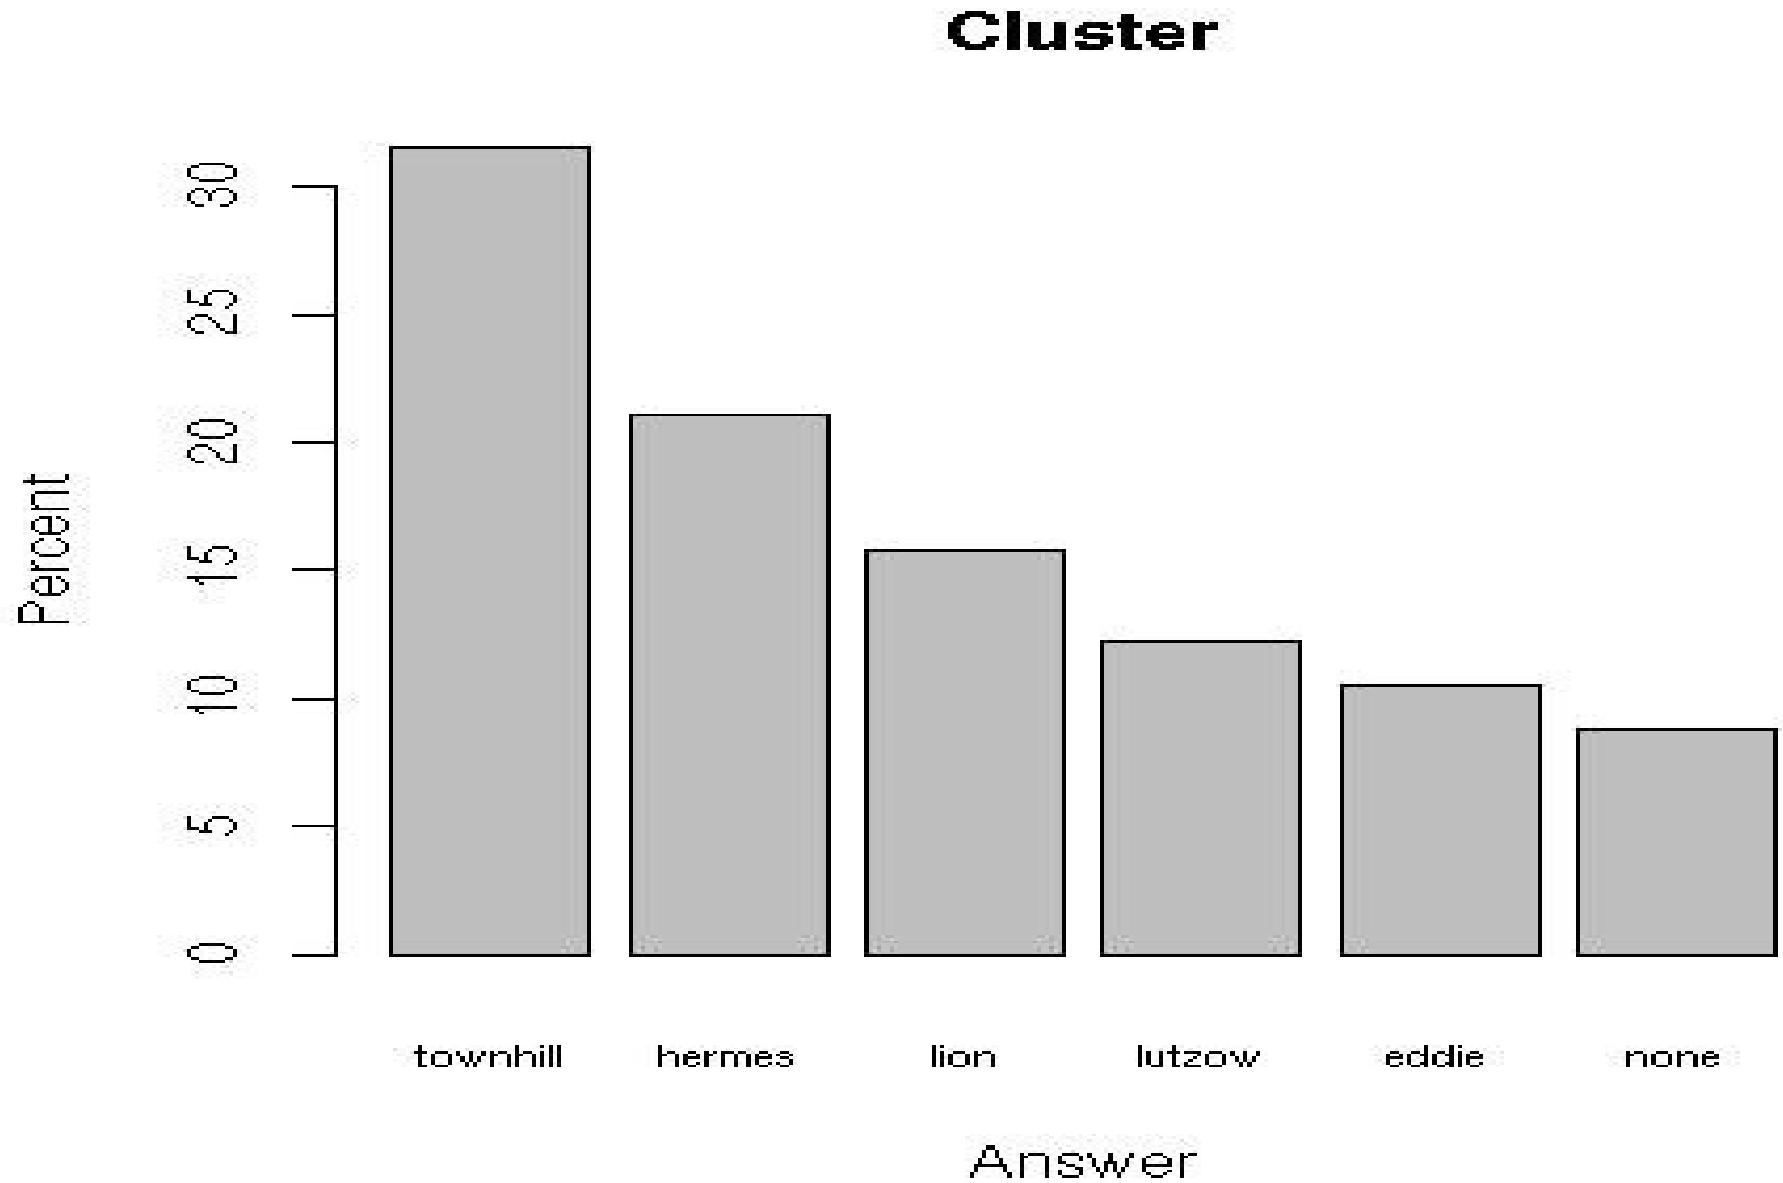
