NAAC GRADE A+ ACCREDITED UNIVERSITY

REVA RACE — REVA Academy for Corporate Excellence

REVA UNIVERSITY
Bengaluru, India

RACE WEEKLY INTELLIGENCE BRIEF

RETINA

AI · CYBERSECURITY · CLOUD · LEADERSHIP



THIS WEEK'S COVER:

THE FDE MOVES IN

The quiet truth of enterprise AI is that most of it does not work. Not because the models are weak, but because a powerful model dropped into a messy real company rarely turns into value on its own. In 2026 the industry settled on an answer, and it was not a better model. It was a person. The forward deployed engineer, who embeds inside your company and builds until the AI actually works, has become the hottest and highest-paid role in technology, with OpenAI, Anthropic, Microsoft and Amazon pouring more than 9 billion dollars into it. Here is what the role is, why it exists, and what it means for how you buy and staff AI.

ALSO INSIDE

The world in 5 · Signal · Hyperfocal · A week in the life · Perspectives · The FDE playbook · Metric of the week · Fun fact · Candid · Book shelf · Research radar · RACE pulse

ISSUE #28 · WEEK OF 9–15 NOVEMBER 2026

Visit our website
race.reva.edu.in

RETiNA
AI · CYBERSECURITY · CLOUD · LEADERSHIP

- **From The Editor's Desk**
- The World In 5
- Signal
- Hyperfocal
- A Week In The Life
- Perspectives
- The FDE Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- RACE Pulse
- RACE Programs

NAAC GRADE A+ ACCREDITED UNIVERSITY
REVA RACE REVA Academy for Corporate Excellence
REVA UNIVERSITY
Bengaluru, India

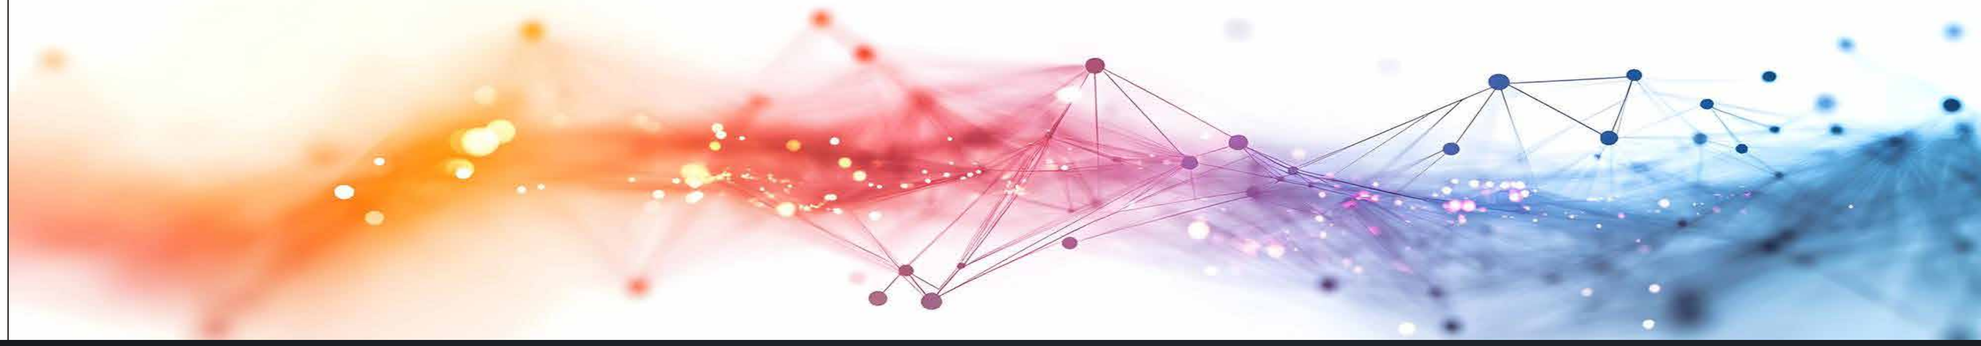
FROM THE EDITOR'S DESK

The most important AI hire of 2026 is not a model. It is a person who moves into your building until the model works.

For three years the industry sold capability. Buy the model, call the API, and the value would follow. It mostly did not. The widely cited finding that the overwhelming majority of enterprise AI pilots show no measurable business impact is not a story about weak models. The models are strong. It is a story about the gap between a capable model and a working outcome, a gap made of messy data, real workflows, cautious users and a hundred small frictions that no demo ever shows. Capability turned out to be the easy part. Deployment is where value is won or lost.

So the most interesting answer the industry found this year was not technical. It was human. It is the forward deployed engineer, a role Palantir invented two decades ago and the rest of the industry suddenly copied in a single year. An FDE is a software engineer who embeds inside a customer's company, lives in its mess, and ships real code until the AI genuinely works for that business, with full accountability for the outcome. Part engineer, part consultant, part diplomat. In 2026, OpenAI, Anthropic, Microsoft and Amazon committed more than 9 billion dollars to building forward deployed teams, and the role became one of the highest-paid in technology.

This matters to a leader in two directions at once. It changes how you buy AI, because the serious vendors now arrive with engineers, not just a licence, and the deal includes people who sit with your team. And it changes how you staff, because the organisations that get real value are building their own embedded engineers rather than buying a tool and hoping. For India, the implications are large and double-edged, and this issue comes to them. First, what the role actually is, and why it has become the whole game.



