SAMKNOWS ANALYSIS OF BROADBAND PERFORMANCE IN CANADA OCTOBER & NOVEMBER 2015

Dated: March 17th, 2016

ELECTRONIC COPY

ISBN: 978-0-660-04966-3
Cat. No.: 978-0-660-04966-3

The Canadian Radio-television and Telecommunications Commission (CRTC)
Ottawa, Ontario
Canada
K1A ON2
Tel: 819-997-0313
Toll-free: 1-877-249-2782 (in Canada only)
https://services.crtc.gc.ca/pub/submissionmu/bibliotheque-library.aspx

Aussi disponible en français

Contents

About The Project

The Canadian Radio-Television and Telecommunications Commission (CRTC) has commissioned SamKnows to conduct a study of the performance of broadband services sold to Canadian consumers. SamKnows is a global leader in broadband measurement and has been working with governments, ISPs, content service providers, application developers, consumer groups, and academics to accurately measure Internet performance since 2009. In doing so, SamKnows has built a global Internet measurement platform, which now spans five continents and conducts many millions of measurements each day.

Data presented in this report was collected between 1st October 2015 and 30th November 2015.

4,486 Whiteboxes were deployed to Canadian volunteers as a part of this study. Data from 3,471 of these Whiteboxes was used in this report. Data from 1,015 Whiteboxes were omitted from this report either due to issues that arose with a third party carrying measurement traffic or due to the Whiteboxes being on ISPs not included in the collaborative group stated below. For purposes of reporting, the data was split into four speed tier "buckets": 5-9Mbps, 10-15Mbps, 16-39Mbps, and 40Mbps+. The Internet Service Providers (ISPs) participating in this project included all the major wireline service providers in Canada other than Sasktel. Specifically Bell, Bell Aliant, Cogeco, Eastlink, MTS, Northwestel, Rogers, Shaw, TELUS and Videotron all participated on a voluntary basis and measurements covered all geographic regions of Canada in a mix of urban and rural settings. These ISPs use technologies such as digital subscriber line (DSL) [1], hybrid-fibre co-axial cable (Cable / HFC) [2] and fibre to the home (FTTH) [3]. Testing has not included any ISPs using satellite or fixed wireless technologies nor did it include resellers of these ISPs' networks.

The test methodology employed is the same as the one SamKnows uses around the globe with other regulators and ISPs. A full description of the test methodology can be found here: https://www.samknows.com/broadband/uploads/methodology/SQ301-005-EN-Test-Suite-Whitepaper-4.pdf

SamKnows typically recommends that a minimum sample of 40 measurement probes should be reporting data per strata in order to provide sufficient statistical accuracy in the results. This is the approach taken in this report. Additional information on sample size methodology can be found here:

[1] This category comprises the technologies used to deliver digital data over copper lines. This includes fibre to the node (FTTN), which refers to the use of optical fibre to the neighbourhood and then a copper line to the customer's home.
[2] This category comprises the technologies used to deliver digital data over a hybrid-fibre co-axial network through DOCSIS platforms. This technology uses an optical fibre to the neighbourhood and then co-axial cable/HFC to the customer's home.
[3] This category comprises the technologies used to deliver digital data through an optical fibre directly to the customer's home.

https://www.samknows.com/broadband/uploads/methodology/SamKnows_Sample_Size_Whitepaper_20150610.pdf.

Any comments on the analysis in this document should be directed to Roxanne Robinson (roxanne@samknows.com).

A

Executive Summary

This report presents the preliminary findings of the measurement study that SamKnows is conducting in Canada on behalf of the CRTC. 4,486 SamKnows Whiteboxes have been deployed in Canadian homes across a range of ISPs and products. Each Whitebox conducts end-to-end performance measurements 24 hours a day, 7 days a week to destinations representative of Canadian consumers' Internet usage.

In this preliminary report, results are reported by technology, speed bucket and region. Technologies have been identified as DSL,[4] Cable/HFC (including all DOCSIS variants)[5] and FTTH[6]. The three regions represented are: (1) West & North (British Columbia, Alberta, Saskatchewan, Yukon, Northwest Territories, and Manitoba), (2) Central (Ontario and Quebec) and (3) East/Atlantic provinces (New Brunswick, Nova Scotia, Prince Edward Island, and Newfoundland and Labrador).

We have experienced some issues during testing with a third party that carried some of the measurement traffic. This affected the reliability of some measurement results. All such data was excluded from this report and the underlying issues have since been resolved. As a result of these difficulties we are providing this preliminary report on an aggregate basis and intend to issue another report later this year which will identify particular ISP's results based on more recent data. This preliminary report demonstrates that:

- Almost all broadband services met or exceeded their advertised speeds, regardless of the access technology in use.
- Even the highest latencies and packet loss exhibited during testing would be more than adequate for any common Internet application at present and latency and web-browsing in Canada compares favourably to those measured in other jurisdictions including the United States.

Unless otherwise stated, all results presented in this report are taken from the peak period, which is defined as 7-11pm local time on weeknights. Moreover, all reported figures are subject to a minimum sample size of 40 Whiteboxes. These provisions are consistent with those used in the FCC's Measuring Broadband America study.

The key findings contained in this preliminary report are as follows:

- The majority of broadband products sold in Canada met or exceeded their advertised download speeds. Cable/HFC and FTTH services delivered download speeds in excess of the rates advertised by ISPs. FTTH services delivered 119% of advertised download speed on average, whilst Cable/HFC services delivered 103%. Most DSL services met or exceeded the advertised rates too.

[4] This category comprises the technologies used to deliver digital data over copper lines. This includes fibre to the node (FTTN), which refers to the use of optical fibre to the neighbourhood and then a copper line to the customer's home.

[5] This category comprises the technologies used to deliver digital data over a hybrid-fibre co-axial network through DOCSIS platforms. This technology uses an optical fibre to the neighbourhood and then co-axial cable/HFC to the customer's home.

[6] This category comprises the technologies used to deliver digital data through an optical fibre directly to the customer's home.

- Performance was largely consistent across all regions, with the vast majority achieving between 109% and 122% of advertised download speed.

- Upload speeds also generally met or exceeded advertised rates. All Cable/HFC and FTTH services exceeded 100% of advertised upload rates. DSL services exhibited more volatility. DSL services in the 5-9Mbps bucket delivered 85% of advertised upload speeds.

- DSL services yielded the highest latency, with results ranging between 30.0ms in the 5-9Mbps bucket to 16.7ms in the 40Mbps+ bucket. Higher latencies with certain DSL products is an expected by-product of the access technology, coupled with the fact that they are often used to deliver broadband services to customers using longer lengths of copper lines. Cable/HFC services were more consistent, with all speed buckets delivering between 17.2ms and 21.7ms. FTTH services yielded the lowest latency at 11.5ms to 11.7ms.

- Latency varied significantly in Canada, with the central regions yielding the lowest latency at 15.2ms on average. The highest round-trip latency was observed in the eastern regions, which averaged 39.6ms. This variation was driven largely by the length of the network path between the client and the test server and the access technologies that were most commonly deployed in those regions.

- Packet loss, which describes how likely it is that a data packet sent from point A will not reach point B, was generally very low across all speed buckets, technologies and regions, although there were exceptions. FTTH services yielded the lowest levels of packet loss, at 0.04% on average. DSL services showed the highest levels of packet loss, with an average of 0.2%. Cable/HFC services averaged 0.11% packet loss. These levels of packet loss are extremely small and would be unnoticeable to any common Internet application.

- Web page loading times to a selection of websites popular in Canada improved as download speeds increased. However, this improvement is not linear. Certain DSL services (those in the 5-9Mbps bucket) loaded web pages in 2.2 seconds on average. The fastest services in the 40Mbps+ bucket load pages in just 0.8 seconds. As has been found in studies in other markets, improvements in page loading time tail off after 10Mbps (at which point latency becomes the dominant factor). There was minimal difference between access technologies when similar speed tiers are compared.

- Figure 1 below shows the fastest product included in this report for each technology, region and speed bucket. FTTH and Cable/HFC services offered the highest speeds, at 300Mbps and 250Mbps respectively. Tested DSL peaked at 50Mbps. The sample plan and deployment of Whiteboxes were designed to reflect the most popular products in each region. The values in Figure 1 are not necessarily reflective of the products available to Canadians in each region; they merely reflect the maximum advertised speeds obtained in the sample.

