



SAMKNOWS TEST METHODOLOGY
Methodology and technical information
relating to the SamKnows testing platform

February 2013 Report

Contents

1	ABOUT THIS DOCUMENT	3
1.1	Limitation of Liability	3
1.2	Copyright	3
2	INTRODUCTION	4
3	PERFORMANCE TESTS	5
3.1	Web Browsing	5
3.2	Video Streaming	6
3.3	Voice Over IP (VOIP)	6
3.4	Speed Tests	7
3.5	UDP Latency and Packet Loss	8
3.6	DNS Resolution	8
3.7	Latency under load	9
3.8	Threshold Manager	9
3.9	Testing Schedule	10

1 **About this Document**

1.1 **Limitation of Liability**

The information contained in this document is provided for general information purposes only. Whilst care has been taken in compiling the information herein, SamKnows does not warrant or represent that this information is free from errors or omissions. To the maximum extent permitted by law, SamKnows accepts no responsibility in respect of this document and any loss or damage suffered or incurred by a person for any reason relying on the any of the information provided in this document and for acting, or failing to act, on any information contained on or referred to in this document.

1.2 **Copyright**

The material in this document is protected by Copyright. No part of the materials in this document may be reproduced for any purpose whatsoever without the written permission of SamKnows.

Introduction

SamKnows Whiteboxes – modified routers acting as an ethernet bridge 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
DNS resolution	The time taken for the ISP's recursive DNS resolver to return an A record for a popular website domain name
Latency under load	The average round trip time for a series of regularly spaced UDP packets sent during downstream/upstream sustained tests
Consumption	A simple record of the total bytes downloaded and uploaded by the router

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 durations are measured in microseconds and all rates in bytes per second.

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 (VOIP)**

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 **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

3.5 **UDP Latency and Packet Loss**

Measures the round trip time of small UDP packets between the Whitebox and a target test node. Each packet contains 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 **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.

3.7 **Latency under load**

The latency under load test operated for the duration of the 30 second downstream and upstream speed tests, with results for upstream and downstream recorded separately. While the speed tests ran, the latency under load test sent UDP datagrams to the target server and measured the round trip time and number of packets lost. Packets were spaced 500 milliseconds (ms) apart, and a 3 second timeout was used. The test recorded the mean, minimum and maximum round trip times in microseconds. The number of lost UDP packets was also recorded.

This test represents an updated version of the methodology used in the August 2011 Report and brings it into line with the methodology for the regular latency and packet loss metrics.

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

3.8 **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. 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.

3.9 Testing Schedule

Test	Target(s)	Duration	Total Est. Daily Volume
Web browsing	9 popular US websites	Est. 30 seconds	80 MB
Video streaming ¹	1 off-net test node	Fixed 10 seconds at 768kbps, 1.25Mbps, 2.25Mbps, 3.75Mbps	60 MB
	1 on-net test node	Fixed 10 seconds at 768kbps, 1.25Mbps, 2.25Mbps, 3.75Mbps	60 MB
Voice over IP	1 on-net test node	Fixed 30 seconds at 64k	1 MB
	1 off-net test node	Fixed 30 seconds at 64k	1 MB
Download speed ^{2,3}	1 off-net test node	Fixed 30 seconds***	4.5 GB at 50Mbps 1.72 GB at 20Mbps 858 MB at 10Mbps 357 MB at 3Mbps 129 MB at 1.5Mbps
	1 on-net test node	Fixed 30 seconds***	4.5 GB at 50Mbps 1.72 GB at 20Mbps 858 MB at 10Mbps 357 MB at 3Mbps 129 MB at 1.5Mbps
Upload speed ^{2,3}	1 off-net test node	Fixed 30 seconds***	174 MB at 2Mbps 87 MB at 1Mbps 44 MB at 0.5Mbps
	1 on-net test node	Fixed 30 seconds***	174 MB at 2Mbps 87 MB at 1Mbps 44 MB at 0.5Mbps
UDP latency	1 off-net test node	Permanent	1 MB
	1 on-net test node	Permanent	1 MB
UDP packet loss	1 off-net test node	Permanent	N/A (uses above)
	1 on-net test node	Permanent	N/A (uses above)
DNS resolution	9 popular US websites	Est. 3 seconds	0.3 MB
DNS failures	9 popular US websites	N/A (As DNS resolution)	
ICMP latency	1 off-net test node	Est. 5 seconds	.3MB
	1 on-net test node	Est. 5 seconds	.3MB
Latency under Load	1 off-net test node	Est. 5 seconds	.3MB
	1 on-net test node	Est. 5 seconds	.3MB
UDP packet loss	1 off-net test node	N/A (As UDP latency)	0.1MB
	1 on-net test node	N/A (As UDP latency)	0.1MB

¹ 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.

² Download/upload daily volumes are estimates based upon likely line speeds. All tests operated at maximum line rate so actual consumption may vary.

³ Frequency: twice every hour. All other tests report results once per hour.

[DOCUMENT ENDS]