

Mistral User Manual 2.9.3

Table of Contents

1 Introduction.....	2
2 Installation.....	2
3 Monitoring an application.....	3
3.1 Configuring contract and log locations.....	3
3.2 Contract specification.....	4
3.2.1 Header.....	4
3.2.2 Comments.....	4
3.2.3 Rules.....	5
3.2.3.1 Monitoring rules.....	8
3.2.3.2 Throttling rules.....	8
3.2.4 Latency Sampling.....	9
3.2.5 Adjusting contracts.....	10
3.3 Log Entries.....	10
3.3.1 Example Log Entries.....	11
4 Example contracts.....	12
4.1 Monitoring Contract.....	12
4.2 Throttling Contract.....	14
5 Plug-ins.....	16
5.1 Update Plug-in.....	16
5.2 Output Plug-in.....	16
5.3 Plug-in Configuration.....	16
5.3.1 PLUGIN directive.....	17
5.3.2 INTERVAL directive.....	17
5.3.3 PLUGIN_PATH directive.....	17
5.3.4 PLUGIN_OPTION directive.....	17
5.3.5 END Directive.....	18
5.3.6 Invalid Configuration.....	18
5.3.7 Example Configuration.....	18

1 Introduction

Mistral is a tool used to report on and resolve I/O performance issues when running complex Linux applications on high performance compute clusters.

Mistral is a small download that allows you to monitor application I/O patterns in real time, and log undesirable behaviour using rules defined in a configuration file called a contract.

2 Installation

Unzip the Mistral product archive that has been provided to you somewhere sensible. Please make sure that you use the appropriate version of Mistral (32 or 64bit) for the machine you want to run it on.

Mistral requires a license, please contact Ellexus if you do not already have a valid Mistral license.

The environment variable **MISTRAL_RLM_LICENSE** must be set to the location of your license, which can be one of:

- a) The pathname of a specific license file.
- b) The pathname of a directory which contains one or more license files.

Mistral will attempt to detect the installation directory correctly on start up however some job schedulers, e.g. Univa Grid Engine, use a spool directory that can break this detection. In this case the environment variable **MISTRAL_INSTALL_DIRECTORY** must be set to the directory used for installation.

There are two flavours of Mistral, designed to be used with either `/bin/bash` or `/bin/[t]csh` as interpreter.

If Mistral is intended to be used with a job scheduler all required environment variables must be available in all interactive and non-interactive shells. It is recommended that global environment variable settings be added to `/etc/bashrc` or `/etc/cshrc` and individual user settings to the user's `.bashrc` or `.cshrc` file.

3 Monitoring an application

Mistral is run using the `mistral` script available at the top level of the installation. To monitor an application you just type “`mistral`” followed by your command and arguments. For example:

```
> ./mistral make all
```

3.1 Configuring contract and log locations

Mistral determines what events to log and/or throttle by using contract files. Mistral uses two types of contracts, local and global. This enables administrators to define global settings for the entire system while also allowing for the creation of tuned settings for specific workloads. The following environment variables configure the locations Mistral uses for contract and log files.

It is not necessary to configure both global and local contracts but at least one valid contract / log pair must be defined. When testing it may be preferable to just use one contract, either local or global, for simplicity.

MISTRAL_CONTRACT_MONITOR_GLOBAL	The name and location of the global monitoring contract file.
MISTRAL_CONTRACT_MONITOR_LOCAL	The name and location of the local monitoring contract file.
MISTRAL_CONTRACT_THROTTLE_GLOBAL	The name and location of the global throttling contract file.
MISTRAL_CONTRACT_THROTTLE_LOCAL	The name and location of the local throttling contract file.
MISTRAL_LOG_MONITOR_GLOBAL	The name and location of the file in which Mistral will log violations of global contract rules.
MISTRAL_LOG_MONITOR_LOCAL	The name and location of the file in which Mistral will log violations of local contract rules.
MISTRAL_LOG_THROTTLE_GLOBAL	The name and location of the file in which Mistral will log throttling events triggered by a violation of a global contract rule.

