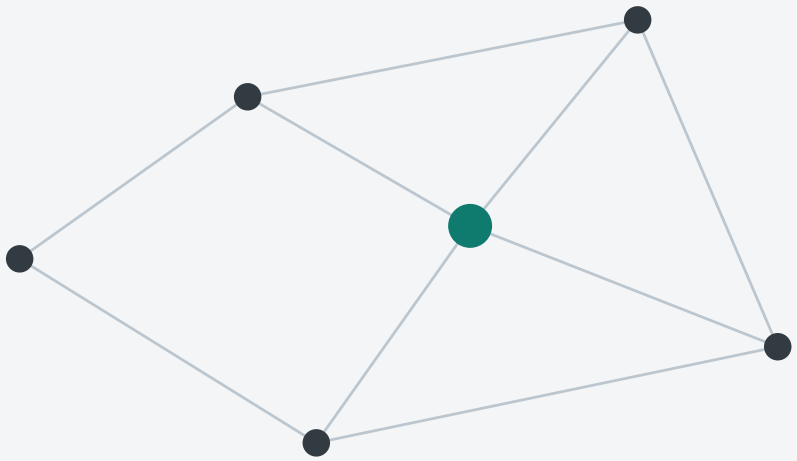
PIVRA

RESEARCH REPORT RP-2026-04

Governing Artificial Intelligence in Pakistan

Institutions, Policies and Regulation



Mamoona Muqadas

November 2026

Pakistan Institute for Vision Research and Action

PIVRA.ORG

Governing Artificial Intelligence in Pakistan

Institutions, Policies and Regulation

AUTHOR

Mamoona Muqadas

PUBLISHED BY

Pakistan Institute for Vision Research and Action

PUBLICATION

Research Report RP-2026-04 | November 2026

RECOMMENDED CITATION

Muqadas, M. (2026). *Governing Artificial Intelligence in Pakistan: Institutions, Policies and Regulation.* PIVRA Research Report RP-2026-04. Pakistan Institute for Vision Research and Action.

ABOUT THIS REPORT

This report maps how artificial intelligence is governed in Pakistan: the institutions that make decisions, the instruments they issue, the sectors that apply them and the evidence that any of it is working. It covers the period from the Digital Nation Pakistan Act 2025 to the first week of November 2026, with earlier material where it explains how the present arrangement came about.

The report is written for people who have to make decisions: ministers and secretaries, regulators, provincial governments, donors, investors, founders, civil society organisations and journalists. It assumes no technical background and does not treat AI as a single technology with a single set of risks.

ABOUT PIVRA

PIVRA is an independent, non-partisan research and policy institute based in Lahore and registered with the Securities and Exchange Commission of Pakistan. It publishes open-access research, analysis and data on Pakistan and provides commissioned research, advisory and training to organisations in Pakistan and abroad.

© 2026 Pakistan Institute for Vision Research and Action.
All rights reserved.

Lahore, Pakistan
www.pivra.org | info@pivra.org

Contents

Executive summary 5

The report in numbers 9

A note on method and evidence 10

01 The governance question 12

02 The state of AI in Pakistan 16

03 The legal baseline: what applies when no AI law does 22

04 The national architecture 26

05 The National AI Policy: ambition and arithmetic 32

06 Data governance, privacy and WASL 38

07 Standards, risk classification and assurance 43

08 Government deployment and procurement 47

09 Sectoral regulation: finance, telecoms, courts, universities, health and markets 52

10 Cybercrime, disinformation and online harms 58

11 Provinces and federalism 61

12 Rights, inclusion and labour 65

13 International alignment, compute and geopolitics 70

14 The implementation audit 74

15 Scenarios to 2030 79

16 Recommendations 83

17 Conclusion 88

Appendix A Timeline of AI governance in Pakistan 90

Appendix B Glossary 92

Appendix C Data notes and calculations 93

Appendix D What the government should publish 95

References 96

Figures

01 The 23 register entries, by category 9

02 The report’s five-question governance test 13

03 AI investment potential: selected scores 20

04 Pakistan’s national AI governance architecture 27

05 The AI allocation and indicative annual cost of the skills targets 35

06 From procurement to continuing accountability 48

07 Illustrative weights for the scenarios to 2030 81

08 A sequence for AI governance reform 84

Executive summary

Pakistan has spent under two years building the architecture of AI governance. It has not yet spent enough time making that architecture work.

Between January 2025 and November 2026 the country acquired a digital state law, a national AI policy, a formal declaration of AI principles, a draft national data governance policy, a national data exchange, sectoral guidance for judges and universities, new public procurement rules and a sovereign AI infrastructure programme. On paper this is one of the more complete governance stacks in South Asia. In practice the binding parts of it are thin. There is no AI statute. No independent AI regulator has been notified. No risk classification for AI systems is in force. No public register of government AI systems exists. No impact assessment has been published. No enforcement decision has been reported. The institutions that are meant to do this work are still being staffed.

That is the central finding of this report. Pakistan's problem is not that it lacks an AI policy. Its problem is that it has confused the production of policy with the construction of governance. The two are related, but they are not the same, and the gap between them is now the most important fact about AI governance in the country.

Twelve findings

1. The legal basis for AI regulation is enabling, not mandatory. The Digital Nation Pakistan Act 2025 gives the Pakistan Digital Authority power to regulate artificial intelligence only in areas assigned by the National Digital Commission. The Act creates a coordinator and a standard-setter, not an AI regulator with a statutory duty to protect the public.[1] Whether the Authority can act on its own initiative on AI remains a question of assignment, not of law.

2. The National AI Policy is a development strategy with a governance annex. Of the 26 actions in its implementation matrix, the large majority are about training, funding, infrastructure and promotion. Only three are explicitly regulatory: an AI directorate, a sandbox and a target for enterprises in it. Rights, redress and independent oversight do not appear as actions with owners and deadlines.[2]

3. Many of the policy's targets fall due in 2025 or 2026. Its action matrix dates the majority of deliverables to those two years. Among the 2026 commitments are public awareness, civil-servant training, AI research, patent registration and a national AI fund. The public record reviewed does not provide sufficient evidence to verify delivery at the stated scale. This is a progress assessment as of 3 November 2026: commitments dated 2025 have passed their stated year, while a target dated only 2026 is not treated as overdue before the end of that year.[2]

4. The money is small relative to the ambition. The 2026-27 public development budget proposes Rs 524.64 million for the National AI Advancement Initiative, against a policy target of training 200,000 people a year and funding 400 AI projects a year. The entire ongoing IT portfolio in the development budget is Rs 16.29 billion, most of it for buildings.[3] This is not a criticism of any single project. It is an arithmetic observation: the targets and the resources are on different scales.

5. Data governance is where AI governance will actually be decided. The draft National Data Governance Policy 2026 sets residency tiers, designates primary data registers, creates a National Data Governance Council and requires AI-specific controls for automated decision-making. The National Data Exchange, WASL, is meant to move records between institutions under consent and audit rules. Both are important. Neither is in force, and the personal data protection bill that would give citizens enforceable rights remains a draft.[4,5,6]

6. The AI Council exists on paper, not yet as a visible institution. The policy creates a council chaired by the federal minister for IT, with a policy implementation cell under the ministry. In early 2026 reporting described the process of forming the first council as beginning. By November 2026 no membership list, meeting record or annual report had been published.[2,7]

7. The Pakistan Digital Authority is the pivot, and its capacity is unproven. The Authority was constituted in August 2025 under a five-year chair and has advertised posts for AI risk classification, impact assessment and lifecycle assurance.[8,9] It has not published an annual report, a budget, a staffing plan or a regulatory agenda. Its standards portal describes a sophisticated publication architecture, but a portal is not a compliance record.

8. Sectoral regulation is uneven, and that unevenness is the real map. The State Bank has surveyed AI use among regulated financial institutions and is the most advanced supervisor on the subject. The Securities and Exchange Commission runs a sandbox. The courts have guidelines. The Higher Education Commission has a draft framework. Health, labour and local government have almost nothing published. A citizen's protection therefore depends on which sector they are standing in.[10,11,12,13]

9. Provincial governments are the implementation frontier. Punjab approved an AI roadmap in March 2026 with a 2029 target and its own delivery unit. Khyber Pakhtunkhwa has a digital strategy and a government language model. Sindh has signed a partnership with Huawei. The federal policy cannot bind any of them, because health, education, police and land are provincial subjects.[14,15,16]

10. Adoption is real but shallow. Mobile broadband reaches about 62 percent of the population by connections, 5G has passed one million users in 23 cities, IT exports reached a record 4.6 billion dollars in FY26, and start-ups raised 133.7 million dollars in 2026 up to the third quarter. But the start-up figure is dominated by one 68 million dollar deal, and the financial sector survey that found about half of institutions deploying or developing AI also found most activity at an early stage.[17,18,19,20,10]

11. The sovereignty strategy depends on foreign hardware and foreign capital. The state has launched GPU-as-a-service through the National Telecommunication Corporation, joined a China-led World AI Cooperation Organization and announced a one billion dollar investment ambition for 2030. It has not published a compute strategy, a chip procurement plan or a position on export controls.[21,22]

12. Rights and redress are the weakest link. There is no AI-specific complaint mechanism, no route for a citizen to learn that a decision affecting them was automated, and no published assessment of any government AI system. The draft data policy promises rights; the personal data bill would create them; neither is law.[4,6]

The argument in one page

The report's argument can be stated simply. Pakistan has built a governance system that is strong on coordination and weak on obligation. The Digital Nation Act coordinates. The AI Policy promotes. The declaration announces. The data policy designs. The standards framework classifies. Almost nothing imposes a duty on anyone to do, or refrain from doing, a specific thing in a specific case, and almost nothing creates a remedy when the duty is missed.

This is not a failure of drafting. It is a choice, made implicitly, to pursue capability before constraint. There is a defensible version of that choice: a country with low AI adoption, weak data institutions and a small fiscal base may reasonably decide that building capability is the priority, and that heavy regulation would entrench incumbents and deter investment. But the choice has costs, and they are not evenly distributed. They fall on people who are subject to automated decisions they cannot see, on workers whose tasks are reorganised without consultation, on citizens whose data moves through an exchange

they cannot inspect, and on the institutions that will one day have to clean up a failure that better governance would have prevented.

The alternative is not a European-style AI Act transplanted into a Pakistani budget. It is a thinner, harder set of obligations: a data protection law with an independent supervisor, a small AI statute that establishes risk tiers and a register of high-impact public systems, mandatory AI clauses in public procurement, sectoral mandates with published guidance, and a redress route that a citizen can actually use. Most of that can be done with the institutions Pakistan already has. What it cannot do without is a decision that implementation matters as much as announcement.

The 2026 scorecard

The table below records the public evidence of progress on selected policy commitments as of 3 November 2026. It is a progress snapshot, not a finding that every listed commitment is overdue. Commitments dated 2025 have passed their stated year; targets dated only 2026 remain within their stated period at the review cutoff. The observations are PIVRA's, based on the records described in the chapters that follow. “No public evidence” means that the review did not locate a published record of delivery; it does not prove that no activity took place.

Table 1. Selected commitments and public evidence of progress as of 3 November 2026

Commitment	Source	Status as of 3 November 2026
National AI Fund established	AI Policy 2025, action 11	No public evidence of a fund, corpus or first call
AI Council notified and meeting	AI Policy 2025, implementation section	Process reported begun; no membership or minutes published
90 percent public AI awareness	AI Policy 2025, action 1	No survey or campaign evaluation published
70 percent of civil servants trained	AI Policy 2025, action 5	Modules announced; no completion data published
1,000 AI research initiatives funded	AI Policy 2025, action 6	No consolidated grant register published
AI patent-registration targets	AI Policy 2025, action 8	No patent register specific to AI published
150 academic institutions given compute access	AI Policy 2025, action 25	GPU service launched; allocation list not published
National data repository	AI Policy 2025, action 19	Data policy still a draft; repository not public
AI Directorate under data commission	AI Policy 2025, action 15	No directorate or commission in operation
AI regulatory sandbox	AI Policy 2025, action 16	SECP sandbox exists; no AI-specific cohort published
Data protection law enacted	Policy commitment since 2018	Bill remains a draft

Commitment	Source	Status as of 3 November 2026
National Digital Masterplan approved	Digital Nation Act 2025	Reported approved in early November 2026; text not public

The pattern is consistent. Where a commitment required an announcement, an announcement was made. Where it required an institution, a budget line, a register or a published decision, the record is thin. This is the implementation gap that the rest of the report documents.

What follows

Chapter 1 sets out the questions the report asks and the method it uses. Chapter 2 establishes the baseline: how much AI Pakistan actually uses, how much it spends, what its compute and data look like, and how independent indices assess it. Chapter 3 describes the constitutional and legal materials that apply when no AI-specific law does. Chapter 4 maps the national institutions. Chapter 5 examines the National AI Policy and its arithmetic. Chapter 6 turns to data governance, privacy and WASL. Chapter 7 explains the standards architecture and why a standard is not a regulation. Chapter 8 examines government deployment and procurement. Chapter 9 covers the sectors: finance, telecoms, courts, universities, health, competition and media. Chapter 10 covers cybercrime, disinformation and online harms. Chapter 11 covers the provinces. Chapter 12 covers rights, inclusion and labour. Chapter 13 covers the international and compute dimension. Chapter 14 is an implementation audit. Chapter 15 sets out scenarios to 2030. Chapter 16 makes recommendations. Chapter 17 concludes.

The report in numbers

Six indicators that frame the governance problem, followed by the composition of the 23-entry governance register.

204.8m
Mobile subscriptions
February 2026

62.3%
Mobile broadband penetration
Connections per 100 people

1.3m
5G users
Across 23 cities

$4.6bn
Technology exports, FY26
Up 21% year on year

Rs 524.64m
AI allocation, FY27
Target: 200,000 trainees a year

98 / 194
AI investment potential
Score 51.11; below global average

23 REGISTER ENTRIES | DOCUMENTARY CATEGORIES



Binding law and rules	2	Published guidance and survey	3
Approved policy and declaration	2	Platforms, services and sandbox	3
Draft frameworks	3	Other bodies and initiatives	10

Figure 1. The 23 register entries, by category. Source: PIVRA classification of Table 9. Each entry is counted once by document type and publication status. Categories do not measure operational readiness. The row-by-row mapping is explained in Chapter 14.

What the numbers show

- Connectivity is high by subscription and shallow by capability: the next-generation penetration rate counts connections, not people who can afford to use them.
- The policy's skills targets cost an estimated Rs 5.2 to 8.2 billion a year against a Rs 0.52 billion AI programme allocation.
- The register includes three draft frameworks and ten other bodies and initiatives. Category counts do not establish whether those entries are operational.

A note on method and evidence

This report is based on publicly available records, including laws, policy documents, consultation drafts, standards, procurement notices, regulator publications, court and committee guidance, official statistics, company and investor disclosures, independent indices, and reporting from Pakistani and international outlets. No interviews were conducted. The assessment is therefore limited to publicly documented information and does not independently verify internal institutional practices or unpublished activity. Findings concerning implementation should be read in this context: the absence of public evidence does not, by itself, establish that no activity has taken place.

What counts as evidence

The report uses a simple hierarchy. A law or a notified rule is the strongest evidence of authority. A policy or a strategy is evidence of intent. A standard or guideline is evidence of a designed process. A procurement notice is evidence of a planned purchase. A press release is evidence of an announcement. A news report is evidence of what was reported. None of these is evidence that a system is operating, that a safeguard has been applied or that a decision has been reviewed.

That hierarchy matters because AI governance produces a great deal of text that looks like implementation. A signed memorandum of understanding is not a deployed system. A training target is not a trained person. A published framework is not a completed assessment. Where the report says that something has been delivered, it relies on a record that shows delivery: an award notice, an annual report, a published register, a completed evaluation or a decision with a date and an author.

Scope

The report covers civilian AI governance: the rules and institutions that apply to AI used by government, businesses, universities, courts and citizens. It does not attempt a full account of military or intelligence AI, which is outside the public record. It treats "AI" as machine-based systems that infer from inputs how to generate outputs such as predictions, content, recommendations or decisions, the definition used by the OECD and adopted in most national frameworks.[23] It uses "governance" to mean the full set of instruments and institutions that shape behaviour, not only statutes. It uses "regulation" more narrowly, for binding rules issued by a body with legal authority.

The report covers the federal government and the four provinces. It notes where Azad Jammu and Kashmir and Gilgit-Baltistan have separate arrangements but does not assess them in detail. It covers the period from January 2025 to 3 November 2026, with earlier material where it explains the present.

A note on numbers

Several figures in this report are contested, provisional or self-reported. Where possible the report gives the source and the date. Currency figures are in the currency used by the source. Fiscal years are identified as FY25 (July 2024 to June 2025), FY26 (July 2025 to June 2026) and FY27 (July 2026 to June 2027). Where a figure is a PIVRA calculation, the text says so and the method is given in the appendix.

Two cautions apply throughout. First, Pakistani data on AI adoption is thin and not collected to a common standard. A survey of 55 financial institutions, a start-up funding total and a count of training participants measure different things and cannot be added together. Second, official announcements

often use round numbers for targets and ambitions. A one billion dollar investment ambition is not a budget line. The report treats it as a statement of intent and says so.

How the report was assembled

The research began with the documents that establish the framework: the Digital Nation Pakistan Act 2025, the National AI Policy 2025, the Islamabad AI Declaration, the draft National Data Governance Policy 2026 and the sectoral instruments. It then worked outward to the institutions that implement them, the budgets that fund them, the procurement that buys them and the evidence that any of it has changed behaviour. The final step was a systematic search for implementation evidence: annual reports, registers, award notices, audit findings, enforcement decisions, evaluation studies and complaint statistics. Where that search produced nothing, the report says so.

The reference list contains the sources used. Each entry gives the issuing body, the title, the date and the address at which the document was consulted. Documents that exist in multiple versions are identified by version or status in the text. A companion data appendix records the indicators used in Chapter 2 and the calculations behind them.

CHAPTER

01

The governance question

A man in Lahore applies for a government loan guarantee. The rejection letter gives no reason. He asks the bank, which says the decision came from the scheme's eligibility system. He asks the scheme, which says the bank submitted the data. He asks for the record, and is told that the record is internal. He has no way of knowing whether a person looked at his file, whether the model used an old address, whether the rule that excluded him was written by the scheme or learned from historical data, or whom to ask for a review. The system that decided his case is not secret. It is simply outside the categories that the law gives him.

IN THIS CHAPTER

Why AI needs its own governance

Pakistan's particular starting point

The five questions

What changed in 2025 and 2026

The argument ahead

Structure

THE ARGUMENT IN ONE LINE

Pakistan has built an architecture of AI governance, but not yet a system of accountability.

FIVE QUESTIONS FOR EVERY AI SYSTEM

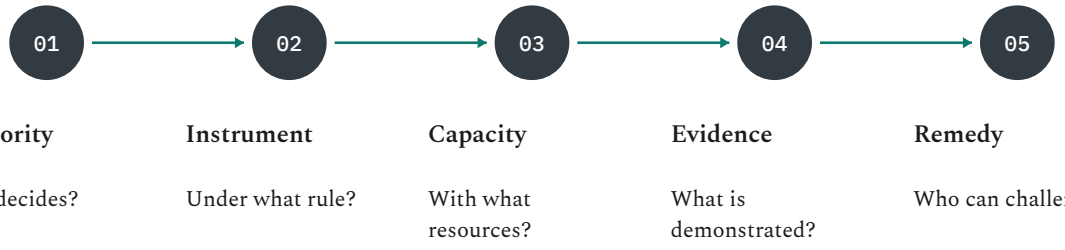


Figure 2. The report's five-question governance test. Source: PIVRA analytical framework, Chapter 1.

This is the ordinary shape of the AI governance problem in Pakistan. It is not usually a robot making a decision. It is a decision that has been partly automated, inside an institution that already had discretion, under a legal framework that was not written with automation in mind. The citizen's difficulty is not that the model is intelligent. It is that the chain of responsibility has become longer than the chain of accountability.

Why AI needs its own governance

Artificial intelligence is not the first technology to test administrative law. Databases, credit scoring, risk profiling and statistical targeting all raised similar questions before large language models made them fashionable. Three features nevertheless make AI distinct enough to require dedicated governance.

The first is generality. A model trained for one task can be adapted to many others, often by people who did not build it and do not fully understand it. A bank's fraud detector and a hospital's triage tool may share the same underlying architecture but produce very different harms. Regulation that targets the technology rather than the use will either be too broad to enforce or too narrow to matter.

The second is opacity, but not in the way it is usually described. The problem is not that a model is a black box. The problem is that the relevant information is distributed across actors: the vendor knows the training data, the deploying institution knows the use case, the auditor knows the test results, and the citizen knows the harm. No single party holds the whole picture. Governance has to assemble it.

The third is speed. A rule that takes three years to write will be applied to a generation of systems that no longer exist. This is why mature jurisdictions have moved towards frameworks that combine binding principles, delegated technical standards and sector-specific supervision, rather than attempting a single comprehensive code. The European Union's AI Act is the most elaborate version of this approach; the NIST AI Risk Management Framework is the most widely used voluntary one; ISO/IEC 42001 turns it into a management system.[24,25,26] None of them can be copied wholesale into a country with Pakistan's institutional capacity. All of them offer design lessons.

Pakistan's particular starting point

Pakistan begins from a position that is unusual in three respects.

It has a large, young, connected population and a small formal AI sector. Mobile broadband reaches about 62 percent of the population by connections, but fixed broadband is limited, electricity is expensive

and unreliable for industrial users, and the IT industry is export-oriented rather than domestically embedded.[17] The country produces talented engineers and loses many of them. It has a vibrant start-up scene in a few cities and a large informal economy that AI will reach late, if at all.

It has a federal constitution in which the subjects that matter most for AI in daily life, health, education, police, land, local government, are provincial. The federal government can set standards, fund programs and coordinate, but it cannot simply order a province to adopt a model, a register or a redress mechanism. Any national AI strategy that ignores this will be a strategy for Islamabad, not for Pakistan.

It has a security establishment that treats data and connectivity as strategic subjects, and a history of internet shutdowns and content blocking that shapes public trust in digital government. That history is not a reason to avoid digital public infrastructure. It is a reason to design it with transparency and independent oversight, because adoption depends on trust and trust depends on what happened last time.

The five questions

This report asks five questions of every institution, instrument and program it examines.

Who decides? Which body has legal authority over the system or the decision, and is that authority explicit or inferred? A ministry that deploys a tool is not necessarily the regulator of that tool. A regulator that supervises a bank is not necessarily the supervisor of the bank's model vendor.

Under what instrument? Is the rule a statute, a regulation, a standard, a guideline, a contract or an announcement? Each has a different legal effect and a different route of challenge. The distinction is not pedantry. It determines whether a citizen can enforce a right or only request a meeting.

With what capacity? Does the responsible institution have staff, budget, technical expertise and data? A mandate without capacity is a press release with a letterhead.

With what evidence? What has actually been tested, published or audited? A framework describing how assessments will be done is not evidence that any assessment was done. This report treats the difference as central.

With what remedy? If the system causes harm, who can the affected person tell, what must that body do, and within what time? A governance system without a remedy is a system for managing reputational risk, not for protecting people.

