HPC & BD Survey
Description and Results

Table of Contents

About this document

This document contains results of the HPC & BD survey, presented during the 2018 spring semester. Additionally, this document contains descriptions of the survey structure, including the explanatory information that is provided within the survey as guidance to participants. This within-survey explanatory information appears throughout this document in text using this font.

Introduction to the survey

The HPC & BD survey gathers information on multiple topics. Participants choose any or all of six topic sections, each of which contains questions tailored to individual participants' interests, expertise, and usage of HPC resources. Larger topic sections contain selectable subsections. All topic sections solicit, but do not require, unstructured comments.

Because individual participants choose among various topic sections, specific content of the survey varies among participants. However, some parts of the survey are consistent: All participants see initial explanatory information on the topic choices, all participants provide information such as their NJIT status, department affiliations, etc, and all conclude the survey with general assessments of the services provided by IST-managed HPC and/or BD resources.

About the results

Participant responses to all survey questions are presented here objectively, that is, without conclusions, observation of trends, etc.; judgment of import is left to the reader. However, finer-grained and targeted analyses, though beyond the scope of this document, can be drawn from the raw data: In particular, responses can be grouped by participant characteristics such as academic department, research interests, HPC usage, etc., allowing for comparisons between groups.

Please send any questions about the survey results, as well as any requests for finer-grained, targeted, or other specialized analyses to: arcs@njit.edu.

Survey participants

58 invitees participated. 56 of these completed the survey, and 2 (one doctoral student and one post-doc) completed several topic sections, but did not take the concluding general-assessment section. An additional 5 invitees who were not direct users of HPC opted to quit the survey.

Participants' NJIT status
The 58 participants consisted of 23 faculty (14 tenured, 8 tenure-track, 1 non-tenure-track), 6 Post-docs, 27 students (24 doctoral, 2 masters, 1 undergrad), and 2 staff (research associate, guest researcher).

Participants' academic departments
Participants' departmental affiliations are shown in the pie chart below.
Members of departments of Mathematical science and Computer science predominated, together accounting for 45% of the 58 participants. The 11 *OTHER* participants consisted of 3 from Informatics, 2 each from Chemistry and environmental sciences, Civil and environmental engineering, and Federated biological sciences, and 1 each from Biomedical engineering and Management.

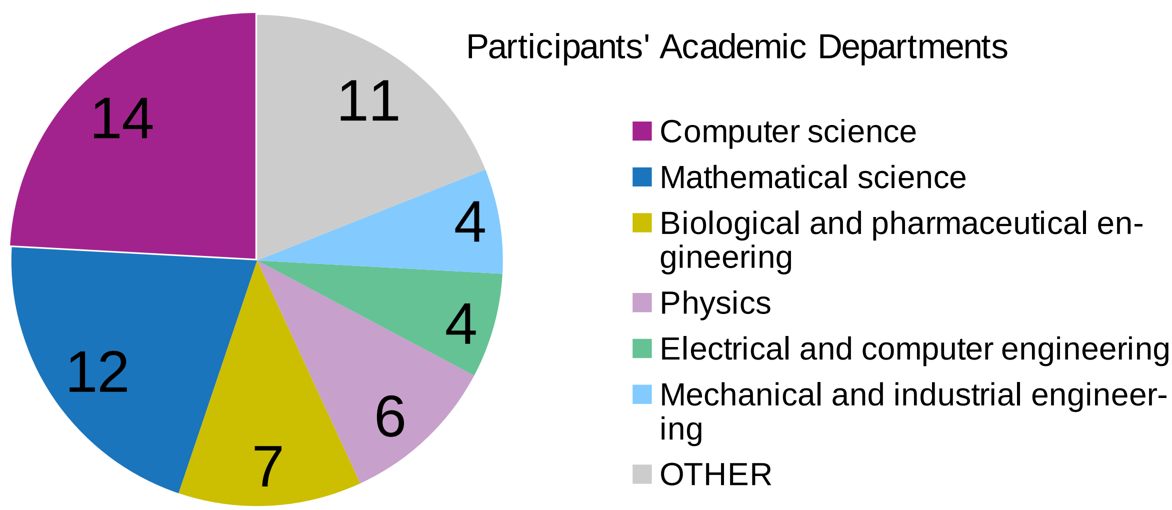

31 participants (53%) reported using IST-managed HPC and/or BD for over 2 years; 11 (19%) for 1 to 2 years; and 16 (28%) for less than a year.

Topic sections: descriptions

Brief topic section descriptions (shown below) were provided.

Section 1 covers High Performance Computing (HPC) hardware resources

Clusters: Kong (general-access (GA)) and Stheno (Dept Mathematical Sciences (DMS)). Both parallel and serial computations can be done on these clusters.

GPU (graphical processing unit) nodes: These nodes contain both CPUs and GPUs. Both parallel and serial computations can be done on the CPUs. The GPUs are suitable only for parallel computations. Both Kong and Stheno contain GPU nodes.

Shared memory machines: Kong "smp" queue (GA), Cnrdp (Center for Natural Resources and Protection), Gorgon (DMS), Phi (GA). Both parallel and serial computations can be done on shared memory machines.

Section 2 covers HPC storage

AFS distributed filesystem: General computational use; accessible from any AFS client computer, including all cluster compute nodes

NFS distributed filesystem: General computational use; accessible only from the cluster on which it is mounted

Parallel file system (PFS) appliance for temporary files

Storage costs

Section 3 covers Big data (BD) computational and storage resources

Hadoop/Spark infrastructure

Storage costs

Section 4 covers Software environment

Open source and commercial applications; libraries; utilities; modules (used for setting the user's environment for a specific application)

Section 5 covers Internet bandwidth and Science DMZ

Internet bandwidth:

This the capacity, including Internet 2, of NJIT's connection to and from the Internet

Science DMZ:

This is a computer subnetwork that is structured to be secure, but without the performance limits that would otherwise result from passing data through a stateful firewall.

The Science DMZ is designed to handle high volume data transfers, typical with scientific and high-performance computing, by creating a special DMZ to accommodate those transfers.

Science DMZ is typically deployed at or near the local network perimeter, and is optimized for a moderate number of high-speed flows, rather than for general-purpose business systems or enterprise computing.

Section 6 covers Consultation with Academic and Research Computing Systems (ARCS)

Consultation with Academic and Research Computing Systems (ARCS) staff: installation of compilers, applications, libraries, and utilities; customized scripts to aid users in their use of HPC resources; assistance in debugging and optimizing code; assistance in getting applications to run; Assistance in running parallel code; assistance in working with and managing big data.

Participants' topic choices are shown in the bar graph below.

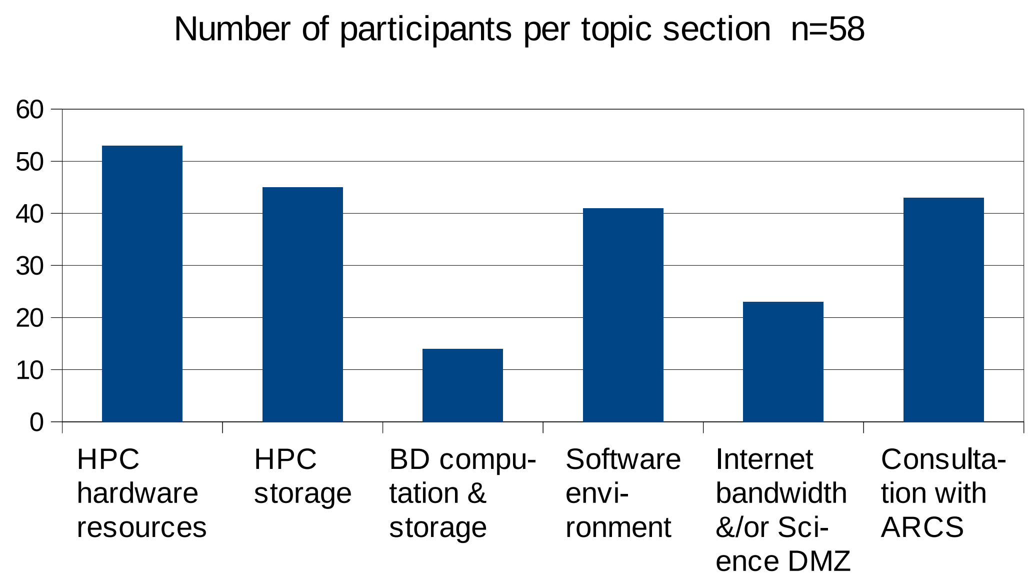

SURVEY RESULTS

Topic Section 1 HPC hardware resources

53 participants (91%)

This topic section contained subsections on **parallel processing** (selected by 65% of the 53 participants), **serial processing** (selected by 73% of the 53 participants), and questions about **HPC hardware resources** relevant to all 53 participants. 37% (19 participants) took both the parallel and serial processing subsections.

