# Use Case 8: AI for green multi-cluster: Intelligent management towards green and low-carbon, large-scale multi-clusters

**Organization:** Institute of Artificial Intelligence (TeleAI), China Telecom

**Country**: China

**Contact Person(S)**:

Primary: Qizhen Weng, wengqzh@chinatelecom.cn

Secondary: Yuankai Fan, fanyk1@chinatelecom.cn

## 1 Use Case Summary Table

| Item | Details |
|---|---|
| Category | Climate Change/Natural Disaster |
| Problem Addressed | Using AI to efficiently managing computing resource and scheduling jobs across multiple large-scale AI clusters. This AI-driven approach aims to reduce energy consumption and carbon emissions, surpasses traditional rule-based multi-cluster management methods in adaptability and effectiveness, and is capable of explaining its decisions to humans via multi-modal generative models. |
| Key Aspects of Solution | • AI-powered job scheduling that improves efficiency through real-time multi-dimensional resource monitoring.<br>• Employing GPU multiplexing and shared resource scheduling algorithms to enhance cluster utilization.<br>• Cross-cluster coordination via multiple AI agents towards overall green and low-carbon goals.<br>• Employing multi-modal generative large models to provide interpretable explanations of scheduling decisions, enhancing transparency for end-users. |
| Technology Keywords | Job Scheduling, Multi-Cluster Management, Multi-Modal Generative Models, GPU Multiplexing, Kubernetes-based Cluster Coordination |
| Data Availability | Private; Public<br>1) https://github.com/alibaba/clusterdata [1]<br>2) https://github.com/InternLM/AcmeTrace [2]<br>3) https://github.com/ml-energy/zeus [3] |
| Metadata (Type of Data) | Text data (job resource requirements and performance logs, energy consumption metrics, energy source information, etc.) and images (cluster metrics represented through graphs) |

(continued)

| Item | Details |
|---|---|
| Model Training and Fine-Tuning | Models include reinforcement learning agents and large language models trained on extensive cluster resource and task operation data to learn optimal scheduling policies. Fine-tuning involves adapting models to dynamic multi-cluster environments, incorporating real-time monitoring data, and improving interpretability through multi-modal generative models. The Fragmentation Gradient Descent (FGD) algorithm is used for efficient GPU-sharing job scheduling, validated through simulation and real cluster deployments. |
| Testbeds or Pilot Deployments | • https://github.com/hkust-adsl/kubernetes-scheduler-simulator [4]<br>• https://github.com/qzweng/clusterdata/tree/master [5]<br>• https://www.usenix.org/conference/atc23/presentation/weng [6]<br>• https://www.usenix.org/conference/nsdi22/presentation/weng [7]<br>• https://sc20.supercomputing.org/proceedings/tech_paper/tech_paper_pages/pap211.html [8]<br><br>The links above contain information regarding deployments around scheduling GPU-sharing workloads, including novel measure of fragmentation to statistically quantify the extent of GPU fragmentation caused by different sources, Metis which is a general-purpose scheduler using deep reinforcement learning (RL) techniques, etc. |

## 2 Use Case Description

### 2.1 Description

The soaring popularity of generative AI models, such as ChatGPT [9], Midjourney [10], and DeepSeek [11], has enriched our lives and created a wealth of new employment opportunities. However, fueling AI requires substantial computational power and energy consumption, usually within AI clusters--data centers equipped with advanced AI processors, such as GPUs.

A recent report from the International Telecommunication Union (ITU) on international standards for AI and its environmental impact [12] highlights that energy consumption is driven primarily by the intensive processing required for training and inference of large AI models. These models are so energy-demanding that the computational power needed to support AI's expansion has been doubling around every 100 days since 2010 [13]. Beyond energy requirements, the vast quantities of fresh water needed to cool AI processors further strain local water resources and ecosystems [14]. This underscores the growing environmental footprint of AI and the urgent need to adopt efficient management practices for AI clusters.