What changed in 2025 and 2026

The period covered by this report is genuinely important. In January 2025 the Digital Nation Pakistan Act created a National Digital Commission chaired by the prime minister, a Pakistan Digital Authority and an Oversight Committee.[1] In July 2025 the federal cabinet approved the first National AI Policy.[2] In February 2026 the government adopted the Islamabad AI Declaration, a nine-principle statement of direction.[27] In June 2026 the Ministry of IT published a draft National Data Governance Policy.[4] In April 2026 the National Judicial Policy Making Committee issued guidelines for AI in the courts.[12] In September 2026 new public procurement rules were notified.[28] In October 2026 the Senate was told that the government is implementing a national AI ecosystem and had joined a China-led World AI Cooperation Organization.[22]

These are not trivial achievements. A country that had no national AI policy two years ago now has a policy, a coordinating authority, a declaration, a data framework in draft, sectoral guidance and a sovereign compute program. The question is what they add up to.

The argument ahead

The report argues that they add up to a governance architecture that is stronger on coordination than on obligation. That is a deliberate and partly defensible choice, but it has reached its limit. The next phase of Pakistan's AI governance needs fewer strategies and more enforceable instruments: a data protection law, a small AI statute with risk tiers, a register of high-impact public systems, procurement conditions with teeth, sectoral guidance that regulators actually apply, and a redress route that a citizen can use without a lawyer.

The report also argues that the binding constraints on AI in Pakistan are not primarily legal. They are data quality, compute cost, electricity, skills and institutional capacity. A government that passes an AI act but does not fix the data registries, the power tariff for data centres or the capacity of sector regulators will have regulated the wrong thing. This is why the report spends as much time on WASL, procurement and budgets as on statutes.

Finally, the report argues that the most important governance decisions are being made outside the AI policy. They are in the data governance framework, which determines who can see and use records; in procurement, which determines what systems government buys and on what terms; in federalism, which determines whether national standards reach schools and hospitals; and in compute policy, which determines whether Pakistani institutions can build or only rent. AI governance is not a separate domain. It is the intersection of these four.

Structure

Chapter 2 establishes the baseline. Chapter 3 sets out the legal materials that apply when no AI-specific law does. Chapters 4 to 8 describe the national architecture, the AI policy, data governance, standards and procurement. Chapter 9 covers the sectors. Chapter 10 covers cybercrime and online harms. Chapters 11 and 12 cover the provinces, rights and inclusion. Chapter 13 covers the international dimension. Chapter 14 audits implementation. Chapter 15 sets out scenarios. Chapter 16 makes recommendations.

The report is long because the subject is. Readers who want the short version can read the executive summary and Chapter 14. Readers who want the argument can read this chapter and Chapter 16. Readers who want the evidence can use the references and the data appendix.

CHAPTER

02

The state of AI in Pakistan

Any account of AI governance has to begin with a baseline, because governance is proportional to use. A country with a handful of pilots needs a different regime from one with millions of automated decisions a day. Pakistan's baseline is distinctive: high connectivity, low formal adoption, a small but globally integrated IT sector, expensive compute, fragmented data and a serious skills bottleneck. This chapter establishes that baseline and explains why it matters for the rules.

IN THIS CHAPTER

Connectivity without depth

The AI economy: exports, start-ups and public money

Adoption: what the surveys show

Compute and power

Data: the real bottleneck

Talent and research

THE ARGUMENT IN ONE LINE

Adoption is real but shallow; data, compute, skills and trust are the binding constraints.

Connectivity without depth

Pakistan's mobile market is large and has grown steadily. Total cellular subscriptions reached 204.77 million at the end of February 2026, up from 202.58 million in January, with teledensity at 82 percent.[17] 3G and 4G subscriptions reached 155.93 million, a next-generation penetration rate of 62.3 percent.[17] 5G services had reached more than 23 cities and about 1.3 million users by late 2026.[18]

These numbers are impressive in aggregate and misleading if read as a measure of useful internet access. A subscription is not a user, a connection is not a device, and a 4G signal is not a data package that a household can afford. Fixed broadband remains a small share of connections, which matters for AI because model development, data annotation and cloud work are done on fixed connections and reliable power, not on a phone. The gender gap in mobile ownership and mobile internet use has been a persistent feature of Pakistan's digital market in GSMA's surveys, and rural coverage remains below urban coverage even where the signal is nominally available.[29]

The policy consequence is straightforward. A government that wants AI adoption cannot treat mobile penetration as evidence that the population is ready to use AI services. The relevant measures are fixed broadband, device ownership, data affordability, digital literacy and trust, and Pakistan does not publish a regular, independent survey of any of them. This is the first data gap that a serious AI governance program should close.

The AI economy: exports, start-ups and public money

Pakistan's technology sector is export-led. Technology exports reached a record 4.6 billion dollars in FY26, a 21 percent increase over the previous year, according to the finance adviser.[19] That figure includes IT services, software, call centres and related exports. It does not measure AI revenue specifically, and no official series does. The commonly cited claim that Pakistan's AI market will be worth a particular sum by 2030 should therefore be treated with caution: most such figures are vendor estimates built from regional assumptions, not measured Pakistani revenue.

The start-up picture is more informative, and less flattering. Pakistani start-ups raised 133.7 million dollars across 10 deals in 2026 up to the third quarter, according to i2i Ventures.[20] Of the 74.7 million dollars raised in the third quarter, 68 million came from a single round by Fasset, a UAE-headquartered fintech for which Pakistan is a key market.[20] Excluding that deal, Pakistani start-ups raised about 6.7 million dollars in the quarter.[20] Early-stage capital remains scarce, and the AI-specific share of what is raised is not disclosed. The ecosystem produces globally competitive founders and a small number of AI-native companies, but it is not yet a capital market that can fund a generation of AI infrastructure firms.

Public money is small relative to the policy's ambitions. The federal public sector development programme for 2026-27 proposes Rs 16.29 billion for ongoing IT and telecom projects, against a total portfolio cost of Rs 79.79 billion.[3] Most of that is physical infrastructure: Rs 6.81 billion for the Islamabad technology park and Rs 3.75 billion for the Karachi IT park.[3] The National AI Advancement Initiative is allocated Rs 524.64 million, cyber security for digital Pakistan Rs 330 million and the semiconductor human resource programme Rs 250 million.[3] The World Bank-supported Digital Economy Enhancement Project receives Rs 3.02 billion.[3]

These are proposed allocations, not expenditures, and the development budget is only one channel of public spending. But the proportions are revealing. The state is investing far more in buildings and in a general digital economy project than in the AI program that the policy describes as the centrepiece of national capability. That is not necessarily wrong: buildings and digital public infrastructure are

prerequisites. It does mean that the policy's training and research targets cannot be met from the AI line alone.

Adoption: what the surveys show

Pakistan has no national AI adoption survey. The best available evidence comes from regulators and industry.

The State Bank's Financial Stability Review 2024 reported a survey of 55 regulated entities. About half reported that they had deployed AI or were developing it.[10] Reported applications included fraud detection, customer service, marketing, credit-risk assessment and process automation.[10] The finding is useful and easy to misuse. It combines deployed systems with systems under development. It does not mean that half of Pakistani financial institutions run operational AI, and it certainly does not mean that half of financial decisions are automated. It is also a 2024 observation, not a November 2026 adoption measure.

Other sectors are less documented. Telecom operators use AI for network optimisation, customer care and fraud, but do not publish adoption data. Universities use generative tools for teaching and research, with the Higher Education Commission's draft framework as the only national guidance.[13] Government agencies have announced chatbots, document-processing tools and analytics platforms, but there is no central inventory of government AI systems, no register of their purpose or risk, and no published assessment of any of them. The provinces are the most active experimenters, as Chapter 11 describes.

The absence of a national survey is itself a governance problem. Without a baseline, the government cannot know whether its training programs are reaching the firms that need them, whether adoption is concentrated in a few large organisations, or whether the benefits are reaching smaller firms and poorer districts. The policy's KPIs include “percentage of industries adopting AI solutions” but no mechanism for collecting it.[2]

Compute and power

AI compute in Pakistan is scarce, expensive and mostly rented. The National Telecommunication Corporation has launched GPU-as-a-service and model-as-a-service from its state-managed data centres, a significant step because it creates a domestic option for public institutions that cannot or will not use foreign cloud.[21] The service's capacity, pricing and allocation rules have not been published in the material reviewed for this report. Several private data centres and telecom operators offer colocation and cloud services, and the Pakistan Cloud First Policy encourages public bodies to use accredited providers.[30]

The binding constraint is not only chips. It is electricity. Commercial and industrial tariffs have been high and volatile, and data centres are energy-intensive. The competitive electricity market, operational since January 2026, allows large consumers to buy power directly through wheeling, and the first auction for wheeling capacity was issued in September 2026.[31] If it works, it improves the economics of domestic hosting. If it does not, the country's AI workloads will continue to migrate to Singapore, Bahrain, Frankfurt and Virginia, taking the data with them.

There is no published national compute strategy. The AI Policy calls for a national AI compute grid, high-performance computing centres and access for at least 100 academic institutions.[2] The 2026-27 budget does not contain a line for a compute grid. The NTC service is a supply-side intervention without a published demand assessment: how much compute the country needs, for what, at what price, and with what public-interest allocation. That assessment should precede any further procurement.

The geopolitical dimension compounds the domestic one. Advanced AI chips are subject to export controls, and Pakistan's access to them depends on decisions made in Washington, Beijing and Taipei. The United States' 2025 AI diffusion framework, which would have placed most countries in tiers with quantitative caps, was rescinded later that year, but the underlying controls on advanced accelerators remain.[32] Pakistan has not published a position on these controls, a procurement plan for restricted hardware or a strategy for obtaining compute through allies. That silence is a risk.

Data: the real bottleneck

If compute is the visible constraint, data is the binding one. AI systems depend on records that are accurate, complete, comparable and lawfully usable. Pakistan's public records are held in silos: identity with NADRA, tax with the FBR, land with provincial boards of revenue, health with provincial departments, education with HEC and provincial commissions, utilities with their own companies. The records are often paper-based, inconsistent across districts and not maintained to a common standard.

The government's answer is WASL, the National Data Exchange, which is designed to let authorised institutions request and receive data without centralising custody.[4] The model is federated: identity data stays with NADRA, land data stays with the provincial authority, and WASL moves a governed request and response. The draft data policy adds residency tiers, a national data catalogue, primary data registers, a National Data Governance Council and a National Data Maturity Index.[4] Chapter 6 examines these in detail.

For the baseline, three points matter. First, WASL is a design and a platform, not yet a national data infrastructure with published coverage and transaction statistics. Second, the legal basis for much of the data sharing that AI would require is unclear, because the personal data protection bill is not law and the data policy is not in force. Third, the data that matters most for Pakistani AI, Urdu and regional language text, speech, local administrative records and domain-specific corpora, is scarce relative to English. The Center for Language Engineering at UET Lahore and the National Center for Artificial Intelligence at NUST have built valuable resources, but they are small relative to the scale that foundation model training requires.

Talent and research

Pakistan has a deep pool of engineering talent and a thin layer of AI specialists. The AI Policy's targets are ambitious: 200,000 people trained a year, 10,000 trainers by 2027, 3,000 postgraduate scholarships a year and 20,000 internships a year.[2] The 2026-27 budget provides Rs 524.64 million for the AI initiative and Rs 250 million for semiconductor human resource development.[3] At any plausible cost per trainee, the training target alone exceeds the budget by an order of magnitude. That does not make the target wrong. It makes it a target that depends on provincial, private and donor funding, and on delivery partners that have not yet been named.

The higher education system is responding. HEC has a draft framework for generative AI in teaching and assessment, and reported training more than 332 faculty members in September 2026, 32 in person and over 300 through university facilities.[13,33] The Civil Services Academy introduced an AI module in January 2026.[34] These are useful capacity-building steps. They are not a national skills system, and they do not measure whether participants can build, procure, audit or govern AI systems.

Research output is difficult to assess without a dedicated bibliometric study. What can be said is that Pakistan's AI research is concentrated in a small number of universities and research centres, that the policy's own patent target implies a very low baseline, and that the incentives for researchers to publish

internationally are stronger than the incentives to build local datasets or evaluate deployed systems. A national AI research strategy should reward the latter, not only the former.

AI INVESTMENT POTENTIAL | SCORE / 100

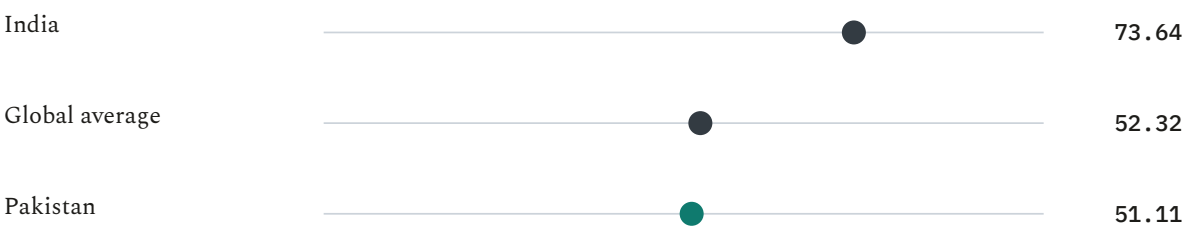


Figure 3. AI investment potential: selected scores. Source: AFD Artificial Intelligence Investment Potential Index 2025. Scores shown on a 0-100 scale.

How the world measures Pakistan

International indices are blunt instruments, but they are useful for locating Pakistan relative to peers.

The Agence Francaise de Developpement's Artificial Intelligence Investment Potential Index 2025 ranked Pakistan 98th out of 194 countries with a score of 51.11, slightly below the global average of 52.32, placing it in Stage 3 of four.[35] India ranked 30th with 73.64.[35] The index weighs economic environment, governance, digital and physical infrastructure, data governance and security, and statistical performance.[35] Its authors describe Stage 3 countries as facing considerable challenges that require targeted policy intervention in infrastructure and human capital.[35]

Oxford Insights' Government AI Readiness Index 2025, which assesses 195 governments, noted Pakistan among the countries that improved, alongside Tanzania and Jamaica, in the context of the Commonwealth AI Consortium.[36] The index's methodology changed in 2025 to emphasise the state's role as buyer, enabler and regulator, which makes year-on-year comparisons hazardous.[36]

The variation across indices is itself informative. Pakistan performs relatively better on measures that reward policy adoption and connectivity, and relatively worse on measures that require institutional depth, data governance and human capital. That is consistent with the argument of this report: the country is good at producing strategies and less good at building the institutions that make strategies real.

Table 2. Baseline indicators for AI in Pakistan

Indicator	Value	Date and source
Cellular subscriptions	204.77 million	February 2026, PTA
Teledensity	82 percent	February 2026, PTA
3G and 4G subscriptions	155.93 million	February 2026, PTA
Next-generation mobile penetration	62.3 percent	February 2026, PTA

Indicator	Value	Date and source
5G users and coverage	About 1.3 million users in 23 cities	2026, PTA data reported in the press
Technology exports	4.6 billion dollars, up 21 percent	FY26, Ministry of Finance adviser
Start-up funding	133.7 million dollars across 10 deals	2026 to Q3, i2i Ventures
AI investment potential	98th of 194, score 51.11	2025 index, published 2026, AFD
Government AI readiness	Improved, no stable rank published	2025, Oxford Insights
Federal IT development portfolio	Rs 16.29 billion proposed	FY27, PSDP
National AI Advancement Initiative	Rs 524.64 million proposed	FY27, PSDP
Financial institutions deploying or developing AI	About half of 55 surveyed	2024, SBP

What the baseline implies for governance

Four conclusions follow from this baseline.

First, Pakistan's AI adoption is shallow and concentrated. Most activity is in a few large banks, telecom operators, technology firms and government agencies in Lahore, Karachi and Islamabad. The policy's benefits will not reach the rest of the economy without deliberate diffusion measures, and the current instruments do not include any.

Second, the constraints are material, not legal. Compute, electricity, data quality and skills are the binding constraints. A government that passes an AI law but does not fix the power tariff for data centres, the quality of land records or the supply of AI teachers will have regulated the wrong thing. This is not an argument against regulation. It is an argument for sequencing it with investment.

Third, the evidence base is weak. There is no national adoption survey, no government AI inventory, no published compute assessment and no independent evaluation of any deployed system. The government cannot govern what it cannot see, and at present it cannot see much.

Fourth, the baseline supports a proportionate approach: strict rules for high-impact systems, light-touch rules for the rest, and the data and compute foundations that make both possible.

CHAPTER

03

The legal baseline: what applies when no AI law does

Pakistan has no artificial intelligence statute. That does not mean AI is unregulated. It means that AI is governed by the general law, applied by institutions that were not designed for it. Understanding what that general law says, and where it runs out, is the necessary foundation for any assessment of the newer instruments.

IN THIS CHAPTER

The Constitution

The general statutes

The data protection gap

Administrative law and the ouster clause

Liability, contracts and consumer protection

What the baseline adds up to

THE ARGUMENT IN ONE LINE

The general law reaches AI, but it was not built for distributed, automated decisions.

The Constitution

Three constitutional provisions do most of the work.

Article 14 provides that the dignity of man and, subject to law, the privacy of home, shall be inviolable.[37] The Supreme Court has read privacy beyond the physical home. In *Mohtarma Benazir Bhutto v. President of Pakistan*, it treated telephone tapping as unconstitutional in the absence of a lawful procedure.[38] In the *Justice Qazi Faez Isa* litigation, privacy was revisited in the context of surveillance and the review of state action.[39] The cases establish that an intrusion into privacy requires a legal basis, a legitimate aim and a proportionate procedure. They do not establish a general data protection code, and they do not give a citizen a direct right to inspect an algorithm.

Article 19A guarantees the right to information in all matters of public importance, subject to reasonable restrictions imposed by law.[37] This is potentially the most important constitutional hook for AI governance, because it grounds a claim to know what systems government uses, what data they process and what decisions they support. The Right of Access to Information Act 2017 gives effect to the right, with exemptions for national security, personal privacy and commercial confidence. In practice, requests about procurement or software are often met with the commercial-confidence exemption, and there is no published jurisprudence on how that exemption applies to AI systems.

Article 4 requires that every person be dealt with in accordance with law and be entitled to due process. A decision that affects a person's rights, made by an automated system without a reasoned explanation or a review, is vulnerable under this provision even if no AI-specific law exists. The difficulty is evidential: the affected person usually cannot show what the system did.

The general statutes

The Electronic Transactions Ordinance 2002 recognises electronic records and signatures and provides a framework for electronic contracting. It is the legal basis for much of digital government, but it says nothing about automated decision-making, explainability or liability for algorithmic harm.

The Prevention of Electronic Crimes Act 2016 is the main criminal statute for digital conduct. It criminalises unauthorised access, copying, interference, electronic forgery and fraud, unauthorised use of identity information and the disclosure of personal information by service providers.[40] The 2025 amendment created the National Cyber Crime Investigation Agency, added an offence of disseminating false information likely to cause fear or unrest, and established a Social Media Protection and Regulatory Authority with powers to order the removal or blocking of unlawful content.[40] The Authority was constituted in March 2026, but reporting in July 2026 indicated that it still lacked a budget, an office and staff, and that its tribunals had not been established.[6]

PECA matters for AI governance in three ways. It is the route for complaints about deepfakes, impersonation and non-consensual imagery, which are among the most immediate AI harms. It imposes data retention and disclosure obligations on service providers, which shape the privacy of AI-mediated communications. And its new false-information offence creates a risk that legitimate reporting about AI systems is chilled, particularly where a journalist or researcher publishes a contested claim about a government deployment.

The Right of Access to Information Act 2017, the Public Procurement Rules 2026 and the sectoral statutes described in Chapter 9 complete the general framework.[28] Each was written for a world of documents, contracts and human officials. Each can be adapted to AI, but none was designed for it.

The data protection gap

Pakistan has been trying to pass a personal data protection law since 2018. A bill was approved by the federal cabinet in February 2022, a revised draft followed in 2023, and as of October 2026 the bill remained a draft.[5,41,6] The 2023 draft proposed a National Commission for Personal Data Protection, breach notification within 72 hours where feasible, restrictions on cross-border transfers and a localisation requirement for critical personal data. None of it is in force.[6]

The result is a patchwork. Banks, microfinance banks, payment system operators and insurers are subject to State Bank and SECP rules on outsourcing, cloud hosting and customer confidentiality.[42] Telecom operators are subject to PTA regulations on spam and unsolicited communications. Companies are subject to the general law of confidence and to their contracts. Citizens have Article 14 and the criminal provisions of PECA, but no supervisory authority, no registration duty, no statutory privacy notice and, outside regulated sectors, no general breach notification obligation.[6]

This gap is the single most important legal fact about AI in Pakistan. Every AI system depends on data. A country without a data protection law cannot give citizens a meaningful right to know what data is used, to correct it, to object to its use or to obtain a remedy when it is misused. The draft National Data Governance Policy 2026 does some of this for public-sector data, but it is a policy, not a statute, and it expressly leaves the supervision of personal data rights to a future authority under a future law.[4]

Administrative law and the ouster clause

AI decisions by public bodies are administrative decisions. They are subject to the ordinary principles of administrative law: legality, rationality, procedural fairness and the duty to give reasons where the decision affects rights. High Courts exercise writ jurisdiction under Article 199 of the Constitution, and the Supreme Court can review under Article 184.

The Digital Nation Pakistan Act 2025 contains a sweeping ouster clause. Section 29 provides that no decision or action taken under the Act, or under rules or regulations made under it, shall be questioned by any agency or challenged in any court or tribunal, and that no injunction shall be granted against such decisions or actions.[1] Read literally, the clause would exclude judicial review of the Pakistan Digital Authority's decisions, including decisions with AI implications. Whether it would survive constitutional challenge is uncertain: Pakistani courts have generally held that ouster clauses cannot bar writ jurisdiction where an authority acts without jurisdiction, in bad faith or in violation of fundamental rights. But the clause creates a chilling effect and an asymmetry. A citizen who wants to challenge an AI-related decision by the Authority must first persuade a court that the ouster clause does not apply.

Section 28 is equally significant. It gives the Act overriding effect over other laws, subject to a carve-out for laws that protect data rights and cybersecurity.[1] That carve-out is important, but it is undefined. Which laws count as data rights laws when the data protection bill is not enacted? What happens when a PDA standard conflicts with a sectoral regulation made under the SBP Act or the PTA Act? The Act says the Authority must consult and that the Commission may resolve conflicts, but it does not say who wins in the meantime.[1] For AI, where sectoral regulators already have mandates over banks, telecoms and securities, this is a recipe for uncertainty.

Liability, contracts and consumer protection

When an AI system causes harm, the legal route depends on the relationship. A patient harmed by a clinical decision support tool may sue in negligence, but must prove a duty, a breach and causation, which is difficult when the tool's behaviour is opaque. A consumer harmed by a misleading chatbot may rely on consumer protection law, but provincial consumer statutes were not written for generated content. A business harmed by a faulty model may rely on its contract, but the contract may exclude liability for exactly the errors that occurred.

There is no strict liability regime for AI, no mandatory insurance, no statutory presumption that a deployer is responsible for the system it uses, and no collective redress mechanism. The law therefore defaults to fault, individual action and contract. That favours large firms with legal departments and disadvantages everyone else.