Region	DSL	Cable / HFC	FTTH
Central	**50**	**250**	**300**
5M-9M	6	7.5	n/a
10M-15M	15	15	15
16M-39M	25	35	25
40M+	50	250	300
East	**n/a**	**250**	**300**
40M+	n/a	250	300
West & North	**50**	**150**	**150**
5M-9M	7	7.5	n/a
10M-15M	15	15	10
16M-39M	25	30	25
40M+	50	150	150

Figure 1: Maximum advertised speed of tiers included in the test results by region and technology

B Key Performance Indicators

B.1 Download Throughput

Download throughput is the measure of the capacity of the user's broadband connection. Higher speeds are more desirable, as it allows the user to retrieve data (be it a web page, music file, or anything else) more quickly.

To characterize the user's maximum access link capacity, measurements were conducted between panelists' homes and the nearest test server. Test servers were deployed in multiple major metro areas throughout Canada.

It is common for broadband providers around the world to differentiate their product offerings by a headline access speed, and this is a key part of their advertising. Speeds are typically expressed in megabits per second (commonly abbreviated to 'Mbps' or 'Mbit/s'). To enable comparability between different products and technologies which may feature vastly different speeds, most regulators conducting measurement studies around the world report on the percentage of advertised speed that products and technologies achieve. In this study, multiple speed tiers are grouped together into 'buckets'.

Figure 2 shows the peak period speeds as a percentage of advertised, broken down by speed bucket and access technology. The vast majority of technologies and speed buckets met or exceeded advertised rates. Cable/HFC and FTTH services all exceeded the advertised rates. FTTH services delivered at least 116% of advertised, with Cable/HFC services ranging between 100% and 106%. DSL services showed more variance, with the 5-9Mbps bucket achieving 88% on average, whilst other DSL buckets very nearly met or exceeded the advertised rates. Variance for DSL is not surprising given the effect copper loop lengths (i.e. how far a customer is from the nearest central office or node) has on line performance.

Download speed as a percentage of advertised speed by technology and speed band

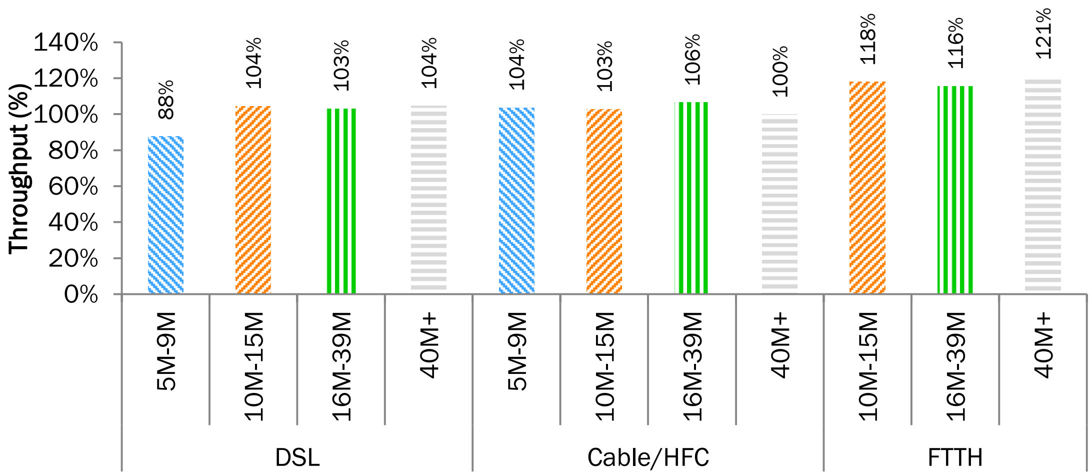

Figure 2: Download speed as a percentage of advertised speed by technology and speed bucket

Figure 3 depicts the average speed as a percentage of advertised achieved across the different regions in Canada. There is minimal variation between the majority of regions, with most achieving between 109% and 122% of advertised. The only exception is the 5-9Mbps bucket in the West & North region, where an average speed of 95% of advertised was reached.

Download speed as a percentage of advertised speed by region and speed band

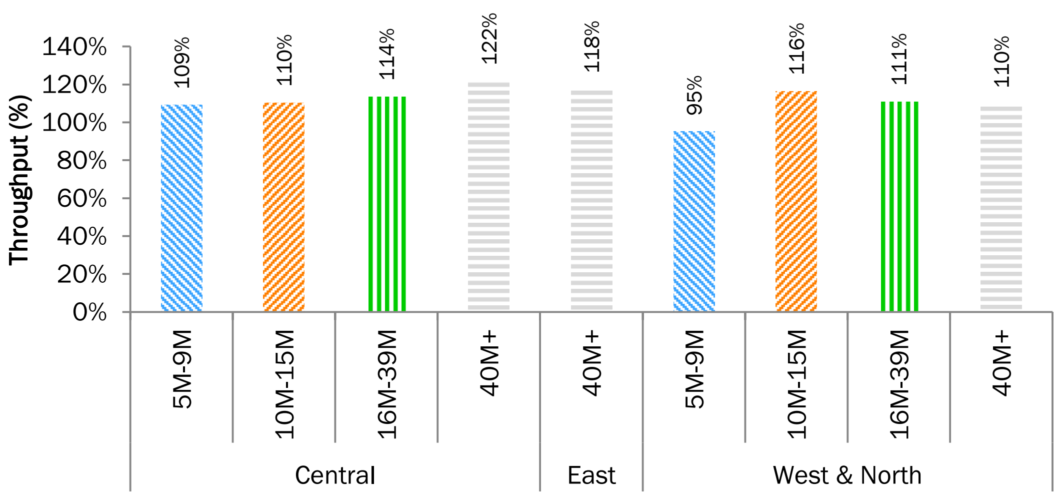

Figure 3: Download speed as a percentage of advertised speed by region and speed bucket

Performance can vary significantly by time of day, as Internet usage increases heavily amongst consumers during peak hours. Figures 4, 5 and 6 depict performance by time of day for DSL, Cable/HFC and FTTH services respectively.

The vast majority of speed buckets on all access technologies showed very stable throughput during all hours of the day. In general, the difference between peak and off-peak performance varied by less than 3 percentage points. However, there are some exceptions. DSL services in the 5-9Mbps bucket fell by 6 percentage points during peak hours. Cable/HFC services in the 40Mbps+ bucket fell from 108% off-peak to 99% during peak hours. In practical terms this fall is minimal, but it is indicative of possible congestion taking place during peak times. FTTH delivered the most consistent speeds, with typically less than 1% variation between peak and off-peak hours.

