OPEN SOURCE ISP / BNG

Are we there yet?

Brody Stockel - AusNOG 2022

WHY ARE WE HERE?

LIKE AS IN, SITTING HERE, IN THOSE CHAIRS, RIGHT NOW.

THE PUB IS VERY ENTICING RIGHT NOW

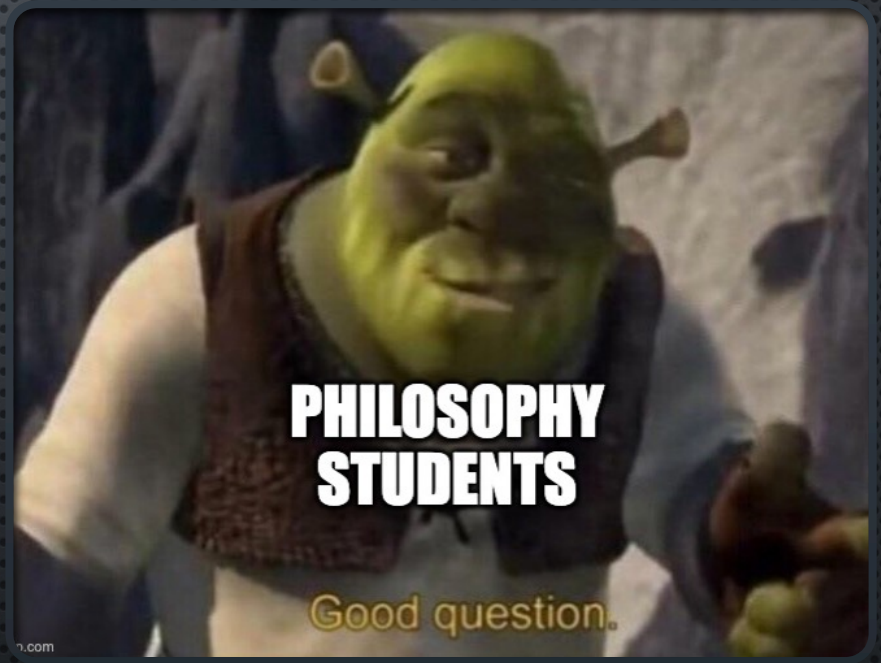

WHO IS THIS MOONBAT?

- WHO: BORDY STOCKTEL
- WHAT:
  - NETWORK ENGINEER @ INTERPHONE
  - ONLINE SH*TPOSTER
  - EDITOR IN CHIEF @ TELCO.NEWS
- WHERE:
  - TWITTER.COM/BRODYSTOCKEL
  - LINKEDIN.COM/IN/BRODYSTOCKEL

- FIRST TIME SPEAKER, PLEASE BE NICE.

- THE FOLLOWING SH*TPOSTS ARE NOT REFLECTIVE OF MY EMPLOYERS VIEWS.

THE CURRENT STATE OF PLAY

Juniper MX204

Juniper MX80

Cisco ASR9901

Nokia 7750-SR

Hot Garbage

THE CURRENT STATE OF PLAY

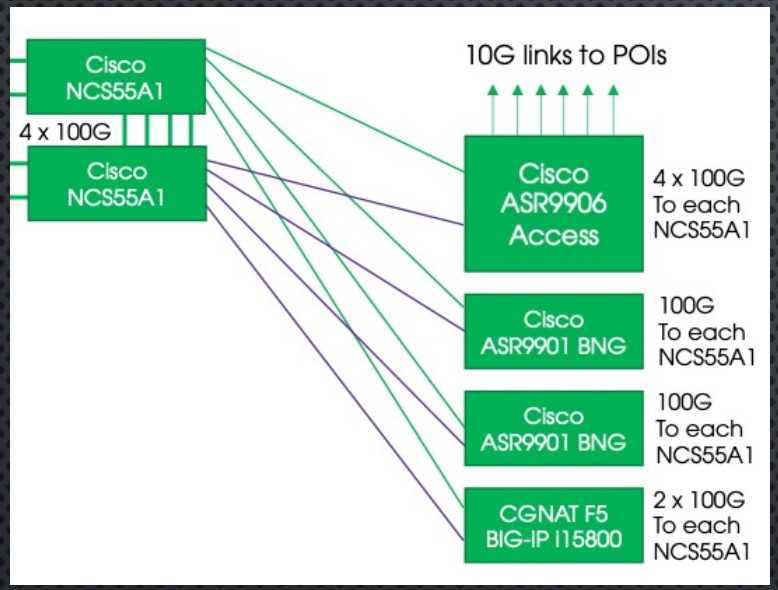

src: AussieBroadband AusNOG 2019

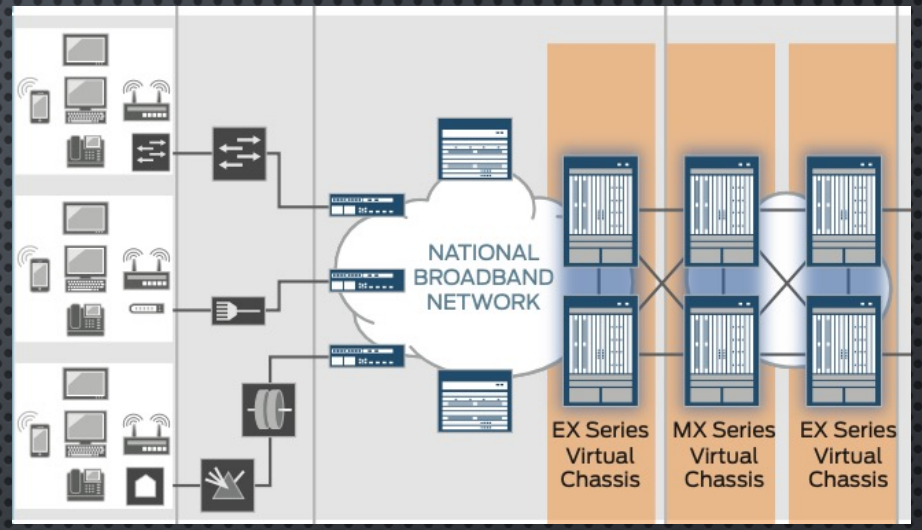

src: Juniper Broadband Edge Solution 2012

> Engineering team of small NBN RSP very glad poi cabinets unlabelled as they install their third Mikrotik BNG

Brody Stockel - AusNOG 2022

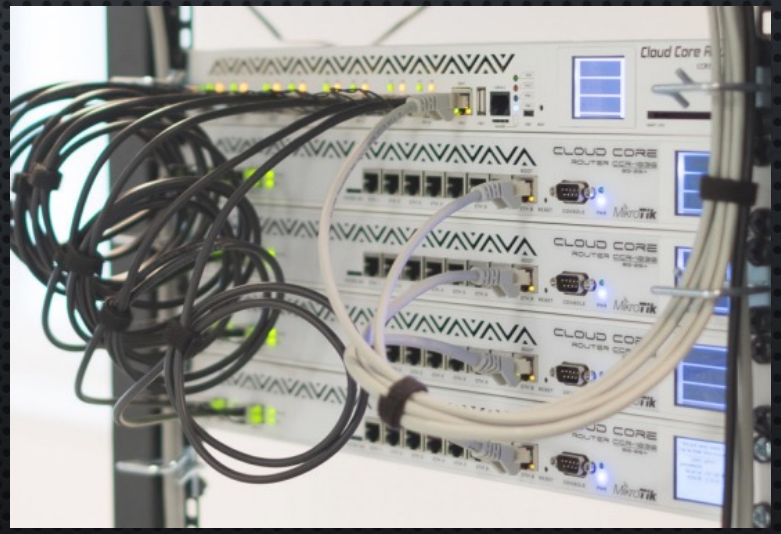

SOMETIMES A LITTLE HARDER TO GET THESE DAYS

CCR1072-1G-8S+: 72 core CPU Cloud Router with Dual Power supply

Brand: Mikrotik
Product Code: CCR1072-1G-8S+

Pricing

Qty	1	5	10	20
Inc.	$5,225.00	$4,703.05	$4,546.47	$4,442.08
Ex.	$4,750.00	$4,275.50	$4,133.15	$4,038.25

Availability:	No/Low Stock!
Further units due:	2022-09-16

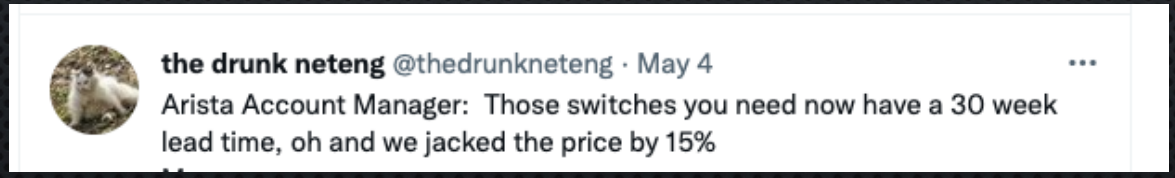

FEATURE SETS

- Deep queues
- Dynamic VLAN detection
- Subscriber DEMUXing
- Dynamic Profiles
- DHCP Server / Relay
- PPPoE Concentrator
- Some sort of VRF support

YEAR OF THE LINUX ~~DESKTOP~~ ROUTER

Or as I have taken to calling it, GNU plus router

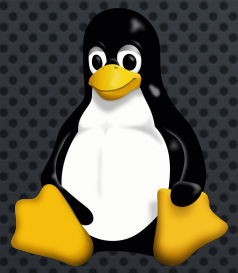

GNU/Linux

- Flexible operating system to go anywhere
- Runs on anything you can find
- Every coding language you can think of

- Free

FRRouting

- "Industry Standard" CLI
- Full BGP / EVPN support
- VRFs for days

- Free

Accel-PPP

- Userland IPoE / PPPoE server
- (Basic) Demux support
- Built in shaping support
- Some English documentation
  …… its mostly in Russian however

- Free

HOLD ON, ACCEL-WHAT?

- Userland IPoE / PPPoE server
- Kernel modules for VLAN auto-discovery
- Able to authenticate IPv4 DHCP sessions via Option82
- Able to manage Linux tc shapers
- Supports DHCPv6 as well (kinda..)
- Semi-complete documentation, Forums are mostly in Russian
- Included with VyOS from 2019 (config implementation is very basic)

WHAT ABOUT THE ACCESS-NNI SIDE?

