

THE CARIBBEAN CANNOT AFFORD TO GET **AI WRONG**

But Can It Afford to Get It Right?

AI Stacking, Enterprise Competitiveness
and Institutional Readiness



PEOPLE
AND SKILLS



INSTITUTIONS
AND GOVERNANCE



PRODUCTIVITY
AND
COMPETITIVENESS

EVIDENCE.
INSIGHT.
PRACTICAL RESPONSES FOR
A STRONGER CARIBBEAN.

JOSEPH COX

Managing Partner
Joseph Cox & Associates Limited

STRONGER INSTITUTIONS
BRIGHTER CARIBBEAN

THE CARIBBEAN CANNOT AFFORD TO GET AI WRONG

But Can It Afford to Get It Right?

CBR Intelligence | October 2026

Joseph Cox

Managing Partner, Joseph Cox & Associates Limited

Publisher, Caribbean Business Review and Caribbean Financial Dispatch

October 21, 2026

Suggested Citation: Cox, J. (2026). *The Caribbean Cannot Afford to Get AI Wrong – But Can It Afford to Get It Right*. CBR Intelligence. Caribbean Business Review.

Keywords:

Artificial Intelligence | Caribbean Economies | AI Stacking | Enterprise Competitiveness | Technology Affordability | Institutional Readiness | Productivity | Digital Transformation | AI Integration | Total Cost of Ownership | Economic Value Realisation | Technology Dependence | Small and Medium-Sized Enterprises (SMEs) | AI Governance | Domestic Value Capture.

BOARD BRIEF

The Caribbean Cannot Afford to Get AI Wrong

But Can It Afford to Get It Right?

Artificial intelligence presents Caribbean enterprises with a consequential strategic choice: how to capture emerging productivity and competitive advantages without assuming financial, technological and institutional commitments that exceed their capacity to generate sustainable economic value.

WHY NOW? THE COST OF WAITING—AND THE RISK OF MOVING TOO QUICKLY

Artificial intelligence is moving from experimental application to an increasingly consequential component of enterprise competitiveness. For Caribbean businesses, the strategic question is no longer simply whether to adopt AI, but whether they possess the financial capacity, institutional architecture and workforce capabilities to convert adoption into sustained productivity gains.

The urgency arises from a convergence of pressures. AI capabilities are expanding rapidly, while subscription-based pricing, overlapping platforms and integration requirements can create recurring expenditure and technological dependencies before measurable economic returns materialise. Smaller enterprises face particular exposure: limited scale, foreign-currency costs and constrained technical capacity can make apparently affordable solutions disproportionately expensive.

Yet postponement carries its own risks. As competitors deploy AI to improve service delivery, reduce processing times and strengthen decision-making, firms that delay may experience a deterioration in their relative competitive position. For Caribbean economies already confronting longstanding productivity constraints, the consequences extend beyond individual enterprises to employment, domestic value creation and the region's positioning within internationally traded services.

The danger is therefore two-sided: premature adoption without institutional readiness can destroy value, while excessive caution can surrender competitive ground. Boards must distinguish technological acquisition from productive capability—and ensure that investment decisions are supported by credible economic analysis, organisational preparedness and measurable accountability.

The Board's Decision

Artificial intelligence is increasingly influencing the conditions of competitive participation, but acquiring more AI tools does not necessarily produce a more competitive enterprise. Directors must establish which capabilities matter, what they cost in full, how benefits will be realised, and whether the organisation can govern the resulting dependencies.

Board proposition. Approve AI investment as a measurable business-capability decision—not as a technology subscription exercise.

Five tests for every material AI proposal

Board test	Question directors should put to management
Competitive exposure	Which cost, service-quality, speed or market-access disadvantage emerges if we do not act?
Total economic cost	Are prerequisite licences, paid seats, implementation, security, training, integration and FX constraints included?
Productive returns	What measurable hours, quality gains, revenue or contribution will be realised—and by whom?
Institutional readiness	Who owns workflow redesign, staff capability, data governance, decision rights and accountability?
Strategic dependence	What is the exit route if vendors change prices, entitlements, access or service conditions?

Three decisions the board should demand

1. Define the capability gap before selecting the tool. Compare bundled, selective-seat, open-source, outsourced and integrated alternatives on equivalent functionality, reliability and safeguards.
2. Require a business case with low, base and high productive-realisation assumptions. Separate labour time released from hours actually converted into economic value; report quality and revenue benefits separately to avoid double counting.
3. Assign one accountable executive and a review date. Require a baseline, a benefits owner, risk limits, procurement controls and a stop/scale decision.

Board Decision Dashboard

Use this dashboard at approval and at the first post-implementation review. A proposal that cannot populate the evidence and ownership columns is not yet investment-ready.

Measure	Minimum evidence required	Accountable owner
Capability and comparator	Target workflow; current baseline; cheapest credible alternative with equivalent output	Business sponsor
Fully loaded first-year cost	Seats + prerequisites + setup + data/security + training + imported services	CFO / CIO
Economic return	Hours × realised value; contribution on new revenue; low/base/high sensitivities	CFO / operating lead
Quality and control	Error rate, first-pass acceptance, human sign-off, privacy and model-risk checks	Risk / compliance
Vendor and FX exposure	Price reset, service availability, switching costs, foreign-currency access	Procurement / treasury
Review and decision gate	30/90/180-day milestones; continue, redesign, scale or stop triggers	CEO / board committee

Worked affordability test: five-person enterprise

Illustrative configuration from the underlying study—not a survey of Caribbean firms. Assumes US\$220 per month in recurring stack costs, US\$1,200 setup, US\$15 per economically realised hour and a 30% contribution margin on incremental revenue.

First-year stacked case	Lean / selective alternative
US\$3,840 total = $12 \times \$220 + \$1,200$	US\$1,220 total = $12 \times \$60 + \$500$
256 economically realised hours to break even	81.3 economically realised hours to break even
US\$12,800 additional revenue at 30% margin	≈US\$4,067 additional revenue at 30% margin

Break-even estimates represent recovery of the illustrative first-year AI expenditure through economically realised labour value. They do not constitute a comprehensive investment appraisal and exclude other benefits, displacement effects and alternative costs.

At 50% productive realisation, the stacked configuration requires 512 gross hours of potential savings to yield 256 economically realised hours. The lean option is not automatically preferable: it must meet the same required quality, capacity, security and service standard.

Decision rule. Approve the lowest total-cost configuration that credibly delivers the required competitive capability and an acceptable risk-adjusted return.

Board Resolution

Management should submit a 90-day AI capability plan identifying:

- (i) two priority workflows;
- (ii) the existing baseline and competitive requirement;
- (iii) at least one credible lean counterfactual;
- (iv) fully loaded costs and three realisation cases;
- (v) security, vendor and FX controls; and
- (vi) named executive accountability, a date for benefits verification, and reporting of realised economic value against the approved business case.

What directors should ask at the next meeting

“What capability are we buying, what alternative have we rejected, what evidence would prove that it is working, and at what point would we stop spending?”

TABLE OF CONTENTS

BOARD BRIEF	3
Five tests for every material AI proposal	4
Three decisions the board should demand.....	4
Worked affordability test: five-person enterprise.....	5
Executive Assessment	7
1 The productivity opportunity and its distribution	8
2 The Economics of AI Stacking.....	8
2.0 Operational workflow test: lead-to-proposal	12
3 Representative AI software price register	17
4 Illustrative enterprise affordability and break-even test.....	19
5 Sensitivity: exchange rates, utilisation & scope	21
6 Three-country evidence: Jamaica, Trinidad and Tobago and Cayman	21
7 Distribution, labour markets and domestic value capture	25
8 Institutional architecture and governance	26
9 A sequenced Caribbean policy agenda	26
10 Analytical scope and interpretation	27
Conclusion	27
Appendix A Pricing sources and validation status	28

Executive Assessment

Global competition rewards enterprises that deliver higher quality, faster response, lower unit costs and sustained innovation. AI expands the capabilities available to competitors, making technological upgrading a strategic concern for firms seeking to scale, protect margins and remain viable. Yet inexpensive entry-level tools do not make competitive capability inexpensive: its full cost includes access, integration, training, verification, governance and continuing use.

