Caribbean AI Framework

A service test for the Caribbean

Caribbean AI Association · AI Framework Committee
Consultation edition · 27 November 2026

Contents

1. A service test for the Caribbean
2. Publication and consultation
3. Who should use the framework
4. Internet access across CARICOM
5. What a readiness rank can tell you
6. What Caribbean institutions have done
7. Where regional policy stands
8. Where the CTU and UNESCO fit
9. What the regional and global figures measure
10. 1. Measure the existing service
11. 2. Give a person authority to stop the trial
12. 3. Prove the service works during an outage
13. 4. Rehearse the supplier exit
14. 5. Share the test and sign the decision locally
15. 6. Check the grid, water supply and cost of closure
16. 7. Publish the results and corrections
17. Regional coordination and national decision authority
18. Data sovereignty and information governance
19. Infrastructure, access and continuity
20. Environmental, ocean and climate safeguards
21. Workforce transitions and regional capabilities
22. Use in public services and industry
23. Children, education and social well-being
24. Creative rights, heritage and local languages
25. Agentic AI, fraud prevention and cybersecurity
26. Scientific integrity and systemic risk
27. International policy references and Caribbean application
28. The first 90 days
29. How to report a result
30. The field guides and annexes
31. How to comment and what happens next
32. Sources and further reading

Figures and tables

Tables 1–2: connectivity and government readiness. Table 3: CTU recommendations. Table 4 and Figure 1: employment exposure and electricity demand. Tables 5–11: acceptance records for seven decisions. Tables 12–13: delivery and results.

A service test for the Caribbean

A flood warning can reach a phone and still send a family to the wrong river. A claims assistant can save a handler time by asking a customer for a document the insurer already holds. In both cases, the service record tells you more than a demonstration of the software.

This consultation draft gives a public office, business or community organisation a way to decide whether to use AI for one task. The decision belongs to the people responsible for that service. Neighbours can share test cases, instructions and measures. Each site must establish whether the tool works for its users, languages, records and rules.

Seven decisions, in order

1. Measure the existing service
2. Give a person authority to stop the trial
3. Prove the service works during an outage
4. Rehearse the supplier exit
5. Share the test and sign the decision locally
6. Check the grid, water supply and cost of closure
7. Publish the results and corrections

Start with the baseline and the right to challenge an error. If either is missing, a pilot cannot show whether people are better off. Facility review applies when the proposal adds a significant physical load.

Publication and consultation

CAIA Research & Policy Series · CAF 2026-01 · 27 November 2026

Author: Adrian Dunkley. Publisher: Caribbean AI Association, AI Framework Committee. CAIA publishes this independent proposal for discussion. It has no legal force. CARICOM, national governments, UNESCO and the CTU issue their own instruments.

How this edition was prepared

The 24 November edition and its linked guides were checked against the 27 November publication review and change register. The revision separates evidence from local proposals, puts seven decisions in order and moves the field material into linked annexes. The evidence includes regional policy reports, public administrative announcements, a research review and dated observations from the ITU series published by the World Bank.

The analysis uses published material. External reviewers, participating offices, translation sign-offs and completed CAIA pilots are not documented for this edition. A data gap is shown as a gap. The next results note should name contributors who consent to attribution and explain how their input changed a decision.

Reuse and citation

Suggested citation: Dunkley, A. (2026). Caribbean AI Framework: a service test for the Caribbean. CAF 2026-01, consultation edition, 27 November. Caribbean AI Association. Permanent publication page.

© 2026 Caribbean AI Association. Contact the publisher for reuse permissions. Third-party material remains subject to its original terms. No DOI has been assigned.

Revision 2. 27 November 2026 · Revision 2. Adds operating annexes and revises the country editions. Release and consultation record

Who should use the framework

Choose one decision a person already makes. A warning editor checks a place name. A claims handler checks which evidence is missing. A teacher decides whether a student should repeat a task. Where the proposed change has no clear owner or observable outcome, narrow it before spending money.

Mission

Help Caribbean communities test an AI service in the language, conditions and institutions where people will use it.

Vision

A resident can find out where AI affected a service, reach a person who can correct an error, and see whether the change delivered the benefit promised.

Geographic coverage

The directory covers 15 CARICOM members and 29 other jurisdictions selected for comparison in the earlier edition. The wider set includes territories and mainland neighbours. These counts describe the publication’s coverage. CARICOM publishes its own membership list. [12]

Five questions for a service owner

- Infrastructure: can the person complete the task during an outage?
- Data: who may use each record and for what purpose?
- Talent: who has the language and subject knowledge to check the work?
- Governance: who can correct an error, hear an appeal and stop the trial?
- Adoption: after rework and complaints, who saves time and who pays?

Internet access across CARICOM

A national average can guide the choice of test sites. It cannot tell you whether a patient can keep a connection long enough to submit a form. Test the complete task on the least reliable channel people use.

