The OpenAI Frontier Playbook

What It Is, What It Costs, Whether You're Ready, and How to Get Started

For CIOs, CTOs, business leaders, and transformation teams.

Leverage Strategies | 2026 | By Dan Albasry

SECTION 01

The Frontier Moment

What changed, why it matters, and why the old approach broke.

A CIO, a business unit leader, and a CISO are sitting in the same room looking at the company's AI portfolio. The CIO sees fourteen different AI experiments running across six departments, none of which share data or governance. The BU leader sees a customer service bot that is great but cannot access the CRM, and a sales assistant that writes emails but has no idea what the pipeline looks like. The CISO sees three shadow AI deployments that nobody approved, two of which have access to customer PII.

They are all looking at the same organization. None of them have the same problem. And the tool they are all waiting for is the one that connects everything.

That is what OpenAI is trying to build with Frontier. Launched on February 5, 2026, Frontier is OpenAI's enterprise platform for AI agents that function as AI coworkers, not chatbots.

It is the company's bet that the bottleneck in enterprise AI has shifted. The models are capable enough. The bottleneck is no longer intelligence. It is infrastructure: how agents are built, connected to real business data, governed, and run at scale inside organizations.

[1] OpenAI, "Introducing Frontier," Feb 5, 2026 [2] OpenAI Frontier product page, openai.com/business/frontier

75%

Enterprise workers say AI helps with new tasks

~3 days

OpenAI's average ship cycle

40%

Enterprise apps embedding agents by end of 2026

[1] OpenAI, 2026 [2] Gartner, December 2025

[1] OpenAI, "Introducing Frontier," Feb 5, 2026 [5] Gartner, Enterprise AI Predictions, Dec 2025

SECTION 02

What Frontier Actually Is

The operating system where your AI coworkers live.

The simplest description: Frontier is the operating system where your AI coworkers live. Not ChatGPT Enterprise with a new name. Not an API. A platform that sits between your AI models, from any vendor, and your business systems, and provides the infrastructure for agents to work across all of them with shared context, proper governance, and measurable performance.

Frontier is built on four pillars.

[2] OpenAI Frontier product page, openai.com/business/frontier [3] Leverage Strategies analysis, Feb 2026

Pillar 1: Business Context. Frontier connects to your data warehouses (Snowflake, BigQuery, Databricks), your CRM (Salesforce, HubSpot), your ticketing systems (Zendesk, ServiceNow), your communication tools (Slack, Teams), and your document stores. It creates a unified semantic layer that gives agents institutional memory.

Pillar 2: Agent Execution. The runtime where agents work. Agents on Frontier can plan multi-step workflows, call tools and APIs, run in parallel, build memory from past interactions, and complete tasks across environments. This is not a chatbot. It is a system that can receive a ticket, pull account history, identify root cause, draft a resolution, and escalate only when the issue requires judgment.

[1] OpenAI, "Introducing Frontier," Feb 5, 2026 [2] OpenAI Frontier product page, openai.com/business/frontier [3] Leverage Strategies analysis, Feb 2026

Pillar 3: Evaluation and Optimization. Built-in feedback loops, metrics, and experiments so agent performance improves over time. Quality becomes observable and governable rather than assumed. This is the piece most companies cobbling together AI tools are missing entirely.

Pillar 4: Enterprise Security and Governance. Agent identity and IAM, explicit permissions, guardrails, audit logs recording every action, and compliance certifications: SOC 2 Type II, ISO 27001, ISO 27017, ISO 27018, ISO 27701, and CSA STAR.

[1] OpenAI, "Introducing Frontier," Feb 5, 2026 [2] OpenAI Frontier product page, openai.com/business/frontier [7] Futurum Group, Frontier Platform Analysis, Feb 2026

Multi-vendor: manages agents built on OpenAI, Anthropic, Google, Microsoft, or open-source models. Not locked to one provider.

Interface-agnostic: agents accessible through ChatGPT, Atlas (OpenAI's AI browser), Slack, Teams, or existing business apps. Not trapped behind one UI.

Open standards: no proprietary formats, no forced replatforming. Designed to integrate with your existing stack.