What the baseline adds up to

The general law is not empty. It provides privacy protection, criminal remedies for some AI harms, access to information, procurement rules and administrative review. But it has four structural weaknesses.

First, it is addressed to human decisions. The duties to give reasons, to hear the affected person and to review a decision are all designed for an official who can explain what they did. An automated system distributes that explanation across a vendor, a data set and a model.

Second, it is sectoral and fragmented. A bank customer has more protection than a patient, a student or a gig worker, not because the harm is greater but because the regulator is more developed.

Third, it lacks a data protection authority. Without one, the rights that the Constitution implies cannot be exercised in practice.

Fourth, it contains an ouster clause that may shield the most important new AI institution from review. That is the opposite of what a new governance regime should do.

Table 3. General legal instruments and their relevance to AI

Instrument	What it covers	AI relevance	Principal gap
Constitution, Articles 4, 14, 19A	Due process, privacy, access to information	Grounds for review and disclosure	No algorithmic explanation duty
Electronic Transactions Ordinance 2002	Electronic records and signatures	Digital government, e-contracts	Silent on automated decisions
PECA 2016, as amended 2025	Cybercrime, content, identity, data disclosure	Deepfakes, impersonation, data misuse	Criminal route only; no regulatory standards
Right of Access to Information Act 2017	Access to public records	Disclosure of AI procurement and use	Commercial-confidence exemption
Public Procurement Rules 2026	Federal procurement	Contract terms for AI systems	No AI-specific clauses
Sectoral laws and regulations	Banking, securities, telecoms, insurance	Supervisory oversight of AI use	Uneven capacity and coverage
Draft Personal Data Protection Bill	Personal data rights	Consent, transfer, breach, redress	Not enacted

CHAPTER

04

The national architecture

Pakistan's AI governance has a centre, but the centre is a coordination body rather than a regulator. The Digital Nation Pakistan Act 2025 created three institutions and a planning instrument. The AI Policy added a council and a cell. The Islamabad Declaration added a set of principles. Understanding how these fit together, and where they do not, is essential to any judgment about who is in charge.

IN THIS CHAPTER

The Digital Nation Pakistan Act 2025

The Authority in practice

The Ministry of IT

The AI Council and the implementation cell

The Islamabad AI Declaration

The wider institutional map

THE ARGUMENT IN ONE LINE

Coordination is crowded. No single body owns the outcomes.

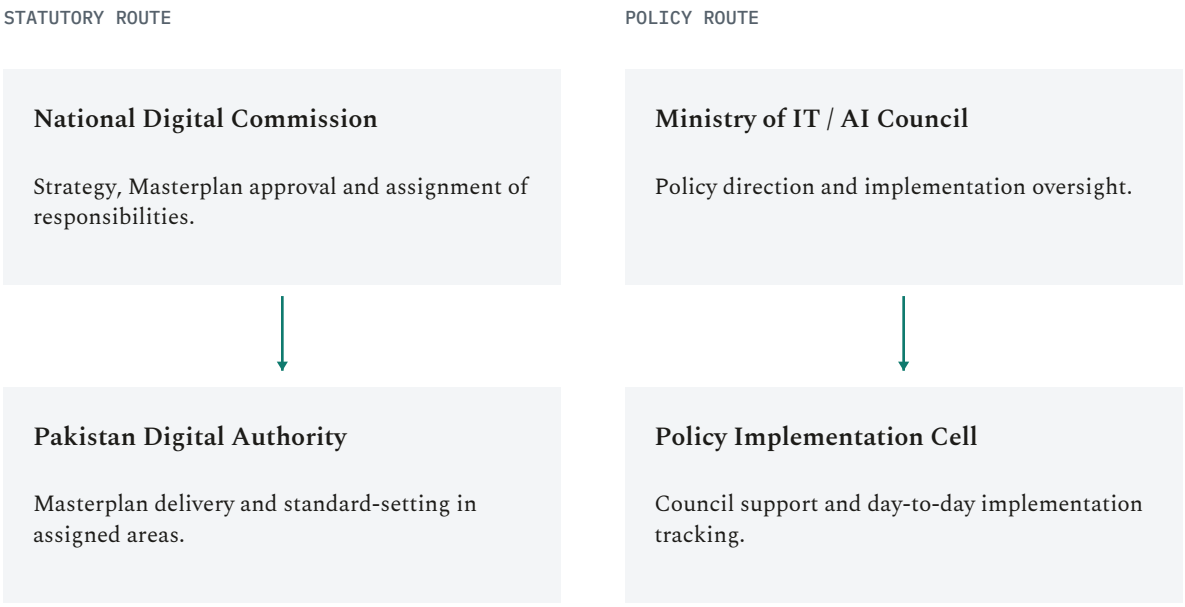


Figure 4. Pakistan's national AI governance architecture. Source: Chapter 4. Simplified institutional map based on the Digital Nation Pakistan Act 2025 and National AI Policy 2025. Policy-created bodies are shown as designed; their operational status is discussed in the text. Arrows denote the two routes, not a comprehensive reporting hierarchy.

The Digital Nation Pakistan Act 2025

The Act was passed on 29 January 2025 and published in the Gazette on 6 February 2025 as Act No. I of 2025.[1] It extends to the whole of Pakistan and came into force at once. Its purpose is to transform Pakistan into a digital nation by enabling a digital society, a digital economy and digital governance.

It creates three bodies.

The National Digital Commission is chaired by the prime minister and includes the four chief ministers, the federal ministers for IT, planning, finance, commerce, interior, economic affairs and information, and the chairs of the FBR, NADRA, PTA and SECP, the Governor of the State Bank and the chair of the Pakistan Digital Authority.[1] It approves the National Digital Masterplan and its implementation plans, gives strategic direction to the Authority, issues directives to align entities with the Masterplan and resolves non-compliance referred to it.[1] It must meet at least once every six months, with a quorum of three quarters of its members including the chair.[1]

The Pakistan Digital Authority is a corporate body with a chair and two members, appointed by the prime minister for five-year terms, with a maximum of five members.[1] It develops and implements the Masterplan, issues regulations, guidelines and standards, coordinates across federal and provincial governments, monitors compliance and advises on legislation.[1] Its functions explicitly include oversight, standard-setting and enforcement in areas assigned by the Commission, including data management and governance, artificial intelligence, virtual assets and other emerging technologies.[1] The qualification is

doing a great deal of work in that sentence: AI is within the Authority's remit only in areas assigned by the Commission, and the Authority must ensure that it does not overlap with other regulators.[1]

The Oversight Committee monitors the Authority's performance and financial management. It is convened by the federal IT minister and includes the IT secretary, the finance and planning secretaries, a representative of the Special Investment Facilitation Council and four independent private-sector members appointed by the prime minister for three-year terms.[1]

The Act also creates a Digital Nation Fund, a non-lapsable fund administered by the Authority, with contributions from the federal government, provincial governments, grants, loans and international sources.[43] The Authority must publish an annual report within 90 days of the end of the financial year, including audited accounts, progress against KPIs and recommendations.[43]

The Authority in practice

The Authority was constituted in August 2025. Dr Sohail Munir was appointed chair and Muhammad J. Sayer a member, both on five-year contracts, after a search committee led by the federal IT minister recommended candidates to the prime minister.[8]

Its public presence is still thin. It maintains a website, a careers page and a standards portal. It has advertised senior posts, including a Chief AI Officer vacancy published in October 2026 with responsibilities for developing governance proposals and assurance processes within the Authority's assigned remit.[9] Those advertisements are evidence of intent to build capability. They are not evidence that the positions have been filled, that assessments have been completed or that the proposed processes are operational.

The Authority has not, in the public record reviewed for this report, published an annual report, a budget, a staffing plan, a regulatory agenda or a decision log. That matters for two reasons. First, the Act requires an annual report, and its absence is a compliance question in its own right. Second, an institution that is meant to coordinate the entire digital state cannot build trust with the public, the provinces or the private sector if its own operations are opaque.

The Ministry of IT

The Ministry of Information Technology and Telecommunication remains the policy owner. It houses the AI Policy, the policy implementation cell, the National AI Council secretariat in practice, and the attached organisations that deliver programs: Ignite, the Pakistan Software Export Board, the National Information Technology Board, the National Telecommunication Corporation and the Pakistan Telecommunication Authority.[2]

The relationship between the Ministry and the Authority is not fully specified. The Act gives the Authority operational authority over the Masterplan and standard-setting, while the Ministry retains policy. In most systems that division works. In Pakistan it is complicated by the fact that the Ministry also chairs the AI Council, controls the budget line for AI programs and supervises the organisations that deliver them. The Authority can set a standard; the Ministry can decide whether anyone gets the money to meet it.

The AI Council and the implementation cell

The AI Policy creates an AI Council as the apex body for strategic direction and implementation oversight. It is chaired by the federal minister for IT and includes the secretaries of IT, science and technology, and planning, the HEC chair, the PDA chair, the provincial chief secretaries, one expert each from academia, industry, civil society and citizen advocacy, and representatives from sectors such as health and agriculture.[2]

The Council is meant to review progress, facilitate federal-provincial collaboration, guide resource allocation from the National AI Fund and provide guidelines for ethical oversight.[2] A Policy Implementation Cell under the Ministry is meant to manage day-to-day implementation, support the Council and track progress.[2]

In early 2026, reporting indicated that the process of forming the first Council had begun.[7] By November 2026 no membership list, terms of reference, meeting schedule or minutes had been published, and the implementation cell had no public presence. It is possible that the Council has met without public disclosure. For governance purposes, an unpublished council is indistinguishable from one that has not met: it cannot be held accountable, and its decisions cannot be tested.

The Islamabad AI Declaration

On 9 February 2026 the government adopted the Islamabad AI Declaration on sovereign, responsible and capability-driven artificial intelligence.[27] It sets out nine principles:

- Sovereign purpose and public value: AI as a sovereign choice to strengthen productivity, governance and measurable public outcomes.
- Human accountability and constitutional authority: AI augments human capacity, does not replace lawful authority, and decisions of public consequence remain under accountable human oversight.
- Pragmatic delivery, use case first: deployment where it demonstrably improves productivity, service delivery and decision support, with scale following proof of impact.
- Sovereign data and trusted stewardship: national, state and citizen data governed under Pakistani law and sovereign interests.
- Explainable, auditable, risk-proportionate systems, particularly in government and where public rights or critical functions are affected.
- Coordinated whole-of-government governance, avoiding fragmentation and vendor dependency.
- Capability, inclusion and responsible innovation.
- Private-sector-led, sovereign compute.
- Global leadership and responsible AI diplomacy.[27]

The Declaration is a statement of direction, not a legal instrument. It creates no duties, no institutions and no remedies. Its value is political and interpretive: it signals that the government's approach is sovereign, use-case-led and private-sector-friendly, and it gives officials a shared vocabulary. Its risk is that it substitutes for binding rules. A declaration of principles is not a substitute for a statute, and the two should not be confused in public communication.

The wider institutional map

Beyond the core, AI-relevant authority is distributed across a large number of bodies:

- The Higher Education Commission regulates universities and has issued a draft generative AI framework.[13]
- The National Judicial Policy Making Committee issued AI guidelines for the courts.[12]
- The State Bank supervises banks and payment systems, and has surveyed AI use.[10]
- The Securities and Exchange Commission supervises capital markets, insurance and non-bank finance, and runs a regulatory sandbox.[11]
- The Pakistan Telecommunication Authority regulates telecoms and the internet, and is the technical regulator for blocking and lawful interception.
- NADRA holds identity data and operates the identity layer of the data exchange.[4]
- The National Information Technology Board builds and operates government platforms.
- Ignite funds research and innovation, and is the proposed home of the National AI Fund.[2]
- The National Telecommunication Corporation operates state data centres and now offers GPU and model services.[21]
- The National CERT issues cybersecurity guidance.[44]
- The National Cyber Crime Investigation Agency investigates cybercrime.[45]
- The Social Media Protection and Regulatory Authority, created in 2025, is meant to regulate online content but is not yet operational.[6]

The result is a dense map with no single AI regulator. That is not necessarily a defect. Many countries deliberately distribute AI supervision across sector regulators, with a central coordinator for standards and cross-cutting risks. The defect in Pakistan's case is that the coordinator's authority is contingent, the sector regulators' AI capacity is uneven and the central policy council is not visibly functioning.

Table 4. Core institutions and their AI responsibilities

Institution	Legal basis	AI role	Accountability
National Digital Commission	Digital Nation Act 2025	Approves Masterplan, assigns AI functions, resolves conflicts	Prime minister, Parliament
Pakistan Digital Authority	Digital Nation Act 2025	Standards, data governance, AI oversight in assigned areas	Commission, Oversight Committee
Oversight Committee	Digital Nation Act 2025	Performance and financial oversight of PDA	Commission
Ministry of IT	Rules of business	AI policy, AI Council, implementation cell, programs	Federal government, Parliament
AI Council	AI Policy 2025	Strategic direction and implementation oversight	Minister for IT
Policy Implementation Cell	AI Policy 2025	Day-to-day implementation and tracking	Ministry of IT
Sector regulators	Sector statutes	Supervision of AI use in their sectors	Their statutes and courts
HEC, NJPMC	Statutory and policy mandates	Guidance for universities and courts	Their constituents

The coordination problem

The architecture has four overlapping coordination mechanisms: the National Digital Commission, the AI Council, the PDA and the sector regulators. Their memberships overlap, their mandates intersect and their relative authority is unclear in several places. The Act says the Authority must avoid overlap with other regulators, but it does not say how. The AI Policy says the Council provides strategic direction, but it does not say how its decisions bind the Authority or the provinces. The Declaration says governance should be coordinated, but it does not create a mechanism.

This is not a purely technical problem. Coordination failures have real consequences. If the Authority issues a data standard that a sector regulator has not adopted, banks may face two regimes. If the AI Council approves a training target that the Ministry cannot fund, the target becomes an unfunded mandate. If the provinces pursue their own AI roadmaps, the national standard may be irrelevant to the services that most citizens use.

The next chapter examines the policy that is meant to resolve these questions. It does not resolve them.

CHAPTER

05

The National AI Policy: ambition and arithmetic

The National Artificial Intelligence Policy 2025 is the most important AI document Pakistan has produced. It is also the most revealing. Read as an industrial strategy, it is coherent, ambitious and in places well designed. Read as a governance instrument, it is thin. This chapter describes what the policy contains, what it promises, what it would cost and what it leaves out.

IN THIS CHAPTER

Origins and status

The six pillars

The money

The skills arithmetic

Infrastructure and sandboxes

The governance actions

THE ARGUMENT IN ONE LINE

The policy is a credible industrial strategy and a weak regulatory instrument.

Origins and status

Pakistan's first AI policy consultation draft was published in June 2022.[46] The approved policy followed three years later, on 31 July 2025.[46] The Ministry's policy register lists the National Artificial Intelligence Policy as approved, with the 2022 consultation draft retained below it.[46] A reader who searches for the policy may therefore find either version. The approved text is the one dated 2025 and titled National Artificial Intelligence Policy; the 2022 document is a consultation draft and should not be cited as current policy.

The policy describes itself as evidence-based, target-oriented, user-centric and forward-looking. It is explicitly linked to the Digital Pakistan Policy, the Cloud First Policy, the draft Personal Data Protection Act, the Digital Nation Pakistan Act and Uraan Pakistan, the national economic transformation plan.[2] That linkage is important: the AI Policy is not a standalone document but one element of a wider digital and economic program.

The six pillars

The policy is organised around six pillars.

The first is the AI innovation ecosystem: a National AI Fund, centres of excellence, an innovation fund, a venture fund, data standards, diaspora engagement and monitoring.[2]

The second is awareness and readiness: a national awareness program, workforce training, scholarships, programs for marginalised women and persons with disabilities, public-sector readiness and intellectual property reform.[2]

The third is a secure AI ecosystem: AI-based cybersecurity, transparency and human oversight, national data security, trust and identity management, a data privacy policy, an open-source AI governance framework, generative AI controls and sandboxes.[2]

The fourth is transformation and evolution: a ranking management system, industrial transformation and public-private sector evolution.[2]

The fifth is AI infrastructure: a national compute grid, data infrastructure and open data, AI hubs, cloud and shared resources, and sandboxes.[2]

The sixth is international partnerships: bilateral and multilateral agreements, participation in global forums, collaborative research, cross-border projects, adoption of international standards, investment, capacity building and trade promotion.[2]

The pillars are sensible. They reflect the standard taxonomy used by countries at Pakistan's stage of development: build supply, build demand, build trust, build connections. The question is not whether the taxonomy is right. It is whether the instruments underneath it are funded, sequenced and enforceable.

The money

The policy's main financial instrument is the National AI Fund, to be established using existing resources under Ignite's research and development fund, with at least 30 percent of that fund allocated to AI on a perpetual basis.[2] A fund committee was to be formed within two months of the policy's enactment.[2] The policy also proposes an innovation fund under the centres of excellence and a venture fund for post-seed financing.[2]

The centres of excellence are to be established in Islamabad, Karachi and Lahore, with auxiliary centres in Peshawar, Quetta, Faisalabad and Hyderabad.[2] They are to receive direct funding from the National AI Fund and to provide computing infrastructure, incubation, curriculum support, training and sandbox assistance.[2]

Whether any of this exists is a separate question. The public record reviewed for this report contains no notification of the National AI Fund, no committee membership, no first call for proposals, no audited corpus and no annual report. It is possible that Ignite has allocated funds internally without a public announcement. But a fund that cannot be found in the budget documents or the gazette is, for practical purposes, not yet a fund.

The skills arithmetic

The policy's skills targets are the most quoted part of the document, and the most revealing when costed.

It promises a National AI Skill Development Program, with NAVTTC, provincial governments, the Ministry of Education and the private sector, to train 200,000 individuals a year in AI, including AI ethics.[2] It promises 10,000 trainers by 2027, 20,000 stipend-based internships a year, 3,000 postgraduate and doctoral scholarships a year, 400 funded AI projects a year with support of up to Rs 1 million each, and 200 funded theses a year with support of up to Rs 200,000 each.[2] The action matrix adds a target of one million people trained by 2027.[2]

These are not compatible with the budget line. The 2026-27 development budget proposes Rs 524.64 million for the National AI Advancement Initiative, which is the program that would have to deliver much of this.[3] A simple calculation shows the scale of the gap.

If the public cost of training one person is Rs 15,000, which is a low estimate for anything beyond a short online course, 200,000 trainees cost Rs 3 billion a year. At Rs 30,000, the cost is Rs 6 billion. The 20,000 internships, even at a stipend of Rs 15,000 a month for three months, cost Rs 900 million. The 3,000 scholarships, at Rs 300,000 a year each, cost Rs 900 million. The 400 project grants, at the maximum of Rs 1 million, cost Rs 400 million. The 200 theses, at the maximum of Rs 200,000, cost Rs 40 million. The total is between Rs 5.2 billion and Rs 8.2 billion a year before any spending on centres, compute or administration.[2,3]

These are PIVRA estimates, not official figures, and the policy assumes co-funding from provinces, universities and the private sector. But the order of magnitude is the point. The program as described cannot be delivered from the AI line alone. It requires either a much larger allocation, a much smaller target, or a financing model that the policy does not describe.

Table 5. Selected AI Policy targets and indicative annual cost

Target	Volume	Unit cost assumption	Indicative annual cost
Trained individuals	200,000	Rs 15,000 to 30,000	Rs 3.0 to 6.0 billion
Internships	20,000	Rs 45,000 stipend	Rs 0.9 billion
Postgraduate scholarships	3,000	Rs 300,000	Rs 0.9 billion
Funded projects	400	Up to Rs 1 million	Up to Rs 0.4 billion
Funded theses	200	Up to Rs 200,000	Up to Rs 0.04 billion
Indicative total			Rs 5.2 to 8.2 billion
AI program allocation, FY27			Rs 0.52 billion

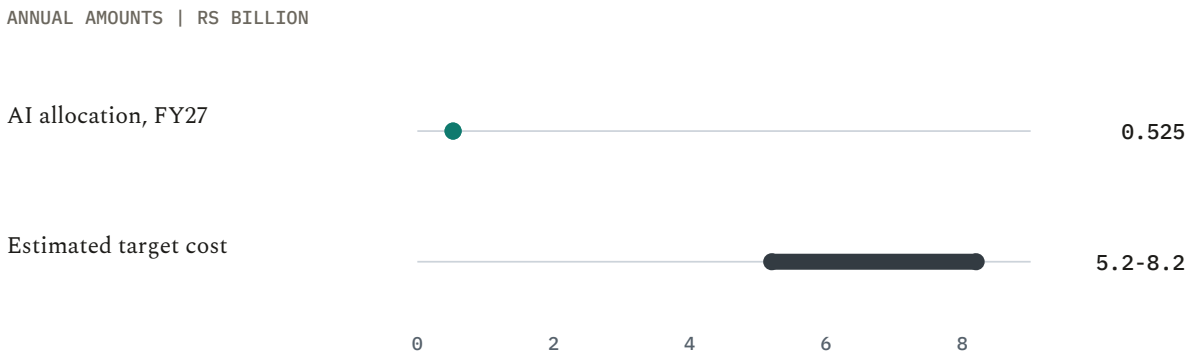


Figure 5. The AI allocation and indicative annual cost of the skills targets. Source: PIVRA calculations, Table 5 and Appendix C. Cost range reflects low and high unit-cost assumptions. FY27 allocation is proposed funding, not expenditure.

Source: PIVRA calculation from targets in the National AI Policy 2025 and the proposed FY27 allocation. Assumptions are illustrative and exclude co-funding.

Infrastructure and sandboxes

The infrastructure pillar is more concrete. It proposes a national AI compute grid, high-performance computing centres and access for at least 100 academic institutions.[2] It proposes a national and provincial AI data repository and the upgrading of public-sector data centres with high-performance computing capability.[2] It proposes a nationwide network of AI hubs in major cities.[2] It proposes that Pakistan contribute about 50 new AI models a year to open platforms.[2] It proposes a regulatory sandbox, with at least 50 enterprises benefiting by 2027.[2]

Some of this is moving. The National Telecommunication Corporation has launched GPU-as-a-service and model-as-a-service from state data centres.[21] The Ministry has tendered for national AI hubs.[47] A consultancy procurement for the hubs was advertised through the electronic procurement system.[48] The Punjab Information Technology Board has separately tendered for a government AI cloud.[49] But a tender is not a hub, and a GPU service without published capacity, price and allocation rules is not yet a compute strategy.

The sandbox target is the clearest test of the policy's governance ambitions. The Securities and Exchange Commission already operates a regulatory sandbox, and its first cohort admitted six solutions for testing.[11,50] But there is no published AI-specific sandbox cohort, no cross-sector sandbox, and no evidence that the 50-enterprise target has been approached. The policy assigns the sandbox to the SECP and the centres of excellence, which is an odd pairing: the SECP regulates financial services, not AI in general, and the centres do not yet exist.[2]

The governance actions

The policy's implementation section creates the AI Council and the Policy Implementation Cell described in Chapter 4.[2] The action matrix adds two governance commitments that deserve particular attention.

