FROM REGIONAL KNOWLEDGE TO DAILY WORK

# LAC AI Playbook

**Adrian Dunkley**
President, Caribbean AI Association

November 2026 · Review edition 1.2

Policy synthesis and proposed implementation method.
This paper does not report a new regional field study.

EXECUTIVE SUMMARY

# Start with the work

Adrian Dunkley's LAC AI Playbook combines a regional context prompt, a business knowledge base and a defined task. Its country and sector guides help readers identify possible uses. Its safety tools ask who can act, what can go wrong and what evidence supports a release decision. [1–3]

This paper develops those components into a proposed adoption method for small firms and the organisations that support them. A government agency, educator or investor can use the same method, but must adapt its controls to the people and decisions affected.

## The central recommendation

Fund and assess a complete workflow. Include preparation, staff training, review, corrections, supervision and a fallback when the tool fails. Compare that workflow with the existing way of doing the task.

The paper draws on the Playbook's practical guides, regional evaluation proposal and safety toolkit. External research supplies context, not proof that the proposed method will work in every country or sector.

Country examples below are possible pilot settings. They are not reports of completed deployments. No financial return, error reduction or time saving is promised.

## Who should use it

Owners and managers can use the pilot plan on page 8. Business-support agencies can use the support priorities on page 7. Technical reviewers should read the testing and permission requirements on pages 5–6.

**Research question.** Can a locally specified, human-reviewed AI workflow reduce total task time without exceeding an agreed error or harm threshold? The proposed comparison is the same task completed with the existing workflow. The outcome must include review and rework time.

| | |
|---|---|
| Regional evidence | 3 |
| Context, knowledge and tasks | 4 |
| Local tests and safety controls | 5–6 |
| Institutional support and pilot plan | 7–8 |
| Evidence limits and references | 9 |

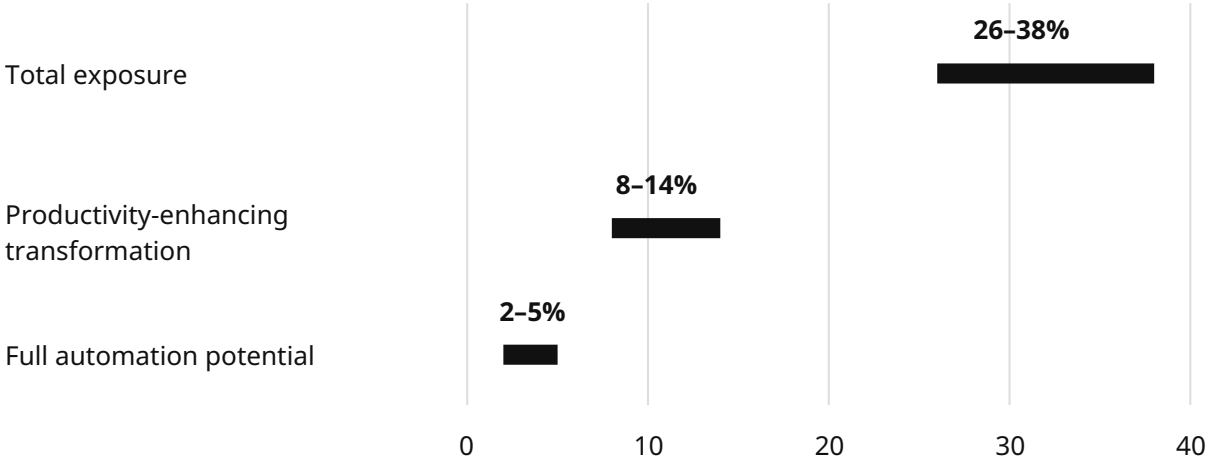
01 · REGIONAL EVIDENCE

# Exposure is not an outcome

The ILO and World Bank's 2024 analysis estimated that 26–38% of jobs in Latin America and the Caribbean could be exposed to generative AI. It distinguished potential productivity gains from the potential for full automation. Digital access could prevent many workers from gaining the benefits. [4]

## Figure 1. Estimated employment exposure to generative AI

Share of employment (%) · Published 2024 estimates




Source: Gmyrek, Winkler and Garganta, ILO Working Paper 121, 31 July 2024. Bars show published ranges, not confidence intervals. The two lower categories are parts of exposure; they do not sum to the total shown. These are model-based estimates of potential exposure, not observed adoption, realised productivity or predicted job losses. They are not separate estimates for every Caribbean economy.

