E4S: The Extreme-scale Scientific Software Stack Release 25.06

Release 25.06 notes

June 6, 2025

High Performance Software Foundation

E4S Team

E4S 25.06: What's New?

- E4S includes 130+ HPC packages on ARM, x86_64, and ppc64le platforms.
- E4S improves support for a cross-platform AI/ML software stack including packages like NVIDIA BioNeMo™, NVIDIA NeMo™, Vllm, HuggingFace CLI, TensorFlow, PyTorch, Google.generativeai (Gemini API), OpenAI (API), TorchBraid, Pandas, Scikit-Learn, JAX, OpenCV, and LBANN with support for GPUs.
- Support for new architecture: NVIDIA Blackwell (sm_120).
- Updates to language and runtime: Chapel with support for AMD and NVIDIA GPUs
- New tools: libCEED
- Applications include: CP2K, DealII, FFTX, GROMACS, LAMMPS, Nek500, Nekbone, NWChem, OpenFOAM, WarpX, WRF, Quantum Espresso, and Xyce
- E4S DocPortal updated with AI/ML tools.
- CUDA upgraded to 12.8 (aarch64, x86_64), ROCm upgraded to 6.3.3, oneAPI upgraded to 2025.1.
- Adaptive Computing's Heidi AI/HPC Cloud on demand data center (ODDC) web-based platform for multi-user, multi-node ParaTools Pro for E4S™ images on AWS, Azure, Google Cloud, and OCI Marketplace with NVIDIA GPUs with VNC based remote desktop and Torque (qsub) and SLURM (sbatch) for multi-node execution:
  - https://adaptivecomputing.com/ and https://paratoolspro.com

PESO

E4S: Extreme-scale Scientific Software Stack

https://e4s.io

About E4S

- E4S is an **ecosystem for science** and a community effort to provide open-source software packages for developing, deploying and running scientific applications on HPC platforms.
- E4S has built a comprehensive, coherent software stack that enables application developers to productively develop highly parallel applications that effectively target diverse exascale architectures.
- E4S provides a curated, Spack based software distribution of 130+ HPC (OpenFOAM, Gromacs, Nek5000, LAMMPS), EDA (e.g., Xyce), and AI/ML packages (e.g., NVIDIA NeMo™, NVIDIA BioNeMo™, Vllm, HuggingFace CLI, TensorFlow, PyTorch, OpenCV, TorchBraid, Scikit-Learn, Pandas, JAX, LBANN with support for GPUs).
- Base images and full featured containers (with GPU support) and DOE LLVM containers.
- Commercial support for E4S through ParaTools, Inc. for installation, maintaining an issue tracker, and ECP AD engagement.
- E4S for commercial clouds: Adaptive Computing's ODDC with ParaTools Pro for E4S™ image for **AWS, GCP, Azure, OCI.**
- With E4S Spack binary build caches, E4S supports both bare-metal and containerized deployment for GPU based platforms.
  - x86_64, ppc64le (IBM Power 10), aarch64 (ARM64) with support for CPUs and GPUs from NVIDIA, AMD, and Intel
  - Container images on DockerHub and E4S website of pre-built binaries of ECP ST products.
- e4s-chain-spack.sh to chain two Spack instances allows us to install new packages in home directory and use other tools.
- e4s-cl container launch tool allows binary distribution of applications by swapping MPI in the containerized app w/ system MPI.
- e4s-alc is an à la carte tool to customize container images by adding system and Spack packages to an existing image.
- E4S 25.06 released on June 6, 2025: https://e4s.io/talks/E4S_25.06.pdf

PESO

Argonne NATIONAL LABORATORY · BERKELEY LAB · Brookhaven National Laboratory · kitware · Lawrence Livermore National Laboratory · Los Alamos NATIONAL LABORATORY · OAK RIDGE National Laboratory · Pacific Northwest · ParaTools · Sandia National Laboratories · SUSTAINABLE HORIZONS INSTITUTE

E4S Download from https://e4s.io

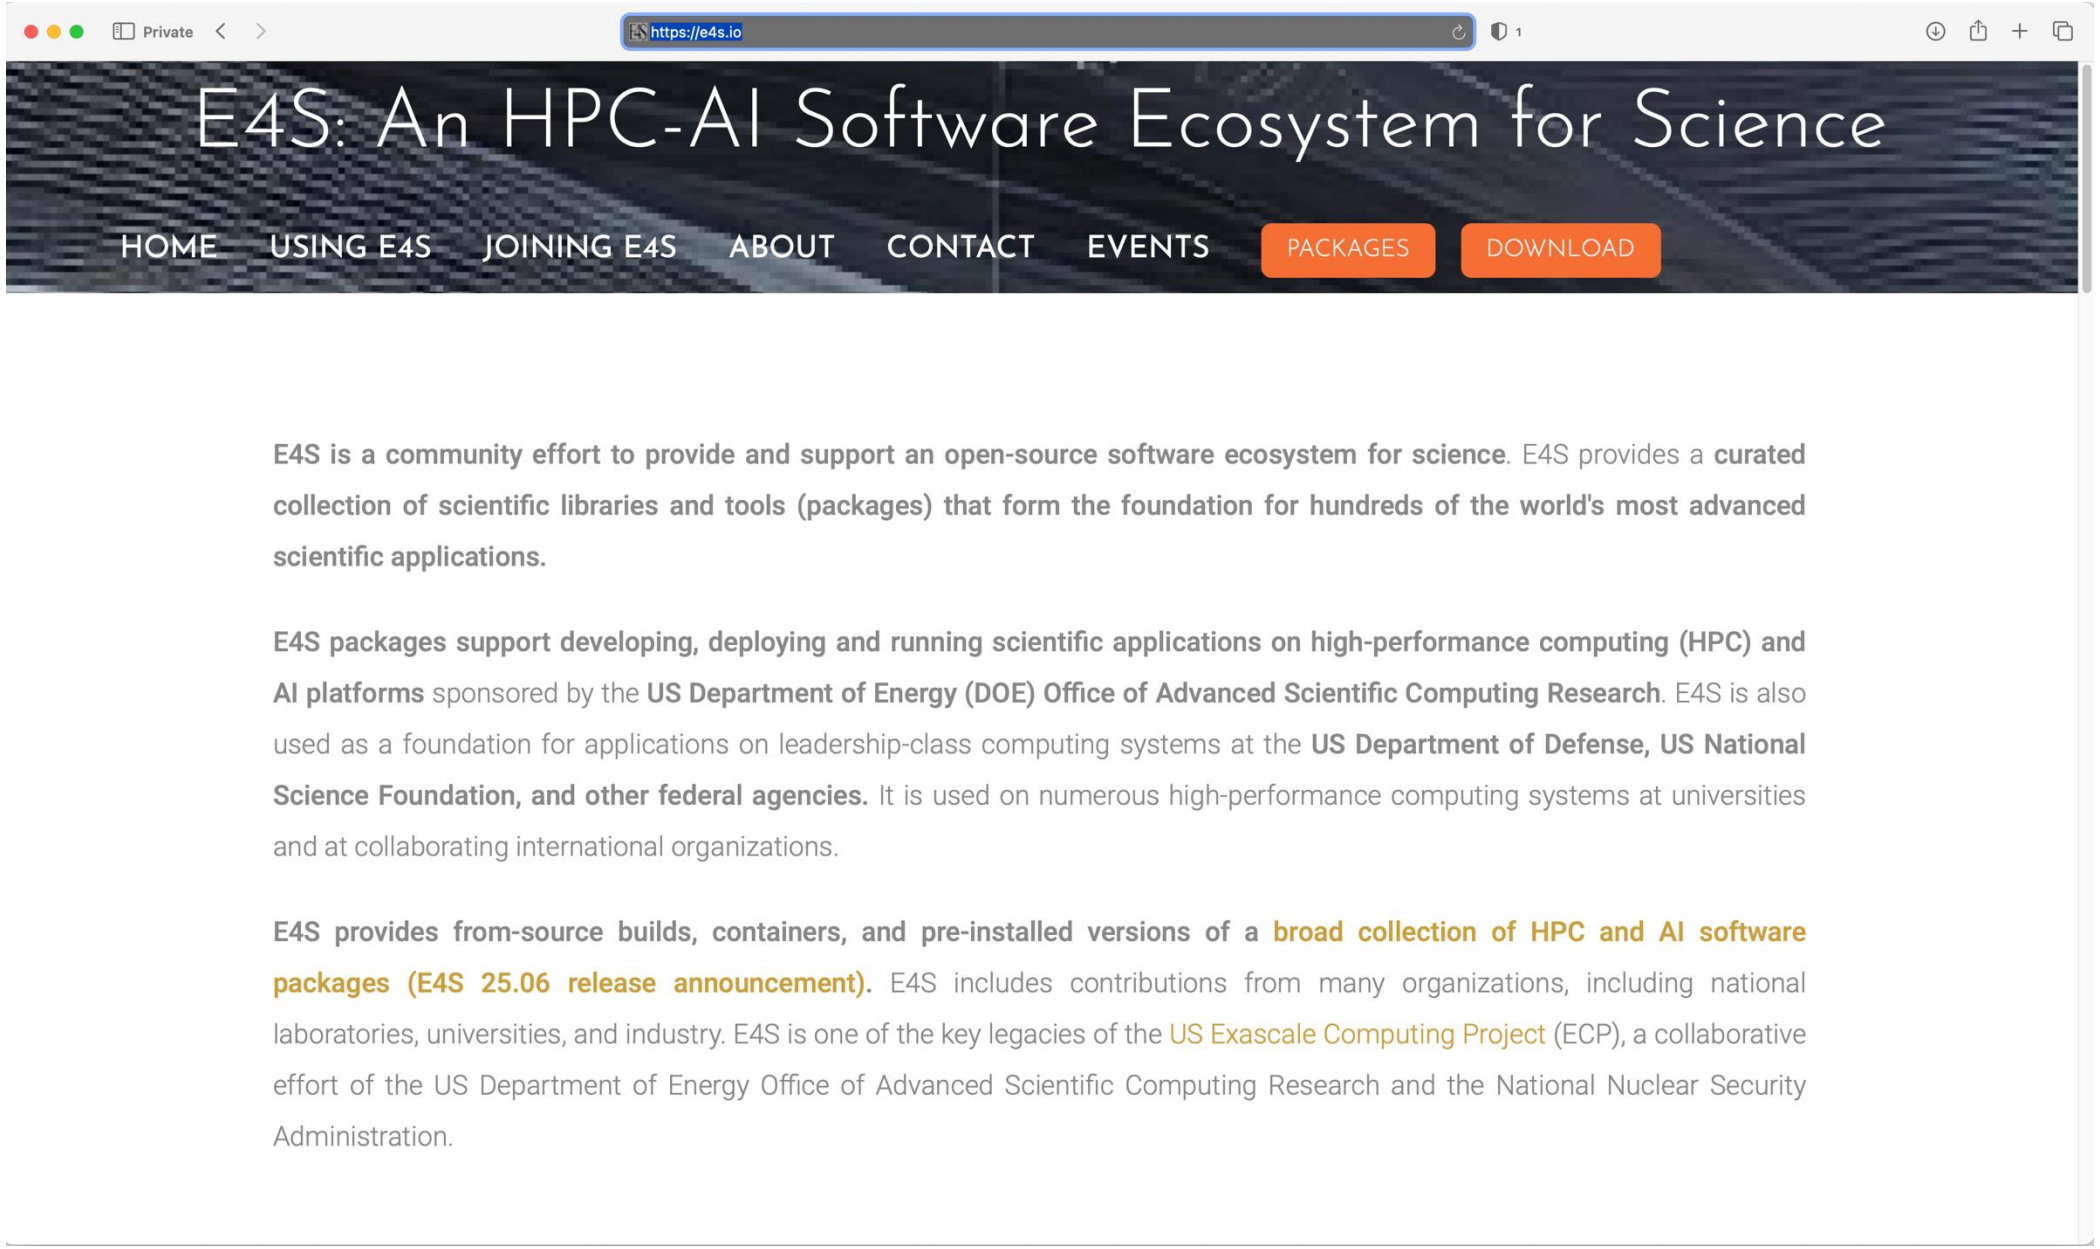

PESO

E4S Container Download from https://e4s.io

- Separate full featured Singularity images for 3 GPU architectures
- GPU full featured images for
  - x86_64 (Intel, AMD, NVIDIA)
  - ppc64le (NVIDIA)
  - aarch64 (NVIDIA)
- Full featured images available on Dockerhub
- 130+ products on 3 architectures

