

# Troubleshooting

## Known problems

### Default G/W for Igor

Igor can't find it's local network:

```
hxd@morgana ~/net> ssh -J igor.lan 192.168.50.159
channel 0: open failed: connect failed: open failed
stdio forwarding failed
kex_exchange_identification: Connection closed by remote host
Connection closed by UNKNOWN port 65535
```

Use web UI to access console and exec as root:

```
ip route add default via 192.168.100.1
```

# Troubleshooting

### Default G/W for Igor

Igor can't find it's local network:

```
hxd@morgana ~/net> ssh -J igor.lan 192.168.50.159
channel 0: open failed: connect failed: open failed
stdio forwarding failed
kex_exchange_identification: Connection closed by remote host
Connection closed by UNKNOWN port 65535
```

Use web UI to access console and exec as root:

```
ip route add default via 192.168.100.1
```

# TCP dump

Installing on OpenWRT:

```
opkg update
opkg install tcpdump
```

## DHCP

```
tcpdump -vvv -i any udp port 67 and port 68
```

# Network Issues

## Wi-Fi slow

Slow access to HOME from laptop on MGMT wi-fi.

### UPDATE: 2023-10-28

Looks like laptop gets very low Rx rate (throughput from haru to Morgana) of ~17Mbit or even 6Mbit:

```
iperf3 -c 192.168.100.20 -p 2345 -R -t 9999
Connecting to host 192.168.100.20, port 2345
Reverse mode, remote host 192.168.100.20 is sending
[ 5] local 192.168.100.161 port 42966 connected to 192.168.100.20 port 2345
[ ID] Interval          Transfer      Bitrate
[ 5]  0.00-1.00    sec  1.75 MBytes  14.7 Mbits/sec
[ 5]  1.00-2.00    sec  1.64 MBytes  13.8 Mbits/sec
[ 5]  2.00-3.00    sec  1.64 MBytes  13.8 Mbits/sec
[ 5]  3.00-4.00    sec  1.55 MBytes  13.0 Mbits/sec
[ 5]  4.00-5.00    sec  1.62 MBytes  13.6 Mbits/sec
[ 5]  5.00-6.00    sec  1.63 MBytes  13.6 Mbits/sec
^C[ 5]  6.00-6.14    sec   252 KBytes  14.4 Mbits/sec
- - - - -
[ ID] Interval          Transfer      Bitrate
[ 5]  0.00-6.14    sec    0.00 Bytes  0.00 bits/sec          sender
[ 5]  0.00-6.14    sec  10.1 MBytes  13.8 Mbits/sec          receiver
```

At the same time I get like 60Mbit sending data out:

```

iperf3 -c 192.168.100.20 -p 2345 -t 9999
Connecting to host 192.168.100.20, port 2345
[ 5] local 192.168.100.161 port 45200 connected to 192.168.100.20 port 2345
[ ID] Interval           Transfer     Bitrate      Retr  Cwnd
[ 5]  0.00-1.00   sec  10.1 MBytes  85.0 Mb/s    0    469 KBytes
[ 5]  1.00-2.00   sec   6.71 MBytes  56.3 Mb/s    0    469 KBytes
[ 5]  2.00-3.00   sec   7.08 MBytes  59.4 Mb/s    0    502 KBytes
[ 5]  3.00-4.00   sec   6.40 MBytes  53.7 Mb/s    0    529 KBytes
[ 5]  4.00-5.00   sec   7.77 MBytes  65.2 Mb/s    0    529 KBytes
[ 5]  5.00-6.00   sec   7.77 MBytes  65.2 Mb/s    0    529 KBytes
[ 5]  6.00-7.00   sec   7.01 MBytes  58.8 Mb/s    0    587 KBytes
[ 5]  7.00-8.00   sec   8.47 MBytes  71.0 Mb/s    0    621 KBytes
^C[ 5]  8.00-8.40   sec   2.41 MBytes  50.4 Mb/s    0    621 KBytes
-----
[ ID] Interval           Transfer     Bitrate      Retr
[ 5]  0.00-8.40   sec  63.8 MBytes  63.7 Mb/s    0          sender
[ 5]  0.00-8.40   sec   0.00 Bytes  0.00 bits/s  0          receiver

```

Haru reports low MCS values of 10 or lower.

