Telecommunications Carriers' Forum

Report on Broadband Performance Monitoring

06 April 2009

Contents Page

1. EXECUTIVE SUMMARY

1.1. The Telecommunications Carriers' Forum (TCF) outlined to the Commerce Commission (the Commission) – in September and December 2008 - some of the issues it had identified with the current broadband performance monitoring. In response, the Commission advised they welcomed feedback and would work with the TCF to ensure that broadband monitoring provides information useful to both the industry and to end users.

1.2. The TCF established a Broadband Performance Monitoring Working Group under the umbrella of its Information Reporting group to identify the current issues with broadband monitoring and to recommend solutions to ensure the information collected was accurate, robust and useful to the industry and end users alike.

1.3. This draft report sets out the current issues with broadband monitoring and recommends options or solutions in the following areas:

1.3.1 identifying issues with current testing and providing options for testing which are more appropriate for providing accurate and useful measurements of broadband performance;

1.3.2 defining issues in gathering sufficient data from a random population sample that would result in robust conclusions reflecting the experience of the wider population;

1.3.3 issues and risks arising from current methodology for collection of broadband performance data, and options for improving methodology;

1.3.4 the challenge of providing timely, flexible and meaningful presentation and disclosure of broadband performance information that meets a variety of consumer needs, and options that take in to consideration approaches currently employed by overseas agencies; and

1.3.5 the importance of educating customers by raising consumer awareness of drivers and influences on broadband performance, referencing suitable current practices employed by overseas regulators.

1.4. The desired outcome from this report is to develop a more robust, representative solution for collecting and presenting the data which the TCF believes will more accurately reflect the Commission's objectives to promote consumer education and understanding of broadband.

1.5. The recommendations at this early stage of the TCF's analysis are:

1.5.1 to test the following influences on broadband performance:

a) throughput;
b) packet delay (latency);
c) packet delay variation (jitter);
d) packet loss;
e) DNS; and
f) availability.

1.5.2 to ensure a sample size that enables statistically significant conclusions.

1.5.3 to use the dedicated end user hardware option (which sits between the end user's modem and computer) to conduct the testing.

1.5.4 to adopt a web based interface to present collected data.

1.5.5 to educate end users on what they should expect from broadband performance.

1.6. The TCF would like the opportunity to discuss the contents of this report and share initial thoughts with Commerce Commission officials. The TCF anticipates that following this meeting, consultation could take place with the Commission, Epitiro, IDC and the TCF, if the Commission sees this as appropriate.

1.7. Until consultation can take place, and solutions to the issues outlined in this report are agreed amongst stakeholders, the TCF requests that the Commission cease the publication of any further Broadband Performance Monitoring Reports.

2. OVERVIEW

2.1. INTRODUCTION

2.1.1 In March 2008, the Commission initiated a monitoring project designed to measure and compare broadband performance in New Zealand (New Zealand Broadband Index). This monitoring was in accordance with Section 9A of the Telecommunications Act, the requisite power for sector monitoring and information dissemination and the Commission's guidelines on the use of Section 9A.

2.1.2 The Commission outsourced the broadband performance monitoring to a private organisation, Epitiro, which benchmarks and ranks Internet Service Providers (ISPs) through the use of eight tests. A market intelligence and advisory service, IDC,

takes the data provided by Epitiro to produce a quarterly New Zealand Broadband Performance report on behalf of the Commission. To date, the Commission has published three reports.

2.1.3 The structure and format of the Broadband Quality report has changed with each quarterly publication as the Commission and Epitiro/IDC have responded to feedback from a wide range of industry participants. While these changes have been welcomed, there are a number of key underlying problems and errors identified by industry participants, including TCF members, which compromise the usefulness of the report. Consequently there is still uncertainty about the veracity of the conclusions drawn from the data collected. Given these ongoing problems the TCF recognised there was a need to assist the Commission with its broadband performance monitoring.

2.1.4 In November 2008 the TCF Board agreed to establish the Broadband Performance Monitoring working group to identify the current broadband monitoring issues and, if required, develop options to resolve those problems. The TCF have now worked through all the identified issues and have developed a set of proposed options to resolve these.

2.2. BACKGROUND

2.2.1 To undertake broadband performance monitoring Epitiro currently uses a service called ISP-I. This service runs eight tests every fifteen minutes across a twenty four hour period. These eight tests are characterised as Key Performance Variables (KPVs). These KPVs are as follows:

a) Synchronisation speed;
b) Time to connect (discontinued at close of Q3 2008);
c) Cached HTTP;
d) Non-cached HTTP;
e) Ping performance;
f) DNS;
g) Email round trip; and
h) Packet loss performance.

2.2.2 These KPV tests are generally focused at the speed of the service being tested. These tests do not capture quality elements or other support features commonly provided with a broadband service.

2.2.3 The Commission selected thirteen ISPs for testing at between one and eleven static sites across Auckland, Hamilton, Wellington, Christchurch and Dunedin.

2.2.4 Not every ISP's premium broadband plans are tested at all eleven sites. Some are tested at only a single site while others are tested at all available sites. This sample design limits the conclusions which can be drawn from the results.

2.2.5 The results from the ISP's tests, obtained over each quarter, are averaged and aggregated by KPV. This enables a ranking and comparison of each ISP by KPV. Overall results for each KPV are then aggregated together to provide a view of all ISPs performance in a particular testing site. This process of result aggregation is undertaken in order to compare the results of ISPs and, in effect, rank each ISP.

2.3. BROADBAND PERFORMANCE MONITORING OBJECTIVES

2.3.1 The TCF recommends that the Commerce Commission's broadband performance monitoring objectives should be to:

"Ensure that customers are educated about broadband services and are able to access timely, accurate and meaningful information in order to make informed decisions about which broadband services best meet their needs.

