Practical AI for Strategy in Financial Services

Archive text edition of the revised manuscript · 11 November 2026

Prepared by the Harbour Think Tank under the leadership of Adrian Dunkley. The original evidence review is dated 26 October 2026.

About this paper

The Harbour Think Tank of the Caribbean AI Association covers the economy, work and enterprise. This paper is one of fourteen in the collection. It looks at how a financial institution prepares strategic choices and tests the assumptions under them. It also looks at how the institution keeps a record of what people agreed. [1]

We wrote this paper for strategy teams, executives, finance leaders, programme owners and board secretariats. Its subject is the preparation of decisions, and nothing in it automates corporate direction. Evidence collection has its own paper on market research, and the risk and compliance papers cover their specialist control processes.

This revision adds a page of build notes to each use case. The notes give a starter prompt, example tools, skills, a method and the limits for an AI agent.

Foreword

The strategies that worry us most are convincing ones that depend on an assumption that nobody owns. A regional growth figure becomes a sales target. A team reads the announcement of a competitor as proof of customer demand. An approved project becomes a promised benefit before anyone defines its measure.

An assistant can help to show these gaps. It can connect a claim to its source and keep disagreements on the page. It can ask what must be true for a proposal to work. We test one thing: that people find the decision easier to challenge afterwards. A good presentation alone does not show that the decision is good.

Start with one live question and a small, authorised evidence set. Keep the original documents and the uncertainty, and keep the name of the person who must decide. Whatever the tool prepares, priorities and resources stay with that person and with the existing authorities of the institution.

Adrian Dunkley, President, Caribbean AI Association

Acknowledgements and terms

The Harbour Think Tank of the Caribbean AI Association prepared this paper, and President Adrian Dunkley led the work. The institutions that we cite supply evidence and methods, and a citation implies no endorsement by them. The ten workflows, the screening scores, the prompts and the agent levels are our proposals for test. We have no results to report from the strategy department of a named firm.

In this paper, artificial intelligence (AI) is software that generates content or classifies information. It can also make predictions. An assistant prepares work for review, and an agent can also do configured steps with tools. These proposals give no authority to allocate budgets, make investment decisions, make external commitments or change the approved strategy.

An assumption is something that the plan relies on and that no one has established as fact. A scenario gives a plausible set of conditions for discussion, and no one must read it as a forecast. A dependency is an input or a decision that another activity needs before it can proceed. A benefit baseline records the start position, and the firm later assesses a result against it.

A prompt is the written instruction that a person gives to an assistant. The word skill has two meanings in this paper. A staff skill is an ability that a person learns. An agent skill is a saved set of instructions that an assistant loads for one task. [5]

Approved source

Draft or retrieval

Human review

Accepted output

The operating boundary

The assistant prepares evidence and questions, with review queues. Authorised people keep the priorities, the budgets, the commitments and the legal judgements. They also keep the official decision record.

Contents

Section | Page

Executive summary | 5

1 The regional business case | 6

2 What the evidence supports | 7

3 The enterprise lens | 8

Process design: BPMN swimlanes | 9

4 Ten practical use cases: how to read each use case | 10

Use cases 01 to 10, each with build notes | 11-30

Priority and responsibility tables | 31-32

5 The foundations lens | 33

6 From pilot to practice | 34

Appendix A Pilot decision canvas | 35

Appendix B Measure value | 36

References | 37-38

Read the summary and the priority tables first. Then select one use case and define a pilot. Complete the decision canvas with the people who do the work.

Executive summary

Start with an assumptions register and a business-case challenge, together with an evidence pack for one defined decision. Each can run on existing documents with named reviewers. None needs a connection to production systems or the authority to commit resources.

Keep observed facts, estimates, assumptions and choices apart. Do this in every summary and chart, and in every board paper that follows. If a source is missing, find it or disclose the gap. Do not treat the unsupported claim as a fact with low confidence. A score that compares initiatives records a judgement under stated criteria. No one can read future returns from it.

Measure four things: unsupported claims, missed dependencies, material omissions and the effort to reach a pack that reviewers can use. After the meeting, check that the agreed actions and benefits are still traceable. Faster drafts help only while decision-makers can still see the evidence and the grounds to disagree with it.

Each use case now has build notes with a starter prompt and example tools. The notes also set a limit for AI agents. In the pilots, no agent acts without the approval of a named person.

Decisions for the sponsor

Select one decision that has a real owner and a deadline. Define the records that the assistant can read. Define the questions that go to finance, risk, legal or board review. Agree in advance on the failure that stops the pilot. Keep the decision itself inside the existing authority structure of the institution.

Choose a bounded process

Measure the baseline

Test with review

Decide from evidence

1 The regional business case

Regional headlines can hide a distinction that matters for a plan. The Caribbean Development Bank (CDB) reported preliminary growth in real gross domestic product (GDP) for 2025. Growth was 4.7% with Guyana and 0.6% without it. The regional definition of the CDB covers its 19 Borrowing Member Countries, and the second figure removes Guyana from the group. The two figures measure changes in aggregate GDP. Neither is an unweighted average of country growth rates. [2]

The chart compares the two aggregates, and each bar compares the real GDP of a group with its real GDP one year earlier. The data use information that was available at 16 February 2026, and the report marks the 2025 data and estimates as preliminary. Neither bar measures demand for a financial product or the addressable market of a firm, and neither measures an AI opportunity. The bars are not shares, and you cannot add them together. [2]

When a statistic becomes a planning assumption, a strategy team must keep its geography and its date with it. Keep the economic coverage also. A business that serves particular islands or sectors, or particular customer groups, needs evidence about that exposure. The regional number only frames the question. A regional review does not supply the evidence to answer it for the book of business of one institution.

