

DICC UM HPC: GPU A100 Utilisation Threshold & Auto-Cancellation Policy

Data-Intensive Computing Centre
Universiti Malaya

16 October 2026



1. OBJECTIVES

- 1.1. The DICC GPU-A100 Utilisation Threshold & Auto-Cancellation Policy is a document that describes the details of a background scheduler that runs on **GPU-A100** nodes (gpu06 & gpu07) every 5 minutes. It watches for GPU jobs that barely utilize their GPU, specifically, jobs using **less than 40% GPU for 1 hour straight**.
- 1.2. To ensure efficient GPU-A100 resource allocation, jobs requiring under 32 GB of VRAM (40% of GPU-A100 VRAM) must be routed to other partitions (i.e., gpu-k40c, gpu-titan, and gpu-v100s), reserving high-capacity partitions for memory-intensive workloads.
- 1.3. If your job stays idle as stated, it will be **cancelled automatically** by the background scheduler.

2. AUTO-CANCELLATION IMPLEMENTATION

- 2.1. Every 5 minutes, the background scheduler looks at every GPU job running on the node.
- 2.2. It measures how much GPU each job is actually using, directly from `nvidia-smi`, not from Grafana. Grafana data can sometimes be unstable or delayed, which may lead to false positives. Using `nvidia-smi` directly gives a more accurate and stable reading.
- 2.3. If your GPU usage drops below 40%, it starts a timer.
- 2.4. If usage stays below 40% for 1 hour straight, your job will be cancelled automatically.
- 2.5. If your GPU usage goes back up above 40% at any point, the timer resets to zero.
- 2.6. **Important:** The background scheduler only looks at a snapshot every 5 minutes. It doesn't watch continuously. If your GPU spikes for 2 minutes between checks, it won't see that.

3. FREQUENTLY ASK QUESTION (FAQ)

3.1. **Q:** Will my job be cancelled?

A: Yes. If your GPU stays below 40% for 1 hour straight, your job will be cancelled automatically on the next check.

3.2. **Q:** I just started my job and it's loading data. Will it get cancelled?

A: Jobs “younger” than 5 minutes are skipped. But if your loading phase takes longer than 1 hour, it could be cancelled. Consider loading data on CPU first, then initialising GPU work.

3.3. **Q:** My job uses GPU intermittently (bursts of work, then idle). Will it be cancelled?

A: It depends. If the idle periods between bursts are longer than 5 minutes, the checker might catch it during an idle moment. But if your GPU usage goes above 40% even once between checks, the timer resets.

3.4. **Q:** Why not use Grafana to check GPU usage?

A: Grafana data can sometimes be unstable or delayed, which may cause false positives. We use `nvidia-smi` directly to get accurate, real-time readings from the GPU driver itself.

3.5. **Q:** I fixed my code and resubmitted. Will the new job be affected?

A: The timer resets for each new job. Your resubmitted job starts with a clean state, if it uses GPU properly, it won't be flagged.

3.6. **Q:** How often does the checker run?

A: Every 5 minutes on each GPU node.