Digital Convergence USP2030 | AI HUB

AI in Social Protection:

DATA SOVEREIGNTY CONSIDERATIONS

AI HUB FOR SOCIAL PROTECTION

The AI Hub for Social Protection (AI Hub) is an initiative under the Digital Convergence Initiative (DCI) that helps social protection institutions to harness the benefits of responsible AI for the public good. It supports the development of strategies, frameworks, digital systems and institutional capacities to design, govern and implement AI solutions that are ethical, effective and uphold data sovereignty.

Visit our website at **spdci.org/ai-hub**.

ACKNOWLEDGMENTS

This report was authored by Robert Worthington and Gilbert Kondani (Senior Systems Architect and Business Analyst, Kwantu). We would like to thank the many colleagues who reviewed the various drafts: Juul Pinxten and Phillippe Leite (World Bank), Rob Worthington (Kwantu), Gabriele Erba (UNICEF), Valentina Barca (independent consultant), Franziska Clausius and Nour Barakat (GIZ). We would also like to thank Susan Sellars (copyeditor) for helping to bring the report into its final form.

REQUIRED CITATION

DCI AI Hub for Social Protection, *AI in Social Protection: Data Sovereignty Considerations*, Deutsche Gesellschaft für Internationale Zusammenarbeit, Germany, 2026.

AI HUB KNOWLEDGE ARCHITECTURE

The AI Hub is committed to advancing the global dialogue on the use of AI in social protection by creating and sharing knowledge on the responsible adoption of AI solutions. This report on Data Sovereignty Considerations is part of a broader portfolio of AI Hub knowledge products that support policymakers, practitioners, and partners in navigating the opportunities and challenges of AI in social protection.

Artificial intelligence is a fast-moving field, with new developments, evidence, and applications emerging continuously. While this report reflects the best available information at the time of publication, it may not include subsequent advances. Readers are encouraged to consult our AI Hub website for the most recent developments, publications, and resources: **spdci.org/ai-hub**.

CONTENTS

ACRONYMS

AI	artificial intelligence
API	application programming interface
AWS	Amazon Web Services
BCR	Binding Corporate Rule
CSL	Cybersecurity Law (China)
DCI	Digital Convergence Initiative
DCRM	Disaster and Climate Risk Management
DPA	data processing agreement
DPG	digital public good
DSL	Data Security Law (China)
EU	European Union
GDP	Gross Domestic Product
GDPR	General Data Protection Regulation
GIZ	Deutsche Gesellschaft für Internationale Zusammenarbeit
GPU	graphics processing unit
HIC	high-income country
IaaS	Infrastructure as a Service
LIC	low-income country
LMIC	low- and middle-income country
MCC-AI	Model Contractual Clauses for AI Procurement
ML	machine learning
NITDA	National Information Technology Development Agency
OECD	Organisation for Economic Co-operation and Development
PaaS	Platform as a Service
PET	privacy-enhancing technology
PII	personally identifiable information
PIPL	Personal Information Protection Law (China)
POPIA	Protection of Personal Information Act
SCC	Standard Contractual Clause
SLA	service level agreement
SL-UDI	Sri Lanka Unique Digital Identity
UMIC	upper middle-income country
UNESCO	United Nations Educational, Scientific and Cultural Organization

GLOSSARY

This glossary provides definitions of key terms used throughout the report.

Admin-of-one-jurisdiction
A policy that restricts system administration of AI infrastructure to personnel operating within national legal boundaries, ensuring that no foreign-based administrator can access or control sovereign data systems.

Artificial intelligence (AI)
A branch of computer science focused on developing systems capable of performing tasks that usually require human intelligence, such as reasoning, prediction, pattern recognition, and natural-language understanding. In social protection, AI includes analysing beneficiary data, predicting vulnerability, automating eligibility determination, and providing conversational interfaces for service delivery. (This definition draws on the *Global Evidence Review*, DCI AI Hub, 2026a.)

Beneficiary data
Personal data held by government social protection agencies on individuals and households who receive or have applied for social protection benefits, including identity records, socio-economic information, vulnerability profiles, and payment histories. As used in this report, beneficiary data is treated as public trust data.

CLOUD Act (Clarifying Lawful Overseas Use of Data Act)
A United States federal law (2018) that enables US law enforcement authorities to compel US-based cloud service providers to hand over data stored outside the United States, regardless of where the data physically resides. This creates jurisdictional exposure risks for governments that host social protection data on infrastructure operated by US companies.

Compliance-as-code
An approach to IT governance in which data residency and sovereignty rules are embedded directly into software deployment pipelines so that compliance is automatically enforced at the point of deployment, preventing any workload containing sovereign data from running on unapproved infrastructure.

Cross-border data transfer
The movement of personal data from one national jurisdiction to another, whether through physical transmission, cloud storage, remote access, or AI model training on infrastructure located abroad. Cross-border transfers are a primary vector of sovereignty risk and are subject to varying regulatory controls across countries.

Data embassy
An arrangement by which a country's data is hosted on physical infrastructure located outside its borders, but the hosting environment is operated under the legal and administrative jurisdiction of the home country. The model (pioneered by Estonia) could enable countries without adequate domestic compute capacity to maintain legal sovereignty over their data even when physical infrastructure is shared or offshore.

Data localisation
Legal or regulatory requirements that mandate that certain categories of data (such as government data, health records, or social protection information) be stored and processed within the national territory. Localisation is a key instrument for asserting data sovereignty, but varies significantly in scope and strictness across countries.

Data processing agreement (DPA)
A contractual instrument between a government or data controller and a third-party vendor or processor that sets out legally-binding obligations as to how personal data may be processed, including requirements on storage location, security standards, sub-processor restrictions, audit rights, and data deletion.

Data protection
The legal and ethical practice of handling personal data in accordance with defined principles that safeguard individuals' rights. This principle requires that data be processed lawfully, fairly, and transpar-

ently, and only when there is a valid legal basis. Personal data must be collected for specific and legitimate purposes, kept to the minimum necessary, and maintained accurately. It should not be retained longer than necessary and must be secured to ensure its integrity and confidentiality against unauthorised use, loss, or damage. (This definition is derived from the European Union's GDPR; see also the DCI-endorsed *Implementation Guide – Good Practices for Ensuring Data Protection and Privacy in Social Protection Systems*, Wagner et al., 2024.)

Data sovereignty

The principle that data generated within a nation's jurisdiction is governed by that nation's laws, policies, and values, ensuring that the country retains local control over data storage, access, and use. As used in this report, data sovereignty is understood as a dimension of data protection that focuses on the right of a country to: (i) determine where and how its sensitive data is stored; (ii) decide who can process that data and under what conditions; and (iii) ensure that cross-border data flows occur only in ways that uphold the country's data protection standards and permit its authorities to enforce those standards. (This definition draws on several sources including Soulé, 2024 and UNESCO's *Recommendation on the Ethics of Artificial Intelligence*, 2021.)

Data transfer impact assessment (DTIA)

A structured assessment conducted before transferring data across borders to evaluate whether the receiving jurisdiction offers equivalent levels of data protection, and to identify and mitigate risks arising from the transfer. DTIAs are a recommended governance mechanism in the report's policy framework.

Derived dataset

A dataset produced by processing, transforming, or aggregating original data – for example, anonymised datasets, synthetic datasets, or feature-extracted datasets used in AI model training. Derived datasets may still carry sovereignty implications when the underlying data originates from sovereign social protection registries.

Differential privacy

A mathematical technique for adding calibrated statistical noise to datasets or AI model outputs in order to protect individuals' privacy while still enabling meaningful analysis. Differential privacy allows governments to participate in collaborative AI training without exposing sensitive individual-level data.

Digital public good (DPG)

Open-source software, data, AI models, or standards that meet criteria for privacy, security, and relevance to the Sustainable Development Goals and that are freely available for use and adaptation. DPGs are recommended in this report as a means of reducing vendor lock-in, strengthening transparency, and enabling countries to build sovereign AI capacity.

Encryption key ownership

The principle that the government or data controller (rather than the cloud or AI service provider) holds exclusive control over the cryptographic keys used to encrypt sovereign data. Sometimes implemented as a 'Bring Your Own Key' (BYOK) arrangement, this ensures that even data hosted on external infrastructure cannot be accessed by the vendor or foreign authorities without government authorisation.

Federated learning

A machine learning (ML) approach in which AI models are trained across multiple decentralised devices or servers (each holding local data) without the data itself leaving those locations. Only model updates (not raw data) are shared and aggregated. Federated learning is presented in this report as a key method for enabling cross-border AI collaboration while preserving data sovereignty.

Hardware attestation

A cryptographic process by which a computing system proves that its workloads are running on certified, approved hardware nodes. Hardware attestation enables governments to verify that sovereign data is being processed only on infrastructure that meets sovereignty and security requirements.

Hybrid cloud

An infrastructure model that combines domestic or sovereign cloud servers with commercial or regional cloud services. In this model, high-sensitivity data and critical workloads are retained on-premises or within national infrastructure, while less sensitive compute tasks are outsourced. Hybrid cloud is a recommended pathway for resource-constrained low- and middle-income countries (LMICs) seeking practical sovereignty without fully forgoing access to external AI tools.

Intellectual sovereignty

The dimension of data sovereignty that concerns ownership and control over knowledge assets derived from national data – particularly AI model

weights trained on government or social protection datasets. When model weights built from domestic data are owned or controlled by foreign vendors, a country loses intellectual sovereignty, even if the original data never leaves the country.

Jurisdictional exposure

The risk that a government's data or AI systems become subject to the laws, enforcement actions, or legal claims of a foreign jurisdiction – typically arising when data is stored on foreign infrastructure, when vendors are incorporated in foreign countries, or when AI models are trained abroad. Jurisdictional exposure is a central concern of this report.

Log residency

A policy that requires all telemetry, audit logs, and system activity records generated by AI or cloud infrastructure to remain within the national territory. Log residency ensures that the forensic and compliance record of AI system operations is subject to domestic jurisdiction and cannot be accessed or removed by foreign authorities.

Low- and middle-income country (LMIC)

A World Bank classification used throughout this report to refer to countries with low or lower-middle gross national income per capita. LMICs typically face greater constraints in relation to digital infrastructure, technical capacity, and regulatory frameworks, creating specific sovereignty risks in AI procurement and deployment.

Model weights

The numerical parameters within a ML model that encode the patterns and knowledge learnt during training. Model weights represent the intellectual output of an AI training process and, when derived from sovereign social protection data, are treated in this report as regulated data assets subject to the same sovereignty protections as the underlying data.

Personally identifiable information (PII)

Any data that can be used to identify a specific individual, directly or in combination with other data. In social protection contexts, PII includes names, national identification numbers, biometric records, addresses, household composition data, and benefit entitlement information. The processing of PII triggers the highest levels of sovereignty protection.

Privacy-enhancing technology (PET)

A category of technical tools and methods (including differential privacy, federated learning, synthetic data generation, and secure multi-party computation) that enable data analysis and AI training while minimising the exposure of sensitive personal data. PETs are recommended in this report as tools for enabling international AI collaboration without compromising sovereignty.

Public trust data

Data held by government agencies on behalf of citizens, where the government acts as custodian, rather than owner, and is obligated to use the data only in ways that serve beneficiaries' interests and uphold their rights. In this report, beneficiary data in social protection systems is characterised as public trust data, grounding a fiduciary duty on the part of government agencies.

Sensitive personal information

A category of personal data that warrants heightened protection due to the particular risks its misuse poses to individuals' rights and dignity. This includes biometric data, health and disability information, household vulnerability profiles, financial data, and information relating to race, ethnicity, religion, or political opinion. Social protection systems routinely process sensitive personal information.

Service level agreement (SLA)

A contractual document that defines the performance standards, availability commitments, and remediation obligations of a technology service provider. In the context of this report, SLAs should include sovereignty-relevant provisions such as data residency guarantees, incident response obligations, and audit cooperation requirements.

Sovereign cloud

A cloud computing environment that is fully subject to the laws, jurisdiction, and administrative control of a single national government, ensuring that data hosted within it cannot be accessed or removed by foreign authorities. Sovereign cloud deployments may be operated domestically by the government itself or by a vendor under strict contractual and legal arrangements that preserve national jurisdiction.

Sovereignty/Infrastructure Quadrant

An analytical framework introduced in this report that classifies AI use cases according to two dimensions – data sensitivity and infrastructure/runtime requirements – to determine the level of sovereignty control required. Quadrant I (high sensitivity, high compute) requires domestic or sovereign-cloud deployment; Quadrant II (high sensitivity, moderate compute) enables federated or hybrid governance; Quadrants III and IV (low sensitivity) permit greater use of commercial cloud services with contractual safeguards.

Sovereignty leakage
A condition in which a government loses effective control over its data or AI systems, even when the original data has not formally left national territory. Sovereignty leakage can occur when model weights trained on national data become proprietary assets of a foreign vendor, when encryption keys are held abroad, or when foreign administrators retain privileged access to domestic systems.

Standard contractual clause (SCC)
A standardised legal provision (often derived from or modelled on EU data protection frameworks) that can be incorporated into contracts with technology vendors to establish binding obligations on data processing, residency, security, and cross-border transfer. In this report, SCCs are a contractual mechanism recommended for governments procuring AI services from foreign providers.

Sub-processor
A third party engaged by a primary vendor to process data on the vendor's behalf, such as a cloud hosting provider or specialist AI service used as part of a broader solution. Sub-processors represent a sovereignty risk, because government data may flow to entities not directly contractually bound to the government. This report recommends that government approval be required for any sub-processor arrangements.

Synthetic data
Artificially generated data that statistically mimics the properties of a real dataset without containing actual personal records. Synthetic data can be used for AI model training, testing, and development in ways that do not expose real beneficiary data, reducing sovereignty risk while maintaining the utility of national datasets.

Vendor lock-in
A condition in which a government becomes dependent on a single technology provider due to proprietary system architectures, non-portable data formats, or exclusive control over AI models, making it difficult or costly to switch vendors or insource capabilities. Vendor lock-in is identified in this report as a key sovereignty risk, and open standards, DPGs, and interoperability requirements are recommended mitigations.

Zero data retention
A contractual and technical requirement that a vendor or service provider process data entirely in-memory during a transaction and delete all data (including logs, caches, and temporary files) immediately upon completion of the operation, with no persistent storage of social protection data on vendor infrastructure.

EXECUTIVE SUMMARY: ENSURING DATA SOVEREIGNTY IN THE CONTEXT OF AI FOR SOCIAL PROTECTION

What is data sovereignty and why is important?

Using artificial intelligence (AI) in social protection systems comes with profound implications for data sovereignty and data protection that governments can no longer afford to ignore. Social protection agencies hold some of the most sensitive personal data in government – identity records, household income, biometric markers, and social histories – and modern AI systems increasingly demand that this data be processed in complex cloud-based infrastructure owned and operated by foreign vendors.

For many low- and middle-income countries (LMICs), this convergence of highly sensitive data and externally controlled digital infrastructure exposes a fundamental governance challenge: *Who ultimately controls the data and the decisions derived from it?* Data sovereignty has, therefore, become a foundational requirement for responsible, trustworthy, and rights-based AI in social protection.

In the context of AI, sovereignty extends far beyond data storage – it encompasses the entire AI lifecycle, from data collection and model training through to inference, updating, and decommissioning. Sovereignty risks appear across this entire lifecycle, creating structural dependencies that undermine national governance and long-term digital autonomy. Governments face vulnerability when model training occurs abroad, encryption keys are controlled by external vendors, foreign administrators retain privileged access, and models trained on national datasets become proprietary assets outside local jurisdiction. Once data (or the derivative models built from it) fall under foreign legal regimes, the ability of national governments to regulate, audit, or intervene becomes significantly constrained and citizens' avenues for redress narrow accordingly.

Beneficiary data constitutes *public trust data*: governments hold it on behalf of citizens and bear a duty of care to protect it from misuse, overreach, or exploitation. When data leave the jurisdiction, or when the underlying AI systems cannot be inspected or controlled, transparency and accountability weaken, exacerbating the risk of exclusion, discrimination, and error. Ultimately, protecting data sovereignty is essential for safeguarding rights and maintaining public trust in digital social protection and cannot be ignored by social protection decision makers considering the use of AI.

This report examines how governments can deploy AI responsibly, while maintaining sovereign control over sensitive beneficiary data. It draws on the AI Hub's *Global Evidence Review* dataset of 99 AI-enabled social protection systems across 44 countries (DCI AI Hub, 2026a), as well as targeted stakeholder consultations with government officials, regulators, and AI practitioners. It forms part of the AI Hub's suite of foundational knowledge products under the Digital Convergence Initiative (DCI). The report focuses on how data sovereignty – the principle that data generated within a nation's borders should be governed by its own laws – intersects with the design, training, and hosting of AI systems in social protection. It is designed to be useful for policymakers, social protection practitioners and digital experts collaborating on planning, developing and deploying AI use cases.

Data sovereignty issues

The conceptual framework for this study consists of three factors that determine data sovereignty: the regulatory context (strictness of localisation requirements and cross-border data transfer models), a county's sovereignty readiness (infrastructure level and governance), and the dimensions of sovereignty risk (data sensitivity, computational demands and model training). These are complemented by a fourth area of analysis that offers an integrated framework to reduce AI related sovereignty risks and consists of four governance levers: policy, infrastructure, contracts, and capacity.

Based on this framework, the analysis identified four issues facing sovereignty that cut across country contexts:

- **Legal frameworks are expanding, but enforcement lags.** Most of the countries reviewed have enacted data protection or localisation laws, but few conduct audits or enforce cross-border data rules. Sovereignty remains largely declarative without operational oversight.
- **Vendor lock-in and model opacity create structural dependency.** Biometric and AI systems frequently rely on proprietary components that governments cannot audit, retrain, or migrate. No verified transitions between major AI vendors were found in LMIC contexts.
- **Sovereignty risks arise across the full AI lifecycle and cannot be confined to storage.** Training on foreign infrastructure can transfer both legal control and intellectual ownership of models derived from sovereign datasets – a dimension of risk routinely overlooked in procurement and governance frameworks.
- **Hybrid and federated models offer practical solutions.** Examples from Estonia, Singapore, Brazil, and Bangladesh demonstrate that sovereign and hybrid clouds, regional data embassy models, and federated learning frameworks can enable AI innovation without relinquishing national control.

Mitigating data sovereignty risk

To mitigate data sovereignty risk, the report proposes an integrated framework to reduce data sovereignty risk built on four interdependent levers, each essential for translating sovereignty from a legal aspiration into an operational reality, via a set of actions:

Lever 1.

Policy instruments – Enact clear localisation and cross-border data transfer laws that explicitly extend to AI training and model ownership. Treat model weights derived from sovereign datasets as regulated assets.

Lever 2.

Infrastructure controls – Invest in sovereign or hybrid cloud infrastructure, encryption key ownership, and compliance-as-code systems that enforce residency rules automatically.

Lever 3.

Contractual mechanisms – Use standardised data processing agreements (DPAs) and Standard Contractual Clauses (SCCs) to ensure local data processing, protect audit rights, and prohibit model reuse.

Lever 4.

Capacity building – Develop institutional, technical, and negotiation capacity to operationalise sovereignty and reduce dependence on external vendors.

