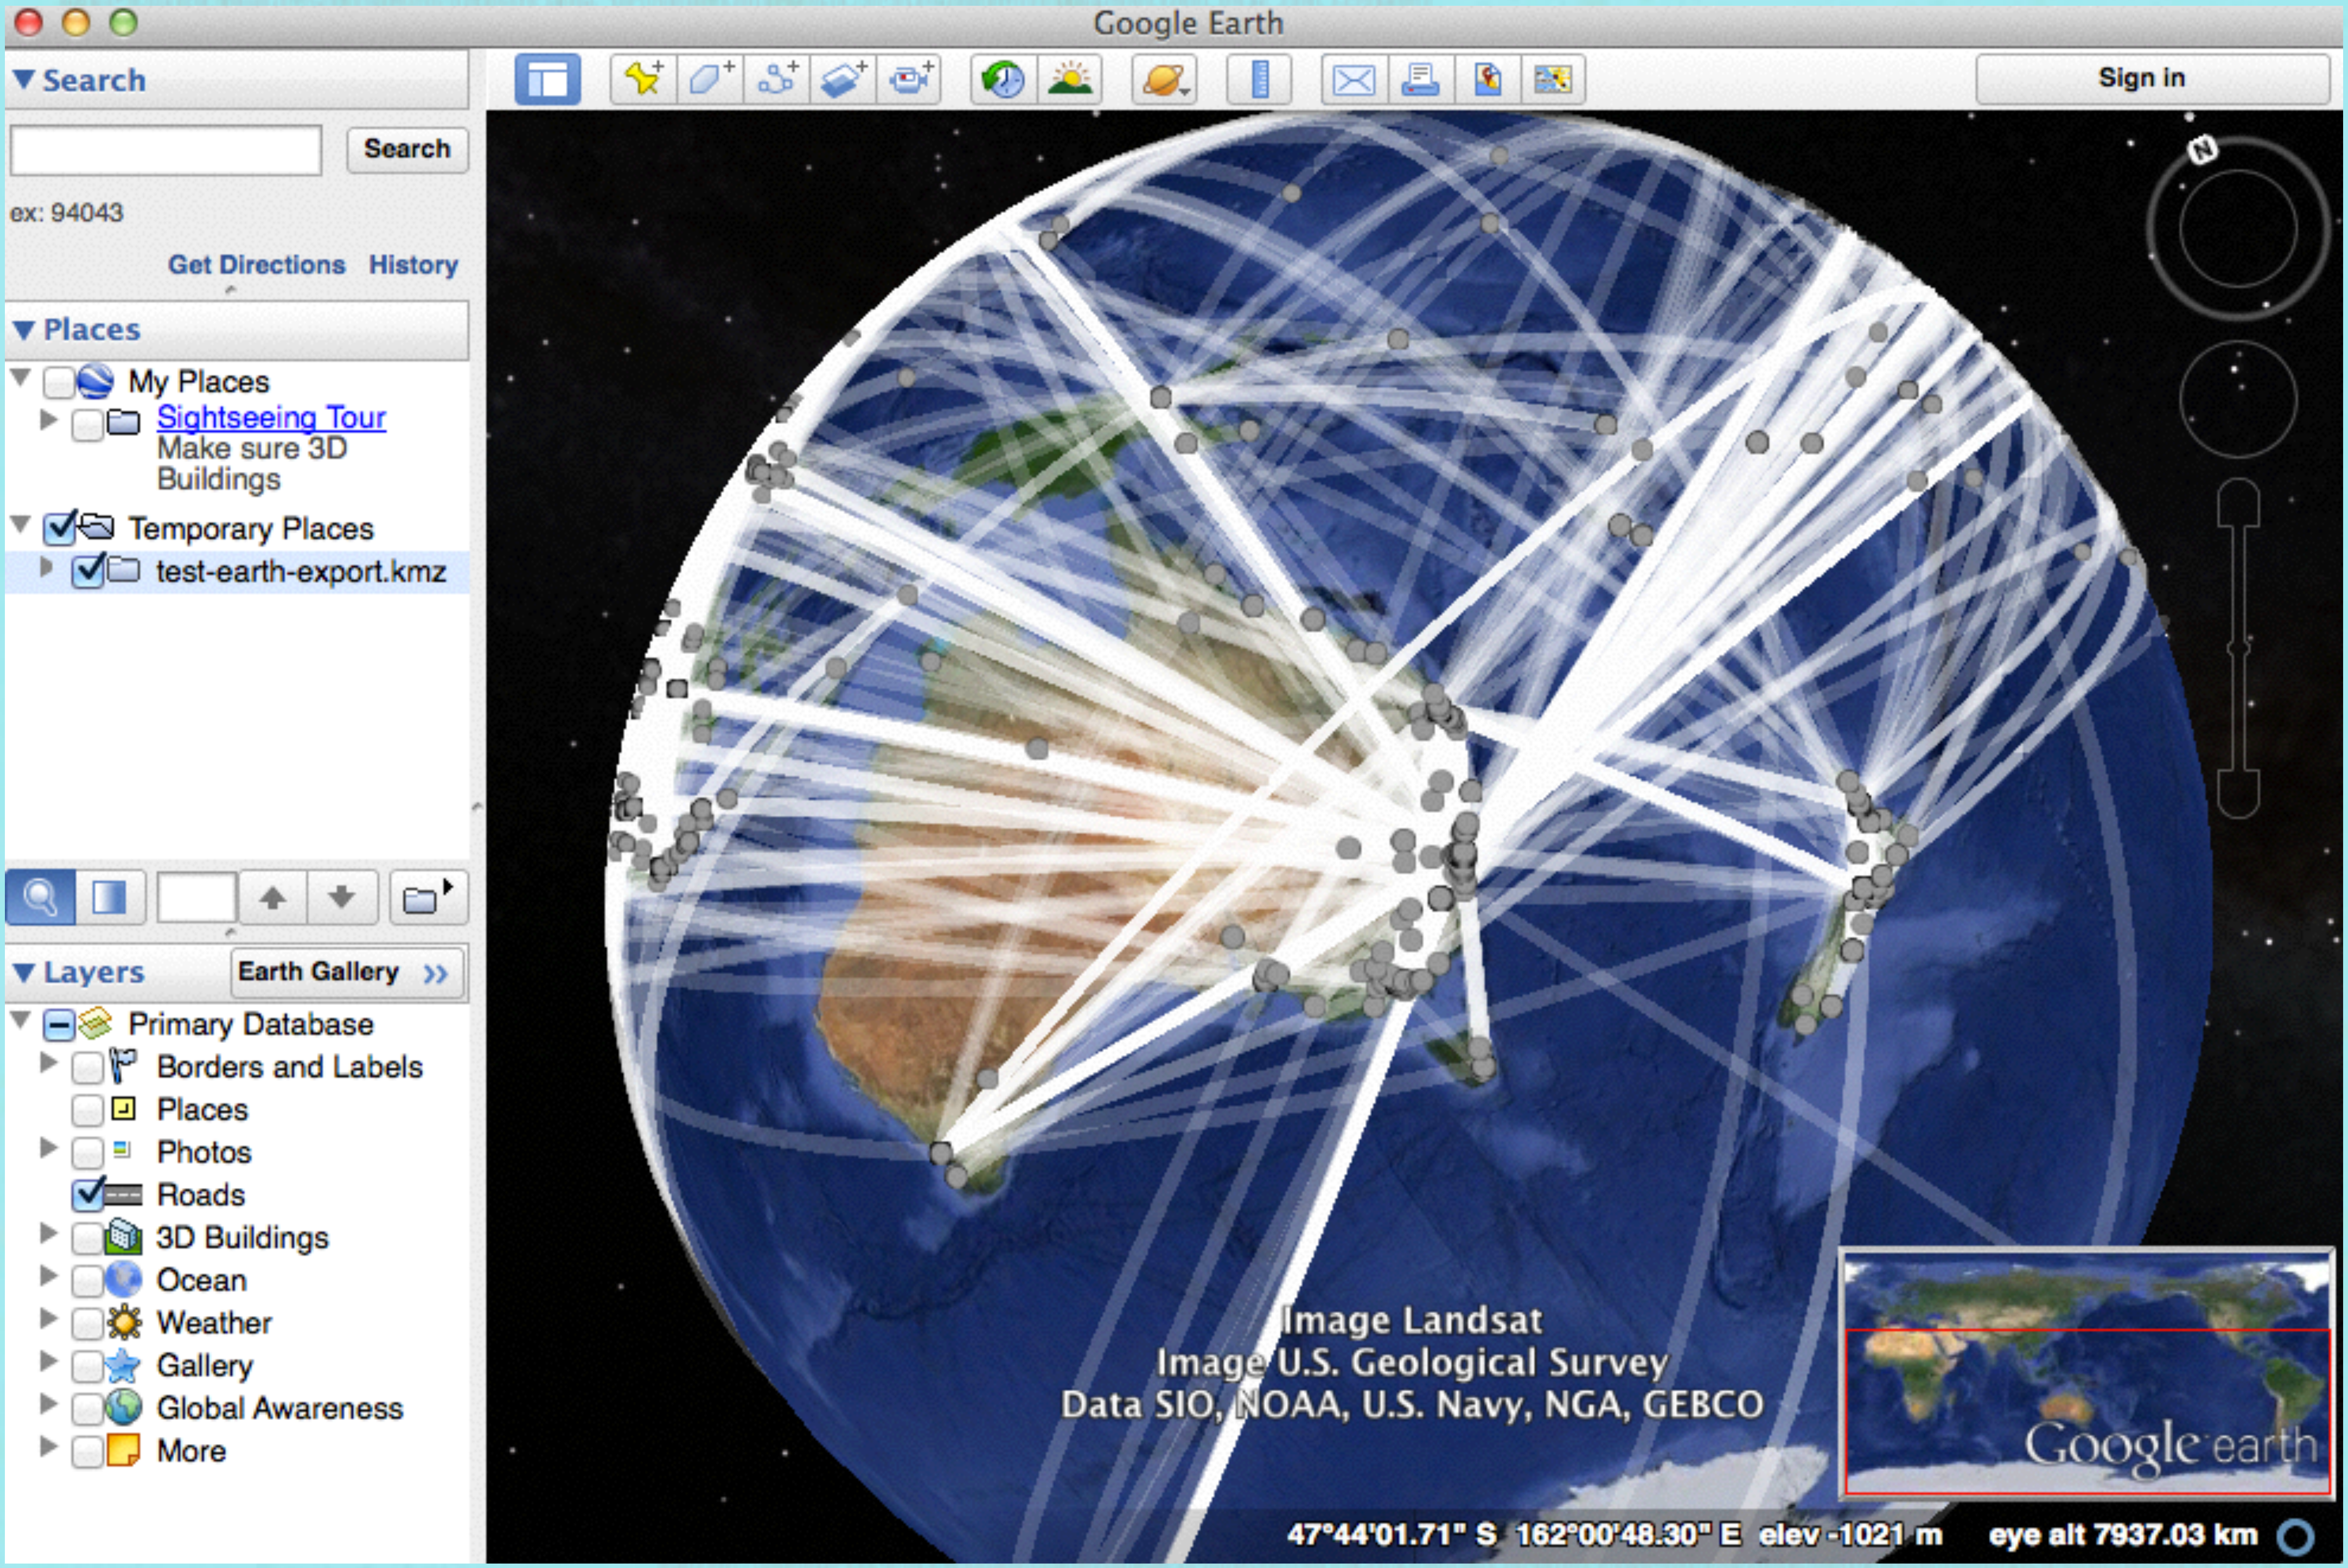
Google Earth
Search
Search
ex: 94043
Get Directions
History
Places
☐ My Places
☐ Sightseeing Tour
Make sure 3D Buildings
☑ Temporary Places
☑ test-earth-export.kmz
Layers
Earth Gallery
Primary Database
☐ Borders and Labels
☐ Places
☐ Photos
☑ Roads
☐ 3D Buildings
☐ Ocean
☐ Weather
☐ Gallery
☐ Global Awareness
☐ More
Sign in
N
Image Landsat
Image U.S. Geological Survey
Data SIO, NOAA, U.S. Navy, NGA, GEBCO
Google earth
47°44'01.71" S 162°00'48.30" E elev -1021 m eye alt 7937.03 km

