

Altair - CSV Plug-in Configuration Guide



Contents

1	Installing the CSV Plug-in	3
2	Configuring Mistral to use the CSV Plug-in	4
2.1	Mistral Plug-in Configuration	4
2.1.1	Path	4
2.1.2	Interval	4
2.1.3	Options	4
2.2	Plug-in Configuration File Options	4
3	Plug-in output format	6
3.0.1	Log labels	7
4	Plug-in summary format	9
5	Example log entries	10
5.1	Example output log	10
5.2	Example summary log	10

1 Installing the CSV Plug-in

Extract the Mistral plug-ins archive that has been provided to you somewhere sensible. Please make sure that you use the appropriate version of the plug-ins for the architecture of the machine on which the plug-in will run.

In addition, if the Mistral Plug-ins package was obtained separately from the main Mistral product, please ensure that the version of the plug-in downloaded is compatible with the version of Mistral in use. Version v7.2.8 of the CSV Plug-in as described in this document is compatible with all versions of Mistral compatible with plug-in API version 7. At the time of writing this is Mistral 2023.1.0 to 2025.2.0.

The CSV plug-in can be found in, for example:

`<installation directory>/output/mistral_csv_v7.2.8/x86_64/`

for the 64 bit Intel compatible version, and in:

`<installation directory>/output/mistral_csv_v7.2.8/aarch64/`

for the 64 bit ARM compatible version.

The plug-in must be available in the same location on all execution hosts in your cluster.

2 Configuring Mistral to use the CSV Plug-in

Please see the Plug-in Configuration section of the main Mistral User Guide for full details of the plug-in configuration specification. Where these instructions conflict with information in the main Mistral User Guide please verify that the plug-in version available is compatible with the version of Mistral in use and, if so, the information in the User Guide should be assumed to be correct.

2.1 Mistral Plug-in Configuration

The Mistral plug-in configuration is in YAML and goes in the same file as the main Mistral Configuration File. The plug-in is declared with the `plugin` mapping and requires at minimum a `path` key-value pair. All of the specified settings are members of the `plugin` mapping.

```
plugin:
  path: plugins/mistral_csv/x86_64/mistral_csv
```

This section describes the specific settings required to enable the CSV Plug-in.

2.1.1 Path

The `path` key must be set to the path of the CSV plug-in executable. This must be either absolute or relative to the `MISTRAL_INSTALL_DIRECTORY` environment variable and needs to be accessible and the same on all hosts. Environment variables are not supported in the value.

```
path: plugins/mistral_csv/x86_64/mistral_csv
```

2.1.2 Interval

The `interval` key takes a single integer value parameter. This value represents the time in seconds the Mistral application will wait between calls to the specified plug-in e.g.

```
interval: 300
```

The value chosen is at the discretion of the user, however care should be taken to balance the need for timely updates with the scalability of the CSV installation and the average length of jobs on the cluster.

2.1.3 Options

The `options` mapping is optional and lists all options to be passed to the plug-in as command line arguments to the executable. A full list of valid options for this plug-in can be found in section [2.2 Plug-in Configuration File Options](#). The order of options is not preserved. These values are passed to the plug-in executable as `--key=value`. For example,

```
options:
  output: out_file
  error: err_log
```

will pass to the plug-in executable the command line arguments `--output=out_file` and `--error=err_log`.

2.2 Plug-in Configuration File Options

The following configuration file key-value options are supported by the CSV plug-in.

```
error: file
```

Specify location for error log. If not specified all errors will be output on `stderr` and handled by Mistral error logging.

```
output: file
```

Specify location for output log. If not specified all output will be output on `stderr` and redirected to the Mistral error log.

summary: file

Specify location for the job summary log. If not specified all output will be output on stderr and redirected to the Mistral error log.

mode: octal-mode

Permissions used to create the error log and output log files specified by the error and output options.

3 Plug-in output format

Each line of output contains the following fields:

```
<TIME-STAMP>, <LABEL>, <PATH>, <FS-TYPE>, <FS-NAME>, <FS-HOST>, <CALL-TYPE>,
<SIZE-RANGE>, <MEASUREMENT>, <MEASURED-DATA>/<TIMEFRAME>, 0/<TIMEFRAME>,
<HOSTNAME>, 0, 0, , , <JOB-GROUP-ID>, <JOB-ID>, , 0
```

Note that several of the comma-separated values are hard-wired empty or 0. These are for backwards-compatibility with previous log formats.

The field definitions are as follows:

<TIME-STAMP> is the end of the time frame when the I/O occurred. The time-stamp is in ISO 8601 format with microsecond precision (YYYY-MM-DDThh:mm:ss.ffffff).

<LABEL> is a textual label describing the category. See the [Log labels](#) section below for more details.

<PATH> is the I/O mount point reported by this log entry, or "*" if this entry is the total across all mount points. For "resource entries" (memory, memory-vsize, user-time, system-time, cpu-time, host-cpu-user-time, host-cpu-user-time, and host-cpu-iowait-time) this is not relevant and is given as /.