**Intended Use:** Utilize artificial intelligence technology to achieve intelligent management and coordinated scheduling of multiple computing clusters, promote the green development of technology, enhance resource efficiency, and reduce energy consumption and carbon emissions.

**Problem to Solve:** With the increase in the number and scale of computing clusters, how to manage resources such as energy, computing power, and bandwidth in multiple clusters, and efficiently schedule and run a large number of computing tasks.

**Limitations of Existing Solutions:** Traditional multi-cluster management methods relying on manual experience and simple rules lack real-time performance and effectiveness when facing a multi-cluster environment with dynamically changing resource supply. They are unable to quickly identify the optimal resource allocation according to the actual situation.

Benefits and Drawbacks of the AI-based Approach:

1) It can adapt to environmental changes 24/7. Based on factors such as the number of running tasks, cluster computing power utilization, energy supply, and carbon emissions, it can flexibly adjust the resource scheduling scheme for multiple clusters to achieve the goal of green and low - carbon development.
2) It can perform performance modeling on computing tasks based on massive resource monitoring and task operation data, etc., thereby improving the overall computing performance.
3) It can use multi-modal generative large models to explain to humans the reasons behind scheduling decisions.

The proposed use case emphasizes the role of AI in enabling innovative and green and economical solutions for sustainable development. By proposing the requirements for intelligent low-carbon management, this system will boost the construction of data collection capabilities and low-carbon capabilities in industrial production. Additionally, by optimizing the utilization of infrastructure, the system enables limited infrastructure to support more tasks, indirectly reducing the need for new infrastructure construction and optimizing resource investment. This system complies with the goal of responsible consumption and production by optimizing the green utilization of resources throughout the AI lifecycle. It will promote green energy consumption patterns in computing clusters such as data centers and reduce waste and carbon emissions.

**Use Case Status:** The use case is part of a larger research project

Partners: N/A

## 2.2 Benefits of the use case

1) It can adapt to environmental changes 24/7. Based on factors such as the number of running tasks, cluster computing power utilization, energy supply, and carbon emissions, it can flexibly adjust the resource scheduling scheme for multiple clusters to achieve the goal of green and low - carbon development.
2) It can do performance modelling on computing tasks based on massive resource monitoring and task operation data, etc., thereby improving the overall computing performance.
3) It can use multi-modal generative large models to explain to humans the reasons behind scheduling decisions.

The proposed use case emphasizes the role of AI in enabling innovative and green and economical solutions for sustainable development. By proposing the requirements for intelligent low-carbon management, this system will boost the construction of data collection capabilities and low-carbon capabilities in industrial production. Additionally, by optimizing the utilization of infrastructure, the system enables limited infrastructure to support more tasks, indirectly reducing the need for new infrastructure construction and optimizing resource investment. This system complies with the goal of responsible consumption and production by optimizing the green utilization of resources throughout the AI lifecycle. It will promote green

energy consumption patterns in computing clusters such as data centers and reduce waste and carbon emissions.

## 2.3 Future Work

**Data Collection:** The dimension of data collection will be expanded to cover more energy and environmental parameters and detailed equipment operation information, so as to provide more comprehensive data support.

**Single-Node Resource Sharing:** Research will focus on developing more efficient resource-sharing algorithms to expand the scope of AI scheduling, including enhancements in GPU virtualization capabilities. The goal is to increase the utilization of GPU resources while minimizing idle waste, all while ensuring that task performance is maintained. Drawing from our operational experience at the Institute of AI, research teams across domains such as Computer Vision (CV), Natural Language Processing (NLP), and Speech have already virtualized a portion of their GPU resources. This GPU-sharing approach has led to an increase of at least 150% in the availability of GPUs for prototype development and debugging, without the need for additional processor purchases.