Apply the same discipline to internal plans. For each figure, identify the legal entity, the currency, the period and the population behind it. Two columns with similar headings can still have different definitions. Do not merge them because the headings agree.

2025 growth depends on the regional definition

Including Guyana: 4.7%

Excluding Guyana: 0.6%

CDB, Caribbean Economic Review and Outlook 2025–2026, pages 8 and 12. [2] These are preliminary 2025 estimates, with information at 16 February 2026. Each bar compares aggregate real GDP with the year before. The first bar is for the 19 Borrowing Member Countries, and the second is for the group without Guyana. The bars are not shares, and they are not an unweighted country average.

2 What the evidence supports

The 2026 foresight toolkit of the United Nations Development Programme (UNDP) describes horizon scanning and alternative scenarios. It also describes how to test a strategy against different future conditions. Its scenario method keeps plausible narratives apart from forecasts. The toolkit gives us a planning method, but it has no evidence that generative AI improves bank strategy or predicts a Caribbean economy. [3]

We use the method because it makes a team show its uncertainty before it selects a response. An assistant can organise source material and draft questions. The people on the team still decide which assumptions are credible and which scenarios need attention. They also decide which risks the institution can accept.

We cannot point to a measured return from AI-assisted strategic planning, and this paper claims none. A convincing narrative can still combine incompatible data or repeat the bias of a source. It can also bury a weak assumption. Test the assistant on a reviewed evidence set, and keep a comparison with the current process.

Even with an uncertain future, some claims have much better support than others. Current facts need sources, and scenarios need explicit assumptions. Recommendations need reasons, and they need a decision-maker who answers for them. Keep those labels visible when someone shortens a summary for a meeting.

A scenario is not a forecast

The UNDP toolkit offers planning methods. We have no measured evidence of AI returns in financial services. Each firm must test source accuracy and retained uncertainty locally, and measure the review effort.

3 The enterprise lens

Follow one strategic choice from the first question to the review that comes after approval. On the way, look at the evidence and the alternatives. Note where the team loses source references or changes definitions. Note also where it treats a pending decision as an approved commitment. A bounded assistant can help at those points.

Start with read-only access, on public material and authorised internal extracts. Keep confidential plans and client records out of tools that the firm did not approve. Do the same for commercially sensitive negotiations. External documents can contain instructions that target the assistant. Treat them as evidence to inspect and never as commands.

The ten use cases cover different parts of the decision process, and each stops at a set boundary. The customer-need synthesis uses existing research and runs no new survey. The competitor analysis reads lawful public disclosures and makes no guesses about secret plans. The regulatory scan builds a review queue, but it does not interpret the law or certify compliance. The prompts in the build notes are start points that each firm must test.

Request and source

Assistant output

Reviewer correction

Accepted result

Keep the requested decision apart from the evidence behind it. Before a proposal becomes a commitment, write down the alternative and the unresolved assumption, with the accountable owner.

Distinguish assumptions from established facts

Business Process Model and Notation (BPMN) swimlanes. Proposed workflow for use case 01.

Strategy planning lead

AI drafting assistant

Confirmed assumption owners

Yes

No

Draft strategy authorised

Provide plan, sources and agreed definitions

Extract assumptions with supporting passages

Check classifications, ownership and gaps

Entries confirmed?

Agree review dates and reconsideration triggers

Register available

Revise classifications and unresolved questions

Save confirmed assumptions and dates

Strategy assumptions register

The assistant invents no evidence and assigns no owner automatically. It does not validate the strategy.

How to read the model. Rounded rectangle: task. X diamond: one decision path. Thin circle: start. Thick circle: end. Solid arrows show the sequence in one process. Lanes show who is responsible.

Notation: Object Management Group, BPMN 2.0.2. https://www.omg.org/spec/BPMN/2.0.2

4 Ten practical use cases

How to read each use case

Each use case has two pages. The first page gives the workflow, the minimum data, the pilot test and the accountable role. The second page gives build notes in five parts.

Starter prompt. Copy the prompt into the approved assistant, and replace the text in square brackets. Each prompt gives the assistant a role, its sources, one task, the rules and the output format.

Tools. The notes name example products. The names are examples and not endorsements. Vendors change product names frequently. Use only a service that the firm approved.

Skills. A staff skill is an ability that the team must build. An agent skill is an instruction file that the assistant loads for the task. Agent Skills is an open format, and many assistants can read it. [5]

Method. The notes name an established method that gives the work a structure.

Agent limits. The notes say what an AI agent can do, and where it must stop.

Four agent levels

Level 0, assist. The assistant answers one request. A person does every step.

Level 1, prepare. The agent collects approved inputs and prepares a draft. It changes no system and sends no message.

Level 2, act with approval. The agent proposes an action. A named person approves each action before it runs.

Level 3, act alone. The agent acts without approval inside fixed limits.

Each pilot in this paper starts at Level 0 or Level 1. A move to a higher level needs a new approval after the pilot. No pilot is at Level 3.

