



EXECUTIVE BUSINESS CASE

Beyond the AI Hype

Turning Artificial Intelligence into Measurable Enterprise Value

Gabriel Guerrero Paredes

Technology transformation, operations, data and analytics, and enterprise AI

A Harvard-style case package for MBA programs, executive education, board discussion, and enterprise AI transformation.

Based on a real enterprise transformation. Case participants are anonymized; names and selected details have been changed to preserve confidentiality.

June 2026

Author Note and Confidentiality Statement

"The cases presented in this document are based on real-life situations adapted and anonymized for educational purposes. While names and certain details have been modified to preserve confidentiality, the leadership challenges, decisions, and lessons reflect actual events."

This case is based on a real enterprise initiative. All case participants are anonymized. Names, selected roles, workforce actions, project references, email addresses, and other traceable details have been changed, removed, or adapted to preserve confidentiality. The organization is described as a diversified automotive and mobility group with several billion dollars in annual revenue, approximately 5,000 employees, and more than 200 operating locations.

Financial estimates in this document are management scenarios, not audited savings. "Released capacity" and "payroll-equivalent value" become realized cash savings only when the organization eliminates or avoids payroll, overtime, contractor expense, or planned hiring. Platform-generated time-saved estimates are treated as directional indicators and are not substituted for process-level validation.

Contents

Section	Purpose
I. Executive Open Letter	A CIO perspective on AI, productivity, and accountability
II. Board Executive Summary	Situation, action, results, risks, and next decision
III. Main Case	Beyond the AI Hype
IV. Exhibits	Adoption, economics, governance, and operating model
V. Teaching Note	Learning objectives, frameworks, and discussion guide
VI. Author Profile and Advisory Proposition	Enterprise AI transformation offering
VII. Web Publication Abstract	Condensed external version

I. EXECUTIVE OPEN LETTER

AI Does Not Create Value Because It Is Deployed

A year ago, I initiated an enterprise AI program without a board mandate, a CEO request, or a vendor-led strategy. I did it because the operating model required it. The organization had too many repetitive tasks, financial pressure demanded more output from a smaller workforce, and employees were spending time on manual reconciliation, research, analysis, reporting, document review, and communication that did not require their highest judgment.

The common response would have been to select one platform, purchase licenses, publish an AI policy, and call the initiative a transformation. I believed that sequence was wrong. Most organizations do not have an AI problem. They have a business problem, a process problem, a data problem, or a management problem - and they use AI as a slogan before deciding what outcome they need.

Core thesis: *The real value of AI is not simply headcount reduction. It is the release of organizational capacity and the disciplined conversion of that capacity into better analysis, faster execution, avoided hiring, and lower operating cost.*

Our approach began with work, not technology. We trained employees by function, mapped the tasks they performed, quantified the hours involved, selected the model suited to the activity, established rules for the data and actions involved, and then measured results. ChatGPT, Claude, Gemini, and Copilot were not competing products in a procurement exercise; they were different capabilities assigned to different business problems.

Education came before scale. Approximately 1,000 employees attended mandatory, function-specific sessions lasting between 90 minutes and two hours. Senior executives received a separate program focused on risk, output quality, accountability, and the expectation that the right use cases should produce materially faster deliverables. Around 300 to 400 employees began using AI for general productivity, while approximately 40 advanced users moved into process optimization.

The most important rule was direct: AI may produce the work, but the professional remains accountable for the result. A plausible answer is not necessarily a true answer. A technically correct command is not necessarily a safe action. A user should not execute code, terminal commands, bulk communications, or system changes that the user cannot understand, validate, or reverse.

That principle became real when a finance team used AI-generated terminal commands to distribute a large group of emails. The workflow worked. No damage occurred. The problem was more fundamental: the users did not understand the mechanism they had executed. The same trust that enabled productivity could later allow an irreversible action. We responded by restricting terminal and Python recommendations for nontechnical users and by treating financial transactions, customer data, bulk communications, production changes, and executable code as high-risk use cases.

From Experimentation to Operating Capability

The first measured results were modest in absolute pesos but significant in design. Two finance processes released approximately 586 hours per year and produced an estimated MXN 59,500 (USD 22,891 on a U.S. labor-cost basis) in annual process-specific capacity value. One eight-hour task was reduced to approximately six minutes. A second activity that had required six employees for three to four hours for 4 days each was reduced to approximately 32 minutes for one employee. Employees built and operated the workflows with limited technology support and without external consultants.

The point was not that two processes transformed the company. The point was that employees had begun learning how to transform their own work. The organization was developing distributed capability rather than dependence on a central technology team or consulting partner.

The current scenario model for 200 advanced licenses estimates a base opportunity of approximately 117,200 hours per year, equivalent to 61.0 full-time employees, and a best-case opportunity of approximately 234,400 hours, equivalent to 122.1 employees. Using an illustrative workforce mix and a 30% employer-cost factor, the net annual payroll-equivalent value ranges from approximately MXN 24.1 million to MXN 49.2 million (USD 4,516,411 to USD 9,094,536 on a U.S.

labor-cost basis). These are not yet cash savings; they are a disciplined estimate of the capacity available to be converted into avoided hiring, cost reduction, backlog removal, or additional productive output.

