FRONTIER.

FRONTIER RESEARCH / DECISION-GRADE ANALYSIS



The CIO in 2030: Running the Enterprise Intelligence Economy

Fourteen operating imperatives for a world where people, agents, software, compute, data, and trust are dynamically allocated to outcomes. In

FR-021 / Dion Hinchcliffe / September 26, 2026

REPORT PROFILE

Built to change a decision.

Fourteen operating imperatives for a world where people, agents, software, compute, data, and trust are dynamically allocated to outcomes. In

REPORT	**FR-021**
AUTHOR	**Dion Hinchcliffe**
PUBLISHED	**September 26, 2026**
PROFILE	**frontier-ro-0.1**

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

The next four years will not merely add AI to the CIO's portfolio. They will change the economi of technology management.

For most of the information age, the CIO assembled relatively scarce technology into durable systems: applications, infrastructure, data platforms, teams, vendors, and controls. The emerg era inverts much of that model. Intelligence can increasingly be summoned on demand. Softw be generated and regenerated. Agents can execute work instead of only assisting with it. Com can be routed dynamically. Model capability and price change continuously. Business users ca create operational technology directly. And the traditional boundary between "the technology organization" and "the business" becomes less useful every year.

Frontier Research believes the CIO of 2030 will be defined less by the systems they own than effectively they allocate intelligence. The role becomes the designer and governor of an enterp intelligence economy: a system in which human expertise, machine intelligence, data, softwar compute, capital, and trust can be combined around outcomes with far greater speed and fluid than today's operating models allow.

That does not mean the future is predetermined. This report is a horizon model, not a literal forecast. The fourteen imperatives are deliberate extrapolations from observable signals in 2026. Regulated industries will move at different speeds. Physical work will change differently from digital work. Some enterprises will remain application-centric well beyond 2030. Technical limitations, liability, regulation, labor markets, and human preference will all constrain what becomes autonomous.

But waiting for certainty is itself a choice. The relevant question for CIOs is not whether every element of this model arrives exactly as described. It is whether the organization will be prepared if even half of it does.



FIG. 01 The CIO in 2030: fourteen interconnected operating imperatives for turning intelligence into value.

Executive perspective: the scarce resource is shifting

A useful way to understand the next CIO mandate is to ask what remains scarce as other things become abundant. Code is becoming less scarce. Generic information is less scarce. Baseline analytical capability is less scarce. Access to increasingly capable models is less scarce. Even sophisticated software production is beginning to move from a capital-like activity toward something closer to a variable input.

Scarcity moves elsewhere. It moves to high-quality proprietary context, trusted data, judgment, accountability, relationships, institutional knowledge, strategic coherence, and the ability to absorb change. It also moves to the organizational mechanisms that determine where autonomy is permitted, where people must remain accountable, and which combinations of intelligence are economically rational.

This shift is already visible. Microsoft's 2026 Work Trend Index, based on a survey of 20,000 AI users across ten countries, found that only 19% were in the strongest combination of individual AI capability and organizational readiness; 10% were "blocked" by organizations that had not caught up to their capabilities, and only 26% said leadership was clearly and consistently aligned on AI.[1] Deloitte's 2026 enterprise AI research likewise found that only one in five surveyed companies had a mature governance model for autonomous AI agents.[2] The technology is not the only bottleneck. Increasingly, the enterprise itself is the bottleneck.

METR's work on AI task-completion horizons illustrates the capability side of the pressure. Its longitudinal research found that the length of software tasks frontier AI agents could complete at a given reliability threshold had been increasing rapidly over multiple years, while the researchers also cautioned that benchmark task horizons do not translate cleanly into whole jobs or messy enterprise work.[3] Anthropic's 2026 Economic Index reports a parallel change in real-world usage: sessions are shifting from conversational assistance toward longer-running agentic work, with higher autonomy in agentic coding environments than in conventional chat.[4]

The strategic implication is straightforward. The dominant CIO problem is moving from access to capability toward organizational absorption: how quickly the enterprise can convert new capability into safe, measurable value.

Six disruptive headwinds

The fourteen imperatives in this report are not isolated predictions. They are responses to six forces arriving at the same time.

The first is capability acceleration. AI systems are moving from answering questions to handling delegated, multi-step work. That does not mean autonomy is universal or reliable, but it does mean the frontier of delegable work is expanding. Microsoft describes the corresponding organizational challenge as a need to rearchitect work, while OECD interviews with organizations deploying agentic AI show that governance and accountability are becoming central practical concerns rather than theoretical ones.[1][5]

The second is intelligence price volatility. The same task can increasingly be served by different models, model sizes, providers, inference strategies, and compute configurations. AWS already provides intelligent prompt routing that selects models based on expected response quality and cost, while Google Cloud spot compute exposes infrastructure pricing that varies with available capacity.[6] [7] These are early examples of a broader shift: intelligence is becoming a variable-priced production input.

The third is software and agent proliferation. AI-assisted software development reduces the cost of creating code, integrations, interfaces, and automation. Anthropic found that 79% of sampled Claude Code interactions in its 2025 study were classified as automation rather than augmentation, although it also emphasized the limits and self-selection in its sample.[8] DORA's 2025 research offers the more important organizational warning: AI acts as an amplifier, magnifying the strengths of high-performing software organizations and the dysfunctions of weak ones.[9] Cheap software does not eliminate technical debt. It can manufacture it faster.

