OCTOBER 2026

BARBARA-CHIARA UBALDI
SOFIA ARON
JOHAN HARVARD
OTT VELSBERG
BENJAMIN WELBY



AI Governance for Delivery: A Practical Guide for Latin America and the Caribbean

TONY BLAIR INSTITUTE FOR GLOBAL CHANGE

Contents

3 Executive Summary

6 Where Does Your Team Fit in the AI-Governance System?

11 Understanding the Government AI-Governance Environment

20 What Does Your Organisation Need to Have in Place?

29 How Should Your Team Govern an AI System Through its Lifecycle?

56 Conclusion: Putting the Lifecycle Into Practice

Executive Summary

Across Latin America and the Caribbean, public institutions are already using or experimenting with artificial intelligence to improve services, support public servants, strengthen decision-making and increase administrative capacity. At the same time, governments are developing the policies, institutions and capabilities needed to govern these technologies responsibly. In many cases, however, implementation is outpacing governance, leaving delivery teams to navigate new technologies while the rules, processes and institutional responsibilities governing their use continue to evolve.

This paper provides a practical guide for the teams responsible for turning AI ambitions into working public-sector solutions. These teams may sit in ministries, agencies, municipalities, digital-government units or other public institutions. They may be responsible for identifying a use case, designing or procuring a system, overseeing implementation or managing a system already in operation. Their common challenge is practical: how do you govern an AI system responsibly while still delivering?

AI governance involves more than laws and policies. It is a system of rules, institutions, decision-making processes and controls that enable governments to develop and use AI responsibly. A functioning governance environment connects these elements with accountable owners, decision rights, operational processes, capabilities and assurance mechanisms, and these need to work together throughout the AI lifecycle.

One of the most important questions delivery teams should ask is where governance decisions sit. Some are made across government, some by individual organisations, and others depend on the characteristics and consequences of a specific AI system. Understanding these levels helps teams know what they inherit, what their organisation should provide, what they must decide themselves and when issues need to be escalated.

The paper sets out our recommendations in four sections:

1. **Where delivery teams fit within the AI-governance system** – differentiating between whole-of-government, organisational, and use-case and operational governance. This helps them determine what is already in place, what needs to be established for the specific use of AI over time and where they may need to seek support or escalate a decision.
2. **The wider national governance environment** – from international principles and standards to their implementation across Latin America and the Caribbean, noting that countries are progressing at different speeds, with differing levels of guidance specific to their country.
3. **AI governance at an organisational level** – delivery teams should not have to create governance processes from scratch for every AI initiative. Organisations must provide a minimum set of functions that make responsible delivery possible, including clear ownership and decision rights, visibility over AI use, reusable controls, proportionate assurance, access to relevant capabilities and lifecycle oversight. These functions do not require a single organisational model or a dedicated AI office; they can be provided through existing functions, shared government services or external expertise.
4. **AI governance at use-case and operational level** – where governance becomes part of day-to-day delivery. The practical core of the guide is a ten-part AI lifecycle that takes teams from identifying the problem and deciding whether AI is appropriate, through classification, assessment, approval, procurement or development, testing and deployment, to monitoring, change, reapproval and retirement, with country-specific case studies. The lifecycle is not intended to be linear: teams may need to return to earlier decisions when new evidence, risks or material changes emerge. Governance should also be proportionate to the potential consequences of the use case rather than applied uniformly to every system.

Focusing on how public institutions should govern the AI systems they develop, procure and use, this paper provides delivery teams with a practical way to understand the governance environment around an AI system and the decisions they need to make throughout its lifecycle. It gives teams a

way to make decisions with the right evidence at the right point in the lifecycle, while creating mechanisms for oversight, learning and adaptation over time. Where governance arrangements already exist, teams should use them. Where they are incomplete, teams should identify the gap, draw on established principles, standards and tools, and escalate questions that cannot appropriately be resolved at the use-case or operational level.

The objective is not to propose a single governance model for governments in Latin America and the Caribbean. It is to advise delivery teams on how to navigate the governance environment they have, understand what it provides, recognise what is missing, and make proportionate decisions that allow AI initiatives to move from experimentation to responsible and sustainable implementation. This reflects the broader message of the paper: countries and organisations do not need to build everything from scratch, but they do need to turn governance principles into systems that work in practice.

01

Where Does Your Team Fit in the AI-Governance System?

For delivery teams, understanding AI governance is crucial because they are responsible for making decisions throughout an AI-initiative's lifecycle, often within governance arrangements that are spread across government, their organisation and the project itself. Governance is not a separate compliance exercise that happens around an AI project. It shapes how a project is selected, designed, procured, tested, deployed and monitored, and determines who is accountable for decisions throughout that process. Having a clear understanding of what governance consists of and why it is critical helps teams identify the requirements, processes and decisions that need to be in place to move an AI initiative forward responsibly.

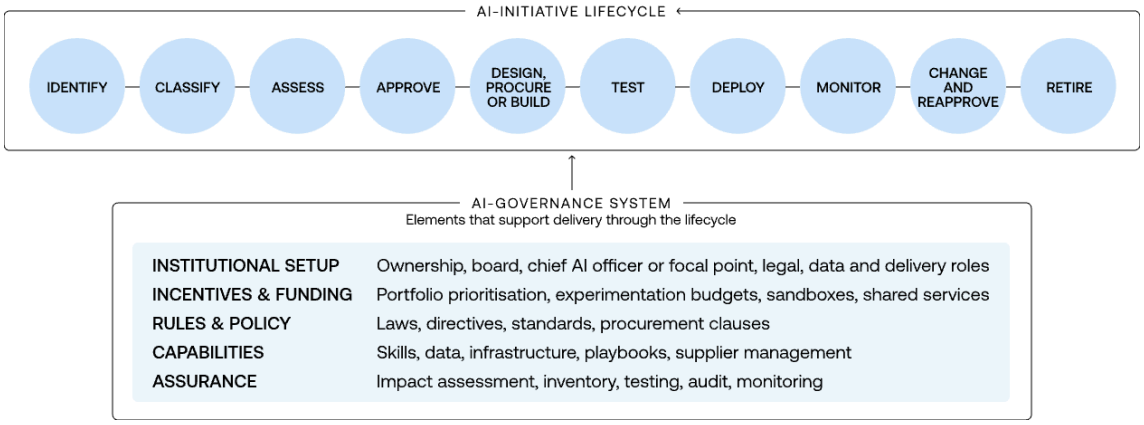
AI governance is the system of rules, institutions, decision-making processes and controls that enables governments to develop, procure and use AI responsibly. A functioning governance system connects rules and policies with accountable owners, decision rights, operational processes, capabilities and assurance mechanisms. These elements need to work together throughout the AI-initiative's lifecycle: rules establish expectations; institutions and accountable owners assign responsibility; processes determine when decisions and approvals are required; and assurance, monitoring and evidence help manage risks, test whether expected benefits and public value are being realised and improve governance over time.

For each use case, teams should be able to state the intended outcome and identify an accountable owner. The baseline, success measures, benefits review and level of documentation should be proportionate to the use case's consequences and the resources committed to it. Reviews should consider operating and oversight costs, who benefits or may be disadvantaged, and whether AI improves on a feasible alternative. This evidence should inform decisions to continue, improve, scale, restrict or retire the use case.

The figure below illustrates the main elements of an AI-governance system and how they support the AI-initiative lifecycle. This will help delivery teams see governance as an operating system rather than a collection of individual requirements. A project may depend on several of these elements at the same time: for example, a team may need a clear policy requirement, an accountable owner, the capabilities to implement the system and an assurance process that generates evidence before deployment.