- From The Editor's Desk
- **The World In 5**
- Signal
- Hyperfocal
- A Week In The Life
- Perspectives
- The FDE Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- RACE Pulse
- RACE Programs

THE WORLD IN 5

GLOBAL HEADLINES A CXO SHOULD BE ABLE TO REFERENCE

Five signals from the past fortnight, worth a note at your next leadership table.

1. A PALANTIR ODDITY BECAME AN INDUSTRY STANDARD

For twenty years, the forward deployed engineer was a title almost nobody used outside Palantir. In a single year, it became the consensus way the AI industry reaches the enterprise. OpenAI, Anthropic, Google, Databricks, Scale AI, Cohere, Mistral and others now hire for it by name, and the big consultancies are scrambling to match the model. The strange job became the standard one.

2. THE FRONTIER LABS PUT BILLIONS BEHIND IT

In May 2026, within days of each other, OpenAI and Anthropic both launched dedicated forward deployed ventures. OpenAI stood up a deployment company to embed engineers inside enterprises, and Anthropic built a multi-billion-dollar enterprise services arm around applied AI engineers. Microsoft and Amazon followed with their own units. More than 9 billion dollars is now committed to putting engineers, not just software, inside customers.

3. IT BECAME ONE OF TECH'S BEST-PAID ROLES

Demand drove pay sharply up. The median forward deployed engineer in the US now earns around 220,000 dollars in total compensation, and senior packages at Palantir reach close to 590,000. These are numbers that rival or beat staff engineers at the biggest tech firms, for a role that did not have a clear market a year ago.

4. IT IS REPLACING THE SALES ENGINEER AND THE CONSULTANT

The labs have largely swapped traditional solutions-engineer headcount, people who demo from a deck, for embedded engineers who build. And because the FDE delivers the actual implementation, it competes directly with traditional consulting. The deployment work that used to go to a system integrator is increasingly done by the vendor's own engineers, sitting with your team.

5. INDIA: TRAINING THE DEPLOYMENT WORKFORCE

In October, Anthropic announced an academy to train 10,000 forward deployed engineers by the end of 2027. India, which has long been the world's technology deployment workforce, has the most to gain and the most to defend here. The role is a chance for Indian engineers to move up from billing hours to owning outcomes, and a direct challenge to the old staff-augmentation model. The country that trains these people fastest wins.

Source:

- *https://en.wikipedia.org/wiki/Forward_deployed_engineer*
- *https://getperspective.ai/blog/palantir-forward-deployed-engineering-playbook-anthropic-openai-copying*
- *https://openai.com/index/openai-launches-the-deployment-company/*
- *https://www.anthropic.com/news/claude-frontier-academy*
- *https://www.levels.fyi/*
- *https://jobsbyculture.com/*



- From The Editor's Desk
- The World In 5
- **Signal**
- Hyperfocal
- A Week In The Life
- Perspectives
- The FDE Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- RACE Pulse
- RACE Programs

SIGNAL

THE FORWARD DEPLOYED ENGINEER, BY THE NUMBERS



$9B+
committed by OpenAI, Anthropic, Microsoft and Amazon to forward deployed ventures in 2026

$220K
median US total compensation for an FDE; senior Palantir packages reach about $590K

10,000
forward deployed engineers Anthropic's Claude Frontier Academy aims to train by end 2027

May 2026
when OpenAI and Anthropic both launched dedicated FDE units within days of each other

8 to 16
weeks a typical FDE spends embedded on a single customer account

24
OpenAI job posts with Forward Deployed in the title on one day, from San Francisco to Seoul

Source:
- *https://openai.com/index/openai-launches-the-deployment-company/*
- *https://www.anthropic.com/news/claude-frontier-academy*
- *https://www.levels.fyi/*
- *https://d2beiqkhq929f0.cloudfront.net/public_assets/assets/000/249/168/original/CMR_Scaler_FDE_Opportunity_Report_19_Oct_2026.pdf*

- From The Editor's Desk
- The World In 5
- Signal
- **Hyperfocal**
- A Week In The Life
- Perspectives
- The FDE Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- RACE Pulse
- RACE Programs

HYPERFOCAL

THE BOTTLENECK WAS NEVER THE MODEL

For a long time the industry behaved as if the hard part of enterprise AI was the intelligence. Build a better model, the thinking went, and adoption would follow. Billions went into capability, and capability duly arrived. And then most of the pilots still failed to produce anything a CFO could measure. That result forced an uncomfortable realisation. The model was never the bottleneck. The bottleneck was everything between the model and the business: the data that is messier than anyone admits, the workflow no two companies run the same way, the users who do not trust it yet, the integration nobody scoped.

None of that is solved by a smarter model. It is solved by a person who will sit inside the specific, awkward reality of one company and make the thing work there. That is what a forward deployed engineer is. Not a salesperson who shows what the AI could do, and not a consultant who writes a plan for someone else to execute, but an engineer who builds the actual solution, in your systems, against your data, with your people, and does not leave until it works. The reason the frontier labs are spending billions on this is simple: it is the only approach that has reliably turned a capable model into a result a customer will pay to keep.

For a leader, the lesson cuts against the instinct of the last three years. The instinct was to treat AI as a product you procure, like software: choose the best one, buy it, roll it out. The forward deployed model says AI in the enterprise behaves far more like a capability you build with help, hands-on, in your own context, one hard-won deployment at a time. The winners will not be the organisations that bought the most advanced model. They will be the ones that got the right engineers, their vendor's or their own, deep enough inside the business to close the gap between what the AI can do and what it actually does.

