

Altair - Mistral Release Notes

Version 2025.2.0



Release 2025.2.0

Highlights:

- Changed the default Mistral configuration to include durations by default.

Other Changes:

- Mistral now measures the following additional libc functions: `splice` and `ustat`.
- Improved start up performance of Mistral.
- Attempting to read or write with a closed file descriptor is treated as I/O on an invalid mount point. #5813
- Fix bug when following binaries launched with GridEngine `qsub`. #6028
- Fix bug when following child jobs on GridEngine queues with `shell_start_mode` set to `posix-compatible`.
- `sendfile` now counts as a seek followed by a read and write. #5975
- Don't count `fdopen` and `fdopendir` as open calls if the open of the file descriptor has already been counted. #2349
- Reports failed I/O to invalid file descriptors. #5813
- Don't treat `fdopen` and `fdopendir` calls as failed just because they set `errno`. (#6052)
- New Mistral ElasticSearch Grafana dashboards. #5925
- Speeds up monitor exit when a plug-in is configured.
- Adds support for Intel and AMD GPU Profiling.
- Removed dependency on `trace-program.sh`.
- Moved container setup scripts to `sbin`. #6009
- Adds ability to disable monitoring for `su` and `sudo` #5746.
- Improve error messages for some common config file problems. #5300
- Added '`sbin/update-dashboard-timeframes.sh`' to automate changing the dashboard interval to match the Mistral timeframe. #6042
- Adds support for `https` as part of InfluxDB quick start. #6001.
- Upgrade version of ElasticSearch to 8.17.3 and Grafana to 11.5.1 in docker compose script.
- Fixed bug where 1s timeframes reported incorrect job cpu and memory usage. #5739.

Release 2025.1.0

Highlights:

- **DEPRECATION NOTICE:** Support for RHEL 5/CentOS 5 is planned to be removed in the next version of Mistral. Please contact support if this will cause you any problems.
- Adds GPU profiling for NVidia GPU's.
- Adds config options controlling Mistral output and licensing. It is no longer required to set environment variables to run Mistral; they are optional.

Other Changes:

- Added RPM and Debian packages for aarch64 architecture.
- Added documentation surrounding Singularity/Apptainer mounting of /tmp and/or /var/tmp.
- Fixed bug in monitor causing it to wrongly output an error when dealing with bind mounts.
- Fixed a problem where a multithreaded program could deadlock if the main thread exits without waiting for the other threads to complete.
- Mistral now measures the following additional libc functions: execveat, _Exit, fctnl64, fstatfs, fs-tatfs64, fstatvfs, fstatvfs64, futimens, preadv, preadv2, preadv64, preadv64v2, pwritev, pwritev2, pwritev64, pwritev64v2, quick_exit, statfs, statfs64, statvfs, statvfs64, utimensat and copy_file_range.
- Record I/O done by glob calls.
- Report the most appropriate mount point when a "bind mount" has been used to mount a device at more than one location.
- Documentation is now available in HTML format as well as PDF #6000.
- Mistral output now uses signed, rather than unsigned, 64-bit integers
- Fixed problem which could cause Mistral to report the wrong total number of reads.
- Fixed problem where environment variables that contained an equals sign were being incorrectly reported.
- Made the Mistral summary consistent with the total data.
- Upgrade version of Elasticsearch to 8.15.1 and Grafana to 10.4.8 in docker compose script.
- Supports Altair Units base licensing.
- Fixed bug that could cause Mistral to segfault if job did not connect with an opened socket.
- Avoid a delay at the end of a job when the Mistral working directory directory is on NFS.
- mistral.sh will report an error if the permissions on /tmp/MISTRAL would stop it from working correctly and this can't be automatically corrected.
- Added support for podman-compose when installing Elasticsearch/InfluxDB and Grafana.
- Add dashboard for network data to InfluxDB.
- Containers started with docker compose will now survive a reboot #5997.
- Grafana containers now use the self signed HTTPS certificate by default #6001.

Release 2024.1.1

- Now also intercepts `close_range` when made as an indirect syscall.

Release 2024.1.0

Highlights:

- Completely reworked directory heirarchy to be more in-line with the FHS /usr heirarchy recommendations.
- Created new RPM and debian packages of breeze, breeze-trace-only, mistral-healthcheck and mistral.

Other Changes:

