

HTC26

Small-Campus Perspectives on the Research Computing Landscape



Jason Simms, Ph.D., M.P.H.

Research Computing Manager, Swarthmore College

THE ARGUMENT

The scale research now requires, or access to newer hardware, may require off campus resources. The OSPool is one of the best on-ramps to national capacity for smaller institutions, but adoption (both of it and similar tools) is limited far less by capability than by **facilitation, workflow fit, and trust.**

01

Why local is no longer enough

Shifting funding and the economics of on-prem compute at small institutions.

02

The realities that shape adoption

Teaching, tempo, data, people, and the analysis bottleneck.

03

The OSPool as a first-order resource

What makes it fit, and what barriers prevent broader use.

These students are the future workforce

~70%

of U.S. undergraduates
attend non-R1 institutions

1

The pipeline to graduate study

Many students from non-R1 institutions go on to R1 graduate programs. We have to prepare them for that environment.

2

The workforce conversation

Governments at all scales cite a need for “workforce development” - here they are!

3

Community college students, especially

Whether entering a vocation or transferring, they need real research methods and hands-on experience.

Single-campus compute grants are drying up

OLD REALITY

CC* / MRI funding a local cluster on your own campus



NEW REALITY

Fewer, larger grants for shared state / regional / national resources

What's getting harder

- Capital grants to buy a cluster for one campus
- Funding spread across many small awards
- Justifying local compute hardware to funders at all

What still (perhaps) gets funded

- Campus networking & connectivity upgrades
- On-ramps that let you reach shared resources
- The implication: plan to consume capacity, not own it

On-prem has become hard(er) to justify

Costs & physical limits

Steep price increases, 6-12 month (or more!) lead times, plus mounting demands on power, cooling, and space.

The right-sizing trap

Too small to be useful, or too large and sitting idle, it's hard to get right.

No charge-back model

Funding often must come from capital budgets, startup funds, or grants, and not recovered usage.

2x–12x

And the cloud is no panacea. For equivalent sustained workloads, cloud runs roughly two to twelve times the cost of on-prem.

Teaching and research are...the same thing?

The line barely exists

Faculty routinely bring active research into the classroom; course projects increasingly need enhanced compute.

Undergrads do the grad-student work

At an R1 a grad assistant fills this role. Here, undergraduates do, but they have less dedicated time to do so.


No buying out of teaching

Faculty can't use grants to buy their way out of the classroom, leaving roughly a ~10-week summer window for intensive projects.

Implication: *Platforms must be dependable and low-friction enough to use in a course and to pick back up after months away - and they want consistency!*

A bursty cycle can result in distrust of non-local services

THE ANNUAL RESEARCH RHYTHM



Long idle or low-intensity stretches, then intense summer bursts; faculty return after months and need to be productive immediately.

WHAT ERODES TRUST IN A REMOTE SERVICE

Services that vanish when a grant ends

Yet another separate account, login, and environment

Allocation requests and renewals to submit

Anything that taxes a tiny research window

Data and networking add their own friction

Datasets are outgrowing us

Storage is often scarce, expensive, and insufficiently performant.

A data-management skills gap

Tools and expertise are thin; staff and researchers lack time and, honestly, motivation to manage data well.

Commodity networking

Limited bandwidth sometimes restricts how effectively shared or cloud resources can be used, especially in pedagogical contexts.

Limited or no controlled-data capability

Despite demand, a secure enclave with dedicated compute is simply out of reach.

“Team science” strains a small shop

Collaboration outruns capacity

Team science pushes toward extramural collaborators — but federation and resource-sharing tools and skills are usually missing.

Required workflows (software, environments, dependencies) can be challenging to stand up and manage locally.

No two researchers alike

By design there's little overlap in expertise, so every project is a one-off: different workflows, packages, environments, dependencies.

A small researcher count kills economies of scale — equipment is very expensive on a per-researcher basis.

New faculty arrive with R1 expectations

We hire from larger institutions — which makes sense. But recent graduates are used to a richer resource ecosystem and want to keep doing research at that scale and style.

Inherited expectations

They often expect a similar resource pool and support cadence they just left behind.

Lingering collaborations

Many keep working with a doctoral advisor, former lab, or other external collaborators.

A hard onboarding


Getting them productive in a leaner environment takes real effort.

More compute doesn't mean faster science

TYPICAL WORKFLOW



REDUCED COMPUTE TIME, BUT INTERPRETATION STILL TAKES JUST AS LONG



Time →

So why scale up? Extra compute can churn a semester's data in weeks, but the time to make sense of the output is unchanged. The real benefit would come from better tools for the interpretation and analysis phase.

