National Competence Center for HPC
Norway

High Performance Computing with Prodrisk

Hydropower Scheduling User Meeting 2025, Oslo

Olaf Trygve Berglihn, Sigma2

2025-05-06

Outline

- Sigma2 - national e-infrastructure in Norway
- A proof-of-concept adapting Prodrisk to High Performance Computing
- A proof-of-concept Software As a Service for Prodrisk - Outlook

National Competence Center for HPC

- Assist research institutes, SMEs, and industry
- *Cost free aid* for Proof of Concepts - TRL 3
- Try-before-you-buy
- Funded by EU: EuroCC2
- Collaboration with SINTEF Industry and NORCE

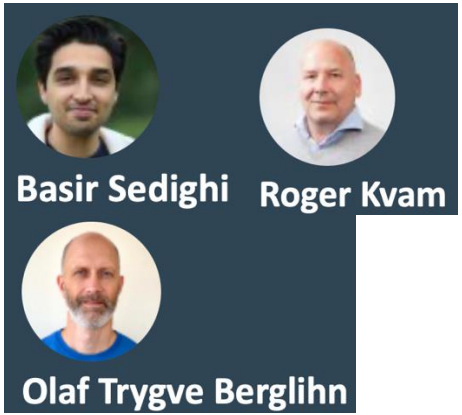

Norwegian research infrastructure services (NRIS)