- Add support for Apptainer. Renamed MISTRAL_SINGULARITY_BIND_PATH to MISTRAL_CONTAINER_BIND_PATH.
- No longer hides monitor PID in logging messages.
- Changed pwrite to no longer count as a seek when writing to a file that was opened in append mode to match Linux kernel behaviour.
- Fixed bug where hiding may not work as intended.
- Many binary names have been changed to include {mistral} to make them more identifiable.
- The monitor doesn't take so long to exit at the end of a job.
- Fixed the default path used by some exec functions to match that found in the running version of glibc.
- Correct the handling of zero length monitor status files.
- Fixed a problem where the monitor could be killed during startup. This resulted in directories being left that would normally have been removed when the monitor exits.
- Fix a problem where a multi-threaded application which runs qrsh could modify the environment in an unsafe way.
- Added docker-compose setup script for InfluxDB 2.
- Improved licensing logs to provide specific warning and error messages in stderr as well as the normal error log locations.
- Renamed the "summary" record, which represents all I/O done by a job, to "jobsummary".
- Added a "mountpointsummary" record, which summarises I/O on each mountpoint used in a job.
- Stopped pread/pwrite calls counting as a seek when doing sequential I/O.
- Avoid double counting file descriptors created by dup in a jobsummary record.
- Used VOV_PROJECT_NAME in addition to VOV_JOBID when identifying Accelerator jobs.
- Sampling of duration measurements is now controlled by entries in the configuration file. The environment variables which used to control this are still supported for backward compatibility. Mistral can also adjust the duration sampling parameters based on the measured overhead.
- Fixed problems with the jobrealtime, accumulatedruntime, accumulateddiotime and iotimepercentage fields in the jobsummary records.
- Renamed mistral script to mistral.sh.
- Added mistral.csh that can be sourced in csh environments.
- Fixed bug where some host cpu measurements were not reported.
- Fixed Grafana memory dashboards to show maximum rather than the sum of records.
- Fixed Grafana memory dashboards to show maximum rather than the sum of processes.
- Fixed job total memory metrics showing incorrect maximum.
- Fixed bug where there could be two mistral-monitor processes started for Altair Grid Engine jobs if not using an array.
- Added job run time and start and end timestamps to all job total records.

Release 2023.1.2

Highlights:

- Fixed issue with ElasticSearch Grafana dashboards not working with time ranges greater than a few hours.

Other Changes:

- Fixed bug in ElasticSearch Grafana dashboard where seeks were not shown.
- Fixed Grafana dashboards so that graph interval matches default Mistral timeframe (20s).
- Fixed Grafana dashboards where mean was totalled, rather than total.
- Fixed Grafana repeated dashboard flickering.

Release 2023.1.1

- Renamed ELLEXUS_DO_NOT_BACKTRACE to ELLEXUS_BACKTRACE and inverted related logic.
- Made backtrace output less confusing, removing libmistrail.so backtrace handler.
- Fixed usage messages from C shell scripts having wrong file extension when referring to self.
- Mistral network measures include connect and accept calls under metadata.
- Disabled libmistrail.so backtrace handler by default.
- Removed the requirement for MISTRAL_INSTALL_DIRECTORY to be set when using plugins with a relative path.
- Increased default memory for elasticsearch and kibana containers when using docker-compose script.
- Fixed bug where a segfault could be thrown for some non-blocking connect operations.
- Upgraded default Grafana (8.5.13->10.0.4) and ElasticSearch (7.16.2->8.8.2) version provided via docker-compose scripts.
- Fix Handling launching mistral_report.sh with \$MISTRAL_LOG as log input path, to cope with hostname being added to file path.

Release 2023.1.0

Highlights:

- Renamed shared object library from “libdryrun.so” to “libmistral.so”. “libdryrun.so” is now a sym-link to “libmistral.so”.
- Remove environment variable MISTRAL_PLUGIN_CONFIG and moved plug-in configuration into the same file as global Mistral configuration.
- Mistral will output JSON instead of CSV (including summary output).
- Mistral now measures and reports network I/O as well as file I/O.
- Renamed mistral.conf to mistral_config.yaml.

Other Changes:

- Now lazily loads libc functions.
- Reduced the amount of stack space used in programs being traced or metered.
- Changed handling of –child argument, now defaults to all on and command that launch child jobs can be removed from the list with - e.g. –child=-qrsh,-qsub.
- Fix problem where a process within Singularity could send information to the wrong monitor process.
- Fixes –relocate option archiving all traces, not just the most recent.
- Sending SIGUSR1 to the monitor process will now cause it to send status output to the error log.
- Removed the MPI implementation specific versions of libdryrun.so.
- Now supports pbsdsh and pbs_tmrsh jobs.
- Added wrappers for sendmmsg and recvmmsg Linux system calls.
- Add singularity-bind program to simplify setting the SINGULARITY_BIND environment variable.
- Mistral_start_reference.sh now describes the additional steps needed when Singularity is run directly, rather than as part of a script.
- Ensure shell builtins, e.g. echo are traced correctly if they are top level commands.
- Fixed bug when tracing csh commands on remote hosts.
- Fixed bug where top level-command could be interpreted by a shell.
- Fixed bug where user names could be truncated in trace directory names.
- Fixed bug for system(3) and popen(3) that could crash some applications which make use of RDMA devices.
- Fixed bug for closedir(2) that caused a segfault when passed NULL.
- Added wrapper for statx(2) Linux system call.
- Fixed bug where stderr was unintentionally open in the monitor, which could alter the behaviour of the Python subprocess module.
- Fixed problem where signal masking within libmistral.so could fail if a signal handler had been installed before libmistral.so was active.
- Fixed bug when tracing sudo that could cause a binary in the current working directory to be run.
- Fixed bug when tracing sudo that could prevent a shell from being used to interpret shebang scripts.
- Added CSV plug-in that outputs the old Mistral format.
- Allow a default mistral_config.yaml file if no config file is specified, or if there's an error parsing the specified config file.
- Added configuration file option waitpid to instruct Mistral which process in a job to wait for before exiting.
- Added support for YAML sequences in MISTRAL config file.
- Improved performance by reducing internal data bandwidth by a factor of 10.
- Fixed problem which could result in data not being reported.
- Fix problem sending long messages to a plugin.
- Fixed sample scheduler scripts to exclude scheduler commands from monitoring.

