

BOF Days

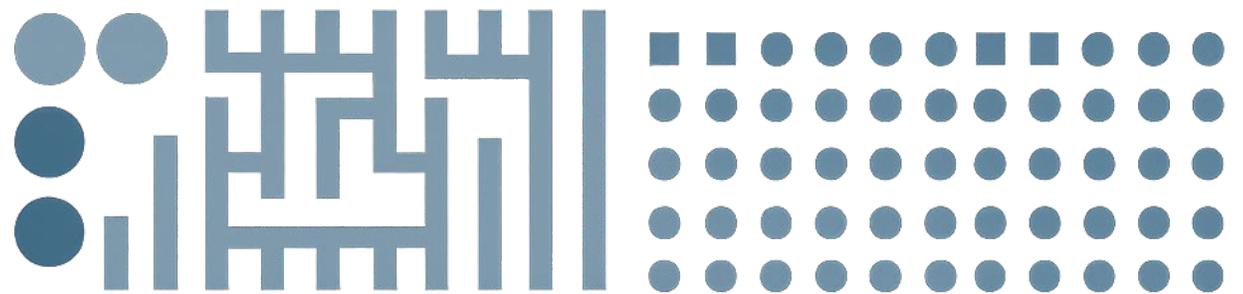
February 10 - 12, 2026

<https://cass.community>

The  
Consortium  
for the  
Advancement of  
Scientific  
Software  
(CASS)



**CASS**



**BOF DAYS**

<https://cass.community/news/2026-02-10-cass-bof-days.html>

# CASS: Stewardship and Advancement of the Scientific Software Ecosystem



- **Inward-facing activities:** Strengthening software products
  - Improve development practices, sustainability, quality, and trustworthiness
  - Enhance user experience and integration within the broader ecosystem
- **Outward-facing activities:** Community engagement and discovery
  - Curate and evolve the software portfolio
  - Help teams connect with and grow their user communities
  - Enable the broader community to discover and adopt useful software

## CASS Members



**CORSA**  
Partnering with foundations to provide sustainable pathways for scientific software

**FASTMATH**  
Stewardship, advancement, and integration for math and ML/AI packages

**PESO**  
Stewarding, evolving and integrating a cohesive ecosystem for DOE software

**RAPIDS**  
Stewardship, advancement, and integration for data, visualization and ML/AI packages

**S4PST**  
Stewardship, advancement and engagement for programming systems

**STEP**  
Stewardship, advancement of software tools for understanding performance and behavior

Sponsored by the Department of Energy, Office of Advanced Scientific Computing Research

# Engage with CASS!

- Learn about CASS:
  - <https://cass.community/about/>
- Join the CASS Announcement list (low-volume):
  - <http://eepurl.com/iRiSnY>
- Find out more about our **software products**
  - Catalog: <https://cass.community/software/>
  - Collected as part of the [Extreme-Scale Scientific Software Stack](#) (E4S)
- Participate in **CASS Working Groups**
  - Impact Framework, Integration, Metrics, Software Ecosystem, User-Developer Experience, Workforce
  - <https://cass.community/working-groups/>

# CASS BOF



## Quickstarting Your Way to Robust, Portable HPC-AI Computing

### Objectives:

- Establish outcome: Develop HPC-AI training resources for interns
- Show E4S QuickStart Approach
- Get Feedback
- Identify some next steps

# How This BOF Will Work

- Brief intro to E4S and its user resources
- Walk through QuickStart tutorial
- Request for community input on user support next steps
- Realtime support for QuickStart users

# E4S: Ecosystem for Science

Release 25.11



E4S

The logo features the letters "E4S" in a large, white, sans-serif font. The background is a dark blue gradient with a subtle hexagonal pattern.