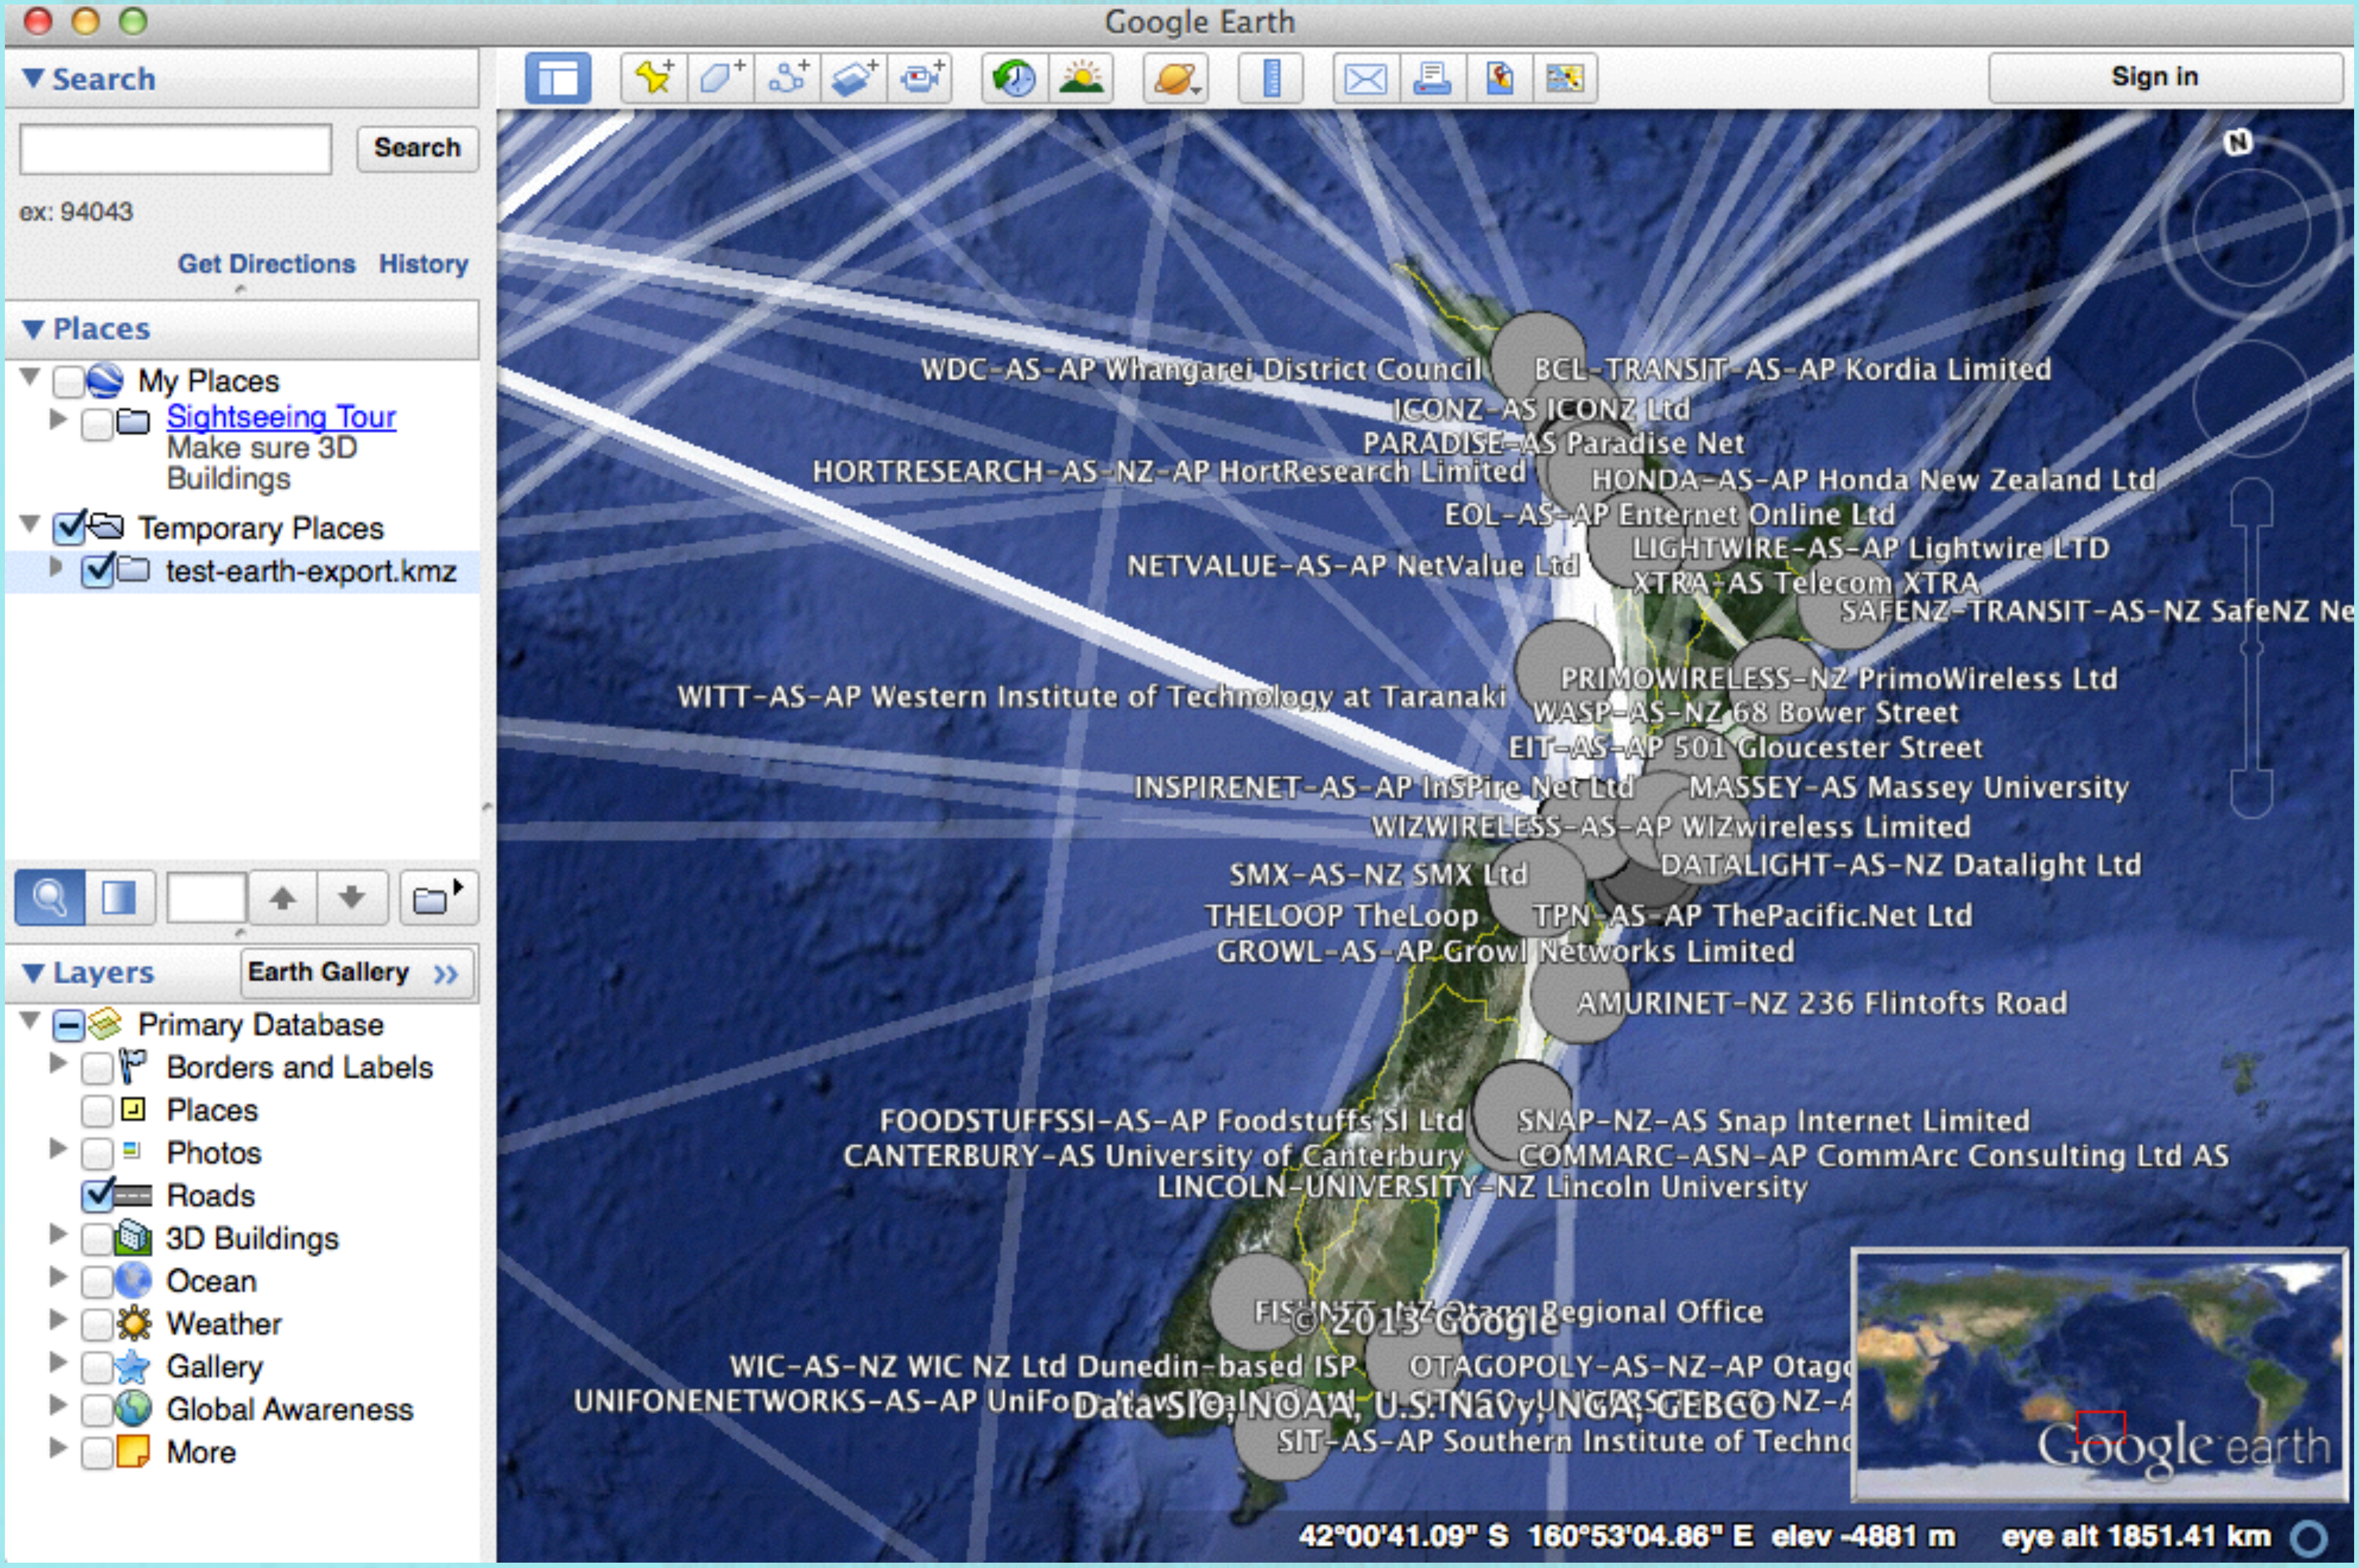
Google Earth
Search
Search
ex: 94043
Get Directions
History
Places
My Places
Sightseeing Tour
Make sure 3D Buildings
☑ Temporary Places
☑ test-earth-export.kmz
Layers
Earth Gallery
Primary Database
□ Borders and Labels
□ Places
□ Photos
☑ Roads
□ 3D Buildings
□ Ocean
□ Weather
□ Gallery
□ Global Awareness
□ More
Sign in
WDC-AS-AP Whangarei District Council
BCL-TRANSIT-AS-AP Kordia Limited
ICONZ-AS ICONZ Ltd
PARADISE-AS Paradise Net
HORTRESEARCH-AS-NZ-AP HortResearch Limited
HONDA-AS-AP Honda New Zealand Ltd
EOL-AS-AP Enternet Online Ltd
LIGHTWIRE-AS-AP Lightwire LTD
NETVALUE-AS-AP NetValue Ltd
XTRA-AS Telecom XTRA
SAFENZ-TRANSIT-AS-NZ SafeNZ Ne
PRIMOWIRELESS-NZ PrimoWireless Ltd
WITT-AS-AP Western Institute of Technology at Taranaki
WASP-AS-NZ 68 Bower Street
EIT-AS-AP 501 Gloucester Street
INSPIRENET-AS-AP InSPire Net Ltd
MASSEY-AS Massey University
WIZWIRELESS-AS-AP WIZwireless Limited
DATALIGHT-AS-NZ Datalight Ltd
SMX-AS-NZ SMX Ltd
THELOOP TheLoop
TPN-AS-AP ThePacific.Net Ltd
GROWL-AS-AP Growl Networks Limited
AMURINET-NZ 236 Flintofts Road
FOODSTUFFSSI-AS-AP Foodstuffs SI Ltd
SNAP-NZ-AS Snap Internet Limited
CANTERBURY-AS University of Canterbury
COMMARC-ASN-AP CommArc Consulting Ltd AS
LINCOLN-UNIVERSITY-NZ Lincoln University
FIS NET NZ Otago Regional Office
© 2013 Google
WIC-AS-NZ WIC NZ Ltd Dunedin-based ISP
OTAGOPOLY-AS-NZ-AP Otag
UNIFONENETWORKS-AS-AP UniFo
Data SIO, NOAA, U.S. Navy, NGA, GEBCO
SIT-AS-AP Southern Institute of Techn
Google earth
42°00'41.09" S 160°53'04.86" E elev -4881 m eye alt 1851.41 km