The resulting paradox is that AI may lower the marginal cost of individual tasks while raising the organisational demands of competing effectively. Firms need not buy every application: selective adoption, bundled services, outsourcing, open-source tools and shared infrastructure may deliver comparable results. The economic test is whether the required capability can be obtained and deployed at a positive return.

Using a 20-platform price register, named AI architectures, a lead-to-proposal integration model, enterprise-scale cost scenarios and selected Caribbean wage benchmarks, this paper evaluates feasible capabilities and conditional returns—not observed technology purchases by Caribbean firms.

Key findings

- Integrated workflows can generate value, but stacking is not inherently superior: lean, modular and shared-service alternatives may deliver comparable capability at lower cost.
- Labour-cost sensitivity is large. The same US\$2,400 annual software outlay requires roughly 41 realised hours at Cayman professional-services salary benchmarks, versus 135–410 hours at selected Jamaican public-service minima and 119–318 hours at selected Trinidad and Tobago public-service salaries.
- Productive realisation is decisive. Time saved without valuable redeployment, and quality improvements without economic conversion, yield limited financial returns.
- Organisational design shapes returns: clear decision rights, reliable data, accountable process owners and verification capacity support productive absorption.
- Policy should favour stack-selection support, pooled evaluation capacity and incident sharing over indiscriminate subscription subsidies.

1 | The productivity opportunity and its distribution

The IMF estimates that roughly 40% of employment globally is exposed to AI, rising to about 60% in advanced economies. Exposure is not equivalent to displacement: jobs may be augmented, transformed or replaced, depending on task composition and institutional choices. The IMF has also described a wide range of possible long-run effects on potential growth, underscoring uncertainty rather than a guaranteed productivity dividend [1–3].

Caribbean economies enter this transition with persistent constraints involving enterprise scale, financing, digital infrastructure, skills, electricity costs, public-sector execution and fragmented markets. AI could alleviate selected bottlenecks—for example by improving research, customer service, forecasting, documentation and quality control. Equally, it could automate inefficient processes without correcting the incentives or accountability failures that produced them.

Three stages must be distinguished: technological availability, productive absorption and domestic value capture. An enterprise may have access to AI without integrating it into a productive process; it may improve its own performance without generating a proportionate financial return; and it may earn a private return without producing commensurate domestic value added after imported software, compute and other inputs. The cost and competitiveness tests therefore examine more than subscriptions.

Analytical propositions. H1 (competitive capability): as AI improves the cost, speed or quality of market offerings, firms face pressure to achieve economically comparable productive capabilities, whether through direct adoption, outsourcing or shared services. H2 (cost paradox): the cost of isolated AI tasks can fall while the cumulative cost of integrated competitive capability remains substantial. H3 (conversion and scale): net returns depend on task fit, utilisation, enterprise scale and organisational architecture. These are conditional propositions tested through transparent scenarios and counterfactual configurations, not claims about observed Caribbean adoption.

CORE ASSUMPTIONS AND ROBUSTNESS | All monetary inputs are stylised US-dollar first-year scenarios, not observed Caribbean MSME expenditure. The base labour-value proxy is US\$15 per economically realised hour; a 30% contribution margin is used only when translating a cost into required incremental sales. A 12-month horizon is assumed; first-year cost equals $12 \times$ recurring monthly charge plus one-off setup. Break-even hours = cost $\div$ US\$15; sales-equivalent break-even = cost $\div$ 0.30. The integration example processes 120 proposal cases/year, saves 240 gross hours at full adoption, and assumes first-pass acceptance rises from 90% to 96%; these are test inputs, not measured outcomes. Sensitivities below vary realisation (25%–100%), cost and exchange rates, and the wage-value proxy. Cheaper alternatives are comparable only where functionality, quality, security and service continuity are adequate.

2 | The Economics of AI Stacking

The emerging AI divide may be less about who has access to artificial intelligence and more about who can afford, integrate and sustain an effective AI ecosystem.

A typical knowledge-intensive workflow may involve an LLM for analysis, a research service for retrieval, a design platform for production, a CRM for distribution, and automation software for hand-offs and reporting. No single tool is necessarily indispensable. But complementary functions often matter more than isolated demonstrations of capability.

Two deployment architectures

The term AI software stack describes two related but economically distinct deployment architectures. The enterprise-builder AI technology stack supports custom AI capabilities through foundation models and hosting, data and knowledge infrastructure, orchestration and agent workflows, application interfaces, and cross-cutting evaluation, security and observability. The functional AI software stack combines ready-to-use applications across research, administration, customer service, marketing, finance, human resources, legal services and analytics. The categories overlap: a firm may use purchased applications alongside custom-built agents or integrations.

Enterprise-builder architecture. The layers include (i) foundation models and inference infrastructure (for example, managed model APIs or locally hosted models); (ii) data preparation, storage and retrieval; (iii) orchestration, agent execution and workflow control; (iv) user-facing applications; and (v) evaluation, observability, cybersecurity and governance spanning the entire system. Illustrative technologies include cloud model services, vector databases, LangGraph or LlamaIndex, and monitoring tools. These are examples of functions and possible suppliers, not a prescribed procurement bundle.

Functional architecture. Enterprises may combine general-purpose assistants and office copilots with CRM and support agents, creative production, meeting intelligence, automation, finance, legal or analytics applications. The relevant unit of analysis is the workflow and its quality-adjusted output, not the number of products purchased. A solo creator, a five-person enterprise and a large regulated institution need not acquire the same functionality, seat count or governance infrastructure.

Table 1. Enterprise-builder AI technology stack: functional layers and illustrative technologies

Layer	Function	Illustrative technologies
Foundation models	Base language and multimodal models	OpenAI GPT models; Anthropic Claude; Google Gemini; Meta Llama
Model infrastructure	API access, inference and hosting	AWS Bedrock; Azure OpenAI / Microsoft Foundry; Google Vertex AI; Anyscale
Data and knowledge	Storage, ETL, retrieval and vector search	Snowflake; Databricks; Pinecone; Qdrant; Chroma; Weaviate
Orchestration and agents	RAG, workflow control, tool use and routing	LangChain / LangGraph; LlamaIndex; Temporal; AutoGen
Application interface	End-user experiences and interfaces	Vercel AI SDK; CopilotKit; Streamlit
Evaluation and observability	Quality, latency, cost, safety and monitoring	LangSmith; Weights & Biases; Arize; Phoenix; Helicone

Table 2. Functional AI software stacks by business activity (illustrative options, not cumulative purchase requirements)

Business function	Illustrative products
Executive productivity and operations	Microsoft 365 Copilot; Google Workspace Gemini; Notion AI; ClickUp Brain; Zapier; Make; UiPath; Fireflies.ai; Otter.ai; Granola
Sales, marketing and customer support	Salesforce Agentforce; HubSpot AI; Intercom Fin; Zendesk AI; Decagon; Jasper; Copy.ai; Anyword; Canva Magic Studio; Synthesia; Runway; Midjourney
Finance, HR and legal	QuickBooks AI; Brex AI; Vic.ai; Salesforce Agentforce for HR; Slack AI; Paradox; Harvey; CoCounsel; Ironclad
Software development and analytics	GitHub Copilot; Cursor; Claude Code; Tableau Pulse; Power BI Copilot; ThoughtSpot

Table 3. Concrete combinations by enterprise scale: illustrative deployment architectures

Scenario	Combined technologies	Economic interpretation
Lean / small business	ChatGPT Plus or Claude + Canva AI + Zapier + Notion AI + Fireflies.ai	Cross-functional subscriptions; verify overlaps, free/bundled entitlements and actual paid seats.
Mid-market	Microsoft 365 Copilot + HubSpot AI + Intercom Fin + GitHub Copilot + Snowflake	Mixed functional and data stack; cost depends on seat counts, CRM/support usage and data consumption.
Enterprise / builder-hybrid	AWS Bedrock or Azure OpenAI + Databricks + LangGraph + Salesforce Agentforce + custom retrieval-augmented generation (RAG) pipelines	Cloud/model consumption, data infrastructure, engineering, governance and agent-service charges; not comparable with a single-seat subscription.