The fourth is workforce recomposition. The International Labour Organization estimates that one in four workers globally is in an occupation with some generative-AI exposure, but concludes that transformation is more likely than wholesale replacement because human input remains necessary across most jobs.[10] The World Economic Forum similarly finds AI and big data among the fastest-growing skill areas through 2030 while analytical thinking, resilience, leadership, creativity, and other human capabilities remain central.[11] The change is therefore not simply fewer people. It is a different division of labor between people and machines.

The fifth is cyber and sovereignty fragmentation. AI is becoming embedded in critical operating processes at the same time that data residency, model provenance, supply-chain exposure, infrastructure sovereignty, and geopolitical constraints are increasing. This makes location, provider concentration, model behavior, and policy jurisdiction part of workload architecture rather than only legal review.

The sixth is trust and regulatory pressure. The EU AI Act moved into active enforcement phases in August 2026, including transparency requirements for certain AI systems and enforcement powers for the AI Office and national authorities.[12] NIST continues to frame AI risk management across the lifecycle, including trustworthiness considerations specific to generative AI.[13] PwC's 2026 work argues that trust-by-design is becoming a scaling capability rather than merely a compliance function.[14]

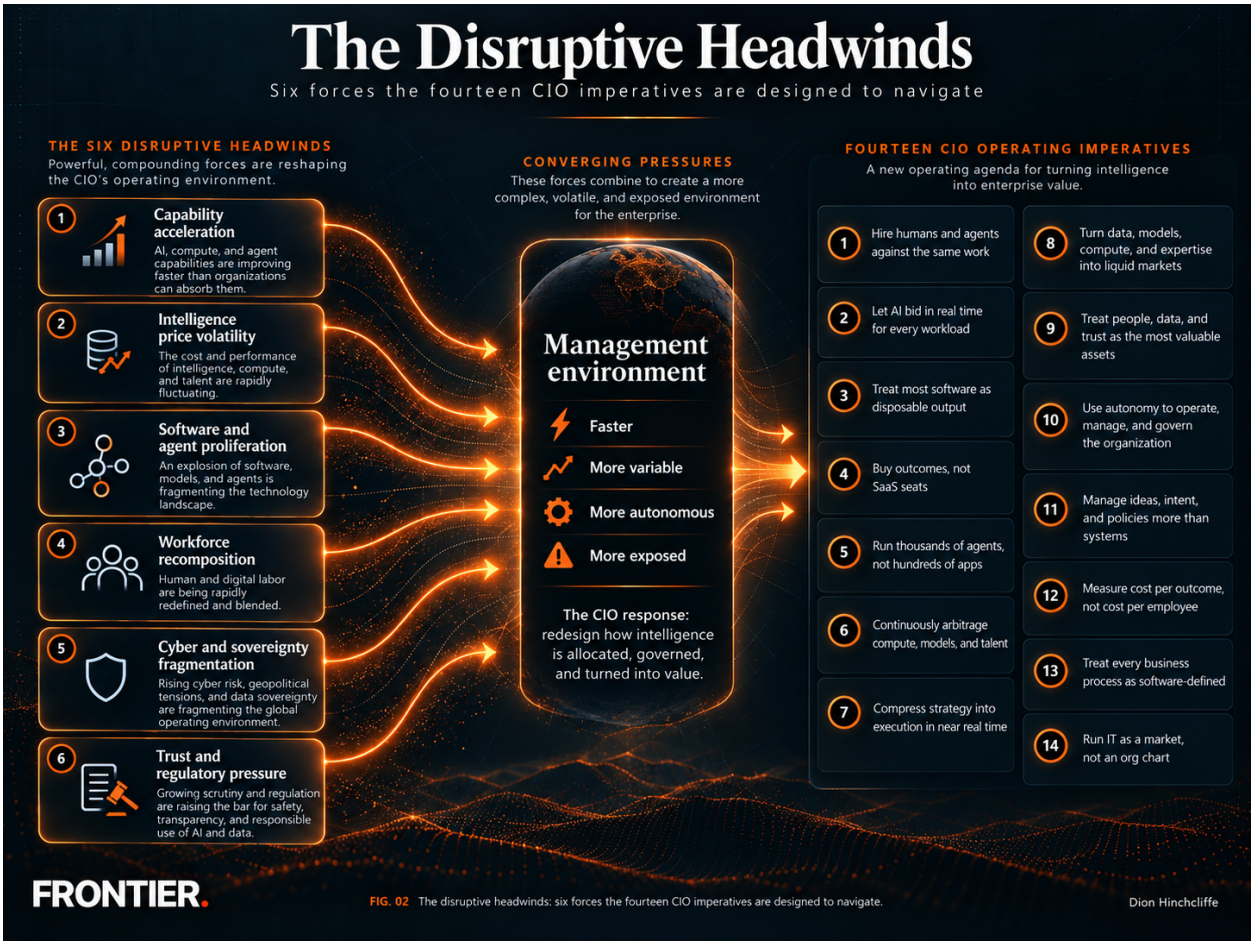


FIG. 02 The disruptive headwinds: six forces the fourteen CIO imperatives are designed to navigate.