FIGURE 1

The AI-initiative lifecycle and how a governance system can support it

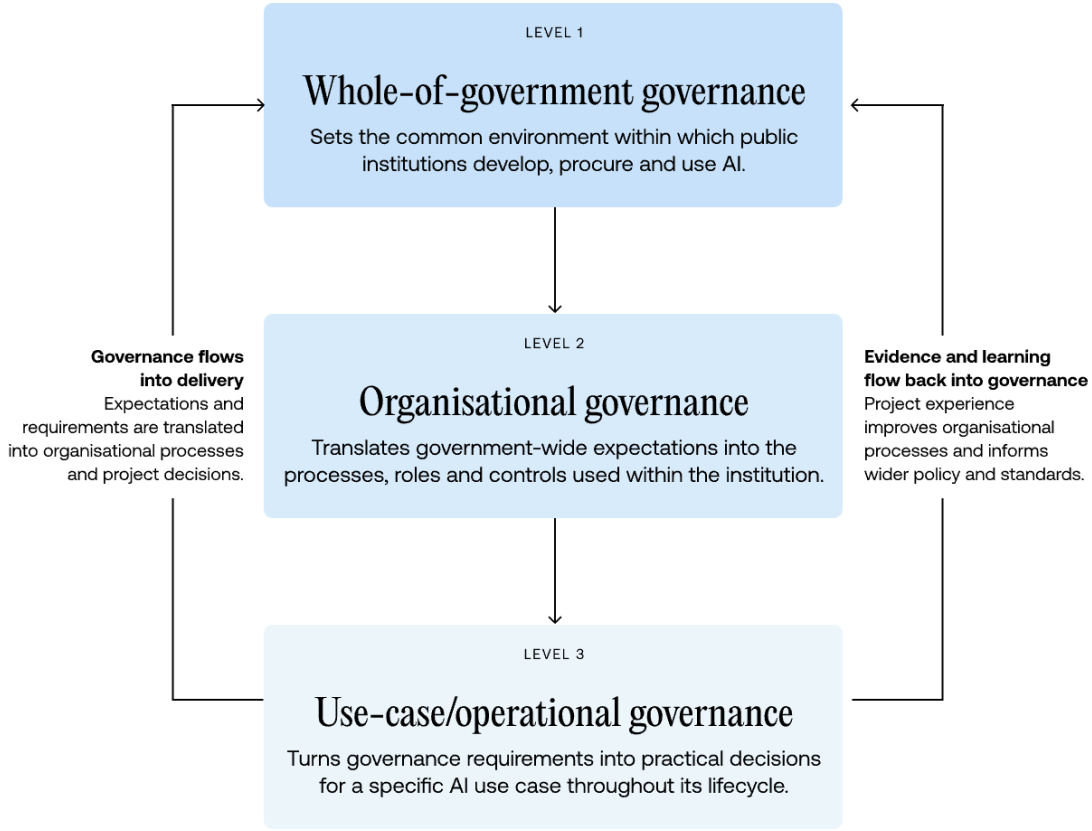


Source: TBI

These elements, however, are not all decided or managed by the same actors. They operate across three connected levels of governance: whole-of-government governance, organisational governance, and use-case and operational governance. Delivery teams need to consider the components of the governance system and the level at which decisions are made. This helps them determine what is already in place, what needs to be established for the specific use of AI over time, and where they may need to seek support or escalate a decision.

FIGURE 2

Three connected levels of governance



Level 1 – Whole-of-government governance: What has government decided? The governance applicable to the whole of government establishes the common environment within which public institutions can deploy, procure and use AI. This can include laws and regulations, national public-sector AI strategies and policies, government-wide common principles, standards and requirements, institutional mandates, and oversight arrangements. These may be established at national or subnational level and may be specific to the public sector or form part of broader sectoral or economy-wide frameworks. Common risk classifications, registers and transparency requirements, standards for procurement, data and security, as well as shared tools, services and capabilities also form part of the governance. For delivery teams, the key

task is to understand which requirements and common services already exist and apply them. They also need to identify which institutions have decision or oversight responsibilities, and whether project-level questions need to be escalated.

Level 2 – Organisational governance: What has the institution put in place? Organisational governance translates government-wide expectations into the processes, roles and controls used within a ministry, agency or other public institution. These arrangements can cover ownership and accountability, decision rights and approval routes, and common risk and assurance processes. They may also include procurement, data and security requirements, AI registers and standard documentation. Delivery teams need to understand which internal processes apply, who needs to be involved, what evidence is required and when to seek additional review or escalate a decision.

Level 3 – Use-case and operational governance: What does this use of AI require? System or use-case governance applies broader requirements to a specific use of AI and the system that supports it. The appropriate approach depends on the use case's purpose, context and potential consequences. It includes decisions about intended use and limits, risk classification and assessment, applicable national, sector-specific and organisational requirements and standards, data and suppliers, testing, human oversight, transparency, monitoring, incident management, changes, reapproval and retirement. For delivery teams, the key task is to make and document proportionate governance decisions for the individual AI system throughout its lifecycle, including once the system becomes part of ongoing service operation. Govern the use case, not only the technology. The same AI model or product can support use cases with very different consequences. A general-purpose model used to summarise internal documents does not require the same governance as the same model used to influence eligibility, enforcement or another consequential public decision. Classification and assurance should therefore normally be applied at the level of the use case and its operating context, while also taking account of risks inherited from the underlying model, supplier and technical architecture.

The three levels are distinct, but they should not operate in isolation. Decisions made at the whole-of-government level need to be translated into organisational processes before they can be applied consistently to individual use cases. In turn, experience from individual projects can expose recurring barriers, new risks or gaps in existing governance arrangements. Governance should therefore work in both directions: requirements move from broader frameworks towards implementation, while evidence and learning from implementation inform how those frameworks evolve.

This feedback loop is particularly important as governments build AI capability through experimentation. Lessons from early pilots should be incorporated into broader national AI strategies and institutional arrangements as governments move from experimentation towards scale.

For delivery teams, this creates a simple practical discipline:

- Apply the laws, standards and organisational requirements that have already been established.
- Decide what needs to be tailored to the purpose, risk and context of the individual AI system.
- Escalate questions that cannot appropriately be resolved at the use-case or operational level, or that reveal a gap in organisational or whole-of-government governance.

The objective is not for every delivery team to become an AI regulator or design its own governance framework. Nor should teams assume that compliance with a national policy is sufficient to govern an individual system responsibly. Effective AI governance depends on each level performing its appropriate role and on delivery teams understanding which decisions are theirs to make, which requirements they need to apply and when a question needs to move beyond the project team.

02

Understanding the Government AI-Governance Environment

Before making decisions about an AI initiative, government delivery teams must understand the government AI-governance environment in which the project sits. There is no single model of AI governance: countries translate international principles into different combinations of laws, policies, standards, institutional arrangements and oversight mechanisms. These differences shape the requirements delivery teams may need to meet, the institutions they may need to engage and the guidance available. This section explains where these national approaches come from, how they differ and what this means for teams implementing AI projects.

How government AI governance takes shape: from global standards to national frameworks

National AI-governance frameworks do not emerge in isolation. Governments draw on international principles and standards that guide AI development, procurement and use, and adapt them to domestic laws, policy priorities and institutional capacities. This shared foundation helps explain both the similarities between national approaches and the differences in how they are implemented.

Common themes include human rights, fairness, transparency, accountability, safety and public value. These principles do not provide detailed operational requirements, but they often inform national and organisational governance frameworks.

Figure 3 introduces selected examples of international frameworks to help delivery teams understand the reference points that often shape national AI governance. It is intended to provide familiarity with key principles rather than a comprehensive inventory.

FIGURE 3

Key international AI-governance frameworks

Framework	Nature and relevance for government teams
OECD AI Principles 2019, updated 2024	A global reference point for trustworthy AI, first adopted in 2019 and updated in 2024. They combine value-based principles with policy recommendations and have influenced several later international processes.
UNESCO Recommendation on the Ethics of AI 2021	A broad global standard on AI ethics, adopted by UNESCO member states. It is particularly relevant for governments seeking to connect AI governance with rights, inclusion, institutional readiness and ethical impact.
G7 Hiroshima AI Process 2023	A set of principles and a voluntary code of conduct focused especially on advanced AI systems. It is relevant for governments because many public-sector tools rely on systems developed by private-sector actors operating across borders.
Council of Europe Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law 2024	The first international legally binding treaty specifically addressing AI. It connects the development and use of AI with human rights, democracy and the rule of law, including requirements around risk and impact management, transparency, procedural safeguards and remedies. Particularly relevant to public-sector uses of AI that may affect people, public services or the exercise of public authority.
UN General Assembly Resolution on AI 2024	A global political statement on safe, secure and trustworthy AI for sustainable development. It reinforces the idea that AI governance is not only a technical issue, but also a development and public-interest issue.
G20 AI Principles 2019	Principles drawn largely from the OECD framework and increasingly shaped by the priorities of emerging and developing economies. They are relevant because they connect AI governance to inclusive growth, development and international cooperation.

Source: TBI

International technical and management standards provide specific requirements, processes and assessment criteria. Developed by bodies such as ISO and IEC, ITU, and IEEE, they can inform procurement specifications, development processes and assessments of systems or suppliers.

Figure 4 presents selected examples to help delivery teams understand how international standards can inform national requirements and guide project implementation. The aim is to familiarise teams with useful reference points, not to provide an exhaustive list or suggest that every standard applies to every project.

