Karlsruher Institut für Technologie

SCC

Survey on scientific needs for the upcoming High-Performance Computer HoreKa-2

Results

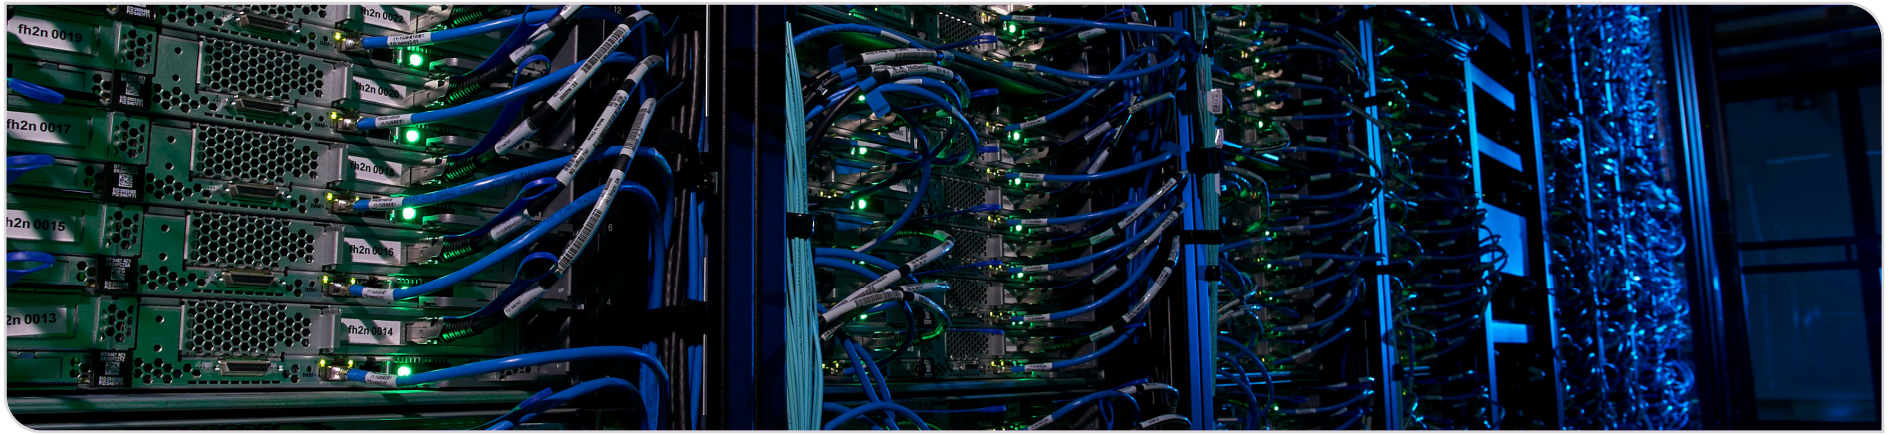

General framework

General framework

- Survey was conducted on Indico
- 182 PIs & PoC were invited
- PIs & PoC received an invitation via email
- 3 emails were sent in total (Invitation, Reminder, Last Call)

→ 37 PIs & PoC responded to the survey

Q0: DFG Subject Area of Respondents

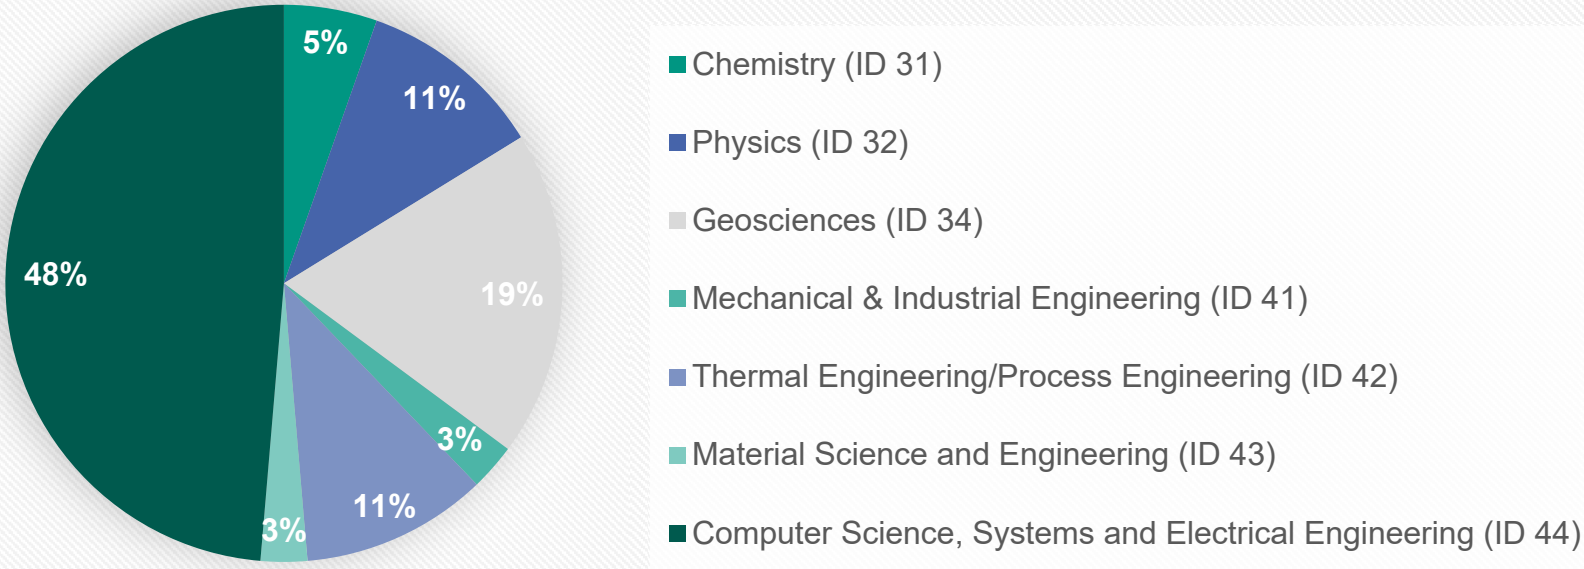

List of Questions (∑15)