In each of the subsections on parallel and serial processing, participants indicated which of the following resources they used to run computations: *Kong (including "smp" queue), Stheno*, *Cnrdp*, *Gorgon*, *Phi*, and, in the parallel subsection, *resources using GPUs*. For each resource chosen, they then rated the adequacy (*Adequate*, *Moderate increase needed*, *Large increase needed*, *Don't know*) of that resource's *number of cores, number of nodes, CPU speed, Max RAM per node*, and (if applicable) *Node interconnect (internal network) speed*. Any participant who rated any aspect as less than adequate then ranked potential remediations for addressing the inadequac(ies). These potential remediations were presented as follows:

> Potential remediation options involve increases in resources that differ in whether they are:
>
> - Paid for by NJIT *versus* paid for by individual researchers
> - On-premise *versus* off-premise
> - Public resources requiring a successful researcher proposal
>
> **Note:** Off-premise providers include *Amazon Web Services*, *Azure*, *Google Cloud Platform*, *IBM Bluemix*, *Oracle Cloud*, *Penguin Computing*

Choices to be ranked: Increase NJIT-provided HPC resources shared among users; Increase user-purchased HPC resources at NJIT; Increase NJIT-provided commercial off-premise resources shared among users; Increase user-purchased commercial off-premise resources; Increase use of publicly available HPC resources (e.g., at a national supercomputing center), successful proposal by researcher is required

Parallel processing

n=35

The number of participants indicating use of each parallel processing resource is shown in the bar graph below.

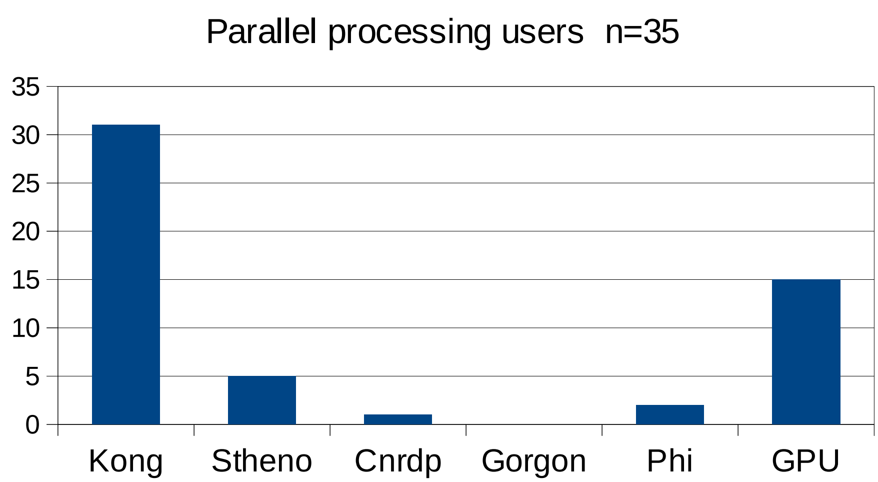

Kong

31 participants (89% of those who use parallel processing) indicated use of Kong. The table below shows the number of participants who provided each particular rating for each aspect.

KONG, parallel n=31	Large increase	Moderate increase	Adequate	Don't know
Number of cores	8	12	9	2
Number of nodes	7	13	9	2
CPU speed	9	12	8	2
Max RAM	12	6	9	4
Node Interconnect speed	8	8	7	8

25 (81%) of the 31 participants who use Kong for parallel processing rated at least one aspect as not adequate. The table below shows the number of participants choosing each ranking option.

KONG, parallel n=25	Rank 1	Rank 2	Rank 3	Rank 4	Rank 5
Increase NJIT-provided HPC resources shared among users	10	7	3	4	1
Increase user-purchased HPC resources	4	3	4	7	7
Increase NJIT-provided commercial off-premise resources shared among users	2	10	3	8	2
Increase user-purchased commercial off-premise resources	1	2	3	6	13
Increase use of publicly available HPC resources (e.g., at a national supercomputing center); successful proposal by researcher is required	8	3	12	0	2

Stheno

5 participants (14% of the 35 participants who use parallel processing) indicated use of Stheno. The table below shows the number of participants who provided each particular rating for each aspect.

STHENO, parallel n=5	Large increase	Moderate increase	Adequate	Don't know
Number of cores	3	0	1	1
Number of nodes	3	0	1	1
CPU speed	3	0	1	1
Max RAM	2	0	1	2
Node Interconnect speed	3	0	1	1

3 (60%) of the 5 participants who use Stheno for parallel processing rated at least one aspect as not adequate. The table below shows the number of participants choosing each ranking option.

Stheno, parallel n=3	Rank 1	Rank 2	Rank 3	Rank 4	Rank 5
Increase NJIT-provided HPC resources shared among users	2	1	0	0	0
Increase user-purchased HPC resources	0	0	1	1	1
Increase NJIT-provided commercial off-premise resources shared among users	0	0	1	1	1
Increase user-purchased commercial off-premise resources	0	1	0	1	1
Increase use of publicly available HPC resources (e.g., at a national supercomputing center); successful proposal by researcher is required	1	1	1	0	0

Cnrdp

Only 1 participant out of 35 (3% of those who use parallel processing) indicated use of Cnrdp. This participant rated all aspects as adequate (note: *Node interconnect speed* was not included in the question), and hence was not queried about remediation rankings.

Gorgon

No participants indicated parallel use of Gorgon.

Phi

2 participants (6% of those who use parallel processing) indicated use of Phi. The table below shows the number of participants who provided each particular rating for each aspect.

PHI, parallel n=2	Large increase	Moderate increase	Adequate	Don't know
Number of cores	2	0	0	0
Number of nodes	2	0	0	0
CPU speed	1	1	0	0
Max RAM	1	1	0	0

Both participants who use Phi for parallel processing rated at least one aspect as not adequate. The table below shows the number of participants choosing each ranking option.

Phi, parallel n=2	Rank 1	Rank 2	Rank 3	Rank 4	Rank 5
Increase NJIT-provided HPC resources shared among users	1	0	1	0	0
Increase user-purchased HPC resources	0	0	1	1	0
Increase NJIT-provided commercial off-premise resources shared among users	0	1	0	1	0
Increase user-purchased commercial off-premise resources	0	0	0	0	2
Increase use of publicly available HPC resources (e.g., at a national supercomputing center); successful proposal by researcher is required	1	1	0	0	0

GPU

15 participants (43% of those who use parallel processing) indicated use of GPUs. The table below shows the number of participants who provided each particular rating for each aspect.

GPU, parallel n=15	Large increase	Moderate increase	Adequate	Don't know
Number of cores	4	6	5	0
Number of nodes	6	5	4	0
CPU speed	1	7	6	1

Max RAM	5	4	5	1

12 (80%) of the 15 participants who use GPUs rated at least one aspect as not adequate. The table below shows the number of participants choosing each ranking option.

GPU, parallel n=12	Rank 1	Rank 2	Rank 3	Rank 4	Rank 5
Increase NJIT-provided HPC resources shared among users	7	1	3	1	0
Increase user-purchased HPC resources	1	1	3	4	3
Increase NJIT-provided commercial off-premise resources shared among users	2	6	2	2	0
Increase user-purchased commercial off-premise resources	0	1	1	2	8
Increase use of publicly available HPC resources (e.g., at a national supercomputing center); successful proposal by researcher is required	2	3	3	3	1

Serial processing

n=38

The number of participants indicating use of each serial processing resource is shown in the bar graph below.

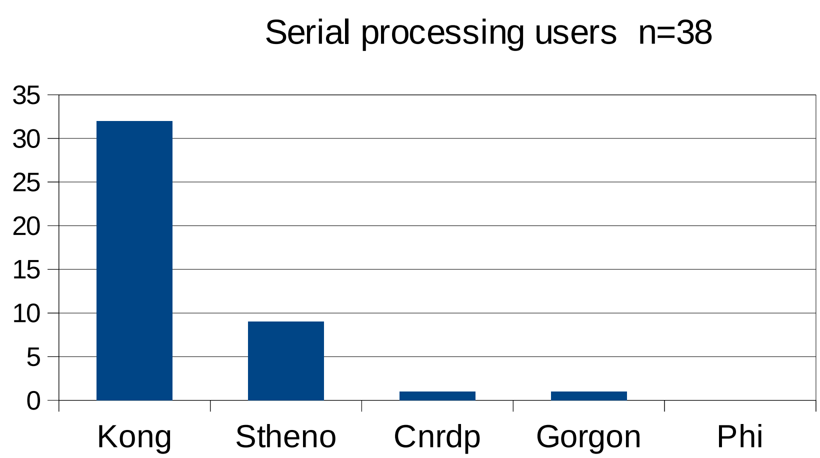

Kong

32 participants (84% of those who use serial processing) indicated use of Kong. The table below shows the number of participants who provided each particular rating for each aspect.

KONG, serial n=32	Large increase	Moderate increase	Adequate	Don't know
Number of cores	7	7	15	3
Number of nodes	7	7	16	2

CPU speed	8	12	10	2
Max RAM	8	5	11	8

22 (69%) of the 32 participants who use Kong for serial processing rated at least one aspect as not adequate. The table below shows the number of participants choosing each ranking option.

KONG, serial n=22	Rank 1	Rank 2	Rank 3	Rank 4	Rank 5
Increase NJIT-provided HPC resources shared among users	10	5	3	2	2
Increase user-purchased HPC resources	3	5	7	4	3
Increase NJIT-provided commercial off-premise resources shared among users	1	7	4	6	4
Increase user-purchased commercial off-premise resources	4	2	1	7	8
Increase use of publicly available HPC resources (e.g., at a national supercomputing center); successful proposal by researcher is required	4	3	7	3	5

Stheno

9 participants (24% of those who use serial processing) indicated use of Stheno. The table below shows the number of participants who provided each particular rating for each aspect.

