External Review of the ROAR Service for the Penn State Institute for Computational and Data Sciences

March 29, 2024

David Hancock, Indiana University
Richard Knepper, Cornell University

Commissioned by: PSU SVPR Andrew Read and ICDS Director Jenni Evans

Executive Summary

In late 2023, SVPR Andrew Read contacted the review team in order to provide an external review of the ICDS ROAR computing resource that would offer suggestions on how to improve the cost recovery performance of the system, make recommendations about architecture and services in support of research computation, and provide an analysis of the organizational structure supporting ROAR aimed at improving the faculty experience of the ROAR resource. The review team met with ICDS leadership, staff, and Penn State faculty who make use of computational resources (both ROAR and other solutions), as well as representatives from other units at Penn State with experience in the cost-recovery and service delivery.

The review team proceeded with collection of additional information and received a large number of ICDS documents for background on the system in addition to meetings which spanned from November 2023 through March 2024. The review team visited Penn State for a site visit in February, which allowed for both free-ranging panel discussions with the ICDS team and faculty with computational needs, as well as face-to-face conversations with stakeholders about ROAR policies, functions, architecture, and utilization. The office of the SVPR conducted a survey instrument which provided broader detail from ROAR users.

There are significant research accomplishments achieved by Penn State being supported by ICDS and ROAR services but the current recovery model and delivery mechanism are in need of adaptation in order to provide long-term sustainability for all stakeholders (ICDS, SVPR, and Faculty). Challenges with the current model have been compounded by insufficient communication with the research community, a lengthy resource transition prolonged by pandemic circumstances, perceived inflexibility of Service Level Agreements (SLAs), and fluctuations in past University leadership and direction (or lack thereof).

The interdisciplinary research aims of ICDS are unique, and through it collaborations are happening that would not naturally occur. The dual mission of fostering such research, as well as operating broad core services is particularly challenging, but appears to be the best fit in terms of research computing services at Penn State. The operational aspects of research computing services need an elevated focus within the center to benefit the wider community in addition to re-engaging the community more directly when setting service direction.

In aggregate the participants in this process have shown a great affinity for the success of Penn State and a desire to improve services and research outcomes. The strategy and tactics derived from these findings and recommendations should yield increased satisfaction and greater opportunity for the institution as a whole.

Introduction

In October of 2023, SVPR Read contacted Knepper in order to discuss challenges at ICDS with recovery and customer satisfaction around the ROAR HPC system, with change at the SVPR level identifying significant budget shortfall for the ROAR service, negative faculty response to a proposed recalculation of ROAR rates and the inclusion of Facilities and Administration rates on ICDS services, and other expressions of faculty dissatisfaction with the ROAR service pricing and capabilities. Based on discussions with SVPR Read and at his request, Knepper and Hancock engaged in a review of the ROAR service within the context of ICDS, in order to deliver this report, with the goal of identifying issues behind the ROAR cost recovery service and providing recommendations intended to help ICDS support a greater range of faculty, align service costs with unit costs, and improve utilization of the resource.

Scope of Charge

For the purposes of this review, SVPR requested an overview of the ROAR service, its cost model, utilization, and responsiveness to faculty needs. Operations within ICDS as the overall organization providing the service were also subject to review inasmuch as these activities support the engagement of new faculty users, support for current users of the system, and partnering opportunities for extending and improving capabilities of the ROAR system.

Based on the charge, the review engaged leadership and staff across ICDS in order to understand a broad range of elements that inform the ROAR service within the context of the institute: the vision and mission of the institute and how they inform the ROAR service; communications that engage new ROAR users and inform them about the use of the cluster; the business administration processes and practices behind starting an SLA, how rates are developed, and the current recovery profile of the service; ROAR system architecture, hardware refresh periods, queue and scheduler configuration, and utilization; operations and support, including helpdesk functions; and relationship between ROAR and the RISE service for research software engineering support.

Review Team

Richard Knepper, PhD is the Director of the Cornell University Center for Advanced Computing (CAC) since 2022. CAC is a computational core facility for Cornell faculty that provides compute, storage, and consulting on a cost recovery basis. Knepper's research focuses on virtual organizations for provision of computational services (such as the XSEDE project), cross-institutional federation and collaboration, and measurement of cyberinfrastructure services, including return on investment.

David Hancock is the Director of Advanced Cyberinfrastructure, Research Technologies of University Information Technology Services at Indiana University. Hancock has considerable experience in research operations and high performance computing architecture, storage, and service management and provisioning. Hancock has published numerous articles on return on

investment in addition to conducting local analysis. He currently leads the Jetstream2 project in addition to local IU HPC and Research Cloud Operations and is Co-PI and Senior Personnel on numerous other NSF-funded projects.

