



CASE STUDY

How CodeContent produced 15 Kubernetes and GPU tutorials averaging 150 code blocks each for Civo

15 Tutorials

Kubernetes and GPU computing tutorials on Civo's Learn platform.

150 Code Blocks

Average per tutorial: Dockerfiles, K8s manifests, Python scripts, configs.

3 Year Partnership

Continuous engagement since May 2023.

◆ About Civo

Civo is a cloud infrastructure provider focused on Kubernetes and GPU computing, competing with DigitalOcean, Vultr, and major cloud providers for developer adoption. Their Learn platform publishes implementation tutorials for infrastructure engineers and ML teams. Learn more at civo.com

◆ The Challenge

Civo's Learn platform needed tutorials showing developers how to run real AI workloads on GPU-enabled Kubernetes clusters. Each tutorial requires spinning up actual clusters, configuring GPU nodes, deploying containers, and validating the workload runs. These are not tutorials you can write from documentation.

Finding writers who could provision infrastructure and deploy AI workloads (not just describe them) was the bottleneck holding back content velocity on the Learn platform.

◆ Why CodeContent

Civo evaluated CodeContent's existing portfolio of Kubernetes and AI infrastructure content. The combination of hands-on infrastructure experience (cluster setup, GPU workloads, container orchestration) and AI/ML implementation skills matched what their Learn platform required.

Civo needed tutorials that were:

✓ Validated on live Civo Kubernetes clusters

✓ Covering GPU computing and AI infrastructure

✓ Implementation-heavy with real Dockerfiles and manifests

✓ Aligned with Civo's product and strategic priorities

”

I have worked with Mostafa for nearly two years. He has consistently produced high-quality technical tutorials. His work stands out for its technical accuracy, clear structure, and depth. Mostafa is proactive in proposing strong tutorial ideas and does a great job shaping topics to align with Civo's products and strategic priorities.



Emma Stewart-Oram

Digital Marketing Executive at Civo

◆ The Work

The engagement required infrastructure + AI skills: provisioning K8s clusters, configuring GPU node pools, deploying containerized AI workloads, and validating everything before writing.

Kubernetes

GPU Nodes

Docker

FastAPI

Llama 4

Gemini

OpenHands

relaxAI

◆ Technical Depth

The [Automated Knowledge-Graph Generator](#) (Feb 2026) uses an open LLM on Civo GPU nodes to extract entities and relationships from text, then visualizes them as an interactive graph interface.

The [Self-Hosted AI Assistant with Llama 4](#) deploys Meta's Llama 4 on Civo Kubernetes with FastAPI, Docker, and the model's 10M token context window, fully containerized and validated on live infrastructure.

The [OpenHands on Civo](#) tutorial sets up an AI agent platform for automated code debugging, refactoring, and testing, running entirely on Civo infrastructure.

◆ Infrastructure Friction

Tutorials depended on Civo's GPU availability, cluster provisioning times, and infrastructure pricing. When Civo released new GPU node types or updated their Kubernetes API, tutorials needed revalidation against the latest infrastructure. The writer spent significant time in the Civo dashboard and CLI, not just in a code editor.

Each tutorial averaged over 2,500 words and 150 code blocks: Dockerfiles, Kubernetes manifests, Python scripts, API endpoints, and infrastructure configuration all in one piece.

◆ Before & After

DIMENSION	BEFORE	AFTER
Content velocity	Limited Learn content	2 tutorials/month, consistent
Writer expertise	Docs-based descriptions	Engineers provisioning real clusters
Technical depth	Overview-level guides	2,500+ words, 150 code blocks avg
Topic coverage	Core K8s docs	GPU, RAG, LLM deploy, agentic AI, cost opt
Validation	Code from docs	Every tutorial on live Civo clusters

◆ Published Tutorials

TUTORIAL	TOPIC
Knowledge-Graph Generator	LLM entity extraction + visualization
Self-Hosted AI with Llama 4	LLM deployment, FastAPI, Docker
OpenHands on Civo	Automated code debugging agent
GPU Cost Optimization	JIT scaling, budget-aware scheduling
RAG + Gemini on K8s	Financial forecasting RAG pipeline
Text-to-SQL on GPU	Natural language to SQL, GPU inference

◆ The Lesson

For cloud infrastructure companies, developers don't trust infrastructure claims; they trust tutorials where someone actually provisioned the cluster, deployed the workload, and showed it working. The code block count matters because infrastructure content that skips the Dockerfile, the K8s manifest, or the config step is content developers abandon halfway through.



Kubernetes-Native

Real clusters, real GPU nodes, real deployments.



AI Workloads

LLMs, RAG, knowledge graphs, agent frameworks.



150 Code Blocks

Dockerfiles, manifests, scripts, configs per tutorial.



Live Validated

Every tutorial tested on Civo infrastructure.