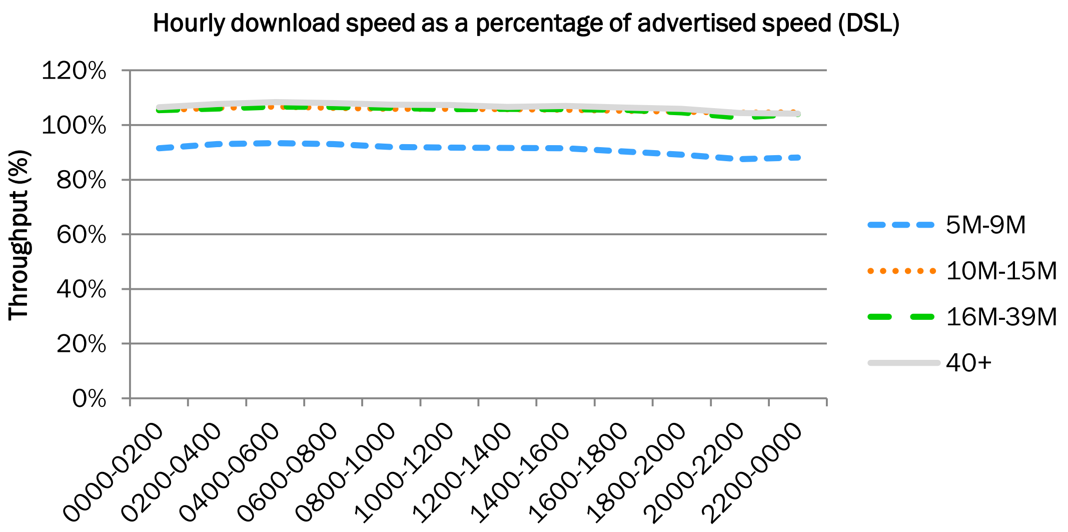

Figure 4: Hourly DSL download speed as a percentage of advertised speed

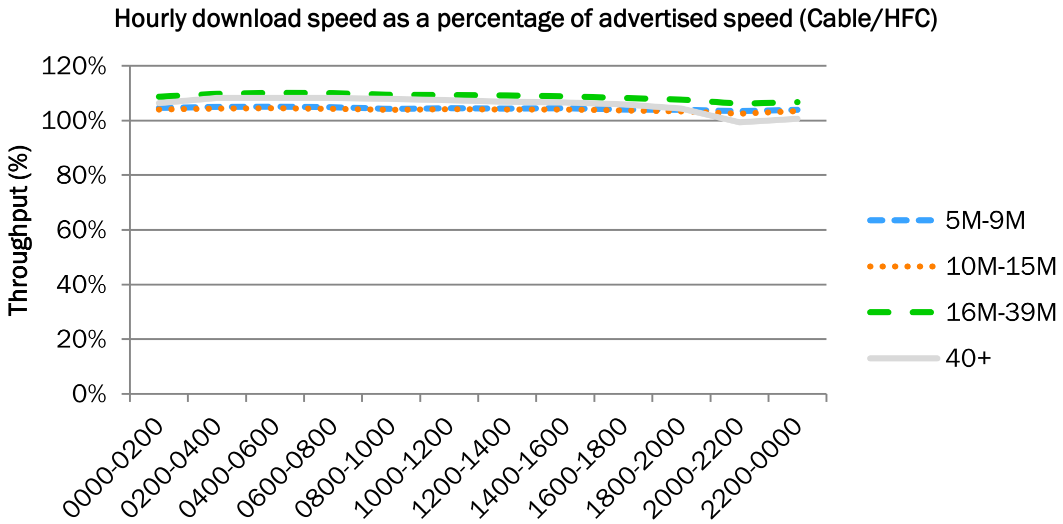

Figure 5: Hourly Cable/HFC download speed as a percentage of advertised speed

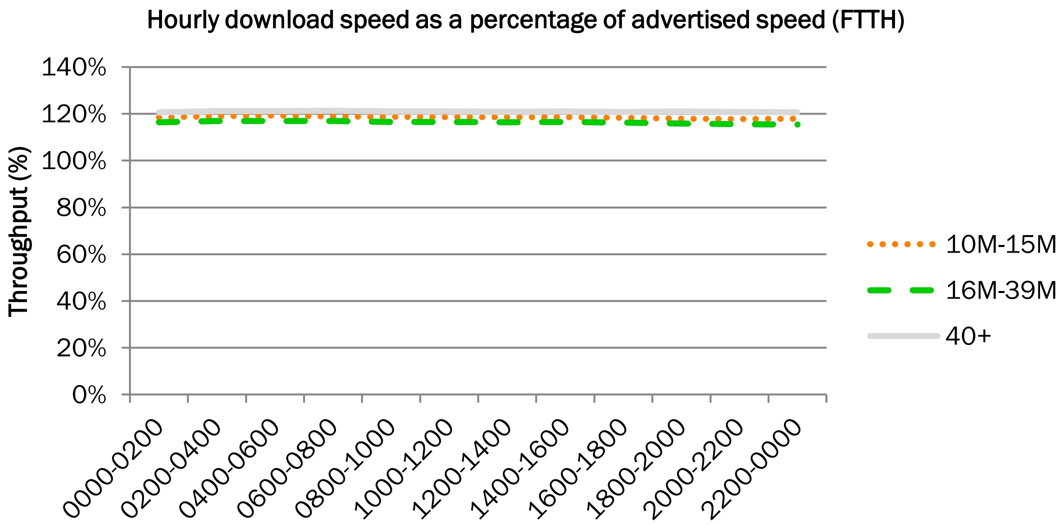

Figure 6: Hourly FTTH download speed as a percentage of advertised speed

B.2 Upload Throughput

Upload throughput is the measure of how fast data can be transmitted from the home to the Internet. Higher speeds allow for pictures, music and documents to be uploaded and shared more quickly.

To characterize the user's maximum access link capacity, measurements were conducted between a nearby test server and panellists' homes.

Historically, the amount of data that users download has vastly outweighed the amount of data that users upload. This has led technologies to be engineered to be asymmetric; i.e. they offer faster download rates than upload rates. As can be seen by comparing the download speeds to upload speeds, this level of asymmetry is falling for new services, such as those delivered using FTTH.

As with the download throughput, results in this section are presented as a percentage of the ISP's advertised product in addition to results for the absolute level of upload speed. This enables comparability between products of vastly different speeds.

Figure 7 depicts upload speed as a percentage of advertised speed for each technology and speed bucket. As with the download measurements, the vast majority of services and technologies met or exceeded the advertised rate. However, the 5-9Mbps DSL, delivered 85% of advertised upload speed.

The 10-15M DSL and 10-15M FTTH results also stand out for the fact that they exceed the advertised rate at 171% and 149% respectively. This is caused by some ISPs choosing to overprovision their upstream speeds, far beyond the advertised upstream rate.

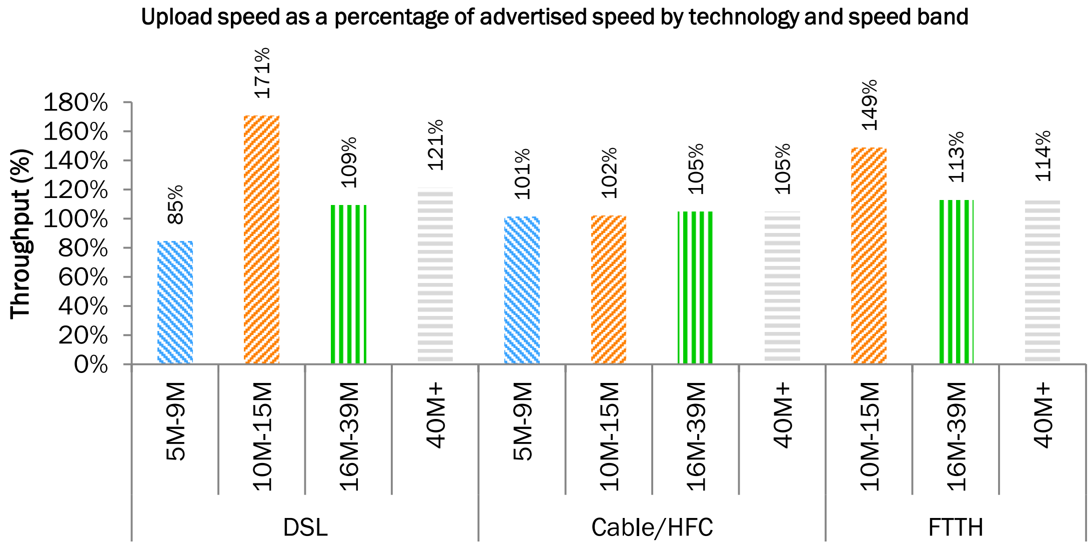

Figure 7: Upload speed as a percentage advertised speed by technology and speed bucket

Advertised upload speeds are far lower than advertised download speeds, reflecting the asymmetric nature of most broadband services.

Figure 8 demonstrates that upload speeds were generally consistent across the regions, with all but one area receiving in excess of the advertised rate. In the West & North, 10-15Mbps services achieved 156% of advertised (again due to certain ISPs overprovisioning upstream rates significantly). However, 5-9Mbps services in the same region only delivered 83% of advertised upstream speed.

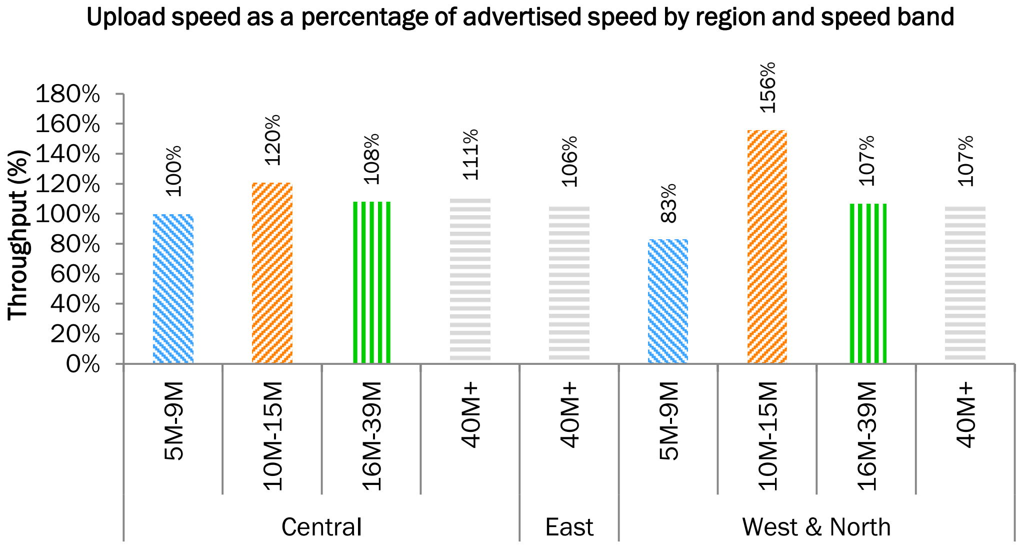

Figure 8: Upload speed as a percentage of advertised speed by region and speed bucket

Upload speeds were very consistent all hours of the day, as shown in the figures 9, 10 and 11 below. Only the 10-15Mbps bucket in DSL technology showed an insignificant decline in throughput during peak hours, with speed falling by three percentage points during peak hours. In all other cases, upload throughput varied by less than 1% throughout the day.

