samknows.com

Performance Monitoring Report

Revision:	Final
Date:	02/08/2008
Author:	Sam Crawford (sam@samknows.com)

0 Executive Summary

The SamKnows Performance Monitoring Network was born out of a desire to demystify the nature of broadband performance in the UK. The project operates completely independently of the ISPs and relies upon keen volunteers kindly donating access to their Internet connections for testing.

The solution employs the use of a small hardware monitoring unit, installed in the volunteers' homes between their home network and their ISP router. By utilising a hardware device it has been made possible to detect other network traffic and defer testing accordingly, which is essential for an accurate result whilst operating in an uncontrolled network.

Measuring broadband performance goes far beyond looking at speed (or "throughput") alone. The testing here reflects this, and the study has examined latency, packet loss, DNS resolution, web page loading, VoIP performance, sending of emails and, of course, speeds. The speed measurements delve deeper too, with comparison between web based speeds and typical peer-to-peer speeds, as well as looking at how running multiple streams affects speed.

The results produced over an initial six week period from 223 monitoring units were certainly interesting. In summary:

- In the majority of metrics there was little discernable difference between *most* ISPs;
- Zen Internet offered the fewest failures across all metrics;
- Virgin Media's cable services and Be/O2's services provided a consistently low latency throughout, whilst Virgin.Net (Virgin's ADSL service) performed poorly;
- BT provided the fastest throughput when measured as a percentage of *implied line speed* (an estimate of the potential maximum speed of the line);
- Be/O2 and Virgin Media produced the greatest raw throughput (in megabits per second), which can likely be attributed to the nature of their products;
- Virgin Media's cable throughput remained consistent on their 2, 4 and 10Mbps products, but was quite variable on their 20Mbps product;
- Testing highlighted the use of traffic shaping in the networks of BT and PlusNet, which resulted in certain classes of traffic slowing significantly during peak hours.

The project will continue into 2009, with improved testing metrics and a greatly increased sample size (to beyond 2000 units) being the focus points for future work. It is hoped that this will provide the necessary level of granularity to enable analysis of performance across ISP products and also across the regions.

A number of third parties have expressed an interest in operating their own deployments of this solution to monitor their own networks or those of the industry as a whole. Any such future deployments would be operated entirely independently (both physically and logically) from the SamKnows Performance Monitoring Network.

Contents

1 Introduction

This report details the results of the first round of our broadband performance monitoring project. Twelve UK based ISPs were tested over a period of eight weeks and this has led to some fascinating findings. The results have been presented with a detailed analysis, in an effort to avoid misinterpretation and/or misrepresentation of the data.

1.1 Motivation

The key motivations behind this project were four fold:

- The lack of a truly *independent* measure of broadband performance
- The lack of a statistically sound methodology to facilitate such performance testing
- The continuing perception that performance equals speed (it does not)
- The desire to dispel some myths

The frustration caused by the points above led us to ultimately develop the solution that has produced the results presented here.

2 Methodology

The testing methodology employed gives us the best combination of accuracy and access to a reasonable sample size of results.

2.1 Technical solution election process

Our key requirements for the methodology were as follows:

- Accuracy of data – Could not have other network traffic interfering with the results;
- Repeatability – Tests should be easily repeatable and we should be able to test a set of connections on a set schedule for the duration of the project;
- Ease of installation – If being deployed far and wide, the solution should be easy to install and thus result in a small support overhead;
- Cost – We do not have infinite resources;
- Adequate sample size – A target of 20 monitoring stations per ISP was set initially.

Installing monitoring hardware and/or software on *dedicated* broadband connections from all of the providers tested would provide a very clean test. However, the cost of this when dealing with any non-trivial sample size quickly ruled this option out. At this point it was realised that volunteers would be required to help with providing connections to test against.

Providing a software installable application for the volunteers seemed the next logical avenue to explore. However, whilst this satisfied the sample size and cost requirements, it fell down badly on the accuracy and repeatability fronts. Other traffic occurring on the network would not be detected by a software application, leading to the results potentially being skewed. Furthermore, many people do not leave their computers on permanently, so acquiring a regular set of time-series based results seemed unlikely.

The chosen solution offers a compromise between the above two. A volunteer's connection is utilised (in order to reach an adequate sample size), and a hardware unit is installed to overcome the accuracy and repeatability issues. By having a dedicated piece of hardware on the network, physically sitting between the volunteer's router and their Ethernet wired PCs allows the unit to detect excess traffic on the network and defer tests accordingly. Similarly, the nearest wireless network (above a certain signal strength) is passively monitored for traffic volume too, meaning that wireless traffic need not interfere with the results either.

The hardware is currently based upon the venerable Linksys WRT54GL [1] wireless router. This provides five Ethernet ports on the rear, as well as two wireless antennas. Volunteers are instructed to connect the WAN port to their existing router and then connect wired PCs to one of the four free LAN ports. Connecting additional switches behind the unit is perfectly acceptable too. The volunteer need not reconfigure their wirelessly connected computers.

The first thing to stress here is that this is not a "silver bullet". The results are only as good as the quality of the sample and outside factors (such as a damaged phone lines, faulty end user routers, etc) are not accounted for by the technical solution. Of course, validation of the results is performed, so suspect units' results will be excluded.

2.2 Software

A customised FreeWRT [2] firmware image was developed and installed on the units. At the point of delivery, this is all that is present on the device. Aside from a single script that checks for the availability of the software component upon boot, the physical unit contains no additional software. This is beneficial both from a security perspective (everything is destroyed when the power is lost) and also from a support perspective (any problems with a unit's configuration can be undone simply by power cycling it). New versions of the software can be delivered remotely without requiring a reboot.

The software itself utilises standard Linux tools (where possible) to perform the tests. Tools such as *ping*, *dig*, *curl*, *iperf* and *tcpdump/libpcap* have been used extensively. By relying upon the years of development and testing that has been poured into these applications we are helping to ensure the accuracy of our own results and can realise a reduced development overhead.

All monitoring units maintain accurate time using *ntp*.

2.3 Volunteer selection process

Within the first two days of promoting the project on the website we had over 1000 volunteers sign up to be involved in the testing – a far greater number than we had first anticipated.

The painstaking process that followed discounted those volunteers that:

- Were using an ISP not on the list (note the exception below)
- Mentioned in the comments the instability of their line
- Belonged to too high a concentration of users of the same ISP in one area
- Were not using an router

Attention was also paid to ensure a fair distribution between users of differing products on the same ISP.

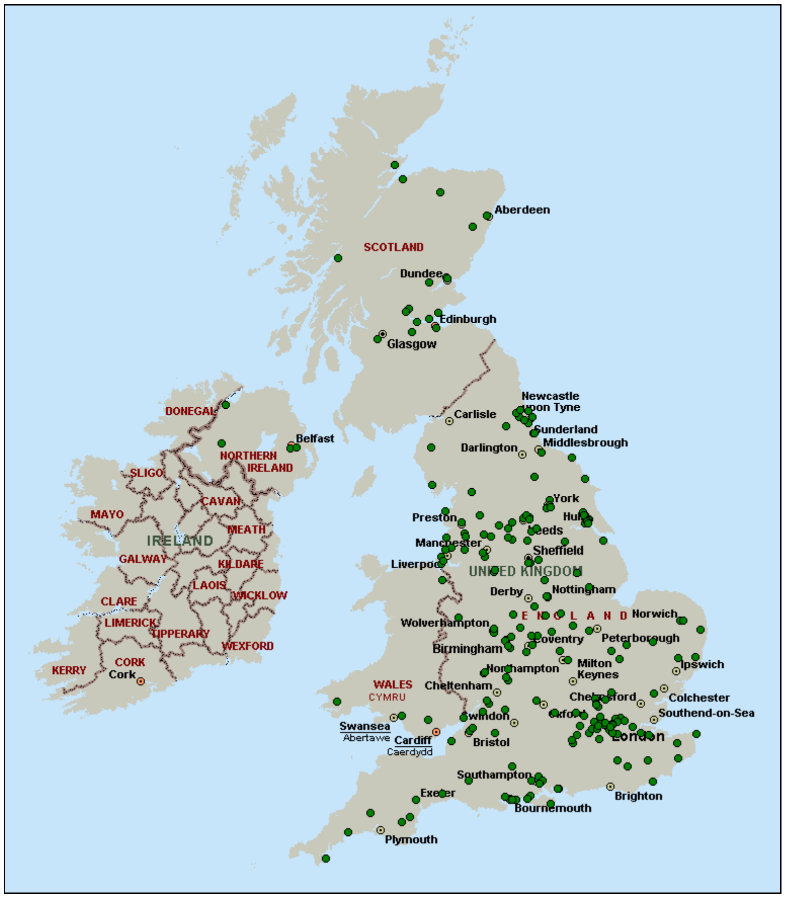

Figure 1 - Location of volunteers around the UK

In total, the results of 223 units were aggregated to produce the results detailed within this document. Some units' results were discounted due to clear configuration or user network issues (these were typically shown by 100% packet loss, or failures that exceeded the norm by a significant margin). The breakdown by ISP is as follows:

ISP	Units
Be Unlimited	24
BT	25
Entanet	18
Karoo	12
Orange	15
Plus Net	18
Sky	20
TalkTalk	16
Tiscali	15
Virgin Media	26
Virgin.net	15
Zen Internet	19
Total	223

Table 1 - Number of monitoring units by ISP

Note: AOL's results have been excluded from this report due to an insufficient number of monitoring units reporting data for the duration of the tests. An increased push for AOL participants is planned for the next round of tests.

Whilst 223 results have been presented here, there were in fact 258 devices involved in the testing. The missing 25 units were excluded from the results because either:

- The ISP was AOL (note above caveat);
- The results deviated massively from the norm of the ISP in question (typified by 50%+ packet loss, or web pages loading in ~20 seconds, etc);
- The volunteer contacted me to state that a fault had developed on the line;
- The unit was never powered on.

In total 9 people changed ISPs during the 2 months that data was collected for. The changes in ISP have been reflected in the results.

2.4 Testing ICMP latency and packet loss

Testing latency and packet loss is most commonly performed using the Unix utility *ping* and this solution is no different. In keeping with good practice, the first ping reply from any host is ignored (due to the delay in potentially having to ARP for the gateway) and an average of the following two is recorded as the result. Indeed, this is how Cisco's IPSLA [3] solution performs its own ping tests.

Four external hosts were "pinged" for the purposes of this test. Three were based in London, with a latency of less than 1 millisecond from a Docklands based server. The fourth was a very popular website based in Europe, with a latency of around 10 milliseconds from the same Docklands based server.

The average round trip time of the tests as well as the number of packets lost is recorded. This test runs every 10 minutes.

2.5 Testing recursive DNS resolver responsiveness and failures

Testing an ISP's recursive DNS resolution can be accomplished using many tools, such as *nslookup, dnsip* and *dig*. For the purposes of our solution, *dig* was chosen for the flexibility it offers the verbosity of the results provided.

Typically an ISP will have two or more recursive DNS resolvers. Rather than using the DNS servers provided by the DHCP leases to the testing units, the software on the units tests the ISP DNS resolvers directly. For example, *Be Unlimited / O2* use *87.194.0.66* and *87.194.0.67*. This allows us to determine failure of a single DNS server. Furthermore, it also overcomes another issue – that of people changing the DNS servers being returned in DHCP leases from their router (this proved quite common with customers on certain ISPs).

The tests record the number of milliseconds for a successful result to be returned. A successful result is deemed to be one when an IP address was returned (the validity of the IP address is not checked). A failure is recorded whenever the DNS server could not be reached or an IP address was not returned. The hostnames of four popular websites were queried every 10 minutes.

2.6 Testing web page loading times

This test utilises the *curl* utility to fetch the main HTML body of a website. Note that additional resources, such as images, embedded media, stylesheets and other external files were not fetched as a part of this test.

The time in milliseconds to receive the complete response from the webserver was recorded, as well as any failed attempts. A failed attempt was deemed to be one where the webserver could not be reached (non HTTP 200 responses were treated as successful results for the purposes of this test).

Three popular UK-based websites were tested every 30 minutes.

2.7 Testing VoIP capability

This test emulates the properties of a Voice over IP phone call in an attempt to determine how suitable the line is for VoIP purposes. Note that an actual VoIP call is not made – but the characteristics of it are emulated.

The test sends a 10 second burst of UDP traffic to one of three target servers residing on our network. Each UDP packet contains 160 bytes, and the traffic is sent at 64kbps. These characteristics match those of the G.711 [4] voice codec.

Please note: This only tests upstream bandwidth. Due to NAT implementation issues on some volunteers' routers, downstream testing proved too unreliable.

The test records the three major characteristics that determine the quality of a VoIP call: delay, loss and jitter. From these an R-value can be derived, and subsequently an estimated MOS (Mean Opinion Score) value. MOS is rated on a level from 1 (poorest) to 5 (perfect audio). The absolute maximum MOS value for G.711 is 4.4.