Hancock and Knepper have worked together in one capacity or another since 2000, supporting faculty with computational services, collaborating on cyberinfrastructure projects and NSF-funded research, and managing teams together and within their respective organizations.

Research Computing Cost Recovery Models

There are three general classes of models within academic research computing: institutions that provide most services without fees (or chargeback only begins after a significant usage level, e.g. Minnesota, Indiana), institutions that attempt to fully recover costs, and hybrid, subsidized approaches. Smith et al[1] provide a brief background of approaches and prior analysis. Since that publication the trend of full cost recovery has been further reduced with no-cost or hybrid models dominating. In a hybrid approach there are two common subsets, institutions that bundle all staff and equipment into a subscription/usage rate with some level of subsidy (e.g. Cornell, Michigan, OSC), and organizations that strictly cover staffing and operating costs with charges for compute/storage costs only (e.g. Purdue, Florida). The University of Michigan is a good example of an organization that started with almost full cost recovery and shifted to a higher level of subsidy and no-fee research packages over time.

Data Collection

Interviews

Hancock and Knepper conducted 39 interviews with respondents both inside and outside of ICDS, beginning with the director, assistant and associate directors, a number of faculty both users of ROAR and those making use of other solutions, managers of teams within ICDS, and representatives of other entities at Penn State that include individuals from the Huck Institute, Social Science Research Institute, Research IT, Center for Applications of Artificial Intelligence and Machine Learning to Industry (AIMI), and School of Medicine. The bulk of these interviews were conducted via zoom, a few of them were conducted face-to-face at the International Conference for High-Performance Computing, Networking, Storage, and Analysis (SC23) where members of the ICDS team were in attendance. Additional discussions occurred during the February 12-13 site visit described below. Interviews were scheduled for, and largely completed within, one hour, although a few discussions went over time if there were particular issues respondents felt should be highlighted in detail.

[1] https://doi.org/10.1145/3332186.3333155

The format of the interviews tended to vary depending on whether the respondent was part of ICDS or other institute leadership, ICDS team members, or faculty, but all interviews started with an overview of the respondent's time at Penn State, functions and activities in terms of context of ICDS. The reviewers allowed conversations to range freely in order to ensure that respondents were able to elevate issues which they felt to be important and if there were certain areas that were mentioned with special interest as regards the external review, directed questions followed in order to reach a sufficient level of understanding on the part of the reviewers. Faculty members were asked to provide a brief overview of their research in terms of application, computational and storage needs as well as their selection of ROAR or alternative resources. Discussions with leadership from other institutes or entities at Penn State focused on the strategies and activities that those organizations implemented to success, including other Co-Hire programs, communications methods that were beneficial for reaching PSU faculty, and interactions and collaborations with ICDS for computational services.

Respondents invariably provided open and direct commentary in their answers. The reviewers assured all respondents that interview discussions were confidential and that findings were to be provided without attribution. Nearly every respondent expressed the desire that their responses would contribute to the review and support the improved performance of the resource. Respondents frequently mentioned the benefit of performing an external review, their hope for positive outcomes for ICDS, and improvements on offerings for Penn State research outcomes via cyberinfrastructure offerings.

Site Visit

The review team came to State College during the week of February 12 and conducted a two-day site visit that allowed for both broader in-person conversations with groups of stakeholders as well as some discussions in-context with individuals. Despite snow delays, ICDS and faculty were well-represented and the reviewers were able to provide progress reports to SVPR Read on-site.

The site visit consisted of:

- Panel session with staff and management of ICDS
- Meetings with ICDS Director Evans, AIMI Director Kumar, SVPR Read
- Finance deep dive session with ICDS administrative staff and leadership
- Open panel meeting with invited faculty members
- Technical deep dive session with ICDS leadership and engineering team

The open format provided opportunity for frank discussion and allowed conversations to range freely across topics. The reviewers were able to use time in between sessions to identify areas that needed to be revisited in order to inquire about further details.

Document analysis

The review team were provided a number of documents upon their request related to ICDS operations, including current organizational charts, budget overviews, rate calculation documents, service rate sheet, sample ROAR SLA documents, and ROAR job and utilization data. These documents supported the bulk of questions for the finance and technical deep dive session during the site visit. Also provided were a previous benchmark analysis comparing offerings at Penn State to peer institutions collected by the RCCI Executive Committee, as well as communications to SVPR Read on the F&A waiver discussion, and SVPR Read's own notes on the situation.

Survey

At the request of SVPR Read a lightweight survey was conducted over two weeks starting on February 23, 2024. The respondents varied from students, staff, and faculty with 249 total responses. The comments and interest were robust with what appears to be demand and high levels of interest to provide input with a significant number of respondents, 25% (63 individuals) agreeing to be contacted for follow-up by ICDS. It is important to note that in our discussions with ICDS staff and leadership that there is currently not a formal survey process for tracking satisfaction with ICDS services. This basic survey instrument was a mechanism for the review team to gather broader input and capture details that could not be accomplished via individual interviews in the allotted time.