The Next Frontier

Commercial AI tools remain the correct environment for drafting, brainstorming, general analysis, public research, creative production, and software development by qualified teams. But repetitive, high-value, data-dependent workflows require something more controlled. The next decision is whether to build internal AI products for financial reconciliation, control reviews, operations analysis, contract review, and recurring reporting.

An internal interface could embed secure system prompts, limit the data available to each department, hide technical complexity, enforce approvals, and create a trusted company experience. It could also create development cost, maintenance obligations, and a new layer of technology that must be governed. The likely answer is not pure build or pure buy, but a hybrid model in which commercial models provide intelligence while the enterprise owns the workflow, controls, context, and user experience.

My position: *AI transformation is not a software deployment. It is an operating-model redesign that combines business problems, process knowledge, people, platforms, policy, and measurable performance.*

This case is an open letter to executives, boards, and transformation leaders. The message is simple: begin with a real problem, design accountability before autonomy, and measure value in the language of the business. AI can accelerate almost any organization. Whether it improves that organization depends on leadership.

II. BOARD EXECUTIVE SUMMARY

Board Executive Summary

Dimension	Executive assessment
Business pressure	Repetitive manual work, financial pressure, and the need to increase output while operating with a smaller workforce and avoiding future hiring.
Starting point	No formal enterprise AI capability; limited executive understanding; inconsistent employee use of public tools; traditional automation had not addressed variable, exception-heavy tasks.
Leadership action	The CIO independently initiated the program, designed training, platform selection, governance, use-case methodology, ROI models, analytics integrations, and measurement dashboards.
Adoption	Approximately 1,000 employees trained; 300-400 general-purpose users; approximately 40 advanced optimization users; 28 advanced licenses in service with expansion planning toward 200.
Operating principle	Select AI according to the business need. Standardize governance, accountability, and measurement - not necessarily the model.
Validated results	Two finance workflows released approximately 586 hours annually and an estimated MXN 59,500 (USD 22,891 on a U.S. labor-cost basis) in process-specific capacity value.
Scenario opportunity	For 200 advanced licenses: base case of 117,200 hours / 61.0 FTE / MXN 24.06M (USD 4.52M U.S.-basis) net annual payroll-equivalent value; best case of 234,400 hours / 122.1 FTE / MXN 49.20M (USD 9.09M U.S.-basis).
Technology department results	120 actual licenses issued to engineering, infrastructure, and support: Level 1 support tickets fell from approximately 1,000 to approximately 100 per month, the support department was reduced by 50% overnight, and faster delivery of partner-facing digital assets contributed approximately MXN 50 million (USD 2,857,143) in new revenue.
Primary risk	Capability without comprehension: users may execute plausible AI recommendations that they cannot validate, reverse, or secure.
Next decision	Continue scaling general-purpose platforms, build internal AI products and connected agents, or adopt a hybrid architecture.

Board-Level Recommendations

1. Adopt a hybrid AI operating model: commercial models for general work; internal products for repetitive, data-dependent, high-value, or high-risk workflows.
2. Require every scaled use case to have a business owner, baseline effort, risk classification, approval model, and value-realization plan.
3. Separate activity metrics from business metrics. Usage proves adoption; hours, quality, cycle time, avoided cost, and revenue impact prove value.
4. Treat released capacity as a management obligation. The board should ask where the hours went and how they changed cost, service, growth, or control.
5. Expand only when governance and implementation capacity can keep pace. License growth without process ownership will increase activity, not transformation.

Decision for the executive team: Should the organization move quickly toward internal AI products and multi-agent workflows, or consolidate the current portfolio until governance, value validation, and product capabilities are more mature?

III. MAIN CASE

Beyond the AI Hype

Turning Artificial Intelligence into Measurable Enterprise Value

Three months after launching the enterprise AI program, the CIO reviewed a set of dashboards. Twenty-four of 28 advanced seats had been active during the prior week. The environment had recorded 466 sessions in 90 days, with usage concentrated in spreadsheets, documents, internal support, design, and structured data. A platform estimate suggested that file operations, conversations, and designs had released roughly 1,000 hours, although the CIO treated that number as directional rather than audited.

The adoption curve was encouraging. The strategic question was more difficult. The organization was considering expansion toward 200 advanced licenses, while teams were already asking for more capable workflows. Finance wanted reconciliation. Operations wanted recurring analysis and data collection. Legal wanted faster contract reviews. Technology teams wanted coding assistance and approved agents. The CIO had to decide whether to continue scaling commercial tools or begin building internal AI products that would place company logic, data access, governance, and approvals behind a simpler interface.

The decision would determine whether the program remained a portfolio of productivity tools or became an enterprise operating capability.

The Company and the Productivity Problem

The organization was a diversified automotive and mobility group with several billion dollars in annual revenue, approximately 5,000 employees, and more than 200 locations. Its operations combined physical assets, customer-facing services, financial controls, fleet activity, legal agreements, digital channels, and distributed operating teams. The scale created substantial administrative work across functions.

The company was also operating through a period of financial pressure. Management needed higher output, faster execution, and lower cost. Some functions had already absorbed workforce reductions. Replacing every manual task with conventional automation was not realistic: many activities involved inconsistent files, changing formats, business judgment, exceptions, and unstructured information. Traditional robotic process automation had addressed stable rules but struggled with work that varied from one cycle to the next.

