⬅️ Connect with CASS
https://tinyurl.com/2024-CASS-BOFS

CASS Community BOI

June 11 – 13, 2024

https://cass.community/bofs

Announcing CASS

The Consortium for the Advancement of Scientific Software

CASS Basics

- A newly-formed organization
- Sponsored by DOE Office of Advanced Scientific Computing Research (ASCR)
- Established by DOE Software Stewardship Organizations (SSOs)

CASS Goals

- Forum for SSO collaboration and coordination
- Bigger than the sum of its parts
- Vehicle for advancing the scientific software ecosystem

CASS Status

- Defining governance structure
- Establishing community awareness
- Building a team of teams
- Collaborating on outreach

Software Stewardship Organization (SSO) Basics

- Each SSO represents a specific software ecosystem concern
- **Product SSOs:** Programming systems, performance tools, math packages, data/viz packages
- **Portfolio SSO:** Curating & delivering software stack to the community
- **Community SSOs:** Workforce, partnerships

Engage with CASS

- Participate in June 11-13 CASS Community BOF Days: https://cass.community/bofs
- Visit https://cass.community

8 Software Stewardship Organizations (SSOs)

DOE Office of Advanced Scientific Computing Research (ASCR) Post-ECP Projects

COLABS

Training, workforce development, and building the RSE community

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, viz, and ML/AI packages

S4PST

Stewardship, advancement and engagement for programming systems

STEP

Stewardship, advancement of software tools for understanding performance and behavior

SWAS

Stewardship and project support for scientific workflow software and its community

Building an Inclusive and Productive Community from Many Organizations to Support Software Stewardship

Suzanne Parete-Koon, ORNL, Lead of HPC-Workforce Community Group
Mary Ann Leung, Sustainable Horizons Institute
Anshu Dubey, ANL
William Godoy, ORNL
Terry Jones, ORNL
Dan Martin, LBNL
Lois Curfman McInnes, ANL
Lavanya Ramakrishnan, LBNL
Elaine Raybourn, SNL

CASS Community BOF Days June 12, 2024

ECP software ecosystem: Powering diverse apps across multiple architectures

Project/PI	EXAALT: Molecular Dynamics Danny Perez
Challenge Problem	Damaged surface of Tungsten in conditions relevant to plasma facing materials in fusion reactors • 100,000 atoms • T=1200K
FOM Speedup	398.5
Nodes Used	7000
ST/CD Tools	Used in KPP Demo: **Kokkos, CoPa**

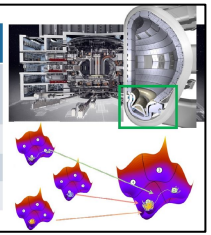

Project/PI	ExaSMR: Small Modular Reactors Steve Hamilton
Challenge Problem	NuScale-style Small Module Reactor (SMR) with depleted fuel and natural circulation • 213,860 Monte Carlo tally cells/6 reactions • 5.12×10^{12} particle histories/cycle, 40 cycles • 1098×10^{6} CFD spatial elements • 376×10^{9} CFD degrees of freedom • 1500 CFD timesteps
FOM Speedup	70
Nodes Used	6400
ST/CD Tools	Used in KPP Demo: **CEED** Additional: Trilinos

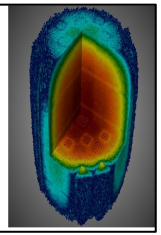

Project/PI	ExaSky: Cosmology Salman Habib
Challenge Problem	Two large cosmology simulations • gravity-only • hydrodynamics
FOM Speedup	271.65
Nodes Used	8192
ST/CD Tools	Used in KPP demo: **none** Additional: CoPa, VTK-m, CINEMA, HDF5.0

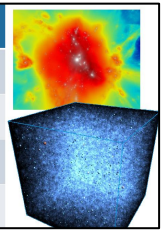

Project/PI	WarpX: Plasma Wakefield Accelerators Jean-Luc Vay
Challenge Problem	Wakefield plasma accelerator with a 1PW laser drive • 6.9×10^{12} grid cells • 14×10^{12} macroparticles • 1000 timesteps/1 stage
FOM Speedup	500
Nodes Used	8576
ST/CD Tools	Used in KPP Demo: **AMReX**, **libEnsemble** Additional: ADIOS, HDF5, VTK-m, ALPINE