Server	SSID	Channel	Client	Location	Client -> Server Mbit	Server <- Client Mbit	
Justine	Haru Legacy	6	L (android)	living room	25	25	
Justine	Haru Legacy	6	futaba	living room	21	9	
Justine	Haru Legacy	6	futaba	kitchen	25	10	
Justine	Haru Legacy	6	futaba	by Haru	26	11	
Justine	Haru Legacy	11	futaba	living room	26	11	
Justine	Haru MGMT	11	morgana	living room	53	23	
Justine	Haru Legacy	6	morgana	living room	50	11	
Justine	Haru	?	morgana	living room	89	72	
Justine	Haru (IE)	52	morgana	living room	95	95	
Justine	Haru MGMT (IE)	11	morgana	living room	92	92	

Justine	Haru MGMT (IE)	11	morgana	kitchen	95	94	
Justine	Haru (IE/1G)	11	morgana	kitchen	53	321	*1
Justine	Haru MTMT (IE/1G)	52	morgana	kitchen	44	92	*1

\*1 - both washing and drying going on

Config changes:

- Set max power to 30dBm
  - channel switched to 5 (was 6): 15-25 on morgana and 13 on futaba
  - channel 11: 10 to 25 on morgana, 16 on futaba
- Setting region to IE(!):
  - 95/95 morgana -> justine over Haru channel 52
  - BINGO!: 92/92 on MGMT channel 11
  - This get me to 100Mbit; but I get much higher rates reported for wifi like 866/780 Mbit on Haru
- eth0 on Speed: 100Mb/s!
  - Looks like faulty cable but I also rebooted Haru so could be that as well, although on secondary port I was getting 1Gbit with laptop before restart.
  - goro reports only lan8 (haru) at 100Mbit/s but others at 1Gbit

```

Mon Oct 30 19:30:27 2023 kern.info kernel: [8584550.111467] RTL8380 Link change:
status: 1, ports 8000
Mon Oct 30 19:30:28 2023 kern.info kernel: [8584550.992936] rtl83xx-switch
switch@1b000000 lan8: Link is Up - 100Mbps/Full - flow control rx/tx
Mon Oct 30 19:30:28 2023 kern.info kernel: [8584551.003123] switch: port 8(lan8)
entered blocking state
Mon Oct 30 19:30:28 2023 kern.info kernel: [8584551.009306] switch: port 8(lan8)
entered forwarding state
Mon Oct 30 19:30:28 2023 daemon.notice netifd: Network device 'lan8' link is up

```

- I have replaced the cable and now have 1Gbps

```

Mon Oct 30 20:32:29 2023 kern.info kernel: [8588272.090512] rtl83xx-switch
switch@1b000000 lan8: Link is Up - 1Gbps/Full - flow control rx/tx

```

## UPDATE: 2023-11-05

There was a packet drop between Haru and Goro. I have replaced the cable that goes from Goro to power supply (for Haru).

After replacing the cable link dropped to 100Mbps, reconnecting it got me 1Gbit.

I have also switched channels for 5GHz radio to use:

- channel: 104 (5520MHz)
- width: 160MHz

Now everything is very fast. In kitchen I got 356/407 Mbps using Haru!

## UPDATE: 2023-11-11

After power issue the network went up in bad state.

I was getting 300Mbit one way and only 46Mbit and high packet loss the other way.

I decided to route another cable to Haru.

Looks like Primary interface is the only one that can take power so I left it connected to the power supply but disconnected it from the switch.

New, longer cable now connects switch to Haru Secondary port which is part of a bridge setup with Primary port so no configuration changes was needed.

Now I am getting up from Morgana in kitchen 657Mbit and down 334Mbit from Justine.