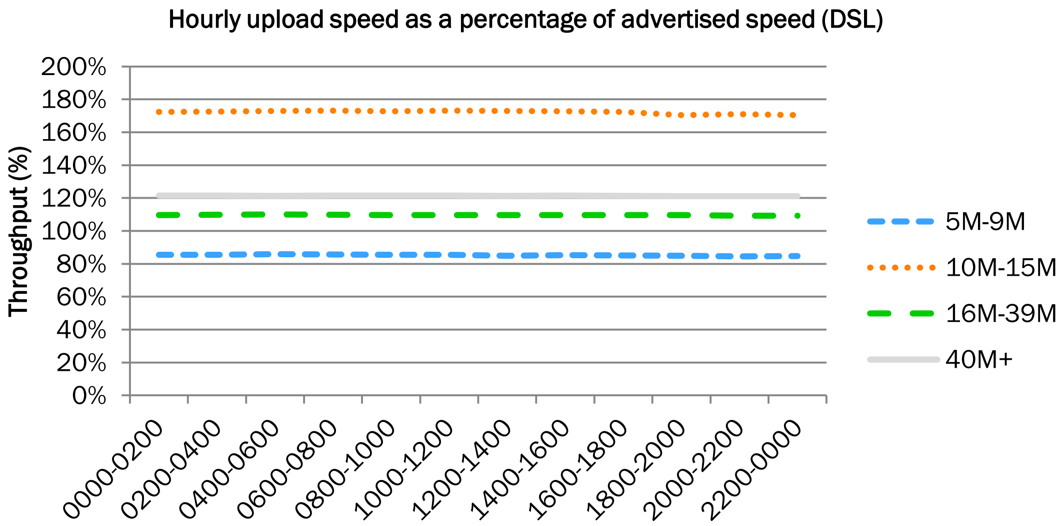

Figure 9: Hourly DSL upload speed as a percentage of advertised speed

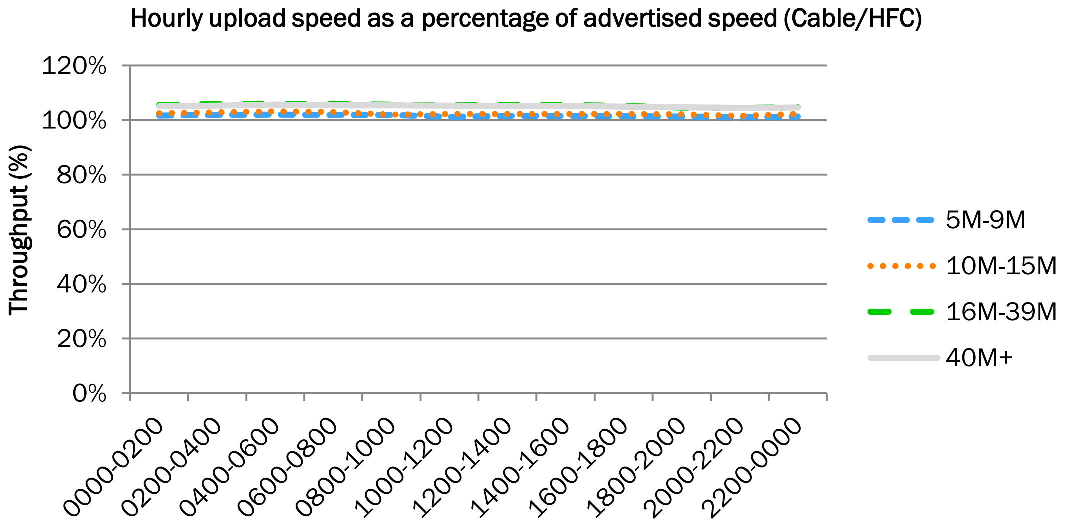

Figure 10: Hourly Cable/HFC upload speed as a percentage of advertised speed

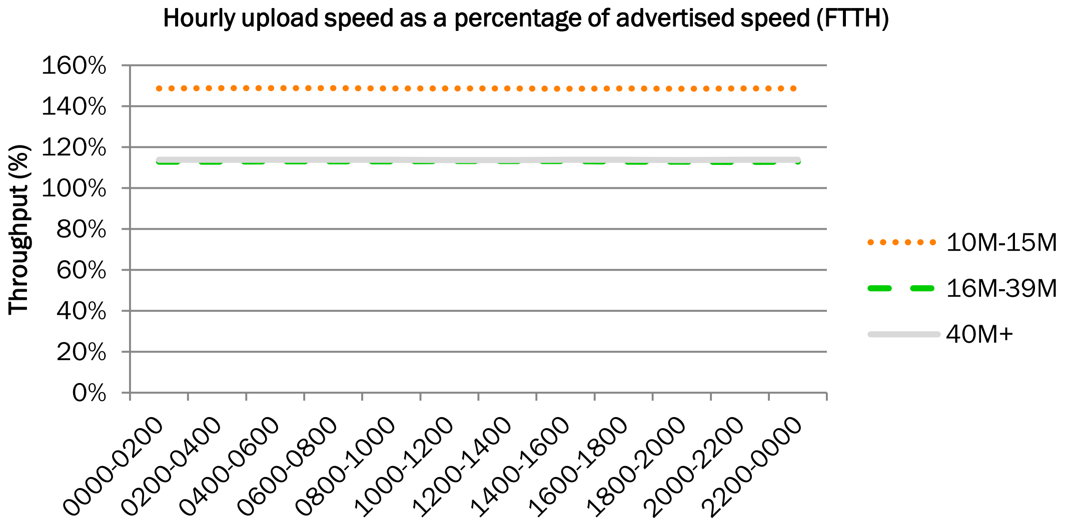

Figure 11: Hourly FTTH upload speed as a percentage of advertised speed

B.3 Latency

Latency is a measure of how long it takes a data packet to travel between point A and point B. It is a significant factor in Internet performance, as latency is a fundamental property of the infrastructure upon which everything else must build. If you have a high latency link, then it does not matter how much capacity your broadband connection has; you will be limited by latency.

The results presented in this section show 'round-trip' latency (i.e. how long it takes for a data packet to travel between point A and point B and then back to point A). While round-trip latency is the most common latency measurement taken (for example, the 'ping' utility captures round-trip latency), the 'round-trip' qualifier is very often omitted. For the remainder of this document 'latency' should be taken to mean 'round-trip latency'. Please note that the proximity of the servers to the end user will affect latency results as these are a measure of distance. The servers that were used for the purpose of testing are located in Halifax, Montreal, Toronto, Vancouver, and Winnipeg.

Latency is almost always expressed in milliseconds. Lower results are better. Latency itself has a lower bound governed by the speed of light in a vacuum, and often there are technological limitations which raise that lower bound. For example, DSL services typically have higher latencies than FTTH services.

Whilst latency is unrelated to capacity (the amount of data that can be transmitted over a broadband connection), increases in latency can have a detrimental effect on achievable speed. Moreover, an increase in latency during peak hours is an early indicator of congestion somewhere on the network path, as routers are taking longer to receive data packets and pass them on. It is worth noting that even the highest latencies exhibited here would more than be adequate for any common Internet application at present. For the majority of use cases, the approximately 20ms latency difference between the best and worst service, as shown in Figure 15 below, would be indistinguishable.

Figure 12 below shows peak period latency by technology and speed bucket. FTTH services delivered the lowest latencies, with this varying between 11.5ms and 11.7ms. Cable/HFC services demonstrated a similar level of consistency, with all speed buckets delivering between 17.2ms and 21.7ms. DSL services showed the widest variance. DSL services in the 5-90Mbps bucket yielded a round-trip latency of 30.0ms. These latency results improved with the faster DSL services, with the fastest services in the 40Mbps+ bucket yielding 16.7ms. This largely reflects the underlying characteristics of the access technology. The fastest DSL services will be VDSL based, meaning that the copper loop length will be very short and interleaving (a mechanism that reduces errors, but at the expense of latency) will typically be disabled. These two factors have the effect of significantly reducing round-trip latency.

