FRONTIER.

FRONTIER RESEARCH / DECISION-GRADE ANALYSIS



The Autonomous Enterprise

What enterprise autonomy can do now, where the economics work, and how CIOs can expand the safe delegation frontier without losing accountability, security, or trust.

FR-026 / Dion Hinchcliffe / September 25, 2026

REPORT PROFILE

Built to change a decision.

What enterprise autonomy can do now, where the economics work, and how CIOs can expand the safe delegation frontier without losing accountability, security, or trust.

REPORT	**FR-026**
AUTHOR	**Dion Hinchcliffe**
PUBLISHED	**September 25, 2026**
PROFILE	**frontier-ro-0.1**

This PDF is a publication rendering of the canonical Frontier Research Object. The web report remains source of truth; citations, figures, methodology, and rights metadata travel with this edition.

Executive brief

The autonomous enterprise is arriving unevenly but unmistakably.

It is not a company with no people. It is an organization in which increasing amounts of cogniti coordination, and execution are delegated to software agents operating inside explicit busines boundaries. People set intent, policy, goals, and exceptions; agents increasingly sense conditi decide among bounded options, act through enterprise systems, and learn from outcomes.

The first wave of enterprise agents was mostly about capability: can a model reason, use tools browse, code, call APIs, and finish a useful task? The second wave is about authority: can the remain present, accumulate context, hold an identity, use credentials, initiate work, coordinate agents, and continue operating when no human is actively prompting it?

That second wave is now visible in product form. Microsoft describes Autopilot as a persistent, proactive agent that lives in the tenant with its own identity, memory, computer, and workspace.[1] Meta's Muse runs inside a dedicated secure VM and is paired with a separate Sentinel agent that mediates outbound actions and permission-sensitive behavior.[2][3] SpaceXAI's Team Bots combine shared context, tools, credentials, and memory around repeatable team roles.[4] Oracle Fusion Claw combines AI reasoning with deterministic enterprise computation inside objectives, permissions, policies, approvals, and risk limits.[5] OpenAI's dots are always-on agents with their own cloud computer, connected tools, recurring work, and user-defined autonomy boundaries.[6] OpenClaw Enterprise is being built as an open, vendor-neutral control plane for persistent agents in sensitive environments.[7]

This matters because the enterprise object being governed is changing. A chatbot is a tool. A persistent agent with memory, credentials, initiative, and write authority is closer to a machine principal.

The practical opportunity in 2026 is therefore not "general autonomy." It is selective autonomy applied to economically attractive work, inside an authority model the enterprise can actually govern.

The broader market evidence already supports the move beyond demos. McKinsey's 2026 global survey finds that 40% of respondents at companies with more than $1 billion in revenue are scaling AI agents, up from 27% a year earlier; 31% of large-enterprise respondents report scaling coding agents, and 32% say agentic coding has already caused them to forgo buying at least one software product or feature.[8] Gartner's review of 107 agentic deployments reaches a complementary conclusion: the strongest value is coming from specialized agents attached to specific business processes, and Gartner expects 80% of tangible agentic-AI ROI in 2028 to come from domain-specific agents rather than general-purpose ones.[9]

Yet the readiness gap is equally real. Deloitte found that only 5% of organizations consider their business processes highly prepared for agentic adoption and only 15% have scaled orchestrated, cross-functional multi-agent systems.[10] The central tension is no longer simply that model capability is advancing faster than enterprise readiness. It is that machine authority is expanding faster than the controls organizations have built to contain it.

Frontier's conclusion is that the limiting factor has shifted. The hardest problems are increasingly not whether an agent can reason, but whether an enterprise can safely give it context, authority, tools, memory, observability, economic limits, and clear accountability.

> **Frontier view: The winner will not be the organization with the most agents. It will be the organization with the largest safe delegation frontier: the greatest amount of valuable work it can delegate without surrendering accountability, security, reliability, or trust.**

This changes the technology leadership agenda. Autonomy is becoming an architecture property, a governance property, and an economic property — not merely a model feature.

What is possible now

Enterprise AI has moved through three practical stages.

The first was assist: summarize, search, draft, classify, recommend. The second is automate: complete a bounded workflow with deterministic systems and a human fallback. The third is orchestrate: coordinate multiple systems, people, and agents toward an outcome over time.

All three are viable now. What changes sharply is the level of effort, reliability engineering, governance, and process redesign required.

The near-term sweet spot is not the most autonomous work. It is work with five characteristics:

- high transaction volume;
- repeatable or semi-structured process logic;
- reliable access to the systems of record;
- explicit policies and decision boundaries;
- a clean path to human exception handling.

  Software engineering is currently one of the clearest examples. McKinsey finds coding agents scaling faster than most other agent categories, particularly in large organizations.[8] IT service management, customer support, security triage, marketing operations, employee service, finance operations, procurement intake, and research-intensive sales workflows are similarly well suited because they combine high volumes with accessible digital systems and measurable outcomes.

  At the other end are broad cross-functional processes: end-to-end order-to-cash, autonomous supply-chain replanning, enterprise command centers, and complex finance-close orchestration. These can create larger value pools, but they cross more systems, policies, owners, and exception paths. The payoff is potentially higher, but so are the integration and governance costs.

