



sigma2

National provider of e-infrastructure

ISC 24 – Norwegian AI landscape

Our strategy

Mission

Strategic responsibility and manages the national e-infrastructure for large-scale data- and computational science in Norway.

for
for
for
kerat
for
for



3 Key strategic Areas

Provide advanced compute and data services addressing users' needs

Facilitator for international research services, cloud-based resources and common components

Combine Sigma2 infrastructure with the capability and competencies of partner institutions to achieve excellent research services

About Sigma2 AS (Ltd)



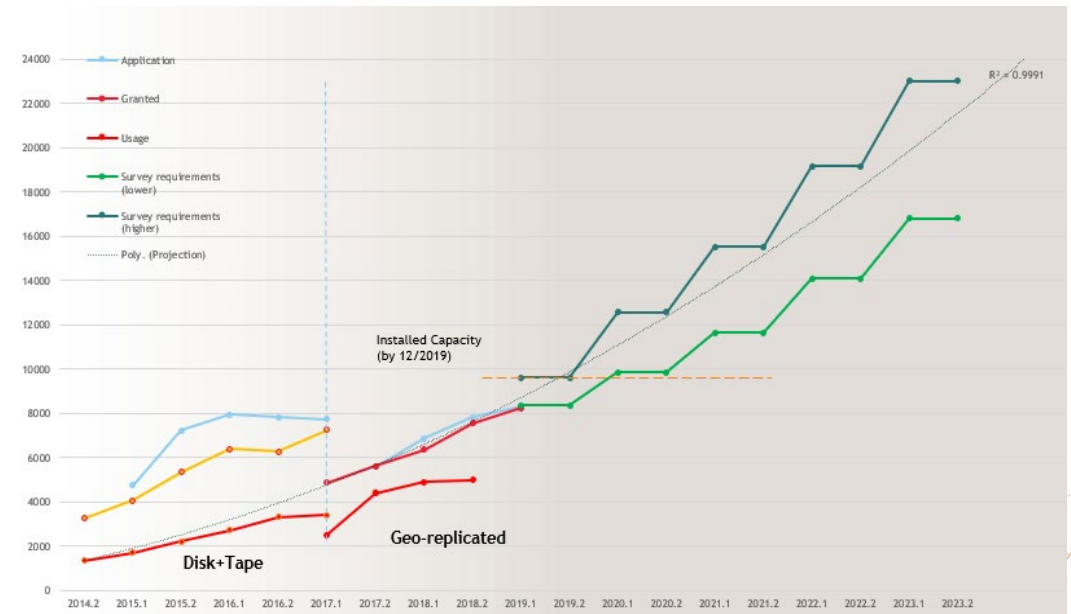


- Established in 2015, all shares now owned by Government (14+3 FTE)
- 6+6 model → evaluation after 6 years (positive evaluation in 2019, next in 2025)
- Governed by a board with 7 members (Major usage organisations and 2 external representatives)
- 50 MNOK yearly from the Norwegian Research Council (RCN)
- 53 MNOK yearly guaranteed user payment from the 4 “main” universities
 - Cooperation agreement
 - Operation and support contract (30 FTE)
 - Frame agreement for project work (5FTE)
- 2200+ users across 600+ projects

RCN competitive infrastructure investment funding

- 2016-2017: 75 MNOK
- 2018-2019: 115 MNOK
- 2020-2021: 337 MNOK*
- 2022-2024: 154 MNOK

User involvement



- Customer and user meetings
- Yearly user satisfaction survey
- Regular user needs surveys



Norwegian research infrastructure services (NRIS)