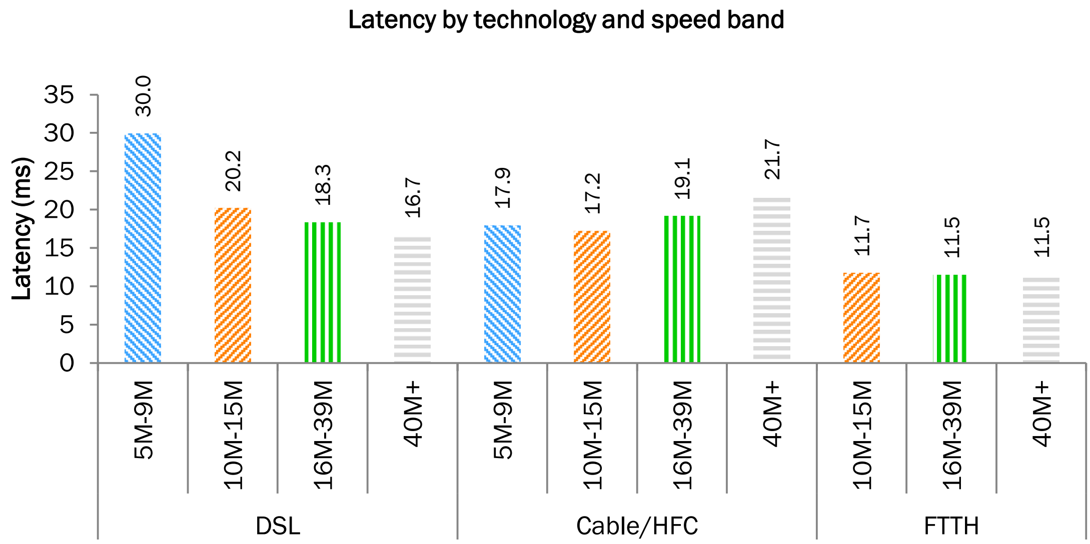

Figure 12: Latency by technology and speed bucket

Figure 13 shows that there is a wide variation in latency between regions. Results from the East showed the highest latencies, at 39.6ms, and Central with the lowest, ranging between 12.2ms and 25.6ms. As latency is largely dependent on the length of the physical network path being measured, plus the access technology in use, it is to be expected that Whiteboxes located in more rural areas (with both longer paths and a higher concentration of DSL) will have higher latencies.

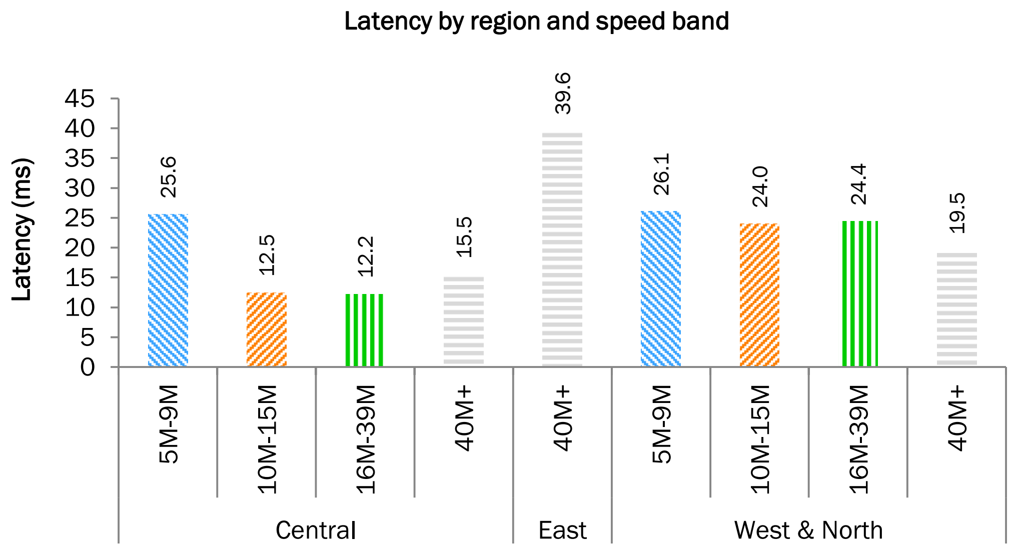

Figure 13: Latency by region and speed bucket

The hourly charts seen in Figures 14, 15 and 16 show that latency across all speed buckets and technologies was generally stable throughput the day. FTTH services typically show less than 0.2ms variation, demonstrating a very consistent performance.

All Cable/HFC services and the 5-9Mbps DSL services showed higher latency during peak hours. However, these increases in latency during peak hours were typically only in the range of 1-3ms, and would be imperceptible to the vast majority of applications.