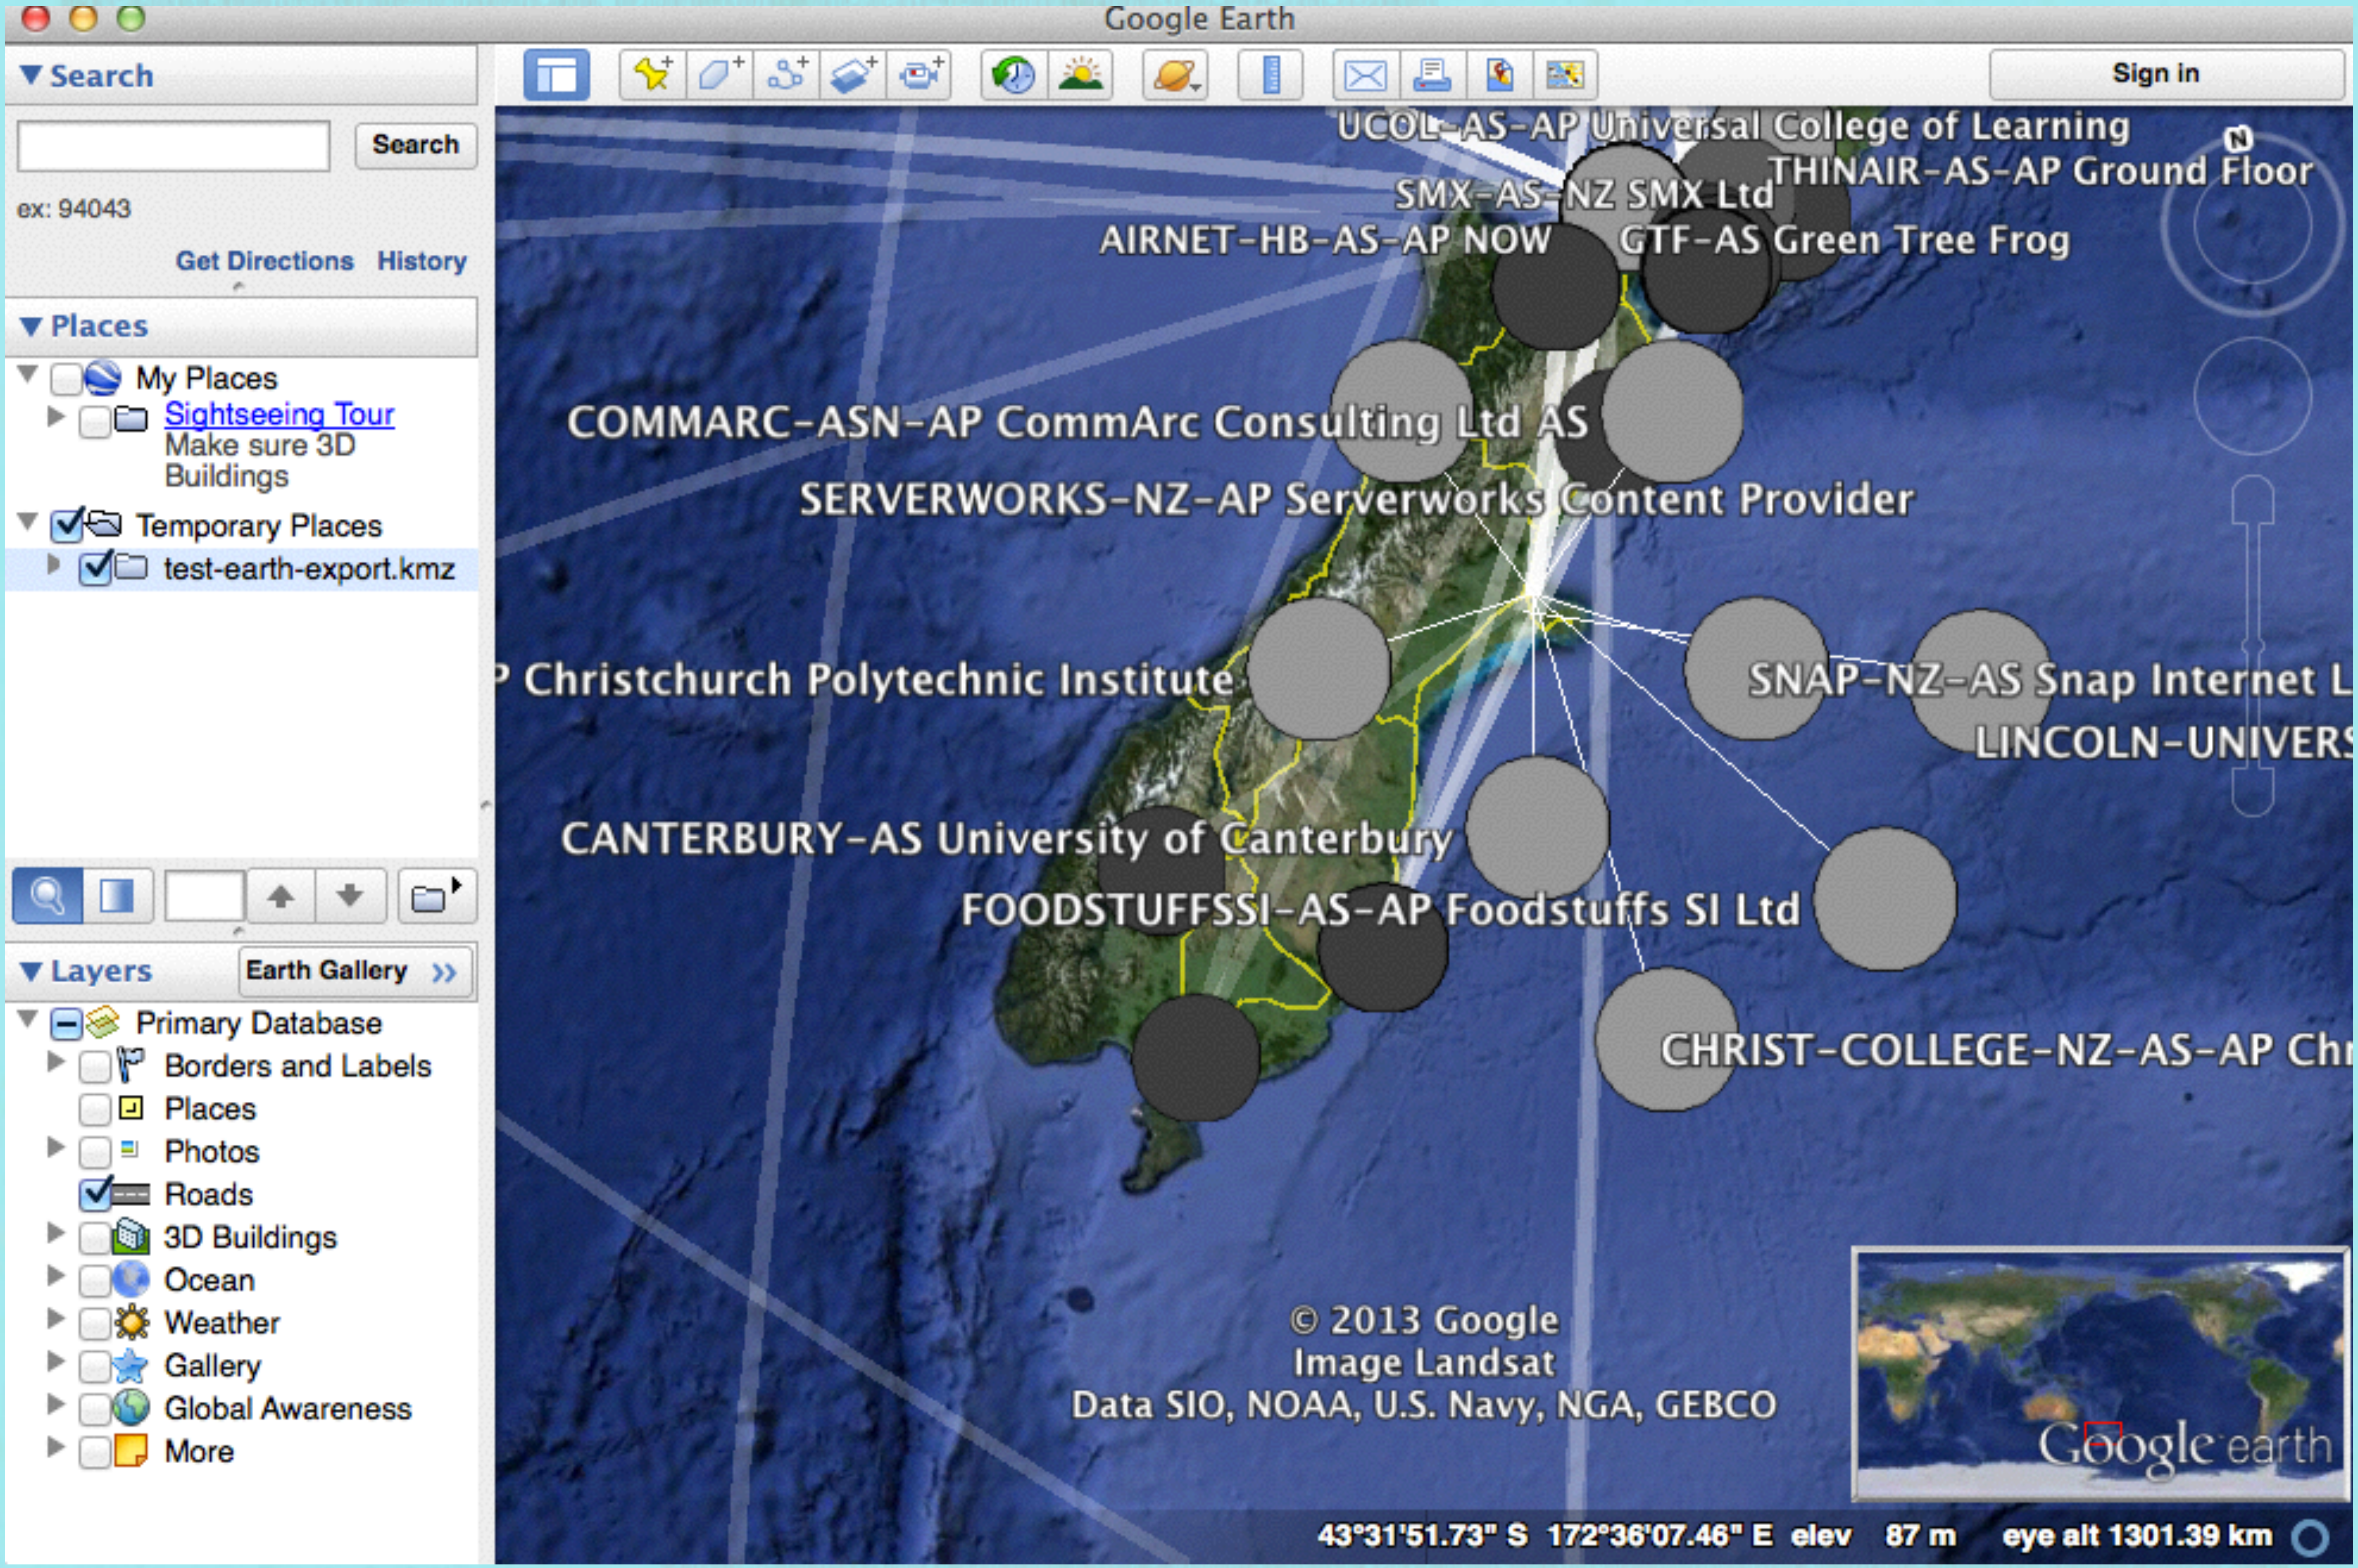
Google Earth
Search
Search
ex: 94043
Get Directions
History
Places
My Places
Sightseeing Tour
Make sure 3D Buildings
Temporary Places
test-earth-export.kmz
Layers
Earth Gallery
Primary Database
Borders and Labels
Places
Photos
Roads
3D Buildings
Ocean
Weather
Gallery
Global Awareness
More
Sign in
UCOL-AS-AP Universal College of Learning
THINAIR-AS-AP Ground Floor
SMX-AS-NZ SMX Ltd
AIRNET-HB-AS-AP NOW
GTF-AS Green Tree Frog
COMMARC-ASN-AP CommArc Consulting Ltd AS
SERVERWORKS-NZ-AP Serverworks Content Provider
P Christchurch Polytechnic Institute
SNAP-NZ-AS Snap Internet L
LINCOLN-UNIVERS
CANTERBURY-AS University of Canterbury
FOODSTUFFSSI-AS-AP Foodstuffs SI Ltd
CHRIST-COLLEGE-NZ-AS-AP Chr
© 2013 Google
Image Landsat
Data SIO, NOAA, U.S. Navy, NGA, GEBCO
Google earth
43°31'51.73" S 172°36'07.46" E elev 87 m eye alt 1301.39 km

