# Cornell

SC17: Cornell Briefing

Cornell Center for Advanced Computing (CAC)

# Profile

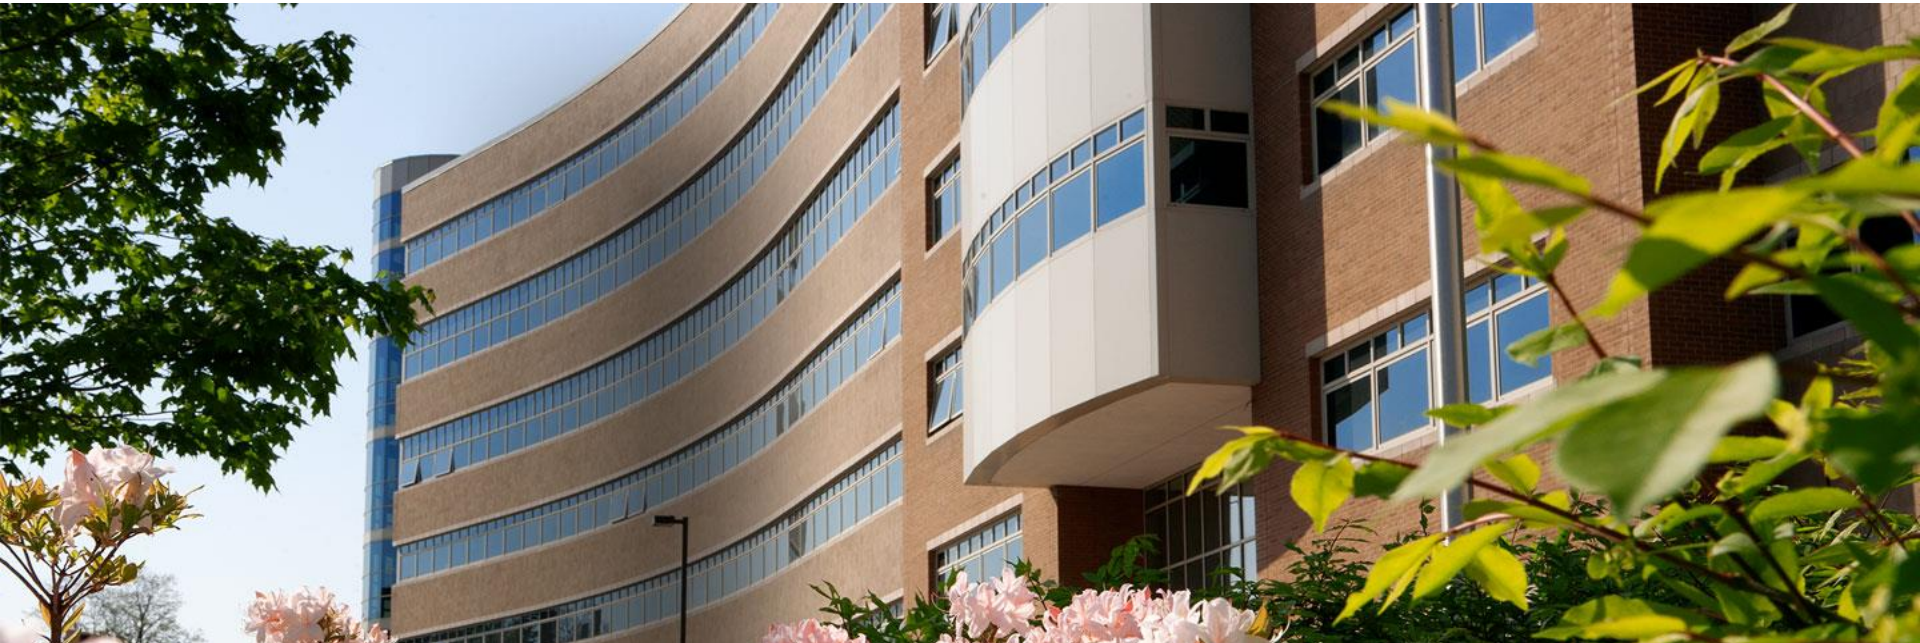

# Cornell CAC overview

- Mission
  - Deliver advanced computing, training, and consulting services that accelerate discovery and broaden impact
- Services
  - HPC systems design/hosting, cloud computing, storage, programming, code optimization, scientific workflow design, db design, visualization
- Impact
  - Supporting Cornell faculty with over $100 million in research funding from NSF, NIH, USDA, DOE, NASA …
  - NSF Computer & Information Science & Engineering (CISE) Advisory Committee

# Historic center firsts

- 1st of 5 U.S. supercomputing centers
- 1st IBM SP supercomputer deployment
- Largest Web server in the world for Olympics
- 1st Dell Top500 supercomputer deployment
- IDC Innovation Award for hepatitis C research
- $5-million NSF award to build federated cloud model (Aristotle Cloud Federation)

