


HPCKP Annual Meeting 2023

# Fusion filesystem



Jordi Deu-Pons

"Moving millions of colossal files from one data storage to another is like herding a herd of hyperactive elephants through a maze of microscopic mouse holes."

**ChatGPT**



# What is Fusion?

- Is a **FUSE filesystem** mounted **inside each container** that gives access to different network data sources and uses a **local file system as a cache** to optimize network data transfers.
- Fusion is tightly integrated with **Nextflow**, a workflow manager widely used to run analysis that involves reading, creating and moving big files.
- Support multiple **data storages**. Currently; S3, Google Storage and Azure Storage.



# Main Fusion features

- Mounted inside the container (on user namespace)
- Local disk cache
- Intra session consistency
- Predictive download
- Asynchronous upload
- Decouple process execution from the data transfer



# Why do we develop Fusion?

Cloud reasons

- Remove staging times (moving files from S3 to local filesystem)
- Simplify infrastructure deployment (no need for EBS auto-scale, FSx, EFS...)
- Increase performance (AWS S3 is fast, NVMe is fast)
- Reduce costs



# Why do we develop Fusion?

HPC reasons

- Bring same user experience on HPC and cloud
- Reduce remote storage pressure
- Get good performance without using parallel file systems
- Simplify infrastructure



# Benchmarks

Running nf-core RNAseq full test pipeline

	<b>WALL TIME</b> hours	<b>BILLABLE</b> CPU hours	<b>RUNTIME</b> CPU hours	<b>STAGING</b> CPU hours	<b>COMPUTE</b> cost run	<b>STORAGE</b> cost month
<b>AWS S3</b> staging all files	11.25h	766.63h	545.23h	221.39h	\$99.05	\$24.56
<b>AWS FSX for Lustre</b> using local scratch	7.19h	406.38h	378.30h	28.08h	\$109.21	\$203.86
<b>AWS FSX for Lustre</b> without scratch	6.86h	358.97h	358.55h	0.42h	\$115.02	\$218.45
<b>Fusion</b> with local NVMe	4.66h	346.10h	342.91h	3.18h	\$49.08	\$19.92

\* All benchmarks were run on AWS Batch on one region using same EC2 on-demand instance families to get comparable results.

**More details:** <https://segera.io/blog/breakthrough-performance-and-cost-efficiency-with-the-new-fusion-file-system/>



# Room for improvement

Future Fusion

- Improve cache garbage collection efficiency
- Collect metrics within the task context
- Transparent compression
- Augmented security
- Support more remote storages








**Jordi Deu-Pons**

Senior Software Engineer

[jordi@seqera.io](mailto:jordi@seqera.io)

Barcelona



# Thank you