FDE Council Professional Capability Framework

Version 0.4 (draft, open for comment)

Published November 20, 2026. Licensed under CC BY 4.0. Current version and comment form: https://www.fdecouncil.com/framework/

This PDF contains the industry-neutral core framework only. Optional industry tracks are published separately; the first is the Mortgage and Fintech track: https://www.fdecouncil.com/framework/mortgage-fintech/

This is a draft. It has not been validated by a job task analysis, it is not a certification standard, and nobody can be assessed or certified against it.

> **This is a draft**
>
> This is the current draft for community review. It has not been validated by a job task analysis, it is not a certification standard, and nobody can be assessed or certified against it. The behavioral indicators are illustrative, for self-reflection and community review. They are not pass criteria. Earlier drafts stay at v0.3 and v0.1. The free self-assessment lets you rate yourself across these 11 domains, for self-reflection only. It does not score AI and agent engineering as a separate domain.

New to the role? Start with what a Forward Deployed Engineer is.

What changed from v0.3

The 11 core domains are the same domains, in the same order, with the same descriptions. No domain was added, removed or merged.

Behavioral indicators for every core domain. Each domain now has illustrative indicators at Associate, Practitioner and Lead, written as things a person does or produces that can be seen in work evidence. They follow the style of v0.1 but are written for the 11 domains. v0.1's 7 domains are not restored.

A new optional Industry tracks layer. The first track is Mortgage and Fintech. A track never replaces a core domain, and it is not required of FDEs who do not work in that industry.

A cross-cutting section on ethics and professional conduct. It applies across all domains and goes beyond the controls in Domain 8. It is not a twelfth domain.

Kept unchanged: the three levels and their scope, the three context-required capabilities, the two example advanced specializations, the design principles, the validation roadmap, the license and the versioning rules.

How the framework is structured

The framework has four layers and one cross-cutting expectation:

11 mandatory core domains. Every FDE should develop capability in all eleven. They are industry-neutral.

Context-required capabilities, used when a deployment calls for them. They do not replace a core domain.

Advanced specializations that build on the core. They are not required of every FDE.

Industry tracks (new in v0.4), which are optional. They describe what the core domains look like in one industry, build on the core and never replace it.

Ethics and professional conduct (new in v0.4), which applies in every domain, layer and level.

The core domains are competency domains, not sequential project phases. An FDE draws on several of them at once, and returns to them throughout a deployment.

What are the FDE proficiency levels?

The expected independence and depth in each domain increase with experience. Three indicative levels describe that progression:

Level	Typical scope (indicative)
Associate	Delivers well-defined parts of a deployment under guidance. Reliable on scoped tasks; asks for help early.
Practitioner	Owns a deployment end to end for one customer or workstream, working independently in ambiguous environments.
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. Trust with customer executives is a description of Lead scope, not a separate skill. The skill of adapting communication for executives sits in Domain 9.

How to read the behavioral indicators

Each indicator describes something a person does or produces that another person could see in real work: a document, a decision record, a test result, a handover, a conversation outcome.

Indicators are cumulative. A Practitioner is expected to show the Associate indicators as well, and a Lead the Practitioner indicators.

Indicators are illustrative, for self-reflection and community review. They are a starting hypothesis for the job task analysis. They are not pass criteria, and no one should be scored, hired, graded or certified against them.

Not every indicator applies to every deployment. Use judgment about context.

DRAFT V0.4

The 11 mandatory core domains of Forward Deployed Engineering

Every FDE should develop capability in all eleven, in any industry. Open a domain to see what an FDE must be able to do and the illustrative behavioral indicators at each level. Domain descriptions are the v0.3 wording, unchanged. For how the domains show up in a real working week, read a day in the life of a Forward Deployed Engineer.

1 Discovery and Problem Framing −

Enter an unfamiliar environment, learn the domain, observe actual work, distinguish symptoms from root causes, identify who decides, and name the problem worth solving. Recommend go, reshape, or stop when data, access, sponsorship, or feasibility is missing.

What an FDE must be able to do. Description unchanged from v0.3.

ASSOCIATE

Prepares and runs discovery sessions from a provided guide, and writes up notes the same day that capture workflows, pain points and open questions.

Observes real users doing the work, and records where actual practice differs from documented process.

