FDE Council Professional Capability Framework

Version 0.5 (draft, open for comment until Friday, December 19, 2026)

Published November 22, 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.

What is the FDE framework?

The FDE framework is an open, vendor-neutral competency framework that defines what a Forward Deployed Engineer should be able to do. Draft v0.5 sets out 11 core domains plus one cross-cutting domain, Regulation, Compliance and Responsible AI, each with behavioral indicators at Associate, Practitioner and Lead. It is free under CC BY 4.0.

Also called the Forward Deployed Engineer framework, the FDE competency framework or the FDE skills matrix. Its formal name is the FDE Council Professional Capability Framework.

Is v0.5 final?

No. FDE Framework v0.5 is a public draft, open for comment until Fri Dec 19, 2026. 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, not pass criteria. Earlier drafts stay online at v0.4, v0.3 and v0.1. The free FDE self-assessment rates you across the 11 core domains, for self-reflection only. It does not yet score the new cross-cutting domain or AI and agent engineering.

On this page

1. What is the FDE framework?
2. What changed from v0.4
3. How the DFW Kickoff shaped v0.5

4. How the FDE framework is structured
5. Are the domains a sequence?
6. The build loop across Domains 3 to 7
7. Internal and consulting FDEs
8. One FDE or a pod
9. Token and AI cost
10. Proficiency levels and the FDE skills matrix
11. The 11 core domains
12. Regulation, Compliance and Responsible AI
13. Enablement and handoff layer
14. Context-required capabilities
15. Advanced specializations
16. Industry tracks
17. Ethics and professional conduct
18. Design principles and validation roadmap
19. How to give feedback
20. Versioning, license and citation
21. Acknowledgements
22. Frequently asked questions

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

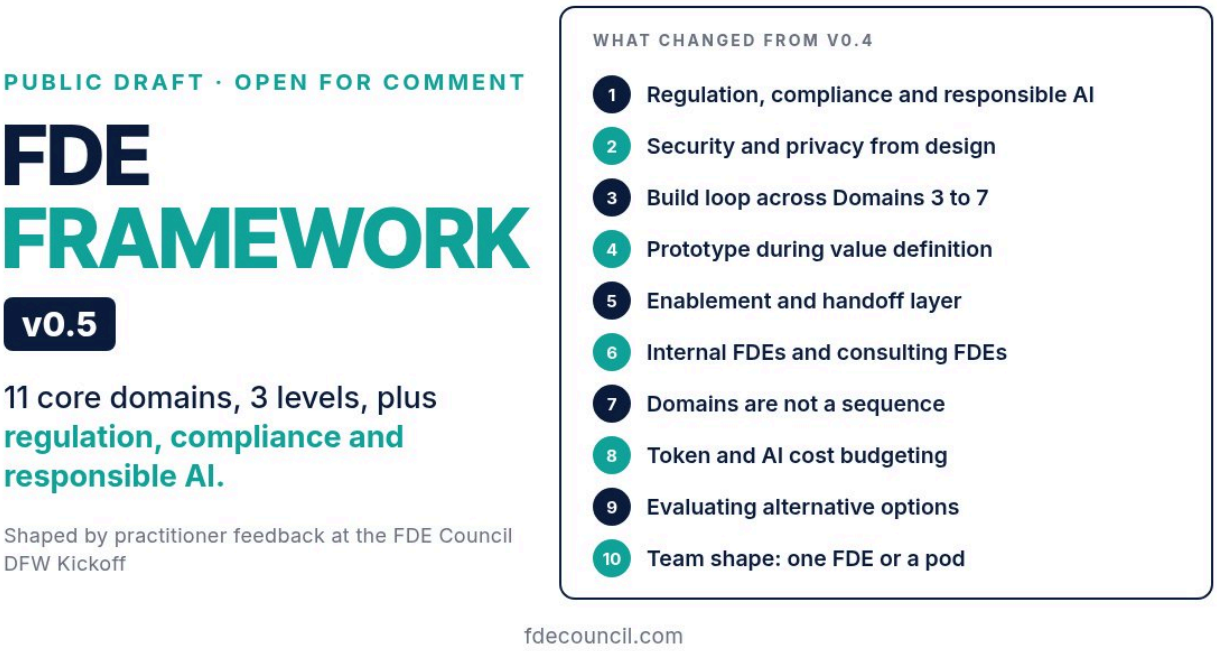

What changed in FDE framework v0.5?