## DHCP no responses, no IP assigned

Looks like Graphene OS uses random MAC for every connection attempt to Haru:

```
13:34:45.185659 IP 0.0.0.0.68 > 255.255.255.255.67: B00TP/DHCP, Request from
b2:7f:5d:03:bf:61, length 288
13:34:46.144502 IP 0.0.0.0.68 > 255.255.255.255.67: B00TP/DHCP, Request from
b2:7f:5d:03:bf:61, length 288
13:34:48.135488 IP 0.0.0.0.68 > 255.255.255.255.67: B00TP/DHCP, Request from
b2:7f:5d:03:bf:61, length 288
13:34:52.237864 IP 0.0.0.0.68 > 255.255.255.255.67: B00TP/DHCP, Request from
b2:7f:5d:03:bf:61, length 288
13:35:00.303513 IP 0.0.0.0.68 > 255.255.255.255.67: B00TP/DHCP, Request from
b2:7f:5d:03:bf:61, length 288
13:35:10.015298 IP 0.0.0.0.68 > 255.255.255.255.67: B00TP/DHCP, Request from
8a:ba:ab:74:33:23, length 288
13:35:11.105965 IP 0.0.0.0.68 > 255.255.255.255.67: B00TP/DHCP, Request from
8a:ba:ab:74:33:23, length 288
13:35:13.292155 IP 0.0.0.0.68 > 255.255.255.255.67: B00TP/DHCP, Request from
8a:ba:ab:74:33:23, length 288
```

```
13:35:17.375826 IP 0.0.0.0.68 > 255.255.255.255.67: B00TP/DHCP, Request from
8a:ba:ab:74:33:23, length 288
13:35:24.972812 IP 0.0.0.0.68 > 255.255.255.255.67: B00TP/DHCP, Request from
8a:ba:ab:74:33:23, length 288
13:43:05.585122 IP 0.0.0.0.68 > 255.255.255.255.67: B00TP/DHCP, Request from
5a:88:7a:60:55:fd, length 288
13:43:06.607297 IP 0.0.0.0.68 > 255.255.255.255.67: B00TP/DHCP, Request from
5a:88:7a:60:55:fd, length 288
13:43:08.473318 IP 0.0.0.0.68 > 255.255.255.255.67: B00TP/DHCP, Request from
5a:88:7a:60:55:fd, length 288
13:43:12.547523 IP 0.0.0.0.68 > 255.255.255.255.67: B00TP/DHCP, Request from
5a:88:7a:60:55:fd, length 288
13:43:21.310413 IP 0.0.0.0.68 > 255.255.255.255.67: B00TP/DHCP, Request from
5a:88:7a:60:55:fd, length 288
13:43:27.942393 IP 0.0.0.0.68 > 255.255.255.255.67: B00TP/DHCP, Request from
ca:9a:d3:13:a3:1c, length 288
13:43:28.860544 IP 0.0.0.0.68 > 255.255.255.255.67: B00TP/DHCP, Request from
ca:9a:d3:13:a3:1c, length 288
13:43:30.774370 IP 0.0.0.0.68 > 255.255.255.255.67: B00TP/DHCP, Request from
ca:9a:d3:13:a3:1c, length 288
13:43:34.697002 IP 0.0.0.0.68 > 255.255.255.255.67: B00TP/DHCP, Request from
ca:9a:d3:13:a3:1c, length 288
13:43:43.292732 IP 0.0.0.0.68 > 255.255.255.255.67: B00TP/DHCP, Request from
ca:9a:d3:13:a3:1c, length 288
```

This may lead to depletion of IP addresses (pool has up to 150 to allocate).

To clean up the pool SSH to Caroline:

```
service dnsmasq stop
mv /tmp/dhcp.leases /tmp/dhcp.leases.bac
service dnsmasq start
```

To mitigate the issue I have reduced leas time from 30 days to 24 hours.

## VPN clients coming from outside are NAT'ed

They will look like **justine**, not their actual VPN IP, since devices can use **caroline** as their default G/W.