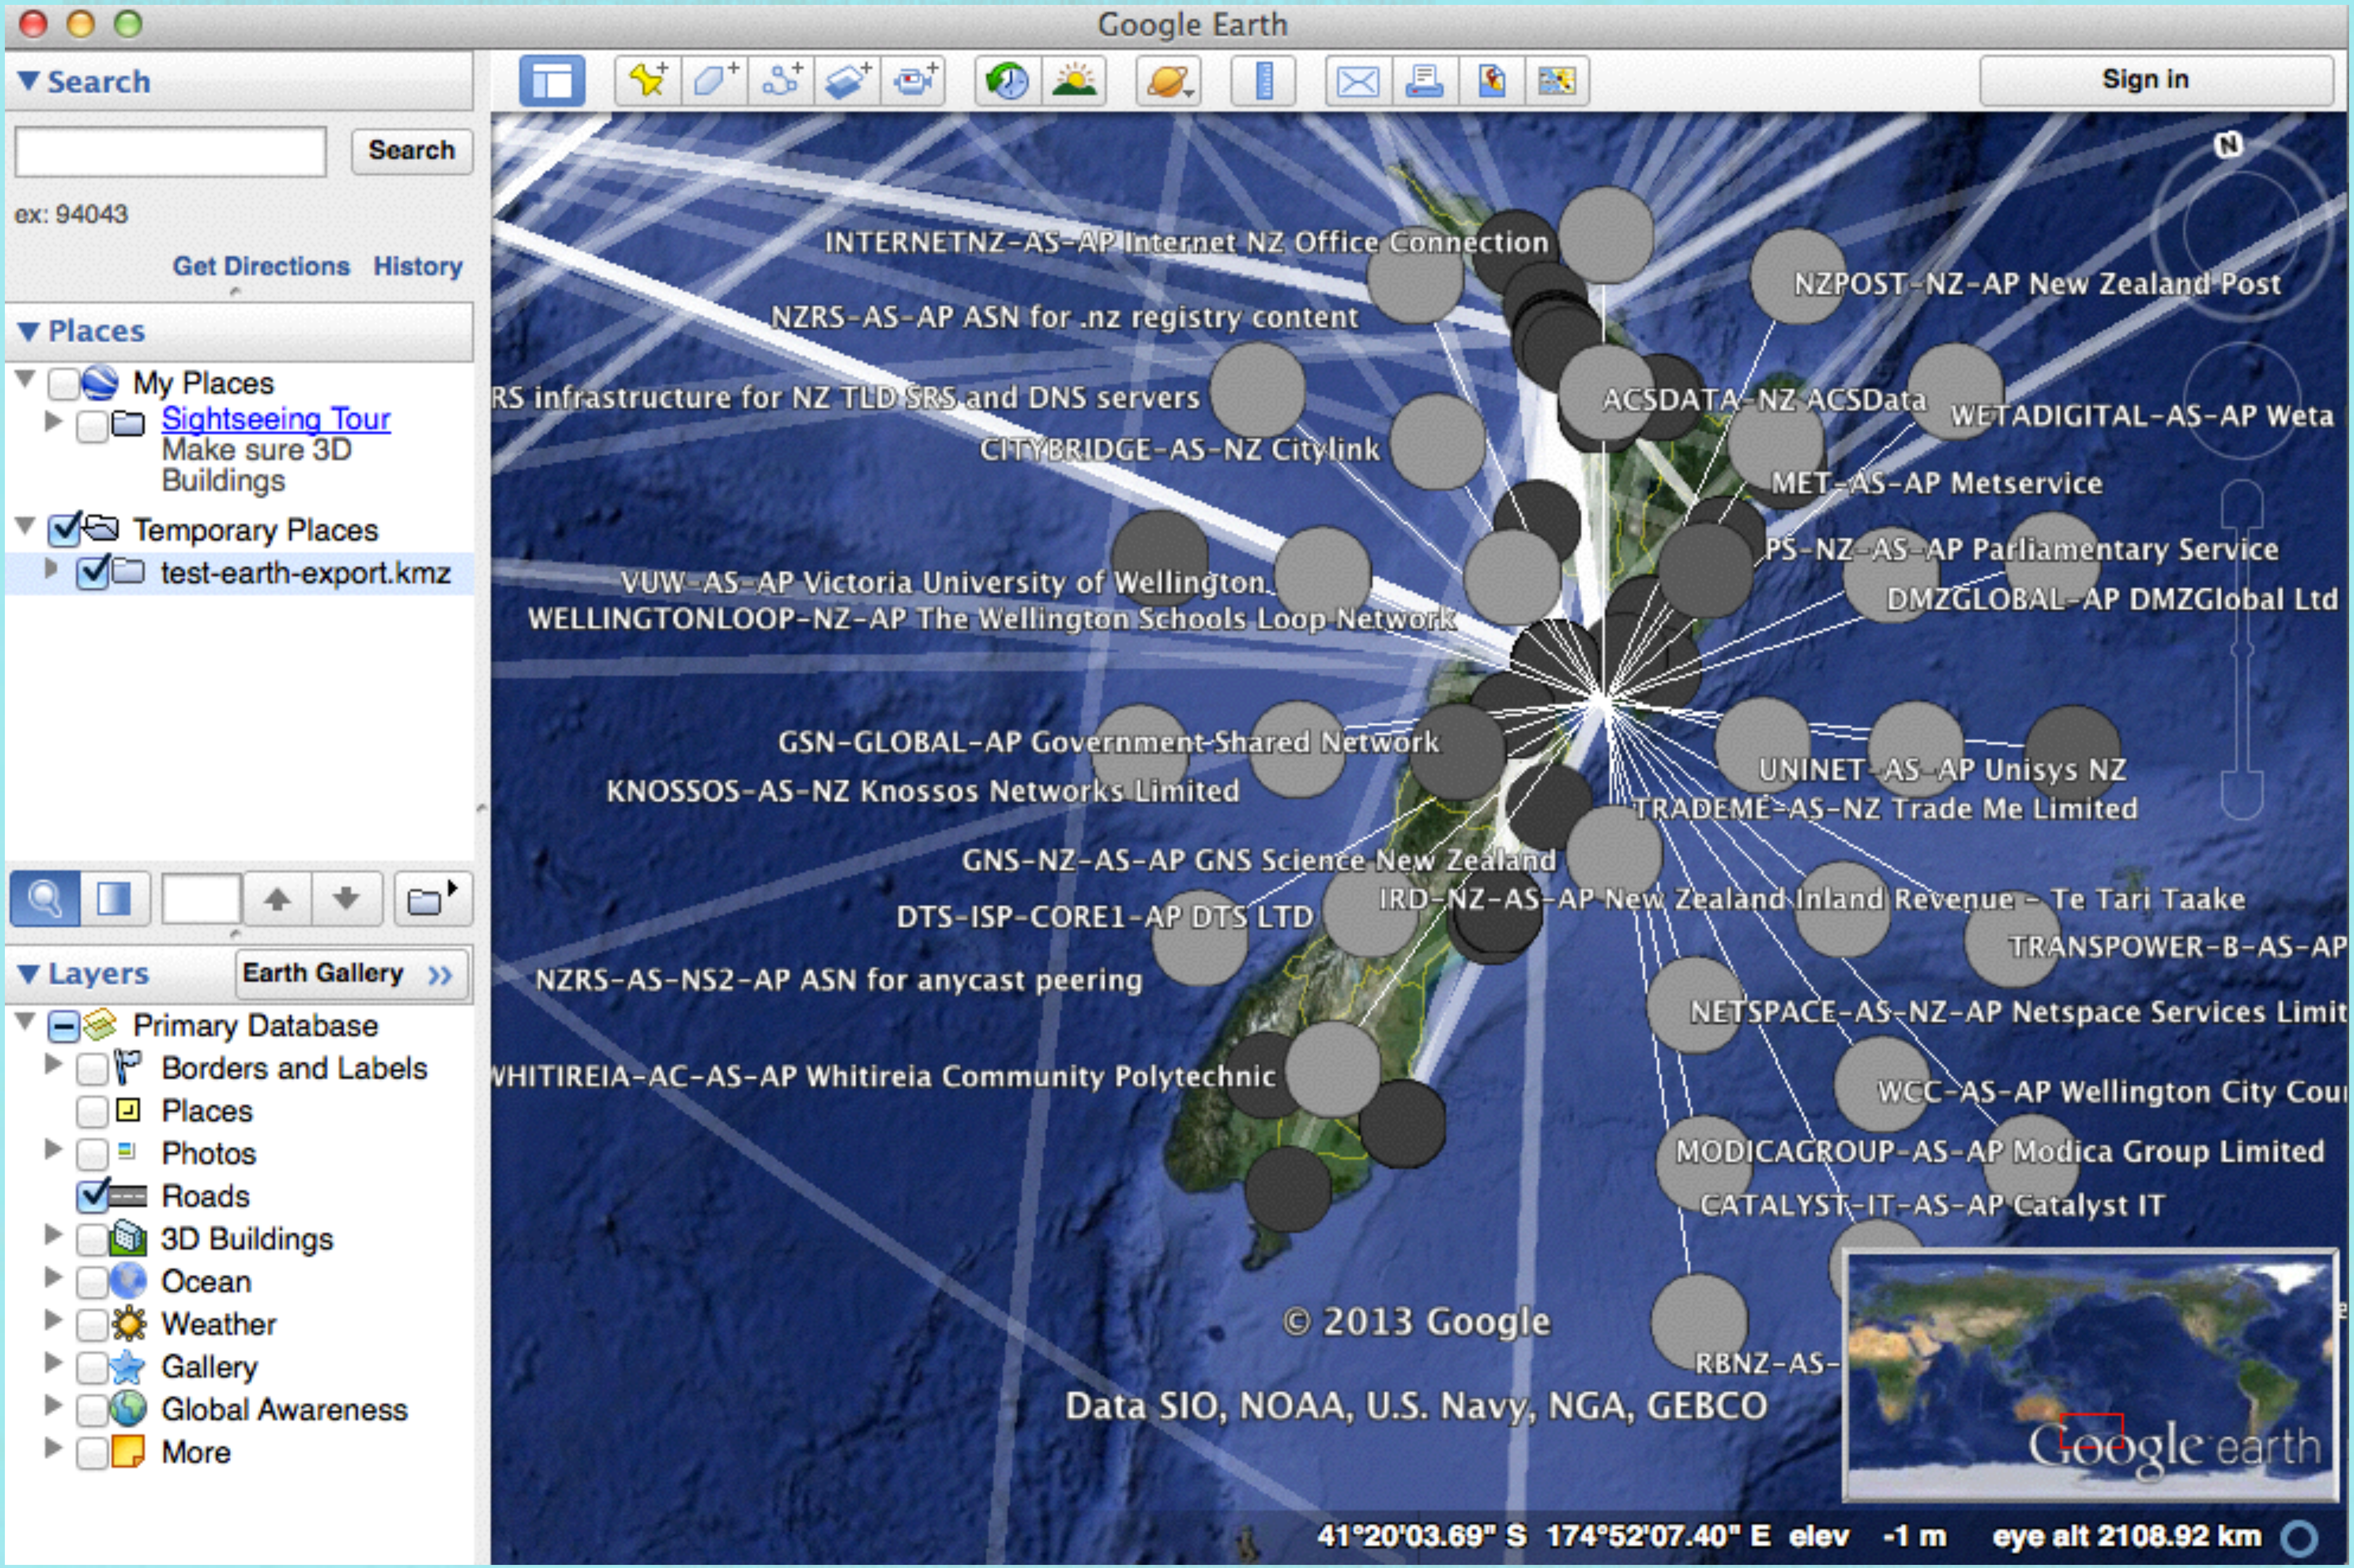
Google Earth
Search
Search
ex: 94043
Get Directions
History
Places
My Places
Sightseeing Tour
Make sure 3D Buildings
Temporary Places
test-earth-export.kmz
Layers
Earth Gallery
Primary Database
Borders and Labels
Places
Photos
Roads
3D Buildings
Ocean
Weather
Gallery
Global Awareness
More
Sign in
INTERNETNZ-AS-AP Internet NZ Office Connection
NZRS-AS-AP ASN for .nz registry content
RS infrastructure for NZ TLD SRS and DNS servers
CITYBRIDGE-AS-NZ Citylink
VUW-AS-AP Victoria University of Wellington
WELLINGTONLOOP-NZ-AP The Wellington Schools Loop Network
GSN-GLOBAL-AP Government Shared Network
KNOSSOS-AS-NZ Knossos Networks Limited
GNS-NZ-AS-AP GNS Science New Zealand
DTS-ISP-CORE1-AP DTS LTD
NZRS-AS-NS2-AP ASN for anycast peering
WHITIREIA-AC-AS-AP Whitireia Community Polytechnic
NZPOST-NZ-AP New Zealand Post
ACSDATA-NZ ACSData
WETADIGITAL-AS-AP Weta
MET-AS-AP Metservice
PS-NZ-AS-AP Parliamentary Service
DMZGLOBAL-AP DMZGlobal Ltd
UNINET-AS-AP Unisys NZ
TRADEME-AS-NZ Trade Me Limited
IRD-NZ-AS-AP New Zealand Inland Revenue - Te Tari Taake
TRANSPOWER-B-AS-AP
NETSPACE-AS-NZ-AP Netspace Services Limit
WCC-AS-AP Wellington City Cou
MODICAGROUP-AS-AP Modica Group Limited
CATALYST-IT-AS-AP Catalyst IT
RBNZ-AS-
© 2013 Google
Data SIO, NOAA, U.S. Navy, NGA, GEBCO
Google earth
41°20'03.69" S 174°52'07.40" E elev -1 m eye alt 2108.92 km
N