<FS-TYPE> is the filesystem type of <PATH>. It is empty for resource entries.

<FS-NAME> is the so-called filesystem "name" of <PATH>, typically a device name or an NFS HOST:PATH specification. This is empty for resource entries.

<FS-HOST> is the host name part of <FS-NAME>, if present. It is empty for resource entries.

<CALL-TYPE> is one of the call types read, write, open, access, create, delete, fschange, mmap; or metadata (for the total of the call types other than read and write). This field is not relevant for resource entries and is given as none.

<SIZE-RANGE> is <MIN>-<MAX> for size-binned read/write data, where <MIN> and <MAX> are the lower and upper operation sizes for the bin being reported, and are both powers of two. Note that the range is exclusive of <MAX>, so that (for instance) if <SIZE-RANGE> is 0-32KiB (as for small/medium/large size binning), then it will include operations up to and including 32,767 bytes, but operations of exactly 32 kibibytes will instead be reported under 32KiB-128MiB. For entries without a size range (such as entries for seek or metadata functions, or resource entries), this appears as 'all'.

When read.sized or write.sized is set to all, the value of <MAX> is always double the value of <MIN>, except for the first bin which is reserved for reads or writes of exactly zero bytes, and has a <SIZE-RANGE> of 0-1B. The second bin has a <SIZE-RANGE> of 1B-2B, so includes single-byte operations only.

Both <MIN> and <MAX> are expressed using the IEC standard prefixes such as "KiB" for "kibibyte" (1024 bytes), "MiB" for "mebibyte" (1,048,576 bytes), and so on.

<MEASUREMENT> is one of the following, depending on the data reported by this log entry:

bandwidth	This entry reports the number of bytes read or written by the stated call type in the stated operation size range on the stated mount point in the timeframe.
count	This entry reports the number of calls of the stated call type with the given operation size range on the stated mount point in the timeframe.
max-latency	This entry reports the maximum latency sampled for the stated call type with the given operation size range on the stated mount point in the timeframe.
mean-latency	This entry reports the mean of latency samples for the stated call type with the given operation size range on the stated mount point in the timeframe.

total-latency	This entry reports an estimate of the total latency for the stated call type with the given operation size range on the stated mount point in the timeframe.
user-time	This entry reports the total CPU time used in user mode during the timeframe by all processes in the job.
system-time	This entry reports the total CPU time used in system mode during the timeframe for all processes in the job.
cpu-time	This entry reports the total CPU time used in either user or system mode during the timeframe for all processes in the job.
memory	This entry reports the total memory (resident set size) for all processes in the job seen during the timeframe.
memory-vsize	This entry reports the total memory (virtual memory size) for all processes in the job seen during the timeframe.
host-cpu-user-time	This entry reports the total CPU time used in user mode on the host during the timeframe.
host-cpu-system-time	This entry reports the total CPU time used in system mode on the host during the timeframe.
host-cpu-iowait-time	This entry reports the total CPU time spent waiting for I/O on the host during the timeframe.

<MEASURED-DATA> is the actual measurement, with suffix B for bytes, us for time (microseconds), or no suffix for a count measurement.

<TIMEFRAME> is the timeframe in seconds.

<HOSTNAME> is the name of the host on which the job was running. The host name includes the domain name.

<JOB-GROUP-ID> is the job group identifier for the job group.

<JOB-ID> is the job identifier for the job.

3.0.1 Log labels

The label of a log entry is intended to be a useful human-readable string summarising the measurement being made, and to work as a key for time-series data in your database. It is composed of the following parts:

- The prefix `job_total_`, if the log entry is the total value accumulated across a job (these log entries are produced at the end of the job if the `totals` configuration item is set).
- The prefix `delayed_`, if this log entry represents an I/O measurement which was somehow delayed and can only be reported substantially later than the end of its timeframe (this should be very unusual).
- The prefix `all_mounts_`, for an I/O log entry giving a total across all mount points, rather than for a single mount point.
- For resource entries, the name of the resource measurement (see <MEASUREMENT> above). Job-total memory resource entries, which report the maximum measurement over the course of the job (rather than a meaningless total of memory usage in each time frame), are prefixed with `max-`.
- For I/O entries, a prefix relating to the measurement type: `count-` for a call count, `bw-` for a bandwidth, `mean-lat-` for mean latency, `max-lat-` for maximum latency, `total-lat-` for total latency.
- For I/O entries, the <CALL-TYPE> (see above).
- For size-binned log entries, an underscore followed by the <SIZE-RANGE> (see above).

Here are some example log labels:

- `bw-read_512B-1KiB`: the number of bytes read in this timeframe, for a particular mount point, in read operations of 512 to 1023 bytes (inclusive).
- `all-mounts_count-open`: the number of calls to open across all mount points in this timeframe.
- `job-total_host-cpu-system-time`: The total system time used on the host over the course of the job.
- `job-total_all-mounts_count-read_all`: the total number of calls to read, across all mount points, over the whole course of the job.

