Performance of Requests over the Tor Network

Karsten Loesing
karsten@torproject.org

Tor Tech Report 2009-09-001
September 22, 2009

1 Introduction

Many users perceive the general performance of the Tor network as rather slow. Little is known, though, about real-world measurements of latency and throughput of requests routed via the Tor network. Such measurements are required to evaluate the effectiveness of possible improvements. This report presents a simple measurement setup to gather response times of three different request sizes as well as the first months of results between June 10 and September 22, 2009.

2 Measurement Setup

The performance (latency and throughput) that Tor users experience depends on numerous factors and is the subject of current research. In order to evaluate progress in improving Tor's performance, we need to continuously measure how fast Tor really is for our users. The primary purpose of these measurements is to compare how performance evolves over time. This measurement setup uses a trivial SOCKS client to download files of various sizes over the Tor network and write down how long substeps take. The following configuration parameters are used:

- 50 KiB file, downloaded every 5 minutes, timeout of 4:55 minutes, `MaxCircuitDirtiness` of 1 minute, no entry guards.
- 1 MiB file, downloaded every 30 minutes, timeout of 29:55 minutes, default `MaxCircuitDirtiness` of 10 minutes, no entry guards.
- 5 MiB file, downloaded every 60 minutes, timeout of 59:55 minutes, default `MaxCircuitDirtiness` of 10 minutes, no entry guards.

A more detailed howto describing the setup can be found here: https://gitweb.torproject.org/torperf.git/blob/HEAD:/measurements-HOWTO

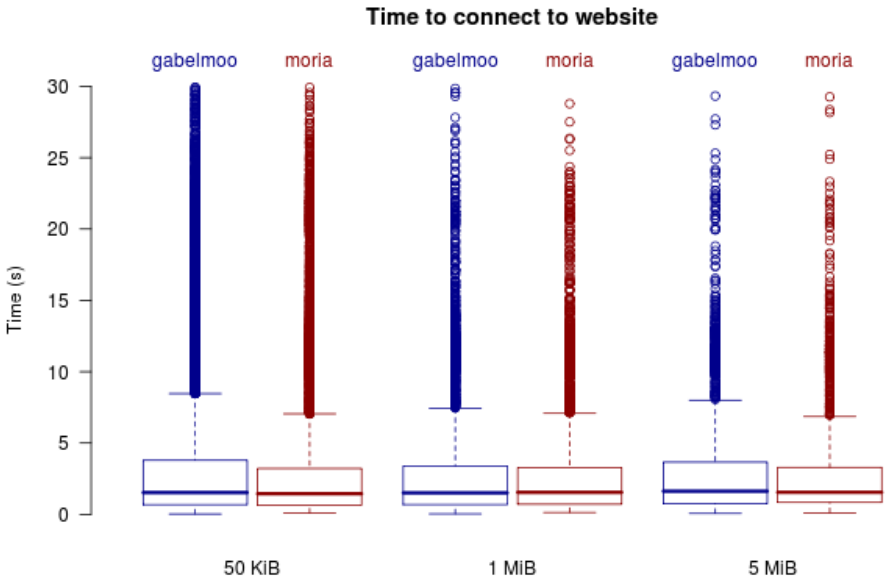

Figure 1: Time to connect to website

3 Results

The measured times include three data points: The first measured time is the time to connect to the website. This step requires Tor to select a pre-built circuit (or create one in rare cases) and send a `CONNECT` command to the exit node. As soon as the exit node replies with a `CONNECTED` cell, the connection to the website is established. Figure 1 shows boxplots[1] of the connection times. The numbers in parentheses denote the total number of runs, the number of runs that timed out, and the number of outliers that did not fit into the graph. The distributions are similar, because the connection time can be assumed independent of the size of the requested file.

The second measured time is the time between starting a request and receiving the first byte of the response. Figure 2 shows these times for the three requested file sizes. Again, the distributions are very similar. The only explanation for deviations comes from the different request numbers.

