

Higher Education HPC Procurement Checklist

A working checklist for research computing teams planning a shared or departmental HPC cluster

Governance and Funding Model

☐	Cluster governance model defined (single-department, shared/condo model, or campus-wide committee-led)
☐	If a condo or buy-in model, per-node or per-share cost structure and priority access rules documented
☐	Faculty or research committee oversight structure established for shared resource allocation decisions
☐	Funding sources identified across departments, grants, and central IT budget, with contribution shares agreed upon in advance

Grant Compliance

☐	Federal grant compliance requirements identified for any NSF, NIH, or other federally funded research relying on this cluster
☐	Data security and compliance requirements confirmed for any grant-funded research involving sensitive or controlled data
☐	Procurement timeline aligned with grant funding cycles and any spend-down deadlines
☐	Documentation process established for demonstrating appropriate use of grant funds in the hardware acquisition

Workload and Capacity Planning

☐	Primary research workloads identified across the departments or research groups the cluster will serve
☐	Workload mix assessed for GPU-intensive AI/ML research vs. traditional CPU-based HPC vs. mixed use
☐	Realistic growth horizon defined (3-5 years) accounting for expected increase in research computing demand
☐	Expansion path confirmed: can the cluster add nodes or capacity incrementally as more departments or grants buy in?

Vendor and Support Consolidation

☐	Number of separate vendor relationships required for the full stack (compute, storage, networking, scheduler, support) evaluated
☐	Single point of contact and escalation path confirmed with primary hardware vendor
☐	Support model matched to internal IT team size and capacity, many research computing teams operate with lean staff relative to system scale
☐	Total support burden estimated realistically, not just initial deployment effort

Technical Specification

☐	Interconnect requirements defined based on workload (standard networking vs. high-speed InfiniBand for tightly-coupled HPC workloads)
☐	Storage tier strategy defined for active research data vs. long-term archival
☐	Scheduler and resource allocation software compatibility confirmed for the shared/multi-department access model
☐	Power and cooling capacity assessed against planned cluster density, particularly for GPU-dense configurations

Lifecycle and Retirement Planning

☐	Expected hardware service life defined (commonly 5 years for shared research clusters)
☐	Retirement and refresh cycle planning built into the initial procurement, not deferred to a future budget cycle
☐	Data migration plan established for the eventual hardware retirement and cluster refresh

→ [Contact Ace Computers Higher Education and Research Computing Team](#)

→ [Contact Ace Computers](#)