Table 1. Latest available internet-use observations by member

CARICOM member	Internet users (% of population)	Observation year
Antigua and Barbuda	72.7%	2024
The Bahamas	92.5%	2024
Barbados	70.4%	2024
Belize	80.0%	2024
Dominica	82.0%	2024
Grenada	70.2%	2024
Guyana	83.0%	2024
Haiti	47.9%	2024
Jamaica	90.1%	2024
Montserrat	No observation	—
Saint Kitts and Nevis	76.9%	2024
Saint Lucia	68.2%	2024
Saint Vincent and the Grenadines	76.2%	2024
Suriname	87.4%	2023
Trinidad and Tobago	82.2%	2024

Source: ITU, via World Development Indicators, November 2026 release. Values rounded to one decimal. A missing observation is not zero. Years differ, so the rows are not a ranking. [3]

What a readiness rank can tell you

The CTU reports the 2025 Oxford Insights government AI readiness ranks for 14 CARICOM members. Its table gives Jamaica rank 93 and Haiti rank 189. The unit is a country. Neither figure tests a single clinic, classroom or flood line.

Table 2. Government-level readiness, as reported by the CTU

CARICOM member	Global rank in the 2025 edition
Antigua and Barbuda	154
The Bahamas	126
Barbados	162
Belize	158
Dominica	176
Grenada	174
Guyana	144
Haiti	189
Jamaica	93
Montserrat	Not listed
Saint Kitts and Nevis	170
Saint Lucia	145
Saint Vincent and the Grenadines	167
Suriname	178
Trinidad and Tobago	122

Source: CTU Caribbean AI Task Force, July 2026, Table 2.2. This paper uses the rankings as reported and does not recalculate the index. [1]

What Caribbean institutions have done

These dated records describe events at named institutions. The proposed uses below are CAIA's interpretation.

Jamaica's power restoration

At the 6 November 2025 briefing after Hurricane Melissa, Jamaica Public Service reported that 55% of customers had power restored. The figure describes a point during recovery. It does not say which customers could use a particular app. [4]

Implication for a pilot. A warning team should rehearse a power loss with the sites that will issue and receive the message. Count the people reached by the fallback route.

Shared investment in the Eastern Caribbean

In June 2020 the World Bank approved an initial US$94 million for Dominica, Grenada, Saint Lucia, Saint Vincent and the Grenadines, and the OECS Commission. That amount was financing approved for a digital transformation project. [5]

Implication for a pilot. For a joint procurement, separate regional project costs from each site's training, integration and ongoing support.

A supervised test in Jamaica

The Bank of Jamaica's FinTech Regulatory Sandbox took effect on 16 March 2020. The Bank sets a process for controlled testing of eligible financial technology. [8]

Implication for a pilot. An insurer can adapt the discipline of a bounded test. A new claims tool still needs its own evidence, approvals and review of the rules that apply.

Where regional policy stands

These dated records describe events at named institutions. The proposed uses below are CAIA's interpretation.

A regional policy endorsement

On 7 July 2026 UNESCO reported that members of COTED-ICT had endorsed its Caribbean AI Policy Roadmap. The endorsement concerns the UNESCO roadmap. [2]

Implication for a pilot. Use the roadmap when setting a regional policy direction. Use a dated local decision record when approving a particular service.

Regional recommendations published in Trinidad and Tobago

The CTU launched its task force's final report at its July 2026 AI Forum. The report sets out seven recommendations, from shared digital infrastructure to regional representation. [1]

Implication for a pilot. The crosswalk in this paper identifies a concrete piece of evidence that a service team can contribute under each recommendation.

Where the CTU and UNESCO fit

The CTU makes seven regional recommendations. CAIA proposes the following records for local teams carrying them out. The local authority keeps its decision even if a neighbouring office supplied the test. [1] [2]

Table 3. Local evidence for the CTU's seven recommendations

CTU recommendation, in brief	CAIA record to prepare
1 · Shared infrastructure	Joint specification, site dependency map and export test.
2 · Harmonised governance	Named approver, applicable rules and appeal route.
3 · Staged economic adoption	Baseline cost, time after rework and a decision about expansion.
4 · Skills at several levels	Paid reviewer hours and a demonstrated task before sign-off.
5 · Public procurement and services	An acceptance test, access rights and an exit rehearsal.
6 · Warning systems and resilience	A drill with realistic language, power and network failures.
7 · Regional representation	A short evidence note recording disagreement and limits on transfer.

UNESCO reports that its roadmap drew on more than 1,000 stakeholders. That consultation informed UNESCO's publication. CAIA will record participation in its own consultation separately. [2]

Three international checks

OECD principle 1.5 supports recording who answers for the system. NIST AI RMF 1.0 organises work into Govern, Map, Measure and Manage. For a comparative legal example, Article 14 of the EU AI Act addresses human oversight for high-risk systems. Confirm separately whether that regulation applies to a particular deployment.

