

Profit-HPC Toolkit

Data collection and storage

The data is collected by the metric collectors (Telegraf or PfiTCollect) and stored in the timeseries database (InfluxDB). The toolkit supports the most popular batchsystems, Slurm and LSF.

Data aggregation and job analysis

The aggregation routine reduces and exports the job metrics. Job analysis aims to detect pathological, low performance or inefficient resource allocation cases by means of defined decision trees.

Reports and recommendations

The toolkit provides *text* and *pdf* reports with performance data and recommendations for improvement based on job analysis. Interactive dashboards show aggregated metrics and indicators. They allow detailed realtime visualization using Grafana.

Structure of the toolkit and dataflow

The modules developed within Profit-HPC are shown in Figure 1 in green.

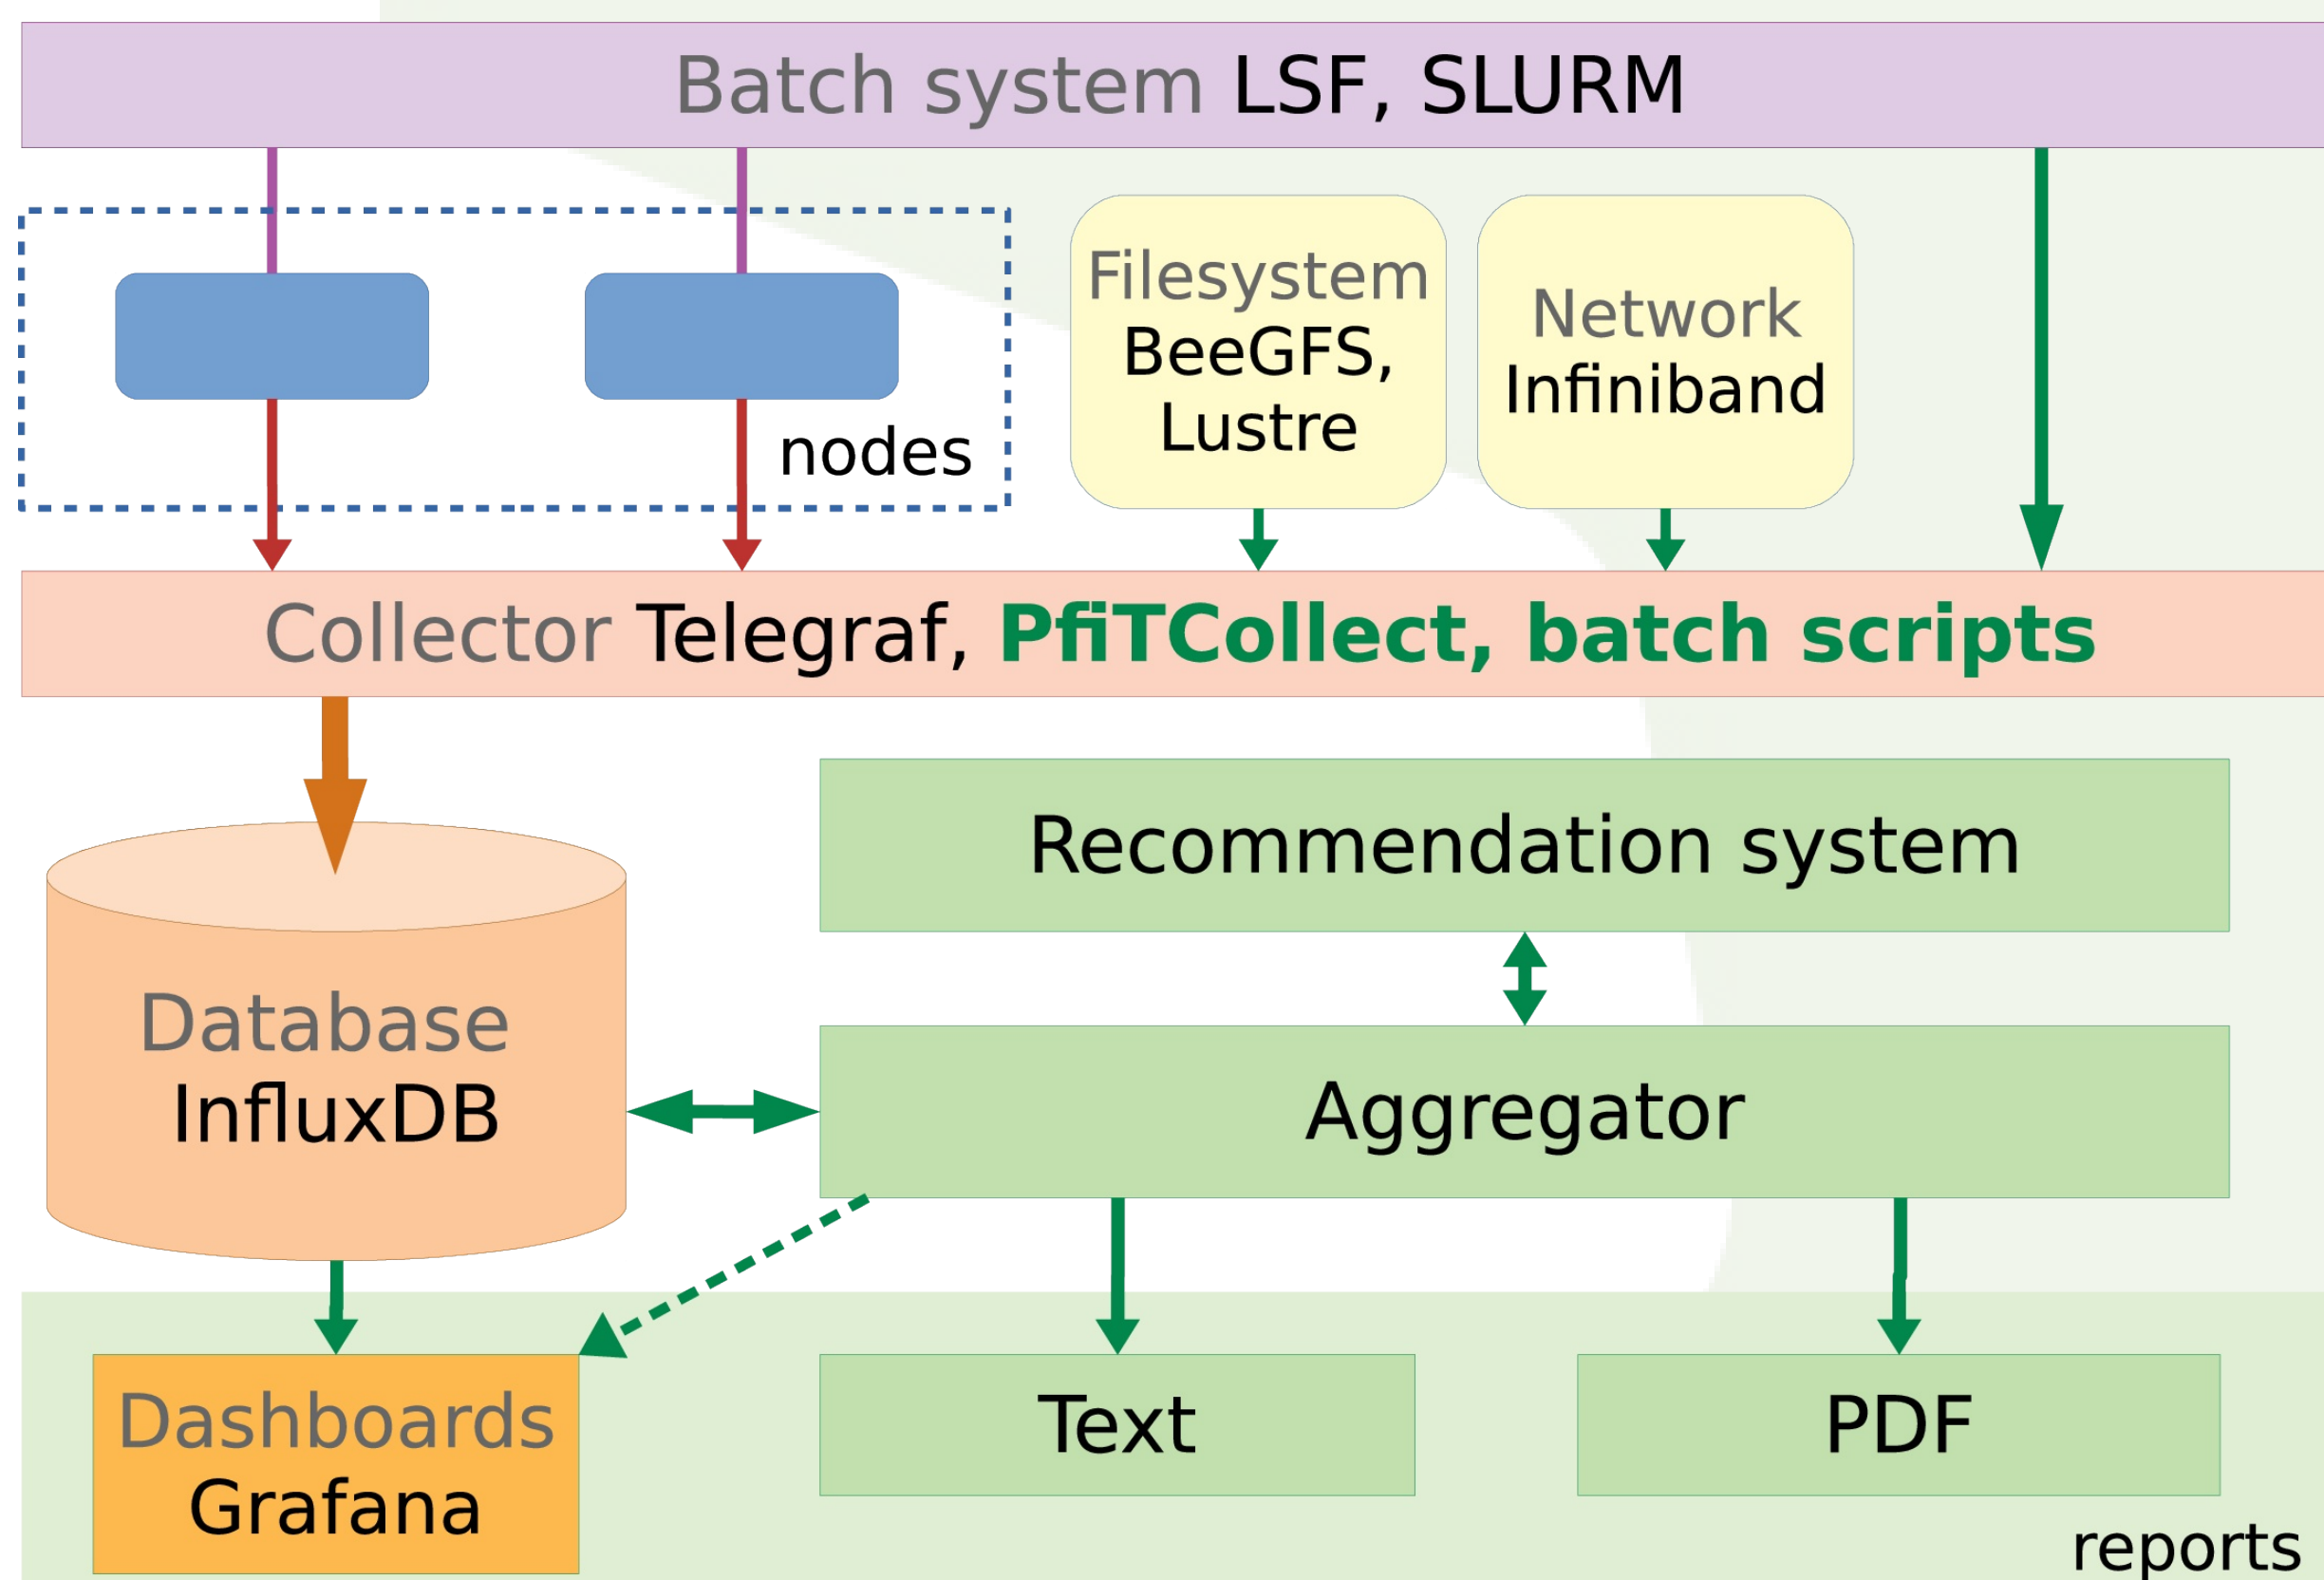


Fig. 1: The structure of Profit-HPC toolkit

Job Analysis

Recommendations

The toolkit uses programmable rules to provide recommendations for users. Rules are formed by *attributes*, which are properties of the job with a fixed set of values they can take. Particular combinations of attributes with specified values form the *rules*.

