

GPU cluster monitoring

Tom Rochette <tom.rochette@coreteks.org>

July 7, 2025 — [dad4e22e](#)

In this article I list the various metrics/alerts one should have when monitoring a GPU cluster to ensure efficient usage.

- Allocated GPUs are used
 - Used to detect jobs that may ask multiple GPUs but end up using 1 or only a few of them
- GPU utilization below threshold (e.g., 10%)
 - Used to detect workloads that do not make full use of the GPU or are allocated to an oversized GPU
- GPU utilization above threshold (90%)
 - Used to detect when the GPU is saturated
- GPU utilization range
 - Used to detect uneven distribution of GPU compute workload
- GPU memory utilization above threshold (e.g., 10%)
 - Used to detect workloads that do not make full use of the GPU or are allocated to an oversized GPU
- GPU memory utilization above threshold (e.g., 95%)
 - Used to detect when a job is about to run out of GPU memory
- If using InfiniBand
 - InfiniBand receive/transmit > 0 when running multi-node workloads
 - * Used to identify workloads that are not properly configured to use InfiniBand