Asks clarifying questions and records stated assumptions instead of filling gaps with guesses.

PRACTITIONER

Runs discovery independently across business and technical stakeholders, and produces a one-page problem statement with a baseline, a target measure and named assumptions that the sponsor signs off.

Names the person who decides, the operative constraint, and the root cause, with the evidence for each.

Checks data, access, sponsorship and feasibility before build starts, and writes a go, reshape or stop recommendation with reasons.

Converts findings into a short list of prioritized, testable hypotheses.

LEAD

Designs discovery for multi-team or executive engagements, and resolves conflicting goals between stakeholders into one agreed problem statement.

Reshapes or declines engagements that fail readiness checks, and explains the decision to customer executives and to their own leadership.

Publishes discovery guides and readiness criteria that other FDEs adopt.

Coaches other FDEs through politically sensitive or hostile discovery.

2 Workflow Design and Value Definition −

Design an improved way of working, including user experience, human decisions, exception handling, measurable benefits, and conditions for success.

What an FDE must be able to do. Description unchanged from v0.3.

ASSOCIATE

Maps the current workflow step by step, including hand-offs, wait times and the systems touched at each step.

Drafts screens, forms or process steps for an agreed design, and tests them with at least one real user.

Lists exceptions and edge cases seen in real work, and flags those the design does not yet handle.

PRACTITIONER

Produces a future-state workflow that shows which decisions stay with people, which are automated, and how exceptions are routed.

Defines measurable benefits with the customer (time, cost, error rate, throughput or risk) and records the baseline before build.

Writes the conditions for success and the conditions that would mean the change should not go ahead.

Revises the design based on user testing, and records what changed and why.

LEAD

Designs workflow changes that span several teams or business units, and gets the owners of each step to agree to them.

Sets a value model that customer finance or operations leaders accept as the basis for judging the deployment.

Sets workflow design patterns and value templates that other FDEs reuse.

3 Solution Architecture and Enterprise Integration −

Design a practical solution within existing systems and constraints. Choose whether to build, configure, or extend the product, including when no new software is the right answer. Make defensible architecture choices and diagnose problems across system boundaries.

What an FDE must be able to do. Description unchanged from v0.3.

ASSOCIATE

Implements integrations against a given design, and documents the interfaces, assumptions and failure modes found.

Reads and navigates unfamiliar customer systems and codebases, and asks for access through the proper channel.

Raises integration risks to the lead early, with the evidence found.

PRACTITIONER

Produces an architecture and decision record that the customer's architects and security team accept.

Records the build, configure or extend decision with explicit cost, speed and maintainability trade-offs, including the option of building nothing new.

Designs to fit the customer's identity, access, network and data-residency constraints without unsafe workarounds.

Diagnoses failures that cross system boundaries, and shows the root cause with logs or traces.

LEAD

Sets reference architectures and integration patterns that are reused across deployments.

Negotiates architecture with senior customer technical leaders and with the vendor's product and engineering teams.

Makes and defends the case that the right answer is not to build, when that is true.

4 Data Foundations and Engineering

Understand what data means, establish permitted access, assess quality, and build dependable data models, transformations, and pipelines.

What an FDE must be able to do. Description unchanged from v0.3.

ASSOCIATE

Writes queries and transformations to a specification, with tests, and documents data issues found.

Handles sensitive data according to the rules provided, and asks before using data in a new way.

Reconciles outputs against source samples, and reports differences instead of hiding them.

PRACTITIONER

Delivers a data readiness assessment covering meaning, quality, lineage, ownership and permitted access, for the data the deployment needs.

Builds the data models and pipelines the deployment needs, with tests, monitoring and documented owners.

Confirms permitted access with the data owner in writing before use, and designs handling that satisfies the customer's governance owners.

Plans migrations from legacy sources, with reconciliation reports the business signs off.

LEAD

Sets data readiness entry criteria for engagements, and designs data approaches for complex or regulated estates.

Resolves data ownership and governance disputes between customer teams.

Builds reusable data patterns and quality checks that are adopted across deployments.

5 Software Engineering and Rapid Experimentation −

Read and modify unfamiliar code, build working solutions, test risky assumptions quickly, debug effectively, and maintain appropriate engineering quality.