OpenOrd

- Found a nice plugin for Gephi called OpenOrd which will do a layout, ignore all the lat/long stuff
- Can use this to end up with an internet tag cloud type thing

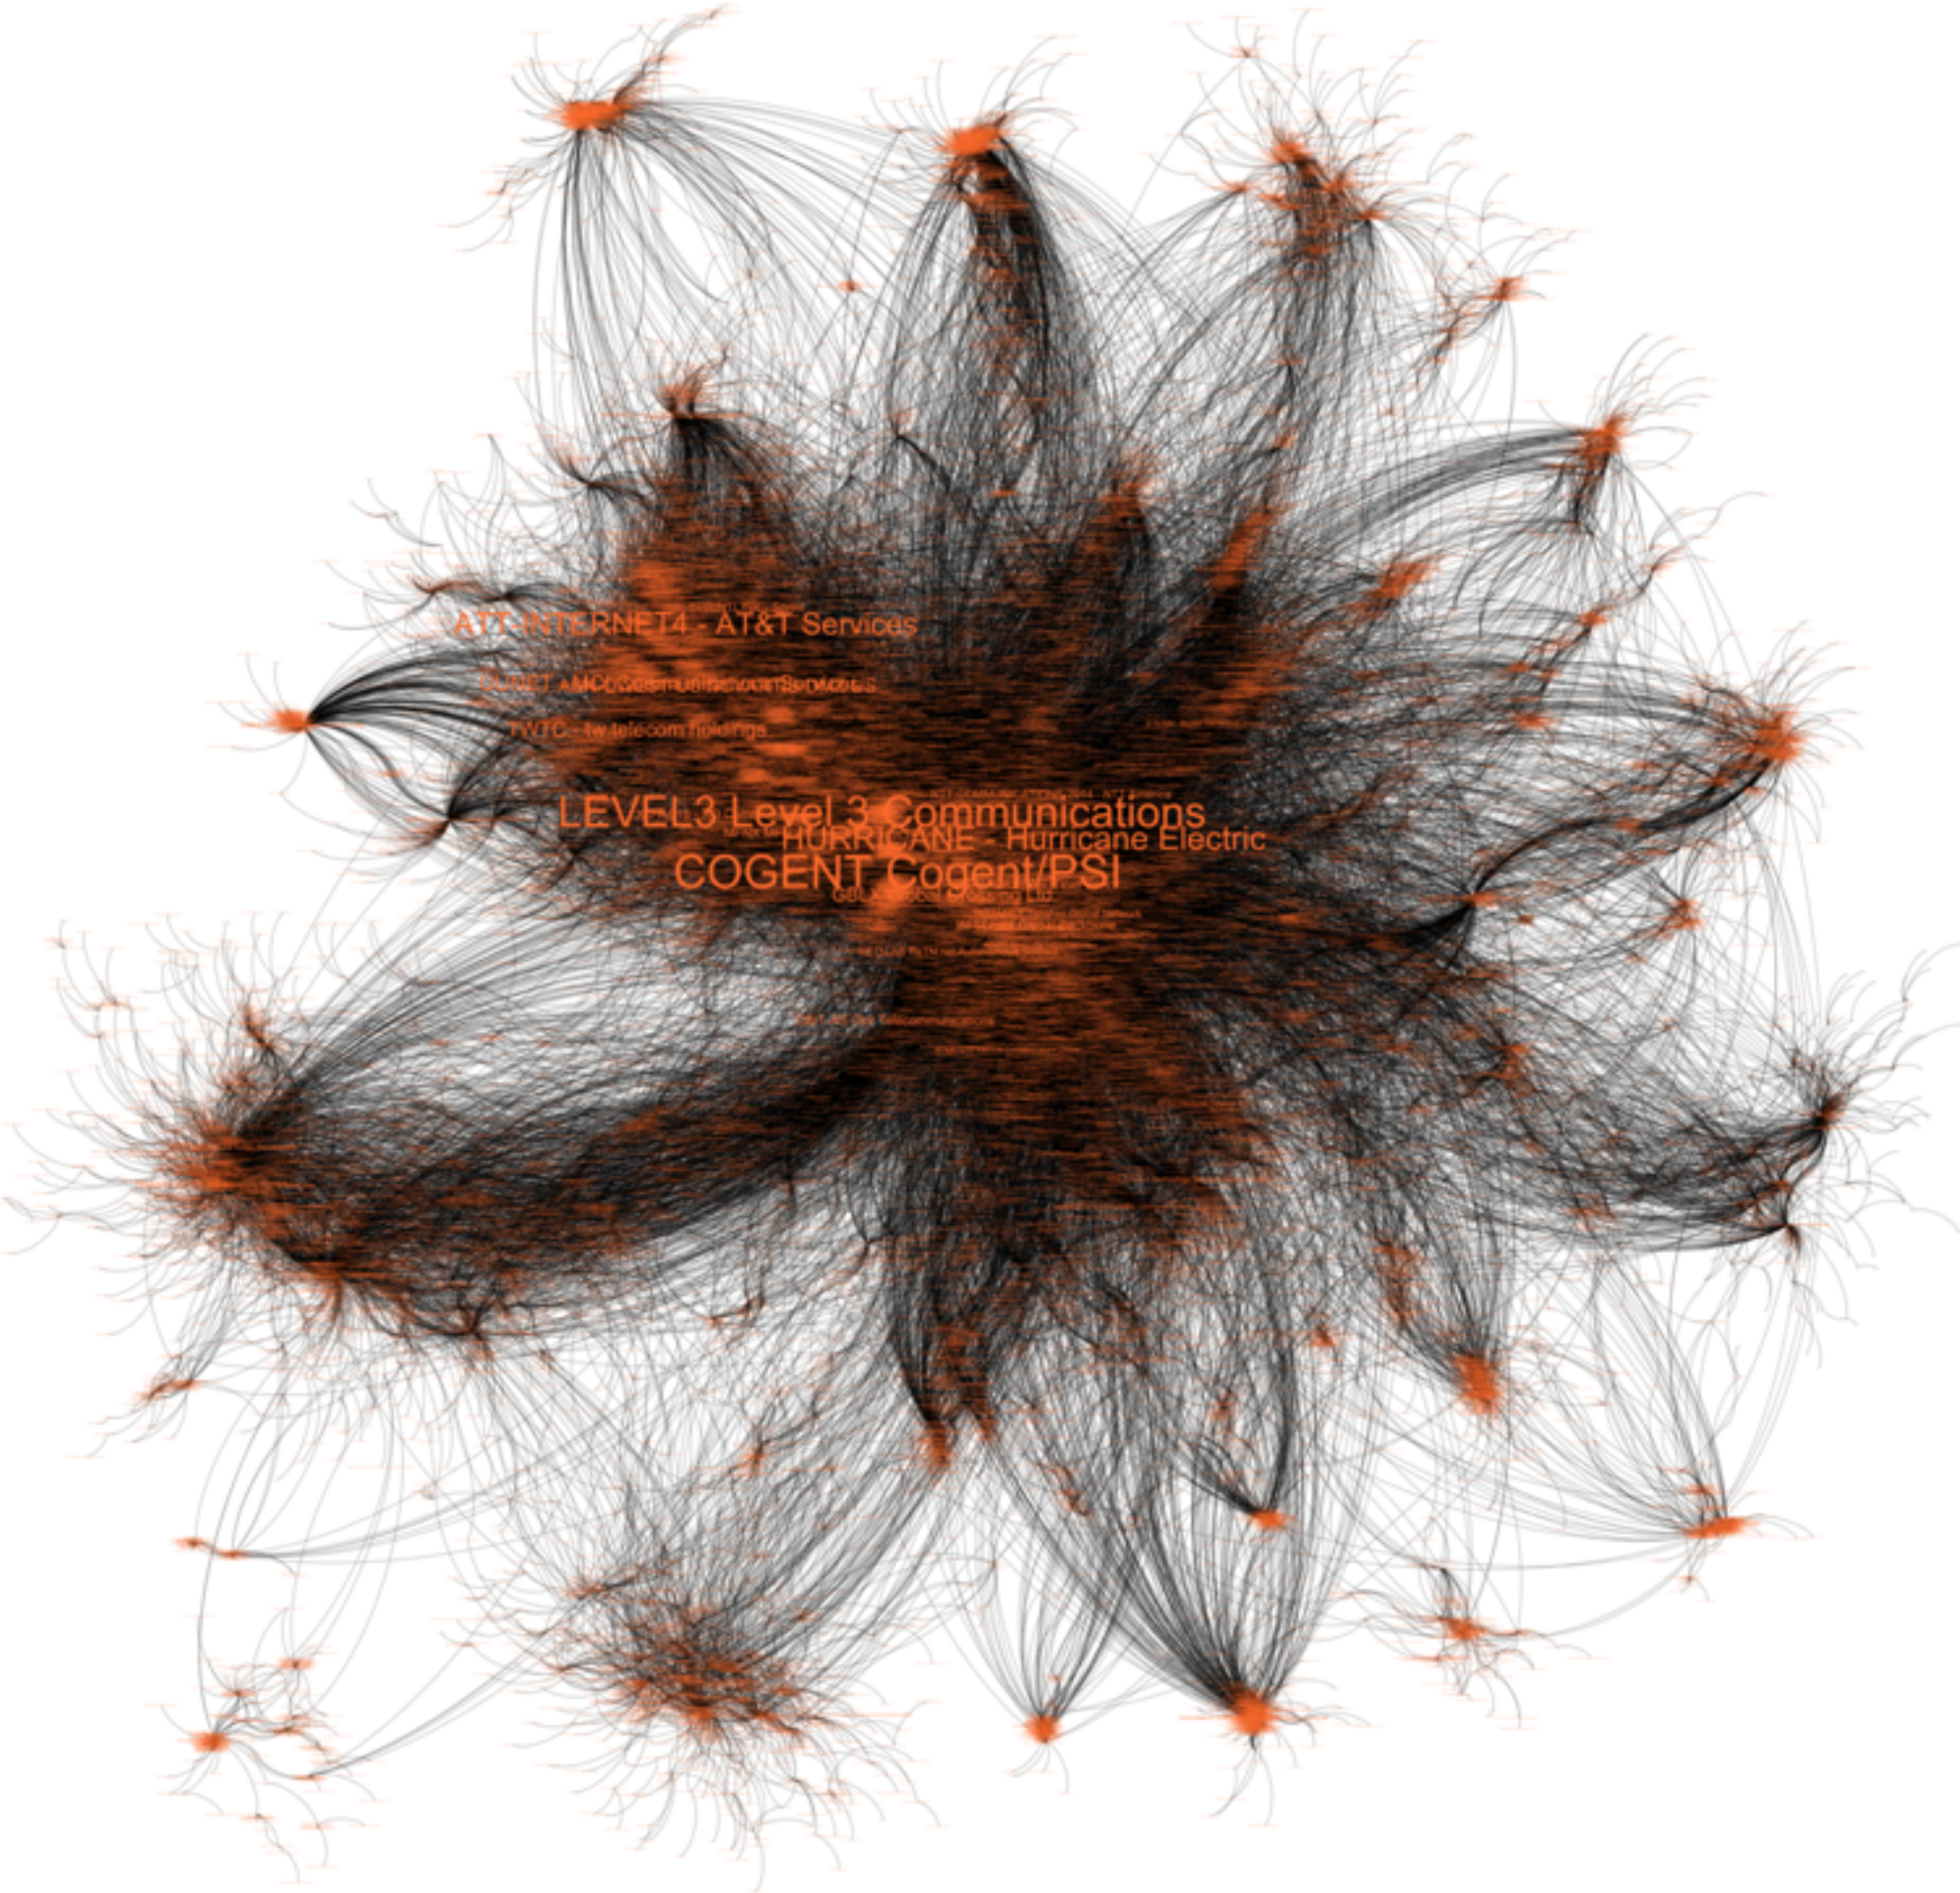
ATT-INTERNET4 - AT&T Services
UUNET
TWTC - tw telecom holdings
NTT-COMMUNICATIONS-2914 - NTT America
LEVEL3 Level 3 Communications
HURRICANE - Hurricane Electric
COGENT Cogent/PSI
GBLX Global Crossing Ltd.

still too big

- Still can’t zoom in on it without making the image gigantic
- Try exporting the giant image to a tile-based system?
- Tried Mapbox. didn’t work.
- Tried some other stuff, job sat processing for 24hrs, still no luck. I think I broke their system.
- Fail again. Got stuck at this point for about 6+ months.
- This map has 45000 nodes, 95000 edges
- Even Gephi seems to struggle a bit, some of the plugins aren’t WEB SCALE

SeaDragon

- Microsoft made a little thing called SeaDragon (yay, Microsoft!)
- Someone wrote an output plugin for this for Gephi
- WIN!
- It's live online right now and you can go and play with it
- http://static.jedi.net.nz/SD/

no internet here

- have a couple of screenshots
- You can zoom in and (kind of) read the names
- Still need to work on readability in the mush, but this is now heaps better
- Can you tell that Level3 and Cogent rule the internet?
- The clusters seem to roughly correspond to geographic areas
- Still tops out at about 12000x10000 before plugin crashes