Findings

Interviews and Site Visit

Interviews with Penn State faculty identified a number of issues around the ROAR service model which represent challenges to success. One of the largest sources of tension arises from issues of communication between ICDS and faculty users. Users reported a number of areas where information about ICDS services and structure was incomplete or there were misconceptions about services; these included misunderstandings about SLA length and support, about the cost of providing operational services, and about the scale of ICDS staffing. A number of the interview respondents indicated that the compute offerings from ROAR are sufficient for their needs, and while some of the respondents indicated that they made use of ROAR's bursting capability, the faculty members that we interviewed did not view the one-hour job start guarantee as a critical factor for use. In discussions with faculty, topics frequently followed computing practices from before the establishment of ICDS or before the implementation of the ROAR cost model, making it clear that researchers tend to anchor their impressions based on their first interactions with the Penn State research computing environment, and that subsequent changes were not fully internalized; some further communication to bring people up to speed seems warranted. Another area that could be helped by further communication is in system architecture and provisioning choices. A number

of faculty felt that ICDS' purchasing and provisioning decisions for ROAR compute and storage were opaque, too costly, or otherwise insufficient for meeting the needs of the user community. As the team begins the next process to refresh hardware resources, it is important to use this process to build trust with users; solicit information that informs requirements and communicate both the benefits and compromises of purchase decisions clearly.

While there were respondents who felt that ROAR costs are too much to pay, a number of respondents noted that they were able to incorporate ROAR service costs into their budgets in order to support necessary research. Certain users noted that they maximized resources by balancing use of the Open Queue with paid SLAs so as to stretch the available resources to the maximum, which seems to be an efficient use of system capabilities; however, not all respondents have workloads that can take advantage of the Open Queue successfully, due to long run times or other factors.

In discussions with ICDS team members, conversations continued to highlight communications as a challenge to successful operations. The team has had a difficult time establishing clear and consistent messaging with the faculty and building trust and understanding across faculty units. While efforts to improve user communications are ongoing, this is a strategic activity that takes a considerable amount of time to show results. Within the Institute, communication appears to be difficult across intra-organizational boundaries, making it hard to get information across team boundaries and up the chain of command, including important feedback. Based on feedback from interviews, this affects operational service delivery activities (such as trouble ticket handling, project and task prioritization, and SLA maintenance) but also internal governance activities. Team members noted that lack of transparency around decisions and processes made for confusing situations; this issue extended into team members being unsure about how they might deal with HR reporting processes. While there are periodic retreats and meetings to provide communication and feedback, there appear to be difficulties building and maintaining a mission-driven culture within ICDS.

Document Analysis

ROAR Core Facility Budget

The ROAR core facility budget provided show shifts from direct purchases to leasing as well as an overall reduction in total expenses from FY23 actuals to FY24 and FY25 projections. The budgets provided are inclusive of RISE services as well. Total income from SLAs are projected to increase from ~$2.3M in FY23 to ~$2.7M in FY24 but well below the original expectation of ~$3.6M. Staff expenditures show a modest decline of $30K over that period with non-personnel expenses showing a $1.4M decline (through reductions in contractors, software & maintenance, and equipment costs).

Rate Setting Worksheets & Explanation

Rate setting at Penn State is challenged by institution specific rules (e.g. unallowable chargeback for lease/equipment investment within the institutes) that adds complexity to the rate and budget detail provided. ICDS was able to produce a static mapping of individual assignments for personnel included in the ROAR or RISE rates that appear appropriate. The challenge in grasping rate detail is that personnel and maintenance expenses are further subdivided into assignment on specific products in the rate worksheet resulting in rates that are not aligned with the equipment cost of a product and do not allow for dynamic fluctuation in effort between services. Rates are currently calculated based on projected sales with a fixed subsidy assigned per product, this can create a negative feedback loop where reductions in projected sales will result in rates increasing (if personnel/equipment cannot be quickly reduced) that may further reduce income. Although the rates have been recalculated and submitted annually as required, the cost to researchers has not changed in 8 years through varying and increasing the subsidy from SVPR. This created a situation where changes in rates to researchers in order to reduce the central subsidy resulted in dramatic increases.

Service Level Agreements

The service level agreement documents are very clear, structured, and professional. The detail provided is at the level that is applicable for external contracts and grants. While in general this is appropriate and welcome it can imply a level of inflexibility that makes researchers uncomfortable or feel locked into long-term agreements. The impression/feeling of the SLA process is not in alignment with actual business practices by the institute (ICDS will gladly modify or create shorter-term agreements and charges are incurred monthly).