- Requirements for the high-performance computer:
  - Q1: To how many nodes does your job scale?
  - Q2: Main memory per task / MPI process
  - Q3: Job profile
  - Q4: Does the program use checkpointing?
  - Q5: Number of files
  - Q6: Storage space (GByte)
  - Q7: Data transfer from/to cluster (GByte)
  - Q8: HPC or AI workloads
  - Q9: Job Profile: CPU & GPU Mix
  - Q10: Porting of Workflows and applications to GPUs
  - Q11: What type of accelerator is needed?
  - Q12a: Can your application run on ARM?
    - Q12b If not, what is the reason?
  - Q13: Used Software
  - Q14: Application support services

Requirements for the High-Performance Computer

Q1: To how many nodes does your job scale?

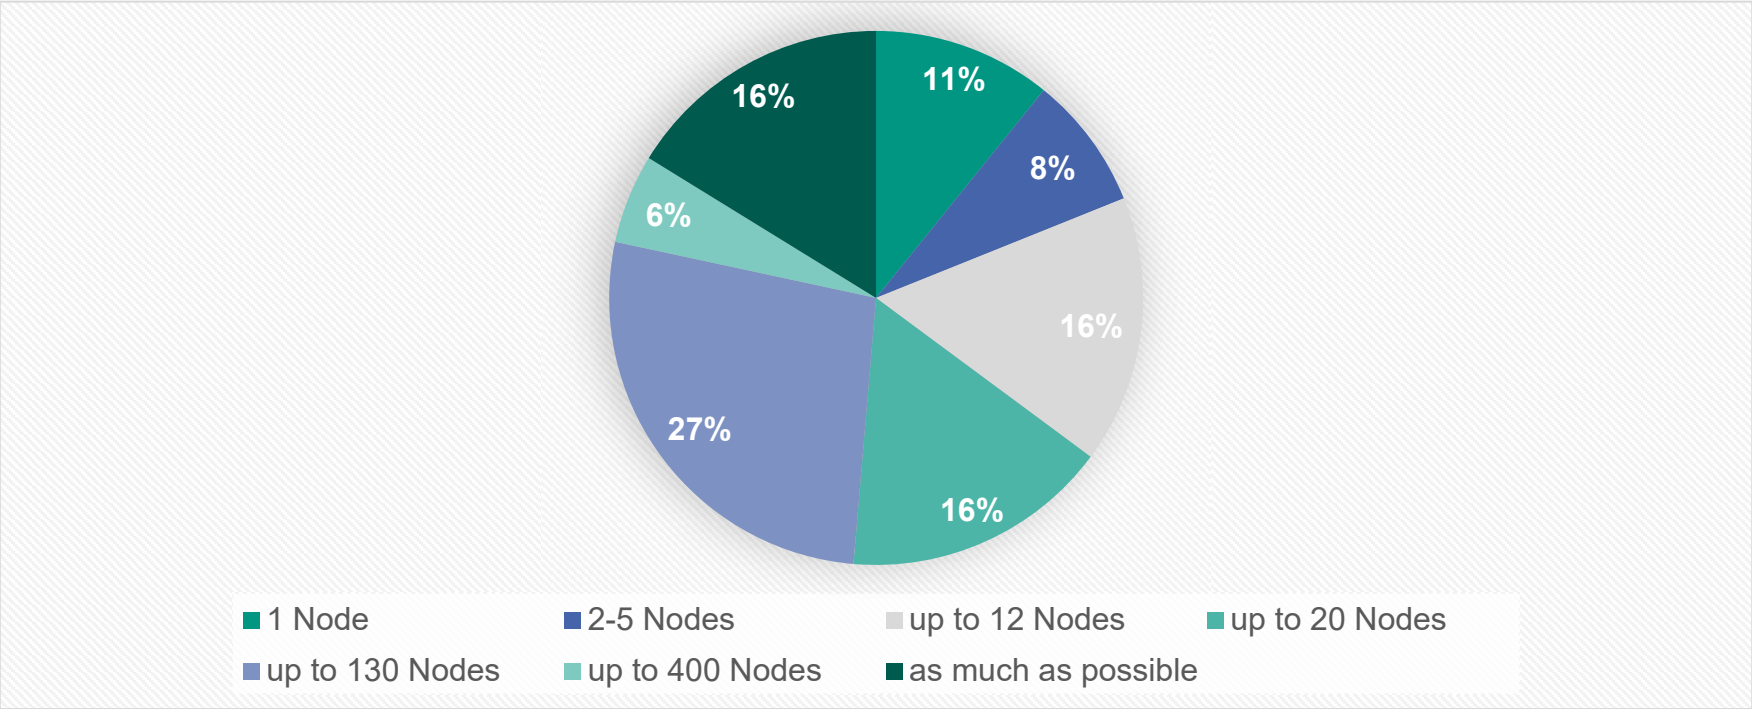

