Advancing Computational Research and Education

For Presentation at HPRC Annual User Meeting

by
Dr. Honggao Liu, Director of HPRC

4/6/2016

Outline

1. Overview HPRC and existing resources
2. Updates on HPRC systems status and x86 platform expansion
3. New HPRC Resource Allocation Committee and policy
4. Plans for growing HPRC at TAMU

HPRC

- **Supports and maintains 5 HPC clusters over 800 TFlops (trillion of floating point operations per second) total computing capacity and over 9 PB storage**
- **Services provided:**
  - User Services: help desk, training, workshops, software support, user support for over 1000 users
  - Systems: system admin functions (Linux), storage (GPFS and NFS), support servers, security, and other cyberinfrastructure support
- **Staff consists of 10 FTE's and 2 student workers.**
- **HPC staffing funds from DOR at TAMU**
- **Access: open to faculty, research staff, collaborators and students.**
- **No charge or usage limitations.**

HPRC Resources

System	SUs	Peak FLOPS
Ada	151,898,400	337 TF
Neumann	287,047,680	419 TF
Curie	6,727,680	26 TF
Crick	3,223,680	12 TF
Eos	27,050,880	35 TF
Lonestar 5	14,000,000 available for TAMU	1.2 PF
TOTAL:	489,948,320	829 TF (TAMU) 1.2 PF (TACC)

Storage	Description	Total Storage
Ada Storage	4 PB raw of IBM GSS storage running GPFS. Shared between Ada and Curie clusters.	4 PB
Crick Storage	22 data nodes each with 14 TB raw storage running GPFS-FPO.	0.316 PB
Tiered storage	3.5 TB total raw storage (10 TB flash storage, 1 PB disk storage, 2.5 PB tape storage) running GFPS.	3.5 PB
Neumann storage	2 PB raw of IBM GSS storage running GPFS.	2 PB

Ada: A 17,340-core IBM x86 Cluster

A 17,340-core system, with 845 20-core nodes equipped with the INTEL 10-core 2.5GHz IvyBridge processor. The other 15 nodes are 1TB and 2TB memory, 4-processor SMPs configured with the INTEL 10-core 2.7GHz Westmere processor. 30 nodes have 2 GPUs each and 9 have 2 Phi coprocessors. The interconnecting fabric is based on the Mellanox SX6536 IB core switch.

Crick: 25-node POWER7+ cluster

For Big Data Analytics Cluster w Storage-rich Nodes
IBM 7R2 Node Characteristics
• 16-core (two 8-core Power7+ processors) 4.2GHz; 537GFLOPS; 32 nm fabrication
• 256 GB memory; 68GB/s per processor socket; Shared (on-chip) L3 80MB per processor
• 10GbE port (link to Wehner Core 10G/40G)
• 4 x 600GB 10K rpm SAS (local disk)
• EXP24S + 24 x 600GB (14.4TB) 10K rpm SAS drives
• REDHAT ENTERPRISE LNX FOR PWR; GPFS Client
• Special Software: IBM InfoSphere BigInsights, Data Explorer Resource; Query Routing for InfoSphere Data Explorer; etc
• Target Areas: Big Data Analytics; Genomic Analysis, Breeding Simulation, mining historical data; Map Reduce; etc

Curie: 50-node POWER7+ cluster

For Special HPC Codes & Applications

IBM 7R2 Node Characteristics

- 16-core (2 8-core Power7+ processors) 4.2GHz; 537GFLOPS; 32 nm fabrication
- 256 GB memory; 68GB/s per processor socket; L3 80MB/ processor
- 10GbE port (link to Wehner Core 10G/40G switch)
- 4 x 600GB 10K rpm SAS (local disk)
- REDHAT ENTERPRISE LINUX FOR PWR ; GPFS Client; IBM Compilers, ESSL, LSF (batch)
- Target areas: applications & codes requiring fast memory and fast cpus

Neumann: 2048-node BlueGene/Q Cluster

32768-core, 419 TFLOPS, Massively Parallel, interconnected in 5D Torus topology

Node: 16 compute 1.6 GHz PowerA2 cores;
1 core for system calls; 1 core for spare
16 GB of DDR3 memory
204.8 GFLOPS (1.6 GHz * 8 Flop/s * 16)

1 I/O drawer per rack: 8 I/O nodes with 1 QDR IB link (via the torus interconnect) for every 2 comp-ute nodes. Smallest job is 128-way

Shared Mass Storage 2: 2 PB of GSS26 con-nects to the I/O drawers via 16 IB connections thru 2 QDR/FDR10 IB switches

4 Front End Nodes: 4-core 3.6 GHz POWER7+ processor, 32 GB RAM, 2 600 GB SAS Drives

OS: Compute Node Kernel (CNK) runs on all compute nodes; a lightweight kernel, similar to Linux, supporting a large subset of its system calls

LoadLeveler (batch); IBM XL and GNU compilers; Engineering & Scientific Library (ESSL)

