



Dynatrace Performance Benchmarks

Methodology

Contents

1. Overview: What are Dynatrace Performance Benchmarks?	3
1.1 How can you best use Benchmarks?	3
1.2 How do we choose who is included?	3
1.3 How do we measure?	3
1.4 What do we measure?	4
2. Ranking, Site Changes and Publication	5
2.1 How do we rank sites?	5
2.2 How do we handle site changes or agent issues?	5
2.3 How do we publish?	5
3. Key Calculations/Terms	6
APPENDIX A: Annotated Sample Performance Benchmark Report	7

1. Overview: What are Dynatrace Performance Benchmarks?

Competitive context is powerful. It can help us know what is possible and where we are falling short.

Dynatrace provides over twenty-five global benchmarks that include hundreds of sites to provide our customers this competitive context. The benchmarks are in key industry verticals around the globe and cover both Homepages and full Clickpath transactions.

Using Dynatrace's Synthetic monitoring product (<https://www.dynatrace.com/platform/synthetic-monitoring/>) we collect technical quality (performance) data on the top sites in each of these key verticals and that data is ranked with results published each week. While the data is captured using our synthetic product it does represent a typical customer experience. To this end, we use real browsers and measure from outside on the Internet. This data provides our customers an understanding of how fast is fast enough, or how an industry is evolving over time.

1.1 How can you best use Benchmarks?

First, it is important to understand their limitations. If you are subscribed to a Homepage benchmark, know that while the Homepage is extremely important it is not the entire site. Further study should be done on other key pages. The same is true for even our Clickpath benchmarks, which will be measuring more pages but do not measure every page on the site. Any good performance management strategy should include additional synthetics measurements on the pages you know to be key.

Secondly, benchmarks do not capture every possible user configuration. We use a standard browser, operating system and connection. Your actual customers are not so consistent. If you want to capture the many variables of your customer experience, consider adding Real User Monitoring (RUM) (<https://www.dynatrace.com/platform/real-user-monitoring/>) from Dynatrace. RUM captures all browsers, all users, all

locations and this wider lens captures the entire experience of customers on your site.

Thirdly, the benchmarks provide objective ranking data focused on performance. They are not meant to be a full analysis or deep dive of the competitive landscape. Said another way, benchmarks give you answers to the "what" questions not the "why" questions. Answering the "why" questions requires a deeper study. Dynatrace would recommend a deeper Competitive Assessment (<https://www.dynatrace.com/services-support/business-insights/#unique-html-id-2>) a few times a year to dive deeper into how your site can improve.

1.2 How do we choose who is included?

We align our benchmarks with our customer's industries and geographic regions. Therefore, you will see such a focus on financial services and retail, two of our top customer segments. Within each region and industry, we choose the top sites based on a combination of metrics that include traffic, sales or innovation. We are always open to feedback around new industries and/or additional sites to include in existing benchmarks. We'd love to hear from you at benchmarks@dynatrace.com.

1.3 How do we measure?

Synthetic measurement is perfectly suited for collecting competitive data. It doesn't require anything to be installed on the site under study or even for the site owner to be aware. Synthetic measurement is simply another user type accessing the site. We are very careful to keep our measurement sampling low and ensure that there is no adverse impact on the site.

We use the Chrome browser because it is the mostly widely used browser in the world. The goal of benchmarking is to capture the typical user experience not every possible user experience. Using the most popular browser allows us to do this. We have also found that for most modern browsers there are not many major differences between performance (if your site is slow in Chrome it is likely slow in IE and Firefox) though of course multi browser testing is very important for QA/functionality.

For our Mobile benchmarks, we use a mobile browser and emulate the connection speed (LTE/4G). We find that connection speed emulation produces fairer results for competitive comparison due to the multiple variables that real mobile networks can introduce. This is especially true for mobile availability where we find that most errors are tied to the carriers themselves and not the site. Using an emulated LTE connection eliminates these errors and allows our benchmark customers to focus on what they have more direct control over.

