Remote Hands

Rebuilding a Pacific Island Internet Service Provider during a Pandemic

Philip Smith

PFS Internet Development

November 2020

Acknowledgements

- This presentation documents the reconstruction of CenpacNet on the Pacific Island nation of Nauru
- Grateful thanks to Ru Garoa and the Cenpac team for allowing me to share their story
- Thanks to the Network Startup Resource Centre for encouragement and inspiration

Nauru

Map by Google

Where?

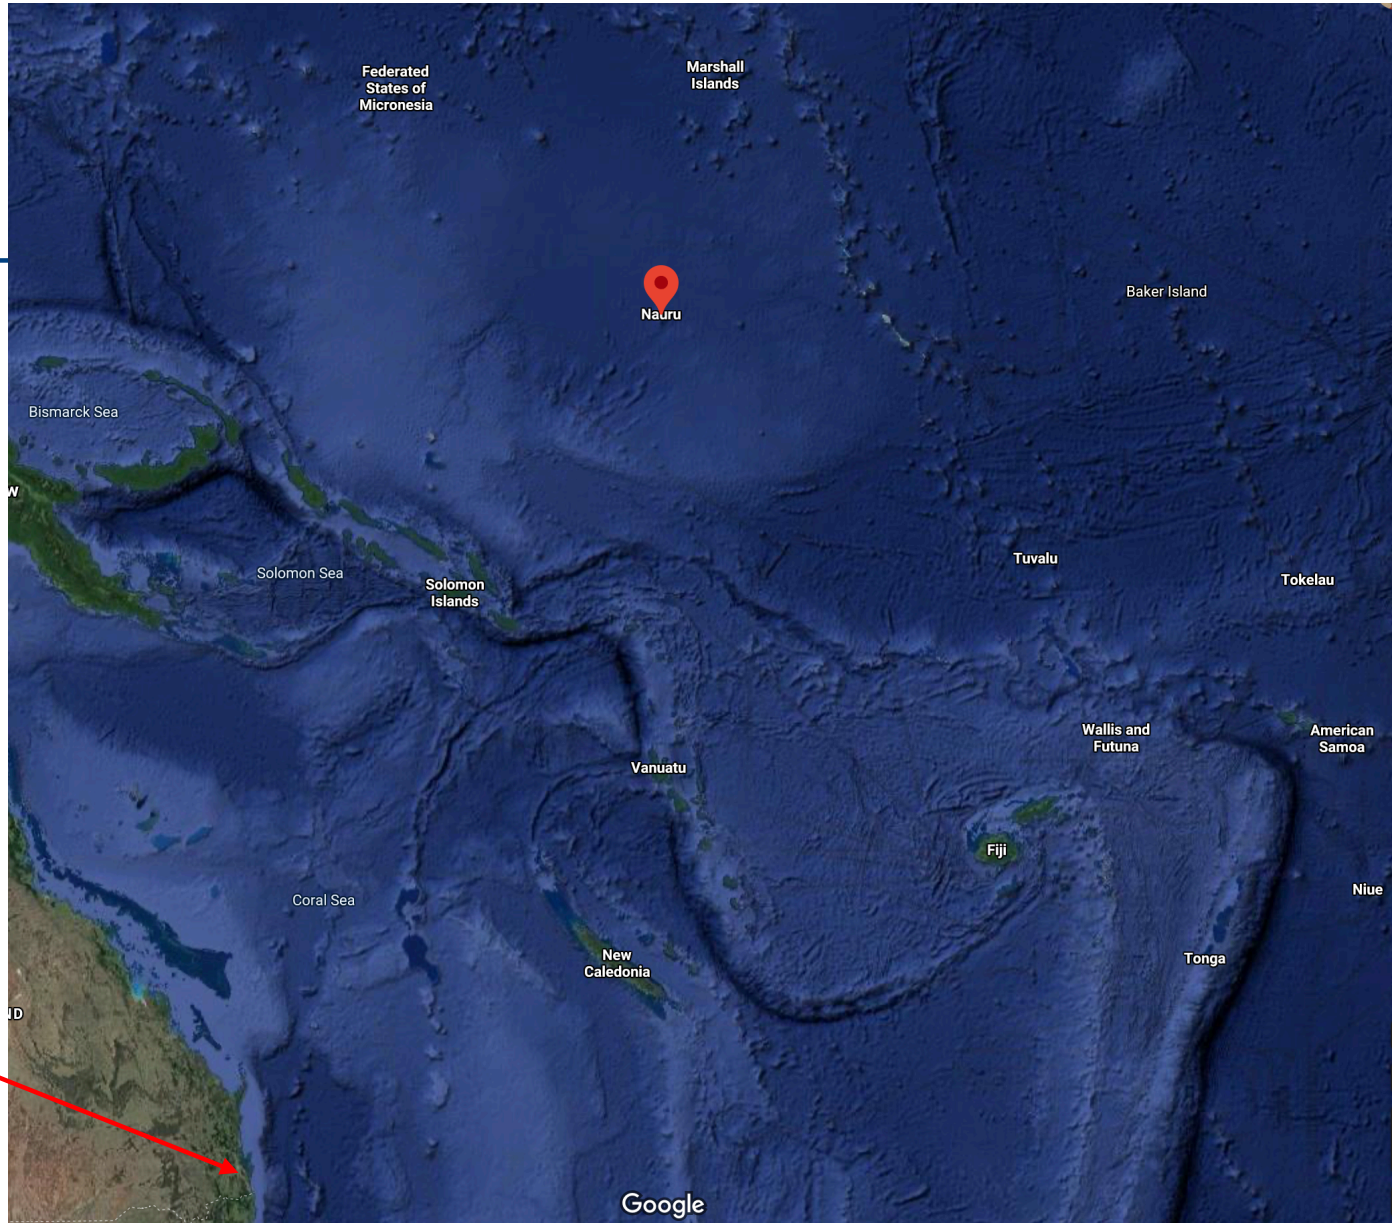

Brisbane, Australia

Map by Google



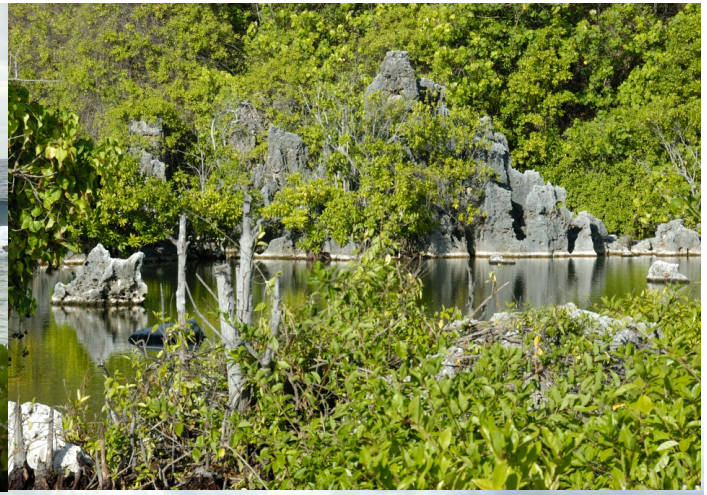



History

- Cenpac (Ru Garoa) came to PacNOG 6 in Fiji
  - November 2009 @ Tanoa Hotel in Nadi
    - https://www.pacnog.org/pacnog6/
- Reason:
  - Seeking assistance from the Pacific community
    - Their lead (only) technician had died
    - The only other personnel were administrative
  - Cenpac operated the .NR ccTLD and ran an Internet Café for the Island
    - https://www.pacnog.org/pacnog6/presentations/nauru-pacnog6.pdf

History

- Network in 2009 was:
  - 384kbps satellite link to Telstra (Australia)
  - Windows NT server running firewall and Internet gateway via a Cisco 1005 router to Newbridge 2601 satellite modem
  - NRWEB, providing ccTLD registry for .NR (by Franck Martin), Webmail and POP3 Email service, and primary DNS for Island
  - DNS1, providing Email SMTP/Anti-Virus/Anti-Spam, and secondary DNS for Island
  - 2x Cisco 2509 with dial-up modems
- Five years of modernisation commenced

History

- By end of March 2014:
  - 7/2.5Mbps Satellite via PacTel from Australia to Nauru
  - Cisco 2811 border router in Sydney
    - AS55722
  - 2x Cisco 3745 routers on Island
    - Customer aggregation (WiFi and ethernet)
    - Local peering with Digicel Nauru (AS55723) – mobile operator
  - Dual-stack network
    - 203.98.224.0/19 and 2403:F600::/32
  - Three “home made” servers running Ganeti
    - NRWEB and DNS1 virtualised
    - Another VM hosting TACACS+, RANCID, Smokeping, Observium, etc

History

- Development work from 2009 to 2014 was implemented by:
  - Remote hands & remote login
  - Site-visit along with Phil Regnauld (NSRC) in July 2012
- From April 2014 to November 2019
  - After change in Government, Cenpac operated by different entities
    - Mostly as ccTLD operator
  - Satellite capacity moved from PacTel to AccLinks
    - AccLinks also built domestic WiFi backbone
  - Hardware changes due to equipment failures
  - Most pre-2014 customers transferred to Digicel Nauru
- Reconstruction & modernisation started in November 2019

Agenda

- Satellite links
- Routing (BGP, IS-IS etc)
- Island-wide WiFi backbone
- Ganeti Cluster
- Router & Switch Infrastructure
- Network Operations
- Observations

Satellite Moves – 2019

- 11th October, AccLinks ordered to cease operations on Nauru
  - Panic on to secure alternative Internet connectivity for the ccTLD and other entities connected to Cenpac
- Original Pactel (now Speedcast) link resumed in December
  - Link used by the local bank branch
  - Negotiations and logistics – Speedcast site visit on 18th December to modify existing connection to also support Cenpac use
- AccLinks slowly reduced the bandwidth provided to Nauru as they withdrew service
  - Deadline end October, but were still there early December

Down down down…

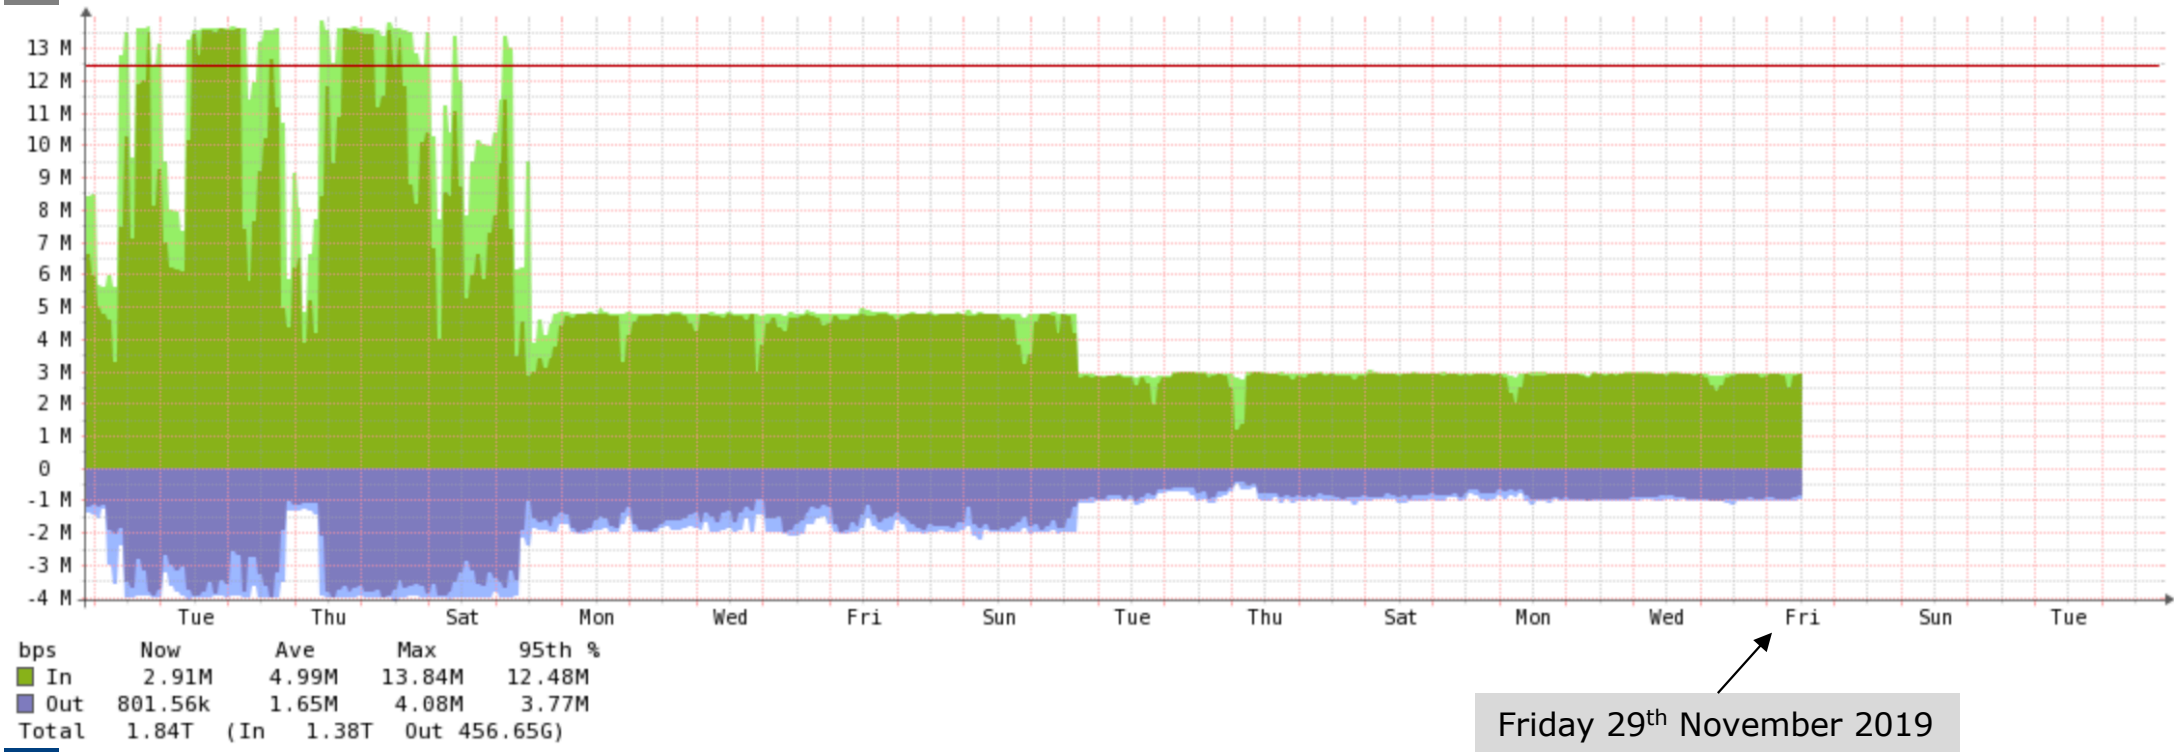

bps	Now	Ave	Max	95th %
In	2.91M	4.99M	13.84M	12.48M
Out	801.56k	1.65M	4.08M	3.77M
Total	1.84T (In	1.38T	Out 456.65G)	

Friday 29th November 2019

Off Air

- Friday 29th November Cenpac and .NR goes off air – AccLinks removes service
  - Service restored to allow for transition – at 128kbps!!
  - Remaining customers disconnected to allow:
    - Router reconfiguration work to accommodate new Speedcast link
    - ccTLD DNS to actually work!
- Wednesday 18th December
  - Speedcast technician arrives on-site to reconfigure satellite link
- Friday 20th December AccLinks link totally cut for final time
  - .NR goes off line, and the entire ccTLD vanishes once DNS TTLs timeout
    - Many important sites no longer resolve: Nauru Govt etc
  - Connectivity removed mid-on-site work by Speedcast
  - Before Cenpac routers reconfigured for new Speedcast connection! 🤯😤😭

Panic to get back online

- Access to routers appeared like chicken & egg problem!
- Options:
  1. New Speedcast connection had public IP address for point-to-point at my insistence
  2. Use Digicel 3G USB modem and AnyDesk
- Solution: Local laptop plugged into Speedcast point-to-point
  - SSH to laptop
  - Serial console to Cenpac router – cable didn't work
  - SSH from laptop via office WiFi to Cenpac network
  - IP address on router interface changed for Speedcast set up
  - Default gateway changed
  - Laptop disconnected, router connected – it worked!! ☺

Back On Air

- 21st December (Saturday)
  - BGP peering brought up by Speedcast
  - But upstreams (Telstra, Vocus, Optus) only update filters every 24 hours!!! 😭🤯😰
  - And AccLinks still originating Cenpac address space from their AS (they had not used AS55722 origin) 😭🤯😰
- 22nd December (Sunday)
  - AccLinks stopped originating Cenpac address space
- 23rd December (Monday)
  - Filter update by Telstra – Internet returns to Cenpac, .NR works again
  - Still no filter update by Vocus or Optus 😭🤯😰

