

Full 2023 HPC SYSPROS Survey Results

Jenett Tillotson
Senior HPC Systems Engineer
National Center for Atmospheric Research
jtillots@ucar.edu

Notes About This Survey

- A copy of this survey can be found [here](#).
- No question was required
- Most questions had an “Other” option
- Several questions were “select all that apply”
- 54 responses
 - 13 respondents did not identify their institution
 - Some duplication of institutions

Institutions (duplicates not removed)

3M Company

Agriculture Victoria Research

Ansys

Brown University

Chan Zuckerberg Biohub

Cornell University

Digital Research Alliance of Canada

Harvard SEAS

Howard Hughes Medical Institute

Indiana University

LBNL/NERSC

Leidos Inc. - Centers for Disease Control & Prevention

Michigan State University

Middle Tennessee State University

Minnesota Supercomputing Institute

Mississippi State University

NAG/bp

NCAR

NCAR/UCAR

NCSA

NIH

NIH

NIH/NHLBI

NIHM

P&G (AFDS)

Pawsey Supercomputing Centre

Penn State Institute for Computational and Data Sciences

TACC

Tennessee Tech University

Texas A&M

The MITRE Corporation

The University of Arizona

UC Davis

UIUC/NCSA

University of British Columbia

University of California at Berkeley

University of Chicago

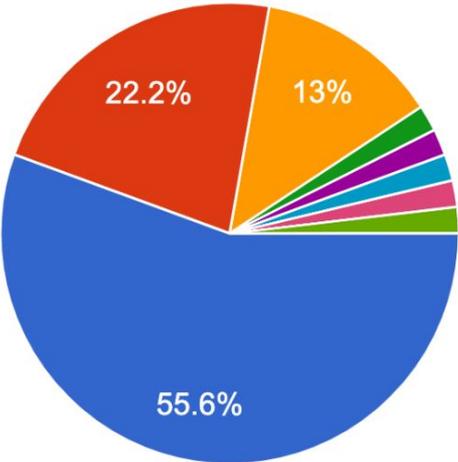
University of Pennsylvania

University of Pennsylvania

University of Utah

What type of institution do you work for?

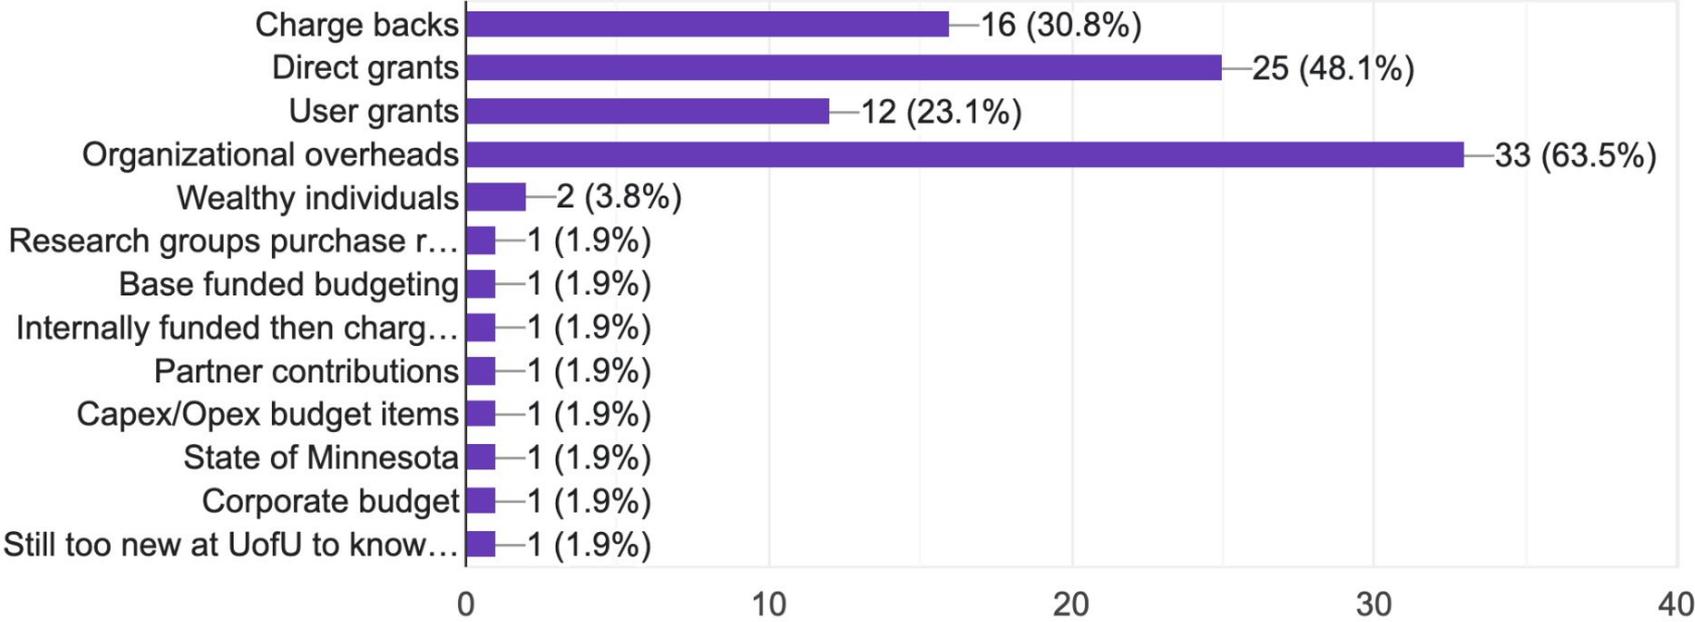
54 responses



- Academic
- Government
- Industry
- Non-profit research nondegree granting
- Non-profit and FFRDC's
- Government funded for academic use - not sure which to tick
- Non-profit academic/research
- Government Owned, Contractor Operated.

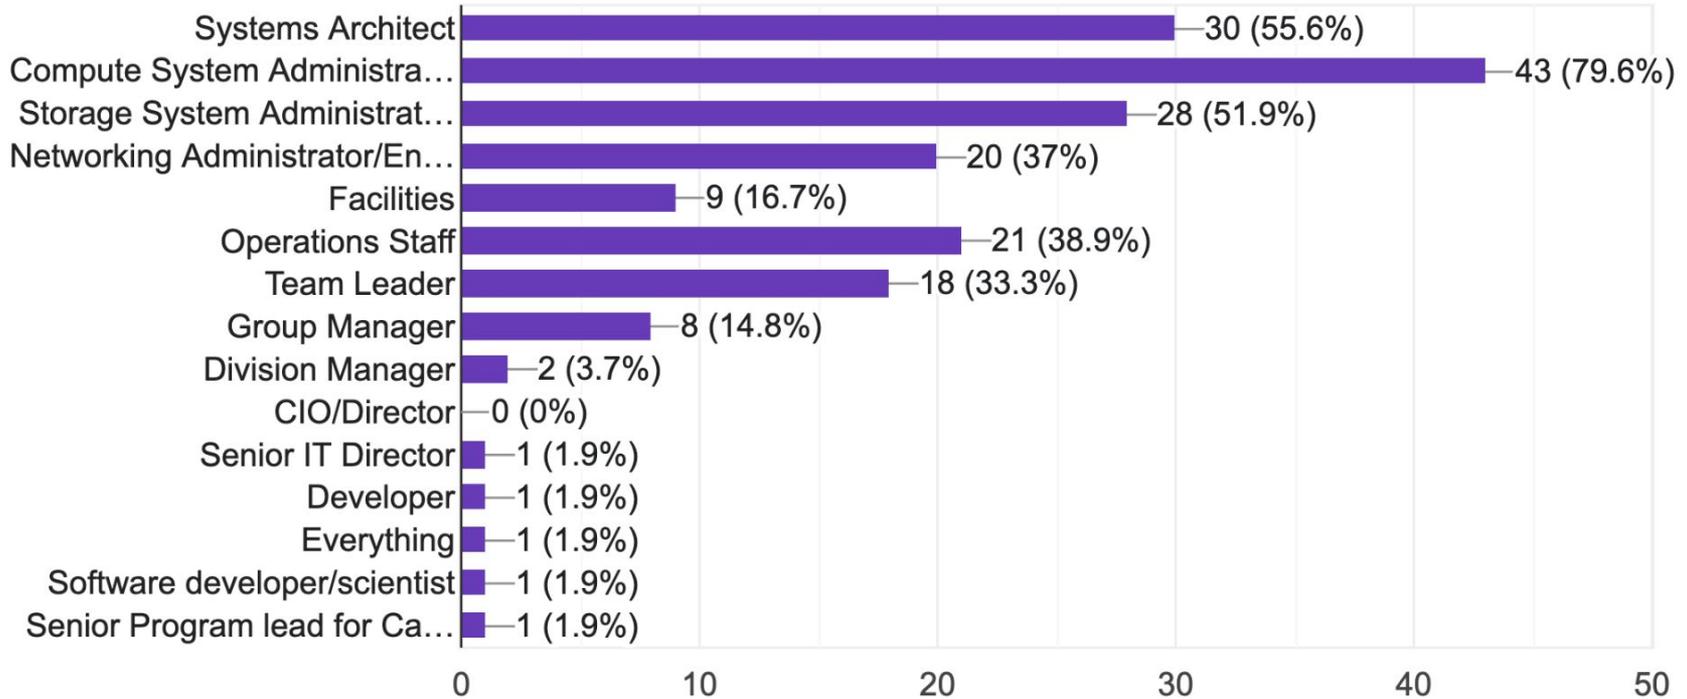
How is your HPC infrastructure funded (select all that apply)?

52 responses

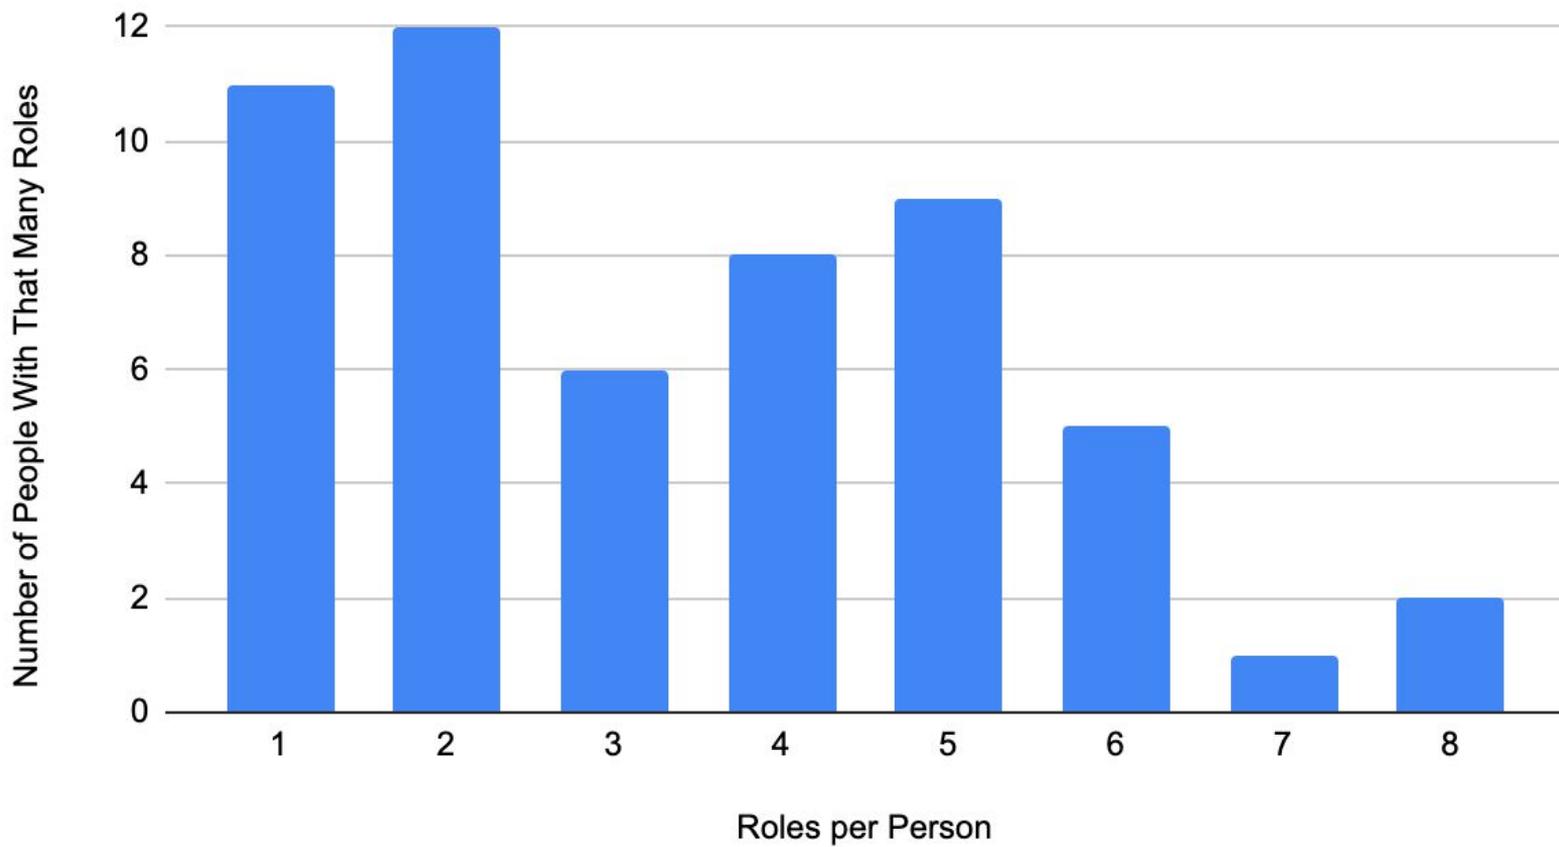


What are your roles at your institution (select all that apply)?

