Distributed Deep Learning by Horovod

Jozsef Kovacs, SZTAKI
jozsef.kovacs@sztaki.hu

ELKH Eötvös Loránd Kutatási Hálózat

SZTAKI

www.neanias.eu

NEANIAS

Novel EOSC Services for Emerging Atmosphere, Underwater & Space Challenges

NEANIAS receives funding from European Union under Horizon 2020 Research and Innovation Programme under grant agreement No. 863448

NEANIAS: Services, Infrastructures, Communities & Business

NEANIAS - Novel EOSC Services for Emerging Atmosphere, Underwater & Space Challenges
Call: INFRAEOSC-02-2019
Project ID: 863448
Duration: 36 months
Finish: 31.10.2022

- Address community-specific needs for underwater, atmosphere and space research sectors
- Onboard communities to the Open Science, EOSC and interdisciplinary research era
- Nurture new business opportunities
- Power-up EOSC

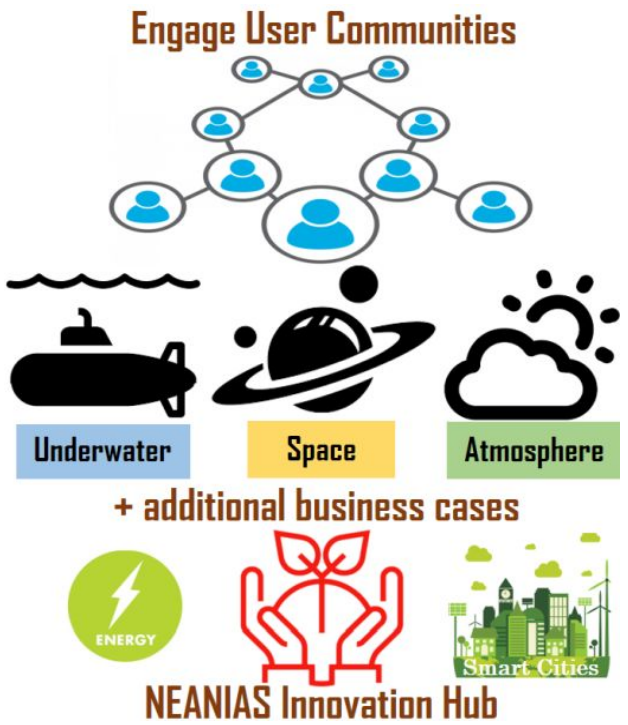

Neanias project structure – Core services

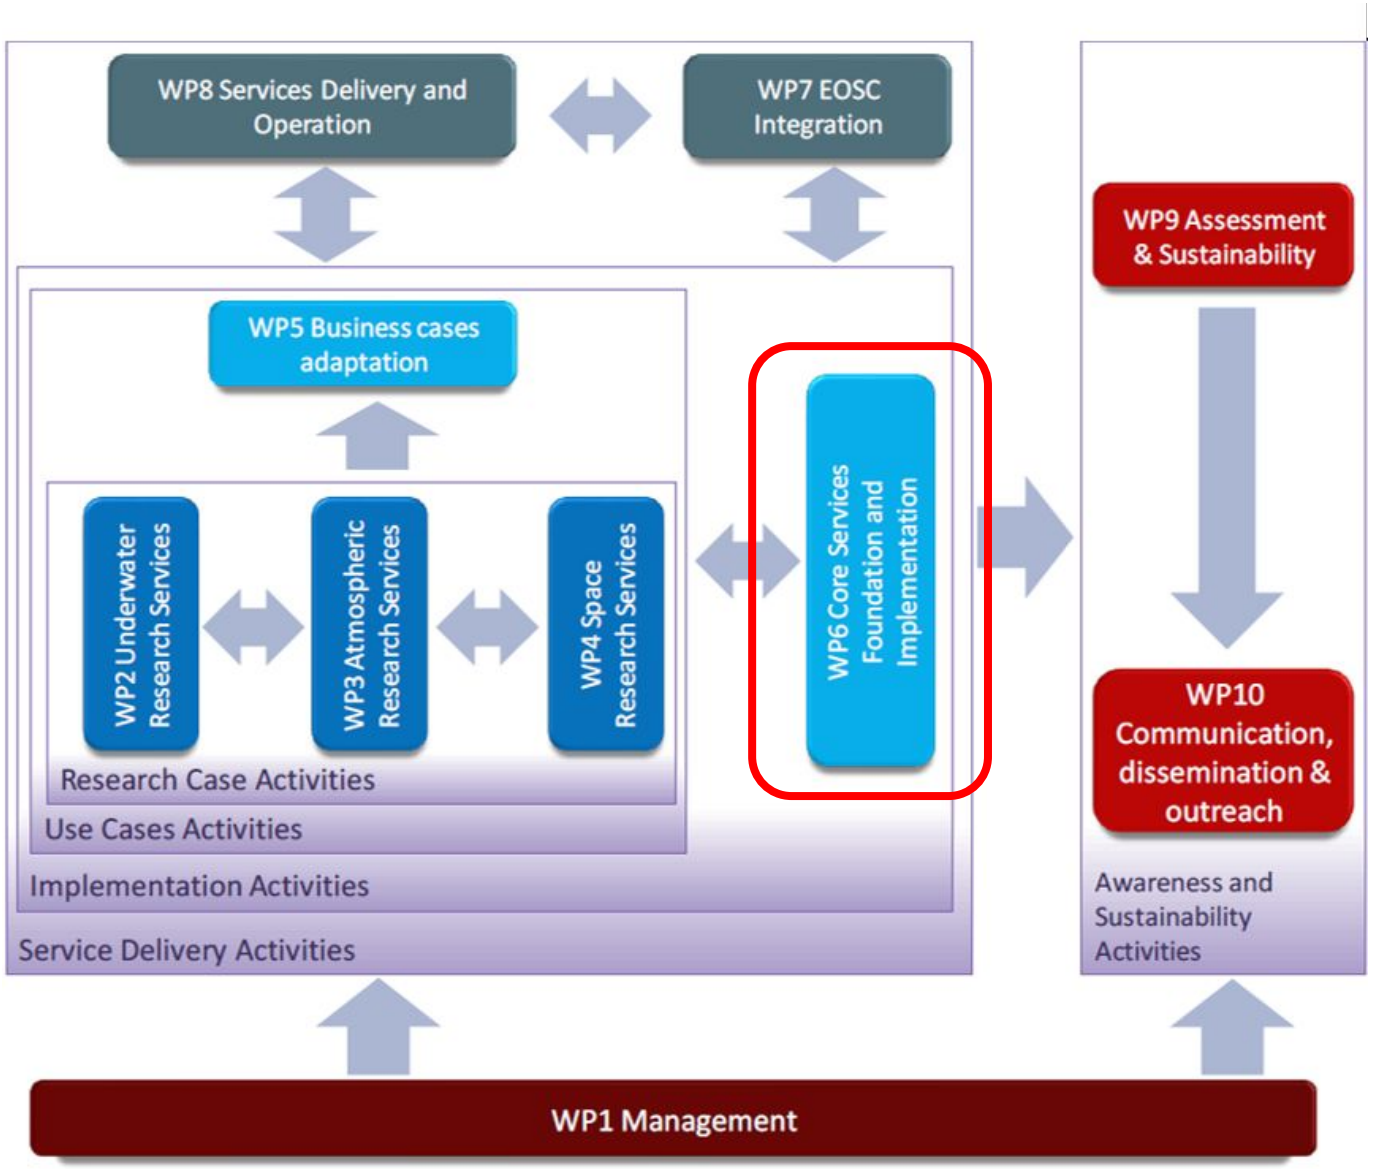

WP6 services

- C1 Open-science lifecycle support services
- C2 EOSC, RI and cloud integration enabling service
- C3 Artificial Intelligence service
- C4 Visualisation services

