

CONTACT

ajayprajapat@live.com
+91-87699-62237
linkedin.com/in/ajayprajapat
Udaipur, India · Open to
Hyderabad

AJAY PRAJAPAT

Senior Forward-Deployed Engineer / Incubation Lead

12+ Years · Full-Stack · AI/GenAI · Rapid Prototyping · Pod
Leadership · Greenfield Products

JD MATCH React + Next.js + TypeScript · Python / FastAPI · AWS Lambda + ECS/Fargate · RAG + Agentic AI · Pod Leadership · Greenfield MVPs · Stakeholder Prototyping

CORE STACK

Frontend

- React.js · TypeScript (7-8 yrs)
- Next.js · Tailwind CSS · MUI
- Redux / Zustand / React Query
- Angular (9-10 yrs), Vue.js

Backend

- Python · FastAPI (AI services)
- Node.js / NestJS / Express (9-10 yrs)
- REST · GraphQL · WebSockets
- Async workers · Serverless APIs

AI / GenAI (3+ yrs)

- LangChain · LangGraph
- OpenAI / Azure OpenAI / Claude
- RAG · Embeddings · Vector Search
- Agentic workflows · Tool Calling
- LlamaIndex · Prompt Engineering

Cloud / AWS

- Lambda · API Gateway · S3
- RDS/Aurora · ECS/Fargate
- IAM · CloudWatch · SQS/SNS
- EKS/Kubernetes · Docker
- AWS CDK · Terraform (exposure)

Databases

- PostgreSQL · MongoDB · Redis
- DynamoDB · pgvector · Pinecone
- OpenSearch · Vector DBs

DevOps

- GitHub Actions · GitLab CI
- Docker · CI/CD · Observability
- Secure SDLC · Logging

PROFILE

Hands-on Senior Engineer and Incubation Lead with 12+ years building fullstack products across React, Next.js, TypeScript, Python/FastAPI, Node.js, AWS, and AI/GenAI stacks. Operates daily in rapid build-demo-feedback cycles — translating ambiguous stakeholder requirements into working prototypes within days, not weeks. 3+ years of production Agentic AI delivery: RAG pipelines, tool-calling agents, LangChain/LangGraph orchestration, and LLM-integrated automation across multi-tenant platforms. Proven pod leader for teams of 4-6 engineers; architect-level design judgment with a greenfield-first mindset. Strong stakeholder communication track record — directly engaging business SMEs, surfacing feasibility risks early, and driving prototypes from concept to pilot-ready.

CAPABILITY SNAPSHOT — JD ALIGNMENT

- ✓ Full-Stack Depth (React · FastAPI · Node · AWS)
- ✓ Agentic AI / GenAI — RAG, Tool Calling, LangChain
- ✓ Pod Mentorship — 4-6 engineers, self-sufficient teams
- ✓ Greenfield Product Thinking — MVP to funded pilots
- ✓ Rapid Prototyping & Build-Demo-Feedback Cycles
- ✓ Architecture Judgment (Proto → Pilot → Prod path)
- ✓ Stakeholder Maturity — SME engagement & trade-offs
- ✓ LLM Integration — OpenAI, Azure OpenAI, Claude

PROFESSIONAL EXPERIENCE

Incubation Lead / GenAI Architect — Agentic AI Products

AppNeural (Independent Platforms & Enterprise AI Products)

AI · INCUBATION
2022 – Present (3+ yrs)

LEADERSHIP

Pod Leadership

- Led pods of 4-6 engineers across greenfield AI & SaaS products
- Daily build-demo-feedback cadence
- Mentored engineers toward architect-track readiness

Stakeholder Engagement

- Translated ambiguous SME requirements into prototypes within 1-2 sprints
- Managed expectations & surfaced trade-offs with business owners

Architecture

- Prototype → Pilot → Production path planning
- Modular service & API contract design
- Event-driven · DDD · CQRS
- Multi-tenant SaaS patterns
- Security-aware design
- Feasibility & risk assessment

EDUCATION

MCA

ISBM University — 2025 (Pursuing)

BCA

BNPG College — 2010-2012

THOUGHT LEADERSHIP

- Applied GenAI & Agentic AI for enterprise
- Architecture-First Development
- Agentic AI: Beyond Chatbots
- Microservices: When and Why
- Engineering for Scale