The first is to establish an AI Directorate under the National Commission for Personal Data Protection by 2026.[2] The Commission does not exist, because the Personal Data Protection Bill has not

been enacted.[5] The action therefore assigns a function to a body that the legislature has not created. It is the clearest example of the policy's tendency to assume institutions into existence.

The second is to create a regulatory sandbox for AI initiatives, with the SECP and the centres of excellence as leads.[2] A sandbox is a useful instrument, but it is not a substitute for rules. It allows a regulator to learn from a controlled deployment. It does not create standards for the systems that are not in the sandbox, which is most of them.

The policy also assigns the National Commission for Data Protection responsibility for making Pakistani data available for 100 commercial models a year.[2] Again, the body does not exist.

The action matrix

The policy's action matrix is the most useful part of the document because it converts aspirations into dated commitments. It contains 26 actions. Twenty carry a 2026 timeline, two are dated 2025 and four are dated 2027. Commitments dated 2025 have passed their stated year. For targets dated only 2026, this report assesses the public evidence of progress at 3 November 2026 without treating the year-end deadline as already missed.[2]

Classified by type, eight are skills, awareness or education measures; thirteen are funding, infrastructure, data or intellectual property measures; two are international; and three are explicitly regulatory.[2] That distribution is the policy's signature. It is a human capital and infrastructure plan with a regulatory appendix.

The matrix also contains internal inconsistencies. The text promises access to computing resources for at least 100 academic institutions, while the matrix promises 150.[2] The text promises 200,000 trainees a year, while the matrix promises one million by 2027.[2] Action 8 describes facilitating over 2,000 AI-led patents but specifies 1,000 patents in its desired-state column.[2] The text describes a National AI Fund under Ignite, while the matrix describes a fund established by 2026 with no named vehicle.[2] These are not fatal, but they matter for accountability: a target that appears in two versions cannot be assessed against either.

The critique

The Digital Rights Foundation's analysis of the policy raises four concerns: enforceability, rights protection, oversight and federal-provincial coordination.[51] Each is well founded.

On enforceability, the policy creates no legal duty on any public body, no penalty for non-compliance and no independent body with authority to investigate. Its targets are commitments, not obligations. A ministry that misses every target faces no consequence beyond a question in Parliament.

On rights, the policy mentions privacy and human rights in its objectives, but the action matrix contains no measure to enact a data protection law, create a redress mechanism, require an impact assessment or protect against automated discrimination.[2] The rights language is aspirational. The instruments are absent.

On oversight, the AI Council is chaired by the minister whose ministry is the policy's main implementer. That is a conflict of interest at the apex. The Oversight Committee created by the Digital Nation Act oversees the Pakistan Digital Authority, not the AI program. There is no independent evaluator.

On federalism, the policy assumes that provincial chief secretaries on the AI Council will deliver provincial action. They may, but they are not bound to, and the provinces have their own politics, budgets

and priorities. The policy contains no mechanism for resolving a disagreement between the federal government and a province over an AI standard, a data-sharing request or a procurement condition.

Assessment

The National AI Policy 2025 is a credible industrial strategy and a weak governance instrument. It correctly identifies the supply-side constraints: skills, compute, data, finance and coordination. It sets targets that are easy to communicate and hard to verify. It creates institutions that are not yet operational. It assigns functions to bodies that do not exist. It contains no rights, no remedies and no enforcement.

None of this makes the policy worthless. A country needs a statement of direction, and this one is better than most in its emphasis on use cases, local languages and inclusion. But the policy should not be mistaken for a regulatory regime, and the government should stop presenting it as one. The next section of this report turns to the instruments that will determine whether the policy's promises become practice: data governance, standards and procurement.

CHAPTER

06

Data governance, privacy and WASL

If AI governance in Pakistan is decided anywhere, it is decided in data governance. Every model depends on records. Every automated decision depends on a data flow. Every right that a citizen might have over an AI system depends on a law that tells them what data is held, who can use it and how to challenge it. Pakistan's data governance framework is therefore not a technical appendix to the AI policy. It is the main event.

IN THIS CHAPTER

The draft National Data Governance Policy 2026

The doctrine

Residency and cross-border transfer

Institutions and instruments

AI and automated decisions

WASL

THE ARGUMENT IN ONE LINE

Data governance is where AI rights and vendor dependence will be settled.

The draft National Data Governance Policy 2026

The Ministry of IT published a draft National Data Governance Policy in June 2026.[4] The Ministry's policy register still lists it as a draft with a date of 26 June 2026.[46] The document describes itself as the primary national instrument for data governance, with binding effect on federal public bodies and their contractors.[4] Its metadata identifies it as DNP-D.001 POL, version 1.1, status FD and maturity PROPOSED, with an effective date pending Cabinet approval and Gazette notification.[4]

The policy's scope is important. It applies to federal ministries, divisions, departments, statutory bodies, regulators, authorities, commissions, autonomous bodies and public-sector companies, and to contractors and partners that process government data.[4] It expressly does not govern personal data held outside the public sector, primary legislation, judicial proceedings or national security, defence, parliamentary and judicial regimes, except where expressly provided.[4] Provinces are encouraged to adopt it or an equivalent framework.[4]

That scope makes the policy a public-sector data law, not a national data law. It can regulate how a ministry shares records with a vendor. It cannot regulate how a bank, a hospital chain, a telecom operator or a social media platform handles personal data. For that, Pakistan still needs the Personal Data Protection Act.

The doctrine

The policy is built on a doctrine with five elements: government data as a national asset held in trust; the public body as custodian rather than proprietor; public value as the test of every data decision; sovereignty as the condition of control; and federation as the architecture.[4]

The doctrine is coherent and, in places, sophisticated. It rejects the idea that a department owns its data, while preserving the department's operational custody. It requires a single authoritative source for each data domain, with other bodies consuming from it rather than duplicating it.[4] It adopts an “open by default, closed by exception” rule, a once-only principle and a minimum-disclosure principle.[4] It requires provenance and lineage, privacy by design, zero-trust access and security proportional to classification.[4]

These are the right principles. The question is whether they will be given legal effect, and by whom.

Residency and cross-border transfer

The policy establishes three residency tiers.[4]

Tier 1 covers restricted, confidential and personal data, including all sensitive personal data and authoritative-source data of national importance. It must be hosted, stored and processed within Pakistan. Cross-border processing or remote access requires specific approval and additional safeguards.[4]

Tier 2 covers internal data not in Tier 1, where there is a credible operational case for offshore processing. It may be processed offshore with prior approval, contractual safeguards, technical controls and audit rights.[4]

Tier 3 covers open data, which may be hosted, mirrored and reused globally subject to attribution and licence terms.[4]

The policy also requires public bodies to notify the Authority of foreign requests for government data, to refrain from voluntary disclosure and to comply only as permitted by Pakistani law and mutual legal assistance arrangements.[4] It allows the Authority to recognise adequate jurisdictions and to approve standard contractual clauses, binding institutional rules or certifications where adequacy has not been determined.[4]

The tiered approach is a sensible compromise between localisation and openness, and it is more sophisticated than the binary localisation proposals that have circulated in Pakistan for a decade. But it is a draft. The Tier 1 rule, if enacted, would require many public bodies to renegotiate cloud contracts, move data onshore and restructure their systems. The policy gives them transition periods, but it does not estimate the cost or identify who pays.

Institutions and instruments

The policy creates a National Data Governance Council, chaired by the Authority, with federal, provincial, regulator, custodian and expert members.[4] It requires a National Chief Data Officer inside the Authority and an Agency Chief Data Officer in every public body.[4] It creates a National Data Maturity Index, annual self-assessments, periodic third-party audits, certifications and public reporting of compliance status.[4] It gives the Authority power to issue binding directions, with reasons and an opportunity for representation, and to require corrective action.[4]

It also lists a large set of supporting instruments, including a data classification standard, a metadata and catalogue standard, a data quality standard, a data security standard, a WASL specification, an open data guideline, a residency and cross-border profile, a privacy by design standard, a citizen identity and credentials standard, an AI and automated decision-making profile, a data value and reuse guideline, a retention and disposal standard, a compliance and audit standard, a rights handling standard, a training guideline and a National Data Strategy.[4]

This is an ambitious architecture. It is also a long list of documents that do not yet exist in final form. A standards portal that describes a publication framework is not evidence that the standards have been issued, consulted on, tested or adopted. Chapter 7 examines that distinction in detail.

AI and automated decisions

The most important part of the data policy for AI is clause 14. It establishes a risk-tiered approach to AI and automated decision-making, with higher-rigor controls for high-risk uses, including pre-production registration, model accountability, explainability commensurate with use, drift monitoring and post-deployment review.[4]

Clause 14.4 requires meaningful human review on request for automated decisions with legal or similarly significant effects, transparent contestability and accessible redress.[4] Clause 14.5A requires public bodies operating such systems to publish a description of each system in a public registry maintained by the Authority, including its purpose, the categories of data inputs and the decision logic at a level of abstraction appropriate to security, intellectual property and operational integrity.[4]

If implemented, these provisions would be the most significant AI governance measures in Pakistan. They would give citizens a right to human review, require a register of high-impact systems and force agencies to describe their decision logic. But they are in a draft policy, not a statute or a notified regulation. They depend on an AI and automated decision-making instrument that has not been issued. And the registry they require does not exist.

The policy also requires privacy impact assessments for high-impact processing, including profiling and automated decision-making, and data sharing impact assessments for new exchanges.[4] It requires breach notification to the Authority and, where there is high risk, to affected individuals.[4] These are good provisions. Their value depends on enforcement, and enforcement depends on a regulator with power to inspect, sanction and publish.

WASL

WASL, the National Data Exchange, is the operational core of the data policy. The official description presents it as a means for authorised institutions to exchange information while the original institutions retain their records.[4] The model is federated: identity data remains with NADRA, land data remains with the provincial authority, and WASL routes a governed request and response.

The technical specification adds detail. It identifies a central platform operated by NADRA and owned by the Authority, with functions including access management, routing, consent coordination and audit records.[52] For requests originating from AI applications, it specifies additional identifiers for the model, version and invoking principal, so that a request can be traced to the system and actor responsible.[52] The consent guideline describes itself as advisory and subordinate to the Act and relevant regulations, and states that delegated consent for AI agents acting for citizens is not operationally enabled in WASL version 1.[53]

These are sensible design choices. The model and version identifiers are more than many data exchanges require, and the decision not to enable delegated consent in the first version is prudent. But the public record contains no published figures on how many institutions are connected, how many transactions have been processed, how many consent requests have been made, how many were refused or how many errors were found. Without those figures, WASL is a design, not an infrastructure.

There is also a governance question. WASL is owned by the Authority and operated by NADRA. NADRA is itself a data custodian and a major user of the exchange. That creates a concentration of roles: operator, custodian and user. The policy's answer is audit and separation of duties, but the audit arrangements are not public. A data exchange that moves identity and land records needs an independent inspector, not only an internal one.

The two-regulator problem

The data policy anticipates a Personal Data Protection Law and a separate authority to supervise it.[4] It says the Authority's mandate is public-sector data as a national asset, while the supervision of personal data rights belongs to the authority under the personal data law.[4] It requires the two to coordinate through a memorandum of understanding, to avoid duplicative directions and to reconcile conflicting guidance.[4]

That is a reasonable division in theory. In practice it creates a risk of two regulators with overlapping jurisdiction over the same records, different cultures and no common appellate route. A citizen complaining about a public hospital's use of their data would have to know whether the issue is a data governance matter, a personal data rights matter or both. The policy does not say.

The risk is compounded by the fact that the personal data authority does not exist. Until the bill is enacted, the Authority is the only data regulator in town, and it has no statutory power over private-sector personal data. That is a vacuum, not a division of labour.

What is missing

Six gaps stand out.

First, the policy has no statutory basis. It is a policy issued by a ministry, not a law passed by Parliament. Its binding effect on public bodies depends on the Digital Nation Act's delegation of powers to the Authority, and that delegation is itself contingent on assignment by the Commission.[1]

Second, it has no independent appeal. The Authority can issue binding directions and require corrective action, but the policy does not create an independent tribunal or ombudsman. A public body that disagrees can seek representation, but a citizen cannot appeal.

Third, it has no penalty regime. The policy refers to enforcement under a compliance instrument that has not been issued. Without penalties, compliance depends on goodwill and ministerial attention.

Fourth, it excludes the private sector. Most personal data in Pakistan is held by banks, telecoms, hospitals, schools, retailers and platforms. The policy does not reach them. The Personal Data Protection Bill would, but it is not law.

Fifth, provincial adoption is voluntary. Health, education, land and police data are provincial. A federal data policy that applies only to federal bodies covers a minority of the records that matter for AI in daily life.[4]

Sixth, parliamentary scrutiny is absent. The policy was published for consultation, but it has not been debated in Parliament, referred to a standing committee or subjected to a regulatory impact assessment. For an instrument that would reshape how the state handles data, that is a significant democratic deficit.

Assessment

The draft data policy is the most sophisticated document in Pakistan's AI governance stack. It understands the issues, borrows intelligently from other jurisdictions and sets out a coherent operating model. It is also, at present, a design for a building that has not been constructed. Its rights depend on a law that has not passed. Its institutions depend on notifications that have not been issued. Its standards depend on documents that have not been finalised. Its enforcement depends on powers that have not been tested.

The policy should be finalised, but not as a substitute for legislation. The right sequence is a Personal Data Protection Act, a statutory basis for the data governance regime, a published set of final instruments, a public registry of high-impact systems, and an independent appeal route. The draft contains the architecture for all of these. What it lacks is the law that would make them real.

CHAPTER

07

Standards, risk classification and assurance

In a country without an AI statute, standards carry unusual weight. They are the only instrument through which a government can specify how a system should be built, tested, documented and monitored without passing a law. They are also easy to mistake for regulation. This chapter explains what Pakistan's standards architecture contains, what it does not, and why the difference matters.

IN THIS CHAPTER

The publication framework

What exists

Framework, regulation, standard, guideline

International standards and their uses

The assurance chain

Why conformance matters more than publication

THE ARGUMENT IN ONE LINE

The standards are intelligently designed and barely populated. Assurance is the gap.

The publication framework

The Pakistan Digital Authority has published a framework for its own instruments, DNP-X.001.[54] It assigns identifiers by series: DNP-A for artificial intelligence, DNP-D for data, and other series for identity, cloud and related subjects. It distinguishes frameworks, policies, regulations, standards, guidelines, technical specifications, reference architectures, profiles and strategies.[54] It also distinguishes working drafts, public comment drafts and published documents, and separates Gazette references from operational document identifiers.[54]

The framework is a genuine contribution. Most governments publish standards as a flat list, with no indication of status, version or legal basis. A reader who finds a PDF can rarely tell whether it is a draft, a proposal, a guidance note or a binding rule. Pakistan's framework tries to solve that problem, and its metadata blocks, with identifiers, status codes, maturity levels, approval dates and review dates, are more transparent than most.[4]

But a framework for publishing documents is not a substitute for the documents. The framework describes how a standard would be identified if it existed. It does not establish that the standard has been issued, consulted on, tested or adopted.

What exists

The public record contains the following instruments in the data and AI series.

The National Data Governance Policy 2026 is DNP-D.001 POL, version 1.1, status FD, maturity PROPOSED.[4] It is a draft awaiting Cabinet approval and Gazette notification.

The WASL specification is DNP-D.300 TS.[52] It is a technical specification for the data exchange, and it is the most detailed operational document in the series.

The WASL consent guideline is DNP-D.250 GDL.[53] It is advisory and subordinate to the Act and regulations, and it states that delegated consent for AI agents is not enabled in version 1.

The AI, automated decision-making and emerging technology profile is DNP-D.401 PRF.[4] It is listed as a required supporting instrument in the data policy's annex, with higher-rigor controls for AI data use, model accountability, explainability, drift, pre-production registration and generative AI.[4] The report could not locate a published version of this profile in the public record.

The AI series, DNP-A, is identified in the framework but the report could not locate published AI-specific standards, a risk classification scheme or an impact assessment methodology in the public record. That absence is significant. The data policy requires risk-tiered controls, pre-production registration and a public registry.[4] None of these can operate without a classification scheme and an assessment method. If those documents exist, they are not public. If they do not exist, the policy's AI provisions cannot be implemented.

Framework, regulation, standard, guideline

The distinction between these instruments is not academic. It determines legal effect, accountability and the route of challenge.

A framework explains an approach. It is useful for orientation and has no binding force.

A policy states the government's position and may bind public bodies through administrative direction. It is not enforceable by a citizen in court unless a statute gives it effect.

A regulation is made under a statutory power, is notified in the Gazette and is enforceable. It can be challenged in court. Its validity depends on the scope of the enabling power.

A standard specifies technical requirements. It may be voluntary or mandatory. It becomes binding when a regulation, a contract or a licence condition incorporates it.

A guideline recommends practice. It is not binding, but it may be evidence of the standard of care in a negligence claim.

A technical specification describes how a system should be built or interfaced. It is binding on parties that adopt it, and it may become binding through procurement.

A reference architecture describes a design pattern. It is not a requirement.

Pakistan's AI stack is heavy on frameworks, policies, guidelines and specifications, and light on regulations and binding standards. That is the pattern of a government that wants to signal direction without creating legal exposure. It is understandable. It is also insufficient for high-impact systems.

International standards and their uses

Three international instruments are relevant to Pakistan's choices.

The NIST AI Risk Management Framework organises risk work into four functions: govern, map, measure and manage.[25] It treats risk management as continuous work across the lifecycle, including testing before deployment and monitoring during operation.[25] Its generative AI profile addresses risks specific to generated content and the wider system in which a model is used.[55] The framework is voluntary and was written for a US context, but its vocabulary is widely used and its structure is adaptable.

ISO/IEC 42001 specifies requirements for an AI management system.[26] It is certifiable, which makes it useful for procurement: a supplier that holds the certificate has been audited against a recognised standard. Certification is not a guarantee of safety, but it is a verifiable signal.

The European Union's AI Act establishes a risk-based regime with prohibited practices, high-risk obligations, transparency duties for general-purpose models and a conformity assessment process.[24] It is the most comprehensive AI law in the world and is being observed closely by regulators in Asia. It is not a template that Pakistan can adopt wholesale, because its institutional requirements, market surveillance system and enforcement budget are far beyond Pakistan's current capacity.

The UNESCO Recommendation on the Ethics of AI and the OECD AI Principles provide the normative baseline that Pakistan's own declaration echoes.[56,23] They are useful for aligning language and for international cooperation. They do not create domestic obligations.

The assurance chain

Assurance is the process by which a claim about a system is tested and verified. A complete chain has seven links:

- A description of the system, its purpose, its users and its boundaries.
- A risk assessment that identifies affected people, potential harms and mitigations.
- A data assessment that covers provenance, quality, representativeness and lawful basis.
- A performance evaluation that tests accuracy, robustness, bias and failure modes in the intended setting.
- An independent review or audit, by a party that did not build the system.
- A deployment decision with conditions, monitoring and incident reporting.
- A periodic re-assessment, including drift, incidents and changes to the system or its context.

Pakistan's public record contains evidence of the first link for some systems, in the form of project descriptions and procurement notices. It contains very little evidence of the other six. There is no published AI audit, no certification body accredited for AI management systems, no testing laboratory, no register of incidents, no post-deployment evaluation and no enforcement decision.

That is the assurance gap. It is more important than the standards gap, because a country can borrow standards. It cannot borrow the institutions that apply them.

Why conformance matters more than publication

The data policy's conformity assessment model is sensible on paper: annual self-assessment by every public body, periodic third-party audits of a sample, certification where an instrument provides for it, and public reporting of compliance status.[4] It is the right model. But it depends on three things that do not yet exist: a cadre of trained assessors, an accreditation body to certify them and a published methodology to apply.

Pakistan has strong professional bodies in accountancy, engineering and law. It has universities with computer science and data science departments. It has a nascent AI research community. It does not have an AI assurance profession. Building one would take years and would require a curriculum, a code of ethics, a certification scheme and a market for the service. The alternative, relying on foreign auditors, is expensive and may not be available for sensitive government systems.

The government should treat assurance capacity as infrastructure. It is as important as compute and more important than another strategy document. A small, well-trained cadre of assessors, operating under a published methodology and reporting to the Authority, would do more for AI governance than a dozen new guidelines.

Assessment

Pakistan's standards architecture is intelligently designed and thinly populated. The publication framework is better than most. The WASL specification is serious technical work. The data policy sets out a coherent conformity model. But the AI-specific standards, the risk classification, the impact assessment method and the registry do not exist in public form, and the assurance institutions that would apply them have not been built.

The next phase should publish the AI risk classification and impact assessment instruments, adopt ISO/IEC 42001 as the reference management standard, create an assessor program, and require conformance for high-impact systems in procurement. Without those, Pakistan will have a standards portal, not a standards system.

CHAPTER

08

Government deployment and procurement

Most AI governance in Pakistan will be exercised through procurement. When a ministry buys a chatbot, a document-processing tool, a facial recognition system or an analytics platform, the contract determines what data the vendor can use, how the system is tested, who can audit it, what happens when it fails and whether the government can leave. A statute sets the boundaries. A contract sets the reality.

IN THIS CHAPTER

The deployment picture

The procurement rules

What an AI contract should contain

The AI hubs procurement

The civil service

Vendor concentration and the sovereignty question

THE ARGUMENT IN ONE LINE

Procurement is the de facto regulator, and it is not performing the role.

A PROPOSED ACCOUNTABILITY PATH FOR PROCUREMENT

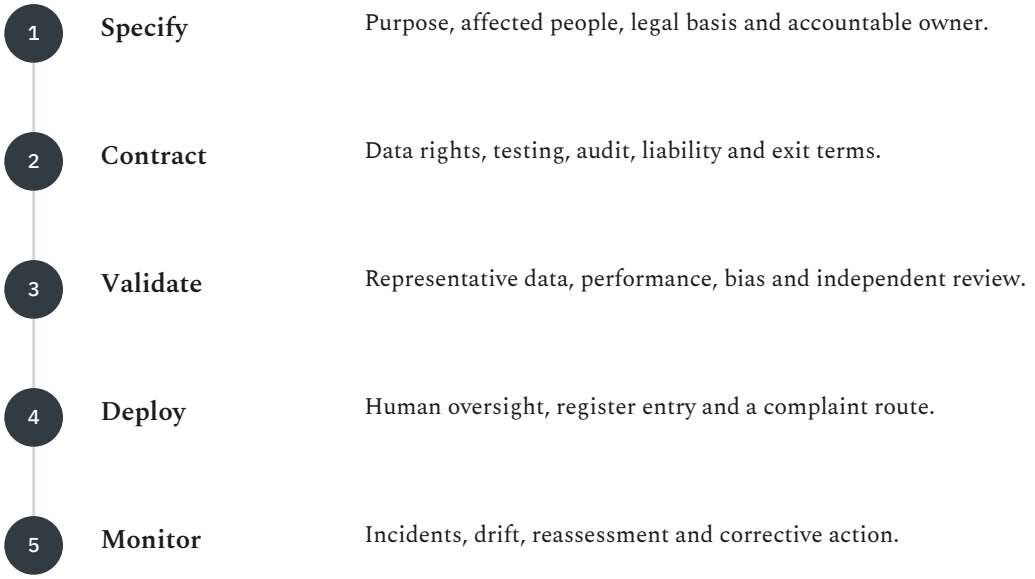


Figure 6. From procurement to continuing accountability. Source: PIVRA synthesis of Chapter 8 recommendations and the assurance chain in Chapter 7. This is a proposed pathway, not a description of current compliance.

