

Reliable computing for teaching

What we have previously relied on can no longer be guaranteed

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.

Planning baseline: 40 data science students + 40 MLIS students; six preparation weeks, 15 teaching weeks and two export/close weeks. Actual courses, term, instructors, PI, resource limits and operating commitments are pending.

Decision requested

Greg has four hours/week total for this work and zero cash budget. TA and course-lead effort is separate; course leads need no salary. Any necessary paid support or resource needs await a confirmed eligible EduCloud grant. Start with components; reduce scope when time or support is insufficient.

One component before a semester

Before piloting · preparation weeks 1–2 - Test components and benchmark GPUs

Exit: Reproducible CPU/GPU component results and a workload envelope, within available time. Local functional checks do not replace AWS/CUDA provider measurements.

After benchmarks · before canary - Submit a sized ACCESS CPU + GPU request

Exit: Complete PI/course/term attachments and approved CPU/GPU resource configuration; request submission and award remain human/provider actions.

After award · preparation weeks 3–4 - Run a ten-person component canary

Exit: Both workflows succeed; actual service limits and unresolved defects are recorded. Stop or revise if core fit fails.

After canary · preparation weeks 5–6 - Rehearse the whole class and recovery

Exit: All thirteen applicable checks have evidence, term support and any needed grant-funded costs confirmed; otherwise narrow or defer the classroom scope.

Weeks 1–15 - Teach the two-course pilot

Exit: Record aggregate completion/access outcomes and actual costs; pause a failing component without removing the CPU course route.

Term end + two weeks - Export, close and decide

Exit: Verified exports, measured restore, closed accounts, final report and a costed continuation decision.

Two CPU workflows; named support

INFO 521: Introduction to Machine Learning - 40 planned seats

Build a reproducible data-cleaning and exploratory-analysis notebook, then fit a small CPU model with held-out validation. Explain missingness, leakage, uncertainty and data provenance.

Inputs: Generate a synthetic civic-service dataset with no real residents; keep core CSV inputs below 250 MB and test peak memory against the awarded workspace. Real/public datasets require a rights and data-classification check.

Environment: Requested: Python, JupyterLab, pandas, scikit-learn, matplotlib and an approved reproducible package manifest; R/RStudio only if the course requires it.

Acceptance: Run the whole assignment from a fresh student workspace; require execution without out-of-memory errors at the approved spec, reproducible key outputs and a complete exported submission.

LIS 671: Digital Curation and Digital Preservation - 40 planned seats

Audit a collection catalog, normalize names/dates/identifiers, document a metadata crosswalk, query an embedded catalog database and export a preservation-ready package with README, schema, provenance and checksums.

Inputs: Use a synthetic 10,000–100,000-record collection catalog with Dublin Core-style fields. Add public/licensed records only after rights review; exclude patron, circulation, donor and restricted collection records.

Environment: Requested: JupyterLab/Python, CSV/JSON/XML parsing and embedded SQLite; optional MARC parsing packages after approval. OpenRefine or other hosted applications require a separate provider fit check.

Acceptance: Execute the complete workflow in a fresh workspace; validate schema, record counts, identifiers and query results; reopen the exported package using the fallback environment.

Staffing coverage

Greg: four hours/week across all allocation, benchmarks, operations, design/data review and grant work. Potential TA effort: 51 hours each; course leads: 35 additional hours each with no salary. Those two roles sit outside Greg's hours. Normal teaching/grading is excluded. Independent backup coverage is unfilled and must be a distinct person.

Publish staffed help windows and a proposed one-business-day first-response target; no 24/7 promise. Escalation: TA/course lead to primary operator, backup/provider and service owner. Primary and backup operators must be different people.

GPU, MESA/CyVerse, NRP and OSPool modules need separate gates. PATH/IDSS eligibility is unconfirmed. Optional resources cannot become required assessment dependencies without their own checks.

