



DOCUMENT REFERENCE:
SQ309-002-EN

SAMKNOWS TEST METHODOLOGY
Web-based
Broadband Performance
White Paper

July 2015

SAMKNOWS QUALITY CONTROLLED DOCUMENT.					
SQ	REV	LANG	STATUS	OWNER	DATED
309	03	EN	FINAL	SC	20150702
REVISION HISTORY					
DATED	REV	AUTHOR	COMMENTS		
20150702	003	LR	General Updates.		
20140106	002	KR	General Updates		
20120411	001	SC	Original.		

Contents

1	IMPORTANT NOTICE	3
2	INTRODUCTION	4
3	THE ARCHITECTURE	5
3.1	Determining the best measurement server	5
3.2	Browser compatibility	6
3.3	Performance	7
3.4	Capturing location and ISP data	7
3.5	Integrating the application into your web site	7
3.6	Capturing test results	8
3.7	Software updates	8
4	PERFORMANCE TESTS	9
4.1	Speed Tests	9
4.2	UDP Latency and Packet Loss	10
5	TEST NODES	11
5.1	On-net and Off-net Nodes	11
5.2	Hardware Specification	11
6	BACKEND SERVICES	12
6.1	Data Collection Service	12
6.2	Data Processing	12
6.3	Data Presentation	12

1 Important Notice

IMPORTANT NOTICE

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.

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.

2 Introduction

SamKnows produces the leading solution for broadband performance measurement and is used by regulators and ISPs around the world. This hardware-based approach delivers unrivalled detail and accuracy, as it is unaffected by factors within the home (such as the PC or home network).

However, there are some situations where this level of certainty over the results is less critical and the focus is instead upon providing instantaneous data to end users. In such situations, deploying a SamKnows Whitebox (a hardware measurement probe) is overkill.

To address this gap SamKnows have developed a web-based application for measuring selected broadband performance metrics. The measurements have been built to the same specification as those used in the SamKnows Whiteboxes, but are designed to run in a web browser on a user's workstation.

Unlike other web-based speed tests, the SamKnows web-based uses pure HTML5 and therefore does not rely on Java or Flash. This allows it to operate in all modern desktop and mobile browsers, as well as many embedded browsers (e.g. in set-top boxes or TVs).

The following key metrics are included in the software-based measurements:

Metric	Primary Measure(s)
Download speed	Throughput in Megabits per second utilising one or more concurrent TCP connections
Upload speed	Throughput in Megabits per second utilising one or more concurrent TCP connections
Latency	Round trip time to the measurement server

3 The Architecture

The measurement application is delivered as a web application via a web browser. In its simplest use case, end users will initiate the measurement process by clicking on the 'Start' button. A short while later the measurement results will be reported to the end user.

A client embedding the software application in their own web page can select specific tests, customise the testing parameters (e.g. the test server list to use), and completely rebrand the user interface.

3.1 Determining the best measurement server

Upon start-up, the application runs a brief latency measurement to a list of nearby test servers. This process allows us to determine the nearest measurement server (in terms of network distance). The measurement server with the lowest round-trip latency is selected as the target for all subsequent measurements (download speed, upload speed and latency).

The list of test servers may be customised by a client embedding the test inside their web site.

3.2 Browser compatibility

The SamKnows web-based speed test relies only upon HTML5 support for optimum performance, and is even operates a compatibility layer for older browsers. Under no circumstances is Java or Flash required.

Browsers fully supporting HTML5 will utilise WebSockets as the transport for all measurements. This effectively gives us raw TCP sockets with which to conduct tests. Internet Explorer 10, Chrome, Safari, Firefox, Android and iOS all support HTML5 and WebSockets.

Older browsers, such as Internet Explorer 8 and 9, will fall-back to using AJAX for carrying out measurements if WebSockets are not supported.

Internet Explorer	Firefox	Chrome	Safari	iOS Safari	Android Browser	Chrome for Android
		31				
		36				
		37				
		39			4.1	
8	31	40			4.3	
9		41	7		4.4	
10	37	42	7.1	7.1	4.4.4	
11	38	43	8	8.3	40	42
Edge	39	44	9	9		
	40	45				
	41+	46+				
Supported via AJAX fallback						
Supported						

Image 1: Compatibility Table

3.3 Performance

SamKnows has tested the web-based speed test on fibre-to-the-premises connections at 1Gbps and has found it to return accurate results. These results were obtained in all major browsers on a 2012-era PC.

Extremely low-speed connections are also catered for, and SamKnows has tested lines as slow as 128kbps.

3.4 Capturing location and ISP data

The server side examines the IP address of the client and utilises geo-location databases such as MaxMind to find the location and ISP of the user. The physical location is typically accurate to city-level and the ISP can be determined with near 100% accuracy.

3.5 Integrating the application into your web site

SamKnows will provide a template HTML page and accompanying resources (Javascript, images, CSS) that can be used as the basis for integration.

The design of the test can be completely customised by the client, although SamKnows provides an attractive, minimalist design by default.

