



Infrastructure Group

Progress since Fall meeting, future issues

Rob Jacob and Chengzhu Zhang
(on behalf of the entire Infrastructure Group)

2021 E3SM All-hands

June 8, 2021

Infrastructure Group responsibilities

- Develop, maintain and support software that is needed for E3SM but is **not** part of the main prognostic models. Configure, build, test, diagnostics, analysis, workflow, driver/coupler
- Manage data sets
- Define, document, manage the process and procedures used in software development within the E3SM Project.

Everything we do should help make the model development, simulation and analysis happen.

“big picture” main model development progress

- Still averaging 1 PR merged per calendar day
- V2 WC developments mostly done (said that in October)
 - Added more metadata to EAM,ELM, MOSART, MPAS output
 - Standard_name for popular CMIP variables.
 - Corrected units (where easy)
 - Several tags made to capture v2 development. Examples:

v2alpha -- Sep 24, 2019 –

all major v2 PRs merged (we thought)

v2alpha2 – April 9, 2020 –

revert to CLUBBv1

v2alpha3 – July 2, 2020 –

back to CLUBBv2

v2alpha5 – Nov 23, 2020 –

the “5_59 fallback”

v2beta – Dec 18, 2020 -

theta-l is default

v2beta3 – March 5, 2021 –

shallow shortwave fixer.

v2beta4 – Apr 9, 2021 –


Fix melt potential calculation

v2rc1 – Apr 23, 2021 –

Change heat exchange from frazil

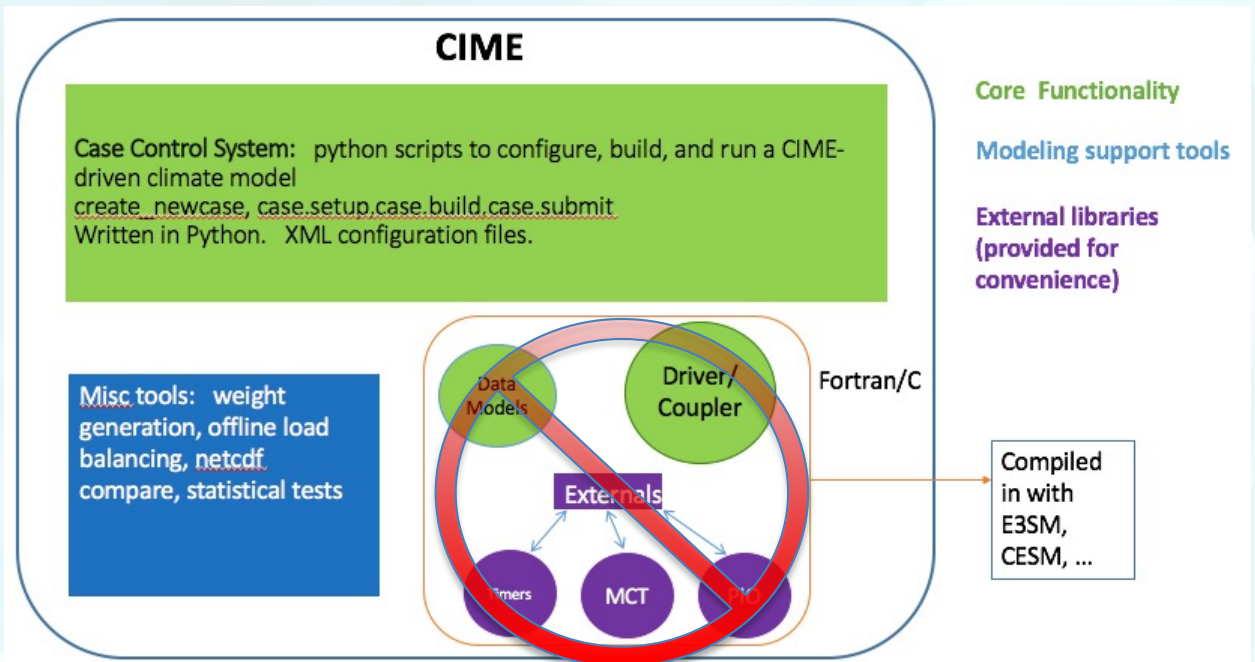
v2rc3 – May 28, 2021 --

Fix heat and mass conservation in frazil




“big picture” main model development progress

- Changes to submodules/code structure.
- CIME: We now have our own copies of coupler, data models and share code
 - E3SM/cime/src/driver/mct -> E3SM/driver-mct
 - E3SM/cime/src/components/data_comps_mct -> /E3SM/components/data_comps
 - E3SM/cime/src/share -> E3SM/share
 - Collection of Fortran utilities needed by 2 or more models (usually ELM, EAM, data models)
Compiles as a separate library before any models.
 - NOTE: CIME paths are still there (for now). Ignore them.
- MPAS: MPAS framework is no longer a submodule
 - E3SM/components/mpas-source -> E3SM/components/mpas-framework