MISTRAL_LOG_THROTTLE_LOCAL	The name and location of the file in which Mistral will log throttling events triggered by a violation of a local contract rule.
-----------------------------------	--

3.2 Contract specification

Monitoring contracts specify expected I/O limits for a process.

If any monitoring rule limit is exceeded a log message is output indicating which process broke the limit and by how much.

If a throttling rule limit is exceeded a log message is output indicating the process that contributed most to breaking the limit and all job processes are rate limited until the job falls within the defined rule.

3.2.1 Header

The first line of the file specifies the contract type and the time frame in the format

```
<VERSION>, <CONTRACT - TYPE>, <TIMEFRAME - PERIOD><TIMEFRAME - UNIT>
```

where:

<VERSION> is the contract format version number which must be “2” for this release of Mistral;

<CONTRACT - TYPE> is either `monitortimeframe` or `throttletimeframe`;

<TIMEFRAME - PERIOD> is the length of the time frame for all rules in this contract, specified as an integer, followed by;

<TIMEFRAME - UNIT> which must be “ms” for milliseconds or “s” for seconds.

For example:

```
2,monitortimeframe,15s
```

3.2.2 Comments

Blank lines and lines starting with “#” are ignored and can be used for adding comments to a contract file.

3.2.3 Rules

Each remaining line specifies a rule in the format

<LABEL>, <PATH>, <CALL - TYPE>, <SIZE - RANGE>, <MEASUREMENT>, <THRESHOLD><UNIT>

where:

<LABEL> is the name of this rule. It appears in log entries related to this rule. It is an arbitrary string of the letters a-z (in lower or upper case), the digits 0-9, the underline character (“_”) or a hyphen (“-”).

<PATH> is an absolute file system path. The rule applies to function calls on paths starting with this value. Mistral de-references all relative paths and symbolic links therefore this path must be fully resolved.

Note that this is a purely string-based comparison. For example, the path “/usr/lib32/libm.a” matches a rule where **<PATH>** is “/usr/lib”. This is useful when many paths share a prefix, so for example it is possible to set **<PATH>** to “/tmp/output/test-” in order to match “/tmp/output/test-1” and so on. To restrict the path to a particular directory, it must be specified with a trailing “/”, for example “/usr/lib/”.

<CALL - TYPE> is the set of call types to which the rule applies. It must specify one or more of these call types:

access	Calls that access file system meta data (stat, readlink, etc.).
create	Calls that create new files (open, creat, mkdir, etc.).
delete	Calls that delete files (remove, rmdir, unlink, etc.).
fschange	Calls that update file system meta data (chmod, rename, etc.).
open	Calls that open existing files (open, fopen, opendir, etc.).
read	Calls that read data from the file system (read, fgets, mmap, readdir, recv, scanf, etc.).
seek	Calls that update the current position within a file (fseek, lseek, rewind, etc.).
write	Calls that write data to the file system (write, error, printf, putc, send, warn, etc.).

When a rule applies to multiple call types, join them with “+” signs. For example, “read+write” matches calls that either read or write data.

<SIZE - RANGE> specifies the range of sizes that match this rule. A size range may only be specified for rules with any combination of the call types “read”, “write”, and “seek” (other types of call have no associated size). A size range is specified in the format:

<SIZE - MIN><SIZE - MIN - UNIT>-<SIZE - LIMIT><SIZE - LIMIT - UNIT>

meaning that a matching size must be at least <SIZE - MIN> but lower than <SIZE - LIMIT>. The <SIZE - MIN - UNIT> and <SIZE - LIMIT - UNIT> are the corresponding units, and must be one of the following:

- “B” Bytes
- “kB” Kilobytes (1,000 bytes)
- “MB” Megabytes (1,000,000 bytes)
- “GB” Gigabytes (1,000,000,000 bytes)

For example, a size range of “1kB-4kB” matches reads (or writes) with 1000 <= size < 4000. (Note the asymmetric bounds: these make it easier to specify non-overlapping ranges.)