01 Build an assumptions register with evidence and owners

Authorised draft strategy

Classify facts and assumptions

Planning lead confirms owners

Evidence-linked register

The assistant reads a draft strategy that the firm authorised for review. It extracts each assumption and links it to the passage that depends on it. Each entry gets one of four labels: sourced fact, estimate, management judgement or unresolved question. The strategy lead confirms the label and the owner.

The assistant needs the current plan, its sources, the decision dates and the agreed definitions. An assumption that has no evidence or no owner shows as a gap. A job title alone never makes a person an owner. Record the responsibility only after the person confirms it.

For the pilot, use a plan that has an outdated market estimate and two assumptions that contradict each other. Reviewers check every entry. Count the missed assumptions and the incorrect classifications, and record the time to a register that the team can use. Stop if the assistant changes an assumption into a fact or invents evidence for it.

The strategy planning lead is accountable. The register makes dependencies visible, but it does not validate the plan. Owners set the review dates and the conditions that reopen each assumption.

First working session

Mark the assumptions in one authorised draft. Record the source and the confirmed owner of each, or flag the gap.

01 Build notes: assumptions register

Starter prompt

"You build an assumptions register from a draft strategy. Use only the attached draft and its sources. List each statement that the plan depends on. Give each one a label: Sourced fact, Estimate, Management judgement or Unresolved question. Quote the passage that depends on it, with the page. Give the source, or write 'No evidence'. Write 'Owner not confirmed' unless a named person confirmed the ownership in the attached record. Mark two

assumptions that contradict each other. Do not change an assumption into a fact. Do not invent a source. Do not say that the strategy is sound."

Tools

Draft and sources: Word documents in a restricted SharePoint library.

Extraction: Claude or Microsoft 365 Copilot, with the draft and the sources attached.

Register: an Excel table or a SharePoint list, with columns for the label, the owner, the review date and the reopen condition.

Skills

Staff skill: tell an assumption from an established fact, and ask who owns it.

Agent skill: an assumptions-register skill with the four labels and the owner rule.

Method

Assumption-based planning: find the assumptions that the plan depends on most, and set a signpost for each one. A signpost is an event that tells the owner to look at the assumption again.

Agent limits

Level 1. An agent can prepare the draft register when the planning lead supplies an authorised draft. It must not assign an owner. It must not change the strategy document.

02 Synthesise existing customer-need evidence without inventing demand

Strategic customer question

Approved research and methods

Themes and conflicting findings

Research lead verifies

Customer-needs evidence

An assistant reads approved studies and groups their findings into themes. Next to each theme, it keeps the date, the population, the method and the limits of each study.

Use authorised reports with their methods, and permitted extracts. Remove identities where possible. If the sample or the question wording is missing, the summary must not estimate how many customers a finding affects. It must not treat the finding as representative.

The hard line is between what participants said and what the institution could do about it. An analyst keeps the two apart. If a stated need becomes an intention to buy, the summary shows demand that no study measured, and the test stops. Test on a small set that has conflicting findings and one older study. Compare the result with a reference that an analyst reviewed.

Count the lost qualifications and the unsupported themes, and count the omitted contradictions. An invented quote ends the test, and so does a participant that a reader can identify. The customer research lead is accountable. New surveys and samples stay with the research process, and so does the validation of demand.

First working session

Compare two research reports. Keep one disagreement between them, with the method and the date of each.

02 Build notes: customer-need evidence

Starter prompt

"You group findings from approved customer research into themes. Use only the attached studies. For each theme, list the findings that support it, with the study, the date, the population and the method. Keep each limit that a study states. Put findings that conflict in one table, and do not resolve the conflict. Use the words of the study for

what participants said. Do not say how many customers have a need unless a study measured it. Do not change a stated need into an intention to buy. Do not invent a quote. Do not include a detail that identifies a participant."

Tools

Studies: approved research reports and de-identified extracts, for example exports from Qualtrics.

Synthesis: a Claude Project that contains only the approved studies.

Evidence table: Excel, with one row for each finding and its study.

Skills

Staff skill: read a sample description and say what a study can and cannot show.

Agent skill: a research-synthesis skill with the evidence table and the rule on stated need.

Method

Thematic analysis (Braun and Clarke) with an evidence table. Two analysts code a sample of findings independently, and they compare their themes.

Agent limits

Level 0. An analyst starts each synthesis from an approved set of studies. The assistant has no access to participant records or to survey tools. It must not contact a customer.

03 Compare public competitor disclosures on a consistent basis

Public report for entity A

Public report for entity B

Align period and definitions

Intelligence lead checks

Factual competitor briefing

Official reports from competitors state priorities and disclosed investments, with operating measures. The assistant extracts these items. It attaches the entity and the period to each item, with the source page. It flags a measure that two firms define differently.

The minimum data is a set of lawful public disclosures and an approved comparison question, for a common period. Anything that a firm did not disclose stays undisclosed. An announced initiative shows intent. The reports cannot tell you if the firm delivered it, or what the firm plans in confidence.

Pilot on a small number of reports that have different fiscal years and segment definitions. A reviewer checks every comparison and calculation. Count the unsupported inferences and the period mismatches, and count the source errors. The test ends if the assistant uses restricted information or invents a market share. The test also ends if the briefing shows an ambition of a competitor as an achieved result.