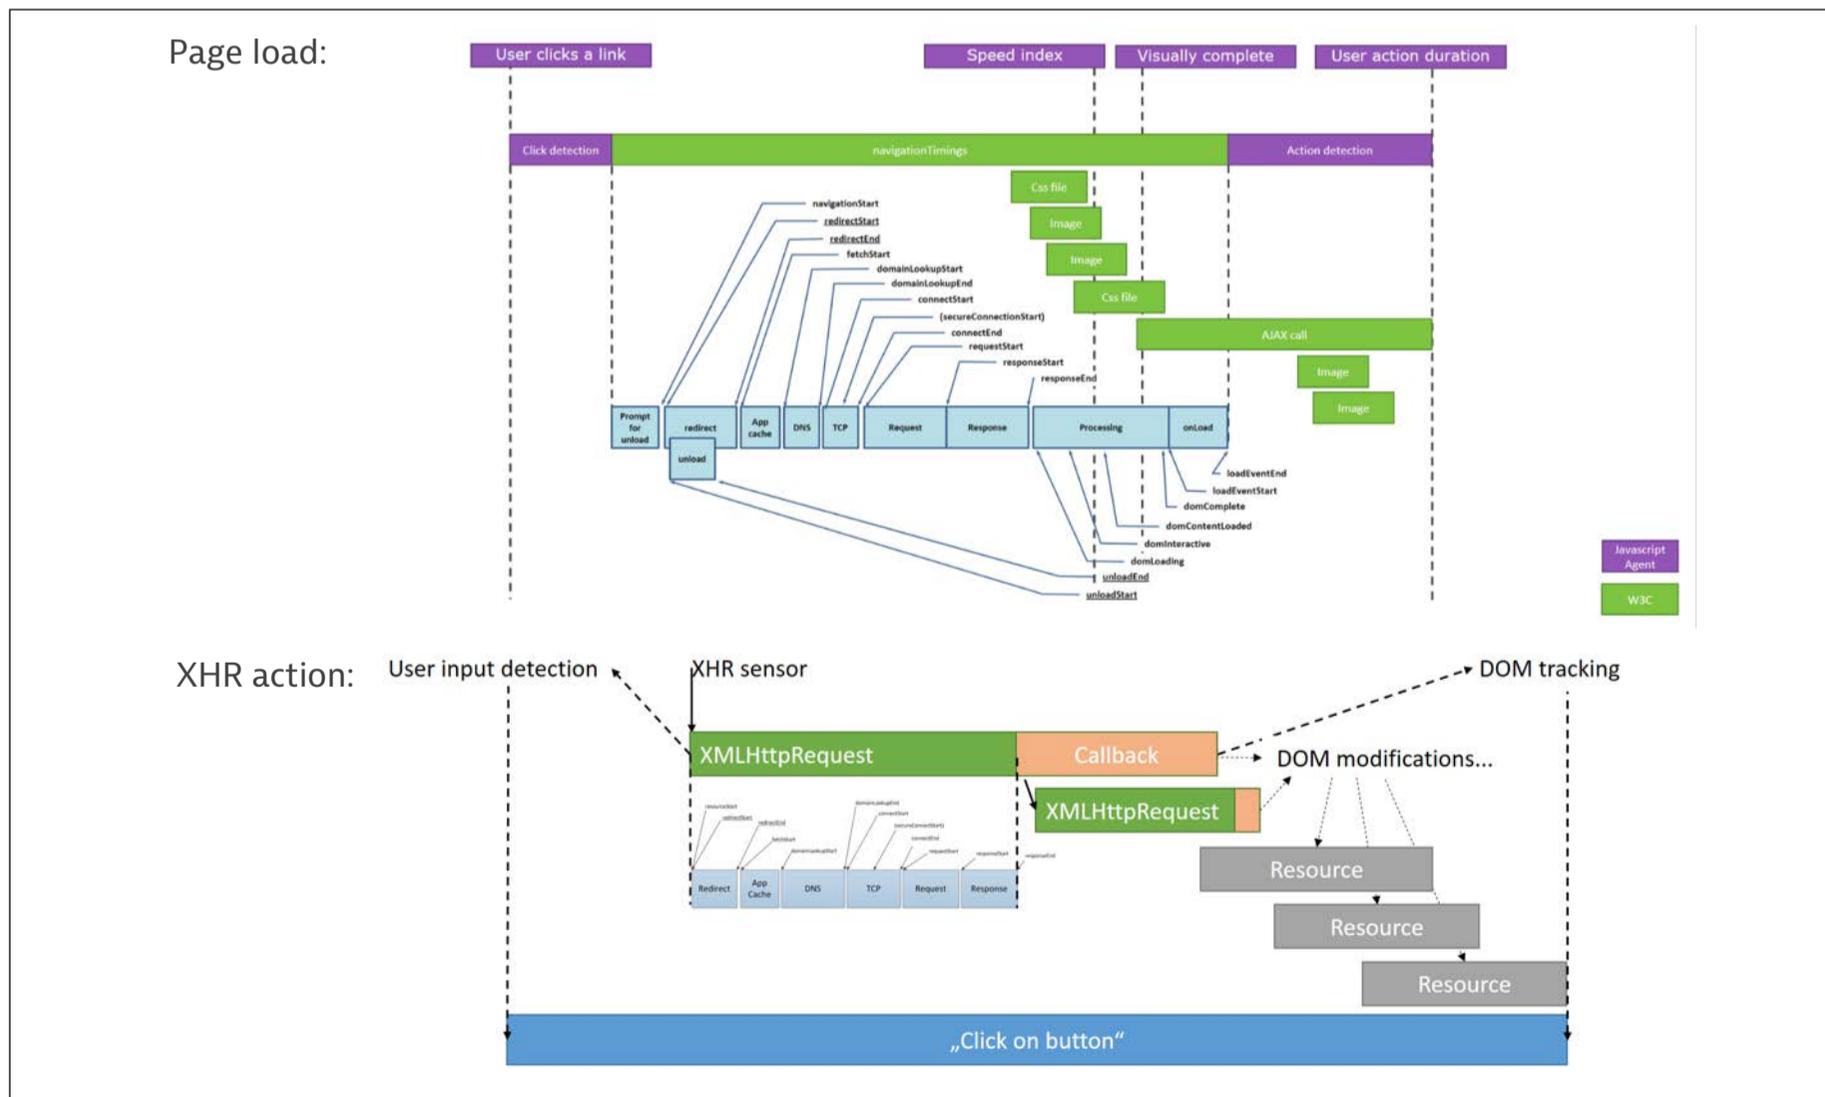
For our Desktop benchmarks, we use very high-speed connections. These connections are faster than a typical user (but not that much faster) but eliminate the "noise" of last mile ISP providers.

