

universität freiburg

Plans and Status in Freiburg

Topic Area-I:
Development of tools for integration of
heterogeneous resources

FIDIUM Collaboration Meeting 2024

September 30th 2024

Michael Böhler, Anton J. Gamel, Stefan Kroboth,
Benjamin Rottler, Dirk Sammel, Raghuvar Vijayakumar,
Markus Schumacher



Bundesministerium
für Bildung
und Forschung

FIDIUM

Physikalisches Institut

Albert-Ludwigs-
Universität Freiburg



Planned contributions

- ▶ TA-I: Development of tools for integration of heterogeneous resources
 - ▶ WP-1: Exploration and efficient integration of opportunistic resources
 - ▶ WP-2: Accounting and Controlling of heterogeneous resources
- ▶ TA-III: Adopting, testing and optimization of production and analysis environments
 - ▶ WP-1: Integration, tests, optimization and deployment of services
 - ▶ WP-3: Support
- ▶ 2 FTEs funded: filled with post-docs from HEP/Computing community
 - ▶ ~ 1 FTE in TA-I

TA-I: Development of tools for integration of heterogeneous resources

WP-1: Exploration and efficient integration of opportunistic resources

MS1 Development of start-stop mechanism ✓

MS2 Development of plugins (interfaces) for the new accounting system ✓

MS3 Development and integration of a container solution ✓



WP-2: Accounting and Controlling of heterogeneous resources

MS1 Design, conception and technology selection ✓

MS2 Implementation of a accounting prototyp ✓

MS3 Further developments of the interfaces and quality assurance ✓


MS4 Consolidation and adjustments in production ✓



Achieved Milestones - WP-1

MS1 Automatic start-stop mechanism

- ▶ COBaID always keeps one drone running
- ▶ some use cases might have empty queues



- ▶ prevents COBaID to start drones
 - ▶ if there are no pending jobs in the OBS
- ▶ available in COBaID: since [v.0.12.1](#)

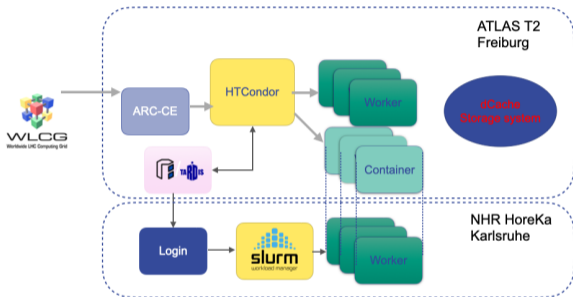
MS2 Interface for accounting system (AUDITOR Plugin in TARDIS)



- ▶ AUDITOR retrieves drone info from TARDIS
- ▶ dedicated AUDITOR plugin in TARDIS project
- ▶ change of drone state triggers action
- ▶ available in TARDIS: since [v.0.7.0](#)



Achieved Milestones - WP-1

MS3 Development and integration of a container solution



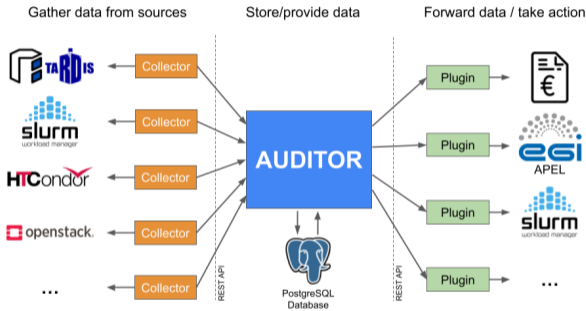
- ▶ Setup dedicated Panda Queue (UNI-FREIBURG_NHR) with
 - ▶ ARC-CE
 - ▶ HTCondor
- ▶ connected to resource of NHR Cluster HoreKa for testing in May

▶ Performed HammerCloud tests and processed



- ▶ Numbers in UserAnalysis and Testing dominated by HC tests (low CPU eff. expected)

