



Contribution ID: 6

Type: **not specified**

## Comparison of High-Performance Distributed File Systems on two Platforms: Linux and Windows

*Friday 15 March 2024 14:45 (15 minutes)*

Subtitle: Superiority of EOS-based Comtrade Distributed File System (CDFFS) for Earth Observation Data Storage.

### Introduction

The rising quantity of data collected necessitates transitioning to the next generation of reliable, high-performance data storage solutions. Despite the clear need for high-performance storage across many industries, there has yet to be a consensus on the optimal high-performance storage technology that would suit all user needs.

The project's goal was speed-testing to provide a possibility to combine them with the currently available state-of-the-art technology results, technological capabilities and existing expertise.

### Methodological approach

We choose representative high-performance storage solutions compared on two different platforms for storage clients.

- Linux: - Ceph - EOS - IBM Spectrum Scale - Hadoop
- Windows: - Ceph - EOS-drive - EOS through Samba - EOS-wnc - Hadoop

An example of specifications of the testing environment is the following.

- *Management node:* - 32 threads - 2x Intel® Xeon® Silver 4208 Processor - 384 GB of RAM - 2x SSDs of 2 TB.
- *Storage nodes:* - 32 threads - 2x Intel® Xeon® Silver 4208 Processor - 64 GB of RAM - 1x SSD of 2 TB for the operation system - 6x HDDs of 2 TB for data.
- *Client nodes:* - 12 threads - 1x Intel® Core™ i5-12400 Processor - 16 GB of RAM - 1x SSD with 1 TB.

### Results

The results are from the testing performed separately in an isolated environment for each high-performance solution but on the same hardware. For each platform, there were three categories of tests related to file sizes: (1) small, (2) medium, and (3) large files.

According to the results, some high-performance file systems have evident advantages, as shown in our presentation. These results should be the starting point for an even more exact comparison between these file systems. They are good starting points in choosing the right high-performance file system.

**Primary author:** MOLAN, Gregor (Comtrade 360's AI Lab)

**Presenter:** MOLAN, Gregor (Comtrade 360's AI Lab)

**Session Classification:** EOS Operations