From Federated Resources to NGA: Bridging the Gap

Aaron Jezghani
Georgia Institute of Technology

Georgia Tech

Testbeds are more than a collection of hardware

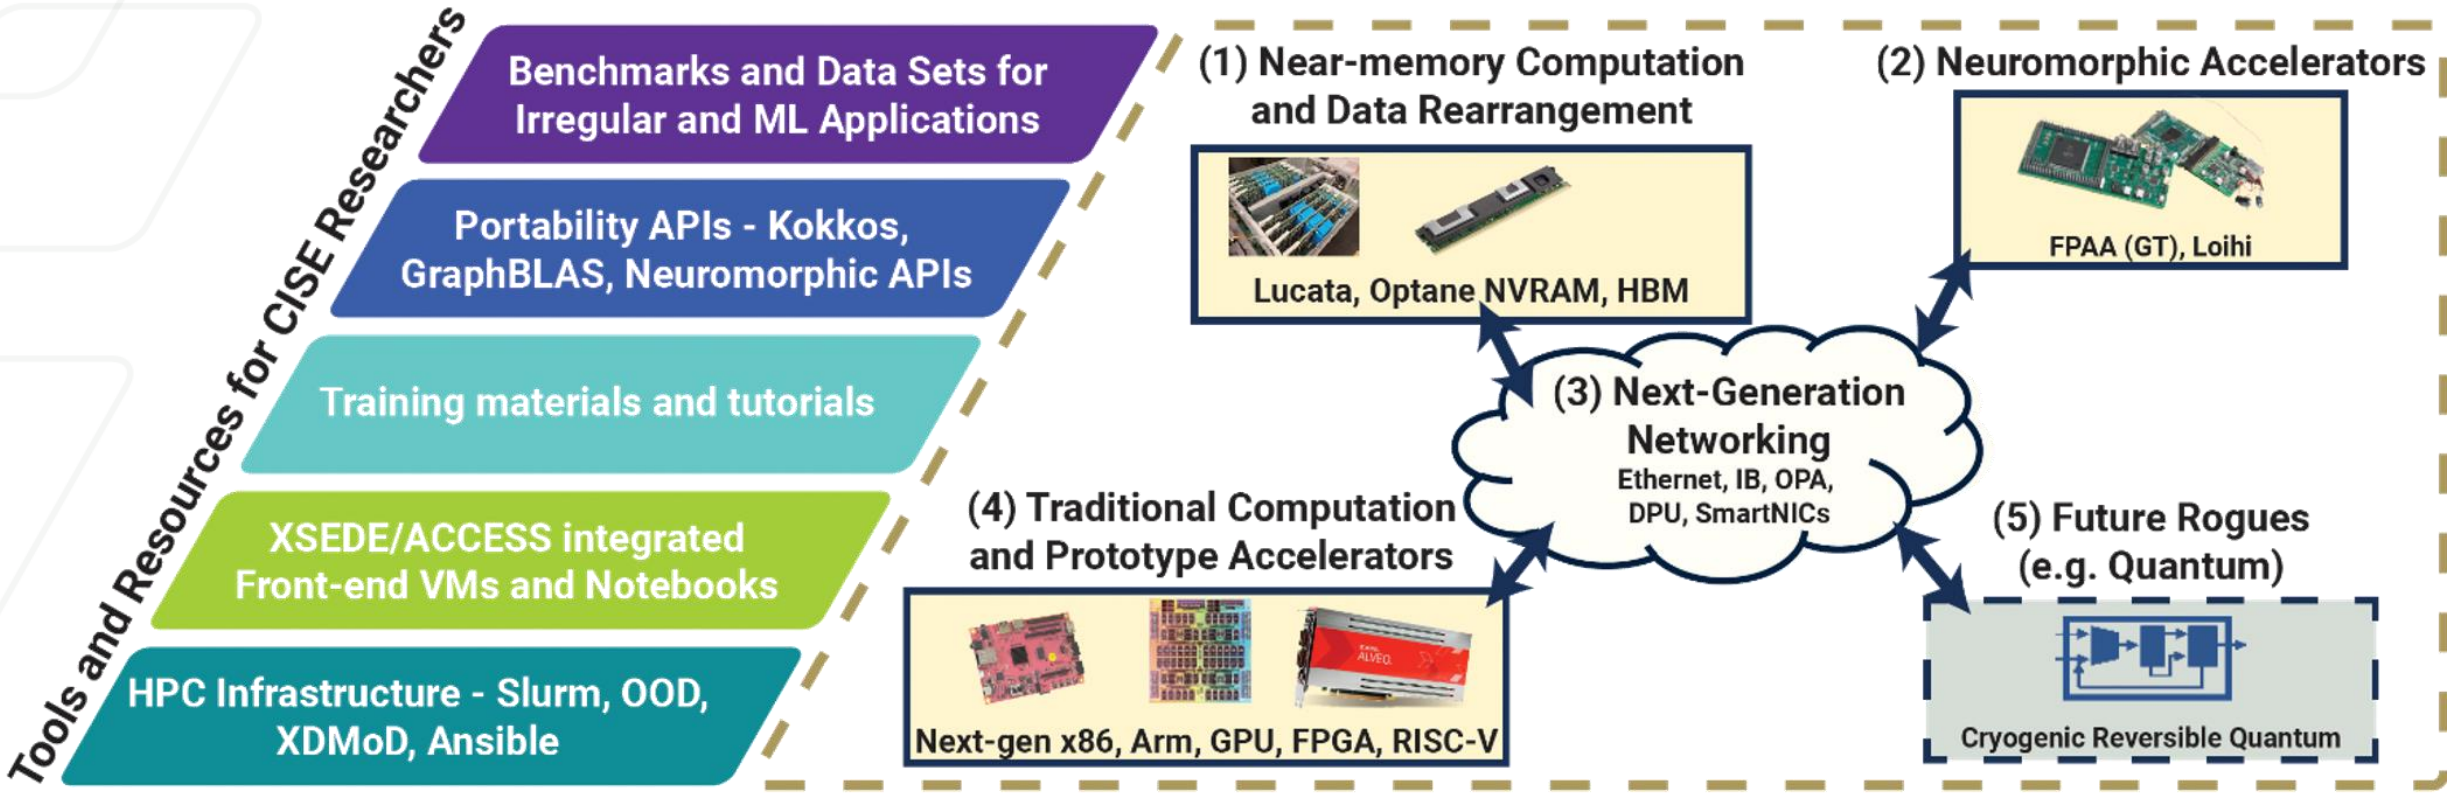

The primary goal of the Rogues Gallery is* to make first-generation start-up or research lab prototypes available to a wide audience of researchers and developers as soon as it is moderately stable for testing, application development, and benchmarking.

**Prior Funding and Support: NSF OAC #2016701, NSF OAC #1828187, Sandia National Labs, Vendor Donations, GT TechFee Funds. Currently operating as a best-effort resource.