## What an owner can infer

Some tasks may be suitable for assistance. That is a reason to run a controlled comparison. It does not justify replacing staff or assuming that a vendor's general benchmark applies to the business.

## What a policymaker must check

Access, training and job quality belong in the same programme. Count who can use the service and who cannot. Report small-island and rural results separately where the sample supports it.

02 · OPERATING METHOD

# Give the tool a local brief

| Input | What to include | What to check |
|---|---|---|
| **Regional context** | Country, language, currency, payment methods, delivery constraints and relevant authorities. | Use facts about this operation. Avoid assumptions about an entire region. |
| **Business knowledge** | Approved product details, opening hours, service terms, price lists and document templates. | Name the owner and review date. Remove records that the tool does not need. |
| **Task instruction** | The requested output, intended reader, allowed sources and decision that stays with a person. | State what the tool must do when information is missing or contradictory. |

## Keep the brief small enough to maintain

A knowledge base becomes a liability when staff cannot tell which price list is current. Start with a short set of approved records. Mark effective dates, superseded versions and the person who can change them.

A local context prompt can reduce avoidable assumptions. It cannot guarantee correct law, accurate arithmetic or appropriate language. Check those outputs against the source or with a qualified reviewer.

## Make uncertainty visible

Require the tool to identify missing details before it produces a customer-ready answer. A booking reply should not invent availability. A quote should not invent tax treatment. A grant application should not invent achievements.

Keep confidential customer or staff records outside a general-purpose tool unless the organisation has approved the specific data use, access and retention arrangements.

**Illustrative guesthouse task.** Draft a reply from the current room list, approved rates and cancellation policy. Use the guest’s requested language. Flag any missing dates. Do not reserve a room, promise a refund or send the reply. The owner checks the draft against the booking system.

This is a proposed workflow derived from the Playbook’s hospitality and context guidance, not an evaluated case study.

03 · EVALUATION

# Test the task in its setting

The Playbook’s proposed regional benchmark covers language, regional context, small-business tasks, cost and low-bandwidth operation. These dimensions are useful for selecting tests. They do not establish a model ranking without inspectable prompts, outputs, scoring records and reviewer procedures. [2]

## Figure 2. The proposed evidence sequence

**1** **Record the baseline**

Time the existing workflow. Retain a sample of correct outputs and common errors. Define the acceptable result before the AI trial.

**2** **Use representative cases**

Include routine work, ambiguous requests, missing information and cases the system should refuse or escalate. Use authorised or synthetic records.

**3** **Run a fair comparison**

Keep task difficulty and scoring rules comparable. Record the tool version, instructions and settings. Separate familiarisation from measured work.

**4** **Review the whole result**

Count preparation, checking and correction time. Check failures by language, task and affected group. Decide whether to stop, revise or extend the trial.

Source: proposed synthesis of the Playbook’s task guidance and regional evaluation design [1–2]. This sequence is an implementation proposal, not an empirically validated intervention.

## Measure completed work

Track total time per accepted task, correction rate and cost per accepted output. Keep serious errors separate from small wording changes. A lower average time does not excuse a harmful failure.

## Use local reviewers

Ask people who know the language and task to assess outputs. Record disagreements and how they were resolved. A translation score alone cannot establish whether a customer-facing answer is appropriate.