For all our benchmarks we choose to measure multiple locations from each region. We always include at least two locations to eliminate the possibility of a single location skewing data. For most of our benchmarks we have at least three locations.

1.4 What do we measure?

We have two types of benchmarks, Homepage and Clickpath. Homepage benchmarks measure a site's Homepage because of how important these pages are to the customer experience. Homepages are still the most common entry point on most sites, set the expectation for performance and their performance is important for SEO rankings. You can tell a lot about a site by the performance of its "front door".

Clickpath benchmarks focus on key user journeys on a site. Almost all site visitors are doing more than just hitting a Homepage. They are accessing their bank accounts or buying products. Our Clickpath benchmarks try to capture these typical user paths. We can't capture every page and we certainly can't emulate the many paths users take through the site. So, we must choose a set of key pages that both represent a typical experience and can be done on EVERY site included in the benchmark. This doesn't mean the number of pages measured will always be equal, as each site has a unique UI design, but the journey will be comparable.



2. Ranking, Site Changes and Publication

2.1 How do we rank sites?

We rank sites based on one key performance metric, Visually Complete. When looking at performance no one metric is perfect. The goal is to find a metric that in most cases represents when a user perceives the page to be complete, and is generic enough to be used fairly on all sites and page load types. We believe that Visually Complete satisfies this goal as it measures the time it takes for the visual area (above the fold) of a page to render. Other page timing metrics may get held up by background page elements that have no bearing on a user's perception of a page being ready to use.

Optimizing Visually Complete timing is more valuable than optimizing other page load timings in terms of user experience as it reflects the amount of time that your real users spend waiting for above-the-fold content to load completely. Above-the-fold is the visible area that extends from the top of the page to the bottom of the browser window.

Visually complete has the following benefits:

- Let's you see exactly how long it takes your end users to see the information they're looking for.
- Provides a true business-relevant metric from the perspective of the user.
- Gives 100% visibility into actual real user experience, regardless of device.
- Aligns IT and business in connecting user experience with business outcomes.
- Accelerates performance improvements.
- Optimizes decisions across development, operations, and user experience.