The briefing is factual, and it gives no basis to coordinate prices or strategy with competitors. Legal specialists review any use that is sensitive for competition law. The competitive intelligence lead is accountable.

First working session

Compare one measure from two public reports. Check the entity and the period, and then the definition.

03 Build notes: competitor disclosures

Starter prompt

"You compare public disclosures from two firms. Use only the attached public reports. For the measure [measure], give the value from each report, with the legal entity, the period, the definition and the page. If the periods or the definitions are different, write 'Not comparable without adjustment' and say why. Do no arithmetic. List announced initiatives under 'Stated intent', and list reported results under 'Reported result'. Write 'Not disclosed' where a report is silent. Do not estimate a market share. Do not guess a plan that a firm did not publish. Do not use

information from outside the attached reports."

Tools

Sources: annual reports and results releases from the website of each firm, and Jamaica Stock Exchange notices.

Extraction: Claude or Microsoft 365 Copilot, with the reports attached.

Comparison table: Excel, where an analyst does each calculation.

Skills

Staff skill: align fiscal years and segment definitions before a comparison.

Agent skill: a disclosure-comparison skill with the comparison table and the "Not disclosed" rule.

Method

A like-for-like comparison table with one row for each measure. Each row has the entity, the period, the definition and the source page. A reviewer checks every figure against the report.

Agent limits

Level 1. An agent can collect new public reports from an approved list and prepare a draft table. It reads public sources only. It must not send the briefing outside the strategy team.

04 Prepare an initiative comparison under explicit criteria

Comparable initiative proposals

Sponsor-approved criteria

Explain scores and gaps

Portfolio owner challenges

Initiative discussion matrix

A scoring matrix can hide the judgements behind its numbers. An assistant sets the proposed initiatives against criteria that the sponsor approved. The criteria are strategic fit, evidence readiness, dependencies, costs and the consequences of failure. The assistant shows the missing inputs, and it gives the reason for any score that it suggests.

The inputs are comparable initiative descriptions, agreed criteria, score definitions and named reviewers. Where costs are uncertain, finance supplies a range and the matrix keeps the range. No one invents precision to fill a cell. An incomplete proposal stays incomplete.

Test on fictional initiatives. Compare the explanations of the assistant with a review by people from more than one discipline. Then change one weight and look at the order. If a small change gives a different order, the ranking was never firm. Count the unsupported scores and the hidden trade-offs, and record any inconsistent treatment. Stop if the model allocates budget or treats the ranking as approval. Stop if it rewards a proposal that left out its risks.

People select the criteria and the final priorities, and the strategy portfolio owner answers for the result. The matrix records a discussion, and the investment-approval process of the institution continues as before.

First working session

Score fictional initiatives under approved criteria. Then change one weight and review the order.

04 Build notes: initiative comparison

Starter prompt

"You prepare a discussion matrix for proposed initiatives. Use only the attached proposals and the criteria that the sponsor approved. For each initiative and each criterion, give the evidence from the proposal, with the page. Suggest a score only where the score definition and the evidence support it, and give your reason. Write 'Input missing' where a proposal is silent. Keep each cost range as finance supplied it. List the trade-offs between the initiatives. Do not calculate a total or a rank. Do not fill a gap with an estimate. Do not recommend a budget. Do

not call the matrix an approval."

Tools

Scoring model: an Excel workbook that holds the weights and calculates the totals.

Explanations: Claude or Microsoft 365 Copilot, with the proposals and the criteria attached.

Sensitivity test: a second sheet in the workbook that changes one weight at a time.

Skills

Staff skill: challenge a score and ask for the evidence behind it.

Agent skill: an initiative-matrix skill with the score definitions and the "Input missing" rule.

Method

A weighted scoring model with a sensitivity test on the weights. If a small change in one weight changes the order, show the two orders to the sponsor.

Agent limits

Level 0. The portfolio owner starts each comparison, and people set every final score. The assistant has no access to budget systems. It must not publish a ranking.

05 Develop scenarios and test the same proposal against each

Sourced starting conditions

Contrasting assumptions

Test the same proposal

Finance and risk challenge

Reviewed scenario discussion

The team selects the uncertain drivers and their plausible ranges. From those, the assistant drafts contrasting scenarios. It asks how one specified proposal performs in each scenario. It records the weaknesses next to the opportunities, and it keeps assumptions apart from current observations.

You need the decision, the time horizon, sourced start conditions and scenario assumptions that the sponsor approved. Label a range that has no support as an exploratory choice. A numerical probability needs an approved analytical basis. Without one, leave the probability out.

Run the pilot as a workshop on a fictional expansion proposal. Reviewers check that each scenario is internally consistent, and that the same decision criteria apply in every scenario. Count the unsupported claims and the missed constraints. The workshop stops if a scenario appears as a forecast or an invented event appears as current news. It also stops if a plausible story appears as a quantified loss estimate.

The strategy scenario lead is accountable. Finance and risk specialists own any validated financial model, and scenarios only support discussion. The method depends on a planning distinction, and we did not test how it performs with an AI assistant. [3]

First working session

Test one fictional proposal against contrasting scenarios, and count the unsupported claims.