Lefdal Mine Datacenter




Sigma2 in collaboration with the partner universities provides and operate NRIS

Hardware located currently on locations

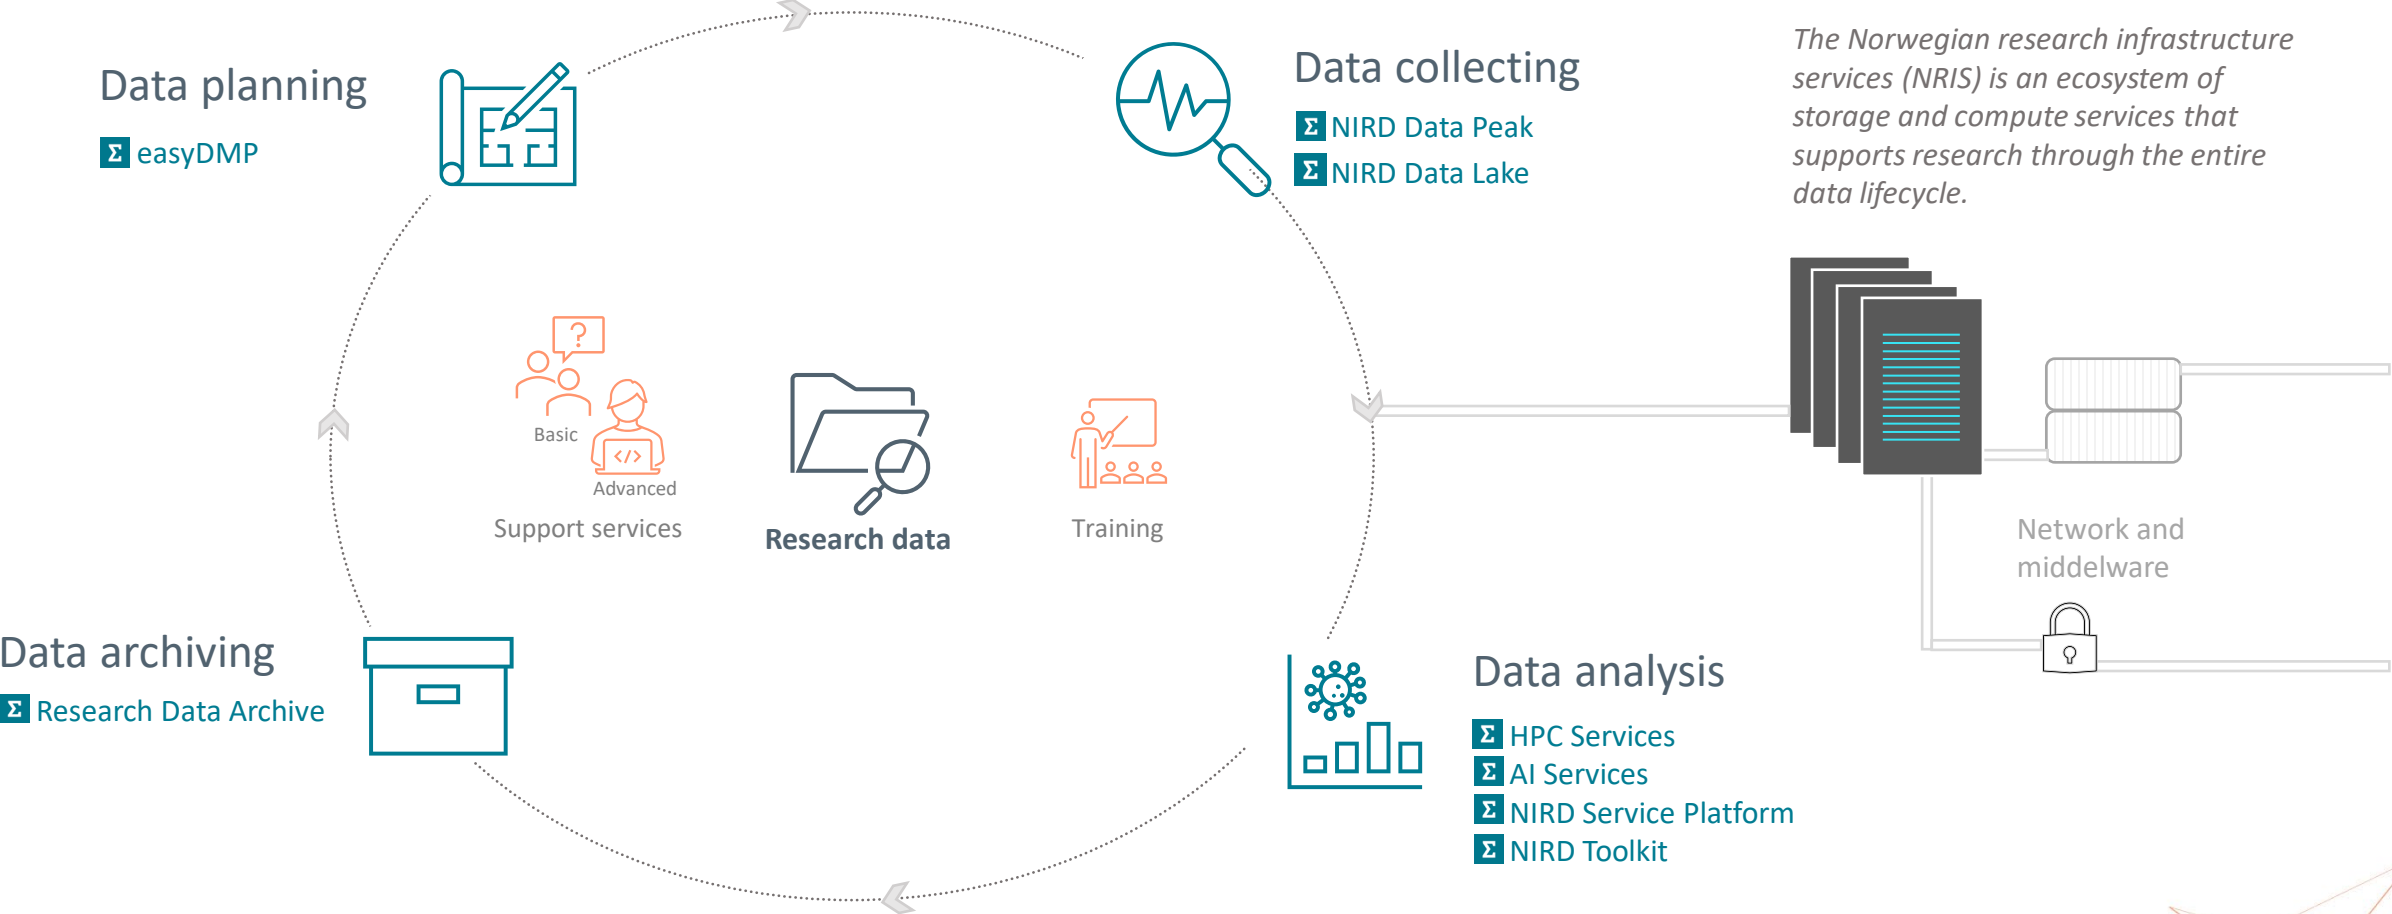
- Tromsø
- Trondheim
- Oslo
- Lefdal Mine Datacenter (LMD)

The infrastructure sits on the Sikt (formerly Uninett) research network backbone (100 Gbps)

