

Web browsing

Website Performance Test

Supported clients: Whiteboxes, Routers

The SamKnows Website Performance Test uses a full headless web browser to accurately capture the user experience of browsing a website from within a user's network. Since the browser is a graphical program that renders the page step by step as more objects are being received from the network the measurements must be made based on what is observable on the screen.

The goal of the test is to highlight how much waiting time is needed for a human to be able to start accessing the contents of the page. It does not necessarily mean that all objects that are included in the page have visual relevance.

Some objects are invisible to the user (e.g. updating page statistics on a server) while others have less impact on the visuals. Thinking visual page load progress against the screen of the fully rendered page is a useful metric which can be used to define when a page has enough objects visible to bring value to the user.

Besides from knowing the time frame of the page load and its visual implications one can also look at object-specific details. This highlights what is loaded, why the page doesn't load fast enough or why a part of the page is broken.

Nowadays many web sites require access to multiple web servers with different physical locations and there might be a bottleneck in just a small part of the network. Some objects (e.g. HTML, CSS or JavaScript) have other objects as dependencies which further

increases the page load time if they are the bottleneck. Being able to identify where time is spent can be used to make better decisions when improving the network or increasing server performance.

In an ideal world using real user equipment would give the most accurate measurements. However, what is available to us is the broadband router in each home or office. This device is not powerful enough to accurately portray the user equipment.

By introducing an offload server it becomes possible to test the browsing experience from the user's own network. The server runs the actual browser while the broadband router only performs time measurements between its network and the needed web servers. What is measured are DNS lookups and object transfers. The offload server makes sure that the web page with all its objects are loaded exactly according to the measured timings from the user's network.

Since web servers may produce content differently depending on the geographical location of the user the offload server never talks directly to any web server (e.g. to decide which objects to measure). Instead a tunnel is established between the user's router and the offload server so that all traffic is originated from the user's own network.

The loading operation involving the browser makes use of a server side data cache of HTTP data to be able to have full control of the timing playback.

The result is user specific test results. When scaling up the test to cover thousands of user networks a good overall picture of browsing

experience over a large geographical area is achieved. The test can potentially be run in any broadband router which is far more than what would be possible when testing with a selected few real user devices.

By periodically testing a large number of user networks the difference in performance can be seen over time. This is very similar to how traffic congestion is monitored for the road network.

As with the Web Browsing Lite test, the primary input parameter to the SamKnows Website Performance Test is the website to measure against. Additional optional configuration parameters can be supplied for:

- User agent to advertise
- Resolution to render the page at
- DNS resolution strategy (DNS, or DNS-over-HTTPS)

There are three different types of test output from the SamKnows Website Performance Test. The main type is results that are stored in a database to enable queries and visualisation in SamKnows One. It includes top-level information as well as per-object statistics. This is useful to see the big picture among thousands of test results.

The table below shows the metrics used for top-level results that are relevant for data analysis:

- Timestamp of the test
- URL fetched
- Total number of DNS requests made
- Total number of objects transferred
- Total download and upload volume
- Number of successful objects
- Number of failed objects
- Peak number of concurrent HTTP requests
- Name resolution service (DNS or DoH)
- Average DNS lookup time
- Distribution of DNS resolution time
- Average TCP connection time
- Distribution of TCP connection time
- Average SSL connection time

- Distribution of SSL connection time
- Average object transfer time
- Distribution of object transfer time
- Time to first element being visible
- Time to browser engine reported it was finished
- Time to the page being visually complete
- Distribution of the page loading time
- Screen resolution
- User agent

Secondly, the SamKnows Website Performance Test also captures information on a per-object basis during a test session. Given that a web page can often reference dozens or even hundreds of objects, this means that we collect a large amount of data! The following information is captured on a per-object basis:

- Timestamp
- Method (e.g. GET, POST)
- URL
- Transfer size
- Transfer speed
- HTTP version
- HTTP response code
- Server IP
- Request start time
- Request end time
- DNS lookup time
- TCP connect time
- SSL connect time
- Time to first byte

Lastly, the SamKnows Website Performance Test also records a video of the browser's rendering of the page (with a statistics bar overlaid onto it) and a HTTP Archive file (HAR), which can be analysed using common open source tools. We will show you an example of this further down.

Web browsing lite

Supported clients: Whiteboxes, Routers, iOS, and Android

Note: This test has been renamed 'Web Browsing Lite' since the introduction of the Website Performance Test, which incorporates a real headless web browser that renders the page and executes dynamic content, such as javascript.

The Web Browsing Lite test 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, allowing for content distribution networks and other performance enhancing factors to be considered.

Each SamKnows agent will visit one or more popular websites on every test run. A blend of multiple sites is commonly used to help smooth out content changes or performance anomalies observed from any one particular website. Alternatively, the test can be configured to test to a reference web page, hosted on SamKnows servers. This has the advantage of ensuring that the content will never change, but is no longer representative of real users' web browsing experience.

The time taken to download the page and its associated 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. Optionally, the test can capture metrics for each object that is fetched too.

The results include the time taken for DNS resolution. The test can be optionally configured to use a DNS-over-HTTPS resolver for its DNS requests.

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. The client advertises itself as a modern version of Mozilla Firefox by default, but can be configured with any user agent.