Forward Deployed Engineers: OpenAI specialists who embed with customer teams to co-build agents in production.

[1] OpenAI, "Introducing Frontier," Feb 5, 2026 [3] Leverage Strategies analysis, Feb 2026 [7] Futurum Group, Frontier Platform Analysis, Feb 2026 [8] OpenAI enterprise documentation, 2026

SECTION 03

Three Deployment Patterns

From AI teammates to strategic agents.

Pattern 1: AI Teammates. Agents that sit alongside individual roles doing analysis, drafting, answering questions, and handling routine work within a specific function. A financial analyst gets an AI teammate that pulls data, generates charts, and drafts weekly reports. A software engineer gets one that reviews code, writes tests, and handles documentation.

What changes when it works: individual contributors get hours back every week. Senior people focus on judgment and strategy instead of data assembly.

[2] OpenAI Frontier product page, openai.com/business/frontier [3] Leverage Strategies analysis, Feb 2026

Pattern 2: Business Process Agents. Agents that orchestrate multi-step workflows across systems of record. A revenue ops agent that monitors new leads, enriches them, scores them, routes to the right rep, and drafts personalized outreach. A support agent that triages tickets, pulls account history, resolves routine issues autonomously, and escalates complex ones with full context.

What changes: entire workflows that require multiple handoffs run end-to-end with minimal human intervention. The human role shifts from executing steps to supervising the process and handling exceptions.

[2] OpenAI Frontier product page, openai.com/business/frontier [3] Leverage Strategies analysis, Feb 2026

Pattern 3: Strategic Project Agents. Multi-department agents deployed against high-value initiatives spanning organizational boundaries. Supply chain optimization across inventory, demand, suppliers, and logistics. Pricing that analyzes market conditions, competitors, elasticity, and margins in real time. Claims processing from intake through investigation, fraud detection, and resolution.

What changes: strategic initiatives that require cross-functional war rooms and months of analyst time become continuously operating systems.

[2] OpenAI Frontier product page, openai.com/business/frontier [3] Leverage Strategies analysis, Feb 2026

SECTION 04

What Early Adopters Are Seeing

Production results from HP, Intuit, Oracle, State Farm, and Uber.

97%

Cycle time reduction in manufacturing

90%+

More customer-facing time for sales

$1B+

Revenue from 5% output gain in energy

Early adopter production results, OpenAI Frontier preview 2026

[1] OpenAI, "Introducing Frontier," Feb 5, 2026 [3] Leverage Strategies analysis, Feb 2026

What the successful deployments share: they started with clearly defined use cases with measurable business value. They had data that was accessible and clean enough for agents to use. They worked closely with OpenAI's Forward Deployed Engineers. And they treated the deployment as a joint build, not a software purchase.

In manufacturing, a production optimization process that took six weeks was reduced to one day. In hardware testing, root cause identification went from approximately four hours per failure to minutes, saving thousands of engineering hours annually.

[1] OpenAI, "Introducing Frontier," Feb 5, 2026 [3] Leverage Strategies analysis, Feb 2026

SECTION 05

Why Most Enterprises Aren't Ready Yet

The prerequisites the marketing materials skip.

The Connectivity Gap. Data is scattered across warehouses, SaaS tools, applications, spreadsheets, and email threads with inconsistent access. Unified metadata layers cataloging data across all systems remain rare. Frontier can connect to your data, but it cannot fix the fact that your data is fragmented. If your customer data lives in four systems with four formats and four definitions of "active customer," the agent will struggle the same way your analysts do.

[16] Atlan, "AI Agent Readiness," 2026

The Semantic Gap. Business definitions live in Confluence pages, Slack threads, spreadsheets, and tribal knowledge, not in machine-readable formats agents can query. "Customer," "revenue," and "active" mean different things across departments.

The Experience Gap. Workflow knowledge lives in people's heads: exceptions, edge cases, judgment calls, escalation procedures. The person who has done the job for ten years knows that when the client says X, they actually mean Y, and you should do Z. That knowledge is invisible to any agent unless someone captures it.

