Performance Analytics for Computational Experiments

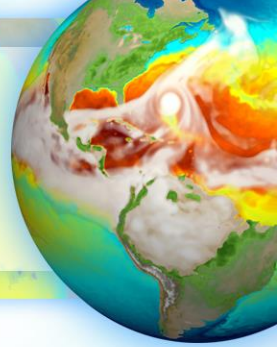

Sarat Sreepathi

Oak Ridge National Laboratory

E3SM All-Hands Webinar

June 4th, 2020

In collaboration with

Youngsung Kim, ORNL

Past Students:

Zachary Mitchell,
Pellissippi State Community College

Gaurab KC,
University of Tennessee Knoxville

Goals

Provide executive summary of E3SM performance to stakeholders

Facilitate:

- Interactive analyses and deep-dives into experiments and application sub-regions, as desired,
- Tracking performance benchmarks and simulation campaigns of interest,
- Facilitating performance research on load balancing and process layouts,
- Identification of bottlenecks to inform targeted optimization efforts.

PACE Architecture

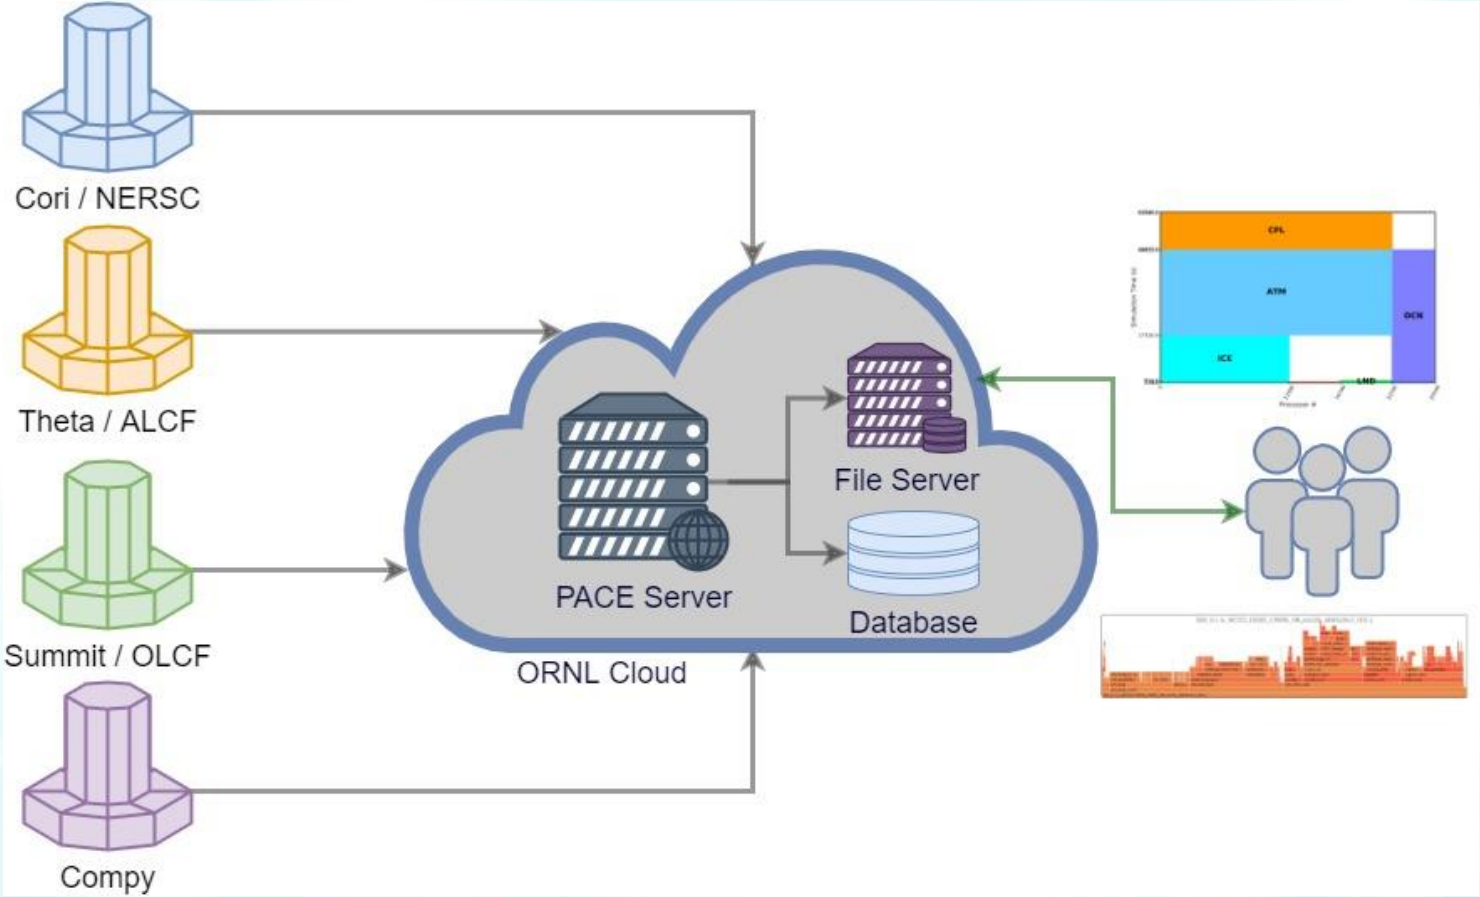

Technology Stack

- Infrastructure
  - ORNL Cloud (CADES)
  - OpenStack VM
- Nginx Web Server + Reverse Proxy
- Python-Flask middleware
