Govt of Pakistan
Civil Services Academy

atomcamp

TRAINING MANUAL

A course on

ARTIFICIAL INTELLIGENCE FOR THE PUBLIC SECTOR

101, 201, 301 & 401



TECH destiNATION PAKISTAN
MINISTRY OF IT & TELECOM

Ministry of Planning Development & Special Initiatives

Message by the Minister for Planning, Development & Special Initiatives



Ahsan Iqbal Chaudhry

Minister of Planning, Development
& Special Initiatives
Government of Pakistan

I welcome this initiative aimed at enhancing understanding of Artificial Intelligence within the public sector. In an era of accelerating technological change and increasing demand for data-driven governance, strengthening the digital and analytical capacity of our civil services is essential for effective and informed public administration.

Pakistan's development strategy, encapsulated in the URAAN Pakistan, 5Es-based Five-Year National Economic Transformation Plan: Exports, E-Pakistan, Environment, Energy, and Equity, provides a roadmap for inclusive and sustainable growth. This framework highlights the importance of digital transformation and human capital development as critical components of socioeconomic progress, and it aligns with ongoing efforts to modernize public service delivery and policy implementation.

The training material presented in this manual will support civil servants in understanding the practical applications of AI, prompt engineering, and emerging technologies in governance. I encourage officers to engage deeply with this resource and apply these insights thoughtfully in their day-to-day roles, improving outcomes for citizens across Pakistan.

I commend the Civil Services Academy and its partners for this timely contribution to capacity building and look forward to the positive impact it will have on our collective efforts to advance responsive and forward-looking public administration.

Message by the Minister for Information Technology & Telecom



Shaza Fatima Khawaja

Minister for Information Technology & Telecom
Government of Pakistan

The formulation of Pakistan's National Artificial Intelligence Policy marks a significant milestone in our national journey toward responsible and inclusive digital transformation. The policy recognizes that while technology will shape the future of governance, real and lasting impact depends on the capacity of our institutions and the preparedness of our people.

A key pillar of the National AI Policy is therefore the upskilling of the public sector—ensuring that government officers at all levels are equipped to understand, use, and govern Artificial Intelligence with confidence and responsibility. Without this human and institutional readiness, even the most advanced technologies cannot deliver meaningful outcomes for citizens.

This training program at the Civil Services Academy represents a practical step in translating that policy vision into action. By introducing probationary officers to the fundamentals of AI, its applications in public administration, and the principles of ethical and responsible use, we are embedding digital competence at the very foundation of civil services training.

The objective is clear: to build AI-literate public leaders—officers who can make informed decisions, safeguard public interest, and leverage emerging technologies to improve governance, efficiency, and service delivery.

I commend the Civil Services Academy and atomcamp for institutionalizing this initiative and for setting a national example of how policy commitments can be transformed into sustainable institutional capacity.

Message by Director General Civil Services Academy (CSA)



Farhan Aziz Khawaja

Director General
Civil Services Academy (CSA)

It is with great pride that I announce the introduction of our new Artificial Intelligence (AI) module at the institute, an important stride in advancing the capabilities of our public sector workforce. In an era where data-driven governance and technological innovation are reshaping the way governments serve their citizens, understanding and applying AI is no longer optional; it is essential.

This module has been carefully designed to provide public officials with the knowledge, analytical tools, and ethical frameworks required to harness AI for evidence-based policymaking, improved service delivery, and enhanced administrative efficiency.

Through this initiative, we aim to strengthen institutional capacity and support the national vision for a digitally empowered and citizen-centric public service.We are living through a decisive shift. Artificial Intelligence is no longer a distant concept but a force already reshaping how economies function, how institutions operate, and how citizens experience the state.

As public servants, you stand at the forefront of transformative change. I urge all participants to approach this program not merely as a training exercise, but as an opportunity to reimagine the future of governance through innovation, collaboration, and integrity. The responsible use of AI can help us anticipate challenges, design inclusive solutions, and ensure that technology serves the greater public good. Let this module inspire you to lead with foresight and purpose, shaping a public sector that is adaptive, transparent, and responsive to the evolving needs of our society.

Message by the Secretary, Ministry of Planning, Development & Special Initiatives



Awais Manzur Sumra

Secretary Ministry of Planning,
Development & Special Initiatives
Government of Pakistan

The Government of Pakistan remains committed to strengthening the capacity of public sector institutions through the responsible adoption of emerging technologies that enhance efficiency, transparency, and quality of service delivery. Artificial Intelligence is no longer a future concept; it is a present capability that must be understood, governed, and applied with care.

This training manual represents a collaborative initiative of the Ministry of Information Technology and Telecommunication, the Civil Services Academy, and atomcamp, developed to introduce Civil Services Academy probationers to the foundational principles and practical implications of Artificial Intelligence in governance.

The focus of this initiative is not experimentation, but structured and informed adoption—positioning AI as a decision-support tool that operates within established governance frameworks, institutional protocols, and principles of public accountability. This document is the first in a planned series aimed at building progressive AI literacy across the civil services. It establishes a baseline understanding of AI and prompt engineering, with emphasis on real administrative use cases, responsible deployment, and the necessity of human oversight.

i

t is my expectation that officers engaging with this material will approach it with professionalism and critical judgment, recognizing that while AI can enhance productivity and analytical capacity, responsibility for decisions and outcomes always rests with the public servant. This initiative marks an important step toward developing a more capable, digitally empowered, and future-ready civil services.

Message by Secretary, Ministry of IT & Telecommunication



Zarrar Hasham Khan
Secretary Ministry of Information
Technology & Telecommunication
Government of Pakistan

This training manual represents a collaborative initiative of the Ministry of Information Technology & Telecommunication (MoITT), the Civil Services Academy (CSA), and atomcamp.

The Government of Pakistan is committed to strengthening the capacity of its public sector institutions by responsibly adopting emerging technologies that enhance efficiency, transparency, and service delivery. Artificial Intelligence is no longer a future consideration; it is a present capability that must be understood, governed, and applied with care.

The focus of this initiative is not experimentation, but structured adoption—ensuring that AI is used as a decision-support tool aligned with governance standards, institutional protocols, and public accountability.

This document is the first in a planned series designed to build progressive AI literacy across the civil services. It lays the foundation by introducing core concepts of Artificial Intelligence and prompt engineering, with an emphasis on real administrative use cases, responsible usage, and human oversight. Subsequent materials will continue to deepen practical capability while addressing policy, security, and implementation considerations.

It is my expectation that officers engaging with this material will approach it with professionalism and critical judgment, recognizing that while AI can enhance productivity, responsibility for decisions and outcomes always rests with the public servant.
This initiative marks an important step toward a more capable, digitally empowered, and future-ready civil services.

Message by Co-Founder atomcamp



Naveed Iftikhar

Co-founder atomcamp

It is a great pleasure to join hands with the Civil Services Academy, Lahore, and the Ministry of Information Technology & Telecom in this important initiative to upskill the next generation of civil servants and decision-makers.

We are living through a decisive shift. Artificial Intelligence is no longer a distant concept but a force already reshaping how economies function, how institutions operate, and how citizens experience the state.

From predictive systems in public health and agriculture to AI-assisted policymaking and data-driven service delivery, emerging technologies are redefining what effective, transparent, and accountable governance looks like. In this era, data is a strategic resource and AI is the capability that turns that data into insight, foresight, and action across every sector of our national life.

This moment demands public leaders who understand AI not as a buzzword, but as a governance priority. They must know how algorithms influence decisions, how to regulate and deploy these technologies responsibly, and how to ensure that AI becomes a driver of inclusive development rather than inequality. Through this collaboration, we are laying the foundation for an empowered, AI-literate leadership that can ask the right questions, make informed decisions, and guide responsible innovation within government. The journey toward a smarter, more resilient, and more sustainable future begins now - and I am honoured to be part of it.

Acknowledgement



Syed Shabbir Akbar Zaidi

Director (CTP/CB)
Civil Services Academy (CSA)
Walton, Lahore

I am pleased to formally extend my profound appreciation to our respected Director General, CSA, for his visionary leadership and strategic foresight in introducing the Artificial Intelligence (AI) module into our training framework. This initiative marks a significant step toward modernizing our curriculum and ensuring that our probationers are equipped with the skills required to navigate an increasingly technology-driven environment.

I would also like to acknowledge the faculty of the Program Wing, CB Wing, and the IT Team for their dedicated efforts, technical competence, and collaborative spirit in supporting the development and implementation of this module. Your contributions reflect commendable professionalism and a strong commitment to academic excellence.

Additionally, I wish to offer special thanks to Dr. Muqeem ul Islam, Director General (KIMS), for his invaluable support, and to Dr. Naveed Iftikhar, Co-Founder atomcamp, for his expert guidance and partnership in strengthening the AI learning experience for our trainees. Their contributions have added immense value to this initiative and further enriched the quality of the program.

Collectively, these efforts demonstrate our institution's commitment to innovation, capacity-building, and continuous improvement in public service delivery and serving the nation. I extend my sincere appreciation to all individuals and departments involved for their unwavering dedication and exemplary teamwork.

Foreword

Artificial Intelligence is no longer a future technology. It is reshaping how governments deliver services, how economies generate value, and how nations exercise power on the global stage. For Pakistan, the question is not whether to engage with AI but how to do so on our own terms, preserving sovereignty over our citizens' data, our critical systems, and our national decision-making.

This document is written for civil servants, policy makers, and public sector leaders who must make decisions about AI infrastructure, procurement, and deployment. It does not assume technical expertise. It does assume that the reader needs to understand the foundational hardware realities that underpin every AI policy decision, from purchasing a GPU cluster to designing a data sovereignty framework.

The document covers the full landscape: what GPUs are and why they matter; how data centers are built, governed, and financed; Pakistan's existing infrastructure and what is missing; the global market of AI chip providers; training versus inference and why that distinction is central to sovereignty; Digital Public Infrastructure and how AI fits into it; and concrete recommendations for Pakistan's path forward.
The authors draw on Pakistan's National AI Policy 2022, the Personal Data Protection Bill, the Digital Pakistan Vision, and international benchmarks from India, Estonia, Brazil, the UAE, and the European Union.

Note on the AI Training Series

This manual is the first volume in a structured, multi-stage national training series on Artificial Intelligence for the public sector. We will publish the next three manuals soon.

The series has been designed as a progressive capability-building pathway —moving from foundational awareness to advanced institutional readiness. Civil servants develop not only understanding, but practical and strategic competence in AI for governance.

Structure of the Series

The training is categorized in four levels:

AI 101
(already published)

Foundations of Artificial Intelligence for the Public Sector

AI 201
(current)

Applied AI Sytems, Governance, & Secure Deployment in the Government

AI 301

Government-grade Models, Fine-tuning, & AI Infrastructure

AI 401

National AI Infrastructure, Data Centers, & Digital Sovereignity

Training Manual
AI 101

Foundations of Artificial Intelligence for the Public Sector

TECH destiNATION PAKISTAN
MINISTRY OF IT & TELECOM

Ministry of Planning Development & Special Initiatives

atomcamp

Training Manual AI 101

Table of Contents

COURSE OVERVIEW ---------- 12

MODULE 1 : INTRODUCTION TO AI ---------- 14

MODULE 2 : INTRODUCTION TO LARGE LANGUAGE MODELS (LLMS) ---------- 31

MODULE 3 : INTRODUCTION TO PROMPT ENGINEERING ---------- 47

MODULE 4 : BEGINNER HANDS-ON AI TOOLS FOR PRACTICAL APPLICATION ---------- 52

Annexure ---------- 59

Course Overview

This two-day beginner-level instructional program provides probationary officers at the Civil Services Academy with a structured introduction to Artificial Intelligence (AI) and prompt engineering, with a specific emphasis on their application in governance and public sector management. While primarily designed for civil servants graduating from the Civil Services Academy, the program is equally applicable for public sector officials across ministries, departments, public sector training academies, and government organizations seeking a foundational understanding of AI for administrative and policy-related work.

The curriculum covers the principles, uses, and limitations of AI and trains officers to apply no-code tools and prompt engineering strategies to generate precise, context-aware, and administratively appropriate outputs. The program emphasizes the responsible and informed use of AI, ensuring that officers understand both the capabilities and constraints of such systems in official settings.

The accompanying manual serves as a guide to support officers and public sector officials in adopting AI concepts and tools in a manner consistent with professional standards, ethical considerations, data confidentiality requirements, and established government protocols.

Duration : 2 days

Target Audience:

- Probationary officers at the Civil Services Academy,
- Civil servants involved in future policy making, administration, office management, and public service delivery,
- Public sector officials from ministries, departments, public sector training academies,
- Autonomous bodies seeking to build foundational capacity in Artificial Intelligence.

Prerequisites: Basic computer literacy and familiarity with text-based interfaces

Learning Objectives

01 AI Understanding & Application

02 Master Prompt Engineering Techniques

03 Execute Government-Specific Tasks Using AI

04 Ensure Ethical, Responsible, and Secure Use

05 Apply Learning Through Practical Exercises

MODULE 1 : INTRODUCTION TO ARTIFICIAL-INTELLIGENCE

What is AI?
Artificial intelligence (AI) is a system that learns from data and uses that learning to make predictions or recommendations or actions. Instead of following fixed rules, AI improves its output by finding patterns in past information and applying them to new situations.

Data is the foundation of AI. This data can be text, images, audio, video, numbers, or records of human activity. AI systems convert this data into mathematical signals and learn which patterns are significant and which are not.

After training, an AI system can classify information, predict outcomes, undertake actions, or support decisions. For example, it can estimate the demand of a commodity, detect unusual behavior in a public space, or suggest the next best action regarding a legal issue based on previous cases.

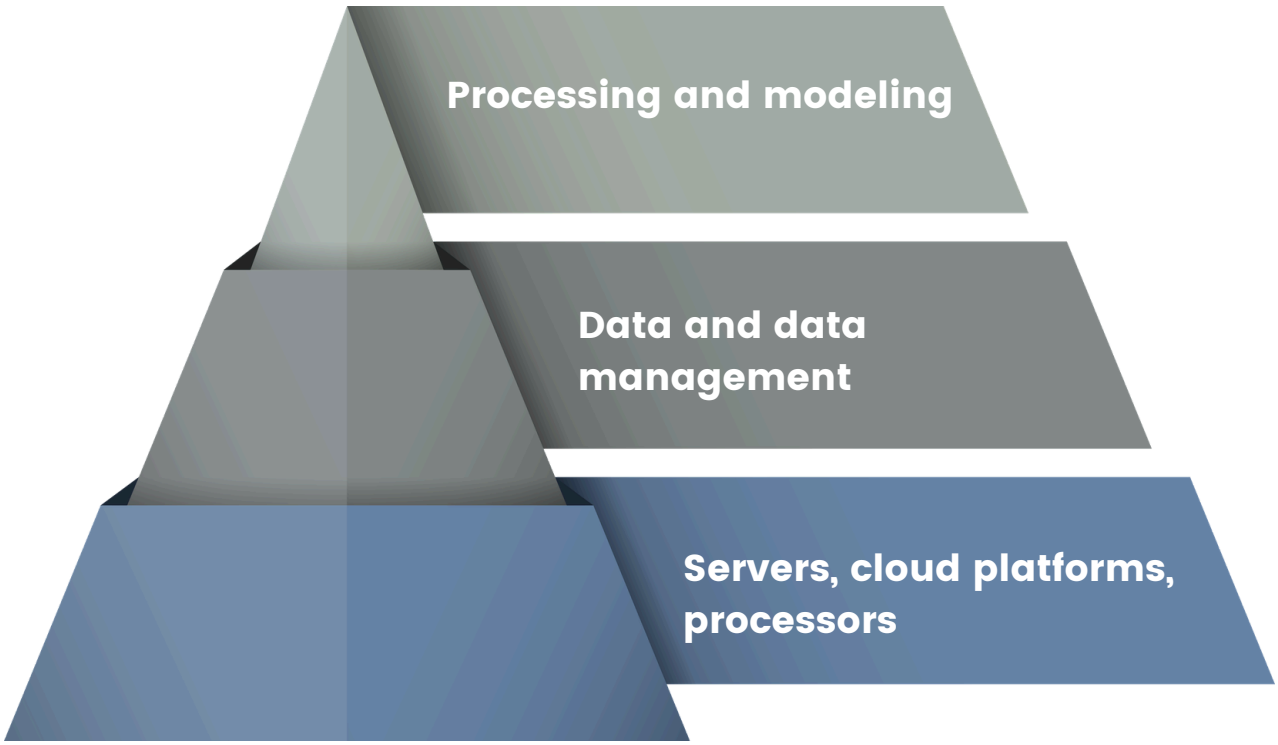
For public service use, it is important to understand that AI does not think or exercise judgment. It produces outputs based on the data it has seen, the assumptions built into its design, and the objectives it is given. This means AI can support officials by improving speed, consistency, and analysis, but responsibility for decisions, accountability, and public outcomes always remains with the human authority.

