

Request for Technology Fee Funds: FY18

NOTE: A separate request should be made for each initiative.

I. Department Number/Department Name:

360 College of Computing

Title of Request (please be brief):

IBM POWER-Based Systems for Data Analysis and HPC Courses

Amount of Request (formula from detailed budget below):

\$153,216

Are there any installation/renovation costs associated with this request?

Yes No

If "Yes" then indicate the source of approved funding:

(Note: Tech Fees are not allowed for installation/renovation)

Executive Summary of Request (100 words or less):

Georgia Tech is at the forefront in research, merging data analysis and high-performance computing with applications ranging from healthcare to national defense. We need to transition this research into the educational program using tools and platforms that will place our students at the same forefront. We propose providing an educationally-oriented system important to Georgia Tech's research directions to ease this transition from research to education. The proposed system provides many aspects otherwise inaccessible to Georgia Tech students

Specific class and/or lab initiative(s) if applicable:

Contact person for this request (incl. phone #):

Jason Riedy, David Bader, Will Powell (4-9301)

Indicate priority per department if applicable:

Number _____ of _____

Indicate priority per college or unit:

Number 4 of 9

II. Impact on Students - Provide course title, course number, and anticipated enrollments:

Titles/Numbers of Course(s)

see attached listing

Anticipated Enrollments

Graduate:	1,918	<input type="radio"/> per sem	<input checked="" type="radio"/> per year
Undergraduate:		<input type="radio"/> per sem	<input checked="" type="radio"/> per year
Total:	1,918	(select one)	

NOTE: Other impacts on students should be described in narrative.

III. Narrative - Provide narrative justification for your intended use of the technology fee funds. Include narrative on how the education or research of the students will be enhanced. Also include how the request aligns with the Strategic Plan of Georgia Tech. Continue in the block below if necessary.

Georgia Tech is at the forefront in research, merging data analysis and high-performance computing with applications ranging from healthcare to national defense. We need to transition this research into the educational program using tools and platforms that will place our students at the same forefront. We propose providing an educationally-oriented system important to Georgia Tech's research directions to ease this transition from research to education. The proposed system provides many aspects otherwise inaccessible to Georgia Tech students.

There are many significant core research areas where Georgia Tech can transition its leading research into a leading education, but for space reasons, we focus on four areas closely related to offered courses:

- using data to drive the materials genome initiative,
- massive data analysis for healthcare,
- bioinformatics for genomic and meta-genomic analysis, and

IV. Detailed Budget - Requested Items by Category List separately any equipment, software, and other allowable expenses (see Tech Fee Guidelines). There is a formula in the "total column" that multiplies the number of items times the unit price. You may enter a figure into the total column if the unit pricing is not applicable. If you need additional rows, contact the Budget Office to receive a modified form.

Supporting documentation is required - Include price justification in some form, such as quotations, published price lists, etc. as a separate PDF attachment. All supporting information should be in a single PDF.

Proposed Number of Items	Estimated Price per Unit	Total (\$)
IBM Power System S822LC with GPU technology	3	\$43,200
Mellanox SB78xx EDR 100Gb/s switch and cabling	1	\$23,616
Total (linked to the total amount of request line above)		\$153,216

Please return form via e-mail in Excel format to: techfees@business.gatech.edu. Supporting information only in a PDF file.

III. Continuation of narrative justification, if necessary

- advanced high-performance computing in support of science and security.

There are many other important and nationally-recognized areas of expertise, but this system has been designed to support these in both need and scale.

The NSF Integrative Graduate Education and Research Traineeship (IGERT) in materials, Flamel, applies machine learning techniques for its students and researchers to ensure that the national materials genome initiative is based soundly in data. Similar machine learning techniques appear in healthcare. Georgia Tech, in conjunction with partners at Children's Healthcare of Atlanta and others, has developed many innovative "explainable" deep learning methods usable in clinical practice. Genomic and meta-genomic analysis translates the massive data generated from "next-generation" sequencers (100s of GB per sample) into relationships between variations and diseases, resistances, and drug interactions. This requires both advanced machine learning (where the NVIDIA accelerators shine) as well as traditional high-performance computing (where the POWER processor shines). Upcoming nation-scale high-performance computing resources are being built around the same technologies as in this proposal. The Department of Energy CORAL system combines POWER processors, NVIDIA GPUs, and Mellanox Infiniband; the proposed system is a result of that program. Georgia Tech has lead applications of such accelerated systems in areas like computational chemistry and cybersecurity, often being the principal component of related industry and federal centers.

In each of these areas, the necessary hardware / execution platforms are available only to groups of researchers and graduate students related to the projects, with occasional special-topics classes having special access. This proposal makes one platform relevant across these and other areas accessible for coursework beyond these specific groups. Another goal is encouraging students (and researchers) to make their tools available for non-experts on the same leading-edge platform, enabling many other students to gain experience. Carefully allocated resource sharing will permit schools and colleges that are not experts in computing to decide resource acquisition based on experience and not press releases while also giving students the experiences that make these students highly desirable.

Background from courses on the proposed systems will expose students to cutting-edge techniques, and the top students (both graduate and undergraduate) will create techniques to advance these new fields. Having experience on this advanced system places students in the top tier for industrial, research, and educational careers. The students also will have relevant and practical background in topics and tools needed by research programs in core Georgia Tech research areas, increasing the pool of data analysis talent for the entire campus.

* Impacted Courses

Note that the data analysis curriculum in the College of Computing is under continuing evolution. Many courses are increasing their use of massive-scale data analysis. The proposed systems can support a variety of one-off class assignments as well as semester-long projects. Regularly offered courses initially targeted by these servers include the following:

- CSE 6140 (CSE Algorithms)
- CSE 6220 (High Performance Computing)
- CSE 6230 (High Performance Computing: Tools and Applications)
- CSE 6643 / MATH 6643 (Numerical Linear Algebra)
- CSE 6730 / CS 4335 (Modeling and Simulation: Fundamentals & Implementation)
- CSE 8803 / CX 4803 (Massive Graph Analysis)
- CX 4242 / CSE 6242 (Data & Visual Analytics)
- CSE 6740 / ISYE 6740 (Computational Data Analysis)
- CSE 6730 (Modeling and Simulation: Fundamentals and Implementation)
- CSE 6250 (Big Data Analytics in Healthcare)

Class projects from many of these classes, particularly CSE 6230 and CX 4242 / CSE 6242, will be encouraged to produce high-performance deep-learning tools accessible to other undergraduate and graduate courses across the campus. The systems will be available through an existing queue system to permit sharing resources without impacting class requirements.

Beyond regularly offered courses, these irregularly offered special topics courses (CSE 8803) are relevant:

- Big Data Analytics in Healthcare (now a regularly offered course): This course applies Georgia Tech's leading research in explainable deep learning methods to healthcare records.
- Massive Graph Analysis: Georgia Tech is a leader in applying NVIDIA GPUs and IBM POWER processors to graph analysis problems. Students will be able to produce cutting-edge results in class, and that is a powerful motivator.
- Advanced Scientific Computing: The proposed system is similar to nodes in the upcoming Department of Energy CORAL next-generation supercomputer, and this course will expose students to real-world supercomputing resources without competing for resources on the nation-