VXLAN-EVPN to the rescue!

- Able to remove the need for BNGs adjacent to the NNI
- Put Access nodes in POIs, haul back via a metro network to the DC
- Avoid flooding BUM traffic, EVPN Type 3 routes to "subscribe" a BNG to an NNI
- EVPN Type 2 routes to handle MAC learning

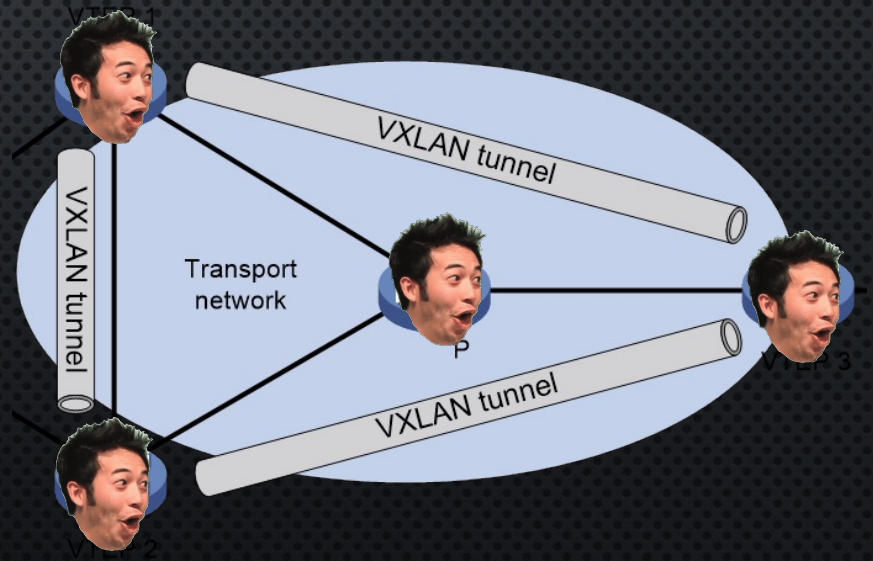

CGNAT

- A sad reality of life
- CPEs with Dual Stack? Ha.
- Masquerade? Too much DR pain.
- Port block assignment? Mebe.

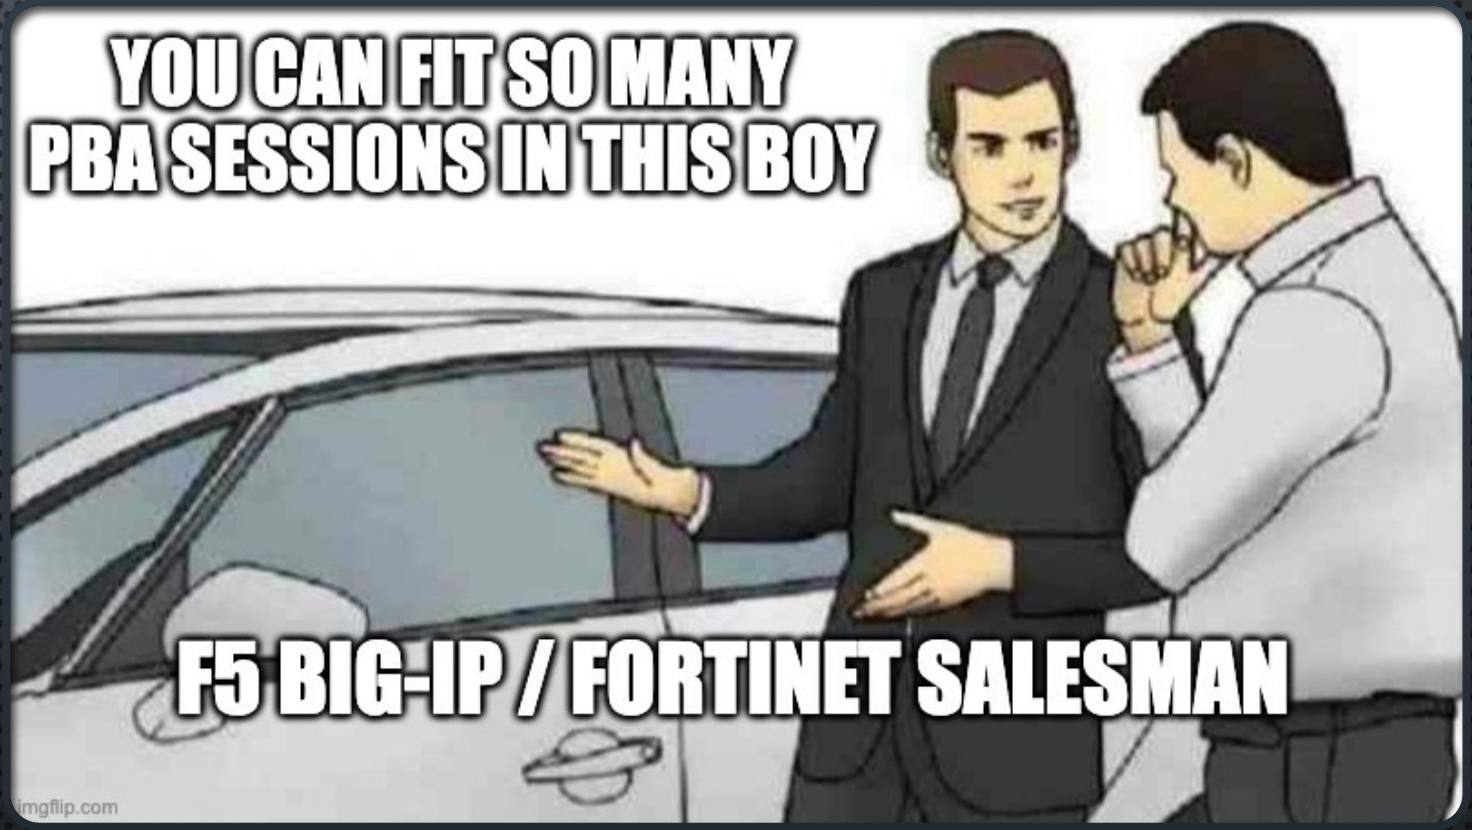

ENTER STAGE RIGHT: DANOS

Forked from the codebase of AT&T dNOS, formally Vyatta vRouter

- Uses FFRrouting under the hood for BGP routing
- Includes a DPDK forwarding dataplane
- Similar config stanza to Junos / VyOS
- Has CGNAT-PBA support
- Free

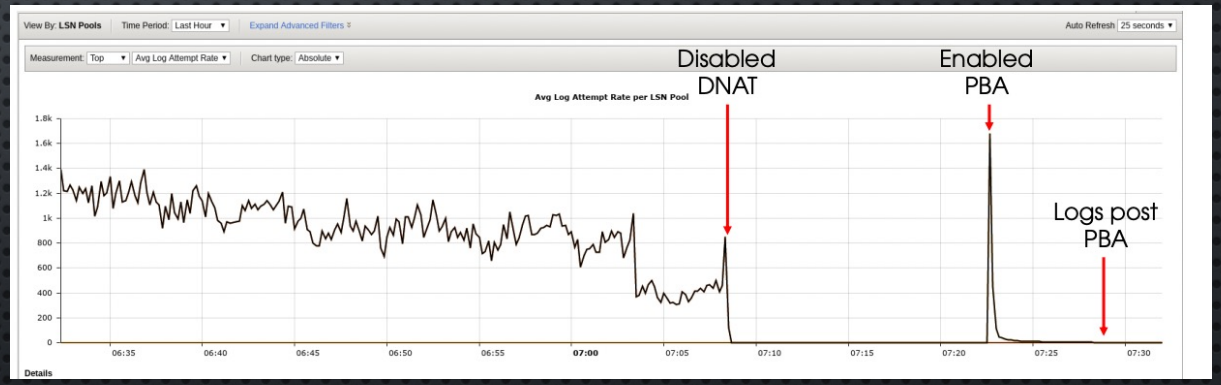

src: AussieBroadband AusNOG 2019

```
/ip firewall nat add chain=srcnat out-interface=nbn-tc4-transit \
  src-address=100.64.0.0/10 action=masquerade log=yes log-prefix=CGNAT
```

