

Gunwant Jain

Infrastructure / Backend Engineer, open to relocation across Europe

+91 9717169269 • mail@wantguns.dev • github.com/wantguns • wantguns.dev

SKILLS

- **Programming Languages:** Go, Rust, C/C++, Python, Bash, Nix
- **Site Reliability:** Kubernetes (EKS, GKE, Kubeadm, K3S, ClusterAPI), Terraform, Ansible, ArgoCD, FluxCD, Helm, Kustomize, Prometheus, Loki, Grafana, Containers (Docker, OCI, LXC), AWS, GCP
- **Operating Systems:** Nix, Linux Kernel, AOSP (Android Open Source Project), Binutils, Cross-Compilation
- **Networking and Protocols:** Cilium, Kong/Nginx, Istio, HAProxy, Ambassador, Caddy, Traefik, BGP, DNS, IMAP, SMTP, TCP, UDP, ARP, WireGuard, SOCKS5, IPinIP
- **Datastores:** PostgreSQL, SQLite, MongoDB, Apache Kafka, S3/Minio, Redis, RabbitMQ, OpenZFS, BTRFS, Linux LVM

WORK EXPERIENCE

Software Engineer (2nd hire)

Apr 2025 — Present

[aion](#)

Remote

Second engineer at a stealth-mode GPU neocloud (H100/H200). Built all cloud infrastructure, CI/CD, and GitOps from scratch.

- Architected aion's **Managed Kubernetes** in Go - a highly secure Kubernetes platform bootstrapped via **Kubeadm** + cloud-init with no SSH or node agents. Customers retain sole cluster access unless they opt into deploying aion's agent.
- Built the Kubernetes Apps Marketplace powered by **FluxCD**, including an **OCI-on-S3** translation backend and **Cue-lang** + ReactJsonSchemaForm for dynamic app-specific configuration forms.
- Designed the full AWS network topology: 4 VPCs with segmented subnets (private, public, observability, database). Wrote all **Terraform** modules (EKS, RDS, ElastiCache, MSK, S3, ECR) with full environment isolation across 3 environments and 6 **EKS** clusters.
- Established **ArgoCD**-based GitOps with a hub-and-spoke model for multi-cluster control. Built all **GitHub Actions** CI/CD pipelines, including a custom bot to close the GitOps loop by auto-updating image tags on push.
- Set up **Tailscale** on clusters with subnet routing for developer access, while restricting PostgreSQL and Kafka access for non-SRE engineers via database-level controls and Kafka ACLs.
- Developed a GPU benchmarking methodology across cloud providers. Optimized InfiniBand bandwidth for **H100** clusters (up to 8x8) to match Nvidia's **465+ GB/s** limits. Created **Ansible** playbooks for multi-node tests (NVBandwidth, DCGM, GPU Burn).

Platform Engineer (SDE 2)

Jan 2024 — Mar 2025

[Ola Krutrim](#)

Bengaluru, India

Sole platform engineer for the **Embodied AI** team. Owned all infrastructure, GPU compute, datastores, and distributed microservices from the ground up.

- Bootstrapped **Kubernetes** clusters on OpenStack VMs using **Kubespray** and **HAProxy**. Implemented GitOps with **FluxCD** and GitHub Actions. Built observability with **Prometheus** and **Loki**, enabling rapid incident response.
- Provided virtually partitioned **Nvidia A100 GPUs** for ML engineers on **Kubeflow** using Nvidia GPU Operator and Multi-Instance GPU (MIG).
- Deployed and maintained a **25TB Minio tenant** on **Cinder CSI**-backed storage as the team's Data Lake, with periodic and incremental **Velero** backups for all persistent state.
- Designed **terabyte-scale, fault-tolerant** sensor-data ingestion infrastructure for edge devices across disparate networks. Developed a highly concurrent **Go** microservice using **MongoDB-based durable queues** for terabyte-scale transfers, with **Windmill** for workflow orchestration.
- Built reconciliation pipelines in Bash and Python for edge device data, using **SQLite** and **Systemd Timers** for coordination, packaged as **Debian packages** for deployment across multiple vehicles.
- Deployed Intel's **CVAT** annotation tool on Kubernetes, with **PostgreSQL tuning** and **disaster recovery**, supporting **100+ concurrent annotators**.