Organizational Charts

Org charts were provided which provide a high level overview of ICDS structure as well as detailed personnel assignment and their percentage commitment, if any, to ROAR and RISE services. A number of vacancies exist across the organization and in general the current level of staffing seems appropriate for the size of the PSU community and resources being managed by ICDS.

Cluster Utilization

Processor core utilization detail was provided for a two year period for ROAR and ROAR Collab in graphical form. While the details are challenging to interpret without the context of ICDS technical staff and understanding the transition of nodes between the systems, it appears that both clusters are underutilized. A variety of factors contribute including guaranteed job start times for SLAs, overestimation of needs by new faculty provisioned via startup packages, SLA purchases sized for peak workflow with bursty usage, and modest limits on Open Queue usage in terms of job duration and size.

Job submission details were also provided for a longer period which factor into our recommendations but the high level nature of those data are not insightful without further, ongoing analysis.

Web accessible material

ICDS provides clear and understandable details for grant facility statements and data management plans that can be leveraged by ROAR consumers. The rate sheet provided to the PSU community is clear in itself, the challenge comes in lack of understanding in how the rates are calculated and what level of subsidy is being provided. There are also basic user guides and system documentation but many of the details need a refresh for the changes to ROAR and ROAR Collab.

Survey

The review team recommends that ICDS leverage the data from the survey instrument as a springboard to develop a regular survey mechanism, and specifically take advantage of the significant number of individuals who are willing to be contacted for follow-up in order to guide the evolution of resources and services on offer.

The survey included questions about respondents' current usage of the ROAR resource as well as, attitudes about costs of ROAR compute and storage resources, the level of communications from ICDS, responsiveness of the i-Ask help desk, and the impact of maintenance periods on work. Survey questions also inquired about whether current costs for compute time and storage are competitive, and how respondents would change usage if the resources were made less expensive. Respondents were also encouraged to leave comments on the system as well.

From the responses, a large number of those surveyed were users of the Open Queue (61%) while 35% had a current SLA, 44% were using resources in their department or lab, 20% using resources at another university or facility, and 9% using public cloud resources (answers were non-exclusive, so that any person answering could be using more than one of the above resources). Based on the comments, there were a significant amount of survey responses from grad students or post-docs within labs, which may have impacted some of the answers, as these respondents don't have control over spending or visibility into costs. This is further borne out by the responses to questions about costs for compute and storage, with around 44% "don't know" responses; 39% of respondents felt that compute costs are either too high or somewhat high; while 30% felt that storage costs are too high or somewhat high, a larger percentage felt that storage costs were just right (12% as opposed to 8% for compute), although satisfaction with the offerings is solid, 50% stated that they agreed with the statement "I am able to get the kind of computational resources needed via the ROAR system." Overall responses characterized a strong tendency towards faster job start times (51% "within an hour" and 15% "within a day") which contrasted with responses gathered in interview sessions, where no one reported the need to start within an hour. The difference in responses may be that our interviews

focused on paying customers of the ROAR resource, while the survey had a large portion of Open Queue users.

Comments from the survey focused on support of further free usage, longer job run times, and increasing core allowances, as well as dissatisfaction with the potential of charging facilities and administration on ICDS costs. Given the large number of users that reported use of the Open Queue, it is likely that answers will tend towards support for more free resources. A number of the users who are paying for resources noted that the SLA model doesn't align well with the cycle of research activities, so that researchers on sponsored funding may be paying for access during periods when there is very little usage, and then require time for further computational analysis after funding wraps up; the cyclical nature of research funding does not line up with the month-to-month operational expense of the SLA model. This misalignment could be improved with a subscription or pay-as-you go model.

In terms of responses to "I am able to get the kind of computational resources needed via the ROAR systems," comments largely fell into a few areas: dissatisfaction with costs (discussed above), issues with software (missing packages or capability needed by researchers), need for additional storage, or need for more GPU resources. A number of these comments referred to issues with training or onboarding materials that would facilitate use of the ROAR resource. A small number of respondents (less than 5) noted that the lack of a secure system for dealing with regulated data was a barrier to research.

Updated cost benchmark analysis

The review team conducted a brief updated benchmark overview based on publicly-available data and discussions with peer institutions that are roughly close to Penn State (public R1 institutions with large populations and similar R&D expenditures). Cost details are described below, with Penn State at the end of the list for comparison purposes.

These data were collected in March 2023 and reflect the most recent fiscal year available. We did not collect all service offerings, only basic/standard CPU nodes (normalized for memory where applicable as month per core for ~8GB of RAM), GPU pricing in full GPU per year increments, and storage (high performance, standard, and archival where available are in terabytes per year).