The economic map in Figure 1 therefore matters more than a generic maturity model. The practical question is not “How autonomous can this become?” It is “Where does additional autonomy create the best completed-work economics at an acceptable risk?”

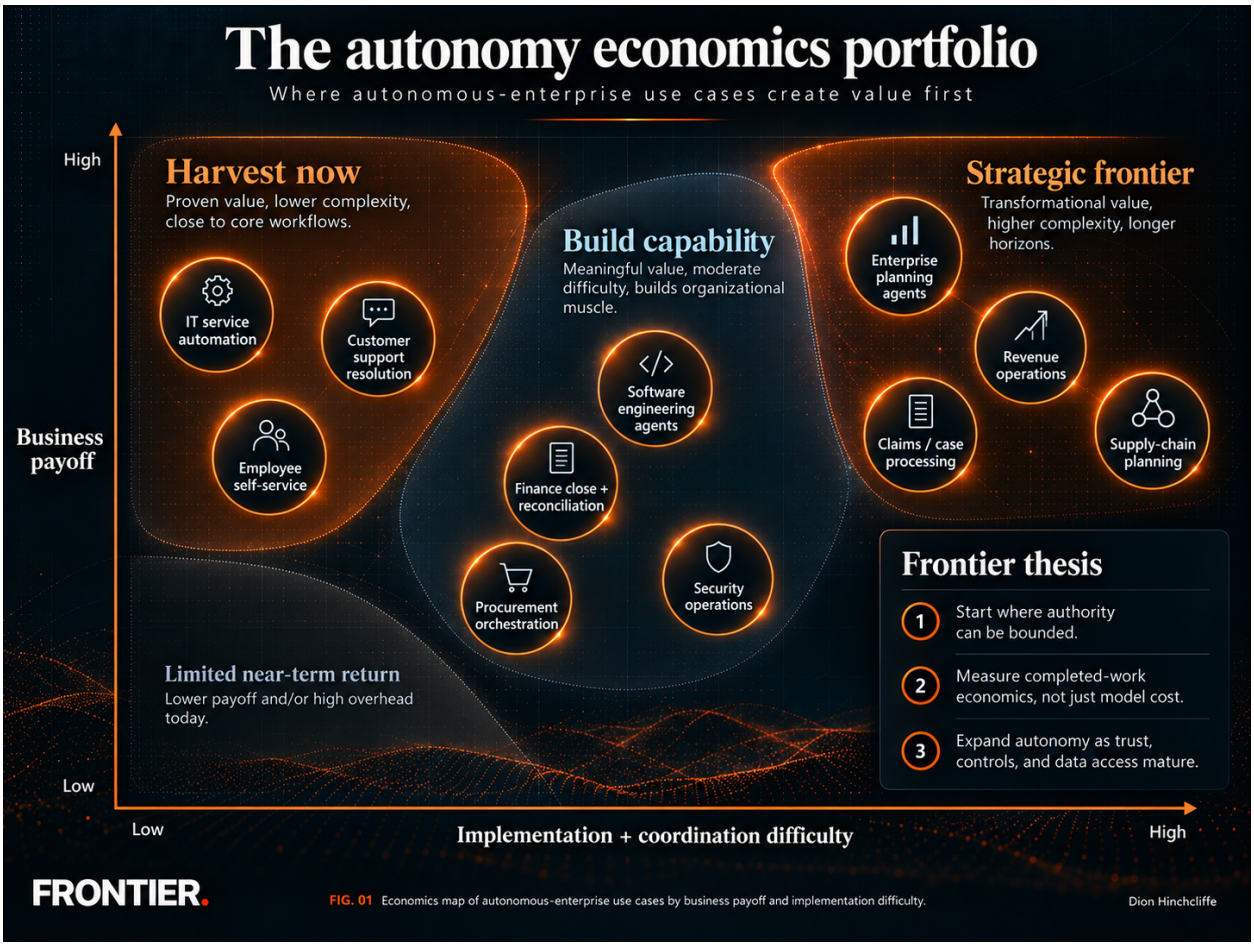


FIG. 01 The autonomy economics portfolio: where autonomous-enterprise use cases create value first.

Opportunity	Typical 2026 viability	Relative payoff	Relative effort
Meeting and document copilots	High	Medium	Low
Enterprise search and knowledge assistants	High	Medium	Low
Contact-center summarization	High	Medium	Low
HR case deflection	High	Medium	Low

IT service-desk automation	High	High	Low–medium
Customer-service assist	High	Medium–high	Medium
Sales proposal and account research	High	Medium	Medium
Software-engineering agents	High	High	Medium
Marketing-operations agents	Medium–high	Medium–high	Medium
Procurement intake and sourcing support	Medium–high	Medium	Medium
Finance close and reconciliation agents	Medium	High	High
Security-operations triage	Medium	High	High
AIOps incident response	Medium	Medium–high	High
Supply-chain planning and exception handling	Emerging	High	Very high
Cross-functional order-to-cash orchestration	Emerging	High	Very high
Autonomous enterprise command center	Emerging	Very high	Very high

The message is not to avoid difficult domains. It is to sequence them. Organizations should harvest bounded, repeatable work first, while building the architecture needed to move into cross-functional autonomy.

The economics of autonomy

Agent economics differ from traditional software economics in an important way: the cost of a workflow remains dynamic after deployment.