*

Georgia Tech

An instructive example: Lucata Pathfinder



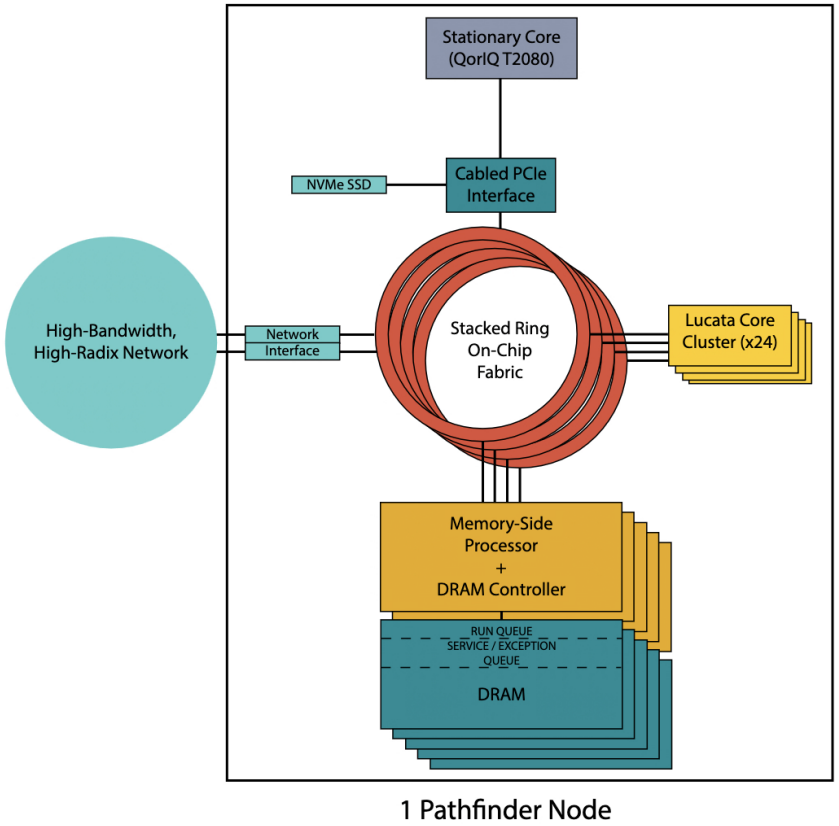

1 Pathfinder Node

Takeaway: The testbed system yielded several research papers, unique instructional opportunities for undergraduate students, and it supported a major govt project (IARPA). *BUT* it was difficult to program, training lagged HW/SW updates, and it was hard to maintain, which ultimately meant that it required support at level that we couldn't maintain it ourselves.