[16] Atlan, "AI Agent Readiness," 2026

Frontier provides the agent platform.
You must supply the context layer.

Unified metadata, machine-readable definitions, captured workflow knowledge.
Without this infrastructure, you cannot deploy reliably.

SECTION 06

Five Challenges Executives Need to See

Go in informed, not naive.

Challenge 1: Data and Context Readiness. You cannot plug Frontier into dirty, fragmented systems and expect coherent agents. The connectivity, semantic, and experience gaps are preconditions, not nice-to-haves.

Challenge 2: Forward Deployed Engineer Dependency. The platform is not yet fully self-service. Every successful early deployment involved FDE support. If you need vendor engineers for every meaningful deployment, how do you expand across business units without linear increases in vendor headcount?

[10] Kore.ai / LinkedIn enterprise analysis, Feb 2026 [16] Atlan, "AI Agent Readiness," 2026

98%

Enterprises deploying agentic AI

79%

Operating without formal AI security policy

40%+

Agentic AI projects canceled by 2027

Gartner, 2025-2026 enterprise AI governance research

[17] Enterprise Management Associates, Dec 2025 [18] Gartner, Agentic AI Predictions, 2025-2026

Challenge 3: Governance and Security Gap. Frontier has strong governance features. But the broader challenge is organizational: decision accountability when agents approve loans or deny claims, shadow AI from employees deploying unauthorized agents, runaway costs from unchecked inference spending, and the absence of policy-as-code where rules are documented but not enforced by the system.

Governance that is too rigid drives teams toward shadow AI. Governance too permissive limits agents to low-impact tasks. The winning platforms strike a narrow balance.

[6] Gartner, Agentic AI Governance Report, 2025 [17] Enterprise Management Associates, Dec 2025 [18] Gartner, Agentic AI Predictions, 2025-2026

Challenge 4: Platform Maturity. Frontier launched February 5, 2026. It is brand new, competing against offerings from Salesforce, Microsoft, and Kore.ai with years of production hardening. Multi-agent orchestration at production scale is largely unproven at Frontier's level.

Challenge 5: Change Management. Introducing AI coworkers requires organizational change beyond technical implementation: role definition, training, success metrics, and feedback channels for reporting agent mistakes. These are leadership problems, not technical problems.

[3] Leverage Strategies analysis, Feb 2026 [10] Kore.ai / LinkedIn enterprise analysis, Feb 2026 [12] OpenAI Frontier sales page, Feb 2026

SECTION 07

How to Get Frontier

Access, process, and realistic timelines.

Frontier is in limited preview as of February 2026. Broader general availability is expected over the coming months. There is no self-serve signup. Go to openai.com/business/frontier/ and click "Contact Sales," or reach out to your existing OpenAI account team.

The sales process follows a standard enterprise pattern: discovery call (30-60 minutes), deep-dive workshop with technical and business stakeholders, proposal with specific agents, integrations, and FDE scope, then pilot deployment with two to three use cases in production before expanding.

[3] Leverage Strategies analysis, Feb 2026 [12] OpenAI Frontier sales page, Feb 2026

Realistic timelines vary by complexity. Simple pilots with well-defined use cases and accessible data: weeks. Complex multi-system integrations with change management: several months. General guidance: plan for two to six months from initial engagement to production deployment for meaningful use cases.

What gets you prioritized: showing up with a concrete story, specific use cases, demonstrated data readiness, and the right stakeholders in the room. OpenAI is triaging demand and the organizations that make the engagement efficient get faster access.

[3] Leverage Strategies analysis, Feb 2026

SECTION 08

What You Should Expect to Pay

Enterprise sales model with custom pricing.

Cost Drivers

Cost Driver	Description
Agent count and complexity	More agents doing complex multi-step workflows cost more
Data volume processed	The amount of data agents consume from connected systems
Model usage	API calls to GPT-5.x, GPT-4.x, or other models per-token
Deployment environment	OpenAI-hosted vs self-hosted affects pricing structure
FDE support level	Forward Deployed Engineer involvement adds to engagement cost

ChatGPT Team at $25-30 per user per month is not the comparator. Frontier sits above that as a platform plus services layer. The mental model is closer to how you budget for a CRM or data warehouse, weighted by scope. Industry observers note that enterprise AI platforms typically involve six-figure to seven-figure annual commitments depending on scale.