And then… (un)Happy Christmas

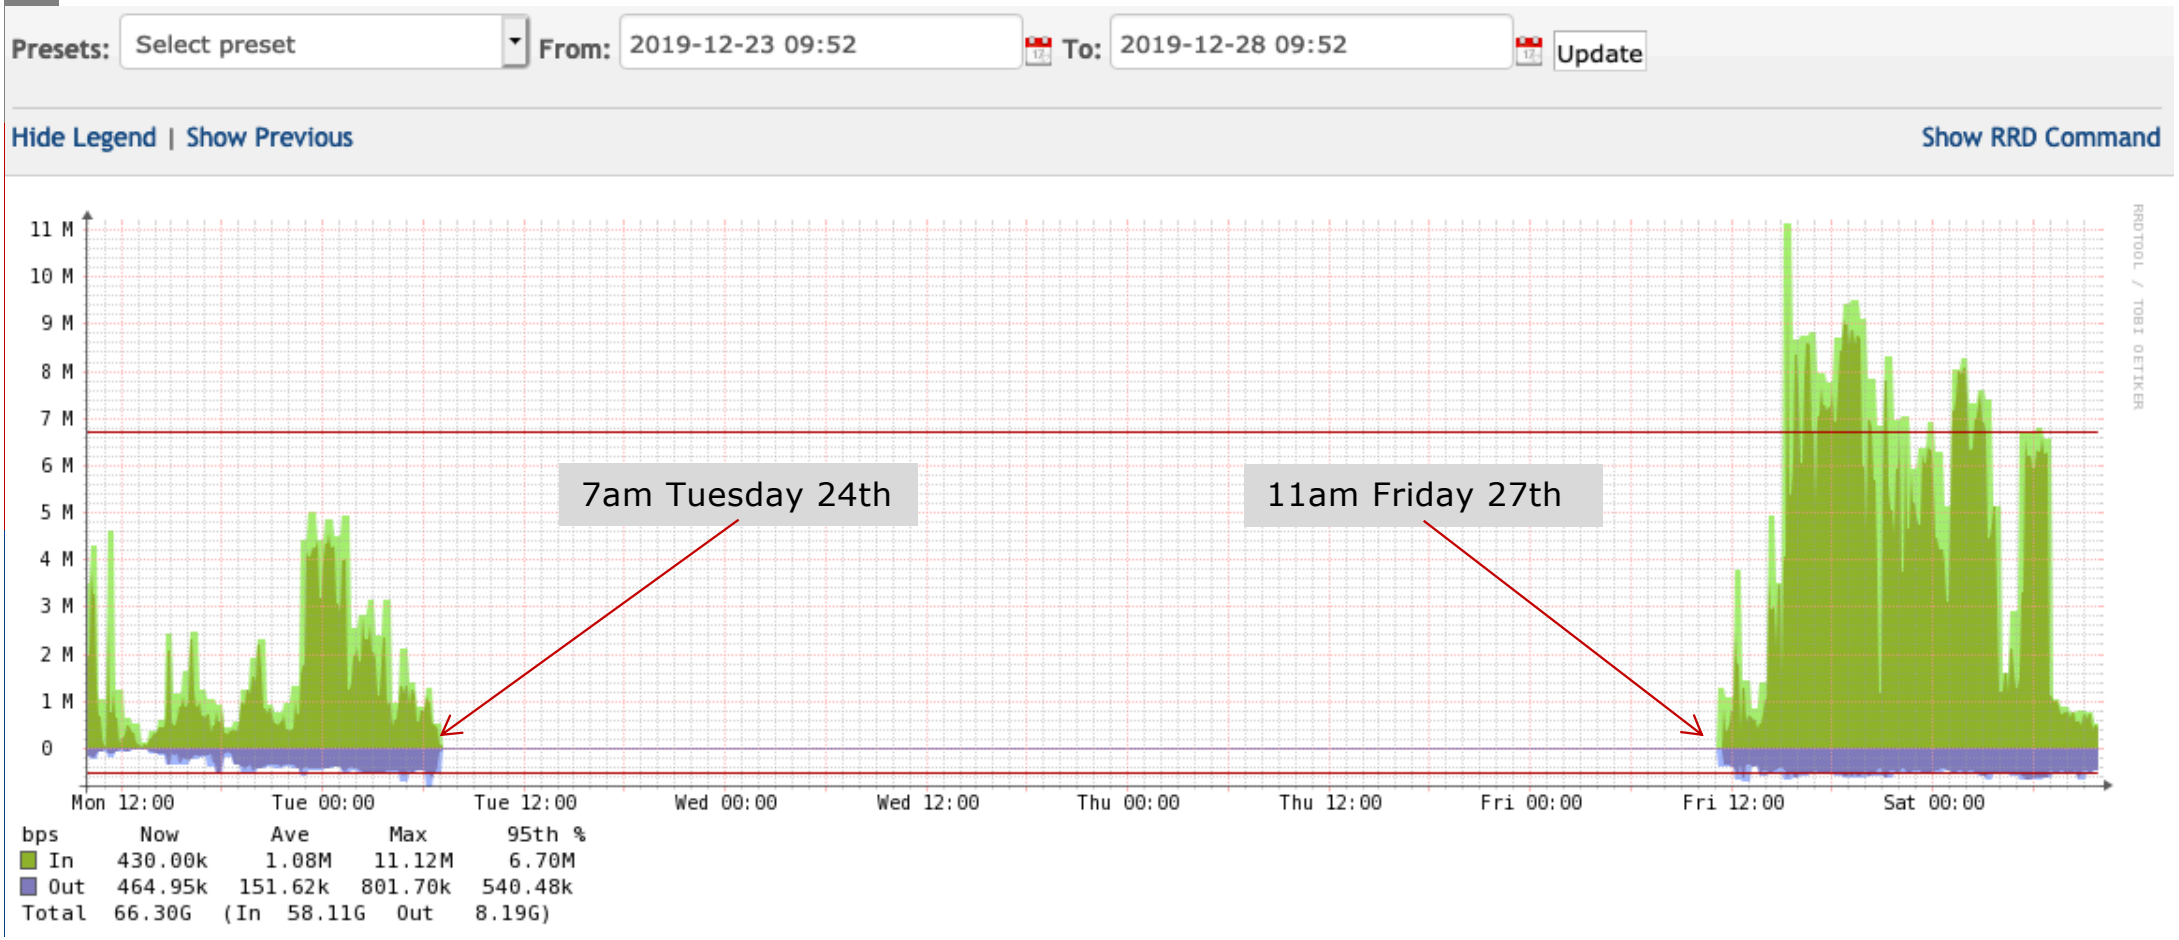

What happened??

- Diagnosis was that the Block UpConverter (BUC) had failed
  - New BUC put on next flight to Nauru (ex-Sydney)
- But a few days later the link came back!
  - So it wasn't the BUC – the new BUC finally turned up months later!

2019-12-27 04:04:16 GMT - Andy C... Additional comments

Hi All,
The link is now back online and performing well.
Service was restored after sending a SMTP command to transmit a CW and then once the carrier was modulated the service was restored.
Regards,
Andy C...

2019-12-27 04:01:58 GMT - Andy C... Additional comments

Hi Rudolph,
Fortunately one of our Operations engineers has managed to restore service on this link.
He found that the Transmit level was extremely low, could barely be seen on the spectrum analyser.
Once he saw this he then sent a command to the modem to transmit a CW which has somehow corrected the issue once the carrier was modulated.
The link is now performing well.
Regards,
Andy C...

Lesson?

- Cutting one connection before the other is commissioned isn't a great idea
  - But what to do when it is mandated by Govt?
  - Caused **major** stress – ccTLD out twice
    - 19th – 22nd December
    - 24th – 27th December
  - Creative ways were needed to access the equipment
    - At least we had two options!

All okay? Ummm…

Just a small problem

Path from Satellite to Cenpac goes via the Bank

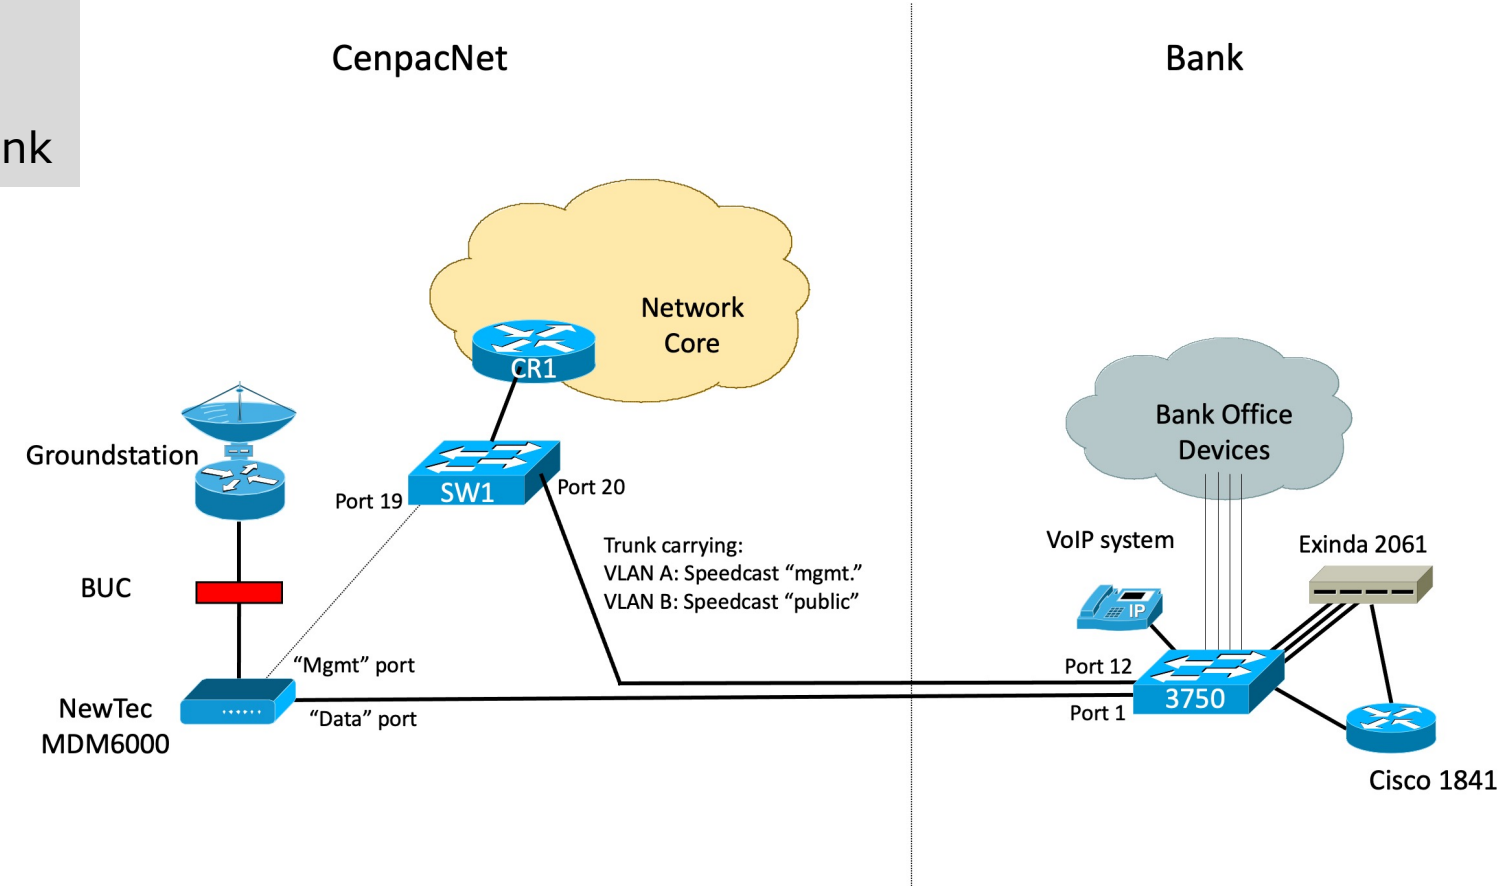

The dog-leg

- Speedcast first visit in late December was urgent trip to get Cenpac re-connected at v short notice
  - Satellite dish on Civic Centre roof, with modem in Cenpac datacentre, and ethernet from modem run to the Bank (in the same building)
  - Added Cenpac connection to existing Bank connection – ethernet cable run from Bank back to Cenpac
    - With a view to fixing this in the new year
- Then SARS-COV-2 struck, travel banned from late March
- Dog-leg finally removed in mid-November 2020, with shipment of new modem (Newtec 2510) with multiple ethernet ports
  - One port for the Bank, another port for Cenpac

No more dog-leg

Fixed!

After 11 months of risk

Including power loss at bank (Cenpac has UPS and Generator, bank does not), equipment turned off, etc, etc.

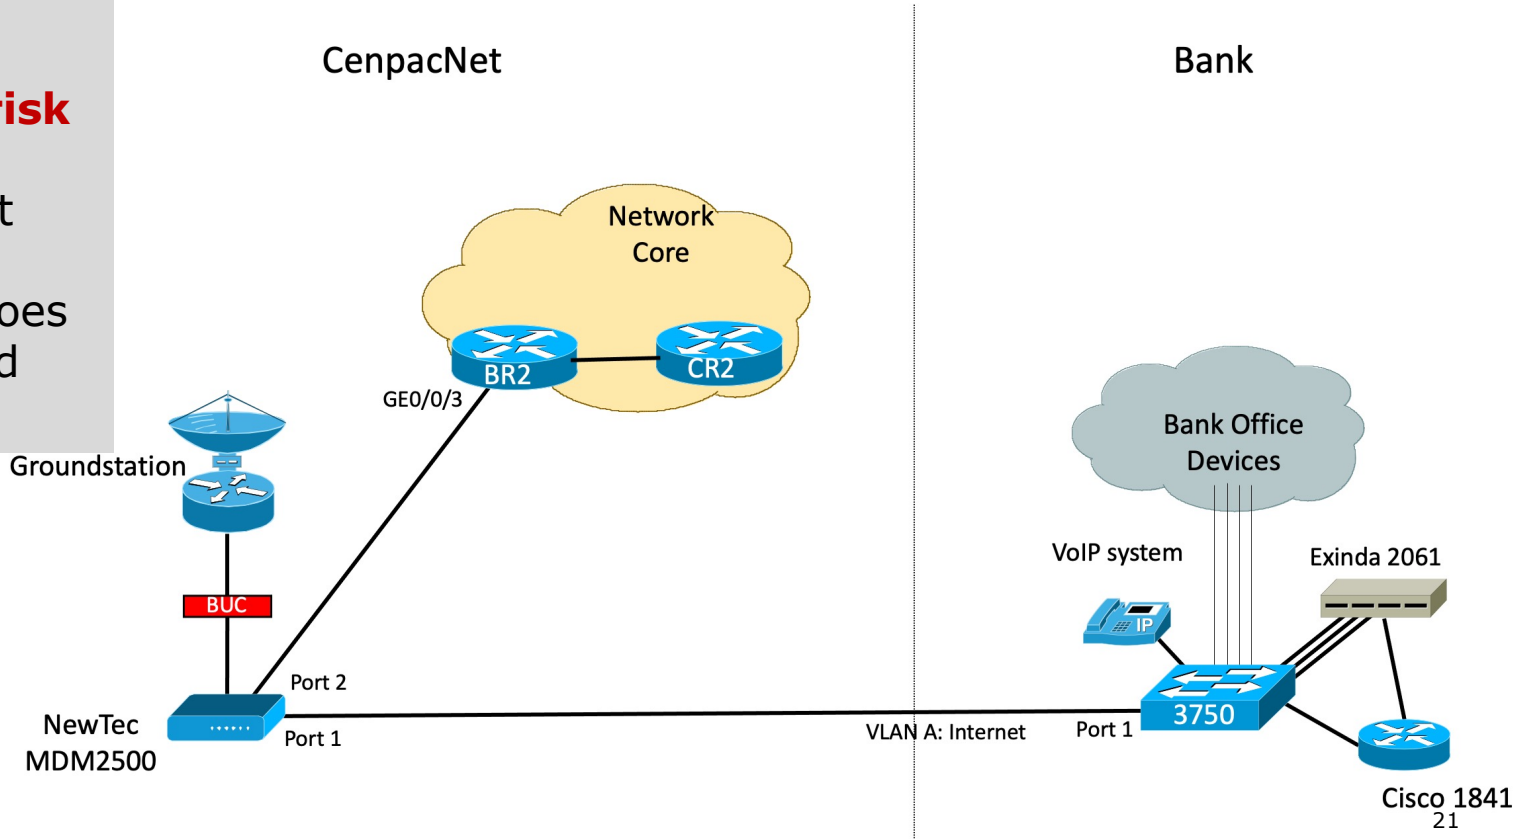

Next Step – Redundant Links

- Second satellite provider (Av-Comm) was contracted
  - SARS-COV-2 spoiled plans here too 😤😭
- Original Plan:
  - Dish and equipment shipped to Nauru
  - AvComm staff come to install everything, and hand over a working service
    - Reason: "Delicate jobs only they can do" – fair enough!
- And then:
  - All Australians and Permanent Residents banned from overseas travel from end of March 2020 – until further notice
  - Shipments delayed due to ensuing COVID-19 shutdowns in Australia
    - Freight flights cancelled, staff work@home, warehouses minimal staffing, etc