  - Application Server
  - Process model inputs/timings
- Minio File Server
  - Object based storage for raw data
- MariaDB database
  - Structured and semi-structured data
  - Flexible Schema
- JavaScript
  - Frontend and visualization

Visualization Pipeline

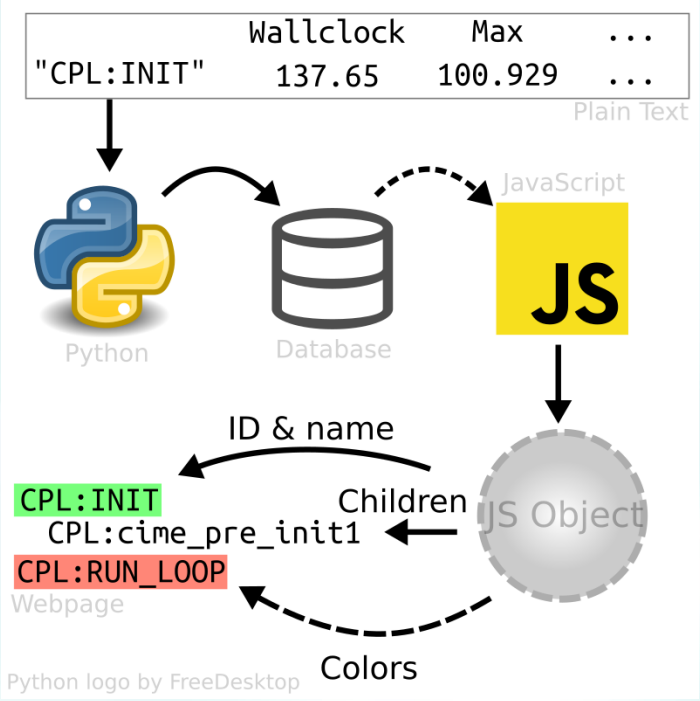

Last but definitely not least:
Cybersecurity compliance at a DOE lab

E3SM Performance Data

- Lightweight performance profiling by default
  - Utilizes General Purpose Timing Library (GPTL) timers
  - Mark start/stop at defined application phases
  - Aggregate statistics for parallel processes
  - Collect computation, communication and I/O performance data
- Performance Archiving
  - Archive performance data in project wide locations
  - Representative provenance data for context
  - System state and various logs
  - Hierarchical directory structure
    - user->case->job
  - Enabled on Anvil, Compy, ALCF, OLCF and NERSC.