AI is therefore best used as a decision-support tool, not a decision-maker. Its value depends on the quality of data, clarity of objectives, and the controls put in place to ensure lawful, fair, and transparent use.

AI functions through a small set of interdependent layers. The first is computing infrastructure, which includes servers, cloud platforms, and processors that provide the raw power needed to train and run AI systems. Without adequate computing capacity, AI models cannot be trained efficiently or deployed at scale.

The second layer is data and data management. This covers how data is collected, cleaned, stored, and governed. High-quality, well-managed data is essential because AI outputs are only as reliable as the data they learn from.

The third layer is processing and modeling, where algorithms analyze data to produce predictions or recommendations. This layer turns raw data into usable outputs, with humans responsible for defining objectives, validating results, and applying judgment in real-world contexts.



AI is not new

AI did not start in 2022 or 2023. It is not something that suddenly came out of nowhere.

AI stands on hundreds of years of work in:

- Mathematics
- Statistics
- Programming
- Logic and reasoning

In 1955, John McCarthy gave the formal term “Artificial Intelligence.” Since then, many researchers have developed ideas about how machines can learn, solve problems, and make decisions.

The real jump happened in the 1990s, when the internet, the web, and computing power grew rapidly. After that:

- Better hardware
- Faster chips
- Cheaper storage
- Cloud computing

All these things made it possible to build AI systems at a large scale. So, AI is the result of decades of work.

What is new today is not the idea of AI, but its accessibility and practical use. Systems that once existed only in research labs are now embedded in everyday tools used by governments, businesses, and citizens, making AI a direct operational factor in policy, regulation, service delivery, and public accountability.

AI will reshape how people work, how services are delivered, and how businesses compete by changing costs, speed, and decision-making across the economy. Policymakers must therefore understand AI to deliberately maximize public value while managing risks related to jobs, equity, safety, and trust.

Digitilisation and AI



Digitisation

Converting analog information into digital form.



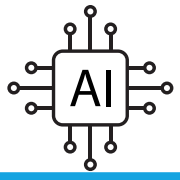
Digitalisation

Using digital technologies to improve existing processes.



Digital Transformation

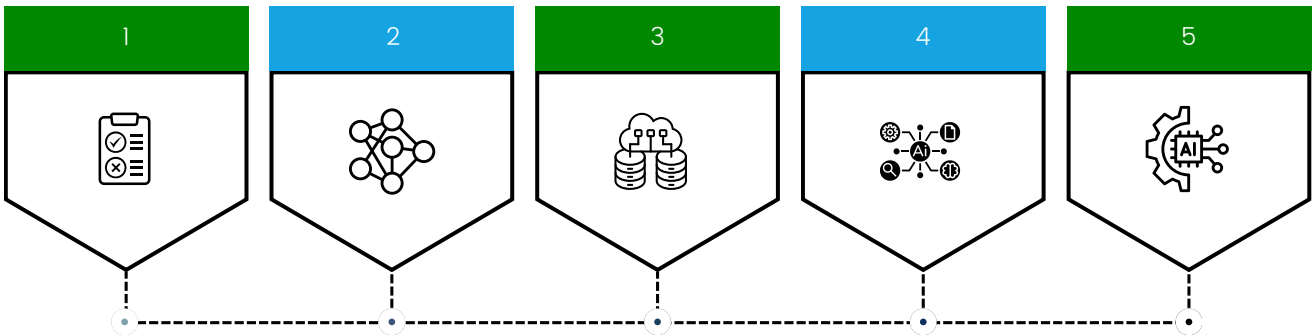
Reinventing how an organisation operates using digital tools — changing culture, processes, and citizen services.



Artificial Intelligence

Machines that can learn, reason, and perform tasks that normally require human intelligence.

Evolution of AI



Rule-Based Systems

- Humans wrote explicit "if-then" rules for machines.
- Examples:Expert systems for medical diagnosis.

Machine Learning Emerges

- Instead of hand-coding rules, computers learn patterns from data.
- Examples:Spam email filters.

Deep Learning & Big Data

- Neural networks with many layers ("deep") learn from massive datasets.
- Global Examples:
- (Google Photos, self-driving cars).

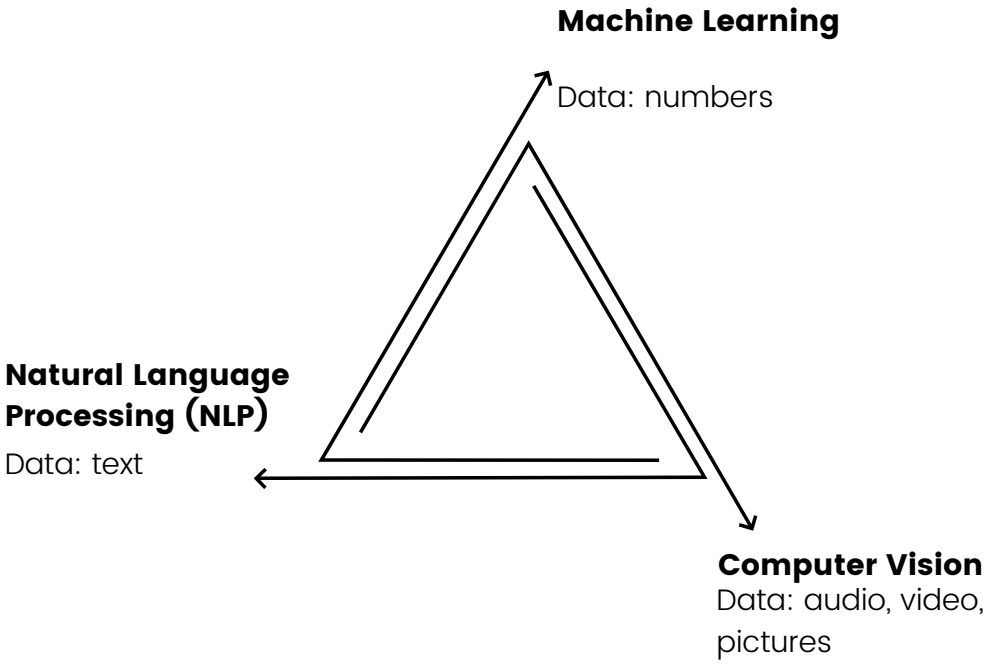
Generative AI

- AI can now create text, images, audio, and even code.
- Global Examples: ChatGPT, DALL·E, Google Gemini.

Agentic AI

- AI systems that not only answer but act on your behalf — connecting across tools, automating workflows.
- Global Examples: AutoGPT, Microsoft Copilot, Lindy.ai.

Subfields of Artificial Intelligence



Different AI sub-fields primarily work with different forms of data. Machine learning often relies on structured numerical data, natural language processing focuses mainly on text and speech, and computer vision works with images and video.

Foundation models & LLMs

Advanced AI systems are built on Foundation Models (FMs). These are large, general-purpose models trained at scale and adapted across many tasks, marking a shift away from narrow, task-specific machine learning.

Large Language Models (LLMs) are one prominent class of foundation models, specialized in understanding and generating human language.

Other equally critical foundation model classes include vision models for interpreting images and video, multimodal models that combine text, images, and video, and scientific models used for climate forecasting, genomics, and materials discovery. Together, these models constitute the current frontier of AI capability.

Why AI Matters for Government, Policy, and Society

Artificial Intelligence is not only changing how tasks are performed; it is reshaping how work is organized, how public services are delivered, and how societies function.

Across government operations, AI is increasingly used to support decision-making, improve service delivery, manage large volumes of information, and increase administrative efficiency. Beyond government, its influence extends to healthcare, education, finance, industry, culture, and the future of work.

For policymakers and civil servants, understanding AI is no longer optional. Effective governance now requires a clear understanding of how AI systems work in practice, where their limitations and risks lie, and how quickly technological change can outpace existing laws, regulations, and institutional processes.

Policymakers need to understand AI's societal and sector-specific impacts, helping participants anticipate change, assess policy consequences, and adapt regulations in a way that protects public interest while remaining practical and future-oriented.

As governments begin to use AI at scale, the role of the state expands beyond adoption to stewardship. Public institutions must decide where AI can assist human judgment, where it should be limited, and where it should not be used at all. These choices affect public trust, accountability, fairness, and long-term state capacity. Clear policy direction is therefore essential to ensure that AI strengthens institutions rather than undermines them.

Why Does Data Matter?

Everything in AI depends on data quality, this theme appears throughout the document. AI is only as good as the information it learns from.

Example: If you train a model on messy customer data, the model will make messy predictions.

The three most important skills for AI (Large Language Models) users
These three skills make anyone effective with AI - no coding required:

a. Asking Clear Questions (Prompting)
Clear instructions lead to clear answers

Example:
Instead of: “Explain finance.”
Try: “Explain basic personal finance to a beginner in 5 bullet points.”

b. Checking and Verifying Responses
AI sometimes produces inaccurate or fabricated details. Users should verify important information.

Example:
If AI generates a medical explanation, cross-check it with a trusted source.

c. Using AI with Your Own Context
Context lets AI read your documents so answers are accurate and grounded.

Example: Upload your company policies → ask questions → AI responds using the exact document.

Everyday Examples of AI

Here are simple, practical examples you can relate to:

- Email Assistance: AI drafts replies, but you review tone and correctness.
- Document Summaries: AI summarizes PDFs, but you confirm important details.
- Language Help: AI rewrites text in simple English for clarity.
- Data Lookup: AI answers questions using your uploaded files

AI Safety Basics

AI literacy also includes safe usage:

- Don’t paste confidential data into public models.
- Always double-check critical information (legal, medical, financial).
- Understand model limitations - AI is not a final authority.

Never upload confidential summaries, meeting minutes and documents on LLMs like ChatGPT, Gemini or Deepseek.

How AI Connects to Machine Learning and Deep Learning

Now that you understand AI at a high level, it's important to see how the different parts of AI fit together. Artificial Intelligence is a broad field. Under it, we have Machine Learning (ML) and Deep Learning (DL), which are simply different ways of building intelligent systems.

Think of it like this:

- AI → The big goal: making computers think, learn, and make decisions.
- Machine Learning → A major method inside AI that allows computers to learn patterns from data.
- Deep Learning → A more advanced type of ML that uses layered neural networks to understand complex patterns like images, speech, and language.

Why Machine Learning and Deep Learning Matter for AI Today

AI used to rely on rules written by humans. But as you saw in earlier sections, modern AI works because systems can learn instead of being manually programmed. That learning ability comes from ML and DL.

What is Machine Learning and Deep Learning?

To make it simple:

Machine Learning

Machine Learning means the computer learns patterns from data instead of us telling it every single rule.

Examples:

- Predicting budget spend
- Predicting population growth

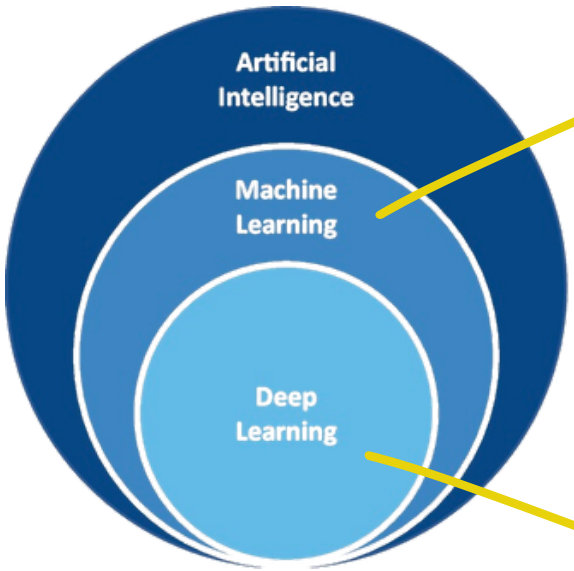
The machine looks at past data and finds patterns.

Deep Learning

Deep Learning is a special type of Machine Learning that uses many layers of computation. These layers help the system understand complex things like:

- Images
- Speech
- Videos
- Natural human language

Deep Learning is why we have modern AI today.



Machine learning (ML) is a type of artificial intelligence (AI) that trains algorithms on data to find patterns, make predictions, and improve performance on tasks without being explicitly programmed for every scenario, enabling systems to learn from experience

Deep learning is a subset of machine learning that uses multi-layered artificial neural networks, inspired by the human brain, to learn complex patterns from vast amounts of data, enabling computers to perform tasks like image recognition, speech transcription, and language understanding with minimal human intervention, powering modern AI applications.

AI Capabilities for Different Use Cases

AI Capability	What Makes It Work
ChatGPT answering questions	Deep Learning + LLMs
Predicting budget spend	Machine Learning
Face detection in safe city cameras	Deep Learning
Email spam detection	Machine Learning
Image generation	Deep Learning

One Simple Example to Connect Them All

Imagine you want a system that identifies whether a photo contains an authentic ID

- Using AI: You want the computer to “recognize” the person using his ID card photo
- Using Machine Learning: You gather thousands of labeled images and teach the system patterns of what a authentic ID card looks like.
- Using Deep Learning: A neural network automatically learns much more complex details -text positions, photo details- without you explicitly programming any of it.

Practical Use of Artificial Intelligence in Government Work

This section explains where and how AI can be practically used in government operations, with a focus on tasks civil servants already perform.

Drafting and Documentation

Civil servants spend a large part of their time preparing:

- Office notes
- Briefs for seniors
- Speeches and statements
- Summaries of meetings

How AI helps

- Producing first drafts of notes and summaries
- Converting bullet points into formal language
- Rewriting documents for clarity and structure

Good practice

- Use AI for the first draft only
- Always review, verify, and finalize yourself
- Never submit AI-generated text without human approval

Research and Policy Support

Officers frequently need:

- Background information on issues
- Comparisons of international practices
- Summaries of laws, reports, and studies

How AI helps

- Summarizing lengthy documents
- Structuring policy briefs
- Listing pros and cons of options
- Generating issue briefs for meetings

Good practice

- Cross-check all facts
- Use official sources for final validation
- Treat AI output as a research assistant, not a source of record

Administrative Efficiency

Routine work such as:

- Email drafting
- Scheduling support
- Preparing standard replies
- Formatting documents

How AI helps

- Drafting routine correspondence
- Creating templates for repetitive tasks
- Structuring reports and presentations
- More time to focus on: analysis, oversight, field engagement, and decision-making

Using AI in Service Delivery

AI can support—not replace—citizen-facing services.

Examples

- Chatbots for basic information
- Automated responses for common queries
- Ticket classification in grievance systems
- Support for helplines and facilitation desks

What Must Always Remain Human

- Final decisions on cases
- Complaint resolution authority
- Disciplinary and legal actions
- Any matter affecting rights, benefits, or penalties

Rule:

AI can assist service delivery.
It must never replace accountability.

From Perception AI → Generative AI → Agentic AI

Module 01: AI 101

AI has evolved in stages. Each stage added new abilities and allowed machines to work in more intelligent ways. Here's a simple breakdown to help you understand how we reached today's modern AI systems.

What it can do:

- See objects in images
- Recognize faces
- Detect suspicious activties from video feeds.
- Hear and transcribe speech
- Classify documents or sounds

Examples:

- Phone unlock with face recognition
- Google Photos detecting "cat," "mountain," or "birthday"
- Security cameras detecting motion
- Speech-to-text apps convert voice into written words

Perception AI understands, but it does not create.

2. Generative AI (The "Creating" Era)

This is the second wave. Generative AI can write, create, and generate new content based on patterns it learned.

What it can do:

- Write essays, emails, and summaries
- Generate images, videos, and voice
- Produce code
- Answer questions
- Translate languages
- Draft reports or documents

Examples:

- ChatGPT is writing a lesson plan or summary
- Midjourney is generating an artwork
- GitHub Copilot writes code automatically
- AI tools are creating marketing posts or product descriptions

Generative AI can produce new content, not just classify it.

3. Agentic AI (The “Doing” Era)

Agentic AI is the newest stage. These systems don’t just create content - they can take actions, use tools, and complete tasks on your behalf.

What it can do:

- Book meetings
- Search the internet and gather information
- Send emails
- Update databases
- Run workflows or automations
- Follow multi-step goals

Examples:

- An AI assistant that reads your email, drafts responses, and sends them when approved
- An AI agent helping/answering calls citizens on motorway and highways.
- An AI agents that routes citizen complaints to relevant departments ,follow up on actions and reports back with progress.
- An AI finance agent that categorizes expenses and generates monthly reports
- AI that takes a goal like “Prepare a 5-slide summary from this PDF” and performs multiple steps automatically

Agentic AI can plan, reason, and act, not just see or create.