STHENO n=9	Large increase	Moderate increase	Adequate	Don't know
Number of cores	3	1	5	0
Number of nodes	3	1	5	0
CPU speed	2	5	2	0
Max RAM	1	5	2	1

7 (88%) of the 8 participants who use Stheno for serial processing rated at least one aspect as not adequate. The table below shows the number of participants choosing each ranking option.

STHENO n=7	Rank 1	Rank 2	Rank 3	Rank 4	Rank 5
Increase NJIT-provided HPC resources shared among users	0	4	3	0	0
Increase user-purchased HPC resources	1	0	3	2	1
Increase NJIT-provided commercial off-premise resources shared among users	3	0	0	2	2

Increase user-purchased commercial off-premise resources	2	0	0	3	2
Increase use of publicly available HPC resources (e.g., at a national supercomputing center); successful proposal by researcher is required	1	3	1	0	2

Cnrdp

Only 1 participant out of 38 (3% of those who use serial processing) indicated use of Cnrdp. The table below shows this participant's rating for each aspect.

CNRDP, serial n=1	Large increase	Moderate increase	Adequate	Don't know
Number of cores	0	0	1	0
Number of nodes	0	0	1	0
CPU speed	0	1	0	0
Max RAM	0	0	1	0

The table below shows the single participant's rankings.

CNRDP serial n=1	Rank 1	Rank 2	Rank 3	Rank 4	Rank 5
Increase NJIT-provided HPC resources shared among users	0	0	1	0	0
Increase user-purchased HPC resources	0	1	0	0	0
Increase NJIT-provided commercial off-premise resources shared among users	1	0	0	0	0
Increase user-purchased commercial off-premise resources	0	0	0	1	0
Increase use of publicly available HPC resources (e.g., at a national supercomputing center); successful proposal by researcher is required	0	0	0	0	1

Gorgon

Only 1 participant out of 38 (3% of those who use serial processing) indicated use of Gorgon. The table below shows this participant's rating for each aspect.

GORGON, serial n=1	Large increase	Moderate increase	Adequate	Don't know
Number of cores	0	0	1	0

Number of nodes	0	1	0	0
CPU speed	0	1	0	0
Max RAM	0	0	1	0

The table below shows single participant's rankings.

GORGON n=1	Rank 1	Rank 2	Rank 3	Rank 4	Rank 5
Increase NJIT-provided HPC resources shared among users	0	0	0	1	0
Increase user-purchased HPC resources	0	1	0	0	0
Increase NJIT-provided commercial off-premise resources shared among users	1	0	0	0	0
Increase user-purchased commercial off-premise resources	0	0	0	0	1
Increase use of publicly available HPC resources (e.g., at a national supercomputing center); successful proposal by researcher is required	0	0	1	0	0

Phi

No participants indicated serial use of Phi.

Use of non-NJIT resources

n=21

Participants were asked whether they used non-NJIT HPC resources, and if so, to list those resources and indicate reasons for their use from a provided list and (optionally) by written description. 21 out of the 53 (40%) of participants in the HPC resources topic section indicated that they used HPC resources outside of NJIT.

Participants' choices of reasons are shown in the bar graph below. The specified non-NJIT resources, along with the reasons for their use, are shown in the subsequent table. Each table row represents responses from a single participant. Hence, resources specified by more than one participant are listed more than once were.

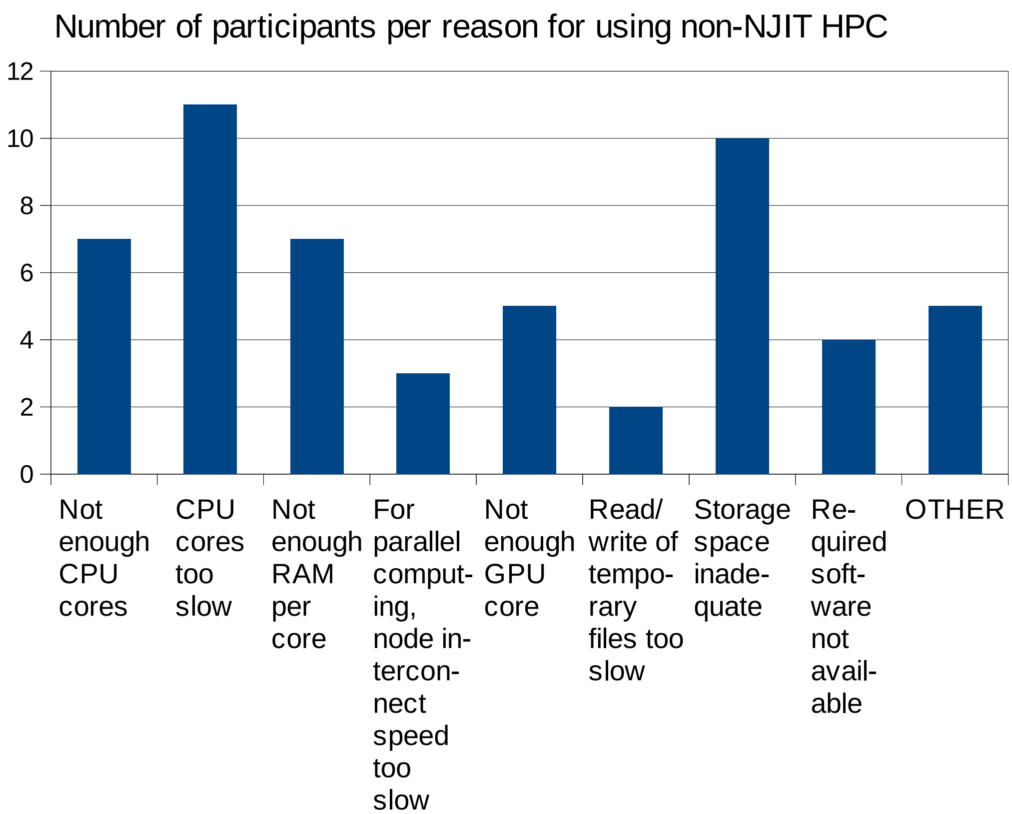

REASONS →	Not enough CPU cores	CPU cores too slow	Not enough RAM per core	For parallel computing, node interconnect speed too slow	Not enough GPU cores	Read/write of temporary files too slow	Storage space inadequate	Required software not available	OTHER
% CHOOSING the REASON →	33%	52%	33%	14%	29%	10%	48%	19%	24%
OUTSIDE RESOURCE									
Xcede		x							
Xsede		x							
Xsede supplementary computation	x	x	x	x	x				
Compute Canada, Xsede		x		x					
Amazon Webservice. Google cloud		x	x		x		x		

platform									
AWS & GCP			x	x		x	x	x	
GPU station with 4 TitanXp GPUs, purchased by professor					x				
lab	x	x							
NFS computer, DOE & Navy computers	x		x				x	x	
NASA Pleides supercomputer							x		
Resources at NYU									x
TACC					x		x		x
CIPRES portal							x		x
Workstation and other universities with better support	x	x	x		x		x	x	
Rutgers Conley3 and Amarel							x		
8-core 4.0Ghz, high performance GPU 2GB workstation		x							
Stanford HPC system, Barcelona Supercomputing Center									x
UIUC cluster machine	x	x	x				x		x
HPC from UPenn							x	x	
unspecified	x	x			x	x			
Unspecified, provided by collaborators	x	x	x						

The table below lists the written descriptions for *OTHER* and for *Additional comments*.

OUTSIDE RESOURCE	Reasons: OTHER	ADDITIONAL COMMENTS
Xsede		Because I'm facing a problem in which my jobs got killed without any reason but thats sure its due to problem in cluster and when I contact the ARCS, they said they are not able to identify the problem.
Xsede supplementary computation		NJIT HPC computational resources are not adequate for the amount of research we perform in our group.
AWS & GCP		data transfer between node and disk to memory and temp space for bd task is not enough. for data science, we don't have the latest version of lib and also I couldn't find cuDNN lib for tensorflow to run the deep learning into GPU. I believe, we may have to improve hpc/bd resources for deep learning for both parallel and serial in terms of lib, driver, development and testing environment
lab		Easier access. Response time.
NFS computer, DOE & Navy computers	some software is just too expensive for the university to consider for one or two users only - so use if can find on nsf or dod computer x	
NASA Pleides supercomputer		For Kong head node, the 1 GB quota is too small
Resources at NYU	More experiments can be run :)	Team is made up of NYU and NJIT students/faculty. Therefore, it is easy for us to use both clusters to run our experiments.
TACC	Inadequate support	Inadequate support and not to up-to-date documentation
Stanford HPC system, Barcelona Supercomputing Center	I am still completing some work on those machines.	I will be using those only until I am done with some old computation. I will switch to NJIT resources very soon, though.;l.pppppp/'
UIUC cluster machine	GLIBC is too old on stheno	For Kong, I was attempting to install miniconda and work on specific project in Python. Unfortunately, the storage for each user is just 5GB and it is totally not enough for a student user who has need in large computation. For stheno, the GLIBC version is way too old (2.5) and it should have up-to-date version for users.

