



HARVARD UNIVERSITY
Information Technology

RC in the Cloud – Survey results

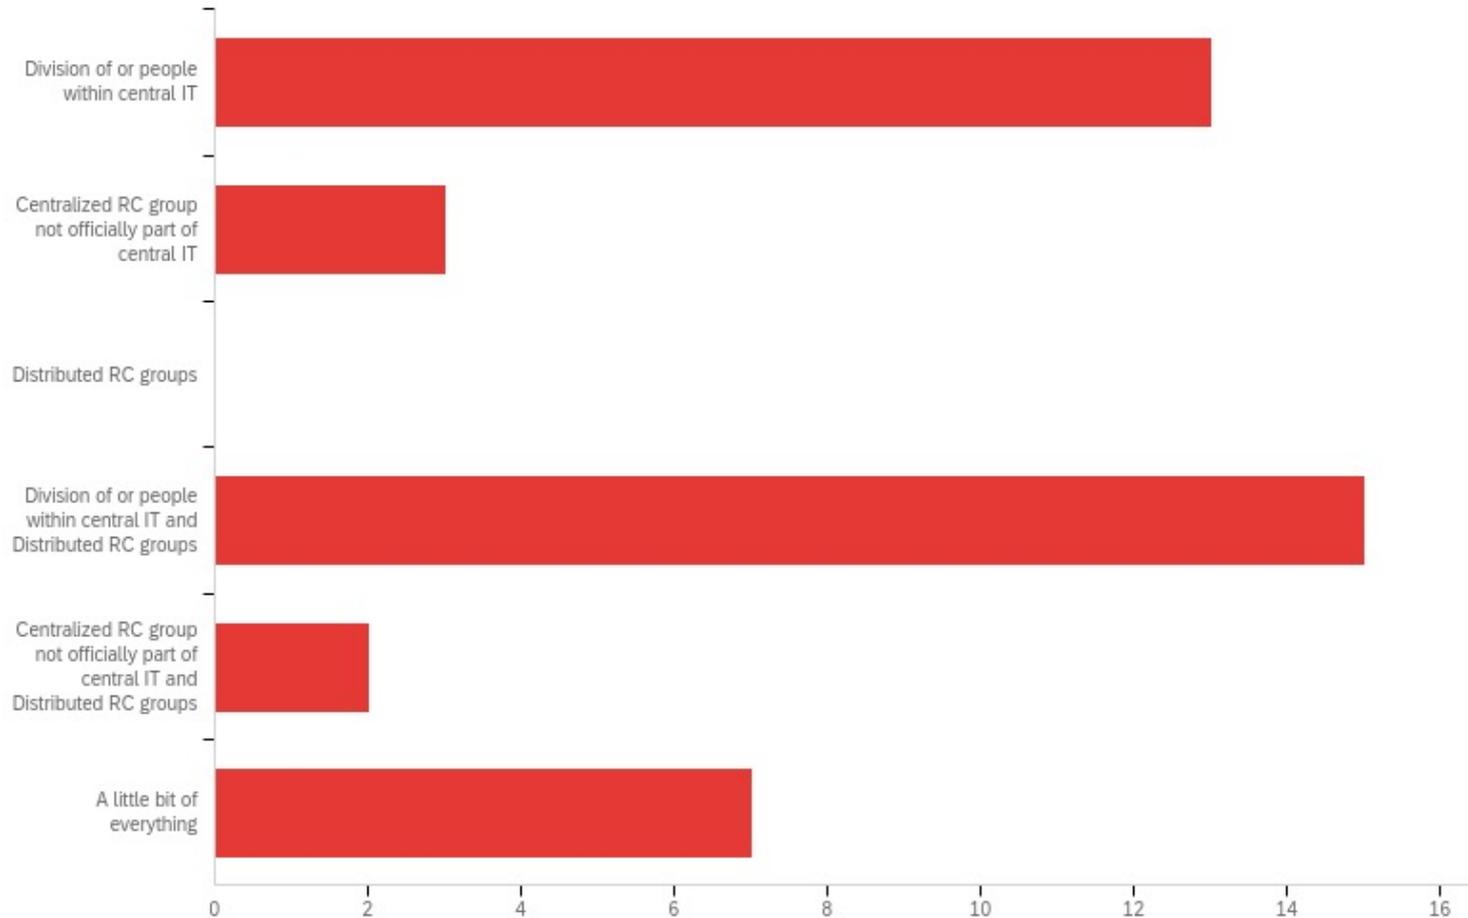
Statistics

- 40 responses, representing 38 institutions
 - Two institutions answered more than once, which has not been corrected for in the data
- Mostly US institutions, plus one in Ireland and two in Canada

How is RC supported generally?