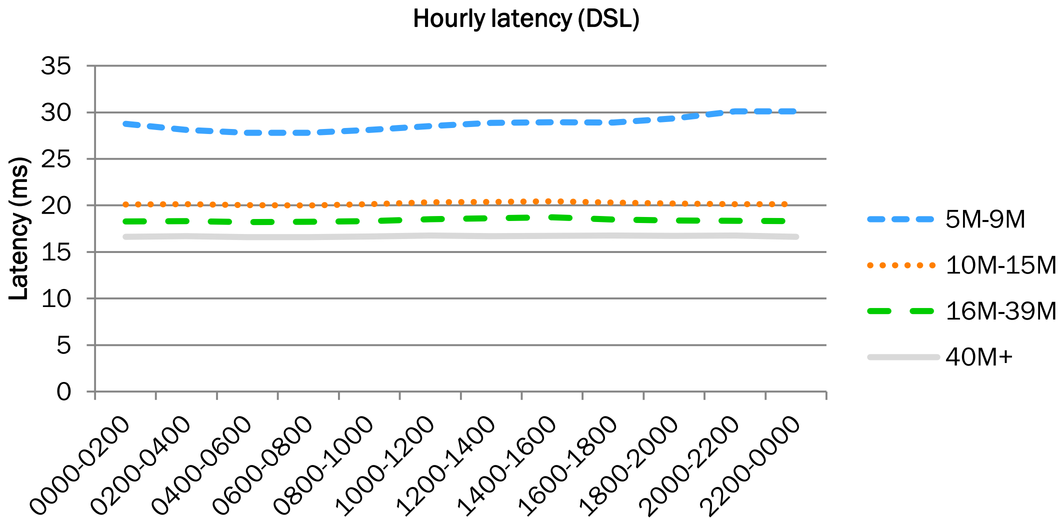

Figure 14: Hourly DSL latency

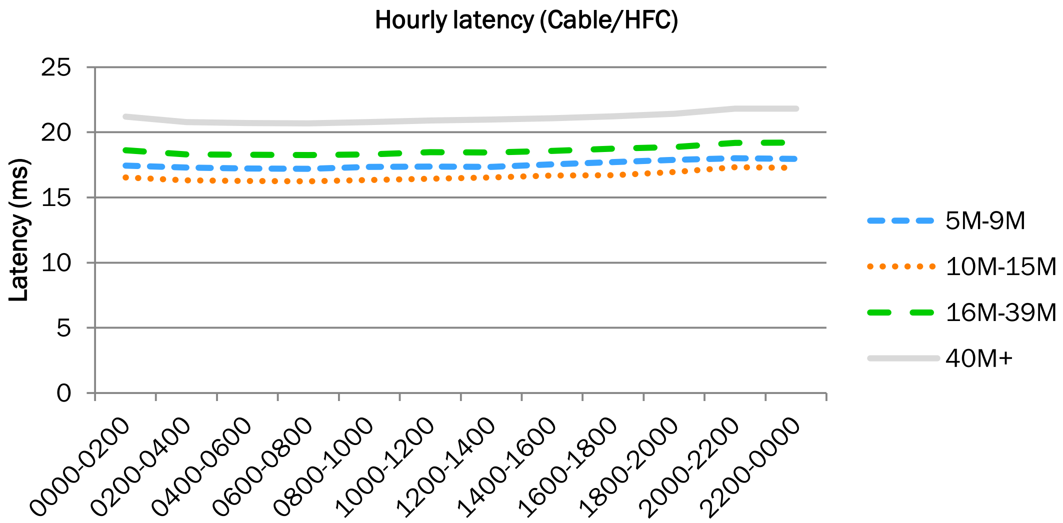

Figure 15: Hourly Cable / HFC latency

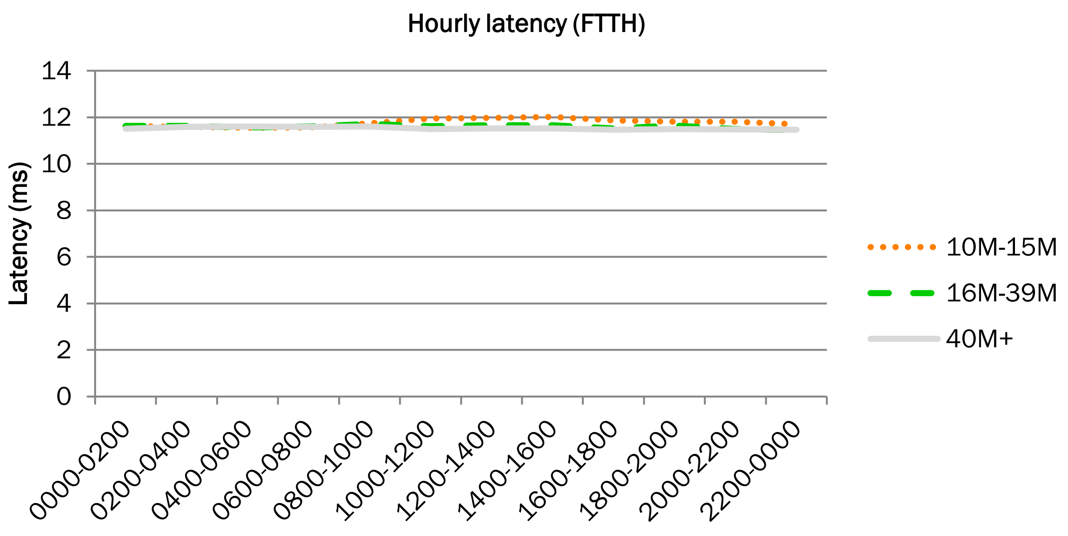

Figure 16: Hourly FTTH latency

B.4

Packet Loss

The rate of packet loss describes how likely it is that a data packet sent from point A will not reach point B. Packet loss is closely related to the latency and is also a fundamental metric in determining how applications perform on a broadband connection. A high rate of packet loss will prevent many applications from working to a satisfactory level. A small increase in packet loss during peak hours is to be expected, as networks are busier and congestion at even one point in a network path may lead to a packet being dropped. Packet loss is generally measured and expressed as a percentage of the overall data packets sent.

Figure 17 below shows that packet loss was extremely low across almost all speed buckets and technologies. Only one speed bucket was shown to exceed packet loss of 0.13%. This was the 5-9Mbps bucket in DSL, which reached 0.46% packet loss. FTTH services delivered the lowest levels of packet loss, with a range of just 0.03% - 0.05%.

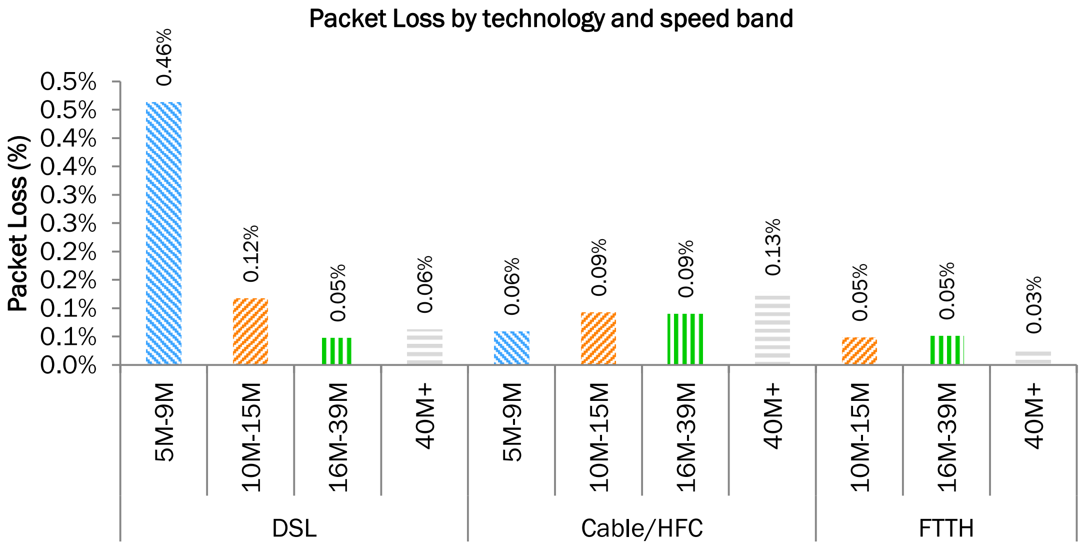

Figure 17: Packet loss by technology and speed bucket

Figure 18 below shows packet loss by region. Packet loss proved very low across all regions, with all but one region delivering 0.12% or lower. Packet loss of the 5-9Mbps bucket in West & North proved high at 0.48%, but only during peak hours. During off-peak hours packet loss is far lower.

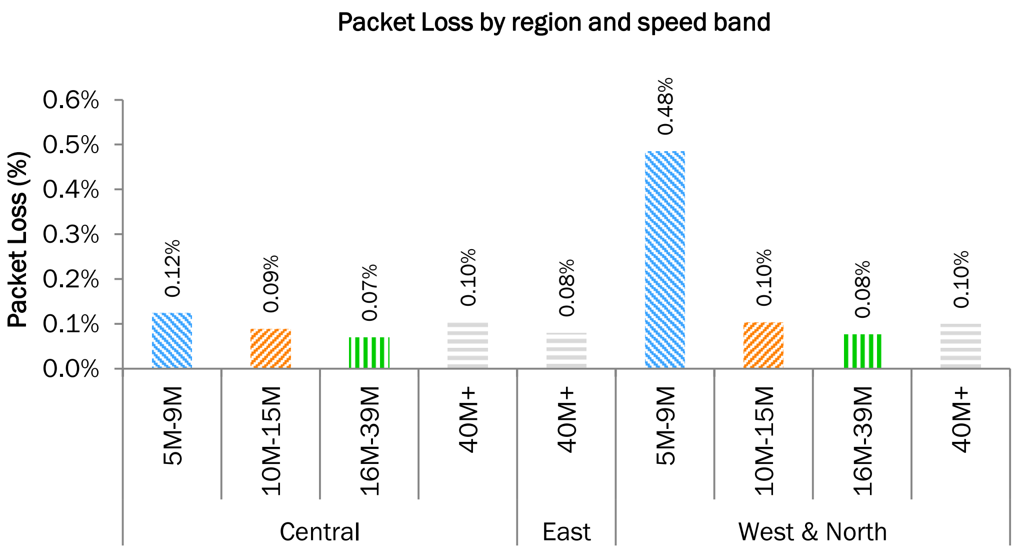

Figure 18: Packet loss by region and speed bucket

Figures 19, 20 and 21 show the variation in packet loss by hour of day across the different access technologies. DSL showed the highest level of variation in packet loss, particularly for the 5-9Mbps bucket. Packet loss increased from 0.12% off-peak to 0.49% for this service. Packet loss on FTTH varied by less than 0.05% during the course of the day, and showed no noticeable rise during peak hours. Cable/HFC services showed a more visible rise in packet loss during peak hours, but still remained less than 0.06%.

Packet loss at these levels, across virtually all technologies, speed buckets and regions would be imperceptible to almost all modern Internet applications.