Also note: Our test assumes a worst case jitter buffer of zero milliseconds. Most VoIP capable routers (those that natively support VoIP channels) incorporate a small ~20ms jitter buffer nowadays.

This test was conducted once per hour.

2.8 Testing SMTP email relaying

Nearly all ISPs offer an SMTP relay for their customers to send email through. This test sends an email through the ISP's relaying SMTP server and records the time it was sent and received at the other end. The times are recorded in the email headers and are synchronised using NTP.

The ISPs' SMTP relay servers were determined from publicly available information on their websites.

All test emails are sent to the same email address, which is directed via MX records to the SamKnows servers.

The email test is conducted every hour.

Please note: All ISPs were tested with the exception of BT and Sky. BT began requiring authenticated connections as well as emails to originate from @bt.com addresses earlier this year. Similarly, Sky have begun using authenticated SMTP following their move to Google Apps based email. Due to lack of user credentials BT's and Sky's SMTP service could not be tested at this time.*

2.9 Speed tests

The project uses a wide variety of speed tests in order to monitor performance under different conditions. The list of such tests is as follows:

- HTTP download on port 80, single thread
- HTTP upload on port 80, single thread
- HTTP download on port 80, multi-thread
- HTTP upload on port 80, multi-thread
- HTTP download on random port over 1024, single thread
- HTTP upload on random port over 1024, single thread
- HTTP download on random port over 1024, multi-thread
- HTTP upload on random port over 1024, multi-thread

The terms "single thread" and "multi-thread" above refer to the number of simultaneous connections to the speed test server. A single threaded test uses a single connection (as one would typically find when downloading a file from a website). The multi-threaded test uses *three* simultaneous connections to complete the download.

All single threaded download tests download a randomly generated 6MB binary file per test. All single threaded upload tests upload a randomly generated 1MB file to the server using an HTTP POST request.

All multi-threaded download tests download a three randomly generated 2MB files simultaneously from the same server. All multi-threaded upload tests upload three randomly generated 500KB files to the server using an HTTP POST request.

Testing of non web-based traffic is emulated by using randomised port numbers in the ranges commonly associated with peer-to-peer traffic. It is acknowledged that HTTP traffic operating over ports other than port 80 can still be detected as HTTP (through the use of deep packet inspection). This method was chosen because it ensured validity in comparison between the two types of speed tests (as they used the same utility to test and the same servers). The future work section notes possible improvements to this specific test.

Additionally, it is understood that some ISPs operate transparent HTTP proxy servers on their networks. To overcome this, our webservers were configured to respond with the following headers, which should disable caching in standards-compliant proxy servers:

```
Cache-Control: "private, pre-check=0, post-check=0, max-age=0"
Expires: 0
Pragma: no-cache
```

All speed tests run every once every six hours (although each unit's tests may occur at any fixed point within that six hour period). This predictability of traffic volumes allowed us to accurately predict the capacity that we would have to cater for – something that online speedtesters do not have the luxury of.

Five speedtest servers were deployed in five different datacenters in and immediately around London to handle the traffic. Each server was monitored constantly for excessive network load and CPU, disk and memory load. Furthermore, the test results gathered by each server were compared against one another daily to ensure no significant variation in the speed attainable per server. Units cycled through the speed test servers in a round-robin fashion when testing.

2.10 Data aggregation

Storing all results in their raw form was an infeasible and unnecessary task. Some data aggregation was clearly necessary. The following details the level of aggregation used for each test.

- ICMP ping tests – Source data every 10 minutes, aggregated every 30 minutes
- DNS tests – Source data every 10 minutes, aggregated every 30 minutes
- Web page tests – Source data every 30 minutes, aggregated every 60 minutes
- VOIP tests – Source data every 60 minutes, aggregated every 60 minutes
- SMTP tests – Source data every 60 minutes, aggregated every 60 minutes
- Speed tests - Source data every 6 hours, aggregated every 6 hours

Of course, the graphs shown in the results may average data at a higher granularity than this (as there would simply be too many data points on the graphs otherwise).

3 Results

3.1 ICMP ping latency and packet loss

Latency is the measure of time for one packet of data to travel from your computer to the destination and back. A connection with a low latency will feel more responsive for simple tasks like web browsing, and certain applications will function far better with lower latencies. Online gamers, for example, will be particularly aware of the latency of their connections, as a lower latency than a counterpart will give them an advantage.

Figure 2 depicts the latency of the connections monitored over six weeks from 1 June to 12 July.

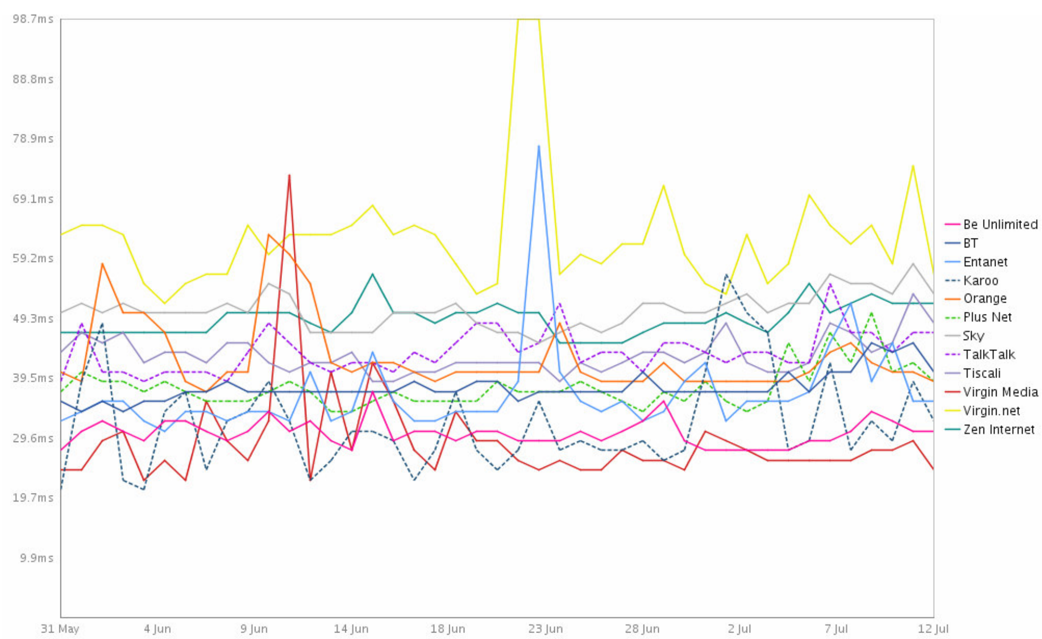

Figure 2 - Latency of all connections over a six week period

This rather chaotic graph actually shows that the vast majority of providers have rather solid performance, at least in terms of latency. Be Unlimited, Virgin Media's cable service and BT do particularly well here.

Sky has a rather high latency, which is perhaps surprising when you consider the majority of connections sampled here were on their LLU platform (which operates over the Easynet network). However, the reason for the latency will quickly become apparent after searching forums that Sky Broadband users frequent: Sky enables interleaving on all of their ADSL customers by default, as a method of increasing line stability (albeit at the cost of latency).

Aside from a few spikes from Entanet and Orange, the only ISP to stand out here for the wrong reasons is Virgin's ADSL service (aka Virgin.net). The below graph from a one week period in July highlights the disparity between the Virgin.net latency and the average of all other providers.

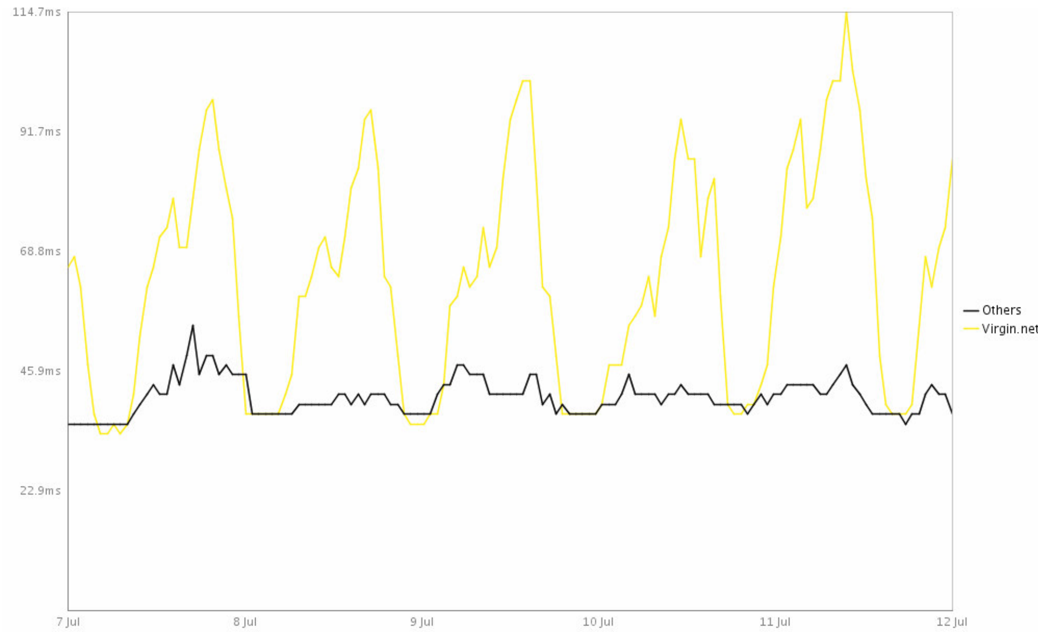

Figure 3 - Latency of Virgin.net connections versus others

Note in particular the cyclic nature of the graph, with latency spiking to around 100ms during the evening hours and then flattening out almost completely during the early hours of the morning.

Packet loss is a relatively rare occurrence in modern networks. Of course, it should be noted that if there is heavy congestion in the network then ICMP packets will be dropped by routers first (as they are deemed as lower priority than other traffic). For this reason a regular level of ICMP packet loss can be seen as an indicator of network congestion.

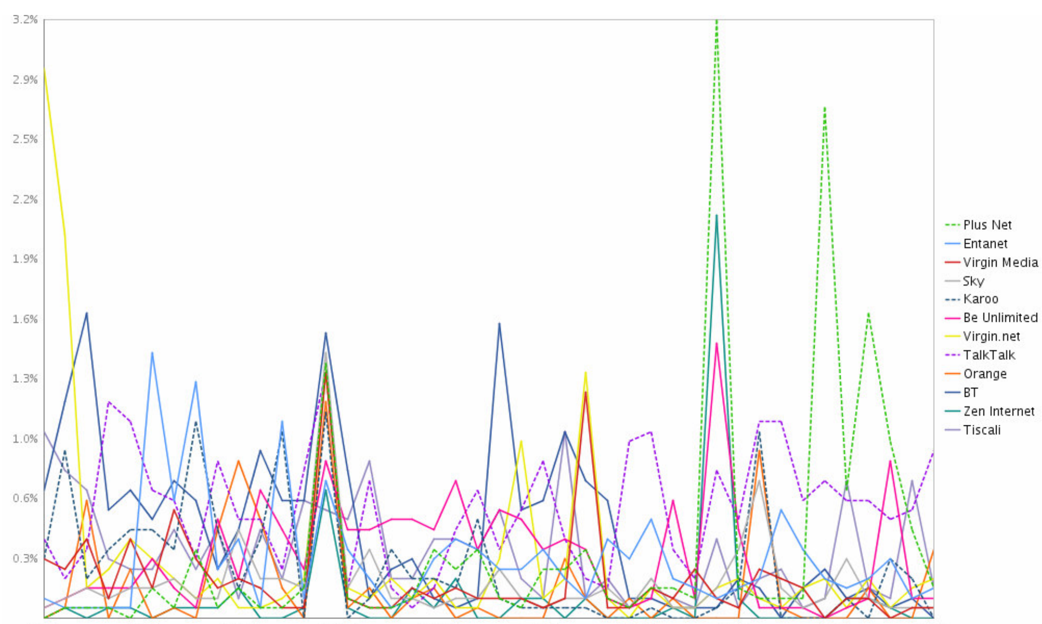

Figure 4 - Packet loss across ISPs (averaged into twelve hour periods)

Figure 4, above, indicates a low average packet loss across all ISPs, with the total loss peaking at about 1% during normal operation.

Some interesting characteristics can be observed from the graph though. During mid to late June some Be Unlimited users were suffering high packet loss at peak times. Whilst the loss never peaks above 1% on the graph, the cyclic nature suggests that this issue was being felt regularly, even if it was limited in scale.

Like Sky, Zen's average latency was rather higher than might have been expected. However, as the data tables below show, Zen have the lowest level of packet loss across all of the providers tested, with Sky not trailing too far behind.

