Deloitte.

Agentic enterprise 2028

A blueprint for cost savings, job creation, and faster growth through agentic AI

September 2025

About the Deloitte AI Institute

The Deloitte AI Institute™ helps organizations connect the different dimensions of a robust, highly dynamic and rapidly evolving AI ecosystem. The Institute leads conversations on applied AI innovation across industries and offers cutting-edge insights, all to promote human-machine collaboration in the "Age of With."

The Deloitte AI Institute aims to promote a dialogue and development of artificial intelligence, stimulate innovation, and examine both challenges to AI implementation and ways to address them. The Institute collaborates with an ecosystem composed of academic research groups, startups, entrepreneurs, innovators, mature AI product leaders, and AI visionaries to explore key areas of artificial intelligence including risks, policies, ethics, future of work and talent, and applied AI use cases. Combined with Deloitte's deep knowledge and experience in artificial intelligence applications, the Institute helps make sense of this complex ecosystem and, as a result, delivers impactful perspectives to help organizations succeed by making informed AI decisions.

No matter what stage of the AI journey you're in, whether you're a board member or a C-suite leader driving strategy for your organization or a hands-on data scientist bringing an AI strategy to life, the Institute can help you learn more about how organizations across the world are leveraging AI for a competitive advantage. Visit us at the Deloitte AI Institute to access the full body of our work, subscribe to our podcasts and newsletter, and join us at our meetups and live events. Let's explore the future of AI together.

www.deloitte.com/us/AIInstitute

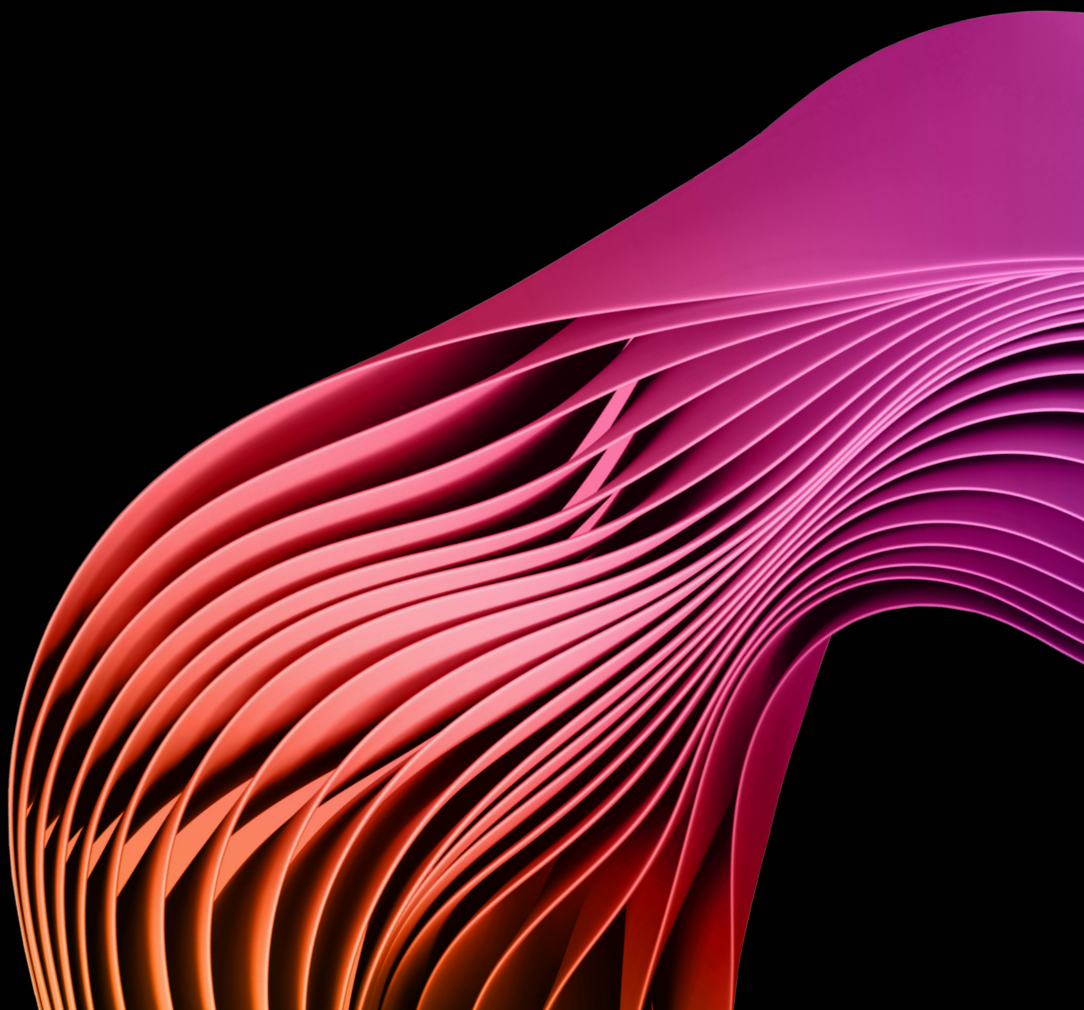

Content

Key takeaways

- Agentic AI isn't just smarter automation—it provides a strategic blueprint for enterprises to achieve cost efficiency, drive revenue growth, and unlock the full potential of their talent.
- While autonomy delivers major benefits, its complexity demands a pragmatic, coordinated evolution of strategy, technology, data, workforce, governance, and change management.
- Human roles will evolve significantly—from monitoring and feedback at early stages to orchestration and strategic oversight as agentic integration advances—demanding dynamic talent pipelines and careful work redesign.
- Enterprises that act early and approach autonomy as a phased transformation will secure lasting advantage and shape industry standards.

