

SDCC/SCDF resource access

May 29, 2025

Tony Wong

Background

- BNL is a multi-disciplinary lab with a few core central services (utilities, administrative support, enterprise computing, site security, etc). The list of core central services has not changed for a long time.
- Previous efforts to organize central support for scientific computing have not lasted, primarily due to an organizational model that fosters few inter-department or inter-directorate activities.
 - At least 3 such initiatives fizzled in the past ~25 years, by my count
- In this environment, the RCF was created by the Physics Dept. after ITD turned down the request to support RHIC's scientific computing needs in the mid-90's.

SDCC/SCDF – current status

- RCF, RACF, SDCC and now SCDF provides scientific computing resources and services to various communities at BNL
 - Historically and (still) primarily optimized to support experimental programs within Physics Dept. (RHIC, ATLAS, Belle-II, DUNE, etc)
 - Also, support for other communities (CAD, CFN, CMPMS, CSI/CSD, NNSD, NSLS-II, NSTD, RIKEN, etc)
 - IC gen2 (and previously IC gen1)
 - Hosting and IT services for group and dept.-sized clusters
 - Websites
 - Authentication & Authorization Infrastructure (AAI)
- Funding for SDCC/SCDF comes primarily (~83%) from Physics programs, and staff effort and targeted activities are aligned accordingly.
 - Majority of the remainder (~17%) support the other listed communities above

(Opportunistic) Access to SDCC/SCDF

- Resources (computing and storage) are usually dedicated to programs. For opportunistic access, the SDCC/SCDF currently has some limited computing and storage resources:
 - Shared (older) pool (~18k cores) available for HTC access
- Services are more accessible
 - Collaborative tools (MatterMost, InvenioRDM, BNLBox, Overleaf, Gitea, etc)
 - Web support
 - Lightweight (no access to core resources) user accounts

Constraints and future activities

- Requests for support have a higher chance of success if it can leverage existing expertise and effort. Some examples:
 - AlmaLinux9 instead Windows or Debian
 - Puppet instead Ansible or Chef
 - Mattermost instead Discord
- Often, available funding > available FTE
 - We normally only hire staff on programmatic (stable and predictable) funding. This limits the usefulness of LDRD, PD, pre-PD, etc funding in expanding staff power.
 - No lab-wide generic support for scientific computing outside of program funds
- I will send out a questionnaire later today on the feasibility of a potential IC gen3
 - Potential to re-establish a true institutional resource
 - Aim to strike a balance
 - Leverage staff expertise and minimize additional effort (if FTE count does not increase)
 - Lower barriers to access (specially to new communities)
 - Provide resources useful to most (all?) communities