<https://e4s.io>

# E4S: Ecosystem for Science

- Open-source HPC ecosystem for scientific software
- Curated, coherent stack for scalable exascale apps
- 125+ HPC, EDA, and AI/ML packages (Spack-based)
- GPU-enabled containers and DOE LLVM images
- Commercial support via ParaTools
- Cloud-ready (AWS, Azure, GCP, OCI, IBM Cloud)
- Bare-metal + container deployment with binary caches
- Multi-arch: x86\_64, ppc64le, aarch64; NVIDIA/AMD/Intel GPUs
- Prebuilt container images (DockerHub, E4S site)
- Tools: spack chaining, MPI swap, container customization
- Latest release: E4S 25.11 (Nov 14, 2025)

# E4S 25.11: What's New?



- New E4S.io website + OpenAI-powered chatbot
- NVIDIA Blackwell & Grace-Blackwell support
- Rocky 9.6, Ubuntu 24.04 LTS support
- Spack 1.0.2 integration
- New build cache: 7500+ optimized binaries
- Python 3.12 AI stack (NeMo, BioNeMo, PyTorch, TensorFlow, JAX, Gemini/OpenAI APIs)
- GPU-enabled HPC apps (CP2K, GROMACS, LAMMPS, OpenFOAM, WRF, etc.)
- Updated toolchains: CUDA 12.9, ROCm 6.4.3, oneAPI 2025.2
- Cloud platform support (Heidi + ParaTools Pro; AWS/Azure/GCP/IBM/OCI; SLURM/Torque)

<https://adaptivecomputing.com/> and <https://paratoolspro.com>

# E4S Spack Integration



- E4S is a curated, Spack based distribution of HPC-AI software.
- Major Changes in Spack 1.0.2
  - Compilers as First-Class Dependencies: Compilers are now treated as proper dependencies in the concretization process, leading to clearer and more reproducible environments.
  - Stable Package API: Spack 1.0 introduces a stable API for package development, improving long-term maintainability and easing contributions.
  - Concurrent Builds: Builds can now run concurrently, leveraging parallel jobs and increasing throughput on multi-core machines.
  - Updated Install Tree Layout: The default install tree format is revamped for better organization and reproducibility.
  - Content-Addressed Build Caches: Binary caches now use content-based addressing, improving the reliability and provenance of shared binaries.
  - Improved Git Provenance: Enhanced mirroring and fetching mechanisms for package sources and dependencies.

# Spack and E4S

## Why Spack is Necessary

- Flexible, source-based HPC package manager
- Fine-grained control (variants, compilers, MPI, GPUs)
- Rapid enablement of new architectures
- Reproducible builds via specs + concretization
- Foundation for custom site configurations

## Why E4S is Necessary

- Curated, validated scientific software stack
- Integrated HPC + AI ecosystem
- Production-ready binaries + containers
- Coherent version alignment across packages
- Reduces integration and validation burden

# Spack and E4S

## How They Complement Each Other

- Spack = *build system*; E4S = *curated distribution*
- E4S leverages Spack for build, concretization, provenance
- Spack enables site-specific extension of E4S
- E4S provides tested defaults; Spack enables customization
- Binary caches + containers accelerate deployment
- Together: flexibility + stability at scale

## Value to the Ecosystem

- Faster time-to-science
- Reduced duplication across labs
- Portable across CPUs, GPUs, and clouds
- Sustainable pathway from R&D to production
- Community-driven evolution of the scientific stack

# Updated E4S website <https://e4s.io>



The screenshot shows the homepage of the E4S website. At the top, there is a navigation bar with the E4S logo and the tagline "E4S - An HPC-AI Software Ecosystem for Science" and "Advancing Science through Trusted Software". To the right of the logo are links for "Use E4S", "Join E4S", "About E4S", and "Contact Us". The main content area has a blue background with the heading "E4S: HPC-AI Software Ecosystem for Science" and the subtitle "Performance-portable libraries & tools for advanced computing". Below this are two buttons: "News & Events" and "Get E4S". The text below the buttons describes E4S as an open-source, community-driven collection of high-quality HPC and AI libraries and tools, supported by the U.S. Department of Energy (DOE). At the bottom, there are four feature cards: "Performance Portability", "Curated Ecosystem", "Reproducible Environments", and "Continuous Validation", each with a corresponding icon and a brief description.