It is worth noting that BT and some of the LLU operators had a noticeably wide spread of latencies, which is largely a result of the xDSL technology in place. Whilst some connections operated with a 10ms round trip time, others ran as high as 70ms. Future work will break these figures down further, so per product latency figures can be studied.

Many will also note the issue on July 1 that saw the results of many ISPs spike to well beyond the norm. Closer examination of the results show that this spike was felt across two of the four hosts being tested against, suggesting that some intermediate route between the ISPs and the target network(s) was at fault temporarily. A similar incident affecting all ISPs occurred on June 14.

Summary data tables for latency and packet loss

ISP	Latency (ms)
Be Unlimited	30.50
BT	37.84
Entanet	37.14
Karoo	32.13
Orange	42.83
Plus Net	37.68
Sky	50.31
TalkTalk	43.99
Tiscali	43.06
Virgin Media	29.03
Virgin.net	62.56
Zen Internet	47.27

Table 2 - Average ICMP latency by ISP

ISP	Loss (%)
Be Unlimited	0.35
BT	0.50
Entanet	0.32
Karoo	0.28
Orange	0.17
Plus Net	0.40
Sky	0.19
TalkTalk	0.58
Tiscali	0.35
Virgin Media	0.21
Virgin.net	0.34
Zen Internet	0.10

Table 3 - Average ICMP packet loss by ISP

3.2 DNS

The DNS (Domain Name System) predates the Internet itself. It allows computers to convert names such as *www.bbc.co.uk* to their associated IP address (e.g. *212.58.251.195*). Indeed, every website a person visits will require a DNS A-record query for the website's hostname. A slow DNS server will not affect download speeds, but it will severely affect the responsiveness of browsing around the Internet.

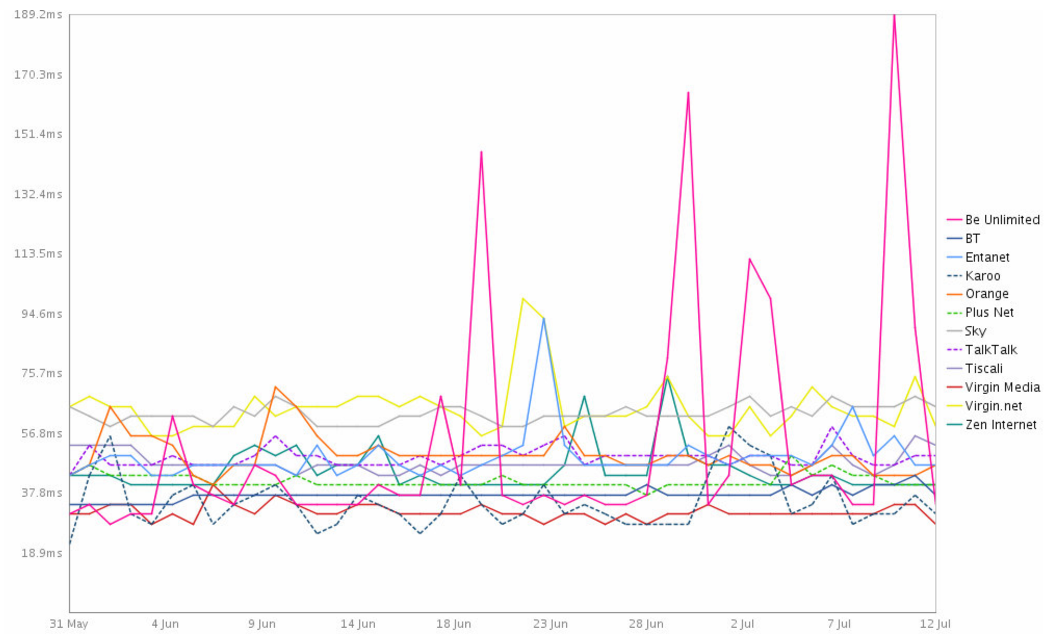

Figure 5 - DNS resolution time by ISP

As with the latency tests, the vast majority of providers perform well, with queries answered in an average of 45.62ms. There is close correlation to the results of those the ICMP latency tests here – BT and Virgin Media both perform very well here, whereas Sky and Virgin.net suffer (due to the inherent latency in the connections).

The only notable exception here is Be Unlimited. Frequent spikes indicate a problematic and intermittent DNS service.

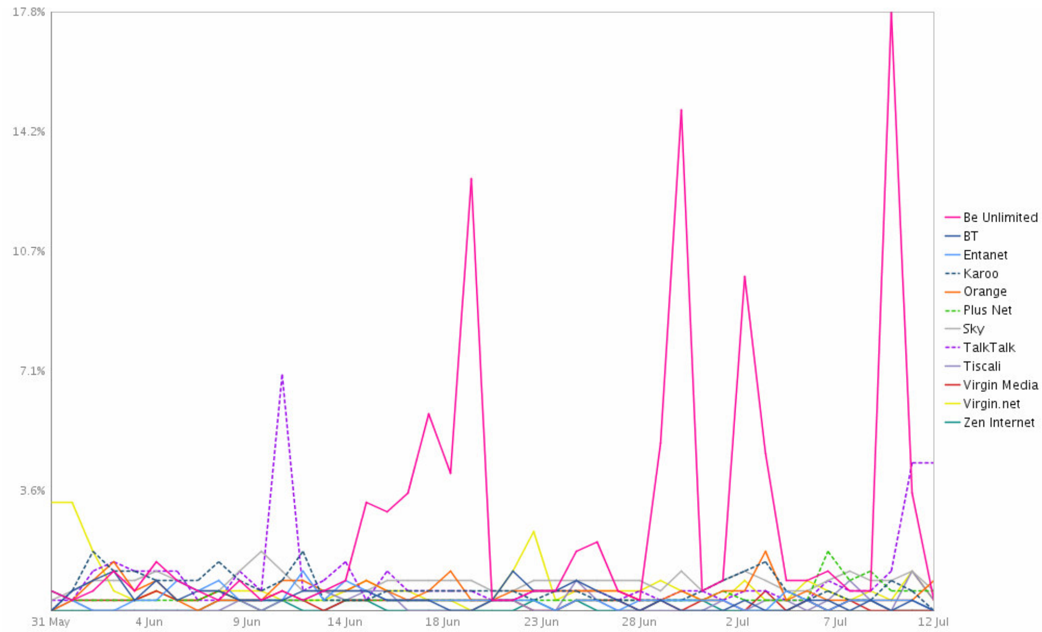

Figure 6 - DNS query failures by ISP

DNS query failures, depicted in the graph above, clearly highlight the issue affecting Be Unlimited users. Whilst the average failure rate across all ISPs is 0.81%, Be come in with a 2.82% failure rate across a six week period.

The problem with Be's DNS has improved though. Our earliest figures for Be, dating back to mid March, show the problem was even more pronounced back then.

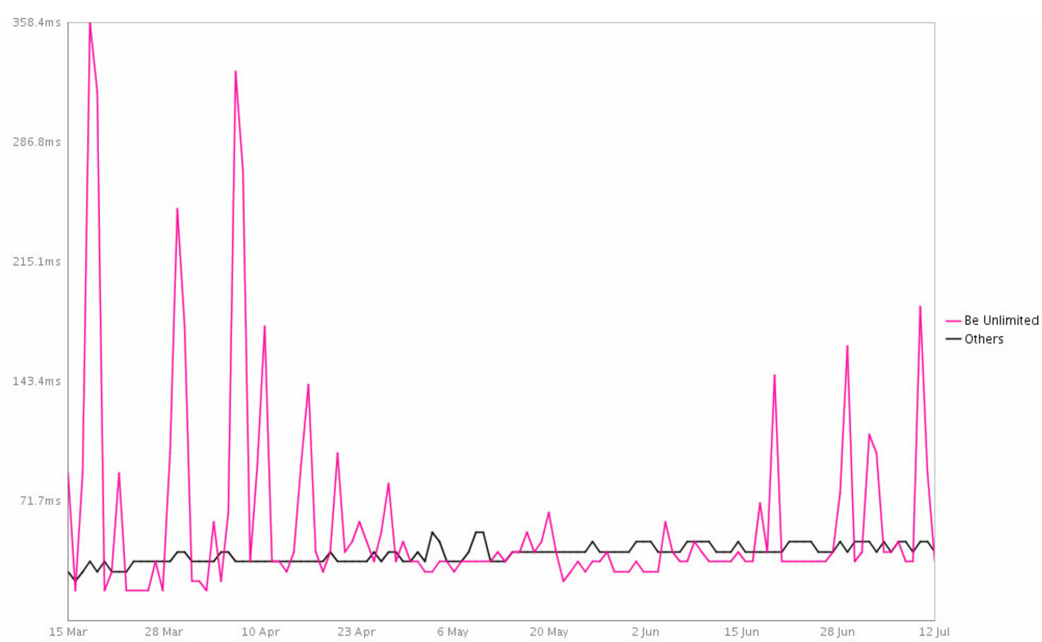

Figure 7 - Be Unlimited DNS resolution times

Putting the Be issue aside, all other providers come in with a sub 1% failure rate, which indicates a fairly robust DNS service. It should be noted Zen Internet lead the pack here by a considerable margin, having over seven times fewer DNS query failures than the average.

Summary data tables for DNS query failure and resolution times

ISP	Failure rate (%)
Be Unlimited	2.82
BT	0.47
Entanet	0.32
Karoo	0.69
Orange	0.59
Plus Net	0.39
Sky	0.87
TalkTalk	0.99
Tiscali	0.24
Virgin Media	0.25
Virgin.net	0.70
Zen Internet	0.11

Table 4 - DNS query failure rate by ISP

ISP	Time (ms)
Be Unlimited	53.90
BT	37.10
Entanet	49.34
Karoo	34.97
Orange	50.32
Plus Net	41.62
Sky	62.52
TalkTalk	49.10
Tiscali	46.94
Virgin Media	31.54
Virgin.net	64.40
Zen Internet	45.36

Table 5 - DNS query resolution time by ISP

3.3 Web page loading

As described in the methodology, this test measures how quickly the front page HTML of four common websites can be fetched. It is important to note that this is purely fetching the HTML – not any associated media resources (e.g. images, embedded flash content, etc).

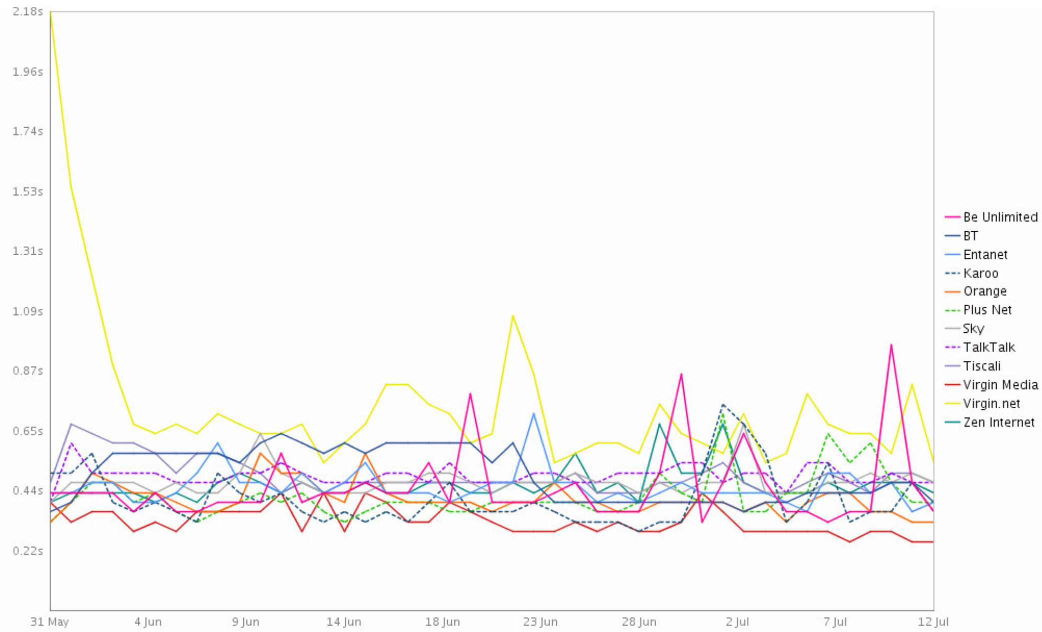

Figure 8 - Web page fetching times

Here we see close correlation to the results of the latency tests. Virgin Media's cable service does particularly well, with most others falling on or around the average. The poor performance of the Virgin ADSL service matches well to the poor performance they exhibited in the latency tests.

Be stands out as a notable exception here too. Whilst they may have had one of the lowest latencies, their web page loading performance is relatively poor. This can be explained by examining the correlation with their DNS performance (See Figure 9).