54 responses

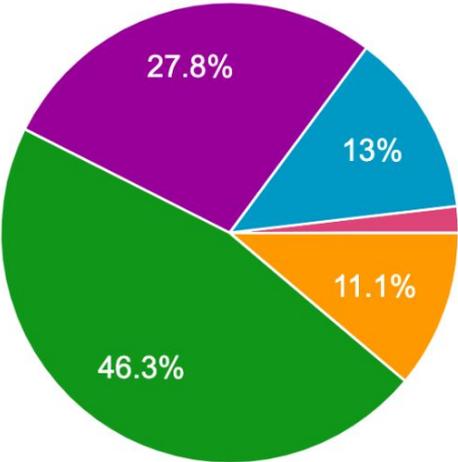


Number of Roles Per Person



What is your highest level of education?

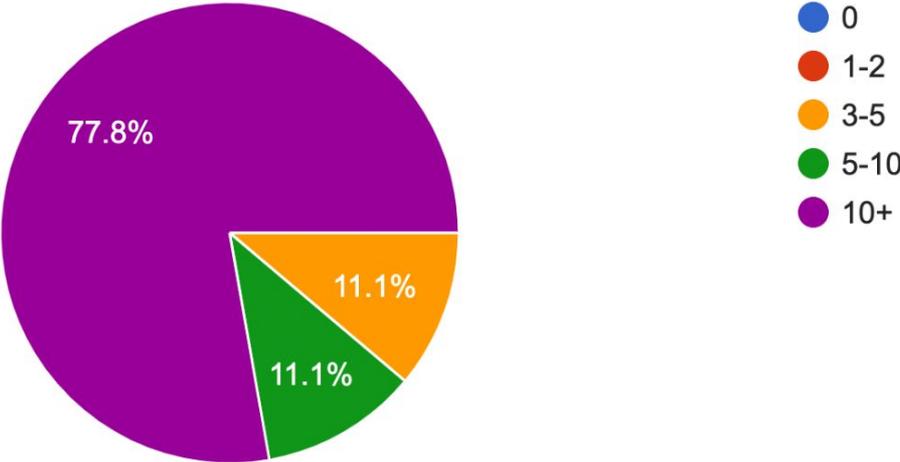
54 responses



- High School Diploma
- Trade School/Technical Certification
- Some College Classes
- Bachelor's Degree
- Masters Degree
- PhD
- All But Dissertation on PhD

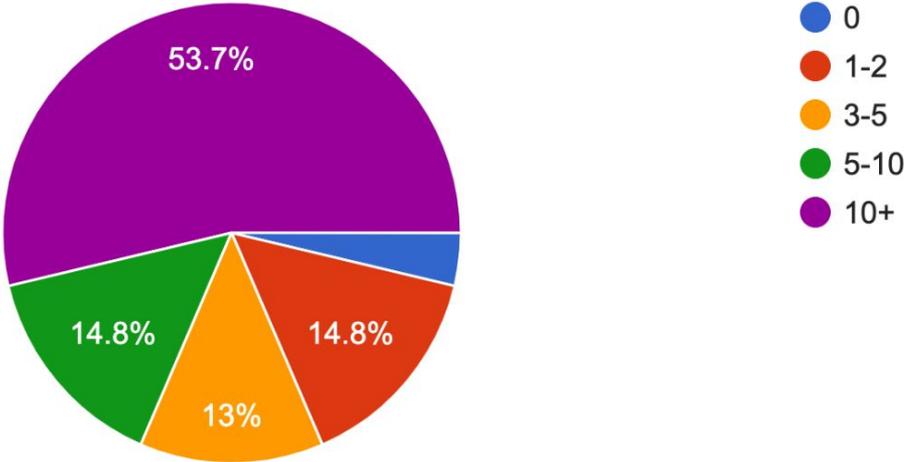
How many years of systems (both HPC, non-HPC, managerial) experience do you have?

54 responses



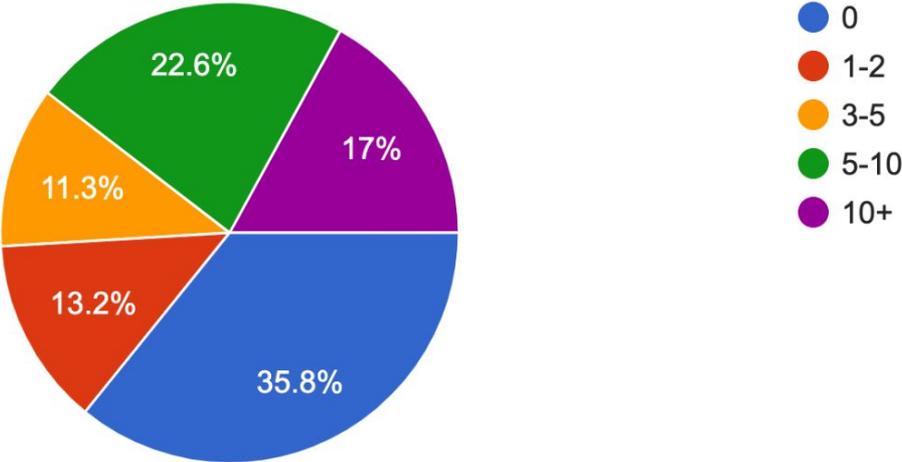
How many years of HPC systems (managerial as well) experience do you have?

54 responses



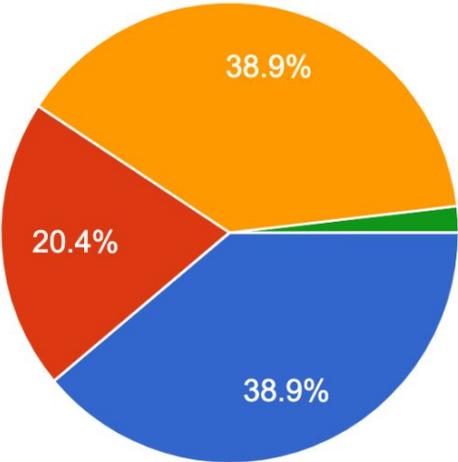
How many years of managerial experience do you have?

53 responses



Do you work from home, at the office, or a hybrid?

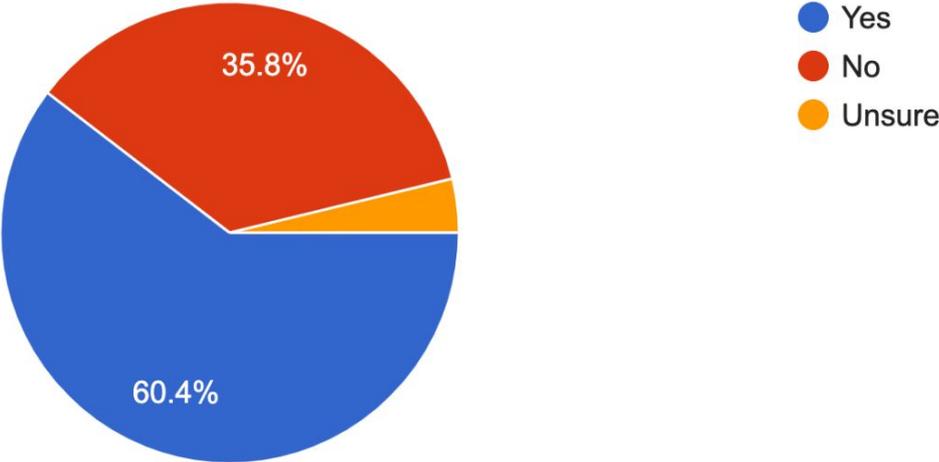
54 responses



- I work from home/remotely most of the time
- I work at the office most of the time
- I work a combination of remotely and at the office
- 100% remote

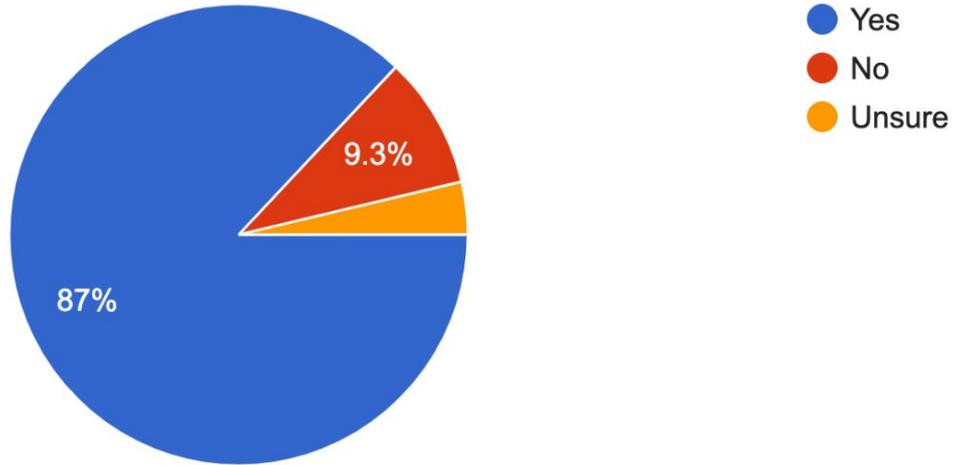
I am a paying member of the ACM SIGHPC SYSPROS Chapter

