Measuring Broadband Australia

Special Report, June 2021

Web Performance Test

Websites: How they work and what can impact the consumer's experience?

Aim

As part of the Measuring Broadband Australia program, this paper provides an analysis of the web performance of two popular reference internet sites.[1]

Why is measuring web browsing performance important?

Web browsing can sometimes be negatively affected by slow and unresponsive webpages. This study aims to uncover the key drivers of website performance and give more clarity to consumers around what happens when browsing websites.

1. All testing was conducted during May 2020. We do not expect there to be material changes to performance that would impact the conclusions of this study were the tests to be repeated.

Contents

How is web browsing performance measured?

The study used the SamKnows Web Performance Test[2] to measure the performance of browsing the web on broadband connections in Australia.

This test captures dozens of metrics, but this study focuses on the following two key metrics:

1. the time taken for the first visual object to load on the webpage; and

2. the time taken for the whole webpage to load.

Both these measures relate directly to the experience of users. Not all objects on a webpage have visual relevance. Some objects on a page are invisible, such as advertising beacons or cookies that track user activity. The time taken for the first visual object to load gives a good insight into how long it takes for the webpage to become useful to a user and is an important metric to track.

The test works by simulating the experience of loading a webpage. This is performed automatically, on an hourly basis, from hundreds of hardware measurement devices (called Whiteboxes) installed in Australian consumers' homes. More detail is provided in the next section which further outlines the testing methodology and key results. All data is recorded by SamKnows Whiteboxes hosted by Measuring Broadband Australia volunteers.

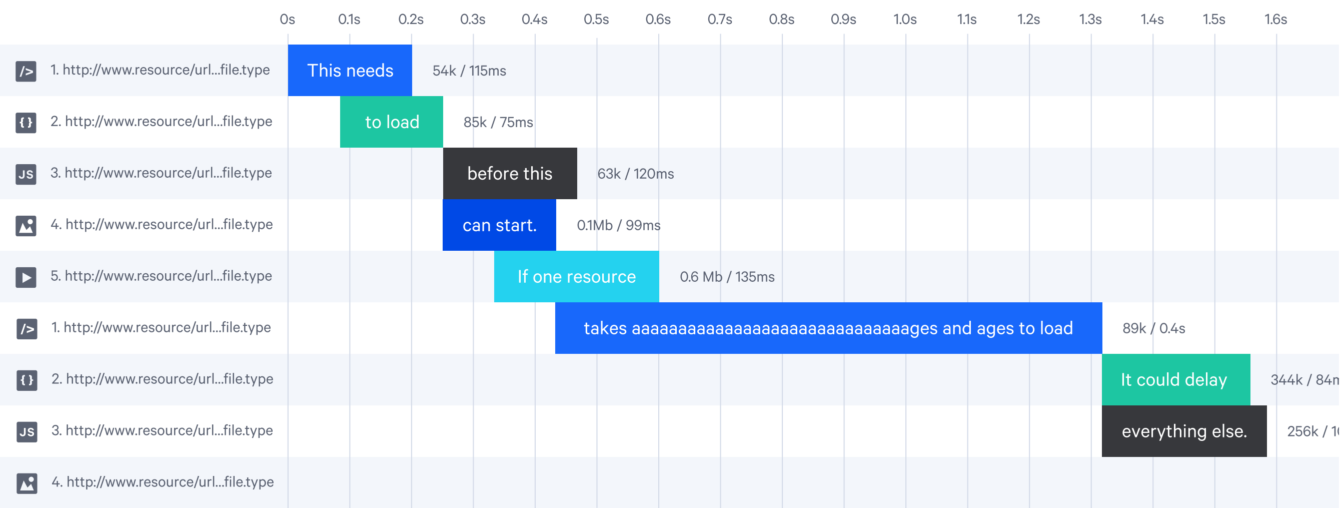

2. https://samknows.com/technology/qoe-tests

Target websites

The two reference websites selected were:

- Website 1 – an online marketplace
- Website 2 – online platform that lists properties for sale and rental

For brevity, throughout the rest of this report the websites will be referred to as Website 1 and Website 2 respectively.

Glossary

Some of the terms discussed might be unfamiliar. This glossary is provided to help explain the acronyms used.

URL – Uniform Resource Location, also called a web address, is a reference to a webpage or file on the internet.

HTML - Hypertext Markup Language, the standard language used for documents to be displayed by a web browser. It provides the basic structure of a webpage and is enhanced with CSS and JavaScript.

CSS – Cascading Style Sheets, a language used to describe the presentation, formatting and layout of a document written in HTML.

JavaScript – A programming language used to control the behaviour of elements in a webpage.

RSP – Retail Service Provider, a company that provides users with their internet subscription.

CDN – Content Delivery Network, a geographically distributed network of servers used to host content.

Objects – the elements on a webpage. This includes things like images, videos, fonts, scripts, CSS and HTML.

Summary of key findings

The study has found that the NBN access speed and provider chosen by the user does not have a material impact on website browsing performance. Both sites reported similar results for all NBN access speeds as demonstrated in the chart below.

Figure 1: Time to first visual and time to visually complete split by NBN speed tier and target website

Table 1: Factors that do not impact web performance

Factor	Explanation
NBN access speed	The study's results show that fibre plans offer reliable web browsing and that the choice of NBN access speed does not have a large impact on web browsing performance. Although higher tiers do record faster web loading times, the increase from a slower plan to a faster plan is small and unlikely to impact the end user. The small size of webpage objects seen on the tested pages are more than capable of being successfully downloaded by each NBN access tier tested. For example, moving from the 50/20 Mbps tier to the 100/40 Mbps tier would lead to almost no noticeable improvement in web performance. This finding suggests for consumers whose main activity is web browsing that plan speed need not be a core consideration when selecting an internet plan.
Choice of RSP	Results for all the major RSPs tested are very satisfactory and suggest there are immaterial differences in performance between providers. All the RSPs studied offer good web performance which suggest that choice of RSP is not likely to have a significant impact on web browsing performance.

This finding means that other factors outside of a consumer's control, such as the design and operability of a website, will have a greater impact on how a website performs and the experience of interacting with it.

For more in-depth analysis of the study's results please refer to the technical appendix.

