SIT Slurm Service Architecture

Docker-based Multi-User GPU Platform on AMD MI210

1. System Architecture

This system is a Slurm cluster running on Ubuntu Server 24.04. Its roles are separated into three main containers to support up to four concurrent users. Workspace, identity, job scheduling, and GPU control are isolated while operating as an integrated system.

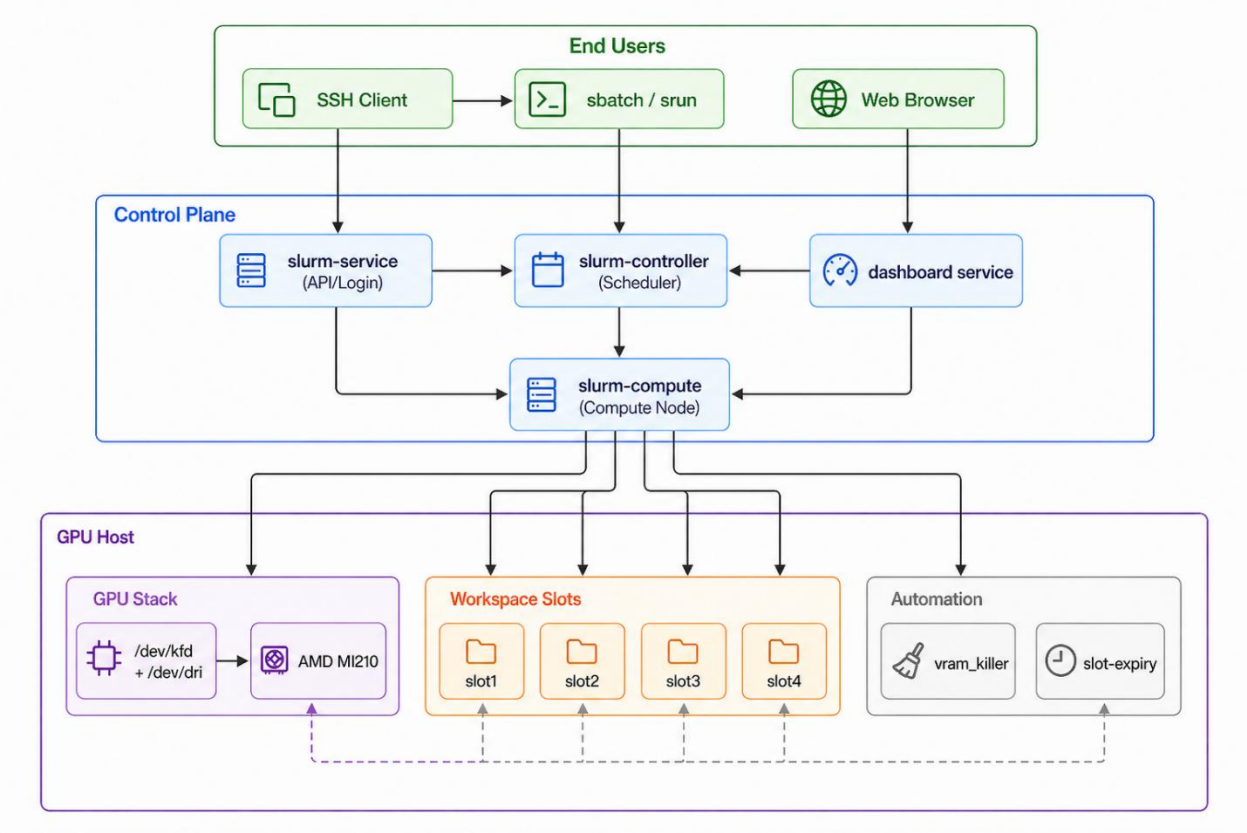

2. Control Plane

- slurm-controller is the central component for job queue management and cluster state management.
- slurm-compute is the node that runs GPU workloads.
- slurm-service is the access point for SSH connections and job submissions.

3. Identity and Slot Model

- The system uses a slot-based identity model as its core design.
- Four slots are available for concurrent users.
- Each slot is assigned a fixed UID in the range 9001-9004.
- Users log in with their own usernames, while their workspace and lifecycle are isolated by slot.

4. Storage and Isolation

- The data for each user slot is stored separately in /workspace/slot[*x*].
- When a slot is revoked, its data is backed up before that slot is cleared.
- The system is designed to prevent data deletion across slots.

5. GPU and VRAM Enforcement

- GPU pass-through is provided to slurm-compute through /dev/kfd and /dev/dri.
- vram_killer.service monitors VRAM usage and enforces a 15 GB cap per user, checking every 5 seconds.
- In production mode, uid is the default grouping method so that all processes owned by the same user are included in the usage check.
- Before terminating processes, the system writes vram_killer_notice and vram_killer.log files to the workspace.

6. Monitoring and User Experience

- The dashboard is a centralized web interface for viewing slot status and utilization at http://slurm-service.sit.kmutt.ac.th:8080.
- The dashboard reads slots.json and combines Slurm job-state data with host metrics.

- The UI displays usernames without showing UIDs.
- The web page refreshes automatically every 10 seconds and includes a Refresh button with a loading state.
- slurm-service displays the system banner and the last login information.

7. Key Limits

- Slot count: 4
- GPU VRAM cap: 15 GB per user
- Storage: 100 GB per slot
- Dashboard refresh: every 10 seconds
- Production VRAM mode: uid