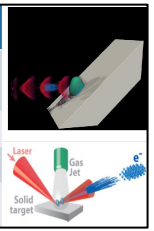

Project/PI	WDMApp: Fusion Tokamaks Amitava Bhatacharjee
Challenge Problem	Gyrokinetic simulation of the full ITER plasma to predict the height and width of the edge pedestal
FOM Speedup	150
Nodes Used	6156
ST/CD Tools	Used in KPP Demo: **CODAR**, **CoPA**, **PETSc**, **ADIOS** Additional: VTK-m

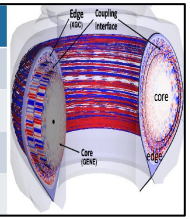

Project/PI	EQSIM: Earthquake Modeling and Risk Dave McCallen
Challenge Problem	Impacts of Mag 7 rupture on the Hayward Fault on the bay area.
FOM Speedup	3467
Nodes Used	5088
ST/CD Tools	Used in KPP Demo: **RAJA, HDF5**

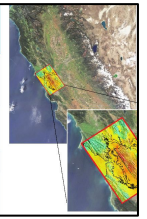

ECP Software Technologies

- Prepare SW stack for scalability with massive on-node parallelism
- Extend existing capabilities when possible, develop new when not
- Guide, and complement, and integrate with vendor efforts
- Develop and deliver high-quality and robust software products

70 software products across 6 technical areas

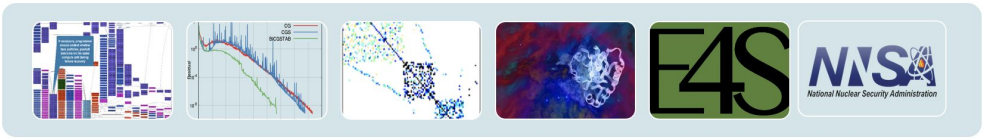

Programming Models & Runtimes	Development Tools	Math Libraries	Data and Visualization	Software Ecosystem	NNSA ST
•Enhance and get	•Continued, multifaceted	•Linear algebra, iterative linear solvers, direct linear solvers, integrators	• I/O via the HDF5 API	•Develop features in Spack necessary to	•Open source NNSA Software projects •Projects that have

ECP = DOE Exascale Computing Project, https://www.exascaleproject.org/

E4S: https://e4s.io

Software Stewardship for Computational Science

- Maintaining a robust software ecosystem for scientific computing includes stewarding
  - Tools for software development and performance
  - Libraries for math, visualization, I/O, data analytics, learning, and more …
  - Applications
  - Standards for all of those, the ecosystem overall
- Research codes run on many different types of computer hardware
  - Laptops and desktops
  - Clusters
  - Supercomputers
- Paradigms are evolving and emerging
  - Heterogeneous architectures (CPUs, GPUs, …)
  - Programing models for high-performance computing (HPC)
  - Machine learning and Artificial intelligence (ML/AI)
  - Quantum computing

We need to foster an inclusive community who collaborate across disciplines … so that we can advance software stewardship as needed for research and scientific discovery.

An Inclusive Community

An inclusive community is one where:

- A **diversity of ideas** is encouraged and represented.
- **People with different career backgrounds and stages** contribute and are respected.
- **All cultures, races, ethnicities, genders and levels of physical and intellectual ability** are welcome and respected.
- Collaboration and competition coexist to drive progress without driving animosity.
- There are structures that enable new people and ideas to find and join the community and veteran members to persist in the community.

Consortium for the Advancement of Scientific Software (CASS) Workforce Initiative

U.S. DEPARTMENT OF ENERGY | Office of Science

Partnership with
Sustainable Horizons Institute
https://shinstitute.org

SUSTAINABLE HORIZONS INSTITUTE

CASS

Fostering and stewarding a workforce community with broad pathways to sustainable HPC & AI careers

Intro to HPC & AI for Science

Accessible introductory material addressing gaps in — and expanding the pipeline of — people with foundational HPC and AI skills. This onramp begins a pathway to build experience and interest in HPC and AI for science.