Employees had begun using AI informally, but there was no consistent method, approved platform model, or shared understanding of risk. Senior executives had limited exposure. Some employees treated AI as a search engine; others assumed that a polished answer was accurate. The organization had no common answer to a basic question: what problem was AI supposed to solve?

A CIO-Initiated Agenda

The program did not originate with the board, the CEO, or a consulting firm. The CIO initiated it after observing the business need. He designed the enterprise training, platform portfolio, governance rules, use-case methodology, ROI model, selected analytical integrations, and adoption dashboards.

His central diagnosis was that most companies started with the desire to use AI and only later searched for a problem. The company would reverse that logic. Teams would first identify work that consumed time or constrained output. The program would then determine whether AI could change the process, what platform was appropriate, what data was required, what controls were necessary, and how value would be measured.




Figure 1. The 6P enterprise AI transformation model: Problem, Process, People, Platform, Policy, and Performance.

Education Before Automation

The CIO began with mandatory, function-specific education. He first trained the technology organization, then finance, operations, legal, fleet, human resources, and other teams. The sessions lasted 90 minutes to two hours and ultimately reached approximately 1,000 employees. Senior executives attended a separate course focused on risks, output quality, accountability, and expected business impact. The CEO, who strongly supported the initiative, was the last senior executive to complete the program.

The content went beyond prompting. Employees learned what AI could and could not do, the differences among tools, business-use examples, privacy risks, hallucinations, validation methods, and the boundary between assistance and automation. The CIO emphasized that the right use cases should produce deliverables four to six times faster, but that speed did not transfer accountability to the model.

Accountability rule: *Employees may use AI to generate, analyze, summarize, compare, and recommend. They remain responsible for the truth, business context, authorization, and consequences of the final output.*

The training also addressed fear. Many employees assumed that AI adoption meant job elimination. The CIO positioned the program as a way to remove repetitive work and strengthen professional capability, while being direct that employees who refused to learn would fall behind colleagues who combined domain expertise with AI.

A Fit-for-Purpose Platform Portfolio

Rather than select one model for every function, the CIO created a fit-for-purpose portfolio. Legal used ChatGPT Team for controlled document work. Finance, operations, IT support, and coding teams used Claude and Claude Code. Marketing used Gemini for images and video. Microsoft Copilot supported general productivity. The choice reflected the nature of the task, data sensitivity, user skill, integration requirements, and the quality of the output required.

Business need	Primary platform	Rationale
Legal analysis and controlled drafting	ChatGPT Team	Enterprise controls and strong general document interaction.
Finance and operational analysis	Claude / Cowork	File-intensive analysis, spreadsheets, structured work, and reusable skills.
Software development and IT support	Claude Code	Coding assistance and technical workflows under engineering controls.
Marketing creative production	Gemini	Image and video generation suited to creative workflows.
General productivity	Microsoft Copilot	Broad access within the productivity suite for routine employee work.

The multi-platform model increased governance complexity, but the CIO believed the alternative - forcing every function into one product - would reduce value and encourage shadow use. Standardization would occur through policy, access, accountability, and measurement rather than through a single model.

Governance Before Scale

Governance was embedded before broad expansion. Employees could use public information and internal nonsensitive documents. Selected tools could access read-only rental, operational, financial, and cost data through approved connections. Customer personal information, employee information, card data, credentials, and privileged information were prohibited. New integrations required approval from the CIO and the security leader.

The program used three operating modes. In the first, AI assisted while the employee performed the action. In the second, AI prepared the output and a person reviewed and approved it before execution. In the third, AI operated

autonomously only within a preapproved, constrained workflow. The risk classification determined which mode was permissible.

Mode	Description	Typical use	Control expectation
Assist	AI drafts or analyzes; the human executes.	Writing, summarization, brainstorming, exploratory analysis.	User validation and accountability.
Supervised automation	AI prepares or changes; a human approves before commit.	Reconciliation, contract review, reporting, bulk communication.	Named approver, traceability, rollback.
Controlled autonomy	AI executes within predefined limits.	Approved agents, scheduled workflows, narrow operational actions.	Technical guardrails, monitoring, exception handling.

The Finance Incident

The program encountered an instructive governance issue when a finance team used AI-generated terminal commands to send between 100 and 200 emails. The commands worked and the messages were distributed correctly. No incident occurred. Yet the users could not explain the commands, identify the systems they touched, or describe how they would reverse an error.

The CIO concluded that technical correctness was not enough. A future prompt could suggest deleting a file, changing a system setting, or executing code that opened a security exposure. The organization introduced guidance that nontechnical users should not receive or execute terminal, shell, or Python recommendations without qualified review. Customer data, financial transactions, bulk communications, production changes, executable code, and database write actions were classified as high risk.

Governance lesson: *The enterprise risk was not malicious AI. It was capability without comprehension: a user could perform an action that was plausible, authorized by habit, and operationally irreversible.*

From Training to Process Redesign

After each training, the CIO scheduled working sessions with the function. Teams identified repetitive tasks, mapped the current process, estimated the number of people and hours involved, selected the appropriate model, developed prompts or workflows, tested the output, and refined the method through three or four sessions. The objective was to transfer ownership to the business user, not create dependency on the CIO or technology team.