Infrastructure Contractor

Tensorfuse (YC W24)

Sep 2024 — Dec 2024

Bengaluru, India

Advised on core architecture for running containerized GPU workloads at scale.

- Benchmarked **Karpenter's** node claim lifecycle on AWS and onboarded **Lambda Labs** as a cloud provider, developing K3S bootstrapping scripts and a custom Karpenter Provider.
- Improved **Buildkit**-based container build times by **50%** by identifying cacheable steps and integrating **Nydus-Snapshotter** for optimized image distribution.

Site Reliability Engineer

Dyte (YC W21, now Cloudflare)

Jun 2022 — Jan 2024

Remote

- Maintained 20+ Kubernetes clusters on AWS/GCP. Led production multi-region upgrades (1.24 to 1.27) using External-DNS for A/B cluster cutover, and reduced cloud costs by **30%** via **Cast AI** binpacking and spot instances.
- Migrated recording backend from ECS to GKE, scaling capacity to **8192 parallel recordings**. Architected cross-cloud authentication using **Workload Identity Federation** and OIDC.
- Built active monitoring probes (Lambda/EventBridge) that report directly to Dyte's public status page, decreasing mean time to detection to **< 1 minute**.
- Designed and brought up the company's staging multi-region environment from scratch. Wrote **Terraform** modules for cross-region VPC peering on AWS.

Research Intern

Ecole Polytechnique de Montreal

Jun 2022 — Sep 2022

Montreal, Canada

- Developed **Ansible** playbooks to automate the provisioning of Multi-Cluster **Hadoop** on **Google Cloud Platform**.
- Created **InjectErr**, a Python framework for injecting faults during **HiBench** workload execution on Hadoop and collecting logs for analysis.

Google Summer of Code Developer

Gentoo Linux

May 2020 — Sep 2020

Remote

- Enhanced **SharkBaitOS**, a **Gentoo Linux** userspace layer between Linux and Android, extending compatibility to System-as-Root devices (Android 9+).
- Wrote a custom C++ initramfs to boot Gentoo off the Linux Kernel and OpenRC scripts for Android LXC containers. Expanded **AOSP's LLVM** toolchain with an **AArch64** native build target.

VOLUNTARY EXPERIENCE

Self-Hosting

- Managing a multi-continent server mesh (Hetzner, Oracle, bare-metal) for 6+ years. All state declared in a [public Nix flake](#) with encrypted secrets, including a custom **WireGuard** mesh module with ACLs.
- Established a **Cilium**-based K3S cluster utilizing WireGuard state for pod subnets, enabling control-plane participation for nodes without public IPs. Set up a mirrored zpool with **OpenZFS** for disk reliability.
- Implemented **ArgoCD** GitOps for 20+ services (Nextcloud, Bitwarden, Forgejo) and transitioned ingress management from Traefik/Caddy to **Cilium Gateway API** for Kubernetes.

Co-Founder, Flowdrive

Oct 2020 — Present

- Co-founded cross-platform ADAS software for Android/SBCs serving **275+ vehicles**, optimizing **Android OS** with low-overhead primitives (**LMDB**, **Cap'n Proto**, **Zeromq**).
- Built the full SaaS backend: a **Flask** API, **SvelteKit** webapp, and a **Go** microservice for converting telemetry into **HLS video streams**.
- Developed **Flowinit**, a Python-based process manager for reliable service coordination on edge devices.

ACHIEVEMENTS

[Globalink Research Intern](#), [Ecole Polytechnique de Montreal](#), MITACS

2022

[National Winner \(Information Security track\)](#), [GRiD 3.0 Hackathon](#), Flipkart

2021

[GSoC Scholar](#), [Gentoo Foundation](#), Google

2020