

Pranav Bhendawade

+1(408)690-7643 | [✉ pranavbhendawade@gmail.com](mailto:pranavbhendawade@gmail.com) | [🌐 LinkedIn](#) | [🐙 GitHub](#) | [📍 Santa Clara, CA](#)

SUMMARY

Software Engineer with 5+ years architecting and operating reliable, scalable distributed systems. Expertise in Kubernetes on AWS (EKS) and GCP (GKE), IaC, GitOps, and service mesh networking. Owned services end-to-end: production code, on-call rotations, incident response, and AI-powered operational tooling. Expanding into AI/ML infrastructure - GPU scheduling, vLLM serving, and agentic observability tooling.

EXPERIENCE

Automation Anywhere Inc.

August 2024 - Present

Software Engineer, Cloud Infrastructure

San Jose, California

- Engineered an Istio service mesh across 90+ production Kubernetes clusters (EKS/GKE) with mTLS, SNI-based routing, and east-west gateways, securing 100% of cross-cluster traffic including AI/ML inference workloads.
- Architected a two-tier IaC architecture with a GitOps-driven control plane for environment lifecycle and a Terraform infra plane for AWS/GCP provisioning, reducing onboarding time by 70% and eliminating manual configuration drift.
- Standardized Helm charts and ArgoCD GitOps pipelines across 90+ clusters, cutting deployment cycle time by 40% and manual configuration errors by 60%.
- Diagnosed production issues (network partitions, DNS failures, pod scheduling, Envoy misconfigs), reducing resolution time by 35% while optimizing fleet-wide autoscaling for 30-40% cost reduction.
- Developed a Python-based incident response agent using FastAPI and MCP that autonomously correlates logs, metrics, and K8s events, reducing MTTD by 60% and cutting manual triage across on-call rotations.

Activision Blizzard

June 2023 - September 2023

Software Engineer Intern, Linux Systems

Irvine, California

- Constructed an infrastructure configuration discovery service in Python that aggregated real-time resource data across multiple sources, providing accurate platform visibility and reducing deployment misconfigurations by 40%.
- Streamlined provisioning of service by containerizing and deploying on Kubernetes with Terraform-provisioned infrastructure on OpenStack, reducing deployment time by 75% and provisioning time by 90%.

Tata Consultancy Services

July 2019 - August 2022

Software Engineer, Cloud Infrastructure

Bangalore, India

- Designed AWS VPC networking for EKS across 3 environments with CNI-aware subnet sizing, pod-level security groups, NACLs, and VPC peering for cross-environment service discovery supporting 100+ node clusters.
- Spearheaded on-call rotation and incident escalation for production infrastructure, established runbooks and drove root-cause analysis, contributing to 45% MTTR reduction using Prometheus/Grafana alerting and log analysis.
- Automated AWS resource inventory and cost reporting across 20+ regions using Python, reducing manual overhead by 80% with real-time utilization dashboards.
- Administered PostgreSQL RDS and ElastiCache Redis serving 15+ production services, tuned parameters and connection pooling for 99.9% uptime. Triage Linux incidents (OOM kills, network stack, systemd) via strace and tcpdump.

CloudYuga Technologies Pvt. Ltd.

December 2018 - May 2019

Software Engineer Intern

Bangalore, India

- Implemented reusable Terraform modules for GKE networking, compute, and storage with GitHub Actions CI/CD checks, reducing infrastructure provisioning time from days to under 2 hours.
- Launched an EdTech platform on GKE using Kubernetes manifests with RBAC, ConfigMaps, Secrets, Ingress, and PVCs, reducing rollout failures by 35%.

SKILLS

Languages	Python (Proficient), Shell/Bash, Go, Java, SQL
Platform & Orchestration	Kubernetes, Helm, ArgoCD, Karpenter, PDB/HPA/VPA, KEDA, Argo Workflows, Docker
Infrastructure & Cloud	Terraform, AWS (EKS, EC2, VPC, IAM, RDS, ElastiCache, ALB), GCP (GKE), GitOps
Observability & Networking	Prometheus, Grafana, OpenTelemetry, ELK Stack, Istio, Envoy, mTLS, TCP/IP, DNS
Data & AI/ML	Redis, PostgreSQL, DynamoDB, GPU orchestration, vLLM, MCP servers, agentic AI
Certifications	Certified Kubernetes Admin (CKA), AWS Solutions Architect, HashiCorp Terraform Associate

EDUCATION

Master's in Computer Science (MS CS)

September 2022 - May 2024

Santa Clara University

Santa Clara, California

Bachelor's in Computer Science (BS CS)

July 2015 - May 2019

Shivaji University

Kolhapur, India

PROJECTS

GPU-Aware Autoscaling Platform on EKS | Kubernetes, Karpenter, KEDA, Prometheus, NVIDIA DCGM, Helm, Terraform

- Orchestrated GPU workload scheduling on EKS using Karpenter and KEDA for autoscaling on GPU utilization and queue depth, reducing GPU idle time by 40% and instrumented Prometheus/Grafana dashboards.

Kubernetes Troubleshooting Agent | Python, FastAPI, MCP, Slack

- Devised an agentic system triggered by webhooks that uses a FastAPI agent with a custom MCP server to query cluster metrics, logs, and events, cutting incident triage time by 50% and posting root-cause analysis to Slack.