05 Build notes: scenario discussion

Starter prompt

"You draft scenarios for a strategy workshop. Use only the attached decision and start conditions, with the drivers that the team selected. Write [number] contrasting scenarios. In each one, state the value of each driver and label it 'Assumption'. Label each current observation 'Observed', with its source. Label a range that has no source as 'Exploratory choice'. Then say how the proposal performs in each scenario against the same criteria. Give weaknesses and opportunities. Do not give a probability. Do not give a loss figure or a return figure. Do not call a

scenario a forecast. Do not invent a current event."

Tools

Drivers and ranges: a shared OneNote page that the team fills in at the workshop.

Scenario drafts: Claude or Microsoft 365 Copilot, with the approved inputs attached.

Numbers: the validated model of finance, which the assistant does not touch.

Skills

Staff skill: tell a scenario from a forecast, and keep the decision criteria constant.

Agent skill: a scenario-draft skill with the labels for each statement and the forbidden figures.

Method

A two-driver scenario matrix, then a strategy test of one proposal in each scenario. The UNDP toolkit describes the two steps. [3]

Agent limits

Level 0. The scenario lead runs the workshop, and a person starts each draft. The assistant has no connection to financial models. It must not send scenarios outside the workshop group.

06 Challenge a business case before approval

Complete business proposal

Costs and source assumptions

Find omissions and double counts

Finance partner reviews

Questions for the sponsor

The assistant challenges a draft business case. It looks for unsupported demand, omitted costs, incompatible assumptions and benefits that the case counts twice. It traces each concern to its page and writes it as a question for the sponsor. It does not rewrite the case to make it pass.

The assistant needs the complete proposal, the cost definitions that finance approved, the calculation workbook and the source assumptions. Reviewers check every calculation in a tool made for that purpose, and never in generated text. A missing cost for implementation or supervision stays open, and so does a missing exit cost.

Put two faults into a fictional case: a duplicated benefit and a cost that someone moved outside the assessment period. Finance reviewers set the reference findings. Count the missed material issues and the unsupported objections, and record the review effort. Stop if the assistant edits the workbook or invents a return estimate. Stop if it hides a result that gets worse when an assumption changes.

The sponsor answers the questions, and authorised people approve the expenditure. A sharp challenge is not an audit opinion, and it gives no assurance of success. The finance business partner is accountable.

First working session

Put a duplicated benefit and a missing cost into a fictional business case.

06 Build notes: business-case challenge

Starter prompt

"You challenge a draft business case before approval. Use only the attached proposal and source assumptions, with the cost definitions from finance. Look for four faults: demand that has no evidence, costs that are missing, assumptions that cannot both be true, and a benefit that the case counts twice. For each concern, give the page and write one question for the sponsor. Check that the costs cover implementation and supervision, and check for an exit cost. Say which result changes most if one assumption changes. Do no arithmetic. Do not rewrite the case. Do not estimate a return. Do not hide a weak result."

Tools

Case documents: the proposal in Word or PDF, and a read-only copy of the Excel workbook.

Challenge: Claude or Microsoft 365 Copilot, with the documents attached.

Question log: a SharePoint list or Planner, with the sponsor as the owner of each answer.

Skills

Staff skill: find a benefit that two initiatives claim, and a cost outside the assessment period.

Agent skill: a case-challenge skill with the four faults and the question format.

Method

A red-team review with seeded errors. Add a pre-mortem question for the sponsor: "The initiative failed. What happened?"

Agent limits

Level 1. An agent can prepare the question list when a sponsor submits a complete case. It has read access only. It must not edit the workbook or approve the case.

07 Map execution dependencies before promising a date

Current initiative plans

Prerequisites and constraints

Map dependent milestones

Owners confirm links

Reviewed dependency map

To help a team assess a delivery date, the assistant reads approved initiative plans for prerequisites. Examples are data access, procurement, specialist capacity, system changes and pending decisions. The assistant proposes a map, with the supporting passage and the owner for each link.

The assistant uses current plans, agreed milestones, known resource constraints and confirmed decision rights. A missing prerequisite becomes a review item, and so does an owner that no one confirmed. Do not invent a delivery date to make the map look complete.

Test on a small programme that has a circular dependency and two projects that compete for one specialist team. Programme managers verify the map. Count the missed prerequisites and the unsupported connections, and record the corrections. The pilot ends if the assistant changes a schedule or assigns staff. It also ends if the assistant marks an unapproved dependency as resolved.

Owners use the map to negotiate with each other. Commitments and resource allocation stay with authorised people, and so do changes to delivery dates. The programme management lead is accountable.

First working session

Map a small programme in which two projects depend on the same specialist team.

07 Build notes: execution dependencies

Starter prompt

"You map the dependencies in a set of approved initiative plans. Use only the attached plans. List each prerequisite: data access, procurement, specialist capacity, system changes and pending decisions. For each link, give the two milestones and the passage that supports the link, with the owner. Write 'Owner not confirmed' where the plan names no owner. Mark each circular dependency. Mark each team that two projects need in the same period. Put missing prerequisites in a review list. Do not give or change a delivery date. Do not assign staff. Do not mark a dependency as resolved."

Tools

Plans: Microsoft Project or Jira, with read access.

Map: Claude or Microsoft 365 Copilot for the draft, and Visio for the reviewed map.

Open items: a RAID log in a SharePoint list.

Skills

Staff skill: confirm a link with the two owners before it goes on the map.

Agent skill: a dependency-map skill with the link format and the circular-dependency check.