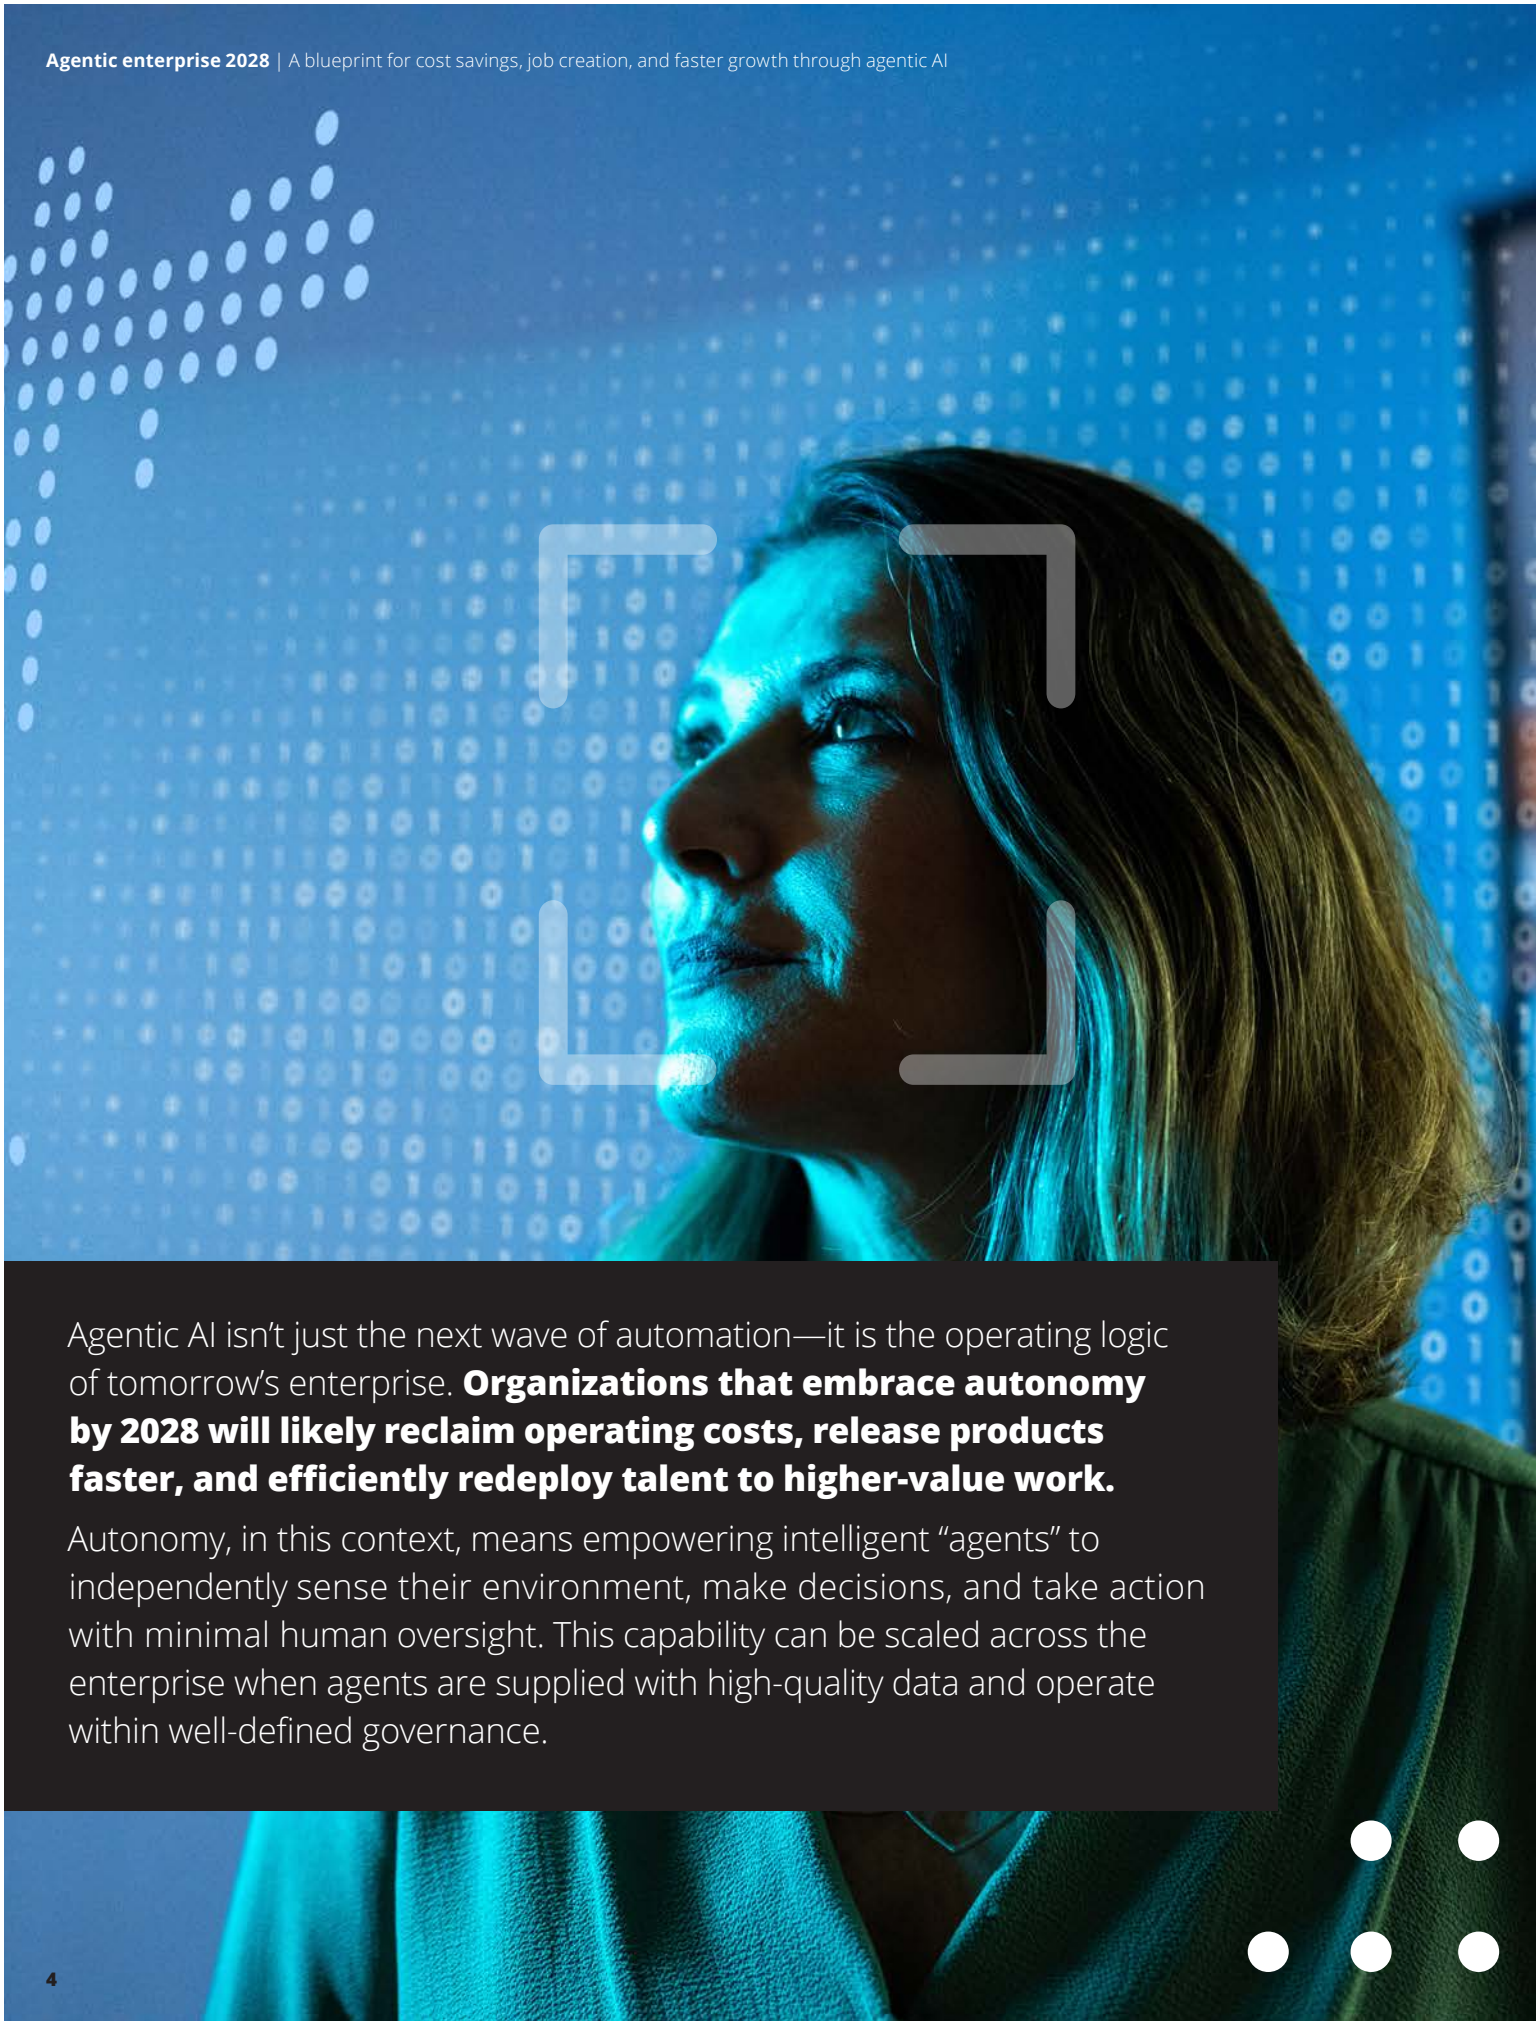

Agentic AI isn't just the next wave of automation—it is the operating logic of tomorrow's enterprise. **Organizations that embrace autonomy by 2028 will likely reclaim operating costs, release products faster, and efficiently redeploy talent to higher-value work.**

Autonomy, in this context, means empowering intelligent "agents" to independently sense their environment, make decisions, and take action with minimal human oversight. This capability can be scaled across the enterprise when agents are supplied with high-quality data and operate within well-defined governance.

Executive summary

In the current landscape, mastering autonomy is critical for leaders aiming to drive resilience and growth. Boards could face a convergence of **margin pressure, talent scarcity, supply chain fragility, and fast-evolving regulations**. Agentic AI—systems that can **plan, act, and learn** on behalf of humans—offers a staged path to enterprise-wide value and unlocks benefits such as:

- **Cost takeout** through accelerated processes and near-zero rework, driving greater operational efficiency.
- **Rapid decision-making** based on real-time agentic monitoring and analysis—adaptive pricing or supply chain moves can happen in minutes instead of days.
- **Revenue uplift** created from 24/7 customer engagement, personalized marketing journeys, and new products aligned with identified consumer trends.
- **Talent leverage and net-new roles created** for orchestration, reliability, and strategic oversight as employees can concentrate on higher-value, innovative work.
- **Faster time-to-market** as autonomous mesh architectures compress product release cycles and enable rapid adaptation to market demands.

But these benefits are matched by complexity. Realizing them depends on climbing a **six-step autonomy ladder** that is grounded on replatforming technology, reskilling talent, revising risk management practices, and redesigning work at scale. Accurately measuring success during the climb is critical to sustaining the upward trajectory. Deloitte's **return-on-autonomy (RoA) dashboard** provides leaders with industry-agnostic, key performance indicators (KPIs) to gauge impact, calibrate investment, and course-correct early to maximize gains.

A pragmatic approach to agentic AI with careful planning and phased implementation helps organizations avoid hype-driven missteps and achieve long-term value. **This paper draws on fresh market evidence, cross-industry use cases, and specialist frameworks to provide executives with a comprehensive blueprint for building and scaling an agentic enterprise.**

The **benefits** of agentic AI are matched by **complexity**.

