

FDE Competency Framework v0.1

v0.1 Draft, open for comment

FDE Council (fdecouncil.com)

October 22, 2026

v0.1 DRAFT, OPEN FOR COMMENT. This is version 0.1 of the framework, written to be reviewed and challenged. It has **not** been validated by a job task analysis (JTA), it is not an assessment standard, and nobody can be assessed against it yet. Every domain, competency and indicator here may change.

Status: v0.1 Draft, open for comment · Published October 22, 2026 ·
Licence: CC BY 4.0 · fdecouncil.com

1. Purpose

Forward Deployed Engineering (FDE) is now a common job title at AI labs, data platforms and enterprise software companies, but nobody agrees on what competence in the role looks like. Each employer writes its own job description, and the early training programs that have appeared each define the role their own way (see Section 4).

This framework aims to give a **shared, vendor-neutral way to describe what competent FDE practice looks like** at three levels of proficiency, so that:

- **practitioners** can see where they are now and plan what to learn next;
- **employers** can write clearer job profiles, interview loops and career ladders;
- **training providers** can map their courses to a common reference; and
- the community can work toward a validated, assessable open standard, built from a later, job-task-analysis-validated version of it.

2. Scope

In scope: the work of an engineer who embeds with a customer (internal or external) to take a technical capability, now usually AI/LLM-based, from an unclear problem to a production outcome the customer owns. The framework covers discovery, design, building, deployment, adoption, handover, and sending what was learned back to the vendor's product.

Out of scope (for v0.1):

- Roles that mainly do pre-sales demonstration, account management, or advisory-only consulting. The FDE sources reviewed explicitly set the role apart from these (Section 4).
- People-management competencies for managers of FDE teams. These are planned for a separate “FDE Leadership” annex.
- Any particular vendor’s platform, model, cloud or tool. The framework describes capabilities, not products.

Industry-neutral by design: Domains 1 to 7 are the **core** and apply to every FDE, in any industry. v0.1 contains **no industry-specific track**. Industry-specific tracks may be added in later versions if the community asks for them (see Section 6, “Future industry tracks”).

3. Design principles

1. **Vendor-neutral.** No competency requires a specific product. Indicators describe observable behaviour and outcomes, not tool names.
2. **Outcome-anchored.** Every source we reviewed measures FDEs by production results, not by demos or hours. The indicators reflect that.
3. **Observable and assessable.** Each behavioural indicator is written so that an assessor could in principle collect evidence of it, for example through a work sample, a portfolio, a structured interview or a simulation.
4. **Designed with ISO/IEC 17024 in mind.** ISO/IEC 17024 is the international standard for bodies that certify persons. FDE Council is not such a body; the framework is written so that a future, validated version could meet that standard’s expectations. The current edition is **ISO/IEC 17024:2026** (Edition 3, published March 2026), which replaced the 2012 edition and added requirements on the use of AI in certification processes ([ISO](#)). Under the standard, a certification scheme must rest on a documented job or practice analysis that identifies the tasks and the competence each task needs, and that analysis must be kept up to date and reviewed systematically ([ISO/CASCO overview of certification of persons](#); [ISO/IEC 17024:2012 text, law.resource.org](#)). **v0.1 is a hypothesis about that analysis, not a substitute for it.** A formal job task analysis, with a practitioner survey and subject-matter-expert panels, is planned before any version is used for assessment (Section 9).
5. **Separation of standard-setting from training.** FDE Council intends to keep the framework (and any future exam) separate from any training offering, including any training of its own, so that no training provider, including any affiliated with FDE Council’s founders, controls what the framework requires.
6. **Specialisation is additive.** Specialty tracks build on the core and never replace it.

7. **Open by default.** Licensed so anyone can use and adapt it (Section 10).

4. Research basis

This draft was informed by public sources reviewed on **October 22, 2026**: FDE job postings and engineering blogs from AI labs, data platforms and enterprise software companies; practitioner writing on the role; and the published outlines of early FDE training programs. **None of their wording is reused in the framework itself**, and the framework is not endorsed by any organisation whose material was reviewed.