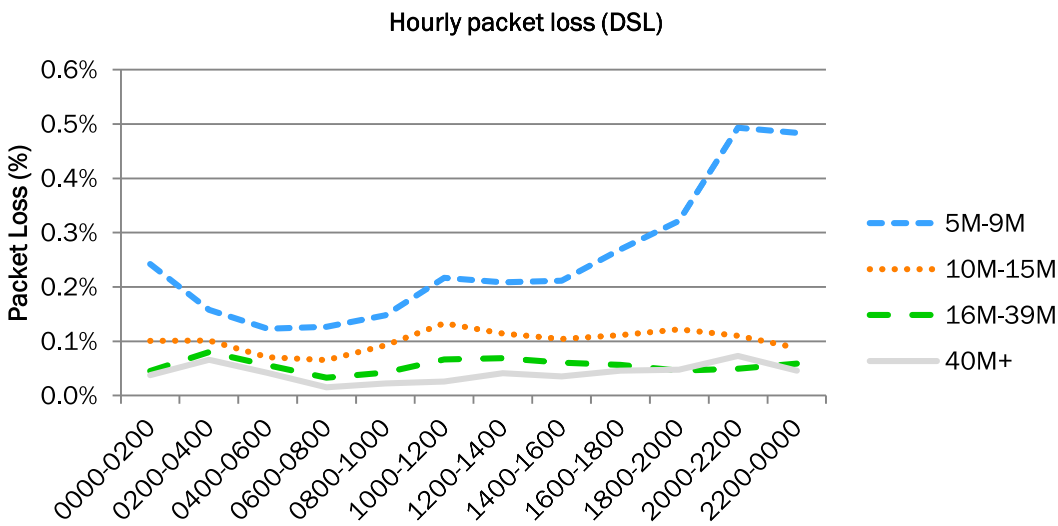

Figure 19: Hourly DSL packet loss

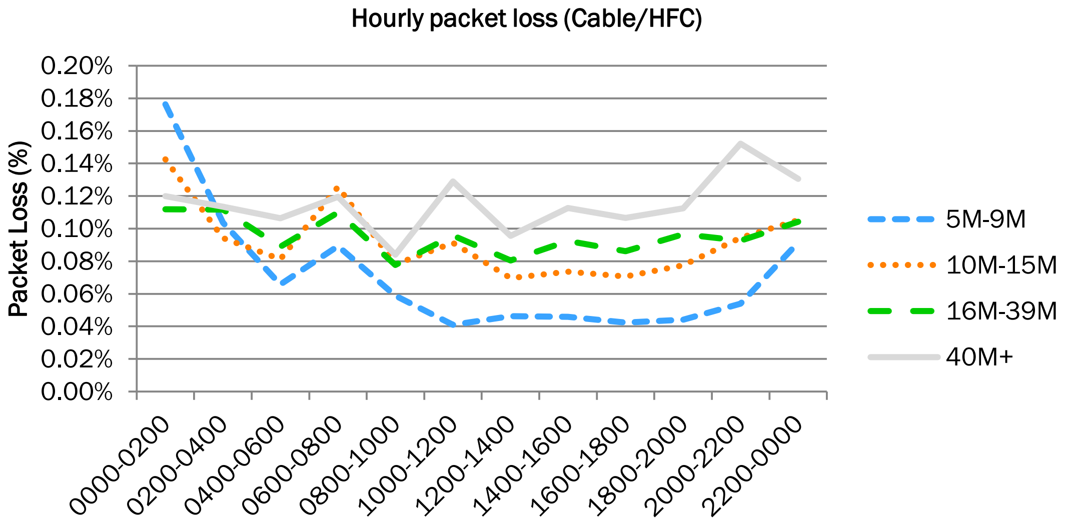

Figure 20: Hourly Cable/HFC packet loss

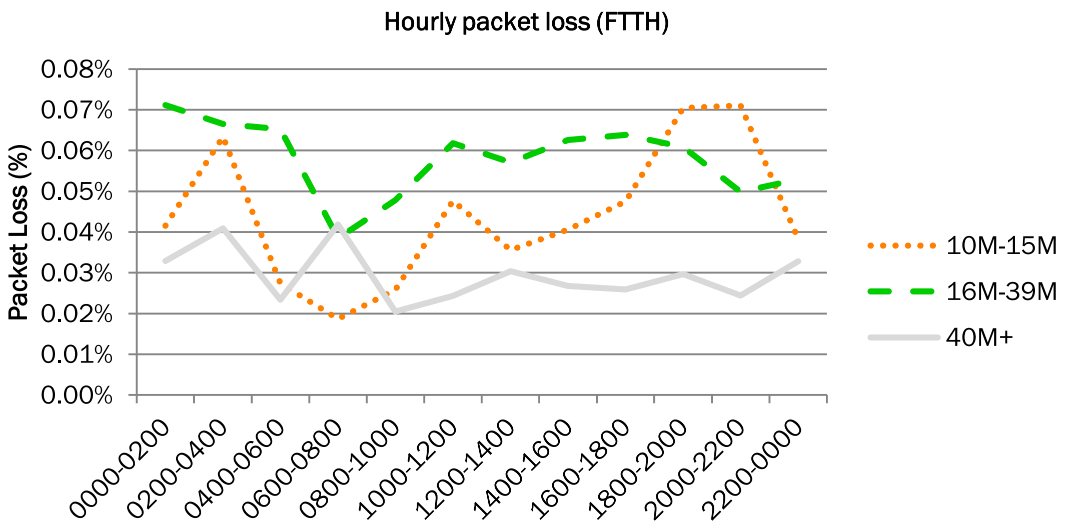

Figure 21: Hourly FTTH packet loss

B.5 Web page loading time

The web page loading time test captures how long it takes for all of the elements of a web page to be received by an end user. Unlike other measurements, this test is conducted against real websites on the Internet, rather than dedicated test servers. The websites tested against were:

- facebook.com/policies
- ca.yahoo.com
- live.ca
- canada.ca
- cbc.ca
- google.ca
- ebay.com
- theweathernetwork.com
- ici.radio-canada.ca
- meteomedia.com

Web page loading time is heavily influenced by both download speed and latency. Studies in other markets have shown that throughput dominates web browsing performance up to approximately 10Mbps, after which latency becomes the dominant factor.

Cable/HFC, FTTH and other DSL services all showed extremely consistent web page loading times between peak and off-peak hours.

Figure 22 depicts average web page loading time by technology and speed bucket. As expected, web page loading time improved as download speed increased, although this improvement tails off rapidly above 10Mbps. DSL services in the 5-9Mbps speed bucket delivered web pages in the longest time, at 2.2 seconds. This improved to 1.1 seconds in the 10-15Mbps bucket. Once testing moved to the 40Mbps+ bucket, this improved further still to 0.7 seconds. Similar behaviour is observed on Cable/HFC services as well, although the 5-9Mbps Cable/HFC services delivered faster performance than DSL at 1.3 seconds.

All 40Mbps+ services load web pages in 0.9 seconds or faster.

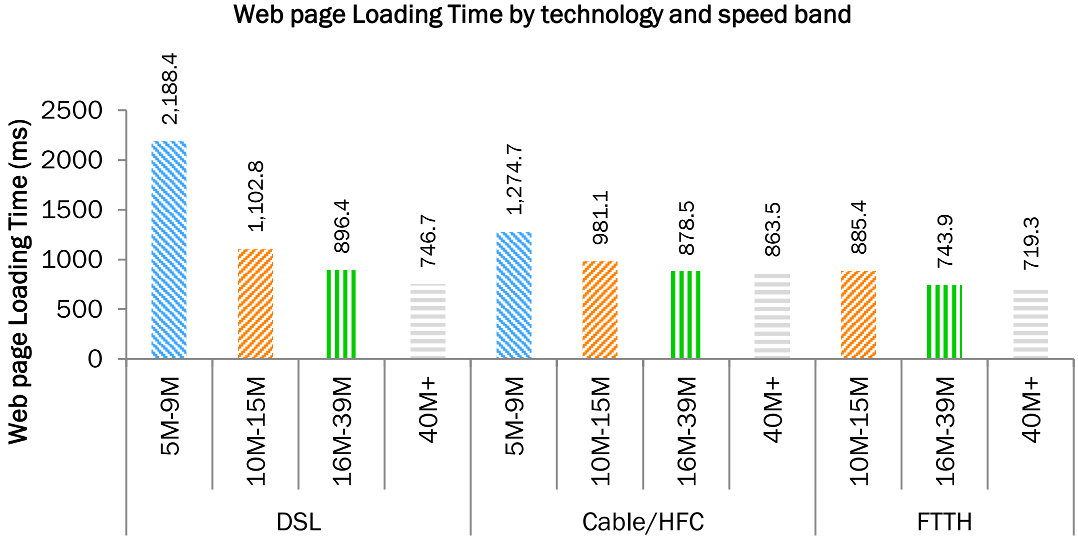

Figure 22: Web page loading time for websites studied across all of the ISPs in question

Figure 23 shows web page loading time by region and speed bucket. In all regions, web page loading times of the 5-9Mbps bucket was slower than the loading times of all other speed buckets.

The faster services in the Central region all delivered the best web page loading times. This is likely due to the prevalence of FTTH deployments in these areas, yielding both high, consistent speeds and low latencies. These conditions deliver the optimum web page loading times.