New processors
n=32

Participants were asked whether they were interested in any of a presented list of new processors, and to give reasons for their interest. 32 (60 %) of participants in the HPC resources topic section indicated interest.

Participant interest in each processor is shown in the bar graph below. Following the graph, the percentage of participants indicating interest in each processor, along with the reasons for their interest, are listed in a table.

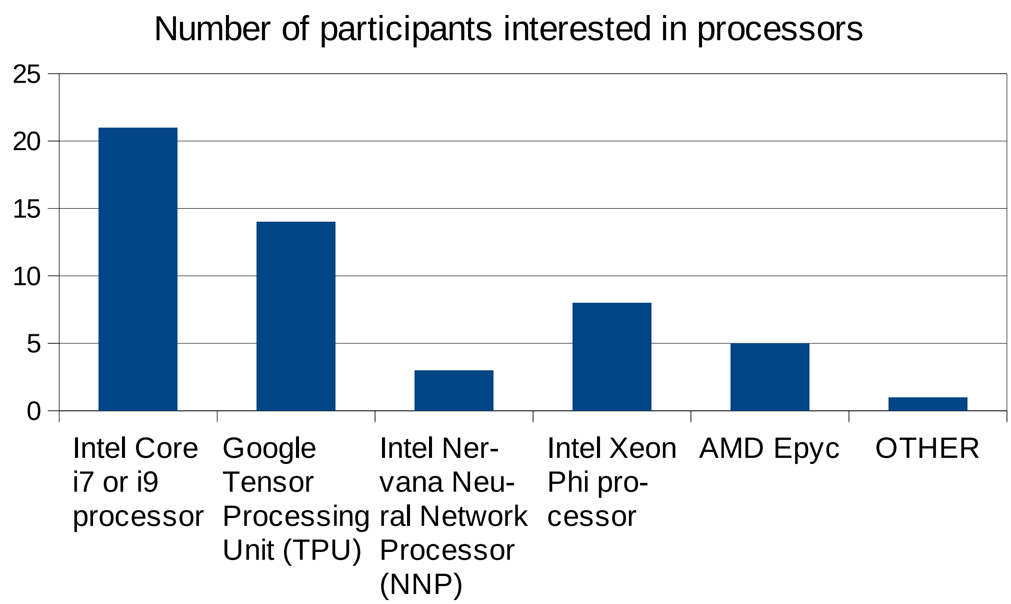

NEW PROCESSOR	% indicating interest	Comments
Intel Core i7 or i9 processor	66%	Could speed up serial job
		computation speed is one of most important issue in my research
		better choice
		Fast CPU

		We need fast processors throughout our work
		Faster and parallel compilation which leads to less code compilation time
		Recent advancements in processor performance provide an increase in computational speed at similar upkeep costs to previous iterations
		highest clock speed
		the simulations require
Google Tensor Processing Unit (TPU)	44%	Sounds great for deep learning applications!
		Deep Learning Projects
		we have p100 but tpu is much faster. I like to use for deep learning
		For deep learning research
		The ASIC provided by the TPU is designed specifically to be used for machine learning, a growing field in computer science.
		I want to try new neural network machine learning algorithms on TPU.
		Efficient for machine learning model training and deployment.
		More neural network research in the future
Intel Nervana Neural Network Processor (NNP)	9%	More neural network research in the future
Intel Xeon Phi processor	25%	Could speed up serial job
		Best compatibility
		It is discontinued; but nice if we had one
		Xeon Phi processors provide the ability to run x86 code on a device that operates much like a GPU.
AMD Epyc	16%	Could speed up serial job
		Best value
		More cores per node
		biggest bang for the buck
OTHER	3%	NVIDIA Volta Series GPU

HPC computational resource documentation
n=53

All 53 participants in the HPC Hardware Resources topic section evaluated the adequacy of HPC computational resources documentation at the HPC and BD Wiki. Participants who indicated a rating other than *Adequate* or *Don't know* where asked to suggest improvements.

The number of participants indicating each adequacy judgment are shown in the pie chart below; suggestions are listed subsequently.

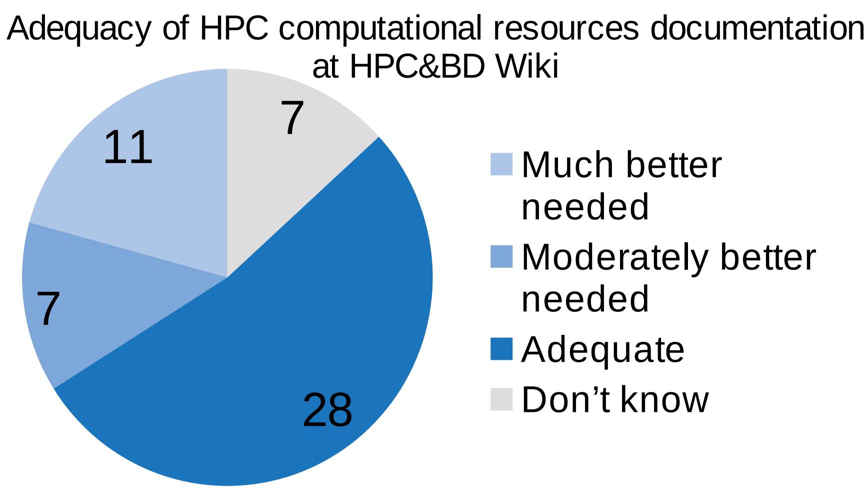

Suggestions by participants who indicated *Much better needed*:

- allowing users to edit the wiki; make the mailing forum available on the wiki
- More detailed resources needed
- Number of nodes, CPU cores, storage. I believe that the size of RAM now is quite enough.
- The links are either hard to find or out of date. There are no instructions how to start working and what needs to be done.
- One needs to experiment with trial and error. Some information is misleading or incorrect. The Wiki does not have a start and end. One needs to click all the entries and figure out what is needed or not.
- Most of the time, the CPU jobs are super slow and take weeks to get finished. GPU, on the other hand, is pretty fast, but the number of GPU nodes are very less. Sometimes, few of the nodes do not work properly and job gets killed.
- It should be up dated frequently

- Better organization of the documentation resources is needed so that more people are aware of the HPC facilities at NJIT. The documentation should be easy to follow. For instance, I did not know the existence of Google's TPU till now.

- A detailed documentation is needed especially a tutorial. Current wiki doesn't contain all detailed command settings.

- Required documentation is mostly unavailable. Available documentation has not been updated since a long time.

- we have lot of information on the wiki. generally, it was scattered around all place.
  1. It would good to have a single page what is available on HPC/BD [from this we can navigate to all resources]. Thanks to whoever created hpc hardware resource sheet[this one really good we can see what type of machine we have and hardware spec for each machine]. but missing piece how do access these machines.example i want to access kong-10 from my afs account.i couldn't find it there.it would be nice.some kind of link how to access those resources.
  2. what are software running on those machines[like hardware sheet]
  3.if i want to install something. what is the process that i need to follow. i know we can someone to help it. but that one is taking time.

- a more extensive documentation is needed. please compare to hpc resources at other schools.

Suggestions by participants who indicated *Moderately better needed*:

- More resources, more cores are needed, just as the research need grows.
- Better descriptions for novice users...maybe provide an online tutorial
- Discussion forum would be helpful. List of some useful scripts would be great as well.
- It is not easy to use. I always have problems finding simple documentation.
- Sometime it takes too long to run the cases. Also, small jobs should run immediately.
- I would like to see usage examples of different scripts on different resources. Also, maybe a better FAQ page that gets updated as questions arise.
- More manuals on data science nodes are needed.

Final HPC resources topic section comments

All participants in the HPC resources topic section were invited to provide further comments. (Participants who reported using both parallel and serial resources could provide comments within each parallel and serial processing subsection, and thus had two opportunities to comment.)

Comments about parallel resources by participants who use both parallel and serial resources:

- As computational chemistry levels and methods increase - use of these higher level and more demanding calculations are needed to keep up with publishing in top level journals.

- GPU nodes not working well with Lammps

- Based on my experience, Generally core and node wise in terms of CPU, GPU is fine but memory wise it is not enough [noticed we have to lower the memory allocation for the variable due to memory error] and data transfer between the node or disk to memory is taking a long time. these are two symptom mostly that i experienced

- We implement and use multi-core programs in most (if not all) of our work otherwise experiments take days to finish. We expect this to increase going forward and expect other researchers to face the same needs.

- Deep learning is popular, so nodes contain Nvidia GPU are needed

Comments from participants who only use parallel resources:

- The staff is excellent and my group has benefited a lot from their assistance and guidance.

- I kindly request to increase the number GPU nodes and speed of the CPU nodes.

Comments about serial resources by participants who use both parallel and serial resources:

- Quality journals are requiring higher and higher level calculations for best accuracy and consequent publication - the computational resource needs are increasing.

- here also same data transfer time and memory is not enough

- We do serial computations in a distributed fashion in large-scale experiments on real and simulated data. We run the same serial program on many multiple inputs across different nodes. Thus our serial work is really distributed computing and we expect the same of other researchers.

- Kong is very bad and very unstable, errors frequently happen when I run simulations

Comments from participants who only use serial resources:

- It would be nice if HPC center offers practical non-credit courses to students and faculty members on how to use HPC resources.

- ok now, but could be better

- They are very slow, and support is not adequate.