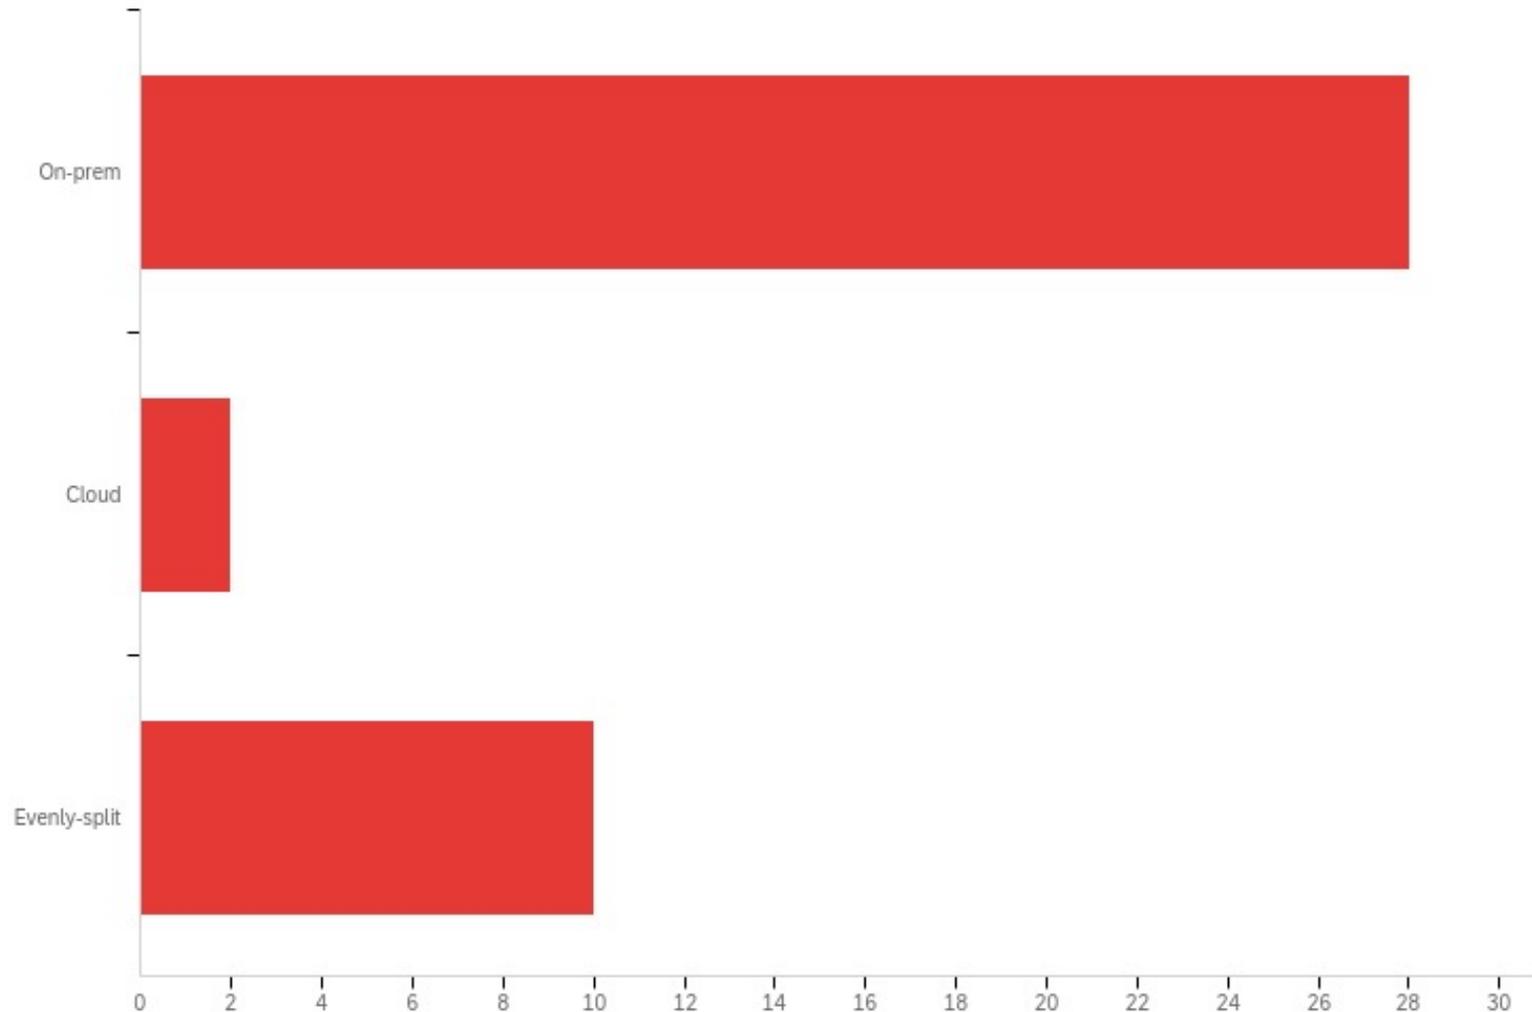
How is Research Computing support provided at your institution?

70% of institutions are supporting RC out of central IT, with a slight majority of those also leveraging distributed RC groups



What is the primary target for research computing in your institution?

