

Cognizant: Governing the AI Builder at Scale

Internal ID: S&F-STR-156

Topic: Risk Management, Capital Allocation, Platform Strategy

Industry: IT Consulting & Other Services

Source: Strategy & Finance

Publication date: 5 October 2026

License: Not specified

A strategy case for MBA students and working professionals

Information cutoff: October 2, 2026, 4:00 p.m. Eastern Time

Company: Cognizant Technology Solutions Corporation (NASDAQ: CTSI)

Before the next cohort

In July 2026, Cognizant said it would scale to 5,000 Frontier Certified Engineers and 10,000 Frontier Business Operators. Its first cohort was expected to be assessed and deployment-ready in the fourth quarter. The company described the roles as part of an AI Builder model that combines industry and technology expertise with client knowledge. It also described a centrally governed Ace Team of experienced engineers and a model in which Frontier talent works within client operations. These are announced structures and targets; public disclosures do not show a completed cohort, settled decision-rights charter, or measured client outcomes. [S1, Executive Summary, printed p. 25; S2, paragraphs 76-99; S3, paragraphs 158-161]

For this classroom exercise, assume a **hypothetical Cognizant management group** is setting operating rules after Q2 2026 and before the expected Q4 cohort. The public record does not establish an actual pending meeting, board vote, or policy choice. The group must decide what to govern centrally, what to delegate to practices and client teams, and what evidence should be required before the model expands.

The choice is consequential because AI Builder is intended to sit inside a global delivery business. Cognizant describes an operating model that combines industry and technology capabilities with client knowledge across client teams and global and regional delivery centers. That model depends on both repeatable engineering and understanding how work is actually performed for each client. The central question is how to make those capabilities reinforce one another without making either the platform or the client account the sole owner of change. [S1, "Executive Summary," printed p. 25]

Two centers of gravity

The July announcement described Frontier Certified Engineers as specialists in AI-enabled engineering and Frontier Business Operators as people who combine process understanding with AI tools. Cognizant said it would use a talent funnel, structured learning, assessment, and deployment readiness. Its June role announcement identified SkillsSpring learning paths. These disclosures describe intended roles and preparation, not enrollment, certification validity, time to proficiency, placement, utilization, retention, or worker outcomes. [S2, paragraphs 76-99; S4, paragraphs 75-86]

The Q2 results release separately described the Ace Team as a centrally governed engineering community. Central engineering could provide common architecture, approved tools, evaluation methods, security and privacy controls, model-risk escalation, and a pool of scarce specialists. Shared components might be reused across accounts and reduce inconsistent practices. But a central group may not know which task is valuable in a particular client workflow, may prioritize reusable work over local fit, or may become an approval queue. [S3, paragraphs 158-161]

Client and practice teams have a different advantage. They know process context, contractual commitments, user needs, and the conditions under which a redesigned workflow will be adopted. They can select tasks, embed Frontier roles, and own service quality after deployment. Yet fully local authority could fragment specialists, duplicate tools, produce uneven control standards, and leave no one responsible for reusable assets or enterprise-wide risk.

These are not mutually exclusive capabilities. The design question is where decision rights sit. Who approves client data and model use? Who can stop deployment? Who chooses and prioritizes use cases? Who assigns scarce engineers? Who owns human review, client commitments, and service failures? Who pays for reusable assets, and who receives credit when they are reused?

Investment, integration, and proof

Cognizant's Project Leap is intended to fund integrated offerings, AI capabilities and partnerships, reshape productivity, and upskill workers. In its June 30 Form 10-Q, the company estimated total program costs of \$230-\$320 million and 2026 savings of \$200-\$300 million, primarily to fund investments. It reported \$84 million of Q2 costs-\$56 million for employee separations and \$28 million of other costs. The savings and costs are management estimates, not realized AI-specific funding or demonstrated returns. The filing does not disclose how much was allocated to AI Builder or how savings were measured. [S1, MD&A, "Executive Summary," printed pp. 25-26; Note 4, "Restructuring Charges," printed p. 14]

The company is also integrating recent acquisitions. It acquired 3Cloud on January 1 and Astreya on June 22, 2026. The 10-Q reported preliminary purchase prices of \$728 million and \$665 million, respectively; the Astreya amount included \$25 million of contingent consideration tied to business outcomes. Cognizant's announced rationale for Astreya emphasized AI infrastructure and operations, hyperscaler relationships, AI OpsHub, and complementarity with AI Builder. That is management's strategic thesis, not evidence of realized integration, retention, cross-selling, or synergies. [S1, Note 3, "Business Combinations," printed pp. 12-13; S5, paragraphs 76-91]

Financial capacity is relevant but does not resolve the design. At June 30, Cognizant reported \$1.051 billion of cash, cash equivalents, and short-term investments, \$1.0 billion drawn on its revolver, and \$850 million of remaining capacity. During the first half it generated \$832 million in operating cash flow, repurchased \$1.6 billion of shares, and spent \$1.3 billion on acquisitions. Q2 bookings declined 6% year over year, while trailing-twelve-month bookings increased 5% to \$29.1 billion. These period-specific figures show competing uses of capital and demand context; they do not establish an AI budget or expected returns. [S1, "Liquidity and Capital Resources," printed p. 41; S3, CFO remarks and capital deployment, paragraphs 103-107]

The company's public disclosures do not provide task-level baselines, error and rework rates, client adoption, time saved, realized price or margin changes, or how value is divided between Cognizant and clients. Hours saved do not automatically become provider margin or client savings: the result depends on redeployment, pricing, contract terms, quality, run costs, and demand. A credible expansion rule therefore needs measurement as well as an organization chart.

