



Test Methodology White Paper

Author: SamKnows Limited

December 6, 2011

Contents

1	INTRODUCTION	3
2	THE ARCHITECTURE	5
2.1	Whiteboxes	5
2.2	Firmware Integration	5
2.3	Deployment	5
2.4	Operation	6
2.5	Communications	6
2.6	Software updates	7
2.7	Kernel parameters	7
3	PERFORMANCE TESTS	8
3.1	Web browsing	8
3.2	Video streaming	8
3.3	Voice over IP	9
3.4	Availability Test	9
3.5	UDP Latency and Packet Loss	10
3.6	Data Usage Test	10
3.7	Speed Tests	10
3.8	ICMP Latency and Packet Loss	11
3.9	ICMP Latency Under Load	12
3.10	DNS resolution	12
3.11	Peer-to-peer	12
3.12	FTP Transfer	12
3.13	Email Relaying	13
3.14	Threshold Manager	13
4	TYPICAL HOURLY CYCLE	14
5	TEST NODES	15
5.1	On-net and Off-net Nodes	15
5.2	Selecting a Test Node	15
5.3	Hardware Specification	15
6	BACKEND SERVICES	16
6.1	Data Collection Service	16
6.2	Data Processing	16
6.3	Data Presentation	16
7	TESTING SCHEDULES	17
7.1	FCC testing schedule (US)	17
7.2	OFCOM testing schedule (UK)	19
7.3	IDA testing schedule (Singapore)	20

Introduction

IMPORTANT NOTICE

This document is confidential and is not for public release. Its intended audience are interested parties with regards to SamKnows broadband performance measurement project.

SamKnows Whiteboxes - consumer grade, home WiFi routers with additional testing software integrated are to be deployed onto the home network of each panellist to test and report on a range of key metrics as outlined in the table below.

Metric	Primary measure(s)
Web browsing	The total time taken to fetch a page and all of its resources from a popular website
Video streaming	The initial time to buffer, the number of buffer underruns and the total time for buffer delays
Voice over IP	Upstream packet loss, downstream packet loss, upstream jitter, downstream jitter, round trip latency
Download speed	Throughput in Megabits per second utilising three concurrent TCP connections
Upload speed	Throughput in Megabits per second utilising three concurrent TCP connections
UDP latency	Average round trip time of a series of randomly transmitted UDP packets
UDP packet loss	Percentage of UDP packets lost from latency test
Consumption	Volume of data downloaded and uploaded by the panellist
Availability	The total time the connection was deemed unavailable
DNS resolution	The time taken for the ISP's recursive DNS resolver to return an A record for a popular website domain name
ICMP latency	The round trip time of five regularly spaced and schedule ICMP packets
ICMP packet loss	The percentage of packets lost in the ICMP latency test
FTP throughput	Throughput in Megabits per second at which a file can be downloaded from or uploaded to an FTP server
Peer-to-peer	Throughput in Megabits per second at which a file can be downloaded

from a peer-to-peer network

Email Relaying

The time taken to relay an email via the ISP's SMTP servers and reach a target mail server

2 The Architecture

The SamKnows Whiteboxes on panellists' home networks execute a series of software tests over their broadband Internet connection. The results of these tests are reported securely up to hosted backend infrastructure.

The majority of tests run against our network of test nodes. These are dedicated servers either on-net (on the local ISP's network) or off-net (on the public Internet). Other tests will execute against real applications hosted on the Internet, mimicking their behaviour and measuring key performance variables.

When a testing cycle has completed, the results are encrypted and transmitted over SSL to hosted backend infrastructure for processing and presentation through a web interface to each panellist and other interested parties.

2.1 Whiteboxes

SamKnows employs a range of different Whiteboxes and is continuously developing support for new hardware. The following is a list of the hardware we currently deploy to consumers:

- Netgear WNR3500L
- Buffalo WZR-HP-G300NH
- Technicolor 587nv2
- TP-Link (multiple models)

2.2 Firmware Integration


SamKnows collaborates with hardware manufacturers to build custom firmware images for their devices with the testing suite embedded.

