



CAPITALIZING ON THE AGENTIC ECONOMY

Unlock enterprise value by harnessing
autonomous agents to reshape business
operations and drive the next wave
of transformation





Executive takeaways

- The **agentic enterprise** is a **structural transformation**, not an incremental upgrade.
- Governance — not model capability — determines whether autonomy is safe.
- Legacy and process debt must be addressed first — without reengineering the core, agentic capability remains confined to pilots and prototypes.
- In the future, agent-native applications will redefine how enterprises run.
- Early adopters will shape the agent-to-agent (A2A) ecosystem and enjoy early-mover advantage.
- DXC Technology provides the engineering, governance and organizational bridge required to deliver safe, scalable autonomous execution.

The agentic imperative

Breakthrough autonomous AI capabilities have triggered a major tech market repricing: AI has evolved from a conversational assistant into an engine of execution.

Historically, enterprise architecture assumed people ran processes supported by IT. Today, the industry is shifting toward the **agentic enterprise**, an operating model where autonomous agents initiate processes, adapt to uncertainty, take action and tee up decisions for human confirmation.

Previous automation waves, such as robotic process automation (RPA), largely papered over the cracks of legacy systems by simply speeding up rigid, human-designed workflows and navigating interfaces built for people.

The agentic shift is fundamentally different. Next-generation operations must be highly adaptable, context aware and outcome driven. With agents now proactively initiating action rather than just following scripts, the strategic mandate for leadership has changed: The challenge is no longer the mapping of static workflows — it is now the establishment of governance guardrails that align autonomous execution with overarching corporate intent.

In large organisations, this is a fundamental business model shift. Autonomy must be constrained from day one by governance, clear decision rights and customer trust. This transformation serves as the foundation for an emerging **agent-to-agent (A2A)** economy, starting with internal collaboration and eventually extending into an interconnected B2B ecosystem.

The primary barrier is no longer AI capability, but reengineering foundational IT and rewiring the enterprise operating model to safely delegate authority. While the immediate challenge is rewiring legacy platforms, the future belongs to agent-native applications. These applications are designed from inception for machine execution rather than for human interaction.



Transitioning manual execution to autonomous workflows enormously reduces baseline costs and drives significant margin expansion. ... But beyond simple cost deflection, multi-agent systems actively accelerate top-line growth by executing market strategies.

The market catalyst and dual wall

Frontier models now feature “native tool use,” allowing them to navigate databases and trigger workflows via APIs. This significantly reduces reliance on traditional user interfaces for routine tasks, a shift that has directly triggered the repricing of seat-based software-as-a-service (SaaS) models.

But this power is hitting a dual wall:

1. Legacy process debt

2. Human-centric operating models

Business workflows rely on employees not only to bridge disconnected systems, but to hold all operational decision authority. The prerequisite for adoption is untangling this legacy process debt while rewiring the business to safely delegate execution away from human-centric operating models.

The agentic equation: Balancing opportunity with risk

Removing humans from the execution loop reshapes unit economics, enabling speed and scale. However, it also means that failures occur at machine speed.

Economic upside

Transitioning manual execution to autonomous workflows enormously reduces baseline costs and drives significant margin expansion. Industry benchmarks (Boston Consulting Group) show cost reductions of 25% to 40% for complex processes. As an example, DXC Technology's deployment of over 30 agents at global infrastructure company Ferrovial successfully reduced manual effort by up to 40%.

But beyond simple cost deflection, multi-agent systems actively accelerate top-line growth by executing market strategies. JPMorgan Chase, for instance, attributes approximately \$2 billion in annual value to AI systems that autonomously hedge portfolios and generate pitch decks.

Risk profile

Moving to this type of model requires corporate/executive board-level accountability and a clear understanding of the systemic risks of human-free execution:

- **Regulatory and data exposure.** Unsupervised access can create vulnerabilities and non-compliance with frameworks like the European Union's General Data Protection Regulation [GDPR] and the California Consumer Privacy Act [CCPA] in the United States.
- **Systemic operational risk.** Machine-speed execution means small errors could propagate quickly across interconnected systems.
- **Reputation and customer trust.** When probabilistic reasoning shifts from harmless text to autonomous action (e.g., incorrect approvals), the potential for reputational damage is immediate.

Executive assurance model

To capture this economic upside without triggering these systemic risks, the agentic enterprise must operate within a hardened, three-layer assurance model:

1. **Policy guardrails** — Autonomy is bound by risk appetite and mapped to explicit decision rights.
2. **Technical safeguards** — Semantic guardrails, validation gates, deterministic fail-safes and audit grade logging prevent unsafe execution.
3. **Operational oversight** — Human-in-the-loop (HITL) checkpoints, kill switches and confidence based escalation ensure that humans remain in control of novelty and risk.

