



FDE AI Agent Frameworks & Tools Guide

A Practical Handbook for Selecting Frameworks, Integrating Tools, and Deploying Production-Ready AI Agents

Learn how Forward Deployed Engineers choose AI agent frameworks, connect business systems, design reliable workflows, and turn enterprise AI requirements into working solutions.

Global Skill Development Council (GSDC)

From Business Problems to Working AI Agents

Start with the business workflow and its failure consequences. Choose an agent only when adaptive reasoning or tool use adds value over simpler automation.



Six responsibilities, one delivery chain

1. Discover the real problem and intended outcome.
2. Map workflows, data, access and operating constraints.
3. Decide whether an agent is the right solution.
4. Select the model, framework and integrations.
5. Prototype, test and validate with users.
6. Deploy, monitor and improve against agreed criteria.


Illustrative financial-services scenario

A firm receives thousands of enquiries, but retrieval across internal knowledge sources is unreliable. The FDE investigates the workflow, identifies data and access requirements, and connects a retrieval-enabled agent to approved systems. Validation compares manual effort with accuracy and security before consequential customer actions are enabled.

The FDE contribution combines business context, engineering trade-offs and reliable delivery.

Inside an Enterprise AI Agent

An enterprise agent is a controlled application around a language model. Reliability depends on the full system, including identity, data access, execution and feedback.



Eight responsibilities

User interface; agent runtime; language model; tool layer; knowledge and data layer; orchestration; guardrails and permissions; evaluation and observability.

Right-size the boundaries

These are logical components, not a requirement for eight separate services. Match deployment boundaries to complexity, ownership and risk.

An LLM response is not a verified business action. Application logic must mediate tool execution through authorization, validated inputs, explicit error handling and controlled commits.

Choosing the Right AI Agent Framework

Compare the execution model and operational fit before adopting an abstraction. These options are not interchangeable products, and no framework is a universal winner.

Framework or SDK	Primary approach	Strengths	Suitable scenarios	Considerations
LangGraph	Graph-based, explicit control	Stateful execution, branching and persistence	Multi-step workflows with recovery or human intervention	Budget effort for graph, state and execution design.
CrewAI	Role-oriented collaboration and delegation	Organizing agent roles and tasks	Work that benefits measurably from delegated responsibilities	Assess control, debugging and operational complexity.
OpenAI Agents SDK	Agent orchestration, tools and handoffs	Agent building within the OpenAI ecosystem	Tool-enabled assistants and agent handoffs	Evaluate provider requirements, portability, integrations and current capabilities.
Microsoft Agent Framework	Agent and multi-agent workflow orchestration	Microsoft's framework for building and coordinating agents	Enterprise agent solutions aligned with supported environments	Verify maturity, capabilities, languages and deployment requirements for the chosen version.


Verify the version you will ship

APIs and features change. Verify these descriptions against current official documentation before publication or implementation. Test required capabilities in the intended runtime rather than assuming feature parity.

A Framework Selection Decision Matrix


Choose against workflow requirements, operating constraints and team needs. Record evidence for each requirement rather than assigning unsupported scores.

Requirement	What to evaluate
Workflow complexity	Sequential steps, branching, loops, long-running work
State management	Persistence between steps or sessions; recovery
Human approval	Review and resumable execution before sensitive actions
Tool integration	APIs, databases and enterprise application contracts
Multi-agent coordination	Measurable value over a single-agent baseline
Observability	Traces, execution history and actionable errors
Security	Authentication, authorization, permissions and isolation
Portability	Required model-provider flexibility
Deployment	Runtime compatibility, scaling and support
Maintainability	Team familiarity, tests, documentation and ownership



Framework selection follows requirements discovery, not the reverse.

The Integration Layer: APIs, MCP, Data, and Enterprise Tools



APIs and function calling

Use typed schemas, validate inputs and outputs, and classify failures. Authentication and authorization remain under application control.

Enterprise data

Connect relational databases, document stores and knowledge repositories. Use embeddings and vector search where useful; RAG grounds responses in retrieved evidence.

Model Context Protocol

MCP standardizes client-server access to tools, resources and prompts. It does not grant authorization or secure an integration automatically. Require trusted servers, access controls, secure transport and careful permissions.

Business applications

CRM, ERP, order management, support platforms and internal search depend on identity and access management boundaries.