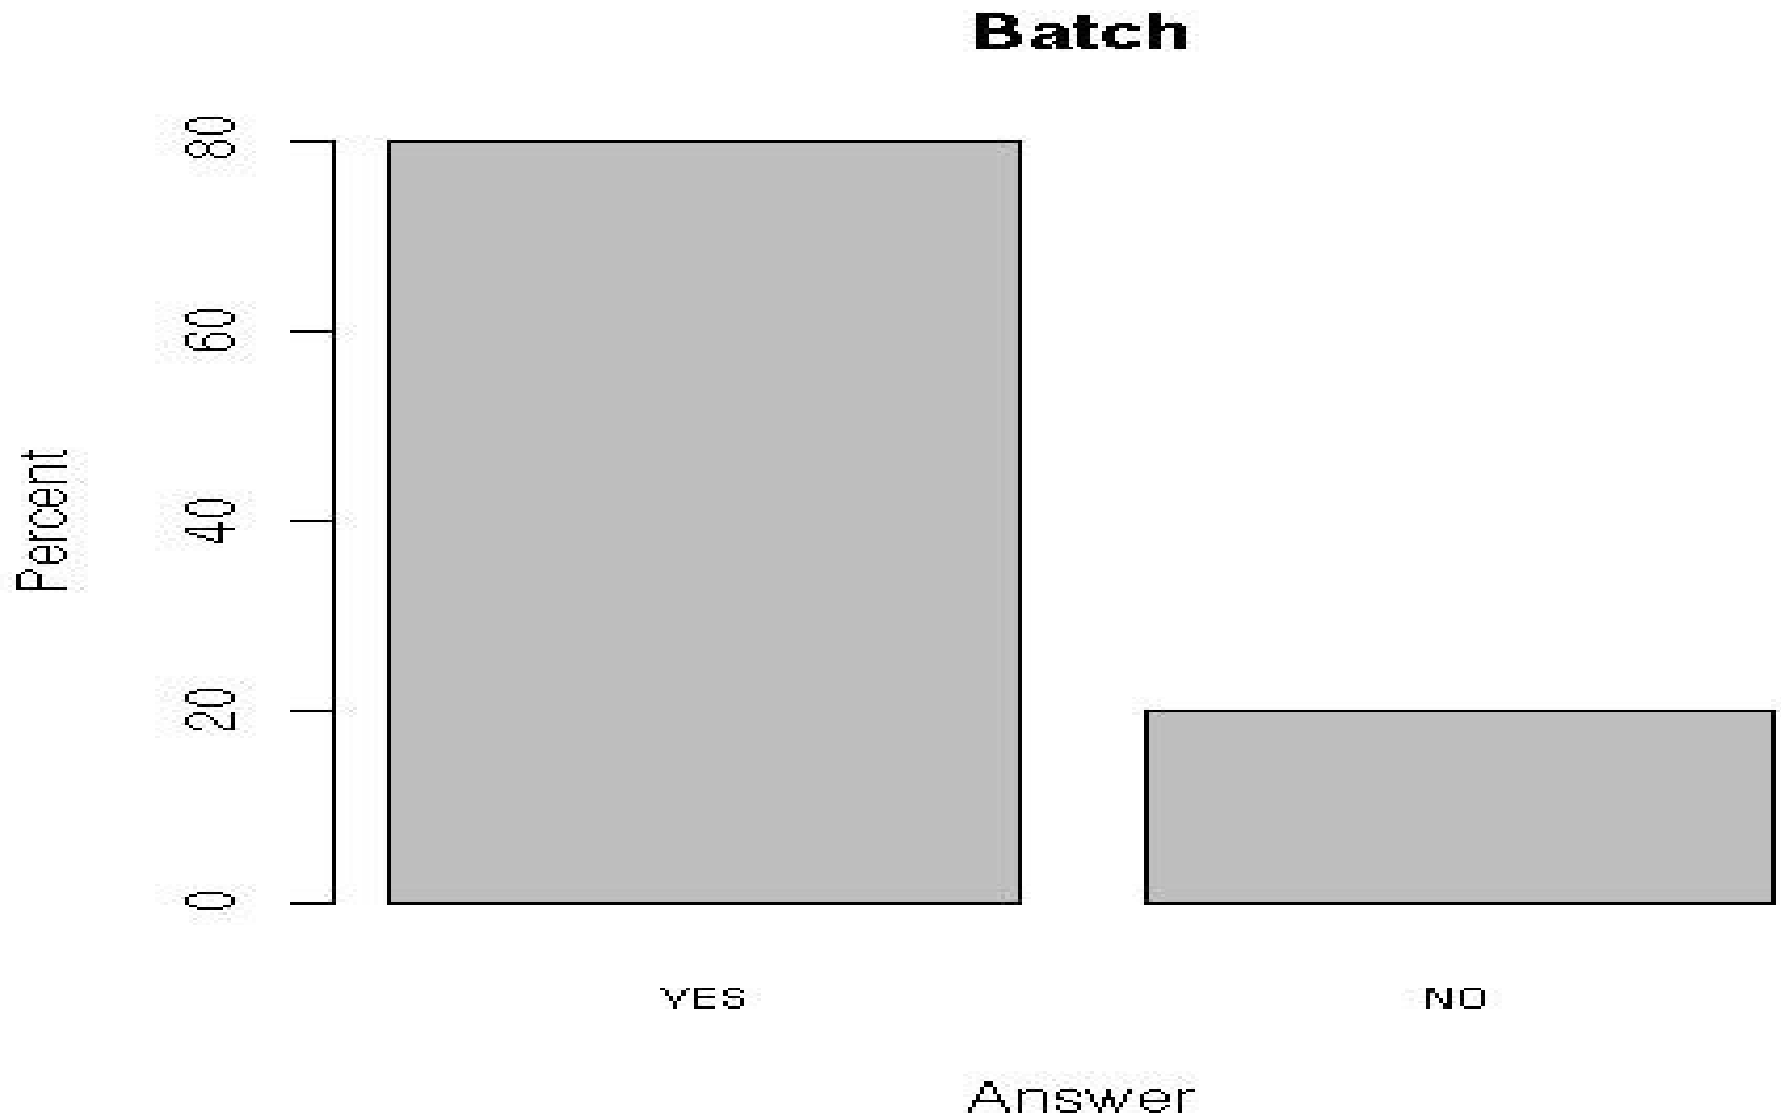