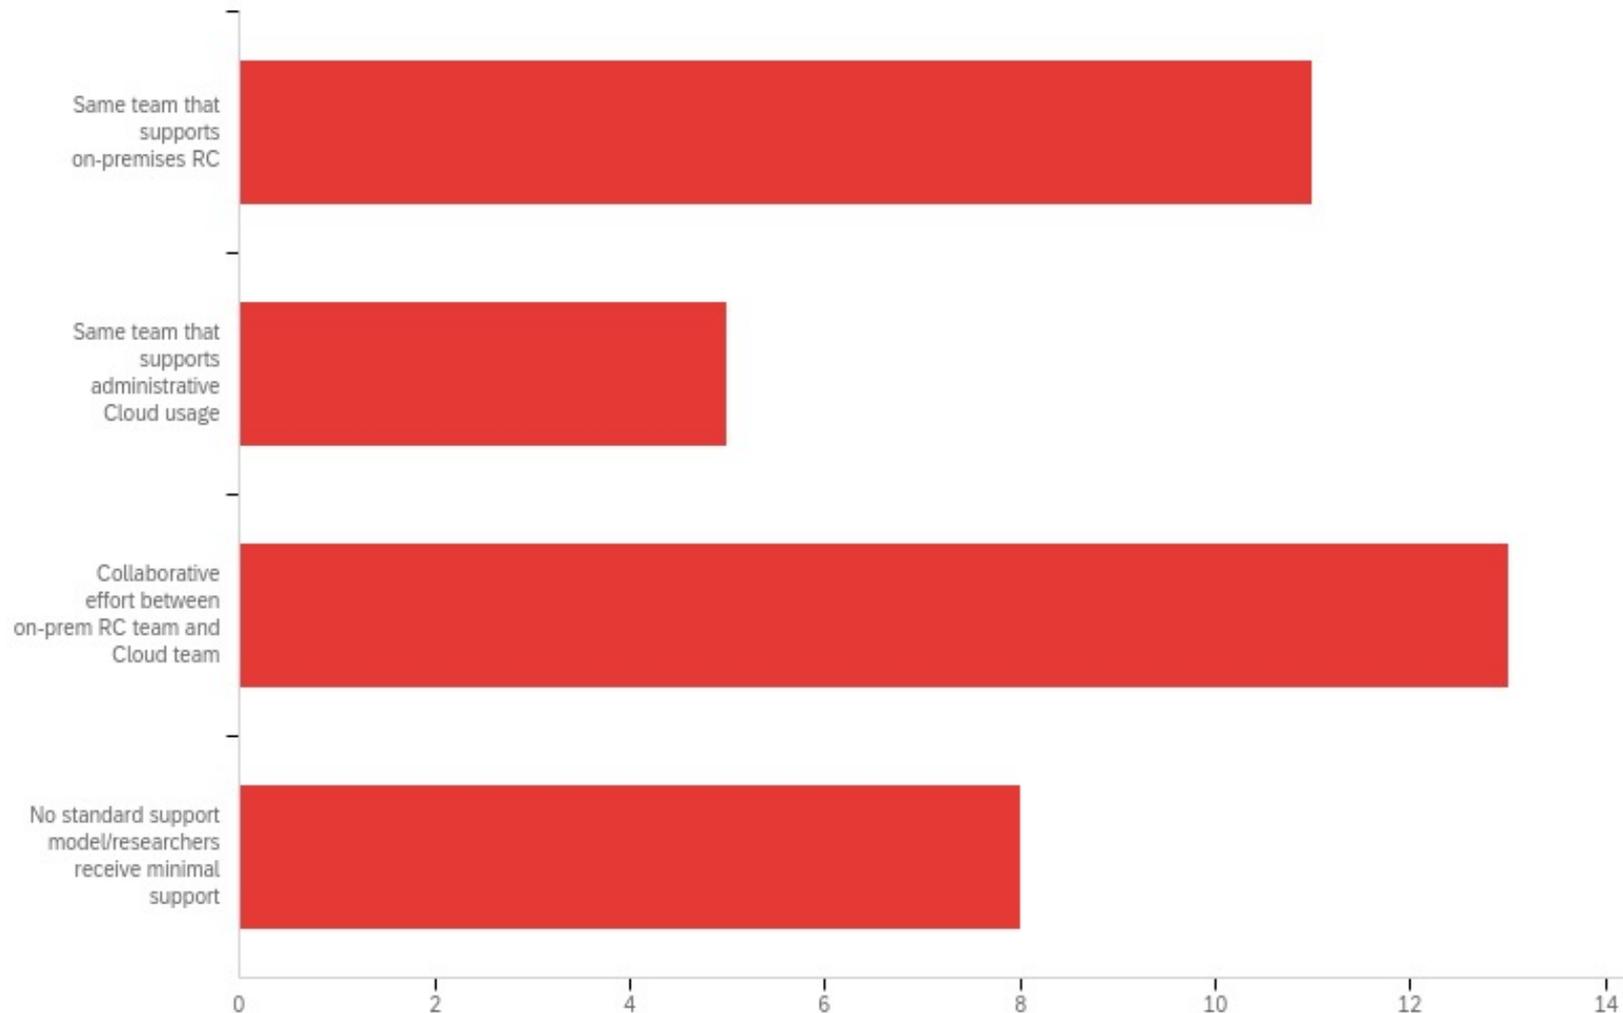
70% of institutions are still primarily on-prem for RC.



How is do you support Cloud RC?

What kind of group provides Cloud RC support?

Support for Cloud RC is most likely to come out of the on-prem RC team or that team collaborating with a Cloud team.