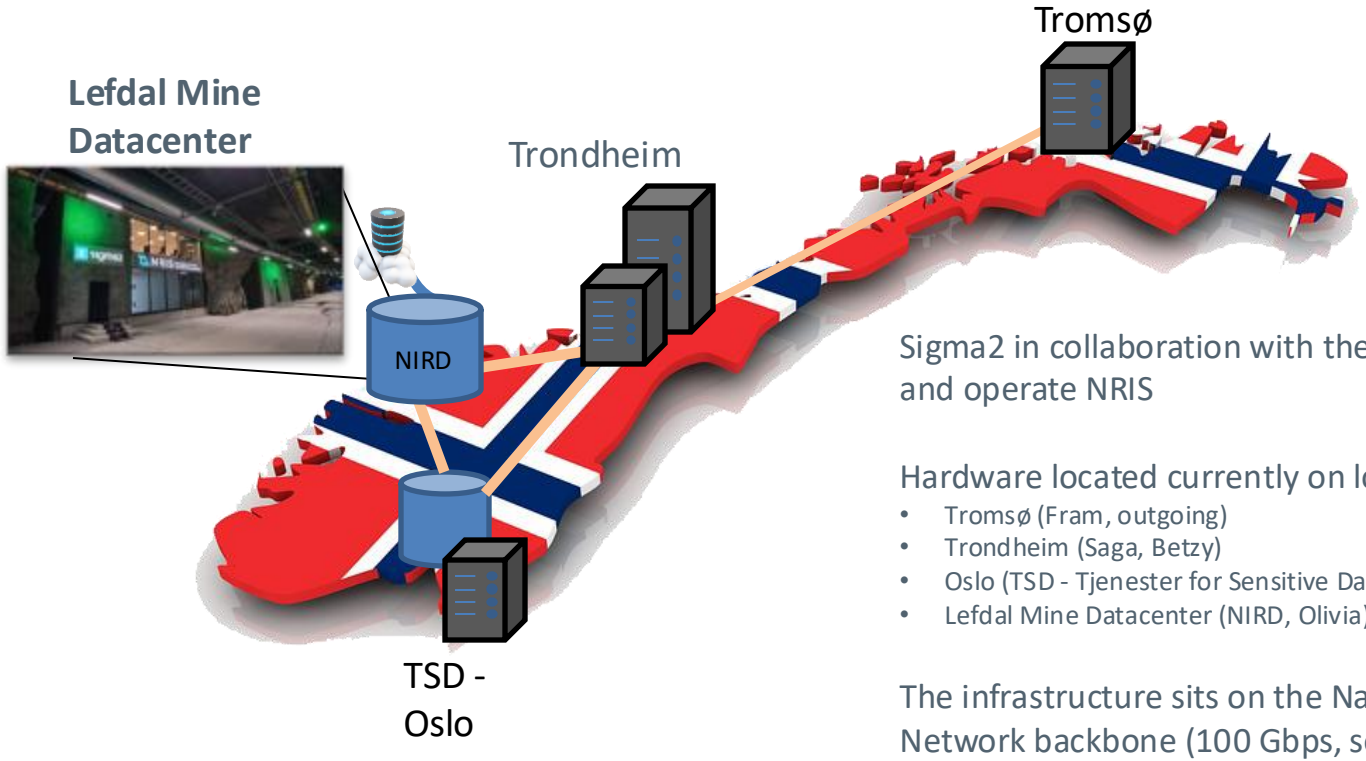

Sigma2 in collaboration with the partner universities provides and operate NRIS

Hardware located currently on locations

- Tromsø (Fram, outgoing)
- Trondheim (Saga, Betzy)
- Oslo (TSD - Tjenester for Sensitive Data / Colossus)
- Lefdal Mine Datacenter (NIRD, Olivia)

The infrastructure sits on the National Research &Education Network backbone (100 Gbps, soon 200 Gbps)