Our new datacenter



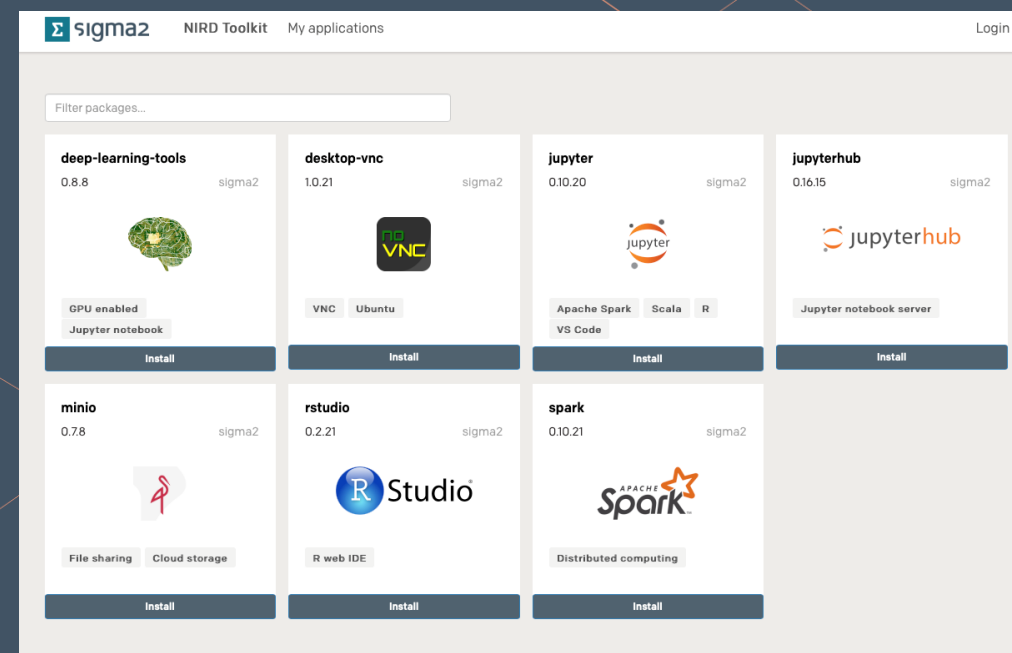
E-infrastructure services landscape for research



NIRD Service Platform GPU hardware

- The NIRD Service Platform (SP) is a Kubernetes based cloud infrastructure directly connected to the NIRD storage infrastructure 44+ PiB (48,5 PB) in Lefdal Mine Datacenter
- Services deployed on the NIRD SP can access / consume **large data volumes**
- **GPUs available on the SP**
 - LMD: 8 worker nodes with: 128 CPU cores 4 x 512 GiB + 4 x 1024 GiB RAM 10 x NVIDIA V100 GPUs
 - Trondheim: 4 GPU nodes with 2 CPUs (40 cores, 768 GiB RAM) and 4 NVIDIA V100 GPUs each

AI/ML applications permanently deployed or running on demand through the **deep-learning toolkit** available on the NIRD SP



Data Portals

The National Infrastructure for Research Data

Use Case: **CO₂ Storage and Capture**

Collaboration between Equinor, SINTEF and Sigma2

“en nettbasert, digital plattform for deling av referanse-datasett fra banebrytende CO₂-lagringsprosjekter”

Klikk for å redigere

<https://co2datashare.org/>

DOI: 10.11582/2020.00005



The screenshot shows the CO₂ DataShare portal website. The header features the logo "CO₂ DataShare" and navigation links for "Datasets", "Organizations", and "About", along with a search bar. The main content area displays a welcome message: "Welcome to the CO₂ DataShare portal" followed by a description: "The CO₂ DataShare online portal is an initiative by the international CO₂ Storage Data Consortium, with the goal to make available curated and well-documented datasets from CO₂ storage projects." Below this is a graphic titled "CO₂ storage data consortium" showing a central database icon connected to various devices (laptop, tablet, smartphone, desktop monitor) and a cloud icon. At the bottom of the graphic, it reads "CO₂ Data Share: The online platform of the CO₂ Storage Data Consortium". The footer includes logos for "Powered by ckan" and "Hosted by UNINETT sigma2".

