

# Performance Analysis of Large Scale HPC Workflows for Earth System Models



**Approved for public release.  
Distribution is unlimited.**

Benjamin Mayer (ORNL)  
Joseph Kennedy (ORNL)  
Katherine Evans (ORNL)  
Jeff Duracha (NOAA)

**November 8, 2017**

#### DOCUMENT AVAILABILITY

Reports produced after January 1, 1996, are generally available free via US Department of Energy (DOE) SciTech Connect.

Website <http://www.osti.gov/scitech/>

Reports produced before January 1, 1996, may be purchased by members of the public from the following source:

National Technical Information Service  
5285 Port Royal Road  
Springfield, VA 22161

**Telephone** 703-605-6000 (1-800-553-6847)

**TDD** 703-487-4639

**Fax** 703-605-6900

**E-mail** [info@ntis.gov](mailto:info@ntis.gov)

**Website** <http://www.ntis.gov/help/ordermethods.aspx>

Reports are available to DOE employees, DOE contractors, Energy Technology Data Exchange representatives, and International Nuclear Information System representatives from the following source:

Office of Scientific and Technical Information  
PO Box 62  
Oak Ridge, TN 37831

**Telephone** 865-576-8401

**Fax** 865-576-5728

**E-mail** [reports@osti.gov](mailto:reports@osti.gov)

**Website** <http://www.osti.gov/contact.html>

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

COMPUTER SCIENCE AND ENGINEERING DIVISION

**PERFORMANCE ANALYSIS OF LARGE SCALE HPC WORKFLOWS FOR EARTH  
SYSTEM MODELS**

Benjamin Mayer (ORNL)  
Joseph Kennedy (ORNL)  
Katherine Evans (ORNL)  
Jeff Duracha (NOAA)

Date Published: November 8, 2017

Prepared by  
OAK RIDGE NATIONAL LABORATORY  
Oak Ridge, TN 37831-6283  
managed by  
UT-Battelle, LLC  
for the  
US DEPARTMENT OF ENERGY  
under contract DE-AC05-00OR22725



**Contents**

1 EXECUTIVE SUMMARY . . . . . 1  
2 INTRODUCTION . . . . . 3  
3 PROGRESS AND RESULTS . . . . . 5  
    3.1 DATA GATHERING . . . . . 5  
    3.2 PERFORMANCE ANALYSES . . . . . 5  
    3.3 DEVELOPING THE WORKLOAD MODEL . . . . . 6  
4 NEXT RECOMMENDED STEPS . . . . . 11  
5 CONCLUSIONS . . . . . 13  
ACKNOWLEDGMENTS . . . . . 15  
REFERENCES . . . . . 17



## 1. EXECUTIVE SUMMARY

Climate assessment reports, such as the Intergovernmental Panel on Climate Change's 5th assessment report [IPCC AR5; IPCC et al., 2014] or the United States National Climate Assessment [U.S. Global Change Research Program, 2014], increasingly rely on many-ensemble climate projections from fully-coupled Earth system models (ESM) to assess possible future climate states. ESM participate in a large set of Coupled Model Intercomparison Projects (CMIP) to better understand their models and the climate system, so that they may provide credible climate projections. However, fully participating in these intercomparison exercises represents an extremely large computational personnel burden for modeling centers [Eyring et al., 2016, Williams et al., 2016a]. However, there has been little attention within the climate modeling community paid to the *in situ* performance of ESMs up until now. Recently a computational performance model intercomparison project (CPMIP) for CMIP6 has begun [Balaji et al., 2017], but there has not yet been an effort to gather, process, and analyze the historical performance data that has been saved within archived log files from previous modeling efforts, nor the development of performance analysis system capable of providing detailed statistical analyses or visualizations of this performance data.

Therefore, we've begun a joint effort to analyze the historical and current performance data on both GFDL and OLCF modeling systems. We have developed a preliminary performance analysis system to capture, analyze, and visualize end-to-end *in situ* performance data and applied it to NOAA/GFDL's post-processing system. Through the course of the project a number of performance improvements have been made, with the steps necessary to handle the expected volume of data (Big Data) likely to come out of the CMIP6 exercises (TB to PB) provided. We have also begun developing a statistical workload model that will allow modeling centers to determine their expected scientific throughput in the context of expected resource (computational and personnel) availability and accounting for typical workflow disruptions. This will allow them to develop a CMIP participation strategy and maximize their level of participation in the exercises.

