

ProfiT-HPC

Design and Test of modular Evaluation and Visualization Tools

Performance Engineering Workshop
Dresden, 25-26 March, 2019

Overview of this Presentation

1. Python3

- PDF Report Generator
- Results summary, comparisons
- Graphics

2. Web-based and Interaktive: Grafana

- Grafana Dashboards
- Global Overview
- Example detail visualizations

PDF Report Generator

- Clearly structured summary including graphics
- No user access through HTTP/HTTPS necessary
- Automatic, Modular and Flexible

PDF Report Generator: Features

- List of important job information
 - Job ID
 - General time details
 - Requested resources
 - Used hardware components
 - Queue
 - Number of nodes
 - Processors / Cores
- Recommendations and tips concerning performance and usage of resources

PDF Report Generator: Features

- Grafical representation
 - Utilization analysis of allocated resources
 - Load distributions
 - Time series plots

PDF-Report-Generator: Requirements

- Important Open Source Python modules
 - json: Data management of job related results received from aggregator
 - time, numpy, pandas: Management of matrices and data types (for example, time formats, data structures, efficient data analysis for python interpreter)
 - matplotlib: Creation of plots (export in SVG format: Scalable Vector Grafics)
 - svglib: Reading and conversion of SVG grafic files
 - reportlab: Creation of printable PDF report

PDF Report Generator: Data Reduction

- Job results from aggregator
 - Reduced data per node
 - global values
 - time series
- Further reduction of graphical representation
 - Global values of the job
 - Node statistics (minimum, maximum, average)

PDF Report Generator: Example

Profit-HPC Report
Report for JobID: 2368599
Time of generation: 28/02/2019 10:01:26




Profit-HPC Report

Job Overview:

Job-ID : 2368599
User name : akhuziyi
Queue : int
Number of nodes : 3
Requested cores : 13
Requested time : 3.00 h
Used time : 1.36 h
Time of job start : 01/10/2018 14:30:00
Time of completion : 01/10/2018 15:51:25

Node Information:

CPU model:
Intel(R) Xeon(R) CPU E5-2650 v3 @ 2.30GHz
Memory per node: 67 GB
Sockets per node: 2
Cores per socket: 10
Threads per core: 1



Global Summary of Resource Usage

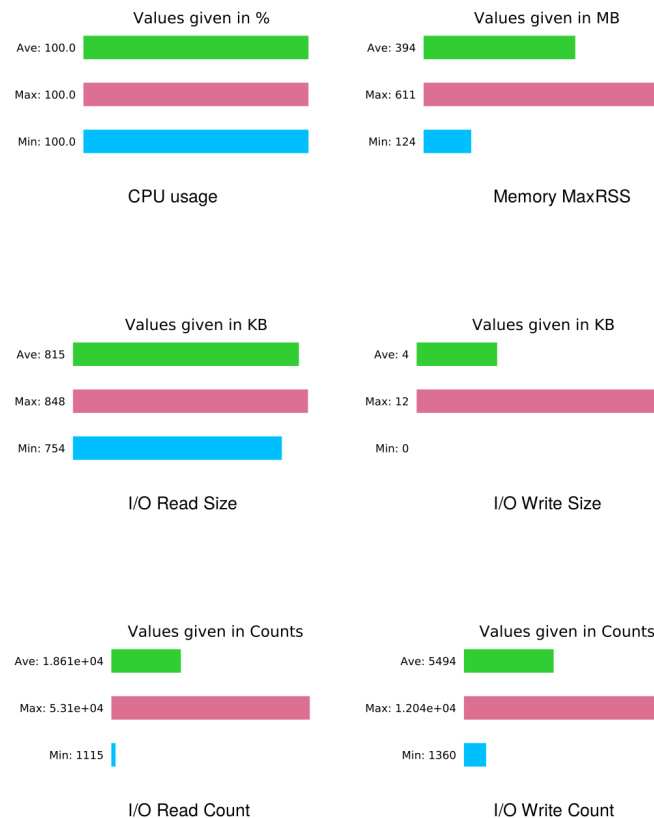
Recommendations:

PDF Report Generator: Example

Profit-HPC Report
Report for JobID: 2368599
Time of generation: 28/02/2019 10:01:26



Node Distributions (average, maximum, minimum):




Gefördert durch die Deutsche Forschungsgemeinschaft (DFG)
KO 3394/14-1, OL 241/3-1, RE 1389/9-1, VO 1262/1-1, YA 191/10-1

2/3

Profit-HPC Report
Report for JobID: 2368599
Time of generation: 28/02/2019 10:01:26



Timeseries Plots:



Gefördert durch die Deutsche Forschungsgemeinschaft (DFG)
KO 3394/14-1, OL 241/3-1, RE 1389/9-1, VO 1262/1-1, YA 191/10-1

3/3

PDF Report Generator

Practical Experience:

- Simple installation of python3 / Loading of additional open-source packages
- High flexibility meaning easily modifiable and extendable code
- Automatic analysis and report generation executes relatively quickly. Test with data for 3 nodes: runtime approx. 4 s
- 3 page printable report with size approx. 130 KByte per job

Web based grafical interface with Grafana

- Administrator view / user view
- Multi-level views of information
 - Job based
 - eye-catching global view of results
 - details on distributions
 - time series details
 - User overview for administrator

Web based grafical interface with Grafana;

Test data volume



Measurements of Zuse Institute Berlin

Cluster: Cray XC30/40
 Number of nodes: 1872 (+ Login, Data, Post)
 Processor types: Intel IvyBridge + Haswell
 Number of cores per node: 24

Databank	Size [GB]	Measurement interval [s]	Duration [days]
zib-db01	5,6	60	3
zib-db02	9	60	1
zib-db03	17	60	1
zib-db04	20	60	1
zib-db05	14	60	1

Datenanzahl:

Number of nodes	Hardware (virtual) number per node	Measurment	Number of metrics	Measurement interval [s]	Number of measurements per day
1872	48	CPU	10	60	1,29E+09
1872	24	PFIT-UPROSTAT	50	60	3,23E+09
1872	1	MEM	10	60	2,70E+07
1872	1	SYSTEM	7	60	1,89E+07
1872	1	SWAP	6	60	1,62E+07
Estimated max. number of measurements per day					4,59E+09
Estimated size [GB]:					18,36

