

Shubham Singh

(+91)-9711778196 | New Delhi, IN | Email - shubham.singh@outlook.com | LinkedIn - [shubham-singh](#)

Member of Technical Staff

Site Reliability Engineering | Platform Engineering | DevOps/DevSecOps | | AI/ML Ops

A pragmatic SRE/Platform Engineer with **6+ years** of experience designing, scaling, and running distributed systems across fast-moving startups and enterprise environments. Builds resilient infrastructure, reduces operational load through automation, and improves reliability through data-driven engineering. Possesses deep hands-on expertise with Kubernetes, AWS services, CI/CD, observability stacks, and high-traffic backend systems. A strong application-engineering background (**Node.js, Python, Java**) enables delivery of production-grade tooling, automation, and well-designed services. Experienced in incident response, performance optimization, and implementing fault-tolerant architectures that improve availability, latency, and overall system robustness.

WORK EXPERIENCE

Salesforce, Inc. – Hyderabad, India.

Oct 2024 – Present

Member of Technical Staff – (Hybrid/Remote)

Empowering platform resilience and scale - owning software development, CI/CD pipeline optimisation, infrastructure automation, proactive monitoring & alerting, security hardening, performance tuning, incident diagnosis, and cross-team collaboration.

- Managing & maintaining **800+ Kubernetes** clusters across hybrid cloud(AWS/Ali/GCP) and on-prem(SAM 1P) environments (RHEL 9.6+), ensuring 99.99% service availability.
- Streamlined CI/CD pipelines using **Terraform, Spinnaker, & ArgoCD**, achieving a **30% reduction** in deployment cycles.
- Enhanced observability and incident response by implementing **Grafana** monitoring dashboards and integrating with **Splunk**, reducing mean time to detect (**MTTD**) issues by **40%**.
- Oversaw critical infrastructure tasks, including **regular OS patching, kernel upgrades, and IPVS module rollouts** to optimize load balancing for on-prem systems.
- Leveraged **AI agentic framework (Warden AI Ops)** framework to significantly enhance automated incident detection & remediation capabilities, achieving a **30% reduction** in incident resolution times, empowering tenants to innovate rapidly.
- Created AI agents like **weekly SQR report**; K8s Agent troubleshooting that uses **K8sGPT** to solve & auto heal the k8s related issues; **Pager Duty alert automation** to describe the alert & suggest remediation.
- Created Kubernetes Operators for automated self-healing workflows, **reducing manual intervention by approximately 40%**.
- Actively engaged in cross-team collaboration, contributing to **architecture design**, security assessments, quarterly reliability initiatives, and performance tuning.
- Building **Customer Engagement Framework (CEF)**, which remediates the Pager Duty alerts by intercepting the metrics and raising the same back to our tenants from the Argus queries, resulting in reducing the ops toil and burden for **unnecessary alerts by 30%**

Airtel International – Gurgaon, India.

May 2022 – Sep 2024

Site Reliability Engineer – (Hybrid/Remote)

Keeping the ship sailing! Responsible for software development, release, deployments, monitoring, automation(CI/CD), security, diagnosis for optimal performance, and supporting the IT infrastructure environment

- Achieved **Einstein Award** in the early span of joining the org & showcased the bar of excellent performance
- Managed and mentored a team of **3 engineers**, driving delivery quality and skill development.
- Rolled out backbone platforms (NPO, KYC 2.0, E-Sim, CLM) across **14 countries**, supporting large-scale telecom provisioning systems.
- Built internal AI/ML Ops tooling for anomaly detection, triage, and resilience improvements.
- Designed and developed the **Affair System** using MEAN, Java, and Python - automating workflows, health checks, reconciliation, and real-time dashboards; reduced functional dependency by **70%**.
- Orchestrated **200+ Micro-services** & their **CI/CD pipelines** using **K8s, Rancher, Flux & Jenkins**
- **Defined and monitored SLOs for critical services, improving reliability standards across org teams.**
- Tuned runtime **DB queries (SQL and NoSQL)** and generating custom metrics from **Grafana & Kibana**

Amway India – Gurgaon, India.