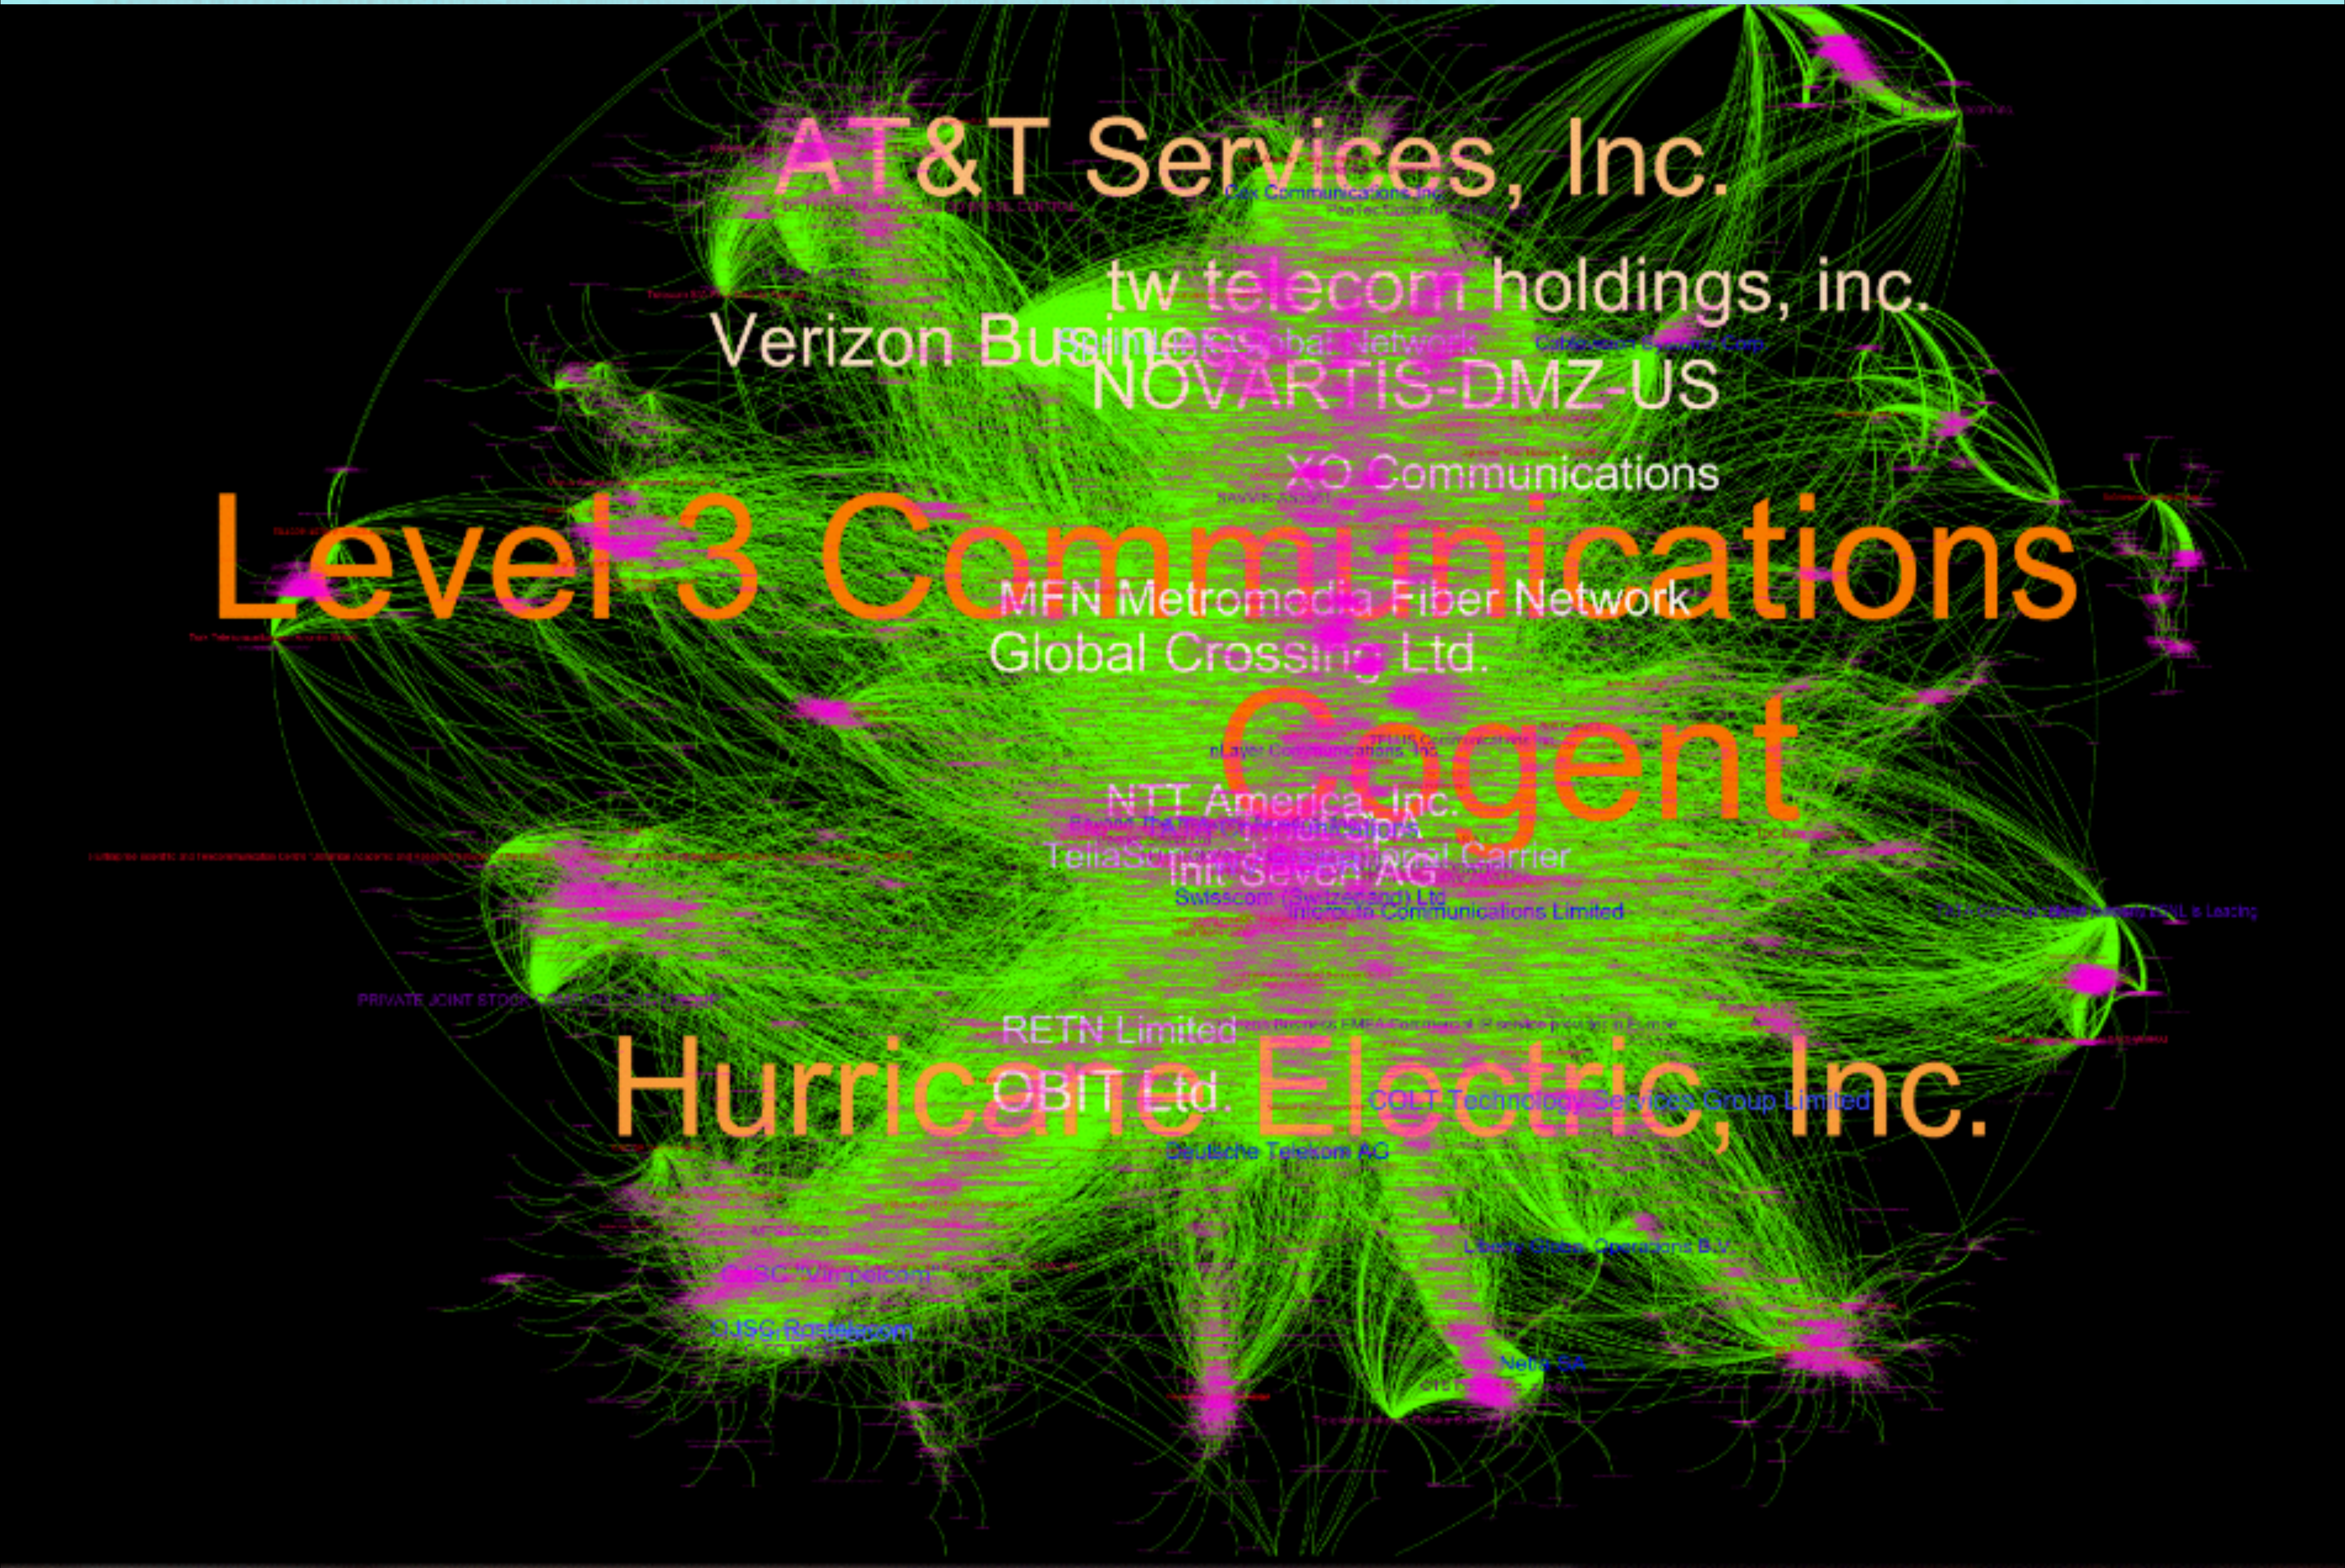
AT&T Services, Inc.
Cox Communications Inc.
tw telecom holdings, inc.
Verizon Business
Sprint Global Network
Cablevision Systems Corp.
NOVARTIS-DMZ-US
XO Communications
Level 3 Communications
MFN Metromedia Fiber Network
Global Crossing Ltd.
Cogent
TELUS Communications Inc.
nLayer Communications, Inc.
NTT America, Inc.
TATA Communications
TeliaSonera International Carrier
Init Seven AG
Swisscom (Switzerland) Ltd
Interoute Communications Limited
PRIVATE JOINT STOCK
RETN Limited
Hurricane Electric, Inc.
OBIT Ltd.
COLT Technology Services Group Limited
Deutsche Telekom AG
OJSC "Vimpelcom"
OJSC Rostelecom
Liberty Global Operations B.V.
Netia SA