The third measured time is the overall time to complete a request. These times include all steps from connecting to the website until having received all bytes of the response. Figure 3 shows the distributions of these times. Obviously, there are differences in the completion times for the three requested file sizes.

Figures 4 to 6 visualize the request completion times over the complete measurement interval. These graphs show the medians and the first and third quartile over time. For every run, the median and quartiles are calculated from the interval starting 1 day before a data point and ending 1 day after the data point. This means that intervals contain 577 (97, 49) values,

[1] Boxplots, or box-and-whisker plots, visualize the distribution of a variable: The box part contains all values within the second and third quartile of the data set with the strong line being the median. The dashed lines contain non-outlier values in an interval of 1.5 times the inter-quartile range below the first quartile and the same distance above the third quartile. Outliers are depicted as circles.

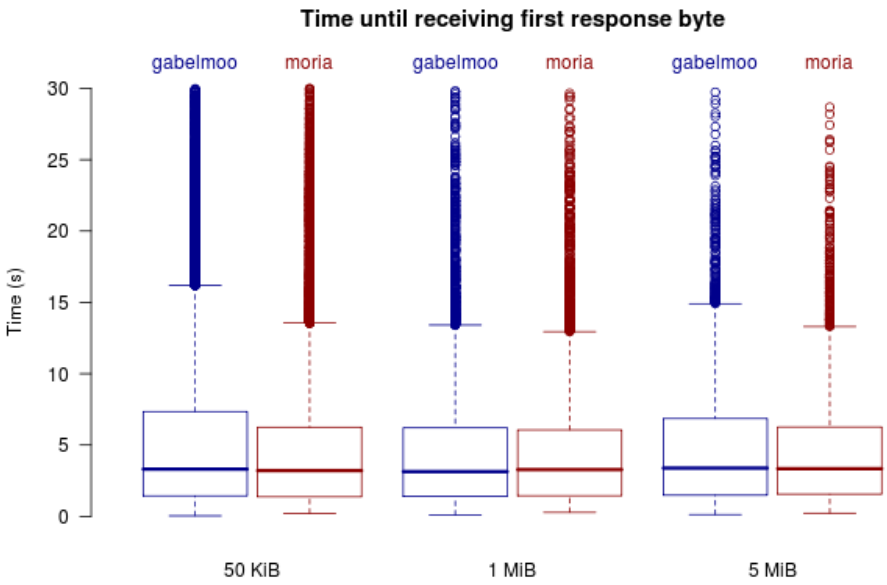

Figure 2: Time until receiving first response byte

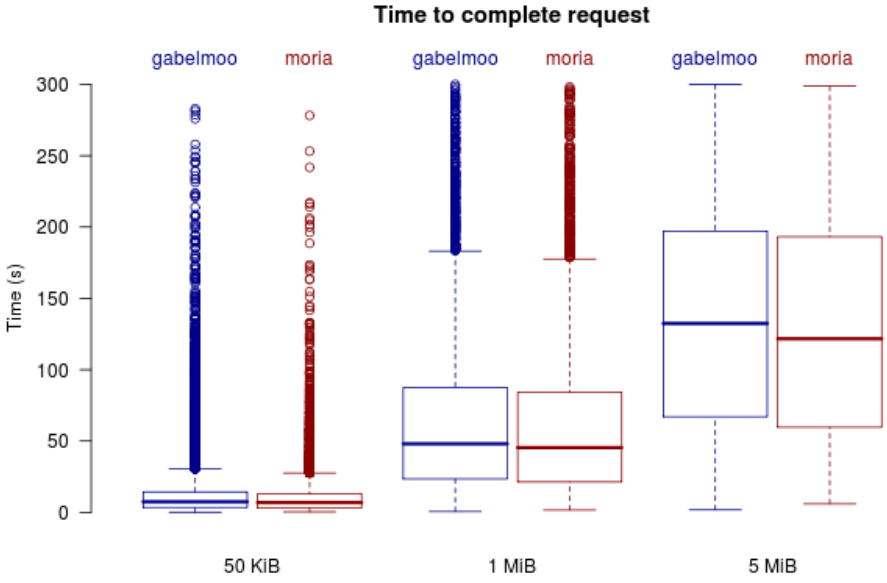

Figure 3: Time to complete request

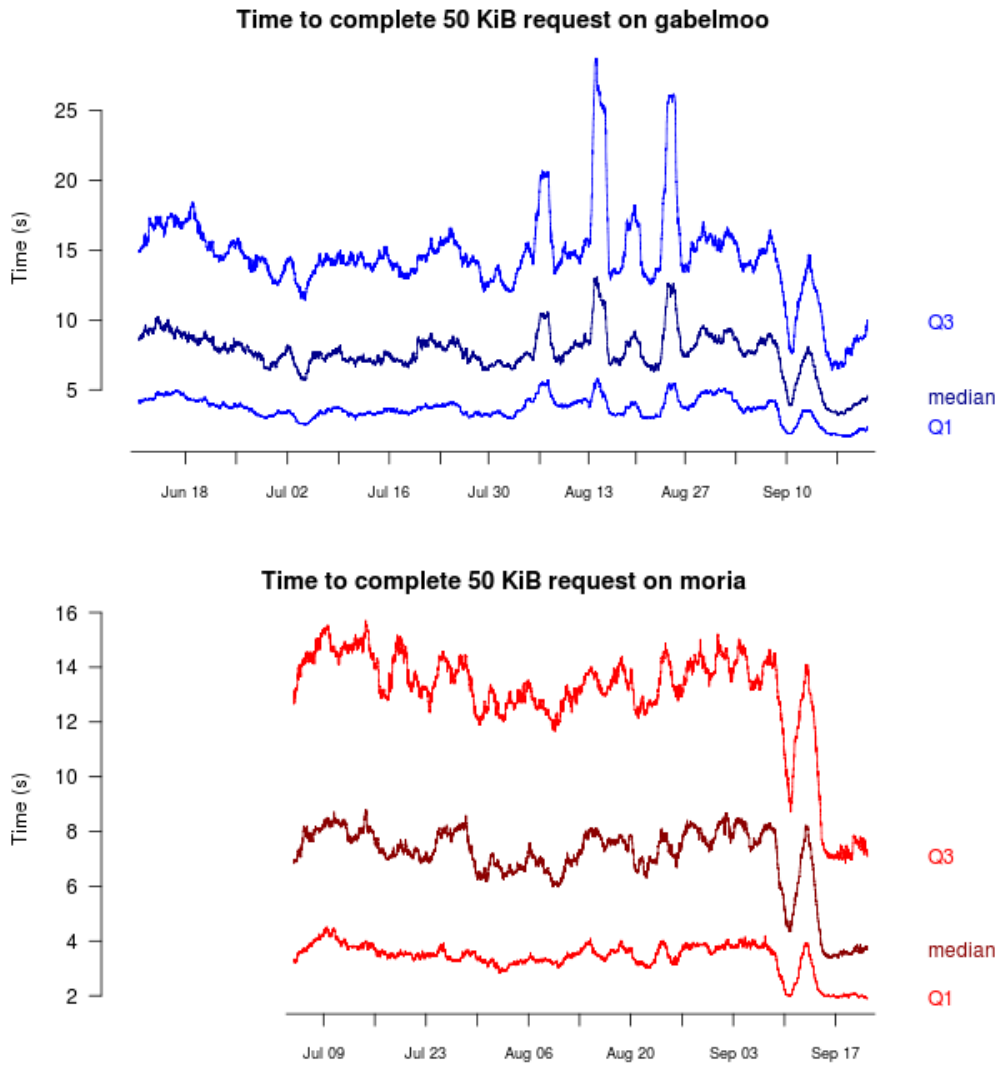

Figure 4: Time to complete 50 KiB request

including the previous 288 (48, 24) data points, the following 288 (48, 24) data points, and the data point itself.