What the regional and global figures measure

Work and the digital divide

The ILO and World Bank estimated that generative AI could affect 26–38% of jobs in Latin America and the Caribbean. Exposure means some tasks in an occupation could change. The range cannot tell an employer how many posts to cut or how many hours a change will save. A service pilot needs a time study and a record of work transferred to other staff or customers. [6]

Electricity and cooling

The IEA estimated that data centres used about 415 TWh worldwide in 2024. Its 2025 base case projected about 945 TWh in 2030. An island utility needs a different measure: the site's coincident peak in MW compared with capacity available at the connection point. [7]

Table 4. Geography, units and type of estimate

Quantity	Scope	Use in a decision
26–38% of jobs	LAC, potential exposure (2024 study)	Select occupations for a task-level baseline.
415 TWh	Global annual consumption, 2024 estimate	Context for energy planning.
945 TWh	Global annual consumption, 2030 projection	Scenario only. Request a signed site load profile.

Human review

A systematic review by Goddard and colleagues examined how people can follow incorrect automated advice. Give reviewers access to the underlying record. Test whether they catch a deliberately incorrect recommendation in a controlled exercise. [11]

Figure 1. Global annual data-centre electricity demand (TWh)



Scale: 0–1,000 TWh. Source: IEA, Energy and AI (2025). [7]

1. Measure the existing service

A supervisor samples closed files from the desk, phone and online queues. Start the clock when a resident first asks for help. Include time spent chasing documents, correcting an answer and returning to the office. A system can look faster simply because it shifts work to the caller.

Decision record

Table 5. Evidence for recommendation 1

Responsibility	Proposed requirement
Accountable role	Service manager
Deadline	Before selecting a supplier
Planning cost class	Low
Evidence to retain	Baseline log and definition of a completed case
Indicator	Completed cases ÷ requests received, by channel. Show the counts as well as the percentage.
Acceptance target	Measure every channel in scope before trial.

Worked test · Check a claims-document request

Proposed sector exercise. Select redacted files with missing pages, unusual document formats and valid alternatives. Give reviewers the policy wording in force on the claim date. Draft a request against a duplicate file. A claims handler compares it with the evidence already held and records every correction before anything reaches a claimant.

Pause condition. A request delays a claim by asking for evidence already accepted, or the tool reveals another claimant’s details.

Sector guide and review procedure

2. Give a person authority to stop the trial

Give the duty officer access to the source record, the time to check it and the power to change an answer. A resident who challenges an outcome needs a contact, an acknowledgement and a decision they can understand. Test that route with an error whose answer is already known.

Decision record

Table 6. Evidence for recommendation 2

Responsibility	Proposed requirement
Accountable role	Service director
Deadline	Before the first live case
Planning cost class	Medium
Evidence to retain	Signed delegation, appeal log and a witnessed reversal
Indicator	Disputed decisions corrected ÷ substantiated errors. Record the wait until correction.
Acceptance target	Every consequential decision has a named reviewer and an accessible challenge route.

Worked test · Check a draft appointment instruction

Proposed sector exercise. Use approved administrative instructions, examples in local languages and changed appointment times. Use test records without patient identifiers. Draft a message in a separate queue. A designated clinician or clinic officer checks meaning and contact details before any use with patients.

Pause condition. The draft changes clinical instructions, identifies the wrong patient or sends a reader to the wrong service.

Sector guide and review procedure

3. Prove the service works during an outage

Run the same task on a basic phone, with a weak signal and with the main network disconnected. Put the paper or voice route in the hands of the people who will staff it. After restoration, reconcile duplicate records before resuming automated work.

Decision record

Table 7. Evidence for recommendation 3

Responsibility	Proposed requirement
Accountable role	Continuity lead
Deadline	Before launch and before the hurricane season
Planning cost class	Medium
Evidence to retain	Timed drill log, roster and reconciliation sheet
Indicator	Completed tasks by channel and minutes to restore a staffed route.
Acceptance target	Meet the local service deadline on the fallback route.

Worked test · Test outage reports when the network is under strain

Proposed sector exercise. Include cell sites with intermittent service, delayed calls and areas where few people could report a fault. Record the time of each observation. Group reports on a shadow incident board. A duty engineer verifies suspected outages and restores sites to the board when automated grouping hides them.

Pause condition. A community disappears from the incident board or a restoration notice precedes engineering confirmation.

Sector guide and review procedure

4. Rehearse the supplier exit

Ask the supplier to export a redacted case set, its history and configuration in a documented format. A different team should restore it without using the supplier’s dashboard. Price the transfer, the staff time and any period when two services must run.

Decision record

Table 8. Evidence for recommendation 4

Responsibility	Proposed requirement
Accountable role	Procurement lead
Deadline	Before signing or renewing
Planning cost class	Medium
Evidence to retain	Contract schedule, export manifest and restore test
Indicator	Records successfully restored ÷ records exported, including corrections.
Acceptance target	Reconcile every test record and identify who bears each exit cost.

Worked test · Check a reading exercise before it reaches pupils

Proposed sector exercise. Use an approved passage, its answer key and examples of locally accepted wording. Keep identifiable pupil work outside the trial. Teachers independently mark the draft exercise, try it with the accessibility tools pupils use and correct the answer key before classroom use.

Pause condition. An answer key penalises an accepted local response or a pupil record enters an unapproved tool.