Georgia Tech

…but they are absolutely collections of hardware



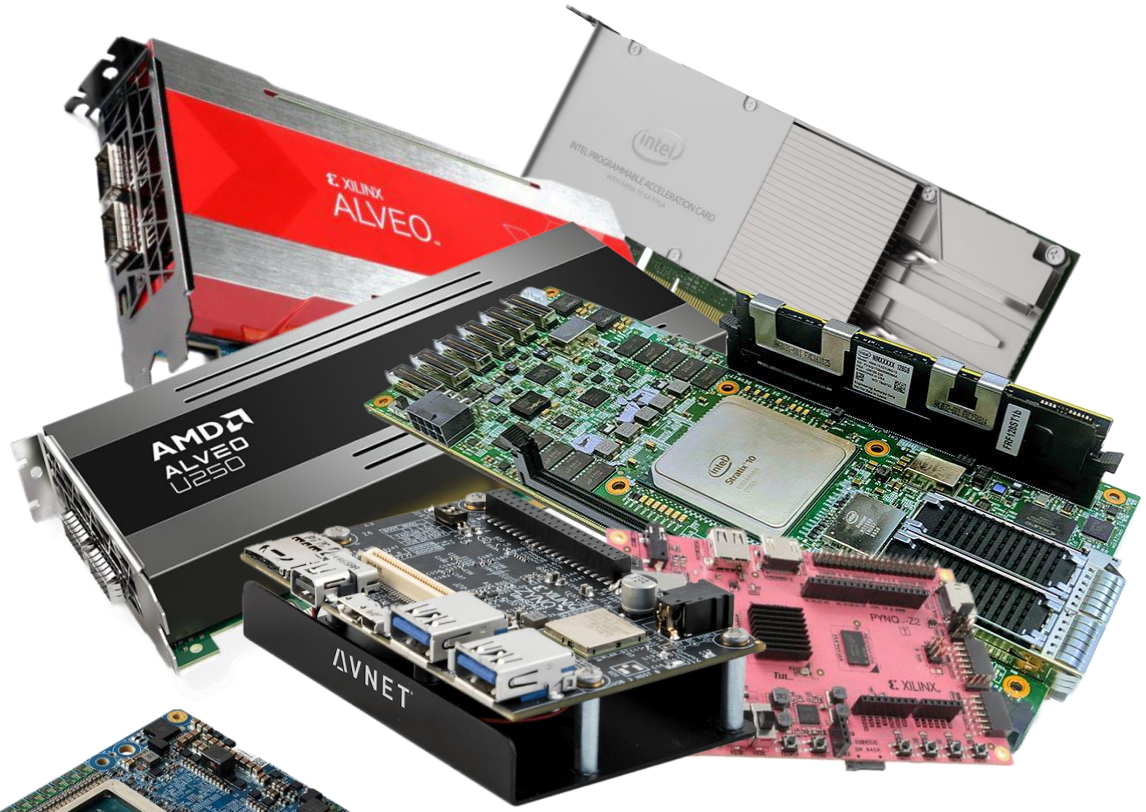

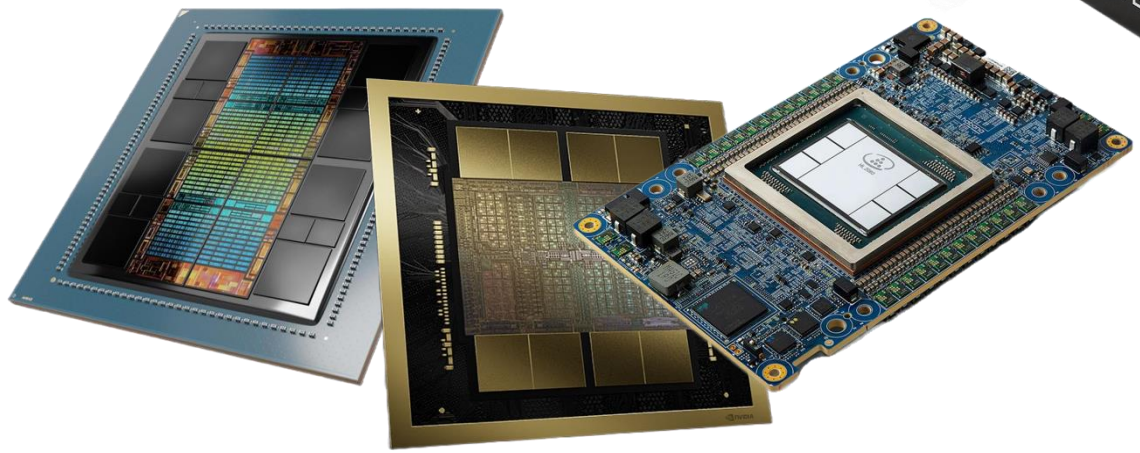

Georgia Tech

Testbeds span the globe

US / North America

Testbed	Sponsor	Who Can Apply	Notes
ACE Testbed (OLCF)	DOE Office of Science, Oak Ridge Leadership Computing Facility	Researchers with an OLCF project allocation	Advanced Computing Ecosystem testbed for emerging HPC/AI hardware, cloud-style workflows and quantum-HPC integration
ACES (Texas A&M)	NSF (#2112356), Texas A&M HPRC	US researchers via NSF ACCESS allocations	Composable accelerator testbed (e.g., Intel PVC GPUs, FPGAs, NEC Vector Engines, Graphcore IPUs)
ALCF AI Testbed	DOE Office of Science, Argonne Leadership Computing Facility	Open-science researchers via ALCF allocation request (e.g., Director's Discretionary)	AI accelerators such as Cerebras, SambaNova, Graphcore and Groq systems
AMD HACC (formerly Xilinx XACC)	AMD University Program	Non-commercial researchers	Heterogeneous Accelerated Compute Clusters with Alveo FPGAs and Instinct GPUs; North American sites at UIUC and UCLA (also hosted in Europe and Asia)
CENATE (PNNL)	DOE Office of Science, Pacific Northwest National Laboratory	Collaborators with PNNL researchers	Center for Advanced Technology Evaluation; early-stage evaluation of emerging architectures, including power and performance measurement
Chameleon Cloud	NSF	US researchers and educators (and their collaborators)	Bare-metal reconfigurable cloud testbed for systems research, with GPUs, FPGAs and edge devices
CRNCH Rogues Gallery	Georgia Tech CRNCH, NSF	Academic, government and industry researchers; request an account	Novel and post-Moore hardware (e.g., near-memory, neuromorphic, FPGA, RISC-V and AI accelerators) with a focus on early-stage evaluation
ExCL (ORNL)	DOE Office of Science, Oak Ridge National Laboratory	ORNL staff and collaborators	Experimental Computing Laboratory for emerging processors, memory and storage technologies
FABRIC	NSF	US researchers and educators (international collaborators may join a US-led project); sign up through the FABRIC portal	Nationwide programmable network testbed with GPUs, FPGAs and SmartNICs at many sites; suited to networking, distributed systems and edge-to-cloud research
HAAPS (Sandia)	NNSA Advanced Simulation and Computing (ASC)	Primarily NNSA/DOE tri-lab researchers and collaborators	Heterogeneous Advanced Architecture Platforms for evaluating future ASC system technologies
Maui High Performance Computing Center	DoD HPC Modernization Program, Air Force Research Laboratory	DoD users and DoD-sponsored researchers	DoD supercomputing center that has hosted novel-architecture evaluations
Open Cloud Testbed	NSF CCRI (#1925464), Mass Open Cloud	US researchers	AMD/Xilinx Alveo U280 FPGAs, networking and virtualization research; see the OCT tutorials