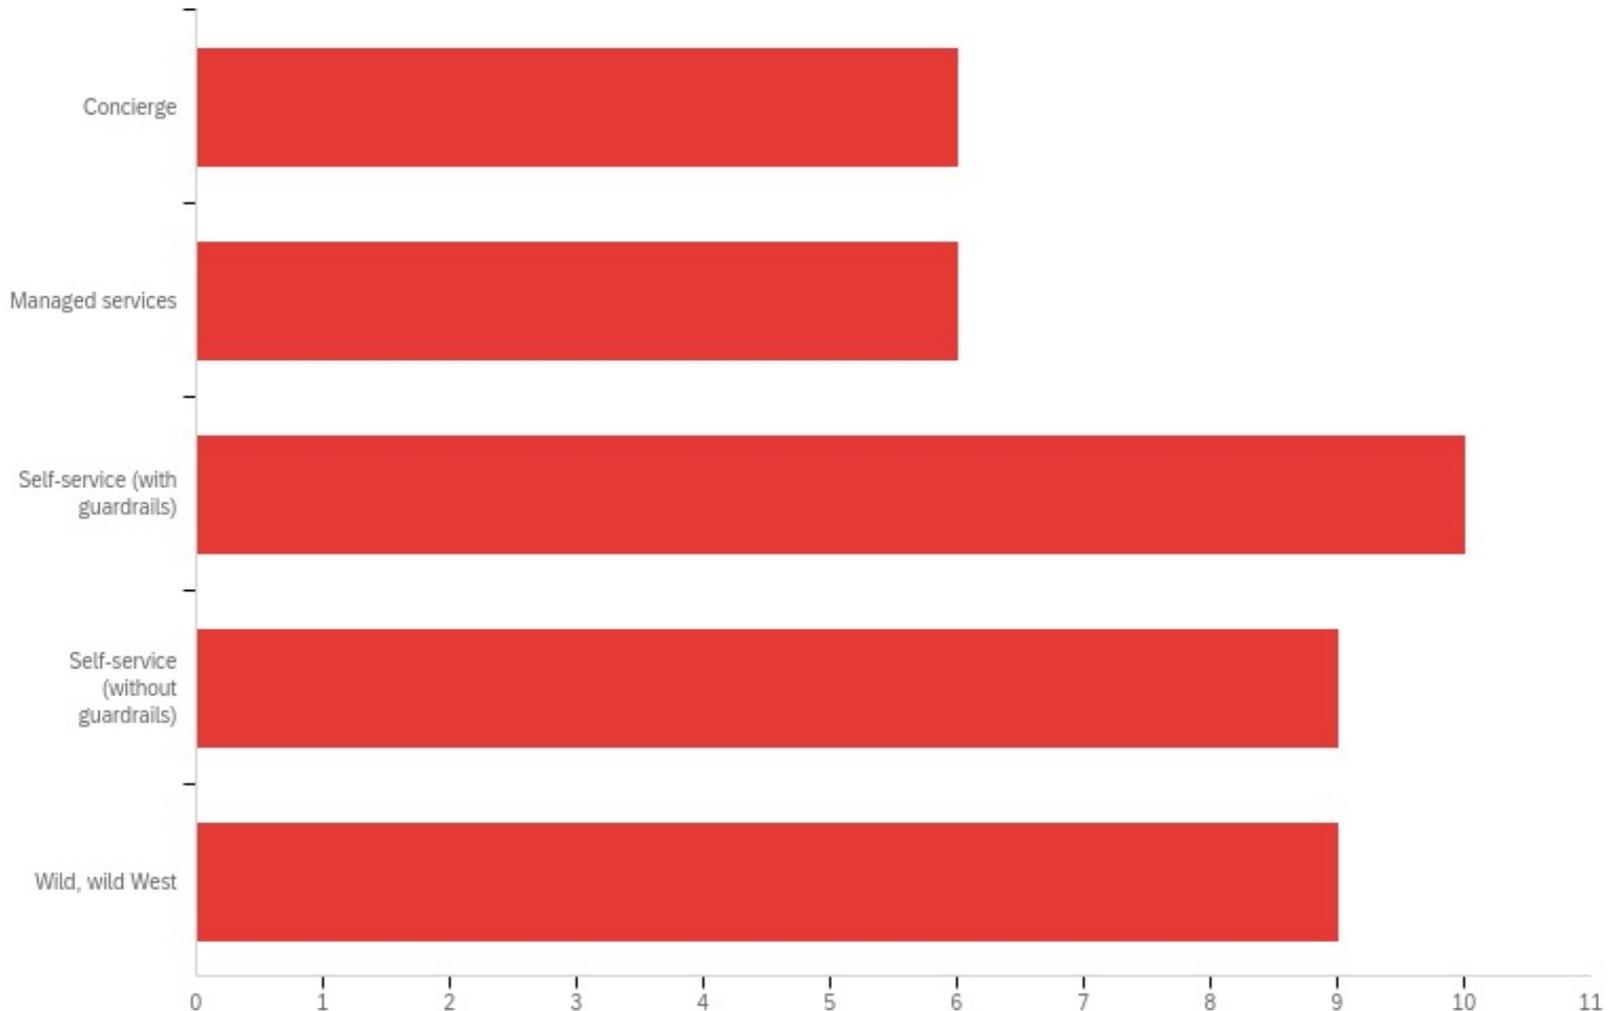
If you do have a central team providing Cloud RC support, what services do they provide?

No one service was common to all 32 respondents, but billing and consulting were nearly universal.

Answer	Count (of 32)
Consolidated billing	25
Consulting on how to best do research in the Cloud (e.g., advice on launching GPU instance families, advice on spot instances, etc.)	24
Cloud Account vending	19
Escalation to Cloud vendors when issues are encountered (e.g., vendor credit issues)	19
Consulting on how to secure research data in the Cloud	18
Cloud instance provisioning by central team for researchers	16
Concierge-level outreach to Cloud vendors (e.g., actively helping researchers identify options for and applying for credits, facilitating architectural sessions with vendors, etc.)	16
Self-service Cloud instance provisioning	13