“big picture” main model plans

- Before v2.0 tag, update README, other docs.
- Remaining BGC and Cryo v2 PRs
 - Additional submodules: GCAM, MARBLE
- Start accepting answer-changing v2.1, v3, v4 developments to existing components



Recent and upcoming changes to testing

- New Circle-CI testing (a free service from Github): every PR is automatically built (using gnu) in a fully coupled case. Replaces Travis-CI.
- Removed sandiatoss3 from testing. Replaced by Chrysalis for nightly testing of e3sm_integration on next/master with baselines.
- Coming soon: Testing on OLCF Ascent (small version of Summit) via ECP.
- Todo: Add maint-1.2 test suite
- Todo: Convert all developer, integration, prod testing to use v2 configurations:
 - Use new MPAS-seaice thermodynamic capability to replace CICE in F-cases
 - More tri-grid configs?
 - If tri-grid is future, bi-grid will be on maint-2.0
 - Low res tri-grid waiting for r2 to work in land model


Chrysalis update

- After a rocky start, Chrysalis is now our fastest machine!
 - 40 years/day for v2 low res
 - 10 years/day for v2 RRM
- Config
 - 512 nodes each with
 - 2 32-core AMD Epyc CPUs 2.4 GHz
 - 256 GB DDR4 RAM
 - 3 PB disk
- After installation, found/fixed several issues
 - Turned off resource-hogging GPFS monitor
 - Switched to network-specific OpenMPI after finding/fixing bug.
 - Updated firmware, changed topology, turned on SMT



MOAB-based coupler and related tools

- MOAB coupler can now...
 - Declare MOAB instances and meshes on
 - EAM: GLL (no longer needed) and physgrid
 - ELM: following atm or river
 - MPAS-ocean, MPAS-seaice
 - MOSART
 - Create a2o, a2i, o2a, i2a, a2l (for trigrd) maps **online**.
 - Read pre-computed maps (r2o) in parallel.
 - Create load-balanced and geographically compatible decompositions in coupler with Zoltan (more later).
 - Still TODO: full coupler data flow, merging, fluxes
- Mbtempest updated on all platforms.
 - New method for finding matrix weights in TempestRemap that works for degenerate cases will change answers for all maps.
 - Fixed a bug for SE -> FV RLL.
 - Can find HYDRO1K to ne1024 map in 10 minutes.
 - Being added to ncremap.



ELM decomposition from ELM MOAB instance

Data Archiving and Publication (Progress over Oct 2020 - May 2021)