RETINA TAKE

Enterprise AI did not have a capability problem. It had a last-mile problem, and the industry's answer was not a cleverer model but an engineer willing to live in the mess until it worked. The value was always in the deployment. We are finally paying for that, and paying a lot.



- From The Editor's Desk
- The World In 5
- Signal
- Hyperfocal
- **A Week In The Life**
- Perspectives
- The FDE Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- RACE Pulse
- RACE Programs

A WEEK IN THE LIFE

WHAT A FORWARD DEPLOYED ENGINEER ACTUALLY DOES

A new segment for this issue. The role in plain terms, told as the loop an FDE runs inside a customer, because the day-to-day is where this model wins or fails.

Discovery

Finding the problem worth solving

The embed starts not with code but with listening. The FDE sits with the people who do the work, watches how things actually happen rather than how the process diagram says they do, and hunts for the specific, high-value problem the AI can genuinely move. This is the most underrated part of the job. Most failed AI projects solved a problem nobody really had, or missed the messy detail that sank them. A good FDE spends real time here before writing a line.

What it shows: The win is chosen before it is built. The scarce skill is finding the right problem in a real company, not building a clever solution to the wrong one.

Build and deploy

Shipping real code in your world

Then the FDE builds, inside the customer's systems, against the customer's real data, with all its gaps and quirks. This is production engineering, not a prototype, and it is done on-site or deeply embedded, not thrown over a wall. The engineer wires the model into the actual workflow, handles the awkward integrations, and gets it in front of real users early, iterating against what breaks. The difference from a demo is total: the demo works once, the FDE makes it work on Tuesday.

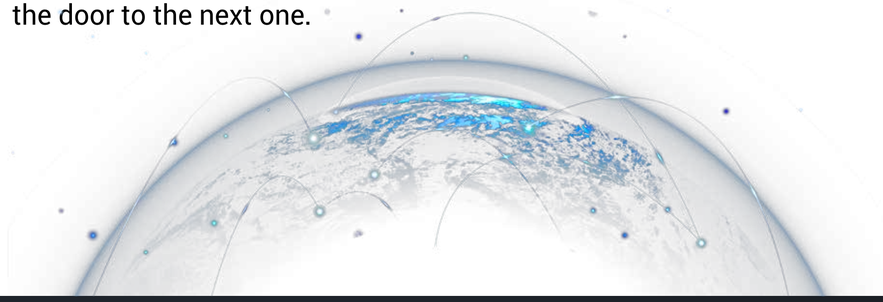
What it shows: The value is in the integration, not the model. Making AI survive contact with a real workflow is the whole job, and it is hard, unglamorous work.

Measure and expand

Owning the outcome, then growing it

An FDE is accountable for a result, not a deliverable. So the loop ends by measuring whether the deployment actually moved the number it promised, and feeding what was learned back, both into the next use case at that customer and, at the labs, into the product itself. A successful embed rarely stops at one win. It earns trust, and trust opens the next problem. This is why the model drives durable customer growth, and why the labs value the field learning as much as the revenue.

What it shows: Accountability for the outcome is the point. The FDE is paid to make a business metric move, and a win it can prove becomes the door to the next one.



- From The Editor's Desk
- The World In 5
- Signal
- Hyperfocal
- A Week In The Life
- **Perspectives**
- The FDE Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- RACE Pulse
- RACE Programs

PERSPECTIVES

DEEP TECH — ANALYSIS BY THE RACE TEAM

RISE OF FDES – NEXT GENERATION MULTIDISCIPLINARY ENGINEERS

The return of the generalist in the age of enterprise AI



Dr. J. B. Simha
Chief AI Mentor and Chief Architect, REVA Academy for Corporate Excellence

https://www.linkedin.com/in/jbsimha/

The technology industry is witnessing a shift in the engineering profession. For more than a decade, specialisation was the dominant career path, with engineers concentrating on areas such as frontend, backend, data engineering, cloud infrastructure or AI. The growing complexity of enterprise AI deployments is now creating demand for a different kind of professional: the Forward Deployed Engineer (FDE). An FDE combines technical depth with multidisciplinary capabilities, business understanding and the ability to deliver solutions in real-world environments. The article reports that FDE job postings rose from 643 to 5,330 in a year, reflecting this changing demand.

The FDE lifecycle: five responsibilities

1. Customer engagement: Work closely with stakeholders to understand operational challenges and translate business concerns into technical problems.
2. Discovery: Examine processes, data, infrastructure and workflows; establish performance baselines and define measurable outcomes before building.
3. Solution design: Balance technical feasibility, budget, timelines, security, compliance and business economics. The most sophisticated solution is not always the best.
4. Development and deployment: Build and integrate production-ready systems in the customer's environment, including data pipelines, APIs, cloud infrastructure and, increasingly, RAG, vector databases, AI agents, evaluation and monitoring.
5. Exit and knowledge transfer. Document, train and equip the customer to operate the solution independently. Success means sustained measurable value, not merely going live.

The multidisciplinary skillset

The ideal FDE has a T-shaped profile: substantial depth in at least one technical discipline, combined with breadth across related domains. The foundation includes Python, SQL, APIs, software engineering, Docker, Kubernetes and cloud platforms. Applied AI skills—RAG, vector databases, agent orchestration and model evaluation—are increasingly important.