<SIZE - MIN><SIZE - MIN - UNIT> may be omitted, in which case the value 0 is used. <SIZE - LIMIT><SIZE - LIMIT - UNIT> may be omitted, in which case there is no upper limit.

If a rule is to apply to all of the specified operations regardless of size, or size is not applicable to one or more of the call types specified in the rule this field must be set to “**all**”.

<MEASUREMENT> is the type of data being measured. The list of valid measurement types differs between monitoring throttling rules. For monitoring rules it must be one of:

bandwidth	Amount of data processed by calls of the specified type in the time frame. This applies only to “read” and “write” calls.
count	The number of calls of the specified type in the time frame.
seek-distance	Total distance moved within files by calls of the specified type in the time frame. This applied only to “seek” calls.
max-latency	The maximum duration of any call of the specified type in the time frame. See section 3.2.4 Latency Sampling.

mean-latency	The mean duration of any call of the specified type in the time frame, provided the number of calls is higher than the value of MISTRAL_MONITOR_LATENCY_MIN_IO . See section 3.2.4 Latency Sampling.
total-latency	The total duration of time spent in calls of the specified type in the time frame, provided the number of calls is higher than the value of MISTRAL_MONITOR_LATENCY_MIN_IO . See section 3.2.4 Latency Sampling.

For throttling rules the only valid measurements are “bandwidth” and “count” as described above.

<THRESHOLD> is the limit for this rule. If the measured data exceeds <THRESHOLD> in <TIMEFRAME>, then the violation is logged.

<UNIT> is the unit for <THRESHOLD>. When <MEASUREMENT> is “bandwidth” or “seek-distance”, this must be one of:

- “B” Bytes
- “kB” Kilobytes (1,000 bytes)
- “MB” Megabytes (1,000,000 bytes)
- “GB” Gigabytes (1,000,000,000 bytes)

When <MEASUREMENT> ends with “-latency”, this must be one of:

- “us” Microseconds
- “ms” Milliseconds
- “s” Seconds

When <MEASUREMENT> is “count”, this must be one of:

- <blank> Exact number of calls
- “k” Thousands of calls
- “M” Millions of calls

For example:

```
red, /mnt/net/abc, write, all, bandwidth, 100MB
```

3.2.3.1 Monitoring rules

Monitoring rules within the same contract are grouped by <PATH>, <CALL-TYPE> and <MEASUREMENT>. If multiple rules in a group have been violated simultaneously, only the rule with the highest <THRESHOLD> is logged.

For example, consider the contract:

```
2,monitortimeframe,1s
#LABEL,PATH,CALL-TYPE,SIZE-RANGE,MEASUREMENT,THRESHOLD
Red,/mnt/net/abc,write,all,bandwidth,1MB
Yellow,/mnt/net,write,all,bandwidth,10MB
Green,/mnt/net/abc,write,all,bandwidth,10kB
#Black,/mnt/net/abc,write,all,bandwidth,1kB
```

In this example, if the application writes more than 10kB/s in subdirectories of “/mnt/net/abc” the Green rule is violated and logged. If it writes more than 1MB/s in “/mnt/net/abc”, the Green and Red rules are violated but only the Red rule is logged. If it writes more than 10MB/s in “/mnt/net/abc”, the Red, Yellow and Green rules all match, but only the Red and Yellow rules are logged. The Black rule is never logged, because it has been commented out with “#”.

3.2.3.2 Throttling rules

If a path matches multiple rules in the throttle contract file, then all of the limitations apply. For example, if the contract file contains:

```
thr_root_r,/,read,all,bandwidth,10000000B
thr_usr_r,/usr,read,all,bandwidth,500000B
thr_usrlib_r,/usr/lib,read,all,bandwidth,90000B
```

Then a read from “/usr/lib/libc.so” is subject to all three bandwidth limitations, whereas a read from “/etc/passwd” is subject only to the first.

3.2.4 Latency Sampling

Latency measurements incur a larger processing overhead than simple count or bandwidth operations. Such measurements are also subject to greater variability in value. To limit the impact of these problems Mistral implements measurement sampling on any latency rules defined in a monitoring contract.

