



Sentiment Analyzer (ENG) v2.3

Performance Assessment Report

[Introduction](#)

[Summary](#)

[Content Samples](#)

[Sample for Small Content:](#)

[Sample for Medium Content:](#)

[Sample for Large Content:](#)

[Overall Assessment](#)

[Comparative Metrics](#)

[Hits Per Second](#)

[Response Time vs Request](#)

[Times vs Threads](#)

Introduction

We conducted a comprehensive load testing using Apache JMeter. The objective was to understand the engine's behavior and responsiveness under different payload sizes, specifically categorized by token dimensions: **small** (~40 tokens), **medium** (~75 tokens), **large** (>= 350 tokens), and a **mixture** of all these sizes.

Each test was run under the following conditions:

- Number of Threads (users): 100
- Ramp-Up Period: 20 seconds (this is the time taken to create all the threads)
- Duration: 300 seconds (each test ran for a total of 5 minutes)
- Loop Control: Infinite loop (the test was set to execute the requests continuously for the entire duration)
- Scheduler: Enabled (the test was bounded by the specified duration)

The server-side environment for these tests was hosted on a machine powered by an **11th Gen Intel(R) Core(TM) i9-11980HK @ 2.60GHz with a boost up to 3.30 GHz, featuring**



ZA Cloud SRL, str. Govora 16A, 400664
Cluj-Napoca, România, Tel: +40 723 175
635, Email: office@zettacloud.ai

16 cores and 32GB of RAM, running the Docker environment on a Windows 11 Pro with WSL2 enabled.

The subsequent sections will delve into the detailed performance metrics gathered from each test run, offering insights into throughput, response times, and system robustness under different scenarios.

Summary

The English- language Sentiment Analysis engine, part of the out-of-the-box Intellidockers text analytics suite (<https://zettacloud.ai/ai-factory/intellidockers/>) is able to reach consistent throughput of **567 requests per second, equivalent to 1.45 Million words per minute**, on a commodity **CPU-only environment** with 16 CPU cores and 32 GB of RAM that costs less than \$4000.

Details on the functionality of the Sentiment Analysis engine is available here:
<https://zettacloud.ai/sentiment-analysis/>

Content Samples

Sample for Small Content:

It's decent for the price point. It's just okay, not great or bad.\n It's an average product. Does what it's supposed to do.\n It's decent for the price point. The product meets the basic requirements. It's just okay, not great or bad.

Sample for Medium Content:

The support team was very helpful. Impressive performance and sleek design. The build quality is top-notch. The build quality is top-notch.\n I absolutely love this product.\n The build quality is top-notch. Works flawlessly and integrates well. The build quality is top-notch. It exceeded my expectations.\n One of the best purchases I've made. Works flawlessly and integrates well. The build quality is top-notch. Works flawlessly and integrates well. I absolutely love this product. It exceeded my expectations.\n

Sample for Large Content:

Neutral experience overall.\n I have no strong feelings about it. It's an average product. It's just okay, not great or bad.\n The product meets the basic requirements. Neutral experience



overall. Neutral experience overall. Does what it's supposed to do. It's decent for the price point. Does what it's supposed to do. I have no strong feelings about it. Does what it's supposed to do. Does what it's supposed to do. \n It's decent for the price point. \n It's decent for the price point. \n Does what it's supposed to do. \n I have no strong feelings about it. \n It's an average product. Does what it's supposed to do. It's decent for the price point. \n It's an average product. \n Does what it's supposed to do. \n It's just okay, not great or bad. It's just okay, not great or bad. It's just okay, not great or bad. Does what it's supposed to do. The product meets the basic requirements. I have no strong feelings about it. \n It's decent for the price point. \n It's just okay, not great or bad. Does what it's supposed to do. The product meets the basic requirements. \n It's decent for the price point. It's an average product. It's an average product. Does what it's supposed to do. It's an average product. Neutral experience overall. I have no strong feelings about it. The product meets the basic requirements. \n I have no strong feelings about it. The product meets the basic requirements. I have no strong feelings about it. The product meets the basic requirements. Does what it's supposed to do. Does what it's supposed to do. It's an average product. It's decent for the price point. Does what it's supposed to do. \n It's an average product. \n It's decent for the price point. It's decent for the price point. \n I have no strong feelings about it. It's just okay, not great or bad. I have no strong feelings about it. Neutral experience overall. It's an average product. \n Neutral experience overall. Neutral experience overall. Does what it's supposed to do. The product meets the basic requirements. The product meets the basic requirements. \n It's just okay, not great or bad. \n Neutral experience overall. It's just okay, not great or bad. Does what it's supposed to do. It's decent for the price point. \n Does what it's supposed to do. \n I have no strong feelings about it. It's an average product.

Overall Assessment

1. **Response Time:** The tests with small and medium token sizes have very similar average response times, around **162ms** and **163ms** respectively. The test with large tokens has a slightly higher average response time of around 178ms, which is expected given the larger payload size. The mixed test had an average response time of **169ms**, falling between the small/medium and large tests.

2. **Throughput:** The throughput is highest for the small tokens test (589.05 requests/sec) and decreases slightly for the medium tokens test (585.34 requests/sec). The throughput for the large tokens test is noticeably lower at 536.68 requests/sec, likely due to the increased processing time for larger payloads. The mixed test has a throughput of **569.76**



requests/sec, which is between the small/medium and large tests.

3. Errors: All tests executed without any errors, indicating the robustness and stability of the system under different payload sizes.


In summary, the system performs well across different content sizes, with a slight increase in response time and a decrease in throughput as the token size increases. The system also proved to be robust, handling all payloads without any errors.

Comparative Metrics

Metric	Small	Medium	Large	Mixed
Average Response Time (ms)	161.98	163.33	177.88	169.07
Throughput (requests/sec)	589.05	585.34	536.68	569.76
Error Count	0	0	0	0

Hits Per Second


The 'Hits Per Second' chart visualizes the rate at which the system received requests over the duration of the test. A consistent high rate indicates stable throughput, whereas fluctuations can indicate variable system demand or potential bottlenecks.





Response Time vs Request

The 'Response Time vs Request' chart illustrates the relationship between individual requests and their respective response times. This helps in identifying patterns or anomalies, such as specific requests that consistently take longer to process.



Times vs Threads

The 'Times vs Threads' chart showcases how the response times vary with the number of active threads. It provides insights into how the system scales and responds as the load increases, which is crucial for understanding concurrency behavior.