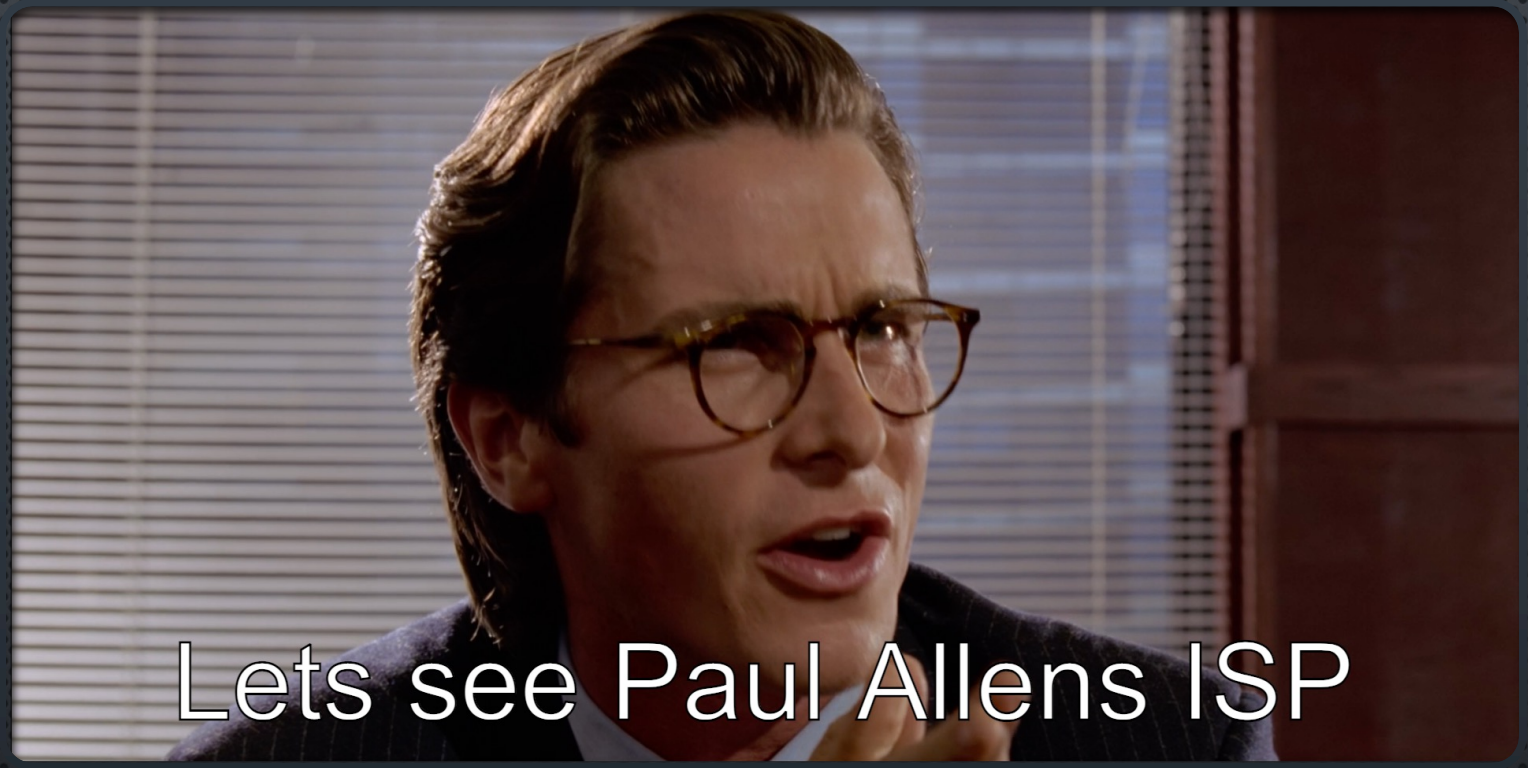

WHAT IT WOULD LOOK LIKE

BRINGING IT ALL TOGETHER

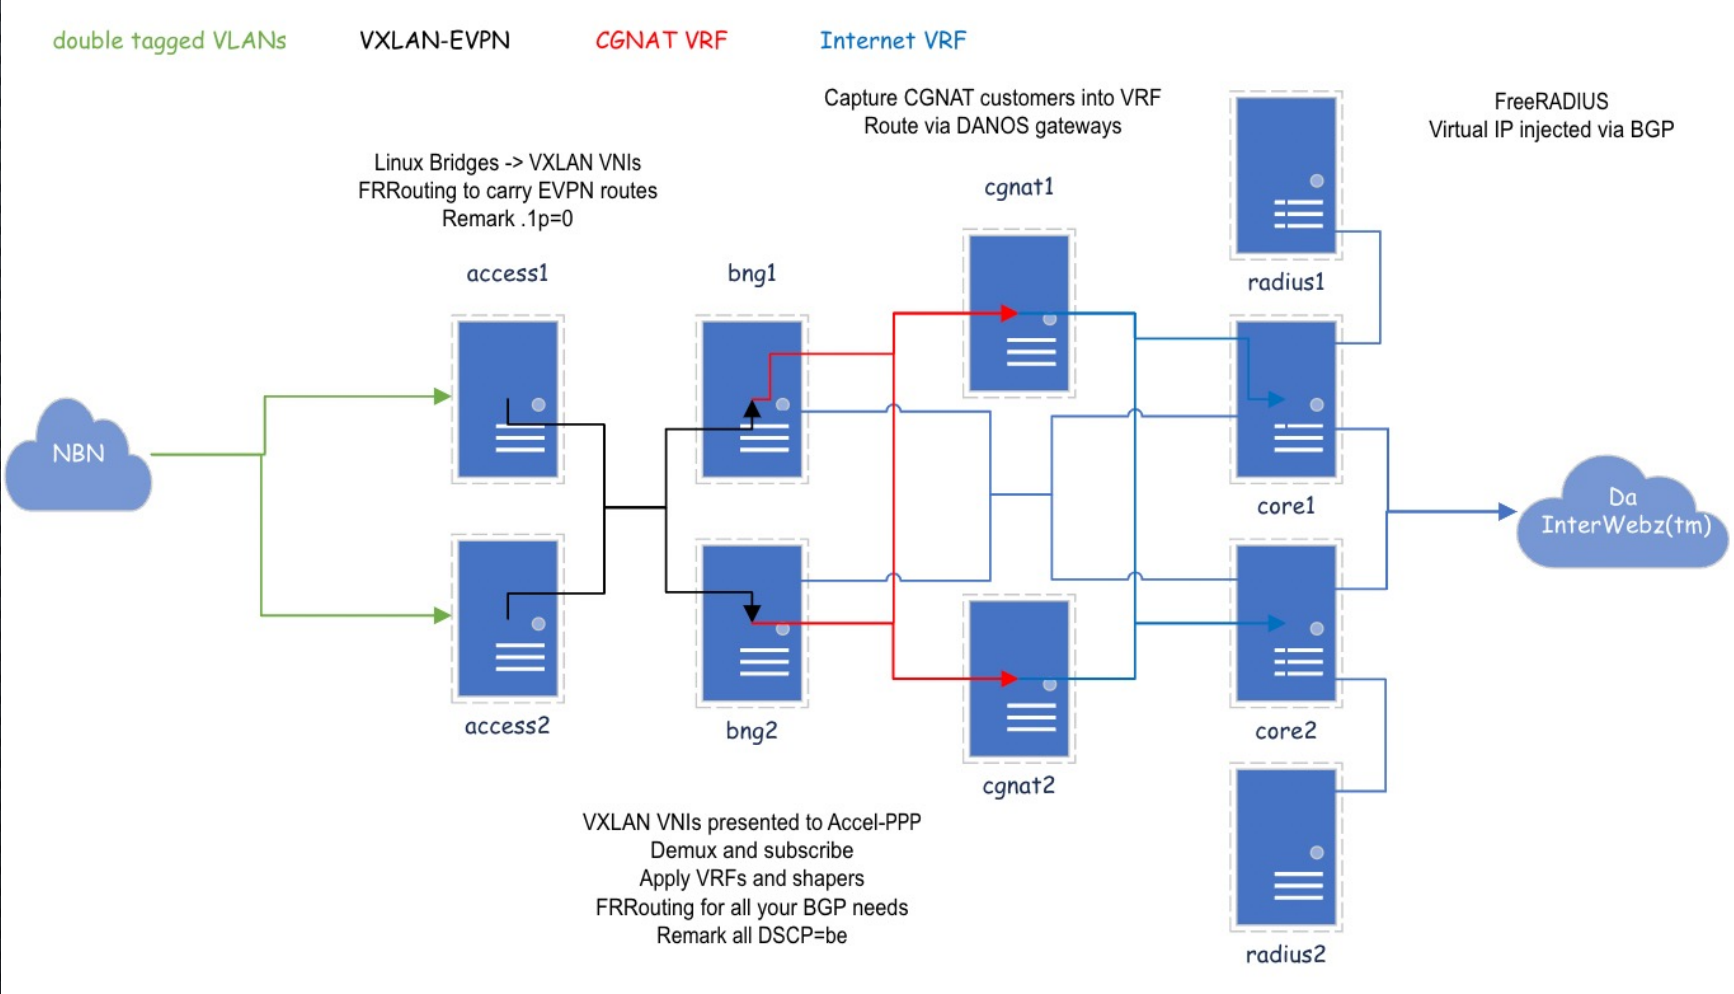

THAT SEEMS A LITTLE TOO EASY?

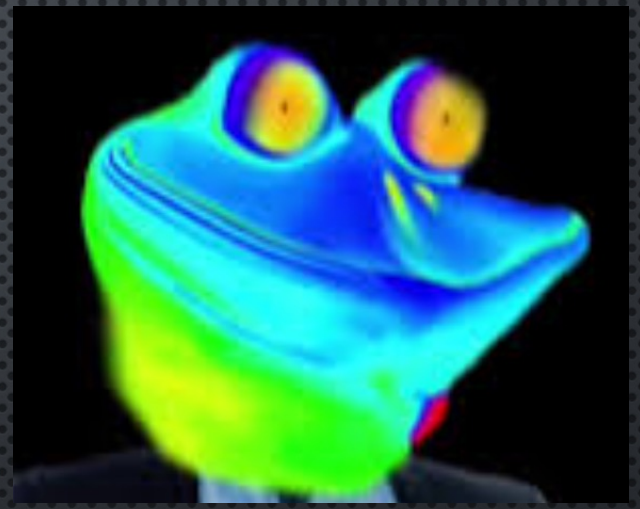

- SURELY IT COULDN'T BE THAT EASY?
- ISPS / VENDORS SPEND YEARS GETTING THIS STUFF TO MARKET
- AND YEARS MORE BEFORE ITS STABLE….