We expect this infrastructure, and the data it produces, to ultimately allow us to identify performance issues and workflow errors that exist across Earth system modeling centers, and possibly HPC modeling centers in general. Center-independent solutions to workflow or performance issues would help alleviate the computational burden for CMIP6+ modeling projects and should realize both labor and cost savings benefits throughout the CMIP community.





## 2. INTRODUCTION

Prediction of future climate states, such as for the IPCC Assessment Reports [IPCC et al., 2014] or the U.S. National Climate Assessment [U.S. Global Change Research Program, 2014], relies on the many ensembles of climate projections provided by fully-coupled Earth system models (e.g., E3SM, CESM, etc.). In order to provide credible projections, Earth system modeling centers around the world participate in a large set of Coupled Model Intercomparison Projects (CMIP) to better understand how the range models respond to changes in radiative forcing, to investigate the sources and consequences of model bias, and to understand the variability of the climate system [Eyring et al., 2016]. These kind of intercomparisons are most effective when there is wide active participation from the entire ESM community, however, because ESMs are (ever increasingly) complex and run at (ever increasingly) high resolution to provide the “best” projections, fully participating in these intercomparisons represents an extremely large computational burden for modeling centers [Eyring et al., 2016, Williams et al., 2016a].

For example, CMIP3 provided 36 TB of data for IPCC AR4 and each modelling group submitted on average 1750 years of model output, yet for CMIP5, the core experiments alone required 2300 years of model output and the data archive is estimated to exceed 3 PB of data [Taylor et al., 2012]. Although CMIP6 recognized this large burden, it is still expected to provide over 50 PB of data [Williams et al., 2016b] and require a minimum of  $\approx 1000$  simulated years to participate in the required core exercises, let alone all 21 endorsed model intercomparison exercises [Eyring et al., 2016]. Since current ESMs range from  $\approx 10^3$  to  $\approx 10^5$  core-hours per simulated year during scientific application [Balaji et al., 2017], minimally participating in CMIP6 will require  $\approx \times 10^7$  core-hours, all but requiring the use of high-performance computers at Leadership Computing Facilities (e.g., Titan at the Oak Ridge Computing Facility (OLCF)).

Further complicating participation in these projects is the intersection of HPC system turnover and these CMIP exercises. In the U.S., 3 major HPC systems are transitioning to a new system during CMIP6, and will likely see the development of 3 more systems during CMIP7 \*. Because the transition between these machines represent a substantial amount of time and effort for both verification and performance tuning, they will likely represent a significant barrier to scientific throughput even though these centers are expected to keep their current systems, they experience more performance issues as they age [Esteban Meneses, 2015, Tiwari et al., 2015]. We can expect this to be common for all large systems at all computing centers charged with producing state-of-the-art computational science. Incumbent with data production and processing is the large support required from workflow science and engineering to manage the simulations and analysis. This support is typically stood-up at every center and can be model dependent, which multiplies the efforts across all CMIP6 participants, even though much of the supporting technology is center-independent.

Due to the large computational and personnel resource burden associated with participating in these CMIP exercises, and the dynamic availability of resources, it is critical that centers develop a coordinated strategy to meet these participation challenges. Fundamentally, these strategies will center around optimization of scientific throughput and allocation of expected resources, both of which require a detailed understanding of the expected end-to-end model performance. Although a computational performance model intercomparison project (CPMIP) for CMIP6 has begun [Balaji et al., 2017], there has been very little *in situ* end-to-end performance data analyzed and published to date. In fact, the preliminary results provided by Balaji et al. [2017] are the only cross-center performance data analyzed and published (of which we are aware). Although, much of this data has been collected for internal-use only at both the Oak Ridge

---

\*<https://energy.gov/downloads/fact-sheet-collaboration-oak-ridge-argonne-and-livermore-coral>

Leadership Computing Facility (OLFC) and the National Climate-Computing Research Center (NCRC) (and likely many others) for years. For example, log files from climate workflows and components are currently being generated, captured, and saved at the OLCF, NCRC, GFDL (post-processing only), NERSC and ALCF. However, robust analyses for this performance data have been prevented by the lack of formal standardization of the output data, the lack of mature processing software to generate the desired data, and the sheer volume of data that exist.