- I haven't used it sufficiently to give an adequate answer to this.

Topic Section 2 HPC storage

45 participants (78%)

AFS and NFS storage

Participants were shown the following background information, and then assessed the adequacy of the base allocations AFS and NFS storage. Participants who rated either allocation as other than *Adequate* (or who chose *Don't know*) were asked to rank four potential remediations.

Storage terminology:

- AFS distributed filesystem: General computational use; accessible from all cluster nodes; accessible from all other HPC servers
- NFS distributed filesystem: General computational use; separate filesystems accessible from Kong and Stheno; not accessible from other HPC servers

The base storage allocations for researchers are shown in the table below.
Base resources are available to all NJIT researchers. If resources beyond the base allocation are needed, the researchers must arrange for the purchase of such resources. Resources that are purchased are dedicated to the purchaser.

Resource type	Resource name	Default allocation	Cost of dedicated resource	Notes
Disk	AFS	500GB max [1]	DiskAndBackupCost	Accessible from all AFS clients – Linux, MacOSX, Windows
Disk	NFS	500GB max [2]	DiskAndBackupCost	Accessible only from the HPC cluster to which it is attached

[1]. Independent of NFS allocation.
[2]. Independent of AFS allocation.

Potential remediation options involve increases in resources that differ in whether they are:

- Paid for by NJIT *versus* paid for by individual users
- On-premise *versus* off-premise

Note: Off-premise providers include *Amazon Web Services*, *Azure*, *Google Cloud Platform*, *IBM Bluemix*, *Oracle Cloud*, *Penguin Computing*

Choices to be ranked: Increase NJIT-provided HPC resources chared among users; Increase user-purchased HPC resources at NJIT; Increase NJIT-provided commercial off-premise resources shared among users; Increase user-purchased commercial off-premise resources

The number of the 45 participants indicating each Adequacy judgment for AFS and NFS storage are shown in the pie charts.

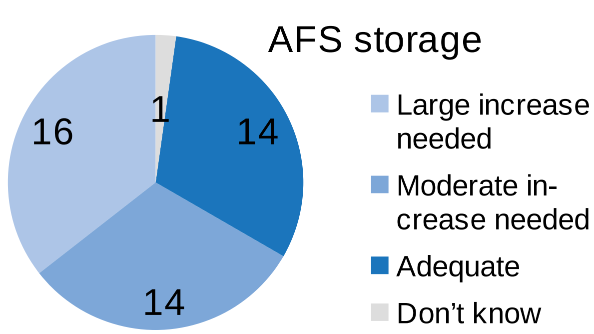

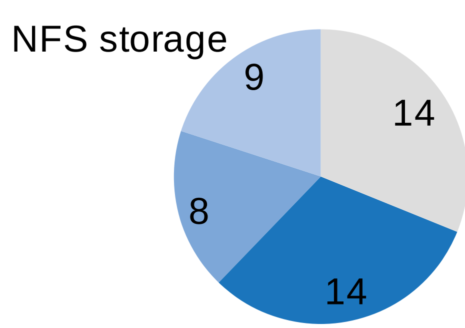

30 participants indicated that at least one of the base allocations was not adequate. Number of participants assigning each remediation to each rank are tabulated below.

n=30	Rank 1	Rank 2	Rank 3	Rank 4
Increase NJIT-provided HPC resources shared among users	18	8	1	3
Increase user-purchased HPC resources at NJIT	6	8	12	5
Increase NJIT-provided commercial off-premise resources shared among users	3	13	12	1
Increase user-purchased commercial off-premise resources	3	1	5	21

Parallel file system (PFS)

Participants were shown the following background information on PFS, and were then asked to indicate the importance, for their research, of having PFS capability in the HPC clusters they used. Participants who chose either *Very important* or *Moderately important* were asked to provide reasons.

Parallel file system (PFS)

Background

- The HPC cluster model uses hundreds of compute nodes, each containing several CPUs, each of which contains several processors (cores) to perform calculations. The two NJIT HPC clusters, Kong (general-access) and Stheno (Dept. Mathematical Sciences only), between them contain 3,448 cores. Jobs running on these cores write/read temporary files to/from disk as part of the computational processes.
- When the temporary files exceed a few gigabytes, writing to and reading from temporary files on disk consumes the most computation time- i.e., disk I/O is the bottleneck. In addition, parallel jobs handling temporary files of any size will encounter this bottleneck.
- The problem is exacerbated by the very large increase in the computational capacity and number of cores in Kong in summer 2015, which has resulted in a very large increase in the amount of large temporary data that the compute nodes are attempting to write to and read from disk.

Implication

- Researchers at NJIT using HPC clusters are dealing with increasingly large sets of data, with the concomitant need for much higher I/O capacity for temporary space.

PFS appliance

- A PFS is a file system that distributes file data across multiple servers and provides for concurrent access by multiple tasks.

- A PFS can be used by both serial and parallel processes running on an HPC cluster.
- A PFS appliance can be connected to multiple clusters.
- PFS examples: IBM General Parallel File System (GPFS), Lustre

The distribution of the 45 participants indicating the importance of PFS capability are shown in the pie chart below.

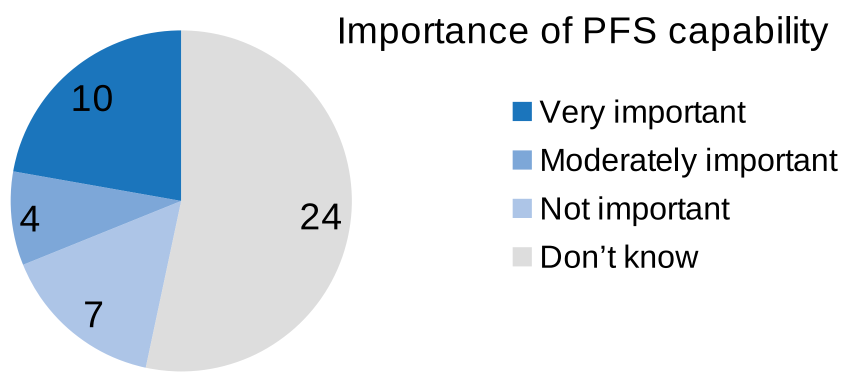

Reasons why 10 participants considered PFS capability *Very important*:

- Recently, we have been trying to move the simulation results out of AFS and it is so difficult. We need to share our simulations output data with the community and we can't do that without moving it out of the cluster.

- We run massive hydrodynamic computation requiring tens of million of nodes. We just published a paper in the prestigious journal Geophysical Research Letters on the nature of oil flow from the Gulf spill of 2010. We could not have done it without the HPC resources at NJIT.

- needed for higher level calculations and reasonable time for product

- Because we often have to share data across multiple users on different servers

- It would partially reduce the I/O bottleneck and would allow to execute my simulations across systems without having to transfer data across.

- We have found that I/O takes up the most time in our experiments. To reduce this time we use the / scratch directories on the nodes but there is still a large overhead when writing to disk. A PFS is important to our work because it would reduce the time spent in reading and writing to disk which all of our jobs need to do.

- Execution time matter during code profiling and can be set off if proper file access is not maintained

- Several of the tools I use for genome evolution is done in parallel; it would be infeasible to do serial runs.
-
- Speed up my training.
- fast

Reasons why 4 participants considered PFS capability *Moderately important*:

- for faster IO
- large set of data, and we would like to share with certain colleagues
- We had a use case where spawning multiple independent threads that write to a single file would have decreased the overall processing time exponentially, Limited IO bandwidth is a huge bottle neck, which limits the amount of processing power you can tap into on kong.
- I have limited knowledge about the technical terms involved. Based on the information provided here "When the temporary files exceed a few gigabytes, writing to and reading from temporary files on disk consumes the most computation time - i.e., disk I/O is the bottleneck". I deal with large datasets - sometimes up to 500 gigabytes. I imagine that this would improved the computing time (ignore it if I am wrong).

Purchasing additional storage
n=8

8 participants out of 45 indicated involvement in decisions regarding purchasing additional storage. They were provided with the information below as background to subsequent questions.

> Researchers can purchase storage in addition to the base allocation of 500GB each of AFS and NFS space.
>
> Additional storage can be either:
>
> - **Tier 1** (very high performance, suitable for high-speed transactional databases)
> - **Tier 2** (high performance, suitable for most HPC applications)
>
> Backup choices are:
>
> - Daily
> - Reduced frequency(two to three times per week)
> - No backup

The 8 participants where asked to indicate which among a given list of costs were suitable to their research needs; if none were suitable, participants were asked to indicate the maximum cost, in dollars per TB per year, at which they would purchase NJIT storage, including backup.

The table below shows the distribution of responses. Each row represents responses of a single participant.

Tier 1, no backup: $870/TB/year	Tier 1, reduced frequency backup: $1010/TB/year	Tier 1, daily backup: $1160/TB/year	Tier 2, no backup: $250/TB/year	Tier 2, reduced frequency backup$390/TB/year	Tier 2, daily backup: $540/TB/year]	None suitable; maximum annual per TB researcher is willing to pay
x			x			
				x		
				x		
x			x	x	x	
						0
						0
						20
						comparable to the cost of storage that I could purchase myself. (Just purchased external drive - 20TB for $1,000.

