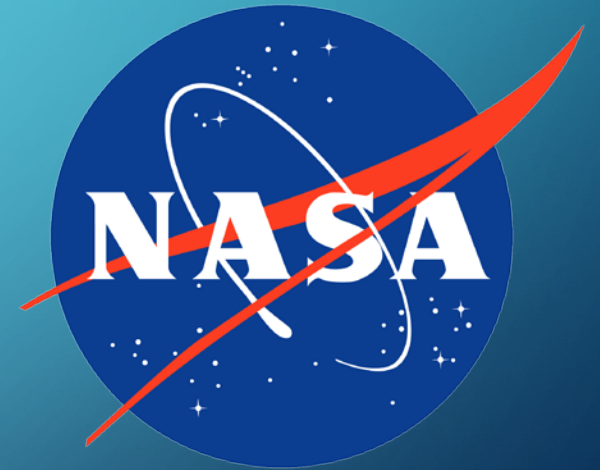


# ANALYZING PERFORMANCE OF CONTAINERIZED CLIMATE MODELS IN SINGULARITY ON A SUPERCOMPUTER

COMPUTER SCIENCE / IT

LAURA CARRIERE, CODE 606.2

NOAH OLLER SMITH, PURDUE UNIVERSITY



# INTRODUCTION AND BACKGROUND

- Software is currently deployed in HPC via modules and package managers
- Containers provide easy software deployment and different environments
- Singularity vs Docker - Singularity is made for HPC and can make use of docker resources.



This is Discover, a supercomputing cluster with over 129,000 cores for NCCS!

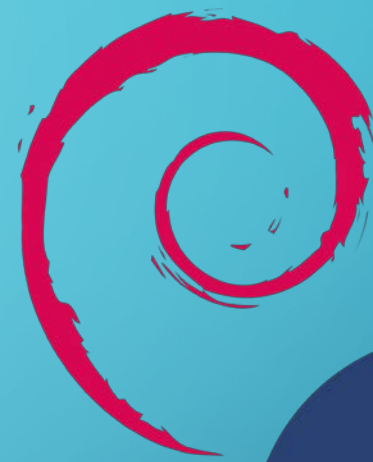
# QUESTION / INVESTIGATION

- Testing to see whether a containerized version of a part of GEOS (Goddard Earth Observing System) can compete with a non-containerized version
- Mainly looking at performance impact on containerized versions
- GEOSgcm and GEOSfvdycore were containerized
- GEOSfvdycore was containerized for testing



# METHODS AND APPROACH

- Had to learn usage of Singularity & Slurm – Used Docker and PBS before NASA
- Created multiple containers for testing, including Debian 10, Fedora 33, and CentOS 7
- Before containerizing GEOSfvdycore, the base libraries had to be containerized first