Therefore, a joint NOAA/GFDL-ORNL effort to gather, process, and analyze both the historical and current performance data at both OLCF and GFDL has been underway. We have developed a flexible and efficient performance analysis system that can be applied to the performance data at these centers and will handle the large volume of data. We are able to generate many descriptive statistics and visualizations key to investigating a center's workflow. We present a description of our preliminary performance analysis system and the results of our initial analysis. Furthermore, we are developing a workload prediction model which will simulate scientific throughput taking into account *in situ* performance, resource availability, and workflow disruptions (e.g., system down-time, system performance degradation, queue backlogs, etc.).

### 3. PROGRESS AND RESULTS

#### 3.1 DATA GATHERING

NOAA/GFDL's workflow management system, FRE (the FMS Runtime Environment), logs  $\approx 1.5$  million lines of text-based output for each simulation, which contains a wide variety of workflow and performance data including queue dwell time, model run time, allocated nodes, and error messages. These logs are left in place on the NCRC file system unless a user takes action to delete them, and are world-readable. In order to gather the historical data, we developed the capability to crawl the file system and generate a list of log file directories to be analyzed. This is given a low execution priority so it does not disturb current operations. Note that, due to the size of the file system, it took  $\approx$  one month to span the system, and it found over 800 GB (40 GB compressed with gzip) of log files.

Through manual examination of a small sample of log files, we determined the overall file structure, identified variables of interest (e.g., queue times, run times, etc.), and noted how those variables could be identified. A python module was developed to parse these log files for variables of interest, which are located in the log files via a regular expression. An associated parsing function is used to pull in the requested data. This setup allows for variables to be added and removed simply by updating the list of regular expressions and their associated parsing functions. Using an associated function allows complicated data (e.g., split across multiple lines) to be handled. As the data is generated, it is stored in an SQLite database, which allows the parsed data to be easily queried, grouped, and associated. Parsing the 800 GB of log files and inserting them into a database takes several weeks.

Although useful, scalability is required, so we realized a significant performance improvement by hashing the first 7 characters of the key strings and building a hash set, which was long enough to prevent most collisions. This is because many of the regular expressions were very simple, so looking for key strings at the beginning of lines was sufficient. Then these key strings can be compared to the hash set instead of running the regular expression. In the event of a match, the matched line can be sent to a function for processing, and in the case of collisions, the function can disambiguate the matched variables. After the hash-comparison is exhausted, the remaining regular expressions that did not lend themselves to hashing is checked. This reduces the time to parse the 800 GB of log files by three orders of magnitude, and should now be scalable to PB of data.

#### 3.2 PERFORMANCE ANALYSES

For NOAA/GFDL's experiments, a particular configuration of a climate simulation is described as a "case", which comprises the numerical methods, model components, model resolution, and initial inputs. A case is, or cases are, then run multiple times to complete the desired simulations for a particular experiment. The outputs of these simulations are considered 'raw' (not yet in the form needed for scientific analysis) and will need to be post-processed (e.g., calculation of derived variables, reordering data to produce time series, providing summary statistics and figures). This post-processing, like the model runs themselves, requires HPC level infrastructure and in fact both NOAA/GFDL and ORNL have separate systems for post-processing model results to enable access to more memory and specialized software. We focus on this class of systems for performance analysis because they also experience ongoing attrition and are bottlenecks during CMIP6 simulations.

Since post-processing transformations are typically constant within a case and are often repeated (for each simulation), a very large set of repeated timing data has been generated. This bolsters our statistical

**Table 1: Example descriptive statistics of the post-processing performance for selected cases.**

Case	N	Diagnostics CHPY			
		Max	Min	Mean	Std dev.
CM4 (all)	3033	2.083	0.001	0.089	0.028
CM4 historical	226	0.762	0.003	0.041	0.019
CM4.5 (all)	58	1.877	0.003	0.271	0.040

analysis of the post-processing performance. This data is primarily composed of the instance when: a post-processing job enters the queue, starts processes, and completes processing, not to mention the substantial amount of data describing the case setup, which is used for classification of our results. Table 1 provides an example of the descriptive statistics that can be produced to understand the general performance of different cases. Table 2 provides an example of the descriptive statistics that can be produced for a specific case to better understand the performance of its diagnostic components.

