From Vision to Practice: A Blueprint for State and Regional AI Computing Infrastructure and Educational Access

October 2026



CRA
Computing Research Association

Authors

Erik Deumens, University of Florida

Ian Fisk, Simons Foundation

Manish Parashar, University of Utah

Barr von Oehsen, Pittsburgh Supercomputing Center

Helen Wright, Computing Research Association

Suggested Citation

Deumens, E., Fisk, I., Prashar, M., von Oehsen, B., Wright, H. (2026). *From Vision to Practice: A Blueprint for State and Regional AI Computing Infrastructure and Educational Access.* Washington, D.C.: Computing Research Association (CRA). http://cra.org/industry/wp-content/uploads/sites/9/2026/10/From-Vision-to-Practice-A-Blueprint-for-State-and-Regional-AI-Computing-Infrastructure-and-Educational-Access.pdf

About the Computing Research Association (CRA)

The Computing Research Association (CRA) represents nearly 300 North American academic units, laboratories, centers, and companies engaged in computing research. Since its founding in 1972, CRA has brought together academia, industry, and government to strengthen the computing research community and its contributions to society.

CRA's mission is to catalyze computing research by leading the community, informing policymakers and the public, and promoting the development of an innovative and responsible computing research workforce. Through its programs and initiatives, CRA supports researchers across career stages and helps shape the future directions of the field.

CRA
Computing Research Association

About the Coalition for Academic Scientific Computation (CASC)

CASC | Coalition for Academic Scientific Computation

CASC envisions a robust, sustainable ecosystem supporting academic research computing and data services, enabled by a vibrant community of professionals.

The core mission of CASC is to: Advocate for public and private investment in research computing and data services; advise federal agencies on relevant funding programs; engage in policy discussions on research computing and data services; foster a community of leaders in this field; and provide a forum for sharing strategic ideas and best practices.

This material is based upon work supported by the U.S. National Science Foundation under Grant No. CNS-2611153. Any opinions, findings, conclusions, or recommendations expressed are those of the authors and do not necessarily reflect the views of the NSF.

Table of Contents

Section 1: Setting the Stage **5**

A National Opportunity, Unevenly Reached 5
A Framework Built on Two Pivots 5
Philanthropy as Catalyst, Not Charity 6
From One Blueprint to Several Models 6
Reading the Four Models Side by Side 7

Section 2: Vignettes **8**

Florida: HiPerGator and the University of Florida 8
Snapshot and Context 8
Governance and Funding 9
Infrastructure and Access 10
Workforce and Educational Development 11
Timeline 11
Lessons Learned 12

Pennsylvania: The Pittsburgh Supercomputing Center and Keystone AI + Quantum Factory 12
Snapshot and Context 12
Governance and Funding 13
Infrastructure and Access 14
Workforce and Educational Development 15
Timeline 15
Lessons Learned 16

Utah: The Redtail AI Factory and the University of Utah 17
Snapshot and Context 17
Governance and Funding 17
Infrastructure and Access 18
Workforce and Educational Development 18
Timeline 19
Lessons Learned 19

New York: Empire AI 20
Snapshot and Context 20
Governance and Funding 20
Infrastructure and Access 21
Workforce and Educational Development 22
Timeline 22
Lessons Learned 22

Section 3: Conclusion **23**

Acknowledgments **25**

Sources **25**

Section 1: Setting the Stage

A National Opportunity, Unevenly Reached

The United States is at a pivotal moment. Artificial intelligence continues to reshape industries, redefine the nature of work, and establish new standards for the creation and application of knowledge. Despite these advances, access to essential computational power and educational frameworks remains inconsistent across states and institutions. This is a disparity that threatens to limit participation in the AI-driven transformation and leave entire regions and communities behind.

Chris Malachowsky's 50-State AI Computing Initiative is his perspective as an NVIDIA cofounder and one of the driving forces behind the University of Florida's AI transformation, which frames closing that gap as a national opportunity. To fully harness AI's transformative capabilities and secure the future of work, the paper argues, AI literacy must become a priority in every discipline. Not to prepare people to compete against AI, but to teach them how to work with it well, so that human judgment, creativity, and context remain decisively in the loop.

The idea from Malachowsky envisions a future in which every university in every state, regardless of geography or size, has access to critical AI infrastructure and educational resources. The goal is not simply to install a powerful new machine at a handful of elite institutions, but to reimagine how knowledge is created, shared, and applied broadly enough that talent is cultivated everywhere.

A Framework Built on Two Pivots

Malachowsky's case for a national initiative rests on two deliberate pivots away from how large computing resources have traditionally been funded and justified:

- **Educational focus over research-first framing.** Rather than justifying high-performance computing purely as a research asset, Malachowsky's argument focuses on educating a broad set of beneficiaries – institutions, students, and the economy – with research benefits accruing indirectly as a wider base of AI-literate users expands the vetted data, collaboration, and discovery the system supports. The paper calls this a "virtuous cycle" since broader education expands the pool of educated AI users, which in turn broadens and improves the research the infrastructure enables.
- **Strategic ownership over federal or centralized asset management.** Rather than a federally owned, isolated research center, the infrastructure becomes a strategic asset owned and operated by the educational institution(s) themselves or consortia of educational and other institutions. This reorientation is directed at institutional, state, and regional priorities as much as at the traditional research mission. In other words, this strategy shifts the institution from steward of a federally funded facility to strategic director of its own.

Taken together, these pivots describe what Malachowsky calls a “hub-and-spoke architecture.” One or more resourced research institutions serve as central hubs, deliberately extending compute access and curriculum outward to their state's broader ecosystem. This interconnected model is meant to drive dense, statewide economic impact and help ensure no community is left behind.

Philanthropy as Catalyst, Not Charity

Malachowsky’s concept treats philanthropy as central to the model, not incidental to it. It describes the initiative as “an unparalleled opportunity for transformative philanthropy” that combines the immediate gratification of tangible impact with the enduring legacy of educational and economic renewal. Unlike a traditional gift to one program, a single or pooled investment radiates outward to every college and community in a state. Malachowsky’s initiative is equally direct that this cannot be a one-sided transaction: a gift succeeds only when the recipient institution also invests so that ownership, not dependency, is what the philanthropy ultimately produces.

The baseline financial model, drawn explicitly from Florida's experience, illustrates the intended split: private partners or benefactors invest to purchase and deliver the AI supercomputer; the university or college system invests roughly a quarter of that amount in facilities and match; and the state invests roughly four times the private contribution in faculty expansion and strategic use. However, as our four state vignettes in Section 2 show, no state has actually landed on that exact ratio. Each state has improvised its own mix of donor, institutional, state, and vendor capital.

