VSB TECHNICAL UNIVERSITY OF OSTRAVA | IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER

SEANERGYS
ENERGY EFFICIENT EXASCALE

SEANERGYS: From Day-0 Datasets to Predictive AI for HPC Resource Management

Ing. Petr Ptáček



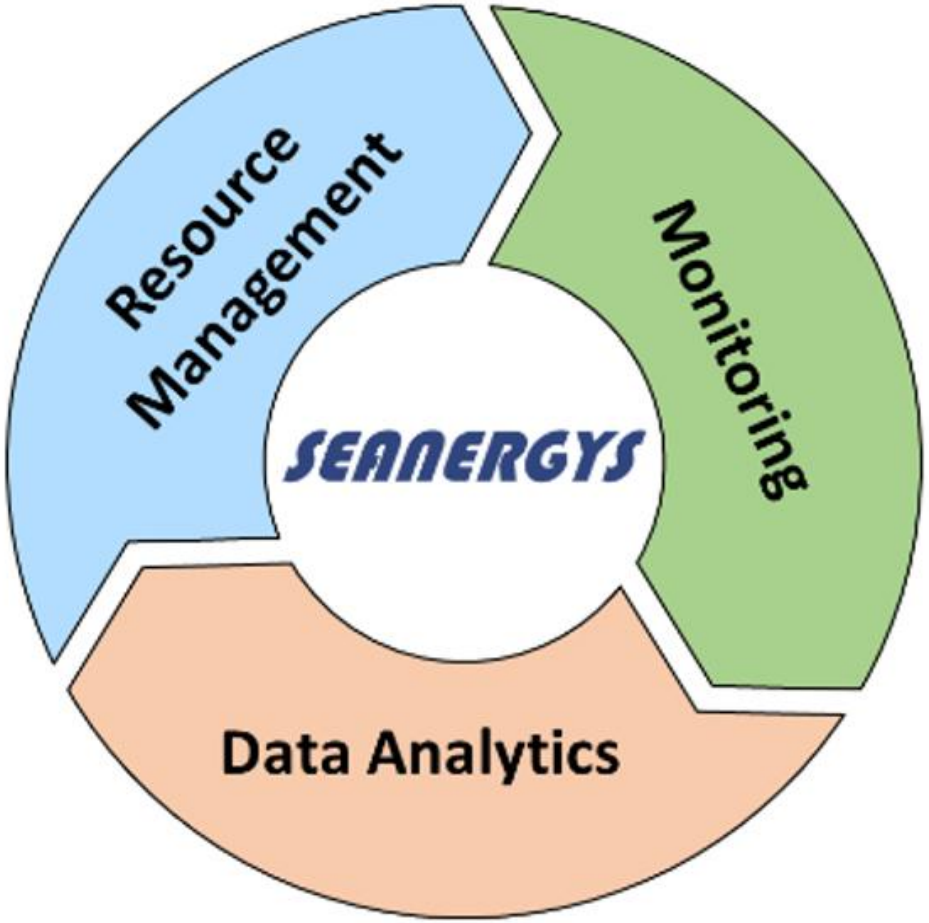
SEANERGYS
Energy Efficient Exascale

EuroHPC
Joint Undertaking

VSB TECHNICAL UNIVERSITY OF OSTRAVA | IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER

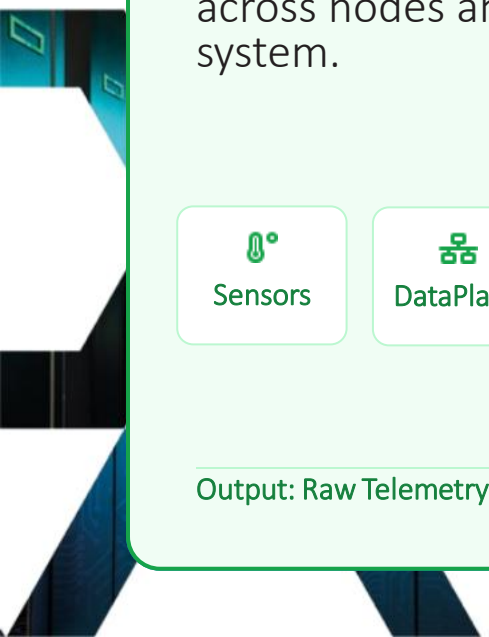
- *„SEANERGYS creates an integrated European software solution that optimises resource utilisation and reduces the energy used for real-world workload mixes. It therefore improves the throughput of HPC systems, generating more R&D results for a given energy budget.“*





SEANERGYS
ENERGY EFFICIENT EXASCALE

VSB TECHNICAL UNIVERSITY OF OSTRAVA | IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER



PILLAR 1

CMI

Comprehensive Monitoring Infrastructure: Collects real-time operational metrics across nodes and whole system.

- Sensors
- DataPlane
- Operation DB

Output: Raw Telemetry Streams

PILLAR 2

AIDAS

Artificial Intelligence Data Analytics System: Transforms data into predictions using the AI production tool pooling AI models from Model Zoo.