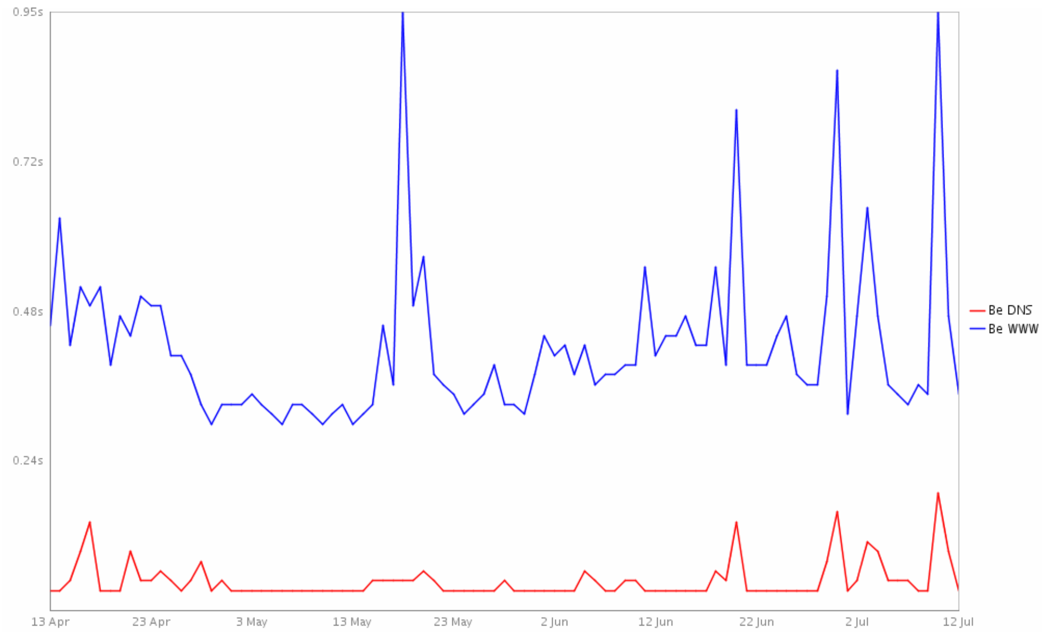

Figure 9 - Correlation between Be's DNS issues and their web page loading times

Note how the web page loading times (in blue) peak as the DNS query time increases (in red). This is a clear example of how poor DNS server performance can impact real world applications such as web browsing.

Failures whilst loading web pages should be quite rare on modern networks. Indeed, the average failure rate for the six weeks from June 1st was 0.34%.

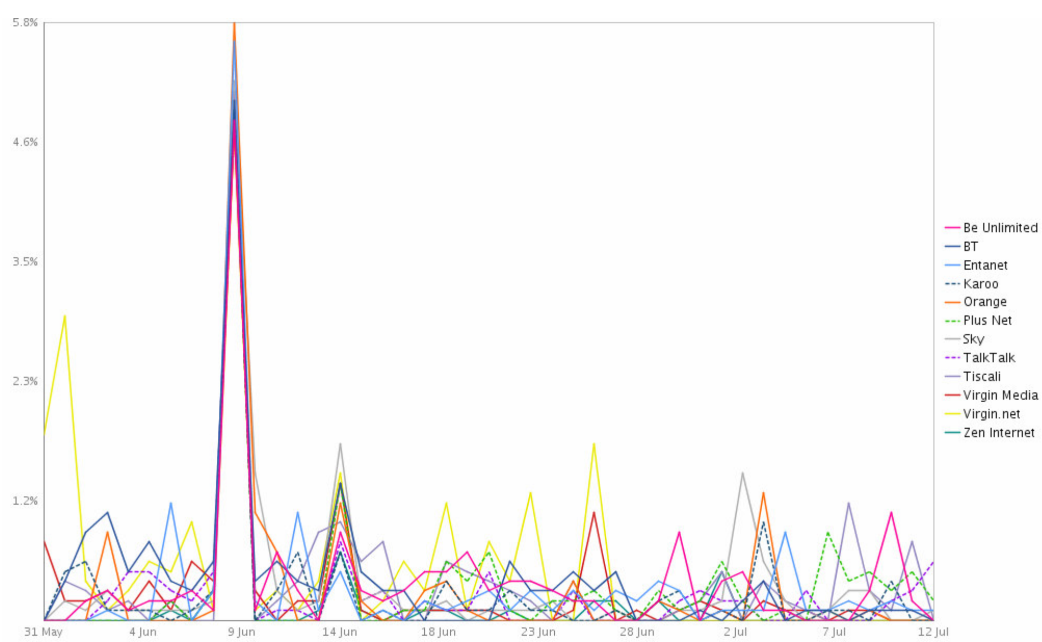

Figure 10 - Web page fetching failure rate

The large spike and small spike in early June (see Figure 10) can be attributed to a popular UK based website going offline for a number of hours. Note how the issues affected all providers simultaneously, indicating that the endpoint itself was at fault and not some intermediate route.

Again, Virgin's ADSL service comes across poorly here whilst Zen's performance tops this chart. Zen's success here is most likely attributable to their low packet loss and DNS query failure rate.

Summary data tables for web page fetching times and failure rate

ISP	Fetch time (ms)
Be Unlimited	450.31
BT	493.43
Entanet	446.82
Karoo	401.24
Orange	406.59
Plus Net	416.37
Sky	470.93
TalkTalk	492.09
Tiscali	486.95
Virgin Media	323.33
Virgin.net	712.64
Zen Internet	458.17

Table 6 – Web page fetching times

ISP	Failure rate (%)
Be Unlimited	0.39
BT	0.34
Entanet	0.32
Karoo	0.27
Orange	0.30
Plus Net	0.33
Sky	0.37
TalkTalk	0.29
Tiscali	0.38
Virgin Media	0.29
Virgin.net	0.47
Zen Internet	0.20

Table 7 – Web page fetching failure rate

3.4 Voice over IP performance

The VoIP test uses a short 10 second burst of UDP traffic designed to emulate a G.711 VoIP call.

Whilst the majority of factors that affect VoIP have already been covered above (e.g. latency, packet loss), this operates under rather different conditions. The latency and packet loss tests described earlier send a few ICMP packets (spaced one second apart) to common hosts on the Internet every 10 minutes. Whilst this has its uses, it does not accurately show how a VoIP call (which sends approximately 50 packets per second) would perform. VoIP packets are UDP based too, so should have a higher priority than ICMP ping packets.

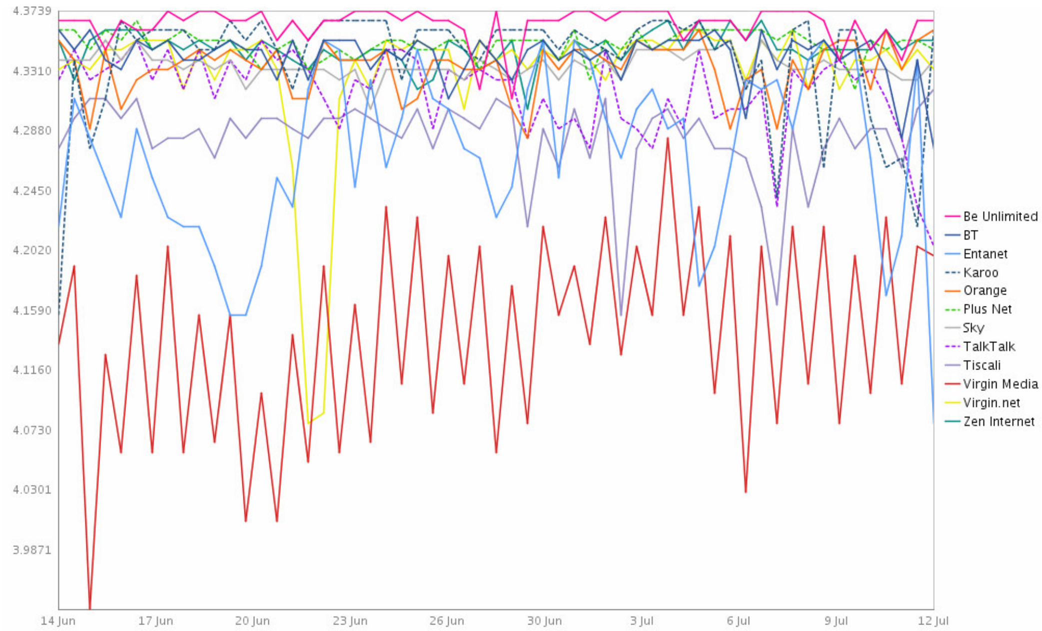

Figure 11 - Voice over IP Mean Opinion Score (MOS)

Figure 11 depicts the calculated Mean Opinion Scores of the simulated VoIP calls. Note that the vast majority are near the theoretical maximum of 4.4, which indicates near zero packet loss, very little jitter and a low packet delay.

The results at first may seem surprising. Virgin Media's cable service appears to perform poorly here, despite the fact that they performed very well in the latency tests. Virgin's ADSL service (labelled Virgin.net) seems to perform about average, in stark contrast to their previous results. The reason behind both of these is jitter. VoIP calls are particularly badly affected by jitter, which is the standard deviation of the packet delay. So, a connection that operates with 10ms delay but frequently spikes to 15ms might not be a problem for normal usage, but this 5ms represents a 50% increase in jitter.

This is what has affected Virgin Media's cable service here – their normally low latency is adversely affected by spikes within the 10 second UDP burst. Conversely, the Virgin.net ADSL service may have poor latency, but it is consistently poor – something that VoIP codecs are capable of dealing with.

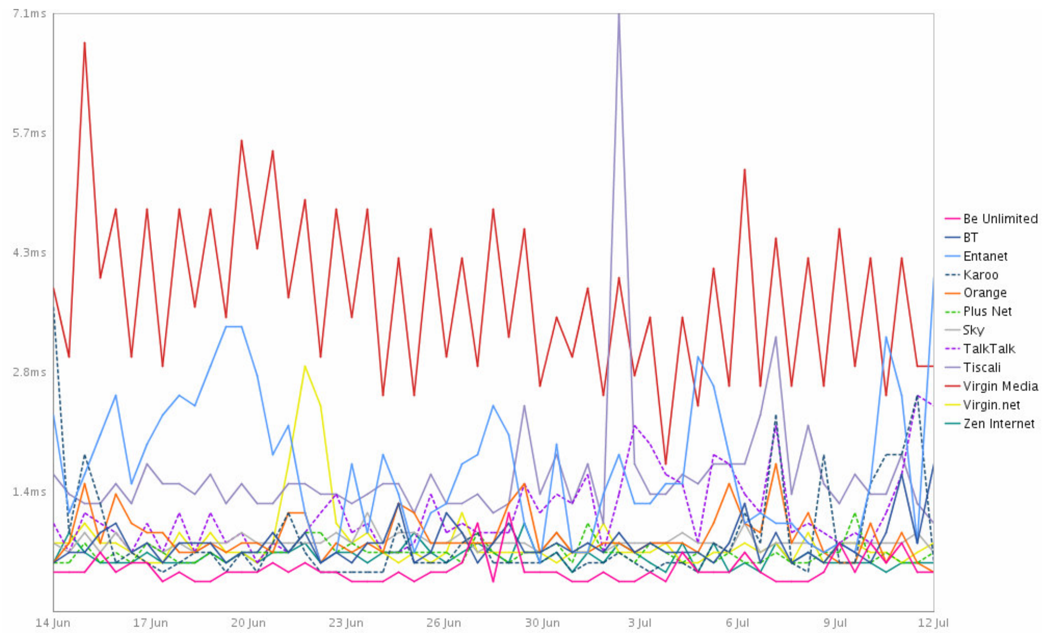

Figure 12 - Voice over IP jitter

Figure 12 shows precisely why Virgin Media and, to a lesser extent, Entanet appear to do badly in the VoIP tests. Their average jitter is considerably higher than other ISPs.

Of course, it is very important to note that we are seeing 6ms jitter at worst here – a level which most voice codecs will have no trouble with at all. Furthermore, many hardware-based VoIP routers now incorporate a "jitter buffer" (typically capable of safely handling jitter of 20ms or less) so the relatively low jitter exhibited here is unlikely to cause any issues. As discussed in the methodology, our test assumes a zero jitter buffer (effectively exaggerating real world results).

Packet loss is also a big factor in VoIP call quality. Whilst the odd dropped packet is acceptable (as each packet in the test here accounts for only 20ms of audio), extended periods of loss will lead to choppy and broken up audio.