AI vs Rule-Based Automation

Many people confuse AI with automation. Rule-based automation follows fixed, human-written instructions: if X happens, do Y. It cannot learn or improve - everything depends on predefined rules.

AI, on the other hand, learns patterns from data instead of relying on hard-coded logic. It adapts, improves with more data, and can handle complex, unpredictable situations where manual rules would fail.

Rule-Based Automation

- You write fixed rules
- If X happens, do Y
- No learning
- No improvement
- Everything depends on human-defined steps

Example - Rule based vs AI:

The rules are already described in government systems to transfer payroll on monthly basis. A simple automation executes this workflow.

AI system is when you use the historical pattern to detect fraud. Over time the system learns how much is transferred from government accounts, how much should have been transferred. By learning this patterns, AI can detect fraudulent activities.

Similarly in Auditing, an AI agent/system can scans huge data by quickly reading millions of records at once, spotting patterns and highlighting anything that looks unusual compared to normal behavior. Humans must review and judge these alerts to confirm real fraud and ensure fairness.

Module 2 : Introduction to Large Language Models (LLMs)

Large Language Models are AI systems that understand and generate human language. They are trained on large amounts of text, such as books, articles, websites, and code. Through this training, they learn how language works in context, which allows them to write, summarize, and answer questions in natural language.

LLMs work by predicting the next word based on previous words. At scale, this enables useful language-based assistance, but it does not mean the model truly understands facts or intent.

How LLMs Are Developed

- Large volumes of text data are collected and cleaned
- The model is trained to predict the next word
- Deep learning layers capture language patterns
- Training requires high-performance computing (GPUs)
- Models are then fine-tuned for safety and quality
- Deployed on cloud or on-premise systems

W**hat Policymakers Need to Understand**

LLMs learn patterns in language, not verified facts. They are powerful support tools, but their outputs must be reviewed and governed, especially in public sector use.

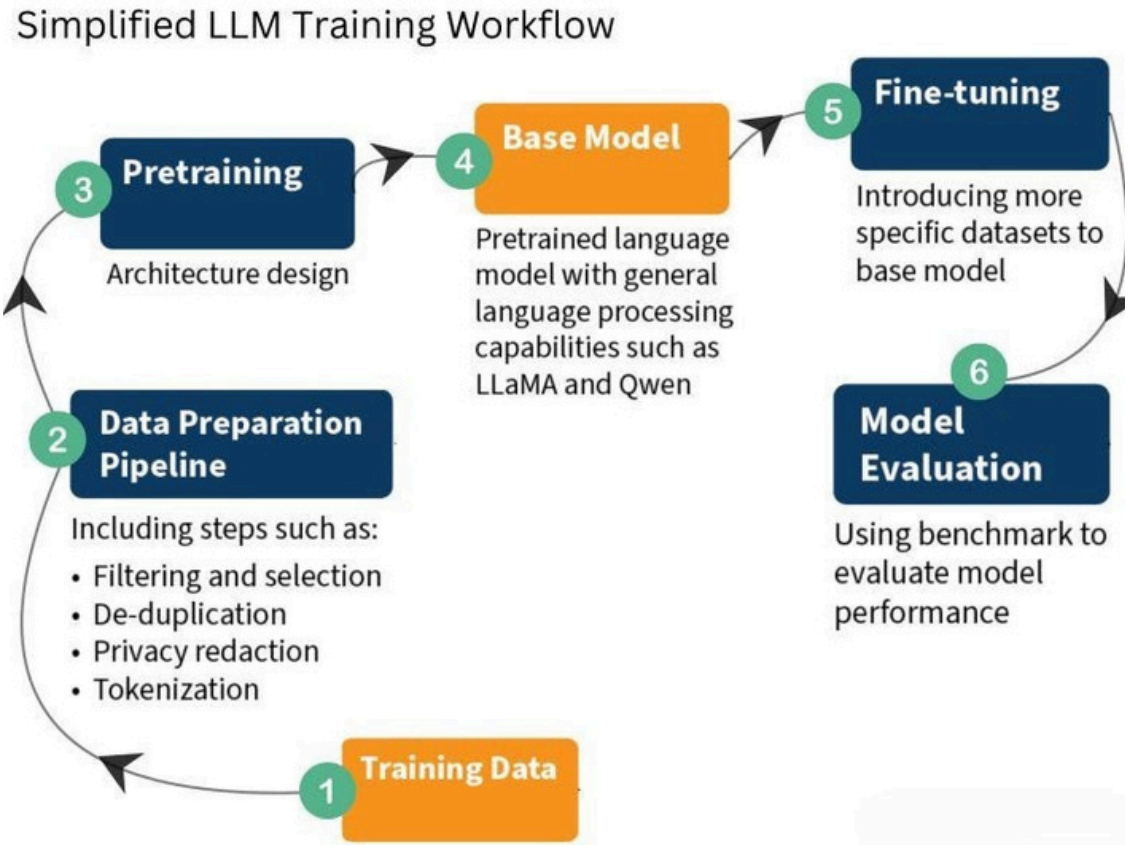


Figure 2 : How LLMs are Trained

Use Case	Key Benefit
Policy Drafting	Faster document creation with a consistent official tone
Legal & Official Letters	Improved grammar accuracy and standardized formatting
Citizen Services	Chatbots and multilingual support for public assistance
Data Analysis	Rapid summarization and trend extraction from large datasets
Training & Upskilling	Personalized learning and skill development content
Digital Transformation	Automation of repetitive documentation and workflows

How Governments Should Use LLMs: Simple Maturity Path

LLMs are trained on generic, global data. Government value comes from how they are used, not from building new models on day one. Risks (hallucination, bias, leakage) reduce as usage maturity increases.

LLM Usage Maturity (Low to High)

1.Better Prompting — Fastest Gains

Clear, structured prompts reduce errors and hallucinations

2. RAG (Context Injection) — Controlled Outputs

Uses official laws, policies, and documents as context. It can improve accuracy and reduce hallucination.

3. Fine-Tuning — Domain Specialization

Trains models on curated government data. It requires more financing and compute . Government departements will have to move towards fine tuning for optimal use of LLMs.

4. Sovereign LLMs — Strategic Endgame

Full control over data and compliance. It requires long-term investment in compute and talent.

Policy takeaway

Start with usage discipline, progress to context control, and pursue model ownership only when readiness exists.This is the strategic path that the government of Pakistan and provincial governments should adopt.

Designing AI Systems for Government Use

Retrieval-Augmented Generation (RAG)

Retrieval-Augmented Generation (RAG) is a method where an AI system retrieves information from approved documents or databases before generating a response.

Instead of relying on general knowledge, the AI:

1. Searches trusted government data sources
2. Retrieves relevant sections
3. Uses only that information to generate its response

This ensures outputs are grounded in official, up-to-date, and verifiable data.

Government Use Examples

- Querying service rules, acts, or regulations
- Answering internal questions using departmental SOPs
- Supporting officers with policy look-ups
- Creating AI assistants for ministries that only reference approved documents

Key Takeaway

RAG ensures AI systems do not “invent” answers, but instead respond based on government-approved knowledge.

Fine-Tuning

Fine-tuning is the process of training an AI model further on domain-specific data so that it consistently reflects a particular style, structure, or subject area.

While general AI models are trained on broad public data, fine-tuning adapts them to:

- Government language
- Sector-specific terminology
- Standard formats used in official work

It shapes how the AI responds, not what policy decisions are made.

Government Use Examples

- Generating standardized notices, letters, or reports
- Assisting with routine documentation in ministries
- Supporting sector-specific functions (health, finance, education)
- Creating internal AI tools trained on historical government documents

Key Takeaway

Fine-tuning helps AI systems sound and behave like a government assistant, not a generic chatbot.

Data Sovereignty, Confidentiality, and In-House AI Systems for Government

Why Data Sovereignty Matters in Government AI

Data sovereignty refers to the principle that government data must remain subject to national laws, regulatory oversight, and institutional control, regardless of how it is processed or analyzed.
For government entities, this includes:

- Citizen records
- Policy drafts and cabinet summaries
- Financial, security, and administrative data
- Internal communications and reports

When such data is processed by externally hosted AI systems, control over storage, usage, and retention may be lost.

Key Risks of Using Foreign-Hosted AI Tools

Using online, externally developed AI platforms for government work can introduce:

- Loss of control over sensitive data
- Unclear data retention and reuse policies
- Jurisdictional risks (data stored outside national boundaries)
- Exposure to model training on confidential inputs
- Non-compliance with national security and privacy laws

Even when tools claim not to store data, verification and auditability are limited.

Principle of In-House and Organization-Specific AI Systems

In-house AI systems are:

- Deployed on government-controlled infrastructure
- Accessed only by authorized personnel
- Trained or configured using approved internal documents
- Governed by existing IT, security, and data policies

Such systems function as internal digital assistants, not public AI services.

Core Design Principles

Government AI systems should follow these principles:

- Data Never Leaves the Organization
  - All documents, queries, and outputs remain within official systems.

- Use Only Approved Internal Sources
  - AI systems must reference:
  - Acts, rules, notifications
  - Departmental SOPs
  - Official reports and datasets

- No Training on External or Public Data by Default
  - The system is contextualized using internal documents only.

- Role-Based Access Control
  - Different officers see only the data relevant to their mandate.

- Human Oversight by Design
  - AI supports officers; final authority remains human.

Building Confidential AI Agents Using Internal Government Data

Step 1: Identify Approved Knowledge Sources

Before any AI system is built, departments must define:

- Which documents are approved for AI use
- Classification levels (public, internal, confidential)
- Update and version-control mechanisms

This ensures data governance precedes AI deployment.

Step 2: Use Retrieval-Based Approaches (Not Public Training)

Instead of uploading documents to public AI tools, organizations should:

- Store documents in secure internal repositories
- Use retrieval-based systems where AI fetches relevant information at runtime
- Prevent models from learning or memorizing confidential data

This approach keeps data contextual but not exposed.

Step 3: Develop Organization-Specific AI Agents

Each department can develop AI agents tailored to its function:

- Policy drafting assistants
- Legal and regulatory lookup systems
- Service delivery support tools
- Internal reporting and briefing assistants

Step 4: Developing Pakistan's Own Local and Domain-Specific LLMs

For true data and linguistic sovereignty, Pakistan should invest in the development of local large language models (LLMs) that:

- Are trained and hosted entirely within Pakistan
- Support local languages and administrative terminology
- Reflect national legal, regulatory, and cultural context
- Are aligned with public-sector use cases rather than consumer applications

Such models would not replace global research efforts but would provide trusted, sovereign AI foundations for government use.

Understanding AI Models for Government and Official Work

Modern governance increasingly integrates Large Language Models (LLMs) to support drafting, policy analysis, research synthesis, and structured communication. This section provides an overview of key AI models - ChatGPT, Grok, Gemini, Claude, and DeepSeek - outlining their core capabilities, practical applications, and known limitations in government and administrative settings.

The objective is to enable officers to make informed, task-specific choices when using AI tools, while maintaining appropriate human oversight and adhering to principles of accuracy, transparency, data security, and ethical responsibility. Used correctly, these models can enhance efficiency and decision support without replacing professional judgment or accountability.

Understanding Limits of LLMs : Bias, Context, and Role of Human Judgment

For Pakistan's public sector, it is essential to recognize that most LLMs are trained predominantly on English-language data and Western knowledge systems. As a result, their outputs may conflict with local culture, societal norms, legal frameworks, or state interests.

Techniques such as better prompt design, RAG, and fine-tuning can reduce risk but do not eliminate underlying bias or epistemic limitations. LLMs are probabilistic systems, not sources of truth; verification and human judgment must remain central and be explicitly embedded in policy design across education, health, and other critical sectors.

Overview of Different LLMs for Government Use

1. ChatGPT (OpenAI)

Overview:

ChatGPT is a widely used LLM developed by OpenAI, built on the GPT architecture. It is optimized for conversational interaction, analytical reasoning, and structured drafting tasks which are commonly required in government offices.

Capabilities:

- Drafting official reports, summaries, memos, and policy briefs
- Converting raw notes into structured Minutes of Meetings (MoMs)
- Preparing citizen-facing communication in multiple languages
- Analyzing text, extracting patterns, and generating recommendations

Limitations:

- May occasionally produce factual inaccuracies (“hallucinations”)
- Output quality depends heavily on clear and well-structured prompts
- Cannot access confidential or internal government databases
- Token limits and knowledge cut-offs may restrict long or updated content

Government Context Examples:

- Drafting policy briefs for ministries and attached departments
- Preparing disaster response notices in Urdu and English
- Structuring office orders, meeting agendas, and summaries
- Reviewing lengthy documents to extract key insights

2. Grok (xAI)

Overview:

Grok is an LLM developed by xAI, designed for speed, real-time information processing, and concise reasoning. It is suitable for fast-moving administrative and communication tasks.

Capabilities:

- Rapid summarization of lengthy documents
- Quick reasoning support for contingency planning
- Generating real-time insights for press statements or talking points
- Interpreting trends in fast-evolving situations

Limitations:

- Tone may be less formal than required for official documentation
- Limited suitability for legal or highly detailed policy drafting
- Not ideal for sensitive or neutral regulatory content

Government Context Examples:

- Preparing rapid talking points for senior leadership
- Summarizing district or field reports
- Drafting short public safety alerts (heatwaves, smog, floods)
- Supporting crisis communication units with quick situational updates

3. Gemini (Google)

Overview:

Gemini is Google's LLM designed for reasoning, drafting, and seamless integration with Google Workspace tools such as Docs, Sheets, and Gmail, making it useful for routine office workflows.

Capabilities:

- Drafting emails, memos, meeting notes, and briefs
- Assisting with data analysis and summaries using Google Sheets
- Translating or summarizing government notifications
- Comparing policy options and drafting evaluations

Limitations:

- Free versions offer limited output length and depth
- Requires internet connectivity and Workspace access
- Less suitable for very long or complex administrative documents

Government Context Examples:

- Drafting inter-departmental correspondence
- Summarizing spreadsheets (budgets, ADP progress, complaint logs)
- Preparing working papers for meetings
- Converting field reports into structured summaries

4. Claude (Anthropic)

Overview:

Claude is designed for accuracy, neutrality, and long-document handling, making it well-suited for detailed administrative, regulatory, and policy-related work.

Capabilities:

- Reading and summarizing long reports (50+ pages)
- Drafting structured inquiry reports and policy options
- Maintaining a formal, neutral, bureaucratic tone
- Supporting multi-step policy and regulatory analysis

Limitations:

- Slower than Grok for quick-response tasks
- May provide overly cautious responses
- Free versions may limit document uploads

Government Context Examples:

- Reviewing lengthy PC-I and PC-II documents
- Drafting inquiry and investigation reports
- Preparing legislative or regulatory briefs
- Conducting comparative policy analysis

5. DeepSeek

Overview:

DeepSeek is a cost-efficient LLM designed for structured outputs, data processing, and large-scale usage, making it suitable for routine administrative automation.

Capabilities:

- Processing structured data and tables
- Generating standardized forms, templates, and checklists
- Automating repetitive drafting tasks (notices, certificates, formats)
- Producing rule-based and consistent outputs

Limitations:

- Less creative or nuanced than ChatGPT or Claude
- Requires carefully structured prompts
- Not suitable for high-level policy or strategic drafting

Government Context Examples:

- Creating standardized formats for office orders and certificates
- Processing large datasets (complaints, attendance, monitoring logs)
- Drafting routine public notices
- Developing templates for administrative workflows

Summary Table: AI Models for Government Use

Model	Strengths	Limitations	Best Government Uses
ChatGPT	Strong drafting ability, versatile, multi-lingual	Risk of factual errors; needs good prompts	Policy drafting, summaries, MoM, public notices
Grok	Fast reasoning, real-time insights	Tone less formal	Rapid updates, crisis comms, quick summaries
Gemini	Google Workspace integration	Free version limited	Email drafts, data summaries, workflow automation
Claude	Excellent with long/technical documents	Sometimes too cautious	Inquiry reports, legislative drafts, briefings
DeepSeek	Efficient, structured outputs	Less nuanced	Templates, notices, formats, dataset processing

MODULE 3 : INTRODUCTION TO PROMPT ENGINEERING

A prompt is the instruction or request given to an AI system to guide its response. Prompt engineering is the practice of designing these instructions in a structured and intentional way to achieve accurate, relevant, and reliable outputs.

Effective use of AI depends largely on how clearly this prompt is written. Prompt engineering refers to structuring instructions so the AI understands its role, the context of the task, and the expected output. For government officers, good prompting enables AI systems to produce precise, relevant, and context-appropriate results, especially when used for analysis, drafting, and decision support. This is especially useful for complex tasks such as policy drafting, analytical reviews, stakeholder communication, and structured decision support, where clarity and consistency are critical.