53 responses

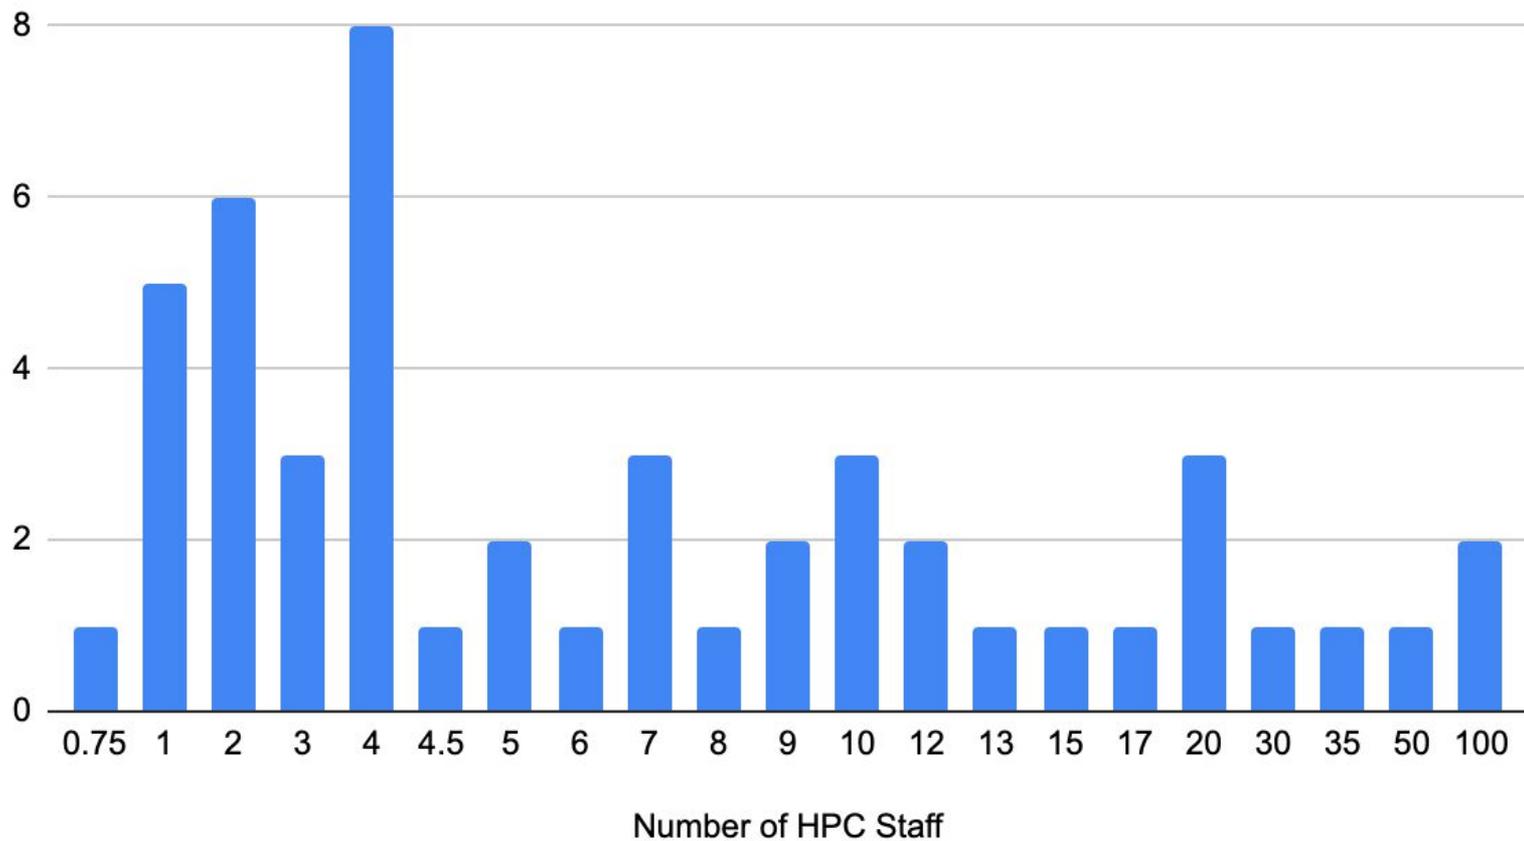


I am a member of the SIGHPC SYSPROS Slack Team

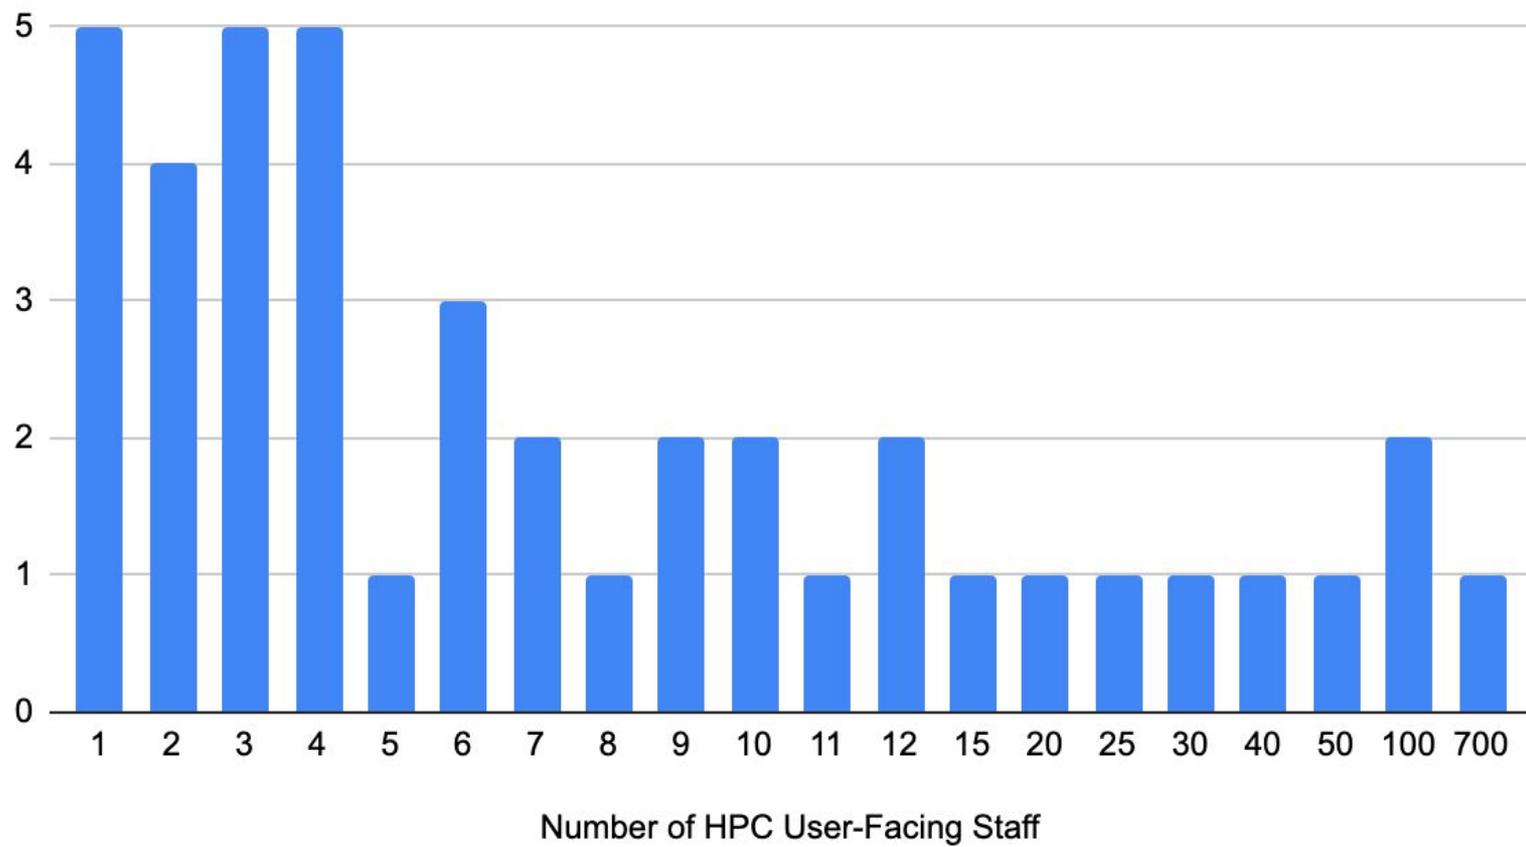
54 responses



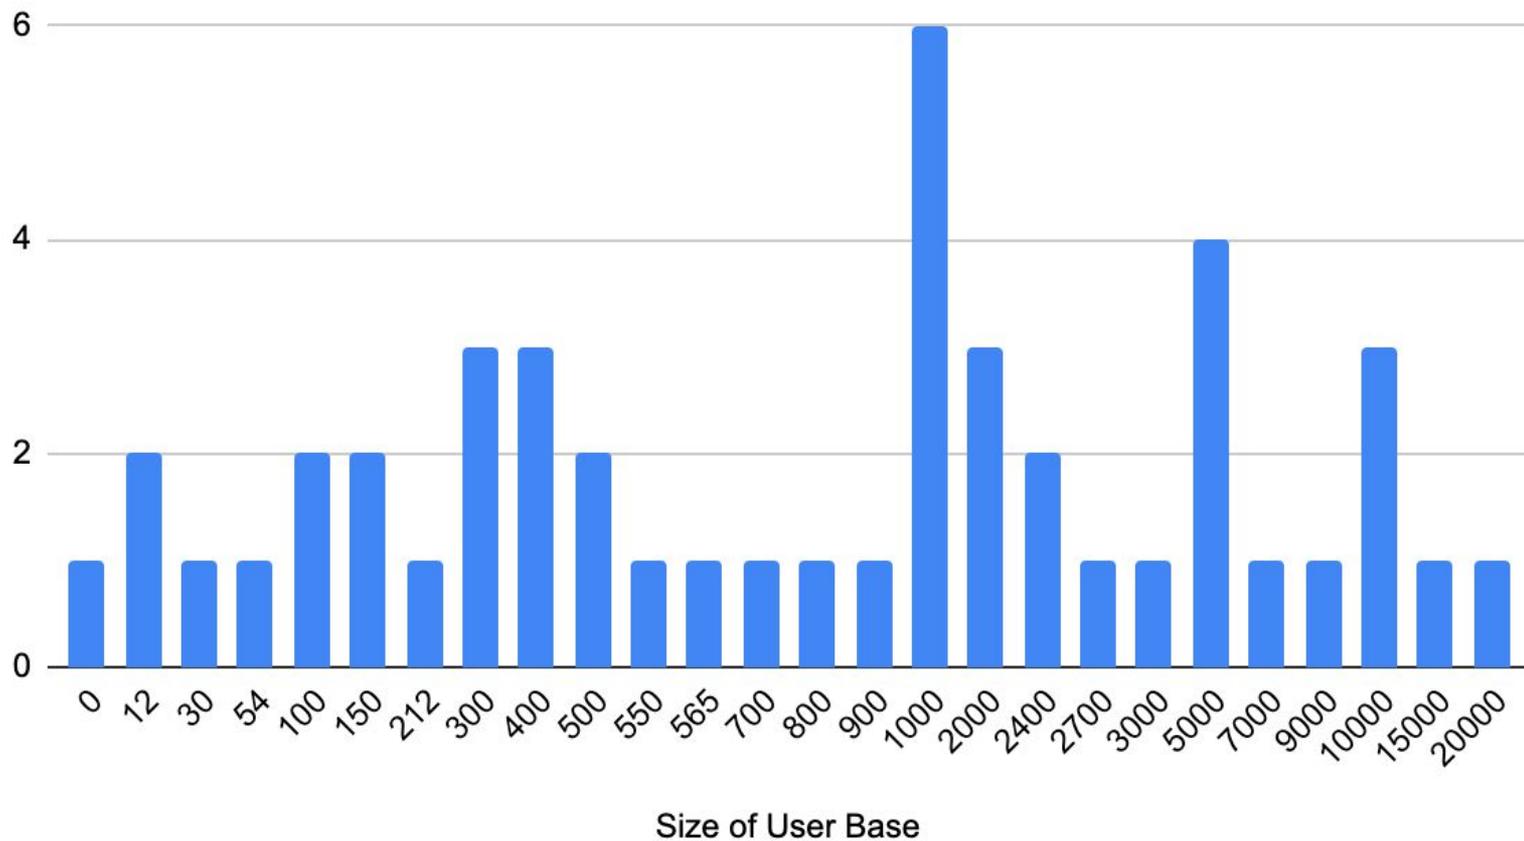
HPC Systems Staff



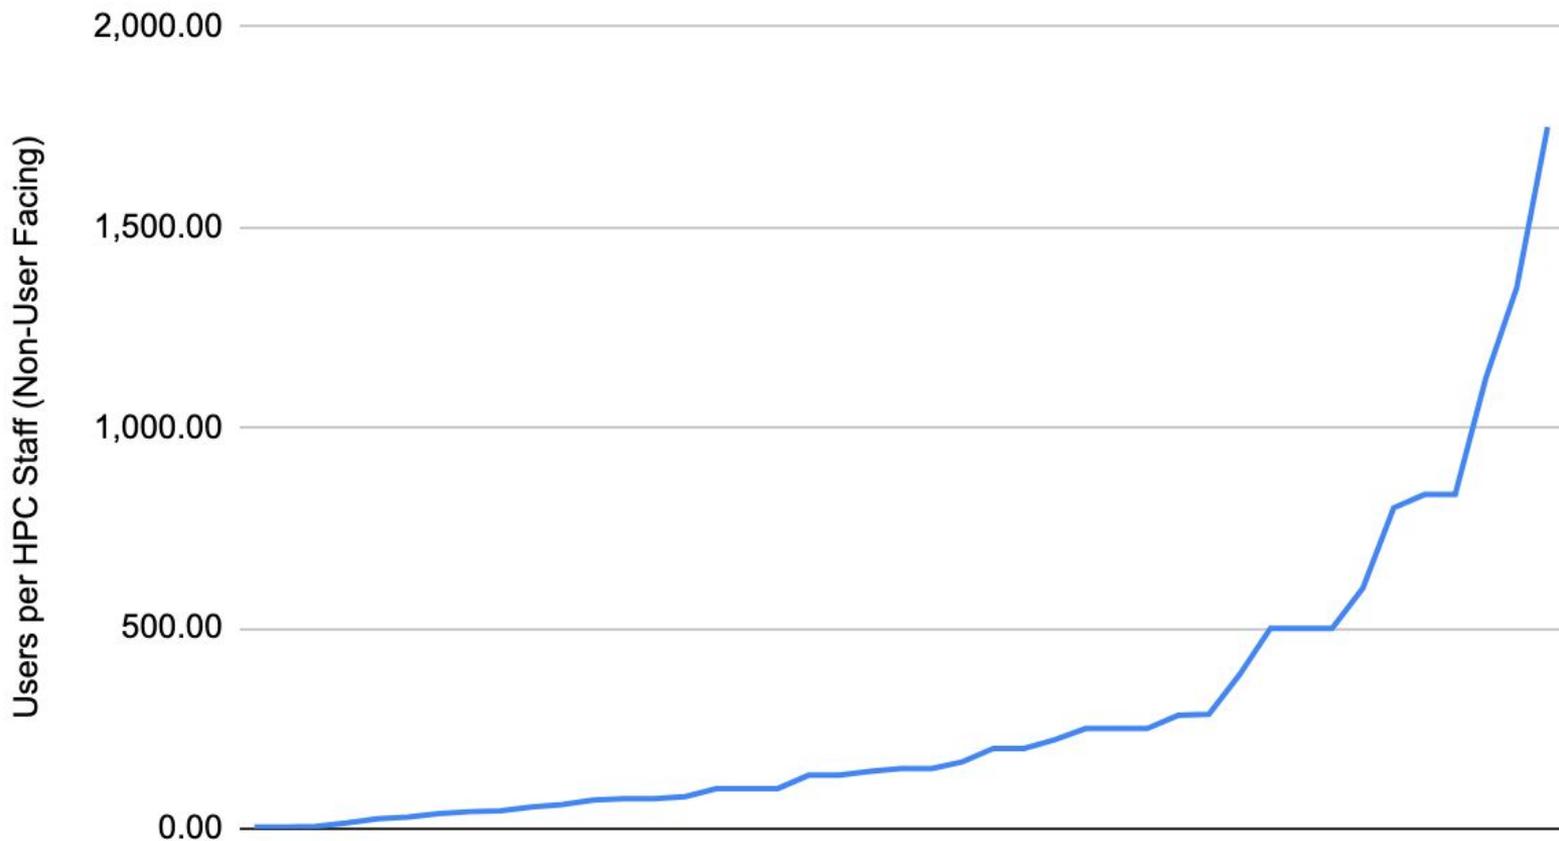
HPC User-Facing Staff



User Base