Decision trees

With such attributes, the rules can be represented as a decision tree. Figure 2 shows an example of a partial decision tree used to generate the recommendations in the presented reports. However, when rules become more complicated, building comprehensive decision trees is not a trivial process; various algorithms, such as *RBDT-1* [1], are used to generate shorter trees.

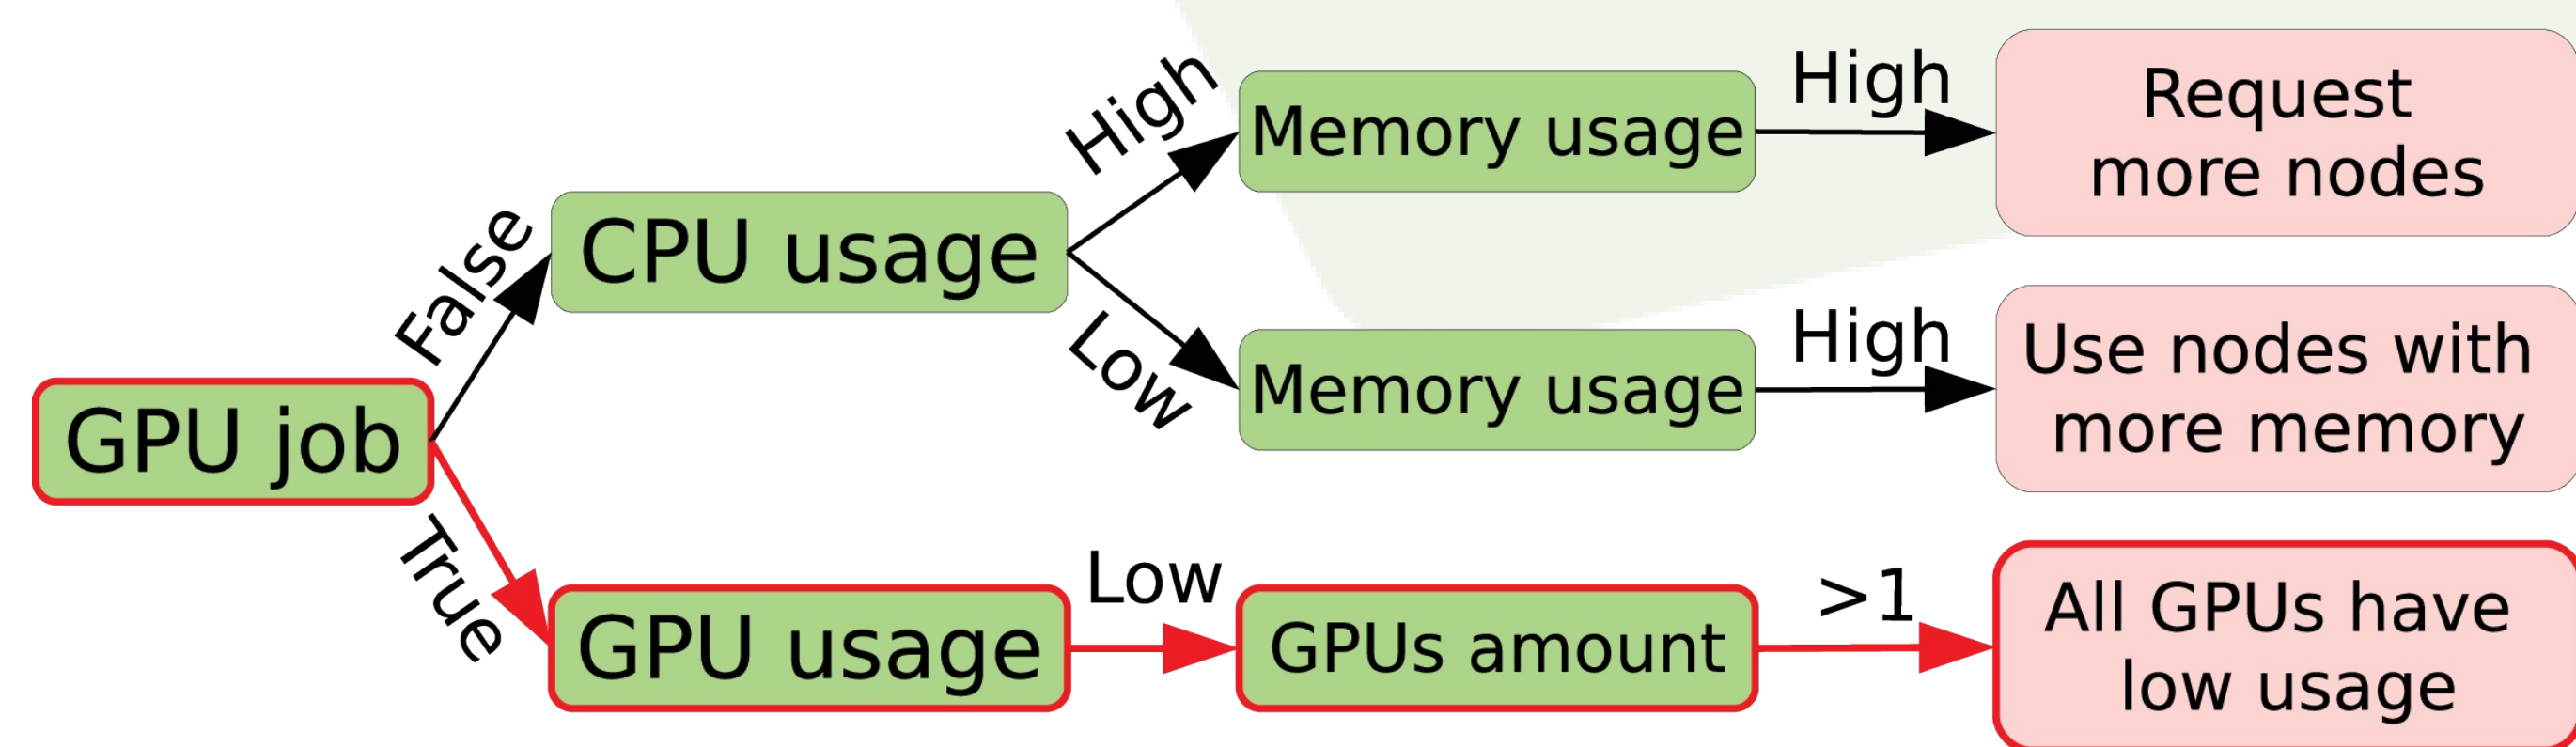


Fig. 2: The partial decision tree used in examples

[1] Amany Abdelhalim, Issa Traore, and Youssef Nakkabi. 2016. Creating Decision Trees from Rules using RBDT-1. *Comput. Intell.* 32, 2 (May 2016), 216-239. doi.org/10.1111/coin.12049

Goals

Increase awareness

of the importance of HPC applications performance (issues and parameters) for the large number of new users typical for Tier 3 systems; provide necessary background for transfer to Tier 2 systems

Automatic metric collection and data storage

with no or minor systems performance impact

Automatic performance analysis

of the job and generating recommendations for optimization

Comprehensive feedback

on the performance and resource consumption of a job with reports of different levels of detail (including optimization recommendations).

Reports and Recommendations

Text report

provides a short summary of the job with various metrics and recommendations.

The report is designed to be configurable. The text report can be printed at the end of the job to user's shell, saved as a text file or sent via email.

Recommendations indicate possible problems in the job and can be used as a starting point of investigating job issues.

Additional time series are available in the other report types.

PDF report

provides additional information in a more comprehensive way.

It includes distributions (5 point boxplots) and time series data.

The first page is sufficient to get a short summary of the job and recommendations.

General Information

User Name:theuser

JobID:2837739

Used Queue/Partition:gpu

Used Nodes:12

Requested Compute Units:12

...

CPU utilization. Mean values + high water mark (hwm) values

nodes	cores	CPU usage total[%]	hwm[%]	memory used hwm[Gib]	memory mean[Gib]	memory alloc[Gib]
dge009	4	100	100	0.60	0.59	20.97
...						

GPU utilization. Mean values + high water mark (hwm) values

nodes	ID [bus]	GPU usage mean[%]	usage hwm[%]	memory hwm[Gib]	# proc total	CPU usage hwm[%]
dge009	02	1	4	4	2	50
...						