Fit the work inside available time

Function / owner	Prep h	Teach h/week	Close h	Total h
Allocation and coordination (Greg)	6	1	1.0	22.0
Benchmark and component work (Greg)	18	1	1.0	34.0
Service operations (Greg)	12	4	6	78
Instructional design and accessibility (Greg)	6	0.5	1.0	14.5
Data, rights and recovery review (Greg)	6	1	4	25
EduCloud grant preparation (Greg)	6	1	2	23
Planning and incident buffer (Greg)	6	1.5	5.0	33.5
Two course support TAs (Not appointed; grant-supported if funded)	4	3	2	102
Two course leads (Course leads; named commitments pending)	16	1	4	70

Prep, teaching and close columns are per person; Total h combines both TAs or both leads. Earlier Greg workload: 230 hours. Current four-hour weekly capacity gives at most 92 hours over 23 weeks: a 138-hour gap. Rescoping or confirmed distinct support is required.

Hours only. Greg has four hours/week total across this work and EduCloud. Over the illustrative six preparation + fifteen teaching + two close weeks, 92 hours is an upper bound, not reserved project time. The earlier role estimates still total 230 hours, a 138-hour gap before other commitments. They are unrevised demand estimates, not a feasible four-hour plan or committed staffing. Narrow the work or confirm distinct support before proceeding. No salary math, cash envelope or finance role.

Available cash budget is zero. Work needing paid support, storage, hosting, development or recovery must have a confirmed eligible EduCloud grant and allowable work package before spending. ACCESS is CPU/GPU resource allocation, not payroll. No grant award is asserted. Existing campus test access and free GPU resources must stay within their authorized no-purchase scope.

The earlier 230-hour demand exceeds the 92-hour upper bound by 138 hours across the modeled 23 weeks. Greg cannot be his own independent backup. The earlier separate backup estimate is 27 hours; no distinct person is committed. TA/course-lead hours are separate and unconfirmed. Institutional privacy/accessibility approval remains separate from Greg's coordination.

[Hours planner and eligible EduCloud grant work packages](#)

Verify before class.

1. Scope | Greg + course leads

Named PI/leads, Greg's four-hour envelope, separate TA hours, distinct backup and any necessary grant-funded term coverage.

2. Benchmark | Greg (operator) + course leads

Campus AWS CPU tests and free GPU benchmarks completed; workload manifest, failures and target ACCESS replay recorded.

3. Provider | Allocation lead + selected ACCESS provider

ACCESS CPU/GPU exchanges, actual quotas/storage, accounting and educational project dates confirmed; CloudBank conditional.

4. Accounts | Operator + course leads

Approved logins, isolation and account revocation demonstrated.

5. Software | Each course lead

Complete assignments run in the pinned environment; required GPU work fits the verified ACCESS device.

6. Capacity | Operator + provider

Provider-approved class load rehearsal; proposed p95 workspace access within five minutes.

7. Save | Course leads + TAs

Save/restart/export demonstrated and outputs reopen elsewhere.

8. Grades | Course leads

Versioned submissions and test grade receipts reconcile in the approved LMS.

9. Data | IT/privacy reviewer + Libraries

Data rights, access roles, destinations and retention approved.

10. Accessibility | Greg (design coordination) + course leads

Essential keyboard/screen-reader workflow and equivalent assessed CPU route demonstrated.

11. Recovery | Primary + backup operators

Independent restore measured against proposed 24-hour RPO and one-business-day RTO.

12. Support | Owner + TAs + operators

Staffed support, escalation and outage fallback rehearsed; no 24/7 promise.

13. Exit | Service owner + course leads + Libraries

Expiry warnings, learner/grade exports, closure/deletion and next-term decision documented.

Evidence, evaluation and exit

Keep private account/grade evidence in the institutional register. Review aggregate access, assignment success, exports, support hours, costs and restore timing. Reconcile grades and complete export, retention and closure before expiry.

[Full readiness register and evidence criteria](#)