2.2 How do we handle site changes or agent issues?

Sites are consistently changing, and this requires us to update our synthetics measurements. This is especially a problem for Clickpath measurements. As sites are

updating, we update our measurements. But this might take us a few days and leave inaccurate data. We have no SLA for when benchmark scripts will be updated and in some rare cases updates might take weeks. This may mean that a site is temporarily removed from our weekly report due to prolonged script issues. For this reason, if you are a Dynatrace customer and are one of our benchmarked sites, it is always recommended you monitor your application separately from the benchmark.

We are committed to data accuracy and as part of our publication process we make every attempt to exclude any data that is not reflecting actual site performance. This "exclusion" process can remove data for periods when the script was failing or for any node issues.

We do not exclude Internet issues, third party issues, or maintenance periods outside of our defined time periods. While some issues are not under our customer's direct control they do impact actual site visitors. Given this publication is a service for all customers, we have the final word on what is excluded or what is not. To maintain the integrity of the benchmarks we cannot exclude data just because a customer asks us, or it might make their site rank more poorly.

We apply exclusions in our reporting system and do not make changes to any raw data in your account.

2.3 How do we publish?

We publish the ranking results every week for the previous week. There are two report types, Homepage and Clickpath. Homepage is a free subscription available to anyone. Clickpath is a paid subscription only available to Dynatrace customers. Both reports contain an industry snapshot of key averages, detailed ranking results for every site monitored, and a csv download of the results (to be sorted or ingested into your own dashboards).

Given customers expect digital properties to be available at all times of the day or night, we measure and compare the sites in the benchmarking using all hours and all days. We do not exclude planned downtime.

3. Key Calculations / Terms

Visually Complete: Visually complete is a point-in-time metric that measures when the visual area of a page has finished loading. Visually complete metrics are typically shorter in duration than comparable metrics (for example, page load time and DOM interactive measures) because users perceive complete page load before 100% of background page elements have loaded. Optimizing visually complete timing is more valuable than optimizing other page load timings in terms of user experience as it reflects the amount of time that users spend waiting for above-the-fold content to load completely. For pages that do not include a Visually Complete, Action Duration is used.

Action Duration: Action Duration is the time required to complete a page load and XHR actions. A page load is an actual page loading in your browser. If you enter a URL in your browser and press enter, a page load occurs. During a page load, many resources are loaded, including images, HTML, and CSS. A XHR action is any action that updates the page DOM via a XMLHttpRequest or fetch().

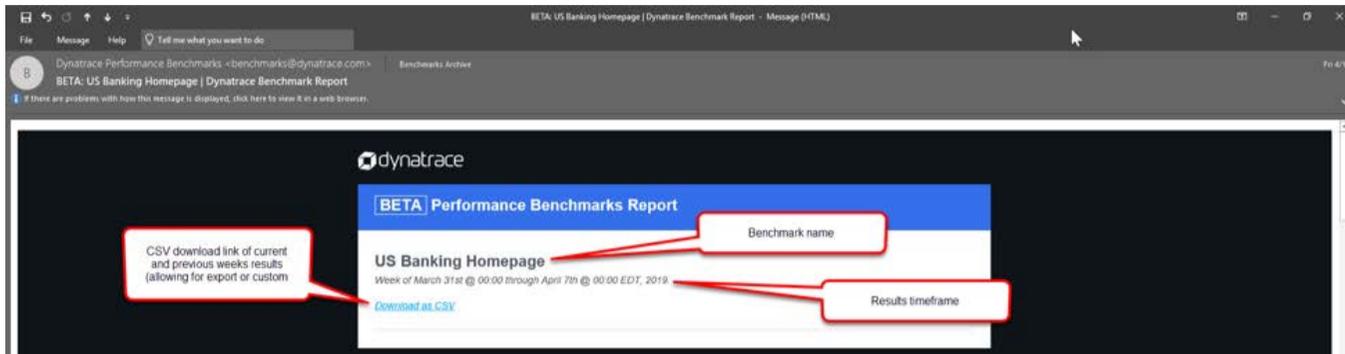
Speed Index: Speed index indicates the average time it takes for all visible parts of a page to display. It's useful for comparing the user experience of various pages. The lower the number, the better the user experience of the page. See the Google WebPagetest documentation page (<https://sites.google.com/a/webpagetest.org/docs/using-webpagetest/metrics/speed-index>) for a detailed definition of speed index and a basic description of the speed index calculation.

