CLOUD BINARY

CB-FDE-2026

SKILL TODAY. LEAD TOMORROW. • LEARN BY DOING..

Forward Deployed Engineer (FDE)

Official 6 Months (24 Weeks) course syllabus and hands-on lab blueprint. Learn to do what Forward Deployed Engineers do: take cloud, DevOps and AI systems into real customer environments and make them work. 6 months, live, hands-on.

TOTAL DURATION	CURRICULUM DEPTH	ARCHITECTURE LAYERS	DELIVERY FORMAT
6 Months (24 Weeks)	**19 Core Modules**	**5 Progressive Layers**	**Live Online + Labs**



LEAD MENTOR

Kesav

AI FDE Architect & Lead Mentor • 17+ Years in Cloud & Systems Architecture • AWS, Azure, FinOps Certified | Mentored 6,000+ Engineers | Active Forward Deployed Engineer

01. 5 Progressive Architecture Layers (6 Months (24 Weeks))

6 months (24 weeks) · 19 modules · 5 progressive architecture layers: Engineering Foundations & Full-Stack (Weeks 1–14), Data & AI Engineering (Weeks 15–20), Integration, Cloud & DevOps (Weeks 19–21), Security, MLOps & Observability (Weeks 22–23), and Customer Delivery & Field Skills (Weeks 23–24). DSA practice and customer-communication drills continue weekly throughout.

LAYER 01 • WEEKS 1-14 • 7 MODULES

Engineering Foundations & Full-Stack

Modules 01-07

From X to FDE transition roadmaps, programming language fluency (Python, TypeScript, JS), algorithmic complexity & system design (HLD/LLD), Git internals, production Linux diagnostics, modern React frontend architectures, and high-throughput Python FastAPI/PostgreSQL/Redis backend services.

LAYER 02 • WEEKS 15-20 • 4 MODULES

Data & AI Engineering

Modules 08-11

Enterprise customer data auditing, ETL/ELT pipelines, Snowflake, PySpark & Airflow orchestration, LLM fundamentals, hybrid RAG retrieval engines, vector databases, multi-agent systems with LangGraph & MCP, and specification-driven Vibe Coding.

LAYER 03 • WEEKS 19-21 • 4 MODULES

Integration, Cloud & DevOps

Modules 12-15

Enterprise legacy modernization & last-mile integrations (mainframes, ERP, CDC), multi-cloud infrastructure (AWS, GCP, Azure) with modular Terraform IaC, hardened production Kubernetes clusters, and automated GitOps CI/CD with ArgoCD.

LAYER 04 • WEEKS 22-23 • 3 MODULES

Security, MLOps & Observability

Modules 16-18

Automated 12-stage DevSecOps quality gates, Zero-Trust IAM & AI governance guardrails, high-throughput model serving engines (vLLM), MLflow registries, and OpenTelemetry distributed tracing with real-time AI observability dashboards.

LAYER 05 • WEEKS 23-24 • 1 MODULES

Customer Delivery & Field Skills

Module 19 + Grand Capstone

The defining discipline of an FDE: technical customer discovery, scoping & sequencing, executive whiteboarding, C4 architecture modeling, business ROI quantification, stakeholder alignment, production handover runbooks, and live capstone defense.

02. Enterprise Toolchain Across 7 Core Pillars

FOUNDATIONS
`Linux` `Windows` `Bash` `Powershell` `Python` `SQL` `Git` `GitHub`

CLOUD
`AWS` `Azure` `Google Cloud`

DEVOPS
`Docker` `GitHub Actions` `Jenkins` `AWS CodePipeline` `AWS CodeBuild` `Azure DevOps` `Kubernetes` `Helm` `Terraform` `Argo CD`

SECURITY
`SAST/DAST tools` `Trivy-style scanners` `secrets managers`

OBSERVABILITY
`Prometheus` `Grafana` `OpenTelemetry`

AI
`OpenAI` `Anthropic Claude` `Google Gemini` `LangChain` `LangGraph` `MCP` `vector databases` `AWS Bedrock` `Azure OpenAI` `Vertex AI`

FINOPS
`cloud billing and cost-explorer tools`

03. Who Should Join & Prerequisites

Fresher or final-year student (B.Tech, MCA, B.Sc, BCA, MBA, BCOM)
A structured path into cloud and AI engineering, with real projects and customer-facing practice

Developer, tester or support engineer (1–5 yrs)
Move from ticket-based work into high-ownership, client-facing engineering roles

DevOps, cloud or SRE engineer (3–10 yrs)
Add GenAI, agents and customer-deployment skills to grow into FDE and solutions roles

Architect or tech lead (10–20 yrs)
Lead AI and cloud deployments for enterprise clients with modern FDE practices

Prerequisites: FDE is not a beginner role; you will most probably be already a Full-Stack Developer stepping into an FDE role. Core foundations and delivery models are introduced in Module 1.