Undertake domestic and international analysis of broadband service performance to provide reasonable expectations about the broadband services which the New Zealand public demands."

2.3.2 The TCF believes that these objectives are consistent with the purpose of the Telecommunications Act and the Sector Monitoring Guidelines which states:

"*The Commission also undertakes monitoring of broadband service levels to promote consumer education and understanding of telecommunication markets. Effective competition will be encouraged if consumers have access to a wide range of information about products in telecommunications markets. The Commission may also monitor service levels in respect of other telecommunication products from time to time."*

3. TCF PERSPECTIVE

3.1. Since the first release of the New Zealand Broadband Performance Monitoring Report the TCF has had growing concerns about the Epitiro methodology and the robustness of its approach. The particular concerns which have been identified to date are:

3.1.1 appropriateness of the independent monitor to supply commercial services to organisations it is monitoring (for example restrictions on auditors);

3.1.2 appropriateness of the independent monitor to make public comment on commercial issues which go beyond its brief;

3.1.3 robustness and statistical validity of sample size and selection;

3.1.4 value of averaging results given the nature of consumption that is strongly influenced by factors such as temporal variability;

3.1.5 lack of transparency regarding the Epitiro methodology (subjective and unknown weightings, and opaque formulas);

3.1.6 bias of Key Performance Variables (KPVs) towards measurement of application and protocol performance over quality and reliability measurement (e.g. availability metrics) and quality of support services;

3.1.7 the adoption of proprietary non-standards based measurement tests such as email RTT as representative of broadband performance or quality;

3.1.8 misrepresentation of source data for some KPVs (e.g. packet loss is a measurement of TCP retransmissions); and

3.1.9 factually inaccurate analysis on causal factors for regional performance differentials (e.g. the citing of the Hamilton Metropolitan Open Fibre Access Network as a reason for best national broadband performance).

3.2. This report outlines the options that the TCF recommends to determine:

3.2.1 what KPVs to test;

3.2.2 how the testing sample is chosen;

3.2.3 the methodology for determining how the tests be conducted;

3.2.4 presentation of the data; and

3.2.5 education of end-users.

4. TESTING

4.1. WHAT TO TEST (KPVs)

4.1.1 Broadband performance is influenced by a range of factors: some factors are influenced by the ISP; others are influenced by the ISP's wholesale infrastructure provider and the end user's home environment. There is also an inherent subjectivity in what each end user will view as key to a good service, from the range of factors contributing to a good broadband service. Investigation of broadband analysis carried out in the UK, Australia and Singapore, by both regulators and industry participants, illustrates that some end users value the speed of their broadband service, others value service reliability, and others are concerned with affordability or value for money.

4.1.2 Consequently, selecting the most relevant factors to test in order to identify broadband performance is inherently a subjective process. As such, it is important in the TCF's view to establish a set of transparent principles in order to guide which performance variables, when tested, would produce the most valuable information to end users. The principles for performance measurement must:

a) **Be objective and quantifiable** - that is statistically significant and independent of subjective judgements;

b) **Be relevant** – utilise recognised industry standards to maintain relevance to broadband end users;

c) **Identify the significance of each measure to end users' overall broadband experience;**

d) **Test broadband infrastructure** – not end user premises equipment, applications or services; and

e) Be provided within a **relevant geographic context**.

4.1.3 **WHAT TO TEST – RECOMMENDED OPTION**

The TCF recommends that the following variables should be the focus of broadband performance testing:

a) ***Throughput****;*

- downstream and upstream in megabits/sec.[1]
- domestic and international in megabits/sec.
- throughput tests using a range of packet sizes.

[1] There is currently a poor understanding of the distinction between the terminology of megabits per second and megabytes per second. The TCF is of the opinion that this should be clarified to a best practice of megabits per second. Bytes are a unit of storage while bits are a transmission unit.

b) ***Packet delay (latency)*;**

- two-way.
- delay in milliseconds.

c) ***Packet delay variations (jitter)*;**

- two-way.
- delay in milliseconds.

d) ***Packet loss*;**

- percentage of packet loss.
- provides an indication of the quality and reliability of a broadband service. Packet loss is typical in packet-based networks.

e) ***DNS*;**

- query time in milliseconds.

f) ***Availability*.**

- percentage of failed tests or non-returned results.

4.1.4 Each of the six recommended variables above align with the principles outlined in section 4.1.2. Each variable can be objectively tested in a manner consistent with international practice. The variables are relevant to the average broadband user and have a major influence on the performance of a good broadband connection.

4.1.5 Other factors such as customer support and value for money, despite being important to consumers, are inherently subjective. As such, the TCF does not recommend testing these factors. However, as part of broadband performance reporting, these should be noted as key factors considered by end users when deciding which ISP to purchase a broadband service from.

4.1.6 The alignment of the principles with the performance variables were analysed as follows:

		Principles			
		Objective & quantifiable	Relevance to end user	Significance of measure	Broadband infrastructure
Performance	Throughput	✓	✓	✓	✓
	Packet delay (latency)	✓	✓	✓	✓
	Packet delay variation (jitter)	✓	✓	✓	✓
	Packet loss	✓	✓	✓	✓
	DNS	✓	✓	✓	✓
	Availability	✓	✓	✓	✓
	Sync Speed	✓	✗	✗	✓
	Time to connect	✓	✗	✗	✓
	Cached HTTP	✗	✓	✗	✗
	Non-cached HTTP	✗	✓	✗	✗
	Email round-trip	✗	✗	✗	✗
	Customer support	✗	✓	✓	✗
	Value for money	✗	✓	✓	✗

