



FREE HIRING GUIDE • 2026 EDITION

The Forward-Deployed Engineer Hiring Guide

What a forward-deployed engineer actually is, how the role differs from sales engineering and solutions architecture, the three FDE profiles, where to find them, how to interview them, and what they cost in 2026.

By the technical recruiting team at Sloane Staffing

Sloane Staffing

Hiring top talent for top companies.

sloane-staffing.com

SLOANE STAFFING · FREE HIRING GUIDE

What's inside

- 01 Introduction

- 02 What a Forward-Deployed Engineer Is

- 03 The State of FDE Hiring in 2026

- 04 FDE vs. Sales Engineer vs. Solutions Architect

- 05 The Three FDE Profiles

- 06 What to Look For & Where to Find Them

- 07 The Interview Loop

- 08 Job-Description Template (Copy-Paste)

- 09 2026 Compensation

- 10 Setting Up Your First FDE to Succeed

- 11 Frequently Asked Questions

- 12 About Sloane Staffing

INTRODUCTION

Before you start

Every applied-AI company eventually hits the same wall. The model works, the demo lands, the contract is signed — and then nothing ships, because the customer's data lives in seven systems nobody has documented in years. The forward-deployed engineer is the role built to get through that wall.

It is also the most misunderstood technical hire on the market. The title is borrowed by sales engineers, solutions architects, implementation consultants and senior backend developers, and those are not the same job. Hiring managers write a req for one, describe another, and interview for a third.

This guide starts with what the role actually is, then covers how to hire one. Use it end to end, or jump to what you need:

If you're deciding whether you need one

Start with chapters 1–3: what an FDE is, the 2026 market, and how the role differs from a sales engineer or solutions architect.

If you're ready to hire

Go to chapters 4–9: the three FDE profiles, where they come from, the interview loop, a JD template, 2026 pay, and how to set up the first 90 days.

CHAPTER 01

What a Forward-Deployed Engineer Is

A forward-deployed engineer embeds with a customer and writes production code against that customer's real data, in the customer's environment, until the product delivers a result. They own the outcome, not the ticket.

Palantir created the role in the early 2010s, when its FDEs were known internally as "Deltas." Until around 2016, Palantir employed more FDEs than regular software engineers. Palantir's own description of the job is still the clearest:

FDEs' responsibilities look similar to those of a startup CTO: you'll work in small teams and own end-to-end execution of high-stakes projects.

— Palantir

One customer, many capabilities

Palantir drew the line between its product engineers and its FDEs in one sentence: a product engineer's focus is "one capability, many customers," while an FDE's focus is "one customer, many capabilities." A product engineer builds a feature that works for everyone. An FDE makes the whole product work for one customer, then feeds back what they learned.

Three jobs in one role

- 1 Build inside the customer**
Scope the problem with the customer's team, connect to their systems and data, and ship working software in their environment. Anthropic's postings describe delivering artifacts like MCP servers, sub-agents and agent skills directly for customers.

- 2 Get it live and prove it works**
Take the build through validation and into production. OpenAI describes its FDE work in three phases: early scoping and prototyping, validation with evaluations, and delivery on the customer's infrastructure.

- 3 Turn the lesson into product**
Identify what will repeat at the next customer and codify it — as a pattern, a template, or a change to the core product. This is what separates an FDE from a consultant: the work makes the next deployment faster.



The FDE is the engineer your customer sees most — and the reason the contract turns into a renewal.

The short definition to use internally. An FDE is a software engineer whose product is a working customer outcome. If the person you are describing does not ship production code, they are not an FDE. If they do not spend real time with the customer, they are not one either.

CHAPTER 02

The State of FDE Hiring in 2026

Demand for FDEs went vertical in about a year. Enterprises are buying AI faster than they can deploy it, and the companies selling AI need engineers who can close that gap on-site. The number of people who have done the job did not grow at the same rate.

729%

Year-over-year growth in FDE job postings, April 2026

\$171.9k

Average FDE base salary across the market

25–50%

Travel to customer sites in frontier-lab FDE postings

4+ yrs

Customer-facing technical experience asked for at a frontier lab

What changed

AI created a deployment gap

A model that works in a demo still has to be connected to a customer's data, permissions and workflows. That work is engineering, it happens on the customer's side, and it decides whether the contract renews.

The biggest names are hiring

Anthropic, OpenAI, Palantir, Stripe and Google Cloud are all recruiting forward-deployed engineers, and consulting firms are recruiting from the same pool. Your offer is competing with theirs whether you can see it or not.