Layer 01: Engineering Foundations & Full-Stack (Part 1: Core Foundations)

Establishing the FDE baseline: career transition roadmap, multi-language fluency (Python & TypeScript), and algorithmic interview gates with system design.

MODULE 01 | LAYER 01 · FOUNDATIONS & FULL-STACK | Weeks 1-2

Introduction

Roadmap Insight: FDE is not a beginner role, you will most probably be already a Full-Stack Developer stepping into an FDE role.

Foundational breakdown of the Forward Deployed Engineer discipline. Explore career transition pathways from traditional software engineering into high-impact customer-embedded engineering, analyze FDE execution tiers, master core responsibilities, and build fluency in essential industry terminology.

CORE TOPICS:

- From X to FDE
- Roles & Responsibilities
- FDE Tiers

FDE TERMINOLOGY:

- Delta
- UAT
- Last-Mile Integration
- System of Record
- Day 2 Operations
- Productized Consulting
- Embedded Engineering

HANDS-ON PRODUCTION DRILL:
Enterprise FDE Discovery & Role Mapping: Map your transition path (From X to FDE), analyze enterprise delivery tiers, and complete a simulated customer delta & last-mile integration assessment.

MODULE 02 | LAYER 01 · FOUNDATIONS & FULL-STACK | Weeks 3-4

Learn a Language

Roadmap Insight: Python is the go-to language for FDEs.

Master the essential programming languages of Forward Deployed Engineering. Deep-dive into Python for rapid customer prototyping, data integration, scripting, and agentic AI workflows, alongside JavaScript / TypeScript for full-stack client integrations and enterprise web services.

CORE TOPICS:

- Python
- JavaScript / TypeScript

HANDS-ON PRODUCTION DRILL:
Enterprise Integration Bridge: Build a full-stack integration connecting a Python backend data pipeline with a TypeScript client interface.

MODULE 03 | LAYER 01 · FOUNDATIONS & FULL-STACK | Weeks 5-6

DSA & System Design

Roadmap Insight: DSA is the first gate in FDE interviews. You won't get to show your AI skills if you don't pass this round.

Master technical interview gates and large-scale architectural problem solving. Rigorous practice across data structures, algorithms, and complexity analysis combined with real-world enterprise system design: high availability, data partitioning, caching strategies, asynchronous processing, and client integration trade-offs.

CORE TOPICS:

- DSA Roadmap
- System Design Roadmap
- Algorithmic Complexity & Optimization
- High-Level & Low-Level Design (HLD/LLD)

HANDS-ON PRODUCTION DRILL:
Enterprise System Design Defense: Design and defend a distributed, fault-tolerant customer platform architecture document (HLD/LLD) with core algorithm implementations.

Layer 01: Engineering Foundations & Full-Stack (Part 2: Systems & Frontend)

Enterprise version control governance, deep Linux kernel diagnostics & socket troubleshooting, and modern customer-facing React architectures.

MODULE 04 | LAYER 01 · FOUNDATIONS & FULL-STACK | Weeks 7-8

Git & GitHub

Master enterprise version control, branching strategies, and collaborative code delivery. Deep-dive into Git internal mechanics (blobs, trees, commits, refs), advanced merge techniques (rebase, squash, cherry-pick, bisect), GitHub pull request workflows, branch protection governance, and pre-commit automation.

CORE TOPICS:

`Git & GitHub Roadmap` `Git Internals & Plumbing Commands` `Branching, Rebasing & Conflict Resolution` `PR Workflows, Code Reviews & Protections` `Git Hooks & Automated Quality Gates`

HANDS-ON PRODUCTION DRILL:
Enterprise Collaboration & Versioning Architecture: Configure a production GitHub repository framework featuring automated PR validations, branch protection rules, semantic versioning, and pre-commit hooks.