- Model ZOO
- ODA Tools
- 3D Digital Twin of Cluster

Output: Actionable Insights

PILLAR 3

DSRM

Dynamic Scheduling and Resource Management System: Enforces energy constraints and optimal job placement dynamically.

- Scheduler
- Efficiency
- Malleable & Moldable Jobs

Output: Dynamic Execution

Day-0 Datasets

7 Supercomputing Centers

Collected anonymized monitoring data from 7 HPC sites (IT4I, LXP, BSC, LRZ, TUD, CINECA, JSC).

Unified Schema Standardization

Normalized telemetry from diverse sources (time-series and tabular data) into a unified format.



Compressed Apache Parquet

Stored in a S3 repository as compressed Parquet files. ~98% size reduction thanks to compression.

Cold-Start Mitigation

Enables AIDAS model training, feature engineering, and validation *before* live monitoring streams from CMI are active.

Day-0 Datasets Data Composition

SEANERGYS
Energy Efficient Exascale

VSB TECHNICAL UNIVERSITY OF OSTRAVA | IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER

Time-Series Data

- Node power consumption & temperatures
- CPU/GPU/Disk utilization metrics
- Network bandwidth & packet rates
- Scheduler event logs & queue states

Example: Node power in Watts

Timestamp	Value	Node ID
2022-04-01 10:00:00	405	cn001
2022-04-01 10:01:00	402	cn001
2022-04-01 10:02:00	395	cn001

Tabular Data

- SLURM job records
- Application trace signatures & module loads
- Cluster metadata (topology, queue limits)

Example: SLURM job metadata

Job ID	User ID	Alloc Nodes	Start Time	End Time	State
132	user_00000005	[cn001]	2024-04-01 02:15:40	2024-04-01 03:05:36	COMPLETED
133	user_00000020	[acn05, acn06]	2024-04-01 05:51:28	2024-04-01 05:59:15	FAILED
144	user_00001011	[cn344, cn345]	2024-04-01 13:23:47	2024-04-01 21:23:47	TIMEOUT

Day-0 Datasets Format 1: Time-Series

VSB TECHNICAL UNIVERSITY OF OSTRAVA | IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER

SEANERGYS
ENERGY EFFICIENT EXASCALE

3-LEVEL HIVE PARTITIONING

Time-Series Format

Optimized columnar format for fine-grained telemetry & job metadata

Nested Hierarchy:
<system-name>/year_month=YY-MM/plugin=<p>/metric=<m>/a_*.parquet

Optimal Compression: Apache Parquet compressed with **Zstandard (ZSTD Level 9)**, dictionary encoding, and 100k row group chunks.

Metadata Package: Includes dataset metadata, queue specifications, and node geo-locations (cluster, row, rack, chassis, blade).

Initial inspiration: UNIBO paper "M100 ExaData: a data collection campaign on the CINECA's Marconi100 Tier-0 supercomputer"

Directory Structure for **IT4I** Time-Series Day-0 Dataset stored in S3 storage

```
karolina/
├── metadata_00_general.json
├── metadata_01_queue_description.csv
├── metadata_02_cluster_description.csv
├── year_month=25-03/
│   ├── plugin=ipmi/
│   │   ├── metric=CPU_0_Temp/
│   │   │   └── a_0.parquet
│   │   └── metric=CPU_1_Temp/
│   │       └── a_0.parquet
│   ├── plugin=slurm/
│   │   └── metric=job_info/
│   │       └── a_0.parquet
│   ├── plugin=nvml/
│   │   ├── metric=GPU_0_AVERAGE_POWER/
│   │   │   └── a_0.parquet
│   │   ├── metric=GPU_0_AVERAGE_POWER/
│   │   │   └── a_0.parquet
│   │   └── metric=...
│   └── plugin=...
└── year_month=...
```

Metadata files

1. Partitioning: months

2. Partitioning: plugins (group of metrics sharing same source)

3. Partitioning: specific metric

The anonymized data are stored in Parquet files

- 6 Months of data
- 720 compute nodes
- 21 plugins
- 1954 metrics
- ~660 GB [compressed]
- ~8 TB [raw]

Day-0 Datasets Format 2: Aggregated

VSB TECHNICAL UNIVERSITY OF OSTRAVA | IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER

SEANERGYS
ENERGY EFFICIENT EXASCALE

3-LEVEL HIVE PARTITIONING

Aggregated Format

Pre-computed statistics from Time-Series format computed over fixed temporal windows

Path Hierarchy:
<system-name>/<split>/<scope>/<format>/<file_name>.<ext>

Supported Splits: Fixed time intervals (hourly, daily, weekly, monthly, quarterly, yearly, global).

Aggregation Scopes:

- by-job
- by-node

Lexical File Naming: Uses split-specific prefixes (e.g., YY_MM_a0.<ext> for monthly split) for automated multi-part loading.