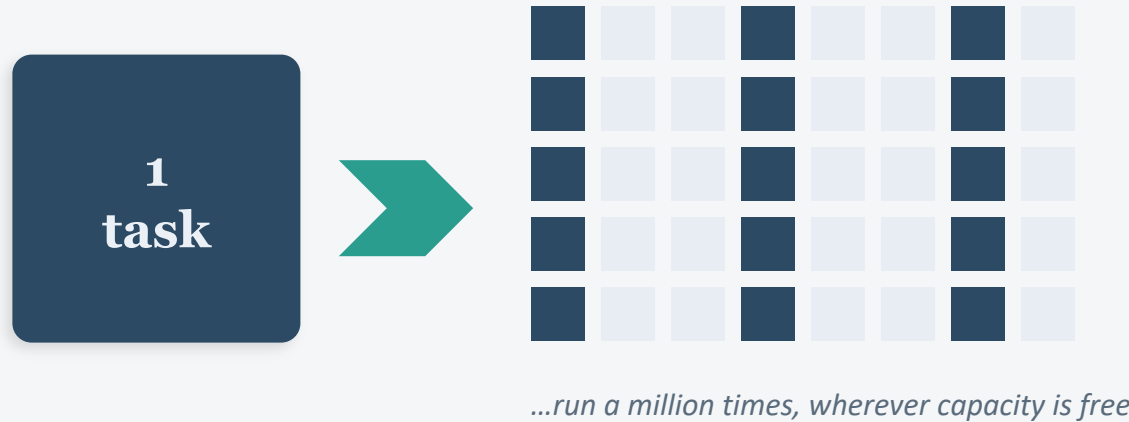
THE CENTRAL TENSION

The best support model for smaller institutions is local. The scale that productive research now requires increasingly is not.

Projects fail, or fall short of their potential, not for lack of compute, but for lack of local facilitation.

Researchers expect bespoke service from known colleagues. But <1 FTE usually covers research computing, and that time often by default skews towards hardware. **We need roles built to facilitate engagement with national resources as a first-order approach.**

Reframe the pitch: one task, a million times



WORKFLOWS THAT FIT

- Parameter sweeps
- Monte Carlo methods
- Bioinformatics read-mapping
- Image / MRI / GIS analysis
- ML training across params or data subsets

The historic myth: “only for huge batches of independent tasks.” True, but it can also run a single task. The message should shift from “millions of tasks” to “one task, a million times.” A subtle, but perhaps meaningful, difference for smaller institutions.

The lowest-barrier national on-ramp

No allocation request

Unlike ACCESS-style resources, there's no proposal to write, shortening time-to-productivity (even if only in perception).

Free at the point of use

No capital outlay, no charge-backs, no cloud bill to absorb.

No hardware to own

Nothing to house, power, cool, or refresh: capacity scales with demand, not budget.

The best first national resource

The gentlest learning curve makes it a natural first step into national cyberinfrastructure.

Barriers that actually limit adoption

Containers are a wall

Perhaps the biggest OSPool-specific blocker. Researchers and facilitators need far simpler guidance for building and using containers.

Proprietary software gaps

Stata, Gaussian, MATLAB, Mathematica, etc., which can be the primary disciplinary tool, aren't generally available.