Which clusters do you use?




Do you use simply for batching jobs?




Do you use for true multi-processor jobs?

Multi-processor

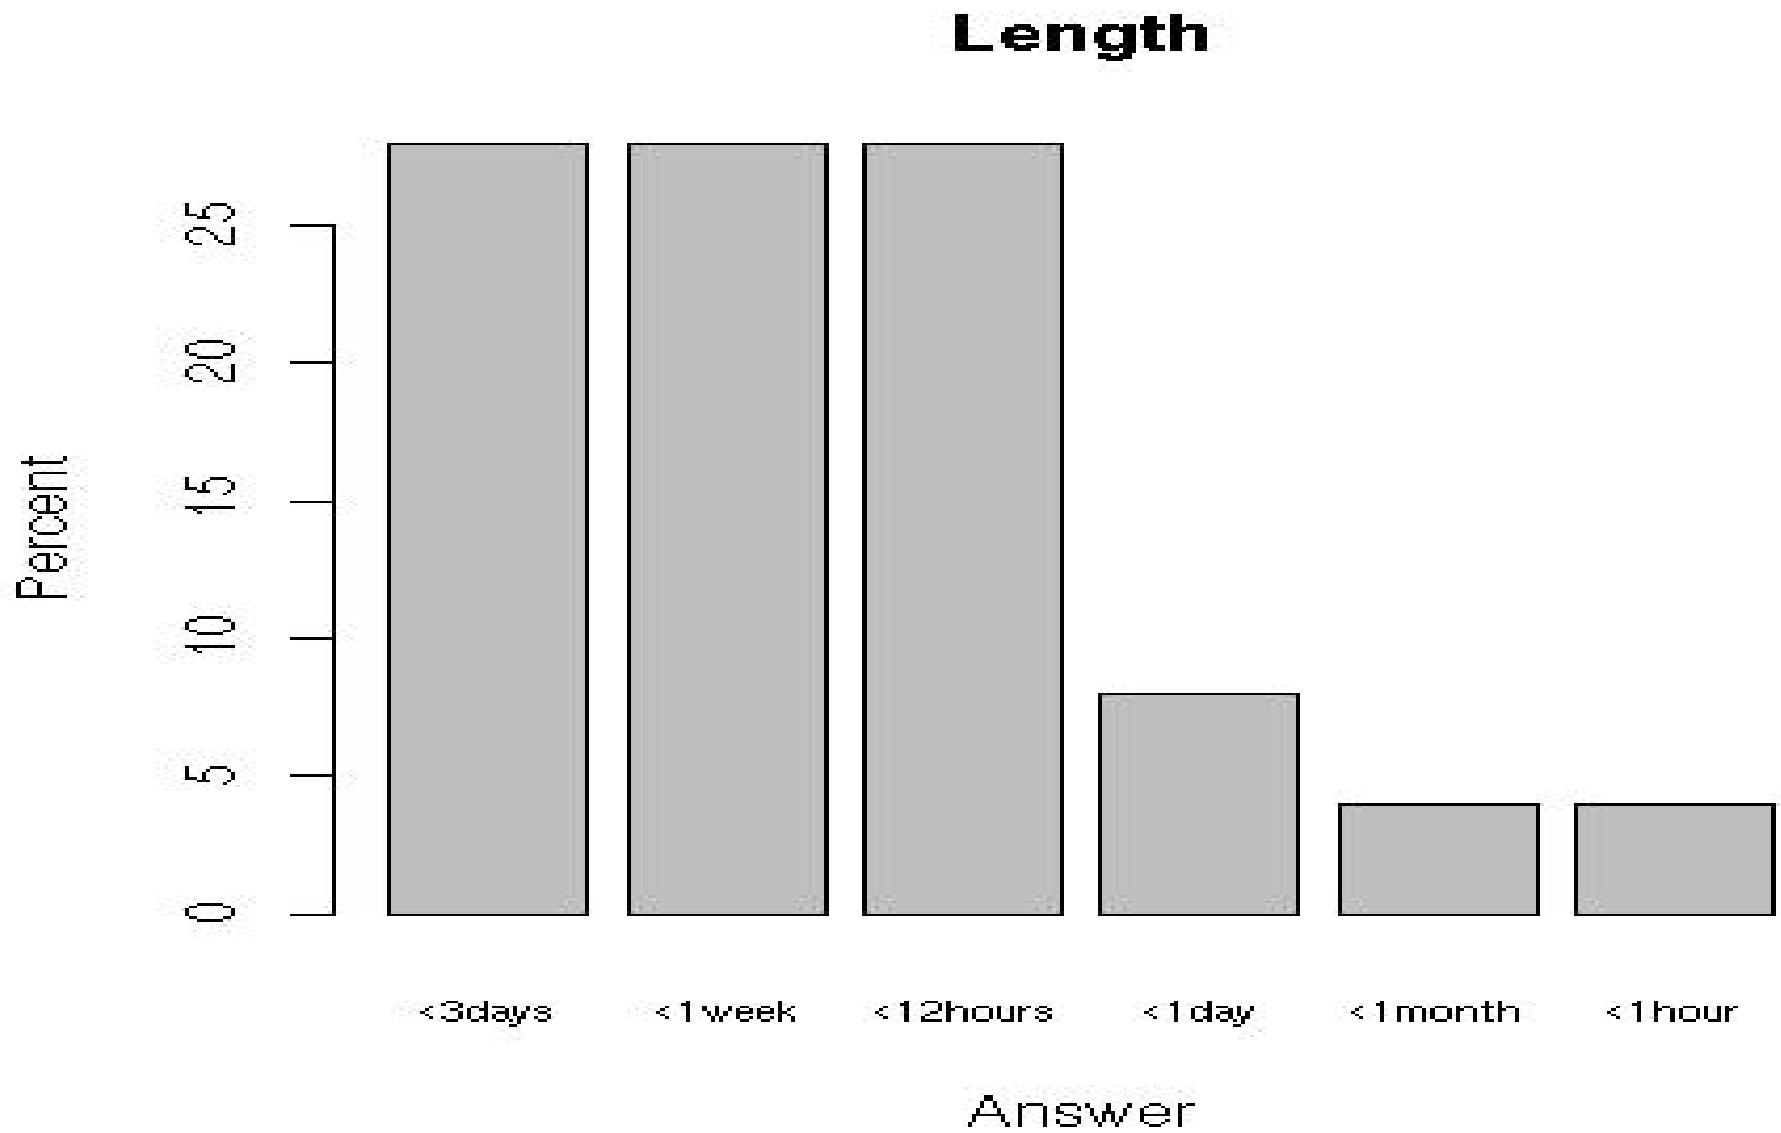


Do you need tight interconnect?

Tight



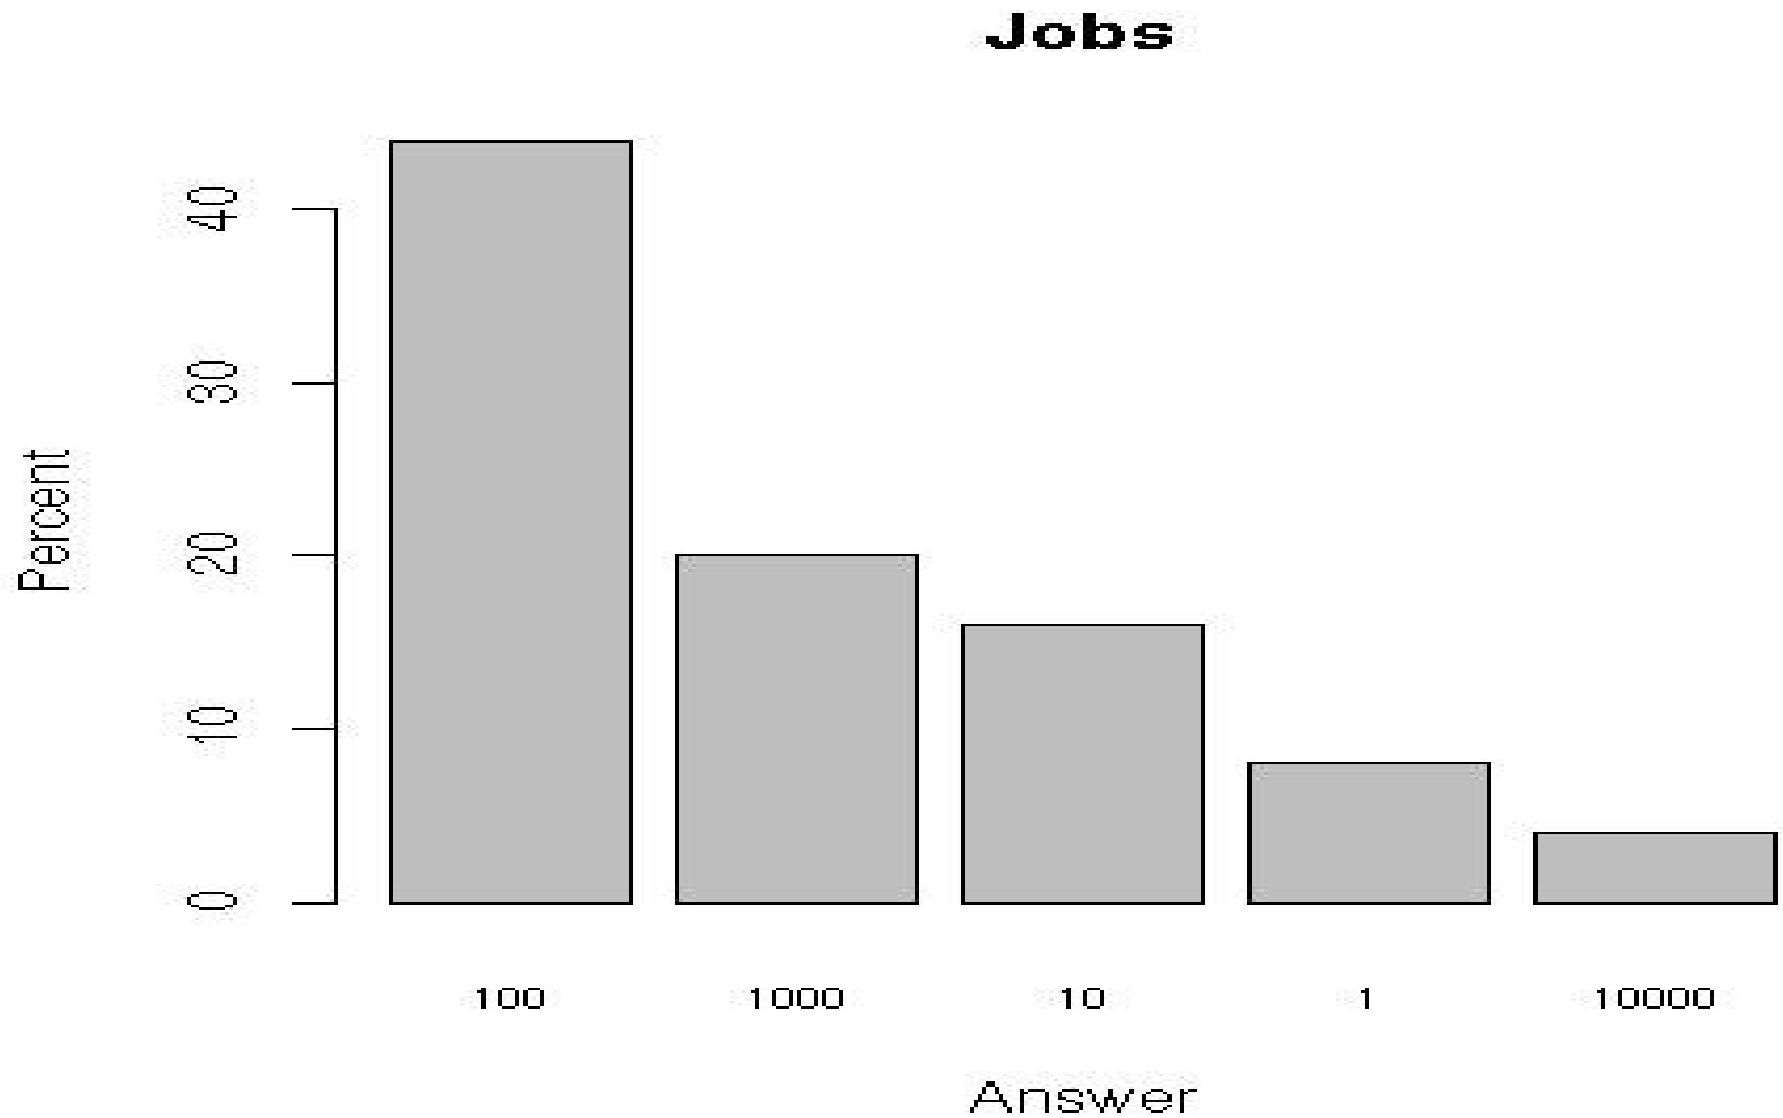
Typical total length of a run?