4.1.7 The Epitiro methodology has been adopted by the Commerce Commission for undertaking broadband performance monitoring. The Epitiro test suite ISP-I focuses on eight key performance variables. Of the eight Epitiro KPVs, five do not align adequately with the principles set out in 4.1.2. Therefore, the TCF recommends not utilising those performance measures. They are:

a) **Synchronisation speed**: testing the line speed after connection to the ISP has been initiated has little overall impact on the performance of a broadband connection. This is because synchronisation speed testing only provides an indication of an upper limit on line speed or on the headline rate, where as the sustained rates of data transmission may actually be lower. It is the sustained rate which is most important to an end user. This is also consistent with the Commission's directive that headline speeds should not be used to sell or promote broadband.

b) **Time to connect**: the average time taken for an ISP to recognise a broadband modem and connect it to the network has minimal relevance to, or impact on, overall broadband performance. This is because DSL broadband connections are generally "Always On" connections. Primarily, a DSL broadband connection will be initiated when the end user turns on the modem.

c) **Cached and non-cached HTTP:** this test is a proxy for throughput. It is useful for identifying whether an ISP uses caching as a matter of managing the issues around data transmission from remote locations. However, it does not give an end user a clear indication about the ISP's provisioning of domestic and international transmission. Transmission provisioning will significantly influence throughput and, therefore, broadband performance for end users. At present the configuration of Epitiro's test is unclear. If it is standards- based, the test would require the ISP to honour the flags set in the content and in the request from a client. If the test is proprietary to Epitiro it is unclear as to how the ISP might handle the request. Consideration should be given to the measurement of throughput to both domestically and internationally dedicated test ends.

d) **Email round-trip:** the testing of email round-trip is excluded because the monitoring of broadband performance focuses on the broadband pipe itself, not the services or applications which may be provided by an ISP. Measurement of this nature is highly proprietary and is in no way indicative of the quality of an ISP's mail service.

4.1.8 The range of KPVs selected by Epitiro differ from variables selected by international regulators and broadband performance monitoring agencies. For example, Singapore's iDA tests and reports only on throughput and latency. So far, the UK's Ofcom has only focused on throughput.

4.2. HOW TO TEST

4.2.1 Undertaking quantitative analysis requires a researcher to give careful consideration to the design of the study. There are a range of aspects to good study design. Firstly, the problem needs to be carefully articulated and understood. Secondly, an appropriate population needs to be identified for testing. Thirdly, the study must then be randomised to ensure that no pre-determined bias is captured in the sample population. Reliable equipment must be used in order to obtain measurable and verifiable observations. Most importantly, but often compromised in research design, is the selection of an adequately-sized sample population. An adequate sample size ensures the statistical significance of conclusions drawn from collected observations is protected.

4.2.2 Failure to address these design issues compromises or limits the ability to make accurate or statistically significant conclusions about the collected observations or data. The study could also be open to considerable challenge regarding the robustness of its conclusions.

4.2.3 **Sample Selection**

a) A key question that researchers have to ask themselves is whether the random sample can explain or predict what the full population actually experiences. This is because the random sample is meant to be representative of the full population. However, small samples can be subject to bias and errors. Consequently, it becomes unclear whether the observed results of a small sample can be used to explain or predict a population. Also, the importance of random selection is often overlooked but is an essential component to research design. Random selection provides each member of the population with an equal opportunity to be selected - thereby reducing the risk of bias in the sample.

b) Statistics or conclusions drawn from an observation that is numerically distant from the rest of the data (outlier) can be misleading because they may not be representative of the real population. Outliers present a cause for concern because they may be the result of faulty data, errors in the extraction of the data, or flaws in the theory being tested. In effect outliers skew the results in a particular direction.

c) The addressable population for broadband performance monitoring is the total active broadband connections (a choice exists around access method) being consumed by end users (a choice exists around residential and commercial customers). It is clearly not possible to test this entire population due to cost and complexity constraints. However, the proportion of the population randomly selected for testing must be sufficient to provide statistical confidence that the resulting observations predict the performance that the remaining population experience.

d) The consequences of incorrect estimates or predictions for a defined population because of inadequate sample size can be significant. If the sample size is too small, conclusions drawn from the results may underestimate or overestimate performance. If the reported performance is less than actual performance (underestimated), the resulting investment to increase performance will drive significant costs into the industry, end users will bear the burden of these costs, and this will be for little marginal benefit. Conversely, if the reported performance is greater than actual investment (overestimated), necessary investment to improve performance may not be made. Both situations will be to the detriment of end users.

e) The challenges in selecting a sample size which is representative of the full population of New Zealand, equally apply to measuring an individual ISP's performance. ISPs with a small end user base may have a particular user type which is not representative of the full population. Consequently there are risks that such a small customer base will skew the ISP's results in a particular direction thereby overestimating or underestimating performance. It is important that the sample size selected for an ISP is verified as an appropriate size in order to make statistically valid conclusions about its performance.

f) The robustness of a selected sample and the ability to use the statistical conclusions to predict overall performance for the population will be a product of how the data is to be used. For example, collected data could be categorised by:

   i) Region (including New Zealand);
   ii) ISP; and
   iii) Access method.

g) As the collected data becomes disaggregated through categorisation the robustness of the sample size will need to be continually checked. For example, the sample size used for the purpose of understanding average New Zealand broadband performance may not necessarily be appropriate for the purpose of drawing statistical conclusions about broadband performance in a particular geographic region, by ISP or by access method.

h) The TCF considers it prudent at this stage to engage a statistics expert to determine what constitutes a statistically significant sample size.