DIY Satellite Build

- Cenpac team had to build the dish themselves
  - May/June 2020 work
  - Instructions provided by Av-Comm
  - Not everything fitted exactly 🤯😤🙄
- Next, cabling up the dish to BUC to satellite modem
- And then link provisioning:
  - Weak signal (dish issue? / BUC issue?)
  - Mid-June – BGP running, upstream filters opened up

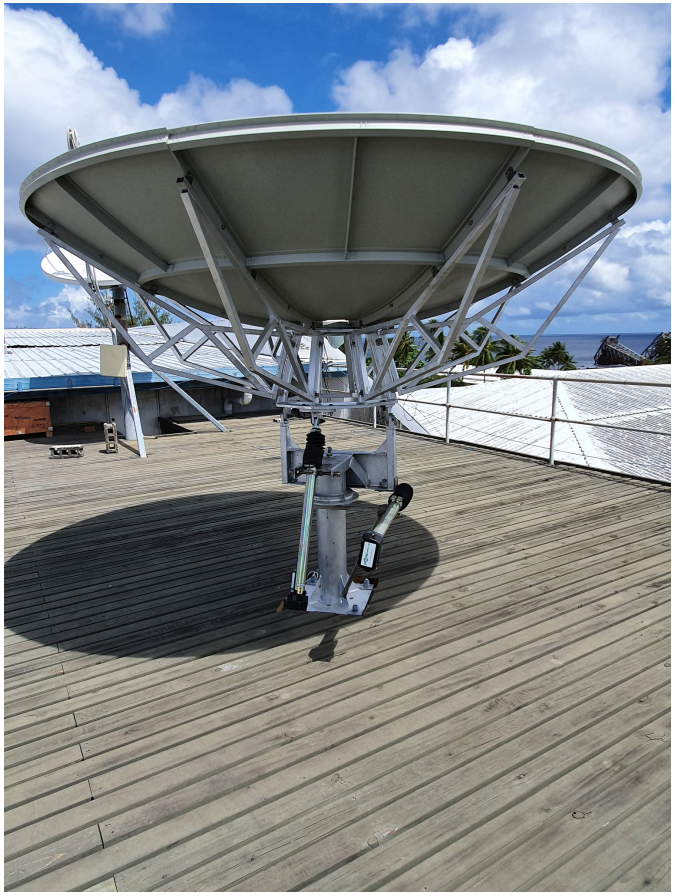


Satellite Systems Corporation
Av-Comm
F: 1450.00 ALARM
SS:0.01V G:873
SYSTEM
ERROR
LOCK
RX1
RX2
TX
UHP
AZ: 244 EL: 301 S: 0 REMOTE
SAT:IS-19
Research Concepts, Inc.
RC2000 DUAL AXIS ANTENNA CONTROLLER
Av-Comm
MODEM PORT 1 / CUSTOMER PORT 2 / MGMT PORT 6-10
RB2 1UiAS-RM
POE
GIGABIT ETHERNET
ETH2
ETH3
ETH4
FAST ETHERNET
ETH6
ETH8
ETH9
ETH10
NAURU
Av-Comm
MikroTik routerboa
SW1
Cisco Systems
SW1-PRT-15 TO V5 SERVER
SW1-PRT-14 TO V4 SE
SW1 PT-21

Link performance tests (1)

- July 24th – something not right (meant to be 62/34Mbps)
  - Inbound

```
On NR host:  iperf3 -s
On US host:  iperf3 -c 203.98.X.X -P 40 -t 250 -O 10

Check on Router: after 4 minutes:
 30 second input rate 34430000 bits/sec, 3460 packets/sec
Conclusion: 32.8Mbps max download
```

  - Outbound

```
On NR host:  iperf3 -s
On US host:  iperf3 -c 203.98.X.X -P 40 -t 250 -O 10 -R

Check on Router: after 4 minutes:
 30 second output rate 21117000 bits/sec, 3443 packets/sec
Conclusion: 20.1Mbps max upload
```

Link performance tests (2)

- July 31st – after tuning and bringing up 2nd channel
  - Inbound

```
On NR host:  iperf3 -s
On US host:  iperf3 -c 203.98.X.X -P 40 -t 310 -O 10

Check on Router: after 3 minutes:
 30 second input rate 49818000 bits/sec, 4459 packets/sec
Conclusion: 47.5Mbps max download
```

  - Outbound

```
On NR host:  iperf3 -s
On US host:  iperf3 -c 203.98.X.X -P 40 -t 250 -O 10 -R

Check on Router: after 4 minutes:
 30 second output rate 21832000 bits/sec, 2674 packets/sec
Conclusion: 20.8Mbps max upload
```

Commentary

- Weak signal blamed for lack of throughout
  - C-band on IS-18
  - Sold as 96Mbps (62/34) but delivering only 68Mbps 😭
  - **`iperf3`** was run to US host
    - But same result when rerun to Asia host and Australia host
- Mid-October 2020:
  - Limited capacity & weak signal on IS-18
  - Changeover to Ku-band on IS-1R
  - New BUC shipped (40W) and realigned the dish
  - A huge difference!

Link performance tests (3)

- October 16th – Ku-band on IS-1R
  - Inbound

```
On NR host:  iperf3 -s
On US host:  iperf3 -c 203.98.X.X -P 40 -t 250 -O 10

Check on Router: after 4 minutes:
 30 second input rate 64873000 bits/sec, 6243 packets/sec
Conclusion: 61.9Mbps max download
```

  - Outbound

```
On NR host:  iperf3 -s
On US host:  iperf3 -c 203.98.X.X -P 40 -t 250 -O 10 -R

Check on Router: after 4 minutes:
 30 second output rate 30157000 bits/sec, 3415 packets/sec
Conclusion: 28.8Mbps max upload
```

Total of 90.7Mbps real throughput on a theoretical 96Mbps capacity link - quite acceptable after allowing for overheads

Today…

- Cenpac has redundant satellite links from two providers
  - Speedcast (JCSAT2B)
    - Nauru to Perth (Western Australia)
  - AvComm (IS-1R)
    - Nauru to Sydney
  - BGP is used, tweaked timers for quick failover
  - Independent satellites
  - Upstream operators have different transit providers
    - Both appear at Equinix Sydney
  - This is as about as good as it can get
  - .NR ccTLD has not been off-line since 27th December 2019 😀

Agenda

- Satellite links
- Routing (BGP, IS-IS etc)
- Island-wide WiFi backbone
- Ganeti Cluster
- Router & Switch Infrastructure
- Network Operations
- Observations

CenpacNet – November 2020

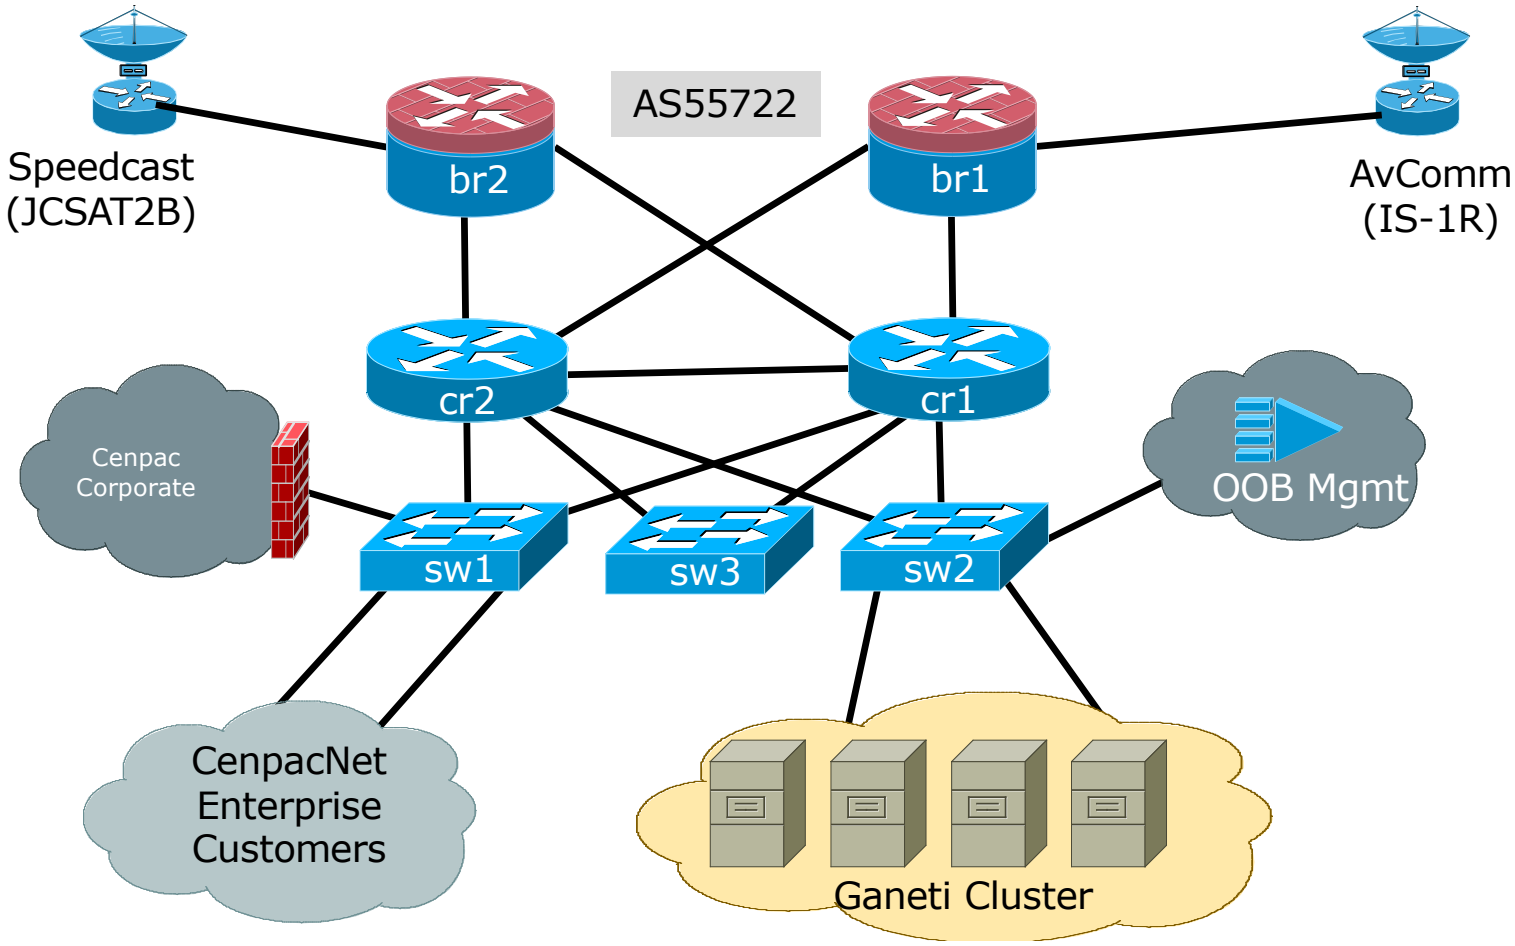

What about EBGP?

- Both satellite links use public IP addresses
  - Private addresses mean the remote border routers are unreachable should BGP fail
  - Each border router has static default route (at high distance) pointing upstream satellite provider
  - BGP failure → both routers are still accessible for remote recovery work as needed
- Both upstreams accept /19 aggregate, and /20 sub-prefixes
  - Av-Comm also accept three specifics for testing purposes

What about EBGP?

- Link from border routers to satellite equipment is ethernet
  - Cannot rely on “interface down” to failover BGP session
- BGP timers changed
  - 10 second keepalive, 30 second holdtime
  - If remote peer disappears for 30 seconds, BGP session is torn down, alternative path is used
  - Outage is around 60 seconds, rather than 5 minutes using default values
  - BFD on a satellite modem?? 😇

EBGP tricks?

- Use Cisco default distances here:
  - EBGP routes are distance 20
  - IBGP routes are distance 200
- Which means each border router has best path default out its own satellite link
  - Backed up by a static default route pointing out the satellite link with distance 254
  - Total BGP failure means remote access is still possible to the public IP address on the WAN side
- IS-IS used to propagate the default route
  - Tagged as **`no-advertise`** community in BGP on border

IS-IS notes

- Anything interesting here?
  - Just usual good practices

```
router isis cenpac
 net 49.0001.0001.0004.00
 is-type level-2-only
 metric-style wide
 set-overload-bit on-startup wait-for-bgp
 log-adjacency-changes
 metric 100000
 passive-interface Loopback0
 default-information originate route-map DEFAULT-ORIG
 !
 address-family ipv6
  multi-topology
  set-overload-bit on-startup wait-for-bgp
  metric 100000
 exit-address-family
```

IS-IS notes

- How do we implement that default originate?
  - IS-IS route-map:

```
route-map DEFAULT-ORIG permit 5
 match ip address prefix-list default
 match ip next-hop BGP-NH
!
```

  - And looking for the correct next-hop in the Global RIB

```
ip access-list standard BGP-NH
 remark Newtec Modem
 permit 110.232.150.113
!
```

  - IS-IS will originate the default route only if the next-hop of the default it sees in the Global RIB is from the EBGP peer
  - Which means this border router will only send default in IS-IS if it hears it from its EBGP peer

Loadbalancing

- Moving target!
  - Nothing new there 😇🙄
- For inbound traffic:
  - Announce /19 and two /20s (no real reason for latter)
    - Prepend towards Speedcast
    - 20Mbps in (Speedcast) vs 62Mbps in (Av-Comm) to try and balance
- For outbound traffic:
  - Default from Speedcast (cannot deliver full table)
  - Default and full routes from Av-Comm
    - 50% discarded (>5 AS hops)
- Summary:
  - Okay for now but will need work as traffic builds

IPv6

- Might have noticed lack of IPv6 mention so far
  - Yet Cenpac first deployed IPv6 in 2010
- Now?
  - Av-Comm cannot provide
  - Speedcast cannot provide
- Local network is dual stack
  - But public services have had their AAAA addresses removed
  - 🤯😤
- Do I really have to look for a GRE tunnel to somewhere?
  - That's so 1990s! 🤯😤

MANRS?

- MANRS is industry initiative led by ISOC to improve routing security globally
  - https://www.manrs.org
- What is Cenpac doing?
  - ROA signing still to be done – hoops to be jumped through
  - ROV is in place
    - Using Routinator 3000 😀😀
    - But default route negates this
  - uRPF on all access interfaces
    - Yes there is some interesting "stuff" coming from customers!
  - peeringDB – to be done

What about IBGP?

- Simple core:
  - 2 border routers (ASR1001)
  - 2 core routers (ASR1002)
- Core are Route Reflectors, Border are clients
- Customer routes go into IBGP
- Aggregates originated in the core
- Standard set up…

Careful what you leak!

- Common problem on Internet today is network operators spraying /24s around at random and claiming they are doing "traffic engineering"
  - Actually they are wrecking their customers' service quality
- While doing performance testing of satellite links:
  - Very slow BGP convergence
  - Withdrawn prefixes being held on to by some operators for >30 mins, blackholing traffic

The "experiment"

- Starting point:
  - Speedcast hears 203.98.224.0/19 le 20 with 2x prepend
  - AvComm hears 203.98.224.0/19 le 20
- The plan (for iperf3 testing)
  - Leak 203.98.252.0/22 via AvComm (filters allow this)
  - Stop the announcement of 203.98.240.0/20 to Av-Comm
- What should happen:
  - All traffic to 203.98.252.0/22 enters via Av-Comm link
  - When 203.98.252.0/22 withdrawn after tests, traffic fails over to Speedcast (following the 203.98.240.0/20 announcement there)
- Sounds simple?