The 8 participants were asked if the current annual payment of storage costs was satisfactory, and if not, why not, and what yearly basis for charges they would recommend.

All 4 participants who indicated a suitable Tier cost in the previous question indicated satisfaction with annual payment.
All 4 participants who indicated that none of the listed costs were suitable likewise indicated that annual payment was not satisfactory.

The table below lists their reasons and recommendations for schedule of storage cost charges. Each row represents responses of a single participant.

Why is annual payment not satisfactory?	**On what early should storage costs be charged?**
Monthly is better	Two year
Availability of funds depends on the grant-specific dates.	Two year
grant money is ebb-and-flow	Five year
This is expensive. That 20TB for $1,000 drive that I just purchased will serve at least 4-5 years.	Five year

Platform-independent HPC storage

All 45 participants in the HPC Storage topic section were asked to indicate the importance of platform-independent access to HPC storage, and were provided with the following background information:

> "Platform-independent access" means that file paths and authorization are independent of which platform - Linux, MacOSX, Windows - is being used to access files.

The distribution of the 45 participants indicating the importance of platform-independent access to HPC storage is shown in the pie chart below.

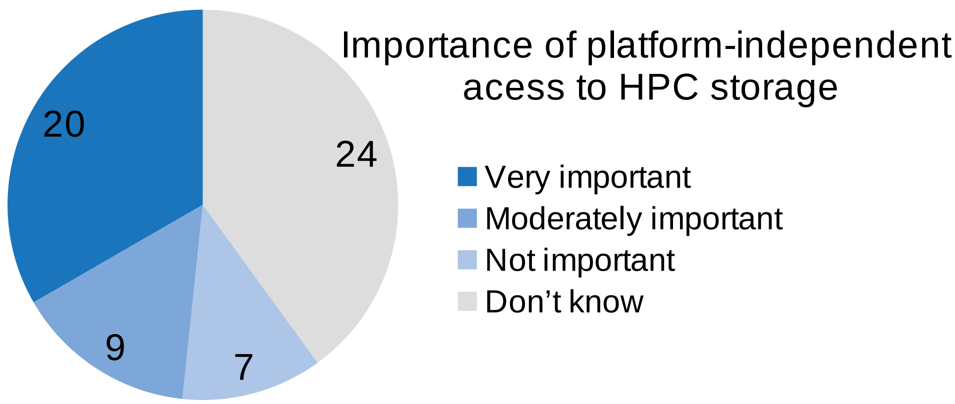

HPC storage documentation Wiki

All 45 participants in the HPC Storage topic section were asked to indicate the adequacy of HPC storage documentation at HPC & BD Wiki. Participants who indicated any inadequacy were asked to suggest improvements. Responses are shown in the pie chart below, followed by comments.

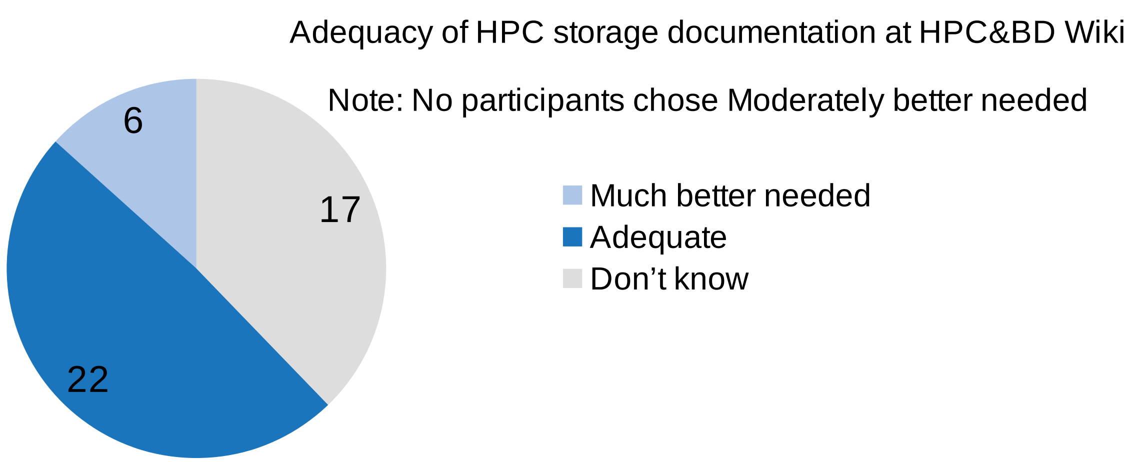

Suggestions from participants who indicated a need for *much better* documentation:

- not sure where to find them

- If the storage can be increased upon students' requests, that would be great.

- Same comment as for Wiki documentation in Section 1, HPC resources: The links are either hard to find or out of date. There are no instructions how to start working and what needs to be done. One needs to experiment with trial and error. Some information is misleading or incorrect. The Wiki does not have a start and end. One needs to click all the entries and figure out what is needed or not.

- The information on the web page should be updated more often

- Have all the information related to HPC at one place. Better web page navigation is needed.

- We have lot of information on the wiki. Generally, it was scattered around all place.
  1.It would good to have a single page what is available on HPC/BD [from this we can navigate to all resources]. Thanks to whoever created hpc hardware resource sheet[this one really good we can see what type of machine we have and hardware spec for each machine]. but missing piece how do access these machines.example i want to access kong-10 from my afs account. I couldn't find it there. It would be nice.some kind of link how to access those resources.
  2. what are software running on those machines[like hardware sheet]
  3.if i want to install something. what is the process that i need to follow. i know we can someone to help it. but that one is taking time.

Final HPC storage topic section comments

All participants in the topic section were invited to provide further comments.

- Nothing critical, but this is the era of big computation and big data. Our computation produces gigabytes of data and we need to have them stored AND easily accessible.

- enough for usage of me

- The "rent" is too damn high! By an order of magnitude.

- The current costs of Tier 1 and Tier 2 disk space is excessively high compared to the market. Many researchers at NJIT feel the same. The current cost for an 8TB drive from BestBuy is $179.99 whereas 8TB of Tier 2 unbackedup space would cost $2000 at NJIT. This is excessively high for big data research. Most researchers need unbackedup terabyte space. We need to revise the costs so that NJIT researchers who need more space can purchase it easily at NJIT.

- storage is good

Topic Section 3 BD computation and storage

14 participants (26%)

Participants saw the following background information:

> The Hadoop/Spark infrastructure is a virtual environment based on VMware Big Data Extensions (BDE).
>
> VMware introduced Big Data Extensions, or BDE, as a commercially supported version of Project Serengeti designed for enterprises seeking VMware support. BDE enables customers to run clustered, scale-out Hadoop applications on the vSphere platform, delivering all the benefits of virtualization to Hadoop users. BDE delivers operational simplicity with an easy-to-use interface, improved utilization through compute elasticity, and a scalable and flexible Big Data platform to satisfy changing business requirements. VMware has built BDE to support all major Hadoop distributions and associated Apache Hadoop projects such as Pig, Hive, and HBase.
>
> The hardware (Horton.njit.edu) associated with BDE is as follows:
>
> - 2 x IBM iDataPlex dx360 M3 nodes, each with:
>   - 2 x Intel Xeon CPU E5-2680 (8 Core)
>   - 16 CPU CORES @ 2.70GHz
>   - 32 Logical Processors with Hyperthreading
>   - 128G RAM
> - 3TB HDFS (Hadoop Distributed File System) disk

Adequacy of big data and storage resources

Participants were asked to indicate the adequacy of aspects of big data computational and storage resources. Participants who indicated that any aspect was not adequate where then asked to rank potential remediations, with the following instructions:

> Potential remediation options involve increases in resources that differ in whether they are:
>
> - Paid for by NJIT *versus* paid for by individual users
> - On-premise *versus* off-premise
>
> **Note:** Off-premise providers include *Amazon Web Services, Azure, Google Cloud Platform, IBM Bluemix, Oracle Cloud, Penguin Computing*
>
> Choices to be ranked: Increase NJIT-provided HPC resources chared among users; Increase user-purchased HPC resources at NJIT; Increase NJIT-provided commercial off-premise resources shared among users; Increase user-purchased commercial off-premise resources

The table below shows the number of participants who provided each particular rating for each aspect of big data computational and storage resources.

n=14	Large increase needed	Moderate increase needed	Adequate	Don't know
Number of cores	2	3	3	6
Max Ram per node	1	3	4	6
Amount of HDFS storage	3	1	4	6

BD comp and storage inadequate n=5	Rank 1	Rank 2	Rank 3	Rank 4
Increase NJIT-provided HPC resources shared among users	4	0	1	0
Increase user-purchased HPC resources at NJIT	0	2	2	1
Increase NJIT-provided commercial off-premise resources shared among users	0	2	2	1
Increase user-purchased commercial off-premise resources	1	1	0	3

Purchasing additional BD storage
n=4

4 out of 14 are involved in purchasing additional storage.

4 participants out of the 14 in this topic section indicated involvement in decisions regarding purchasing additional BD storage.
They were provided with the information below as background to subsequent questions.

> Researchers can purchase storage in addition to the base allocation of 500GB each of AFS and NFS space.
>
> Additional storage can be either:
>
> - **Tier 1** (very high performance, suitable for high-speed transactional databases)
> - **Tier 2** (high performance, suitable for most HPC applications)
>
> Backup choices are:
>
> - Daily
> - Reduced frequency(two to three times per week)
> - No backup

The 4 participants where asked to indicate which among a given list of costs were suitable to their research needs; if none were suitable, participants were asked to indicate the maximum cost, in dollars per TB per year, at which they would purchase NJIT storage, including backup.

