

Gaurav Lahoti

Senior Cloud & AI Architect

✉ gaurav.lahoti25@gmail.com 📞 +91-9148448463 📍 Gurugram 🌐 <https://www.linkedin.com/in/glahoti/>

Summary

11+ years of engineering leadership specializing in Distributed Systems and Cloud-Native Architecture (GCP/AWS). Driven by a passion for technical excellence, maintaining a razor-sharp focus on the rapidly evolving AI landscape. Leveraging a decade of enterprise integration expertise (Apigee, Microservices, Event-Driven Architecture) to design and deploy production-grade Agentic AI systems. Currently bridging the gap between legacy enterprise stacks and modern AI strategy within Deloitte's Engineering practice. Expert in architecting autonomous workflows using LangGraph, Google ADK, and MCP, transforming traditional middleware patterns into intelligent, self-governing operations that ensure deterministic execution for LLMs.

Work History

Deloitte Manager	Gurugram Sep '19 - Present
EY Consultant	Bengaluru Jun '18 - Aug '19
Accenture Application Development Senior Analyst	Hyderabad Apr '15 - Jun '18

Projects

Client: Fiber Broadband Service Provider (US) June '23 - Present

- Led a distributed engineering squad of 20+ developers to deliver a cloud-native integration fabric on GCP, successfully migrating the client from a monolithic legacy stack to a decoupled, event-driven microservices architecture.
- Architected a reusable Resiliency & Observability Framework on GCP, standardizing error handling for 75+ microservices. This initiative reduced technical debt by 40% and cut Mean Time to Resolution (MTTR) by 2x.
- Engineered the enterprise adoption of Google's Application Integration (IPaaS), establishing the firm's first reference architecture for event-driven workflows.
- Secured stakeholder buy-in for a multi-year digital transformation roadmap by demonstrating the feasibility of 10+ complex API patterns, directly influencing the client's decision to greenlight the implementation phase.
- Partnered with executive leadership to define the Phase 2 technical roadmap, translating complex architectural requirements into precise resource estimation and capacity planning models.

Client: Leading Data & SSD Solutions Provider (US) Jun '21 - May '23

- Directed a cross-functional engineering team of 20+ developers across 3 time zones, delivering a net-new integration landscape from ground zero to production within aggressive M&A timelines
- Architected a scalable API Catalog & Governance Model, streamlining the delivery of 75+ critical integrations. This standardization reduced development cycle time by 20% and minimized code redundancy.
- Designed the "Lead-to-Cash" data fabric, orchestrating complex data flows between CRM and ERP systems that enabled real-time inventory visibility and streamlined order processing for high-volume semiconductor sales.
- Instituted Engineering Best Practices (CI/CD), mentoring 5 new hires on cloud-native development standards, which accelerated the team's velocity and ensured consistent delivery quality.

Client: Data Analytics and Consumer Intelligence (US) Sep '19 - May '21

- Guided an 8-person development squad to engineer the critical "Quote-to-Cash" pipeline, integrating Salesforce CPQ with NetSuite ZAB. This automation eliminated manual data entry errors and reduced order processing time by 30%.
- Spearheaded R&D for NetSuite ZAB (Zone Advanced Billing) integration, establishing a firm-wide benchmark for subscription billing architectures that was later adopted as a standard reference pattern within Deloitte.
- Diagnosed and resolved critical performance bottlenecks in the legacy middleware layer, re-architecting high-volume interfaces to handle 500 TPS peak loads while maintaining sub-200ms latency.
- Enforced Code Quality & Design Standards acting as the primary code reviewer; authored the project's comprehensive "Integration Design Principles" guide to ensure maintainability and scalability of the codebase.

Client: Food Processing Giant (Multi-Regions) Apr '15 - June '18

- Formulated and implemented enterprise-grade SOA patterns using Oracle SOA Suite, OSB, and BPM, facilitating real-time data exchange across a distributed global supply chain.
- Developed high-availability interfaces for critical financial reporting (ODI), ensuring 100% data consistency between legacy mainframes and the new ERP landscape.
- Managed the "Hypercare" Stability Phase, utilizing root cause analysis (RCA) to permanently fix production incidents, resulting in a smooth transition from UAT to stable business-as-usual operations.
- Optimized the SDLC by introducing Test Driven Development (TDD) methodologies, reducing defect leakage into production by 15%.

Skills

- Agentic Frameworks: LangChain, LangGraph, Google ADK
- LLMs & Model Garden: Gemini 3 Pro/Flash, Claude Haiku 4.5, GPT-4.1, GCP Vertex AI, AWS Bedrock
- AI Protocols: Model Context Protocol (MCP), Agent-to-Agent (A2A), OpenAPI (OAS) for Tool Use
- RAG & Vector Data: Pinecone, Embeddings (Gecko/Titan).
- Cloud Native: GCP (Cloud Run, Pub/Sub, Firestore, BigQuery, VertexAI, Cloud Storage, Gemini Enterprise), AWS (Lambda, EventBridge, S3, SageMaker), Kubernetes, terraform.
- Security: OAuth2, IAM, Zero Trust Architecture, DLP (Data Loss Prevention)
- Integration & API: Apigee X, Microservices, Boomi, Event-Driven Architecture (EDA), Oracle SOA/OSB, ODI
- Languages & Tools: Python (Pydantic, FastAPI), SQL, Cursor, VS Code

Certifications

AWS Certified Machine Learning Specialty	AWS Certified Machine Learning Eng. Associate	AWS Certified AI Practitioner
Microsoft Azure AI Fundamentals	Google Cloud Certified Generative AI Leader	Google Cloud Digital Leader
Google Cloud Professional Security Engineer	Google Cloud Associate Cloud Engineer	

Strategic Leadership & Innovation

- Pioneered the practice's adoption of Agentic AI by authoring foundational whitepapers and POVs on operationalizing Google & AWS Agent engines, establishing the technical playbook that guided practice-wide upskilling and delivery standards.
- Steered technical pursuits for 25+ global clients (US, Nordics, India), achieving a 60% win rate by pivoting proposals from "Staff Augmentation" to "Outcome-Based Delivery".
- Modernized commercial estimation by quantifying AI efficiency across the full SDLC (Code, QA, Migration Scripts), enabling 20-30% faster delivery timelines and securing greenfield wins via competitive cost structuring.