PESO

Download E4S 25.06 GPU Container Images: AMD, Intel, and NVIDIA

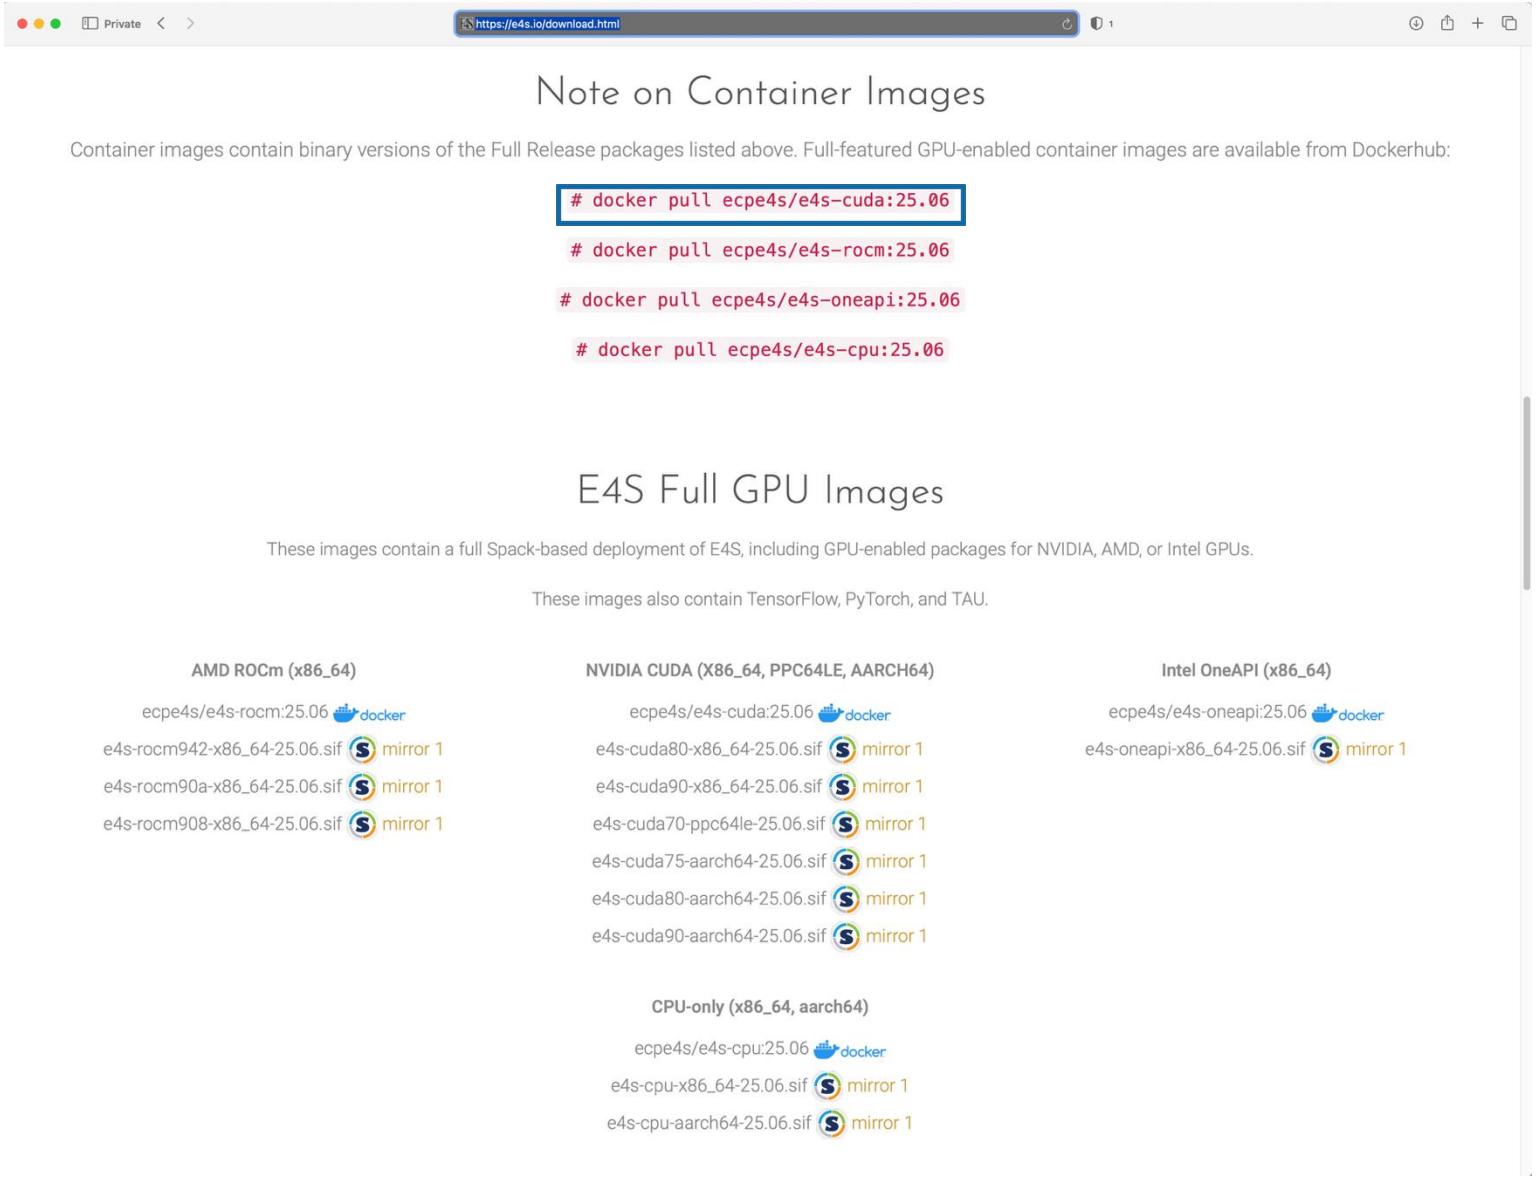

https://e4s.io

E4S base container images allow users to customize their containers

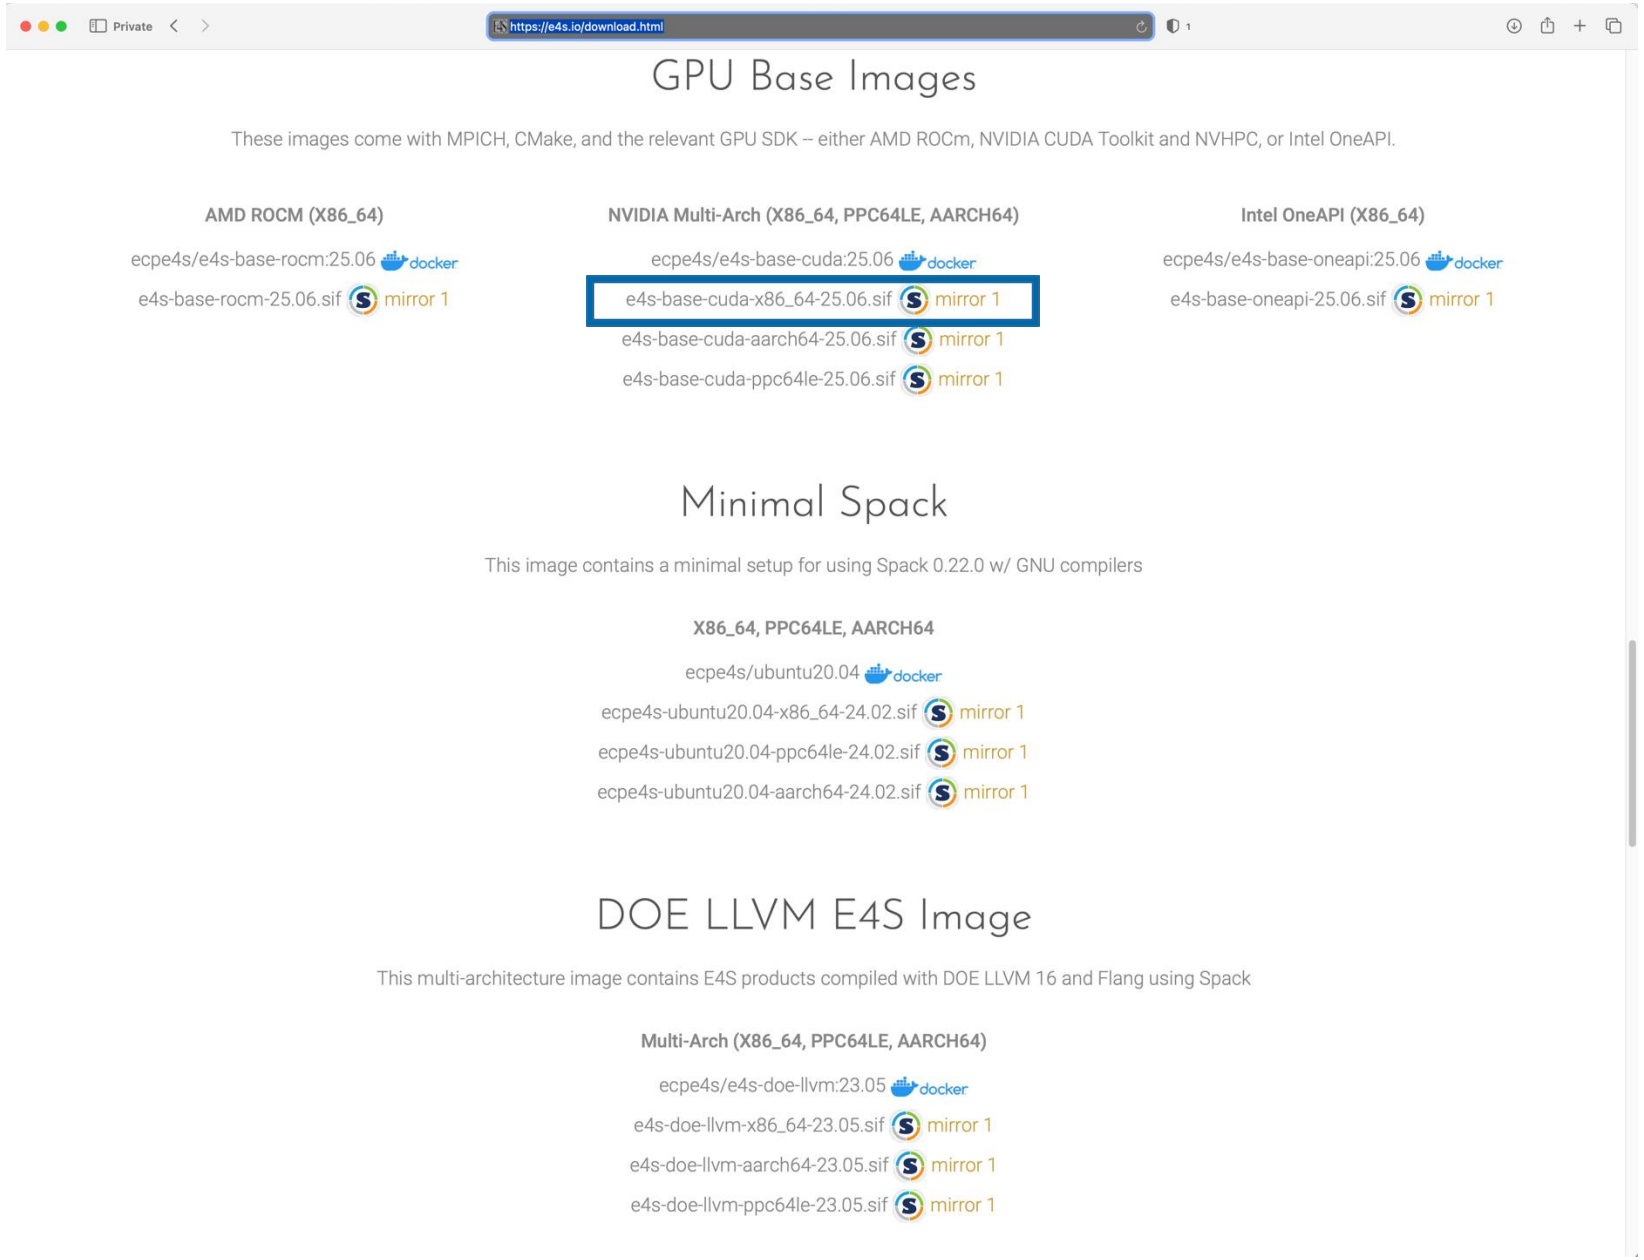

- Intel oneAPI
- AMD ROCm
- NVIDIA CUDA

E4S Application Specific CI and minimal CI images

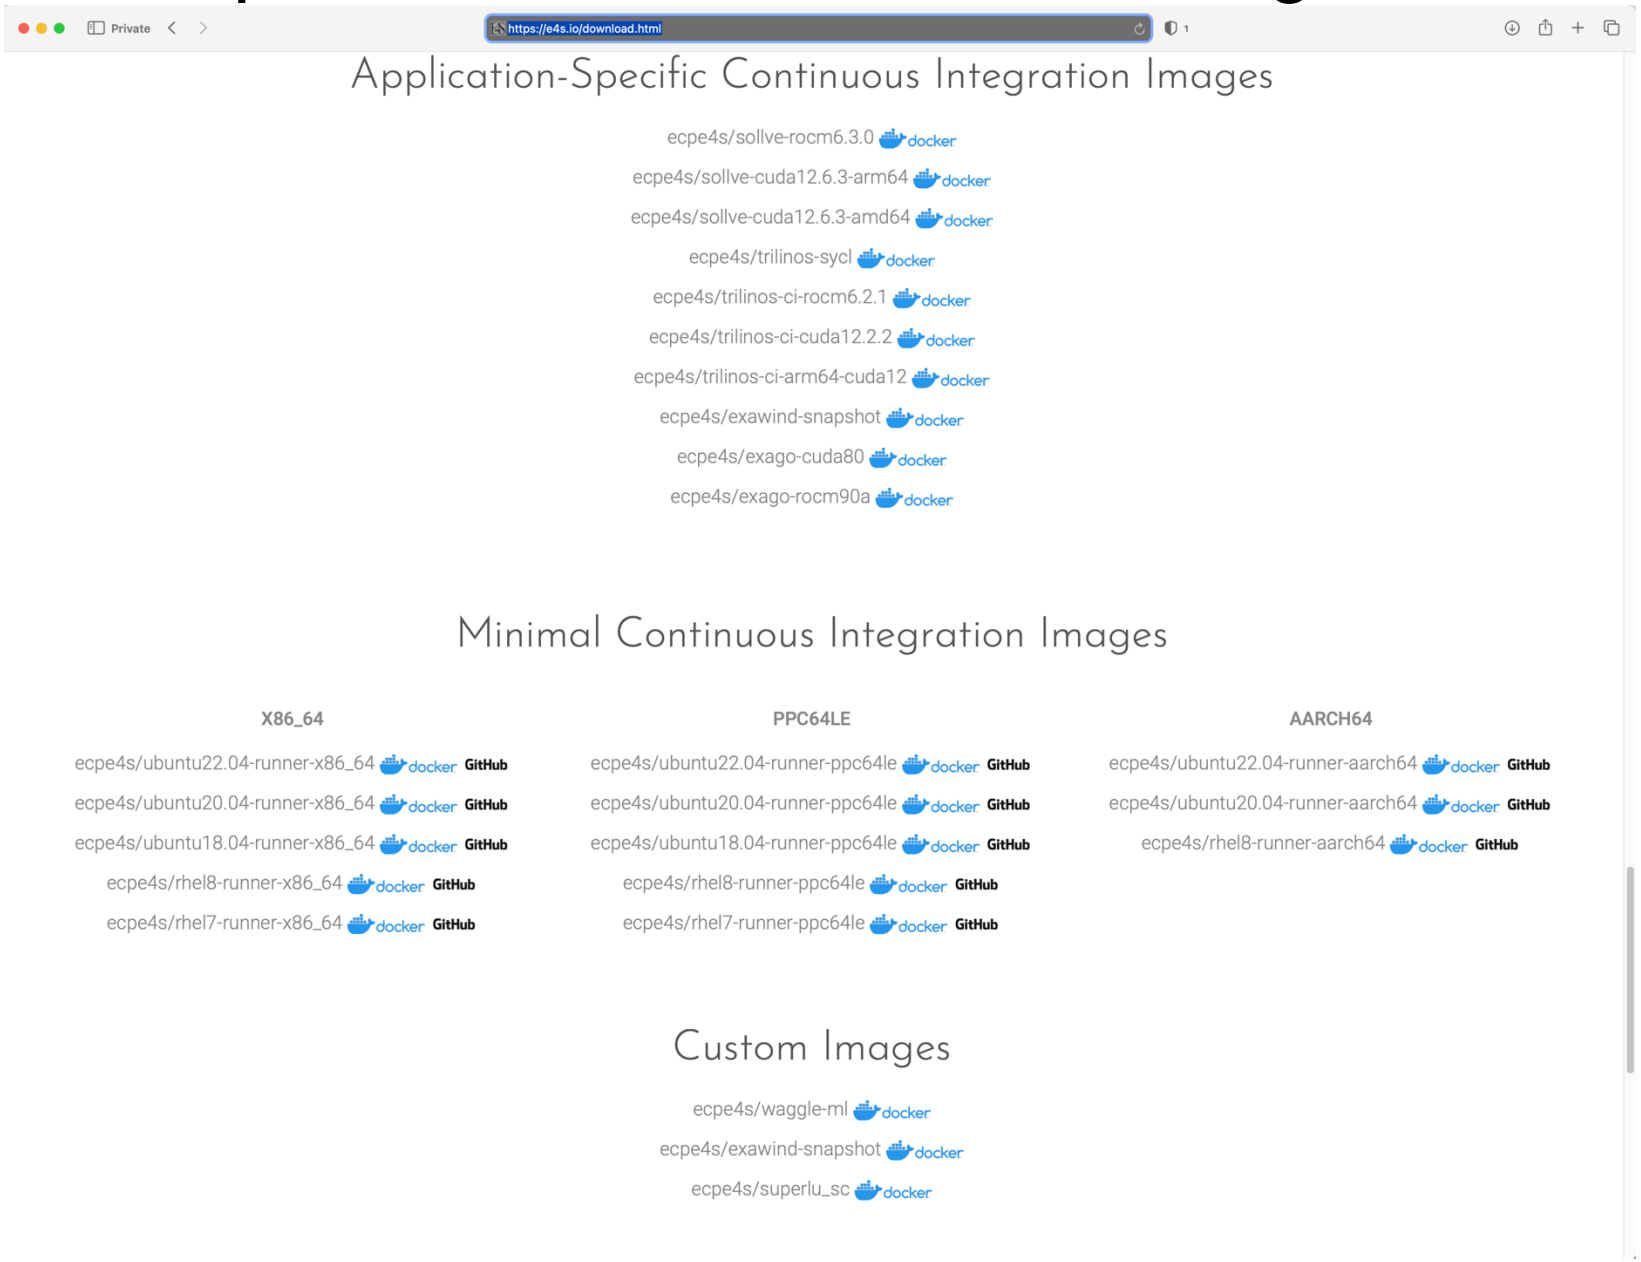

E4S Facility Deployment and AWS EC2 Image

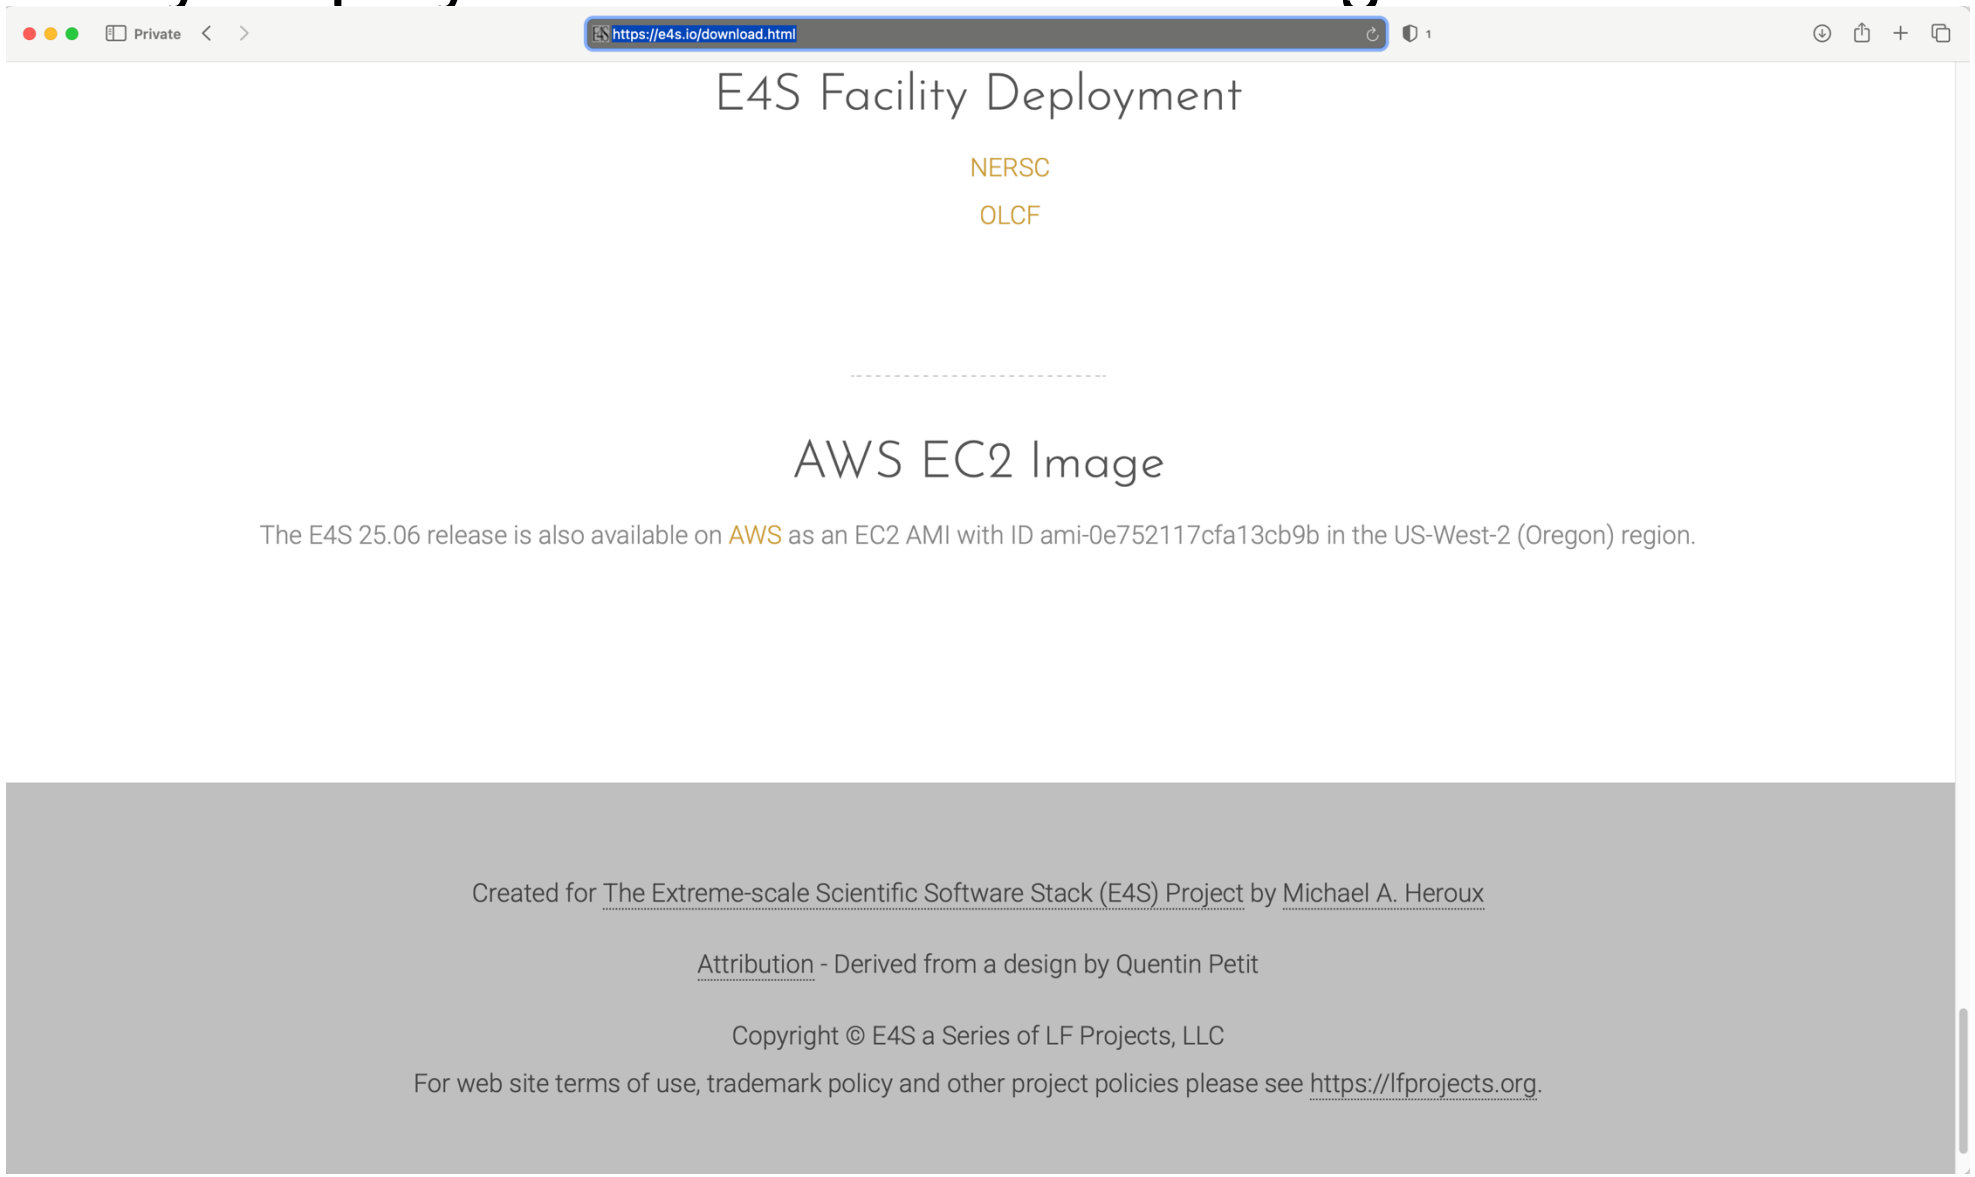

E4S 25.06 Detailed Documentation for Bare-metal Installation

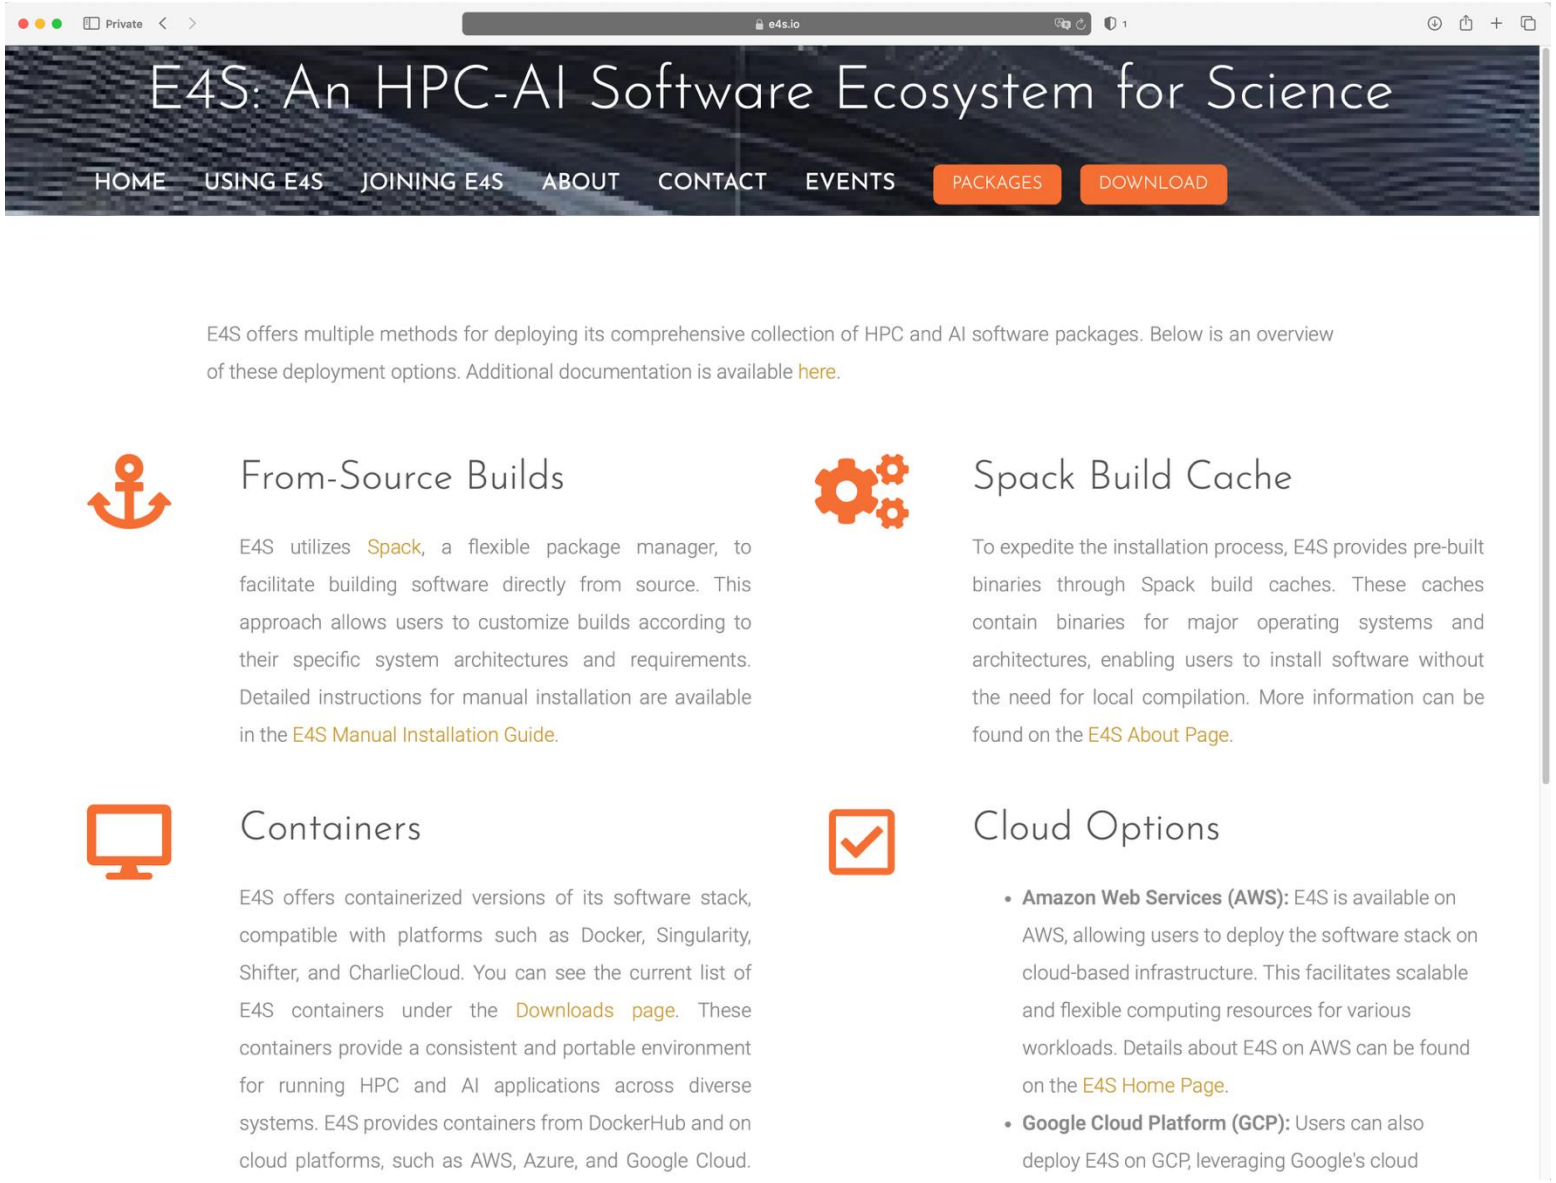

https://e4s.io

E4S DocPortal updated with new applications and AI/ML tools

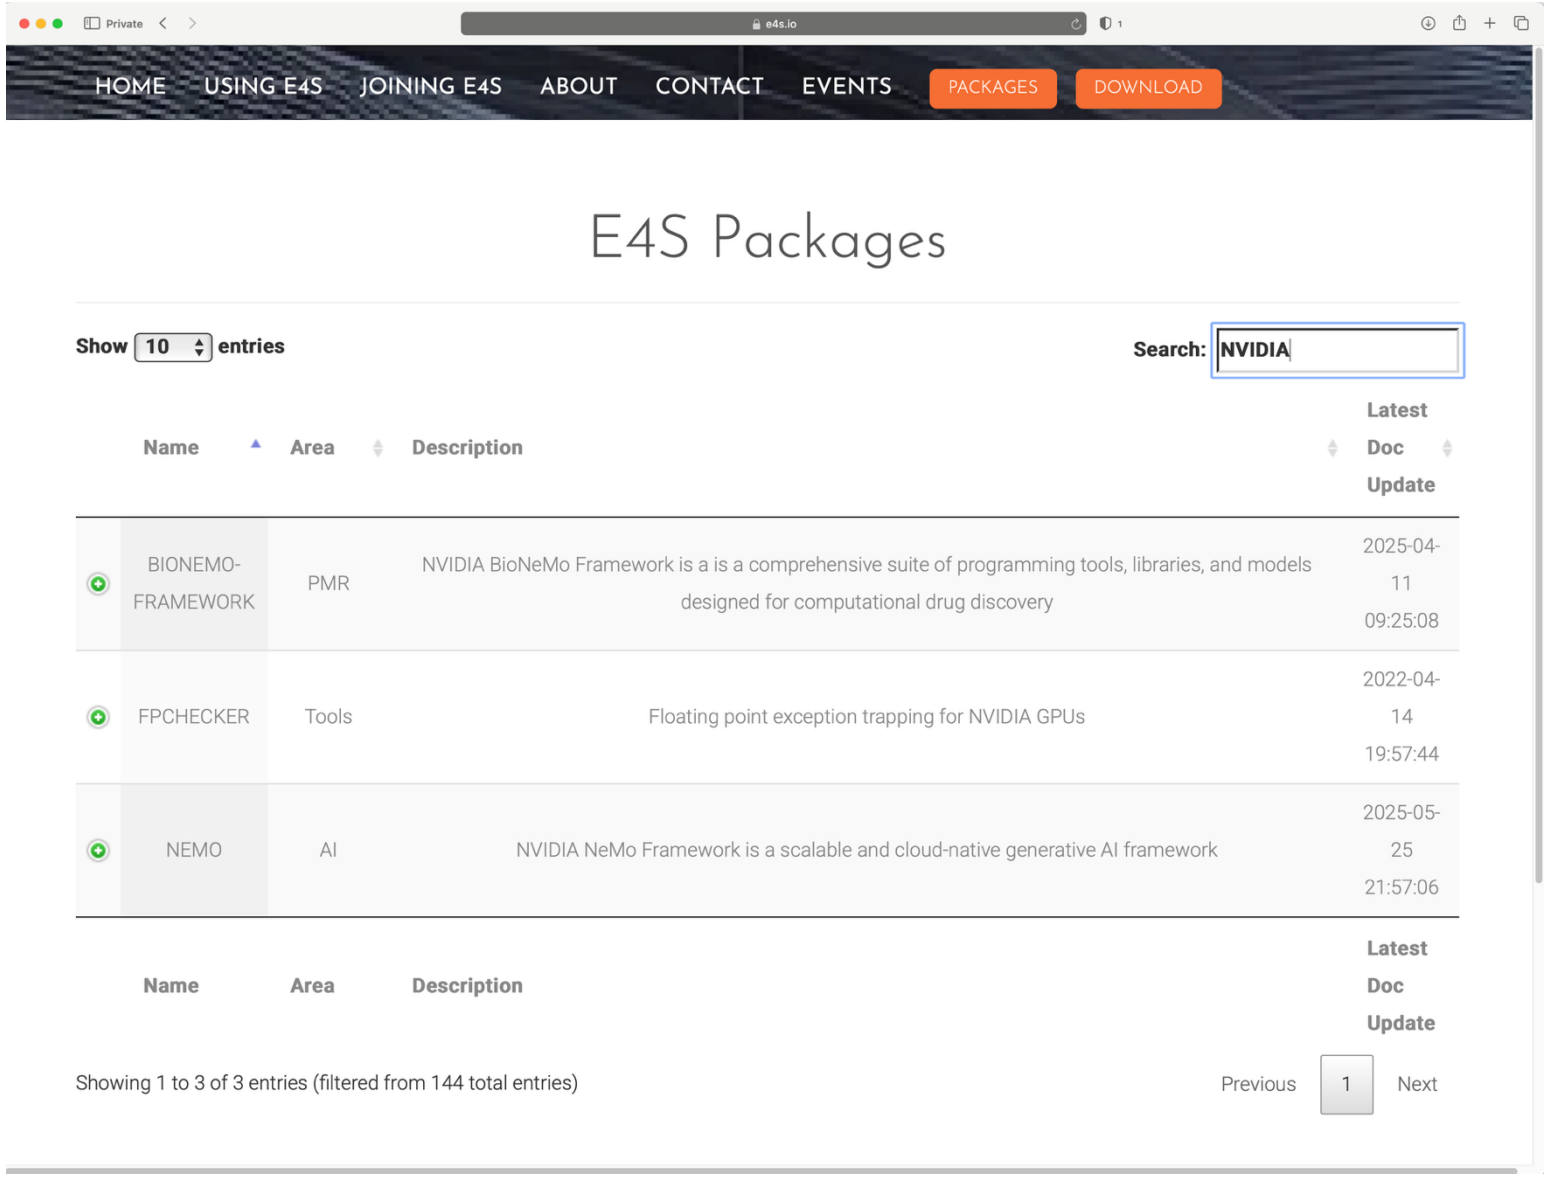

https://e4s.io

E4S Tools: e4s-chain-spack.sh to customize software stack

```
sameer@mothra:~$ ls ~/images
e4s-cuda80-x86_64-25.06.sif
sameer@mothra:~$ singularity run --nv ~/images/e4s-cuda80-x86_64-25.06.sif
Singularity> /etc/e4s/e4s-chain-spack.sh ~/spack
Cloning into '/home/sameer/spack'...
remote: Enumerating objects: 686113, done.
remote: Counting objects: 100% (976/976), done.
remote: Compressing objects: 100% (463/463), done.
remote: Total 686113 (delta 772), reused 518 (delta 510), pack-reused 685137 (from 3)
Receiving objects: 100% (686113/686113), 230.82 MiB | 37.06 MiB/s, done.
Resolving deltas: 100% (326280/326280), done.

------
Configuration SUCCESS!

Downstream: /home/sameer/spack
Upstream: /spack

To use the downstream Spack instance, run the following command in your shell:
. /home/sameer/spack/share/spack/setup-env.sh
------

Singularity> . /home/sameer/spack/share/spack/setup-env.sh
Singularity> spack find valgrind
==> Error: No package matches the query: valgrind
Singularity> spack install valgrind
[+] /usr/local/mpich/install/mpich (external mpich-4.2.3-47excoypwhfmhx57rfs6reouvninugcf)
[+] /usr (external glibc-2.35-a7drdl4tlx4bu3mzhor75pskvd3pdot6)
[+] /spack/opt/spack/linux-ubuntu22.04-x86_64_v3/gcc-11.4.0/gcc-runtime-11.4.0-f63c77kavzjtpmnhucd2oyfaxagwjzla
[+] /spack/opt/spack/linux-ubuntu22.04-x86_64_v3/gcc-11.4.0/boost-1.86.0-6qkv24gbidwxhllgah6jrkym5ev2cng5
[+] /spack/opt/spack/linux-ubuntu22.04-x86_64_v3/gcc-11.4.0/gmake-4.4.1-qp5blvcyuzgzhqsrp2ew6gq2nlos34b2
==> Installing valgrind-3.23.0-feuxx36lsqp7quzmhmo4opbiadwpsars [6/6]
==> No binary for valgrind-3.23.0-feuxx36lsqp7quzmhmo4opbiadwpsars found: installing from source
==> Fetching https://mirror.spack.io/_source-cache/archive/c5/c5c34a3380457b9b75606df890102e7df2c702b9420c2ebef9540f8b5d56264d.tar.bz2
==> Ran patch() for valgrind
==> valgrind: Executing phase: 'autoreconf'
==> valgrind: Executing phase: 'configure'
==> valgrind: Executing phase: 'build'
==> valgrind: Executing phase: 'install'
==> valgrind: Successfully installed valgrind-3.23.0-feuxx36lsqp7quzmhmo4opbiadwpsars
  Stage: 3.78s.  Autoreconf: 0.01s.  Configure: 48.56s.  Build: 37.71s.  Install: 2.97s.  Post-install: 0.60s.  Total: 1m 33.97s
[+] /home/sameer/spack/opt/spack/linux-ubuntu22.04-x86_64_v3/gcc-11.4.0/valgrind-3.23.0-feuxx36lsqp7quzmhmo4opbiadwpsars
Singularity> spack load valgrind
Singularity> which valgrind
/home/sameer/spack/opt/spack/linux-ubuntu22.04-x86_64_v3/gcc-11.4.0/valgrind-3.23.0-feuxx36lsqp7quzmhmo4opbiadwpsars/bin/valgrind
```

Specify location of downstream Spack installation directory

Source downstream Spack's setup-env.sh

Install a new Spack package in downstream Spack directory

Load new package (valgrind) using spack load

PESO

https://e4s.io

E4S Tools: e4s-chain-spack.sh to customize software stack

```
Singularity> which valgrind
/home/sameer/spack/opt/spack/linux-ubuntu22.04-x86_64_v3/gcc-11.4.0/valgrind-3.23.0-feuxx36lsqp7quzmhmo4opbiadwpsars/bin/valgrind
Singularity> valgrind --help
usage: valgrind [options] prog-and-args

  tool-selection option, with default in [ ]:
    --tool=<name>             use the Valgrind tool named <name> [memcheck]
                              available tools are:
                              memcheck cachegrind callgrind helgrind drd
                              massif dhat lackey none exp-bbv

  basic user options for all Valgrind tools, with defaults in [ ]:
    -h --help                 show this message
    --help-debug              show this message, plus debugging options
    --help-dyn-options        show the dynamically changeable options
    --version                 show version
    -q --quiet                run silently; only print error msgs
    -v --verbose              be more verbose -- show misc extra info
    --trace-children=no|yes   Valgrind-ise child processes (follow execve)? [no]
    --trace-children-skip=patt1,patt2,...    specifies a list of executables
                              that --trace-children=yes should not trace into
    --trace-children-skip-by-arg=patt1,patt2,...   same as --trace-children-skip=
                              but check the argv[] entries for children, rather
                              than the exe name, to make a follow/no-follow decision
    --child-silent-after-fork=no|yes omit child output between fork & exec? [no]
    --vgdb=no|yes|full        activate gdbserver? [yes]
                              full is slower but provides precise watchpoint/step
    --vgdb-error=<number>     invoke gdbserver after <number> errors [999999999]
                              to get started quickly, use --vgdb-error=0
                              and follow the on-screen directions
    --vgdb-stop-at=event1,event2,... invoke gdbserver for given events [none]
         where event is one of:
           startup exit abexit valgrindabexit all none
    --track-fds=no|yes|all    track open file descriptors? [no]
                              all includes reporting stdin, stdout and stderr
    --time-stamp=no|yes       add timestamps to log messages? [no]
    --log-fd=<number>         log messages to file descriptor [2=stderr]
    --log-file=<file>         log messages to <file>
    --log-socket=ipaddr:port  log messages to socket ipaddr:port
    --enable-debuginfod=no|yes query debuginfod servers for missing
                              debuginfo [yes]

  user options for Valgrind tools that report errors:
    --xml=yes                 emit error output in XML (some tools only)
    --xml-fd=<number>         XML output to file descriptor
    --xml-file=<file>         XML output to <file>
    --xml-socket=ipaddr:port  XML output to socket ipaddr:port
    --xml-user-comment=STR    copy STR verbatim into XML output
    --demangle=no|yes         automatically demangle C++ names? [yes]
    --num-callers=<number>    show <number> callers in stack traces [12]
    --error-limit=no|yes      stop showing new errors if too many? [yes]
    --exit-on-first-error=no|yes exit code on the first error found? [no]
    --error-exitcode=<number> exit code to return if errors found [0=disable]
    --error-markers=<begin>,<end> add lines with begin/end markers before/after
                              each error output in plain text mode [none]
    --show-error-list=no|yes|all  show detected errors list and
                              suppression counts at exit [no].
                              all means to also print suppressed errors.
    -s                        same as --show-error-list=yes
```

Downstream Spack's package is loaded in your environment

e4s-chain-spack.sh helps customize the software stack using upstream /spack (read-only in the container) for package dependencies while installing a new package in the downstream Spack in your writable home directory.