Large Language Models do not reason in the human sense. They generate responses by predicting likely patterns based on training data, which means they can produce outputs that sound confident but are incorrect or incomplete (hallucinations). Thoughtful prompting can reduce this risk by narrowing the scope, adding context, and requesting structured outputs, but it cannot eliminate hallucinations.

Effective use of AI depends on the quality of prompts. A clear definition of the AI's role, relevant context, expected output, and task clarity significantly improve results, especially when using AI agents in government work. While many prompt frameworks exist, no single framework will remain sufficient as LLMs continue to improve. Understanding the principles of good prompting matters more. The example provided on the next pages can be adapted and refined to improve productivity over time.

Anatomy of an Effective AI Prompt

Act as a senior policy analyst in the Ministry of Commerce.	Role
Develop a policy note for the Commerce Minister outlining the key advantages and disadvantages of bilateral free trade agreements.	Task
The policy note is required for internal briefing purposes ahead of a cabinet-level discussion. The intended audience is senior decision-makers who require a clear, balanced overview rather than technical economic detail.	Context
Explain the purpose of bilateral free trade agreements, outline three potential benefits for the national economy, and highlight three key risks or sectoral challenges, using a neutral, evidence-informed approach.	Reasoning
Maximum 300 words, formal and objective in tone, with a short introduction, bullet-pointed benefits, bullet-pointed risks and challenges, and a brief conclusion.	Output Format
Do not advocate for or oppose any agreement, avoid naming specific countries unless requested, and end with a summary paragraph.	Stop Conditions

Sample Prompt

You are a public health policy analyst supporting a district administration in Pakistan.

Use only authentic Government of Pakistan data, including population and demographic data from official sources such as the Pakistan Bureau of Statistics (PBS) and other publicly available government datasets. Do not rely on private surveys or estimates unless clearly stated as assumptions.

Analyze the district Mardan's population demographics with a focus on children under five years of age, including population size, gender distribution, urban–rural split, and any relevant socio-economic indicators available in government data.

Based on this analysis, identify priority areas and population segments for a child vaccination campaign. Highlight gaps, risks, and practical considerations for outreach.

Present the output as a short policy brief with:

1. Key demographic insights
2. Implications for vaccination coverage
3. Actionable recommendations for campaign design

Keep the language clear and suitable for senior government decision-makers.

KEY LESSONS FOR PROMPT ENGINEERING

A key lesson in prompt engineering is that clarity matters more than complexity. AI systems respond best when instructions are direct, specific, and free of ambiguity. Long or complicated prompts are not necessarily better; what matters is whether the task, purpose, and expected outcome are clearly stated.

One widely used approach is to define the role the AI should play, such as an analyst, policy drafter, or administrative assistant. Assigning a role helps the AI align its response with the user's intent and reduces irrelevant or unfocused output, which is especially important in government contexts.

Another important element is context. Providing background information, constraints, or reference material allows the AI to tailor its response to the policy, sector, or institutional setting. Without context, the AI may produce generic answers that are less useful for public-sector decision-making.

Clearly specifying the expected output is also critical. Stating whether the response should be a summary, a set of options, a comparison, or a draft policy note improves both accuracy and usability. This is particularly valuable when AI is used to support reviews, briefings, or internal reports.

Finally, users should recognize that prompt frameworks are guides, not fixed rules. As AI models continue to improve, rigid formulas will become less important than understanding core principles: clear intent, sufficient context, and iterative refinement. Effective prompting is an ongoing practice that improves with experimentation and feedback.

Ten Government-Centric Practice Exercises

1. Rewrite a vague prompt to address a specific economic issue in Punjab (e.g., agriculture productivity or SME development).

2. Draft a contextualized smog-related public safety alert as the Assistant Commissioner of Gujranwala.

3. Create a few-shot prompt to classify citizen complaints (Sanitation, Water Supply, Encroachment, Revenue, Health, Education).

4. Prepare a structured 150-word bullet-point summary of the Annual School Census Report with three actionable recommendations.

5. Refine a summary of polio campaign progress for a briefing to the Minister for Health.

6. Develop a chain-of-thought prompt outlining monsoon flood preparedness steps for District Khairpur.

7. Use role-based prompting to propose a three-point municipal waste management plan as a Deputy Commissioner.

8. Design a prompt chain for digitizing land records in Khyber Pakhtunkhwa, including a policy brief, press release summary, and bilingual awareness post.

9. Draft a neutral 100-word summary of a land dispute in South Punjab for submission to the Commissioner.

10. Complete a capstone by designing an AI-assisted workflow for a government task (e.g., PC-I summary, MoM, divisional review, or CMO briefing note).

MODULE 4 : BEGINNER HANDS-ON AI TOOLS FOR PRACTICAL APPLICATION

This module provides probationary officers with guided, beginner-friendly hands-on exposure to selected AI tools that demonstrate how Large Language Models can be applied in practical, low-risk, non-technical workflows.
The focus of this module is experiential learning—allowing participants to see, build, and interact with AI-powered systems without coding, while reinforcing responsible use, human oversight, and relevance to administrative work.

Module Objectives

- Understand how modern AI tools are used beyond chat interfaces
- Create structured documents using Claude Artifacts
- Build a simple no-code chatbot for informational use
- Design a basic automation for repetitive administrative tasks
- Understand the concept of local / on-premise AI models using Ollama

This module equips probationary officers with practical exposure to modern AI tools while reinforcing governance-first thinking. Participants leave with a clear understanding of:

- What AI tools can do
- What they should not be used for
- How AI fits into structured administrative workflows rather than replacing them

Creating Structured Documents Using Claude Artifacts

To demonstrate how AI can be used to generate well-structured official documents such as briefs, notes, checklists, and summaries in a controlled format.

Tool Overview

Claude Artifacts allow users to generate and edit long-form structured outputs (documents, tables, templates) in a separate workspace, rather than conversational chat.

Hands-On Activity

Task: Create a structured briefing note using Claude Artifacts.

Participants will:

1. Open Claude and select the Artifact feature
2. Use a role-based prompt such as:
3. "Act as a Section Officer. Create a one-page briefing note on improving complaint redressal timelines in district offices."
4. Generate an artifact structured as:
    - Title
    - Background
    - Key Issues
    - Recommendations
5. Edit and refine the document within the artifact workspace

Building a No-Code Chatbot Using Botsonic

To demonstrate how AI-powered chatbots can be created without coding for basic informational or internal support use cases.

Tool Overview

Botsonic is a no-code platform that allows users to build chatbots trained on uploaded documents or predefined information.

Hands-On Activity

Task: Build a simple informational chatbot.

Participants will:

1. Create a Botsonic account
2. Upload a small sample document (e.g., office procedures, FAQs, training notes)
3. Configure a chatbot to answer basic questions such as:
    a. Office timings
    b. Required documents
    c. Process explanations
4. Test the chatbot by asking sample questions

Governance Note

- Not to upload confidential or sensitive data
- That such chatbots are suitable only for non-sensitive, informational use

Creating a Simple Automation Using Make.com

To introduce the concept of workflow automation, showing how AI and automation reduce repetitive administrative effort.

Tool Overview

Make.com is a no-code automation platform that connects apps and triggers actions based on events.

Hands-On Activity

Task: Build a basic automation workflow.

Participants will:

1. Create a simple scenario in Make.com
2. Connect two basic modules (trigger + action)
    a. Trigger: New entry added to a Google Sheet (e.g., complaints log)
    b. Action: Automatically generate a formatted summary using AI
    c. Output: Send the summary via email or store it in a document
3. Run and test the automation

Understanding Local AI Models Using Ollama for On-Prem Use

To introduce the concept of local/on-premise AI, especially relevant for government environments with data sensitivity concerns.

Tool Overview

Ollama allows users to download and run small language models locally on their computers without internet dependency.

Hands-On Activity

Task: Download and run a small language model locally.

Participants will:

1. Observe a live demonstration of installing Ollama
2. Download a lightweight model (e.g., LLaMA-based model)
3. Run a basic prompt locally such as:
    a. Summarize this paragraph in formal language."
4. Compare local model behavior with cloud-based AI tools

Creating Custom AI Gems for Professional Productivity

This session introduces participants to the concept of custom AI Gems designed to support specific professional and administrative tasks.

Tool Overview

Google Gems are custom-configured AI assistants designed to perform specific, recurring professional tasks with consistency and structure.

Hands-On Activity

Task: Creating a “Meeting Minutes (MoM) Assistant” Gem

Participants will:

1. Open Google Gemini using the link https://gemini.google.com/app
2. Create a new Gem with the purpose of converting raw meeting notes into formal Minutes of Meeting.
3. Set the Gem’s role as a government documentation assistant using a neutral and factual tone.
4. Define a fixed structure for output: meeting details, discussion summary, decisions, and action items.
5. Test the Gem by pasting rough notes and review the generated MoMs before final use.

Data Analysis Using Julius

To demonstrate how officers can use a no-code AI data analysis tool to quickly extract insights from data without writing formulas, code, or queries.

Tool Overview

Julius is an AI-powered, no-code data analysis tool that allows users to upload datasets (CSV or Excel) and interact with them using plain language questions.

Hands-On Activity

Task: Analyze a small dataset and generate insights for a briefing note.

Participants will:

1. Upload a simple dataset (e.g., complaints data, attendance records, service delivery statistics) into Julius.
2. Ask Julius to explain what the dataset contains and identify key variables.
3. Ask questions such as:
    a. “What are the main trends in this data?”
    b. “Which category has the highest volume?”
    c. “Are there any noticeable changes over time?”
4. Generate one chart or visual summary using natural language.
5. Ask Julius to produce a short, plain-language insight summary suitable for an internal briefing.

Annexure-General Suggestions

1. Optimize Prompt Structure:
    - Use numbered lists or bullet points in prompts to enforce structured outputs (e.g., “include 1) objective, 2) strategies”).
    - Incorporate specific constraints (e.g., “use budget data”) to avoid generic responses.
    - Experiment with temperature settings (if available, e.g., via API) to balance creativity and precision (low temperature for CoT, higher for creative summaries).
2. Maximize Efficiency:
    - Minimize token usage by avoiding redundant context in chained prompts (e.g., reference prior outputs instead of repeating).
    - Use batch testing to compare multiple prompt variations quickly (e.g., test CoT with/without step labels).
    - Leverage Grok‘s free access on grok.com or x.com for rapid prototyping, noting usage limits.
3. Enhance Effectiveness:
    - Combine techniques (e.g., CoT within role-based prompts) for complex tasks like policy analysis.
    - Include few-shot examples in prompts to guide output style (e.g., “format like this: [sample policy summary]”).
    - Validate outputs against real-world data (e.g., cross-check with health or tourism reports).

4. Iterative Refinement:
    - Analyze LLM outputs for biases or inaccuracies, especially for sensitive issues (e.g., community disputes).
    - Use feedback loops in prompt chains to refine outputs iteratively (e.g., “revise based on stakeholder feedback”).
    - Log prompt-output pairs to identify patterns and optimize future prompts.
5. Ethical Considerations:
    - Ensure prompts avoid generating biased or harmful content, critical for diverse communities.
    - Verify factual accuracy in outputs, especially for legal or financial recommendations, using official data sources.

Specialised AI Tools, Apps, and Websites for Different Professions

No	Profession / Domain	Representative AI Tools & Primary Use
1	Advertising & Marketing	Jasper (marketing content), Surfer SEO (content optimisation), Albert.ai (automated media buying)
2	Art & Design	Midjourney (text-to-image), Adobe Firefly (image & text effects), Invoke (custom model training)
3	Content Creation & Writing	Grammarly (grammar & style), Jasper (idea generation), Hemingway App (readability analysis)
4	Education	Google AI Essentials (education tools), Khanmigo (AI tutoring), Slides AI (presentation generation)
5	Engineering	GitHub Copilot (code suggestions), Lovable (coding assistant), Adept (AI software task automation)
6	Finance & Banking	IBM Watsonx (risk & fraud detection), Daloopa (financial data extraction), Tesseract (investment analytics)
7	Healthcare	Google AI (Breast Cancer Detection), Babylon Health (symptom analysis), Tempus (genomic analytics)
8	Human Resources	Paradox (AI screening), Jobscan (resume optimisation), Talentsoft (employee learning)
9	Legal	Harvey AI (legal research), Casetext (case & statute search), Evisort (contract review)
10	Logistics & Supply Chain	FedEx AI (routing), Freightos (freight pricing), Locus (route optimisation)
11	Manufacturing	Siemens AI (predictive maintenance), Augury (failure prediction), Seebo (quality prevention)
12	Media & Entertainment	Descript (text-based audio/video editing), Runway (AI video), ElevenLabs (voiceovers)
13	Project Management	Asana Intelligence (task management), Motion (auto scheduling), Reclaim AI (focus optimisation)
14	Real Estate	Zillow AI (home recommendations), Realtor.com AI (market insights), Restb.ai (property image tagging)

No.	Profession / Domain	Representative AI Tools & Primary Use
15	Retail & E-commerce	Amazon AI (product recommendations), Dynamic Yield (personalisation), Criteo (targeted ads)
16	Sales	Relevance AI (lead generation), Salesforce Einstein (predictive scoring), Gong (sales-call analytics)
17	Science & Research	Perplexity (cited answers), Scite (paper evaluation), Elicit (research summarisation)
18	Software Development	GitHub Copilot, Amazon CodeWhisperer, Tabnine (code completion)
19	Customer Service	Intercom (chatbots), Zendesk (ticket routing), Drift (visitor engagement)
20	Data Science & Analytics	Numerous.ai (spreadsheet analytics), Tableau AI (visualisation), DataRobot (AutoML)
21	Photography	PhotoRoom (background removal), Luminar Neo (AI photo enhancement), Remove.bg (background removal)
22	Audio Production	Descript (dialogue editing), LALAL.AI (vocal separation), Altered Studio (voice modification)
23	Human Resources & Recruiting	Jobscan (resume analysis), Eightfold AI (talent matching), HiredScore (applicant ranking)
24	Administrative & Productivity	Notion AI (summaries & organisation), Microsoft Copilot (M365 assistance), Zapier (workflow automation)

The Artificial Intelligence 101 for Public Sector program is a joint initiative by the Ministry of Information Technology and Telecommunication, Ministry of Planning, Development and Special Initiatives, Civil Services Academy, and atomcamp, aimed at equipping Civil Services Academy probationers with essential AI awareness for effective governance and policy-making.

For further information, please contact the Civil Services Academy, Walton Lahore

Training Manual
AI 201

Applied AI Sytems, Governance, & Secure Deployment in the Government

TECH destiNATION PAKISTAN
MINISTRY OF IT & TELECOM

Ministry of Planning Development & Special Initiatives

atomcamp

Training Manual AI 201

Table of Contents

COURSE OVERVIEW ---------- 66

MODULE 1 : AI GOVERNANCE AND THE ROLE OF CIVIL SERVANTS ---------- 68

MODULE 2 : DESIGNING AI SYSTEMS AND WORKFLOWS ---------- 79

MODULE 3 : AI DEPLOYMENT AND SECURE IMPLEMENTATION ---------- 105

MODULE 4 : PUBLIC SECTOR APPLICATIONS AND POLICY CONTEXT ---------- 129

CONCLUSION ---------- 142

Course Overview

This two-day intermediate-level instructional program equips probationary officers at the Civil Services Academy with a structured transition from AI literacy to the applied use of Artificial Intelligence (AI) in government. Building on concepts introduced in AI 101, the program focuses on how AI is planned, governed, and deployed in public sector contexts, with emphasis on administrative decision-making, policy implementation, and service delivery. While designed for graduating probationary officers, it is equally relevant for public sector officials across ministries, departments, training academies, and government organizations involved in managing AI initiatives.

The curriculum covers practical government use cases, including workflow automation, decision-support systems, and AI-enabled service processes. Officers learn to work with no-code tools, understand AI-driven workflows, and engage effectively with vendors and technical teams without requiring programming skills. The program emphasizes responsible AI adoption, enabling participants to assess risks, understand system limitations, and ensure outputs remain accurate, contextual, and appropriate for official use.

The accompanying manual serves as a practical guide for planning, supervising, and governing AI projects in accordance with professional standards, ethical requirements, data confidentiality obligations, and established government protocols, while aligning AI adoption with Pakistan's regulatory and policy framework.

Duration : 2 days

Target Audience:

- Probationary officers at the Civil Services Academy
- Civil servants involved in policy making, administration, office management, and public service delivery;
- Public sector officials from ministries, departments, public sector training academies,
- Autonomous bodies seeking applied capacity in Artificial Intelligence.