What the evidence suggests-and leaves open

Studies of professional services give reasons to consider task design and organizational interfaces. A comparative study of four UK law and accountancy firms found that AI systems can automate repetitive processing and augment professional inference; speed and quality gains were observed, while cost benefits were less uniform and could be offset by data preparation, training, or licensing. The study is qualitative and concerns earlier AI systems, not Cognizant or a GenAI productivity coefficient. [S6, Methods §3; Table 2; Figure 2; findings §§4.2-4.4; discussion §§5.2-5.5]

A multiple-case study of three Australian audit firms, based on 15 informants and documents, describes adoption as shaped by internal resources, technology constraints, readiness, regulation, competition, champions, and risk or compliance roles. It provides possible mechanisms for linking central standards with local adoption, not proof that a federated model will outperform at Cognizant. [S7, Methods §§3-3.3, pp. 3-6; Table 1; findings §4, pp. 6-9; Table 3 and discussion §5, pp. 9-12]

A 2026 conceptual research agenda raises a further question: if AI removes junior workers' formative tasks or reduces scrutiny of generated work, what happens to learning, expertise, accountability, and professional trust? It also allows for new tasks and demand that could expand work. These are hypotheses to examine, not forecasts of Cognizant's headcount or productivity. [S8, accepted manuscript, discussion and research agenda pp. 13-23, especially learning pp. 15-16 and workforce/accountability pp. 17-20; Table 2]

The sources support no single organizational answer. Cognizant would need to test the effects of authority, controls, and learning in its own work. A useful scorecard could track client outcomes and service quality; cycle time, error and rework; adoption and repeat use; implementation and run costs; reuse across accounts; security or model incidents; human review and escalation; certification-to-placement conversion, learning and retention; acquisition integration milestones; and the division of value between client and provider. Baselines, denominators, time windows, and accountable owners would need to be specified before comparison.

Three designs for discussion

1. **Central platform and control.** Put platform architecture, approved data and model environments, security, evaluation, specialist engineering, and risk escalation under a central mandate. Practices propose use cases and provide client context. This may increase consistency and reuse; it could also slow work and weaken local accountability.
2. **Practice-led delivery.** Give practices and client teams authority to choose tasks, redesign workflows, embed Frontier roles, and own client outcomes, with minimum enterprise safeguards. This may increase fit and adoption; it could also duplicate tools, fragment scarce talent, and create inconsistent controls.
3. **Federated decision rights.** Centrally fund and govern shared foundations, specialist standards, and risk limits; delegate task choice, workflow redesign, client deployment, and operational accountability to practices within those limits. Specify service levels, escalation, reuse/IP rules, and a scale gate. This may combine consistency and context, but can add coordination costs or blur accountability if interfaces are vague.

Any design must also decide how Frontier workers learn through delivery, how senior reviewers remain accountable for AI-assisted work, and how the acquired capabilities of 3Cloud and Astreya connect to the model. Organic development may build adaptable talent and

institutional knowledge; acquisitions may accelerate specialized capabilities and relationships. Integration capacity, employee and customer retention, overlap, and total cost remain uncertain in the public record.

Assignment

Prepare a two-page decision memo for the hypothetical management group. Recommend a governance design for the next AI Builder cohort and state:

- which decisions and capabilities belong centrally and which belong to practices/client teams;
- how use cases enter, receive safeguards, and earn permission to scale;
- who owns client outcomes, model or service failures, reusable assets, and scarce talent;
- how learning, apprenticeship, human review, and acquired capabilities fit the design; and
- a compact scorecard with baselines, owners, and evidence that would change your recommendation.

You may combine the three designs. Clearly distinguish reported facts, management targets or estimates, and your assumptions. Do not infer a company-specific productivity, margin, or employment effect from the studies. Do not treat Project Leap's estimated savings as realized or AI-specific, or assume that announced training targets have been met.

Source note

This case uses information public by October 2, 2026, 4:00 p.m. ET. Sources below identify the evidence and exact locators. Company statements about targets, expected savings, strategic rationale, and anticipated benefits remain attributed claims.

- **S1** Cognizant Q2 2026 Form 10-Q: [SEC filing](#), Executive Summary pp. 25-26; Note 3 pp. 12-13; Note 4 p. 14; Liquidity and Capital Resources p. 41.
- **S2** Cognizant, "Cognizant to scale to 5,000 Frontier Certified Engineers and 10,000 Frontier Business Operators," July 9, 2026: [issuer release](#), paragraphs 76-99.
- **S3** Cognizant, "Cognizant Reports Second Quarter 2026 Results," July 29, 2026: [issuer release](#), paragraphs 103-107 and 158-161.
- **S4** Cognizant, "Cognizant Develops Frontier Certified Engineer and Frontier Business Operator Roles," June 1, 2026: [issuer release](#), paragraphs 75-86.
- **S5** Cognizant, "Cognizant to Acquire Astreya," April 29, 2026: [issuer release](#), paragraphs 76-91.
- **S6** Spring, Faulconbridge, and Sarwar (2022), *Journal of Operations Management* 68(6-7), 592-618: [DOI/full text](#), Methods §3; Table 2; Figure 2; §§4.2-4.4, 5.2-5.5.
- **S7** Yang, Blount, and Amrollahi (2024), *Technological Forecasting and Social Change* 201, 123251: [DOI · full-text PDF](#), Methods §§3-3.3, pp. 3-6; Table 1; §4, pp. 6-9; Table 3 and §5, pp. 9-12.
- **S8** Faulconbridge et al. (2026), *Journal of Professions and Organization* 13(2), joaf018: [published article · accepted manuscript](#), discussion/research agenda pp. 13-23; Table 2.

The research dossier contains the full source register, evidence extraction, and literature review. No post-cutoff result is used.