#62 SUMMER 2026

Journal

The Capco Institute Journal of Financial Transformation

A new world order

Strategic pivots

From global factories to sovereign AI pods: how geopolitics and GenAI rewire software delivery

Gerhardt J. Scriven

KING'S College LONDON
KING'S BUSINESS SCHOOL

CAPCO
a wipro company

The Capco Institute Journal of Financial Transformation

Journal

Recipient of the Apex Award for Publication Excellence

Table of

contents

Colliding worlds

Monetary machinations

Strategic pivots

Private markets

2026, Edition 62

Journal

Welcome to the 62nd edition of the Capco Journal of Financial Transformation, the second to be published in partnership with King's College London, a world-renowned leader in education and research.

This year marks the Journal's 25th anniversary and the beginning of **Capco30**, our strategic vision as we build towards Capco's landmark 30th anniversary.

At the heart of Capco30 is The Expert Advantage: our commitment to be the trusted partner clients turn to when the stakes are highest. We are proud to have built an exceptional business by obsessing about our deep expertise, our strong client relationships, our innovative culture and our reputation for impact.

The Journal has embodied that ambition since its launch in 2001. Over the past quarter century, it has brought together leading academics and industry practitioners – including Nobel laureates, international regulators and board members of major institutions – to examine the forces shaping financial services and the wider global economy. We are equally proud to see Capco experts sharing their perspectives alongside such distinguished contributors.

As Professor Crawford Spence, our editor from King's College, observes in his introduction, this edition might aptly have been titled 'how to survive and thrive in turbulent times'.

After decades of globalization, firms are both refocusing on local market shifts and new forms of cross-border collaboration. Technology is driving integration as trade barriers rise; financial innovation is accelerating amid monetary realignment and volatile geopolitics; and the green transition is redrawing both energy markets and fiscal priorities.

Edition 62 focuses on the **new world order** emerging from these shifts – and how leaders across financial services and energy are rethinking established assumptions, navigating new risks and identifying opportunities for growth.

The coming decade could bring one of the most consequential reconfigurations of global business and economic systems in generations.

This environment reinforces the importance of expertise, judgement and the ability to turn insight into action. At Capco, we help clients connect vision models and strengthening resilience to reimagining customer experiences and the future of work.

We bring together industry depth, technology and data capabilities, and disciplined execution to help clients navigate the moments that matter most with confidence.

The Journal is one expression of that Expert Advantage: convening leading perspectives, challenging established thinking and helping shape the agendas that will define the future of our industries.

I hope you will find the articles in this edition both thought-provoking and valuable for informing your organization's future vision and strategic decisions.

My sincere thanks and appreciation to our contributors, readers and clients.

Annie Rowland, Capco CEO

2026, Edition 62

Editor's note

KING'S BUSINESS SCHOOL

We could equally have labelled this volume of the Journal of Financial Transformation 'how to survive and thrive in turbulent times', given that each of the articles effectively attempts to tackle this issue in their own unique way. Times are indeed turbulent. Recent months and years have borne witness to increased military conflict, heightened economic shocks and the emergence of new political forces.

These disruptions are often captured in popular discourse by reference to, variously, the decline of the rules based international order, the rise of authoritarianism and the re-emergence of nationalism and protectionism as challenges to globalization. While putting clear conceptual parameters around what is a very unsettled geopolitical period is difficult, there is at the very least little doubt that the world is changing and that a new order is emerging. To paraphrase an old political aphorism: the old is dying and the new is struggling to be born.

This geopolitical reshuffling is the backdrop to the current issue of the Journal. Contributors were asked to think about how financial services and energy might respond to, and indeed themselves help reshape, the new global order that is emerging. What they have presented here can be grouped into four conceptual clusters, each of which makes up a distinct sub-theme in this issue of the journal: colliding worlds; monetary risks; strategic pivots; and private markets.

The first group of articles we refer to here as colliding worlds. The worlds that collide here specifically are those of finance and energy. Each of these forces us to zoom out from financial transactions and investments and think about the wider environmental and material basis for these. For example, Ragland's fascinating look at the arms race for computing power shows how the rapid expansion of data centers is at odds with the way in which existing energy infrastructure is built. Those who will be successful in rolling out advanced computing power are those who will be able to finance and build out alternative energy infrastructure at the same time.

This has strong parallels with Xu and Preda's article looking at the infrastructure of bitcoin mining, which demonstrates in some detail how supposedly digital assets have heavy materiality in terms of both the physical infrastructure that supports them and the energy consumption that goes into their mining and processing. Finally, Savignon takes the energy transition head on by exploring the mechanics and possibilities of carbon credits in Brazil, a country which is characterised by a dominance of renewables in its energy mix. Showing real heterogeneity in how carbon credits are deployed in different sectors in Brazil, this article concludes by outlining how those who are able to integrate thinking around finance, data and energy will be able benefit from and contribute to the energy transition in the Brazilian market.

Monetary risks refer to the various ways in which currencies are embroiled with risk at both the national and transnational levels. As the fundamental means via which deals are transacted, it is perhaps no surprise that we find money, in its various guises, to have so many political inflections. Cryptoassets are often heralded as some kind of safe haven ('digital gold') during times of economic and political turmoil and there are many central banks across the globe who are now pursuing digital currency projects of their own. These have significantly lagged the development of private crypto initiatives. One of the reasons for this might be the perceived financial crime risk that is associated with digital currencies in general, something which Tjeransen's article outlines in detail.

Fiat currencies are also deeply imbricated in geopolitical factors. This is the subject of Hou et al's article which draws attention to the ways in which currency risk premia are associated with trade imbalances between nations. Such factors are of increasing importance to, not only forex traders, but any organization engaging in currency hedging. Finally, nowhere is the importance of currency to geopolitics more evident than in Zayim's article looking at how central banks engage in currency swaps with other central banks. Focusing on both China and the US, Zayim shows how each has sought to build up spheres of influence with different counterparties as part of larger attempts to accrue power and minimize geopolitical risk for themselves.

Our third conceptual 'bucket' looks at how individual organizations respond to geopolitical shifts. In other words, these articles look at the various strategic pivots that organizations make in order survive and thrive in turbulent times. These strategic pivots can take different forms depending on the geography and industry in question. For example, Schnyder shows how banks in East Central Europe can mitigate risk in response to rising anti-FDI policies in that region. In Scriven's article the author shows how software delivery needs to be reconfigured in order to overcome increasing regulatory divergence and export controls.