**E4S** - An HPC-AI Software Ecosystem for Science  
Advancing Science through Trusted Software

Use E4S   Join E4S   About E4S   Contact Us

## E4S: HPC-AI Software Ecosystem for Science

Performance-portable libraries & tools for advanced computing

[News & Events](#)   [Get E4S](#)

Welcome to **E4S**, the *HPC-AI Software Ecosystem for Science* — an open-source, community-driven collection of high-quality HPC and AI libraries and tools. E4S is an **Extreme-scale Scientific Software Stack** that enables scientists, developers, and institutions to develop, deploy, and run performance-portable applications across CPUs and GPUs from NVIDIA, AMD, Intel, and Arm. Ready to adapt to new architectures as they emerge.

Supported by the **U.S. Department of Energy (DOE)** and its partners, E4S accelerates scientific innovation on systems ranging from laptops to exascale supercomputers.

- Performance Portability**  
Run anywhere—from laptops to exascale supercomputers using modern GPU and CPU architectures.
- Curated Ecosystem**  
Over 120 interoperable libraries and tools for HPC and AI workflows, validated for research and production use.
- Reproducible Environments**  
Deploy via Spack, pre-built containers for Docker, Singularity, and commercial cloud environments.
- Continuous Validation**  
E4S packages are tested across DOE leadership and cloud platforms for quality and reproducibility.

# Chatbot integration in E4S Website

A screenshot of a web browser displaying the E4S website. The browser's address bar shows 'https://e4s.io'. The website layout includes four feature cards at the top: 'Performance Portability', 'Curated Ecosystem', 'Reproducible Environments', and 'Continuous Validation'. Below these are two larger cards: 'Chat with the E4S Bot' and 'Visit the E4S Product Catalog', each with a blue button. At the bottom, there is a section titled 'Trusted by the HPC-AI Community' with descriptive text and a blue footer containing copyright information.

Performance Portability  
Run anywhere—from laptops to exascale supercomputers using modern GPU and CPU architectures.

Curated Ecosystem  
Over 120 interoperable libraries and tools for HPC and AI workflows, validated for research and production use.

Reproducible Environments  
Deploy via Spack, pre-built containers for Docker, Singularity, and commercial cloud environments.

Continuous Validation  
E4S packages are tested across DOE leadership and cloud platforms for quality and reproducibility.

Chat with the E4S Bot  
Ask questions, explore tools, or get recommendations from the E4S Guide Bot.  
[Chat with the E4S Guide Bot](#)

Visit the E4S Product Catalog  
Browse member products, product families, with extensive links, via the E4S Product Catalog  
[Browse the E4S Product Catalog](#)

**Trusted by the HPC-AI Community**

E4S is developed and supported by contributors from U.S. national laboratories, universities, and industry partners. It is available as both open-source and commercial distributions such as [ParaTools Pro for E4S™](#) for commercial cloud platforms such as AWS, Azure, GCP, OCI, and IBM Cloud.