Methodological note: These are illustrative combinations submitted for research, not recommendations to buy every product or evidence that each is necessary. Prices in the existing 20-platform register cannot simply be summed across these examples: licence terms, included features, seats, API consumption, service tiers and integration work differ. Product names and entitlements should be rechecked before publication. A stack is economically justified only where its incremental, quality-adjusted benefits exceed its full incremental costs.

Economic consequences of architecture choice

Functional stacks typically expose firms to recurring per-user, per-feature and usage-based charges; builder stacks may require more substantial engineering, compute, data-management and evaluation expenditure. Neither architecture is invariably more costly. The efficient choice depends on transaction volume, task

complexity, regulatory exposure, reuse of infrastructure, staff capabilities and the availability of bundled, free or open-source alternatives. A hybrid architecture may minimise total cost for some firms.

The study therefore distinguishes cost per seat, cost per transaction, cost per integrated workflow and cost per quality-adjusted unit of output. It will test whether the incremental gross value of integration exceeds the additional engineering, coordination, verification and maintenance costs. At the regional level, the ability to amortise fixed costs across small markets and to preserve strategic autonomy becomes part of the competitiveness question.

Builder-stack cost schedule (separate from the functional price register)

The builder-stack cost schedule distinguishes model inference/API usage and compute hosting; data preparation, storage and retrieval; orchestration and workflow engineering; interface development; evaluation, observability and security; technical staffing and training; and recurring maintenance, vendor migration and exit provisions. Published tariffs establish priced components, while explicitly stated scenarios capture implementation and operating requirements not available as standard tariffs. These inputs support comparative affordability tests without implying observed enterprise expenditure.

Price-register scope and validation

Table 8 is a register of 20 representative functional-platform offers, not a full enterprise-builder bill or an additive shopping list. Monthly figures are per-user charges on monthly billing; annual figures are the separate, prepaid annual-plan charges supplied following a per-user pricing check. Differences between twelve monthly invoices and an annual prepayment are not arithmetic discrepancies. Plan eligibility, limits, included features and underlying licence requirements should be recorded when applying the register to a specific configuration.

AI may democratise access to individual technologies while simultaneously creating new barriers to enterprise competitiveness through the cumulative costs of integration, technological sophistication and scale. Individual business tasks may become cheaper and faster, even as the investment required to compete effectively across an entire enterprise increases.

This distinction is fundamental. AI expenditure is not inherently an economic burden. It becomes one when the cost of acquiring, integrating, operating and maintaining the technology exceeds the economic value an enterprise can capture from it.

The economics of AI adoption must therefore be examined beyond individual subscription prices. Four interconnected cost categories are particularly relevant.

1. Access costs. Subscriptions, licences, usage credits and application programming interface (API) charges across multiple applications. These costs may increase with the number of users, the volume of activity and the sophistication of the services required.

Adding paid seats across functional AI tools without accounting for prerequisite software licences, such as qualifying Microsoft 365 subscriptions required for Microsoft 365 Copilot, can materially understate the true cost of deployment.

2. Integration costs. Connecting platforms, transferring and preparing data, establishing workflows, maintaining interoperability and ensuring that different technologies operate effectively within existing business systems.

3. **Capability costs.** Developing the workforce skills, managerial competencies, organisational processes and governance arrangements necessary to exploit the technologies productively. These investments may require changes in organisational design, incentives, decision rights and accountability mechanisms.

4. **Dependency costs.** Exposure to changing prices, service conditions, vendor restrictions, technological obsolescence and the expense of switching platforms. Dependence may also arise where enterprises lack the internal capabilities to maintain, modify or replace critical technological systems.

These costs interact. A relatively inexpensive application may generate substantial value when deployed within an enterprise possessing the necessary skills, data infrastructure and organisational capacity. The same application may produce limited returns in an enterprise that must first undertake extensive complementary investments.

Conversely, free tiers, open-source models, existing-licence AI features and shared-service arrangements can reduce financial entry costs. Their implementation, oversight, support and sustained productive use may nevertheless entail material expenditure. The scenarios below therefore describe alternative architectures rather than a necessary progression toward a costly stack.

The AI integration premium is the additional economic value created by coordinating applications into an end-to-end workflow, compared with using the same functions independently. It must be evaluated net of incremental integration expenditure. The following calculation operationalises the concept using transparent scenarios rather than unavailable firm-level data.

Operational measure: gross integration premium (GIP) = $V(\text{integrated}) - V(\text{stand-alone})$. Incremental integration cost (ΔC) = $C(\text{integrated}) - C(\text{stand-alone})$. Net integration premium (NIP) = $GIP - \Delta C$. The integration decision breaks even when $GIP = \Delta C$.

2.0 | Operational workflow test: lead-to-proposal

A stylised professional-services firm processes 120 comparable lead-to-proposal cases annually (10 per month). Both configurations use the same underlying research, drafting, CRM and document-production functions, handle the same lead mix, and are assessed at the same utilisation, service specification and 12-month horizon. The stand-alone configuration requires manual transfer and reconciliation between applications; the coordinated stack adds automated hand-offs, common client records and controlled review checkpoints. No change in lead volume or selling price is assumed. All time, quality and implementation values below are scenario assumptions, not observed Caribbean or US enterprise performance.

Table 4. Lead-to-proposal workflow: common utilisation and quality assumptions

Workflow stage	Isolated tools: minutes/lead	Coordinated stack: minutes/lead	Mechanism of change
Lead intake and qualification	45	30	Shared CRM fields; reduced duplicate entry
Research and client context	75	45	Reusable retrieval and source hand-off
Proposal drafting and costing	90	55	Structured prompts and reusable templates
Human quality/compliance review	60	35	Consistent checks; human approval retained
Formatting, delivery and CRM update	30	15	Automated assembly, dispatch and logging
Total per lead	300 (5.0 h)	180 (3.0 h)	120 minutes saved per lead
Annual throughput	120 leads	120 leads	Same utilisation and task mix
Annual labour input	600 h	360 h	240 hours released at full realisation
First-pass quality acceptance (assumed)	90%	96%	108 vs 115.2 expected accepted proposals

The quality metric is the share of proposals meeting a predefined first-pass specification for factual accuracy, completeness, commercial consistency and compliance. The assumed 90% and 96% acceptance rates are illustrative, not measured improvements. Expected accepted proposals (108 and 115.2) are statistical expectations rather than fractional deliverables. Review time is included in both configurations; quality gains are not assigned a monetary value in the base case.

Table 5. Lead-to-proposal: first-year economic integration premium (US\$; 120 leads/year)

Measure	Isolated tools	Coordinated stack	Increment / comparison
Annual technology and implementation cost	\$2,460	\$3,840	+\$1,380
Annual labour hours	600	360	-240 h
Labour opportunity cost at \$15/h	\$9,000	\$5,400	-\$3,600
Total modelled economic cost	\$11,460	\$9,240	-\$2,220
Expected first-pass accepted proposals	108	115.2	+7.2 expected
Economic cost per accepted proposal	\$106.11	\$80.21	-\$25.90
Gross integration value: 240 h × \$15	—	—	+\$3,600
Net integration premium: \$3,600 – \$1,380	—	—	+\$2,220

Interpretation. Holding annual lead volume constant, coordination reduces modelled handling time from five to three hours per lead. If all 240 released hours have a genuine \$15/hour economic use—through redeployment, avoided expenditure or incremental contribution—gross value is \$3,600. After the \$1,380 incremental first-year stack cost, the net integration premium is +\$2,220. The premium is an incremental comparison; the \$2,460 stand-alone expenditure is already in the baseline and must not be counted again. Cost per accepted proposal includes the same assumed labour valuation and software/implementation cost in each configuration; the quality advantage is shown as a separate output metric, not added again as revenue.