A weakness in any single lever undermines the whole architecture. A strong localisation law without technical attestation remains unenforceable; sovereign infrastructure without skilled administrators is insecure; and contractual safeguards without domestic audit capacity are merely symbolic. These levers form the foundation of the Sovereignty Implementation Matrix (see Table 4, Chapter 8), which aligns practical actions such as hosting policies, encryption standards, and contractual clauses with specific sovereignty risks.

The Integrated Framework to Reduce AI-Related Sovereignty Risks in the following figure ensures that AI in social protection is implemented in a way that strengthens rather than erodes data sovereignty, by strengthening legal and governance frameworks; adopting sovereignty-aligned hosting and infrastructure pathways; transforming procurement and contracting; and building institutional and technical capacity.

AN INTEGRATED FRAMEWORK TO REDUCE AI-RELATED SOVEREIGNTY RISKS

POLICY INSTRUMENTS

Laws, regulations, and institutional authority

1. Data protection and localisation laws and policies should explicitly extend to AI training and model management.
2. Mandate data transfer impact assessments and regulator approval for cross-border processing of sensitive personal information.
3. Empower data protection or AI oversight bodies with jurisdiction to audit AI systems used in social protection.
4. Adopt beneficiary data as public trust clauses in social protection law.

INFRASTRUCTURE CONTROLS

Hosting, encryption, and technical assurance

1. Sovereign and hybrid cloud models are critical for sensitive data.
2. Implement encryption key ownership, hardware attestation and log residency policies.
3. Adopt compliance as code so deployment pipelines automatically enforce residence policies.
4. Use sovereign clouds or explore data embassy models for sovereign model training and hosting.

CONTRACTUAL MECHANISMS

Standard clauses and vendor accountability

1. Insert standard contractual clauses mandating local processing and deletion after use.
2. Negotiate data processing agreements forbidding model training on government data.
3. Require zero retention for sub processer approval and audit rights.
4. Define model ownership and IP clauses for AI models trained on sovereign data.

CAPACITY BUILDING

Institutions, skills, and collaborative ecosystems

1. Establish sovereign data or AI oversight units.
2. Develop technical expertise in cloud encryption and federated learning governance.
3. Strengthen legal negotiation skills for digital contracts.
4. Foster regional federated learning consortia to share models without sharing data.
5. Enhance citizen literacy and grievance mechanism on data rights.

Conclusion

AI can strengthen social protection delivery only when it reinforces – not undermines – national sovereignty. Sovereignty is not a static legal claim, but an operational capability: the ability of a state to decide, enforce, and adapt how its data and AI systems are governed. By embedding sovereignty principles into law, infrastructure, contracts, and institutions, governments can ensure that digital transformation enhances both efficiency and autonomy.

Responsible AI adoption requires sovereign AI capacity. Without it, social protection systems risk trading short-term technological gains for long-term dependency. With it, AI becomes a tool for empowerment – anchoring digital governance in the rights, dignity, and trust of the citizens it is meant to serve.

CHAPTER 1. UNDERSTANDING DATA SOVEREIGNTY

Artificial intelligence is rapidly reshaping social protection systems (Henman, 2020; ISSA & UNU-EGOV, 2024; OECD, 2025a, 2025b; DCI AI Hub, 2026a), yet its deployment is advancing faster than the capacity of governments to maintain control over the data and models that underpin these services. Social protection agencies have long understood the importance of safeguarding personal data within their jurisdictions, but the emergence of advanced, computation-intensive AI tools is making this significantly harder to achieve. Modern AI systems depend on vast quantities of data, often including highly sensitive information such as identification records, household income, biometric markers, and social histories (Eubanks, 2018; Dobson, 2019). They may also require complex cloud-based infrastructure owned and operated by foreign vendors.

For many low- and middle-income countries (LMICs), this convergence of highly sensitive data and externally controlled digital infrastructure exposes a deeper governance challenge: *Who ultimately controls the data and the decisions derived from it?* (Rikap et al., 2024).

Data sovereignty is, therefore, a foundational requirement for responsible, trustworthy, and rights-based AI in social protection. In the context of AI, sovereignty extends beyond storage to the entire AI lifecycle – covering data collection, model training, inference, and maintenance. It ensures that data used for social protection remains subject to national laws, ethical standards, and oversight mechanisms. Without such control, governments risk dependence on external actors, exposure to foreign jurisdictions, and the erosion of citizens' trust in public institutions.

Sovereignty risks appear across the entire AI lifecycle, creating structural dependencies that undermine national governance and long-term digital autonomy. Governments face vulnerability when model training occurs abroad, encryption keys are controlled by external vendors, foreign administrators retain privileged access, and models trained on national datasets become proprietary assets outside the local jurisdiction. Once either data or derivative models fall under the control of foreign legal regimes, the ability of governments to regulate, audit, or intervene becomes significantly constrained and citizens' avenues for redress narrow accordingly.[1]

Protecting data sovereignty is essential for safeguarding rights and maintaining public trust in digital social protection. Beneficiary data are public trust data: governments hold it on behalf of citizens and have a responsibility to protect it from misuse, overreach, or exploitation (World Bank, 2021; Wagner et al., 2024). When data leave the jurisdiction – or when the underlying AI systems cannot be inspected or controlled – transparency and accountability weaken, exacerbating the risk of exclusion, discrimination, and error. Global governance norms such as the European Union's (EU's) General Data Protection Regulation

1 As Microsoft's Legal Director for France noted to the French Parliament in 2025: Microsoft "cannot guarantee" that French citizens' data held through public contracts would remain protected from U.S. authorities without local consent (Wodecki, 2025).

(GDPR), UNESCO's Recommendation on the Ethics of AI, and emerging frameworks in Africa, Asia, and Latin America affirm that governments must retain authority over public-sector data to uphold the rule of law in digital systems (European Union, 2016; UNESCO, 2021).

Ultimately, the issue of data sovereignty when developing AI systems for the social protection sector is too risky to be ignored (see Box 1).

BOX 1. RISKS OF IGNORING DATA SOVEREIGNTY

- **Loss of control over citizens' most sensitive data:** Exposure to foreign jurisdictions and compelled access undermines national authority and privacy protections.
- **Opaque and unaccountable automated decisions:** External hosting or training limits governments' ability to audit or contest AI-driven eligibility and fraud decisions.
- **Long-term dependency on external vendors:** Proprietary platforms, non-portable models, and remote administrative rights entrench reliance and reduce digital autonomy.
- **Erosion of public trust in social protection systems:** Perceived or actual loss of national control over identity and beneficiary data can delegitimise reforms and amplify social harms.
- **Fiscal exposure and forced discontinuation:** Escalating licensing, hosting, or support costs, combined with limited exit options, can leave governments unable to sustain arrangements, risking abrupt service disruption for vulnerable beneficiaries when alternatives are unavailable.

This report provides a practical roadmap for implementing AI in social protection in a way that preserves national authority over data, models, and infrastructure. It explores how governments can implement AI systems for social protection that are both technically feasible and sovereignly governed. It examines the criteria for selecting and contracting AI vendors, hosting models that align with national jurisdictional requirements, and safeguards to prevent data misuse or cross-border leakage.

Chapter 2 introduces the **Methodology and Scope** for the study, followed by Chapter 3 which sets out the **Conceptual Framework**. Chapter 4, **Global Evidence Review: How AI is Currently Deployed in Social Protection**, provides a synthesises of the global evidence, giving a snapshot of how countries are faring in terms of data sovereignty. Chapter 5 on **The Regulatory Context** defines how countries legally assert control over their data and under what conditions it may be stored, processed, or transferred abroad. Chapter 6, **Data Sovereignty Readiness**, categorises countries as advanced, emerging or nascent depending on their infrastructure and data governance framework, data protection laws, legal basis for algorithmic accountability, explainability, and redress, as well as their capacity to negotiate with vendors. Chapter 7, **Dimensions of Data Sovereignty Risk**, identifies where risk is the highest based on three interrelated dimensions: data sensitivity, infrastructure intensity, and model training location and control. It introduces the Sovereignty/Infrastructure Quadrant, which is a practical tool for policymakers and implementers to assess the trade-offs between data sovereignty and technical feasibility. The Quadrant links the regulatory context with risk mitigation strategies, translating legal principles into operational choices that social protection agencies can apply when planning to deploy AI in social protection. Chapter 8, **An Integrated Framework for Reducing Data Sovereignty Risk**, provides practical ways of mitigating data sovereignty risks via coordinated action across four interdependent governance levers – policy, infrastructure, contracts, and capacity – each of which must operate effectively for governments to retain control over data and AI systems. This chapter concludes with considerations on how the four levers can be turned into a **Sovereignty Strategy** via a **Sovereignty Implementation Matrix**. Finally, Chapter 9, **Conclusion**, provides a brief conclusion. The Annexes offer resources for practitioners including a **Global Typology of Data Sovereignty Regimes** (Annex 1), a **Vendor Checklist for AI Procurement in Social Protection** (Annex 2) and **Sample Contractual Clauses** (Annex 3).

CHAPTER 2. METHODOLOGY AND SCOPE

2.1 Methodology

A mixed-method approach was used to assess how AI can be implemented in social protection while upholding data sovereignty. Legal, technical, and institutional analysis was undertaken through a focused desk review, supplemented by targeted stakeholder consultations and applied use-case assessments to ensure both analytical depth and practical relevance.

The desk review, including comparative legal mapping, identified the main regulatory factors shaping sovereignty risks in AI systems. The review examined global and regional frameworks on data protection, localisation, cross-border transfers, and AI regulation, including the EU's GDPR, China's Personal Information Protection Law (PIPL), and African and Association of Southeast Asian Nations (ASEAN) initiatives. It identified key sovereignty regimes, localisation models, and transfer mechanisms that affect AI deployment, and assessed emerging concepts such as sovereign cloud options and privacy-preserving AI for their relevance in LMIC contexts.

Stakeholder consultations provided practical insights into infrastructure constraints, enforcement challenges, and trade-offs in sovereign AI adoption. Five key informant interviews were conducted with government agencies, social protection institutions, regulators, and AI experts, highlighting feasibility issues in hosting AI systems, difficulties enforcing sovereignty mandates, and tensions between innovation, cost, and national control. These insights helped refine the assessment of real-world sovereignty risks such as vendor lock-in, opaque data flows, and limited governance capacity.

To understand real-world constraints, the study analysed use-cases from the AI Hub's *Global Evidence Review* of 99+ AI-enabled social protection systems across 44 countries, focusing on vendor patterns, hosting choices, and portability limitations (DCI AI Hub, 2026a). This empirical grounding served as a methodological tool to assess structural dependencies, including: reliance on foreign infrastructure or humanitarian platforms; prevalence of proprietary biometric or other AI models; absence of verifiable vendor transitions; common gaps in data protection agreements, standard contractual clauses, or data-location clauses; and frequency of cross-border data flows without documentation.

By combining literature, stakeholder perspectives, and system-level evidence, the methodology situates sovereignty considerations in actual implementation practices rather than theoretical assumptions.

AI tools were used selectively to support efficiency and consistency, with all outputs fully verified by human researchers. ChatGPT-5 and Claude 4.5 supported targeted literature searches, terminology checks, and preliminary drafting, but all content was reviewed, validated, and edited by the research team. No unverified AI-generated material was included, and AI use complemented rather than replaced expert judgment.

2.2 Scope

This report adopts an institutional perspective on data sovereignty. In other words, the primary unit of focus is agencies responsible for delivering social protection programmes, which may map to different levels depending on the context.

The term 'data sovereignty' is understood to be a dimension of data protection that focuses on national control over sensitive data and its processing; in other words, the right of a country to determine where and how its sensitive data is stored, who can process that data and under what conditions, and how cross-border data flows occur so that the country's data protection standards are upheld and its authorities can enforce those standards. Hence, this study concentrates only on those aspects of data protection that intersect with data sovereignty. It does not cover broader areas of data protection as they relate to AI (such as the re-purposing of data, algorithmic design or the use of AI for decision-making).

This framing complements:

- The *AI Hub Risk Assessment Framework for Social Protection: A Practical Guide to Responsible AI Deployment* (DCI AI Hub, 2026b), which considers the broader set of potential risks of AI, including those related to aspects of data protection
- The DCI-endorsed ***Implementation Guide – Good Practices for Ensuring Data Protection and Privacy in Social Protection Systems*** (Wagner et al., 2024), which discusses the overarching framing for data protection

The analysis does not consider broader questions of AI sovereignty, but acknowledges links to national capacity-building. Discussions on the development of sovereign AI capacity intersect with this report, but are considered only in relation to the importance of building technical capacity to understand and use AI responsibly.

CHAPTER 3. CONCEPTUAL FRAMEWORK

The conceptual framing for this study hinges on three interlocking areas that help to assess data sovereignty risk: the regulatory context, dimensions of sovereignty 'readiness' and dimensions of sovereignty risk. These are complemented by a fourth area of analysis consisting of four governance levers: policy, infrastructure, contracts, and capacity.

3.1 The regulatory context: Shaping what is legally allowed

To understand how sovereignty obligations shape AI implementation options, the study incorporated a structured regulatory-mapping exercise covering data protection, localisation, cross-border data transfer, and AI-specific legislation across multiple jurisdictions. This component classified countries into strictness tiers (from full localisation to unrestricted transfer regimes) and legal models (e.g. adequacy, permit-based, blacklist systems). This enabled a comparative analysis of how national laws determine permissible hosting arrangements, model-training locations, vendor obligations, and minimum safeguards for AI in social protection. The mapping also informed the risk classification of AI use cases, particularly where regulatory asymmetries create vulnerabilities for LMICs.

3.2 Sovereignty readiness

The assessment of legal framings is accompanied by a structured assessment of national digital infrastructure across three functional levels, enabling the analysis of how technical capacity determines sovereignty risks and potential implementation pathways. Drawing on the *Global Evidence Review* (DCI AI Hub, 2026a), countries were categorised as advanced, emerging, or nascent in terms of compute availability, cloud maturity, government data centres, and digital ID ecosystems. This tiering approach was used to consider in which cases domestic hosting is technically feasible for LMICs and, therefore, which AI use cases may require sovereign or hybrid cloud models.

3.3 Dimensions of sovereignty risk

A structured framework was used to classify AI use cases according to data sensitivity, computational demands, and model training, enabling a systematic assessment of sovereignty risks across the AI lifecycle. This risk-based approach aligns with global data-protection practices and grounded the study's recommendations in proportional, context-appropriate safeguards.