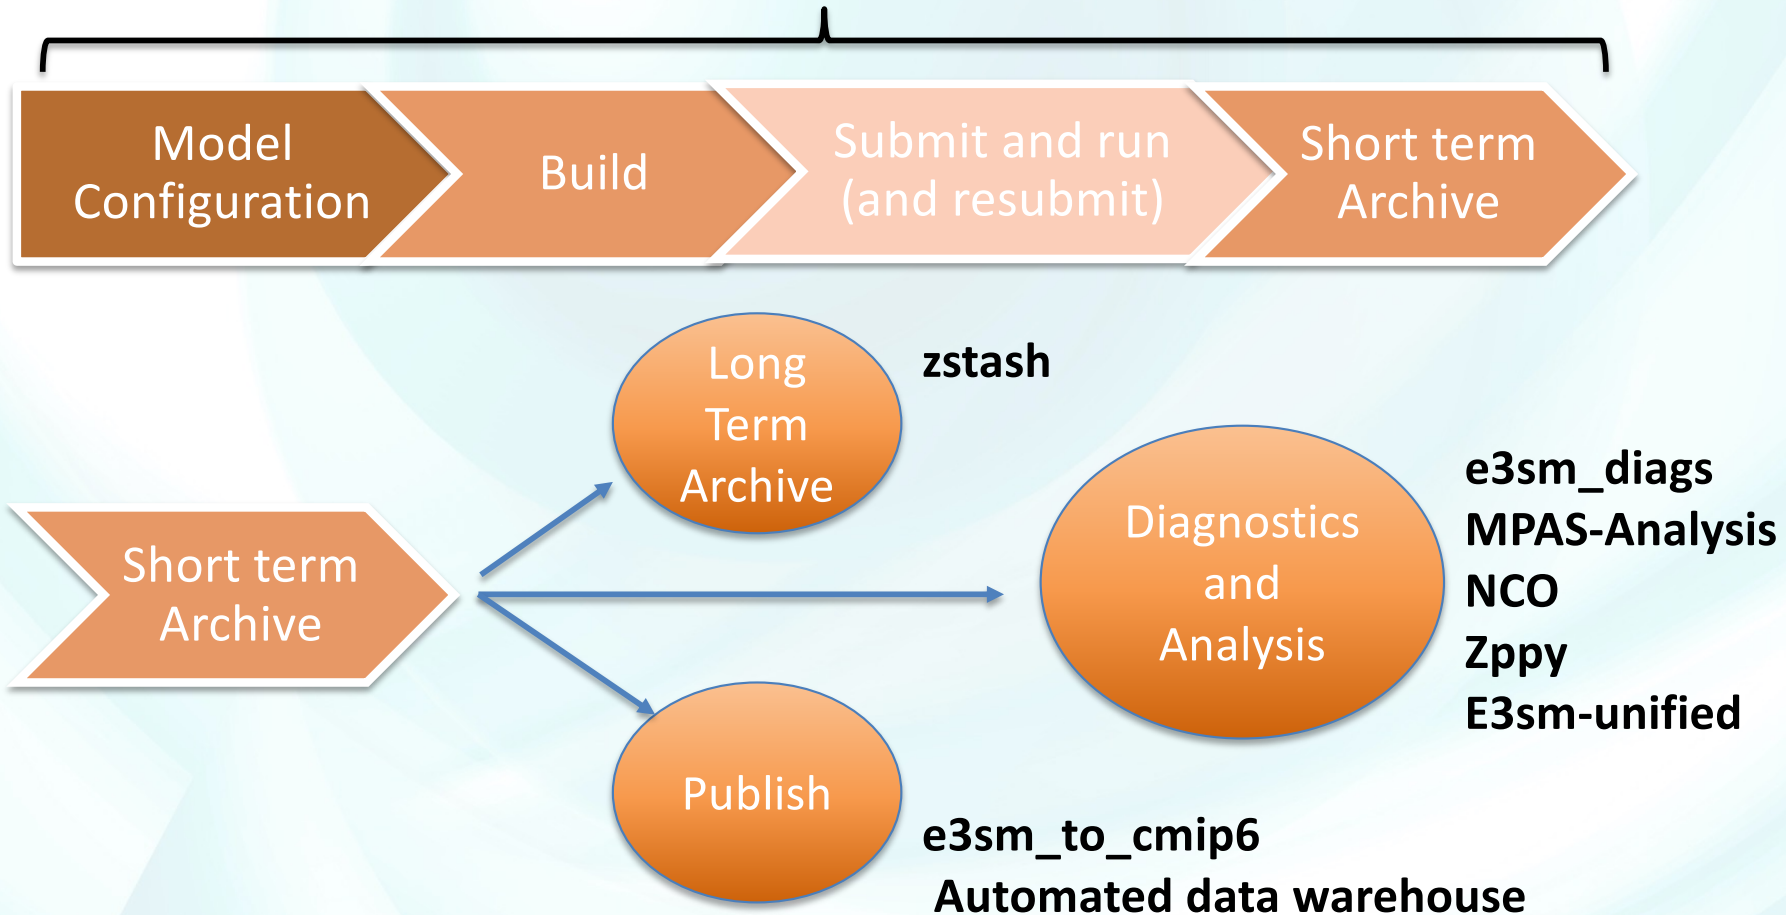
	E3SM	CMIP6
Datasets	563 (195 this year)	5,542 (3,734 this year)
Total Files	694,229 (183,758 this year)	66,379 (26,829 this year)
Data Size	596TB (236TB this year)	12.3TB (5.8TB this year)

Data volume doubled this year.
Data holding ~ 50:1

- The LLNL archive now contains raw data for every major simulation.
- Recent publication: Cryosphere v1, 21-year HRv1 high-volume, BGC SSP585, COVID MIP.
- Validation (including file integrity, temporal continuity, dataset completeness testing), and data publication are run automatically. Tools are currently in development to handle large ensembles.
- Assisting CMIP publication for external projects: **CFMIP** at LLNL, **PAMIP** at UCI
- More output will be put into CMIP regridded time-series on standard grids and submit to CMIP: e.g. Sub-monthly variables, additional monthly variables.
- New ESGF node at ANL is ready to publish datasets.