Frame the ROI around early adopter results: FTE hours saved, cycle time reduction (97% in manufacturing), error reduction, revenue deltas ($1B+ for one energy producer), and the cost of the alternative: continuing to run fragmented, ungoverned AI experiments that do not compound.

[3] Leverage Strategies analysis, Feb 2026 [14] OpenAI pricing page, 2026

SECTION 09

Frontier vs. Everyone Else

The competitive landscape.

Platform Comparison

Platform	Focus	Strength	Limitation
OpenAI Frontier	Cross-system agent mgmt	Multi-vendor, FDE support	New, limited availability
Salesforce Agentforce	CRM-centric agents	Deep SF integration	Ecosystem-locked
Microsoft Copilot	Productivity + cloud	Enterprise trust, M365	Complex licensing
Google Vertex AI	Cloud-native dev	Strong data/ML infra	Less enterprise exp.
LangChain	Open-source frameworks	Dev community, flexibility	No enterprise governance
Kore.ai	Enterprise VA	450M daily interactions	Less cutting-edge models

Salesforce and OpenAI announced a partnership in October 2025 where OpenAI models power Agentforce and Salesforce apps integrate into ChatGPT, suggesting complementary positioning rather than pure competition.

Both OpenAI and Anthropic are preparing for public offerings, making enterprise revenue traction and platform stickiness critical for IPO narratives. The competitive insight that matters most: the first platforms to convince enterprises they can safely delegate real authority to agents will gain disproportionate traction. Trust compounds.

[3] Leverage Strategies analysis, Feb 2026 [7] Futurum Group, Frontier Platform Analysis, Feb 2026

SECTION 10

How to Get Your Organization to Say Yes

The methodology for building internal alignment.

The technology works. The use cases are clear. The early adopter results are compelling. But none of that matters if you cannot get your own organization to commit the budget, the team, and the political will to move forward.

The single most game-changing thing you can do is pre-sell the idea before the formal budget conversation. If your CFO is hearing about Frontier for the first time in the budget meeting, you have already lost.

[3] Leverage Strategies analysis, Feb 2026

Pre-selling with partners. In your next one-on-one with the CTO, mention what you have been reading about Frontier's multi-vendor agent management and ask how they are thinking about governing the AI tools already in production. In your CISO sync, bring up the shadow AI problem. These are not asks. They are conversations that build shared awareness.

The goal is that by the time you bring the formal proposal forward, two or three of your peers can say, unprompted, that they have been discussing this with you and see the value.

Pre-selling with approvers. Start with competitive context two to three weeks before your formal request: share an article about Uber deploying Frontier, a Gartner quote, or the 40% adoption forecast. Then share internal data one to two weeks before: your AI tool inventory, shadow AI exposure, governance gaps.

If your budget conversation contains any information your approver is seeing for the first time, you have not pre-sold enough.

The Stakeholder Map

Stakeholder	What They Own	What They Will Ask
CEO / BU Leader	Strategic rationale	Business impact? What are competitors doing?
CIO / CTO / Data	Integration, architecture	How does it integrate? Dependencies?
CISO / Risk / Legal	Governance, compliance	SOC 2? ISO? Audit trails?
Ops / Function Heads	Workflows agents run	Will this work for my process?
CFO / Finance	ROI, TCO vs alternatives	What does it cost? Payback period?

The CFO. Your CFO is not hostile to investment. They are hostile to unquantified risk and vague ROI promises. Frame Frontier in terms they already use: show them the fully loaded cost of your current AI portfolio, the engineering time maintaining integrations, the compliance risk. Then show what a unified platform looks like.

Do not lead with technology. Lead with the financial model. The CFO does not care about multi-agent orchestration. They care about whether this consolidates spend, reduces risk exposure, and shows measurable ROI within a defined timeframe.