The starting and finishing point – Time = 0

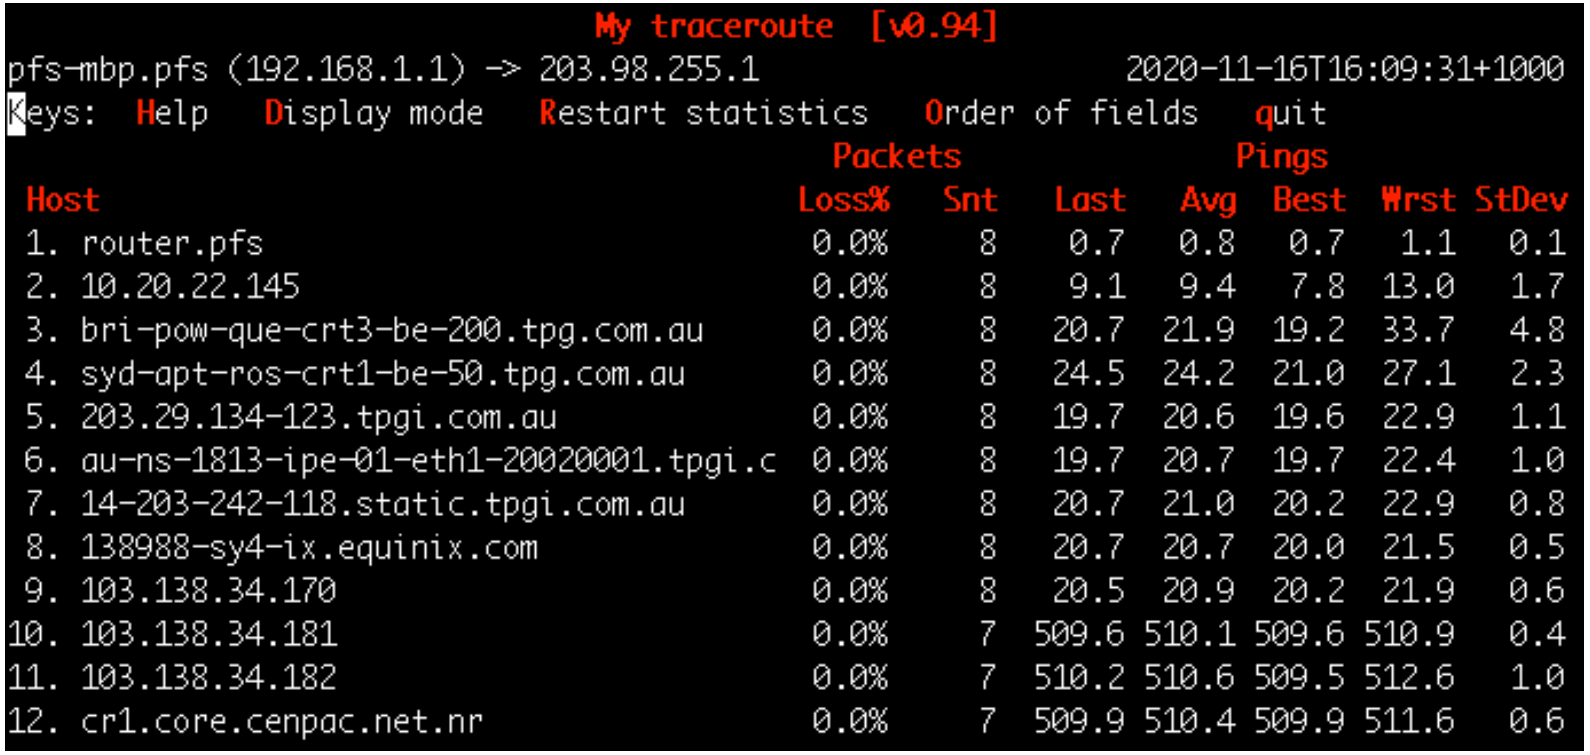

```
                                My traceroute  [v0.94]
pfs-mbp.pfs (192.168.1.1) -> 203.98.255.1                    2020-11-16T16:09:31+1000
Keys:  Help   Display mode   Restart statistics   Order of fields   quit
                                          Packets                Pings
 Host                                   Loss%   Snt   Last   Avg  Best  Wrst StDev
 1. router.pfs                           0.0%     8    0.7   0.8   0.7   1.1   0.1
 2. 10.20.22.145                         0.0%     8    9.1   9.4   7.8  13.0   1.7
 3. bri-pow-que-crt3-be-200.tpg.com.au   0.0%     8   20.7  21.9  19.2  33.7   4.8
 4. syd-apt-ros-crt1-be-50.tpg.com.au    0.0%     8   24.5  24.2  21.0  27.1   2.3
 5. 203.29.134-123.tpgi.com.au           0.0%     8   19.7  20.6  19.6  22.9   1.1
 6. au-ns-1813-ipe-01-eth1-20020001.tpgi.c  0.0%  8   19.7  20.7  19.7  22.4   1.0
 7. 14-203-242-118.static.tpgi.com.au    0.0%     8   20.7  21.0  20.2  22.9   0.8
 8. 138988-sy4-ix.equinix.com            0.0%     8   20.7  20.7  20.0  21.5   0.5
 9. 103.138.34.170                       0.0%     8   20.5  20.9  20.2  21.9   0.6
10. 103.138.34.181                       0.0%     7  509.6 510.1 509.6 510.9   0.4
11. 103.138.34.182                       0.0%     7  510.2 510.6 509.5 512.6   1.0
12. cr1.core.cenpac.net.nr               0.0%     7  509.9 510.4 509.9 511.6   0.6
```

Before we start, best path is via Av-Comm 😀

Time +1minute

My traceroute [v0.94]

pfs-mbp.pfs (192.168.1.1) -> 203.98.255.1 2020-11-16T15:38:46+1000

Keys: Help Display mode Restart statistics Order of fields quit

	Host	Packets Loss%	Snt	Pings Last	Avg	Best	Wrst	StDev
1.	router.pfs	0.0%	77	0.7	0.9	0.6	3.1	0.4
2.	10.20.22.145	0.0%	77	7.7	9.1	7.6	15.2	1.3
3.	bri-pow-que-crt3-be-200.tpg.com.au	0.0%	77	20.0	20.6	19.4	25.1	0.9
4.	syd-apt-ros-crt1-be-50.tpg.com.au	0.0%	77	28.2	24.3	19.8	32.4	2.7
5.	203.29.134-123.tpgi.com.au	1.3%	77	21.2	21.3	19.2	41.5	3.5
6.	203.29.134-68.tpgi.com.au	0.0%	77	23.2	22.7	20.0	27.2	1.6
7.	be18-140.br02.hkg15.pccwbtn.net	0.0%	77	163.6	162.7	161.4	166.1	0.9
8.	hundredge0-0-0-24.br04.hkg12.pccwbtn.n	3.9%	77	164.7	163.2	162.0	172.9	1.4
9.	be3492.rcr51.hkg01.atlas.cogentco.com	0.0%	77	250.2	250.6	249.2	253.8	0.8
10.	be2414.ccr21.hkg02.atlas.cogentco.com	0.0%	77	274.0	273.7	272.3	278.4	1.2
11.	be3691.ccr21.sin01.atlas.cogentco.com	0.0%	77	437.1	436.5	435.1	440.3	0.9
12.	be2428.ccr51.per01.atlas.cogentco.com	0.0%	76	487.7	487.6	486.3	490.4	0.8
13.	be2429.ccr51.syd01.atlas.cogentco.com	0.0%	76	311.6	309.9	308.5	313.3	0.9
14.	be2092.agr12.syd01.atlas.cogentco.com	0.0%	76	334.0	334.5	333.1	339.0	1.3
15.	154.18.96.106	0.0%	76	153.3	162.9	152.1	355.8	32.2
16.	138988-sy4-ix.equinix.com	0.0%	76	152.3	153.2	149.5	158.5	1.1
17.	103.138.34.170	0.0%	76	156.1	153.5	152.2	158.4	1.0
18.	103.138.34.181	0.0%	76	642.0	642.9	641.2	646.9	1.1
19.	103.138.34.182	0.0%	76	642.7	643.1	641.8	646.7	0.9
20.	cr1.core.cenpac.net.nr	0.0%	76	643.6	643.4	641.4	647.5	1.2

This should not have happened.

203.98.252.0/22 was leaked to the current best path provider. It is globally visible.

Why did the path change??

Note: Brisbane → Sydney → Hong Kong → Singapore → Perth → Sydney → Nauru 🤪

Do operators actually peer in Sydney??

Time +15 minutes

```
route-views>sh ip bgp 203.98.240.0/20 | i 55722
  1221 4637 38456 38456 38456 5666 55722 55722 55722
<snip>
  3257 4826 38456 5666 55722 55722 55722
  701 1299 7545 7545 17559 138988 55722          <--------------
  852 6939 4826 38456 5666 55722 55722 55722
  6939 4826 38456 5666 55722 55722 55722
  1239 1299 7545 7545 17559 138988 55722         <--------------
  4901 6079 6939 4826 38456 5666 55722 55722 55722
<snip>
  24441 4826 38456 5666 55722 55722 55722
  3333 1273 1299 7545 7545 17559 138988 55722    <--------------
  2497 4826 38456 5666 55722 55722 55722
  7018 1299 7545 7545 17559 138988 55722         <--------------
  53767 6939 4826 38456 5666 55722 55722 55722
<snip>
```

This should not be happening!!

The highlighted paths were withdrawn 15 minutes ago. Yet still visible at RouteViews!

What is AS7545 doing??

Time +20 minutes

```
                    My traceroute  [v0.94]
pfs-mbp.pfs (192.168.1.1) -> 203.98.255.1                2020-11-16T15:52:34+1000
Keys:  Help   Display mode   Restart statistics   Order of fields   quit
                                 Packets               Pings
 Host                          Loss%   Snt   Last   Avg  Best  Wrst StDev
 1. router.pfs                  0.0%    12    0.7   0.8   0.6   1.3   0.2
 2. 10.20.22.145                0.0%    12    8.4   8.5   7.5  10.9   0.9
 3. bri-pow-que-crt3-be-200.tpg.com.au  0.0%  11  20.4  20.6  19.5  27.1  2.2
 4. syd-apt-ros-crt1-be-50.tpg.com.au   0.0%  11  25.6  24.0  21.0  27.1  2.1
 5. 203.29.134-123.tpgi.com.au  0.0%    11   20.0  20.4  19.5  24.8   1.5
 6. 203.29.134-68.tpgi.com.au   0.0%    11   21.9  22.7  19.9  27.7   2.3
 7. be18-140.br02.hkg15.pccwbtn.net  0.0%  11  161.9 162.7 161.7 165.8  1.1
 8. hundredge0-0-0-24.br04.hkg12.pccwbtn.n  0.0%  11  162.1 163.1 161.9 166.3  1.3
 9. be3492.rcr51.hkg01.atlas.cogentco.com  0.0%  11  251.2 250.9 249.4 254.2  1.3
10. be2414.ccr21.hkg02.atlas.cogentco.com  0.0%  11  275.5 273.7 272.8 275.5  0.9
11. be3691.ccr21.sin01.atlas.cogentco.com  0.0%  11  435.7 437.9 435.5 454.7  5.6
12. be2428.ccr51.per01.atlas.cogentco.com  0.0%  11  487.1 487.3 486.9 488.3  0.4
13. be2429.ccr51.syd01.atlas.cogentco.com  0.0%  11  309.2 309.6 308.7 310.4  0.5
14. be2092.agr12.syd01.atlas.cogentco.com  0.0%  11  334.2 334.1 332.8 336.5  1.1
15. 154.18.96.106              0.0%    11   152.9 166.7 152.5 299.9  44.2
16. as4826.nsw.ix.asn.au       0.0%    11   155.1 154.1 153.0 155.8   0.9
17. be108.cor02.syd04.nsw.vocus.network  0.0%  11  301.8 301.3 300.6 303.0  0.6
18. be202.bdr03.sjc01.ca.us.vocus.network  9.1%  11  303.6 301.9 301.0 303.6  0.8
19. 173.205.35.1               0.0%    11   306.0 307.3 306.0 312.0   1.7
20. ae1.cr3-sjc1.ip4.gtt.net   0.0%    11   306.5 307.0 305.4 310.5   1.5
21. be3258.ccr41.sjc03.atlas.cogentco.com  0.0%  11  307.1 308.9 305.3 330.7  7.3
22. be3670.ccr22.sfo01.atlas.cogentco.com  0.0%  11  309.5 306.5 305.2 310.0  1.7
23. be3694.ccr21.pdx01.atlas.cogentco.com  0.0%  11  349.4 350.4 349.4 351.5  0.7
24. be2216.ccr51.pdx02.atlas.cogentco.com  0.0%  11  350.6 351.9 350.5 355.1  1.4
25. be2237.ccr51.syd01.atlas.cogentco.com  0.0%  11  459.9 459.2 458.3 461.4  0.9
26. be2092.agr12.syd01.atlas.cogentco.com  0.0%  11  485.5 484.3 483.0 485.5  1.0
27. 154.18.96.106              0.0%    11   302.2 317.3 302.1 448.6  46.1
28. as4826.nsw.ix.asn.au       0.0%    11   306.3 305.4 302.6 314.6   3.4
29. be108.cor02.syd04.nsw.vocus.network  0.0%  11  451.6 451.4 450.3 452.4  0.5
30. be202.bdr03.sjc01.ca.us.vocus.network  20.0%  11  451.5 451.0 450.1 452.2  0.8
```

203.98.252.0/22 now withdrawn from AvComm link

As expected, reachability vanishes, and there is a huge routing loop between Australia and the entire US

Brisbane → Sydney → Hong Kong → Singapore → Perth → Sydney → Equinix Sydney → San Jose → San Francisco → Portland → Sydney → Equinix Sydney → San Jose → …

This should clear in a minute or two, right?

Time +22 minutes

Time +23 minutes

My traceroute [v0.94]

pfs-mbp.pfs (192.168.1.1) -> 203.98.255.1 2020-11-16T15:54:42+1000

Keys: Help Display mode Restart statistics Order of fields quit

	Host	Packets Loss%	Snt	Pings Last	Avg	Best	Wrst	StDev
1.	router.pfs	0.0%	6	0.9	0.8	0.7	1.0	0.1
2.	10.20.22.145	0.0%	6	7.9	8.5	7.7	9.2	0.7
3.	bri-pow-que-crt3-be-200.tpg.com.au	0.0%	6	21.6	21.1	19.5	25.0	2.1
4.	syd-apt-ros-crt1-be-50.tpg.com.au	0.0%	6	20.2	23.2	20.2	27.1	3.0
5.	203.29.134-123.tpgi.com.au	0.0%	6	22.2	20.8	19.6	22.2	0.8
6.	au-ns-1813-ipe-01-eth1-20020001.tpgi.c	0.0%	6	20.7	20.3	19.6	20.8	0.5
7.	14-203-242-118.static.tpgi.com.au	0.0%	6	37.0	23.6	20.3	37.0	6.6
8.	as4826.nsw.ix.asn.au	0.0%	6	21.9	22.2	21.3	23.6	0.8
9.	as38456.bdr01.syd03.nsw.vocus.net.au	0.0%	6	20.7	21.4	20.7	22.5	0.6
10.	(waiting for reply)							
11.	spd-0004.10026.telstraglobal.net	0.0%	6	24.6	22.6	21.3	24.6	1.2
12.	i-91.sydp-core04.telstraglobal.net	0.0%	5	23.4	24.2	23.4	24.6	0.5
13.	i-20802.eqnx-core02.telstraglobal.net	0.0%	5	183.9	184.6	183.7	186.1	1.0
14.	i-92.eqnx03.telstraglobal.net	0.0%	5	184.1	194.5	183.0	228.0	19.3
15.	(waiting for reply)							
16.	64.124.57.71.ipyx-284545-900-zyo.zip.z	0.0%	5	183.5	182.9	182.5	183.5	0.4
17.	(waiting for reply)							
18.	64.124.57.71.ipyx-284545-900-zyo.zip.z	0.0%	5	183.5	189.4	182.4	213.5	13.5
19.	(waiting for reply)							
20.	64.124.57.71.ipyx-284545-900-zyo.zip.z	0.0%	5	184.7	183.7	182.4	185.5	1.4
21.	(waiting for reply)							
22.	64.124.57.71.ipyx-284545-900-zyo.zip.z	0.0%	5	183.7	184.4	182.7	186.5	1.7
23.	(waiting for reply)							
24.	64.124.57.71.ipyx-284545-900-zyo.zip.z	0.0%	5	183.7	184.0	183.1	187.2	1.8
25.	(waiting for reply)							
26.	64.124.57.71.ipyx-284545-900-zyo.zip.z	0.0%	5	182.9	183.5	182.9	183.9	0.4
27.	(waiting for reply)							
28.	64.124.57.71.ipyx-284545-900-zyo.zip.z	0.0%	5	183.1	183.5	183.0	184.6	0.6
29.	(waiting for reply)							
30.	64.124.57.71.ipyx-284545-900-zyo.zip.z	0.0%	5	182.9	183.5	182.7	185.1	1.0

My traceroute [v0.94]

pfs-mbp.pfs (192.168.1.1) -> 203.98.255.1 2020-11-16T15:55:29+1000

Keys: Help Display mode Restart statistics Order of fields quit

	Host	Packets Loss%	Snt	Pings Last	Avg	Best	Wrst	StDev
1.	router.pfs	0.0%	17	0.7	0.8	0.6	1.0	0.1
2.	10.20.22.145	0.0%	17	9.2	8.7	7.5	13.5	1.3
3.	bri-pow-que-crt3-be-200.tpg.com.au	0.0%	17	19.6	20.2	19.6	21.8	0.5
4.	syd-apt-ros-crt1-be-50.tpg.com.au	0.0%	17	23.3	25.1	20.3	30.0	2.8
5.	203.29.134-123.tpgi.com.au	0.0%	17	20.0	20.6	19.2	22.7	0.9
6.	au-ns-1813-ipe-01-eth1-20020001.tpgi.c	5.9%	17	20.1	20.6	19.7	21.4	0.5
7.	14-203-242-118.static.tpgi.com.au	0.0%	17	20.1	38.2	19.8	116.9	37.3
8.	as4826.nsw.ix.asn.au	0.0%	16	24.6	21.8	20.4	25.1	1.4
9.	as38456.bdr01.syd03.nsw.vocus.net.au	0.0%	16	21.2	21.7	20.6	25.2	1.1
10.	(waiting for reply)							
11.	spd-0004.10026.telstraglobal.net	0.0%	16	21.5	22.0	21.1	23.3	0.6
12.	i-91.sydp-core04.telstraglobal.net	0.0%	16	24.0	24.2	22.9	26.1	0.8
13.	i-20802.eqnx-core02.telstraglobal.net	0.0%	16	184.6	184.6	183.0	186.9	1.1
14.	i-14808.ny8a-core01.telstraglobal.net	0.0%	16	247.5	247.5	245.6	250.5	1.4
15.	i-14803.ny8a-core01.telstraglobal.net	0.0%	16	295.7	295.6	294.8	296.6	0.5
16.	unknown.telstraglobal.net	0.0%	16	294.7	295.1	294.0	297.5	1.0
17.	ldn-b7-link.telia.net	0.0%	16	304.1	301.0	298.8	311.5	3.3
18.	ldn-bb4-link.telia.net	26.7%	16	309.4	308.0	306.9	309.8	0.9
19.	nyk-bb3-link.telia.net	37.5%	16	291.9	292.7	291.3	295.7	1.4
20.	rest-bb1-link.telia.net	0.0%	16	322.5	321.9	320.6	325.3	1.3
21.	nyk-bb3-link.telia.net	62.5%	16	310.9	307.7	306.6	310.9	1.6
22.	palo-b24-link.telia.net	18.8%	16	306.2	306.5	305.8	308.3	0.8
23.	las-b24-link.telia.net	0.0%	16	312.8	307.8	306.5	312.8	1.6
24.	vocus-ic-324263-las-b24.c.telia.net	26.7%	16	316.3	316.4	315.9	317.3	0.4
25.	be101.bdr04.sjc01.ca.us.vocus.network	26.7%	16	315.9	316.2	315.4	318.5	1.0
26.	be203.cor01.syd11.nsw.vocus.network	0.0%	16	315.6	316.2	315.6	317.3	0.5
27.	be100.bdr01.syd03.nsw.vocus.network	0.0%	16	318.8	317.1	315.9	321.8	1.5
28.	as38456.bdr01.syd03.nsw.vocus.net.au	0.0%	16	317.1	316.9	315.9	322.2	1.5
29.	(waiting for reply)							
30.	spd-0004.10026.telstraglobal.net	0.0%	16	317.8	317.2	316.5	318.3	0.6

Time +24 minutes

```
route-views>sh ip bgp 203.98.252.0 | i 55722
  8283 57866 3491 174 174 174 17559 138988 55722
  1221 4637 1299 174 174 174 17559 138988 55722
  3333 12859 2914 174 174 174 17559 138988 55722
```

203.98.252.0/22 was withdrawn 4 minutes ago

The world did the right thing, apart from AS174. Why are these still here?