Sector guide and review procedure

5. Share the test and sign the decision locally

A partner may borrow the questions and the scoring guide. It still needs local place names, a person authorised to approve the service and a record of the errors its own users encountered. Share counts and test instructions when the underlying records cannot travel.

Decision record

Table 9. Evidence for recommendation 5

Responsibility	Proposed requirement
Accountable role	Local accountable officer
Deadline	Before sharing a test or expanding it
Planning cost class	Low
Evidence to retain	Versioned test brief, permissions and local decisions
Indicator	Participating sites with a signed local decision ÷ sites proposing launch.
Acceptance target	Every participating site signs its own decision.

Worked test · Test credits on a small recording catalogue

Proposed sector exercise. Choose recordings with documented permission for this trial. Include shared authorship, variant names and tracks whose permissions exclude online release. Generate a draft credit sheet in a private workspace. Send each entry to the named rights holder for review before the curator changes the catalogue.

Pause condition. A recording goes beyond its agreed use or a credit appears publicly before approval.

Sector guide and review procedure

6. Check the grid, water supply and cost of closure

Request peak load in megawatts and annual consumption in megawatt-hours. Ask what happens on the hottest dry day and who pays for network upgrades. The closure estimate should cover equipment removal, battery handling, site restoration and the operator’s insolvency.

Decision record

Table 10. Evidence for recommendation 6

Responsibility	Proposed requirement
Accountable role	Project sponsor with utility and water-provider input
Deadline	Before a site commitment
Planning cost class	High
Evidence to retain	Dated utility study, water balance and closure estimate
Indicator	Peak MW as a share of local available capacity, water by source and funded closure cost.
Acceptance target	Resolve peak, dry-season and closure questions before commitment.

Worked test · Test how a utility prioritises leak reports

Proposed sector exercise. Use closed work orders, including repeat calls, low-pressure areas and faults reported during communications outages. Mask customer identifiers. Run a second queue alongside dispatch. An operator compares priorities with field evidence. Dispatch continues through the existing control room.

Pause condition. A confirmed urgent fault drops out of view or staff cannot recover the original queue during an outage.

Sector guide and review procedure

7. Publish the results and corrections

Compare the pilot with the same task, channels and case types in the baseline. Report staff and user time, missed urgent cases and unresolved complaints. If a rule changed during the trial, give its date. Explain the decision to expand, revise or stop.

Decision record

Table 11. Evidence for recommendation 7

Responsibility	Proposed requirement
Accountable role	Service owner
Deadline	Six months after launch
Planning cost class	Low
Evidence to retain	Dated six-month results note and correction log
Indicator	Completion, cost per completed case and correction time by channel.
Acceptance target	Publish the decision, measures and limitations together.

Worked test · Check a shift handover before the next crew starts

Proposed sector exercise. Choose past handovers from day and night shifts with open repairs, changed equipment states and contractor notes. Remove worker identifiers. Draft a handover from a copy of the log. Both crews compare the draft with the signed record before the supervisor approves it for the trial.

Pause condition. A draft presents an unresolved hazard as cleared, or a manager uses trial data to rank individual workers.

Sector guide and review procedure

Regional coordination and national decision authority

Caribbean countries can share evidence and negotiate common contract terms. Each government must still explain the decisions it makes under its own laws and remain answerable to the people affected.

Share the work countries repeat

When two public agencies consider the same AI service, they can ask for the same evidence about performance, security and data handling. Pool the test methods and share findings that contain no personal or protected information.

Name the authority that answers for each decision

Apply that principle to a joint purchase: agree on a common checklist, then require each participating authority to sign its own decision. The record should name the person who can suspend the service, the body that hears complaints, and who pays if the supplier fails. Publish the reason for the decision, with lawful redactions. Invite affected people and smaller jurisdictions to comment before terms are set.

Protect the option to leave

In deals with overseas governments, cloud providers or advanced AI developers, ask who controls the contract, who can change the service and what happens if supplies are restricted. Require an export test and a workable exit clause before signing. Use a shared negotiating position where it helps, without promising to buy as a bloc or treating political alignment as evidence that a service works.

A cross-border test

If two agencies use the same supplier, they can agree on an error definition. Match the cases by task and channel. Record any field a partner cannot lawfully share and exchange an aggregate or a locally created example instead. [1]

Field guide and related tests

Data sovereignty and information governance

Before an AI system reads a record, trace who can see it, which laws apply and how the organisation can take back control when a contract ends.

Map where the record goes

A hosted AI tool can send a prompt, uploaded file, usage log or backup to different companies and countries. Ask for a map of each stage, including access by the supplier’s staff and subcontractors. Record why each transfer is needed, which legal rules apply and which authority can investigate a complaint. A local server address alone does not answer who can access the records, reuse them to train a model or respond to an order from another jurisdiction.

Set a separate rule for sensitive use

Check the law in force where the service runs and for the sector using it. When public benefits, credit or insurance are at stake, identify the minimum information the system needs. Test whether it penalises people for absent records or language differences. Require a human review route, and let people correct relevant errors through the existing service.

Put permission limits into the contract

Use those limits in the purchase terms: separate permission to store or process records from any permission to train on them. Name who may grant an exception. Rehearse export, suspension and deletion before going live; ask for evidence of what can be deleted and what must be retained under law. For community archives, verify who can grant permission before reuse.