In times of geopolitical uncertainty, capital markets can become less efficient in allocating capital and so firms have to focus more on their internal efficiencies, which is what Gkikopoulos and Stathopoulos illustrate in their article exploring the relationship between firms and banks; such relationships have implications for tax planning and productivity. Political turbulence also impacts hiring strategies, which Nguyen and Sila demonstrate in relation to corporate hiring practices during times of election uncertainty. Finally, Confalone makes a strong case for a strategic pivot towards developing 'enterprise intelligence' in the age of AI, a shift which has governance implications for financial services organizations.

The final sub-theme of the issue reflects an ongoing interest that we have had in private markets. The shift from public to private markets is one of the big secular shifts of modern finance, which was dealt with explicitly in issue 61. In this issue, we build on this with two further articles. One of these, by Phalippou, tackles the shift of private equity towards retail investors, outlining the potential risks in selling private equity products that were designed with institutional investors in mind to a less informed and less powerful constituency. Preda and Chen's article pursues another tack, showing how venture capital funding in Asia and elsewhere has to be cognisant of surrounding government policy imperatives if it is to be successful.

Taken together, these articles argue that financial services are not simply being buffeted by geopolitical change. They are becoming one of the arenas in which that change is negotiated. Through payment systems, liquidity backstops, energy and computing infrastructure, corporate strategy and the evolving boundary between public and private markets, we can see how geopolitical forces manifest themselves and adapt to existing institutional constraints and imperatives.

Professor Crawford Spence
King's College London

From global factories to sovereign AI pods:

how geopolitics and GenAI rewire software delivery

Author | **Gerhardt J. Scriven** | Executive Director, Capco

Abstract

While geopolitics has always shaped the global order, it is now exerting a far more direct and constraining influence, transforming regulatory divergence, data sovereignty, sanctions and export controls (including restrictions on technology transfer), and critical-infrastructure resilience into binding constraints on cross-border service delivery in financial services and energy. In parallel, GenAI is reshaping global service delivery models by compressing routine tasks into AI-assisted workflows and deepening technical integration even as jurisdictions fragment. Using the software development lifecycle (SDLC) as a high-signal microcosm, this paper shows how geopolitical constraint and AI intensity shift delivery from labor arbitrage factories toward smaller, higher seniority, jurisdiction bounded AI augmented crews, while reshaping education pathways and high skill mobility. A 2×2 framework maps four plausible futures – global factory, augmented global delivery, localized duplication and sovereign AI pods – and the conditions that produce each. It closes with instrumentable indicators and a sequenced response to avoid localized duplication by hardening AI-aware SDLC controls, redesigning sourcing around data sovereignty, and rebuilding talent pipelines so that systems thinking and secure engineering become early-career defaults. This enables staffing for the control plane – practitioners who can validate AI outputs and hold accountable decision rights in regulated and critical environments – while also supporting selective mobility for scarce, high-accountability roles.

1. Introduction

After decades in which globalization expanded cross-border operating models and reinforced large-scale labor arbitrage, the operating environment for financial services and energy firms is tightening in two directions at once.

First, geopolitical contraction is becoming operational. In practical terms, this increasingly manifests as regulatory divergence across jurisdictions, data localization and sovereignty expectations, export controls and sanctions regimes, intensified third-party scrutiny as well as renewed emphasis on critical infrastructure protection [Airaudo et al. (2025), Gopinath et al. (2024), PIFS (2024)]. For example:

- For the financial services domain – under the E.U.'s Digital Operational Resilience Act (E.U. DORA) banks and insurers cannot treat cloud/tooling merely as procurement; they must be able to evidence resilience, governance, and concentration-risk controls across their critical providers [European Union (2022)].

- For the energy domain – in the U.S., critical infrastructure protection reliability standards translate directly into measurable control obligations that shape what operational telemetry exists, who can access it, and how changes are governed – all of which feeds back into delivery geography and operating models (NERC CIP-010-4; NERC CIP-013-2).

Note that financial services and energy are treated together here because both are regulated, resilience-critical sectors where technology delivery is inseparable from auditable controls, accountable decision rights, and third-party risk – making them structurally sensitive to the same 'how/where' pressures.

These regulatory pressures do not merely add compliance overhead; they constrain where sensitive work can be performed, how systems are architected, and which vendors and locations remain viable for delivery at scale. It should be noted that fragmentation does not necessarily suggest a collapse of globalization in aggregate: flows can remain substantial overall, while being reweighted and re-routed along geopolitical alignment and policy constraints.

Second, generative AI (Gen AI) is changing the mechanics of technology delivery. In practical terms, GenAI accelerates and partially automates routine knowledge-work tasks embedded in the software development lifecycle (SDLC): translating requirements into artefacts; drafting and refactoring code; generating tests and documentation; and improving incident triage and runbook execution [DORA (DevOps Research and Assessment) (2024), Cui et al. (2025)]. For example, reporting on Meta's internal coding assistant describes capabilities including analyzing failed tests and submitting proposed fixes for human review [Dixit (2025)]. Similarly, Google Cloud documentation describes how Gemini Cloud Assist can be used directly inside Cloud Logging to summarize log entries during incident triage [Google Cloud (2025)]. Within Capco Brazil's AI lab, a publicly available requirements document was analyzed from seven distinct viewpoints (inspired by the Systems Engineering Methodology) using GenAI, and the corresponding resultant assets were then used to generate a working and bug-free application prototype using Agentic AI – all within a single day.

Notably, this evolution is not suggesting a wholesale replacement of the engineering discipline, but rather a reconfiguration of how work is produced, reviewed, and governed – often compressing cycle times while increasing the premium on verification, risk controls, and senior technical judgment.

The combination of geopolitical contraction and AI-infused delivery creates a perfect storm in both how work is done (team composition, controls, cadence, and accountability) and where work is done (the geography of delivery, the viability of offshore factory models, and the demand for jurisdiction-specific expertise) for the financial and energy sectors. These particular sectors act as bellwethers for how geopolitical constraint and GenAI-driven software delivery compression translate into binding operating constraints on tooling, data access, and accountable decision rights. The SDLC itself provides a precise lens for analyzing these shifts because it sits at the intersection of productivity, risk, and cross-border delivery – and it is structurally central to both financial services and energy, where complex technology estates must meet stringent resilience, safety, and regulatory expectations.

