ProfiT-HPC: Profiling Toolkit for High Performance Computing

Azat Khuziyakhmetov

Gesellschaft für wissenschaftliche Datenverarbeitung mbH Göttingen

9. HPC-Status-Konferenz der Gauß-Allianz (Paderborn)
17.-18.10.2019

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

Overview

Organization

Goals

Motivation

Comprehensive performance analysis of jobs.
Users get reports with necessary information about the job to evaluate its efficiency and tweak submission parameters or programs if needed.

Goals

- Auto generated reports with following data:
  - chosen metrics and performance indicators
  - recommendations to reach higher efficiency
- Improved resources utilization of HPC
- Raising awareness for performance parameters among users

Architecture

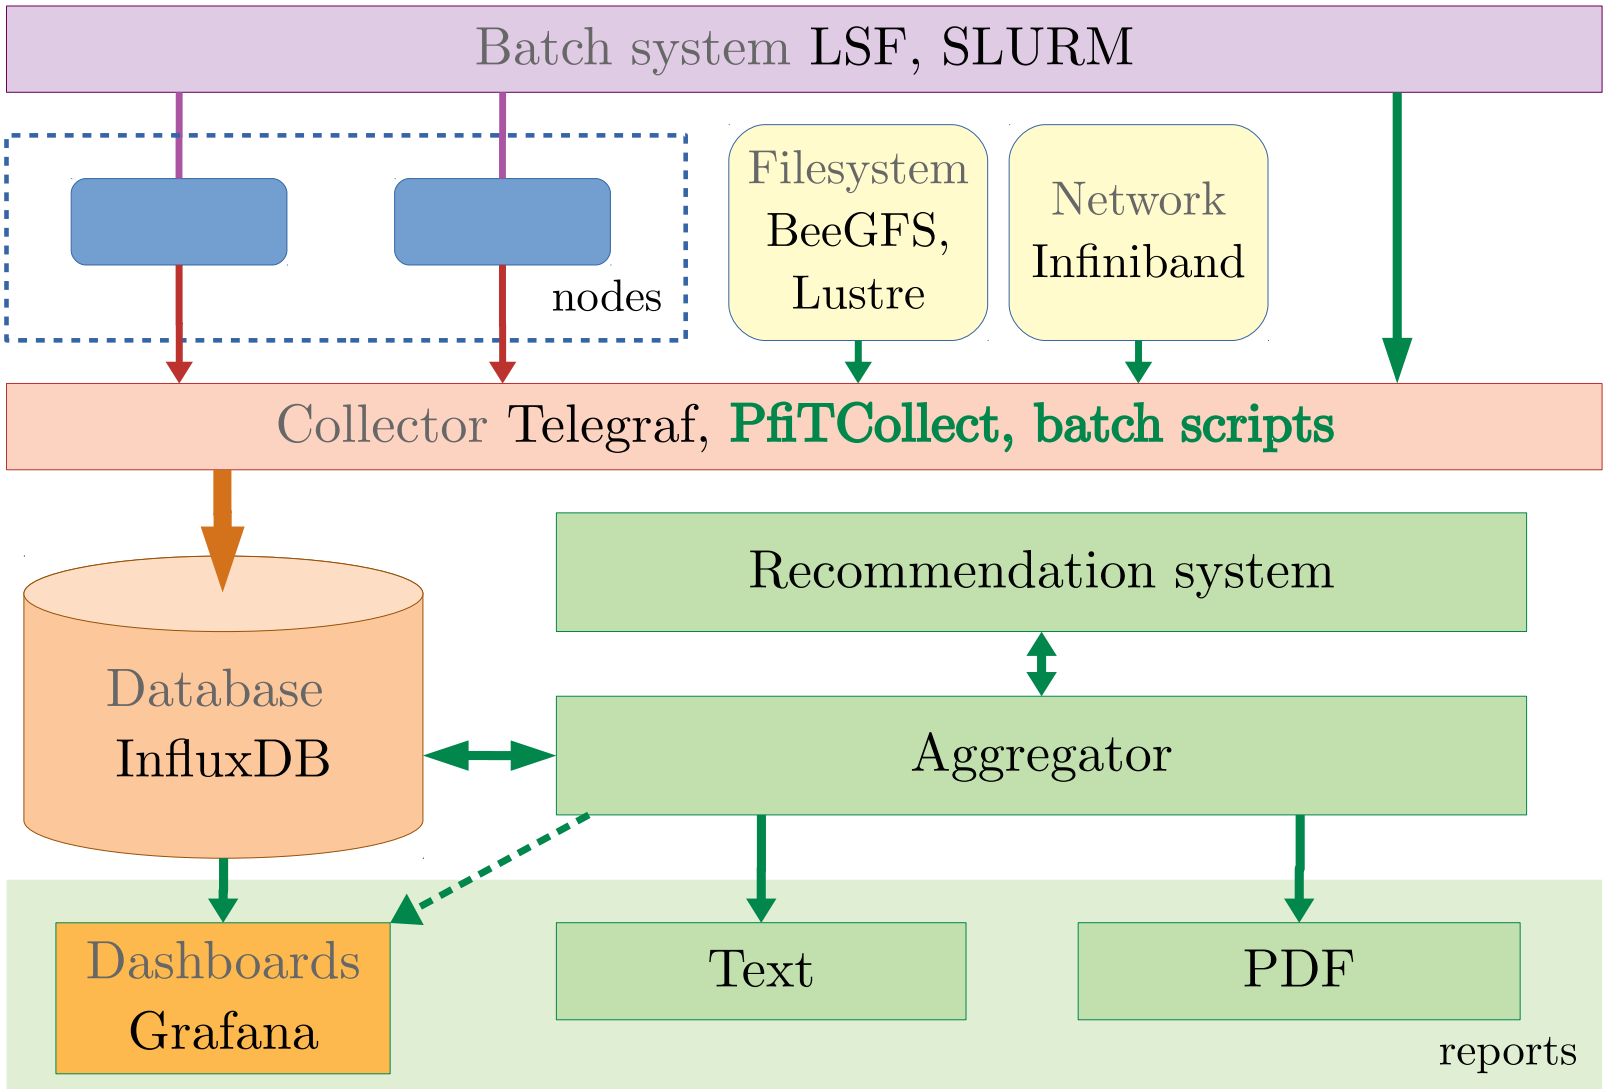

Main focuses

shared env we gather process based metrics which allow to use the toolkit with shared nodes

analysis the recommendation system analyzes gathered metrics to deliver comprehensive performance information

lightweight components of the toolkit like *collectors* and *report generators* are designed to be lightweight

reports multiple report types (*text*, *pdf*, *dashboards*) ensure that users get convenient depth of information based on their requirements

Current state

- Installed on the HPC systems hosted by project partners
- Reports are available for admins and consultants
- Recommendations are being tested and implemented
- Released docs and sources (including **recommendations**)
- Compact and comprehensive *text* and *pdf* reports

Metrics

Nodes

- CPU
- Memory
- System

Process

- CPU times
- Memory
- IO

Filesystem

- Lustre
- BeeGFS
- Local

Others

- Job
- GPU
- Network (Infiniband)

PfiTCollect

Installing Telegraf was refused by some HPC administrators because of security and resource usage reasons. Alternative?!

Key characteristics

- A metric collector written in C
- A lightweight alternative to Telegraf ...
- ... with strongly reduced resource usage and no root access
- Collects metrics every predefined time step
- Stores metrics in InfluxDB database
- Runs stable in pre production

Database (InfluxDB)

Works with high load

- 1 month retention policy
- 10 seconds interval
- 400 GB of data
- 500 nodes
- 25 million series

Another high load scenario

- 1 week retention policy
- 1 minute interval
- 70 GB of data
- 2000 nodes
- 7 million series

Recommendation system

Definition

Recommendations are suggestions for users to help with improving efficiency of their jobs

Recommendation system is designed to be completely separated from other parts of the toolkit. Only input data should be provided.

