



Load Testing and Benchmarking EOS

Andrej Čop; Supervisor: Emmanouil Bagakis

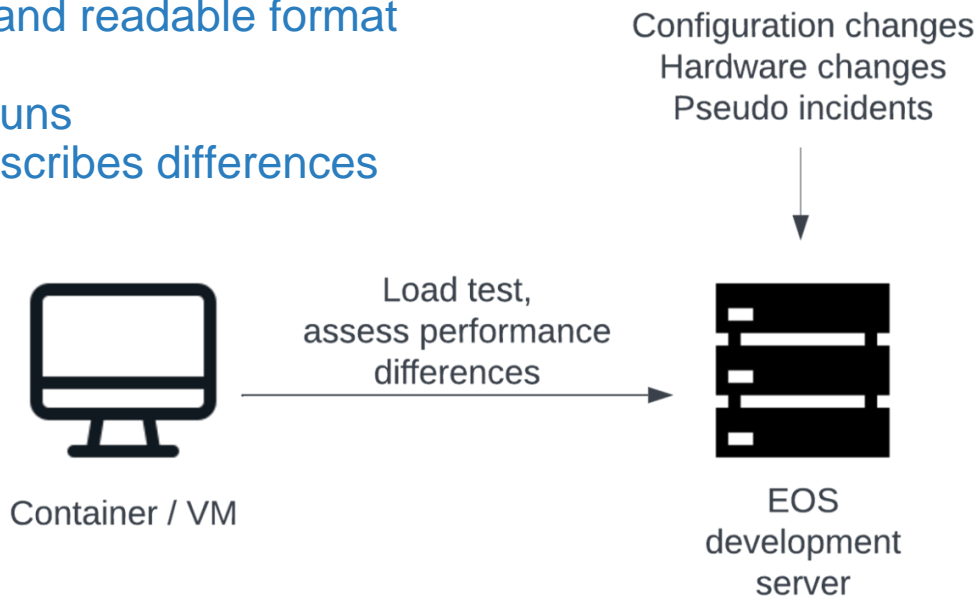
CERN Openlab summer student project 2023

IT-SD-GSS EOSHPM

Purpose of this project

Develop easy to use modular tool that can:

- Load test and assess performance of EOS
- Display metrics in a simple and readable format
- Compare metrics between runs
- Create output report that describes differences

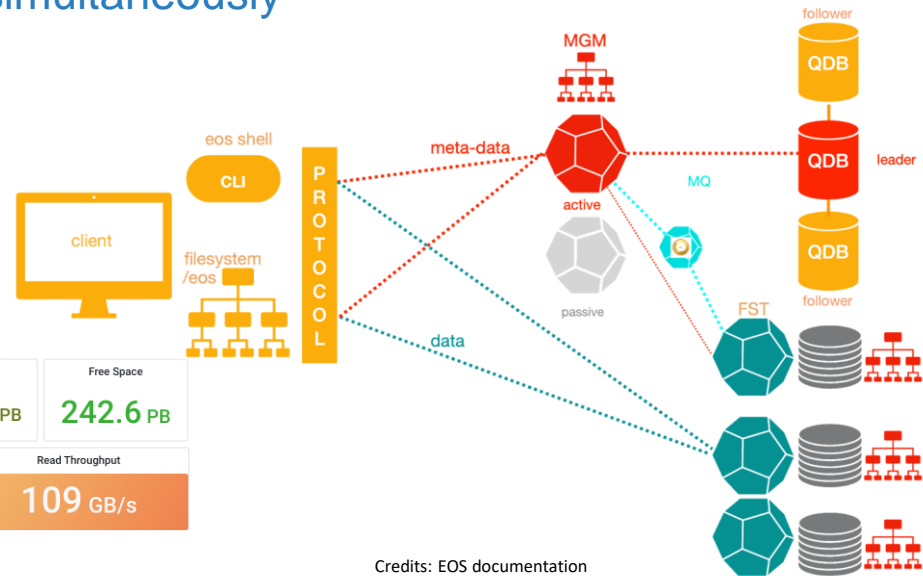


EOS Open storage

- Service for storing large amounts of physics data and user files
- Developed by CERN
- Supports thousands of users simultaneously
- Hierarchical namespace /eos
- Petabytes of data

Number of Files	Number of Directories	Total Space	Used Space	Difference	Free Space
7.87 Bil	643 Mil	774.68 PB	532.09 PB	-12.50 PB	242.6 PB
Current Writers	Current Readers	IOPS	Write Throughput	Read Throughput	
8.19 K	46.6 K	417 K io/s	28.7 GB/s	109 GB/s	