E4S Tools: e4s-alc: Customize container images

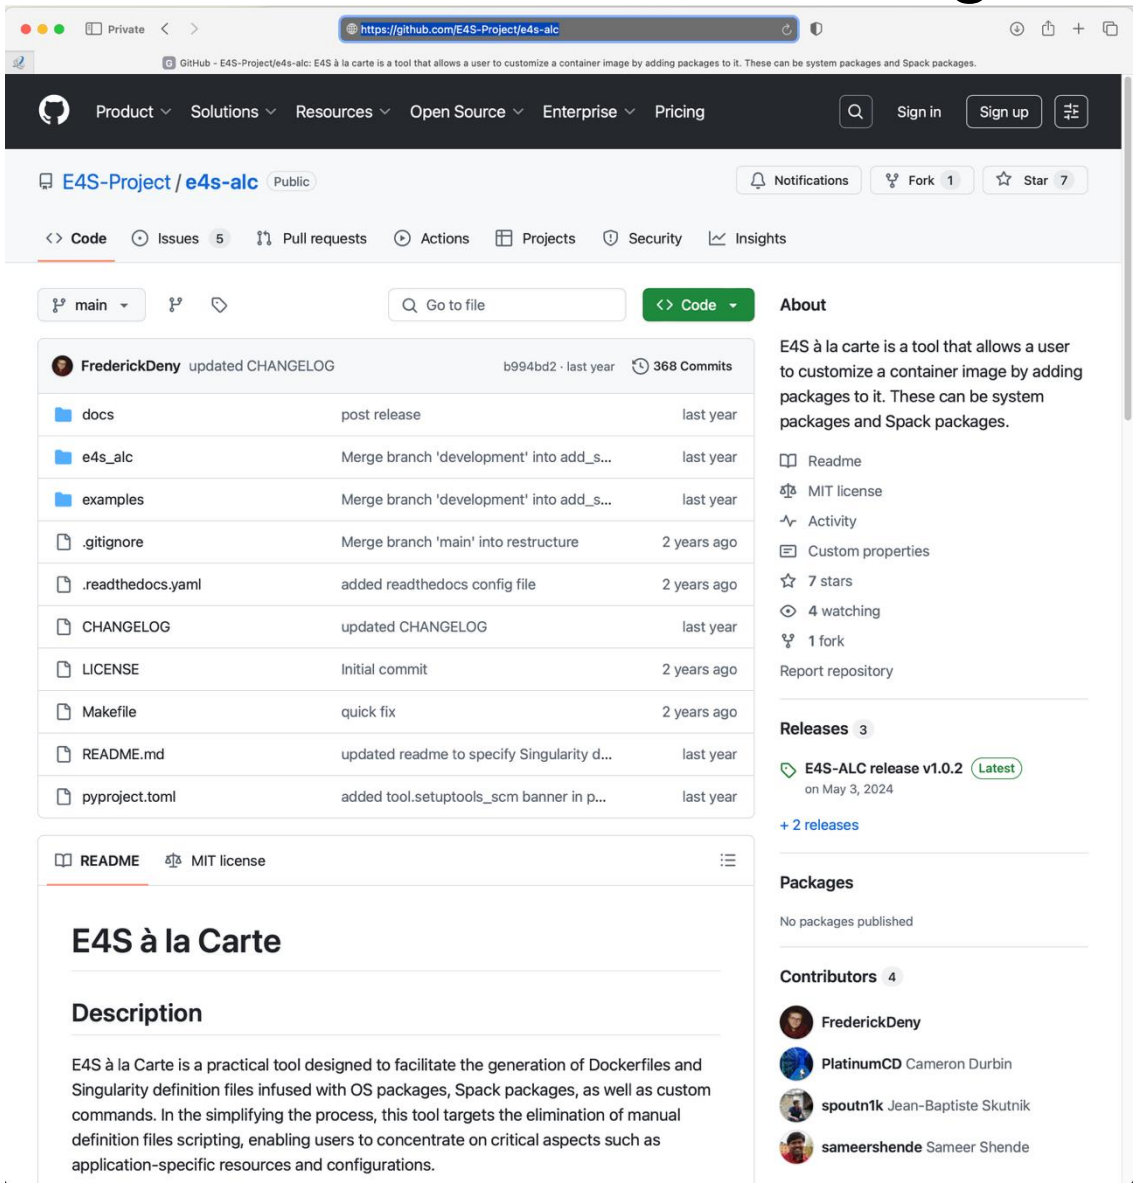

- Add new system packages
- Add new Spack packages
- Add new tarballs
- Customize the container image
- Start with a base image
- Add packages
- Create a new container image!

https://github.com/E4s-Project/e4s-alc

E4S Tools: e4s-cl: Container Launch tool for MPI applications

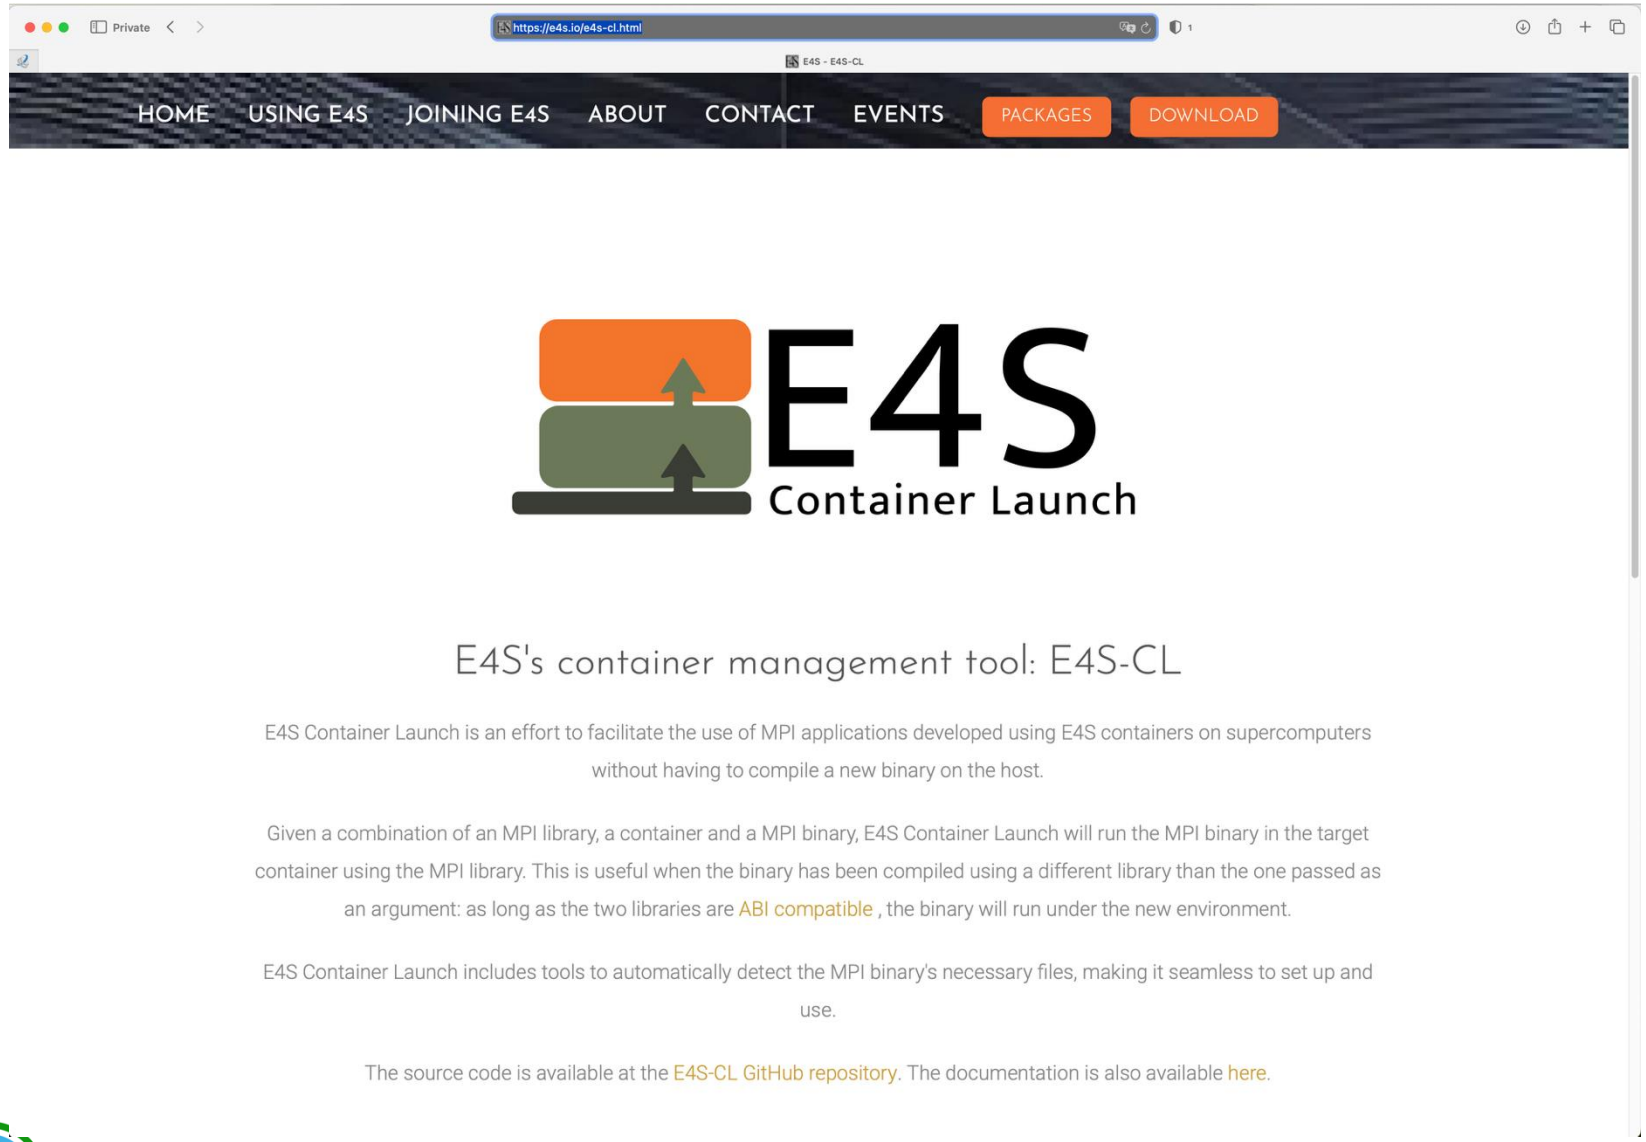

- Distribute your MPI application as a binary with an E4S image
- While deploying on a system substitute the embedded containerized MPI in application with the system/vendor MPI
- Use inter-node network interfaces efficiently for near native performance!

https://e4s.io/e4s-cl.html

e4s-cl: A tool to simplify the launch of MPI jobs in E4S containers

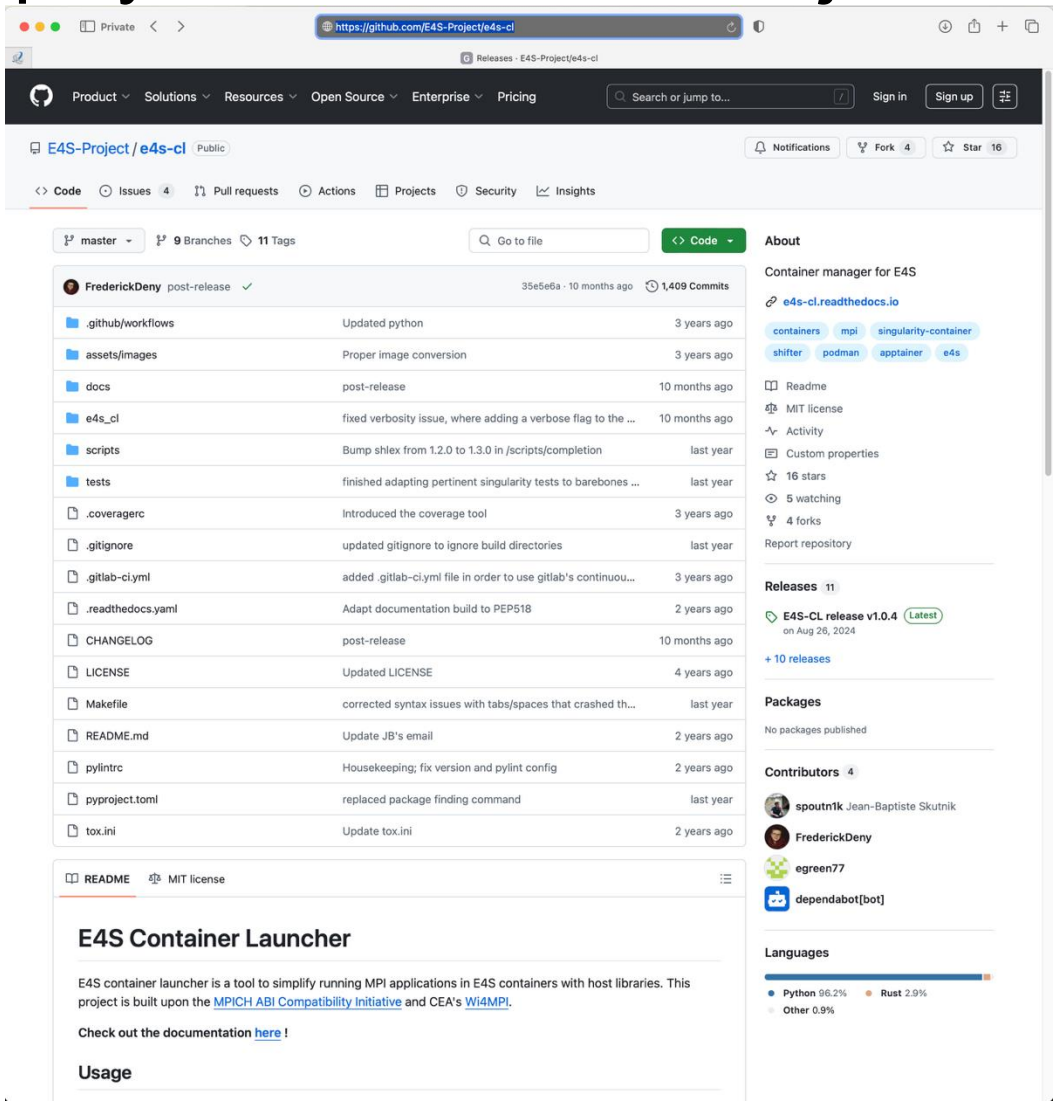

https://github.com/E4S-Project/e4s-cl

e4s-cl: A tool to simplify the launch of MPI jobs in E4S containers

- E4S containers support replacement of MPI libraries using MPICH ABI compatibility layer and Wi4MPI [CEA] for OpenMPI replacement.
- Applications binaries built using E4S can be launched with Singularity using MPI library substitution for efficient inter-node communications.
- e4s-cl is a new tool that simplifies the launch and MPI replacement.
  - e4s-cl init --backend [singularity|shifter|docker] --image <file> --source <startup_cmds.sh>
  - e4s-cl mpirun -np <N> <command>
- Usage:

```
% e4s-cl init --backend singularity --image ~/images/e4s-gpu-x86.sif --source ~/source.sh
% cat ~/source.sh
   . /spack/share/spack/setup-env.sh
   spack load trilinos+cuda cuda_arch=90
% e4s-cl mpirun -np 4 ./a.out
```

https://github.com/E4S-Project/e4s-cl

PESO

E4S on Commercial Cloud Platforms: ParaTools Pro for E4S™

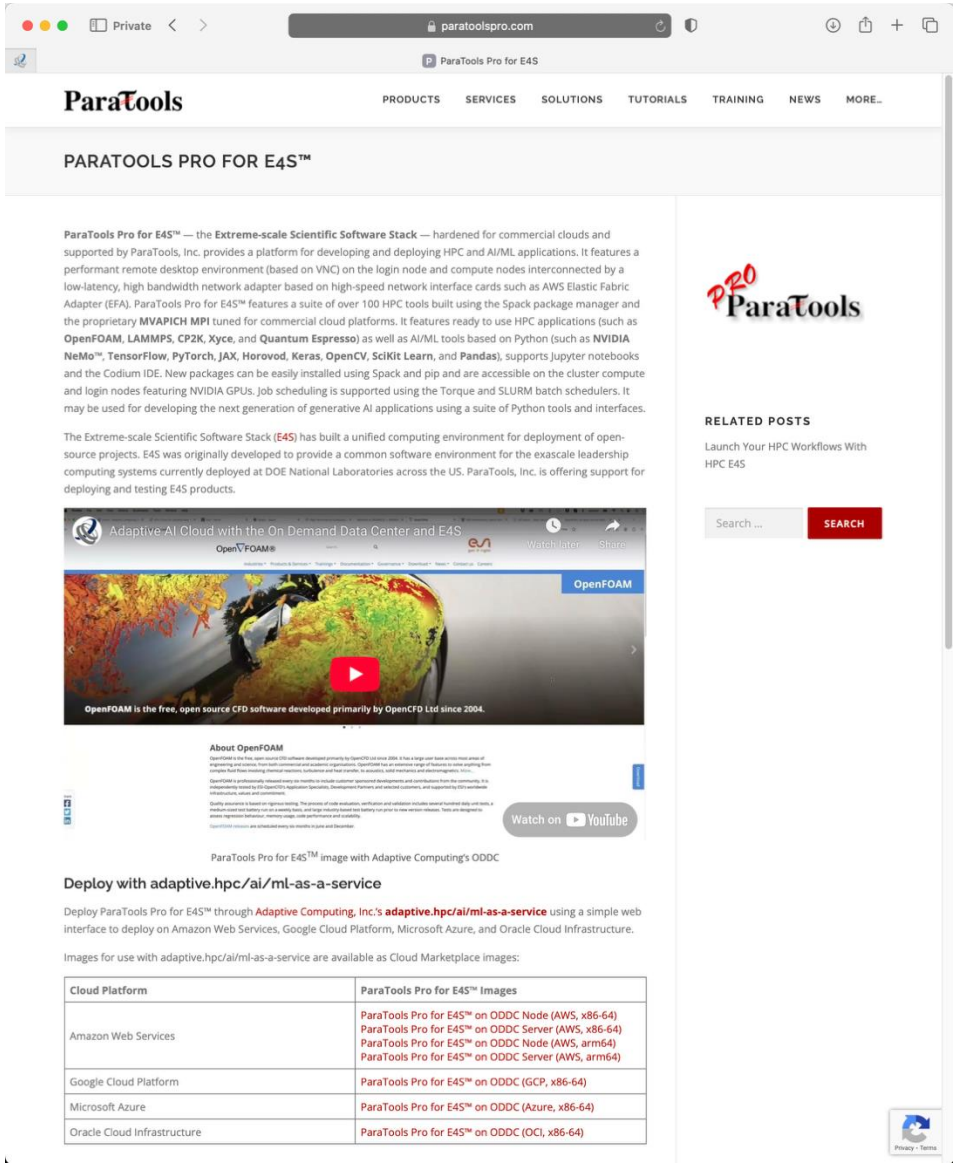

- ParaTools Pro for E4S™* images in vendor marketplaces support:
  - AWS
  - Azure
  - Google Cloud (GCP)
  - Oracle Cloud Infrastructure (OCI)
- Supports SLURM and Torque for scheduling jobs on multi-node GPU accelerated nodes
- Shared GPU accelerated login node with a VNC based remote desktop
- Adaptive Computing's Heidi/ODDC
- AWS PCS and PC (x86, ARM64)
- Azure Cyclecloud
- Google GCluster

* Acknowledgment:
Supported by
DOE SBIR Phase I and II
DE-SC0022502

U.S. DEPARTMENT OF ENERGY | Office of Science

PESO

https://paratoolspro.com and https://www.energy.gov/technologytransitions/sbirsttr

E4S on Adaptive Computing's Heidi AI/On Demand Data Center (ODDC)

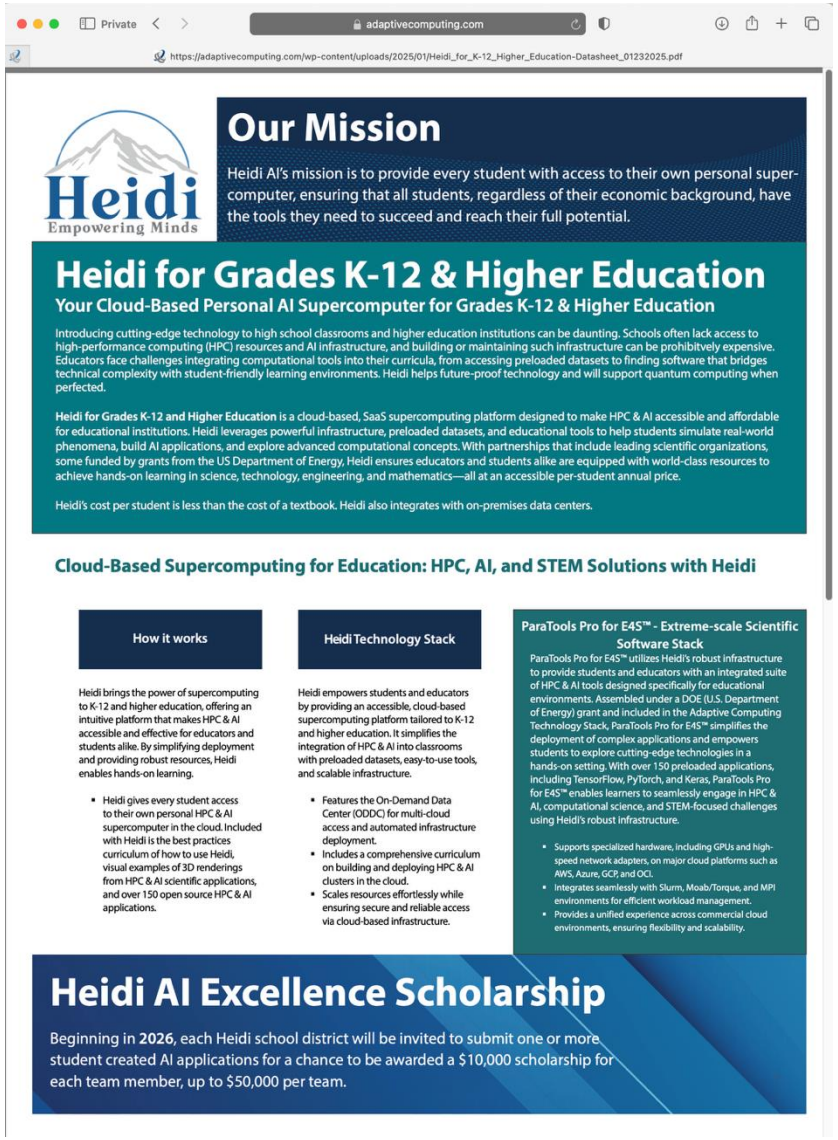

Private adaptivecomputing.com

https://adaptivecomputing.com/wp-content/uploads/2025/01/Heidi_for_K-12_Higher_Education-Datasheet_01232025.pdf

Heidi
Empowering Minds

Our Mission

Heidi AI's mission is to provide every student with access to their own personal super-computer, ensuring that all students, regardless of their economic background, have the tools they need to succeed and reach their full potential.

Heidi for Grades K-12 & Higher Education

Your Cloud-Based Personal AI Supercomputer for Grades K-12 & Higher Education

Introducing cutting-edge technology to high school classrooms and higher education institutions can be daunting. Schools often lack access to high-performance computing (HPC) resources and AI infrastructure, and building or maintaining such infrastructure can be prohibitvely expensive. Educators face challenges integrating computational tools into their curricula, from accessing preloaded datasets to finding software that bridges technical complexity with student-friendly learning environments. Heidi helps future-proof technology and will support quantum computing when perfected.

Heidi for Grades K-12 and Higher Education is a cloud-based, SaaS supercomputing platform designed to make HPC & AI accessible and affordable for educational institutions. Heidi leverages powerful infrastructure, preloaded datasets, and educational tools to help students simulate real-world phenomena, build AI applications, and explore advanced computational concepts. With partnerships that include leading scientific organizations, some funded by grants from the US Department of Energy, Heidi ensures educators and students alike are equipped with world-class resources to achieve hands-on learning in science, technology, engineering, and mathematics—all at an accessible per-student annual price.

Heidi's cost per student is less than the cost of a textbook. Heidi also integrates with on-premises data centers.

Cloud-Based Supercomputing for Education: HPC, AI, and STEM Solutions with Heidi

How it works

Heidi brings the power of supercomputing to K-12 and higher education, offering an intuitive platform that makes HPC & AI accessible and effective for educators and students alike. By simplifying deployment and providing robust resources, Heidi enables hands-on learning.

- Heidi gives every student access to their own personal HPC & AI supercomputer in the cloud. Included with Heidi is the best practices curriculum of how to use Heidi, visual examples of 3D renderings from HPC & AI scientific applications, and over 150 open source HPC & AI applications.

Heidi Technology Stack

Heidi empowers students and educators by providing an accessible, cloud-based supercomputing platform tailored to K-12 and higher education. It simplifies the integration of HPC & AI into classrooms with preloaded datasets, easy-to-use tools, and scalable infrastructure.

- Features the On-Demand Data Center (ODDC) for multi-cloud access and automated infrastructure deployment.
- Includes a comprehensive curriculum on building and deploying HPC & AI clusters in the cloud.
- Scales resources effortlessly while ensuring secure and reliable access via cloud-based infrastructure.

ParaTools Pro for E4S™ - Extreme-scale Scientific Software Stack

ParaTools Pro for E4S™ utilizes Heidi's robust infrastructure to provide students and educators with an integrated suite of HPC & AI tools designed specifically for educational environments. Assembled under a DOE (U.S. Department of Energy) grant and included in the Adaptive Computing Technology Stack, ParaTools Pro for E4S™ simplifies the deployment of complex applications and empowers students to explore cutting-edge technologies in a hands-on setting. With over 150 preloaded applications, including TensorFlow, PyTorch, and Keras, ParaTools Pro for E4S™ enables learners to seamlessly engage in HPC & AI, computational science, and STEM-focused challenges using Heidi's robust infrastructure.

- Supports specialized hardware, including GPUs and high-speed network adapters, on major cloud platforms such as AWS, Azure, GCP, and OCI.
- Integrates seamlessly with Slurm, Moab/Torque, and MPI environments for efficient workload management.
- Provides a unified experience across commercial cloud environments, ensuring flexibility and scalability.

Heidi AI Excellence Scholarship

Beginning in **2026**, each Heidi school district will be invited to submit one or more student created AI applications for a chance to be awarded a $10,000 scholarship for each team member, up to $50,000 per team.