Importantly, descriptive statistics cannot provide a complete description of the data and can be misleading in even simple datasets [e.g., Anscombe’s quartet; [Anscombe, 1973](#)]. Therefore, we also provide many visualizations of the data to help develop performance insight. For example, in [fig. 1](#), many interesting features of the data are readily apparent (e.g., gaps in the time series, vertical spikes of increased run times, outliers, and an apparent shift in the performance around the 2016–2017 rollover), but would get overlooked by basic descriptive statistics. A root-cause analysis of these interesting features was performed, and some of the features can be explained by correlated system events. For example, the gaps in the time series could be caused by reasons including:

- The model was not being run
- Log files may have been moved, hidden, or deleted
- The post-processing job failed
- System failures

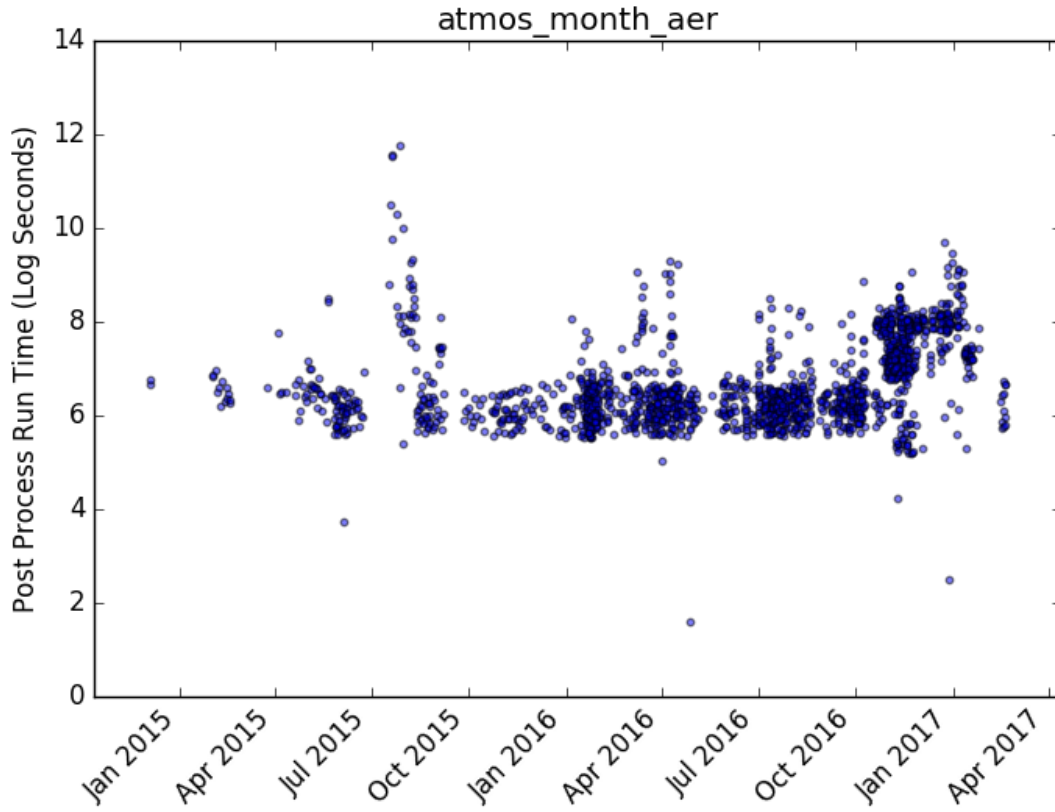
In order to determine the cause of the data gaps, we developed a dataset of post-processing system issues. By comparing the performance data and the system data, we found that the data gap in July 2015, which was followed by a large spike in processing time, was caused by a tape-system failure. This failure caused many jobs already in the queue, or entering the queue during the failure, to sit idle until the failure was resolved. When the tape system resumed normal operation, the post-processing jobs that had been sitting idle were able to retrieve their required data and finish processing, leading to excessively long run times logged. However, it took some time to flush out the accumulated backlog of jobs on the system. Thus, the gap in the data shows the duration of the system downtime and the width of the spike shows the how long it took the post-processing machine to clear the backlog. These types of downtime events are not uncommon at HPC centers and will likely impact the overall throughput during a large modeling project, although they are likely randomly distributed in time. By combining visualization, statistical analysis, and investigation in an iterative approach, it is our view that a much deeper understanding of performance and expected throughput will be gained.