Factors which can have a large impact on web performance

The study shows that NBN access speed and choice of RSP are not factors that impact website performance for the reference sites tested. This means that other things are happening outside of the consumers control that could be affecting the experience. These include the following factors:

Table 2: Key Factors that can influence how a webpage performs

Advertising and Tracking Software	Most modern websites use tracking software to learn more about their users and provide targeted advertisements.
Webpage complexity	Websites which require many objects are likely to be slower to visualise. This is due to the increased time needed to download all separate objects and the possibility these objects could be stored on many different servers with many dependencies.
Webpage design	Design choices can either improve or worsen a website's loading performance. For example, if Javascript files are referenced at the top of a web page, then the rest of the webpage cannot be loaded until this Javascript has been loaded and executed. This may have no bearing on the webpage layout initially, so it is typically recommended to defer loading of Javascript to the end of the page.
Where the site is hosted in relation to the user	Latency is a major factor in web browsing performance. The lower the latency, the faster the objects that form the web page can be downloaded. For example, if a website is hosted on servers in Australia, it will be much faster for Australians to load than if the website was hosted in the United Kingdom. Major websites tend to utilise Content Delivery Networks (CDNs) to help move content closer to the users.

These factors are explained in further detail below.

Advertising and tracking software

A further feature of modern websites is the presence of advertising and tracking software. Tracking software keeps a record of information about a user such as where they click on a webpage and how long they spend interacting with the site. Moreover, it can also enable tracking across browsers and devices. In most cases this information is then used to display targeted advertisements to users, for instance in advertising banners. This software is the reason that after viewing a product on an e-commerce site[3], users may continue to see advertisements for that exact product on other pages as they browse.

The study has uncovered that the bulk of tracking software employed by Website 1 is proprietary software and is used to track users as they interact with the webpage. This information is then used to provide the user with more targeted advertisements on other sites.

This software can make websites more complex and impact web performance.

3. E-commerce is defined as 'the activity of electronically buying or selling of products on online services or over the Internet'- https://en.wikipedia.org/wiki/E-commerce.

Webpage complexity

The first webpage[4] available on the internet was approximately 4KB in size and composed of plain text. Web browsers at the time could not display images or videos so as a result webpages only supported HTML. In 1993 browsers began to support images and since then additional objects have also been supported, such as videos. This has meant that webpages have become more and more complex. The average webpage size as of April 2021 is approximately 2MB in size.[5] The image below provides an example of the different objects that could make up a webpage.

Figure 2: Objects that make up the average webpage

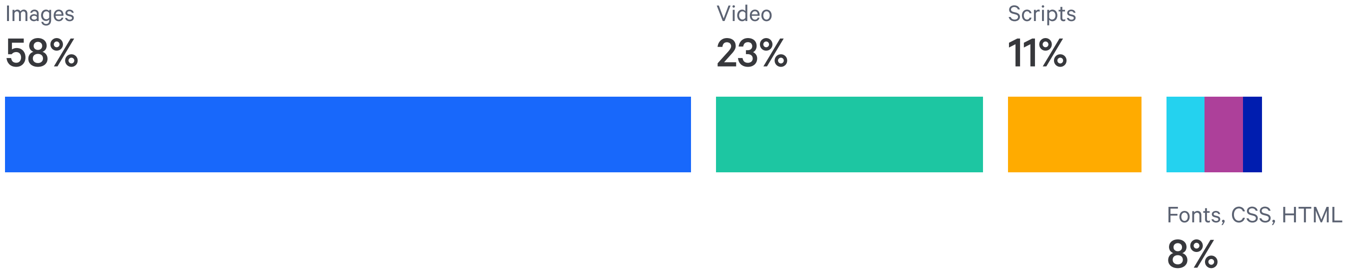

Larger webpages with lots of different objects, can take longer to load as high resolution images can consume lots of bandwidth whilst loading. Website 1 is a good example of this, the website loads with many images as well as a carousel advertising different products across the top.

Compared with the first webpage built below, it is easy to see why the number of objects a webpage has to load may affect the performance of the website.

Figure 3: The first webpage created

World Wide Web

The WorldWideWeb (W3) is a wide-area hypermedia information retrieval initiative aiming to give universal access to a large universe of documents.

Everything there is online about W3 is linked directly or indirectly to this document, including an executive summary of the project, Mailing lists , Policy , November's W3 news , Frequently Asked Questions .

What's out there?
Pointers to the world's online information, subjects , W3 servers, etc.
Help
on the browser you are using
Software Products
A list of W3 project components and their current state. (e.g. Line Mode ,X11 Viola , NeXTStep , Servers , Tools , Mail robot , Library)
Technical
Details of protocols, formats, program internals etc
Bibliography
Paper documentation on W3 and references.
People
A list of some people involved in the project.
History
A summary of the history of the project.
How can I help ?
If you would like to support the web..
Getting code
Getting the code by anonymous FTP , etc.

4. The first webpage on the internet was created by Tim Berners Lee and can be found here: http://info.cern.ch/hypertext/WWW/TheProject.html
5. Source: https://httparchive.org/reports/page-weight#reqTotal

Webpage design

Website designers have many tools at their disposal to help improve website performance. These can include things like optimising image sizes, so a website is not trying to load very heavy images. Webpages can also remove auto-play multimedia, again something that slows down websites and causes them to lag. Design choices such as using white space more frequently can help reduce the amount of data needed to serve the page. White space means any unused space on the page – where there is no colour, no images, no code and no videos.

Where the site is hosted

Major websites make use of CDNs to host their content. Sometimes multiple CDNs are used by a single website, in tandem with their own hosting. An issue between the user's web browser and just one of these server locations could significantly increase webpage loading times. Some objects (e.g. HTML, CSS, or JavaScript) have other objects as dependencies (for example, CSS files can reference external 'web fonts'), and if they are the bottleneck, these other objects can further impact the page load time. By identifying where the slowdown occurs, website developers can take measures to mitigate these issues. Some issues can be solved with page optimisation alone, whilst others will require working with the hosting provider for the website.

Technical Appendix

This section is intended to provide a more detailed and technical discussion of the methodology and results from the Web Performance testing of Website 1 and Website 2.