Draft v0.5 makes ten changes to v0.4. The 11 core domains keep their names and numbers. New behavioral indicators are marked “New in v0.5” inside each domain, and changed descriptions say what was added.

What changed from v0.4 to v0.5

Change	What v0.5 says	Where to read it
Regulation, compliance and responsible AI	A new mandatory cross-cutting domain with five checkpoints and its own indicators at three levels, including use-case sign-off by accountable owners such as legal, marketing and portfolio owners.	Regulation, compliance and responsible AI
Security and privacy before production	Domain 8 work starts at design, is checked on every pass of the build loop, and is approved before go-live. New indicators in Domains 3 and 8.	Domain 8: Security, Privacy, Safety and Governance
An iterate loop across Domains 3 to 7	Architecture, data, build, evaluation and production readiness repeat as a loop, with an iterate, promote, reshape or stop decision on each pass.	The build loop

Change	What v0.5 says	Where to read it
Rapid prototype in value definition	Domain 2 allows a quick, clearly labelled prototype or demo when the customer needs to see the value before approving a use case.	Domain 2: Workflow Design and Value Definition
Training and handoff layer	A new enablement and handoff layer: training by role, a runbook, a named owner and agreed exit criteria, so the customer team can run the solution without the FDE.	Enablement and handoff layer
Internal FDEs, not only consultants	One framework for internal and consulting FDEs, with a reading guide that shows what changes in each context.	Internal and consulting FDEs
Domains are not a sequence	Stated plainly, with a diagram. Domain numbers are stable identifiers for citation and the self-assessment, not steps.	Are the domains a sequence?
Token and AI cost budgeting	AI cost is estimated in Domain 3, measured in Domain 6 and tracked in Domain 7, with new indicators in Domains 3 and 6.	Token and AI cost
Evaluating alternative options	Domain 2 compares alternatives before choosing one, including a process change with no new software and automation without AI.	Domain 2: Workflow Design and Value Definition
Team shape: one FDE or a pod	The framework describes capability, not headcount. A single FDE or a pod of two or three can cover the domains, and Domain 10 records the choice.	One FDE or a pod

Kept from v0.4: the names and numbers of the 11 core domains, the three levels and their scope, every v0.4 behavioral indicator, the context-required capabilities, the two example advanced specializations, the Mortgage and Fintech industry track, the ethics and professional conduct section, the design principles, the validation roadmap and the license. The v0.4 text is archived and was not rewritten.

How did the DFW Kickoff shape v0.5?

Practitioners at the FDE Council DFW Kickoff in Lewisville, Texas, on November 21, 2026, reviewed draft v0.4 in person. They proposed eight changes and raised two more questions, on evaluating alternative options and on team shape. All ten are in this draft as proposed wording, listed in the table above. The strongest agreement was on making regulation, compliance and responsible AI explicit.

Nothing in v0.5 is final. The comment period closes on Fri Dec 19, 2026, and the comments will be reviewed before the next draft. Tell us which changes you support, oppose or would word differently through the framework feedback form or the FDE Council LinkedIn group.

How is the FDE framework structured?

Draft v0.5 has five layers, one cross-cutting domain and one cross-cutting expectation:

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

One mandatory cross-cutting domain (new in v0.5): Regulation, Compliance and Responsible AI, applied at checkpoints across the core domains.

An enablement and handoff layer (new in v0.5), so the customer team can run the solution after the FDE leaves.

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, 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, which applies in every domain, layer and level.

Are the FDE framework domains a sequence?

No. The 11 core domains are competency domains, not project phases. The numbers are stable identifiers for citation and for the self-assessment, and they do not set an order. An FDE draws on several domains at once and returns to them throughout a deployment.