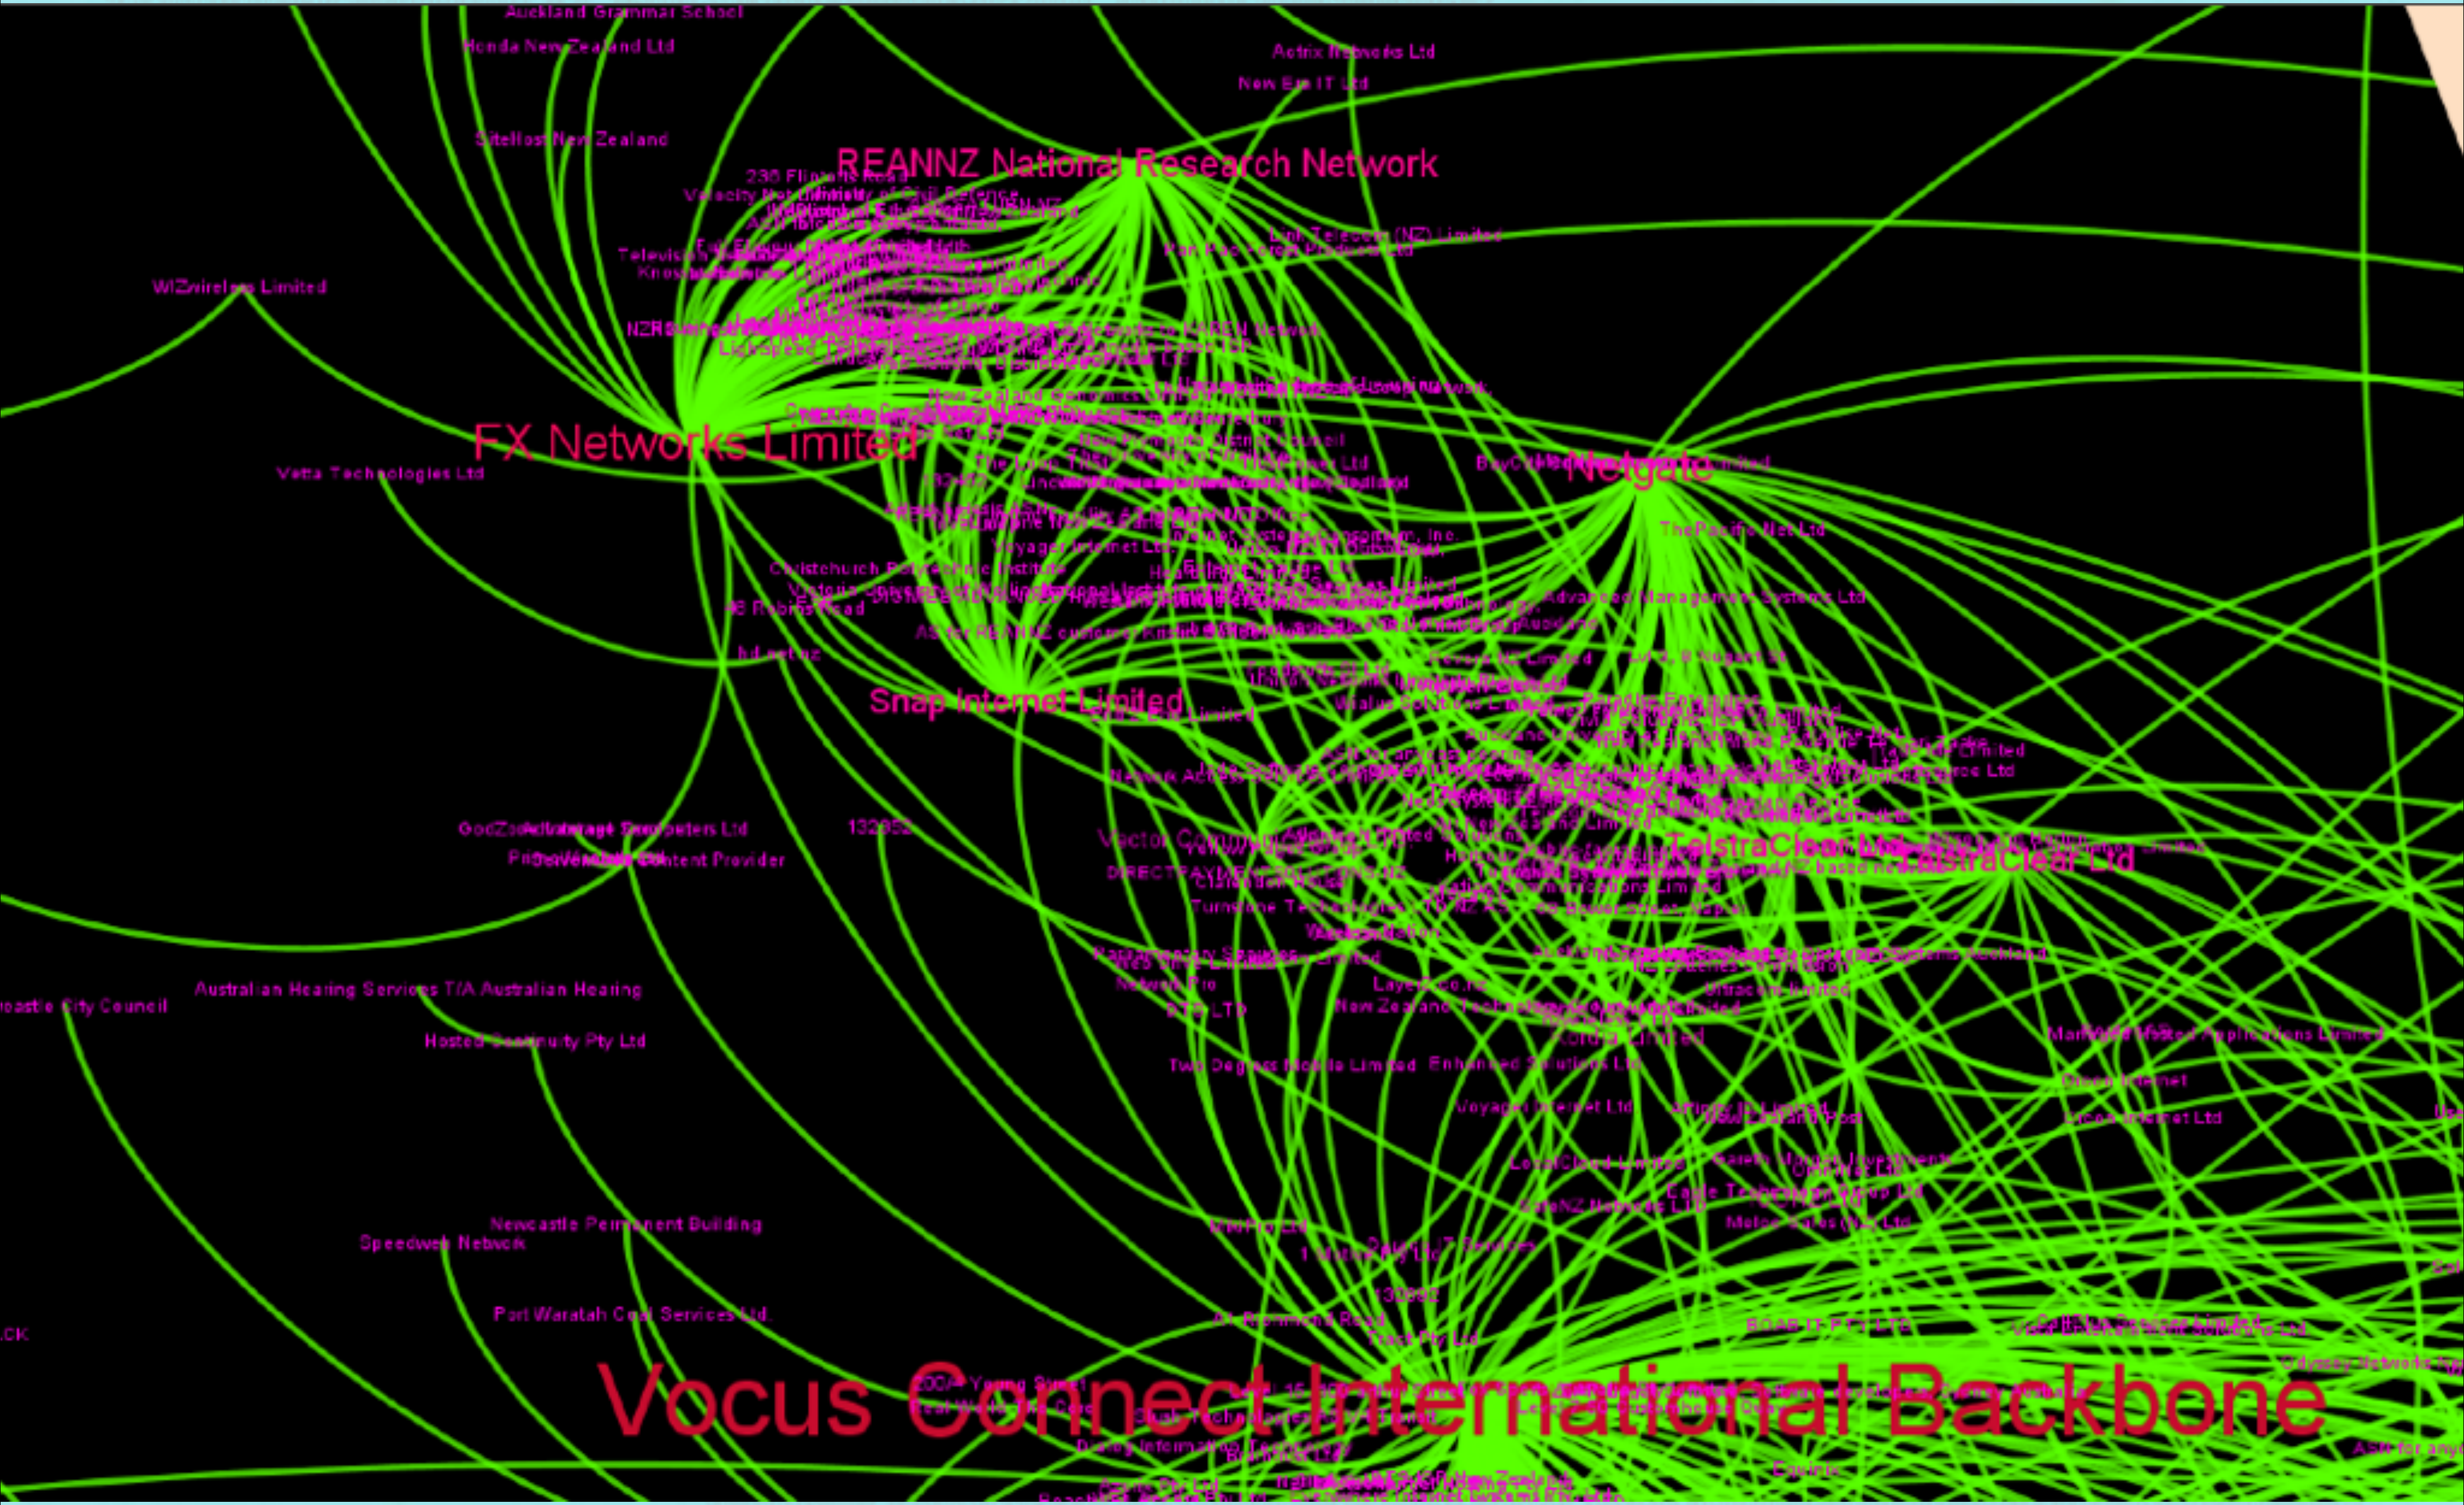

REANNZ National Research Network
FX Networks Limited
Netgate
Snap Internet Limited
TelstraClear Ltd
Vocus Connect International Backbone
Auckland Grammar School
Honda New Zealand Ltd
SiteHost New Zealand
WIZwireless Limited
Vetta Technologies Ltd
Actrix Networks Ltd
New Era IT Ltd
236 Flinders Road
Television New Zealand
NZRS
Lightspeed
Link Telecom (NZ) Limited
Pan Pac Forest Products Ltd
New Zealand Genomics Limited
The Loop Trust
Voyager Internet Ltd
Christchurch Polytechnic Institute
48 Robins Road
hd.net.nz
AS for REANNZ customer
Advanced Management Systems Ltd
ThePacific.Net Ltd
BayCity
Revera NZ Limited
Lvl 2, 8 Nugent St
Vialus Solutions Limited
Auckland University of Technology
ASN for anycast peering
Network Access
Vector Communications
DIRECTPAYMENTSOLUTIONS-NZ
Turnstone Technologies LTD NZ AS
Parliamentary Service
Network Pro
DTS LTD
Two Degrees Mobile Limited
Enhanced Solutions Ltd
Layer2.co.nz
New Zealand Technology
Kordia Limited
Ultracom Limited
Managed Hosted Applications Limited
Orcon Internet
Orcon Internet Ltd
Voyager Internet Ltd
New Zealand Post
LocalCloud Limited
Gareth Morgan Investments
Cable Technologies Group Ltd
DataNZ Networks LTD
Melco Sales (NZ) Ltd
WinPro Ltd
132952
132882
41 Richmond Road
Tract Pty Ltd
BOAB IT PTY LTD
Odyssey Networks
ASN for anycast
Equinix
GodZone
Advantage Computers Ltd
Content Provider
132952
Australian Hearing Services T/A Australian Hearing
Newcastle City Council
Hosted Continuity Pty Ltd
Newcastle Permanent Building
Speedweb Network
Port Waratah Coal Services Ltd.
200A Young Street
Real World The Core
Dialog Information Technology