- ParaTools Pro for E4S™ images in commercial cloud marketplaces launched using Heidi
- Supports Torque for scheduling jobs on multi-node GPU accelerated nodes
- Shared GPU accelerated login node with a VNC based remote desktop

https://adaptivecomputing.com

ParaTools Pro for E4S™ on Commercial Clouds: AWS Marketplace

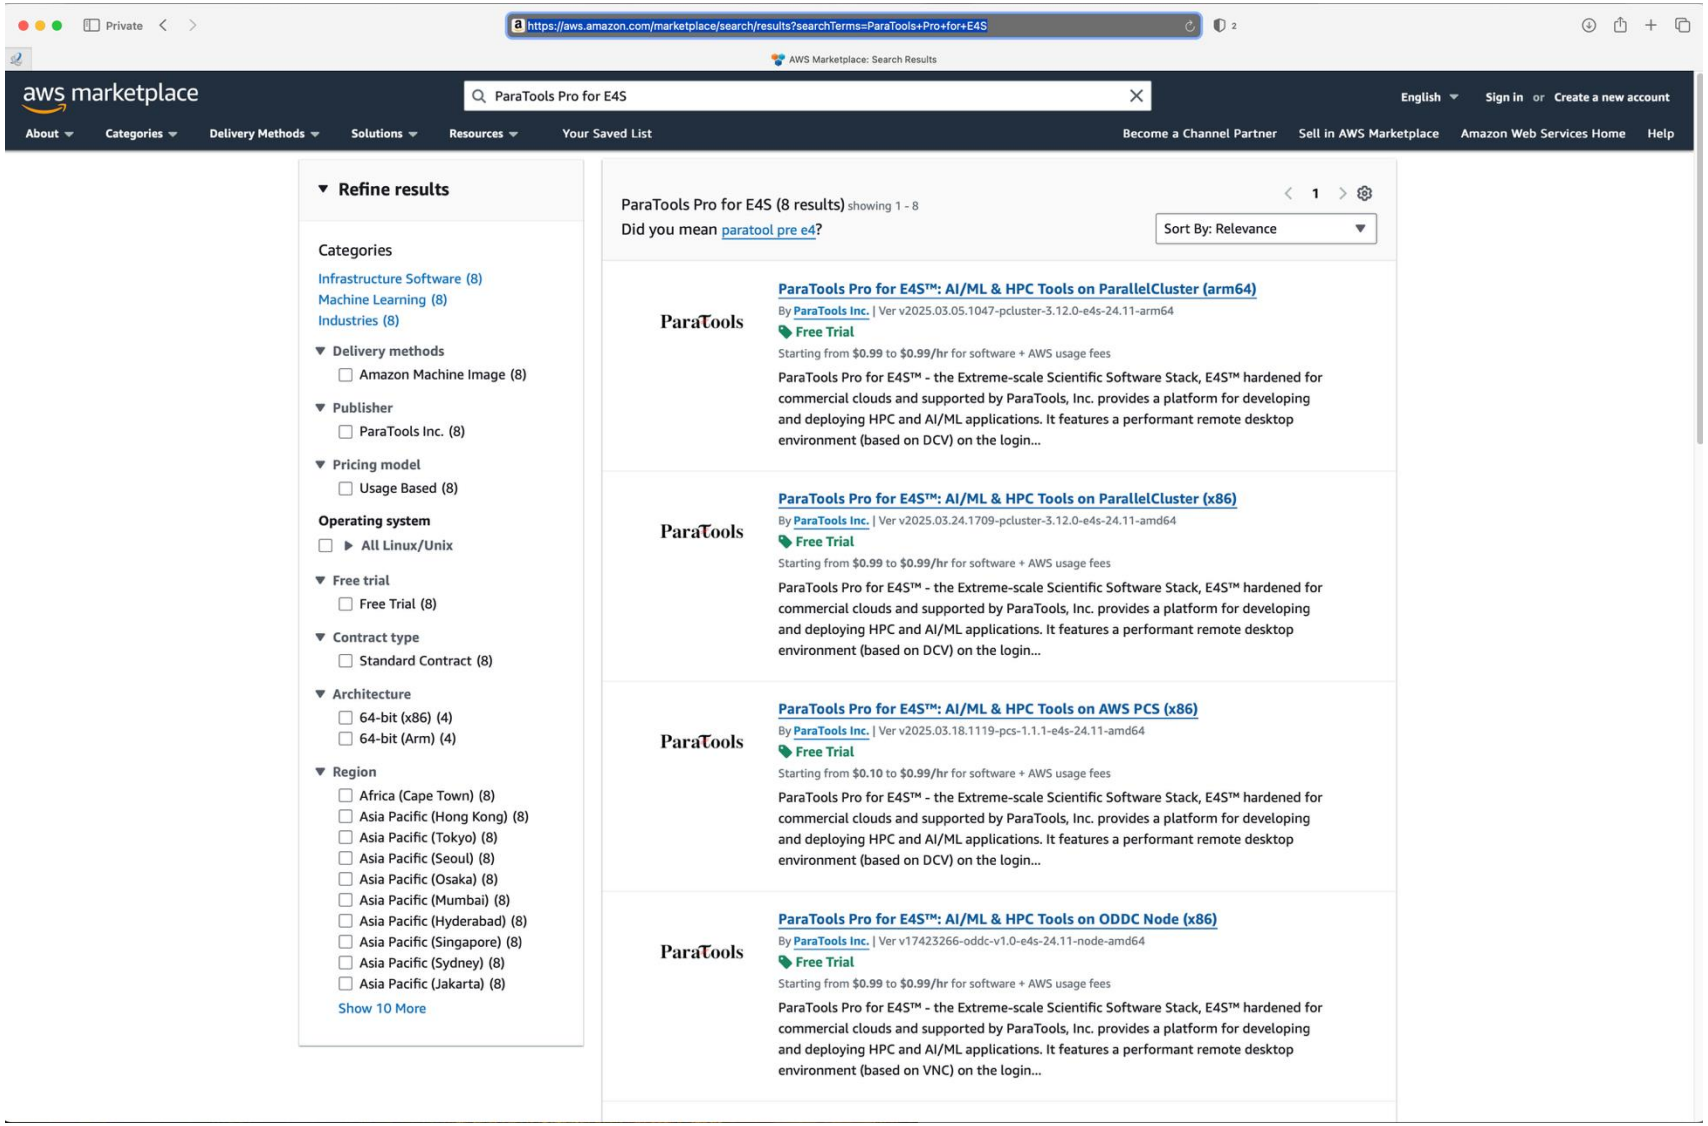

ParaTools Pro for E4S™ on AWS supports

- AWS Trainium and Inferentia custom AI hardware with NeuronX SDK
- AWS PCS and PC on x86_64 and aarch64 nodes
- NVIDIA GPUs
- SLURM (PCS and PC) and Torque (ODDC node/server)
- Also on AWS Marketplace in GovCloud (US East & West)
- Elastic Fabric Adapter (EFA)
- MVAPICH MPI [X-ScaleSolutions, LLC and The Ohio State University]

ParaTools Pro for E4S™ on Commercial Clouds: Azure Marketplace

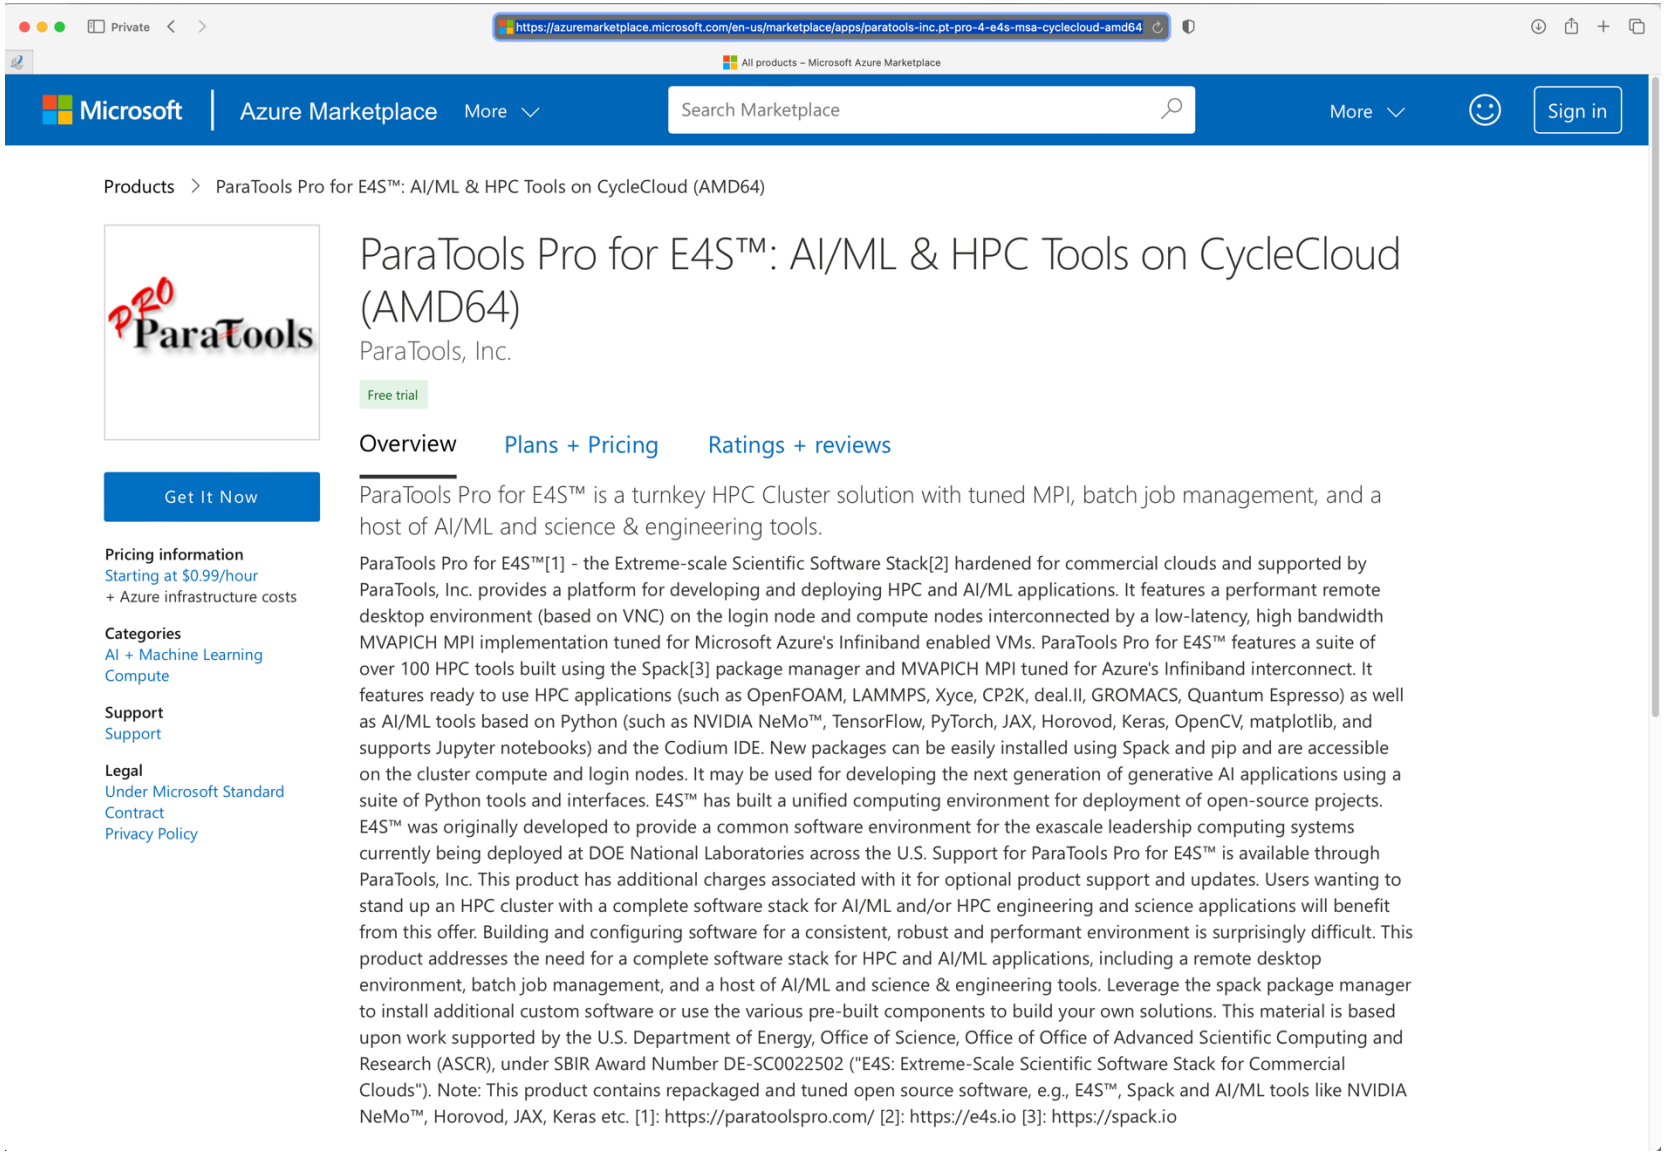

ParaTools Pro for E4S™ on Azure Marketplace supports

- SLURM (Azure CycleCloud) and Torque schedulers (Adaptive Computing's ODDC)
- Support for Infiniband Network adapter

PESO

ParaTools Pro for E4S™ on Google Cloud Marketplace

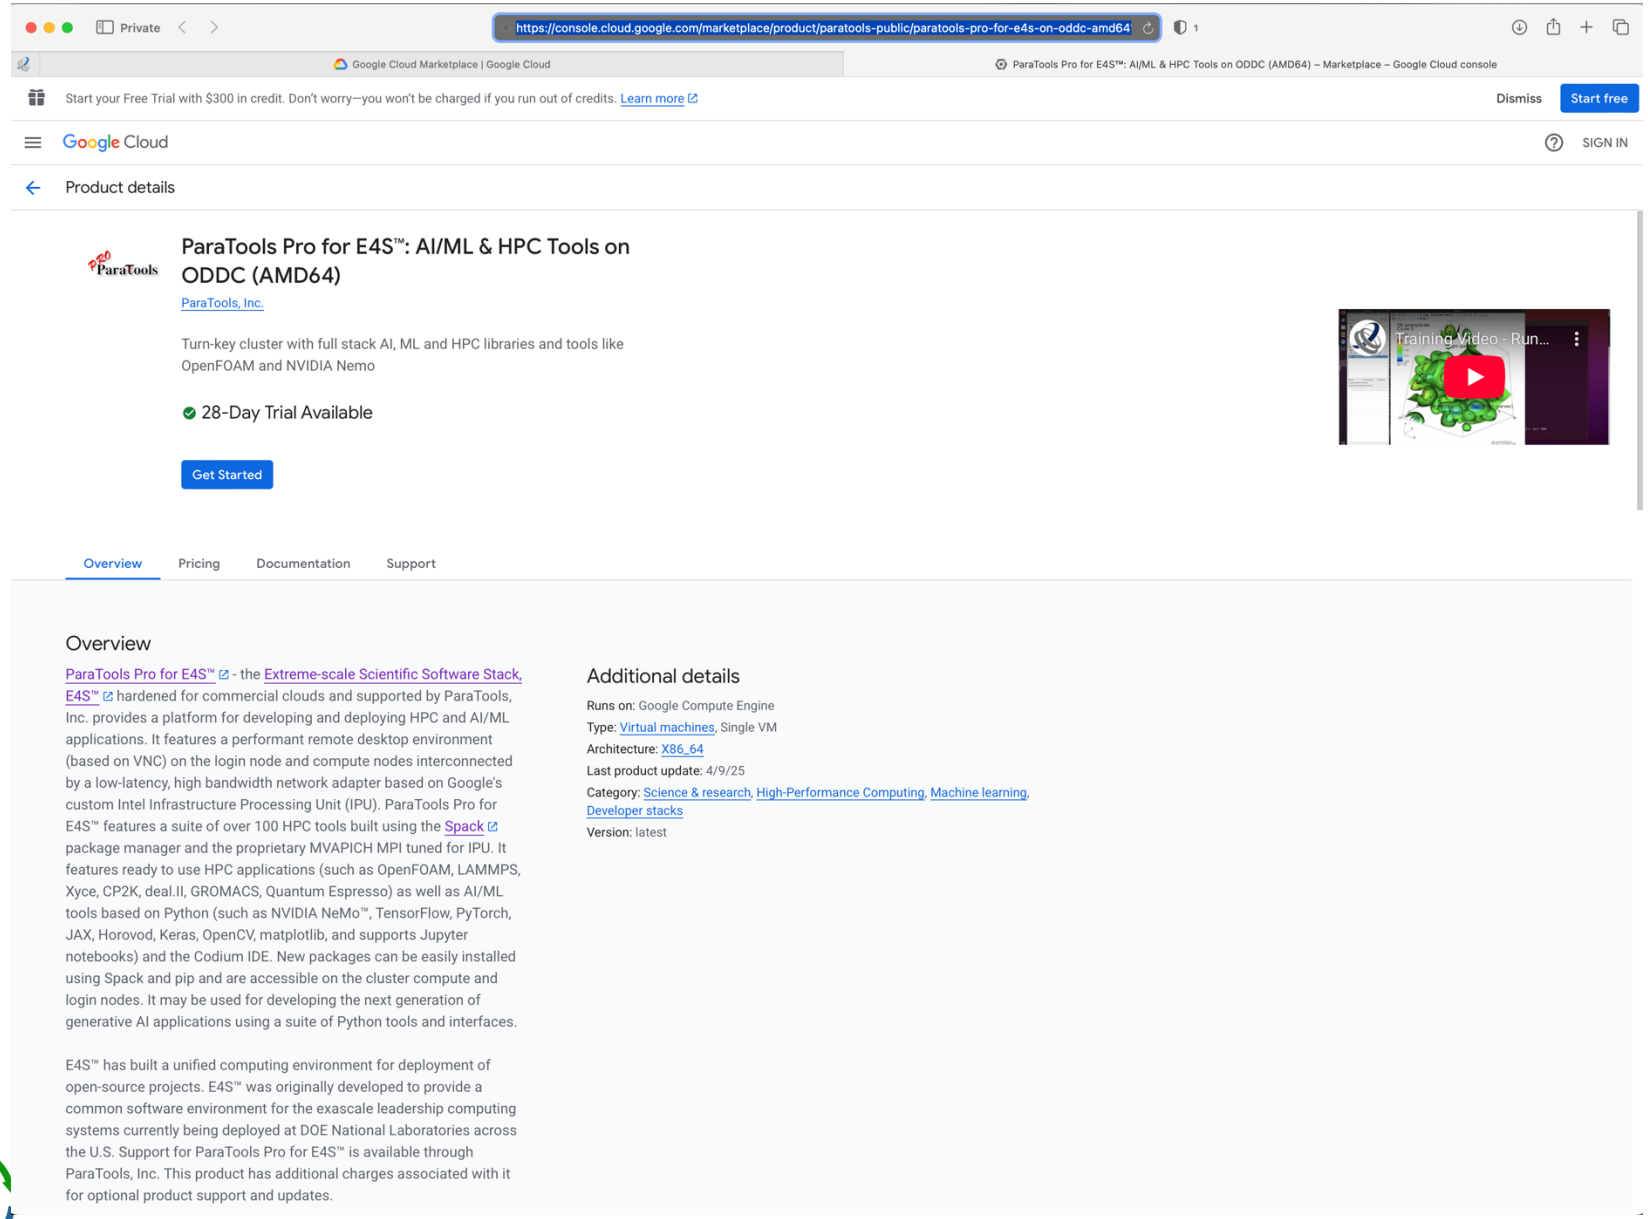

ParaTools Pro for E4S™ Google Cloud Marketplace supports

- SLURM (GCluster) and Torque schedulers (Adaptive Computing's ODDC)
- Support for Google IPU network adapters

ParaTools Pro for E4S™ on Oracle Cloud Marketplace

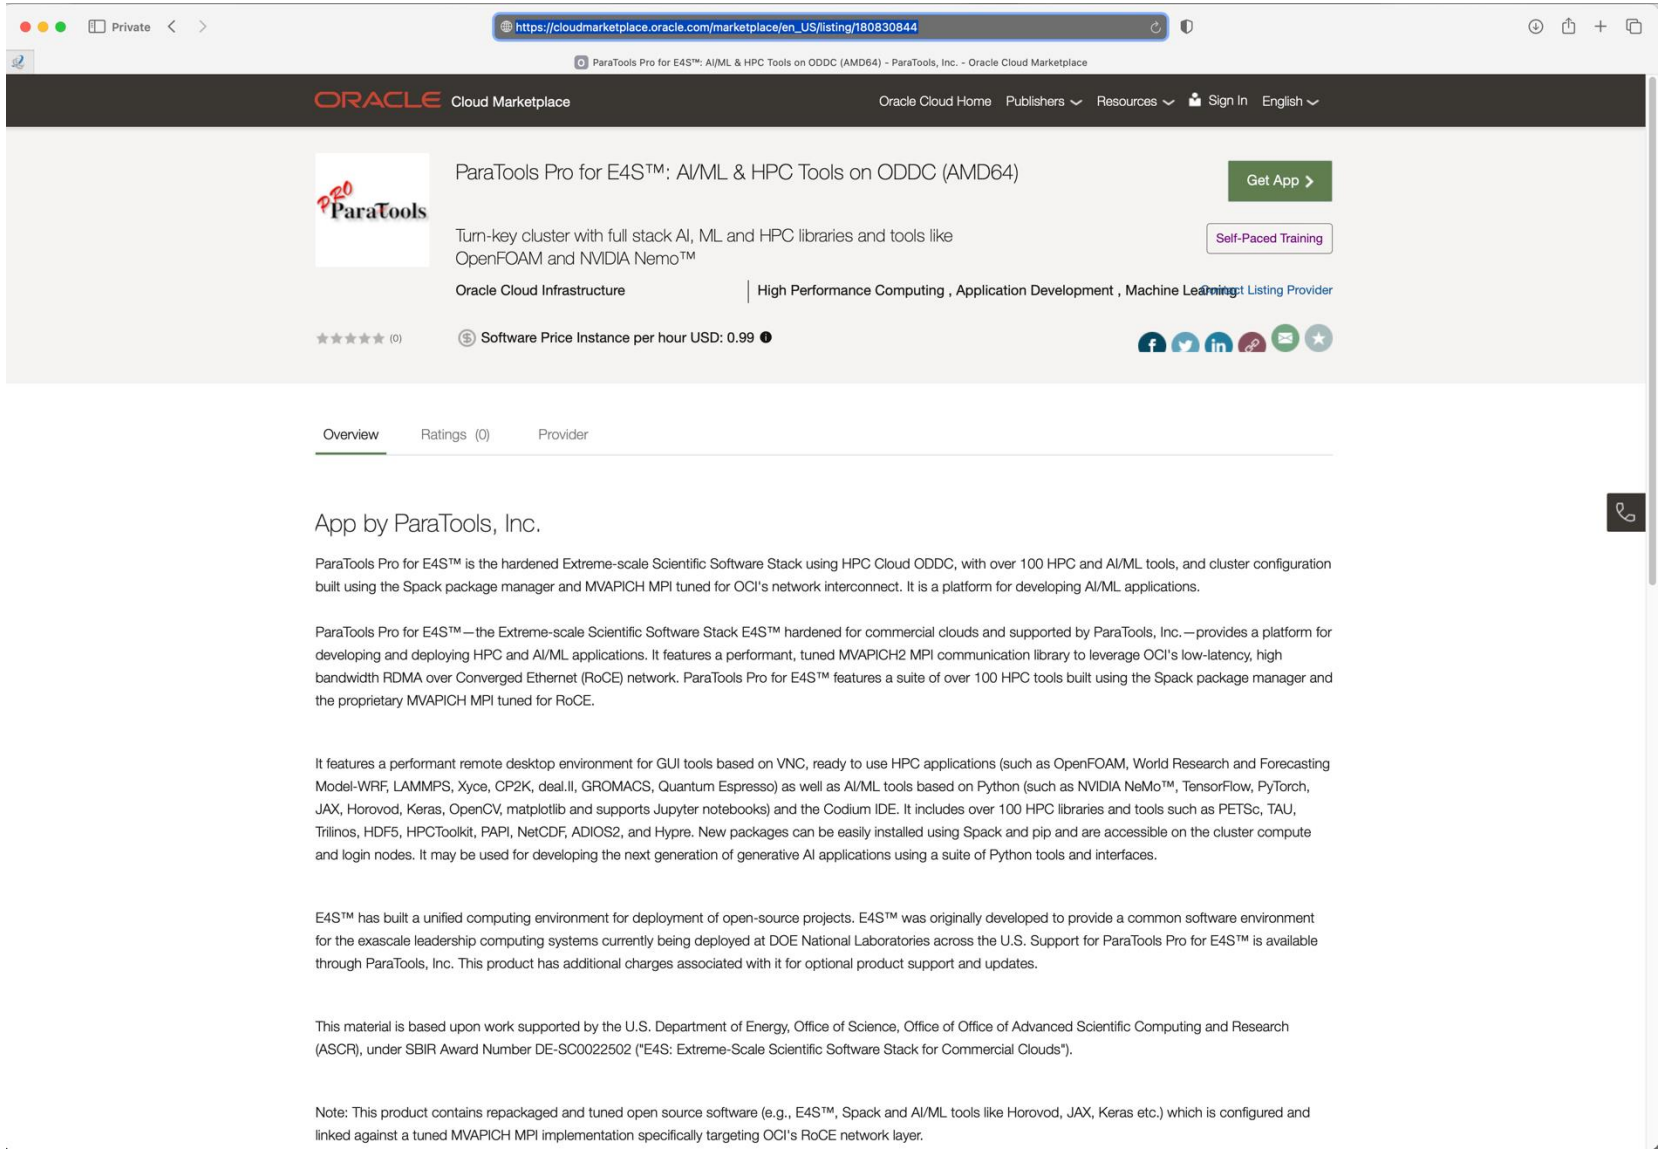

ParaTools Pro for E4S™
for Oracle Cloud Infrastructure
(OCI) Marketplace
supports Torque (ODDC) and
RDMA over Converged Ethernet
(RoCE) network adapters and
GPUs on login and compute nodes