All firmware is designed and built by engineers at SamKnows, and then passed through QA testing with the hardware manufacturer. All devices operate a 2.6 series Linux kernel.


2.3 Deployment

The Whitebox has two modes of operation:

- 1) Operate as a router, replacing the user's existing Ethernet router. All wired and wireless devices should connect through the Whitebox.



- 3) Operate as an Ethernet bridge, co-existing with an existing modem/router. All wired devices should connect through the Whitebox. Wireless devices should continue to connect to their existing router.



In order to determine when it is safe to execute tests, the end user's traffic levels are monitored continuously. This is the reason for connecting all Ethernet devices through the Whitebox (and for connecting wireless devices through it in the first scenario).

In the second scenario above, the Whitebox will passively monitor the strongest wireless access point in the vicinity for traffic.

2.4 Operation

Panellists have no ability to disable, reconfigure or influence the SamKnows software in any way through normal usage.

Panellists are, as part of the terms of service, required to leave their Whitebox and other networking equipment permanently powered on and connected to ensure consistent testing.

2.5 Communications

All communications between the Whitebox and the Data Collection Service on the backend hosted infrastructure are initiated by the Whitebox, encrypted over SSL and subject to authentication

The Whitebox communicates with the target test nodes over a variety of TCP and UDP ports. The Whitebox will also communicate with some unmanaged services over both TCP and UDP. Please see the testing section for more details.

2.6 Software updates

The SamKnows software suite has the ability to auto-update itself, downloading updated binaries and testing schedules from the Data Collection Service and storing locally in RAM or flash.

2.7 Kernel parameters

A range of kernel parameters are employed to improve the performance of the Whitebox on high latency and high throughput links. Note that the Whitebox uses TCP Reno (with the SACK option) as its congestion avoidance algorithm of choice. A listing of the kernel parameters can be found below.

```
net.ipv4.tcp_congestion_control=reno
net.core.rmem_default=65536
net.core.rmem_max=1048576
net.core.wmem_default=32768
net.core.wmem_max=262144
net.ipv4.tcp_rmem="4096 65536 1048576"
net.ipv4.tcp_wmem="4096 32768 262144"
net.core.netdev_max_backlog=2500
net.ipv4.tcp_timestamps=0
net.ipv4.tcp_window_scaling=1
net.ipv4.tcp_sack=1
net.ipv4.tcp_no_metrics_save=1
net.ipv4.tcp_tw_recycle=1
net.ipv4.tcp_fin_timeout=10
```

3 Performance Tests

SamKnows has designed and developed its performance tests in house, ensuring adherence to relevant RFCs. All performance tests are written in C, for performance and portability across a range of hardware platforms.

SamKnows performance tests do not incorporate any third party commercial or free or open source (F/OSS) code. Some tests do however dynamically link to F/OSS libraries.

All times are measured in microseconds.

To provide metrics on the key performance indicators requested, the following performance tests are utilised:

3.1 Web browsing

Measures the time taken to fetch the HTML and referenced resources from a page of a popular website. This test does not test against centralised testing nodes; instead it tests against real websites, ensuring that content distribution networks and other performance enhancing factors may be taken into account.

Each Whitebox will test ten common websites on every test run. The time taken to download the resources, the number of bytes transferred and the calculated rate per second will be recorded. The primary measure for this test is the total time taken to download the HTML page and all associated images, Javascript and stylesheet resources.

The results include the time taken for DNS resolution. The test uses up to eight concurrent TCP connections to fetch resources from targets. The test pools TCP connections and utilises persistent connections where the remote HTTP server supports them.

The test may optionally run with or without HTTP headers advertising cache support (through the inclusion or exclusion of the "Cache-Control: no-cache" request header). The client advertises the user agent of Microsoft Internet Explorer 8.