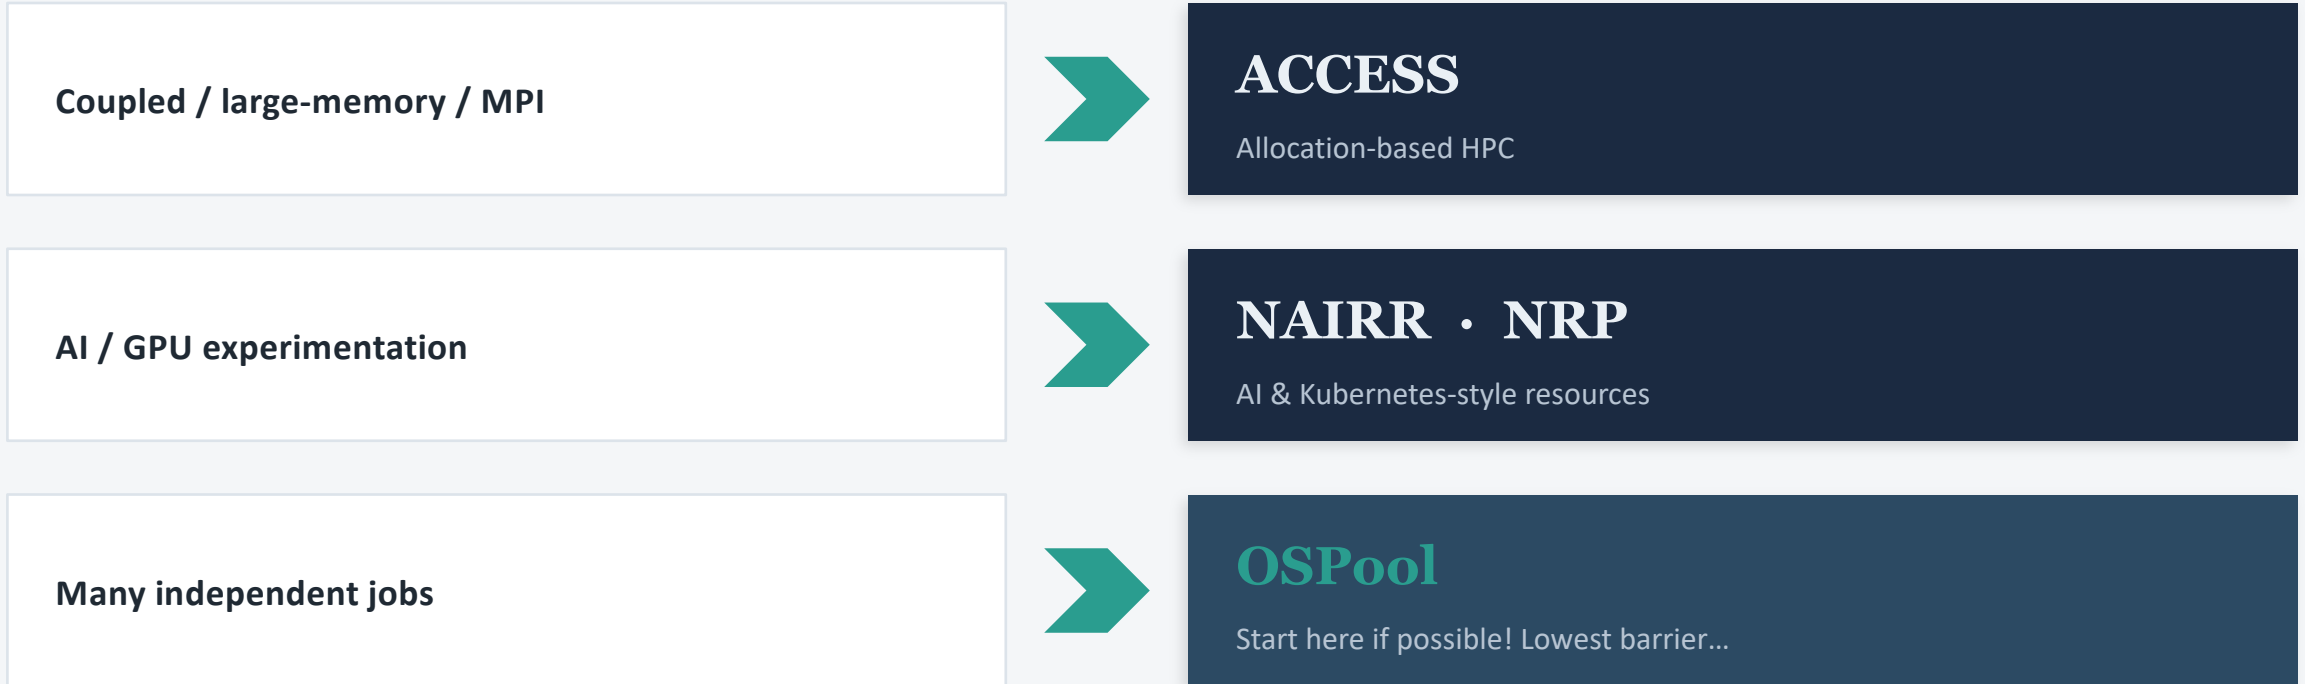
Not for every job

Node-spanning MPI and tightly-coupled, large-memory work aren't good candidates.

Researcher resistance & risk

Redesigning a workflow (for example, to implement checkpointing) risks productivity loss in a tiny research window, a trade many won't make.

Route the work by its shape



Stop trying to solve every need locally. Right-size the local cluster for commercial software, coupled jobs, and benchmarking to identify rate-limiting factors, and treat national resources as a first-order solution. Not an easy paradigm shift.

Facilitation is the real lever - not hardware

Reframe the role

With <1 FTE for research computing, the scarce resource is often human facilitation, not compute capacity.

Hire and define roles around facilitating engagement with national resources, as a primary job responsibility.

Be willing to do more

The usual rule is “don't do it for them.” With a tiny research window, more hands-on help may be exactly what gets a workflow off the ground.

Get creative: a paid student support role assigned to a researcher can carry the migration that staff can't.

Portability contributes to reproducibility

Hyperscalers, proprietary hardware, and the retreat from local compute can leave a project reproducible only in one specific environment.

The risk

Science that can't be re-run outside a single vendor or environment isn't fully reproducible, and that gap is widening.

The (partial) answer

Favor portable, container-based workflows on shared infrastructure where we can, though acknowledging how hard full portability remains.

Six moves for the community

- 1 Reassert the OSPool message** “One task, a million times.” And yes, even a single task.
- 2 Make containers easy/easier** Far simpler, friendlier guidance and assistance for building and using containers.
- 3 Fund facilitation, not just hardware** Define roles that facilitate collaboration and national engagement as a primary role.
- 4 Treat national as first-order** Right-size local systems; route work by shape across ACCESS, NAIRR/NRP, OSPool, etc.
- 5 Build better analysis tools** (Somehow) attack the real bottleneck: the time it takes to make sense of results.
- 6 Improve data staging and sharing** Critical to leveraging data in collaborative and pedagogical contexts that don’t tax local envs.

IN CLOSING

The compute is within reach. Closing the gap to broad adoption is the work (and it's mostly human!).

If we invest in facilitation, reframe the message, and meet researchers where their time and trust actually are, the OSPool and other resources and collaborative efforts can be a primary solution for smaller institutions.

Thank you · Questions welcome

Jason Simms · jsimms1@swarthmore.edu