GOING OFF THE DEEP END

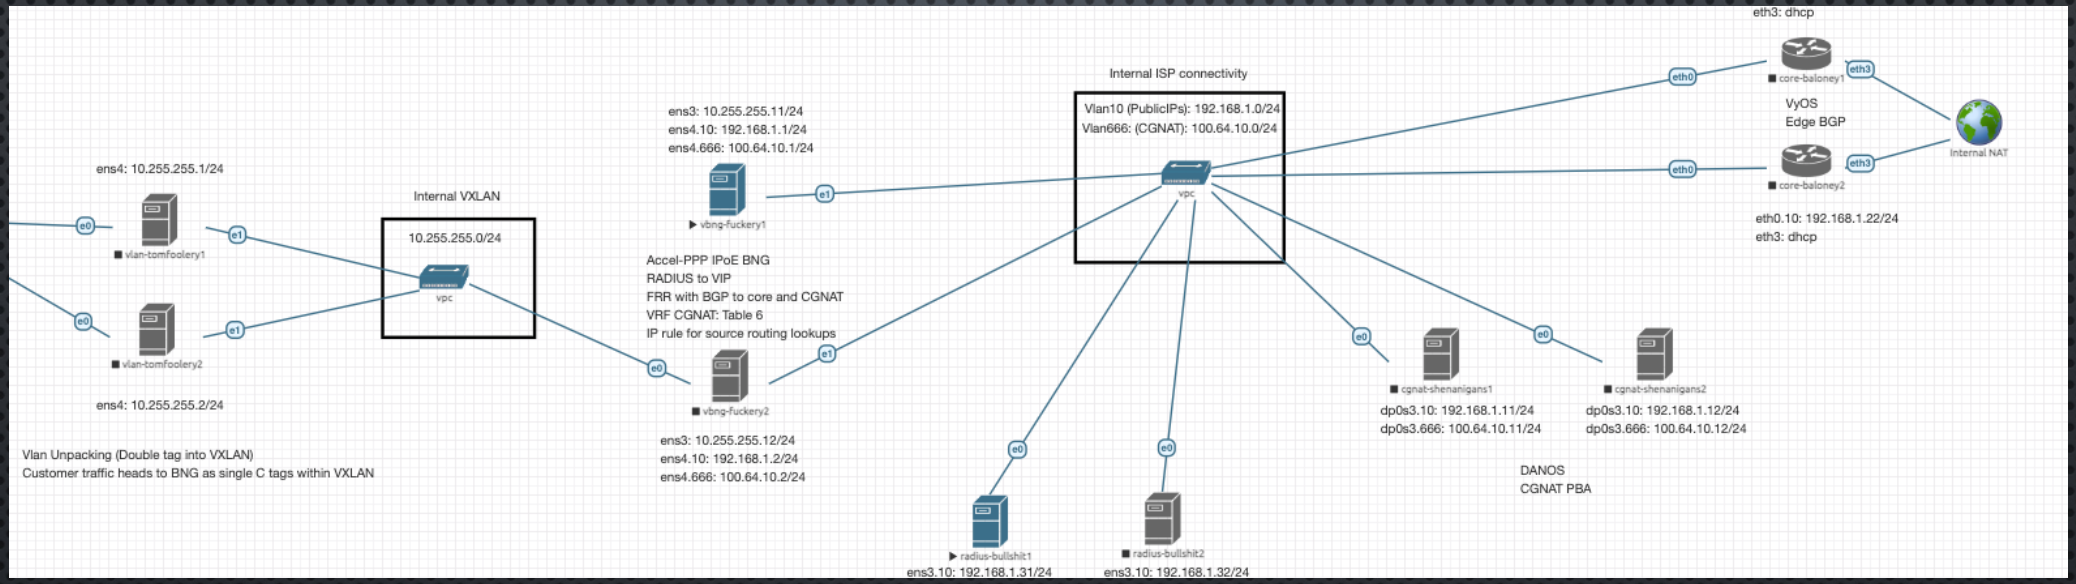

Brazilian guys doing CGNAT: bit.ly/3c5EZmN

GOING OFF THE DEEP END

- MAYBE I SHOULD ACTUALLY LAUNCH THIS?
- PUT SOME EFFORT INTO IT AND SEE HOW IT PANS OUT?

GOING OFF THE DEEP END

- SOME GROUND RULES
- NATIVE IPV6 IS A MUST.
- IPV4 IS LIMITED, TRY AND USE V6 WHERE POSSIBLE