3.2 Video streaming

This generic streaming test can be configured to model the characteristics of a variety of voice and video protocols. For the purpose of the video streaming test, the intention is to simulate an end user viewing a streaming video over one of the many popular websites that provide this service (e.g. YouTube).

The test operates over TCP and uses a proprietary client and server side component. The client and server negotiate the test parameters at the start of each test.

A three second playout buffer is configured and the client will attempt to download data from the server at the maximum rate necessary to ensure that this buffer is never empty. A separate thread is reading data from this buffer at a fixed rate, looking for buffer underruns (which would

manifest themselves to users as a pause in video). The client will record the time to initial buffer, the total number of buffer underruns and the total delay in milliseconds due to these underruns.

It is expected that the bitrate of the streaming will vary according to the access line speed being tested.

3.3 Voice over IP

This test utilises the same generic streaming test as the video test, albeit with different configuration. The test operates UDP and, unlike the video streaming test, utilises bi-directional traffic.

The client initiates a UDP stream to the server and a fixed-rate stream is tested bidirectionally. A de-jitter buffer of 25ms is used to reduce the impact of jitter. The test measures this disruption by monitoring throughput, jitter, delay and loss. These metrics are measured by subdividing the stream into blocks, and measuring the time taken to receive each block (as well as the difference between consecutive times).

The test uses a 64kbps stream with the same characteristics and properties (i.e. packet sizes, delays, bitrate) as the G.711 codec.

Jitter is calculated using the PDV approach described in section 4.2 of RFC5481. The 99th percentile will be recorded and used in all calculations when deriving the PDV.

3.4 Availability Test

Measures the availability of the network connection from the Whitebox to multiple target test nodes by sending and receiving TCP segments to a receiving server located on each test node.

The client establishes long lived TCP connections to the server on each test node, periodically sending TCP packets containing a timestamp in microseconds.

The server echoes back the same data to the client and if it fails to respond or the connection is reset via TCP RST or FIN then the client will attempt to re-establish the connection. If the client is unable to re-establish the connection to all 3 servers simultaneously, it is inferred that Internet connectivity is at fault, the test records a failure locally, along with a timestamp to record the time of failure.

To aid in diagnosing at which point in the route to the target test nodes the connectivity failed, a traceroute is launched to all target test nodes, the results of which are stored locally until connectivity is resumed and the results can be submitted.

This test is executed when the Whitebox boots and runs permanently as a background test.

3.5 **UDP Latency and Packet Loss**

Measures the round trip time of small UDP packets between the Whitebox and a target test node. Each packet consists of an 8-byte sequence number and an 8-byte timestamp. If a packet is not received back within three seconds of sending, it is treated as lost. The test records the number of packets sent each hour, the average round trip time of these and the total number of packets lost. The test will use the 99th percentile when calculating the summarised minimum, maximum and average results.

As with the availability test, the test operates continuously in the background. It is configured to randomly distribute the sending of the echo requests over a fixed interval, reporting the summarised results once the interval has elapsed.

3.6 **Data Usage Test**

Measures the amount of inbound and outbound data traversing the WAN interface of the Whitebox over a one hour period. Data usage arising from SamKnows testing is taken into account and is reported separately.

3.7 **Speed Tests**

Measures the download and upload speed of the given connection in bits per second by performing multi-connection GET and POST HTTP requests to a target test node.

Binary non-zero content herein referred to as the payload is hosted on a web server on the target test node. The test operates for either a fixed duration (in seconds) or a fixed volume (in MB). It can also output average throughput at multiple intervals during the test (e.g. once every 5 seconds). The client will attempt to download as much of the payload as possible for the duration of the test. The payload and all other testing parameters are configurable and may be subject to change in the future.

Four separate variations on the test are supported:

- Single connection GET
- Multi connection GET
- Single connection POST
- Multi connection POST

Note that SamKnows recommends the usage of the multi connection test for all faster service tiers, and typically uses 3 concurrent connections. Each connection used in the test counts the numbers of bytes of the target payload, transferred between two points in time and calculates the speed of each thread as Bytes transferred / Time (seconds).