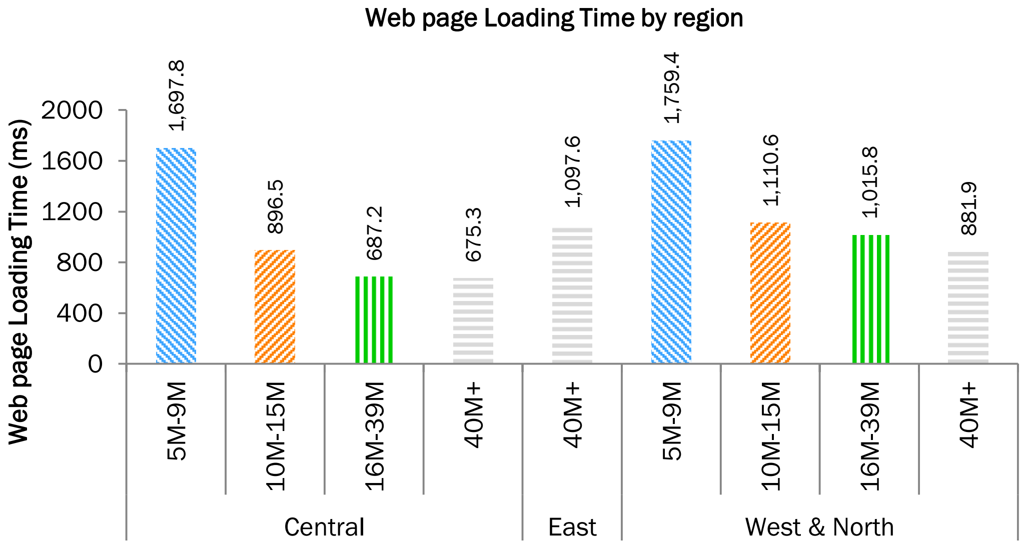

Figure 23: Web page loading time during peak hours

Figures 24, 25 and 26 show web page loading times by hour of day.

All technologies and speed buckets showed an interesting rise in page loading time outside of the night time hours (after 8am). This increase is typically between 50-100ms. Given that this affects all access technologies and speed buckets, and is not observed in the throughput or latency charts, it suggests that the cause of this rise might be the websites themselves. Load on the websites will likely increase during daylight hours, which is exactly the period where the loading time increases are observed here.

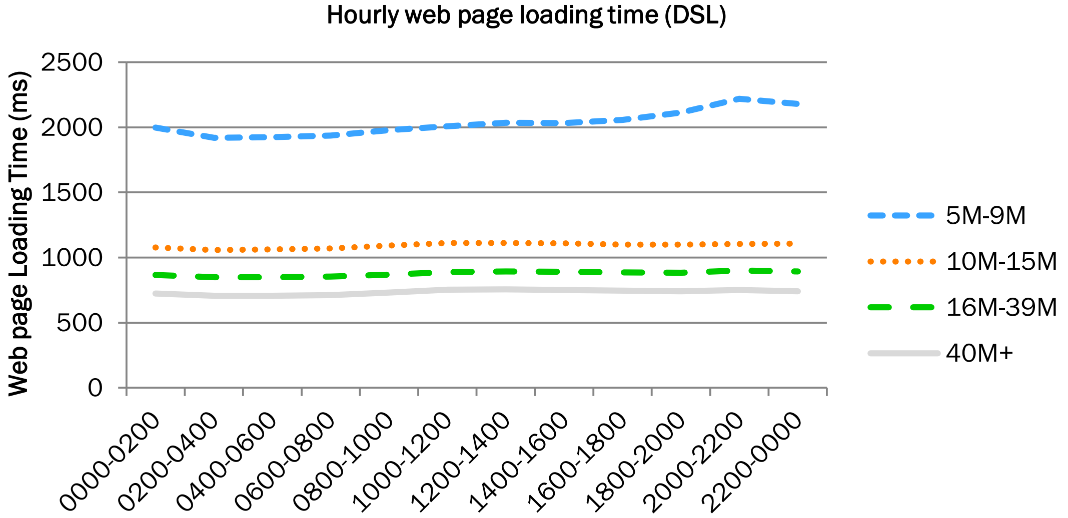

Figure 24: Hourly DSL web page loading time

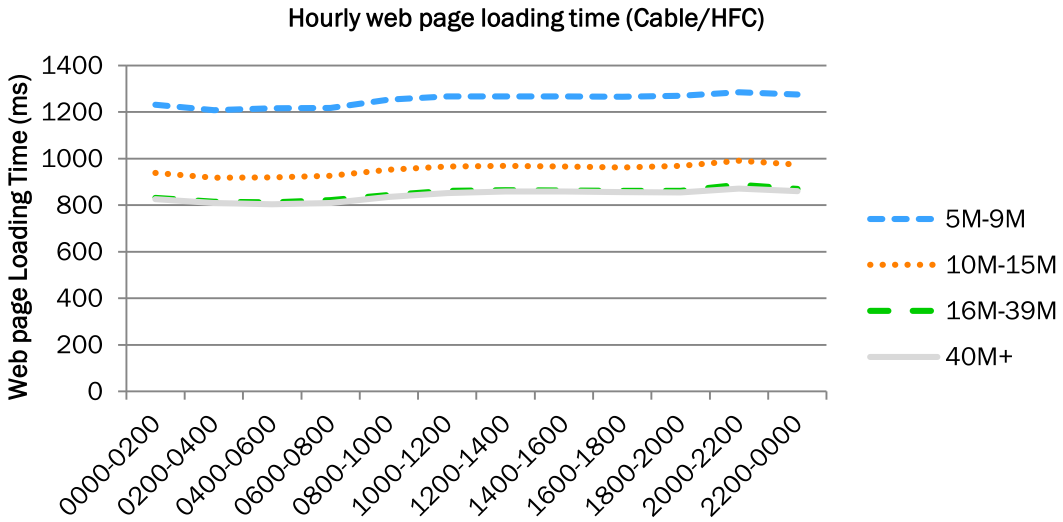

Figure 25: Hourly Cable/HFC web page loading time

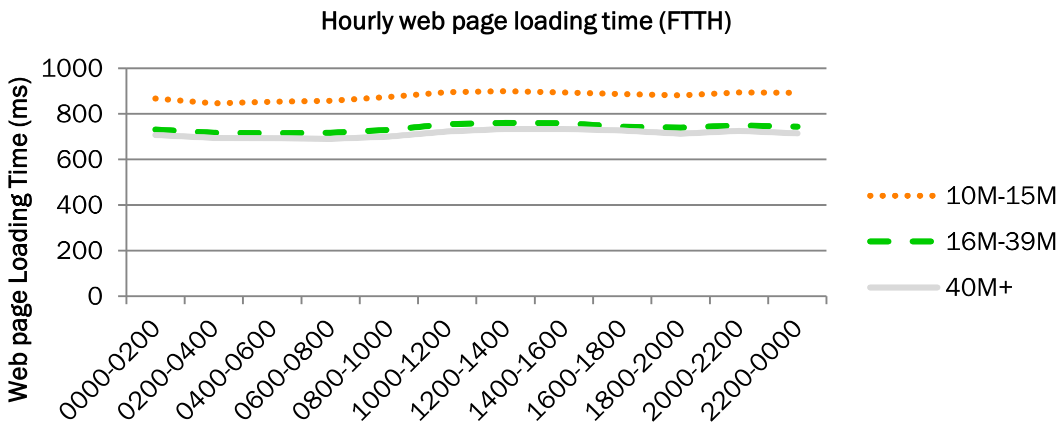

Figure 26: Hourly FTTH web page loading time

B.6 Conclusion

This report represents the first step in the CRTC's efforts to measure and report on the performance of Canadian broadband networks.

Almost all broadband services met or exceeded their advertised speeds, regardless of the access technology in use. FTTH and Cable/HFC services in particular delivered reliable download and upload speeds at all hours of the day. FTTH services achieved 119% of advertised download speeds on average, and Cable/HFC services achieved 103%. DSL showed more variance. DSL's overall average was 97% of advertised speeds with the majority of DSL services exceeding advertised rates.

Throughput (as a percentage of advertised) does not vary significantly by regions, with all regions averaging between 109% and 117% of advertised download speed. However, some services such as those in the 5-9M bucket in West & North fall slightly below advertised speeds during peak hours.

In the latency, packet loss and web browsing metrics, Cable/HFC and FTTH services delivered the most consistent results. DSL services exhibited more variance. Even the highest latencies exhibited during testing would be more than adequate for any common Internet application at present and latency and web-browsing in Canada compares favourably to those measured in other jurisdictions including the United States.

Later in 2016 the CRTC will publish a second report that will expand upon this study to focus on individual ISPs. This report represents the first in an ongoing effort by the CRTC to better understand the true state and performance of broadband Internet access services available to Canadians.