Engineer the failure path

Define timeouts, rate limits, retries with backoff, idempotency, error classification and fallbacks. Check permissions at execution and maintain audit logs.

A support agent can retrieve knowledge and query order status after authentication. Refunds require a separate policy check and approval before a write operation.

Designing Reliable Agent Workflows

Prefer controlled workflows to unconstrained loops. Make allowed transitions, tool budgets, stopping conditions and recovery behavior explicit.



Seven patterns to combine deliberately

Sequential execution, branching, parallel work, routing, human-in-the-loop review, retry and recovery, and stateful execution each solve a different coordination problem.


Persist only state needed for execution and recovery. Protect sensitive fields, validate workflow state and define timeouts. Retries must not duplicate transactions.

Avoid multiple agents where deterministic code suffices, unlimited tool-call loops, model-generated plans treated as trusted instructions, unvalidated state, and retries of non-idempotent actions without safeguards.

Make AI Agents Safe, Testable, and Reliable

Quality extends beyond a convincing final answer. Evaluate decisions, permissions, tool calls and recovery as well as response content.

Evaluation area	Example test question
Task completion	Did the workflow achieve the intended outcome?
Factual grounding	Do accessible sources support the claims?
Tool selection	Was the appropriate approved tool chosen?
Argument accuracy	Were identifiers, types and values correct?
Authorization	Was unauthorized data or action blocked?
Policy compliance	Were required approvals enforced?
Failure handling	Did an outage trigger safe recovery or fallback?
Latency and cost	Did execution stay within agreed budgets?
Human handoff	Was escalation timely and sufficiently contextual?



Testing methods

Use representative, protected datasets. Include normal, ambiguous, incomplete and adversarial requests, malformed tool inputs and outages. Repeat trials, track failure rates and rerun regression tests after prompt, model, tool or workflow changes.

Security controls

Enforce least privilege and validated arguments. Treat retrieved content as untrusted; defend against injection and exfiltration. Require approval for high-impact actions, redact logs and maintain audit trails and incident response.

Guardrails reduce risk but do not guarantee safety.

Monitor What Happens After Deployment

Operations dashboard

Illustrative sample data layout · values intentionally unpopulated

Success rate

—

Completed workflows

Latency

— / —

Median / p95

Tool failures

—

Error rate

Cost per completed task

—

Human escalation rate

—

Capture the execution chain

Trace request and workflow IDs, model and tool calls, step durations, failure points, handoffs, retries and outcomes. Measure end-to-end latency, task completion, token usage, cost and evaluation scores.

Make signals actionable


Use structured events and correlation IDs with redaction, retention and access policies. Alert on rising failures, slow dependencies, cost spikes, repeated tool loops and quality drift.

LangSmith, Arize Phoenix and OpenTelemetry serve different purposes and require different integrations. OpenTelemetry supplies instrumentation and telemetry standards, not a complete AI evaluation platform.

Knowledge source changes → incomplete answers → traces reveal retrieval failures → logs identify the integration → regression tests verify the fix.

From Customer Discovery to Production Deployment

Scenario: an enterprise wants an agent to answer employee questions using internal policies and approved systems.



01 Discover

Agree objectives, users, workflows, data, permissions, failure consequences and success criteria. Deliver requirements and a risk assessment.

02 Design

Choose agent versus automation, model, framework, retrieval and tool interfaces. Document security boundaries and evaluation strategy in the architecture.

03 Build

Implement request validation, retrieval, tools, orchestration, output checks, error handling and tracing. Produce a working prototype.

04 Evaluate

Test typical, ambiguous, missing-data, unauthorized, injection and outage cases. Measure latency and cost; produce results and remediation actions.

05 Deploy

Prepare configuration, authentication, release pipeline, monitoring, rollback and operations documentation. Confirm the solution is deployment-ready.


06 Improve

Monitor outcomes, review failures, collect authorized feedback and make controlled releases. Maintain a continuous improvement plan.

A working prototype is not a release decision

Production readiness needs evidence from testing, security review, operational planning and business acceptance.

Reference Architecture: Enterprise Knowledge and Support Agent



Separate retrieval from action

The knowledge path returns evidence with source references. The business-tool path executes through controlled interfaces with explicit read/write boundaries.