The test

The SamKnows Web Performance Test uses the engine of a real web browser to measure how long a webpage takes to render from a real person's home. This allows us to accurately simulate a consumer's actual web browsing experience. The testing also considers many factors that a test run from a data centre would be unable to measure, helping to give us real insights into end users' experiences.

For example, the test allows us to consider all of the following, including:

- the user's internet connection performance
- the user's RSP's core network and routing
- the user's RSP's DNS infrastructure[6]
- any CDN relationships that the user's RSP and/or the content provider has
- how the user's RSP connects with the content providers

The purpose of the SamKnows Website Performance Test is to capture the user experience of a website by testing from inside the user's own network. Since the browser is a graphical program that renders the page step by step as objects are received from the network, the results must reflect what can be seen on the screen.

The goal of the test is to highlight how long someone has to wait before they can start accessing the contents of the page. However, not all objects for a webpage have visual relevance. Some objects are invisible, such as an advertising beacon, while others have less impact on the visuals overall. Capturing a "visual loading progress" is a useful metric that allows us to track how much of a page is visible to the user throughout the few seconds that a page is loading.

In addition to recording the overall webpage loading time, we can also measure the performance of specific objects, which can then reveal why the page is slow to load, or even broken (e.g. a slow-to-download JavaScript file may slow down the entire page load, but a slow-to-download GIF may not).

The target websites

For the test to produce completely comparable results the sites to be tested must be broadly consistent in content over the reporting period. This discounts common popular sites such as news and social media sites, which are highly dynamic in content. The two chosen websites both met the criteria to allow comparison.

6. The DNS infrastructure may be critical in selecting a geographically relevant server, which may impact the page rendering time

Our chosen metrics

As the test aims to determine end users' experience, the metrics we will focus on measure the time taken for objects to become visible on the page. The two metrics we will use are:

- Time to first visual – this measures the time taken for the first visual object to load. This covers the time the screen is all white until the initial page structure is displayed.
- Time to visually complete – this measures the time taken for all objects with visual impact to be loaded. This is what the user experiences, while background objects resulting in no screen updates may still be ongoing.

Technical Specification

The test runs on a Whitebox representing a user running a desktop browser (not a mobile device). This is controlled by specifying a user agent string to the web server.

Since the number of fetched objects depends on what can be rendered on screen the resolution of the browser needs to be consistent. The pixel size of 1440 x 900 is used which represents the median of the available screen sizes. Thus, page load results should only be compared against test runs with the same screen resolution.

Discovery of web objects is done by tunnelling traffic from an offload server via the Whitebox to the web servers (relay phase). However, the connections that are involved in the actual measurement of web object transfers are originated from the Whitebox. Name resolution requests are managed in a similar way.

The overall page load operation is measured on the offload server using the timing characteristics of name resolution and individual object transfers. It is here 'time to first visual' and 'time to visually complete' is determined by inspecting periodic screenshots of a browser session with replayed objects that are given to the browser with the same timing as the measured Whitebox.

All results are based off test data that was recorded by SamKnows Whiteboxes hosted by Measuring Broadband Australia volunteers.

Results

The following results examine the performance of both sites for all Whiteboxes and under a number of splits for the period of May 2020.

Summary of key findings

- Plan speed does not have a huge impact on web browsing and is not the key driver of web performance as outlined in Table 1.
- Performance across different RSPs is broadly consistent with all RSPs offering a satisfactory service for end users.
- The impact of advertising objects on web performance is multifaceted. This report features a breakdown of the approaches used by the two target websites.

Overview of performance

The table below presents the average performance for the two target sites for all of May 2020.

Table 3: Average performance for the two target websites for May 2020

Website	Time to First Visual (s)	Time to Visually Complete (s)
Website 1	1.9	2.9
Website 2	2.0	2.7

Note: These results include testing of all 12/1 Mbps, 25/5 Mbps, 50/20 Mbps and 100/40 Mbps plans for all tested RSPs.

The average time to first visual for both sites is similar with figures of 1.9 seconds and 2.0 seconds for Website 1 and Website 2 respectively. The time to visually complete are by definition larger with the results coming in at 2.9 seconds and 2.7 seconds for Website 1 and Website 2 respectively.

The results show that on average both sites load all visual objects in less than 3 seconds. The results also allow inference about the structure of the respective pages. Website 1 has a lower average time to first visual than Website 2 which suggests the objects needed for its first visual are quicker to load. This could be due to there being fewer necessary dependencies (e.g. Javascript files) to load Website 1 first visual or that the objects are on less servers and with less bottlenecks. This pattern reverses for the time to visually complete with Website 2 needing less time on average to move from first visual to visual completion.

While bulk averages are useful to give an overall representation of performance, they fail to provide much information on users suffering consistent poor performance. Cumulative distribution plots can provide a more comprehensive measure of performance in accessing sites.

Figure 4: Cumulative distribution function showing 'time to first visual' and 'time to visually complete' for two websites