For example, building trust with executives (Domain 9) often starts before architecture (Domain 3), and security and privacy (Domain 8) start at design, not after production (Domain 7). For a step-by-step delivery process, from discovery to handover, see the FDE delivery lifecycle.

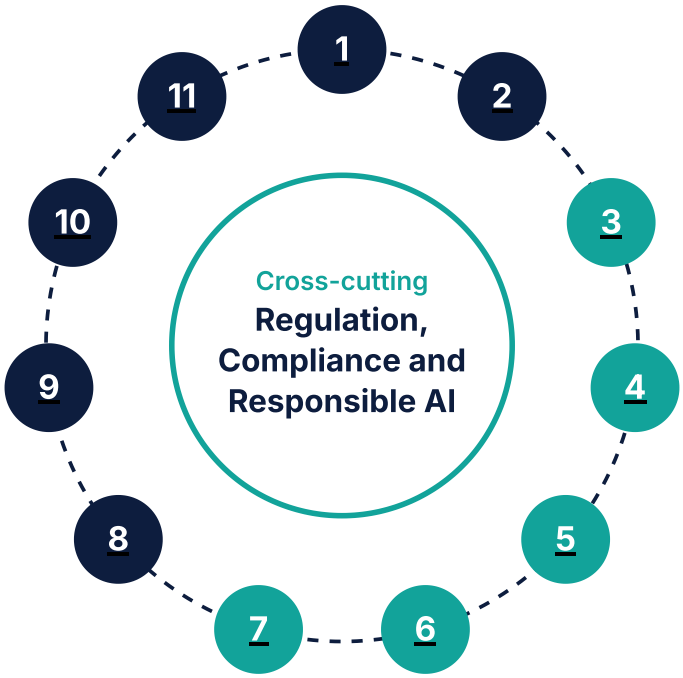

Teal: the build loop, Domains 3 to 7. Numbers are identifiers, not steps.

The 11 core domains as a cycle around the cross-cutting domain. Select a number to open that domain.

The build loop across Domains 3 to 7

Before a solution reaches production, Domains 3 to 7 repeat as a loop: design the architecture (Domain 3), prepare the data (Domain 4), build and test the riskiest assumption (Domain 5), evaluate against the acceptance measures (Domain 6), and check production readiness (Domain 7). Security and privacy (Domain 8) and the compliance checkpoints run inside every pass.

Each pass ends with one recorded decision: iterate again, promote to production, reshape the use case, or stop. Many deployments stall at the proof of concept. The loop makes the decision to move past it, or to stop, explicit.

Does the FDE framework apply to internal FDEs?

Yes. The framework applies to FDEs who work inside their own company as well as to consulting and vendor FDEs. The competencies are the same. What changes is who the customer is, how the work is funded, and what exit means.

How to read the FDE framework as a consulting FDE or an internal FDE

Area	Consulting or vendor FDE	Internal FDE
Discovery (Domain 1)	Learns a new customer's business from the outside.	Already knows the business, and still observes the actual work and challenges assumptions from the inside.
Embedded collaboration (Domain 9)	Earns trust as an outsider, usually through a sponsor.	Serves a department whose leader is often the main stakeholder. Trust comes from shared goals and track record.
Delivery and commercial judgment (Domain 10)	Works to a contract, a scope and a price.	Works to an internal business case, a budget and the department's priorities.
Enablement and handoff	Leaves at exit, and the customer team takes over.	Decides whether to stay embedded, rotate to the next team, or train the department to run it.
Field learning and reuse (Domain 11)	Needs permission to reuse work across customers.	Can reuse across departments more easily, and still respects data, ownership and confidentiality rules.

v0.5 keeps one framework with two readings, instead of two separate variants, so that citations and the self-assessment stay consistent. If you think two variants would serve you better, say so in your comments.

Is a Forward Deployed Engineer one person or a pod?

Either. The FDE framework describes capability, not headcount. One FDE can cover the domains, or a pod of two or three people can cover them together.