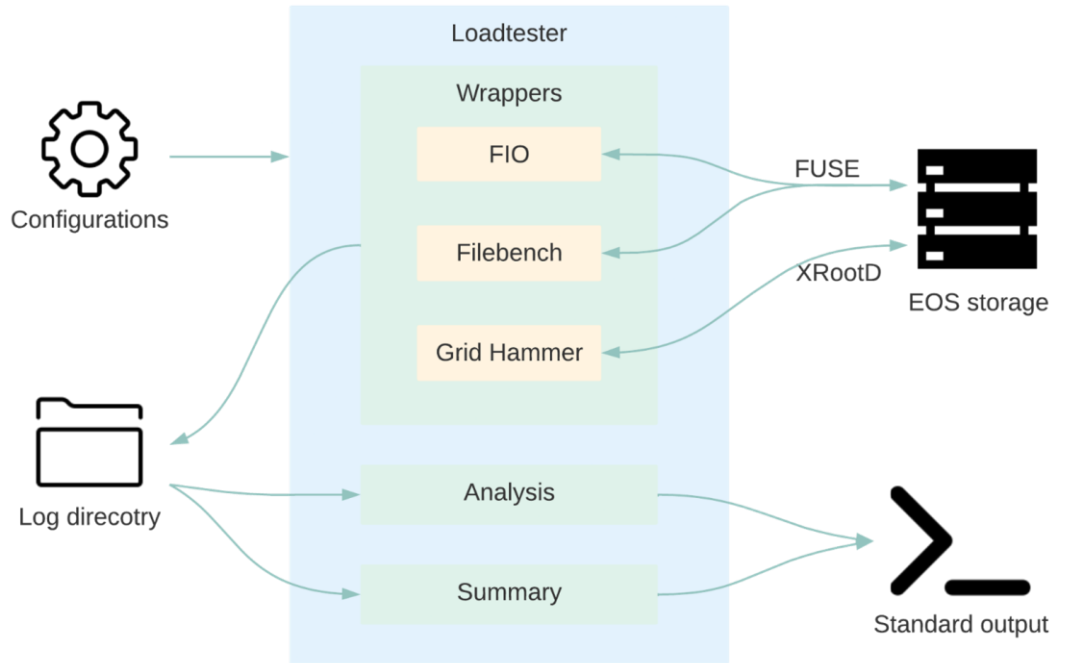
Data collected on 14. 8. 2023



Credits: EOS documentation


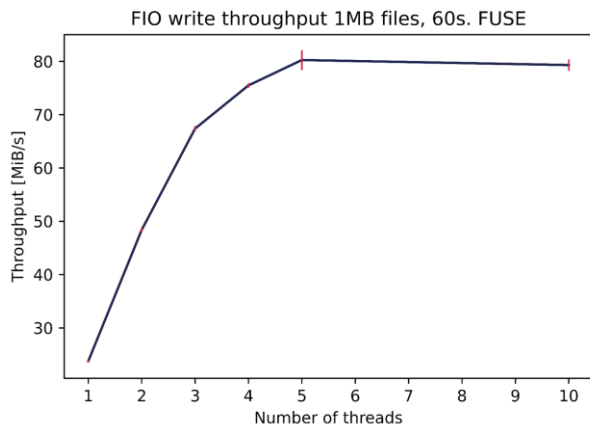
Toolkit and development

- Wrappers over existing tools
- Analysis software
- Automation scripts
- Organized log storing
- Summary report



Results

- Already proven useful with configuration changes



Filebench changes:

Read throughput: 0.54%
 Write throughput: 0.54%
 Read latency: 5.56%
 Write latency: -4.55%

Fio average changes:

Read throughput: 0.42%
 Write throughput: 0.40%
 Read latency: 0.19%
 Write latency: 0.55%

Hammer average changes:

Read rate: 2.49%
 Write rate: 2.27%

Challenges and errors

Open files limit

- default on VM was too low (1024)

```
1076 operations failed out of 2000.
```

```
[ERROR] Error response: No route to host
```

```
[ERROR] Error response: Numerical argument out of domain
```

DNS 114 requests/sec

- dns caching/dnsmasq

Results inconsistencies, big variance

- time based tests
- containerized environment
- multiple runs

Thank you!

andrej.cop@cern.ch

[linkedin.com/in/andrejcop](https://www.linkedin.com/in/andrejcop)

"The only source of knowledge is experience." ~ Albert Einstein