There is still no connectivity to/from the 203.98.252.0/22 address block. The routing loop from Time +23 minutes has stabilised

Time +30 minutes

```
route-views>sh ip bgp 203.98.240.0/20 | i 55722
  1221 4637 38456 38456 38456 5666 55722 55722 55722
<snip>
  3257 4826 38456 5666 55722 55722 55722
  701 1299 7545 7545 17559 138988 55722    <--------------
  852 6939 4826 38456 5666 55722 55722 55722
  6939 4826 38456 5666 55722 55722 55722
  1239 1299 7545 7545 17559 138988 55722    <--------------
  4901 6079 6939 4826 38456 5666 55722 55722 55722
<snip>
  24441 4826 38456 5666 55722 55722 55722
  3333 1273 1299 7545 7545 17559 138988 55722    <--------------
  2497 4826 38456 5666 55722 55722 55722
  7018 1299 7545 7545 17559 138988 55722    <--------------
  53767 6939 4826 38456 5666 55722 55722 55722
<snip>
```

This definitely should not be happening!!

The highlighted paths were withdrawn 30 minutes ago. Yet still visible at RouteViews!

What is AS7545 really doing?? Do they use BGP?

Time +35 minutes

```
                                My traceroute  [v0.94]
pfs-mbp.pfs (192.168.1.1) -> 203.98.255.1                    2020-11-16T16:05:33+1000
Keys:  Help   Display mode   Restart statistics   Order of fields   quit
                                          Packets               Pings
 Host                                   Loss%   Snt   Last   Avg  Best  Wrst StDev
 1. router.pfs                           0.0%     8    0.7   0.8   0.5   1.0   0.2
 2. 10.20.22.145                         0.0%     8    8.4   8.9   7.8  11.2   1.1
 3. bri-pow-que-crt3-be-200.tpg.com.au   0.0%     8   19.9  20.0  19.1  21.0   0.6
 4. syd-apt-ros-crt1-be-50.tpg.com.au    0.0%     8   22.7  23.5  20.3  27.4   2.3
 5. 203.29.134-123.tpgi.com.au           0.0%     8   20.4  20.4  19.6  20.9   0.4
 6. au-ns-1813-ipe-01-eth1-20020001.tpgi.c 0.0%   8   20.0  20.4  19.8  21.5   0.6
 7. 14-203-242-118.static.tpgi.com.au    0.0%     8   22.2  21.0  20.0  22.4   0.8
 8. as4826.nsw.ix.asn.au                 0.0%     8   22.5  22.3  20.7  27.8   2.3
 9. as38456.bdr01.syd03.nsw.vocus.net.au 0.0%     8   21.0  21.4  21.0  22.3   0.5
10. (waiting for reply)
11. (waiting for reply)
12. (waiting for reply)
13. 110-232-150-86.pacificteleports.net  0.0%     7  582.5 582.5 569.2 591.6   8.1
14. cr1.core.cenpac.net.nr               0.0%     7  587.0 580.5 567.6 608.2  14.8
```

Finally we have connectivity to 203.98.252.0/22 address range again.

The global routing system has converged on the 203.98.240.0/20 address block being announced only to Speedcast

That has taken 15 minutes exactly!

Is BGP really this slow???

Time +35 minutes

```
route-views>sh ip bgp 203.98.252.0 | i 55722
  1221 4637 1299 174 174 174 17559 138988 55722
```

203.98.252.0/22 was withdrawn 15 minutes ago

This bogus path is still hanging on in there. In fact, it took a further 5 minutes to disappear.

- This meant that 203.98.252.0/22 was **totally** offline for 15 minutes
  - No connectivity from several global vantages points – all showed lengthy routing loops
  - Paths all showed AS174 holding on to the prefix even though it had long been withdrawn

Time +40 minutes

My traceroute [v0.94]

pfs-mbp.pfs (192.168.1.1) -> 203.98.255.1 2020-11-16T16:09:31+1000

Keys: Help Display mode Restart statistics Order of fields quit

	Host	Packets Loss%	Snt	Pings Last	Avg	Best	Wrst	StDev
1.	router.pfs	0.0%	8	0.7	0.8	0.7	1.1	0.1
2.	10.20.22.145	0.0%	8	9.1	9.4	7.8	13.0	1.7
3.	bri-pow-que-crt3-be-200.tpg.com.au	0.0%	8	20.7	21.9	19.2	33.7	4.8
4.	syd-apt-ros-crt1-be-50.tpg.com.au	0.0%	8	24.5	24.2	21.0	27.1	2.3
5.	203.29.134-123.tpgi.com.au	0.0%	8	19.7	20.6	19.6	22.9	1.1
6.	au-ns-1813-ipe-01-eth1-20020001.tpgi.c	0.0%	8	19.7	20.7	19.7	22.4	1.0
7.	14-203-242-118.static.tpgi.com.au	0.0%	8	20.7	21.0	20.2	22.9	0.8
8.	138988-sy4-ix.equinix.com	0.0%	8	20.7	20.7	20.0	21.5	0.5
9.	103.138.34.170	0.0%	8	20.5	20.9	20.2	21.9	0.6
10.	103.138.34.181	0.0%	7	509.6	510.1	509.6	510.9	0.4
11.	103.138.34.182	0.0%	7	510.2	510.6	509.5	512.6	1.0
12.	cr1.core.cenpac.net.nr	0.0%	7	509.9	510.4	509.9	511.6	0.6

One minute earlier, the configurations were returned to the status quo:

Speedcast get 203.98.224.0/19 with 2x prepend

AvComm get 203.98.224.0/19 le 20

It took under one minute for the path from Brisbane to Nauru to change from via Speedcast to go via AvComm

Careful what you leak!

- Advice: never leak a sub-prefix on a single path
  - If the path goes, it'll take 15+ minutes for backup via covering aggregate 🤯😤🤬
  - Some providers are holding on to BGP prefixes even if they have been withdrawn
    - Why??
    - This causes routing loops
    - Breaks connectivity for the affected networks

- Questions (rhetorical!):
  - Why is it fashionable for operators to leak specifics on single paths when this is the result?
  - Do operators actually know how BGP works?

Bizarre Behaviour

- How is this for a traceroute:

```
br2.core#trace shell.nsrc.org
Type escape sequence to abort.
Tracing the route to shell.nsrc.org (128.223.157.52)
VRF info: (vrf in name/id, vrf out name/id)
  1 110-232-150-113.pacificteleports.net (110.232.150.113) [AS 64565] 1 msec 0 msec 0 msec
  2 shell.nsrc.org (128.223.157.52) [AS 5666] 566 msec 554 msec 546 msec
  3 169.254.4.1 [AS 5666] 561 msec 573 msec 566 msec
  4 100.80.0.1 [AS 5666] 559 msec 558 msec 538 msec
  5 10.151.164.74 [AS 5666] [MPLS: Label 5443 Exp 1] 591 msec 619 msec 609 msec
  6 10.151.164.76 [AS 5666] 599 msec 612 msec 595 msec
  7 SPD-0004.10026.telstraglobal.net (203.192.130.57) [AS 4637] 604 msec 597 msec 598 msec
```

NewTec Modem "feature"

- And in the other direction:

```
philip@pfs-mbp ~ % traceroute 203.98.248.1
traceroute to 203.98.248.1 (203.98.248.1), 64 hops max, 52 byte packets
 1  router (192.168.1.254)  1.036 ms  0.673 ms  0.642 ms
 2  10.20.22.145 (10.20.22.145)  8.418 ms  8.566 ms  8.127 ms
 <snip>
13  as38456.bdr01.syd03.nsw.vocus.net.au (175.45.107.6)  22.019 ms  23.761 ms  24.273 ms
 <snip>
17  host1-248.cenpac.net.nr (203.98.248.1)  71.634 ms  69.832 ms  71.471 ms
18  * * *
19  110-232-150-118.pacificteleports.net (110.232.150.118)  574.427 ms *  591.623 ms
```

NewTec Modem "feature"

Agenda

- Satellite links
- Routing (BGP, IS-IS etc)
- Island-wide WiFi backbone
- Ganeti Cluster
- Router & Switch Infrastructure
- Network Operations
- Observations

Island WiFi Backbone

- Cenpac purchased the AccLinks WiFi backbone
  - Aim is to provide Island-wide enterprise internet access
- Equipment:
  - Ubiquiti PowerBeam 5AC & airFiber 5U
  - Configured and (mostly) operational
- But:
  - Management IP addresses on 192.168.1.0/24 and 10.20.0.0/16 address block on VLAN that was also used at Cenpac for OOB access to satellite modems!
  - How to get access from Brisbane??

AccLink WiFi Network

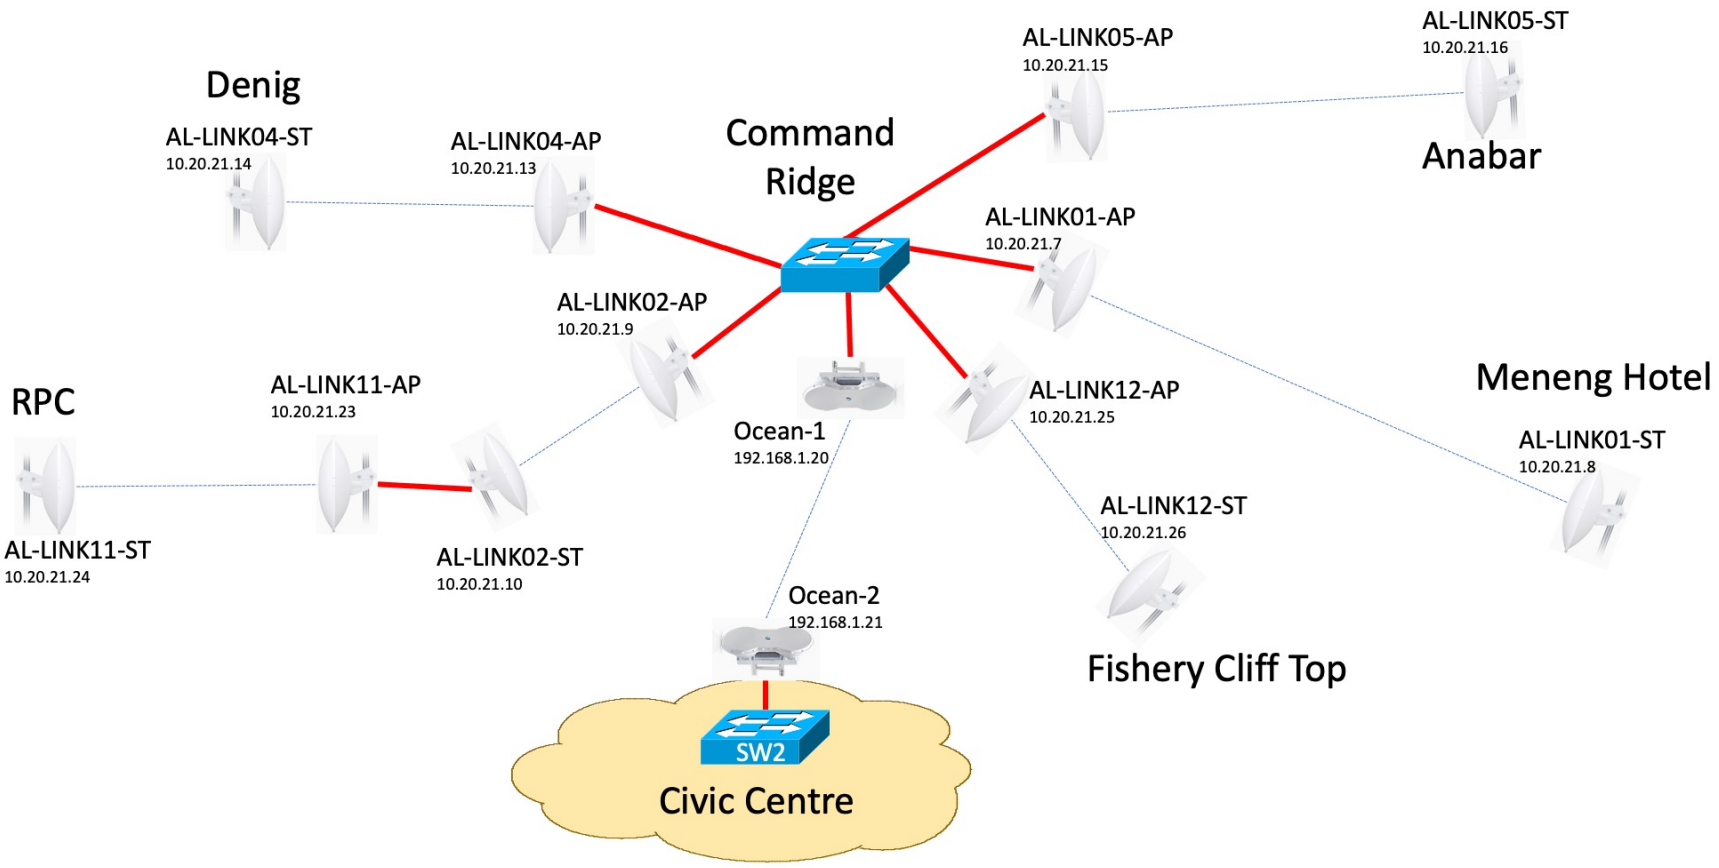

Reconfiguring the WiFi backbone

- Cenpac has no ”VPN service”
- Solution:
  - SSH proxy via their Jumphost
  - Switch terminating WiFi link was not connected to satellite modem
  - dot1Q interface on router introducing 10.20.0.0/16 into BGP
  - Another dot1Q interface on router introducing 192.168.1.0/24 into BGP
  - VLANs trunked to the switches hosting the Ubiquiti devices
  - WiFi management addresses now accessible in Cenpac’s core network
  - Web browser access from @home (192.168/16 net) via Jumphost
- Anyone tried this over a 550ms RTT link recently?

Reconfiguring the WiFi backbone

- Problem:

```
philip@pfs-mbp ~ % ifconfig en12
en12: flags=8863<UP,BROADCAST,SMART,RUNNING,SIMPLEX,MULTICAST> mtu 1500
            options=40b<RXCSUM,TXCSUM,VLAN_HWTAGGING,CHANNEL_IO>
            ether 64:4b:f0:10:07:5b
            inet 192.168.1.1 netmask 0xffffff00 broadcast 192.168.1.255
<snip>
            media: autoselect (1000baseT <full-duplex>)
            status: active
```

  - How to access 192.168.1.0/24 net on Nauru??