4.2.4 **Methodology**

a) **Repeatable:** An ISP's performance will vary across a twenty four hour period due to the range of influences on broadband performance. As a result, it is important that tests are conducted on a regular basis. To enable regular testing each test must be repeatable. Conducting repeatable tests at regular intervals will provide a view of the temporal changes in broadband performance. Repeating tests will enable accurate comparisons.

b) **Accurate:** Because the results will educate end users about which ISP's broadband performance best meets their needs it is important that the test results be accurate. The analysis should not include an averaging of results. Due to the natural temporal differences in performance an averaging process merely smoothes out performance variants such as peaks, troughs or outliers, therefore failing to give an accurate view of performance. The standard industry practice employs an aggregation approach using medians.

c) **Verifiable:** Given the importance of accuracy the test results must be verifiable. This will involve transparency and ongoing review of the conducted tests and resulting data to minimise the implications of faults and errors.

d) **Cost Effective:** The cost of conducting broadband performance monitoring should be weighed against the delivery of the defined objectives. This would involve quantifying the benefits end users obtained from the monitoring of broadband performance.

e) **Exclude "external" factors from measurement:** Tests need to be carefully designed to ensure non-

measurement related load on an end user's broadband connection does not impact measurable performance. The same is applicable for an end user's LAN and PC. The performance or load of the end user's wired or wireless LAN or PC should not impact on the measurement of broadband performance. Excluding "external" factors will result in only the broadband connection being tested.

f) **Minimise the risk of gaming:** Public access to information pertaining to the location of test nodes should be restricted. Information relating to end users selected for testing and their location, or network test sites, should be retained as confidential. When testing ports and packet sizes a random selection process helps minimise the ability for ISPs to control performance in order to influence results. Undertaking these activities will help protect the accuracy of the data which is collected. The risk of artificially skewing performance will be constrained as the selected sample size grows.

g) **Dispute resolution:** When discrepancies arise about collected data there should be scope to initiate a review process – to verify the reliability of the data under question.

h) **Ongoing evaluation:** The performance of the monitoring regime should be assessed on a periodic basis to ensure the delivery of the objectives remains cost effective. This should keep the regime current and relevant while still meeting the objective of enabling customers to make informed choices about broadband performance.

i) From the international analysis undertaken by the TCF, it is clear that there is no commonality of approach.[2] Some approaches to broadband performance testing, such as downloadable software, capture a range of influences or drivers of broadband performance. This makes it difficult to effectively identify or isolate which of those influences or drivers is acting as a constraint on performance. Other approaches to broadband performance monitoring, such as network test satellites, can control the influence of a number of performance drivers, which helps identify which driver or influence is a performance constraint.

j) Of these different approaches, none are either wholly right or wrong. It is a matter of identifying which most

[2] Reference UK, Singapore & Australian analysis.

optimally delivers the objective of *"ensuring that customers are educated about broadband services and are able to access timely, accurate and meaningful information in order to make informed decisions about which broadband service best meet their needs"* in a way which is consistent with the methodological principles established as part of the research design.

4.2.5 Options for Carrying Out Testing

a) The TCF explored the advantages and disadvantages of four approaches that it believes are most closely aligned with the objective of broadband performance monitoring and the methodological principles.

4.2.6 Network Satellite

a) Dedicated network satellites involve constructing sanitised laboratory conditions at sites across the country. These laboratory conditions provide the best means of controlling the factors which relate to the customer's premises and equipment - in a manner consistent for each tested ISP.

Advantages	**Disadvantages**
Dedicated hardware resource	Limited hosting facilities within residential areas
Influence of external factors can be managed: > Variability in configuration parameters > Hardware and operating system performance > Routing management	Costs
Providers at the same location allows for comparisons: > Same loop length > Same conditions	Easier to game the process

4.2.7 End User Hardware

a) Providing an end user with dedicated hardware which sits between the end users router/modem and computer has advantages. The hardware can be configured to detect excessive traffic on the network and defer tests accordingly. Results collected from dedicated end user hardware will still be influenced by the wiring in the end user's premises.

Advantages	Disadvantages
Dedicated test device	Influence of house wiring (and uncertainty about the magnitude of this constraint on overall performance)
PC & OS influences removed	Does not identify the different responsibilities of the ISP and end user
Flexibility - detailed geographic view by loop length, unbundled areas and access technologies	Requires more tech savvy users in order to maintain (e.g. filters, manual install, power)
Cost effective relative to ISP-I	Cost of $75 to $100 per unit
	Not available in New Zealand today
	Temporal problem: P2P may be on for ten hours so delays in tests

4.2.8 **Downloadable Software**

a) Making software available for the end user to download is straight forward. However, results collected from this approach would have limited accuracy; they can be readily skewed by other traffic occurring in the end user's premises and computers may not be left on.

Advantages	**Disadvantages**
Low cost	No control of CPE
Potential for representative test sample	Limited range of tests
Easy installation	Not dedicated
	ISPs gaming without controlled participant selection
	Customer wireless LAN may be used
	Can't control when tests can be run

4.2.9 **Passive Network Tests**

a) Passive network testing involves using downloadable software that monitors actual traffic on a network. This approach concentrates on analysing live traffic going through the network interface of an end user's PC. Inspection of the packets will reveal the parameters of the network connection while based on end user-generated traffic. A passive network test option differs from the above options which involve constructing specific tests, based on international standards, to understand a broadband connection performance.

Advantages	**Disadvantages**
Low cost	Not dedicated
Potential for representative test sample	Results show actual performance (poor quality CPE or poor wireless coverage can affect the results)
Easy installation	
Based on user-generated traffic (i.e. measures end-user's activities)	
Measured performance is directly linked with end-user perceived performance	
No need to create sets of tests (any activity can be monitored)	
Very difficult to game by ISPs	

4.3. <u>HOW TO TEST - RECOMMENDED OPTION</u>

4.3.1 The TCF acknowledges, at this early stage of analysis, that if it wasn't for the anticipated costs of establishing a robust sample size, the network satellite option would be the most ideal for monitoring broadband performance. If it can be demonstrated that costs of ensuring a sufficient sample size of network satellites are outweighed by the benefits delivered to end users, the TCF would consider changing its recommendation to this option.