Achieved Milestones - WP-2: The Auditor Accounting Ecosystem



Modular accounting ecosystem

► Collectors

- Accumulate data

► Core component

- Accept data
- Store data
- Provide data

► Plugins

- Take action based on stored data

Documentation and code

<https://github.com/ALU-Schumacher/AUDITOR>

AUDITOR: AccoUning Data handling Toolbox for Opportunistic Resources

Presented at 26th International Conference on Computing in High Energy & Nuclear Physics (CHEP)
May 8-12, 2023, Norfolk, VA, USA: [AUDITOR: Accounting for opportunistic resources](#)

Available Collectors and plugins

▶ TARDIS Collector

- ▶ Collect drone information

▶ SLURM Collectors (2 types)


- ▶ Collect information about SLURM jobs via SLURM CLI commands

▶ HTCondor Collector (dev @ KIT)

- ▶ Equivalent of SLURM collector for HTCondor

▶ Kubernetes Collector (dev @ Wup)

- ▶ accounts resources from kubernetes clusters
- ▶ will be available with auditor v0.7.0



▶ Priority plugin

- ▶ trigger command based on delivered resources (e.g. adjust group priority)

▶ APEL accounting plugin

- ▶ Reports accounting data to the WLCG accounting service (APEL) - also sub clusters

▶ Utilization report (FIDIUM ext)

- ▶ Analyse requested vs consumed resources of a user

Numbers from AUDITOR git repo

▶ Extensive documentation

Auditor

- Features

Running Auditor

- Migrating the database
- Using Docker
- Configuration files
- Metrics exporter for Prometheus

Compiling Auditor

Packages

Collectors

- SLURM Collector
- SLURM Epilog Collector
- HTCondor Collector

Plugins

- APEL Plugin
- Priority Plugin

Auditor Clients

License

- Contribution

Overview

Auditor

Auditor stands for Accounting Data Handling Toolbox For Opportunistic Resources. Auditor ingests accounting data provided by so-called *collectors*, stores it and provides it to the outside to so-called *plugins*.

It comes with a well-defined REST API which allows for the implementation of application-specific collectors and plugins. This makes it well suited for a wide range of use cases.



Overview of the AUDITOR ecosystem. AUDITOR accepts records from collectors, stores them in a PostgreSQL database and offers these records to plugins which take an action based on the records.

- ▶ 8 contributors
- ▶ from 3 universities
 - ▶ Freiburg (main development), KIT, Uni Wuppertal
- ▶ 15 releases - latest **v0.6.2**
- ▶ Continuous improvements: **Commits**



- ▶ Also on <https://zenodo.org/records/13239266>

The screenshot shows the Zenodo record page for the accounting ecosystem AUDITOR. The page includes a search bar, navigation links for 'Communities' and 'My dashboard', and a 'Log in' / 'Sign up' button. The record is published on July 29, 2024, and is version v0.6.2. It has 71 views and 2 downloads. The title is 'The accounting ecosystem AUDITOR'. Below the title, there is a list of contributors: Boehler, Michael; von Cube, Florian; Fischer, Max; Giffels, Manuel; Kleinhans, Raphael; Krobath, Stefan; Rottler, Benjamin; Sammel, Detl; Schreyer, Matthias; Vijayakumar, Raghav. There is a 'Show affiliations' button. The 'Versions' section shows a table with columns for Version, DOI, and Date. The table lists two versions: v0.6.2 (DOI: 10.5281/zenodo.13239266, Date: Jul 29, 2024) and v0.5.0 (Date: Jul 4, 2024).

► AUDITOR paper submitted to Computing and Software for Big Science

Research Article

Auditor: Accounting Data Handling Toolbox for Opportunistic Resources

Michael Boehler, Florian von Cube, Max Fischer, Manuel Giffles, and 7 more

This is a preprint; it has not been peer reviewed by a journal.

<https://doi.org/10.21203/rs.3.rs-4741479/v1>
This work is licensed under a CC BY 4.0 License

Abstract

Increasing computing demands and concerns about energy efficiency in high-performance and highthroughput computing are driving forces in the search for more efficient ways to use available resources. Sharing resources of an underutilised cluster with a high workload cluster increases the efficiency of the underutilised cluster. The software COBaID/TARDIS can dynamically and transparently integrate and disintegrate such resources.