- Solution – use home WiFi subnet!!

```
philip@pfs-mbp ~ % ifconfig en0
en0: flags=8863<UP,BROADCAST,SMART,RUNNING,SIMPLEX,MULTICAST> mtu 1500
            options=400<CHANNEL_IO>
            ether f8:ff:c2:22:f5:28
            inet 192.168.4.1 netmask 0xffffff00 broadcast 192.168.4.255
<snip>
            media: autoselect
            status: active
```

Reconfiguring the WiFi backbone – Task List

- Software updates:
  - All APs were running the code they shipped with in 2016 & 2017
  - Many caveats!! First upgrade did local AP first, but then link did not return
    - Had to downgrade the local AP to regain access
    - Read Release Notes and user experiences online 😇
    - Find interim compatible code and upgrade the remote – cross fingers!
    - Then upgrade local to interim compatible
    - And then complete the upgrade to current code! Phew!
  - Seems like some updates were not compatible with older code
  - Solution: tread carefully, read Release Notes FIRST, upgrade remote AP, then local, etc.

Reconfiguring the WiFi backbone – Task List

- Country was set to Russia (??)
  - Changed to Australia – just because
- Channel choices:
  - Reviewed due to bad interference on some links
  - Channel widths not as documented – APs won't associate with mismatch
- Username/passwords changed
- Management access made consistently a separate VLAN
- SSH password access disabled, keys uploaded
- NTP configured
- NMS access set up (for LibreNMS)
- Added to Nagios & Smokeping

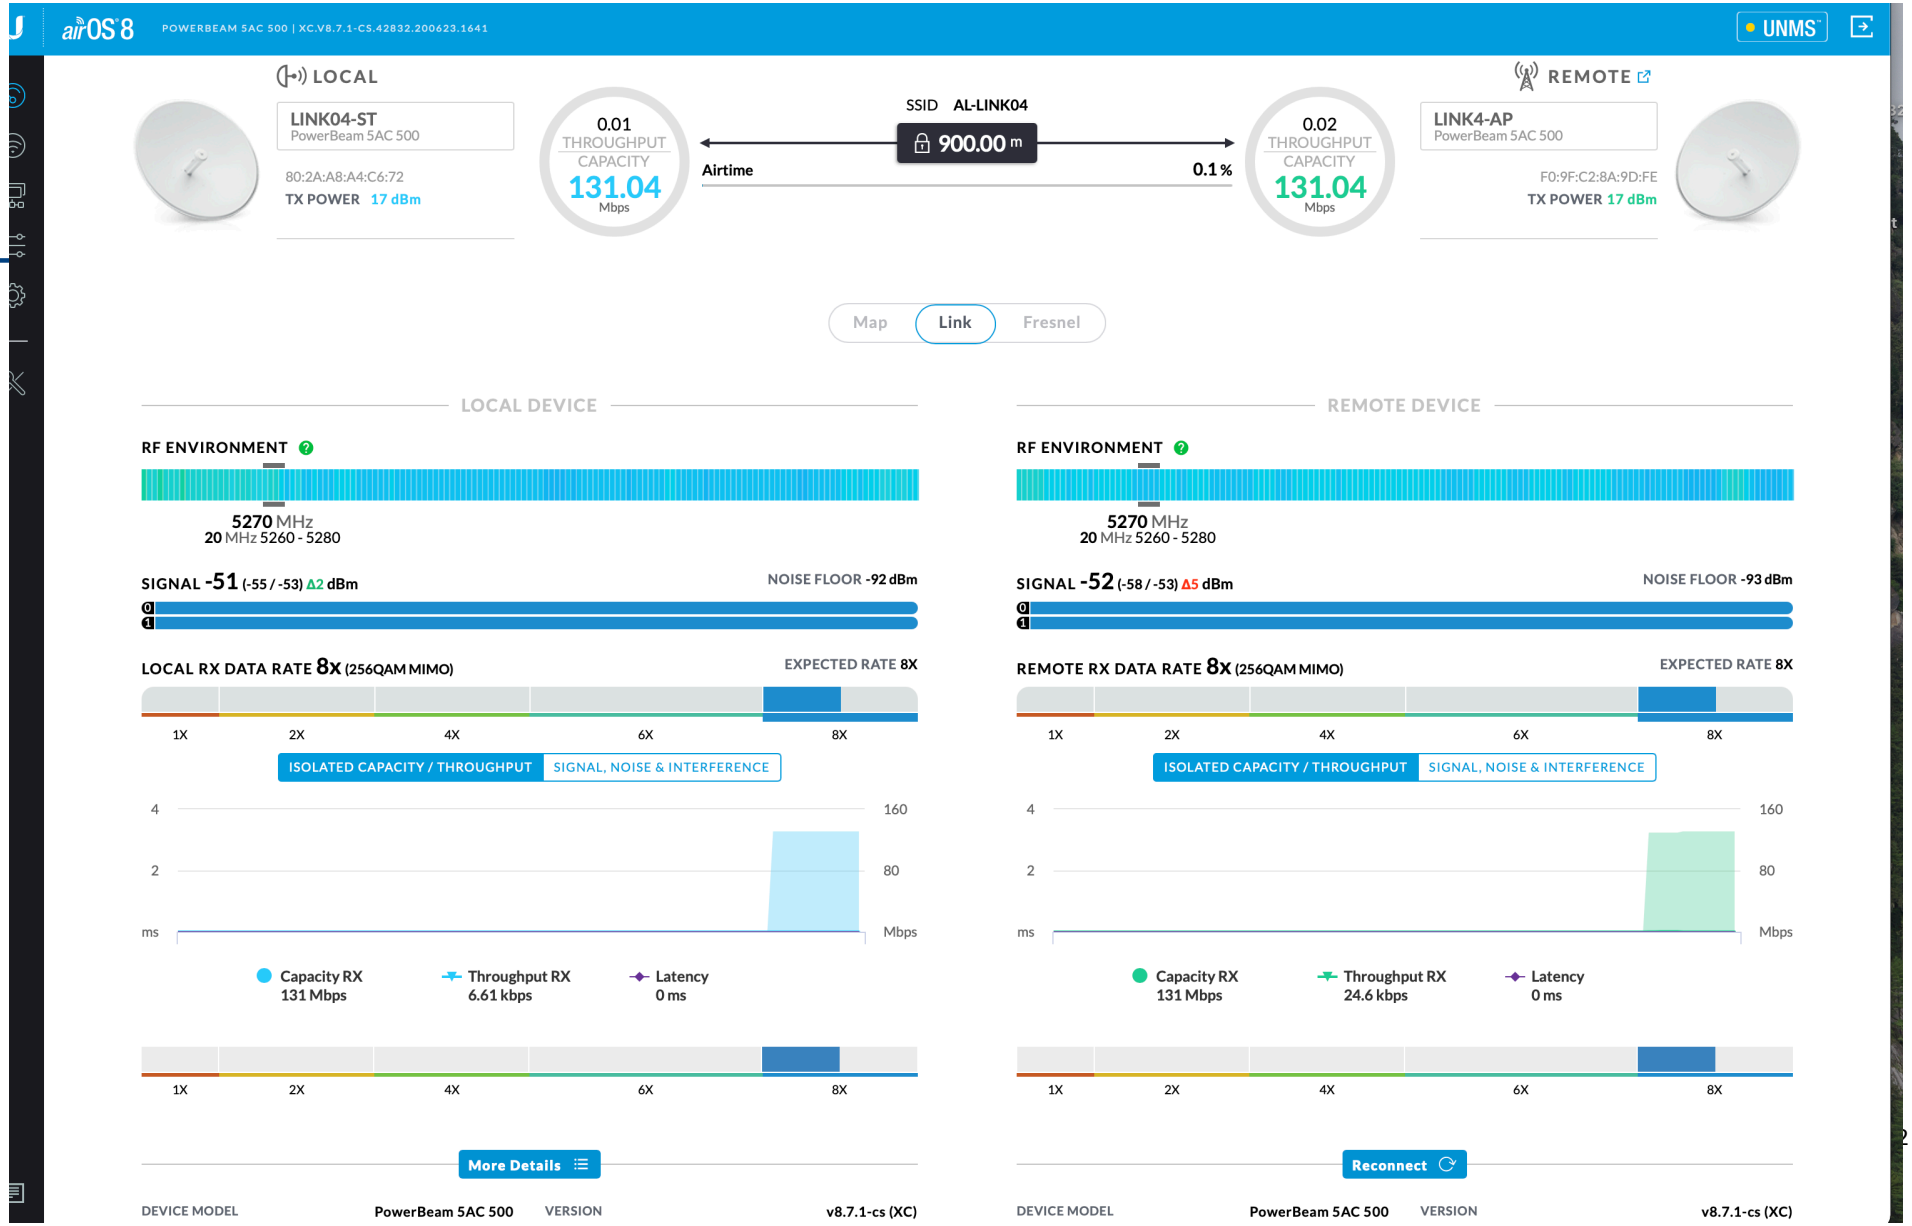
airOS 8
POWERBEAM 5AC 500 | XC.V8.7.1-CS.42832.200623.1641
UNMS
LOCAL
LINK04-ST
PowerBeam 5AC 500
80:2A:A8:A4:C6:72
TX POWER 17 dBm
0.01
THROUGHPUT CAPACITY
131.04
Mbps
SSID AL-LINK04
900.00 m
Airtime
0.1 %
0.02
THROUGHPUT CAPACITY
131.04
Mbps
REMOTE
LINK4-AP
PowerBeam 5AC 500
F0:9F:C2:8A:9D:FE
TX POWER 17 dBm
Map
Link
Fresnel
LOCAL DEVICE
RF ENVIRONMENT
5270 MHz
20 MHz 5260 - 5280
SIGNAL -51 (-55 / -53) Δ2 dBm
NOISE FLOOR -92 dBm
LOCAL RX DATA RATE 8X (256QAM MIMO)
EXPECTED RATE 8X
1X
2X
4X
6X
8X
ISOLATED CAPACITY / THROUGHPUT
SIGNAL, NOISE & INTERFERENCE
4
2
ms
160
80
Mbps
Capacity RX
131 Mbps
Throughput RX
6.61 kbps
Latency
0 ms
More Details
DEVICE MODEL
PowerBeam 5AC 500
VERSION
v8.7.1-cs (XC)
REMOTE DEVICE
RF ENVIRONMENT
5270 MHz
20 MHz 5260 - 5280
SIGNAL -52 (-58 / -53) Δ5 dBm
NOISE FLOOR -93 dBm
REMOTE RX DATA RATE 8X (256QAM MIMO)
EXPECTED RATE 8X
Capacity RX
131 Mbps
Throughput RX
24.6 kbps
Latency
0 ms
Reconnect
DEVICE MODEL
PowerBeam 5AC 500
VERSION
v8.7.1-cs (XC)

Reconfiguring the WiFi Trunk

- Main airFiber trunk from Command Ridge (highest point on Nauru, above Civic Centre) to Civic Centre (Cenpac HQ) was set to 10MHz channel width and half-duplex
  - Giving about 65Mbps throughput, 1.6ms latency, 0.6ms jitter
- Changed from half-duplex to full-duplex means using different transmit and receive channels
  - No change in the 65Mbps throughput, but now 1ms latency, 0.1ms jitter
  - Picked channels 40MHz apart to allow for future change to 40MHz channel width
  - **NB: Change remote device first!**

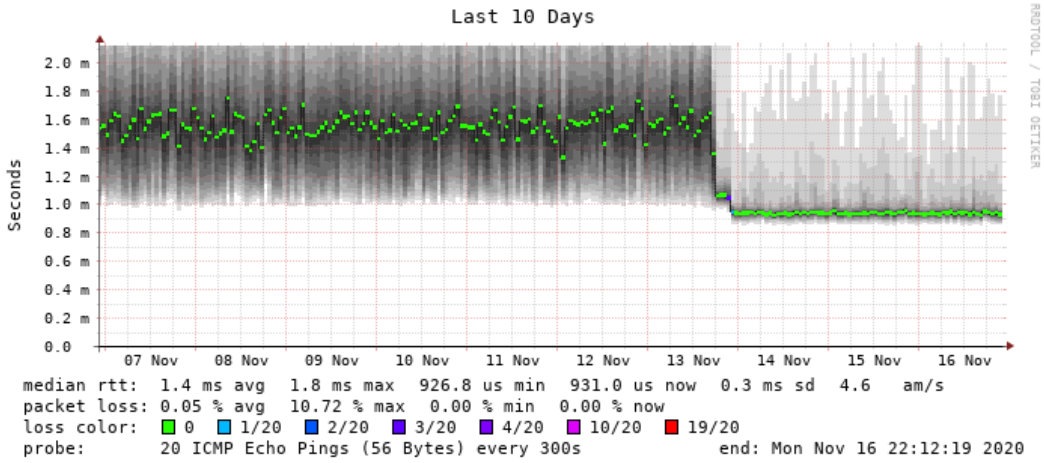

Reconfiguring the WiFi Trunk

- Now to increase the bandwidth available between Command Ridge and Civic Centre
  - Change remote device first, then the local one
- Changed to 20MHz channel width
  - Had to change transmit and receive frequencies as airFiber was using top end of the permitted spectrum
  - Resulted in 130Mbps throughput
- Changed to 40MHz channel width
  - Had to change transmit and receive frequencies again!
  - Resulted in 200Mbps throughput
  - Latency now 900ns, jitter around 50ns

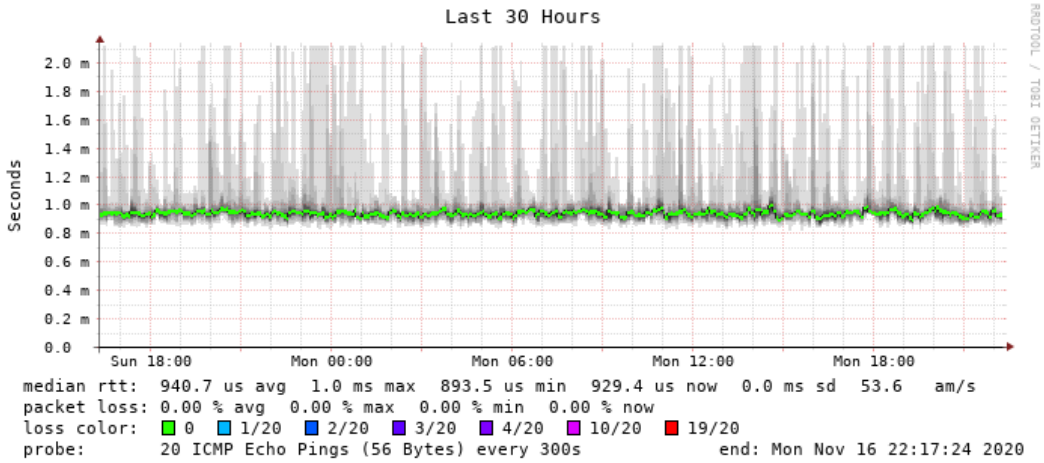

airFiber AF-5U V4.0.5

LOCAL

DEVICE NAME
LINK1-AP

MAC ADDRESS
24:A4:3C:38:A0:DD

LINK NAME CR-WL-1

1.120 km (3675 ft)

RECEIVE CAPACITY
200.71 Mbps
TX POWER (EIRP) 36 dBm

RX FREQUENCY
5790 MHz

RX CHANNEL WIDTH
40MHz

MASTER

TX FREQUENCY
5830 MHz

TX CHANNEL WIDTH
40MHz

SLAVE

RECEIVE CAPACITY
191.90 Mbps
TX POWER (EIRP) 36 dBm

REMOTE

REMOTE IP
203.98.254.3

MAC ADDRESS
24:A4:3C:38:A1:9E

LOCAL RX
IDEAL POWER -51 dBm .
CHAIN IMBALANCE 1 dB .

