SYED HUZAIFA BIN AFZAL

AI Engineer | Secure LLM Deployment | AI Platform Engineering | Cloud Infrastructure

Tacoma, WA | +1 (206) 750-5956 | huzaifaafzal10@gmail.com | linkedin.com/in/syed-huzaifa-bin-afzal

PROFESSIONAL SUMMARY

AI-focused cloud infrastructure and cybersecurity professional with 6+ years of experience building secure cloud platforms, automation workflows, CI/CD pipelines, and production-ready infrastructure. Currently completing a Master of Cybersecurity & Leadership at the University of Washington - Expected June 2026, with applied research in secure enterprise generative AI, AI governance, Shadow AI risk, role-based access control, model control, and data residency. Hands-on experience deploying and evaluating self-hosted LLM platforms using Ollama, Open WebUI, Docker, and Windows Server, with strong foundations in Python, cloud engineering, observability, access control, and responsible AI adoption.

CORE TECHNICAL SKILLS

AI / LLM Platforms: Ollama, Open WebUI, local LLM deployment, secure LLM operations, model allow-listing, model versioning concepts, RBAC for AI platforms, evaluation of self-hosted model behavior, AI governance

AI Security & Governance: NIST AI RMF, ISO/IEC 27001:2022 control mapping, Shadow AI research, data residency, human oversight concepts, separation of duties, threat modeling, enterprise AI risk analysis

Programming & Prototyping: Python, Bash, PowerShell, JavaScript, C++, SQL/analytics exposure through Redshift/Superset, API Gateway and Lambda automation

Cloud & Platform: AWS primary production experience; Azure working knowledge; GCP working knowledge; Oracle Cloud; Docker; Kubernetes exposure; Windows Server; Linux; Terraform; CloudFormation

DevOps & MLOps-Adjacent Practices: CI/CD workflows, GitHub Actions, GitLab CI/CD, Jenkins, deployment automation, rollback planning, runbooks, platform administration, monitoring, post-deployment validation

Observability & Data Infrastructure: CloudWatch, Datadog, OpenSearch, S3 Inventory, Redshift, Superset, dashboards, logs, metrics, alerting, incident triage, storage analytics

AI RESEARCH & PLATFORM PROJECTS

Secure Enterprise Generative AI Platform / GovernAI | University of Washington June 2025 - Present

- Co-authored accepted IEEE SVCC 2026 paper: "Are You Aware of Shadow AI? GovernAI for Addressing Emerging Risks."
- Designed and deployed an on-premises generative AI platform using Ollama, Open WebUI, Docker, and Windows Server to support enterprise data residency and governed AI adoption.
- Implemented a three-role RBAC model for Standard User, AI Platform Admin, and System Administrator functions, supporting separation of duties and controlled AI platform administration.
- Conducted a 70-respondent workplace AI survey across public and private sector users: 71% reported using unsanctioned public AI tools at work, while 97% indicated they would use a governed organization-provided alternative.
- Validated platform network behavior using host-based firewall rules and Wireshark traffic analysis, confirming no unintended outbound communication during normal operation.
- Benchmarked the platform against ChatGPT, Microsoft 365 Copilot, and Gemini Pro across drafting, summarization, and code-generation workflows.
- Mapped implemented controls to the NIST AI Risk Management Framework and ISO/IEC 27001:2022 Annex A, connecting technical AI deployment choices with governance and security requirements.

AI Communication & Executive Stakeholder Reporting | University of Washington September 2025 - Present

- Serve as a part-time student news reporter covering campus news, student life, and developments in IT and AI affecting the university community.
- Interview executives and university stakeholders, including podcast conversations with executives, and translate technical AI topics into clear, accessible public-facing writing.
- Published reporting on UW Purple and student collaboration with UW-IT around next-generation AI; active participant in UW Purple Monthly Training and the UW AI Community of Practice.

PROFESSIONAL EXPERIENCE

DevOps Engineer | Harri (contracted by Kalam 4 Solutions) March 2022 - May 2025

Production AWS platform supporting enterprise SaaS customers.

- Built secure and reliable cloud infrastructure using Terraform, AWS, CI/CD workflows, monitoring tools, and automation patterns that support production-scale SaaS operations.
- Designed reusable Terraform modules for API Gateway, CloudFront, Route 53, OpenSearch, and Redshift, reducing environment provisioning time by about 40% and improving deployment consistency.

- Implemented IAM/RBAC and SSO patterns, refined least-privilege access, and centralized operational access through AWS Systems Manager to strengthen auditability and access governance.
- Built S3 Inventory, Redshift, and Superset infrastructure analytics to surface usage patterns, storage growth, and cost drivers for SRE and engineering teams.
- Improved reliability through Redis upgrades, OpenSearch cross-cluster replication, disaster recovery runbooks, multi-region readiness, monitoring improvements, and incident follow-up.
- Reduced AWS cloud spend by about 30% through Trusted Advisor, Cost Optimization Hub, Reserved Instance planning, rightsizing, and FinOps practices.
- Mentored engineers and contributed to infrastructure design reviews, Terraform refactoring, release workflow improvements, and secure platform operations.

Cloud Infrastructure / Software Engineer | Visionet Systems (contracted by Systems Limited) July 2019 - March 2022

Cloud automation, data platform, and infrastructure modernization for enterprise customers.

- Automated AWS operations using Systems Manager, Lambda, Python, Bash, and PowerShell for patching, scheduling, post-deployment validation, and routine cloud maintenance.
- Built Jenkins CI/CD pipelines for Terraform-based AWS EMR deployments, including cluster creation/destruction workflows, tagging validation, state handling, and end-user notifications.
- Supported Spark-as-a-Service migration from EC2/HDP to AWS EMR, applying Terraform, CloudWatch, Docker, EC2, AWS networking, and autoscaling patterns for big-data processing workloads.
- Supported application modernization from EC2 to ECS using Docker, ECR, Fargate, ALB, Route 53 service discovery, CloudWatch, Terraform, and containerized microservices patterns.
- Integrated CloudWatch and Datadog monitoring to improve real-time visibility, alerting, and incident response for cloud-hosted applications.
- Worked directly with client teams in Agile environments to gather requirements, present status, troubleshoot issues, and deliver infrastructure automation improvements.

AI-RELEVANT STRENGTHS FOR GENERIC APPLICATIONS

- Best-fit roles: AI Engineer, AI DevOps Engineer, AI Infrastructure Engineer, AI Security Engineer, Secure GenAI Platform Engineer, AI Governance Engineer, and MLOps-adjacent Platform Engineer.
- Strongest AI value: secure deployment and operation of AI platforms, enterprise GenAI governance, AI risk controls, platform administration, cloud infrastructure, and production reliability.
- Growth areas to avoid overstating: deep model training, PyTorch/TensorFlow production research, large-scale VLM/diffusion training, vector database ownership, and fine-tuning unless supported by a specific project.

EDUCATION

University of Washington - Master of Cybersecurity & Leadership - Expected June 2026

GPA: 3.99 | Beta Gamma Sigma Business Honor Society | Upsilon Pi Epsilon Honor Society

Ghulam Ishaq Khan Institute of Engineering Sciences and Technology - B.S. in Computer Science, 2015-2019

CERTIFICATIONS & HONORS

- AWS Certified Solutions Architect - Associate
- Oracle Cloud Infrastructure Architect - Professional; OCI Architect Associate; OCI Developer Associate; OCI Foundations Associate
- Foundations for Cybersecurity Analytics - University of Washington
- Top Performer - Harri (2023, 2024); Best Team Award - Systems Limited (2020)