Latency sampling is controlled via three environment variables. All three variables must be set to a positive integer if defined.

MISTRAL_MONITOR_LATENCY_SAMPLE	The sampling rate. If set to n Mistral will only measure the latency of every n^{th} operation. Setting this to 1 will cause the latency of all I/O operations to be measured. Defaults to 10.
MISTRAL_MONITOR_LATENCY_MIN_IO	The minimum number of I/O operations that must be seen in a single time frame for the sample to be considered statistically significant. Defaults to 100.
MISTRAL_MONITOR_LATENCY_MAX_IO	The maximum number of I/O operations to measure in a single time frame before it is considered not to add statistically significant information. Defaults to 1000.

Latency measurements are not made if no latency rules are defined. If latency rules are defined but any of the environment variables described above are not explicitly defined Mistral will output a message stating the default values that will be used.

The minimum and maximum counts defined above are applied individually to each <CALL-TYPE> class. For example, using default configuration where a minimum of 100 operations must be seen before rules are applied, if 150 read operations are sampled in a single time frame any defined latency rules against read events will be applied. If, during that same time frame, only 75 write operations are seen any latency rules defined against write events will not be applied.

It is important to note that any the total-latency rules defined are only compared to the sampled I/O operations and are subject to the maximum sample size count limit defined.

3.2.5 Adjusting contracts

It is possible to update contracts for running jobs and can be particularly useful to increase thresholds to prevent excessive logging. How this is done differs between global and local contracts.

Global contracts are assumed to be configured with high “system threatening” rules that should not be frequently changed. These contracts are intended to be maintained by system administrators and will be polled approximately once a minute for changes on disk.

Local contracts can be update dynamically during a job execution run by the use of an update plug-in. Using an update plug-in is the only way to modify the local contracts in use by a running job. If an update plug-in configuration is not defined Mistral will use the same local configuration contracts throughout the life of the job.

Please see section 5 for details on the configuration and use of plug-ins.

3.3 Log Entries

Log entries are output in the following format:

```
<TIME-STAMP>, <LABEL>, <PATH>, <CALL-TYPE>, <SIZE-RANGE>, <MEASUREMENT>,
<MEASURED-DATA>/<TIMEFRAME-PERIOD><TIMEFRAME-UNIT>,
<THRESHOLD>/<TIMEFRAME-PERIOD><TIMEFRAME-UNIT>, <PID>,
<COMMAND-LINE>, <FILE-NAME>, <JOB-GROUP-ID>, <JOB-ID>
```

Where the field definitions are as follows:

<TIME-STAMP>	is the time the violation was logged, in ISO 8601 format.
<LABEL>	is copied from the violated rule.
<PATH>	is the path that caused <MEASURED-DATA> to exceed <THRESHOLD>.
<CALL-TYPE>	is copied from the violated rule.
<SIZE-RANGE>	is copied from the violated rule.
<MEASUREMENT>	is copied from the violated rule.
<MEASURED-DATA>	is the data rate of the job that exceeded the limit.
<THRESHOLD>	is copied from the violated rule.

<TIMEFRAME-PERIOD>	is copied from the violated rule.
<TIMEFRAME-UNIT>	is copied from the violated rule.
<PID>	is the id of the process in the job that performed the most I/O that contributed to violating the rule.
<COMMAND-LINE>	is the name of the process in the job that performed the most I/O that contributed to violating the rule. It includes the parameters for the execution.
<FILE-NAME>	is the path to the file being accessed by the operation that violated the rule.
<JOB-GROUP-ID>	is the job group identifier for the job group that violated the rule.
<JOB-ID>	is the job identifier for the job that violated the rule.