# No access to SERVER VLAN from HOME with justine G/W

## VPN clients can access justine MGMT interface IP

```
11:06:10.811023 vpn In IP 172.17.1.10 > 192.168.100.2: ICMP echo request, id 44951, seq 814, length 64
11:06:10.811041 vpn Out IP 192.168.100.2 > 172.17.1.10: ICMP echo reply, id 44951, seq 814, length 64
```

I have set up more strict forwarding rules:

```
Chain FORWARD (policy DROP 52 packets, 4368 bytes)
num pkts bytes target prot opt in out source destination
9 35 5441 ACCEPT all -- enpls0 vpn 192.168.0.0/24 0.0.0.0/0
10 32 5244 ACCEPT all -- vpn enpls0 0.0.0.0/0 192.168.0.0/24
```

but this does not help.

This is because it is not going through FORWARD but through INPUT:

```
iptables -A INPUT -i vpn -d 192.168.100.0/24 -j LOG
```

```
[604557.640819] IN=vpn OUT= MAC= SRC=172.17.1.10 DST=192.168.100.2 LEN=84 TOS=0x00 PREC=0x00
TTL=64 ID=56812 DF PROTO=ICMP TYPE=8 CODE=0 ID=44951 SEQ=1064
```

Adding the above rule will block it. But is this normal that any IP assigned to any interface can be used when routing to a G/W?

**FIXED:** I have added the following rule:

```
iptables -A INPUT ! -i mgmt -d 192.168.100.0/24 -j DROP
```

**FOLLOWUP:** Is this the same for other devices? Try static route to them and access MGMT IP.

```
ip route add 192.168.100.50 via 192.168.0.50 dev wlan0
```

And it will ping on 192.168.100.50. So services bound only to 192.168.100.50 will be exposed to HOME VLAN devices.

**CONCLUSION:** Binding to selected IP address does not protect service from being accessed from another network interface without extra firewall rule to prevent this!

What makes localhost (`lo`) interface special - services bound to it are not accessible from other interfaces? It has `local` routing tables set up by default.

## Same problem for caroline

Same with **caroline**:

```
ip route add 192.168.100.1 via 192.168.0.1 dev wlan0
```

I can now access <https://192.168.100.1/cgi-bin/luci/> while on HOME VLAN, probably will work form SERVER and any other as well!

**FIXED:** DROP all INPUT on all interfaces apart from MGMT. Added ACCEPT rules for LAN, GUEST and SERVER VLANs for DHCP (UDP 67), DNS (TCP/UDP: 53) and ICMP.

I could not use the negative interface setup as in case of justine. Probably should use default INPUT DROP on justine as well and only allow mgmt interface traffic as well.

## Would this be same for goro and haru?

They don't have IP on HOME network but they have interface. Injecting packet for MGMT IP to their HOME interface may be (using MAC/ARP) possible but they would not respond since they have no routing to HOME network?

## Access to web services from internal network

Need to manually add static route for devices using default DHCP G/W (justine) when going to local server services like video.hexadust.net.

```
ip route add 46.7.126.16 dev eth0 via 192.168.0.1
ping video.hexadust.net
```

## Things to try:

1. Try to push static routes from DHCP - this did not work for some reason

2. Set up static route on justine
3. Set up SNAT on justnie
4. Use bridge layer DNAT: [https://ebtables.netfilter.org/br\\_fw\\_ia/br\\_fw\\_ia.html](https://ebtables.netfilter.org/br_fw_ia/br_fw_ia.html)
5. Set up split DNS
  1. `video.hexadust.net 192.168.50.159`
  2. This would require justine to forward to SERVER VLAN since ann uses 0.2 as default G/W
6. Set up separate DNS for internal access

## No more IPs on DHCP

Looks like Horizon box eats up leases: <https://www.boards.ie/discussion/2057720465/my-new-virgin-media-stb-issued-two-lan-i-ps>

I have removed leases from `/var/lib/dhcp.lease` file on caroline.

## Justine access is laggy; DNS is slow

Happens after OpenWRT updates.

Try disconnecting and connecting network cables for Haru and uplink on goro.