E3SM Tools for Every Phase of Model Development

CIME Case Control System, run_e3sm



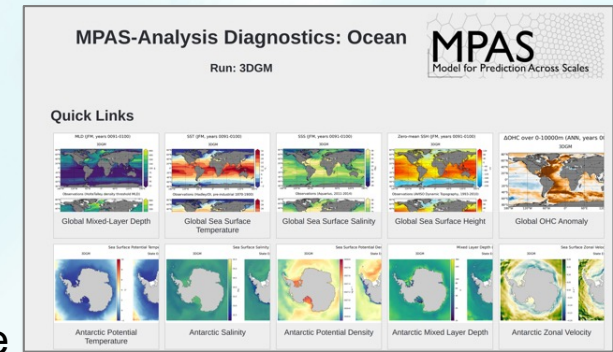
Case Control System, run script updates

- New, shorter version of run_e3sm script available.
- create_newcase: new optional argument “–case-group” to indicate case is part of larger group of simulations.
- Provenance collection:
 - New “case hash” created for each create_newcase invocation.
 - More info on status of git repository collected (submodule status, diffs).
 - “replay.sh” created as you create/modify case with CIME commands to reproduce.
 - Capture run_script if used.
- Performance tracking
 - Include detailed I/O timers from SCORPIO
 - More memory use reporting

Analysis and Diagnostics Tools

Tools included in standard post-processing workflow:

- **NCO**: Regidding, Climatology Generation and Timeseries Extraction, support **all E3SM components**
- **E3SM-Diags**: Diagnostics package for **Atmosphere**, extending the support to **Land and River**
- **MPAS-Analysis**: Diagnostics package for **Ocean and Sea-ice**
- **Zstash**: Long term archiving tool
- **E3sm_to_cmip**: convert output to **CMIP** format
- **E3SM-Unified**: Conda package for distributing all the above and more
- **Zppy***: Post-processing tool chain to automate most commonly analysis



More tools used in E3SM evaluation /connected to E3SM:



ILAMB



CMEC



AMWG
CVDP

Analysis and Diagnostics Tools: Updates

Upcoming July release of E3SM-Unified



- E3SM-Unified: Conda environment (and system-specific tools) for analysis
Available on all supported platforms (**anvil/chrysalis, compy, cori, theta, summit**)

Updates of **main** E3SM packages:

- **e3sm_diags**: TC analysis, Annual cycle zonal mean (ARM Diags from last release)
- **mpas-analysis**: Performance and configurability improvements; better on-the-fly mapping file and region mask generation (especially for high-res meshes)
- **nco**: Support ncclimo monthly mth_srt/end (also for high frequency); S1D PFT/MEC analysis; Mask support end-to-end; fixes for machine and component specific issues...
- **zstash**: Support tar files-based checksums
- **e3sm_to_cmip**: Support high frequency variable (3hourly, daily..) converter.

New packages being included:

- **zppy** 1.0.0: Automatically run c;imo, time-series, e3sm_diags, mpas-analysis, global time series plots.
- Tempest-extremes: automated pointwise feature tracking on unstructured grids