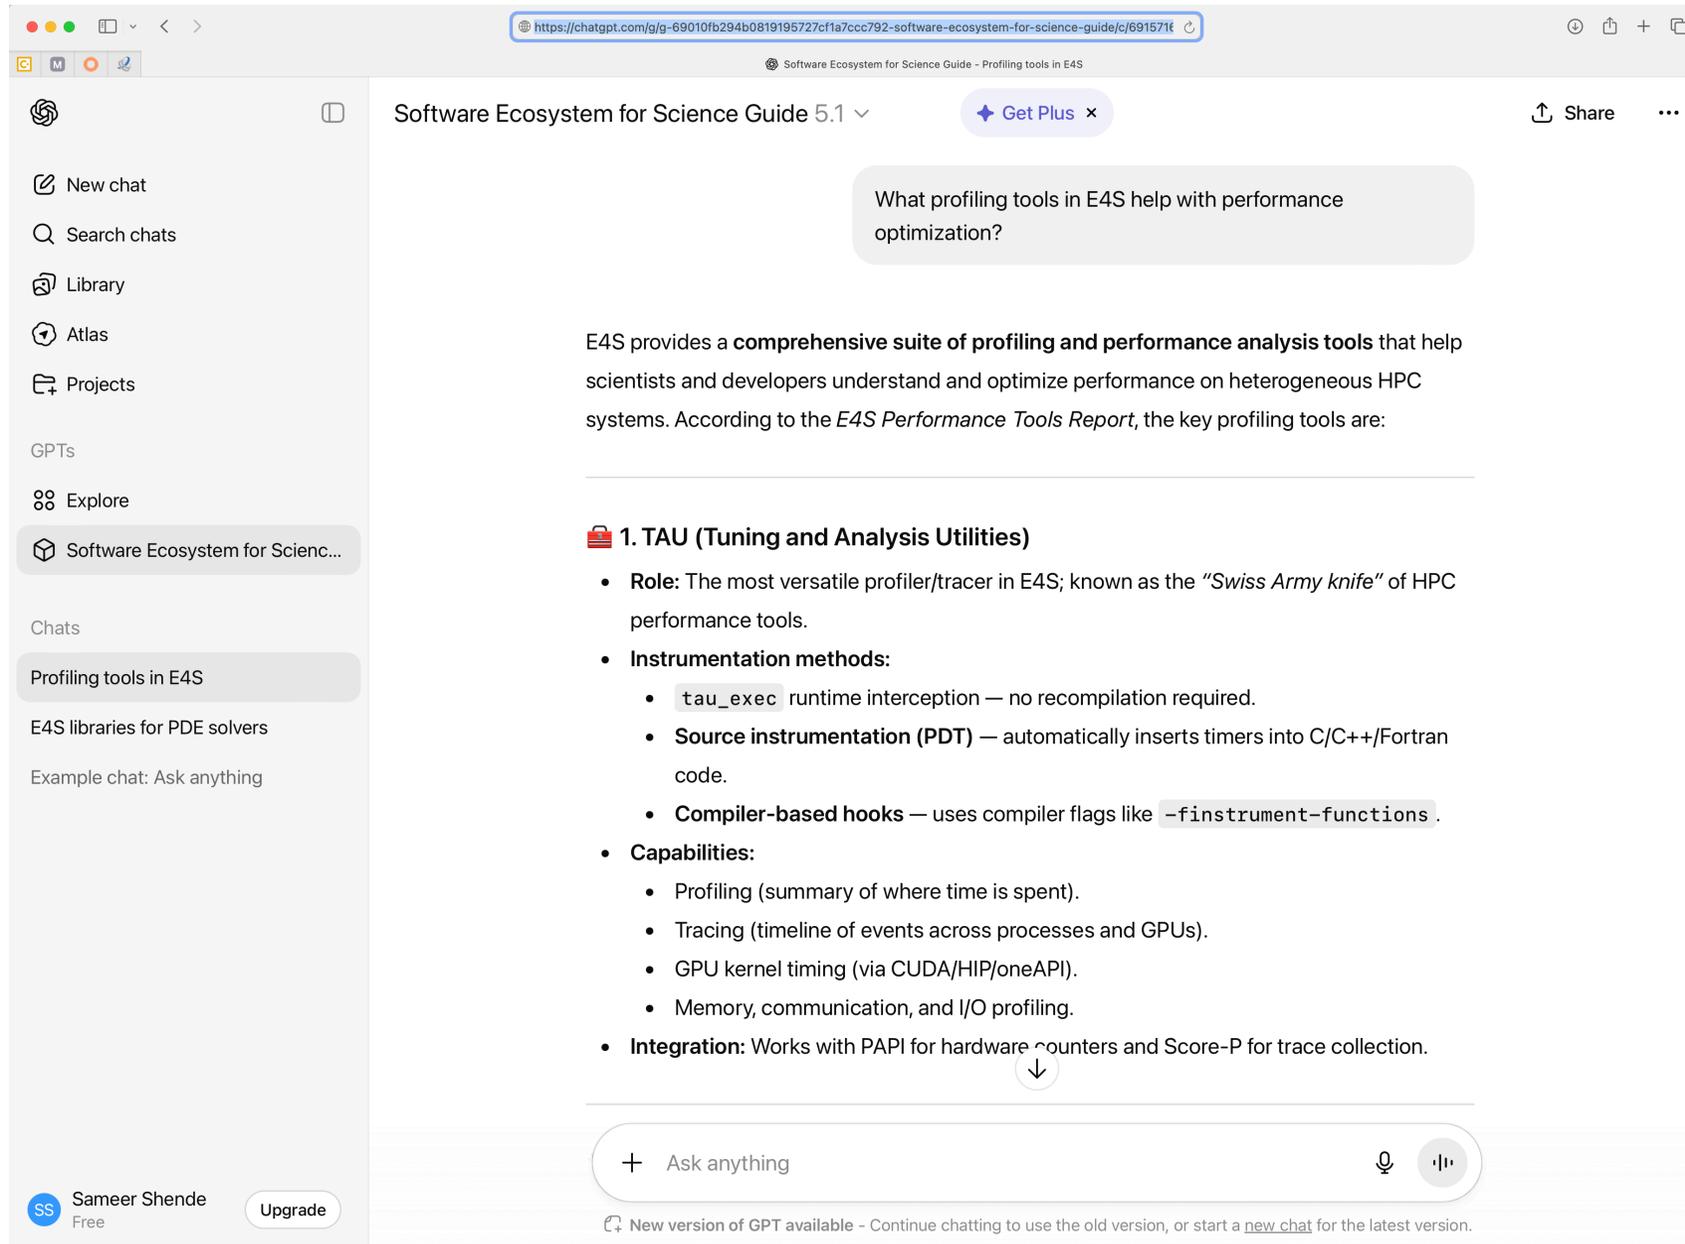
© 2025 E4S Project, a Series of LF Projects, LLC. For website terms of use, trademark policy and other project policies please see [lfprojects.org](#).  
Supported by the U.S. Department of Energy Office of Advanced Scientific Computing Research.