3.3.1 Example Log Entries

The following is an example of a rule violation log entry:

```
2015-01-30T14:30Z, red, /mnt/net/abc, write, all, bandwidth, 102MB/15s, 1MB/15s, 1234, /mnt/tool/bin/abc -d -e, /mnt/net/abc/file, 5, 5
```

If the <PATH> in one rule is a subdirectory of the <PATH> in another rule, then a single process accessing the subdirectory may violate both rules e.g.

```
2015-01-30T14:30Z, red, /mnt/net/abc, write, all, bandwidth, 102MB/15s, 1MB/15s, 1234, /mnt/tool/bin/abc -d -e, /mnt/tool/abc/file2, 5, 5
2015-01-30T14:30Z, yellow, /mnt/net, write, all, bandwidth, 102MB/15s, 10MB/15s, 1234, /mnt/tool/bin/abc -d -e, /mnt/tool/abc/file1, 5, 5
```

Although violated throttling rules will cause Mistral to slow the I/O operation of all processes within a job, any I/O operation that is already in progress when throttling is applied will complete without any modification by Mistral.

As a result the I/O rate measured may still exceed the defined limit even under throttling. The actual I/O rate that was achieved when applying the throttle is output in the <MEASURED-DATA> field.

4 Example contracts

4.1 Monitoring Contract

Consider the following contract:

```
2,monitortimeframe,1s
#LABEL,PATH,CALL-TYPE,SIZE-RANGE,MEASUREMENT,THRESHOLD
High_reads,/usr,read,all,bandwidth,1MB
High_reads_bin,/usr/bin,read,all,bandwidth,5MB
Higher_reads_bin,/usr/bin,read,all,bandwidth,50MB
High_create_lat,/tmp,create,all,mean-latency,10ms
High_num_w,/home/,write,all,count,750
```

Examining each line individually:

2,monitortimeframe,1s

This line identifies the contract as containing monitoring rules that are applied over a time frame of 1 second.

High_reads,/usr,read,all,bandwidth,1MB

This line defines a rule name “High_reads” and tells Mistral to generate an alert when the total amount of data read from /usr exceeds 1MB within the 1s time frame.

If a monitored process were to read a 2MB file in /usr/share/doc/ in less than a second, for example, this rule would be violated and a log message of the following form would be output:

```
2015-07-30T14:30Z,High_reads,/usr,read,all,bandwidth,2MB/1s,
1MB/1s,15392,/mnt/tool/bin/python script.py,
/usr/share/doc/glibc-common-2.17/README.timezone,3,6
```

High_reads_bin,/usr/bin,read,all,bandwidth,5MB

Higher_reads_bin,/usr/bin,read,all,bandwidth,50MB

These two lines define two additional rules named “High_reads_bin” and “Higher_reads_bin” respectively.

All reads in “/usr/bin” will be tested against all three rules currently defined as a read under “/usr/bin” is also a read under “/usr”.

If a process read 6MB of data in less than 1 second both the “High_reads” rule and the “High_reads_bin” rule would be violated. As the rules are defined on different paths a log message for both rule 1 and rule 2 will be output:

```
2015-07-30T14:30Z,High_reads,/usr,read,all,bandwidth,6MB/1s,
1MB/1s,15392,/bin/bash script.sh,/usr/bin/data,3,6
2015-07-30T14:30Z,High_reads_bin,/usr/bin,read,all,bandwidth,
6MB/1s,5MB/1s,15392,/bin/bash script.sh,/usr/bin/data,3,6
```

If a process read 60MB of data in less than 1 second all three currently defined rules would be violated, but only the first and third rule would be logged. This is because Mistral only logs the largest threshold violated when multiple rules are defined on the same “path”, “call-type” and “measurement” as is the case with the “High_reads” and “High_reads_bin” rules:

```
2015-07-30T14:30Z,High_reads,/usr,read,all,bandwidth,60MB/1s,
1MB/1s,15392,/bin/bash script.sh,/usr/bin/data,3,6
2015-07-30T14:30Z,Higher_reads_bin,/usr/bin,read,all,bandwidth,
60MB/1s,50MB/1s,15392,/bin/bash script.sh,/usr/bin/data,3,6
```

High_create_lat,/tmp,create,all,mean-latency,10ms

The rule labelled “High_create_lat” is only concerned with function calls that create file system objects (“create”) under “/tmp”. In this case the latency of each call made during the time frame is accumulated and averaged over the total number of these calls, provided the number of calls within the time frame is higher than the value of **MISTRAL_MONITOR_LATENCY_MIN_IO**.