SIGNAL STRENGTH -63 / -64 dBm

REMOTE RX
IDEAL POWER -51 dBm
CHAIN IMBALANCE 10 dB

SIGNAL STRENGTH -64 / -74 dBm

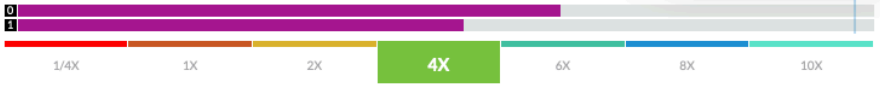

LOCAL RX REAL-TIME CAPACITY / SPEED
MODULATION RATE 4x (16QAM MIMO)

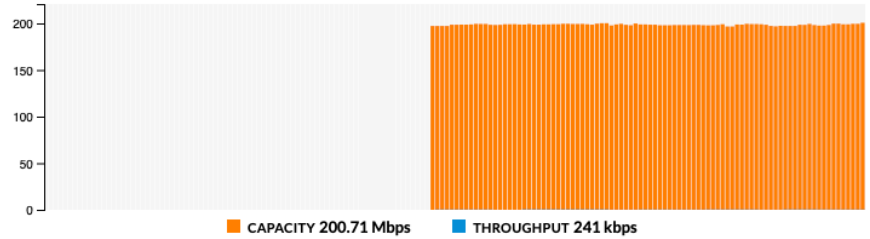

REMOTE RX REAL-TIME CAPACITY / SPEED
MODULATION RATE 4x (16QAM MIMO)

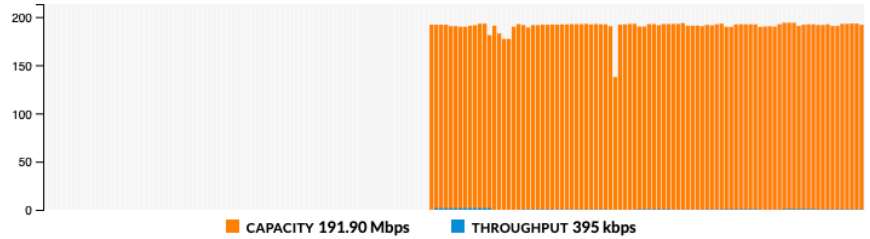

DEVICE

OPERATING MODE	Master	UPTIME	86 days 02:17:34	LATITUDE_LONGITUDE	-0.538695, 166.911942	CONDUCTED TX POWER	14 dBm
RF LINK STATUS	Operational	LINK UPTIME	00:09:11	ALTITUDE	24 m (79 ft)		
VERSION	v4.0.5	REGULATORY DOMAIN	ACMA	SYNCHRONIZATION	Disabled		
DATE	2020-11-13 22:12:46	SECURITY	AES-128	SATELLITES TRACKED	12		
RADIO MODE	MIMO			GPS SIGNAL QUALITY			
DUPLEX	Full Duplex				100 %		

Deploying WiFi to end-sites

- Common issues:
  - Ubiquiti NanoStation M5s used only had username/password changed
  - Channels left to "auto" → lots of frequency conflicts
  - SSIDs not unique per link (APs mis-associating)
  - Management IP address "default" or DHCP
  - Operating system "as shipped"
  - Wide open to potential abuse
- Remediation:
  - As for WiFi backbone: all devices updated, channels fixed, and brought under structured management control
  - "Default" Ubiquiti management IP address is 192.168.1.20 making initial remote access more "interesting"

Deploying WiFi to end-sites

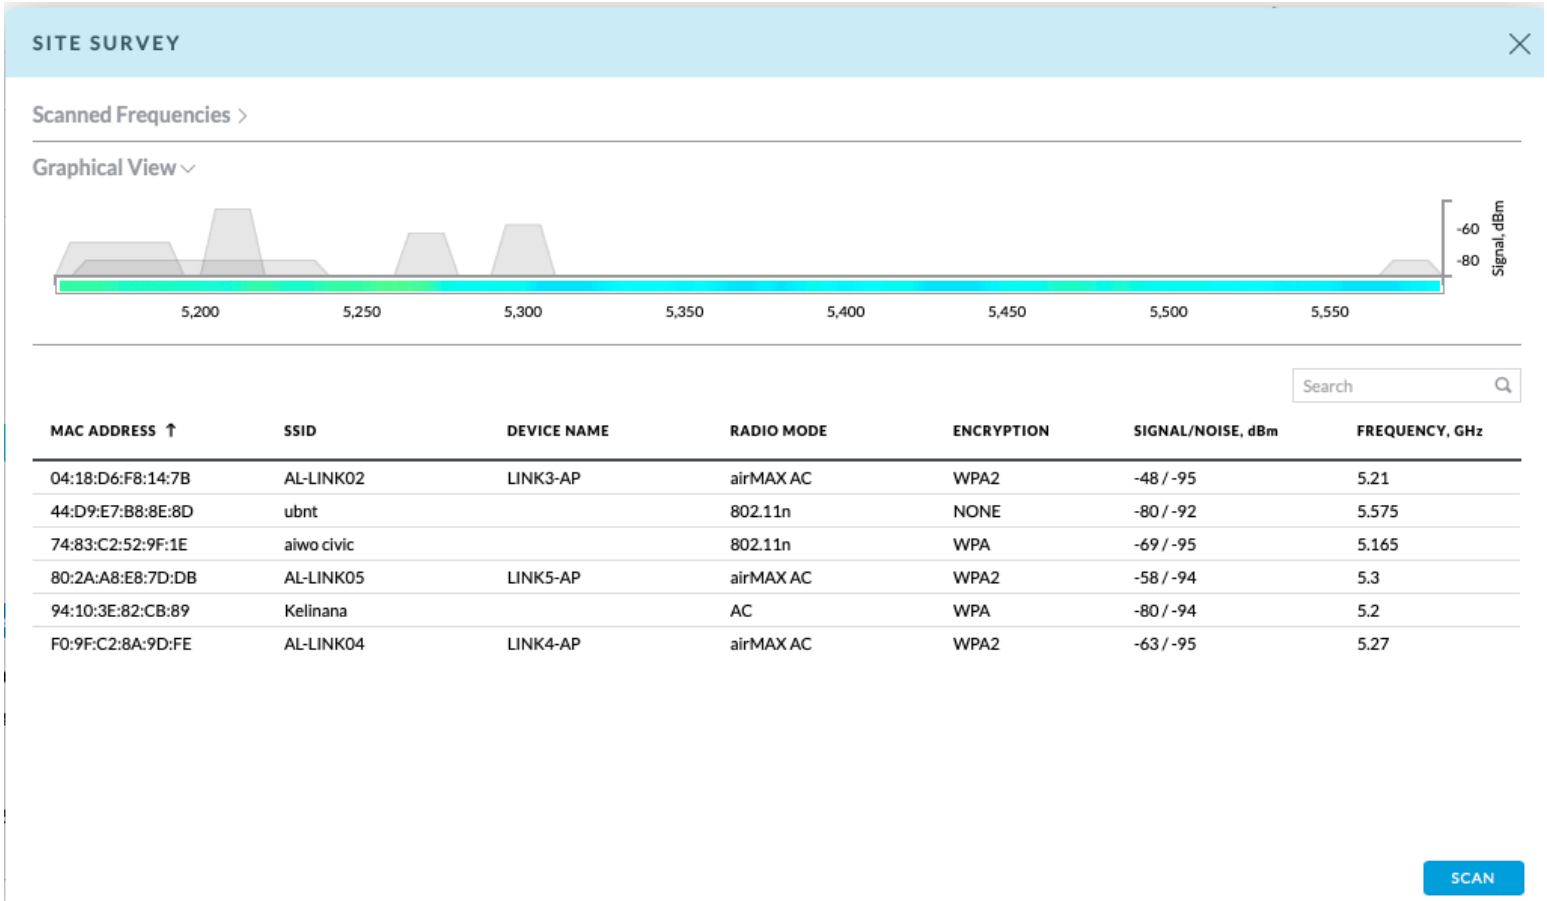

Not all WiFi is Cenpac – anyone can throw up a point to point link!

Challenging when trying to run a professional network infrastructure – always have to keep watch for channel conflicts!

Agenda

- Satellite links
- Routing (BGP, IS-IS etc)
- Island-wide WiFi backbone
- Ganeti Cluster
- Router & Switch Infrastructure
- Network Operations
- Observations

Ganeti Cluster

- Major achievement of 2012 visit was to set up a Ganeti Cluster
  - Retired the 2005 PC hardware that was in service
  - In 2018, the 2012 hardware Phil installed was itself replaced with two Dell R330s

  - The photo shows the challenge – to be sorted, from Brisbane! 🤓

Ganeti Cluster – operating services

- Outcome of 2012 site visit, separate VMs for
  - NRWEB (migrated from physical hardware)
  - DNS (migrated from physical hardware)
  - NOC, which hosted:
    - Observium
    - NfSen
    - TACACS+
    - RANCID
- And this remained until November 2019
  - Only “care and maintenance” for 7 years
    - Which meant NfSen tidy-up every 6 months (**`nfexpire`** not configured)

Ganeti Cluster

- Two tasks:
  - Expand cluster beyond two physical servers
  - Migrate from VMs to Containers
- Reasoning:
  - Better cluster resilience & support for future expansion
  - Containers are vastly more resource efficient
  - Resource upgrade/modernisation of the cluster hardware

Ganeti Cluster

- Dell R330 hardware extended
  - Rack rails procured, servers mounted, replacing the polystyrene blocks
  - Disks upgraded (1TB to 2TB), storage expanded (2TB to 8TB), caddies introduced
  - Second PSU added
  - Memory max'ed out (64GBytes)
  - iDRAC access enabled (from remote) allowing full remote management of the cluster
    - More intuitive than the command line "racadmin"
  - Two more Dell R330s procured and made identical (2nd hand of course)
    - Ever tried procuring 2nd hand parts during a lock down?
    - Currently waiting to be added to cluster

Ganeti Cluster

- Two VMs created by Phil – to host LXD containers
  - Ganeti takes care of mirroring/redundancy
- All the Network Services built fresh, one per container
  - LibreNMS, NfSen, Nagios, Smokeping, RANCID
  - Jumphost, NOC, TACACS+
  - Nauru mail services (using Zimbra)
  - Primary and Secondary DNS for .NR
  - Two Caching DNS
  - Routinator 3000 Validator
  - Speedtest
- Only other VM remaining is the "NRWEB" ccTLD manager
  - 2012 NOC and 2006 DNS turned off

Ganeti Cluster

- Future plans:
  - Two new VMs on the two new cluster members
  - Redistribution of containers across the four LXDs to improve load balancing
    - NfSen is a disk space hog (if we want decent history)
    - Zimbra is a resource hog
  - Replacement of ccTLD manager with modern ccTLD registry system

Speedtest?

- Cenpac customers were running Speedtest on their links and complaining they weren't getting the bandwidth they were paying for
  - Nearest Speedtest node from Nauru is in Sydney
  - Speedtest does very badly on high latency links
  - The only useful test tool is something like iperf3
  - But the "typical user" thinks that Speedtest is the gold standard of performance testing
- Speedtest node on Nauru became essential!
  - Now no argument about bandwidth – domestic bandwidth is limited by physical access speed for all Cenpac customers!

Typical Speedtest result on Nauru

Typical Speedtest result prior to local node going live

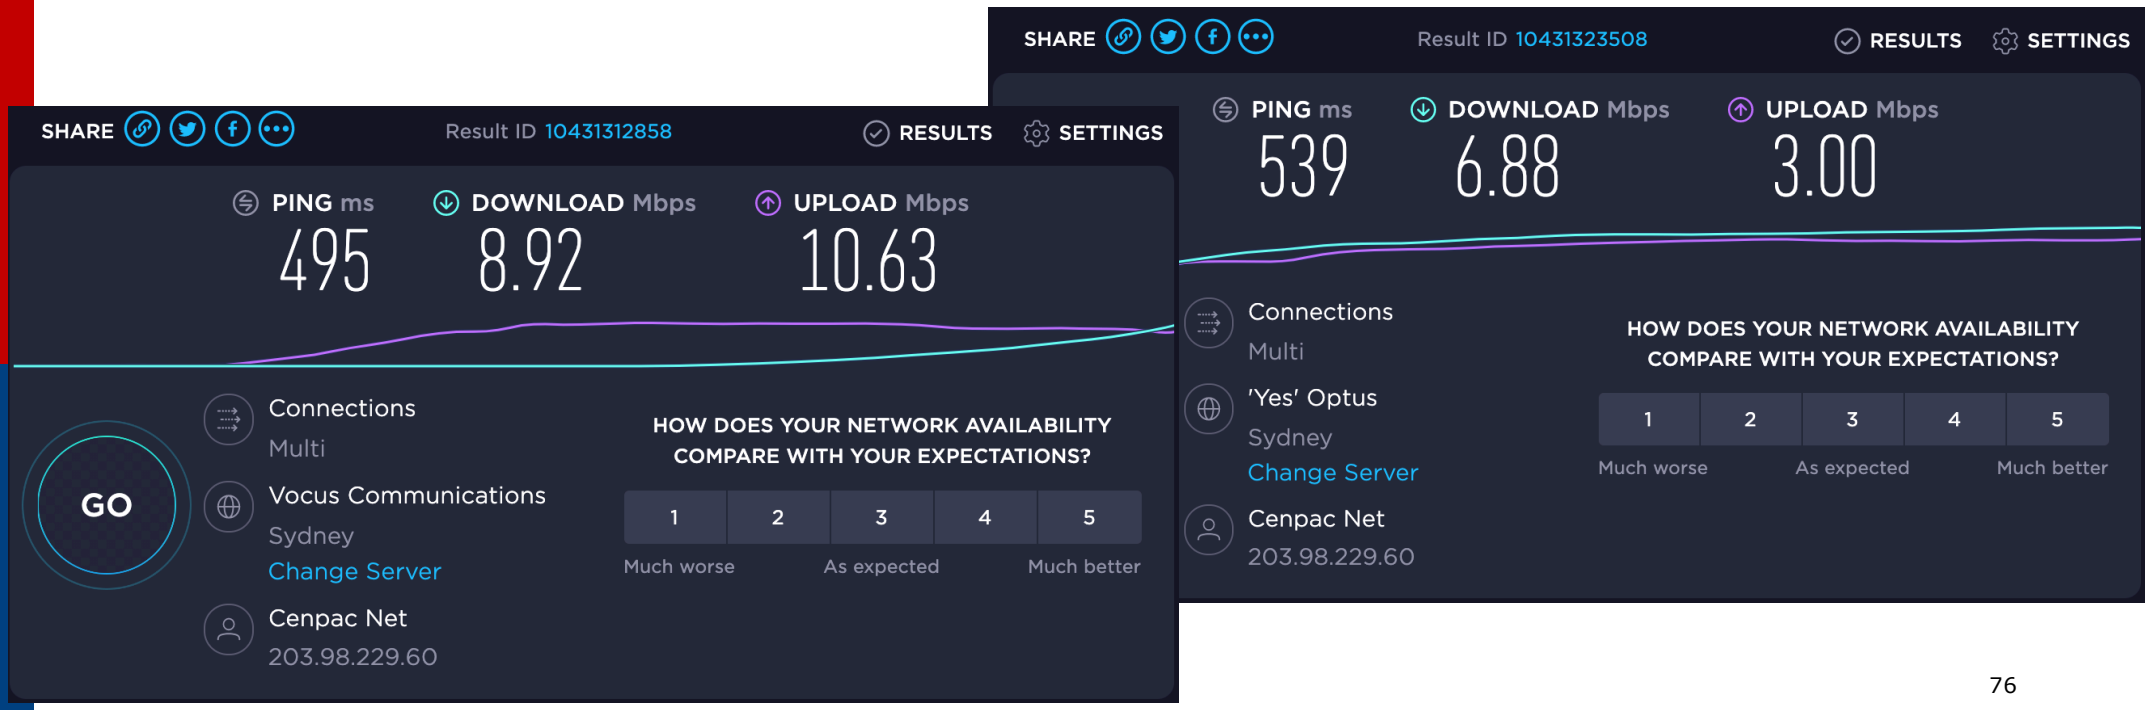

Typical Speedtest result on Nauru

Typical Speedtest result after the local node went live

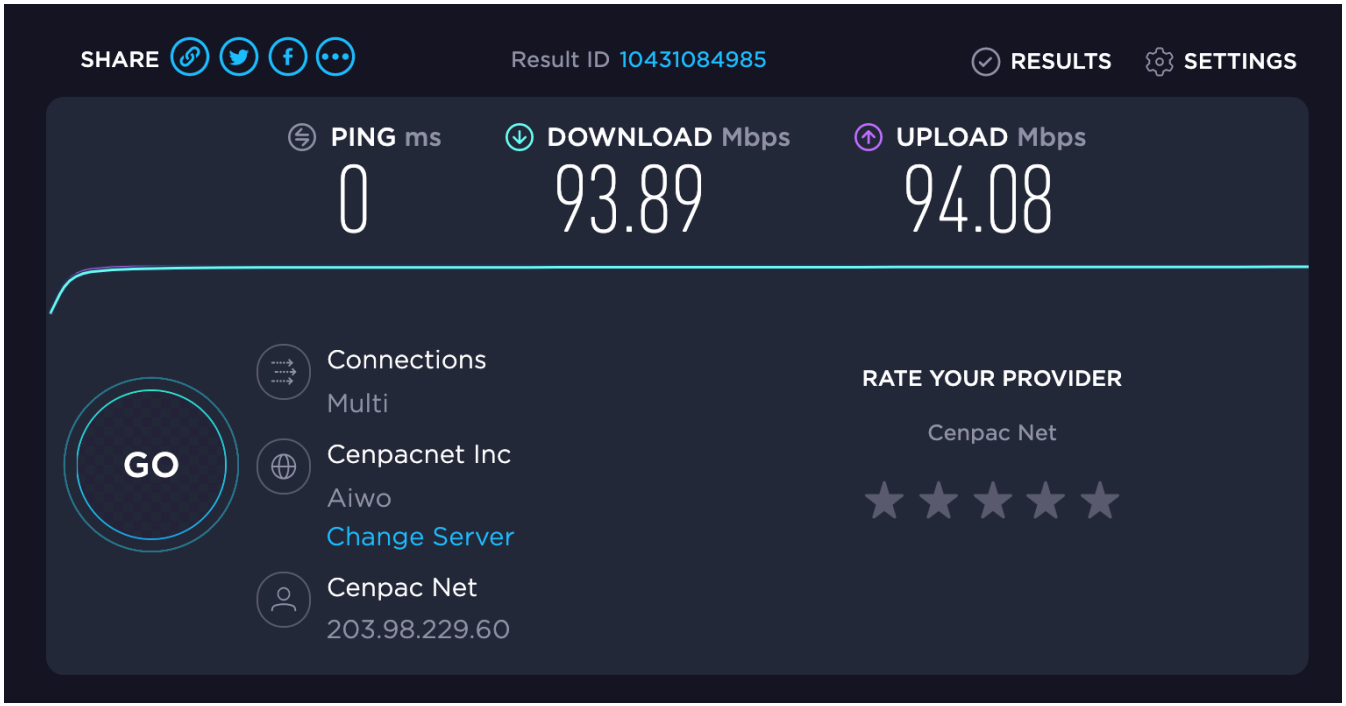

Agenda

- Satellite links
- Routing (BGP, IS-IS etc)
- Island-wide WiFi backbone
- Ganeti Cluster
- Router & Switch Infrastructure
- Network Operations
- Observations