Prerequisites: Completion of the AI 101 course, along with basic computer literacy and familiarity with text-based interfaces.

Learning Objectives

01 Define Civil Servant's Role as AI Product Owners and System Architects

02 Understanding Core Technical Concepts for Government Decision-Makers

03 Understanding integration of AI systems into real workflows

04 Deployment of AI Models in Government

05 Reflecting upon the Public Sector Use Cases and District-Level Applications

06 Policy and Regulatory Context for AI in Pakistan

MODULE 1 : AI GOVERNANCE AND THE ROLE OF CIVIL SERVANTS

Artificial Intelligence in government is not a purely technical project, rather it is a governance responsibility. Civil servants act as AI product owners who define objectives, supervise implementation, and ensure lawful and ethical use. Rather than building models, officers design workflows, oversee vendors, and remain accountable for outcomes. This module introduces the mindset and core concepts needed to manage AI systems responsibly in public administration.

Learning Outcomes

By the end of this module, participants will be able to:

- Define the role of a civil servant as an AI system owner and decision supervisor
- Differentiate automation, AI, machine learning, DevOps, and MLOps in government projects
- Evaluate how AI fits into administrative workflows while preserving human authority

1.1 The Role of Civil Servants as AI Product Owners and System Architects

In government AI initiatives, civil servants act as product owners and system architects, not hands-on developers.

This means that the responsibilities are to:

- Define what the AI system should achieve
- Ensure alignment with laws, policies, and public interest
- Oversee implementation and decision-making use

Technical teams such as vendors or IT staff:

- Build models
- Write code
- Handle technical deployment

You provide:

- Domain knowledge such as land revenue rules and service workflows
- Governance oversight
- Operational requirements

Think of your role as the bridge between governance and technology. You translate administrative needs into AI system requirements and ensure outputs are used appropriately.

Accountability and Human Oversight

AI systems assist decisions but do not replace responsibility.

- Final accountability remains with the department
- AI recommendations must not be treated as final decisions
- AI should function as a support tool, not an autonomous authority

In practice, this requires:

- Human-in-the-loop validation for critical processes
- Documentation of how AI outputs are reviewed
- Transparent appeal mechanisms for citizens

If a citizen is affected by an AI-assisted decision, they must have a clear path to human review.

Thinking Like a System Architect

As a system architect, you design the high-level workflow, not the code.
Example: AI complaint management system

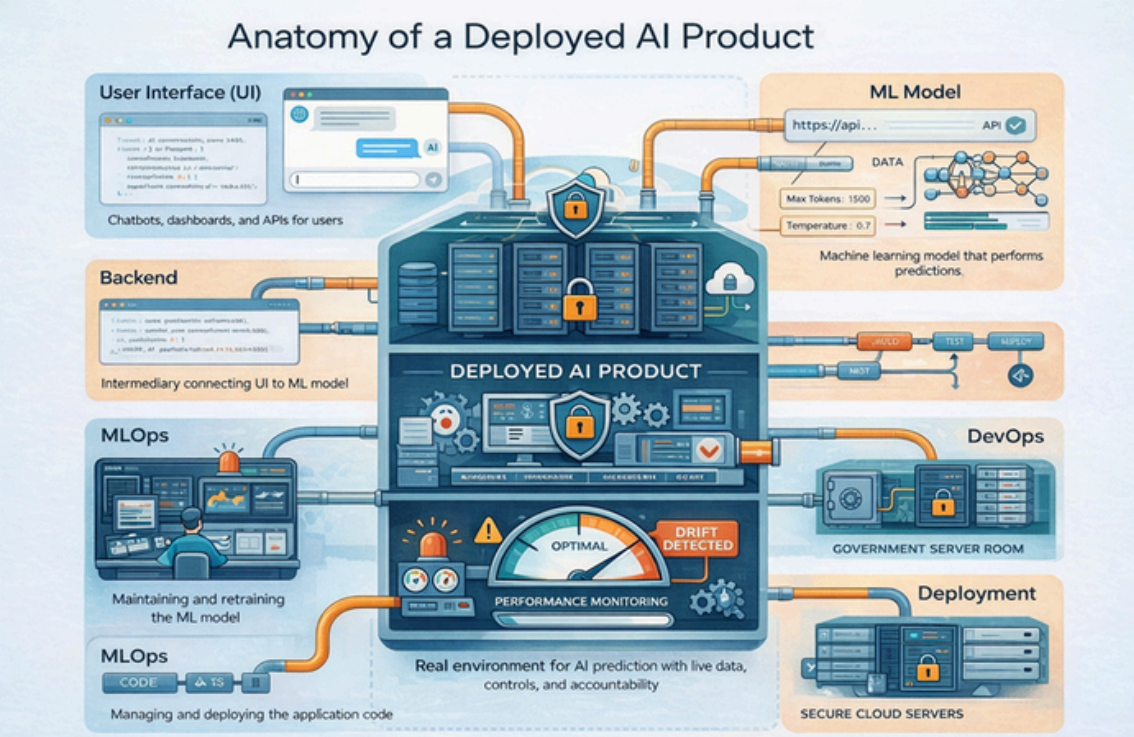
- Chatbot collects complaint
- NLP model categorizes issue
- System routes case to department

Key questions you should ask:

- Where is data stored?
- Is sensitive data sent to external cloud services?
- How is privacy protected?
- What happens if the AI makes an error?
- Is there a human review step?

To succeed in this role, it's important to build a basic understanding of how AI systems are constructed.

1.2 Understanding an AI System (System Layers)



The diagram shows how a complete AI system works from user request to final decision support. Read it top to bottom as the journey of a citizen query.

1. User Interface

Where people interact with the system.
Examples: chatbot, portal form, or officer dashboard.
It only collects input and shows results.

2. Backend Layer

Connects the interface to the AI model.
It validates data, applies rules, and sends requests to the model securely.

3. AI / ML Model

The intelligence of the system.
Performs tasks such as classification, text extraction, or recommendations.
Outputs assist decisions but do not replace human authority.

4. DevOps and MLOps
Keeps the system running after deployment.
Includes updates, retraining, version control, and security maintenance.

5. Monitoring and Oversight
Tracks accuracy and detects model drift.
If performance drops, the model must be reviewed or retrained.

AI systems are not just models. They are governed workflows requiring continuous supervision. As an AI product owner, you ensure each layer is controlled, accountable, and aligned with policy.

Your Oversight Responsibilities as an AI Product Owner

The diagram highlights that DevOps (deployment and infrastructure) and MLOps (model training, monitoring, retraining) are essential parts of any AI system. As a project lead, your responsibility is to ensure vendors or technical teams have clear processes for both.

You should be able to ask practical questions such as:

- How will the model be updated if performance drops or patterns change?
- Who retrains the model and what data will be used?
- Where will the system be deployed: government servers or cloud?
- Is there a rollback plan if the system fails?

The key principle: successful AI is not just a model rather it is a maintained system.

Rule-based

Governance and Ethical Responsibility
You must also ensure AI use complies with policies, laws, and public trust requirements.

For example:

- A Safe City surveillance system must respect privacy and legal approvals.
- A land record AI can assist decisions but cannot replace the authorized officer.
- Some technically possible use cases may still be rejected if ethically risky.

Government AI deployment is therefore an act of stewardship that requires deciding not only how AI is used, but also where it should not be used.

Your Role in the Project

You act as:

- Visionary : identify useful AI opportunities
- Supervisor : coordinate teams and vendors
- Guardian : enforce governance and legality

You are effectively coordinating multiple stakeholders to deliver an AI system that is reliable, lawful, and beneficial to citizens.

Before proceeding further, it is important to understand the core technical concepts that government decision-makers should be familiar with.

1.3 Core Technical Concepts for Government Decision-Makers

Government AI projects often involve multiple technical terms that sound similar but refer to different responsibilities in a system. As a project lead, understanding these distinctions helps you communicate clearly with vendors, evaluate proposals, and set correct expectations.

1. Automation: Rule-Based Task Execution

What it does: Performs repetitive tasks using predefined rules.

How it works

- Follows fixed instructions: if X happens → do Y
- No learning or adaptation
- Stops working if conditions chang

Example

Automatically sending an acknowledgement email after a citizen submits a form.

Automation improves efficiency but cannot handle variation or judgement.

2. Artificial Intelligence (AI) : Decision Support from Patterns

What it does: Performs tasks requiring human-like reasoning, classification, or judgement.

How it works

- Analyzes data and patterns instead of fixed rules
- Handles ambiguity and complex situations
- Can prioritize, recommend, or predict outcomes

Example

Analyzing applications to identify high-priority cases without explicit rules.

Automation follows instructions. whereas AI determines what the instructions should be.

3. Machine Learning (ML) — Learning from Data

What it is: A subset of AI that learns patterns from examples instead of programming logic.

How it works

- Train model using labeled historical data
- Model predicts future outcomes
- Improves when retrained with new data

Example

Detecting fraudulent transactions or classifying land use from satellite imagery.

AI is the goal (intelligence), whereas ML is the primary method used to achieve it.

4. DevOps — Reliable Software Deployment

What it manages: The lifecycle of software applications.

Purpose

- Faster and safer updates
- Minimal service disruption
- Collaboration between developers and IT operations

Includes

- CI/CD pipelines
- Automated testing
- Monitoring
- Infrastructure management

Why it matters

Ensures AI applications remain available and stable when updated.

5. MLOps — Reliable Model Deployment

What it manages: The lifecycle of machine learning models.

Purpose

- Keep models accurate over time
- Retrain when data changes
- Track model versions and performance

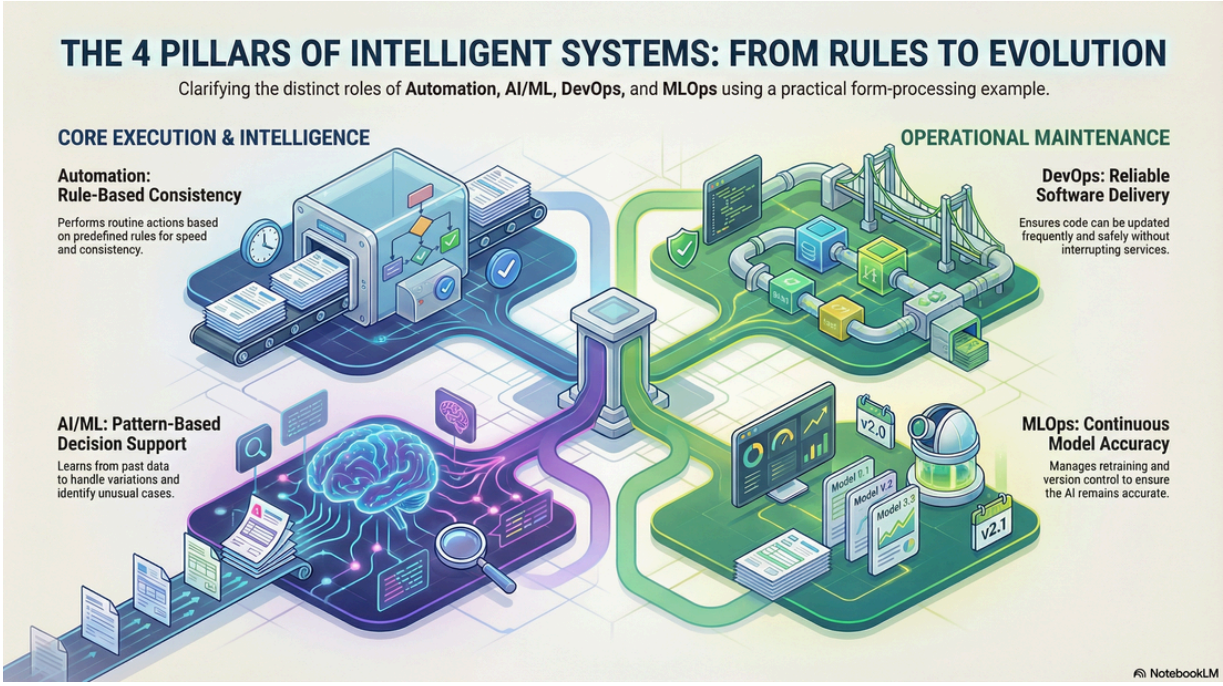
Includes

- Data preprocessing
- Training and validation
- Deployment
- Monitoring model drift
- Scheduled retraining

Why it matters

An AI model degrades over time if not maintained.

1.4 Illustrative Example: Form Checking System



Consider a department that manually checks submitted forms before approval.

- **Automation:** checking steps are programmed into the system using predefined rules. For example, if a specific field contains a certain value, the system performs a defined action automatically every time. This improves speed and consistency, but the system cannot handle forms that differ from the expected format or contain unusual information.
- **AI/ML:** system is trained on past forms and their outcomes. Instead of relying only on fixed rules, it learns patterns and can extract information even when forms vary slightly, such as interpreting handwritten comments. The AI can also identify unusual submissions and flag them for human review because it has learned what typical cases look like.

To support this system operationally,

- **DevOps** practices ensure that the software running the form checking process can be updated frequently and safely. New rules or improved software versions can be deployed without interrupting services for citizens.

- **MLOps** practices manage the machine learning model itself. As new labeled examples accumulate over time, the model is retrained periodically. Each version of the model is evaluated and stored so that if performance decreases, the system can revert to a previous reliable version.

This example demonstrates how automation performs routine execution, AI/ML provides intelligent decision support, DevOps maintains reliable software delivery, and MLOps ensures the AI model remains accurate and up to date.

As a civil servant project lead, you are not expected to implement DevOps or MLOps yourself, but you should ensure your technical team or vendor follows these practices. Doing so reduces the risk of project failure because the system is treated as an evolving service rather than a one time deliverable.

Key Takeaways

AI systems support decisions but never replace responsibility. Civil servants translate policy needs into system requirements, ensure human oversight, and maintain accountability. A successful government AI project is not just a model, it is a governed, monitored, and continuously maintained public service system aligned with law and public trust.

Moving Forward

With these concepts understood, the next step is practical implementation. No-code and low-code workflow platforms help connect AI, automation, and systems together, allowing teams to deploy solutions without deep programming expertise.

MODULE 2 : DESIGNING AI SYSTEMS AND WORKFLOWS

After understanding your role as an AI owner, the next step is learning how AI systems actually work together in practice. Government AI projects depend on integration; connecting data sources, workflows, and decision processes securely. This module explains how APIs, custom AI assistants, no-code workflows, and agentic processes combine to form usable public sector solutions, while ensuring systems rely on authoritative data and remain under human control.

Learning Outcomes

By the end of this module, participants will be able to:

- Explain how APIs connect AI systems with government databases and services
- Design structured workflows using no-code tools and controlled AI assistants
- Assess when and how agentic workflows can safely support administrative processes

2.1 No-Code Workflows in Government

No-code and low-code platforms use visual workflows built from triggers, actions, and APIs. Tools like make.com and n8n are useful for learning AI-enabled automation because they make workflow steps visible (triggers, data mapping, rules, approvals, exceptions, and audit trails) and enable quick individual productivity use cases.

Example workflow:

- A complaint is submitted in an online portal (trigger)
- AI summarizes the complaint (action)
- An email is sent to the concerned officer (action)

Users simply connect modules and configure settings (API keys, conditions) instead of writing scripts.

These tools are included in this manual as capability demonstrators to help civil servants understand how intelligent automation is structured end-to-end. In government, however, automation typically must run on on-prem or sovereign infrastructure and integrate with custom and legacy systems under strict security and governance controls.

Make.com (Cloud-Based Automation)
You previously used Make.com in AI 101 to connect a data source to an AI summary and automated email.

Key characteristics:

- Easy interface and quick prototyping
- Large library of integrations
- Runs on external cloud infrastructure

Because processing occurs outside government servers, it should only be used for non-sensitive data workflows.