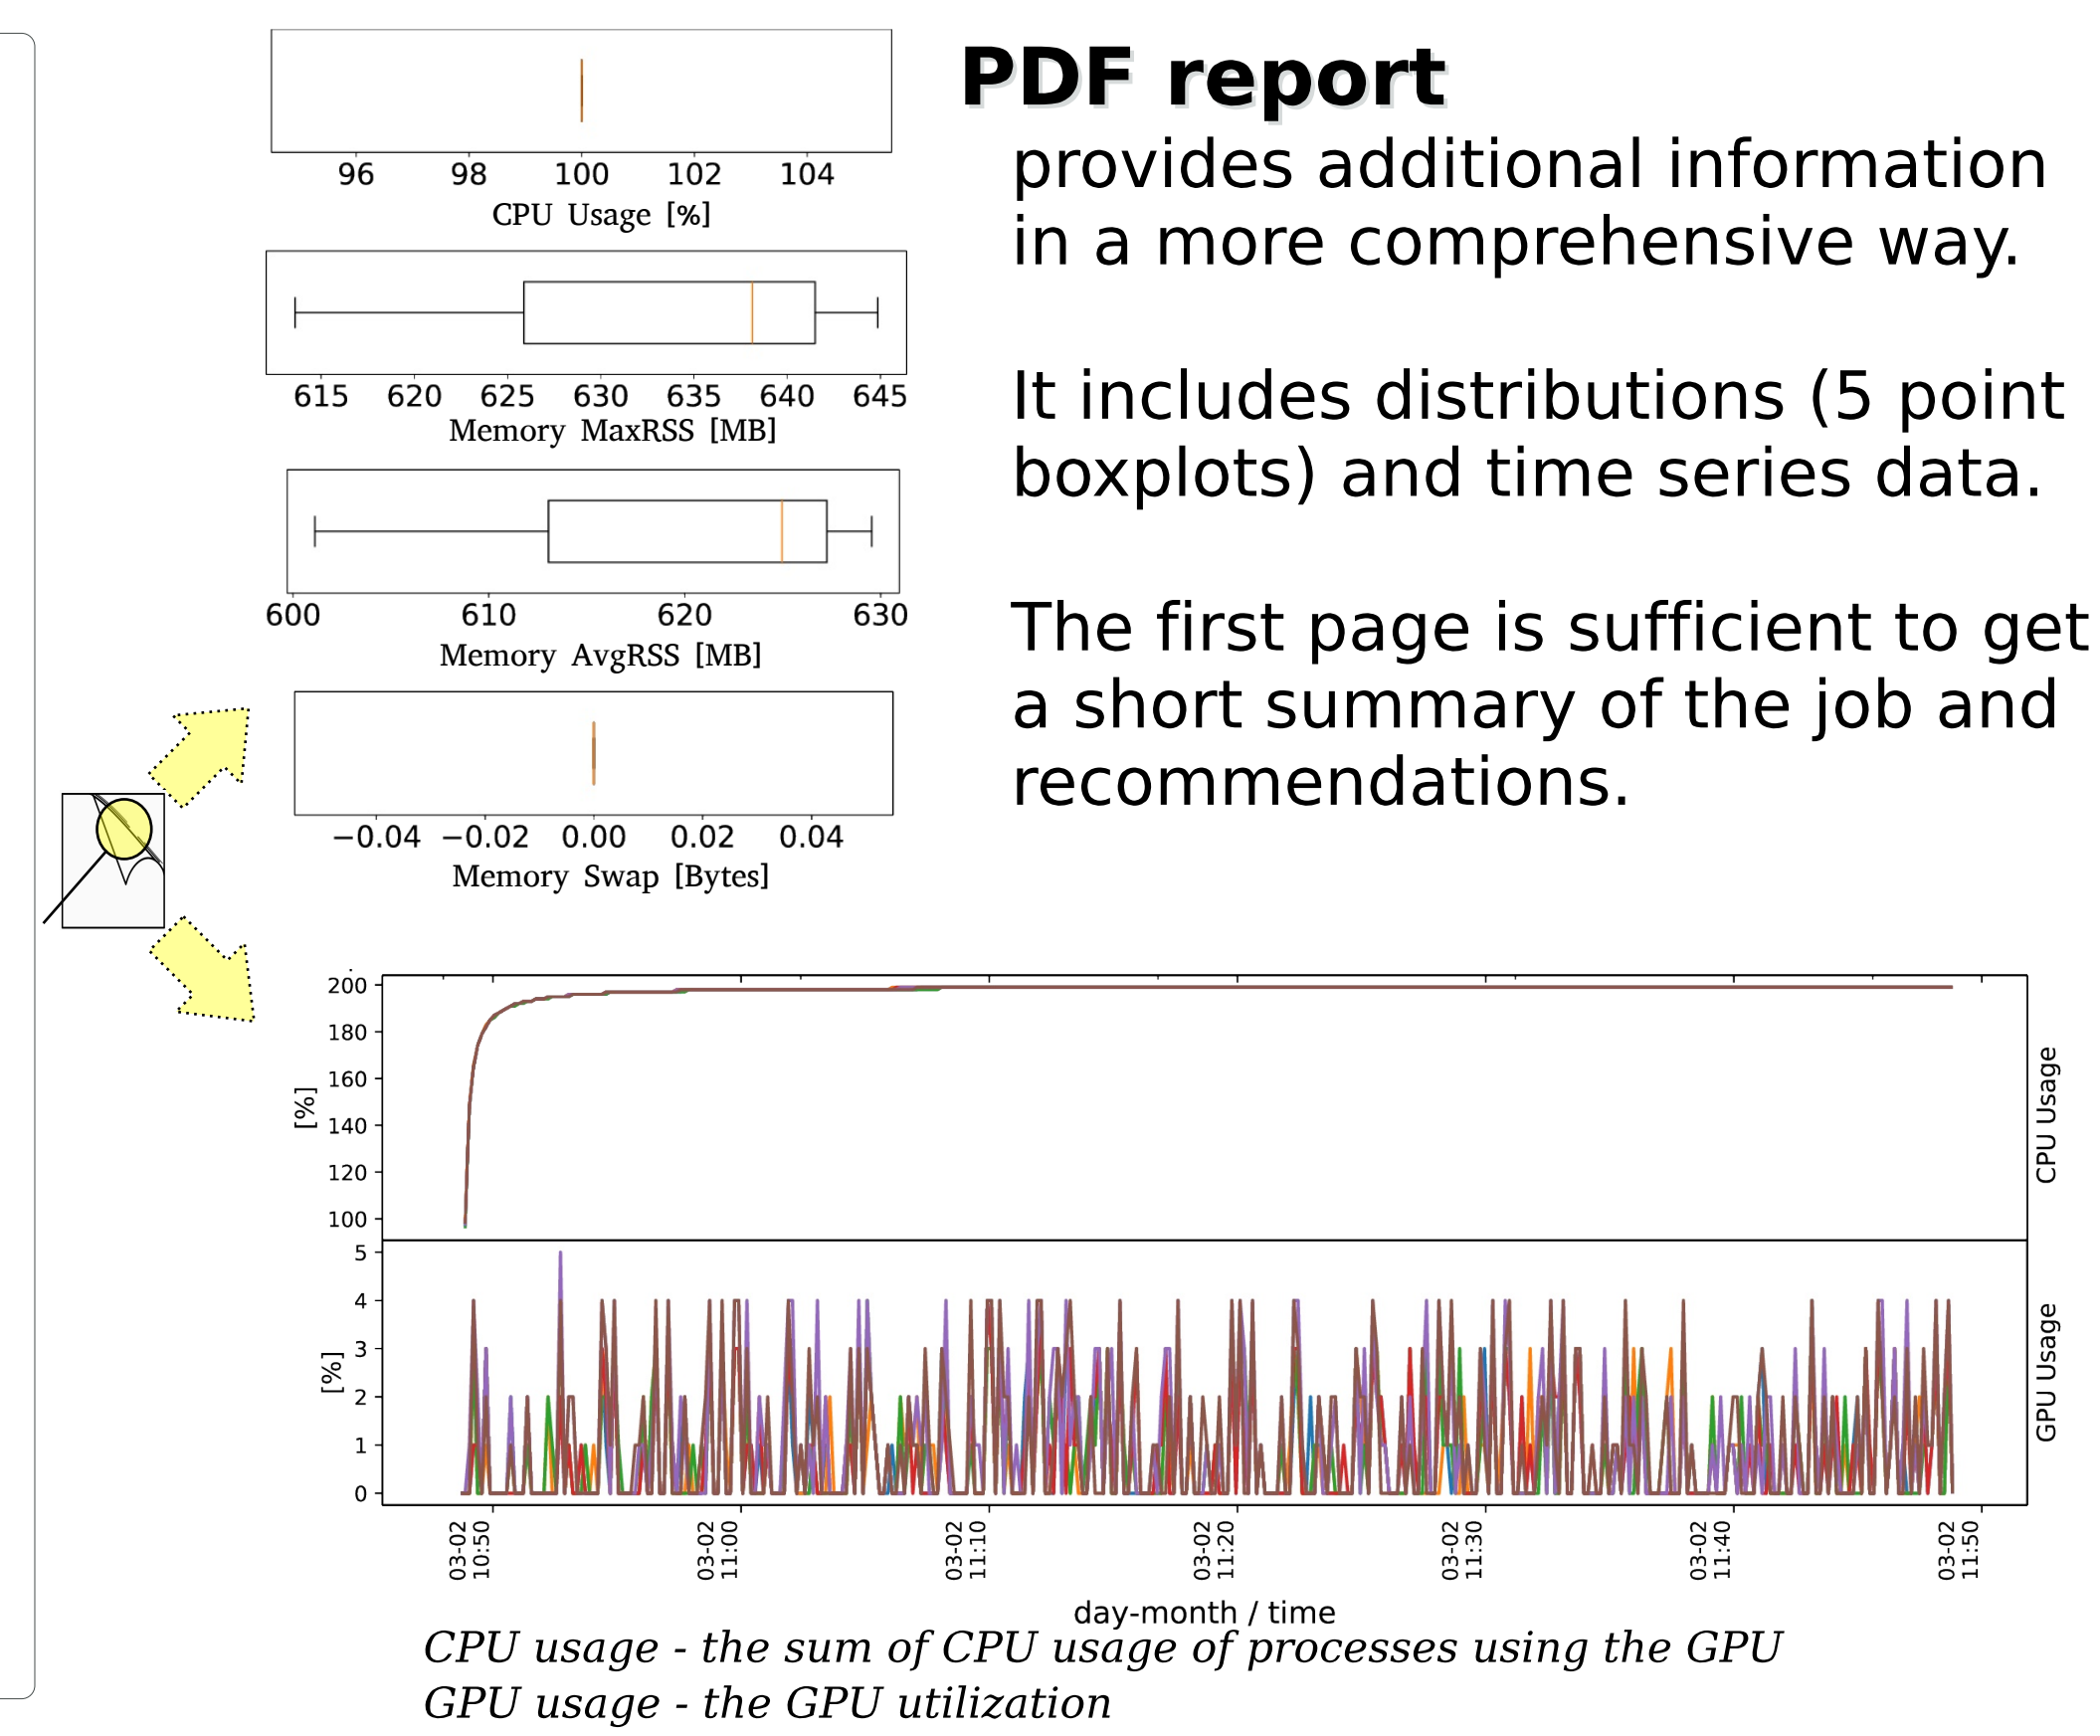
