# Altair - InfluxDB Plug-in Configuration Guide

# Contents

**1 Installing the InfluxDB Plug-in** **3**

**2 Configuring Mistral to use the InfluxDB 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
2.1.4 Switches . . . 4
2.2 Plug-in Configuration File Options . . . 5
2.3 Plug-in Configuration File Switches . . . 5

**3 Mistral's InfluxDB data model** **6**
3.1 Measurements . . . 6
3.2 Series Timestamps . . . 6
3.3 Mountpoint . . . 6
3.3.1 Tags . . . 6
3.3.2 Fields . . . 7
3.4 Network . . . 7
3.4.1 Tags . . . 7
3.4.2 Fields . . . 7
3.5 Resources . . . 8
3.5.1 Tags . . . 8
3.5.2 Fields . . . 8
3.6 Summary . . . 8
3.6.1 Tags . . . 8
3.6.2 Field . . . 9

# 1 Installing the InfluxDB 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 InfluxDB 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 InfluxDB plug-in can be found in, for example:

`<installation directory>/output/mistral_influxdb_v7.2.8/x86_64/`

for the 64 bit Intel compatible version, and in:

`<installation directory>/output/mistral_influxdb_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 InfluxDB 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_influxdb/x86_64/mistral_influxdb
```

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

### 2.1.1 Path

The `path` key must be set to the path of the influxdb plguin-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 in the value are not supported.

```
path: plugins/mistral_influxdb/x86_64/mistral_influxdb
```

### 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 InfluxDB 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:
    database: dbname
    error: filename
```

will pass to the plug-in executable the command line arguments `--database=dbname` and `--error=filename`.

### 2.1.4 Switches

The `switches` mapping is optional and lists all switches to be passed to the plug-in as command line arguments to the executable. A full list of valid switches for this plug-in can be found in section 2.3 Plug-in Configuration File Switches. The order of switches is not preserved. Switches not present are presumed to be off. These switches are passed to the plug-in executable as `--key`. For example,

```
switches: ["job-as-tag"]
```

will pass to the plug-in executable the command line argument `--job-as-tag`.

## 2.2 Plug-in Configuration File Options

The following command line options are supported by the InfluxDB plug-in.

`bucket: bucket_name`

Set the name of the bucket the plug-in should connect to. Defaults to “mistral”

`database: dbname`

See *bucket* option. Provided for backwards compatibility with InfluxDB 1.x.

`error: filename`

Set the location of the file which should be used to log any errors encountered by the plug-in. Defaults to sending messages to *stderr* for handling by Mistral.

`host: hostname`

Set the location of the InfluxDB host the plug-in should connect to. Defaults to “localhost”

`influx-version: num`

Set the major version number of the InfluxDB server being connected to. Defaults to 2.

`mode: octal-mode`

Permissions used to create the error log file specified by the -e option.

`org: org_name`

Sets the name of the organization the bucket to be written to belongs to. InfluxDB 2.x or greater.

`orgid: org_id`

Sets the id of the organization the bucket to be written to belongs to. InfluxDB 2.x or greater.

`password: passwd`

Sets the secret to be used for authentication. If *user* is also specified, then basic http authentication is used. If not, then the secret is assumed to be an API token. In this case, if secret begins with *file:*, then the plug-in will attempt to read the token from the first line of the supplied file.

`port: number`

Set the port the plug-in should use for the connection. Defaults to “8086”.

`user: username`

Set the username the plug-in should use to authenticate its connection to InfluxDB. This will cause the plug-in to attempt basic http authentication.

## 2.3 Plug-in Configuration File Switches

`job-as-tag`

Use job-id and job-group as a tag and not a field.

`ssl`

Use HTTPS to connect to InfluxDB rather than plain HTTP.

`skip-ssl-validation`

Causes the plug-in to not validate any self-signed CA certificates.

# 3 Mistral's InfluxDB data model

This section describes how the Mistral InfluxDB plug-in stores data within InfluxDB. It should be noted that there is some duplication of terminology between Mistral and InfluxDB.

## 3.1 Measurements

Both products use the term "measurement". In Mistral's case this refers to the type of statistic being recorded e.g. bytes written versus duration of calls; whereas for InfluxDB it is the primary identifier of related data (analogous to a table in a standard RDBMS). It will be specified which measurement is being referred to.

This is dealt with by mapping a Mistral record's `type` to an InfluxDB measurement. As such, the following list of all possible InfluxDB measurement values aligns with the possible values for `type` output by Mistral. As above where this list conflicts 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.

| InfluxDB Measurement | Description |
|---|---|
| mountpoint | Mistral data associated to the I/O done on a particular mountpoint. |
| network | Mistral data associated to the I/O sent or received on a particular IP address. |
| resources | Mistral data associated cpu and memory usage. |
| summary | Mistral data that is output at job end and summarises the file I/O done. |

## 3.2 Series Timestamps

Mistral produces log entries with time stamps to a precision of one second. The Mistral InfluxDB plug-in will attempt to normalise any timezone information to derive the correct value for the "microseconds since epoch" required by InfluxDB.

The next sections detail the different tags and fields for each measurement.