A separate Coalition for Academic Scientific Computation (CASC) framework paper, published in June 2026, on regional research-computing collaboration called “Regional Collaboration as Strategic Infrastructure: A Framework for the Future of Research Computing” arrives at a similar list for what makes these partnerships durable, see Section 3.

From One Blueprint to Several Models

Malachowsky’s effort did not begin as a policy proposal. It emerged from what he describes as the organic, on-the-ground experience of the University of Florida, one of the earliest institutions to integrate AI at scale into its academic and research life following a 2020 gift of an AI supercomputer from Malachowsky and NVIDIA. UF's independent efforts (building infrastructure, revising curricula, forming partnerships, shifting institutional culture) revealed, after the fact, a blueprint that other states and institutions could follow without waiting on federal coordination, legislation, or partisan alignment.

In the roughly five years since UF's founding gift, several other states and regions have begun building their own versions of this model. While they have not converged on a single design, they are in alignment on the ultimate goal of developing local AI infrastructure. That is, funding mixes, governance structures, and even the basic question of who owns the infrastructure differ meaningfully from one state to the next. Section 2 of this report profiles four of those

efforts in full: Florida, Pennsylvania, Utah, and New York, drawn from interviews and source documentation gathered around CRA and NSF's 2026 Industry/Philanthropy Summit on Potential NAIRR State Hubs for Educational Transformation.

Reading the Four Models Side by Side

No two of these states have made the same choices about ownership, funding, or allocation, and each is at a different point on the path. The following comparison summarizes where each stands as of August 2026:

Model	Anchor & Governance	Funding Sources	Allocation Method	Stage (as of Aug. 2026)
Florida	Single anchor (UF); central UFIT Research Computing owns the infrastructure; statewide access via Board of Governors agreement	$60M Malachowsky/NVIDIA gift + $15M UF match; $20M state legislative appropriation for faculty; user-fee cost recovery	Priority Access Model; teaching free, research at subsidized rates	Mature: 4th hardware generation (2025 Blackwell refresh); ~8,000 active users
Pennsylvania	Multi-university consortium (7 R1s) under Keystone AI + Quantum Factory; PSC (Pitt/CMU) anchors; Team PA is neutral convener/LLC	University contributions to a shared coordinator pot; state (Commonwealth) contributions; industry and philanthropy still being negotiated	"Condo model" in design — pay for capacity used, unused cycles pooled; not yet operational	Building: launch summit Oct. 2025, all players have signed Letter of Understanding (LOU) and contributed funds, Keystone AI + Quantum Factory is now an LLC
Utah	Single anchor (University of Utah / SCI/CHPC); vendor-owned infrastructure under an HPE GreenLake agreement, colocated off-campus	State + higher-ed contributions; private philanthropic donor (cancer/health AI); possible fee-for-service; NVIDIA workforce MOU	Mixed: startup allocations, proposal-based allocations, and credit-based on-demand access	Very early: system online Summer 2026; enablement Challenge launched July 2026

Model	Anchor & Governance	Funding Sources	Allocation Method	Stage (as of Aug. 2026)
New York	Legislated 10-member consortium (incl. Columbia, Cornell, CUNY, NYU, SUNY, Flatiron); board with institutional + external members; anchored at Univ. at Buffalo	New York State investment; founding Simons Foundation gift; 10-year institutional operating commitments from every partner	Equal per-institution "service unit" allocation, independent of use case or funding source	Building: Alpha (2024) through Beta (June 2026); permanent facility due late 2027/early 2028

Section 2: Vignettes

The table above summarizes the current status of each state, while the vignettes that follow provide context on how those positions were reached. Each vignette follows a common structure: snapshot and context, governance and funding, infrastructure and access, workforce and educational development, timeline, and lessons learned. This consistency enables direct comparison across states and with the framework presented in Section 1.

Florida: HiPerGator and the University of Florida

Snapshot and Context

The University of Florida operates HiPerGator through UFIT Research Computing, with access extended to the entire State University System of Florida, other Florida institutions, state agencies, and SEC universities. The anchor investment was a $60M donation of an AI supercomputer, an NVIDIA A100 SuperPOD, from UF alumnus and NVIDIA cofounder Chris Malachowsky and NVIDIA in 2020, matched by $15M from UF for data-center power and HVAC. A $20M/yr recurring state legislative appropriation in 2021 funded 100 new AI faculty positions across all 16 colleges. Ongoing costs are funded through a blend of the Provost's Office, the Office of Research, central IT, and faculty grant charges.

The initiative started with teaching: the provost's vision for an AI university aimed at all UF students, not narrowly at training AI engineers for tech companies, but at preparing graduates across every discipline (law, English, history, fine arts) to use AI appropriately as citizens and domain experts. The first cohort under the 2020 plan graduated in May 2026. From there, the effort grew outward from students, to faculty research, to university operations (digital twins of buildings, facilities management, an AI grants-search tool for the Office of Research), to the whole state. HiPerGator now functions as the de facto supercomputer of the State of Florida.

The 2020 gift offer was the immediate catalyst, paired with UF's unprecedented commitment to teach AI Across the Curriculum and framed against the national AI-workforce gap identified by the National Security Commission on AI. Readiness made the gift possible in the first place: UF had operated HiPerGator since 2013, and its CIO had insisted in 2010–2012 on building a data center larger than then-current needs. By 2020, half the machine room stood empty, existing capacity that made the donated system straightforward to site.

The educational goal is that every student in every major can become AI-literate, AI-competent, or AI-expert; AI ethics coursework is required, and individual colleges own their own disciplinary AI curricula. Key stakeholders span the Provost (vision owner), the CIO and UFIT Research Computing (operations), the VP for Research, the Board of Trustees, the state Board of Governors, the Florida Legislature, NVIDIA and Chris Malachowsky personally, and the deans of all 16 colleges.

Governance and Funding

Ownership is deliberately centralized: the donated machine went to central UFIT Research Computing, not to any college, and support staff are centralized rather than distributed across colleges. A Research Computing Advisory Group advises but makes no allocation decisions, and there is no multi-institution allocation committee. Statewide access instead rests on an agreement with the Florida Board of Governors: any Florida public university's faculty can use HiPerGator under the same conditions as UF faculty.

Financial sustainability has been proven over two replacement cycles (HiPerGator 3 and 4): user charges fund replacement of storage, CPUs, and lower-end GPUs, with the most recent CPU/storage refresh ($9M) paid entirely from five years of cost recovery. The flagship GPU system is treated separately: when the donated A100 SuperPOD aged out, the Board of Trustees provided $24M (offset by resale of the old machine) to buy the Blackwell B200 system, deployed in 2025. Staff and operations (roughly $4–5M/year) are deliberately not cost-recovered; the university funds them as overhead. The arrangement functions as an investment rather than a cost center: of UF's roughly $1B research portfolio, $500–600M is supported by HiPerGator, roughly a hundredfold return on the university's annual multi-million-dollar operating contribution.