- LOGGING INTO THE PLATFORM IS UNWISE. EVERYTHING SHOULD BE AUTOMATED.
- CONFIGURATION AUTOMATION IS A **MUST**

Thankfully, Linux has some very handy functions for this.

FRRouting calls it 'extended-nexthop'
Using IPv6 addresses for IPv4 destinations.
Almost completely removes the need for IPv4 internally.

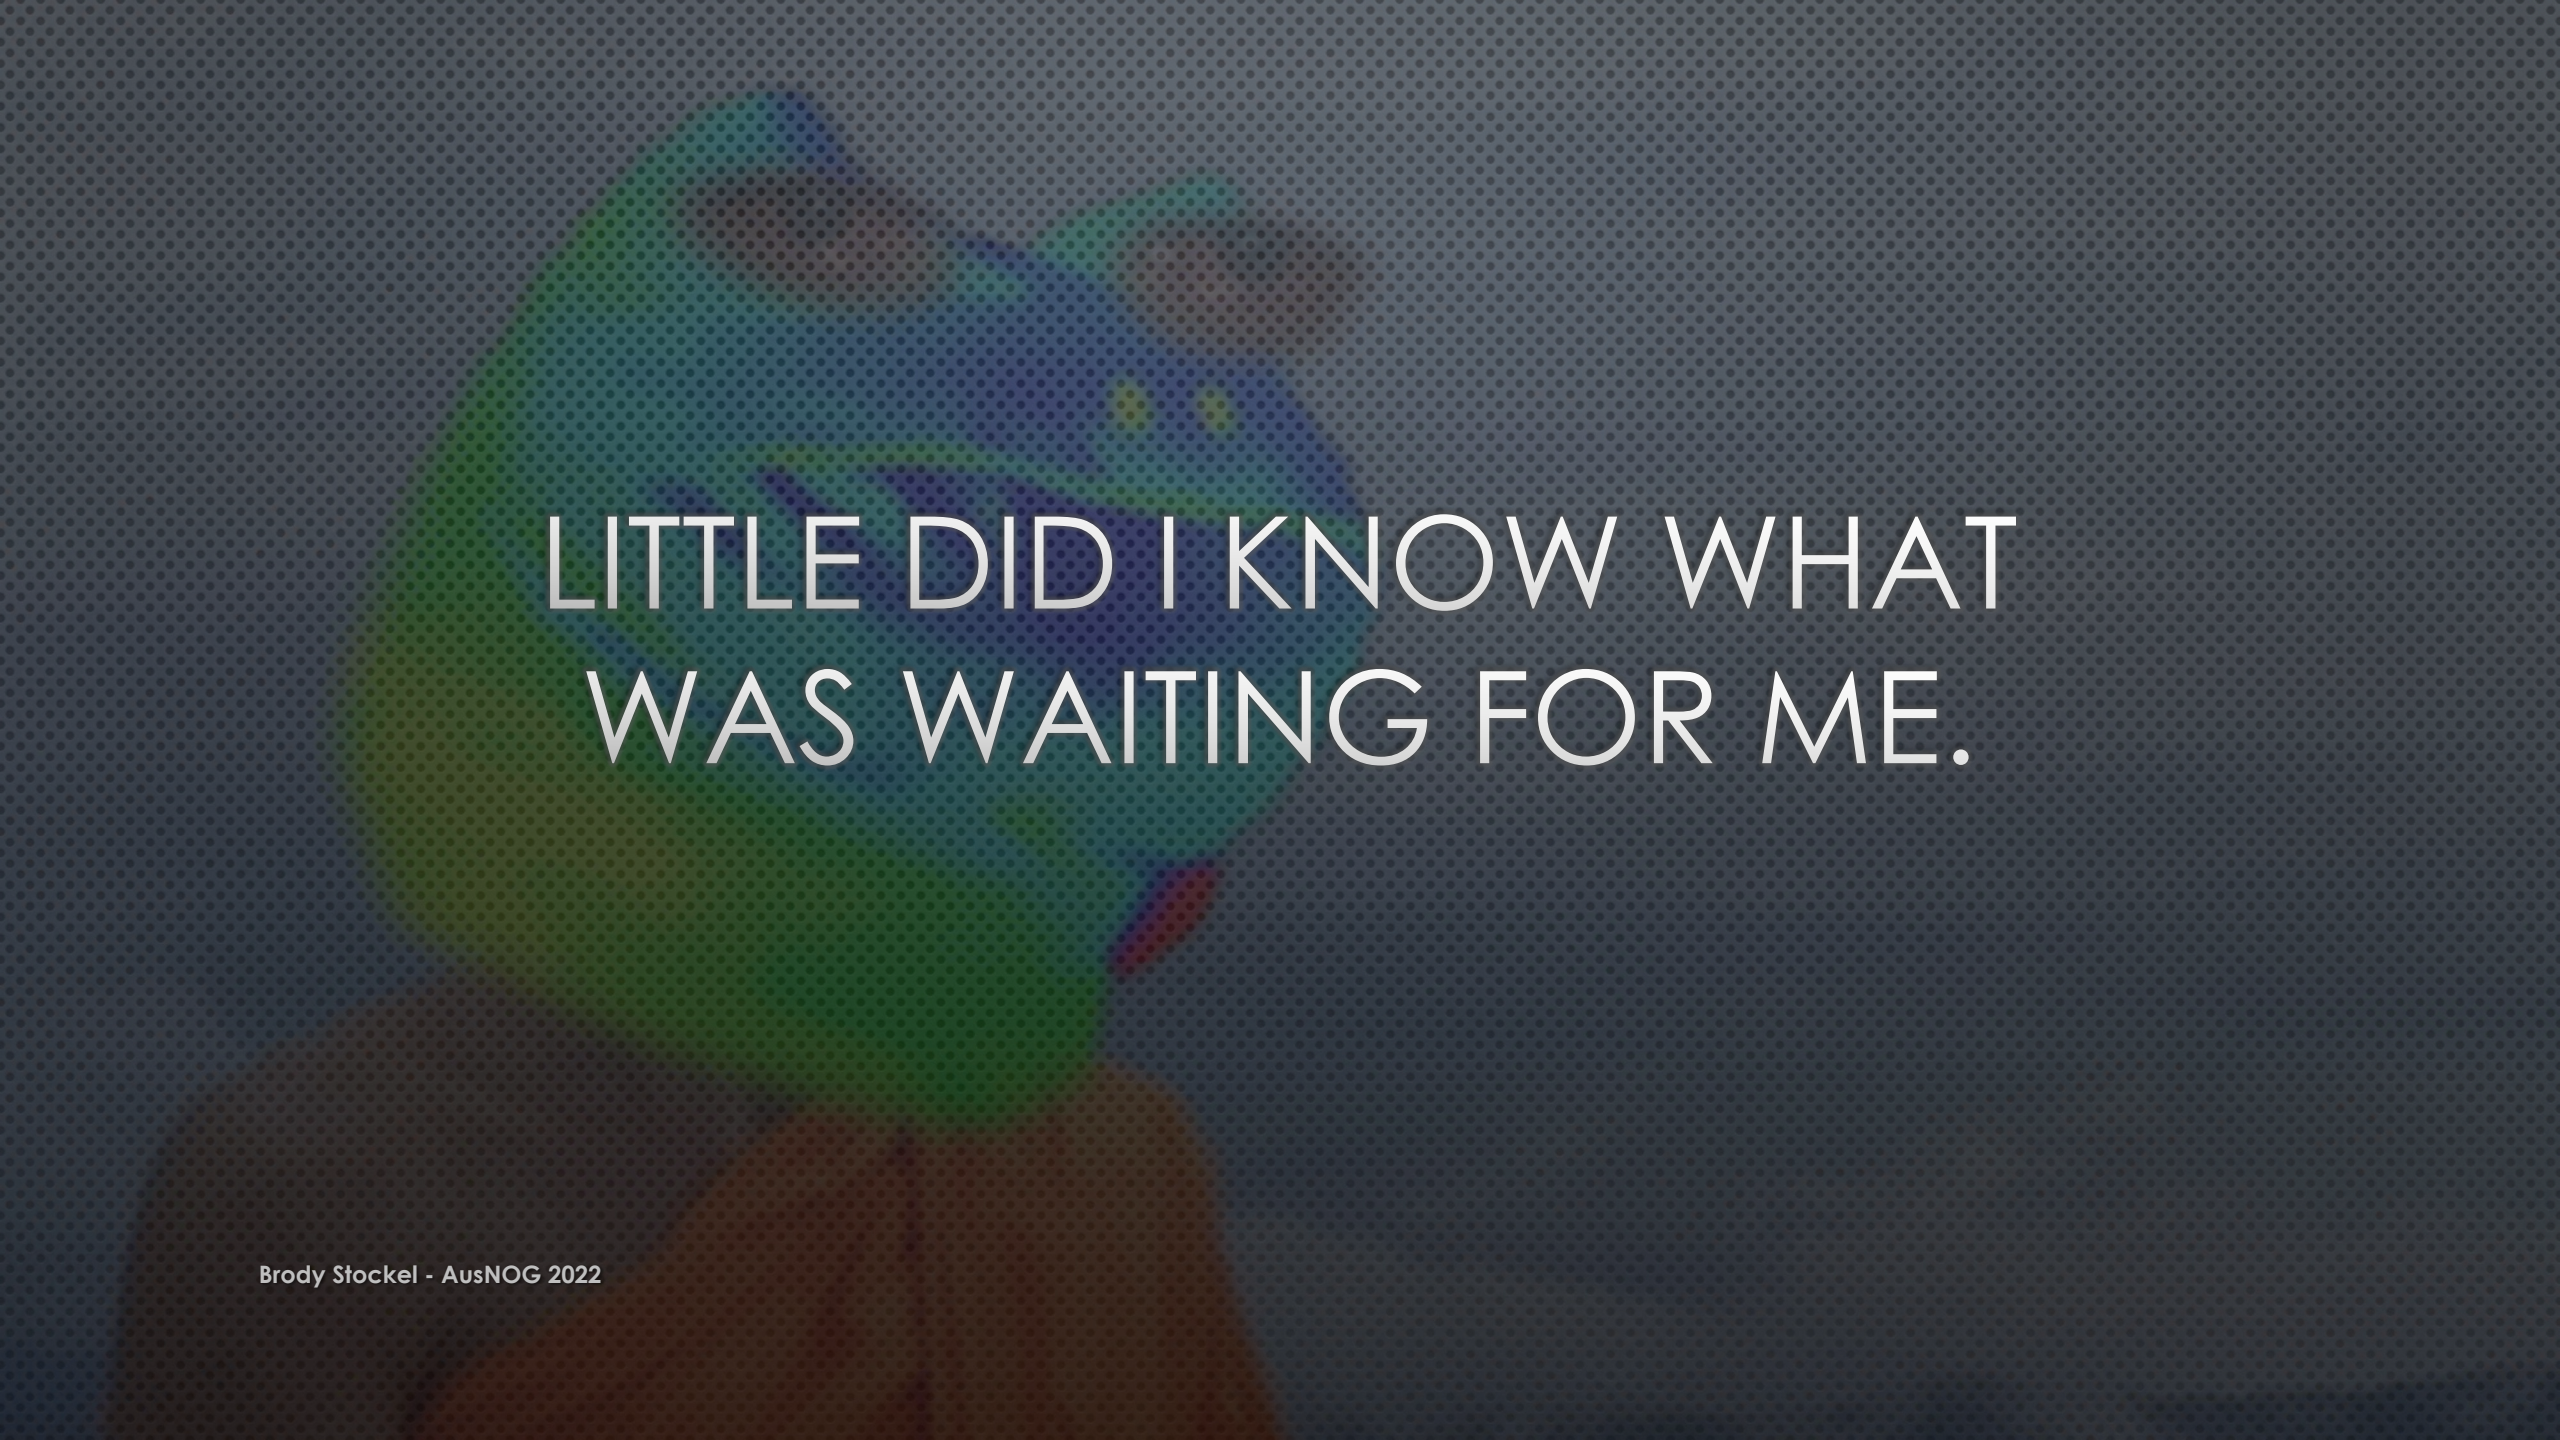
LITTLE DID I KNOW WHAT
WAS WAITING FOR ME.

