Measuring Broadband New Zealand

Fibre Max Review

December 2020

This document describes in detail the investigation carried out by ComCom, SamKnows and the industry into the performance of Fibre Max lines in New Zealand. It also details the changes made and the subsequent improvements in performance seen.

Background

Measuring Broadband New Zealand was launched in 2018 with the aim of providing consumers with accurate information on how their internet speeds were performing.

A range of RSPs, products and technologies are measured as part of the study and reports detailing performance over selected months are released on a quarterly basis.

One plan in particular has stood out in the reports published so far – Fibre Max.

The fastest plan available to New Zealanders has a very varied distribution of speeds as shown in the chart below. Only 40.6% of all tests recorded were above 800Mbps for Fibre Max compared to 74% of tests being above 100Mbps on Fibre 100. These results were a discussion point in the collaborative meetings, and it was decided to form a working group to investigate what might be causing lower than expected performance.

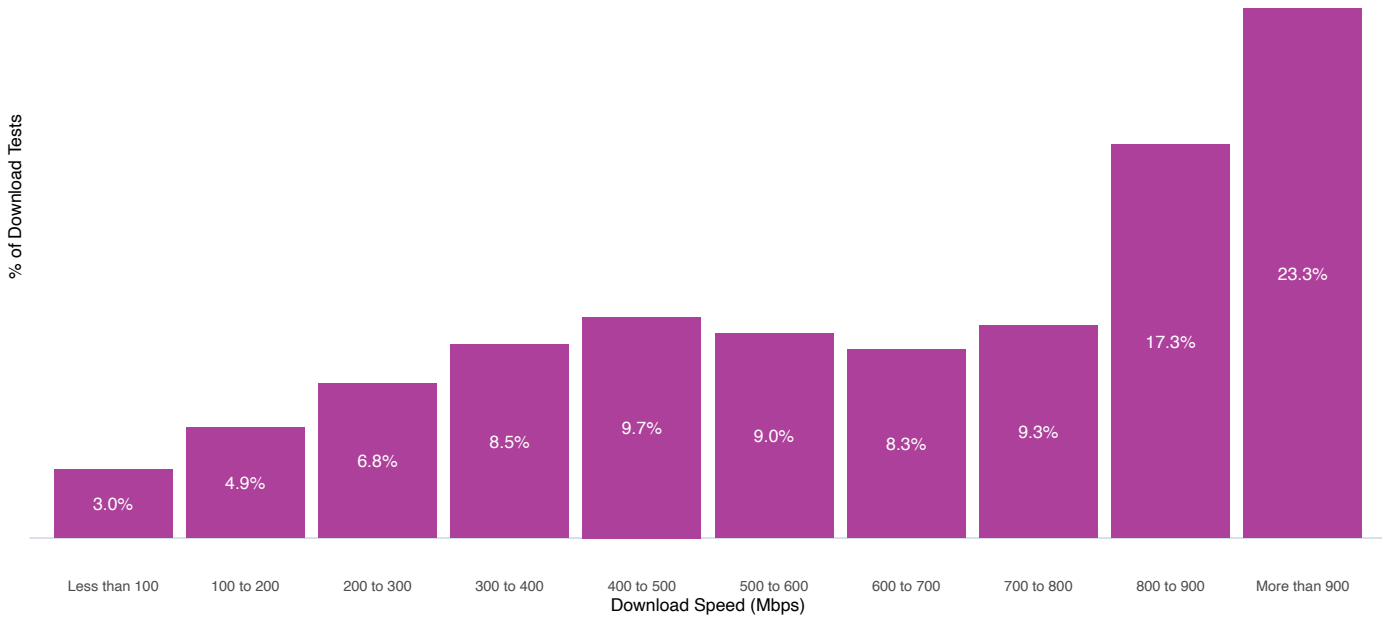

Figure 1: Distribution of download speed test results on the Fibre Max plan (from Winter Report August 2020)

The working group participants included RSP X and LFC Y. The Commission was happy to provide Whiteboxes for setup in a lab environment and SamKnows allocated a point of contact to interpret results and manage any changes required in the test environment.

Executive Summary of findings

The investigation has been a success and multiple parties have identified changes that could be made to improve performance for Fibre Max subscribers.

SamKnows would like to thank the members of the working group for the time and effort spent to identify the changes required and address these. Without the commitment from the industry and working group these may have remained unresolved, harming both customer experience and perception of the Fibre Max product.

A number of issues have been identified during the course of the investigation, including differing ONT performance, peering issues and a potential Linux kernel bug. Most of these have been resolved already and fixes are under development for the remaining ones.