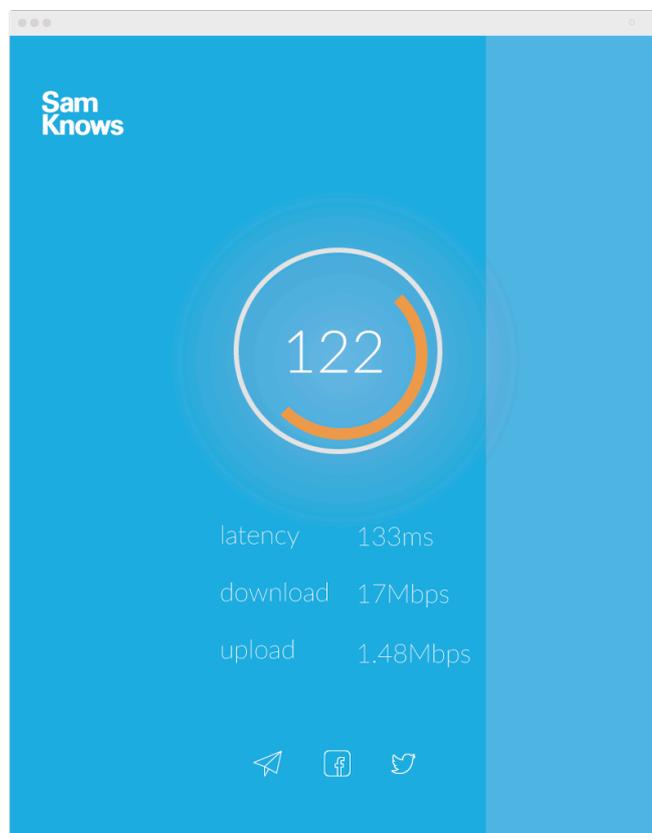


Image 2: SamKnows Web-Browsing Speed Test

If the client wishes, the entire frontend can be replaced. In this scenario, SamKnows is effectively just supplying a Javascript API to the client for them to integrate as they wish. Documentation is available regarding the methods, the callbacks expected, and the data structures in use.

3.6 **Capturing test results**

Once measurements have been completed, the results are submitted securely back to SamKnows's infrastructure for analysis in SamKnows Analytics. All results are transmitted over SSL/TLS, and take the form of a JSON blob. A client may wish to operate a results collection endpoint within their own network too.

3.7 **Software updates**

The software application may be remotely updated by SamKnows to incorporate new features or bug-fixes. No changes are required to the end-user's computer to effect this change.

4 Performance Tests

SamKnows has designed and developed its performance tests in house, ensuring adherence to relevant RFCs. All times are measured in microseconds and all speeds in bytes per second.

4.1 Speed Tests

Measures the download and upload speed of the broadband connection in bits per second. The transfer is conducted over one or more concurrent TCP connections.

If the user's web browser supports WebSockets, then these are used as the transport for all measurement traffic. This is far more efficient than using repeated HTTP requests and allows for extremely high speeds to be reached. If the user's web browser is too old to support WebSockets, then repeated HTTP requests (using the same TCP connection) are made.

In the download speed test the client will repeatedly fetch chunks of data from the target test node. The content is discarded as soon as it is received.

In the upload test the client will generate the payload itself to send to the server. Similarly, the client will repeatedly upload chunks of this data to the target test node, which will immediately discard the content.

The speed tests (both download and upload) operate for a fixed-duration specified in seconds. A maximum limit on transfer volume may be imposed if data volumes are of concern.

Both the download and upload tests will dynamically scale the number of TCP connections in use. By default, six concurrent connections are used. However, this may increase up to 32 concurrent TCP connections on the very fastest broadband lines (such as 1Gbps links).

Factors such as TCP slow start are accounted for through the use of a "warm-up" period. This period begins as soon as the test starts and seeks to remove the impact of TCP slow start from the results and also serves to determine the number of TCP connections to be used for the remainder of the test. It is important to note that the data transferred in the warm-up period is excluded from the main test results.

The speed test client will record the throughput, bytes transferred and time taken at the end of the test.

4.2 **UDP Latency and Packet Loss**

Measures the round-trip time of the smallest possible packet between the user's workstation and the target test node.

If the user's web browser supports WebSockets then these are used as the transport for the latency measurement packets. If the client's browser does not support WebSockets then HTTP HEAD requests are sent via an AJAX call. The WebSockets approach typically delivers much more stable measurements as there are no additional HTTP overheads.

The round trip time is recorded in microseconds. Due to the test operating over TCP (both in the WebSockets and HTTP variants), it is not possible to capture packet loss.

5 Test Nodes

SamKnows test clients, including the Smartphone App, target dedicated, bare metal servers that are configured as the end point for the various performance measurements.

5.1 On-net and Off-net Nodes

SamKnows maintains a global network of test nodes that the measurement application tests 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 ISP’s 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 Hardware Specification

Test nodes must meet the following minimum specification:

- CPU: Dual core Xeon (2GHz+)
- RAM: 4GB
- Disk: 80GB
- Operating System: CentOS/RHEL 6.x
- Connectivity: Gigabit Ethernet connectivity, with gigabit upstream link.

6 Backend Services

SamKnows employs a fully managed infrastructure for the purposes of data collection from the Whiteboxes and software applications, data processing, data presentation and measurement 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 Whiteboxes and software applications 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

Users of the software measurement application will be presented with their results as soon as they have completed the tests.

Additionally, users may register with the website and request that their measurements be recorded together. Users logging in to the website will then be able to see historical measurement results, presented in both tabular and graphical form. They may also download a spreadsheet of all of their measurement results.

[DOCUMENT ENDS]