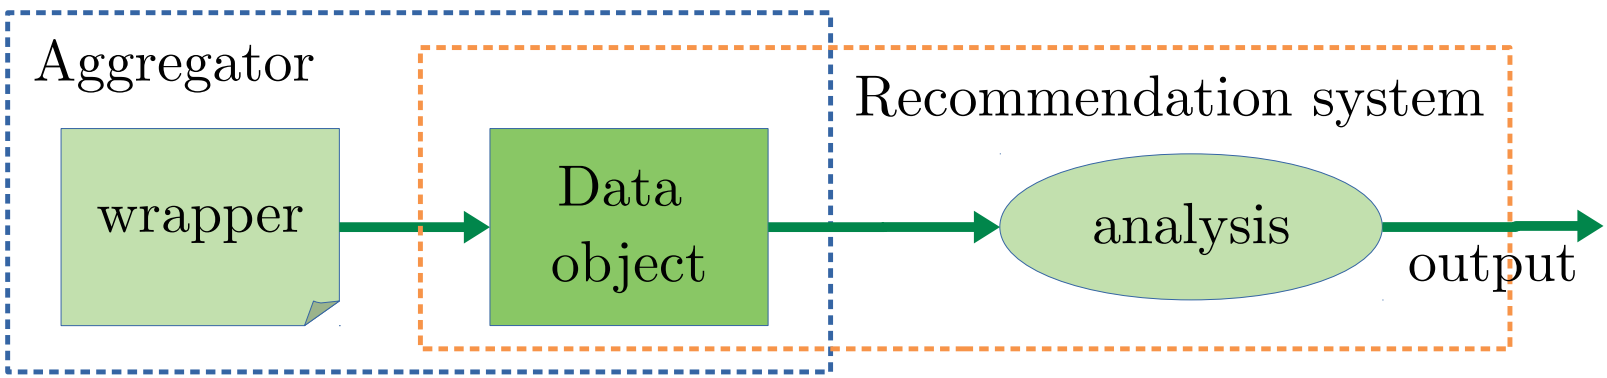

Rules and Attributes

Attribute is a property of the job with predefined values

$$A \in \{V_1, V_2, ..., V_n\}$$

For example, *requested_walltime* attribute might be `HIGH` or `NORM`

Rule is a set of attributes with specified values

$$R : A_1 = V_1 \ \& \ A_2 = V_2 \ \& \ ... \ \& \ A_k = V_k$$

For example, *fix_requested_walltime* rule might be formed by attributes like *requested_walltime* = `HIGH` & *cpu_usage* = `NORM`