GOING OFF THE DEEP END

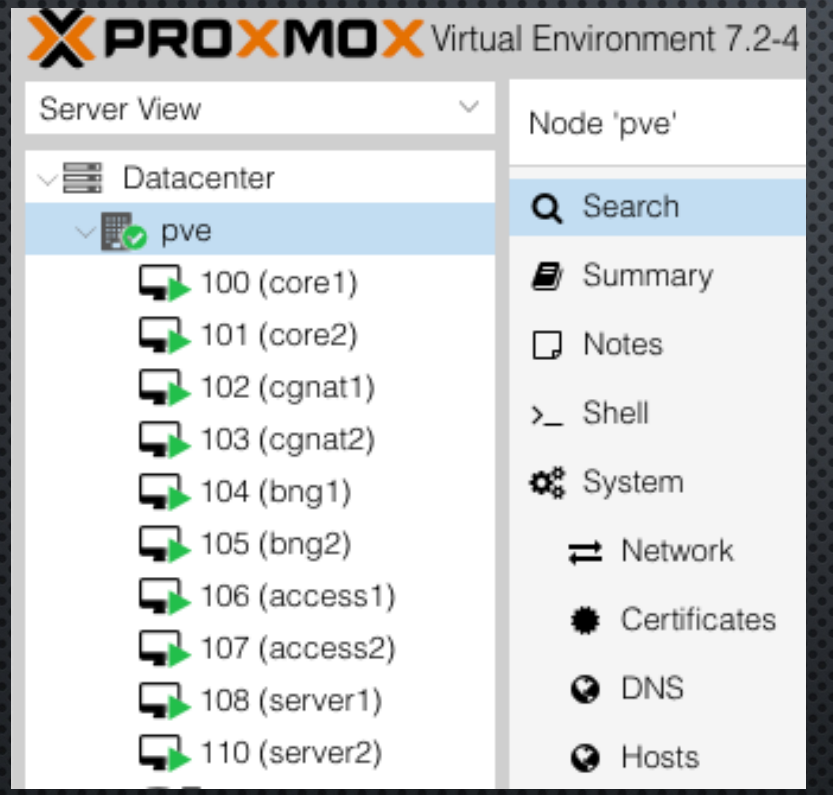

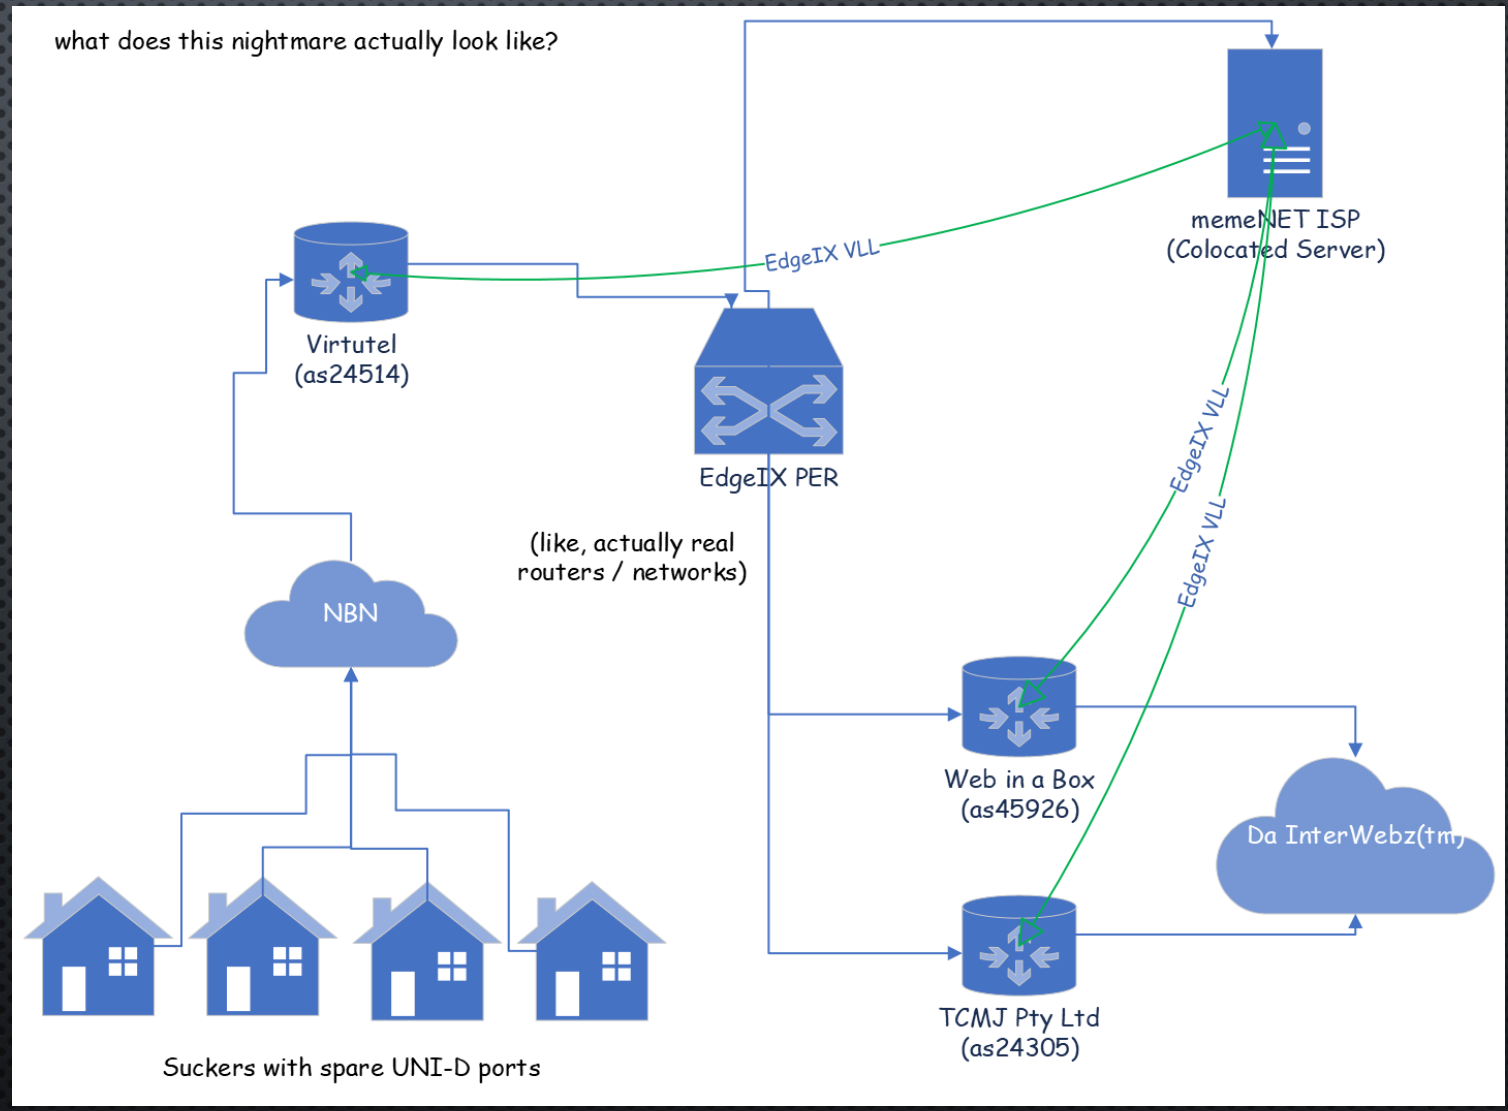

SOMEONE PLEASE HELP…

As expected, not everything goes according to plan.

- FRR-EVPN DOESN'T SUPPORT IPV6 VTEP ADDRESSES
- DANOS DOESN'T HAVE EXTENDED-NEXTHOP SUPPORT
- ACCEL-PPP DOESN'T SUPPORT IPV6 DESTINATIONS FOR RADIUS
- ACCEL-PPP DOESN'T NATIVELY SUPPORT AUTHENTICATING VIA DHCPV6

Just to name a few……

OH GOD, IT HURTS SO BAD

- DANOS DOESN'T SUPPORT BFD
- LINUX ECMP AND CONNTRACK GET REALLY WEIRD AT TIMES
- ACCEL-PPP DOESN'T SUPPORT GOING MORE THAN 1 VLAN DEEP
- LINUX BRIDGES AREN'T TRULY ISOLATED, THE BRIDGE HOST WILL TRY AND RESPOND TO ARP REGARDLESS OF IP CONFIGURATION

The Journey for IPv6

- Accel-PPP by default only accepts DHCPv6 relay-fwd packets if they conform:
  - From sport 546
  - To dport 547
- NBN however change the sport to 547, so Accel-PPP will ignore it.

- Accel-PPP also doesn't deal with Option18 Interface-ID
  - So NBN will drop the responding packet without this option present

This can be fixed via a couple patches #opensource

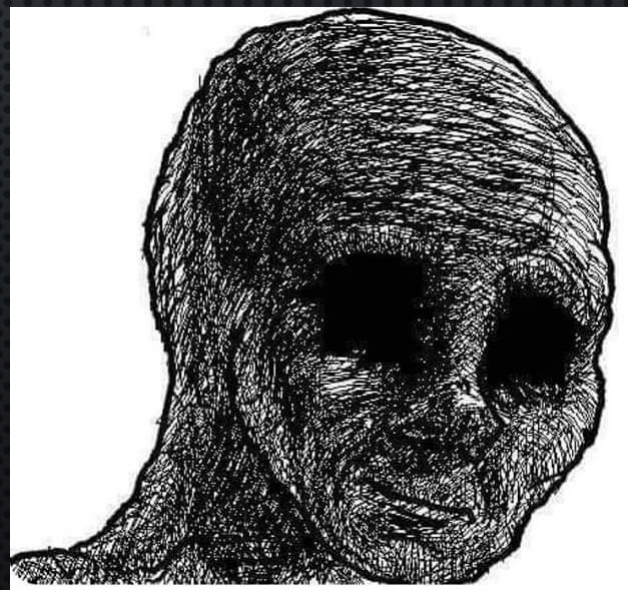

HOLY SH*T, IT WORKS

```
→  ~ tsh ssh root@bng1 "accel-cmd show sessions"
     ifname      |    username     |       ip       |                  ip6                   |        ip6-dp          | type | state  |
-----------------+-----------------+----------------+----------------------------------------+------------------------+------+--------+
 br-6BSD-1.2000 | br-6BSD-1.2000 | 103.176.61.65  | 2400:3120:fff0:7:d2dd:49ff:feef:d3ff/64 | 2400:3120:fff1:300::/56 | ipoe | active |
 br-4BDB-1.2000 | br-4BDB-1.2000 | 100.64.144.104 | 2400:3120:fff0:8:250:56ff:fe9e:73e/64   |                         | ipoe | active |
→  ~ tsh ssh root@bng2 "accel-cmd show sessions"
     ifname      |    username     |       ip       |                  ip6                   |        ip6-dp          | type | state  |
-----------------+-----------------+----------------+----------------------------------------+------------------------+------+--------+
 br-2LID-1.2000 | br-2LID-1.2000 | 100.64.144.102 | 2400:3120:fff0:14:76ac:b9ff:fea7:9520/64 | 2400:3120:fff1:600::/56 | ipoe | active |
→  ~
```

```
brodys@cgnat2:~$ show cgnat public
 Public Address  Port Range #Prts Blks Used  Blk  Port range   TCP   UDP Other
 103.176.61.193  1025-65534 64510     1/251    0   1025-1280    57     1     0
 103.176.61.194  1025-65534 64510     1/251    0   1025-1280    31     3     0
brodys@cgnat2:~$ show cgnat subscriber
     Subscriber      Paired Addr           Sessions  Blks  Ports    Map Reqs  Map Fails  Dura
 100.64.144.102   103.176.61.193                92     2     92       11737          0  03:(
brodys@cgnat2:~$
```