LUMI



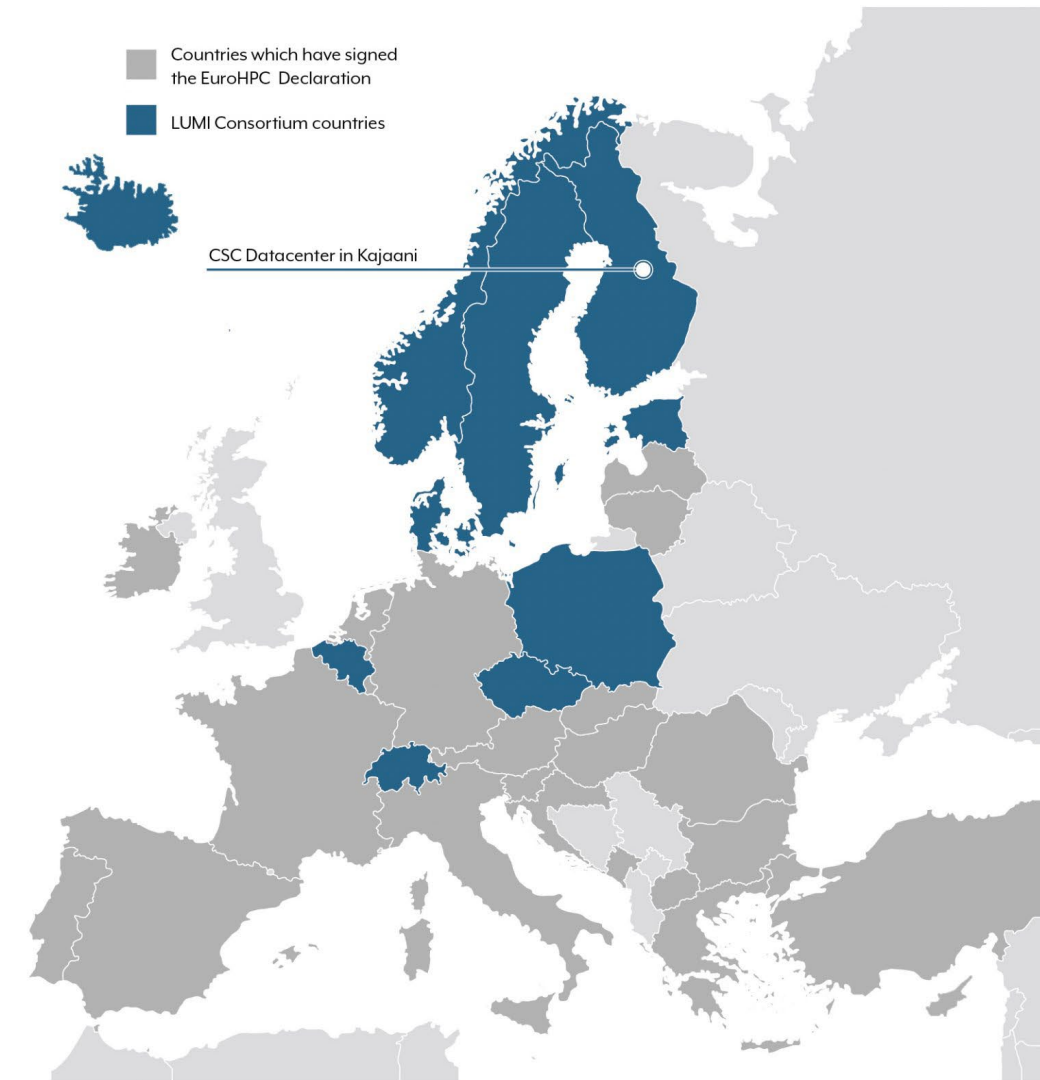
The Queen of the North

The LUMI Consortium

- Unique consortium of 10 (now 11 with NL joining) countries with strong national HPC centers and competence gives a unique opportunity for knowledge transfer and sharing and providing user support for the system
- Norwegian researchers and industry get dedicated access to top tier machine (375 PFLOPS)
- The total budget of the LUMI EuroHPC pre-exascale system is over 202 MEUR

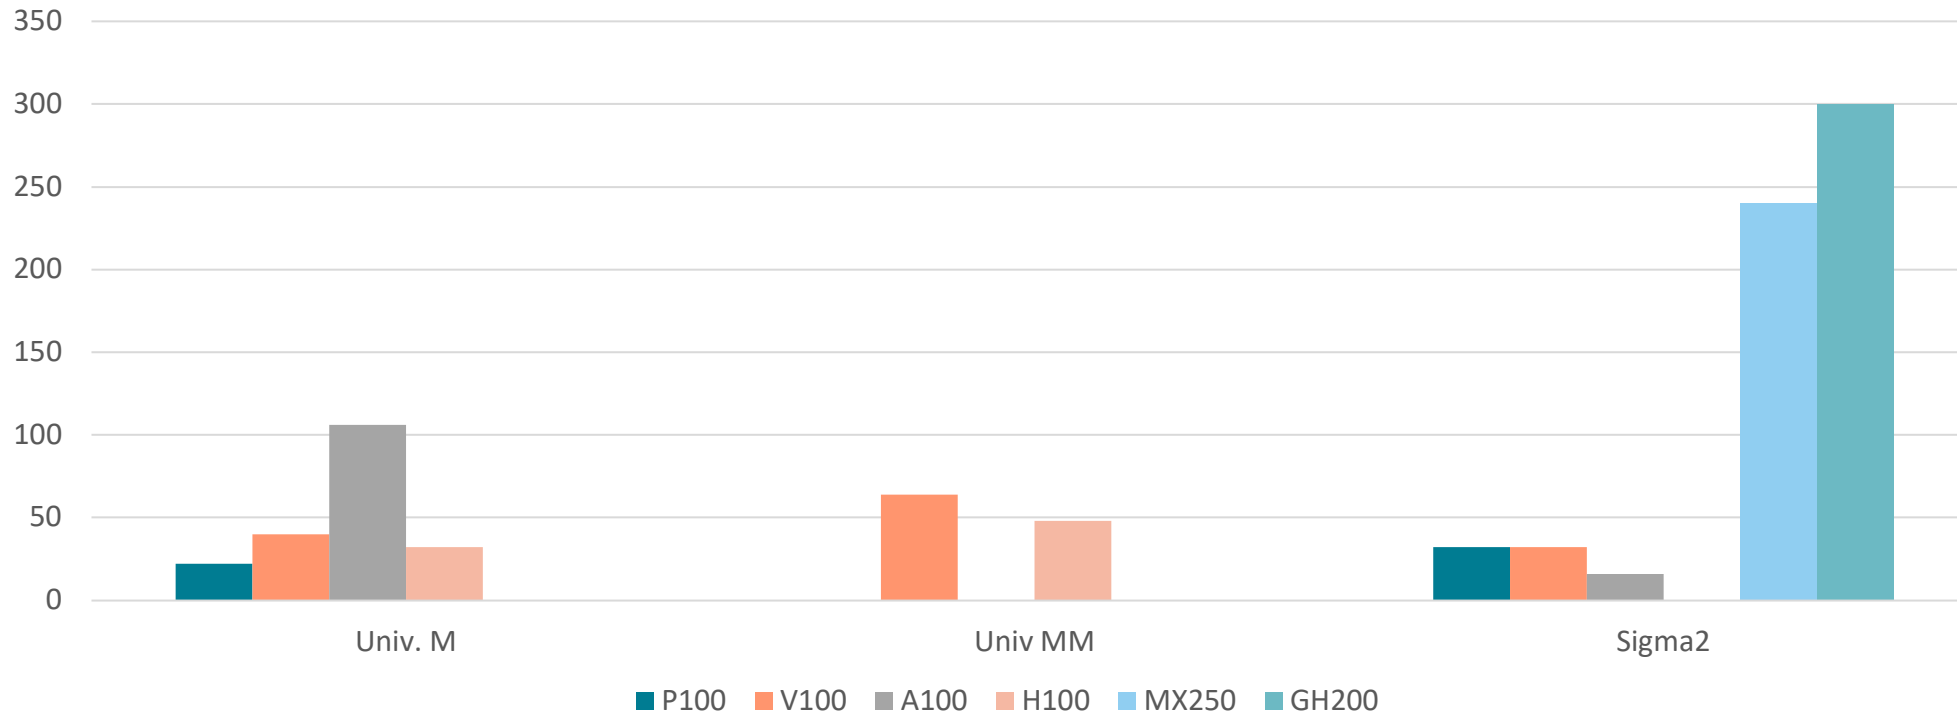
Finland	50 M€	Norway	4 M€
Belgium	15.5 M€	Poland	5 M€
Czech Republic	5 M€	Sweden	7 M€
Denmark	6 M€	Switzerland	10 M€
Estonia*	2 M€	EU	102.5 M€

