

Checkpoint/Restart Options on Cori



Acknowledgements:

Rebecca Hartman-Baker, Tiffany Connors, Steve Leak, and Chris Samuel at NERSC;
Gene Cooperman, Rohan Garg, Twinkle Jain, and Prashant Chouhan at Northeastern University;
Harsh Khetawat at Carolina State University,

Zhengji Zhao
NERSC User Engagement Group
NERSC User Group Monthly Webinar
September 24, 2020

What is checkpointing?

- Saves the state of a running process to a file
- The process can later be restarted from the file, continuing from where it left off from any computer
- Checkpoint approaches
 - Internal checkpoint/restart in applications
 - External transparent checkpointing

Why Use C/R?

- Enable long running jobs
- Better queue turnaround
- Prepare for system failures, including PSPS
- Prepare for future
 - Increasing real-time workloads in the N10 time frame, which may require preempt partial or full system running jobs
- Enable a 75% charging discount with flex QOS on Cori KNL
- Variable-time job scripts make it easier to use C/R

Variable-Time Job Scripts

- Add a few `sbatch directives` and `bash functions` in your job scripts
- Can be used with applications with internal C/R support or can be checkpointed with external C/R tools
- Allow longer jobs to run in multiple shorter jobs
 - User specifies a minimum and maximum time for the job
 - System then finds the best time slots for the job + automates job resubmissions
 - User just submits one job script
- Allow improved queue turnaround by better utilizing backfill opportunities

Available C/R Options on Cori

- For applications with internal C/R support
 - Variable-time jobs + flex QOS
- For applications **without** internal C/R support
 - DMTCP (Distributed MultiThreaded CheckPointing) - a transparent C/R tool
 - + Variable-time jobs + flex QOS
 - Works with **serial and threaded applications only** on Cori

To run

```
#SBATCH ...  
  
#for c/r with dmtcp  
module load dmtcp nersc_cr  
  
#checkpointing once every hour  
start_coordinator -i 3600  
  
#run job under dmtcp control  
dmtcp_launch ./a.out
```

To restart

```
#SBATCH ...  
  
module load dmtcp nersc_cr  
  
#checkpointing once every hour  
start_coordinator -i 3600  
  
#restart the job from dmtcp  
checkpoint files  
./dmtcp_restart_script.sh
```

What's Next for C/R at NERSC

- MANA for MPI and MPI + OpenMP Applications
 - MANA: MPI-Agnostic Network-Agnostic Transparent Checkpointing tool, implemented in DMTCP
 - Works with VASP, Gromacs and HPCG currently
 - Will be available to users for testing soon
- C/R user trainings
- Community building

First International Symposium on Checkpointing for Supercomputing (SuperCheck21)

- Date: **Feb 4-5, 2021** (8:00 am - 12:45 pm PST)
- Online workshop featuring the latest work in checkpoint/restart research, tools development, and production use
- Participation highly encouraged!
 - We especially encourage abstract submissions on adopting C/R tools in production workloads.
- More details at <https://ckpt-symposium.lbl.gov>

References

- MANA:
 - R. Garg, G. Price, and G. Cooperman, “[MANA for MPI: MPI-Agnostic Network-Agnostic Transparent Checkpointing](#)”, Proceedings of the 28th International Symposium on High-Performance Parallel and Distributed Computing, June 2019, Pages 49-60.
 - [Code](#)
- DMTCP:
 - [Website](#)
 - [Code](#)
- NERSC Documentation
 - [NERSC website checkpoint/restart \(DMTCP\)](#)
 - [Variable-time job scripts](#)
- Training Materials:
 - [Variable-time job scripts for VASP users](#)
 - [Variable-time job scripts](#)
 - [DMTCP for users running serial and threaded applications](#)

The logo for the National Energy Research Scientific Computing Center (NERSC). It features the acronym "NERSC" in a bold, white, sans-serif font, centered within a dark blue rectangular background. The background has a subtle radial light effect emanating from behind the text.

Thank you!