Finance focused on spreadsheet validation, cross-referencing, payment documents, and supplier communications. Operations used AI-assisted web research and created simple HTML interfaces that allowed employees to upload a spreadsheet and refresh recurring reports. Legal accelerated contract reviews. Technology teams used AI for coding, IT support, documentation, and design. Across functions, the strongest adoption occurred where business pressure was most immediate.

Initial Results

Two finance workflows provided the first validated process-level evidence. Updating payment receipts fell from eight hours of one employee to approximately six minutes. Sending payment receipts to suppliers fell from six employees spending three to four hours each to one employee spending approximately 32 minutes. Together, the processes released approximately 586 hours per year and produced an estimated MXN 59,500 (USD 22,891 on a U.S. labor-cost basis) in annual process-specific capacity value under the original labor-cost assumptions.

Process	Before	AI-enabled state	Annual hours released	Annual capacity value (MXN)	Annual capacity value (USD)
Payment receipt update	1 employee x 8 hours	1 employee x approximately 6 minutes	95	MXN 9,600	USD 3,711

Process	Before	AI-enabled state	Annual hours released	Annual capacity value (MXN)	Annual capacity value (USD)
Supplier receipt distribution	6 employees x 3-4 hours each	1 employee x approximately 32 minutes	491	MXN 49,900	USD 19,180
Total	-	-	586	MXN 59,500	USD 22,891

Source: Internal before-and-after process measurements. Values represent released capacity under process-specific payroll assumptions.

Other outcomes were visible but not yet fully monetized: reporting moved from hours to minutes; recurring operational reports could be refreshed through a simple interface; external research and data collection accelerated; contract reviews became faster; and analytical outputs improved in depth and speed.

Adoption Metrics: Activity Versus Value

The usage dashboards confirmed sustained engagement. Twenty-four of 28 advanced seats were active in a recent week, and weekly use had increased as licenses expanded. In 90 days, the platform recorded 466 Cowork sessions. File and document activities dominated usage, indicating that employees were applying AI to operational work rather than only conversational experimentation.

The dashboards also showed low daily stickiness for some tools. The CIO did not interpret this automatically as failure. Many finance, legal, and operational activities occurred weekly or monthly, not every day. The more important question was whether each active user had a validated process and whether the hours released were converted into business outcomes.

The Economic Model

The CIO developed base and best-case scenarios for the 200 advanced licenses. The model used the average hours released by the first two validated processes, assumed two optimized processes per license in the base case and four in the best case, and converted the released hours into full-time-equivalent capacity using 1,920 productive hours per employee per year.

To estimate payroll-equivalent value, management used an illustrative workforce mix of 20 analysts, four managers, and one director. Monthly base salaries were modeled at MXN 20,000 (USD 6,250 on a U.S. analyst-benchmark basis) for analysts, MXN 40,000 to MXN 60,000 (USD 12,500 to USD 18,750 on the same basis) for management roles, and a 30% employer cost factor. The midpoint blended loaded cost was approximately MXN 411,840 per employee per year (USD 75,000 under the U.S. standard used for comparison below). Annual license cost was estimated at MXN 151,200 (USD 8,640).

Metric	Base case (MXN)	Base case (USD)	Best case (MXN)	Best case (USD)
Processes optimized annually	400	—	800	—
Annual hours released	117,200	—	234,400	—
FTE-equivalent capacity	61.0	—	122.1	—
Gross payroll-equivalent value	MXN 25.14M	USD 4.58M	MXN 50.28M	USD 9.16M
Annual license cost	MXN 1.08M	USD 0.06M	MXN 1.08M	USD 0.06M
Net annual payroll-equivalent value	MXN 24.06M	USD 4.52M	MXN 49.20M	USD 9.09M
Illustrative ROI	2,228%	7,318%	4,555%	14,737%

Financial discipline: The model does not claim that 61.0 to 122.1 positions will be eliminated. It estimates the labor capacity that could be released. Management must decide

	<i>whether to convert that capacity into avoided hiring, reduced payroll, more analysis, faster service, lower backlog, or additional growth. The USD columns are not currency conversions of the MXN figures; they show what the same released capacity would be worth if valued at a U.S. fully loaded labor cost of USD 75,000 per FTE per year, so the gap between the MXN and USD columns reflects the wage differential, not exchange-rate movement.</i>
--	--

Executive Sponsorship

The CEO supported the initiative and asked the CIO to move quickly. No executive challenged the investment because the use cases were tied to visible business problems. A senior finance executive became the strongest champion after independently studying prompt engineering. His enthusiasm created momentum, but it also illustrated the risk of overconfidence. At one point, he wanted to use persona prompting to advise other departments. The CIO emphasized that a well-structured prompt could not substitute for the professional who understood the operating context.

The program had been presented to the CEO but not yet to the board. The CIO was therefore managing two narratives: an internal adoption program that needed speed, and an external governance and value story that needed evidence strong enough for board scrutiny.

Technology Department Impact: Engineering, Infrastructure, and Support

In addition to the 200-license scenario modeled for business functions, the CIO issued 120 actual licenses to the Technology department's development, cloud infrastructure, and support teams. These teams applied the same platforms to engineering workflows rather than clerical ones: software deployment, website and module development, partner-facing digital assets, design and mockup work, and security code review.