What an FDE must be able to do. Description unchanged from v0.3.

ASSOCIATE

Ships well-specified changes in at least one mainstream language, with tests, and follows the team's code review and version control practice.

Reads unfamiliar code well enough to make a safe, small change and explain what it affects.

Debugs with logs, tests and a written hypothesis, and records what was tried.

PRACTITIONER

Builds a working prototype that tests the riskiest assumption first, and records what the result showed.

Owns the code for a deployment from prototype to production, at an engineering quality matched to its risk.

Decides, and writes down, which prototype code is thrown away and which is hardened for production.

Leaves code, tests and documentation that the customer's engineers can maintain without the FDE.

LEAD

Sets engineering quality bars for field code that balance speed with risk across deployments.

Builds reusable components or tools that other FDEs adopt, and that reduce time to a first working result.

Reviews other FDEs' code and experiment designs, and raises the quality of their work.

6 Verification, Evaluation and Outcome Measurement

Establish whether the solution works, how it fails, whether users can use it, and whether it improves the target workflow, with evidence appropriate to the risk. When the system uses a model, include an agreed test set, allowance for answers that are not deterministic, and a regression check after the model changes.

What an FDE must be able to do. Description unchanged from v0.3.

ASSOCIATE

Runs existing test and evaluation suites, and reports results accurately, including failures.

Writes test cases from real examples supplied by users or domain experts.

Records known failure modes and the conditions that trigger them.

PRACTITIONER

Agrees acceptance measures with the customer before build, and uses them to make a go or no-go decision.

Builds an evaluation set with the customer's domain experts. When the system uses a model, the set allows for answers that are not deterministic and is rerun as a regression check after a model change.

Runs usability checks with real users, and fixes what blocks them.

Measures the outcome against the Domain 2 baseline after go-live, and reports it honestly, including when it falls short.

LEAD

Defines evaluation methods and quality bars for a portfolio of deployments or for a vertical.

Sets the level of evidence required in proportion to risk, and holds a release when the evidence is not there.

Turns recurring field failures into evaluation cases and product feedback that the vendor acts on.

7 Production Deployment, Reliability and Operations

Deploy safely, monitor performance, recover from failures, and manage operating costs. When the system calls a model, operating cost includes token and inference cost, and the runbook names the model version, watches for version drift, and says what replaces the model if the provider retires it. Leave the customer able to operate the system: runbooks, knowledge transfer, support, and a named owner.

What an FDE must be able to do. Description unchanged from v0.3.

ASSOCIATE

Deploys using established pipelines and change processes, and confirms the release works after deployment.

Adds logging, metrics and alerts to components as specified.

Reports incidents promptly and accurately, with a timeline of what happened.

PRACTITIONER

Owns production readiness for a deployment: deployment pipeline, monitoring, alerting, rollback and a tested recovery path.

Writes a runbook that names the system owner and, when a model is used, the model version, how version drift is watched, and what replaces the model if the provider retires it.

Tracks operating cost, including token and inference cost when a model is used, against the budget agreed with the customer.

Completes a handover in which the customer team operates the system without the FDE, and records the named owner at exit.

LEAD

Defines deployment, reliability and handover standards for FDE teams.

Leads the response to serious incidents across vendor and customer teams, and publishes a blameless review with follow-up actions.

Sets cost and capacity guardrails that keep deployments affordable to operate at scale.

8 Security, Privacy, Safety and Governance

Identify risks, implement appropriate controls, respect authorization and data boundaries, secure required approvals, and escalate work beyond personal competence or authority. Apply proportionate responsible-use controls: transparency, a record of how decisions were made, and human review where a person must be able to intervene.

What an FDE must be able to do. Description unchanged from v0.3.

ASSOCIATE

Follows secure coding and secrets handling practice, and fixes issues found in scans or reviews.

Works only within granted access and data boundaries, and asks before crossing one.

Escalates a suspected security, privacy or safety issue the same day, through the agreed channel.

PRACTITIONER

Produces a risk assessment for the deployment, with proportionate controls, including threat modeling for AI-specific risks when a model or agent is used.

Prepares the evidence for the customer's security, privacy, risk and compliance reviews, and gets the required approvals before go-live.

Implements responsible-use controls matched to the use case: transparency to affected people, a record of how decisions were made, and human review where a person must be able to intervene.

