

# Guido Maria Serra

Principal Infrastructure Engineer | AI/ML Systems & GPU Platforms

Berlin, Germany

+49 173 7349 795 | guido@whitehat.berlin | linkedin.com/in/zeph1ro | github.com/zeph

## Profile

---

Principal Infrastructure Engineer with ~30 years of experience building and operating large-scale distributed systems, from early Linux infrastructure to modern GPU-based Kubernetes platforms for AI/ML workloads. Specialized in stabilizing critical infrastructure, optimizing costs (e.g., \$300K/year savings at OLX), and coaching SRE/DevOps teams through operational challenges.

*Highlights:* Operated infrastructure at 55k-host scale (12 DCs); coached teams through production crises.

## Selected Engagements

---

### PlayStation Now — Global Infrastructure Operations

Operated multi-datacenter platform (55k hosts, 12 DCs); built Ceph storage, PKI-based SSH. Designed custom metrics DSL (TextX) integrated with Grafana. Gentoo maintainer for Ganeti/KVM. *Result:* Stabilized platform under global scaling pressure; enabled reliable launch of PlayStation Now.

### HelloFresh — Security & Infrastructure Scaling

Built security team 1→5 FTEs in 2 weeks. Rolled out vuln scanning, Elasticsearch SIEM, incident response. *Result:* SIEM operational; AWS/Cloudflare edge monitoring established.

### OLX Group — Infrastructure Cost & Reliability

Led APM strategy across 4 hubs. CDN/SSL/DNS vendor selection. NewRelic vendor manager (\$2M budget). Load testing and performance validation (custom tooling). *Result:* \$300K/year savings through FinOps; improved platform reliability.

## Professional Experience

---

**Security & DevOps Consultant** | DFND Security / Forward Earth / 3T Labs 2025  
Remote — Multiple clients

- Fastly (Signal Sciences) WAF and edge security deployments for EMEA clients.
- ISO 27001 certification project; Vanta auditor training and compliance automation.

**Software Engineer** | FinTech (stealth) 2025  
Tokenized financial systems (Python, TypeScript, microservices).

- Developed backend components for tokenized financial systems.
- Evaluated and applied AI-assisted coding workflows to accelerate iteration cycles.
- Performed load testing and reverse engineering of tools/libraries to validate system behavior.

**Head of SRE** | Octostar 2024–2025  
GPU Kubernetes platform for ML workloads; ClickHouse/MinIO data platform.

- Designed GPU-backed Kubernetes clusters for ML training/inference.
- Optimized cost and performance for ML workloads (mixed on-prem/cloud GPU nodes).
- Validated workload portability across local Kubernetes (kind/minikube) and cloud (Hetzner).

**DevOps Lead (Advisory)** | Contensi 2024–present  
CDKTF infrastructure automation; mentoring junior engineers.

**Staff SRE (Fractional)** | Prima Assicurazioni 2022–2023  
Coached 12-engineer SRE team out of operational deadlock; AWS multi-region rearchitecture (London, Barcelona, Milan).

**Cloud Architect** | ClovrLabs 2021–2023  
 Built 5-person security ops team; Kubernetes/GitOps platform (ArgoCD, Traefik, Datadog).

**Site Reliability & Security Engineer** | Max Planck Institute 2022–2024  
 Hybrid OpenStack/Kubernetes; 1.5PB storage; federated IAM (LDAP/SAML/OAuth2).

**Security / DevOps Engineer** | HelloFresh 2020–2021

**DevOps Engineer** | Mobivia / ATU 2018–2022  
 4-year modernization: Cloudflare/Traefik migrations, Keycloak SSO, ArgoCD GitOps.

**DevOps Team Lead** | Urban Mobility Int. (WeShare/VW) 2019  
 1-week pre-launch engagement; GCP Functions and Kubernetes.

**Start-up CTO** | SiWeGO (Trento, Italy) 2017–2019  
 Solution design for low-OPEX operations; algorithmic design of graph overlaps; suppliers management.

**Senior Performance Engineer** | OLX Group 2016–2018

**Site Reliability Engineer** | Sony PlayStation (Gaikai → PlayStation Now) 2013–2016

### Earlier Career (2008–2013)

---

**Rocket Internet** — Chief Troubleshooter & Senior Data Architect (2013) | Berlin  
**txtr (3M)** — Backend QA Manager (2011–2012) | Berlin  
**ProfitBricks (→ IONOS)** — Senior System Engineer (2011) | Berlin  
**Nokia Maps (Gate5 → HERE)** — Senior System Engineer (2010–2011) | Berlin  
**Vodafone Group** — QA Architect (2008–2010) | Düsseldorf

### Foundational Era (1996–2008)

---

Solution Architect for SMEs in Italy. Designed LDAP/VoIP/firewall solutions; released patches for GOsa<sup>2</sup>, PHP Seagull, GNU TLA. Clients included ISPs, banks (BPU Banca), and telcos —including the Stendhal digital library for Biblioteca Sormani (still online: digitami.it/stendhal). Stager under Alessandro Rubini for embedded Linux R&D. Lab Assistant at Politecnico di Milano (C/Linux socket programming, 2002). Active in FOSDEM since 2007.

### Core Technical Domains

---

<b>Distributed Systems &amp; Infrastructure</b>	Large-scale Linux, Ceph, distributed storage, fleet management
<b>Cloud &amp; Platform Engineering</b>	AWS, GCP, Kubernetes, Terraform, GitOps
<b>Observability &amp; Reliability</b>	Prometheus, Grafana, ELK, OpenTelemetry, SLO/SLI
<b>Security &amp; Compliance</b>	WAF, IAM/SSO, ISO 27001, SIEM
<b>AI/ML Workloads</b>	GPU environments, data pipelines, local/cloud portability
<b>Programming</b>	Go, TypeScript, PHP, SQL; hands-on, strongest in debugging and systems analysis
<b>Low-Level Systems</b>	C, embedded Linux, kernel-level performance

### Education & Credentials

---

**Bachelor of Engineering in Computer Science** (final year)

— Politecnico di Milano | 2000–2008, resumed 2025

**Training:** Red Hat Ceph Storage (CEPH125), NYC 2016; ISO 27001 / Vanta Auditor, 2025

### Additional

---

**Languages:** Italian (Native), English (Fluent), German (Working)

**Security:** DEFCON CTF finalist; iCTF World Champion; SOC GOON (2025)

**Open Source:** FSFE Fellow; co-founder of Politecnico Open \*nix Labs (POuL); wicd Debian package maintainer

**Licenses & Interests:** Driving (A, B); Offshore Sailing (racing cert., maritime radio); Skiing; Trekking