Communication, problem-solving, stakeholder management and comfort with ambiguity are equally critical. Enterprise deployments also require evaluation and governance: defining benchmarks, testing failure cases, measuring reliability and addressing privacy, security, auditability, regulatory requirements and human oversight.

- From The Editor's Desk
- The World In 5
- Signal
- Hyperfocal
- A Week In The Life
- **Perspectives**
- The FDE Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- RACE Pulse
- RACE Programs

PERSPECTIVES CONTD...

Different FDE profiles

The role includes five overlapping profiles: Builder FDEs develop and deploy customized solutions; Sales-Engineer+ FDEs lead discovery, demonstrations and proofs of concept; AI/Data FDEs implement AI and data systems; Advisory/Translator FDEs connect business needs with technical design; and Ops/Reliability FDEs manage production performance and incidents. Organizations may combine these responsibilities or distribute them across teams. A related distinction is between product builders, who customize solutions, and platform operators, who configure and operationalize existing products.

From services to scalable technology businesses

FDE delivery can take three forms: a dedicated engineer embedded with a customer; an FDE-led implementation with defined scope and outcomes; or an FDE-enabled product model that turns repeated customer problems into reusable products and platforms. The third offers the strongest scalability. Custom integrations can become reusable connectors, discovery can become standardized diagnostics, and models can become configurable services. The article suggests aiming to reuse roughly 70–80% of a solution while retaining room for customer-specific requirements. This shifts the business from bespoke consulting towards productized delivery.

The global opportunity and India's position

Technology companies increasingly recognize that selling an AI platform is only the beginning: customers also need help with data preparation, integration, workflow redesign, governance, deployment and adoption. Although the FDE title is more established in the US, similar work appears globally under titles such as Deployment Engineer, Solutions Engineer, Field Engineer and Customer Engineer. India has significant potential, given its engineering workforce, technology ecosystem and domain expertise across manufacturing, banking, retail, logistics and healthcare.

The future: specialists who can cross boundaries

The rise of FDEs does not signal the end of specialization; deep technical expertise remains the entry requirement. What distinguishes an FDE is the ability to connect disciplines, navigate organizational constraints and take responsibility for outcomes beyond a narrow technical role. The next generation of engineers will combine software, data, AI, infrastructure and business understanding to turn ambiguous problems into working, governed and economically viable systems.

RACE perspective

The scarce hire is the engineer who can build and also sit inside the customer's mess until it works. That blend of deep code and human range is hard to find and harder to fake, and it is exactly where India can move up, from billing hours to owning outcomes. Train for both, not just the code.

- From The Editor's Desk
- The World In 5
- Signal
- Hyperfocal
- A Week In The Life
- Perspectives
- **The FDE Playbook**
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- RACE Pulse
- RACE Programs

THE FDE PLAYBOOK

GETTING REAL VALUE OUT OF AI, NOT JUST A LICENCE

Five moves, ordered from fastest to most strategic, that any leadership team can act on.

MOVE 1

STOP BUYING TOOLS, START BUYING DEPLOYMENT

The era of buying an AI licence and hoping it sticks is over, because the evidence says it mostly does not. When you buy AI now, buy the deployment with it. Insist that the vendor puts engineers alongside your team to make it work in your environment, and judge the deal on the outcome delivered, not the capability demonstrated. The serious vendors already work this way, and the ones who will not are selling you the easy half.

THE MOVE: Make embedded deployment part of every serious AI purchase. Buy a working outcome in your context, not a capable product in a demo.

MOVE 2

FIND THE PROBLEM BEFORE YOU FUND THE BUILD

Most AI failures are chosen, not caused. They pick a problem that looked good in a slide and sank in reality. Borrow the FDE's discipline: before you fund a build, put a capable person inside the actual workflow to find the specific, high-value problem worth solving. Time spent on discovery is the cheapest insurance against an expensive, impressive deployment of the wrong thing.

THE MOVE: Fund discovery before you fund building. The right problem, found in the real workflow, is worth more than a clever solution to the wrong one.

MOVE 3

BUILD YOUR OWN EMBEDDED ENGINEERS

Relying entirely on a vendor's FDEs leaves your AI capability, and your institutional knowledge, outside the building. The organisations getting durable value are growing their own: engineers who sit with the business, own outcomes, and learn your context deeply. Start small, with one or two people embedded in a real problem, and treat it as a capability to build, not a project to finish.

THE MOVE: Grow an internal forward deployed capability, not just a vendor dependency. Your own embedded engineers keep the learning, and the advantage, in-house.

- From The Editor's Desk
- The World In 5
- Signal
- Hyperfocal
- A Week In The Life
- Perspectives
- **The FDE Playbook**
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- RACE Pulse
- RACE Programs

THE FDE PLAYBOOK CONTD...

GETTING REAL VALUE OUT OF AI, NOT JUST A LICENCE

MOVE 4

HIRE AND TRAIN FOR THE RARE BLEND

The role needs production engineering and real human range in the same person, and that combination is scarce and hard to interview for. Do not screen only for coding. Test for whether a person can find the real problem, earn a user's trust, and own a result. And invest in training, because the blend can be built, and the talent market for it is far tighter than the demand.

THE MOVE: Hire and develop for code plus customer range, not code alone. The deployment half of the skill is the scarce half, so train for it deliberately.