Permission for each use

A custodian may lend a recording to an archive while reserving the right to refuse voice imitation. Put access, reuse, training and deletion in separate entries. A supplier change reopens each permission. [2]

Field guide and related tests

Infrastructure, access and continuity

Measure the connection people actually have. Check peak electricity needs and the backup route before relying on an AI service or investing in a data centre.

Test beyond the capital

Ask people to complete a real service task on a basic phone, with weak reception and without a paid subscription. Include disabled people and readers who need text in a different language. Record which users could not complete the task and why. If the service handles essentials such as benefits or clinic appointments, keep a staffed alternative. Local tests establish how one service performs.

Compare the cost with the power available

A data centre’s annual electricity bill does not tell a utility what the site needs on its busiest day. Ask for peak site demand in megawatts, annual use in megawatt-hours, the cost and timing of grid upgrades, and who bears those costs. Compare a local build with leased regional capacity and smaller tools that may need fewer resources.

Make backup service part of the design

Run a drill in which power, internet or the hosting provider fails. Start with the most urgent service action, such as finding an existing appointment. Ask staff to switch channels and recover the records without exposing private data. Put recovery targets and the person in charge in the operating plan. Upgrade only when a tested option can meet the service standard.

A service during a storm

The restoration figure in Jamaica’s 6 November briefing covered customers with restored power. It did not measure internet access or service uptime. Record all three separately during the drill. [4]

Field guide and related tests

Environmental, ocean and climate safeguards

A power-saving design can still use more water. Set limits by site and season, then measure impacts on the communities and coast near the project.

Ask where the water comes from

Ask the proposer to state the source, withdrawal and consumption separately for a typical month and for a hot, dry period. State where the accounting starts and ends, over what period, and when and where water returns. Check who else uses the source, how much usable storage exists and what happens during restrictions. Dry and evaporative cooling can trade electricity use for water use. Use it to design questions for a bidder, not as a substitute for measurements at the proposed site. Use the same workload, weather and accounting period for every design in the comparison.

Follow the route from site to sea

Map stormwater, cooling discharge and fuel storage against waterways, settlements and sensitive coastal areas. The Cartagena Convention and its land-based pollution protocol offer wider-Caribbean context; check which instruments and local permit conditions apply to the site. Ask the responsible authorities to check sensitive habitats, flood risks and the discharge path before approval. Set monitoring points upstream and downstream where practicable. Publish the permitted limits and who will investigate an exceedance.

Count what the project adds beyond its fence

Apply it by separating the facility’s own electricity, water and emissions from those linked to power generation and equipment. State the boundary before comparing options; shifting consumption elsewhere does not make it disappear. Where AI is proposed for hurricane alerts or marine monitoring, check it against observations from the area and have staff rehearse how to respond when the model or sensor fails.

The site boundary

A global electricity projection does not answer the neighbour who asks about the well. Put the proposed intake, return flow, dry-season limit and monitoring point on the site plan. Record the organisation that will read each meter. [7]

Field guide and related tests

Workforce transitions and regional capabilities

Record a shift before AI changes it. Use that starting point to check what happens to pay and work quality. Count lasting technical roles separately from short-term jobs.

Find which tasks change

A call centre may use AI to summarise a customer call while staff still interpret the policy and make decisions. Ask workers which tasks disappear, which take longer to verify and which new tasks enter their shift. The International Labour Organization’s exposure index measures how much work might change; it does not predict layoffs.

Create reasons to stay

Training can increase a worker’s options without creating a local job that uses their new skills. Tie major public or publicly supported purchases to a plan for who will learn, operate and improve the system after launch. Negotiate paid placements where people can work on real tasks, agreed access to training material, and a mentoring commitment. If a university and a company share a project, write down who can reuse the code and publish the findings. Separate local technical work from jobs tied to construction or short contracts.

Make the gains and costs visible

Agree with managers and worker representatives on what happens to time released by a pilot. It could reduce the number of unresolved cases or give workers time to train. Compare correction time, quality and workload before and after. Monitor whether workers lose routes to gain judgment and experience, especially if entry-level tasks are cut. Use local hiring and retention evidence, and identify what remains unknown. Do not infer who will leave the country from a course completion or one employer’s results.

Work that moves

Follow a case from the first request through the final correction. A shorter central queue can coincide with longer visits at the branch. Ask staff who answer follow-up calls to record how much time each one takes. [6]

Field guide and related tests

Use in public services and industry

Test a specific service against the way it works today. A pilot should make the service better for the people who use it and leave a workable route when the software fails.

Give small teams a place to start

A guesthouse can draft answers from its own booking terms and have a staff member check refunds or storm closures. A farm cooperative can compare crop or rainfall summaries with field observations and official forecasts before circulating them. A public agency can try the same method with benefit enquiries: test against current rules, local language and incomplete records, then keep a staffed route for people the tool cannot help. Similar buyers can share test cases and still check their own sites, languages and rules.