04 · SAFETY AND ACCOUNTABILITY

# Control what the system can do

Adrian Dunkley’s S4 method organises review into Specify, Stress-test, Safeguard and Scrutinise. It is a proposed educational method. The related Playbook toolkit uses a human-entered risk index to organise review; its score is not a probability of harm or a security certificate. [3, 5]

| S4 stage | Required record | Example decision |
|---|---|---|
| **Specify** | Purpose, users, data, allowed actions and named owner. | The tool may draft replies but cannot send them. |
| **Stress-test** | Cases, outputs, failures and unresolved questions. | An external document must not change payment instructions. |
| **Safeguard** | Access limits, approval steps and a tested manual fallback. | Restrict access to the records needed for this task. |
| **Scrutinise** | Review dates, incidents, changes and retest results. | Repeat relevant tests after a model or permission change. |

## Separate content from authority

OWASP describes indirect prompt injection as instructions entering through external material such as a webpage or file. A connected assistant may encounter such content during ordinary work. Limit its privileges and require independent approval for consequential actions. [6]

A sentence in a prompt is not an access control. The application must enforce which records and tools are available.

## Make the stop decision usable

Name the person who can suspend the workflow. Record how to remove access, preserve relevant logs and return to the manual process. Protect personal data in incident records.

NIST’s voluntary AI Risk Management Framework supports context-specific risk management across the system lifecycle. It does not replace local legal duties. Apply the relevant national and sector requirements separately. [7]

**Proposed release condition.** Do not enable a consequential action while its authority, review process or failure response remains unresolved. A low average score cannot cancel a critical control failure.

05 · INSTITUTIONAL PRIORITIES

# Build support around real constraints

## Small-business support agencies

Offer task clinics where owners bring a recurring problem and leave with a bounded pilot. Help them prepare approved records, compare total costs and select a review method. Avoid training that counts tool demonstrations as business outcomes.

## Governments and regulators

Publish usable procurement requirements: data handling, access, performance evidence, human review and exit arrangements. Ask suppliers to show how a service works under local language and connectivity conditions.

Keep the route for complaints and redress clear. An AI-generated explanation does not remove the public body’s responsibility for a decision.

## Educators and training providers

Teach learners to define tasks, inspect evidence, identify errors and know when to seek help. Assess their independent judgement as well as their ability to produce an output with AI.

## Investors and lenders

Request evidence from completed workflows rather than headline model scores. Examine recurring costs, customer retention, supervision effort and dependence on a single supplier. Distinguish a successful demonstration from a repeatable operating process.

## Regional institutions

Coordinate reusable test materials and reviewer training. Maintain country-specific entries with dates and source links. A shared format can reduce duplicated work without pretending that laws, infrastructure or customer expectations are uniform.

## Resilience and access

Include a low-bandwidth route and an offline fallback where feasible. A locally hosted model may reduce some data transfers, but still requires secure devices, maintenance, access controls and task testing. Smaller does not automatically mean safer or cheaper.

**Where to begin.** Select a service that already has reliable records and a manageable review load. Possible settings include booking enquiries, product descriptions or document formatting. Keep clinical decisions, credit decisions and other consequential judgements outside a simple starter pilot.

These are proposed programme choices. They require local feasibility and legal review; they are not findings that one sector should always adopt first.

06 · IMPLEMENTATION

# Run a pilot with a clear decision

| Stage (illustrative) | Work to complete | Evidence for the decision |
|---|---|---|
| **Weeks 1–2**<br>**Define** | Select one task. Name the owner. Map data and permissions. Record the manual baseline. | Approved scope, baseline records, cost limit and stop conditions. |
| **Weeks 3–4**<br>**Prepare** | Clean the knowledge base. Train reviewers. Assemble routine and failure cases. | Dated sources, test register and agreed scoring rules. |
| **Weeks 5–8**<br>**Compare** | Run a limited, supervised trial. Record review time, corrections and incidents. | Paired or otherwise comparable task records, with exclusions explained. |
| **Weeks 9–12**<br>**Decide** | Review results and costs. Retest changes. Stop, revise or extend the workflow. | A signed decision, unresolved limits, next review date and rollback owner. |

## Set thresholds before the trial

Adjust the schedule to risk, staffing and the availability of representative cases.