Sustainable Research Pathways

A multi-lab workforce development program with students and faculty working side-by-side with DOE lab teams on world-class projects using HPC and AI for science.

HPC Workforce Community Group

Enabling staff of DOE national labs to share their collective insight for inclusive and equitable workforce development and retention for careers in HPC and AI for science.

CASS Workforce Initiative pursues a multi-layered strategy, with both crosscutting and project-specific activities

Leveraging and extending the ECP Broadening Participation Initiative:

- *Building a diverse and inclusive HPC community mission-driven team science,* IEEE CiSE, Nov 20 https://doi.org/10.6084/m9.figshare.24563371
- *Intro to HPC Bootcamp: Engaging new communities through energy justice projects*, J. Comp. Science Edu., 2024 https://doi.org/10.22369/issn.2153-4136/15/1/10
- *A multipronged approach to building a diverse workforce and cultivating an inclusive professional environment for DOE high-performance computing*, 2021, https://doi.org/10.6084/m9.figshare.17192492

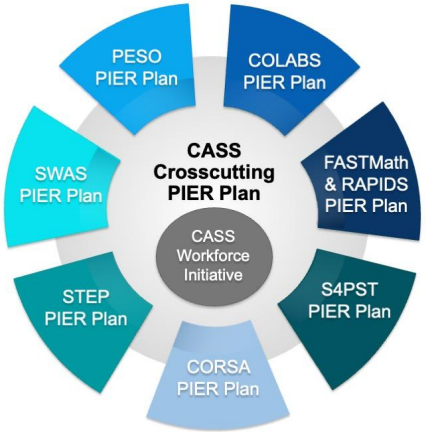

Why the CASS Workforce Initiative? Unique multilab partnership across DOE computing sciences

- Strength in spanning multiple institutions / strength in numbers / network beyond what any individual lab could do
- Proactive outreach and deployment of DOE scientific software tools and technologies to communities beyond traditional targets

Example: Formal Inclusive Community Scaffolding

CASS Crosscutting PIER Plan: Preliminary brainstorming … Now prioritizing initial focus:

- Commit to participate in the Sustainable Research Pathways (SRP) Program
- Define and use actionable metrics that quantify diversity in collaboration, community, and authorship of publications
- Outline an aggressive training and reward plan for people who proactively work to promote inclusive and equitable research
- Craft a vision and strategy for how senior personnel can mentor, promote, and empower the career growth of diverse early-career community members
- Support a strategy for the inclusion of people executing research software engineering (RSE) work in publications, and pursue community activities for RSE career growth
- Encourage all members of the consortium to proactively engage in activities that improve the diversity, equitability, inclusivity and accessibility of our community – to avoid “invisible work” in which only members of underrepresented groups are involved in such activities
- Engage all members to refine the consortium code of conduct
- Pursue the CASS Workforce Initiative (extending the ECP Broadening Participation Initiative)
    - HPC Workforce Community Group
    - Sustainable Research Pathways
    - Intro to HPC and AI for Science

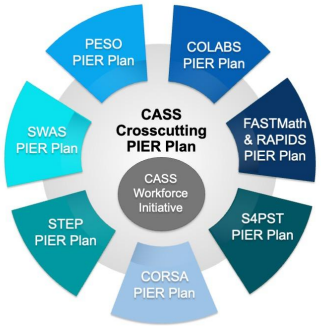

CASS community formally agree to support specific practices and projects.

Example: Informal Inclusive Community Scaffolding

HPC Workforce Community Group - formerly the HPC Workforce Development and Retention Action Group (HPC-WDR)

Mission: Enable the DOE National Laboratories and their related computing communities to share their collective insight for inclusive and equitable workforce development and retention for careers related to scientific computing.

We foster a community, within the DOE High-Performance Computing communities, that comes together on a regular basis to share ideas, catalog best practices, and develop recommendations and strategies for improvement.