Our data centre at Lefdal Mine Datacenter

Green computing power for the future

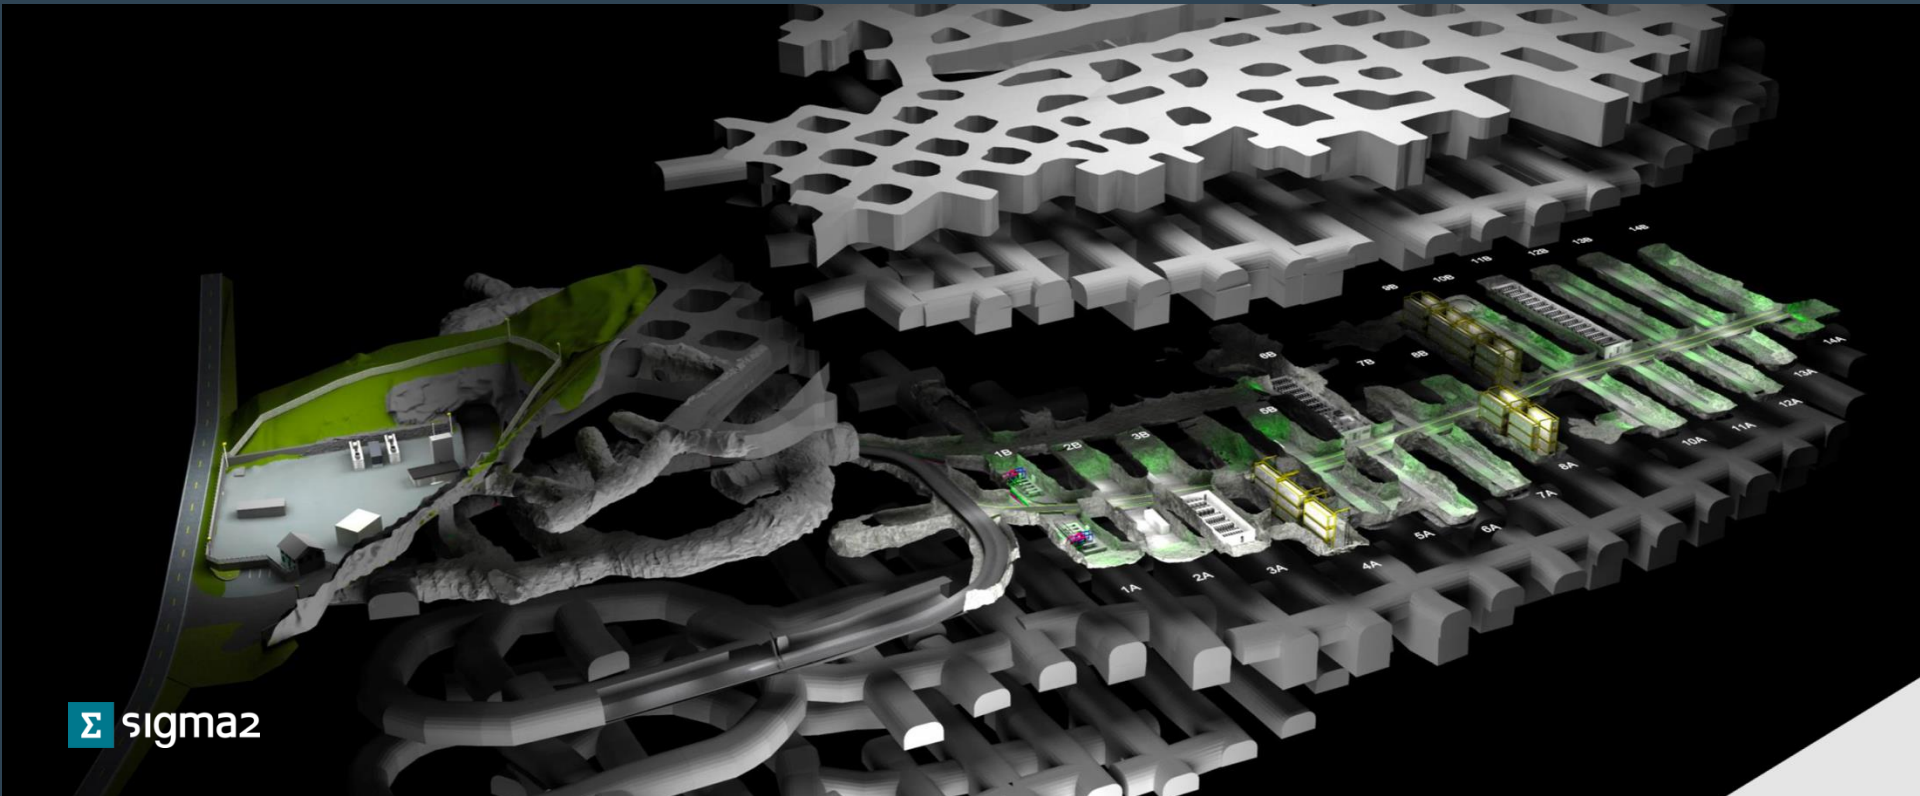


sigma2
NRIS
NIRD - the National
Infrastructure for Research Data
sigma2
NRIS
NIRD - the National
Infrastructure for Research Data


LEFDAL MINE
DATACENTER
sigma2
NRIS
Norwegian research
infrastructure services
6A

Services for computing

System	BullSequana XH2000
Max Floating point performance, double	6.2 Petaflops
Number of compute nodes	1344
CPU type	AMD® Epyc™ 7742 2.25GHz
CPU cores in total	172032
CPU cores per node	128
Memory in total	336 TiB
Memory per node	256 GiB

GPU type	Nvidia A100 40 GB with NVLink
GPUs per node	4
GPUs in total	16
Operating System	Red Hat Enterprise Linux 7
Total disc capacity	2.5 PB
Interconnect	InfiniBand HDR 100, Dragonfly+ topology
Top500 June 2020	55th place @ 1250 nodes, 76% efficiency

----- 7.7 PB

Upcoming: Olivia supercomputer

- HPE Cray Supercomputing EX
- 252 nodes with CPU
- Each node has 2x AMD Turin CPUs
- Each CPU has 128 cores
- **64,512 CPU cores**
- **304 GPUs**
- I**nterconnection:** HPE Slingshot

The LUMI AI factory consortium

- LUMI-AI, a multi-exaflop range for ML training tasks
- AI Factory project for services and support– 100 FTEs (2025-2028)
- LUMI IQ – quantum computer accelerated partition, 300qb, 2400qb in 2028/29
- Dedicated access to top tier machine designed for AI
- Data spaces
- Pretrained foundation models
- Dataset catalog