Factors such as TCP slow start and congestion are taken into account by repeatedly downloading small chunks (default 256KB) of the target payload before the real testing begins. This “warm up” period is said to have been completed when three consecutive chunks were downloaded at the

same speed (or within a small tolerance (default 10%) of one another). In a multi connection test, three individual connections are established (each on its own thread) and are confirmed as all having completed the warm up period before timing begins.

Content downloaded is output to /dev/null or equivalent (i.e. it is discarded), whilst content uploaded is generated and streamed on the fly from /dev/urandom.

The following is an example of the calculation performed for a multi connection test utilising three concurrent connections.

S = Speed (Bytes per second)
B = Bytes (Bytes transferred)
C = Time (Seconds) (between start time point and end time point)
S1 = B1 / C1 (speed for Thread 1 calculation)
S2 = B2 / C2 (speed for Thread 2 calculation)
S3 = B3 / C3 (speed for Thread 3 calculation)
Speed = S1 + S2 + S3

Example values from a 3MB payload:

B1 = 3077360 C1 = 15.583963
B2 = 2426200 C2 = 15.535768
B3 = 2502120 C3 = 15.536826
S1 = B1/C1 = 197469.668017
S2 = B2/C2 = 156168.655454
S3 = B3/C3 = 161044.475879
S1 + S2 + S3 = Total Throughput of the line = 197469.668017 + 156168.655454 + 161044.475879 =
514682 (Bps) * 0.000008 = 4.12 Mbps

3.8 ICMP Latency and Packet Loss

This test measures the round trip time (RTT) of ICMP echo requests in microseconds from the Whitebox to a target test node. The client sends 5 ICMP echo requests of 56 bytes to the target test node awaiting up to three seconds for a response to each. Packets that are not received in response are treated as lost. The minimum, maximum and standard deviation of successful results are also submitted.

3.9 **ICMP Latency Under Load**

The latency under load test operates for the duration of the 30 second downstream and upstream speed tests. Whilst the speed tests run, the latency under load test sends 20 ICMP echo packets to the target server and measures the round trip time. Packets are spaced 500ms apart, and a 3 second timeout is used.

The test records the mean, minimum and maximum round trip times in microseconds. The number of lost ICMP echo requests is also recorded. The test operates both during the downstream and upstream speed tests, and the results are recorded separately.

This is an early version of the latency under load test and was incorporated following input from MIT. Enhancements to the test (such as making it use UDP datagrams rather than ICMP packets, in order to maintain parity with our existing latency/loss test) will be incorporated into a future version.

3.10 **DNS resolution**

This test measures the DNS resolution time of a selection of common website domain names. These tests will be targeted directly at the ISPs recursive resolvers. A list of appropriate servers will be sought from each ISP in advance of the tests.

3.11 **Peer-to-peer**

The peer-to-peer test emulates the BitTorrent transfer of a 10MB binary file (recommended minimum size). The various test servers around the world act as dedicated seeders and are permanently registered with the BitTorrent tracker, seeding this sample file. Upon startup, the client communicates with the tracker and obtains a list of peers (the seeds), handshakes with each, and then begins transferring the payload. The client is

expressly prohibited from sharing peer-to-peer traffic with other test clients and will only communicate with the super seeders. BitTorrent encryption is not used in this approach.

Metrics that are capture include (1) the average and peak throughput during the transfer, (2) the number of connections established to peers, (3) the total number of pieces transferred and (4) the number of TCP connections that were reset during transfer.

The total throughput would broadly match the results of the downstream speed test in a network without traffic shaping. The presence of TCP connections being reset may also indicate the presence of other forms of traffic management.

3.12 **FTP Transfer**

This test measures the throughput of performing an FTP transfer to/from a designated FTP server. The test can transfer a file of an arbitrary size (although a 10MB minimum is recommended). In

both the download and upload modes the test uses a single TCP connection to perform the transfer.