E4S 25.06 Release on DockerHub: Base images

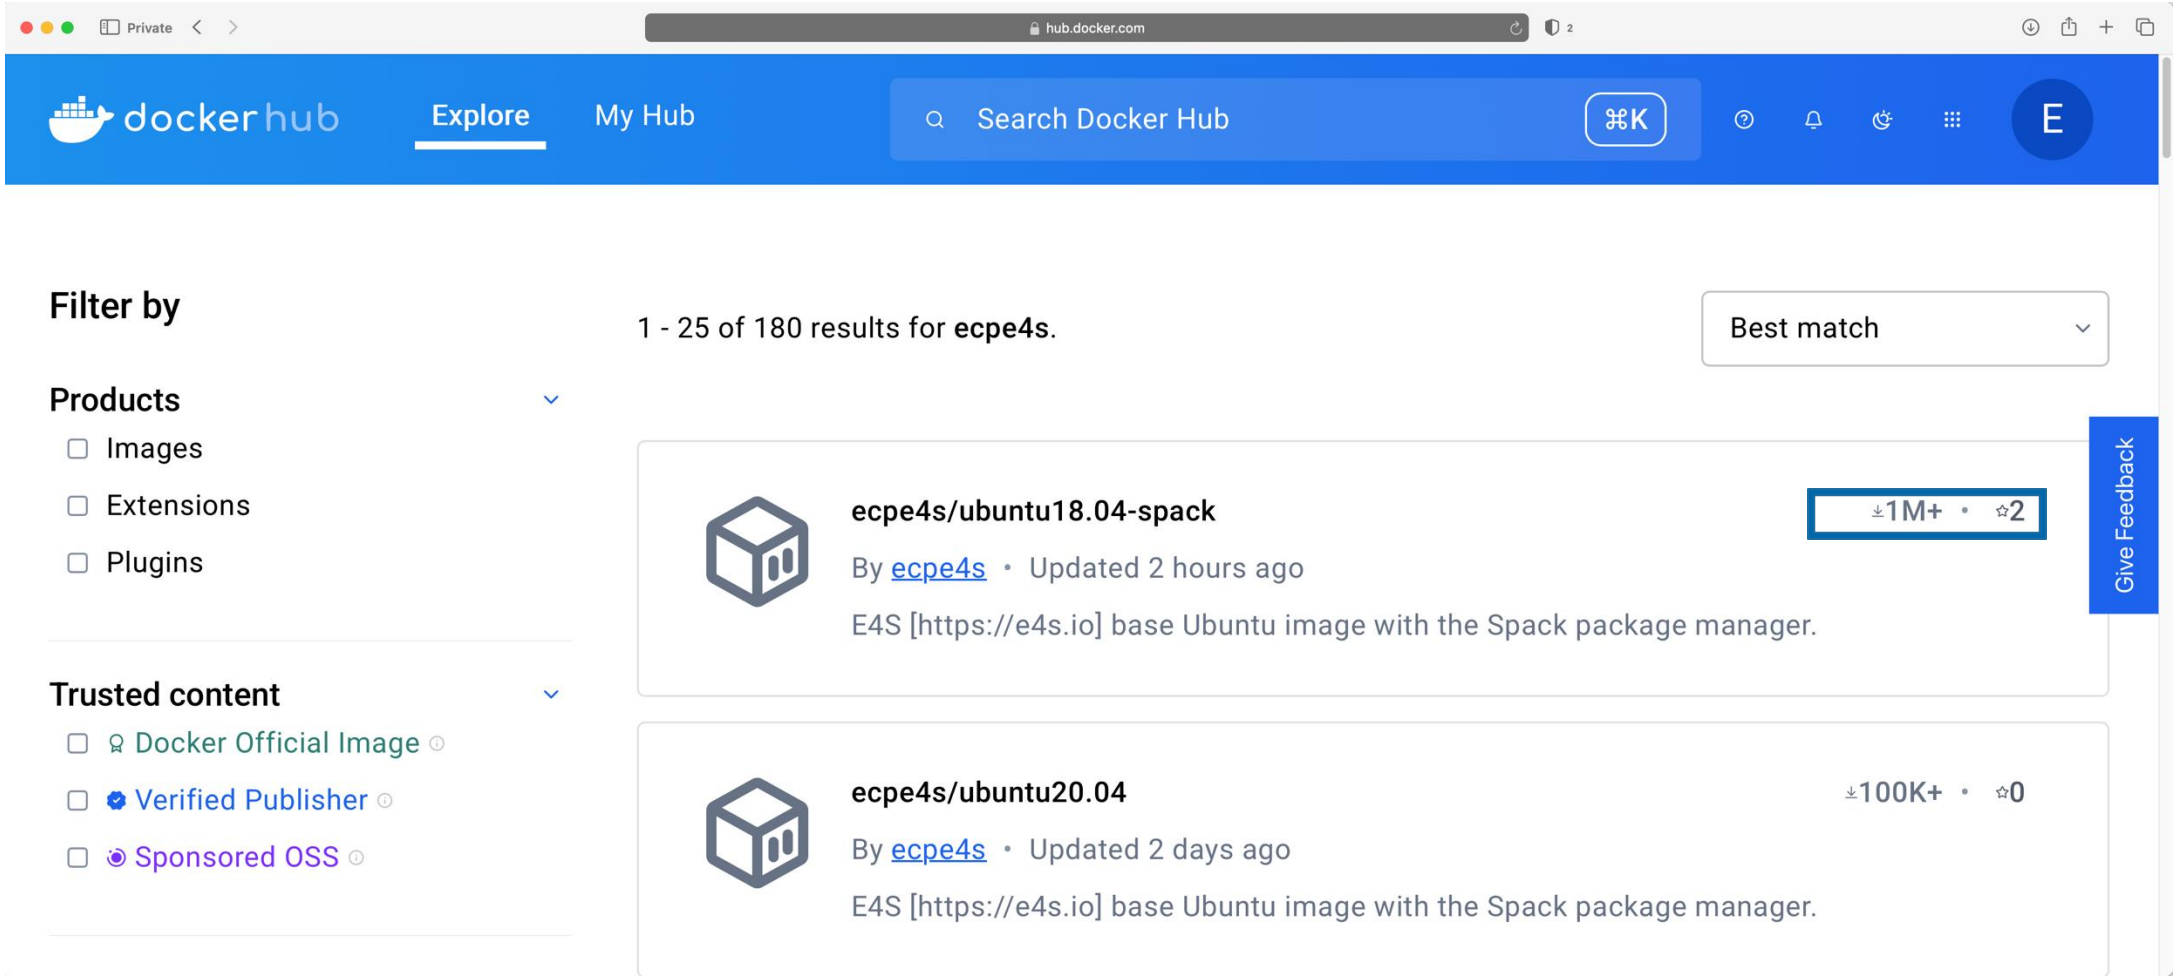

https://e4s.io

E4S 25.06 Release on DockerHub: Full Featured CUDA image

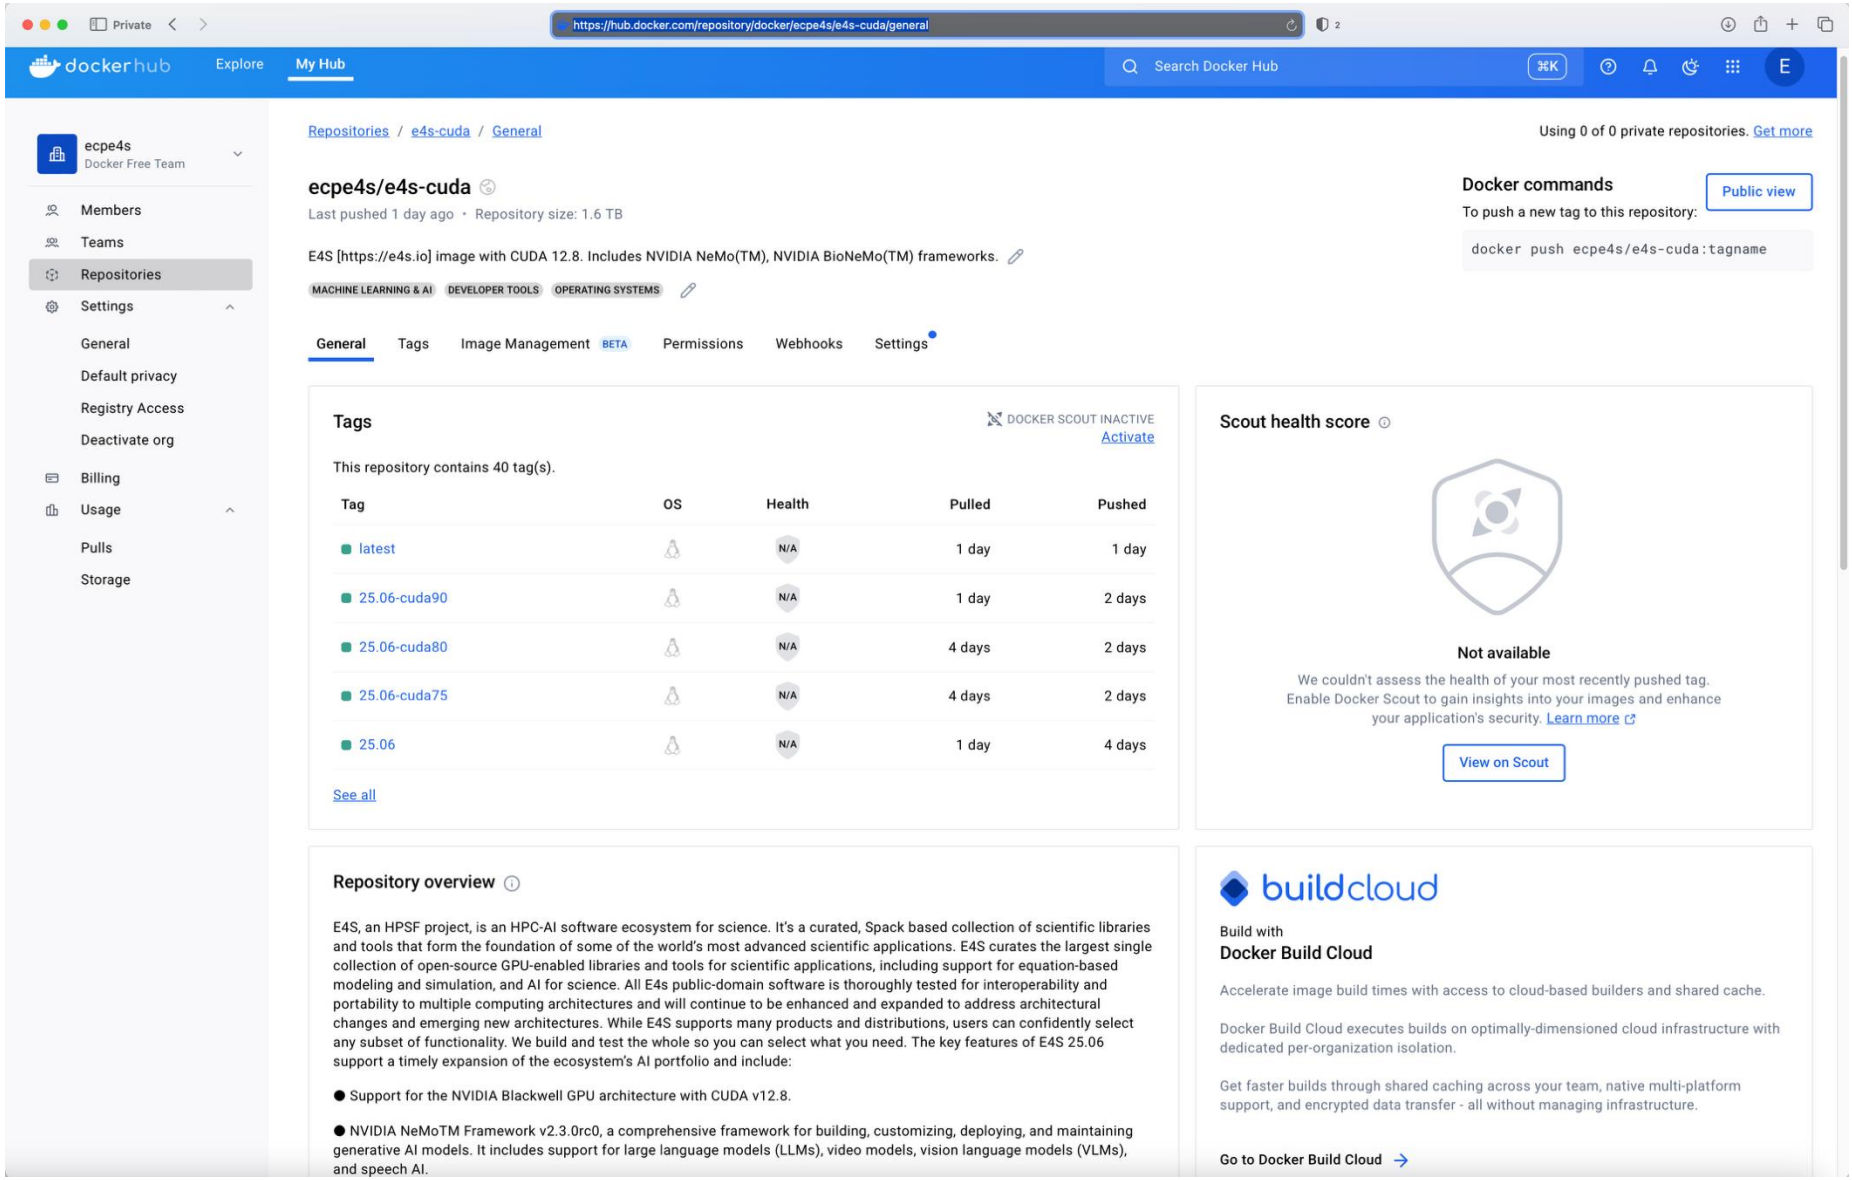

https://e4s.io

E4S 25.06 Release on DockerHub: Full Featured ROCm image

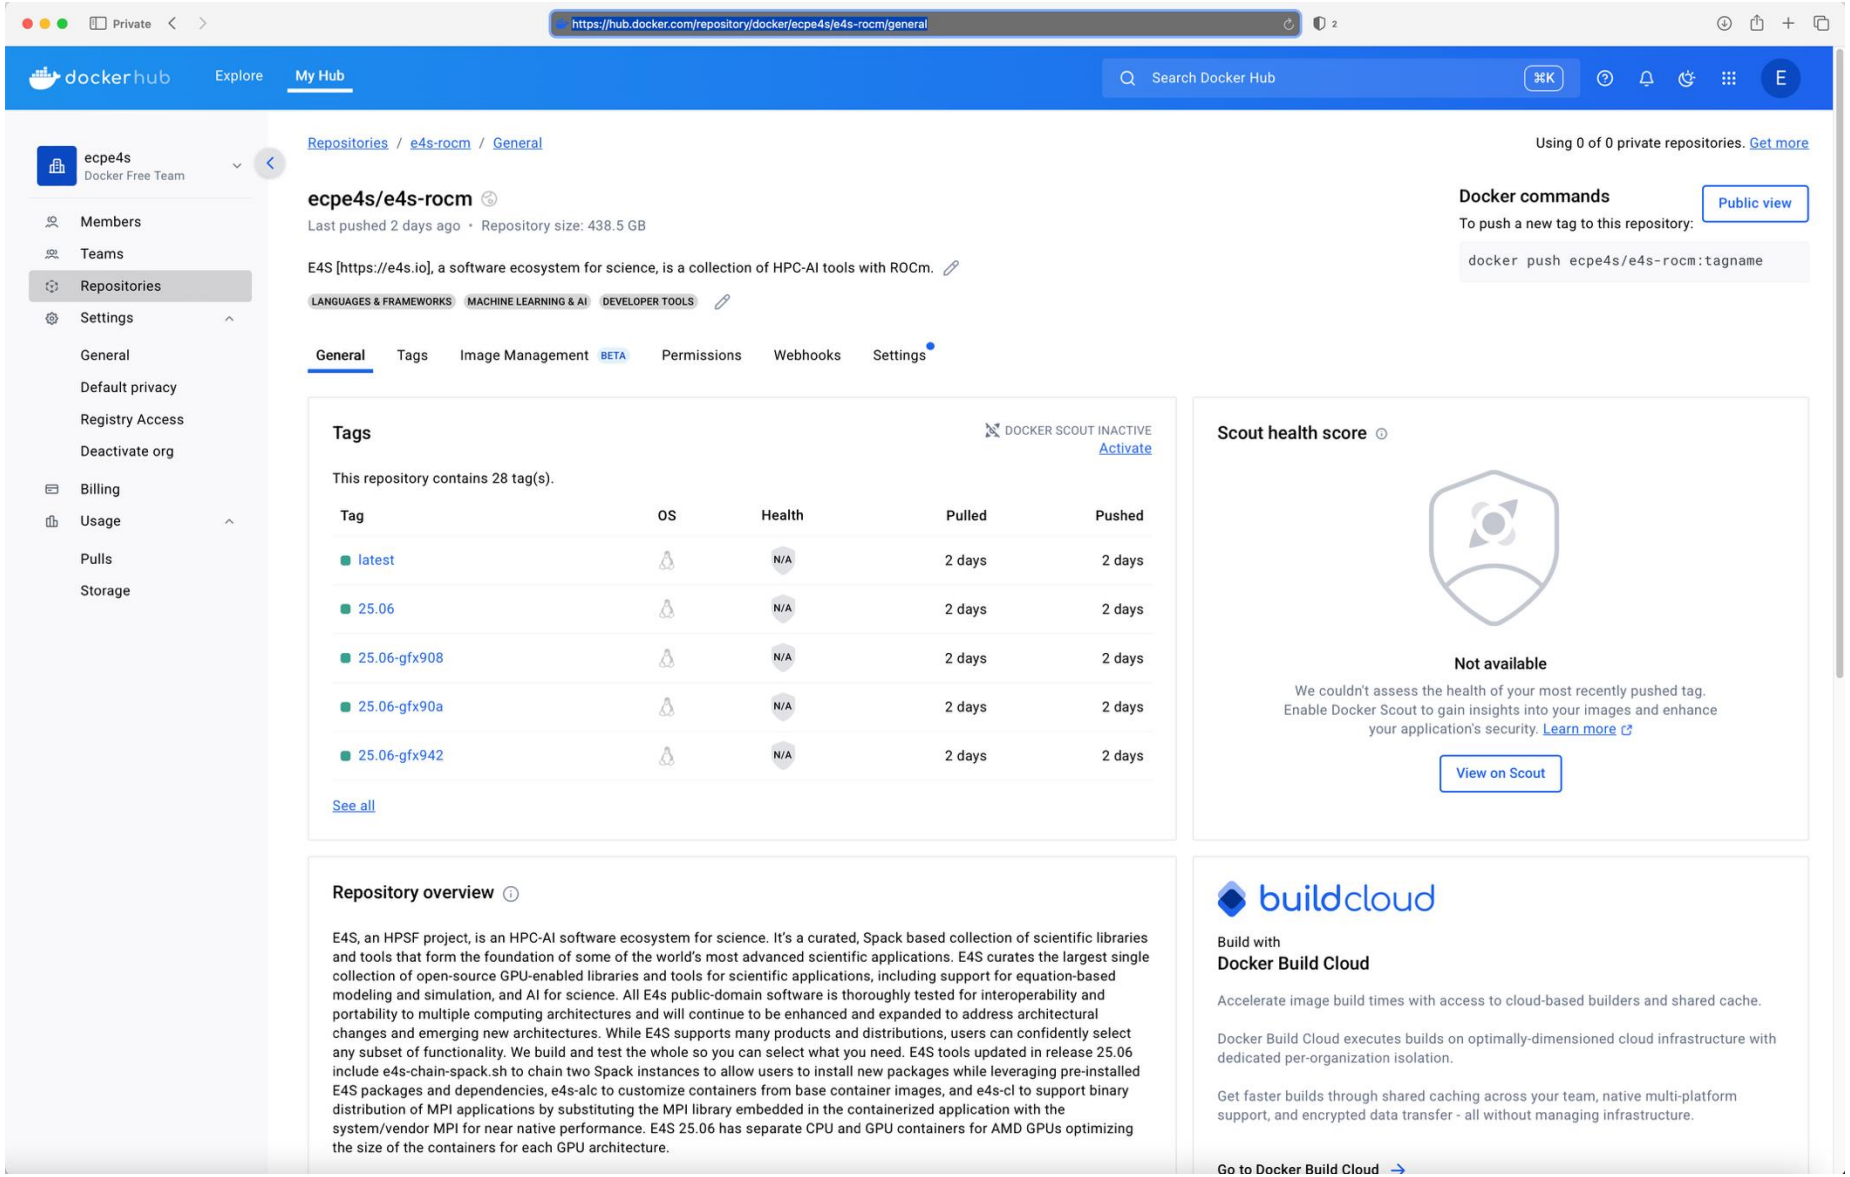

https://e4s.io

E4S 25.06 Release on DockerHub: Full Featured oneAPI image

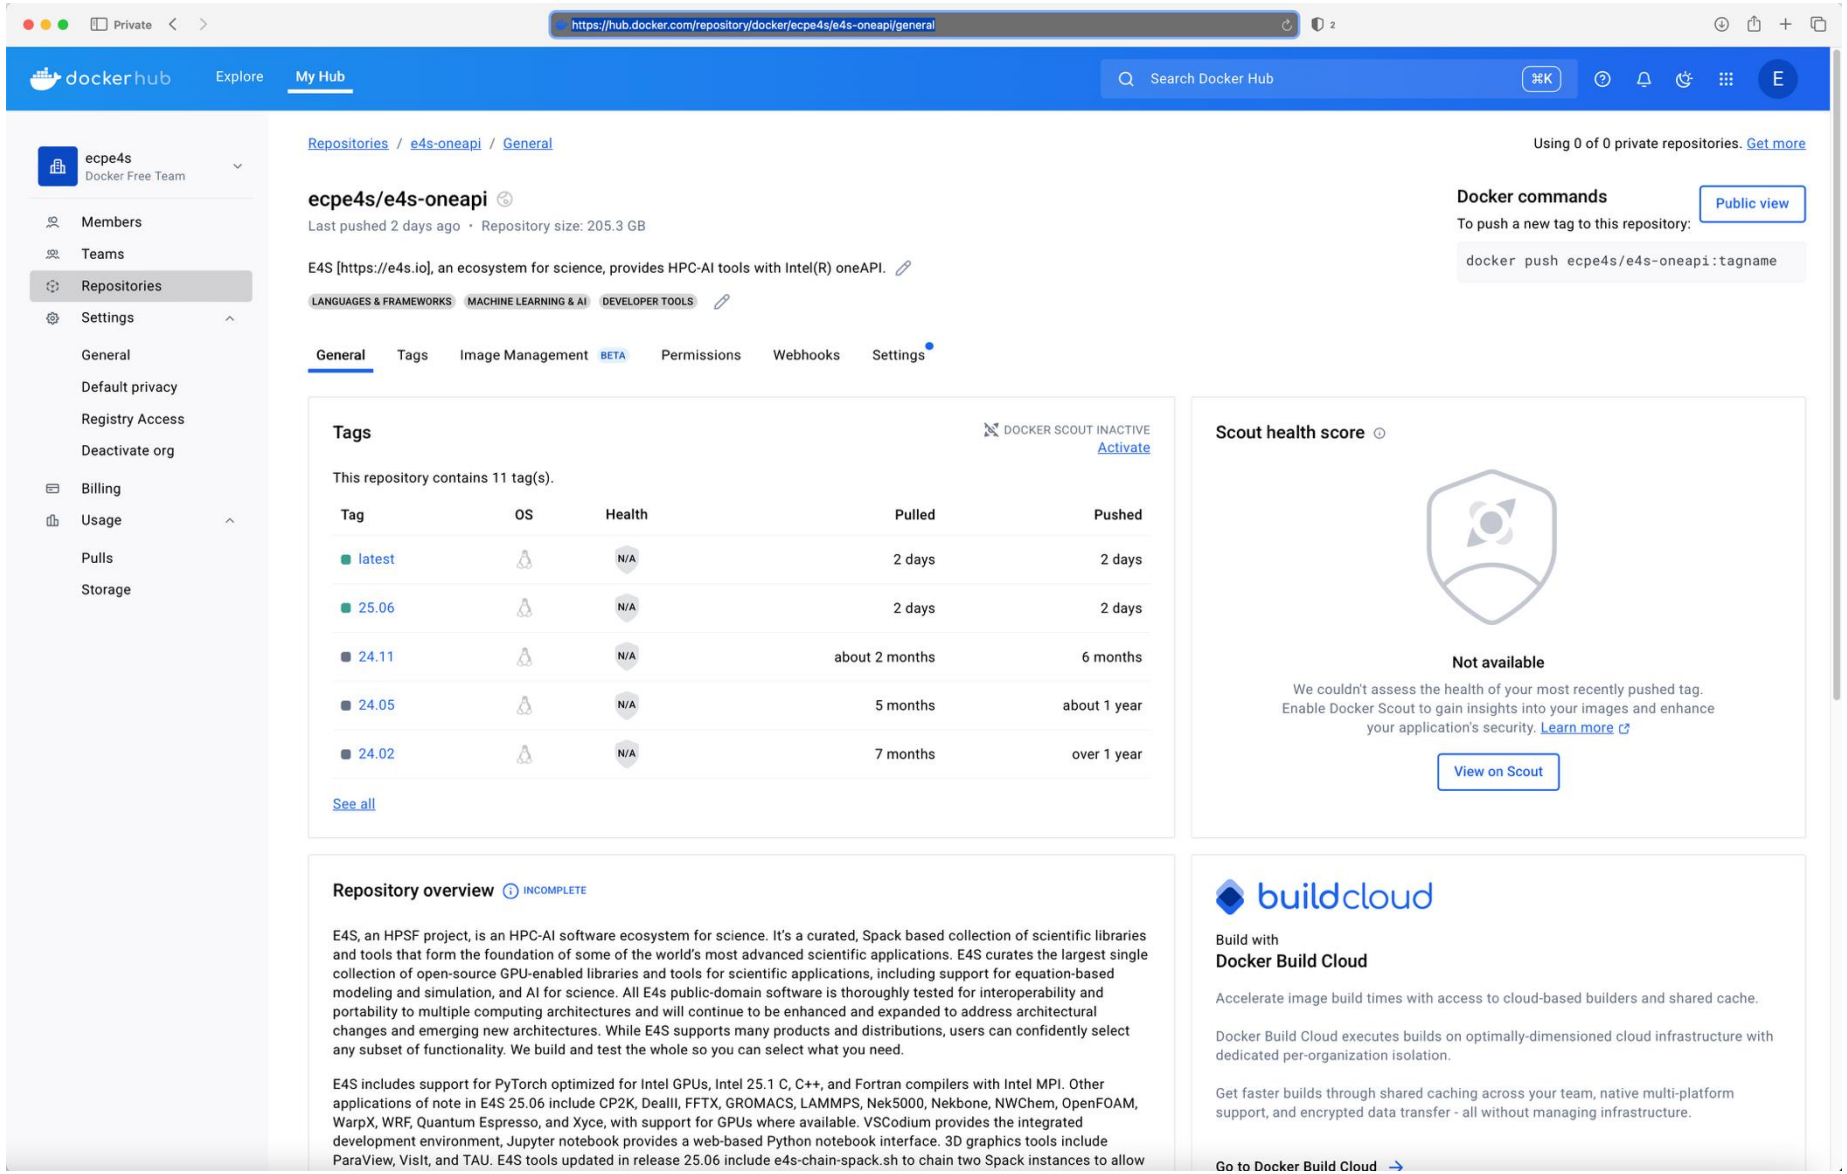

https://e4s.io

E4S 25.06 image for NVIDIA H100 GPU on x86_64

```
$ singularity run --nv e4s-cuda90-x86_64-25.06.sif
Singularity> ls /opt/demo/e4s-cloud-examples/
bionemo       cuda          julia-cuda       lammps            nalu                openfoam       pytorch                  superlu-dist-cpu  vllm
clean-all.sh  fetch-all.sh  julia-mpi        machine-learning  nemo                osu-benchmarks pytorch-gpu              tau               xyce
CoMD          horovod       jupyter-notebook matmult           nemo-speech_to_text petsc-cpu      pytorch-image-classifier tensorflow
containers    jax           laghos           mpi-procname      neuronx             petsc-cuda     qe                       visit
Singularity> ls /opt/demo/e4s-cloud-examples/machine-learning/
clean.sh  gemini  openai  perplexity  pytorch  tensorflow
Singularity> ls /opt/demo/e4s-cloud-examples/vllm
gradio_openai_chatbot_webserver.py  llama2_template.jinja  README.md  run.sh  run-smaller.sh
Singularity> python
Python 3.10.12 (main, Feb  4 2025, 14:57:36) [GCC 11.4.0] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>> import nemo
>>> import bionemo
>>> import torch
>>> import openai
>>> import google.generativeai
>>> import huggingface_hub
>>> import jax
>>> import pandas
>>> import cv2
>>> import sklearn
>>> import mpi4py
>>> import matplotlib
>>> import seaborn
>>> import plotly
>>> import vllm
```

```
>>> vllm.__version__
'0.8.3.dev0+g25f560a62.d20250520'
>>> nemo.__version__
'2.3.0rc0'
>>> tensorflow.__version__
'2.19.0'
>>> torch.__version__
'2.6.0'
>>> torch.cuda.get_arch_list()
['sm_80', 'sm_90', 'sm_120']
>>>
```

https://e4s.io

E4S 25.06 image for NVIDIA GPUs (x86_64)

```
Singularity> spack find
-- linux-ubuntu22.04-x86_64_v3 / gcc@11.4.0 ---------------------
```

abseil-cpp@20240722.0	expat@2.7.0	libceed@0.12.0	openssl@3.4.1	py-urllib3@2.1.0
adiak@0.4.1	fftw@3.3.10	libdwarf@0.11.0	otf2@3.0.3	py-wheel@0.45.1
adios2@2.10.2	fftw@3.3.10	libedit@3.1-20240808	papi@7.1.0	python@3.10.12
amrex@25.03	fftx@1.2.0	libevent@2.1.12	papi@7.1.0	python-venv@1.0
arborx@1.5	findutils@4.10.0	libfabric@1.22.0	parmetis@4.0.3	raja@2024.07.0
arpack-ng@3.9.1	flecsi@2.3.2	libffi@3.4.6	parmetis@4.0.3	raja@2024.07.0
asio@1.32.0	flex@2.6.3	libiberty@2.41	parsec@3.0.2209	re2c@3.1
autoconf@2.72	flux-core@0.67.0	libiconv@1.17	pcre2@10.44	readline@8.2
automake@1.16.5	fmt@11.1.4	libidn2@2.3.7	pdt@3.25.2	sed@4.9
axom@0.10.1	fmt@11.1.4	libint@2.9.0	perl@5.40.0	slate@2024.10.29
bc@1.07.1	gasnet@2024.5.0	libmd@1.1.0	perl@5.40.0	slepc@3.22.2
berkeley-db@18.1.40	gcc-runtime@11.4.0	libmonitor@2023.03.15	perl-data-dumper@2.173	snappy@1.2.1
binutils@2.43.1	gdbm@1.23	libpciaccess@0.17	petsc@3.22.4	spiral-package-fftx@1.3.0
bison@3.8.2	gettext@0.23.1	libpng@1.6.39	petsc@3.22.4	spiral-package-jit@1.1.0
blaspp@2024.10.26	ginkgo@1.9.0	libsigsegv@2.14	pigz@2.8	spiral-package-mpi@1.1.0
blt@0.7.0	git@2.48.1	libsodium@1.0.20	pkgconf@2.3.0	spiral-package-simt@1.1.0
blt@0.7.0	glibc@2.35	libtool@2.4.7	protobuf@3.28.2	spiral-software@8.5.1
boost@1.86.0	gmake@4.4.1	libunistring@1.2	protobuf@3.29.3	sqlite@3.46.0
boost@1.86.0	gmp@6.3.0	libunwind@1.8.1	py-calver@2022.6.26	strumpack@8.0.0
boost@1.86.0	gperftools@2.15	libxc@7.0.0	py-certifi@2023.7.22	sundials@7.2.1
boost@1.86.0	gromacs@2024.4	libxcrypt@4.4.38	py-cffi@1.17.1	superlu-dist@9.1.0
boost@1.86.0	hdf5@1.8.23	libxml2@2.13.5	py-charset-normalizer@3.3.0	superlu-dist@9.1.0
bricks@2023.08.25	hdf5@1.14.5	libyaml@0.1.7	py-cython@3.0.11	sz@2.1.12.5
butterflypack@3.2.0	hdf5@1.14.5	libyaml@0.2.5	py-editables@0.5	tar@1.35
bzip2@1.0.8	heffte@2.4.1	libzmq@4.3.5	py-flit-core@3.10.1	tasmanian@8.1
c-blosc2@2.15.1	hpctoolkit@2024.01.1	lizard@2.0	py-fypp@3.1	tau@2.34.1
ca-certificates-mozilla@2025-02-25	hpcviewer@2025.01	llvm@19.1.7	py-hatchling@1.25.0	texinfo@7.1
cabana@0.7.0	hpx@1.10.0	lua@5.3.6	py-idna@3.4	trilinos@16.1.0
caliper@2.12.1	hwloc@2.11.1	lua@5.4.6	py-meson-python@0.16.0	umpire@2024.07.0
camp@2024.07.0	hwloc@2.11.1	lua-luaposix@36.1	py-numpy@2.2.4	umpire@2024.07.0
camp@2024.07.0	hypre@2.32.0	lz4@1.10.0	py-packaging@24.2	umpire@2024.07.0
chai@2024.07.0	hypre@2.32.0	lzo@2.10	py-pathspec@0.11.1	unzip@6.0
chapel@2.4.0	icu4c@74.2	m4@1.4.19	py-pip@24.3.1	upcxx@2023.9.0

