CPMIP performance metrics for CMIP6

For a complete description of the metrics, please check Balaji et al. 2017

Metric	Description
Model	Model name and version
CMIP6 Config	Descriptive tag of CMIP6 experiment
Resolution	Grid resolution for each component of the model given by the number of grid points (NXxNYxNZ)
Platform	Machine used
Institution	Institution launching the experiment
Complexity	It is measured as the number of prognostic variables per component of the model. If not available directly from the model configuration or code, it can be computed by dividing the size of the restart file (containing the complete state) per component, measured in words (e.g., 8 bytes for double precision) divided by the number of grid points (NXxNYxNZ)
SYPD	Simulated years per day for the model in a 24 h period. Suggested 1 year chunk/leg size as common standard to calculate it
ASYPD	Actual SYPD: This number should be lower than SYPD due to interruptions, queue wait time, data transfer or issues with the model workflow. This is collected by measuring the time between first submission and the date of arrival of the last history file on the storage file system. Communicate which history file will be used for calculations
CHSY	Core hours per simulated year. This is measured as the product of the model runtime for 1 SY and the number of cores allocated.
Parallelization	Total number of cores allocated for the run
JPSY	JPSY is the energy cost of a simulation, measured in Joules per simulated year. Communicate which methods are available for each machine
Coupling Cost	Overhead caused by coupling. Measured as the normalized difference between the time-processor integral for the whole model vs. the sum of individual concurrent components, including runtime and parallelization for the whole model and each component
Memory bloat	It is the ratio of the actual memory size to the ideal memory size. The measured runtime memory usage on the system (often called “resident set size”, RSS, often published in the job epilog) is divided between instructions and data. The portion of memory devoted to instructions is measured by taking the size of the executable files produced during compilation
Data output cost	For synchronous I/O where the computational PEs also perform I/O, it requires a separate “No I/O” run to collect CHSY. For models using asynchronous I/O it may be possible to measure it by looking at the allocation fraction of the I/O server. Suggested “No IO (writing process)” run for all cases as common standard
Data intensity	It is the measure of data produced per compute hour, in GB/CH. This is measured as the quotient of data produced per SY, examining the

	output directories and divided by CHSY

HPC-TF Performance metrics for CMIP6

1) Global numbers for the whole CMIP6 (see the "Global Resource HPC-TF Metrics" tab)

Metric	Description
Useful Simulated Years	Number of years simulated for CMIP6 for the useful runs only, i.e. runs that produced useful scientific data used in analysis and/or published on ESGF
Total Simulated Years (**optional**)	Total number of years simulated for the whole CMIP6 experiment, including tuning and runs that were finally discarded for any reason
Useful Data Produced Size	Total data output volume produced for the whole CMIP6 for the useful runs only, i.e. runs that produced useful scientific data used in analysis and/or published on ESGF
Total Data Produced Size (**optional**)	Total data output volume produced for the whole CMIP6 including tuning and runs that were finally discarded for any reason
Useful core hours	Number of core hours used for the whole CMIP6 for the useful runs only, i.e. runs that produced useful scientific data used in analysis and/or published on ESGF
Total core hours (**optional**)	Total number of core hours used for the whole CMIP6, including tuning and runs that were finally discarded for any reason
Total **(NEW!)** Person/Months	Total number of human resources used for the setting up the CMIP6 experiments for each specific platform, running them until the end of the simulation (including time spent to rerun due to errors during the execution) and managing the data produced by these experiment until publication (this excludes the time spent to assemble and tune the model used)https://docs.google.com/document/d/18wNcx3TmWbuk6eK8m-qqGfSUhw3IoHGCgA_gEFXbIAk/edit

2) Numbers for the CMIP6 experiments run with a specific model (see the bottom of the "Performance-modelA" tab).

Metric	Description
Useful Simulated Years	Number of years simulated for CMIP6 experiments run with this model for the useful runs only, i.e. runs that produced useful scientific data used in analysis and/or published on ESGF
Useful Data Produced Size	Total data output volume produced for CMIP6 experiments run with this model for the useful runs only, i.e. runs that produced useful

	scientific data used in analysis and/or published on ESGF
Useful core hours	Number of core hours used for CMIP6 experiments run with this model for the useful runs only, i.e. runs that produced useful scientific data used in analysis and/or published on ESGF

3) Number for the platform (see the bottom of the Machine tab)

Pflops-hour peak and linpack is the equivalent Peak and Linpack PFlop-hours. For a given computer used look up the total number of cores (N), the total peak performance (RPeak) and the Linpack Performance (Rmax) in the Top 500 list (http://www.top500.org/list). Equivalent Peak and Linpack performance are then given by C*Rpeak/N and C*Rmax/N respectively where C is the number of core hours given (total core hours) by the experiment.