General tendency - Question 1

Prominent answers

- the more the better really..
- codes are typically extremely scalable

Q2: Main memory per task / MPI process

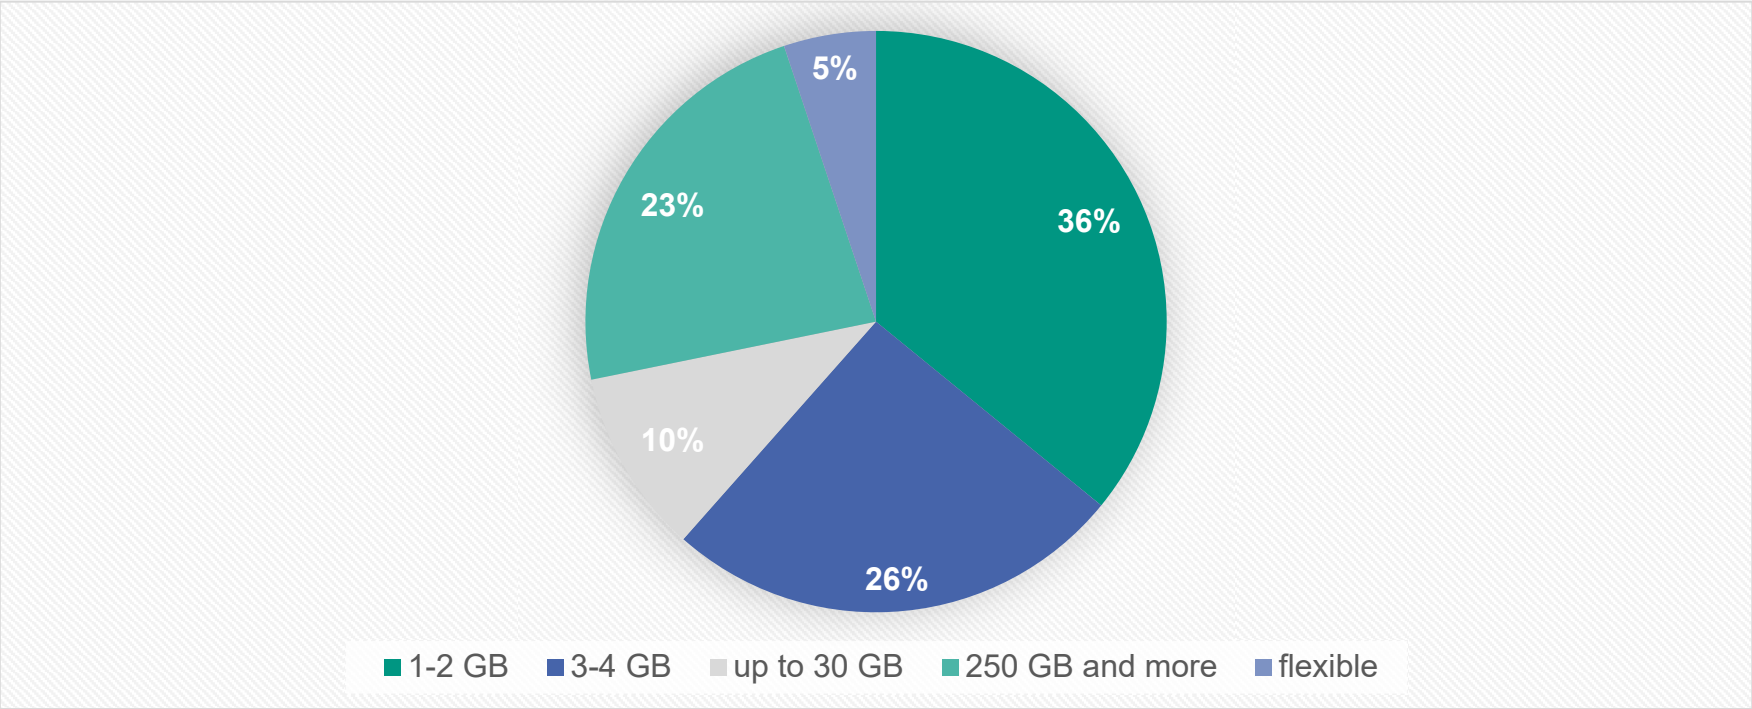

General tendency - Question 2

Prominent answers

- 1-2 GB is usually enough
- Varies on setup and simulation size
- bigger bottleneck is the GPU memory
- Please add more A100 or H100 nodes that support large GPU memory

Q3: Job profile

A. **Memory bound**: 13 (35.14%)

B. **Compute bound**: 13 (35.14%)

C. **Unclear**: 11 (29.73%)

Q4: Does the program use checkpointing?

A. **Yes**: 30 (81.08%)

B. **No**: 7 (18.92%)

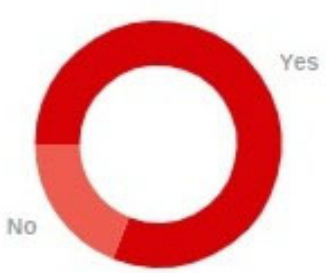

Q5: Number of files (storage system)