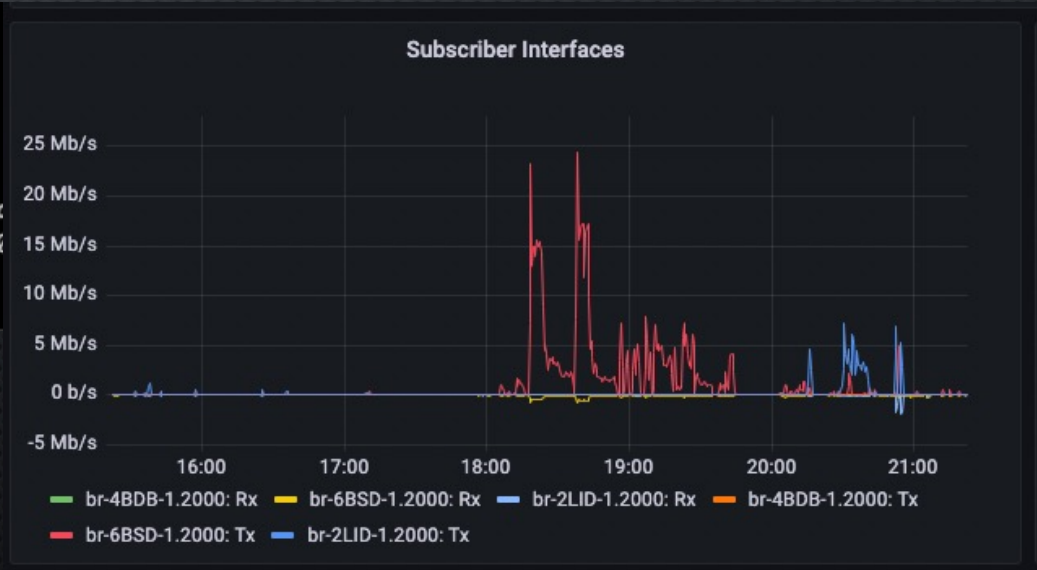

FULL SEND

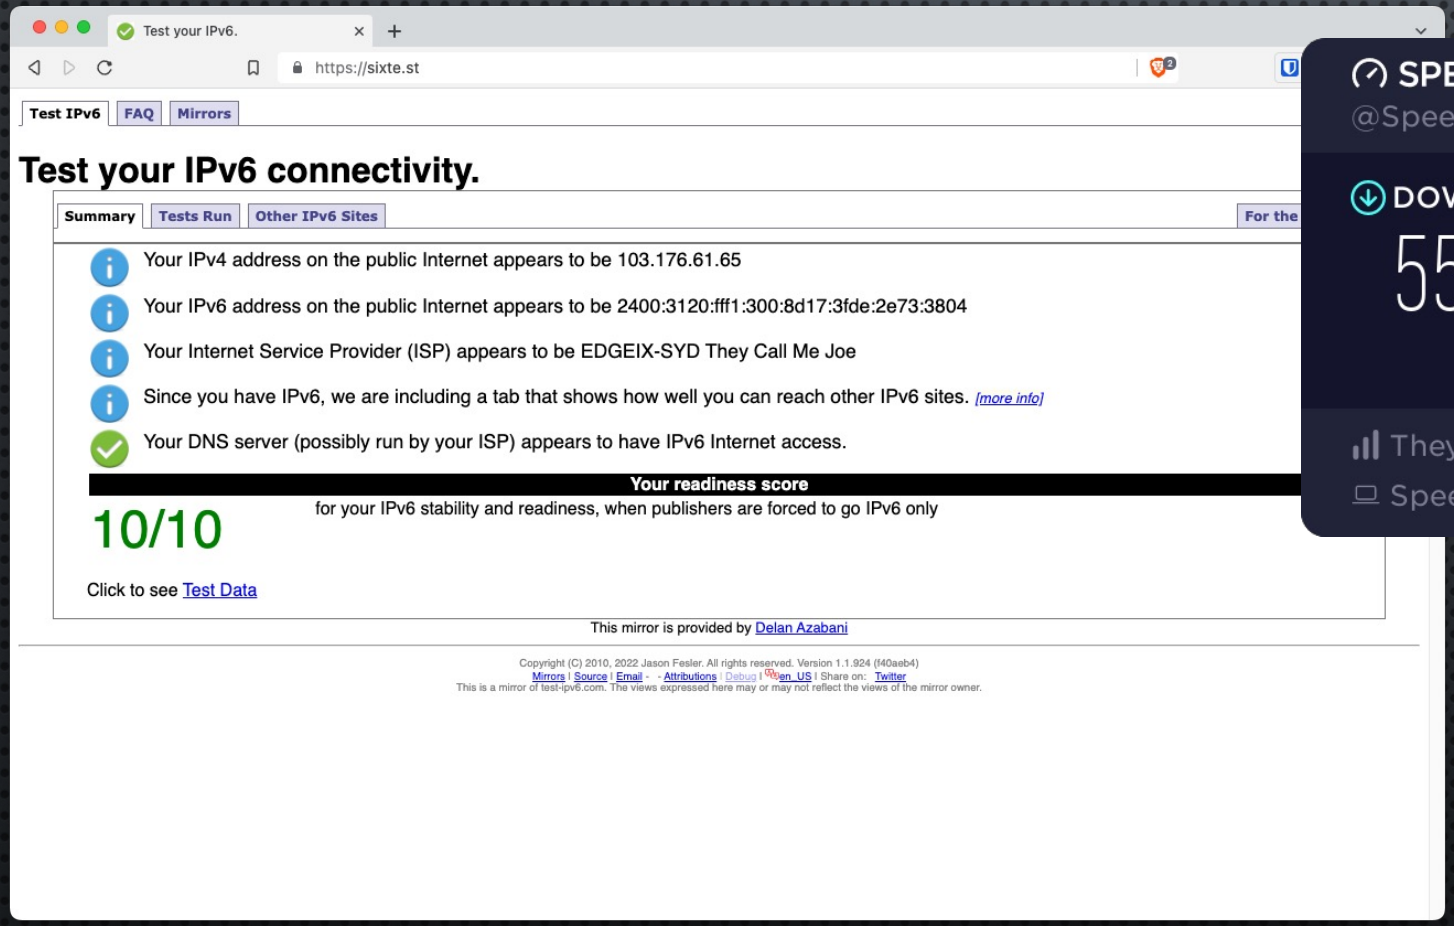
Test your IPv6.
https://sixte.st
Test IPv6
FAQ
Mirrors
Test your IPv6 connectivity.
Summary
Tests Run
Other IPv6 Sites
For the
Your IPv4 address on the public Internet appears to be 103.176.61.65
Your IPv6 address on the public Internet appears to be 2400:3120:fff1:300:8d17:3fde:2e73:3804
Your Internet Service Provider (ISP) appears to be EDGEIX-SYD They Call Me Joe
Since you have IPv6, we are including a tab that shows how well you can reach other IPv6 sites. [more info]
Your DNS server (possibly run by your ISP) appears to have IPv6 Internet access.
Your readiness score
10/10
for your IPv6 stability and readiness, when publishers are forced to go IPv6 only
Click to see Test Data
This mirror is provided by Delan Azabani
Copyright (C) 2010, 2022 Jason Fesler. All rights reserved. Version 1.1.924 (f40aeb4)
Mirrors | Source | Email - - Attributions | Debug | en_US | Share on: Twitter
This is a mirror of test-ipv6.com. The views expressed here may or may not reflect the views of the mirror owner.

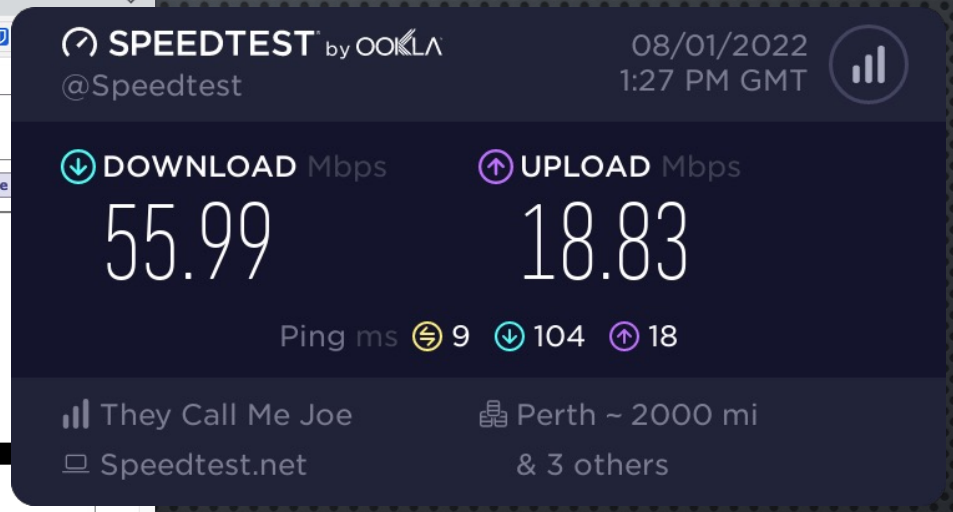
SPEEDTEST by OOKLA
@Speedtest
08/01/2022
1:27 PM GMT
DOWNLOAD Mbps
55.99
UPLOAD Mbps
18.83
Ping ms 9 104 18
They Call Me Joe
Speedtest.net
Perth ~ 2000 mi
& 3 others

JUST WHEN WE LOOKED IN THE CLEAR

```
[ ID] Interval           Transfer     Bitrate         Retr
[  5]   0.00-10.00  sec  1.52 GBytes  1.31 Gbits/sec  677             sender
[  5]   0.00-10.00  sec  1.52 GBytes  1.30 Gbits/sec                  receiver
```

Before

```
[ ID] Interval           Transfer     Bitrate         Retr
[  5]   0.00-10.00  sec  6.42 GBytes  5.52 Gbits/sec  3102            sender
[  5]   0.00-10.00  sec  6.42 GBytes  5.51 Gbits/sec                  receiver
```

After

Turns out, looping through Open vSwitch ~9 times in each direction
Isnt the best way to do some things

```
5: enp131s0f1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 4000 qdisc mq state UP mode DEFAULT group default qlen 1000
    link/ether 64:9d:99:b1:5f:97 brd ff:ff:ff:ff:ff:ff
    vf 0     link/ether ae:27:59:b8:b7:98 brd ff:ff:ff:ff:ff:ff, vlan 1000, spoof checking off, link-state auto, trust on
    vf 1     link/ether b2:70:00:ba:b1:00 brd ff:ff:ff:ff:ff:ff, vlan 1001, spoof checking off, link-state auto, trust on
    vf 2     link/ether 9a:58:96:e7:31:3a brd ff:ff:ff:ff:ff:ff, vlan 1002, spoof checking off, link-state auto, trust on
```

Virtual Functions to the rescue!

JUST LOOK AT THESE TESTIMONIALS

saul i'm not sure i have anything to complain about with this halfass nbn connection anymore

seeker

paul allen i have never been to that restaurant

BUT HOW DOES IT COMPARE?

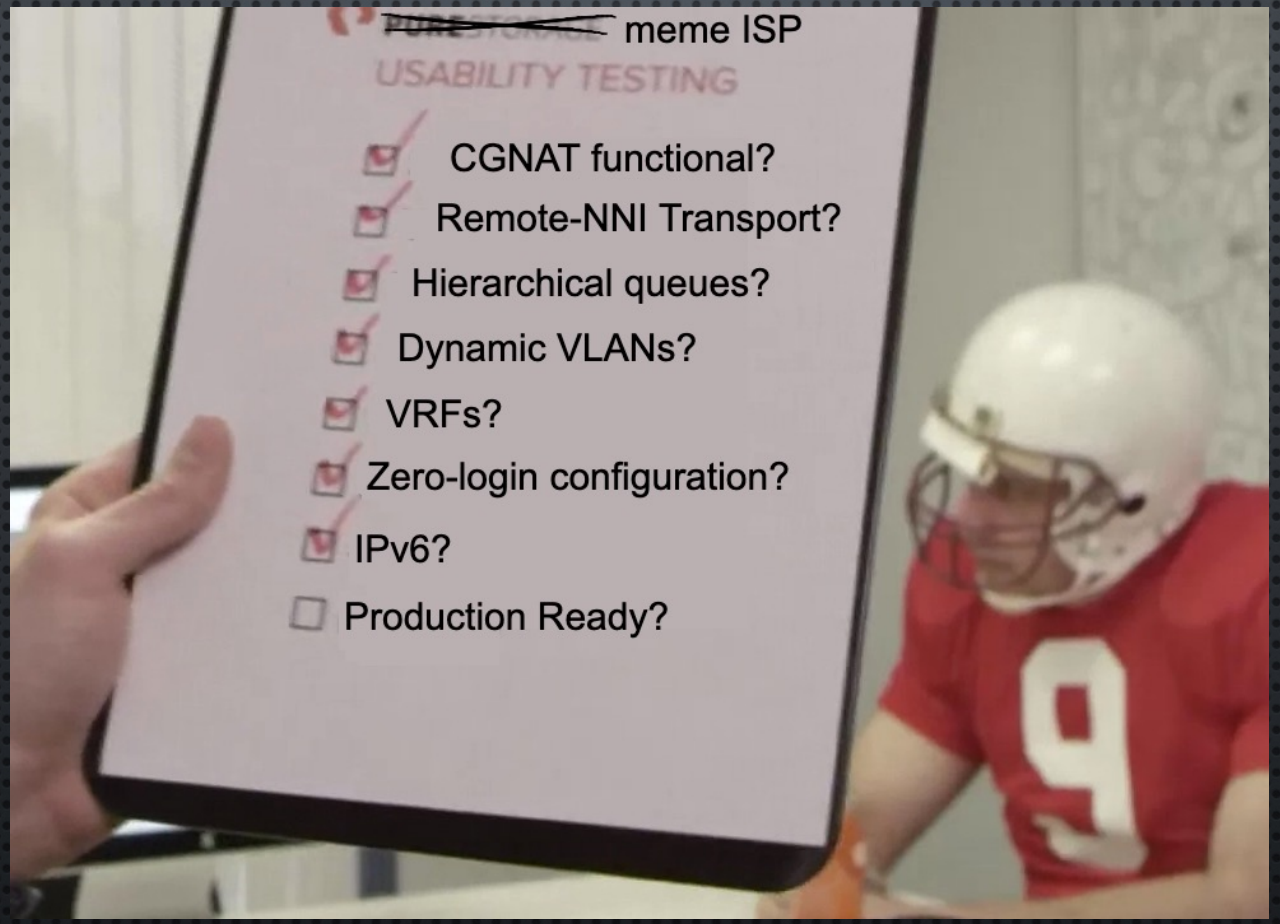

SO, ARE WE THERE?