Themes that recur across the sources (and therefore became core domains): embedded, end-to-end ownership; discovery and problem framing under ambiguity; hands-on production engineering, including LLM and agent systems; working with messy or governed customer data; security and compliance as part of delivery; evaluation tied to the customer's acceptance criteria; adoption, handover and a planned exit; and feeding field learning back into the product.

Gaps this framework tries to fill:

- **Proficiency levels defined by observable behaviour.** The early programs reviewed tend to define levels by course or market segment rather than by what the person does.
- **An explicit commercial-judgment and engagement-design domain.** Practitioner commentary on entry, exit and kill criteria for engagements points to this need.
- **A vendor-neutral, openly licensed reference** that is kept separate from any one training product.

5. Structure of the framework

- **7 core domains**, all industry-neutral. No industry-specific tracks in v0.1.
- Each domain has a **definition**, **4 to 6 competencies**, and **behavioural indicators at three proficiency levels**:

Level	Typical scope (indicative, not a requirement)
Associate	Delivers well-defined parts of a deployment under guidance. Is reliable on scoped tasks and asks for help early.
Practitioner	Owns a deployment end to end for one customer or workstream, working independently in ambiguous

Level	Typical scope (indicative, not a requirement)
Lead	Owns several or high-stakes deployments, sets the approach for others, shapes the vendor's product and playbooks, and is trusted by customer executives.

Years of experience are deliberately **not** used to define levels. If years are used at all, the job task analysis should decide that.

6. The domains

Domain 1: Customer Discovery and Problem Framing

Definition: Turning a customer's unclear, often contested request into a specific, valuable and feasible problem statement with agreed success measures, before significant build effort is committed.

Competencies

- 1.1 Plans and runs discovery with users, operators, technical owners and decision-makers.
- 1.2 Maps the current workflow, systems of record and pain points as they really are, not as documented.
- 1.3 Separates the stated request from the underlying need, and identifies who actually decides.
- 1.4 Frames the problem as testable hypotheses, with baseline metrics and target outcomes.
- 1.5 Assesses feasibility and readiness (data, access, sponsorship, risk), and recommends go, reshape or stop.

Associate	Practitioner	Lead
Prepares and documents discovery sessions using a provided template; captures workflows and pain points accurately.	Runs discovery independently across business and technical stakeholders, and produces a one-page problem statement with baseline, target metric and assumptions that the sponsor signs off.	Designs discovery for multi-team or executive-level engagements, resolves conflicting stakeholder goals, and reframes or declines engagements that fail readiness criteria.
Flags gaps between what	Identifies the real	Builds reusable

Associate	Practitioner	Lead
users say and what the systems show.	decision-maker and the operative constraint, and adjusts scope accordingly.	discovery playbooks and qualification criteria that other FDEs adopt.
Asks clarifying questions rather than assuming requirements.	Converts findings into prioritised, testable hypotheses.	Coaches others on handling hostile or politically sensitive discovery.

Domain 2: Solution Architecture and Systems Integration

Definition: Designing a solution that fits the customer’s real environment (legacy systems, identity, network limits, existing platforms) and choosing the simplest architecture that will deliver the outcome reliably.

Competencies

- 2.1 Designs end-to-end architectures across the vendor platform, customer systems and third-party services.
- 2.2 Integrates with customer systems through APIs, events, files and databases, including legacy and undocumented interfaces.
- 2.3 Makes and documents build-vs-configure-vs-extend-the-product decisions and trade-offs.
- 2.4 Designs for the customer’s identity, access, network and data-residency constraints.
- 2.5 Chooses the right kind of solution, including deterministic code, classical ML, LLM or agent, or no software at all.