System information (ZIB): <https://www.hlrn.de/home/view/System3/CrayHardware>

Web based grafical interface using Grafana: Administrator view

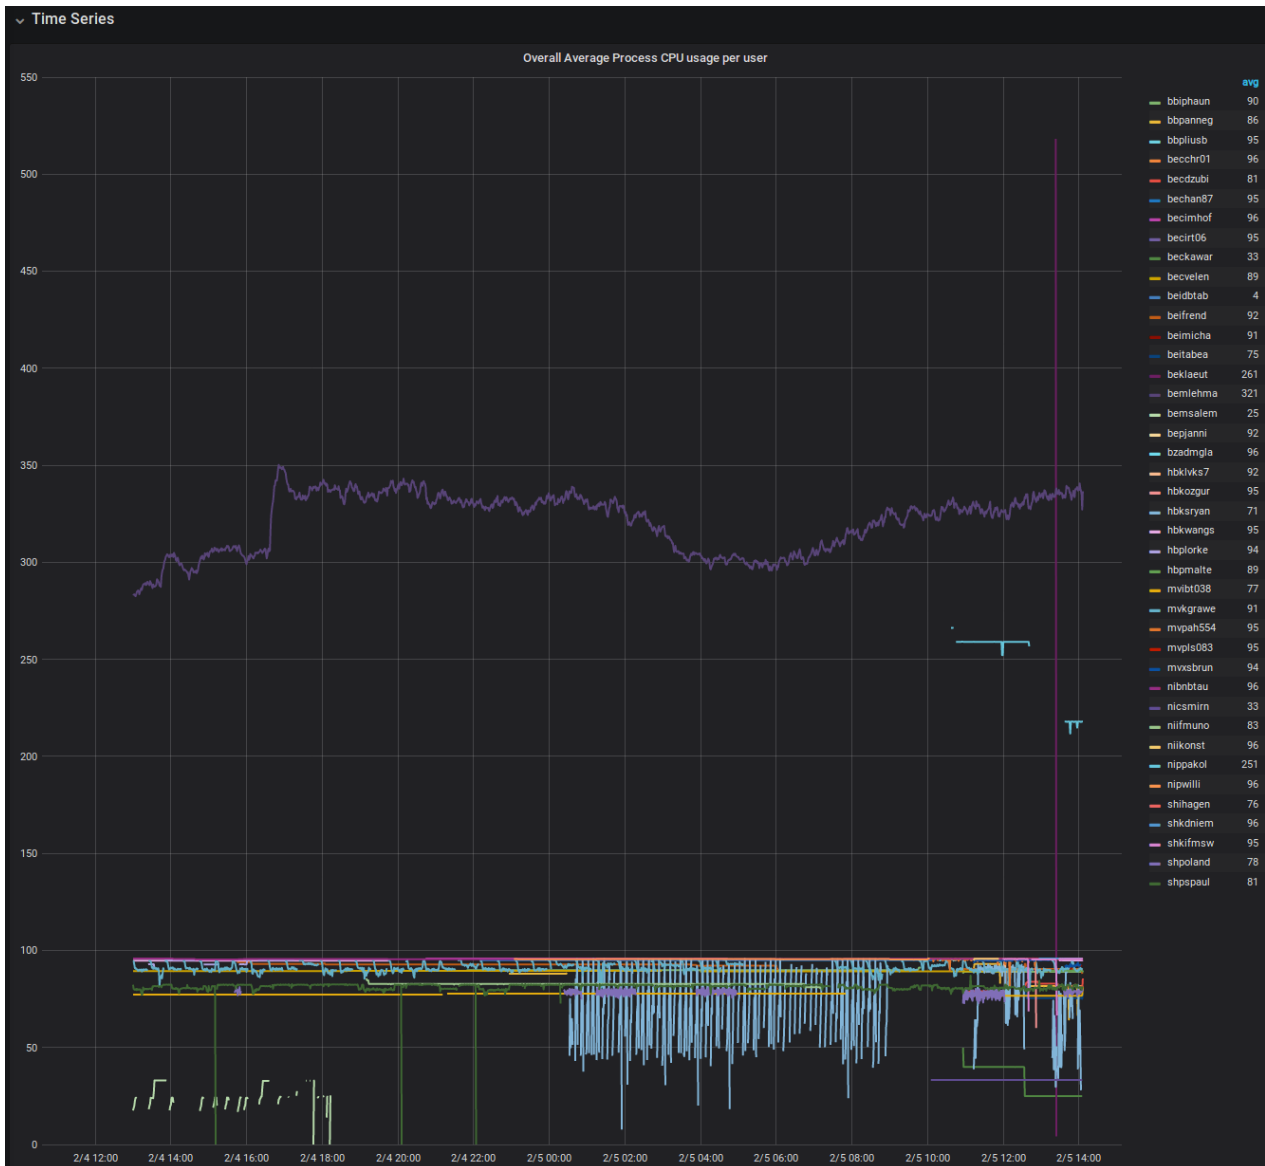
List of users | Average Process CPU usage

Metric	Avg
16435	559
16448	91
16449	90
16480	0
16495	11
16532	0
16539	0
16560	92
16615	0
16618	0
16619	0
16642	4
16654	79
16698	25
16702	0
16751	79
16790	83
16805	95
16806	2 K
16832	80
16836	93