MOVE 5

MEASURE ENGINEERS ON OUTCOMES, NOT OUTPUT

The FDE model works because the engineer is accountable for a business result, not a shipped feature. Bring that discipline inside. Point your best technical people at a named business metric, give them the access and the mandate to move it, and judge them on whether it moved. It changes what they build, and it is the surest way to close the gap between what your AI can do and what it actually does.

THE MOVE: Hold your embedded engineers to a business number, not a deliverable. Outcome accountability is what turns capability into value.

Source:

- *https://openai.com/index/openai-launches-the-deployment-company/*
- *https://www.anthropic.com/news/claude-frontier-academy*
- *https://getperspective.ai/blog/palantir-forward-deployed-engineering-playbook-anthropic-openai-copying*
- *https://getperspective.ai/blog/the-forward-deployed-engineer-playbook-how-to-structure-run-and-scale-an-fde-function-in-2026*
- *https://joinplank.com/state-of-fde*



- From The Editor's Desk
- The World In 5
- Signal
- Hyperfocal
- A Week In The Life
- Perspectives
- The FDE Playbook
- **Metric Of The Week**
- **Fun Fact**
- Candid
- Book Shelf
- RACE Pulse
- RACE Programs

METRIC OF THE WEEK

ONE NUMBER, PROPERLY UNDERSTOOD

729% **growth in forward deployed engineer job postings in the year to April 2026. A role, not a tool, is the hottest thing in AI.**

WHAT IT IS: Analyses of the hiring market put the growth in forward deployed engineer postings at around 729 percent in the twelve months to April 2026, with some counts even higher. In a labour market that has otherwise cooled for engineers, demand for this one role exploded, and the pay climbed with it.

WHY IT MATTERS NOW: The single hottest hire in AI is not a researcher or a model, it is the person who makes AI work inside a real company. That tells you where the industry now believes the value, and the difficulty, actually lives. It is a market-sized vote that deployment, not capability, is the binding constraint on AI value.

HOW LEADERS USE IT: Read it as a signal about your own org chart. If your AI strategy is all tools and no embedded people who own outcomes, you are investing where the market has decided the value is not. The scarce, expensive hire is the tell. Match your spending to it.

Source:

- *https://d2beiqkhq929f0.cloudfront.net/public_assets/assets/000/249/168/original/CMR_Scaler_FDE_Opportunity_Report_19_Oct_2026.pdf*
- *https://www.scaler.com/topics/forward-deployed-engineer-jobs-grew-1165-percent/*
- *https://analyticsindiamag.com/visual-stories/three-words-two-indias*

FUN FACT

SOMETHING TO REPEAT AT DINNER

For almost twenty years, forward deployed engineer was a quirky job title that essentially one company, Palantir, used, and most of the industry quietly thought was a bit odd. Engineers who embed with customers, live on-site, own the outcome: it did not fit the clean, scalable, ship-the-software model everyone else admired. Then, in the space of a single year, OpenAI, Anthropic, Google, Microsoft, Amazon and the big consultancies all adopted it, by name, and it became the hottest job in technology. A twenty-year-old idea nobody copied turned into an overnight sensation, which is a useful reminder that the obvious, unglamorous answer is sometimes just early.

Source: https://getperspective.ai/blog/palantir-forward-deployed-engineering-playbook-anthropic-openai-copying | https://en.wikipedia.org/wiki/Forward_deployed_engineer | https://openai.com/index/openai-launches-the-deployment-company/ | https://www.anthropic.com/news/claude-frontier-academy

- From The Editor's Desk
- The World In 5
- Signal
- Hyperfocal
- A Week In The Life
- Perspectives
- The FDE Playbook
- Metric Of The Week
- Fun Fact
- **Candid**
- Book Shelf
- RACE Pulse
- RACE Programs

CANDID



Pradeepta Mishra
Vice President of Artificial Intelligence,
Beghou Consulting

https://www.linkedin.com/in/pradeepta/

The Views are Strictly Individual/Personal, Not representing my company, should not be quoted:

Q1. **The industry's answer to failing AI pilots turned out to be a person, the forward deployed engineer, not a better model. What does that tell a board about where AI value really comes from?**

Ans: I think this tells us something fundamental: AI value doesn't come from having the best model; it comes from making AI work within the business.

We've spent the last few years obsessing over model capabilities, benchmarks, and impressive demos. But the real challenge begins when we try to integrate AI into actual business processes.

Take life sciences, for example. Generating AI-driven recommendations for sales representatives is relatively easy. Making those recommendations work within CRM systems, compliance requirements, and everyday field operations is where the real value lies.

That's why forward deployed engineers are becoming important. They connect technological capabilities with business realities.

For boards, I believe the conversation needs to shift from "How many AI projects are we running?" to "How much measurable business impact are we creating?"

Ultimately, AI success isn't about model intelligence. It's about business outcomes.

Q2. **Serious vendors now arrive with embedded engineers, not just a licence. How should that change how a company buys and evaluates AI?**

Ans: I think enterprises need to fundamentally rethink how they buy AI. Traditionally, we evaluated software based on features, pricing, scalability, and vendor reputation. With AI, that's no longer enough.

Companies should evaluate not just what the technology can do, but how effectively the vendor can make it work in their environment.

A great AI demo doesn't necessarily translate into a successful enterprise deployment.

