

CS554 Project Ideas

FusionFS:CKPT - Efficient Checkpointing with Distributed File Systems

Overview

Checkpointing is the de facto mechanism to achieve fault tolerance on large-scale distributed systems. This project is comprised of three parts: (1) a thorough literature review of state-of-the-art techniques in checkpointing; (2) design and implement a checkpointing mechanism, and integrate it to the FusionFS distributed file system [1]; (3) evaluate the checkpointing system with benchmarks (e.g. BLCR[2]).

Because FusionFS is optimized for write-intensive workloads, and because checkpointing is highly write-intensive, we expect the checkpointing performance on FusionFS is significantly higher than conventional approaches where checkpoint data are stored remotely on NAS such as GPFS (for extreme-scale supercomputers) and S3FS (for Amazon EC2 cloud). The testbed will be Amazon EC2 cloud platform, where you will run experiments up to 128 instances and compare the checkpointing performance of your approach over FusionFS to BLCR. The outcome of this project will likely lead to a conference or workshop (short) paper.

Relevant Systems and Reading Material

Please read the following papers (and their references) before submitting your proposal:

- [1] Dongfang Zhao, Zhao Zhang, Xiaobing Zhou, Tonglin Li, Ke Wang, Dries Kimpe, Philip Carns, Robert Ross, and Ioan Raicu. "FusionFS: Towards Supporting Data-Intensive Scientific Applications on Extreme-Scale High-Performance Computing Systems", IEEE International Conference on Big Data, 2014.

Available online: <http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=7004214>

- [2] Paul H. Hargrove and Jason C. Duell. Berkeley Lab Checkpoint/Restart (BLCR) for Linux Clusters, *Technical Report, Lawrence Berkeley National Laboratory*, 2006.

Available Online: http://crd.lbl.gov/assets/pubs_presos/CDS/FTG/Papers/2006/LBNL-60520.pdf

Preferred/Required Skills

- Principles: operating system, distributed systems, computer network
- Programming: Shell Script, Perl/Python, C, C++, PThread, sockets, FUSE
- Operating systems: Linux

Project Mentor

- Dongfang Zhao
- Email: dzhao8@iit.edu