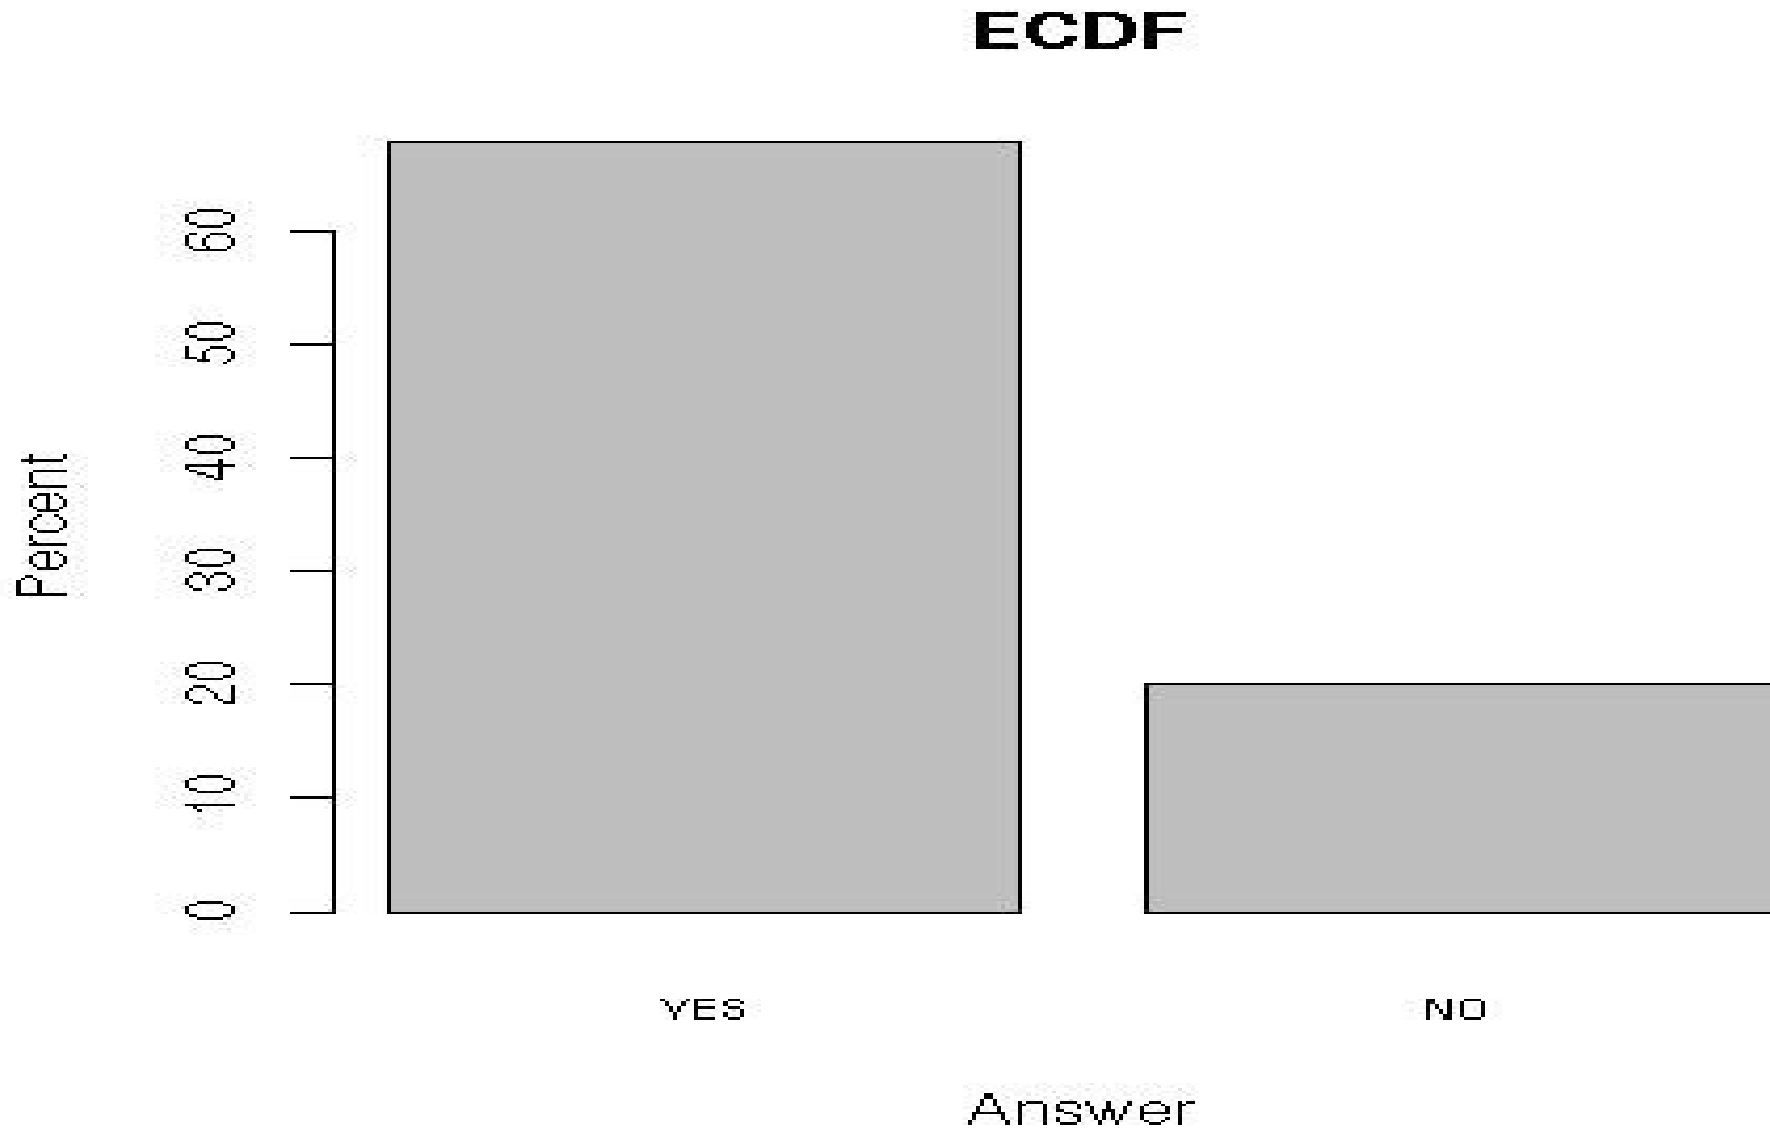
Typical length of each component in a run?