- 2% Norwegian ownership in LUMI (240 GPUs)
- Actually 30% of national investment funding
- Scale of economy



GPU systems i Norway 2024/25

GPUs for research





Advanced User Support (AUS)

- Applications accepted continuously
- Evaluation by the NRIS management with help from domain expertise from the Resource Allocation Committee (RFK)
- Optimal NRIS resource person and is allocated time to work on the project
- Free of charge

More information on:

www.sigma2.no/advanced-user-support

AI related NCC project portfolio

- AI for determining sea biomass from fleet sonar data
- ML to detect production fault in glass fibre production
- LLM on scientific articles for research «ChatGPT»
- AI for automatic cleaning of 10's of Millions custom declarations
- AI for novelty detection of new trends in cleaned customs traffic
- AI for detecting field deliniation for increased crop yield



Success story, from small scale to LUMI: Digifarm



- 40 % of global fields are over-fertilized, leading to production drop between 15-20%
- How much seeds are required per field?
- What seeds are better suited for each field?



Accurate boundary data can help the estimated 570 million farmers and growers worldwide boost their yields and lower their input costs.




Image courtesy: Digifarm, ESA

Success story, from cloud to LUMI: Digifarm

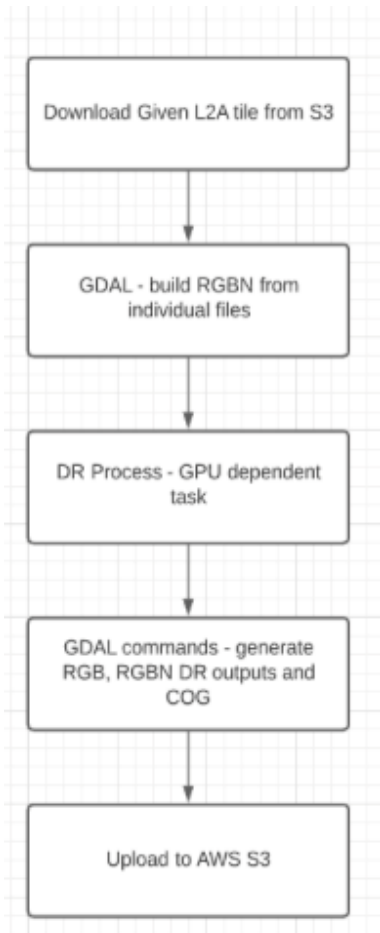


- Delineated Field Boundaries
 - Automated detection of fields boundaries and seeded acres
- Productivity zones
 - Historical biomass accumulation in low, medium and high zones
- S2 Time Series
 - Long-term and in-season cloud filtered vegetation indices on the same chart
- Deep resolution Imagery
 - Sentinel-2 images up to 1 meter deep resolution
- Sustainability index
 - Multi-year analysis and vegetation indices using proprietary algorithm
- Crop classification
 - Based on extensive ground-data, as well as in growt season