- Deployments: compressed from approximately three hours to about fifteen minutes.
- Website and module builds: compressed from approximately three months to about one week; once tooling matured, partner-facing websites could be built in days.
- Design and mockup presentations: reduced to a few hours once the tools were correctly configured.
- Security code scanning: compressed from approximately one week to about one hour.
- Level 1 support tickets: callable AI agents handled routine requests (connectivity, network access, user registration, password resets), cutting monthly Level 1 ticket volume from approximately 1,000 to approximately 100.

Applying the same modeling approach used for the 200 advanced licenses, management averaged the hours released by two representative technology workflows — a full website/module delivery cycle (approximately 480 hours compressed to approximately 40 hours) and a security code scan (approximately 40 hours compressed to approximately 1 hour) — for an illustrative average of approximately 240 hours released per optimized activity, and again assumed two optimized activities per license annually in the base case and four in the best case.

Metric	Base case (MXN)	Base case (USD)	Best case (MXN)	Best case (USD)
Optimized activities annually	240	—	480	—
Annual hours released	57,600	—	115,200	—
FTE-equivalent capacity	30.0	—	60.0	—
Gross payroll-equivalent value	MXN 12.36M	USD 2.25M	MXN 24.71M	USD 4.50M
Annual license cost	MXN 0.65M	USD 0.04M	MXN 0.65M	USD 0.04M
Net annual payroll-	MXN 11.71M	USD 2.21M	MXN 24.06M	USD 4.46M

Metric	Base case (MXN)	Base case (USD)	Best case (MXN)	Best case (USD)
equivalent value				
Illustrative ROI	1,807%	5,977%	3,713%	12,056%

Source: Illustrative management scenario, using the same modeling logic as the 200-license case (average hours released by representative workflows, 2/4 optimized activities per license annually, 1,920 productive hours per FTE per year). The blended loaded cost and license cost per seat match the assumptions used for the 200-license scenario; the USD columns are the U.S. labor-cost-basis equivalent, not a currency conversion.

	Technology department success: This program — platform selection, governance, training, and the 120-license engineering rollout — was designed, built, and run by the Technology department, without an external mandate or consulting partner.
--	--

Faster delivery of partner-facing websites and digital assets also had a commercial effect. Because the Technology department could stand up partner integrations and digital assets in days rather than months, the company was able to bring new partner-facing developments to market faster and capture revenue that slower delivery would have delayed or lost. Technology-enabled developments contributed approximately MXN 50 million (USD 2,857,143 at the FX rate used throughout this case) in new revenue through partner integrations and digital assets during the period, without naming individual projects.

The support team's result was different in kind from the scenarios above: it was already realized, not projected. As Level 1 ticket volume fell by roughly 90% (approximately 1,000 to approximately 100 per month), the support department was reduced by 50% overnight — a smaller cut than the ticket-volume drop because the remaining Level 3 and external-interconnectivity tickets are individually more complex and time-consuming to resolve, leaving the remaining team focused on those escalations. At an illustrative average handling time of 20 minutes per Level 1 ticket, the volume reduction alone represents approximately 300 hours per month (3,600 hours per year) of support capacity no longer required for repetitive issues.

The Next Phase: Internal AI Products

The CIO was now considering internal AI products for financial reconciliation, control reviews, operations analysis, contract review, and recurring reporting. The applications would use commercial models but provide a company-designed interface. Secure system prompts could limit each function to approved actions. Data access could be restricted to the specific tables required. Technical commands could be hidden. Approval steps could be embedded. Usage and value could be measured by workflow rather than by chat activity.

The proposal would require developers, product owners, AI engineers, cybersecurity specialists, data integration support, and business process owners. It would also introduce a maintenance obligation and the risk of building custom applications around rapidly evolving vendor capabilities.

Strategic option	Advantages	Risks
Scale commercial tools	Fast deployment, low development burden, continuous vendor innovation, broad flexibility.	Fragmented experience, greater training burden, weak process standardization, users may access functions they do not understand.
Build internal AI products	Trusted interface, embedded controls, company context, standardized workflows, precise data access.	Development and maintenance cost, slower delivery, product-management demands, risk of recreating vendor features.
Hybrid model	Commercial intelligence with company-owned workflows, approvals, context, and user experience.	Requires architectural discipline, prioritization, governance capacity, and a capable product team.

The Decision

The CIO had already demonstrated adoption, executive support, governance controls, and initial process results. The company could expand quickly toward 200 advanced licenses and allow functions to continue developing their own

methods. It could pause to validate more processes and strengthen benefit realization. Or it could invest in a hybrid operating model that used commercial models underneath internal, governed AI products and connected agents.

Each choice carried an opportunity cost. Moving slowly could waste momentum and allow uncontrolled usage to return. Moving quickly could create fragmented workflows, inflated savings claims, and new operational risks. Building internal tools could establish a defensible enterprise capability, but it would compete for scarce development and cybersecurity capacity.