The three peaks in comletion times on gabelmoo in August, that are most visible in the 50 KiB times, are most likely a local problem due to overloading the measuring node. The results of moria do not exhibit any such effects.

The two drops in September are the result of including results of active bandwidth scanners in the directory votes. With these new bandwidth information, the measuring clients make use of optimized load balancing.

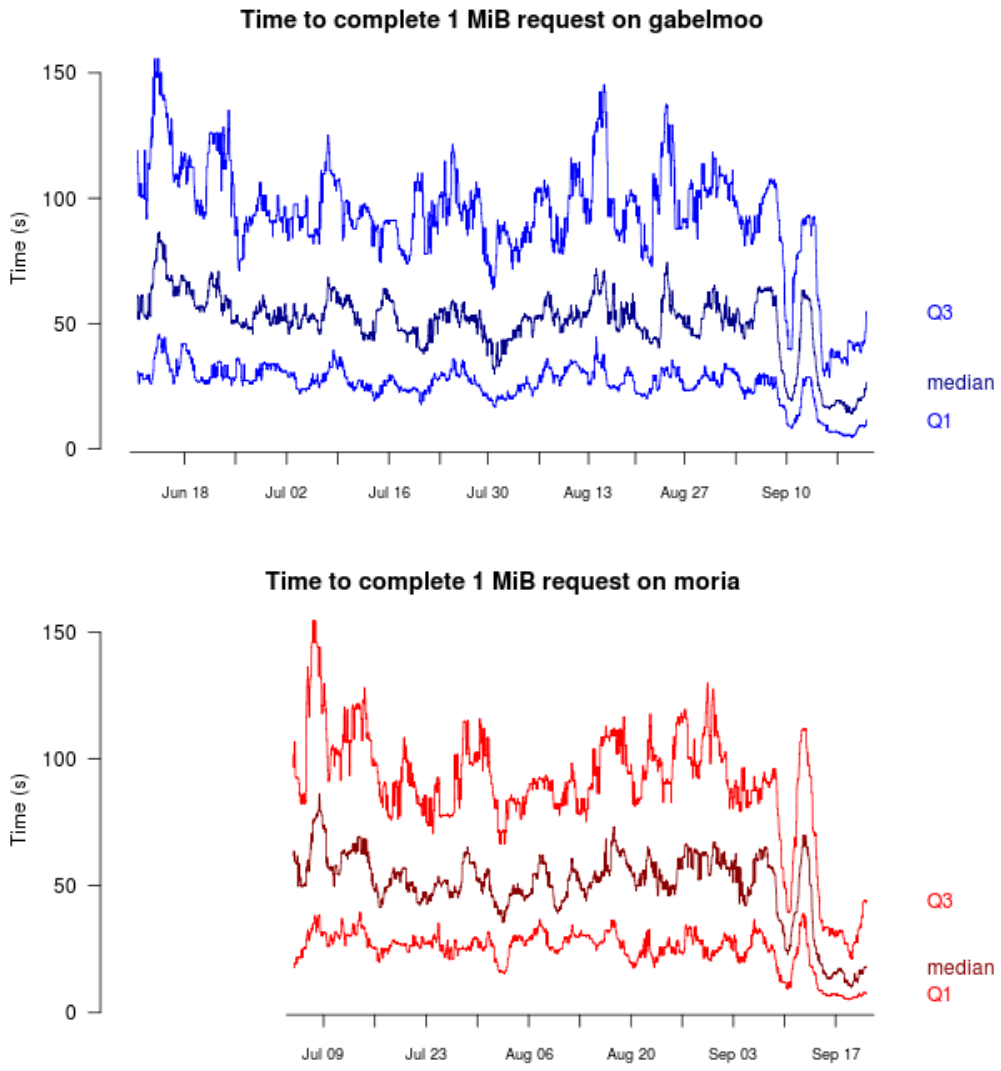

Figure 5: Time to complete 1 MiB request

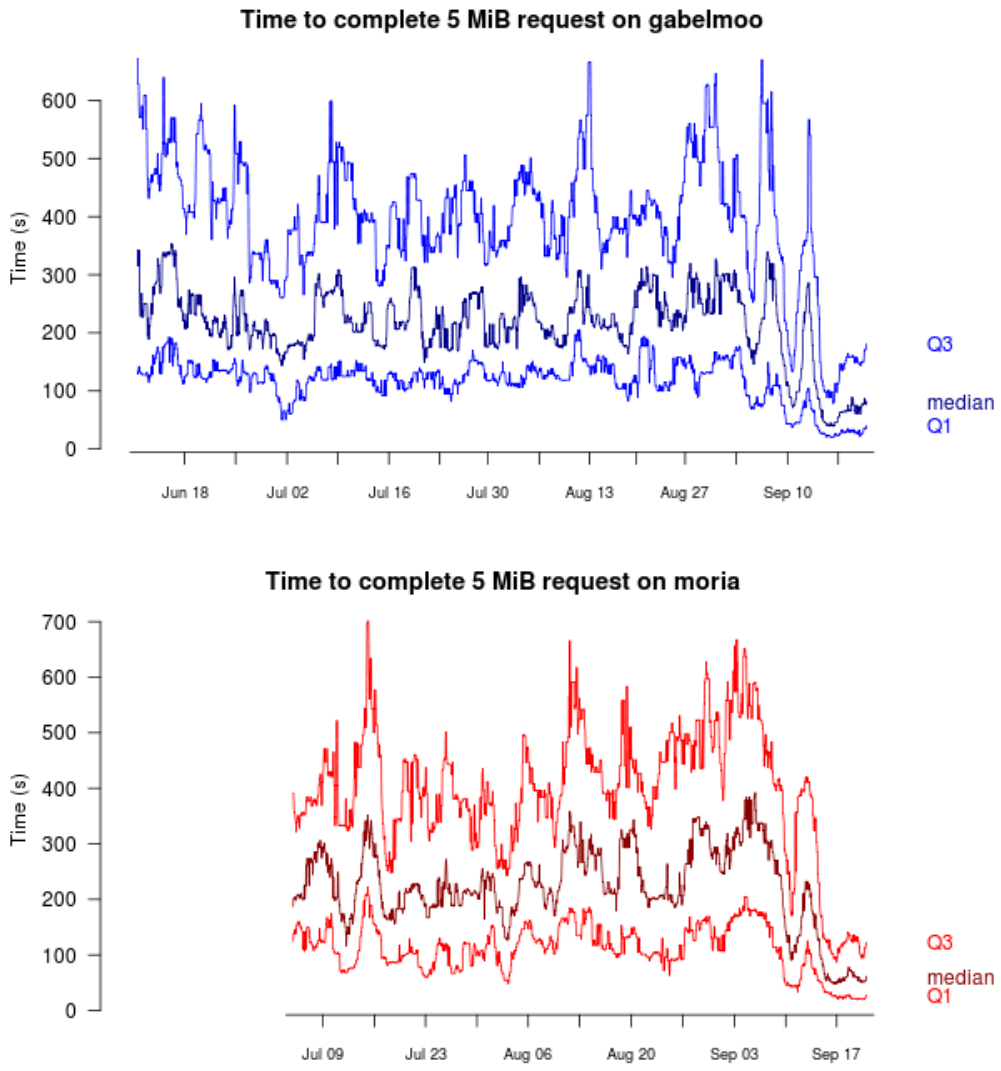

Figure 6: Time to complete 5 MiB request

4 Discussion

The presented measurements of downloading files over the Tor network help us learn average latencies and throughput in the current Tor network. In the future, the gathered data can help us evaluate performance improvements.