Mass Storage and Connectivity

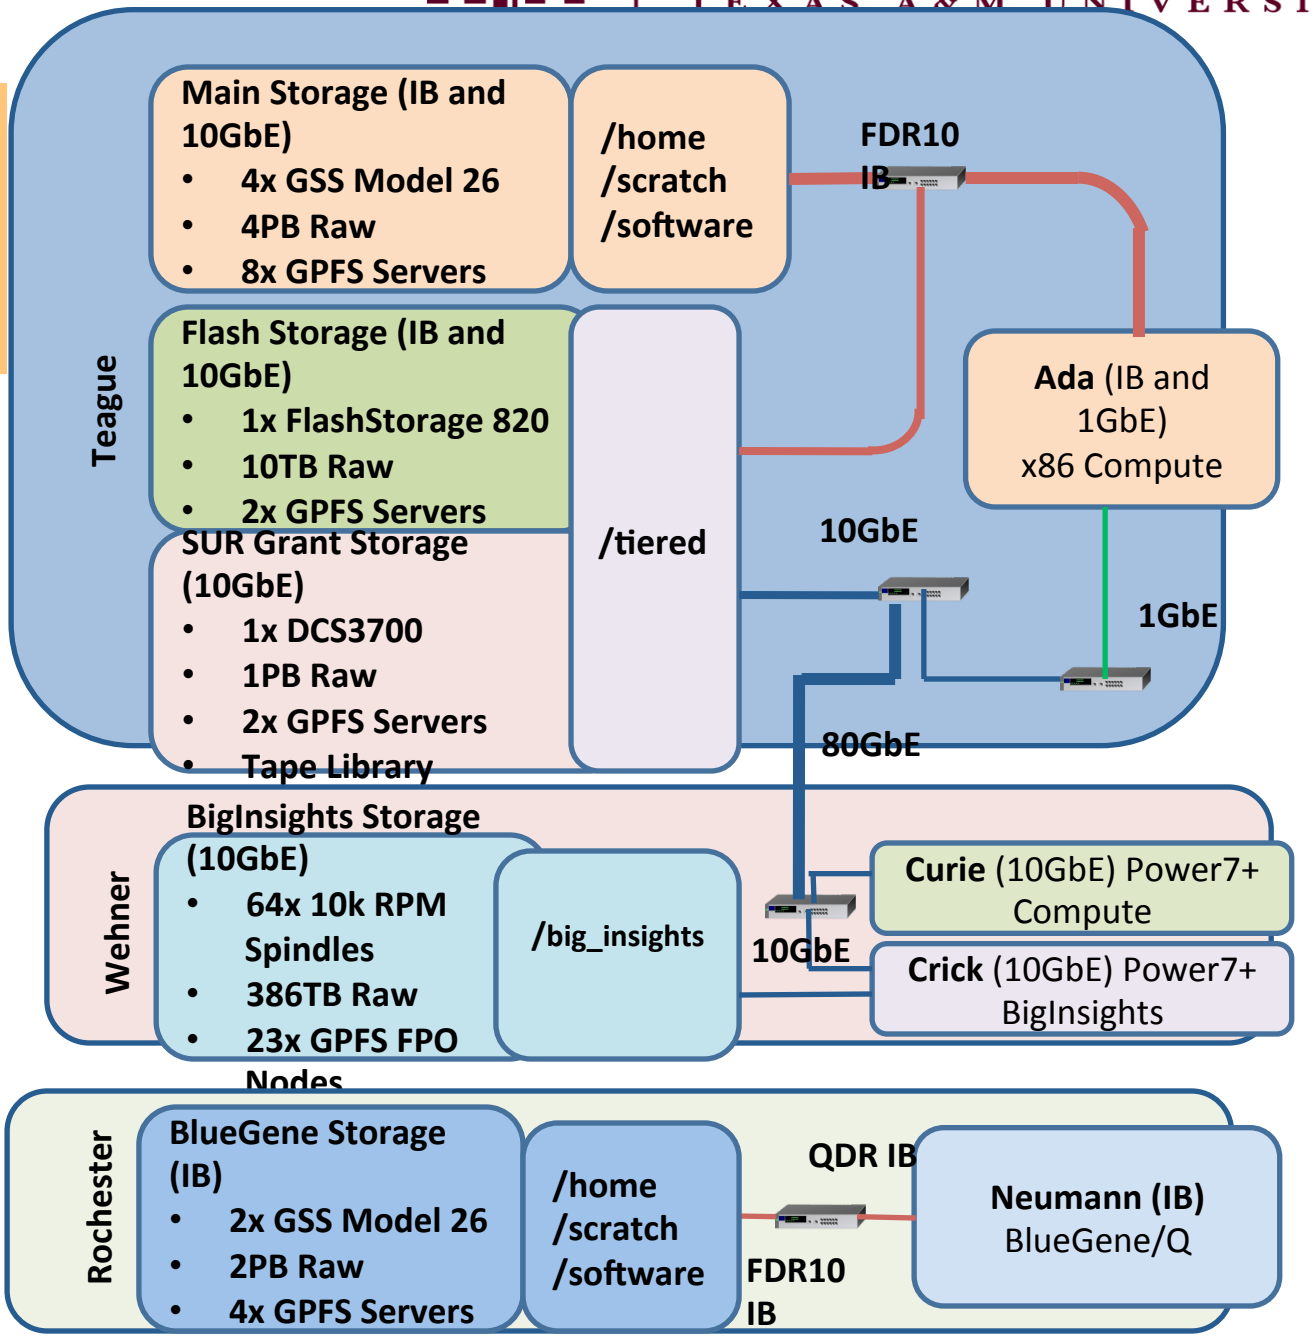

HPRC Users

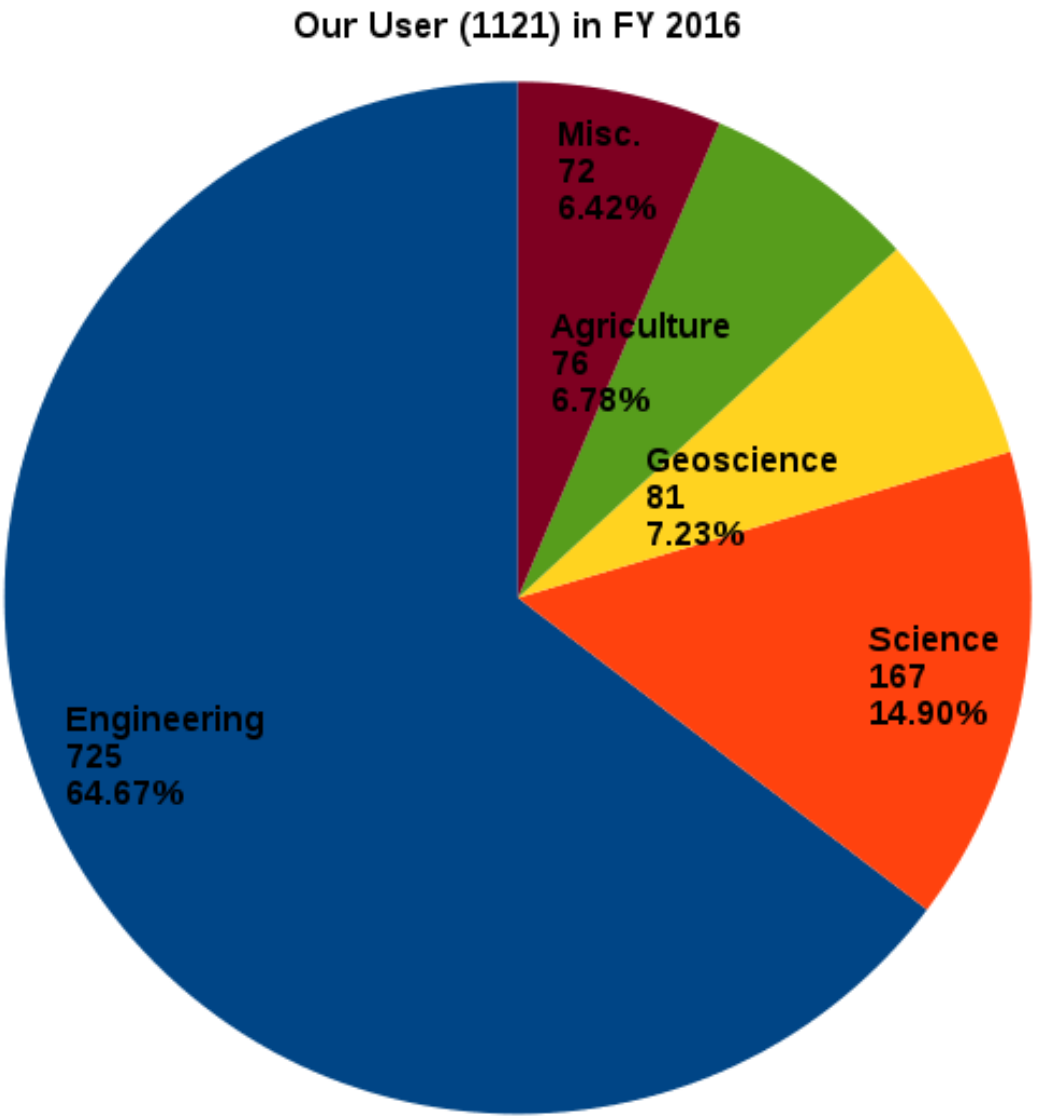
Our User (1121) in FY 2016
Misc.
72
6.42%
Agriculture
76
6.78%
Geoscience
81
7.23%
Science
167
14.90%
Engineering
725
64.67%