### 3.3 DEVELOPING THE WORKLOAD MODEL

From the post-processing results, we were able to develop three successively more complex mean performance models (MPM) of NOAA/GFDL’s model. The first and simplest model ( $MPM_1$ ) is to simply

**Table 2: Example descriptive statistics of the post-processing performance for the diagnostic components of the CM4 historical case.**

	Diagnostic	N	CHPY			
			Max	Min	Mean	Std. dev.
1	atmos	10	0.762	0.227	0.330	0.151
2	atmos_diurnal	7	0.048	0.022	0.032	0.008
3	atmos_month_aer	10	0.085	0.020	0.044	0.023
4	atmos_scalar	7	0.017	0.003	0.007	0.004
5	ice	10	0.142	0.026	0.051	0.037
6	land	10	0.521	0.094	0.159	0.123
7	land_cubic	7	0.189	0.149	0.171	0.012
8	land_instant_cubic	7	0.035	0.017	0.022	0.007
9	ocean_5daily	13	0.011	0.003	0.005	0.003
10	ocean_Agulhas_section	2	0.021	0.009	0.015	0.006
11	ocean_Barents_opening	2	0.016	0.007	0.012	0.005
12	ocean_Bering_Strait	2	0.014	0.007	0.011	0.004
13	ocean_Davis_Strait	2	0.015	0.008	0.011	0.003
14	ocean_Denmark_Strait	2	0.018	0.011	0.015	0.003
15	ocean_Drake_Passage	2	0.016	0.008	0.012	0.004
16	ocean_English_Channel	2	0.016	0.008	0.012	0.004
17	ocean_Faroe_Scotland	2	0.016	0.008	0.012	0.004
18	ocean_Florida_Bahamas	2	0.017	0.008	0.012	0.005
19	ocean_Fram_Strait	2	0.015	0.008	0.012	0.004
20	ocean_Gibraltar_Strait	2	0.016	0.008	0.012	0.004
21	Iceland_Faroe_U	2	0.036	0.008	0.022	0.014
22	Iceland_Faroe_V	2	0.024	0.008	0.016	0.008
23	ocean_Iceland_Norway	2	0.020	0.008	0.014	0.006
24	ocean_Indonesian_Throughflow	2	0.036	0.008	0.022	0.014
25	ocean_Mozambique_Channel	2	0.038	0.008	0.023	0.015
26	ocean_Pacific_undercurrent	2	0.018	0.008	0.013	0.005
27	ocean_Taiwan_Luzon	2	0.017	0.008	0.012	0.004
28	ocean_Windward_Passage	2	0.016	0.008	0.012	0.004
29	ocean_annual	10	0.167	0.029	0.063	0.047
30	ocean_annual_z	10	0.120	0.021	0.042	0.033
31	ocean_daily	7	0.042	0.030	0.036	0.004
32	ocean_monthly	10	0.261	0.040	0.077	0.063
33	ocean_scalar_annual	10	0.011	0.004	0.008	0.003
34	ocean_scalar_monthly	10	0.025	0.004	0.014	0.008
35	refineDiag	50	0.183	0.048	0.098	0.027



**Figure 1: Logged run time (job start to job completion) of atmospheric diagnostics over time of a small group of CM4 cases.**

take the total compute time of all the diagnostic components for each of the runs from a particular case, and then average the runs to find expected run time ( $t_1^e$ ). This model, however, doesn't account for different configurations within a case which may cause certain components to be turned on and off, or the distribution of the component costs.