Method

A dependency map with a check for circular links, then the critical path method on the confirmed links. Keep open items in a RAID log.

Agent limits

Level 1. An agent can prepare a draft map when a plan changes. It has read access to the plans. It must not change a schedule or an assignment.

08 Assemble a board evidence pack without suppressing disagreement

Requested board decision

Authorised papers and evidence

Options, questions and dissent

Secretariat checks completeness

Draft board evidence pack

From an authorised paper set, the assistant builds an index and a draft decision brief. The brief has separate sections for the requested decision, the supporting evidence, the alternatives, the unresolved questions and the dissenting views. Each substantive claim links to its source.

The minimum data is the final submission set, the board-paper rules, the access permissions and the requested decision. The brief flags a missing paper, and it flags a paper whose approval status is not clear. Never record silence in meeting notes as agreement.

The pilot checks that dissent survives a summary. Use a fictional pack that has conflicting estimates and a late revision. The secretariat checks completeness and source versions, and checks access. Count the material omissions and the unsupported statements of approval. Stop if a dissent disappears or privileged material reaches an unauthorised reader. Stop if the assistant records a decision that the board did not make.

The board secretariat lead is accountable. Directors keep their responsibilities. The authorised minutes stay the record of the decision, whatever the brief said.

First working session

Assemble a fictional board brief that contains dissent and a late source revision.

08 Build notes: board evidence pack

Starter prompt

"You build an index and a draft decision brief from a board paper set. Use only the attached papers. Make five sections: Requested decision, Supporting evidence, Alternatives, Unresolved questions and Dissenting views. Link each claim to its paper and version, with the page. Quote each dissenting view in the words of its author. List each paper that is missing or has no clear approval status. Where two papers give different estimates, show the two. Do not record agreement that the papers do not state. Do not record a decision. Do not remove or soften a dissent."

Tools

Papers: a restricted SharePoint library or the board portal of the firm, with Purview sensitivity labels.

Brief: Microsoft 365 Copilot or Claude in a restricted workspace, with access for the secretariat only.

Version check: Word Compare between the late revision and the paper before it.

Skills

Staff skill: check that each dissent in the papers is also in the brief.

Agent skill: a board-brief skill with the five sections and the dissent rule.

Method

A traceability matrix from each claim to its paper and version. Keep a dissent log, and check each entry against the brief.

Agent limits

Level 0. The secretariat starts each brief from the final submission set. No agent has continuing access to board papers. The assistant must not send the pack or write minutes.

09 Track benefits against a defined baseline

Signed-off benefit baseline

Periodic outcome evidence

Check definitions and causality

Benefit owner reviews

Qualified benefit findings

The assistant compares approved benefit definitions with periodic reports. It flags missing evidence and changed denominators, and it flags delivery milestones that a report shows as results. Completed activity stays apart from demonstrated benefit, and estimates stay apart from realised results.

Use the signed baseline, the measure definition, the owner, the period and the supporting records. If the baseline is absent, the report states that limit. It must not reconstruct a favourable start point. When a measure changes, keep its original definition next to the new one.

Pilot on one completed fictional initiative that has delayed benefits and a cost increase. Finance and programme reviewers check the conclusions. Count the incorrect benefit classifications and the unsupported attributions, and record the correction effort. The pilot stops if a report counts the same benefit twice. It also stops if a report calls released capacity a cash saving without financial evidence.

Attribution is the weak point. A trend over time cannot show alone that the initiative caused an improvement. For this reason, report the other changes that can explain it. The benefits realisation owner is accountable.

First working session

Check one fictional benefit report against its baseline and definition, and against its supporting evidence.

09 Build notes: benefits tracking

Starter prompt

"You compare a periodic benefit report with the approved benefit definitions. Use only the attached baseline and definitions, with the report. For each benefit, give the definition, the baseline, the reported value and the evidence. Give each item one label: Activity completed, Benefit estimated, Benefit demonstrated or No evidence. Say if the denominator or the definition is different from the baseline. If no baseline exists, write 'No baseline' and do not construct one. List other changes in the period that can explain the result. Do not say that the initiative caused a result. Do not call released time a cash saving. Do not count one benefit twice."

Tools

Baseline and definitions: a benefits register in a SharePoint list, with an owner for each benefit.

Reports: Power BI or Excel, with figures that finance checked.

Comparison: Claude or Microsoft 365 Copilot, with the register and the report attached.

Skills

Staff skill: tell completed activity from demonstrated benefit, and ask what else changed.

Agent skill: a benefits-check skill with the four labels and the "No baseline" rule.

Method

Benefits realisation management with a profile for each benefit: the definition, the baseline, the owner and the measurement date. Before a causal claim, list the rival explanations.

Agent limits

Level 1. An agent can prepare the comparison for each periodic report. It has read access to the register. It must not change a baseline or a definition, and it must not publish findings.

10 Maintain a regulatory horizon review queue

Official regulatory sources

Record status and dates

Status and scope clear?

Specialist review queue

Preserve unresolved flag

Yes

No

An assistant watches an approved list of official regulatory publications. It logs new consultations, final instruments and effective dates, with possible strategic questions. Each entry links to the original and proposes a review owner. The assistant never declares that a rule applies, because specialists decide that.

