Review Responses for: Improved Advection, Resolution, Performance, and Community Access in the New Generation (Version 13) of the High Performance GEOS-Chem Global Atmospheric Chemistry Model (GCHP)" by R. V. Martin et al.

We thank editor for the comments, which have helped improve the quality and clarity of our manuscript. We have responded to each comment below. The original comments are in black, our responses are in blue and the changes to the manuscript text are in *blue italics*.

Comments to the author:

One of the reviewers has now checked your revised submission, and now both recommend publishing your manuscript. After reading it, I agree. However, before accepting it, I think a few issues must be fixed. I list them next:

- First, Spack is a core part of your deployment, and you have adapted a version to be used with GCHP; therefore, I think you must publish it in Zenodo, as you did with the model code. In this way, it is important that you name it with a version number.

We added to the Code Availability section, the link to the Spack post on Zenodo. We also updated the GCHP version link.

- I find it highly interesting that you have tested the model in AWS, and you spend a whole subsection of the manuscript describing it and discussing this capability. However, I think readers would benefit from putting your results into context. A behaviour that has been observed before is that cloud environments can outperform local supercomputers for a low number of cores (see, for example, Montes et al., 2019; I understand that this could be self-serving from my side, but I am not aware of many of these evaluations for climate models). This is probably related to the selection of the cloud infrastructure, where it is possible to allocate a few cores connected with a fibre channel, and when the number of cores exceeds the capability of servers in the same rack, cores allocated could be in a different data centre. Something in this way can be observed in your figure 5 for AWS C90 and C180 (independently of the use of IMPI or OMPI). I think it is important that this is noted in the paper, as it is something that eventually could have a more significant impact on performance than, for example, others that you mention, such as the use of containers.

We added the following at the introduction to this section.

Cloud computing has been able to outperform local supercomputers for a low number of cores (e.g., Montes et al., 2020), but HPC applications with intensive internode communication have previously not scaled well to a large cluster on the cloud (Mehrotra et al., 2016;Coghlan and Katherine, 2011).

Please, when you mention Git, cite the work explaining it: Torvalds (2014)

Reference added at line 395 where Git is first mentioned.