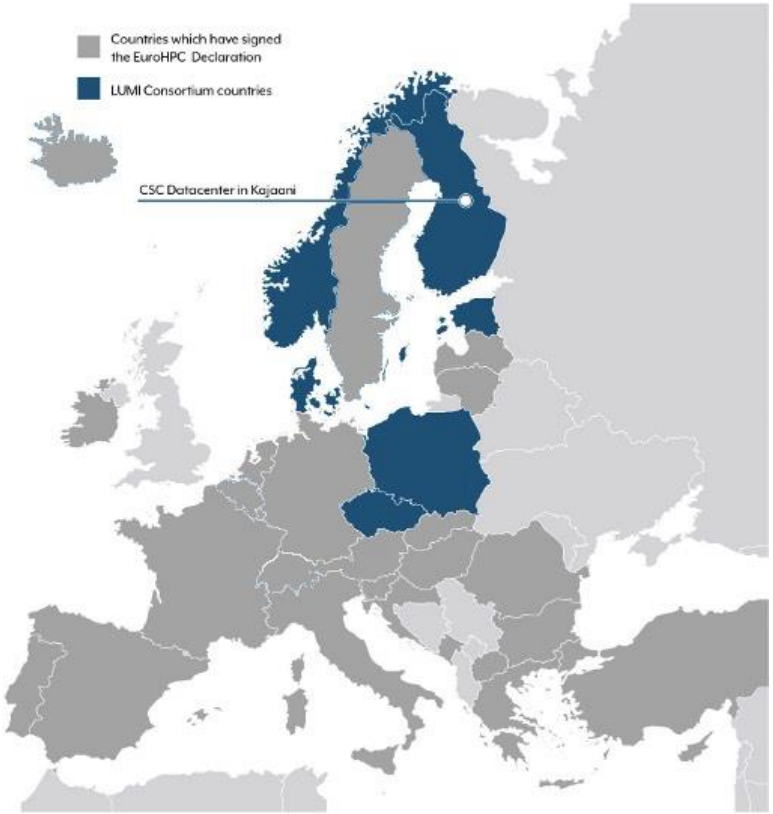

Finland	250 M€	Norway	20,4 M€
Czech Republic	11 M€	Poland	10 M€
Denmark	10 M€	EU	290 M€
Estonia	5 M€	**Total**	**580M€**

Sigma2 – commercial use

Prices for category C

Service	Price
HPC *(valid from 1 October 2024)*	0.13 NOK per CPU-hour
NIRD Data Peak *(valid from 1 October 2024)*	2 070 NOK per TB/year
NIRD Data Lake *(valid from 1 October 2024)*	1 110 NOK per TB/year
Backup *(valid from 1 October 2024)*	1 110 NOK per TB/year

All prices are stated in Norwegian kroner (NOK)

Moving Prodrisk to HPC

Collaboration with SINTEF Energy and Knut Skogstrand Gjerden

- Adaption to HPC tooling: lmod, CMake, Slurm
- Refactoring of libraries
- Removal of proprietary compiler directives and oddities