*Early Center success led to NYS grant that built Rhodes*

ARISTOTLE
CLOUD FEDERATION

- NSF CC*DNI DIBBs project (2015-2020)
  - PI Cornell (Lifka); co-PIs U. at Buffalo (Furlani) and UCSB (Wolski)
  - Red Hat, Amazon, Globus, Dell, HPE, Ace Computers
- Federated cloud model goals
  - Share resources between institutions
  - Transparent instance movement between institutions
  - Cross-institution allocations
  - Burst to public cloud during peak usage
  - Open XDMoD cloud accounting & metrics
    - e.g., DrAFTS spot instances bid value predictor

*Reduce time to science*
federatedcloud.org

# Aristotle science use cases

- Developed webGlobe interface to visualize & analyze geospatial data (Chandola)
  - To model changes in future climate simulations using machine learning algorithms (video)
- Computing market efficiency on intraday data for all U.S. stocks (Roesch)
  - To enable policymakers to assess the impact of high-frequency trading on markets
- Mapped transcriptome data to metabolic models of gut microbiota for xylem-feeding insects (Douglas)
  - To exploit animal microbe symbiosis to produce pest-resistant crops (video)
- Used IoT, edge computing & Aristotle for irrigation scheduling (McCurdy/Wolski)
  - To create SmartFarms (e.g., 66% less water used at model vineyard)
  - Also, using Jetstream and Centaurus K-means as a service to rapidly analyze soil electrical conductivity data that indicates health of soil
- Evaluating wind speed variability and other atmospheric properties using Docker-installed WRF physics (Pryor)
  - To enable industry to reduce wind energy project risks and cost

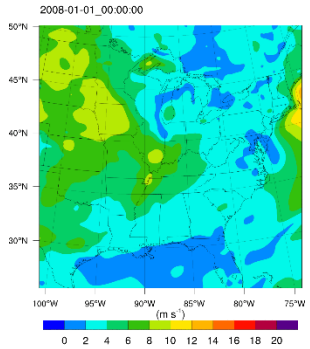


*Pryor visualization*

# Other recent Cornell awards

- NSF CC*DNI DIBBs supplemental awards
  - Organized/Chaired first NSF DIBBs PI Workshop (67-page report)
  - 6 REU students made meaning contributions to Aristotle use cases
  - Transitioning Aristotle to OpenStack; investigating container technologies
- NSF CC*Networking Infrastructure award to upgrade research connectivity
- TACC Stampede2 and IU Jetstream subawards
  - To develop national training (Cornell Virtual Workshops)
- NSF NANOGrav Physics Frontiers Center (co-lead CyberI)
- NSF XSEDE 2.0 (U.S. cyberinfrastructure program)
  - Manager for Community Resource Integration; Community Infrastructure Lead; Training Lead; XSEDE Campus Integration Software


# Staff expertise

- HPC
- Cloud computing
- Data
- Disaster recovery
- IT audit
- Software architecture and development
- Web services and web hosting for research applications
- Visualization
- Training, education & outreach (including eCornell courses)

# Collaborative research support

- Simulating carbon cycling enzymes (Aristilde) case study
- Investigating multi-core technologies for CMS trigger (Wittich) seminar
- Discovering complex material properties (Fennie) video
- Distributing tools to design TAL effectors (Bogdanove) tools
- Searching for gravitational waves in pulsar timing (Cordes) NANOGrav PFC
- Predicting turbulent reacting multiphase flows (Desjardins)
- Developing water management modeling tools (Reed)
- Neuro modeling of sentence comprehension (Hale)

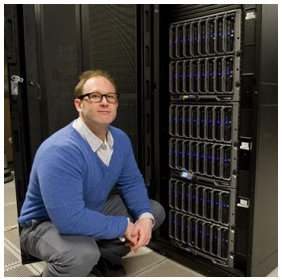

*Fennie cluster*

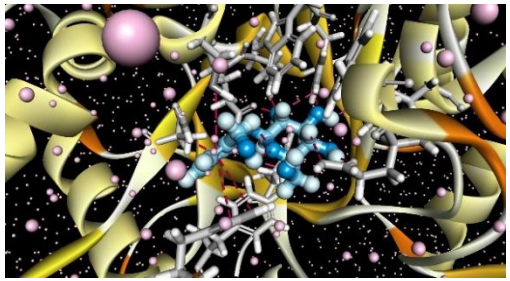

*Aristilde visualization*

# New services

- MATLAB Distributed Computing Server (MDCS) on Red Cloud
  - Red Cloud subscriptions available
  - Researchers must be familiar with MATLAB Parallel Computing Toolbox client to use MDCS
  - MDCS on Red Cloud tutorial available
  - Fee-based assistance if needed
- Private/hybrid clouds for universities and industry
  - Architecture design consulting (including HIPAA systems)
  - Testbed development
  - Deployment

red cloud

Cornell Tech

# New York City Campus

# Cornell Tech mission, focus

- Create pioneering leaders, technologies for the digital age
  - Focus on research and technology commercialization
- Current NYC campus
  - Serving academic community of 600
  - Masters in CS and ORIE, MBAs, PhDs
- When fully completed
  - 2,000 grad students
  - 280 faculty and staff
- For more information
  - "This is Cornell Tech" video

*The Bridge at Cornell Tech: home to start-ups and companies collaborating with Cornell Tech students*

Thank you for visiting!

# Cornell University

*CAC is supported by the Office of the Vice Provost for Research*