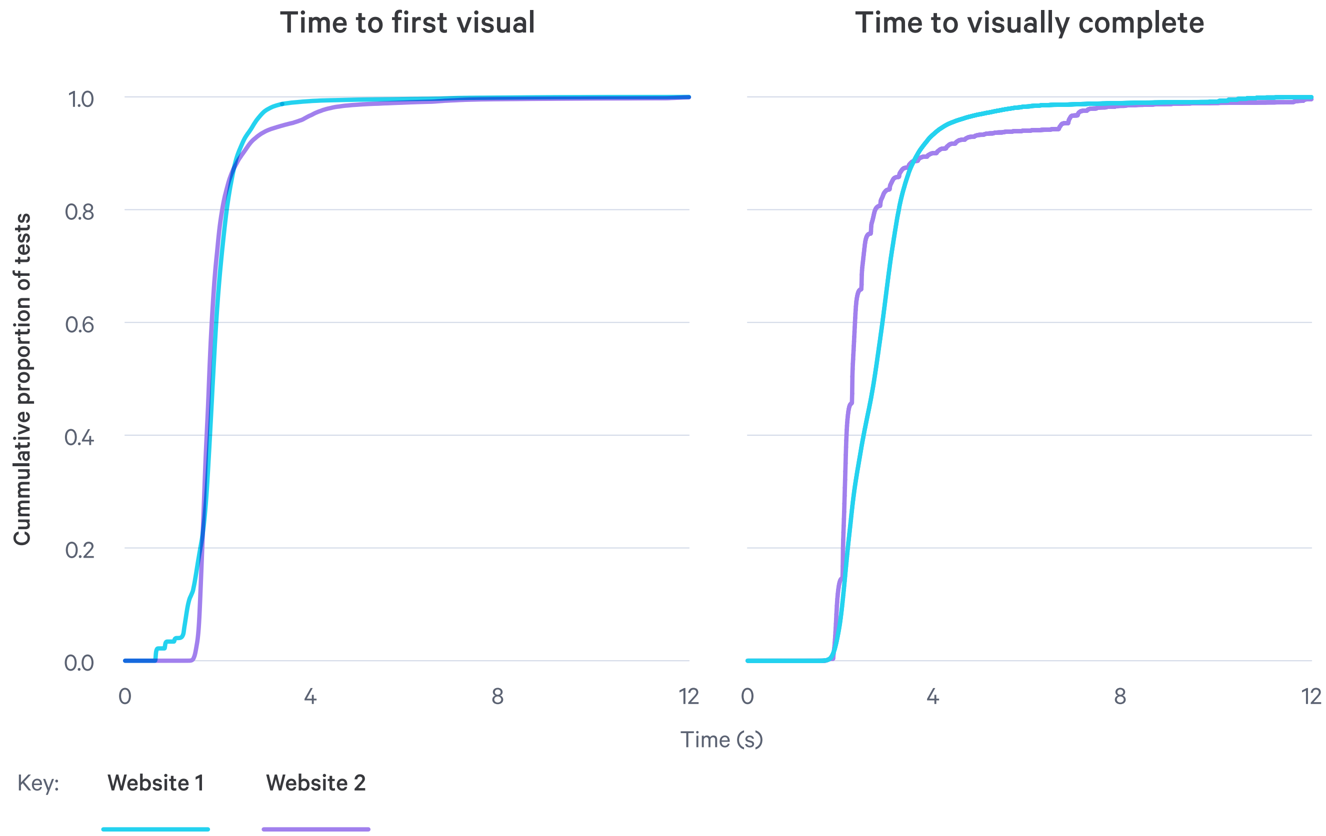

The above cumulative distribution plots show the proportion of tests which are below a given time value for each of the two metrics. Both plots show that for both sites very few tests are seeing large times to first visual or time to visually complete. Performance for Website 2 on time to visually complete shows results getting worse around the 90th percentile. Around 5% of tests took 7 seconds and above to reach visual completion.

The variability of performance over the measurement period is another outcome of interest and would serve to highlight if either site experiences periodic degradation in performance. A simple measure of this is to consider daily average values for all Whiteboxes for both target websites for time to first visual and time to visually complete. A chart of this for all tests in May 2020 is presented below:

Figure 5: Daily average values for all Whiteboxes for both target websites for 'time to first visual' and 'time to visually complete'

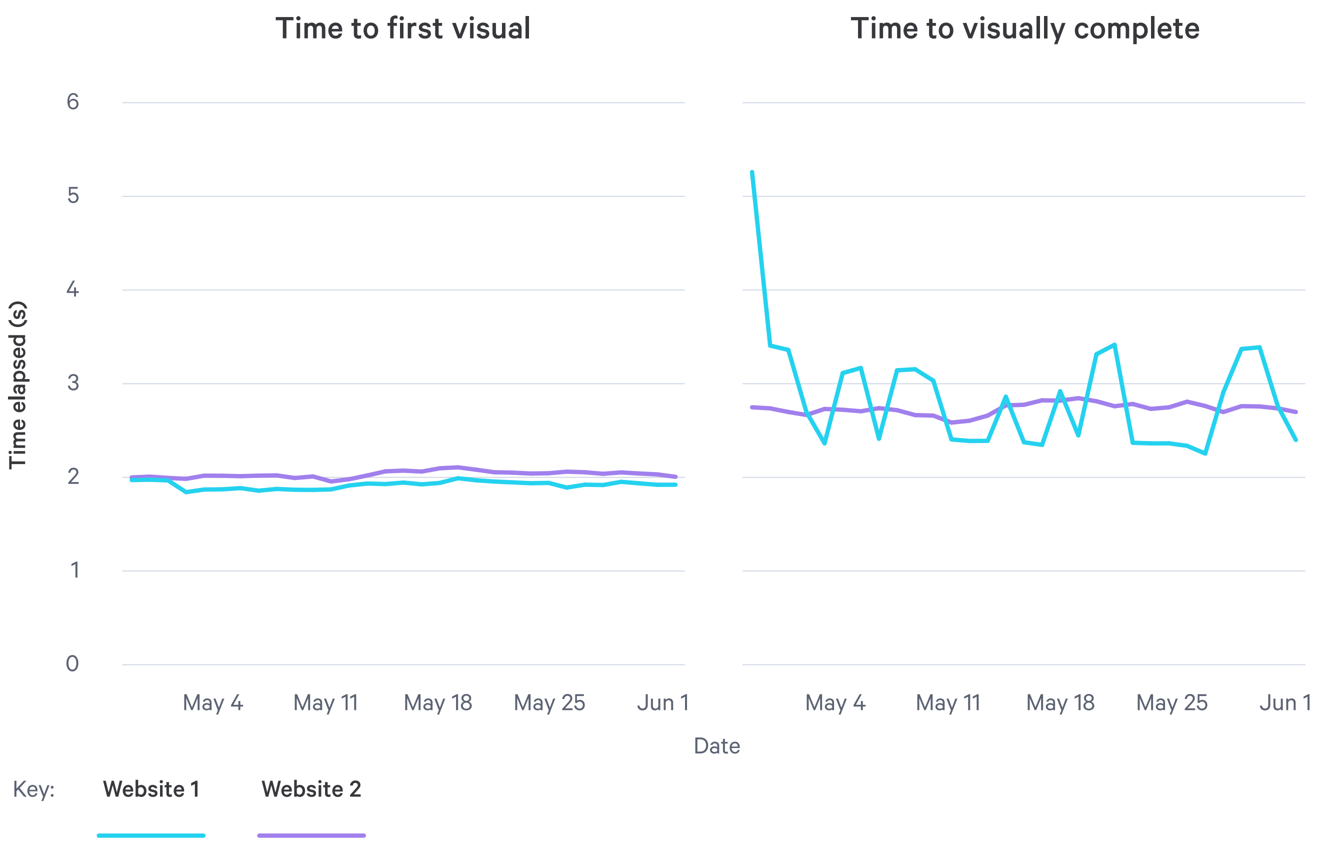

The chart shows that for Website 2 the average daily time to first visual and time to visually complete are stable at around 2 seconds and 2.8 seconds respectively. While the daily average time for first visual for Website 1 was also consistent at around 2 seconds, there was more variability in the daily average time to visually complete. At the start of May 2020 we observed a spike in this value above 5 seconds, a level which was never repeated as the daily average then oscillated around the 3 second mark with mild variation.