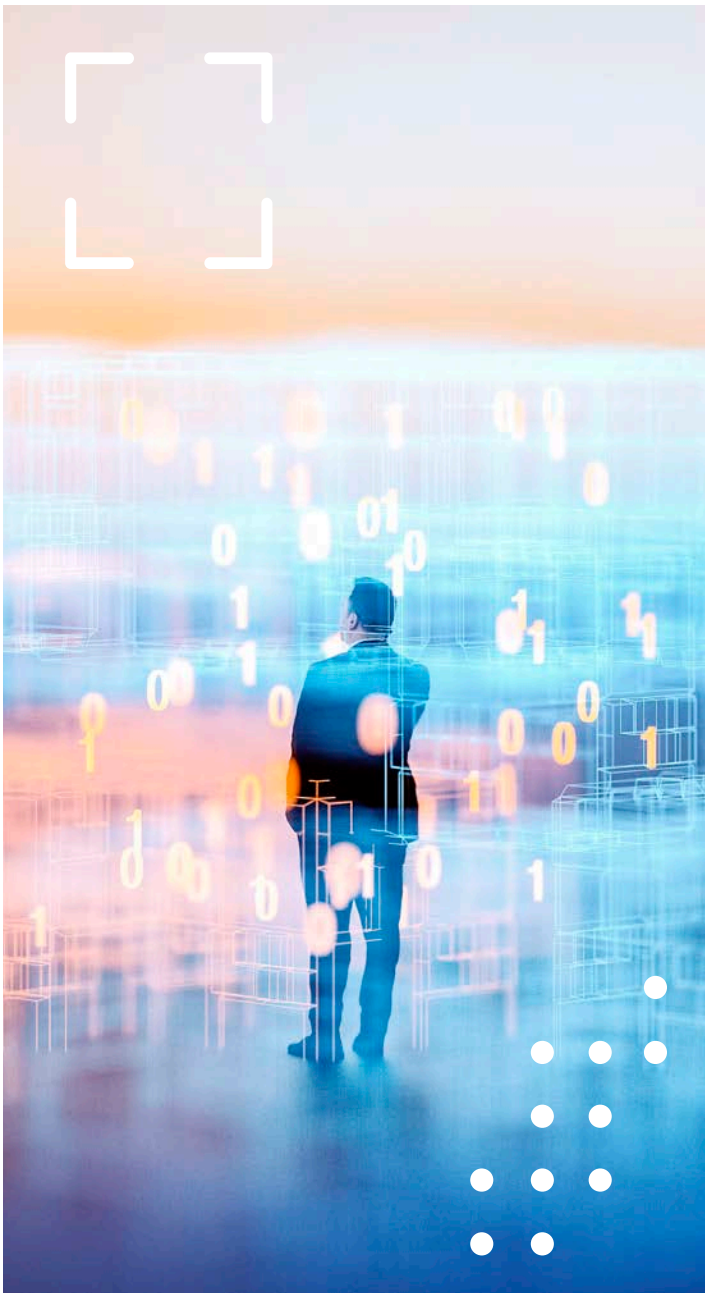

Why now?
The imperative for autonomy

The strategic clock is ticking faster than ever, as **three macro forces compress the time frame for making critical organizational decisions** and moves related to AI.

1 Competitive pressure

- 77% of CEOs say AI will "shape the future of business," yet two-thirds concede their **current business model isn't ready**.[1]

2 Regulatory scrutiny

- The EU AI Act and recent US executive orders codify obligations for high-risk AI systems used in areas like government, health care, HR, and finance, making **proactive governance and compliance a competitive advantage**.
- Agentic AI turns compliance from cost center to differentiator—firms that bake in auditability and reporting can **fast-track market entry and secure lower cost of capital**.

3 Tech momentum and evolution

- Gartner forecasts one-third of enterprise applications will embed autonomous agents by 2028.[2] The maturity of agentic capabilities will continue to **unlock new use cases and drive value** for those organizations willing to invest.

Here's a comparative look at the anticipated capabilities

While 2025 sees AI agents capable of handling a variety of structured tasks and acting as sophisticated assistants, by 2028 they are projected to become autonomous partners that tackle complex, multistep problems and proactively shape decision-making.

Capability	2025	2028 (projected)
Autonomy and decision-making	Primarily human-in-the-loop, rule-bound choices	Significantly more autonomous and capable of making proactive decisions; human role shifts to reviewer and monitoring agents emerge
Reasoning and planning	Can handle structured problems with short-term planning	Strategic, real-time adaptive planning in ambiguous situations
Natural language understanding	Context-aware chat and commands	Near-human level of understanding across nuances and multiple languages
Task complexity and execution	Automates repetitive, rule-based tasks	Orchestrates complex, multistep workflows with dependencies based on a defined goal or outcome
Multi-agent systems	Early, small-scale collaborations within structure frameworks (e.g., LangGraph)	Widespread adoption of sophisticated multi-agent systems with specialized skills, mirroring a human team
Integration and adoption	Point solutions (e.g., customer service or IT help desk); tech-heavy rollout	Enterprise-wide applications driven by low/no-code integration capabilities
Learning and adaptation	Learns from explicit feedback	Self-improving via reinforcement and human observation

The autonomy ladder
Understanding levels of agentic AI

Agentic capabilities span a spectrum of autonomy, from basic task automation to highly independent agents capable of complex, end-to-end processes. This progression offers opportunities to create new efficiencies, drive innovation, and transform business operations. **To realize the full potential of agentic AI, business leaders should proactively prepare their organizations to navigate and scale across these different levels of autonomy.**

Our Autonomy Ladder aligns to six maturity plateaus

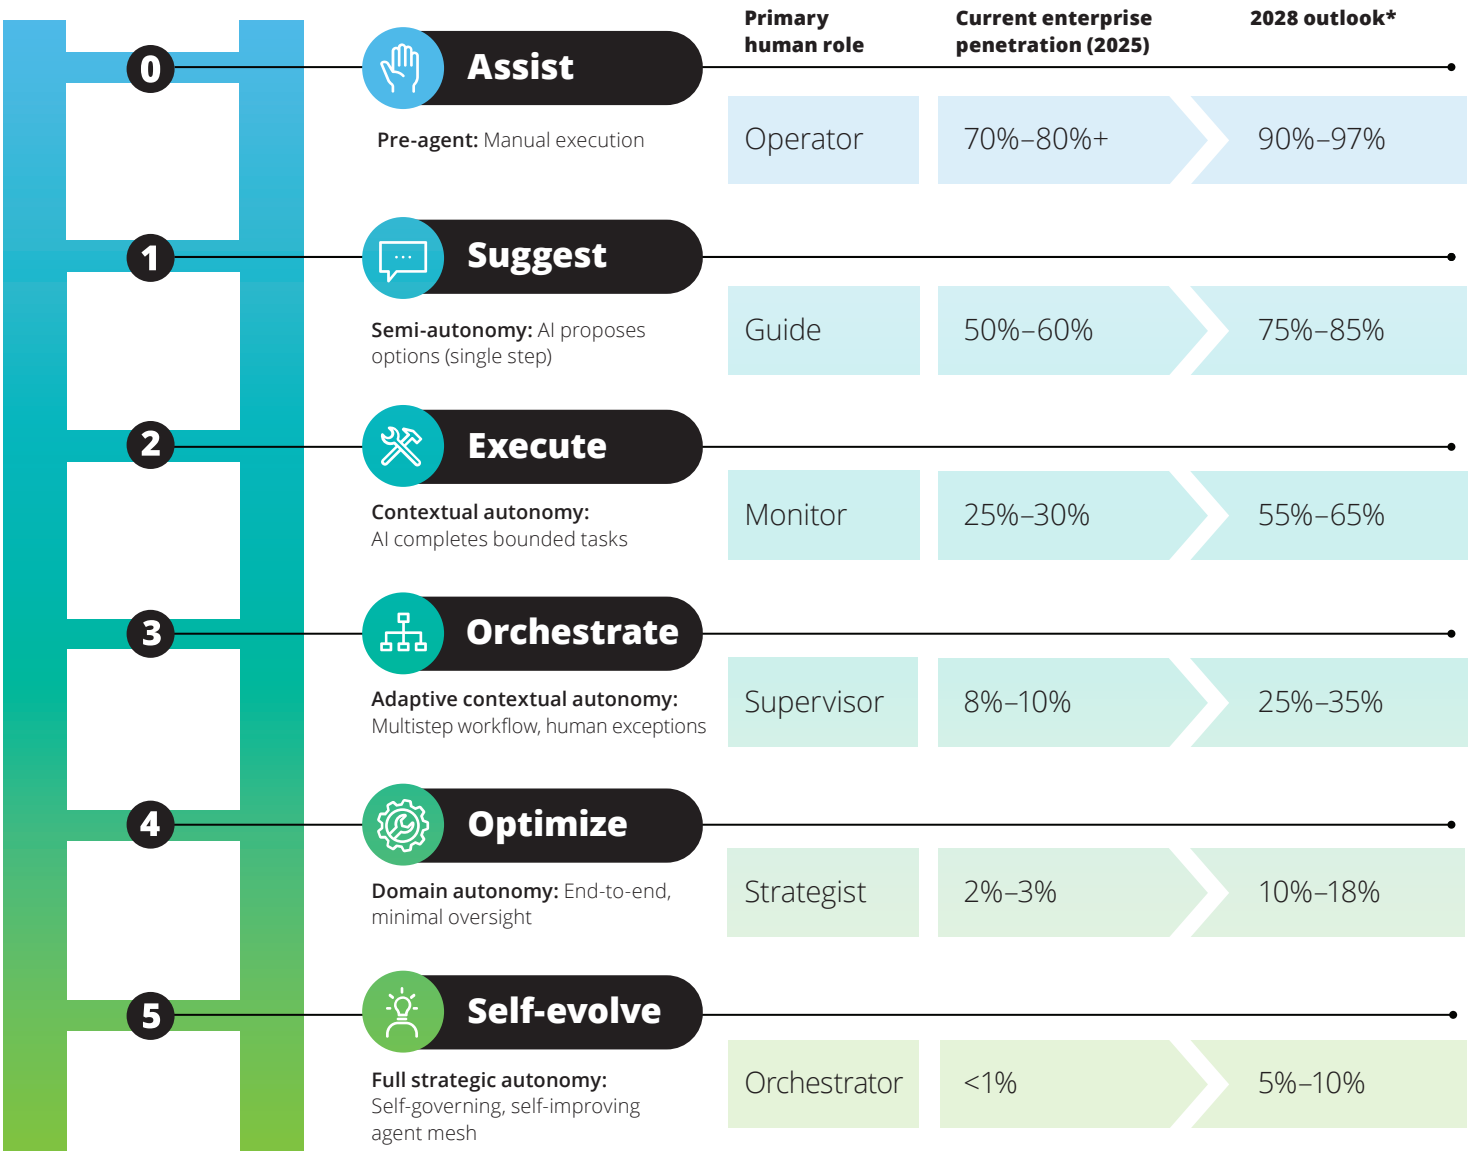

- Sample frame = global Fortune-2000-type organizations
- "Penetration" = at least one production-grade deployment at scale (not pilot/proof-of-concept)

*Adoption bands reflect sector intensity, regulatory clarity, data fabric maturity, compute elasticity, and change-management capacity. 2028 outlook is scenario-based and illustrative, not predictive.