MODULE 05 | LAYER 01 · FOUNDATIONS & FULL-STACK | Weeks 9-10

Linux Skills

Acquire hands-on mastery of enterprise Linux systems, diagnostics, and terminal automation. Deep-dive into Linux kernel architecture, userspace diagnostics, systemd service lifecycle, and process management. Master production Bash automation, regex stream transformations (awk, sed, grep), storage architectures (LVM, ext4, XFS), and network socket troubleshooting.

CORE TOPICS:

`Linux Roadmap` `Shell & Bash Roadmap` `Kernel Internals & Memory Diagnostics` `Process Lifecycle & systemd Services` `POSIX Permissions & SSH Hardening` `Network Sockets & Troubleshooting (ss, netstat, lsof, tcpdump)`

HANDS-ON PRODUCTION DRILL:
Linux Break-Fix & Automation Lab: Triage a live simulated production server failure (CPU lockup, disk/inode exhaustion, and network socket degradation), restore system health, and deploy an automated Bash watchdog daemon.

MODULE 06 | LAYER 01 · FOUNDATIONS & FULL-STACK | Weeks 11-12

Frontend Skills

Roadmap Insight: FDEs need to be proficient in full-stack development, have high agency and be able to deliver full apps for customers.

Master modern client-side software engineering to deliver demo-ready, production-grade applications for enterprise customers. Build deep fluency across HTML5, CSS responsive design systems (Flexbox, Grid, Tailwind), core JavaScript async mechanics (Event Loop, Promises), and React component architecture (Hooks, Context, Virtual DOM, and state management).

CORE TOPICS:

`Frontend Roadmap` `HTML & Semantic Structure` `CSS & Responsive Design (Tailwind, Flexbox, Grid)` `JavaScript (ES6+, Event Loop, Async/Await)` `React (Hooks, Component Architecture, State Management)` `Frontend Project Ideas`

HANDS-ON PRODUCTION DRILL:
Enterprise Client Telemetry Dashboard: Architect a responsive React + TypeScript analytics portal with real-time WebSocket data updates, dynamic filtering, and production edge deployment.

Layer 01 & Layer 02: Resilient Backend, Data Engineering & AI Foundations

High-throughput Python FastAPI/PostgreSQL backends, enterprise data audits with Snowflake/PySpark/Airflow, and production RAG retrieval engines.

MODULE 07 | LAYER 01 · FOUNDATIONS & FULL-STACK | Weeks 13-14

Backend Skills

Engineer resilient, high-throughput backend services and enterprise data layers with Python and FastAPI. Master asynchronous programming, relational database schema design with SQL and PostgreSQL (indexing, query planning, transactions, ACID), NoSQL document and key-value stores (MongoDB, Redis), RESTful API contract design, OpenAPI specifications, and modern GraphQL schemas.

CORE TOPICS:

- Backend Roadmap
- Python (AsyncIO, type hints & application structure)
- FastAPI (routing, dependency injection, validation & middleware)
- SQL & PostgreSQL (Schema Design, Indexing, Transactions)
- NoSQL DBs (MongoDB, Redis, Document & Key-Value Stores)
- APIs Design (RESTful, OpenAPI Specs, Rate Limiting)
- GraphQL (Schemas, Resolvers & DataLoader)
- Backend Project Ideas

HANDS-ON PRODUCTION DRILL:
Enterprise High-Throughput Core API: Architect and benchmark a distributed Python FastAPI + PostgreSQL backend with Redis caching, ACID transaction rollbacks, OpenAPI documentation, and a unified GraphQL gateway.

MODULE 08 | LAYER 02 · DATA & AI ENGINEERING | Weeks 15-16

Data Engineering Skills

Roadmap Insight: A large share of FDE work starts with a data audit. If you can't untangle a customer's schema, you can't build on top of it.

Audit, transform, and operationalize enterprise data at scale. Master data modeling paradigms (Star/Snowflake schemas, Normalization, Data Vault), robust ETL/ELT pipeline design, modern cloud data warehouses (Snowflake, BigQuery, Redshift), distributed batch and stream processing with Apache Spark, and workflow orchestration with Apache Airflow.