Realisation sensitivity. If only 50% of the 240 nominal hours are economically captured, $NIP = (120 \times \$15) - \$1,380 = +\$420$. At 25% realisation, $NIP = (60 \times \$15) - \$1,380 = -\$480$. The break-even threshold is 92 economically realised hours, equivalent to 38.3% of the 240 nominal hours. Thus faster processing alone does not establish a positive premium: management must convert released capacity into productive value. A different assumed first-pass acceptance rate would change quality-adjusted unit costs, but not the base NIP unless its financial value is separately evidenced or explicitly modelled.

Comparator discipline. The \$2,460 isolated-tools baseline is the cost of delivering the same 120-case workflow using the same functional components without automated hand-offs; the \$3,840 integrated option adds coordination. By contrast, the \$1,220 five-person lean configuration in the enterprise affordability table is a minimum-scope alternative with different assumed service and implementation inputs. It is not the controlled comparator for this workflow. It should be tested as a third option only after demonstrating comparable task coverage, throughput, confidentiality, reliability and first-pass quality. The positive incremental integration premium establishes a return relative to the specified isolated baseline, not that the integrated configuration is the cheapest way to obtain competitive capability.

To prevent double counting, V includes only incremental, economically captured value: monetisable labour capacity, avoided expenditure, incremental contribution from revenue, or quality-adjusted output, net of error-correction and verification losses. Unused time savings and gross sales are not automatically counted as economic benefits. Compare alternatives at the same task mix, time horizon and scale; where quality differs, report it explicitly and value it only once.

Quality-value sensitivity (separate from the base integration premium). At 120 proposals, moving from 90% to 96% first-pass acceptance produces 7.2 additional expected first-pass acceptances. If each avoids \$30 of independently documented rework cost, the additional annual value is \$216; if each instead generates \$100 of incremental contribution from genuinely additional successful conversions, the additional value is \$720. These are alternative illustrations, not cumulative benefits. They must be excluded if the 240 released hours already incorporate the same rework reduction, or if acceptance only measures quality compliance and does not increase final sales. The base \$2,220 net premium therefore remains unchanged absent separate evidence of monetisable quality effects.

The first-year premium is not necessarily the steady-state premium: one-off implementation costs should be amortised over a declared useful life for a multi-year comparison; ongoing maintenance, usage and switching costs must still be included. A negative premium is a valid result and supports selective or stand-alone adoption instead of integration.

The resulting economic question is not whether AI makes individual tasks cheaper. It is whether the total investment required to achieve and sustain competitive performance remains proportionate to the value enterprises can generate.

Could AI widen the Caribbean's competitiveness gap even as it makes individual business tasks cheaper and faster?

For Caribbean economies, this question extends beyond individual enterprises. If the benefits of AI accrue disproportionately to organisations with greater financial resources, stronger institutional capabilities, superior data and larger markets, technological diffusion could reinforce existing structural inequalities rather than overcome them.

The regional policy question is whether shared capabilities, credible institutions and firm-level integration can convert access to AI into competitively priced, domestically captured output.

2.1 | From named combinations to measurable economic exposure

Tables 1–3 identify available technological configurations, not observed purchasing patterns. Table 6 separates published vendor charges from explicitly modelled implementation, governance and operational costs. Because seats, tokens, compute hours and service tiers use different pricing units, the analysis evaluates defined configurations under transparent utilisation assumptions rather than treating advertised prices as a complete measure of productive AI capability.

Table 6. Stack combinations and cost mechanisms: a measurement schedule

Configuration and components	Cost drivers to measure	Counterfactual / economic test
Lean functional: ChatGPT or Claude + Canva + Zapier + Notion + Fireflies	Paid versus free/bundled seats; automation runs; premium design features; staff time for verification; cross-platform hand-offs	Compare with existing office software, free tiers and manual workflow; value only verified output improvements
Mid-market functional/data: Microsoft 365 Copilot + HubSpot + Intercom + GitHub Copilot + Snowflake	Eligible Microsoft base licence; CRM/contact tiers; support resolution volume; developer seats; Snowflake compute/storage; integration and data stewardship	Compare modular adoption by department with all-at-once integration; test marginal contribution of each component
Enterprise builder/hybrid: AWS Bedrock or Azure OpenAI + Databricks + LangGraph + Salesforce + custom RAG	Tokens/inference and hosting; ingestion and vector retrieval; engineering; evaluation; security; observability; CRM licences; maintenance and vendor switching	Compare custom deployment with purchased software, managed services and regional shared infrastructure

These are combinations, not prescriptions: a firm may omit components, substitute providers or purchase equivalent functionality through existing bundles. Named services differ in their pricing units. A seat licence, token charge, compute hour and CRM contact tier must not be added as though they were interchangeable units.

2.2 | Worked comparative affordability test: assumptions, not observed expenditure

Table 9 is the single enterprise-level comparison: three functional configurations and a custom-builder/hybrid scenario, each paired with a lean, open-source or shared-service alternative. Table 7 does not repeat those rows; it stress-tests only the custom-builder configuration under lower and higher infrastructure and implementation assumptions. All amounts are transparent scenarios, not observed regional expenditure or vendor quotations.

Table 7. Custom-builder sensitivity only: low, base, high and managed/open-source alternatives (US\$)

Custom-builder scenario	Recurring / month	One-off setup	Year-one total	Hours at \$15 realised value	Revenue at 30% margin	Interpretation
Low	\$500	\$6,000	\$12,000	800	\$40,000	Lower usage and limited engineering
Base	\$1,000	\$12,000	\$24,000	1,600	\$80,000	Same builder baseline as Table 9
High	\$2,000	\$25,000	\$49,000	3,266.7	\$163,333	Higher inference and engineering load
Lean managed/open-source	\$300	\$4,000	\$7,600	506.7	\$25,333	Credible only with equivalent service standards

Comparison assumptions (Tables 7 and 9): identical 12-month horizon, US\$15 per economically realised hour and 30% contribution margin. First-year cost = 12 × monthly recurring charges + one-off setup; break-even hours = cost / 15; required incremental revenue = cost / 0.30. Lower-cost options qualify as competitive counterfactuals only if they meet the same required functionality, throughput, reliability, privacy and quality standards. These are alternative technology pathways, not observed purchases or verified market minima.

Builder stress-test decomposition: US\$1,000/month is assumed to comprise US\$400 for inference/compute, US\$250 for data/storage/retrieval, US\$150 for monitoring/evaluation and US\$200 for security and operations; the US\$12,000 one-off amount represents assumed engineering, integration and validation. The breakdown is an arithmetic illustration only. It excludes financing, taxes, foreign-exchange conversion, hardware, connectivity and ongoing staff opportunity costs not already priced into the inputs. Table 7 varies both recurring and one-off cost inputs, not observed usage rates.

For any stack, first-year total cost $C_1 = 12R + F$, where R is monthly recurring cost and F is initial implementation expenditure. If a realised hour yields v dollars of incremental economic value, break-even realised hours $H^* = C_1/v$. If incremental revenue has contribution margin m , revenue required to cover the cost is C_1/m . At $v = \text{US\$15}$ and $m = 0.30$, the illustrative custom-builder scenario requires 1,600 economically realised hours or US\$80,000 in incremental revenue. These are alternative ways of expressing the threshold, not benefits that can be added together without accounting for overlap.

