

CESM Migration (Porting) to Yellowstone – A Retrospective

- 1) CESM Science (Complexity) and Yellowstone
- 2) CESM Porting Process
- 3) Pre-Yellowstone (Feb 2012 → Oct 2012)
- 4) Early Access (Oct 2012 → Feb 2013)
- 5) Advanced Scientific Discovery (Nov 2012 → Feb 2013)

Mariana Vertenstein, CESM Software Engineering Group Head

Thanks to CESM Software Engineering Group and Computer Information System Lab



CESM System Is Growing in Complexity

- **Components**

- 6 component models: atmosphere, land, sea-ice, ocean, **land-runoff**, and land-ice plus coupler (**2 new components coming soon**)
- Each component can be active or prescribed
- Multiple physics options in each component – Chemistry, Biogeochemistry and High-altitude capability
- Various component versions and backwards compatibility (CAM4, CAM5, CLM4.0, **CLM4.5**, CISM1, **CISM2**)
- Multi-instance capability for DART data assimilation and SPCAM

- **Configurations (Component Sets)**

- ~140 **supported** model configurations

- **Resolutions**

- ~50 **supported** resolutions across components
- **New refined CAM-SE and MPAS atm, land and ocean grids**



The Porting Process

- **Verify functionality**
 - Run hundreds of tests (e.g. restart) across supported model configurations and resolutions
- **Validate climate**
 - Target several popular model configurations
 - *Completely new validation mechanism created in collaboration with CISL that provides a success/fail validation metric for first time (now viable with Yellowstone capacity)*
- **Load balance and performance tune**
 - Maximizing throughput and minimizing idle cycles depends on resolution, configuration and scalability

Pre-Yellowstone : University of Colorado's Janus Computer

CESM verified, validated and performance tuned

Positive

- resolved many problems before Yellowstone access (e.g. new intel compiler settings, high-resolution memory and performance issues)

Negative

- long queue wait times made debugging difficult
- difficult to obtain accurate computational allocation costs for Yellowstone



Yellowstone : Early Access

- Machine performance allowed quick verification tests
- Rapid turnaround sped up debugging
- Example: new validation test
 - 100 1-year runs of standard CESM in 13 hours, caught a subtle issue with a compiler function



3 Releases On Feb. 1

- **cesm1.1.1** (support for new model features)
 - CAM-SE/CAM5 development code base
 - CAM5/CLM4.0 targeted defaults
 - Targeting “scientific support” for FV 1°, FV 2° 1850, 20th Century and RCP CMIP5 simulations
 - Out of box support for yellowstone, titan
- **cesm1.0.5** (support for CMIP5 science)
 - Out of box support for yellowstone, titan
 - Scientific support for all CAM4 FV CMIP5 simulations
- **ccsm3** (university community demand)

Yellowstone: ASD High Resolution Simulations (Justin Small PI)

– Resolutions

- 25km CAM5/CLM4.0 (spectral element grid)
- 10km POP2 (62 level), CICE (tripole grid)
- 1/2° RTM (lat/lon grid)
- Resolutions allow for resolving ocean eddies and atmospheric hurricanes


– Performance

- ~2 model- years/day on 16,000 processors
- Previous LLNL atlas run of 25km CAM4/ 10KM POP2 (42 levels) had only 1/2 model-years/day on 4000 processors


– Target: multi-decade simulations (grand challenge goal)



3 months of hourly precipitation



3 months of hourly water vapor (TMQ)



1 year of daily SST, sea-level height and sea-ice coverage