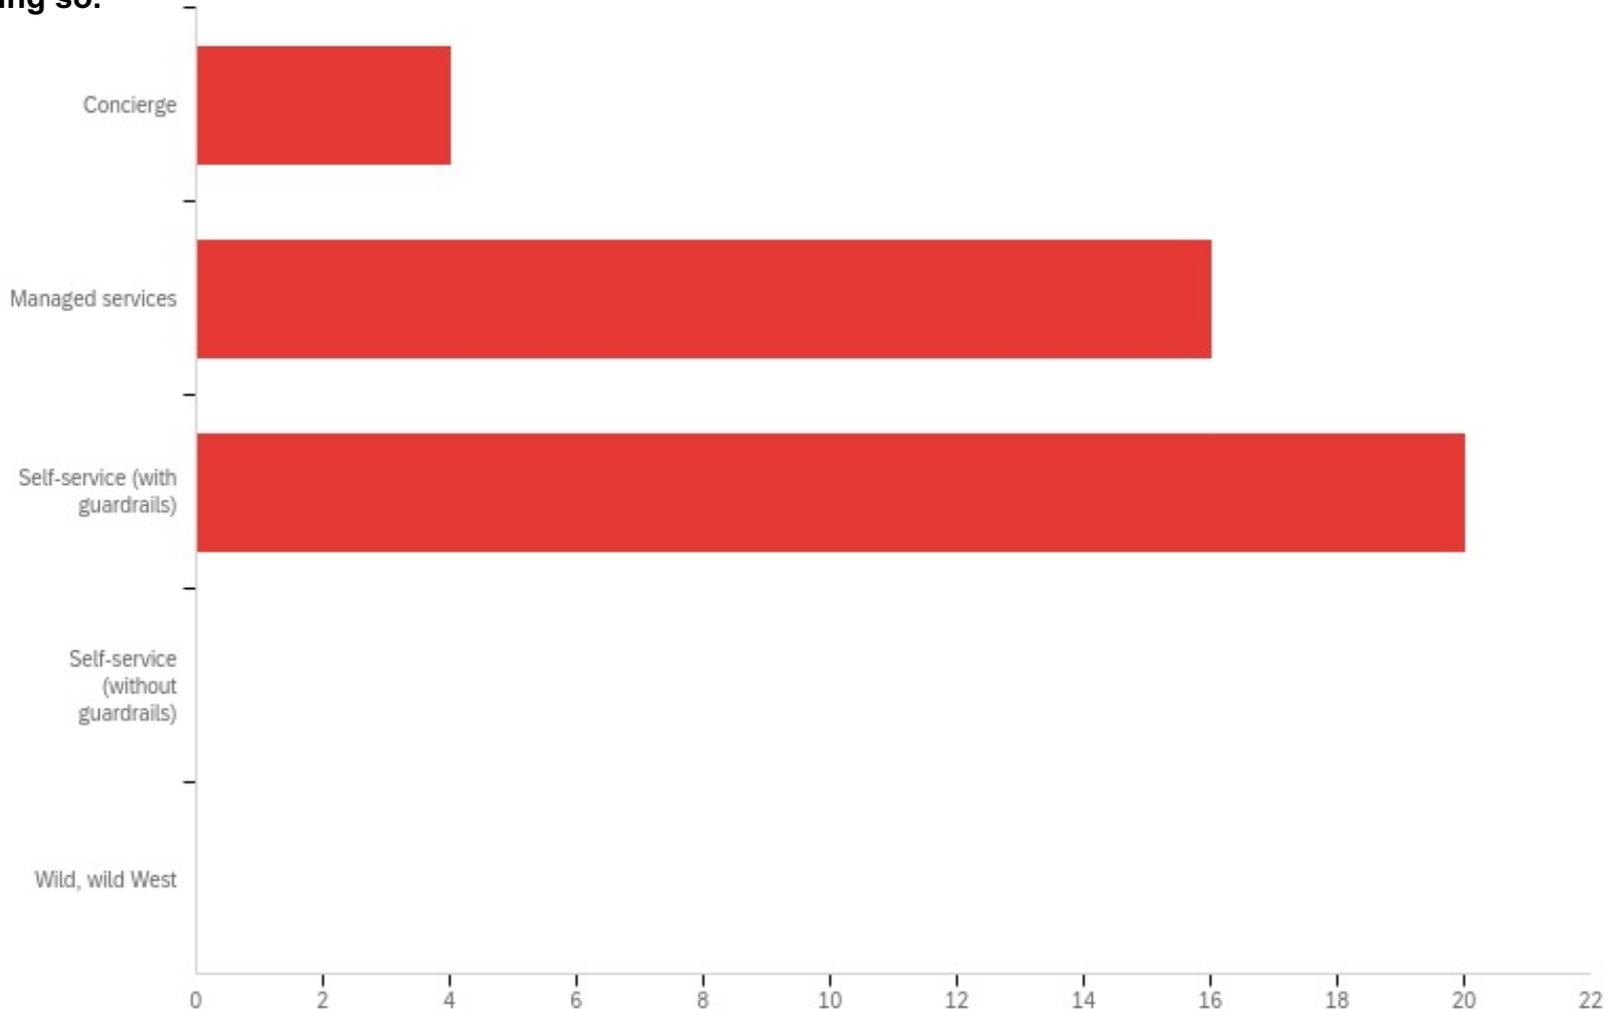
How would you characterize your institution's level of RC Cloud support in general?

Levels of support vary widely, though self-service with guardrails is the largest pattern.