PESO

https://e4s.io

E4S 25.06 image for NVIDIA GPUs (x86_64)

```
Singularity> spack find +cuda
-- linux-ubuntu22.04-x86_64_v3 / gcc@11.4.0 ----------------------
adios2@2.10.2       camp@2024.07.0    gromacs@2024.4        kokkos@4.6.01          nvcomp@2.2.0      strumpack@8.0.0     upcxx@2023.9.0
amrex@25.03         chai@2024.07.0    heffte@2.4.1          kokkos-kernels@4.6.01  papi@7.1.0        sundials@7.2.1      vtk-m@2.2.0
arborx@1.5          chapel@2.4.0      hpctoolkit@2024.01.1  lammps@20240829.1      parsec@3.0.2209   superlu-dist@9.1.0  zfp@1.0.0
axom@0.10.1         cp2k@2025.1       hpx@1.10.0            lapackpp@2024.10.26    petsc@3.22.4      tasmanian@8.1
blaspp@2024.10.26   cusz@0.14.0       hwloc@2.11.1          legion@24.12.0         petsc@3.22.4      tau@2.34.1
bricks@2023.08.25   fftx@1.2.0        hypre@2.32.0          libceed@0.12.0         raja@2024.07.0    trilinos@16.1.0
cabana@0.7.0        flecsi@2.3.2      kokkos@4.5.01         magma@2.9.0            raja@2024.07.0    umpire@2024.07.0
caliper@2.12.1      flux-core@0.67.0  kokkos@4.6.01         mfem@4.7.0             slate@2024.10.29  umpire@2024.07.0
camp@2024.07.0      ginkgo@1.9.0      kokkos@4.6.01         mgard@2023-12-09       slepc@3.22.2      umpire@2024.07.0
==> 57 installed packages
Singularity> nvcc --version
nvcc: NVIDIA (R) Cuda compiler driver
Copyright (c) 2005-2025 NVIDIA Corporation
Built on Fri_Feb_21_20:23:50_PST_2025
Cuda compilation tools, release 12.8, V12.8.93
Build cuda_12.8.r12.8/compiler.35583870_0
Singularity> which huggingface-cli
/usr/local/bin/huggingface-cli
Singularity> which firefox
/usr/bin/firefox
Singularity> which codium
/usr/bin/codium
Singularity> which jupyter
/usr/local/bin/jupyter
Singularity> nvidia-smi
Sat Jun  7 22:02:16 2025
+-----------------------------------------------------------------------------------------+
| NVIDIA-SMI 570.124.06             Driver Version: 570.124.06     CUDA Version: 12.8     |
|-----------------------------------------+------------------------+----------------------+
| GPU  Name                 Persistence-M | Bus-Id          Disp.A | Volatile Uncorr. ECC |
| Fan  Temp   Perf          Pwr:Usage/Cap |           Memory-Usage | GPU-Util  Compute M. |
|                                         |                        |               MIG M. |
|=========================================+========================+======================|
|   0  NVIDIA H100 PCIe               On  |   00000000:E1:00.0 Off |                    0 |
```

https://e4s.io

NVIDIA® BioNeMo™ Framework on E4S 25.06 CUDA x86_64

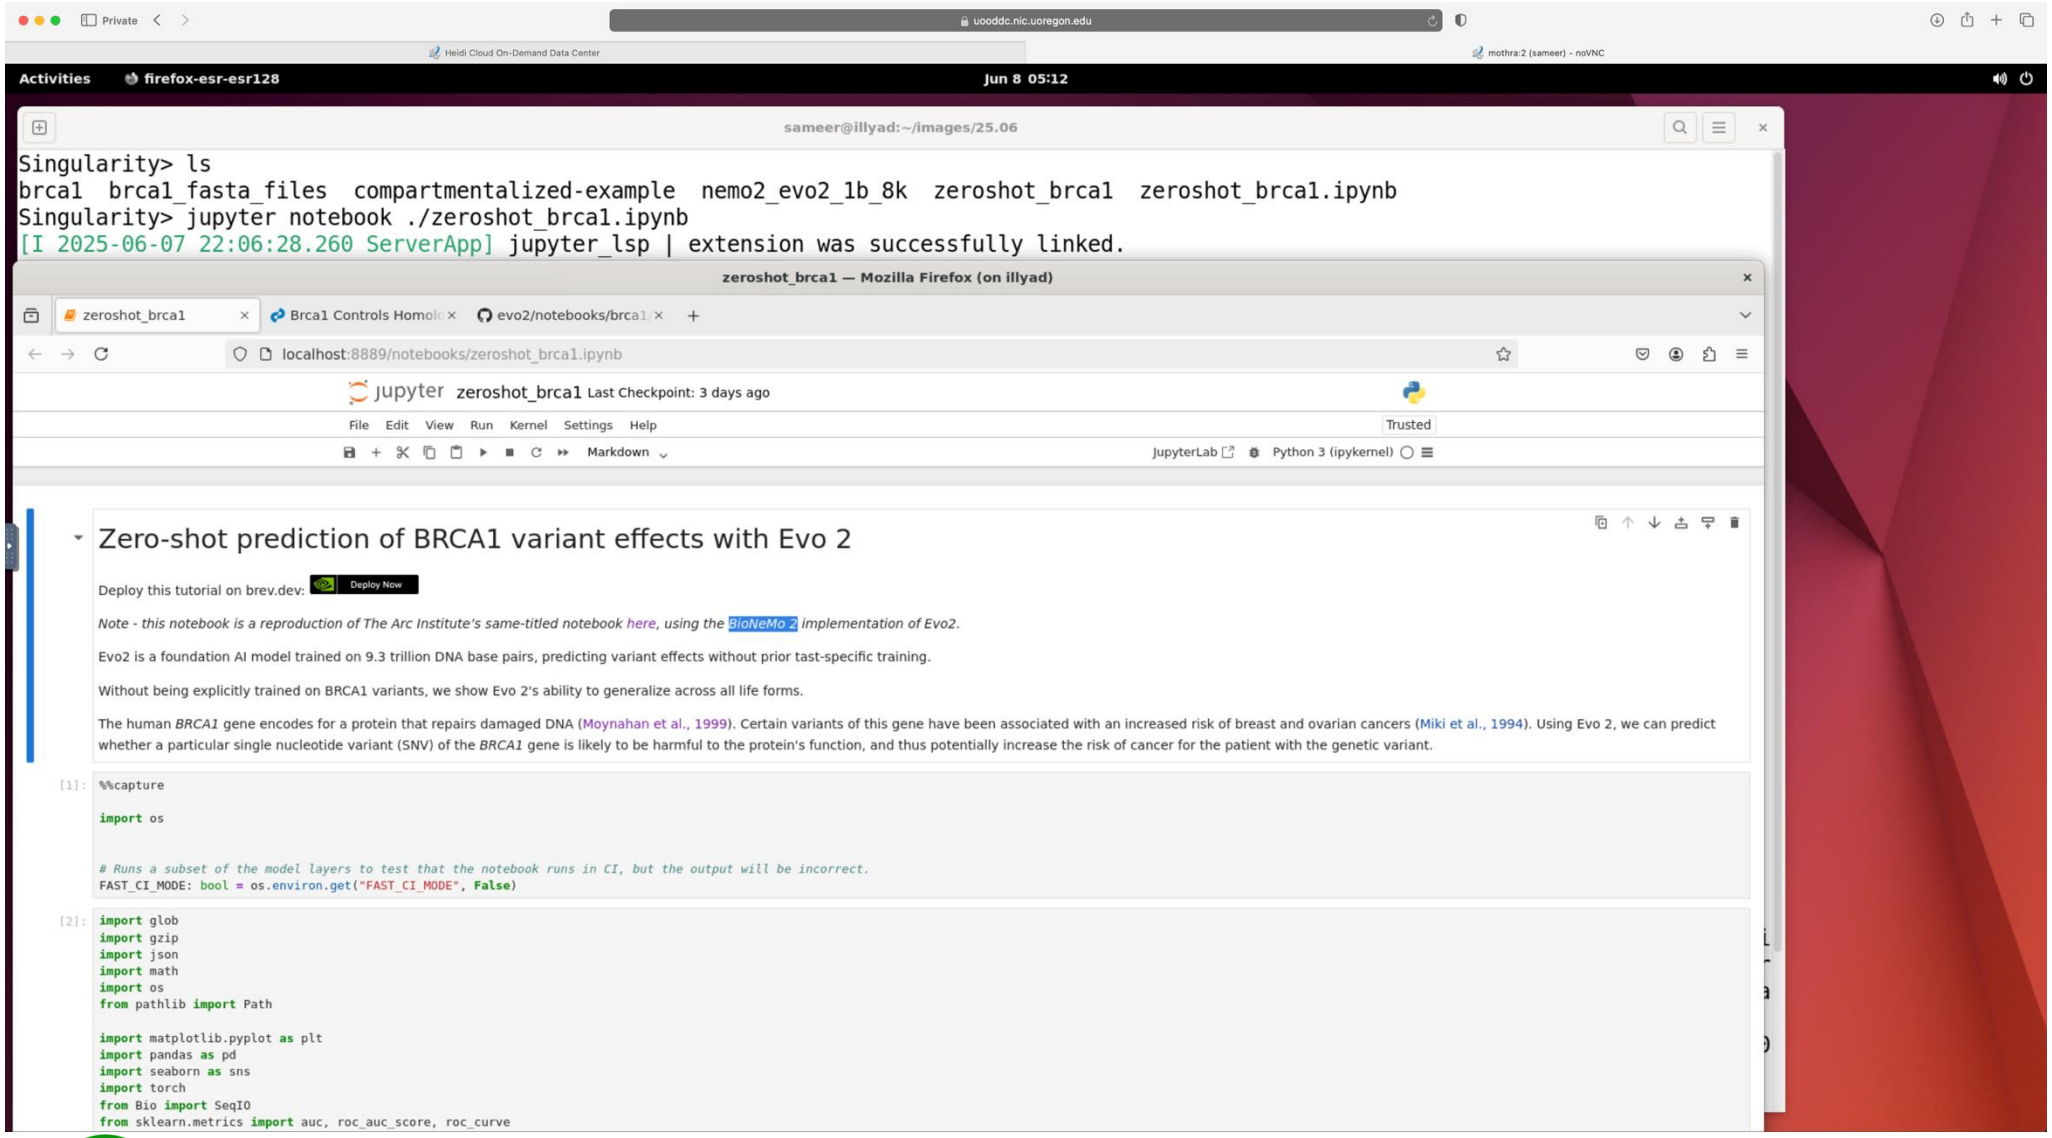

E4S 25.06 image for CUDA and x86_64 with VSCodium IDE

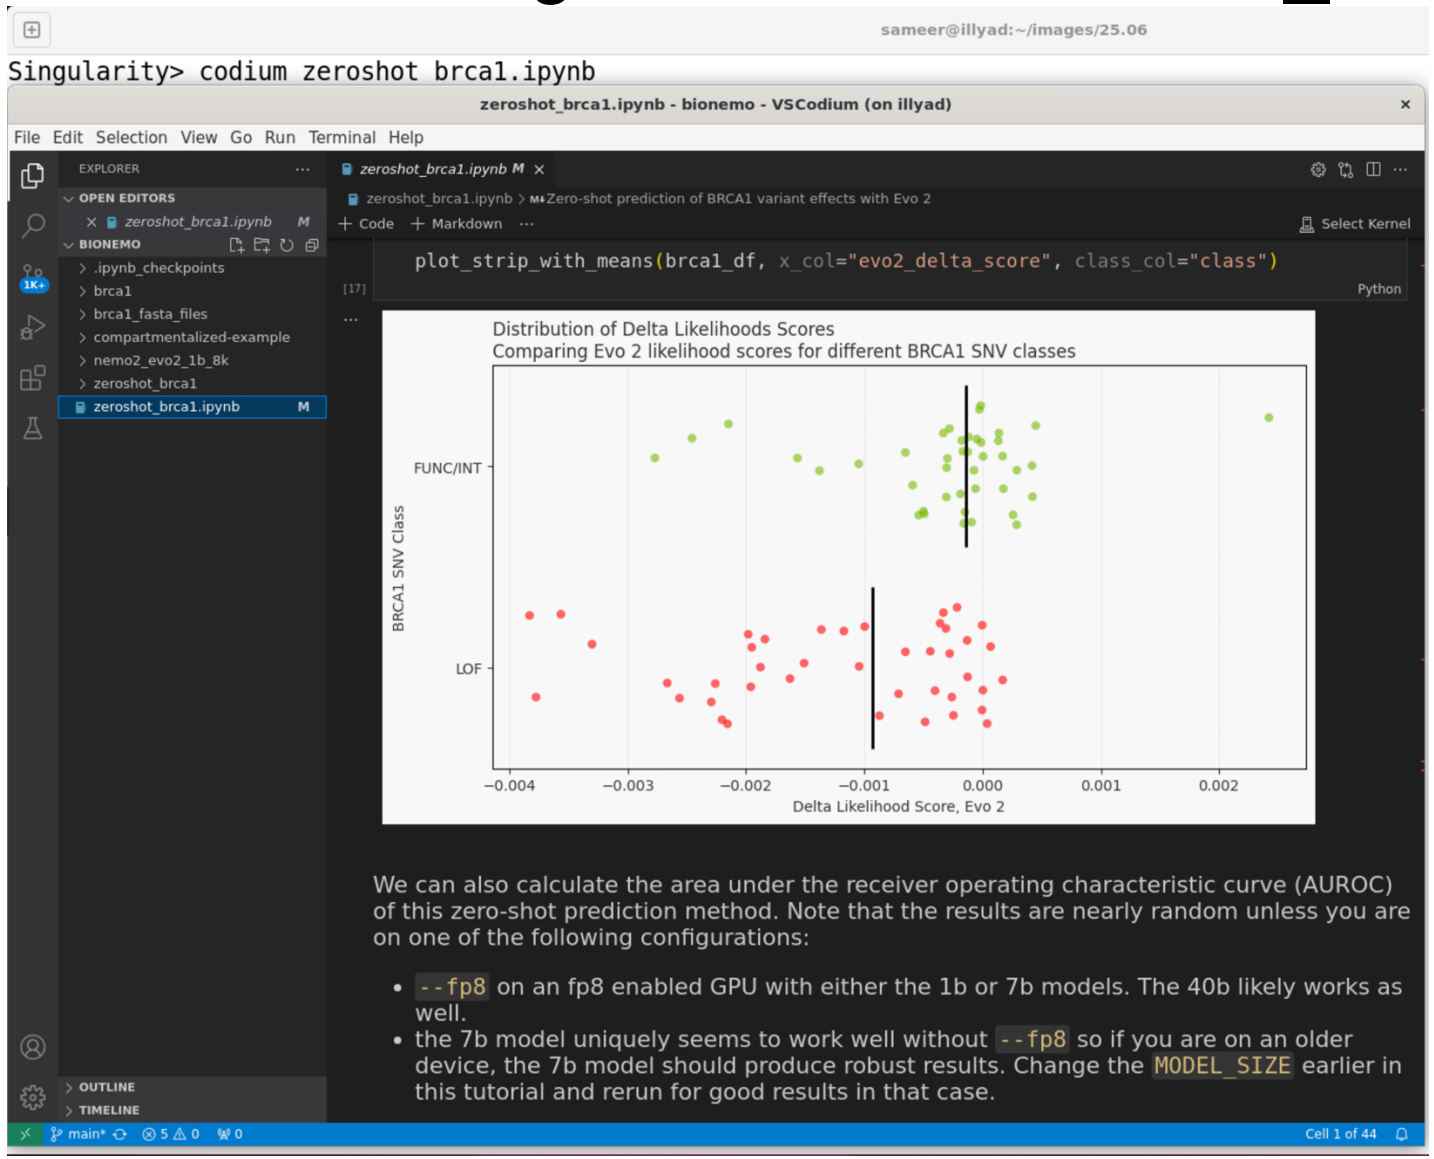

- NVIDIA H100 (cuda90) GPU on x86_64
- Jupyter Notebook in VSCodium IDE
- Running NVIDIA® BioNeMo™ Framework for biopharma workflows

https://e4s.io

E4S 25.06 image for CUDA and x86_64 with VSCodium IDE

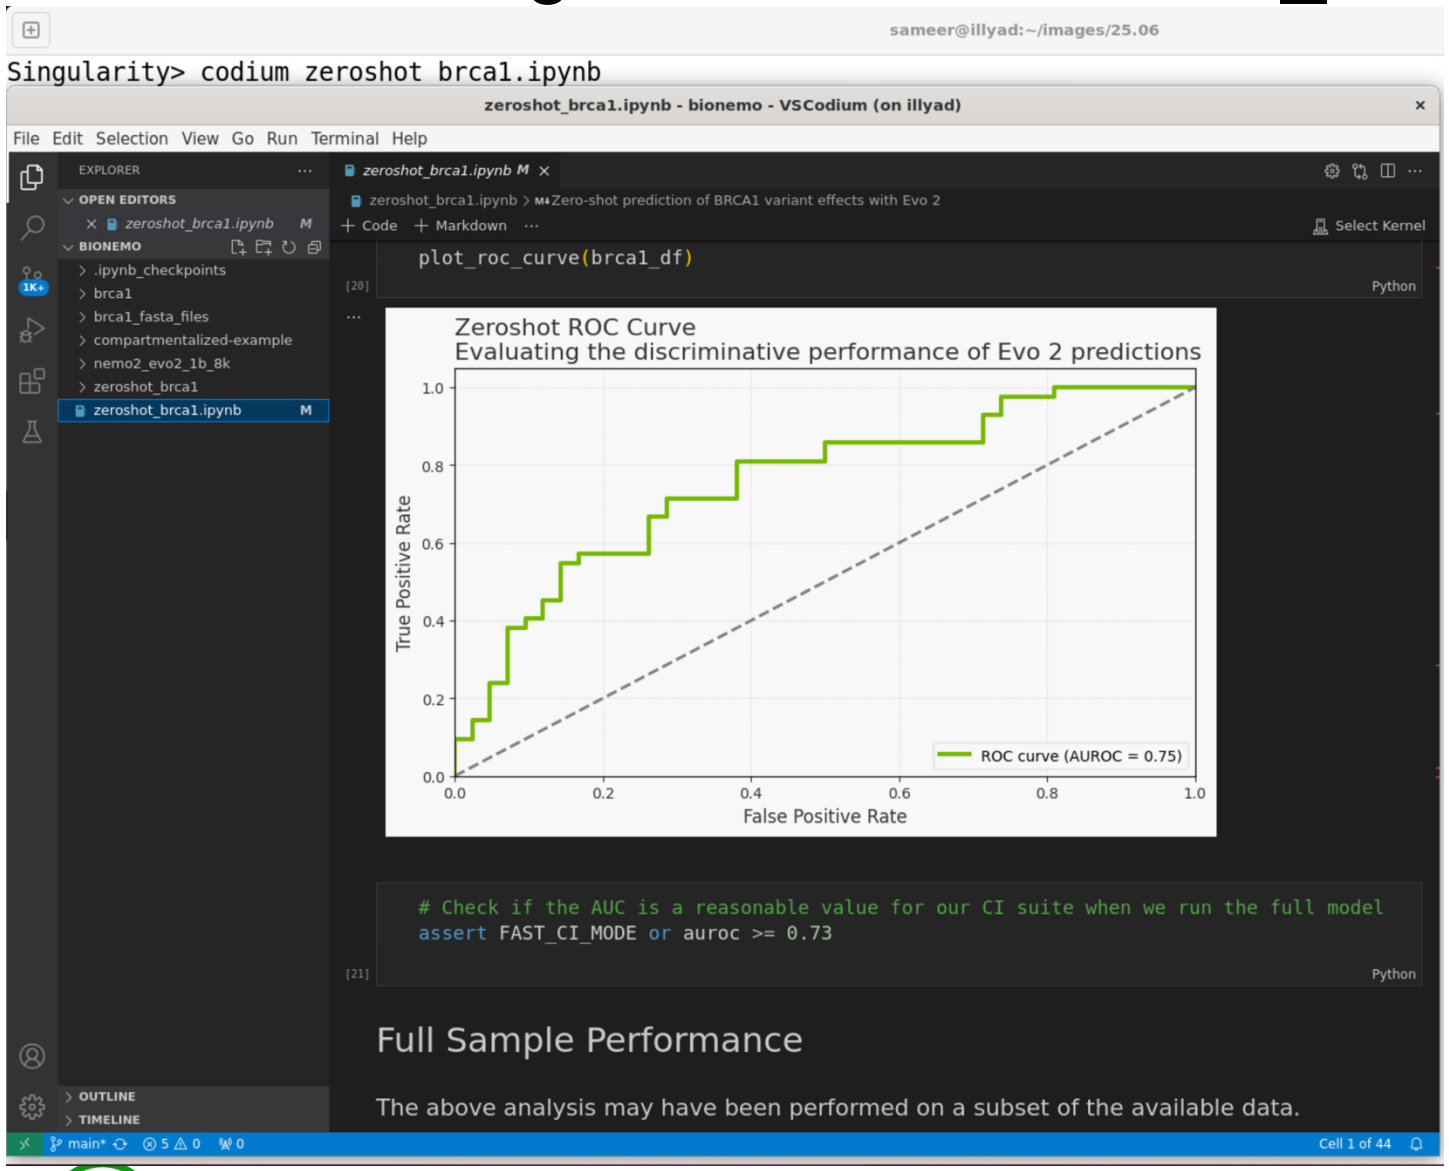

- NVIDIA H100 (cuda90) GPU on x86_64
- Jupyter Notebook in VSCodium IDE
- Running NVIDIA® BioNeMo™ Framework for biopharma workflows