Keep review close to consequential decisions

A bank can use AI to organise documents for assessment; an insurer can use it to sort claims for staff review. Neither workflow should turn an unverified inference into a refusal. Financial literacy tools can explain fees, fraud warnings and basic terms using public materials so people can make their own choices. Avoid personalised product recommendations. Clinics can first try reviewed administrative tasks; schools can start with teacher-checked materials. Clinical or grading decisions need stronger evidence and review by the qualified staff responsible for them.

Match safeguards to consequences

Check the rules that govern the service before a pilot and before wider use.

Who sees the wrong request?

The IAIS treats governance and conduct as relevant when insurers use AI. In a claims trial, give the handler the policy version that applied on the loss date. Compare the draft against evidence already accepted. [9]

Field guide and related tests

A review for consequential decisions

Use the assessment before approving a system that can change someone’s rights, safety or access. Record the legal authority, tests, staffed remedy, reviewer and decision to stop or proceed.

High-impact use and approval record

Children, education and social well-being

AI can explain a lesson or help organise everyday tasks. Schools and families should check whether it builds a person’s ability to act without it and protects their privacy.

Keep the part of the task a learner should practise

A secondary school class can ask students to predict how a change in bus fare affects a weekly budget. Give everyone the same figures. After students explain their first answer, let the tool offer a hint that the teacher has checked. Then give students a similar problem without AI. If students finish faster with the tool but cannot explain their choices in the second task, change how the tool is used. The UN Global Digital Compact calls for accessible digital skills in diverse languages and formats. Offer a paper or teacher-led version to students without a device or reliable connection.

Give children a way to reach a person

UNICEF’s guidance covers privacy, safety, fairness and children’s development. A school or youth service should minimise the personal information it collects, explain the tool in language children understand and provide a clear route to a trusted adult. Do not ask a chatbot to diagnose a mental health condition or act as the only crisis contact. If a youth-facing service allows personal conversation, check that a young person can leave easily and reach human support.

Treat claims about well-being with care

An OpenAI and MIT Media Lab research collaboration found mixed effects across its study conditions. Very high use was associated with self-reported indicators of dependence, an association that does not show that use caused it. Do not use session length as proof that a service is helping. Ask users whether the tool makes it easier to do what they chose, and make feedback optional.

A test for the teacher

Have the learner attempt a task independently after a supported exercise. Compare the reasoning as well as the grade. Agree with the teacher what remains private and who may see the feedback. [10]

Field guide and related tests

UNICEF guidance on AI and children

Creative rights, heritage and local languages

AI can help an artist edit audio or a community transcribe an archive. Decide who may authorise the use before giving a tool access to the material.

Separate the recording from permission to use it

Suppose a community archive wants transcripts of recorded stories in a local Creole. Start with a small collection that the archive and relevant speakers or custodians have agreed can be processed. Name who can listen to the recordings, whether the supplier may keep them, and who can read the transcript. Ask people who know the language to check names, changes in meaning and places where the software filled a gap. Make corrections possible before wider release. An agreement to make a transcript need not give a supplier permission to reuse a recording to train a model.

Record the terms before production

For music, film, design or performance, record the source, intended use, attribution, any agreed payment and a contact for disputes. When commissioning a synthetic voice, confirm permissions for the person’s voice and for any recordings used. Tell audiences when a synthetic performance could reasonably be mistaken for the real person. If a use involves culturally restricted material, consult the people who hold responsibility for it and respect their restrictions. WIPO describes several ways existing intellectual-property rules may apply to traditional cultural expressions; protections vary by jurisdiction. This chapter recommends an ethical practice and does not determine ownership or permissions under local law.

Return the work to its contributors

An archive can put this into practice by giving approved transcripts back to the community and making sure contributors can correct them. A project agreement should say what the community receives and what happens when a vendor changes its terms.

The place name people hear

Have a listener repeat the action after hearing the warning once. Record incorrect locations separately from incorrect actions. A fluent sentence counts as a failure when either error changes what a resident would do. [2]

Field guide and related tests

Agentic AI, fraud prevention and cybersecurity

An AI agent can use tools to act. Give it the smallest permission it needs, and keep a person able to check and stop consequential actions.

Start with read-only help

An insurer can let an agent prepare a fraud referral or group duplicate alerts for an analyst. The analyst can see the records and verify a discrepancy before using it. A fraud flag should prompt checking; it should not be treated as proof. Do not let an agent freeze an account, reject a claim, send a payment or disclose records on its own. Verify unexpected payment or account-change requests through an established contact channel before taking action.

Check what the agent can do with its permissions

Test the permission boundary on the actual tools and data routes before relying on it. In a test using approved sample records, check that the agent leaves an ambiguous case in the queue and cannot change the record or pass information outside the authorised system. Treat instructions inside messages and documents as untrusted. Keep a record of approvals and make sure a named operator can remove the agent’s access. If the model or connected tools change, repeat the relevant checks.

Plan for an incident and a return to service

CARICOM’s updated Cyber Security and Cybercrime Action Plan provides a regional coordination reference. Use local reporting rules and approved routes. In a tabletop drill, have a named lead pause the workflow, preserve the relevant records and restore the manual queue without losing unresolved cases.

