[Slurm User Group Meeting 2022]

Burst Buffer Lua Plugin For Lustre

Kota Tsuyuzaki / Rikimaru Honjo / Yusuke Kaneko / Kohei Tahara

NTT Computer and Data Science Laboratory / NTT TechnoCross Corporation

NTT AI Research and Slurm

- 13 Individual Laboratories and 2300 employees
- Slurm Workload Manager for AI Training Clusters
- >5 International Papers/Workshops on AI Research Field
  - Reported in GTC 2021 “How to Improve Your Data Center GPU Utilization Efficiency”
  - https://www.nvidia.com/en-us/on-demand/session/gtcspring21-s31410/?playlistId=playList-8821a8c4-205c-488c-95f0-55826822e1a8

The Scaling Problem

- AI Training DataSet and Cluster Users Increasing
- Storage Scalability
  - Peta Bytes Scaling Distributed File Storage
- Scalable Cluster Architecture
  - Data Transfer between Storage and Compute Nodes

generated by stable diffusion

Lustre Filesystem

- Open Source Parallel File System
  - POSIX-compatible
- Used in TOP500 Clusters
- Scale both at Throughput and at Volume
  - PB Scale Volume in Production
- Information
  - https://www.lustre.org/
  - https://wiki.lustre.org/Main_Page
  - https://www.opensfs.org/events/lug-2022/

Lustre Persistent Client Cache

- Feature of Lustre Client Software Since 2.14
- Attach Cache Devices via Lustre Client Tool
  - Transparent Access
  - Any Cache Devices
- Read-Write/Read-only Mode

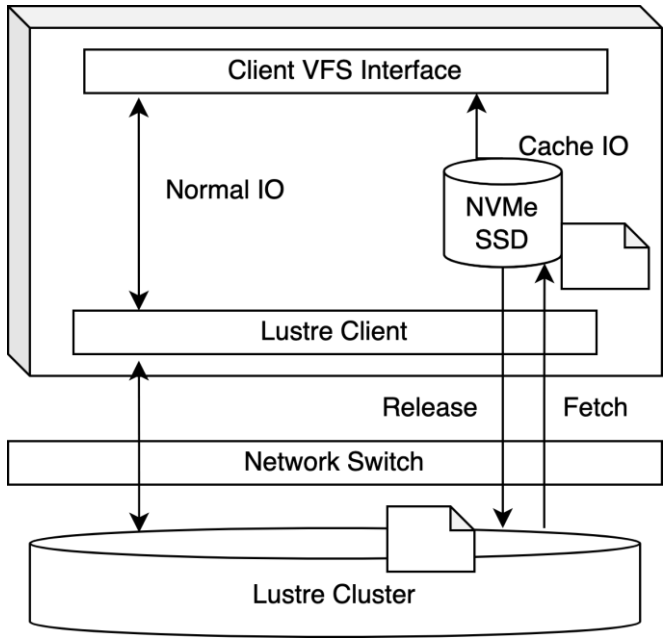

Burst Buffer Lua Plugin

- Generic Burst Buffer Interface for Slurm Workload Manager
  - Since Slurm Workload Manager 21.08 (https://slurm.schedmd.com/burst_buffer.html)
- Capability to Adapt Custom Script for Each Job State

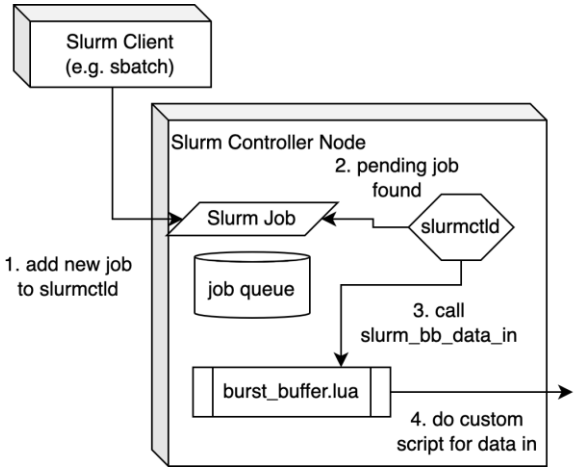

example: slurm_bb_data_in

Job State	burst_buffer.lua function
Job submission	slurm_bb_job_process
Job pending	slurm_bb_data_in
Job allocated resources not running yet	slurm_bb_pre_run
Job finished, completing state	slurm_bb_data_out
Job completed	slurm_bb_teardown