4.3.2 Due to the potential cost constraints of the network satellite option it seems practical to consider the option of deploying dedicated end user hardware, as described in paragraph 4.2.7. The hardware is located between an end user's router and operating system. Although this option captures the effects of household wiring in broadband performance it still isolates the influence of a number of performance parameters produced by an end user's equipment and configuration. This option has recently been selected by Ofcom in the United Kingdom. Ofcom rejected the network satellite option and the downloadable software option.

5. DISCLOSURE OF RESULTS

5.1. If end users are to be educated about broadband performance it is important they are able to access timely, accurate and meaningful information. This will aid end users in making more informed decisions about which broadband services best meet their needs. Therefore, careful consideration must be given on how to disclose collected data. The TCF has developed the following set of principles which can determine the most appropriate approach to disclosing broadband performance. These principles are:

5.2. **Simple and Easily Understood**

5.2.1 The disclosure of broadband performance should be targeted at those who have limited understanding of the technical nature of broadband delivery and the key influences on performance. This group is the least likely to be able to source the necessary information to make informed decisions about which broadband services best meet their needs. Consequently, for the information to be of use to this group, the presentation of results should be simple and easily understood. In order to aggregate the results into simple analysis, it will likely require a degree of interpretation based on agreed assumptions about what end users seek from their broadband connection.

5.2.2 End users with a strong grasp of the technical aspects of broadband should equally be able to access information which is simple and easily understood. This will more likely involve access to more detailed information involving minimal interpretation. Providing some interpretation is important because these end users are able to make informed choices about the tradeoffs which they want to make.

5.3. **Flexible**

5.3.1 It is important that there is flexibility in disclosing performance results given the diversity of understanding about broadband performance amongst end users. Failure to provide flexibility will result in the disclosure of broadband performance being of little use to large groups of end users. This would be inconsistent with the objective of monitoring broadband performance leading to informed decision-making amongst end users.

5.3.2 When trying to integrate flexibility in to the publishing of collected results, detail in the analysis becomes important. Providing detailed data/information enables end users to take the results, interpret these, and make their own tradeoffs based on their individual broadband performance preferences.

5.3.3 On the other hand, there is a tension between providing flexibility and simplicity. The more granular the disclosure becomes the more difficult it is for end users to interpret. The more basic the information the greater the assumptions required about what end users want from their broadband connection. This removes the ability for end users to make their own tradeoffs.

5.3.4 Jurisdictions currently undertaking broadband performance monitoring have adopted a range of approaches to publishing results[3]. Some issue reports with varying depths of analysis. The most obvious constraint in producing a report is the balance between detail, length and usefulness. An alternative to producing a report is to deliver the material in a web based interface – enabling the raw data to be collated and aggregated in a range of ways by an end user.

5.3.5 Adopting a web based interface allows an end user to access the raw data in a manner consistent with their level of knowledge and how they use their broadband connection. Such an approach has been successfully adopted in Singapore by the government's Infocomm Development Authority (iDA). End users have a web interface where they can select from a number of tabs. The end user is able to determine the desired flexibility and simplicity. It is the TCF's view that this is a more optimal approach for delivering on the education objective of broadband performance monitoring in New Zealand as it should address the needs of a greater cross-section of end users. It enables end users to make their own tradeoffs based on personal preferences.

[3] These jurisdictions are the UK, Australia, and Singapore.

5.4. In addition to end users, industry bodies will benefit from having greater access to more detailed data from the web based interface option. Industry bodies will be able to use their own internal expertise to analyse the data and provide analysis and interpretation for their particular membership. For expert users able to undertake their own analysis, they too can undertake similar assessments to inform themselves on broadband performance in their neighbourhood.

5.5. **Meaningful**

5.5.1 The TCF have considered three methods for making technical information meaningful to end users.

a) **Ranking**

i) Ranking results is a straightforward exercise that is readily understood by the public. End users will understand the difference between who comes first or last, and will factor in other elements to their broadband connection decision such as price and customer support, before deciding whether to switch provider or not. At face value this creates desirable incentives in the market. Each broadband provider is incentivised to improve their performance to reach the "top of the table".

ii) Unfortunately there is a significant disadvantage to this approach which must be carefully considered - ranking may lead to undesirable outcomes for the market. For example, there are minimal technical specifications which need to be in place for a particular broadband application to function or perform well. Performance that is greater than minimal technical specifications can be desirable. However, if the quality or performance is significantly greater than the technical specifications this will come at an additional cost but with no equivalent incremental benefit to the end user. It is in effect, a "gold-plated service". Consequently, if the industry provides a level of quality greater than the technical specifications while being subject to performance rankings, broadband providers will invest to out perform others, driving unnecessary costs into the industry that end users will bear for little or no additional benefit.

b) **Rating System**

i) A rating system is a commonly used tool which can minimise the aforementioned risk of quality improvements that create no additional benefit for end users. However, it requires the construction of a range of thresholds around technical performance in order to allocate ranking "stars" in an objective and transparent manner.

c) **Graphical Disclosure**

i) An alternative approach is to graphically present performance plotted against the technical specifications that the end user should expect. See Section 5.7.

5.6. Timeliness

5.6.1 Broadband performance is not static. Over a twenty-four hour period there are peak and off-peak periods where performance varies. As investment in broadband infrastructure continues to grow broadband performance will also change. With increasing uptake of broadband connections the growth in the customer base of ISPs will influence the performance of their broadband connections.

5.6.2 These temporal changes need to be captured and conveyed to end users in a timely manner. Failure to do so will impede the ability for end users to make optimal decisions about which broadband services best meet their needs.