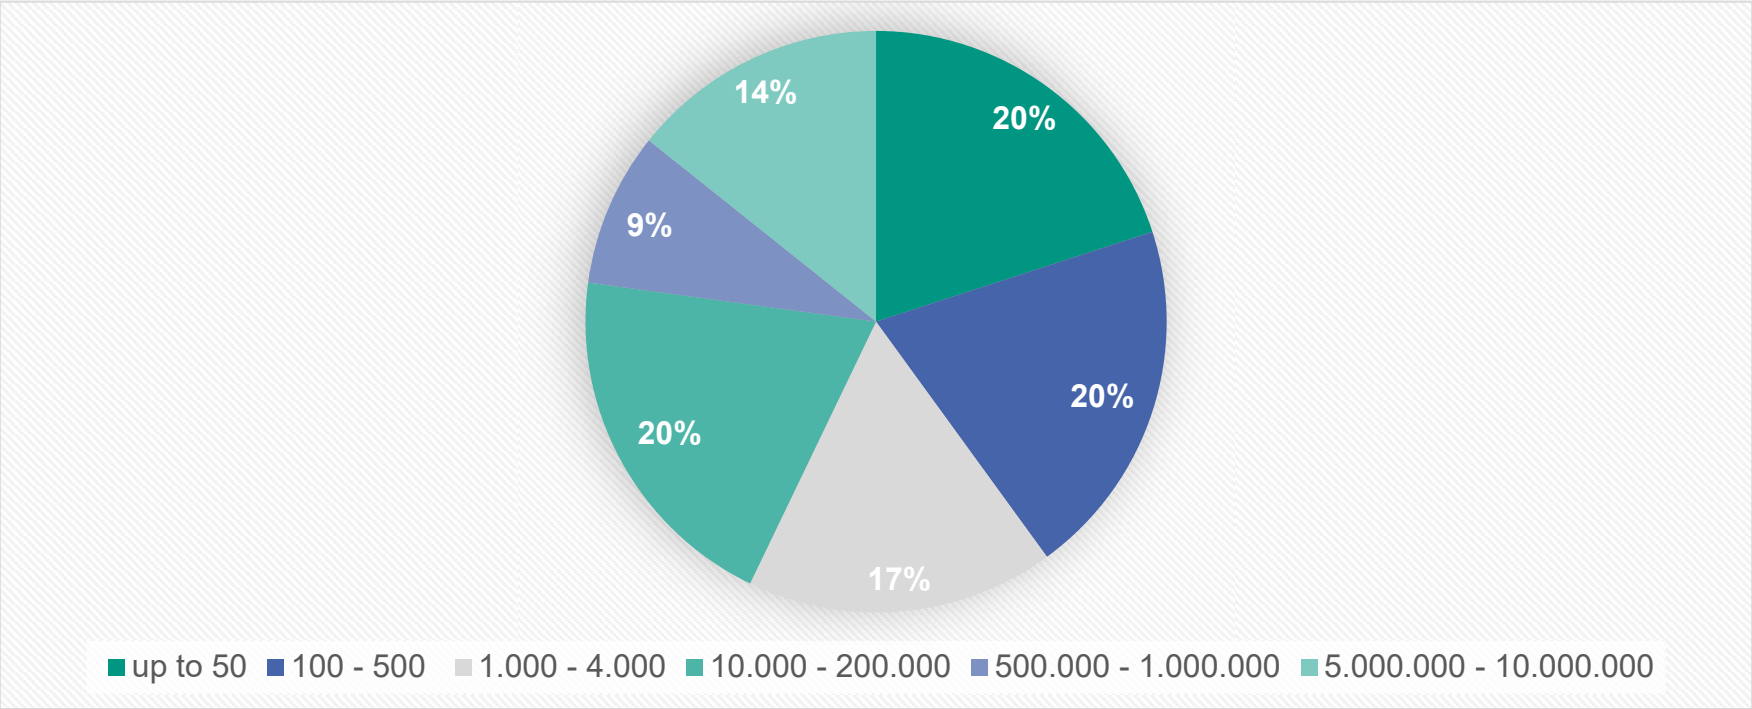

General tendency - Question 5

Prominent answers

- With dataset the “number of files” are even larger as some datasets contain many small files
- Larger datasets are usually partitioned into zip files

Q6: Storage space (GByte)