FIGURE 4

Examples of international AI standards relevant to implementation

Illustrative standard	Nature and relevance for government teams
ISO/IEC 42001:2023 – AI-management systems	Establishes requirements for setting up, maintaining and improving an organisation's AI-management system. Helps teams understand the organisational processes supporting responsible AI, rather than assessing an individual application's technical performance.
ISO/IEC 23894:2023 – Guidance on AI risk management	Provides guidance on identifying and managing AI-related risks and integrating risk management into organisational activities. Can inform the approach used to assess and manage risks within individual projects.
ISO/IEC 5338:2023 – AI-system lifecycle processes	Defines processes for managing AI systems throughout their lifecycle. Provides a reference for organising development and operational activities, rather than treating governance as a one-off review before deployment.
IEEE 7001-2021 – Transparency of autonomous systems	Defines measurable and testable levels of transparency. Helps teams translate a general commitment to transparency into criteria against which a system can be assessed.
IEEE 3119-2025 – Standard for the Procurement of Artificial Intelligence and Automated Decision Systems	Provides a structured process for procuring AI and automated-decision systems, including problem definition, solicitation preparation, vendor and solution evaluation, contract negotiation and contract monitoring. It helps procurement teams incorporate AI-specific risk, transparency, accountability and governance requirements into the procurement lifecycle.
ISO/IEC 42005:2025 – AI-system impact assessment	Provides guidance for assessing the potential impacts of AI systems on individuals and society throughout the lifecycle. Particularly relevant to teams conducting impact assessments and documenting how identified impacts, risks and mitigation measures are addressed.

Source: TBI

How do governments translate AI principles into practice?

Although many countries draw on the same AI principles, there is no single approach to national AI governance, which must reflect local contexts and be fit for purpose. As a result, governments translate international principles into different combinations of strategies, policies, institutional arrangements, laws and implementation tools, based on their national priorities and capacities.

These choices shape the requirements delivery teams need to meet, who is responsible for oversight and what support is available for implementation.

These differences are often reflected in legislation and regulation. While regulation is only one part of a broader AI-governance system, it can have a direct effect on how AI systems are developed, procured and used, including which requirements are binding, which authorities have oversight and what organisations need to demonstrate before and after deployment. This is why it is useful for delivery teams to have a basic understanding of the various regulatory approaches governments are taking.

Governments can translate shared AI principles into national governance through different combinations of rules, institutions and implementation mechanisms. Figure 5 compares selected approaches to operationalising AI governance. It is not intended to classify countries into fixed models, but to illustrate the different ways governments can translate shared principles into rules and requirements that may ultimately affect AI projects.

FIGURE 5

Approaches and instruments governments use to operationalise AI governance

Approach or instrument	How it works
Cross-sector AI legislation	AI-specific legislation establishes binding requirements across sectors, often differentiated by risk and the roles of developers and users. The EU AI Act illustrates this approach.
Sector-specific and regulator-led approaches	Regulators apply or clarify requirements within their existing mandates. Examples include the UK's regulator-led approach outlined in 2023 and Colombia's Circular Externa 002 of 2024, addressing how data-protection requirements apply to AI.
Standards, codes and non-binding guidance	Standards and guidance help translate principles into management processes and operational practices. The US NIST AI Risk Management Framework, for example, provides voluntary guidance for managing AI risks.
Assurance mechanisms	Assessments, testing and audits generate evidence about whether a system meets specified expectations or requirements. The UK's Introduction to AI Assurance describes techniques that support implementation and demonstrate compliance.
Regulatory sandboxes	Time-limited, supervised environments help participants test approaches and clarify compliance questions. Colombia's Privacy by Design and Default sandbox for AI projects, launched in 2021, illustrates this use of regulatory experimentation.
Government-wide administrative policy and controls	Governments can establish common or mandatory requirements for public institutions through central administrative policies, standards and guidance, without relying on comprehensive AI-specific legislation. These can include accountable owners, AI registers, impact assessments, transparency requirements, training, review and reporting obligations. Canada's Directive on Automated Decision-Making, Australia's Policy for the Responsible Use of AI in Government and the UK's Algorithmic Transparency Recording Standard illustrate this approach.

Source: TBI

For delivery teams, the practical lesson is that countries can take different approaches to AI governance – including regulation and supporting instruments – and their own governance environment may reflect one or a combination of these approaches. A country may have adopted comprehensive AI-specific legislation, rely primarily on existing sectoral regulators and rules, combine legislation with standards and guidance, or still be developing its approach. Understanding where their country sits helps teams anticipate how AI governance is likely to be structured, which institutions may play a role and where more detailed requirements or guidance may come from.

This does not mean that every country fits neatly into one model. Different approaches can coexist and governance frameworks continue to evolve as governments gain experience with AI. Delivery teams don't need to determine which model is "best", but rather understand the approach their government is taking and how that broader environment may shape the governance of their own AI projects.

What should delivery teams know about AI governance in Latin America and the Caribbean?

Latin America and the Caribbean (LAC) are not following a single path to AI governance. Countries are progressing at different speeds and through different combinations of national strategies, legislation, institutions, standards and implementation tools. The Latin American Artificial Intelligence Index (ILIA) 2025[1] captures this variation across 19 countries, classifying Chile, Brazil and Uruguay as pioneers, eight countries including Colombia, Costa Rica and the Dominican Republic as adopters, and more than one-third as explorers, with more nascent ecosystems and more limited institutional capacities. At the same time, the index shows strong and accelerating AI adoption across the region.

This variation means that the governance environment around an AI project will depend on the national context in which the delivery team operates. Some teams may have access to relatively mature national strategies, institutions, requirements and implementation guidance. Others may operate in environments where some of these elements are still being developed. The recent Inter-American Development Bank (IDB) reference model for public-sector AI in LAC similarly identifies the gap between strategic ambition and implementation capacity as a central regional challenge.[2]

The regional context also matters because governments are often building AI governance on top of existing digital government, data, procurement, cyber-security, and sectoral institutions and capabilities. These foundations are not equally developed across countries, and the resources available to establish new AI-specific arrangements also vary. Earlier analysis by the

Organisation for Economic Co-operation and Development (OECD) and Development Bank of Latin America and the Caribbean (CAF) found differences in leadership, institutional capacity and the availability of data, funding, expertise and digital infrastructure across LAC governments.[3] This means that the institutional pathway towards AI governance may look very different from one country to another.

For many governments, effective AI governance will be built by extending and connecting existing state capabilities rather than creating a parallel AI-governance architecture from scratch. Data-protection authorities, digital-government bodies, procurement functions, cyber-security teams, audit institutions and sector regulators already hold many of the capabilities required. The practical question is not merely which new AI-specific institutions are needed, but how existing responsibilities, processes and shared capabilities should be adapted and connected.

At the same time, the region is generating practical approaches that are worth learning from. Governments are adapting international principles for their own institutional contexts and developing tools and arrangements that offer useful lessons beyond their borders. For example:

- **Brazil has developed an Ethical Impact Self-Assessment Framework for AI in the federal public sector.**[4] The framework is designed as a practical, consultative tool to help public institutions identify ethical impacts and mitigate risks associated with AI systems. It includes supporting tools such as an AI-system record, transparency guidance and a RACI matrix, illustrating how broad ethical principles can be translated into practical governance instruments.
- **Chile has linked its national AI policy to a concrete government-wide implementation plan.**[5] Its Action Plan contains 177 initiatives coordinated across 14 ministries, providing a mechanism for translating national priorities into actions and assigning implementation responsibilities across government.
- **Colombia has used a national AI-readiness assessment to identify governance and capacity gaps.**[6] Its UNESCO AI-Readiness Assessment involved more than 3,380 stakeholders and the findings were used to

inform a policy-implementation roadmap. The country's emerging AI legislation has also drawn on the assessment as part of its evidence base.

- **Trinidad and Tobago is conducting a national-readiness assessment to diagnose existing gaps and design an AI-governance framework.**[7] To ensure the government's assessment is aligned with international best practices, it has partnered with UNESCO and the United Nations Development Programme to examine institutional capacity, governance, skills, infrastructure and broader readiness for AI adoption. The findings will inform policy development, investment priorities and long-term national AI strategy.
- **Uruguay has incorporated governance, implementation and monitoring directly into its national AI strategy.**[8] Its 2024–2030 strategy establishes governance as one of its three main objectives and calls for appropriate institutional arrangements, regulatory frameworks, implementation processes, and mechanisms for monitoring and updating the strategy.

These examples illustrate that national AI governance can take different forms and can develop through different pathways. Some governments are strengthening regulation, others are building institutional capacity and implementation tools, and others are embedding AI governance within broader digital-government and data policies.

The practical implication is twofold for delivery teams. First, they need to understand the national governance environment they are operating in and what it already provides. Second, they should not assume that the national framework will answer every governance question. Where national arrangements are incomplete, teams may need to draw on international principles, standards, guidance and experience from other governments to identify what is missing and inform proportionate approaches at the use-case or operational level.

The objective is not to replicate another country's model. It is for each country to understand the starting point of its own governance environment, recognise both the strengths and the gaps, and use available international and regional experience to navigate those gaps responsibly.