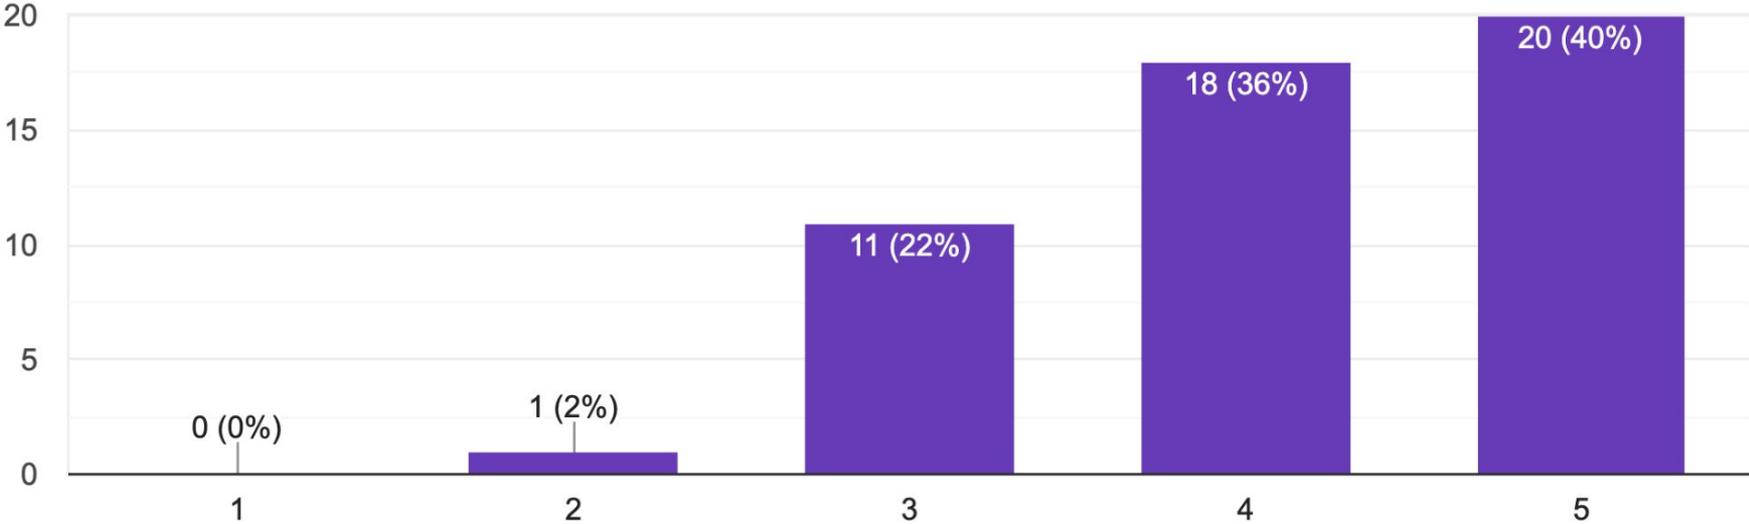


Ratio Users/Systems Staff



How easy is it to find staff to hire?

50 responses

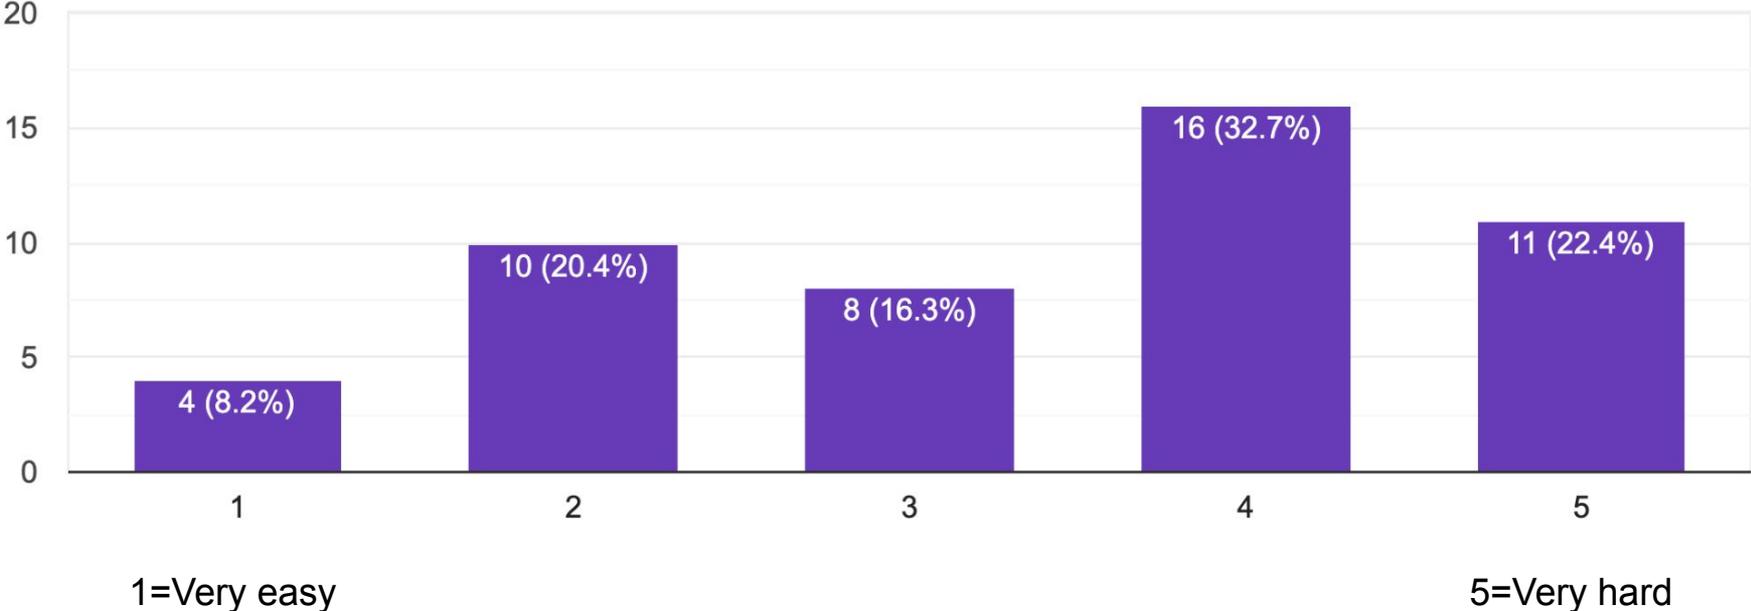


1=Very Easy

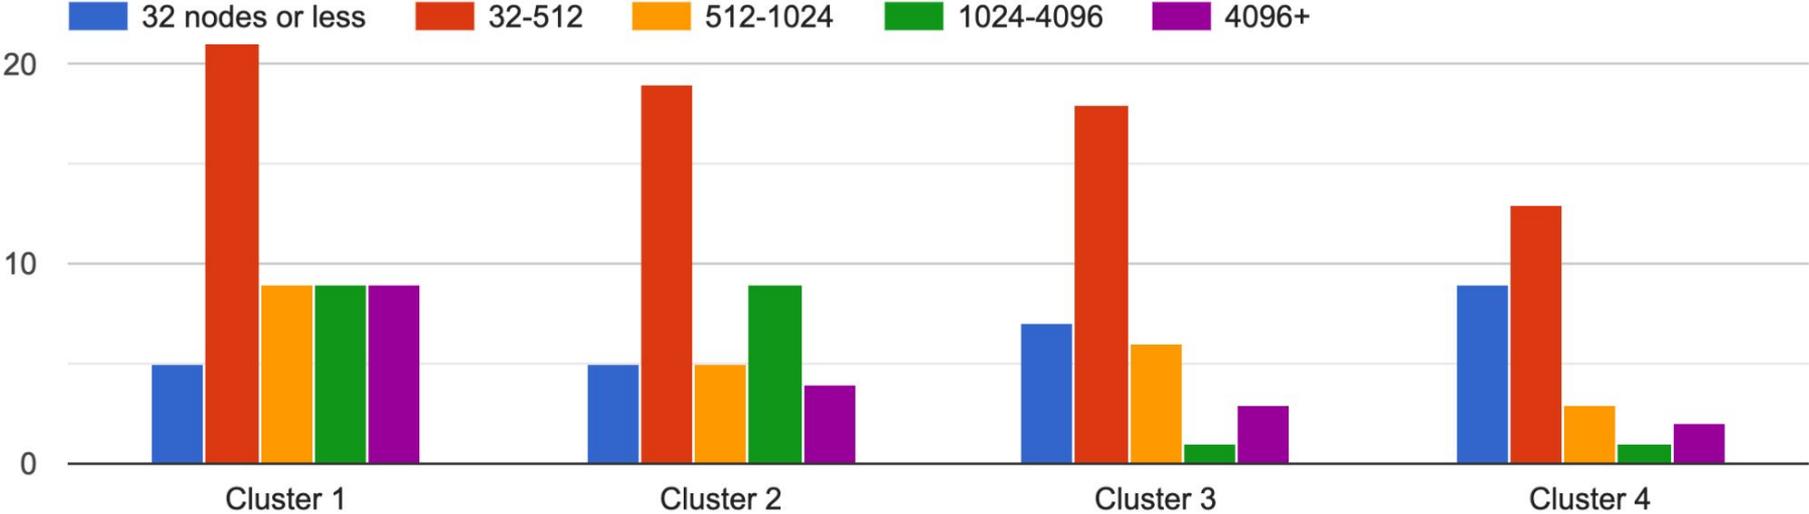
5=Very Hard

How easy is it to find places to get HPC Systems training?