Who can reverse an action?

A permission to prepare a request differs from permission to issue it. During a controlled rehearsal, have a second person check the proposed action and then withdraw the permission. Keep the audit record. [10]

Field guide and related tests

Stop the action and keep the service open

The incident annex gives duty leads a severity guide, notification record and restart checks. It asks each service to record the deadlines and competent authorities that actually apply.

Incident response and reporting · Risk register

NIST Generative AI Profile · Singapore model framework for agentic AI

Scientific integrity and systemic risk

Use AI to ask useful scientific questions, and measure its performance before relying on the answer. Prepare for disruptions that could affect several services at once.

Compare results against observations

A Caribbean climate team can compare an AI rainfall model with its current forecast and quality-checked observations from storms the model did not train on. Report the places and events where the model performs poorly, including gaps in station coverage. Publish enough information about data and tests for another team to check the result where licences permit. Early research can inform decisions about what to test next. An experimental forecast should not replace an authorised public warning.

Check who shares the same weak point

A common model or supplier can make failures spread: an outage could interrupt tools used by several public agencies. This is a systemic risk, because the harm extends beyond one team. List shared suppliers and demonstrate that an essential service can switch to a manual process or an independent backup. Provide regulators and authorised partners with incident information through agreed channels. The UN Global Digital Compact calls for inclusive global cooperation on AI; the EU AI Act provides a separate legal reference for systemic-risk obligations for certain general-purpose AI models within its scope.

Use dated evidence when planning for extreme risk

The February 2026 International AI Safety Report describes disagreement about whether loss of control could lead to human extinction. It does not establish a probability. It found early signs of relevant capabilities in the systems it reviewed, below the level it judged sufficient for loss of control. This finding applies to the systems reviewed at that time. Rehearse how to suspend consequential actions, and address documented harms such as fraud, exclusion and service outages.

Dependencies across services

List which essential services depend on the same identity provider, data feed or network. A shared backup on that same dependency will fail at the same time. Test an independent route. [1]

Field guide and related tests

International AI Safety Report 2026

International policy references and Caribbean application

Human rights guide what the team needs to protect, and local law determines what it can do. Published risk frameworks can help the team decide what evidence to collect.

UNESCO and the OECD: rights and accountability

Its Caribbean roadmap turns those concerns toward regional capacity, culture and resilience. A Caribbean purchaser can apply both by involving people affected by a service, giving them a route to challenge a decision and naming the officer who can change it.

NIST and Singapore: evidence before permission

The US National Institute of Standards and Technology's voluntary AI Risk Management Framework helps teams organise risk work, identify who could be harmed, test the system and manage failures over time. Singapore's agentic AI framework asks teams to limit the actions an agent can take and assign meaningful human oversight. Together they suggest evidence to collect, not a guarantee that a system is safe.

The EU, the UN and CARICOM: different roles

The EU AI Act is a regulation whose requirements depend on scope, role and use. Its approach to matching duties to the seriousness of harm can inform the region; verify applicable laws and cross-border obligations separately. The UN Global Digital Compact calls for wider participation and cooperation. It leaves operating choices to the bodies implementing it. The CARICOM Single ICT Space roadmap and CTU's Caribbean AI Task Force add regional context. Governments could share test methods and specialist capacity while each competent authority retains its decision under its own mandate.

The authority behind a rule

The UNESCO roadmap supplies regional direction. NIST provides a voluntary risk-management process. A local law, permit or contract governs a specific obligation. Write the jurisdiction and applicable provision next to each duty. [10]

Field guide and related tests

Canada Algorithmic Impact Assessment · EU AI Act Article 27 on fundamental rights impact assessment · Singapore model framework for agentic AI

The first 90 days

Time starts when a service office accepts responsibility for a proposed pilot. No organisation is recorded here as having agreed to take part.

Table 12. Proposed delivery timetable

When	Who acts	Decision and record
Days 1–15	Service manager and frontline staff	Select a task. Record who uses it, which channel they use and how to correct an error.
Days 16–30	Operations analyst and user representative	Sample the current service. Agree the denominator and which errors would require an immediate pause.
Days 31–60	Trial lead and independent checker	Run matched cases in shadow mode. The existing service continues to make decisions. Log disagreements and missing evidence.
Days 61–75	Continuity and procurement leads	Witness a fallback drill and a supplier export. Reconcile the restored records.
Days 76–90	Local accountable officer	Publish a reasoned decision to stop, amend or move to a bounded live trial. Record objections and follow-up dates.
Six months after live launch	Service owner	Publish outcomes, costs, corrections and the decision on continuation.

Planning assumptions

Low cost means existing staff can complete the work in a scheduled review. Medium cost requires dedicated staff time, paid participation, training or a drill. High cost requires a site study, capital expenditure or specialist assessment. These are planning categories. Each sponsor should attach its own estimate in local currency and identify who pays.

Who would pay for a pilot

The delivery register names proposed owners, what each use needs to cost and how to measure results. Record the sponsor, funding source and signed decision before entering a start date.