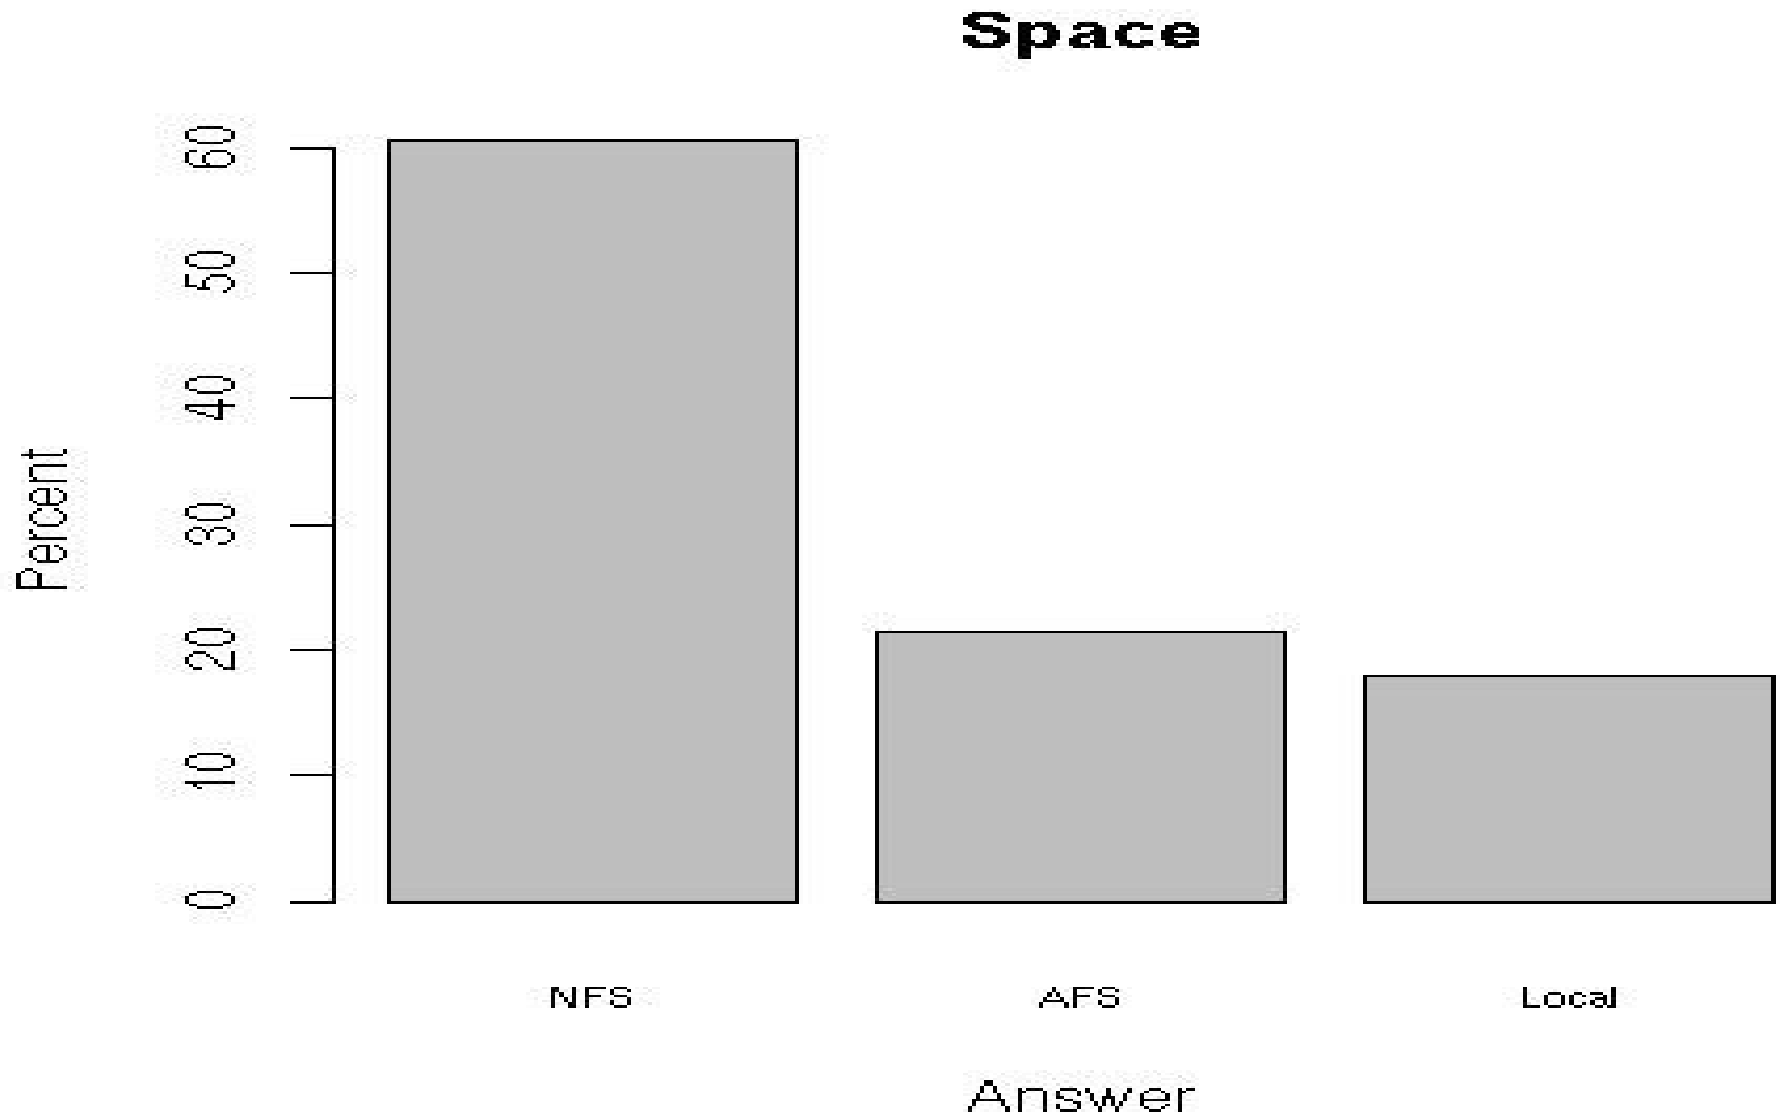
Typical number of component jobs making up each run?



Would you like file space shared with the ECDF cluster?




Type of file space currently used?



For your shareable data which is more important:


- shareable between Informatics and ECDF clusters (Answer=ecdf)
- shareable between Informatics clusters and your desktop (Answer=local)



Describe how your data is accessed by your job:


- accessible by every node (Answer=all)
- partitioned by node (Answer=node)