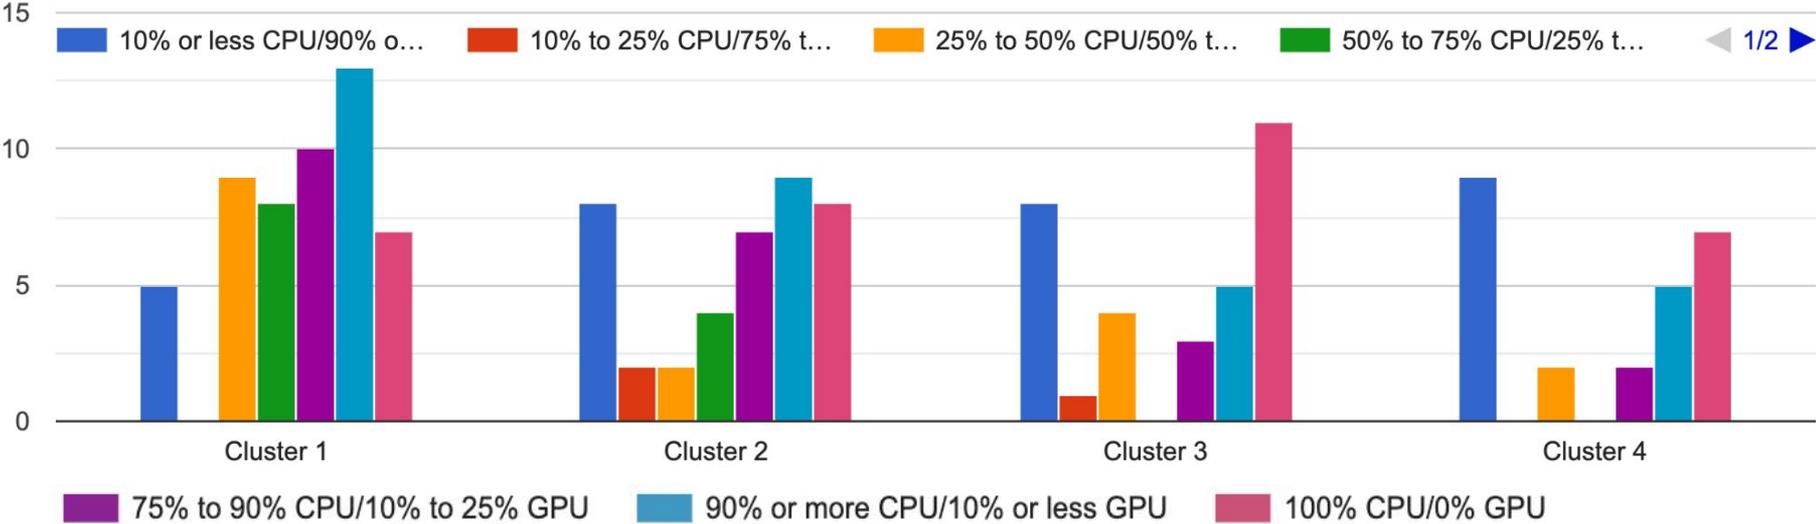
49 responses



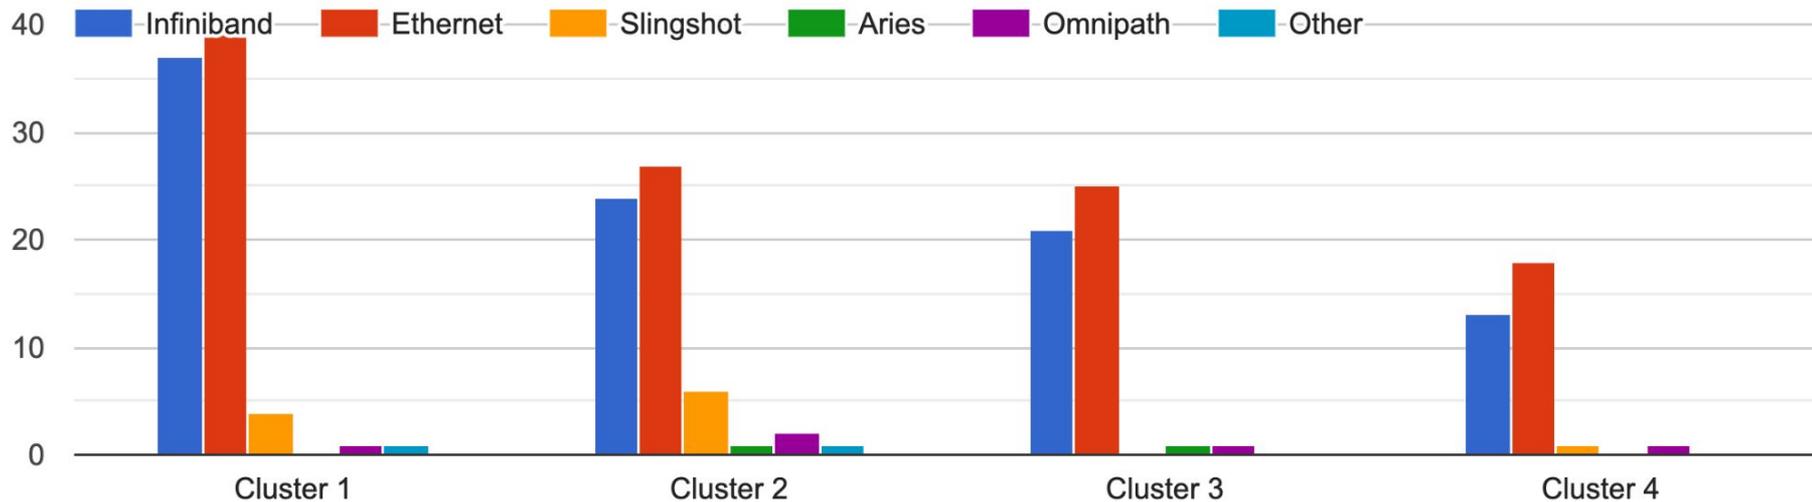
What are the sizes in nodes of the top four on-premise compute clusters at your institution?



What percentage of the CORES are CPU vs GPU in the top four on-premise computer clusters at your institution?

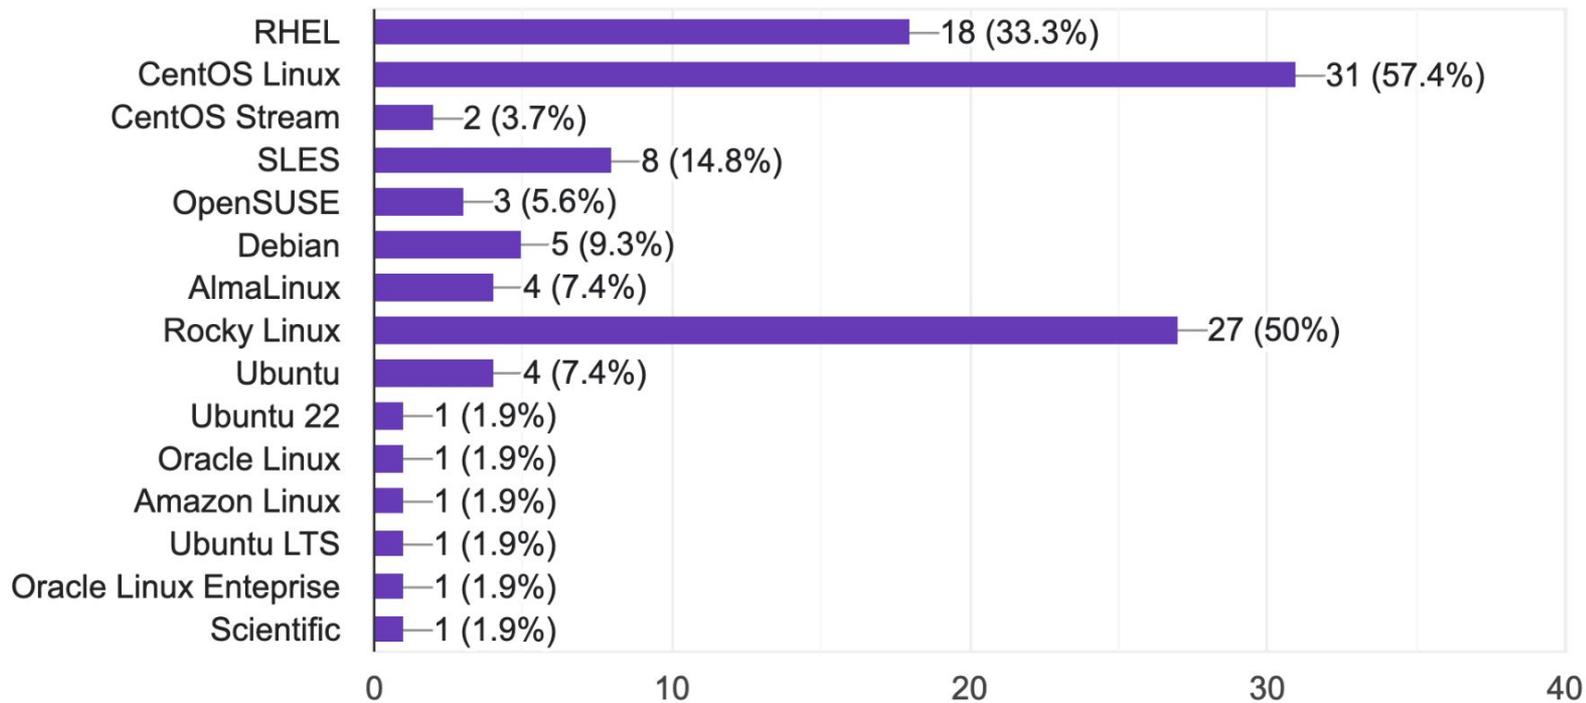


What networking technologies does your site use on the top four on-premise computer clusters (select all that apply)?



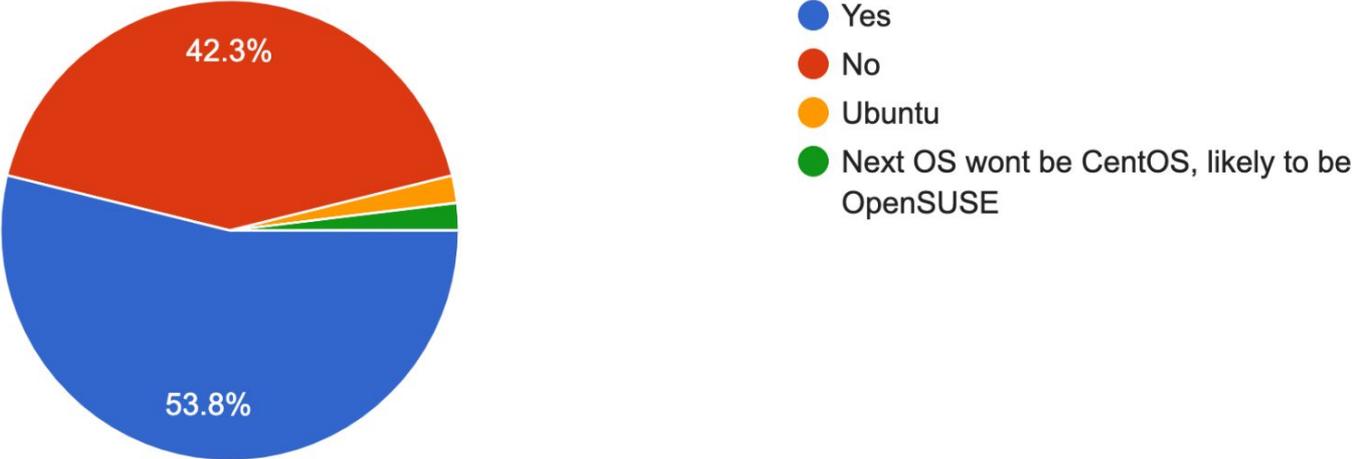
What operating systems do you run at your site on your compute nodes (select all that apply)?

54 responses



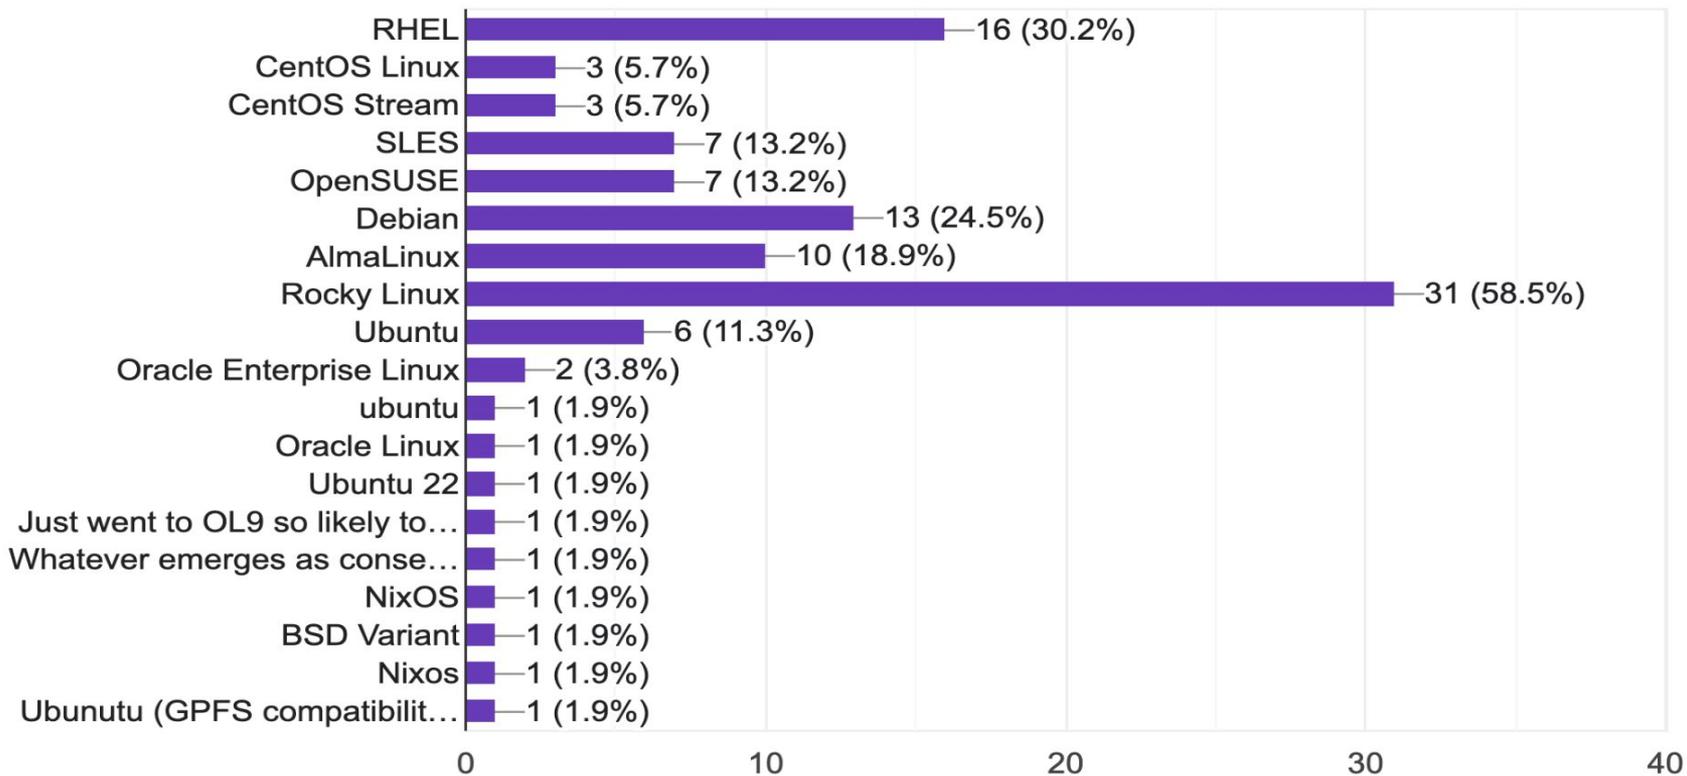
Have you recently decided to run an operating system on an HPC compute system when you haven't run that operating system on an HPC compute system in the recent past?