Color-coded JobID list of CPU usage

- blue: Processes using more than 1 CPU core (mean CPU usage >> 105 %)
- green: 20 % < mean CPU usage < 105 %
- red: < 20 %

Web based grafical interface using Grafana: Administrator view

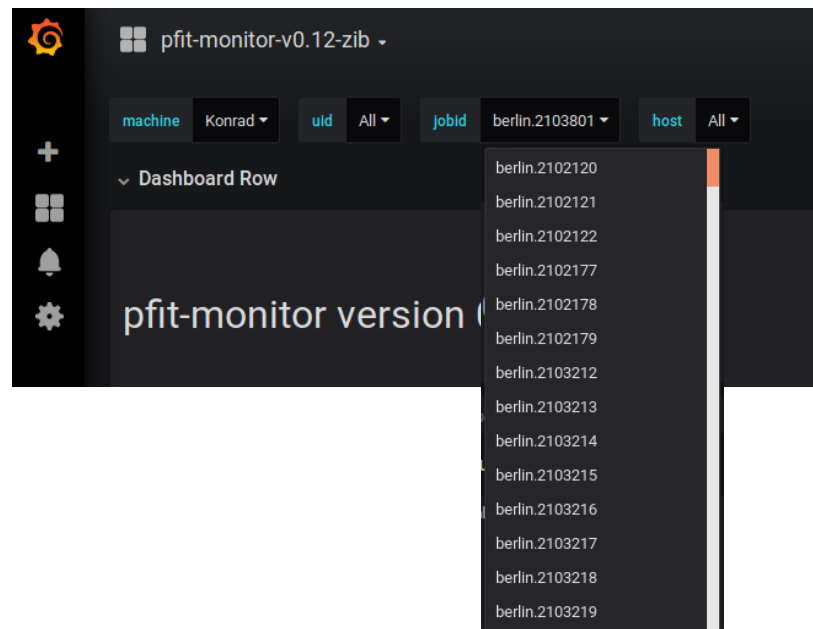


Problem with overview of many jobs and users:

The data base is so large, that the evaluation of the table and grafics causes notable delay.

Suggestion for improvement:
Global values could be extracted and saved in a separate data base with the aggregator for longer term access.