Jun 2020 – Apr 2022

DevSecOps Engineer – (Hybrid) (Jun 2021 – Apr 2022)

Responsible for development, deployments, and engaged with multiple activities to provide user a frictionless shopping experience. Worked as a Technical SEO Manager, Technical DevSecOps Engineer across various functional teams and as a Product Owner & Engineer to mitigate the production support incidents.

- Supported e-commerce, order-management, and production systems through development, automation, monitoring, and incident management.
- Delivered **10+ end-to-end order journeys** and built performance-testing scripts using headless Chrome, Node.js, InfluxDB, AWS Lambda, Jenkins, and Grafana.
- Built **ACUTE**, a centralized reconciliation, sales, tax, and incident-automation platform; reduced monthly and annual production incidents by **100+ and 400+** respectively.
- Automated key metrics, log processing, and alerting workflows to improve operational insight.
- Gained expertise in **Amstack**, Amway's internal CI/CD automation system, reducing dependency on external service providers.

DevOps Engineer – (Remote) (Jun 2020 – May 2021)

- Promoted to DevSecOps within a year due to high-impact contributions.
- Received **Flame & Ignite Awards** three times for outstanding performance.
- Reduced Docker build times by **40%** by optimizing pipelines and fixing AWS integration issues.
- Built new CI/CD pipelines and reconciled large data sets across **10+ microservices**.

Octonius Inc. – Palo Alto, CA, USA; Barcelona, Spain.

Jun 2018 – Aug 2021

Chief Product Officer – (Hybrid) (Jul 2019 – Aug 2021)

Built the entire product from Scratch, started as a Software Engineer, and paved my way up to as a founding member. Worked on the Core System of the application (MEAN Stack), built and maintained the CI/CD pipelines, managed the whole database, and created the entire infrastructure to ship the product to our direct customers.

- Fostered & collaborated with **CEO**; Spearheaded prototyping & implementation of the **15+ services**
- Successfully architecture and led the product development to deliver a highly scalable & maintainable app
- Effectively refactored the existing & delivered an app that was **40% faster, 50% lighter, and 30% robust**
- Recognised twice as the **Best of Palo Alto Awards 2020** and **TOP REMOTE WORK TECH** by **CIO Reviews**

Front-End Chapter Lead, Software Engineering – (Hybrid) (Dec 2018 – Jun 2019)

- Scrutinised with the founders and directly managed **10+** people in the product development team
- Containerised the application for deployment using **Docker Swarm** and slashed the dependency by **80%**
- Prepared the **official docs** for all **15 services** their corresponding contracts, modules & **150+ REST APIs**
- Built Folio – a real-time collaboration document system using ShareDB, Web Sockets, and Quill.JS supporting up to **1000 concurrent users**.

Software Engineer – (Remote) (Jun 2018 – Nov 2018)

- Implemented new features in the app using **Angular 8, Node.js, and MongoDB**
- Integrated **Google Drive Cloud** service in the app
- Ensured the technical feasibility of **UI/UX** designs & optimised app for **scalability** and **maximum speed**
- Created the **Kanban Boards** (Trello like functionality for tasks management) in the app

EDUCATION

The University of Delhi, Netaji Subhas Institute of Technology (NSIT) – New Delhi, India

Bachelor of Engineering (B.E.) – Computers Science, (Aug 2016 – Jun 2020) – **Division 1st**

- Chairperson of the **IEEE-NSIT- CS Chapter** (<https://www.ieeensut.com/>)
- Head of Operations of **HACKEAM'18** & **TEAMC'18** which were a **domestic hackathon** and an **international conference** respectively
- Core-committee member of **Rotaract - NSIT**, where we organised various events based on menstrual hygiene, plantation & cleanliness drive, cancer awareness camps, sandwich making and distribution drive, aiding marginalised and vulnerable groups, donating books, visiting and serving NGOs

PUBLICATIONS

MorCVD | A Unified Database using Advance Machine Learning

2019

Nature Journal – Scientific Reports

This database is a curated resource that combines 23,377 protein interactions between the human host and 432 unique pathogens involved in CVDs in a single intuitive web application. *It has an impact factor of 42.78* (<https://www.nature.com/articles/s41598-019-40704-5>)

This project is related to the conversion of facial description to the possible matches of faces using GANs and machine learning (<https://ieeexplore.ieee.org/abstract/document/9557583>)