Website 1's variation in time to visually complete

Case study: Website 1's time to visually complete shows two factors which can be investigated further due to the detail provided by the web performance test. Specifically:

- the spike in time to visual completion at the beginning of May
- the oscillation of average time to visual completion between 2.5 and 3.5 seconds

The first of these factors was experienced by a selection of tests run for each Whitebox on 1 May 2020. On looking at a breakdown of each individual visual object loaded during each test, the cause of the spike in time to visual completion was found to be due to an almost unnoticeable small graphical update to the top banner of the page. This small change is unlikely to have had an impact on user experience as to an end user the page would seem complete and have full functionality.

For the variability, a similar phenomenon is experienced. Individual tests with longer times to visually complete see the bulk of visual process occur before 3 seconds, but small visual changes occur later. These are unlikely to have much impact on user experience.

The test allows for visual progress to be measured by percentile, so we can investigate the time until 90% of the page is loaded. In the case of Website 1, the time to 90% of the page to be loaded is far less variable than the full time to visually complete. This is demonstrated in the figure below which for Website 1 plots time to first visual and time to 90% completed.

Figure 6: 'Time to first visual' and time to 90% completed for Website 1

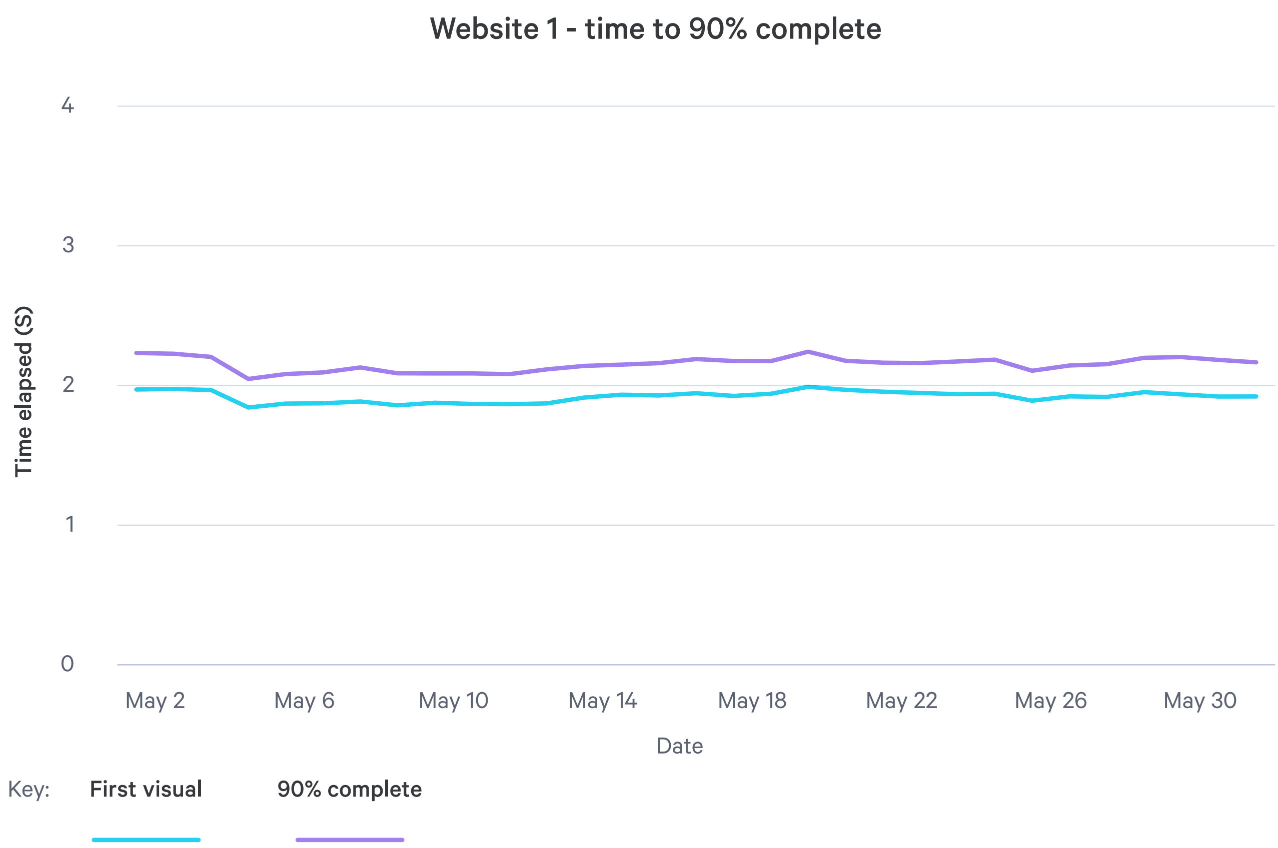

Figure 6 shows the complete removal of the variation seen in the chart of time to visually complete. This suggests the impact of the variability in time to visual completion is unlikely to have significant impact on end users. The spikes in time to being visually complete only covered the last 10% of the page to be loaded and users would be able to interact with a page with 90% of its visuals loaded.

Impact of plan speed on web performance

Web browsing makes up a significant amount of an average consumer's internet usage and a desire to ensure low load times could be a factor in how consumers select their speed tier.

Our results highlight that download speed does not have a significant impact on web performance. The reason for this is that primarily each individual object needed to render a webpage is itself small and thus most broadband plans are well equipped to manage these downloads. The table below presents the average time to first visual and average time to visually complete for each NBN speed tier and each target website.