I would look at three things: the strength of the AI platform, the expertise of the engineers working alongside our teams, and the vendor's willingness to take accountability for business outcomes.

I would also move away from endless proofs of concept toward smaller, production-focused deployments with clearly defined success metrics.

Most importantly, vendors should transfer knowledge rather than create long-term dependency.

We shouldn't just buy AI technology. We should invest in partnerships that deliver measurable business value.

Q3. **Should an enterprise rely on its vendors' forward deployed engineers, or build its own embedded capability, and how do you decide?**

Ans: I strongly believe in a hybrid approach. Use vendor expertise to

- From The Editor's Desk
- The World In 5
- Signal
- Hyperfocal
- A Week In The Life
- Perspectives
- The FDE Playbook
- Metric Of The Week
- Fun Fact
- **Candid**
- Book Shelf
- RACE Pulse
- RACE Programs

CANDID CONTD...

accelerate AI adoption, but build internal capabilities to sustain it.

Vendors understand their technology better than anyone else. They know the models, platforms, limitations, and emerging capabilities. But they don't necessarily understand our business, customers, processes, and organizational complexities.

That's where internal teams become essential.

For capabilities that directly influence competitive advantage, such as pricing, customer engagement, or strategic decision-making, I would invest in building strong internal expertise.

For more standardized applications or highly specialized technical challenges, vendor engineers can play a bigger role.

The important thing is to establish clear ownership from the beginning. Vendor engineers should work alongside internal teams, not independently of them.

My philosophy is simple: leverage vendors for speed and specialized expertise, but never outsource your understanding of your own business.

Q4. **India is the world's technology deployment workforce. Is the forward deployed model a threat to that business, an opportunity to move up from it, or both?**

Ans: I see it as both, but I'm more excited about the opportunity.

India has built a remarkable technology services industry around engineering talent, delivery excellence, and scale. However, much of our traditional business model still revolves around effort, billable hours, and team size.

AI is changing that equation. A small team of highly capable engineers can increasingly deliver outcomes that previously required much larger teams.

That's certainly a threat to the traditional services model.

But here's the opportunity: India can move from being the world's technology implementation partner to becoming the world's AI transformation partner.

We already have strong engineering capabilities and deep industry experience across healthcare, banking, manufacturing, and other sectors.

What we need now is to combine technical expertise with domain knowledge, consulting skills, and ownership of business outcomes.

The future isn't about how many engineers we deploy. It's about how much business value those engineers create.

Note: CANDID features direct contributions from CXO leaders.
To contribute, contact: retina@reva.edu.in

- From The Editor's Desk
- The World In 5
- Signal
- Hyperfocal
- A Week In The Life
- Perspectives
- The FDE Playbook
- Metric Of The Week
- Fun Fact
- Candid
- **Book Shelf**
- RACE Pulse
- RACE Programs

BOOK SHELF

RETINA'S READ OF THE WEEK

RETINA RECOMMENDS · THIS WEEK'S THEME

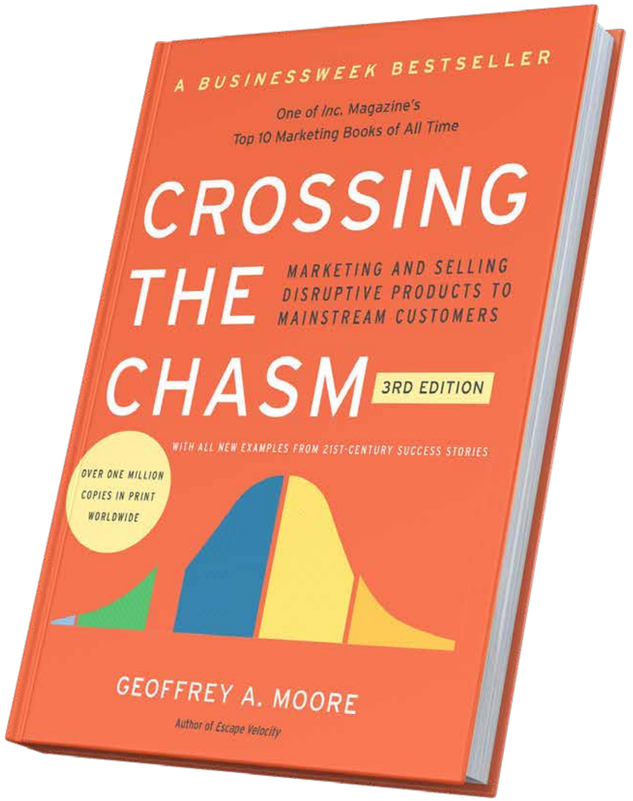
Crossing the Chasm: Marketing and Selling Disruptive Products to Mainstream Customers

Geoffrey A. Moore · HarperBusiness, 1991 (revised 2014)

Moore's classic is about the gap, the chasm, between the early enthusiasts who will tolerate a rough new technology and the mainstream customers who need it to just work. Most promising technologies die in that gap. Read against this week's cover, the fit is clean. The forward deployed engineer is, in effect, a human bridge across the chasm, the person who does the unscalable, hands-on work of making a powerful but unfinished technology deliver real value to a cautious mainstream buyer.

What it gives a leadership team is a frame for why capability alone never crosses the gap. Moore's insight is that mainstream adoption is won with the whole product, the technology plus everything around it that makes it usable in a specific context, and that winning it requires getting your hands dirty in one beachhead at a time. That is the FDE motion described thirty years early, and it explains why buying the model is never enough.