Closing question: *What should the CIO scale next: licenses, processes, governance, or internal products - and in what sequence?*

Discussion Questions

1. Was the CIO correct to adopt a multi-platform AI portfolio rather than standardize on a single vendor?
2. How should the company distinguish between released capacity, avoided cost, and realized savings?
3. What governance controls are required when nontechnical users can generate executable actions?
4. Should the company expand toward 200 advanced licenses before validating more process-level value?
5. Which workflows justify an internal AI product, and which should remain in commercial tools?
6. Who should own the next phase: the CIO, a business-led AI center of excellence, individual functions, or a hybrid structure?

IV. EXHIBITS

Exhibit 1 - Enterprise Context and Three-Month Rollout

Dimension	Anonymized case facts
Enterprise	Diversified automotive and mobility group
Scale	Several billion dollars in annual revenue; approximately 5,000 employees; more than 200 locations
Program start	Three months before the case decision
Employees trained	Approximately 1,000
General-purpose users	Approximately 300-400
Advanced optimization users	Approximately 40
Advanced licenses	28 in service; expansion planning toward 200
Functions engaged	Technology, finance, operations, legal, fleet, human resources, marketing, and senior executives
Primary objective	Reduce cost, release capacity, avoid future hiring, accelerate execution, and build AI literacy




Figure 2. Adoption funnel after the initial three-month rollout. General-purpose users shown at midpoint of the reported 300-400 range.

Exhibit 2 - Platform Portfolio and Governance Boundaries

Platform	Primary use	Permitted data / action boundary
ChatGPT Team	Legal and controlled document work	Public and approved nonsensitive internal documents; no personal, card, credential, or privileged data.
Claude / Cowork	Finance, operations, file analysis, reporting	Selected read-only operational and financial data; approved files and connectors.
Claude Code	Software development and technical support	Qualified technical users; code review and engineering controls; nontechnical terminal use restricted.
Gemini	Marketing images and video	Approved creative content; no confidential or personal information.
Microsoft Copilot	General productivity	Routine employee work under corporate access and data policies.

Risk tier	Examples	Required control
Low	Drafting, summarization, brainstorming, public research	User review and accountability.
Moderate	Internal document analysis, spreadsheets, reporting, contract review	Approved platform, business owner, data classification, human approval.
High	Customer or employee data, transactions, bulk communications, executable code, production changes, database writes	CIO and security approval, qualified reviewer, restricted access, logging, rollback, exception handling.

Exhibit 3 - Advanced AI Adoption Dashboard




Figure 3. Point-in-time 90-day dashboard: 24 weekly active members of 28 seats and 466 Cowork sessions.

Source: Enterprise AI platform administration dashboard. Total spend shown in the dashboard reflects platform usage charges and does not replace annual license cost.

Interpretation: weekly active use increased as seats expanded. The dashboard is evidence of adoption, not proof of financial value. Process-level measurement remains necessary.

Exhibit 4 - How Advanced Users Applied AI




Figure 4. Adoption profile, skills, and connectors for the advanced user environment.

Source: Enterprise AI platform administration dashboard. Identifiable user-level information excluded.

Observed activity	90-day count
File operations	204
Conversations	877
Designs created	14
MCP writes	82
Platform-estimated time saved	Approximately 1,000 hours

The time-saved estimate is based on configured minutes-per-task assumptions. It is directional and should be reconciled against validated before-and-after process measures.

Exhibit 5 - Validated Process Results

Process	Before	Now	Annual hours released	Annual process-specific value (MXN)	Annual process-specific value (USD)
Update payment receipts	1 employee, 8 hours	1 employee, approximately 6 minutes	95	MXN 9,600	USD 3,711
Send payment receipts to suppliers	6 employees, 3-4 hours each	1 employee, approximately 32 minutes	491	MXN 49,900	USD 19,180
Total	-	-	586	MXN 59,500	USD 22,891

Source: Internal process measurements. Employees executed the improvements with limited technology supervision and no external implementation partner.

Measurement note: *Realized process result: approximately 586 hours released annually. The MXN value uses process-specific payroll assumptions; the USD figure is the equivalent value if the same hours were released under U.S. labor costs, not a currency conversion. Neither should be combined with the enterprise scenario model.*

Exhibit 6 - Base and Best-Case Economic Scenarios




Figure 5. Illustrative net annual payroll-equivalent value for the 200-license environment.

Metric	Base case (MXN)	Base case (USD)	Best case (MXN)	Best case (USD)
Processes per license per year	2	—	4	—
Total optimized processes	400	—	800	—
Annual hours released	117,200	—	234,400	—
FTE-equivalent capacity	61.04	—	122.08	—
Gross payroll-equivalent value - midpoint	MXN 25,139,400	USD 4,578,125	MXN 50,278,800	USD 9,156,250
Annual license cost	MXN 1,080,000	USD 61,714	MXN 1,080,000	USD 61,714
Net payroll-equivalent value	MXN 24,059,400	USD 4,516,411	MXN 49,198,800	USD 9,094,536

Payroll assumption	MXN	USD
Workforce mix	20 analysts, 4 managers, 1 director	—
Analyst monthly salary	MXN 20,000	USD 6,250 (U.S. benchmark)
Manager monthly salary range	MXN 40,000-60,000	USD 12,500-18,750 (U.S. benchmark)
Employer cost factor	30% planning assumption	—
Productive hours per FTE	1,920 per year	—
Midpoint blended loaded cost	MXN 411,840 per FTE per year	USD 75,000 per FTE per year (U.S. standard)

Source: Management scenario. The 30% factor is a planning assumption for employer contributions and statutory/benefit costs; Finance should replace it with the company-specific loaded payroll factor for formal budgeting.