The teams are small on purpose

OpenAI had just over 10 FDEs across 8 cities on 3 continents in 2025, and Ramp ran about 15 in pods. When the category leaders compete over double-digit headcount, every experienced FDE is in demand.

The supply did not move

By April 2026 postings were 5,230% above January 2025 levels. The pool of engineers who have actually shipped inside a customer's environment grows by experience, not by job postings.

What this means for your search. Almost nobody with real FDE experience is applying to job boards. Filling the role is an outbound search into a small, heavily recruited pool, so the time you spend defining the job before you start is the cheapest time in the whole process.

CHAPTER 03

FDE vs. Sales Engineer vs. Solutions Architect

Most failed FDE searches are a definition problem. The req says forward-deployed engineer, the hiring manager describes a sales engineer, and the panel interviews for a services consultant. Settle which job you have before you write anything.

	Forward-deployed engineer	Sales engineer	Solutions architect	Professional services
Ships production code for the customer	Yes	No	Rarely	Sometimes
Main work happens	After signature	Before signature	Around design and architecture	After signature
Owns	Whether the customer gets a result	The technical win	The design of the solution	Scope delivered on time
Reports into	Engineering or deployment	Sales	Sales or customer success	A services P&L
Measured by	Deployments live, outcomes shipped	Win rate	Adoption and design quality	Utilization, billable hours
Feeds the core product	Directly	Through feedback	Through feedback	Rarely

Do you actually need an FDE yet?

You probably do if...

Deals close but deployments stall; customers need custom integration work before they see value; renewals depend on what gets built after signature; and the lessons from each deployment should change the product.

You probably don't if...

Deals stall before signature on technical questions (hire a sales engineer); the product works out of the box after configuration (hire implementation or CS); or the work is billable scope with no product feedback loop (hire services).

One question settles it. Ask "who owns whether this customer succeeds after signature, and do they write code?" If the answer is someone in sales, you need a sales engineer. If the answer is nobody, and the work needs code, you need an FDE.

CHAPTER 04

The Three FDE Profiles

"FDE" covers three different hires. They share the core job — ship inside the customer, own the result — but they need different backgrounds, and a candidate who is excellent at one can struggle at another.

PROFILE

Platform FDE

The original Palantir model. Deploys a data or software platform into a customer's systems: integrations, data pipelines, workflows and custom applications on top. Strongest on data engineering and systems integration.

PROFILE

Applied AI FDE

The fastest-growing version. Builds LLM applications and agents inside customer environments: retrieval, tool use, evaluations and guardrails. Strongest on production LLM work and turning a demo into something reliable.

PROFILE

Regulated-Industry FDE

Either of the above, inside government, healthcare or financial services. Adds the domain knowledge and compliance constraints the customer operates under, and sometimes a security clearance.

How they differ

	Platform FDE	Applied AI FDE	Regulated-Industry FDE
Core technical skill	Data engineering, integration	LLM apps, agents, evaluations	Either, plus compliance constraints
Typical languages	Python, SQL, TypeScript/Java	Python, TypeScript	Python, plus whatever the customer runs
Best source	Palantir alumni, data platform SAs	Founding and AI-product engineers	Engineers from GovTech, health-tech, fintech
Hardest part of the job	Undocumented legacy systems	Making non-deterministic output reliable	Moving fast inside slow procurement and security review

Regulated work narrows the pool further. Anthropic's 2026 federal civilian FDE posting required prior experience with US federal, state or local agencies, and listed a security clearance as a plus. Each requirement like that removes most of an already small market, so separate the true must-haves from the nice-to-haves before you post.

CHAPTER 05

What to Look For & Where to Find Them

The technical bar is the easy half of this hire. The candidates who fail in the seat usually fail on judgment and communication, not on code.

Five traits that predict success

Trait	What it looks like
Production engineering	Has shipped and maintained real systems, not just prototypes. Can talk about what broke after launch.
Customer judgment	Can hold a room with a skeptical customer, re-scope on the spot, and say no without losing the relationship.
Comfort with ambiguity	Starts building before the requirements are clean, and finds the real problem behind the stated one.
Product instinct	Notices what will repeat at the next customer and pushes it back into the product.
Ownership	Talks about outcomes they delivered, not tickets they closed. Has a story where they killed their own solution.