One person with deep skill in every domain is rare and costly, so many engagements use a small pod, for example a delivery or product lead paired with an engineer who prototypes quickly. When a pod is used, the pod as a whole should cover every domain, and each member should still be developing capability in all eleven. Domain 10 records the team shape, who owns depth in each domain, and the cost trade-off. For how to staff a practice, read how to build a Forward Deployed Engineering team.

How does the FDE framework handle token and AI cost?

AI cost is planned before build, not discovered in production. Token and inference cost is estimated, and the model is chosen, in Domain 3. It is measured alongside quality in Domain 6, and tracked against the agreed budget in Domain 7.

Estimate (Domain 3): cost per transaction and per month, and the right model, or no model, for each step. Not every step needs AI, and a smaller or in-house model can be the right choice.

Measure (Domain 6): cost per completed task, reported with quality, against the estimate.

Track (Domain 7): operating cost against the budget agreed with the customer, including model version changes.

What are the FDE proficiency levels?

The expected independence and depth in each domain increase with experience. Three indicative levels describe that progression. Read with the indicators in each domain below, they form an FDE skills matrix by level:

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 V0.5

What are the 11 core domains of the FDE framework?

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. In v0.5, five descriptions (Domains 2, 3, 6, 8 and 10) gain new sentences, and new indicators are marked. The other descriptions are the v0.3 wording, unchanged. The numbers are identifiers, not steps. 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. Compare alternative options before choosing one, including changing the process without new software and automating without AI. When the customer needs to see the value before approving a use case, build a quick, clearly labelled prototype or demo of the future workflow.

What an FDE must be able to do. Description updated in v0.5: evaluating alternative options and a rapid prototype during value definition were added. The rest is the v0.3 wording.

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.

New in v0.5: Compares at least two alternative options, including a process change with no new software and automation without AI, and records why the chosen option was picked.

New in v0.5: When the customer needs to see the value first, builds a quick prototype or demo of the future workflow, labels it clearly as a prototype, and records what it does not yet do.

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.

New in v0.5: Decides when a value-definition prototype is worth its cost, and keeps it separate from the production build in Domain 5.

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. When the system calls a model, estimate token and inference cost before build, and choose the right model, or no model, for each step, including smaller or in-house models. Agree security, privacy and compliance requirements during design, not after build.

What an FDE must be able to do. Description updated in v0.5: AI cost estimation, model choice, and security, privacy and compliance requirements at design were added. The rest is the v0.3 wording.

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.

New in v0.5: Produces an AI cost estimate before build, covering tokens, inference and hosting per transaction and per month, and compares it with the value agreed in Domain 2.

New in v0.5: Gets security, privacy and compliance requirements agreed with their owners during architecture, before build starts.

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.

New in v0.5: Sets model-choice and AI cost guidelines that keep deployments affordable, including when not to use a model.

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. In v0.5, a value-definition prototype can come earlier, in Domain 2. Domain 5 is where the riskiest assumptions are tested and the chosen code is hardened for production.

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. Measure cost as well as quality, including token and inference cost per completed task when a model is used.

What an FDE must be able to do. Description updated in v0.5: cost is measured alongside quality. The rest is the v0.3 wording.

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.

New in v0.5: Reports quality and cost together, including cost per completed task when a model is used, and compares the cost with the Domain 3 estimate.

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. In v0.5, AI cost is estimated in Domain 3 and measured in Domain 6 before it is tracked here, and the handover is set out in the new enablement and handoff layer.

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. This work starts at design, not after production: requirements are set during Domains 1 to 3, checked on every pass of the build loop, and approved before go-live.

What an FDE must be able to do. Description updated in v0.5: security and privacy start at design. The number 8 is an identifier, not a step after Domain 7. Rules that apply to the use case sit in the cross-cutting Regulation, Compliance and Responsible AI domain.

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.

New in v0.5: Gets security and privacy requirements agreed during architecture (Domain 3), and checks them on every pass of the build loop, not only before go-live.

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. For an internal FDE, the customer is usually the department being served, and its leader is often the main stakeholder. See internal and consulting FDEs.

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. Choose a team shape, a single FDE or a small pod, that covers every domain at a cost the engagement can carry.