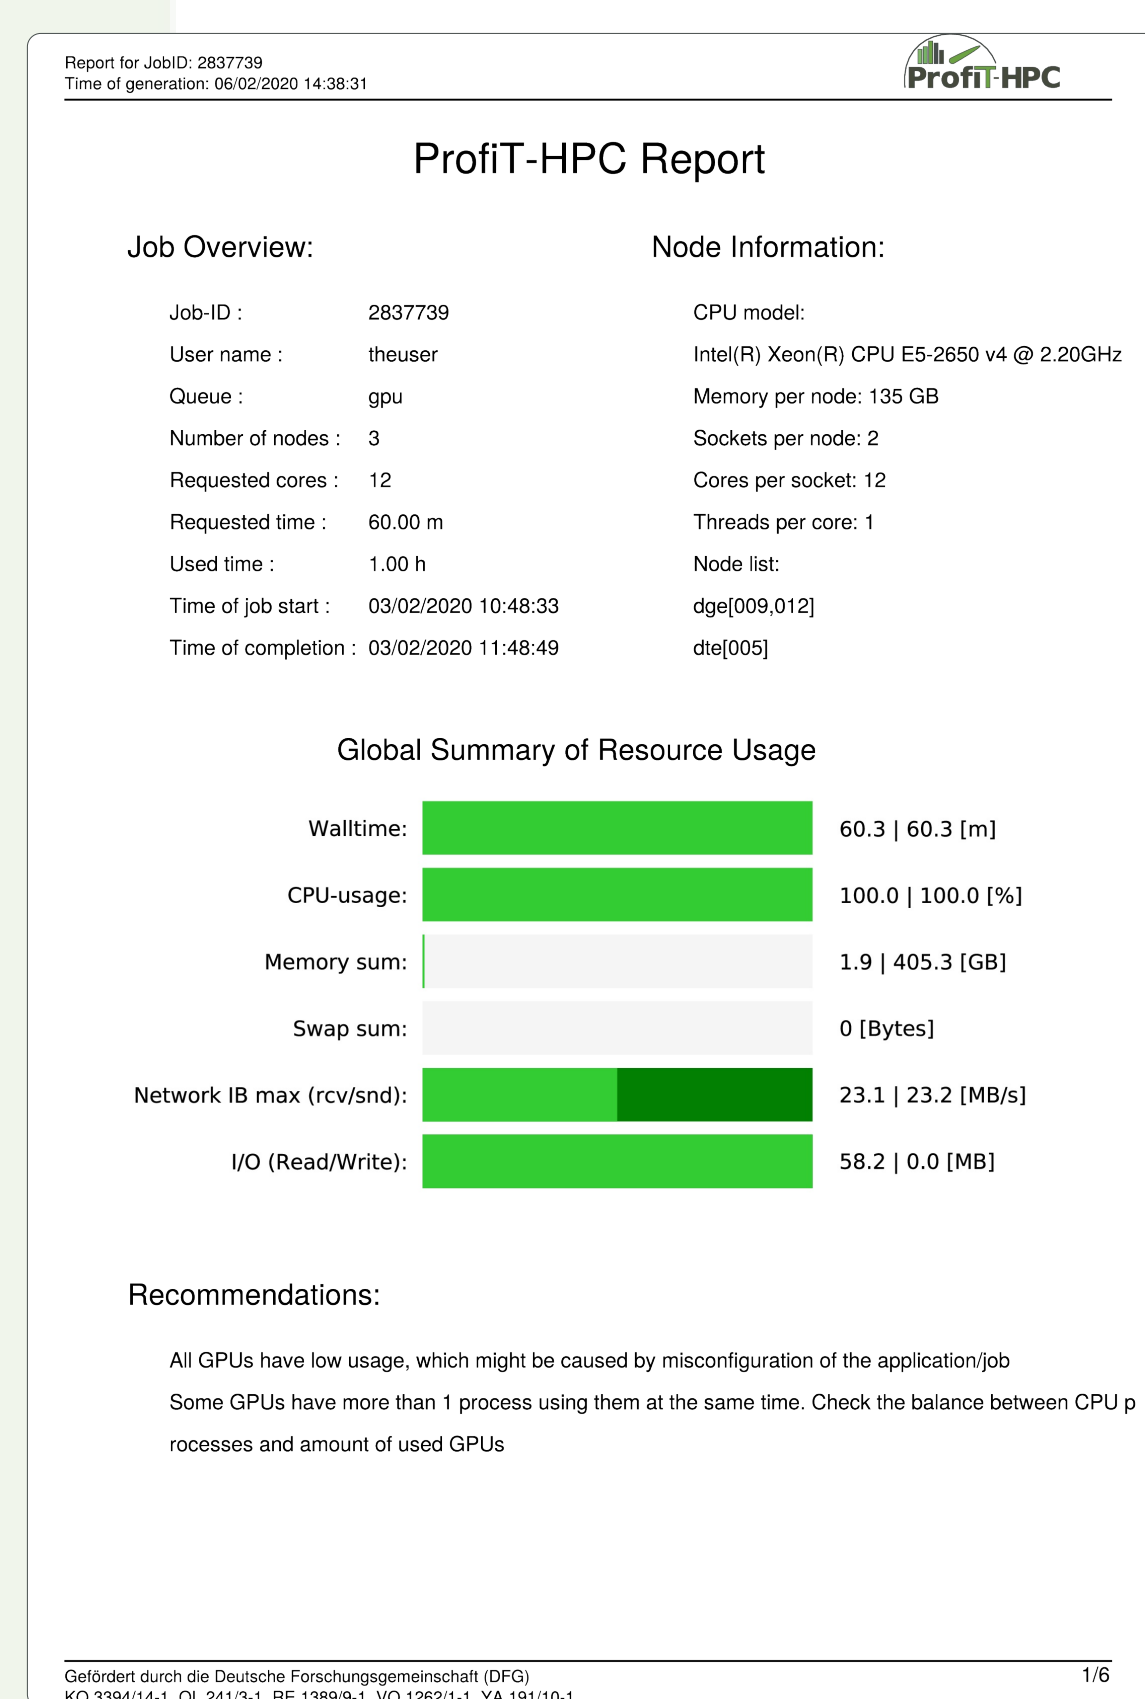
Per node IO and network. Sums and rates