Define acceptable accuracy for each task type and what counts as a serious error. Set a maximum review burden and a budget. Choose enough cases to cover important failure modes; there is no universal sample size for all uses.

If the sample is too small to estimate rare harm, say so. A period without incidents is not proof of safety.

## Use a short decision record

State what was tested, how it compared with the baseline, what failed and what remains unknown. Explain who benefited and whether anyone faced extra work or risk.

For an expansion, specify the added users, data and permissions. Check whether that change invalidates the previous evidence.

**Decision rule.** Continue only if the workflow meets the pre-agreed quality and control requirements and offers a worthwhile benefit after checking, correction and operating costs. Otherwise revise it or return to the existing process.

The next useful regional evidence will come from documented local trials with comparable outcomes. Publishing failures and boundary conditions will help other organisations avoid repeating them.

SOURCES AND LIMITS

# Evidence and references

## What this paper establishes

This is a conceptual policy synthesis and implementation proposal. It combines the Playbook's context/knowledge/task method, regional evaluation design and S4 safety approach. The employment chart reproduces dated external estimates. No new experiment, systematic review or regional adoption survey was conducted for this paper.

## What still needs evidence

Claims of model superiority, universal time savings and country-wide readiness require separate validation. The regional benchmark is used here as a test-design proposal, without its unverified performance tables. Proposed pilot stages and release conditions are author recommendations, not official regulatory standards.

**[1] Dunkley, A. AI Playbook for the Caribbean and Latin America.** Practical guide, regional context prompt and business knowledge-base template. Author-maintained resource; source reviewed 11 November 2026. aiplaybooklac.netlify.app. Source basis: guide/README.md, guide/part-1.md, guide/context-prompt.md and guide/knowledge-base.md in the author's repository.

**[2] Dunkley, A. The LAC Benchmark, version 0.1, April 2026.** Author-maintained evaluation design, rankings/lac-benchmark.md. Its proposed test dimensions inform page 5. This paper does not reproduce or validate its reported model scores. Available through the Playbook's rankings resource.

**[3] Dunkley, A. AI Safety & Alignment Toolkit.** Human-reviewed ordinal screening method. Documentation and engine inspected 11 November 2026. aiplaybooklac.netlify.app/safety-alignment.

**[4] Gmyrek, P., Winkler, H. and Garganta, S. (2024).** *Buffer or Bottleneck? Employment Exposure to Generative AI and the Digital Divide in Latin America.* ILO Working Paper 121, joint ILO–World Bank publication, 31 July. doi.org/10.54394/TFZY7681. Primary publication page verified 11 November 2026. Modelled exposure is not observed job loss.

**[5] Dunkley, A. S4 AI Safety Framework.** Specify, Stress-test, Safeguard, Scrutinise. Author's framework and Playbook editorial source ledger dated 3 November 2026. Proposed educational method; no independent effectiveness claim is made. Original issue: aibosshq.com/safety-alignment/issue-01.html.

**[6] OWASP Gen AI Security Project (2025).** LLM01:2025 Prompt Injection. Primary guidance page verified 11 November 2026. genai.owasp.org/llmrisk/llm01-prompt-injection/.

**[7] NIST (2023; 2024).** AI Risk Management Framework 1.0 and Generative Artificial Intelligence Profile, NIST AI 600-1. Voluntary guidance. The NIST overview, checked 11 November 2026, states that AI RMF 1.0 is being revised. nist.gov/itl/ai-risk-management-framework; doi.org/10.6028/NIST.AI.600-1.

Design follows the typographic hierarchy, A4 page structure and publication furniture of CAIA's *Learning on Our Terms* editorial white paper, adapted to black and white. No United Nations affiliation or endorsement is implied.