Resource allocation follows a simple principle: nothing is free, so there is no allocation committee. Teaching allocations are free to instructors, paid by the provost and negotiated per semester; research and graduate education pay subsidized rates. Paying removes the classic governance problem: no panel decides whose science is worthy of access, as with NSF ACCESS, and no single group can monopolize the machine, because the market self-regulates. For external users the value proposition is speed: rates are comparable to or modestly above home-institution costs, but paid jobs start immediately instead of waiting weeks in a queue, and early adopters have stayed with the system since.

NVIDIA, while not directly involved in shaping the initiative, remained a willing resource UF could call on, and watched the program's progress with genuine interest. UF staff made pre-production hardware production-ready, including resolving firmware and fan failures on

early Blackwell nodes with NVIDIA engineers. Motherboards have also been redesigned based on UF findings. UF serves as a reference customer, effectively part of NVIDIA's sales motion to other universities, and receives early hardware access and software at no cost in exchange. For the donor, the arrangement is regarded as proof of impact and ownership, and Malachowsky now visits governors nationwide to promote the model. Rather than fund the replacement machine directly, he left that investment to UF's board, which is strong evidence that the model is sustainable rather than dependency-forming. The broader lesson is that a gift must be a two-sided partnership, where the recipient's obligations, matching facilities investment, staffing, and execution, are what make the donor's investment succeed. UF Press's Warrington Press imprint published a fuller academic account of this arc in May 2026, Michael Carrillo's *Building an AI University: From Infrastructure Gift to Systemic Transformation*, for readers who want the complete institutional case study.

Infrastructure and Access

As of January 2026, HiPerGator comprises 60,000 CPU cores and a 1,552-GPU ecosystem, including a 504-GPU NVIDIA B200 SuperPOD, a digital-twin hub, and dedicated inference capacity for campus AI services. It's backed by more than 70 PB of tiered storage, NDR-400 InfiniBand, and dual 400 Gbps links to Florida Lambda Rail. Compliance was designed in from the start: the system is HITRUST-certified and NIST 800-171 compliant, with HiPerGator-RV supporting ITAR/EAR/CUI at CMMC Level 2, enabling AI work on protected health information for precision-medicine research. A single central facility, with no lab clusters, is enforced by a 2015 VPR policy barring faculty from building their own clusters; facilities will not approve power or HVAC requests for them, and startup compute funds flow to the central system instead.

All Florida public university faculty use HiPerGator on UF's terms (teaching free, research at UF-subsidized rates), effectively meaning UF subsidizes the whole state. The scaling condition is that first-line support and training must come from each institution's own local research computing/IT staff; UF anchors a weekly call of research computing leads across Florida's universities to coordinate this. Users extend well beyond academia: startup companies and state agencies (including a Florida Department of Commerce AI project) sponsor projects, and earlier outreach included a $1M congressional earmark to train Palm Beach State College faculty and a statewide K-12 teacher AI program.

Scale, as of the most recent count, runs to roughly 8,000 active users under about 800 sponsors (faculty PIs across Florida's universities, startup CEOs, and state employees, roughly ten users per group) with thousands of enrollments per semester in AI-content courses at UF alone.

Workforce and Educational Development

UF's human infrastructure has grown into a layered, roughly 40-FTE staff: DevOps (9), support (7, plus a student-staffed help desk), training (2), regulated-research facilitators (5), and a 10-person AI consulting team (AI consultants, research software engineers, digital-twin engineers). Staffing grew in step with demonstrated demand: seven positions were requested up front alongside the 2020 gift, followed by five more, then three more, as usage climbed. It

was an incremental, trust-based approach rather than one large ask. The consulting workflow runs full-stack: a hallway conversation leads to an intake, then a statement of work, weekly check-ins, and a deployed application. One proof point: a research software engineer cut a projected three-year NIH 400-brain dataset analysis down to roughly five weeks.

On the student side, a university AI certificate (covering fundamentals, ethics/data/technology, and disciplinary application) feeds into college-owned curricula across all 16 colleges: a BS in data science (2020), an MS in applied data science (2022), an MS in AI systems (2023), and a Quality Enhancement Plan with defined learning outcomes. The first full cohort under this plan graduated in May 2026.

Timeline

- 2010–2012: UF's CIO builds an oversized data center against skepticism, the readiness that later captured the gift.
- 2013: First HiPerGator; the central research computing model begins.
- 2015: No-private-clusters policy adopted; NIST 800-53 secure research environment built.
- 2020: $60M Malachowsky/NVIDIA gift plus $15M UF data-center match; AI Across the Curriculum commitment announced; Board of Governors statewide-access agreement follows.
- 2021: HiPerGator 3 (A100 SuperPOD) goes live; GatorTron launched; statewide and SEC access extended; $20M from the legislature funds 100 AI faculty.
- 2022–2024: AI2 Center established; HITRUST certification achieved (87 findings closed in 6 weeks); AI-inspired Quality Enhancement Plan (QEP) adopted; Malachowsky Hall opens; Blue Sky Task Force convened; first version of NaviGator AI provides unified access to generative AI chat capabilities approved for work with UF data.
- 2025: A100 SuperPOD replaced with a Blackwell B200 system ($24M Board of Trustees funding plus resale proceeds); $9M CPU/storage refresh self-funded from cost recovery.
- 2026: First graduating class of the 2020 plan; continued GPU/storage expansion; NaviGator campus AI services keep expanding to cover agentic AI.

Lessons Learned

The biggest cultural/organizational hurdle was centralization: persuading faculty to give up lab clusters (enforced through the 2015 policy and facilities approvals) and colleges to accept one central machine and one central support team. On the regulatory side, HITRUST auditors found 87 gaps across 288 controls; UF closed them in six weeks only because of five years of prior NIST 800-53 experience, at a cost of roughly $170K/year for the annual review. For institutions starting from scratch, UF's experience suggests planning for roughly five years of runway. Technically and operationally, hardware without people fails: every capability increase was paired with staffing and training investment, and first-line support has to be local to scale statewide. Financially, staff costs cannot realistically be cost-recovered; sustainability requires

separating the renewal business model (user charges) from operations (institutional investment).

The honest answer on what UF would do differently is that the real mistakes were made in 2013–2020, before the AI initiative: allocation governance that didn't work, a distributed computing culture, and compliance gaps. By 2020, those fixes were already in place, so the gift landed on seven years of preparation.