CORE TOPICS:

- Data Engineer Roadmap
- Data Modeling (Star/Snowflake, Normalization, Data Vault)
- Data Pipelines (ETL / ELT, Change Data Capture, Idempotency)
- Data Warehouses (Snowflake, BigQuery, Redshift, Lakehouse)
- Apache Spark (Distributed DataFrames, PySpark, Optimization)
- Apache Airflow (DAG Orchestration, Operators, Monitoring)

HANDS-ON PRODUCTION DRILL:
Enterprise Customer Data Audit & ELT Pipeline: Ingest messy, heterogeneous client data sources into a modern cloud data warehouse with an automated Airflow DAG, performing deduplication, schema validation, and PySpark transformations for downstream analytics and AI consumption.

MODULE 09 | LAYER 02 · DATA & AI ENGINEERING | Weeks 17-18

AI Engineering Skills

Master foundation models, token economics, and enterprise LLM integration. Deep-dive into transformer architectures, advanced prompt engineering, context engineering, semantic search with vector databases (Pinecone, pgvector, Qdrant), production Retrieval-Augmented Generation (RAG) pipelines, model provider evaluation, and prompt lifecycle management.

CORE TOPICS:

- AI Engineer Roadmap
- LLM Fundamentals (Transformers, Attention, Token Economics)
- Prompt Engineering (Few-Shot, CoT, System Prompts, Structured JSON)
- Context Engineering (Window Management, Compression, Dynamic Context)
- RAGs (Document Ingestion, Chunking, Hybrid Search & Reranking)
- Vector DBs (Pinecone, Qdrant, ChromaDB, pgvector)
- Choosing Model Providers (Frontier vs Open Weight, Cost, Latency & Privacy)
- Prompt Management & Evals (Evaluation Metrics, Guardrails & Langfuse)

HANDS-ON PRODUCTION DRILL:
Production Document Intelligence RAG Engine: Build an enterprise-grade retrieval pipeline featuring semantic chunking, hybrid vector search (Pinecone + BM25), Cohere reranking, structured JSON enforcement, and latency/token-cost observability.

Layer 02 & Layer 03: Autonomous Agents, Vibe Coding & Enterprise Integration

Architecting stateful LangGraph & MCP agent swarms, spec-driven Vibe Coding delivery, and bridging enterprise legacy mainframes, ERPs & CDC.

MODULE 10 | LAYER 02 · DATA & AI ENGINEERING | Weeks 19-20

AI Agents

Architect, build, and deploy autonomous enterprise AI agents and multi-agent systems. Master cognitive agent architectures (ReAct, Plan-and-Solve), tool and function calling integration, short-term and persistent memory systems, Model Context Protocol (MCP) server/client ecosystems, and stateful multi-agent orchestration using LangGraph.

CORE TOPICS:

- Agent Architectures (ReAct Loops, Planning, Reflection, Human-in-the-Loop)
- Tools & Functions (Function Calling, Tool Schemas, SQL/API Execution)
- Memory & State (Working Memory, Vector Persistence, State Graphs)
- Model Context Protocol (MCP Clients, Servers, Dynamic Tools)
- Multi-Agents (Supervisor Patterns, Swarms, LangGraph Workflows)
- AI Agent Project Ideas

HANDS-ON PRODUCTION DRILL:
Enterprise Multi-Agent Operations Swarm: Architect a stateful agent network using LangGraph and Model Context Protocol (MCP) connecting live enterprise SQL databases, customer APIs, persistent memory, and supervisory approval gates.

MODULE 11 | LAYER 02 · DATA & AI ENGINEERING | Week 19

Vibe Coding

Accelerate enterprise shipping speed with AI-assisted software engineering and specification-driven development. Master next-generation coding agents and LLM copilot workflows using Claude Code, OpenAI Codex, Google Gemini Code Assist, and Cursor to iterate at 10x velocity while maintaining strict code quality, type safety, and test coverage.

CORE TOPICS:

- AI-Assisted Coding Tools: Claude Code, Codex, Gemini, Cursor
- Specification-Driven Development: Authoring PRDs and technical specs for AI coding agents
- Interactive Pair Programming: Context steering, repository indexing, and multi-file code refactoring
- Automated Test Generation & Regression Verification with AI Copilots
- Vibe Coding Governance: Guarding against AI hallucination, security anti-patterns, and license risks

HANDS-ON PRODUCTION DRILL:
Rapid Prototyping & Greenfield Delivery Drill: Build, test, and ship a full-stack production microservice from PRD to passing end-to-end test suite in under 4 hours using Claude Code and Cursor.