3.13 Email Relaying

This test measures how long it takes for an email to be sent via the ISP's public SMTP relay and to reach a third-party (SamKnows-hosted) email server. Each email sent by the Whitebox has a timestamp embedded within it of when the email was sent. Once the Whitebox has sent an email, it records the message ID in a temporary file to be checked later. Every 10 minutes after, the Whitebox will connect to the SamKnows mail server and attempt to find each of the message IDs listed from its temporary file. Upon finding a matching email, the timestamp found in the email is subtracted from the current timestamp in order to find the total time for delivery.

Should an email not be found on the mail server after 3 days, the test is marked as having failed.

3.14 Threshold Manager

Prior to and during testing, a threshold manager service monitors the inbound and outbound traffic across the WAN interface of the Whitebox to calculate if a panellist is actively using the Internet connection. The threshold for traffic is set to 64kbps downstream and 32kbps upstream. Additionally, a threshold is also defined on the acceptable load average of the Whitebox, should these thresholds be breached prior to the test starting, the test will be delayed for a minute and the process repeated. Should this occur during a test, the current test will be cancelled and the result discarded. If the connection is being actively used throughout, this pause and retry process will occur up to 5 times before the entire test cycle is abandoned.

4 **Typical Hourly Cycle**

A test cycle on the Whitebox occurs once an hour every hour, 24 hours a day. The timing of the testing is randomised per Whitebox to ensure an even spread across the panel.

A scheduling service on the Whitebox manages the following tasks:

- Execute tests
- Send Test results
- Check the backend service for a new testing schedule
- Check the backend service for updated performance tests

The availability and data usage tests run permanently in the background as soon as the Whitebox has booted.

A test cycle may last up to 15 minutes depending on the performance of the host Internet connection. Once testing is complete, results are securely transmitted to the Data Collection Service on the backend infrastructure.

5 Test Nodes

Whiteboxes target dedicated, bare metal servers configured as the end point for the speed, streaming, VoIP, jitter, latency, packet loss and availability tests.

5.1 On-net and Off-net Nodes

SamKnows maintains a global network of test nodes that the Whiteboxes test against. Many of these are built upon the Measurement Labs infrastructure and their locations can be found at <http://code.google.com/p/ndt/wiki/MLabOperations>. These nodes are said to be “off-net”, as they do not reside directly on any one ISPs network.

ISPs may contribute hardware for the purposes of hosting “on-net” test nodes. These are nodes which are hosted within the ISP’s network. The purpose of these nodes is to allow the ISP to determine what (if any) degradation in performance occurs outside of their network.

5.2 Selecting a Test Node

At startup, Whiteboxes retrieve a list of all active test nodes from the SamKnows infrastructure. The Whitebox then uses a simple series of ICMP pings to measure approximate latency to each. The node with the lowest latency is said to be the “closest” and will be used from that point on.

Whiteboxes will then perform tests against the closest off-net node and the closest on-net node for that ISP (assuming the ISP has provided one).

Should the selected test node become unavailable for an extended period then SamKnows will ask the Whitebox to re-select its closest targets.

5.3 Hardware Specification

Test nodes must meet the following minimum specification:

- Dual-core CPU of 2Ghz
- 4GB RAM
- 20GB disk space
- Gigabit Ethernet connectivity
- Redhat Linux 5.x / CentOS 5.x

6 Backend Services

SamKnows employs a fully managed infrastructure for the purposes of data collection from the Whiteboxes, data processing, data presentation and Whitebox management.

Currently hosted directly in the US and United Kingdom, the backend makes use of dedicated hardware firewalls, load balancers and bare metal hardware.

SamKnows operations oversee the management of the backend infrastructure, adhering to industry standard practices for security and operational management.

The backend can be broken down into 3 distinct areas:

6.1 Data Collection Service

The data collection service or DCS, is the gateway for the Whitebox to communicate with the backend for sending tests results and requesting configuration updates. Communication with the DCS is over TCP 443 with all communications encrypted via SSL.