Delivery and consultation register

How to report a result

A faster average can hide a worse service. Publish the number of eligible requests, cases reviewed, completed tasks and unresolved cases. Keep the original result when a correction is made so the reader can trace the change.

Table 13. Six measures for the results note

Measure	Calculation and reporting rule
Completion	Tasks completed ÷ attempts. Report the numerator and denominator for each channel.
Error	Substantiated errors ÷ reviewed cases. Separate missed urgent cases from routine corrections.
Time	Report the median time from request to completion and the slowest tenth where sample size allows.
Cost	Staff, supplier and rework cost ÷ completed cases. Report user travel or telephone cost separately.
Remedy	Days to acknowledge and resolve a challenge. Include open challenges at the cut-off date.
Access	Compare channels and languages where lawful. Suppress small groups where a person could be identified.

A worked calculation

Illustrative figures: if 42 of 50 requests reach completion on the staffed route, completion is 84%. If eight cases remain open at the reporting date, keep them in the denominator. The example supplies arithmetic only. It gives no evidence about a Caribbean service.

A participant should set its baseline before testing. Record whether a target is a mandatory safeguard, such as a working appeal route, or a performance goal that stakeholders may revise.

The field guides and annexes

The annexes are maintained online so that a regulator link, test case or stop procedure can be corrected without changing the numbered recommendations.

Annex A. Risk register

Filter example risks by tier and accountable role. Download the register to assign a local owner, score evidence and set a review date. Open risk register

Annex B. Data centre conditions

Record peak load, water source, storm continuity, community impacts and closure arrangements. Open data centre conditions

Annex C. Sector guides

Each of the 12 guides links a bounded pilot with its review and stop procedure. Open sector guides

Annex D. Country pilot notes

Use the 15 CARICOM notes first. The wider directory keeps the other 29 discussion notes available for comparison. Open country pilot notes

Topic guides · Language summaries · Glossary · Quick read

Find the records for this revision

The register contains 123 scenarios. The approval, incident and delivery templates are available for local service teams.

Risk register · Approval record · Incident log · Delivery plan

How to comment and what happens next

Send comments to info@caribbeanaiassociation.com by **31 January 2027**. Use “Caribbean AI Framework consultation” as the subject. Name the section, the service or jurisdiction you know, the change you want and a public source if one is available. Please leave names of clients, pupils and patients out of examples.

Questions for readers

- Which of the seven decisions would you change, and what would the service owner do differently?
- Which regional figure, rule or example needs correction? Give the date and source.
- Which service could test a recommendation, in which language, and who would have authority to stop?

Publication schedule

The proposed response date is 28 February 2027. The committee should group submissions by issue, publish the decision on each proposal and give a reason for deferral. Comments may be summarised without identifying the writer. Seek consent before publishing a name or an attributed quotation.

Participation record

Suggested participants include service users, disability advocates, language practitioners, frontline staff, sector regulators, utilities and national digital offices. No invitation or endorsement is recorded in this edition. Updates should distinguish invitations, replies and verified pilot results.

Terms

Baseline: the measured performance of the current service. Shadow mode: a trial in which proposed outputs cannot change a live decision. Stop authority: the person empowered to suspend the service change. Fallback: the route people can still use when the main service fails.

Track the response

The public record gives the timetable and a template for recording each decision on a submission. Add named independent reviewers and language sign-off when the committee has their evidence and permission.

Delivery and consultation record

Sources and further reading

Sources were consulted on 27 November 2026. Internet-use years appear beside the observations. External evidence describes the scope and date stated by its publisher.

1. Caribbean AI Task Force (2026). Toward Harmonized AI Policies and Recommendations for the Caribbean. CTU, pp. 15–16, 30–40.
2. UNESCO Office for the Caribbean (7 July 2026). UNESCO’s Caribbean AI Roadmap wins regional backing. United Nations Caribbean.
3. International Telecommunication Union. Individuals using the Internet (% of population), IT.NET.USER.ZS. World Development Indicators, November 2026 release. Observation years vary.
4. Jamaica Information Service (November 2025). JPS restores power to 55 per cent of customers. Briefing of 6 November.
5. World Bank (22 June 2020). First-time financing by World Bank for digital economy in the Eastern Caribbean approved for US$94 million.
6. ILO and World Bank (2024). Buffer or Bottleneck? Employment Exposure to Generative AI and the Digital Divide in Latin America.
7. International Energy Agency (2025). Energy and AI. Executive summary, 2024 estimate and 2030 base case.
8. Bank of Jamaica. FinTech Regulatory Sandbox. Effective 16 March 2020.
9. International Association of Insurance Supervisors (2 July 2025). Application Paper on the Supervision of Artificial Intelligence.
10. NIST (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1.
11. Goddard K, Roudsari A, Wyatt JC (2012). Automation bias: a systematic review of frequency, effect mediators, and mitigators. Journal of the American Medical Informatics Association, 19(1), 121–127.
12. CARICOM. Member States and Associate Members. Membership directory.

Earlier detailed policy · 24 November 2026 PDF · Field guide source catalogue