burst_buffer/lua and slurmscriptd
https://youtu.be/q5j5meDDPpA
https://slurm.schedmd.com/SLUG21/burst_buffer_lua.pdf

Design of Burst Buffer Lua for Lustre

- Develop as Open Source Eco System
  - lfsccm (lustre file system client cache manager) package Installation on Controller Node
  - Slurm like configuration for resource
  - Few Lustre Native Configuration for Persistent Client Cache (*1)

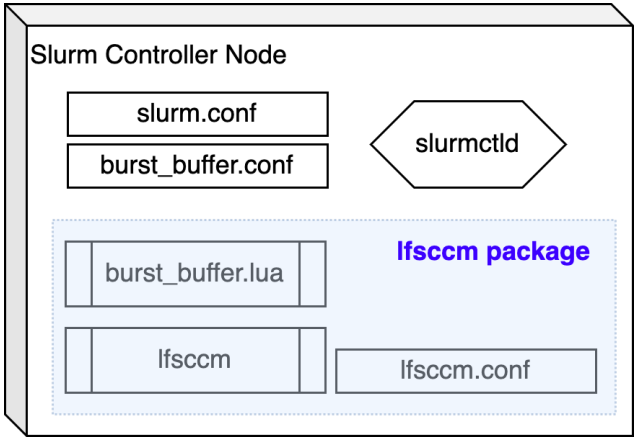

Install:
pip install lfsccm

lfsccm.conf:
NodeName=<compute_node> rwid=<cache id> roid=<cache id>

*1: https://doc.lustre.org/lustre_manual.xhtml#pcc

Burst Buffer for Lustre: How it works

- Stage In with “#PCC” directive jobs
- slurmctld calls BB Lua plugin and control the cache via lfsccm