6.2 Data Processing

A cluster of database servers utilising a specialized column based storage engine to process and store results data. All publicly identifiable information (PII) is encrypted and is only accessible by panellists themselves and SamKnows.

6.3 Data Presentation

Data is made available via a Web 2.0 style reporting interface, accessible over SSL with granular access controls.

7 Testing Schedules

The following schedule provides a breakdown of test durations and indicative impact on monthly bandwidth usage.

7.1 FCC testing schedule (US)

Test	Target(s)	Frequency	Duration	Est. Daily Volume
Web browsing	10 popular US websites	Hourly, 24x7	Est. 30 seconds	80MB
Video streaming*	1 off-net test node	Hourly, 24x7	Fixed 10 seconds at 768kbps, 1.25Mbps, 2.25Mbps, 3.75Mbps	60MB
	1 on-net test node	Hourly, 24x7	Fixed 10 seconds at 768kbps, 1.25Mbps, 2.25Mbps, 3.75Mbps	60MB
Voice over IP	1 on-net test node	Hourly, 24x7	Fixed 30 seconds at 64k	1MB
	1 off-net test node	Hourly, 24x7	Fixed 30 seconds at 64k	1MB
Download speed**	1 off-net test node	Every other hour, 24x7	Fixed 30 seconds***	4.5GB at 50Mbps 1.72GB at 20Mbps 858MB at 10Mbps 357MB at 3Mbps 129MB at 1.5Mbps
	1 on-net test node	Every other hour, 24x7	Fixed 30 seconds***	4.5GB at 50Mbps 1.72GB at 20Mbps 858MB at 10Mbps 357MB at 3Mbps 129MB at 1.5Mbps
Upload speed**	1 off-net test node	Every other hour, 24x7	Fixed 30 seconds***	174MB at 2Mbps 87MB at 1Mbps 44MB at 0.5Mbps
	1 on-net test node	Every other hour, 24x7	Fixed 30 seconds***	174MB at 2Mbps 87MB at 1Mbps 44MB at 0.5Mbps
UDP latency	1 off-net test node	Hourly, 24x7	Permanent	1MB
	1 on-net test node	Hourly, 24x7	Permanent	1MB
UDP packet loss	1 off-net test node	Hourly, 24x7	Permanent	N/A (uses above)

	1 on-net test node	Hourly, 24x7	Permanent	N/A (uses above)
Consumption	N/A	24x7	N/A	N/A
Availability	3 off-net test nodes	24x7	Permanent	1MB
DNS resolution	10 popular US websites	Hourly, 24x7	Est. 3 seconds	0.3MB
ICMP latency	1 off-net test node 1 on-net test node	Hourly, 24x7	Est. 5 seconds	0.3MB
ICMP packet loss	1 off-net test node 1 on-net test node	Hourly, 24x7	N/A (As ICMP latency)	N/A (uses above)

* Video streaming rates: Lines will only stream the rates they are capable of, according to the latest speed test results. If a rate is deemed unreachable (e.g. a 3.75Mbps rate on a 1Mbps line), then it will be skipped.

** Download/upload daily volumes are estimates based upon likely line speeds. All tests will operate at maximum line rate so actual consumption may vary.

*** The speed tests run for a fixed duration of 30 seconds, but output the cumulative average transfer speed every 5 seconds (i.e. the first average is for seconds 0-5, the second is for 0-10, etc). This allows us to measure performance over multiple durations, whilst running only one test.