Associate	Practitioner	Lead
Implements integrations against a given design, and reads and navigates unfamiliar customer codebases.	Produces an architecture and decision record that the customer’s architects and security team accept.	Sets reference architectures and integration patterns that are reused across deployments.
Documents interfaces, assumptions and failure modes encountered.	Chooses between custom build, configuration and product extension, with explicit cost, speed and maintainability trade-offs.	Negotiates architecture with senior customer technical leaders and the vendor’s product and engineering teams.
Raises integration risks early.	Designs to fit identity, access and data-	Decides when the right answer is “don’t build

Associate	Practitioner	Lead
	residency constraints without unsafe workarounds.	this," and makes that case.

Domain 3: AI/LLM Application Engineering and Evaluation

Definition: Building production-quality applications on AI models (including retrieval, tool use and agentic workflows) and proving their quality with evaluations the customer trusts.

Competencies

- 3.1 Writes production-grade code (tested, reviewed, maintainable) in at least one mainstream language.
- 3.2 Builds LLM-based components: prompting and structured outputs, retrieval over customer content, tool and function calling, and multi-step or agent workflows.
- 3.3 Designs evaluations: test sets drawn from real cases, acceptance thresholds agreed with the customer, regression checks, and human review where needed.
- 3.4 Analyses errors and improves quality, cost and latency systematically.
- 3.5 Applies guardrails and keeps a human in the loop to match the risk of the use case.
- 3.6 Keeps up with changing model capabilities and judges when a change matters to a deployment.

Associate	Practitioner	Lead
Builds and tests well-specified components, and follows team code and evaluation conventions.	Owens the AI system for a deployment, from prototype to production, against agreed acceptance criteria.	Defines evaluation methodology and quality bars for a portfolio of deployments or a vertical.
Runs existing evaluation suites and reports results accurately, including failures.	Builds an evaluation set with the customer's domain experts and uses it to make go/no-go decisions.	Turns recurring field failures into product or model feedback that the vendor acts on.
Explains how model behaviour affects user experience in their component.	Diagnoses failure modes (retrieval, reasoning, tool use, data) and fixes the cause, not the symptom.	Sets policy on where autonomy is and isn't appropriate in high-risk workflows.

Domain 4: Data Engineering and Data Readiness

Definition: Getting the customer's real data (messy, siloed, governed, sometimes legacy) into a state where a solution can use it safely and reliably.

Competencies

- 4.1 Profiles and assesses data quality, lineage and fitness for purpose.
- 4.2 Builds and operates pipelines (batch and streaming as needed) and data models.
- 4.3 Works within data governance: classification, consent, retention, residency, and least-privilege access.
- 4.4 Prepares unstructured content (documents, images, transcripts) for retrieval and extraction.
- 4.5 Plans migration and reconciliation from legacy sources, with checks the customer can verify.

Associate	Practitioner	Lead
Writes queries and transformations to spec, and documents data issues found.	Delivers a data readiness assessment and the pipelines a deployment needs, with tests and monitoring.	Sets data-readiness entry criteria for engagements, and designs data approaches for complex or regulated estates.
Handles sensitive data according to the provided rules.	Designs access and handling that satisfy the customer's governance owners.	Resolves data-ownership and governance disputes between customer teams.
Reconciles outputs against source samples.	Plans and executes migrations with reconciliation reports the business signs off.	Builds reusable data patterns adopted across deployments.

Domain 5: Deployment, Operations, Security and Responsible AI

Definition: Shipping solutions into the customer's production environment and keeping them secure, observable, compliant and dependable, with security and compliance treated as part of delivery rather than an afterthought.

Competencies

- 5.1 Packages and deploys software through repeatable pipelines (CI/CD and infrastructure as code) in customer-controlled environments, including restricted or air-gapped ones.

- 5.2 Instruments systems for monitoring, logging, alerting and AI-specific observability (quality drift, cost, latency).
- 5.3 Applies secure engineering: secrets handling, access control, vulnerability management, privacy, and threat modelling for AI systems (for example prompt injection and data exfiltration through tools).
- 5.4 Takes solutions through customer security, risk and compliance reviews.
- 5.5 Documents and runs responsible-AI controls appropriate to the use case: transparency, auditability, bias checks, and incident response.