The CEO cares about competitive positioning over two to three years. Show them that HP, Intuit, Oracle, State Farm, and Uber are already in production. Where are your direct competitors? The gap between early movers and laggards widens as compounding effects take hold.

The CISO is not trying to block innovation. They are trying to prevent the $670,000 breach premium from ungoverned AI. Show them Frontier's governance architecture, then show them what you have today: 79% of enterprises without formal security policies for autonomous tools.

[1] OpenAI, "Introducing Frontier," Feb 5, 2026 [3] Leverage Strategies analysis, Feb 2026 [17] Enterprise Management Associates, Dec 2025

The Three-Bucket Rationale. Every dollar your organization spends can be categorized into one of three buckets: it drives revenue, it reduces cost, or it manages risk. The most defensible investment proposals show impact across all three. Enterprise AI platforms are one of the rare investments that genuinely do.

Bucket One: Drives Revenue. Frontier enables agents that accelerate revenue-generating workflows. The early adopter proof points: 90% more customer-facing time for sales teams, 97% cycle time reduction in manufacturing. These are not theoretical projections. They are measurements from production deployments.

Bucket Two: Reduces Cost. Most enterprises are running a fragmented AI portfolio of eight to fifteen tools that do not share data, governance, or context. A unified platform consolidates much of this. The tool consolidation alone often pays for a meaningful portion of the platform investment.

[1] OpenAI, "Introducing Frontier," Feb 5, 2026 [3] Leverage Strategies analysis, Feb 2026

Bucket Three: Manages Risk. This is the bucket most internal proposals miss, and it is often the one that tips the decision. The risk is twofold: the security risk of ungoverned AI (shadow deployments, unaudited agent actions, compliance exposure) and the competitive risk of falling behind (40% enterprise agent adoption by end of 2026 means your competitors are moving now).

Framing the risk bucket explicitly gives your approver permission to invest proactively rather than reactively. It shifts the conversation from "can we afford to do this" to "can we afford not to."

[5] Gartner, Enterprise AI Predictions, Dec 2025 [17] Enterprise Management Associates, Dec 2025

Present Options, Not a Single Proposal

	Option A: Full Deploy	Option B: Phased Pilot	Option C: Status Quo
Scope	3-5 use cases, dedicated team, FDE support, 90-day build	Single use case, one team, 60-day pilot, expand if results warrant	Continue fragmented portfolio, manual governance
Investment	6-7 figure annual (platform + FDE + team)	Low 6 figures first year	No new cost, but growing shadow AI risk
Outcome	Unified AI infra, multiple agents compounding	Proof of concept data in 60 days	Silos persist, no compounding, risk grows
Risk	Higher upfront, mitigated by consolidation savings	Slower to compound, single use case may not show full value	Shadow AI grows, competitors deploy, talent leaves

"We already have AI tools." There is a difference between AI tools scattered across departments and an AI platform that governs them. One gives individuals productivity gains. The other gives the organization a capability that compounds.

"This sounds expensive." Add up your current AI tool portfolio spend, engineering time maintaining integrations, compliance risk from ungoverned deployments, and opportunity cost of agents that cannot share context. The platform alternative typically consolidates a meaningful portion of that spend.

"What if it doesn't work?" This is why the phased pilot exists. One use case, one team, sixty days. Clear success metrics before you start. If it misses, you have spent a defined amount to learn something important.

"OpenAI is too new at enterprise." Legitimate concern. The tradeoff: Salesforce and Microsoft are ecosystem-locked. If your agents need to work across systems, Frontier's multi-vendor design is more aligned.

"We should wait." The companies deploying now are building compounding advantages. Waiting is not neutral. It is a decision to let competitors build while you stand still.

The Internal Approval Narrative. Build it like a story with six beats: (1) Problem: ungoverned AI bots, no shared context, no auditability. (2) Opportunity: a single control plane for AI coworkers with governance built in. (3) Why now: competitors are embedding agents into core workflows. (4) Proof points: 97% cycle time reduction, 90% more customer time, $1B+ revenue impact. (5) Risk mitigation: SOC 2 Type II, ISO certifications, full audit trails. (6) Investment frame: six-figure annual commitment, three options for commitment level.