MODULE 12 | LAYER 03 · INTEGRATION, CLOUD & DEVOPS | Week 19

Integration Engineering

Master last-mile integration and enterprise systems connectivity. FDEs spend significant time embedding into customer environments with legacy mainframes, ERPs, SOAP/XML endpoints, and proprietary data stores. Learn patterns for legacy modernization, resilient API adapters, message brokering, idempotency, data synchronization, and backward-compatible cutover strategies.

CORE TOPICS:

- Working with Legacy Systems: SOAP/XML, Mainframe bridges, monolithic data stores
- Last-Mile Integration Architecture: Sidecars, proxies, and API translation layers
- Enterprise Systems of Record (SoR) Integration: SAP, Salesforce, Oracle, ServiceNow
- Data Transformation & CDC: Event-driven synchronization, Kafka Connect, Debezium
- Resilience & Fault Isolation: Circuit breakers, retry budgets, dead-letter queues (DLQ), and UAT sign-off

HANDS-ON PRODUCTION DRILL:
Legacy System Bridge & Modernization Proxy: Architect an event-driven adapter proxying legacy enterprise XML/RPC endpoints to modern REST/GraphQL APIs with schema validation, Redis caching, and circuit-breaker fault tolerance.

Layer 03: Multi-Cloud, Kubernetes & GitOps Platform Delivery

Multi-cloud infrastructure provisioning with modular Terraform, production Kubernetes microservice orchestration, and declarative GitOps via ArgoCD.

MODULE 13 | LAYER 03 · INTEGRATION, CLOUD & DEVOPS | Week 20

Cloud Platforms

Deploy and manage enterprise infrastructure across AWS, GCP, and Azure. Master multi-cloud compute and networking, managed machine learning services (AWS Bedrock & SageMaker, Google Cloud Vertex AI, Azure OpenAI), and declarative Infrastructure as Code (IaC) with Terraform.

CORE TOPICS:

- DevOps & Cloud Skills Overview
- Multi-Cloud Foundations: AWS, Google Cloud (GCP), Microsoft Azure
- Managed ML Services: Amazon Bedrock, SageMaker, Google Vertex AI, Azure OpenAI
- Infrastructure as Code (IaC): Terraform providers, state management, modular architectures
- Cloud Networking & Security: Multi-tier VPCs, private endpoints, IAM roles, and secret management

HANDS-ON PRODUCTION DRILL:
Multi-Cloud AI Platform Provisioning: Write modular Terraform configurations to provision secure VPCs across AWS and GCP, wiring IAM access to Amazon Bedrock and Google Vertex AI endpoints with automated state locking.

MODULE 14 | LAYER 03 · INTEGRATION, CLOUD & DEVOPS | Week 20

Containers

Containerize and orchestrate enterprise microservices for mission-critical client deployments. Master Docker multi-stage image optimization, non-root security, local orchestration with Docker Compose, and production Kubernetes (K8s) deployment patterns (Pods, Deployments, Services, Ingress, ConfigMaps, and Helm).

CORE TOPICS:

- Docker Fundamentals: OCI images, Dockerfile multi-stage builds, layer caching, distroless images
- Docker Networking & Storage: Bridge/overlay networks, volume persistence, Docker Compose
- Kubernetes Architecture: Control Plane, worker nodes, kubelet, Pod lifecycle, Deployments & ReplicaSets
- K8s Service Discovery & Routing: ClusterIP, NodePort, LoadBalancer, Ingress Controllers, cert-manager
- Configuration & Secret Management: ConfigMaps, Kubernetes Secrets, RBAC, and Helm package management
- Enterprise Production Deployment: Zero-downtime rolling updates, resource requests/limits, and health probes

HANDS-ON PRODUCTION DRILL:
Production Kubernetes Microservices Deployment: Package a multi-tier AI/Full-Stack application into hardened Docker containers and deploy to a Kubernetes cluster using Helm with automated horizontal pod autoscaling (HPA) and zero-downtime rolling updates.