SamKnows believes that one factor in particular, identified by both RSP X and LFC Y, is the main driver for the varied performance seen on the Fibre Max tier. This relates to packet loss occurring under high-burst conditions on certain models of ONT. This was caused by burst parameter configuration, rather than any congestion factors. This affected broadband connections with higher round-trip latencies to the test servers much more than connections with lower latencies.

RSP X has made two key changes to their network in response to this discovery. This has resulted in RSP X's Fibre Max performance improving to levels comparable with 1Gbps services in other countries. LFC Y is looking at a design change that, for some RSPs depending on their specific network configuration, may improve Fibre Max performance.

Figure 2 below demonstrates the impact of RSP X's changes. The chart shows individual download speed test results from three Auckland-based Whiteboxes in RSP X homes that were identified as underperforming. These are from the MBNZ measurement panel and are not lab devices. No changes were made to the test configuration of these Whiteboxes or test servers during the period shown. The step-change in download speed from an erratic 400-700Mbps prior to RSP X's changes in early September to a reliable 900Mbps+ afterwards is clearly visible.

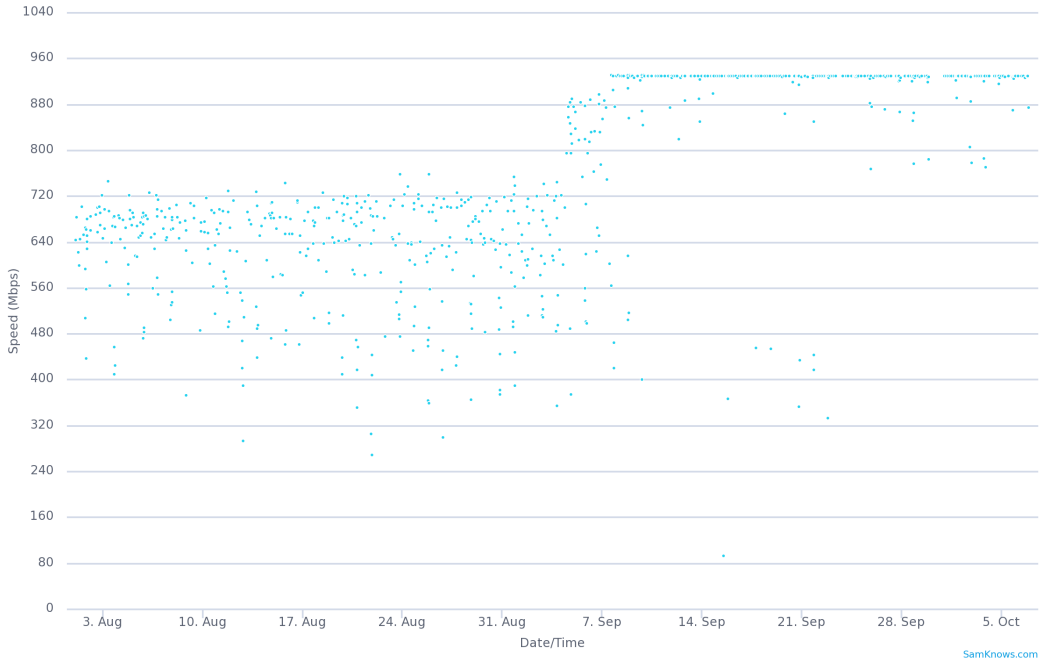

Figure 2: Download speeds from three Auckland Whiteboxes on RSP X Fibre Max RSP X's changes occurred on September 4th and 7th.

Executive Summary of findings (continued)

These network configuration changes have also improved the performance of real applications too. One of the tests that the SamKnows Whitebox carries out is a download test from Netflix's content delivery network, which often utilises servers hosted by the RSPs. In this test, SamKnows has no control over the server-side configuration (such as the TCP congestion control algorithm) or server hosting locations. The Netflix download speed measurements have shown a similarly significant improvement since the changes in early September.

Additional issues have been identified and resolved by multiple parties in the course of this investigation. RSP X implemented a routing change in late June 2020 that ensured that traffic from the South Island had a more optimal traffic path to the Wellington SamKnows server at REANNZ. A peering issue was identified at APE in September 2020 and resolved later the same month. SamKnows identified a likely bug with the Linux kernel running on the Wellington test server, and deployed a kernel update in November 2020 that resolved the issue. A number of other potential causes of the variable performance of Fibre Max were also explored in the course of this investigation, and these are discussed in the full report below.

SamKnows' Fibre Max performance investigation

The first area of investigation was the SamKnows measurement platform itself. Whilst the same SamKnows hardware, software and configuration is used in multiple other countries with consistent 940Mbps results, it is still important to acknowledge the possibility that some previously unseen combination of factors could be causing the variable results seen in New Zealand.

