



BARROCO
CONSULT

QUALITY DIAGNOSIS

Diagnosis report

QuickPizza (Grafana public load-testing environment)

Demonstration report

October 18, 2026
Date

4
Pages reviewed

Desktop + Mobile
Devices

Performance test
public load environment

Executive summary

We ran a three-stage performance test with k6: a smoke test with 1 virtual user, a light load ramping to 25 users over a few minutes and a short spike to 55 users. We measured latency (p50, p95, p99), throughput, error rate and the slowest endpoints, and assessed the front with Lighthouse. The result is solid: no request failed and latency stayed stable from smoke to spike, with no relevant degradation point. The only attention point is the initial mobile load.

Findings by severity



Top risks

- 1 Security headers missing** MEDIUM
Raises exposure to interception, clickjacking and content injection.
- 2 Initial mobile load could improve** LOW
On mobile, the first visit takes a bit longer to become usable.
- 3 The generation endpoint is the slowest** LOW
Under high traffic, this is the endpoint most likely to become a bottleneck.

What is already good

- No request failed: 0% error across 15,928 requests
- Latency stable from smoke to the 55-user spike (p95 around 124 ms), with no degradation point
- Throughput of about 69 requests per second
- Very fast front on desktop (largest content in 0.8 s)

Measurements

home page (Lighthouse)

Metric	Desktop	Mobile	Good up to
LCP	800 ms	3.6 s	2.5 s
CLS	0	0	0.1
TBT	0 ms	0 ms	200 ms

Lab measurement: real visitor experience may differ.

Behavior under load

69
req/s
Throughput


0%
Error rate

25 ms
p50


124 ms
p95

140 ms
p99

55 VUs
Peak



Latency (p50/p95/p99) across the test.



Virtual users: smoke, ramp to 25, short spike to 55 and recovery.

Latency by endpoint

Endpoint	p50	p95	p99	Max
GET /	25 ms	26 ms	30 ms	297 ms
POST /api/pizza	93 ms	139 ms	150 ms	400 ms
GET /api/quotes	25 ms	29 ms	36 ms	304 ms
GET /api/names	25 ms	30 ms	38 ms	283 ms

Lab measurement in a public environment built for load testing; the values do not represent a client's production infrastructure. 15,928 requests across 3,982 iterations, with no errors.

Findings by category

Performance

2

Initial mobile load could improve

LOW

On mobile the largest content appears at 3.6 s; on desktop at 0.8 s. The other loading indicators are good.

Business impact: On mobile, the first visit takes a bit longer to become usable.

Recommendation: Optimize images and prioritize the main content for the first mobile load.

PERF-FRONT-MOBILE · Where: home page (mobile) (Mobile) · Occurrences: 1 · Effort: Medium · Evidence: LCP 3.6 s on mobile (lab measurement) (threshold good LCP ≤ 2.5 s)

The generation endpoint is the slowest

LOW

The endpoint that builds the pizza responds at p95 139 ms, about four times the read endpoints (~27 ms). This is expected for this endpoint, but it is the point to watch under heavier load.

Business impact: Under high traffic, this is the endpoint most likely to become a bottleneck.

Recommendation: Track this endpoint in larger load tests and consider caching or optimization if volume grows.

PERF-ENDPOINT-PIZZA · Where: POST /api/pizza · Occurrences: 1 · Effort: Medium · Evidence: p95 139 ms vs ~27 ms on the read endpoints

Security headers

1

Security headers missing

MEDIUM

Responses carry no HSTS, Content-Security-Policy, X-Content-Type-Options, X-Frame-Options or Referrer-Policy; one cookie is sent without the Secure attribute.

Business impact: Raises exposure to interception, clickjacking and content injection.

Recommendation: Add these headers at the server or CDN and mark all cookies as Secure.

```
HTTP 200 https://quickpizza.grafana.com/
content-type: text/html; charset=utf-8
set-cookie: AWSALB=****; Path=/ (one cookie without Secure)
strict-transport-security: (absent)
content-security-policy: (absent)
x-content-type-options: (absent)
x-frame-options: (absent)
referrer-policy: (absent)
```

SEC-HEADERS · Where: all responses · Occurrences: 5 · Effort: Low

30-day plan

Week 1 Fix what blocks usage and sales	Week 2 Close the remaining high risks	Week 3 Quality improvements • Security headers missing	Week 4 Prevent regressions • Regression load tests in the pipeline • Performance budget for the mobile front
--	---	--	--

Limitations of this diagnosis

- Test run on Grafana's public environment, built for load testing; values do not represent a client's production
- Light load and a short spike (up to 55 virtual users), without prolonged stress
- Front measured in a lab; real experience may differ

Checks not run

- Prolonged stress test: outside the demo scope and respectful use of the public environment

Next steps

- On a real product, repeat with load aligned to expected traffic and agreed latency targets
- Monitor the generation endpoint under larger volumes

Suggested package

40h

hours per month

A demo application with solid performance under light load and a short spike; no relevant problem.

Continuous performance monitoring and regression load tests.

Pricing on request.

Shall we fix these together?

Book a 30-minute call with the Barroco team: we walk you through every finding and build the plan for your product.

barrocoqa.com · contactus@barrocoqa.com

[Book a call](#)

Diagnostic produced with AI assistance and validated by a human analyst. Sample report, for public distribution.