**UPDATE: 2023-11-05:** I have replaced both cables that connect Haru and Goro. The link tends to drop to 100Mbit if Goro end is disconnected, reconnecting fixes it.

## Igor does not set it's default route after boot

While it is configured in web UI it is not taking effect on boot and results in it unable to find local network. VMs stuff will work OK though.

## TODO

## YunoHost access to internet via default IP instead of VPN

Not good for IRC etc.

1. Clone the VM and put copy in HOME VLAN, remove public web stuff from it, remove web clients from SERVER VLAN one.
2. Better yet. Create DMZ VLAN that has only access to Caroline and will run `www` server. Use SERVER VLAN to only have access to Justine, so out traffic goes out of the VPN, add interface to YunoHost that is in DMZ VLAN so that it can get requests from www server, but have it default G/W to use Justine over DMZ?
3. Create SANDBOX VLAN that only has access to justine as G/W.
4. Create VM that connects to Justine via WireGuard and pit it in HOME network and in virtual network that it will be G/W for, put other VMs in that sandbox virtual network.
5. Create VM that uses SERVER VLAN to make connection to Mullvad (dedicated key) and act as a G/W for sandbox virtual network for other VMs.

## 24/09: WWW access very slow

Downloading file from www server (Caddy) is slow.

Phone over 5G with VPN (to justine): ~1-2 MB/s

Phone over Haru with VPN: ~2-3 MB/s

From laptop (VPN + Mullvad): ~3 MB/s [NOTE: Even using local SERVER IP it goes over Mullvad IP!]

```
> curl --insecure -o /dev/null --connect-to 192.168.50.159:443
https://jpastuszek.net/links/data --http1.1
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload  Total  Spent    Left  Speed
  1   412M    1 5216k    0     0   50859      0  2:21:39  0:01:45  2:19:54 63055
```

From laptop but using port 8080 (darkhttpd) and no TLS: ~38 MB/s [NOTE: NAT'ed by 192.168.0.2 justine]

```
> curl --insecure -o /dev/null http://192.168.50.159:8080/links/data
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload  Total  Spent    Left  Speed
 100   412M  100  412M    0     0   37.9M      0  0:00:10  0:00:10  --:--:-- 38.0M
```

From SDF: ~1 MB/s

```
$ curl -o /dev/null https://jpastuszek.net/links/data --http1.1
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
```

```

          Dload Upload  Total  Spent  Left  Speed
  5  412M    5 22.8M    0    0  884k      0 0:07:57 0:00:26 0:07:31 957k^C

$ curl -o /dev/null https://jpastuszek.net/links/data
% Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
          Dload  Upload  Total  Spent  Left  Speed
  2  412M    2 11.9M    0    0  796k      0 0:08:49 0:00:15 0:08:34 1007k

```

From Igor VM in the same VLAN (SERVER): ~85 MB/s

```

hxd@void ~> curl --insecure -o /dev/null --connect-to 192.168.50.159:443
https://jpastuszek.net/links/data --http1.1
% Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
          Dload  Upload  Total  Spent  Left  Speed
100  412M  100  412M    0    0  84.8M      0 0:00:04 0:00:04 --:--:-- 87.0M
hxd@void ~> curl --insecure -o /dev/null --connect-to 192.168.50.159:443
https://jpastuszek.net/links/data
% Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
          Dload  Upload  Total  Spent  Left  Speed
100  412M  100  412M    0    0  83.7M      0 0:00:04 0:00:04 --:--:-- 85.6M

```

Same but to dakrhttpd (no TLS): ~700 MB/s:

```

> curl --insecure -o /dev/null http://192.168.50.159:8080/links/data
% Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
          Dload  Upload  Total  Spent  Left  Speed
100  412M  100  412M    0    0  691M      0 --:--:-- --:--:-- --:--:-- 691M