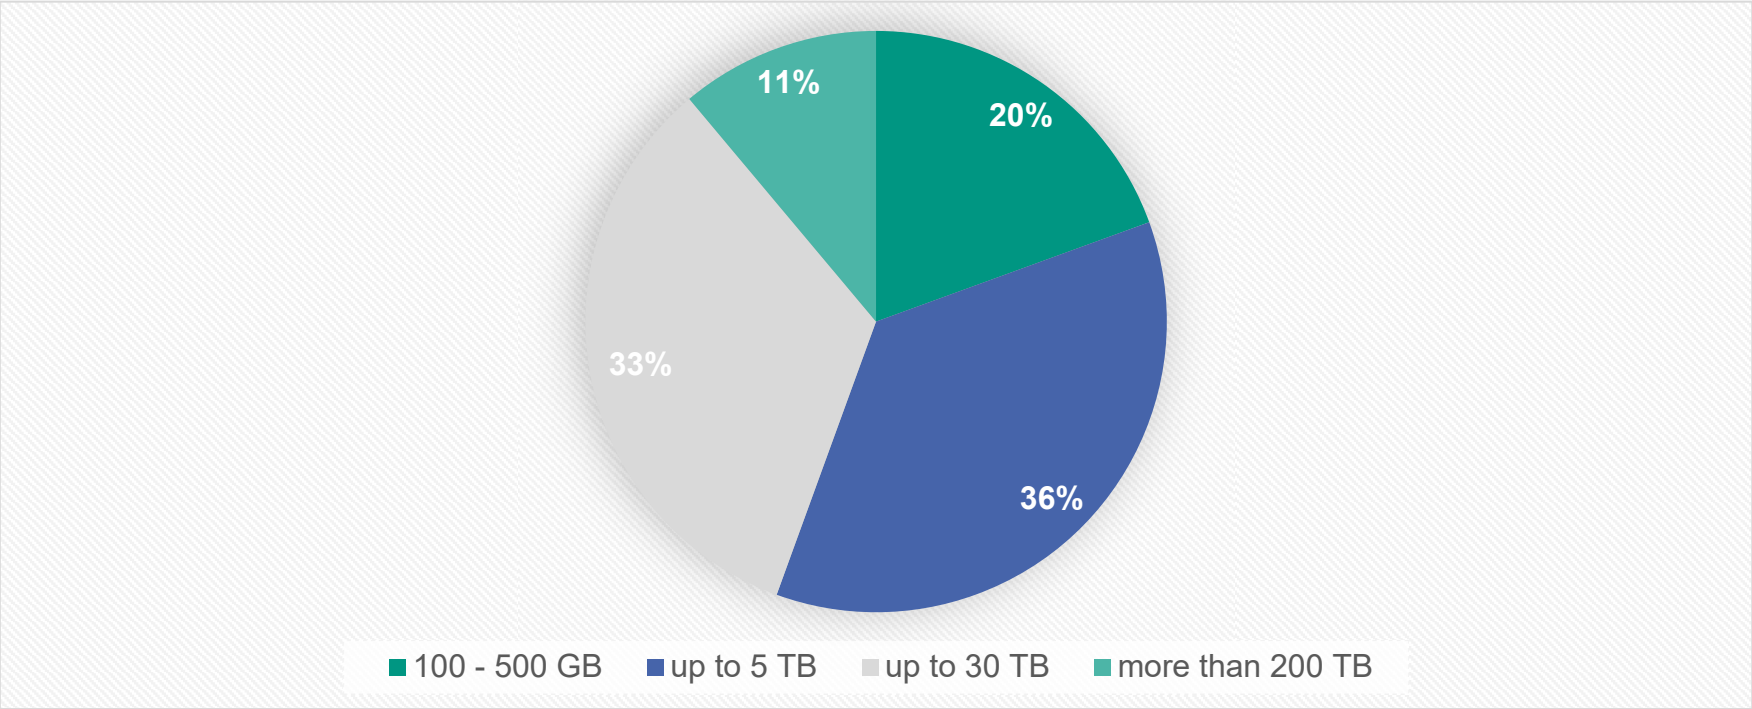

General tendency - Question 6

Prominent answers

- Highly depends, up to several dozen terabytes
- With workspaces up to a few terabyte for storing large datasets

Q7: Data transfer from/to cluster (GByte)

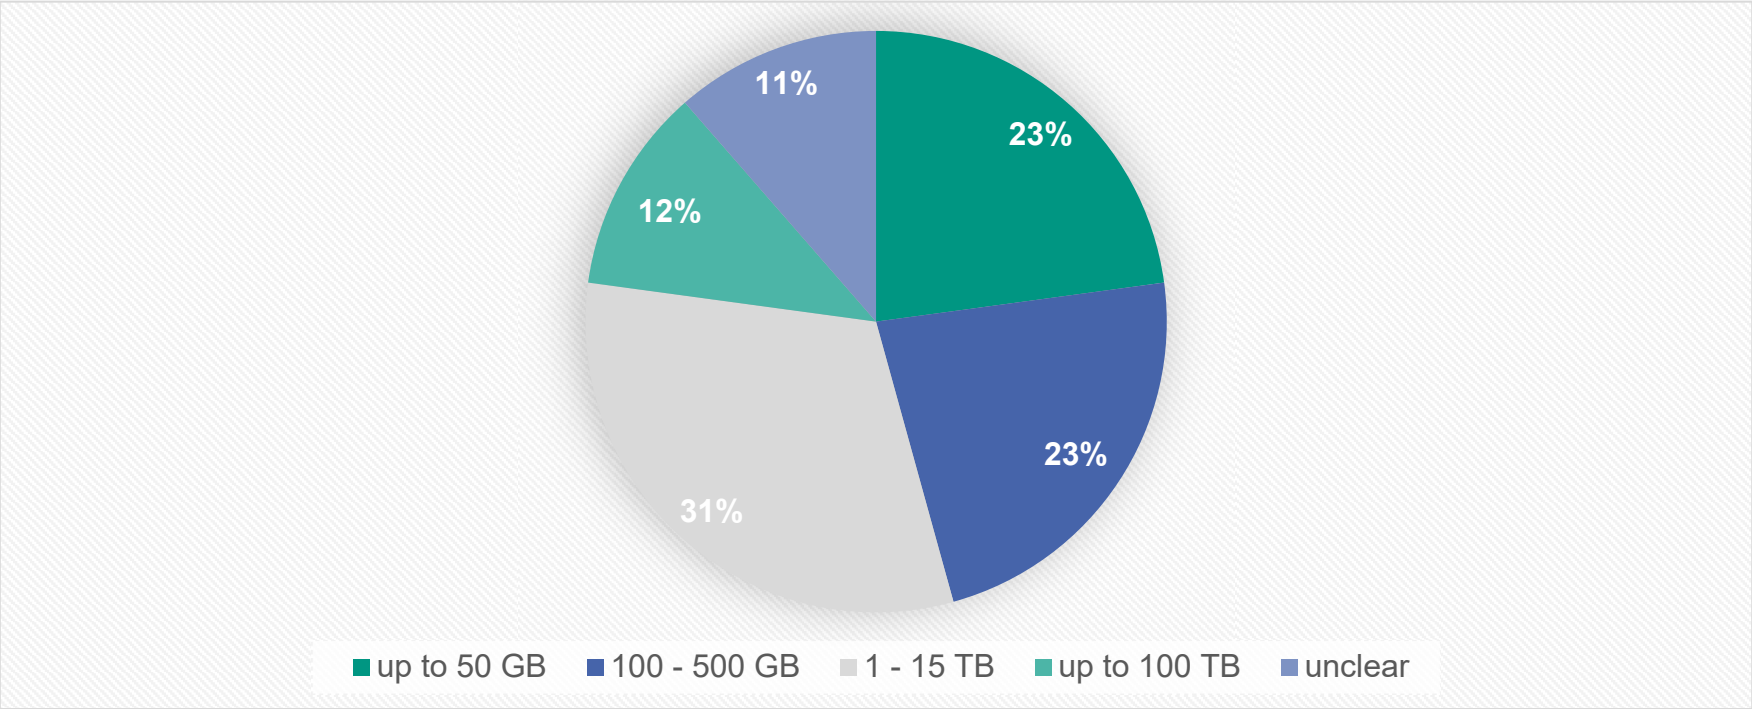

General tendency - Question 7

Prominent answers

- transfer to the cluster <10 GB
  transfer from cluster if trajectories are still needed ~500 GB
  transfer from cluster if only finally evaluated data/ datasets are stored < 10 GB
- Usually, the data transferred are checkpoints, results, and curated datasets ~50GB