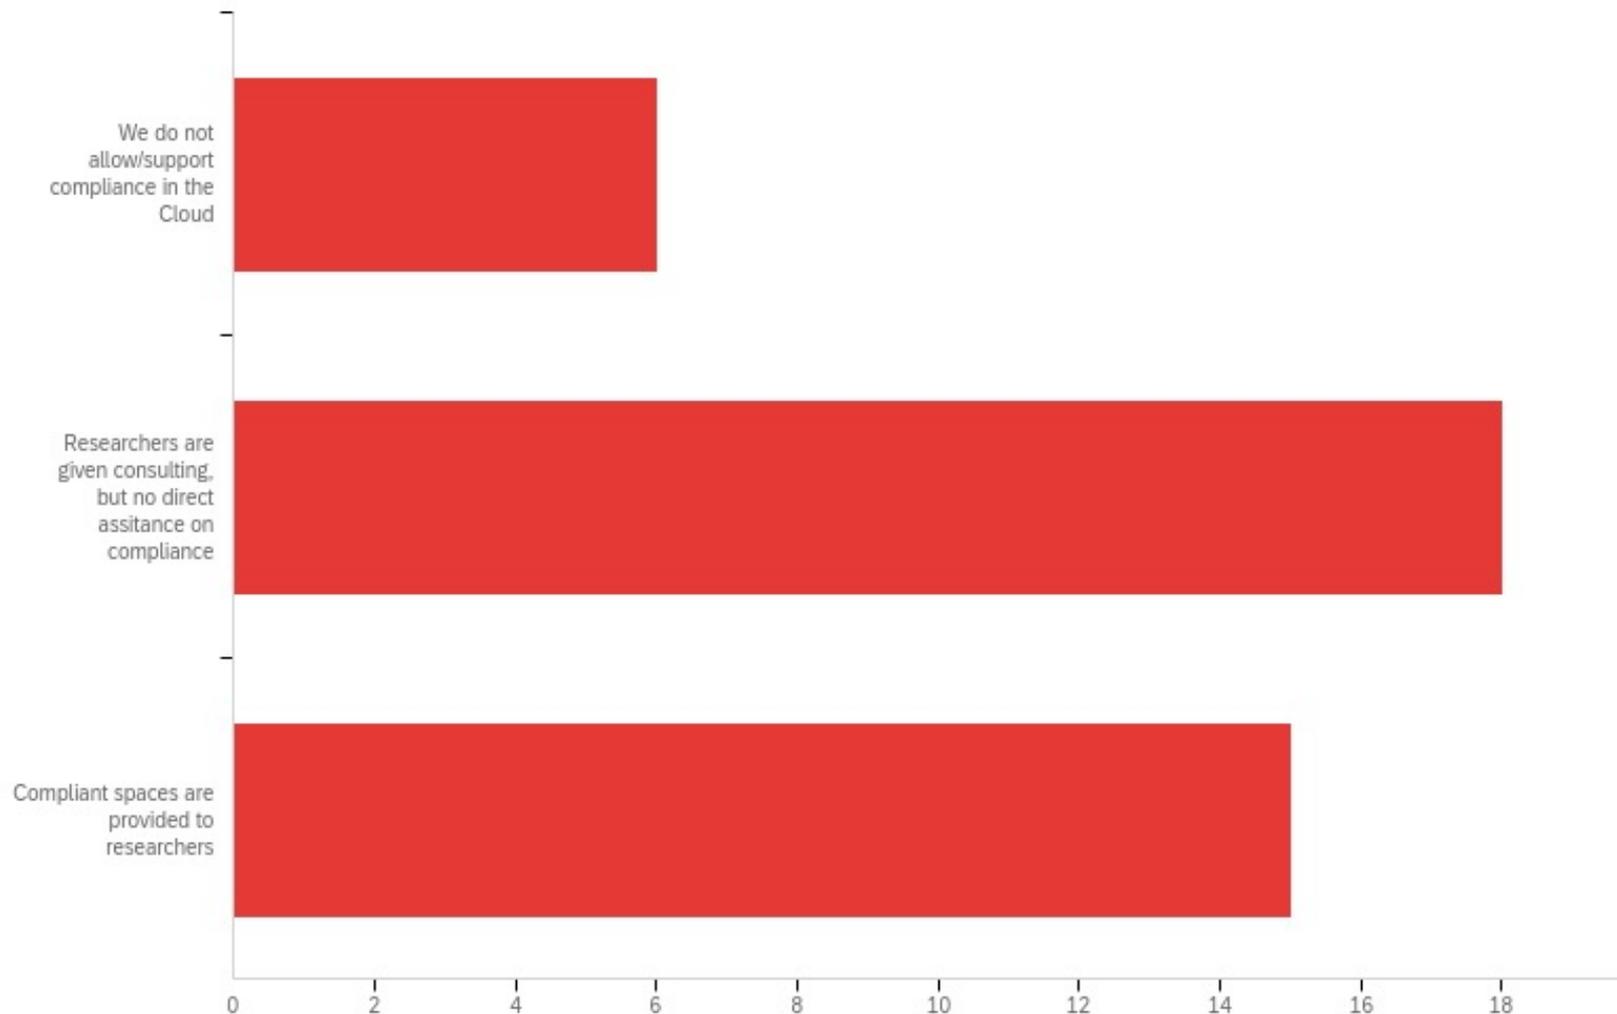
What level of Cloud RC support do you want to (realistically) achieve?

Nobody has their sights set any lower than self-service with guardrails, with half looking to provide an even higher level of service. Some groups already providing concierge service would like to stop doing so.



How does your institution handle compliance regimes in the Cloud (e.g., ITAR, CUI, HIPAA, etc.)?

Elaboration comments suggested that a number of groups are actively trying to upgrade from the “consulting” to the “compliant spaces” levels



What drives your researchers to the Cloud?