Table 4: Average time to first visual and average time to visually complete for each NBN speed tier and target website

	Website 1		Website 2	
NBN Speed Tier (Mbps)	Time to First Visual (s)	Time to Visually Complete (s)	Time to First Visual (s)	Time to Visually Complete (s)
12/1	2.3	3.3	2.9	3.6
25/5	2.0	2.8	2.2	3.1
50/20	1.9	2.8	2.0	2.7
100/40	1.9	2.8	2.0	2.6

The results show clearly that while higher speed tiers do have faster loading times, the impact of an increase in download speed is met with a far more marginal decrease in loading time. Doubling your download speed by moving from the 50/20 Mbps tier to the 100/40 Mbps tier would lead to almost no noticeable improvement in web performance for both of the sites tested.

An examination of cumulative density plots to compare the speed tiers similarly reveals only marginally better performance for higher speed tiers.

Figure 7: Cumulative Distribution Function comparing 'time to first visual' and 'time to visually complete' split by NBN access speeds

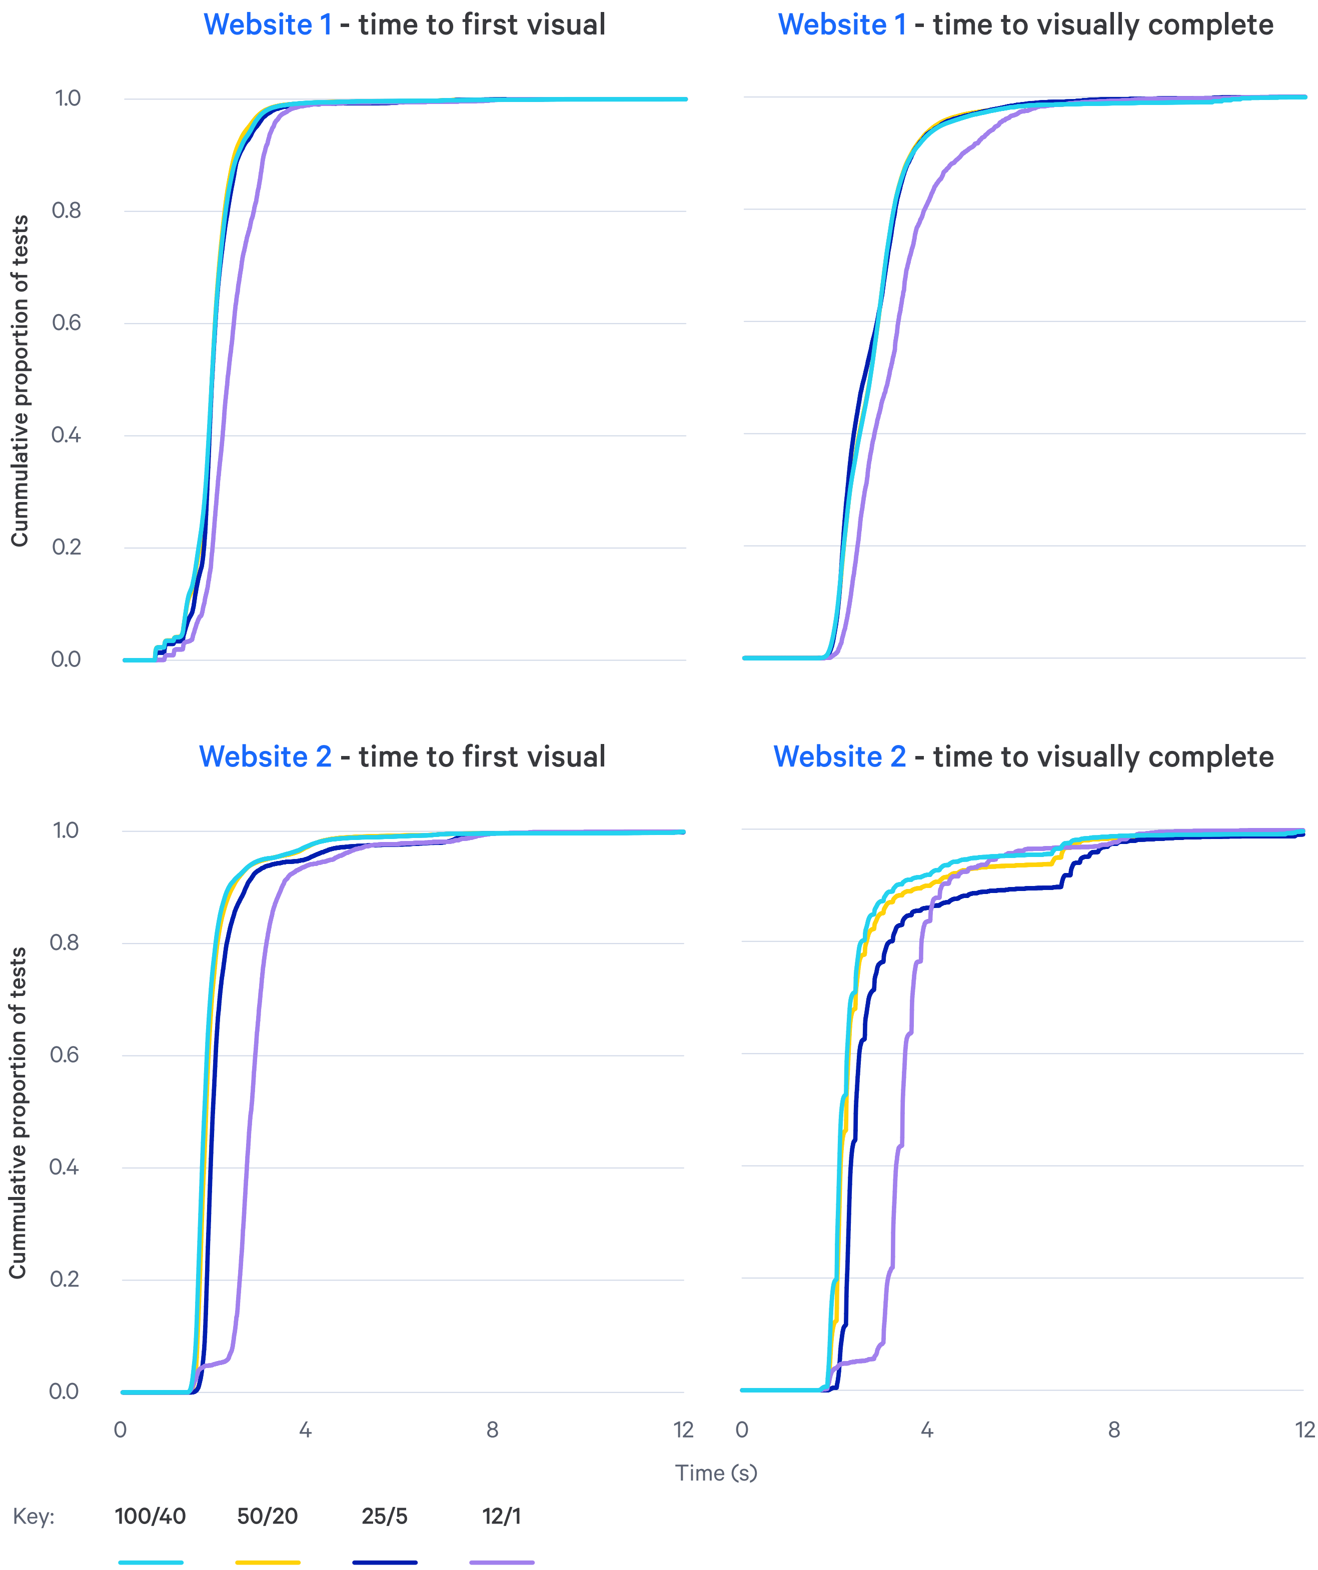

The cumulative density charts show that on both measures for both sites the 25, 50 and 100 Mbps speed tiers have similar performance, for all RSPs tested.

Performance of different RSPs

RSPs can have significant influence on web browsing performance through configuration choices for their networks such as the deployment of web caches. These decisions can then impact the load time of the many objects needed to render a webpage for an end user.

