

# Active Benchmarking RGW

Marcel Lauhoff | Staff Software Engineer  
October, 2025

## Active Benchmarking

With *active benchmarking*, you analyze performance while the benchmark is still running (...), using other tools. You can confirm that the benchmark tests what you intend it to, and that you understand what this is. Data becomes Information.

<https://www.brendangregg.com/activebenchmarking.html>

# Stack

1. Workload Generator
2. Observability Stack
3. System Under Test

k6

Prometheus + Grafana

mtail, ceph-exporter


RGW

# k6


Run JavaScript  
in a loop  
with instrumentation  
at scale

# Let's Look at Dashboards

## Load: Light -&gt; Marginal -&gt; Saturated




## Performance Overview



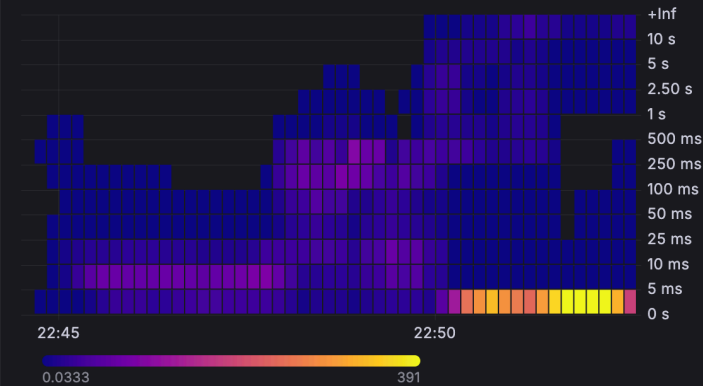
### K6 HTTP Req Duration




### HTTP Latency Stats




### RGW Beast Latency




### RGW Op Latency (mtail rgw\_op\_latency)



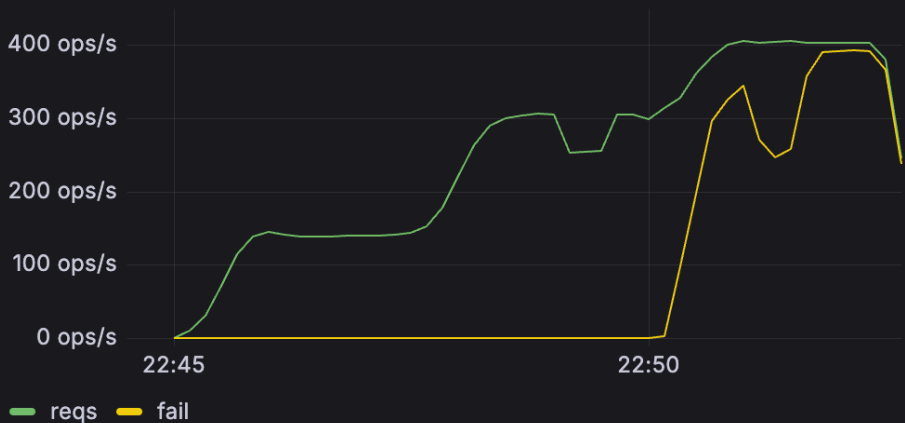
### RGW Req In Flight




### Latency Histogram All Requests




### Req Rate



### HTTP Status Rate (mtail rgw\_ops\_status\_total)



### HTTP Verb + RGW Error (mtail rgw\_ops\_status\_total)



# Active Benchmarking

**Do it Live:** Don't wait for the report. There is real value in watching live dashboards as the system crosses from healthy, to constrained, to saturated.

Averages Hide the Truth.

Native RGW Histograms.

Thank you!

Marcel Lauhoff <[marcel.lauhoff@clyso.com](mailto:marcel.lauhoff@clyso.com)>

# Bonus

# mtail

```
histogram rgw_beast_latency_seconds buckets 0.005, 0.01, 0.025, 0.05, 0.1, 0.25, [...] by verb
histogram rgw_beast_content_length_bytes buckets 0, 1024, 2048, 4096, 16384, [...]
counter rgw_beast_ops_total by verb, http_status
counter rgw_beast_user_agents_total by verb, http_status, user_agent
counter rgw_beast_addr_total by verb, http_status, addr
```

```
/beast: \S+: (?P<addr>\S+) - \S+ .* \?"(?P<verb>[A-Z]+) \S+ \S+\?" (?P<http_status>\d+) (?
P<content_length>\d+) - \?"(?P<user_agent>[^\\"]*)\?" - latency=(?P<latency>\d+\.\d+)s/ {
  rgw_beast_latency_seconds[$verb] = float($latency)
  rgw_beast_content_length_bytes[$verb] = int($content_length)
  rgw_beast_ops_total[$verb][$http_status]++
  rgw_beast_user_agents_total[$verb][$http_status][$user_agent]++
  rgw_beast_addr_total[$verb][$http_status][$addr]++
}
```

## k6

Run JavaScript  
in a loop  
with instrumentation  
at scale

```
export default function () {  
  http.get('https://ceph.com');  
}
```


k6

Run JavaScript  
in a loop  
with instrumentation  
at scale

<https://github.com/grafana/k6-operator>

## k6

# Run JavaScript in a loop with instrumentation at scale



Source: <https://grafana.com/docs/k6/latest/testing-guides/test-types/>

# k6

## Run JavaScript in a loop with instrumentation at scale

[...]

```
checks_total.....: 90    13.122179/s  
checks_succeeded.....: 100.00% 90 out of 90  
checks_failed.....: 0.00%  0 out of 90
```


### HTTP

```
http_req_duration.....: avg=140.36ms  min=119.08ms med=140.96ms max=154.63ms p(90)=146.88ms p(95)=148.21ms  
{ expected_response:true }.....: avg=140.36ms  min=119.08ms med=140.96ms max=154.63ms p(90)=146.88ms p(95)=148.21ms  
http_req_failed.....: 0.00%  0 out of 45  
http_reqs.....: 45    6.56109/s
```

[...]

k6

# Run JavaScript in a loop with instrumentation at scale



Stream  
metrics live to  
Prometheus  
and others