Introduction to Slurm and k8s

Ziqin Li

System Research Association @ Sichuan University

November 7, 2024

Cluster Management: Shared HPC Cluster's view

- Multi-tenancy: Admission Control
- Assume all tenants are able to handle with environment issues: Container is not the default method for running applications
- Scheduler: limited resources with batch of jobs

Slurm's Architecture

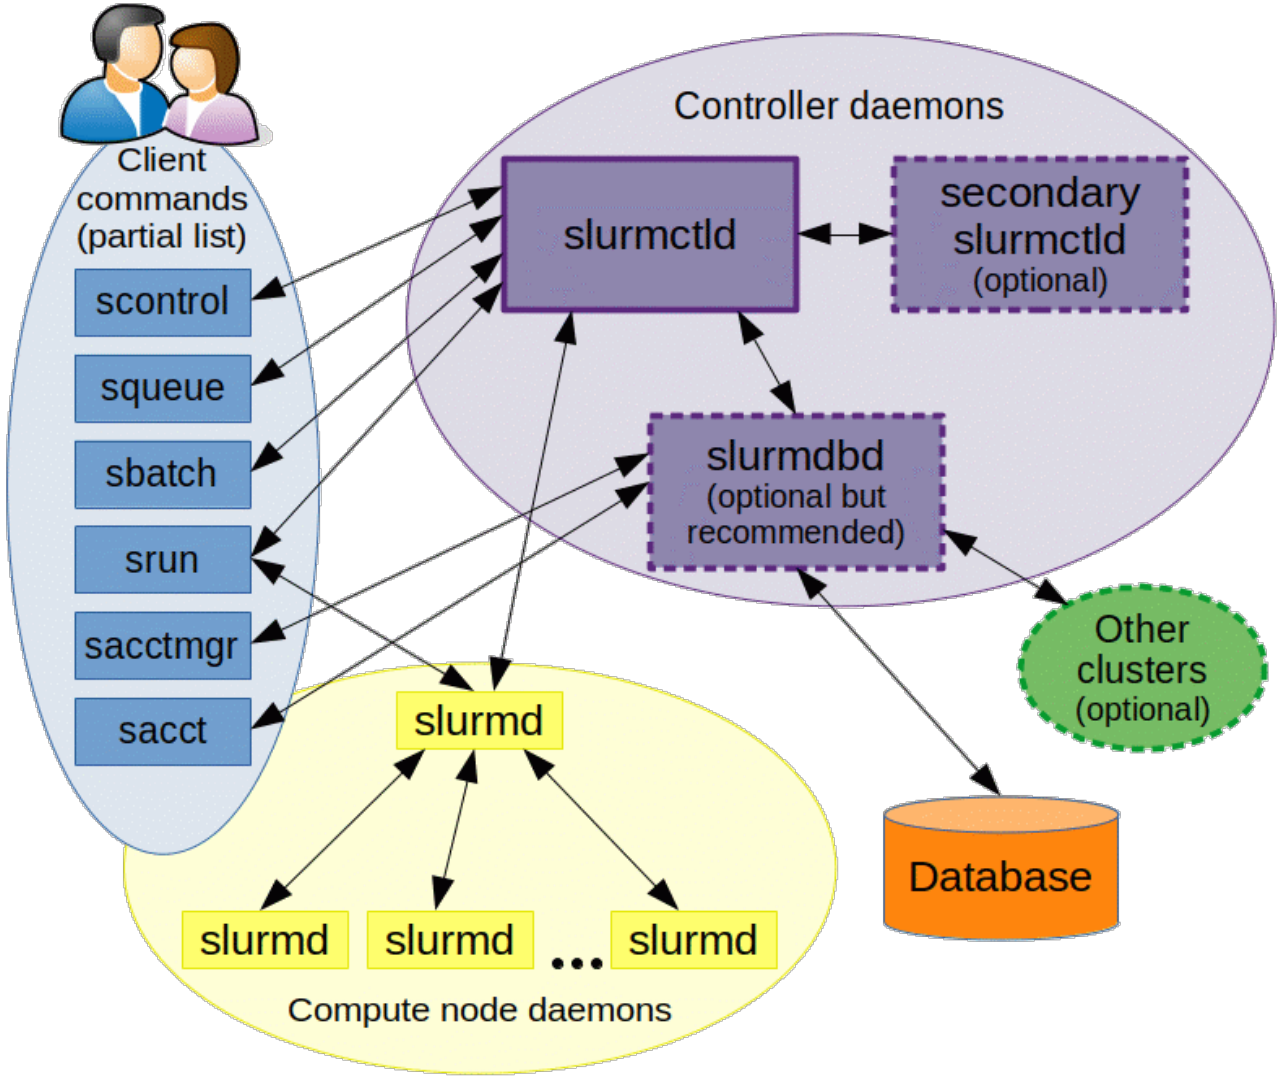

Daemons: slurmd, slurmctld, slurmdbd

- slurmd: recieve and executes tasks
- slurmctld: serve as the control plane of clusters: send control messages (start, stop, change node's state etc)
- slurmdbd: interfaces for recording informations (including running time, job details ..)

Usages

- sinfo: show nodes' states
  cluster is divided into several partitions
  sinfo -n nodename show node information
- squeue: list the waitlist for jobs
- srun: run jobs by command line
- sbatch: submit batch job

Cluster Management: Data Center's View

- Orches
- fault-tolerent
- scheduler: infinity resources and limited resources request
- extensiblity
- auto-run jobs: container

Kubernetes

- Key: Orchestrate pods, nodes, and svcs
- Pods: Basic component in K8s, consisted of multi containers
- Deployments / Replicas: Another indirect layer of pods
- Services: Interfaces between Pods; Ingress / Gateway can export them to Internet

References

Thank You!