4 Plug-in summary format

Mistral will collect aggregated statistics about the type of I/O that was performed. This has been split into three categories, good (green), medium (yellow) and bad (red). Please refer to the Mistral Manual for how I/O is categorised.

Note that this has been called “traffic light” output in previous versions of Mistral.

Log entries are output in the following format with one entry per job:

```
<TIME-STAMP>, <RUN-TIME>, <IO-TIME>, <%IO-TIME>, <IO-CALLS>,  
<RED-TIME>, <%RED-TIME>, <RED-CALLS>, <%RED-CALLS>,  
<YELLOW-TIME>, <%YELLOW-TIME>, <YELLOW-CALLS>, <%YELLOW-CALLS>,  
<GREEN-TIME>, <%GREEN-TIME>, <GREEN-CALLS>, <%GREEN-CALLS>,  
<JOB-GROUP-ID>, <JOB-ID>
```

Where the field definitions are as follows: <TIME-STAMP> Time when this log entry was created. We use ISO 8601 format with microsecond precision: YYYY-MM-DDThh:mm:ss.ffffff

<RUN-TIME> Wallclock runtime of this job (μ s).

<IO-TIME> Time spent doing I/O calls (μ s).

<%IO-TIME> % of runtime that was spent on I/O.

<IO-CALLS> Total number of I/O calls.

<RED-TIME> Time spent doing bad I/O (μ s).

<%RED-TIME> % of total I/O time that is bad I/O.

<RED-CALLS> Number of bad I/O calls.

<%RED-CALLS> % of total I/O calls that are bad I/O.

<YELLOW-TIME> Time spend doing medium I/O (μ s).

<%YELLOW-TIME> % of total I/O time that is medium I/O.

<YELLOW-CALLS> Number of medium I/O calls.

<%YELLOW-CALLS> % of total I/O calls that are medium I/O.

<GREEN-TIME> Time spent doing good I/O (μ s).

<%GREEN-TIME> % of total I/O time that is good I/O.

<GREEN-CALLS> Number of good I/O calls.

<%GREEN-CALLS> % of total I/O calls that are good I/O.

<JOB-GROUP-ID> Job group identifier.

<JOB-ID> Job identifier.

5 Example log entries

5.1 Example output log

The following are examples of log entries for the output file:

```
2022-07-20T14:27:06.000000,bw-read_4B-8B,/home,ext4,/dev/nvme0n1p5,
,read,4B-8B,bandwidth,36543B/1s,0/1s,dev371.example.com,0,0,,,,,0
2022-07-20T14:27:31.000000,count-access,/dev,devtmpfs,udev,
,access,all,count,22/1s,0/1s,dev371.example.com,0,0,,,,,0
2022-07-20T14:28:14.000000,count-write_64KiB-128KiB,/home,ext4,/dev/nvme0n1p5,
,write,64KiB-128KiB,count,185/10s,0/10s,dev371.example.com,0,0,,,,,0
2022-07-20T14:28:16.000000,max-lat-read_256B-512B,/home,ext4,/dev/nvme0n1p5,
,read,256B-512B,max-latency,14us/10s,0/10s,dev371.example.com,0,0,,,,,0
2022-07-20T14:29:10.000000,mean-lat-write_1KiB-2KiB,/home,ext4,/dev/nvme0n1p5,
,write,1KiB-2KiB,mean-latency,12us/10s,0/10s,dev371.example.com,0,0,,,,,0
2022-07-20T14:29:25.000000,total-lat-read_all,/home,ext4,/dev/nvme0n1p5,
,read,all,total-latency,9264us/10s,0/10s,dev371.example.com,0,0,,,,,0
2022-07-20T14:29:30.000000,vsize,/,,,
,none,all,memory-vsize,2643369984B/1s,0/1s,dev371.example.com,0,0,,,,,0
2022-07-20T14:25:35.000000,all-mounts_bw-read_8KiB-16KiB,*,,,
,read,8KiB-16KiB,bandwidth,19808465B/10s,0/10s,dev371.example.com,0,0,,,,,0
2022-07-20T14:25:43.000000,all-mounts_max-lat-open,*,,,
,open,all,max-latency,1580us/10s,0/10s,dev371.example.com,0,0,,,,,0
```

5.2 Example summary log

The follow are examples of log entries for the summary file:

```
2022-07-20T14:34:53.000000,707us,108us,15.28%,70,54us,50.00%,50,71.43%
,39us,36.11%,17,24.29%,15us,13.89%,3,4.29%,
2022-07-20T14:38:29.000000,2428us,479us,19.73%,101,126us,26.30%,62,61.39%
,96us,20.04%,30,29.70%,257us,53.65%,9,8.91%,2022.pbs,2022.pbs.1
```