It is a marketing book, not an engineering one, which is the point. It reframes AI adoption as a go-to-market and deployment problem rather than a capability race. For a CXO trying to understand why the most advanced model did not win, and why the embedded engineer did, it is the sharpest lens on the shelf.



- From The Editor's Desk
- The World In 5
- Signal
- Hyperfocal
- A Week In The Life
- Perspectives
- The FDE Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- **RACE Pulse**
- RACE Programs

RACE PULSE

EVENTS · COMMUNITY · WHAT'S COMING

How RACE's M.Sc. in Business Analytics develops the skills of a forward deployed engineer

It's Monday morning at an automotive dealership. The parts manager has an order to place and a familiar dilemma: buy too much and cash sits on the shelf; buy too little and a customer's vehicle waits another day. A forecast can help, but the manager still has questions. Which parts are likely to run short? What's already on order? Can the team get those answers before the supplier's cut-off? Two related RACE capstones explored pieces of this problem. One examined inventory forecasting at a dealership. The other explored a generative AI tool that translates everyday questions about stock into SQL queries. Put together, they raise the question a forward deployed engineer (FDE) would ask: how do we get the right information into the hands of the person placing the order?

Where the engineering begins

In our M.Sc. in Business Analytics at RACE, REVA University, we ask learners to start with a decision like that one. Who makes it? What information do they use today? What would count as an improvement? In the dealership, the answer might be fewer urgent orders or fewer delays for customers. Once learners have a baseline, they can get to work with Python and SQL, checking whether records agree, looking for gaps in the data and building a pipeline they can run again. The business question gives the technical work a purpose. It tells learners which data matters, which errors would be costly and whether the solution can help in time.

Would you trust it with the next order?

A capstone review brings that question into the room. A forecast may look accurate overall and still miss the slow-moving part that holds up a repair. A tool that turns plain-language questions into SQL may give a plausible answer while counting stock already reserved for another job. Learners have to show how they tested for failures like these, explain where the system needs human judgement, and consider who should have access to the data. They also have to explain what would need to happen before a team could rely on the solution. That is where the programme's work on data engineering, AI evaluation and governance comes together.

Many working professionals arrive at RACE already knowing where a process breaks down. A planner knows when a forecast arrives too late. A finance analyst knows which figures must reconcile before anyone acts. The master's gives them practice turning that knowledge into a system they can build, test and explain to the people who will use it. Back at the parts counter, the test is simple. When the next order is due, can the manager make a better decision and explain why?

To know more, call:
+91 89040 58866

enquiry@race.reva.edu.in ·
https://race.reva.edu.in/pg-diploma-msc-in-business-analytics/

- From The Editor's Desk
- The World In 5
- Signal
- Hyperfocal
- A Week In The Life
- Perspectives
- The FDE Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- **RACE Pulse**
- RACE Programs

RACE PULSE CONTD...

Someone has to train 10,000 of them

When a frontier lab says it wants to train ten thousand forward deployed engineers, it is naming a skills gap at national scale. The FDE blend, real engineering plus the range to deploy it inside a business and own the outcome, is exactly the capability RACE was built to develop. Our programmes turn working professionals into people who can build AI and make it work in the real world, which is the whole of this role.

Certified Agentic AI Engineer

The closest thing we run to an FDE track: building and deploying real agentic systems, with governance and a capstone that ships something that actually works, not a demo. Built for engineers who want to move from writing code to owning outcomes.

To know more, call:
+91 89040 58866

enquiry@race.reva.edu.in ·
https://race.reva.edu.in/certified-agentic-ai-engineer

Master's and PG programmes

Applications are open across RACE's Master's and PG programmes: Business Analytics (ranked #1 in India), Artificial Intelligence (#3), Cybersecurity (#5), and Cloud Architecture and Security. Weekend format, AICTE and UGC approved, built for professionals who want to build AI that delivers in production.

To know more, call:
+91 89040 58866

enquiry@race.reva.edu.in
https://race.reva.edu.in/programs/

RACE Sandbox and CXO Circle

RACE Sandbox, the on-campus career-simulation open day, returns through the term. For senior leaders, the CXO Circle breakfast continues as a closed-door peer conversation on the shifts RETINA covers.

Next dates and slots at race.reva.edu.in, or call +91 89040 58866.

RETiNA
AI · CYBERSECURITY · CLOUD · LEADERSHIP

- From The Editor's Desk
- The World In 5
- Signal
- Hyperfocal
- A Week In The Life
- Perspectives
- The FDE Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- RACE Pulse
- **RACE Programs**

NAAC GRADE A+ ACCREDITED UNIVERSITY

REVA RACE REVA Academy for Corporate Excellence

REVA UNIVERSITY
Bengaluru, India

Master's & PG Programmes

New Batch: November 2026 · 12–24 months · Weekend · For Working Professionals

#1 IN INDIA · ANALYTICS INDIA MAGAZINE

ANALYTICS

PG Diploma / M.Sc. in
BUSINESS ANALYTICS
AWS Academy · Microsoft Azure

12–24 Months

#3 IN INDIA · TECHGIG

ARTIFICIAL INTELLIGENCE

M.Tech. in
ARTIFICIAL INTELLIGENCE
AWS Academy · Microsoft Azure

24 Months

#3 IN INDIA · TECHGIG

ARTIFICIAL INTELLIGENCE

PG Diploma / M.Sc. in
ARTIFICIAL INTELLIGENCE
AWS Academy · Microsoft Azure

12–24 Months

#5 IN INDIA · ANALYTICS INDIA MAGAZINE

CYBERSECURITY

M.Tech. in
CYBERSECURITY
EC-Council · Microsoft Azure · Terralogic

24 Months

#5 IN INDIA · ANALYTICS INDIA MAGAZINE

CYBERSECURITY

PG Diploma / M.Sc. in
CYBERSECURITY
EC-Council · Microsoft Azure · Terralogic.

