ACCESS and Major Facilities

Benedikt Riedel
UW-Madison

CI4MF
16 January 2024

Content Warning

These are my opinions, based my experience and observations with ACCESS and its resource providers

Context

- ACCESS - Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support
  - Allocation, common user support, monitoring, etc. for **NSF-funded HPC resources**
    - PSC Bridges-2, SDSC's Expanse, NCSA's Delta, TACC's Stampede, Purdue's Anvil
    - Except for leadership class facility – Frontera, Vista, Horizon
  - No control of CI resources beyond allocating time – By design
  - Replaced XSEDE in 2022
  - INCITE is the DOE equivalent
- Major Facilities
  - Large scale (>$100M) NSF-funded research facilities
  - Multi-user, multi-science facilities

Context - MFs and Computing

- MFs require a certain level of “dedicated” CI resources to complete their mission
  - Who provides that CI depends a lot on the MF
  - “Normal” ACCESS model with an allocation and shared resources doesn’t fit
- There are many different MFs with different missions and operating principles
  - From “vertically integrated” (experiment to final result/paper) to a resource/data provider
- Question of mission: LHC vs. LSST
  - LHC has a science mission and is closed community
    - Vertically integrated from experiment to final result/paper
    - LHC experiments need to provide resources to create data products and perform data analysis
  - LSST is a science/discipline community MF
    - LSST is a resource like an ACCESS resource – Astronomers request time, request data products, etc.
    - LSST computing needs are focused on providing data products and alerts
    - ACCESS will support individual scientists using LSST data, not necessarily the “LSST collaboration”

What will we cover?

- Experience of MFs with ACCESS
- How could ACCESS and MFs work together
- What is missing in the ACCESS portfolio from a MF perspective

MFs & ACCESS – Computing Paradigms

- MFs are not “classical” HPC
  - Limited to no tightly coupled workloads that require >1 node
  - Ingeniously/pleasingly/embarrassingly parallel workloads
  - AI (and distributed AI training) is a recent development
- MFs are data-intensive and (mostly) distributed
  - MF workloads can be organized chaos – Flying the plane while replacing the engines
  - Big data in, big data out – MFs move data in and out for every workload
  - ACCESS orients data movement around “do computation locally and move output data once you are done”

MFs & ACCESS

- Non-standard user conundrum
  - Large userbase hidden behind MF, see LHC, IceCube, LIGO
  - MFs are odd balls in the NSF portfolio – Not your standard single-few-several PI projects
  - Using > 1 ACCESS resource needs to be well motivated
- Non-homogeneity of MF application
  - MF application behavior changes significantly between workloads
  - Containers, containers, and more containers
- Non-homogeneity of ACCESS resources
  - Every ACCESS resource is slightly different – Login (MFA), software stack, available resources, policies, etc.
  - Large complex software stacks and need to be moved around and are centrally managed
  - Takes MFs effort to integrate, is the effort worth it?

Working Together

Lots of Opportunity!

ACCESS resource providers have all resources to build CI for an MF

Working Together

Need a virtualization environment? - JetStream2, (CyVerse, Chamelon)

Need storage? - Open Storage Network, etc.

Need GPUs? – Delta, Expanse, Bridges-2, etc.

Need CPUs? – “All”

Working Together

- What is the hold up?
  - Mission
    - ACCESS are **multi-user HPC** facilities
    - Computing Paradigms
    - Funded through CISE/OAC and not MF directorates
    - NSF internal question of funding and funding cycles – MFs are multi-decade instruments vs. 5-10 years for ACCESS resources
  - “Another user”
    - MF comes in as a “user” – Multi-user facilities
    - MFs are 24/7/365 in terms of science operations – IceCube uptime is > 99% (over the 12 years)
    - Are MFs different then other NSF PIs? – Not to ACCESS
  - Question of support
    - MF researchers expect dedicated support that is more than what ACCESS providers are used to or can provide without additional resources

Working Together

- Better cross-resource support and homogeneity – Low Hanging Fruit
  - Why can I use apptainer/cvmfs on resource A and not B?
  - MFA policies more suited towards remote submission
    - Collaboration with ACCESS, PATh, and MFs
- AI Resources
  - Most MFs can't afford/have access to/justify large scale AI hardware, i.e. more than single 8x A/H100 machine
- Hosting CI
  - ACCESS resource providers could be be hosts of MF CI
    - Example OSDF caches at SDSC, Internet2 backbone, used by LIGO, LHC, etc.
    - MF buys extra server(s) and hosted/admined to ACCESS resource
  - Economies of scale
  - Requires dedicated support

What is missing in the ACCESS portfolio

- Archival
  - Economies of Scale
  - NERSC in DOE – One big tape archive is easier than 10 smaller ones
- Easier support for “distributed” workloads
  - Solutions for IceCube, LHC, and LIGO are there just not the personpower to implement them
- Common policies/setups across ACCESS sites

Thank you!

Questions?