Model selection, token use, context size, tool calls, retries, human review, exception rates, latency requirements, and run frequency all change the unit economics. A workflow that is unattractive today can become viable after model costs fall, routing improves, or exceptions are redesigned out of the process.

McKinsey's August 2026 analysis makes the point especially clearly. In one customer-service example, token costs represented only 20% to 25% of variable run costs while human oversight represented 70% to 75%.[11] The implication is profound: simply waiting for cheaper tokens does not solve the economics of autonomy. The real economic lever is reducing expensive human intervention without allowing error costs to explode.

That requires a different ROI model.

Frontier recommends measuring cost to completed work, not cost per token, agent, seat, or interaction. A completed outcome might be a resolved incident, closed account reconciliation, onboarded customer, processed claim, qualified lead, or replenished inventory exception. The denominator must be the business result.

A useful operating equation is:

Autonomy value = (human effort avoided + cycle-time value + quality lift + revenue or risk value) " (AI run cost + integration cost + oversight cost + error cost + governance cost).

This creates four practical design principles.

First, volume matters. Fixed implementation and governance costs amortize best over large pools of repeatable work.

Second, reuse matters. Common agent capabilities—research, document understanding, policy checking, reconciliation, outreach, scheduling, exception classification—should be reusable across workflows rather than rebuilt for each one.

Third, routing matters. The most capable model should not be the default model. Organizations need model and tool routing based on task difficulty, risk, latency, and cost.

Fourth, exceptions matter most. Every workflow has a point where uncertainty, authority, or risk requires a human. The percentage of cases crossing that boundary often determines whether the economics work.

The largest opportunity is therefore not "labor replacement." It is work compression: fewer handoffs, less waiting, faster decisions, lower coordination cost, and higher capacity. McKinsey describes this as reducing the "coordination tax" embedded between functions and systems; its 2026 survey found widespread individual productivity gains but much less EBIT impact, highlighting how task productivity can fail to become enterprise value when workflows remain unchanged.[12]

BCG reaches the same conclusion from operating-model work: large gains appear when organizations redesign end-to-end processes around agentic execution instead of placing AI on top of existing workflows.[13]

Why autonomous enterprise programs stall

The most common failure mode is architectural optimism: an organization sees an impressive agent demonstration and assumes production autonomy is primarily a model-selection problem.

It is not.

Processes were designed for people

Most enterprise processes contain undocumented judgment, informal workarounds, handoffs, duplicate approvals, and tacit knowledge. Agents expose that debt immediately. Deloitte's research found that only one in five leaders believes their organization is prepared to redesign processes for autonomous operation.[10]

Automating a poor process can make it faster without making it better. The more autonomous the system, the more dangerous that becomes.

Identity and authority are harder than intelligence

An agent that can reason but cannot safely authenticate, receive scoped authority, or prove what it did is not enterprise-ready.

NIST's August 2026 guidance argues that agentic systems are repeating a familiar security pattern: shipping capability faster than identity foundations. It highlights the need to treat agents as attributable actors with strong authentication and authorization rather than as anonymous model calls.[14]

This means every production agent needs an owner, identity, permission boundary, credential strategy, allowed-tool set, audit trail, and revocation path. Standing broad credentials are the wrong pattern. Authority should be least-privilege, contextual, and time-bounded wherever possible.

Enterprise context is fragmented

Agents need more than retrieval. They need semantic context: which customer is meant, which product hierarchy applies, which policy is current, which source is authoritative, which metric definition the organization uses, and which transaction is canonical.

Knowledge graphs, semantic layers, master data, governed retrieval, and high-quality metadata become part of the execution architecture, not just analytics infrastructure.

Cross-functional autonomy has organizational owners

An agent can cross systems more easily than a company can cross silos.

Order-to-cash spans sales, legal, finance, operations, customer success, and often external partners. Supply-chain exception management crosses planning, logistics, procurement, finance, and manufacturing. Autonomous workflows force organizations to answer questions that traditional applications can avoid: Who owns the outcome? Who owns the agent? Who can change its policy? Who absorbs its errors? Who approves expanded authority?

Reliability must be measured continuously

Agents are probabilistic systems operating inside dynamic environments. Testing once before launch is inadequate.

Production autonomy requires evaluations, traces, replay, policy checks, anomaly detection, cost telemetry, outcome metrics, and controlled rollout. The appropriate unit is not merely model accuracy. It is reliable completed work within policy.

Gartner's 2026 analysis warns specifically about overestimating agent reliability, agent sprawl, unmanaged token costs, weak data foundations, and insufficient change management.[9]

Architecture of an autonomous organization

The autonomous enterprise should not be designed as a single super-agent.

It is better understood as a governed mesh of domain agents, deterministic workflows, enterprise systems, people, and shared control services. Figure 2 shows the merged business and IT architecture.