USE CASES

Climbing the autonomy ladder across domains

LEVEL	CUSTOMER/ CUSTOMER EXPERIENCE	OPERATIONS & SUPPLY CHAIN	FINANCE & RISK	HUMAN RESOURCES & TALENT	INFORMATION TECHNOLOGY & CYBERSECURITY
0 Assist	Frequently asked questions (FAQ) returns chatbot	Optical character recognition (OCR) invoice capture	General ledger line item summarizer	Job description drafter	Log file text explainer
1 Suggest	Cross-sell prompt during live chat	Low-stock automatic reorder alert	Cash-flow forecast "what-if" scenarios	Offer letter compensation suggestions	Software patch priority ranking
2 Execute	End-to-end refund processing	Robotic process automation plus Generative AI to create advanced shipping notices	Accounts payable three-way match and posting	Automated benefits-enrollment packet generation	Automated software test-case generation
3 Orchestrate	Proactive outreach to customers predicted to cancel	Order-to-cash exception handling and smart rerouting	Continuous "know your customer" remediation workflow	Zero-touch employee onboarding (accounts, equipment, mandatory training)	Security incident triage followed by containment playbook
4 Optimize	Full-funnel marketing campaign design and launch with drift checks	Autonomous demand planning and inventory balancing	Month-end financial close with automated anomaly remediation	Workforce capacity and schedule optimization	Self-healing patch deployment and scaling engine
5 Self-evolve	Real-time catalog and offer evolution by a mesh of collaborating agents	Self-negotiating procurement and supply chain network	Autonomous treasury management and hedging strategies	Autonomous high-performer flight-risk identification to prevent turnover and promote retention	Adaptive adversarial testing with continuously evolving defenses

Architecting the autonomous operating system

How can leaders **successfully build organizations** fueled by autonomous systems?

As organizations shift from traditional automation toward deploying intelligent agents that can act with increasing autonomy, they may face a new class of opportunities and challenges. Agentic systems, powered by large language models (LLMs) and intelligent decision engines, offer the potential to transform how work is performed across domains. To scale responsibly and effectively, enterprises should evolve their strategy, technology, data, workforce, governance, and change approach concurrently. **This paper outlines a comprehensive framework with six core dimensions for increasing agentic autonomy, accompanied by actionable recommendations for each.**

Six pillars for architecting the autonomous operating system

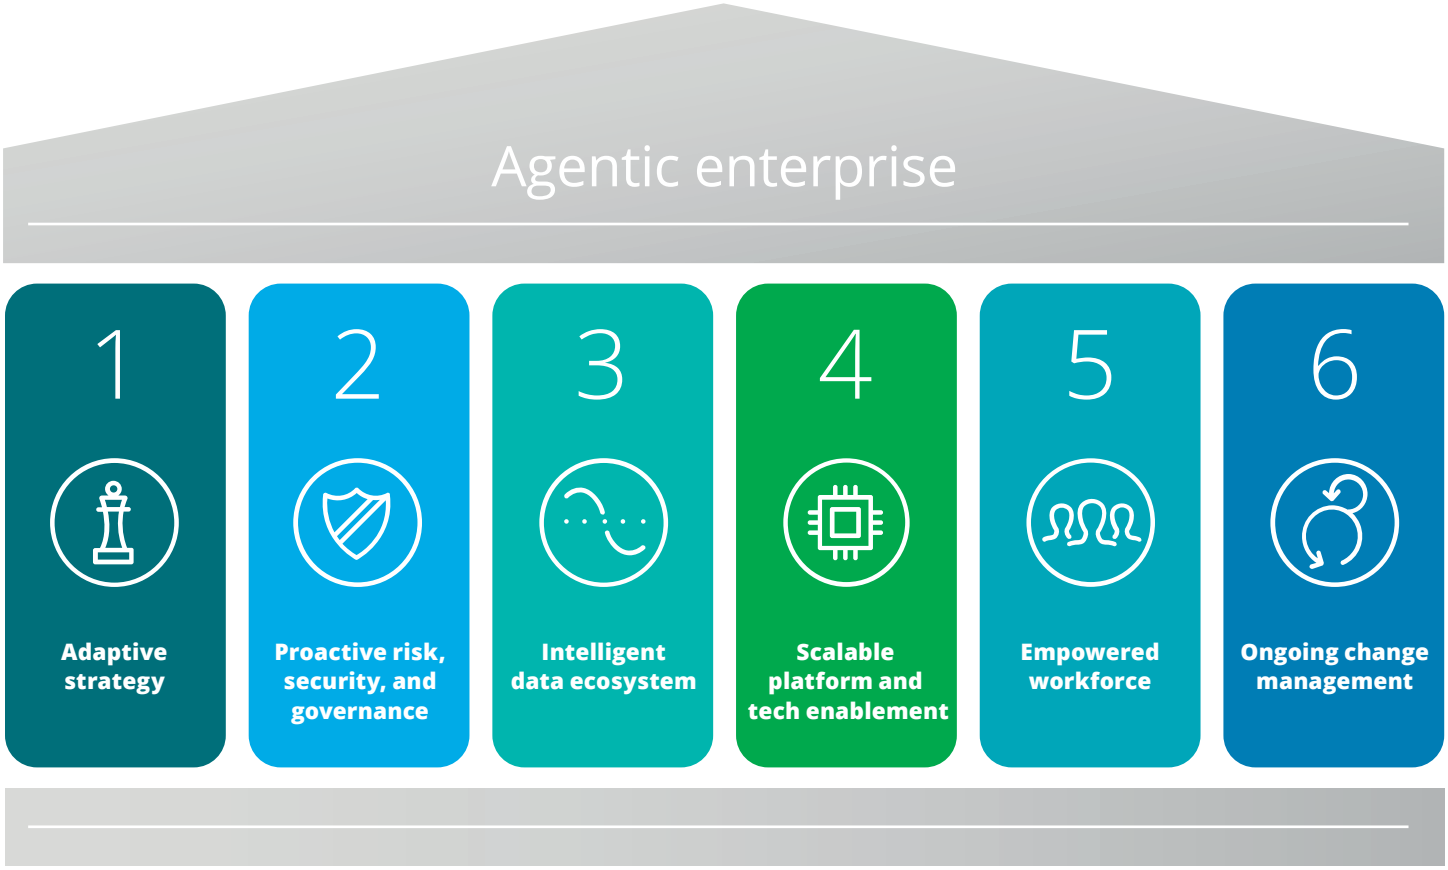

1 Adaptive strategy

Aligning autonomy with business value

Strategic planning lays the foundation for agentic transformation. Organizations often face important choices about where and how to integrate agentic capabilities into their operations. Successfully scaling autonomy requires a pragmatic shift from a use case "build-on" strategy to embedding agentic systems in the very foundation of an organization. Deploying agents alone will not fix long-standing process inefficiencies; true transformation needs both technology and systemic change. Organizations must commit to being an AI-first enterprise—entirely reimagining the organization across people, processes, and technology—to realize the full benefits of autonomy. A key step is making business processes agent-ready: structured and data-rich so autonomy adds value, not inefficiency.

In practice, agentic integration typically follows three main patterns:

Agents "overlay": Agents are added as enhancements to existing workflows, tools, and infrastructure, providing incremental automation and intelligence without fundamentally changing underlying workflows—for example, by deploying bots to handle repetitive tasks within established business operations.

Agents "as-a-service": Agentic capabilities are increasingly available as default features from leading enterprise platform providers, enabling organizations to adopt and scale agents rapidly. Pre-built, cloud-based solutions offer embedded agentic functionalities—such as AI-driven analytics, workflow automation, or virtual assistants—directly within their existing SaaS platforms (e.g., CRM, ERP, HR systems).

Agents "by design": Organizations move beyond incremental enhancements or pre-packaged solutions and instead architect or re-architect their systems and workflows from the ground up via autonomous agents. This may involve redesigning entire business processes as networks of collaborating agents, each responsible for specific tasks, decision-making, or optimization. While this approach requires significant investment in engineering, change management, and organizational alignment, it offers the greatest potential for competitive differentiation and long-term value creation by making autonomy and intelligence intrinsic to the business's digital fabric.

Selecting the right integration approach(es) is itself a strategic decision, as it determines the pace, scale, and impact of autonomy within the enterprise. Different areas of the business may require different strategies. To maximize value, organizations should not treat integration and use case identification as separate exercises. When integration strategy, targeted use-case selection, and a clear framework for strategic traceability and success metrics are developed in tandem, organizations can drive adoption, accelerate learning, and prove value. High-potential use cases are those in which agentic autonomy can reduce operational burden, enhance responsiveness, or improve service quality in ways that directly support business objectives.

This paper will explore the key components of organizational readiness required for effective agentic integration and how these components relate to each integration pattern, enabling enterprises to chart a clear and pragmatic path toward scalable autonomy.

Key considerations

- Agentic integration approaches and readiness
- Use-case identification and prioritization
- Alignment with business objectives
- Strategic traceability and success metrics

POTENTIAL FOCUS AREAS

1. **Conduct a multidimensional readiness assessment across the enterprise:** Evaluate processes, technology, data, and workforce maturity to determine where integration patterns (overlay, as-a-service, or by-design) would best fit across the business. (Suggestions for conducting a maturity assessment are provided in the final section of this paper.)
2. **Design a layered rollout strategy:** Begin with agentic overlay pilots and as-a-service capabilities embedded in existing SaaS platforms for quick wins, and leverage the insights gained to re-architect core processes around autonomous agents.
3. **Create a dynamic, agentic use-case scoring and selection system:** Keep momentum on ideas with strong potential for scalability and interoperability to effectively deploy resources.
4. **Establish an Agentic Center of Excellence (CoE) to maintain the portfolio of agentic initiatives:** A purpose-built CoE turns strategic intent into a renewable engine of autonomous value—prioritizing only what matters and establishing success measures.