52 responses



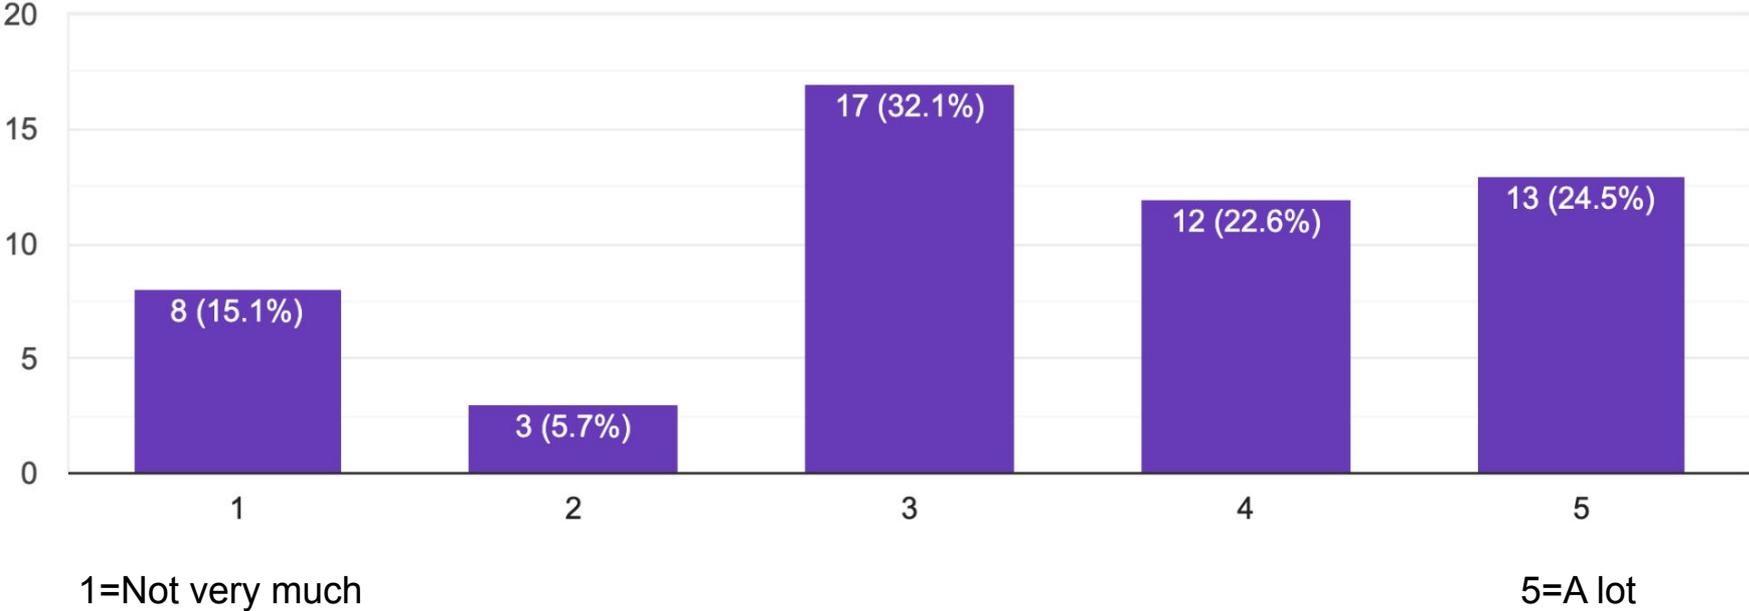
What operating systems are you considering running on your HPC compute systems (select all that apply)?

53 responses



How much does hardware support impact your choice of operating systems?

53 responses



Workload managers/Job Schedulers/Resource Managers

TORQUE
3.8%

HTCondor

3.8%

IBM Spectrum LSF
6.4%

Kubernetes
10.3%

Maui

1.3%

Moab

2.6%

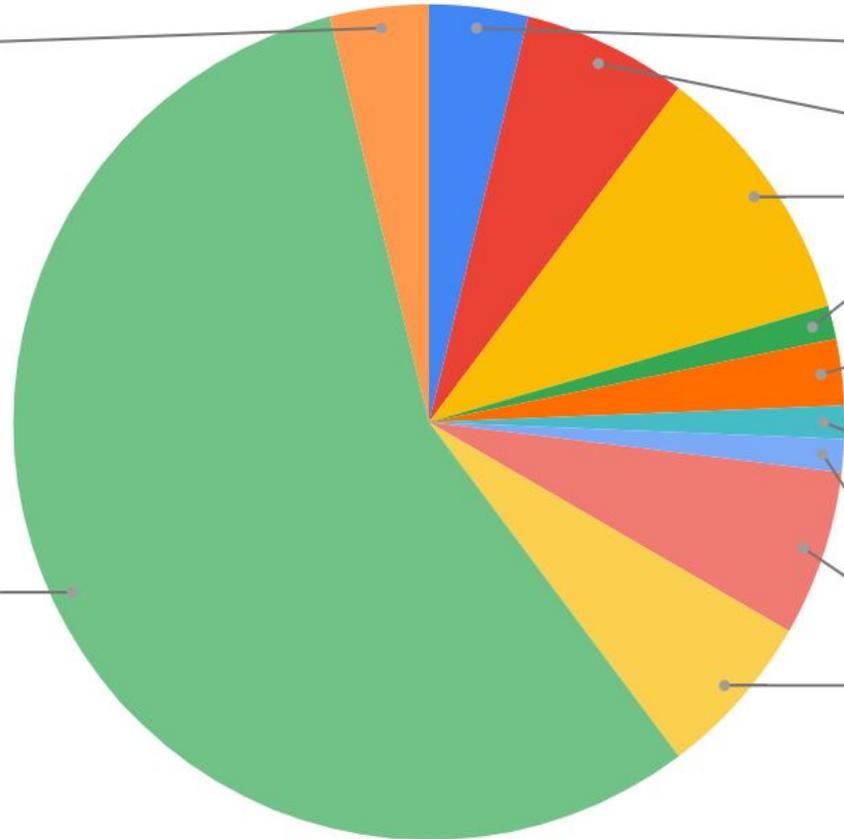
OpenPBS
1.3%

Openstack
1.3%

PBSPPro
6.4%

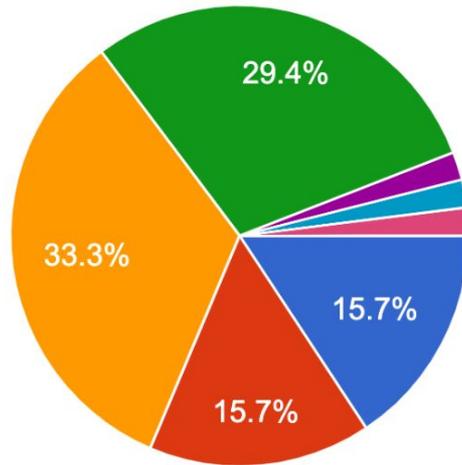
SGE/UGE/AGE/OGE
6.4%

SLURM
56.4%



How many applications does your site support?

51 responses



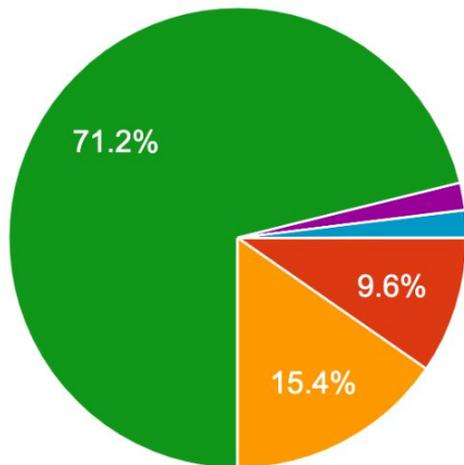
- Our staff supports a select few applications (less than 10)
- We have several staff supported applications (10 to 50)
- We have lots of staff supported applications (over 50)
- Our staff supports whatever users ask...
- Depends on your definition of "support..."
- Our staff supports whatever users ask...
- We support whatever, to differing level...

How Many Applications - Other Responses

- Depends on your definition of "support", but we put best effort into applications at the 10-50 range.
- Our staff supports whatever users ask us to support, where "support" usually just means install and break/fix. Deeper application support (usage guidance, optimization, etc) is very limited.
- We support whatever, to differing levels of support depending on how widely used

What applications does your site allow on the compute clusters?

52 responses



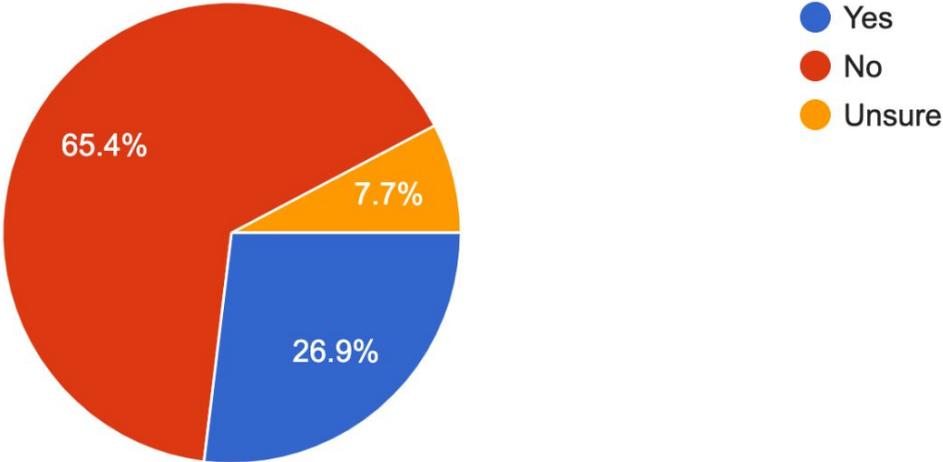
- We allow only a few applications and tightly control what is run.
- We have a number of applications all centered around a particular field of st...
- We have a number of applications from a wide variety of fields of study.
- We allow users to run whatever they like and do not control what is run (outside...)
- We allow users to run whatever they li...
- compute nodes do not have applicatio...

What Applications Do You Allow - Other Responses

- We allow users to run whatever they like and do not control what is run (outside of security or abuse concerns), all centered around a particular field of study -- medical imaging.
- Compute nodes do not have applications, they are all run in login nodes

Does your site track the energy used by HPC workloads?

52 responses



What is your site doing to reduce the energy consumption of HPC workloads?

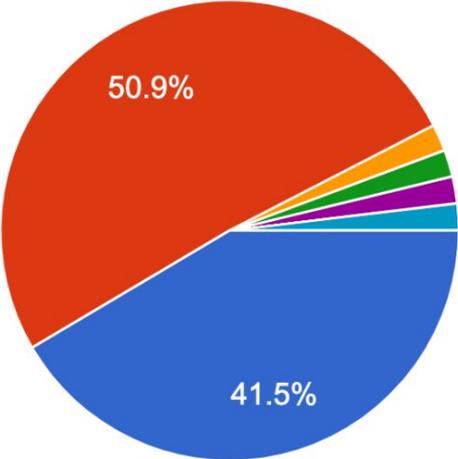
- Considering water cooling options
- Adjusting fan speeds (It surprisingly saves a lot of power!)
- Water cooled racks, lower power CPU chips
- Nothing presently but looking at ideas for the future
- Running them as fast as possible and leveraging frequency scaling/turbo for applications that don't use all cores
- CPU throttling idle cores
- We use free-air cooling and direct liquid cooling with high density racks
- More efficient CPUs, moving workloads to GPUs
- Working to retire older/less energy-efficient hardware in place of more efficient equivalents
- LEED certified data center, powered by hydro. I'm researching more.
- Switching to ARM for as much CPU compute as possible
- Regular hardware upgrades include some inherent efficiency improvements. That aside, no intentional energy reduction efforts.
- Energy consumption is not a concern currently.
- Efficient architecture, choosing sites based on efficient cooling availability
- It's not possible when clusters are at 99% utilization all the time.

What is your site doing to reduce the energy consumption of HPC workloads?