The table below shows the distribution of responses. Each row represents responses of a single participant.

Tier 1, no backup: $870/TB/year	Tier 1, reduced frequency backup: $1010/TB/year	Tier 1, daily backup: $1160/TB/year	Tier 2, no backup: $250/TB/year	Tier 2, reduced frequency backup$390/TB/year	Tier 2, daily backup: $540/TB/year]	None suitable; maximum annual per TB researcher is willing to pay
x			x			
				x		
						0
						20

The 4 participants were asked if the current annual payment of storage costs was satisfactory, and if not, why it was unsatisfactory, and what yearly basis for charges they would recommend.

Both participants who indicated a suitable Tier cost in the previous question indicated satisfaction with annual payment.
Both participants who indicated that none of the listed costs were suitable also indicated that annual payment was not satisfactory.

The table below lists their reasons and recommendations for schedule of storage cost charges. Each row represents responses of a single participant.

Why is annual payment not satisfactory?	**On what early should storage costs be charged?**
Monthly is better	Two year
grant money is ebb-and-flow	Five year

Final BD computation and storage comments

All participants in the topic section were invited to provide further comments.

- Universities are building 1,000 nodes clusters with 100 GB for each user. We need to stay competitive.

- NJIT resources are excellent for learning and proof of concept prototyping.   Massive and temporary computational resources require public/community cloud service providers.

- The "rent" is too damn high! Like by an order of magnitude.

- We have lot of information on the wiki. generally, it was scattered around all place.
  1.It would good to have a single page what is available on HPC/BD [from this we can navigate to all resources]. Thanks to whoever created hpc hardware resource sheet[this one really good we can see what type of machine we have and hardware spec for each machine]. but missing piece how do access these machines. example I want to access kong-10 from my afs account. I couldn't find it there. It would be nice for some kind of link on how to access those resources.
  2. what are software running on those machines[like hardware sheet]
  3.if I want to install something. what is the process that I need to follow? I know we can someone to help it. but that one is taking time.

- The costs are excessively high compared to current market rates. AT BestBuy 8TB costs $179.99 whereas the same Tier 2 unbacked-up costs $2000 at NJIT. This is unaffordable for most researchers and so we need to bring the costs down.

Topic Section 4 Software environment

41 participants (71%)

The information below, along with a list of approximately 200 HPC and BD software options currently available, was shown.

> The software environment is the combination of applications - open source and commercial, libraries, utilities, and modules.
> (*Modules* are used for setting the user's environment for specific software.)

Suitability of software environment

Participants were asked to rate the suitability of the software environment for their work. The table below shows the number of participants who provided each particular rating.

All 41 participants in the Software environment topic section were asked to indicate the suitability of the software environment for their work. Responses are shown in the pie chart below.

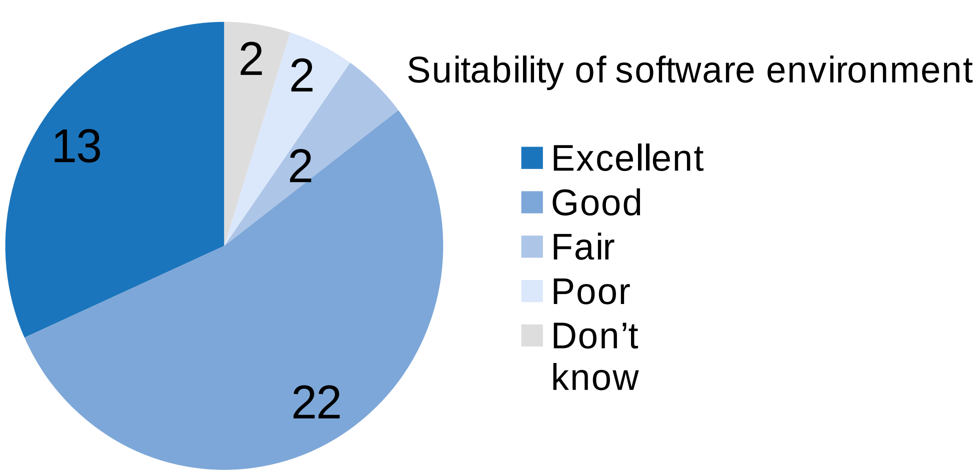

Participants were asked (but not required) to name software not already available that they would like to use, and to indicate whether there were associated costs, along with anticipated use (*high, medium, low, none*) for research and for teaching. Each participant could name up to 5 items.

Additional software

14 out of 41 (34%) requested new software, shown in the table below.
Note that some items (e.g., Tensorflow) were requested by multiple participants. Each individual request is listed, as participants varied in their proposed usage.

NAME of REQUESTED SOFTWARE	Associated costs	Research use	Teaching use
Matlab 2017b	yes	high	high
Matlab Computer Vision Toolbox	yes	high	high
Updated version of Lammps; Working GPU acceleration in Lammps, Lammps with most of the packages preinstalled	no	high	low
Updated Anaconda Python 2 (and maybe 3) packages/libraries	no	high	high
Too many options for compiler lead to problems; simplify the list	no	medium	low
cclib python library; can be installed from anaconda by "conda install -c omnia cclib"	no	high	high
Quantum Espresso; Open source alternative to Gaussian 16	no	high	high
GAMESS; Open source alternative for Gaussian	no	high	high
Polyrate; Software for rate constant calculations; has interfaces for Gaussian, NWChem, GAMESS	no	high	high
Open babel; like cclib it's a useful tool to have	no	high	high
Tensorflow	no	high	low
Tensorflow	no	medium	medium
Tensorflow	no	high	medium
tensorflow	no	medium	high

Scikit Learn (python library)	no	high	low
Keras	no	medium	low
Keras	no	high	low
keras	no	high	medium
Nektar++	no	high	no
cuDNN. Deep learning for GPU	no	high	high
git latest version	no	high	high
Stata	yes	high	no
SAS	yes	high	no
http://liulab.dfci.harvard.edu/MACS/Download.html (MACS)	no	high	no
BETA: http://cistrome.org/BETA/	no	medium	no
STAR: https://github.com/alexdobin/STAR	no	medium	no
pytorch	no	high	medium

Final Software environment comments

All participants in the topic section were invited to provide further comments.

- It would be nice to have a website with all softwares available and a brief description of each software.
- Update the OS!
- The staff of HPC has been very accommodating to us adding software when needed.
- Overall software environment is excellent. Current suggested software are some reputable open source alternatives. The cclib python library should be added as it's a great tool for post-processing computational chemistry data. Open babel is also a very useful tool.
- You forgot the new mono installation!

- Ok
- Please keep update the existing packages/libraries
- Standard C++ based code on Unix environment; https://www.nektar.info/

Topic Section 5 Internet bandwidth and Science DMZ

24 participants (41%)

Internet bandwidth

n=22

The following introductory information was provided:

> Internet bandwidth is the capacity of NJIT's connection to and from the Internet.

Participants rated the suitability of Internet bandwidth, including Internet 2 if applicable, to their work.
The pie chart below shows the number of participants who provided each particular rating.

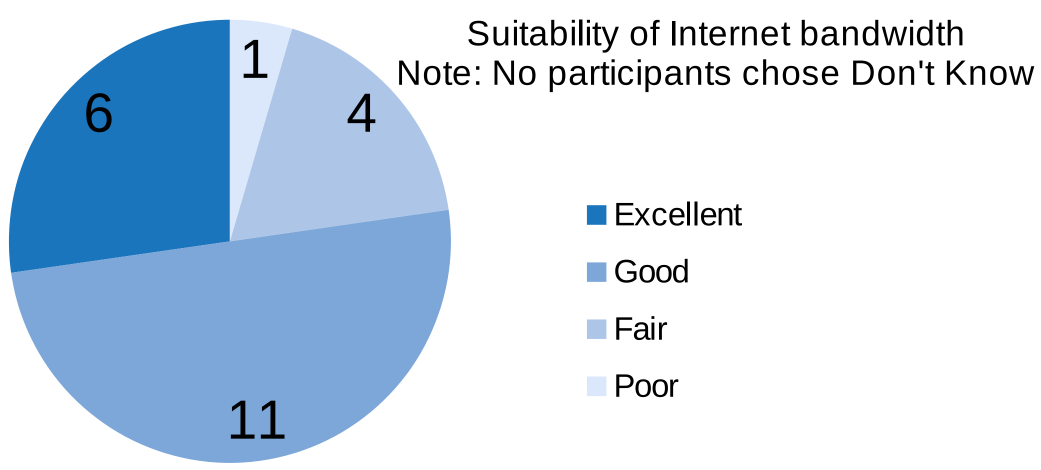

Participants indicated whether their research has been hampered by difficulties in transferring data from or to the Internet.
6 out of 22 (27%) reported such difficulties, and were asked to describe the difficulties. Their descriptions are listed below.