03

What Does Your Organisation Need to Have in Place?

Six functions every organisation needs to govern AI

Delivery teams should not have to build the governance environment for every AI project from scratch. Organisations need a minimum set of functions that clarify who is accountable, which processes apply, what evidence is required and where decisions should be escalated.

This section identifies six minimum organisational functions – found across established AI-governance approaches, including guidance from the OECD and the National Institute of Standards and Technology (NIST)[9,10] – that should be available to support responsible AI delivery. Together, they provide a common operating environment for teams governing individual AI systems.

Meeting this minimum does not require a particular organisational structure or level of maturity. These functions can be provided through dedicated teams, existing institutional functions, shared government services or external support. What is essential is that each function can be performed effectively and that delivery teams know how to access it.

1. OWNERSHIP AND DECISION RIGHTS – WHO IS ACCOUNTABLE AND WHO CAN DECIDE?

Every organisation using AI should have clear ownership, responsibilities and decision rights. Delivery teams should know who is accountable for the use of AI, who owns an individual use case, who can approve or stop an initiative and which decisions require additional review within the organisation. Responsibilities across business, technology, legal, data, procurement, cyber-security and risk functions should also be clear.

These roles do not need to sit within a dedicated AI office or follow a single organisational model. What matters is that accountability and authority are explicit, documented and accessible to teams rather than dependent on informal relationships. NIST similarly emphasises documented roles, responsibilities, communication lines and accountability structures for managing AI risks.[11]

Where a system is developed or operated by a supplier, responsibilities can be allocated contractually, but the public organisation remains accountable for deciding whether and how the system is used in carrying out its functions.

Minimum organisational function: A named accountable owner, clear decision rights, and defined approval and escalation routes.

2. VISIBILITY AND CLASSIFICATION – WHAT AI IS BEING USED AND HOW MUCH GOVERNANCE DOES IT REQUIRE?

Organisations need a basic mechanism to know what AI systems are being developed or used, who owns them and their significance in terms of potential impact, risk or organisational importance. An AI inventory or register, combined with a common intake process, can provide visibility across the portfolio and help ensure that governance is applied consistently, rather than only to projects that attract attention.

At a minimum, this mechanism should record the use case, accountable owner, lifecycle status, classification or risk level, relevant supplier or model where applicable, and the status or date of any required review.

Organisations should also have a consistent way to classify use cases according to their potential impact or risk. Not every AI system requires the same level of governance, but teams should not have to develop the criteria for making that distinction each time. NIST explicitly calls for organisations to maintain an inventory of AI systems and determine the appropriate level of risk-management activity.[12]

Minimum organisational function: A mechanism for maintaining visibility over AI use cases, a common intake process, and criteria for classifying use cases and determining the appropriate level of governance.

3. CONTROLS – WHAT COMMON REQUIREMENTS CAN TEAMS REUSE?

Organisations should provide reusable requirements and controls for recurring governance issues. Depending on the institution, these may cover data and privacy, cyber-security, procurement, documentation, human oversight and other requirements that should be considered across projects.

The purpose is not to impose identical controls on every AI system. Rather, organisations should provide a common starting point that teams can apply proportionately and adapt to the context and risks of each use case. Reusable controls reduce duplication, improve consistency and allow governance requirements to be considered earlier in the lifecycle. This is consistent with OECD and NIST approaches, which emphasise organisational policies, processes and controls for managing AI risks.[13,14]

Minimum organisational function: Reusable requirements, templates and controls for the most common governance issues affecting AI projects.

4. ASSURANCE – HOW IS IT DETERMINED THAT THE REQUIREMENTS HAVE BEEN MET?

Organisations should establish how AI projects obtain review, assurance and approval, including what evidence is required, who reviews it and at what stage during the lifecycle review is needed. AI assurance is the process of testing, evaluating and monitoring AI systems, then communicating the evidence so decision-makers and the public can make informed judgements about their trustworthiness. The level of assurance should be proportionate to the risk and characteristics of the use case.

Assurance is different from controls. Controls establish what teams are expected to do; assurance establishes how the organisation determines whether those expectations have been adequately addressed. Higher-risk

or novel systems may require additional assessment or independent review, while lower-risk uses may follow a lighter process. For higher-consequence systems, the organisation should also determine when assurance needs to be independent of the project team or system owner.

The objective is not to require a central committee to approve every AI use case. Instead, organisations should provide predictable review routes and clear escalation mechanisms. OECD research shows that formal mechanisms such as pre-deployment risk assessment, internal review and post-deployment audit remain unevenly established across governments.[15]

Minimum organisational function: Defined assurance requirements, proportionate review and approval routes, and clear escalation mechanisms.

5. CAPABILITY – CAN TEAMS ACCESS THE EXPERTISE THEY NEED?

Governance arrangements only work when organisations have access to the capabilities needed to apply them. Delivery teams may require technical, legal, data, procurement, cyber-security, policy or risk expertise at different points in the lifecycle. They may also need access to appropriate data, infrastructure and implementation support.

These capabilities do not need to reside within the project team or even within the organisation itself. Institutions can draw on central government functions, shared services, external expertise or partnerships. The minimum requirement is access to the capability, rather than the creation of a fully staffed internal function. OECD identifies skills, data, digital infrastructure, investment, procurement and partnerships as key enablers of AI adoption in government, while the IDB places governance and institutional capacity, infrastructure and data, and talent at the centre of its LAC framework.[16,17]

Minimum organisational function: Access to the expertise, data, infrastructure and other resources needed to assess, procure, implement and oversee AI.

6. LIFECYCLE OVERSIGHT – HOW WILL AI BE MONITORED, REVIEWED AND CHANGED OVER TIME?

Organisations should establish how AI systems are monitored, reviewed and managed after deployment. This includes expectations for performance monitoring, benefits realisation, incident reporting, escalation, periodic review and reassessment when a system materially changes.

At the organisational level, the focus is on establishing the mechanisms and thresholds that make this possible: what needs to be monitored, who receives reports, when reassessment is triggered, and who has authority to restrict, change or retire a system. A service owner should be identified early. Where responsibility transfers from a delivery team, the handover should be documented and accepted before live use, with resources, monitoring, incident response and user-review arrangements in place.

AI systems can change over time through model or data updates, changes in suppliers or changes in how the system is used, and NIST explicitly includes ongoing monitoring, periodic review and safe decommissioning within its governance function.[18]

Minimum organisational function: Agreed measures and review arrangements for performance, outcomes and benefits; incident and escalation mechanisms, and a clear process for material changes, reapproval and retirement.

WHAT THIS MEANS FOR DELIVERY TEAMS

These six functions are intended to enable use-case and operational governance. At the organisational level, governance provides the common operating environment: it establishes who is accountable, what processes and controls are available, how assurance is obtained, where expertise can be accessed and how systems are governed over time. These functions do not all need to be performed internally. Organisations can draw on shared government services, external expertise or independent third-party assurance and audit providers where appropriate, while retaining

accountability for how AI is used. The delivery team then works within that environment to govern the specific use of AI and its operation over time, drawing on the capabilities and resources available to it.

For a delivery team, the practical test is straightforward:

- Can we identify who is accountable and who can make the relevant decisions?
- Can we see what AI systems and use cases are in place and how they function?
- Do we know which requirements, controls and processes apply?
- Do we know what evidence, assurance and approval are required?
- Can we access the expertise, data, infrastructure and implementation support required to do this effectively?
- Do we know how incidents, monitoring, changes, reapproval and retirement are handled?

Where the answer to these questions is yes, this means the organisation has the core functions needed to support AI governance. Where important functions are missing, the gap should be made visible and addressed at the appropriate organisational level rather than being left for individual project teams to solve informally.

How are organisations in Latin America and the Caribbean putting AI governance into practice?

Across Latin America and the Caribbean, organisations are at different stages of building the institutional capacity needed to govern AI in practice. Regional assessments highlight differences in leadership, institutional roles, data and digital capabilities, access to expertise and implementation capacity. But they consistently emphasise the importance of sustained leadership and coordination at both central and institutional levels. More

recent regional work similarly identifies the challenge as one of building the state capacity needed to implement, oversee and adapt AI, rather than establishing new rules or strategies (OECD/CAF;[19] IDB;[20] IDB[21]).

This means that while organisations do not need to follow a single institutional model, they do need to ensure that the core functions of AI governance are covered. In some institutions, these functions may be supported by dedicated AI teams; in others, they may sit across existing digital, data, legal, procurement, cyber-security or risk functions, supported by central government or external expertise. The examples that follow illustrate how public organisations in the region are beginning to translate these functions into practical arrangements.

