

EnergyScopium Monitor user documentation

DENERGIUM

DENERGIUM is an Inria spinoff. DENERGIUM is a software publisher and a positive environmental company. DENERGIUM's mission is to improve the efficiency of the use of IT infrastructures, particularly data centers.

DENERGIUM publishes the EnergyScopium software suite, which enables cluster managers to optimize energy resources and enables users and developers to integrate energy criteria into their usage.

The EnergyScopium software suite

The EnergyScopium software suite is available in several versions:

- Version Monitor: installed by the cluster manager. The cluster manager has access to information on the cluster's energy behavior. Users have access to energy information about their work. All resources and data are stored in the manager's space (On-Premise version).
- Optimize On-Premise version: installed by the cluster manager. The user can activate the capture of energy information about his work. All resources and data are stored in the manager's space (known as the On-Premise version).
- Optimize version in SaaS (Cloud) mode: Minimal installation by user or manager. The user initiates an account on the EnergyScopium SaaS platform (www.energyscopium.com). The user can activate the capture of energy information on his work. The data acquired is transferred to a secure server, which returns the job's energy data. On EnergyScopium SaaS, you can have a global and detailed view of your work (known as an energy workshop).

Your manager has installed the Monitor version of the EnergyScopium software suite.

The rest of this document enables the user to use the services provided by the software suite.

User Documentation

The services provided by the EnergyScopium Monitor software suite are :

- An end-of-job banner (this service can be activated and deactivated by the manager and user)
- An energy report in html format
- Energy data in open json format
- Access to the EnergyScopium SaaS platform at no extra cost
- Access to support

NB: all the API commands described and the proposed documents are available under

```
/gpfs/home/energyscopium/energyscopium
```

The banner

How to use the service

Service enabled or disabled by manager

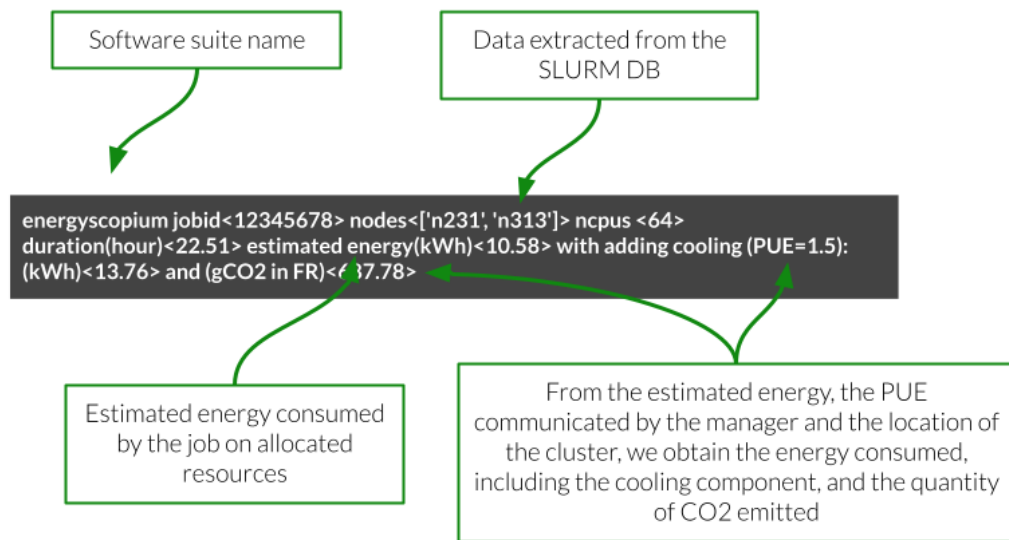
If the service has been enabled by the manager, the user can disable this service for his current job by inserting an export.

```
export ENERGYSKOPIUM_OFF=1
```

To re-enable the service, and if the environment variable is persistent

```
unset ENERGYSKOPIUM_OFF
```

Interpreting data



www.denergium.com

Note: the estimate takes into account jobs running simultaneously on the servers, providing a granular estimate of the work involved.

Energy report and energy data

The energy report and energy data are available on D+1

The aim of the energy report is to provide a synthetic, user-friendly view of the work's energy behavior.

The aim of the energy data via the eprofile and estat files is to provide an exhaustive and precise view of the work's energy behavior.