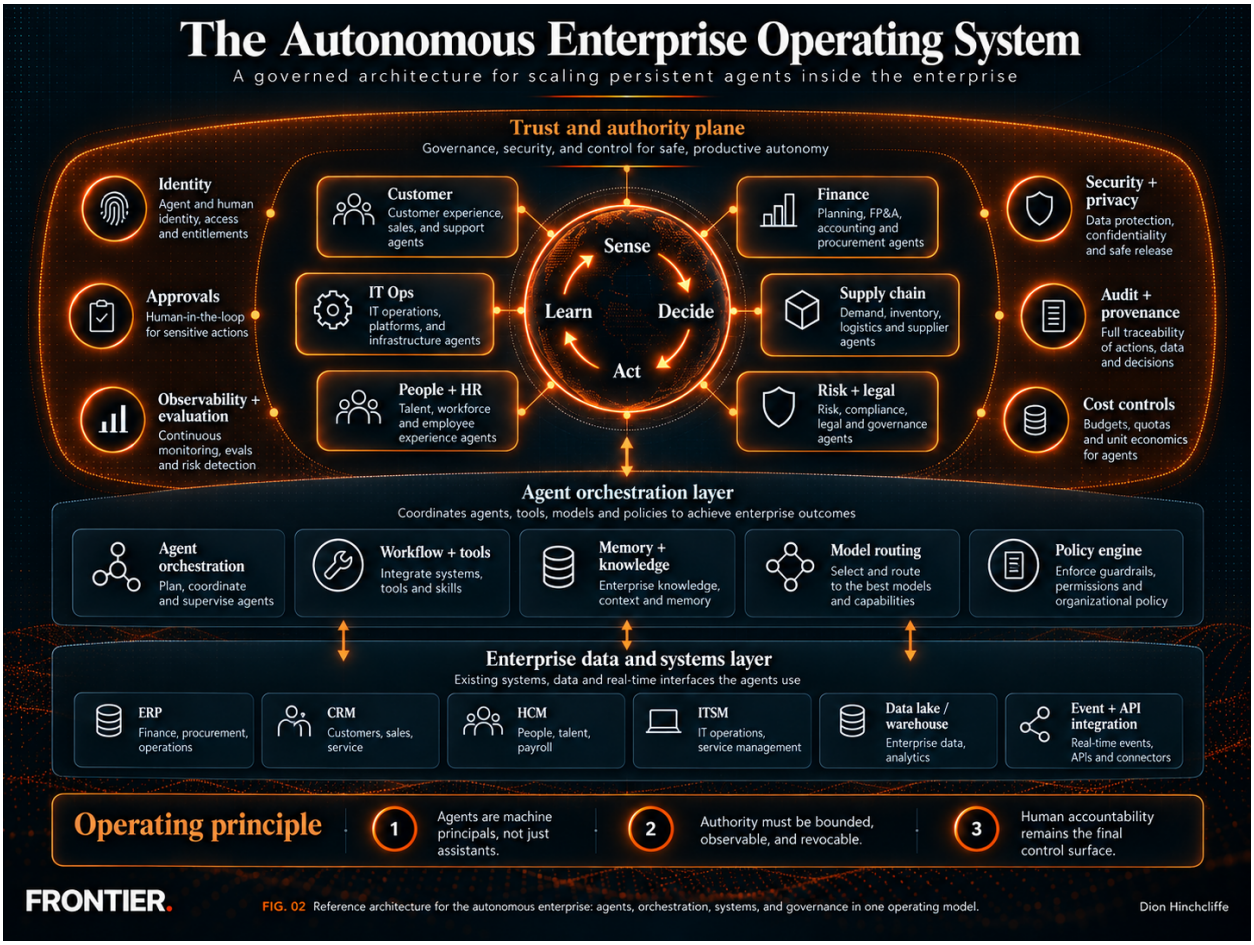


FIG. 02 The Autonomous Enterprise Operating System: agents, orchestration, systems, and governance in one operating model.

The architecture begins with a simple operating loop:

Sense !’ Decide !’ Act !’ Learn.

That loop is implemented through several layers.

1. Channels and work surfaces

Employees, customers, partners, field workers, collaboration tools, mobile applications, and conversational interfaces are where intent enters and outcomes return.

The goal is not to force every interaction through a chatbot. Autonomy should surface inside the work context where the user already operates.

2. Domain agents and digital workers

The most effective agents are increasingly specialized around a business domain or role: sales and marketing, customer service, operations and supply chain, finance and procurement, HR and workplace, IT and security.

This specialization is consistent with Gartner's deployment evidence: domain-specific agents are easier to bound, evaluate, govern, and improve than a universal enterprise agent.[9]

3. Orchestration and decisioning

This is the critical middle layer.

It contains planning, workflow, business rules, approvals, policy guardrails, exception routing, memory/context, and optimization. It determines what an agent may do, when another agent or system should act, and when a person must intervene.

This layer is where autonomy becomes enterprise-grade. Without it, organizations have agents. With it, they begin to have an autonomous operating system.

4. Knowledge and intelligence

Enterprise search, knowledge graphs, semantic layers, analytics, process mining, model hubs, and simulation/digital twins turn organizational information into usable context.

This layer is increasingly strategic because agent quality depends on the quality, freshness, and semantics of the context presented at decision time.

5. Data and integration fabric

APIs, event buses, integration platforms, RPA, identity, vector stores, lakehouses, and master data management connect agents to the operational enterprise.

The key architectural shift is from "AI reads enterprise data" to AI participates in enterprise transactions. That requires transaction-safe interfaces, permission-aware access, idempotency, and observability.

6. Systems of record and execution

ERP, CRM, HCM, SCM, ITSM, collaboration suites, industry applications, and content repositories remain essential.

The autonomous enterprise does not replace systems of record. It changes their role. They become trusted transactional substrates underneath a more dynamic layer of agentic coordination.