Deployment overview prior to the investigation

Before discussing the investigation in-depth, it is worth recapping the deployment and configuration of the measurement platform.

SamKnows Whiteboxes - small embedded Linux devices - are deployed to New Zealand broadband consumers nationwide. Consumers are asked to install the Whitebox inside their home, connected to their router via an ethernet cable. The Whitebox is aware of 'cross-traffic' inside the home and does not run measurements when other users are using the broadband connection.

The Whiteboxes run frequent measurements, including download speed tests, to measurement servers in New Zealand as well as real applications such as Netflix and YouTube. The speed test measurements operate over TCP using eight parallel TCP sessions, and last for ten seconds.

The speed tests used in the public reporting run against powerful dedicated servers hosted by REANNZ in Wellington and Auckland. These servers are provisioned at 10Gbps.

Whitebox 8.0 hardware

The current generation of Whitebox (8.0) is capable of measuring 1Gbps downstream and upstream over both TCP and UDP.

A minority of Whiteboxes deployed in New Zealand already reported speeds of up to 940Mbps consistently, which suggested that the Whitebox itself was not the limiting factor. The same Whitebox 8.0 hardware, firmware and software is used across all deployments globally. This hardware is also used in multiple other countries with FTTH and DOCSIS 3.1 services, and many of these consistently report speeds of up to 940Mbps.

An earlier investigation into underperformance in 2019 looked at the possibility that insufficiently powerful routers and other in-home factors could be driving the low speeds on Fibre Max. There was no conclusive evidence to demonstrate the routers were the cause of the speeds seen. In order to rule out other in-home factors, SamKnows requested that RSPs and LFCs host a Whitebox in their labs on a well-provisioned Fibre Max connection behind a representative set of CPE (Customer Premises Equipment). RSP X, LFC Y and LFC Z volunteered to assist and they were sent Whiteboxes to install in their network. The same underperformance was seen on the lab Whiteboxes, which provided a stable basis to continue the investigation.

Whitebox 8.0 hardware (continued)

Before the Whitebox was finally ruled out as a potential factor, SamKnows also requested that RSPs and LFCs provide access to a modern laptop or embedded device installed alongside the Whitebox upon which we could run comparable measurements. RSP X provided access to an Odroid C2 embedded device, installed alongside the Whitebox, which SamKnows ran a comparable set of measurements on to the same REANNZ test servers. The Odroid C2 saw the same underperformance as the Whitebox, which confirmed that the Whitebox was not the limiting factor.

SamKnows' measurement configuration

SamKnows experimented with a wide range of configuration changes to the testing to see if any yielded a significant improvement in performance. These were tested using the lab hosted Whiteboxes from RSP X and the LFCs. Briefly, the changes tested were:

- Testing using alternative clients, including a simple cURL call and iperf3
- Increasing the number of concurrent TCP connections from our standard of 8 up to 64 (and various increments between)
- Using alternative TCP congestion control algorithms
- Testing to alternative test servers, including Cloudflare (Auckland) and various Speedtest.net servers. Cloudflare uses TCP BBR across all of its servers.
- Using non-Whitebox hardware, such as the Odroid C2 referenced in the previous section
- Using shorter and longer measurement durations than our standard ten second test; and
- Using different port numbers to rule out any unexpected DPI (Deep Packet Inspection) behaviour

Whilst some changes improved download speed and others harmed it, none yielded the desired step-change to the expected stable >=930Mbps throughput.

It was observed that connections with lower round-trip time to the test servers (2-5ms) tended to yield significantly better results than connections with higher round trip time (5-20ms). The lower latency connections would often yield performance between 800Mbps and 900Mbps. Whilst latency is certainly a factor in speed test results, it is not unreasonable to expect a consistent 940Mbps at 20ms RTT - SamKnows sees this routinely in other countries. A number of factors can impact latency, including physical distance (which cannot be overcome), peering/routing relationships and congestion factors.

Test Servers

Speed tests in the Measuring Broadband New Zealand project make use of test servers hosted by REANNZ in Wellington and Auckland. These have 10Gbps interfaces.

Traffic statistics that SamKnows collect every minute from these servers demonstrated that they were nowhere near capacity. They would very rarely burst above 1Gbps, and never more than 2Gbps.

Moreover, SamKnows ran measurements between the REANNZ-hosted test servers and other servers elsewhere which are also provisioned at 10Gbps. These tests yielded results in excess of 3.5Gbps using a single TCP flow, demonstrating that the servers and the immediate connectivity provided by REANNZ was sufficient to deliver at least 1Gbps.