If at the end of the time frame this “mean-latency” is higher than “10ms” then a log message will be output, for example:

```
2015-07-30T15:10Z,High_create_lat,/tmp,create,all,
mean-latency,22ms,10ms,15537,/bin/bash script.sh,/tmp/data,3,6
```

High_num_w,/home/,write,all,count,750

The rule labelled “High_num_w” is violated if the number of write calls exceeds 750.

```
2015-07-30T15:10Z,High_num_w,/home,write,all,count,863,
750,15537,/bin/bash script.sh,/home/data,3,6
```

4.2 Throttling Contract

Consider the following contract:

```
2, throttletimeframe, 1s
#LABEL, PATH, CALL-TYPE, SIZE-RANGE, MEASUREMENT, ALLOWED
High_reads, /usr, read, all, bandwidth, 5MB
High_reads_bin, /usr/bin, read, all, bandwidth, 1MB
High_num_r, /home/, read, all, count, 750
```

Examining each line individually:

2, throttletimeframe, 1s

This line identifies the contract as containing throttling rules that are applied over a time frame of 1 second.

High_reads, /usr, read, all, bandwidth, 5MB

If a monitored job were to try and read a 6MB file in /usr/share/doc/ in less than a second, for example, this rule would be violated. When Mistral identifies an I/O operation that would violate a throttling rule it will introduce a sleep long enough to bring the observed I/O back down to the configured limit and a log message of the following form will be output:

```
2015-07-30T14:30Z, High_reads, /usr, read, all, bandwidth, 1MB/1s,
1MB/1s, 15392, /mnt/tool/bin/python script.py, 3, 6
```

High_reads_bin, /usr/bin, read, all, bandwidth, 1MB

The second rule in this contract is very similar to the first. Again it is monitoring read bandwidth but this time is only interested in reads that occur under “/usr/bin” and will allow up to “1MB” of data to be read before the rule is violated.

In this case all reads in “/usr/bin” will be tested against both the “High_reads” and “High_reads_bin” rules as a read under “/usr/bin” is also a read under “/usr”.

If a process were to attempt to read 3MB of data in “/usr/bin” in less than 1 second it would violate the “High_reads_bin” rule but would not violate the “High_reads” rule. Therefore the process would be limited to “1MB/1s” and up to three log messages similar to the one described above will be output depending on the number of function calls that are used to read the 3MB of data.

If the process instead attempted to read 6MB of data in less than 1 second both the currently defined rules would be violated. Even though the rules are defined on

different paths in this case the most restrictive rule applies and again the process will be throttled to “1MB/1s” and up to 6 log messages generated by violations of the “High_reads_bin” rule will be logged.

High_num_r,/home/,read,all,count,750

The third rule does not care about how large each operation is, it is simply interested in the total number of times a call is made to a “read” operation. If a total of more than “750” read operations are performed within the time frame of 1 second under “/home/” then on the 751st read Mistral would introduce a sleep long enough to bring the data rate under “750/1s” and a log message of the following form would be logged:

```
2015-07-30T16:45Z,High_num_r,/home/,read,all,count,750/1s,750/1s,16601,/usr/lib64/firefox/firefox,/home/data,1,1
```

5 Plug-ins

Currently two different plug-ins are supported.

5.1 Update Plug-in

The update plug-in is used to modify local Mistral configuration contracts dynamically during a job execution run according to conditions on the node and / or cluster. Using an update plug-in is the only way to modify the local contracts in use by a running job.

Global contracts are assumed to be configured with high “system threatening” rules that should not be frequently changed. These contracts are intended to be maintained by system administrators and will be polled periodically for changes on disk as described above. Global contracts cannot be modified by the update plug-in in any way.

If an update plug-in configuration is not defined Mistral will use the same local configuration contracts throughout the life of the job.