The Internal Champion Checklist. Weeks 1-2: Build the foundation. Inventory your current AI portfolio: every tool, every bot, every shadow deployment. Document the spend, governance gaps, integration engineering time. Gather competitive intelligence: which companies in your market are deploying enterprise AI platforms?

Weeks 2-3: Plant seeds with partners. Have informal conversations with your CTO, CISO, and business unit leaders. Not pitching. Just asking questions and sharing observations.

Weeks 3-4: Frame the world for approvers. Share competitive intelligence and market data with your CFO and CEO. Forward an article. Send a one-paragraph email with the Gartner data. You are not asking for anything yet.

Week 4-5: Share the internal data. Frame without editorializing: here is our AI inventory, here is the spend, here is the governance gap.

Week 5-6: Deliver the proposal. By now your approver has the context and has heard from your peers. Present three scenarios. The approval rate for pre-sold proposals is dramatically higher than for cold pitches.

SECTION 11

Are You Frontier-Ready?

An honest readiness assessment.

Must-haves before engaging OpenAI: two to three clearly defined agent roles with measurable business value. Systems that expose APIs agents can consume. Secure authentication mechanisms established. Data governance policies defined. Data quality audited for target use cases. Internal team identified to work alongside FDEs.

If you have all six, you are ready to engage. If you have four or five, start the conversation while closing gaps. Fewer than four: invest in readiness first.

[3] Leverage Strategies analysis, Feb 2026 [16] Atlan, "AI Agent Readiness," 2026

Should-haves before production: unified metadata layer across core systems, machine-readable business glossary, workflow knowledge captured in queryable format, agent authorization structures defined, audit and monitoring requirements specified, change management plan.

Red flags that mean "not ready yet": data scattered across 10+ systems with no unified catalog, business definitions only in people's heads, no internal team with bandwidth, compliance team not briefed, budget conversation not started, "we want to try AI" without specific use cases.

[3] Leverage Strategies analysis, Feb 2026 [16] Atlan, "AI Agent Readiness," 2026

SECTION 12

The 90-Day Readiness Plan

Be worth Frontier's time.

Days 1-30: Map and Clean. Inventory every AI tool, bot, and experiment running in the organization. Include shadow AI. Kill unauthorized deployments or bring them under governance. Choose two to three candidate use cases that are high-value, clearly defined, and have accessible data.

Start building the context layer: catalog data sources, begin a machine-readable business glossary, capture workflow knowledge for your candidate use cases. This is the hardest and most valuable work.

[3] Leverage Strategies analysis, Feb 2026 [16] Atlan, "AI Agent Readiness," 2026 [17] Enterprise Management Associates, Dec 2025

Days 31-60: Audit and Govern. Run a readiness audit against the checklist with your Data, IT, and CISO teams. Define governance: who approves new agents, what data they can access, what actions they can take autonomously, how actions are logged, what the escalation path is, and who is accountable.

Establish success metrics for each candidate use case: the before state and the target state. These become the ROI hypothesis you present to Finance and to OpenAI.

[3] Leverage Strategies analysis, Feb 2026 [16] Atlan, "AI Agent Readiness," 2026

Days 61-90: Package and Engage. Build the internal brief: the problem, the opportunity, use cases with ROI hypotheses, readiness assessment, governance framework, and investment frame. Prepare the same for OpenAI: use cases, systems map, data readiness, governance posture, specific questions.

Request the demo. Come with the brief, the use cases, and the team. Bring the CIO, a business line owner, and someone from security. The organizations that get fastest access are the ones that make the engagement efficient for both sides.

[3] Leverage Strategies analysis, Feb 2026 [12] OpenAI Frontier sales page, Feb 2026

About Leverage Strategies

Leverage Strategies helps organizations evaluate, prepare for, and deploy enterprise AI platforms. From readiness assessments and governance design to use case identification and implementation planning, we work with leadership teams navigating the shift from AI experiments to AI infrastructure.

No pitch deck. Just an honest assessment of where you are and what the next step looks like.

Leverage Strategies | 2026 | By Dan Albasry