Gyanaranjan Pradhan

AI Infrastructure Engineer, DevOps & SRE

Delhi NCR, India · gyansemail@icloud.com · +91-9938373626 · iamgyan.com · linkedin.com/in/elord404

SUMMARY

I have been on call for production servers since **February 2019**, first at **ServerGuy**, then **DigitalOcean**, then **One.com**. The work is Linux, AWS, and hosting platforms, usually under an SLA and usually while something is already broken. I hold **99.99% uptime** on the platforms I manage and get priority incidents back **under 10 minutes**.

When open models got good enough to run locally, I started serving them myself rather than waiting for someone else to. **Qwen 3.6** and **Gemma 4** run on hardware I manage for anything that cannot leave the network, and I reach for **Amazon Bedrock** when a harder model is worth paying for. I built **PAKAD.AI** on the same idea. It turns a server's own telemetry into context, then writes the root cause of an incident in plain English along with the steps to fix it.

CORE SKILLS

AI & Models	Qwen 3.6, Gemma 4, Mistral, Anthropic Claude, Amazon Nova, DeepSeek, OCR (Tesseract, docTR)
Model Serving	Ollama, vLLM, LM Studio, llama.cpp, quantization, GPU sizing, batching, VRAM planning
AI Integration	MCP servers, RAG pipelines, multi-model routing, judge/consensus evaluation, Open WebUI, AnythingLLM, CMDB mapping
Cloud & Hosting	AWS, Amazon Bedrock, DigitalOcean, Plesk, cPanel, Proxmox, KVM/XEN
Observability	Prometheus, Grafana, Zabbix, VictoriaMetrics
Security & SOC	VAPT, Incident Response, Wazuh, OpenVAS, WireGuard, TLS, system hardening
Platform	Linux, Docker, CI/CD, Go, Node.js, PostgreSQL, Bash

PROFESSIONAL EXPERIENCE

AI Solutions Consultant, DevOps & SOC — Apr 2026 to Present

Independent Consulting, Remote

- Integrated AI models into support and ticketing platforms, resolving and self-verifying L1 and L2 incidents automatically and escalating L3 and critical tickets to engineers with diagnosis and remediation steps attached.
- Architected a multi-model triage pipeline in which three open models run independent RAG passes over logs, metrics, and parameters, with a judge model consolidating their outputs into a single verdict.
- Enforced read-only MCP access for all model tooling and CMDB-based resource mapping, so automated diagnosis carries no write path into production.
- Advise engineering teams on adopting AI inside existing security and compliance boundaries.
- Run VAPT engagements and lead incident response within SOC operations.

DevOps Engineer III — Apr 2025 to Jan 2026

Group.one (One.com), Remote, Denmark

- Operated scalable, secure infrastructure for large-scale hosting and cloud platforms, including Plesk estates.
- Rebuilt monitoring and alerting to surface faults before customer impact, supporting sustained 99.99% uptime.
- Resolved critical production incidents end to end, covering triage, root-cause analysis, and cross-team remediation.

Senior Cloud Engineer II Mar 2023 to Jan 2025

DigitalOcean, Remote, US

- Delivered L3 support for premium partners across the application, hosting, and platform layers.
- Executed migrations and performance and security remediation on live customer workloads.
- Authored RCA documentation and knowledge-base articles that enabled L1 and L2 to close repeat issues without escalation.

Senior Linux Server Engineer, Support & Operations Manager Feb 2019 to Mar 2023

ServerGuy, New Delhi, India

- Managed Linux and Windows hosting estates across WordPress, Magento, cPanel, and Plesk.
- Led patching, scaling, backups, migrations, and automation of recurring operational work.
- Mentored engineers and maintained SLA-bound delivery across the support organisation.

PROJECTS

PAKAD.AI, AI Infrastructure Engineer (SaaS) pakad.ai

Go, Node.js, PostgreSQL, VictoriaMetrics, Claude

- Converts live infrastructure telemetry into context and returns plain-English root cause with actionable remediation steps.
- Lightweight Go agent streams telemetry while agentless checks monitor every site; multi-tenant dashboard, NOC wall, and public API.
- Requires human approval before any recommended action reaches production.

Managed Wazuh SOC, Multi-Region Security Platform

Wazuh, nginx, WireGuard, OpenVAS, TLS

- Built HA Wazuh managers with a two-node indexer, nginx L4 and TLS edge, WireGuard inter-region tunnels, and controlled agent enrollment.
- Paired continuous OpenVAS scanning with VAPT engagements and documented incident-response runbooks.

Infrastructure & SOC Monitoring Panels

Prometheus, Zabbix, Grafana, Wazuh, OpenVAS, Kali

- Delivered availability and capacity-planning visibility across the estate, plus security telemetry and vulnerability validation workflows.

CERTIFICATIONS

AWS SysOps Administrator · AWS Cloud Practitioner · Red Hat RHCSA · Plesk Obsidian Professional · CompTIA Network+ · CompTIA A+ / ITF · Ethical Hacking Expert

EDUCATION

B.Sc. Information Technology, Centurion University of Technology & Management 2014 to 2017

CHNA, Computer Hardware & Network Administrator, IANT 2017 to 2019

LANGUAGES

English · Hindi · Odia · Russian