Thesis: The combined impact of regulatory fragmentation and GenAI is reshaping the operating environment for the new global order, shifting delivery away from large labor-arbitrage models toward smaller, higher-seniority, jurisdiction-bounded AI-augmented crews, with material knock-on effects for education pathways and high-skill mobility. This thesis is demonstrated through changes in SDLC mechanics and the resulting reconfiguration of delivery geography.

The remainder of this paper translates these two forces into a practical operating-model lens: a 2×2 framework that distinguishes four delivery archetypes and a small set of indicators that show which archetype an organization is drifting towards. The SDLC is used throughout as the microcosm to show where these forces manifest and what a sequenced response might look like in practice.

2. A simple framework: global delivery archetypes by constraint and AI intensity

Towards establishing a compact model for this complex operating environment, we utilize a simple 2×2 framework, as is illustrated in Figure 1.

In Figure 1, the horizontal axis represents AI augmentation intensity (low → high): the extent to which delivery work is routinely accelerated or partially automated by GenAI across SDLC activities. The vertical axis represents geopolitical/regulatory constraint (low → high): the degree to which delivery is bounded by jurisdictional rules, supervisory expectations, and critical-infrastructure risk posture. Note that in energy

Figure 1: Global delivery archetypes by regulatory constraint and AI intensity

Geopolitical/regulatory constraint		
High	**3. Localized duplication** • Regional builds • Duplicate controls • Higher unit cost since AI does not compress cost **Financial services example:** E.U./U.S./APAC delivery splits into parallel regional stacks: separate cloud tenants/environments, CI/CD, SOC/SIEM, and release approval boards; limited AI adoption. **Energy example:** Critical infrastructure delivered under multiple regimes runs separate change-control/compliance evidence and incident-response authority/evidence stores; operational technology constraints limit integration of modern tooling.	**4. Sovereign AI pods** • Small senior crews • Jurisdiction-bound, viable delivery • Risk-led governance **Financial services example:** Small, senior in-region credit decisioning crew: AI runs policy-to-logic checks, scenario/fairness analysis, and exception triage; humans approve cutoffs, overrides, and adverse-action reasons; data and logs stay local. **Energy example:** Small, senior plant/grid reliability pod with in-region tool/data boundaries; AI supports code/test generation and incident summaries, while deployment authority, change approval, and safety/resilience evidence remain explicitly human-accountable platform guardrails.
Low	**1. Global factory** • Offshoring and pyramid economics • FTE scale • Standardized playbooks with incremental AI **Financial services example:** Low-risk apps (marketing/public content/internal tools) built by large offshore pyramid teams using standard CI/CD and conventional QA. **Energy example:** Non-critical corporate apps (HR/procurement) outsourced offshore under standard controls and limited safety/regulatory exposure.	**2. Augmented global delivery** • AI-enhanced cross-border teams • Higher throughput per FTE • Centralized enablement with manageable boundaries **Financial services example:** Distributed squads use AI copilots on a centralized engineering platform; cross-border delivery remains viable because data is synthetic/sanitized or boundary-contained. **Energy example:** Global teams deliver asset-performance analytics; operational data is processed locally/abstracted and only aggregates flow globally; AI accelerates code/tests/docs within platform guardrails.
	Low	High

AI-augmented Intensity

– and increasingly in systemically important financial services – critical-infrastructure reliability and resilience requirements translate directly into enforceable control expectations, making delivery design (tooling, access, evidence, and change authority) a first-order operating constraint rather than an implementation detail.

The 2x2 framework yields four operating-model archetypes:

1. **Global factory** (low constraint, low AI) reflects classic offshoring and pyramid economics optimized for scale and unit cost, i.e., business as usual.
2. **Augmented global delivery** (low constraint, high AI) retains cross-border teaming but retools it around AI copilots and centralized platforms, typically increasing throughput while preserving global integration at the workflow layer.
3. **Localized duplication** (high constraint, low AI) emerges where jurisdictional barriers force regional builds and duplicated control stacks, while AI augmentation intensity is low (or adoption is constrained by policy). It results in multiple mini-factories, one per jurisdiction, each with its own tooling, governance, and risk control raising cost and reducing change velocity.
4. **Sovereign AI pods** (high constraint, high AI) combine tight jurisdictional boundaries with AI-enabled productivity, resulting in smaller, higher-seniority crews, stronger governance, and tooling/data that is explicitly bounded to a regulatory perimeter.

Crucially, this framework is not a forecast of a single future but describes the conditions under which different delivery models become economically and operationally rational – and it provides a consistent lens for assessing implications for workforce formation, education pathways, sourcing strategy, and high-skill mobility.

3. The SDLC case: what changes in how work is done

The SDLC is a structured set of activities organizations use to translate an idea or requirements into working software that is released into day-to-day business use [ISO/IEC/IEEE (2017)]. While several delivery methodologies exist (e.g., sequential waterfall, iterative agile, or hybrid variants), the SDLC exists to make change repeatable and defensible by establishing clear handoffs, decision rights, and quality controls across planning, design, build, test, deployment, and operations phases. While SDLC operations and governance practices may differ across institutions, the underlying lifecycle principles and disciplines, especially those tied to risk, traceability, and operational resilience, remain persistent. As such, the SDLC is an appropriate micro-foundation for understanding the operational impact of GenAI because it is where productivity, risk management, and organizational accountability intersect.

At an industry level, DORA's 2024 research explicitly frames AI as having a significant impact on software development practices and outcomes [DORA (DevOps Research and Assessment) (2024)]. At the level of measured productivity, large-scale randomized field experiments with software developers show that access to an AI coding assistant can increase completed tasks, with the combined analysis across three real-world trials reporting a 26.08% increase in weekly completed tasks among developers using AI [Cui et al. (2025)].

Note that this does not eliminate the SDLC; it changes the balance of work inside it – shifting effort from first-draft production toward verification, governance, and higher-order engineering judgment.

What follows is a pragmatic stage-by-stage decomposition: SDLC stage, what AI changes, what it does not change, implications for team shape and controls. The intent is to make the AI-augmented crew concrete and operational rather than aspirational.