The lab testing conducted with RSP X also provided confidence that the REANNZ servers were not the cause of the underperformance of Fibre Max. These lab Whiteboxes were configured to test against the REANNZ servers, Cloudflare (hosted in Auckland, using TCP BBR) and a small selection of speedtest.net (Ookla) test servers. These all showed similarly poor results.

Analysis of the timeline, resolution and other factors

Note: RSP X have requested details of the changes made to their network be redacted from the report and as such these changes are not described in detail here.

In late August 2020, RSP X discovered a higher than expected rate of packet discards under certain conditions that they suspected could be harming performance. They made a change based on this theory and deployed it to a single RSP X broadband connection, which resulted in its measured download speeds improving from 300-500Mbps to around 900Mbps.

On September 4th 2020, RSP X deployed the change across their whole customer-base (not just their lab devices). We will refer to this below as Change 1.

On September 7th 2020, RSP X deployed a further change for their whole customer-base. We will refer to this below as “Change 2”.

Before looking at any changes after September 7th and discussing any other factors, we will review the impact of these two changes. We will do this by looking at a small selection of RSP X Whiteboxes from the wider MBNZ panel, which RSP X does not have access to. By using these devices rather than the lab devices, we can see the real-world impact of the changes. No changes to the measurements, configuration or test servers were made during this period, so it is only the impact of RSP X’s changes that we will be examining.

Figure 3 below shows download speed test results for three Auckland-based RSP X Whiteboxes between August 2020 and October 2020. These three Whiteboxes had between 5ms and 12ms RTT to the Auckland test server. We can observe two clear step changes in the results. The first starts around 3pm local time on the 4th September (after ‘Change 1’), and results in the majority of speed tests improving to between 800 and 900Mbps. The second starts around 4pm on the 7th September (after ‘Change 2’), and results in almost all speed tests recording between 929 and 931Mbps.

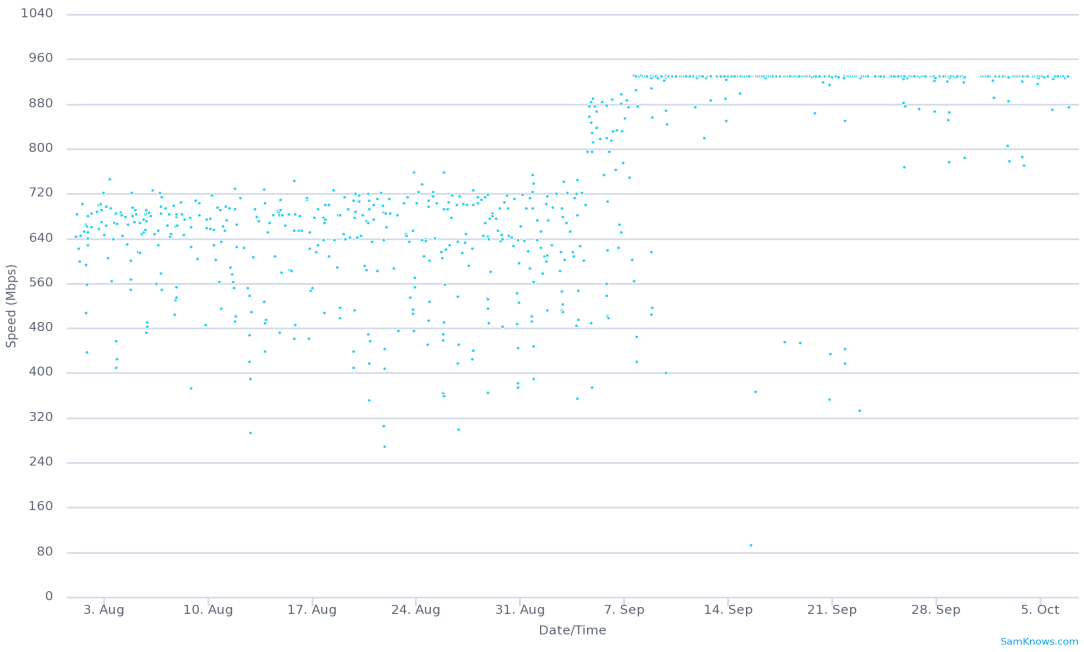

Figure 3: Download speeds from three Auckland Whiteboxes on RSP X Fibre Max. These three Whiteboxes had RTTs of 5-12ms.

Analysis of the timeline, resolution and other factors (continued)

We can observe two clear step changes in the results. The first starts around 3pm local time on the 4th September (after 'Change 1'), and results in the majority of speed tests improving to between 800 and 900Mbps. The second starts around 4pm on the 7th September (after 'Change 2'), and results in almost all speed tests recording between 929 and 931Mbps.