The table below presents the average time to first visual and average time to visually complete for all tests over the month of May 2020 split by RSP. As these testing results were undertaken in May 2020 they may not be fully reflective of the current performance of RSPs or targeted websites.

Table 5: Average time to first visual and average time to visually complete split by RSP for May 2020

	Website 1		Website 2	
RSP	Time to First Visual (s)	Time to Visually Complete (s)	Time to First Visual (s)	Time to Visually Complete (s)
Aussie Broadband	1.9	2.8	2.0	3.0
Dodo & iPrimus	2.0	2.9	2.3	3.0
Exetel	1.8	2.7	1.9	2.5
iiNet	2.0	2.9	2.0	2.5
MyRepublic	1.9	2.8	2.1	2.6
Optus	1.8	2.7	1.9	3.4
Telstra	1.8	2.6	1.9	2.3
TPG	2.2	3.3	2.0	2.6
Vodafone	1.8	2.6	1.9	2.4

Note: These results include testing of all 12/1 Mbps, 25/5 Mbps, 50/20 Mbps and 100/40 Mbps plans for all tested RSPs.

Some key takeaways from the results split by RSP are:

- The results for all RSPs were very satisfactory with immaterial differences in performance between providers.
- The majority of RSPs mirror the overall results for each metric presented earlier in this report.
- The headline differences in these metrics are unlikely to have a noticeable impact on the end user.
- Visualising these results through distribution plots leads to the same overall conclusions.

Figure 8: Average time to first visual and average time to visually complete split by RSP for May 2020