Names the work that is beyond personal competence or authority, and hands it to the right owner in writing.

LEAD

Defines security and governance standards for field deployments, and leads the response to serious security or safety incidents.

Works with customer security, risk and compliance leaders as a peer, and resolves disagreements about acceptable risk.

Sets responsible-use controls for high-impact use cases, and holds or stops a deployment when the controls are not in place.

9 Embedded Collaboration, Trust and Adoption

Work alongside customer teams. Adapt communication for executives, operators, and engineers. Handle disagreement and resistance, support changes to roles and workflows, build the customer's capability, and help the solution become part of normal operations.

What an FDE must be able to do. Description unchanged from v0.3.

ASSOCIATE

Writes clear status updates and documentation for technical peers and users.

Supports user training and collects user feedback in a form the team can act on.

Escalates concerns early through the right channel.

PRACTITIONER

Runs regular readouts adapted to executives, operators and engineers, and keeps the sponsor's confidence through setbacks.

Delivers bad news early, with options and a recommendation.

Designs and runs the adoption plan, including training, champions and changes to roles, and measures usage against the plan.

Builds the customer team's capability so that they can extend the solution, and records what they can now do on their own.

LEAD

Is relied on by customer executives as an advisor, and resolves escalations and conflicts across organizations.

Leads change across several teams or business units, including resistance to changes in roles and workflows.

Mentors other FDEs in executive communication and difficult conversations.

10 Delivery and Commercial Judgment

Manage scope, priorities, dependencies, uncertainty, costs, and commitments. Make sound delivery trade-offs without overpromising or bypassing controls. Know when to narrow or stop, and understand the commercial model well enough to judge cost against value.

What an FDE must be able to do. Description unchanged from v0.3.

ASSOCIATE

Estimates own tasks, delivers them reliably, and makes blockers and slips visible early.

Keeps a written record of commitments made to the customer and their status.

Raises scope requests to the lead instead of agreeing to them alone.

PRACTITIONER

Owns the engagement plan, delivers a working result early, and adjusts scope to protect the agreed outcome.

Records trade-offs between scope, speed, cost and quality, and gets the sponsor’s agreement to them.

Writes entry, exit and stop criteria for the engagement, and recommends narrowing or stopping when they are met.

Explains the cost to deliver and the value delivered in terms the sponsor and the vendor’s account team both accept.

LEAD

Shapes which engagements are taken on, and sets qualification and stop criteria used across a portfolio.

Structures engagements around intellectual property boundaries and commercial outcomes, with the vendor's sales and legal teams.

Stops or reshapes an engagement when cost exceeds value, and explains the decision to customer and vendor leadership.

11 Field Learning, Product Feedback and Reuse

Convert field experience into actionable feedback, reusable assets, and improvements to future deployments while respecting confidentiality, ownership, and product boundaries.

What an FDE must be able to do. Description unchanged from v0.3.

ASSOCIATE

Logs product issues and feature gaps with steps to reproduce and the customer impact.

Documents what was built so that others can find and reuse it.

Removes customer-confidential details before sharing any lesson outside the engagement.

PRACTITIONER

Writes product feedback with evidence, frequency and business impact that the product team can prioritize.

Turns deployment work into reusable assets (templates, components, runbooks) with clear ownership and permission to reuse.

Runs a retrospective at the end of a deployment, and records what future deployments should do differently.

Checks intellectual property and confidentiality boundaries with the right owner before reusing customer-derived work.

LEAD

Turns repeated field lessons into product, playbook or roadmap changes that the vendor adopts.

Runs a feedback loop between FDEs and product teams, and tracks which feedback was acted on.

Sets the rules for what may and may not be reused across customers.

Design, build, verify, operate

Domains 2, 5, 6 and 7 answer four different questions:

Design (Domain 2): What should improve, and how should the workflow operate?

Build (Domain 5): Can we construct the proposed solution?

Verify (Domain 6): What evidence shows it works and delivers the intended benefit?

Operate (Domain 7): Can it continue working safely and reliably after the FDE leaves?

Evaluation is a requirement for all FDE work, including systems that do not use AI. When the system uses a model, Domain 6 also includes an agreed test set, allowance for answers that are not deterministic, and a regression check after the model changes. For a practical way to judge a live deployment, read how to measure whether an FDE deployment succeeded.