2 Proactive risk, security, and governance

Managing emerging autonomy risks

Agentic systems can only thrive if stakeholders trust them. To earn and maintain this trust, organizations should embed proper usage guidelines, transparency, and shared accountability at every stage of the AI agent life cycle. By proactively addressing uncertainty and fostering a culture of continuous learning and adaptation, organizations can enable the scalable and responsible deployment of agentic systems.

However, the increased autonomy of agents introduces a new spectrum of potential risks—including operational and regulatory concerns—against a backdrop of rapid technological evolution and marketplace uncertainty. In this dynamic environment, a robust governance framework is essential not only to anticipate and mitigate current risks, but also to enable organizations to pivot quickly as new threats and opportunities emerge. Compliance now goes beyond meeting current standards and requires anticipating future risks and designing flexible guardrails that can adapt as regulations and business needs evolve.

As organizations scale agents, maintaining human oversight can become increasingly challenging. One emerging solution is the introduction of "guardian agents" designed to monitor, validate, and manage other agents. These supervisory systems can reinforce trust and resilience by ensuring that autonomous decisions remain auditable and compliant even at scale.

Key considerations

- Agile risk assessment and management strategies
- Proper usage guidelines, safeguards, and controls
- Continuous testing and audits
- Transparency and explainability
- Policy-as-code

Illustrative agentic use-case risk levels

A structured assessment framework is essential for evaluating risk and determining the right safeguards and governance for each use case. The table below illustrates how organizations can categorize agentic system use cases by risk level across key dimensions.

Business impact	Output exposure	Data sensitivity	Autonomy level	Regulatory burden
LOW				
Non-critical workflow (e.g., research support)	Internal only	Low risk or public	Pre-defined, repeatable steps	None
MEDIUM				
Semi-critical (e.g., IT help desk)	Internal; may reach clients/customers	Moderate, potentially sensitive	Semi-autonomous with human checkpoints	Some industry and regulatory rules
HIGH				
Mission-critical (e.g., finance, HR)	Direct external use/client- or customer-facing	Confidential/regulated	Fully autonomous, multi-agent handoffs	Strict industry and regulatory oversight

POTENTIAL FOCUS AREAS

1. **Institute agent-specific risk management frameworks and monitoring:** Establish structured methodologies to evaluate the risks unique to autonomous agents, and transform risk management into an ongoing process to proactively detect and address emergent failures before they escalate.
2. **Establish cross-functional AI risk governance:** Form a committee of legal, IT, risk, compliance, and business leaders to help ensure shared accountability, co-designed policies, and balanced oversight mechanisms as autonomy increases.
3. **Embed policy-as-code and automated controls:** Implement machine-readable risk, compliance, and governance policies within agentic solutions, enabling automated enforcement, real-time guardrails, and comprehensive audit trails.
4. **Enhance transparency and human oversight:** Clearly document and communicate agent limitations, guardrails, and escalation pathways, and maintain human intervention for high-impact or uncertain scenarios.

3 Intelligent data ecosystem

Foundation for intelligent autonomy

AI agents are only as effective as the data that fuels them. To reason, predict, and act reliably, they need seamless access to accurate, complete, and context-rich information—transaction records, sensor feeds, knowledge articles, customer interactions—in addition to the nuanced, experience-based insights of human experts. When data pipelines are fragmented or lineage is opaque, agents inherit those weaknesses. A modern, trusted data foundation—built on architectures that unify operational insights, analytical data, and human expertise—is a non-negotiable prerequisite for agentic success. Even the most robust data infrastructure is incomplete without the integration of human expertise. To achieve true autonomy by 2028, AI agents must learn from the human experts who perform the work today, capturing the nuances, subtleties, and even institutional or cultural knowledge that are often lost through workforce turnover. This tacit knowledge gained over years of experience cannot be fully replicated by data alone.

To unlock the full potential of agentic systems, organizations should embrace **"convergence architectures"**—bringing together both operational and analytical data into unified, accessible platforms. This convergence ensures that AI agents can draw on real-time operational signals as well as deep analytical insights, enabling richer reasoning and more adaptive decision-making. Building these architectures requires clear business objectives, unified data governance, and the adoption of modern data platforms that support both transactional and analytical workloads. Real-time data integration through event-driven architectures and streaming technologies is essential, as is ensuring scalability, flexibility, and robust security across all data flows. Automated data lineage, monitoring, and standardized, reusable data services further strengthen the architecture, while incremental adoption helps demonstrate value and refine the approach.

Yet even the most advanced data architectures are incomplete without the transfer of human expertise. Data provides breadth, but expertise provides depth. Enabling this transfer is essential to help ensure AI agents evolve into truly organizationally intelligent systems. Organizations should design meaningful incentives for employees to collaborate with AI systems, transforming individual know-how into a lasting corporate asset. This is a critical dependency for agentic integration that should be addressed alongside data modernization efforts.

Key considerations

- Data readiness
- Data quality and richness
- Data accessibility
- Data fabric and vector context

POTENTIAL FOCUS AREAS

1. **Enhance data quality management:** Implement continuous data validation and cleansing processes.
2. **Establish a convergence architecture backbone:** Unify operational and analytical data on a single, cloud-native platform—leveraging event-driven and streaming technologies—to give agents real-time, end-to-end visibility across the enterprise.
3. **Make data accessible:** Create secure integrations between agents (e.g., MCP/A2A) to unify access across enterprise systems. MCP (Model Context Protocol) and A2A (Agent-to-Agent) architectures will be addressed in the next section on technology enablement.
4. **Stand up a vectorized data fabric for context-rich retrieval:** Agents gain low-latency, context-aware access to the *right* slices of knowledge—boosting precision, reducing bias, and speeding decision cycles.
5. **Operationalize data governance and lineage as code:** With transparent provenance and enforceable policies, you curb systemic bias, satisfy auditors, and give agents a self-serve way to pick high-trust data.

4 Scalable platform and tech enablement

Infrastructure for scale and integration

Agentic systems must operate across a highly interconnected ecosystem of tools, platforms, and databases. Agents themselves will originate from a diverse range of sources: natively built into hyperscale cloud platforms or enterprise application platforms, third-party agentic providers, and highly bespoke solutions. Organizations should be prepared to integrate agents across this broad spectrum of sources, ensuring interoperability and cohesive operation.

Technology readiness includes not just raw compute power, but also frictionless integration with enterprise workflows. Agentic systems ingest data, query ERP and CRM platforms, trigger downstream actions, and surface insights through user-facing apps. Technology readiness therefore goes far beyond GPUs; it requires a service-oriented architecture that is MCP-enabled, A2A-friendly, and MAS-powered (multi-agent systems), allowing agents to interoperate with core business applications.

Traditional API-driven (application programming interface) architectures, while effective for point-to-point integrations, often create bottlenecks and rigid dependencies that limit the flexibility and scalability needed for autonomous agentic systems. To enable scalable and sustainable autonomy by 2028, organizations should evolve from traditional API-driven architectures to modern MCP, A2A, and MAS-powered architectures—laying the foundation for agentic systems to operate autonomously and at scale.

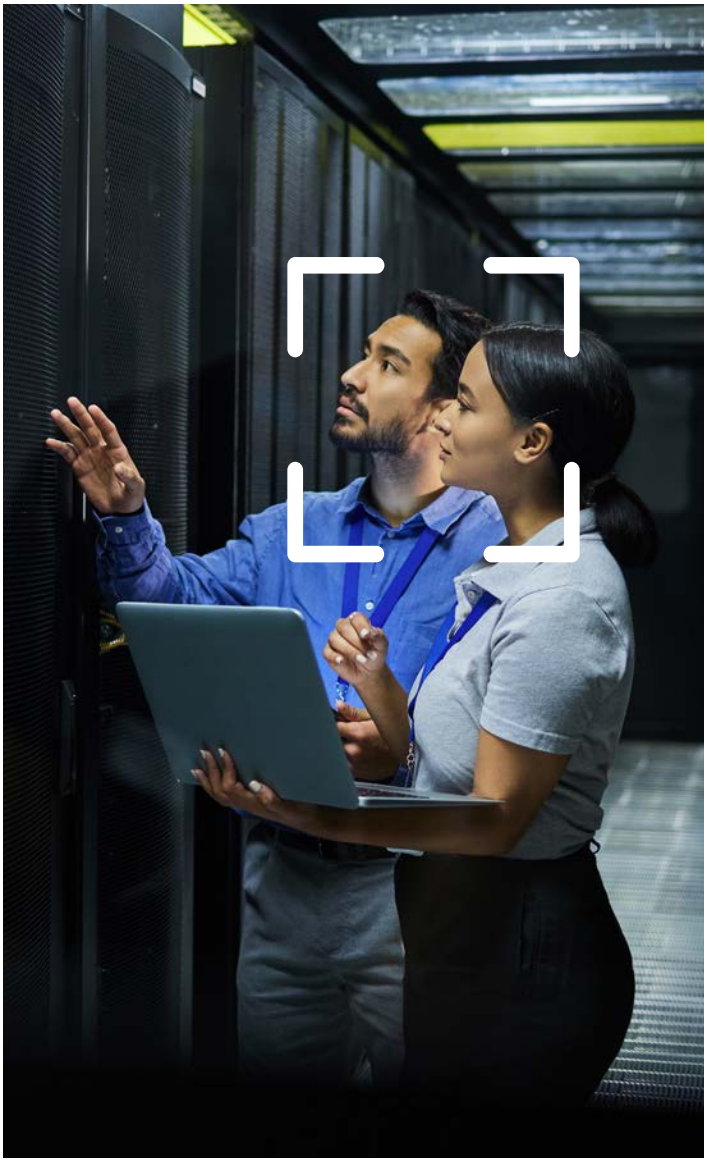

Key considerations

- Infrastructure readiness
- Compute infrastructure
- Integration capabilities
- Agent mesh and orchestration graphs

POTENTIAL FOCUS AREAS

1. **Invest in scalable infrastructure:** Ensure hardware and cloud environments can support increasing agent loads.
2. **Enable external integration:** Shift from traditional API-driven approaches to building robust MCP and A2A integrations to embed agents into business-critical applications and workflows.
3. **Deploy robust deployment management:** Anticipate and manage technical shifts brought on by agentic integration.
4. **Establish an "agent-mesh" orchestration layer:** Treat the fleet of agents like a cloud-native microservices mesh, enabling secure discovery, context exchange, and telemetry through uniform, policy-driven interfaces.