Q8: HPC or AI workloads

A. **Mainly HPC**: 24 (64.86%)

B. **Mainly AI**: 9 (24.32%)

C. **Mix of both**: 4 (10.81%)

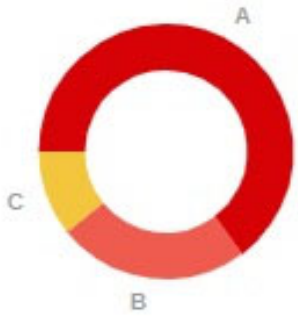

Q9: Job Profile: CPU & GPU Mix

A. **Only CPU**: 15 (40.54%)

B. **Only GPU**: 7 (18.92%)

C. **CPU & GPU**: 15 (40.54%)

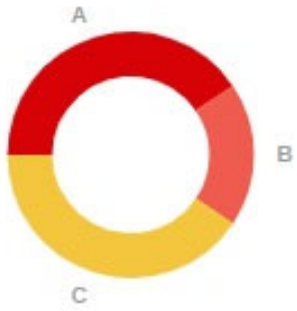

Q10: Porting of Workflows and applications to GPUs (within the next 3 years)

A. Porting to GPU intended: 8 (21.62%)

B. Help needed: 0 (0.00%)

C. No porting intended: 10 (27.03%)

D. Already running on GPU: 19 (51.35%)

Q11: What type of accelerator is needed?

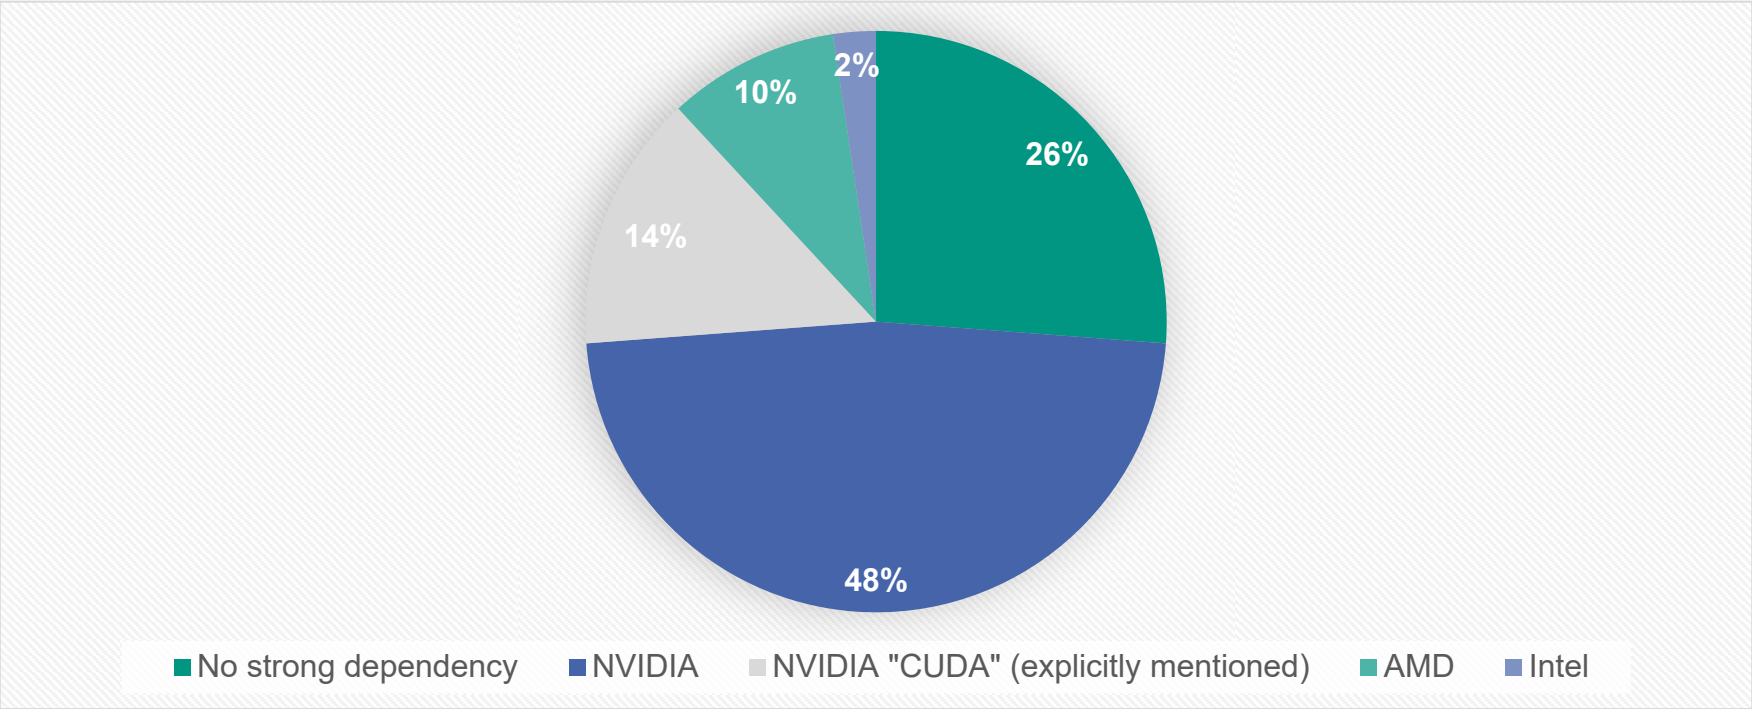

General tendency - Question 11

Prominent answers

- not always needed, some algorithms do not fit into accelerators' specific architecture
- Most of the codes are strongly dependent on CUDA

Q12a: Can your application run on ARM?