AI Gateway and its integrated components

› Basic core services

› Development and training of ML models

› Model serving of ML models

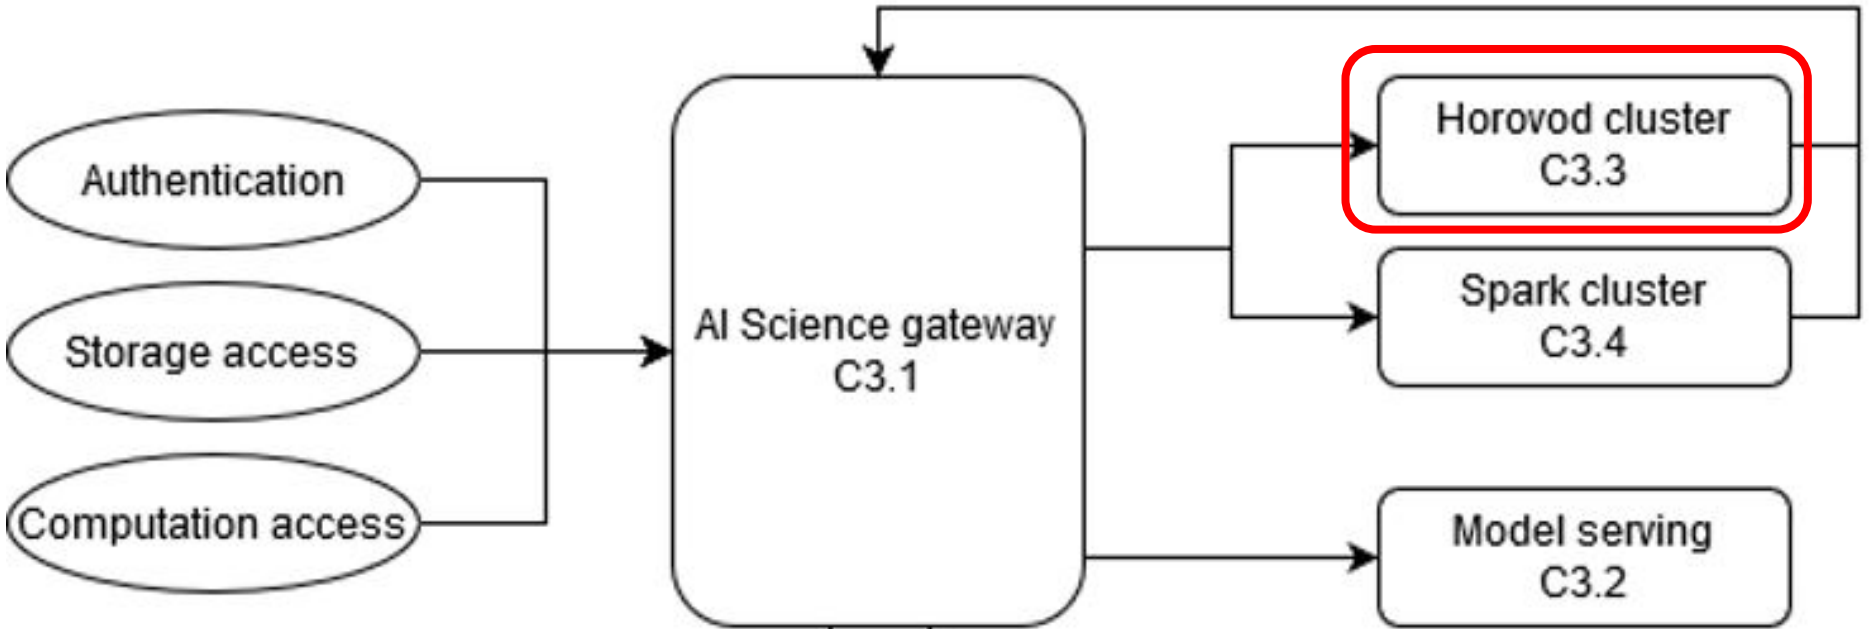

Horovod

- Open-source distributed deep learning framework from Uber
- Supports TensorFlow, Keras, PyTorch, Apache MXNet and Spark
- Provide an easy-to-use framework for distributed training
  - Execute on hundreds of GPUs with just a few lines of additional code
  - Great scaling efficiency
- Data parallel execution
  - Training data is divided into subsets
  - Train the same replicated model on each node
  - Model parameters are synchronised between the workers
- Ring-Allreduce strategy
  - Horovod utilizes Message-Passing Interface (MPI)
  - Each node communicates with two of its peers 2*(N-1) times
  - NVIDIA NCCL 2.0 for intra-node communication