```

iperf3 between laptop and www: 300 Mbit/s, 37 MB/s [NOTE: it is NAT'ed by 192.168.0.2 justine]

```

hxd@morgana /tmp [1]> iperf3 -c 192.168.50.159 -t 9999 -R
Connecting to host 192.168.50.159, port 5201
Reverse mode, remote host 192.168.50.159 is sending
[ 5] local 172.17.1.10 port 42964 connected to 192.168.50.159 port 5201
[ ID] Interval          Transfer      Bitrate
[ 5]  0.00-1.00  sec  34.2 MBytes  287 Mbits/sec
[ 5]  1.00-2.00  sec  33.1 MBytes  278 Mbits/sec
[ 5]  2.00-3.00  sec  35.1 MBytes  295 Mbits/sec
[ 5]  3.00-4.00  sec  35.0 MBytes  294 Mbits/sec
^C[ 5]  4.00-4.17  sec  5.75 MBytes  289 Mbits/sec
- - - - -

```

```

[ ID] Interval          Transfer    Bitrate
[ 5]  0.00-4.17    sec  0.00 Bytes  0.00 bits/sec          sender
[ 5]  0.00-4.17    sec  143 MBytes  288 Mbits/sec         receiver
iperf3: interrupt - the client has terminated
hxd@morgana /tmp [1]> iperf3 -c 192.168.50.159 -t 9999
Connecting to host 192.168.50.159, port 5201
[ 5] local 172.17.1.10 port 34608 connected to 192.168.50.159 port 5201
[ ID] Interval          Transfer    Bitrate      Retr  Cwnd
[ 5]  0.00-1.00    sec  48.1 MBytes  403 Mbits/sec   35   508 KBytes
[ 5]  1.00-2.00    sec  51.8 MBytes  434 Mbits/sec    0   577 KBytes
[ 5]  2.00-3.00    sec  55.4 MBytes  465 Mbits/sec    0   643 KBytes
^C[ 5]  3.00-3.69    sec  38.0 MBytes  462 Mbits/sec    0   684 KBytes
- - - - -
[ ID] Interval          Transfer    Bitrate      Retr
[ 5]  0.00-3.69    sec  193 MBytes  439 Mbits/sec   35          sender
[ 5]  0.00-3.69    sec  0.00 Bytes  0.00 bits/sec          receiver
iperf3: interrupt - the client has terminated

```

```

root@www ~# iperf3 -s
-----
Server listening on 5201 (test #1)
-----
^[[1;5CAccepted connection from 192.168.0.2, port 42948
[ 5] local 192.168.50.159 port 5201 connected to 192.168.0.2 port 42964
[ ID] Interval          Transfer    Bitrate      Retr  Cwnd
[ 5]  0.00-1.00    sec  36.2 MBytes  304 Mbits/sec   99   514 KBytes
[ 5]  1.00-2.00    sec  33.4 MBytes  280 Mbits/sec   64   419 KBytes
[ 5]  2.00-3.00    sec  34.5 MBytes  289 Mbits/sec    0   477 KBytes
[ 5]  3.00-4.00    sec  35.0 MBytes  294 Mbits/sec    0   528 KBytes
[ 5]  3.00-4.00    sec  35.0 MBytes  294 Mbits/sec    0   528 KBytes
- - - - -
[ ID] Interval          Transfer    Bitrate      Retr
[ 5]  0.00-4.00    sec  145 MBytes  304 Mbits/sec  163          sender
iperf3: the client has terminated
-----
Server listening on 5201 (test #2)
-----
Accepted connection from 192.168.0.2, port 34604
[ 5] local 192.168.50.159 port 5201 connected to 192.168.0.2 port 34608
[ ID] Interval          Transfer    Bitrate

```

```

[ 5] 0.00-1.00 sec 44.8 MBytes 375 Mbits/sec
[ 5] 1.00-2.00 sec 52.2 MBytes 438 Mbits/sec
[ 5] 2.00-3.00 sec 54.5 MBytes 457 Mbits/sec
[ 5] 2.00-3.00 sec 54.5 MBytes 457 Mbits/sec
- - - - -
[ ID] Interval          Transfer      Bitrate
[ 5] 0.00-3.00 sec    191 MBytes  533 Mbits/sec          receiver
iperf3: the client has terminated
-----
Server listening on 5201 (test #3)
-----