nodes	read (sum)	IO write (sum)	receive[sec]	Infiniband send[sec]
dge009	23.88 MiB	0 B	17.37 MiB	17.28 MiB
...				

Recommendations

- All GPUs have low usage, which might be caused by misconfiguration of the application/job

- Some GPUs have more than 1 process using them at the same time. Check the balance between CPU processes and amount of used GPUs



Interactive dashboards

Dashboards are accessible during or after the job execution, for instance, as a color coded global job indicators or as process (or node) usage distribution as bar chart diagrams shown in the Figure 3.

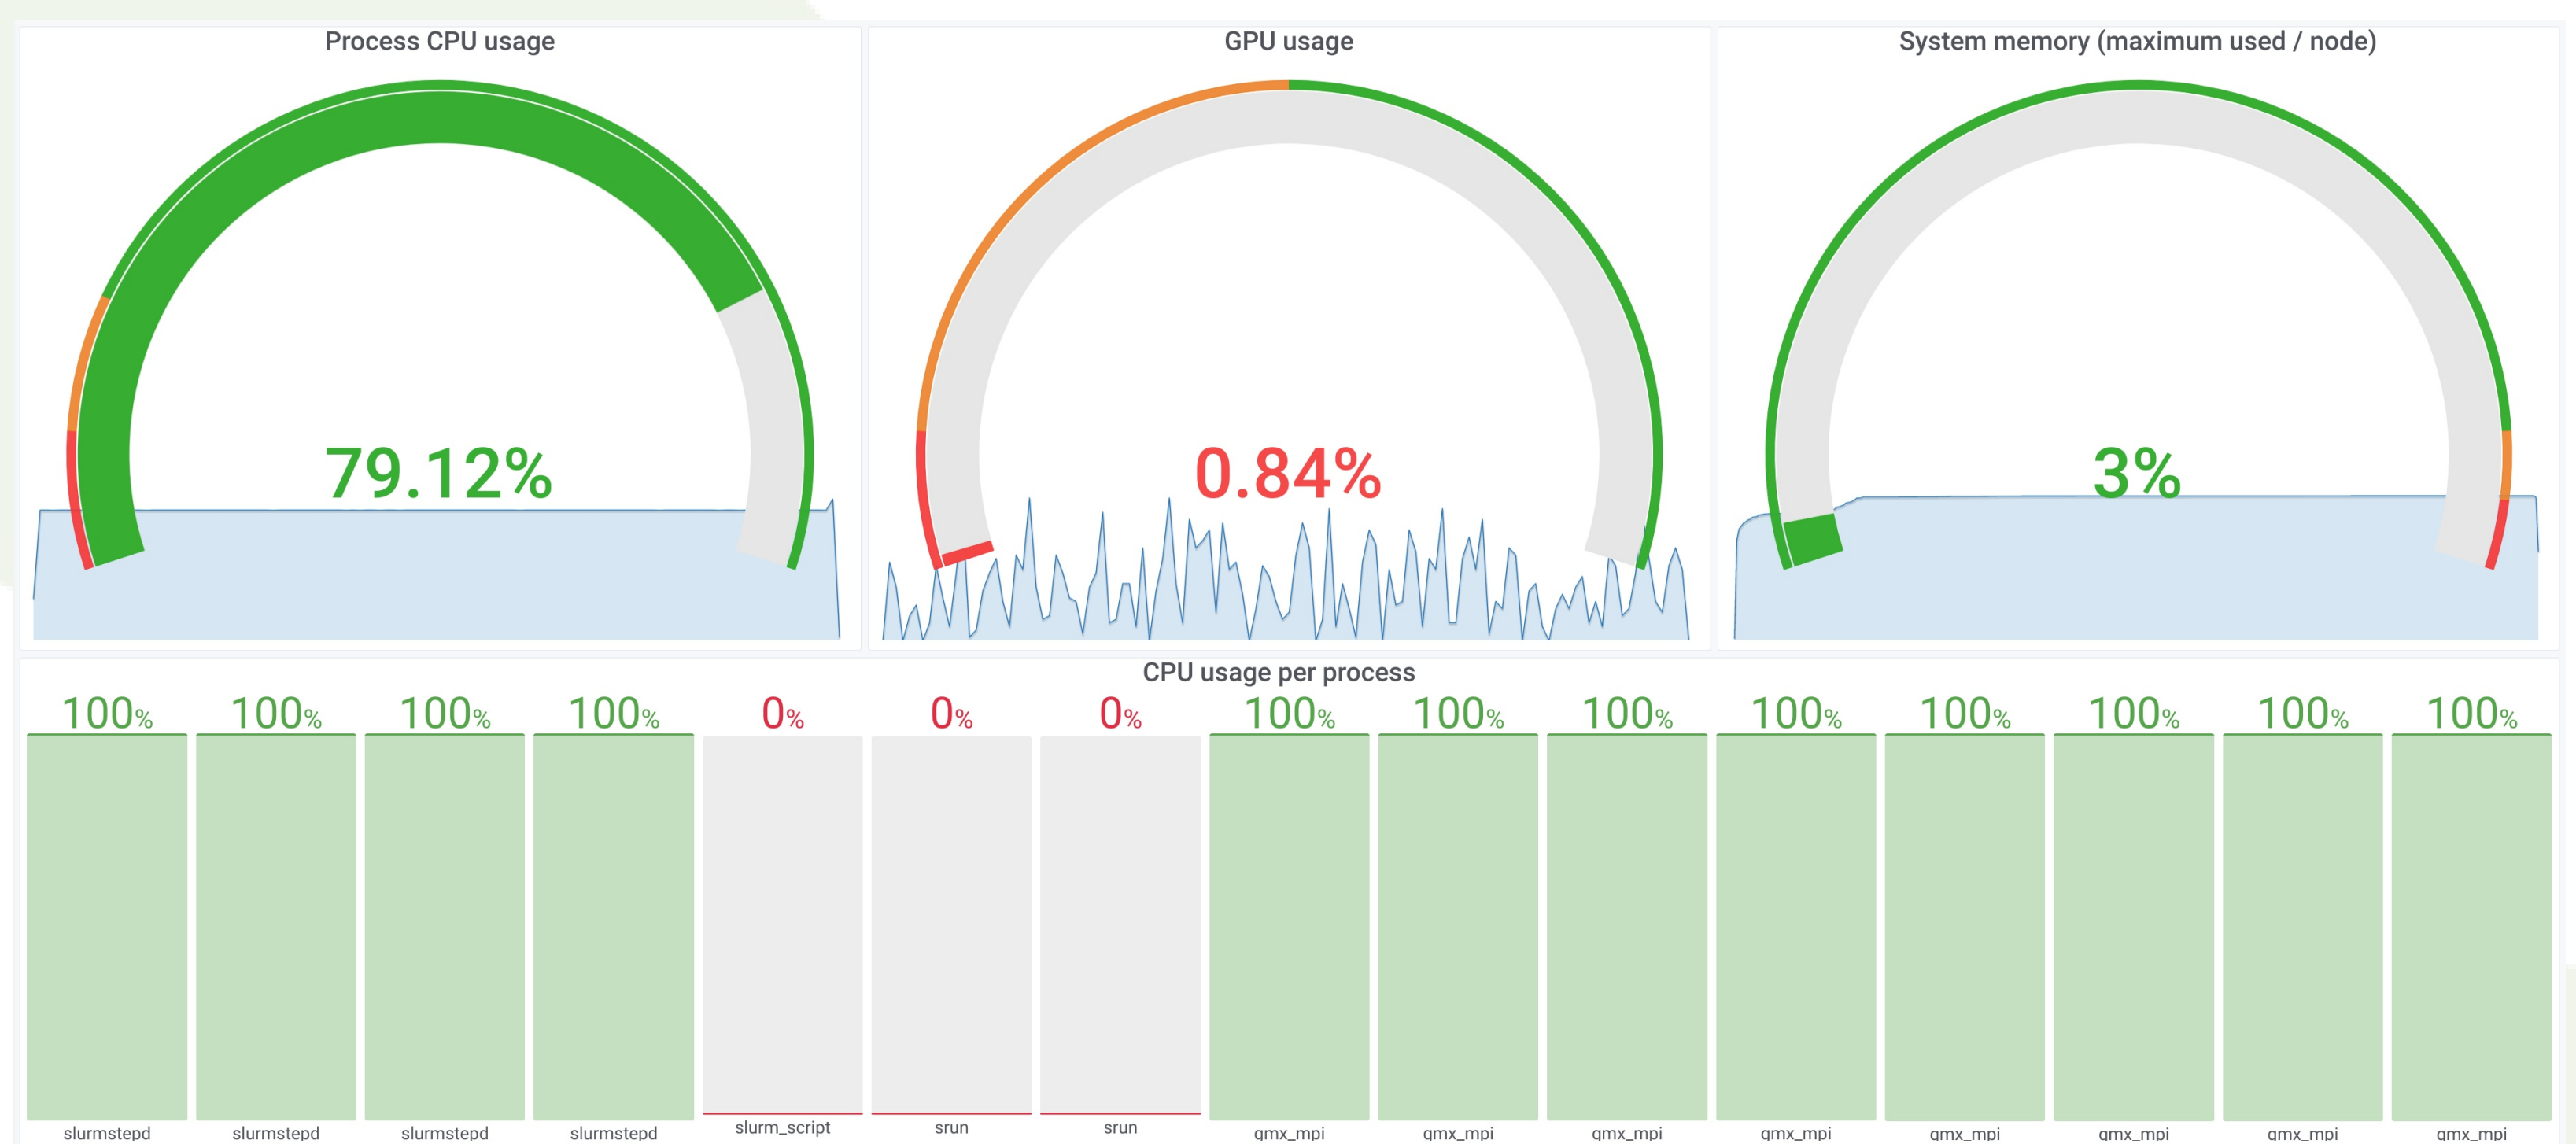


Fig. 3: An example of Grafana dashboards. Charts show low GPU usage and nearly 100% CPU usage of all processes of the job

Achievements

Development of complete monitoring, analysis and reporting toolkit including metric collector, job analysis, recommendation system and report generation

Automatic collection of metrics from various sources process and node based (for shared and non-shared nodes)

Automatic job performance analysis detects pathological and low performance cases

Automatic report generation with recommendations offering different levels of details for review of performance and resource consumption of a job and formulation of recommendations for improving code or job parameters

Tool has been successfully deployed on the targeted systems