How to use the service

```
energyscopium_monitor_getjobedata.sh $jobid
```

Si le \$jobid n'est mis en argument
energyscopium: error: must provide JOBID

Un répertoire de réception peut être spécifié avec un second argument
`energyscopium_monitor_getjobedata.sh 9081340 ./reports/`
(le répertoire choisi doit exister)

The call returns:

- if the report is available

```
energyscopium_monitor_getjobedata.sh 9081340  
(attente éventuelle le temps de générer le rapport)  
es: info: data available  
es: info: saving <./energy scopium_result_9081340.tar.gz>
```

- if the jobid is not in the database

```
energyscopium_monitor_getjobedata.sh 010101010  
es: warning: bad data returned for this jobid
```

- if the jobid has no energy data (for the moment if the job is not finished, or never if the collection system was not activated)

```
energyscopium_monitor_getjobedata.sh 9658517  
es: warning: jobid not ended and/or not started
```

- if the job has been completed but energy data are not available (energy data for reports are available on D+1)

```
energyscopium_monitor_getjobedata.sh 9658516  
es: info: not yet possible
```

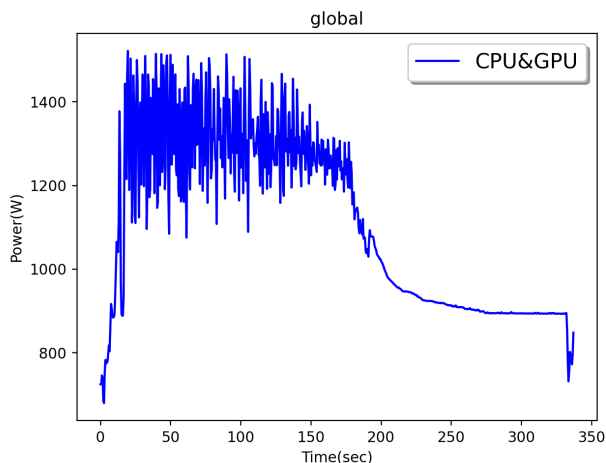

Interpreting data

Tarball contents

```
tar tvf energyscopium_result_12345678.tar.gz  
12345678/energyscopium_estat_12345678.txt  
12345678/energyscopium_eprofile_12345678.txt  
12345678/energyscopium_report_12345678.html  
12345678/image/energyscopium_plot_12345678_global.png  
12345678/image/energyscopium_plot_12345678_n035-1.png  
12345678/image/energyscopium_plot_12345678_n035-0.png
```

To view the html report, simply open it in your favorite browser.

- For further information, please refer to the appendix.
- A label is used to position the job on an efficiency scale.
- Plots show how the job behaves during execution, at the grain of the processors and GPUs.



The profile and estat files are in ASCII, JSON format.

- Elements of understanding are given in the appendix.
- An example code makes it easy to insert these forms into your own processes.

```
python3 energyscopium_check_e.py
```

Limitations

- Manager's rule: only the owner of the job and the manager have the right to access the job's energy data. Protection is in place against accessing the energy data of others, and bypassing this protection may result in sanctions.
- The current version of the EnergyScopium solution delivers biased data when several jobs are running simultaneously on the servers.

Access to the EnergyScopium SaaS platform for data mining

Users of the EnergyScopium Monitor software suite can access the EnergyScopium platform (SaaS) at no extra cost to :

- deposit their energy data
- explore their energy data

To do this, users must :

- open an account on the EnergyScopium Server (SaaS) platform
- download the user identification file generated by the platform
- upload the job's energy data tarball to EnergyScopium Server (SaaS)

```
EnergyScopium Server (SaaS) : www.energyscopium.com
```

Dedicated documentation guides users through account opening and initial use.

EnergyScopium_saas_onboarding_latest.pdf

Support access

For help or problems concerning the use of the EnergyScopium solution on the cluster, we strongly advise you to use the cluster support system. The manager will relay the request to EnergyScopium support if necessary.

For help or problems interpreting energy data, please use support@energyscopium.com, indicating (at least) the cluster name.

For help or problems using the EnergyScopium SaaS platform, please use support@energyscopium.com

Annexes

Elements for understanding an energy report

- TDP: (source Intel) TDP stands for Thermal Design Power, in watts, and refers to the power consumption under the maximum theoretical load. Power consumption is less than TDP under lower loads. The TDP is the maximum power that one should be designing the system for. This ensures operation to published specs under the maximum theoretical workload.
- Application energy consumption measurement: the sum of all measurements made on each processor and GPU present during application execution
- Global application energy consumption estimation: Based on the data read, a model developed by DENERGIUM estimates the server's energy consumption during application execution. The model is parameterized for computing servers. The thermal part (cooling system) is not included in the model.
- The energy efficiency label was introduced by DENERGIUM in 2023. The label is based on:
 - Measure: the ratio of energy consumed by Processors (packages) and GPUs (all) divided by the sum of TDP (Thermal Design Power)
 - Indicator: The ratio is a relevant indicator for measuring energy efficiency
 - Limitation: The ratio is not sufficient to assess whether an application is energy-optimized
- Period (ms): mean of the energy acquisition period
- Acquisition quality (low, medium, high): indicates the quality of the acquisition in relation to the acquisition period. High means that the initial period was obtained during the entire application run, Low means that several periods were detected without energy data. This is usually due to system overload. The consequence may be minor errors in the data generated. As acquisition quality is not normalized according to the duration of the application, a long-running application may have low acquisition quality, even if the acquisition has been running nominally for 99% of the runtime.
- Information dumped: the energy data available depends on the hardware technology. Intel processors provide the CPU (package) and DRAM, AMD processors provide the CPU (package), NVIDIA GPUs provide an overall module value