```

Some requests from Morgana to WWW are very slow, like 45KB/s while most are fast 3MB/s:

```

146.70.189.27:36190 192.168.50.159:443 ESTABLISHED 0s 82 KB/s

hxd@morgana /tmp> time curl -o /dev/null https://wiki.hexadust.net/attachments/15
% Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
           Dload  Upload   Total      Spent    Left     Speed
100 3460k  100 3460k    0     0  53464      0  0:01:06  0:01:06  --:--:--  78502

-----
Executed in 66.28 secs    fish           external
   usr time 100.61 millis  56.00 micros  100.55 millis
   sys time  60.88 millis 705.00 micros   60.17 millis

hxd@morgana /tmp> time curl -o /dev/null https://wiki.hexadust.net/attachments/15
% Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
           Dload  Upload   Total      Spent    Left     Speed
100 3460k  100 3460k    0     0  2588k      0  0:00:01  0:00:01  --:--:--  2590k

-----
Executed in 1.35 secs    fish           external
   usr time  52.86 millis   0.00 micros   52.86 millis
   sys time  29.61 millis 748.00 micros  28.86 millis

hxd@morgana /tmp> time curl -o /dev/null https://wiki.hexadust.net/attachments/15
% Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
           Dload  Upload   Total      Spent    Left     Speed
75 3460k  75 2598k    0     0  55139      0  0:01:04  0:00:48  0:00:16 28232^C

```

---

```
Executed in 48.58 secs fish external
usr time 71.12 millis 567.00 micros 70.56 millis
sys time 54.87 millis 196.00 micros 54.67 millis
```

Requests to kernel.org are fine.

Requests from sanbox-gw are also slow - they go via Mullvad that is deployed on sandbox-gw (not via justine):

```
146.70.189.27:35562 192.168.50.159:443 ESTABLISHED 0s 48 KB/s
> time curl -o /dev/null https://wiki.hexadust.net/attachments/15
% Total % Received % Xferd Average Speed Time Time Time Current
Dload Upload Total Spent Left Speed
73 3460k 73 2544k 0 0 59686 0 0:00:59 0:00:43 0:00:16 62609^C
```

---

```
Executed in 44.33 secs fish external
usr time 67.40 millis 676.00 micros 66.72 millis
sys time 37.12 millis 0.00 micros 37.12 millis
```

Using local route to go directly to caroline does fix the issue, so going through the internet is where the slowdown happens.

```
ip route add 46.7.126.16 dev wlp52s0 via 192.168.0.1
```

Things to try:

1. Change VPN exit node
2. Try access from non-VPN connection

Fast.com (over VPN) shows 300Mbit down and 1.2 Mbit up - this explains slow download from the server (as it is upload to the internet). Without VPN I get 300Mbit / 48Mbit - so the slow down is due to VPN.

Ping with max MTU:

```
ping -4 -M do -c 20 -s 1392 jpastuszek.net
```

Detailed connection info:

```
ss -ntpi
```

Update 2024-10-11:

- MTU looks good and is 1420 with MSS 1368.
- Ping packet loss is 30-40%
- When connections is slow we get 16% bytes re-transmitted; good connection has 0.2%
- Bad ping packet loss is on both IE endpoints

FIX: Changing VPN server did the trick.

## Slow Jitsi Meet transmission

Sending out 170KB/s and receiving stream of 64KB/s (512kbit/s).

- Ann is set up to statically route to video.hexadust.net (public IP) to caroline (default G/W is justine)
- I have noticed that UDP 10000 traffic goes directly to SERV VLAN (yunohost/192.168.50.137) so it bypasses the static route and ends up going to justine
  - Added additional static route to forward directly to caroline and confirmed with tcpdump (ether host) that it works
  - This did not improve traffic/quality
  - It would have been going to justine and there to caroline I suppose, so only extra hop; it would be NAT'ed on justine though: `iptables -t nat -A POSTROUTING ! -d 192.168.0.0/24 -o enp1s0 -j MASQUERADE`
  - How did Chrome know about local SERV IP? WebRTC shares the IP?