Web based grafical interface using Grafana: Administrator view



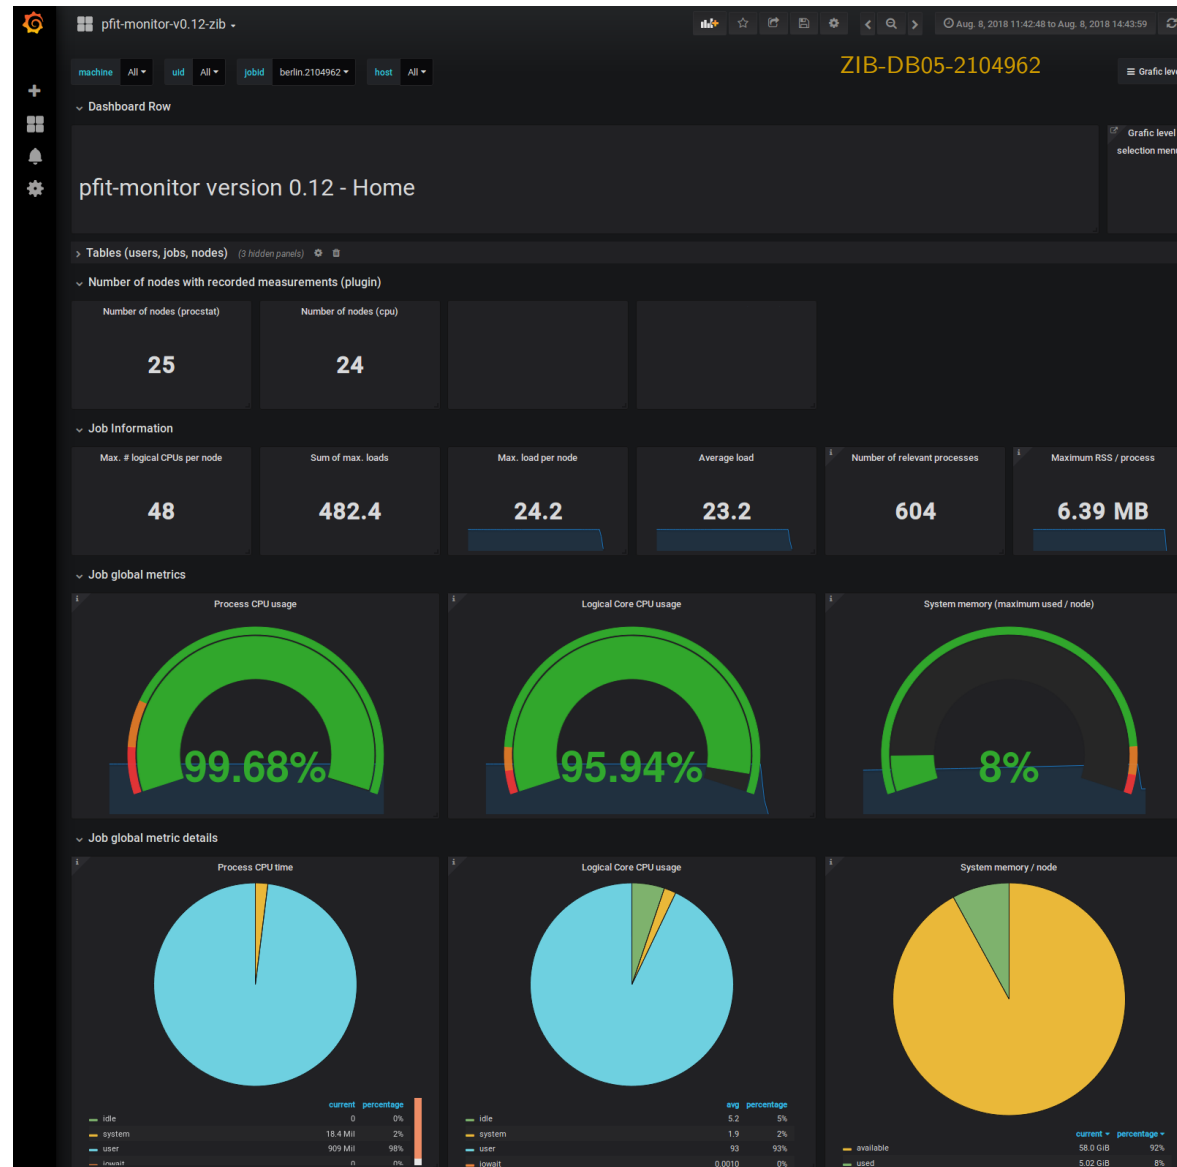
Administrator menu prompt:

- Overview of all users via UID
- Overview of all jobs via JobID

Web based graphical interface using Grafana: administrator view (single job)



Example 1



Web based grafical Interface using Grafana



Practical Experience:

- Visualization of the metrics of a single job is fast and clearly structured
- Visualization delays if
 - Large overview of users chosen (too many jobs being reduced at one time)
 - Too many nodes chosen, for example, displaying time series plots of all processes on all nodes for very large jobs