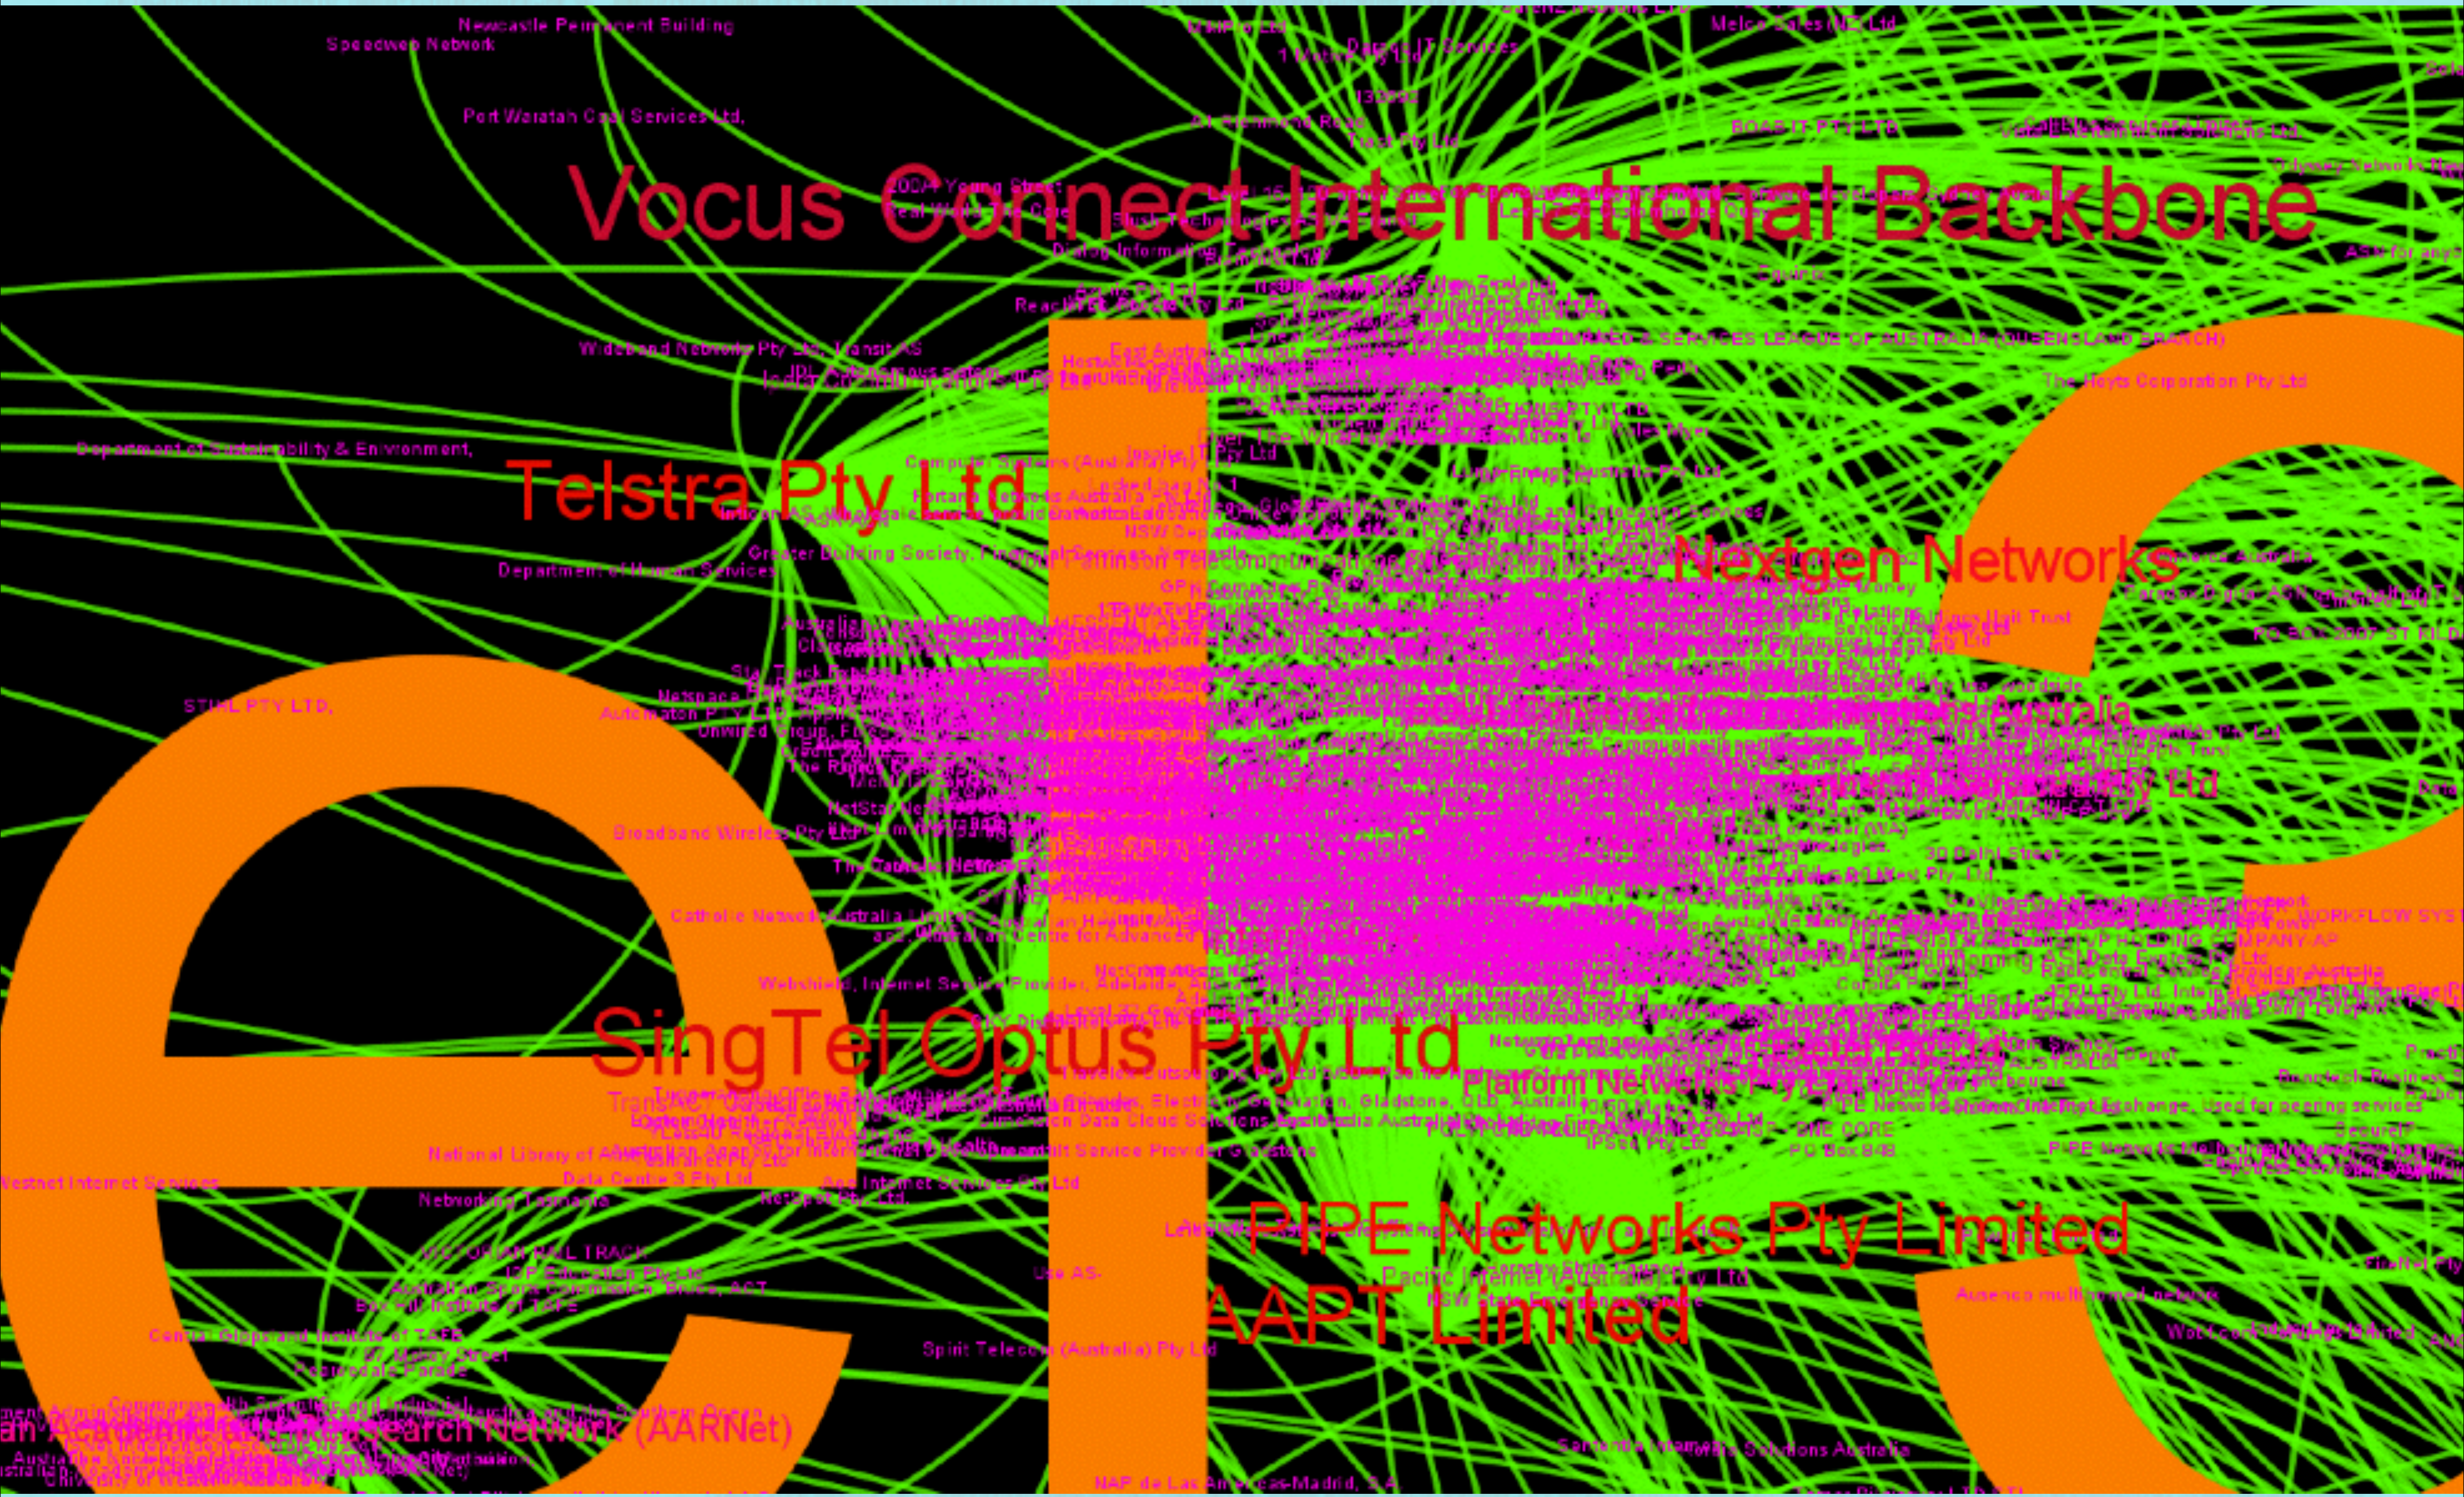
Vocus Connect International Backbone
Telstra Pty Ltd
Nextgen Networks
SingTel Optus Pty Ltd
PIPE Networks Pty Limited
AAPT Limited
Platform Networks Pty Ltd
Pacific Internet (Australia) Pty Ltd
an Academic and Research Network (AARNet)
Newcastle Permanent Building
Speedweb Network
Port Waratah Coal Services Ltd.
Melco Sales (NZ) Ltd
Widebend Network Pty Ltd, Transit AS
Department of Sustainability & Environment,
Department of Human Services
Greater Building Society, Financial Services, Newcastle
STUPL PTY LTD.
Broadband Wireless Pty Ltd
Catholic Network Australia Limited
Webshield, Internet Service Provider, Adelaide, Australia
National Library of Australia
Data Centre 3 Pty Ltd
Networking Tasmania
Central Gippsland Institute of TAFE
Spirit Telecom (Australia) Pty Ltd
NAP de Las Americas-Madrid, S.A.
The Hoyts Corporation Pty Ltd
RETURNED & SERVICES LEAGUE OF AUSTRALIA (QUEENSLAND BRANCH)
Computer Systems (Australia) Pty Ltd
30 Dalhi Street
PO Box 848
Ausenco multihomed network
Spirit Telecom

In Summary

- Don’t buy lots of disk and expect to be able to graph it easily
- Don’t buy weird equipment
- Don’t try and do weird shit with data or even normal stuff with SNMP without being prepared to fail a LOT
- Don’t pretend you can sysadmin, hire a real one if you don’t know what you’re doing :)
- Pictures of cats

Bonus Slide - Image host fail

- most image hosts will let you hide zip files in your jpegs
- just cat file.zip >> file.jpg
- unzip file.jpg works fine
- look at image in image software, works fine
- this is not new, and exploits “features” of each file format
- They could fix this but don′t seem to be interested, or don’t think they need to - simple image reprocessing rather than allow the raw image to be downloaded would fix it
- In the mean time, enjoy using flickr’s free 1TB of storage for your lolcat collection

- Thank you!
- twitter.com/trogs