Figure 4 below looks at three different RSP X Fibre Max Whiteboxes. These three Whiteboxes had between 1ms and 2.5ms RTT to the test server, and had far better results prior to September than the Whiteboxes from Figure 3. Whilst it is much harder to discern than in Figure 3, there is a clear stabilising of the results on the 4th of September and again on the 7th of September following RSP X's changes.

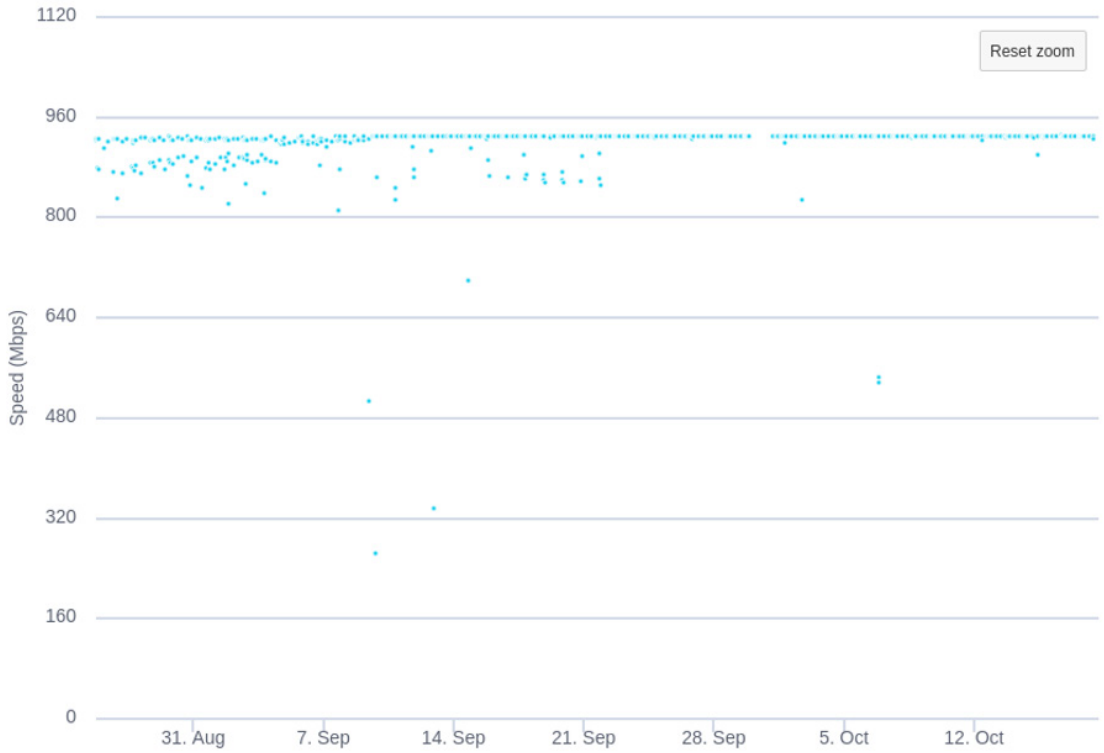

Figure 4: Download speeds from three different Auckland Whiteboxes on RSP X Fibre Max. These three Whiteboxes had RTTs of 1-2ms.

We suspect that the lower RTT for the Whiteboxes in Figure 4 is the reason why they performed better prior to September than the higher RTT Whiteboxes in Figure 3. A lower RTT alone does not provide improved speed. Instead, we believe this to be related to the packet loss issue driven by the burstiness parameters. When a TCP receiver detects packet loss, it will send a TCP retransmission request to the sender (in the form of a duplicate acknowledgement). Because the Whiteboxes in Figure 4 had a lower RTT, their retransmission requests reached the sender (the server) faster than the Whiteboxes from Figure 3, which allowed the missing data to be retransmitted faster, which ensured that throughput remained quite high.

To be clear, the takeaway here should not be that 2.5ms RTT is good and 10ms RTT is bad - both are perfectly capable of sustaining 940Gbps throughput. But in the presence of packet loss or other impairments, the link with the lower RTT will likely recover faster.

Analysis of the timeline, resolution and other factors (continued)

Moreover, we can use a peering issue that occurred later in September to demonstrate that latency does not play a significant role now that 'Change 1' and 'Change 2' have been applied. Figure 5 below shows the same as Figure 3, except split by test server. Again, this is based on measurements from three RSP X Auckland Whiteboxes on Fibre Max. We can see that measurements switch to using the Wellington test server between September 8th and September 23rd. This was driven by RSP X investigating an ongoing issue in the APE (Auckland Peering Exchange) network which was confirmed by temporarily shutting down the link. This issue was subsequently identified and resolved by APE, and the link re-established.