Ada Usage

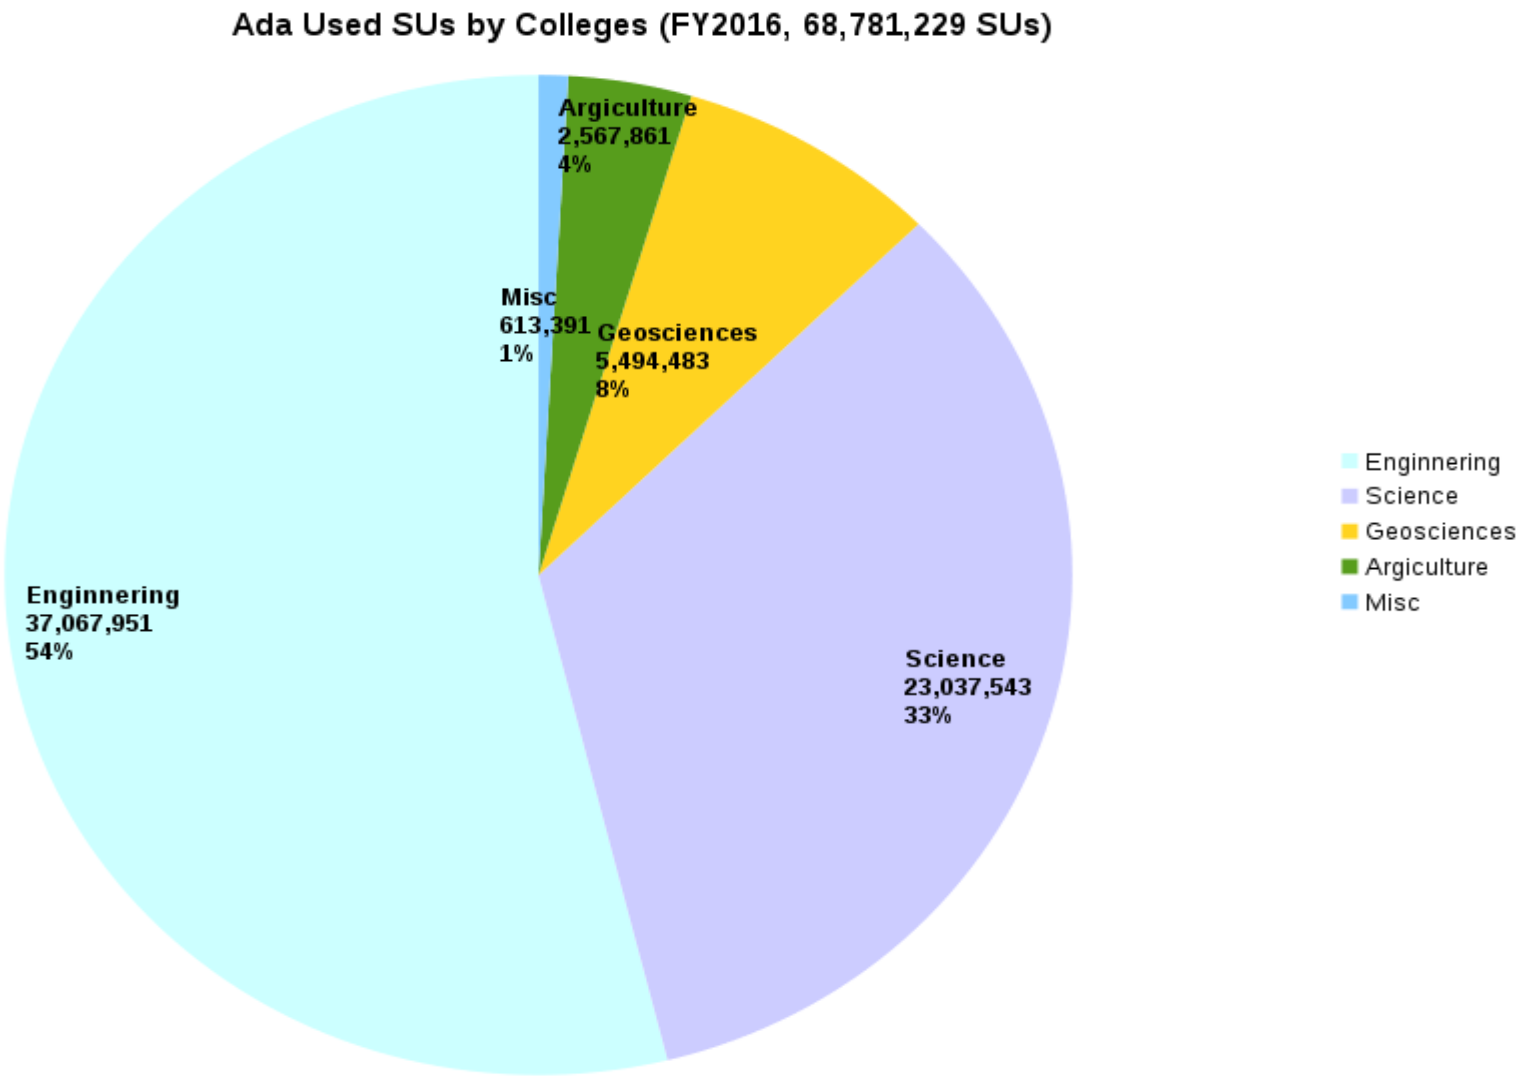
Ada Used SUs by Colleges (FY2016, 68,781,229 SUs)
Argiculture
2,567,861
4%
Misc
613,391
1%
Geosciences
5,494,483
8%
Enginnering
37,067,951
54%
Science
23,037,543
33%
Enginnering
Science
Geosciences
Argiculture
Misc