FIGURE 1. CONCEPTUAL FRAMEWORK: REGULATORY CONTEXT, SOVEREIGNTY READINESS AND SOVEREIGNTY RISK

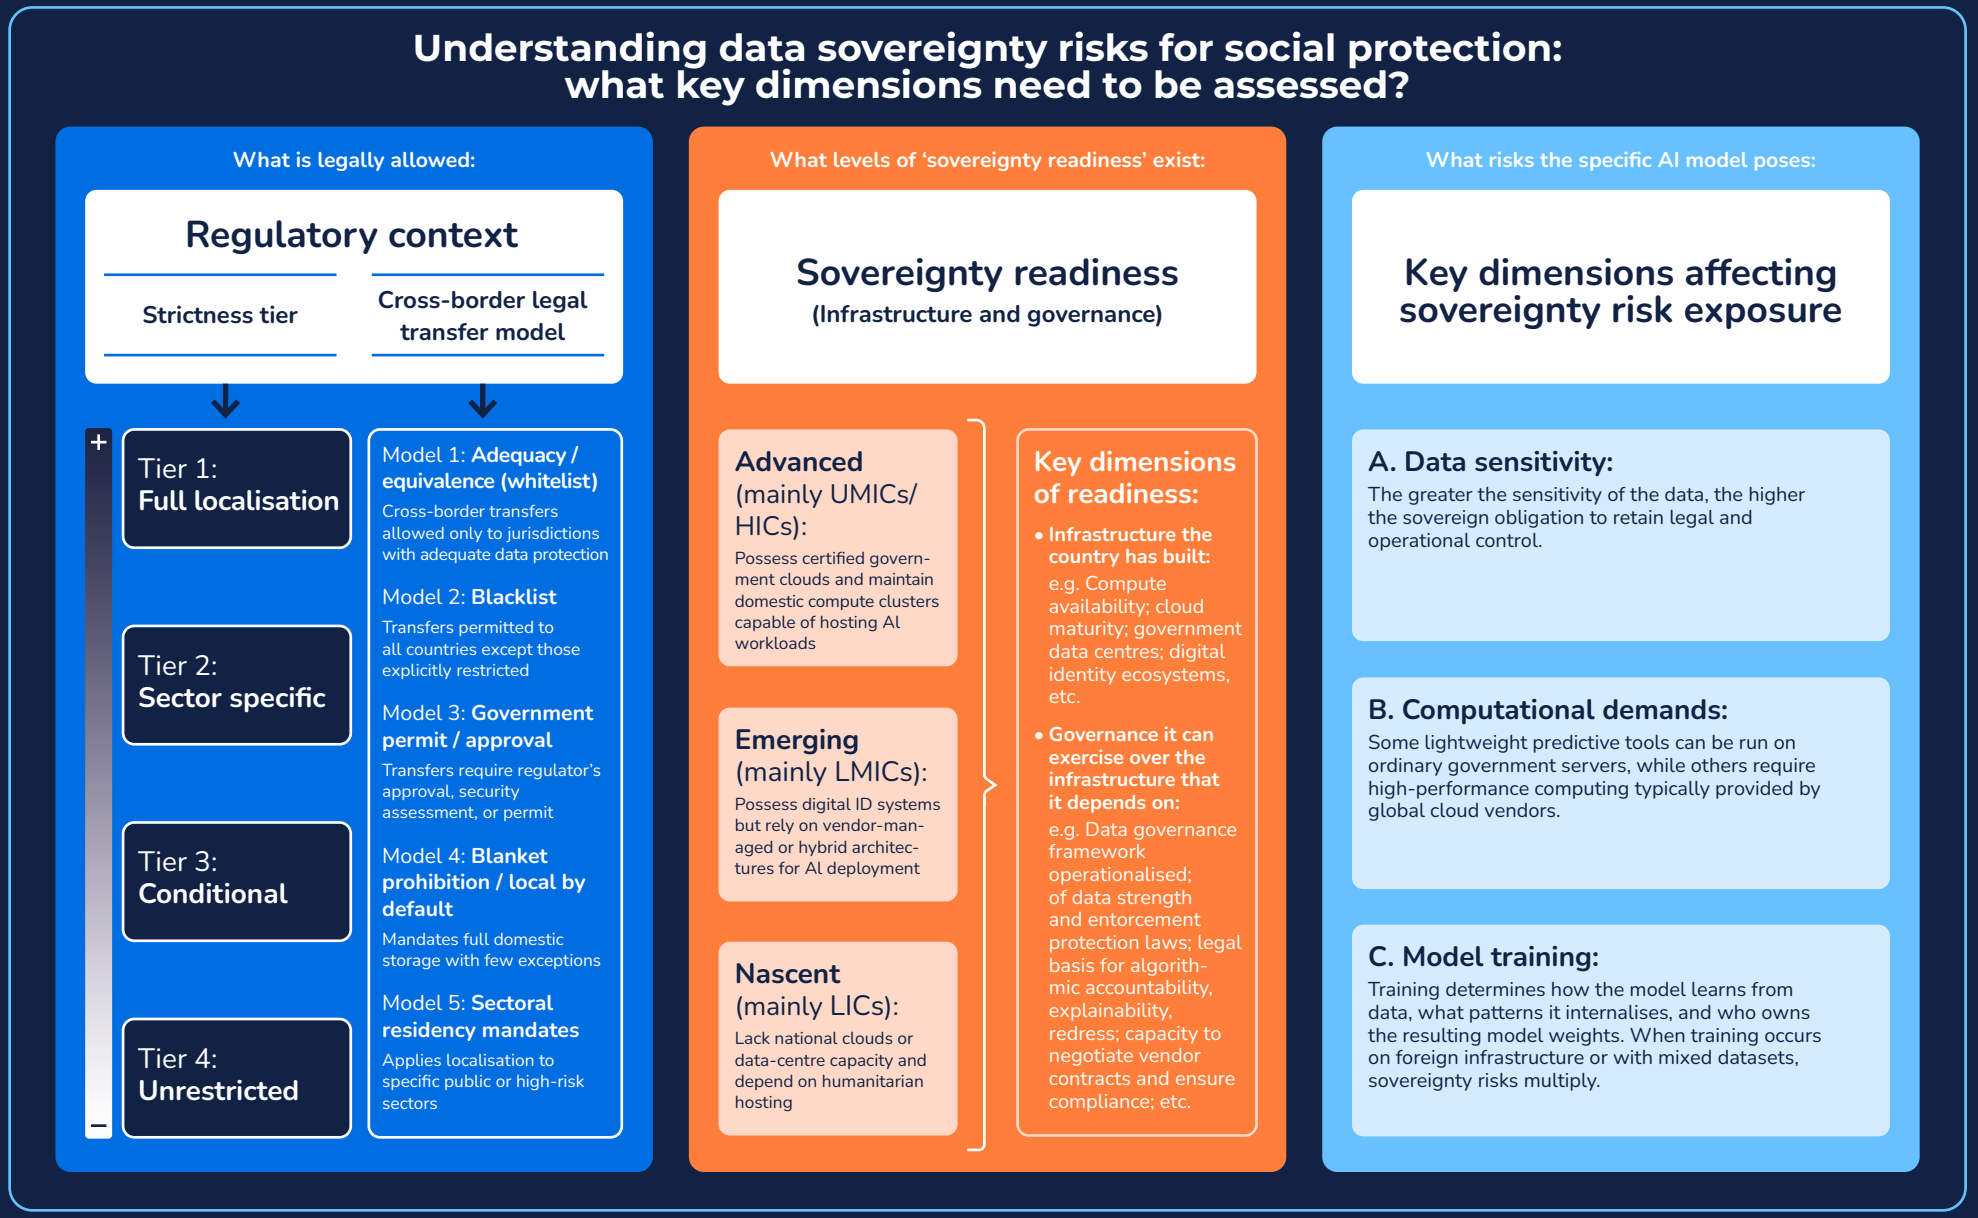

3.4 Four governance levers of data sovereignty

Finally, the three dimensions of data sovereignty are complemented by four governance levers: policy, infrastructure, contracts, and capacity. These levers are areas in which action must be taken to reduce data sovereignty risk. Each lever reinforces the others and must operate effectively for governments to retain control over data and AI systems. Together, they determine whether or not governments can retain legal, operational, and institutional control over the data and AI systems that underpin social protection.

FIGURE 2. THE FOUR GOVERNANCE LEVERS OF DATA SOVEREIGNTY

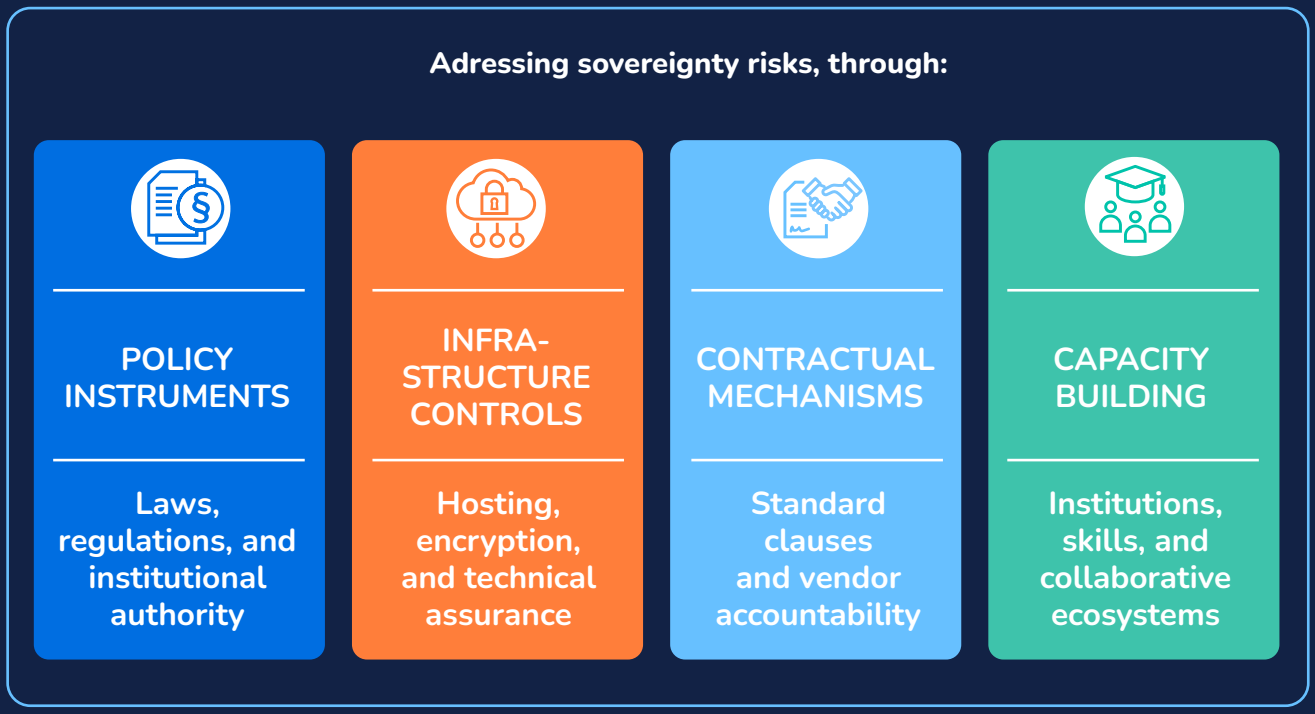

CHAPTER 4. GLOBAL EVIDENCE REVIEW: HOW AI IS CURRENTLY DEPLOYED IN SOCIAL PROTECTION

This chapter gives a snapshot of how AI is currently deployed in social protection systems around the world and the implications for data sovereignty. Drawing on the *Global Evidence Review* (DCI AI Hub, 2026a), the analysis examines patterns in technology adoption, hosting architectures, vendor dependency, and oversight arrangements. The evidence provides an empirical basis for understanding how sovereignty risks materialise in practice, linking directly to the conceptual framework.

4.1 Patterns in AI use

The *Global Evidence Review* found that AI adoption in social protection is dominated by traditional analytical models, but there are key differences in the deployment of biometric and generative AI. Nearly three-quarters (74%) of the 99 cases studied involve traditional machine learning (ML) tools such as predictive targeting, poverty mapping, fraud detection, and risk scoring – a trend consistent across income levels. In contrast, all 16 biometric AI deployments identified (16%) occur exclusively in low-income countries (LICs) and LMICs, accounting for nearly one-third of the cases identified. These focus primarily on identity verification, social insurance authentication, and humanitarian distribution controls (DCI AI Hub, 2026a).

The adoption of generative AI is concentrated in high-income countries (HICs), reflecting disparities in infrastructure, governance capacity, and model availability. Ninety percent of all generative AI use cases – mainly virtual assistants and chatbots – are implemented in countries such as Singapore, Estonia, the United Kingdom, and the United States. Of the 99 use cases identified, only 2% are in LMICs. The pattern reflects both capability gaps and the recency of large language model integration into government systems.

BOX 2. KEY STATISTICS FROM THE AI HUB'S GLOBAL EVIDENCE REVIEW

Of the 99 AI-enabled social protection systems studied:

- 74% use AI use traditional ML; the use of biometric AI (16%) is exclusive to LMICs.
- 90% of generative AI use cases occur in HICs.
- In LMICs, 80–97% of systems have zero hosting transparency.
- Only one-third of systems show confirmed domestic/sovereign hosting.
- At least 8% exhibit signs of vendor lock-in; real levels are likely higher.
- No verified vendor transitions from LMIC governments have been documented.
- No publicly available algorithmic audits or bias assessments were found across cases.
- The top cloud hyper-scalers (e.g. AWS, Azure and Google Cloud) control ~70% of the global cloud market, reinforcing structural dependency.
- Between 80 and 97% of cases in LICs and LMICs provide no public hosting information, and almost none demonstrate sovereign environments.

Source: DCI AI Hub (2026a)

4.2 Hosting, transparency, and vendor dependence

Across the 44 countries reviewed, data-hosting arrangements for AI systems remain largely opaque, posing significant risks to sovereignty and compliance. Two-thirds of the cases studied offer no public documentation of where data is stored or processed and only one-third show evidence of local or sovereign hosting (DCI AI Hub, 2026a). The absence of confirmed overseas hosting does not imply domestic control; rather, it suggests chronic under-reporting, particularly in the commercial cloud environments commonly used in AI deployments. Without clarity on data location, governments cannot evaluate compliance with data-residency laws or assess exposure to foreign jurisdictions.

Hosting transparency gaps are most pronounced in LMICs, where sovereignty controls are weakest. HICs account for nearly all cases of verified sovereign or domestic hosting (61%). In contrast, between 80 and 97% of cases in LICs and LMICs provide no public hosting information, and almost none demonstrate sovereign environments. This suggests potential hidden dependence on external vendors and foreign cloud services in settings least equipped with safeguards or regulatory oversight.

Vendor opacity mirrors hosting opacity, limiting the ability of governments to evaluate dependency and sovereignty risks. Nearly one-third of cases do not identify the vendor or technology provider. Among those disclosed, providers range from government-built platforms (GovTech Singapore, Estonia's Bürokratt, Germany's DRV Bund) to a small set of proprietary firms (SAS, Wysa, Ignitia). The gaps are particularly acute in low-income regions, where reliance on external vendors is likely higher than the dataset reveals.

Even based on conservative criteria, there is clear evidence of vendor lock-in. At least 8% of systems exhibit confirmed indicators of lock-in, primarily overseas hosting or reliance on proprietary platforms with non-portable architectures (also see Rikap et al., 2024). Given incomplete hosting and vendor data, this figure almost certainly underestimates the extent of structural dependency across systems.

BOX 3. MARKET CONCENTRATION OF AI CLOUD PROVIDERS

As of mid-2025, the global cloud infrastructure (IaaS/PaaS) market remains heavily concentrated among a few hyperscalers, who also lead in AI-cloud services. The top five providers control roughly 70% of the global cloud infrastructure market, with the top three (AWS, Azure, Google Cloud) responsible for approximately 63% of the market. This illustrates how AI-driven demand is reinforcing concentration around major hyperscalers.

Source: Synergy Research

4.3 Oversight and governance gaps

Operational oversight of AI systems remains extremely limited, with legal frameworks rarely translating into practice. Across all use cases, there is scant publicly available evidence of algorithmic audits, bias testing, human-rights assessments, or performance monitoring. Systems deployed through humanitarian actors further complicate visibility, as they likely operate under internal organisational protocols rather than national regulatory regimes. As a result, most AI systems in social protection run without documented safeguards or accountability mechanisms.

Oversight gaps are most acute in LICs, where institutional mandates, coordination mechanisms, and regulatory capacity remain weak. Few LICs or LMICs demonstrate formal governance structures, such as data protection authority involvement, AI-specific procurement rules, or classification frameworks for high-risk systems. HICs show emerging examples, including Estonia's audited X-Road infrastructure and Singapore's integration of AI governance into procurement. These are exceptions rather than the norm.

Taken together, these patterns highlight that governance, not only technology, shapes sovereignty outcomes. Weak transparency, absent oversight, and fragmented accountability create conditions under which AI systems can operate outside national jurisdiction, even when laws nominally require local control.

4.4 Infrastructure levels and sovereignty readiness

In relation to sovereignty readiness, countries can be categorised as advanced, emerging, or nascent, depending on their infrastructure model. Compute availability, cloud maturity, sovereign hosting, and digital ID integration predicts sovereignty capacity more accurately than income level (World Bank, 2024; DCI, 2025). Several LMICs (India, Indonesia, Brazil, South Africa, Thailand, Malaysia, Vietnam, the Philippines) demonstrate advanced capabilities through sovereign or hybrid clouds, graphics processing unit (GPU) clusters, and advanced government platforms. Others with similar income levels (Bangladesh, Pakistan, Kenya, Rwanda) remain at the emerging or nascent level due to lack of sovereign compute or reliance on external hosting. This confirms that sovereignty readiness reflects deliberate investment and governance choices rather than gross domestic product (GDP).

4.5 Emerging sovereign and hybrid pathways

Despite constraints, several countries demonstrate feasible pathways to strengthen sovereign control over AI systems. Estonia's X-Road architecture shows how open, federated infrastructure with transparent APIs, domestic hosting, and routine audits can institutionalise sovereignty. Singapore's National Digital Identity (NDI)/SingPass ecosystem embeds AI governance and local residency directly into design and procurement. Bangladesh's Disaster and Climate Risk Management (DCRM) platform illustrates how humanitarian-led systems can transition into nationally-owned architectures through capacity building and phased stewardship transfer.

4.6 Implications for data sovereignty

The evidence demonstrates that data sovereignty risks are fundamentally shaped by uneven infrastructure maturity and dependency on external compute environments. Emerging and nascent countries depend heavily on foreign hosting and training environments, creating exposure to external jurisdictions, particularly for high-sensitivity AI applications. Legal localisation mandates cannot be enforced operationally without investment in sovereign or hybrid cloud capacity, national key management, and transparent vendor architectures.

Limited domestic infrastructure also heightens the risk of vendor lock-in, opaque data flows, and loss of intellectual sovereignty over models trained on sovereign datasets. These risks underline that high-sensitivity AI use cases must not be deployed on infrastructure outside national control, and design choices must be based on actual – not assumed – sovereignty capacity.

BOX 4. THE OTHER SIDE OF THE EQUATION – VENDOR HOST COUNTRY OBLIGATIONS

The governance burden this report addresses falls almost entirely on recipient countries: build better procurement frameworks, negotiate stronger contracts, invest in sovereign infrastructure. Yet the structural power in these relationships sits elsewhere – with the vendors and with the countries that host them.

Most AI systems deployed in social protection in LMICs originate from a small number of high-income jurisdictions. These vendor host countries typically have robust domestic data protection and AI governance regimes. What they largely lack are mechanisms that extend those obligations extraterritorially – to govern how their companies behave when operating in markets with weaker regulatory capacity.

This gap is not unique to AI. Analogous debates have played out in extractive industries (France's Duty of Vigilance Law), financial services (Financial Action Task Force [FATF] home-country obligations), and pharmaceutical exports. The logic in each case is the same: where there is structural asymmetry between a powerful exporter and a less-regulated recipient market, home country accountability helps fill the enforcement vacuum. For AI vendors, this could take several forms: mandatory human rights and data protection due diligence before deploying high-risk AI systems abroad; extraterritorial application of domestic AI regulations to overseas government contracts; export notification or approval requirements for biometric and predictive systems; and home-country liability where practices would be unlawful domestically.

Until such frameworks emerge, the responsibility for protection remains disproportionately (and unfairly) with the countries least equipped to exercise it.

CHAPTER 5. THE REGULATORY CONTEXT

When evaluating the risk posed to data sovereignty by AI in social protection, the first important factor to consider is the regulatory context: how countries legally assert control over their data and under what conditions it may be stored, processed, or transferred abroad. For social protection systems adopting AI, this context is the foundation upon which technical and operational decisions must rest. It determines whether AI models can be hosted in-country or rely on cross-border services, what contractual safeguards must be in place, and how national authorities can enforce compliance. Understanding these frameworks is, therefore, essential for governments seeking to implement AI in ways that preserve both privacy and sovereignty.

FIGURE 3. WHAT DOES THE REGULATORY CONTEXT LEGALLY ALLOW?

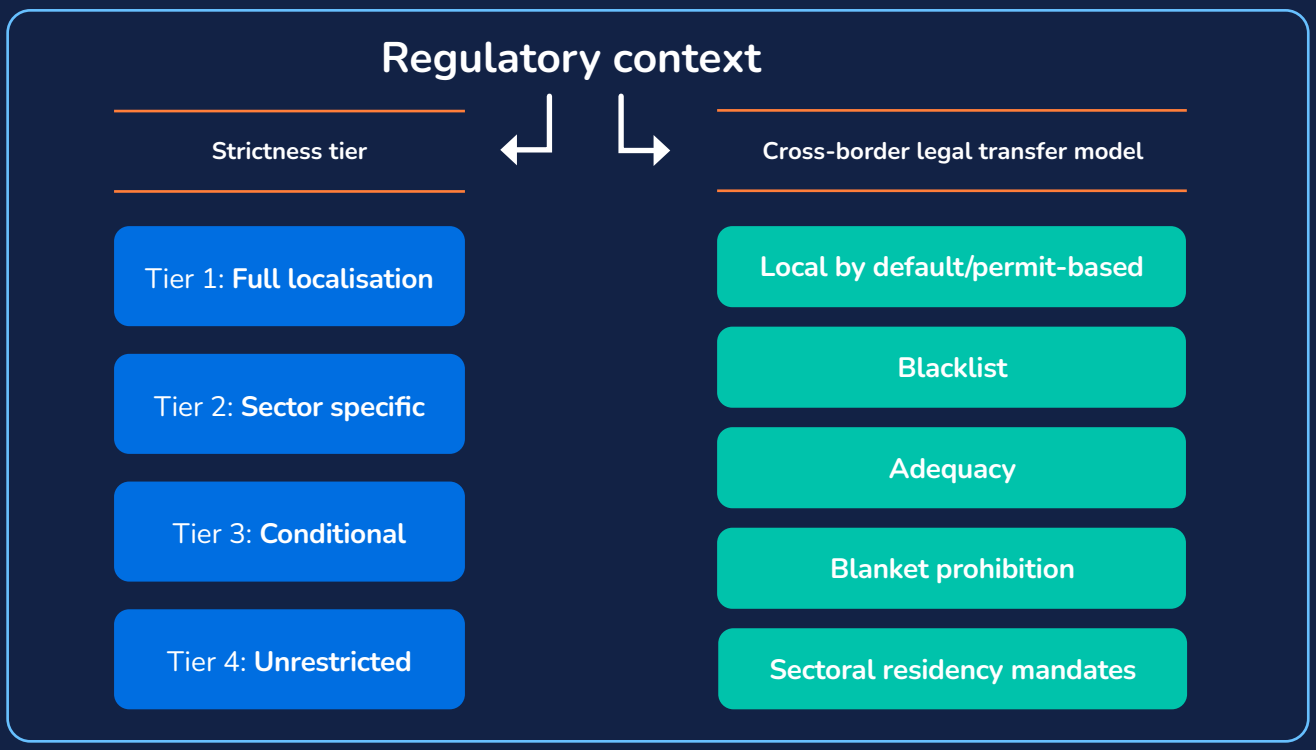

5.1 Country typologies

Our review found that several countries have implemented AI-specific legislation that is relevant to data sovereignty, but this is not yet widespread. To interpret these variations systematically, this study applied typologies of data sovereignty regimes. Countries were classified according to: the strictness of localisation requirements (see Table 1), and the legal mechanism governing cross-border transfers (see Table 2). The country examples on which the typologies are based are provided in Annex 1.