## 3.3 Mountpoint

### 3.3.1 Tags

| Tag Key | Description |
|---|---|
| path | Copied from `mountpoint.path` field from the Mistral log unchanged. |
| fstype | Copied from `mountpoint.fstype` field from the Mistral log unchanged. |
| fsname | Copied from `mountpoint.fsname` field from the Mistral log unchanged. |
| fshost | Copied from `mountpoint.fshost` field from the Mistral log unchanged. |
| cumulative | Copied from `cumulative` field from the Mistral log if present, or `false`. |
| hostname | Copied from `hostname` field from the Mistral log unchanged. |
| jobtotal | Copied form `jobtotal` field from the Mistral log unchanged. |
| jobid | Copied from `jobid` field from the Mistral log unchanged only if the `--job-as-tag` argument is specified. |

| Tag Key | Description |
|---|---|
| jobgropid | Copied from `jobgroup` field from the Mistral log unchanged only if the `--job-as-tag` argument is specified. |
| environment#ENV | There are potentially multiple tags of this kind, where each `ENV` is a key in the `environment` object from the Mistral log, equal to the corresponding value. |

#### 3.3.2 Fields

| Field Key | Description |
|---|---|
| jobid | Copied from `jobid` field from the Mistral log unchanged only if the `--job-as-tag` argument is **not** specified. |
| jobgropid | Copied from `jobgroup` field from the Mistral log unchanged only if the `--job-as-tag` argument is **not** specified. |
| io#calltype#measurement#value | There are potentially multiple fields of this kind, corresponding to the value of `io.calltype.measurement.value` from the Mistral log. For example, `io#read_all#bytes#mean/s` is the value for `io.read_all.bytes.mean/s` from the Mistral log. |

### 3.4 Network

#### 3.4.1 Tags

| Tag Key | Description |
|---|---|
| address | Copied from the `address` field form the Mistral log unchanged. |
| cumulative | Copied from `cumulative` field from the Mistral log if present, or `false`. |
| hostname | Copied from `hostname` field from the Mistral log unchanged. |
| jobtotal | Copied form `jobtotal` field from the Mistral log unchanged. |
| jobid | Copied from `jobid` field from the Mistral log unchanged only if the `--job-as-tag` argument is specified. |
| jobgropid | Copied from `jobgroup` field from the Mistral log unchanged only if the `--job-as-tag` argument is specified. |
| environment#ENV | There are potentially multiple tags of this kind, where each `ENV` is a key in the `environment` object from the Mistral log, equal to the corresponding value. |

#### 3.4.2 Fields

| Field Key | Description |
|---|---|
| jobid | Copied from `jobid` field from the Mistral log unchanged only if the `--job-as-tag` argument is **not** specified. |

| Field Key | Description |
|---|---|
| jobgropid | Copied from `jobgroup` field from the Mistral log unchanged only if the `--job-as-tag` argument is **not** specified. |
| io#calltype#measurement#value | There are potentially multiple fields of this kind, corresponding to the value of `io.calltype.measurement.value` from the Mistral log. For example, `io#read_all#bytes#mean/s` is the value for `io.read_all.bytes.mean/s` from the Mistral log. |

## 3.5 Resources

### 3.5.1 Tags

| Tag Key | Description |
|---|---|
| hostname | Copied from `hostname` field from the Mistral log unchanged. |
| jobtotal | Copied form `jobtotal` field from the Mistral log unchanged. |
| jobid | Copied from `jobid` field from the Mistral log unchanged only if the `--job-as-tag` argument is specified. |
| jobgropid | Copied from `jobgroup` field from the Mistral log unchanged only if the `--job-as-tag` argument is specified. |
| environment#ENV | There are potentially multiple tags of this kind, where each `ENV` is a key in the `environment` object from the Mistral log, equal to the corresponding value. |

### 3.5.2 Fields

| Field Key | Description |
|---|---|
| jobid | Copied from `jobid` field from the Mistral log unchanged only if the `--job-as-tag` argument is **not** specified. |
| jobgropid | Copied from `jobgroup` field from the Mistral log unchanged only if the `--job-as-tag` argument is **not** specified. |
| resources#collection#measurement#value | There are potentially multiple fields of this kind, corresponding to the value of `resources.collection.measurement.value` from the Mistral log. For example, `resoruces#host#rssmem#max` is the value for `resources.host.rssmem.max` from the Mistral log. |

## 3.6 Summary

### 3.6.1 Tags

| Tag Key | Description |
|---|---|
| hostname | Copied from `hostname` field from the Mistral log unchanged. |
| jobtotal | Copied form `jobtotal` field from the Mistral log unchanged. |
| jobid | Copied from `jobid` field from the Mistral log unchanged only if the `--job-as-tag` argument is specified. |
| jobgropid | Copied from `jobgroup` field from the Mistral log unchanged only if the `--job-as-tag` argument is specified. |
| environment#ENV | There are potentially multiple tags of this kind, where each `ENV` is a key in the `environment` object from the Mistral log, equal to the corresponding value. |

### 3.6.2 Field

| Field Key | Description |
|---|---|
| jobid | Copied from `jobid` field from the Mistral log unchanged only if the `--job-as-tag` argument is **not** specified. |
| jobgropid | Copied from `jobgroup` field from the Mistral log unchanged only if the `--job-as-tag` argument is **not** specified. |
| iosummary#total#measurement | There are potentially multiple fields of this kind, corresponding to the value of `iosummary.total.measurement` from the Mistral log. For example, `iosummary#total#jobrealtime` is the value for `iosummary.total.jobrealtime` from the Mistral log. |
| iosummary#category#measurement#value | There are potentially multiple fields of this kind, corresponding to the value of `iosummary.category.measurement.value` from the Mistral log. For example, `iosummary#yellow#calls#iocount` is the value for `iosummary.yellow.calls.iocount` from the Mistral log. |