- **Ownership and decision rights – Brazil:** Brazil's federal government has developed a RACI matrix for AI systems that defines who is responsible, who approves, who provides technical guidance and who is informed at each stage of the AI lifecycle, from initial scoping through deactivation. The tool can be adapted by individual public bodies and applies to both internally developed and procured AI systems. It provides a practical example of how organisations can translate broad accountability requirements into explicit roles and decision rights.[22]
- **Visibility and classification – Peru:** Peru's digital-government programme has developed common AI resources for the public sector, including guidance on AI risk assessment, performance evaluation and data quality, alongside ethical standards and an implementation strategy. These resources provide public entities with common methods for identifying and assessing AI use cases, reducing the need for individual project teams to develop their own approaches to recurring governance questions.[23]
- **Controls – Colombia:** Colombia has developed an Ethical Guide for the Implementation, Development and Use of AI Systems in Public Entities, providing practical guidance to public institutions on issues such as responsible development and use, rights, transparency and other considerations that can be incorporated into institutional and project

processes. It illustrates how organisations can translate high-level principles into common guidance that delivery teams can apply across projects.[24]

- **Assurance – Uruguay:** Uruguay's public-sector AI-governance approach includes mechanisms for project selection and prioritisation, criteria for assessing proposed AI initiatives, and arrangements for implementation and monitoring. This illustrates how organisations can establish common decision and review processes rather than treating assurance as an ad hoc exercise for individual projects.[25]
- **Capability – Uruguay:** Uruguay's national AI strategy and earlier public-sector AI work identify skills, data, infrastructure and capacity building as foundations for responsible AI adoption, including the development of technical guidance and capabilities within the public sector. This provides an example of how organisations can build access to relevant expertise and supporting capabilities alongside governance arrangements.[26]
- **Capability – Jamaica:** Jamaica's government asked its existing multi-sector National AI Task Force to develop recommendations for a government-wide AI-literacy programme. The initiative aims to build AI awareness, competence and responsibility across the public service, where AI is already being used, going beyond technical training to cover issues such as risk, privacy, security, accountability and human oversight.[27]
- **Lifecycle oversight – Brazil:** Brazil's RACI approach covers the full lifecycle of AI systems, from prospecting and experimentation through deployment, sustainment and deactivation. By extending accountability beyond deployment, it illustrates how monitoring, change and eventual retirement can be incorporated into organisational governance rather than treating approval as the end point.[28]

These examples illustrate that the functions required for responsible AI can be built incrementally and through different institutional arrangements. Organisations can draw on existing functions, shared services and external expertise depending on their size, mandate and resources.

For delivery teams, the practical test is therefore not whether an organisation has a sophisticated AI-governance function, but whether it provides a clear and usable way to perform each function. Where a function is not yet in place, the gap should be made visible and addressed at the appropriate organisational level, rather than leaving individual project teams to develop informal solutions on their own.

04

How Should Your Team Govern an AI System Through its Lifecycle?

Once teams understand the national governance environment and the organisational functions available, they need to translate these into decisions for a specific AI system. This is where AI governance becomes part of delivery: teams need to decide whether a use case is appropriate, what level of governance it requires, what evidence is needed for approval, what conditions should be met before deployment, and how risks, performance, outcomes and benefits will be monitored and reviewed over time.

The lifecycle provides a practical way to organise these decisions. It is not a rigid sequence. Teams may move backwards and forwards between stages: an assessment may change the proposed use case or design, testing may require further work, or a material change after deployment may trigger reassessment or reapproval. Governance should also be proportionate to the potential consequences of the use case rather than applied uniformly to every system (OECD[29]; NIST[30]).

Some governance requirements run across the lifecycle. The ten stages detailed below set out the main decisions a team needs to make – but accountability and documentation, data governance and cyber-security, transparency and contestability, stakeholder engagement, assurance and evidence management need to be considered throughout. The extent to which they apply should depend on the use case and its potential consequences. Where an AI system materially affects rights, benefits, enforcement or access to public services, teams should establish clear routes for affected people to challenge outcomes, obtain meaningful human review and access appropriate redress.

Benefits realisation also runs across the lifecycle. Teams should define the intended outcomes and benefits at the outset, establish a baseline and proportionate success measures, and identify who is responsible for reviewing the results. After deployment, they should assess whether those

benefits are being realised in practice, for whom, at what cost and with what unintended effects. This evidence should inform decisions to continue, improve, scale, restrict or retire the system.

At each stage, teams should distinguish between decisions they can make themselves and decisions that require approval from an authorised owner, assurance route or oversight body, and record the outcome: proceed, proceed with conditions, escalate, pause or stop.

These stages are decision questions, not ten mandatory approval meetings. In order to avoid making the process too slow or bureaucratic, teams should reuse existing approval processes and evidence where appropriate. Each decision record should identify the authorised decision-maker, evidence reviewed, outcome and conditions. Each condition needs an owner, deadline and verification point. Approval to develop or procure does not constitute approval for live use. The minimum evidence can be a short record for a low-consequence use case. More consequential use cases will require more extensive assessment, testing, assurance and approval.

1. Identify – what problem are we solving, and is AI the right solution?

Governance should begin with the problem to be solved, not the technology available. Teams should define the problem, the intended outcomes and public value, the users and affected groups, and the role AI would play in the process. Before proceeding, teams should also ask a more basic question: is AI actually needed? A simpler digital, data or process intervention may achieve the intended outcome with lower risk, cost and complexity. This “question zero” helps prevent teams from adopting AI simply because the technology is available. The team should also establish the intended uses, limitations and accountable owner of the initiative. Where proportionate, it should document the current baseline, the expected benefits and how success will be assessed.

WHAT SHOULD YOUR TEAM CONSIDER?

- The problem and intended outcomes
- The current process or baseline
- Expected benefits and public value
- Who is expected to benefit and who may be disadvantaged
- How success will be measured
- Intended users and affected groups
- The role AI will play in the process
- Decisions or actions the system will support
- Intended uses, limitations and prohibited uses
- An accountable business or service owner
- A named owner for benefits realisation

Key governance decisions

Is the use case clearly defined, and is AI an appropriate and proportionate way to address the problem?

Minimum evidence

A documented use case, intended purpose, baseline or current process, expected benefits and success measures, key affected groups, considered alternatives, review points, an accountable owner and a named benefits owner.

Owner/approval route

Accountable business or service owner; delivery team prepares the use-case record and options.

CASE STUDY

Colombia: A Problem-First Approach to AI Adoption

Colombia's Ministry of Information and Communication Technologies has developed a guideline to support public officials in identifying and addressing opportunities to improve the processes and services they lead.[31] The guideline establishes a clear sequence for solution identification: first, define the problem and improvement objective; second, identify the specific process or function to improve; third, assess whether traditional technological solutions can address the need. Only after these steps are completed are officials encouraged to consider the adoption of emerging technologies as a solution, assess its technical and operational feasibility and finally proceed to solution design and piloting. For delivery teams, the practical implication is to start with the problem to be solved and assess the most appropriate solution, rather than starting with the technology or AI system and looking for a problem to which it can be applied.

2. Classify – how much governance is needed?

Not every AI system requires the same level of scrutiny. In practice, proportionality affects the depth of assessment, evidence required, assurance and approval route, testing, monitoring and review frequency: an internal productivity tool should not face the same process as a system informing eligibility, enforcement or allocation decisions.

Teams should assess the potential consequences of failure, misuse or error, rather than relying primarily on the type or sophistication of the technology. Where a national or organisational risk taxonomy exists, teams should use it.

Classification should be applied to the specific use case rather than the technology in isolation and should be treated as provisional. It should be revisited if assessment, design, testing or subsequent changes materially alter the system's potential impact or operating context. Where one does not, a practical classification can help determine the intensity of governance required. Classification should also consider the system's degree of autonomy and authority to act, including the tools, systems, permissions or actions it can access or execute. As AI systems become more capable of taking actions within workflows, the consequences of an incorrect action may differ substantially from those of an incorrect recommendation. The organisation's risk tolerance should also be clear, particularly for higher-consequence use cases.

WHAT SHOULD YOUR TEAM CONSIDER?

- Whether the system affects rights, benefits or access to public services
- Whether outputs influence consequential decisions
- Whether meaningful human review is possible
- Who is affected and whether vulnerable groups are involved
- The scale and reversibility of potential impacts
- The system's degree of autonomy and the actions, tools or permissions it is authorised to use
- The organisation's risk tolerance

Key governance decisions

What level of assessment, assurance and oversight is proportionate to the potential consequences of this use case?

Minimum evidence

A documented classification, the rationale for it and the governance requirements triggered by that classification.

Owner/approval route
Delivery team proposes classification; organisational risk, AI or assurance function confirms material or uncertain cases.