Eos Usage

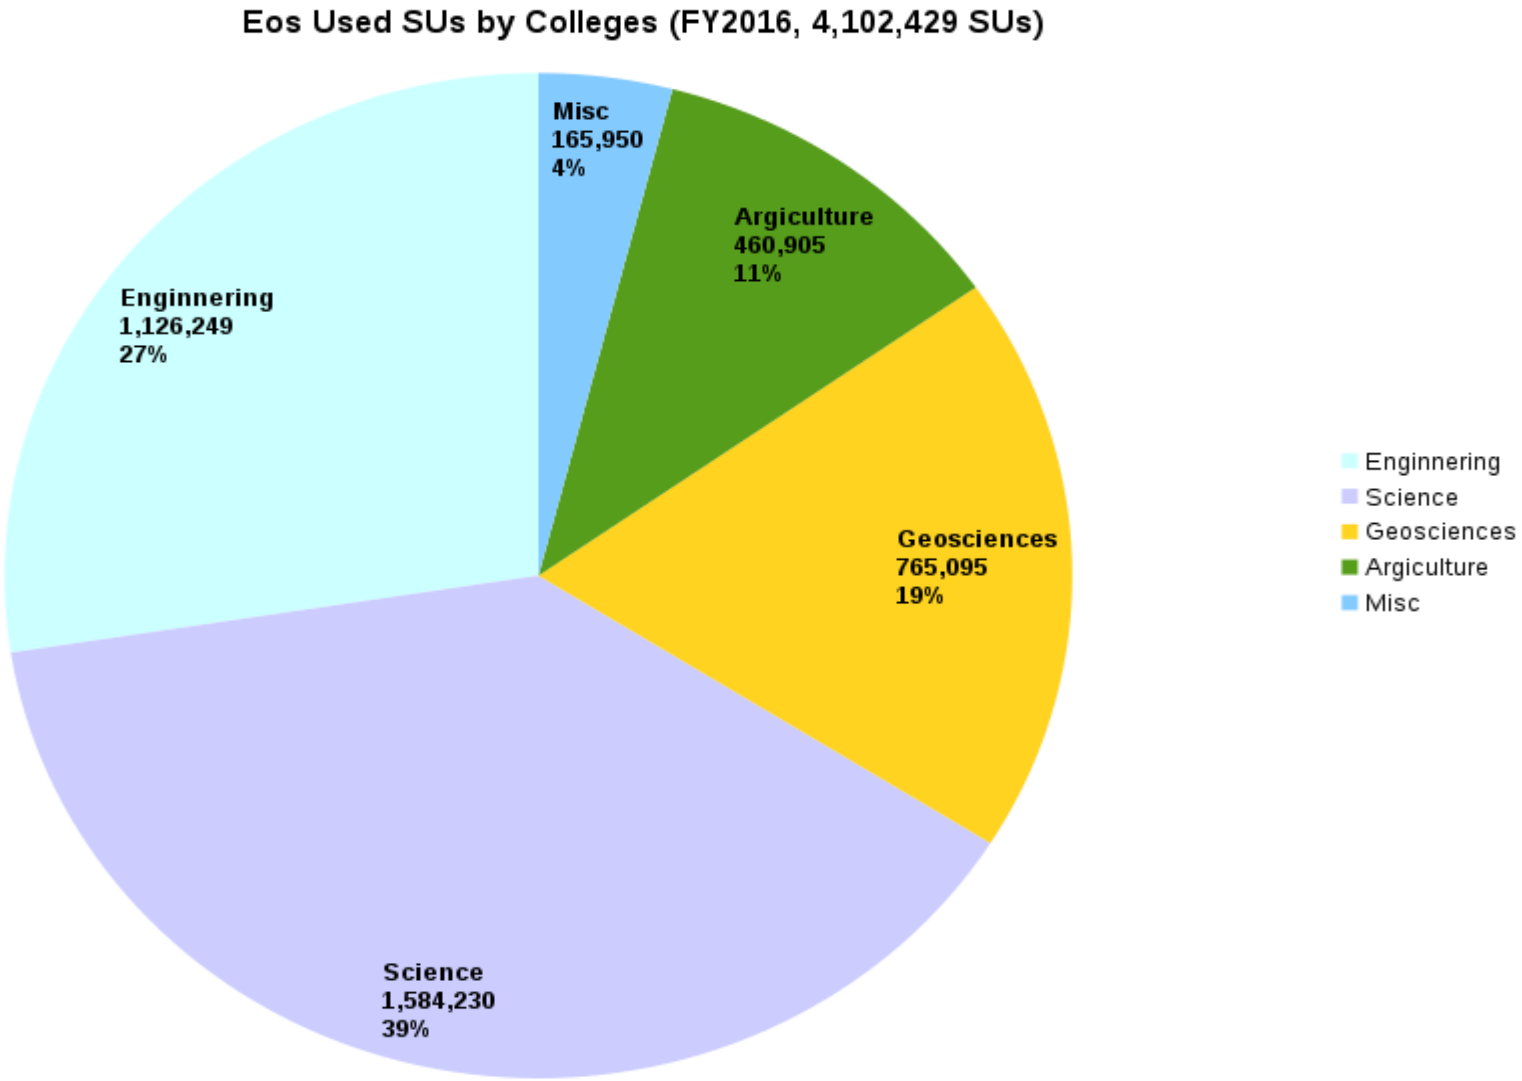
Eos Used SUs by Colleges (FY2016, 4,102,429 SUs)
Misc
165,950
4%
Argiculture
460,905
11%
Enginnering
1,126,249
27%
Geosciences
765,095
19%
Science
1,584,230
39%
Enginnering
Science
Geosciences
Argiculture
Misc