3.1 Plan/discover

What AI changes. In planning and discovery, GenAI can improve the speed and breadth of discovery artifact synthesis, such as summarized stakeholder inputs, initial requirement specifications, as well as comparable solution options. In practice, this increases throughout into the delivery pipeline – more candidate options, faster initial articulation, and quicker iteration on scope.

What AI does not change. In regulated, safety- and availability-critical environments, GenAI cannot own prioritization, accountability, or regulatory interpretation. The decision rights around what to build (and what not to build), why it is justified, and how it aligns with risk appetite remain inherently human because they are governance decisions, not text-generation tasks.

Implications for team shape and controls. The AI-augmented crew typically shifts from requirements authorship towards requirements stewardship. Controls tighten around (i) traceability (what sources informed the requirement), (ii) approval (who signs off), and (iii) reproducibility (whether the same artifact can be regenerated and validated). This naturally pushes teams toward smaller groups with clearer accountable owners, supported by AI for drafting and analysis.

3.2 Design

What AI changes. In design, GenAI accelerates pattern selection and option exploration: feeding off reference architectures, producing structured design documents, and prompting threat modeling questions that broaden coverage (e.g., common misconfigurations and data handling paths). Qualitative evidence from software engineering professionals suggests GenAI can increase productivity but also surfaces adoption barriers that are sociotechnical (how tools fit into routines and organizational practices), not merely technical [Russo (2024)].

What AI does not change. Architectural decisions remain bounded by operational realities (latency, resilience, observability, audit evidence, and integration complexity), none of which can be dismissed by faster drafting. In energy-sector delivery, this validation burden is often amplified by the interface between IT delivery and operational technology environments, where safety and availability expectations make failure modes and change impact analysis materially less forgiving.

Implications for team shape and controls. Design becomes more explicitly governed. An AI-augmented crew benefits from standardizing design checkpoints (architecture review, security review, data-handling review) and instrumenting design artifacts for accountability. This is also where trustworthy AI disciplines begin to appear as delivery controls: NIST's Generative AI Profile, for example, emphasizes governance, pre-deployment testing, and incident disclosure as key considerations for managing risks novel to or exacerbated by generative systems [Autio et al. (2024)].

3.3 Build

What AI changes. In build activities, GenAI speeds scaffolding, refactoring, and translation across languages/frameworks. It can propose implementations, generate boilerplate code, and assist in navigating legacy estates by explaining unfamiliar code paths. The empirical signal here is increasingly difficult to dismiss: the large-scale field experiments cited earlier found not only higher completed task volumes but also increases in intermediate production indicators such as code updates (commits) and compilation frequency – consistent with acceleration of iterative development work [Cui et al. (2025)].

What AI does not change. Correctness, secure-by-design decisions, and dependency hygiene remain human responsibilities. The faster code can be produced, the more important it becomes to ensure it is not merely plausible, but correct, maintainable, compliant with enterprise standards, and resilient to security weaknesses.

Implications for team shape and controls. This is the point in the SDLC where the economics of the junior-heavy engine starts to change. When first-draft code production becomes cheaper, the bottleneck shifts to expert review, integration, and risk controls. The pragmatic control response is a stronger review and validation spine: code review standards that explicitly account for AI assistance, secure SDLC automation (e.g., scanning and policy gates), and stronger traceability of generated artifacts. NIST's emphasis on lifecycle risk management provides a useful external anchor for why velocity must be paired with control [Autio et al. (2024)].

3.4 Test

What AI changes. GenAI expands test generation capacity: proposing unit and functional tests together with regression scenarios and generating synthetic test data derived from requirements and prior defect patterns. It can also draft test documentation and translate between business acceptance criteria and technical assertions, improving alignment between intent and verification.

What AI does not change. Test strategy, quality gates, and evidence suitable for audit remain human-owned. In financial services and energy, testing is not simply about coverage – it is about demonstrating controlled change, reliability, and defensible operation under stress. That requires judgment about what must be proven, at what rigor, and how to document that proof.

Implications for team shape and controls. AI pushes teams toward a more formalized notion of evaluation as a first-class artifact. In addition to code tests, teams increasingly maintain evaluation suites for AI-assisted workflows (e.g., prompt libraries, retrieval configurations, and acceptance thresholds for model-assisted outputs). This aligns with risk-management emphasis on pre-deployment testing and governance for generative systems [Autio et al. (2024)].

3.5 Deploy/operate

What AI changes. In operations, GenAI improves triage and diagnosis workflows: summarizing incidents, proposing likely root causes, generating runbook steps, and accelerating knowledge retrieval across fragmented documentation or directly from the source code itself. This can reduce mean time to understand (MTTU) and make operational teams more effective – particularly in complex estates.

What AI does not change. Production risk decisions and change approval logic remain human responsibilities, even if workflow automation is extensive. High-impact operational actions – those affecting safety, systemic stability, or regulatory reporting – require explicit accountability and controlled decision rights. In energy and other critical infrastructure domains, this constraint is often reinforced by operational separation and formalized change authority for high-impact environments, where evidence of safe operation and controlled change can be as material as functional delivery velocity.

Implications for team shape and controls. Human-in-the-loop becomes a design principle, not a slogan. AI outputs are treated as decision support; approval rights and escalation paths are clarified; and incident response is instrumented for audit and disclosure expectations. NIST's profile explicitly includes incident disclosure as a key consideration in GenAI risk management, reinforcing the need to connect operational AI use to governance and reporting disciplines [Autio et al. (2024)].

3.6 The central 'how' impact

The SDLC is not dead; it is reconfigured. As GenAI compresses routine production work, delivery value shifts toward architecture, verification, provenance, security, and accountable approvals, i.e., a thicker control plane and higher seniority at decision points. This is because speed without controls magnifies risk and reduces defensibility. This is the operating-model logic behind the move from traditional pyramid delivery (junior-heavy execution) through diamond models (more experienced blended onshore/offshore teams) toward AI-augmented crews: smaller teams, higher effective seniority in decision points, and a thicker control plane embedded throughout delivery.