5.6.3 Providing timely information depends on the time taken to conduct and collate the data, the complexity of the reporting, and the corresponding time required to have the report published. Naturally, delivering a robust report requires an investment of time in interpreting and analysing data. This delay may render the analysis useless. The advantage of adopting a web based approach to data collection and presentation only needs to be populated with the most recent information set. This ease of providing a predefined framework may facilitate more frequent reporting.

5.7. DISCLOSURE – RECOMMENDED OPTION

5.7.1 Applying the above criteria, the TCF initially concludes that a web based interface, adopting a predefined framework, is most optimal for the disclosure of results to end users. Collected data can be disclosed in a manner that meets the education needs of the greatest cross section of end users. Frequency of

reporting can readily be adjusted because the web based interface merely needs to be populated with the most current data. A graphical presentation, such as that below, provides the flexibility for end users to make their personal tradeoffs around performance accounting for temporal changes in performance. Baseline specifications for these common broadband activities are outlined in Section 6.9.

Performance Variable "X" – results between date "A" and date "B".

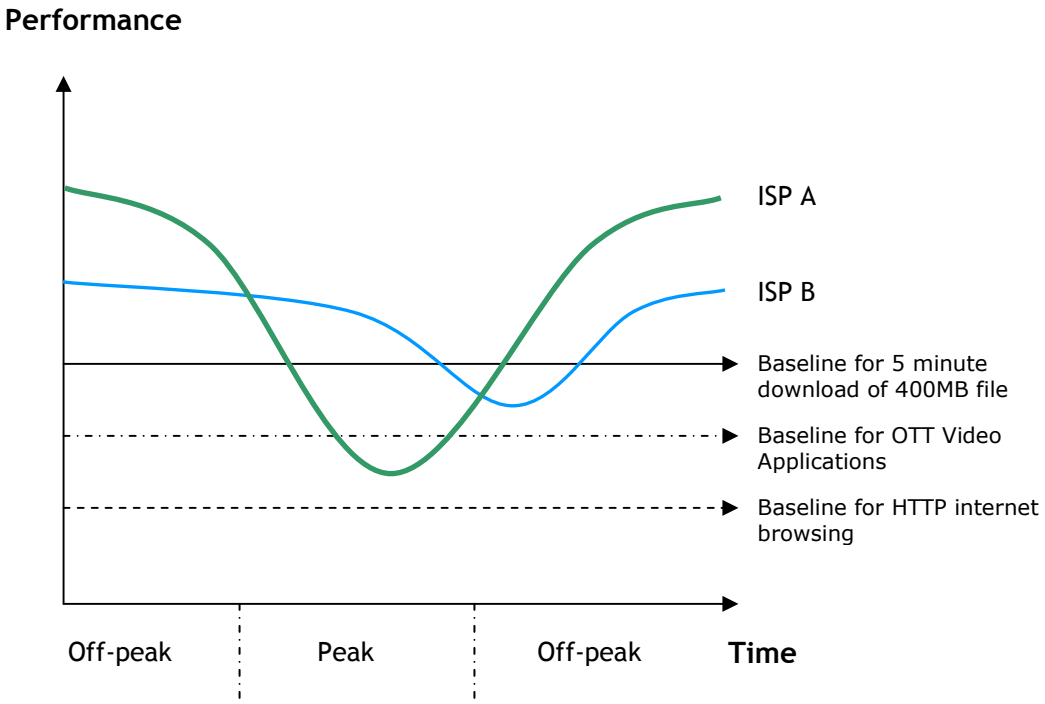

6. EDUCATION

6.1. In addition to educating end users about actual broadband performance it is equally useful to inform them about the key drivers of performance. Aiding end users to understand the magnitude of influence on performance criteria will help end users get more out of their broadband connection. The international research conducted by the TCF to date shows that the general public's awareness of key drivers of broadband performance is low. This limited understanding makes it difficult for end users to understand how they can optimise their broadband performance.

6.2. For example, the following factors were identified by Ofcom as factors which influence broadband speed[4]:

6.2.1 How near you live to an exchange;
6.2.2 The speed of router/modem;
6.2.3 Choice of ISP;
6.2.4 Speed of individual websites;
6.2.5 Number of people in the area using the Internet;
6.2.6 Processing speed of computer;
6.2.7 Choice of package;
6.2.8 Sharing connection with others; and
6.2.9 Quality of household wiring.

6.3. Limited understanding of these performance drivers results in end users being constrained in their ability to optimise performance. It also creates and perpetuates ill-informed expectations amongst end users about how ISPs and end users are able to enhance broadband performance. For example, if an end user's home has poor household wiring and antiquated equipment, such as an older PC, it may constrain the performance of their broadband connection. They are unlikely to experience the full benefits of a change regardless of which ISP they may switch to as a result of reported performance.

6.4. The physical nature of infrastructure and drivers, such as the length of the copper loop, will naturally be a constraint. Making spot comparisons of broadband performance between neighbours or friends may be meaningless because of the location of each person's house relative to the exchange or cabinet.

[4] Ofcom: UK Broadband Speeds 2008: Research Report, UK January 2009.