- University of Florida (no indirect cost (ICR) charged), prices escalate by ~10% for 1-year term instead of 5-year term.
    - $200 per core w/8GB ram (5 years) -> $3.33/mo/core
    - $1200 per GPU for RTX2080ti or 6000 GPU (5 years) -> $20/mo/GPU ($240/yr)
    - $125/TB for archive/long term storage (5 years) -> $25/yr/TB
    - $625/TB for high performance storage (5 years) -> $125/yr/TB
    - Secure system (HiPerGator-RV) has >2x premium for cores and 8x premium for storage (although it is replicated so this cost is an upcharge of 4x in reality)

- University of Michigan (they note a 35-40% subsidy on the rates below)
    - As part of a recent UM Research Computing Package that provides some central funding to most faculty members the following resources are available before incurring fees:
        - 80K CPU hours annually
        - 10TB of replicated high performance storage
        - 100TB of archival storage, access to a secure enclave with up to 16GB of RAM.
        - Some schools further buy in to share cost increasing these quotas (e.g. Medical School matches PIs 9:1).
    - Basic core (same size as UFL) $10.82-12.53/mo
    - Colocation service for direct integration into a shared cluster $743.07/node/yr ($61.92/node/mo)
    - GPU (V100 or A40) $118.33 - $121.64 per mo -> $1420 - $1460 per year
    - High performance storage $156.22/TB/yr replicated, $77.88/TB/yr unreplicated
    - Long term storage $54.37/TB/yr replicated, $27.36/TB/yr unreplicated
    - Long term archive replicated $20.04/TB/yr
- University of Minnesota
    - The compute resource is provided without fee as it's centrally funded, 125 million hours free annually in aggregate
    - Max free storage tiers are ≤20TB per group high performance, ≤120TB second tier before charges
    - Internal costs are in full node increments for 1 or 5 years (similar to UFL) but break down for comparison as follows (charges below ICR exempt):
        - Basic w/4GB (UFL/UMich) $1.6/mo/core for 5 years or $1.71/mo/core for 1 year.
        - Normalized Basic for 8GB RAM $3.20/mo/core for 5 years or $3.42/mo/core for 1 year.
        - GPU (A100, A40 is about ½ the A100 cost) $2286/GPU/yr or $190.48/GPU/mo for 5 years $2330/GPU/yr or $194.19/GPU/mo for 1 year
        - High performance storage $122.58/TB/yr
        - Second tier storage (disk but not archive, S3 via Ceph) $40.53/TB/yr
        - They no longer offer archival storage services
- Ohio Supercomputing Center (serving higher education institutions in the state)
    - OSC academic members receive $1000/yr credit per PI, the bulk of funding for OSC comes from the State of Ohio and commercial clients. The rates represent only 20% of the actual cost (subsidized at approximately 80% from those other sources)
        - Basic core $2.1/mo/core but charged as a per node hour (full node increments), also offer a subscription at a fraction of the cost ($.003 per core hour, 45x less)
        - GPUs (V100, rates do not reflect newer A100/H100 available) $1402/GPU/yr or $116.80/GPU/mo

        - Project storage (GPFS mid-tier performance but backed up) $38.40/TB/yr
- Purdue University
    - No ICR is charged on these rates, prices are per node or per GPU for 5-year terms. You must purchase a node to have access to the community cluster(s).
    - These rates are not public but shared by the RCAC director in confidence for this exercise. The institution provides a subsidy for all the operations and support staff in addition to data center space, racks, and networking to reduce the cost per node.
        - Basic node $0.55/mo/core for 5 years (2GB of memory per core)
        - Normalized Basic for 8GB $2.20/mo/core for 5 years
        - GPU (A100) $1400/GPU/yr for 5 years (H100 likely to jump to $5400/GPU/yr after their next rate approval)
        - Storage (mid-tier disk) $70/TB/yr
- PSU
    - Standard core (compares to normalized memory profile) $7/mo/core
    - GPUs (A100) $4800/GPU/yr
    - High Performance Storage $80.04/TB/yr
    - Archive Storage $14.40/TB/yr

Conclusions of benchmark analysis

The majority of peer institutions do not charge indirect cost on their research computing services and provide some level of access at no cost. UMN and UMich provide the most resources without fees followed by OSC. Purdue and OSC are subsidizing their rates most significantly for computation with UFL providing the lowest subsidy but still significant, their latest compute system (HiPerGator) was provided by a cost matching donation to the institution ($25M donated, $25M from NVIDIA, and $20M contributed by UFL).

Cost for GPU resources is rising significantly (both on-premise and in the cloud). This can be observed by comparing rates of different GPU generations listed (P100/V100/A100/H100) in the benchmark analysis. Note that there is a high degree of GPU cost variance which reflects the different generations of hardware as the cost per NVIDIA GPU has risen 2-4x each generation (UFL, UMich, & OSC have older or non-enterprise cards).