https://e4s.io

Creating a Chatbot using Vllm using E4S 25.06 image for x86_64

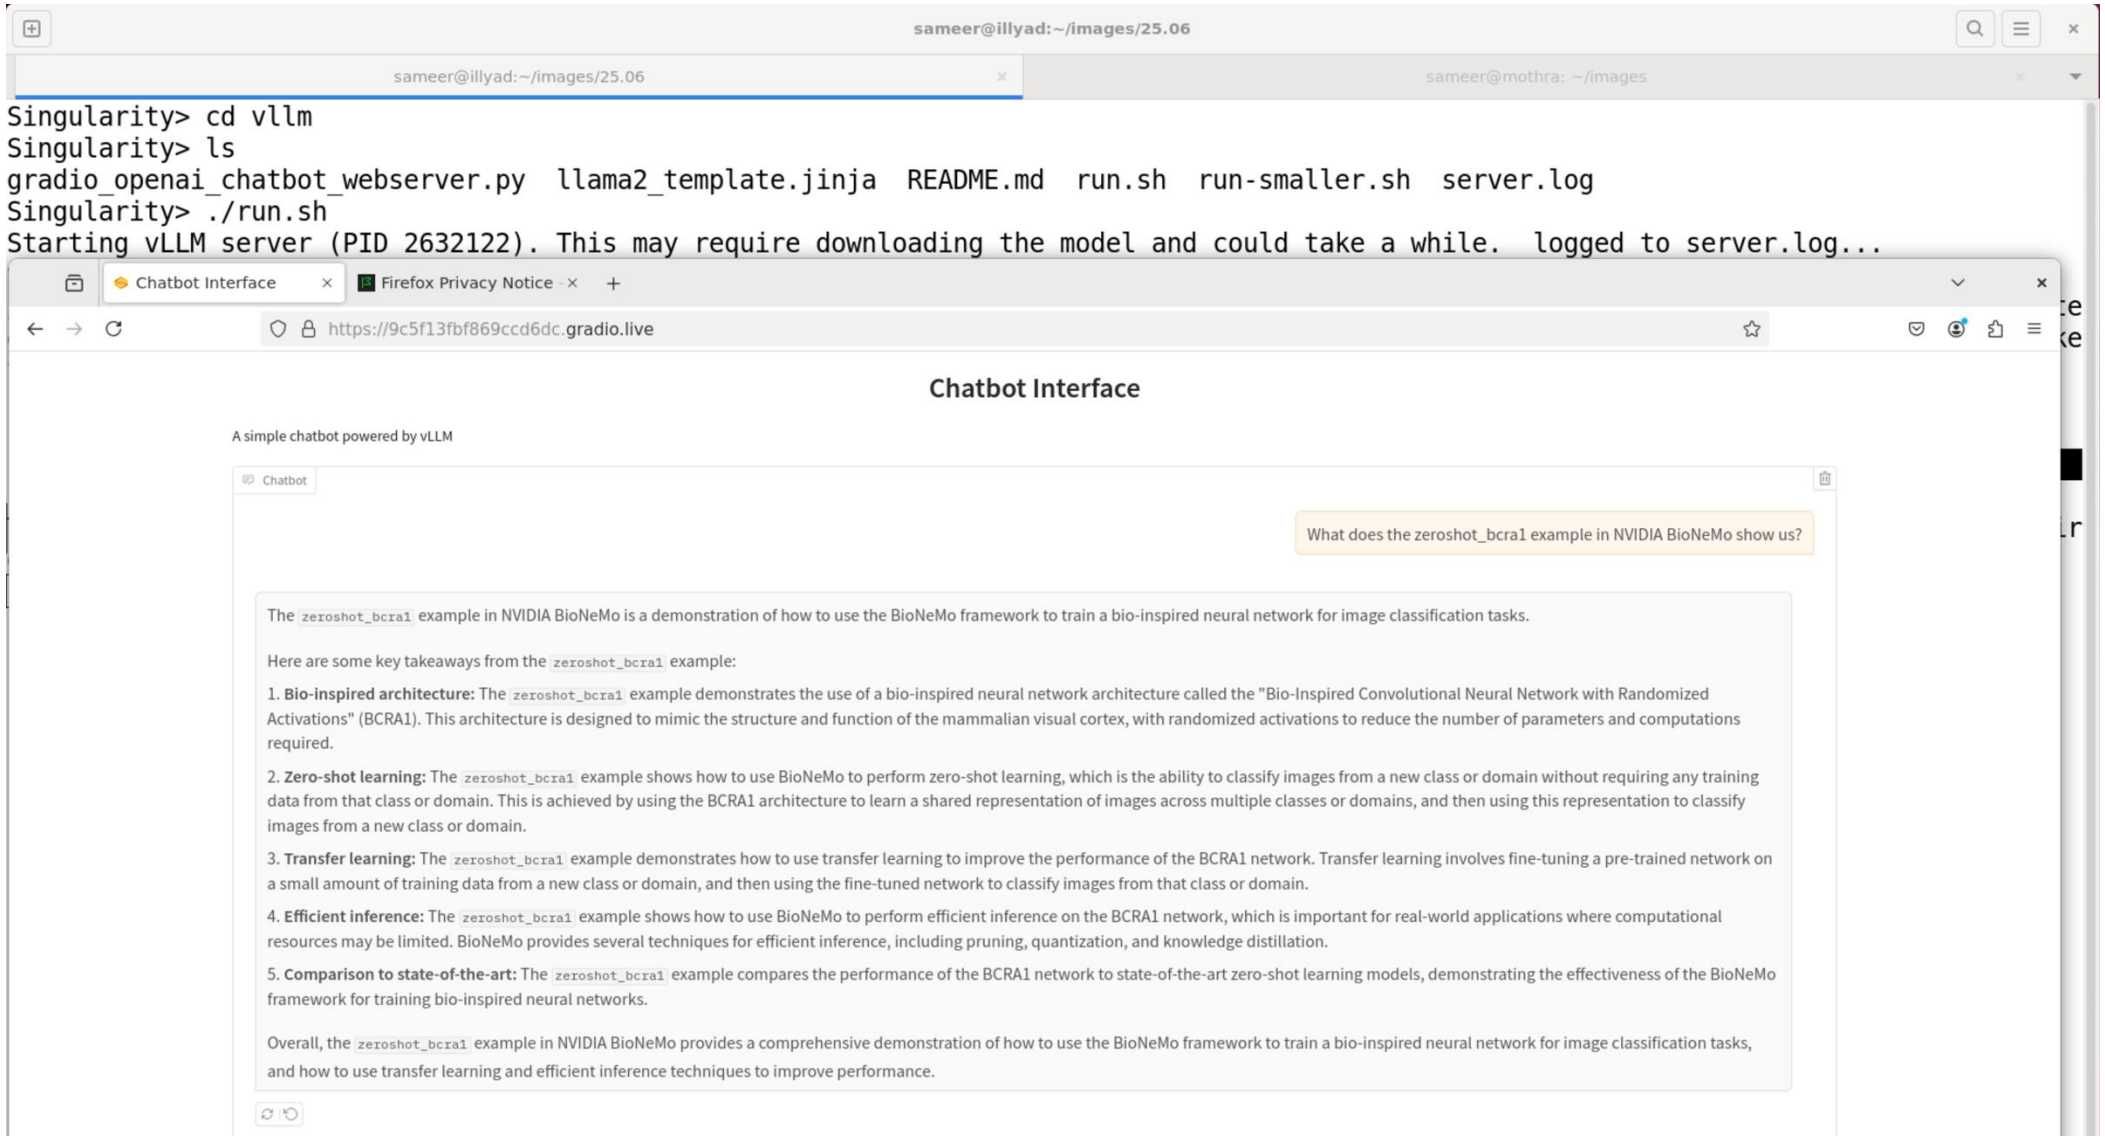

- NVIDIA H100 (cuda90) GPU on x86_64
- Vllm chatbot running after huggingface-cli login
- Using local H100 GPU

GPU accelerated 3D graphics using E4S 25.06 image for x86_64

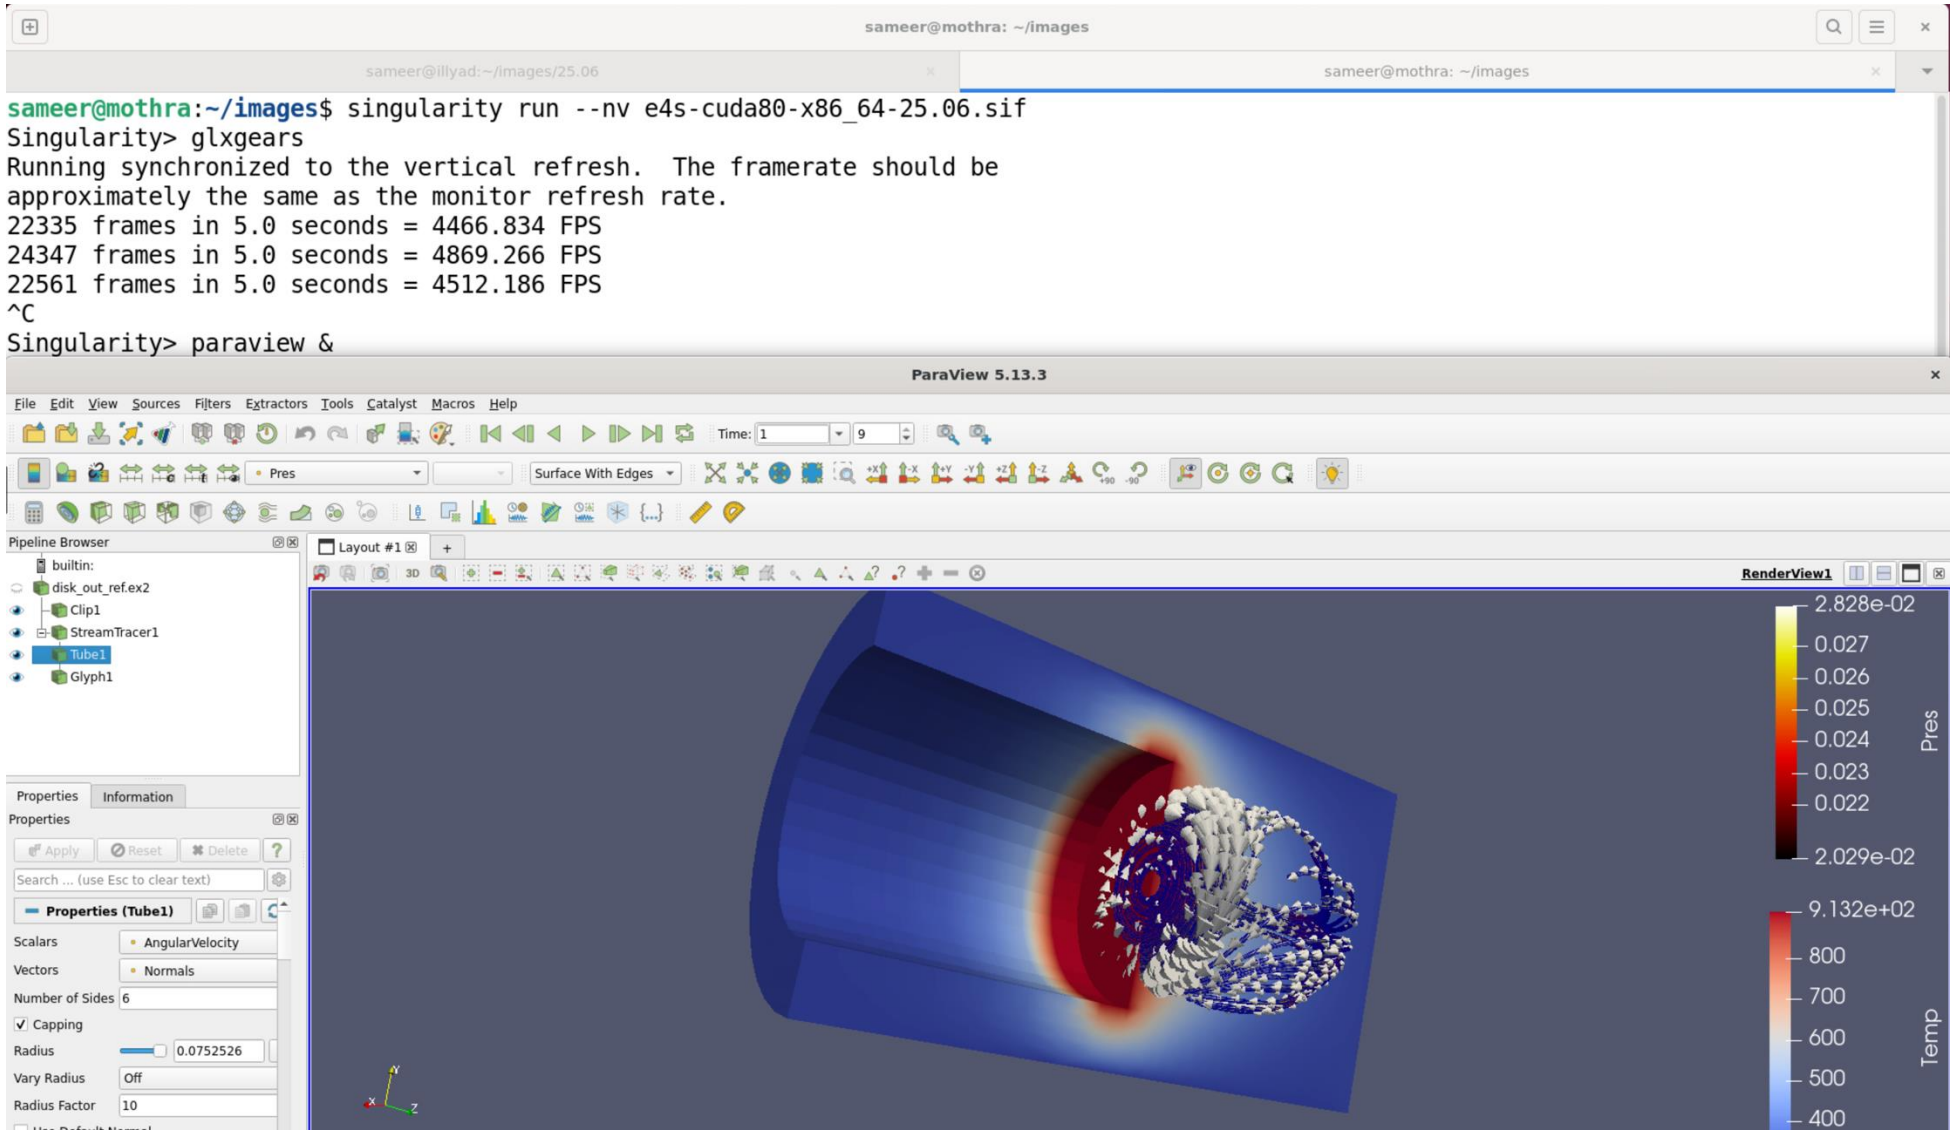

- Rendering on an NVIDIA A100 (cuda80) GPU on x86_64
- ParaView
- Using Adaptive Computing's Heidi/ODDC remote desktop

VisIt: GPU accelerated 3D graphics in E4S 25.06 image for x86_64

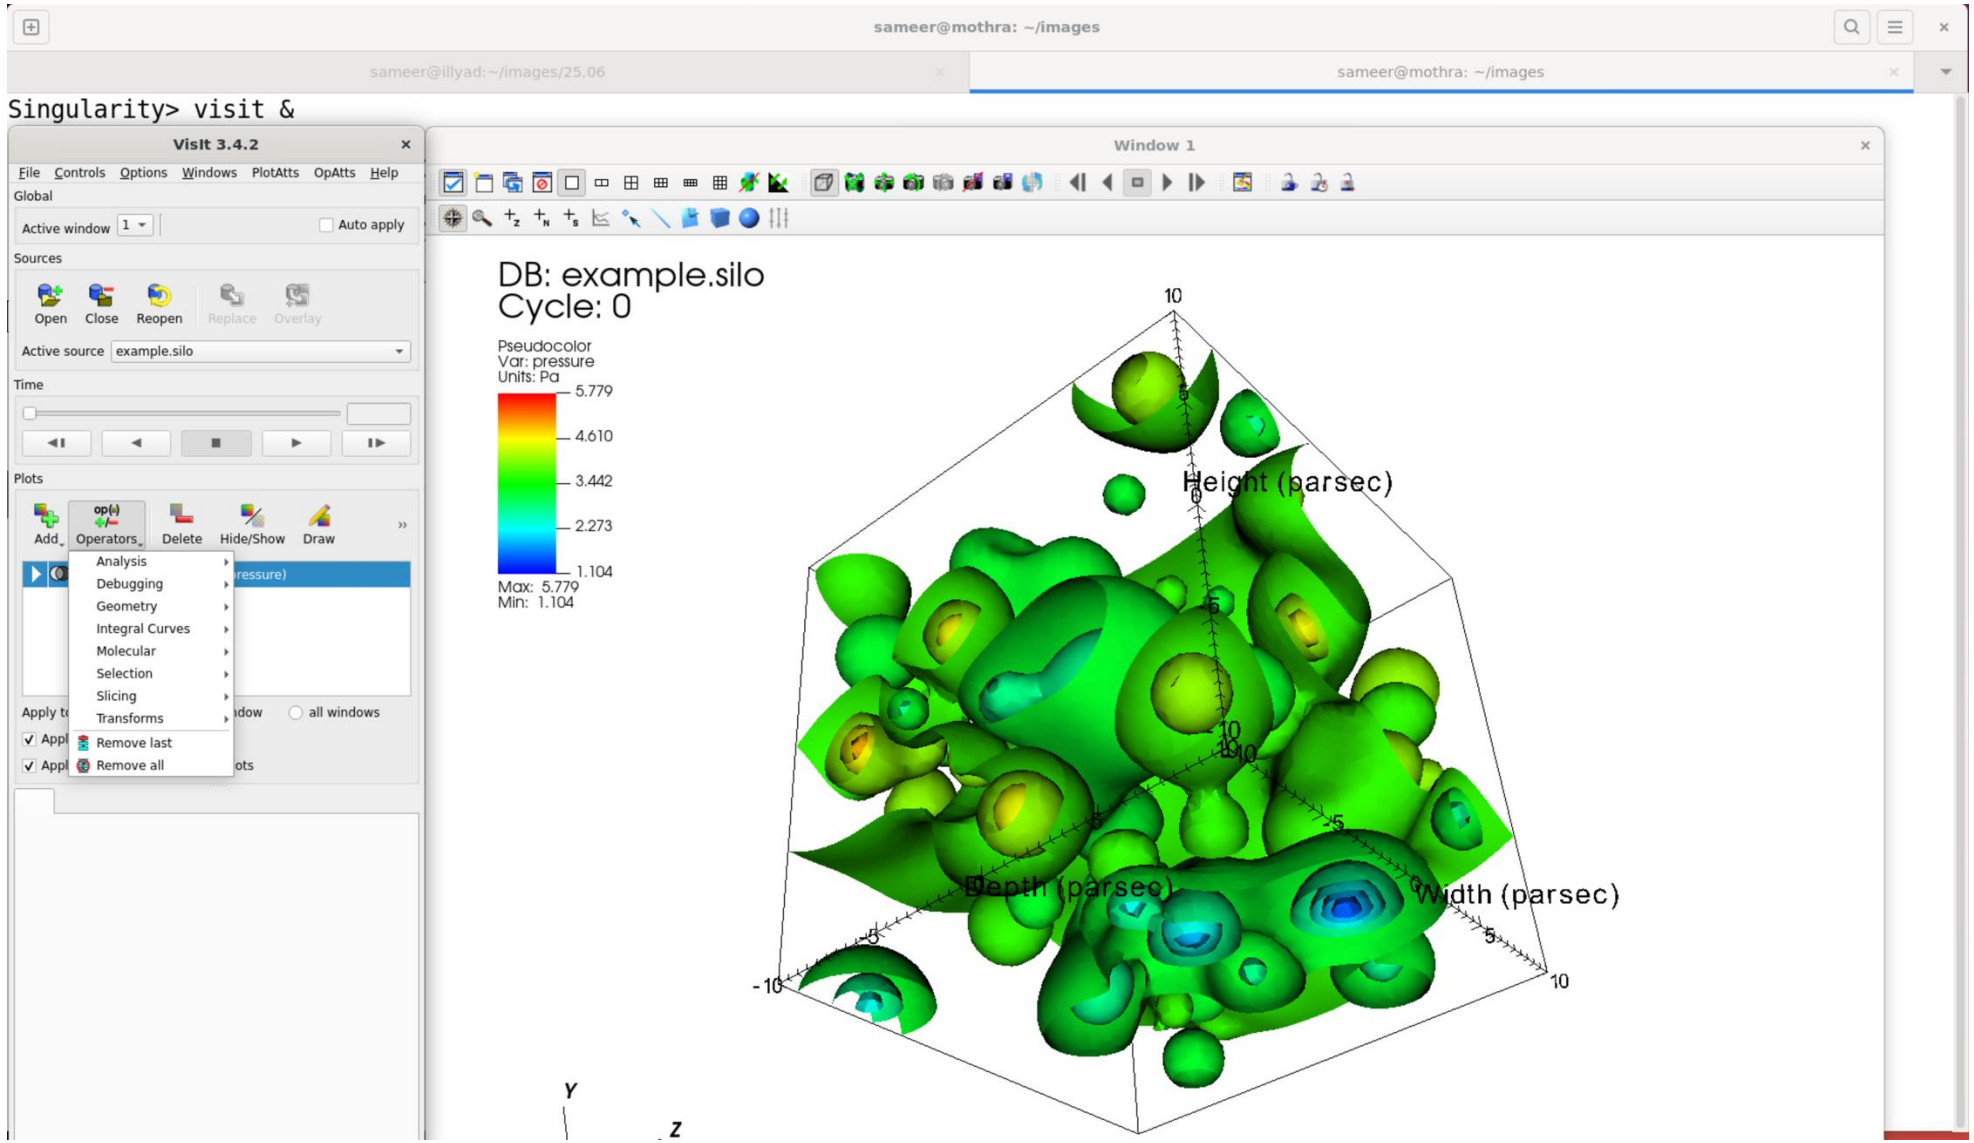

- Rendering on an NVIDIA A100 (cuda80) GPU on x86_64
- VisIt
- Using Adaptive Computing's Heidi/ODDC remote desktop

https://e4s.io

E4S 25.06 image for NVIDIA Grace-Hopper (cuda90, aarch64)

```
$ singularity run --nv e4s-cuda90-aarch64-25.06.sif
Singularity> uname -m
aarch64
Singularity> spack find +cuda
-- linux-ubuntu24.04-aarch64 / gcc@13.3.0 ----------------------
adios2@2.10.2      camp@2024.07.0    ginkgo@1.9.0          kokkos@4.6.01          libpressio@0.99.4      petsc@3.22.4         tasmanian@8.1     zfp@1.0.0
amrex@25.03        chai@2024.07.0    gromacs@2024.4        kokkos@4.6.01          magma@2.9.0            raja@2024.07.0       tau@2.34.1
arborx@1.5         chapel@2.4.0      heffte@2.4.1          kokkos@4.6.01          mfem@4.7.0             raja@2024.07.0       trilinos@16.1.0
axom@0.10.1        cp2k@2025.1       hpctoolkit@2024.01.1  kokkos-kernels@4.6.01  mgard@2023-12-09       slate@2024.10.29     umpire@2024.07.0
blaspp@2024.10.26  cusz@0.14.0       hpx@1.10.0            lammps@20240829.1      nvcomp@2.2.0           slepc@3.22.2         umpire@2024.07.0
cabana@0.7.0       fftx@1.2.0        hwloc@2.11.1          lapackpp@2024.10.26    omega-h@10.8.6-scorec  strumpack@8.0.0      umpire@2024.07.0
caliper@2.12.1     flecsi@2.3.2      hypre@2.32.0          legion@24.12.0         parsec@3.0.2209        sundials@7.2.1       upcxx@2023.9.0
camp@2024.07.0     flux-core@0.67.0  kokkos@4.5.01         libceed@0.12.0         petsc@3.22.4           superlu-dist@9.1.0   vtk-m@2.2.0
==> 57 installed packages
Singularity> nvcc --version
nvcc: NVIDIA (R) Cuda compiler driver
Copyright (c) 2005-2025 NVIDIA Corporation
Built on Fri_Feb_21_20:26:18_PST_2025
Cuda compilation tools, release 12.8, V12.8.93
Build cuda_12.8.r12.8/compiler.35583870_0
Singularity>
Singularity> nvidia-smi
Sat Jun  7 22:35:36 2025
+-----------------------------------------------------------------------------------------+
| NVIDIA-SMI 570.124.06             Driver Version: 570.124.06     CUDA Version: 12.8     |
|-----------------------------------------+------------------------+----------------------+
| GPU  Name                 Persistence-M | Bus-Id          Disp.A | Volatile Uncorr. ECC |
| Fan  Temp   Perf          Pwr:Usage/Cap |           Memory-Usage | GPU-Util  Compute M. |
|                                         |                        |               MIG M. |
|=========================================+========================+======================|
|   0  NVIDIA GH200 480GB             On  |   00000009:01:00.0 Off |                    0 |
| N/A   27C    P0            122W /  700W |   89270MiB /  97871MiB |      1%      Default |
|                                         |                        |             Disabled |
+-----------------------------------------+------------------------+----------------------+
```

https://e4s.io

E4S 25.06 image for x86_64 for AMD MI300A GPU (ROCm 6.3.3)

```
$ singularity run e4s-rocm942-x86_64-25.06.sif
Singularity> uname -m
x86_64
Singularity> lscpu | grep AMD
Vendor ID:                            AuthenticAMD
Model name:                           AMD Instinct MI300A Accelerator
Virtualization:                       AMD-V
Singularity> ls /opt
amdgpu  demo  rocm  rocm-6.3.3
Singularity> spack find +rocm
-- linux-ubuntu22.04-x86_64_v3 / gcc@11.4.0 ----------------------
amrex@25.03        chai@2024.07.0            hip@6.3.2        hpctoolkit@2024.01.1  kokkos@4.6.01        petsc@3.22.4      superlu-dist@9.1.0  vtk-m@2.2.0
arborx@1.7         chapel@2.4.0              hipblas@6.3.2    hpx@1.10.0            lapackpp@2024.10.26  raja@2024.07.0    tasmanian@8.1
blaspp@2024.10.26  fftx@1.2.0                hipfft@6.3.2     hypre@2.32.0          legion@24.12.0       slate@2024.10.29  tau@2.34.1
cabana@0.7.0       gasnet@2025.2.0-snapshot  hiprand@6.3.2    kokkos@3.7.01         libceed@0.12.0       slepc@3.22.2      trilinos@16.1.0
caliper@2.12.1     ginkgo@1.9.0              hipsolver@6.3.2  kokkos@4.5.01         mfem@4.7.0           strumpack@8.0.0   umpire@2024.07.0
camp@2024.07.0     heffte@2.4.1              hipsparse@6.3.2  kokkos@4.6.01         paraview@5.13.2      sundials@7.2.1    upcxx@2023.9.0
==> 43 installed packages
Singularity> which huggingface-cli
/usr/local/bin/huggingface-cli
Singularity> which codium
/usr/bin/codium
Singularity> which jupyter
/usr/local/bin/jupyter
Singularity> python
Python 3.10.12 (main, Feb  4 2025, 14:57:36) [GCC 11.4.0] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>> import torch
>>> torch.cuda.get_arch_list()
['gfx900', 'gfx906', 'gfx908', 'gfx90a', 'gfx942', 'gfx1030', 'gfx1100', 'gfx1101', 'gfx1102', 'gfx1200', 'gfx1201']
>>> torch.__version__
'2.7.1+rocm6.3'
>>> import vllm
INFO 06-07 22:39:54 [__init__.py:248] Automatically detected platform rocm.
>>>
```

E4S 25.06 image for CPU x86_64 without GPU runtimes (GCC)

```
$ singularity run e4s-cpu-x86_64-25.06.sif
Singularity> uname -m
x86_64
Singularity> spack find -x
-- linux-ubuntu22.04-x86_64_v3 / gcc@11.4.0 ----------------------
adios@1.13.1            dyninst@13.0.0            hdf5-vol-async@1.7     mfem@4.7.0              petsc@3.22.4            sundials@7.2.1
alquimia@1.1.0          e4s-alc@1.0.2             hdf5-vol-cache@v1.1    mgard@2023-12-09        phist@1.12.1            superlu@7.0.0
aml@0.2.1               e4s-cl@1.0.4              hdf5-vol-log@1.4.0     mpark-variant@1.4.0     plasma@24.8.7           superlu-dist@9.1.0
amrex@25.03             exago@1.6.0               heffte@2.4.1           mpifileutils@0.12       plumed@2.9.2            swig@4.0.2-fortran
arborx@1.7              exaworks@0.1.0            hpctoolkit@2024.01.1   nccmp@1.9.1.0           precice@3.1.2           sz3@3.2.0
argobots@1.2            fftx@1.2.0                hpx@1.10.0             nco@5.3.2               pruners-ninja@1.0.1     tasmanian@8.1
axom@0.10.1             flecsi@2.3.2              hypre@2.32.0           nek5000@19.0            pumi@2.2.9              tau@2.34.1
bolt@2.0                flit@2.1.0                kokkos@4.6.01          nekbone@17.0            py-h5py@3.12.1          trilinos@16.1.0
boost@1.87.0            flux-core@0.67.0          kokkos-kernels@4.6.01  netcdf-fortran@4.6.1    py-jupyterhub@1.4.1     turbine@1.3.0
bricks@2023.08.25       fortrilinos@2.3.0         laghos@3.1             netlib-scalapack@2.2.2  py-libensemble@1.4.3    umap@2.1.1
butterflypack@3.2.0     fpm@0.10.0                lammps@20240829.1      nrm@0.1.0               py-petsc4py@3.22.4      umpire@2024.07.0
cabana@0.7.0            gasnet@2025.2.0-snapshot  lbann@0.104            nwchem@7.2.3            qthreads@1.18           upcxx@2023.9.0
caliper@2.12.1          ginkgo@1.9.0              legion@24.12.0         omega-h@10.8.6-scorec   quantum-espresso@7.4.1  variorum@0.8.0
chai@2024.07.0          globalarrays@5.8.2        libceed@0.12.0         openfoam@2412           raja@2024.07.0          wannier90@3.1.0
chapel@2.4.0            glvis@4.2                 libnrm@0.1.0           openmpi@5.0.6           rempi@1.1.0             warpx@25.03
charliecloud@0.38       gmp@6.3.0                 libquo@1.4             openpmd-api@0.16.1      scr@3.1.0               wps@4.5
conduit@0.9.3           gotcha@1.0.7              libunwind@1.8.1        papi@7.1.0              slate@2024.10.29        wrf@4.6.1
cp2k@2025.1             gptune@4.0.0              loki@0.1.7             papyrus@1.0.2           slepc@3.22.2            xyce@7.8.0
datatransferkit@3.1.1   gromacs@2024.4            mercury@2.4.0          parsec@3.0.2209         stc@0.9.0
dealii@9.6.2            h5bench@1.4               metall@0.30            pdt@3.25.2              strumpack@8.0.0
==> 118 installed packages
Singularity>
```