CASE STUDY

Brazil: What Proportionality Can Look Like in Practice

Brazil's federal Ethical Impact Self-Assessment is divided into two questionnaires designed for implementing teams to complete.[32] The first questionnaire screens the risk of an AI system or use case and classifies it into four categories, from low to excessive risk. The resulting classification determines the required depth of the subsequent Main Assessment Questionnaire. Low-risk AI uses are subject to a simplified assessment, while high-risk systems must complete the full version. The final results then guide implementing teams on their next steps, ranging from reassessing the fundamentals or improving ethical adherence to maintaining continuous monitoring. For delivery teams, the practical implication is not that every government should follow Brazil's framework, but that classification should materially change the depth of assessment required and, more broadly, the governance burden that follows.

3. Assess – what risks and requirements need to be addressed?

Once the use case has been classified, the team should identify the risks and requirements that need to be addressed before the project proceeds. This may include AI-specific assessment as well as existing privacy, data protection, cyber-security, legal, procurement or sectoral processes. The classification should determine which assessments are relevant and how much evidence is required. Teams should use the organisation's common controls and assurance processes rather than creating new approaches from scratch. Assessment should also consider whether the system is capable of delivering the benefits originally identified, not only whether it can perform technically. This includes examining whether the necessary adoption, process change, data, staffing and operational conditions are in place, and whether the expected benefits remain credible once implementation and operating costs are considered.

WHAT SHOULD YOUR TEAM CONSIDER?

- Data quality, suitability and representativeness
- Privacy and data protection
- Legal basis
- Security and misuse
- Fairness and discriminatory impacts
- Transparency and explainability
- Reliability and performance
- Human oversight, review and redress
- Supplier and third-party dependencies

- Expected outcomes and public value
- Conditions and dependencies required to realise the benefits
- Implementation and operating costs
- Whether benefits are likely to be distributed equitably

Key governance decisions
Can the identified risks and requirements be adequately addressed within the proposed design and operating model?

Minimum evidence
Relevant assessments, documented risks and mitigations, expected outcomes, and a record of unresolved issues requiring escalation.

Owner/approval route
Delivery team and relevant specialists prepare the assessment; the accountable owner confirms that the evidence and unresolved issues are ready for approval.

CASE STUDY

Dominican Republic: Integrating AI Risk Assessments With Existing Governance Frameworks

The Dominican Republic's National AI Strategy seeks to build on the country's existing digital governance architecture, including cyber-security, privacy, data protection and regulatory frameworks.[33] As part of its governance pillar, the strategy sets objectives to update existing legal frameworks on intellectual property and personal data protection, for example, to account for AI systems. Moreover, it calls for collaboration with the country's National Cybersecurity Center to develop cyber-security and privacy standards for AI, while the broader Digital Agenda already establishes mechanisms for cyber-risk management, incident response, data governance and regulatory oversight. For delivery teams, the practical implication is to build AI governance into existing processes rather than creating parallel ones from scratch. Teams can apply established safeguards and assessment processes for areas such as cyber-security, privacy, data protection or legal compliance to AI systems, adapting them where necessary to address AI-specific risks.

4. Approve – should the project proceed, and under what conditions?

Assessment should lead to a decision. The appropriate decision-maker should determine whether the project can proceed and any conditions that should apply. The decision should consider whether the risks can be managed and whether the initiative can deliver sufficient public value to

justify its costs, risks and operational demands. Where evidence remains uncertain, approval may be conditional on a pilot or time-limited deployment with defined evaluation criteria and a subsequent decision point. The approval route should reflect risk. Lower-risk systems may follow a lightweight process, while higher-risk or novel systems may require more senior, cross-functional or independent review. The objective is not for a central committee to approve every AI use case, but for teams to know who can make the decision, what evidence is required and when escalation is necessary.

WHAT SHOULD YOUR TEAM CONSIDER?

- Identify the authorised decision-maker
- Confirm that required assessments and controls are complete
- Determine whether conditions are needed before proceeding
- Record unresolved risks and escalation decisions
- Define conditions that must be checked at later stages

Key governance decisions
Is there sufficient evidence to authorise the initiative to move into design, procurement or development, and are its expected benefits, costs, risks and conditions for success sufficiently clear?

Minimum evidence
Assessment findings, classification, baseline, expected benefits, success measures, resource and operating requirements, unresolved risks, proposed conditions and a proportionate benefits-realisation and evaluation plan. Record the decision, authorised scope and decision-maker separately.

Owner/approval route
Authorised organisational decision-maker or governance route; delivery team prepares the decision pack.

CASE STUDY

Brazil: Escalation and Approval Process for Public-Facing AI Systems

Brazil's National Telecommunications Agency (ANATEL) has established its own governance structure for the institution's AI systems.[34] This structure includes a Permanent Thematic Forum under ANATEL's Executive Management Committee (CGE), responsible for validating AI risk classifications and supporting monitoring and decision-making, while ANATEL's AI Research Group (IA.lab) provides the technical expertise and assessments that inform these decisions. For public-facing AI systems, the policy establishes a concrete escalation process: the responsible team first requests a technical assessment from IA.lab, which is then submitted to the Permanent Thematic Forum. Where relevant risks are identified, the matter is escalated to the CGE for final deliberation and approval. For delivery teams, the practical implication is to understand the applicable escalation and approval process, including when additional technical review or senior approval is required, and ensure that the relevant evidence reaches the appropriate decision-making authority.

5. Design, procure or build – how will governance requirements be embedded in the solution?

Governance requirements need to move from assessment documents into the design, procurement and implementation of the system. This is particularly important in government because the choice between buying and building affects not only cost and speed, but also data access, supplier

dependency, transparency, security, interoperability and control over future changes. Teams should therefore make an explicit buy-versus-build decision and embed relevant governance requirements into technical specifications, contracts or internal development processes. Governance should be part of delivery, not a separate compliance exercise. Design and procurement decisions should also preserve the data, access rights and reporting arrangements needed to measure outcomes, costs and benefits after deployment.

WHAT SHOULD YOUR TEAM CONSIDER?

- Buy-versus-build rationale
- Data ownership, access and permitted use
- Privacy and security requirements
- Model and version control
- Sub-processors and downstream dependencies where relevant
- Logging, fallback and business-continuity arrangements
- Supplier due diligence and responsibilities
- Documentation, auditability and access to evidence
- Access to the data needed to assess outcomes and benefits
- Supplier reporting obligations relevant to benefits monitoring and evaluation
- Incident notification and change-management obligations
- Interoperability, portability and exit arrangements
- Responsibilities when models or services are updated

Key governance decisions

Have the governance requirements been translated into the technical, procurement and contractual arrangements for the system?

Minimum evidence

Documented design or procurement requirements and clear allocation of responsibilities between government and suppliers.

Owner/approval route

Delivery, procurement and technical leads embed requirements; accountable owner retains decision authority over major trade-offs.

CASE STUDY

Chile: Include AI Governance and Ethical Considerations in Supplier Contracts

Chile's public procurement agency, ChileCompra, has issued a directive for the procurement of data science and AI solutions.[35] The directive provides recommendations for public officials structuring AI contracts and distinguishes between acquiring existing solutions and commissioning tailor-made systems. In particular, it recommends incorporating specific contractual clauses covering algorithmic audits, data-protection impact assessments and information-security measures to guarantee that procured solutions meet standards of transparency, fairness and accountability. For tailor-made systems, it also recommends ensuring that public officials have access to the source code and technical documentation needed to understand, reproduce, operate and maintain the solution. For delivery teams, the practical implication is to consider AI-governance requirements upfront during procurement and incorporate them, where appropriate, into contractual obligations or considerations that preserve institutional oversight and accountability throughout the system lifecycle.

6. Test – is the system fit for its intended use?

Testing should establish whether the system performs sufficiently well in the context in which it will actually be used. Technical benchmarks alone may not be enough. Teams should test realistic tasks, representative data and relevant edge cases, and assess whether human oversight works in

practice. Testing should be linked to predefined acceptance criteria. Where those criteria are not met, the team should remediate, redesign or reconsider the use case.

Acceptance criteria should be defined before final testing and should reflect the actual operating context and consequences of failure. Where relevant, teams should compare performance against the existing process or non-AI alternative rather than assessing the AI system in isolation. Testing should also confirm whether the data and measurement arrangements needed to assess outcomes and benefits after deployment are workable.

WHAT SHOULD YOUR TEAM CONSIDER?

- Accuracy and reliability
- Performance on representative tasks and data
- Failure modes and edge cases
- Security and robustness
- Fairness or differential performance where relevant
- Usability and human-AI interaction
- Effectiveness of human oversight
- Ability to meet transparency and documentation requirements
- Misuse, abuse and adversarial testing where relevant

Key governance decisions

Does the evidence show that the system is sufficiently reliable, safe and fit for its intended use?