Together, these headwinds create a management environment that is faster, more variable, more autonomous, and more exposed. The fourteen imperatives are a way to navigate it without confusing speed with progress.

1. Hire humans and agents against the same work

The workforce plan should increasingly begin with work rather than job titles. For every material outcome, the CIO and business leaders should ask which components require human judgment, accountability, empathy, negotiation, creativity, physical presence, or deep contextual understanding; which components can be delegated to agents; and which become stronger when humans and agents work together.

This is not a euphemism for replacing employees with software. It is a more rigorous form of organizational design. The ILO's exposure research explicitly distinguishes between exposure and redundancy, and the World Economic Forum's skills data suggests that human-centered capabilities remain important even as AI skills grow quickly.[10][11] A field experiment involving professionals at Procter & Gamble found that individuals using generative AI could match the performance of traditional two-person teams on the studied product-innovation tasks, while AI also helped bridge functional expertise boundaries.[15] That does not generalize to every kind of work, but it demonstrates why the unit of workforce design is changing.

The practical shift is from staffing positions to assembling capability. A future requisition may describe the outcome, accountability boundary, skills, context, and decision rights required, then determine the right mixture of people and agents. In some cases the human will execute and AI will assist. In others the agent will execute and a human will review. In still others a human will own the outcome while supervising a fleet of specialized agents.

The change-management risk is obvious. If employees hear "human-agent workforce design" only as code for headcount reduction, they will rationally resist the instrumentation and knowledge-sharing required to make it work. CIOs need an explicit social contract: automation will change work, some tasks and roles will disappear, new roles will emerge, and the enterprise will invest in mobility and capability rather than pretending disruption is painless. Credibility matters more than reassurance.

2. Let AI bid in real time for every workload

For decades, architecture decisions were comparatively static: select a platform, negotiate a contract, place the workload, and optimize it over time. AI creates a more dynamic decision surface. The best execution environment for a unit of work may vary by model quality, inference price, latency, data sensitivity, regional availability, energy characteristics, context-window requirements, and risk.

AWS's intelligent prompt routing already chooses among models within a family based on predicted quality and cost.[6] Google Cloud's spot infrastructure provides a separate example of dynamically priced compute, with discounts tied to available capacity and an explicit tradeoff around preemption.[7] These mechanisms are primitive compared with what an enterprise intelligence market could become, but the direction is clear.

"Bid" should not be interpreted narrowly as a literal auction. The important idea is competition among execution options. The enterprise should be able to evaluate approved combinations of models, compute, providers, locations, and policies against a workload's requirements and route accordingly.

CIOs should begin building this abstraction now. New AI workloads should expose measurable requirements for quality, latency, risk, residency, and unit cost rather than hard-coding permanent dependence on one model. Model routing, policy enforcement, observability, and evaluation should become platform capabilities. The strategic asset is not a favorite model. It is the ability to switch intelligently as the market changes.

3. Treat most software as disposable output

Software has historically been expensive enough to become an asset enterprises protect for decades. AI changes the economics of producing it. When an interface, integration, workflow, report, or application can be generated and modified quickly, some implementation should be treated less like a building and more like an artifact that can be regenerated when requirements change.

The word "disposable" is intentionally provocative. It does not mean careless, untested, insecure, or undocumented. It means the enterprise should stop assuming that every implementation deserves a long life simply because it was expensive to create. DORA's 2025 research is especially relevant: AI does not rescue weak engineering systems; it amplifies them.[9] McKinsey's 2026 work on AI-driven product development likewise finds that organizations seeing meaningful gains are redesigning the whole product-development system, including roles, verification, control, and change management, rather than merely adding coding assistants.[16]

The durable layer therefore moves upward. Architecture, interfaces, identity, data contracts, security controls, policy, provenance, evaluation, and tests become more important because the implementation beneath them can change more often. The code becomes more replaceable; the constraints around the code become more valuable.

This will be culturally difficult for engineering organizations whose status has long been attached to owning systems. The new prestige moves toward designing durable boundaries, creating reusable platforms, encoding intent, and retiring obsolete implementation aggressively. A good 2030 architecture may be judged partly by how safely it allows software to die.

4. Buy outcomes, not SaaS seats

The seat is an artifact of human-centric software economics. Agentic systems consume value through actions, tasks, credits, conversations, tokens, tool calls, and compute. Microsoft Copilot Studio already supports usage-based billing through credits, while Salesforce Agentforce prices important capabilities through action-oriented Flex Credits and consumption models.[17][18]

This creates an opening for a more important procurement shift. Instead of asking only how many people need access to a product, CIOs should increasingly ask what the enterprise is paying to achieve a business result. What does it cost to resolve a service case, reconcile an account, onboard a supplier, remediate a vulnerability, qualify a lead, or release a compliant feature?

Outcome buying is harder than seat buying because the enterprise must define the outcome, measure quality, allocate responsibility for exceptions, and avoid contracts that reward activity rather than value. But that difficulty is exactly why the CIO should not let vendors define the unit of value by default.

