SCDF Stakeholders and Users Technical Meeting

June 12, 2025

Action items from May 8th meeting

- Some general considerations with regards to access to SCDF/SDCC resources and services
  - Most SCDF resources (computing and storage) are dedicated to programs. The SDCC/SCDF currently has limited (~18k computing cores) available for opportunistic (via HTCondor) access. There are no storage resources available for opportunistic access.
  - Services (including some licensed software) are more accessible
    - Collaborative tools (MatterMost, InvenioRDM, BNLBox, Overleaf, Gitea, etc)
    - User-facing support (new accounts, web, OpenShift vm's, etc)
    - Lightweight (no access to core resources) user accounts
  - Remote access to new platforms (AI/ML, non-x86, etc) possible through existing collaborative partnerships with Cloud service providers (ie, Google) or hardware manufacturers (ie, HPE, Supermicro, etc)
    - If interested, discuss with SCDF management
- Questionnaire sent to stakeholders on May 30
  - Compiling responses to help guide potential requisition of IC gen3
  - Preliminary results shown today – fuller picture if more stakeholders participate
  - Currently writing a draft for PD proposal.

Action Items from May 8th meeting

- How can PI's obtain guidance on funding for SCDF/SDCC services in future proposals? Some considerations:
  - Hiring additional staff on 'soft' money (LDRD, PD, pre-PD, etc) is difficult, so expectations of support for short-term projects need to be realistic.
  - Generally speaking, leveraging existing staff expertise and SCDF/SDCC services minimize cost and delays. Proposals with exotic hardware architectures or requirements for new services increase costs and incur higher risks of delayed/unavailable support. Some examples:
    - AlmaLinux instead of Debian or Windows
    - Gitea instead of GitLab
    - Apptainer instead of Docker
  - Please discuss with the SCDF/SDCC contact for your organization (see https://www.sdcc.bnl.gov/information/sdcc-organizational-chart) to request assistance (or contact SCDF management, if your organization is not listed).

Questionnaire results (preliminary)

- Received feedback from CFN, ATLAS, DUNE and HET/USQCD
- Preferred configuration is Intel/AMD x86 with gpu's and low-latency interconnect (ie, Infiniband)
  - Nvidia (preferred) or AMD gpu's
  - 4-8 GB of RAM/core
  - Modest (up to ~100 GB) transient disk storage (per job) needs
  - Moderate (up to ~1 PB) of persistent disk storage needs
  - Wide range of computing needs (~140k – 30M core-hours per year)
  - Limited interest in institutional tape archival storage
  - Software infrastructure
    - HEPIX/WLCG software stack
    - MPI + CUDA/HIP
    - AI/ML support
  - Near-future needs (Fall 25 and onwards)
- This fits well within current SCDF/SDCC resource deployment and expertise.