THOR: The Neuromorphic Commons	NSF, MATRIX AI Consortium at UT San Antonio (with UT Knoxville and UC San Diego)	Community researchers; request access via the project website or thor@utsa.edu	Open neuromorphic computing testbed hosted at UTSA, including a SpiNNcloud SpiNNaker2 system

Europe

Testbed	Sponsor	Who Can Apply	Notes
AMD HACC (formerly Xilinx XACC)	AMD University Program	Non-commercial researchers	European sites at ETH Zurich (Switzerland) and Paderborn University (Germany)
BSC RISC-V Testbed / Software Development Vehicle	Barcelona Supercomputing Center, EuroHPC / European projects	European research collaborators; contact BSC	Spain. RISC-V hardware and emulation platforms for HPC software development
CSCS	Swiss National Supercomputing Centre (ETH Domain)	Swiss and international researchers through CSCS calls and collaborations	Switzerland. Hosts early-access and test systems for new architectures alongside production machines
EPCC Testbeds	EPCC, University of Edinburgh; UKRI	UK researchers and collaborators; contact EPCC	UK. Novel-architecture testbeds (e.g., Cerebras, RISC-V and FPGA systems)
JUNIQ (Jülich)	Forschungszentrum Jülich, German federal and state funding	German and European researchers via JUNIQ access calls	Germany. Jülich UNified Infrastructure for Quantum computing; quantum annealers, simulators and quantum-HPC integration
UCL ARC Testbeds (RCNIC, Contender)	University College London, Advanced Research Computing	UCL researchers and collaborators	UK. Research Computing and Networking Innovation Centre and the "Contender" testbed for evaluating candidate hardware

Asia

Testbed	Sponsor	Who Can Apply	Notes
AMD HACC (formerly Xilinx XACC)	AMD University Program	Non-commercial researchers	Asian site at the National University of Singapore
RIKEN R-CCS (Fugaku)	RIKEN Center for Computational Science, MEXT (Japan)	Researchers worldwide via HPCI and Fugaku proposal calls	Japan. A64FX Arm-based Fugaku, plus research into future architectures and quantum-HPC hybrid platforms

https://github.com/caatb/testbed-resources/blob/main/README.md

Georgia Tech

But how do we make testbeds practical?

- Funding is becoming scarcer
- Rising hardware costs (especially memory) are real
- Security is becoming a primary challenge
- Regular IT infrastructure may be incompatible
- Functionality may not align with normal policy

Allocations for users are often preferred…

Federated Identity

ACCESS

NRP NATIONAL RESEARCH PLATFORM

OSG

PATh

Institutional Identity

GEORGIA TECH

OPEN RESEARCH COLLABORATION



Georgia Tech

Gateways are great for tutorials…

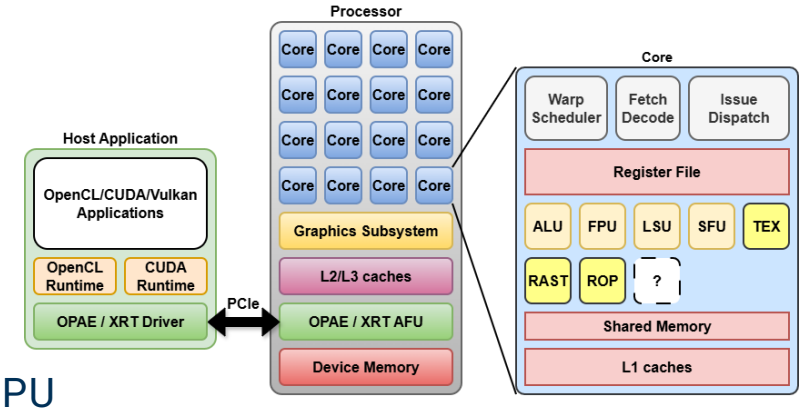

jupyterhub + GitHub

Hardware
- Fully configurable multi-core GPU
- PCIe host-device interface
- Floating-point Units, caches & shared Memory
- 32 and 64-bit ISA