- In my opinion, for the right use case, yeah..
- Some limitations, no deal breakers however
- General NBN RSP duties, no worries.

Costs:
- 40-50 hours of dev time to get going
- 1 server bought from eBay
- 1 borrowed NIC
- Several under the table deals

- $300-400 of Suntory Lemon
- $150 of Soju

= one ISP strung together with hopes and dreams, and operating stable.

SO, ARE WE THERE?

- Cloud BNG? (Access switches @ POIs, BNGs in the public cloud)
- Containerised BNG? (running the stack under k8s)
- Micro BNG (Small BNGs located @ the POIs)

HOW ABOUT A MORE REALISTIC DESIGN?

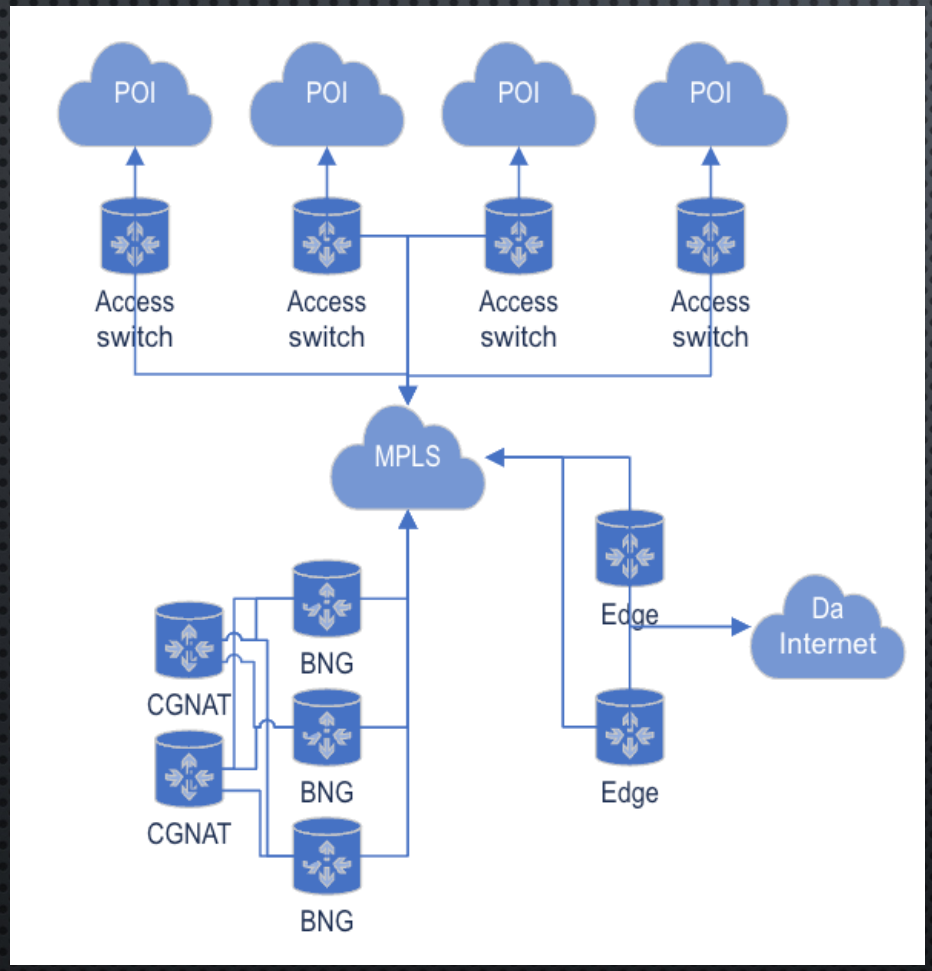

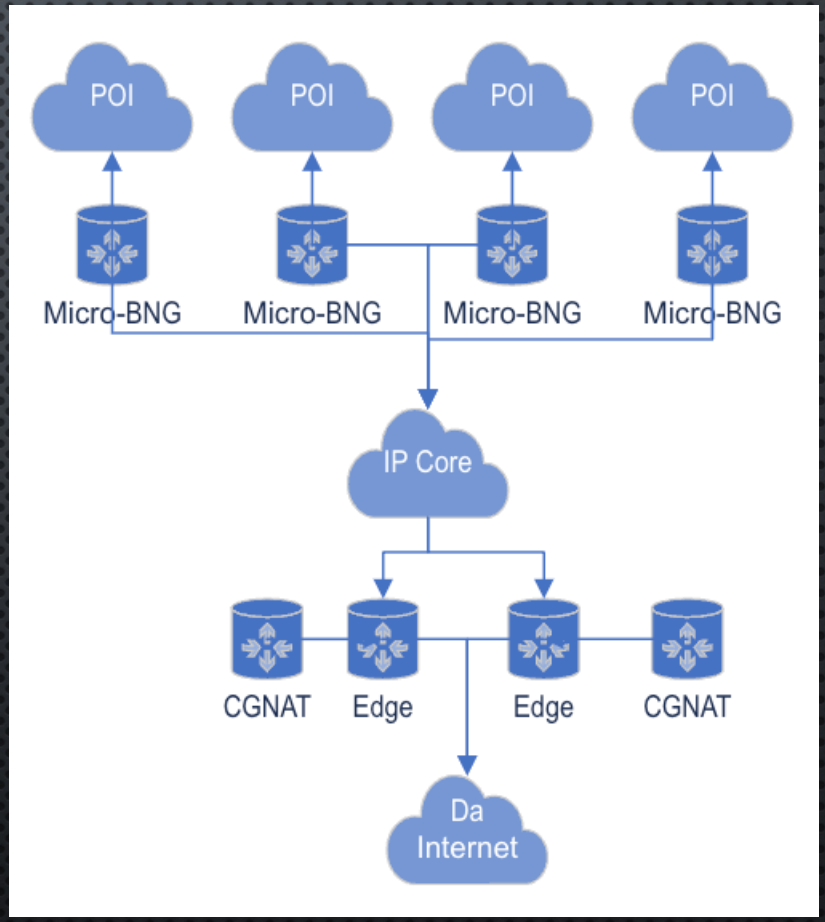

DID YOU KNOW?

- IT ONLY TAKES 48 HOURS FROM UPDATING PEERINGDB TILL COGENT EMAILS YOU.
- HE.NET WAS THE FOLLOWING DAY

Introducing Cogent (ISP/AS174)– presence in Perth!

From Wijsman, Ron <rwijsman@Cogentco.com> Date Fri 20:12

Dear Sir

Good day

I would like to introduce to you Cogent Communications – we are a global IP carrier.
We are keen to learn more about your business and explore if we can collaborate with you.

Updated	
Contact Info Updated	2022-07-06T10:52:51

I am one again asking to discuss how Cogent Communications could partner with you and help lower the cost of your Internet, Colo, and IP needs

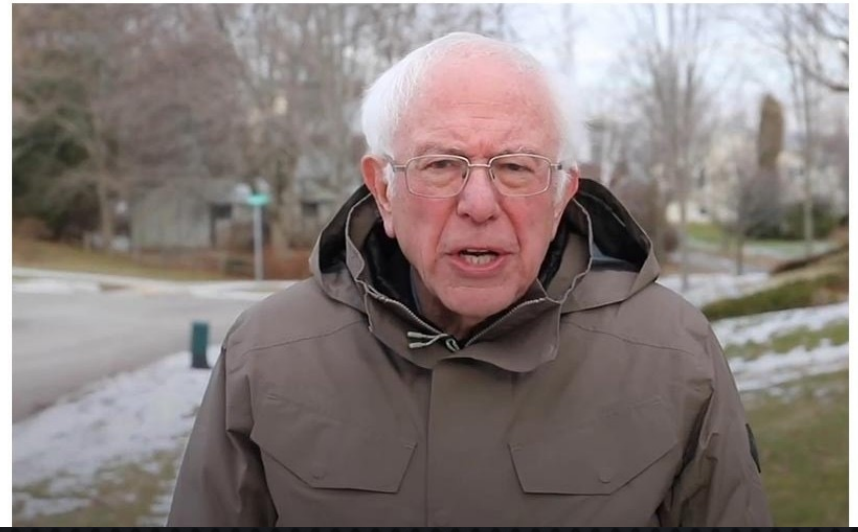

THIS PROJECT WAS SPONSORED BY:

For providing:
An ASN, rack space,
IP transit, and a peering port

For providing:
NBN Aggregation Connectivity

For providing:
IPv4 and IPv6 space

For providing:
Additional IP transit capacity

THANK YOU.

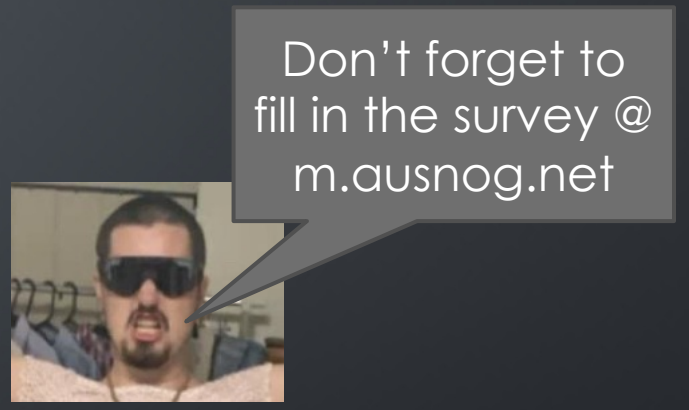

- TWITTER.COM/BRODYSTOCKEL
- LINKEDIN.COM/IN/BRODYSTOCKEL

- GITHUB.COM/VSQUAR3/MEMEISP <-- GITHUB LINK WITH A BUNCH OF STUFF I WROTE