Speed comes from autonomy — but safety comes from design, not from hope.

The control spectrum: Right-sizing autonomy

An agentic-powered enterprise does not necessarily mean the use of generative AI for everything. Instead, safe autonomy relies on an architectural separation of concerns, routing tasks based on the required level of control:

- **Deterministic logic (absolute control)**. Hard-coded rules and APIs for high-stakes execution (e.g., financial transactions). Outcomes are 100% predictable.
- **Machine learning (high control)**. Bounded statistical models for specific tasks like risk scoring and fraud detection.
- **Generative AI (high flexibility)**. Large language models (LLMs) provide the cognitive glue to interpret intent and orchestrate workflows, but the LLMs hand off final, critical actions to deterministic or machine learning-driven gates.

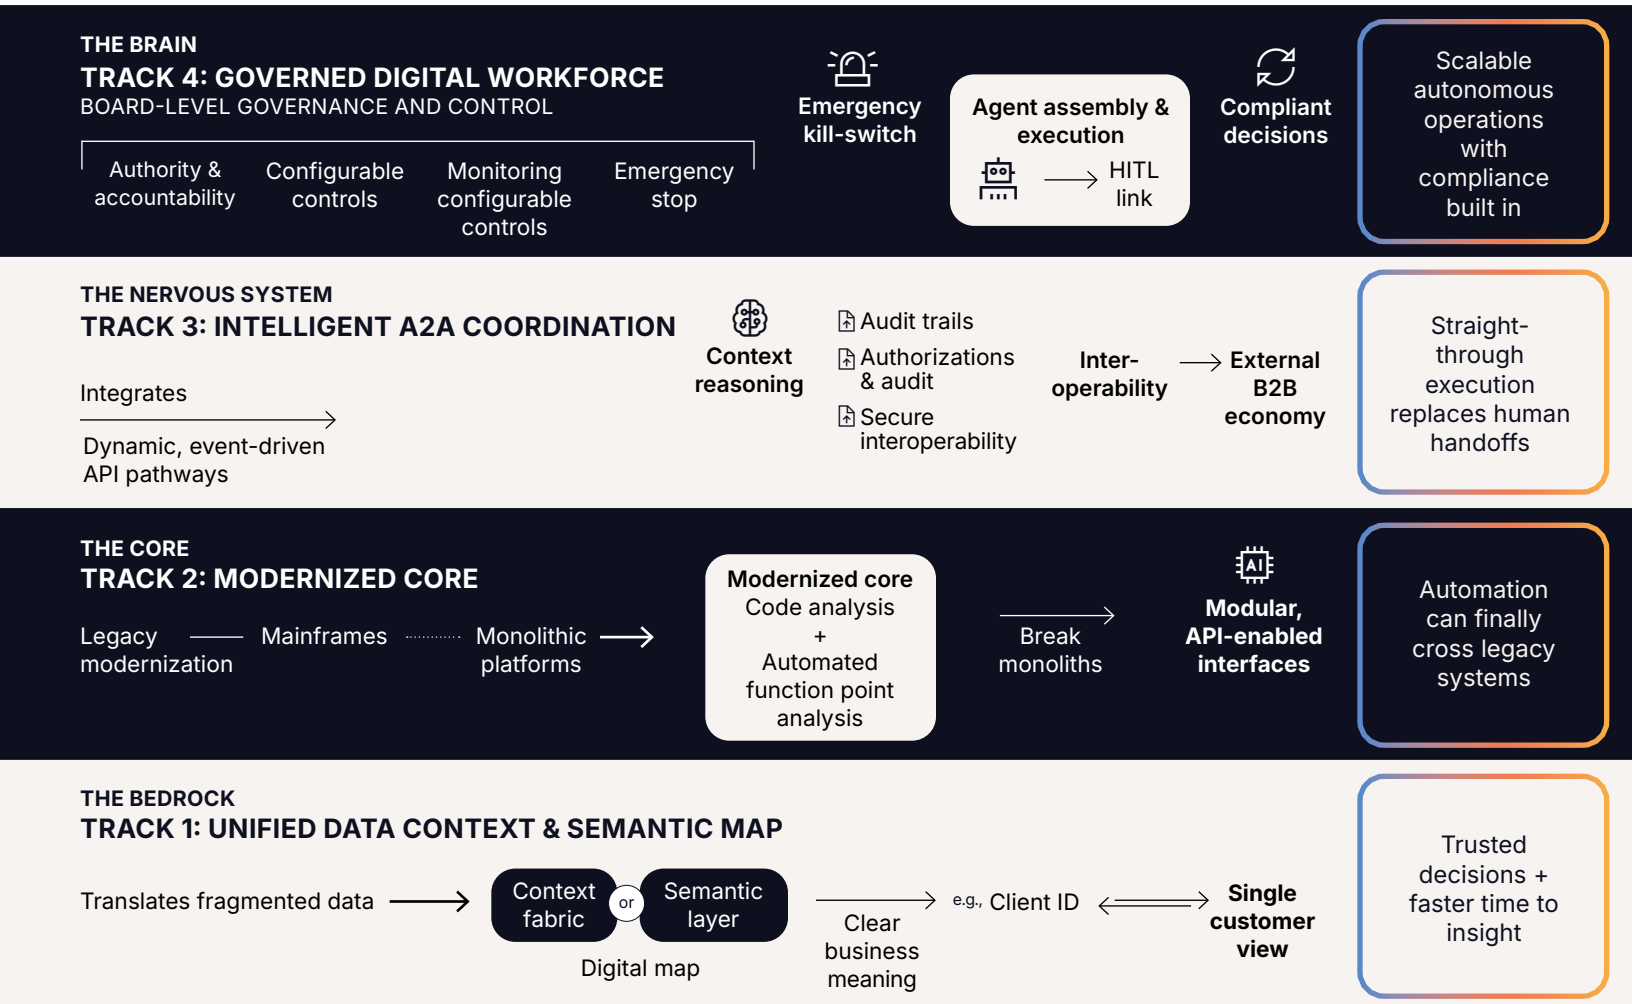


The DXC Xponential roadmap for autonomous execution

DXC Xponential operationalizes this shift through an incremental engineering approach that evolves existing investments, addressing targeted business value streams. Crucially, safe and controlled autonomy is engineered from the outset by embedding policy constraints directly into the data, execution and orchestration layers.

We identify **four non-sequential transformation tracks**. The first three tracks build on longstanding modernization principles, but their purpose has shifted: They are now the non-negotiable foundation for agentic operations. Without these tracks, agents remain trapped in prototypes; with the tracks, enterprises achieve architectural readiness for safe machine execution. The fourth track consists of activating and deploying the digital workforce.

- **Track 1 – the bedrock: Unified data context** prepares the data in readiness for AI
- **Track 2 – the core: Agent-reachable core** addresses legacy limitations that prevent agent actions
- **Track 3 – the nervous system: Agent-to-agent (A2A) ecosystem** establishes dynamic pathways for agents to collaborate
- **Track 4 – the brain: Governed digital workforce** provides the governance protocols for safe execution



Track 1: Creating a unified data context

AI models need deep, structural context to act responsibly. Rather than just pooling raw data, we establish a **context fabric**, a digital map of entities, rules and decision points. This semantic layer translates fragmented data into clear business meaning, ensuring that humans and machines share a unified understanding of how work is done (for example, mapping "Client ID" in a CRM system and "A_Num" in a mainframe to the exact same customer). With privacy controls embedded directly into this fabric, it becomes the critical bridge from isolated AI experiments to a cohesive, enterprise-grade system.