A more demanding test uses the incremental cost and benefit of integration: $\Delta NPV = \sum_t [(\text{incremental contribution and verifiable cost savings} - \text{incremental operating costs})/(1+r)^t] - \text{incremental upfront investment}$. The comparator must be the best feasible lower-cost configuration, including free, bundled, modular or shared-service alternatives—not an artificial no-AI baseline. A positive gross productivity effect does not establish a positive net return.

2.3 | Scale, utilisation and falsification

The scale hypothesis is conditional, not automatic. Fixed setup expenditure per unit of output falls with higher throughput only if quality, compliance, latency and demand are held sufficiently constant. Usage-based inference, human verification and governance expenditure can rise with volume. Accordingly, measure cost per completed quality-adjusted workflow, not merely subscriptions per employee. Test whether pooled services, open-source components or lower-cost modular adoption overturn the apparent scale disadvantage.

A credible empirical result could reject the paper's strongest interpretation: if smaller firms achieve comparable quality-adjusted output at equal or lower full cost using selective tools, or if integration costs decline rapidly enough to offset added stack complexity, the proposed rising competitiveness threshold is not supported for those firms or tasks. The study must report that outcome rather than assume stacking necessarily raises costs.

3 | Representative AI software price register

Table 8 presents the per-user pricing register supplied for this assessment. Monthly charges are monthly-billed prices; annual charges are separately quoted prepaid annual prices, not twelve times the monthly charge. These are alternative purchase terms. The independently specified enterprise budgets in Table 9 and workflow costs in Tables 4–5 are modelling inputs, not sums of the Table 8 products. Applicable plan entitlements, seat minimums and billing terms remain relevant to any actual procurement.

Table 8. Selected AI platform charges per user: monthly billing versus prepaid annual plan (USD)

Platform	Function	Monthly billing (USD / user)	Prepaid annual plan (USD / user)*
ChatGPT Plus	Research and general assistance	\$20.00	\$240.00
Claude Pro	Analysis and writing	\$20.00	\$200.00
Perplexity Pro	Research and search	\$20.00	\$204.00
Microsoft 365 Copilot Business	Office productivity	\$25.20	\$302.40
Google Workspace Business Standard	Productivity and Gemini	\$14.00	\$168.00
Notion Business	Knowledge management	\$20.00	\$240.00
Grammarly Pro	Writing	\$30.00	\$360.00
Jasper Pro	Marketing content	\$69.00	\$708.00
HubSpot Starter	CRM and marketing	\$20.00	\$240.00
Zapier Professional	Automation	\$29.99	\$239.88
Make Core	Automation	\$12.00	\$144.00
ClickUp Brain AI	Project management	\$9.00	\$108.00
Otter.ai Pro	Meeting notes	\$16.99	\$99.96
Fireflies.ai Pro	Meeting notes	\$18.00	\$120.00
Descript Creator	Audio and video	\$35.00	\$288.00
ElevenLabs Creator	Voice generation	\$22.00	\$220.00
Adobe Firefly Pro	Generative imagery	\$19.99	\$168.00
Synthesia Starter	Video	\$29.00	\$216.00
Semrush AI Visibility Base	Search visibility	\$99.00	\$1,188.00
Canva Pro	Visual design	\$18.00	\$144.00

Vendor prices represent the identified offers reviewed for the October 2026 assessment, with actual procurement requiring confirmation of current prices, entitlements and applicable billing conditions.

Price interpretation matters. Annual prepaid plans may involve a lower effective monthly rate than month-to-month billing. For example, Zapier Professional is shown at \$29.99 monthly or \$239.88 for an annual prepayment. Microsoft Copilot may require an underlying Microsoft 365 licence; HubSpot and automation platforms may impose usage or credit limits. The prices are per-user where applicable, and must not be summed without matching seat counts, tier entitlements and workload requirements.

Low-cost baseline and alternative pathways. A lean pathway may combine free AI tiers, existing productivity-suite entitlements and selectively deployed open-source software. It can avoid multiple subscriptions, but free quotas, privacy restrictions, support needs, hosting, model evaluation and staff time can still create costs. A selective pathway purchases a small number of paid tools for demonstrably valuable tasks; an enterprise-wide pathway adds identity controls, security, workflow integration and ongoing governance. The study should compare these pathways on total cost and quality-adjusted economic value, not presume that a costly stack is inevitable.

4 | Illustrative enterprise affordability and break-even test

Table 9 consolidates the enterprise affordability scenarios; Table 7 separately tests the custom-builder sensitivity range. Each configuration uses a 12-month horizon and specified first-year setup costs. These scenario totals are not constructed by summing the prepaid annual prices in Table 8. Comparisons are conditional on equivalent quality and productive capability; taxes, FX spreads, hardware and connectivity are excluded from the stated inputs.

Table 9. Unified enterprise-scale affordability comparison: stacked and lean configurations (US\$)

Enterprise	Staff	Stack / mo	Setup	Year-one cost	Hours \$15/h	@ Revenue @ 30% margin	Lean / open-source / shared-service alternative (first-year assumptions and thresholds)
Sole proprietor	1	\$95	\$300	\$1,440	96.0	\$4,800	Free/bundled + 1 paid tool: \$20/mo; \$120 setup; \$360 total; 24 h; \$1,200 revenue
Five-person firm	5	\$220	\$1,200	\$3,840	256.0	\$12,800	Selective seats/shared automation: \$60/mo; \$500 setup; \$1,220 total; 81.3 h; \$4,067 revenue
Twenty-person firm	20	\$650	\$5,000	\$12,800	853.3	\$42,667	Selective licences/managed shared service: \$180/mo; \$1,800 setup; \$3,960 total; 264 h; \$13,200 revenue
Custom builder / hybrid	Variable	\$1,000	\$12,000	\$24,000	1,600	\$80,000	Managed API + open-source + outsourced support: \$300/mo; \$4,000 setup; \$7,600 total; 506.7 h; \$25,333 revenue

Table 10. Realisation-rate robustness across three firm sizes (gross annual hours required at US\$15 per realised hour)

Firm size	First-year stacked cost	Realisation rate	Gross hours required
Sole proprietor	\$1,440	25%	384.0
Sole proprietor	\$1,440	50%	192.0
Sole proprietor	\$1,440	75%	128.0
Sole proprietor	\$1,440	100%	96.0
Five-person firm	\$3,840	25%	1,024.0
Five-person firm	\$3,840	50%	512.0
Five-person firm	\$3,840	75%	341.3
Five-person firm	\$3,840	100%	256.0
Twenty-person firm	\$12,800	25%	3,413.3
Twenty-person firm	\$12,800	50%	1,706.7
Twenty-person firm	\$12,800	75%	1,137.8
Twenty-person firm	\$12,800	100%	853.3

The base Table 9 thresholds are economically realised hours. At 50% realisation, a five-person firm must free 512 gross hours to obtain 256 hours of economic benefit; at 25%, it must free 1,024. The same logic applies to all three sizes. These are hours across the firm, not per employee, and are not cash savings unless hours are eliminated or redeployed to economically valuable output.

Counterfactual assumptions (Table 9): alternatives are stylised minimum-scope configurations, not verified lowest market prices. They retain the same 12-month utilisation horizon and US\$15/hour and 30% contribution-margin assumptions as the main scenarios. A lower cost is not proof of equivalent capability: compare only where quality-adjusted output, service continuity, data protection and workflow coverage are adequate. Open-source software is not assumed to eliminate hosting, verification or staff costs. These rows model feasible substitution, not observed Caribbean purchases.

Break-even condition: annual incremental gross profit plus verified labour and operating cost savings must equal or exceed annualised subscription, integration, training, verification and maintenance costs. A US\$1,440 first-year cost requires US\$1,440 in attributable net operating benefit. At US\$15 of genuine value per hour saved, this implies 96 hours; if the benefit instead comes from new sales with a 30% contribution margin, US\$4,800 of additional sales are required. These are alternative mechanisms and must not be added together when they represent the same benefit.