- Scientific Computing Workforce Webinars
  Last Webinar: HPC Culture: https://hpc-workforce-development-and-retention.github.io/hpc-wdr/events/event-hpc-culture

- Website with blogs and shared events
  https://hpc-workforce-development-and-retention.github.io/hpc-wdr/

People volunteer to participate and share Ideas and best practices and work on projects

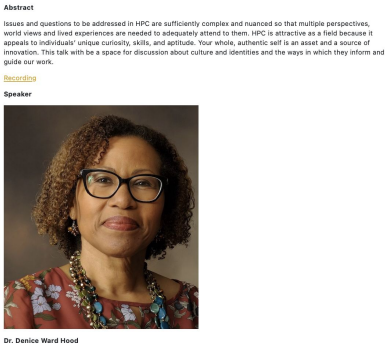

Our Purpose Today

We each represent our own different organization, experiences, and ideas related to building a research ecosystem.

- We will explore those backgrounds with ***Same and Different Exercises*.**
- We will discuss the results as a group and think about how we can align this wealth of experience for building a software stewardship community.
- We will discuss some inclusive community building structures proposed in the CASS Plan for Promoting Inclusive and Equitable Research (PIER) and build on those with our shared experience of ideas.

Same and Different: We ourselves

We will join breakout rooms, each with a leader (starting in a minute).

First, I'll explain how the game works:

- **Part 1:** Each person should take 5 minutes to list three things about themselves on the google sheet for their room.
- **Part 2:** During the next 5 minutes, your group should sort the list for what is the **same** and what is **different**.
    - Also list **rare assets** - **unique or less represented ideas/skills/interests**

Same and Different: We ourselves example

Example:

(I am a mother; I won my first figure skating competition medals last week; My first programming language was Fortran)
(I play soccer; I have two dogs; I like Mexican food)
(I paint miniatures, I take care of my parents; Python was my first programming language)
(I like Korean food; I like to cook; cats are my favorite animals)

Same	**Different**
Love of sports	Different sports with different kinds of skills
Love of animals	Different animals are represented
Care taking	Different kinds of dependents get care
Foodies	Different foods are represented
People program	Compiled and non-compiled languages, vintage and new languages are represented

Rare Assets: Art-related hobbies

Breakout Instructions

Join a breakout room (that has fewer people in it already). Don’t intentionally try to join rooms with your close collaborators or co-workers, but don’t worry if they are there by chance.

- Your room leader will have the google sheet link.
- When you enter the room, get the link from them.
- Remember your room number.

Any questions?

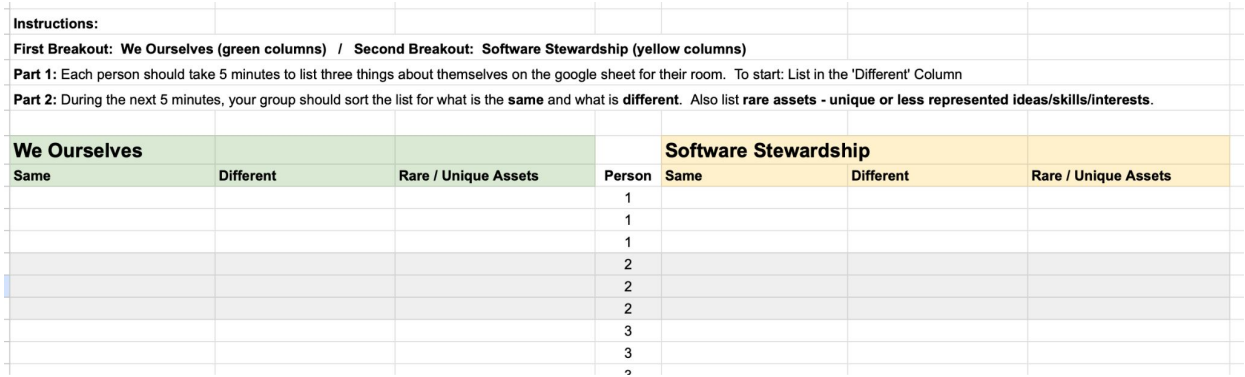

Instructions:

First Breakout: We Ourselves (green columns) / Second Breakout: Software Stewardship (yellow columns)

Part 1: Each person should take 5 minutes to list three things about themselves on the google sheet for their room. To start: List in the 'Different' Column

Part 2: During the next 5 minutes, your group should sort the list for what is the **same** and what is **different**. Also list **rare assets - unique or less represented ideas/skills/interests**.