Moved

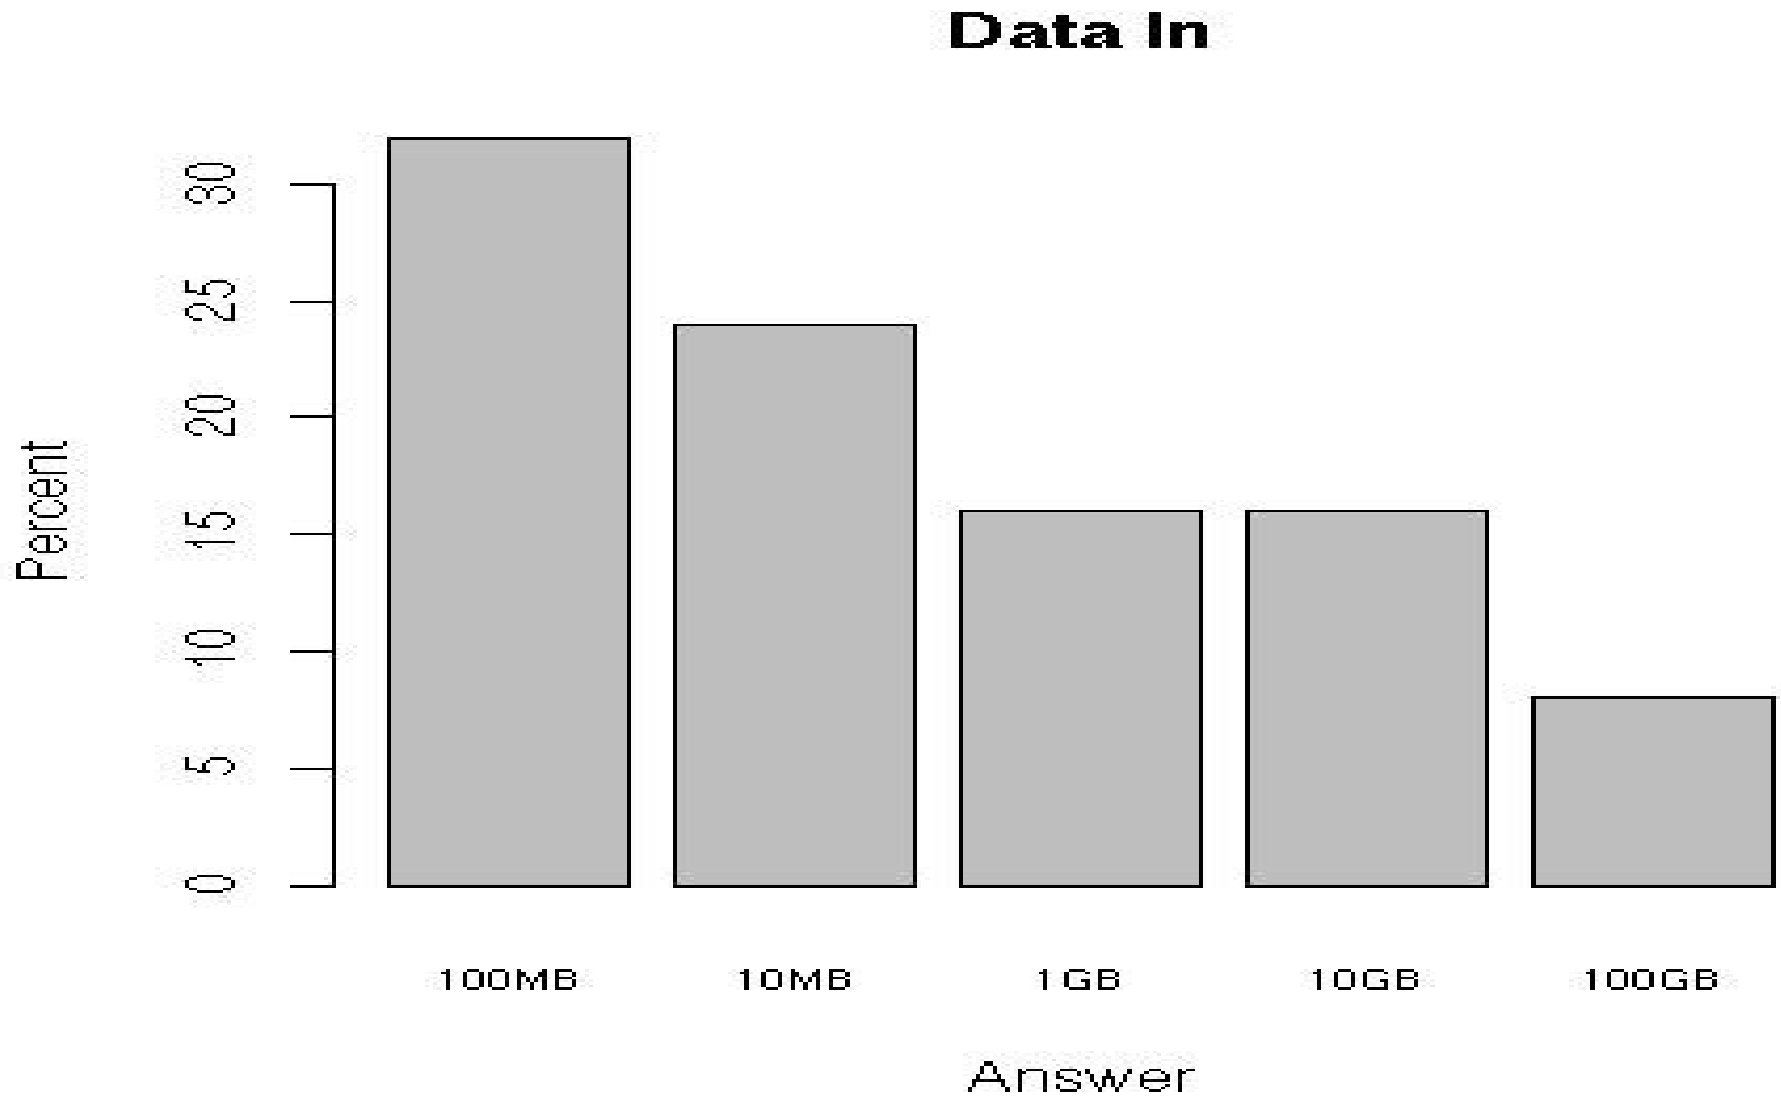


Describe the nature of disk access:

