

Mistral is an I/O profiling tool that runs at the application level. It is lightweight enough to be deployed to your whole cluster to measure performance and to locate rogue jobs with bad I/O patterns. The analytics from Mistral are designed to be combined with the data you get from your file system to give a whole-system overview.

It's important to maximise throughput on your shared storage in order to maintain good system performance, but it's the application performance and job throughput that really count. By benchmarking storage at the application level you can see the impact of each action on the jobs as they run.

Sampling I/O latency with Mistral

As well as measuring how many reads, writes or meta data operations an application performs, Mistral can record or sample I/O latency. Mistral can alert you if the mean, peak or total latency goes over a limit that you set. This simple measure is enough to tell you everything you need to know.

How does the performance of the file system vary over time?

You can use regular jobs or specific benchmark applications to compare performance at different times of the day or on different days of the week.

What is the impact of file system performance on my applications?

Often users complain about slow performance and you see slow I/O access times on your file system, but these are not always connected. Mistral gives you per-application I/O latency data so you can see the impact of slow I/O on that application. It might only be a small percentage of run time.

Meta-data performance

Mistral lets you measure the I/O latency of several classes of I/O operations: open, create and delete, as well as stat and file-system-change operations such as move or rename.

Small reads and writes

You can look at the performance of reads and writes in general, but it's important to also measure the size of the operation. Large reads or writes might be high-performance, but take a long time for a single operation. Similarly, you might see lots of small reads and writes served quickly, but the overall performance is really low. Poor performance is often blamed on bad I/O patterns such as small reads and writes. With Mistral you can measure the performance of those small reads and writes and compare that with the performance metrics of the file system to gauge the impact of those I/O patterns on the storage and the resulting impact on the performance of other jobs on the cluster.