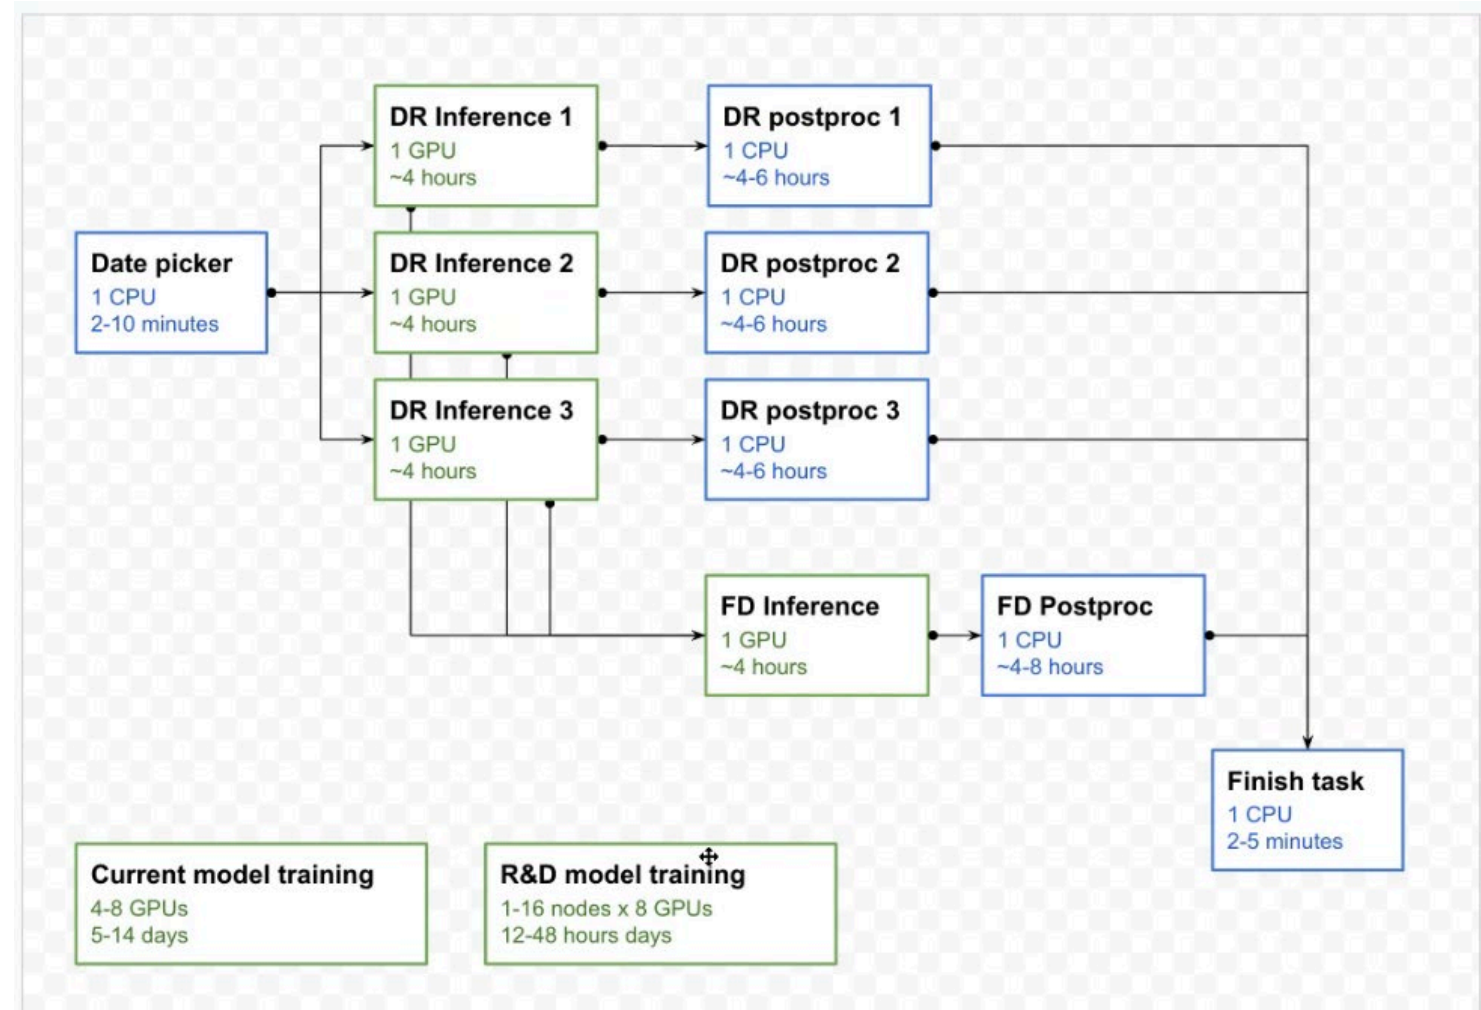


Success story, from cloud to LUMI: Digifarm

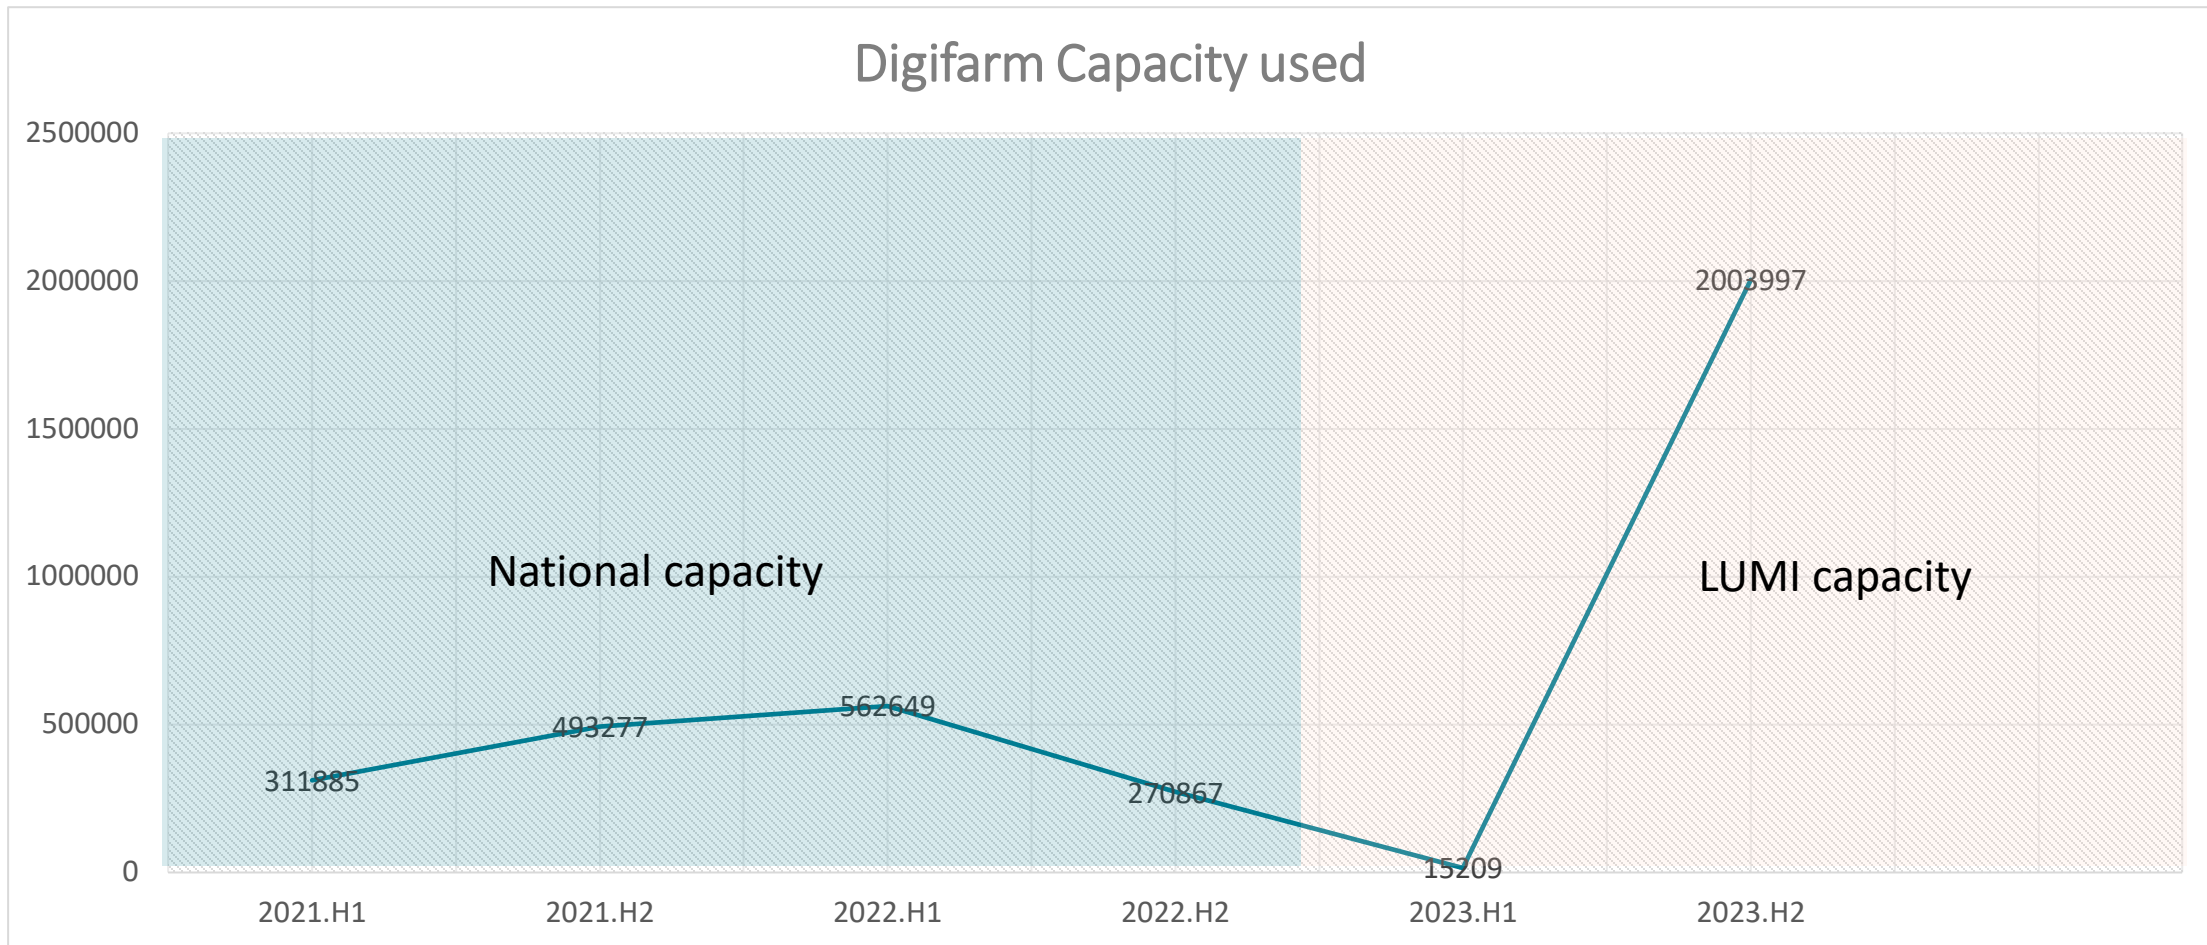
Serial cloud workflow



HPC ready workflow



Larger resources promote increased usage