Routers and Switches

- Nothing "exciting" here
  - Upgrading those remotely has been something we've all been doing for 30+ years!
  - Top tip: don't saw off the branch you are sitting on – it hurts when you hit the ground
- Upgrades:
  - Cisco 3845 and 2911 routers replaced with two ASR1001s (border routers) and two ASR1002s (core routers)
  - Catalyst 2950 aggregation switches replaced with Catalyst 2960S-48 switches
    - 4 fibre ports for connections to ASR1002 core routers

Equipment Room

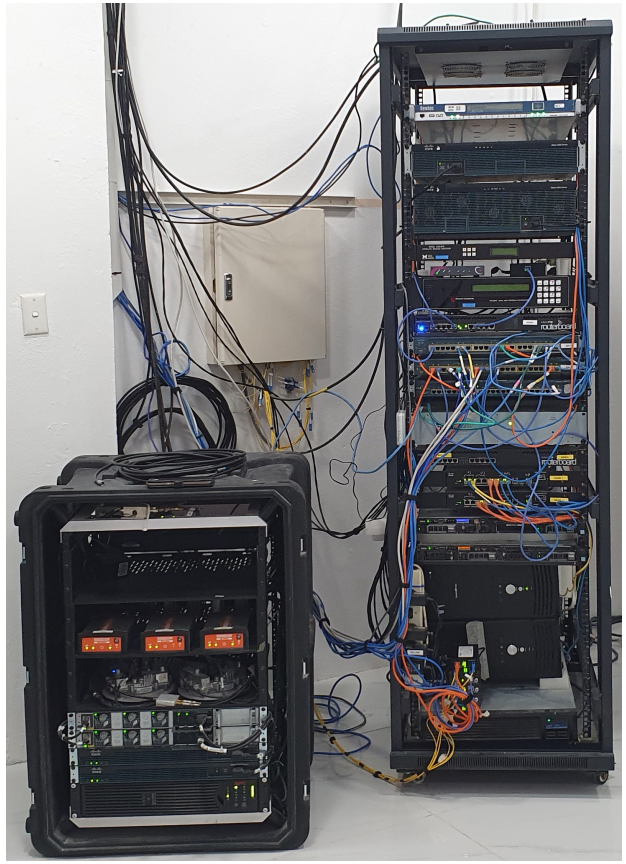

May 2020

August 2020

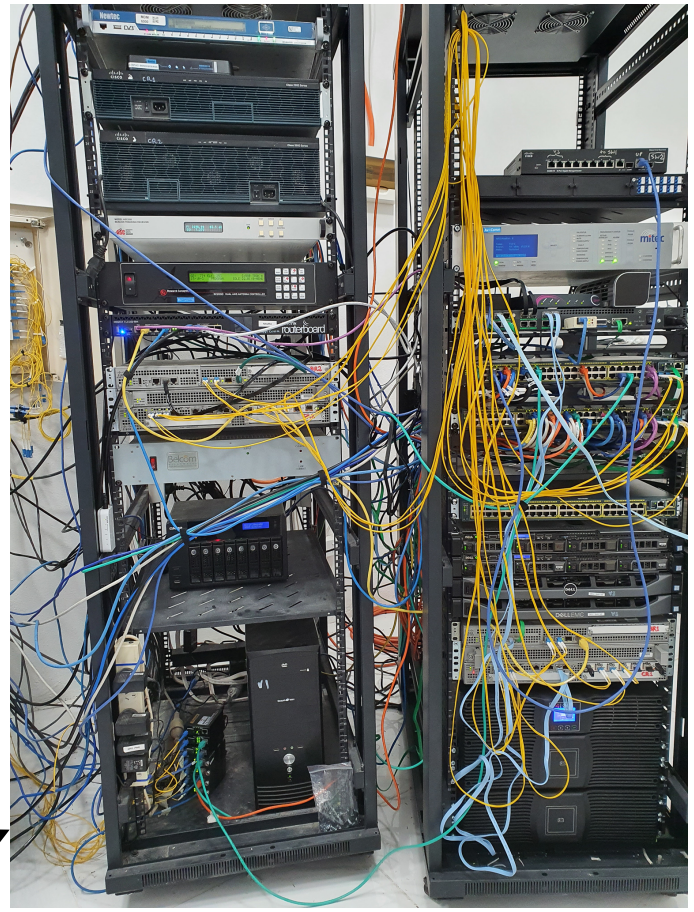

Upgrade process

- The plan:
  - Build the whole new core mesh
    - Four ASRs, two Cat2960S switches
    - Interconnected by fibre
  - And then transfer each external connection and customer connection from the old core to the new core
- Easier said than done
  - How to maintain connectivity from Brisbane to help guide local team??
  - Think about it… 🤪🤓

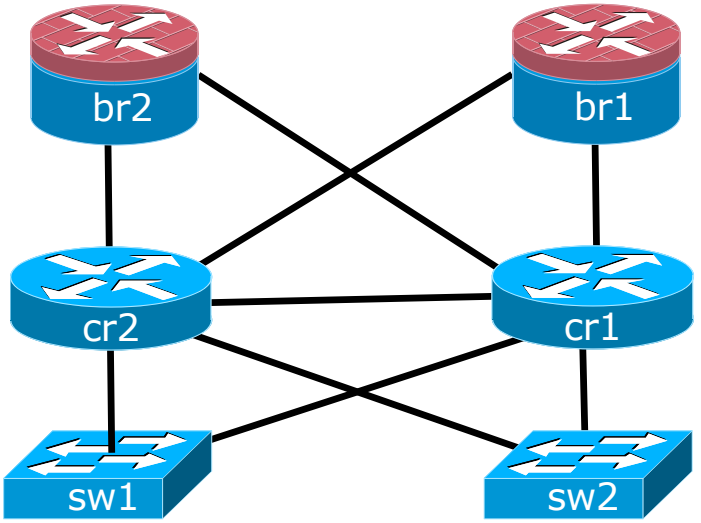

Routers racked up

BR1 and CR1
(right hand rack)

BR2 and CR2
(left hand rack)

Routers and Switches

- Building the new core:
  - New routers and switches installed in racks, as normal
  - Cabling diagram and port allocation provided
    - Single mode SFPs, 3m fibre patch leads
  - Installation guided over WhatsApp video call
  - Cisco 2511RJ terminal server sitting on old/live core handled console access to this new core
- How to swap live network on to new core??
  - Move one satellite link means connectivity could be to old core or new core, depending on upstream provider, best path, etc

Core transition (1)

- Disconnect one satellite link completely
  - This guarantees no asymmetric routing or blackholing of traffic
  - Make sure border router access on new core was feasible using public IP address of upstream provider (see earlier!)
- Then move the other satellite link to the new core
  - IPv4 address block originated from new core now
  - Access to border router successful
  - But WhatsApp video call dropped (obviously)
- Move Cenpac office WiFi to new switch port
  - Back on to WhatsApp video again 😀

Core transition (2)

- Next, moving Ganeti cluster on to the new switch infrastructure
  - Cluster used 2x Cisco Small Business SG300-10 for replication and general connectivity
  - Our idea was to move from SG300s directly onto the Cat2960S
  - Failed – was too ambitious
    - SG300 way of doing VLAN trunking is v different from IOS, more different than expected 🤪😉🤯
    - Not helped by one of the servers in the cluster having trunking configured in a different way from the other 🙃😐
  - Backed out, and simply moved the SG300 uplink ports to the Cat2960S pending future resolution – success!

Core transition (3)

- Finally:
  - Move 2511RJ terminal server for console access
  - Added in the other satellite link
  - Moved all customer connections from old switch
  - All guided over WhatsApp Video call
  - And after that, updating LibreNMS/Smokeping/Nagios
- Summary:
  - Instructions to remote hands need to be crystal clear without any assumptions made
  - Diagrams help! Good diagrams help even more!
  - Customer outage was only for about an hour too – all work carried out in late evening

Agenda

- Satellite links
- Routing (BGP, IS-IS etc)
- Island-wide WiFi backbone
- Ganeti Cluster
- Router & Switch Infrastructure
- Network Operations
- Observations

Network Operations

- Local team monitors:
  - LibreNMS for traffic usage by customers and on satellite links
    - Lots of customers “testing” their links to make sure they are getting what they are paying for – and the network delivers every time 👏
  - Nagios for connectivity issues
    - Especially WiFi network
  - Smokeping for packet loss, jitter, and round trip time
    - Especially WiFi network and satellite links
  - Speed of Zimbra
    - How quickly Zimbra works (given it is a resource hog) shows the cluster’s general health 😇

Watching everything!

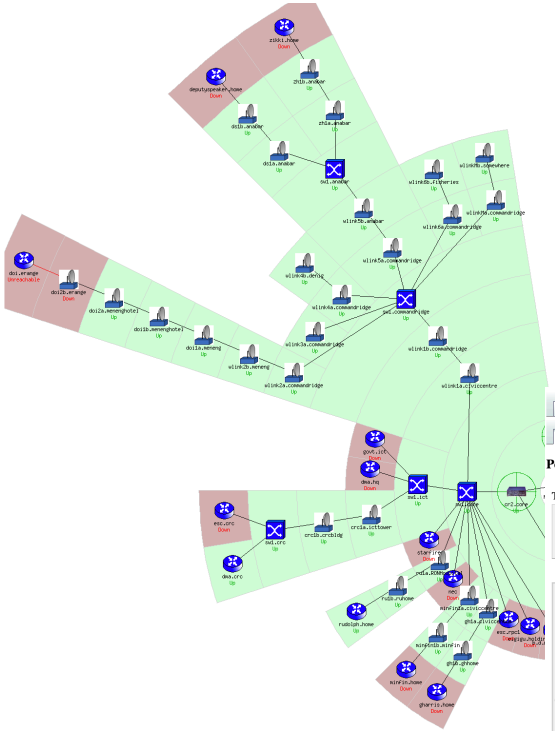

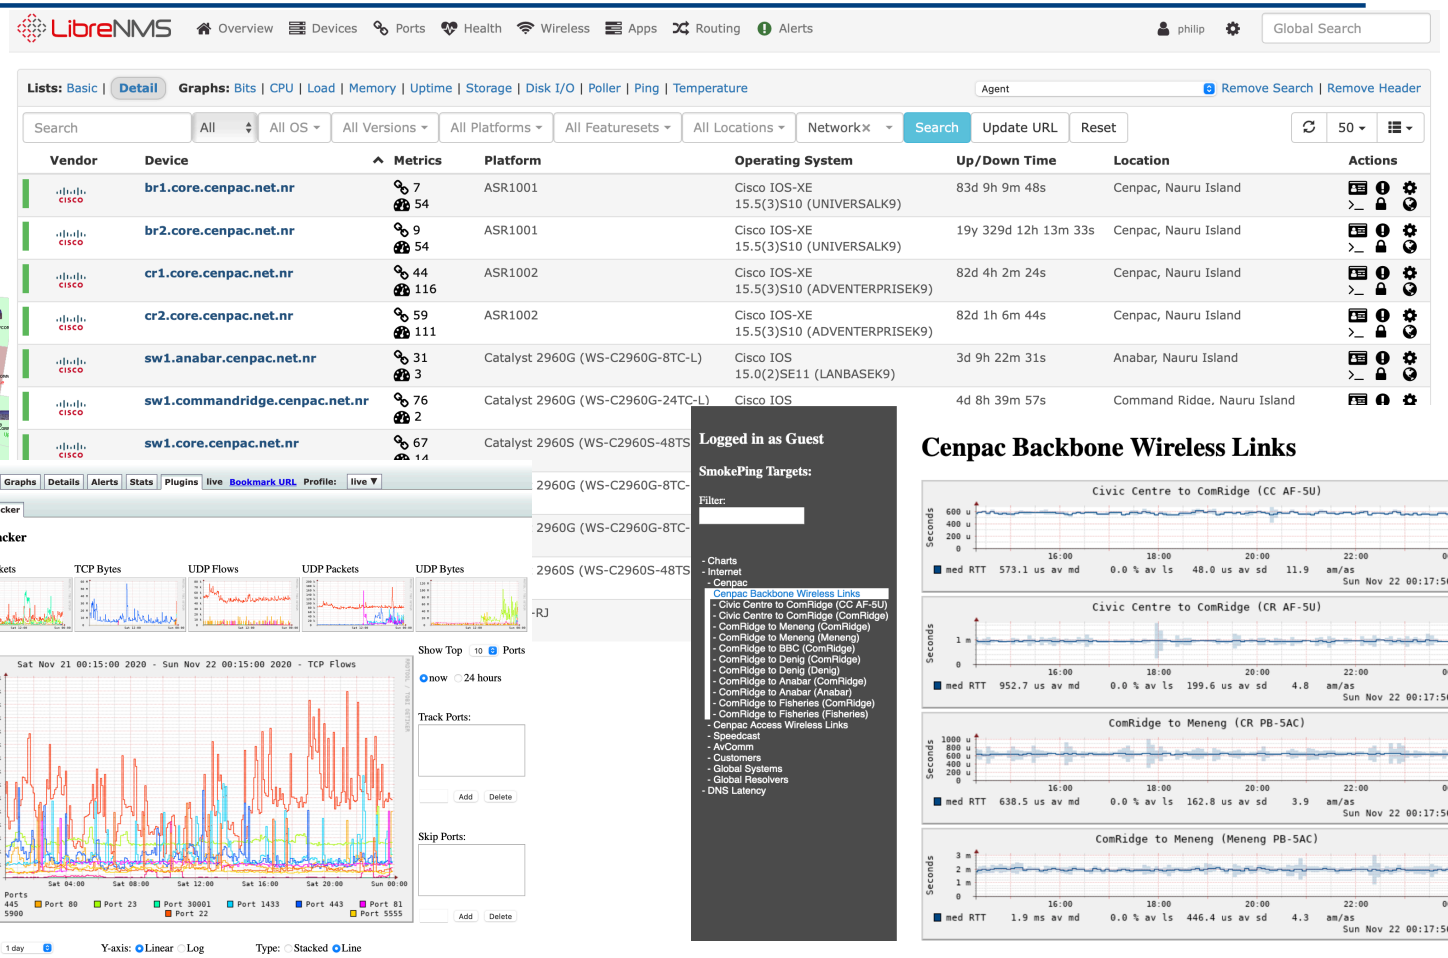

Installations

- New customer installations:
  - Fibre (when it works)
    - WiFi Router goes in to customer site (usually)
    - Media converters (BiDi) – big sigh!
  - WiFi
    - WiFi backbone – and NanoStations from the Island Towers to end customers
    - Direct links from Civic Centre (if line of sight) to customers
    - From fibre termination points around Island to customers
    - NanoStations either preconfigured in office (remotely from Brisbane) or re-configured once installed in the field (see earlier)
  - Diagrams help! Lots!

Agenda

- Satellite links
- Routing (BGP, IS-IS etc)
- Island-wide WiFi backbone
- Ganeti Cluster
- Router & Switch Infrastructure
- Network Operations
- Observations

Observations

- Procuring Equipment
  - If not available locally (in country), it may as well not exist
    - And even during April-June lockdown, inter-state was taking 2-3 weeks!! 😐
  - International shipments to Australia took 3+ months in mid-year
    - Now quotes are down to 6-8 weeks 🤯😤🙄
- Shipping Equipment
  - Everything has slowed down during the pandemic
  - Freight forwarders have:
    - Reduced staffing
    - Introduced social distancing
    - Temperature checks and Hand Sanitizer everywhere
  - Careful with your Air Way Bill details!
    - Receipts, proof of shipping, confirm flights, etc

Observations

- Never Assume Anything!
- Communications
  - Avoid lengthy emails to explain a task – diagram or photo!
  - One subject per email, especially if there are lots of tasks needed to achieve a particular outcome
- Documentation
  - Diagrams and Photos are worth thousands of words
  - Diagrams and Photos avoid frustration and mis-understanding
  - Countless pictures taken, countless diagrams swapped
  - Update monitoring systems immediately tasks are completed

Observations

- Language:
  - Explain work required in clear and plain language
  - Don’t use jargon – don’t assume your jargon is known by others
- Remote Hands
  - When asking clients and colleagues to do work, don’t assume they have your own or your regular remote hands’ skillset
- Video tools
  - Use your favourite video tool and get on a video call – seeing with your own eyes saves emails, voice calls, messaging,…
- Planning
  - More detailed planning needed – how to stay connected when doing major transitions

Thank you!