

- 
1. Website
 2. Timing comparison
 3. Discussion ShimmerCat effects
 4. Measurement data

Case Study: Edge Accelerator Network

Executive Summary

We did an analysis of performance gains when using ShimmerCat Edge Accelerator Network for the domain <https://www.swedol.se>, particularly on the page <https://www.swedol.se/produktsortiment.html>. We could observe improvements in both measured loading time metrics:

Metric 1: Mean time between end-of-HTML loading and DOMContentLoaded event

- ✓ 48% improvement for desktop
- ✓ 39% improvement for mobile

Metric 2: Mean time between end-of-HTML loading and first contentful view

- ✓ 40% improvement for desktop
- ✓ 83% improvement for mobile

Contact

If you have any questions, don't hesitate to contact us:

ShimmerCat AB

henrik.frienholt@shimmercat.com

+4670-4851947

1 Website

Site analyzed:

- <https://www.swedol.se>

Page analyzed:

- <https://www.swedol.se/produktsortiment.html>

1.1 Measurements and metrics

We created a test setup where we performed automatic measurements to compare the currently online version of the site, running with ShimmerCat, with the origin, which has the original version of the site. In the test setup we created a virtual machine in *cloud.google.com*, that was configured to serve as a proxy reducing bandwidth¹.

The tests were done using traffic control, which is a very useful Linux utility that gives you the ability to configure the kernel packet scheduler². To account for possible differences in network connectivity, both series of tests were configured to have a latency of 400ms. The traffic control settings were:

```
tc qdisc add dev eth1 root tbf rate 2mbit burst 32kbit latency 400ms
```

These network parameters are close to what mobile users experience, and at the same time give a minimum common denominator which makes the comparison possible.

Two basic metrics were measured:

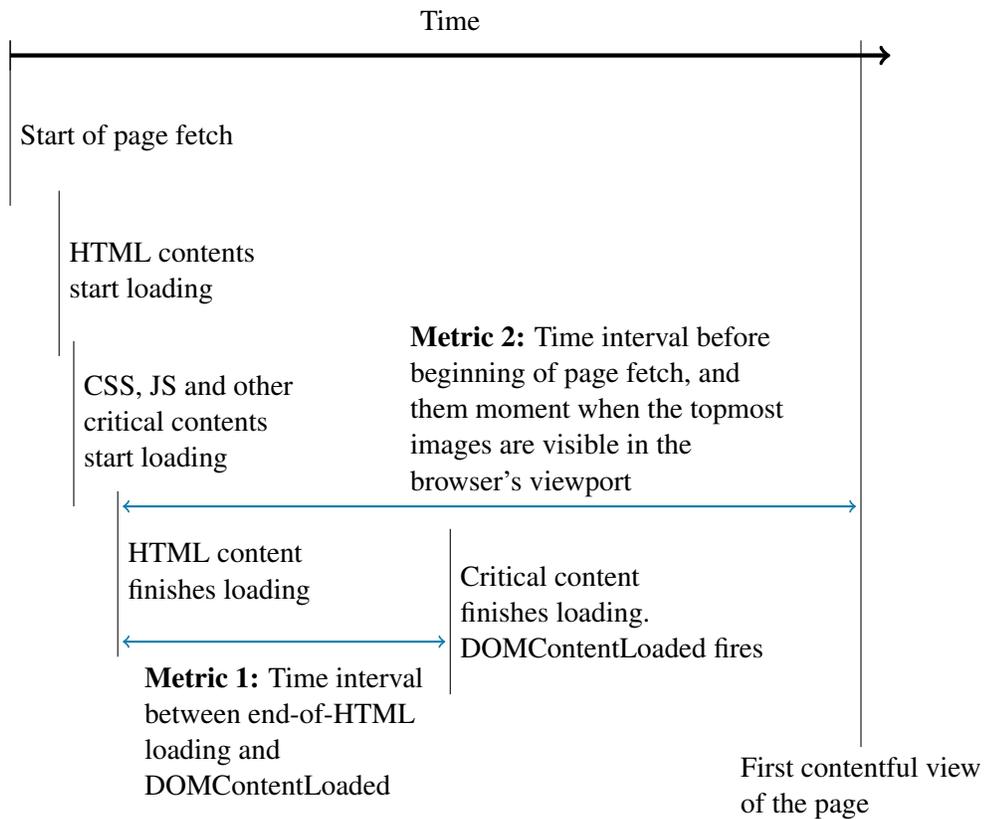
1. Time interval between end-of-HTML loading and DOMContentLoaded event.
2. Time interval between end-of-HTML loading and first contentful view.

The two metrics and the events in time as a web page is fetched can be seen in Figure 1 below. All measurements are repeated several times to account for variation and to establish confidence intervals.

¹<https://github.com/shimmercat/SOCKS5-proxy>

²<https://netbeez.net/blog/how-to-use-the-linux-traffic-control/>

Figure 1: Page load and schematic of metrics.