Caching: For our measurements we always begin with a clear cache, to capture a first-time site visitor experience. For transactional measurements, we handle caching as a user would as we move from page to page. This means that if the core site framework (JS and CSS) are loaded on the first page in a transaction and cache control is set correctly we will not redownload this content on the next page.

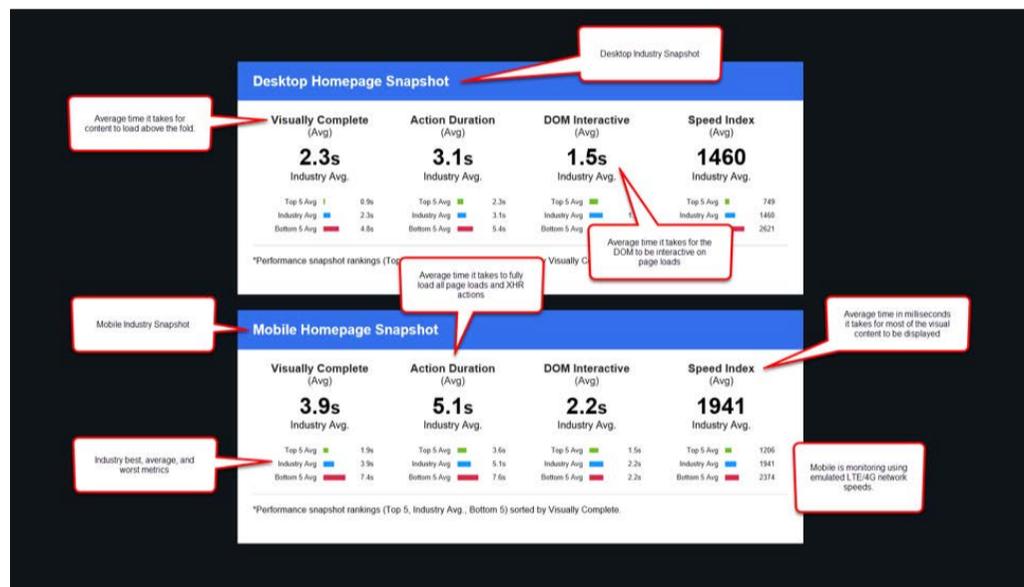
Monitoring Locations: Dynatrace's public synthetic locations are hosted on major global cloud providers such as Amazon EC2, Azure, Alibaba, and Google.

APPENDIX A

Annotated Sample Performance Benchmark Report



For subscribers, we include a link to download a CSV file that include the underlying data for the current and previous week. This allows for sharing, sorting and long-term archiving of the data.



The Industry snapshot attempts to provide not only industry averages but also details on the top and bottom sites. When available, Desktop and Mobile performance are shown in the same report so you get a quick view into their differences.

Performance ranking based on Visually Complete

Sorted on Visually Complete

Previous weeks ranking

Ranking trending (up or down and how many spots)

Rank	Previous	Trending	Participant	Visually Complete	Action Duration	DOM Interactive	Speed Index
1	1	—	Wells Fargo Bank	0.6s	1.2s	0.6s	488
2	—	—	Zions First	0.9s	3.5s	1.0s	631
3	3	—	Eastern Ba	1.0s	2.3s	0.9s	855

Performance rankings are currently set to use Visually Complete. As Dynatrace's synthetic metric collection grows, we will continue to add more metrics that show how sites perform.

The average values are shown given that some sites require fewer or more steps to complete the same business transaction. A high-ranking site with a high average likely indicates that they have fewer steps/pages in their transaction.

[Learn more at dynatrace.com](https://www.dynatrace.com)

Dynatrace has redefined how you monitor today's digital ecosystems. AI-powered, full stack and completely automated, it's the only solution that provides answers, not just data, based on deep insight into every user, every transaction, across every application. The world's leading brands, including, 72 of the Fortune 100, trust Dynatrace to optimize customer experiences, innovate faster and modernize IT operations with absolute confidence

06.16.17 2327_WP_jw

 [@Dynatrace](https://twitter.com/Dynatrace)  fb.com/dynatrace