Labour time saved is not automatically a cash saving. If no paid hours are reduced and employees cannot redeploy the released time productively, the financial return may be negligible. Similarly, faster content production does not guarantee additional demand or customer conversion. Evaluations should track quality-

adjusted output, turnaround time, revenue conversion, customer retention, error rates and staff workload before attributing benefits to AI.

5 | Sensitivity: exchange rates, utilisation & scope

Suppose the five-person scenario incurs US\$3,840 in first-year costs. A 20% increase in subscription prices raises its US\$2,640 annual recurring component by US\$528, bringing first-year costs to US\$4,368.

A 10% depreciation defined as a decline in the domestic currency's US-dollar value raises the local-currency cost of an unchanged US-dollar invoice by approximately 11.1%; a 10% increase in the quoted rate (local currency per US dollar) raises that cost by 10%. The three currencies, however, operate under distinct regimes. Jamaica maintains a managed float in which the rate is market-determined with Bank of Jamaica intervention; depreciation risk is therefore ongoing, although the Jamaican dollar has been relatively stable recently. Trinidad and Tobago operates a de facto stabilised arrangement, with the official rate held near TT\$6.75–6.80 per US dollar for an extended period; a material exchange-rate adjustment would more likely reflect a discrete policy decision, while foreign-exchange rationing can impose availability costs independently of the official rate. The Cayman Islands dollar is hard-pegged to the US dollar (CI\$1 = US\$1.20), limiting direct exchange-rate exposure on US-dollar invoices under the maintained peg. The actual invoice should therefore be modelled using the observed exchange rate and the relevant regime's risk profile.

If only half the anticipated productive hours are realised, the required benefit per realised hour doubles. If only a fraction of employees need a paid seat, costs can fall sharply. Bundles, free tiers, targeted licences and process redesign can also reduce expenditure. The appropriate policy response is therefore not to subsidise indiscriminate subscription accumulation but to improve informed selection, integration and measured utilisation.

6 | Three-country evidence: Jamaica, Trinidad and Tobago and Cayman

This section replaces the single illustrative US\$15-per-hour labour-value assumption with occupational and sectoral benchmarks from three Caribbean jurisdictions. It is a scenario analysis, not a pooled regional wage estimate. The sources measure different populations: Jamaican public-service pay scales; selected Trinidad and Tobago public-service job salaries; and advertised salaries in Cayman Islands job postings. They are therefore analysed separately, with identical financial assumptions to expose the effect of labour-cost differences. The Jamaican and Trinidad and Tobago public-service figures are selected government salary schedules, not samples or lower-bound estimates of private MSME pay; they must not be interpreted as representative of private-enterprise compensation, even where they serve as conservative illustrative reference points.

6.1 | Harmonised calculation and limitations

Salary conversion uses 160 hours per month and 1,920 hours per year, the study's four-week monthly convention. The alternative 2,080-hour paid-year convention is tested below. Neither convention deducts leave or nonproductive time. Salary is a proxy for the value of a productive hour, not necessarily marginal output, avoidable expenditure or fully loaded employment cost.

Standardised subscription scenario: US\$200 per month, US\$2,400 annually, for one defined software bundle or enterprise subscription allocation. This is an assumed comparison unit, not a claim that every firm requires that subscription or that it covers multiple employee seats. Illustrative exchange rates, held fixed for modelling rather

than represented as spot quotations: J\$160/US\$1; TT\$6.80/US\$1; CI\$0.82/US\$1. Thus annual local-currency costs are J\$384,000, TT\$16,320 and CI\$1,968, respectively. Implementation, training, tax, data, cybersecurity, payment charges and additional seat costs are excluded from this narrow subscription-only comparison.

Required economically realised hours = annual local-currency AI cost / hourly salary equivalent. Required gross time savings at 50% realisation = twice that amount. Realisation is the share of nominal hours saved that becomes actual economic value; it is not a prediction about the efficiency of any tool.

Hours convention sensitivity. The 1,920-hour denominator reflects 160 hours per month for twelve months. An alternative 40-hour × 52-week convention yields 2,080 hours. Holding annual pay and software costs constant, the 2,080-hour assumption reduces the implied hourly salary by 7.7% and increases the required breakeven hours by 8.3%. Neither convention deducts paid leave, holidays or nonproductive time; firm-level evaluation should use actual paid and productive hours.

Fully loaded employment-cost sensitivity. Salary-only values omit employer social contributions, benefits, paid leave, overhead and other labour-related costs. If a firm can genuinely avoid or redeploy fully loaded labour cost equal to 1.25 × base salary, the hours needed to recover a fixed AI bill fall by 20% (for example, 256 to 204.8 realised hours in the five-person US\$3,840 case). This is a sensitivity parameter, not an observed regional employer-cost ratio; if staff costs remain fixed and freed hours cannot be productively redeployed, the higher accounting labour cost does not create a cash saving.

6.2 | Jamaica: published public-service salary bands

The Ministry of Finance and the Public Service salary scales for monthly paid staff, 2025–2028, give annual amounts by pay band and step. The values below use the MINIMUM of the “With effect from 1/4/2026” row, not an average salary, actual payroll observation or private-sector wage. These are useful lower-step public-service benchmarks for professional and management functions; they cannot describe the typical Jamaican microenterprise.

Table 11. Jamaica – selected 2026 minimum scales and break-even hours

Selected 2026 minimum scale	Annual J\$	J\$/hour	Break-even realised h	Gross h at 50%
Administrative Management GMG/AM 2 (Band 4)	1,796,613	936	410.4	820.7
Administrative Management GMG/AM 3 (Band 5)	2,299,817	1,198	320.6	641.2
Senior Executive GMG/SEG 1 (Band 7)	3,676,602	1,915	200.5	401.1
Senior Executive GMG/SEG 3 (Band 9)	5,457,937	2,843	135.1	270.2

Table 11 note: Jamaica dollars (J\$), selected public-service scale minima; 1,920 paid hours/year and assumed J\$160/US\$. Not representative private-sector wages.

Interpretation: at the selected GMG/AM 2 minimum, the assumed subscription costs approximately 410 hours of salary-equivalent economic value; at the GMG/SEG 3 minimum, approximately 135 hours. This does not imply that managers are more automatable than administrative employees. The economically achievable return depends on task mix, service quality, the ability to redeploy labour, and workflow redesign. Source: Jamaica Ministry of Finance and the Public Service, Compensation Unit–ERB/SHRMD, “Salary Scales for Monthly Paid Staff for the Period April 1, 2025–March 31, 2028”, dated 2 September 2026, pp. 1–2; row effective 1 April 2026.

6.3 | Trinidad and Tobago: selected public-service occupations

The amounts are drawn from selected occupation listings on the Government of Trinidad and Tobago Public Service Careers portal. For this illustration, the listed TT\$ amounts are treated as monthly salaries; the pay-period metadata for each listing has not been independently confirmed. Accordingly, the resulting hourly values and break-even hours are conditional on that interpretation. The five deliberately selected public-service positions are not an occupation-weighted national sample and are not representative of private MSME pay.

Table 12. Trinidad and Tobago – selected occupations and break-even hours

Selected occupation	Monthly TT\$	TT\$/hour	Break-even realised h	Gross h at 50%
Facilities Technician	8,200	51.25	318.4	636.9
Accountant II	9,698	60.61	269.3	538.5
Facilities Manager	15,900	99.38	164.2	328.5
Programme Manager	21,100	131.88	123.8	247.5
Programme Director	22,000	137.50	118.7	237.4

Table 12 note: Trinidad and Tobago dollars (TT\$), selected public-sector job salaries; 160 hours/month and assumed TT\$6.80/US\$. Not nationally representative.