Architectural impact: Avoids massive data migrations. Accelerates cross-silo AI value by mapping data where it resides.

Track 2: Unlocking the agent-reachable core

For AI agents to add business value, they need to connect with your core systems, including legacy code and enterprise platforms. We speed this up by using autonomous engineering agents that analyze old code and automate function point analysis to provide clear transformation estimates. Through this data-driven process, we break down monolithic systems into modular, API-enabled interfaces that enforce transaction limits and strong security.

Engineering impact: Legacy modernization becomes iterative, lowering risk. Automated discovery replaces manual code analysis, making it more practical and cost effective to address technical debt.

Track 3: Building the agent-to-agent (A2A) ecosystem

We remove operational friction between systems by deploying enterprise orchestration platforms to serve as a unified agentic layer. Unlike traditional RPA, which depends on fragile UI and fixed rules, this layer leverages enterprise-grade integration to build dynamic, event-driven API pathways. This enables an internal agent-to-agent ecosystem with contextual reasoning, letting agents coordinate processes, manage exceptions, and ensure that actions are securely authorized and tracked with tamper-proof audit trails. Once validated internally, these secure channels can connect to the wider B2B economy, allowing your digital workforce to interact independently with external enterprise agents.

Operational impact: Moves from the "RPA trap" of automating broken processes to fundamentally redesigning them for intelligent, straight-through processing via a centralized orchestration engine.



As agentic capabilities mature, these same foundations support the progression from decision support to fully autonomous execution — demonstrating [DXC] Xponential blueprint's replicability in a highly complex environment.

Track 4: Deploying the governed digital workforce


With the legacy core exposed, processes reengineered and data unified, we activate the governed digital workforce. This track is divided into two fundamental components: strict enterprise governance, and actual assembly of the autonomous agents.