7. The trust plane

Identity, security, privacy, compliance, model governance, evaluations, audit, resilience, FinOps, and human override must span every layer.

The emerging market confirms this architectural direction. Microsoft Agent 365 is explicitly positioned as a cross-agent control plane for discovery, identity, governance, security, and observability.[15] AWS AgentCore places identity, gateway policy, runtime isolation, and observability around agents regardless of model or framework.[16] ServiceNow's AI Control Tower similarly emphasizes discovering, securing, governing, observing, and measuring agents across the enterprise.[17]

The architectural implication is clear: governance is moving from policy documents into the runtime path of autonomous work.

8. The human governance loop

Humans remain responsible for goals, high-impact approvals, exception handling, outcome monitoring, and changing the system.

This is not a temporary concession to immature models. It is the organizational control structure that makes scalable delegation possible.

The more capable agents become, the more important it is to define where human judgment is economically and ethically valuable rather than keeping people in every loop by default.

The safe delegation frontier

The right objective for a CIO is not maximum autonomy. It is maximum safe agency.

Every organization has a boundary beyond which delegated machine authority becomes unacceptable. On one side of that boundary, agents can act with sufficient reliability, evidence, reversibility, and control. On the other, the same agent behavior creates unacceptable operational, financial, legal, security, or reputational exposure.

Frontier calls that boundary the Safe Delegation Frontier.

The frontier is not fixed. It expands as the enterprise gets better at identity, policy, observability, evaluation, data governance, human escalation, and economic controls. A company with weak machine identity and poor runtime evidence may safely use copilots and bounded task agents but should hesitate before giving autonomous systems transactional authority. A company with mature controls can delegate more consequential work because it can prove who or what acted, constrain the action before it happens, observe it while it runs, stop it when conditions change, and reconstruct the decision afterward.

This is why governance should be understood as an enabler of autonomy, not merely a brake on it.

The progression is roughly:

- Copilots augment a person who remains continuously in the loop.
- Task agents complete discrete units of work with explicit success criteria.
- Bounded workflow agents orchestrate multi-step work inside defined guardrails.
- Domain agents operate across tools and data within a well-defined business domain.
- Transactional agents execute consequential business processes with measurable outcomes.
- Fleet orchestration coordinates many agents toward enterprise objectives under shared policy and economic constraints.

  Moving up that curve should require evidence, not enthusiasm. Each step increases the amount of authority delegated and therefore increases the value of identity, policy enforcement, runtime evaluation, provenance, data controls, rollback, and human escalation.

