



DAWGEN GLOBAL THOUGHT LEADERSHIP

The Caribbean AI Realisation Series · Edition 07

The Caribbean AI Realisation Playbook.

Every framework, ladder, and checklist from seven editions — consolidated into one board-ready reference.

A position paper by **Dr. Dawkins Brown**

Executive Chairman & Founder, Dawgen Global

PUBLISHED June 2026 **AUDIENCE** Boards, CEOs, CFOs, CIOs, CHROs

DAWGEN GLOBAL THOUGHT LEADERSHIP

The Caribbean AI Realisation Series • Edition 07 • The Series Finale

The Caribbean AI Realisation Playbook.

Every framework, ladder, and checklist from Editions 01–06 — consolidated into one board-ready reference for governing AI in a Caribbean enterprise.

Keep this document beside the board table. It is built to be navigated and used, not merely read once.

A position paper by **Dr. Dawkins Brown**

Executive Chairman & Founder, Dawgen Global

PUBLISHED June 2026 **JURISDICTION** Caribbean & Global **AUDIENCE** Boards, CEOs, CFOs, CIOs, CHROs

HOW TO USE THIS PLAYBOOK

Six papers asked six questions. This one puts the answers on a single table, in the order a board actually meets them.

This is the seventh and final paper in Dawgen Global’s Caribbean AI Realisation Series. The previous six were written to be read in sequence; this one is written to be kept and consulted. It reproduces, in one place, every framework, ladder, and checklist the series produced, organised so a board can find the right tool for the decision in front of it without re-reading six documents.

The playbook follows the life of an AI initiative as a Caribbean enterprise actually experiences it: deciding what to attempt first (Edition 01), choosing where to begin (02), funding what comes next (03), authorising systems that act (04), transitioning the workforce those systems change (05), and choosing who advises through all of it (06). Each section below restates the core framework from its edition in reference form, followed by the director’s questions that close it. A consolidated master checklist appears at the end.

A note on how to read the tables. The ladders — the Authority Ladder and the Workforce Transition Ladder — are colour-coded by escalating consequence: green for the lowest-stakes stage, amber and red as the stakes rise, navy at the top. The two ladders are designed to move together, and the playbook is built around that pairing: every increase in what a machine is allowed to do should be matched by a deliberate decision about what the people beside it now do.

“A framework read once is an idea. A framework kept beside the table is a discipline. This playbook exists to turn six editions of ideas into one standing discipline for the Caribbean boardroom.”

— Dr. Dawkins Brown, Executive Chairman, Dawgen Global

The series at a glance

Edition	Title	The question it answers
01	From AI Ambition to Measurable Outcomes	How do we run a disciplined first pilot?
02	Five Caribbean AI Use Cases Worth Funding First	Where should we actually start?
03	Wave Two: What Comes Next	What do we fund second — and are we ready?
04	Before the First Agent Goes Live	Should we authorise an AI that acts?
05	When the Work Changes Hands	What happens to the people on the payroll?
06	Who Sits Beside You	Who should advise us through all of it?
07	The Caribbean AI Realisation Playbook	How do we keep all of it on one table?

Part 1 — The disciplined first pilot (Edition 01)



The series began with the most common cause of failed Caribbean AI pilots: not the technology, but an imprecise definition of the use case and no agreed measure of success. Edition 01 set a discipline for the first pilot — narrow scope, a named owner, a baseline, and a pre-agreed measure of whether it worked. The reference test is the five-part definition every pilot should satisfy before funding.

The first-pilot discipline

Element	The standard a fundable pilot meets
Specific use case	A single, narrowly defined task — not “adopt AI” but “draft first-pass responses to a defined class of customer query.”
Measurable outcome	A baseline captured before the pilot and a pre-agreed metric of success, so the result is judged against evidence, not impression.
Named owner	One accountable person, senior enough to act on the result — not a committee and not “IT.”
Bounded cost & time	A fixed budget and a fixed window, so a pilot that does not work is stopped rather than quietly extended.
Decision rule	Agreed in advance: what result leads to scale, what leads to stop, what leads to iterate. Decided before the data arrives, not after.

Director's questions — the first pilot

1. Is the use case a single defined task, or a vague ambition?
2. What is the baseline, and what metric decides success — agreed before we start?
3. Who, by name, owns this pilot and can act on the result?
4. What are the fixed cost and time bounds, and who can stop it?
5. Have we written down, in advance, what result leads to scale, stop, or iterate?

Part 2 — Where to begin: five use cases (Edition 02)




Edition 02 named the categories of use case that most reliably repay a Caribbean enterprise's first investment — chosen because they combine high frequency, tolerance for human review, and a clear baseline to measure against. They are starting points, not a ranking; the right first use case is the one in this set that maps to the enterprise's own highest-volume, lowest-risk work.