AWARDS AND HONOURS

- Achieved the title of **Country Winners** at **Reboot the Earth** presented by the United Nations and represented **India** in the **United Nations** at UN Headquarters, New York to promote women empowerment and how women are affected due to climate change and how it can also affect other 150 Million people
- Featured as the **Inventor of the Month** in the Month of August 2018 by a Japan Based Company - **Progate**
- Supervised as the **Product Development Engineer** at **CareZilla** and accomplished the MVP, which **connects the caregivers and caretakers** (patients - where caregiving is most used to address impairments related to old age, disability, a disease, or a mental disorder) over the **intuitive PWA**
- Taught **15+ children** as a Technical Mentor at **Mentors without Borders** from Middle East and Africa
- Recognised as the **Youth Climate Summit Representative** at **United Nations** and worked on Projects:
 - Aimed to provide **employment to women**; Connect women with diverse **skill sets & STEAM skills**
 - Empower women to **work for climate, child protection**, impact of violence, exploitation, & abuse
 - Decrease the number of **Greenhouse Gases in the environment** – to address **UN Global SDGs**

PROFESSIONAL SKILLS

- **Client-Side / Front End** – HTML; CSS; TypeScript / JavaScript; Angular; React.js; Electron.js; Svelte
- **Server-Side / Back End** – C#; Java; Python; Node.js; TypeScript; Go (Golang)
- **Databases** – MySQL; MongoDB; SQL Server; PostgreSQL; Redis; InfluxDB; Prometheus (TSDB)
- **Programming Languages** – TypeScript / JavaScript; Java; Python; C / C++; Go (Golang)
- **Cloud Technologies** – AWS CloudFormation; Terraform (modular setups, blue/green deployments, for_each, dynamic blocks); AWS CDK; Crossplane (multi-cloud provisioning); IAM; Cognito; Security Groups; NACL; EC2; Lambda; App Runner; Auto Scaling; S3; EBS; EFS; CloudFront; Route 53; Amazon EKS; AWS Systems Manager (SSM); AWS Secrets Manager;
- **Monitoring & Alerting** – Dynatrace; Prometheus (custom exporters, AI-driven alerting); Grafana; Kibana (ELK); CloudWatch; Splunk (Kubernetes-integrated); Loki
- **Security** – GitHub Actions; SonarQube; GitLab CI; IAM; AWS Security Hub; GuardDuty; AWS Shield; WAF; Amazon Inspector; Puppet (node hardening); OPA/Gatekeeper; HashiCorp Vault
- **Logs Management** – Logstash; Fluentd; Splunk; CloudTrail; CloudWatch Logs; Loki; Vector
- **CI/CD & Config Management** – GitLab CI/CD; Jenkins; Bitbucket Pipelines; Travis CI; Flux (GitOps); Spinnaker (canary/rollback); ArgoCD (GitOps pipelines); Helm; Kustomize
- **Search Platforms** – Apache Solr; Elasticsearch (ELK Stack); OpenSearch
- **Messaging Queues / Pub-Sub** – Amazon SQS; RabbitMQ; Kafka; Apache Pulsar; Redis Streams; MQTT
- **Linux / Unix** – File Systems; DNS; Networking; System Calls; Kernel Internals; Performance Tuning for High-Scale Kubernetes; System Troubleshooting; Logging; eBPF; iptables/nftables

- **DevOps / Containerization** – Kubernetes (Operators, CRDs, multi-tenant 1,000+ clusters); Rancher; MinIO; KEDA; Flux; Docker Compose / Swarm; Istio (Service Mesh); GitOps best practices; Python / Go-based self-healing automation; Buildah; Podman; Containerd
- **AI / ML / Agentic Systems** – AI/ML Ops (TensorFlow, PyTorch, scikit-learn); LangChain; OpenAI APIs; Amazon Bedrock; Hugging Face Transformers; Weaviate; Pinecone; Vector Databases; Prompt Engineering; Autonomous Agents; RAG Pipelines; Agentic Framework Design; AI-driven Monitoring and Remediation; Natural Language Interfaces; LLM Orchestration; Generative AI Tools (CodeWhisperer, Copilot, TabNine, Cursor, CodeGenie)
- **Miscellaneous** – Socket Programming; Nginx; Tomcat; GraphQL; Computer Networking; Operating Systems