While cloud spot pricing can be attractive for preemptable CPU workloads, the GPU rates overall are an order of magnitude above academic prices. Current EC2 pricing for a p4d.24xlarge with 8xA100's is $35,910/GPU/yr for an on-demand instance, $26,529/GPU/yr for a 1-year reservation, and $11,895/GPU/yr for spot instances that will be frequently interrupted.

The storage rates are more comparable across institutions when one factors in replication and performance levels. PSU compares very favorably to both high performance and archival storage rates at peer institutions.

There is not a standard rate model among these institutions, nor are the authors aware of a research computing standard rate process across institutions. It is clear from discussions with

multiple institutions that no-cost allocations for compute and storage tend to get used to the utmost capacity available and that structuring free allocations should be done in a judicious manner that balances the goals of the allocation (typically providing enough support for researchers to complete some research that is sufficient to justify external funding for additional compute and storage if necessary) with the overall resources allocated to the operation.

Recommendations

The following recommendations are not in a prioritized order but arranged topically and can often be considered individually allowing ICDS leadership to prioritize implementation based on institutional direction and strategic goals.

- **Service Rates (multiple variations listed below)**: The annual rate process should be an evolutionary one with a willingness to tweak the model based on researcher input. Static rates and lack of guidance from prior leadership at PSU have created a disconnect between the original intent/model of ICDS and current research needs. In general, the SLA process is understood by faculty to be a rigid one and ICDS to be a flexible one; it is important that the SLA process be based on reaching agreement and finding solutions rather than holding positions.
    - **Modest Change**: When setting rates, even if hardware/lease costs cannot be included in the rate, the process should take into account the cost of a given type of compute or storage (e.g. if a P100 and A100 take the same amount of effort to manage, the chargeback rate should mirror the difference in cost of the equipment).
    - **Moderate Change**: Calculate rates using total inventory, instead of projected sales, and make clear what the subsidy is actually providing. This can be combined with the prior recommendation.
    - **Extreme Change**: Make a subscription model available, perhaps as the only model. Researchers are charged based on what they consume and can buy ahead in bulk if desired, with a system similar to a Starbucks card.
        - The subscription model would mitigate unused startup capacity for new faculty that results in over-provisioning of hardware resources.
        - Startup subscriptions could expire after a set number of years or be carried forward to help seed mid-career faculty. Based on the experience at Cornell, capacity in terms of overall core hours outstrips the overall liability in committed subscriptions before usage occurs, so that subscriptions do not have to expire.
    - **Extreme Change**: In this version, ICDS might explore a product segmentation model:
        - An Open Queue free tier, provided thanks to subsidy by Senior Vice Provost
        - Priority Execution, the default model for most usage

            - Flat hour calculation with weighting based on CPU/GPU capability. Time stays constant, but allocation is decremented faster based on more expensive hours incurred on faster CPUs/GPUs.
            - No different rate per product but a normalization function that establishes the difference between hardware - placing higher costs on more expensive equipment, similar to the normalized Service Unit used for national center allocations
        - Guaranteed Execution: pay additional for the 1-hour execution guarantee through an increased burn rate
        - There remains a risk of having a too-complex matrix of pricing (option for 1-hour, option for 6x burst, etc), keep it as simple as possible, sell the subscription in hours or standardized service units as described in the survey results above, and then increase the burn rate for use of higher-end node types or guaranteed execution times, e.g. it would be possible to tie GPU hour usage to the number of cores per GPU, so for a 56-core, 4-GPU system, an hour of GPU time incurs 14 core-hours.
- **Services**
    - The team should clearly communicate that SLAs are not required to be multi-year with a goal of eliminating SLAs or making the process self-service if possible (see subscription models above). ICDS will currently support shorter SLAs but it is not well advertised and many researchers are not aware of the fact.
    - Continue supporting an option for researchers to bring their own systems if their needs do not fit into existing products.
        - Have a fixed (or template) log-scaling structure for bring-your-own systems architected with ICDS, e.g. X hours per month for up to Y size, XX hrs for up to Z size. Examples should be clearly published so the model, and equity of individuals that utilize this option is clear.
        - This is the case where the SLA process is more appropriate, as there are multiple teams involved in engineering, operations, and support.
        - Clear discussions about hardware lifecycle and operating length of agreement are critical to ensure that researchers have rational expectations.
    - ICDS should collaborate with central IT on services that are already robust (e.g. leveraging VM services for persistent services and non-standard solutions). It could be possible for ICDS to provision in bulk for re-allocation (e.g. ICDS contributes to the infrastructure cost of the VM farm and receives a set allocation to re-sell). This is an arrangement institutions frequently make to provide secure enclave services or research web services.
    - Storage
        - The storage and archival rates of ICDS compare favorably to many institutions, particularly considering the high performance nature of project storage offered. At the time of the review the rates for archive are reasonable at $14.40 per year per TB, $80 per year per TB for VAST (flash). See the benchmark analysis for further comparisons.