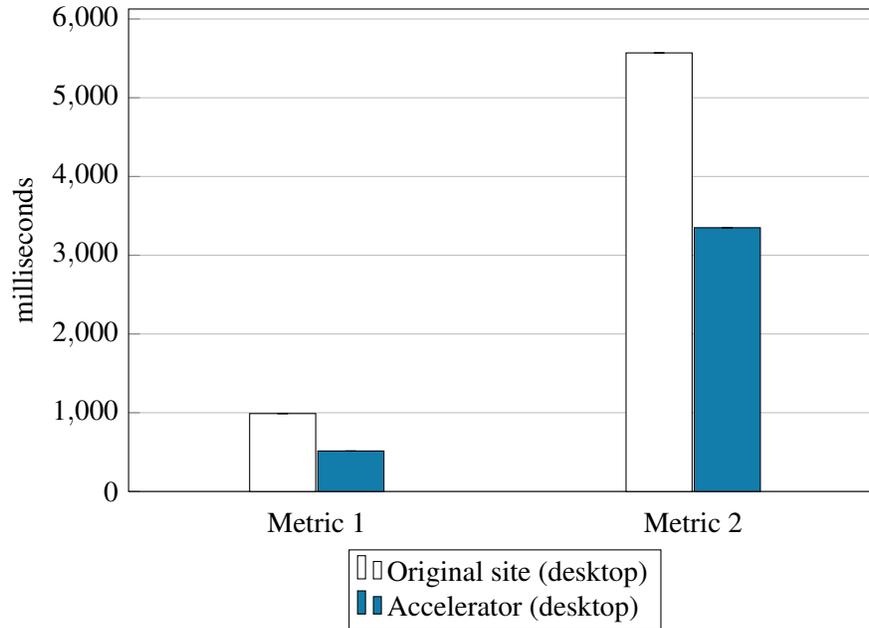
1.2 What is first contentful view?

First contentful view is the moment when all the over-the-fold images are visible in the browser's viewport. This can be seen as the moment where the visitor can start interacting with the website.

1.3 What is DOMContentLoaded?

DOMContentLoaded is an important moment in the lifecycle of a web page, since it's when most Javascript becomes active.

Figure 2: Desktop measurement results for metric1 and metric2.



2 Timing comparison

Figure 2 and 3 below show the compared mean times for desktop and mobile respectively.

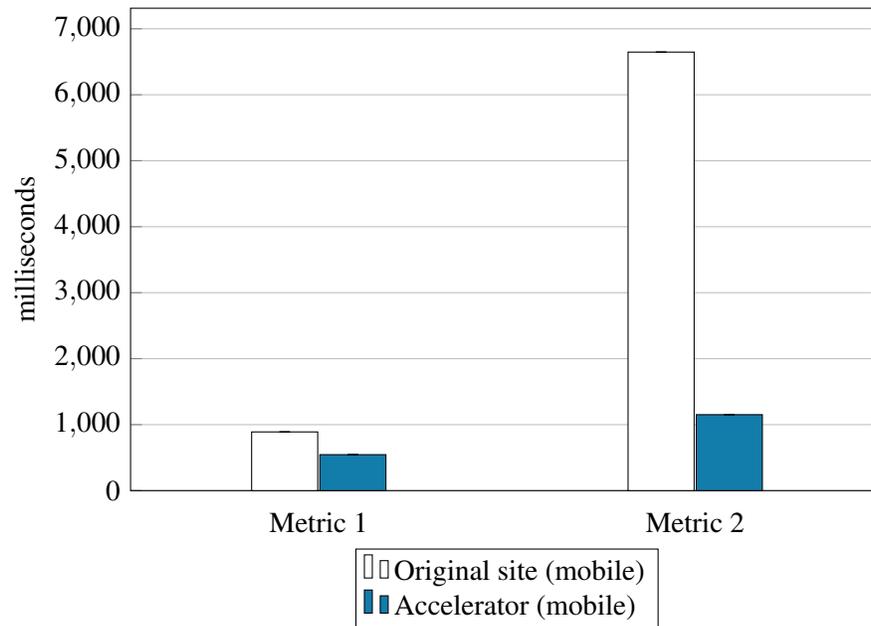
For metric 1 we find a 48% improvement for desktop, and a 39% improvement for mobile. For metric 2 we find a 40% improvement for desktop, and a 83% improvement for mobile.

3 Discussion ShimmerCat effects

The original site uses HTTP/1.1, and it is clear that there are substantial performance benefits by switching to HTTP/2. In addition to the protocol change, three main characteristics of take the results to the level displayed above:

- HTTP/2 priorities: sending critical before non-critical resources.
- Brotli compression: contributes to a 25% reduction in size for static assets compared to origin site gzip
- Image optimization: which converts origin server images in png format

Figure 3: Mobile measurement results for metric1 and metric2.



to next-gen formats, such as JPEG 2000, and WebP. The resulting file size reduction effects are around 70%.

4 Measurement data

Each of the measurements compiled in this report is available in the digital data folder³. The raw data table is indexed on the following components:

- Data series, e.g. *accelerator_on* and *accelerator_off*
- Samples. For a given series, each sample is indexed with a characteristic *experiment_no*.

You will find the data as Pandas python pickle, a format which is popular in modern data-engineering.

You will also find the screenshots used to measure the time to first contentful render, cataloged by series and by experiment-no. The name of each screenshot image contains a number which is the number of milliseconds since the

³<https://drive.google.com/open?id=1KoO9YuQfx-ybU26veQEXsm0UO8vBZ57>

beginning of the page load.