Procurement, finance, architecture, and product teams should begin building a common vocabulary for outcome units. The transition will be gradual, and many categories will retain subscription components, but the direction is from access pricing toward work pricing. Enterprises that can measure work economically will negotiate from a stronger position than those that cannot.

5. Run thousands of agents, not hundreds of apps

Applications bundle interface, workflow, logic, permissions, and data access into relatively durable packages. Agents begin to unbundle those functions. A future enterprise will still rely on systems of record, but more work will be initiated, coordinated, and completed by specialized agents operating across them.

The challenge is that agents are not merely another software inventory. They may act, spend, communicate, create data, modify systems, invoke other agents, and adapt behavior. BCG's 2026 work on the enterprise AI control plane argues that agent proliferation is already outrunning fragmented governance and that enterprises need unified identity, policy, visibility, and control.[19] Deloitte's finding that only one in five surveyed organizations has mature autonomous-agent governance reinforces the gap.[2]

CIOs should treat the agent estate as a first-class production environment. Every production agent should have a durable identity, owner, purpose, permission envelope, spend limit, data boundary, evaluation record, telemetry, version history, escalation path, and kill mechanism. The enterprise should know not only what agents exist, but what they are authorized to do and what evidence demonstrates that they remain fit for purpose.

The organizational danger is recreating SaaS sprawl at machine speed. If every function independently creates agents with no shared control plane, the enterprise gets thousands of invisible micro-applications with delegated authority. The right goal is not maximum agent count. It is maximum useful autonomy per unit of governable complexity.

6. Continuously arbitrage compute, models, and talent

Enterprises already arbitrage capital, suppliers, and capacity. They will increasingly arbitrage intelligence. The same outcome may be achievable through a premium model, a smaller model with better orchestration, a human expert, a human-plus-agent team, an external specialist, more compute, or a redesigned process.

Anthropic's 2026 Economic Index provides an early empirical signal that users apply different levels of autonomy and model capability to different kinds of outputs, with more compute associated with more valuable artifacts in its analysis.[4] Cloud model routers already optimize cost and quality at request time.[6] The next step is enterprise-level optimization across more than models.

The CIO should therefore build an intelligence cost curve. For important classes of work, quantify human cost, model cost, orchestration cost, latency, failure cost, rework, supervision, and business value. The objective is not the cheapest execution. It is the best risk-adjusted economics for the outcome.

This imperative requires care in how "talent" is treated. People are not interchangeable compute units. Human capability carries judgment, relationships, accountability, learning, and organizational memory. The point is not to commoditize people; it is to stop treating the current division of labor as fixed. In a mature intelligence economy, the enterprise deliberately decides where scarce human attention produces the highest differentiated value.

7. Compress strategy into execution in near real time

Annual strategy cycles assume a meaningful delay between sensing, deciding, funding, building, and measuring. AI compresses each stage. Signals can be detected faster. Scenarios can be generated continuously. Policy can constrain decisions automatically. Agents can prepare or execute approved changes. Telemetry can reveal results almost immediately.

McKinsey's 2026 Global Tech Agenda finds top technology organizations moving toward closer business-technology strategy integration and product/platform operating models in which decisions happen much faster.[20] Its September 2026 operating-model research makes the broader point that AI value depends on redesigning how decisions and work flow through the organization, not simply adding technology to existing structures.[21]

The CIO opportunity is to build a strategy-to-execution nervous system. Strategic priorities should connect to portfolio decisions, architecture, funding, delivery, operational telemetry, customer outcomes, and financial results. When evidence changes, the organization should be able to adjust without waiting for the next annual planning ritual.

This does not mean delegating strategy to AI. Strategic choices involve values, ambiguity, risk appetite, and accountability that remain leadership responsibilities. The goal is to remove unnecessary latency between an intentional decision and the enterprise mechanisms that carry it out.

8. Turn data, models, compute, and expertise into liquid markets

Most enterprises still allocate scarce resources through meetings, annual budgets, ticket queues, internal politics, and organizational ownership. The result is friction: valuable capabilities exist but are difficult to discover, price, trust, or reuse.

A more liquid enterprise makes capabilities discoverable and composable. A business team seeking an outcome should be able to identify approved data products, models, agents, APIs, compute pools, expert communities, and reusable software components, along with their ownership, quality, availability, policy constraints, and economic cost.

This is the core of Frontier's "enterprise intelligence economy" thesis. The word market does not require that every internal interaction carry a literal transfer price. It means that resources have enough metadata and governance to move toward demand. Scarce capabilities become visible. Reuse becomes measurable. Bottlenecks become economically legible. Supply can respond to demand faster than hierarchy alone allows.

CIOs should begin with capability catalogs that include economic and policy metadata. The goal is to make the enterprise's intelligence assets easier to discover and combine without weakening governance. Liquidity without trust produces chaos; governance without liquidity produces queues.

9. Treat people, data, and trust as the most valuable assets

As generic intelligence and generic code become cheaper, durable differentiation moves toward what competitors cannot simply buy from the same vendor. That includes distinctive people, proprietary data, institutional context, relationships, reputation, and permission to act.