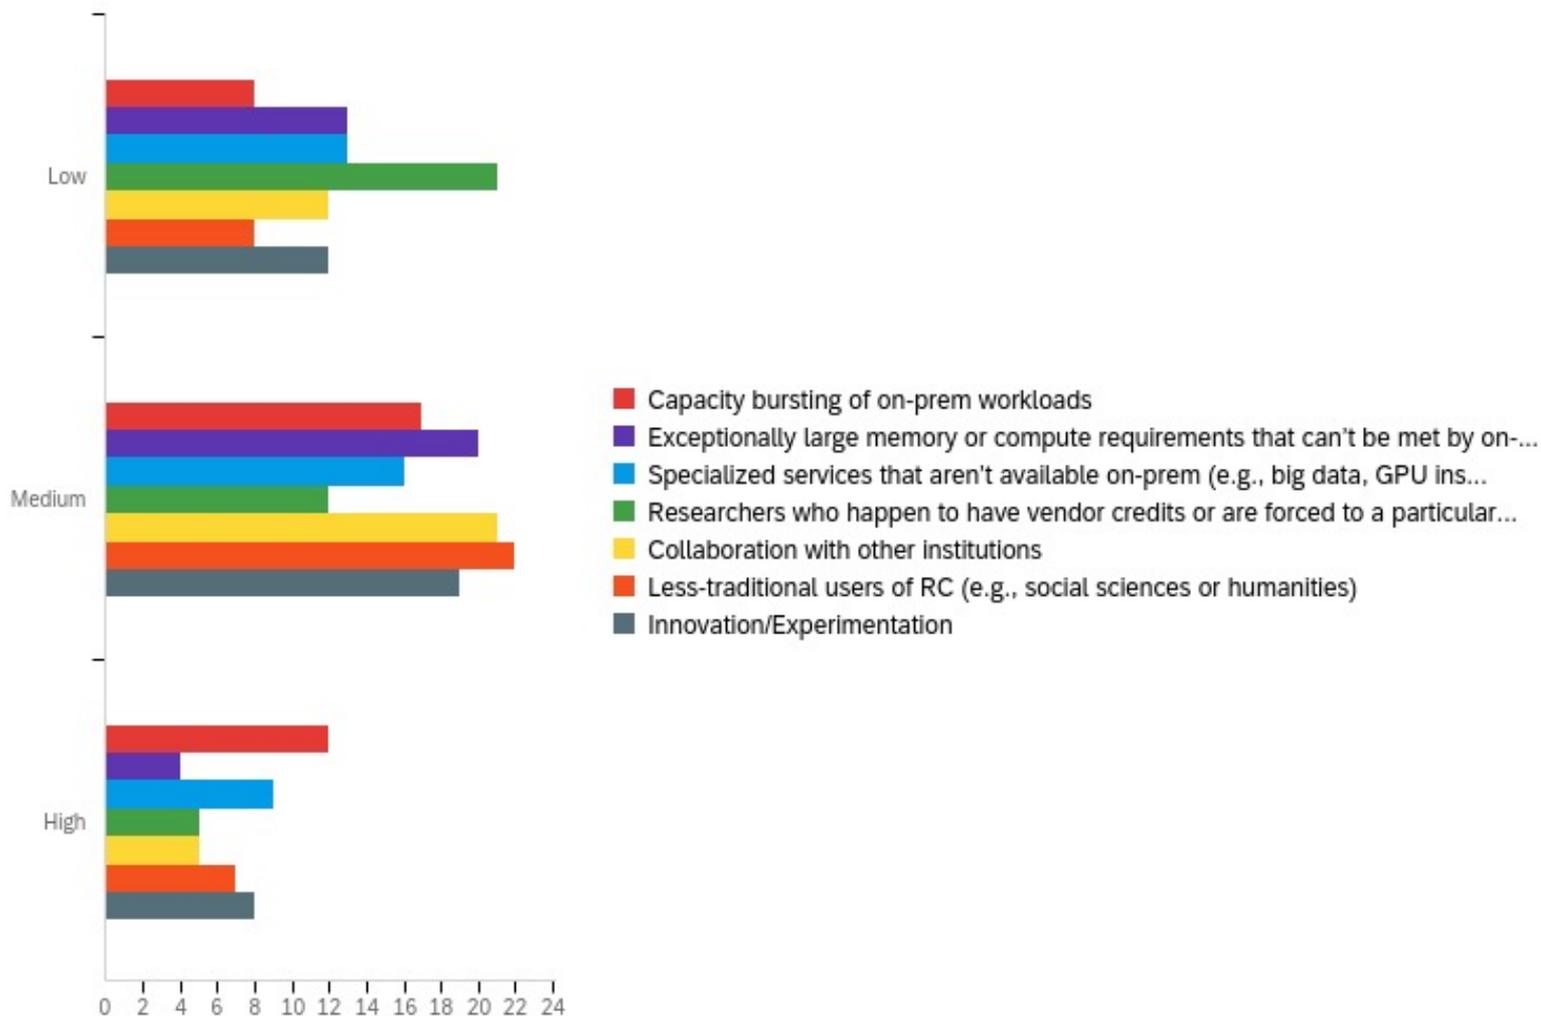
If Cloud is not the primary target for RC, when is it generally used?

No one reason was common across all respondents, but Innovation/Experimentation, Specialized Services, and Vendor Credits led the pack.

Answer	Count (of 38)
Innovation/Experimentation	28
Specialized services that aren't available on-prem (e.g., big data, GPU instances)	27
Researchers who happen to have vendor credits or are forced to a particular vendor by terms of a grant	25
Collaboration with other institutions	16
Exceptionally large memory or compute requirements that can't be met by on-prem	15
Less-traditional users of RC (e.g., social sciences or humanities)	12
Capacity bursting of on-prem workloads	11

Where are the challenges in the Cloud RC space?