        - Other institutions informally surveyed at CASC have seen significant uptake of storage centralization through rate reduction. Two other public institutions noted that long-term project storage costs of $30-50/TB/yr greatly increased centralization.
    - ICDS should consider offering long-term project storage that is less performant than current project storage. This offering would be similar to second/mid-tier storage of other institutions in the benchmark analysis. If this does not make sense from an architecture and services expansion perspective ICDS should collaborate with central IT to provide such a service. This would meet a significant research need, and in many cases obviate the need for many departmental NAS devices if it provides NFS/SMB services. These solutions can typically be provided for under $50TB/yr with some level of local data protection.
- **Personnel**
    - The split of the operations and engineering teams does not align with the organizational design of most centers, nor does it seem to be ideal for ICDS. Moving the operations team under the technical director would align personnel with implementation/engineering and allow greater cross-training, personnel mobility, and potentially a modest reduction in staffing over time.
    - ICDS needs to provide training and information on the escalation process for Human Resources related issues and identify the contacts inside and external to the institute who can escalate issues. While PSU appears to publish required Title IX details, the reorganization of HR personnel has created a gap for the institute without a clear path of escalation. The vast majority of individuals interviewed do not know who to contact if they have a personnel issue that cannot be, or is not addressed, by the chain of command. This could be accomplished through an Institute-level code of conduct with a set of trusted individuals to contact for escalation.
    - The dual mission of ICDS in fostering interdisciplinary research while also providing computing services to the entire university makes it difficult for the institute director to have the depth and breadth of knowledge to offer high quality research computing services. ICDS should consider creating a chief operating officer role. Such a role would provide singular focus on developing and implementing a comprehensive strategy for research computing, data, and consulting services. This type of role should have significant leadership responsibility within the institute as well as researcher-facing exposure outside the institute.
- **Business team**
    - ICDS Administration should streamline tools where possible, particularly aligning with central PSU systems where available, e.g. more reliance on ServiceNow and iLab. ICDS might consider the retirement of INFOR CRM and leverage ServiceNow CRM capabilities for managing SLAs and other contact details to be in the same environment as the ticketing system.

    - Provide self-service options where possible (this would be much easier for both ICDS to manage, and for faculty members to navigate with a subscription-based model or a model with fewer products).
    - The growth of ICDS over the past 7 years has outstripped the knowledge transfer and formal processes for the organization; it is critical that ICDS focus on building the operational capacity for sustainable support as opposed to its current model. One such example of a capacity-building effort would be the implementation of centralized user and project information; clear metrics for support ticket delegation and closure; and a change management process that involves researcher-facing and communications teams.
- **Communication**
    - Regarding advisory committees, ICDS should include HPC working group leads on campus within the to-be-formed ICDS CI faculty advisory committee.
    - ICDS needs to be more transparent to the PSU community in sharing how rates are set, providing regular public utilization reports, and conducting regular external surveys.
        - Specifically ICDS should clearly communicate the intent and use of the SVPR subsidy to the community in this process. Most other institutions benchmarked provided such detail.
        - Likewise, the administration team should create a cost-comparison sheet to external services keeping in mind the performance level.
    - ICDS should conduct more informal outreach to other campuses within the Commonwealth system, this should be a distributed effort, it should not require a high level event or invitation from the chancellor.
    - ICDS should develop case studies with faculty, this could be made as a requirement of co-hire arrangements. These case studies would also help highlight that these co-hires are paying same rates for services once startup costs have been expended.
    - ICDS should conduct a needs assessment for future investments regularly, as well as more timely assessment of service-level activities (per-interaction or annually).
        - This could be supplemented by ongoing interviews with SLA owners (targeting ~20% per year for in-depth discussions).
        - ICDS should regularly publish a plan of the ROAR-Next roadmap and storage refresh cycle; potentially coupling it with an annual town hall or other sessions for discussion of other surveys, communication, and/or performance metrics.
- **Metrics & Reporting**
    - ICDS should report contract and grant income from ROAR/ROAR Collab/Storage users (both for paid and unpaid members) to help demonstrate the overall impact on PSU. These data should highlight co-hires in a distinct category, to show the contributions to interdisciplinary research.

- Utilization reports and graphs need to be made available to the ROAR and ROAR Collab users at a minimum, these should include overall system utilization and job wait time. Ideally job success/failure metrics should be added over time.
- Regular reporting of open queue usage for both research and education. Both the SVPR and the PSU community as a whole need to understand the value delivered through the open queue to the institution and to the PI, it represents a dollar equivalent value provided to the community at no direct cost.

Conclusion

ICDS provides considerable support to the Penn State research community and has the potential to continue to offer computational and storage services in a sustainable fashion moving forward. ICDS staff and leadership have both interest and resolve to improve service offerings and support for researchers; the recommendations put forth in this report reflect input from faculty, staff and colleagues as well as the review team's experiences with research computing and data organizations.