**Single-Cluster Resource Scheduling:** Monitoring and task data will be regarded as the state space, and scheduling optimization strategies will be regarded as the action space. Technologies such as reinforcement learning and large language models will be used to establish intelligent scheduling strategies, making them more suitable for complex and variable task requirements and further improving the resource utilization efficiency within a single cluster. The de facto standard for cluster coordination is Kubernetes, which is widely adopted by leading AI companies such as OpenAI [12]. Building on Kubernetes, we have open-sourced a scheduling algorithm known as the Fragmentation Gradient Descent (FGD) policy [13]. This algorithm enables highly efficient scheduling decisions within hundreds of milliseconds on clusters comprising over 1,200 servers. Simulations have demonstrated that FGD outperforms traditional bin-packing algorithms in GPU-sharing scenarios, reducing unallocated GPUs by up to 49% [14].

**Multi-Cluster Collaborative Scheduling:** With the help of technologies such as multi-agent large models, a more intelligent cross-cluster collaboration mechanism will be constructed. According to the real-time status of each cluster and task priorities, the optimal cross-cluster resource allocation will be achieved, promoting the achievement of overall green and low-carbon goals.

**Cluster Scheduling Result Interpretability:** Multi-modal generative large models will be utilized to present the scheduling decision-making process to users in an understandable form, enhancing trust in artificial intelligence scheduling results and facilitating subsequent strategy adjustment and optimization.

# 3 Use Case Requirements

- **REQ-01**: It is required to monitor multi-dimensional resource usage, profile typical job resource requirements, and collect environmental and infrastructural data from AI clusters.
- **REQ-02**: It is required to develop data-driven scheduling algorithms capable of adapting to dynamic resource availability, and conduct a preliminary evaluation of their performance using a cluster simulator.

- **REQ-03**: It is recommended to enable resource sharing to maximize hardware utilization.
- **REQ-04**: It is recommended to coordinate multiple clusters with intelligent agents to make wise scheduling decisions.
- **REQ-05**: It is recommended to use multi-modal generative models to explain scheduling decisions to human users.