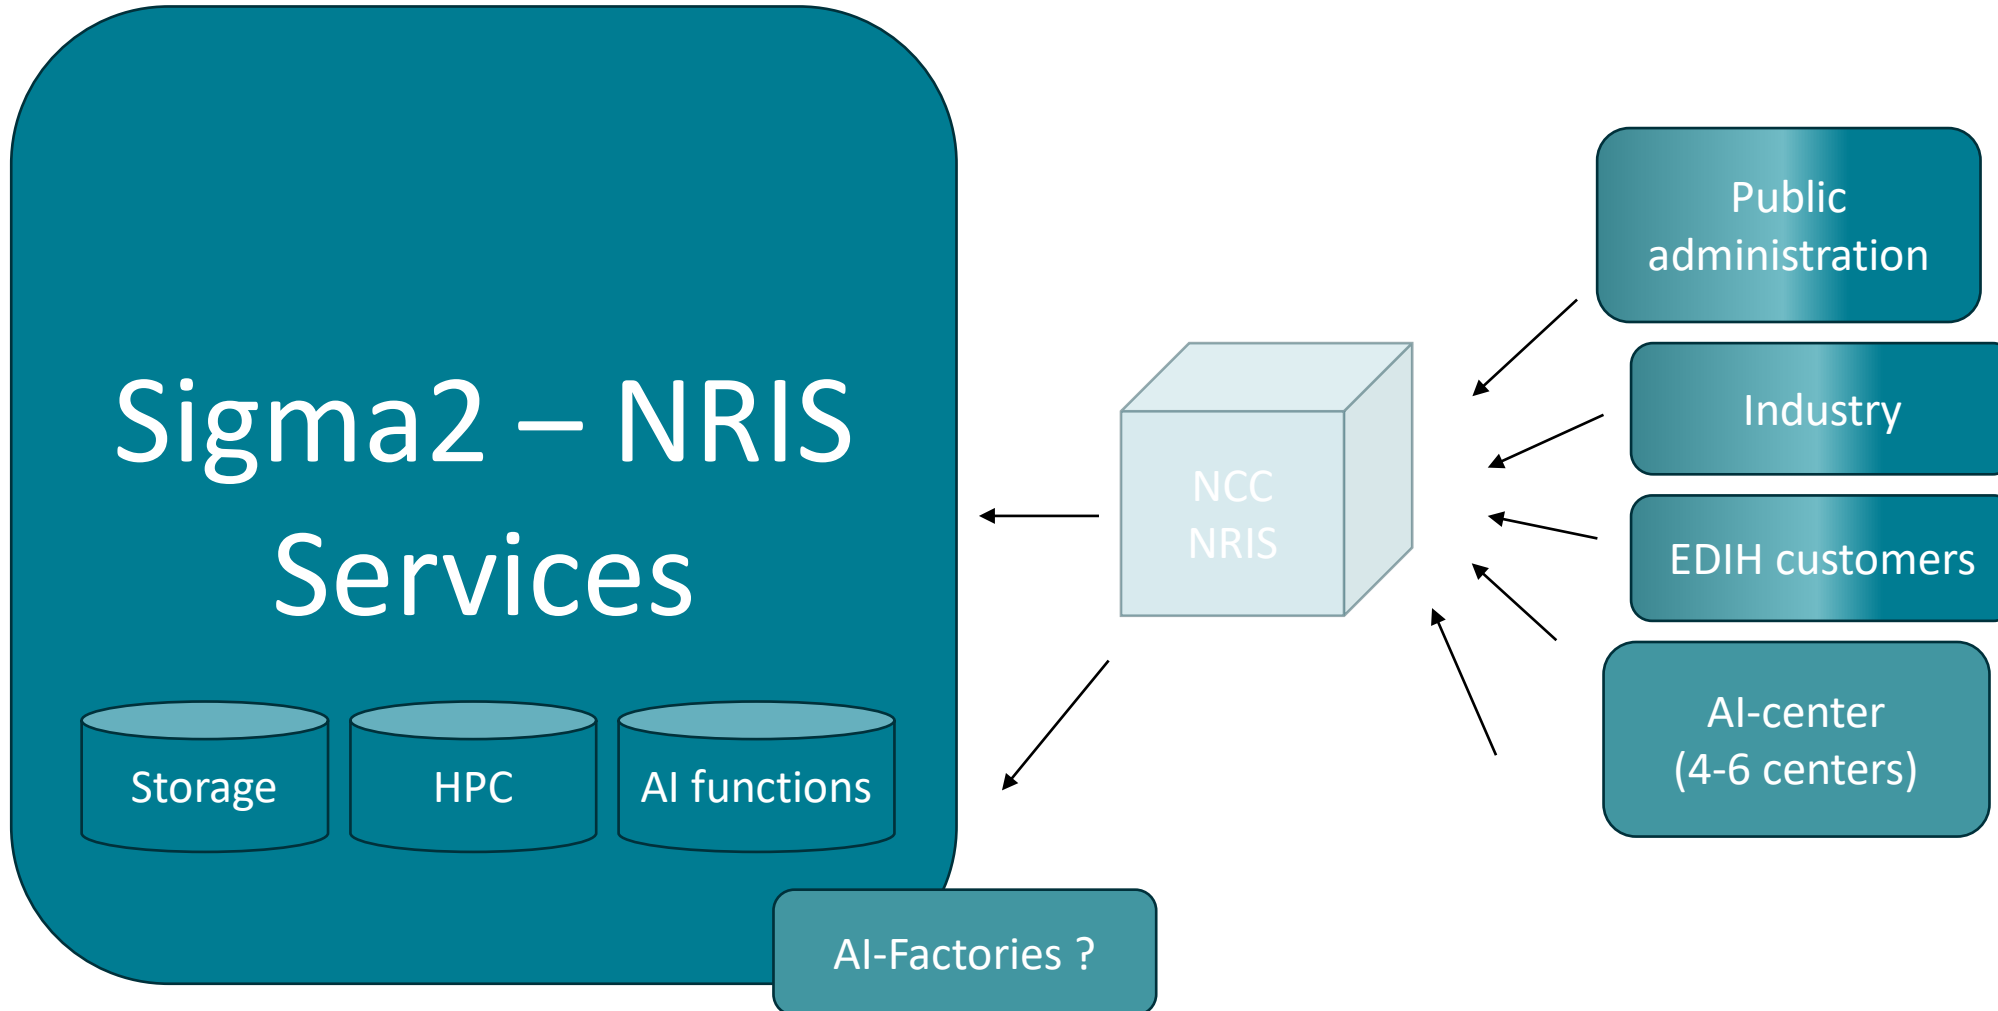
*Consumed resources in CPU hour equivalents

Larger resources promote better models



- Increased training data to 4 million acres
- Leveraging all 12 spectral bands from Sentinel-2
- Training model from 4 classes to 12 classes
 - Crop fields
 - Tree crops
 - Pastures
 - Water ways
 - Urban area
 - Etc
- Can now support EUDR by detecting deforestation

NCC as an enabler for new user groups



Thank you for your attention

www.sigma2.no