ICDS leadership has the best sense of the cultural fit and feasibility of the review team's recommendations, moreover, ICDS leadership has the ability to establish goals that will support changes towards sustainability and greater buy-in from the Penn State community. The choice of actions from the list above will depend on strategic discussions that should involve members of ICDS, the office of the SVPR, and faculty stakeholders. ICDS should seek out regular input from institutes that have more mature processes, such as the Huck institute, and that input should be carefully weighted in the resulting approach. The changes that will support ICDS' further development will be built upon greater formalization and standardization of operations, as well as transparency of practices and services to Penn State Faculty. This transformation will not be an immediate process, but will require cultural shifts within the Institute that take considerable time to implement and reify over the long term.

Appendix 1: Review Charge Sent to Faculty and Staff

Penn State HPC community:
Prompted by the discussion around rate changes for our High Performance Computing (HPC) services at UP, we have initiated an external review to evaluate and help shape our HPC going forward. High performance computing is critical for growing our interdisciplinary research excellence across the university, a foundational element of President Neeli Bendapudi's Goal 2. It is vital we meet user needs in a cost-effective manner in this rapidly moving and critically important space.

The external review team consists of Richard Knepper, PhD, Director of the Cornell University Center for Advanced Computing, and David Hancock, Director of Advanced Cyberinfrastructure, Research Technologies of University Information Technology Services at Indiana University. Dr. Knepper specializes in cost recovery operations, organizational collaborations, and large-scale cyberinfrastructure management and policy; Mr. Hancock has considerable experience in research operations and high performance computing architecture, storage, and service management and provision. They will work with ICDS staff, faculty leaders and faculty users, and consult widely.

We have a wide range of users and administrators for our HPC (>2,000), so it will be impossible for our reviewers to contact everyone. If they reach out to you, please do engage with them positively and share your expertise and experiences. For those who are not contacted, we will also be providing a small survey, so you are able to provide your feedback. Gathering of everyone's thoughts will be an important part of our growth in this area. Feedback from the survey will be provided to the reviewers.

The review is intended to be complete by March 31, 2024. We will share the findings widely and look forward to shaping our investments going forward.

Sincerely,
Andrew Read, Interim Senior Vice President for Research
Jenni Evans, Director, Institute for Computational and Data Sciences.

Wolf Hey
ICDS Director of Technology-Based Research Projects and Partnerships
Interim Communications Manager

The Pennsylvania State University
214 Computer Building
University Park, PA 16802
+1-814-867-1467
wolfhey@psu.edu

Appendix 2: List of Respondents

Amit Amritkar
Chad Bahrmann
Carrie Brown
Phil Canakis
Keith Cheng
Guido Cervone
Jim Coder
Doug Cowen
Doug Dodson
Eric Donnel
Kathy Drager
Deb Ehrenthal
Jenni Evans
Wayne Figurelle
Adam Focht
Eric Ford
Steven Greybush
Cindy Yuexing Gulis
Chad Hanna
Wolf Hey
Vasant Honavar
Jason Hughes
David Hunter
Mahmut Kandemir
Gretta Kellogg
Soundar Kumara
Galen Lentz
Derek Leydig
Emily Martell
Jessica Menold
Scarlett Miller
Scott Milner
Robert Nicholas
Patrick North
Edward O'Brien
Chuck Pavloski
Justin Petucci
Lesley Shaffer
Gary Skousen
Mike Stedelin
Adri Van Duin
Lindsay Wells

Appendix 3: Survey Instrument

1. Are you aware of the Institute for Computational and Data Services (ICDS) and the associated ROAR service?
2. Regarding your use of computational resources for your research are you (select all that apply)
3. "I feel fully informed about the ICDS resources and their costs." Do you: (Likert scale on agreement)
4. Computing costs for the ROAR system are: (Likert scale: Too Expensive, Somewhat expensive, Just right, Somewhat inexpensive, Too inexpensive (not reflective of true costs))
5. Storage costs at ICDS are: (Likert scale as previous)
6. “I am able to get the kind of computational resources needed via the ROAR system” Do you: (Likert scale on agreement)
    6.1. Why or why not?
7. “Support via the i-ASK helpdesk for my issues on ROAR is effective at solving my problems” Do you: (Likert scale on agreement)
    7.1. Why or why not?
8. ROAR Maintenance outages are: (Likert scale on frequency)
9. If the cost of ROAR were reduced I would: (Likert scale on usage)
10. I want to see my jobs on ROAR start: (Likert scale: Within an hour, Within a day, Within a few days, Not important as long as all my jobs run)
11. Please describe your appraisal of the ROAR system, including any of the areas discussed in the previous questions or areas that are important to you.
12. Do you agree to be contacted by ICDS for a brief discussion?