12–24 Months

CLOUD

PG Diploma / M.Sc. in
CLOUD ARCHITECTURE & SECURITY
AWS Academy · Microsoft Azure

12–24 Months

Short-Term Certification Programmes

100 Hours · Project-Driven · Customisable for Organisations

AGENTIC AI | 4–6 weeks
AGENTIC AI SYSTEMS DESIGN
Multi-agent systems using LangChain, LangGraph, CrewAI. Capstone: deployed autonomous agent.

AI ENGINEER | 8–12 WEEKS
AI ENGINEERING CERTIFICATION
Data pipelines, model training, LLMs, deployment. Real production projects every batch.

CLOUD | 8–12 WEEKS
CLOUD ARCHITECTURE & DEVSECOPS
Multi-cloud, IaC, secure CI/CD on real AWS/Azure/GCP environments. Not demos.

CYBERSECURITY | 8–12 WEEKS
CYBERSECURITY & SOC OPERATIONS
Threat detection, incident response, SIEM/SOAR. Hands-on in RACE's live Security Operations Centre.

DATA & AI | 12 WEEKS
DATA SCIENCE & BUSINESS ANALYTICS
Statistics to ML to BI. Python, SQL, Power BI, Tableau on real business datasets.

LEADERSHIP | 8–12 WEEKS
TECHNO-FUNCTIONAL LEADERSHIP
Executive presence, decision-making, leading tech-driven teams. Built for mid-to-senior managers.

Enroll with us · race@reva.edu · race.reva.edu.in

NAAC GRADE A+ ACCREDITED UNIVERSITY

REVA RACE

REVA Academy for Corporate Excellence

RETiNA

AI · CYBERSECURITY · CLOUD · LEADERSHIP

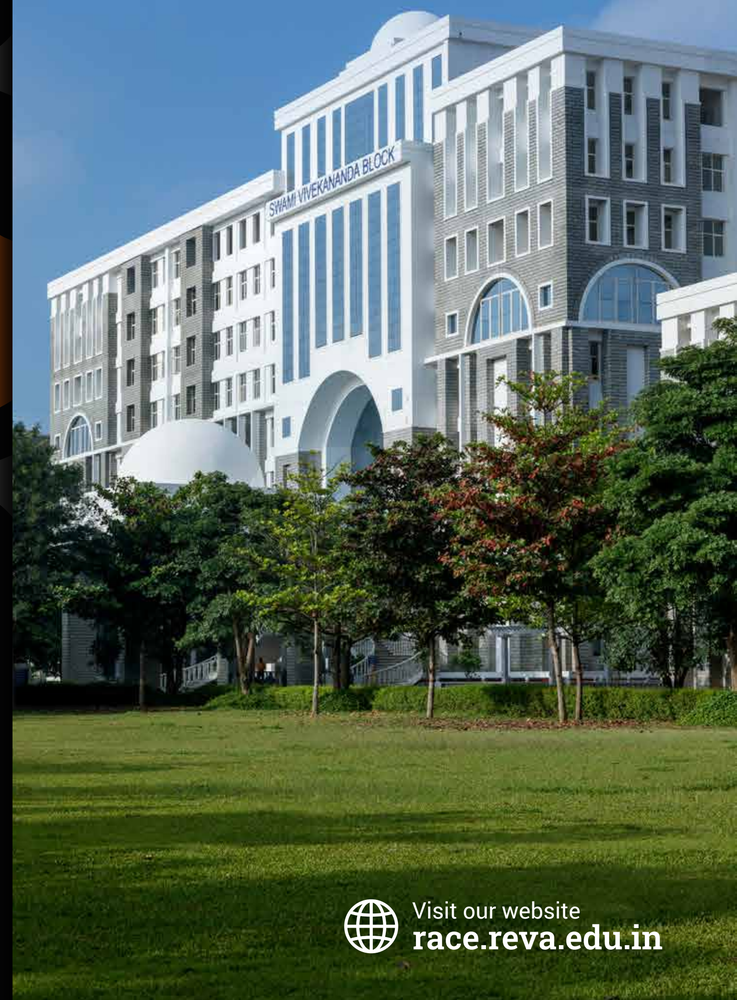
THIS WEEK, ASK YOUR LEADERSHIP TEAM

"On our most important AI effort, do we have a person accountable for the outcome, embedded deep enough in the real workflow to make it work, or did we buy a capable tool and hope?"

REVA Academy for Corporate Excellence,
REVA University
Rukmini Educational Charitable Trust
Kattigenahalli, Yelahanka, Bangalore, Karnataka, India, 560 064
Phone: +91 89040 58866

RETINA PUBLISHES EVERY SUNDAY · FORWARD TO A CXO YOU RESPECT

REVA UNIVERSITY
Bengaluru, India



Visit our website
race.reva.edu.in