Open-source toolchain
- LLVM: Compiler
- POCL: OpenCL
- CuPBoP: CUDA
- OPAE: Intel FPGA API
- XRT: Xilinx FPGA API
- Verilator: RTL simulation
- Yosys: RTL Synthesis

FPGA Synthesis
- Altera: Arria10, Stratix 10
- Xilinx: Alveo U50, U250, U280, U55

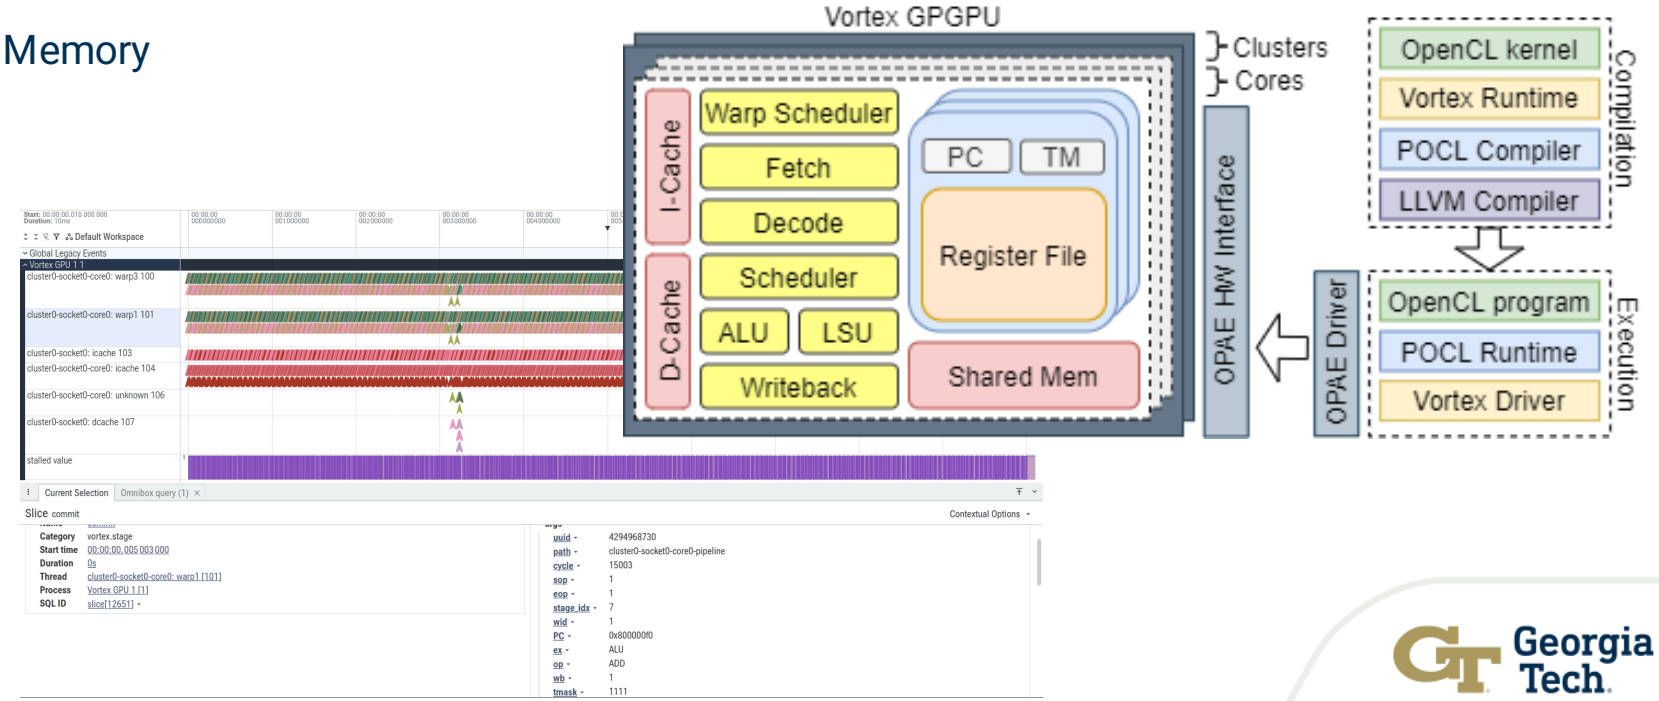

More efforts considering FaaS or similar

Remote Kernels

```
In [1]: %%javascript
        var kernel = Jupyter.notebook.kernel;
        kernel.execute('kernel_name = "' + kernel.name + '"');

In [2]: print(kernel_name)
        python3

In [3]: !hostname
        login-phx-dev.pace.gatech.edu

In [4]: a=5

In [5]: print(a)
        5

In [6]: %store a
        Stored 'a' (int)

In [1]: %%javascript
        var kernel = Jupyter.notebook.kernel;
        kernel.execute('kernel_name = "' + kernel.name + '"');

In [2]: print(kernel_name)
        ssh__sshtest

In [3]: !hostname
        atl1-1-03-002-35-0.pace.gatech.edu

In [4]: print(a)
        ---------------------------------------------------------------------------
        NameError                                 Traceback (most recent call last)
        Cell In[4], line 1
        ----> 1 print(a)

        NameError: name 'a' is not defined

In [5]: %store -r a

In [6]: print(a)
        5
```