PwC's 2026 AI Performance Study found that leading companies were more likely to use responsible-AI frameworks and cross-functional governance, and that employees in those organizations were more likely to trust AI outputs.[22] PwC's separate trust-by-design work argues that reusable controls and accountability mechanisms can accelerate scaling rather than merely constrain it.[14] The competitive point is important: trust is not only a defensive asset. It determines how much autonomy customers, employees, regulators, and leaders will tolerate.

People belong in this same category. The World Economic Forum's 2030 skills outlook continues to emphasize analytical thinking, leadership, resilience, creativity, curiosity, and other human capabilities alongside AI and technical skills.[11] As machines perform more execution, the value of judgment and accountability becomes more visible, not less.

CIOs should stop treating people programs, data governance, cybersecurity, privacy, responsible AI, and brand trust as separate support functions around a technology core. They are parts of the production system. An enterprise that cannot be trusted with autonomy cannot scale autonomy.

10. Use autonomy to operate, manage, and govern the organization

Automation follows predefined steps. Autonomy pursues an outcome within constraints. That distinction changes the operating model because agents can monitor, investigate, plan, coordinate tools, escalate exceptions, and take bounded action without waiting for a human at every step.

OECD's 2026 research on agentic AI in organizations documents that organizations are already working through practical questions of delegation, oversight, accountability, and control.[5] NIST's AI risk framework emphasizes trustworthiness across the AI lifecycle, while the EU AI Act increasingly makes transparency, human oversight, documentation, and accountability enforceable requirements for relevant systems.[13][12]

The operating principle should be progressive autonomy. Define explicit levels: recommend, prepare, act with approval, act within bounded authority, and act independently. Tie each level to evidence, observability, rollback, financial limits, data access, and named human accountability.

Autonomy is not abdication. The more work systems can perform independently, the more precise the enterprise must become about intent and authority. CIOs who treat autonomy as "turning people off" will create brittle systems. CIOs who treat it as an engineered distribution of decision rights can create enormous leverage.

11. Manage ideas, intent, and policies more than systems

When software can increasingly be generated from natural-language intent, the instruction can become more durable than the implementation. What should the business achieve? What constraints may never be violated? What tradeoffs are acceptable? Who can approve an exception? What evidence is required before an action? What does good look like?

Those questions are becoming technology artifacts.

Policy today often lives in documents, committee minutes, and institutional memory. That is insufficient for an environment where agents can act at machine speed. NIST's lifecycle approach and the EU's emphasis on documentation, traceability, transparency, and oversight both point toward governance that can be operationalized rather than merely declared.[13][12]

The CIO should begin converting critical enterprise intent into explicit, testable policy. Security policy, data policy, financial authority, customer-treatment rules, architecture constraints, risk thresholds, regulatory requirements, and escalation paths should increasingly be machine-readable and enforceable at runtime.

This changes the CIO's management surface. The enterprise may have fewer enduring application decisions and more enduring policy decisions. Systems will change. Intent should survive them.

12. Measure cost per outcome, not cost per employee

Traditional productivity measures were built around labor because labor was the dominant variable input. That becomes misleading when the same result can be produced by changing combinations of people, agents, models, software, and compute.

The economic unit must move upward. What does it cost to resolve a claim, close the books, restore a service, onboard an employee, ship a feature, remediate a vulnerability, qualify a customer, or produce a dollar of margin? Microsoft and Salesforce are already exposing consumption-oriented AI pricing primitives, while BCG's 2026 work on technology value argues that measurement must distinguish different kinds of technology investment and connect them to economic outcomes.[17][18][23]

CIOs should select a small set of important business outcomes and establish fully loaded unit economics now, before agents make the underlying resource mix harder to interpret. Include failure, supervision, risk, rework, and change costs rather than counting only token or license spend.

This measure also improves the human conversation. Cost per employee naturally frames people as the expense to minimize. Cost per outcome creates room to ask a better question: what combination of human and machine capability produces the best result at the right quality and risk?

13. Treat every business process as software-defined

A process described only in a slide deck or employee handbook cannot adapt at machine speed. A software-defined process has explicit state, roles, rules, interfaces, data, controls, exceptions, telemetry, and evidence. It can be inspected, simulated, measured, redesigned, and executed by combinations of people and agents.

McKinsey's work on the agentic organization argues for end-to-end process redesign rather than layering AI on top of legacy workflows.[24] Its 2026 research on product development makes the same point in a specific domain: meaningful gains come from rewiring the whole system and its verification mechanisms, not just deploying tools.[16] DORA's findings reinforce why this matters: the organizational system determines whether AI amplifies strength or dysfunction.[9]

The CIO should therefore elevate process architecture to the same seriousness as application architecture. Choose high-value processes and model them as executable systems. Identify decision points, exception paths, human-accountability boundaries, information dependencies, control points, and outcome metrics.

The mistake is to automate the current process and declare success. Software-defined processes are valuable because they can be redesigned continuously. The objective is not faster bureaucracy. It is less bureaucracy.

14. Run IT as a market, not an org chart

This is the synthesis of the other thirteen imperatives.

Traditional IT allocates resources through hierarchy: budgets, teams, projects, standards boards, procurement cycles, and queues. Those mechanisms were rational when technology was expensive, capacity was relatively fixed, and software changed slowly. They become less effective when intelligence, software, and compute are elastic and when agents can act across organizational boundaries.