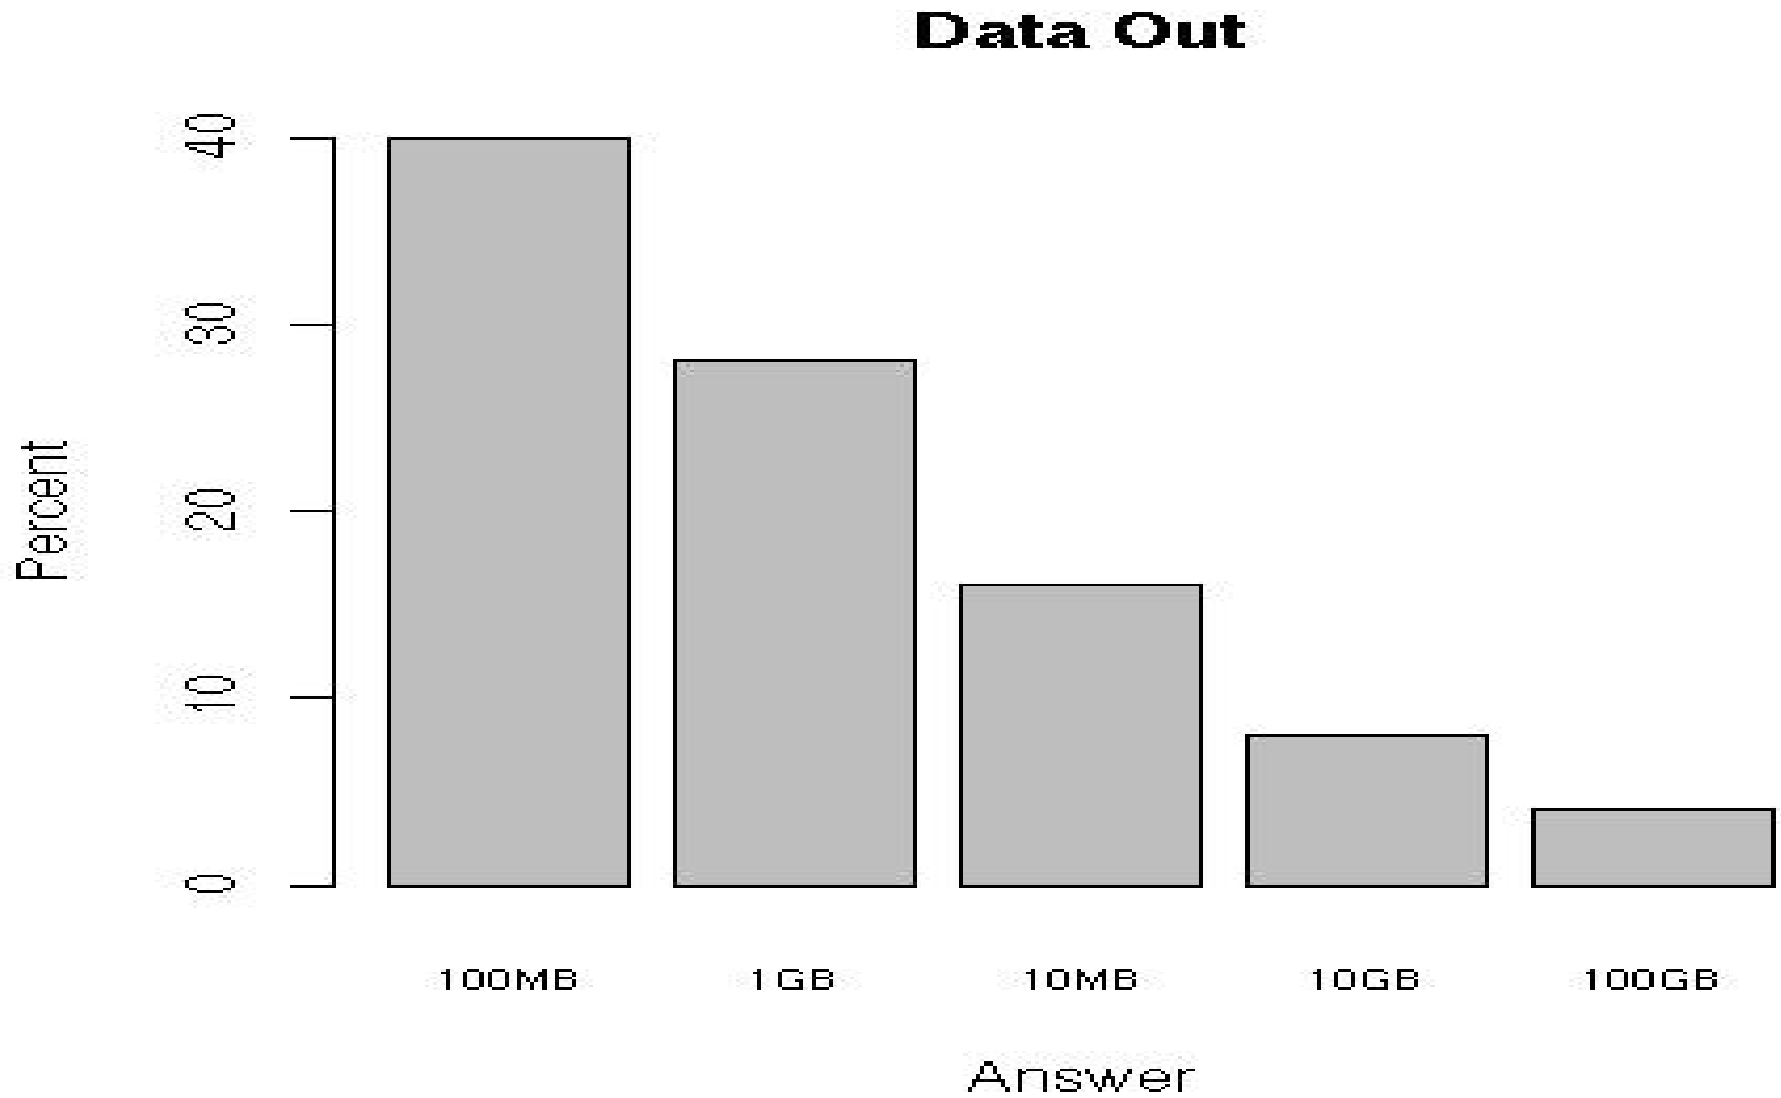
- proportion of data set loaded to node at start of job run and then processed locally (Answer=local)
- whole data set can be accessed by any node during any part of a job run (Answer=full)



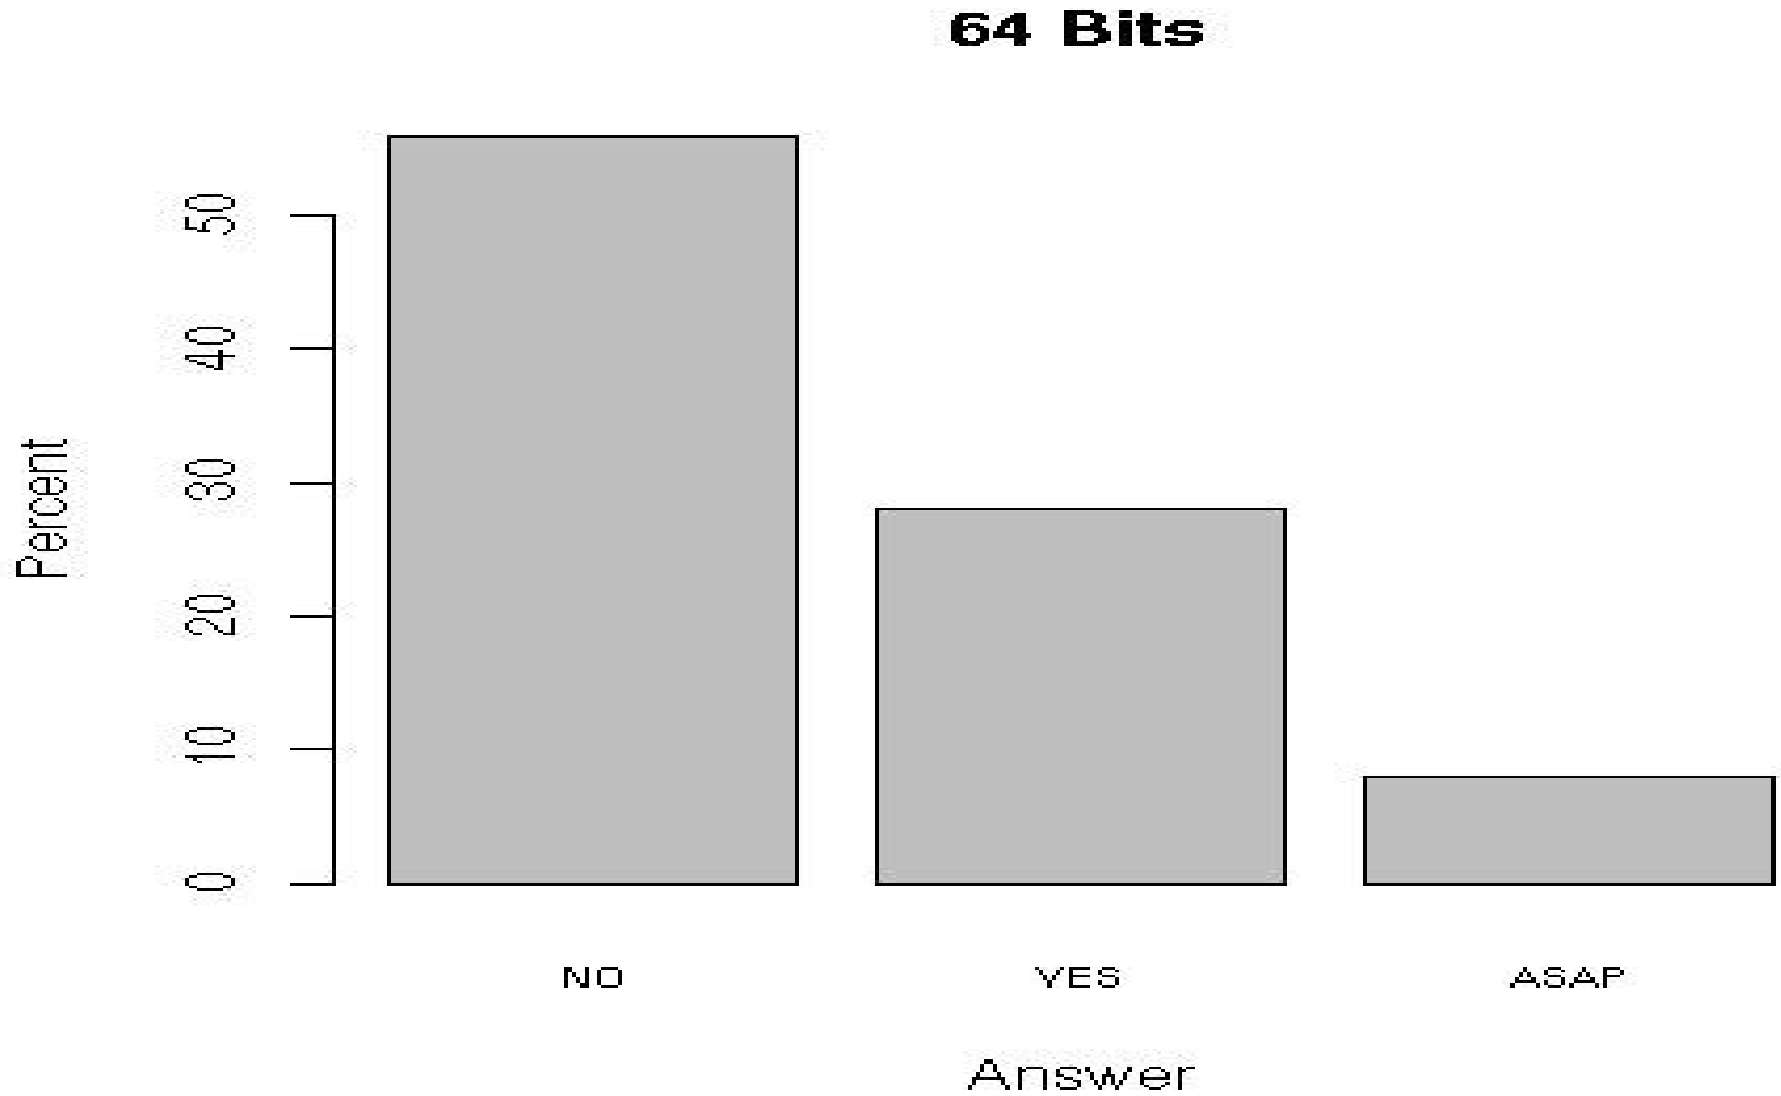
Typical amount of data processed per-run?