For morgana UDP video traffic goes:

1. To VPN connection to Justine
2. On exits on justine and is NAT'ed through HOME VLAN with destination IP of yunohost
3. It should go to caroline
4. Caroline routes it over SERV VLAN to yunohost

With static route I can bypass VPN tunnel to justine and go directly to caroline to be VLAN routed to yunohost

## No 5GHz WiFi

```
phy0-ap0: DFS-RADAR-DETECTED freq=5520
```

I got radar detection triggered and after 30 minutes it tried to turn it back off but was failing; manually restarting the interface fixed the issue:

```
Mon Nov 11 20:31:03 2024 daemon.notice hostapd: phy0-ap0: DFS-NOP-FINISHED freq=5500
ht_enabled=0 chan_offset=0 chan_width=0 cf1=5500 cf2=0
```

```
Mon Nov 11 20:31:03 2024 daemon.err hostapd: could not get valid channel
Mon Nov 11 20:31:03 2024 daemon.notice hostapd: phy0-ap0: interface state DFS->DFS
Mon Nov 11 20:31:03 2024 daemon.notice hostapd: phy0-ap0: DFS-NOP-FINISHED freq=5520
ht_enabled=0 chan_offset=0 chan_width=0 cf1=5520 cf2=0
Mon Nov 11 20:31:03 2024 daemon.notice hostapd: phy0-ap0: interface state DFS->DFS
Mon Nov 11 20:31:03 2024 daemon.notice hostapd: phy0-ap0: DFS-CAC-START freq=5520 chan=104
sec_chan=-1, width=0, seg0=114, seg1=0, cac_time=60s
Mon Nov 11 20:31:03 2024 daemon.err hostapd: 20/40 MHz: center segment 0 (=114) and center
freq 1 (=5510) not in sync
Mon Nov 11 20:31:03 2024 daemon.err hostapd: Can't set freq params
Mon Nov 11 20:31:03 2024 daemon.err hostapd: DFS start_dfs_cac() failed, -1
Mon Nov 11 20:31:03 2024 daemon.notice hostapd: phy0-ap0: DFS-NOP-FINISHED freq=5540
ht_enabled=0 chan_offset=0 chan_width=0 cf1=5540 cf2=0
Mon Nov 11 20:31:03 2024 daemon.notice hostapd: phy0-ap0: interface state DFS->DFS
Mon Nov 11 20:31:03 2024 daemon.notice hostapd: phy0-ap0: DFS-CAC-START freq=5520 chan=104
sec_chan=-1, width=0, seg0=114, seg1=0, cac_time=60s
Mon Nov 11 20:31:03 2024 daemon.err hostapd: 20/40 MHz: center segment 0 (=114) and center
freq 1 (=5510) not in sync
Mon Nov 11 20:31:03 2024 daemon.err hostapd: Can't set freq params
Mon Nov 11 20:31:03 2024 daemon.err hostapd: DFS start_dfs_cac() failed, -1
```

This is a reported issue:

- <https://forum.openwrt.org/t/5ghz-vht160-on-channel-100-breaks-on-radar-dfs-event-and-never-comes-back/165274>
- <https://forum.openwrt.org/t/center-segment-0-106-and-center-freq-1-5510-not-in-sync/164185>

I have changed from channel `104` to `100` which is the start of 160MHz wide channel that I use: [https://en.wikipedia.org/wiki/List\\_of\\_WLAN\\_channels](https://en.wikipedia.org/wiki/List_of_WLAN_channels). Not sure if this will help, the Wikipedia table suggest the center frequency for 160MHz wide channel is 114 but it cannot be selected. Looks like channel list is made for 20MHz wide channels. So selection 100 puts it as the first 20MHz channel for the 160MHz channel. *Channel Analysis* shows that my radio uses channels 100 through to 128 which matches the table on Wikipedia.

---

Revision #64

Created 2023-04-10 18:02:04 IST by hxd

Updated 2024-11-11 21:53:53 GMT by hxd