Exhibit 7 - Next-Phase Decision Framework

Criterion	Scale commercial tools	Build internal products	Hybrid model
Speed	High	Medium to low	Medium
User flexibility	High	Low to medium	High within approved boundaries
Governance embedded in workflow	Medium	High	High
Company-specific context	Medium	High	High
Development burden	Low	High	Medium to high
Vendor dependence	High	Medium	Medium
Best fit	General work and experimentation	Stable, repetitive, strategic workflows	Portfolio approach across the enterprise

Prioritization Criteria for Internal AI Products

1. Financial value and avoided cost potential.
2. Frequency and repeatability of the process.
3. Hours released and number of employees affected.
4. Security, control, and data sensitivity.
5. Ability to scale across functions without recreating the workflow.

V. TEACHING NOTE

Teaching Note: Beyond the AI Hype

This teaching note is designed for MBA courses, executive education, board workshops, and senior leadership programs in digital transformation, technology strategy, operations, governance, and organizational change.

Case Synopsis

A CIO at a large, diversified enterprise independently launches an AI program in response to repetitive work, financial pressure, and the need to operate with fewer resources. The program combines function-specific education, a multi-platform tool portfolio, governance controls, process redesign, and financial measurement. Within three months, adoption is strong and two processes show measurable time savings. The CIO must decide whether to scale commercial tools, invest in internal AI products, or pursue a hybrid operating model.

Learning Objectives

- 1. Distinguish AI adoption from AI transformation and identify the operating-model changes required for value realization.
- 2. Evaluate the strategic trade-offs between single-platform standardization and a fit-for-purpose AI portfolio.
- 3. Design governance for AI-assisted actions, especially when users cannot fully understand executable recommendations.
- 4. Assess financial claims by separating released capacity, payroll-equivalent value, avoided cost, and realized savings.
- 5. Determine when an enterprise should build internal AI products rather than rely exclusively on commercial interfaces.
- 6. Examine the CIO role as an architect of enterprise productivity, not only an owner of technology.

Suggested Assignment Questions

- 1. What was the CIO's most important decision during the first three months, and why?
- 2. Does the multi-platform strategy create more value than the governance complexity it introduces?
- 3. How should the finance terminal-command incident change the control model?
- 4. Are the base and best-case financial scenarios credible? What evidence would be required before the board relied on them?
- 5. Which next-phase option should the CIO recommend to the CEO?
- 6. How should the company convert released capacity into a measurable business outcome?

Analytical Framework 1 - The 6P Model

Element	Key question	Application in the case
Problem	What business constraint must change?	Manual work, cost pressure, limited capacity, slow execution.
Process	What activity, baseline, frequency, and exception pattern exist?	Function-level process mapping before tool selection.
People	Who owns the outcome and can validate it?	Business users trained; advanced users coached until independent.
Platform	Which model and interface fit the task?	ChatGPT, Claude, Gemini, Copilot selected by use case.
Policy	What data, action, approval, and accountability rules apply?	Risk tiers, read-only access, prohibited data, technical review.
Performance	How is value measured and realized?	Hours, cycle time, FTE-equivalent capacity, avoided cost, quality.

Analytical Framework 2 - Value Realization Ladder

Level	Evidence	Management implication
Activity	Licenses, sessions, prompts, files, active users	Confirms access and engagement; does not prove value.
Output	Reports, contracts, analyses, code, communications	Shows work produced faster; quality must be validated.

Level	Evidence	Management implication
Process	Cycle time, effort, error rate, throughput	Demonstrates operational improvement.
Capacity	Hours and FTE-equivalent capacity released	Creates an opportunity; management must redeploy or remove cost.
Financial impact	Avoided hiring, payroll reduction, lower contractors, revenue, working capital	Converts productivity into economic value.

Analysis of the Central Dilemma

The case should not be taught as a simple build-versus-buy decision. The deeper question is sequencing. The company must scale four capabilities together: employee adoption, process redesign, governance, and product engineering. Expanding licenses faster than process ownership creates unused or superficial capacity. Building internal products before use cases stabilize risks automating the wrong workflow. Overengineering governance can slow adoption and push employees back to unapproved tools.

A strong recommendation is a hybrid portfolio governed by thresholds. Commercial tools remain the default for general, low-risk, nonrepeatable work. Internal products are justified when a workflow is high-frequency, economically material, data-dependent, risk-sensitive, and stable enough to standardize. Each internal product should have a business owner, product owner, security review, baseline metrics, and a value-realization plan.

Financial Discussion

Participants should challenge the distinction between the two validated processes and the enterprise scenario. The 586-hour result is based on observed before-and-after work. The base and best cases extrapolate from limited evidence and use a broader workforce cost model. They are useful for investment sizing, but they should not be reported as realized savings.

The board should require management to classify value in four categories: released capacity, avoided cost, realized cash savings, and growth or quality value. It should also ask how many identified processes have reached pilot, validation, and scale. A benefits ledger should track the owner, baseline, implementation cost, realized value, and disposition of the released hours.