TABLE 1. TYPOLOGY BY STRICTNESS OF LOCALISATION REQUIREMENTS (4 TIERS)

TIER	STRICTNESS	DESCRIPTION	EXAMPLES (2024–2025)	IMPLICATIONS FOR AI IN SOCIAL PROTECTION
Tier 1 **Full/blanket localisation (local-by-default)**	Highest strictness	Personal or critical data must be stored domestically; exports are allowed only through government authorisation or in anonymised form.	**China** – PIPL (2021), DSL (2021), CSL (2017), CAC Measures (2022–2024); **Russia** – Federal Law No. 242-FZ (2015); **Uzbekistan** – Law on Personal Data (2021 amendment)	AI models must be hosted on sovereign or national infrastructure. External training or inference involving personal data is generally prohibited.
Tier 2 **Conditional/ sector-specific localisation**	Moderate strictness	Localisation is required for select sectors (public, health, finance, or telecom).	**Indonesia** – GR 71/2019; **Nigeria** – NITDA Cloud Policy (2019); **Kenya** – Cloud Policy (2025	Social protection data commonly falls within public-sector residency mandates. Hybrid or sovereign clouds are suitable for compliance.
Tier 3 **Conditional cross-border transfer regimes**	Controlled transfers	Transfers are allowed only if recipient jurisdiction provides adequate protection or safeguards.	**EU** – GDPR (2016); **South Africa** – POPIA (2020); **India** – DPDP Act (2023, blacklist model); **Vietnam** – PDP Decree (2023)	Enables use of regional or international cloud services if contractual clauses (e.g. SCCs, BCRs) ensure equivalent protection.
Tier 4 **Unrestricted data flows**	No overall controls	There are no general restrictions on transfers; sector specific or sub-national legislation may apply.	**United States** – HIPAA (1996), United States (2024–2025)	Review related legislation (e.g. HIPAA) which may impose restrictions.

Note: Personal Information Protection Law (PIPL); Data Security Law (DSL); Cybersecurity Law (CSL); Cyberspace Administration of China (CAC); Government Regulation (GR); National Information Technology Development Agency (NITDA); General Data Protection Regulation (GDPR); Protection of Personal Information Act (POPIA); Standard Contractual Clause (SCC); Binding Corporate Rule (BCR); Digital Personal Data Protection (DPDP); Personal Data Protection (PDP); Health Insurance Portability Act (HIPPA)

As well as the strictness of localisation requirements, it is important to look at the legal mechanisms governing cross-border transfers.

TABLE 2. TYPOLOGY OF LEGAL MECHANISMS GOVERNING CROSS-BORDER TRANSFERS (5 MODELS)

MODEL	DESCRIPTION	REPRESENTATIVE EXAMPLES	POLICY LOGIC AND RELEVANCE FOR AI
Model 1 **Adequacy / equivalence (whitelist)**	Cross-border transfers are allowed only to jurisdictions with adequate data protection or through binding contractual safeguards (e.g. SCCs, BCRs).	**EU** (GDPR); **South Africa** (POPIA); **Argentina** (PDPA); **Brazil** (LGPD)	Provides a high level of protection, but demands strong regulatory capacity; appropriate for Tier 3 frameworks balancing sovereignty with trade.
Model 2 **Blacklist**	Transfers are permitted to all countries except those explicitly restricted.	**India** (DPDP Act, 2023); **United States** (2025)	Business-friendly; retains flexibility while maintaining sovereign control through selective prohibition.
Model 3 **Government permit / approval**	Transfers require regulator's approval, security assessment, or permit.	**Egypt** (PDPL, 2020); **Vietnam** (PDP Decree, 2023); **China** (CAC measures, 2022–2024)	Maximises government oversight; suitable for sensitive sectors such as social protection and defence.
Model 4 **Blanket prohibition / local-by-default**	Mandates full domestic storage with few exceptions.	**China, Uzbekistan, Russia**	Guarantees sovereignty, but limits access to global AI infrastructure and innovation ecosystems.
Model 5 **Sectoral residency mandates**	Applies localisation to specific public or high-risk sectors.	**Indonesia, Nigeria, Philippines, Kenya**	Effective transitional approach for LMICs – protects critical data while enabling hybrid AI use.

Note: General Data Protection Regulation (GDPR); Protection of Personal Information Act (POPIA); Personal Data Protection Law (PDPA); General Personal Data Protection Law (LGPD); Digital Personal Data Protection (DPDP); Personal Data Protection Law (PDPL); Personal Data Protection (PDP); Cyberspace Administration of China (CAC)

5.2 Global trends and regional differentiation

The global trend is toward conditional localisation – balancing national control with participation in international digital ecosystems (OECD, 2023). Countries increasingly recognise that strict isolation (Tier 1) can limit access to advanced AI tools, while unrestricted data flows (Tier 4) can undermine public trust and state authority. This has led to the rise of hybrid approaches (Tiers 2 and 3), in which governments define *which* categories of data – such as biometric, financial, or social protection information – must remain in-country, while allowing other types to move under regulated conditions.

BOX 5. KENYA – REGULATORY ENFORCEMENT AND AI GOVERNANCE: FROM ENFORCEMENT ACTION TO ANTICIPATORY ARCHITECTURE

Kenya's experience across three interconnected developments offers the most multi-dimensional LMIC case study in the regulatory enforcement of AI and data governance. The centrepiece is the Worldcoin enforcement action. In 2023, Tools for Humanity collected biometric iris and facial data from over 350,000 Kenyans (roughly 25% of the company's global customer base at launch) without an adequate legal basis. The Office of the Data Protection Commissioner found the company liable for privacy breaches and, in May 2025, the High Court declared all data collection unlawful and ordered permanent deletion of all biometric data within seven days under Office of the Data Protection Commissioner (ODPC) supervision. Deletion was confirmed in January 2026. The International Commission of Jurists Kenya described the ruling as a global precedent.

This enforcement backdrop is significant because Kenya is simultaneously deploying AI in social protection at scale. The Social Health Authority's 2025 rollout of biometric verification across Level 4–6 health facilities, covering over 22 million registered members, operates directly under the Data Protection Act 2019 and its General Regulations. The regulatory enforcement regime that constrained Worldcoin applies equally to the Social Health Authority system, demonstrating a governance architecture in which the same legal framework disciplines both commercial actors and state-operated AI systems.

The Kenya Cloud Policy 2025, effective May 2025, completes the picture by establishing the most granular LMIC data classification system found in the global evidence base, with four sensitivity tiers determining hosting requirements. Top secret and secret data must remain within government or dedicated private cloud infrastructure inside Kenya; restricted data may use public cloud subject to domestic residency. This represents a shift from reactive enforcement toward anticipatory regulatory architecture.

For LMICs, a hybrid approach is particularly relevant. Many have adopted residency mandates for government or public-sector data as part of their data protection frameworks. For instance, Indonesia and Nigeria require that all government or public body data (which would include social protection registries) be hosted locally, while permitting international data processing for less sensitive functions. Kenya's 2025 Cloud Policy formalises a similar tiered classification, aligning hosting decisions with data sensitivity levels. Such frameworks offer an incremental path to sovereignty, enabling governments to protect critical data domains while gradually developing domestic capacity for full localisation.

BOX 6. EUROPE UNION – SHAPING GLOBAL NORMS IN DATA PROTECTION

European legislation continues to shape global norms through what scholars have termed the 'Brussels Effect' (Bradford, 2020). For example, the General Data Protection Regulation (GDPR) established a global benchmark for cross-border data governance, requiring that personal data exported from the EU remain subject to equivalent protection. This has encouraged other jurisdictions to adopt adequacy-based or hybrid models to maintain data trade with the EU.

More recently, the Artificial Intelligence Act (AI Act) and the proposed Cloud and AI Development Act (CADA) extend this influence by linking data sovereignty directly to AI system governance and infrastructure capacity. Together, they illustrate a regulatory model that couples privacy, accountability, and infrastructure sovereignty – principles increasingly echoed in LMIC digital strategies. For countries developing AI-enabled social protection systems, the EU framework demonstrates how localisation, contractual safeguards, and sovereign cloud investment can coexist within a single, coherent governance architecture.

5.3 Implications for AI in social protection

The legal environment in which AI systems operate vary according to each regime's position on the localisation/transfer continuum, and this has implications for social protection agencies. In Tier 1 and 2 countries, AI models processing personal or sensitive social protection data must be hosted domestically, typically on government, sovereign, or hybrid clouds. For Tier 3 countries, cross-border hosting may be permissible if legal safeguards (such as SCCs or data transfer impact assessments) are in place.

Sovereignty depends not only on law, but on enforceability. Even in regimes with robust legislation, weak institutional capacity or lack of technical control can leave governments unable to implement their sovereignty mandates. Empowerment and capacity building can close this enforcement gap by strengthening the institutional mechanisms that give legal sovereignty practical effect (see Chapter 8).

Ultimately, data sovereignty regimes set the parameters for what is technically and legally possible in AI deployment. They determine not only *where* data can reside, but also *how* governments can exercise control, enforce accountability, and protect citizens' rights.

REGULATORY CONTEXT STRICTNESS TIERS

Tier 1. Strict localisation

Tier 2. Sector specific

Tier 3. Conditional

Tier 4. Unrestricted

BOX 7. TOGO'S NOVISSI – INNOVATION BEFORE FORMAL AI REGULATION

In April 2020, Togo's government launched Novissi at exceptional speed to deliver fully digital emergency cash transfers during the COVID-19 crisis. Subsequent phases incorporated ML methods using mobile phone call detail records (CDRs) and satellite data to estimate relative consumption or wealth proxies across roughly 5.7 million mobile subscribers, representing a large share of the population. These models were used to identify approximately 57,000 of the poorest rural individuals for targeted support.

Independent evaluations reported sizeable gains in targeting performance compared with simple geographic approaches, with studies indicating roughly a 40% improvement in precision. By early to mid-2021, Novissi had reached hundreds of thousands of beneficiaries, with published estimates ranging from about 800,000 to nearly one million depending on timing and source.

Togo did have a personal data protection law prior to Novissi's deployment (adopted in 2019), although enforcement institutions were still maturing. The country did not (and still does not) have AI-specific regulations. The programme, therefore, reflects innovation ahead of dedicated AI governance frameworks, rather than deployment in a complete legal vacuum.

Novissi underscores a broader governance lesson: when formal regulation is limited, system outcomes depend heavily on data quality, evaluation, institutional capacity, and implementation choices.

CHAPTER 6. DATA SOVEREIGNTY READINESS

Data sovereignty readiness is not a single condition, but the product of two distinct capacities: the infrastructure that a country has built, and the governance it can exercise over the infrastructure that it depends on. This chapter looks at data sovereignty readiness. Readiness, as used here, is about capability: the technical and institutional capacity to translate legal permissions into operational realities. A country with a strict localisation mandate and no domestic compute infrastructure has effectively declared sovereign control, but cannot enforce. A country with strong infrastructure and weak governance is in the opposite position.

6.1 Infrastructure readiness

Infrastructure readiness predicts sovereignty capacity more reliably than income level. Countries can be categorised as advanced, emerging, or nascent based on compute availability, cloud maturity, government data-centre capacity, and digital identity ecosystems (see Section 3.2 and Chapter 4). Several LMICs (India, Indonesia, Brazil, South Africa) demonstrate advanced sovereignty capacity through sovereign or hybrid cloud deployments and GPU clusters. Others with comparable income levels (Bangladesh, Pakistan, Kenya, Rwanda) remain in the emerging or nascent level due to gaps in sovereign compute or reliance on external hosting (World Bank, 2024; DCI, 2025). The pattern confirms that infrastructure readiness reflects deliberate investment choices, not GDP.

A national data centre alone does not confer sovereignty if critical layers of the technology stack remain under external control. The stack underpinning modern AI deployment is deep: hardware, orchestration tools, model training frameworks, and the GPU clusters on which heavy workloads run. When any of these layers are controlled by a foreign entity (a hyperscaler providing compute-as-a-service, a hardware vendor retaining remote administration rights, or a proprietary API provider) sovereignty is partial at best. For social protection specifically, dependency on foreign GPU capacity and vendor-managed model training creates structural risk, even when beneficiary data is nominally stored within national territory. The three largest hyperscalers collectively hold approximately 63% of the global cloud market (as noted in Chapter 4), and Oxford Insights' Government AI Readiness Index consistently finds LMIC governments scoring lowest on the infrastructure sub-indices that determine sovereignty capacity (Oxford Insights, 2024). These are structural features of the global AI economy, not temporary gaps.

Infrastructure readiness is a necessary condition for operational sovereignty, not a sufficient one. A nascent country cannot achieve sovereignty over high-sensitivity AI systems, regardless of how strong its governance frameworks are, because it lacks the compute capacity to host

them domestically. An advanced country with poor governance can squander its infrastructure advantage through weak procurement or contracts that hand effective control to vendors. Infrastructure and governance investments must, therefore, advance together, and the combination of strong infrastructure with weak governance (common in ambitious LMIC digitalisation programmes) represents a particularly acute risk (explored in Chapter 7).

6.2 Governance readiness

Governance readiness is the institutional and legal capacity to exercise the sovereignty that infrastructure makes possible – and it is this dimension that most consistently falls short in LMIC contexts. Where infrastructure readiness asks whether a country can host and control its AI systems, governance readiness asks whether institutions exist to enforce that control: to classify data, authorise processing, audit vendors, challenge non-compliant contracts, and provide citizens with meaningful redress. Governance readiness has four interconnected components, each of which must function for sovereignty to move from declarative to operational.

6.2.1 Data governance frameworks operationalised

The central distinction is not between countries that have frameworks and countries that do not, but between frameworks that operate on paper and those that operate in practice. What is absent in most LMIC contexts is the combination of: an empowered institution with a clear mandate; a documented data classification system; systematic mapping of data flows within and across government; and a functioning compliance machinery. The World Bank's *World Development Report 2021* (World Bank, 2021) found that no low-income country surveyed maintained a dedicated data governance entity with adequate technical capacity and independent enforcement powers. For AI in social protection specifically, the governance gap is best measured at the point of citizen interaction – whether the agency knows where beneficiary data goes, who can access it, and under what conditions – not at the level of strategy documents. Between 80% and 97% of AI cases in LICs and LMICs provide no public hosting information (DCI AI Hub, 2026a), and almost none demonstrate documented compliance with national data-governance standards (World Bank, 2021).

This is the operational expression of the public trust obligation. Agencies cannot protect data held on behalf of citizens if they do not know where it is, who can access it, or under what conditions it is processed. The governance gap documented above is a breach of the custodial duty that social protection agencies carry on behalf of the people whose data they hold.

6.2.2 Data protection laws

The coverage of data protection law has expanded rapidly – roughly 137 countries now have some form of relevant legislation – but a law that cannot be enforced against a foreign AI vendor offers only declarative sovereignty. Most data protection authorities in African and lower-income Asian contexts operate with limited staff and budgets, which constrains their capacity (Tech Hive Advisory, 2024). A comparative review of Kenya and Nigeria (two of the more advanced DPA structures in Africa) found that in practice, enforcement remains inconsistent, hindered by limited capacity, fragmented regulation, and low public awareness (Juma & Faturoti, 2025). This confirms that the challenge is often not legal architecture, but institutional capacity: countries have lawful processing obligations and data subject rights on paper, but lack the investigative tools to audit AI systems or compel foreign vendors to remediate breaches. The EU's GDPR represents the most developed enforcement architecture globally, but even it has required years and significant resources to act against major technology firms. AI adoption is accelerating faster than enforcement capacity is growing.

6.2.3 Legal basis for algorithmic accountability, explainability, and redress

Algorithmic accountability is the thinnest area of governance readiness globally. In most LMIC contexts, citizens whose welfare eligibility is determined by an AI system have no formal legal basis to challenge that decision, demand an explanation, or seek redress. The Dutch SyRI case (in which the District Court of The Hague ruled a welfare fraud-detection platform unlawful in 2020) illustrates what accountability jurisprudence looks like when it exists: a legal hook in the GDPR framework, a court willing to give it substantive effect, and civil-society organisations with standing to litigate (Van Bekkum & Zuiderveen Borgesius, 2021). Few LMIC contexts have all three. The Telangana Samagra Vedika case in India (where ML-based eligibility verification produced widespread wrongful exclusions with no formal redress mechanism) illustrates the alternative: accountability through investigative journalism rather than law (Amnesty International, 2024). The EU AI Act and GDPR Article 22 represent the most advanced architecture globally for algorithmic accountability, but such frameworks are largely absent in LMIC contexts, and most countries have not yet defined what a high-risk AI system is, let alone established rights for those affected by one (UN Special Rapporteur on Extreme Poverty and Human Rights, 2019).

The accountability deficit documented is a breach of fiduciary duty. Social protection agencies hold beneficiary data in trust for the citizens who provided it. When an AI system makes a consequential decision about welfare eligibility and no mechanism exists to challenge or explain that decision, the agency has failed in its custodial role. This is why redress is not an optional feature of governance readiness: it is what distinguishes a government exercising sovereign stewardship from one merely exercising sovereign control.

6.2.4 Capacity to negotiate vendor contracts and ensure compliance

The most practically important dimension of governance readiness is the capacity to negotiate AI vendor contracts that actually protect sovereignty. A country can have excellent laws and a functioning DPA and still sign away most of its sovereignty protections through an uninformed procurement process. Standard hyperscaler contracts are typically non-negotiable on the provisions that matter most: jurisdiction, the scope of sub-processor chains, audit rights, data egress terms, and the effect of the vendor's home-country legal obligations including compelled disclosure under US CLOUD Act provisions. Few LMICs have the model contracts, in-house expertise, or negotiating leverage to deviate from these defaults, and sub-processor chains (through which social protection data may flow to multiple foreign jurisdictions) are a particular vulnerability (Ada Lovelace Institute et al., 2021). Chapter 8 enumerates the specific protections sovereignty-protective contracts must include, such as data localisation clauses, source-code escrow, algorithmic audit rights, exit and portability clauses, zero data retention requirements, and provisions neutralising vendor-home-jurisdiction override.

The negotiating gap documented here reflects a structural asymmetry that individual capacity-building cannot fully resolve. Most AI systems deployed in LMIC social protection originate from a small number of vendors incorporated in high-income jurisdictions – jurisdictions that typically maintain strong domestic data protection and AI governance regimes, but lack extraterritorial mechanisms extending those obligations to overseas government contracts. The regulatory burden of protection falls disproportionately on the countries with the least negotiating leverage, a pattern with direct analogues in extractive industries and financial services before extraterritorial accountability frameworks emerged (see Chapter 4, Box 4). Until vendor home-country accountability mechanisms develop (through mandatory due diligence obligations, export notification requirements, or home-country liability for practices that would be unlawful domestically), contractual protections and institutional capacity remain the primary instruments available to recipient governments, and their limitations should be understood as features of the current global governance environment, not failures of individual countries.