The four pools FDEs come from

- 1 Palantir and its alumni**
The original pool. They understand the model without explanation. It is also the most heavily recruited.
- 2 Solutions architects who still write code**
At data and infrastructure companies like Databricks and Snowflake. They already work with customers and convert unusually well.
- 3 Founding and early startup engineers**
Engineers who were effectively a customer's CTO at a small company. They ship fast and tolerate ambiguity well.
- 4 Consulting-firm engineers**
They bring client skills already built and need product depth. Consulting firms are also hiring from this pool, so move quickly.

Red flags

Only ever advised

Their customer work was slides and architecture diagrams. Nothing they designed went to production under their name.

Defends every decision

Cannot name a time they were wrong. At your customer, that turns into defending an architecture instead of fixing the problem.

Surprised by the travel

Frontier-lab postings ask for 25–50% travel. If this is news in the final round, the process was set up wrong.

Wants to avoid customers

A great engineer who wants to be left alone to build is a product engineer. Hire them for that seat instead.

CHAPTER 06

The Interview Loop

Most FDE loops over-invest in algorithm rounds. An algorithm round tells you nothing about whether a candidate can walk into a hostile stakeholder meeting and re-scope a project. Test the job instead.

① Recruiter screen (30 min)

Motivation, customer-facing history, and the travel conversation. Ask about travel in the first fifteen minutes. It is the most common reason late-stage FDE processes collapse.

② Live scoping exercise (60 min)

Give the candidate a deliberately messy customer brief with a contradiction buried in it. Play the customer. The signal is whether they find the contradiction and say so, and what they propose to build first.

3

Practical build (90 min)

A realistic integration or AI task using your stack or a close proxy — connect to a messy data source, build a small working flow, explain the trade-offs. Grade what works, not what is clever.

4

Deployment deep-dive (60 min)

Have them walk through a real system they shipped for a customer: the stakeholders, what broke, what they would change, and what they fed back into the product.

Questions for every FDE candidate

- Tell me about a deployment that went badly. What was your part in it, and what did you change afterwards?
- Describe a time you killed your own solution. How did you know it was wrong?
- A customer's executive asks for something you think is the wrong thing to build. What do you do?
- What did you build for one customer that later became part of the product?
- How do you decide when a customer's request is a one-off and when it is a pattern?

Applied AI FDE

- How do you know an LLM application is good enough to put in front of a customer's users?
- Walk me through the evaluations you built for your last AI project. What did they catch?
- An agent works 90% of the time in testing. What do you do about the other 10% before launch?

Platform and regulated-industry FDE

- Tell me about the worst legacy system you have had to integrate with. How did you get the data out?
- How have you worked inside a customer's security review without stalling the project?
- What would you do in your first week at a customer whose systems nobody has documented?

CHAPTER 07

Job-Description Template (Copy-Paste)

Written for an Applied AI FDE. Swap the bracketed sections for the profile you are hiring. Stating the travel, the reporting line and the pay in the posting filters out the wrong candidates before the first call.

Title

Forward-Deployed Engineer — [Applied AI / Platform / Public Sector]

About the role

You will work directly with [customer type] to deploy [product] into their environment and make it deliver results. You will scope problems with customer teams, build production software on their systems and data, and turn what you learn into patterns and product improvements that make the next deployment faster. You will report to [Head of Deployment / VP Engineering].

What you'll own

- End-to-end technical delivery for [N] customer deployments at a time
- Production integrations, [agents / data pipelines / applications] and evaluations in customer environments
- The technical relationship with customer engineering and business teams
- Reusable deployment patterns and feedback into the core product

Must-have

- [4+] years building production software, including customer-facing technical work

- Strong [Python / TypeScript], and experience with [LLMs and agents / data engineering / systems integration]
- A track record of shipping in ambiguous environments with incomplete requirements
- Willingness to travel to customer sites [25–50]% of the time

Nice-to-have

- Experience in [industry / public sector], or a [security clearance]
- Prior FDE, founding-engineer or technical consulting experience

Compensation

[\$X–\$Y] base, plus [equity details]. [Travel expenses covered; location or remote policy].

Don't post it as a senior backend role. Customer-facing engineers do not read backend postings, and the ones who do often drop out at the job description. Put the customer work, the travel and the product feedback loop in the first paragraph.

CHAPTER 08

2026 Compensation

Only base pay is reliably published for FDE roles. Equity is where much of the rest of a package sits, and it is almost never posted.