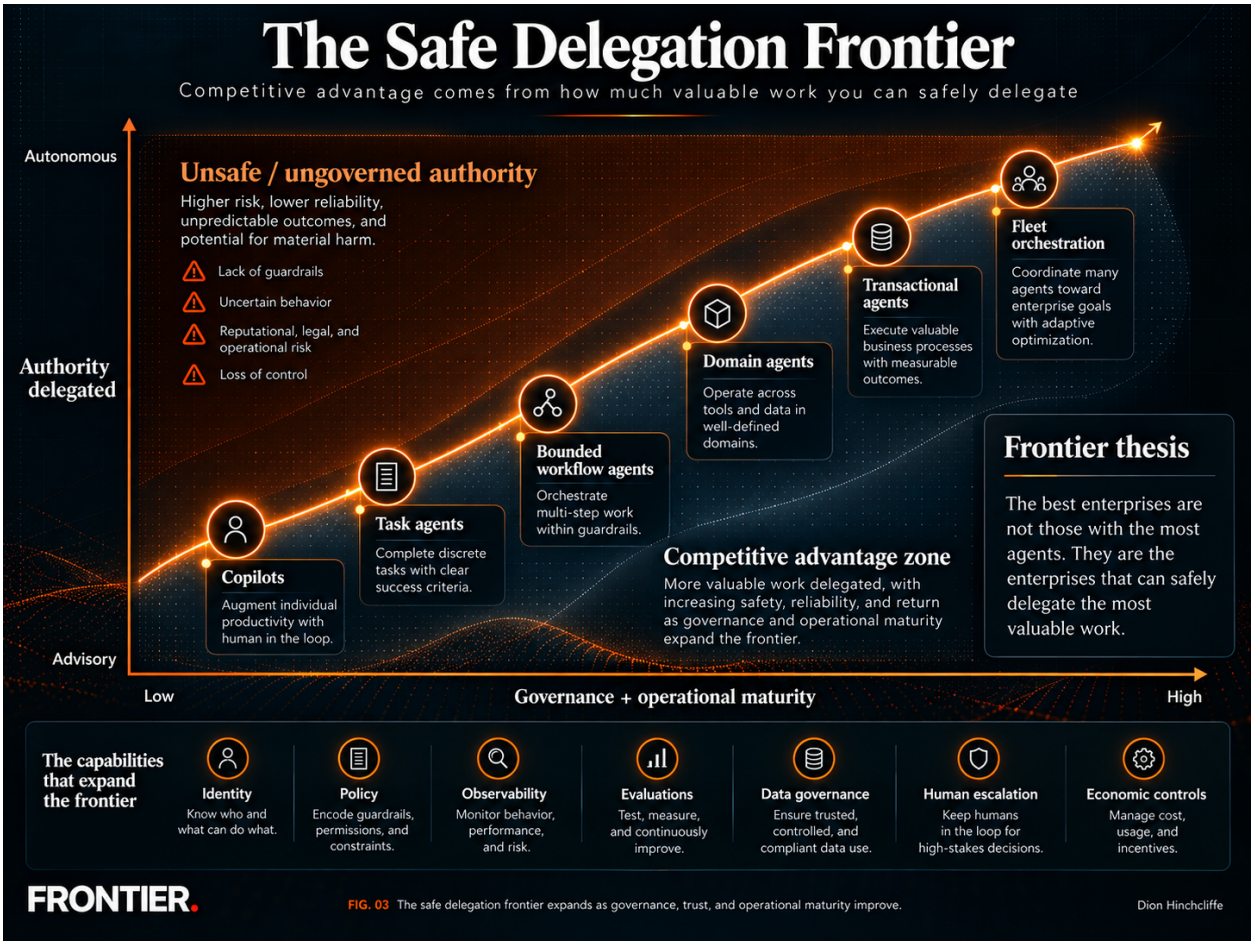


FIG. 03 The safe delegation frontier expands as governance, trust, and operational maturity improve.

The strategic implication is profound: trust becomes productive capacity. Better governance does not merely reduce downside. It increases the amount of valuable work the enterprise can hand to machines. In that sense, identity systems, evaluation infrastructure, policy engines, observability, and data controls become part of the production system for autonomy.

That creates a different competitive metric. The autonomous enterprise is not the enterprise with the highest agent count. It is the enterprise that can move the most economically valuable work across the safe delegation frontier, faster than competitors can, without allowing the control environment to fall behind.

The platform battle is moving up the stack

The strategic contest is moving beyond model quality and assistant UX toward the control plane for machine labor: identity, memory, execution, tools, policy, approvals, observability, evaluation, economics, and authority.

The late-September wave makes that shift unusually visible. Microsoft Autopilot is persistent and tenant-resident, with its own identity, memory, computer, and workspace; Microsoft is building it on OpenClaw.[1][18] Meta is taking Muse into Meta Enterprise Platform.[19] SpaceXAI's Team Bots combine shared context, tools, credentials, and memory around reusable team roles.[4] Oracle Fusion Claw puts governed agentic execution directly inside transactional business applications.[5] OpenAI's dots add always-on agents with their own cloud computer and recurring work.[6] OpenClaw Enterprise is pursuing an open, vendor-neutral control plane for persistent agents.[7]

Strategic control point	Representative platforms	What they are trying to own
Persistent work agent	Microsoft Autopilot, Meta Muse, OpenAI dots	A durable agent that remains present and continues work across time.
Shared digital coworker	SpaceXAI Team Bots, Microsoft Autopilot	Reusable team context, tools, memory, and role knowledge.
Transactional authority	Oracle Fusion Claw, SAP, Salesforce, ServiceNow, Workday	Governed execution close to systems of record and approvals.
Agent substrate	AWS AgentCore, Google, Microsoft, OpenAI	Runtime, orchestration, tools, models, identity, and policy.
Cross-agent governance	Microsoft Agent 365, Google, ServiceNow, IBM, UiPath	Inventory, lifecycle, policy, security, evaluation, and observability.
Open infrastructure	OpenClaw Enterprise	Portability across model and application vendors.
Enterprise context	Salesforce, SAP, Oracle, Workday	Proprietary semantics, workflow state, and transactional data.

Three implications follow.

First, persistent agents are becoming machine principals. Meta's Muse required a dedicated secure VM, protected credentials, a separate Sentinel agent, permission checks, and an audit trail because it can use connected applications and operate while the user is not actively supervising it.[2][3] Microsoft gives Autopilot its own tenant identity and workspace.[1] Identity and authority are no longer peripheral controls; they are part of the product architecture.

Second, application vendors retain a powerful advantage where autonomy meets transactions. Oracle's Fusion Claw explicitly combines adaptive AI reasoning with deterministic enterprise computation while bounding execution through permissions, approvals, objectives, and risk limits.[5] SAP, Salesforce, ServiceNow, and Workday are pursuing the same strategic territory from different application domains.[20][21][17][22]

Third, portability becomes more valuable as authority rises. OpenAI, Google, and AWS are all stretching from models into runtime and governance.[23][24][25][16] OpenClaw's enterprise move — and Microsoft's decision to build Autopilot on OpenClaw — shows why enterprises may want the agent runtime separated from the model and application surface.[18][7]

The CIO question is therefore no longer simply "Which agent platform wins?"

It is which control points must remain portable, which can be delegated to a strategic platform, and where proprietary context is valuable enough to justify platform gravity.

Frontier recommends keeping six things conceptually separable even when one vendor supplies several of them: agent identity, policy, tool contracts, memory/context, evaluation/telemetry, and model/runtime selection.

How to enable the autonomous enterprise

Frontier recommends a staged enablement model.

1. Start with an autonomy portfolio, not an agent catalog

Inventory candidate workflows and score them on business payoff, execution volume, process variability, data readiness, integration effort, risk, exception rate, and measurability.

Do not begin by asking where to “use agents.” Begin with work that is expensive, slow, coordination-heavy, or constrained by scarce expertise.

2. Redesign the workflow before automating it

Map the current process, but do not preserve it by default.

