YARN-5864 queue priority preemption performance report - V1

Wangda Tan

Test Environment

Use SLS (Scheduler Load Simulator) to run following tests:

- Mocked 4K nodes, each node has 128G memory, so in total the cluster manages 500TB memory resources.
- Submitted 600 applications to the cluster, size of application varies
- Resource of containers is ranged from 8G to 64G, subject to uniform distribution.
- Lifespan of containers is ranged from 19 sec to 96 sec, subject to uniform distribution
- Node heartbeat interval is 1 sec.
- Run the SLS test for 5 mins, and get reserved / allocated resource of queues and cluster.

Test machine configuration:

- OSX 10.12.2
- 2.5 GHZ i7
- 16GB memory
- SSD

Test configuration (All use Capacity Scheduler)

- Queue configuration
    - 4 Queues, A / B / C / D
    - A.priority=2, capacity=5%
    - B.priority=1, capacity=10%
    - C.priority=1, capacity=25%
    - D.priority=0, capacity=60%
- Conf 1: Enable queue priority preemption, allow move reserved container
- Conf 2: Enable queue priority preemption, disallow move reserved container
- Conf 3: Disable queue priority preemption

Test Result

Cluster total reserved memory

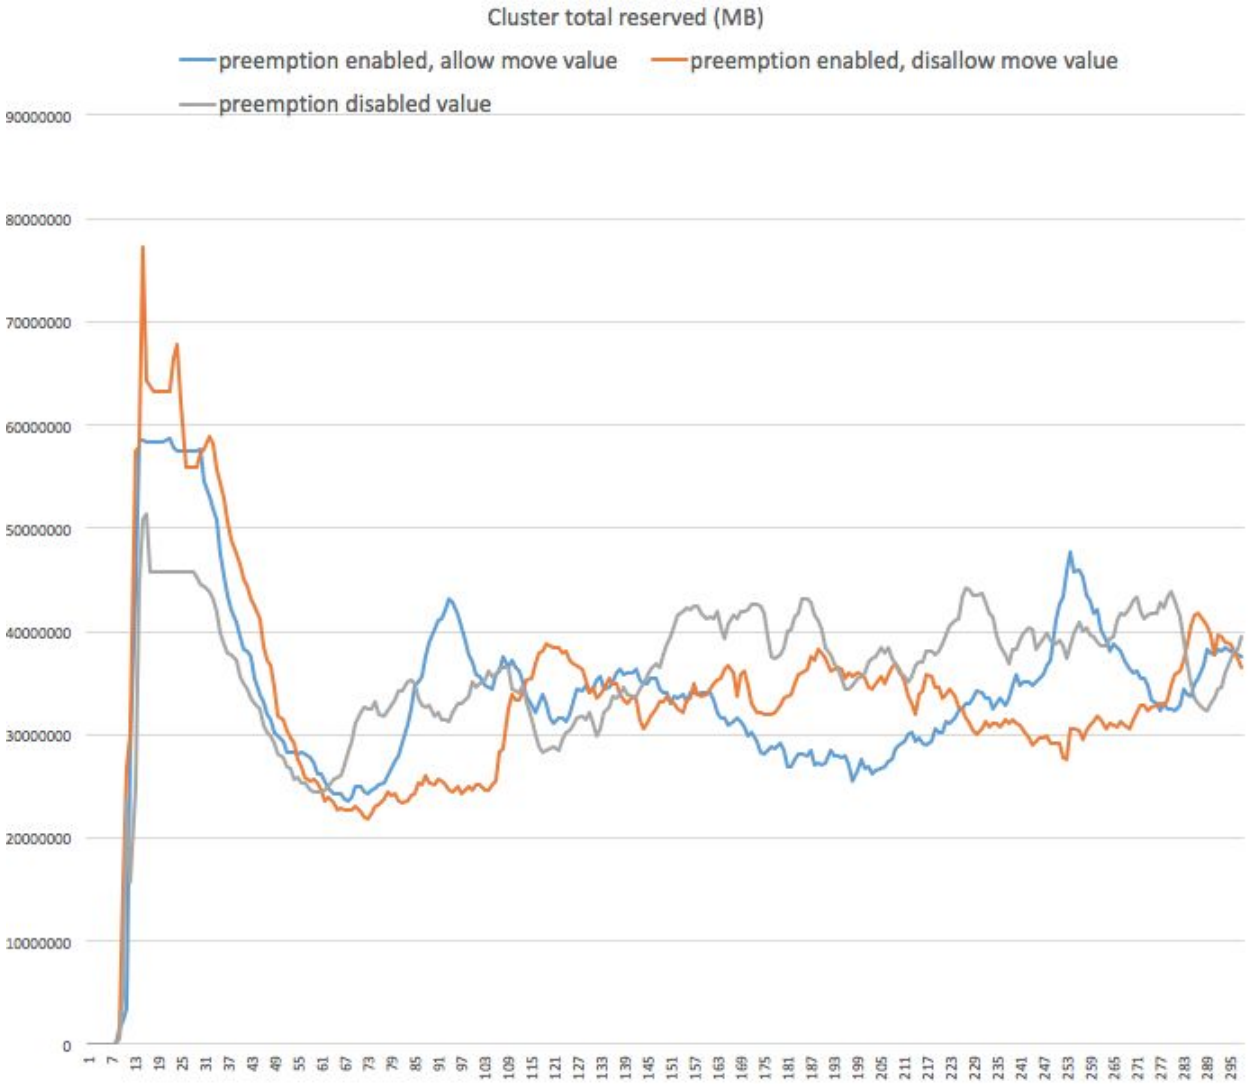

Chart 1: Cluster Total Reserved Resource (in MB), **lower is better**

Above chart shows total reserved resource in the cluster over time. When queue priority preemption is disabled, cluster has higher reserved resource, which means lower utilization.

Queue's total allocated memory (deducted reserved memory)

Queue-A's total allocated memory

Chart 1: Queue A Total Allocated Resource (in MB), **higher is better**
Queue A has highest priority, so by expectation it should get higher allocated memory. Queue A can get higher allocated resource when preemption enabled.