Minimum evidence
A testing plan, test results, identified failures, remediation actions and a documented readiness decision.

Owner/approval route
Delivery and technical leads provide test evidence; assurance reviewers assess it; the authorised decision-maker determines readiness and accepts any permissible residual risk.

CASE STUDY

Jamaica: Test the Service, Not Only the Model

Jamaica's UNESCO AI-readiness assessment highlights the coexistence of strong digital uptake with urban-rural connectivity gaps and power interruptions that can undermine service continuity.[36] The government is actively working to enhance digital connectivity in education, in particular, aligned with broader sectoral plans to provide equitable access to technology that allows preparation for a digitally driven future. These findings underline why readiness for AI-enabled public services cannot be judged through model performance alone. For delivery teams, the practical implication is to test the intended operating conditions: who can access the service, what happens when connectivity or power fails, and whether a workable alternative remains available. Deployment evidence should demonstrate continuity and inclusion as well as technical performance.

7. Deploy – are the conditions for live use in place?

Deployment should be treated as a governance decision, not just the final technical step. Before going live, teams should confirm that the conditions attached to approval have been met and that the arrangements required to operate the system responsibly are in place. For higher-risk systems, deployment may require explicit sign-off following assurance and testing. The team should also be clear about how users will be informed, how human oversight will operate and how incidents will be handled. The baseline, success measures, data-collection arrangements, review dates and responsibility for benefits realisation should also be confirmed before live use. A system should not proceed with broad benefit claims but no practical way to test them.

WHAT SHOULD YOUR TEAM CONSIDER?

- Named system ownership
- Documented intended uses and limitations
- Human oversight arrangements
- User and operator training and operating procedures
- Appropriate user information and transparency
- Incident and escalation procedures
- Monitoring indicators and review arrangements
- Baseline and benefit indicators
- Benefits owner and review dates
- Data-collection and evaluation arrangements
- Operating and oversight costs

- Feedback or challenge mechanisms where relevant
- Required records and documentation
- Fallback, override or rollback arrangements where relevant
- Complaint, challenge and human-review routes where relevant

Key governance decisions
Is the system ready to operate in the real-world context for which it was approved?

Minimum evidence
Deployment approval, confirmation that required controls are in place, documented ownership and monitoring arrangements, and an agreed plan for reviewing performance, risks, outcomes, costs and benefits.

Owner/approval route
Named service or system owner signs off live use, with senior or cross-functional sign-off where required.

CASE STUDY

Brazil: Make AI Use Transparent and Auditable Through a Public Inventory

Brazil's Federal Court of Accounts (TCU) has published a website containing an inventory of its operational AI systems.[37] The website documents each system's intended purpose, responsible units, results, identified risks and safeguards in place. TCU has also issued guidelines for the responsible use of AI within the institution, which establish AI as a support tool rather than a replacement for human analysis, validation and accountability, and emphasise the need for trained personnel and human oversight. For delivery teams, the practical implication is to clearly document intended uses, known risks, limitations, ownership and human-oversight arrangements throughout the AI system's lifecycle. Making this information publicly available strengthens accountability and enables both government counterparts and citizens to scrutinise how AI is used and identify potential misuses, while creating an auditable record of the conditions and safeguards governing deployment.

8. Monitor – is the system delivering the expected results and operating as intended?

Governance does not end at deployment. The named service or system owner should monitor whether the system continues to operate safely, reliably and within its approved purpose. The named benefits owner should review whether it is producing the intended outcomes and public value.

Monitoring benefits requires more than reporting model performance or activity levels. Teams should compare results with the agreed baseline and, where feasible, with the previous process or another credible alternative. They should consider the full costs of implementation and operation, whether expected efficiencies have translated into improved services or productive capacity, who is benefiting, and whether any groups are experiencing poorer outcomes or unintended effects.

Benefits and impacts should be reviewed at defined points, not only when an incident occurs. The evidence should inform decisions to improve, scale, restrict or discontinue the system. Where benefits are not being realised, teams should examine the causes. Monitoring is therefore both a risk-management activity and a way to test whether the original case for using AI remains valid.

WHAT SHOULD YOUR TEAM CONSIDER?

- System performance and reliability
- Changes in data or model behaviour
- Incidents and errors
- User feedback and complaints
- Unexpected outcomes or new risks
- Changes in affected populations or operating conditions
- Compliance with human oversight
- Achievement of intended outcomes and public value
- Results compared with the baseline and feasible alternatives
- Who is benefiting and whether outcomes differ across groups
- Adoption and changes to staff or user behaviour

• Full operating, supplier, assurance and oversight costs

• Whether expected efficiencies have translated into better outcomes or usable capacity

• Whether the original assumptions and benefits case remain valid

Key governance decisions
Is the system continuing to operate within its approved purpose and risk tolerance, and is there sufficient evidence that it is delivering public value to justify continued use, wider adoption or further investment?

Minimum evidence
Defined performance-outcome and benefit indicators; an agreed baseline; periodic review and evaluation records; actual costs and benefits where measurable; incident and user-feedback records; and documented corrective, scaling or escalation decisions.

Owner/approval route
The service owner oversees operational performance and risk. The named benefits owner reviews outcomes, costs and benefits, drawing on analytical or evaluation expertise where needed. The accountable owner acts or escalates where results fall below agreed thresholds or the original case for using AI is no longer valid. Authority to restrict or suspend use must be defined before deployment.

CASE STUDY

Brazil: Monitoring, Auditing and Incident Reporting Procedures

Brazil's Federal District AI Governance Policy establishes a clear monitoring, reporting and incident escalation framework for AI systems.[38] The policy assigns the Secretariat of Government and Innovation (SGDI) responsibility for the permanent monitoring of AI governance, while the agencies and entities operating AI systems must submit a semi-annual report to SGDI. These reports must include the list of systems in operation, their risk classification, incidents recorded and measures taken, performance indicators for high-risk systems, the status of public officials' training and any updates to the agency's AI inventory. In case of a serious incident involving risk to the life, physical integrity or health of people, violations of fundamental rights, interruption of essential public services, data leakage, or systemic discriminatory results, agencies must notify SGDI within a pre-established timeframe. SGDI has the authority to assess the severity of the incident and recommend corrective measures to the agency or escalate further to the General Coordination of Information and Communication Technologies for the possible suspension of the system's operation. For delivery teams, the practical implication is having monitoring systems in place with clear incident-escalation procedures so that the service owner has a clear escalation route in case of an emergency.

TBI has been working with the Federal District and Instituto Hardware BR to strengthen its approach to AI governance, including through a study visit to London in March 2026 to capture international best practices and

benchmarking, and the development of a green paper with recommendations for the responsible adoption and governance of AI across the Federal District.

9. Change and reapprove – has anything changed enough to alter the governance requirements?

AI systems do not remain static. Models may be updated, suppliers may change their services, new data may be introduced, functionality may expand or the system may begin to be used for a different purpose. Not every change requires a new approval, so teams should establish what constitutes a material change and what governance response it triggers. This creates an important feedback loop: governance should be able to move backwards through the lifecycle when circumstances change.

WHAT SHOULD YOUR TEAM CONSIDER?

- Changes to purpose or intended use
- Changes to populations affected
- Changes to data or sensitivity of data
- Changes to the role AI plays in decision-making
- Changes to the underlying model or significant architecture
- Changes to suppliers or key service providers

• Changes to the operating context or risk profile

• Changes to system prompts, retrieval sources or knowledge bases where these materially affect behaviour

• Changes to tools, permissions or actions the system is authorised to take

• Changes that materially affect expected outcomes, benefits, costs or the assumptions on which the original case was based

• Material supplier or model-version updates

• Expansion to new user groups, decisions, services or workflows

Key governance decisions
Has the system changed enough to require reassessment, retesting or reapproval?

Minimum evidence
A documented change-management process, materiality thresholds and records of reassessment or reapproval where required.

Owner/approval route
Service owner and change authority determine materiality; material changes return to the relevant approval route.

CASE STUDY

Mexico: Defining Which Changes Trigger Reassessment and Require Revalidation

Mexico's National Health System guideline for AI technology establishes clear procedures for managing changes in AI systems that can occur after deployment.[39] The guideline distinguishes between regular software updates and adaptive models that update automatically through continuous learning. The latter require system teams to document and revalidate these updates locally. Additionally, changes in the context that the system is used in, like the population served, input data or clinical practices, can also trigger an extraordinary revaluation of the model. System-performance deterioration, detected through continuous monitoring and evaluation, can also trigger recalibration, restricted use or suspension of a system, which can be reactivated through a predetermined technical validation process. For delivery teams, the practical implication is to define in advance what constitutes a material change and establish proportionate responses, ensuring that significant changes to a system or its operating context trigger reassessment and, where necessary, renewed approval rather than allowing the original approval to apply indefinitely.