Remove redundant approvals, eliminate unnecessary handoffs, expose hidden decision rules, define authoritative data, and specify exception paths. BCG’s 2026 work shows that end-to-end redesign is the dividing line between incremental task gains and much larger operating improvements.[13]

3. Define an autonomy envelope

For every production agent or agent team, specify:

- the outcome it owns;
- events that may trigger it;
- systems and data it may access;
- actions it may take;
- spending or transaction limits;
- decisions requiring approval;
- prohibited actions;
- escalation conditions;
- service-level objectives;
- evaluation thresholds;
- accountable human owner.

This is the executable contract of enterprise autonomy.

4. Make agent identity a first-class architecture domain

Every agent should be discoverable and attributable.

Create lifecycle management for registration, ownership, credentials, permissions, policy, monitoring, and retirement. Use delegated or workload identities instead of shared credentials. Apply least privilege and short-lived authority wherever the platform permits it.

NIST's guidance makes this foundational rather than optional.[14]

5. Build a shared orchestration layer

Avoid hard-wiring every agent directly to every system.

Use a governed mediation layer for tools, workflows, rules, approvals, memory, exception routing, and inter-agent coordination. This is the point where heterogeneous agents can participate in consistent enterprise controls.

6. Treat enterprise context as production infrastructure

Create trusted retrieval, semantic definitions, knowledge services, metadata, and business context that agents can consume consistently.

Agentic systems magnify ambiguity. If three systems disagree about a customer, policy, product, or metric, autonomy makes the inconsistency operational.

7. Instrument everything

Trace prompts, plans, tool calls, policy decisions, transactions, exceptions, latency, model usage, costs, and business outcomes.

An enterprise should be able to reconstruct why a significant agent action occurred, what context was used, what policy allowed it, and what happened afterward.

8. Evaluate completed work continuously

Create test suites and production evaluations at three levels:

1. task quality — did the agent perform the immediate task correctly?
2. workflow quality — did the end-to-end process reach the desired outcome?
3. enterprise quality — was the outcome compliant, economical, secure, and aligned with policy?

The third level is where enterprise autonomy succeeds or fails.

9. Make AI economics an operating discipline

Track cost per completed outcome, exception cost, human review time, model/tool spend, error cost, and realized business value.

Re-evaluate routing and workflow design regularly because the economics are changing rapidly. McKinsey argues that agentic unit economics should become part of recurring business reviews rather than a one-time business case.[11]

10. Create a human-agent operating model

Roles will change before org charts do.

Employees will increasingly supervise, direct, correct, and collaborate with agents. Managers will allocate work across humans and digital workers. Process owners will become autonomy owners. Security, architecture, data, finance, legal, and business operations will share responsibility for runtime policy.

This requires training, incentives, job redesign, and clear accountability—not merely tool adoption.

11. Design for multi-vendor autonomy

Assume that important business processes will include agents from application vendors, cloud platforms, model providers, automation suites, and internal teams.

Standardize identity, tool access, observability, evaluation, metadata, policy, and agent-to-agent interaction where possible. Avoid making the control plane inseparable from a single model.

12. Expand autonomy only when evidence supports it

Use progressive delegation.

Start with recommendation, move to human-approved action, then to bounded autonomous action, and finally to wider exception-based oversight when reliability and economics support it.

Autonomy should be earned empirically.

The CIO mandate changes

For decades, CIOs were responsible for systems that stored information, ran transactions, and enabled people to perform work.

The autonomous enterprise adds a new responsibility: systems that perform work themselves.

That turns several familiar disciplines into something larger.

Enterprise architecture becomes the architecture of human and machine agency.

Identity management expands from people and applications to digital workers.

Integration becomes tool access and action mediation.

Data governance becomes decision-context governance.

IT operations become the operations of probabilistic workers.

FinOps expands from cloud resources to machine cognition.

Cybersecurity moves from protecting applications to constraining autonomous actors.

Portfolio management expands from applications and projects to agents and delegated business capabilities.

And governance moves from periodic review to runtime control.

This does not reduce the importance of the CIO. It pulls technology leadership closer to the operating model of the company.

The most advanced organizations will eventually stop treating “AI agents” as a separate technology category. They will become one execution resource among several: people, deterministic software, robots, APIs, and autonomous digital workers.

The enduring architectural question will be simple:

Who—or what—should perform this work, with what authority, under which controls, at what economics?

Frontier outlook

The autonomous enterprise is not a 2030 concept. Important pieces are production-ready now, but autonomy will arrive process by process rather than company by company.

The organizations that move fastest will be those that can repeatedly find economically attractive work, redesign it, expose reliable context, assign bounded authority, observe execution, measure completed outcomes, and expand delegation as evidence improves.

Each successful workflow produces reusable agents, tools, policies, context, integrations, evaluations, and operating knowledge. Those assets lower the cost and risk of the next autonomous workflow. In that sense, governance maturity expands the Safe Delegation Frontier, and the expanded frontier allows more valuable work to move safely into machine execution.

The right metric is not agent count. It is how much valuable work can be delegated safely, economically, and accountably.

The autonomous enterprise is the accumulation of thousands of well-governed delegations — each one extending the frontier of what the organization can trust machines to do.

SOURCES / CITATIONS

Trace the argument.

Every source below is linked. Source numbering follows first appearance in the report.