MODULE 15 | LAYER 03 · INTEGRATION, CLOUD & DEVOPS | Week 21

DevOps Skills

Master end-to-end continuous integration, continuous delivery (CI/CD), and GitOps deployment automation. Build automated pipeline quality gates using GitHub Actions, container registries, automated testing, and declarative GitOps continuous delivery with ArgoCD.

CORE TOPICS:

- DevOps Roadmap: CI/CD Pipelines & Automation Fundamentals
- Continuous Integration (CI): GitHub Actions, automated test execution, linting, matrix builds
- Artifact Management: Container image registries, semantic versioning, and release tagging
- Continuous Delivery & Deployment (CD): Staged environments (Dev, Staging, Prod), branch protections
- GitOps Automation: Declarative continuous delivery with ArgoCD, drift detection, and auto-sync
- Release Strategies: Blue/Green deployments, Canary rollouts with Argo Rollouts, and automated rollbacks

HANDS-ON PRODUCTION DRILL:
End-to-End GitOps Deployment Pipeline: Configure an automated GitHub Actions CI workflow that tests, builds, and pushes container images, triggering declarative ArgoCD GitOps deployment to a Kubernetes cluster with Canary validation.

Layer 04: DevSecOps, High-Throughput MLOps & AI Observability

12-stage automated security gates, high-throughput vLLM model serving, MLflow registries, and OpenTelemetry distributed tracing for agent swarms.

MODULE 16 | LAYER 04 · SECURITY, MLOPS & OBSERVABILITY | Week 22

Security

Embed defense-in-depth security across every layer of the enterprise application lifecycle. Master DevSecOps automated shift-left security scanning, API security (OWASP Top 10 API), modern authentication (OAuth 2.0, OIDC, JWT, SAML), AI governance and safety guardrails, and enterprise data privacy compliance (GDPR, HIPAA, SOC 2).

CORE TOPICS:

- DevSecOps: Shift-left security, SAST (Semgrep/SonarQube), SCA (Snyk), and container vulnerability scanning (Trivy)
- API Security: OWASP API Security Top 10, rate limiting, request validation, mutual TLS (mTLS), and WAF rules
- Authentication & Authorization: OAuth 2.0, OpenID Connect (OIDC), JWT verification, RBAC, and ABAC policies
- AI Governance: Prompt injection defense, jailbreak detection, toxicity filtering, and automated guardrails
- Data Privacy & Compliance: Data anonymization, PII masking, encryption at rest/in transit, SOC 2, HIPAA, GDPR

HANDS-ON PRODUCTION DRILL:
Automated 12-Stage DevSecOps & AI Guardrail Gate: Build a bulletproof CI/CD security pipeline enforcing SAST, SCA, container scanning, secret detection, and an AI proxy gateway with automated PII redaction and prompt-injection filtering.

MODULE 17 | LAYER 04 · SECURITY, MLOPS & OBSERVABILITY | Week 23

MLOps

Bridge machine learning engineering with production cloud operations. Master the MLOps Roadmap: scalable model serving architectures (vLLM, Ollama, Triton Inference Server), automated pipeline orchestration (Kubeflow, MLflow), continuous experiment tracking, model registry governance, and automated model retraining triggers.

CORE TOPICS:

- MLOps Roadmap: Production Machine Learning Lifecycle & Operational Architecture
- Model Serving Engines: vLLM, Ollama, HuggingFace TGI, Triton Inference Server, continuous batching
- Model Registry & Versioning: MLflow, Weights & Biases, artifact storage, and model lineage
- Pipeline Orchestration: Kubeflow Pipelines, automated feature stores, and batch inference workflows
- Model Packaging & Quantization: GGUF, AWQ, GPTQ, INT8/INT4 precision for GPU memory optimization

HANDS-ON PRODUCTION DRILL:
High-Throughput Model Serving Engine: Deploy a quantized open-weight foundation model using vLLM on a GPU-enabled cluster with continuous batching, streaming response APIs, and MLflow model registry tracking.

MODULE 18 | LAYER 04 · SECURITY, MLOPS & OBSERVABILITY | Week 23

Observability

Gain deep, real-time operational visibility into production distributed systems and AI applications. Master automated evaluation pipelines, deep debugging of autonomous agent execution traces, token latency profiling, inference cost optimization, distributed tracing (OpenTelemetry), and LLM telemetry (Langfuse, Arize Phoenix).