Choose search deliberately

Use a vector database when semantic retrieval helps. Conventional search or structured queries may be sufficient.

Enforce access before use

Filter retrieved documents by user permissions. Check authorization before protected data or actions and require human approval before sensitive commits.

Keep business rules outside the model

Application rules validate proposed actions. Model judgment alone cannot authorize a transaction.

Trust and operations span every layer. Preserve execution evidence and recovery paths across retrieval, tool use and deployment.

Your AI Agent Deployment Readiness Checklist

Check items only when supporting evidence exists. Record unresolved gaps, owners and remediation before the release decision.

Business discovery

- ☐ Problem and desired outcome agreed
- ☐ Users and boundaries defined
- ☐ Baseline metrics recorded
- ☐ Agent justified over simpler options

Architecture and integrations

- ☐ Architecture choices documented
- ☐ Data sources and tools mapped
- ☐ API contracts and failures defined
- ☐ State, retries and timeouts specified
- ☐ Engineering ownership assigned

Security and governance

- ☐ Authentication and authorization tested
- ☐ Least privilege enforced
- ☐ Approval controls verified
- ☐ Prompt-injection defenses tested
- ☐ Logging and retention approved

Testing and evaluation

- ☐ Representative test cases ready
- ☐ Tool selection and arguments tested
- ☐ Failure recovery verified
- ☐ Regression suite runs successfully
- ☐ Acceptance thresholds satisfied

Deployment and operations

- ☐ Monitoring and alerts validated
- ☐ Rollback and incident response ready
- ☐ Controlled releases supported
- ☐ Business owner review complete
- ☐ Post-deployment evaluation planned

Release Decision

- ☐ Ready for controlled deployment
- ☐ Needs remediation
- ☐ Not approved for deployment

This checklist supports engineering review. It does not replace formal security, compliance or business approval.

Plan Your Next FDE AI Agent Project

Project / owner / review date

Define the problem

Business problem · target users · current process · desired outcome

Data sources · agent actions · actions needing human approval

Set acceptance criteria

Task success · factual accuracy · tool-call success · maximum latency · cost · security and authorization · escalation conditions

Select the architecture

Single or multi-agent · framework considered and reason · model and environment

Retrieval needs · APIs and integrations · state persistence and recovery

Document the decision

Selected architecture · alternatives considered · outstanding dependencies · approval requirements · next milestone

Record risks

Risk	Business impact	Mitigation	Owner	Validation status

Turn AI Engineering Knowledge into Practical FDE Skills

The GSDC Certified Forward Deployed Engineer Certification is for professionals who want to see how business requirements, AI capabilities and engineering practices come together in real implementations.

Connect the disciplines

- Understand the FDE role and working approach.
- Translate customer requirements into solution designs.
- Connect AI applications, integrations and implementation workflows.
- Apply practical deployment and evaluation considerations.
- Develop a structured approach to business problems.

Program coverage

The official program page describes discovery and problem framing, business translation, and three build-and-deploy projects: a rule-based service, an ML service, and a RAG plus agents service.

Ready to Explore Forward Deployed Engineering?

Build your understanding of the skills, technologies and practices used in customer-focused engineering.

Explore the GSDC Certified Forward Deployed Engineer Certification

<https://www.gsdouncil.org/certified-forward-deployed-engineer>



Globally Recognised Certification Body
190+ countries · 2,50,000+ Certified Professionals



Certified Forward Deployed Engineer



EBOOK

Extensive and exclusive Ebook created by world's experts to help you with understanding core concepts.



LEARNING MATERIALS

Get access to learning materials such as videos, ebooks, templates, and practice exams, which will help you clear the certification exam.



CREATED BY EXPERTS

GSDC certifications are created and authored by world's leading experts in the field.

LEARNING OBJECTIVE

- Gain a strong understanding of Forward Deployed Engineering concepts, roles, tools, technologies, and development practices.
- Identify business and customer challenges and design AI-driven solutions aligned with real-world requirements.
- Apply modern AI frameworks, APIs, data integration, and deployment practices to build and implement solutions within existing business systems.
- Apply testing, evaluation, security, monitoring, and optimization practices to deliver reliable, scalable, and production-ready AI solutions.

Enroll now with the code **LEARN20** To avail **20%** discount

Enroll Now

www.gsdcouncil.org 