Associate	Practitioner	Lead
Deploys using established pipelines, and follows secure-coding and secrets practices.	Owens production readiness for a deployment: runbooks, monitoring, rollback, and on-call handover.	Defines deployment and security standards for FDE teams, and leads response to serious incidents.
Fixes issues found in security scans or reviews.	Prepares evidence for and passes customer security, risk and compliance reviews.	Engages customer CISOs, risk and compliance leaders as a peer, and shapes the vendor's security posture for field deployments.
Reports incidents promptly and accurately.	Threat-models AI-specific risks and puts in proportionate controls.	Sets proportionate responsible-AI controls for high-impact use cases.

Domain 6: Stakeholder Communication, Adoption and Change Management

Definition: Building trust across the customer organisation and making sure the solution is actually adopted, by explaining clearly, delivering bad news early, and managing the human side of change.

Competencies

- 6.1 Adapts technical communication to executives, operators, engineers and risk owners.
- 6.2 Maintains relationships and trust, including when delivering bad news or pushing back.
- 6.3 Plans and supports adoption: training, workflow redesign, champions, and feedback loops.
- 6.4 Measures adoption and impact and reports them honestly.

- 6.5 Runs demos and readouts that tie technical progress to business outcomes, without overselling.

Associate	Practitioner	Lead
Writes clear status updates and documentation for technical peers.	Runs regular stakeholder readouts and keeps the sponsor's confidence through setbacks.	Is trusted by customer executives as an advisor, and handles escalations and conflicts across organisations.
Supports user training and gathers user feedback.	Designs and executes the adoption plan, and measures usage and impact against the baseline.	Leads change programmes that span multiple teams or business units.
Escalates concerns early through the right channel.	Delivers bad news early, with options.	Mentors others in executive communication and difficult conversations.

Domain 7: Delivery Leadership and Commercial Judgment

Definition: Running an engagement as a bounded, outcome-focused investment: scoping it, sequencing it, knowing when to stop, handing it over cleanly, and turning what was learned into reusable value for the vendor.

Competencies

- 7.1 Scopes and sequences work to show value early (smallest proof first), and manages trade-offs between scope, speed and quality.
- 7.2 Defines and enforces entry, exit and stop criteria for an engagement.
- 7.3 Understands the commercial model of the engagement (cost to deliver, value delivered, intellectual-property boundaries, expansion potential) and acts accordingly.
- 7.4 Plans handover so the customer can run the solution without the FDE.
- 7.5 Codifies field learning into playbooks, reusable components and product feedback.

Associate	Practitioner	Lead
Estimates and delivers own tasks reliably, and makes blockers visible.	Owens the engagement plan, delivers a working proof early, and adjusts scope to protect the outcome.	Shapes the engagement portfolio and investment decisions, and sets qualification and stop criteria.
Documents what was built so others can	Executes a planned handover and verifies the	Structures engagements around intellectual-

Associate	Practitioner	Lead
maintain it.	customer team can operate independently.	property boundaries and commercial outcomes, together with the vendor's sales and legal teams.
Contributes field observations to team retrospectives.	Turns at least one deployment pattern into a reusable asset or a product request.	Builds the feedback loop from field to product that shifts the vendor's roadmap.

Future industry tracks (not part of v0.1)

v0.1 is deliberately industry-neutral: every competency above applies to FDE work in any sector. Industry-specific tracks may come later, but only if they are proposed and driven by the community, for example when enough practitioners in one sector say the core domains miss something essential to their work. Any future track would build on the core (Principle 6), never replace it, and would go through the same open review as the core framework.

7. How the levels relate across domains (guidance)

- A person does **not** need the same level in every domain. A typical Practitioner profile might be Practitioner in Domains 1, 3, 5, 6 and 7 and Associate in Domain 4.
- Any future assessment built on a validated version would define a **minimum profile** per level. That will come from the job task analysis, not from this draft.