- ▶ Led rapid prototype delivery across 6+ greenfield AI products — from blank-slate concept to demo-ready build in days, operating a daily build-demo-feedback model aligned to incubation sprint cycles.
- ▶ Mentored and directed pods of 4-6 engineers across frontend (React/Next.js), backend (Python/FastAPI, NestJS), and AI layers — maintaining ~70% team self-sufficiency through reusable patterns and clear architectural guardrails.
- ▶ Engaged directly with business stakeholders and SMEs to translate ambiguous product ideas into technical scope, prototype plans, and feasibility assessments — surfacing trade-offs early and driving pilots to funding evaluation.
- ▶ Architected Python/FastAPI + LangChain/LangGraph agentic services with planner-executor-critic orchestration, tool-calling, Redis-backed memory, RAG pipelines, and OpenAI/Claude model routing for production AI automation.
- ▶ Deployed AI microservices on AWS (Lambda, API Gateway, ECS/Fargate, S3, RDS/Aurora, CloudWatch) with containerized delivery, CI/CD guardrails, and observability patterns supporting multi-tenant AI platforms.
- ▶ Built RAG pipelines with pgvector, Pinecone, and OpenSearch over enterprise document corpora; designed embedding workflows, retrieval quality controls, and context-aware response generation for knowledge-intelligence products.
- ▶ Evaluated rapid-build tooling including AI coding assistants, low/no-code workflow automation platforms, and Replit-style environments for accelerating prototype iteration and reducing engineering overhead.

Senior Forward-Deployed Engineer — Full-Stack & AI

LOGISTICS · CLOUD
2021 – Present (5 yrs)

Apple Inc → Infosys → Awign (Contract Delivery)

- ▶ Owned rapid prototype and feature delivery across React/Next.js frontends and Node.js/TypeScript backend services for enterprise logistics and cloud infrastructure platforms — working directly with product stakeholders to convert requirements into working demos within sprint cycles.
- ▶ Designed and shipped AWS-hosted services using Lambda, API Gateway, S3, IAM, ECS/Fargate, and RDS/Aurora; standardized containerized deployments with Kubernetes, GitHub Actions CI/CD, and OpenTelemetry observability.
- ▶ Built React and Next.js UI surfaces with TypeScript, responsive component systems, Redux/Zustand state management, and API-driven UX patterns aligned to backend service contracts.
- ▶ Led Kafka-based event-driven architecture and CQRS workflow design for shipment processing — increasing throughput, improving traceability, and isolating failures across distributed logistics services.
- ▶ Drove backend engineering standards across APIs, messaging (Kafka/SQS), testing, logging, performance optimization, and cost-aware reliability for shared platform services.
- ▶ Shaped secure execution patterns for containerized workloads, including isolation boundaries and auditability requirements for higher-risk automation and AI-generated inputs.

Software Architect / Senior Full-Stack Developer

Aug 2019 – Sep 2021

Smart Factory (Industrial IoT)

- ▶ Led architecture and delivery of real-time industrial monitoring prototypes — rapidly converting operational requirements from plant SMEs into working Angular dashboards and NestJS/WebSocket telemetry services on AWS.
- ▶ Designed Dockerized AWS workloads (EC2, S3, ELB) with monitoring, alerting, and threshold-based detection for manufacturing floor operations — prototype to production path with clear operational handoff.
- ▶ Managed cross-functional delivery pod across frontend, backend, and cloud layers; reviewed designs and guided engineers to maintain velocity without sacrificing architecture quality.

Software Engineer / Project Manager

Oct 2017 – Aug 2019

Neon TSS (IoT / Streaming)

- ▶ Led greenfield migration from PHP stack to Node.js/NestJS and Angular — coordinating backend, device firmware, and frontend engineering through iterative production release cycles.
- ▶ Bridged hardware-aware delivery with cloud backend architecture for a Wi-Fi video streaming platform, improving streaming stability and driving platform to production-ready quality.

Technical Research Assistant — Prototyping & R&D

2015 – 2017

GITS (Academic R&D Lab)

- ▶ Built backend and web prototypes for ERP, LMS, IoT, and content-delivery systems — proof-of-concept delivery with Node.js and Angular/React under academic and production constraints.

SIGNATURE PROJECTS — INCUBATION & AI DELIVERY

Agentic Business Operating System (BizXOS / BizlyUp)

Greenfield AI Platform · 2022–2026

Stack: Python/FastAPI, NestJS, LangChain, LangGraph, OpenAI, Anthropic Claude, Redis, pgvector, AWS Lambda, ECS

- ▶ Designed planner-executor-critic multi-agent architecture with Claude/OpenAI model routing, tool-calling, Redis-backed working memory, and audited human-review checkpoints for autonomous business operations.
- ▶ Built FastAPI AI service layer integrated with LangChain/LangGraph orchestration, RAG retrieval over pgvector, and serverless AWS execution patterns (Lambda + API Gateway + ECS/Fargate).
- ▶ Defined execution guardrails, state transitions, and observability patterns for multi-tenant agent workflows — keeping business automation governable, auditable, and safe for operations teams.