- Added script to simplify elasticsearch setup.
- Improved performance of multi-process jobs by reducing the number of times the config file is parsed.
- Added config option to specify whether data should go only to a plug-in or also to the MISTRAL_LOG file.
- The Mistral output now includes a number of rate measurements. For example, as well as the number of bytes read during a time interval, the mean, minimum, median and maximum number of bytes per second are also reported. The default time frame has been increased to 20 seconds.
- The time taken by a traced function is now referred to as a “duration” rather than a “latency”. Two configuration environment variables have been renamed accordingly. (MISTRAL_MONITOR_LATENCY_MAX_IO becomes MISTRAL_MONITOR_DURATION_MAX_IO and MISTRAL_MONITOR_LATENCY_SAMPLE becomes MISTRAL_MONITOR_DURATION_SAMPLE.)
- Fixed bug causing plug-in to report wrong environment variables values.
- Mistral no longer outputs a special “delayed” record if it receives data when it is too late to be included in the relevant time frame. Such data is however still included in the job total.
- Disabled remote tracing of scheduler programs by default.
- Fixed bug where monitor would get stuck in a loop on fast machines when running a large number of small processes.
- Added the number of CPUs used by a job to the output.
- Mistral timestamps now have a precision of one second.
- Mistral no longer assumes that a job has ended if it doesn't do any I/O during the first few seconds.
- Improve Singularity PBS Integration.
- Fixed problem where some calls could be left out of the summary record and others counted multiple times.
- Improve Docker PBS Integration.
- Samples durations for summary output.

Release 2022.2.1

- Reduced the amount of stack space used in programs being traced or metered.
- Fix bug setting strace output directory with remote traces.
- Improved performance by reducing internal data bandwidth by a factor of 10.
- Fix bug in Grafana dashboards where the ad-hoc filter has no data source and doesn't work by default.

Release 2022.2.0

Highlights:

- **Contracts have been removed.** By default, Mistral now automatically meters and reports all I/O on all mount-points, this can be changed with a new configuration file - the path to which is specified with `MISTRAL_CONFIG`.

Other Changes:

- The latency of I/O on unix domain sockets is no longer counted as file system latency.
- Fix problems seen on Fedora 35 and glibc 2.34.
- The error log command line option has been changed from `-log` to `-errlog`.
- Fix a problem which caused I/O on a pipe, or a file descriptor which is not associated with an inode, to be counted as I/O to a file in the current working directory.
- The start-monitor utility can now be run without setting `LD_PRELOAD`.
- Fix problem if Mistral is run with the path name of a command which creates a child job.
- Make the `-child-job` option work correctly when Mistral is used with Singularity.
- Fix problem when Mistral is used with an MPI job running on Altair Grid Engine.
- Fix problem where child jobs over-wrote user child job (`ELLEXUS_CHILD`) setting.
- Correctly identify mount points when running in a Singularity container.
- Fix problems shown by compiling with gcc version 12.1.
- Fix bug in path resolution triggered by flush in some circumstances.
- Fix bug where strace option wasn't preserved in child jobs.
- `ELLEXUS_OUTPUT_DIRECTORY` renamed `MISTRAL_OUTPUT_DIRECTORY` and is no longer changed whilst being used.
- Fix problem where a program calls `close_range` or `closefrom`.
- Add documentation of the integration with Altair Grid Engine 2022.2.
- The script `mistral.csh` was removed.
- The environment variable which represents the name of the error log has been changed to `MISTRAL_ERR_LOG`, though `ELLEXUS_ERR_LOG` is still supported for backward compatibility.
- Error log messages are sent to syslog unless a log file has been configured either by using the `-errlog` command line option, or by setting the `MISTRAL_ERR_LOG` environment variable.
- Make Grid Engine starter script work in binary mode
- Changed branding and improved structure of Healthcheck report.
- Configuration file added to allow more control over which data is reported.
- Fix problem where the mistral script could exit before the mistral log file had been written.
- Fix problem which could result in Mistral output being unintentionally delayed.
- Improved handling of unwritable user specified paths.
- Remove MPI-specific libraries (as all MPI implementations now use the same `libdryrun.so` library).
- `MISTRAL_LITE` mode is no longer supported as this has been superseded by the `MISTRAL_CONFIG` file.
- Add sample file for Altair Grid Engine 2022.2 integration.

Release 2022.1.3

- Fixed bug in path resolution triggered by flush in some circumstances.
- Fixed bug in argmangling where ELLEXUS_CHILD was not respected in child jobs.
- Fix problem when Mistral is used with an MPI job running on Altair Grid Engine.