Component 1: Enterprise governance frameworks and safety controls. To ensure absolute control over autonomous execution, technical guardrails must be driven by robust, non-technical enterprise governance. We embed the Executive Assurance Model directly into the architecture to satisfy board-level scrutiny by answering critical leadership questions:

- **Delegated authority:** What authority is delegated? Autonomy is bound by risk appetite, mapping exactly which decisions an agent can make independently and what requires authorization.
- **Accountability:** Who is accountable? Agents act as delegates of human owners. We establish clear business ownership for every autonomous process, ensuring that compliance, legal and risk teams maintain continuous oversight of the policies governing agent behavior.
- **Monitoring:** How is it controlled? HITL validations and configurable escalation points ensure that execution pauses when agents encounter high liability, novelty or low confidence.
- **Emergency authority:** Who can stop the system? Per-agent “kill switches” and predefined rollback procedures instantly arrest errant behavior and restore the last known good state.

Component 2: Agent assembly and execution. Following established guidelines, the digital workforce is assembled and deployed via enterprise orchestration platforms, and role-specific agents take charge of the new digital processes.

- **Role definition and system access.** Rather than writing traditional code, we define the agent's specific persona, its operational boundaries, and its exact level of access to the enterprise systems unlocked in Track 2.
- **Execution logic and decision pathways.** The agent's workflow clearly distinguishes routine business rules from cases needing adaptive AI reasoning. The system flexibly handles exceptions while keeping decisions policy based, risk assessed and transparent to human oversight.
- **Business impact.** The component compresses multistakeholder approval cycles from days to minutes, with every decision logged and explainable.



Proof at scale: Italy's SIM platform

DXC is successfully delivering this type of transformation across numerous private and public sector organizations. One example of where DXC's complete blueprint has been applied is Italy's Sistema Avanzato ed Integrato di Monitoraggio e Previsione (SIM), where disconnected public administrations have been transformed into a single, autonomous engine for climate risk management.

With SIM, by applying DXC Xponential principles, we have delivered:

- **Scale.** Unified 40 disconnected public administrations and fused 3,000+ live data streams into a single operational view
- **Decision support.** Ability to independently generate climate-risk forecasts, replacing manual analysis cycles
- **Speed.** Real-time alerting enables emergency response within minutes, directly protecting physical infrastructure
- **Governed reliability.** Mission-critical workloads execute within strict architectural and security guardrails

SIM validates the foundational tracks of the Xponential blueprint. As agentic capabilities mature, these same foundations support the progression from decision support to fully autonomous execution — demonstrating the Xponential blueprint's replicability in a highly complex environment.

The DXC Xponential advantage: Engineering the bridge

DXC's differentiation lies not in access to frontier models but in our ability to safely operationalize autonomy at enterprise scale.

- **Agent-led engineering.** Engineering agents accelerate code discovery, architecture mapping, requirements tracking and transformation planning — reducing the required effort from months to days.
- **Platform and architecture first, not tool first.** We rebuild the foundational data, integration and governance layers — deploying robust enterprise orchestration platforms that serve as the engine for safe autonomy, rather than relying on fragmented point solutions.
- **Organizational adoption that lands.** We guide the workforce transition from manual "doers" → supervised "reviewers" → strategic "orchestrators" of digital workers, ensuring that cultural adoption sticks.

The result is a transformation blueprint that is **faster, safer and enterprise grade**, from day one.

About the authors



Karim Jeribi is senior vice president at DXC Technology, where he leads Consulting & Engineering Services for Europe. He is responsible for regional growth, portfolio strategy and delivery execution across consulting, software engineering, data & AI, platform modernization, and digital transformation services. Connect with Karim on [LinkedIn](#).



Matt Kay serves as an executive leader in DXC Technology's Consulting & Engineering Services in Europe. He is responsible for shaping DXC's strategic direction throughout the region, with a particular focus on AI implementation and its integration into client service offerings. Matt has a strong background in strategic transformation programs for both public and private sector clients. Connect with Matt on [LinkedIn](#).

Sources

[DXC Technology Newsroom – Ferrovial AI Workbench Announcement, April 2025](#)

[AlInvest Analyst Article – DXC & Ferrovial Efficiency Impact, April 2025](#)

[CNBC – JPMorgan Agentic AI & Pitch Deck Automation, Sept. 2025](#)

[LinkedIn – JPMorgan GenAI Case Study \[\\$2B Annual Value\], Sept. 2025](#)

[Boston Consulting Group – Feb. 2026](#)



DXC Technology (NYSE: DXC) is a leading enterprise technology and innovation partner delivering software, services, and solutions to global enterprises and public sector organizations — helping them harness AI to drive outcomes at a time of exponential change with speed. With deep expertise in Managed Infrastructure Services, Application Modernization, and Industry-Specific Software Solutions, DXC modernizes, secures, and operates some of the world's most complex technology estates. Learn more on [dxc.com](#).

[LinkedIn](#) | [Instagram](#) | [TikTok](#) | [YouTube](#)