Workflow Overview

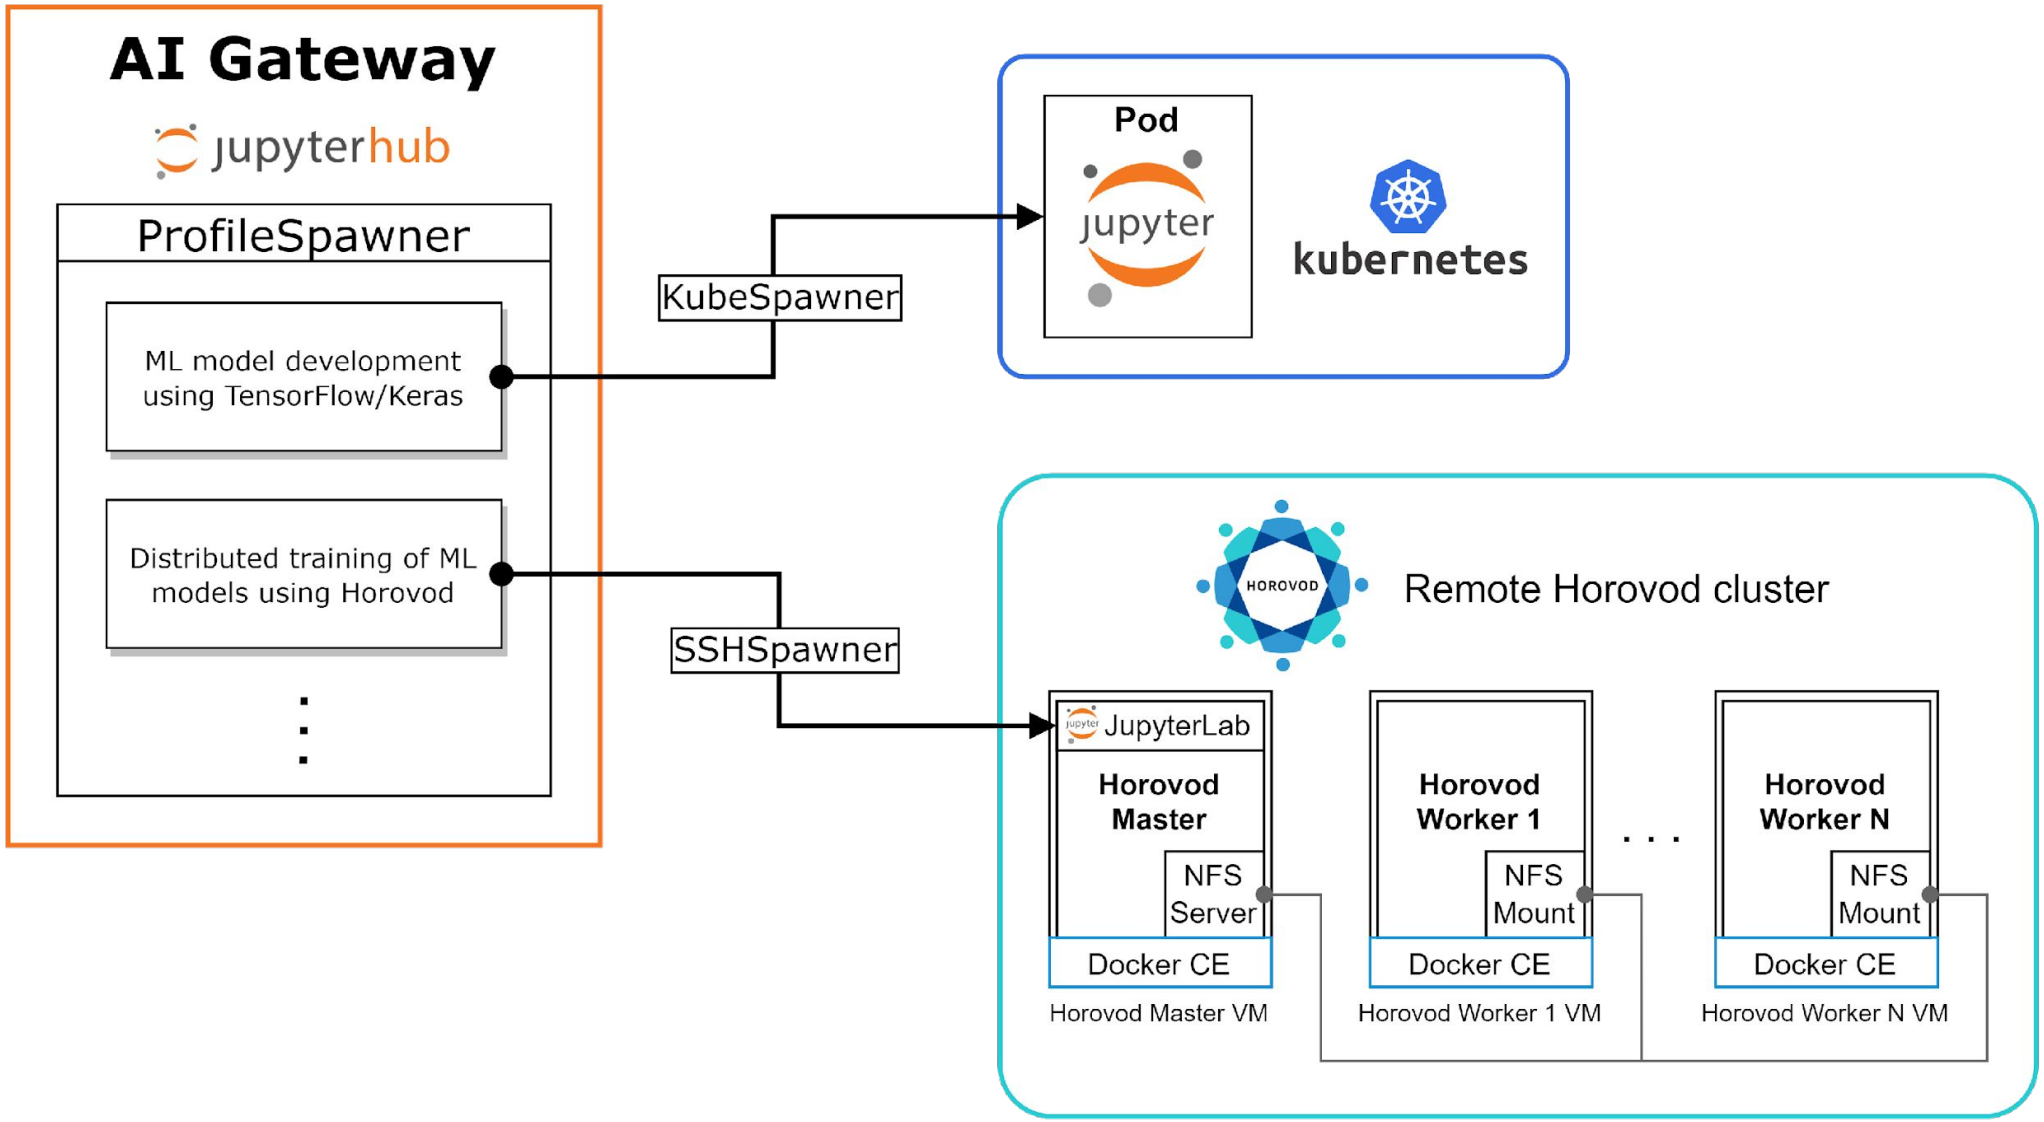

Monitoring with Prometheus Grafana stack

Integration into NEANIAS AI-Gateway

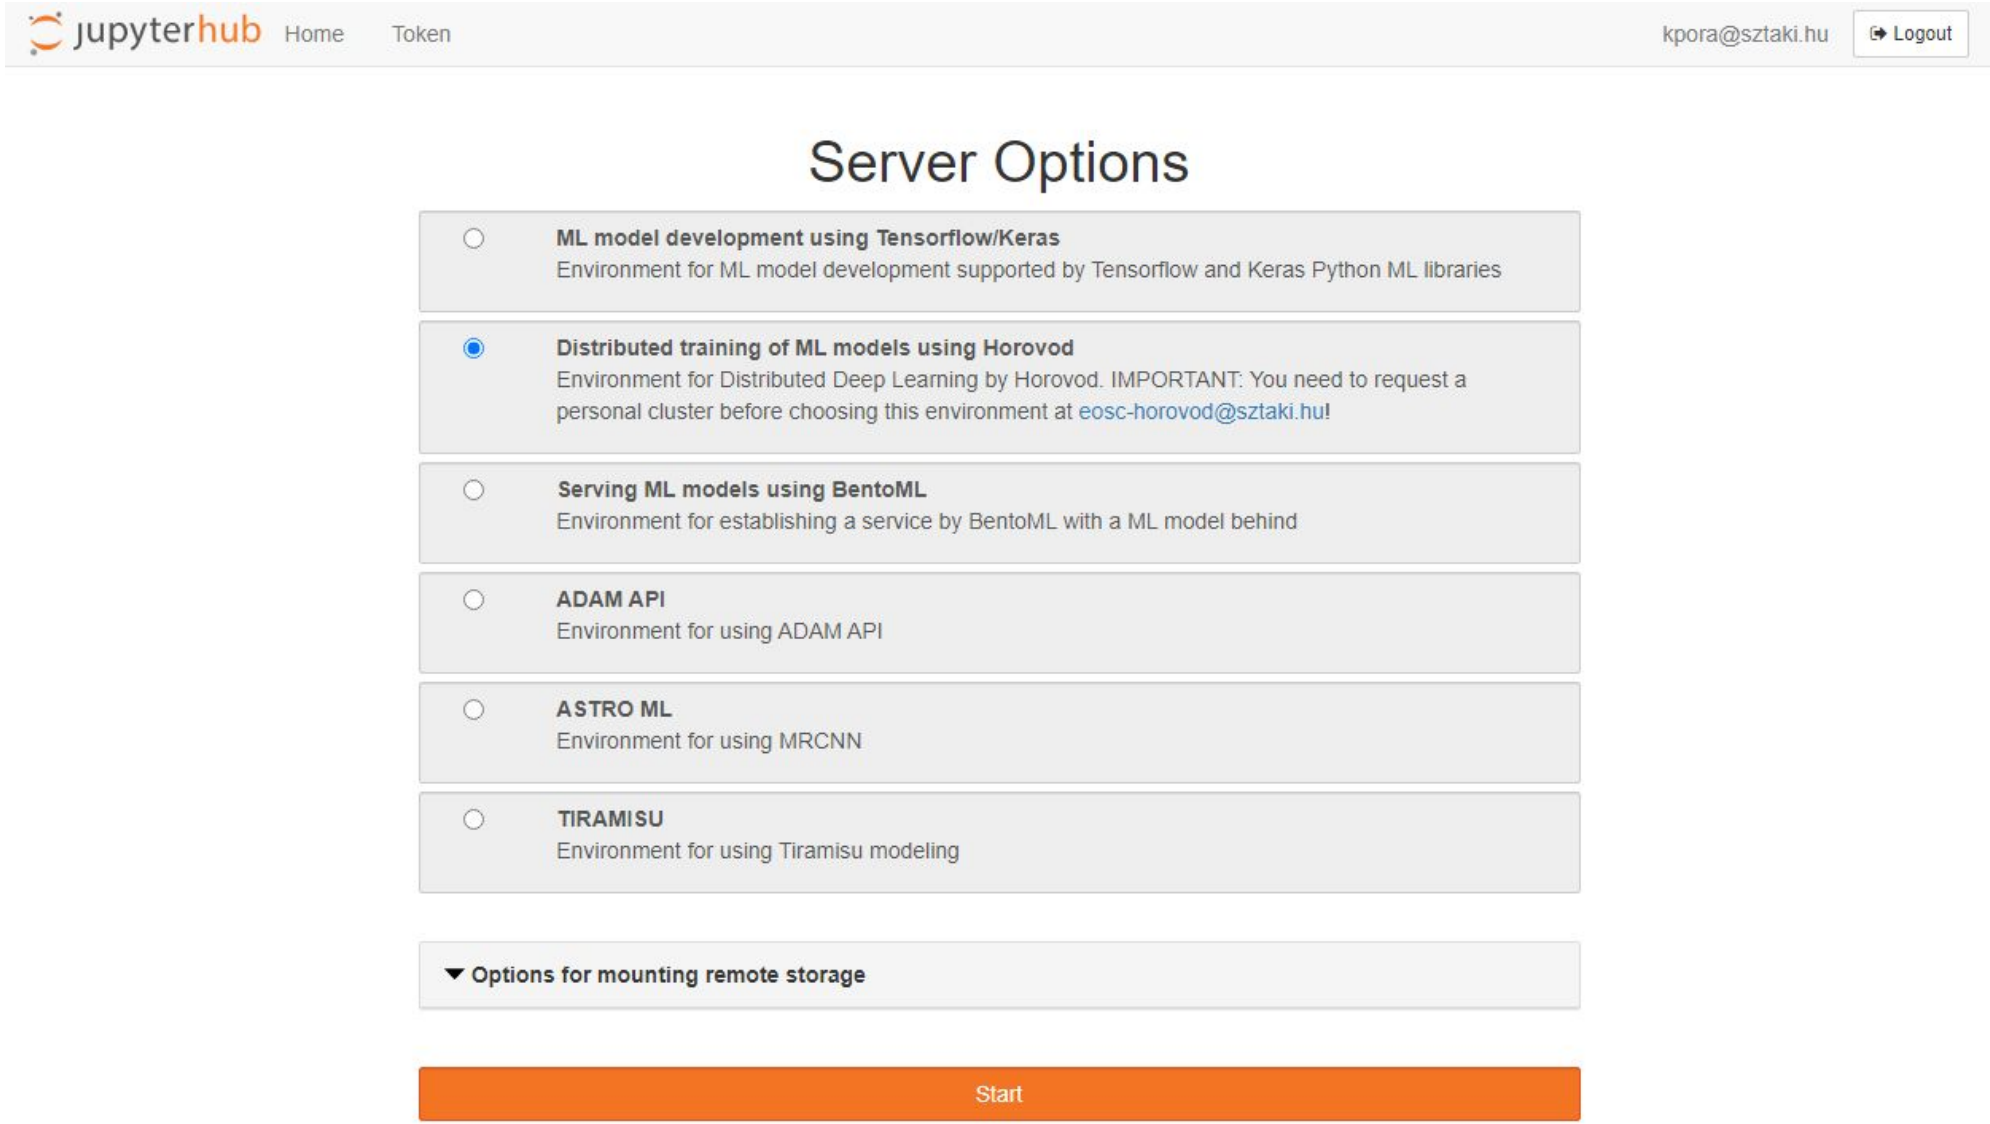

Intro page

AI-Gateway Environments

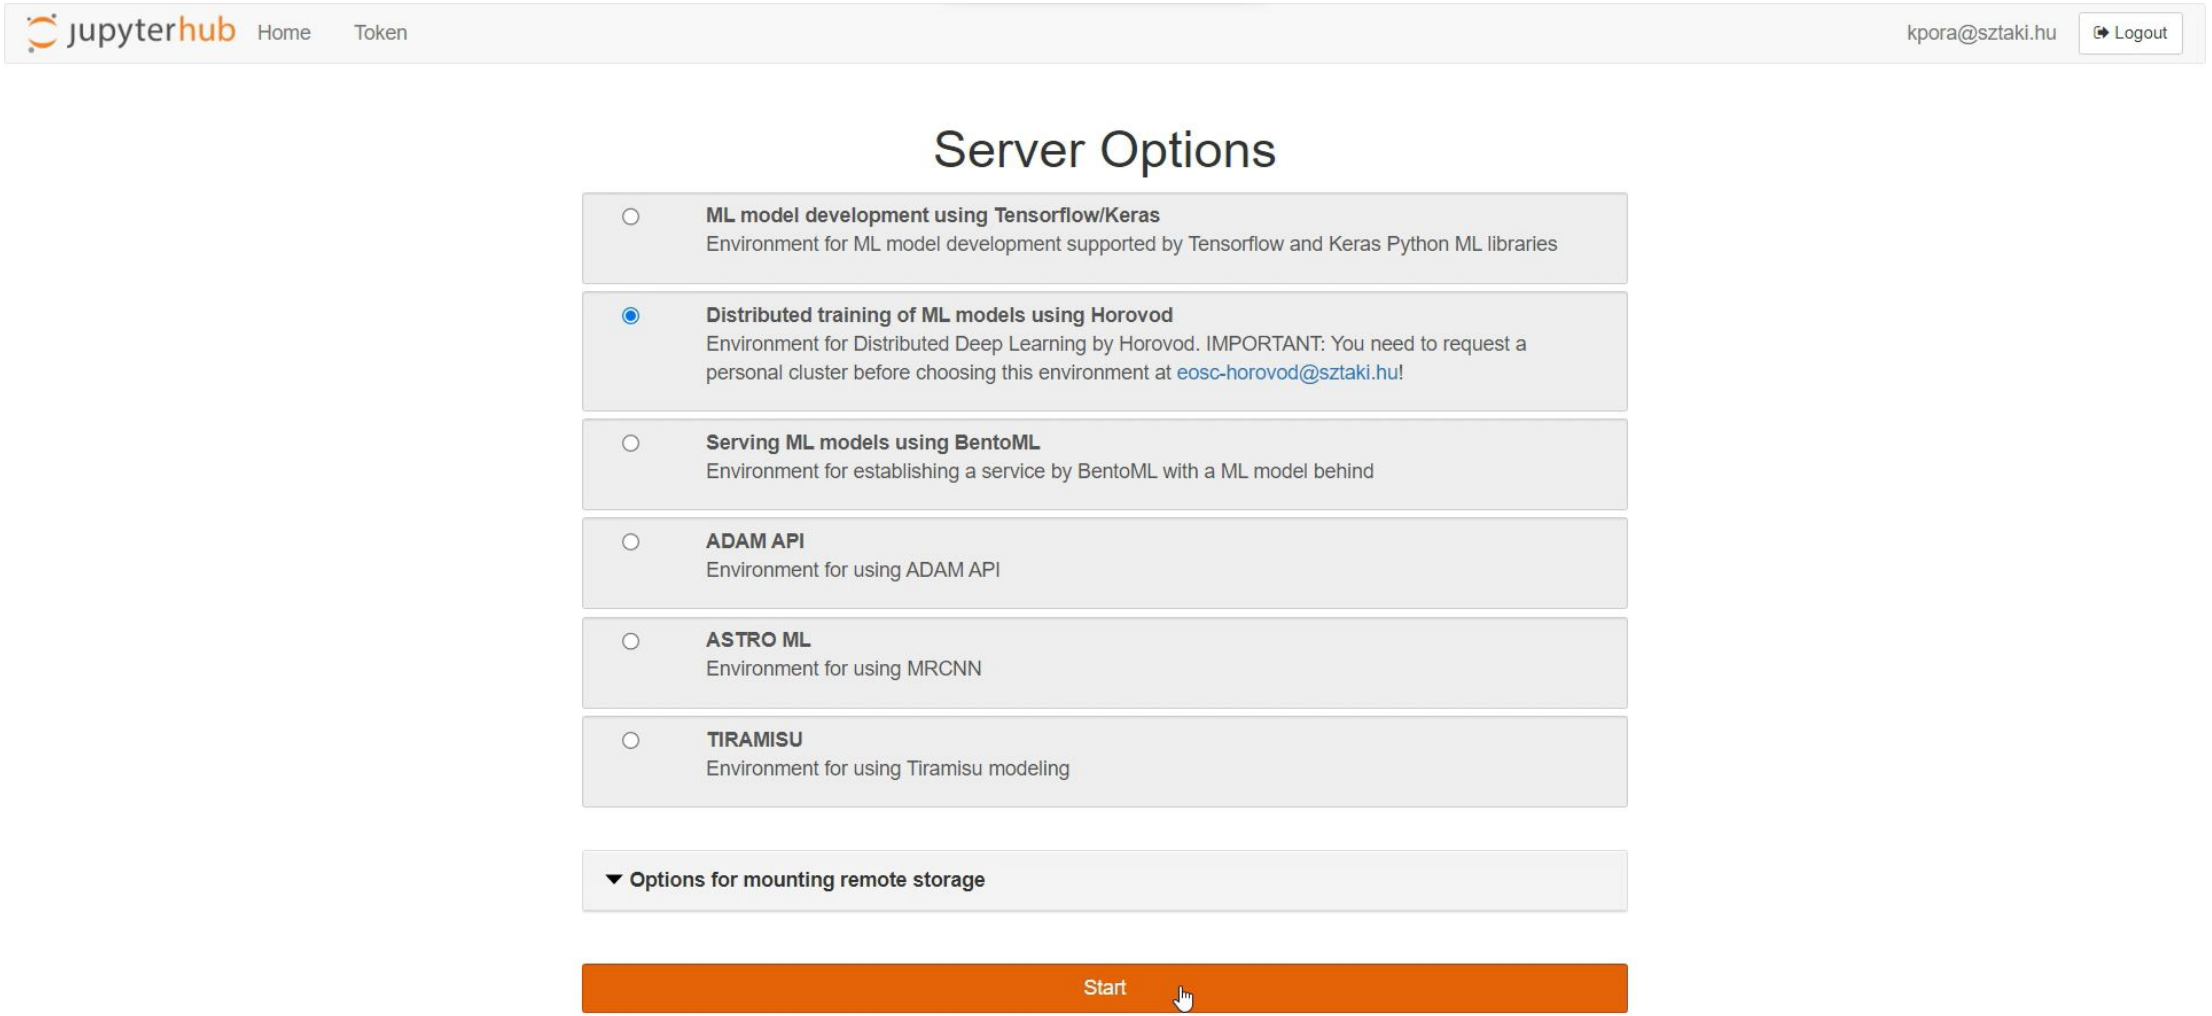

Starting the single-user server using SSHSpawner

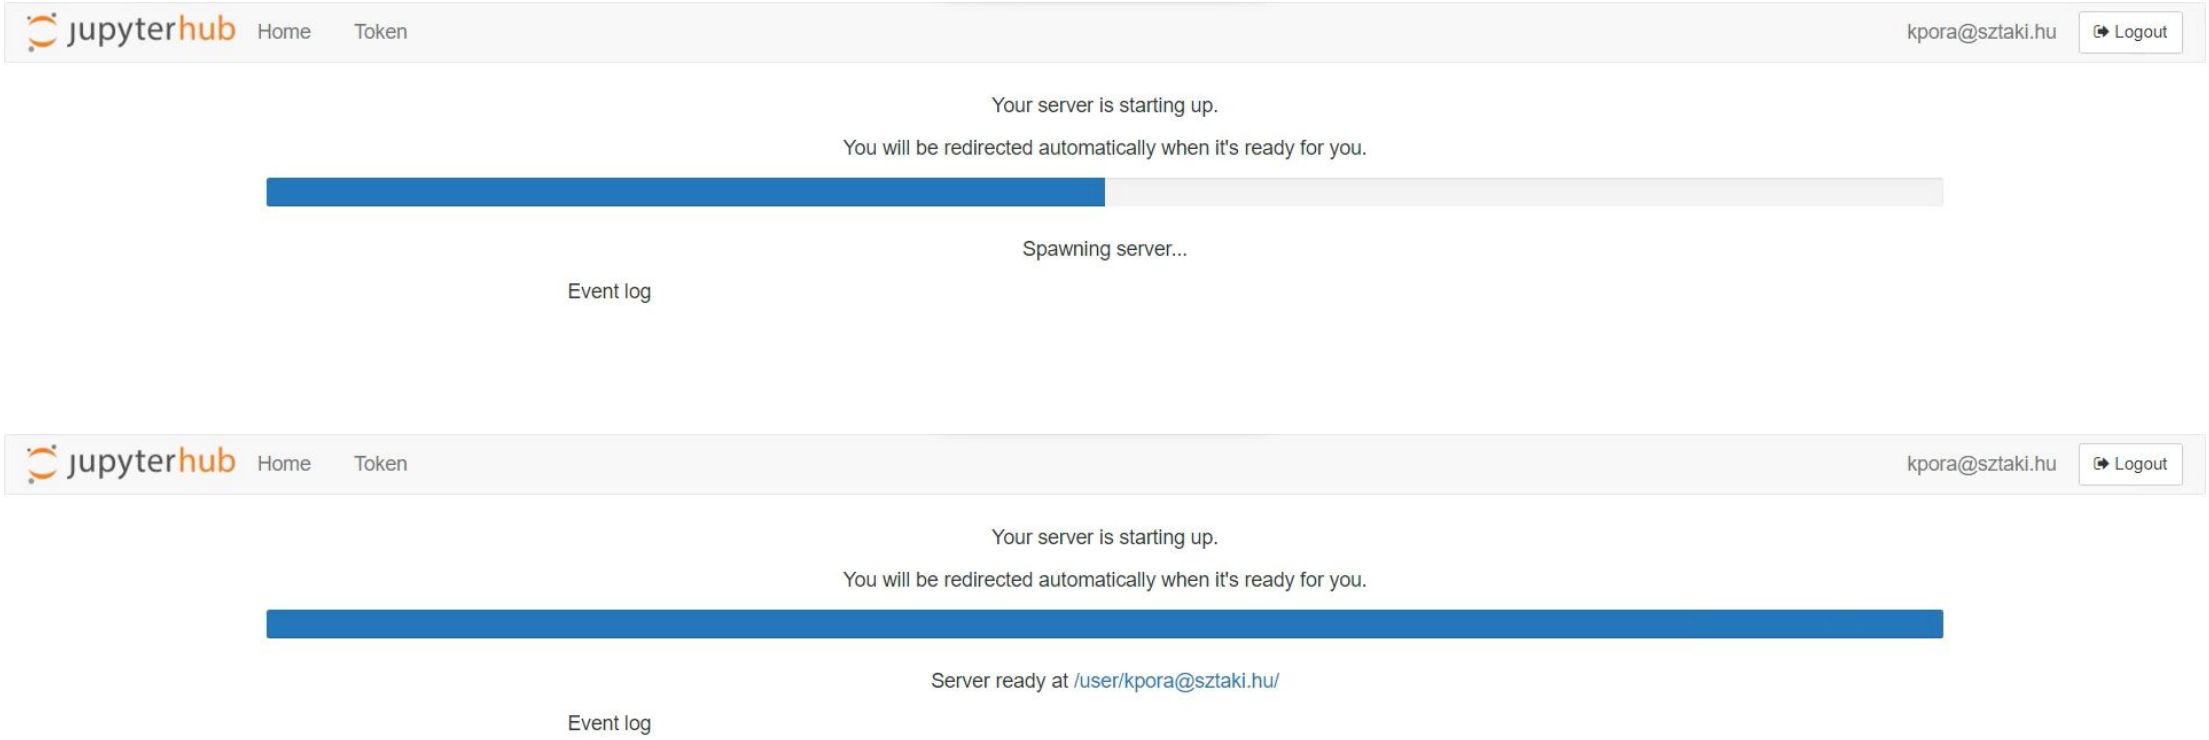