The inputs are the jurisdictions of the firm, its legal entities, the official source list and the definitions of publication status. Where the status or the date is not clear, the entry keeps an unresolved flag. The same applies to an unclear scope. Log a failed website check as a coverage gap, because the failure says nothing about a change.

Use a historical set that has a consultation and a final rule, with a superseded notice. Compliance specialists verify classification and completeness inside that defined source set. The pilot cannot establish coverage outside it. Count the missed items and the false effective dates, and count the incorrect entity matches. Stop if a proposal becomes a binding rule in the summary. Stop if the assistant certifies compliance.

The regulatory strategy liaison is accountable. Legal and compliance specialists decide which rules apply and what the obligations are. The strategy team considers only the implications that those specialists approved. The compliance register stays the record of obligations of the institution.

First working session

Classify a historical consultation and a final rule from official sources, with a superseded notice.

10 Build notes: regulatory horizon

Starter prompt

"You keep a review queue of regulatory publications. Use only the attached items from the approved official source list. For each item, give the title, the publisher, the publication date and the link. Give one status: Consultation, Final, In force or Superseded. Give the effective date only if the document states it. If the status or a date is not clear, write 'Unresolved' and say what is missing. Do the same for the scope. Suggest a review owner and one question for the strategy team. Do not say that a rule applies to the firm. Do not call a consultation a rule. Do not say that the firm complies."

Tools

Sources: the publications pages of the approved regulators, for example the Bank of Jamaica and the Financial Services Commission.

Scheduled check: Power Automate or n8n, with a log entry for each failed check.

Queue: a SharePoint list with the status and the review owner, and a link to the original.

Skills

Staff skill: read the status of a publication and find its effective date.

Agent skill: a horizon-queue skill with the four status labels and the "Unresolved" rule.

Method

Horizon scanning on a fixed source list, with four status labels and a coverage log. [3] Compliance specialists check a sample of the entries.

Agent limits

Level 1. An agent can check the approved sources on a schedule and add draft entries to the queue. It must not decide that a rule applies. It must not change the compliance register.

Choose the first pilot

These screening scores assume authorised document sets, fictional test cases, named reviewers, and no production changes or external commitments. The scores are our first assessments, and no company measured them. Score them again for your institution and your decision.

Case | Value | Feasible | Data | Effort | Risk

01 Assumptions register | 5 | 5 | 4 | Small | Moderate

02 Customer-need evidence | 4 | 4 | 3 | Medium | Higher

03 Competitor disclosures | 4 | 4 | 4 | Small | Moderate

04 Initiative comparison | 4 | 3 | 3 | Medium | Higher

05 Scenario discussion | 4 | 3 | 3 | Medium | Higher

06 Business-case challenge | 5 | 4 | 4 | Medium | Higher

07 Execution dependencies | 4 | 3 | 3 | Medium | Higher

08 Board evidence pack | 5 | 4 | 4 | Medium | Higher

09 Benefits tracking | 4 | 3 | 3 | Medium | Higher

10 Regulatory horizon | 4 | 3 | 3 | Medium | Higher

How to read the table

Value is the possible improvement in decision preparation or the possible decrease in rework. Feasible shows how well the firm can run a controlled test. Data shows how ready the evidence is. Each score is from 1 (weak) to 5 (strong).

Small effort is a bounded review of a source or a document. Medium effort adds specialist calculations, more than one owner, sensitive board material or a workshop. Risk rates the consequence of an error, and it does not rate the probability.

Use cases 01 and 06 are the first choices, together with use case 08. Use authorised documents and fictional test packs. Scenarios and rankings need specialist challenge, and so do regulatory implications. A high score changes no decision authority. A failed check on evidence or confidentiality still stops the use case.

Name an owner and a measure

Case | Accountable role | Pilot measures | Agent level

01 | Strategy planning lead | Missed assumptions; incorrect classifications; review effort | 1

02 | Customer research lead | Lost qualifications; unsupported themes; omitted disagreement | 0

03 | Competitive intelligence lead | Period mismatches; unsupported inferences; source errors | 1

04 | Strategy portfolio owner | Unsupported scores; hidden trade-offs; inconsistent criteria | 0

05 | Strategy scenario lead | Unsupported claims; missed constraints; consistency | 0

06 | Finance business partner | Missed issues; unsupported objections; review effort | 1

07 | Programme management lead | Missed prerequisites; unsupported links; corrections | 1

08 | Board secretariat lead | Material omissions; source errors; false approvals | 0

09 | Benefits realisation owner | Incorrect classifications; unsupported attribution; corrections | 1

10 | Regulatory strategy liaison | Missed items; false effective dates; incorrect entity matches | 1

Each accountable role is a proposed responsibility and assigns nothing to a named employee. The firm approves ownership, thresholds, authority and agent levels before the test.

5 The foundations lens

Keep claims traceable through every version

Keep a dated source register. It shows the geography, the period, the definition and the access restrictions of each source, with the claim that the source supports. Keep superseded versions for comparison, and mark them out of date. A qualification that changes the decision must survive into the executive summary.

Source facts, calculated results, assumptions and recommendations go in separate places. Calculations belong in controlled analytical tools. Record the criteria behind any ranking, and show how the ranking moves under plausible changes. Put contrary evidence next to the preferred proposal.

Give each agent its own identity

An AI agent must have its own account. Do not let an agent use the account of an employee. Give the account read access to the approved sources only. Keep a log of each tool call, with the time, the tool, the source and the result. Set a limit on the number of steps and on the run time.