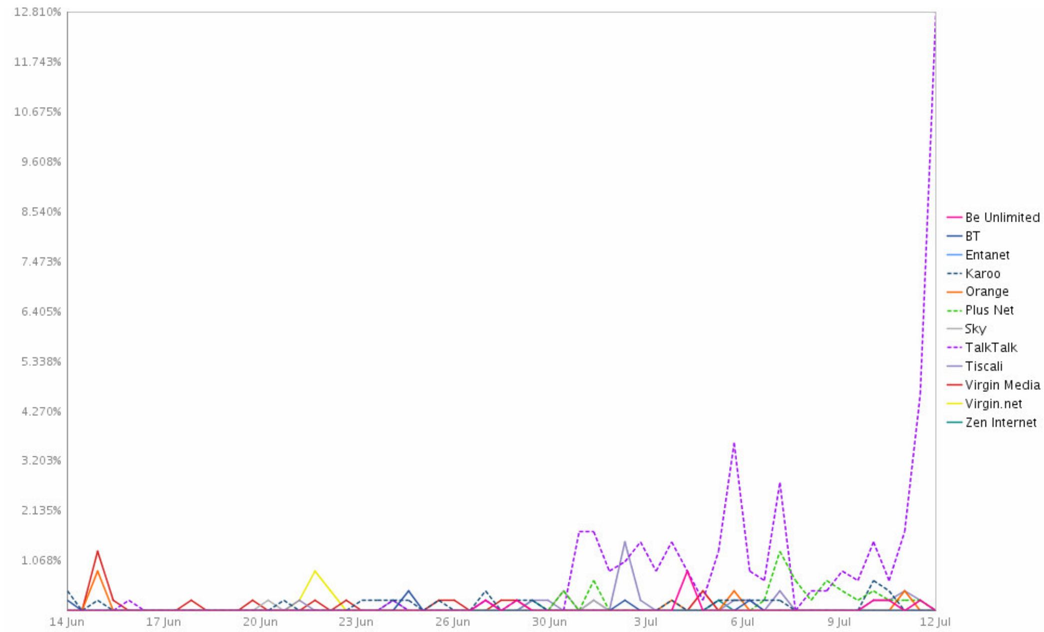

Figure 13 - Voice over IP packet loss

Nearly all ISPs operate at near-zero packet loss for this test, which is impressive. However, TalkTalk begin to feature in the graph for the wrong reasons as of early July. Of particular interest is the spike on July 12 that sees TalkTalk hit 18% packet loss – a very significant amount. Their previous results would tend to indicate that this is a temporary issue, but at the time of writing no newer data was available so we are unable to confirm this.

Summary data tables for VoIP tests

ISP	Calculated MOS
Be Unlimited	4.36
BT	4.34
Entanet	4.27
Karoo	4.34
Orange	4.33
Plus Net	4.35
Sky	4.33
TalkTalk	4.32
Tiscali	4.30
Virgin Media	4.14
Virgin.net	4.33
Zen Internet	4.35

Table 8 – VoIP Mean Opinion Score by ISP

ISP	Jitter (ms)
Be Unlimited	0.53
BT	0.85
Entanet	1.78
Karoo	0.82
Orange	0.89
Plus Net	0.69
Sky	0.85
TalkTalk	1.08
Tiscali	1.43
Virgin Media	3.74
Virgin.net	0.79
Zen Internet	0.66

Table 8 – Jitter when sening 501 UDP packets

ISP	Packet loss (%)
Be Unlimited	0.03
BT	0.05
Entanet	0.02
Karoo	0.11
Orange	0.05
Plus Net	0.10
Sky	0.18
TalkTalk	0.38
Tiscali	0.08
Virgin Media	0.11
Virgin.net	0.02
Zen Internet	0.01

Table 9 – Packet loss when sending 501 UDP packets

3.5 SMTP relay performance

Unfortunately not all ISPs could be tested for this specific test, due to lack of SMTP credentials that are required for BT and Sky connections. For this reason these two ISPs have been excluded from this test. This is something we are working to resolve for the next report. Entanet also do not feature in this test, as their resellers typically provide their own SMTP servers for customers to use (so there is no common test for us to run here).

Figure 14, below, shows the percentage of emails sent via the ISP's SMTP relay to reach a common destination that took over 3 minutes.

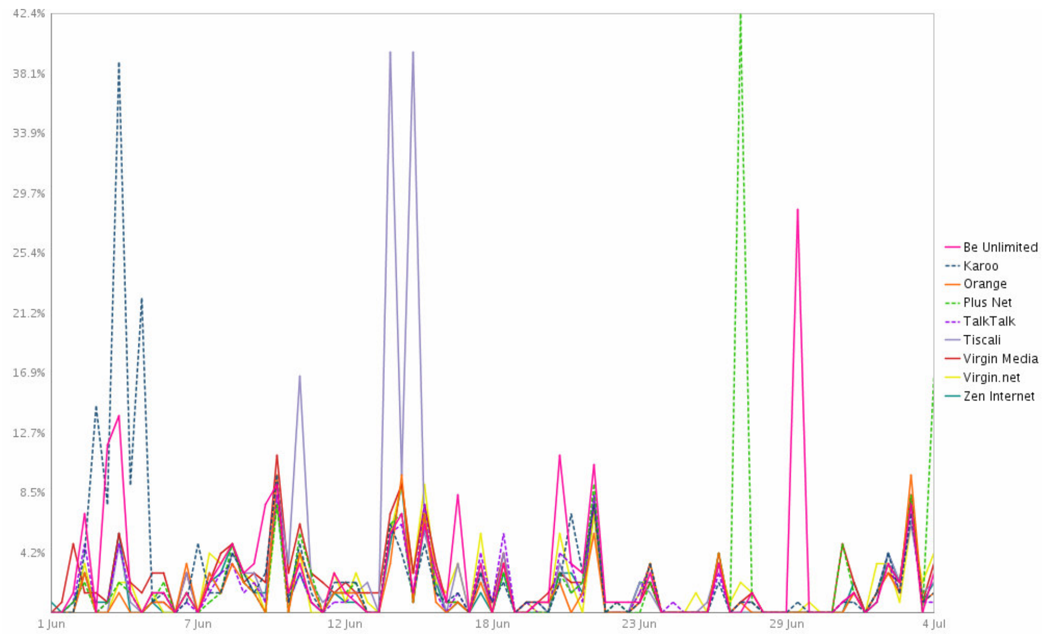

Figure 14 - Percentage of emails taking longer than 3 minutes to be delivered

One important thing to note is the occasional spikes that affect all ISPs, suggesting that the destination mail server itself (or the route between) is somehow at fault. These spikes were investigated and an over-zealous spam filter was discovered to be the root cause.

Ignoring this glitch, we see that most providers are delivering the vast majority of email within three minutes. There are clear spikes affecting Karoo, Tiscali, Plus Net and Be, but the infrequency of these suggests that it is not a recurring problem and can most likely be attributed to internal problems or temporarily high loads on the mail servers.

3.6 Speed tests

It is an unfortunate truth that the broadband industry is obsessed by speed. It seems that a week cannot pass without someone announcing the results of a study that shows ISP X to be the fastest, based upon Y thousand speed tests over a Z month period. Whilst the methodology (and sometimes motives) behind such studies may be questionable, their eager consumption by the public, ISPs and regulators alike suggests that the issue will be with us for some time yet.

Not wishing to disappoint, we have included our own speed study as a part of this larger project. Rather than presenting a simple table stating "ISP X is fastest with 6.2Mbps" we have examined the performance of the providers in depth and compared them not only to each other, but also to themselves.

It would be unfair to compare the performance of two ISPs directly using raw speed (in megabits per second) as the measure. Be Unlimited, for example, have a headline speed of 24Mbps and will only provide service to those within a certain distance of the exchange whilst BT Broadband only offer up to 8Mbps to anyone. Comparing the two side by side in terms of raw speed will almost certainly result in Be Unlimited outperforming them every time. Whilst this is still a valid metric (and one that we will not ignore), it does not tell us much about how well the connections perform relative to their maximum throughput.

For this reason the reader will note that the majority of our speed test results are expressed in percentage terms. This is a percentage of the *implied line speed*, which is defined as being the maximum throughput achieved across all speed tests within a two day period. The *multi-threaded* speed test tends to push a connection to its limit, and it is often the result of running this test in the early hours of the morning that produces the *implied line speed*.

By using this percentage scale we can directly compare two ISPs against one another on the same graph in a fair and consistent manner.

3.6.1. Port 80 HTTP download speed tests

This test emulates a task we are all familiar with – downloading files from a website.

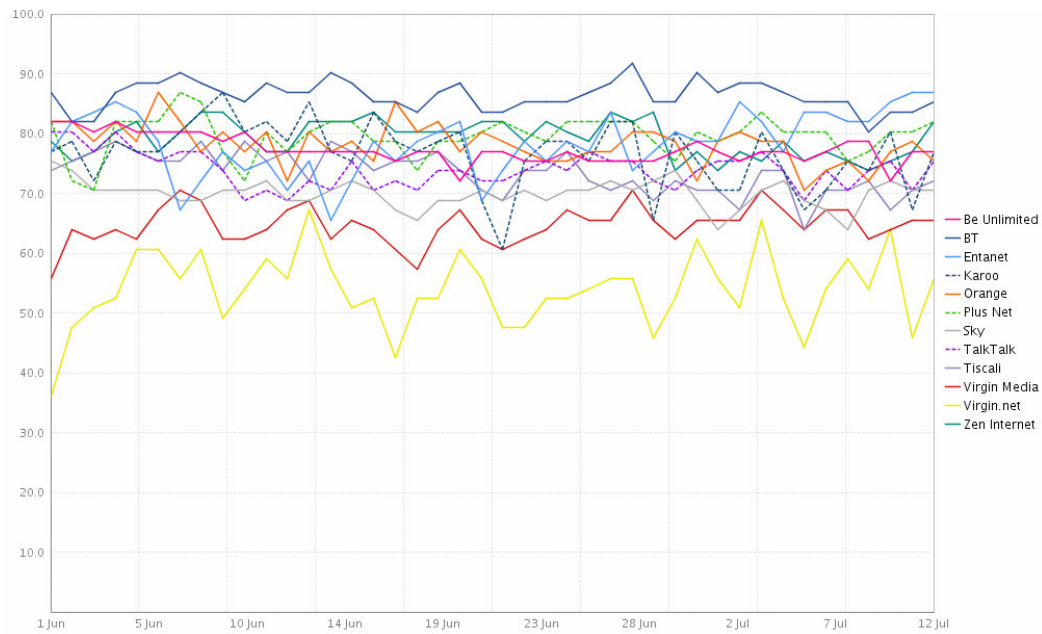

Figure 15 - Port 80 downstream throughput as a percentage of implied line speed

Figure 15, above, depicts the daily averaged download speeds of each ISP, expressed as a percentage of the implied line speed. Most providers hover around the 75% mark, meaning that if you can reach 6Mbps under optimum conditions you might expect to achieve 4.5Mbps on average.

BT noticeably outperform the rest of the ISPs in this test, whilst Virgin.net and even Virgin Media's cable service both suffer somewhat. BT's strong performance here may partly be due to the fact that they use the BT Central Plus [5] product from BT Wholesale, which provides LNS-based access to their core network via ten region PoPs. ISPs using the standard BT Centrals would typically take an L2TP handoff at one location. BT Retail are the only BT Wholesale customer, to our knowledge, that utilise the BT Central Plus product.

A subset of the above data, averaged into six hour data points, was used to produce Figure 16 (below).

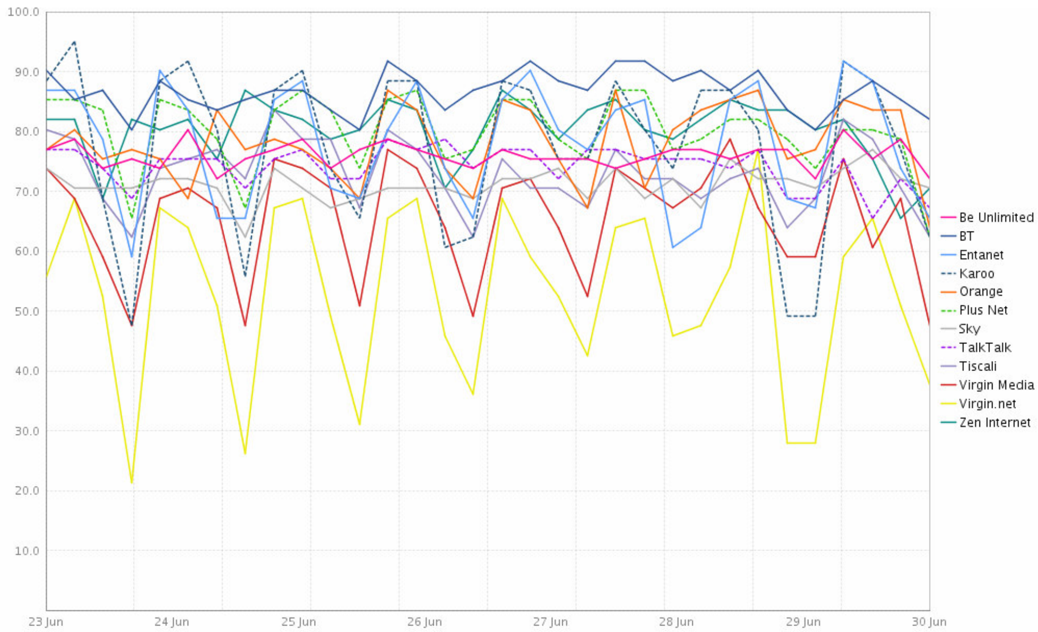

Figure 16 - Port 80 downstream throughput as a percentage of implied line speed (one week)

Here we see the reason for Virgin's relatively low performance in the previous graph. Both their ADSL and Cable service drop significantly in speed during peak hours, then ramp back up to the average in quieter hours.