A. **Yes**: 15 (40.54%)

B. **No**: 4 (10.81%)

C. **Unclear**: 18 (48.65%)

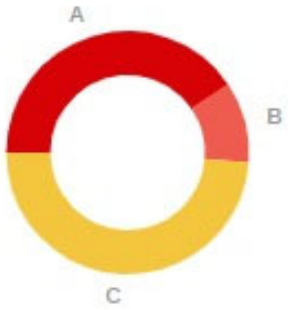

Q12b: If not, what is the reason?

- Some of the codes run on ARM but very slow. This cannot be improved because of the ARM architecture.
- SVE (The Arm Scalable Vector Extension) would be interesting to us, we'd have to port our SIMD kernels, though.
- ARM version are not fully tested and do not have all the physics features. It will take more time to transition algorithms for ARM architectures.
- 2x CUDA dependencies
- Framework support is missing
- using of proprietary software like StarCCM+ and Hyperworks

Q13: Used Software

- 6x OpenMPI (provided by SCC)
- 3x NetCDF & 3x HDF5 (provided by SCC)
- 2x CUDA (provided by SCC)
- 5x PyTorch (provided by SCC)
- 2x VASP & 1x Quantum Espresso & 1x Yambo (provided by SCC)
- 4x OpenFOAM & 3x LAMMPS & 2x Gromacs (provided by SCC)
- 2x VRE
- 2x Compiler (provided by SCC)
- 1x Starccm+
- 1x Hyperworks
- 1x OpenMM
- 1x nvhpc (provided by SCC)
- 1x eccodes (provided by SCC)
- 1x CDO (provided by SCC)
- 1x ADIOS2
- 1x Lapack & 1x FFTW (provided by SCC)
- 1x C++ (provided by SCC)
- 1x Conda (provided by SCC)
- 1x Salvus
- 1x SPECFEM
- 1x NWChem
- 1x Turbomole
- 1x ASE
- 1x Fireworks
- 1x CMS (CernVM-File System)
- 1x Comsol
- 1x Glib/Gfortan (provided by SCC)
- 1x paraview (provided by SCC)

Q13a: Used Software - What percentage of the applications you use were created by you?