8. Example uses (non-normative)

- **Self-assessment:** rate yourself per domain and pick two indicators to work toward.
- **Hiring:** pick the 8 to 12 indicators that matter most for a role, and build interview questions or work samples around them.
- **Career ladders:** align internal levels to Associate, Practitioner and Lead behaviours.
- **Curriculum mapping:** training providers can show which competencies their course covers. FDE Council does **not** endorse courses on that basis.

9. Validation roadmap (working toward ISO/IEC 17024 alignment)

1. **Community review of v0.1** (open comment period).
2. **Job task analysis:** a structured survey of practising FDEs and FDE leaders, rating each task and competency for importance and frequency, plus subject-matter-expert workshops across company types and industries. The first round will sample practitioners in North America (US, Canada and Mexico); other regions can follow. The planned *State of FDE 2026* survey can provide early input, but it is not a substitute for a formal job task analysis.
3. **v0.2 to v1.0:** revise the domains and indicators, and publish the job task analysis method and results.
4. **Assessment design (later, separate, only after validation):** competence requirements per level, assessment methods, psychometric review, and separately governed oversight, so standard-setting and training stay apart as ISO/IEC 17024 expects.

10. Versioning and licensing

Versioning. Releases use the form *MAJOR.MINOR*. 0.x releases are drafts for review. 1.0 will be the first job-task-analysis-validated release. Every release will carry a change log, and earlier versions will stay available.

Licence: Creative Commons Attribution 4.0 International (CC BY 4.0) ([license deed](https://creativecommons.org/licenses/by/4.0/)). Why this licence:

- **Adoption is the goal.** CC BY 4.0 lets anyone share and adapt the material for any purpose, including commercially, as long as they credit the source. Employers can paste it into job ladders and trainers can map courses to it without asking permission, which helps it spread as a common reference.
- **Credit is preserved.** Users must give appropriate credit, link the licence and indicate changes. The licence also says users must not suggest the licensor endorses them, which helps protect FDE Council's neutrality.
- **It keeps the text open and the name separate.** The framework text is open. The FDE Council name and logo are not covered by the licence; a CC licence does not grant trademark rights.

Notice: FDE Competency Framework v0.1 © 2026 FDE Council, licensed under CC BY 4.0 (<https://creativecommons.org/licenses/by/4.0/>). Attribution: "FDE Competency Framework v0.1, FDE Council, fdecouncil.com".

11. How to give feedback

During the v0.1 review period, FDE Council plans to accept feedback through:

- the feedback form at fdecouncil.com/join/?topic=framework;
- the FDE Council LinkedIn group discussion threads; and
- invited reviewer sessions with practitioners and FDE leaders.

Helpful feedback names the domain and competency number (for example "3.3"), says what's missing, wrong or unclear, suggests alternative wording, and notes your context (company type, region, level). Please do not share confidential customer information.

12. What v0.2 needs

- ☐ Feedback from at least 25 to 50 practising FDEs and 10 FDE leaders, across AI labs, data and enterprise platforms, startups and consultancies (target; to be confirmed).
 - ☐ Job task analysis design: task inventory, rating scales, sampling plan, SME panel composition.
 - ☐ A decision on whether "Deployment Strategist"-type (Palantir "Echo") roles are in scope or get their own profile.
 - ☐ A simple, open process for the community to propose an industry-specific track, and the evidence a proposal would need (none is planned for v0.2 unless the community asks).
 - ☐ Indicators for the planned FDE Leadership annex (hiring, team design, engagement portfolio).
 - ☐ An evidence guide: example artefacts that would show each indicator.
 - ☐ Accessibility and plain-language review, plus translation needs.
 - ☐ A conflict-of-interest and governance note for the framework committee.
 - ☐ Mapping tables against public job postings (and, if useful, against other programs' published outlines) to test coverage, without copying their text.
-

Disclaimer: This draft is for discussion. It is not legal, regulatory or compliance advice, and is not endorsed by any organisation cited.