CORE TOPICS:

- Observability Fundamentals: Metrics, Logs, Traces (OpenTelemetry, Prometheus, Grafana)
- Building Eval Pipelines: Ragas framework, synthetic test datasets, LLM-as-a-Judge evaluations
- Debugging Agent Logic: Execution trace visualization, cyclical loop detection, tool-call debugging (Langfuse)
- Inference Optimization: TTFT (Time to First Token), TPOT (Time per Output Token), caching, prompt compression
- Cost & Telemetry Dashboards: Real-time token consumption, per-user billing attribution, and SLA alert triggers

HANDS-ON PRODUCTION DRILL:
Full-Stack AI Observability & Evaluation Platform: Instrument a production LangGraph agent with OpenTelemetry and Langfuse, tracking execution traces, token costs, latency percentiles, and automated Ragas evaluation benchmarks.

Layer 05: Customer Delivery, Field Acumen & Capstone Simulation

The defining capstone discipline of an FDE: stakeholder discovery, executive whiteboarding, business ROI quantification, and shipping verified client outcomes.

MODULE 19 | LAYER 05 · CUSTOMER DELIVERY & FIELD SKILLS | Week 24

Customer Delivery & Field Skills

The defining discipline of a true Forward Deployed Engineer. Master the trifecta of Discovery & Scoping, Business Acumen, and High-Impact Technical Communication. Learn to extract unstated customer pain points, navigate enterprise workflows, manage executive stakeholders, author production Statements of Work (SOWs), and deliver verified business ROI.

CORE TOPICS:

- Discovery & Scoping: Requirements gathering, scoping & sequencing, balancing Scope, Speed, and Quality
- Business Acumen: Enterprise workflows, ROI & AI impact quantification, product feedback loops, Time-to-Value
- Communication: Technical writing, presenting architectural options, executive reporting, de-escalation
- Stakeholder Management: Partnering with customer CTOs, product managers, security teams, and central engineering
- Production Handover: Architecture Decision Records (ADRs), operations runbooks, and Day-2 knowledge transfer

HANDS-ON PRODUCTION DRILL:
The Grand FDE Customer Deployment Simulation: Lead a simulated enterprise deployment from discovery brief and executive whiteboarding to architecture spec (C4 model), working multi-tier implementation, live executive demo, and production handover runbook.

CULMINATING PROGRAM SIMULATION (WEEK 24) | LIVE CLIENT DEMO & HANDOVER

Production Capstone Projects

In the final phases, you work through realistic client engagements modeled on live enterprise products. You build, secure, containerize, deploy, and monitor full-stack AI and cloud solutions, finishing with an architectural defense and live client handover.

REPRESENTATIVE CAPSTONE BRIEFS:

- #01: Cloud Binary LMS Student AI Companion
- #02: C3Ops Career Co-Pilot
- #03: C3Ops FinOps Agent

04. Forward Deployed Engineering Career Pathways

FDE compensation and impact across industry experience brackets:

0-2 yrs	Associate Cloud Engineer · Junior DevOps Engineer · Associate Solutions Engineer · Junior AI Engineer
2-5 yrs	Forward Deployed Engineer · DevOps Engineer · Cloud Engineer · Solutions Engineer
5-10 yrs	Senior Forward Deployed Engineer · SRE · Platform Engineer · AI Deployment Engineer
10-15 yrs	FDE Lead · Solutions Architect · Cloud/AI Architect
15-20 yrs	Head of Field Engineering · Principal Architect · Director of Solutions

Career Acceleration Support: Career support included — resume and LinkedIn review, GitHub portfolio review, mock interviews with FDE-style case rounds, and demo day.

Ready to Step into Forward Deployed Engineering?

Live classes led by Kesav, AI FDE Architect & Lead Mentor.
Live Online & In-Person Hybrid (WeWork, K Raheja Mindspace, Madhapur, Hyderabad) with 1-on-1 Mentor Code Audits.

WhatsApp: +91 91000 73006 | Email: info@cloudbinary.in | Web: https://cloudbinary.in

Live Online across India + In-Person (WeWork, K Raheja Mindspace, Madhapur, Hyderabad) + 24/7 Cloud Range