10. Retire – should the system continue, change or stop?

Retirement should not be an afterthought. Teams should consider from the beginning what would justify continuing, replacing or discontinuing the system. Continuation and retirement decisions should draw on evidence of

realised outcomes and benefits, as well as risk and technical performance. A system should not continue indefinitely just because it remains operational or because significant resources have already been invested. A system may no longer be justified because it does not deliver the expected value, its risks cannot be adequately managed, its technology has become obsolete or a better alternative is available. Retirement should preserve appropriate records and lessons while ensuring that decommissioning does not unintentionally disrupt the public service.

WHAT SHOULD YOUR TEAM CONSIDER?

- Who can authorise decommissioning
- What happens to data, records and documentation
- How supplier access and contracts are closed
- How technical and operational dependencies are removed
- How continuity of the public service is maintained
- How affected users are informed where relevant
- What lessons should be captured for future projects

Key governance decisions
Does the evidence on outcomes, benefits, costs and risks justify continuing, changing, scaling, replacing or retiring the system?

Minimum evidence
A documented continuation or retirement decision, evidence of realised outcomes, benefits and costs, a decommissioning plan where applicable, and a record of relevant information and lessons learned.

Owner/approval route
Accountable owner recommends retirement; authorised governance route approves where service impact or risk is material.

CASE STUDY

Brazil: Planning, Executing and Learning From the AI System's Decommissioning

Brazil's Unified AI Guide supports the governance of AI systems throughout their whole lifecycle, including decommissioning.[40] For this final stage, the guide recommends documenting the reasons for retirement and identifying any related or dependent systems so key stakeholders can be notified in time, and the transition can be planned. When retiring the model, responsible teams must assess whether a replacement solution is in place, where applicable, ensure sensitive data is removed or appropriately archived, and disconnect endpoints and release technical resources. Finally, the guidelines stress that teams should preserve the project's technical knowledge and ensure accountability after shutdown by archiving key information and documenting lessons learned. Across these stages, a RACI matrix provides clear guidelines on who is responsible for executing, approving, consulting on and being informed of each activity. For delivery teams, the practical implication is to plan for system retirement with clear ownership, approval processes, procedures and clear timelines to safely remove technical dependencies while preserving lessons learned and accountability.

05

Conclusion: Putting the Lifecycle Into Practice

The lifecycle should be understood as a repeatable decision process, not a checklist that is completed once. Governance should begin before a system is procured or built, continue through deployment and operation, and return to earlier stages when new evidence or material changes require it. The objective is not to apply the greatest possible amount of governance, but to ensure that the right decisions are made, with proportionate evidence, at the right point in the lifecycle.

Evidence about risks, performance, costs, outcomes and realised benefits should also feed back into organisational governance. It can help organisations improve templates and controls, compare use cases, direct resources towards approaches that demonstrate public value, and stop or redesign those that do not. This closes the loop between governing individual AI systems and improving the organisation's wider approach to AI. This reinforces the principle that governance is part of delivery rather than a separate compliance exercise. Following this approach can help LAC countries govern AI systems responsibly, as they continue to develop and evolve, while still delivering for their citizens.

Endnotes

1 https://www.cepal.org/en/publications/84471-latin-american-artificial-intelligence-index-ilia-2025-executive-summary

2 https://publications.iadb.org/es/modelo-de-referencia-para-politicas-de-ia-en-el-sector-publico-de-america-latina-y-el-caribe

3 https://www.oecd.org/en/publications/2022/03/the-strategic-and-responsible-use-of-artificial-intelligence-in-the-public-sector-of-latin-america-and-the-caribbean%5F17c90e5e/full-report/component-8.html

4 https://www.gov.br/governodigital/pt-br/infraestrutura-nacional-de-dados/inteligencia-artificial-1/publicacoes/framework-aie

5 https://minciencia.gob.cl/areas/inteligencia-artificial/Inteligencia-Artificial/politica-Nacional-Inteligencia-Artificial-Plan-Accion/

6 https://www.unesco.org/en/articles/ai-readiness-assessment-methodology-colombia

7 https://mpaai.gov.tt/news/launch-tts-national-artificial-intelligence-ai-assessment

8 https://www.gub.uy/agencia-gobierno-electronico-sociedad-informacion-conocimiento/comunicacion/publicaciones/estrategia-nacional-inteligencia-artificial-del-uruguay-2024-2030

9 https://doi.org/10.1787/795de142-en

10 https://doi.org/10.6028/NIST.AI.100-1

11 https://doi.org/10.6028/NIST.AI.100-1

12 https://doi.org/10.6028/NIST.AI.100-1

13 https://doi.org/10.1787/795de142-en

14 https://doi.org/10.6028/NIST.AI.100-1

15 https://doi.org/10.1787/795de142-en

16 https://doi.org/10.1787/795de142-en

17 https://doi.org/10.18235/0014418

18 https://doi.org/10.6028/NIST.AI.100-1

19 https://doi.org/10.1787/1f334543-en

20 https://doi.org/10.18235/0013801

21 https://doi.org/10.18235/0014418

22 https://www.gov.br/governodigital/pt-br/infraestrutura-nacional-de-dados/inteligencia-artificial-1/publicacoes/matriz-raci?

23 https://www.gob.pe/institucion/pcm/informes-publicaciones/6701585-informes-finales-sobre-la-estrategia-estandares-y-guias-de-la-ia-para-gobierno-digital-en-el-peru

24 https://www.dnp.gov.co/Prensa%5F/Noticias/Paginas/gobierno-impulsa-guia-etica-de-inteligencia-artificial.aspx

25 https://www.gub.uy/agencia-gobierno-electronico-sociedad-informacion-conocimiento/comunicacion/publicaciones/estrategia-inteligencia-artificial-para-gobierno-digital/estrategia-4

26 https://www.gub.uy/agencia-gobierno-electronico-sociedad-informacion-conocimiento/comunicacion/publicaciones/estrategia-nacional-inteligencia-artificial-del-uruguay-2024-2030

27 https://opm.gov.jm/minister-wheatley-issues-urgent-directive-to-national-ai-task-force-to-advance-ai-literacy-across-the-government-of-jamaica/

28 https://www.gov.br/governodigital/pt-br/infraestrutura-nacional-de-dados/inteligencia-artificial-1/publicacoes/matriz-raci

29 https://www.oecd.org/en/publications/governing-with-artificial-intelligence%5F795de142-en.html

30 https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf

31 https://gobiernodigital.mintic.gov.co/692/articles-179148%5FGuia%5FTecnologias%5FEmergentes.pdf

32 https://www.gov.br/governodigital/pt-br/infraestrutura-nacional-de-dados/inteligencia-artificial-1/publicacoes/framework-aie

33 agendadigital.gob.do/wp-content/uploads/2023/10/Final_ENIA-Estrategia-Nacional-de-Inteligencia-Artificial-de-la-Republica-Dominicana-.pdf

34 https://informacoes.anatel.gov.br/legislacao/resolucoes-internas/2152-resolucao-interna-anatel-554

35 https://www.chilecompra.cl/2026/03/chilecompra-actualiza-directiva-para-la-contratacion-de-proyectos-de-ciencia-de-datos-e-inteligencia-artificial/

36 https://bcj.gov.jm/sites/default/files/2026-02/Jamaica%20AI%20Readiness%20Assessment%20Report%20-%20UNESCO.pdf

37 https://portal.tcu.gov.br/imprensa/noticias/novo-site-amplia-transparencia-e-apoia-uso-responsavel-de-ia-no-setor-publico

38 https://www.sinj.df.gov.br/sinj/Norma/2c3aea13e3b14bf6bc81721d8278a586/Decreto%5F48901%5F03%5F07%5F2026.html

39 https://dgmoss.salud.gob.mx/dgmoss/uploads/Lineamiento%5FImplemt%5FIA.pdf

40 https://www.gov.br/governodigital/pt-br/infraestrutura-nacional-de-dados/inteligencia-artificial-1/publicacoes/guia-unificado-de-inteligencia-artificial-para-o-setor-publico-1/GuIA

TONY BLAIR
INSTITUTE FOR
GLOBAL CHANGE

Follow us

facebook.com/instituteglobal

x.com/instituteGC

instagram.com/institutegc

General enquiries

info@institute.global

Copyright © October 2026 by the Tony Blair Institute for Global Change

All rights reserved. Citation, reproduction and or translation of this publication, in whole or in part, for educational or other non-commertial purposes is authorised provided the source is fully acknowledged Tony Blair Institute, trading as Tony Blair Institute for Global Change, is a company limited by guarantee registered in England and Wales (registered company number: 10505963) whose registered office is One Bartholomew Close, London, EC1A 7BL.