A market allocates through signals. In the enterprise context, the CIO sets the market rules: approved suppliers, identity, interoperability, architecture boundaries, data access, quality thresholds, risk limits, pricing visibility, observability, and accountability. Demand can then pull the appropriate mix of capabilities toward the outcome.

McKinsey's 2026 technology research describes top-performing organizations moving toward product/platform models and stronger strategic integration between technology and the business.[20] BCG's control-plane work adds the governance layer required as autonomous agents proliferate.[19] Frontier's extrapolation is that the endpoint is not merely a faster IT department. It is an enterprise in which technology allocation becomes increasingly dynamic, governed, and outcome-driven.

The IT organization does not disappear. Its highest-value responsibilities become more important: architecture, platforms, policy, data, security, economics, trust, talent, and market design. The CIO stops being the executive through whom every technology decision must pass and becomes the executive who makes good technology decisions possible everywhere.

The human paradox: more automation makes the human layer more consequential

The easiest reading of this report is that people matter less. That is the wrong conclusion.

People become less valuable for some forms of repetitive execution at the same time that they become more valuable for direction, accountability, judgment, originality, relationships, negotiation, empathy, leadership, and trust. The ILO's conclusion that most exposed work is more likely to be transformed than eliminated and the World Economic Forum's continued emphasis on human skills support this more nuanced view.[10][11]

The P&G field experiment is instructive because AI did not simply substitute for teamwork. It changed the performance of individuals and teams and altered how expertise crossed functional boundaries.[15] That is a redesign signal. If AI changes what one person or one small team can accomplish, it changes the optimal shape of teams, the role of managers, the value of specialists, and the career path by which people learn.

This creates a serious leadership obligation. CIOs cannot present AI primarily as a labor-removal program and then expect employees to enthusiastically redesign work around it. If every productivity gain is translated into an immediate threat to employment, employees learn to hide gains, protect tasks, resist instrumentation, and distrust transformation.

A more durable social contract is explicit. Use AI aggressively to remove low-value work, expand human capability, increase capacity, and redesign roles. Be equally explicit about where work will genuinely disappear. Invest in mobility, apprenticeship, training, and transition. Do not promise that nobody will be affected, and do not treat uncertainty as an excuse to say nothing.

Change management is not the communications package around an AI deployment. It is the design of the new social contract for work.

The technology curve is outrunning the organizational curve

This is the most important practical finding in the report.

Technology capability can change in months. Organizations change through identity, incentives, decision rights, budgets, skills, power, process, trust, and management behavior. Those move more slowly. Microsoft's 2026 data showing substantial misalignment between individual AI capability and organizational readiness is one expression of the problem.[1] BCG's transformation work makes the same point from a different angle, arguing that the majority of AI transformation effort belongs in people, process, and operating-model change rather than the algorithm itself.[25]

McKinsey's September 2026 operating-model research is even more direct: value flows when organizations redesign the way work and decisions are structured around the technology.[21] Its recent work on coordination costs shows why individual productivity is insufficient; the value can disappear in the handoffs between teams, functions, and systems.[26]

The CIO therefore has to run two transformations at different speeds. The technology transformation increases capability. The organizational transformation increases absorption capacity. The second is likely to determine the realized value of the first.

The CIO must change personally

The CIO who treats AI as another platform selection will misunderstand the shift. The role itself has to adapt.

First, the CIO must become an active user of agents rather than merely their executive sponsor. Leaders cannot intelligently redesign work they have never personally delegated. They need firsthand experience of what models do well, where they fail, how supervision changes, and how quickly capability moves.

Second, the CIO must become comfortable managing probabilistic systems. Traditional enterprise technology management rewards determinism. AI introduces confidence thresholds, evaluation, variability, model drift, bounded autonomy, human escalation, and continuous learning. The goal is not to make probabilistic systems pretend to be deterministic. It is to engineer acceptable behavior around uncertainty.

Third, the CIO must become economically fluent in intelligence. Model cost, compute cost, human cost, agent cost, failure cost, rework, latency, risk, and business value increasingly belong in the same decision. FinOps becomes part of a broader discipline: Intelligence Economics.

Fourth, the CIO must become a policy designer. When execution accelerates, policy is one of the few mechanisms that can scale intent. The future CIO should spend more time asking whether a decision boundary is explicit and enforceable, and less time asking who owns the server.

Fifth, the CIO must become a talent architect. The question is no longer only what skills to hire. It is what combination of people and machines should own an outcome, how human capability grows instead of atrophies, and how managers learn to lead hybrid teams.

Sixth, the CIO must become a trust broker. Customers, employees, regulators, boards, and partners need to know where authority resides when AI acts. Trust must be engineered through evidence, transparency, control, and accountability rather than asserted through policy statements.

Finally, the CIO must give up being the center of technology decision-making. That may be the hardest change. The future CIO creates a system in which many good technology decisions happen automatically or locally because architecture, policy, economics, and governance are sound.

The organizational operating model must change with the technology

There are five organizational systems CIOs should redesign in parallel.