Directory Structure for public **Fugaku** Job-Centric Day-0 Dataset stored in S3 storage

```
fugaku/
└── monthly/
    └── by-job/
        └── parquet/
            ├── 21_03_a0.parquet
            ├── 21_04_a0.parquet
            ├── 21_05_a0.parquet
            ├── 21_06_a0.parquet
            ├── ...
            └── 24_04_a0.parquet
```

1. Partitioning: by split (months)

2. Partitioning: by scope (by job)

3. Partitioning: by data format (parquet)

The anonymized data stored in Parquet files

- 38 Months of data
- 24 million jobs
- 45 features per job
- ~26 GB [compressed]

IT4I Anonymization Methodology for Day-0 Datasets

VSB TECHNICAL UNIVERSITY OF OSTRAVA | IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER

SEANERGYS
ENERGY EFFICIENT EXASCALE

Job IDs & Timestamps

Global Job ID Offset: Shifted by random integer offset

Global Time Shift: Shifted by random time offset

Job ID: 18946 → 2428729
Job Start Time: 2025-03-24 → 2023-01-01

Temporal Preservation

Identities & Queues

User ID: Replaced with *user_[8-char]*

Project ID: Replaced with *project_[6-char]* (stricter for special projects)

Queue / QoS: Replaced with *queue_[4-char]*, preserves walltimes & node limits

User ID: it4i-franta → user_00000001
Project ID: open-32-19 → project_002002
Queue ID: qcpu → queue_0004

Identity Protection

Nodes & Topology

Node Ids & Geolocation: Preserved for 3D digital twin visualization (thermal hotspot detection)

Topology Retention

File Paths & Command lines

Left-Part Prefixes: Retained (**/mnt, /scratch , /home**).

Middle & Right Parts: Replaced with *path_[8-char]* & *file_[8-char]*.

Command Lines: Commands anonymized like paths; arguments replaced with placeholders.

File Path:
/mnt/proj3/usr/main.cpp →
/mnt/path_00000002/file_00000002

Command Line:
/home/franta/my_exe_file arg1 arg2 →
/home/path_00010020/file_00001200 arg arg

Storage & Code Safety

AIDAS Model Zoo

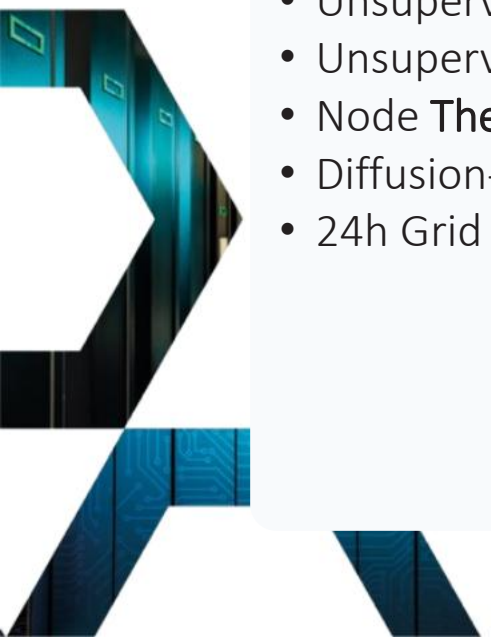
SEANERGYS
Energy Efficient Exascale

VSB TECHNICAL UNIVERSITY OF OSTRAVA | IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER

System-Centric Models

Focuses on overal cluster health, telemetry quality, localized thermal hotspots, and energy grid context

- Unsupervised node **Anomaly Detection**
- Unsupervised node **Anomaly Prediction**
- Node **Thermal Anomaly Prediction**
- Diffusion-based **Telemetry Repair** and **Fill**
- 24h Grid **Carbon Intensity Forecasting**



Job-Centric Models

Focuses on workload characterization and job metric predictions prior to and during job execution. Provides also post-execution feedbacks

- Job **Mean power** & **Energy** prediction
- **CPU vs Memory-Bound** profile classification
- **Zombie Job Detection**
- Job **Scalability** & **Co-Location** prediction
- Job **Misconfiguration Analysis** & **User Feedback**
- **Walltime** & **Optimal Resources** recommendation

Conclusion & Future Works

Day-0 Datasets Takeaways

Unified formats for datasets from **7 European HPC sites**.

Enabled pre-training and **validation** of AI models prior to live stream deployment.

Protected user/project privacy while **preserving spatial context** for 3D Digital Twins.

Transition to CMI Live Data

Unified Metric Naming Scheme across all centers

On-the-Fly anonymization and Parquet files creation.

Fine-tuning pre-trained AIDAS models for real-time streaming inference.

Acknowledgement: This project has received funding from the European High Performance Computing Joint Undertaking (EuroHPC JU) under Grant Agreement No. 101177590 (SEANERGYS).

EUROHPC JU GA #101177590

VSB TECHNICAL UNIVERSITY OF OSTRAVA | IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER

SEANERGYS
ENERGY EFFICIENT EXASCALE

THANK YOU FOR YOUR ATTENTION

Ing. Petr Ptáček
petr.ptacek@it4i.cz

IT4Innovations National Supercomputing Center
VSB – Technical University of Ostrava
Studentská 6231/1B
708 00 Ostrava-Poruba, Czech Republic
www.it4i.eu