JupyterLab Environment

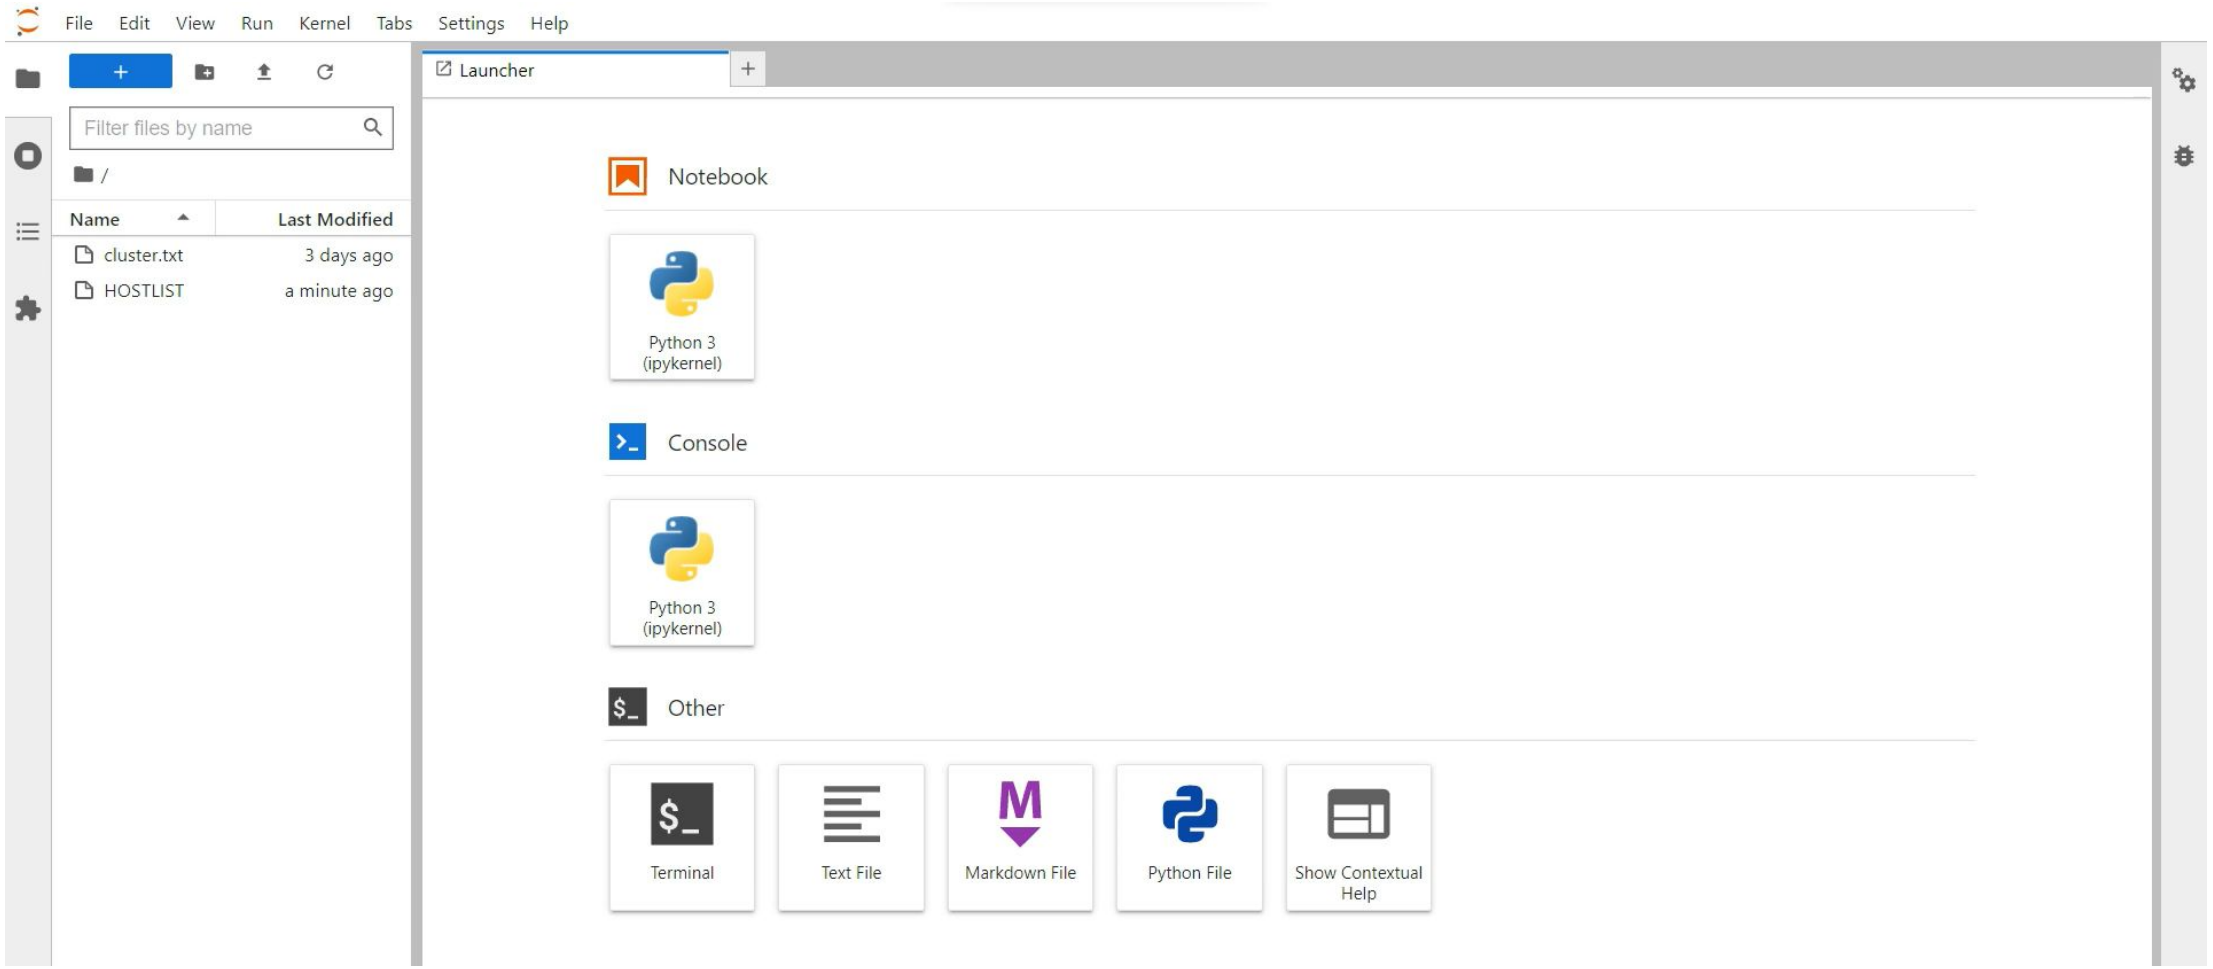

File Browser - /horovod directory

Cluster Summary

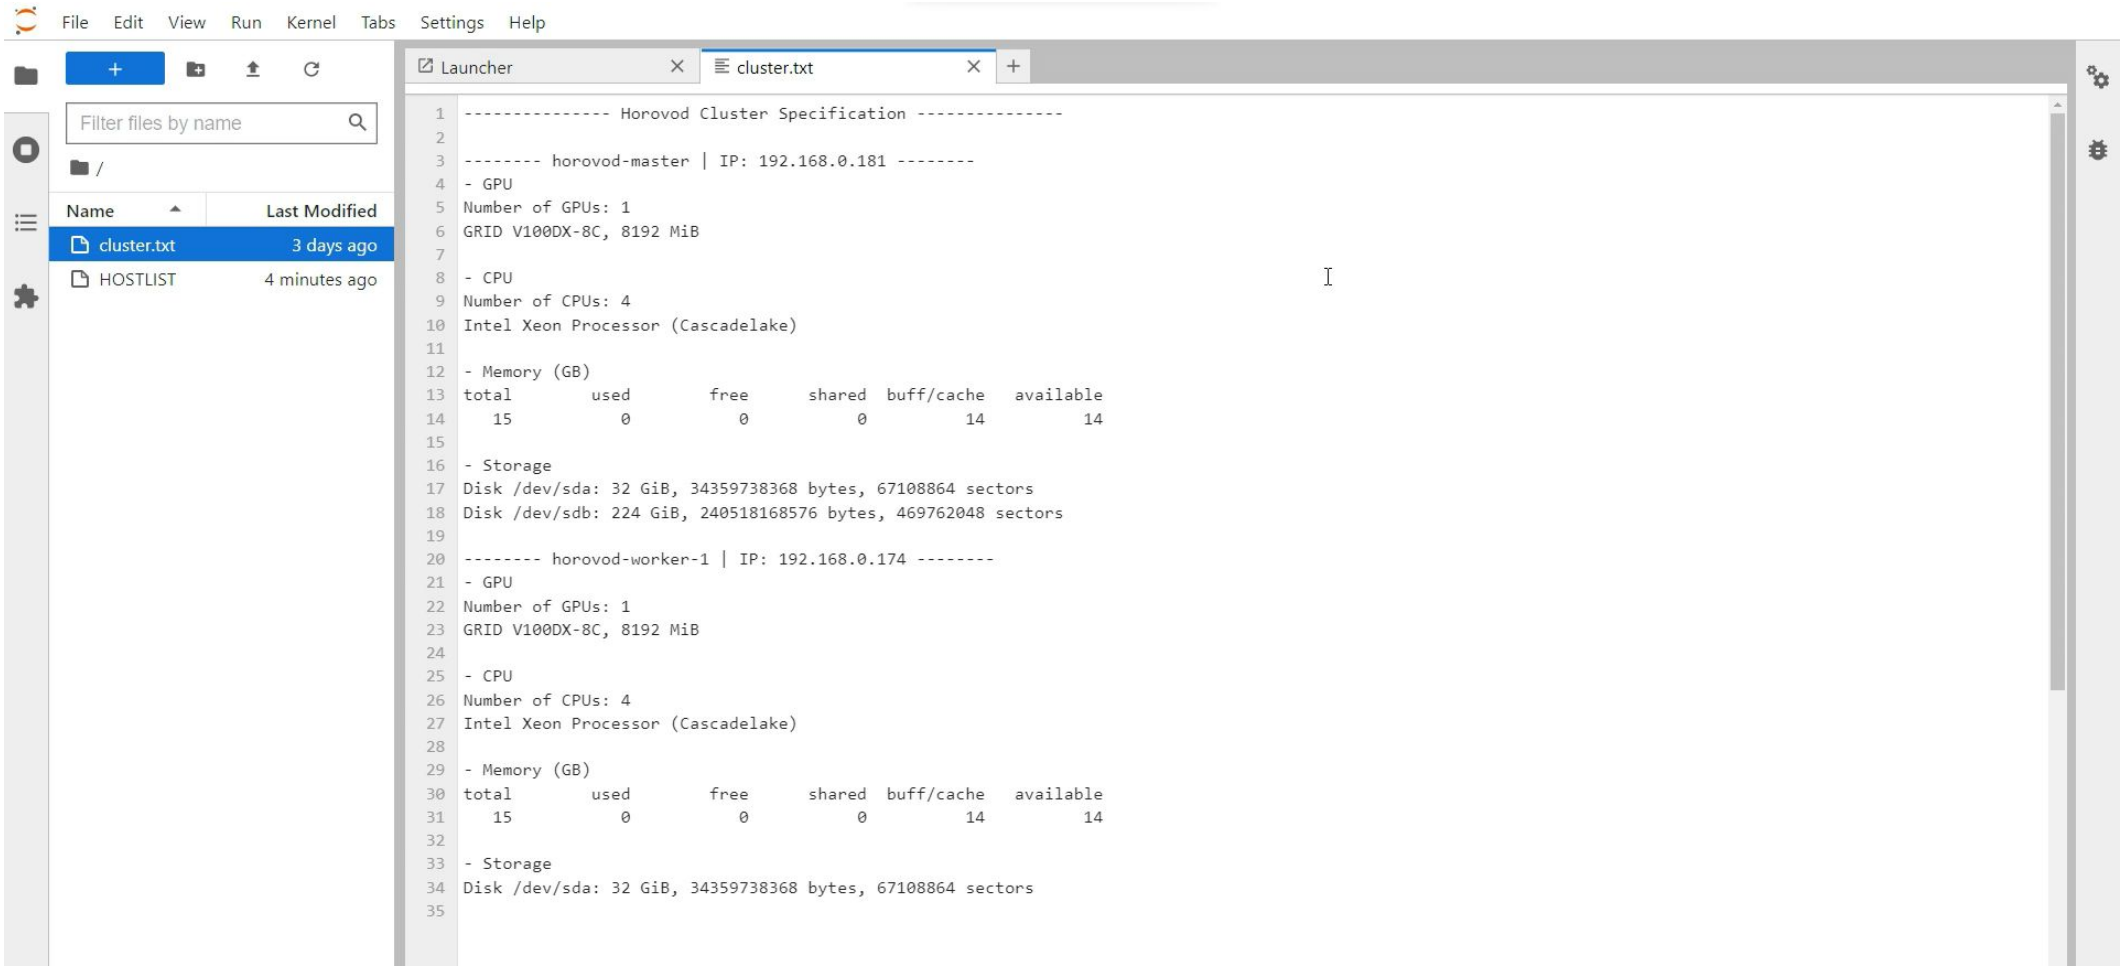

Host list for training

Example training script

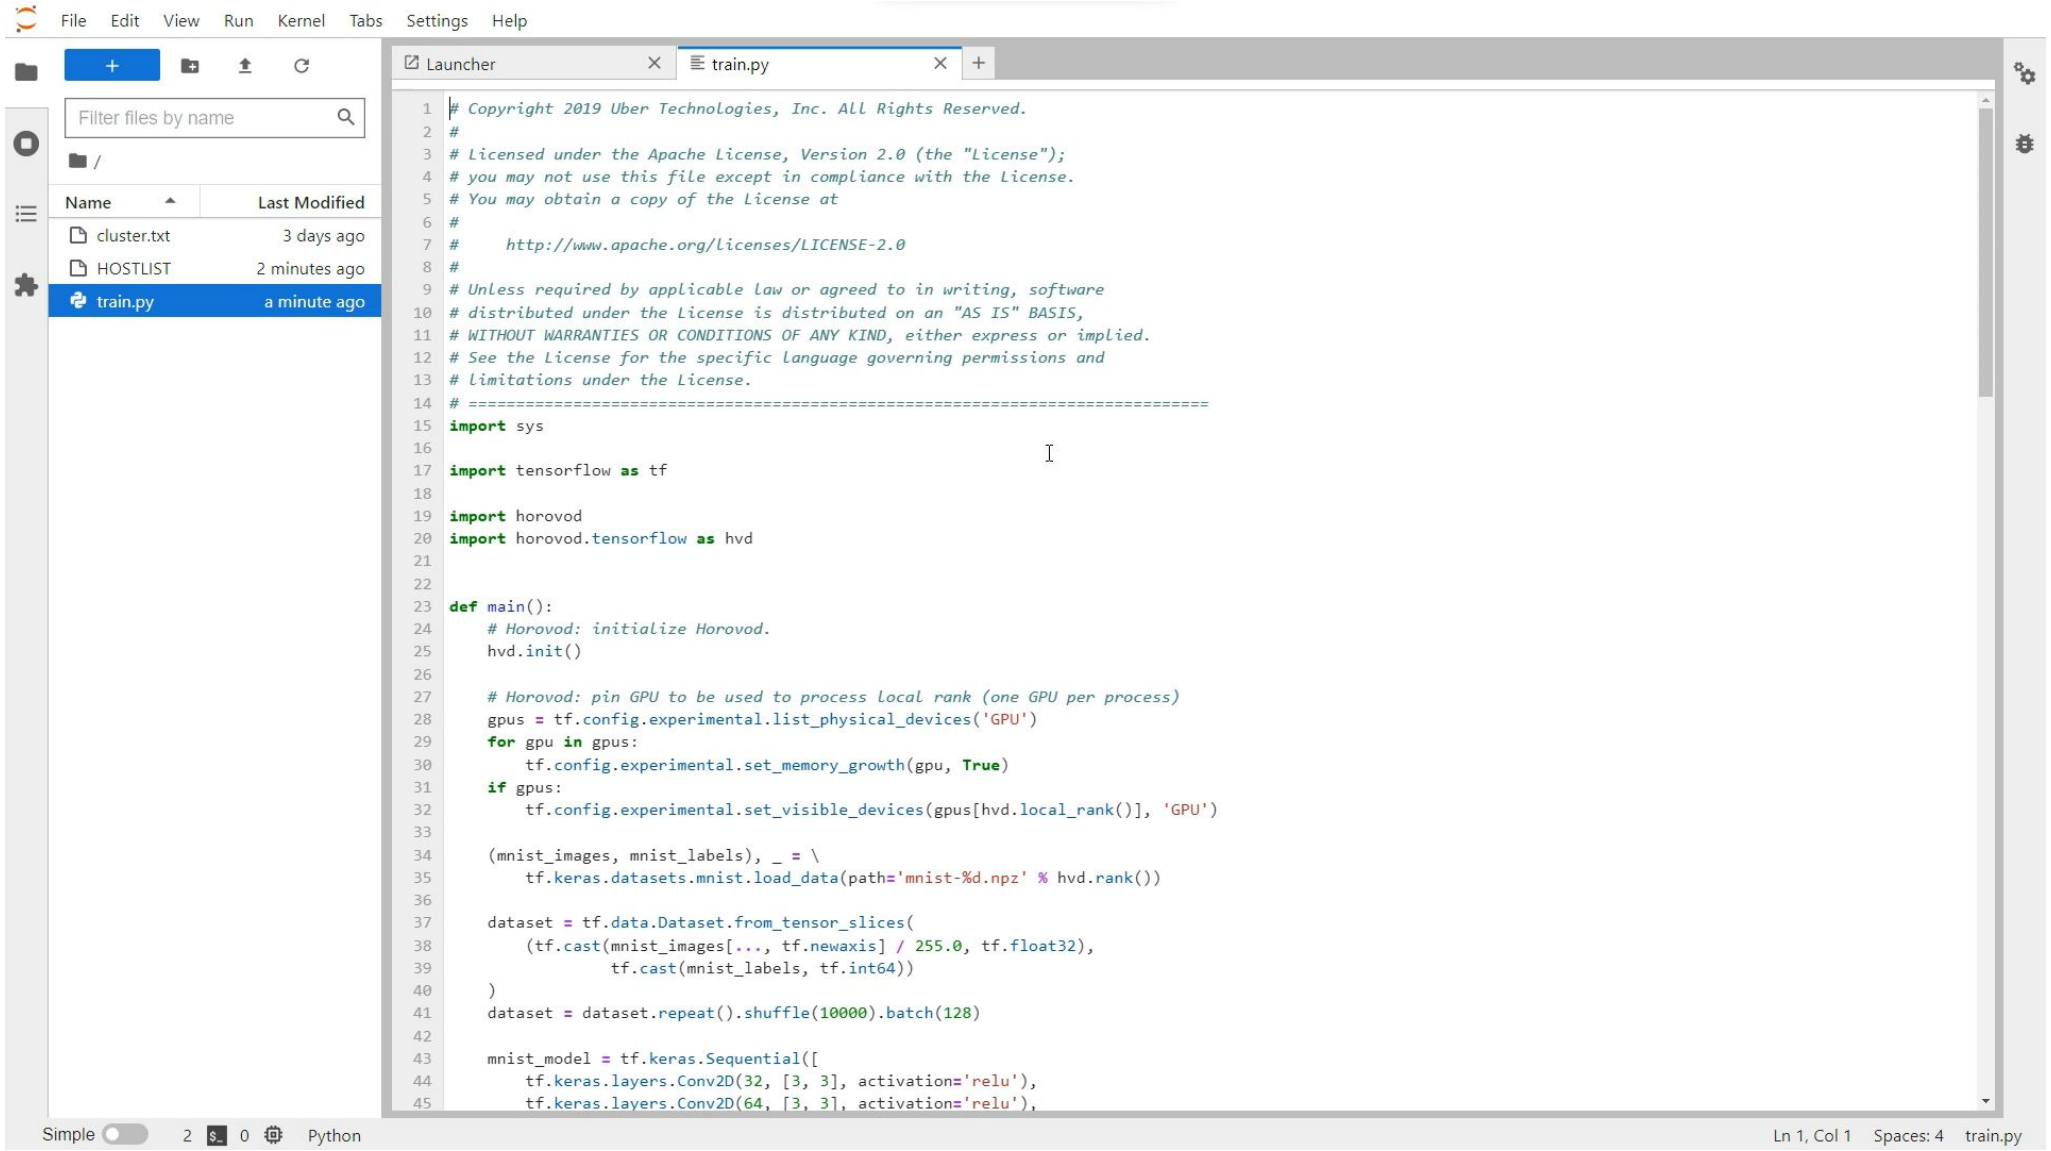

Onboarding to the EOSC Marketplace

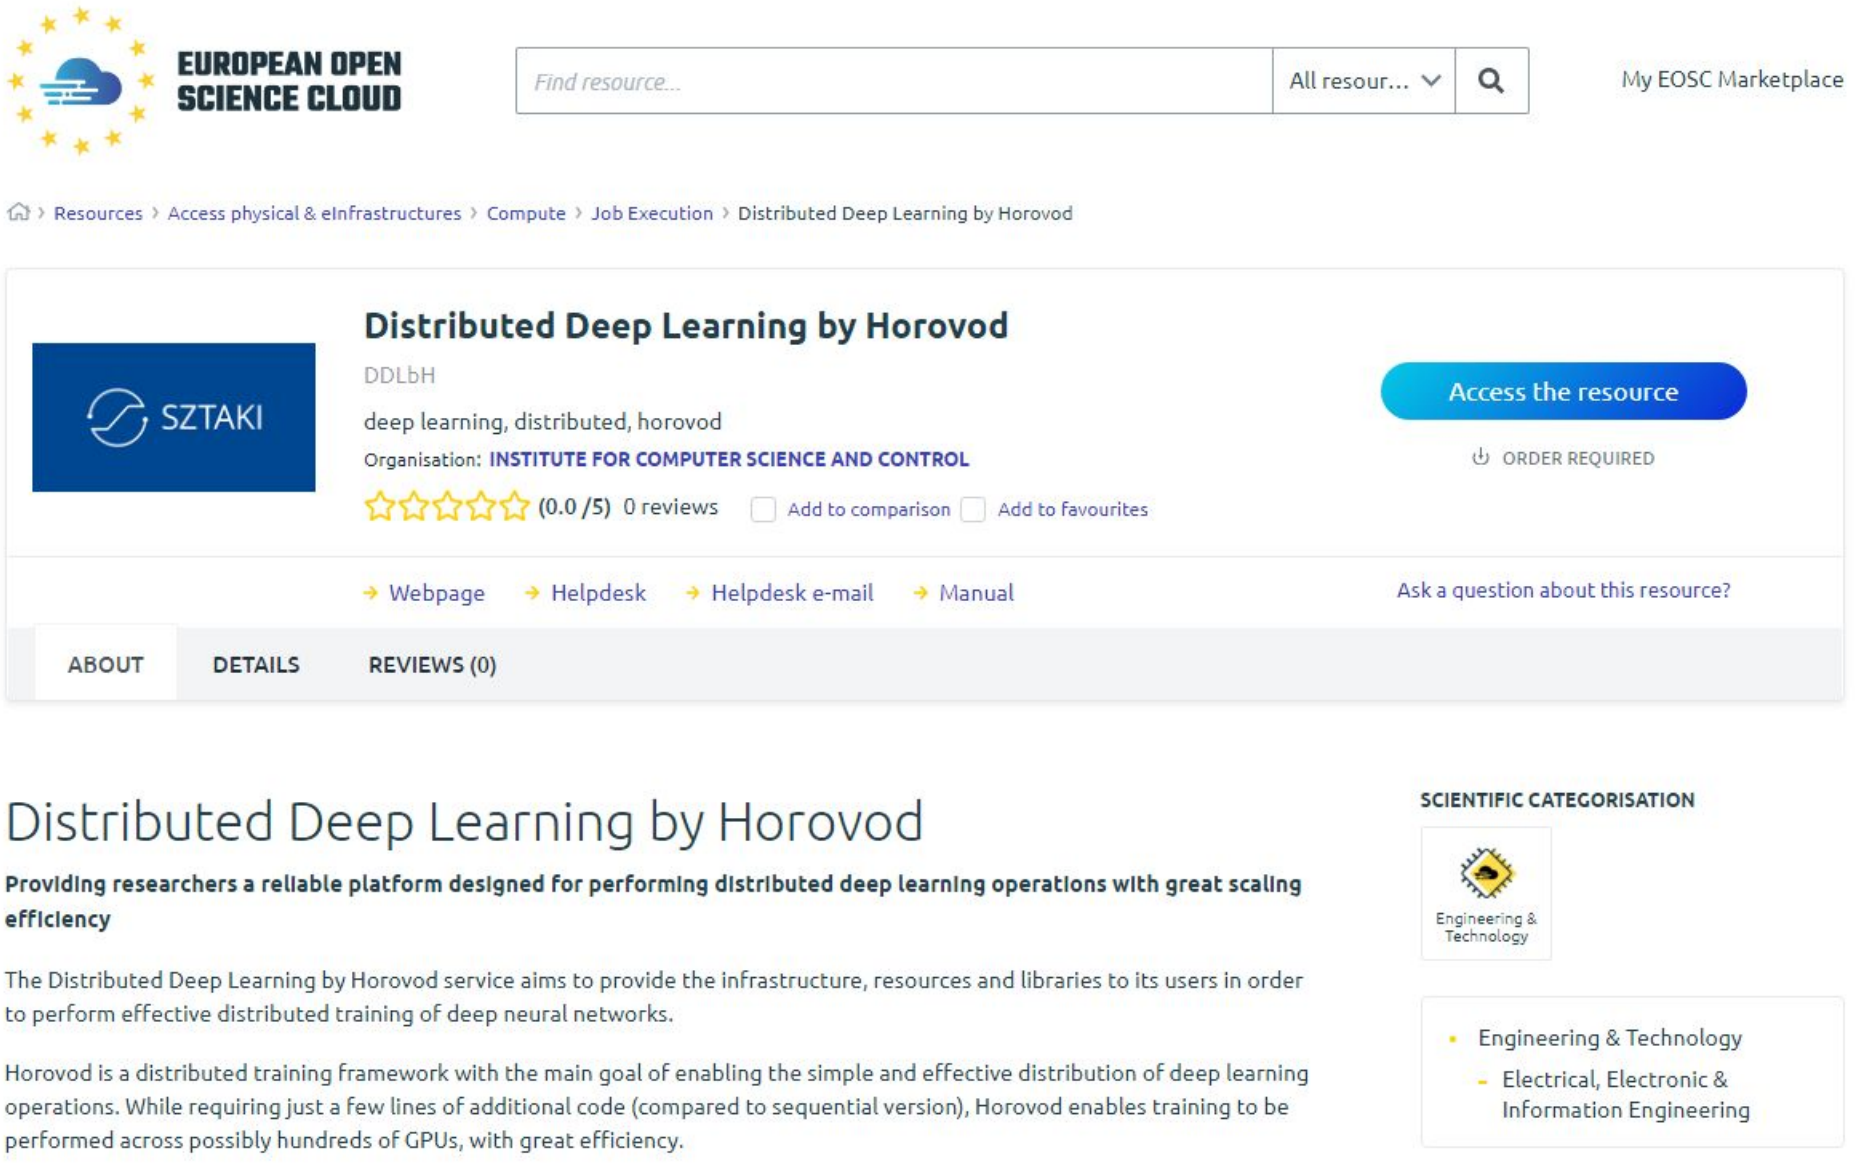