01 **Microsoft — Introducing the new Copilot with Home, Code and Autopilot**
https://blogs.microsoft.com/blog/2026/09/25/introducing-the-new-copilot-with-home-code-and-autopilot/

02 **Meta — Introducing Muse**
https://about.fb.com/news/2026/09/introducing-muse-personal-ai-agent/

03 **Meta AI Research — How We Built Safety Into Muse**
https://research.meta.ai/blog/security-and-safety-for-ai-agents-our-approach-with-muse

04 **SpaceXAI — Team Bots**
https://x.ai/news/team-bots

05 **Oracle — Introducing Fusion Claw**
https://blogs.oracle.com/cx/introducing-fusion-claw

06 **OpenAI — Meet your dot**
https://help.openai.com/en/articles/6825453-chatgpt-release-notes

07 **OpenClaw — OpenClaw Enterprise**
https://openclaw.ai/blog/openclaw-enterprise

08 **McKinsey & Company — The state of AI in 2026: On the road to ROI**
https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai

09 **Gartner — Mastering Agentic AI: Multimillion-Dollar ROI Lessons**
https://www.gartner.com/en/articles/agentic-ai-roi

10 **Deloitte — AI Agents are Only the Beginning: The AI Readiness Gap**
https://www.deloitte.com/us/en/about/press-room/deloitte-survey-examines-ai-readiness-agentic-ai-success.html

11 **McKinsey & Company — Where AI agents pay off: A practical guide to the economics of agentic workflows**
https://www.mckinsey.com/capabilities/quantumblack/our-insights/where-ai-agents-pay-off-a-practical-guide-to-the-economics-of-agentic-workflows

12 **McKinsey & Company — Cutting the coordination tax: How agentic AI can reshape workflows**
https://www.mckinsey.com/industries/industrials/our-insights/cutting-the-coordination-tax-how-agentic-ai-can-reshape-workflows

13 **BCG — AI-First Enterprise Operations: Reinventing the Operating System of Work**
https://www.bcg.com/publications/2026/reinventing-the-operating-system-of-work-with-ai

14 **NIST — Back to the Future: Why Agentic AI Needs a Strong Identity Foundation**
https://www.nist.gov/blogs/cybersecurity-insights/back-future-why-agentic-ai-needs-strong-identity-foundation

15 **Microsoft — Agent 365: The Control Plane for Agents**
https://www.microsoft.com/en-us/microsoft-agent-365

16 **AWS — Amazon Bedrock AgentCore FAQs**
https://aws.amazon.com/bedrock/agentcore/faqs/

17 **ServiceNow — AI Control Tower expansion for enterprise AI governance**

https://newsroom.servicenow.com/press-releases/details/2026/ServiceNow-expands-AI-Control-Tower-to-discover-observe-govern-secure-and-measure-AI-deployed-across-any-system-in-the-enterprise/

18 **OpenClaw — Microsoft Autopilot is built on OpenClaw**

https://openclaw.ai/blog/microsoft-autopilot-openclaw

19 **Meta — Launching Meta Enterprise Platform**

https://about.fb.com/news/2026/09/launching-meta-enterprise-platform/

20 **Salesforce — Introducing the Trusted Enterprise AI Harness**

https://www.salesforce.com/ap/news/press-releases/2026/09/11/salesforce-introduces-the-trusted-enterprise-ai-harness/

21 **SAP — Autonomous Enterprise**

https://www.sap.com/products/autonomous-enterprise.html

22 **Workday — Workday and Google Cloud expand partnership around Agent System of Record**

https://newsroom.workday.com/2026-05-28-Workday-and-Google-Cloud-Expand-Strategic-Partnership-to-Bring-AI-Agents-for-HR-and-Finance-Into-Employees-Daily-Workflows

23 **OpenAI — Introducing the Agents API**

https://openai.com/index/introducing-the-agents-api/

24 **OpenAI — Frontier: Enterprise platform for AI agents**

https://openai.com/business/frontier/

25 **Google Cloud — Gemini Enterprise Agent Platform overview**

https://docs.cloud.google.com/gemini-enterprise-agent-platform/overview

26 **Oracle — Introducing Fusion Agentic Applications**

https://www.oracle.com/news/announcement/oracle-introduces-fusion-agentic-applications-2026-03-24/

27 **UiPath — Maestro September 2026 release notes**

https://docs.uipath.com/maestro/automation-cloud/latest/release-notes/september-2026

28 **IBM — Agentic Control Plane in watsonx Orchestrate**

https://www.ibm.com/new/announcements/introducing-the-agentic-control-plane

29 **Anthropic — Introducing Claude Sonnet 5**

https://www.anthropic.com/news/claude-sonnet-5

FRONTIER.

TAKE THE NEXT STEP

Push this further.

Research is useful when it changes a decision. Bring us the question, the constraint, or the argument you want tested.

RESEARCH / BRIEFINGS

inquiries@thisisfrontier.ai

Briefings, advisory work, executive discussions, and research inquiries.

REPRINTS / LICENSING

reprints@thisisfrontier.ai

Reprints, enterprise distribution, excerpts, and licensing.

FRONTIER

www.thisisfrontier.ai

Research, Data Insights, News, and the live Frontier corpus.

All rights reserved. Contact Frontier for redistribution or reprint rights.

SEE FURTHER. MOVE SOONER.