A donated supercomputer is not a prerequisite, but readiness is critical. Data-center capacity, centralized research computing, a no-clusters policy, and compliance experience are what turn an opportunity into a program. The hub-and-spoke statewide model scales: one anchor institution subsidizes access, sets uniform terms (teaching free, research paid), and requires local first-line support at spoke institutions. The priority access model eliminates allocation politics and prevents monopolization: governance by market rather than committee. The roughly hundredfold return on HiPerGator's operating cost suggests the ask should be framed as investment, not cost. Finally, industry partnership must be bidirectional and technical (co-engineering, reference-customer value), and philanthropy must come with recipient ownership.

Pennsylvania: The Pittsburgh Supercomputing Center and Keystone AI + Quantum Factory

Snapshot and Context

The Pittsburgh Supercomputing Center (PSC) is a joint center of the University of Pittsburgh and Carnegie Mellon University, founded in 1986. PSC's Director since May 2023 also serves as Vice Chancellor for Research Computing at Pitt, with affiliate appointments in CMU's ECE and Biomedical Informatics departments. Funding is primarily NSF (Bridges-2, Neocortex, Bridges-3, ACCESS, NAIRR Pilot) and NIH (BIL, HuBMAP, SenNet), plus Commonwealth of Pennsylvania and industry participation through the Keystone AI + Quantum Factory. PSC also pays for an outside government-relations firm, Buchanan, to represent it in Harrisburg, since neither university has the capacity to focus specifically on PSC's state relationships.

The target audience began with the seven Pennsylvania R1 universities that laid the foundation, with plans to expand to non-R1 institutions; other stakeholders include the state, the federal government, industry (particularly manufacturing), and philanthropists.

What triggered the project was largely organic timing rather than a single planned effort, driven by aging data centers. On the state side, PSC's new director found the center received comparatively little state funding versus peer national centers, and began actively lobbying in Harrisburg via Buchanan to raise PSC's profile with the Commonwealth. Separately, several Pennsylvania universities (Penn State, Penn, Temple, and others) were independently approaching the state's Department of Community and Economic Development (DCED) for help funding data-center upgrades to meet AI power and cooling needs, and for community-college AI/data-science workforce funding (PSC had already run a successful $500K Adapt PA grant for this and wanted to expand it). A Buchanan-arranged meeting

connected PSC with Jen Gilburg at DCED, a former Intel employee who recognized the overlap in what she was hearing from multiple schools and asked PSC to draft a proposal for what it would cost to support all seven PA R1 universities statewide. That proposal, which came in at just under $1 billion once a shared data center, network, and staffing were included, became the basis for what turned into the Keystone AI and Quantum Factory.

The educational goal is to broaden participation in computational and data-intensive fields across levels, from undergraduates to experienced researchers. Confirmed stakeholders include Pitt, CMU, and the other five PA R1s (Drexel, Lehigh, Penn State, Temple, Penn); the PA Department of Community and Economic Development; Team Pennsylvania (the nonprofit now hosting the Keystone Factory as an LLC); NSF; industry partners including AWS, NVIDIA, and Dell; and Pennsylvania's congressional delegation, notably Senator Dave McCormick.

Governance and Funding

There is no central infrastructure yet. The Keystone Factory is still in its infancy, and no single entity owns it: as of mid-2026, the initiative exists as a coordinating structure rather than a built system. PSC itself remains a joint center between CMU and Pitt. Above PSC, the Keystone Factory has a Governance Committee made up of the seven R1 universities' leadership (VPRs/CIOs), with PSC's Director, Jen Gilburg (DCED), and Team PA sitting on it ex officio (Team PA's seat is temporary, until the Keystone Factory becomes fully independent). Ownership was deliberately routed through Team Pennsylvania rather than PSC, CMU, or Pitt, specifically to avoid the perception that the Keystone Factory is a Western Pennsylvania project. Real historical friction between Pittsburgh, Philadelphia, and Penn State (the land-grant institution) factored into that decision. Team PA set the Keystone Factory up as its own LLC within the past year; the Keystone Factory is now registered in SAM.gov and has become an independent 501(c)(3), able to pursue grants directly rather than through Team PA or a single university, and is also registered with NSF to submit proposals directly.

On financial sustainability, meaning what happens when the initial grant ends, each of the seven universities put money into a shared pot roughly a year ago to fund a coordinator and project-management function, covering part of that role's time and others'. That money is running out, and the group is now negotiating what comes next. Rather than a flat per-university assessment, they are leaning toward a condo model: each institution pays for the infrastructure capacity it actually needs (Lehigh and CMU, for example, have different needs), with unused cycles going into a shared pool, plus industry contributions layered on top to help fund capabilities for startups and manufacturing partners. This is expected to be a multi-way partnership among universities, the state, industry, and philanthropy, with no single funder controlling it; the state has also put money into the same pot to date. On the earlier roughly $1B proposal, it was meant to cover a new shared data center plus per-campus needs. Individual R1s were separately estimated to be looking at anywhere from $10M to $50M-plus each on their own campuses, with the data center itself running roughly half that total again. The Commonwealth balked at that scale, which is part of why the Keystone Factory shifted toward a phased, multi-source approach rather than one lump state appropriation.

Resource allocation today is handled primarily through NSF ACCESS and the NAIRR Pilot, peer-reviewed national allocation rather than local discretion, but that applies to PSC only, not the state initiative. For the Keystone Factory specifically, an Infrastructure Committee is mapping near-term, 1–5 year, and 10-year build-out phases, while a separate Funding Committee is working industry, federal, state, and philanthropic funding streams. Day-to-day allocation logic going forward is the condo model described above: pay for what you use, with unused cycles pooled.

Industry and philanthropy partners get clear value in return. For industry, the appeal is simultaneous access to talent, students, and faculty across seven R1s plus four PA medical schools (Pitt, Penn State/Hershey, Temple, Penn) and a partnership with UPMC, a scale of access no single-university relationship can offer, plus direct feedback on new technologies deployed through the factory that helps with product design, tested across a broad academic demographic. Pennsylvania itself is framed as a selling point: a small-scale version of the rest of the country, rural and urban, a net energy exporter, making it a realistic testbed for solutions meant to scale nationally. A stated goal is also to keep PA-born startups from relocating, the common pattern of university spinouts moving to California, by giving them a reason and the infrastructure to stay and hire locally. On IP, NSF's Industry-University Cooperative Research Center model and engineering research centers serve as templates, while doing this across seven R1s at once, rather than one university and one company, is acknowledged to be more complex. The approach has been to deliberately defer detailed IP and data-use agreements until they are actually needed, rather than let those negotiations stall the initiative. On philanthropy specifically, the framing is societal good, avoiding job losses to AI, and using the four medical schools as a base for life-sciences and healthcare-focused projects philanthropists could get behind.