During this period, RTT to the Wellington server was lower than Auckland, so the Whiteboxes switched to using Wellington. The RTTs were consistently between 15ms and 22ms to Wellington. As we can see from Figure 5 below, even with the extra 10ms RTT, throughput remains at around 930Mbps between September 8th and 23rd.

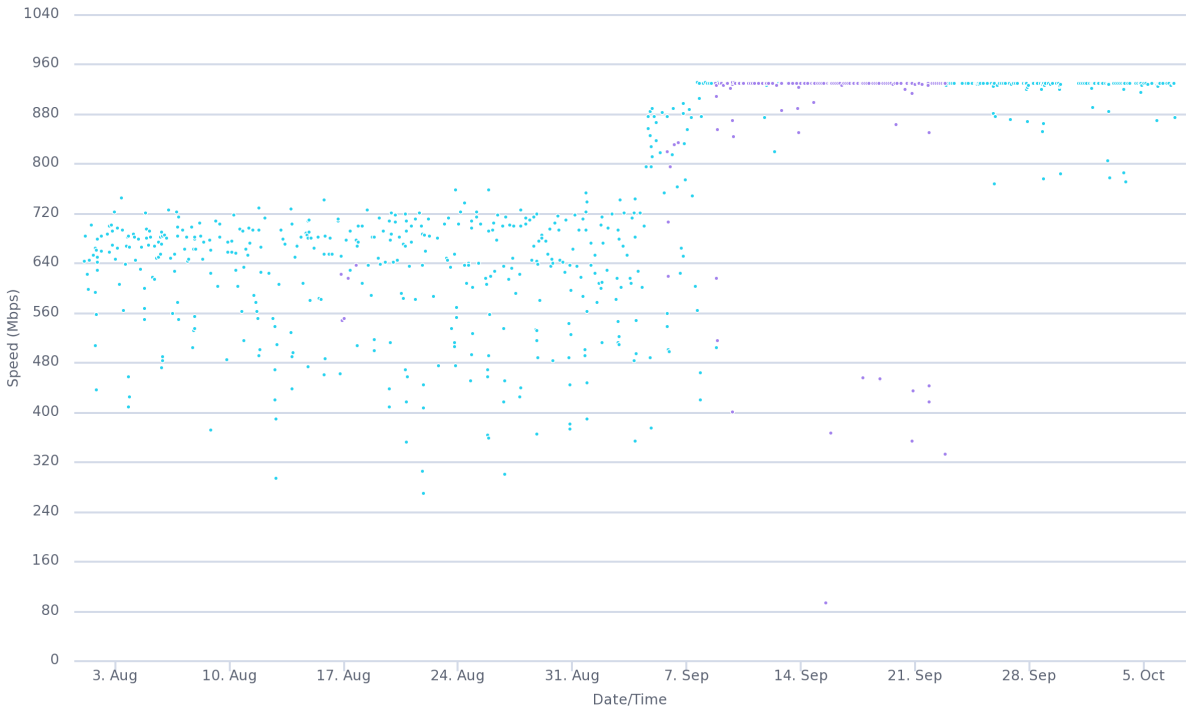

Figure 5: Download speeds from three Auckland Whiteboxes on RSP X Fibre Max. These three Whiteboxes had RTTs of 5-12ms to Auckland, and 15-22ms to Wellington. After the changes in early September, a stable 930Mbps was achieved to both Auckland and Wellington, even with the higher RTT.

It is worth stressing that whilst there was a peering issue in Auckland between September 8th and 23rd, this is tangential to the main outcome of this investigation. The performance improvements were seen a day before September 8th and remained consistent after local connectivity was restored on the 23rd.
To be clear, the takeaway here should not be that 2.5ms RTT is good and 10ms RTT is bad - both are perfectly capable of sustaining 940Gbps throughput. But in the presence of packet loss or other impairments, the link with the lower RTT will likely recover faster.

Lastly, we look to see how these changes have affected Netflix performance. SamKnows have no control over the Netflix server configuration (TCP congestion control algorithm or other parameters, peering, or placement of the server), so this can be considered an even cleaner test than those carried out above.

Analysis of the timeline, resolution and other factors (continued)

Figure 6 below shows the percentage of download speed results over 800Mbps from the nearest Netflix cache across all RSP X Fibre Max Whiteboxes between August 15th-31st and between September 15th-30th. These dates ranges were chosen because they straddle the period when RSP X made their changes (September 4th and 7th).