HPRC Resources Status

System	SUs	Peak FLOPS
Ada	151,898,400	337 TF
Neumann	287,047,680	419 TF
Curie	6,727,680	26 TF
Crick	3,223,680	12 TF
Eos	27,050,880	35 TF
Lonestar 5	14,000,000 available for TAMU	1.2 PF
TOTAL:	489,948,320	829 TF (TAMU) 1.2 PF (TACC)

Storage	Description	Total Storage
Ada Storage	4 PB raw of IBM GSS storage running GPFS. Shared between Ada and Curie clusters.	4 PB
Crick Storage	22 data nodes each with 14 TB raw storage running GPFS-FPO.	0.316 PB
Tiered storage	3.5 TB total raw storage (10 TB flash storage, 1 PB disk storage, 2.5 PB tape storage) running GFPS.	3.5 PB
Neumann storage	2 PB raw of IBM GSS storage running GPFS.	2 PB

HPRC Resource Allocation Committee (HPRC-RAC)

Name	Department	Contact Email
James C. Hu (Chair)	Biochemistry and Biophysics	jimhu@tamu.edu
Wolfgang Bangerth	Mathematics	bangerth@tamu.edu
Christodoulos A. Floudas	Chemical Engineering	floudas@tamu.edu
Ping Chang	Oceanography	ping@tamu.edu
Raymundo Arroyave	Materials Science and Engineering	rarroyave@tamu.edu

Allocation Category

Allocation Category	Percentage of Common Pools	Allocation Authority
Program Development	10%	VPR
Discretionary	10%	HPRC Director
Basic Allocation (5,000 SUs)	up to 80%	HPRC Staff
Startup Allocation (5,000~200,000 SUs)	up to 80%	HPRC Director
Research Allocation (> 200,000 SUs)	up to 80%	HPRC-RAC

Allocation Classes

Class	Description
Program Development	A proposal to assist with the adoption of high performance research computing as part of the research enterprise. Awarded by the VPR from the Program Development resources.
Discretionary	A proposal determined to be of value, but lying outside the normal allocation process. Awarded by the HPRC Director from the Discretionary resources.
Basic	Every user account, upon creation, receives a Basic allocation. This allows sufficient time for conventional processing on the systems, and developing information for potential formal proposals. These allocations are limited to 5,000 SUs. Basic allocations are not renewable.
Startup	A proposal requesting an allocation for the purpose of exploring the value of high performance computing for a new project. Only 2 startup allocations may be active at any time per PI.
Research	A request for a large amount of time for a significant research project. May be awarded by simple majority agreement of the HPRC-RAC. Research requests are limited to 3 million SUs, and a PI may not receive more than 5 million SUs from the common pool through Research Allocations during any given academic year to ensure equitable distribution of resources among all PIs.