Infrastructure and Access

PSC's current systems include Neocortex, a Cerebras Wafer Scale Engine system whose WSE carries roughly 900,000 processing elements, used for large-model work (Llama, BERT, transformer-class models) and one of six AI systems selected for the original NAIRR Pilot cohort; Bridges-2, a general-purpose HPC/AI workhorse and leading contributor to ACCESS and the NAIRR Pilot; and Bridges-3, featuring NVIDIA B200 GPU servers, high-core-count AMD CPU nodes, an all-flash Lustre file system, and XDR/NDR InfiniBand. It's heterogeneous by design so users can mix AI training, accelerated simulation, and large-memory CPU work in one workflow.

Bridges-3 is explicitly justified on serving institutions without local HPC, emerging research programs, and classroom use, and Learning Lab modules, training, workshops, and internships were developed under Bridges-2. On the Keystone Factory side, the approach has been a deliberate sequencing rather than a specific access program: Pennsylvania has close to 300 colleges and universities, and the group felt bringing everyone in at once would be too overwhelming. They chose to get the seven R1s pointed in the same direction first as the first major hurdle, and are now starting to loop in R2s and other non-R1 institutions.

On scale, there are no numbers yet to report, since the initiative is still in its infancy and does not yet have centralized infrastructure. What does exist is a small pilot and proof of concept in which several universities agreed to make their own computing resources available to researchers at other PA campuses (Pitt researchers using Penn's resources, Penn researchers using PSC's), intended to demonstrate that the universities can actually collaborate. The Keystone Factory also issued an RFI to cloud-service and quantum-computing providers describing the initiative and asking how they would want to structure a partnership; based on responses, a follow-up RFP has now been issued focused on the financial model for cloud and quantum access. In parallel, PSC is in active conversations with NVIDIA (about its 50-state strategy), Dell, AWS, and smaller PA-focused investors.

Workforce and Educational Development

Training, workshops, internships, and Learning Lab modules were refined through Bridges-2 and will be expanded under Bridges-3. The anchor program is ADAPT PA, PSC's statewide workforce initiative introducing AI and data science across Pennsylvania's higher-education landscape. What started as a $500K pilot focused on community colleges has grown to roughly 450 active learners across 19 institutions, spanning community colleges, R2s, and minority-serving institutions. ADAPT PA is now the workforce foundation the Keystone Factory Workforce Development committee is building on, and it addresses directly what the state and philanthropic funders have signaled they care most about: economic development, workforce pipelines, and broadening participation beyond the R1 tier.

Workforce Development is the fifth and most recent Keystone Factory governance committee, added after PSC found itself increasingly in conversation with Pennsylvania's Department of Education and Department of Labor and Industry, and concluded the Keystone Factory had been too research-focused relative to what the state and philanthropic funders actually prioritize.

Timeline

- 1986: PSC founded as a Pitt/CMU joint center.
- 2020–21: Neocortex deployed (Cerebras WSE).
- May 2023: PSC names a new Director, who shortly after begins active outreach to Harrisburg via Buchanan Ingersoll & Rooney, having found PSC under-resourced by the Commonwealth relative to peer national centers.
- January 2024: Neocortex selected for the NAIRR Pilot; Bridges-2 joins the NAIRR allocation pool.
- October 2, 2025: Keystone AI and Quantum Factory launch summit at Penn State. The seven R1s formally commit to the initiative and agree to contribute to a shared coordinator and project-management pot.
- Within the past year: Team Pennsylvania stands up the Keystone Factory as its own LLC, since no single university could sponsor it without raising conflict-of-interest concerns. Governance, Infrastructure, and Funding committees are stood up in that

order, followed by Government Relations; Workforce Development is added most recently (2026).

- April 21, 2026: Public press launch.
- 2026: NSF grants extend Neocortex to May 2027 and Bridges-2 to September 2027.
- July 2026: $10M NSF award for Bridges-3.
- 2026 (ongoing): RFI issued to cloud/quantum-computing providers; follow-up RFP issued on the financial model for cloud/quantum access; conversations underway with NVIDIA, Dell, AWS, and smaller PA-focused investors.
- Early 2027: Bridges-3 build begins at PSC's new data center; online summer 2027.
- Next step (as of August 2026): getting university presidents and the chancellor to direct their institutional foundations offices to actively engage existing donors on the Keystone Factory's behalf.

Lessons Learned

The barriers span every category. State government does not move at the pace or scale this required, since Pennsylvania's default is to spread funding evenly across the state, a pattern likened to spreading peanut butter evenly, rather than concentrating funding behind one high-impact initiative. Academic institutions are also slow-moving by nature, and standing up the LLC pushed the universities well outside their comfort zone. Federal funding patterns have changed fundamentally, and most universities have not adjusted; the Keystone Factory needs to be agile enough to second people quickly and respond fast when a competitive grant opportunity appears. University internal bureaucracy, including legal and development offices and a layered administrative hierarchy, puts the wrong people in the room for fast decisions. State politics and timing add a fifth barrier: Pennsylvania's governor is currently running for re-election, and AI and data centers are not a politically popular thing to champion right now, so major state funding is not expected until after the election. Tennessee's $50M commitment stands in contrast; Pennsylvania has not made a similar move. By contrast, Senator Dave McCormick, who is not in an election cycle, has been very responsive and has opened doors to new communities and events, including a Defense and Innovation Summit. Pennsylvania's status as a swing state, and the fact that both McCormick and the governor have national political ambitions, factors into how the Keystone Factory positions itself politically.

On what is transferable to other states and regions, the core lesson drawn from PSC's experience is to get the major research universities aligned first before trying to bring in the full landscape of institutions at once. Pennsylvania has close to 300 colleges and universities; the Keystone Factory deliberately started with just the seven R1s to keep the effort manageable, with plans to loop in R2s and non-R1s next. Using a neutral third-party convener (Team PA, not PSC or any one university) to hold the initiative was also central to defusing regional rivalries between Western and Eastern Pennsylvania, and Penn State's distinct position as the land-grant institution. The value of an independent executive director role, rather than someone embedded inside a single university, is also transferable: it opens access to state government conversations that would not be available from inside a faculty or administrative role at any one institution.

Utah: The Redtail AI Factory and the University of Utah

Snapshot and Context

The Redtail AI Factory is supported by the State of Utah and led by the University of Utah, with philanthropic funding from the Huntsman Family Foundation and Brynn Ballard and Peter Huntsman, as detailed in a university news release. The factory builds on existing AI infrastructure developed through prior investments and grant funding. It will serve researchers, educators, students, businesses, and public servants throughout Utah, reflecting a statewide mission that extends beyond the university campus.