Use-case category	Why it repays a first investment
Document-heavy drafting	High volume, repetitive, and naturally reviewed by a human before it leaves — risk stays low while time saved is visible.
First-line customer response	Large query volumes with a long tail of routine questions; AI drafts, a person approves, and resolution time falls.
Summarisation & search	Turning large internal document stores into answerable knowledge — immediate productivity, contained risk.
Data extraction & entry	Moving structured data from forms and documents into systems; clear baseline, easily measured accuracy.
Analysis & first-draft reporting	Producing first-pass analyses and reports for human refinement — leverage on scarce senior analytical time.

The selection test

From these five, the right first use case for your enterprise is the one that sits at the intersection of three things: it is high-frequency work, it tolerates a human reviewing the output before it takes effect, and it has a baseline you can measure today. If a candidate fails any of the three, it is a worse first pilot than one that passes all three — however exciting it sounds.

Part 3 — Wave two: what to fund second (Edition 03)



Edition 03 addressed the moment a first pilot has worked and the enterprise asks what to fund next. Wave two is more ambitious — deeper integration, more autonomy, higher stakes — and the discipline is a readiness test before scaling, not after. The reference is the four readiness conditions an enterprise should satisfy before committing to wave two.

Readiness condition	What it asks
Proven wave-one result	Did the first pilot actually meet its pre-agreed measure — on evidence, not enthusiasm?
Data foundation	Is the data the next use case depends on accurate, accessible, and governed — or will the system inherit a mess?
Capability to operate	Can the enterprise run, monitor, and maintain a more autonomous system — or only buy one?
Governance to match	Does oversight scale with the autonomy being added — the bridge into Edition 04?

The wave-two principle

Wave two should be funded on demonstrated readiness, not on the momentum of a successful pilot. A first win creates appetite; appetite is not a foundation. The four conditions above are the difference between scaling a capability and scaling a liability.

Part 4 — Authorising AI that acts (Edition 04)



Edition 04 drew the central line of the series: between AI that recommends and AI that acts on the enterprise’s behalf. Crossing it is a delegation of authority, and a board decision. The two reference tools are the Authority Ladder — how much autonomy is delegated — and the six governance additions an agentic deployment requires.

The Authority Ladder

Level	What the system is allowed to do	Board control required
L1 — Read	Observes and reports; takes no action and changes nothing.	Minimal — data access and privacy review.
L2 — Recommend	Proposes actions for a human to accept or reject; the human acts.	Review of recommendation quality and bias.
L3 — Draft	Prepares actions ready to execute; a human approves before they take effect.	Approval workflow; the human must be able to challenge, not just click.
L4 — Act w/ oversight	Executes autonomously within set limits; a human monitors and can intervene.	Kill switch, audit trail, defined limits, named owner.
L5 — Act autonomously	Executes within its mandate without routine human checkpoints.	Full governance stack; board-level authorisation and continuous assurance.


The six governance additions for agentic deployments

Addition	What it requires
Defined authority limits	Explicit written bounds on what the system may do, to what value, in what domain — beyond which it must escalate.
A kill switch	A tested means to stop the system immediately, owned by a named person, not buried in the vendor’s console.
An audit trail	A complete, reviewable record of what the system did and why — sufficient to reconstruct any decision after the fact.
A named system owner	One accountable executive who owns the system’s behaviour, performance, and assurance — not the vendor.
Escalation & exception handling	A defined path for what the system does when it reaches the edge of its mandate — and who receives the exception.
Continuous assurance	Ongoing monitoring of accuracy, drift, and behaviour — not a one-time sign-off at go-live.

Director’s questions — authorising AI that acts

1. Where on the Authority Ladder does this system sit — and where will it sit in a year?
2. What exactly can it do without a human, and what are the written limits?
3. Who owns the kill switch, and when was it last tested?
4. Could we reconstruct any single decision it made, from the audit trail alone?
5. Who, by name, is accountable for its behaviour in production?

Part 5 — The workforce transition (Edition 05)



Edition 05 governed the human half of the same authorisation: what automation does to the people on the payroll. Its tools are the Task-Reallocation Map, the Workforce Transition Ladder, and the six workforce commitments. The Workforce Ladder is built to move in step with the Authority Ladder of Part 4.

The Task-Reallocation Map

Destination	What it means	The board’s concern
Absorbed	Now done by the system end to end.	Lost capability — who can still do it by hand?
Reshaped	Still human, but checking rather than doing.	Richer role, or a hollow rubber stamp?
Created	New oversight, exception, and data-stewardship work.	Who is trained and funded — before go-live?
Removed	Disappears entirely; was only ever a workaround.	Capacity released — redeploy to neglected work.