eprofile format

```
{
  "date": "2023/02/15 18:47:02",
  "pgm": "energyscope@inria",
  "duration(sec)": 105,
  "dateofrun": "2023/02/15 18:04:57.736851",
  "from_appli": {
    "jobid": "0",
    "user": "",
    "command": ""
  }
}
```

```
"comment": "",
"project": "",
"cluster": "",
"partition": ""
},
"data": {
  "data": {
    "ebuild": {
      "cutoff_b": [],
      "cutoff_e": [],
      "cutoff_be": [],
      "exclusive": "",
      "method": "none"
    },
    "eprofile_period(ms)": 500,
    "nb_sample": 211,
    "tags": {
      "es_total": {
        "start": "2023/02/15 18:04:57.736851",
        "stop": "2023/02/15 18:06:43.736851",
        "start(sec)": 0,
        "stop(sec)": 105,
        "joule(J)": 483.976
      },
      "es_appli_total": {
        "start": "2023/02/15 18:05:01.000000",
        "stop": "2023/02/15 18:06:35.000000",
        "joule(J)": 426.586,
        "start(sec)": 3,
        "stop(sec)": 97
      },
      "point1": {
        "start": "2023/02/15 18:05:55.000000",
        "stop": "2023/02/15 18:06:21.000000",
        "joule(J)": 117.441,
        "start(sec)": 57,
        "stop(sec)": 83
      }
    }
  },
  "etotal(W)": [
    4.405,
    4.405,
    4.944,
    4.047,
    3.924,
    6.469,
    ...
    1.442,
    1.299,
    1.246,
    1.275
  ]
}
```



```
}  
}  
}  
}  
}  
}  
}
```

estat format

```
{  
  "date": "2023/02/15 18:47:02",  
  "pgm": "energyscope@inria",  
  "duration(sec)": 105,  
  "dateofrun": "2023/02/15 18:04:57.736851",  
  "from_appli": {  
    "jobid": "0",  
    "user": "",  
    "command": "",  
    "comment": "",  
    "project": "",  
    "cluster": "",  
    "partition": ""  
  },  
  "arch": {  
    "data": {  
      "total": {  
        "node": 1,  
        "phy": 1,  
        "core": 4,  
        "gpu": 0  
      },  
      "node_list": [  
        "SEDBOR-Mathieu"  
      ],  
      "joule(J)": {  
        "etotal(J)": 484,  
        "ecpu(J)": 351,  
        "edram(J)": 133,  
        "etotal(%)": 100.0  
      },  
      "energy_est": {  
        "etotal(J)": 3634,  
        "etotal(kWh)": 0.001,  
        "model": "joule_model_202203 and pue=1.3",  
        "etotal_pue(kWh)": 0.001  
      }  
    }  
  },  
}
```

```
"carbon":{
  "country": "FR",
  "carbon(gCO2)": 0.051,
  "model": "carbon_model_202203"
},
"acquisition_period":{
  "start": "2023/02/15 18:04:57.736851",
  "stop": "2023/02/15 18:06:43.736851"
}
},
"arch":{
  "SEDBOR-Mathieu":{
    "data":{
      "system":{
        "os": "",
        "fgovernor": "",
        "fmin": 0.0,
        "fmax": 0.0
      },
      "acquisition(ms)":{
        "min": 511.0,
        "max": 542.0,
        "mean": 533.97
      },
      "acquisition_hole": [],
      "power(W)":{
        "min": 3.743,
        "max": 17.302,
        "mean": 4.616
      }
    }
  },
  "arch":{
    "0":{
      "arch":{
        "0":{
          "arch": {},
          "data": {}
        },
        "1":{
          "arch": {},
          "data": {}
        },
        "2":{
          "arch": {},
          "data": {}
        },
        "3":{
          "arch": {},
          "data": {}
        }
      }
    },
    "data":{
      "modelname": "Intel(R)Core(TM)i7-8665UCPU@1.90GHz",
```

```
"core": 4,  
"joule(J)": {  
  "etotal(J)": 484,  
  "ecpu(J)": 351,  
  "edram(J)": 133,  
  "etotal(%)": 100.0  
},  
"tdp(W)": 15  
}  
}  
}  
}  
}  
}
```