The project's animating need is captured in University of Utah President Taylor Randall's framing: “Our goal is ensuring the state is awash in computing power by building a robust and scalable AI ecosystem that benefits our entire higher education system, powers research to find new healthcare cures, and ignites Utah's entrepreneurial spirit.” The educational goal centers on the Utah System of Higher Education (USHE) institutions and the broader Utah community as key stakeholders.

Governance and Funding

The University of Utah has entered into a multi-year HPE GreenLake agreement under which HPE provides the AI infrastructure, lifecycle management, and associated services that power the Redtail AI Factory. The system is housed at a commercial data center rather than on campus. User-side operations, allocations, and enablement are handled by the University of Utah's Center for High-Performance Computing (CHPC) at the Scientific Computing and Imaging (SCI) Institute, complementing CHPC's existing on-premises computing, data, and AI infrastructure at the university-owned Downtown Datacenter.

On financial sustainability, the university is exploring a combination of contributions from state government and state higher-education institutions, private donors, and a possible fee-for-service model for certain use cases and user segments. Resource allocation will support multiple approaches, including immediate start-up allocations, traditional proposal-based allocations, and cloud-like, credit-based on-demand access. The Redtail system came online in Summer 2026; early enablement was launched through a Challenge program starting in July 2026, with normal operations and allocations beginning in Fall 2026.

The philanthropic donor behind the initiative is motivated by accelerating research and development of cancer treatments using AI. For HPE, the Utah AI Factory is its first partnership with a state and an academic customer for its sovereign AI infrastructure solution (GreenLake), and the company is looking to build out similar partnerships elsewhere. Separately, Utah has signed an MOU with NVIDIA for statewide workforce training and research in AI and advanced technology.

Infrastructure and Access

Redtail is an HPE GreenLake sovereign AI solution powered by NVIDIA technology, chosen for its proven architectures and state-of-the-art support for both current and future AI workloads

as well as existing HPC workloads. Speed of deployment was the primary consideration behind acquiring a full vendor-provided at-scale solution rather than self-building, and colocation at a commercial data center was driven by footprint, power and cooling constraints, and the cost of expanding those at the university's own data center.

Access is structured as a statewide system available to researchers, educators, and innovators across Utah. Initial enablement has focused on USHE institutions (all public universities, community colleges, and technical colleges in the state) with K-12 support expanding through a partnership with the Utah State Board of Education. The university has launched an initial challenge call for projects open to all USHE institutions and is considering a gateway mechanism to enable scalable institution-level access to the system going forward.

The ultimate scale the Redtail AI Factory will reach is not yet determined, given that the system launched recently, though the aim is broad impact across the state. More than 4,000 active users are already served by existing SCI and CHPC advanced computing and data resources, and Redtail is expected to expand this further to new user segments.

Workforce and Educational Development

The university is establishing a program of training, access, and facilitation for users at all levels, initially focused on research and education faculty across domains and campuses. This builds on current mechanisms including regular training and user support from CHPC, office hours and community-of-practice convenings from the Office of AI, AI special interest groups facilitated by the Responsible AI Initiative, upskilling programs from the Kahlert School of Computing, and initiatives such as the weekly AI Foundations and Tinker Lab hosted by Digital Learning Technologies.

For students, a parallel program of training, access, and facilitation covers student and professional development and entrepreneurship, building on the faculty-facing mechanisms above along with an annual student-led AI symposium and industry credentials embedded in AI-related coursework.

Beyond formal coursework, the university is building a coordinated AI learning, development, and innovation ecosystem with multiple entry points for students, faculty, researchers, entrepreneurs, industry, and government partners. This includes AI Ambassadors who foster AI literacy and community engagement; AI Innovation Labs providing collaborative space for hands-on learning and interdisciplinary experimentation; an AI Learning, Development, and Evaluation Sandbox for safely prototyping and testing AI applications before production deployment; and AI Engineers who provide technical expertise to help translate ideas into working AI solutions. Together, these programs are meant to create experiential pathways that accelerate AI adoption, workforce development, innovation, and responsible AI practice across the university and the state.

Timeline

- 1989-1994: Utah Supercomputing Institute (USI) was launched to become CHPC; SCI Institute was founded.

- October 2023: Launch of the University of Utah One-U Responsible AI Initiative with a University investment of $100M to advance societal impact and translation of AI.
- July 2024: Utah Office of AI Policy established to provide learning labs and regulatory mitigation and shape AI policy in Utah.
- April 2025: University of Utah Office of AI established to support statewide academic-industry AI coordination under the direction of the Chief AI Officer.
- September 2025: Board of Trustees approval for the public and private partnership for the AI supercomputer.
- February 2026: Legislative support of $15 million for an AI factory and broader AI ecosystem.
- June 2026: The Redtail AI Factory is online and reached rank 159 on the TOP500 list and rank 76 on the Green500 list, which measures performance per unit of power.
- Summer-Fall 2026: Utah higher education-wide enablement Challenge was released in early July 2026, and the first selected projects are expected to onboard in late summer. Follow-on enablement activities will target additional user segments across the state. Redtail builds on and empowers an evolving sovereign Utah AI infrastructure ecosystem of long-standing advanced computing, data, and network infrastructure.

Lessons Learned

Because the Redtail AI Factory is in its early stages, the lessons learned so far are broad ones: the growing importance and wide need for enablement, the impact of a rapidly changing AI technology landscape, and the complexity of developing effective sustainability models. The university stresses the importance of clearly articulating requirements up front — the types and range of workloads and workflows to be enabled, connectivity and storage needs, and the various trade-offs among technology and architecture generation, capacity and scale, physical constraints, connectivity speeds, compliance needs, and overall budget envelope. Working with a large system vendor also puts a premium on customer-side experience with large project contracting, project management, scheduling and milestone setting, requirements definition, change-order processes, and verification and acceptance — and negotiating all of these well.

Regarding what could be done differently, the university is still learning as it proceeds. Much of the current work has gone into enablement plans meant to help the community prepare for an AI factory before it fully ramps up. The university's framing is that people make the difference in successful outcomes, and it plans to actively monitor the user experience and assist users “where they are.”

On what's transferable to other states and regions, Utah offers a model for public-private partnership-funded infrastructure that supports statewide needs, which other states can adapt, along with an emerging demonstration of how a single piece of infrastructure can simultaneously support research, education, and innovation.

New York: Empire AI

Snapshot and Context

Empire AI is a New York State consortium that gives public and private research institutions access to advanced computing resources for artificial intelligence research. Led by New York State in partnership with the consortium's member institutions, it represents the first effort of its scale in the United States to place world-class AI computing infrastructure in the hands of universities and nonprofit research organizations. The consortium includes Columbia University, Cornell University, the City University of New York (CUNY), the Icahn School of Medicine at Mount Sinai, New York University (NYU), the Rochester Institute of Technology (RIT), Rensselaer Polytechnic Institute (RPI), the University of Rochester, the State University of New York (SUNY), and the Flatiron Institute.