Ada Allocation

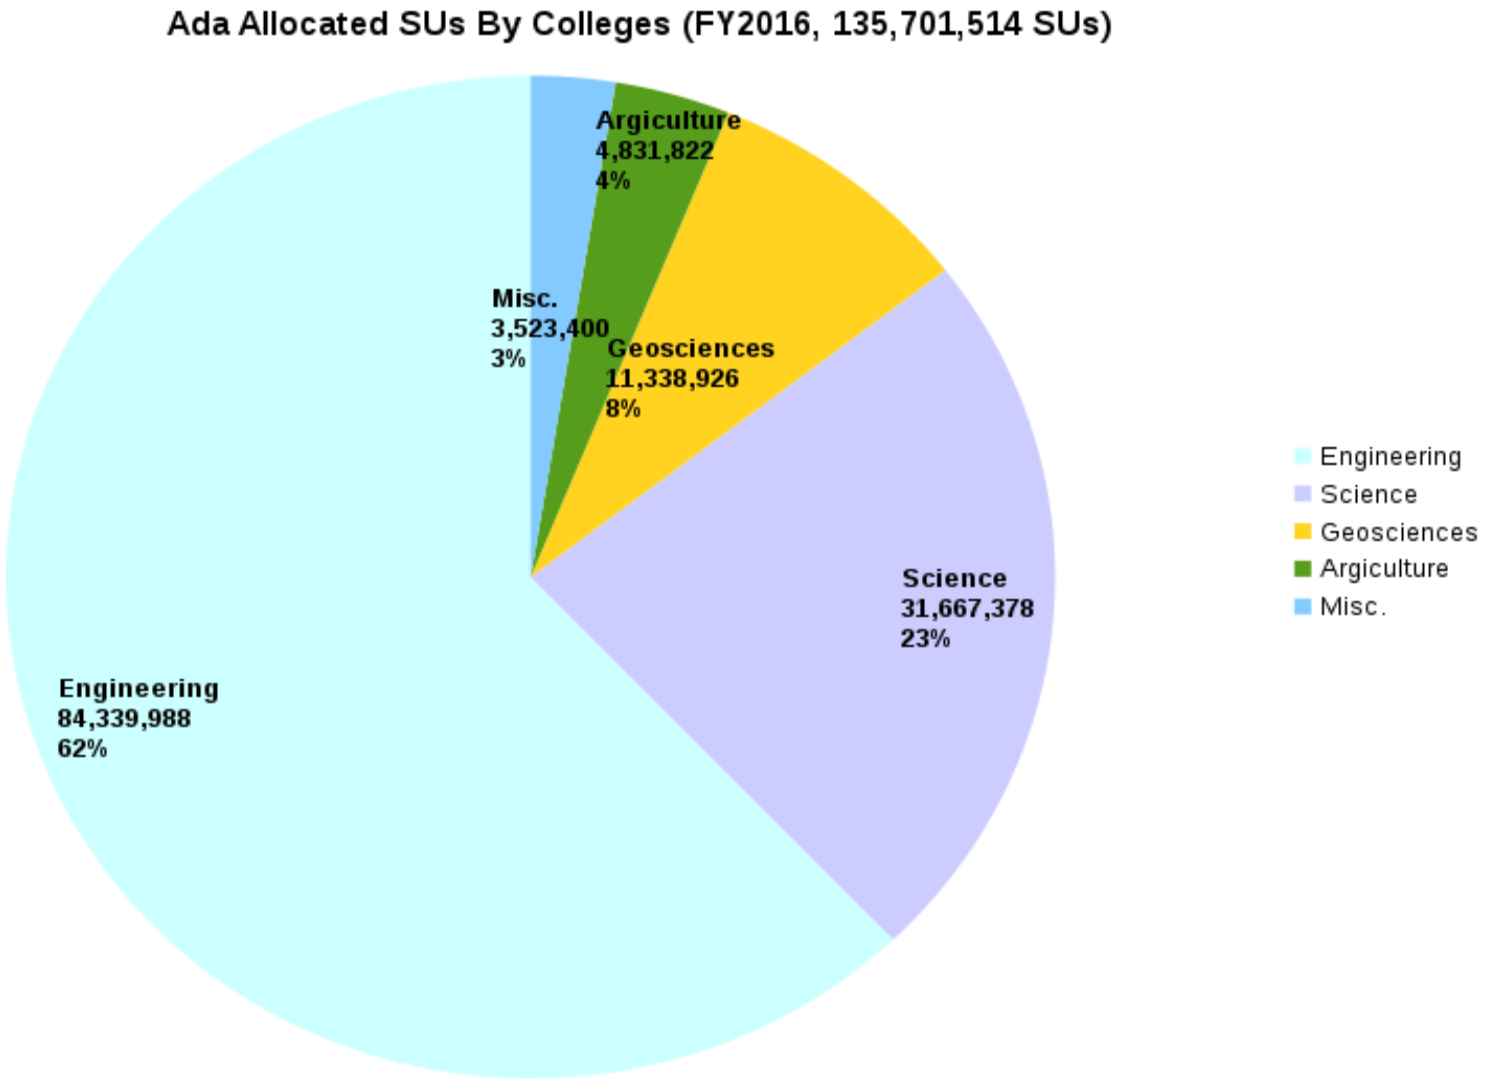

Eos Allocation

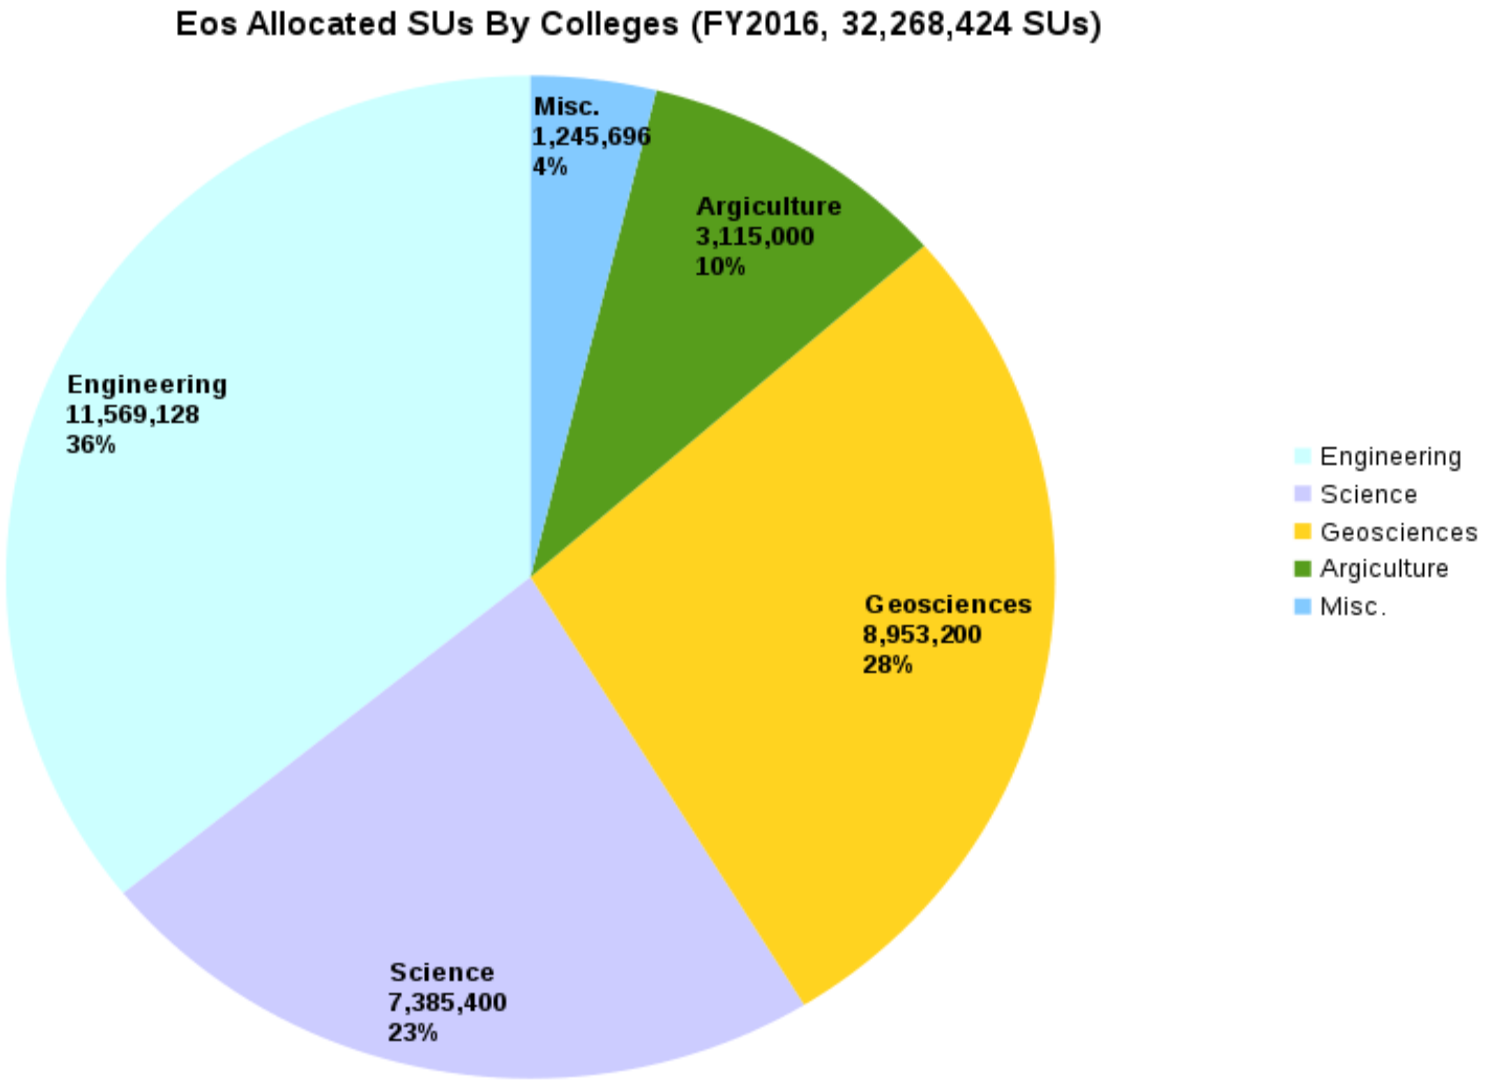

Plans for Growing HPRC at TAMU

- **Increase support effort to help more people in more disciplines use HPC/Cyberinfrastructure to advance their research**
  - **Add applications (research scientists) staff to support and research on a broader range of compute-intensive and data-intensive projects**
  - **Reach out to non-traditional HPC user communities in the areas of arts, humanities, social sciences and biosciences.**
- **Explore grant and research opportunities and participate potential collaborations with faculty and researchers**
  - **Encourage HPRC staff to be involved in more research projects and participate grant proposals writing**
  - **Encourage Faculty to include HPRC staff time in their proposals**
  - **Facilitate opportunities for interdisciplinary and multidisciplinary networking and collaborations**

Plans for Growing HPRC at TAMU

- **Help people move from using campus resources to using nationally available resources**
  - **Aggressively participate XSEDE programs such as Campus/ Regional Champions, Service Provider Forum.**
  - **Work with big HPC users to submit successful allocation proposals**
- **Support and advance computational education across the university**
  - **Increase EOT activities such as offer more workshops, camps, tutorials.**
  - **Encourage HPRC staff to teach courses and pursue adjunct research faculty with academic departments**

Plans for Growing HPRC at TAMU

- **Attract faculty to be affiliated with HPRC**
  - **Provide an open environment and platform for multidisciplinary and interdisciplinary collaborations**
  - **Offer dedicated or priority access to HPRC resources**
  - **Provide some dedicated HPRC staff time supporting faculty's research projects**
  - **Provide 100% of project credits to faculty's home department**

Plans for Growing HPRC at TAMU

- **Establish life-cycle replacement funding to sustain HPRC environment**
  - **Provide a broad spectrum of advanced computing resources to campus faculty-researchers**
  - **Sustain leading edge cyberinfrastructure and support researchers to advance TAMU research**
- **Increase collaborations and partnership with other HPC or research computing centers**

HPRC

An interdisciplinary research environment that advances computational and data-enabled sciences and engineering, and the disciplines we touch.

- **HPRC was recognized as an Intel Parallel Computing Center**
- **NVIDIA GPU Programming Workshop on April 27-28**

Thank You!

Q&A