This chapter examines how government buys and uses AI, what the procurement rules require, and where the gaps are.

The deployment picture

Pakistan has no central inventory of government AI systems. That absence is itself a finding. The data policy would require a public registry of automated decision systems with legal or similarly significant effects, but the registry does not exist.[4] The AI Policy's KPIs include the percentage of industries adopting AI, but not a count of government systems.[2]

What can be observed is a pattern of pilots and platforms. Federal agencies have deployed chatbots for citizen information, document processing for files and applications, analytics for tax and customs risk, and identity verification tools. Provincial governments have been more experimental, particularly Punjab, which has an AI roadmap, an AI delivery unit and a set of service projects.[14] Khyber Pakhtunkhwa has a government language model and a digital strategy.[15] Sindh has signed a partnership with Huawei covering smart cities and government technology infrastructure.[16]

These are announcements and program documents, not evaluation records. The public record does not show how many systems are in production, what they cost, how they perform, what data they use, whether they have been assessed or what happens when they produce a wrong result. A government that cannot answer those questions about its own systems cannot govern them.

The procurement rules

Pakistan's federal procurement framework changed in September 2026, when the Public Procurement Rules 2026 were notified.[28] The rules address evaluation, records, contract management and public access to award documents, subject to specified provisions.[28] They preserve the 2004 rules for proceedings initiated before commencement, which means that an AI procurement started before the change is governed by the old rules.[28]

The rules are an improvement on their predecessor. Public access to award documents is a meaningful transparency measure, and contract management requirements give agencies a basis for holding suppliers to delivery. But the rules are generic. They were not written for AI, and they do not require the things that AI procurement most needs: a data protection impact assessment, a model documentation package, testing on representative data, bias evaluation, human review arrangements, incident reporting, audit rights, restrictions on training the vendor's model on government data, and an exit plan that includes data and model portability.

Those terms have to be written into the contract. If they are not, the supplier's standard terms will govern, and the standard terms of most AI vendors are designed to protect the vendor.

What an AI contract should contain

A model AI contract for a public body should contain at least the following.

A clear description of the system, its purpose, its users, its boundaries and the decisions it supports. The description should distinguish between the system's output and the decision that follows, because the public body, not the vendor, remains responsible for the decision.

A data schedule that identifies the data to be used, its source, its lawful basis, its retention period, its location and the persons who can access it. The schedule should prohibit the vendor from using government data to train or improve its own models without explicit written consent.

A performance schedule that specifies the metrics, the test data, the acceptance thresholds, the evaluation method and the consequences of failure. For a classifier, the schedule should require disaggregated performance across relevant groups. For a document system, it should require retrieval quality, unsupported-answer rates and source attribution.

An assurance schedule that gives the public body the right to audit the system, to commission an independent evaluation, to receive model documentation and to inspect subprocessors. It should require incident notification within a defined period.

A human oversight schedule that specifies which decisions require human review, who conducts it, what information they see and how the affected person can request it.

A security schedule that covers access control, encryption, logging, vulnerability management, incident response and data breach notification.

An exit schedule that requires the return or destruction of data, the migration of configurations and embeddings, the continuation of services for a defined period and the transfer of know-how. Without it, the public body is locked in.

A liability schedule that allocates responsibility for harm, including harm caused by model error, data breach and discriminatory output. Vendors will resist, but a public body that accepts unlimited liability for a vendor's model is not procuring, it is indemnifying.

Most Pakistani AI contracts are not public, so it is not possible to say how many contain these terms. The absence of standard clauses, model contracts or guidance from the Authority or the Public Procurement Regulatory Authority suggests that practice varies widely.

The AI hubs procurement

The Ministry's tender for national AI hubs illustrates the kind of primary record that is available.[47] It identifies a ministry procurement and a reference number for selecting a firm.[48] A separate capacity-building corrigendum extended a submission deadline to 1 June 2026.[48]

Those records document procurement steps. They do not show contract award, payment, completed facilities or delivered training. A reader who wants to know whether the hubs exist must look for an award notice, a completion certificate, a site visit report or a budget execution figure. None of those is in the public record reviewed for this report.

The hubs are important because they are the physical layer of the AI policy. If they are built, they could provide compute, training and testing capacity to universities and firms that cannot afford their own. If they are not, the policy's infrastructure pillar is a slide.

The civil service

The civil service is where AI meets the citizen most directly. The National CERT's Cybersecurity Handbook 2026-27, aimed at federal, provincial and public-sector personnel, is reported to instruct staff to use department-approved tools, avoid entering classified material and citizens' personal information into public AI platforms, protect credentials and review generated outputs.[44] The reported instructions also address accidental disclosures and unapproved extensions.

These are sensible controls. They are also conduct rules, not governance. They tell an employee what not to paste into a chatbot. They do not tell the department what it must do before deploying a system that affects a citizen's rights, how it must test the system, or how it must record the decision.

The Civil Services Academy introduced an AI module in January 2026.[34] That is a useful step. But a one-off module does not create the capability to write a data schedule, evaluate a model or audit a vendor. Those are specialist skills, and the government will need to hire or contract for them.

Vendor concentration and the sovereignty question

Pakistan's AI market is small, and the number of vendors that can deliver a large government system is smaller still. That gives suppliers leverage. It also creates a sovereignty problem: if the model, the cloud and the integration are all provided by one foreign company, the government's ability to switch, audit or continue operating in a crisis is limited.

The government's response has been to build state-owned capacity through the National Telecommunication Corporation and to encourage domestic cloud providers.[21] That is a reasonable start. But sovereign infrastructure is only useful if it is competitive on price, performance and support. A state data centre that cannot match the latency, tooling and reliability of a hyperscaler will not attract the workloads it needs to be viable.

The alternative to full sovereignty is not dependence. It is a portfolio: some workloads on sovereign infrastructure, some on trusted foreign clouds under contract, and some on a hybrid arrangement with data residency and audit rights. That requires a classification of workloads and a procurement capability that Pakistan does not yet have.

Assessment

Procurement is the de facto AI regulator in Pakistan, and it is not performing the role. The rules are generic, the contracts are opaque, the standard clauses do not exist and the inventory does not exist. The result is that the most consequential governance decisions, about data access, vendor lock-in, testing, audit and redress, are made project by project, by officials who may not have the expertise or the leverage to insist on them.

The remedy is not a new law. It is a set of standard AI contract clauses, issued by the Authority and the Public Procurement Regulatory Authority, mandatory for federal procurements above a threshold; a requirement for a data protection impact assessment and a model documentation package before award; a public register of AI systems above a risk threshold; and a small central team that can advise agencies on difficult procurements. Those measures would do more for AI governance in the next two years than any new strategy.

CHAPTER

09

Sectoral regulation: finance, telecoms, courts, universities, health and markets

Pakistan does not have a single AI regulator. It has a set of sector regulators, each with its own statute, culture and capacity, encountering AI at different speeds. This chapter examines the most important sectors. The pattern that emerges is consistent: the sectors with strong supervisors are beginning to act, the sectors with weak or divided supervision are not, and the citizen's protection depends on which sector they happen to be in.

IN THIS CHAPTER

Financial services

Telecommunications

The courts

Higher education

Health

Competition and consumer protection

THE ARGUMENT IN ONE LINE

A citizen's protection depends on which sector they happen to be standing in.

Financial services

The State Bank of Pakistan is the most advanced AI supervisor in the country, and the bar is not high.

Its Financial Stability Review 2024 reported a survey of 55 regulated entities. About half had deployed AI or were developing it, with applications in fraud detection, customer service, marketing, credit-risk assessment and process automation.[10] The same review described responsible-AI guidelines as being finalised at the time.[10] This report found no published issuance of those guidelines in the public record as of November 2026. It is possible that they were issued in a circular that the report did not locate, but the absence of a public document is itself a governance problem: banks cannot comply with guidance they cannot find.

The State Bank's existing rules already reach some AI risks. Its framework on outsourcing to cloud service providers requires approval for offshore hosting of material workloads, due diligence on providers, audit rights and exit arrangements.[42] Its rules on customer confidentiality restrict disclosure of account information.[42] Its outsourcing framework addresses personally identifiable information and cross-border arrangements.[42] These are not AI-specific, but they cover the infrastructure on which AI runs.

What is missing is model risk management. Banks increasingly use machine learning for credit scoring, anti-money laundering alerts, fraud detection and customer segmentation. A model that is wrong in a systematic way can deny credit to a group, miss a fraud pattern or generate thousands of false alerts that exhaust the compliance team. Prudential regulators in other jurisdictions require model inventories, validation, documentation, ongoing monitoring and board accountability. The State Bank has not published an AI-specific model risk standard. It should.

The Securities and Exchange Commission regulates capital markets, insurance and non-bank finance. Its regulatory sandbox provides a controlled environment for limited live testing over a defined period, particularly where an innovation does not fit existing arrangements or its effects are uncertain.[11] Its first cohort admitted six solutions for testing.[50] The sandbox is a useful learning mechanism. It is not a substitute for rules, and the SECP has not published AI-specific guidance for capital markets, asset managers or insurers. Its 2020 cybersecurity guidelines for insurers address cyber risk, not AI.[6]

The financial sector's AI governance gap is therefore a gap in the middle: infrastructure rules exist, sandbox exists, but the model-level standards that would make AI safe in credit, insurance and trading do not.

Telecommunications

The Pakistan Telecommunication Authority regulates the pipe, not the model. Its mandate covers licensing, spectrum, quality of service, tariffs, interconnection, content blocking and lawful interception. AI enters its domain in three ways.

First, telecom operators use AI for network optimisation, predictive maintenance, customer care and fraud detection. The PTA does not publish adoption data or AI-specific requirements for operators. Its quality-of-service rules apply to outcomes, not methods.

Second, AI is used to manage content. The PTA is the technical agency through which blocking orders are implemented, and the 2025 amendments created a Social Media Protection and Regulatory Authority with content removal powers.[40] The interaction between the two is unclear, and the new authority is not yet operational.[6]

Third, AI is used in lawful interception and metadata analysis. The legal framework for interception is spread across the telecom law, PECA and the security agencies' own rules. It is not transparent, and the safeguards are not published.

The PTA's most important AI decision may be about spectrum and 5G. AI workloads depend on high-capacity, low-latency connectivity, and Pakistan's 5G rollout has been slow: about 1.3 million users in 23 cities by late 2026.[18] The pace of rollout affects the viability of edge AI, smart city applications and industrial automation. It is a governance decision, even if it is not framed as one.

The courts

The National Judicial Policy Making Committee issued National Guidelines for the Use of Artificial Intelligence in Judicial Institutions on 29 April 2026, following approval at its 57th meeting.[12] The guidelines describe AI as an aid to legal research, case management, document processing and administration, while preserving human judgment and judicial independence.[12] They cover transparency, accountability, bias, sensitive information and training.[12] They recognise the constitutional and administrative autonomy of the High Courts, leaving implementation to the relevant jurisdictions.[12]

The guidelines are a serious document. They correctly distinguish between AI as an aid and AI as a decision-maker, and they preserve judicial authority. But they are guidance, not law, and their implementation depends on each High Court. There is no published evidence that any court has deployed a system, adopted a policy, trained its staff or audited a tool. There is no disclosure requirement for litigants, no route to challenge an AI-assisted decision and no register of tools in use.

The risk is not that a judge will hand down a sentence written by a chatbot. It is that AI-assisted research, summarisation and case management will introduce errors that are hard to detect, that litigants will not know a tool was used, and that the efficiency gains will be captured by courts while the risks are borne by the people appearing before them. The guidelines acknowledge these risks. They do not yet manage them.

Higher education

The Higher Education Commission hosts a framework on generative AI in higher education covering disclosure, consent, privacy, academic integrity and institutional responsibilities.[13] Its cover calls it a draft, version 1.0, despite the word Final in the file name, a detail that illustrates why document status must be read from the document and not from its URL.[13]

The framework envisages university and departmental processes for approval, oversight and handling misuse.[13] That is the right level: academic judgement about what AI use is permitted belongs to the institution and the teacher, not to a national regulator. But the framework is a draft, and universities have adopted it unevenly. Some have detailed policies; others have none. Students in different institutions face different rules for the same conduct.

The capacity picture is mixed. HEC reported a faculty program on 3 September 2026 with 32 participants attending in person and more than 300 joining through university facilities.[33] That is a useful start, but it is a fraction of the faculty who need to understand AI in assessment, research and administration. The bigger gap is in AI literacy for students outside computer science, who will use AI in every profession and are rarely taught how to evaluate its outputs.

Health

Health is the sector with the highest potential benefit and the weakest governance.

AI applications in Pakistani health include radiology and imaging, tuberculosis screening, maternal and child health, telemedicine triage and hospital administration. Some are piloted by provincial health departments, some by private hospitals, some by donors and non-governmental organisations. There is no national inventory, no evaluation standard, no clinical validation requirement specific to AI and no guidance on liability when a tool contributes to a misdiagnosis.

The Drug Regulatory Authority of Pakistan regulates medical devices and, increasingly, software as a medical device. Its capacity to assess AI-enabled devices is limited, and its processes are not transparent. Provincial health departments own most service delivery and have no AI-specific units. The Pakistan Medical and Dental Council regulates medical practice, not tools.

The result is a sector where innovation is outpacing oversight. A hospital can procure an AI triage tool without a published clinical evaluation, deploy it without telling patients, and have no clear route for a patient harmed by its output. That is not a hypothetical: it is the default in the absence of rules. The remedy is a national health AI assessment pathway, a registry of approved tools, a requirement for clinical validation on local data and a clear liability rule. None of that requires a new statute. It requires the health ministry and the regulator to act.

Competition and consumer protection

The Competition Commission of Pakistan has a mandate over anti-competitive agreements, abuse of dominance and mergers. AI raises familiar and unfamiliar questions: algorithmic pricing that facilitates tacit collusion, self-preferencing by platforms that both operate a marketplace and sell on it, mergers that concentrate data, and access to compute or model distribution.

The Commission has not published AI-specific guidance. Its merger review powers could be used to examine data concentration, but the threshold and methodology are not adapted to digital markets. On the consumer side, provincial consumer protection laws cover misleading claims and defective goods, but they were not written for generated content, automated recommendations or algorithmic pricing. A consumer who is harmed by a chatbot's false statement may have a claim, but proving reliance and causation is difficult.

Media and content

Pakistan's media sector is regulated by the PTA, the Pakistan Electronic Media Regulatory Authority for broadcast, and now the Social Media Protection and Regulatory Authority for online content.[40] The 2025 amendments created new offences for false information likely to cause fear or unrest, and gave the new authority power to order removal or blocking.[40]

AI enters this domain through deepfakes, synthetic audio and video, coordinated inauthentic behaviour and algorithmic amplification. The legal tools are blunt: criminal offences after the harm, and blocking orders that may remove legitimate content along with the harmful. There is no provenance standard, no labelling requirement, no rapid redress route for a person whose image is synthesised, and no transparency requirement for platform recommendation systems.

The risk to governance is asymmetric. A deepfake of a politician or a judge can spread faster than a correction, and the legal process takes months. The absence of a rapid, proportionate remedy is a gap that will be exploited.

Labour

Pakistan has no AI-specific labour rules. The Factories Act, provincial industrial relations statutes and social security laws were written for a different economy. AI affects labour through task automation, algorithmic management, gig work and surveillance. None of these is addressed.

The most immediate issue is algorithmic management of gig workers, who are classified as independent contractors and have little protection against opaque ratings, deactivation or wage setting. The second is displacement in clerical, customer service and data-processing jobs, which employ large numbers of young Pakistanis. The third is workplace surveillance, which AI makes cheaper and more pervasive.

The policy response so far is skills training, which is necessary but not sufficient. It does not address the workers who lose their jobs, the conditions of those who remain or the bargaining power of those who are managed by algorithms. A national AI adjustment strategy, with income support, retraining entitlements and portable benefits, would be a more serious response. It does not exist.

Table 6. Sectoral AI governance at a glance

Sector	Lead regulator	AI-specific instrument	Status	Main gap
Banking and payments	State Bank	Responsible AI guidelines reported in 2024	Not published in public record	Model risk standard
Securities and insurance	SECP	Regulatory sandbox; cyber guidelines	Operational sandbox	AI guidance for markets and insurers
Telecoms	PTA	None specific	Not applicable	AI in networks, content and interception
Courts	NJPMC and High Courts	National guidelines, April 2026	Guidance only	Court-level policies, disclosure, audit
Higher education	HEC	Generative AI framework, draft v1.0	Draft	Binding policy, student AI literacy
Health	DRAP and provinces	None specific	Absent	Clinical validation, registry, liability
Competition	CCP	None	Absent	Algorithmic collusion, data concentration
Media and content	PTA, PEMRA, SMPRA	Content offences and blocking powers	Partly operational	Deepfake redress, provenance, transparency
Labour	Provinces	None	Absent	Algorithmic management, adjustment

The pattern

Three conclusions follow from the sectoral review.

First, AI governance in Pakistan is sectoral by default, and the sectors are unevenly equipped. A bank customer has a supervisor with a survey, a sandbox and cloud rules. A patient has a regulator with limited AI capacity. A gig worker has no relevant regulator at all.

Second, the strongest instruments are indirect. The State Bank's cloud outsourcing rules, the SECP's sandbox and the PTA's spectrum decisions affect AI more than any AI-specific guidance. Governance through infrastructure is often more effective than governance through principles.

Third, the sectors that have issued guidance have not built enforcement. The courts, the universities and the financial sector all have documents. None has a published assessment, audit or enforcement decision. The gap between guidance and practice is the same gap that runs through the national architecture.

CHAPTER

10

Cybercrime, disinformation and online harms

AI's most visible harms in Pakistan are not abstract. They are fraud, impersonation, non-consensual imagery, disinformation and online harassment. These harms are governed by criminal law, content regulation and platform policy, not by AI law. This chapter examines whether those tools are adequate.

IN THIS CHAPTER

The threat picture

The legal framework

Deepfakes and impersonation

Disinformation

Election integrity

Assessment

THE ARGUMENT IN ONE LINE

Criminal law is slow. Prevention, provenance and rapid redress are missing.

The threat picture

Pakistan's digital threat environment has three features that AI amplifies.

The first is scale. A country with 204 million mobile subscriptions and a young, mobile-first population is a large market for scams, phishing and fraudulent investment schemes. Generative tools lower the cost of producing convincing messages, voices and images in Urdu and regional languages.

The second is trust. Public institutions are not universally trusted, and information about politics, religion and security is highly contested. A synthetic video of a public figure, a fake notification from a bank or a fabricated court order can spread quickly and be believed by people who have limited access to verification.

The third is weak redress. The cybercrime investigation agency is centralized, its capacity is limited relative to the volume of complaints, and most offences are non-cognizable, bailable and compoundable.[40] A victim of online fraud or impersonation may wait months for a response, and the platform on which the harm occurred may be outside Pakistani jurisdiction.

The legal framework

The Prevention of Electronic Crimes Act 2016 is the core statute. It criminalises unauthorised access, copying, interference, electronic forgery, electronic fraud, unauthorised use of identity information and disclosure of personal information by service providers.[40] The 2025 amendment created the National Cyber Crime Investigation Agency, added the false-information offence and established the content authority.[40]

The National Cyber Crime Investigation Agency's public guidance connects online harassment protections with PECA and directs users to its complaint service.[45] That is a useful route for some harms. It is not a route for a person who wants to challenge an automated decision, correct a data error or find out whether a government system used their data.

The National CERT's Cybersecurity Handbook 2026-27 sets out baseline security practices for federal, provincial and public-sector personnel, including rules on AI tool use, classified material, personal information and review of generated outputs.[44] The handbook is a valuable operational document. It is not a regulatory instrument, and it does not apply to the private sector.

Deepfakes and impersonation

Pakistan has no deepfake-specific offence. A synthetic video may be prosecuted under PECA's forgery, fraud, identity or harassment provisions, depending on the facts.[40] A civil claim may lie in defamation or privacy. Neither route is quick, and both require the victim to identify the perpetrator, which is often impossible.

The absence of a rapid remedy is the critical gap. A person whose image is used in non-consensual synthetic pornography needs removal within hours, not a criminal trial within months. Platforms have their own reporting mechanisms, but they are inconsistent, slow and not accountable to Pakistani regulators. A statutory takedown route, with a court order or an administrative direction and a right of appeal, would be more effective than a new criminal offence.

Provenance is the other half of the solution. Content credentials, watermarking and signed metadata can help platforms and users identify synthetic media, but they are not yet standard in Pakistani media or government communication. The government could require provenance disclosure for official content and encourage platforms to preserve and display provenance signals. It has not.

Disinformation

The 2025 false-information offence has attracted criticism from journalists and rights groups, who argue that it can be used to criminalise legitimate reporting and political criticism.[40] The concern is well founded. The offence turns on knowledge or reason to believe that information is false and likely to cause fear, panic or unrest. Those are elastic concepts, and the enforcement record will determine whether the provision targets harmful disinformation or legitimate dissent.

The content authority adds another layer. Its powers to order removal or blocking are broad, its rules are not final and its tribunals are not established.[6] Until they are, content disputes continue to run through the PTA and the courts, with limited transparency and no consistent appeal route.

A better approach would combine three elements: a narrow, clearly defined offence for coordinated inauthentic behaviour and synthetic impersonation; a transparent takedown process with reasons and appeal; and support for independent fact-checking and media literacy. Pakistan has the first element in over-broad form, the second in fragmented form and the third hardly at all.

Election integrity

AI-generated content is a risk to elections, referendums and public consultations. Pakistan's next general election is not due until 2029, but by-elections, local elections and political campaigns will use AI-generated content before then. There are no rules on disclosure of AI-generated political advertising, no requirement for platforms to maintain ad libraries accessible to Pakistani researchers and no standing mechanism for rapid response to synthetic content.

The Election Commission of Pakistan has a mandate over campaign conduct but has not issued AI-specific guidance. The PTA has content powers but is not an election regulator. The gap is institutional, not technical.

Assessment

Pakistan's response to AI-enabled harm is heavy on criminal law and light on prevention, redress and transparency. The criminal route is necessary for serious conduct, but it is slow, centralized and ill-suited to the volume and speed of synthetic content. The content authority is not yet operational. The platforms are not accountable to Pakistani regulators in any meaningful way.

The priority measures are practical: a rapid takedown route for non-consensual synthetic imagery, with judicial or administrative oversight; provenance requirements for official content; a transparent and appealable content process; a narrow offence for coordinated inauthentic behaviour; and a national digital literacy program that reaches beyond English-speaking urban users. None of these requires a comprehensive AI act. All of them require institutions that can act quickly and explain what they did.

CHAPTER

11

Provinces and federalism

Pakistan's AI governance is federal in form and provincial in consequence. The federal government sets policy, funds programs and coordinates. The provinces deliver the services that AI will most directly affect: health, education, police, land records, local government, agriculture and social protection. A national AI strategy that does not reach the provinces is a strategy for the federal capital.