4. The delivery geography case: what changes in where work is done

When GenAI reduces the quantity of routine build work and simultaneously increases the importance of governance, auditability, and data controls, delivery location becomes more constrained and more strategic as a direct outcome. Put differently, AI shifts the binding constraint from "how many developer-hours can we buy in the cheapest market?" to "where can we legally, securely, and accountably execute an SDLC that now depends on sensitive data, model usage controls, and regulator-facing evidence?". This is consistent with evidence that geoeconomic fragmentation is increasingly expressed through policy-driven frictions, including a wider use of interventions that shape trade and financial flows across geopolitical lines [Airaudo et al. (2025)].

4.1 From labor arbitrage to risk and jurisdiction arbitrage

In classic global factory delivery, geography decisions optimize labor cost and scale. Under higher-constraint conditions (upper half of Figure 1), the dominant optimization changes: the highest-value SDLC decisions – architecture, security posture, model/tool approvals, and release accountability – move closer to (i) the data, (ii) the regulated or critical-infrastructure entity, and (iii) the executives who carry institutional accountability.

This is not simply cultural preference. In financial services, cross-border data movement is increasingly constrained by localization requirements and supervisory expectations that shape viable cloud architectures and operating models [PIFS (2024)]. In energy and critical infrastructure industrial contexts, enforceable cybersecurity and supply-chain requirements similarly harden expectations around configuration change control, vulnerability assessment, and vendor access – placing practical limits on where production-adjacent work and evidence generation can occur [NERC CIP-010-4; NERC CIP-013-2; ISA/IEC 62443].

Once GenAI tooling becomes embedded in SDLC execution, the control plane (data access paths, telemetry, provenance, and evidence) becomes integral to delivery itself – not an overlay.

4.2 From 'factory scale' to 'crew leverage'

The shift from pyramid economics toward smaller, higher-seniority AI-augmented crews changes the geometry of delivery – spanning how software is developed as well as how it is implemented and then maintained over a longer period. A ten-person pod led by a senior engineer, architect, and risk/security counterpart can be localized (or duplicated regionally) far more easily than a 200-person offshore factory that depends on high volumes of junior build capacity. In the language of Figure 1, this represents one way in which organizations can migrate from global factory toward sovereign AI pods when constraints rise: AI compresses routine throughput, making smaller teams viable without necessarily losing productivity.

It should be noted, however:

- that duplication remains expensive without strong platform foundations (common pipelines, shared controls, reusable reference architectures) – which is why the Augmented Global Delivery quadrant often depends on centralized platforms and standardized governance patterns that can be deployed across regions;
- while AI intensity remains low – because tools cannot be adopted at scale or governance prevents sustained use – organizations can land in Localized Duplication: regional builds and duplicated control stacks, with higher cost and slower change.

4.3 What happens to offshore delivery

Offshoring does not end; it changes form. As routine build work becomes less economically central, offshore delivery is more likely to concentrate in:

- **Specialized engineering centers** (platform engineering, security engineering, test automation, SRE/observability, tooling enablement), where capabilities can be industrialized and reused; and
- **Tightly bounded operational roles,** including follow-the-sun operations, but with stricter segmentation: hardened toolchains, controlled datasets, and auditable handoffs rather than open-ended access to production-adjacent environments.

The practical effect is a more explicit partitioning of SDLC activities into sovereign work (must be local) versus portable work (can be distributed), and that partition often maps cleanly onto data access and change authority.

4.4 Vendor model implications

As AI changes unit economics, FTE-based contracting weakens: the buyer is no longer purchasing linear capacity; they are purchasing outcomes delivered through a mix of humans, automation, and vendor platforms. As point in case, Deloitte's 2024 Global Outsourcing Survey reports broad adoption of AI within outsourced services and highlights increased emphasis on value/outcome-based relationships, alongside heightened governance needs for the extended workforce ecosystem [Deloitte Insights (2024)].

Market data also indicates continued enterprise spend shifting towards cloud-first platforms in the context of AI adoption [ISG (2025)]. The governance burden rises in parallel: model usage policies, provenance logging, audit trails, and evidence packages become contractually material, because they determine whether the delivery model is regulator-acceptable and operationally defensible.

This is one of the clearest ways financial services and energy actively shape the new global order: by embedding jurisdictional accountability and evidentiary standards into procurement and delivery contracts, they indirectly reshape vendor roadmaps, managed-service designs, and the practical rules of engagement for cross-border digital work. Where integration used to be driven primarily by cost and platform scale, it increasingly becomes conditional on auditability, provenance, and control planes. In the new world a vendor or platform provider often cannot realize outcomes through a remote factory alone; it needs senior engineers embedded close to the domain, stakeholders, and operational constraints.

5. Second-order effects

5.1 Education and the career pipeline

The fundamental implication for education is not simply that skills must change, but that the SDLC learning ladder changes when routine tasks are automated. Historically, early-career software engineers progressed by producing high volumes of first-draft work (boilerplate code, straightforward tickets, basic test cases, and documentation) while learning standards and risk instincts through repetition and review. As GenAI tools measurably increase throughput on routine software tasks [Cui et al. (2025)] and become embedded in delivery practices [DORA (2024)], that rung of the learning ladder compresses. The same is true for non-technical finance professionals: in many finance roles, early-career progression has historically depended on high-repetition tasks – building spreadsheets, drafting memos, preparing reconciliations, compiling management reporting, writing policy/control narratives, etc. GenAI can now generate credible first drafts of many of these outputs. That changes the learning ladder in the same way it does for engineers: less time spent producing, more time required supervising, validating, and exercising judgment.

The result is an acute training challenge: how to develop sound engineers when the practice surface area shifts away from first drafts and toward validation, integration, and control. This pushes a concrete rebalancing of curricula and early-career development toward four emphases:

1. **Systems thinking and architecture fundamentals.** Graduates must reason across dependencies, service boundaries, data lifecycles and failure modes, i.e., systems thinking – and develop skills that remain non-trivial even when code is faster to draft [World Economic Forum (2025)]. In safety- and availability-critical contexts, systems thinking is not only a mindset; it is frequently operationalized through systems engineering and model-based systems engineering (MBSE) disciplines aligned to lifecycle standards (INCOSE, Systems Thinking).
2. **Verification and validation discipline.** Testing strategy, reasoning about edge cases, and evidence generation for auditability become core competencies, not specialist add-ons [Autio et al. (2024)].
3. **Secure coding and dependency risk.** As code generation accelerates, so does the risk of propagating insecure patterns or fragile dependencies; security must be taught as an engineering default, not a downstream gate [Autio et al. (2024)].
4. **Working effectively with AI tools.** Prompting is secondary. The primary skill is supervision and evaluation: assessing output quality, detecting failure modes, enforcing policy, and maintaining provenance and accountability [Autio et al. (2024), UNESCO (2023)].