Context-required capabilities

Use a context-required capability when the deployment calls for it. Skip it when it does not. These capabilities sit on top of the core domains. They do not replace Domain 5, Domain 6, or Domain 8.
Unchanged from v0.3.

AI and agent engineering

An agent, in this draft, is software that takes a goal and carries it out over more than one step, usually by calling tools or a model. A language model is in this same capability when the deployment uses one, including prompting and structured outputs. Many FDE deployments never use either. When a deployment does, the FDE has to be able to build or integrate it. For how this looks on an engagement, read what an AI Forward Deployed Engineer does.

That means six things:

Prompting and structured outputs, when a language model has to return something the rest of the system can use.

Retrieval over the customer's own content, when the agent needs facts the model was not trained on.

Tool calling, including a multi-step agent that chooses the next tool from what the last step returned.

A person who can intervene where the risk requires it, before the agent acts on a customer system or a customer decision.

Mitigate prompt injection. Control customer data leaving through a tool. Prompt injection cannot be fully prevented.

The model version in production: watch for version drift, and say what replaces the model if the provider retires it.

The surrounding engineering stays in Domain 5. Whether it works, and whether the workflow improved, stays in Domain 6. For a model, that includes an agreed test set, allowance for answers that are not deterministic, and a regression check after the model changes. Token and inference cost stays in Domain 7. Security, privacy, and approvals stay in Domain 8. The row below is this same capability.

Context-required capabilities in draft v0.4

Capability	When it applies	What an FDE must be able to do
AI and agent engineering	The deployment uses models, retrieval, tool calling, or multi-step agents.	Build or integrate the model or agent: prompting, structured outputs, retrieval over customer content, tool calling, and multi-step workflows. Keep a person able to intervene where the risk requires it. Mitigate prompt injection, and control customer data leaving through a tool. Name the model version, watch for version drift, and plan for the provider retiring the model. Whether it works remains Domain 6, including an agreed test set, allowance for non-deterministic answers, and a regression check after a model change. Token and inference cost remains Domain 7.
Restricted and air-gapped environments	The customer limits installation, network access, or ongoing operation, including isolation from the public internet.	Deploy, monitor, and hand over the solution inside those limits, without a path that assumes open network access.
Regulated-data constraints	The data brings residency, consent, retention, or sector rules beyond ordinary access control.	Work inside those rules, say which rule applies, and escalate what is outside personal authority. Ordinary access control remains Domain 4 and Domain 8.

Advanced specializations

An advanced specialization builds on the core. It never replaces a core domain. Since v0.3, the framework names two examples so the layer is no longer only a label. It is not a closed list, and it is not required of every FDE. Unchanged from v0.3.

Example advanced specializations in draft v0.4

Specialization	Builds on	What it adds
Portfolio practice leadership	Domains 9, 10, and 11, at Lead scope	Owns several deployments, sets methods other FDEs reuse, and turns repeated field lessons into product and playbook changes.
AI evaluation lead	Domain 6 and AI and agent engineering, at Lead scope	Owns how model-based systems are tested: the agreed test set, allowance for answers that vary, and the regression check when a model changes. Not required on deployments that do not use a model.

NEW IN V0.4

Industry tracks

Industry tracks are optional: a track builds on the 11 core domains, never replaces a core domain, and is not required of FDEs who do not work in that industry. Each track describes what the core domains look like in one industry. The first track is the Mortgage and Fintech track, published on its own page.

NEW IN V0.4

Ethics and professional conduct for Forward Deployed Engineers

These expectations apply across every domain, layer and level. They are not a twelfth domain, and they go beyond the security, privacy and governance controls in Domain 8. Domain 8 is about the controls a system needs. This section is about how the FDE behaves.

> **Honesty about capability and limits.** Show what the solution actually does today. Do not stage a demo that implies capability the product does not have, and say clearly what is a prototype, what is simulated, and what is not yet built. Report results, including failures, as they are.

Conflicts of interest and vendor neutrality. Disclose any interest that could affect a recommendation, including commissions, partnerships or personal stakes. When the honest recommendation is a different tool, a smaller scope, or no software at all, say so.