5 Empowered workforce

Designing roles with AI in mind

The human workforce will not be replaced by agents, but it will be redefined. Human roles and skills shift as agentic adoption matures—from early monitoring and feedback to advanced orchestration and strategic oversight. Dynamic talent pipelines can help organizations build the right capabilities at the right time, aligning workforce design with agentic integration. "Human-in-the-loop" is evolving into more sophisticated, deeply integrated human-machine collaboration.

The impact of agents will likely differ by work type: routine tasks may be largely automated, analytical and creative roles will be augmented by agent-driven insights and ideation, and interpersonal or safety-critical work will continue to rely heavily on human judgment, empathy, and oversight. As a result, organizations should tailor agent integration to the unique demands of each role to maximize value and help ensure effective, context-sensitive collaboration.

Key considerations • Role and work redesign • Job creation • AI fluency and skill-building programs

The agentic enterprise will likely impact different types of work differently

KNOWLEDGE WORK: ROTE/STRUCTURED
Very high automation of repetitive, rules-based tasks. Humans shift to exception triage, quality oversight, and orchestration.

KNOWLEDGE WORK: DOMAIN/COMPLEX
Selective automation of synthesis and research; agents act as co-analysts while humans lead judgment and scenario design.

TECH
High automation of setup and boilerplate coding tasks (e.g., project configuration, repetitive test cases, basic integrations). Human work shifts toward system design, orchestration, and governance.

CREATIVE
Routine creative production (copy, imagery, layouts) is automated. Value shifts to direction, brand alignment, and curation.

FRONTLINE: GENERAL
High automation of repetitive frontline tasks (checkout, inventory, routing, scheduling). Humans focus on exceptions and customer experience.

FRONTLINE: SPECIALIST
Administrative burden is reduced (documentation, checklists, readiness). Humans focus on complex judgment, safety-critical tasks, and service.

SKILLED TRADE
Limited displacement; workflows are heavily augmented with diagnostics, safety checks, and predictive maintenance.

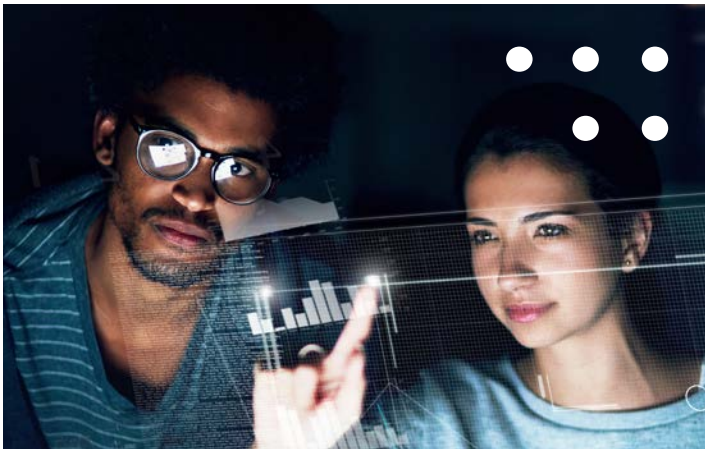

The agentic economy will fuel a boom in skilled-trade jobs

AI may live in the cloud, but its performance and carbon footprint still depend on very real hardware, facilities, and logistics networks. Global AI infrastructure spending—including compute, networking, and data centers—is projected to surpass **$200 billion by 2028**,[3] driven heavily by the expansion of GPU clusters and edge AI deployments.

The surge in GPU-dense data centers, edge devices, autonomous robots, and AI-enabled logistics is already creating a parallel surge in skilled-trade jobs. There are **five trade-oriented roles**—high-density data center electrician, immersion-cooling HVAC technician, edge-compute field engineer, green-microgrid operator, and industrial sensor and IoT installer—many of them adaptations of classic crafts. These roles will be critical to building and running the physical backbone of an "agentic" economy over the next three to five years.

Want to learn more?
Take a look at the *Future-ready workforce* infographic.

The agentic enterprise will also create new roles and job functions

As enterprises embrace agentic and autonomous systems, the nature of work is undergoing a fundamental transformation. Traditional roles centered on manual execution and routine tasks are rapidly evolving. In their place, new opportunities are emerging that focus on strategy, oversight, and orchestration—requiring a workforce that can guide, supervise, and optimize the performance of intelligent agents. The rise of agentic autonomous systems will create entirely new roles and job functions, such as agent process strategists and continuous-learning engineers, dedicated to guiding, governing, and collaborating with these digital coworkers. Agentic strategy leads will map the business landscape, pinpoint high-value use cases, and architect multi-agent workflows that align with corporate objectives and risk appetite. Agentic "supervisors" will sit on the front lines within functions, monitoring agent outputs, validating edge-case decisions, and triggering overrides when quality or safety thresholds are breached. Supporting both layers, a cohort of tech and talent specialists—combining change management, upskilling, and lightweight technical knowledge—will be essential to integrate agent platforms with legacy systems, rewrite job descriptions, and coach employees on productive human-agent collaboration. Together, these roles can help ensure that autonomy amplifies rather than undermines organizational performance and trust. The reality is that the impact of autonomous agents will likely give rise to entirely new roles that we cannot yet envision—just as positions like social media manager or influencer were unimaginable before the advent of social media.

	Strategy and business integration	Engineering and ops	Trust, risk, and accountability	People, adoption, and experience
KEY NEW ROLES	• Agentic Process Architect • AI Business Translator • Value-Stream Product Owner	• Multi-agent Systems Engineer • Prompt & Knowledge Designer • Autonomous Reliability Engineer • Continuous-Learning Steward	• Autonomy Auditor • Trust Lead • Agentic Risk & Compliance Partner • Human Escalation Officer	• Human-AI Interaction Coach • Adoption Analytics Manager • AI Personality & Experience Designer • Partner Ecosystem Lead
PRIMARY VALUE DELIVERED	Aligns autonomous workflow roadmaps to enterprise strategy, translates tech jargon into P&L impact, and prunes low-ROI use cases.	Builds and keeps the technical backbone stable—choreographs agents, designs reusable prompts, ensures "five-nines" uptime, and governs self-improving models.	Provides guardrails: forensic decision traceability, red-team testing, policy alignment, and empathetic human intervention when algorithms overstep.	Upskills staff for supervisory roles, tracks sentiment vs. productivity, curates brand-consistent agent "voice," and brokers external agent integrations.

POTENTIAL FOCUS AREAS

1. **Redesign roles and responsibilities:** Map today's org chart to future roles, redesign roles around human-machine collaboration, and develop new career ladders.
2. **Pilot new roles and job functions within an Agentic CoE:** Use trial positions to refine responsibilities of new roles (Agentic Process Strategist, Orchestration Engineer, Risk Specialist).
3. **Build fluency with training programs and create upskilling pathways:** Equip teams with the skills to understand, supervise, and collaborate with agents.
4. **Create a dynamic talent pipeline:** Forecast demand over 6-, 12-, 24-, and 36-month horizons for newly created and redesigned roles.

6 Ongoing change management

Evolving human-agent collaboration

An agentic-driven transformation demands more than traditional change management—it calls for a dynamic, continuously adaptive approach that is embedded into daily operations. **Organizations should foster a culture of experimentation, agility, and trust, where continuous learning and adaptation are the norm.** However, this requires a whole new level of behavioral change—one that is unfolding in an especially challenging environment:

- **Record-low levels of trust** are undermining adoption of AI tools.[4]
- **Employee willingness to support organizational change** has dropped sharply, from 74% in 2016 to just 44% in 2024.[5]
- **Workforce tensions are rising:** 75% of workers crave stability, while 85% of leaders say organizations must become more agile.[6]
- **Confusion around AI is widespread:** 54% of workers are concerned about the blurred lines between human and machine contributions.[7]

Successfully navigating this landscape requires organizations to go beyond the old change management playbook, embracing new strategies that address trust deficits, clarify the human-technology relationship, and support both stability and agility for the workforce.

Beyond the workforce, agentic autonomy will spill over into how communities experience your brand. As decision-making shifts to algorithms, every automated customer interaction, pricing adjustment, or supply chain reroute becomes a public touchpoint that can strengthen or erode reputation in real time. Communities will likely scrutinize whether autonomous systems amplify local economic opportunity or hollow it out, whether they reduce carbon footprints or strain resources, and whether they treat vulnerable populations fairly or encode bias at scale. In short, community benefit and brand trust should be designed into the agentic roadmap, not bolted on after deployment.

Key considerations

- Continuous change management that enables rapid experimentation
- Building trust for agentic systems
- Brand and community impacts

POTENTIAL FOCUS AREAS

1. **Deploy robust change management to support unprecedented change and challenges:** Prepare stakeholders through communication, feedback loops, and adaptive planning.
2. **Co-design guardrails with the workforce—not just for them:** Run rapid-fire "human-in-the-loop" design sprints with frontline employees and risk leads to map when an agent may act autonomously, when it must seek human confirmation, and how it overrides work.
3. **Assess brand, community, and talent impacts:** Understand how workforce and organizational changes driven by agentic AI could affect your brand, customers, communities, and future talent pipelines.
4. **Understand and action root causes for lack of workforce trust:** Given the impact trust levels can have on adoption, understand where you may face additional challenges within your organization when deploying agentic AI. Getting to the root causes of the lack of trust will help you to develop targeted solutions and interventions, and increase the likelihood for successful adoption.

