

DEXTERITY VENTURE INSTITUTE | CFDA - CERTIFIED FORWARD-DEPLOYED AI SPECIALIST

The Complete FDE Journey - 25 Learning Blocks, 75 Hours | Batch commencing last week of August 2026 | Track A: Mon/Wed/Fri evenings | Track B: Sat/Sun | Both tracks cover identical content

A Forward-Deployed AI Specialist is not a deployment engineer. It is an engineer who works end-to-end on the client's side - understanding the business problem, designing the AI solution, building it, shipping it into the client's environment, and making it stick. The curriculum below is organised by that journey.

Journey stage	Block	Block title	Hrs	What you will be able to do	What you will produce	Industry lens
1. UNDERSTAND THE CLIENT (6 hrs)	B01	The FDE Role and the Enterprise Environment	3.0	Distinguish the FDE role from Solutions Architect, Sales Engineer and Consultant; map a client's stakeholders, approval gates and constraints before proposing anything technical.	FDE role-positioning statement; enterprise constraint register	All four tracks - role variants compared
	B02	Problem Structuring, Customer Discovery and the POC Charter	3.0	Convert an ambiguous executive ask into a structured issue tree, run client discovery interviews, and produce a POC charter with success criteria and a kill switch.	Issue tree + POC Charter (graded); working dev environment	Four parallel client asks - one per track
2. ENGINEER LIKE A PROFESSIONAL (12 hrs)	B03	Production-Ready Python	3.0	Restructure exploratory code into a typed, configured, tested Python package that someone else can verify and run.	Refactored package with test suite (70%+ coverage)	Common baseline
	B04	Git, GitFlow and the Client Repository	3.0	Operate confidently inside a shared client repository - branching strategy, reviewable pull requests, clean handover hygiene.	Team repo with protected main + merged PRs (graded)	Cross-track teams own a client repo
	B25	AI-Assisted Development: Coding Agents in the FDE Workflow	3.0	Drive coding agents (Claude Code, Cursor, Copilot) as an engineering discipline - scope work into agent-executable specs, supply repository context deliberately, and review agent-authored diffs as rigorously as a human pull request.	Repository context file + a feature shipped agent-first with a review log (graded)	Common baseline; client code-confidentiality and IP limits per track
	B08	Data Pipelines, Quality and Lineage	3.0	Build the ingestion layer every RAG and agent system depends on - with quality gates, PII masking and source-level traceability.	Reusable ingestion pipeline with quality + PII gates (graded)	Healthcare claims, IT Ops telemetry, Consulting documents
3. BUILD AI SYSTEMS (18 hrs)	B09	ML Foundations and Model Lifecycle in Production	3.0	Frame a business problem as an ML problem with a metric the client will accept; operate the model past deployment day.	Metric justification memo; registered model with promotion workflow	Healthcare billing accuracy, IT Ops incident triage
	B10	Transformers, NLP and the Hospital Billing Fine-Tuning Case	3.0	Settle the fine-tune vs prompt vs retrieve decision on evidence, and integrate a fine-tuned model into a live claims workflow.	Fine-tuned model + evaluation report (graded)	Healthcare leads; analogous case per track
	B11	Model Selection and Prompt Engineering as Engineering	3.0	Select an LLM on cost, latency and data-residency evidence; treat prompts as versioned, tested, evaluated artefacts.	Model selection scorecard; versioned prompt library with evals	Residency constraints: Healthcare vs Consulting
	B12	LLM Workflow Patterns and AI Systems Architecture	3.0	Design the harness around the model - context assembly, tool surface, memory, control loop - which determines whether the system works in production.	System architecture pack + working eval harness (graded)	Each team architects for its own client
	B13	RAG Fundamentals and Advanced Retrieval	3.0	Build a retrieval pipeline whose failures you can explain; upgrade it with hybrid search, re-ranking and measured evaluation.	Evaluated RAG system with before/after metrics	Each team indexes its own track corpus
	B14	Rapid Prototyping and the Financial Knowledge Assistant	3.0	Deliver an end-to-end retrieval + generation system with citations and a credible interface - fast enough to matter in a client meeting.	Deployed knowledge assistant with citations (graded)	Consulting leads; parallel corpora per track
	B15	Agent Anatomy and Multi-Agent Architecture	3.0	Define what makes a system agentic, screen where agency adds value vs risk, and choose a multi-agent topology deliberately.	Agentic-fit screening matrix; topology design	Four track use cases screened