What's published

Source	Published base
Market-wide average (Indeed, via Business Insider)	\$171,911
Posted range across employers (Indeed)	~\$170,000 to over \$200,000
Anthropic — FDE, Applied AI (2026 posting)	\$200,000–\$300,000
Anthropic — FDE, Federal Civilian, New York (2026 posting)	\$280,000–\$320,000

Read the spread, not the midpoint. A general-market FDE role and a frontier-lab FDE role can be \$100,000 or more apart on base before equity is discussed. A budget built only from published base figures will lose candidates in the final round.

If you can't win on cash, win on scope

A Series B company will not match a frontier lab's package. What it can offer is ownership: the first FDE who gets to design how the whole company deploys will often take less cash than an offer where they would be one of many. Make that case in the first conversation, not at the offer.



Budget for travel as part of the offer. At 25–50% travel, how you handle travel expenses, time off after heavy travel and location flexibility matters to candidates almost as much as base. Put it in writing early.

Sources: Business Insider (via AOL), Indeed data on forward-deployed engineer postings and pay, May 2026; LeadDev, "The rise of the forward-deployed engineer," July 2026; The Pragmatic Engineer, "Forward deployed engineers," August 2025; Anthropic job postings, Forward Deployed Engineer (Applied AI) and Forward Deployed Engineer (Federal Civilian), 2026. US figures. Accessed September 2026.

CHAPTER 09

Setting Up Your First FDE to Succeed

A strong FDE placed in the wrong structure will look like a weak hire within a quarter. Decide these four things before the start date.

1 Reporting line

Put FDEs under engineering or a dedicated deployment leader, not sales. If they report to sales, they will be pulled into pre-sale work and measured on deals, and the deployments will slip.

2

How they are measured

Deployments live, time to first value, and patterns contributed back to the product. Not hours, tickets or revenue closed.

3

The path back to product

Agree how FDE findings reach the product team: a regular review, a shared backlog, or rotation into product engineering. Palantir moved many FDEs into product engineering after 2016, carrying their field experience with them.

4

Team shape

Small pods beat a pool. Ramp runs about 15 FDEs in pods. Pairing FDEs with named customers or segments builds context that a shared queue throws away.

A realistic first 90 days

Period	Focus	What good looks like
Days 1–30	Product depth and a shadowed deployment	Can demo and explain the product's limits honestly
Days 31–60	Owens one customer deployment with support	A scoped plan agreed with the customer; first build in their environment
Days 61–90	Ships to production	A live deployment, plus one pattern or product change fed back

CHAPTER 10

Frequently Asked Questions

What is a forward-deployed engineer?

A software engineer who embeds with a customer and writes production code against the customer's real data and systems until the product delivers a result. Palantir created the role in the early 2010s and describes it as closest to being a startup CTO: small teams and end-to-end ownership of high-stakes projects.

How is an FDE different from a sales engineer?

A sales engineer owns the technical win before the contract and usually reports into sales. An FDE owns the outcome after it, ships production code in the customer's environment, and is measured on whether deployments go live. Hiring one when you needed the other is the most common mistake we see.

What does an FDE cost in 2026?

The market-wide average base salary is \$171,911, and Indeed's posted range runs from about \$170,000 to over \$200,000. Frontier labs publish higher bands — Anthropic listed \$200,000–\$300,000 for an Applied AI FDE in 2026. Equity is on top and rarely published.

How much do FDEs travel?

A lot. Anthropic's 2026 postings asked for 25–50% travel to customer sites, and OpenAI's lists up to 50%. Discuss it on the first call.

How fast can you introduce qualified FDEs?

In days, not weeks. We recruit FDEs outbound from the four pools in this guide, and screen them on production engineering, customer judgment and travel fit before you see them.

Who can we introduce you to today? Whether you are hiring your first FDE or building a deployment team, book a consultation at sloane-staffing.com/contact — we will help you decide which profile you need and what it will cost.

ABOUT

About Sloane Staffing

We're a go-to-market and technical recruiting firm that places the people who build and grow modern companies — from marketing, sales, and GTM engineering to the specialized data and AI talent behind them.

Specialist, not generalist

We recruit within focused practice areas and qualify candidates on the problems they've solved — not keyword-matched résumés.

We reach the people who aren't looking

We go outbound to pre-qualified, pre-interviewed candidates, so you typically see a vetted shortlist within days, not weeks.

Contract & permanent

One partner for both — whether you need a project-based resource or a long-term hire.

Founder-led expertise

Built by operators from the GTM and martech world who know these roles from the inside.

Let's find your next hire. Book a free consultation at sloane-staffing.com/contact — tell us who you need, and we'll show you how we'd source them.