Empire AI was created in response to a widening infrastructure gap: training and developing advanced AI models increasingly requires computing systems that are prohibitively expensive for individual universities to acquire and operate, while the global market for advanced AI processors has become concentrated among a small number of major technology companies. Even well-resourced universities can struggle to obtain enough computing capacity, and smaller institutions, public universities, and researchers in less commercially prominent disciplines face still greater barriers. Empire AI addresses this by pooling public investment, institutional expertise, and philanthropic support into a single shared statewide resource rather than requiring each institution to build its own AI computing center.

Governance and Funding

Empire AI is structured as a consortium rather than as a computing facility belonging exclusively to a single university. The computing systems are located at the University at Buffalo, which provided the initial data center space; the final, permanent data center is being constructed on land on the Buffalo campus. The consortium is governed by a board with representatives from each of the participating partners: nine institutional members mapping to individual institutions or combinations of institutions, plus external board members who bring outside expertise. The board oversees the project, ensures consortium and state funding is used responsibly, and confirms the project is meeting its goals and obligations to the community. Day-to-day operations are handled by a dedicated team with an executive director and, beneath the director, teams for system administration, scientific and user support, and facility construction and operations, split between Manhattan and Buffalo.

The funding model combines state funding, institutional participation, and philanthropic support. The initial Alpha system was made possible in part through a major contribution from the Simons Foundation. New York State's investment supports development of the larger computing center and the purchase and upgrade of computing infrastructure, and each consortium partner provides an annual operations contribution with a ten-year support commitment, giving Empire AI a stable funding model for the next decade, one long enough to justify both the facility investment and the effort of relocating research programs onto shared infrastructure.

Computing resources are allocated across the consortium through an accounting system based on service units: each institution is awarded the same number of service units annually to divide among its own research projects, with each type of resource (GPU, CPU, storage) representing a specific number of service units per unit time. This ensures all members receive an equal share of the resources while allowing each to make its own local priority choices. Service units also carry a standard monetary value. It is not used for billing within Empire AI, but available for consortium members to charge back to grants if they choose.

Infrastructure and Access

Empire AI is being built in phases, allowing researchers to begin using the infrastructure while the full facility is still under construction. The first system, Alpha, came online in October 2024 at the University at Buffalo's Center for Computational Research, providing access to NVIDIA H100 accelerators for training and evaluating advanced AI models. It has since expanded steadily: Alpha entered service in October 2024 with approximately 100 H100 GPUs; Alpha+ expanded the system to 144 H100 GPUs in June 2025; Alpha++ increased the total to 192 GPUs (144 H100 and 48 H200) in December 2025; and Beta entered service in June 2026 with 288 B200 GPUs configured in NVL72 systems. Alongside processing resources, Empire AI has deployed 20 PB of high-performance storage and a dedicated network connection to NYSERNet.

This phased approach gives users early access to scarce computing resources, lets the consortium develop its operating and allocation procedures incrementally, and creates opportunities to learn from actual workloads before the full facility is complete. The permanent Empire AI computing center is scheduled for completion in late 2027 or early 2028 on the University at Buffalo's North Campus, with sustainability built into its design: the facility is expected to run on renewable hydropower from Niagara Falls and to recover heat generated by the computing equipment for use in campus dormitories, a deliberate response to the growing concern that advanced computing access must expand without letting the energy and environmental costs of AI grow unchecked.

All partner institutions are treated equally and have equal access to the infrastructure, which helps level the field for lower-resourced institutions. Access alone, however, is not sufficient: institutions with limited local expertise may need help preparing data, adapting software, managing accounts, using accelerators efficiently, and complying with security and responsible-AI requirements. Empire AI's operations budget includes dedicated effort for user support and the broader scientific program, on the premise that the consortium has its greatest statewide impact when computing allocations are paired with technical consulting and development.

Workforce and Educational Development

Empire AI is also explicitly a workforce-development initiative. By giving students and faculty direct experience with advanced AI infrastructure, the consortium aims to help New York develop the technical talent needed by universities, government, nonprofit organizations, and AI-enabled industries. The shared infrastructure creates an additional educational benefit beyond simply using AI models: students learn how large-scale AI systems are designed,

operated, evaluated, and governed, including training in data stewardship, reproducibility, cybersecurity, energy efficiency, algorithmic bias, model evaluation, and the societal consequences of AI deployment.

Timeline

- 2024: New York establishes the Empire AI consortium and begins development of a shared statewide AI-computing capability.
- October 2024: The Alpha system begins operation at the University at Buffalo, with approximately 100 H100 GPUs.
- June 2025: Alpha+ expands available H100 capacity to 144 GPUs.
- December 2025: Alpha++ introduces additional H100 and H200 capacity, reaching 192 GPUs total.
- June 2026: The Beta system comes online with 288 B200 GPUs in NVL72 configurations.
- 2026–2028: User programs, governance processes, technical support, and educational activities continue to expand.
- 2028: The permanent Empire AI computing center is scheduled for completion on the University at Buffalo's North Campus.

Lessons Learned

Empire AI demonstrates that advanced AI infrastructure can be treated as a shared regional capability rather than a collection of institution-specific purchases. It’s a model that lets a state aggregate demand, secure scarce hardware, build specialized operational expertise, and make advanced resources available to a broader academic community. The consortium has also learned that the gap between what's available to academia and what's available to industry is both financial and technical: some of the resources being deployed are unique in academia, with no existing community knowledge base to draw on. Empire AI is rapidly building up technical expertise, operational experience, and the start of best practices. This is knowledge the consortium expects will be valuable to the next initiative that attempts something similar.

The most transferable element of Empire AI is not a particular processor or facility design, but the idea that a state or region can create a shared AI research utility by combining public leadership, university collaboration, philanthropy, and professional computing operations. For other states, the central lesson is that infrastructure should be organized around a clearly defined public mission: a successful regional initiative must combine computing capacity with fair access, expert support, faculty development, student pathways, sustainable funding, and governance that represents institutions of different sizes and missions. Empire AI's own framing is that New York is seeking not merely to participate in the AI economy, but to shape how AI is developed and used, by building shared infrastructure and placing it in service of research, education, and public benefit.

