



NPP/CSI SDCC Operations meeting

May 26, 2023

    @BrookhavenLab

Agenda – for Today

- What are CSI plans for new systems: IC gen2, SciServer ?
- Retirement of old systems and transition to new systems needs planning
- SDCC effort to support CSI systems
- Data Center associated costs (Space, power and cooling) and implications
- CURL requests

Agenda – for Next meeting

- Midway HPC planning and operations
- CURL update
- TCO update

CSI plans for IC gen 2 & SciServer?

- SDCC effort for these activities have been *temporarily frozen* until further clarification

Retirement / Transitions

- Retirement of old systems and transition to new systems need planning and effort
 - List of CSI systems to be retired in backup slides
 - IC gen1 to gen2 – requires operational overlap (of at least 2 months)
 - Orderly software and user migration
 - Admin effort during transition
 - Future SciServer dependency on retirement of IC gen1 (repurposing hardware)

SDCC estimated operational effort

- SDCC effort required to support IC Gen2 & SciServer
 - IC gen2 – cluster administration (compute, disk and tape storage)
 - 0.75 FTE (additional 0.25 FTE for disk/tape storage)
 - Effort must grow if IC gen2 grows
 - SciServer – procurement, hosting and installation
 - 0.25 FTE (procurement, hosting and installation)
 - 1.5 FTE (operations)
- Need to forecast beyond FY24
 - Align staff effort with other priorities
 - Advanced forecasting is essential for staff management

Data Center Costs and Occupancy

- Data Center Costs (space, power & cooling) - New for CSI

- TCO calculations depend on physical (725 or 515) placement and timelines of new CSI-initiated resources

- IC gen2
- SciServer

CSI - B725	Space charge and electricity
FY22	\$0
FY23	\$130,000

- Implications for data center floor availability – Planning is needed

- Placement of SciServer and/or Midway Partnership resources may impact future expansion plans
 - IC gen2 (FY24 and beyond)
 - sPHENIX (FY24-25) and ATLAS in the LHC-HL regime (FY28 and beyond) if Power System #4 is not deployed in B725/MDH as of FY25
 - Back-up slides provide more details

CURL – SDCC list

B/725 - Infrastructure

- 1) B/725, Backup Diesel Generator - ~\$1.5M
- 2) B/725, Power, and Cooling Infrastructure Expansion - ~\$10M

B/515 – Network room (ITD) and Sigma-7 room (SDCC) – Electrical work

- 3) Decommission APC MegaWatt Battery UPS - ~\$250k
- 4) Swap Electric Feeders from Mega Watt UPS to Fly Wheel UPS - ~\$450k
- 5) Replacement of Bard AC units (was an F&O CURL that did not get funded) - ~\$150k

Back-up slides

Transition Details

- IC gen1 and gen2 must operate simultaneously for a period of time (at least 2 months)
 - Finalization of NIC configuration for IC gen.2 master and submit nodes
 - Migration of software to gen2
 - Migration of data (1 PB) from IC gen1 GPFS to IC gen2 GPFS
 - Migration of IB connected Globus DTN from IC gen1 GPFS rack (B515/BGL) to IC gen2 GPFS rack (B725/MDH)
 - Migration (and debugging) of user activity to gen2
- SciServer is dependent on retirement of IC gen1 / ML cluster equipment
 - Retirement of gen1 (hard deadline of Sep 30, 2023) is dependent on gen2 operational availability
 - Composition and size of SciServer cluster impacts availability timeline
 - Need to estimate power and cooling
 - Purchase of additional hardware is conditional on size of SciServer cluster

CSI equipment in 515

- Systems retiring by Sep 30, 2023:
 - K80 GPU based racks of IC gen1 (6 racks in B515/BGL)
 - Skylake CPU racks of IC gen1 (3 racks in B515/BGL)
 - KNL cluster along with KNL / IC gen.1 infrastructure rack (5 racks in B515/BGL)
 - HPC-1 legacy system (3 racks in B515/BGL)
 - Legacy portion of KBASE setup (2 racks in B515/QCDOC)
- Systems (being kept past Sep 30, 2023)
 - Test CSI systems (in Sigma-7)
 - CSI 2-frame IBM TS4500 tape library (with 8 LTO8 drives) and associated HPSS mover systems (in CDCE)
 - SciServer system (converted portions of IC gen1 / ML clusters - to be placed in CDCE)
- Implications for data center charges
 - CDCE and Sigma-7 are currently 100% charged to NPP, so modifications will be made to accommodate CSI resources starting from FY24

CSI equipment in 725

- Systems:
 - Composition of IC Gen2
 - KBASE
 - Recently purchased hardware currently used as SciServer testbed located in B515/QCDOC proposed to be moved in B725/MDH as requested by Shinjae's group (needs further discussion)
 - SciServer system
 - Part of the system based on new hardware purchased in FY23 to be placed directly into B725/MDH

Implications for data center floor space and charges

- Increased footprint of CSI-initiated activities will have budgetary impact
- Need mid-term CSI plans for occupancy of MDH – may impact support for near and mid-term NPP priorities (sPHENIX, ATLAS, etc.)

Atos system

2 Atos Sequana XH3000 Racks (284 A each = ~136 kW max each)

1 Management Rack (64 A = ~13 kW max)

2.5-3 FTE to operate

Estimated space charge: \$4,200 w/o OH per year (11 rack positions equivalent in the MDH HD area)

Estimated electric charge: ~150 kW @ \$0.085kWh = ~\$9,300/month (\$112k w/o OH per year) assuming 100% load