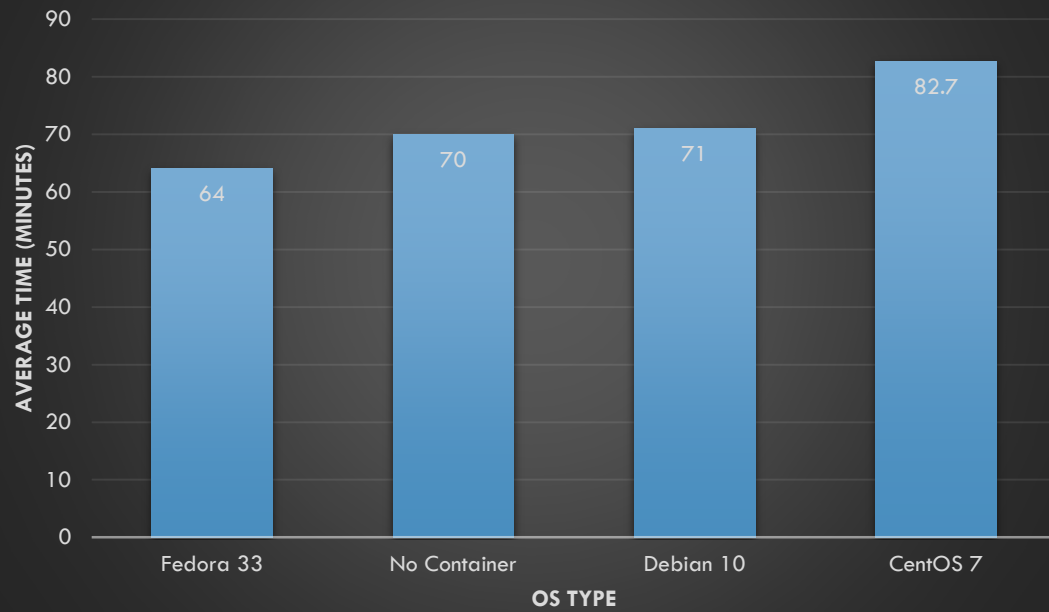
# METHODS AND APPROACH

- Two parts to a run: Actually setting up the experiment and then running.
- Wanted to see longer and shorter compute times.
- Horizontal resolution of experiments was changed to get a better idea of efficiency
- Each container and the non-containerized version got 3 runs on a single node for longer and shorter compute times

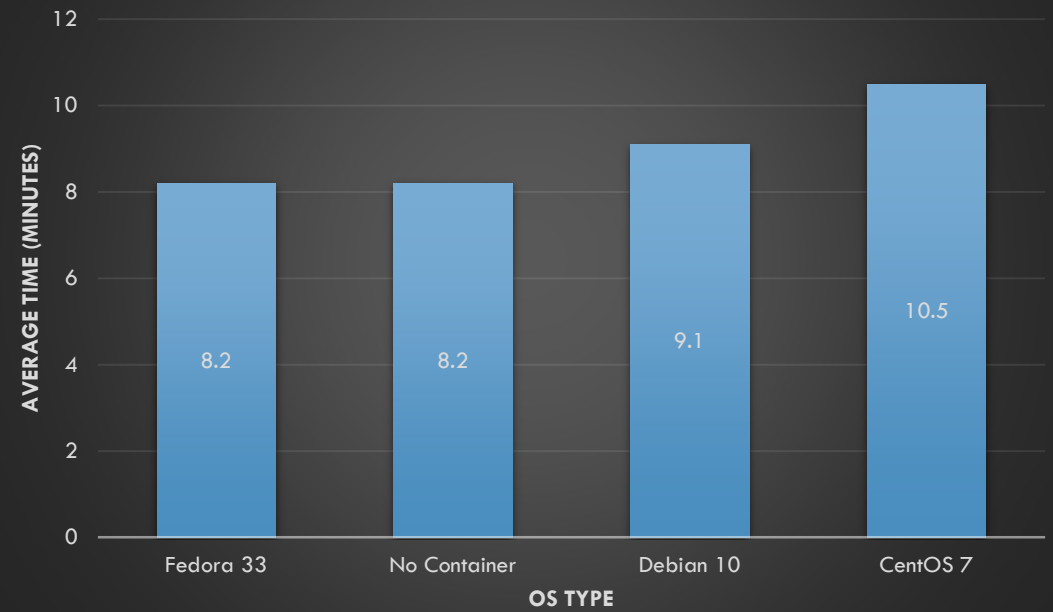


# RESULTS AND GRAPHS

**OS Type vs. Average Time for Longer Computations**



**OS Type vs. Average Time for Shorter Computations**



# DISCUSSIONS AND CONCLUSIONS

- Tests definitely indicate that containers can compete with non-containerized versions
- Containers could have done well from having up-to-date software libraries that GEOSfvdycore depends upon compared to Discover
- Topics to move forward on this project would include minimizing the container size and improving how models can be executed

# REFERENCES AND ACKNOWLEDGEMENTS

- Singularity documentation <https://sylabs.io/docs/>
- Slurm documentation <https://slurm.schedmd.com/documentation.html>
- MPI Tutorial <https://mpitutorial.com/>
- GEOS-5 wiki [https://geos5.org/wiki/index.php?title=GEOS-5\\_Earth\\_System\\_Modeling\\_and\\_Data\\_Assimilation](https://geos5.org/wiki/index.php?title=GEOS-5_Earth_System_Modeling_and_Data_Assimilation)
- Special thanks to Laura Carriere, Matt Thompson, Kenneth Peck, Jordan Caraballo-Vega, and Bruce Pfaff