Prodrisk on HPC - results

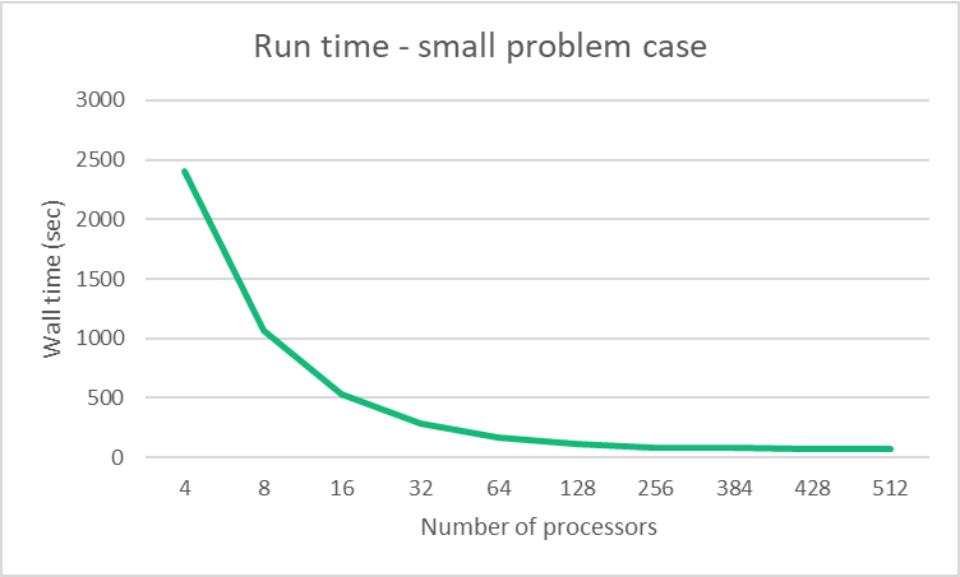

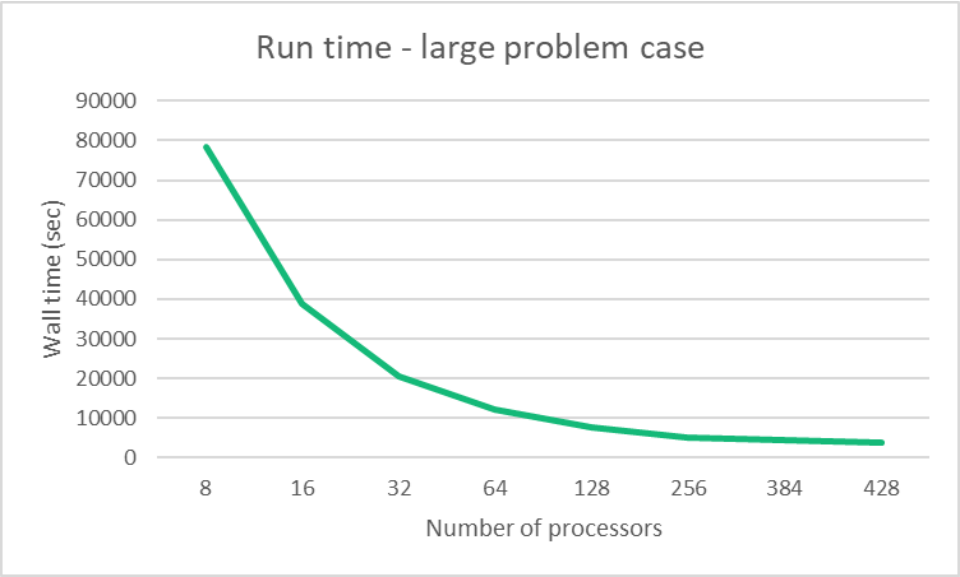

Prodrisk on HPC – CPU-efficiency

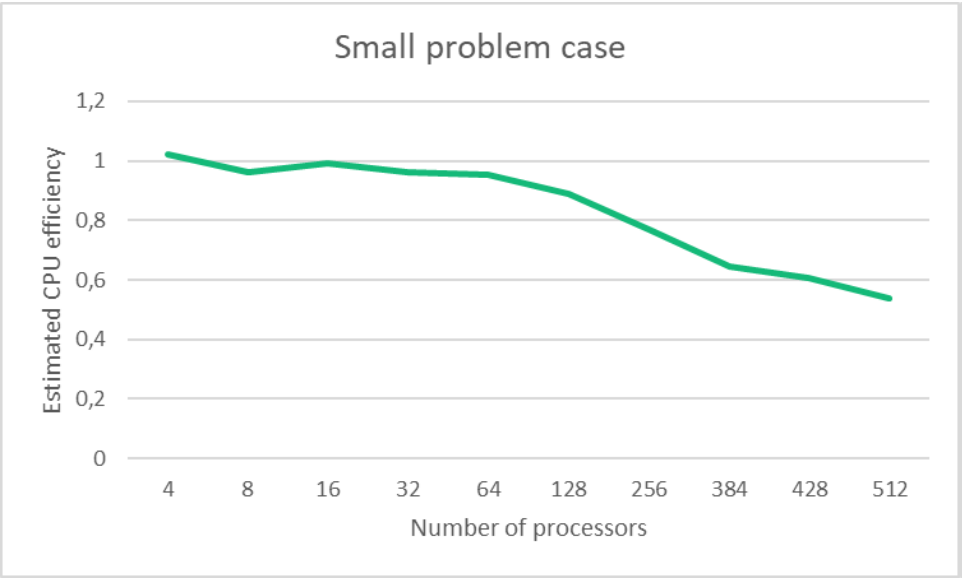

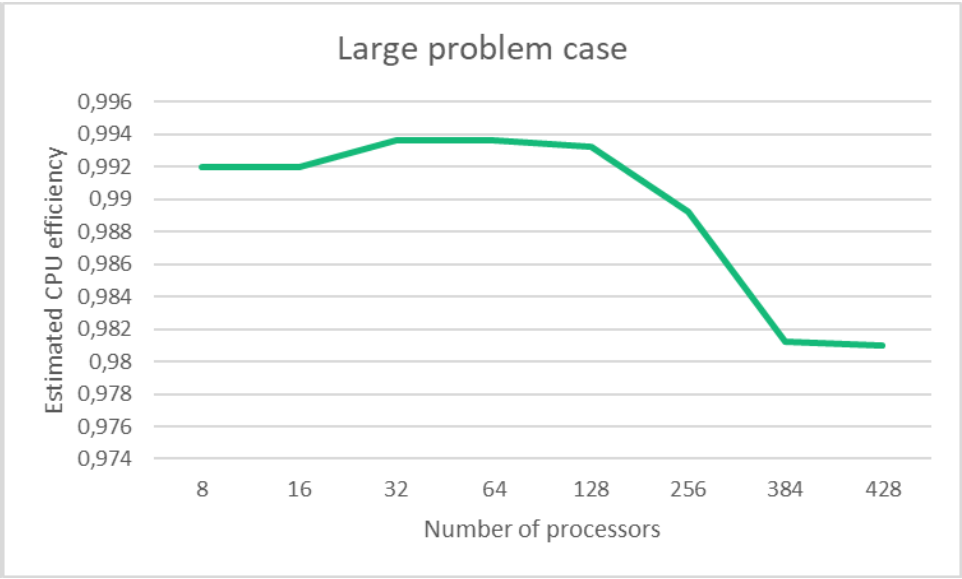

Profiling MPI and CPU usage

Summary: prodrisk is Compute-bound in this configuration

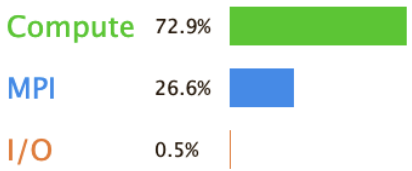

Time spent running application code. High values are usually good.
This is **high**; check the CPU performance section for advice

Time spent in MPI calls. High values are usually bad.
This is **low**; this code may benefit from a higher process count

Time spent in filesystem I/O. High values are usually bad.
This is **very low**; however single-process I/O may cause MPI wait times

This application run was Compute-bound. A breakdown of this time and advice for investigating further is in the CPU section below.

As little time is spent in MPI calls, this code may also benefit from running at larger scales.

CPU

A breakdown of the 72.9% CPU time:

The per-core performance is memory-bound. Use a profiler to identify time-consuming loops and check their cache performance.

Little time is spent in vectorized instructions. Check the compiler's vectorization advice to see why key loops could not be vectorized.

MPI

A breakdown of the 26.6% MPI time:

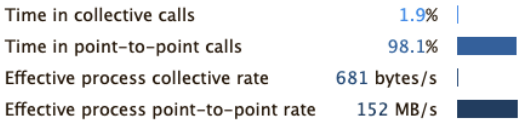

Most of the time is spent in point-to-point calls with an average transfer rate. Using larger messages and overlapping communication and computation may increase the effective transfer rate.