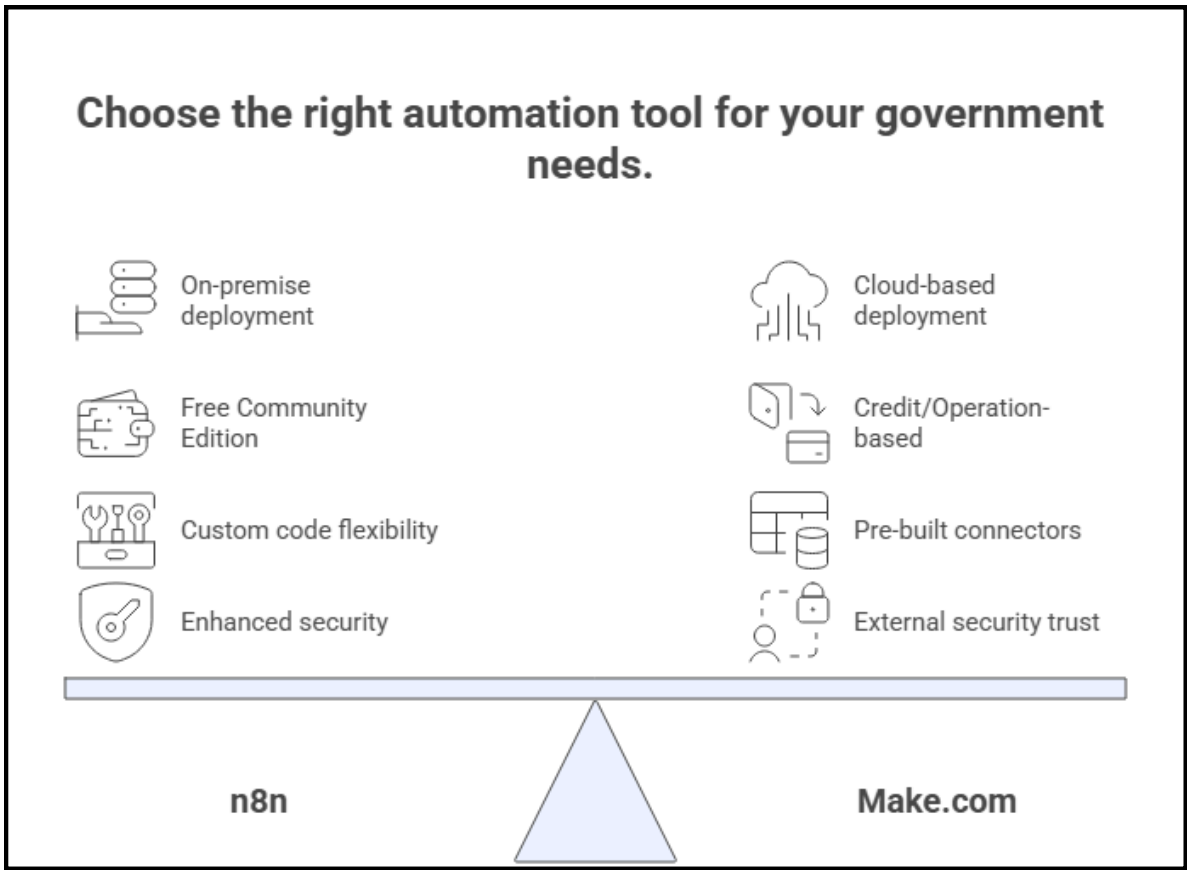
n8n (Self-Hosted Government-Friendly Automation)
n8n provides similar workflow automation but can be installed on internal infrastructure.

Key advantages:

- Can run on local servers or private cloud
- Data remains within government control
- Supports internal databases and secure systems
- Allows custom logic when required

Use case:
A department can automate records processing, connect internal databases, and call AI APIs while ensuring no data leaves the official network.

Let's compare Make and n8n on a few key points:

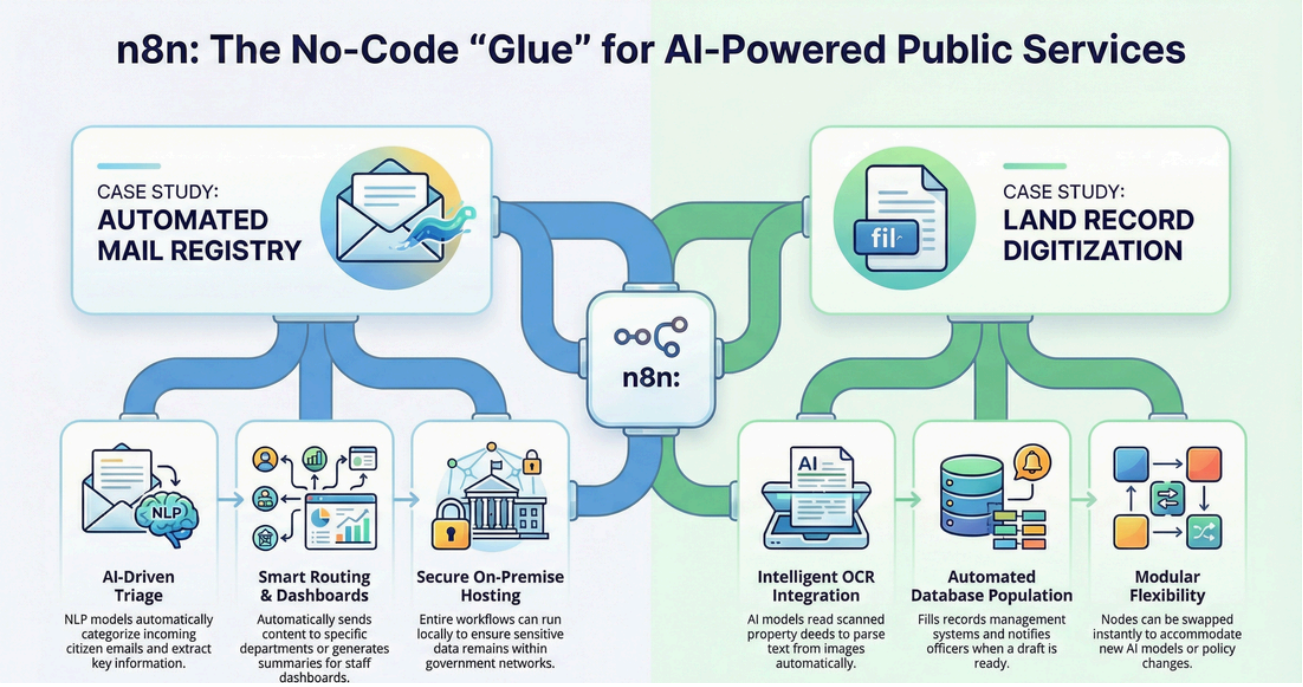


Make.com and n8n empower you to become solution architects, integrating AI and automation to streamline your workflows. Use Make.com for quick wins in non-sensitive areas or for early experiments, and embrace n8n for serious, secure deployments within your own environment. Emphasize to your teams that these tools are designed to augment productivity, not to bypass IT policies. When used responsibly, they can be game changers in delivering faster public services integrated with AI capabilities.

In summary, Make.com is ideal for quick, cloud-based prototypes, while n8n is better suited for secure, on-premise deployments of automated workflows. As civil service managers, you might start experimenting with Make.com, particularly for non-sensitive public data, to test its functionality. However, for production systems that involve sensitive citizen data or require integration with internal government systems, it is advisable to switch to n8n on a controlled server. This approach allows you to benefit from the ease of no-code tools without compromising on data governance and security.

2.2 Workflow Design Examples

Let's consider a practical examples of how n8n can be used.



Example 1: Automating Incoming Mail Processing (District Office)

Objective: Automatically categorize and route citizen emails.

Workflow steps using n8n:

- Monitor the official inbox for new emails (trigger).
- Send the email text to an AI NLP service for classification or information extraction.
- Route the categorized case to the relevant department.
- Store the summary in a database or Excel dashboard.
- Notify staff if required.

Deployment options:

- Run fully on-premise using a local AI model.
- Or call an approved external AI API through n8n's HTTP Request node.

Benefit:
Visual workflow replaces custom software development and can be easily modified.

Example 2: Land Record Digitization Support (PLRA Scenario)

Objective: Assist officers in preparing digital land records.

Workflow steps using n8n:

1. Detect upload of a scanned property deed (trigger).
2. Send the image to an OCR AI service to read Urdu text.
3. Extract key fields and populate the Land Records Management System.
4. Notify an officer to verify the drafted record.

Benefit:
Configuration replaces complex integrations. If the OCR model changes, only that workflow step needs updating, not the whole system.

After exploring how to build AI workflows, next we will explore how AI systems are integrated into real workflows and what concepts such as APIs and agentic workflows mean in practical government use.

2.3 APIs and System Integration

AI systems in government rarely operate alone. They must connect with databases, portals, and existing software to retrieve data and perform actions. Integration is therefore a core part of any practical AI deployment.

What is an API?
An API (Application Programming Interface) is a structured way for software systems to communicate automatically.

Instead of a human manually entering information:

- One system sends a request
- Another system returns a response

Example:
An identity authority provides a verification API. A government AI application can send CNIC details and instantly receive confirmation.

Many AI services work the same way:

- Send text → receive generated response (language model)
- Send image → receive detected objects or extracted text (vision model)

Even no-code tools like Make.com and n8n use APIs behind the scenes when connecting email, databases, and AI services.
System Integration in Government

Government departments often use legacy systems that were not designed for automation. To use AI effectively:

- Create secure API access to official databases
- Work with vendors/IT teams to build an integration layer
- Ensure real-time and authorized data access

Without proper integration, AI systems cannot reliably operate on official government data.

2.4 Retrieval-Augmented Systems (Practical AI Architecture)

In government settings, AI must answer using verified records, not general knowledge. Retrieval-Augmented Generation (RAG) enables this by grounding AI responses in official data.

Example: Land Records Assistant
An officer asks about a property's ownership history. The system:

1. Authenticates the request
2. Retrieves relevant records from the official database via secure API
3. Sends the retrieved data to the AI model
4. Generates a clear, human-readable explanation

The AI does not guess. It explains what exists in the database.

Why RAG is Appropriate for Government

- Responses are grounded in authoritative data
- Outputs reflect current records
- AI acts as an interpreter, not a decision-maker

Governance Requirements

- Authenticate all requests before access
- Retrieve only necessary information (least privilege)
- Protect sensitive data during transmission and processing
- Maintain logs for accountability
- Keep humans responsible for final decisions

APIs and Integration as Governance
RAG relies on APIs to connect AI with government systems. This is not only technical integration but a governance responsibility. Proper integration ensures secure data access, compliance with policy, and reliable outputs aligned with official records.

After securely integrating AI systems with official government data, the next step is controlling how the AI interprets and communicates that information. A general-purpose model will respond generically unless configured for departmental rules, tone, and requirements. This is where model customization becomes necessary.

2.5 Customizing AI Models for Government Use (Custom GPTs)

The term GPT (Generative Pre-trained Transformer) commonly refers to models such as GPT-4 that power conversational AI systems. A custom GPT means adapting a general-purpose AI model to perform a specific government function.

There are several ways to customize such models:

- **Fine-tuning**

A base model is trained further on a specialized dataset so it learns domain-specific knowledge.

For example, training on past citizen helpline responses allows the model to adopt official communication style and terminology.

This approach requires large datasets and technical expertise. In government settings, fine-tuning can help a model understand legal terminology or generate responses aligned with official correspondence.

- **Prompt Engineering / Few-shot Prompting**

Instead of retraining the model, instructions and examples are provided each time it runs.

For example, the prompt may define the model as an assistant for the Punjab Land Records Authority and require it to cite record identifiers in responses.

This creates a controlled behavior without modifying the model itself. Departments can embed policies, tone, and reference material directly into prompts.

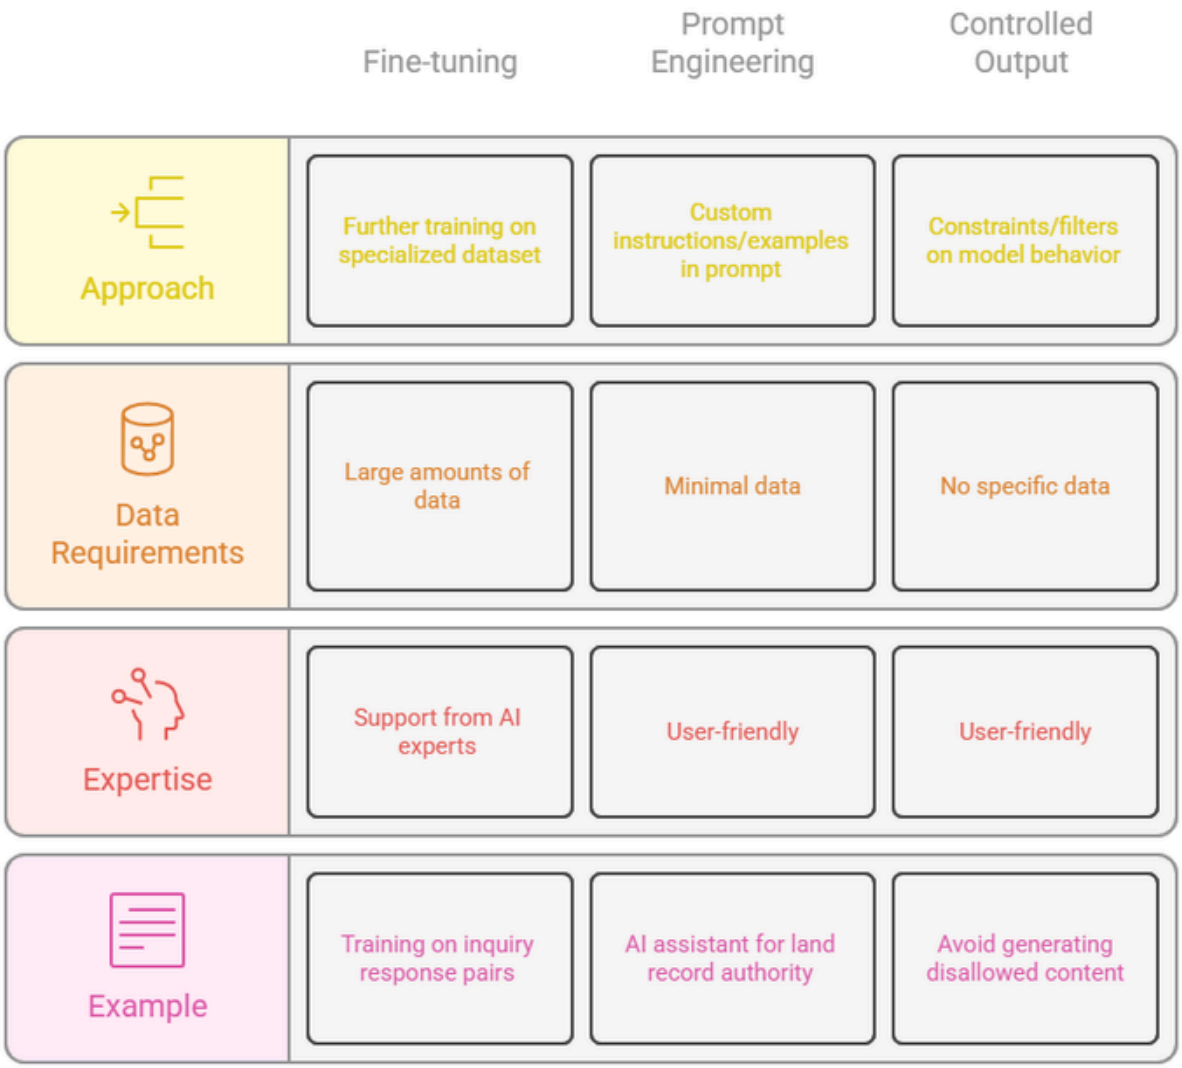
- **Controlled Output (Guardrails)**

Constraints are added to ensure responsible behavior.

Examples include:

- enforcing structured formats such as official letters
- preventing prohibited content
- requiring disclaimers when confidence is low
- avoiding fabrication of uncertain information

Many platforms allow such rules to ensure outputs remain consistent with government accountability standards.



Purpose of Custom GPTs in Government

Custom GPTs function as specialized assistants.

Examples include:

- TaxGPT: Helps officers interpret FBR tax regulations
- HealthGPT: Assists doctors with medical guideline queries

These systems rely on general AI capabilities but are enhanced with local data and operational rules.

Security Considerations

Government data must only be used in approved environments.

- If internal documents are uploaded to external services, the platform must be formally vetted.
- Enterprise platforms may offer encryption and no-training guarantees, but approval is still required.
- Alternatively, open-source models can be hosted internally to avoid external dependencies.
- Local models (e.g., LLaMa-based) can run on-premises using tools like Ollama and integrate with n8n for fully offline deployments.