# Chatbot integration in E4S Website



A screenshot of a web browser displaying a ChatGPT interface. The browser's address bar shows the URL: https://chatgpt.com/g/g-69010fb294b0819195727cf1a7ccc792-software-ecosystem-for-science-guide. The page title is 'Software Ecosystem for Science Guide 5.1'. The main content area features a central heading 'Software Ecosystem for Science Guide' by Michael Heroux, described as an expert on scientific libraries and tools in the E4S ecosystem. Below this, there are four interactive buttons with questions: 'Which E4S libraries should I use for parallel PDE...', 'How do I install and configure PETSc with Spack?', 'Can you explain how Kokkos interacts with Trilinos?', and 'What profiling tools in E4S help with performance...'. The left sidebar contains navigation options like 'New chat', 'Search chats', 'Library', 'Atlas', 'Projects', 'GPTs', and 'Explore'. At the bottom, there is a chat input field with a microphone icon and a status bar for 'Sameer Shende Free' with an 'Upgrade' button.

# E4S Bot trained on E4S content augmenting OpenAI's ChatGPT



https://chatgpt.com/g/g-69010fb294b0819195727cf1a7ccc792-software-ecosystem-for-science-guide/c/691571f

Software Ecosystem for Science Guide 5.1 [Get Plus](#) x Share

What profiling tools in E4S help with performance optimization?

E4S provides a **comprehensive suite of profiling and performance analysis tools** that help scientists and developers understand and optimize performance on heterogeneous HPC systems. According to the *E4S Performance Tools Report*, the key profiling tools are:

### 1. TAU (Tuning and Analysis Utilities)

- **Role:** The most versatile profiler/tracer in E4S; known as the “*Swiss Army knife*” of HPC performance tools.
- **Instrumentation methods:**
  - `tau_exec` runtime interception — no recompilation required.
  - **Source instrumentation (PDT)** — automatically inserts timers into C/C++/Fortran code.
  - **Compiler-based hooks** — uses compiler flags like `-finstrument-functions`.
- **Capabilities:**
  - Profiling (summary of where time is spent).
  - Tracing (timeline of events across processes and GPUs).
  - GPU kernel timing (via CUDA/HIP/oneAPI).
  - Memory, communication, and I/O profiling.
- **Integration:** Works with PAPI for hardware counters and Score-P for trace collection.

Ask anything

Sameer Shende Free Upgrade

New version of GPT available - Continue chatting to use the old version, or start a [new chat](#) for the latest version.

# E4S Bot



chatgpt.com

Launch GPT-OSS 120B E4S

Get Plus x

Share ...

How can I launch OpenAI's 120B parameter OSS model using E4S?