6.3 The Sovereignty Readiness Matrix

The Sovereignty Readiness Matrix (Figure 4) is a diagnostic tool whose purpose is to identify the right intervention pathway for each readiness profile, not to rank countries. Infrastructure readiness sets the ceiling on what is technically possible; governance readiness determines whether that ceiling is ever reached. Infrastructure is a necessary condition: without adequate compute, connectivity, and digital public infrastructure, even the strongest governance framework cannot achieve operational sovereignty. But infrastructure alone is not sufficient. Countries that have built domestic data centres can still find themselves exposed to foreign legal jurisdiction, proprietary vendor dependencies, or contractual terms that quietly waive the very protections the infrastructure was supposed to enable.

Four profiles are analytically most important. Where both infrastructure and governance are strong (Estonia and Singapore are the clearest examples), data sovereignty is operationally feasible across all quadrants – the priority is active vendor management and adaptive legal reform. Where infrastructure is strong but governance is weak, data may be technically onshore while remaining legally exposed: no institution has documented the data flows, no contract specifies what the vendor may do with derivative data, and no audit mechanism verifies compliance. Indonesia illustrates this tension, with advanced GPU infrastructure, but persistent gaps in enforcement capacity and cross-agency data governance – the intervention priority is governance, not more infrastructure. Where governance is strong, but infrastructure is nascent (Uzbekistan presents elements of this profile, with a Tier 1 localisation mandate, but limited domestic AI compute), the law demands sovereignty that current capability cannot deliver – the priority is sequenced infrastructure investment alongside interim compliance pathways such as approved regional hosting and federated arrangements. Where both conditions are weak, which characterises many LICs and fragile states, interim safeguards are the only realistic near-term option: contractual protections negotiated through multilateral frameworks, digital public goods that reduce vendor opacity and lock-in, and human-in-the-loop oversight arrangements that maintain accountability even when algorithmic accountability law is absent.

FIGURE 4. SOVEREIGNTY READINESS MATRIX

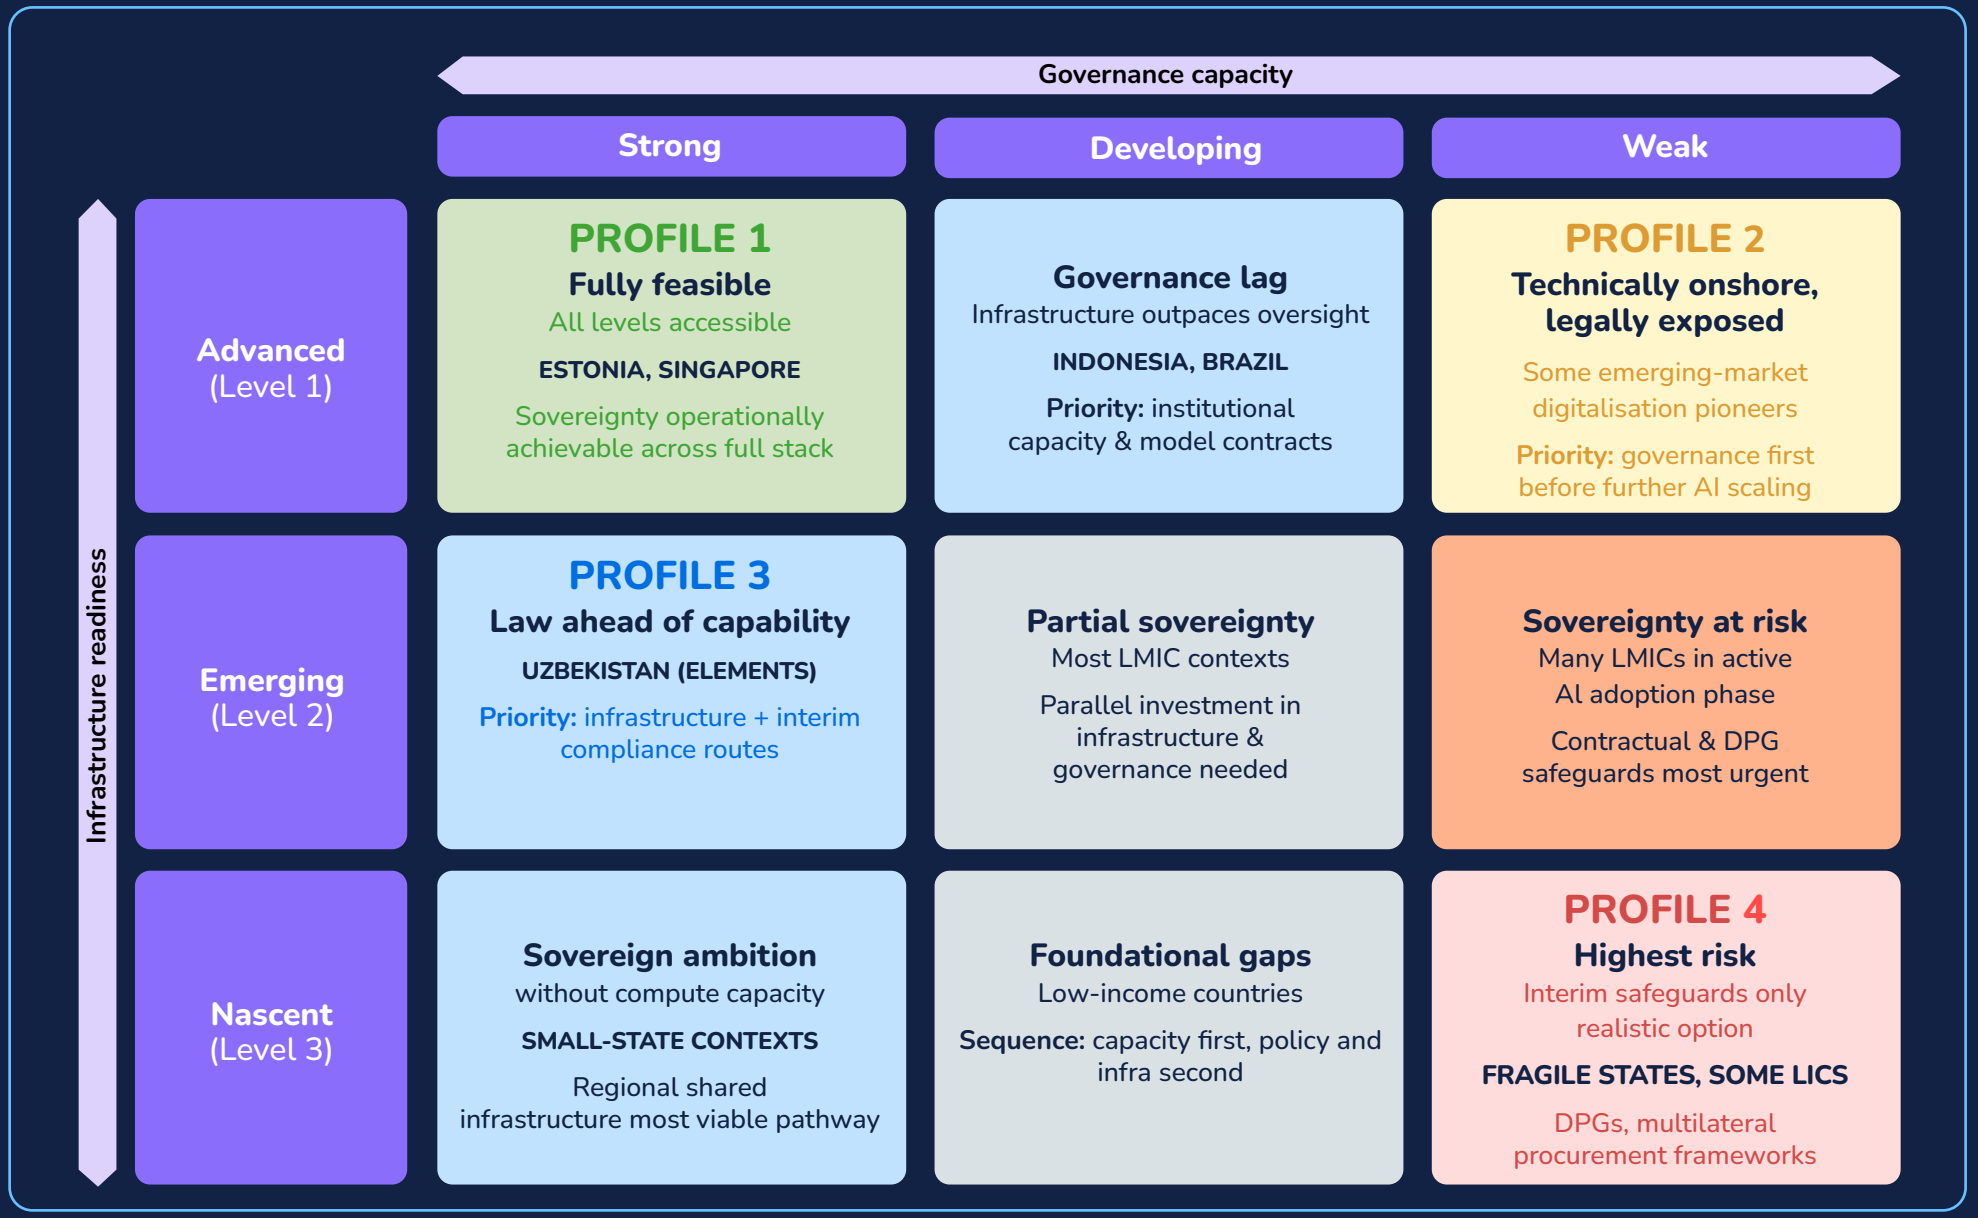

Key profiles: **PROFILE 1** – Fully feasible **PROFILE 2** – Onshore, legally exposed **PROFILE 3** – Law ahead of capability **PROFILE 4** – Highest risk: interim safeguards

6.4 Implications for risk assessment

A country's readiness profile is the analytical bridge between the regulatory context established in Chapter 5 and the dimensions of sovereignty risk examined in Chapter 7. It determines which risks are most acute and which mitigation pathways are realistically available. For countries with nascent infrastructure and weak governance, the acute risks are not regulatory non-compliance, but permanent dependency: vendor lock-in, foreign control of model weights derived from national data, and sub-processor chains that route beneficiary data through jurisdictions the government has no legal relationship with. For countries with advanced infrastructure and strong governance, the risks are different: regulatory arbitrage by sophisticated vendors and the gradual erosion of sovereignty through contract renewals that dilute original protections.

The four governance levers – policy, infrastructure, contracts, and capacity – are not equally available to all countries. The readiness profile determines which lever to pull first. A country with weak governance and nascent infrastructure needs capacity investment before policy reform, so that when legislation is enacted, institutions exist to implement it. A country with strong infrastructure and weak governance does not need more data centres, it needs a procurement framework and empowered oversight institutions. A country with strong governance, but limited infrastructure should prioritise contractual levers and the legal architecture for interim sovereignty measures, while sequencing infrastructure investment toward the highest-sensitivity use cases. Getting this sequencing right is a governance question, and the Sovereignty Readiness Matrix is designed to make the logic explicit. Chapter 7 examines the specific dimensions of risk that make the gap between readiness aspiration and readiness reality consequential for beneficiaries.

CHAPTER 7. DIMENSIONS OF DATA SOVEREIGNTY RISK

The next step is to understand when sovereignty is most at risk, which requires examining three interrelated dimensions: data sensitivity, computational intensity, and model training (location and control).

7.1 Data sensitivity

The first dimension of data sovereignty is data sensitivity. AI systems used in social protection rely on data that ranges from aggregate public indicators to highly personal information about personal income, household composition, or identity (DCI, 2025). The greater the sensitivity of the data, the higher the sovereign obligation to retain legal and operational control (World Bank, 2021; UNESCO, 2021).

7.2 Computational intensity

The second dimension is computational intensity. AI models vary widely in their computational demands. Some lightweight predictive tools can be run on ordinary government servers, while others – particularly deep-learning systems for image, voice or natural-language analysis – require high-performance computing typically provided by global cloud vendors. This infrastructure intensity determines whether hosting must occur domestically or can safely rely on cross-border arrangements under adequate safeguards.

7.3 Model training

The third dimension is model training, which is not often adequately considered in discussions on sovereignty. Training determines how the model learns from data, what patterns it internalises, and who owns the resulting model weights (NIST, 2023; OECD, 2019). When training occurs on foreign infrastructure or with mixed datasets, sovereignty risks multiply as the original data may leave the jurisdiction, the trained model may embed sensitive features derived from that data, and ownership or control over retraining may rest with an external entity.

FIGURE 5. KEY DIMENSIONS AFFECTING SOVEREIGNTY RISK EXPOSURE

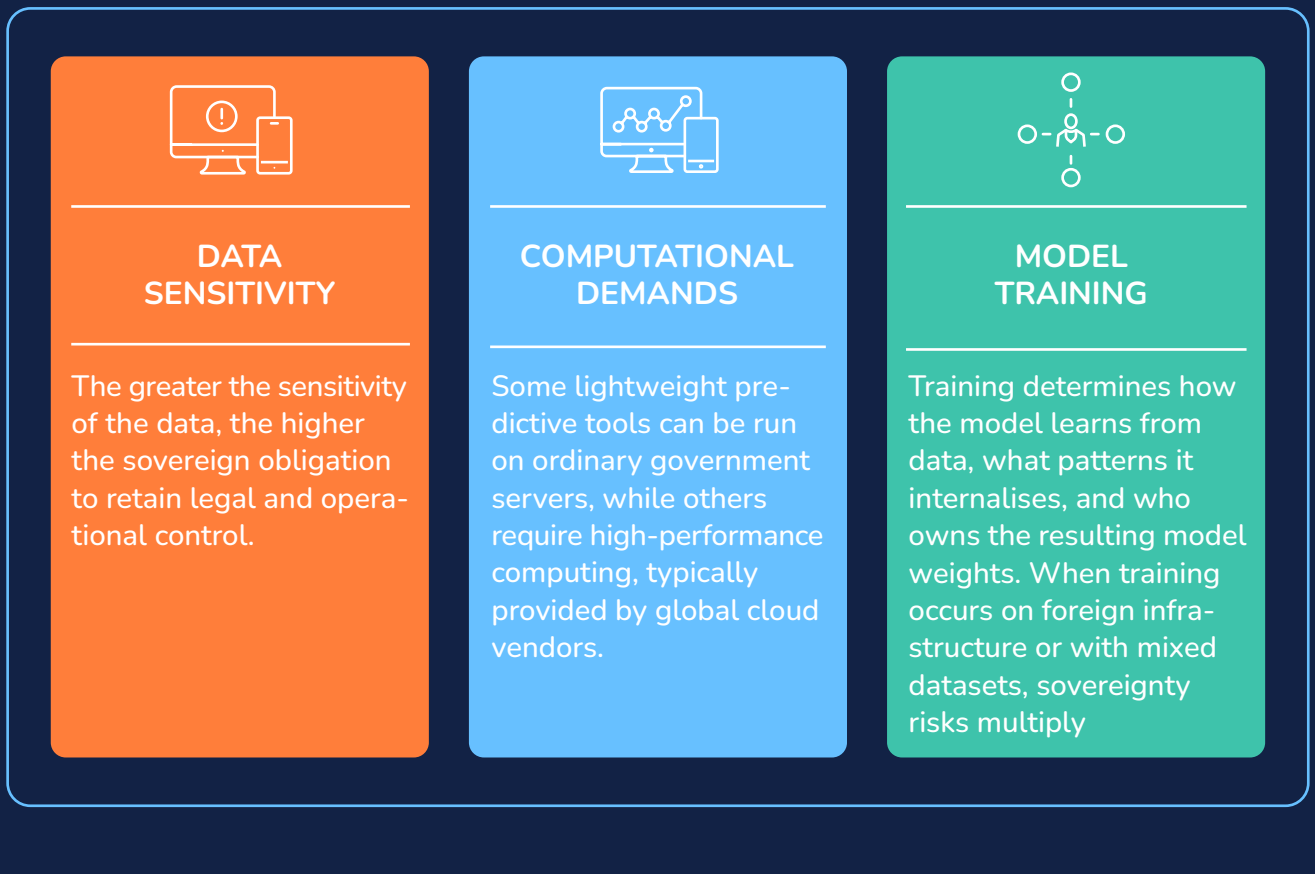

BOX 8. WHY MODEL TRAINING MATTERS FOR SOVEREIGNTY

Including training as a core dimension shifts the perspective from a static view of data storage to a dynamic understanding of how AI systems evolve and reproduce sovereignty risks over time. This is important for three reasons.

- **First, embedded knowledge is a sovereign asset.** Trained models effectively encode the data and values on which they are built. If model weights or architectures trained on domestic social-protection data are controlled abroad, this constitutes a loss of intellectual sovereignty: the knowledge derived from citizens' data becomes a foreign asset.
- **Second, it impacts lifecycle compliance and jurisdictional reach.** Data protection obligations extend beyond the dataset to any derivative artefact. When training occurs on infrastructure subject to foreign law, even anonymised data may fall under external jurisdiction, undermining national enforcement capacity.
- **Third, it creates dependency and capability gaps.** Externalised training creates technical dependency. Governments unable to retrain, refine or audit models domestically remain reliant on vendors, weakening both sovereignty and accountability. Integrating training into sovereignty assessments, thus, reinforces the capacity-building priorities.

At the same time, over-restricting training can impede innovation. Many LMICs lack sufficient compute or data diversity to train advanced models locally. A more nuanced approach would be to retain domestic control for sensitive or personally identifiable data, but allow regulated participation in regional or global training collaborations using anonymised or synthetic datasets. This follows the federated and hybrid governance models described earlier and aligns with privacy-enhancing techniques.

7.4 The Sovereignty/Infrastructure Quadrant

Together, these three dimensions define the sovereignty profile of an AI system. High-sensitivity data, intensive compute requirements, and externally-managed training environments create the highest-risk configuration. Conversely, anonymised data, modest runtime needs, and locally-governed training present minimal sovereignty exposure.

Plotting these dimensions together into the Sovereignty/Infrastructure Quadrant classifies AI applications according to the level of control governments must retain to ensure compliance and accountability, thereby assessing sovereignty risk (Table 3). The quadrant provides a practical tool for policymakers and implementers to assess the trade-offs between data sovereignty and technical feasibility. It links directly to the mitigation strategies presented in Chapter 8, which outline how to mitigate risks through infrastructure design, legal safeguards, and institutional capacity building.

- **Quadrant I: Sovereign AI Zone – is the 'no go' category for cross-border transfers.** These systems handle data that, if exposed or transferred abroad, would violate legal and ethical obligations. AI applications in this space, such as national biometric systems or fraud detection engines, should be fully sovereign, hosted on domestic or legally controlled infrastructure, and managed by national entities. The cost of local infrastructure is justified by the high stakes of data exposure and loss of control.

- **Quadrant II: Federated or Hybrid Governance Zone – provides opportunities for collaboration without compromising sovereignty.** Governments may retain personal data domestically while participating in regional or global federated learning initiatives, where model parameters, not raw data, are shared. This allows AI innovation to continue while preserving legal control. Federated or hybrid approaches are especially relevant for LMICs with limited

domestic computing power, but growing regional cooperation frameworks.

- **Quadrant III: Compute-Intensive Cloud Zone – includes applications where data sensitivity is low, but computational demand is high.** Here, sovereignty risks can be managed through contractual and technical safeguards, including encryption, Standard Contractual Clauses, and clear data processing agreements. Outsourcing such applications to regional or global clouds can improve efficiency and access to advanced AI capabilities, as long as oversight mechanisms are robust.
- **Quadrant IV: Shared Innovation Zone – encompasses low-sensitivity, low-infrastructure applications suitable for open experimentation.** This is an ideal entry point for social protection agencies beginning their AI journey, as pilots can be conducted with minimal risk while building institutional learning and digital maturity. Open-source or digital public goods approaches are particularly effective here, laying the groundwork for sovereign capacity.