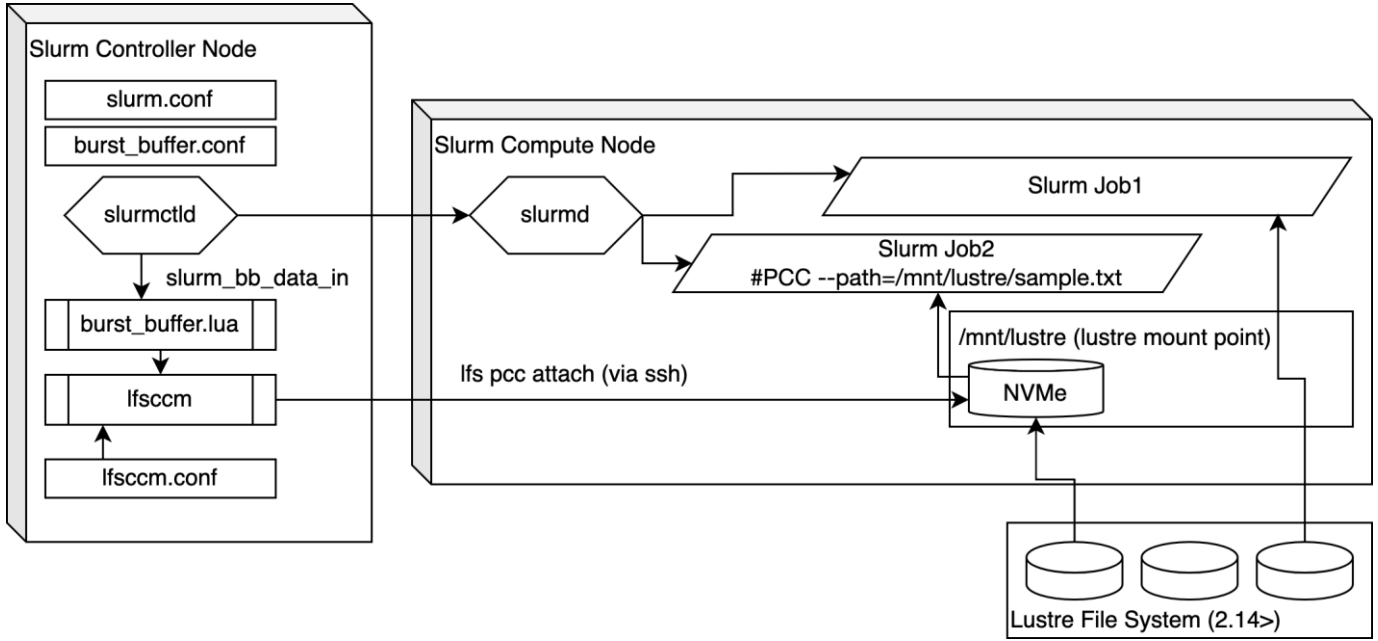

Demo

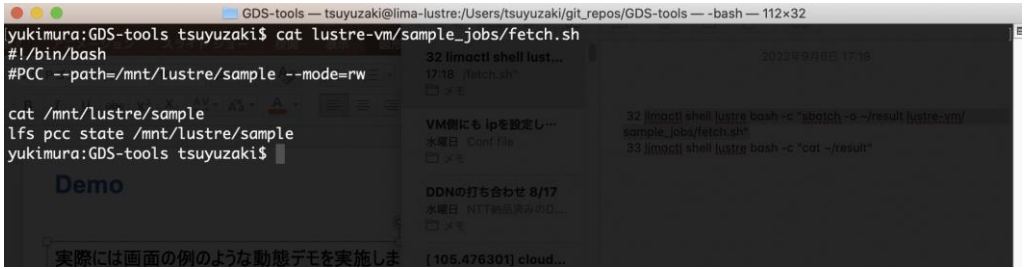
GDS-tools — tsuyuzaki@lima-lustre:/Users/tsuyuzaki/git_repos/GDS-tools — -bash — 112×32
yukimura:GDS-tools tsuyuzaki$ cat lustre-vm/sample_jobs/fetch.sh
#!/bin/bash
#PCC --path=/mnt/lustre/sample --mode=rw
cat /mnt/lustre/sample
lfs pcc state /mnt/lustre/sample
yukimura:GDS-tools tsuyuzaki$
Demo

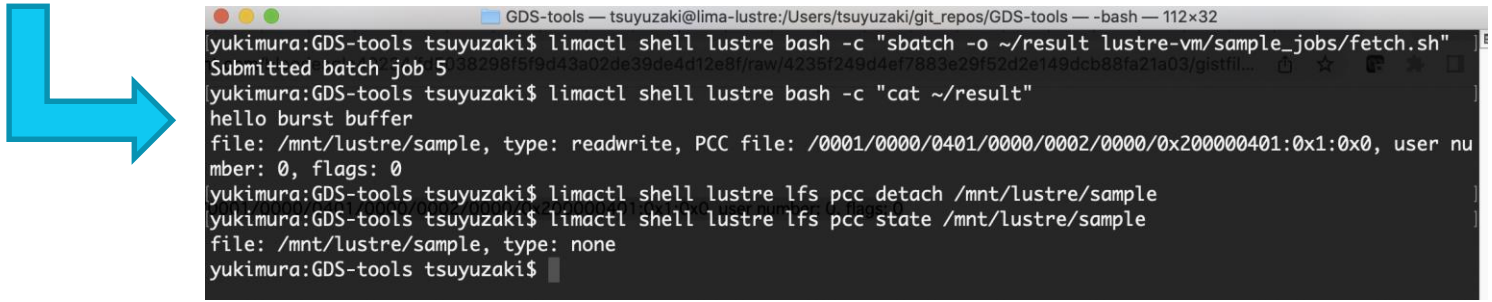
GDS-tools — tsuyuzaki@lima-lustre:/Users/tsuyuzaki/git_repos/GDS-tools — -bash — 112×32
yukimura:GDS-tools tsuyuzaki$ limactl shell lustre bash -c "sbatch -o ~/result lustre-vm/sample_jobs/fetch.sh"
Submitted batch job 5
yukimura:GDS-tools tsuyuzaki$ limactl shell lustre bash -c "cat ~/result"
hello burst buffer
file: /mnt/lustre/sample, type: readwrite, PCC file: /0001/0000/0401/0000/0002/0000/0x200000401:0x1:0x0, user nu
mber: 0, flags: 0
yukimura:GDS-tools tsuyuzaki$ limactl shell lustre lfs pcc detach /mnt/lustre/sample
yukimura:GDS-tools tsuyuzaki$ limactl shell lustre lfs pcc state /mnt/lustre/sample
file: /mnt/lustre/sample, type: none
yukimura:GDS-tools tsuyuzaki$

