

# Analytics module

- [Install analytics module for SAYMON ISO-image](#)
- [Analytics module configuration](#)
- [Analytics module monitoring](#)
- [Analytics module performance](#)

## Install analytics module for SAYMON ISO-image

By default, the analytics module used to determine splashes and predictions is not included in the SAYMON ISO image. You can install it in 3 steps:

1. download [analytics-create-and-run.sh](#) script and place it in the file system on the virtual machine where the SAYMON ISO image is installed;
2. run the script:

```
sudo ANALYTICS_BRANCH=x.y.z ./analytics-create-and-run.sh
```

3. check that the saymon-analytics container is up and running:

```
docker ps -a
```

Next you can configure the connection of the analytics module to SAYMON in the [server configuration file](#).

## Analytics module configuration

To connect the analytics module to the SAYMON server, you need to enter additional parameters into the SAYMON configuration file:

```
/etc/saymon/saymon-server.conf
```

Open the configuration file and add the following values to the "server" section:

```
{
  ...
  "server": {
    ...
    "analytics_enabled": true, # enables analytics module
    "analytics_processes": 2 # the number of logical cores allocated for
analytics processing (by default, all available system cores are used)
  }
}
```

Save configuration file and restart the SAYMON server:

```
service saymon-server restart
```

## Analytics module monitoring

You can use the "saymon-analytics" container logs to monitor the operation of the analytics module. The steps are:

1. connect to the container:

```
docker exec -ti saymon-analytics bash
```

2. open log file:

```
tail -f /opt/analytics/analytics.log
```

Log file contains information about the settings with which the module is started (socket and the number of processes involved in the pool), the metrics processed and the results of their processing, errors that occur during the module's operation.

## Analytics module performance

Analytics for each of the objects is carried out in a separate stream of calculations using one logical core of the system. All metric values for the same object are processed sequentially, which allows sequential update of the forecast model in accordance with the order of data receipt. Nowadays, the performance of modern processors allows you to process each incoming value in real time (with a maximum frequency = 1 metric value per second). However, on superheated or overloaded machines the performance of the cores can decrease, that leads to the accumulation of tasks for the analytics module and increasing the reaction time of the system to new incoming data. To increase system performance in tasks with a high frequency of incoming metrics, it is recommended to allocate one logical core for each object for analytics. If there are additional cores, the system can also use them for preprocessing analytic tasks, since locking on a predictive model takes only a fraction of the total processing time of the analytic task.