TABLE 3. THE SOVEREIGNTY/INFRASTRUCTURE QUADRANT

DATA SENSITIVITY	INFRASTRUCTURE REQUIREMENT: HIGH	INFRASTRUCTURE REQUIREMENT: LOW
High	**Quadrant I – Sovereign AI Zone** • Critical systems process sensitive or identifiable social-protection data • Requires domestic or sovereign-cloud hosting, national encryption-key control, and full legal jurisdiction **Examples:** Biometric ID verification, fraud detection, dynamic eligibility scoring ***AI training position:*** *Training using real beneficiary data should occur locally or under sovereign agreement; foreign training only with equivalent protections.*	**Quadrant II – Federated/Hybrid Governance Zone** • Sensitive data, but moderate compute needs; enables federated or regional collaboration **Examples:** Case-worker decision support, social-risk prediction, proactive eligibility ***AI training position:*** *Federated learning and privacy-preserving computation allow shared model training without moving data across borders (NIST, 2023).*
Low/ moderate	**Quadrant III – Compute-Intensive Cloud Zone** • Low-sensitivity data, but high runtime demands; cross-border cloud use acceptable with legal and technical safeguards (SCCs, encryption) **Examples:** Conversational AI, translation systems, large-scale job-matching ***AI training position:*** *Model training on synthetic or public datasets can occur in trusted global or regional clouds if IP ownership remains sovereign.*	**Quadrant IV – Shared Innovation Zone** • Low data sensitivity and low compute needs; minimal sovereignty risk **Examples:** Document automation, policy simulation, operational forecasting, climate and disaster analytics ***AI training position:*** *Open innovation and digital public goods training encouraged to build capacity and regional knowledge.*

This Sovereignty/Infrastructure Quadrant not only simplifies policy decision making, but also aligns with the risk-based perspective adopted by many global data protection frameworks, such as the EU's GDPR and AI Act. It emphasises proportionality: the stricter the sovereignty requirements and the greater the sensitivity of data, the stronger the need for national control.

7.5 Policy implications

The Sovereignty/Infrastructure Quadrant translates legal principles into operational choices that social protection agencies can apply when planning AI deployments.
The following implications emerge:

- **The regulatory context determines boundaries, but design determines control.** A country's legal tier (Chapter 5) sets the outer limits of permissible data transfers, but sovereignty is ultimately operationalised through design choices – how data are stored, encrypted, and processed in each quadrant.

- **AI training and deployment require distinct governance strategies.** While training may occur under different infrastructure conditions than deployment, it must be subject to clear ownership and jurisdictional rules to prevent leakage of sovereign data or model control.

- **Capacity underpins sovereignty.** The Quadrant highlights that maintaining control – especially in Quadrants I and II – depends on the domestic capacity to manage infrastructure, negotiate vendor contracts, and enforce data protection law. This links directly to the empowerment of citizens, as well as staff of institutions, to have greater agency in engaging with AI.

- **Sovereignty should be embedded in every stage of the lifecycle.** Jurisdictional control must extend from data collection to model training, deployment, and updating. Contracts and governance frameworks should specify not only where data reside, but where models are trained and who owns the outputs.

- **Federated and privacy-preserving training offers a potential model.** Federated learning, synthetic data generation, and differential privacy can enable collaboration without data export – a pragmatic path for LMICs building capacity while maintaining control.

- **Infrastructure planning must anticipate training needs.** Investment in sovereign or regional high-performance compute clusters should consider future AI training demands, not just operational deployment. Partnerships under regional cloud or data-embassy frameworks (see Box 10) can mitigate cost barriers.

- **Governance capacity determines enforcement.** Effective oversight of training arrangements requires skilled regulators and auditors capable of verifying data provenance, monitoring compliance, and assessing model.

By analysing AI use cases through the combined lenses of *data sensitivity*, *computational intensity*, and *training control*, governments can make informed, risk-based decisions on how and where AI should be built and deployed in social protection systems.
This lifecycle approach links the regulatory foundations established in Chapter 5 with the practical mitigation strategies detailed in Chapter 8, ensuring that data sovereignty remains an operational principle, not a static rule, throughout the evolution of AI in social protection.

CHAPTER 8. AN INTEGRATED FRAMEWORK FOR MITIGATING DATA SOVEREIGNTY RISK

This chapter proposes an integrated framework for mitigating data sovereignty risks through coordinated action across four interdependent governance levers – policy, infrastructure, contracts, and capacity. These levers work together to enable governments to retain legal, operational, and institutional control over data and AI systems. The balance among these levers varies by country and capacity level, but the objective remains constant – to ensure that data and AI models serving public functions remain subject to the jurisdiction, values, and oversight of the state and its citizens.

8.1 Policy instruments: Laws, regulations, and institutional authority

Strong legal frameworks are the cornerstone of data sovereignty. The typologies in Chapter 5 show how different countries define the perimeter of sovereignty, ranging from strict localisation mandates to conditional transfer models. For AI in social protection, these frameworks must evolve from abstract privacy principles to operationally enforceable requirements that address AI-specific risks.

First, data protection and localisation laws and policies should explicitly extend to AI training and model management, ensuring that personal or sensitive data used to train models remains under national jurisdiction or within a trusted regional framework. Clauses can require that model weights derived from sovereign datasets be treated as *regulated data assets*, subject to the same rules as the underlying data.

Second, governments should empower regulators, such as data protection authorities or AI oversight councils,

with mandates to review cross-border AI contracts and data-transfer impact assessments. This institutional authority ensures that any movement of data or models is legally sanctioned and auditable.

Third, policy instruments should embed transparency and provenance obligations for AI models. Emerging provisions, e.g. the EU AI Act, point toward a shift from static compliance to real-time accountability. Sovereignty should depend on the ability to prove – through digital evidence – that data and models are stored, processed, and updated within authorised jurisdictions.

Finally, sectoral policy frameworks for social protection should recognise that citizens' data collected for social protection are held *in trust* for the people – not owned or controlled at will. Public agencies, therefore, have a duty to protect, use, and manage that data only in ways that serve the beneficiaries' interests and uphold their rights.

8.2 Infrastructure controls: Hosting, encryption, and technical assurance

Policy without control over infrastructure remains symbolic. Physical and logical hosting arrangements determine whether or not sovereignty can be enforced in practice. The Sovereignty/Infrastructure Quadrant in Chapter 7 illustrates that AI use cases with high data sensitivity or intensive runtime requirements (Quadrants I and II) demand stronger infrastructure sovereignty than low-risk cases (Quadrants III and IV).

Sovereign and hybrid cloud models are, therefore, critical. As discussed in Chapter 4, African, Asian, and Latin American governments are investing in national or regionally managed AI clouds – such as Brazil's Nuvem de Governo, Singapore's RE:AI, and Africa's Smart Africa Cloud – to host sensitive workloads under domestic control (World Bank & Partners, n.d.). For resource-constrained contexts, hybrid models – in which critical data and workloads remain on domestic servers while less sensitive compute is outsourced – provide a practical compromise.

BOX 9. INDONESIA – SOVEREIGN AND HYBRID CLOUD: FROM REGULATORY MANDATE TO AI INFRASTRUCTURE

Indonesia's GPU Merdeka initiative, launched in mid-2024 by Lintasarta (an Indosat subsidiary) in partnership with NVIDIA, represents one of the most advanced sovereign AI compute infrastructure in the LMIC category. The facility deploys clusters of NVIDIA H100 SXM GPUs purpose-built for AI model training and inference, rather than general cloud hosting. This provides the foundation for Indonesia to train frontier-scale AI models under domestic legal control, rather than merely storing data onshore while computational sovereignty leaks offshore.

The regulatory driver is Indonesia's Government Regulation 71/2019, which mandates domestic storage and processing for public electronic systems, creating the compliance imperative that induced AWS to open a Jakarta data centre in 2022. Regulation preceded and created the commercial market for sovereignty-compatible infrastructure, confirming that policy instruments and infrastructure investment are mutually reinforcing rather than sequential.

Technical assurance complements physical control. Encryption, key management, and attestation mechanisms ensure that even if data are processed on external infrastructure, governments retain exclusive control of access. *Admin-of-one-jurisdiction* policies (restricting system administrators to national legal boundaries), *log residency* (keeping all telemetry within country), and *hardware attestation* (cryptographic proof that workloads run on certified nodes) are emerging as the technical foundations of operational sovereignty.

Future-ready sovereign architectures can also embed compliance-as-code, enabling automatic enforcement of data-residency rules in deployment pipelines. Under this model, no workload containing 'sovereign-only' data can run on unapproved infrastructure – a design that anticipates real-time regulatory audit.

For high-sensitivity AI training, countries should explore data embassy models that maintain legal sovereignty even when physical infrastructure is shared. Estonia's data-embassy approach demonstrates how encrypted, remotely-hosted environments can preserve national jurisdiction, an option increasingly viable for LMICs through regional or multilateral partnerships (Robinson et al., 2019).

BOX 10. DATA EMBASSY MODELS: AN EMERGING PATHWAY FOR SHARED AND SOVEREIGN AI INFRASTRUCTURE

Data embassies offer a novel approach to preserving digital sovereignty by hosting critical government data and services in secure, extraterritorial locations. First pioneered by Estonia, the data-embassy model stores replicated government systems in foreign jurisdictions under special bilateral agreements that grant the data the same legal protections as domestic infrastructure. This model demonstrates how trusted partners, treaty-based safeguards, and controlled hosting environments can provide states with a high degree of operational sovereignty even when data is physically located abroad.

While data embassies have not yet been extended to sovereign AI compute, they offer a viable template for shared, regionally governed AI cloud infrastructure. The underlying principles of data embassies – treaty-based jurisdictional control, secure hosting with diplomatic-level protections, pooled investment, and jointly-governed standards – could be adapted for sovereign AI needs. A 'compute embassy' model would allow multiple countries, particularly LMICs, to co-invest in trusted, geopolitically neutral AI-compute facilities that operate under co-defined sovereignty rules. Such shared infrastructure could provide access to an AI cloud that would otherwise be financially or technically unattainable. As AI accelerates demand for secure, sovereign compute, data-embassy principles provide one of the most promising, and currently under-explored, strategic pathways for collective capacity building (Robinson et al., 2019).

8.3 Contractual mechanisms: Standard clauses and vendor accountability

Even with sovereign infrastructure, most governments will continue to depend on external providers for cloud, AI, or analytical services. In these cases, contracts become the operational expression of sovereignty. Standardised legal instruments – such as Standard Contractual Clauses (SCCs), data processing agreements (DPAs), and service level agreements

(SLAs) – are the key tools for retaining control and ensuring redress (see Annex 3 for examples). Effective contracts should include the following minimum guarantees:

- **Data residency and processing location:** Explicit clauses should specifying that all data and backups remain within the jurisdiction or a designated sovereign region.
- **No training or reuse of government data:** Vendors must be prohibited from using social-protection data to train proprietary or third-party models.
- **Zero data retention:** Providers must process data in-memory only and delete all copies immediately after task completion.
- **Audit and inspection rights:** Governments should retain the right to conduct periodic audits or request real-time evidence of compliance, including hardware attestation logs.
- **Sub-processor restrictions and jurisdictional immunity:** Sub-contractors may be used only with written consent, and foreign legal claims (such as under the US CLOUD Act) must be contractually neutralised through indemnification or data-segmentation measures.

In the context of AI training, contracts must also specify ownership and intellectual-property rights over model weights and derived datasets. This prevents 'sovereignty leakage', where governments lose control over models built from public data. For lower-risk applications (Quadrants III and IV), contractual mechanisms can allow controlled use of commercial APIs, provided that governments implement measures such as masking personally identifiable information (PII), anonymisation, and secure API gateways. These mechanisms create layered protection, ensuring that even when data are processed by third-party AI services, sensitive information never leaves the sovereign perimeter.

BOX 11. SRI LANKA – STANDARD CONTRACTUAL CLAUSES: SOVEREIGNTY THROUGH PROCUREMENT DESIGN

Sri Lanka's Unique Digital Identity (SL-UDI) programme illustrates how contractual sovereignty protections can be embedded at the point of procurement design, before a vendor relationship is established. Funded through a grant from the Indian government, SL-UDI explicitly rejected proprietary identity platforms on the ground of high costs, limited flexibility, and vendor lock-in risk, adopting instead the MOSIP open-source platform. This provided for source code access under a Mozilla Public License 2.0, removing the dependency risk that escrow arrangements attempt to address retrospectively.

Beyond the platform choice, Sri Lanka imposed a set of contractual conditions that operationalise core sovereignty protections. All citizen data must be stored exclusively in Sri Lanka, addressing the data residency gap. The Managed Services Provider must be of Sri Lankan legal origin, restricting sub-processor jurisdiction to the domestic legal environment and neutralising foreign compelled-access risk. A pre-deployment security audit by Sri Lanka CERT was required before any citizen data could be entered into the system, establishing an enforceable audit-rights mechanism at the outset rather than as a post-hoc inspection right.

Taken together, Sri Lanka's approach demonstrates that SCC-equivalent protections are achievable without a formal SCC framework, by structuring procurement requirements directly. The case is instructive for LMICs negotiating with both commercial vendors and development partners.

8.4 Capacity-building interventions: Institutions, skills, and collaborative ecosystems

Achieving data sovereignty in AI-enabled social protection systems depends as much on people and institutions as it does on laws and technologies. Sovereignty requires the active capacity to govern, enforce, and sustain control over data and the algorithms that process it. Empowering stakeholders, therefore, means enabling governments, institutions, and citizens to exercise that control in practice, by building the knowledge, authority, and tools necessary to manage AI systems responsibly, negotiate with vendors on equal terms, and ensure that sensitive social protection data remains under domestic jurisdiction and aligned with public values.

Institutional empowerment begins with strengthening the technical and organisational capabilities of social protection agencies. A government cannot meaningfully assert data sovereignty if its institutions lack the skills to manage, audit, or maintain the very systems that embody sovereign control. Building technical capacity for data governance (including data management, cloud administration, and AI oversight) is, thus, an essential precondition for sovereignty.

Social protection ministries and their digital units should be equipped to classify data by sensitivity, manage encryption keys, and operate or supervise sovereign or hybrid cloud environments. These capacities extend beyond IT functions – they require strategic understanding among policymakers of how AI architectures intersect with data localisation laws, cross-border transfer restrictions, and international contractual obligations. When agencies possess in-house expertise to oversee these dimensions, they can better enforce compliance, resist vendor lock-in, and adapt AI tools to local priorities.

Regulatory and governance capacity are equally vital. Data sovereignty frameworks depend on strong domestic regulators capable of enforcing localisation requirements and overseeing international data flows. In many LMICs, data protection authorities or ICT regulators have limited jurisdiction or technical capacity to supervise complex AI ecosystems. Strengthening these institutions – through training, legal mandates, and cooperation with regional bodies – ensures that sovereignty is not merely declared, but is operationally enforceable (DCI, 2025).

Governments must also develop the capacity to negotiate robust contractual safeguards with external vendors, including clauses that guarantee local jurisdiction, data residency, and government ownership of encryption keys. Establishing inter-ministerial AI governance councils that bring together experts in law, technology, and social policy can provide oversight over high-risk AI systems used in social protection, review vendor agreements, and ensure accountability to both national legislation and citizen rights.

At the same time, data sovereignty must be understood as extending beyond state control to include social legitimacy. Empowering communities and beneficiaries plays a critical role in sustaining the social license for sovereign data governance. When people understand how their data is used, where it is stored, and how it is protected, trust in digital social protection systems increases. Participatory data governance mechanisms – such as community consultation forums,[2] citizen representation in oversight bodies, and feedback mechanisms for grievance and redress – can help align technological choices with societal expectations.

Capacity building for data sovereignty also involves developing a skilled workforce capable of designing, auditing, and maintaining sovereign AI systems. This requires multi-level investments in education and training, from short courses on data protection and AI ethics for policymakers, to technical programmes on secure cloud management, anonymisation, and privacy-preserving computation for IT professionals.

2 See, for example, the UK Confidentiality Advisory Group an independent body that provides expert advice on the use of confidential patient information for research purposes (https://www.hra.nhs.uk/about-us/committees-and-services/confidentiality-advisory-group/).

Universities and public sector academies can play a key role in embedding these competencies in national digital training curricula. Partnerships with international and regional institutions, as well as South-South cooperation initiatives, can accelerate knowledge transfer and help countries learn from comparable experiences. Over time, such initiatives contribute to a domestic talent base that can reduce dependence on foreign expertise and sustain national AI infrastructure under local control.

Open-source technologies and interoperable standards offer another pathway to empowerment. Proprietary, opaque AI platforms often restrict the ability of governments to inspect, modify, or migrate systems, creating technical dependence, which undermines sovereignty (Rikap et al., 2024). In contrast, open-source AI components and digital public goods (DPGs) enable transparency, adaptability, and shared innovation. Governments can audit code, tailor algorithms to local data, and pool resources with neighbouring countries to develop regional AI solutions that comply with shared data sovereignty principles. By investing in open standards and interoperable data architectures, countries can ensure that sovereignty does not lead to isolation, but instead supports federated collaboration, in which national systems remain autonomous yet capable of secure data exchange for collective learning and benchmarking.

Sustaining data sovereignty requires a culture of continuous learning and collaboration. AI technologies evolve rapidly, and the competencies needed to govern them must evolve in parallel. Establishing communities of practice among social protection agencies, data protection authorities, and civil society organisations allows for the ongoing exchange of experiences, identification of emerging risks, and co-creation of governance tools. Development partners and donors can support this process not by substituting external expertise for local decision-making, but by investing in long-term institutional capacity and regional networks for sovereign digital governance. By doing so, they help build resilience against external dependencies and ensure that sovereignty principles endure beyond individual projects or administrations.

Ultimately, stakeholder empowerment and capacity building are the foundation of data sovereignty. Laws and infrastructures can establish the boundaries of control, but only empowered institutions and informed citizens can exercise that control effectively and sustainably. In the context of AI for social protection, sovereignty is realised not through isolation, but through agency – the ability of national actors to decide how data is governed, how technology is used, and how digital systems serve the public good. By investing in institutional competence, participatory governance, and open sovereign technologies, countries can ensure that AI becomes a tool for empowerment rather than dependency, reinforcing the autonomy, rights, and dignity of the people whom social protection systems are designed to serve.

8.5 Integrating the four levers into a sovereignty strategy

The four levers – policy, infrastructure, contracts, and capacity – must operate as an integrated system. Each reinforces the others. Policy defines the legal perimeter while infrastructure enforces that perimeter technically. Contracts operationalise it in external relationships and capacity sustains and adapts it over time (see Figure 6).

FIGURE 6. AN INTEGRATED FRAMEWORK TO REDUCE AI-RELATED SOVEREIGNTY RISKS

A weakness in any lever undermines the whole architecture. For example, a strong localisation law without technical attestation remains unenforceable; sovereign infrastructure without skilled administrators is insecure; and contractual safeguards without domestic audit capacity are merely symbolic.

The actions under each lever are elaborated on in Table 4 – the Sovereignty Implementation Matrix – with illustrative examples of each action applied in a country context that governments can follow to operationalise their sovereignty strategy.