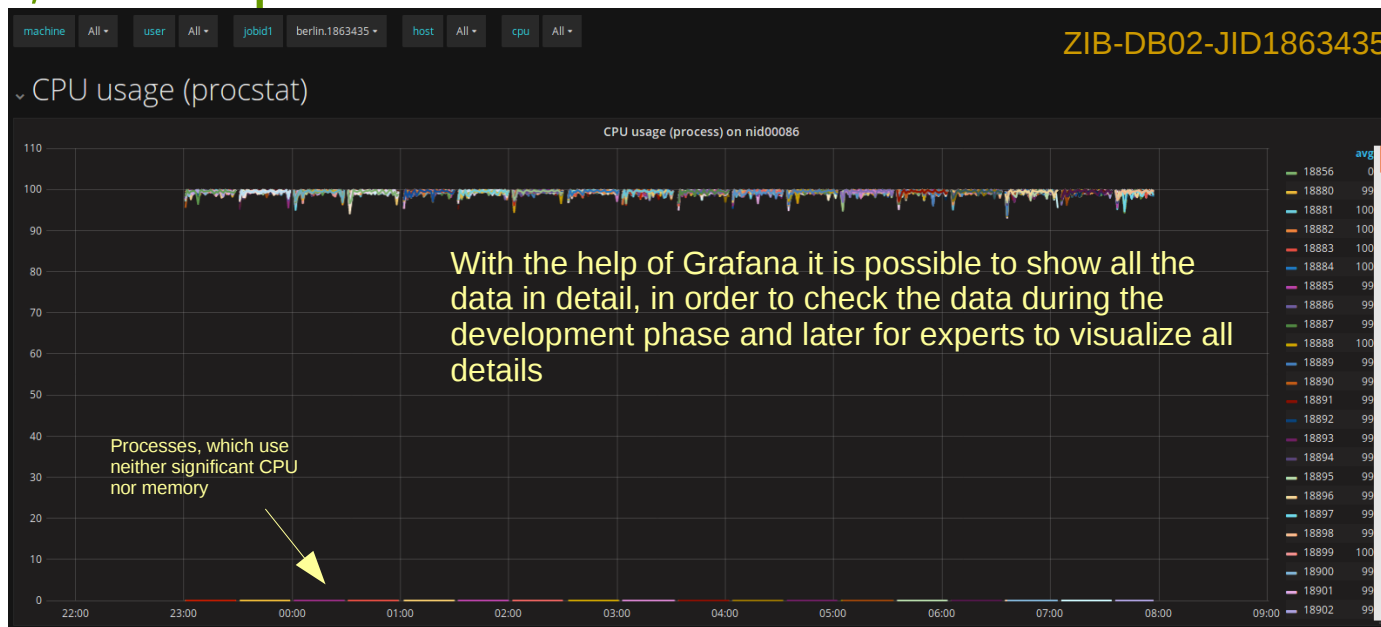
Web based grafical Interface using Grafana, Example 2