Thought for 16s >

### Quick steps — launch OpenAI's gpt-oss-120b on E4S with vLLM

Short answer: since recent E4S releases include **vLLM**, you can run the gpt-oss 120B model directly with vLLM's `serve` command on a machine in your E4S environment (recommended: single 80 GB GPU like an H100 / A100-80GB or a multi-GPU setup). The cookbook from OpenAI gives a one-line server command and examples for calling it with the OpenAI-compatible API. OACISS +1

---

#### 1) Preconditions (hardware & E4S)

- GPU:** gpt-oss-120b is intended for  $\geq 60-80$  GB VRAM (fits on a single H100 / A100-80GB or multi-GPU). If you only have smaller GPUs you'll need model parallel / multi-GPU setup or run the 20B model instead. OpenAI Cookbook +1
- E4S image:** E4S 25.06+ lists vLLM in the stack — meaning vLLM may already be in the image or available as a module/container. If vllm is already on `PATH` inside your E4S environment, you can skip pip installs. OACISS

---

#### 2) Minimal, working flow (commands you can run right now in an E4S compute node)

If `vllm` is in the E4S image and you have an appropriate GPU:

```
bash
```

```
# make sure you are on the node with the target GPU (and modules/drivers loaded by you)
# then just:
vllm serve openai/gpt-oss-120b
```

That will (by default) download the model from HuggingFace and spin up an OpenAI-compatible server at `http://localhost:8000/v1`. Example client call (Python OpenAI SDK pointed at your local server):

+ Ask anything

Sameer Shende Free Upgrade

ChatGPT can make mistakes. Check important info.

# E4S Documentation



Private < > <https://e4s.io/about/> 1

**E4S** E4S - An HPC-AI Software Ecosystem for Science  
Advancing Science through Trusted Software

Use E4S Join E4S About E4S Contact Us

**ABOUT E4S**

- Overview
- E4S Introduction
- First-Time Users
- E4S Product Families
- E4S FAQs
- News & Events

## About E4S

E4S provides a curated collection of scientific libraries and tools supporting scientific discovery through advanced computing technologies. The primary value E4S provides is bringing together a large collection of independently developed libraries and tools and enhancing the overall enterprise of developing and using them in aggregate. This portion of the E4S website focuses on topics that help understand what E4S is, how to contribute to it, and how it can be used.

Topic	Description
Overview - About E4S	High-level summary of the E4S project, mission, scope, and target audiences.
E4S Introduction	Short primer on E4S goals, components, and how it supports HPC-AI scientific workflows.
E4S First-Time Users	Quickstart guidance for new users: getting started, basic concepts, and common workflows.
E4S Product Families	Catalog and explanation of E4S product groups (programming systems, libraries, tools, etc.).
E4S FAQs	Frequently asked questions covering usage, licensing, support, and common troubleshooting.
E4S News & Events	Announcements, release notes, and upcoming events related to E4S.

© 2025 E4S Project, a Series of LF Projects, LLC. For website terms of use, trademark policy and other project policies please see [lfprojects.org](https://lfprojects.org).  
Supported by the U.S. Department of Energy Office of Advanced Scientific Computing Research.

# E4S Product Catalog



Private e4s.io

**E4S** E4S - An HPC-AI Software Ecosystem for Science  
Advancing Science through Trusted Software

Use E4S Join E4S About E4S Contact Us

**USE E4S**

- Overview
- E4S Guide Bot
- E4S Documentation
- E4S Product Catalog**
- Finding and Installing with Spack
- E4S-Specific Spack Installation
- E4S Container Download
- E4S Container Installation
- E4S Container Launch
- Test with E4S
- ParaTools Pro for E4S™

### E4S Product Catalog

The E4S Product Catalog provides a convenient searchable and sortable table that provides per-product information for all primary E4S products.