The results demonstrate that a much higher percentage of tests achieved 800Mbps or greater from Netflix since the changes in early September. 12.6% of tests achieved 800Mbps or greater from Netflix prior to the changes, and this increased to 46% after the changes.

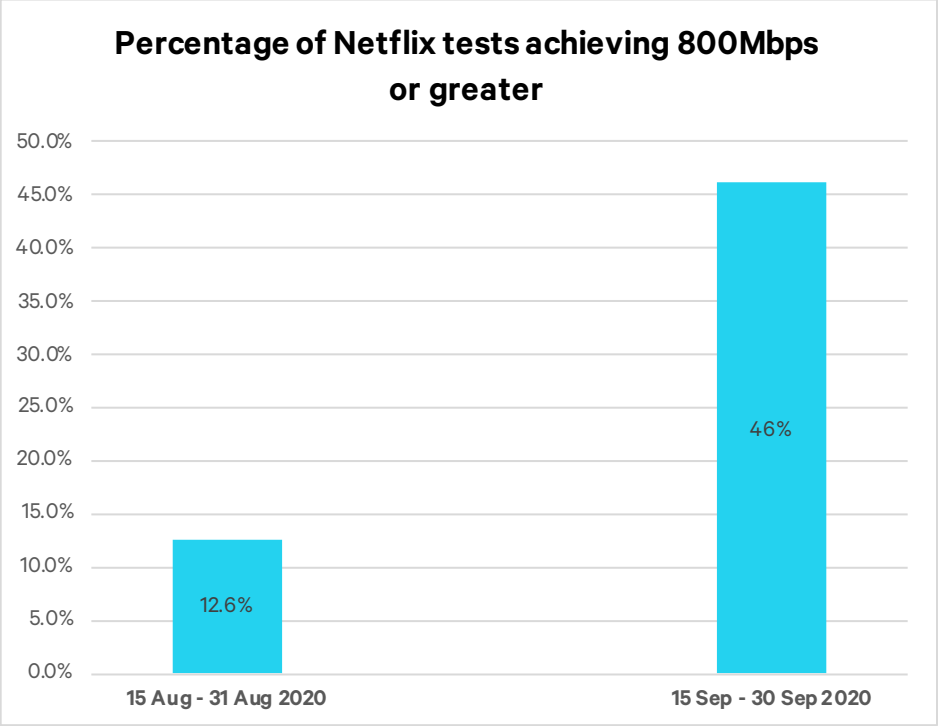

Figure 6: Percentage of Netflix tests achieving 800Mbps or greater for RSP X Fibre Max Whiteboxes between 15th August 2020 - 31st August 2020 (before the changes) and 15th September 2020 – 30th September 2020 (after the changes). The changes by RSP X in early September significantly improved throughput from Netflix.

It is worth stressing the difference between the regular SamKnows multi-TCP connection speed tests and the Netflix speed measurement discussed above. Whilst at first, 46% of Netflix speed tests achieving 800Mbps or greater may still sound poor, this is not the case. This measure of download speed from Netflix is conducted over a single TCP connection, and is therefore much more sensitive to latency and packet loss than the regular SamKnows multi-TCP connection speed tests. Moreover, the traffic is served from the same Netflix caches that real Netflix users use, so there could be server-side congestion or other factors impairing performance that we are unaware of.

The increase from 12.6% to 46% of Netflix tests achieving 800Mbps or greater demonstrates that RSP X's changes have led to significant increases in throughput for third party services, and not just speed test servers.

Single TCP performance to the Wellington Test Sever

As a part of the joint investigation, RSP X was carrying out single TCP connection speed tests to the Wellington and Auckland REANNZ test servers. Whilst single TCP connection speed tests do not form part of the Measuring Broadband New Zealand study, they are very useful for investigating performance problems. This is because testing with multiple TCP connections tends to mask performance drops (any one TCP connection experiencing packet loss and its subsequent slow-down will be offset by the other TCP connections filling the newly available capacity).

RSP X reported seeing unexplained speed drops when testing against the Wellington SamKnows server with a single TCP connection. These speed drops were seemingly random but visible with repeat testing. SamKnows investigated with the help of REANNZ, who provided additional infrastructure to track down this issue. With REANNZ's help, any issues with their peering or their network hardware were ruled out, and the problem was eventually reproduced reliably.

Experimentation revealed that our higher-than-default TCP send buffer size (wmem max) value was the likely cause of the speed drops. These had not been seen anywhere else globally. The default for the Linux kernel version was 4 MiB, whilst the SamKnows test server configuration increased this to 8 MiB to allow for broadband connections with a very high bandwidth-delay-product. Reducing the value back to 4 MiB seemed to eliminate the issue, but this was an unsatisfactory resolution, and we expected a possible TCP stack bug.