## 4 Sequence Diagram

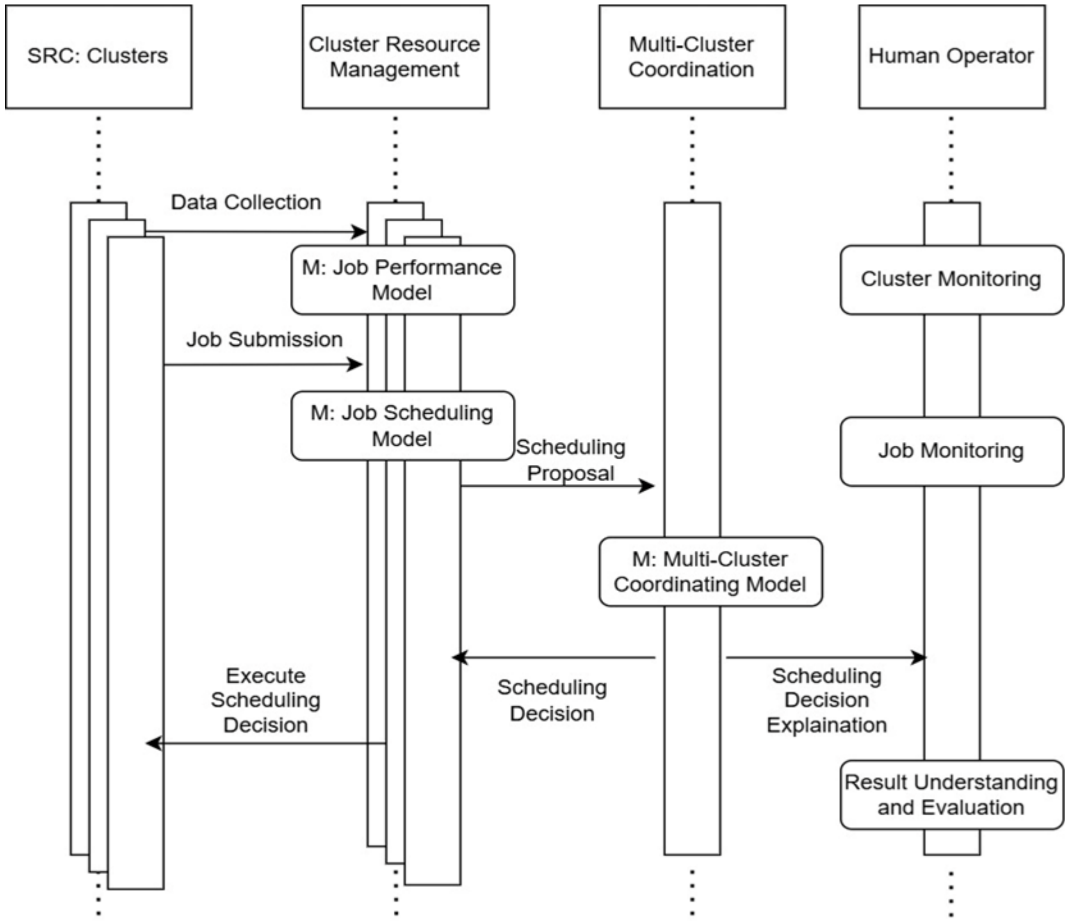


## 5 References

[1] "Alibaba Cluster Trace Program," *GitHub*, Feb. 04, 2023. https://github.com/alibaba/clusterdata

[2] InternLM, "GitHub – InternLM/AcmeTrace," *GitHub*, 2023. https://github.com/InternLM/AcmeTrace (accessed Jun. 09, 2025).

[3] ml-energy, "GitHub – ml-energy/zeus: Measure and optimize the energy consumption of your AI applications!," *GitHub*, May 17, 2025. https://github.com/ml-energy/zeus (accessed Jun. 09, 2025).

[4] hkust-adsl, "GitHub – hkust-adsl/kubernetes-scheduler-simulator: Kubernetes Scheduler Simulator," *GitHub*, May 11, 2023. https://github.com/hkust-adsl/kubernetes-scheduler-simulator.

[5] qzweng, "GitHub – qzweng/clusterdata: cluster data collected from production clusters in Alibaba for cluster management research," *GitHub*, 2021. https://github.com/qzweng/clusterdata/tree/master (accessed Jun. 09, 2025).

[6] W. Gropp, Torsten Hoefler, Rajeev Thakur, and Ewing Lusk, *Using advanced MPI modern features of the Message-Passing Interface*. Cambridge, Mass. The Mit Press, 2014.

[7] Q. Weng *et al.*, "{MLaaS} in the Wild: Workload Analysis and Scheduling in {Large-Scale} Heterogeneous {GPU} Clusters," *www.usenix.org*, 2022. https://www.usenix.org/conference/nsdi22/presentation/weng

[8] "Metis: Learning to Schedule Long-Running Applications in Shared Container Clusters at Scale," *Supercomputing.org*, 2025. https://sc20.supercomputing.org/proceedings/tech_paper/tech_paper_pages/pap211.html (accessed Jun. 09, 2025).

[9] OpenAI, "ChatGPT," *ChatGPT*, 2025. https://chatgpt.com/

[10] Midjourney, "Midjourney," *Midjourney*, 2024. https://www.midjourney.com/

[11] deepseek-ai, "GitHub – deepseek-ai/DeepSeek-V3," *GitHub*, 2024. https://github.com/deepseek-ai/DeepSeek-V3

[12] Wang L., Weng Q., Wang W., Chen C., Li B. (2020). Metis: Learning to Schedule Long-Running Applications in Shared Container Clusters at Scale. *SC*. Retrieved from https://sc20.supercomputing.org/proceedings/tech_paper/tech_paper_pages/pap211.html

[13] Weng Q., Xiao W., Yu Y., etc. (2022). MLaaS in the Wild: Workload Analysis and Scheduling in Large-Scale Heterogeneous GPU Clusters. *USENIX NSDI*. Retrieved from https://www.usenix.org/conference/nsdi22/presentation/weng

[14] Weng Q., Yang L., Yu Y., etc. (2023). Beware of Fragmentation: Scheduling GPU-Sharing Workloads with Fragmentation Gradient Descent. *USENIX ATC*. Retrieved from https://www.usenix.org/conference/atc23/presentation/weng