Work design must move from roles and tasks toward outcomes, human-agent boundaries, and explicit accountability. Management must evolve from supervising people who execute to supervising systems of people and agents that deliver. Economics must move from projects, licenses, and headcount toward consumption, capability, and outcome unit costs. Governance must move from periodic review toward embedded policy and runtime evidence. Incentives must reward redesign, reuse, delegation, learning, and measurable value rather than preservation of existing territory.

The hard part is that these systems interact. A company can deploy agents but fail because managers are rewarded for team size. It can create a model router but fail because procurement requires a three-year commitment to one supplier. It can make processes software-defined but fail because policy is still trapped in documents. It can encourage experimentation but punish the first visible failure.

This is why AI transformation is fundamentally an operating-model program. BCG's 10/20/70 framing emphasizes that people and process dominate the effort required to scale AI, and its 2026 workforce work finds future-built companies investing much more aggressively in AI learning and organizational capability.[25] The specific ratios will vary by organization, but the direction is right: most of the hard work is not in model selection.

A practical CIO agenda

The next 90 days

Establish the baseline. Inventory material AI and agent activity, including unsanctioned activity, and identify who owns each production deployment. Select ten important business outcomes and establish current unit economics for them. Define an enterprise agent identity and ownership standard. Choose three end-to-end processes for genuine redesign rather than incremental automation. Define progressive autonomy levels and the human accountability attached to each. Create a cross-functional operating coalition that includes technology, finance, people, security, legal/risk, and major business leaders.

At the same time, make the leadership team use agents on real work. Executive AI literacy cannot remain conceptual. The organization will take cues from whether leaders are personally changing how they work or merely asking everybody else to do so.

The next 12 months

Build the common infrastructure for scale. Establish an enterprise control plane for agents, model routing, policy, identity, evaluation, observability, and spend. Create a catalog of approved models, agents, data products, APIs, reusable capabilities, and expert pools. Move at least one material workflow from application-centric design to agent-centric orchestration. Begin shifting suitable vendor agreements from seats toward usage, actions, transactions, or outcomes.

Redesign the human system in parallel. Establish mobility paths for roles that are materially changing. Train managers to supervise hybrid work. Change performance measures so experimentation, reuse, judgment, and measurable value are rewarded. Encode critical policy into machine-enforceable controls where feasible. Measure trust and adoption alongside ROI.

The next 24 to 36 months

Make the new operating model normal. Allow approved workloads to route dynamically across models and compute providers. Expand the software portfolio's use of generated and replaceable implementation atop durable enterprise foundations. Operate a large agent estate with automated lifecycle, evaluation, security, cost, and policy enforcement. Expand internal markets for reusable intelligence assets. Move a meaningful share of technology funding toward outcome-based allocation.

At this stage the enterprise should begin to feel different. Strategy reaches execution faster. Capabilities are easier to discover and reuse. Autonomy expands within explicit boundaries. Technology becomes less about owning systems and more about orchestrating intelligence.

What not to do

Do not turn agents into another SaaS-sprawl problem. Do not confuse agent count with value. Do not automate broken processes without redesigning them. Do not hardwire every important workload permanently to one model provider. Do not use AI savings as the only economic objective. Do not let business-led agent creation become invisible shadow infrastructure. Do not create governance that requires a committee meeting for every autonomous action. Do not centralize AI expertise into a permanent priesthood. Do not tell employees that people are the most important asset while designing transformation exclusively around headcount removal.

Most of all, do not preserve the current IT operating model merely because AI can be bolted onto it.

The metrics that will matter

Traditional CIO metrics will not disappear. Uptime, security, delivery performance, service quality, resilience, cloud spend, and technical debt remain essential. But they will be insufficient.

The 2030 dashboard should also measure cost per business outcome; time from strategic decision to operational execution; human and agent contribution by outcome; percentage of agent actions executed within policy; intervention and exception rates; cost and quality variance across models; reuse of data, agents, and other intelligence assets; time to retire obsolete implementation; percentage of critical policy enforced at runtime; employee AI capability and mobility; trust in autonomous systems; and value produced per unit of intelligence consumed.

These measures describe a fundamentally different enterprise. They treat intelligence as a production resource, autonomy as an engineered property, and trust as part of operating capacity.

The real CIO in 2030

The biggest mistake is imagining the CIO of 2030 as today's CIO with better AI.

The role changes because the objects being managed change. When software becomes easier to create, owning software matters less. When intelligence becomes purchasable by the unit, allocating intelligence matters more. When agents execute work, identity and authority matter more. When implementation changes constantly, architecture and policy matter more. When machine capability expands, distinctly human judgment matters more. When autonomous systems operate continuously, trust becomes productive capacity. When technology becomes inseparable from every business process, the boundary of "IT" stops describing the job.

The CIO becomes the architect of how the enterprise converts human and machine intelligence into outcomes.

The CIO stops running tech.

They transform intelligence into value faster than their competition.

Research basis and method

This report synthesizes observable technology, labor, management, software-engineering, cloud-economics, governance, and regulatory signals available through September 26, 2026, then extrapolates them toward a 2030 operating horizon. It is not a survey of Frontier clients, a forecast with claimed probability, or an assertion that all enterprises will converge on the same operating model.

