

ACCESS CPU + GPU allocation request

ACCESS allocation; tier selected from measured CPU/GPU request (Discover/Accelerate/Maximize as justified)

Draft for review — not submitted

Project and team

Title: UArizona Information Science CPU + GPU Teaching Allocation — [APPROVED TERM / YEAR]

Institution: University of Arizona · College of Information Science

Pi: Greg Chism (proposed PI; confirm ACCESS profile, affiliation and eligibility before submission)

Co Pis: [APPROVED COURSE LEAD(S), IF CO-PIs]

Dates: [PROJECT START: SIX WEEKS BEFORE COURSE] through [PROJECT END: TWO WEEKS AFTER FINAL COURSE DATE]

Students: 80 planned students: 40 data science and 40 MLIS; confirm actual enrollments

Staff: Up to six instructional/technical staff accounts; confirm role mapping and whether these count against provider limits

Public overview

The University of Arizona College of Information Science seeks an ACCESS CPU and GPU allocation to support reproducible data science and digital-curation teaching in [TERM / YEAR]. The initial proposed pilot includes INFO 521 and LIS 671, approximately 80 students, with later expansion contingent on measured demand and approval. Students will analyze data, evaluate appropriately sized models, normalize collection metadata, query catalogs and export documented outputs. Before piloting, we will test CPU components through campus UITS/AWS and benchmark GPU workloads on free resources. The classroom pilot will use approved ACCESS resources, with named local support, accessible CPU-equivalent activities, independent recovery and end-of-term exports. GPU amounts and provider fit will be justified by measurements rather than free-tier quotas.

CPU + GPU scope and other resources

INITIAL PILOT CPU BASELINE: 80 students × 8 active workspace-hours/week × 15 weeks = 9,600 workspace-hours / 19,200 CPU core-hours at 2 vCPU per workspace; 25% allowance gives 12,000 workspace-hours / 24,000 CPU core-hours before staff, idle billing and orchestration. Preferred 6 GB RAM/student; initial homes 5 GB/student + 20 GB/course shared data = 440 GB before staff/overhead. These are workload descriptions, not universal ACCESS units.

GPU REQUEST: quantity pending pre-pilot free GPU benchmarks and target-resource verification. For each approved GPU workload, calculate measured charged run-hours × planned runs/student or group × participants/groups × repetitions, then add separately justified rerun/checkpoint allowance, staff tests and measured overhead. Record GPU model, VRAM, precision, software and any device fraction; session time is not automatically full-device GPU time. Convert with the selected provider's current accounting/exchange rate. Do not request zero GPU by default: CPU and GPU both belong in the intended allocation.

LARGE ALLOCATION: confirm actual courses/terms/cohorts beyond the 80-seat pilot before extrapolating. Sum distinct CPU, GPU and storage demands across approved cohorts and subtract workloads already covered elsewhere. Do not double count campus AWS component tests as classroom CPU demand. Existing Jetstream2 scenarios (36,600 / 516,600 / 617,400 CPU SUs and an illustrative 1,920 fractional GPU-resource SUs) are comparison models, not this final request. The talk's 480-seat calculator is also illustrative. Tier, numeric credits/resource units and concurrency/storage quotas remain incomplete until measured and provider-reviewed.

Other available resources

Campus UITS/AWS: Greg confirms access via requests, restricted in this plan to CPU component testing; exact limits/charges/dates need confirmation. Free Kaggle and Colab GPU resources: staff pre-pilot benchmarks only, subject to actual account eligibility, quotas and hardware availability. ACCESS provides the intended CPU and GPU classroom allocation. CloudBank is conditional through ACCESS. CyVerse/MESA, NRP and OSPool remain optional separate components; MESA funding is not automatically an instructional allocation or operations budget.

Before submission and piloting

Use ACCESS for the large CPU and GPU allocation and the classroom pilot. Use campus UITS/AWS only to test CPU components, and free GPU resources to benchmark workloads before piloting. Size the ACCESS request from measured demand; consider CloudBank through ACCESS only if its managed service can meet the confirmed requirements.

Operating constraint: Greg has four hours/week total and zero cash budget. TA/lead hours are separate; course leads need no salary. Any necessary paid needs depend on confirmed eligible EduCloud grant funding.

Required completion

- Actual course/term/enrollment and any additional cohorts for the large request — incomplete.
- Eligible designated PI/co-PI profiles and current required CV PDFs — incomplete.
- Tier-appropriate education/resource plan and required course syllabi; CloudBank specifically requires syllabus PDFs if selected — incomplete.
- Dated CPU component tests and free GPU benchmarks with target-resource accounting/scaling justification — not run.
- Provider-confirmed CPU/GPU units, storage/concurrency and educational project structure — missing.
- Named operators, confirmed term coverage and any eligible grant support, support/data/closure plan — pending.

Review and submission steps

1. Complete CPU component tests and free GPU benchmarks before any student pilot; retain failed controls and reproducible artifacts.
2. Confirm the classroom/education scope and whether multiple courses/terms require separate projects. Keep unrelated MESA research work separate unless explicitly approved.
3. Choose the ACCESS tier from measured demand: Discover explicitly supports large classes; Accelerate or Maximize requires the relevant justification. "Large" alone does not select Maximize.
4. Complete the appropriate resource plan, course attachments and PI/co-PI CV PDFs using current portal guidance; this PDF is a supporting draft, not the final tier-specific submission.
5. Submit for separate CPU/GPU resource exchanges and needed quota/storage. If CloudBank proves feasible, use its documented ACCESS Classroom route and actual accounting instead.
6. Before piloting, reproduce benchmark workloads on the awarded resources and pass applicable classroom readiness checks. No award or reserved capacity is presumed.

Supporting draft, not a final tier-specific resource document, syllabus, CV or completed portal submission. All benchmark/readiness checks remain planned.

[ACCESS request preparation](#)

[Full request and conditional CloudBank branch](#)