Prodrisk Software as a Service

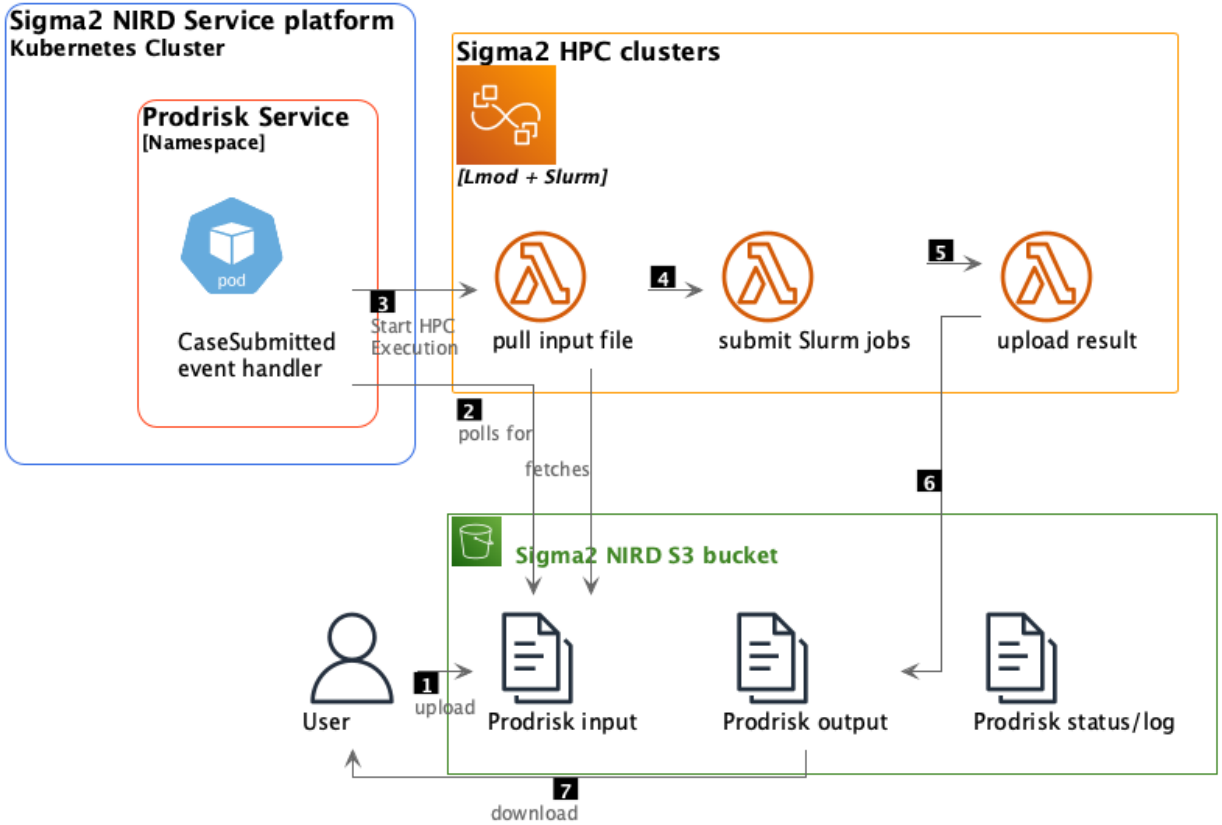

Outlook

- Safe and secure Prodrisk Software as a Service?
- AI-factory
- Sigma2 as a strategic provider for HPC-services for R&D in the energy industry

Acknowledgements

This project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 101101903. The JU receives support from the Digital Europe Programme and Germany, Bulgaria, Austria, Croatia, Cyprus, Czech Republic, Denmark, Estonia, Finland, Greece, Hungary, Ireland, Italy, Lithuania, Latvia, Poland, Portugal, Romania, Slovenia, Spain, Sweden, France, Netherlands, Belgium, Luxembourg, Slovakia, Norway, Türkiye, Republic of North Macedonia, Iceland, Montenegro, Serbia

EuroHPC
Joint Undertaking

Sigma2's role as a provider outside academia