Decision trees

The set of the rules can be represented and visualized as a decision tree

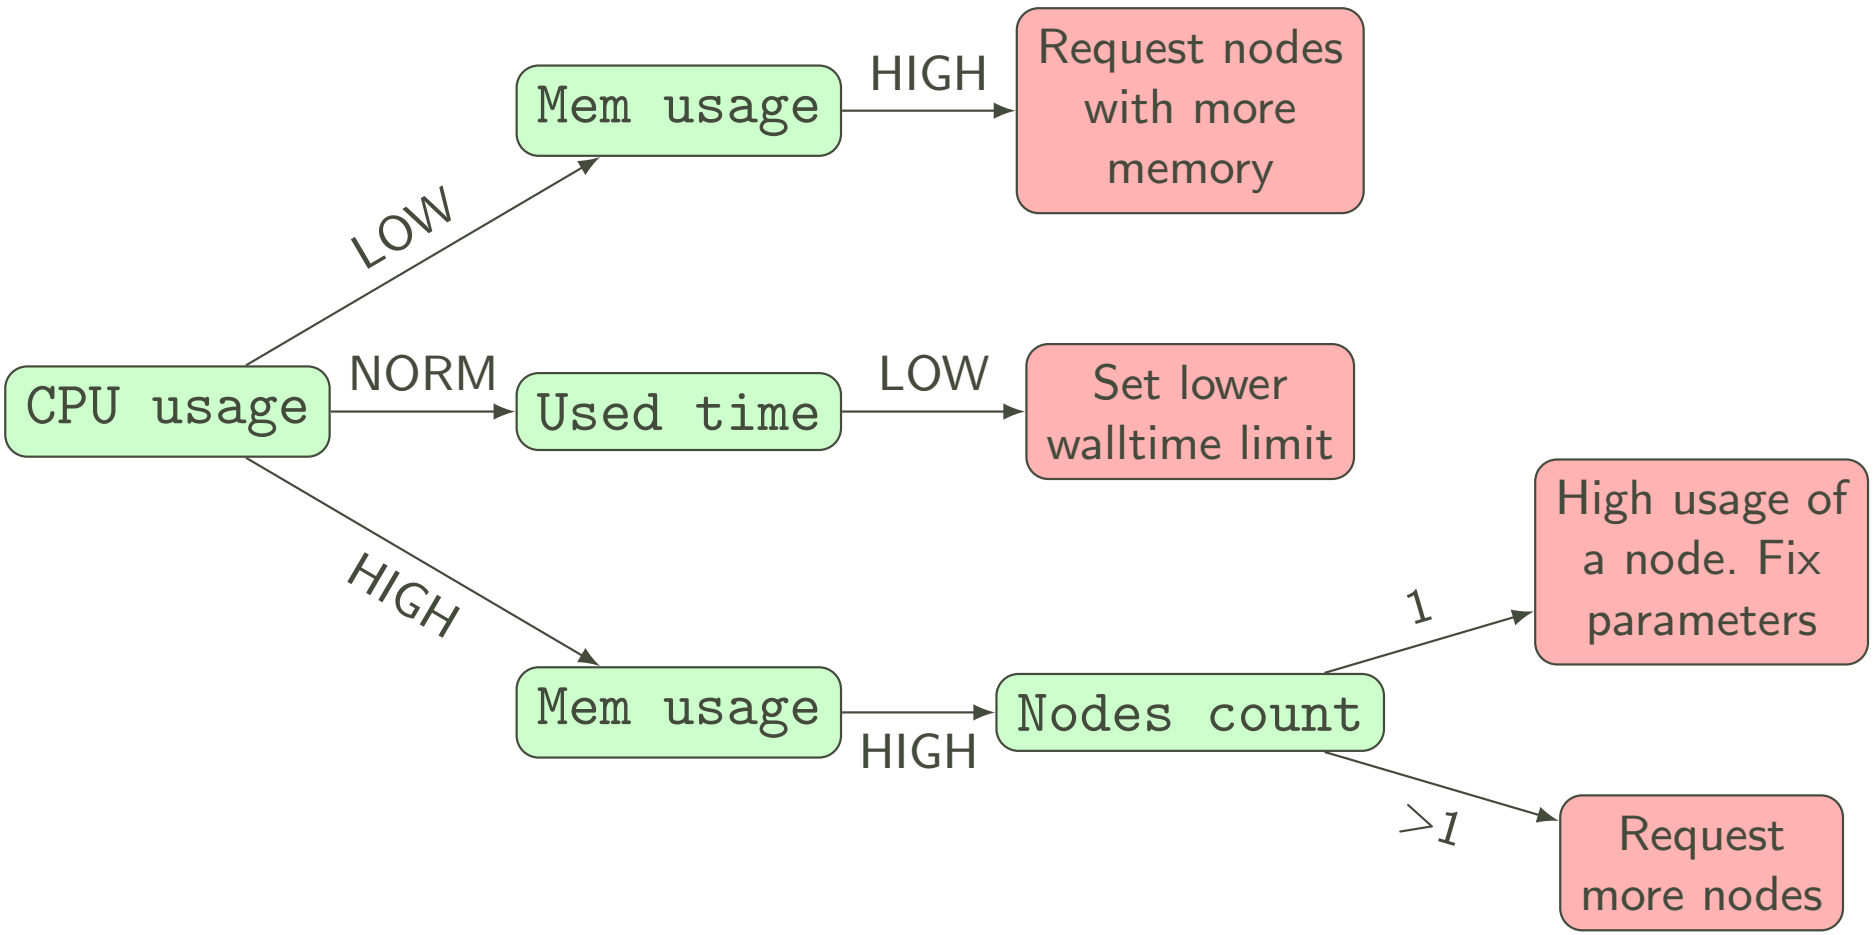

text report

The main purpose of the text report is to be sent by email or be outputted right into user's shell.

Key characteristics

- fixed width
- compact formatting
- contains only summary + recommendations

pdf report

The main purpose of the pdf report is to give the user a grafical overview of the requested and utilized resources.

Key characteristics

- Global overview of job
- Charts with global and node based utilization
- List of recommendations
- Detailed time series plots
- Easier to print and share than dashboards

Links

`profit-hpc.de`
: The main page where you can find all these links

`gitlab.gwdg.de/profit-hpc`
: Aggregator module along with the recommendation system

`profit-hpc.de/wp-content/uploads/2019/10/userguide.pdf`
: User's documentation

`profit-hpc.de/wp-content/uploads/2019/10/adminguide.pdf`
: Administrator's documentation

Conclusion

- Installation
  - Robust database server is the key component
  - Python makes installation much easier
- Workload
  - Toolkit works with high load on $\sim$ 2000 of nodes
- Reports
  - The data in all reports should come from same calculations
  - Compact and comprehensive notation is important
- Recommendations
  - Should be transparent how they are calculated
  - Important to cover many possible cases

Thank you!

Questions?

https://profit-hpc.de info@profit-hpc.de