Apprenticeship models therefore need to be redesigned. Juniors can learn by validating AI output, expanding tests, improving guardrails, and conducting structured reviews – work that directly builds judgment while producing defensible artifacts for regulated environments. At a policy level, this aligns with broader calls for more agile, data-driven skills and lifelong learning as skill demand shifts faster than traditional education cycles are able to adapt [OECD (2025a)], alongside concerns that training supply may lag emerging AI skill needs [OECD (2025b)].

There is also a need to understand the feasibility of workers acquiring newly valuable expertise as AI changes task composition [National Academies of NASEM (2025)]. Ultimately, the objective is not to advocate labor substitution, but to ensure the talent pipeline remains viable as the SDLC reorganizes around AI-augmented production and control.

5.2 Migration and talent mobility

The ‘where’ implications of SDLC reconfiguration are ultimately expressed through people flows, but in this paper, they are best framed as a workforce composition and location problem driven by operating constraints rather than as a general debate about migration. As delivery models shift away from large labor-arbitrage factories and toward smaller, jurisdiction-bounded AI-augmented crews, the center of gravity moves toward roles that carry explicit accountability for risk, controls, and regulatory defensibility. Three implications follow:

- First, organizations should expect increased demand for high-accountability roles in-region – including architects, security and resilience engineering leadership, AI governance leads, and risk owners who can sign off controls and evidence packages. This follows from the way GenAI changes the SDLC: it accelerates production, but raises the premium on governance, provenance, pre-deployment testing, and incident disclosure disciplines, which are inherently tied to accountable decision rights [Autio et al. (2024)]. Where data localization and supervisory expectations constrain architectures and operating models, accountable roles predictably cluster closer to regulated entities, sensitive data, and primary regulators [PIFS (2024)].
- Second, continued demand for global talent is to be expected, but with more selective profiles. As routine work compresses, firms will disproportionately value adaptable technology generalists, senior engineers, and IT professionals who combine domain context with risk fluency – capabilities aligned toward analytical, systems, and creative thinking, technology literacy, and continuous learning [World Economic Forum (2025), OECD (2025a)]. The practical effect is not less global talent, but a reweighting toward roles that can supervise AI-assisted workflows, validate outputs, and operate within tighter governance constraints.
- Third, a plausible organizational response is the emergence of talent corridors: formalized mobility pathways for scarce, high-accountability roles as firms rebalance where accountability must sit. Recent OECD analysis describes how OECD countries are increasingly tailoring labor migration channels to specific labor market needs, adjusting eligibility thresholds, and expanding bilateral and regional cooperation mechanisms that facilitate labor mobility [OECD (2025c)]. In practice, these corridors are likely to be explicitly aligned to roles that sit at the intersection of engineering, governance, and regulated delivery.

6. Conclusion: what leaders should watch and do

The analysis in the previous sections translates a complex macro narrative into operational mechanics where two forces are colliding inside delivery organizations:

- GenAI is compressing the SDLC and shifting the work mix toward higher-order judgment. As a result, delivery value moves from sheer volume of build capacity to control-plane capabilities: verification, provenance, security, architecture, and accountable approvals.
- Jurisdictional friction is increasing. Re-concentration of control activities in-region is consistent with data localization and supervisory constraints, and with broader evidence that fragmentation increasingly operates through policy-driven frictions.

Importantly, financial services and energy do not merely adapt – through themselves hardening evidentiary expectations for digital operations, third-party dependence, and critical infrastructure resilience, they actively influence the operating standards and commercial designs that will govern cross-border digital delivery in the new world order.

The 2×2 framework (Figure 1) makes the operating-model implications concrete. AI creates a rightward pull that can increase throughput and standardize workflows while regulation creates an upward push that limits where work, tooling, data, and decision rights can reside. However, this is the new global order contradiction in delivery form: technical integration can deepen even as legal and political perimeter fragment. Depending on where a firm sits on these two axes, it will rationally converge on one of four archetypes (Global Factory, Augmented Global Delivery, Localized Duplication or Sovereign AI Pods) – each with distinct economics, risk posture, and delivery geometry, i.e., how delivery is physically and organizationally arranged which, in practical terms, captures:

- **Footprint and distribution:** how many delivery sites exist, where they are, and whether delivery is centralized or dispersed.
- **Team topology:** how work is split across onshore/nearshore/offshore; the mix of local versus global teams; and whether teams are pods, centers, or factories.
- **Flow of work and handoffs:** the number and type of handoffs across borders, time zones, and legal entities; and whether work is sequential (assembly-line) or end-to-end (pod-based).
- **Placement of decision rights:** where approvals, accountability, and risk ownership sit (e.g., local sign-off versus centralized governance).
- **Data/tooling residency constraints:** where data is processed, where environments run, and how tightly they are partitioned by jurisdiction.

Second-order effects follow operational mechanics directly. As routine drafting of SDLC deliverables becomes cheaper and faster, the SDLC learning ladder compresses: junior progression shifts from producing first drafts to validating AI output, expanding test evidence, and strengthening guardrails. In parallel, the talent location equation changes: accountable roles (risk, security, architecture, model governance) concentrate closer to regulated entities and critical environments, while global talent demand persists but becomes more selective in terms of the type of work that is done and credential heavy.

For leaders, the central question is not whether GenAI will change delivery, but whether the organization can industrialize GenAI without drifting into Localized Duplication – a state where rising constraints force regional replication of tooling and controls while productivity gains fail to offset the cost and complexity. The practical objective is to land in a stable operating model – typically Augmented Global Delivery and/or Sovereign AI Pods by explicitly designing the control plane, the sourcing footprint, and the talent pipeline. Towards this, a viable approach is to instrument a small set of observable indicators that reveal which quadrant the organization is moving towards and then act in a sequenced way to steer toward the desired operating model.

