We thank the editor Dr. Farneti in handling our manuscript, and Dr. Griffies, Dr. Bell, Dr. Hogg and Dr. Hirschi for their positive and constructive comments. We have acknowledged their work in the Acknowledgements section. Please find our point-by-point reply below in red text.

Referee #3 (Andy Hogg)

This paper advocates for a cloud-based strategy to address the problems in sharing and analysing the large volumes of data that emerge from high-resolution ocean model simulations. I found the paper to be interesting and well-written, and concur with two previous reviewers that this manuscript is a worthwhile contribution to the literature. I have some minor comments, which the authors may like to take into account, listed below. I would be happy to recommend publication if these issues are addressed.
We thank the referee's positive and constructive comments.

- Line 8-9 - Consider deleting the sentence naming the 5 models from the abstract?
  Done.

- Section 2 — I found this description of the process of sharing data, and the ARCO format, to be particularly useful. But one thing I don't understand is whether the authors are arguing the Zarr files produced here are optimal for all operations. For example, if I wanted to filter with FFTs, average in time or average in space, would the Zarr chunking remain optimal for all of these operations? Or is there a trade-off between operations?
  We indeed advocate for the Zarr format for all datasets where parallelized analyses are anticipated. Zarr format scales much better (depending on the computational architecture, up to orders of magnitude) than the NetCDF format for parallelized computation without any trade off.

- Line 115 — Improved GS separation is a nice feature, but the global ocean is bigger than just the North Atlantic and there are many more processes revealed by resolution than WBC separation. I'm not convinced that separation is more "key" than other processes that are improved with resolution. Maybe just back away from this statement a little?
  We have added: "in the North Atlantic" in the sentence acknowledging that the GS separation is not a global feature.

- Line 122 - "this will…" is a little ambiguous.
  We have rephrased this as: "the 3D diagnostics will".

- Line 146 — my recollection is that the tides in LLC4320 had a bug in the tidal forcing which overestimates the tidal magnitude (but I apologise that I can't put my hands on the appropriate reference). I suggest the authors check on this issue as they revise the manuscript.
  The reviewer is correct. We have added in lines 159: "Also note that tidal forcing in the LLC4320 simulation was inadvertently overestimated by a factor of 1.1121."

- Line 180 — "… the two …" - also ambiguous.
  We have rephrased this as: "between the submeso- and meso-scale".

- Line 191 - there is a case made about daily-averaged submesoscale fields, but it wasn't clear (to this reader) where these daily-averaged fields were used in this paper?
  The 3D diagnostics are based on the daily-averaged fields. We have added this as: "... using the daily-averaged outputs" in line 173.

- Line 222 — "This presents …" ambiguous …
  We have rephrased this as: "The smaller predicted values presents…"

- Figure 6 — On a first read, I was amazed at the similarities between the parameterised submesoscale fluxes and measured buoyancy flux. Actually, it looked too good to be believable. But when I looked at D1 the comparison was underwhelming. I suspect the use of the spatial median in Fig 6 is unfairly favouring the comparison. I would prefer the authors to show D1 as the main figure, or perhaps show both in the main text, for a warts-and-all comparison of the parameterisation.
  We have put both figures 6 and D1 (now Figure 7) in the main text. We have also added in the Conclusions that the agreement between the submesoscale flux and its prediction from the parametrization are "in the spatially averaged sense" in line 310.

- Section 5 — The authors make some good points here and I agree with most of them. But I found the approach to be slightly evangelical. Fundamentally, the argument seems to be "we have found the best approach, but if the scientists/funders don't back us then it will fail". I agree that the approach espoused here is good, and I would like to advocate for it myself. But a more dispassionate discussion of the pros and cons would probably be an advantage here. For example, a significant disadvantage here is the risk that the Google Cloud Platform is discontinued or unavailable to researchers in some nations, for whatever reason. That is not such an outlandish proposition, but could be catastrophic for an open platform like this. There are other risks of equal access, long term funding, etc. I am just asking here for a more objective analysis of the risks here — which would be a greater service to the reader than the advocative approach.
  A core design principle of both Pangeo Forge and (IIUC) JupyterHub is being cloud vendor agnostic. So while the JupyterHub for this project happened to be hosted on GCP, and the data happened to be hosted by AWS (i.e. OSN), there is nothing about the underlying technologies — Pangeo Forge and Jupyter — which require these vendors. Unless the future comes down to no major cloud vendor providing such services, we believe the cloud-based framework to be robust.
  The JupyterHub that 2i2c runs is also intentionally designed to be cloud-agnostic and none of that technology is dependent on Google. This is codified in 2i2c's "Right to Replicate" principles: https://2i2c.org/right-to-replicate/, and by the constraint that all of the technology behind the hub is community-driven and vendor- and platform-agnostic. We have noted this in lines 279-284 as: "We would like to note that while we have chosen GCP and OSN for the cloud platform, the core design principle and technology behind Pangeo Forge and JupyterHub operated by 2i2c are non-proprietary and cloud vendor agnostic (for example, as defined in 2i2c's "Right to Replicate", https://2i2c.org/right-to-replicate). We could re-deploy the entire cloud platform on a

different cloud provider with relative ease. This lets the users of this platform benefit from the flexibility and efficiency of the cloud, while minimizing the risk of lock-in and dependence on proprietary technology.”