

DEO SHANKAR

Multi-Cloud Architect

CONTACT

jha.deo1771@gmail.com

+91-9891698505

Noida, India

deos.dev

LinkedIn

GitHub

TECHNICAL SKILLS

Cloud Platforms

AWS (EKS, SageMaker, Bedrock, Lambda, DynamoDB) • GCP (Vertex AI, BigQuery) • OCI

AI/ML & Data

LangChain • LangGraph • RAG • MLflow • OpenSearch • Kafka • Kinesis • Redis

DevOps & Tools

Kubernetes • Terraform • Docker • Prometheus • Grafana • Jenkins • GitLab CI

Languages

Python • Java • Spring Boot • FastAPI • Go

IMPACT METRICS

5M+ Daily Transactions

99.9% System Availability

15TB Daily Data Processing

40% Cost Reduction Achieved

CERTIFICATIONS

AWS

• Solutions Architect Pro
• Machine Learning Specialty

Google Cloud

• Cloud Architect Pro

Databricks

• Solutions Architect Champion

Solutions Architect with **13+ years** designing distributed systems at scale. Specialized in **AWS cloud architecture**, **ML infrastructure**, and **event-driven systems**. Led architecture for systems handling **5M+ daily transactions** with **99.9% availability**.

PROFESSIONAL EXPERIENCE

ARCHITECT @ Tiger Analytics

Feb 2023 - Present

- Built agentic RAG with **LangGraph + MCP + Bedrock** for **2M documents**, autonomous agents self-correct to achieve **87% F1**
- **Failed:** GraphQL federation across 8 services. Reverted to REST after 3 months - boring tech wins
- Built ML model to optimize Docker images for **100+ clients**. Chose heuristics (95% optimal, 30sec) over integer programming (100% optimal, 5min)
- Built MCP-based agent orchestration allowing **30+ AI agents** to collaborate via standardized protocol, **60% reduction** in hallucinations

SR. CLOUD TECHNICAL ENGINEER @ Oracle

Sep 2020 - Dec 2022

- Migrated **50TB Oracle DB** to OCI using GoldenGate, **zero downtime**, **30% cost reduction**
- Built facial recognition training pipeline on **OKE with 4x A100s** for Australian university (**500K images**), handled bias mitigation
- Achieved **SOC2 compliance**: KMS encryption, mTLS, CloudTrail + Splunk audit logging

TECHNICAL LEAD @ Paytm

Jul 2019 - Jun 2020

- **Crisis:** Black Friday outage (8x traffic spike). Implemented ProxySQL for connection multiplexing + circuit breakers in 2 hours
- Built fraud detection with **SageMaker + Kinesis**, solved DynamoDB **40K WCU** limit via sharding
- Reduced cart abandonment **15%** (**\$2M ARR**) by optimizing API **800ms→200ms**

SENIOR CONSULTANT @ Xebia

May 2018 - Jul 2019

- Modernized **12 legacy applications** to AWS, reduced infrastructure costs by **45%** using **Lambda + Fargate**
- Migrated **200+ VMs** to AWS using **AWS MGN**, achieved **zero data loss** across 3 data centers
- Designed disaster recovery for banking client: **RTO 15min**, **RPO 1min** using **Route53 + RDS Multi-AZ**

ASSOCIATE LEAD, TECHNOLOGY @ Nagarro

Jul 2016 - Jan 2018

- Built HPC cluster with **OpenMPI + Slurm** for financial modeling, achieved **10x speedup** on Monte Carlo simulations
- Designed IoT platform processing **1M sensor events/sec** using **MQTT + Cassandra + Spark Streaming**
- Implemented edge computing for **5000+ IoT devices**, reduced latency from **500ms to 50ms**

SOLUTION INTEGRATOR @ Ericsson

Jan 2016 - Jun 2016

- Built telecom inventory management system with **Spring MVC + Hibernate**, tracking **10K+ network elements**
- Implemented real-time sync with network elements using **SNMP + JMS**, reduced discrepancies from **15% to 2%**

ASSOCIATE CONSULTANT @ Wipro

Aug 2012 - Apr 2015

- Built ETL pipelines with **DMExpress + Oracle**, processing **500M CDRs daily**
- **Learning:** No CDC in initial design caused production DB crash. Lesson: incremental > full scans

EDUCATION

M.S. Data Science

IIM & IIT Indore | 2022-24

B.Tech Computer Science

JUET | 2008-12

RESEARCH & OPEN SOURCE

secscan-cli

Security scanning tool on PyPI

MCP-A2A Protocol

Paper submitted to FGCS Journal

PHILOSOPHY

I believe in "boring technology" for critical paths - proven patterns over bleeding edge. Perfect isn't always best: chose 95% optimal in 30sec over 100% optimal in 5min for Docker optimization.

My reviews focus on: "How does this fail?" before "How does this work?"

I choose trade-offs explicitly: DynamoDB over Aurora (10ms vs 50ms), eventual consistency over distributed transactions (10x simpler).

KEY PROJECTS

High-Scale Payment Platform Architecture (Paytm)

Challenge: System crashing at 1M txn/day during flash sales. **Architecture:** Event-sourcing with **Kinesis + DynamoDB + Lambda**. Used CQRS pattern, write to DynamoDB streams, async to RDS. **Scale:** **5M txn/day**, **99.9% uptime**. **Trade-off:** Accepted eventual consistency for 10x throughput.

AI-Driven Docker Optimization Platform (Tiger Analytics)

Challenge: Millions of inefficient Docker images across 100+ clients causing slow deployments. **Solution:** Synthesized 50K Dockerfiles to train optimization model. Compared **integer programming vs heuristics**. Chose heuristics: 95% optimal in 30sec vs 100% optimal in 5min. **Impact:** Average image size reduction **65%**, deployment time **3x faster**, saved clients **\$2M in cloud costs**.

Agentic AI Document Processing Platform (Tiger Analytics)

Challenge: 2M legal documents requiring multi-step reasoning and validation. **Architecture:** **LangGraph agents + MCP protocol + Bedrock Claude**. Agents self-orchestrate: extraction → validation → correction loops. Custom MCP servers for tool isolation. **Results:** **87% F1**, **60% less hallucinations**, agents autonomously handle **90% of L1 queries**.

TECHNICAL LEADERSHIP

- **Architecture Reviews:** Established weekly design reviews at Tiger Analytics, preventing 3 major architectural mistakes
- **Mentorship:** Trained 15+ engineers across AWS, Kubernetes, and ML infrastructure
- **Documentation:** Created architectural decision records (ADRs) for 20+ critical decisions
- **Crisis Management:** Led incident response for 5 P1 outages, established RCA culture
- **Open Source:** Maintain security scanning tools and protocol translation libraries with active community engagement