6.1 Signals to watch

1. **Delivery seniority mix.** Track the senior-to-junior ratio in delivery teams, especially in decision roles (architecture approvals, security sign-off, release accountability). If this ratio rises steadily, it indicates that AI is reducing routine production work.

2. **Speed versus quality debt.** Instrument cycle time using lead time for change and deployment frequency and pair this with rework and stability measures (e.g., defect escape rate, rollback rate, and change failure rate). Acceleration is only net-positive if it does not convert into downstream incident load or audit-defensibility gaps.

3. **Contract construct shift.** Track the share of spend and critical workloads delivered under FTE/time-and-materials constructs versus outcome-, platform-, or managed-service constructs. Separately, track whether contracts explicitly require AI usage controls, which become material as governance expectations rise.

4. **Geographic redistribution of control versus execution.** Track where control plane work is performed: risk acceptance, security exceptions, model/tool approvals, change authority, incident disclosure decisions, and audit evidence production – versus where execution plane work is performed (build, automation, lower-sensitivity operations).

Using these signals, a simple scoring system (that turns the four signals into indicators) can be devised to calculate AI augmentation versus constraint scores and then correlate this back to the 2x2 framework. This datapoint – a snapshot in time – can be refreshed quarterly.

6.2 What to do

Actions that steer the organization toward its chosen delivery posture (i.e., the intended mix of archetypes in the 2x2 framework) by workload sensitivity and jurisdictional constraint):

1. **Redesign SDLC controls for AI-augmented workflows**

   Embed AI-related evaluation gates, provenance, secure usage policies, and incident disciplines into delivery rather than treating them as compliance overlays.

   How this helps:

   This step builds the control plane that makes higher AI augmentation scalable and defensible – these disciplines convert AI use from tool adoption into an auditable operating capability. Without this, organizations may see speed gains but will struggle to sustain them under audit, resilience scrutiny, or supervisory challenge – especially in high-criticality environments.

2. **Redesign sourcing footprint and vendor models around jurisdictional constraints**

   Once controls and evidence requirements are explicit, partition delivery into sovereign work (must sit in-region due to data, accountability, or critical-infrastructure constraints) versus portable work (can remain distributed).

   How this helps:

   Geography becomes a design constraint rather than an emergent problem. Furthermore, it aligns vendor commercials with the new unit economics, i.e., contracting for outcomes as opposed to raw capacity.

3. **Redesign talent pipelines**

   With the target operating model defined through steps 1 and 2 above, retool apprenticeship, curricula partnerships, and internal academies toward supervision, verification, and systems thinking – so that the pipeline can reliably produce engineers who can operate AI-augmented workflows under regulated constraints.

How this helps:

This rebuilds the pipeline around the new early-career practice surface area so the organization can staff AI-augmented workflows without overpaying for scarce senior labor indefinitely.

Over time, the industry will not converge on a single archetype; rather, the market share of the four operating models will shift as geopolitical constraint and AI intensity evolve. In the near term (through ~2026), the distribution is likely to remain weighted toward the left-hand quadrants: Global Factory persists for low-sensitivity work, while Localized Duplication expands as organizations respond to sovereignty and policy frictions faster than they can industrialize GenAI control planes and obtain governance/assurance comfort.

By approximately 2030, the center of gravity is likely to move rightward, as GenAI becomes embedded in delivery practice and as evaluable, auditable control patterns (provenance, pre-deployment testing, incident discipline) become standardized and productized into platforms. Augmented global delivery should grow where data can be bounded technically; sovereign AI pods should become the default for high-sensitivity and critical environments.

This transition is unlikely to be linear: expect an S curve (see Figure 2 as illustrative example) – incremental adoption while controls, permissions, and skills mature, followed by more rapid reconfiguration once defensible control planes are deployable at scale.

Ultimately, organizations in tightly regulated sectors such as finance and energy are best positioned for the dual shift in geopolitical contraction. GenAI reshaping delivery models will be those that treat GenAI not as a productivity add-on, but as a catalyst to re-architect delivery – aligning controls, geography, and talent into a coherent, auditable operating model that remains scalable under fragmentation.

Figure 2: Illustrative share of right-hand archetypes in the 2x2 framework between 2026 and 2030

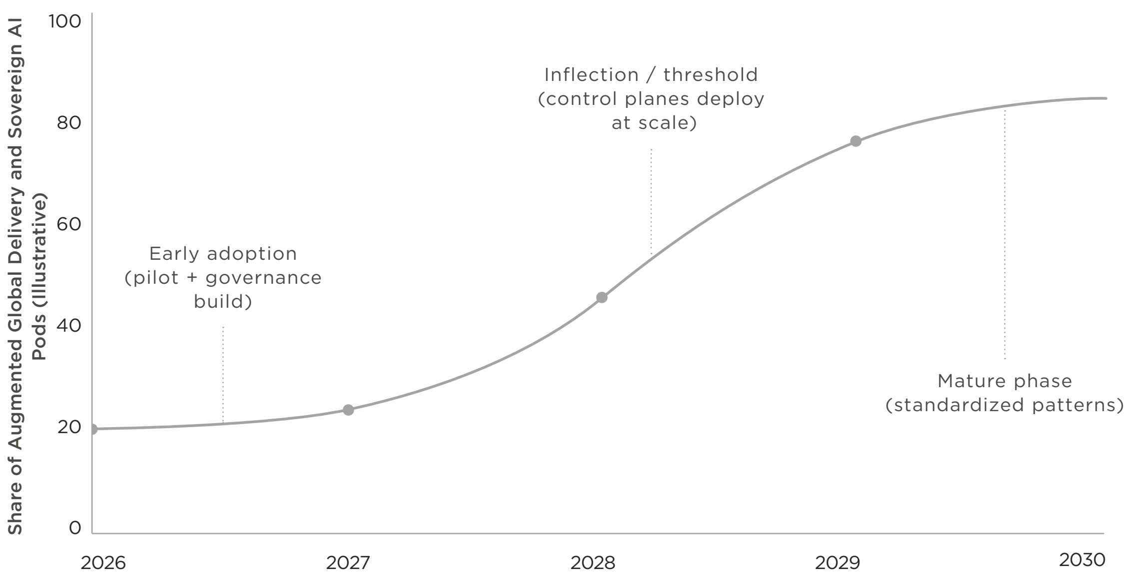

References

Airaudo, F., F. De Soyres, K. Richards and A.M. Santacreu, 2025. Measuring Geopolitical Fragmentation: Implications for Trade, Financial Flows, and Economic Policy, International Finance Discussion Papers No. 1408, Board of Governors of the Federal Reserve System