- Large increase in recently deployed hardware's power efficiency, work with facilities management to use water and air cooling more efficiently, work with users to improve job efficiency, begin work on energy use per job reporting
- Renovating/expanding the data center with better HACs to allow increasing operating temps
- Nothing hpc-specific outside of general data center efficiency
- Moving to all liquid-cooled infrastructure
- Keep infrastructure utilization as high as possible to minimize energy lost to idle hardware; also additions of liquid cooling
- We are starting to have conversations with researchers about options.
- Trying to make systems more efficient
- During times of high summer heat and electrical demand when we are asked by facilities to cut power use, we work with projects and users to see what resources could be suspended for a short time. This might only be 5-10% of a resource, but watts are watts.
- We are using processed chilled water to cool IT racks
- New hardware, improvements in pipelines, user training

Has your site implemented water cooling for compute resources or is planning to do so in the near future (5 years)?

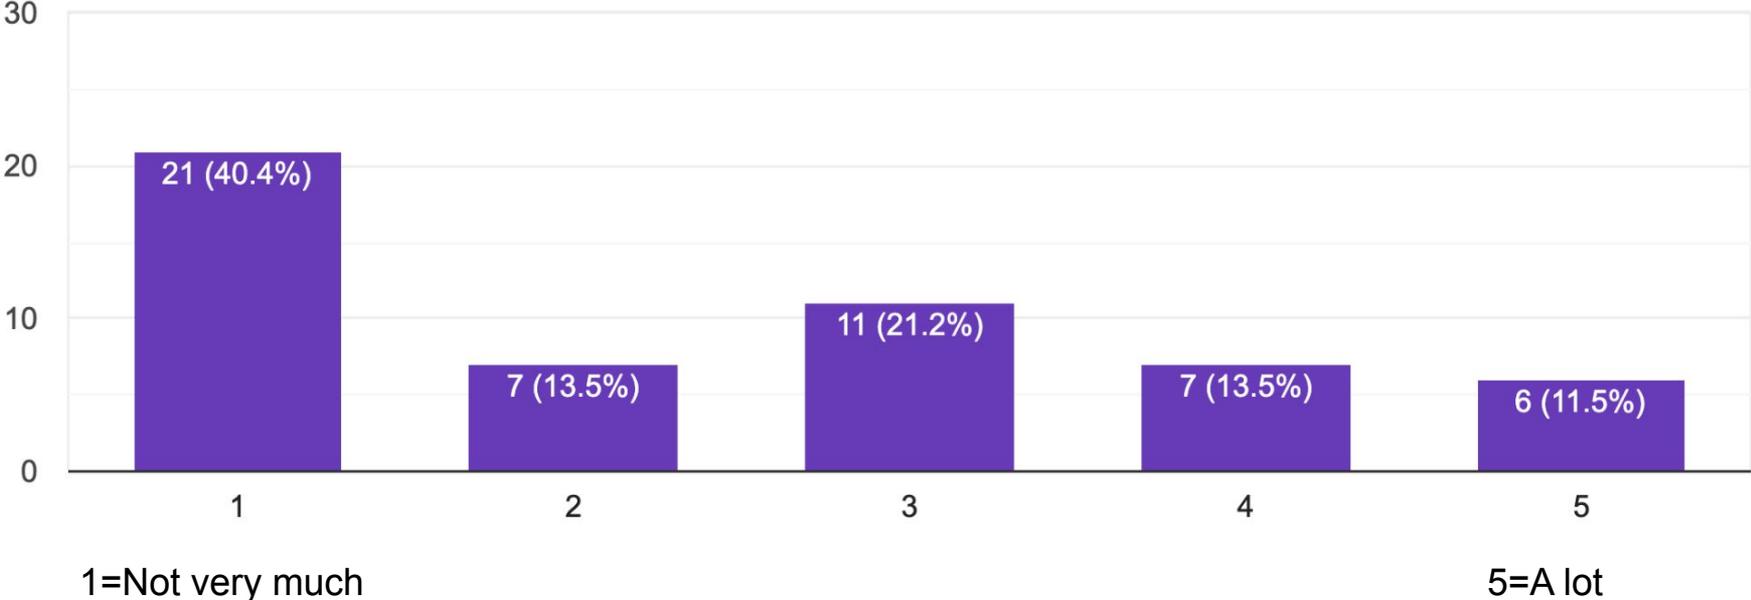
53 responses



- No
- Yes
- it is POSSIBLE that in 5 years we may have begun water cooling (need is burgeoning), but not on any planned h...
- definite maybe
- We are investigating water cooling and full immersion for GPU deployments
- Data center built with the assumption majority of load would be direct water...

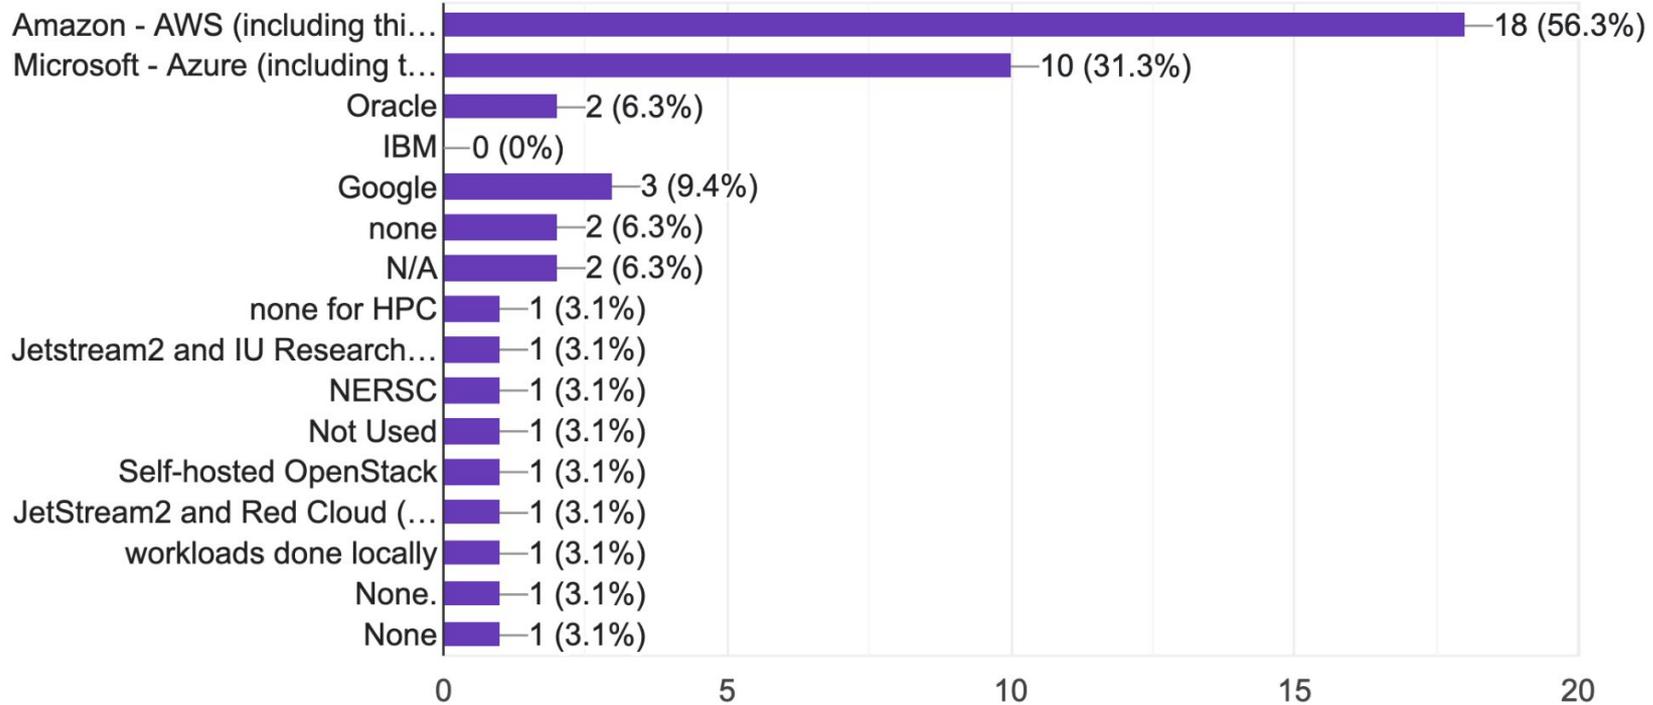
How has the recent rise of AI models, particularly large language models such as ChatGPT, impacted your hardware procurement decisions?

52 responses



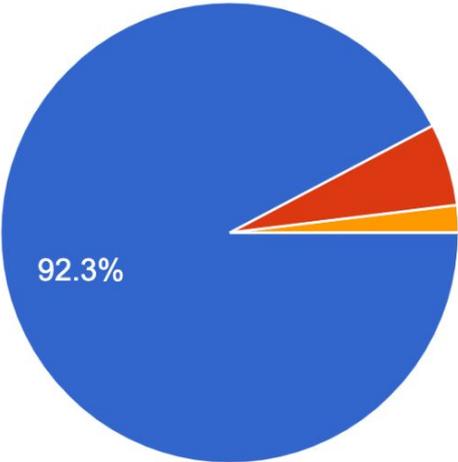
What clouds services do you use for HPC workloads (select all that apply)?

32 responses



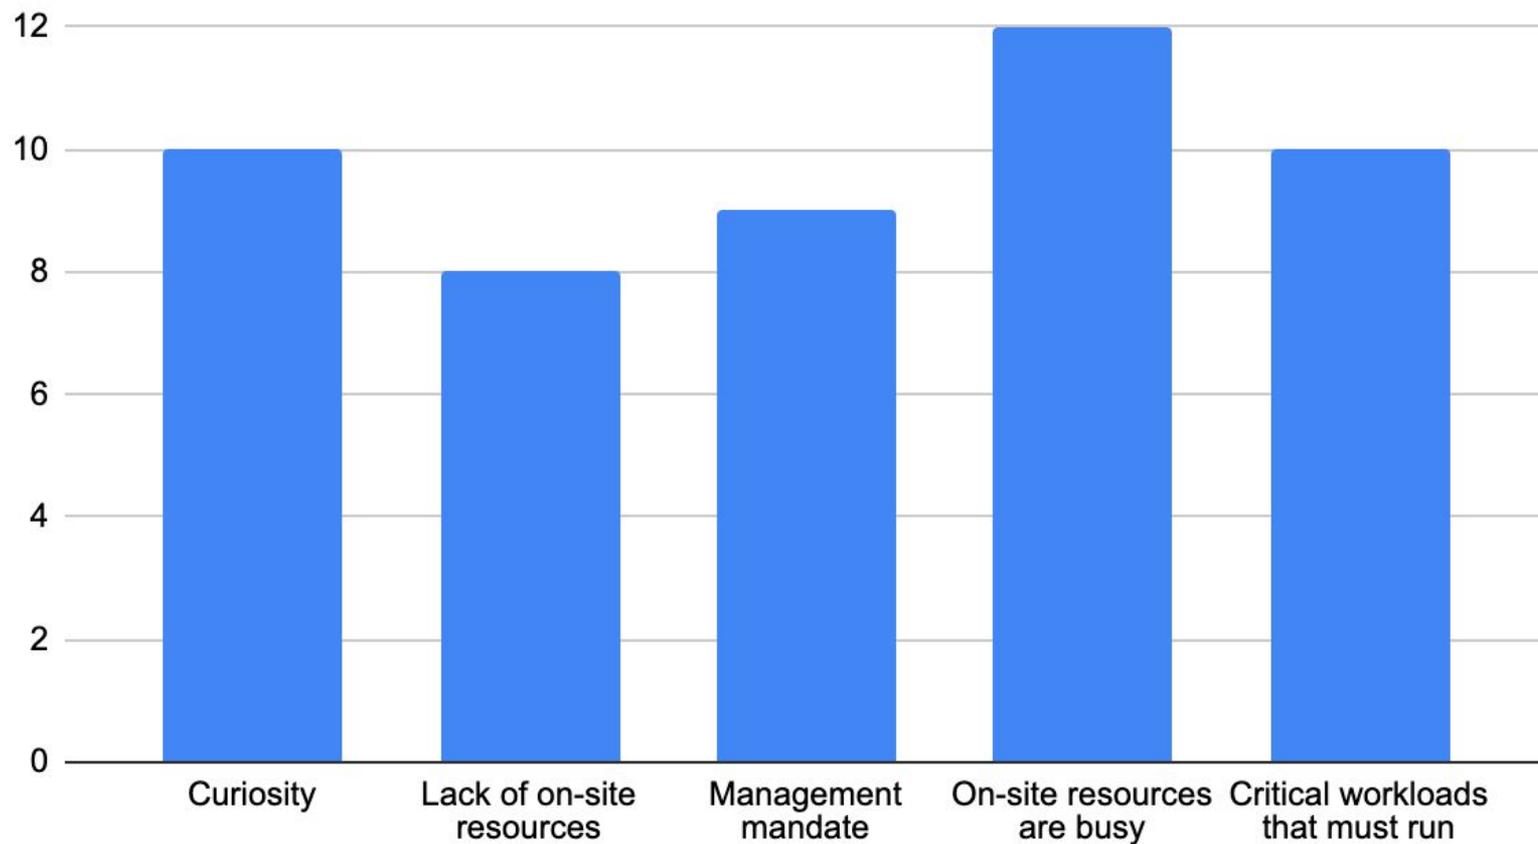
What percentage of the HPC workloads at your institution are run in the cloud vs. on-premise?