Providers such as Be, BT, TalkTalk and Sky all maintain a near constant speed throughout the day. Given the higher average connection speed of their customers, this is a particularly impressive feat on the Be network.

The raw throughput (in Mbps) graph is shown below, again averaged over 24 hours.

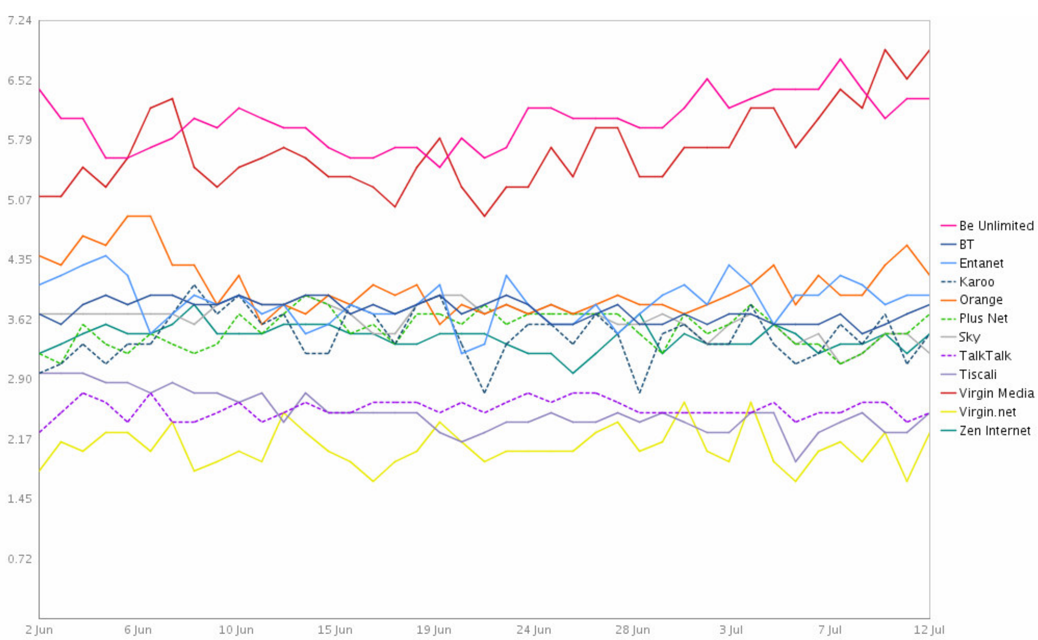

Figure 17 - Port 80 downstream raw throughput (in Mbps)

Unsurprisingly Virgin Media and Be Unlimited lead the pack in this view, which can be accounted for by their higher proportion of customers on faster connections (See the section on Variation in Implied Line Speeds for more).

The contrasting results for Virgin Media highlight how much difference data presentation can make. Whilst Figure 17 shows them to have one of the fastest average raw throughput speeds, the variability in the performance is high, with it falling to nearly 50% during peak hours (see Figure 16).

The Virgin Media results are particularly interesting once you begin to break the results down by product.

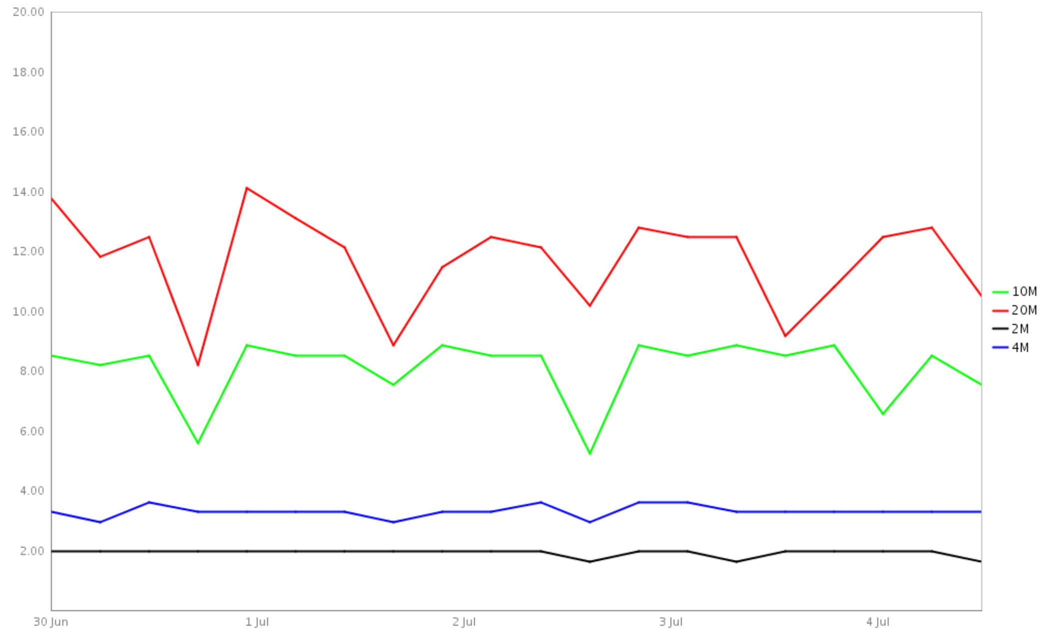

Figure 18 - Virgin Media cable HTTP speeds (Mbps) by headline product speed

Figure 18, above, shows the reason for Virgin's dips in speed very clearly. The 2Mbps and 4Mbps products maintain their full speed consistently, and the 10Mbps product does a good job too, apart from the occasional dip. The 20Mbps product is where the issue lays, with significant dips seen every evening (the above graph being shown over a 5 day period).

It should be noted that whilst the maximum speed above for the 20Mbps product is shown to be around 14Mbps, this is of course an average over all of the 20Mbps lines monitored. Significant variation was seen between individual 20Mbps lines, with many running at 20Mbps during off-peak hours, but some dipped down to below 10Mbps (hence the average presented above).

Summary data tables for port 80 based download speed tests

ISP	Percentage of implied line speed
Be Unlimited	77.35
BT	86.24
Entanet	78.34
Karoo	76.60
Orange	78.31
Plus Net	79.66
Sky	69.99
TalkTalk	74.04
Tiscali	73.57
Virgin Media	64.83
Virgin.net	53.95
Zen Internet	79.37

Table 10 – Port 80 downstream throughput
(percentage of implied line speed)

ISP	Throughput (Mbps)
Be Unlimited	6.32
BT	3.94
Entanet	4.01
Karoo	3.60
Orange	4.22
Plus Net	3.69
Sky	3.82
TalkTalk	2.66
Tiscali	2.64
Virgin Media	5.88
Virgin.net	2.15
Zen Internet	3.68

Table 11 – Port 80 downstream throughput
(raw Mbps)

3.6.2. Port 80 HTTP upload speed tests

As might be expected, upstream bandwidth does not suffer anywhere nearly as badly as downstream bandwidth.

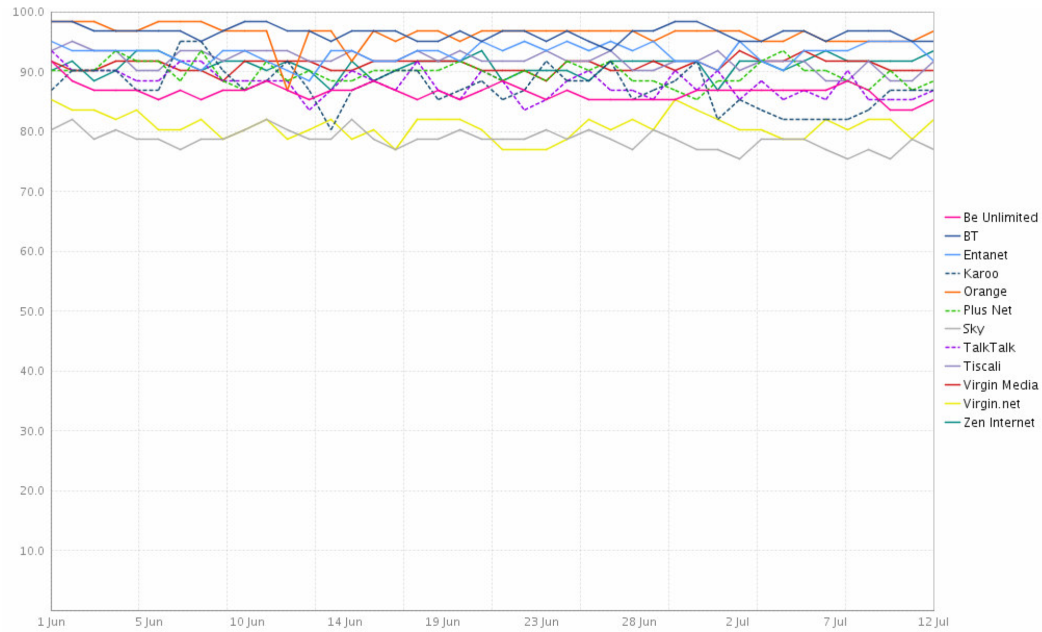

Figure 19 - Port 80 upstream throughput as a percentage of implied line speed

Figure 19 shows nearly all providers hitting near full upstream speed all of the time. Orange and BT perform particularly well, with Virgin.net and – surprisingly – Sky falling behind in this test. Again, this could be due to the fact that interleaving is enabled on all Sky lines as standard, although this should not significantly affect sustained transfer speeds.

Summary data tables for port 80 based upload speed tests

ISP	Percentage of implied line speed
Be Unlimited	86.54
BT	96.35
Entanet	92.96
Karoo	87.35
Orange	96.31
Plus Net	89.75
Sky	78.84
TalkTalk	87.86
Tiscali	91.90
Virgin Media	91.11
Virgin.net	80.91
Zen Internet	90.93

Table 12 – Port 80 upstream throughput (percentage of implied line speed)

ISP	Throughput (Mbps)
Be Unlimited	0.86
BT	0.35
Entanet	0.48
Karoo	0.34
Orange	0.37
Plus Net	0.45
Sky	0.50
TalkTalk	0.32
Tiscali	0.34
Virgin Media	0.50
Virgin.net	0.33
Zen Internet	0.39

Table 13 – Port 80 upstream throughput (raw Mbps)

3.6.3. Download tests over other ports

Downloading files from sources other than websites is becoming more and more prevalent nowadays. Peer to peer is often cited as consuming vast quantities of Internet traffic, and increasingly the same can be said for streaming video services (such as YouTube, BBC iPlayer and 4 on Demand).

This test downloads files from the same servers as the Port 80 tests discussed previously, but instead downloads them from a range of different port numbers that would normally be associated with peer to peer traffic.

Rather than presenting a single graph that compares providers against one another, a single graph will be shown here per ISP, comparing the Port 80 results (shown previously) to the non port 80 results.

Note again that all speeds are expressed as a percentage of the implied line speed. On each graph the light blue line represents the same port 80 based speedtest results as discussed earlier, and the darker blue line represents the non port 80 speedtest results.