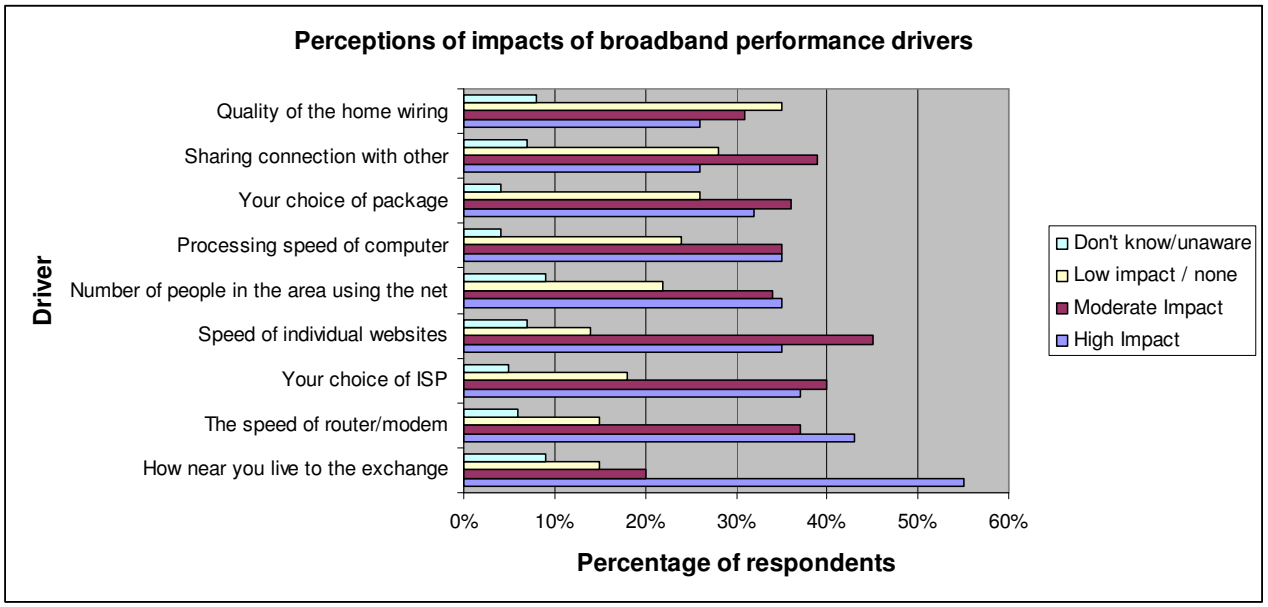

Source: Ofcom: *UK Broadband Speeds 2008: Research Report*, UK Jan 2009.

6.5. In Singapore the government's Infocomm Development Agency (iDA) has published a range of tips on how end users can get more from their broadband connection. The iDA has listed six areas where end users can optimise their broadband connections performance. These are:

6.5.1 **Monitoring the consumption of PC resources:** internet performance is dependent on the number of applications being launched and the processing load on their machines.

6.5.2 **Modem/Router configuration:** modems/routers can also under perform if they are not properly configured.

6.5.3 **Wireless LAN could be a bottleneck:** if using a Wi-Fi or wireless LAN connection, the raw data rate of 11Mbps, which can drop down to 6Mbps in actual data rate, may end up as a bottleneck to your broadband connection.

6.5.4 **Updating and Protecting your PC:** PCs infected by viruses or spyware can suffer from degradation in performance.

6.5.5 **Getting more from your broadband connection:** many bandwidth-intensive applications can far exceed upper limits on TCP. There are simple methods for managing this constraint.

6.6. End user expectations vary substantially about the broadband performance required to successfully run applications or use the internet. This is often because end users are generally unaware of the technical features which need to be in place in order to successfully utilise their broadband connection.

6.7. In Australia, the Department of Broadband, Communications and the Digital Economy (DBCDE) has published a guide on the type of broadband service an end user should purchase[5]. The DBCDE guide identifies a number of online activity categories for which end users may use a broadband connection, it defines light, medium and heavy user types, and then outlines what each user type should look for in a service. A similar approach has been adopted in Singapore by the iDA.

6.8. The Singapore and Australian jurisdictions are educating end users about the broadband performance they should expect in order to effectively use their connection for the online activities they are interested in. This helps mitigate the effects of mis-informed perceptions about overall performance. It also minimises the risk of ISPs engaging in excessive investment with no additional benefit to end users.

6.9. EDUCATION - RECOMMENDED OPTION

6.9.1 The TCF considers that further investigation should be made into adopting a similar approach to those used in Singapore and Australia. Defining technical specifications that an end user should consider when purchasing a broadband connection (based on categories of online activity) will add significant value to the outcomes from monitoring of broadband performance. Possible user category options are outlined below.

Web surfing (light)

Downstream throughput 256Kbps[6]
Upstream throughput 128Kbps
DNS < 100ms
Rationale: • Assumes average internet page of 130KB[7] • Assumes page should load in < 6s [8]

[5] Australian Government: http://www.dbcde.gov.au/communications_for_consumers/internet/broadband_for_consumers/choosing_a_broadband_service last accessed 29 Jan 2009.

[6] Given the above, the "Web surfing light profile" is based on loading a 130KB page in less than 6s. This gives:

Min speed = 130KB * 8 / 6 = 173 Kbps.
The above was rounded up to 256 Kbps.

[7] "The Average Web page", http://www.optimizationweek.com/reviews/average-web-page/, May, 2008

[8] "Speed Up Your Site: Web Site Optimization" by Andrew B. King, ISBN: 0-7357-1324-3

• Using multiple socket TCP

Web surfing (heavy)

Downstream throughput 1Mbps[9]
Upstream throughput 128Kbps
DNS < 100ms

Rationale: • Assumes average internet page of 130KB with addition of Flash heavy sites and similar • Assumes page should load in < 1s [10]

Video streaming (low quality)

Downstream throughput 512Kbps
Upstream throughput 128Kbps
DNS < 100ms

Rationale: • Assumes YouTube at standard quality[11]

[9] Given the above, the "Web surfing heavy profile" is based on loading a 130KB page in less than 1s. This gives:

Min downstream speed = 130KB * 8 / 1 = 1040 Kbps ~= 1Mbps.

An upstream speed of 128Kbps is sufficient to support the above downstream speed.