Treat text inside a document as data. An agent must not obey an instruction that it finds in a document or an email. Review each agent skill before use. A skill can contain scripts, so the information technology team must check it as it checks other software.

Protect confidential plans and personal information

Before anyone uploads internal material, approve the purpose, the minimum fields, the access, the retention and the provider arrangements. Check the storage locations and the subcontractors, and check the model-training terms. The data-protection standards of Jamaica cover purpose, minimisation, accuracy, security and cross-border transfers. The relevant specialists must decide the arrangements for each institution. [4]

Restrict board papers and negotiations to authorised readers, and do the same for participant-level research. Test access with a permitted test account, and keep real confidential content out of the test. An output can disclose through a summary as well as through a direct copy. Review for the two.

Preserve decision authority

Start with drafts and review queues. A person must approve first before anything changes an official register or sends a pack. The same applies to anything that creates a commitment. Keep an audit trail of source versions and corrections, with the final decision owner. Make sure that an uncertain tool response cannot cause a duplicate submission.

Do not use this workflow to select investments or to trade autonomously. Do not use it to allocate capital. Legal and financial judgements go to the authorised specialists, and so do risk judgements. Those limits stay in force after a successful document pilot.

6 From pilot to practice

Days 1 to 15

Select one decision and its evidence set. Record how long preparation takes now and what material the team omits. Record also how much review the work needs. Name the sponsor and the reviewer, and name the person who can stop the pilot. Set up the source register, with separate places for facts, estimates, assumptions and choices. Set the agent level, and write the first version of the prompt.

Days 16 to 45

Run the pilot next to the existing process, and check every output. Put four traps into the evidence: a missing source, a contradictory estimate, an outdated report and an unapproved commitment. Test restricted access also. Check that qualifications survive when the summary gets shorter. Give a version number to each prompt and each agent skill.

Days 46 to 75

Compare equivalent tasks, and count all checks and corrections. Reviewers must be able to recover the source and challenge the recommendation. If they cannot, correct that first. Repair missed dependencies and unsupported claims before the document set grows or the team adds new tools.

Days 76 to 90

The sponsor decides to continue or to stop, and a decision to continue can include revisions. Record the monitoring owner and the source-update process, with the rollback route. Keep unresolved assumptions open in the decision record when the pilot ends.

Pilot decision canvas

Complete this canvas with the process owner before you start the assistant. Attach the approved source list and the test results to the decision record.

Strategic decision and accountable sponsor:

Legal entity and geography:

Planning horizon:

Sources, versions, definitions and access rights:

Known facts and assumptions:

Unresolved questions:

Permitted assistant actions:

Decisions reserved for authorised people:

Agent level and tool permissions:

Prompt version and agent skill version:

Baseline and comparison tasks:

Review effort:

Quality threshold and specialist reviewers:

Stop condition and escalation route:

Decision date and evidence required:

Measure value without inventing savings

Record the time for preparation, checks, correction and administration on comparable decisions. Count the unsupported claims, the material omissions, the missed dependencies and the source-version errors. A shorter pack is useful only if it keeps the information that the decision needs.

After the meeting, check that the approved actions and the benefit definitions are still clear. Do not treat board approval or a project launch as proof of commercial value. A completed milestone is not proof either. Keep

forecasts and observed results in separate columns, with a third column for the evidence that attributes a change.

Released time is capacity, and we count it as a cash saving only with evidence of lower expenditure. Include the costs of software, integration, training, source maintenance and supervision. For an agent, add the usage cost of each run. Continue only where quality and confidentiality hold. A good average time saving cannot offset one serious error in the decision record.

Record | What to enter

Current process | Preparation; review; correction effort; material omissions

Pilot process | Same measures on comparable decisions; all checking included

Costs | Software; integration; training; source upkeep; supervision; agent usage

Decision quality | Supported claims; visible uncertainty; missed dependencies

Decision | Continue (with or without revisions) or stop; owner; review date

References

[1] Caribbean AI Association. Think Tanks. Harbour programme remit. Accessed 26 October 2026. https://caribbeanaiassociation.com/think-tanks

[2] Caribbean Development Bank. 2026. Caribbean Economic Review and Outlook 2025–2026. Notes, page 8; Borrowing Member Country economic performance, page 12 and footnote 1. Preliminary 2025 data; information available at 16 February 2026. https://www.caribank.org/sites/default/files/publication-resources/CDB-AnnualRegionalEconomicReview-2025-26-Apr20.pdf

References continued

[3] United Nations Development Programme. 2026. Foresight for Strategic Planning Toolkit. Horizon scanning, pages 16–17; scenarios, pages 56–57; strategy testing, pages 106–108. https://www.undp.org/sites/g/files/zskgke326/files/2026-07/undp_foresight_for_strategic_planning_toolkit.pdf

[4] Office of the Information Commissioner, Jamaica. Data Protection Standards. Accessed 26 October 2026. https://oic.gov.jm/page/data-protection-standards

[5] Agent Skills. Agent Skills Overview. Accessed 10 November 2026. https://agentskills.io

Evidence and version

Evidence reviewed on 26 October 2026. Text revised and build notes added on 11 November 2026. Prepared by the Harbour Think Tank, Caribbean AI Association. Proposed workflows and prompts need local validation before deployment.