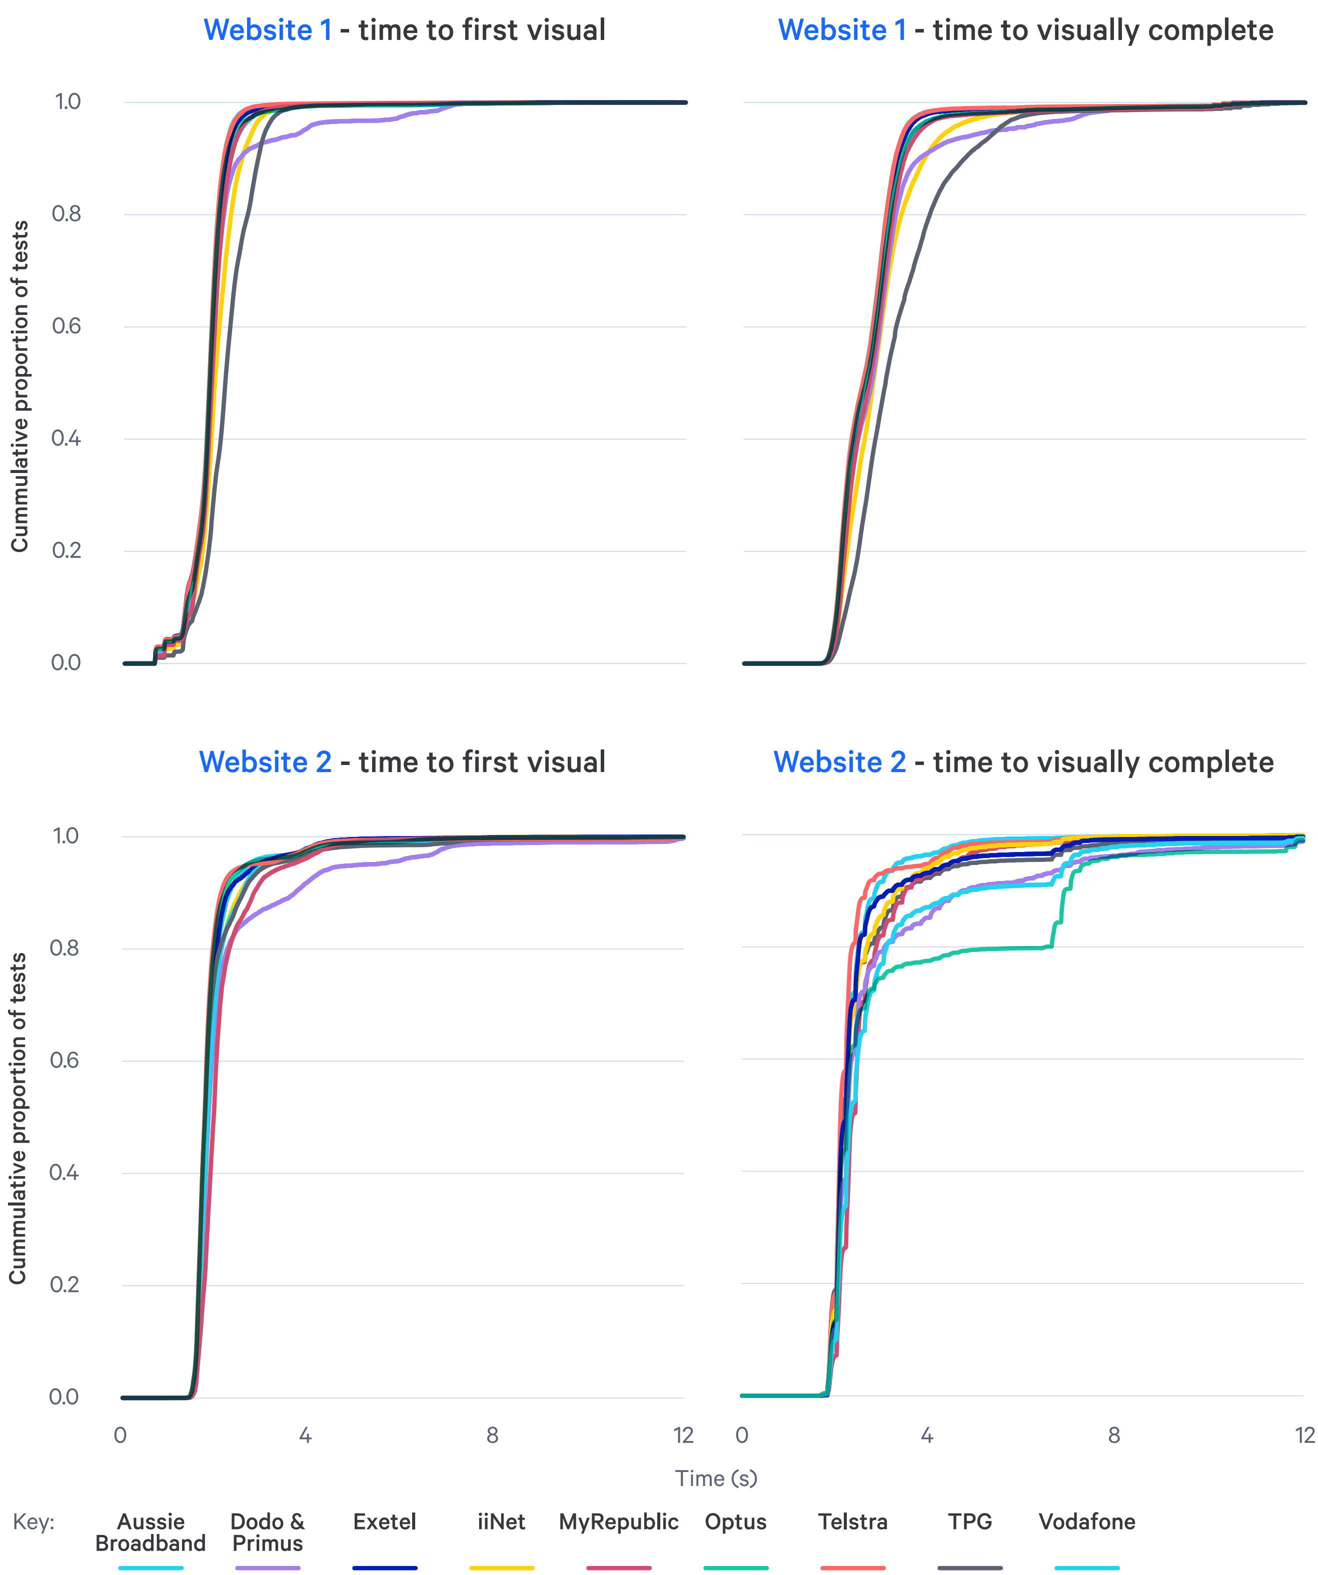

For time to first visual most RSPs show a consistent distribution of performance. The key takeaways are:

- Most providers have no unusual issues with performance at this level for both sites.

For time to visually complete, the chart is similar and most RSPs show similar distribution of times. In summary:

- The bulk of RSPs are having no unusual issues with performance at this level for both sites.

The impact of advertising

The process of loading a webpage can involve the additional features on top of the files used solely to provide the visual content. A common feature of modern sites and especially those based on e-commerce is the presence of advertising and/or tracking objects. These keep a record of information about a user such as where they click on a webpage and how long they spend interacting with the site and can also enable tracking across browsers and devices.

Both target websites employ these techniques. The addition of these objects can impact load times and poor implementation or design with the use of multiple third-party objects stored across many server locations could lead to poor performance.

The following section provides information on advertising and/or tracking objects for each of the two sites studied in this report.

A crude initial measure of a page's complexity is the average number of objects contained within it. The table below provides this figure for both sites for all tests during May 2020.

Table 6: Average number of objects per webpage

Site	Average number of objects
Website 1	58.3
Website 2	91.9

On this measure, Website 2 site attempts to load over 50% more objects than Website 1. This result is a good starting point for expanding upon the relative structure of both sites and how they incorporate advertising and tracking objects.

Website 1

As stated above, Website 1, on the crude measure of the number of objects loaded, the less complex of the two sites. In general, Website 1 makes use of fewer third-party applications with the majority of the objects loaded coming from a small number of services. These include Website 1's own proprietary advertising software.

Website 2

The Website 2 site contains a large number of individual objects and in contrast to Website 1 many of these objects are provided by third parties for advertising and/or tracking purposes. In most cases, these objects are loaded after the site's visuals and have little impact on time to first visual or time to visually complete.

Conclusion

The way websites work has changed as the internet has evolved, as outlined by the comparison between the first website ever made and the two websites tested. This trend towards feature rich websites looks set to continue as networks get faster. The study aimed to assess factors that impacted website performance and if any could be controlled by consumer choice.

The results show that for the target websites tested the factors a consumer has a control over, such as choice of RSP and access speed, do not have a material impact on the experience of browsing webpages. The results indicate that a consumer's website experience is more likely to be impacted by factors outside the consumer's control. For example, website design, website complexity, the presence of advertising software and where the website is hosted geographically can all affect website performance.