

Pablo del Arco

📍 Valencia, Spain ✉ pablodelarco1@gmail.com ☎ (+34) 628 854 252 in pablo-del-arco 🌐 pablodelarco

Cloud & DevOps Engineer — AI Agent Architect

Cloud & DevOps Engineer and **AI Infrastructure & Agent Architect** who builds production Kubernetes platforms (K3s, RKE2, Terraform, ArgoCD) and the autonomous agent systems that run on top of them. Deploys **AI inference** and federated learning stacks (llama.cpp, LiteLLM, Ray, Flower.ai) on **NVIDIA H100/A100 GPUs** across sovereign European cloud environments. Speaker at **FOSDEM 2026**, **Foss-North 2026**, and **IPCEI-CIS GA 2026**. Trilingual: Spanish (native), English (C1), French (B2).

Experience

Cloud & DevOps Innovation Engineer

Valencia, Spain

OpenNebula Systems

February 2024 – Present

- Build and deploy production VMs and **Kubernetes** platforms (RKE2, K3s) using **GitOps** (ArgoCD), CI/CD (GitHub Actions), and IaC (**Terraform**, **Ansible**).
- Technical Lead across **4 EU-funded projects** (IPCEI-CIS, O-CEI, EMERALD, CASANDRA), coordinating 20+ partners including SAP, Telefónica, Siemens, and DB InfraGO across federated **GPU infrastructure** deployments in Spain, Netherlands, and Poland.
- Architected and deployed **Fact8ra** [🔗](#), a federated **European Virtual Hyperscaler** spanning 10 zones across 8 countries with 5 sovereign cloud providers, deploying **NVIDIA H100 and A100 GPUs** for multi-tenant AI training and inference workloads with data sovereignty by design.
- Integrated **Flower.ai federated learning** [🔗](#) framework into the OpenNebula Marketplace, enabling **distributed ML workloads** across sovereign edge environments for the IPCEI-CIS European project.
- Deployed **EuroCopilot** [🔗](#), a fully local **AI inference appliance** leveraging **llama.cpp** and **LiteLLM**, optimizing **load-balanced inference workloads** for CPU-only infrastructure to maximize resource utilization.

Technical Content Creator: Cloud, DevOps & AI

Remote

Medium (Freelance)

January 2025 – Present

- Grow the **Kubernetes Homelab Blog Series** [🔗](#) to 3,000+ readers, featured by the Cilium project. Produce technical articles about container orchestration, cloud-native and AI systems.

5G Field Verification (R&D) Engineer Internship

Espoo, Finland

Nokia

May – Nov. 2023

- Developed **Python automation** scripts for **5G New Radio (NR)** KPI analysis, boosting data processing efficiency by **~20%** across large-scale field datasets.
- Automated **Raspberry Pi** test workflows for 5G RAN field verification, reducing test cycle times by **15%**.

Projects

AgentBloc – AI Agent Team Designer

2026 – Present

- Build and open-source **AgentBloc** [🔗](#), an AI agent team designer for Claude Code. Users describe their workflow in conversation; AgentBloc designs the agent team, discovers integrations (APIs, MCP servers, browser automation), and generates all deployment artifacts with contract-first, phase-gated architecture.

Kubernetes Homelab – AI Infrastructure Lab

2024 – Present

- Maintain a production **K3s homelab** [🔗](#) (26 stars, 800+ commits) with GitOps (ArgoCD), Prometheus, Grafana, Longhorn, and Cilium. Hosts **AI inference workloads** (llama.cpp, LiteLLM) as a testbed for **MLOps experimentation**.

ClaudeClaw – AI Personal Agent Team

2024 – Present

- Built **ClaudeClaw**, an autonomous **AI agent team** orchestrated via Telegram bot interface: 30+ scheduled agents handling personal automation, health data tracking, and workflow coordination, running 24/7 on the homelab cluster.

Education

Double MSc. Computer, Communication & Information Sciences
Aalto University – EURECOM

France & Finland
Sept 2021 – June 2023

- Thesis: LoRaWAN Smart Agriculture System on Microsoft Azure (4/5).

BSc. Telecommunication Systems Engineering
UCAM University of Murcia

Murcia, Spain
Sept 2016 – June 2021

Speaking

- **FOSDEM 2026** (Brussels): Speaker on edge cloud deployment with OpenNebula on Raspberry Pi.
- **Foss-North 2026** (Gothenburg): Speaker on open-source cloud-edge infrastructure.
- **IPCEI-CIS GA 2026** (Rotterdam): Speaker on federated AI infrastructure with EuroCopilot/Virt8ra.

Technologies

AI/MLOps & Inference: llama.cpp, LiteLLM, Flower.ai (Federated Learning), Ray, Hugging Face, NVIDIA GPU ecosystem (H100, A100, H200/B200 familiarity), NemoClaw/OpenClaw integration, model serving, CPU inference optimization

AI Agents: Claude Code, MCP, multi-agent orchestration, prompt engineering, n8n

Containers & CI/CD: Docker, Kubernetes (K3s, RKE2, Cluster API), Helm, ArgoCD, Cilium, GitHub Actions, Trivy

IaC & Cloud: Terraform, Ansible — AWS, Azure, OpenNebula, OVHcloud, IONOS, Scaleway

Programming & Tools: Linux, Python — Prometheus, Grafana, Tailscale, Raspberry Pi, Home Assistant

Certifications: OpenNebula Certified Expert (2024), LFS258 Kubernetes Fundamentals (2024), Azure AZ-900 (2023)

Languages

Spanish (native) **English** C1 (Cambridge) **French** B2 (DELF)