Power up your AI pipelines with CircleCI's new GPU compute

Hello,

Today we introduced an update to our compute offering that will let you build, test and deploy AI applications without managing your own hardware.

New GPU Resource Classes:
To further support developers building AI-enabled applications, we're introducing three new GPU resource classes:

- gpu.nvidia.small.gen2: 1 NVIDIA Tesla A10G GPU, 4 CPU, 16 GB RAM
  gpu.nvidia.small.multi: 2 NVIDIA Tesla T4 GPUs, 4 CPU, 16 GB RAM
- gpu.nvidia.medium.multi: 4 NVIDIA Tesla T4 GPUs, 8 CPU, 30 GB RAM
-

Built on the Nvidia chipset, these AWS G5 compute options provide the performance needed for compute-intensive AI jobs. This is currently available in the Scale plan.
The full resource classes pricing details and conditions are available on our website.

Preloaded Linux CUDA Image for AI/ML:
CircleCI's Linux CUDA image now includes essential AI/ML-specific libraries, such as PyTorch, TensorFlow, and Hugging Face transformers, right out of the box.
This image is available exclusively to jobs using the gen2 GPU resource class.
Explore the details in our documentation.

Thank you for being an essential part of the CircleCI community.
Best regards,
Alex, Benny & Webster
CircleCI Team

P.S. Our product team is actively looking for engineers building AI-enabled projects with AI to discuss our next roadmap choices. If you are interested in taking part in the future of CircleCI, please book a 30-minute slot with our AI team.

—

Share this newsletter with your peers building AI-powered applications! They can subscribe here.