7.2 OFCOM testing schedule (UK)

Test	Target(s)	Frequency	Duration	Est. Daily Volume
Web browsing	OFCEM test website	Hourly, 24x7	Est. 3 seconds	8.4MB
Voice over IP	1 on-net test node	Every other hour, 24x7	Fixed 10 seconds at 64k	1.92MB
Download speed*	1 off-net test node	Once 12am-6pm Once 6am-12pm Once 12pm-6pm Every hour 6pm-11pm	6MB fixed size	67.5MB
Upload speed*	1 off-net test node	Once 12am-6pm Once 6am-12pm Once 12pm-6pm Once 6pm-12pm	3MB fixed size	12MB
UDP latency	1 off-net test node	Hourly, 24x7	Permanent	1MB
UDP packet loss	1 off-net test node	Hourly, 24x7	Permanent	N/A (uses above)
DNS resolution	3 popular UK websites	Hourly, 24x7	Est. 1 seconds	0.1MB
ICMP latency	3 popular UK websites	Hourly, 24x7	Est. 5 seconds	0.1MB
ICMP packet loss	3 popular UK websites	Hourly, 24x7	N/A (As ICMP latency)	N/A (uses above)

* Download/upload daily volumes are estimates based upon likely line speeds. All tests will operate at maximum line rate so actual consumption may vary.

7.3 IDA testing schedule (Singapore)

Test	Target(s)	Frequency	Duration	Est. Daily Volume
Web browsing	5 popular websites IDA website	Hourly, 24x7	Est. 30 seconds	40MB
Video streaming	SingTel, StarHub, M1, US West coast, US East coast, UK, Shanghai	Every six hours	Fixed 10 seconds at 2.25Mbps	300MB
Voice over IP	SingTel, StarHub, M1, US West coast, US East coast, UK, Shanghai	Every six hours	Fixed 30 seconds at 64k	5MB
Download speed (single connection)	SingTel, StarHub, M1, US West coast, US East coast, UK, Shanghai	Every six hours	Fixed 10 seconds	7GB at 200Mbps 1.75GB at 50Mbps 70MB at 2Mbps
Download speed (multi connection)	SingTel, StarHub, M1, US West coast, US East coast, UK, Shanghai	Every six hours	Fixed 10 seconds	7GB at 200Mbps 1.75GB at 50Mbps 70MB at 2Mbps
Upload speed (single connection)	SingTel, StarHub, M1, US West coast, US East coast, UK, Shanghai	Every six hours	Fixed 10 seconds	7GB at 200Mbps 1.75GB at 50Mbps 70MB at 2Mbps
Upload speed (multi connection)	SingTel, StarHub, M1, US West coast, US East coast, UK, Shanghai	Every six hours	Fixed 10 seconds	7GB at 200Mbps 1.75GB at 50Mbps 70MB at 2Mbps
UDP latency	SingTel, StarHub, M1, US West coast, US East coast, UK, Shanghai	Hourly, 24x7	Permanent	2MB
UDP packet loss	SingTel, StarHub, M1, US West coast, US East coast, UK, Shanghai	Hourly, 24x7	Permanent	N/A (uses above)
DNS resolution	SingTel, StarHub, M1, US West coast, US East coast, UK, Shanghai	Hourly, 24x7	Est. 3 seconds	0.3MB
ICMP latency	SingTel, StarHub, M1, US West coast, US East coast, UK, Shanghai	Hourly, 24x7	Est. 5 seconds	0.3MB
ICMP packet loss	SingTel, StarHub, M1, US West coast, US East	Hourly, 24x7	N/A (As ICMP latency)	N/A (uses above)

	coast, UK, Shanghai			
Peer-to-Peer	SingTel, StarHub, M1, US West coast, US East coast, UK, Shanghai	Every six hours, 24x7	Fixed 10MB transfer	300MB
FTP download	SingTel, StarHub, M1, US West coast, US East coast, UK, Shanghai	Every six hours, 24x7	Fixed 10MB transfer	300MB
Email round-trip	ISP SMTP server	Every six hours, 24x7	Est. 3 seconds	0.3MB

Note: The five websites are facebook.com, google.com.sg, youtube.com, yahoo.com and wikipedia.org

Note: The “Best SG test node” may refer to the SingTel, StarHub or M1 test node, depending on the ISP that the Whitebox is connected to.

Note: Assuming a sample base of 50 units per ISP/product, if we run tests once every 6 hours on each unit and have a normal distribution of units to hours, then we have approximately 8 samples per hour per ISP/product combination.