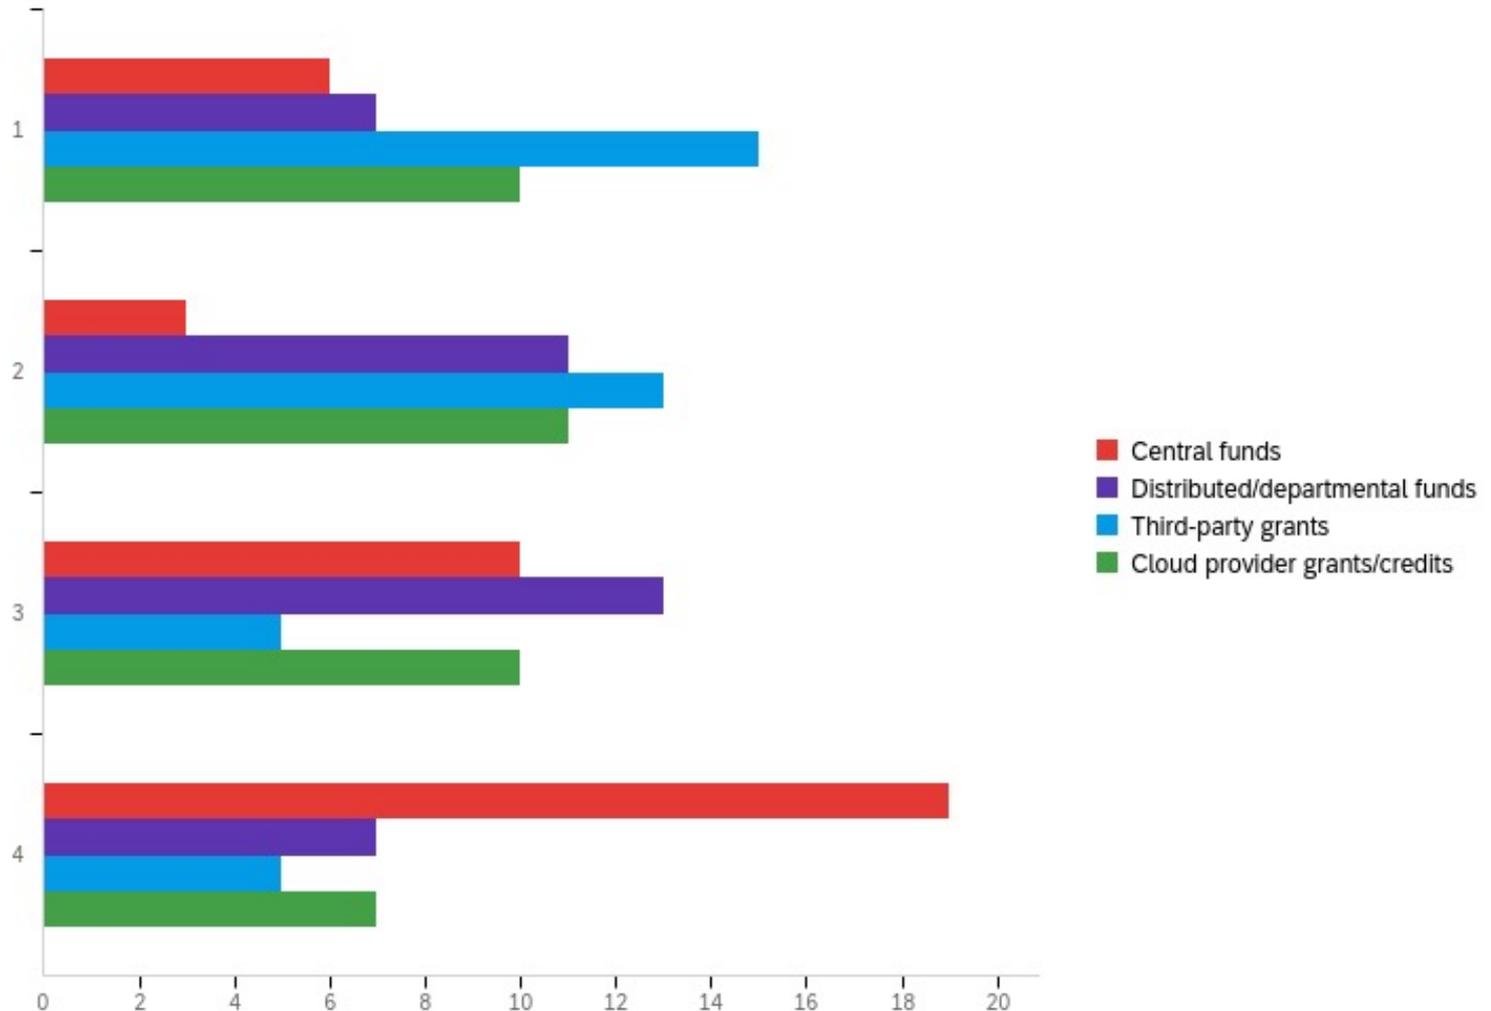
Capacity bursting was the largest challenge, and vendor credits was the smallest, but there is clearly plenty of challenge across many areas.



Funding

Please rank how RC in the Cloud is generally funded at your organization

Ranking the options by their mean values, they rank 1) Third-party grants, 2) Cloud provider credits, 3) Departmental funds, 4) Central funds



What does your institution do about indirect cost recovery for research computing?

While some institutions have indirect cost recovery and some do not, nearly all have the same policy for Cloud and on-premises spend.