Journey stage	Block	Block title	Hrs	What you will be able to do	What you will produce	Industry lens
4. BUILD AGENTIC SYSTEMS (12 hrs)	B16	LangGraph and CrewAI Hands-On	3.0	Build the same workflow twice - as a state graph and as a role-based crew - and argue for one on controllability, cost and debuggability.	Two working implementations + comparison memo (graded)	Cross-track teams
	B17	AutoGen, PydanticAI, Interoperability and Agent Memory	3.0	Orchestrate conversational agent teams, impose type safety, make frameworks interoperate, and give agents memory that can be purged on request.	Typed agent + cross-framework demo; agent with purgeable memory	Framework recommendation per track
	B18	Google ADK and the Shipment Rerouting Agent	3.0	Build an enterprise-grade agent on Google ADK and apply multi-agent design to a live logistics disruption with human escalation.	ADK agent + multi-agent rerouting system (graded)	IT Ops and AI in IT lead; Automobile variant
5. SHIP INTO THE CLIENT'S ENVIRONMENT (15 hrs)	B05	Jenkins, GitHub Actions and CI/CD Pipelines	3.0	Wire source control to automated build, test and deploy; choose the right CI stack for a client's on-premise or cloud constraints.	Working CI pipeline triggered on pull request	IT Ops and AI in IT lead; Healthcare on-prem rationale
	B06	Containerisation with Docker	3.0	Package a Python AI service into a lean, security-scanned, reproducible image published through an automated pipeline.	Published container image + CI rebuild job (graded)	Cross-track teams containerise their own service
	B07	Kubernetes via Minikube and API Contract Design	3.0	Deploy, scale and roll back a containerised AI service; publish an API contract a client engineering team can integrate against unaided.	Deployed, load-tested AI endpoint + OpenAPI spec (graded)	IT Ops leads the platform conversation
	B19	Cloud Topologies, Amazon Bedrock and AWS Copilot	3.0	Pick a deployment topology that survives the client's security review; deploy a managed-LLM application with platform guardrails.	Topology decision record; Bedrock-backed deployed service	Healthcare private, Consulting multi-tenant, IT Ops hybrid
	B20	On-Premise Inference and Agentic Deployment on GCP	3.0	Run capable models inside the client perimeter and quantify the cost; take your agent system to a monitored cloud deployment.	Locally served model; cloud-deployed monitored application (graded)	Healthcare and BFSI-adjacent Consulting lead
6. GOVERN, SECURE AND OBSERVE (6 hrs)	B21	OWASP LLM Top 10 and Governance under the DPDP Act	3.0	Attack your own system against a recognised threat taxonomy; design controls for DPDP Act 2023, RBI, NABH, NMC and CDSO obligations.	Threat assessment with reproduced exploits; governance control matrix	Sectoral overlays per track
	B22	LLM Observability and Arize Phoenix Instrumentation	3.0	Define the metrics that reveal degradation before the client does; trace every answer end-to-end and triage a production incident.	Fully instrumented pipeline + incident triage report (graded)	Track-specific KPI sets
7. MAKE IT STICK WITH THE CLIENT (6 hrs)	B23	From POC to Production and the Solution Design Document	3.0	Run the transition most AI projects fail at; produce the written artefacts a client architecture board will actually approve.	Production readiness pack; complete Solution Design Document (graded)	Team's own client system
	B24	Client Engagement, Capstone Demo Day and Final Assessment	3.0	Handle the human half of the role - stakeholder communication, scope creep, executive updates - then present and defend your full system.	Capstone demo before a client-persona panel + final assessment	Panel includes a persona from each track

THE SKILL MIX AT A GLANCE			
Building AI and agentic systems (Stages 3-4)	30 hrs	40.0%	
Shipping and deployment engineering (Stage 5)	15 hrs	20.0%	
Client-embedded consultative work (Stages 1 + 7)	12 hrs	16.0%	
Professional engineering, data and AI-assisted development (Stage 2)	12 hrs	16.0%	
Governance, security and observability (Stage 6)	6 hrs	8.0%	
Every block produces a concrete artefact, and 14 of the 25 blocks carry a graded deliverable - so by Demo Day you hold a portfolio of client-ready work, not just certificates of attendance. Deployment engineering is one stage of the journey (about 20%); the rest is the full-stack, client-embedded work that defines the forward-deployed role. Block B25 (AI-assisted development) was added in v1.1 following Batch 2 learner feedback.			