Measuring impact and value
Return-on-Autonomy (RoA) KPIs

Redesigning enterprises with autonomous, **agentic AI is now the clearest way to achieve lasting advantage** in any sector.

As routine tasks migrate to agents, human talent pivots to judgment, creativity, and strategic orchestration—shrinking product cycles and amplifying differentiation. However, AI autonomy develops incrementally; it advances through discrete maturity steps. Each step delivers its own benefit, but—more importantly—lays the technical, data, and organizational groundwork for the next. The result is a compounding effect: Every new level multiplies the value of the preceding one.

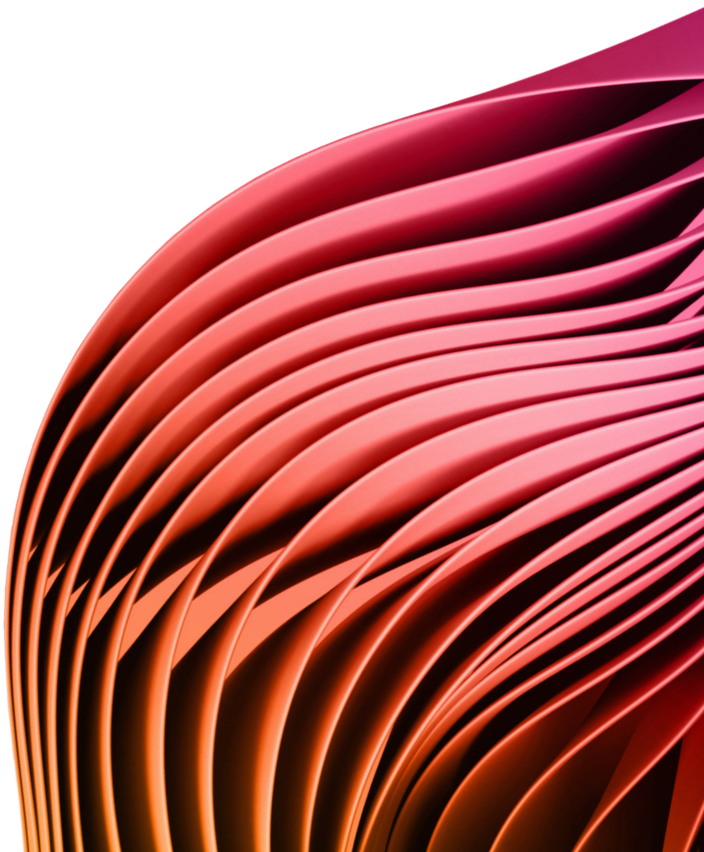

At-a-glance RoA dashboard (illustrative KPIs and targets)

Agentic AI programs can scale rapidly, but value creation is not automatic. Clear, outcome-oriented KPIs are crucial to 1) prove the business case and 2) decide when to accelerate, pause, or pivot on agentic solutions. The table below provides a balanced scorecard—cost, speed, productivity, quality, and trust—that links technical progress via increased agentic autonomy to enterprise results. The scorecard is structured to capture impacts both **at the process level** (how individual workflows or business units are affected) and **at the portfolio level** (how the organization's overall mix of processes, products, or services is transformed).

How to use these KPIs

1. **Establish a 2025** baseline that captures today's reality—before autonomy meaningfully expands.
2. **Set an "aspiration line" for 2028** that is aggressive but evidence-based; treat it as a living target refreshed every planning cycle.
3. **Instrument each metric with automated data feeds** where possible.
4. **Review performance quarterly** so leaders can reallocate funding or recalibrate ambitions.

Cost	Speed	Productivity	Quality	Trust
		Illustrative process/ workflow-level KPIS		
Unit cost/transaction *% reduction in cost per transaction*	**Median process cycle time** *hours/days per process*	**Autonomy rate** *% of tasks completed end to end (excludes human confirmation clicks)*	**First-pass yield** *% of tasks completed with no or minimal rework needed*	**In-workflow satisfaction** *user rating*
		Illustrative portfolio/ enterprise-level KPIS		
Total operating cost *% change from baseline*	**Speed-to-market** *average weeks to deployment*	**Output per FTE** *index at 100*	**Enterprise defect/ audit-finding rate** *% reduction in findings year over year*	**Agent reliability** *% of workflows where agent delivered the expected output without escalation*

Pragmatic implementation roadmap

The **future of automation** belongs to AI agents—but this shift will not happen overnight.

Successfully embedding autonomous systems within any enterprise requires careful consideration of a wide range of factors—including current technology maturity, existing processes, and the organization's readiness for adoption. Leaders must recognize that each enterprise is unique; the path to autonomy will depend on where teams stand today in terms of infrastructure, skills, and openness to change.

These suggested maturity assessment criteria can assist leaders in asking the right questions to determine which integration strategies—agents overlay, as-a-service, or by design—are the best starting point (or next step) on their autonomy journey.

If the majority fall under **quick win**, launch agents overlay/as-a-service. If most fall under **transformation**, begin planning for agents by design and deeper process redesign.

Consideration	Agentic quick wins (agents overlay or as-a-service)	Agentic transformations (agents by design)
Process clarity	• Clear, standardized • Partially defined (may be viable but agentic by design recommended)	• All processes, well-defined or ambiguous
System and data readiness	• Disconnected systems, early-stage data • Moderately connected, developing governance	• Highly connected, robust data fabric
Compute elasticity	• Rigid/fixed infrastructure • Somewhat elastic, cloud-enabled (may begin incremental agentic by design)	• Fully elastic, cloud-native
Risk and regulatory	• Low–moderate stakes • Minimal–partial compliance automation	• Mission-critical, highly regulated • Policy-as-code required
Talent readiness and work redesign	• Early-stage or developing AI talent • Incremental role shifts, focus on reskilling and redistributing tasks	• Robust, experienced AI and AgentOps teams • Work redesign is central, full role and process redesign with agents as core collaborators
Cost to integrate vs. replace	• Cheaper to integrate legacy or costs balanced	• Cheaper to replace legacy; redesign cost-justified

The following steps provide a glimpse of a **high-level roadmap** for an enterprise on the autonomy journey; however, the specific details and priorities should be tailored through further discussion and assessment to ensure alignment with organizational context and goals.

The bottom line is that leaders should act now to set the foundation for success in embedding autonomous systems within their organizations.

0 Assist

Level 0 → 1

First 30 days

Assist to suggest

AI moves from tools that support human operators, like chatbots and text summaries, to AI agents that suggest options for tasks.

Primary value unlock: AI proposes single-step options

Indicative enterprise impact*: Noticeable reduction in handoffs, intake latency, and data-entry errors

1 Suggest

Level 1 → 2

First 90 days

Agentic averlay and as-a-service pilots

AI pilots address more complex, end-to-end processes with human oversight and quality data supporting the effort.

Primary value unlock: AI executes bounded tasks end to end

Indicative enterprise impact*: Material OpEx relief through labor hours removed and straight-through processing

2 Execute

Level 2 → 3

Next 180 days

Hybrid expansion to agentic-by-design flows

Agent autonomy expands and the first agent "by design" process flows are engineered with supervisor agents managing high-impact work and humans managing exceptions. Humans begin upskilling.

Primary value unlock: Multistep workflows with human-in-the-loop exceptions

Indicative enterprise impact*: Step-change in cycle time and quality control across functions

3 Orchestrate

Level 3 → 4

By month 12

Agentic-by-design pilots

AI agents operate in teams across functional areas with automated testing for resilience. A core process is redesigned with agents. New human roles emerge to monitor and guide success, while new "guardian agents" supervise other agents.

Primary value unlock: Closed-loop optimization of entire value chains

Indicative enterprise impact*: Meaningful margin lift and incremental revenue from dynamic pricing/asset-yield gains

4 Optimize

Level 4 +

Years 2–3

Enterprise-wide agentic network

Autonomous systems graduate from supervised pilots to self-managed feedback loops. Company uses insights to launch new products or services. Humans embrace continuous learning and oversight.

Primary value unlock: Self-improving agent mesh

Indicative enterprise impact*: Transformational productivity shift; unlocks new data-driven offerings and roles

* Median, cross-industry observations; adjust for sector intensity and starting baseline.

Key takeaway

Treat autonomy as a staged journey, not a plug-in feature. Each phase builds unique assets (data, talent, guardrails) that prepare the enterprise for the next leap, ensuring that autonomy delivers durable advantage rather than siloed efficiency gains.

Conclusion
Act early, scale confidently

Agentic AI isn't a singular rollout; it's a phased transformation, much like the shift from steam to electricity on the factory floor. As companies ascend the autonomy ladder, they should strategically upgrade their technology—building modular platforms, real-time data fabrics, and guardrail orchestration—while elevating talent into new roles within a culture of experimentation and continuous upskilling. Organizations that start this refit now will likely do more than unlock incremental efficiency; they will likely create adaptive operating models that learn, self-optimize, and compound value over time. By demonstrating safe, scalable models early, these leaders will shape regulations, define industry standards, and create barriers that challenge later entrants.

The message is clear: Consider agentic-driven autonomy as a journey with well-defined stages, measurable milestones, and cross-functional ownership. Those who move first—methodically but decisively—are better positioned to capture outsized gains and help define the playbook for the agentic enterprise era.