A slightly more complex model ( $MPM_2$ ) takes the average compute time of each individual components for all the case's runs, and then adds them together to get the total expected runtime ( $t_2^e$ ). This method will provide a better estimate of  $t^e$  for a case with all components active, but isn't able to take into account a particular case configuration or set of them. (Note,  $MPM_2$  reduces to  $MPM_1$  in the case that there are an equal number of executions for each component.) However, if the expected configuration(s) of a case are known, we can use  $MPM_3$  to determine  $t_3^e$ , which multiplies the average component runtime by an execution weight, and sums them. For an individual run, the weights will be one for each active component, or zero for the inactive components. For a distribution of cases, the weights will be number of times the case is active, divided by the total number of runs (Note: if all components are active, this reduces to  $MPM_2$ ). Table 3 shows the expected runtime of the CM4 Historical case using each model.  $MPM_3$  produces a significantly shorter expected runtime because many of the diagnostic components are run infrequently (see Table 2). These estimates can be used determine the total required compute time ( $t^c$ ) needed for a particular modeling project or intercomparison exercise.

**Table 3: Expected runtime of the CM4 historical case using the three different workload models.**

Case	CHPY		
	$t_1^e$	$t_2^e$	$t_3^e$
CM4 historical	1.435	1.428	0.305






#### 4. NEXT RECOMMENDED STEPS

Our performance analysis system has provided key insights into, and a solid platform for, developing a robust cross-center performance analysis system. The next step to developing such a system will be to locate all available historical performance data and finish developing a set of robust parsers for the log files produced on: NOAA/GFDL's compute system, ORNL's compute system (Titan/Summit), and ORNL's post-processing system (Rhea). This will provide two independent end-to-end performance datasets, which will allow investigation into localized performance issues and cross-center performance issues. Furthermore, this dataset would allow us to produce the CPMIP metrics for historical simulation efforts (CMIP3-5) allowing for an investigation into performance trends which is not currently possible. A further important step would be to build the parser into current workflow management systems so that performance data from every run is automatically stored in a performance database (for each centers, or all centers).

As such a system is developed and deployed at modeling centers, the SQLite database will quickly grow very large (TB to PB by ISMIP7), leading to slow inserts, updates, and queries. To alleviate this issue, a distributed database designed for "Big Data" storage and queries should be stood up (e.g., Spark DB, Scality). Additionally, our python analysis methods could be adapted to a Big Data processing platform such as Apache Spark (via PySpark), and a real-time dynamic visualization system with a web interface could be developed to speed up analysis. We have developed a preliminary example of such a system using Dash by Plotly, which is a promising candidate even at scale and shown in Figure 2. Importantly, all of the aforementioned technology is available at ORNL through OLCF and CADES resources.



**Figure 2: Screenshot of an interactive visualized of the logged run time (job start to job completion) of atmospheric diagnostics using Dash by Plotly.**

Because an iterative development cycle lends to a more detailed performance analysis, many database operations will be performed repeatedly and as the database grows and these repeated operations will be more expensive. To reduce this burden, we've devised and tested a preliminary cache layer that can optionally be turned on and off and is placed between the database and analysis code. The cache layer saves a hash of any SQL query performed and the results of the query to a JSON file. Then, before any query is performed, the new query hash is checked against the previous hash set, and if it has already been performed, the results are loaded directly from its associated JSON file. Because hash look-ups and loading a JSON file is very efficient, the cache layer can significantly out perform the database. To fully implement this in production, the cache layer would likely need to be reset according to some rule (e.g., every database insert or update, at some frequency).

Furthermore, the workflow models can be significantly improved and expanded on, as the current model only relies on coarse statistics. A more robust model could actually track the entire workflow including the queue dynamics across different machines, apply system events at some frequency, and account for the performance variability that arises from the assigned processor layout on different machines. A robust estimation of the cost of a modeling project, in the context of other ongoing projects at the center, could then be better estimated from this model. The use of statistical methodologies like Monte-Carlo methods and machine learning algorithms should also be explored for these purposes.

## 5. CONCLUSIONS

Though the course of this project, we have developed a preliminary performance analysis system that is able to gather, process, analyze, and visualize end-to-end modeling performance data. This system has successfully been applied to NOAA/GFDL's post-processing system and is ready to be expanded to NOAA/GFDL's simulation system, and both ORNL's simulation and post-processing systems. Because of the large volume of data expected through the course of CMIP6, we have developed and outlined a number of methods to harden this system for "Big Data" (TB to PB). Furthermore, the development of a statistical workload prediction model, which may help modeling centers simulate their expected scientific throughput by taking into account *in situ* performance, resource availability, and workflow disruptions, has begun.