- using external resources that are faster

- We attempt to run our remote site in real time remotely, which works fine at home, but at NJIT it is often too slow and laggy to be reliable. We would like to transfer our stored data to NJIT, but would not even attempt it given NJIT's slow network.
- slow connection to nearby universities, Rutgers in particular. Not sure about the source of the problem.
- While working on the remote connection to NJIT servers from outside the US, the speed slows down greatly. The speed of the internet service at is fast however.
- Just one host downloading data from Europe managed to bring down the Internet 2 connection.
- Downloading multiple large files from the NIH Cancer databases.
- Some times it is slow

Final Internet Bandwidth comments

All participants in the topic section were invited to provide further comments.

- Pretty good, though it varies depending on my location on campus

Science DMZ

n = 8

The following introductory information was provided:

> **"Science DMZ" refers to a computer subnetwork that is structured to be secure, but without the performance limits that would otherwise result from passing data through a stateful firewall.**
>
> The Science DMZ is designed to handle high volume data transfers, typical with scientific and high-performance computing, by creating a special DMZ to accommodate those transfers.
>
> Science DMZ is typically deployed at or near the local network perimeter, and is optimized for a moderate number of high-speed flows, rather than for general-purpose business systems or enterprise computing.

Desirability of Science DMZ

Participants were asked to rate the desirability of implementing a Science DMZ at NJIT as it related to their work.
Distribution of choices are shown in the pie chart below.

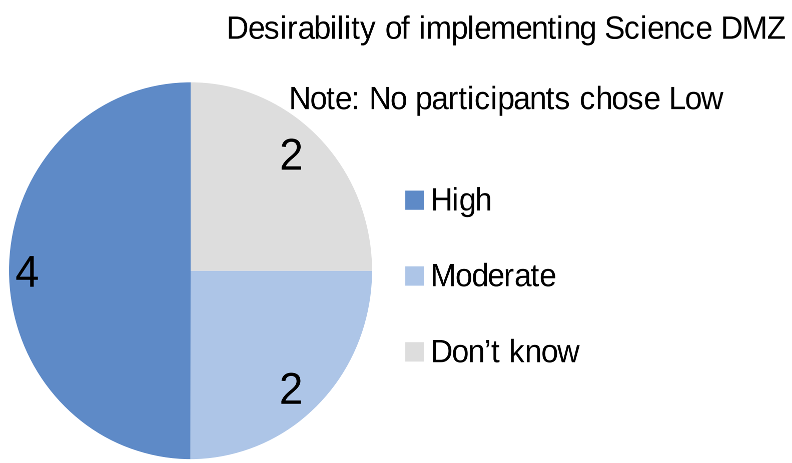

Final Science DMZ comments

All participants in the topic section were invited to provide further comments.

- It would be helpful to researchers in big data who download big datasets.

Topic Section 6   Consultation with ARCS

43 participants (74%)

The following introductory information was provided:

> *Consultation* is interaction with Academic and Research Computing Systems (ARCS) staff in areas such as getting started in HPC and BD, problems encountered when running jobs, optimizing throughput, running parallel jobs, managing disk space, and assistance in working with and managing big data.

Effectiveness of ARCS consultation

The 43 participants in the Consultation topic section rated the effectiveness of consultation in their research.
The pie chart below shows the number of participants who provided each particular rating.

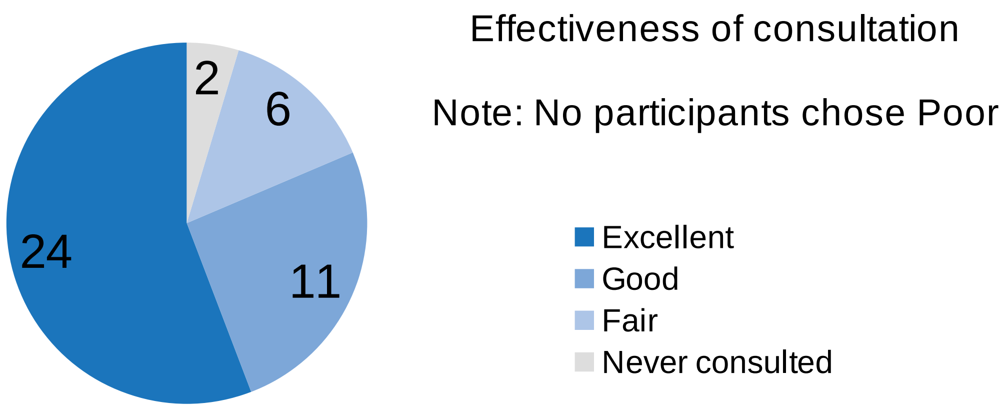

Final Consultation comments

All participants in the topic section were invited to provide further comments.
Note: Except as indicated, all comments are from participants who rated consultation as *excellent*.

- Requests for help have been promptly replied and resolved fast. The consultations have been very helpful.
- [*Staff member G. Wolosh*] is great. But so is practically every person we interacted with.
- Consultation and interaction with ARCS has been very good to excellent - All of my interaction and interactions of my research group have been successful in terms of product use and availability.
- Excellent work on helping me get my code running. [Staff member G. Golosh] is awesome.
- ARCS staff is excellent.
- Impressive in fact.
- Really good
- These guys are top-notch!
  - They are experts
  - They are extremely reliable, available, and answer our requests very promptly.

  I couldn't be happier.
- Excellent consultation
- Response is always great; fast and accommodating.
- Cannot say more good things about you! Excellent.
- *Good* rating: ARCS has been very responsive to our requests for consultation. Unfortunately, the usual result of that consultation is that our needs do not fit the model ARCS follows, so it has been somewhat frustrating.
- *Fair* rating: It could have been more speedy and more effective.

Final Section Satisfaction with IST-managed HPC/BD

56 participants

This final section was presented to all participants; that is, it was not among the selectable topic sections. Two of the initial 58 participants had quit the survey before this point, hence the ***n*** of 56.

Ratings of IST HPC/BD

Participants rated their and/or their research group's satisfaction with their use of IST-managed HPC and/or BD resources.

The pie chart below shows the number of participants who provided each particular rating.

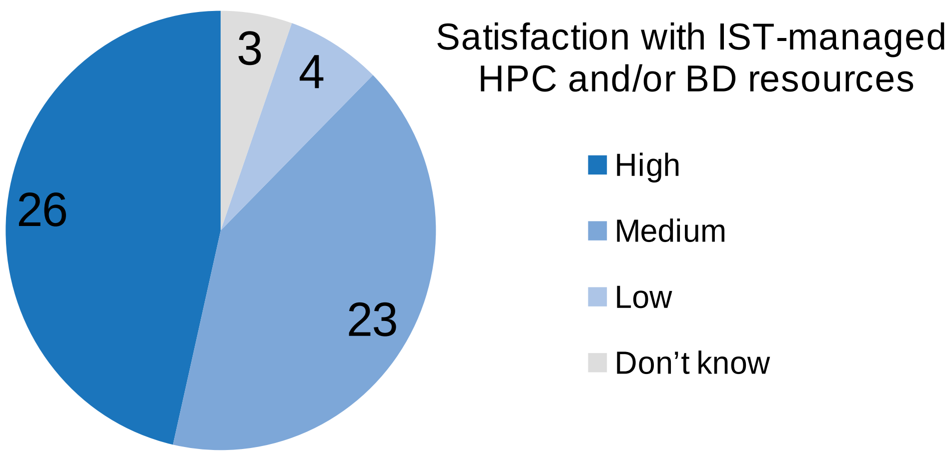

Participants were invited to suggest up to five changes to IST-managed NJIT HPC and/or BD.

Suggestions for IST HPC/BD

Suggestions from participants who rated their satisfaction as ***high***:

- Regularly updated modules such as Lammps and Python
- please provide latest git environment for kong
- installation of certain dependency related software/libraries should be allowed
- Please provide more support on data science nodes/NVIDIA GPUs since we will need to run deep learning programs on them.
- Make it support globus transfer: https://www.globus.org/

Suggestions from participants who rated their satisfaction as ***medium***:

- More storage space in the account
- Increase bandwidth for data transfer for Kong

- Provide some sort of tutorial on the website. It took me awhile to locate ARCS/find someone who could help me regarding Kong. A detailed classification of who does what would help.
- Better documentation and ease of getting access to new softwares
- Please increase the number of nodes and GPU nodes. It's better to separate users by department and assign resources to different department.
- Documentation, RAM and data transfer.
- More multi-core nodes (at least 20 nodes if not more)
- More RAM nodes
- Faster I/O to disk
- Greater Internet bandwidth for simultaneous download of several large datasets

Suggestions from participants who rated their satisfaction as ***low***:

- Increase speed and improve support.
- Storage should be increased significantly, both temporary and permanent.
- Load management should be improved. It is not clear how is this actually managed.

Final IST-managed HPC/BD comments

All participants were invited to provide further comments.

- It might be useful if there are are occasional informal 1-2 hour gatherings between HPC/BD staffs and users over tea and computers to exchange tips and demos of using the facility better. Alternatively, monthly office hour might be also helpful.
- The research group is very satisfied. We would not be productive if we did not have this facility and the ARCS group
- XSEDE is much better than NJIT HPC.
- Please add more space quota to a user (especially for a student). 5 GB cannot accommodate all input data given the size of data we are dealing nowadays. Let alone some module already take a lot of disk space.
- Hope we can use nodes with newest GPUs and RAM memory.
- kong is unstable, and it does not support the new version of matlab