IN THIS CHAPTER

Punjab

Khyber Pakhtunkhwa

Sindh

Balochistan

The federal-provincial mechanism

What a federal settlement would look like

THE ARGUMENT IN ONE LINE

Provinces are the implementation frontier, and the federation has no common floor.

The 18th Amendment devolved most of these subjects in 2010. The federal government retains IT and telecommunications, but it cannot direct a province to adopt a model, a register or a procurement condition. It can offer funding, standards and coordination. It can also create a national framework that provinces may choose to join. It cannot legislate for them.

Punjab

Punjab is the most active province on AI, and the most consequential.

On 12 March 2026 the Chief Minister approved the province's first AI Roadmap, setting a 2029 deadline for Punjab to become the leading AI province of South Asia.[14] The meeting approved an AI Technical Advisory Board and directed the inclusion of high-level experts.[14] The roadmap includes a manifesto to be launched in March and an AI Delivery Unit whose international launch was planned for June 2026.[14]

The language is ambitious, and the institutional design is more interesting than the headline. An AI Delivery Unit inside the chief minister's office is a model used in other countries to drive cross-departmental projects with central authority. If it has a mandate, a budget and technical staff, it can do what a line department cannot: set standards, run pilots, measure results and stop projects that fail. If it is a coordination forum without authority, it will produce meetings.

Punjab also has the largest digital program. The World Bank approved 70 million dollars for Connected Punjab on 30 June 2026, combining broadband, digital services, computing infrastructure and payments, as part of a broader 278 million dollar government programme including 208 million dollars in provincial counterpart funding.[57] Those are programme financing figures covering several components. They cannot be described as an AI-only budget, but they show the scale of provincial digital investment.

The Punjab Information Technology Board has separately tendered for a consultancy to establish a government AI cloud.[49] That is a significant decision: a provincial cloud for AI workloads could reduce dependence on federal infrastructure and on foreign hyperscalers, or it could duplicate capacity that already exists. The business case is not public.

Khyber Pakhtunkhwa

Khyber Pakhtunkhwa published a Digital Transformation Strategy 2030 that sets out proposed AI governance measures, local-language models and a government-focused language model, KP GovGPT.[15] The strategy is more detailed than most provincial documents and recognises the province's language diversity, which is both a design requirement and a political one.

KP's challenge is implementation capacity. The province has a capable IT board and a growing tech sector in Peshawar, but its fiscal space is limited and its security environment is difficult. A government language model is a useful symbol of sovereignty, but its value depends on whether it is used in service delivery and whether it performs well in Pashto, Hindko and other local languages.

Sindh

Sindh has taken a different route. In 2026 it signed a memorandum of understanding with Huawei to develop an AI-native digital province, covering smart cities and government technology infrastructure.[16] The partnership is significant because it brings a major Chinese vendor into the province's digital stack at a time when the federal government is also deepening its relationship with China on AI.

The advantage is speed and finance. The risk is dependence: a single vendor providing infrastructure, platforms and applications can lock a province into a technology stack that is difficult to leave, and can create data access arrangements that are not transparent. The memorandum's terms have not been published.

Karachi also has the country's largest concentration of technology firms, financial institutions and data centres. Its provincial government has been less active on AI policy than Punjab, but its private sector is more developed. That creates a different governance challenge: the province's AI use is led by firms, not by the state, and the provincial regulator's capacity to supervise it is limited.

Balochistan

Balochistan is the least connected province and the least advanced on AI. It has a small IT sector, limited broadband and pressing security and development challenges. In November 2026 the federal government reported the completion of a modern IT hub in Dera Bugti under the Balochistan Social Development Initiative.[58] That is a welcome investment, but a single hub does not create an AI ecosystem.

The province's priority is not an AI strategy. It is connectivity, electricity, education and basic digital services. A national AI policy that treats Balochistan as a site for a centre of excellence, without addressing those foundations, will produce a building with no users.

The federal-provincial mechanism

The Constitution provides for the Council of Common Interests to resolve federal-provincial disputes over subjects in the concurrent list, but IT and AI are not concurrent subjects in a way that gives the Council a clear role. The Digital Nation Act includes the four chief ministers in the National Digital Commission, which is a sensible coordination device.[1] The AI Policy includes provincial chief secretaries in the AI Council.[2]

These are useful forums. They are not decision rules. The Commission can approve a Masterplan, but it cannot compel a province to implement it. The AI Council can agree a target, but it cannot fund it. The result is a system that depends on consensus, and consensus depends on politics, money and capacity.

The provinces also compete. Punjab wants to be the leading AI province in South Asia. KP wants a government language model. Sindh wants an AI-native digital province. Competition can be healthy: it produces laboratories, attracts investment and creates peer pressure. It can also produce duplication, incompatible standards and vendor lock-in, particularly when each province signs its own deal with a different technology supplier.

What a federal settlement would look like

A workable federal settlement would have four elements.

First, a minimum national baseline of rights and standards, agreed by the National Digital Commission and enacted where necessary, that applies to any AI system used in a public service, regardless of which level of government operates it. The baseline should cover data protection, human review, non-discrimination, transparency and redress.

Second, provincial flexibility on delivery. Provinces should be able to choose their own vendors, models and service designs, provided they meet the baseline and publish the required information.

Third, fiscal incentives. The federal government can condition digital and AI grants on adherence to the baseline, publication of system registers and participation in national interoperability standards. That is more effective than a legal mandate that the federation cannot enforce.

Fourth, a dispute resolution mechanism. When a province refuses to share data, adopts an incompatible standard or signs a contract that conflicts with national security rules, there needs to be a route to resolve it. The National Digital Commission is the obvious forum, but it needs a secretariat, a technical committee and a published decision procedure.

Assessment

Provincial AI governance is where the action is, and where the fragmentation is. Punjab is building a delivery model, KP is building a language model, Sindh is building a vendor partnership and Balochistan is building basic infrastructure. Each is rational in its own terms. Together they risk a patchwork of standards, rights and data arrangements that no citizen can navigate and no national regulator can oversee.

The federal government's role should be to set the floor, fund the public goods and coordinate the interfaces, not to attempt to run provincial programs. The provinces' role should be to experiment, deliver and report. That division is implicit in the Constitution and explicit in no AI document.

Table 7. Provincial AI and digital initiatives

Province	Lead initiative	Date	Status and main risk
Punjab	AI Roadmap, AI Technical Advisory Board, AI Delivery Unit	March 2026	Approved; delivery capacity and vendor dependence are risks
Punjab	Connected Punjab, 70 million dollars	June 2026	World Bank approved; multi-component, not AI-specific
Punjab	Government AI cloud consultancy	2026	Tendered; business case not public
Khyber Pakhtunkhwa	Digital Transformation Strategy 2030, KP GovGPT	2026	Published; implementation capacity is the constraint
Sindh	Huawei AI-native digital province memorandum	2026	Signed; terms not published
Balochistan	IT hub in Dera Bugti	November 2026	Completed; connectivity and skills remain the gap

CHAPTER

12

Rights, inclusion and labour

AI governance is ultimately about people. This chapter examines who is protected, who is excluded and who bears the costs of adoption. It covers privacy and redress, the digital divide, language, gender, disability, children and work.

IN THIS CHAPTER

Rights

Redress

The digital divide

Language

Children

Disability

THE ARGUMENT IN ONE LINE

Rights are aspirational, inclusion is underfunded and labour has no AI agenda.

Rights

The Constitution gives Pakistanis dignity, privacy, due process, equality and access to information.[37] Those rights apply to AI, but the routes to enforce them are weak.

Privacy is the clearest example. Article 14 protects the privacy of home and the dignity of the person, and the courts have read it to require a lawful procedure for surveillance.[38] But there is no data protection authority, no registration of databases, no statutory privacy notice and no general breach notification duty.[6] A citizen who wants to know what data a government department holds about them can use the Right of Access to Information Act, but the process is slow and the exemptions are broad. A citizen who wants to know what an AI system did with that data has no route at all.

Non-discrimination is the second gap. Article 25 prohibits discrimination on grounds of sex, religion, race, caste or place of birth. An AI system that reproduces historical bias in credit, hiring, policing or service allocation may violate that guarantee, but proving it requires data and expertise that the affected person does not have. There is no statutory duty to test for bias, no requirement to publish disaggregated performance and no regulator with a mandate to investigate algorithmic discrimination.

Due process is the third. A decision that affects a person's rights must be made in accordance with law, with a hearing and reasons. An automated decision can satisfy those requirements in form while defeating them in substance, if the official simply adopts the model's output. The draft data policy's right to meaningful human review is the right response.[4] It is not yet law.

Redress

There is no AI-specific complaint mechanism in Pakistan. A person harmed by an AI system must find a route through the sector regulator, the courts, the cybercrime agency or the deploying institution. Each route has limits.

The sector regulator can investigate a bank, an insurer or a telecom operator, but its mandate is prudential or sectoral, not individual redress. The courts can review an administrative decision, but the person must know that AI was involved and must be able to show what went wrong. The cybercrime agency can investigate a criminal offence, but most AI harms are not crimes. The deploying institution can review its own decision, but it is judging its own system.

The draft data policy would create a rights-handling process, with notices, intake channels, response timelines and complaint handling.[4] It would also require a public registry of high-impact systems.[4] Those are the right foundations. They need a statute, an independent appeal body and published statistics on complaints, response times and outcomes.

The digital divide

Pakistan's connectivity is high by subscription and shallow by capability. Mobile broadband reaches about 62 percent of the population by connections, but fixed broadband, device ownership, affordability and digital skills are unevenly distributed.[17] The groups least likely to be connected are rural women, low-income households, persons with disabilities and older people.[29]

AI can widen or narrow that gap. It can widen it if services are designed for smartphone-literate, English-speaking users and if the cost of connectivity rises. It can narrow it if AI is used to provide services in local languages, by voice, on low-end devices, with human assistance where needed. The policy documents mention inclusion, but they do not contain a concrete inclusion strategy with targets, budgets and evaluation.

The gender gap deserves particular attention. GSMA's surveys have consistently found Pakistan among the countries with the widest gaps between men and women in mobile ownership and mobile internet use.[29] An AI service that assumes a smartphone, a data package and a bank account will exclude many women before it excludes anyone else. The policy's target of training marginalised women is welcome, but training is not the same as access.

Language

Pakistan's official language is Urdu, and its government communicates in Urdu and English. Most Pakistanis speak a regional language at home. AI systems perform best in English, well in Urdu and poorly in Pashto, Sindhi, Balochi, Punjabi and other languages, because the data and the models are thinner.

The Center for Language Engineering at UET Lahore and the National Center for Artificial Intelligence at NUST have built Urdu corpora, speech tools and models.[2] The policy proposes local-language models and a government language model in KP.[15] These are important investments. But a language model is not a language policy. The state needs to fund the collection, curation and licensing of text and speech in each language, set evaluation benchmarks, and require that public services work in the languages people actually speak.

The risk is that AI becomes an English-language layer on top of an Urdu-language state, with regional languages left out entirely. That would deepen the divide between those who can use the system and those who cannot.

Children

Children are a distinct case. They use AI chatbots, recommendation systems and educational tools, often without adult supervision. They are also the subject of data collection by schools, platforms and applications. The draft data policy requires enhanced protections for children, including age-appropriate notice, restrictions on profiling and parental engagement where required.[4] The personal data bill would add more. Neither is in force.

There is no specific regulation of AI chatbots for children, no age verification requirement for generative tools, no mandatory safety testing for educational AI and no route for a parent to complain about a system that harms their child. The online safety debate in Pakistan has focused on content, not on the design of AI systems that interact with children.

Disability

AI has significant potential for persons with disabilities: speech-to-text, text-to-speech, image description, translation into sign language and assistive navigation. It also has significant risks: biased speech recognition, inaccessible interfaces, automated decisions that exclude people with atypical profiles.

Pakistan's disability law, the Persons with Disabilities Act 2020, requires accessibility in public services. It was not written for AI. There is no accessibility standard for AI systems, no requirement for testing with disabled users and no guidance on procurement. The policy mentions persons with disabilities in its training target.[2] That is a start, not a strategy.

Labour

Pakistan's labour market is young, large and informal. AI will affect it through automation, augmentation and algorithmic management. The effects will differ by sector: clerical and customer service work is exposed to automation; software development and data work are exposed to augmentation; gig work is exposed to algorithmic control.

The policy response has been skills training, which is necessary but not sufficient. It does not address workers who lose their jobs, the quality of new jobs, the bargaining power of workers managed by algorithms or the social protection needed during transition. There is no redundancy scheme tied to automation, no portable benefits for gig workers, no right to explanation for an algorithmic dismissal and no requirement for employers to consult before deploying systems that change work.

The gig economy is the most immediate case. Ride-hailing, delivery and freelance platforms use algorithms to assign work, set prices and rate workers. The workers are classified as independent contractors, which limits their access to social security and collective bargaining. A rating algorithm can deactivate a worker without a hearing. The legal framework has not caught up, and the platforms are not required to explain their decisions.

A serious labour agenda would have four elements: transparency and appeal rights for algorithmic management, portable social protection for gig workers, adjustment support for displaced workers and a national skills strategy that is linked to actual employer demand. None of these requires an AI act. All of them require a government that treats AI as a labour market issue, not only a technology issue.

Table 8. Rights and inclusion: instruments and gaps

Issue	Existing instrument	Gap
Privacy	Constitution, Article 14; PECA; sector rules	No data protection law or authority
Access to information	Article 19A; Right of Access to Information Act 2017	No register of AI systems
Non-discrimination	Constitution, Article 25	No bias testing or disaggregated reporting
Human review	Draft data policy, clause 14	Not in force
Redress	Sector regulators; courts; NCCIA	No AI-specific complaint route
Gender	National policies; GSMA evidence	No AI inclusion targets or budget
Language	Urdu and English official use	No funded multilingual data program
Children	Draft data policy; PECA	No AI-specific child safety rules
Disability	Persons with Disabilities Act 2020	No AI accessibility standard
Labour	Provincial labour laws	No algorithmic management rights

Assessment

Pakistan's rights and inclusion framework for AI is aspirational, fragmented and underfunded. The Constitution provides the principles, the draft data policy provides some of the design, and sector regulators provide some of the enforcement. But there is no single route for a citizen to exercise a right, no published data on who is harmed and no mechanism to ensure that the benefits of AI reach the people who need them most.

The most important single measure would be a data protection law with an independent authority and a redress function. The second would be a public register of high-impact AI systems, with a plain-language description and a complaint route. The third would be a funded inclusion program, with targets for women, rural users, persons with disabilities and regional languages. These are not expensive relative to the cost of exclusion. They are the difference between AI governance for the few and AI governance for the country.

CHAPTER

13

International alignment, compute and geopolitics

Pakistan's AI governance is shaped as much by what happens abroad as by what happens in Islamabad. The chips come from abroad, the cloud is mostly abroad, the standards are set abroad and the capital is increasingly foreign. This chapter examines Pakistan's international posture, the compute question and the risks of a bifurcating global AI order.

IN THIS CHAPTER

The strategic posture

The World AI Cooperation Organization

The United States, export controls and chips

The European Union and the standards question

Multilateral engagement

Investment and the one billion dollar ambition

THE ARGUMENT IN ONE LINE

Sovereignty rhetoric depends on foreign hardware, foreign capital and foreign standards.

The strategic posture

Pakistan's public position is sovereign and non-aligned. The Islamabad Declaration commits the state to a sovereign approach, private-sector leadership and international partnerships consistent with the Constitution and national interests.[27] The AI Policy commits it to bilateral and multilateral agreements, participation in global forums, collaborative research, cross-border projects, adoption of international standards and international investment.[2]

That posture is sensible for a country that needs access to technology from both China and the West. It is also difficult to sustain. The global AI order is splitting into competing technology stacks, standards bodies and export-control regimes. A middle power can hedge, but it cannot avoid choosing on specific issues: which vendor supplies its data centres, which standards it adopts for high-risk systems, which jurisdiction's law governs its cloud data and which international body it supports when rules conflict.

The World AI Cooperation Organization

In October 2026 the Deputy Prime Minister told the Senate that Pakistan had joined the World AI Cooperation Organization, a China-led initiative with around 50 members, headquartered in Shanghai, whose founding document Pakistan had signed and whose first council meeting was expected in mid-December 2026.[22]

The organization is new, and its rules, membership obligations and decision procedures are not yet public in detail. Its stated aims, transparency, openness and cooperation, are broad enough to accommodate many positions. Its significance is structural: it gives China a multilateral vehicle for AI governance outside the OECD and the Western-led institutions, and it gives Pakistan a seat in a body where it can shape norms rather than only adopt them.

The risk is that membership becomes a signal of alignment in a bifurcating system. If the organization develops its own standards, certification schemes or data-sharing arrangements that are incompatible with those used by Pakistan's export markets, Pakistani firms could face a choice between compliance regimes. If it remains a dialogue forum, the risk is low and the diplomatic value is real.

The United States, export controls and chips

Access to advanced AI chips is the most consequential international constraint on Pakistan's AI ambitions.

The United States controls the export of advanced accelerators and related technology. In January 2025 it published an AI diffusion framework that would have divided the world into tiers, with quantitative caps on compute exports to most countries. The framework was rescinded later in 2025, but the underlying controls on advanced chips and chip-making equipment remain.[32] Pakistan's access to the most advanced hardware therefore depends on US licensing policy, the availability of alternative suppliers and the willingness of vendors to sell.

Pakistan has not published a position on export controls, a chip procurement strategy or an assessment of the compute it needs. That silence is a strategic weakness. Without a public assessment, the government cannot negotiate with suppliers, prioritise public investment or explain to universities and firms why compute is scarce. It also cannot prepare for the scenario in which access becomes more restricted.

The alternative supply route is Chinese. Huawei and other Chinese vendors offer AI accelerators, cloud platforms and data centre equipment, and they are actively marketing to governments in Asia and Africa. Chinese hardware may be more available and cheaper, but it comes with its own dependencies:

software ecosystems, security concerns, interoperability with Western tools and the risk of secondary sanctions. Pakistan's choice is not between dependence and independence. It is between different forms of dependence, managed with different degrees of transparency.

The European Union and the standards question

The European Union's AI Act is the most developed regulatory regime in the world.[24] It matters to Pakistan for three reasons.

First, Pakistani firms that sell AI-enabled products or services into the EU, or that process EU personal data, will have to comply with EU rules. That is already true for data protection, where the absence of an adequacy decision means Pakistani firms rely on contractual safeguards. It will become more true for AI, where high-risk systems will require conformity assessment and documentation.

Second, the EU's risk-based approach is becoming a reference model for other jurisdictions. Pakistan's draft data policy already uses risk tiers, pre-production registration and human review.[4] Aligning those concepts with the EU framework, while adapting the obligations to Pakistan's capacity, would make it easier for Pakistani firms to operate in both markets.

Third, the EU is a source of development finance and technical assistance for digital governance. Its requirements for data protection, transparency and independent oversight are likely to shape the projects it funds.

The United Kingdom, Singapore, Japan, South Korea, India and the UAE have all published AI governance frameworks. Pakistan can learn from them without adopting any of them wholesale. The most useful comparison is India, which has a similar federal structure, a large IT services sector and a comparable capacity problem. India's approach has been to use existing regulators, issue guidance and invest in compute and skills, rather than to create a single AI regulator. Pakistan's architecture is closer to that model than to the EU's.

Multilateral engagement

Pakistan participates in the UN, UNESCO, the ITU, the OECD's AI Policy Observatory as an observer or partner in some activities, the Commonwealth AI Consortium and other forums.[56,23] It has aligned with the UNESCO Recommendation on the Ethics of AI and cites the ITU's AI for Good initiative in its policy.[2]

These engagements are useful for norm-setting, technical assistance and reputation. They do not substitute for domestic institutions. A country can sign every declaration and still have no data protection law, as Pakistan demonstrates.

Investment and the one billion dollar ambition

The government has announced a one billion dollar investment ambition for Pakistan's AI ecosystem by 2030.[22] The figure is a statement of intent, not a budget line, and no public pipeline, investor list or financing plan accompanies it. The 2026-27 development budget proposes Rs 524.64 million for the National AI Advancement Initiative, which is a small fraction of the ambition.[3]

Foreign investment in Pakistani AI is real but modest. The World Bank is financing digital public infrastructure and the digital economy.[59] Chinese vendors are supplying cloud and data centre equipment. Gulf investors are active in Pakistani start-ups and digital infrastructure. The diaspora is a

source of capital, expertise and market access. None of this yet adds up to one billion dollars, and the government should either publish a credible pipeline or describe the figure as an aspiration.

Data flows and foreign access

The draft data policy's Tier 1 residency rule would keep restricted, confidential and personal data in Pakistan.[4] That is a significant restriction on cloud use, and it will affect the government's relationships with foreign providers. It also raises a question that the policy does not answer: what happens when a foreign court or regulator demands access to data held by a Pakistani public body or its contractor?

The policy requires notification to the Authority and prohibits voluntary disclosure, with compliance only as permitted by Pakistani law and mutual legal assistance arrangements.[4] That is the right principle. In practice, Pakistani bodies have limited ability to resist a demand from a foreign jurisdiction, particularly when the data is held by a subsidiary of a foreign company. The legal framework for mutual legal assistance is slow, and the political asymmetry is significant. A government that wants sovereign control of data needs not only a residency rule but also a cloud access law, contractual protections and a diplomatic strategy.

Climate, energy and the cost of compute

AI is energy-intensive, and Pakistan's energy system is expensive and carbon-intensive. The competitive electricity market, operational since January 2026, is an attempt to let large consumers, including data centres, buy power directly at lower cost.[31] If it succeeds, it could make Pakistan a competitive location for regional compute. If it fails, the country will continue to lose workloads to the Gulf and Southeast Asia.

The climate dimension is also a development opportunity. AI can improve grid management, water allocation, crop forecasting and disaster response, all of which matter enormously in Pakistan. The policy mentions agriculture and climate in its sectoral ambitions but does not fund a dedicated program. A national AI for climate and food security mission, with public data, open models and provincial delivery, would be a better use of scarce resources than another general-purpose hub.

Assessment

Pakistan's international strategy is long on membership and short on leverage. It has joined the right forums, signed the right declarations and attracted useful partners. It has not published a compute strategy, a standards alignment plan, a position on export controls or a credible investment pipeline. The result is a sovereignty rhetoric that depends on decisions made elsewhere.

The remedy is not to pick a side. It is to be explicit about the choices: which standards Pakistan will adopt for high-risk systems, which jurisdictions it will treat as adequate for data transfers, which vendors it will trust for sovereign infrastructure, how it will respond to export controls and what it will offer international partners in return for access. Those choices can be made in a public strategy, with trade-offs stated. They cannot be avoided.

CHAPTER

14

The implementation audit

This chapter brings the report's evidence together in a single assessment. It asks a simple question of each instrument and commitment: what exists, in what legal form, with what operational evidence? The answer is summarised in the table and explained in the sections that follow.

IN THIS CHAPTER

The instrument register

What the register shows

Budget and staffing

Procurement and delivery

Sectoral delivery

Why the gap persists

THE ARGUMENT IN ONE LINE

The gap is produced by incentives that reward documents over decisions.