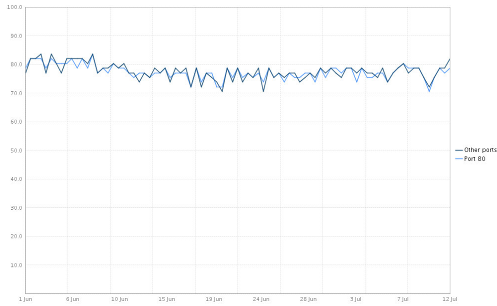

Figure 20 - Be Unlimited traffic

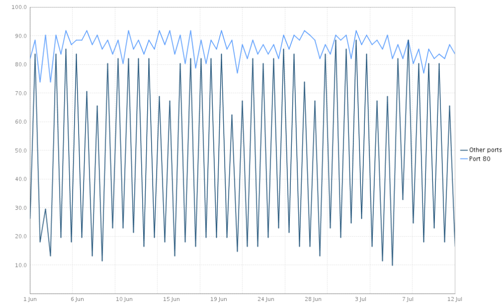

Figure 21 - BT traffic

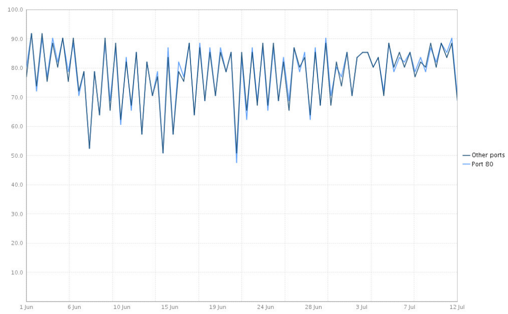

Figure 22 - Entanet traffic

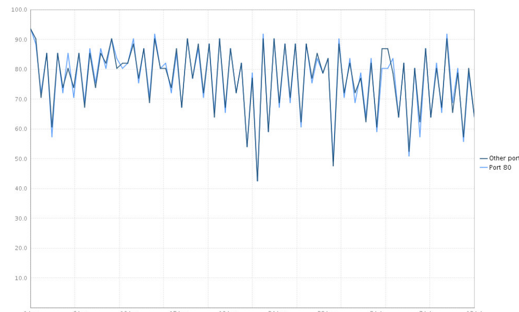

Figure 23 - Karoo traffic

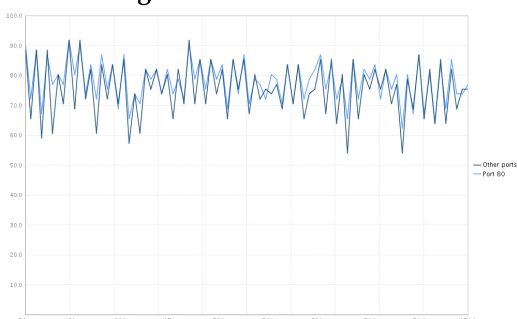

Figure 24 - Orange traffic

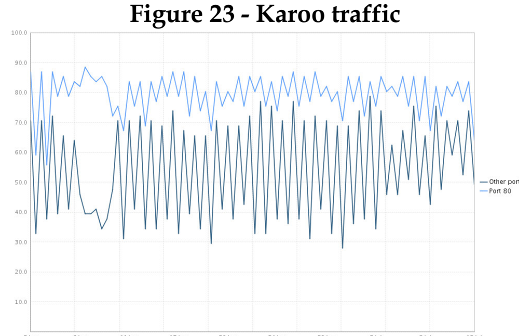

Figure 25 - Plus Net traffic

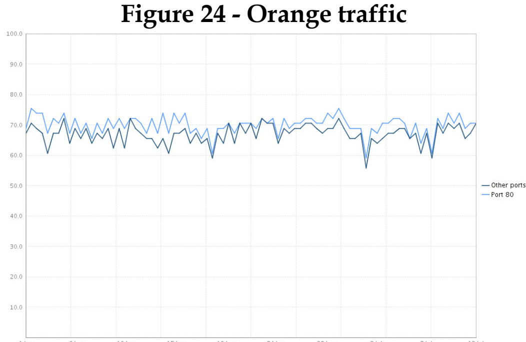

Figure 26 - Sky traffic

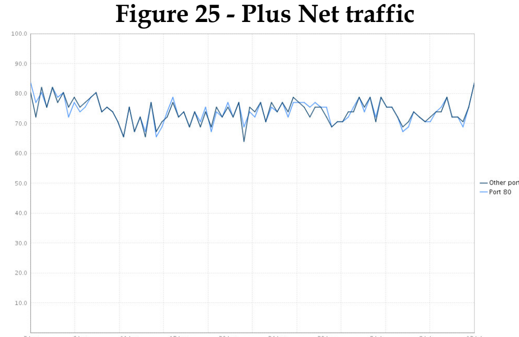

Figure 27 - TalkTalk traffic

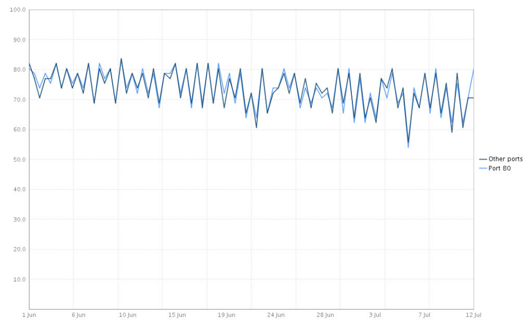

Figure 28 - Tiscali traffic

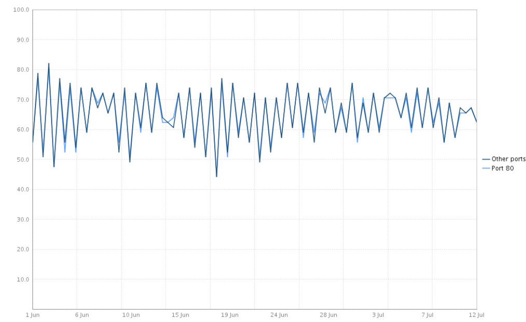

Figure 29 - Virgin Media traffic

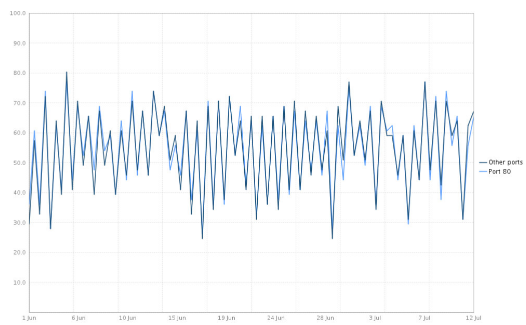

Figure 30 - Virgin.net (ADSL) traffic

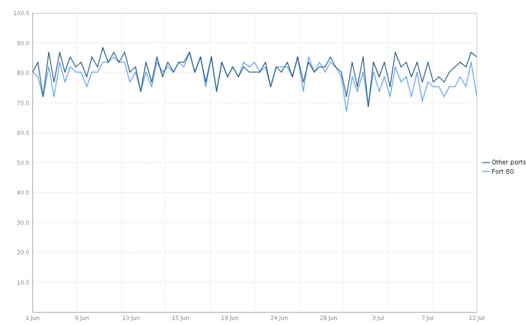

Figure 31 - Zen Internet traffic

In all but two cases we see that the two lines very nearly match, and given that both tests are being run consecutively, this is what one would expect on a normal network.

However, in the case of both BT and Plus Net we see something rather different. Whilst the port 80 test performed very well at all times, non port 80 traffic drops to nearly 15% of the line speed on BT connections during peak hours. The cause of this can without doubt be attributed to traffic shaping – the practice of prioritising one type of traffic over another. The term "traffic management" is also frequently used by some providers.

Plus Net openly admits and advocates their traffic shaping policies on their website, so these results were to be expected for them. However, BT is not so forthcoming, and the breadth and scale of the activity is rather surprising. In fact, the same characteristics can be seen when looking at any of the BT connections we monitored individually – including business connections and also including two connections that were used *only* by the monitoring units we installed. This suggests that the policy is applied universally, regardless of product and regardless of usage volume.

However, it should be noted that it is this shaping that likely helps their port 80 (HTTP) speedtest results to perform so well. Whilst some may see any form of traffic shaping or traffic management as a bad thing, if you are not a peer-to-peer user or a heavy downloader then BT's and PlusNet's practices will actually benefit you. Future work will examine how these policies affect other "interactive" applications, such as SSH, VoIP and video streaming.

It may also surprise some to note that certain other ISPs are not demonstrating similar dips for non port 80 traffic. The possible reasons for this are numerous:

- They could be traffic shaping based upon volume rather than traffic type (as Virgin Media do);
- Their traffic analysers could be more intelligent or configured differently to those of BT's and recognise that this was not real peer to peer traffic, and thus not shape it;
- They could not be employing traffic shaping at all – the equipment required to do this *properly* is very expensive.

Summary data tables for non port 80 download speed tests

ISP	Percentage of implied line speed
Be Unlimited	77.48
BT	48.13
Entanet	77.98
Karoo	76.84
Orange	75.64
Plus Net	54.45
Sky	66.99
TalkTalk	73.96
Tiscali	73.53
Virgin Media	65.04
Virgin.net	54.05
Zen Internet	81.40

Table 14 – Non port 80 downstream throughput (as a percentage of implied line speed)

ISP	Throughput (Mbps)
Be Unlimited	6.29
BT	2.12
Entanet	3.99
Karoo	3.60
Orange	4.05
Plus Net	2.43
Sky	3.57
TalkTalk	2.65
Tiscali	2.63
Virgin Media	5.92
Virgin.net	2.15
Zen Internet	3.70

Table 15 – Non port 80 downstream throughput (raw Mbps)

3.6.4. Speed tests using multiple connections (threads)

So called "download accelerators" have been available on the Internet for many years now. Whilst their techniques may vary, nearly all of them rely on a characteristic of TCP/IP itself to increase performance.

In a nutshell, TCP/IP treats every stream equally. So, if person A and person B are on a shared network link and person A opens up 100 simultaneous connections whilst person B has only one running, person A will be allocated 100 times the bandwidth through the link. The issue was well covered earlier this year George Ou [6] on ZDNet.

Peer to peer applications by their very nature open up many connections, which produces the same net result – you use more than an equal share of resources.

Some ISPs also suggest that in order to receive maximum throughput it is necessary to run multiple connections, and they were keen for our testing to reflect the maximum throughput by doing precisely this. (Note: The accuracy of this statement is questionable at best – in a perfect network there is nothing preventing a TCP/IP stream from consuming the entire link).

But is the speed achieved whilst downloading using multiple connections really representative of the speed you will receive? Ultimately it depends entirely upon the applications in use. For this reason we have included this section, which compares the results of download speed tests using a single connection and the same download but using multiple connections (three).

Again the results have been presented as a comparison between the ISPs' own results, and not with each other. The red line represents the tests running with multiple connections, whilst the orange line represents the single connection tests.