[10] "Designing Web Usability" by Jakob Nielsen, ISBN: 1-56205-810-X.

[11] YouTube System Requirements,
http://help.youtube.com/support/youtube/bin/answer.py?answer=78358&topic=17174

- Specifies minimum requirements as "Broadband connection with 500+ Kbps".

Video streaming (high quality)

Downstream throughput 2Mbps
Upstream throughput 128Kbps
DNS < 100ms
Rationale: • Assumes YouTube at High Definition quality[12]

Voice calling

Downstream throughput 128Kbps
Upstream throughput 128Kbps
DNS < 100ms
Packet loss < 0.1%
Latency < 100ms (one way)
Jitter < 10ms (one way)
Rationale: • Assumes G.711 encoding at 90Kbps • Follows ITU-T Y. 1541 recommendation • In line with Enhanced UBS specification • The latency and jitter relate to the ITU standards for PSTN replacement measurement points that would sit within an access seekers network • Internet voice services may be achieved to an acceptable level with less stringent specifications for latency and jitter

[12] YouTube HD http://www.mydigitallife.info/2008/11/20/how-to-embed-and-play-720p-hd-high-definition-youtube-videos-fmt22-code-hack/

- YouTube is currently providing 720p HD video in "beta". This activated by appending a "&fmt=22" to the end of URLs. However, YouTube is likely to launch this in the short to medium term.

- Codec details: 1280×720 (720p), H.264 video @ 1024Kbps; audio @ 44.1KHz 232Kbps Stereo.

- In other words, around 1024 + 232 = 1256 Kbps required.

Video calling

Downstream throughput 512Kbps
Upstream throughput 512Kbps
DNS < 100ms
Packet loss < 0.1%
Latency < 100ms (one way)
Jitter < 10ms (one way)

Rationale:

- Assumes G.711 encoding at 90Kbps
- Follows ITU-T Y. 1541 recommendation
- In line with Enhanced UBS specification

P2P file sharing

Downstream throughput 3 Mbps
Upstream throughput 0.5Mbps
DNS < 100ms
Packet loss < 0.1%

Rationale:

- Assumes 1GB file
- Assumes end user wants the download to complete in < 1 hour
- Assumes BitTorrent needs ¼ of downstream going upstream for good performance

7. CONCLUSION

7.1. Monitoring broadband performance is not a simple exercise. International analysis highlights that broadband performance monitoring is in its infancy with no convergence on a particular approach. This is a product of a range of factors. The technical nature of broadband, the diversity in domestic influences on broadband infrastructure, challenges to optimal study design and diversity in end user understanding, all make the construction and implementation of a successful monitoring regime challenging. For this reason the TCF strongly encourages the Commerce Commission to work with the industry to continue developing a broadband performance monitoring regime specifically tailored for New Zealand.

7.2. As a supporter and advocate of the Commerce Commission's desire to educate end users about broadband performance in New Zealand, the TCF would like to work with the Commerce Commission as it strives towards designing the optimal regime for monitoring broadband performance in New Zealand to ensure users have information from which to make informed choices in selecting the most appropriate broadband service.

7.3. In order to prepare the TCF to provide such support to the Commission, TCF members have dedicated considerable resources to constructing a transparent framework that could be used as a foundation for a broadband performance monitoring regime. The TCF is of the view that such a framework is essential to the effective design and implementation of a monitoring regime. In its absence there is a real risk that the objective of "ensuring that customers are educated about broadband services and are able to access timely, accurate and meaningful information in order to make informed decisions about which broadband service best meets their needs" will not be met.

7.4. The TCF's analysis has identified that the current approach of the Commerce Commission to broadband performance monitoring and reporting must be reconsidered. The ongoing presence of some substantial flaws and errors means that, in the TCF's view, adhering to the status quo fails to deliver on the Commission's objectives. In addition, the current approach raises real risks that actual broadband performance is underestimated or overestimated. The outcomes from such underestimation or overestimation are detrimental to end users and the industry as a whole.

7.5. In order to overcome risks such as "gold-plating" or providing incentives for under-investment, the TCF has identified a range of options which can help address these risks. As such, the TCF strongly encourages the Commission to consider these options. Considerable expertise and time

has been given to developing some viable alternatives that the Commission could use as a basis for consultation with the industry. Until such consultation has taken place, and subsequent solutions agreed amongst stakeholders, the TCF again requests that the Commission cease the publication of any further Broadband Performance Monitoring Reports.

8. NEXT STEPS

8.1. The TCF would like the opportunity to discuss the contents of this report and share initial thoughts with Commerce Commission officials as soon as practicable. Following this initial discussion, members of the TCF are keen to understand Epitiro and IDC's views on how the monitoring regime can be enhanced in order to improve end users' understanding of broadband performance.

8.2. As there is a variety of participants in the broadband sector - ranging from customers, to user groups, to ISPs and infrastructure providers – the TCF suggests that it would be useful and important for the Commission to hold a series of consultations and or workshops. This consultation process should ensure that all relevant stakeholders have the opportunity to participate in a collaborative process to test the TCF's initial thinking while informing and shaping the Commission's current approach to broadband performance monitoring.

9. GLOSSARY OF TERMS

Cached means high-speed storage mechanism.

DNS means Domain Name System/Service**.**

KPVs means Key Performance Variables.

Latency means a time delay between the moment something is initiated and the moment one of its effects begins or becomes detectable.

Packet Delay Variations (PDV) or (Jitter) means the difference in end-to-end delay between selected packets in a flow with any lost packets being ignored.

Packet Loss means the ratio of total lost packet outcomes to total transmitted packets in a population of interest.

Ping Performance means a utility to determine whether a specific IP address is accessible.

Throughput means the amount of data transferred from one place to another or processed in a specified amount of time.