However, sharing resources also requires accounting. AUDITOR (Accounting Data Handling Toolbox for Opportunistic Resources) is a modular accounting ecosystem that is able to cover a wide range of use cases and infrastructures. Accounting data are gathered via so-called collectors, which are designed to monitor batch systems, COBaID/TARDIS, cloud schedulers or other sources of information. The data is stored in a database, and access to the data is handled by the core component of AUDITOR, which provides a REST API along both Rust and a Python client libraries. So-called plugins can take actions based on accounting records.

Depending on the use case, one simply selects a suitable collector and plugin from a growing ecosystem of collectors and plugins. To facilitate the development of collectors and plugins for yet uncovered use cases by the community, libraries for interacting with AUDITOR are provided.

accounting heterogeneous resources high-performance computing high-throughput computing

Status: **Under Review**

Springer
Computing and Software for Big Science

Version 1
posted 19 Aug, 2024

👤	Reviewers agreed at journal	02 Sep, 2024
✉	Reviewers invited by journal	02 Sep, 2024
👤	Editor assigned by journal	23 Jul, 2024
🕒	Submission checks completed at journal	16 Jul, 2024
📅	First submitted to journal	15 Jul, 2024

You are reading this latest preprint version

📄 Citations **See more**

👁 Engagement **11 views**

💬 Comments **0**


Schema adjustments - DB Performance improvements

MS4 Consolidation and adjustments in production

- ▶ reading and writing tests identified inefficiencies in the database structure

AUDITOR v0.5.0

AUDITOR v0.6.2



- ▶ adjustments to the database schema fixed these performance losses

Query Type	10k		50k		100k		1 mil	
	0.5.0	0.6.2	0.5.0	0.6.2	0.5.0	0.6.2	0.5.0	0.6.2
All records	267.7	4.8	1007.3	25.0	2929.9	42.2	791457.6	421.1
Time range	100.5	0.6	686.6	2.5	1814.3	24.4	41948.1	188.7
Time range and site_id	88.8	1.0	814.6	4.1	2500.0	25.2	70142.9	232.6
Component	99.7	7.8	564.5	42.6	2366.5	31.1	41905.9	1648.8
Time range + meta + comp	77.0	0.5	499.6	3.2	1159.0	25.0	29114.5	496.0
Time range + 2 meta + comp	80.5	0.6	567.6	2.5	2285.6	25.7	23803.7	187.3
One record	92.3	0.1	1136.4	0.1	1164.6	0.1	34045.7	1.0

all numbers in milliseconds

WP-1: Exploration and efficient integration of opportunistic resources

MS1 Development of an Utilization Report Plugin (Q3-2025)

- ▶ Analyse requested vs consumed resources of a local user
- ▶ Send a weekly report with possible savings and CO₂ footprint



WP-2: Accounting and Controlling of heterogeneous resources

MS1 Development of a stress test collector/plugin (Q4-2024)

- ▶ Allows for stress testing and performance measurements
- ▶ Collector for testing uploads to AUDITOR
- ▶ Plugin for testing data access at AUDITOR

MS2 Further development and optimization of the auditor db (Q3-2025)

- ▶ Based on the results of the stress test analysis





Conclusion & Outlook

Conclusion

- ▶ all milestones achieved
- ▶ AUDITOR v0.6.2 is a stable accounting ecosystem ready for production
 - ▶ smaller open issues and user requests → work never ends...

Outlook

- ▶ still exciting milestones for FIDIUM extension period ahead
 - ▶ Utilization Report Plugin
 - ▶ AUDITOR Stress test Collector and Plugin



References



Michael Boehler
University of Freiburg
Telefon +49 761 203 8409
michael.boehler@physik.uni-freiburg.de

Website: <https://alu-schumacher.github.io/AUDITOR/>
GitHub: <https://github.com/ALU-Schumacher/AUDITOR>
FIDIUM: <https://fidium.erumdatahub.de>
Email: auditor@physik.uni-freiburg.de