Governance Discussion

The finance incident illustrates why user accountability alone is insufficient. The business user remains accountable for the business outcome, but the enterprise is accountable for platform security, permissible functionality, training, access controls, and technical review. The correct model is shared accountability.

Accountability	Primary owner
Business objective and accuracy	Business user and process owner
Output validation and approval	Qualified functional reviewer
Platform security and access	Technology and cybersecurity
Data classification and permitted use	Data owner, legal, security, and business
Executable code and system actions	Qualified technical reviewer
Benefit realization	Functional executive and finance

Recommended 90-Minute Teaching Plan

Time	Discussion
0-10 minutes	Opening vote: scale licenses, consolidate, build, or hybrid?
10-25 minutes	Diagnose the initial business problem and the CIO's sequencing choices.

Time	Discussion
25-40 minutes	Debate multi-platform strategy and governance complexity.
40-55 minutes	Analyze the terminal-command incident and shared accountability.
55-70 minutes	Test the financial model and value classifications.
70-85 minutes	Design the next-phase operating model and product criteria.
85-90 minutes	Re-vote and summarize leadership lessons.

Key Teaching Insights

- AI literacy is a prerequisite, not a transformation outcome.
- Business problems should determine models; governance should unify the portfolio.
- The highest risk may be a correct answer executed by a user who does not understand it.
- Adoption metrics and financial value are different measurement systems.
- Released capacity has no economic value until management decides what to do with it.
- The CIO can use AI to expand from technology leadership into enterprise productivity and operating-model design.

VI. AUTHOR PROFILE AND ADVISORY PROPOSITION

Author Profile

Gabriel Guerrero Paredes, MBA, is a Group Chief Information Officer with executive experience across technology, operations, data and analytics, digital transformation, and enterprise productivity. His credentials include an MBA, PMP certification, Lean Six Sigma Black Belt, prior operations leadership, U.S.-LATAM business experience, and advisory-board participation. He leads technology in a multibillion-dollar, multi-location operating environment.

His executive positioning is that of a technology-transformation leader who converts innovation into measurable operating value, with a trajectory toward broader Chief Digital and AI Officer, COO, and CEO responsibilities.

Enterprise AI Transformation Proposition

The methodology demonstrated in this case can be applied to large Mexican groups, LATAM companies, multinational subsidiaries, and operationally intensive organizations with approximately 3,000 to 10,000 employees.

Phase	Client outcome
AI maturity assessment	Current-state view of adoption, data readiness, risks, skills, and operating capability.
AI strategy and portfolio	Prioritized business problems, platforms, roadmap, investment thesis, and build-buy-hybrid choices.
Governance design	Data rules, risk tiers, accountability, approvals, access, monitoring, and incident response.
Executive and workforce education	Board, C-suite, functional, and advanced-user programs tied to actual work.
Use-case and process transformation	Process mapping, prompt/workflow design, pilots, validation, and ownership transfer.
ROI and value validation	Hours, cycle time, quality, capacity, avoided cost, realized savings, and revenue impact.
Internal AI products and agents	Company-specific interfaces, controlled data access, secure prompts, approvals, and connected workflows.
Scaling operating model	AI center of excellence, product teams, governance cadence, benefits ledger, and continuous improvement.

Differentiation: *Unlike traditional AI consulting, this approach is based on an enterprise implementation led from inside the operating company. It begins with real business problems, embeds governance before scale, selects technology according to the use case, builds internal capability, and measures results in hours, capacity, avoided cost, and financial return.*

This document is intended to support boardroom discussion, executive interviews, advisory conversations, and senior leadership searches in which enterprise AI, digital transformation, productivity, and governance are strategic priorities.

VII. WEB PUBLICATION ABSTRACT

Beyond the AI Hype: A CIO Framework for Enterprise Value

Most companies start their AI journey by buying licenses. We started with the work.

A large, diversified enterprise faced financial pressure, repetitive administrative tasks, and the need to increase output with fewer resources. The CIO launched an AI program without an external mandate. Approximately 1,000 employees received function-specific education, 300 to 400 began using AI for general productivity, and around 40 advanced users moved into process optimization. Different platforms were selected for legal, finance, operations, technology, marketing, and general productivity.

Governance was built before scale. Personal information, customer data, card data, credentials, and privileged information were prohibited. Selected operational and financial data were available through controlled, primarily read-only connections. Users remained accountable for outputs, and high-risk actions required review.

Two early finance processes released approximately 586 hours annually. The 200-license environment now represents an illustrative opportunity of 117,200 to 234,400 hours per year, equivalent to 61.0 to 122.1 full-time employees and approximately MXN 24.1 million to MXN 49.2 million (USD 4,516,411 to USD 9,094,536 on a U.S. labor-cost basis) in net annual payroll-equivalent value. These figures are capacity scenarios, not yet audited cash savings.

The next question is not whether to use AI. It is whether to continue scaling commercial tools or build internal, governed AI products around company workflows and data. The strongest answer is likely a hybrid: commercial models underneath company-owned context, controls, approvals, and user experience.

The 6P rule: *Begin with the problem. Redesign the process. Train the people. Select the platform. Govern the action. Measure the performance.*