

GOALS & OBJECTIVES

Increasing energy and power demands are the key challenges for HPC/AI systems to meet the ever-increasing computational requirements of European industry and research – more critical in the Exascale and post-Exascale era.

New compute and memory/storage devices (like processors, accelerators, high-bandwidth memory, etc.) can improve the energy efficiency of applications, and new interconnects will speed up communication and thus performance of multi-node workloads. Hardware alone, however, cannot address inefficiencies caused by sub-optimal assignment of workloads to resources, such as over- or under-provisioning, running applications at scales where their parallel efficiency drops, or mismatch of workload characteristics with specific node architectures. Dynamic workflows and applications with characteristics that change over time pose particular challenges.

SEANERGYS will design and develop a production-quality software suite for energy-efficient operation of European HPC and AI systems by:

- Creating a holistic monitoring infrastructure and common data repository for operational HPC/AI data.
- Developing an advanced AI-based data analytics framework for HPC and AI operational data.
- Implementing a dynamic resource management system to optimise use of resources and energy efficiency in heterogeneous HPC/AI systems and to adapt to dynamic workloads.
- Validating the SEANERGYS software suite in real operational environments and making it available under permissive licenses.

The software stack is planned to be evaluated and deployed on production class HPC and AI systems at participating top tier European supercomputer sites.

IN A NUTSHELL

BUDGET:

- Overall budget € 33 Million,
EuroHPC JU funding € 16.5 Million

PROJECT TERM:

- June 2025 – May 2029

CONSORTIUM:

- Coordinator: Forschungszentrum Jülich
- 16 Partners, 7 European countries



GET CONNECTED

Email: seanergys-pmt@fz-juelich.de

Web:



FOLLOW US



SEANERGYS

ENERGY EFFICIENT EXASCALE

KEY COMPONENTS

Unified Holistic Monitoring (CMI)


- Collects and consolidates data from sensors, nodes, control systems, and applications.
- Supports structured and unstructured telemetry (e.g., resource usage, cooling, scheduler metadata).
- Data flows through a standardized data plane — replacing fragmented tools with a unified source and interface for real-time and historical analysis.

AI-Based Data Analytics (AIDAS)

- Analyzes system behavior at node, job, and system level based on the CMI monitoring data.
- Offers real-time insights and long-term modelling for power, energy, and performance optimization.
- Utilizes a shared “Model Zoo” for adaptive model selection.
- Includes powerful visual tools for administrators, SW developers and system end-users to drive efficient use and operation of HPC and AI supercomputers.

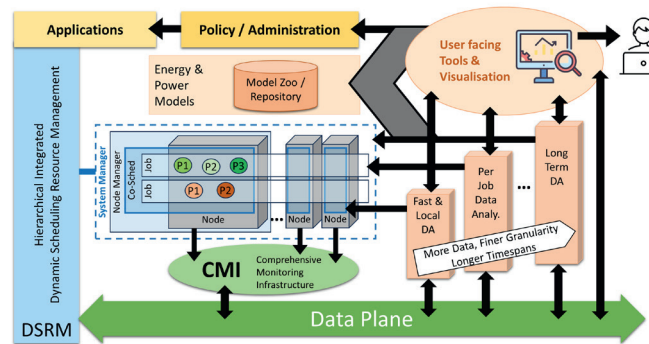
Dynamic Scheduling & Resource Management (DSRM)

- Ingests AIDAS analytics to manage and optimize workload execution across system, node, and job layers.
- Can adapt resource allocation and configuration to dynamic workload behaviour.
- Integrated feedback loop with AIDAS will enable intelligent, self-adjusting operations.



INTEGRATED SOFTWARE STACK

For the three core components to effectively interact, SEANERGYS defines, documents and standardises a series of interfaces and shared data models/formats. They will allow a clear compartmentalisation, smooth integration, streamlined testing and deployment and ultimately efficient execution across many different European HPC and AI sites.



SEANERGYS Software Development Methodology

- The project follows a stringent, quality-oriented SW development method.
- It uses a common development operations platform and a sophisticated CI/CD system to deploy and test on a plethora of HPC and AI systems.
- Advanced Gitlab functionality is used to enable end-to-end tracking of use cases, requirements, SW design elements and code created.

DEPLOYMENT & EVALUATION

Deployment on HPC staging systems

A number of systems are used to test deployments, train operational teams, and gather feedback to improve the software before final production rollout.

Evaluation of the developed solution in production

Post-deployment, the software is evaluated using defined benchmarks and key performance indicators, with lessons learned compiled into a guide to support external HPC deployments and knowledge effective transfer.

Deployment on Production Systems

The final stage deploys the tested software stack on production HPC and AI systems, after a successful, thorough validation by the participating supercomputer centers.

FUNDING

The SEANERGYS project has received funding from the European High Performance Computing Joint Undertaking (JU) under grant agreement No101177590. The JU receives support from the European Union’s Horizon Europe research and innovation programme and from the Czech Republic, France, Germany, Greece, Italy, and, Spain.

Views and opinions expressed are those of the author(s) only and do not necessarily reflect those of the European Union or of the granting authority. Neither the European Union nor the granting authority can be held responsible for them.