52 responses



- 10% or less cloud/90% or more on-premise
- 10% to 25% cloud/75% to 90% on-premise
- 25% to 50% cloud/50% to 75% on-premise
- 50% to 75% cloud/25% to 50% on-premise
- 75% to 90% cloud/10% to 25% on-pre...
- 90% or more cloud/10% or less on-pr...

Driving Cloud Usage

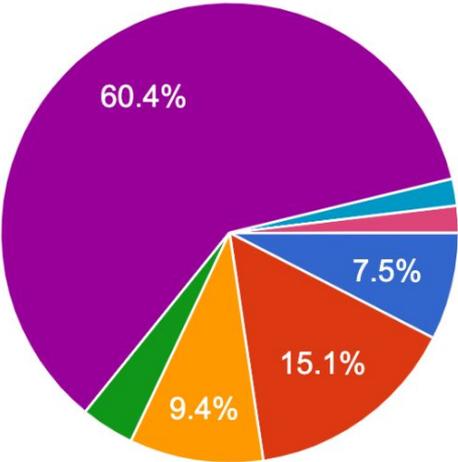


What is driving your institution to run workloads in the cloud?

- Ease of accessing data
- Wider variety of compute technologies
- People think cloud is sexy and one CSuite went to a Gartner conference so we needed a cloud strategy
- Services pending migration to on-prem resources
- GPU shortage
- Diversifying due to reaching power consumption limits on premise
- Reproducible pipelines
- Nothing. We strive to be fixed cost and competitive.
- Workloads more suited to cloud environments
- Potential to share pipelines with other organizations
- Batch scheduling isn't suitable for all workloads
- There are other management reasons why we cannot or are not allowed to provide services in house.

Does your site "burst" HPC workloads to the cloud?

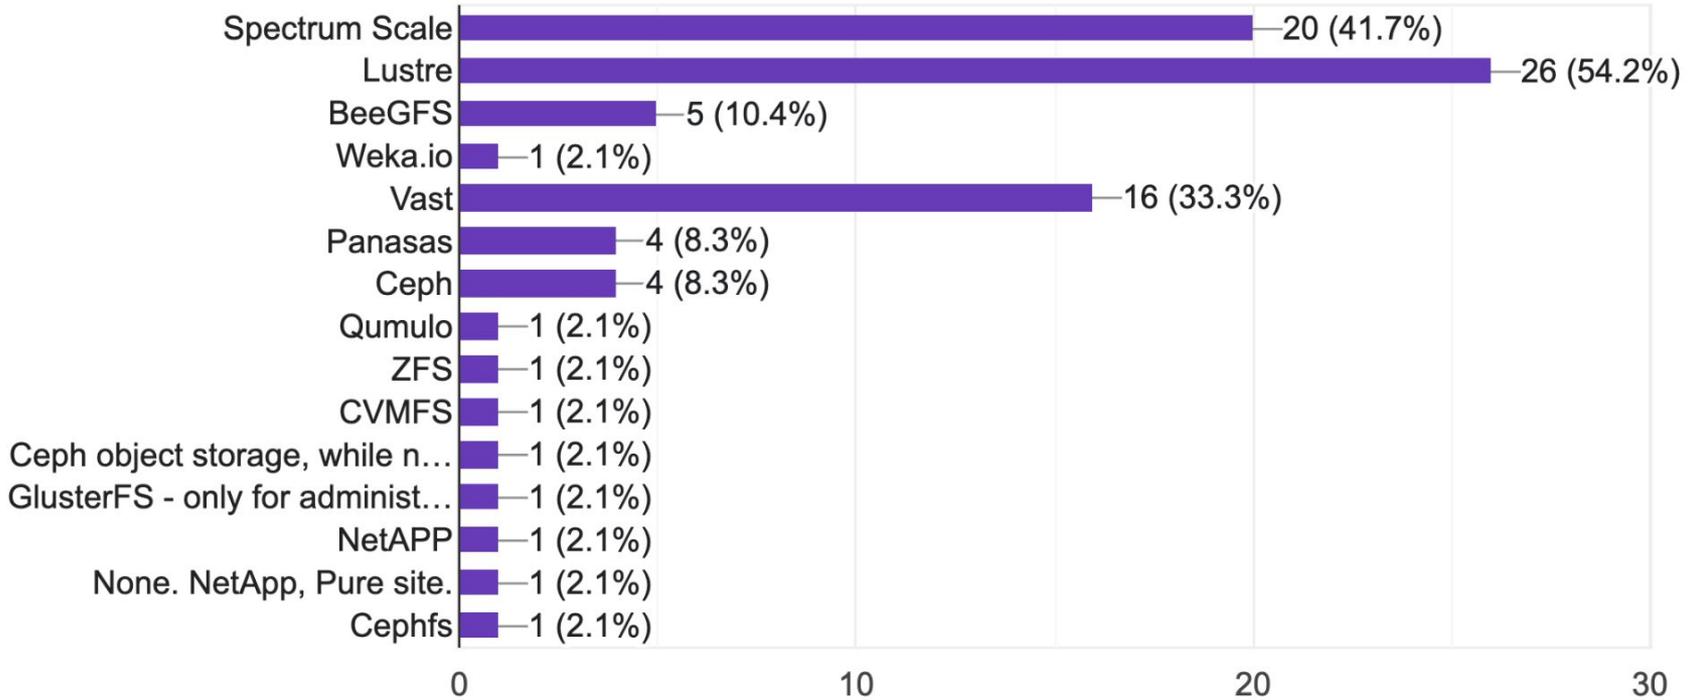
53 responses



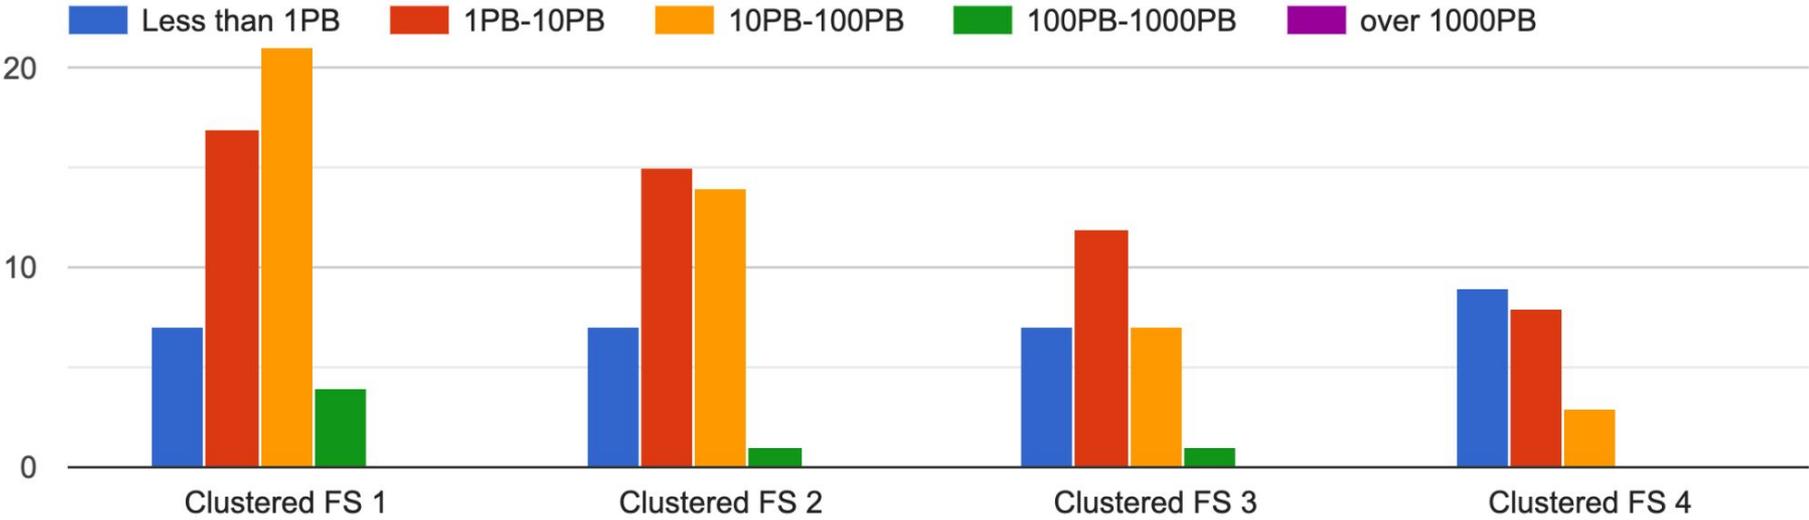
- Yes
- Our site would like to, but haven't implemented anything
- Our site tried to implement something, but has not been successful
- Our site tried to implement something, but has given up
- No
- We do not at this time, but other sites...
- We don't, but we could.

What clustered file systems do you have at your institution (select all that apply)?

48 responses

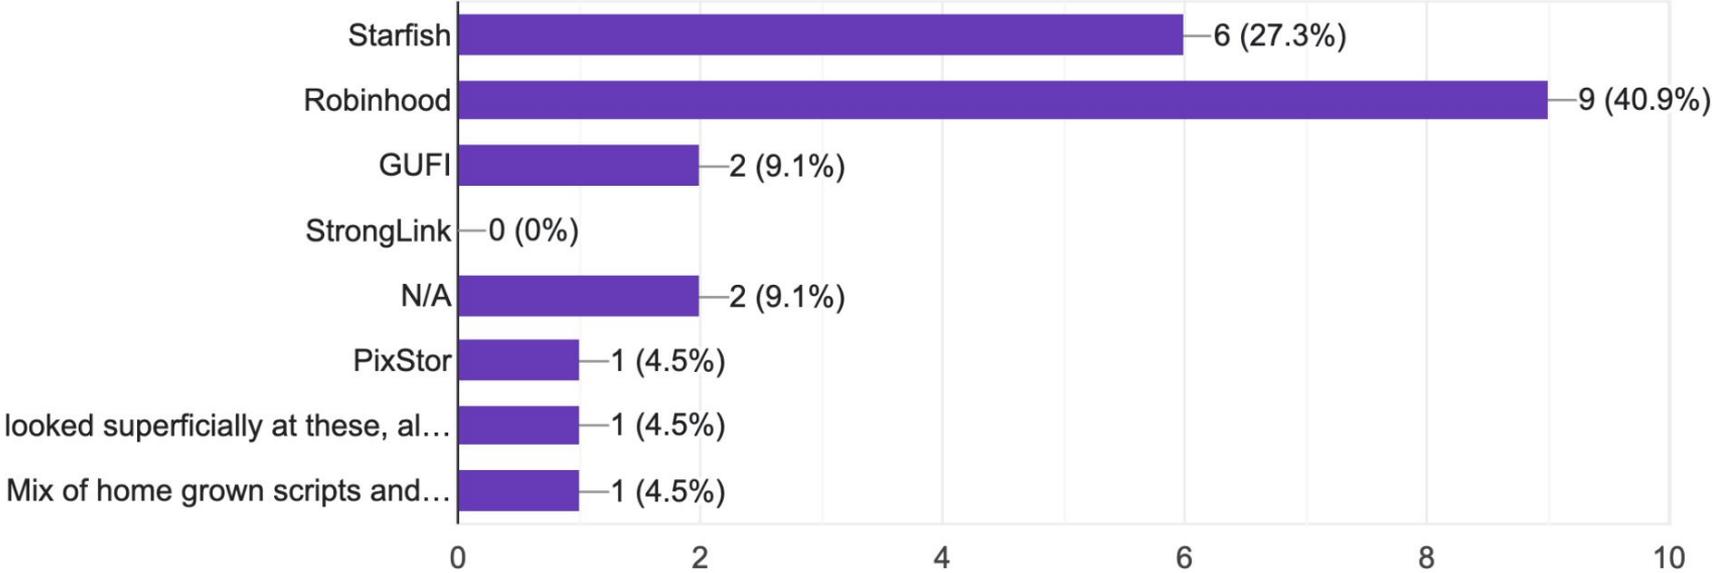


What are the sizes (in PetaBytes) of the four largest on-premise clustered storage systems at your institution?



Do you support management of data on filesystems managed at your site with any of the following software?

22 responses



Is there anything else you'd like to tell us?

- Hi!
- Moving toward VAST + Ceph most likely in very near future, moved to Rocky from Fermi SL when Rocky first released, lots of XSEDE / ACCESS usage ("cloud" == "someone else's resources"), lots of OGS (is that what "SGI" was supposed to be under schedulers?)
- Getting funding for permanent staffing is impossible, so we have to rely on temporary contract labor.
- Thanks for being awesome.
- Thank you for providing this Slack instance!!!
- Love the slack channels!
- Tired... so tired.
- looking forward to the results!
- I am new so my answers may not all be accurate