Cyber-Shuttle

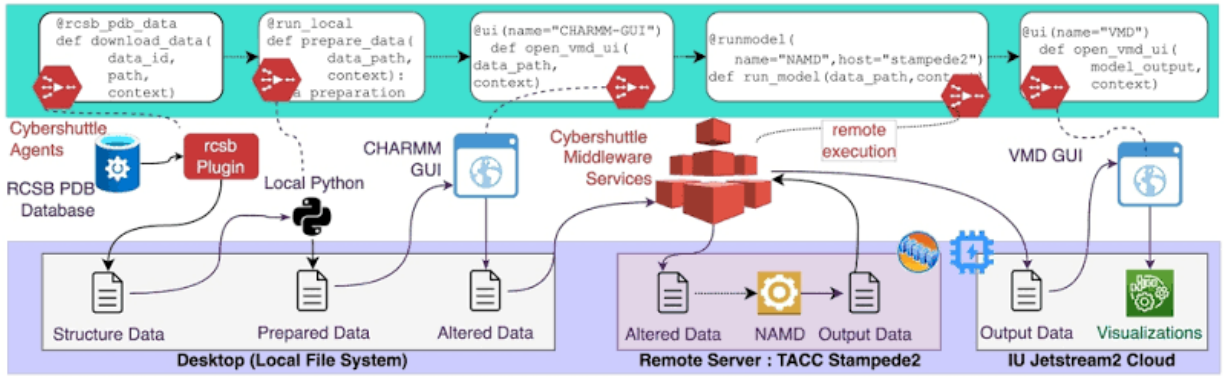

Globus Compute

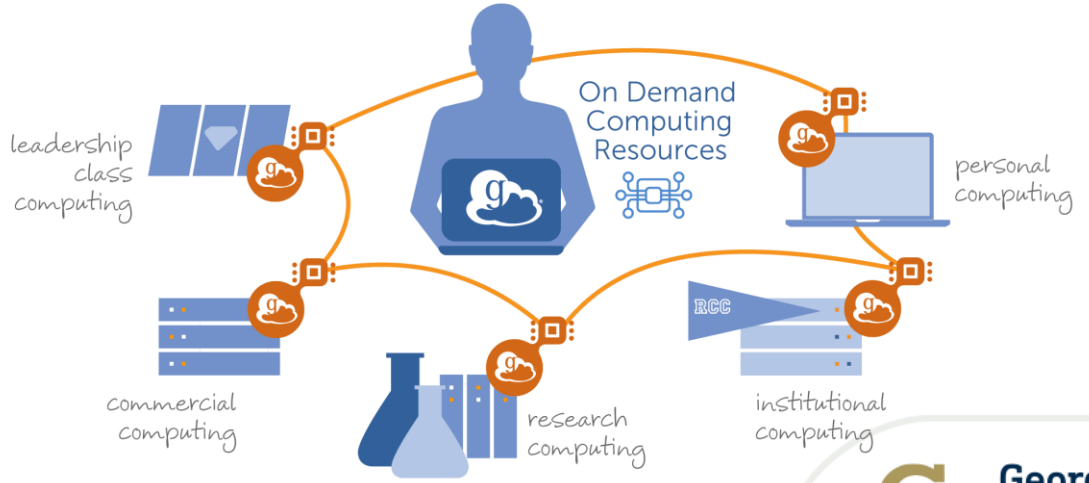

Georgia Tech

The shift in US scientific funding

Genesis Mission





Georgia Tech

NSF Programmable Cloud Laboratories

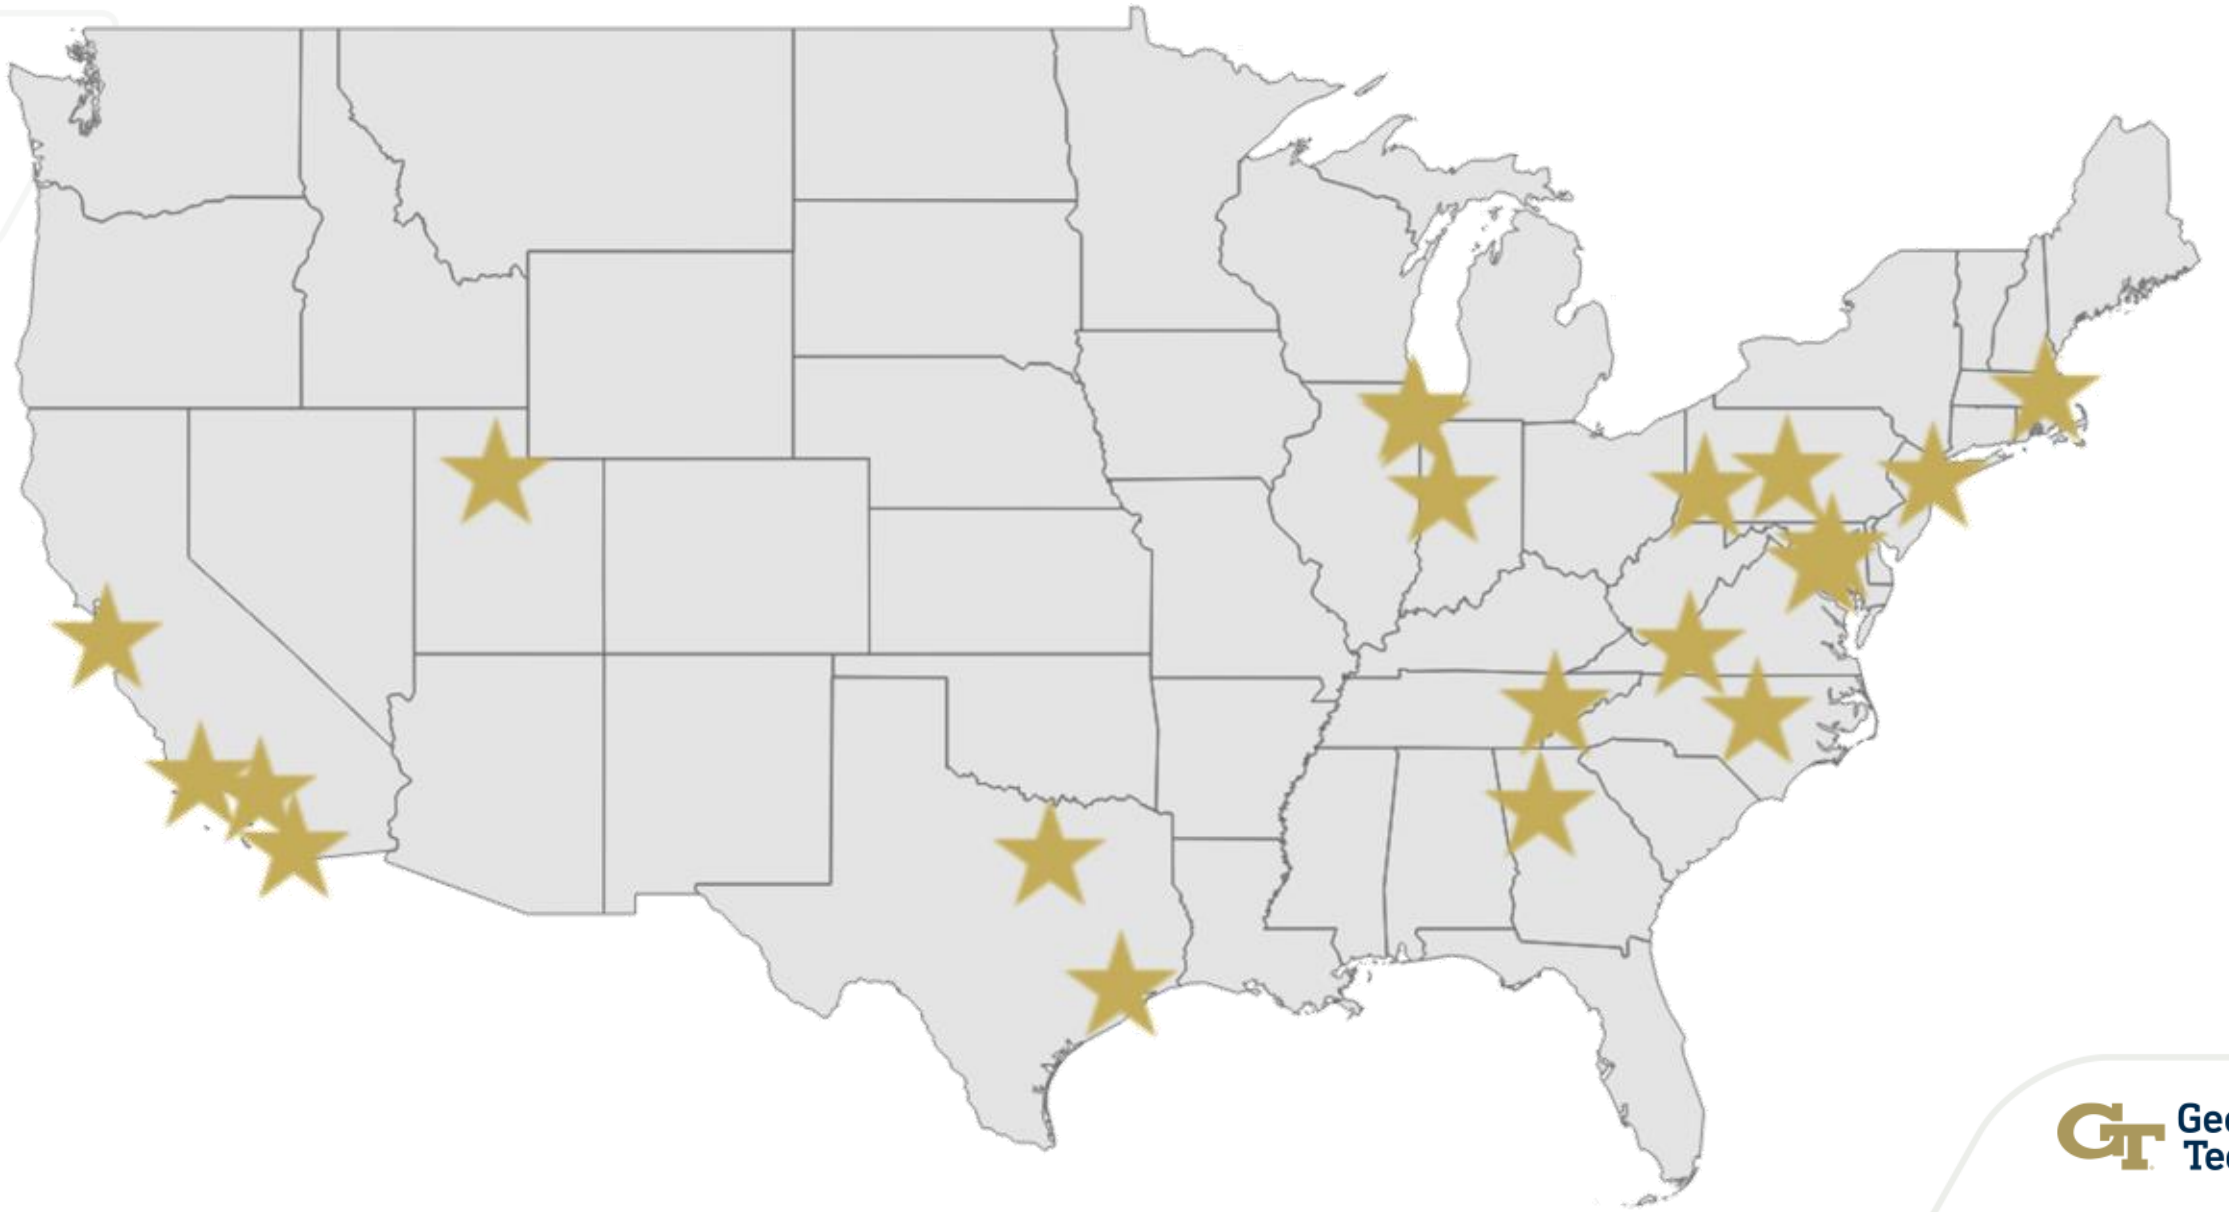

Georgia Tech

PCL Focus Areas

- **Biotechnology, Biology, & Chemistry Nodes**
  - CLAIRE (Boston University), DREAM (Northwestern), The CRAB Lab (University of Maryland + NIST), AI-BIOME (CUNY), AI Science Foundry (Carnegie Mellon + ANL + LLNL), AI-Driven High-Throughput Reaction Discovery Platform (Scripps)
- **Materials Science & Manufacturing Nodes**
  - SPEED (NCSU), COAST PCL Node (UCSB), PoLARIS (University of Chicago + ANL), PCL-AIMMPF (Georgia Tech), AIMD-NET (Johns Hopkins), ATHENA (University of Tennessee), ICoN-PCL (Purdue)
- **Semiconductors, Electronics, & Quantum Nodes**
  - AI-Enabled, Cloud-Based Self-Driving Nanofabrication Node (Virginia Tech), LATTICE (Penn State), Bringing Electronics Prototyping Home (Adom Industries), READINESS (Rice), AURORA (University of Utah)
- **Characterization & Cross-Network Standards Nodes**
  - ELECTRA (Caltech), GEMSTONE (Stanford + Emerald Cloud Lab)