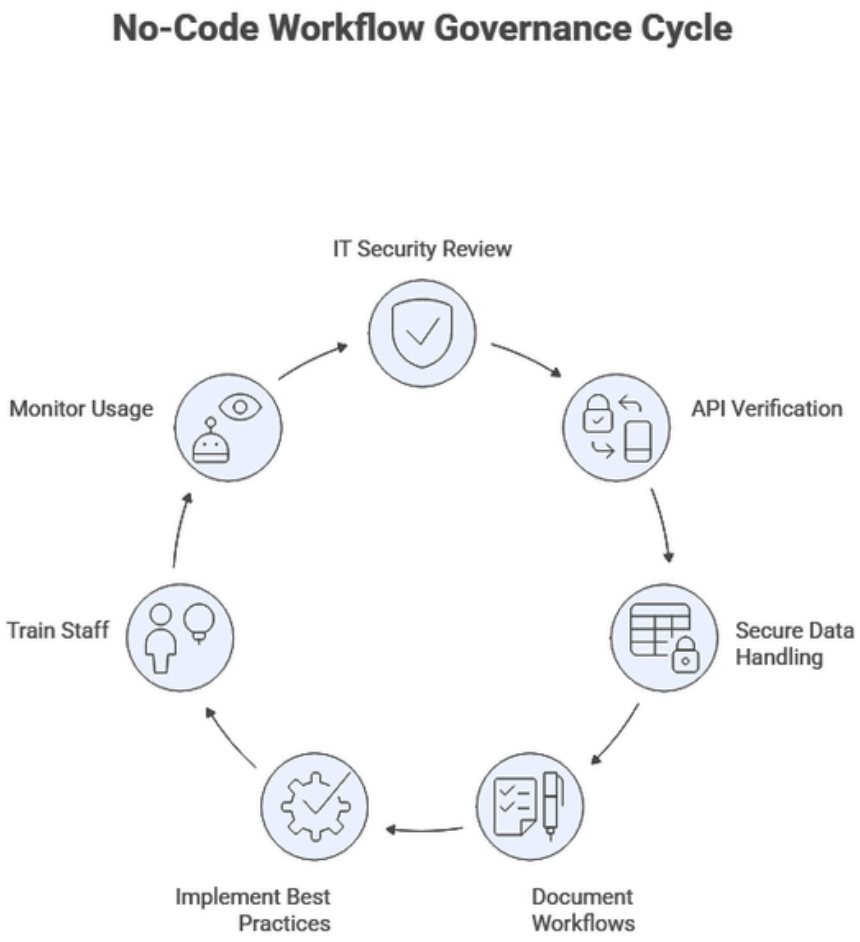
Customizing AI models allows government departments to control how an AI understands policies, communicates with citizens, and follows official procedures. However, even a well-configured assistant still primarily responds to requests.

The next step in AI capability is enabling systems not only to answer questions but to actively carry out tasks. This introduces the concept of agentic workflows, where AI systems can plan actions, interact with multiple systems, and complete multi-step processes with minimal human intervention.

Best practices and oversight

No-code workflows must be handled responsibly. The ease of automation does not remove the need for planning and governance. Always ensure the following:

- Involve the IT team for a security review, especially when workflows call external APIs.
- Store sensitive credentials (e.g., API keys) securely using encrypted storage features such as those provided in n8n.
- Maintain clear documentation describing each workflow's purpose and behavior.
- Treat no-code automations with the same rigor as traditional software systems, since they operate official processes.



2.6 Agentic Workflows

Agentic workflows refer to AI systems that can plan and execute a sequence of actions to achieve a goal. Instead of only answering questions, the AI breaks a task into steps, interacts with tools, and completes parts of a process automatically.

What this means in practice
An agent can:

- Read a request
- Retrieve information from databases through APIs
- Apply rules or checks
- Generate a document
- Trigger another system (for example scheduling a meeting)

This goes beyond a simple chatbot and becomes a task-performing assistant.
Government example
For a citizen complaint:

1. The agent reads the complaint
2. Checks existing records in a database
3. Drafts an official response
4. Schedules a meeting if required

For a land mutation application:

1. Verify required documents
2. Check rules or calculations through connected systems
3. Prepare a recommendation note for the officer

The officer still makes the final decision.

How it works

Agentic systems rely on:

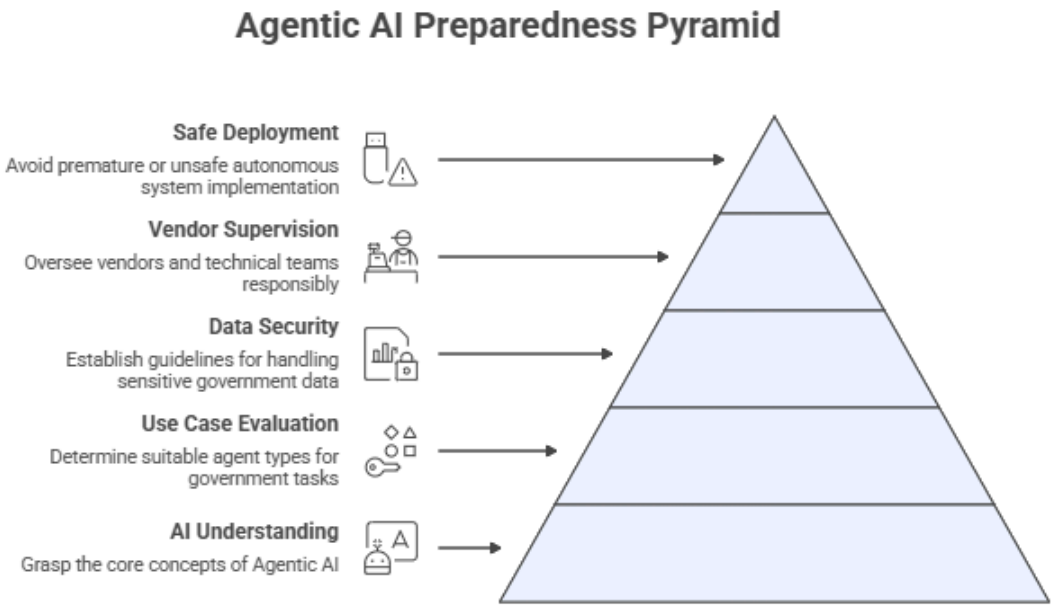
- APIs to access official data
- Tools or scripts for calculations or checks
- Frameworks such as LangChain or LangGraph
- Workflow tools like n8n to control what actions the agent is allowed to perform

2.7 Agentic AI in Government Systems: From Prototypes to Government-Grade Agents

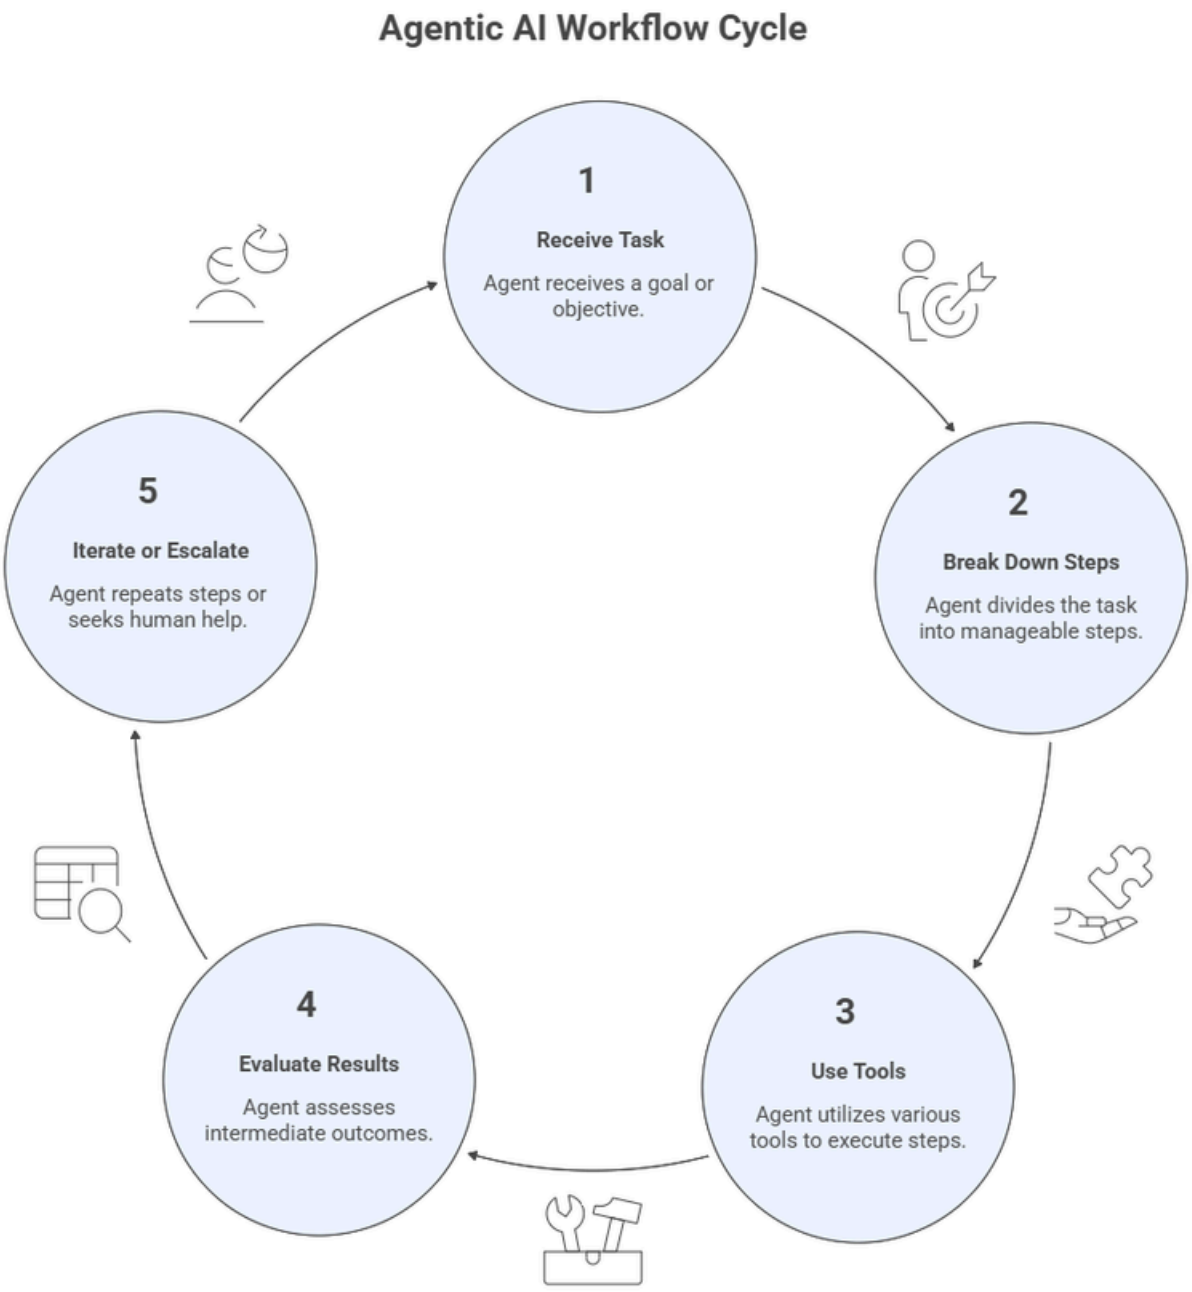
In recent years, the term “Agentic AI” has rapidly entered public discourse. Vendors, technology platforms, and consultants increasingly promote “AI agents” as autonomous digital workers capable of performing complex tasks with minimal human involvement. While the concept holds genuine promise for government systems, it is also highly fragmented, inconsistently defined, and frequently misrepresented.

This section is included to provide civil servants with a clear, grounded, and operational understanding of Agentic AI—free from hype—and to distinguish experimental prototypes from government-grade, deployable agent systems.

This manual does not prepare officers to build agents themselves. It prepares them to:



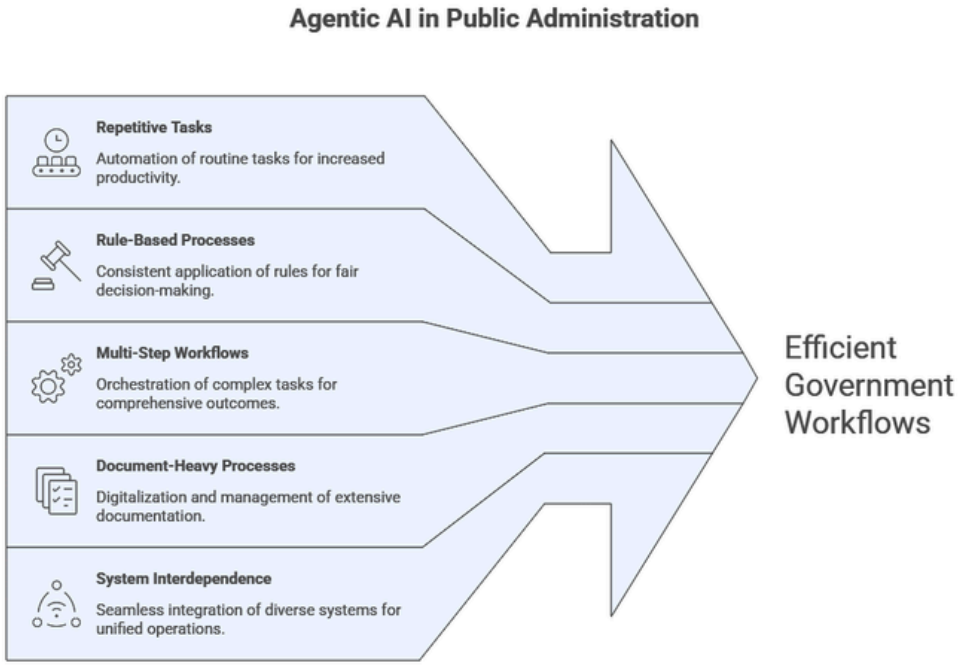
Unlike simple chatbots or single-prompt AI tools, an agent:



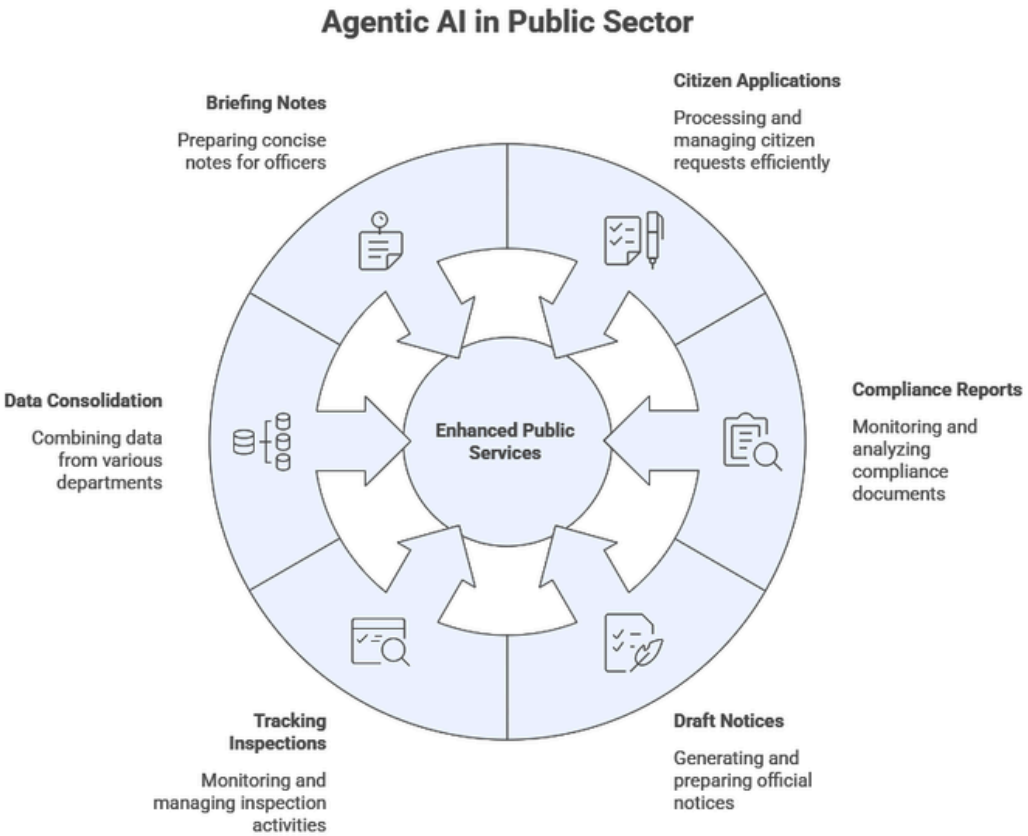
In government terms, an agent is best understood as a digital process executor, not a decision-maker.An agent should prepare information, execute workflows, and recommend actions, but final authority must always rest with a human officer.