This includes:

- Codifying sovereignty requirements in law
- Implementing compliant and attested infrastructure
- Binding all vendors through standardised sovereign contracts
- Building continuous institutional and technical capacity

Applied together, these measures allow governments to harness AI responsibly, ensuring that innovation strengthens rather than erodes sovereignty: strengthening legal and governance frameworks; adopting sovereignty-aligned hosting and infrastructure pathways; transforming procurement and contracting; and building institutional and technical capacity. In the social-protection domain – where the state acts as custodian of citizens' most sensitive data – such an integrated approach is not optional, but essential.

TABLE 4. SOVEREIGNTY IMPLEMENTATION MATRIX FOR AI-ENABLED SOCIAL PROTECTION

GOVERNANCE LEVER	PURPOSE	KEY ACTIONS TO REDUCE DATA SOVEREIGNTY RISK	ILLUSTRATIVE EXAMPLES/ APPLICATIONS
Lever 1. Policy instruments – laws, regulations & institutional authority	Define the legal perimeters of data and AI control.	• Update data-protection and localisation laws to explicitly cover AI training, model weights, and derivative data assets. • Mandate Data Transfer Impact Assessments and regulator approval for cross-border processing. • Empower data-protection or AI oversight bodies with jurisdiction to audit AI systems used in social protection. • Adopt beneficiary-data-as-public-trust clauses in social-protection laws to codify fiduciary responsibility.	• California Assembly Bill 2013 (Effective 1 January 2026) – requires developers to post documentation about data used to train generative AI systems • EU AI Act – requires high-risk AI systems to be developed using training, validation and testing data sets that meet quality criteria

GOVERNANCE LEVER	PURPOSE	KEY ACTIONS TO REDUCE DATA SOVEREIGNTY RISK	ILLUSTRATIVE EXAMPLES/ APPLICATIONS
Lever 2. Infrastructure controls – hosting, encryption & technical assurance	Ensure operational and technical enforcement of legal sovereignty.	• Deploy **sovereign or hybrid clouds** for high-sensitivity use cases (Quadrants I–II). • Implement **encryption-key ownership, hardware attestation**, and **log-residency** policies. • Adopt **compliance-as-code** so that deployment pipelines automatically enforce residency rules. • Use **data-embassy or regional GPU clusters** for sovereign model training under legal control.	• Brazil – Government Cloud *(Serpro/Dataprev)*. • Singapore – RE:AI sovereign AI cloud • Kenya – Cloud Policy 2025 introduces tiered data-sensitivity hosting • Estonia – Data-embassy model for off-site sovereign backups
Lever 3. Contractual mechanisms – Standard Contractual Clauses (SCCs), data processing agreements (DPAs) & vendor accountability	Translate sovereignty principles into enforceable obligations with third parties.	• Insert **SCCs** mandating local processing and deletion after use (see Annex 2 for examples). • Negotiate **DPAs** forbidding model training on government data. • Require **zero-retention, sub-processor approval**, and **audit rights**. • Define **model-ownership and IP clauses** for AI trained on sovereign datasets. • Apply **PII-masking and secure API gateways** when using commercial AI services.	• Azure OpenAI Enterprise DPA – zero data retention and local residency • Google Vertex AI – regional processing options • National procurement templates embedding SCCs and audit clauses for cloud vendors
Lever 4. Capacity-building interventions – institutions, skills & collaboration	Create enduring capability to implement and enforce sovereignty.	• Establish **sovereign-data or AI-oversight units** within ministries. • Develop **technical expertise** in cloud, encryption, and federated-learning governance. • Strengthen **legal negotiation skills** for digital contracts. • Foster **regional federated-learning consortia** to share models without sharing data. • Enhance **citizen literacy and grievance mechanisms** on data rights.	• Smart Africa Alliance – regional sovereign-cloud blueprint • Federated-learning pilots for eligibility and fraud detection across LMICs • National AI academies and digital-negotiator training programmes • Public communication campaigns on beneficiary data rights

CHAPTER 9. CONCLUSION

AI is transforming how social protection systems operate, enhancing efficiency, responsiveness, and inclusion. Yet, as this report demonstrates, these gains come with significant implications for data sovereignty and protection. This report has shown that the potential benefits of using AI for social protection carry implications for data sovereignty and protection. Across LMICs, the adoption of AI in social protection is outpacing the establishment of robust governance safeguards. Dependency on external cloud services, foreign vendors, and opaque proprietary systems means that governments risk losing effective control over how sensitive beneficiary data is stored, processed, or reused.

This report identifies four levers that can reduce systemic sovereignty risks: policy, infrastructure, contracts, and capacity. Together, they reveal that sovereignty is not achieved through legislation alone. It must be operationalised through deliberate investment in infrastructure, enforceable contractual frameworks, and empowered public institutions capable of managing and auditing AI systems across their full lifecycle.

At the same time, absolute isolationism is neither feasible nor desirable. The approach taken by countries such as Estonia, Singapore, and Brazil demonstrates that hybrid and federated models combining domestic hosting, regional cloud collaboration, and stringent contractual safeguards, can reconcile sovereignty with innovation. Similarly, federated learning and privacy-preserving computation show that collaboration does not have to entail loss of control.

Ultimately, data sovereignty in AI-enabled social protection is a dynamic capacity: the ability of a state to decide, implement, and enforce how its data and models are governed, while ensuring that these decisions reflect the rights and interests of its citizens. Sovereignty is, therefore, not an obstacle to digital transformation, but a precondition for trustworthy AI adoption and long-term digital independence.

This report identifies four priorities for implementing AI in social protection in a way that strengthens, rather than erodes, data sovereignty: strengthening legal and governance frameworks; adopting sovereignty-aligned hosting and infrastructure pathways; transforming procurement and contracting; and building institutional and technical capacity. For each priority, the Sovereignty Implementation Matrix in Chapter 8 (Table 4) sets out the key actions and country examples that illustrate what these look like in practice.

If implemented collectively, the measures outlined in this report will enable countries to shift from dependency to autonomy, transforming AI from a source of risk into a vehicle of sovereign digital empowerment. In doing so, social protection systems can harness the benefits of AI while ensuring that the rights, dignity, and data of their citizens remain firmly within their control.

REFERENCES

Ada Lovelace Institute, AI Now Institute and Open Government Partnership, *Algorithmic accountability for the public sector: Learning from the limitations of the private sector*, Ada Lovelace Institute, London, 2021

Amnesty International, *Use of entity resolution in India: Shining a light on how new forms of automation can deny people access to welfare*, Amnesty International, 30 April 2024, https://www.amnesty.org/en/latest/research/2024/04/entity-resolution-in-indias-welfare-digitalization/ (accessed 18 May 2025)

Bradford, A., *The Brussels effect: How the European Union rules the world*, Oxford University Press, 2020

DCI, *A data governance framework for digital social protection systems*, Digital Convergence Initiative, GIZ, Germany, 2025

DCI AI Hub for Social Protection, *AI adoption in social protection: A global evidence review (2020–2025)*, Deutsche Gesellschaft für Internationale Zusammenarbeit, Germany, 2026a

DCI AI Hub for Social Protection, *AI Hub risk assessment framework for social protection: A Practical Guide to Responsible AI Deployment*, Deutsche Gesellschaft für Internationale Zusammenarbeit, Germany, 2026b

Dobson, K., 'Welfare fraud 2.0? Using big data to surveil, stigmatize, and criminalize the poor', *Canadian Journal of Communication*, 44(3), 331–342, 2019

Eubanks, V., *Automating inequality: How high-tech tools profile, police, and punish the poor*, St Martin's Press, 2018

EU Public Buyers Community, *AI model contractual clauses for public procurement*, 2023–24, 2025

Henman, P., 'Improving public services using artificial intelligence: Possibilities, pitfalls, governance', *Asia Pacific Journal of Public Administration*, 42(4), 209–221, 2020

ISSA and United Nations University Operating Unit on Policy-Driven Electronic Governance (UNU-EGOV), *Artificial intelligence in social security organizations*, International Social Security Association (ISSA), Geneva, 2024

Juma, I. and Faturoti, B., 'Enforcing data privacy in Kenya and Nigeria: Towards an African approach to regulatory practice', *International Review of Law, Computers & Technology*, 1–26, 2025, https://doi.org/10.1080/13600869.2025.2506918

NIST, *AI risk management framework (AI RMF 1.0)*, National Institute of Standards and Technology (NIST), US Department of Commerce, Gaithersburg, MD, 2023

OECD, *Recommendation of the Council on Artificial Intelligence (OECD AI Principles)*, OECD Publishing, Paris, 2019

OECD, *The nature, evolution and potential implications of data localisation measures*, OECD Publishing, Paris, 2023

OECD, *AI and the future of social protection in OECD countries*, OECD Artificial Intelligence Papers, OECD Publishing, Paris, 2025a

OECD, *Harnessing artificial intelligence in social security: Use cases, governance and workforce readiness*, OECD Digital Government Studies, OECD Publishing, Paris, 2025b

Oxford Insights, *Government AI Readiness Index 2024*, Oxford Insights, Oxford, 2024, https://oxfordinsights.com/ai-readiness (accessed December 2025)

Rikap C., Durand, C., Paraná, E., Gerbaudo, P., and Marx, P., *Reclaiming digital sovereignty: A roadmap to build a digital stack for people and the planet*, University College London Institute for Innovation and Public Purpose's (UCL IIPP), 2024, https://www.ucl.ac.uk/bartlett/public-purpose/Reclaiming-Digital-Sovereignty

Robinson, N., Kask, L., and Krimmer, R., 'The Estonian Data Embassy and the applicability of the Vienna Convention: An exploratory analysis', In: *Proceedings of the 12th International Conference on Theory and Practice of Electronic*, 2019

Shein, E., 'The dangers of automating social programs', *Communications of the ACM*, 61(10), 17–19, 2018

Soulé, Folashadé, *Digital sovereignty in Africa: Moving beyond local data ownership*, Policy Brief No. 185, Centre for International Governance Innovation, Waterloo, ON, Canada, June 2024, https://www.cigionline.org/static/documents/PB_no.185_eRCHbMI.pdf

Synergy Research, *Cloud Market Share Trends - Big Three Together Hold 63% while Oracle and the Neoclouds Inch Higher*, Synergy Research, 19 November 2025. https://www.srgresearch.com/articles/cloud-market-share-trends-big-three-together-hold-63-while-oracle-and-the-neo-clouds-inch-higher

Tech Hive Advisory, *Roundup on data protection in Africa – 2024*, Tech Hive Advisory, 2024, https://cdn.prod.website-files.com/641a2c1dcea0041f-8d407596/67e51af4813be0d929d4b8ee_Roundup%20on%20Data%20Protection%20in%20Africa_%202024%20.pdf (accessed 18 May 2025)

UK Government, *Guidelines for AI procurement* (GOV.UK), 2020

UNESCO, *Recommendation on the ethics of artificial intelligence*, UNESCO, Paris, 2021

UN Special Rapporteur on Extreme Poverty and Human Rights, *Report of the Special Rapporteur on extreme poverty and human rights: Digital technology, social protection and human rights*, UN Doc. A/74/493, United Nations General Assembly, New York, 2019

Van Bekkum, M. and Zuiderveen Borgesius, F. J., 'Digital welfare fraud detection and the Dutch SyRI case', *European Journal of Social Security*, 23(4), 323–340, 2021

Wagner, B., Ferro, C. and Stein-Kaempfe, J., *Implementation guide – Good practices for ensuring data protection and privacy in social protection systems*, SPIAC-B Working Group on Digital Social Protection, GIZ, 2024

Wodecki, B., 'Microsoft Tells French Lawmakers It Can't Protect User Data from US Demands', SDxCentral, 21 July 2025, https://www.sdxcentral.com/news/microsoft-tells-french-lawmakers-it-cant-protect-user-data-from-us-demands/

World Bank, *World development report 2021*: Data for better lives, World Bank, Washington, DC, 2021

World Bank, *Scaling up social assistance where data is scarce: Opportunities and limits of novel data and AI*, World Bank, Washington, DC, 2024

World Bank and Partners, *Digital public goods and government cloud guidance for LMICs*, World Bank, Washington, DC, n.d.

World Economic Forum, *AI procurement in a box: AI government procurement guidelines*, 2020

Legislation, regulations and decisions

Brazil, *Lei Geral de Proteção de Dados Pessoais (LGPD) (Law No. 13.709/2018, as amended)*, Government of Brazil, Brasília, 2018

Brazil, *Bill No. 2,338/2023 – Legal framework for artificial intelligence*, Senate of Brazil (pending enactment), Brasília, 2024

China (People's Republic of China), *Cybersecurity Law* (amended), National People's Congress Standing Committee, Beijing, 2017

China (People's Republic of China), *Data Security Law (DSL)*, National People's Congress Standing Committee, Beijing, 2021

China (People's Republic of China), *Personal Information Protection Law* (PIPL), National People's Congress Standing Committee, Beijing, 2021

China (People's Republic of China), *Provisions on algorithmic recommendation for Internet information services*, Cyberspace Administration of China, Beijing, 2022

China (People's Republic of China), *Interim measures for the management of generative artificial intelligence services*, Cyberspace Administration of China, Beijing, 2023

Department of Information and Communications Technology – Philippines (DICT), *Department Circular No. 10: Philippine Government Cloud-First Policy*, DICT, Manila, 2020

Egypt, *Personal Data Protection Law (Law No. 151/2020)*, Arab Republic of Egypt, Cairo, 2020

European Commission, *Implementing Decision (EU) 2021/914: Standard Contractual Clauses for transfers of personal data to third countries*, European Commission, Brussels, 2021

European Commission, *Data Governance Act (Regulation (EU) 2022/868)*, European Commission, Brussels, 2022

European Commission, *Digital Services Act (Regulation (EU) 2022/2065) and Digital Markets Act (Regulation (EU) 2022/1925)*, European Commission, Brussels, 2022

European Union, *General Data Protection Regulation (Regulation (EU) 2016/679)*, Official Journal of the European Union, 2016

European Union, *Artificial Intelligence Act (Regulation)*, European Parliament and Council, Brussels, 2024

Government of India, *Digital Personal Data Protection Act, 2023 (Act No. 22 of 2023)*, Government of India, New Delhi, 2023

Government of Indonesia, *Government Regulation No. 71/2019 on Electronic Systems and Transactions (GR 71)*, Government of Indonesia, Jakarta, 2019

Government of Indonesia, *Government Regulation No. 28/2024 on Health Information Systems*, Government of Indonesia, Jakarta, 2024

Government of Kenya, *Data Protection Act No. 24 of 2019 and Data Protection (General) Regulations, 2021*, Government of Kenya, Nairobi, 2019

Government of Kenya (MICDE/ICT Authority), *Kenya Cloud Policy* (Government Gazette 2025), MICDE/ICTA, Nairobi, 2025

Government of Malawi, *Data Protection Act, 2024*, Government of Malawi, Lilongwe, 2024

Government of Nigeria (NITDA), *Nigeria Cloud Computing Policy & NITDA Guidelines*, NITDA, Abuja, 2019

Government of Rwanda, *Law No. 058/2021 of 13/10/2021 on the protection of personal data and privacy*, Government of Rwanda, Kigali, 2021

Government of Russia, *Federal Law No. 242-FZ amending Federal Law No. 152-FZ on Personal Data*, Government of the Russian Federation, Moscow, 2015

Government of South Africa, *Protection of Personal Information Act (POPIA) No. 4 of 2013*, Government of South Africa, Pretoria, 2013/2021

Government of Uzbekistan, *Law on Personal Data No. ZRU-547 (2019) with 2021 amendments on localization of citizens' data*, Oliy Majlis, Tashkent, 2019/2021

Government of Vietnam, *Personal Data Protection Decree (PDP Decree)*, Government of Vietnam, Hanoi, 2023

Saudi Arabia Communications and Information Technology Commission (CITC), *Global Artificial Intelligence Hub Law (Draft)*, CITC, Riyadh, 2025

Singapore, *Personal Data Protection Act (PDPA)*, Personal Data Protection Commission, Singapore, 2021

United States, *Health Insurance Portability and Accountability Act of 1996*, 1996

United States, *Executive Order 14117 (Feb. 28, 2024) and DOJ Final Rule (Jan. 8, 2025) establishing the Data Security Program for transactions involving bulk sensitive personal and government-related data*, Federal Register, Washington, DC, 2024–2025

ANNEX 1. GLOBAL TYPOLOGY OF DATA SOVEREIGNTY REGIMES

This annex provides a detailed overview of global regulatory approaches to data sovereignty, expanding on the summary introduced in Chapter 5. The comparative Country Matrix highlights the implications for social protection data where relevant. The analysis undertaken interprets the legislation from the perspective of a government social protection data system. Different interpretations may apply for other contexts. While every effort has been made to ensure the accuracy of the interpretation, we recommend consulting a legal specialist to verify before using.

TABLE A1-1. COUNTRY MATRIX: STRICTNESS VS LEGAL APPROACH

COUNTRY/ REGION	STRICTNESS TIER	LEGAL MODEL	KEY PROVISIONS AND NOTES	RELEVANCE TO SOCIAL PROTECTION
China	Tier 1	Local-by-default/ permit-based/sectoral residency	PIPL & DSL require domestic storage; government review for outbound transfers. Cybersecurity Law (CSL) – effective 1 June 2017 Data Security Law (DSL) – effective 1 September 2021 Personal Information Protection Law (PIPL) – effective 1 November 2021	Health data, medical health information, and genetic data are explicitly classified as 'sensitive personal information' under the PIPL, Article 28, requiring separate consent and stricter protective measures. Social protection data would likely fall under the 'important data' classification based on GB/T 43697-2024 criteria, as data affecting 'social stability' and 'public interests' is classified as important data. Core data definition specifically includes 'data related to important people's livelihood and major public interests'.

COUNTRY/ REGION	STRICTNESS TIER	LEGAL MODEL	KEY PROVISIONS AND NOTES	RELEVANCE TO SOCIAL PROTECTION
Russia	Tier 1 or 2	Local-by-default/ permit-based/ adequacy model	All personal data of Russian citizens must be collected and stored in-country. Federal Law No. 152-FZ 'On Personal Data' of 27 July 2006 (Personal Data Law) Federal Law No. 23-FZ 'On Amendments to the Federal Law On Personal Data and Certain Legislative Acts of the Russian Federation' (signed 28 February 2025) – amending Article 18(5) of 152-FZ Federal Law No. 266-FZ 'On Amendments to the Federal Law On Personal Data' (14 July 2022)	Applies to any welfare data linked to national ID systems.
Uzbekistan	Tier 1 or 2	Local-by-default/ permit-based/ sectoral residency	Citizen data cannot be stored abroad. Law of the Republic of Uzbekistan 'On Personal Data' No. ZRU-547 (2 July 2019)	Prohibits external hosting of beneficiary databases, but also mentions 'adequate protection' for transfers.
Indonesia	Tier 2 or 3	Sectoral residency/ adequacy	Public electronic system operators (ESOs) must store data domestically.	Social assistance and insurance systems must host data locally.
Nigeria	Tier 2 or 3	Sectoral residency/ permit-based	Government ministries departments and agencies must use local cloud providers.	Social registers treated as sovereign data under National Information Technology Development Agency (NITDA) policy.
Philippines	Tier 3	Sectoral residency/ adequacy	Social protection data is regarded as personal data; cross-border flows allowed, but require safeguards.	Processing of sensitive personal information (which includes health records, social security numbers) requires specific consent prior to processing. Higher security standards apply.
Kenya	Tier 2 or 3	Sectoral residency/ adequacy	Cloud Policy (2025) classifies public data by sensitivity; mandates domestic hosting for restricted datasets.	Social protection registries classified as 'restricted'.
European Union	Tier 3	Adequacy model	GDPR limits transfers to jurisdictions with equivalent protection.	Permits use of global cloud AI if SCCs or adequacy decisions apply.