The evidence base prioritizes public institutions, universities, primary research, technology-provider telemetry and documentation, and management consultancies. Competing technology analyst firms are intentionally excluded. Vendor evidence is used primarily for observable product, pricing, and usage signals rather than as independent validation of vendor claims.

Important limitations apply. Agent capability remains uneven. Benchmark performance does not equal whole-job automation. Vendor telemetry reflects each provider's user population and product design. Consultancy research is subject to sample and respondent effects. Labor-market exposure is not the same as displacement. Regulation and pricing will continue to change. The fourteen imperatives should therefore be read as directional operating hypotheses that help CIOs make better decisions under uncertainty.

SOURCES / CITATIONS

Trace the argument.

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

01 **Microsoft 2026 Work Trend Index Annual Report**
https://www.microsoft.com/en-us/worklab/work-trend-index/agents-human-agency-and-the-opportunity-for-every-organization

02 **Deloitte — The State of AI in the Enterprise 2026**
https://www.deloitte.com/us/en/what-we-do/capabilities/applied-artificial-intelligence/content/state-of-ai-in-the-enterprise.html

03 **METR — Measuring AI Ability to Complete Long Software Tasks**
https://metr.org/blog/2025-03-19-measuring-ai-ability-to-complete-long-tasks/

04 **Anthropic Economic Index: Cadences**
https://www.anthropic.com/research/economic-index-june-2026-report

05 **OECD — Agentic AI in organisations: Early insights from practitioner interviews**
https://www.oecd.org/en/publications/agentic-ai-in-organisations_1257a26f-en.html

06 **AWS — Intelligent prompt routing in Amazon Bedrock**
https://docs.aws.amazon.com/bedrock/latest/userguide/prompt-routing.html

07 **Google Cloud — Spot VMs**
https://docs.cloud.google.com/compute/docs/instances/spot

08 **Anthropic Economic Index: AI's impact on software development**
https://www.anthropic.com/research/impact-software-development

09 **DORA — State of AI-assisted Software Development 2025**
https://dora.dev/research/2025/dora-report/

10 **ILO — Generative AI and Jobs: A Refined Global Index of Occupational Exposure**
https://www.ilo.org/publications/generative-ai-and-jobs-refined-global-index-occupational-exposure

11 **World Economic Forum — Future of Jobs Report 2025, Skills Outlook**
https://www.weforum.org/publications/the-future-of-jobs-report-2025/in-full/3-skills-outlook/

12 **European Commission — The enforcement framework of the AI Act**
https://digital-strategy.ec.europa.eu/en/policies/enforcement-ai-act

13 **NIST AI RMF: Generative Artificial Intelligence Profile**
https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-generative-artificial-intelligence

14 **PwC — How trust by design accelerates AI value**
https://www.pwc.com/gx/en/issues/technology/trust-by-design.html

15 **The Cybernetic Teammate: A Field Experiment on Generative AI Reshaping Teamwork and Expertise**
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5188231

16 **McKinsey — Beyond the copilot: Scaling the agentic product development life cycle**
https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/beyond-the-copilot-scaling-the-agentic-product-development-life-cycle

17 **Microsoft Copilot Studio plans and pricing**

https://www.microsoft.com/en-us/copilot/pricing/copilot-studio

18 **Salesforce Agentforce pricing**

https://www.salesforce.com/agentforce/pricing/

19 **BCG — Enterprise AI Control Plane: The CIO's Guide to Governing and Accelerating AI Agents**

https://www.bcg.com/publications/2026/how-cios-govern-ai-agents-at-scale

20 **McKinsey Global Tech Agenda 2026**

https://www.mckinsey.com/capabilities/mckinsey-technology/our-insights/mckinsey-global-tech-agenda-2026

21 **McKinsey — The key to AI value is hiding in plain sight: Your operating model**

https://www.mckinsey.com/capabilities/people-and-organization/our-insights/the-key-to-ai-value-is-hiding-in-plain-sight-your-operating-model

22 **PwC — 2026 AI Performance Study**

https://www.pwc.com/gx/en/news-room/press-releases/2026/pwc-2026-ai-performance-study.html

23 **BCG — How CIOs Can Prove the Value of Tech in the Age of AI**

https://www.bcg.com/publications/2026/how-cios-can-prove-the-value-of-tech-in-the-age-of-ai

24 **McKinsey — The agentic organization**

https://www.mckinsey.com/capabilities/people-and-organization/our-insights/the-agentic-organization-contours-of-the-next-paradigm-for-the-ai-era

25 **BCG — AI Transformation Is a Workforce Transformation**

https://www.bcg.com/publications/2026/ai-transformation-is-a-workforce-transformation

26 **McKinsey — Cutting the coordination tax: How agentic AI can reshape workflows**

https://www.mckinsey.com/industries/industrials/our-insights/cutting-the-coordination-tax-how-agentic-ai-can-reshape-workflows

FRONTIER.

TAKE THE NEXT STEP

Push this further.

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

RESEARCH / BRIEFINGS

inquiries@thisisfrontier.ai

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

REPRINTS / LICENSING

reprints@thisisfrontier.ai

Reprints, enterprise distribution, excerpts, and licensing.

FRONTIER

www.thisisfrontier.ai

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

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

SEE FURTHER. MOVE SOONER.