We Ourselves				Software Stewardship		
Same	Different	Rare / Unique Assets	Person	Same	Different	Rare / Unique Assets
			1			
			1			
			1			
			2			
			2			
			2			
			3			
			3			

Breakout Instructions

Join a breakout room (that has fewer people in it already). Don’t intentionally try to join rooms with your close collaborators or co-workers, but don’t worry if they are there by chance.

- Your room leader will have the google sheet link.
- When you enter the room, get the link from them.
- Remember your room number.

Any questions?

Rooms are now open.

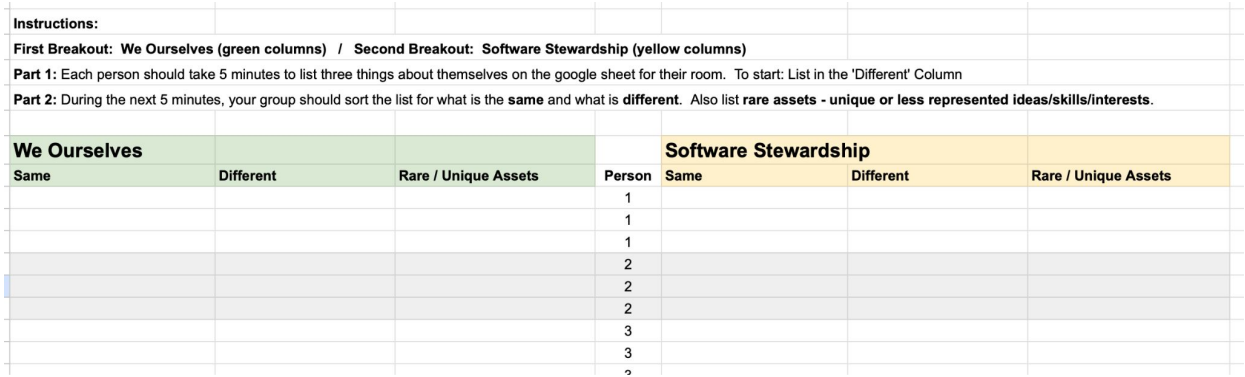

Instructions:

First Breakout: We Ourselves (green columns) / Second Breakout: Software Stewardship (yellow columns)

Part 1: Each person should take 5 minutes to list three things about themselves on the google sheet for their room. To start: List in the 'Different' Column

Part 2: During the next 5 minutes, your group should sort the list for what is the **same** and what is **different**. Also list **rare assets - unique or less represented ideas/skills/interests**.

We Ourselves				Software Stewardship		
Same	Different	Rare / Unique Assets	Person	Same	Different	Rare / Unique Assets
			1			
			1			
			1			
			2			
			2			
			2			
			3			
			3			

Same and Different: We ourselves: Sharing

What are some of the commonalities from each group?

What are some differences represented?

What are some of the rare/unique assets of each group?

Same and Different: Software Stewardship

Next Exercise

Each person should take 5 minutes to list three things in their work that they do for software stewardship in the google sheet for their room.

During the next 10 minutes, your group should sort the list for what is the same and what is different.

Also list rare/unique assets.

Same and Different: Software Stewardship Example

(I teach coding and tool usage for HPC, Technical writer, I help users)

(I build profiler tools, I design tests for those tools, I train end users)

(I develop Julia, I test Julia, I teach people how to use Julia)

Same	**Different**
Teaching of ends users	Tools and languages, Different tools/Langues
Developers	Users/Trainer
Testing	Different methods of testing for different things
User training	Generalist / For Specific tools

Unique Asset: technical writing

Discussion Software Stewardship

What are some of the commonalities from each group?

What are some differences represented ?

What are some of the Unique Assets?

Discussion

Why are the commonalities common?

Would some of those commonalities have been rare assets a decade ago?

How can we use our commonalities to collaborate?

Why are the *rare assets* rare?

How do we build a community that supports an influx of rare assets (new ideas)?

Thank you!

We welcome you to join us for further discussions on advancing the HPC workforce …

QR code: *Interest Sheet for HPC Workforce Community Group*