What an FDE must be able to do. Description updated in v0.5: team shape was added. The rest is the v0.3 wording.

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.

New in v0.5: Records the team shape for the engagement, a single FDE or a pod, names who owns depth in each domain, and explains the cost trade-off to the sponsor.

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, except where AI cost is covered.

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 is estimated in Domain 3, measured in Domain 6 and tracked in Domain 7 (see token and AI cost). Security, privacy, and approvals stay in Domain 8. The row below is this same capability.

Context-required capabilities in draft v0.5

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 is estimated in Domain 3, measured in Domain 6 and tracked in 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.

Capability	When it applies	What an FDE must be able to do
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. Which rules apply to the use case, and who signs it off, sits in the cross-cutting Regulation, Compliance and Responsible AI domain.

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 v0.5

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 V0.5 · CROSS-CUTTING DOMAIN

Where do regulation, compliance and responsible AI fit in the FDE framework?

In v0.5, Regulation, Compliance and Responsible AI is a mandatory cross-cutting domain. It is not numbered, because it does not happen at one point: it applies at five checkpoints across the core domains, and it has its own behavioral indicators at Associate, Practitioner and Lead.

Cross-cutting domain: Regulation, Compliance and Responsible AI

Identify the laws, sector rules and internal policies that apply to a use case from the start, and keep checking them as it changes. Apply responsible AI principles: fairness, transparency, accountability, privacy and inclusion. Tell the people affected when AI is used. Get use-case sign-off from the accountable owners, such as legal, marketing and portfolio owners, before build and again before go-live. Recognize when the solution itself can help the customer meet an obligation, and bring in legal or compliance experts when a question is beyond personal competence.

What an FDE must be able to do. New in v0.5. This is not legal advice, and FDEs are not expected to act as lawyers.

The five compliance checkpoints

1. **After discovery** (with Domain 1). Name the laws, sector rules and internal policies that apply, and whether AI use must be disclosed. Include them in the go, reshape or stop recommendation.
2. **During value definition** (with Domain 2). Put compliance cost and constraints into the value case. In a regulated enterprise they can change whether a use case is worth doing. Note when the solution could itself help meet an obligation.
3. **After design, before build** (with Domains 3 and 8). Confirm that the design meets the rules that apply and the security and governance requirements, before development starts.
4. **Before go-live** (with Domains 6, 7 and 8). Get written use-case sign-off from the accountable owners, and put disclosure to affected people in place.
5. **In operation** (with Domains 7 and 11). Check again when a rule changes, or when the use case grows beyond what was approved, for example new product types or new groups of users.

Behavioral indicators by level

ASSOCIATE

Records which rules and policies the team was told apply, and asks before using data or AI in a new way.

Labels AI-generated output where people could mistake it for a human decision or a verified fact.

Raises a possible compliance issue the same day, through the agreed channel.

PRACTITIONER

Produces a map of the rules and policies that apply to the use case after discovery, confirmed by the customer's compliance owner.

Runs the checkpoint after design and before build, and records the result.

Gets written use-case sign-off from the accountable owners before go-live, and puts AI-use disclosure in place.

Applies responsible AI principles in proportion to the impact, and records how fairness, transparency, accountability, privacy and inclusion were handled.

LEAD

Sets use-case governance for a portfolio: who signs off, when, and with what evidence.

Works with customer legal, risk and compliance leaders as a peer, and stops or reshapes a use case that cannot meet its obligations.

Turns repeated compliance findings into checklists and patterns that other FDEs reuse.

How it differs from Domain 8, ethics and regulated-data constraints

Domain 8 covers the security, privacy and safety controls a system needs. Ethics and professional conduct covers how the FDE behaves. Regulated-data constraints remain a context-required capability for sector data rules such as residency, consent and retention. This domain covers which external rules and internal policies apply to the use case, how AI use is disclosed, and who must approve it.

Why a cross-cutting domain and not Domain 12?