AI Documentation & Knowledge Intelligence (InteraDoc)

GenAI Platform · 2022–2025

Stack: Python/FastAPI, NestJS, LangChain, OpenAI, Anthropic Claude, RAG, Vector Search (pgvector / Pinecone), Access Control

- ▶ Built LangChain-based RAG pipelines over code, documents, and diagrams with embedding workflows, retrieval quality controls, and Claude/OpenAI response generation for enterprise knowledge-intelligence use cases.
- ▶ Designed context-aware Q&A experiences for onboarding, support, and internal knowledge retrieval — rapidly prototyped and iterated with stakeholder feedback cycles to validate retrieval relevance.

GenAI Hiring & Expert Marketplace (GenAIxperts)

GenAI Product · 2022–2025

Stack: NestJS, Python/FastAPI, OpenAI APIs, LangChain, RAG, Semantic Search, pgvector

- ▶ Prototyped LLM-powered resume parsing, semantic search, RAG pipelines, tool-calling workflows, and AI-generated interview question automation — from concept to demo-ready in under 2 sprints.
- ▶ Designed AI-assisted matching flows improving talent-to-role relevance; reduced manual screening effort through structured extraction and AI-ranking workflows.

Enterprise SaaS Platform (Multi-Tenant)

SaaS · 2021–2022

Stack: React, Angular, NestJS, Node.js, AWS (EC2/S3/IAM/ELB), Docker, Kubernetes, RBAC APIs, CI/CD

- ▶ Architected multi-tenant SaaS platform with React/TypeScript frontend, NestJS backend, RBAC/ABAC enforcement, DDD-aligned service boundaries, and AWS-hosted containerized delivery.
- ▶ Standardized release pipelines (GitHub Actions, Docker, Kubernetes) and observability patterns across client environments — creating a repeatable model for greenfield feature delivery.

Agentic IT Support Automation (LocalITCare AI)

Agentic AI · 2022–2025

Stack: NestJS, Python/FastAPI, LangChain, Anthropic Claude, OpenAI, Tool Calling, RAG, Ticket Triage

- ▶ Built issue classification, RCA, technician-routing, and recommendation agents with controlled tool use and auditable support decisions — dispatching AI-ready recommendations before technician assignment, reducing MTTR.

ACHIEVEMENTS & IMPACT

- ▶ Delivered 6+ greenfield AI products from zero to pilot-ready across full stack and AI layers, owning prototype-to-demo cycles end-to-end.
- ▶ 3+ years of production Agentic AI delivery: RAG pipelines, tool-calling agents, LangChain/LangGraph orchestration, multi-agent architectures with observability and guardrails.
- ▶ Mentored and led engineering pods of 4-6 across React, Python/FastAPI, NestJS, and AWS stacks — maintaining 70%+ team self-sufficiency through reusable patterns and clear architecture direction.
- ▶ Directly engaged business stakeholders to scope prototypes, surface feasibility risks, and drive rapid iteration — operating daily build-demo-feedback cadence outside traditional requirement-heavy delivery models.
- ▶ Standardized AWS deployment patterns (Lambda, ECS/Fargate, API Gateway, CloudWatch) with Docker/Kubernetes CI/CD across multiple platform and product teams.
- ▶ Improved engineering throughput through reusable component systems, clean service boundaries, and stronger frontend-backend API contract discipline.

ATS: Senior Forward-Deployed Engineer · Incubation Lead · React.js · Next.js · TypeScript · Python · FastAPI · Node.js · NestJS · AWS Lambda · API Gateway · ECS · Fargate · S3 · RDS · Aurora · RAG · LangChain · LangGraph · LlamaIndex · OpenAI · Anthropic Claude · Azure OpenAI · Agentic AI · GenAI · Tool Calling · Vector Search · pgvector · Pinecone · PostgreSQL · MongoDB · Redis · DynamoDB · Docker · Kubernetes · GitHub Actions · CI/CD · Rapid Prototyping · Pod Leadership · Stakeholder Engagement · Greenfield Products · MVP Delivery · System Design · API Design · Microservices · Event-Driven Architecture · MUI · Tailwind CSS · Redux · Zustand