Georgia Tech

The AI Manufacturing and Materials Pilot Facility Node (GT)

"[E]mpower U.S. researchers to invent, deploy, de-risk, and use AI to more rapidly discover and deploy new materials and manufacturing technologies using autonomous research workflows of programmable experiments and simulations"

1) Create the ability to program research workflows from the cloud and by AI agents
2) Develop open metadata and data standards for self-driving labs
3) Develop an innovative AI-orchestrated digital twin that uses LLMs to assist users who plan, monitor, and analyze their research goals
4) Boost AI abilities to translate materials and manufacturing scientific queries into autonomous research and development workflows
5) Pioneer AI-agents that continuously optimize the planning and scheduling operations of PCL-AIMMPF workflows in uncertain environments
6) Create and deploy data assimilation AI-agents that use sensors from the facility to estimate the state of the scientific research progress
7) Increase the PCL-AIMMPF equipment and simulations capable of use in autonomous workflows
8) Develop technologies and standards that allow scientific workflows from one node to be combined with those of others to form trans-nodal research pathways.

Georgia Tech

In short

- Next-generation hardware and software needs more than just a few great minds
- Democratization is critical – sustainable operations, extended applications, broader adoption
- NSF and DOE programs are relevant and provide opportunity for innovation

Some relevant SC Events

SmartNIC Tutorial
Sunday, Nov 15 | 8:30-5:00 CT | Room W193b
Hands-on introduction to programming and using SmartNICs.
https://github.com/gt-crnch-rg/smartnic-tutorial-sc26

Advanced Architectures Birds of a Feather (BoF)
Tuesday, Nov 17 | 12:15-1:15 CT | Room W194a
Come learn the latest from testbeds that support emerging HPC architectures.
https://caatb.github.io/aatb-bofs/sc-2026-bof.html

Correctness Testing of HPC Scientific Software in the Age of AI-Generated Code (BoF)
Wednesday, Nov 18 | 12:15-1:15 CT | Room W176abc
As AI-generated code makes its way into scientific software, how do we make sure it's correct? Join us as we try to figure it out!
https://sc26.conference-program.com/presentation/?id=bof160&sess=sess389

Georgia Tech