This infrastructure, and the datasets it can produce, will allow us to identify performance variability issues and workflow errors in large scale climate modeling that exist across these centers, which are likely to exist within all large compute centers. This will allow for the creation of center-independent solutions to workflow or performance issues and provide critical information to developing an optimal participation strategy for CMIP and similar modeling projects. This is expected to realize both labor and cost savings benefits across all facilities and benefit the wider climate modeling community.



## **ACKNOWLEDGMENTS**

This research was supported by a collaborative effort between the U.S. Department of Energy, Office of Science and the National Oceanic and Atmospheric Administration. This manuscript has been authored by UT-Battelle, LLC and used resources of the National Center for Computational Sciences at Oak Ridge National Laboratory, both of which are supported by the Office of Science of the U.S. Department of Energy under Contract No.DE-AC05-00OR22725. This research also used resources of the joint NOAA/DOE Partnership Project: The National Climate-Computing Research Center (NCRC). The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes.



## REFERENCES

- F. J. Anscombe. Graphs in statistical analysis. *The American Statistician*, 27(1):17–21, 1973. ISSN 00031305. URL <http://www.jstor.org/stable/2682899>.
- V. Balaji, E. Maisonnave, N. Zadeh, B. N. Lawrence, J. Biercamp, U. Fladrich, G. Aloisio, R. Benson, A. Caubel, J. Durachta, M.-A. Foujols, G. Lister, S. Mocavero, S. Underwood, and G. Wright. CPMIP: measurements of real computational performance of Earth system models in CMIP6. *Geoscientific Model Development*, 10(1):19–34, 2017. ISSN 1991-9603. doi: 10.5194/gmd-10-19-2017. URL <http://www.geosci-model-dev.net/10/19/2017/>.
- T. J. D. M. Esteban Meneses, Xiang Ni. Analyzing the interplay of failures and workload on a leadership-class supercomputer. *Cray User Group Meeting*, 2015.
- V. Eyring, S. Bony, G. A. Meehl, C. A. Senior, B. Stevens, R. J. Stouffer, and K. E. Taylor. Overview of the Coupled Model Intercomparison Project Phase 6 (CMIP6) experimental design and organization. *Geoscientific Model Development*, 9(5):1937–1958, 2016. ISSN 19919603. doi: 10.5194/gmd-9-1937-2016.
- IPCC, Core Writing Team, R. K. Pachauri, and L. A. M. (eds.). Climate change 2014: Synthesis report. contribution of working groups i, ii and iii to the fifth assessment report of the intergovernmental panel on climate change, 2014.
- K. E. Taylor, R. J. Stouffer, and G. A. Meehl. An overview of CMIP5 and the experiment design. *Bulletin of the American Meteorological Society*, 93(4):485–498, 2012. ISSN 00030007. doi: 10.1175/BAMS-D-11-00094.1.
- D. Tiwari, S. Gupta, J. Rogers, D. Maxwell, P. Rech, S. Vazhkudai, D. Oliveira, D. Londo, N. DeBardeleben, P. Navaux, L. Carro, and A. Bland. Understanding gpu errors on large-scale hpc systems and the implications for system design and operation. In *2015 IEEE 21st International Symposium on High Performance Computer Architecture (HPCA)*, pages 331–342, Feb 2015. doi: 10.1109/HPCA.2015.7056044.
- U.S. Global Change Research Program. Climate change impacts in the United States : U.S. national climate assessment, 2014. URL <http://purl.fdlp.gov/GPO/gpo48682>.
- D. N. Williams, V. Balaji, L. Cinquini, S. Denvil, D. Duffy, B. Evans, R. Ferraro, R. Hansen, M. Lautenschlager, and C. Trenham. A global repository for planet-sized experiments and observations. *Bulletin of the American Meteorological Society*, 97(5):755–766, 2016a. ISSN 00030007. doi: 10.1175/BAMS-D-15-00132.1.
- D. N. Williams, M. Lautenschlager, S. Denvil, L. Cinquini, R. Ferraro, D. Duffy, C. DeLuca, V. Balaji, B. Evans, and C. Trenham. 6th Annual Earth System Grid Federation Face-to-Face Conference Report. Technical report, Lawrence Livermore National Laboratory, 2016b.