The additional server provided by REANNZ was upgraded to the latest Linux kernel (5.9.6), retaining the same TCP configuration, and was re-tested. The intermittent performance drops had disappeared. Following this, the live Wellington and Auckland REANNZ servers were updated to the new Linux kernel on November 10th 2020.

Figure 7 below shows the impact on individual multi-TCP measurements from three Wellington based Whiteboxes on RSP X Fibre Max. One of the three Whiteboxes intermittently saw results around 750Mbps prior to the change, and these disappeared after the kernel upgrade.

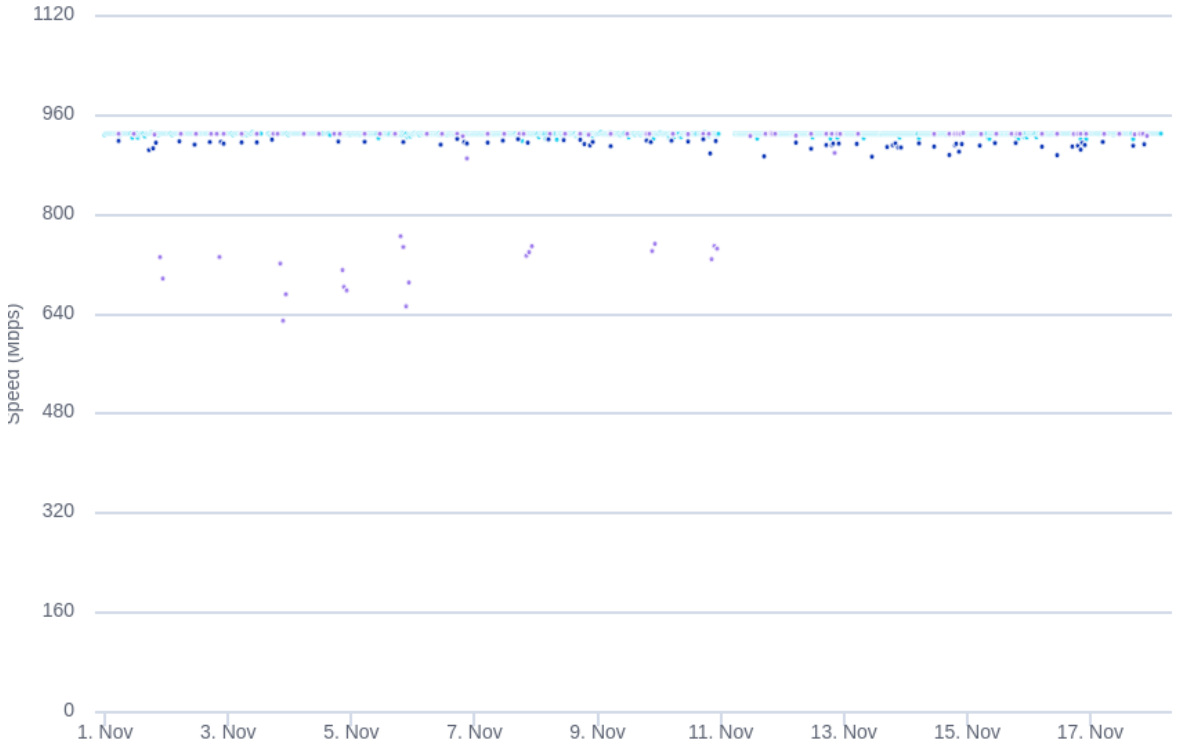

Figure 7: Download speeds from the Wellington test server for three RSP X Fibre Max Whiteboxes in early November 2020. A Linux kernel upgrade on November 10th appears to have eliminated intermittently low results for one of the Whiteboxes.

Future Work

SamKnows understands from LFC Y that they have identified an issue with packet discards under certain traffic conditions on their first and second-generation ONT. An update to rectify this is in development. LFC Y will implement a design change to shape downstream traffic at the OLT and this is expected to improve the Fibre Max performance for some RSPs depending on their specific network configuration.

Additionally, SamKnows is currently working with REANNZ to arrange the deployment of a third test server in Christchurch. This should reduce latency for users on the South Island, who will currently be using the Wellington test server primarily.

SamKnows is also happy to experiment with TCP congestion control algorithms to assess the impact of these on performance.

As stated, one of the aims of the MBNZ project is to improve the internet for New Zealand consumers. Additional work on investigations may be carried out in the future to assess other factors that may impact performance.