The instrument register

Table 9. AI governance instruments and implementation status, November 2026

Instrument	Type	Status	Operational evidence
Digital Nation Pakistan Act 2025	Statute	In force	PDA constituted; Commission and Committee not publicly reported
National AI Policy 2025	Policy	Approved	AI Council not publicly constituted; no annual report
Islamabad AI Declaration	Declaration	Adopted	No legal effect; no implementation plan
National Data Governance Policy 2026	Draft policy	Draft	Not in force; supporting instruments not issued
Personal Data Protection Bill	Draft bill	Draft	Not enacted; no authority
WASL	Platform and specification	Operating design	No published coverage or transaction statistics
AI risk classification and impact assessment	Standards	Not located	No published instrument
Public registry of AI systems	Registry	Not created	Required by draft policy only
National AI Fund	Fund	Not established	No corpus, committee or call published
AI Council	Coordination body	Process begun	No membership or minutes published
Policy Implementation Cell	Unit	Not visible	No public record
National Digital Masterplan	Plan	Reported approved	Text not published
National AI Advancement Initiative	Program	Budgeted	Rs 524.64 million proposed; delivery not reported
National AI Hubs	Procurement	Tendered	No award or completion published
GPU and model services, NTC	Service	Launched	Capacity and pricing not published
National Guidelines for AI in Judicial Institutions	Guidance	Issued	No court-level policy or audit published

Instrument	Type	Status	Operational evidence
HEC generative AI framework	Guidance	Draft	No compliance data
National CERT handbook	Guidance	Issued	No compliance audit published
PPRA Rules 2026	Regulation	Notified	Applies to new procurements
SBP AI survey	Survey	Published	55 entities; 2024 data
SBP responsible AI guidelines	Guidance	Reported in 2024	Not published
SECP regulatory sandbox	Mechanism	Operational	Six solutions in first cohort
Social Media Protection and Regulatory Authority	Statutory body	Constituted	No budget, office or staff reported

What the register shows

The register contains 23 entries spanning legal instruments, institutions, platforms and initiatives. Figure 1 groups each entry once, by document type and publication status. These categories describe the register; they are not a score of implementation or a count of operational institutions.

- **Binding law and rules** (2): the Digital Nation Pakistan Act and PPRA Rules 2026.
- **Approved policy and declaration** (2): the National AI Policy and Islamabad AI Declaration.
- **Draft frameworks** (3): the National Data Governance Policy, Personal Data Protection Bill and HEC generative AI framework.
- **Published guidance and survey** (3): the judicial AI guidelines, National CERT handbook and SBP AI survey.
- **Platforms, services and sandbox** (3): WASL, NTC GPU and model services, and the SECP regulatory sandbox. Inclusion records the platform, service or mechanism described; it does not establish full deployment or independent assurance.
- **Other bodies and initiatives** (10): AI risk classification and impact assessment, the public AI registry, National AI Fund, AI Council, Policy Implementation Cell, National Digital Masterplan, National AI Advancement Initiative, National AI Hubs, proposed SBP responsible AI guidelines, and the Social Media Protection and Regulatory Authority. These entries have different reported statuses, as Table 9 records; the category does not imply that all are inactive or legally equivalent.

No AI-specific regulation was identified among the registered entries. The review did not locate an AI-specific enforcement decision, audit report or compliance assessment published by the institutions covered. Missing public evidence is not proof that no internal activity occurred.

The pattern is not random. It reflects a system in which the production of documents is rewarded and the production of decisions is not. A ministry that publishes a policy receives credit. A regulator that publishes guidance receives credit. An institution that runs a survey receives credit. An institution that

refuses a procurement, fines a vendor, orders a system withdrawn or publishes an adverse audit receives attention, criticism and legal risk. The incentives favour announcement over enforcement.

Budget and staffing

The budget picture is thin. The FY27 development budget proposes Rs 524.64 million for the National AI Advancement Initiative and Rs 16.29 billion for the wider IT portfolio, most of it for buildings.[3] There is no published allocation for the AI Council, the policy implementation cell, the data governance council, the national data officer, the AI registry or the assurance function. A governance system without a budget is a governance system that depends on the goodwill of existing staff.

The staffing picture is similar. The Authority has advertised senior posts, including a Chief AI Officer.[9] It has not published its organizational chart, its headcount, its salary scales or its vacancy rate. The AI Council has no published secretariat. The data policy creates a National Chief Data Officer and Agency Chief Data Officers in every public body, but no appointments have been announced.[4]

Procurement and delivery

The AI hubs procurement is the clearest test of delivery. A tender was issued, a reference number was assigned and a deadline was extended.[47,48] No award notice, completion certificate or budget execution figure has been published. The same is true of the Punjab government AI cloud consultancy and the national AI platforms proposed in the policy.[49,2]

It is possible that these projects are proceeding and that the information is available on request. But a governance audit can only assess the public record, and the public record does not show delivery. The government's own transparency commitments, in the Digital Nation Act and the data policy, require more.[1,4]

Sectoral delivery

The sectoral picture is more varied. The State Bank has done the most with the least: a survey, a commitment to guidance and a set of cloud outsourcing rules.[10,42] The SECP has a functioning sandbox.[11] The courts have guidelines.[12] HEC has a draft framework and a training program.[13,33] The National CERT has a handbook.[44]

But none of these institutions has published an assessment of a deployed system, an audit of a vendor, a complaints statistic or an enforcement decision. The guidance exists; the evidence of effect does not.

Why the gap persists

Five explanations are more plausible than the others.

The first is incentives. Announcements are politically valuable; enforcement is not. A minister can cut a ribbon at an AI hub. A regulator who blocks a system creates an enemy. The system rewards the former and punishes the latter.

The second is capacity. The Authority is new and small. The sector regulators have limited AI expertise. The civil service has few people who can read a model card, write a data schedule or evaluate a vendor's claims. Capacity is not built by hiring a few advisers; it is built by training, retaining and empowering a cadre.

The third is legal caution. Officials are unsure what they can do without a statute. The Digital Nation Act gives the Authority power in assigned areas, but the assignment is not public. The data policy is a

draft. The personal data bill is not law. In the face of uncertainty, the safest course is to publish guidance and wait.

The fourth is federalism. The federal government cannot compel provinces, and the provinces have their own priorities. Coordination requires negotiation, and negotiation takes time.

The fifth is the donor dynamic. International partners fund strategies, workshops and dashboards more readily than they fund salaries, audits and enforcement. That skews the system towards planning and away from operations.

What would change the picture

Three interventions would make the biggest difference in the shortest time.

First, publication. Require every federal body to publish an AI system register, a data protection impact assessment for high-impact systems and an annual compliance statement. The act of publication creates accountability even without enforcement.

Second, budget linkage. Tie the AI program's funding to the delivery of specific governance outputs: a functioning registry, a published risk classification, a trained assessor cadre, an audited procurement. Money should follow evidence, not announcements.

Third, parliamentary scrutiny. Refer the annual AI governance report to the relevant standing committees, require the Authority to appear before them and publish the questions and answers. Parliaments are better than regulators at creating political cost for failure.

The next chapter sets out scenarios for how the gap might close, or widen, by 2030.

CHAPTER

15

Scenarios to 2030

Scenarios are not forecasts. They are structured stories about how the future could unfold, used to test decisions against more than one plausible outcome. The three scenarios below are built from two uncertainties that matter most for AI governance in Pakistan: the capacity of public institutions to implement rules, and the openness of the international environment to Pakistani access to compute, capital and markets.

IN THIS CHAPTER

Scenario A: managed convergence

Scenario B: fragmented adoption

Scenario C: control-first

Leading indicators

Implications

THE ARGUMENT IN ONE LINE

Fragmented adoption is the most likely path. It is stable until it is not.

For comparative planning, PIVRA assigns illustrative weights of 25 percent to managed convergence, 50 percent to fragmented adoption and 25 percent to control-first. Fragmented adoption receives twice the weight of either alternative because it extends the institutional fragmentation described in this report. Managed convergence requires coordinated reform; control-first depends on a substantial shift towards security-led restrictions. The equal weights for those two alternatives express no stronger preference between them. These are qualitative judgements expressed numerically, not statistically estimated probabilities or predictions, and should be reconsidered as the leading indicators change.

Scenario A: managed convergence

In this scenario, the gap closes. The Personal Data Protection Act is enacted in 2027, with an independent authority and a transitional period for public bodies. A short AI statute follows in 2028, establishing risk tiers, a register of high-impact systems, a right to human review and a redress route. The Pakistan Digital Authority becomes operational, with a published budget, staff and annual report. The AI Council meets quarterly and publishes its decisions. The provinces sign a minimum baseline agreement and adopt compatible registers. Sector regulators issue AI guidance, and a small assurance unit begins auditing high-impact systems.

The conditions for this scenario are demanding: political continuity, fiscal space, a functioning Parliament, a cooperative federation and continued access to foreign technology. It is assessed as less likely than fragmented adoption, with the same illustrative weight as control-first. The architecture is already in place; what is missing is the decision to make it work.

The benefits would be substantial. Citizens would have a route to challenge automated decisions. Investors would have a clearer regulatory environment. Government procurement would be more competitive and less corruptible. Pakistan would be better placed to attract the data centre, cloud and AI services investment that it wants. The cost would be modest relative to the benefits: a data protection authority, a small AI secretariat, an assurance unit and a register.

Scenario B: fragmented adoption

In this scenario, AI spreads faster than governance. Banks, telecoms, hospitals, universities and provincial governments adopt tools vendor by vendor. The federal government publishes more strategies but does not enforce them. The provinces go their own way, signing deals with different vendors and building incompatible systems. The data protection bill is reintroduced but not passed. The AI Council is constituted but meets irregularly. A register is promised but not published.

This is the most likely scenario, because it is the continuation of the present. It delivers real benefits: better fraud detection, faster services, new firms, new jobs. It also delivers real harms: opaque decisions, data breaches, vendor lock-in, exclusion of those without connectivity or English, and a patchwork of rights that depends on where a citizen lives.

The scenario is stable in the short run and unstable in the long run. Each incident erodes trust. Each breach creates pressure for restrictive rules. Each vendor lock-in raises costs. The absence of a governance framework does not prevent governance; it defers it to a crisis, when the rules are likely to be written hastily and badly.

Scenario C: control-first

In this scenario, security concerns dominate. A major incident, a political crisis or a conflict leads to tighter controls on content, data and connectivity. AI is treated primarily as a surveillance and censorship technology. Data localisation is enforced without the safeguards in the draft policy. Foreign cloud providers are restricted. The content authority becomes operational with broad powers and limited appeal. Investment slows, talent leaves and international partners disengage.

This scenario is not the most likely, but it is not remote. Pakistan has a history of internet shutdowns, content blocking and surveillance controversies. The 2025 amendments to PECA created a content authority with broad powers and a false-information offence that can be used against journalists.[40] A security shock could turn the current balance towards control.

The costs would be severe. AI governance would become a subset of security policy, with no independent oversight, no rights and no redress. The economy would lose access to foreign cloud and compute. The best firms and researchers would leave. The country would be less able to use AI for public benefit, at exactly the moment when the technology is most consequential.

ILLUSTRATIVE PLANNING WEIGHTS | NOT A FORECAST

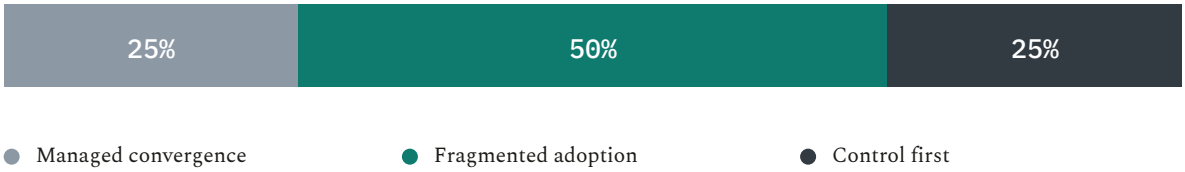


Figure 7. Illustrative weights for the scenarios to 2030. Source: PIVRA qualitative assessment, Chapter 15. Fragmented adoption is weighted twice either alternative because it extends the current institutional pattern. These are planning weights, not statistically estimated probabilities.

Leading indicators

The following indicators would tell an observer which scenario is unfolding. They are chosen because they are observable, dated and difficult to fake.

- Enactment of the Personal Data Protection Act, with an independent authority and a commencement date.
- Notification of the AI Council, with published membership, terms of reference and minutes.
- Publication of the Pakistan Digital Authority's first annual report, budget and staffing plan.
- Publication of the AI risk classification and impact assessment instruments.
- Creation of a public register of high-impact government AI systems, with entries.
- Issuance of AI-specific guidance by the SBP, SECP, PTA, DRAP and HEC.
- Publication of WASL coverage, transaction and consent statistics.
- A federal-provincial agreement on a minimum AI rights and standards baseline.
- Enactment of an AI statute with risk tiers, redress and enforcement powers.

- A published compute strategy with a demand assessment and a public-interest allocation.
- A published position on export controls and a chip procurement plan.
- An annual State of AI Governance report debated in Parliament.

Table 10. Scenario comparison

Dimension	Managed convergence	Fragmented adoption	Control-first
Data protection law	Enacted 2027	Not enacted	Enacted but weak
AI statute	Enacted 2028	Absent	Security law, not AI law
AI Council	Operational and public	Nominal	Subordinated to security
Registry of systems	Published	Promised	Classified
Sector guidance	Issued and enforced	Issued, not enforced	Restrictive
Provincial coordination	Baseline agreement	Patchwork	Centralised
Compute access	Diversified	Vendor-dependent	Restricted
Rights and redress	Meaningful	Uneven	Absent
Investment	Growing	Moderate	Declining
Illustrative weight, PIVRA assessment	25 percent	50 percent	25 percent

Implications

For government, the scenarios argue for acting early. The cost of a data protection law and a register is low; the cost of a crisis-driven response is high. The window for a managed transition is open now, but it will not stay open indefinitely.

For business, the scenarios argue for preparing for more regulation, not less. Firms that can document their data, test their models, explain their decisions and report incidents will be better placed in every scenario. Firms that rely on opacity will be exposed in the first two scenarios and may thrive briefly in the third, at the cost of market access.

For citizens and civil society, the scenarios argue for demanding transparency now. A register of AI systems, a right to human review and a complaints route are achievable in the next two years if there is political pressure. They will be much harder to secure after a crisis.

For donors and international partners, the scenarios argue for funding institutions, not only strategies. A trained assessor, a functioning registry and an independent authority are less photogenic than a launch event, but they are what determines whether AI governance works.

CHAPTER

16

Recommendations

The recommendations below are sequenced by feasibility and impact. They assume that the current government remains in office and that no major fiscal or security shock occurs. They are deliberately modest in legal ambition and ambitious in institutional effect. Pakistan does not need a 200-page AI act. It needs a small number of enforceable rules, a few functioning institutions and a lot of publication.

IN THIS CHAPTER

Do these five first

Build the institutions, 6 to 24 months

Deepen accountability, 12 to 36 months

What not to do

Sequencing

THE ARGUMENT IN ONE LINE

Five measures in the next year would change the trajectory.

THE REPORT'S RECOMMENDED SEQUENCE

FIRST YEAR

Establish enforceable foundations

Data protection law; a public AI register; procurement clauses; an operational Council; sector work plans.

6-24 MONTHS

Build institutional capacity

Risk assessment and assurance; WASL audit; compute and skills plans; federal-provincial coordination; multilingual data.

12-36 MONTHS

Deepen accountability

A short AI statute; an annual governance report; standards alignment; public AI literacy.

Figure 8. A sequence for AI governance reform. Source: Chapter 16. The phases overlap; individual deadlines and responsible institutions are specified in Table 11.

Do these five first

1. Enact the Personal Data Protection Act. This is the single most important measure. Without it, every other AI governance reform is built on sand. The Act should create an independent authority with investigation, audit and sanction powers, a registration or notification regime for high-risk processing, breach notification, data subject rights and a redress route. It should include a transitional period for public bodies and a clear interface with the Pakistan Digital Authority.[5,6]

2. Publish a register of high-impact government AI systems. The draft data policy already requires one.[4] The government should not wait for the policy to be finalised. It should issue an interim instruction requiring every federal ministry and agency to list the AI systems it uses, their purpose, the data they process, the vendor, the risk classification, the human review arrangement and the complaint route. The register should be public, machine-readable and updated quarterly. A provincial register should follow.

3. Issue standard AI procurement clauses. The Authority and the Public Procurement Regulatory Authority should publish a model contract schedule covering data rights, training restrictions, performance testing, audit, incident reporting, human review, liability and exit. The clauses should be mandatory for federal AI procurements above a threshold and should be incorporated into the PPRA rules or issued as binding guidance.[28]

4. Notify the AI Council and publish its work. The Council should be constituted with the membership set out in the policy, a published terms of reference, a quarterly meeting schedule, published minutes and an annual report to Parliament.[2] If the Council is not the right vehicle, the government should say so and designate another. A nominal council is worse than no council because it creates the appearance of oversight without the substance.

5. Give the sector regulators an AI mandate. The SBP, SECP, PTA, DRAP, HEC and the NJPMC should each publish an AI work plan, a timeline for guidance and a point of accountability.[10,11,12,13] The plans should be specific: model risk guidance for banks, AI guidance for capital markets, AI in network operations for telecoms, a clinical validation pathway for health AI, a binding academic integrity policy for universities and court-level AI policies for the judiciary.

Build the institutions, 6 to 24 months

6. Publish the AI risk classification and impact assessment instruments. The DNP-A series should be issued for public consultation, with a risk tier structure, a screening questionnaire, an assessment methodology and a template for publication.[60] The methodology should be compatible with the NIST AI Risk Management Framework and ISO/IEC 42001 so that Pakistani assessments are recognised abroad.[25,26]

7. Create an AI assurance unit and an assessor cadre. The unit should sit within or alongside the Authority, with a mandate to audit high-impact systems, commission independent evaluations, publish results and maintain a register of qualified assessors. It should be funded from the Digital Nation Fund and staffed with a mix of engineers, lawyers, data scientists and sector specialists.[43]

8. Publish WASL statistics and commission an independent audit. The Authority should publish the number of connected institutions, the volume and type of transactions, the number of consent requests and refusals, the error rate and the number of security incidents. It should commission an independent audit of the exchange's governance, including the separation of NADRA's roles as operator, custodian and user.[52]

9. Adopt a national compute strategy. The strategy should assess demand, map existing capacity, set a public-interest allocation for universities and public research, address the power tariff and the competitive wheeling framework, and set out a position on export controls and vendor diversification.[31,21] It should be published for consultation and updated annually.

10. Cost and sequence the skills program. The government should publish a costed delivery plan for the training, internship and scholarship targets, with named delivery partners, completion measures and a funding model that includes provincial and private contributions.[2] It should report completion, not enrolment, and it should track employment outcomes.

11. Create a federal-provincial AI compact. The National Digital Commission should agree a minimum baseline for AI in public services, covering data protection, human review, non-discrimination, transparency and redress. Provinces should retain flexibility on delivery, but federal digital and AI grants should be conditional on adherence to the baseline and publication of registers.[1]

12. Fund a multilingual data program. The government should fund the collection, curation, licensing and evaluation of text and speech in Urdu and the major regional languages, with open benchmarks and public licences. The program should be led by a consortium of universities and language technology centres, not by a single vendor.[2]

Deepen accountability, 12 to 36 months

13. Enact a short AI statute. The statute should establish principles, risk tiers, a register of high-impact systems, a right to human review, a redress route, audit powers and penalties. It should not create a licensing regime for general-purpose models, which would be unenforceable and would entrench incumbents. It should build on the institutions and instruments already in place, not duplicate them.

14. Institute an annual State of AI Governance report. The Authority should publish an annual report with a register of systems, compliance statistics, audit findings, incident data, enforcement decisions and recommendations. The report should be debated in the relevant standing committees of Parliament, with the Authority's leadership appearing in person.

15. Adopt a standards alignment strategy. The government should publish which international standards it will recognise for AI management, risk management, testing and documentation, and how it will mutual-recognise assessments with major trading partners. It should engage with the World AI Cooperation Organization, the OECD, UNESCO and the ISO processes with a coherent position, not a series of ad hoc statements.[61,23]

16. Build public AI literacy. The government should fund a national AI literacy campaign, in Urdu and regional languages, covering how AI systems work, how to spot synthetic content, how to protect personal data and how to complain. It should be delivered through schools, vocational institutions, libraries, community organisations and telecom operators.[2]

Table 11. Recommendations, owners and timing

Recommendation	Lead	Timing	Success measure
Personal Data Protection Act	Parliament, MoITT	2027	Act in force; authority appointed
Interim AI system register	PDA, MoITT	6 months	Public register with entries
Standard AI procurement clauses	PDA, PPRA	6 months	Mandatory schedule issued
AI Council notification	MoITT	3 months	Members, terms and minutes published
Sector AI work plans	SBP, SECP, PTA, DRAP, HEC, NJPMC	12 months	Guidance issued and applied
Risk classification and impact assessment	PDA	12 months	Instruments consulted and published
Assurance unit and assessor cadre	PDA	18 months	First audits published
WASL statistics and audit	PDA, NADRA	12 months	Coverage and audit published
National compute strategy	MoITT, PDA	12 months	Strategy published
Costed skills plan	MoITT, HEC, NAVTTC	9 months	Completion and employment data
Federal-provincial compact	National Digital Commission	18 months	Baseline agreed

Recommendation	Lead	Timing	Success measure
Multilingual data program	HEC, universities	24 months	Corpora and benchmarks released
AI statute	Parliament	2028	Act in force
Annual State of AI Governance report	PDA	Annual	Report debated in Parliament
Standards alignment strategy	MoFA, MoITT	18 months	Strategy published
National AI literacy campaign	MoITT, provinces	24 months	Reach and comprehension measured

What not to do

The report also has advice on what to avoid.

Do not create a new AI regulator before the data protection authority exists. Two overlapping regulators with no clear jurisdiction would make the current coordination problem worse.

Do not license general-purpose models. Pakistan does not have the capacity to evaluate every model, and a licensing regime would mainly burden domestic start-ups while foreign models enter through the cloud.

Do not write a comprehensive AI act by copying the EU. The EU's institutional machinery, market surveillance system and enforcement budget are not available in Pakistan. A shorter statute with delegated standards is more likely to be implemented.

Do not rely on voluntary codes for high-impact systems. Voluntary codes are useful for experimentation and for low-risk uses. They are not adequate for systems that allocate benefits, deny services, determine eligibility or affect liberty.

Do not treat a declaration, a strategy or a memorandum of understanding as implementation. The test is a published decision, a register entry, an audit finding or a budget execution figure.

Sequencing

The five first measures are mutually reinforcing. The data protection act creates the authority and the rights. The register creates visibility. The procurement clauses create leverage. The Council creates coordination. The sector mandates create enforcement in the places where AI is already used.

The institutional measures take longer because they require money, people and organisational change. They should begin now, in parallel, so that the legal framework does not arrive before the capacity to apply it.

The accountability measures are the hardest because they require political will. They are also the ones that make the rest sustainable. A government that publishes an annual report, appears before Parliament and explains its failures is a government that is serious about governance. A government that does not is managing appearances.