Performance Summary (Lustre's IO Traffic)

Case: Large Scale GAN Training

Lustre Traffic (Burst Buffer ON)

only 1 peak for fetching all data by "data in" process

Lustre Traffic (Burst Buffer OFF)

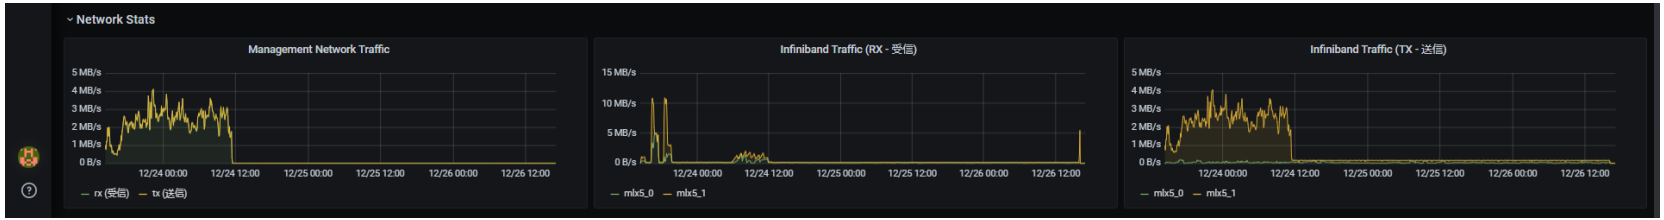

BB Lua for Lustre: Limitations

- Only RW mode is available in Lustre 2.14
  - Lustre 2.16 is candidate for RO mode
  - lfsccm already supports both RO and RW
    - › Use development version in our testing
- No Custom Message Interface of Generic BB Lua Specification
  - lfsccm FAILS IMMEDIATELY with invalid configuration or with job parameters
  - Emitting WARNING message to users without FAIL would be suitable in several situations

Project Repository and Testing Environment

- Project Repository
  - https://github.com/DDNStorage/lfsccm

- Demo/Testing
  - Sample Configuration and Setup All-In-One Scripts available in lfsccm project
    - Require Lima (https://github.com/lima-vm/lima)
  - Launch `lfsccm/demo.sh`
    - lfsccm/lustre-vm/lustre.yaml # lima project file for VM
    - lfsccm/lustre-vm/lustre/* # Lustre setup scripts
    - lfsccm/lustre-vm/slurm/* # Slurm and BB lua for lustre setup scripts

Summary

- Slurm Burst Buffer is for Cluster Scaling
- Burst Buffer Lua for Lustre (lfsccm project)
  - Easy Deployments
  - Native Slurm/Lustre Supports
  - Efficient Caching

Enjoy!!
Your Burst Buffer!!

E-mail:
Kota Tsuyuzaki (kota.tsuyuzaki.pc@hco.ntt.co.jp)
Repository:
https://github.com/DDNStorage/lfsccm