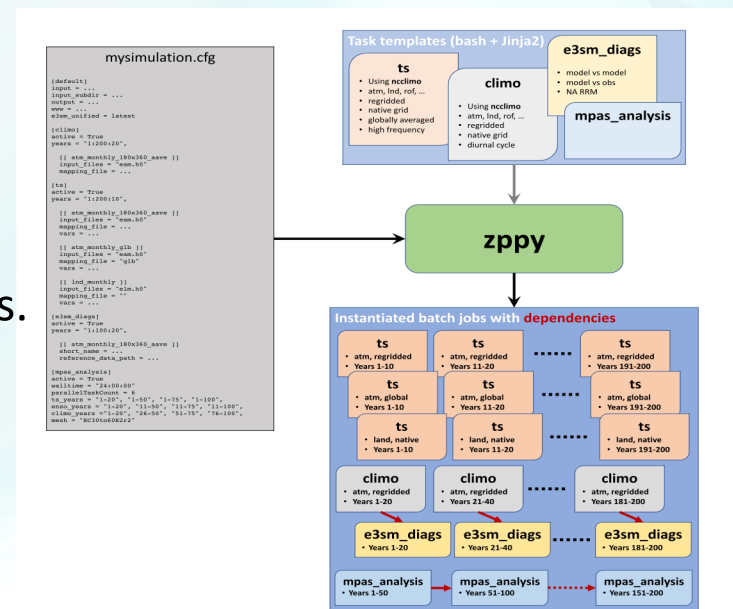
- Current status:

- Release candidates under testing within developers and will be open to all of E3SM in June

Zppy: Automated post-processing workflow

A new Joint WaterCycle and Infrastructure effort

- Goal: Speed up the post-processing of E3SM simulations by automating commonly performed tasks. Serve as a standard tool for simulation campaigns, promoting consistent post-processing, results curation and simulation documentation.
- Built around tasks:
 - Generating/regridding **time series/climatology** with **NCO**,
 - Generating analysis figures with **E3SM Diags** or **MPAS-Analysis**.
 - Generating global time series.
 - More to be included.
- A framework to simplify tools use, esp. for long simulations.
- A single configurable file (“mysimulation.cfg”) (see fig).
- Zppy parses the configuration file, handles dependencies and generates/submit batch jobs for execution by SLURM.



Analysis and Diagnostics Tools: Plan

- Keep integrating analysis and priority metrics from science groups
- Support new tool development addressing requests from science groups (i.g. zppy)
- Set up and document **best practice** running the model and post-processing for v2
- Jointly working with Performance team to interface zppy and PACE for simulation documentation automation and provenance.
- Leverage DOE-funded external projects: ARM-diags, iLAMB, PMP, CMEC. to extend current capability.
- Jointly working with PCMDI/RGMA scientists to replace CDAT by extending xarray (**xcdat**).
- Performance enhancement: multi-node task parallelism
- Increased support for external projects on post-processing and analysis tools with documentation/slack channels/tutorials

Encouraging V2 simulation campaign standards

- V1 campaigns had too much variability in details of how they were run.
 - Huge problems in post-processing and ESGF publishing from inconsistent archiving.
 - Some errors were un-recoverable.
 - Simulation campaigns run by other groups have similar, easily avoidable, problems.
- To avoid this, all v2 campaigns should follow the same process, procedures for doing runs.
 - All case names formed consistently (so file names will be consistent)
 - Recover from errors and restart the same way
 - Archive results with the same process.

E.g. Start with group leads filling out a campaign planning table:

Fill out after run starts



Case	Case-id	--case-group	Case runner	Machine	Total years	Code version	Date started	casedir

Then each case runner follows the same procedure to actually do the runs.

Common procedure prototyped by WaterCycle

- Use **run_e3sm script**
 - Configure, checkout, compile, basic regression tests, production simulation.
- Store on disk with **Short term archive** (using CIME)
- Post-process with **zppy**
 - /Time series/climatologies/Regridding (**ncclimo**), **E3SM Diags**, **MPAS-Analysis**. More to come; automate as much as possible (cmorization with **e3sm_to_cmip?**)
- Long term archive with **zstash**
- Simulation page documentation (by hand, will later be mostly automated)

For details, see [Running E3SM step-by-step guide](#) on Confluence (thanks to Ryan).

Standardize post-processing tools and centralize datasets

Additional challenges

- Disk management:
 - Copy full disk issue: Scrubber installation. Designated space on NERSC HPSS/disk recommended for intermediate simulations.
 - Best to use NERSC cfs disk space for routinely analyzed data
 - Chrysalis: leave data there. Can publish and archive locally.
- Additional data issues:
 - Where to use data compression?
 - output more/higher frequency data you normally wouldn't?
 - Losslessly compress everything?
 - Reduce default output on EAM h0 file (with large ensemble in mind).
 - Define/implement metadata requirements for **input** data and **analysis** data.
- Tool use:
 - Train/encourage/require all production run managers to use run-script/zstash
 - Get all ecosystem projects to use zstash.
 - Automatic addition of provenance as runs are processed. Auto-creation of run-summary pages.