Ready to get started?
Here are a few strategic questions

Strategy

How can we leverage agentic autonomy to reinforce our core strategic objectives, brand promise, and competitive advantage?

Risk, security, and governance

What controls and oversight should we establish to identify, monitor, and mitigate risks from autonomous decisions?

Data

Do we have the scalable, high-quality data pipelines and feedback loops required to train, monitor, and continuously improve agents?

Platform and tech enablement

Which platform upgrades, integration patterns, and orchestration capabilities are critical to widely deploy and scale autonomous agents?

Workforce and change management

How should we reskill, redeploy, and engage talent so the workforce thrives alongside increasing levels of agent autonomy?

Get in touch

Vivek Kulkarni
Managing Director,
AI Transformation
Deloitte Services LP
vivkulkarni@deloitte.com

Prakul Sharma
Principal, AI & Data
Deloitte Consulting LLP
praksharma@deloitte.com

Greg Vert
Principal, Human Capital
Deloitte Consulting LLP
gvert@deloitte.com

Caroline Ritter
Manager,
AI Transformation
Deloitte Services LP
carritter@deloitte.com

Contributors to this report:
Beena Ammanath, Jim Rowan, Scott Holcomb

Endnotes

1. Maya Derrick, *Gartner Report: Only 44% of CIOs Seen as AI-Savvy by CEOs,* Technology Magazine, May 8, 2025, https://technologymagazine.com/ai-and-machine-learning/gartner-report-only-44-of-cios-seen-as-ai-savvy-by-ceos.
2. Emma Keen, *Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027,* Gartner, June 25, 2025, https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027.
3. *Artificial Intelligence Infrastructure Spending to Surpass the $200Bn USD Mark in the Next 5 years, According to IDC*, IDC, February 18, 2025, https://my.idc.com/getdoc.jsp?containerId=prUS52758624.
4. *Trust in Organizational Generative AI,* informed by the Feb-Apr 2025 Deloitte TrustID® Study.
5. *Strategies to Build Readiness for Transformational Changes*, Gartner, August 14, 2024, https://www.gartner.com/en/documents/5676155.
6. Shannon Poynton, Sue Cantrell, Nic Scoble-Williams, David Mallon, and Gaurav Lahiri, *Stagility: Creating stability for workers for organizations to move at speed*, Deloitte Insights, March 24, 2025, https://www.deloitte.com/us/en/insights/topics/talent/human-capital-trends/2025/organizational-agility-and-creating-stability-at-work.html.
7. Sue Cantrell, Jason Flynn, and Nic Scoble-Williams, *AI is revolutionizing work. You need a human value proposition for the age of AI.*, Deloitte Insights, March 24, 2025, https://www.deloitte.com/us/en/insights/topics/talent/human-capital-trends.html#ai-revolutionizing-work.

Further reading

- **Deloitte AI Institute**
  Prompting for Action (2024)
- **The Business Operations Room | Deloitte's Executive Blog**
  The Autonomous Enterprise (2025)

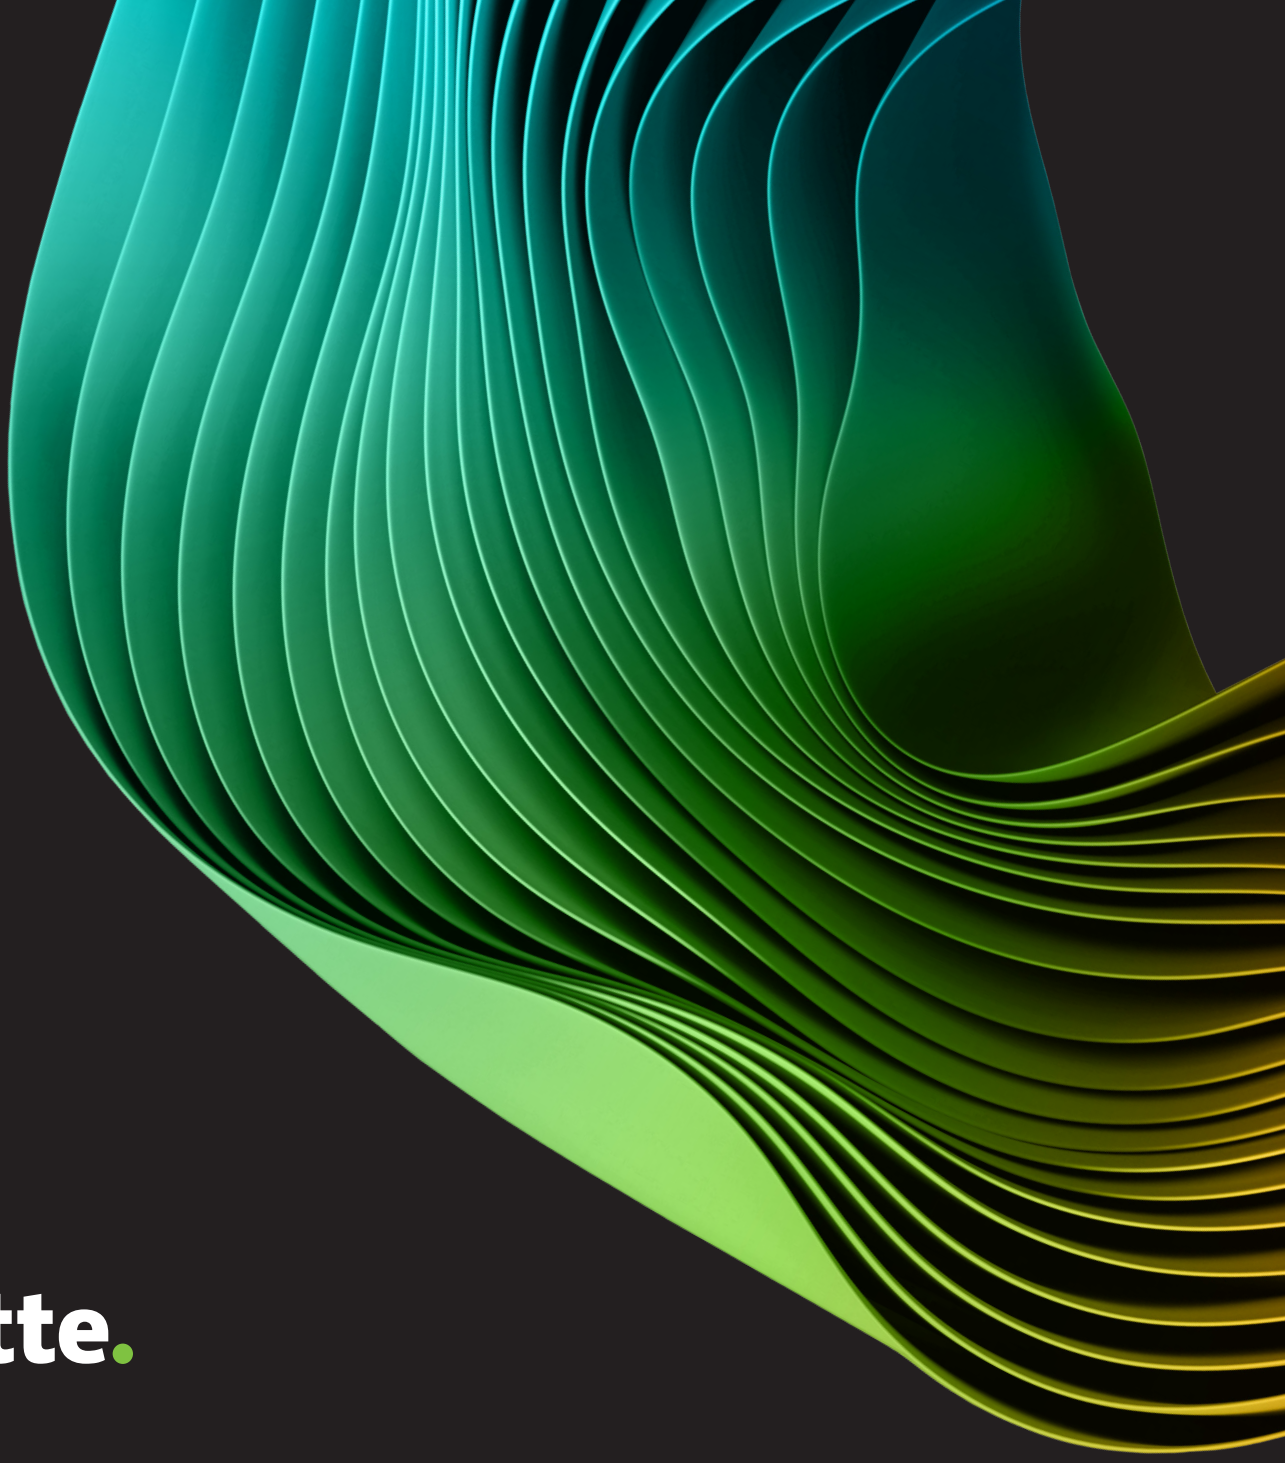

Deloitte.

About Deloitte

As used in this document, "Deloitte" means Deloitte Consulting LLP, a subsidiary of Deloitte LLP. Please see www.deloitte.com/us/about for a detailed description of our legal structure. Certain services may not be available to attest clients under the rules and regulations of public accounting.

Deloitte provides industry-leading audit and assurance, tax and related services, consulting, financial advisory, and risk advisory services to nearly 90% of the Fortune Global 500® and thousands of private companies. Our people deliver measurable and lasting results that help reinforce public trust in capital markets, enable clients to transform and thrive, and lead the way toward a stronger economy, a more equitable society, and a sustainable world. Building on its 175-plus year history, Deloitte spans more than 150 countries and territories. Learn how Deloitte's approximately 457,000 people worldwide make an impact that matters at www.deloitte.com.