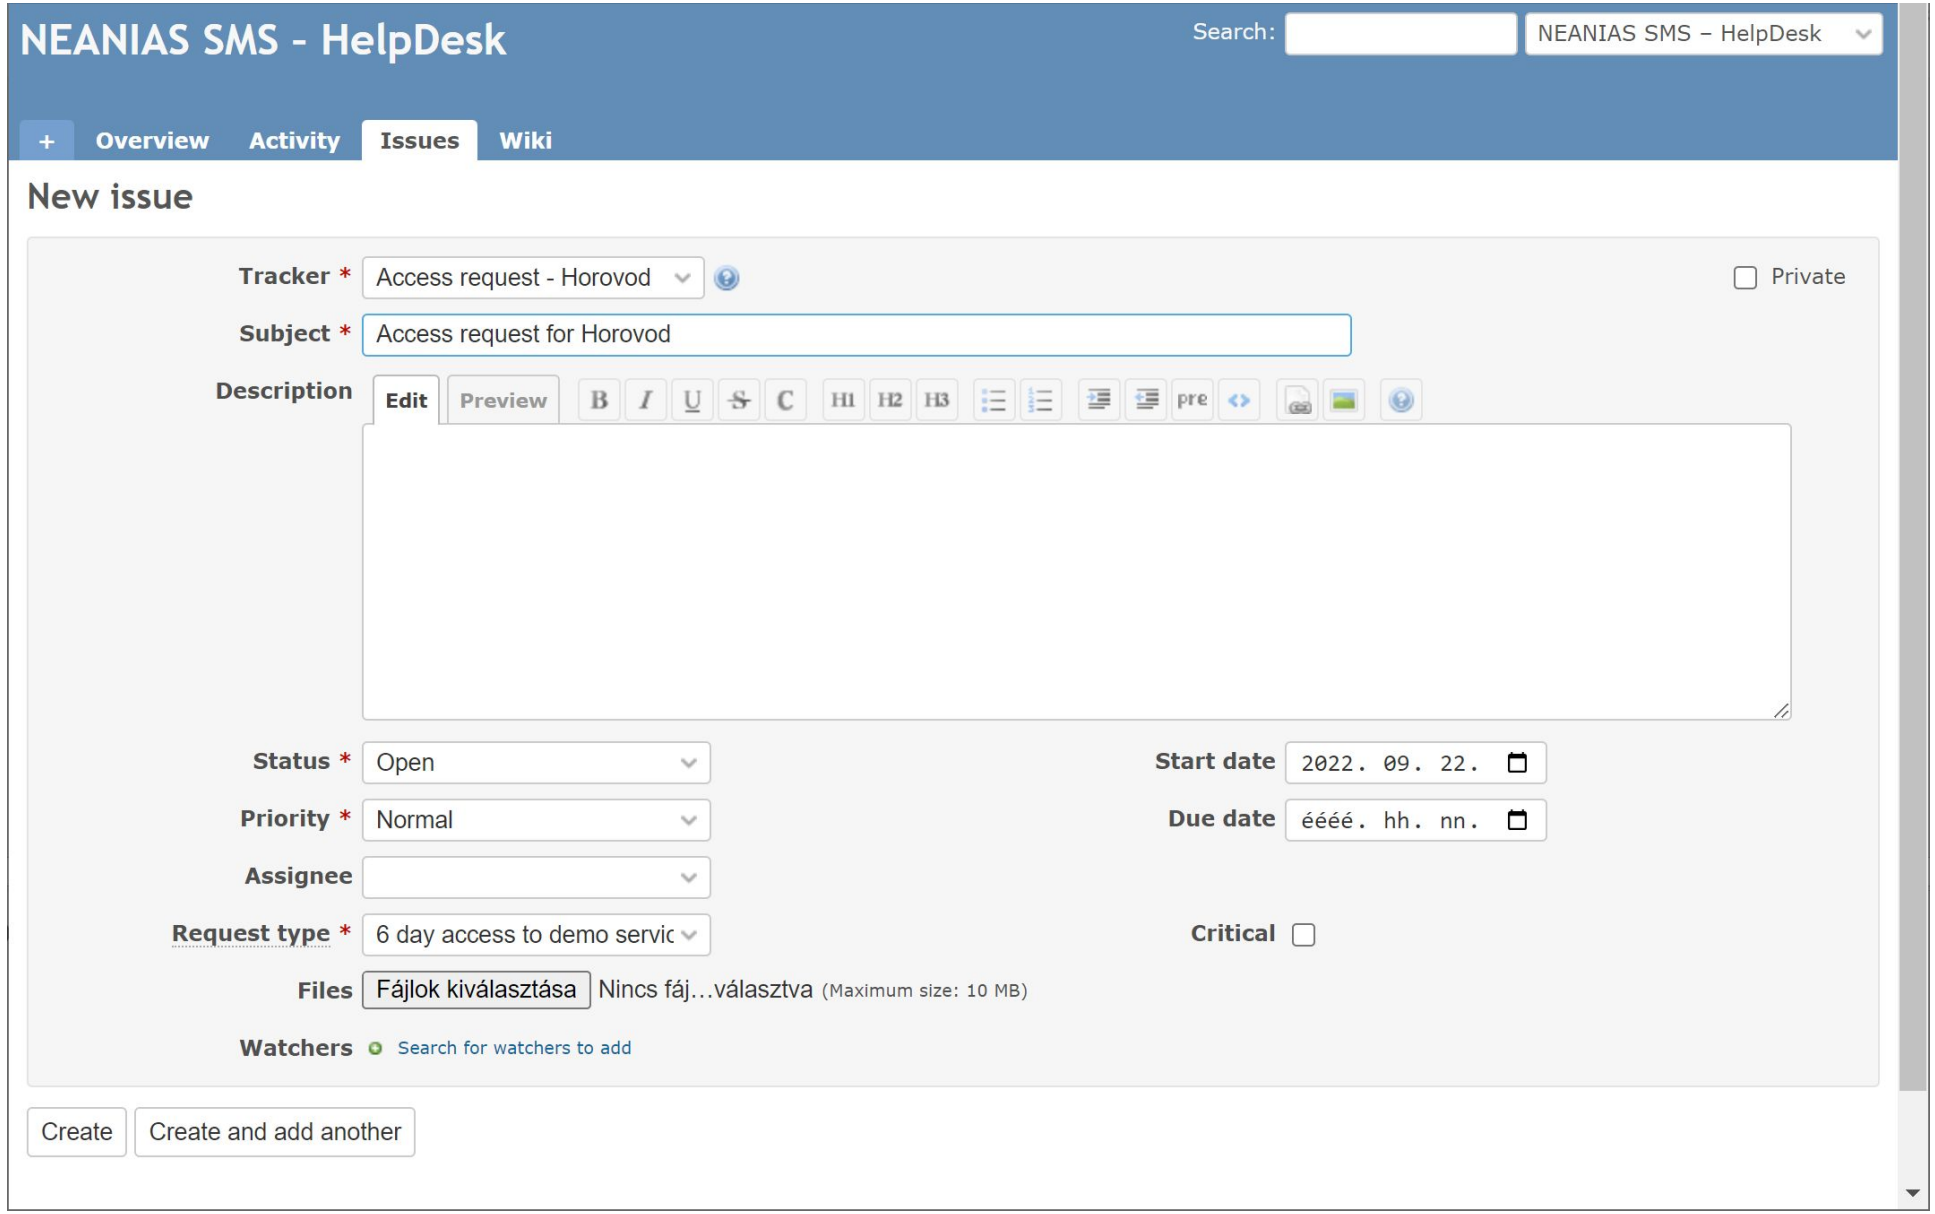
NEANIAS SMS - HelpDesk
Search:
NEANIAS SMS – HelpDesk
+
Overview
Activity
Issues
Wiki
New issue
Tracker *
Access request - Horovod
Private
Subject *
Access request for Horovod
Description
Edit
Preview
Status *
Open
Start date
2022. 09. 22.
Priority *
Normal
Due date
éééé. hh. nn.
Assignee
Request type *
6 day access to demo servic
Critical
Files
Fájlok kiválasztása
Nincs fáj…választva
(Maximum size: 10 MB)
Watchers
Search for watchers to add
Create
Create and add another

Access modes

Limited-time demo

- Gain short term access to a demo cluster
- Hosted on ELKH Cloud
- 4 GPU enabled nodes

ELKH Cloud

Request deployment on EOSC resources

- Long term access
- Exact period length and node count is up to negotiation

EUROPEAN OPEN SCIENCE CLOUD

Self-hosted

- User manual
- Technical consultation

NEANIAS

Novel EOSC Services for Emerging Atmosphere, Underwater & Space Challenges

NEANIAS receives funding from European Union under Horizon 2020 Research and Innovation Programme under grant agreement No. 863448

ELKH Eötvös Loránd Kutatási Hálózat

Thank you for your attention!

József Kovács, SZTAKI
jozsef.kovacs@sztaki.hu

Follow us:

http://www.neanias.eu

https://twitter.com/Neanias_eu

https://www.facebook.com/neanias.eu/

https://www.linkedin.com/groups/13786081/

ELKH | Eötvös Loránd Research Network