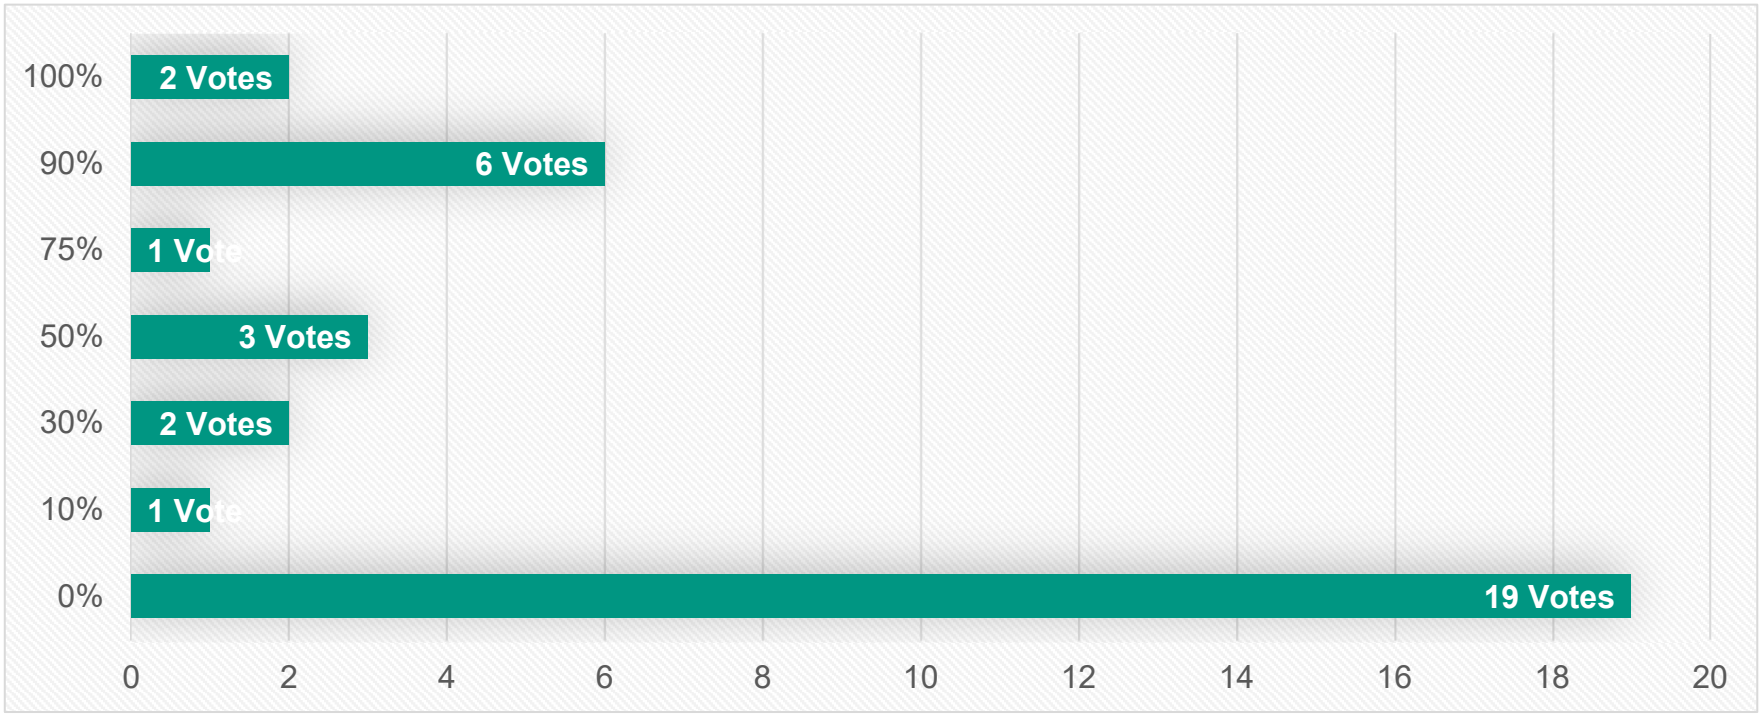

Q13b: Used Software - Do you use the software we offer or do you compile your own software?

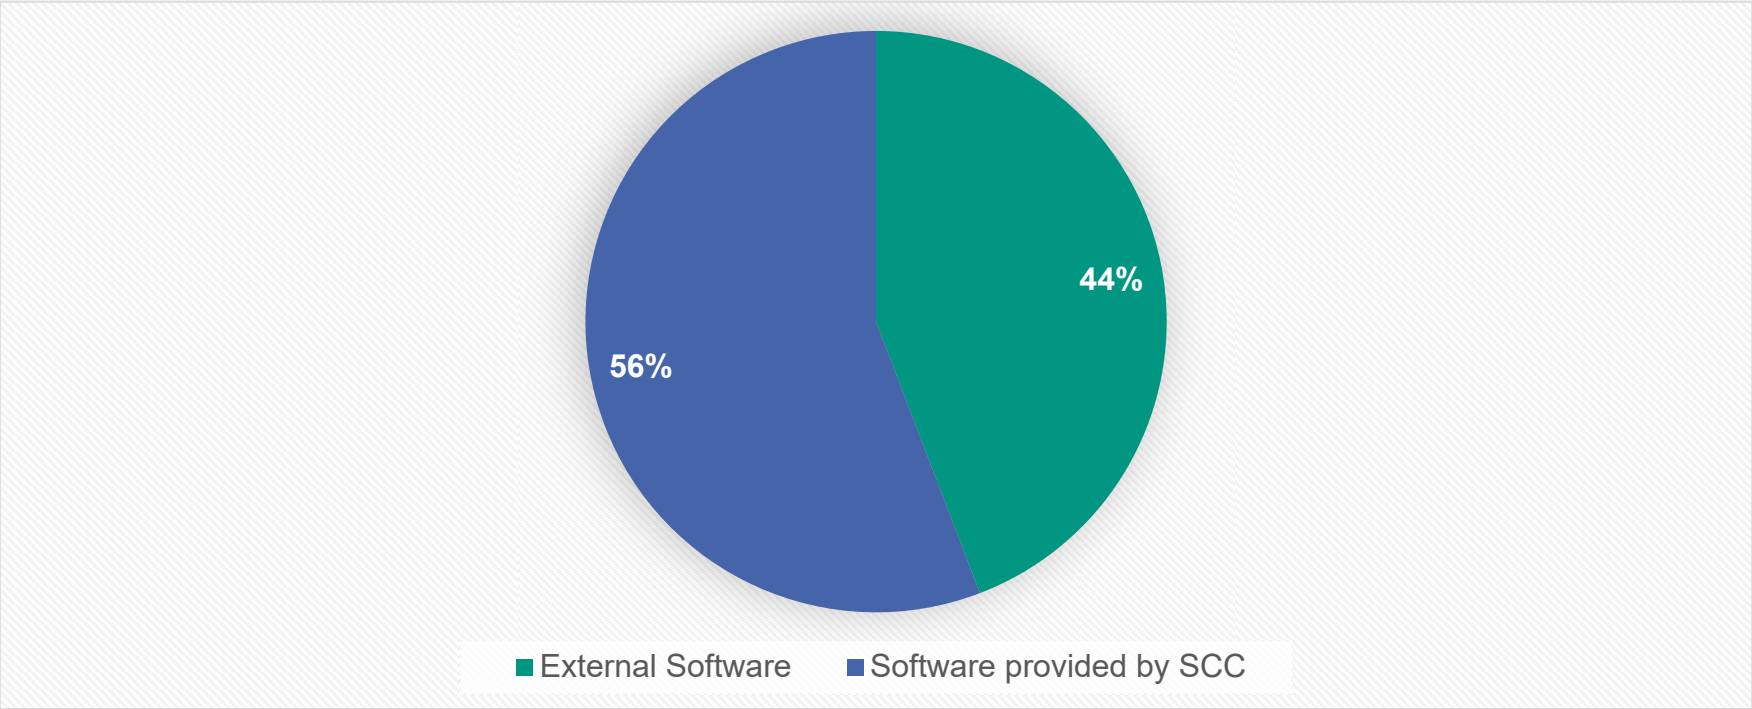

Q14: Application support services

A. **No**: 7 (18.92%)

B. **Yes, in the context of questions and answers**: 21 (56.76%)

C. **Yes, in the context of up to one-day workshops**: 3 (8.11%)

D. **Yes, as part of smaller projects**: 1 (2.70%)

E. **Yes, as part of projects lasting up to several months**: 5 (13.51%)