CHAPTER

17

Conclusion

Pakistan's AI governance has a paradox at its centre. The country has produced more policy in under two years than in the previous twenty, and it has produced very little governance. It has a law that creates institutions, a policy that sets targets, a declaration that states principles, a data framework that designs a system and sectoral guidance that advises professionals. It does not have a data protection law, a functioning AI council, a risk classification, a register of systems, an audit or an enforcement decision.

THE ARGUMENT IN ONE LINE

The test is whether a citizen can find, challenge and understand an automated decision.

The paradox is not unique to Pakistan. Many countries produce strategies faster than institutions. What makes Pakistan's case worth studying is the clarity of the gap: the documents show what a working system would look like, and the implementation shows what is missing.

The report's central argument is that the gap is a choice, not a destiny. Pakistan has the legal materials, the institutions and the talent to close it. It lacks the decision to prioritise implementation over announcement, and the political mechanisms that would make implementation matter. A data protection law, a register, a set of procurement clauses, a functioning council and a handful of sector mandates would transform the picture within two years. None requires a new constitutional settlement or a large budget, only a government willing to be held to account.

The report's second argument is that AI governance cannot be separated from the material conditions of AI. Compute, electricity, data quality, skills and connectivity are not background factors. They determine what is possible. A government that writes rules for systems that cannot be built, or that ignores the systems that are already being used, will have governed a fiction. The most useful reforms are often unglamorous: a land registry that works, a power tariff that makes data centres viable, a statistics office that can measure adoption, a university that can teach AI in Urdu.

The report's third argument is that rights are not a luxury to be added after growth. They are the mechanism through which growth is made legitimate and durable. A citizen who cannot challenge an automated decision, a worker who cannot understand an algorithmic dismissal, a woman who cannot get online, a patient who cannot find out whether a tool contributed to a misdiagnosis: these are not edge cases. They are the test of whether AI governance is for the country or for its vendors.

Pakistan has an opportunity that few countries at its income level have: a young population, a large diaspora, a growing IT sector and a government that has stated its ambition. It also has a closing window: the choices made in the next two to three years will determine whether AI becomes a general-purpose technology that improves public services and creates jobs, or another enclave technology that benefits a few firms in a few cities while the rest of the country watches.

The report ends with a simple test. In two years, a citizen should be able to find out whether an AI system was used in a decision about them, ask for a human review and get a reply. A parliamentarian should be able to read a register of government AI systems and a report on their performance. A regulator should be able to point to an audit, a condition or a penalty. A province should be able to adopt a national standard without surrendering its priorities. A student in Multan should be able to learn AI in Urdu and get a job that uses it.

If those things are true, Pakistan will have governed intelligence. If they are not, it will have governed the appearance of intelligence, which is a different and less useful achievement.

APPENDIX A

Timeline of AI governance in Pakistan

Table A1. Selected events, 2016 to 2023

Date	Event	Significance
2016	Prevention of Electronic Crimes Act	Criminal law for digital conduct; later amended
2018	Digital Pakistan Policy	First broad digital policy; no AI-specific rules
2021	National Cyber Security Policy	Cybersecurity framework; AI mentioned as emerging technology
2022	Pakistan Cloud First Policy	Cloud adoption rules for public sector
2022	Personal Data Protection Bill approved by cabinet	Never enacted
2023	Revised Personal Data Protection Bill	Still a draft

Table A2. Selected events, 2025

Date	Event	Significance
29 January 2025	Digital Nation Pakistan Act 2025	Creates Commission, Authority and Oversight Committee
January 2025	Prevention of Electronic Crimes (Amendment) Act	NCCIA, false information offence, content authority
31 July 2025	National AI Policy approved	First national AI policy
August 2025	Pakistan Digital Authority constituted	Chair and member appointed

Table A3. Selected events, 2026

Date	Event	Significance
January 2026	AI module at Civil Services Academy	Civil service capacity building
9 February 2026	Islamabad AI Declaration	Nine principles; no legal effect
February 2026	AI Council formation process reported	No membership published
12 March 2026	Punjab AI Roadmap approved	Provincial AI strategy and delivery unit
March 2026	Social Media Protection and Regulatory Authority constituted	Not operational by July 2026
29 April 2026	National Guidelines for AI in Judicial Institutions	Guidance for courts
June 2026	Draft National Data Governance Policy	Public consultation; not in force

Date	Event	Significance
30 June 2026	World Bank approves Connected Punjab	70 million dollars; part of 278 million dollar programme
18 July 2026	Technology exports reach 4.6 billion dollars in FY26	Record export figure
3 September 2026	HEC trains over 332 faculty in AI	Capacity building
28 September 2026	Public Procurement Rules 2026 notified	New federal procurement framework
20 September 2026	First wheeling auction request for proposals	Competitive electricity market opens to large consumers
27 October 2026	Senate informed of AI ecosystem and WACO membership	One billion dollar ambition; China-led body
November 2026	National Digital Masterplan reported approved	12 priority sectors; text not public
3 November 2026	This report	Public record assessment

APPENDIX B

Glossary

Algorithmic management. The use of automated systems to assign, monitor, rate or discipline workers.

Artificial intelligence system. A machine-based system that infers from inputs how to generate outputs such as predictions, content, recommendations or decisions. The definition used by the OECD and adopted in this report.

Automated decision-making. A decision that is wholly or partly generated by a system, without meaningful human involvement in the specific case.

Conformity assessment. The process of demonstrating that a system meets specified requirements, through testing, inspection, certification or a combination.

Data protection impact assessment. A structured assessment of the risks to individuals from a data processing operation, and the measures to mitigate them.

Data sharing impact assessment. The draft data policy's term for an assessment of a proposed inter-agency or third-party data exchange.

Digital public infrastructure. Foundational digital systems, such as identity, payments and data exchange, that enable public and private services.

High-impact use. A use whose failure, misuse, inaccuracy or unauthorised disclosure may materially affect rights, services, security, public trust or the continuity of essential functions.

Large language model. A model trained on large text corpora to generate or analyse language.

Model risk. The risk of loss or harm from decisions based on incorrect or misused model outputs.

Privacy-enhancing technology. A technical method that protects personal data while enabling analysis or verification, such as differential privacy, secure computation or zero-knowledge proofs.

Red teaming. Adversarial testing of a system to find failure modes, vulnerabilities or harmful outputs.

Regulatory sandbox. A controlled environment in which a regulator allows limited live testing of an innovation, with conditions and supervision.

Risk tier. A category that determines the level of obligation applied to a system, usually based on the severity and likelihood of harm.

WASL. Pakistan's National Data Exchange, a governed platform for inter-agency data sharing.

APPENDIX C

Data notes and calculations

This appendix records the assumptions behind the report's calculations and the caveats that apply to the data.

Training cost estimate

The report estimates the annual cost of the AI Policy's skills targets in Table 5. The calculation is:

- 200,000 trainees multiplied by a unit cost of Rs 15,000 to Rs 30,000, giving Rs 3.0 billion to Rs 6.0 billion.
- 20,000 internships multiplied by a stipend of Rs 15,000 a month for three months, giving Rs 900 million.
- 3,000 scholarships multiplied by Rs 300,000 a year, giving Rs 900 million.
- 400 projects multiplied by a maximum grant of Rs 1 million, giving up to Rs 400 million.
- 200 theses multiplied by a maximum grant of Rs 200,000, giving up to Rs 40 million.

The total range is Rs 5.24 billion to Rs 8.24 billion. The assumptions are illustrative. The actual cost depends on the length and quality of training, the stipend level, the scholarship value and the share funded by provinces, universities and employers. The comparison with the Rs 524.64 million proposed allocation is intended to show the order of magnitude, not to provide a budget estimate.

SBP survey

The State Bank's Financial Stability Review 2024 reported a survey of 55 regulated entities. About half reported deploying AI or developing it. The survey combines deployed and developing systems, covers regulated entities only and reflects 2024 conditions. It should not be read as a measure of operational AI across the financial sector, still less across the economy.

Start-up funding

The 133.7 million dollar figure for 2026 to the third quarter comes from i2i Ventures, as reported by TechJuice. It covers 10 disclosed deals. The third-quarter total of 74.7 million dollars includes a 68 million dollar round by Fasset, a UAE-headquartered fintech. Excluding that round, third-quarter funding was about 6.7 million dollars. The figures are disclosed deals only and understate informal and angel investment.

AFD index

The Artificial Intelligence Investment Potential Index 2025 ranked Pakistan 98th of 194 with a score of 51.11, against a global average of 52.32. India ranked 30th with 73.64. The index weighs economic environment, governance, infrastructure, data governance and statistical performance. It is a composite index and should be read alongside other measures, not as a definitive ranking.

Fiscal figures

FY26 runs from July 2025 to June 2026. FY27 runs from July 2026 to June 2027. The PSDP figures for FY27 are proposed allocations, not expenditures. Currency figures are in the currency used by the source. Where dollar figures are given, they are the source's figures, not conversions at a particular exchange rate.

Survey and index caveats

Composite indices use different definitions, weights and data vintages. Pakistan's rank can move for reasons unrelated to policy, including changes in other countries' scores. The report uses indices to locate Pakistan relative to peers, not to grade specific policies.

APPENDIX D

What the government should publish

The following documents would allow a citizen, journalist, researcher or investor to assess AI governance without relying on press releases. Most of them are already required by law or policy.

- A register of government AI systems, with purpose, owner, data, vendor, risk tier, human review and complaint route.
- The membership, terms of reference, minutes and annual report of the AI Council.
- The Pakistan Digital Authority's annual report, budget, staffing plan and decision log.
- The National Digital Masterplan, with implementation plans and budgets.
- The AI risk classification and impact assessment instruments.
- Data protection impact assessments for high-impact systems, with redactions.
- WASL coverage, transaction, consent and incident statistics.
- The National AI Fund's corpus, committee, calls and awards.
- AI procurement awards, contract values and model contract schedules.
- Sector regulators' AI guidance, inspection findings and enforcement decisions.
- The annual State of AI Governance report, debated in Parliament.

A government that publishes these documents will be easier to hold to account. It will also be easier to invest in, because uncertainty about the rules is one of the main reasons investors avoid emerging AI markets. Transparency is not only a democratic virtue. It is an economic one.

References

1. Government of Pakistan. *The Digital Nation Pakistan Act, 2025 (Act No. I of 2025)*. 29 January 2025.
   https://pakistancode.gov.pk/pdffiles/administrator2c2fd2fe7da1a657c40589a0705b3e20.pdf
2. Ministry of Information Technology and Telecommunication. *National Artificial Intelligence Policy 2025*. Approved 31 July 2025.
   https://moitt.gov.pk/SiteImage/Policies/National%20AI%20Policy.pdf
3. Computerworld Pakistan, reporting the federal PSDP 2026-27. *Federal PSDP 2026-27 Allocates Rs 16.29 Billion for IT Parks and AI*. 2026.
   https://cwpakistan.com/federal-psdp-2026-27-allocates-rs-16-29-billion-for-it-parks-and-ai/
4. Ministry of Information Technology and Telecommunication. *National Data Governance Policy 2026 (public consultation draft, DNP-D.001 POL, version 1.1)*. June 2026.
   https://moitt.gov.pk/SiteImage/Policies/Data%20Governance%20Policy%20June%202026.pdf
5. Ministry of Information Technology and Telecommunication. *Legislation register, including the draft Personal Data Protection Bill*. Accessed 3 November 2026.
   https://www.moitt.gov.pk/Legislations
6. Irfan Mir Halepota and Associates. *Data Protection for Businesses in Pakistan in 2026: PECA, Sector Rules and the Unfinished Personal Data Protection Bill*. 19 October 2026.
   https://irfanlaw.com/blog/data-protection-peca-businesses-pakistan.html
7. Jang. *Work begins on forming the first AI Council to implement the artificial intelligence policy*. 2026.
   https://jang.com.pk/news/1554039
8. Daily Times. *Pakistan forms digital authority, appoints Dr. Sohail Munir as chairman*. 20 August 2025.
   https://dailytimes.com.pk/1356603/pakistan-forms-digital-authority-appoints-dr-sohail-munir-as-chairman/
9. Pakistan Digital Authority. *Careers and vacancies, including Chief AI Officer*. Accessed 3 November 2026.
   https://pda.gov.pk/careers.php
10. State Bank of Pakistan. *Financial Stability Review 2024*. 2024.
    https://www.sbp.org.pk/FSR/2024/FSR-2024.pdf
11. Securities and Exchange Commission of Pakistan. *Regulatory Sandbox: what it is and how it operates.* Accessed 3 November 2026.
    https://www.secp.gov.pk/regulatory-sandbox/what-is-regulatory-sandbox/
12. National Judicial Policy Making Committee. *National Guidelines for the Use of Artificial Intelligence in Judicial Institutions*. 29 April 2026.
    https://ljcp.gov.pk/naig
13. Higher Education Commission. *Framework for the Ethical Use of Generative AI in Higher Education Institutions (draft, version 1.0)*. 23 April 2025.
    https://www.hec.gov.pk/english/services/faculty/Plagiarism/Documents/Final_Framework%20for%20Ethical%20use%20of%20Generative%20AI%20in%20HEI%27s-%2023April 2025.pdf
14. Radio Pakistan. *CM Punjab approves AI Roadmap for province*. 12 March 2026.
    https://www.radio.gov.pk/12-03-2026/cm-punjab-approves-ai-roadmap-for-province
15. Khyber Pakhtunkhwa Information Technology Board. *Digital Transformation Strategy 2030*. 2026.
    https://www.kpitb.gov.pk/Downloads/DigitalKP.pdf
16. The Express Tribune. *Sindh taps Huawei to build an AI-powered future*. 2026.
    https://tribune.com.pk/story/2635271/sindh-taps-huawei-to-build-an-ai-powered-future
17. Pakistan Telecommunication Authority data, reported by TechJuice. *Pakistan Mobile Subscribers Cross 204 Million as 4G Usage Rises*. 2026.
    https://www.techjuice.pk/pakistan-mobile-subscribers-cross-204-million-as-4g-usage-rises/
18. Computerworld Pakistan. *Pakistan 5G User Base Reaches 1.3 Million Across 23 Cities As Rollout Continues*. 2026.
    https://cwpakistan.com/pakistan-5g-user-base-reaches-1-3-million-across-23-cities-as-rollout-continues/
19. Associated Press of Pakistan. *Pakistan's technology exports hit record $4.6bn in FY26: Khurram*. 18 July 2026.
    https://www.app.com.pk/business/pakistans-technology-exports-hit-record-4-6bn-in-fy26-khurram/

20. i2i Ventures, reported by TechJuice. *Pakistan Startups Raise $133.7 Million in Funding During 2026*. October 2026.
https://www.techjuice.pk/pakistan-startups-raise-133-7-million-in-funding-during-2026/

21. Associated Press of Pakistan. *NTC launches sovereign AI infrastructure with GPUaaS, MaaS*. 2026.
https://www.app.com.pk/national/ntc-launches-sovereign-ai-infrastructure-with-gpuaas-maas/

22. Radio Pakistan. *Govt implementing national AI ecosystem, Senate informed*. 27 October 2026.
https://www.radio.gov.pk/27-10-2026/govt-implementing-national-ai-ecosystem-senate-informed

23. Organisation for Economic Co-operation and Development. *OECD Principles on Artificial Intelligence*. 2019, amended 2024.
https://oecd.ai/en/ai-principles

24. European Parliament and Council of the European Union. *Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (AI Act)*. 2024.
https://eur-lex.europa.eu/eli/reg/2024/1689/oj

25. National Institute of Standards and Technology. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. January 2023.
https://www.nist.gov/itl/ai-risk-management-framework

26. International Organization for Standardization and International Electrotechnical Commission. *ISO/IEC 42001:2023, Information technology, Artificial intelligence, Management system*. 2023.
https://www.iso.org/standard/81230.html

27. Ministry of Information Technology and Telecommunication. *Islamabad AI Declaration on Sovereign, Responsible, and Capability-Driven Artificial Intelligence*. 9 February 2026.
https://moitt.gov.pk/Detail/YmY4ZGFkODktYzI5Yy00OGUxLWE4ZDEtZDJjNzcwMWZhZGU1

28. Pakistan Public Procurement Regulatory Authority. *Public Procurement Rules, 2026*. Notified 28 September 2026.
https://epms.ppra.gov.pk/public/procurement-rules

29. GSMA. *Mobile Gender Gap Report 2025*. 2025.
https://www.gsma.com/mobilefordevelopment/mobile-gender-gap-report-2025/

30. Ministry of Information Technology and Telecommunication. *Pakistan Cloud First Policy*. 18 February 2022.
https://moitt.gov.pk/SiteImage/Policies/Pakistan%20Cloud%20First%20Policy-Final-25-02-2022%20(1).pdf

31. TechX Pakistan. *Pakistan Competitive Electricity Market Opens and Tech Businesses Can Now Buy Power Directly*. 2026.
https://techx.pk/pakistan-competitive-electricity-market-tech-data-centres/

32. United States Bureau of Industry and Security. *Framework for Artificial Intelligence Diffusion (rescinded May 2025)*. January 2025.
https://www.bis.gov/

33. Higher Education Commission. *HEC Trains Over 332 Faculty*. 3 September 2026.
https://www.hec.gov.pk/english/news/news/Pages/HEC-Trains-Over-332-Faculty.aspx

34. Press Information Department. *Artificial intelligence module at the Civil Services Academy*. January 2026.
https://pid.gov.pk/site/press_detail/31547

35. Agence Francaise de Developpement, reported by The News. *Artificial Intelligence Investment Potential Index 2025: Pakistan's AI potential stuck in third gear*. 2026.
https://www.thenews.com.pk/amp/1297009-pakistan-s-ai-potential-stuck-in-third-gear

36. Oxford Insights. *Government AI Readiness Index 2025*. 2025.
https://oxfordinsights.com/ai-readiness/government-ai-readiness-index-2025/

37. Government of Pakistan. *Constitution of the Islamic Republic of Pakistan, Articles 14 and 19A*. 1973, as amended.
http://www.pakistani.org/pakistan/constitution/

38. Supreme Court of Pakistan. *Mohtarma Benazir Bhutto v. President of Pakistan, PLD 1998 SC 388*. 1998.
https://www.supremecourt.gov.pk/

39. Supreme Court of Pakistan. *Justice Qazi Faez Isa case, PLD 2021 SC 1*. 2021.
https://www.supremecourt.gov.pk/

40. Government of Pakistan. *Prevention of Electronic Crimes Act, 2016, as amended to 2025*. 2025 consolidation.
https://www.sja.gos.pk/assets/Updated_Laws/The%20Prevention%20of%20Electronic%20Crimes%20Act,%20Rules%20Final%20Index%20(%20Upto%20date%202025).pdf

41. DLA Piper. *Data Protection Laws of the World: Pakistan*. Last modified 17 March 2026.
https://www.dlapiperdataprotection.com/?c=PK

42. State Bank of Pakistan. *BPRD Circular No. 01 of 2023: Framework on Outsourcing to Cloud Service Providers*. 2023.

https://www.sbp.org.pk/bprd/2023/C1.htm

43. Government of Pakistan. *Digital Nation Pakistan Act, 2025, sections 13 to 16 (Digital Nation Fund, accounts and audit)*. 29 January 2025.
https://pakistancode.gov.pk/pdffiles/administrator2c2fd2fe7da1a657c40589a0705b3e20.pdf

44. National Cyber Emergency Response Team of Pakistan. *Cybersecurity Handbook 2026-27*. 2026.
https://pkcert.gov.pk/handbook.asp

45. National Cyber Crime Investigation Agency. *Online harassment complaints and PECA guidance*. Accessed 3 November 2026.
https://www.nccia.gov.pk/harrasment.php

46. Ministry of Information Technology and Telecommunication. *Policy register, including approval status and dates for the National AI Policy and Data Governance Policy 2026*. Accessed 3 November 2026.
https://www.moitt.gov.pk/Policies

47. Ministry of Information Technology and Telecommunication. *Tender notice: Establishment of National Artificial Intelligence Hubs across Pakistan*. 2026.
https://www.moitt.gov.pk/SiteImage/Tenders/AI%20HUBS%20Advertisement.pdf

48. Electronic Procurement and Asset Disposal System. *Procurement notice P71677, consultancy services for national AI hubs*. 2026.
https://pa.epads.gov.pk/procurement/consultancy-services/62310/sbd

49. Punjab Information Technology Board. *Tender register, including consultancy for a government AI cloud*. 2026.
https://pitb.gov.pk/tendernotices?page=4

50. Securities and Exchange Commission of Pakistan. *First regulatory sandbox cohort: six solutions admitted for testing*. 2025.
https://www.secp.gov.pk/document/fintech-framework-for-the-non-banking-sector-in-pakistan/

51. Digital Rights Foundation. *DRF Analyses Pakistan's National AI Policy 2025*. 2026.
https://digitalrightsfoundation.pk/wp-content/uploads/2026/05/DRF-Analyses-Pakistans-National-AI-Policy-2025.pdf

52. Pakistan Digital Authority. *WASL National Data Exchange Specification (DNP-D.300 TS)*. 2026.
https://standards.pda.gov.pk/dnp-d.300ts

53. Pakistan Digital Authority. *WASL consent guideline (DNP-D.250 GDL)*. 2026.
https://standards.pda.gov.pk/dnp-d.250gdl

54. Pakistan Digital Authority. *Publication framework for DNP instruments (DNP-X.001 FWK)*. 2026.
https://standards.pda.gov.pk/dnp-x.001fwk

55. National Institute of Standards and Technology. *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile (NIST AI 600-1)*. July 2024.
https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf

56. UNESCO. *Recommendation on the Ethics of Artificial Intelligence*. 2021.
https://www.unesco.org/en/ethics-artificial-intelligence/recommendation-ethics

57. World Bank. *New Support for Punjab's Digital Transformation Will Help Expand Broadband Access and Improve Service Delivery*. 30 June 2026.
https://www.worldbank.org/en/news/press-release/2026/06/30/pakistan-new-support-for-punjab-s-digital-transformation-will-help-expand-broadband-access-and-improve-service-delivery

58. Radio Pakistan. *Modern IT hub completed in Dera Bugti under the Balochistan Social Development Initiative*. 3 November 2026.
https://www.radio.gov.pk/03-11-2026/modern-it-hub-completed-in-dera-bugti-under-bsdi

59. World Bank. *Pakistan country programme: digital economy and digital public infrastructure financing*. 2026.
https://www.worldbank.org/en/country/pakistan

60. Pakistan Digital Authority. *Standards and publications portal, DNP-X.001 publication framework and DNP-A AI series*. Accessed 3 November 2026.
https://standards.pda.gov.pk/dnp-x.001fwk

61. Radio Pakistan, statement by the Deputy Prime Minister to the Senate. *Pakistan joins the World AI Cooperation Organization*. 27 October 2026.
https://www.radio.gov.pk/27-10-2026/govt-implementing-national-ai-ecosystem-senate-informed

PIVRA

INDEPENDENT RESEARCH ON PAKISTAN

Evidence on Pakistan, for decisions that matter.

PIVRA is an independent, non-partisan research and policy institute based in Lahore, publishing open-access research, analysis and data on Pakistan and providing advisory and training services.

www.pivra.org
info@pivra.org
Lahore, Pakistan

© 2026 Pakistan Institute for Vision Research and Action.
All rights reserved.