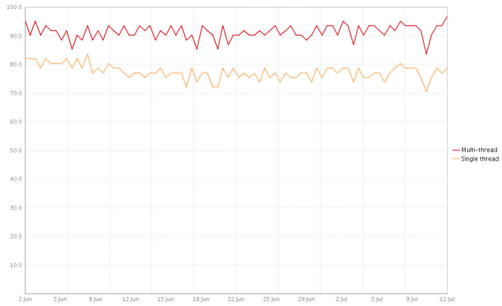

Figure 32 - Be Unlimited traffic

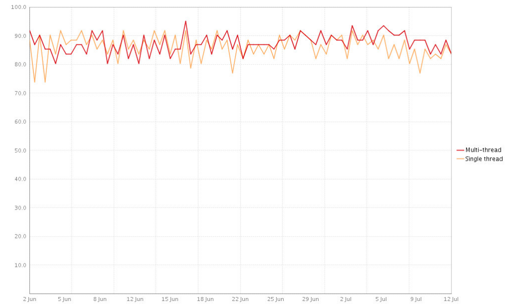

Figure 33 - BT traffic

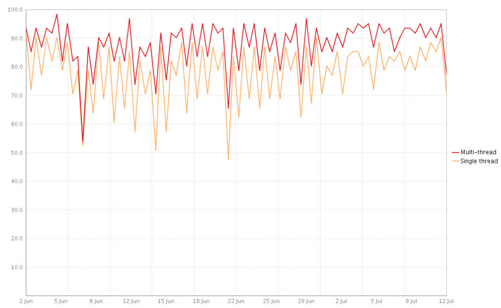

Figure 34 - Entanet traffic

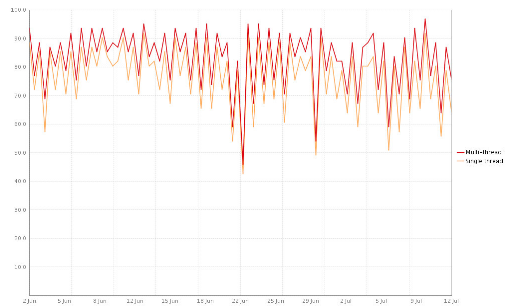

Figure 35 - Karoo traffic

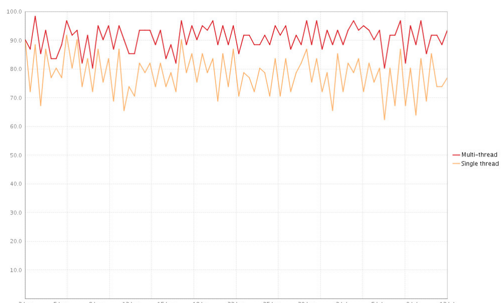

Figure 36 - Orange traffic

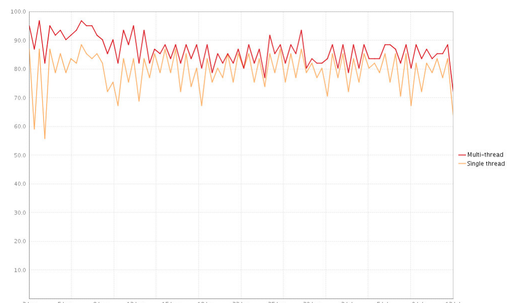

Figure 37 - Plus Net traffic

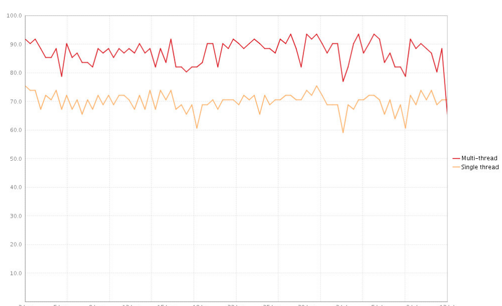

Figure 38 - Sky traffic

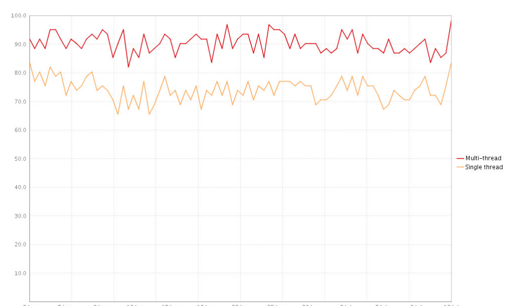

Figure 39 - TalkTalk traffic

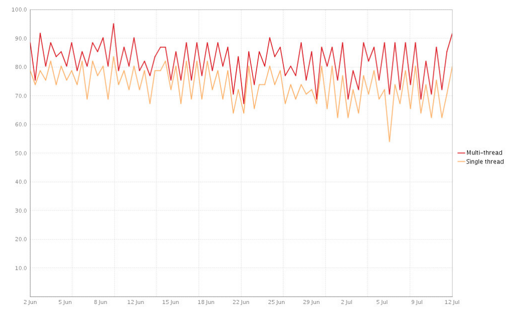

Figure 40 - Tiscali traffic

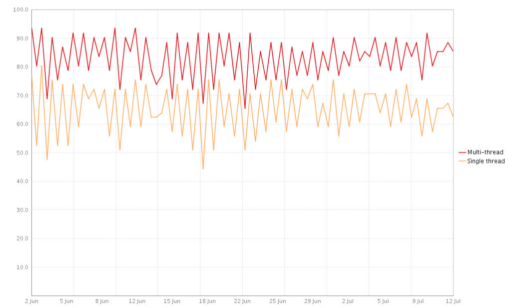

Figure 41 - Virgin Media traffic

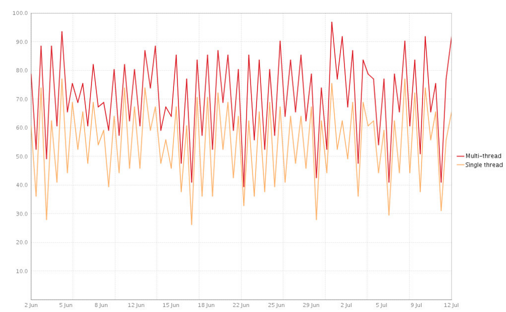

Figure 42 - Virgin.net (ADSL) traffic

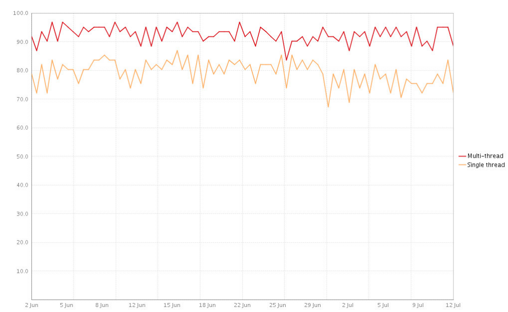

Figure 43 - Zen Internet traffic

The results are interesting if nothing else. Some providers demonstrate a large increase when utilising multiple connections (e.g. Virgin Media and Sky), whereas others show almost no change (e.g. BT).

Whilst it is tempting to immediately blame network contention for the sudden increase obtained when using multiple connections, this is unlikely to be the root cause in all cases. The simple fact is that some variance is to be expected and many other factors could just as easily be the cause of larger variance, including:

- Client-side modem/router architecture;
- Virtual Path contention at the exchange (if on a BT Wholesale ADSL based product);
- Multiple streams can be more easily load balanced between busy routes;
- Traffic shaping (or lack of), although there is no evidence presented here that traffic shaping is being applied to multiple traffic streams.

3.6.5. Variation in implied line speeds

As discussed earlier in this section, nearly all of the speed test results have been displayed as a percentage of the *implied line speed* – a speed that the system recalculates on a daily basis based upon the maximum speed achieved in the past two days.

In order to demonstrate that the implied line speeds were largely consistent across the sample, we have included Figure 44 below.

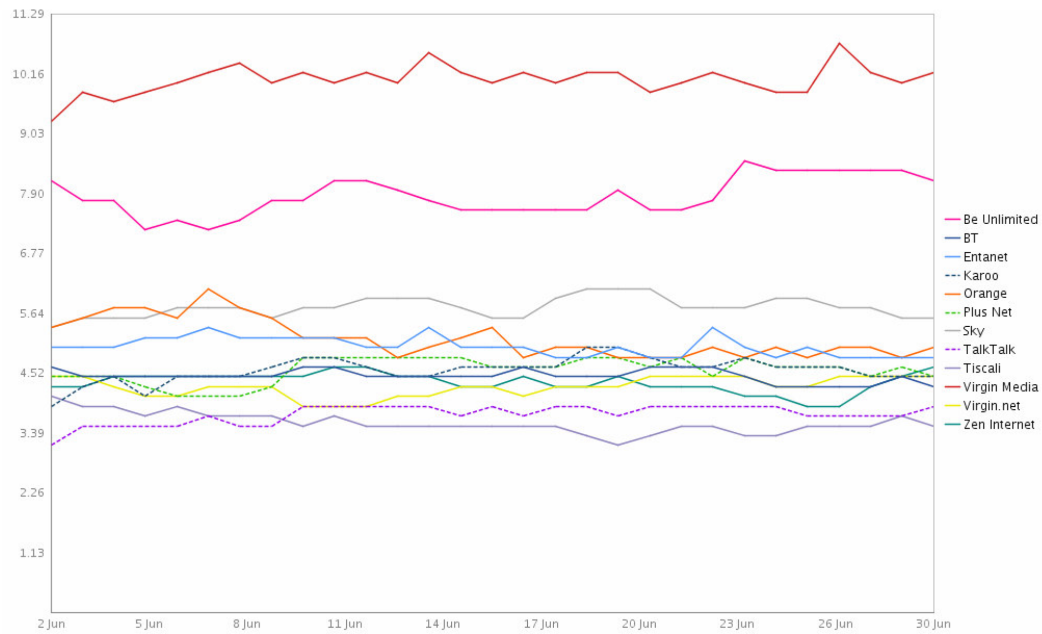

Figure 44 - Implied line speed by ISP

Whilst we have been warning off using raw throughput as a fair metric for comparison in this document, it is still an interesting value to look at. Of particular note here is Be's and Virgin's relatively very high implied line speeds.

We noted earlier in the speed test results that Virgin's cable speeds fluctuated significantly during peak hours, whilst BT's were very stable. But is there more to this?

With an average implied line speed of 10.55Mbps Virgin Media would have to handle 9.5Mbps on each line consistently to reach a 90% throughput speed. BT, for example, had an implied line speed of 4.67Mbps, meaning that they need only handle 4.2Mbps to hit 90%.

Whilst many of the ADSL providers appear to do better in the speed tests than Virgin's cable product, the fact that Virgin have to handle far more traffic to hit the same percentage mark cannot be simply ignored.

4 Future work

There are many possible avenues of improvement that we are keen to explore.

In order for us to be able to draw firmer conclusions about some tentative results (which have not been presented here) and to begin to regionalise results we need to increase our sample size significantly. Regionalising the results is of particular interest to us, as it will allow us to begin to show areas in which some providers perform brilliantly and some in which they do not. With the boom in LLU (where the LLU operators have to provide their own backhaul) we suspect that significant differences in regional results can be seen. To date, no formal study of regional differences in broadband performance has been carried out.

Improving and expanding the suite of tests is also high on the to-do list. Clearly the SMTP test, which is still in its infancy, needs to be expanded to include *all* of the ISPs being tested. Improvements to the non port 80 speed tests are also being planned. These will likely involve rewriting/obfuscation of HTTP headers to avoid detection by deep packet inspecting traffic shapers (this will then appear as bulk binary traffic).

Some have requested that we include BitTorrent-specific tests in the suite. If this were to be conducted it would have to operate within a controlled environment (i.e. communicate with a managed tracker with a fixed number of high bandwidth seeds). Initial research suggests that the operation of the protocol would likely make results difficult to compare. One possible solution is to record the two sides of a BitTorrent conversation and use an application such as *tcpreplay* to repeat the stream between client and server. Whether this is a true test of BitTorrent (which makes heavy use of multiple connections) is questionable.

Finally, we are acutely aware of next generation broadband services on the horizon. Virgin Media are launching their 50Mbps product later this year (or early next), and H2O have their 100Mbps products on the horizon too. The recent announcement from BT regarding their £1.5bn fibre-optic investment has not gone unnoticed either. Active development is underway on the next generation monitoring unit, which will be capable of accurately testing all of these new products.

5 Conclusion

The results detailed in this report have given a sound statistical grounding to some beliefs, dispelled others, and unearthed some interesting facts too.

Observant readers will notice that only a few of the ISPs tested were discussed extensively in the results. Virgin, BT, Be, Sky and Zen all received focus in the text surrounding the results, but the other seven ISPs went largely unmentioned. The fact is that the other ISPs all performed around average – their results were neither excellent nor were they poor, but instead "middle of the road".

This may surprise some - particularly if you are a frequent reader of the various broadband related forums online. Users in Hull, where there is no choice for fixed-line broadband (with Karoo being the only option), were particularly keen to have their provider represented here. Many of these users were convinced their provider would stack up poorly against the competition, so the fact they are on-par should be the source of some comfort that they are no worse off than users in other areas.

But there were many observable traits, and it's worth highlighting the most prominent of these again here:

- Zen excelled in terms of reliability, with the fewest recorded failures across all tests;
- Be Unlimited and Virgin Media's cable service shone in the latency tests, albeit for different reasons. Be Unlimited accept orders from customers with lines of 5km or less, so are more likely to have a faster connection, whilst Virgin Media have their cable network which is not dependent on distance from the exchange, so latency is nearly constant;
- Sky suffered somewhat in the latency tests, most likely due to interleaving being enabled on lines by default (customers can request for it to be removed);
- Be Unlimited fell down on the DNS and web page tests due to their ongoing DNS resolution issues;
- Virgin.net, the ADSL product from Virgin Media, performed rather poorly in the majority of tests. Highly variable latency that bottomed out during off-peak hours suggest significant contention as being the cause;
- All providers offered a sufficiently reliable network for carrying VoIP traffic;
- BT's port 80 (world wide web) speeds were excellent and could not be topped (when measured as a percentage of implied line speed);
- Be and Virgin Media's cable service offered the fastest raw throughput speeds in Mbps, largely owing to the technologies used. Significant variation in speed was demonstrated on Virgin's 20Mbps products, whilst other products performed relatively consistently;
- BT and PlusNet were both demonstrated to use traffic shaping at peak times on non port 80 traffic, although, and perhaps surprisingly, no other ISPs demonstrated such usage. Future work will be conducted in this area to see how this affects other applications;
- A clear disparity can be seen between running speed tests with a single stream and running them with multiple streams. When testing your own connection, be sure you are testing on a level playing field.

Future work will now focus on developing the solution further (i.e. improved tests and faster testing units for next generation connections) and increasing the sample size. The sample size is the key though, and we plan to expand this to over 2000 units by Spring 2009. This will provide approximately 200 units per ISP, which should allow us to begin comparing an individual provider's products and regional variations in performance.

6 References

1. **Linksys WRT54GL wireless router**
http://www.linksys.com/servlet/Satellite?c=L_Product_C2&childpagename=US%2FLayout&cid=1133202177241&pagename=Linksys%2FCommon%2FVisitorWrapper

2. **FreeWRT**
http://www.freewrt.org

3. **Why is the first ping lost?**
http://blog.ioshints.info/2007/04/why-is-first-ping-lost.html

4. **G.711 voice codec**
http://en.wikipedia.org/wiki/G.711

5. **BT Central Plus product description**
http://www.btwholesale.com/pages/static/Products/Internet/BT_IPstream/BT_central_plus.html

6. **Fixing the unfairness of the TCP congestion protocol**
http://blogs.zdnet.com/Ou/?p=1078&page=1