Altman, S. A. and C.R. Bastian, 2024. DHL Global Connectedness Report 2024, DHL Group (in partnership with NYU Stern)

Autio, C., R. Schwartz, J. Dunietz, S. Jain, M. Stanley, E. Tabassi, P. Hall,and K. Roberts, 2024. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile (NIST AI 600-1), National Institute of Standards and Technology. https://doi.org/10.6028/NIST.AI.600-1

Cui, Z. (Kevin), M. Demirer, S. Jaffe, L. Musolff, S. Peng, and T. Salz, 2025. The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers, working paper/preprint), Microsoft Research

Deloitte, 2024. Global Outsourcing Survey 2024: Multidimensional sourcing – Orchestrating the extended workforce ecosystem, Deloitte Insights

Dixit, P., 2025. Meta is using rival AI models like Anthropic's Claude for better coding efficiency, Business Insider, June 13, https://www.businessinsider.com/meta-anthropic-ai-llama-coding-efficiency-2025-6

DORA (DevOps Research and Assessment), 2024. Accelerate State of DevOps Report 2024, Google Cloud

European Union, 2022. Regulation (EU) 2022/2554 on digital operational resilience for the financial sector (Digital Operational Resilience Act – DORA), Official Journal of the European Union

Google Cloud, 2025. Gemini Cloud Assist / Cloud Logging documentation: log summarisation during incident triage, Google Cloud Documentation

Gopinath, G., P.O. Gourinchas, A.F. Presbitero, and P. Topalova, 2024. Changing Global Linkages: A New Cold War?, IMF Working Paper WP/24/76, International Monetary Fund

Information Services Group (ISG), 2025. AI Drives Tech Services, Software Spending to New High in Q3: ISG Index™, press release / ISG Index), ISG

INCOSE (International Council on Systems Engineering), n.d. INCOSE Systems Engineering Handbook (Version 5) – product page, INCOSE

ISA (International Society of Automation), n.d. ISA/IEC 62443 series of standards (Industrial Automation and Control Systems Security)

ISO/IEC/IEEE, 2017. ISO/IEC/IEEE 12207:2017 – Systems and software engineering – Software life cycle processes, International Organization for Standardization / International Electrotechnical Commission / IEEE

ISO/IEC/IEEE, 2023. ISO/IEC/IEEE 15288:2023 – Systems and software engineering – System life cycle processes, International Organization for Standardization, ISO

National Academies of Sciences, Engineering, and Medicine (NASEM), 2025. Artificial Intelligence and the Future of Work, National Academies Press https://doi.org/10.17226/27644

NERC. (CIP-010-4). Cyber Security – Configuration Change Management and Vulnerability Assessments

NERC. (CIP-013-2). Cyber Security – Supply Chain Risk Management.

OECD (Organisation for Economic Co-operation and Development), 2025a. OECD Skills Outlook 2025, OECD Publishing

OECD (Organisation for Economic Co-operation and Development), 2025b. Bridging the AI skills gap: Is training keeping up? OECD Publishing

OECD (Organisation for Economic Co-operation and Development), 2025c. International Migration Outlook 2025, OECD Publishing, https://doi.org/10.1787/ae26c893-en

Program on International Financial Systems (PIFS), 2024. Data Localization, Cloud Adoption, and the Financial Sector, PIFS International.

Russo, D., 2024. Navigating the Complexity of Generative AI Adoption in Software Engineering, ACM Transactions on Software Engineering and Methodology, 33(5), Article 135, https://doi.org/10.1145/3652154

UNESCO, 2023. Guidance for generative AI in education and research, United Nations Educational, Scientific and Cultural Organization

World Economic Forum (WEF), 2025. The Future of Jobs Report 2025, World Economic Forum.

About Capco

Capco, a Wipro company, is a global technology and management consultancy shaping change in the financial services and energy industries. For almost 30 years, we have been trusted to help our clients adapt, transform and create long-term value across capital markets, banking, payments, insurance, wealth and asset management, and the energy and utilities sectors. Our deep industry expertise, partnership mindset and award-winning Be Yourself At Work culture are amplified by our strengths in advisory, technology, data and AI innovation. We support our clients to establish clear priorities, connect vision to value and deliver measurable impact when the stakes are highest.

Expert-led, AI-infused, impact-focused – we do not just respond to change, we help shape it. To learn more, visit www.capco.com or follow us on LinkedIn, Instagram, Facebook and YouTube.

Worldwide offices

APAC
Bengaluru – Electronic City
Bengaluru – Sarjapur Road
Bangkok
Chennai
Gurugram
Hong Kong
Hyderabad
Kuala Lumpur
Mumbai
Pune
Singapore

MIDDLE EAST
Dubai

EUROPE
Berlin
Bratislava
Brussels
Edinburgh
Frankfurt
Geneva
Glasgow
London
Milan
Paris
Vienna
Warsaw
Zurich

NORTH AMERICA
Charlotte
Chicago
Dallas
Houston
Montreal
New York
Orlando
Toronto

SOUTH AMERICA
São Paulo
Rio de Janeiro

About King's Business School

King's Business School is a triple-accredited business school in the heart of London. It ranks amongst the top 10 business schools for research in the UK, based on the Research Excellence Framework 2021, and is consistently ranked in the top ten in the UK for business by the Times and Sunday Times Good University Guide, the Complete University Guide and the Times Higher Global University rankings.

The School advances world-leading education and research across four knowledge frontiers: (1) developing a modern workplace, (2) creating sustainable and socially responsible business, (3) Improving health and public services organisation and (4) strengthening economic and financial systems.

It is the ninth faculty of King's College London, leveraging interdisciplinary potential across the College. It drives pioneering thinking and engages the limitless energies of the city's businesses, policymakers, entrepreneurs and change-makers. King's Business School puts its commitment to drive positive change at the heart of its life-changing research and "London as a classroom" experiential education.

FSC
www.fsc.org
MIX
Paper from responsible sources
FSC® C004309