Eye-catching
evaluation of metrics
and indicators



Web based grafical Interface using Grafana, Example 2



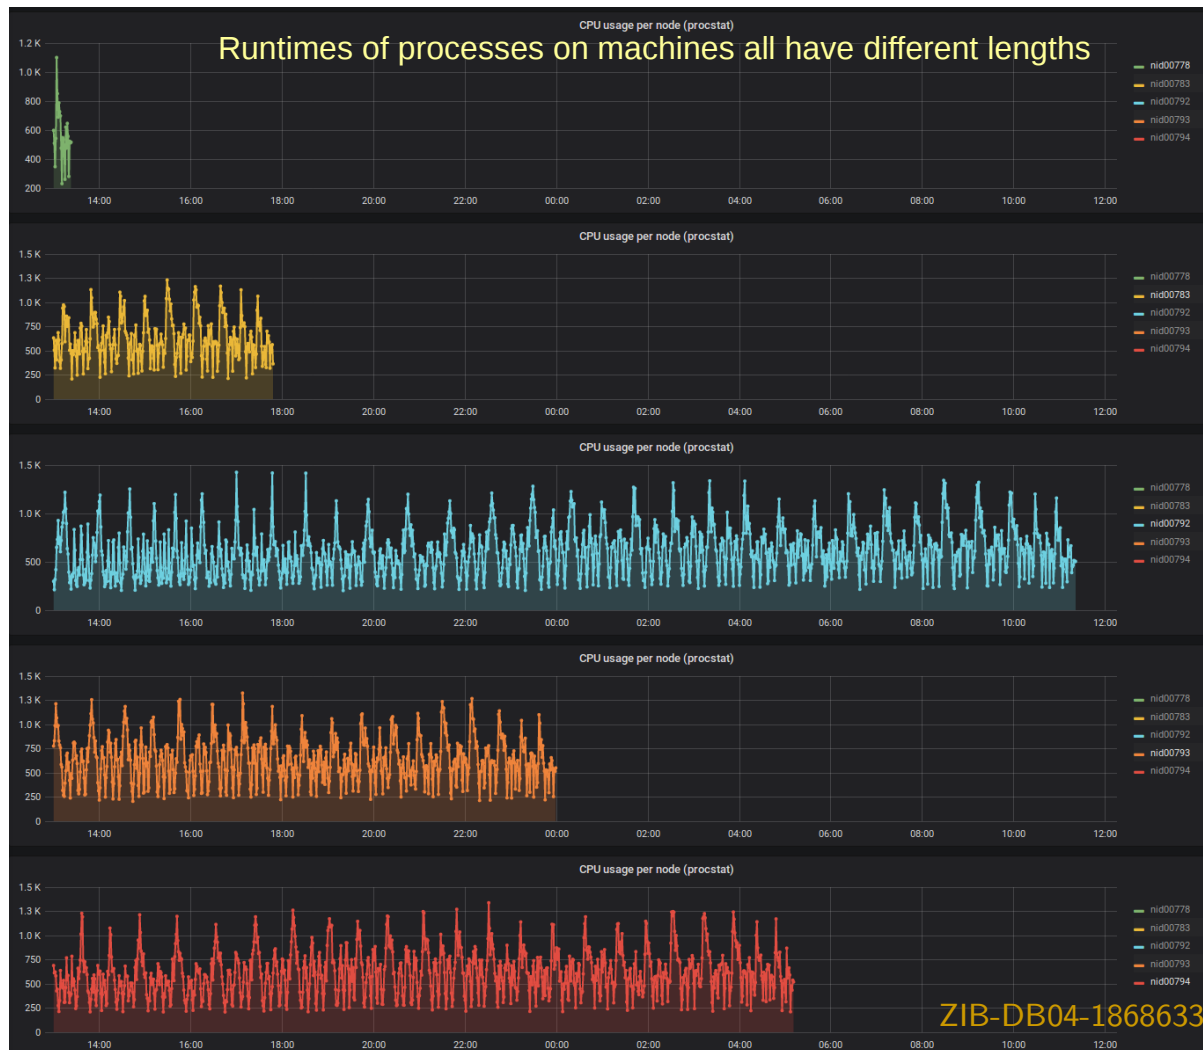
Web based grafical Interface using Grafana, Example 3

Color coding for
conspicuous metrics
and indicators



Web based grafical Interface using Grafana, Example 3

CPU usage




Web based grafical Interface using Grafana, Example 4



Web based grafical Interface using Grafana, Example 4

CPU usage and max. Memory RSS per Process



Summary

- Summary of fulfilled tasks and practical insight
 - Structured graphical representation of metrics and indicators allows quick and easy interpretation
 - PDF report offers a printable summary of the job and performance analysis without intricate infrastructure
 - Fast, stable, easy to install, extendable
 - Grafana offers an interactive and deep examination of the results
 - Administrator/User view (global interpretation, bar diagrams, time series)
 - Helped also during development phase of the project
 - Simple and comprehensible visualisation of the utilization of components for users without expert knowledge of HPC
 - Quick, stable, flexible, real-time and archived data (accessible both during and after job execution)

- Short-term perspectives
 - Additional metrics and indicators
 - Further testing and extension of PDF report generator
 - Implementation of automatic recommendations and tips