5.2 Output Plug-in

The output plug-in is used to record alerts generated by the Mistral application. All event alerts raised against any of the four valid contract types are sent to the output plug-in. The four contract types are:

- Global Monitoring
- Global Throttling
- Local Monitoring
- Local Throttling

If an output plug-in configuration is not defined Mistral will default to recording alerts to disk as described above.

5.3 Plug-in Configuration

On start up Mistral will check the environment variable **MISTRAL_PLUGIN_CONFIG**. If this environment variable is defined it must point to a file that the user running the application can read. If the environment variable is not defined Mistral will assume that no plug-ins are required and will use the default behaviours as described above.

The expected format of the configuration file consists of one block of configuration lines for each configured plug-in. Each line is a comma separated pair of a single configuration option directive and its value. Whitespace is treated as significant in this file. The full specification for a plug-in configuration block is as follows:


```
PLUGIN,<OUTPUT|UPDATE>
INTERVAL,<Calling interval in seconds>
PLUGIN_PATH,<Fully specified path to plug-in>
[PLUGIN_OPTION,<Option to pass to plug-in>]
...
END
```

5.3.1 PLUGIN directive

The `PLUGIN` directive can take one of only two values, “UPDATE” or “OUTPUT” which indicates the type of plug-in being configured. If multiple configuration blocks are defined for the same plug-in the values specified in the later block will take precedence.

5.3.2 INTERVAL directive

The `INTERVAL` directive 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.

5.3.3 PLUGIN_PATH directive

The `PLUGIN_PATH` directive value must be the fully qualified path to the plug-in to be run e.g. `/home/ellexus/bin/output_plugin.sh`. This plug-in must be executable by the user that starts the Mistral application. The plug-in must also be available in the same location on all possible execution host nodes where Mistral is expected to run.

The `PLUGIN_PATH` value will be passed to `/bin/sh` for environment variable expansion at the start of each execution host job.

5.3.4 PLUGIN_OPTION directive

The `PLUGIN_OPTION` directive is optional and can occur multiple times. Each `PLUGIN_OPTION` directive is treated as a separate command line argument to the plug-in. Whitespace is respected in these values.

As whitespace is respected command line options that take parameters must be specified as separate `PLUGIN_OPTION` values. For example if the plug-in uses the option “`--output /dir/name/`” to specify where to store its output then this must be specified in the plug-in configuration file as:

```
PLUGIN_OPTION, --output
PLUGIN_OPTION, /dir/name/
```

Options will be passed to the plug-in in the order in which they are defined.

Each `PLUGIN_OPTION` value will be passed to `/bin/sh` for environment variable expansion at the start of each execution host job.

5.3.5 END Directive

The `END` directive indicates the end of a configuration block and does not take any values.

5.3.6 Invalid Configuration

Blank lines and lines starting with “#” are silently ignored. All other lines that do not begin with one of the configuration directives defined above cause a warning to be raised.

5.3.7 Example Configuration

Consider the following configuration file; line numbers have been added for clarity:

```
1  # Automatically generated by gen.sh. DO NOT EDIT
2
3  PLUGIN, OUTPUT
4  INTERVAL, 300
5  PLUGIN_PATH, /home/ellexus/bin/output_plugin.sh
6  PLUGIN_OPTION, --output
7  PLUGIN_OPTION, /home/ellexus/log files
8  END
9
10 PLUGIN, UPDATE
11 INTERVAL, 60
12 PLUGIN_PATH, $HOME/bin/update_plugin
13 END
```

The configuration file above sets up both update and output plug-ins.

Lines 1-2 are ignored as comments. The first configuration block (lines 3-8) defines an output plug-in (line 3) that will be called every 300 seconds (line 4) using the command line

```
/home/ellexus/bin/output_plugin.sh --output "/home/ellexus/log files"
```

(lines 5-7). The configuration block is terminated on line 8.

The blank line is ignored (line 9).

The second configuration block (lines 10-13) defines an update plug-in (line 10) that will be called every 60 seconds (line 11) using the command line `/home/ellexus/bin/update_plugin`, (line 12), assuming `$HOME` is set to `/home/ellexus`. The configuration block is terminated on line 13.