Section 3: Conclusion

Malachowsky's premise is that a state's advanced computing capability should be viewed like public infrastructure similar to electricity, clean water, and transportation. The computing infrastructure can be owned locally, funded through a blend of philanthropy, industry, state, and institutional investment, and extended outward on a hub-and-spoke model. Florida, Pennsylvania, Utah, and New York are four live tests of that premise, though no two have run the experiment the same way. Florida centralized ownership in one anchor institution and a priority access model governs access. Pennsylvania is still negotiating how seven independent universities can share ownership without any one of them dominating it. Utah bought speed by handing infrastructure delivery to a vendor rather than building it in-house. New York legislated a ten-institution consortium into existence and allocates compute time in equal shares precisely so no partner can be seen as more important than another. Each choice was shaped by that state's own politics, institutional history, and funding reality. The benefit of profiling four states in this report, rather than one, provides a variety of examples for other states working to develop their own path.

The CRA-NSF Summit's breakout sessions, synthesized in the Thematic Findings document, raised the following concerns, and the four vignettes discussed here make them concrete rather than abstract. Summit participants warned that hardware without dedicated human infrastructure (partnership staff, research software engineers, data librarians) becomes underutilized; Florida's roughly 40-person staffing build-out, and its warning that "hardware without people fails," is that lesson already lived out. Summit participants warned against conflating elite research infrastructure with basic educational infrastructure; Florida's split between free teaching access and priced research access, and Utah's still-forming tiered allocation model, are two different attempts to keep those tiers decoupled. Summit participants flagged IP and contracting delays as a leading cause of stalled momentum; Pennsylvania's deliberate decision to defer detailed IP agreements until they are actually needed, rather than let them block the Keystone Factory's launch, is a direct response to that exact failure mode. Finally, Summit participants worried about geographic inequity: that only already-wealthy institutions would benefit (a concern New York's equal-service-unit allocation and Florida's statewide Board of Governors access agreement were each designed, in their own way, to avoid).

CASC's *Regional Collaboration as Strategic Infrastructure* paper (mentioned in Section 1) also reached the same conclusion the Summit's Santa Clara breakout-session participants and this report's four state contributors arrived at on their own: durable regional collaboration comes down to four requirements, namely a sustained catalyst, institutional buy-in, sustained funding, and trusted governance, the same four visible across all four vignettes in this report. When a Santa Clara convening of practitioners, this report's four vignettes, and a separate cross-institutional research-computing analysis all converge on the same short list of what actually determines success, that list is worth taking seriously.

None of the four states are done in their development and implementation. The director of the Pennsylvania effort is candid that centralized infrastructure "doesn't exist yet." Utah's system is months old. New York's permanent facility will not be finished until 2027 or 2028. Even Florida,

the most mature of the four and the one Malachowsky draws its blueprint from directly, is still building pieces of what the white paper envisions, five years after its founding gift. We hope this reassures rather than discourages a state just starting out. Willingness, not completeness, is what each of these efforts actually had in place at the moment their opportunity arrived. Florida's CIO built a data center a decade before anyone asked for one. The director of the Pennsylvania effort began lobbying Harrisburg years before DCED called back. Utah's university had existing CHPC infrastructure and a Responsible AI Initiative already running before Redtail was ever proposed. New York's consortium had governance and allocation rules designed before Alpha came online, not after.

For a state or institution weighing whether and how to start, the four vignettes herein suggest a short, practical checklist drawn directly from what the authors of this report said worked and what they wished they'd done sooner.

- First, build data-center and compliance readiness before an opportunity arrives, rather than after.
- Second, decide early whether ownership will sit with one anchor institution or a multi-institution consortium, since that choice determines almost everything else about governance and allocation.
- Next, treat philanthropic and industry investment as a partnership requiring matching institutional commitment, not a gift with no strings attached. It is also critical to invest in human infrastructure (such as training, support staff, and research software engineers) at the same time as the hardware.
- Finally, expect governance and funding sustainability to take longer to resolve than the technology does.

The four states in this report are early evidence that the movement toward computing public infrastructure is already underway, unevenly and imperfectly, and that each new state joining the movement has real, documented predecessors to learn from rather than a blank page.

Acknowledgments

This report draws on philanthropic efforts led by Chris Malachowsky (Cofounder, Fellow, and Senior Vice President, NVIDIA), whose framing of a 50-state AI computing initiative and direct role in the University of Florida's AI transformation shaped Section 1 of this report.

Tracy Camp (Executive Director and CEO, Computing Research Association) and Kathryn Kelley (Executive Director, Coalition for Academic Scientific Computation) were close partners throughout the process, and we are grateful for their valuable contributions.

This report also draws on discussion and findings from the CRA-NSF Industry/Philanthropy Summit on Potential NAIRR State Hubs for Educational Transformation, held February 17, 2026, in Santa Clara, California, and jointly convened by the Computing Research Association (CRA) and the National Science Foundation (NSF). The Summit's breakout-session findings are synthesized separately in the companion Thematic Findings report and its associated CRN article. We appreciate the conversations that came from that event, as well as a later conversation between Tracy Camp and Manish Parashar, which ultimately led to this report.

Sources

- *Building an AI University: From Infrastructure Gift to Systemic Transformation*, University of Florida Press, Warrington Case Series, 2026.
- Coalition for Academic Scientific Computation, "Regional Collaboration as Strategic Infrastructure: A Framework for the Future of Research Computing," June 2026.
- CRA-Industry and the National Science Foundation. "Thematic Findings from the CRA-NSF Industry/Philanthropy Summit on Potential NAIRR State Hubs for Educational Transformation," 2026.
- CRA-Industry. "Building AI Infrastructure for All: Key Takeaways from the CRA-NSF Industry/Philanthropy Summit," May 2026.
- Chris Malachowsky. "The New Frontier of Education: The 50-State AI Computing Initiative to Empower America for the AI Era," Rev 2.4, 6/14/26.
- Coalition for Academic Scientific Computation, in partnership with the National Science Foundation. "NSF–CASC Workshop on Envisioning NAIRR State and Regional Hubs for Educational Transformation" (Executive Summary), March 9, 2026, Washington, DC.
- *From Compute to Competitiveness: The Vice President for Research*, University of Florida Press, Warrington Case Series, 2026.
- *From Infrastructure Stewardship to Institutional Capability: The Chief Information Officer*, University of Florida Press, Warrington Case Series, 2026.
- Interviews and contributed source documentation gathered around CRA and NSF's 2026 Industry/Philanthropy Summit on Potential NAIRR State Hubs for Educational Transformation, including:
    - Erik Deumens (Senior Director, Research Computing, University of Florida)
    - Ian Fisk (Chief Technology Officer, Flatiron Institute, Simons Foundation)

- Manish Parashar (Chief AI Officer, Executive Director & Chair in Computational Science and Engineering, University of Utah)
- Barr von Oehsen (Director, Pittsburgh Supercomputing Center, and Vice Chancellor for Research Computing, University of Pittsburgh)