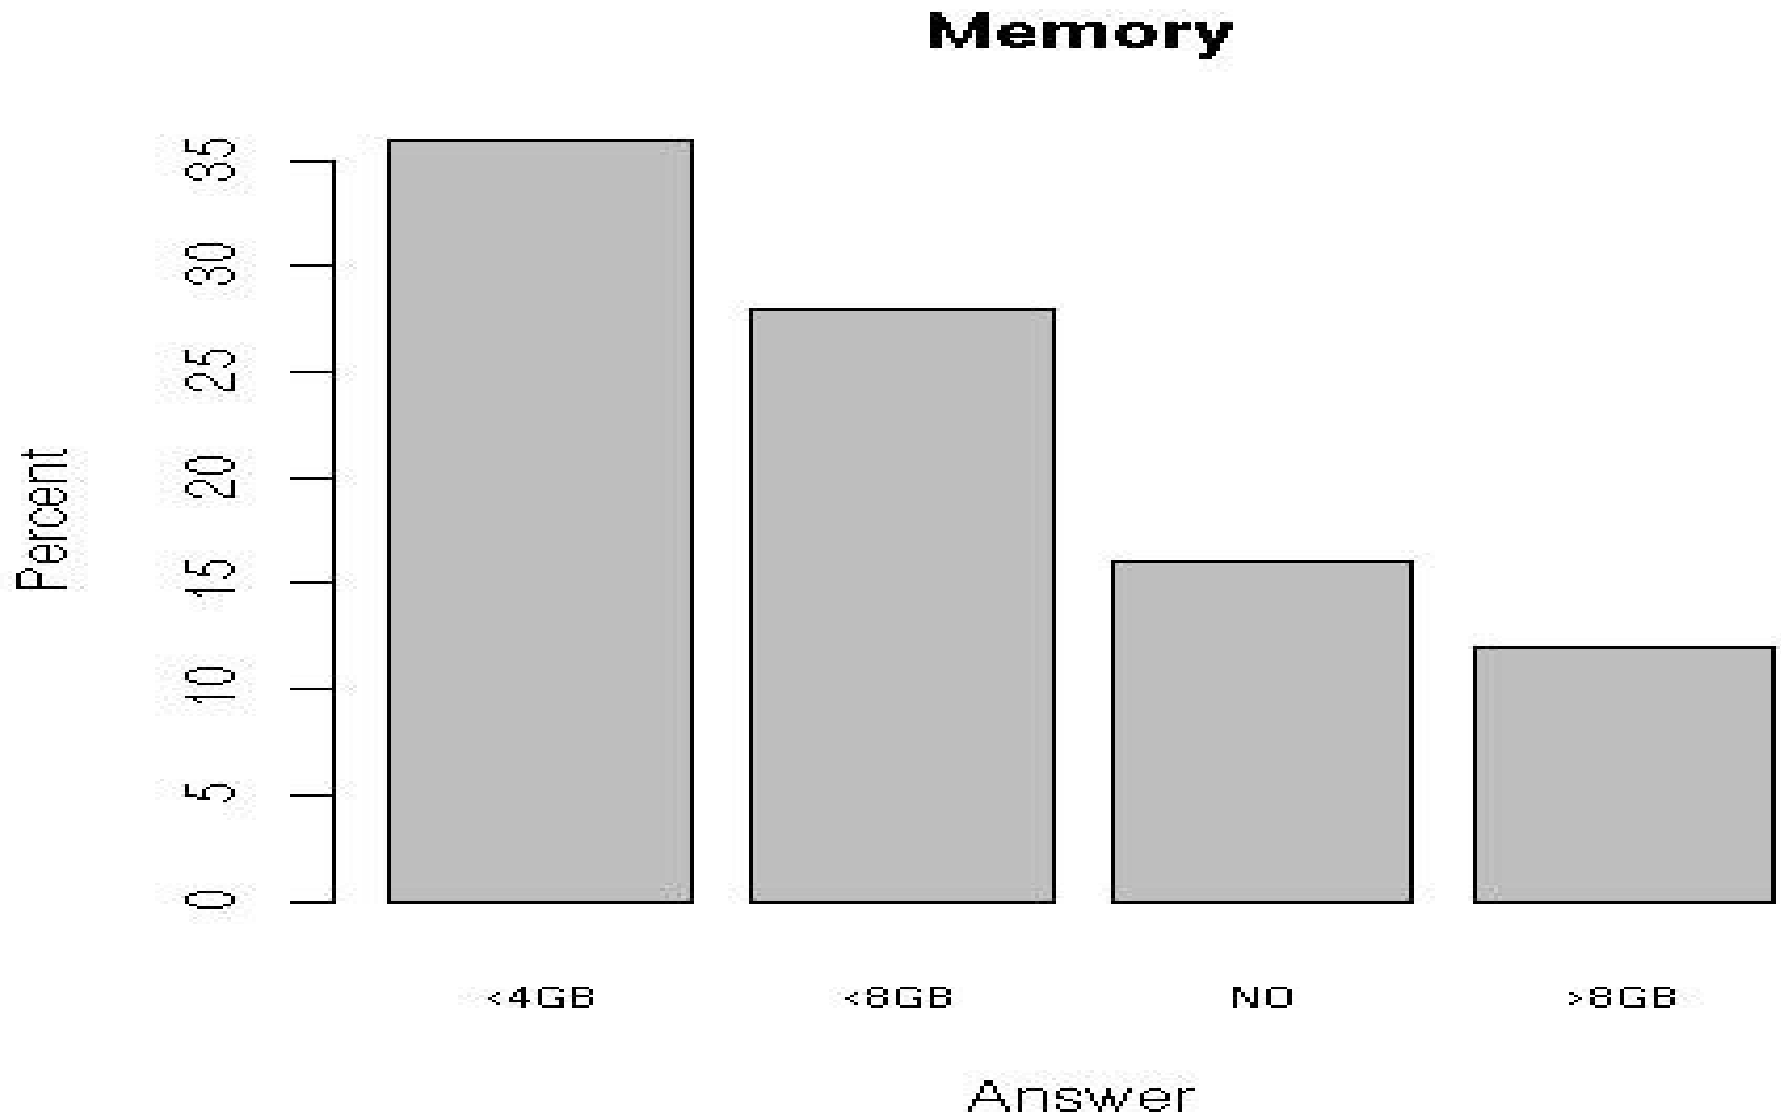
Typical amount of data generated per-run?



Do you need 64bit?




Do you need large memory nodes (Answer is per-processor core)?



Would you use monitoring tools (eg. Ganglia)?

Monitoring



Have you used Condor?

Condor

