

# PROTECTING SHARED STORAGE FROM BAD I/O PATTERNS WITH ALTAIR MISTRAL™

Enterprise workload managers for high-performance computing (HPC) efficiently allocate jobs and orchestrate CPU, memory, license, and GPU sharing — but they're not designed to control the workloads' I/O patterns, an area that may be overlooked when tuning HPC environment performance. It's easy for a job with bad I/O patterns to overload shared storage, which commonly happens when a user has tested a workflow on one or two nodes, then scales the workload without understanding how the I/O patterns scale.


## Early Detection Enables Quick Action

Finding out immediately which job and processing stage are causing a problem is key to solving it before it gets out of hand. The problem application may be written in-house or it could be a third-party tool that has been deployed inefficiently. We call applications that have bad I/O patterns “rogue jobs” or “noisy neighbors.”


HPC infrastructure with critical execution time constraints (e.g., HPC for patient-specific genomics medicine, weather prediction for airports, high-speed trading in financial markets) can all see value from this kind of protection against downtime and severely reduced throughput. As distributed systems and compute clusters become more complex, the need for monitoring is increasingly important. To run compute efficiently and to forecast and design for the future, it's important to have visibility into what's being run today.

## Altair Mistral – Per-Job Monitoring for I/O, CPU, and Memory

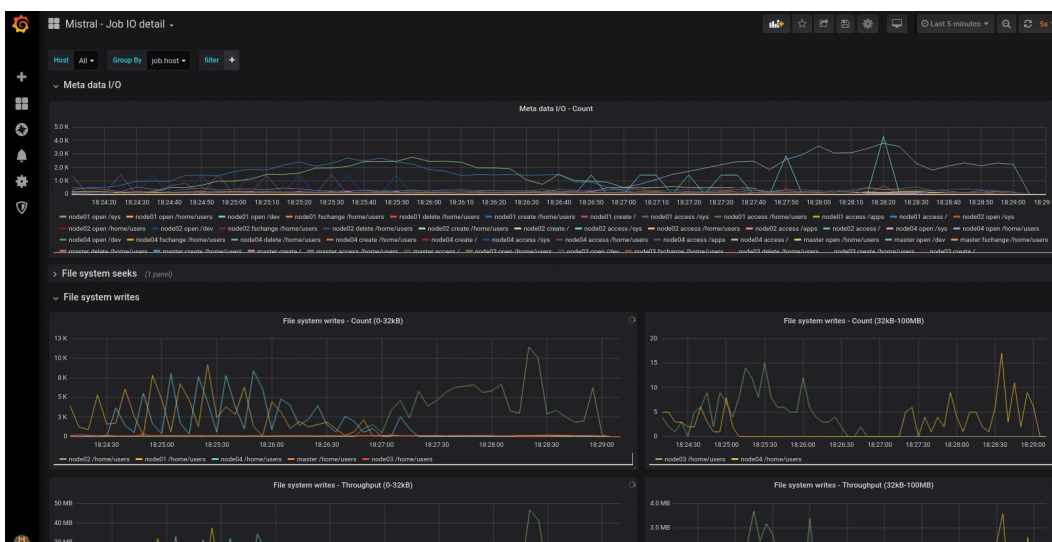
[Altair Mistral](#) is the leading application monitoring tool for HPC and scientific computing, with the unrivaled ability to track I/O patterns across an HPC cluster. Altair Mistral monitors I/O, CPU, and memory, quickly locating rogue jobs and storage bottlenecks and keeping track of what's running on clusters day-to-day.



Pinpoint problem jobs with bad I/O patterns



Act quickly to solve problems that affect performance



Altair Mistral is lightweight enough to run in production, and flexible enough to ensure that you get the most from on-premises HPC and have all the information you need to manage your hybrid cloud.



### Find Noisy Neighbors that Overload Shared Storage

In HPC systems, it's easy for a single user or application to overload shared storage with bad I/O. Applications are often tested on only one machine, and their I/O patterns don't scale well to another. This can harm file system performance and slow down an HPC cluster for all users.

Altair Mistral makes it easy to find these noisy-neighbor applications. It shows drops in file system performance, overall load on the file system, and which jobs are contributing to it. Within seconds, you can identify the user, job, and host causing the problem and can act quickly to solve it.



Quickly identify the user, job, and host causing an I/O problem

### Get Insight into I/O Patterns that Harm Shared Storage

System- and storage-agnostic with a lightweight footprint for production monitoring, Altair Mistral livestreams read, write, and metadata use, as well as storage performance for every job and mount point or file system. It delivers per-job I/O and compute metrics and helps you understand storage requirements for procurement and planning. In all, Altair Mistral:

- Shows I/O per job, user group, and host
- Shows I/O per mount point
- Tracks read(), write(), and metadata operations such as open() and delete()
- Tracks file system performance, looking at the time spent in I/O for each application
- Tracks the total, mean, and max time spent doing I/O per job
- Tracks CPU and memory utilization

### Easy Integration into Your HPC Environment

Altair Mistral is usually integrated with the scheduler and is injected into each job through a prologue or starter method. It comes with sample integration for Altair® Grid Engine®, Altair® Accelerator™, and Altair® PBS Professional®, as well as third-party workload managers like LSF and Slurm. You can also run Altair Mistral on login nodes and other machines without a scheduler. Because it's scheduler- and system-agnostic, you can use a single solution for both on-premises and cloud analytics.

Altair Mistral integrates with existing monitoring data pipelines so you can combine Altair Mistral data with the metrics from your scheduler or file system. It works out-of-the-box with Elasticsearch, InfluxBD, Grafana, and Splunk. Through an integration with Fluent Bit, Altair Mistral can push data to a wide range of databases and messaging systems.

For more information, please visit [altair.com/mistral](https://altair.com/mistral).