The annual subscription is equivalent to approximately 318 salary-value hours at the Facilities Technician benchmark, but 119 hours at the Programme Director benchmark. The result illustrates the labour-value sensitivity of the same software expenditure; it does not establish that automation of programme-director responsibilities is technically feasible, desirable or safe. Source: Government of Trinidad and Tobago, publicservicecareers.gov.tt, user-selected Facilities Technician, Accountant II, Facilities Manager, Programme Manager and Programme Director listings, supplied 20 October 2026.

6.4 | Cayman Islands: advertised salary evidence

WORC's September 2026 Job Postings Report provides year-to-date 2026 mean advertised annual salaries by industry (Appendix Table A5, report p. 16) and selected occupations (Appendix Table A6, p. 17). This is the broadest industry coverage among the three source sets, but it captures advertisements on WORC's online portal rather than all vacancies or salaries actually paid. The following table retains a cross-section relevant to AI adoption and employment distribution.

Table 13. Cayman Islands – selected industry means and break-even hours

Industry	Annual CI\$	CI\$/hour	Break-even realised h	Gross h at 50%
Professional, scientific and technical activities	92,180	48.01	41.0	82.0
Financial and insurance activities	91,832	47.83	41.1	82.3
Public administration and defence	77,262	40.24	48.9	97.8
Information and communication	62,329	32.46	60.6	121.2
Construction	32,890	17.13	114.9	229.8
Administrative and support services	29,757	15.50	127.0	254.0
Accommodation and food services	25,936	13.51	145.7	291.4

Table 13 note: Cayman Islands dollars (CI\$), year-to-date 2026 advertised industry salary averages from WORC; 1,920 hours/year and assumed CI\$0.82/US\$. Job advertisements are not observed earnings.

At the reported professional-services salary benchmark, the subscription requires approximately 41 economically realised hours; at the accommodation and food services benchmark, about 146 hours. This 3.6-fold difference reflects salary values, not demonstrated differences in actual AI productivity. Higher-value sectors may also face more expensive compliance, verification, confidentiality and integration requirements. Source: Workforce Opportunities and Residency Cayman (WORC), Labour Market Demand Unit, Job Postings Report, September 2026, Appendix Table A5, p. 16.

6.5 | What the three cases establish—and what they do not

The scenarios establish a mechanical but economically consequential relationship: for a fixed foreign-currency software bill, lower salary-equivalent hourly values imply a higher number of economically realised hours to recover costs. They do not establish the profitability of AI, the rate of job displacement, national average adoption costs, or a ranking of countries' readiness. Because labour cost is not the same as marginal product, a complete business case must model (a) incremental output and contribution margins; (b) the value of redeployed time; (c) actual cash savings; (d) software seat and usage costs; and (e) implementation and governance costs.

The distributional mechanism is straightforward: a fixed foreign-currency software bill requires more realised hours to recover when the relevant hourly value is lower. But salary is only a proxy; higher output quality, additional sales and contribution margins may change the result. The model therefore tests alternative realisation rates, task values, configurations and shared-service arrangements without assuming that wage comparisons measure observed enterprise productivity.

Two return channels must therefore be tested separately. The labour-substitution channel values economically captured reductions in hours, rework or operating costs; it depends on redeployment, avoidable expenditure and relevant task value, not simply salary. The revenue-expansion channel values additional orders, better customer conversion, higher prices justified by quality or access to export markets at the incremental contribution margin. Where the same released capacity generates the additional sales, count the realised contribution once rather than adding a hypothetical wage saving. The three-country comparisons isolate the first channel as an affordability benchmark; they cannot rank the second channel without sector- and market-specific assumptions.

6.6 | Comparative techno-economic assessment

The assessment compares identifiable technology configurations and alternative delivery pathways rather than attempting to estimate the prevalence of AI use among Caribbean firms. Principal measures are first-year and recurring total cost, break-even realised hours, incremental revenue required at a specified contribution margin, cost per quality-adjusted output where assumptions permit, and sensitivity to scale, utilisation, exchange rates and vendor terms. Each scenario separates observed vendor prices and published wage benchmarks from hypothetical implementation and productivity assumptions.

Reference configurations include a lean functional stack (ChatGPT or Claude, Canva, Zapier, Notion and Fireflies); a mid-market configuration (Microsoft 365 Copilot, HubSpot, Intercom, GitHub Copilot and Snowflake); and an enterprise-builder architecture combining cloud/model services, data infrastructure, orchestration, business applications and custom retrieval. These are illustrative combinations, not a claim that representative US or Caribbean firms deploy them. Compare each with leaner bundled, free, open-source, outsourced and shared-service alternatives. For each, identify billing units, fixed and variable costs, seats, inference and usage charges, engineering and integration effort, skills, verification, maintenance, security and switching exposure. Apply the

stated Caribbean salary and currency benchmarks to calculate conditional affordability and break-even thresholds. Do not assign unverified market prices to implementation work; use explicitly labelled low/base/high assumptions.

The critical economic test is whether a firm can secure the productive capability required by its competitive environment at a sustainable cost—not whether it owns a particular stack. Organisational design, incentives, decision rights, data quality and accountable process ownership determine how much of the potential benefit is realised. This study evaluates those mechanisms conceptually and through sensitivity analysis; it does not present scenario outputs as measured Caribbean business performance.

7 | Distribution, labour markets and domestic value capture

The likely distributional concern is not that small enterprises can never benefit from AI. It is that fixed setup costs, minimum licence commitments, specialised skills and management attention may be spread more easily across a larger customer base. The relevant comparison is cost per quality-adjusted unit of output, not nominal subscription price alone.

The Caribbean's BPO and shared-services activities face both exposure and opportunity. AI may automate repetitive customer interactions, transcription, documentation and standardised analysis. It may also support higher-value services, quality assurance and multilingual delivery. The net employment effect will depend on client demand, service repositioning, retraining and firms' capacity to move up the value chain. No Caribbean-wide employment displacement estimate is inferred here from global exposure statistics. Two concrete upgrading routes are AI-assisted customer-experience operations—complex-case escalation, multilingual support, supervised quality assurance and customer-retention analytics—and domain-specialised knowledge-process services, including compliance documentation, claims review and financial-services operations with human validation. Both require sector expertise, secure data access, documented accuracy standards and client acquisition; neither follows automatically from acquiring a chatbot.

Imported AI subscriptions are debits for cross-border services, but AI-enabled activity may also generate export receipts, substitute imports and increase domestic operating contribution. These are distinct measures: private operating contribution concerns the enterprise's financial return; net foreign-exchange effects concern additional receipts less foreign-currency payments; and national value added depends on domestic output less imported intermediate consumption, with careful treatment of labour redeployment and transfers. The illustration below reports the first two channels and does not claim to measure the third.

Illustrative contribution accounting. Let ΔR denote incremental sales attributable to AI-enabled activity; m the enterprise contribution margin after ordinary non-AI variable costs but before separately identified imported inputs; S verified incremental enterprise operating-cost savings not already included in m ; F incremental imported software, inference, hosting and payment services; and M other incremental imported intermediate inputs. Define $\Delta C = (\Delta R \times m) + S - F - M$ as an indicative retained enterprise operating contribution before any separately identified domestic setup expense. This is not an estimate of national value added: savings that represent displaced domestic wages or purchases may redistribute income rather than increase economy-wide output, and imported inputs already included in m must not be deducted twice. Separately, the direct foreign-exchange effect of an export transaction is incremental export receipts less the associated imported payments.