Client confidentiality and intellectual property. Protect customer data, code and business information, inside and after the engagement. Do not reuse customer-derived work, or share a customer's name or lessons publicly, without permission.

Consent and the dignity of affected workers. Observe and record people's work only with their knowledge and the customer's agreement. Treat the people whose jobs change as stakeholders, not obstacles. Be truthful with them about what the change means for their work.

Raising concerns. When something is unsafe, unlawful, or harmful to the people the system affects, raise it promptly through the right channel, and keep a record. Do not continue work that bypasses a required control while the concern is open.

Crediting others. Credit the customer staff, colleagues and community members whose work contributed, in readouts, case studies and reusable assets.

Staying within competence. Take on work you can do or can learn safely under supervision. Say when a question needs a lawyer, a compliance officer, a security specialist or another expert, and bring them in.

Design principles

Vendor-neutral

Outcome-anchored

Observable and assessable. Indicators are written so that a future, validated process could assess them. Nobody can be assessed against this draft.

Built to be ready for ISO/IEC 17024 later (this draft is a starting hypothesis for the job task analysis, not a substitute for it)

Standard-setting kept separate from training, including any training affiliated with FDE Council's founder. The founder's own training is held to the same rule: it may say "aligned to the open framework" and nothing stronger.

Specialization is additive: advanced specializations and industry tracks build on the core domains and never replace them

Open by default

Validation roadmap

1. Community review of the draft (open comment period).
2. Job task analysis: a structured survey of practicing FDEs and FDE leaders plus subject-matter-expert workshops. (State of FDE 2026 gives early input; it is not a substitute.)
3. v0.x to v1.0: revise, and publish the job-task-analysis method and results.
4. Later and separate: any credential scheme, under a separately governed committee.

How to give feedback

Name the domain number, the context-required capability, or the industry track, say what is missing, wrong, or unclear, suggest wording, and tell us your context (company type, region, level). Feedback on the behavioral indicators is especially welcome: tell us which ones match real work and which do not. Please do not share confidential customer information. If you are mapping a learning path to this draft, read how to map a learning path.

Use the feedback form

Join the discussion in the LinkedIn group

Versioning and license

Releases use MAJOR.MINOR. 0.x releases are drafts; 1.0 will be the first release validated by a job task analysis. Earlier versions are kept online where possible, and their framework text is not rewritten when a later draft is published. v0.2 is not currently online. This draft is also available as a PDF (274 KB).

FDE Council Professional Capability Framework v0.4 © 2026 FDE Council, licensed under CC BY 4.0. You may share and adapt it for any purpose, including commercially, with credit. Use does not imply endorsement by FDE Council.

Editor: Noor Surani. Reviewers will be listed as the advisory board forms.

How to cite this framework

> FDE Council. (2026). *FDE Council Professional Capability Framework* (Version 0.4, draft). https://www.fdecouncil.com/framework/. Licensed under CC BY 4.0.

To cite a single domain or track, give its name as well, for example “FDE Council Professional Capability Framework v0.4, Domain 6.” Each industry track page gives its own citation. Earlier drafts: v0.3 and v0.1.

Change log

v0.4 November 2026

Current draft. Same 11 core domains, names, order and descriptions. Illustrative behavioral indicators at Associate, Practitioner and Lead for every core domain, for self-reflection and community review, not pass criteria. New optional Industry tracks layer, starting with a Mortgage and Fintech track. New cross-cutting section on ethics and professional conduct. Context-required capabilities, advanced specializations, design principles and validation roadmap kept. The v0.1, v0.2 and v0.3 texts were not rewritten.

v0.3 November 2026

Earlier draft. Same 11 core domains. Seven descriptions made explicit, including model evaluation in Domain 6 and token and inference cost in Domain 7. Context-required capabilities: AI and agent engineering, restricted environments, and regulated-data constraints. AI and agent engineering mitigates prompt injection and names model version drift. Two example advanced specializations, including AI evaluation lead. The v0.1 and v0.2 texts were not rewritten.

v0.2 October 2026

Earlier draft. 11 mandatory core domains. Evaluation is its own domain and applies to all FDE work. Not currently online.

v0.1 October 2026

First public review draft: 7 domains with competencies and behavioral indicators at 3 levels.