Search  Area

Name ↑	Area ↓	Description ↓	Last Updated ↓	Details
DEEPHYPER AI	AI	DeepHyper is a powerful Python package for automating machine learning tasks	2025-10-27	<a href="#">Open</a>
HOROVOD AI	AI	a distributed deep learning training framework	2023-06-12	<a href="#">Open</a>
JAX AI	AI	Autograd and XLA, brought together for high-performance numerical computing	2025-11-11	<a href="#">Open</a>
KERAS AI	AI	a multi-backend deep learning framework	2025-10-30	<a href="#">Open</a>
LBANN AI	AI	an open-source, HPC-centric, deep learning training framework	2025-05-09	<a href="#">Open</a>
NEMO AI	AI	NVIDIA NeMo Framework is a scalable and cloud-native generative AI framework	2025-11-10	<a href="#">Open</a>
OPENAI-PYTHON AI	AI	The OpenAI Python library provides convenient access to the OpenAI REST API	2025-11-10	<a href="#">Open</a>
OPENCV AI	AI	Open Source Computer Vision Library	2025-08-01	<a href="#">Open</a>
PANDAS AI	AI	powerful Python data analysis toolkit	2025-11-01	<a href="#">Open</a>
PYTORCH AI	AI	Provides accelerated tensor computation and deep neural networks	2025-11-11	<a href="#">Open</a>
SCIKIT-LEARN AI	AI	a Python module for machine learning	2025-11-05	<a href="#">Open</a>
TENSORFLOW AI	AI	An end-to-end open source platform for machine learning	2025-10-28	<a href="#">Open</a>

# E4S: Bare-metal installation, Containers, and Cloud images



Private < > e4s.io

**E4S** E4S - An HPC-AI Software Ecosystem for Science  
Advancing Science through Trusted Software

Use E4S Join E4S About E4S

## GET E4S

### Get E4S

Overview  
E4S Spack Caches  
E4S Container Download  
ParaTools Pro for E4S™

E4S offers multiple methods for obtaining its collection of HPC and AI software packages. Below is an overview of these deployment options. Additional documentation is available from the side menu.

Get Option	Description
Spack Build Cache	Pre-built binaries provided via Spack build caches enable faster installations without local compilation. More information available on the <a href="#">E4S Buildcache Page</a> .
Containers	E4S offers containerized versions of its software stack, compatible with Docker, Singularity, Shifter, and CharlieCloud. See the <a href="#">Container Download Page</a> for available containers. See side menu for additional container information.
Commercial Cloud Options	E4S containers are available via <a href="#">ParaTools Pro for E4S™</a> , including <a href="#">AWS Marketplace</a> . E4S can be deployed on Amazon Web Services (AWS) for scalable cloud-based HPC/AI workloads and Google Cloud Platform (GCP) with access to high-performance VMs and storage.

© 2025 E4S Project, a Series of LF Projects, LLC. For website terms of use, trademark policy and other project policies please see [lfprojects.org](https://lfprojects.org). Supported by the U.S. Department of Energy Office of Advanced Scientific Computing Research.

Let's Visit <https://e4s.io>

- Website Tour

- About E4S

- Intro, First-Time Users, Product Families
    - E4S Learning: New resource – under construction – Your input requested

- E4S Documentation

- Product Catalog

- E4S Guide Bot

- Product Family Prompt Guides

- E4S Website Feedback:

[Ask a question \(GitHub Discussions\)](#)

[Report an issue for this page](#)

[Propose page edits](#)

- E4S QuickStart – Let's build and run an MPI code on your laptop!

# We Need Your Input

- We are preparing E4S Learning:
  - A resources for new E4S users to get started with portable HPC-AI computing
  - E4S QuickStart is just the start
- Goal: Prepare content, on-demand (video) resources for self-paced learning, webinars
- Timeline: Prepare for 2026 summer intern season
- Motivation:
  - Interns at labs must often acclimate to HPC-AI software ecosystems during on-ramping
  - We want to be ready with E4S Learning as a resource for on-ramping

# Questions

- Who among your community can benefit from something like E4S Learning?
- Is the approach we highlighted in E4S QuickStart a good starting point?
- If so, what should we do next?
- Ideas:
  - MLOps workflow with a simple example
  - Other workflows: solvers, performance optimization
    - Caution: Don't want to duplicate content already available from product teams, e.g., we don't need a PETSc-only example
- Who can we engage at your institution to prepare E4S Learning?

## Last Activity: Help with E4S QuickStart issues

- We will stick around for anyone who wants help with the E4S QuickStart process.
- Others: Thank you for your participation!