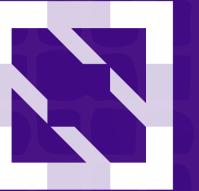




KubeCon



CloudNativeCon

 OPEN SOURCE SUMMIT

China 2019



CloudNativeCon



KubeCon



OPEN SOURCE SUMMIT

China 2019

在Kubernetes上运行AI/DL作业

Da Ma (@k82cn)

AI/DL, Spark

Service Workload



Frameworks



Infra



A vertical dashed line separates the Service Workload from the High Performance Workload.

Recommendation

Search

High Performance Workload



Gaps for AI/DL, BigData
and so on



kubernetes



- **Job/Queue Management**

- Queue status/configuration
- Hierarchical queue
- Job with multiple pod template
- Lifecycle management of Job, e.g. restart, suspend/resume
- Error Handling, e.g. restart job if pod failed (MPI, TFJob)
- Indexed Job
- Task dependency, e.g. Spark (executor/driver)
- Delay Pod Creation
- ...

- **Runtime**

- Singularity
- ...

- **Scheduler**

- Coscheduling
- Faire-share
- Queue
- Preemption/Reclaim
- Reserve/Backfill
- Topology (network, accelerator)
- ...

- **Others**

- Throughput
- Round-trip
- Data locality (Data Aware Scheduling)
- ...

Volcano: A Kubernetes Native Batch System



Website: <https://volcano.sh>

Github: <http://github.com/volcano-sh/volcano>

Twitter: https://twitter.com/volcano_sh

Slack: <http://volcano-sh.slack.com>

Email: volcano-sh@googlegroups.com



Job/Queue Demo

- Queue status/configuration
- Hierarchical queue
- Job with multiple pod template, (MPI, TensorFlow)
- Lifecycle management of Job, e.g. restart, suspend/resume
- Error Handling, e.g. restart job if pod failed
- Indexed Job
- Task/Job dependency, e.g. Spark (executor/driver)
- Delay Pod Creation

Scheduler Demo

- Gang-Scheduling (e.g. MPI, TensorFlow training)
- Fair-Sharing between jobs (e.g. Spark)
- Fair-Sharing between namespaces cross queue
- Task priority
- Proportion between queues
- Preemption/Reclaim
- Backfill
- Reservation

Q & A