Worked export-services illustration (linked to Table 9). The five-person stacked configuration costs US\$220/month plus US\$1,200 setup, or US\$3,840 in year one. Assume AI-enabled service upgrading generates

US\$20,000 in additional export receipts at a 30% contribution margin before separately identified imported inputs, together with US\$1,200 in verified, non-overlapping enterprise operating-cost savings. Assume US\$2,640 of recurring charges and US\$500 of other incremental intermediates are imported. The indicative retained operating contribution before domestic setup cost is $\Delta C = (US\$20,000 \times 0.30) + US\$1,200 - US\$2,640 - US\$500 = US\$4,060$. If the US\$1,200 setup payment is made to a domestic service provider, the purchaser's first-year contribution after setup is US\$2,860, while that payment is domestic revenue rather than an automatic foreign-exchange leakage. If setup is instead imported, the purchaser's contribution is likewise US\$2,860 but the additional US\$1,200 becomes an external payment. With domestic setup, gross export receipts of US\$20,000 less US\$3,140 of specified imported charges imply a direct net foreign-exchange inflow of US\$16,860; with imported setup, it is US\$15,660. Neither figure is national value added. The calculation excludes any other imported inputs already captured in the margin, and does not presume that domestic cost savings translate one-for-one into national output.

8 | Institutional architecture and governance

AI does not remedy unclear decision rights, weak internal controls, dysfunctional incentives or absent accountability merely by producing faster outputs. On the contrary, poorly supervised automation can increase the speed and scale of error. Enterprises and governments should map the business process first, identify accountable decision-makers, define escalation and human override points, and only then choose tools.

For public administration, procurement should specify auditability, data retention, vendor exit rights, independent testing, incident reporting, security controls and measurable service outcomes. For regulated finance, healthcare, energy and telecommunications, consequential AI use should receive risk-based assessment and independent verification rather than reliance solely on vendor assurances.

Three governance approaches warrant comparison. Voluntary vendor commitments may move quickly but are weak where verification and disclosure are discretionary. Blanket pauses may be difficult to enforce, restrict beneficial local uses and entrench established suppliers. Binding risk-based obligations can improve accountability, but poorly designed compliance costs may disadvantage smaller innovators. Caribbean governments should prioritise interoperable procurement requirements, pooled testing capacity and proportionate safeguards.

9 | A sequenced Caribbean policy agenda

Immediate: establish the baseline and protect critical functions

Map AI-related supplier dependencies and high-consequence automated decisions in critical infrastructure and public agencies; require named human accountability, override procedures and incident logs. Publish a periodically updated vendor-price and stack-comparison register with transparent billing assumptions, capability categories and low-cost alternatives. This is a market-information and procurement exercise, not a prerequisite survey of Caribbean MSME invoices.

Near term: reduce integration friction and information asymmetry

Develop a regional incident-sharing protocol and a pooled independent evaluation facility for small states. Publish neutral procurement and interoperability guidance. Through chambers of commerce, development

agencies and universities, offer stack-selection clinics and practical workflow redesign support for MSMEs. Track before-and-after business outcomes rather than the number of licences purchased.

Medium term: build productive absorption and fair competition

Embed AI capability in enterprise upgrading and workforce transition programmes, particularly in BPO, tourism, professional services and public administration. Strengthen competition and data portability arrangements, assess concentration and platform switching costs, and examine regional procurement or shared-service arrangements where scale economies are demonstrable. Publish periodic estimates of domestic value capture and distributional effects.

10 | Analytical scope and interpretation

This comparative techno-economic assessment uses identifiable AI architectures, the existing 20-platform price register, selected Caribbean labour-cost and exchange-rate benchmarks, and transparent implementation, utilisation and productivity assumptions. It does not seek to measure the technologies currently purchased by Caribbean firms. Its results establish conditional affordability, break-even and competitiveness implications for defined configurations; they are not estimates of regional adoption prevalence or realised firm-level productivity. Publication checks should focus on vendor billing units and entitlements, benchmark sources, calculation accuracy and sensitivity to plausible alternatives. No proprietary Caribbean invoices, engineering time sheets, production telemetry or adoption survey are prerequisites for this analysis.

Conclusion

The Caribbean cannot afford to get AI wrong—but getting it right does not mean purchasing every new application. As competition changes the capabilities firms need, the relevant question is how to secure and sustain those capabilities at a viable cost through software, organisational redesign, specialist skills and, where appropriate, outsourcing or shared services. Lower task costs alone do not guarantee positive enterprise returns or greater domestic value added. Business strategy and public policy must therefore evaluate productive deployment, economically realised benefits and domestic value capture rather than licence counts or unsubstantiated adoption claims.

Appendix A | Pricing sources and validation status

Pricing evidence (20 representative platforms): Table 8 distinguishes per-user monthly billing from separately quoted prepaid annual-plan charges. The independently specified workflow and enterprise-stack scenarios do not mechanically aggregate those prices. Before an actual purchase, confirm the vendor plan, jurisdiction, taxes, minimum seat commitments, included credits and prerequisite licences. Canva Pro is a solo-creator offer, not an enterprise-wide licence.

IMF (2024), Gen-AI: Artificial Intelligence and the Future of Work. <https://www.imf.org/-/media/files/publications/sdn/2024/english/sdnea2024001.pdf>

IMF (2026), Leveraging Artificial Intelligence and Enhancing Countries' Preparedness. <https://www.imf.org/en/news/articles/2026/02/03/sp-md-leveraging-artificial-intelligence-and-enhancing-countries-preparedness>

IMF (2026), Europe and the Global AI Race. <https://www.imf.org/en/news/articles/2026/09/20/sp091926-europe-and-global-ai-race>

Microsoft (2026), Copilot Business pricing. <https://www.microsoft.com/en-us/copilot/pricing/business>

Microsoft (2026), SMB bundle announcement. <https://techcommunity.microsoft.com/blog/partnernews/partner-blog--partner-led-momentum-broader-availability-for-smb-microsoft-365-bu/4523663>

HubSpot, Customer Platform pricing. <https://www.hubspot.com/pricing/suite>

Zapier, Plans and Pricing. <https://zapier.com/pricing>

OpenAI, ChatGPT pricing. <https://openai.com/chatgpt/pricing>

Further platform pages for validation: claude.com/pricing; perplexity.ai/enterprise/pricing; workspace.google.com/pricing; notion.com/pricing; grammarly.com/plans; jasper.ai/pricing; make.com/en/pricing; clickup.com/brain/pricing; otter.ai/pricing; fireflies.ai/pricing; descript.com/pricing; elevenlabs.io/pricing; adobe.com/products/firefly/plans.html; synthesia.io/pricing; semrush.com/pricing/ai/; canva.com/pricing.

Country benchmark sources and limits:

Jamaica—Ministry of Finance and the Public Service, Compensation Unit—ERB/SHRMD, “Salary Scales for Monthly Paid Staff for the Period April 1, 2025–March 31, 2028”, dated 2 September 2026, pp. 1–2, minima effective 1 April 2026;

Trinidad and Tobago—Government Public Service Careers portal, selected Facilities Technician, Accountant II, Facilities Manager, Programme Manager and Programme Director listings supplied 20 October 2026 (monthly-pay interpretation remains conditional on listing metadata);

Cayman—WORC, Job Postings Report, September 2026, Appendix Table A5 (p. 16) and A6 (p. 17).

The Jamaican and Trinidadian public-sector salary scales provide occupational benchmarks for assessing AI affordability, recognising that public- and private-sector remuneration often respond to common labour-market conditions, particularly through middle management. The Cayman Islands figures, by contrast, are drawn from employer-reported salary offers in WORC recruitment data and therefore provide a more direct indication of

prevailing recruitment-market remuneration. None of these datasets should be interpreted as a comprehensive national wage distribution, but each offers a legitimate basis for the specified affordability comparisons. The analysis combines published software subscription prices and identifiable Caribbean exchange-rate and remuneration benchmarks with explicitly modelled enterprise configurations, implementation costs and productive-utilisation assumptions. Exchange-rate depreciation, changes in subscription prices and variations in productive realisation are examined through sensitivity tests rather than presented as observed outcomes.