Practitioners at the DFW Kickoff agreed that regulation and compliance must be explicit, and that "governance" alone is too broad. They proposed different places for it: a new Domain 12, a Domain 0 before discovery, part of value definition, or a check after design and before build. Each proposal is right about one moment. A cross-cutting domain with checkpoints covers all of them, and it avoids implying that compliance comes last (Domain 12) or only at the start (Domain 0).

It is mandatory and has its own indicators, so it carries the weight of a core domain. Keeping the 11 core domain numbers stable also keeps existing citations, such as "v0.4, Domain 6", and the self-assessment valid. Whether it should become a numbered core domain before v1.0 is an open question for this comment period.

NEW IN V0.5 · LAYER

How does the FDE framework handle training and handoff?

Every deployment should end with a team that can run, change and own the solution without the FDE. The enablement and handoff layer makes that an explicit expectation, with exit criteria agreed at the start.

The layer draws on Domain 7 (runbooks, support and a named owner), Domain 9 (building the team's capability) and Domain 11 (reusable assets), and adds what they leave implicit:

Training by role. Operators, engineers and managers each learn what they need to run, change or oversee the solution.

A runbook and a named owner before the FDE leaves.

A supervised run. The customer team operates and changes the solution without help, while the FDE is still there to catch problems.

Agreed exit criteria, written at the start of the engagement, not at the end.

For internal FDEs, an embed-or-train decision. Decide whether the FDE stays embedded, rotates to the next team, or trains the department to run it.

Behavioral indicators by level

ASSOCIATE

Prepares training material and runbook sections for the parts they built, and checks them with a real user.

PRACTITIONER

Agrees exit criteria with the sponsor at the start, runs training by role, and confirms the customer team has operated and changed the solution without help before the FDE leaves.

LEAD

Sets handoff standards across deployments, and decides with customer leaders whether to stay embedded, rotate, or train the team to run the solution.

SINCE V0.4

What are FDE framework 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.

SINCE V0.4

What ethics and professional conduct does the FDE framework expect?

These expectations apply across every domain, layer and level. They are not a domain, and they go beyond the controls in Domain 8 and the rules covered by the cross-cutting Regulation, Compliance and Responsible AI domain. Domain 8 is about the controls a system needs, and the cross-cutting domain is about the rules a use case must meet. 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.

What are the FDE framework 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 do I give feedback on the FDE framework?

Name the domain number, the cross-cutting domain, the layer, 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, license and how to cite the FDE framework

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. A PDF of draft v0.5 is available. The v0.4 PDF (274 KB) stays online with the archived v0.4 page.

FDE Council Professional Capability Framework v0.5 © 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 0.5, public 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 v0.5, Domain 6” or “v0.5, cross-cutting domain: Regulation, Compliance and Responsible AI.” Each industry track page gives its own citation. Earlier drafts: v0.4, v0.3 and v0.1.

Change log

v0.5 November 2026

Current public draft, open for comment until Fri Dec 19, 2026. Same 11 core domain names and numbers. New cross-cutting domain: Regulation, Compliance and Responsible AI. New enablement and handoff layer. Build loop across Domains 3 to 7, domains stated as not a sequence, internal FDEs covered, and new description sentences and indicators in Domains 2, 3, 6, 8 and 10 for alternative

options, rapid prototyping, AI cost, security at design and team shape. Shaped by practitioner feedback at the DFW Kickoff. The v0.4 text was not rewritten.

v0.4 November 2026

Superseded by v0.5. Same 11 core domains, names, order and descriptions as v0.3. Illustrative behavioral indicators at Associate, Practitioner and Lead for every core domain. New optional Industry tracks layer, starting with a Mortgage and Fintech track. New cross-cutting section on ethics and professional conduct.

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. Two example advanced specializations, including AI evaluation lead.

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.

Acknowledgements

FDE Framework v0.5 was shaped by the **Participants of the FDE Council DFW Kickoff, Lewisville, Nov 21, 2026**, who reviewed draft v0.4 in person and proposed the changes listed in what changed from v0.4. Thank you for your time, your challenges and your field experience.

Individual contributors are named here only with their written permission. If you took part and would like to be credited, email nsurani@fdecouncil.com.