E4S 25.06 image for Intel GPUs with Intel compilers (x86_64)

```
$ singularity run e4s-oneapi-x86_64-25.06.sif
Singularity> uname -m
x86_64
Singularity> clinfo -l
Platform #0: Intel(R) OpenCL
 `-- Device #0: Intel(R) Xeon(R) Silver 4410T
Platform #1: Intel(R) OpenCL Graphics
 `-- Device #0: Intel(R) Data Center GPU Max 1100
Singularity> spack find -x
-- linux-ubuntu22.04-x86_64_v3 / oneapi@2025.1.0 ----------------
adios@1.13.1          darshan-runtime@3.4.6     hdf5-vol-cache@v1.1    mgard@2023-12-09        precice@3.1.2          sz@2.1.12.5
adios2@2.10.2         darshan-util@3.4.6        hdf5-vol-log@1.4.0     mpark-variant@1.4.0     pruners-ninja@1.0.1    sz3@3.2.0
alquimia@1.1.0        datatransferkit@3.1.1     heffte@2.4.1           mpifileutils@0.12       pumi@2.2.9             tasmanian@8.1
aml@0.2.1             e4s-alc@1.0.2             heffte@2.4.1           nccmp@1.9.1.0           py-cinemasci@1.7.0     tau@2.34.1
aml@0.2.1             e4s-cl@1.0.4              hpctoolkit@2024.01.1   nco@5.3.2               py-h5py@3.12.1         tau@2.34.1
amrex@25.03           exaworks@0.1.0            hpx@1.10.0             nekbone@17.0            py-jupyterhub@1.4.1    trilinos@16.1.0
amrex@25.03           faodel@1.2108.1           hypre@2.32.0           netcdf-fortran@4.6.1    py-libensemble@1.4.3   turbine@1.3.0
arborx@1.7            flecsi@2.3.2              kokkos@4.6.01          netlib-scalapack@2.2.2  py-petsc4py@3.22.4     umap@2.1.1
arborx@1.7            flit@2.1.0                kokkos@4.6.01          nrm@0.1.0               qthreads@1.18          umpire@2024.07.0
argobots@1.2          flux-core@0.67.0          kokkos-kernels@4.6.01  nwchem@7.2.3            quantum-espresso@7.4.1 unifyfs@2.0
ascent@0.9.3          fortrilinos@2.3.0         laghos@3.1             omega-h@10.8.6-scorec   raja@2024.07.0         upcxx@2023.9.0
axom@0.10.1           gasnet@2025.2.0-snapshot  lammps@20240829.1      openfoam@2412           rempi@1.1.0            upcxx@2023.9.0
bolt@2.0              ginkgo@1.9.0              legion@24.12.0         openmpi@5.0.6           scr@3.1.0              variorum@0.8.0
boost@1.87.0          ginkgo@1.9.0              libcatalyst@2.0.0      openpmd-api@0.16.1      slate@2024.10.29       veloc@1.7
bricks@2023.08.25     globalarrays@5.8.2        libceed@0.12.0         papi@7.1.0              slepc@3.22.2           vtk-m@2.2.0
butterflypack@3.2.0   gmp@6.3.0                 libnrm@0.1.0           papyrus@1.0.2           stc@0.9.0              wannier90@3.1.0
cabana@0.7.0          gotcha@1.0.7              libquo@1.4             parallel-netcdf@1.14.0  strumpack@8.0.0        wrf@4.6.1
cabana@0.7.0          gptune@4.0.0              libunwind@1.8.1        parsec@3.0.2209         sundials@7.2.1         xyce@7.8.0
caliper@2.12.1        gromacs@2024.4            loki@0.1.7             pdt@3.25.2              sundials@7.2.1         zfp@1.0.0
chai@2024.07.0        h5bench@1.4               mercury@2.4.0          petsc@3.22.4            superlu@7.0.0
charliecloud@0.38     hdf5@1.14.5               metall@0.30            phist@1.12.1            superlu-dist@9.1.0
conduit@0.9.3         hdf5-vol-async@1.7        mfem@4.7.0             plumed@2.9.2            swig@4.0.2-fortran
==> 129 installed packages
Singularity> which ifx
/opt/intel/oneapi/compiler/2025.1/bin/ifx
```

PESO

https://e4s.io

E4S 25.06 image for Intel® oneAPI built with Intel MPI and compilers

```
Singularity> spack find -dl -v amrex +sycl
-- linux-ubuntu22.04-x86_64_v3 / oneapi@2025.1.0 ----------------
vibmei2 amrex@25.03~amrdata~ascent~catalyst~conduit~cuda+eb+fft~fortran~hdf5~hypre~ipo+linear_solvers+mpi~openmp~particles~petsc~pic~plotfile_tools~rocm~sha
red~sundials+sycl~tiny_profile build_system=cmake build_type=Release dimensions=1,2,3 generator=make precision=double
pkhuomn     cmake@3.31.6~doc+ncurses+ownlibs~qtgui build_system=generic build_type=Release
2kk4mv5         curl@8.11.1~gssapi~ldap~libidn2~librtmp~libssh~libssh2+nghttp2 build_system=autotools libs=shared,static tls=openssl
vizwdiu             nghttp2@1.65.0 build_system=autotools
izwrsao                 diffutils@3.10 build_system=autotools
y2xtcez                     libiconv@1.17 build_system=autotools libs=shared,static
avu5xay             openssl@3.4.1~docs+shared build_system=generic certs=mozilla
v4unosi                 ca-certificates-mozilla@2025-02-25 build_system=generic
27y67gw                 perl@5.40.0+cpanm+opcode+open+shared+threads build_system=generic
mie6fpa                     berkeley-db@18.1.40+cxx~docs+stl build_system=autotools patches=26090f4,b231fcc
eeheopa                     bzip2@1.0.8~debug~pic+shared build_system=generic
cpcojya                     gdbm@1.23 build_system=autotools
kud5giw                         readline@8.2 build_system=autotools patches=1ea4349,24f587b,3d9885e,5911a5b,622ba38,6c8adf8,758e2ec,79572ee,a177edc,bbf97f1,
c7b45ff,e0013d9,e065038
rv3c6cm         ncurses@6.5~symlinks+termlib abi=none build_system=autotools patches=7a351bc
jzthfaz         zlib-ng@2.2.3+compat+new_strategies+opt+pic+shared build_system=autotools
7qebqtu     glibc@2.35 build_system=autotools
nx5zgqc     gmake@4.4.1~guile build_system=generic
lzycjnj     intel-oneapi-mkl@2025.1.0~cluster+envmods~gfortran~ilp64+shared build_system=generic mpi_family=none threads=none
nxlzht7     intel-oneapi-mpi@2021.15.0~classic-names+envmods~external-libfabric~generic-names~ilp64 build_system=generic
3obsopm     intel-oneapi-runtime@2025.1.0 build_system=generic
ncwe67l         gcc-runtime@11.4.0 build_system=generic
2ianxzp     pkgconf@2.3.0 build_system=autotools

==> 1 installed package
Singularity>
```

Intel oneAPI based packages in E4S 25.06 image for x86_64

```
$ singularity run e4s-oneapi-x86_64-25.06.sif
Singularity> module avail

------------------------------- /spack/share/spack/lmod/linux-ubuntu22.04-x86_64/intel-oneapi-mpi/2021.15.0-nxlzht7/Core --------------------------------
   adios/1.13.1                     ginkgo/1.9.0-sycl-openmp (D)    nccmp/1.9.1.0                   rempi/1.1.0
   adios2/2.10.2                    globalarrays/5.8.2              nco/5.3.2-openmp                scr/3.1.0
   alquimia/1.1.0                   gptune/4.0.0                    nekbone/17.0                    slate/2024.10.29-openmp
   amrex/25.03-sycl                 gromacs/2024.4-openmp           netcdf-fortran/4.6.1            slepc/3.22.2
   amrex/25.03                (D)   h5bench/1.4                     netlib-scalapack/2.2.2          stc/0.9.0
   arborx/1.7-sycl                  hdf5-vol-async/1.7              nwchem/7.2.3                    strumpack/8.0.0-openmp
   arborx/1.7                 (D)   hdf5-vol-cache/v1.1             omega-h/10.8.6-scorec           sundials/7.2.1-sycl
   ascent/0.9.3-openmp              hdf5-vol-log/1.4.0              openfoam/2412                   sundials/7.2.1             (D)
   axom/0.10.1-openmp               hdf5/1.14.5                     openpmd-api/0.16.1              superlu-dist/9.1.0
   bricks/2023.08.25                heffte/2.4.1-sycl               papyrus/1.0.2                   tasmanian/8.1
   butterflypack/3.2.0-openmp       heffte/2.4.1               (D)  parallel-netcdf/1.14.0          tau/2.34.1-level-zero
   cabana/0.7.0-sycl                hpctoolkit/2024.01.1            parsec/3.0.2209                 tau/2.34.1                 (D)
   cabana/0.7.0               (D)   hpx/1.10.0                      petsc/3.22.4                    trilinos/16.1.0
   caliper/2.12.1                   hypre/2.32.0                    phist/1.12.1-openmp             turbine/1.3.0
   chai/2024.07.0                   laghos/3.1                      plumed/2.9.2                    umpire/2024.07.0
   conduit/0.9.3                    lammps/20240829.1-openmp        precice/3.1.2                   unifyfs/2.0
   darshan-runtime/3.4.6            libcatalyst/2.0.0               pruners-ninja/1.0.1             upcxx/2023.9.0-level-zero
   datatransferkit/3.1.1            libnrm/0.1.0                    pumi/2.2.9                      upcxx/2023.9.0             (D)
   exaworks/0.1.0                   libquo/1.4                      py-cinemasci/1.7.0              veloc/1.7
   faodel/1.2108.1                  mercury/2.4.0                   py-h5py/3.12.1                  vtk-m/2.2.0
   flecsi/2.3.2                     metall/0.30                     py-libensemble/1.4.3            wannier90/3.1.0
   fortrilinos/2.3.0                mfem/4.7.0                      py-petsc4py/3.22.4              wrf/4.6.1
   ginkgo/1.9.0-openmp              mpifileutils/0.12               quantum-espresso/7.4.1-openmp   xyce/7.8.0

------------------------------------------------- /spack/share/spack/lmod/linux-ubuntu22.04-x86_64/Core -------------------------------------------------
   aml/0.2.1-level-zero        e4s-cl/1.0.4                      kokkos/4.6.01-openmp             nrm/0.1.0              swig/4.0.2-fortran
   aml/0.2.1             (D)   flit/2.1.0                        kokkos/4.6.01-sycl-openmp (D)    openmpi/5.0.6          sz/2.1.12.5
   argobots/1.2                flux-core/0.67.0                  legion/24.12.0                   papi/7.1.0             sz3/3.2.0
   bolt/2.0                    gasnet/2025.2.0-snapshot          libceed/0.12.0                   pdt/3.25.2             umap/2.1.1
   boost/1.87.0                gmp/6.3.0                         libunwind/1.8.1                  py-jupyterhub/1.4.1    variorum/0.8.0
   charliecloud/0.38           gotcha/1.0.7                      loki/0.1.7                       qthreads/1.18          zfp/1.0.0
   darshan-util/3.4.6          intel-oneapi-mpi/2021.15.0   (L)  mgard/2023-12-09-openmp          raja/2024.07.0
   e4s-alc/1.0.2               kokkos-kernels/4.6.01-openmp      mpark-variant/1.4.0              superlu/7.0.0

------------------------------------------------------------ /opt/intel/oneapi/modulefiles -------------------------------------------------------------
   compiler-intel-llvm/latest          compiler/latest          dev-utilities/latest          mkl/latest           tbb/latest
   compiler-intel-llvm/2025.1.1 (D)    compiler/2025.1.1 (D)    dev-utilities/2025.1.0 (D)    mkl/2025.1  (D)      tbb/2022.1 (D)
   compiler-rt/latest                  debugger/latest          dpl/latest                    mpi/latest           umf/latest
   compiler-rt/2025.1.1         (D)    debugger/2025.1.0 (D)    dpl/2022.8             (D)    mpi/2021.15 (L,D)    umf/0.10.0 (D)

  Where:
   L:  Module is loaded
   D:  Default Module
```

https://e4s.io

PyTorch built for Intel XPUs in E4S 25.06 Intel® oneAPI image

```
$ singularity run e4s-oneapi-x86_64-25.06.sif
Singularity> uname -m
x86_64
Singularity> python
Python 3.10.12 (main, Feb  4 2025, 14:57:36) [GCC 11.4.0] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>> import torch
>>> import intel_extension_for_pytorch as ipex
[W607 22:52:26.365045146 OperatorEntry.cpp:154] Warning: Warning only once for all operators,  other operators may also be overridden.
  Overriding a previously registered kernel for the same operator and the same dispatch key
  operator: aten::geometric_(Tensor(a!) self, float p, *, Generator? generator=None) -> Tensor(a!)
    registered at /pytorch/build/aten/src/ATen/RegisterSchema.cpp:6
  dispatch key: XPU
  previous kernel: registered at /pytorch/aten/src/ATen/VmapModeRegistrations.cpp:37
       new kernel: registered at /build/intel-pytorch-extension/build/Release/csrc/gpu/csrc/gpu/xpu/ATen/RegisterXPU_0.cpp:186 (function operator())
>>> torch.xpu.is_available()
True
>>> torch.__version__
'2.7.0+xpu'
>>>
```

E4S 25.06 image for ppc64le with NVIDIA GPUs (IBM Power 9)

```
$ singularity run --nv e4s-cuda70-ppc64le-25.06.sif
Singularity> uname -m
ppc64le
Singularity> spack find -x
-- linux-ubuntu20.04-ppc64le / gcc@9.4.0 ------------------------
adios@1.13.1          darshan-runtime@3.4.6     hdf5@1.14.5           mercury@2.4.0          pdt@3.25.2             sundials@7.2.1
adios2@2.10.2         darshan-util@3.4.6        hdf5-vol-async@1.7    metall@0.30            petsc@3.22.4           superlu@7.0.0
alquimia@1.1.0        datatransferkit@3.1.1     hdf5-vol-cache@v1.1   mfem@4.7.0             petsc@3.22.4           superlu-dist@9.1.0
aml@0.2.1             dyninst@13.0.0            hdf5-vol-log@1.4.0    mfem@4.7.0             phist@1.12.1           superlu-dist@9.1.0
amrex@25.03           e4s-alc@1.0.2             heffte@2.4.1          mgard@2023-12-09       plasma@24.8.7          swig@4.0.2-fortran
amrex@25.03           e4s-cl@1.0.4              heffte@2.4.1          mgard@2023-12-09       plumed@2.9.2           sz@2.1.12.5
arborx@1.5            exago@1.6.0               hpctoolkit@2024.01.1  mpark-variant@1.4.0    precice@3.1.2          sz3@3.2.0
arborx@1.7            exago@1.6.0               hpctoolkit@2024.01.1  mpifileutils@0.12      pruners-ninja@1.0.1    tasmanian@8.1
argobots@1.2          exaworks@0.1.0            hpx@1.10.0            nccmp@1.9.1.0          pumi@2.2.9             tasmanian@8.1
ascent@0.9.3          faodel@1.2108.1           hpx@1.10.0            nco@5.3.2              py-cinemasci@1.7.0     tau@2.34.1
axom@0.10.1           fftx@1.2.0                hypre@2.32.0          nek5000@19.0           py-h5py@3.12.1         tau@2.34.1
axom@0.10.1           fftx@1.2.0                hypre@2.32.0          nekbone@17.0           py-jupyterhub@1.4.1    trilinos@16.1.0
bolt@2.0              flit@2.1.0                kokkos@4.6.01         netcdf-fortran@4.6.1   py-libensemble@1.4.3   turbine@1.3.0
boost@1.87.0          flux-core@0.67.0          kokkos@4.6.01         netlib-scalapack@2.2.2 py-petsc4py@3.22.4     umap@2.1.1
bricks@2023.08.25     flux-core@0.67.0          kokkos-kernels@4.6.01 nrm@0.1.0              qthreads@1.18          umpire@2024.07.0
bricks@2023.08.25     fortrilinos@2.3.0         kokkos-kernels@4.6.01 nwchem@7.2.3           quantum-espresso@7.4.1 umpire@2024.07.0
butterflypack@3.2.0   fpm@0.10.0                laghos@3.1            omega-h@10.8.6-scorec  raja@2024.07.0         unifyfs@2.0
cabana@0.7.0          gasnet@2025.2.0-snapshot  lammps@20240829.1     omega-h@10.8.6-scorec  raja@2024.07.0         upcxx@2023.9.0
cabana@0.7.0          ginkgo@1.9.0              lammps@20240829.1     openfoam@2412          rempi@1.1.0            upcxx@2023.9.0
caliper@2.12.1        ginkgo@1.9.0              lbann@0.104           openmpi@5.0.6          scr@3.1.0              variorum@0.8.0
caliper@2.12.1        globalarrays@5.8.2        legion@24.12.0        openpmd-api@0.16.1     slate@2024.10.29       veloc@1.7
chai@2024.07.0        glvis@4.2                 libcatalyst@2.0.0     papi@7.1.0             slate@2024.10.29       visit@3.4.1
chai@2024.07.0        gmp@6.3.0                 libnrm@0.1.0          papi@7.1.0             slepc@3.22.2           vtk-m@2.2.0
chapel@2.4.0          gotcha@1.0.7              libpressio@0.99.4     papyrus@1.0.2          slepc@3.22.2           wannier90@3.1.0
charliecloud@0.38     gptune@4.0.0              libquo@1.4            parallel-netcdf@1.14.0 stc@0.9.0              wps@4.5
conduit@0.9.3         gromacs@2024.4            libunwind@1.7.2       paraview@5.13.2        strumpack@8.0.0        wrf@4.6.1
cp2k@2025.1           gromacs@2024.4            loki@0.1.7            parsec@3.0.2209        strumpack@8.0.0        xyce@7.8.0
cp2k@2025.1           h5bench@1.4               magma@2.9.0           parsec@3.0.2209        sundials@7.2.1         zfp@1.0.0
==> 168 installed packages
Singularity>
```

```
$ singularity run --nv e4s-cuda70-ppc64le-25.06.sif
Singularity> spack find +cuda
-- linux-ubuntu20.04-ppc64le / gcc@9.4.0 ------------------------
amrex@25.03          camp@2024.07.0     gromacs@2024.4        kokkos@4.6.01          nvcomp@2.2.0           raja@2024.07.0      umpire@6.0.0
arborx@1.5           camp@2024.07.0     heffte@2.4.1          kokkos@4.6.01          omega-h@10.8.6-scorec  slate@2024.10.29    umpire@2024.07.0
axom@0.10.1          chai@2024.07.0     hiop@1.0.0            kokkos-kernels@4.6.01  papi@7.1.0             slepc@3.22.2        umpire@2024.07.0
blaspp@2024.10.26    cp2k@2025.1        hpctoolkit@2024.01.1  lammps@20240829.1      parsec@3.0.2209        strumpack@8.0.0     umpire@2024.07.0
bricks@2023.08.25    exago@1.6.0        hpx@1.10.0            lapackpp@2024.10.26    petsc@3.22.4           sundials@7.2.1      upcxx@2023.9.0
cabana@0.7.0         fftx@1.2.0         hwloc@2.11.1          magma@2.9.0            petsc@3.22.4           superlu-dist@9.1.0
caliper@2.12.1       flux-core@0.67.0   hypre@2.32.0          mfem@4.7.0             raja@0.14.0            tasmanian@8.1
camp@0.2.3           ginkgo@1.9.0       kokkos@4.6.01         mgard@2023-12-09       raja@2024.07.0         tau@2.34.1
==> 53 installed packages
Singularity> which codium
/usr/local/codium/bin/codium
Singularity>
```

https://e4s.io

E4S: An HPC-AI Software Ecosystem for Science!

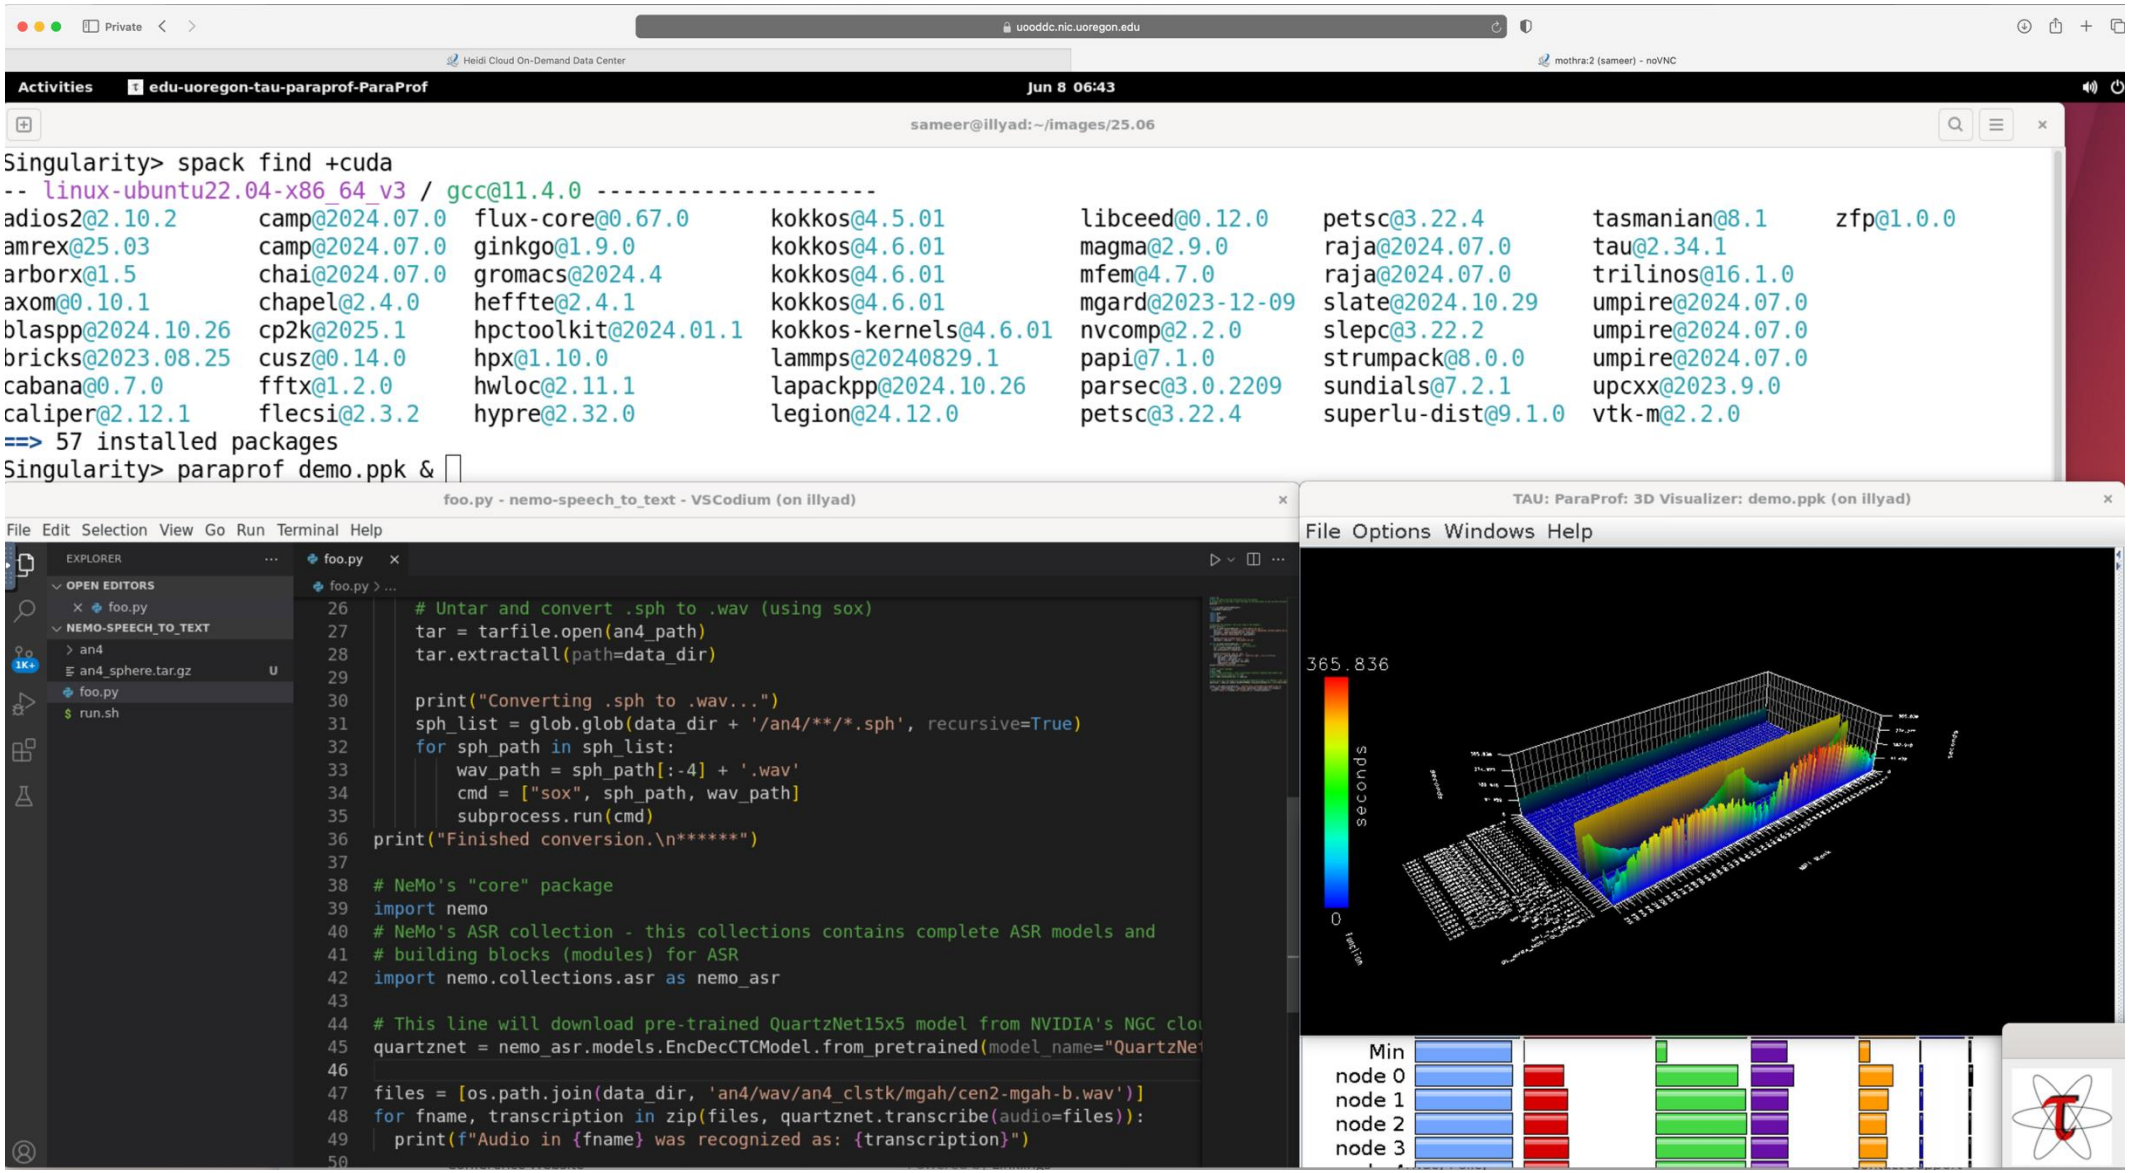

https://e4s.io

Acknowledgment

- *This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Next-Generation Scientific Software Technologies program, under contract numbers DE-AC02-AC05-00OR22725 and DOE SBIR DE-SC0022502.*

- https://science.osti.gov/ascr
- https://pesoproject.org
- https://ascr-step.org
- https://hpsf.io
- https://www.energy.gov/technologytransitions/sbirsttr

PESO

Thank you

https://www.exascaleproject.org

This research was supported by the Exascale Computing Project (17-SC-20-SC), a joint project of the U.S. Department of Energy's Office of Science and National Nuclear Security Administration, responsible for delivering a capable exascale ecosystem, including software, applications, and hardware technology, to support the nation's exascale computing imperative.

Thank you to all collaborators in the ECP and broader computational science communities. The work discussed in this presentation represents creative contributions of many people who are passionately working toward next-generation computational science.

PESO