New Capabilities

- Parse Model Inputs
  - XMLs: 187,682
  - Namelists: 440,374
  - RCs: 24,187
  - Raw data (.zip): 24,468
  - Total experiments: 24,468
- Automated nightly uploads
  - Using Jenkins
  - Deployed on Compy and Cori
  - Anvil, Theta – on the way

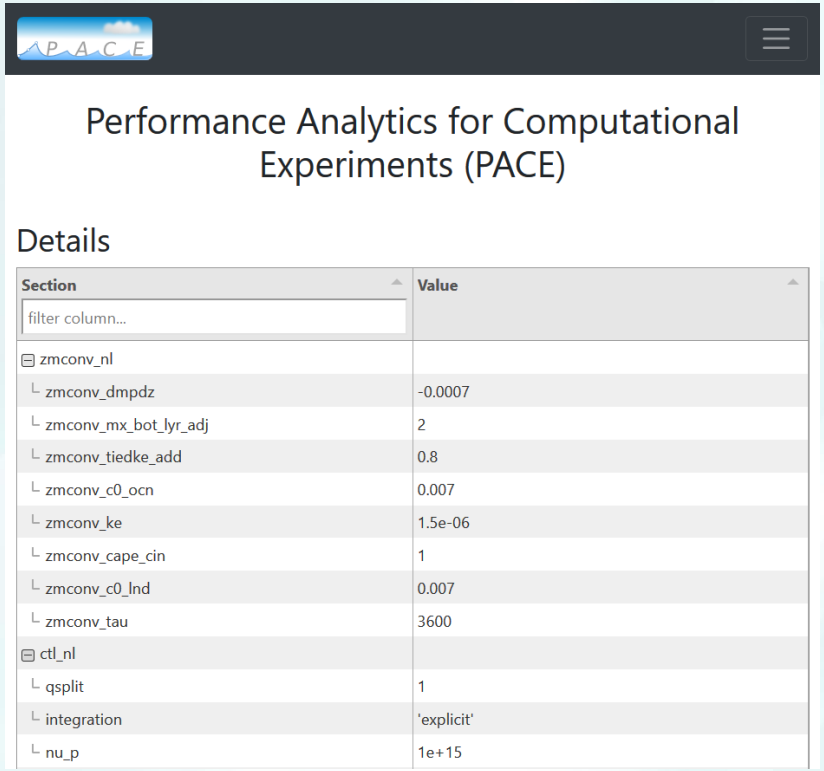

View and filter namelists

PACE Statistics

- 151 users
- 10 platforms
- Performance Benchmarks
  - High-res Atmosphere
  - High-res Ocean
- Simulation Campaigns
  - DECK v1
  - High-res Water Cycle
  - MMF Early Science
  - BGC: BCRC, BCRD, BDRC, BDRD

- 24,468 experiments
- 652,243 model input files
- 59,560 timing files

Aggregate statistics and reports can inform:

- INCITE and other compute allocation proposals
- Computing procurements

Tree Graph

Summarize time taken by model components

Recursively explore time taken by model sub-regions

Flame Graph

High-level overview of a parallel process execution time

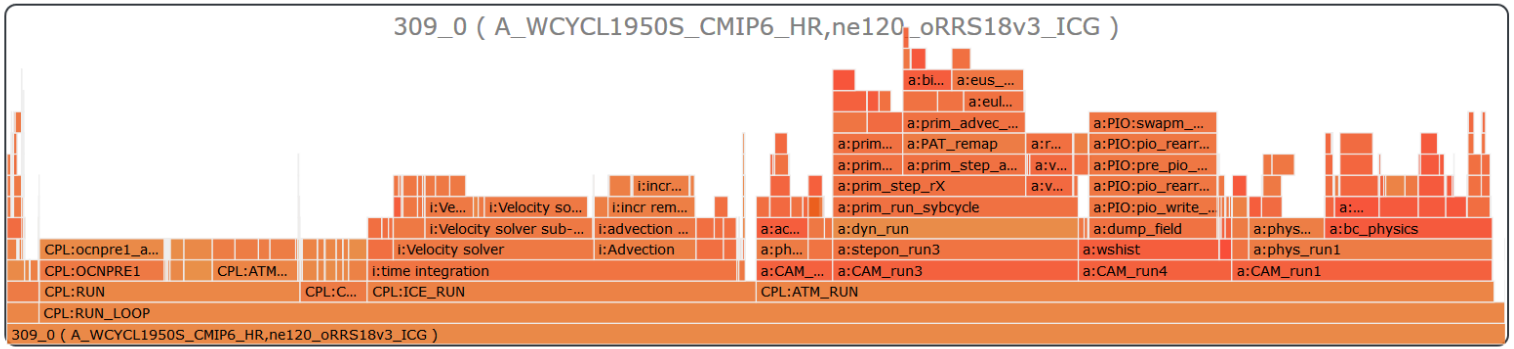

Performance Research Directions

Load Balancing

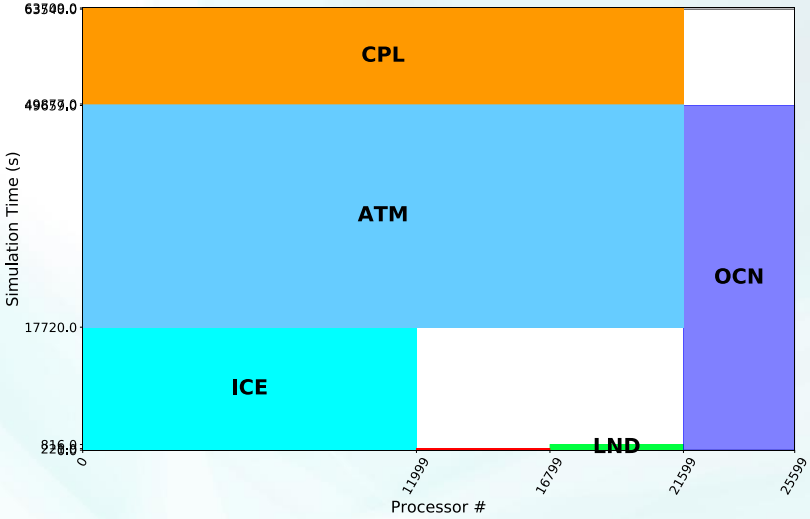

MPI Task Mapping

Optimization

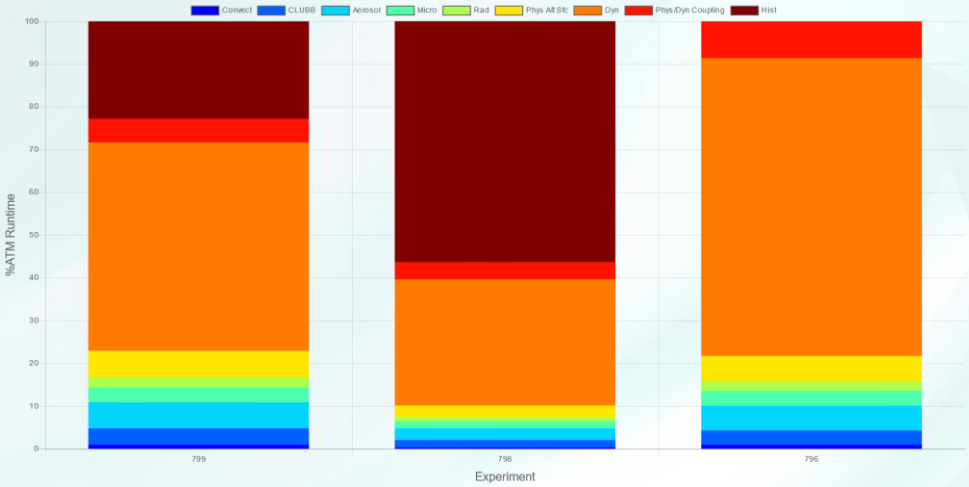

Atmosphere model time distribution

Acknowledgments

Links

- BER Highlight
- PACE Portal
- Reference Page on E3SM.org
- Video - Web Portal Features
- Video - How to Upload Data

https://pace.ornl.gov

Thanks!

Contact:
sarat@ornl.gov