Queue-B's total allocated memory

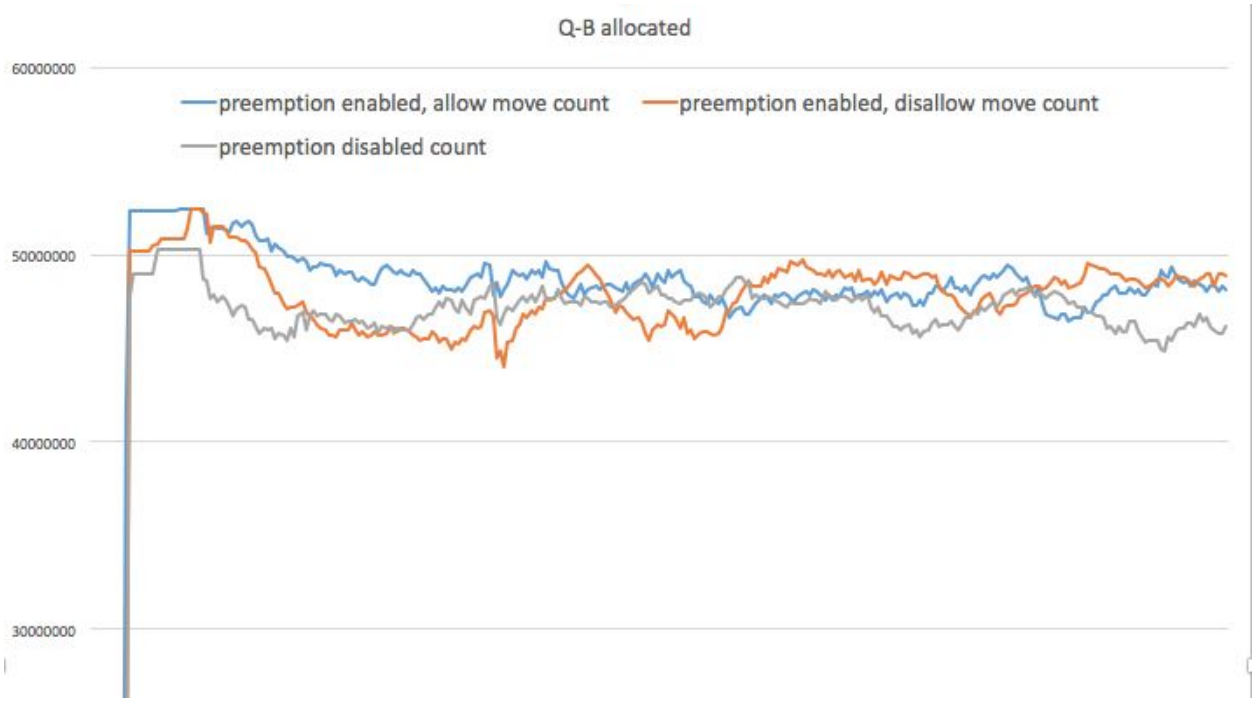

Queue-C's total allocated memory

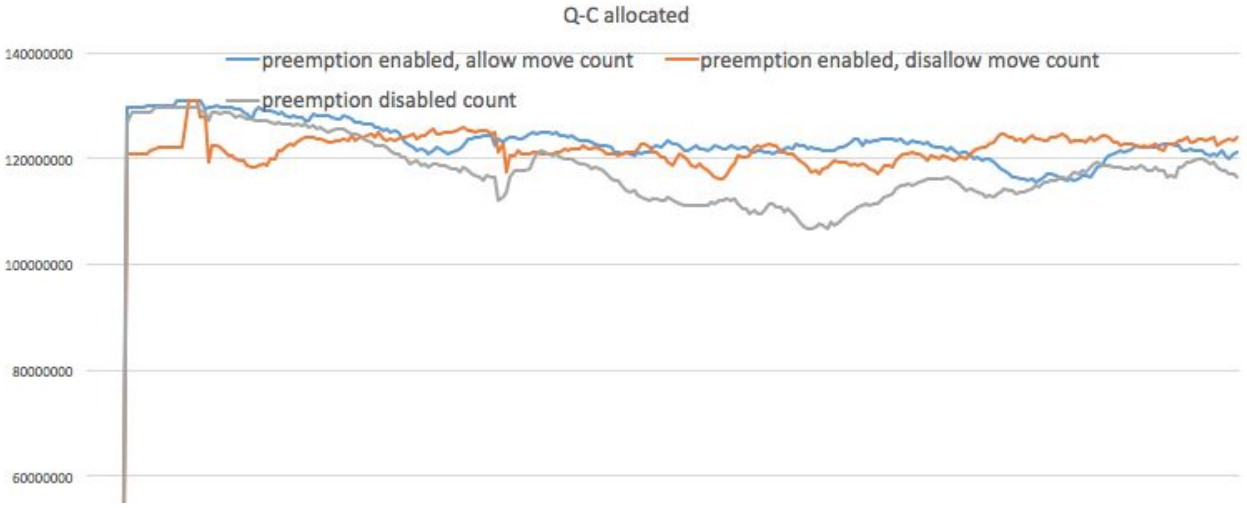

Queue-D's total allocated memory

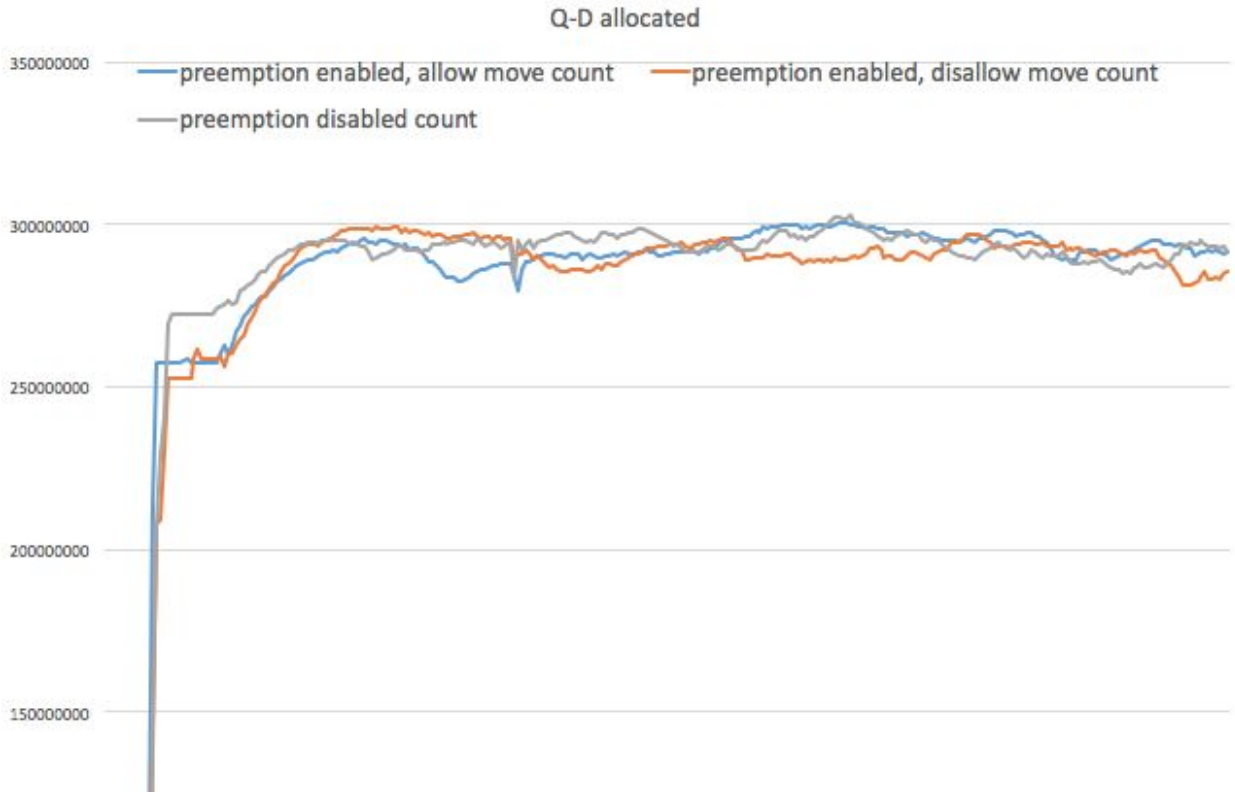

Time spent on preemption calculation (per round)

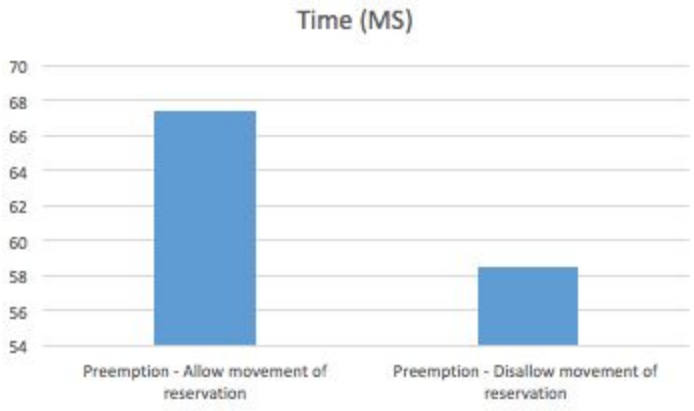

Above chart is the time spent to calculation preemption for priority queue preemption. (In ms)

Summary

According to the test report:

- When queue priority preemption is enabled:
    - Queue with higher priority can get better utilization.
    - Queue with lower priority doesn’t have significant regression of resource utilization
- Time spent on queue priority preemption is reasonable.