The Workforce Transition Ladder

Stage	What the human does	Board priority
S1 — Augmented	Does the work; the system assists. Accountability stays human.	Fluency — people use the tool well.
S2 — Supervisory	Reviews and approves the system’s work before it takes effect.	Reskilling — reviewing is harder than doing.
S3 — Exception	Handles only what the system escalates; most of the old role is gone.	Role redesign and redeployment planning.
S4 — Stewardship	No longer does the work; owns the system that does.	A new, high-value role — a deliberate destination.

The six workforce commitments

Commitment	In one line
Decide redeploy / retrain / reduce	Choose the path for each group and say so — ambiguity is not kindness.
Fund reskilling before the system	A released line item before go-live, or it will not happen.
Protect institutional memory	Name who can still do the automated task by hand if the system fails.
Communicate before the rumour	The workforce hears it from leadership, in a planned way, first.
Redeploy by default	Released capacity goes to neglected work; reduction is the considered exception.
Name an executive owner	One accountable owner — CHRO or COO — for the whole transition.

Director’s questions — the workforce transition

1. Has this been mapped at the task level, not just headcount?
2. For each group, have we decided redeploy / retrain / reduce — and will we tell them?
3. Is the reskilling funded and scheduled before go-live, in this business case?
4. Who retains the ability to do the automated work by hand if the system fails?
5. Will our people hear this from us, in a planned way, before they hear it as a rumour?

Part 6 — Who advises through all of it (Edition 06)



Edition 06 turned the question on the advisor: the AI era rewards a model that is independent of the vendors, integrated across the disciplines, and accountable to the enterprise. The reference tools are the three properties and the five questions a board should put to any advisor — including Dawgen Global.

The three properties

Property	What it means	The failure it prevents
Independence	Economics not tied to any platform being evaluated.	Advice shaped by the advisor’s product margin.
Integration	All disciplines under one engagement, reconciled before the board sees them.	Seam failures and the reconciliation tax.
Accountability	Answers to the enterprise for the outcome, across the system’s life.	The orphaned decision — live, but owned by no advisor.

Director’s questions — the advisor

1. Do you make money if we choose a particular platform or model?
2. Who, by name, is accountable for the whole decision — not just one discipline?
3. Will you show us the alternatives you rejected, and why?
4. Who owns this relationship in three years, when the system is live?
5. Can the same team advise on the governance, the workforce, and the contracts together?

The consolidated board checklist

One page to take into the meeting. If the enterprise cannot answer these, it is not ready to proceed to the next stage — whichever stage it is at.

Stage	The board must be able to say...
Before the first pilot	“We have a single defined use case, a baseline, a named owner, fixed bounds, and a written scale/stop/iterate rule.”
Choosing where to begin	“Our first use case is high-frequency, human-reviewed, and measurable against today’s baseline.”
Before wave two	“Wave one met its measure; our data, capability, and governance are ready for more autonomy.”
Before AI that acts	“We know its ladder level, its written limits, its kill-switch owner, its audit trail, and its named owner.”
Before production (people)	“We have the task map, the path for each group, funded reskilling, memory protection, and a communication plan.”
Choosing the advisor	“Our advisor is independent where it matters, integrated across the disciplines, and accountable for years.”

The single principle beneath the whole playbook

Every framework in these seven editions is a way of making one habit routine: decide deliberately, in advance, and in writing — about the machine and the people together. The Authority Ladder and the Workforce Ladder climb in step; the governance of the system and the transition of the workforce are one authorisation, considered twice. A board that holds to that single discipline does not need to remember every table in this document. The tables are only there to make the discipline easy.

How Dawgen Global works with Caribbean enterprises

Dawgen Global is an independent, integrated, multidisciplinary professional services firm headquartered in New Kingston, Jamaica, operating across more than fifteen Caribbean territories. The firm spans the disciplines an AI decision actually touches — Business Advisory, Risk Management, Cybersecurity, HR Advisory, Tax, Audit & Assurance, IT & Digital Transformation, and Legal Process Outsourcing — and selects technology through a curated network of global partners and vendors rather than selling a platform of its own.

This playbook is the discipline the firm brings to Caribbean boards: one engagement accountable for the whole AI decision, the system and the people governed together, every framework above applied to the enterprise’s own situation. Boards wishing to put the playbook to work — or to pressure-test where they currently stand against the consolidated checklist — can contact Dawgen Global at info@dawgen.global.

This completes the Caribbean AI Realisation Series. Seven editions, one discipline, kept on one table.

Dawgen Global • Big Firm Capabilities. Caribbean Understanding.

The Caribbean AI Realisation Series • Edition 07 • “The Caribbean AI Realisation Playbook.”