COUNTRY/ REGION	STRICTNESS TIER	LEGAL MODEL	KEY PROVISIONS AND NOTES	RELEVANCE TO SOCIAL PROTECTION
India	Tier 2 or 3	Blacklist model	Transfers are permitted unless destination is blacklisted.	Allows international cloud use with safeguards; sectoral rules may tighten.
Vietnam	Tier 2 or 3	Permit-based/ sectoral residency	Cross-border transfers require data protection impact assessment (DPIA) and regulator notification.	Public sector data is subject to oversight by the Ministry of Public Security.
South Africa	Tier 2 or 3	Adequacy model/ sectoral residency	Transfers are allowed under equivalence or consent provisions.	Social grant systems may use regional cloud if protections match POPIA.
United States	Tier 4	Sectoral residency/ targeted national-security export controls	No federal localisation; transfers that fall within scope of EO 14117 are prohibited or restricted, regardless of foreign privacy law, if they would give countries of concern or covered persons access to covered bulk data.	The United States has no general cross-border data transfer regime or localisation requirement, but EO 14117 (2024) and the Department of Justice Data Security Program (2025) impose targeted national-security restrictions on transfers of bulk sensitive personal data and government-related data to designated 'countries of concern' and covered persons. Public sector localisation limited to defence/national security. While not required for social protection, stringent requirements for risk assessments, security safeguards, and third-party contracts (like business associate agreements) create de facto restrictions, as transfers are only permissible if equivalent protection can be guaranteed.
Singapore	Tier 2 or 3	Contractual/adequacy	Transfers are allowed with comparable protection under contract.	Enables flexible AI cloud deployment; relies on vendor compliance.

ANNEX 2. VENDOR CHECKLIST FOR AI PROCUREMENT IN SOCIAL PROTECTION

The following questions are provided as a checklist for use when procuring AI services from vendors. Each of the sections can be included as an annex in requests for proposal (RFPs) to guide vendors. By posing stringent procurement questions, LMIC governments can safeguard data sovereignty and require GDPR-grade protections when procuring AI systems for social protection.

This checklist synthesises requirements from the GDPR, the EU AI Act, as well as international AI procurement guidance (World Economic Forum, 2020; UK Government, 2020; EU Public Buyers Community, 2025).

Q1. Data residency and processing location

Procurement questions:

- Will all data and backups be stored and processed exclusively within the country (or an agreed sovereign region)?
- Does the vendor's architecture prevent any transfer of social protection data outside the jurisdiction without explicit government approval?

Q2. No reuse or training on government data

Procurement questions:

- Will the vendor commit not to use government-provided data for AI model training or product improvement beyond this contract?
- Does the vendor agree not to share or monetise social protection data in any form (including anonymised) for purposes outside the contracted solution?

Q3. Zero data retention (in-memory processing)

Procurement questions:

- Does the solution ensure no persistent storage of personal data during processing (i.e. only in-memory or ephemeral processing)?
- Will temporary data (cache, logs, temp files) be immediately deleted or wiped as soon as each session or operation is completed?

Q4. Audit and inspection rights

Procurement questions:

- Will the vendor grant the government the right to conduct audits (including periodic scheduled audits and surprise inspections) of all relevant systems and data handling processes?
- Does the vendor maintain detailed security and processing logs (including attestation logs for system integrity) that the government can review for compliance verification?

Q5. Sub-processor restrictions and jurisdictional immunity

Procurement questions:

- Will the vendor disclose all sub-processors (cloud providers, sub-contractors, etc.) involved in processing government data, and agree that no sub-processor will be used without prior written approval from the government?
- Are any vendor entities or sub-processors subject to foreign laws (e.g. the US CLOUD Act) that could compel data disclosure? If so, what contractual or technical measures will be in place to prevent unlawful foreign access or claims of jurisdiction over the data?

Q6. Encryption and key management

Procurement questions:

- Does the solution enforce encryption at rest and in transit for all sensitive data using industry-standard algorithms (e.g. AES-256, TLS 1.2+)?
- Will the government retain exclusive ownership of encryption keys (e.g. via a Bring Your Own Key model), so that the vendor cannot access plain-text data without authorisation?

Q7. Legal compliance and certifications

Procurement questions:

- Does the vendor comply with all applicable national data protection laws and regulations (e.g. data privacy acts, cybersecurity requirements) in handling social protection data?
- Can the vendor demonstrate GDPR-level compliance through certifications or audits (for example, ISO/IEC 27701 for privacy information management, ISO/IEC 27001 for security, or any GDPR-approved certification or code of conduct)?
- Will the vendor contractually commit to ongoing legal compliance, including adapting to new laws or regulations and maintaining relevant certifications throughout the contract duration?

Q8. Model explainability and transparency

Procurement questions:

- Will the vendor provide documentation and disclosures explaining the AI system's decision-making logic (e.g. describing algorithms, model features, and how outputs are generated)?
- Does the vendor agree to enable algorithmic transparency and audits, such as allowing the government or a third-party auditor to inspect the AI model's behaviour, verify the inputs/outputs, or review source materials (subject to IP protections) to ensure decisions can be explained?
- Are there tools or interfaces for explainability that can be used to interpret how the AI model/system arrived at a given decision or recommendation?

Q9. Citizen data rights

Procurement questions:

- How will the vendor assist in upholding data subject rights under GDPR-like principles or local law (such as the right of access to one's data, rectification of inaccuracies, erasure of data, restriction or objection to processing, and data portability)?
- Does the system have built-in capabilities to retrieve, correct, or delete an individual's data upon request by the government or the citizen?
- Will the vendor support the government in handling consent withdrawal if the AI system relies on consent (e.g. immediately ceasing processing of that person's data)?

Q10. Capacity building and technology transfer

Procurement questions:

- Does the proposal include a training and capacity-building plan for government personnel (e.g. training staff to use and administer the AI system, interpret its outputs, and manage it post-deployment)?
- Will the vendor provide knowledge transfer or technology handover at the end of the contract (such as documentation, source code escrow, model files, or continued technical support) to enable the government to sustain and possibly scale the solution independently?
- How will the solution avoid vendor lock-in? For example, are there provisions for data and model portability, open standards, or licensing that allow the government to migrate the system or integrate improvements locally?

ANNEX 3. SAMPLE CONTRACTUAL CLAUSES

This annex presents three reference instruments that governments and development partners can use as templates when contracting with external AI service providers. Together, they form a layered contractual architecture that operationalises the sovereignty principles set out in Chapter 8 of this report: a data processing agreement (DPA) that governs the controller-processor relationship, Standard Contractual Clauses (SCCs) that safeguard cross-border data transfers, and Model Contractual Clauses for AI Procurement (MCC-AI) that address the AI-specific risks (model ownership, bias testing, audit rights, and human oversight) that generic data protection agreements do not cover (EU Public Buyers Community, 2025).

All three instruments are publicly available, officially published by European Union institutions, and represent the current international standard for responsible AI procurement. They are suitable as benchmark references even for governments outside the EU: many LMICs have adopted GDPR-equivalent frameworks, and even those that have not will find these instruments a practical starting point for negotiation with international vendors.

The instruments should be read together and adapted to the specific regulatory context of the procuring government. They do not replace legal advice, and governments should ensure that any contract is reviewed in light of applicable national law.

No single instrument provides complete coverage. The three documents below address different, but complementary, dimensions of sovereign AI contracting. The table below sets out how the three instruments fit together.

TABLE A3-1. REFERENCE INSTRUMENTS FOR USE WHEN CONTRACTING WITH EXTERNAL AI SERVICE PROVIDERS

INSTRUMENT	DOCUMENT TYPE	PRIMARY FOCUS	WHEN TO USE
EU DPA (2021/915)	Standard Contractual Clauses	Data processing obligations between controller and processor	All AI vendor contracts where vendor processes government data
EU SCCs (2021/914)	Standard Contractual Clauses	Safeguards for cross-border data transfers to third countries	When vendor infrastructure is located outside the procuring country
MCC-AI (2025)	Model Contractual Clauses	AI-specific risks: model ownership, transparency, audit, human oversight	All AI system procurement, especially high-risk use cases in social protection

Instrument 1.
EU Standard DPA Clauses (Controller-Processor)

Commission Implementing Decision (EU) 2021/915 establishes the Standard Contractual Clauses for the relationship between a **data controller** (the government agency) and a **data processor** (the AI vendor). It is the foundational instrument for ensuring that a vendor's processing of social protection data is lawfully governed under a written contract with binding obligations.

The clauses implement the requirements of Article 28 of the GDPR, which defines the minimum content of any controller-processor agreement. In the social protection context, this decision governs the relationship in which, for example, a national social registry authority (controller) engages a vendor to provide an AI-based eligibility assessment system (processor).

TABLE A3-2. INSTRUMENT 1: EU STANDARD CONTRACTUAL CLAUSES FOR CONTROLLER-PROCESSOR RELATIONSHIPS (DPA)

Purpose	Governs the legal relationship between a government agency as data controller and an AI vendor as data processor. Establishes binding obligations on the vendor regarding data use, security, sub-processors, and compliance.
Source	EUR-Lex: Commission Implementing Decision (EU) 2021/915
Key provisions	• Vendor may only process data on documented instructions from the controller • All persons with access to personal data are bound by confidentiality • Vendor must implement appropriate technical and organisational security measures • Sub-processors may only be engaged with prior written authorisation from the controller • Vendor must assist the controller in fulfilling data subject rights (access, rectification, erasure) • Vendor must return or delete all personal data at the end of the contract • Controller retains the right to audit the vendor's compliance • Vendor must notify the controller without undue delay of any personal data breach

Note on adaptation: The clauses provide two options for sub-processor authorisation (specific versus general written authorisation). For high-sovereignty contexts such as Uzbekistan (Tier 1) or other countries with strict data residency requirements, Option 1 (specific prior written authorisation for each sub-processor) is recommended. Governments should also ensure that Annex II of the DPA – which specifies the technical and organisational security measures – is completed with explicit reference to data residency and encryption key management requirements.

Instrument 2. EU Standard Contractual Clauses for International Data Transfers

Commission Implementing Decision (EU) 2021/914 provides the Standard Contractual Clauses for the transfer of personal data to third countries, that is, countries outside the European Economic Area that do not have an adequacy decision. For governments in LMICs procuring AI services from international vendors (cloud providers, software companies, or AI platform providers headquartered in Europe, the United States, or elsewhere), these SCCs provide the internationally recognised benchmark for governing such transfers.

The 2021 SCCs replaced three earlier sets of clauses and introduced a modular structure covering four transfer scenarios. The most relevant for social protection AI procurement are Module 2 (Controller to Processor), applicable where a government transfers data to a vendor for processing, and Module 3 (Processor to Processor), applicable where a primary vendor engages sub-processors in different jurisdictions.

TABLE A3-3. INSTRUMENT 2: EU STANDARD CONTRACTUAL CLAUSES FOR INTERNATIONAL DATA TRANSFERS (SCCS)

Purpose	Provides legally binding safeguards for transfers of personal data from a government agency to an AI vendor located outside the country or outside the EEA. Establishes enforceable obligations that survive any change in applicable law.
Source	EUR-Lex: Commission Implementing Decision (EU) 2021/914
Key provisions	• Module 2 (Controller to Processor): covers transfers where the government is controller and vendor is processor • Module 3 (Processor to Processor): covers onward transfers to sub-processors • Data exporter must verify adequacy of protection in the destination country before transfer • Vendor must notify the government of any binding legal demands for disclosure (e.g. under CLOUD Act) • Vendor must challenge legally binding requests for access where possible • Government retains audit rights and right to suspend or terminate transfers if safeguards are breached • Vendor must maintain records of all processing activities carried out on behalf of the government

Note on adaptation: Governments outside the EU that are not legally required to use these SCCs will still find them a powerful negotiating benchmark with international vendors, many of whom are already familiar with these obligations through their EU business. The European Commission also publishes Q&A guidance on the SCCs, which clarifies how the clauses function in practice. Word versions of all four modules are available on the European Commission publications page.

Instrument 3.
EU Model Contractual Clauses for AI Procurement (MCC-AI)

The EU Model Contractual Clauses for AI Procurement (MCC-AI) were developed by the EU Public Buyers Community (a European Commission initiative) and updated in March 2025 to align with the EU AI Act. They represent the most comprehensive publicly available contractual template specifically designed for the procurement of AI systems by public sector organisations.

Critically, the MCC-AI address aspects of AI procurement that generic DPAs and SCCs do not cover: model risk management, dataset governance, algorithmic transparency, human oversight requirements, and post-deployment monitoring. They are available in two versions:

- MCC-AI High Risk – for AI systems with significant potential impact on individuals' rights or welfare, such as automated eligibility determination, fraud scoring, or predictive targeting of social assistance. This version is directly applicable to most high-value AI use cases in social protection.
- MCC-AI Light – for AI systems that do not qualify as high-risk, but may still affect citizens, such as administrative workflow automation, document processing, or back-office analytics.

TABLE A3-4. INSTRUMENT 3: EU MODEL CONTRACTUAL CLAUSES FOR AI PROCUREMENT (MCC-AI, 2025 EDITION) (AI-SPECIFIC MODEL CLAUSES)

Purpose	Provides AI-specific contractual obligations for public sector procurement of AI systems. Covers model risk management, data governance, transparency, human oversight, conformity assessment, and post-deployment monitoring. Designed to supplement (not replace) a DPA.
Source	EU Public Buyers Community: Updated MCC-AI (March 2025)
Key provisions	• 12-step risk management framework spanning the full AI system lifecycle, including pre-deployment testing • Data governance requirements: datasets must be relevant, accurate, and appropriate to the deployment context • Technical documentation and instructions for use must be provided by the vendor • Mandatory automatic logging of AI system operations, retained for appropriate periods • Conformity assessment requirements before deployment for high-risk systems • Cybersecurity measures against adversarial attacks, data poisoning, and model evasion • Human oversight mechanisms: operators must be able to intervene, override, or suspend the AI system • Data portability: datasets must be returned in a common file format at contract termination • Vendor must destroy all government datasets at contract end unless otherwise agreed • Vendor prohibited from using government data to train proprietary or third-party models

Note on using the MCC-AI: The clauses are designed to be annexed to an existing service agreement, not used as a standalone contract. They do not address intellectual property, payment, liability, or applicable law – these must be covered in the main agreement. A detailed commentary document is published alongside the clauses to explain how individual provisions should be interpreted and applied. All translations into EU languages are also available at the link above.

Guidance on applying these instruments in practice: Layering approach

For a typical AI vendor engagement in social protection, for example, procuring a cloud-hosted machine learning system for beneficiary targeting or case management, governments should deploy all three instruments in combination:

- Include the EU DPA clauses (2021/915) as the core controller-processor agreement, ensuring Annex II specifies data residency, encryption standards, and key management requirements.
- Attach the EU SCCs (2021/914), Module 2, for any cross-border transfers, with Clause 14 transfer impact assessment completed to document the adequacy of protection in the vendor's jurisdiction.
- Attach the MCC-AI High Risk clauses as a further schedule governing AI-specific obligations, selecting provisions proportionate to the risk profile of the specific system being procured.

Relationship to the vendor checklist (Annex 2)

The procurement checklist provided in Annex 2 of this report maps directly onto the contractual instruments described here. Governments should use the checklist to evaluate vendor responses during procurement, and then incorporate the relevant minimum guarantees (data residency, zero data retention, audit rights, sub-processor restrictions) as binding obligations in the DPA annexes and MCC-AI schedules. This ensures that pre-contract commitments are converted into enforceable contractual terms.

Jurisdiction and adaptation

Governments operating outside the EU are not legally required to use these instruments, but are strongly encouraged to treat them as a minimum benchmark for negotiation. Several considerations apply:

- For Tier 1 sovereignty regimes (such as Uzbekistan), data residency clauses in the DPA Annex II and sub-processor restriction clauses must be strengthened to explicitly prohibit processing outside national territory without explicit ministerial authorisation.
- For Tier 2 regimes, the SCCs (2021/914) provide the appropriate framework for cross-border transfers where adequacy has not been formally established.
- For lower-income countries with limited legal capacity, development partners should support governments in completing the mandatory annexes (particularly the technical and organisational measures in DPA Annex II and the risk management framework in MCC-AI Annex A) as part of procurement support missions.
- The International Association of Privacy Professionals (IAPP) provides a practical guide on applying the MCC-AI that is accessible to non-specialists: **https://iapp.org/news/a/eu-model-contractual-clauses-for-ai-procurement-a-practical-guide**.

Quick Reference: Document Links

DOCUMENT	TYPE	LINK
EU Standard DPA Clauses (Controller-Processor)	DPA / SCCs	EUR-Lex: 2021/915
EU SCCs for International Data Transfers	SCCs	EUR-Lex: 2021/914
EU SCCs – Word versions (all modules)	SCCs (editable)	EC Publications
EU SCCs – Q&A Guidance	Guidance	EC Q&A PDF
EU MCC-AI (updated 2025, High Risk + Light)	AI model clauses	Public Buyers Community
MCC-AI Non-High Risk Template (PDF)	AI model clauses	PDF Download
IAPP Practical Guide to MCC-AI	Commentary	IAPP

IMPRINT

Published by
Deutsche Gesellschaft für Internationale Zusammenarbeit (GIZ) GmbH

Registered offices
Bonn and Eschborn

Friedrich-Ebert-Allee 32 + 36
53113 Bonn, Germany
T +49 228 44 60-0
F +49 228 44 60-17 66
info@giz.de

Dag-Hammarskjöld-Weg 1–5
65760 Eschborn, Germany
T +49 6196 79-0
F +49 6196 79-11 15

Required citation
DCI AI Hub for Social Protection, *AI in Social Protection: Data Sovereignty Considerations*, Deutsche Gesellschaft für Internationale Zusammenarbeit, Germany, 2026.

Authors
Robert Worthington
Gilbert Kondani

Responsible for the content
Bernd Schramm, GIZ

Editing
Susan Sellars

Layout
EYES-OPEN, Berlin

Bonn, August 2026

AI Hub for Social Protection

The AI Hub is a global partnership, with close cooperation among nine core implementing partners: German Social Accident Insurance Institution for the Construction Sector (BG BAU), Government of Canada – Employment and Social Development Canada (ESDC), Dataprev (the Social Security Information and Technology Enterprise of the Brazilian government), Federation of German Pension Insurance Institutions (DRV Bund), Deutsche Gesellschaft für Internationale Zusammenarbeit (GIZ), International Labour Organization (ILO), International Social Security Association (ISSA), Organisation for Economic Co-operation and Development (OECD) and the World Bank.

The AI Hub is implemented by:

BG BAU Berufsgenossenschaft der Bauwirtschaft

Government of Canada / Gouvernement du Canada

DATAPREV

Deutsche Rentenversicherung

giz Deutsche Gesellschaft für Internationale Zusammenarbeit (GIZ) GmbH

issa

OECD

THE WORLD BANK

International Labour Organization