Why Agentic AI Is Beneficial for the Public Sector

Agentic AI is appealing in public administration because many government workflows are:



Examples of Agentic AI in Public sector includes:



Tiers of Agentic AI

Civil servants must learn to distinguish three distinct tiers of Agentic AI, each with different risk profiles, technical requirements, and acceptable use cases. The three levels are as follows:



Level 1: No-Code / Cloud-Based Agentic AI (Exploratory & Learning Stage)

Level 2:Low-Code, On-Prem-Capable Agentic Systems (Operational but Controlled)

Level 3: Government-Grade Agentic AI Systems (Advanced & Regulated)

Lets discuss them in detail

Level 1: No-Code / Cloud-Based Agentic AI (Exploratory & Learning Stage)

- Purpose

Use no-code tools such as Custom GPTs, Make.com agents, Gemini demos, or similar SaaS platforms to understand how AI agents work. These systems are externally managed and offer limited control over data, logs, and model behavior.

- How to use

Create agents through natural language instructions and prebuilt integrations. Use only for demonstrations, training, and quick prototypes. Work strictly with dummy or synthetic data to simulate workflows, for example categorizing sample complaints and suggesting departments.

- Where appropriate

Use for officer training, concept validation, and internal experimentation. Treat as a learning environment, not an operational system.

- What to avoid

Do not upload or process real citizen data such as CNIC records, land records, health data, or law enforcement information. Do not use for official decisions or legally defensible processes. Do not rely on these tools for audit trails or accountability.

- Rule

Level 1 tools are for learning and prototyping only. Never deploy them in real government operations.

Level 2: Low-Code, On-Prem-Capable Agentic Systems (Operational but Controlled)

- Purpose

Use workflow orchestration tools such as n8n to build AI systems inside government infrastructure. This level provides control over data, execution, and audit trails, making it suitable for real operational support.

- How to implement

Build agents as step-by-step workflows. Connect databases, APIs, and AI models through defined logic. Keep processing on government servers or private cloud where possible. Add mandatory human approval points before any action. Log every step for auditability and compliance.

- Where appropriate

Use for operational assistance where AI supports staff but does not decide independently. Example: an inspection support agent that reads field reports, summarizes findings, flags possible violations, and sends a draft to an officer for review. The officer remains the decision maker.

- Key safeguards

Keep sensitive data on-premises. Ensure transparent logic and traceable actions. Restrict automation to recommendations and drafts. Define escalation paths and approval checkpoints.

- Advantages

Allows integration with existing MIS systems. Supports local or approved AI models. Maintains records for audits. Enables SOP-based automation while preserving accountability.

- Organizational requirements

Officers do not need coding skills, but the department must have technical support and governance rules in place. This is the minimum level appropriate for operational government AI use.

- Required Skills (Organizational, Not Individual):

While officers do not need to code, the organization must have some technical capacity, including IT partners or entities like PITB. Clear governance rules are also necessary to ensure the responsible use of these systems.

Level 3: Government-Grade Agentic AI Systems (Advanced & Regulated)

- Purpose

Build fully programmable AI agents for high-impact government functions. Use frameworks such as LangChain, LangGraph, or similar systems with Retrieval-Augmented Generation. This level is used only when strong legal, technical, and institutional controls exist.

- How to implement

Develop agents through code with defined reasoning steps and modular logic. Retrieve information only from approved data sources. Host models in controlled government infrastructure. Ensure outputs include explanations, citations, and confidence indicators. Maintain full audit logs and enforce mandatory human approval before any action.

- Where appropriate

Use for analytical decision support in sensitive domains. Example: a land record intelligence agent that checks records, compares metadata, and prepares an analytical brief for officers. The system never modifies records or contacts citizens without human authorization.

- Key safeguards

Keep systems in secure environments. Use verified datasets only. Provide explainable outputs. Implement escalation and failure handling procedures. Never allow autonomous legal or enforcement decisions.

- Operational requirements

Requires advanced engineering expertise, security architecture, and continuous maintenance. Deploy only for high-impact cases where regulation, accountability, and long-term support are assured.

As AI moves from experimental tools to operational and government-grade systems, it begins handling real and sensitive records. At this point, privacy, security, and accountability become critical.

2.7 Governance of Agentic AI

Agentic systems can perform actions across software systems, so strong controls are required. An incorrect automated action can have administrative or legal consequences.

Control principles

- Restrict capabilities through sandboxing
- Allow only approved or read-only queries by default
- Require confirmation before irreversible actions
- Display reasoning steps so users can review planned actions
- Keep humans responsible for final decisions

Fully autonomous decision-making agents are not advisable in government environments. Treat agents as junior assistants that gather information and propose actions, while officers approve execution.

Government APIs as Public Services
Departments may both consume and publish APIs.

Examples:

- Safe City systems exposing incident detection APIs
- Land record authorities providing ownership verification APIs
- Translation or document analysis models shared across departments

Designing AI as reusable services improves interoperability and aligns with digital government goals.

Practical Design Guidelines

Integration

- Apply least privilege access (limited read access instead of full control)
- Log all system interactions
- Prefer APIs over manual or UI automation

Custom GPTs

- Ground responses in official documents and policies
- Review outputs before treating them as official decisions
- Host internally or use approved secure platforms for sensitive data

Agentic workflows

- Start testing in controlled environments
- Keep human approval in critical steps
- Monitor actions and maintain audit trails

AI, APIs, and RPA Together

- APIs provide reliable integration
- RPA can temporarily connect legacy systems where APIs do not exist
- AI can read data while RPA enters it into older software

Prefer APIs whenever possible, but use RPA as a transitional solution.

Key Takeaway

- APIs act as bridges between systems
- Custom AI models act as trained specialists
- Agents act as assistants using tools

Your role is to design workflows where these work together safely and transparently.

The next section explains how to safely design and operate such systems: Building Confidential AI Agents Using Internal Government Data.

2.8 Building Confidential AI Agents Using Internal Government Data

Step 1: Identify Approved Knowledge Sources

Before any AI system is built, departments must define:

- Which documents are approved for AI use
- Classification levels (public, internal, confidential)
- Update and version-control mechanisms

This ensures data governance precedes AI deployment.

Step 2: Use Retrieval-Based Approaches (Not Public Training)

Instead of uploading documents to public AI tools, organizations should:

- Store documents in secure internal repositories
- Use retrieval-based systems where AI fetches relevant information at runtime
- Prevent models from learning or memorizing confidential data

This approach keeps data contextual but not exposed.

Step 3: Develop Organization-Specific AI Agents

Each department can develop AI agents tailored to its function:

- Policy drafting assistants
- Legal and regulatory lookup systems
- Service delivery support tools
- Internal reporting and briefing assistants

Step 4: Developing Pakistan's Own Local and Domain-Specific LLMs

For true data and linguistic sovereignty, Pakistan should invest in the development of local large language models (LLMs) that:

- Are trained and hosted entirely within Pakistan
- Support local languages and administrative terminology
- Reflect national legal, regulatory, and cultural context
- Are aligned with public-sector use cases rather than consumer applications

Such models would not replace global research efforts but would provide trusted, sovereign AI foundations for government use.

Key Takeaways

Effective government AI is not a standalone chatbot but an integrated system built on reliable data, controlled automation, and supervised decision-making. Proper integration ensures accuracy, custom AI ensures relevance, and human oversight ensures accountability. When workflows are thoughtfully designed, AI becomes a structured assistant that improves efficiency without compromising governance.

Way Foward

After exploring how to build AI workflows, the next step is deciding where they should run. In government environments, deployment matters as much as development because systems must protect data, comply with policy, and remain reliable. This leads to deployment architecture: selecting the hosting model based on data sensitivity and operational needs.

MODULE 3 : AI DEPLOYMENT AND SECURE IMPLEMENTATION

Designing an AI system is only half the responsibility; deciding where and how it runs is equally critical in government environments. Deployment architecture must protect citizen data, comply with national policies, and ensure reliability of public services. This module introduces deployment models (on-premises, cloud, and hybrid), security practices, and governance requirements under Pakistan's regulatory framework, helping officials make informed and compliant implementation decisions.

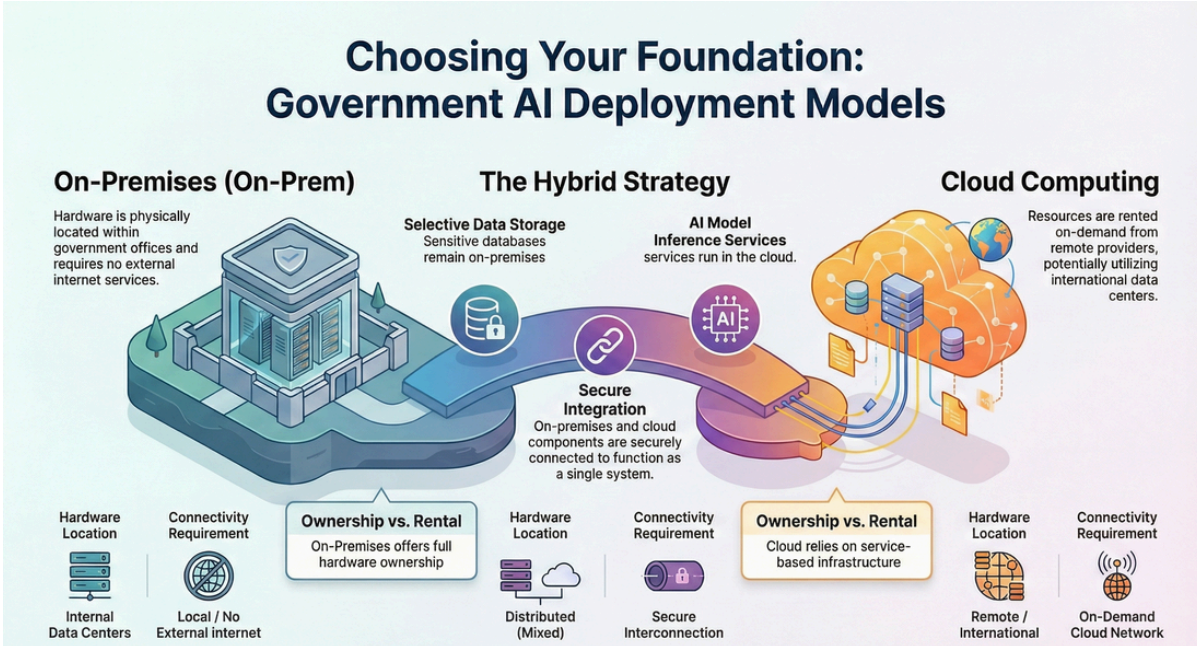
Learning Outcomes

By the end of this module, participants will be able to:

- Evaluate appropriate deployment models based on data sensitivity and operational needs
- Apply security and governance controls throughout the AI lifecycle
- Plan responsible scaling strategies that maintain governance, oversight, and human accountability

3.1 Choosing Deployment Architect

One of the first architectural decisions for any government AI system is determining where it will be deployed and run. In simple terms, there are three options:

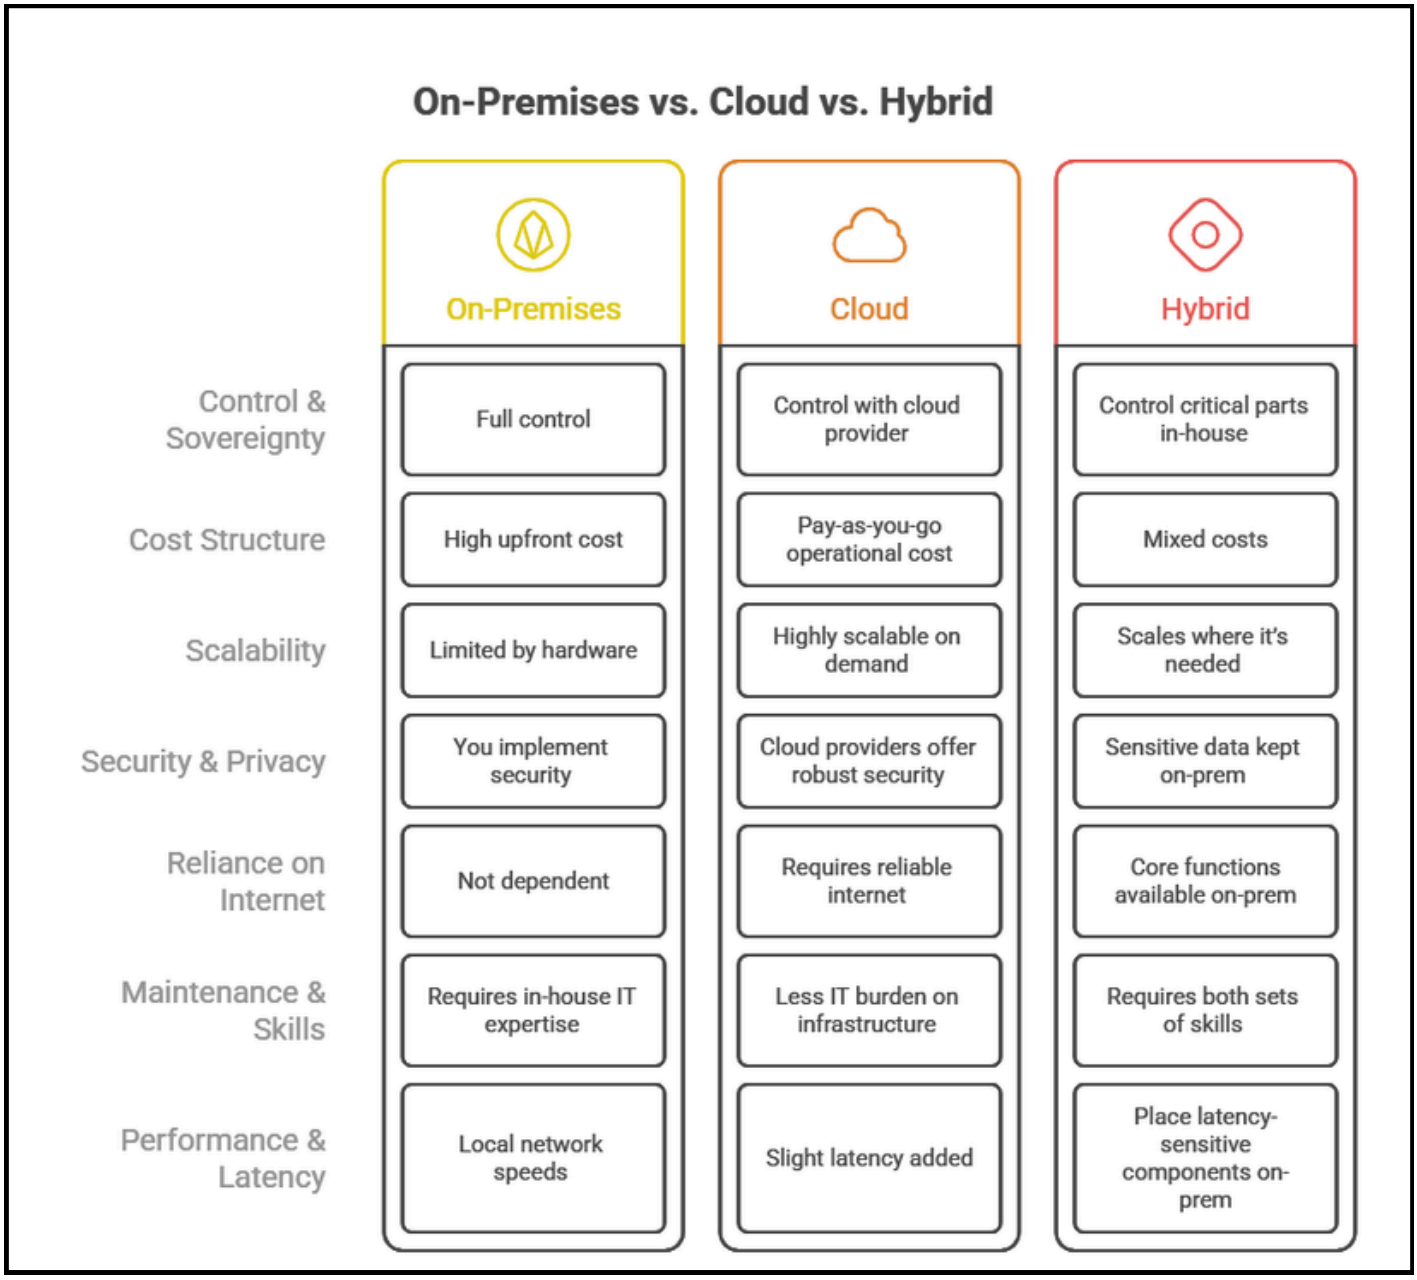


- **On-Premises** (On-Prem): The AI software runs on servers or computers physically located within the organization's data center or office. The government fully owns and manages the hardware and network, and no external internet services are required for core functionality.

- **Cloud:** The AI system runs on cloud computing platforms such as Amazon Web Services, Microsoft Azure, Google Cloud, or local cloud providers. The servers may be in remote data centers, potentially even outside Pakistan if international providers are used. In this case, the government rents computing resources and services from the cloud provider on-demand.

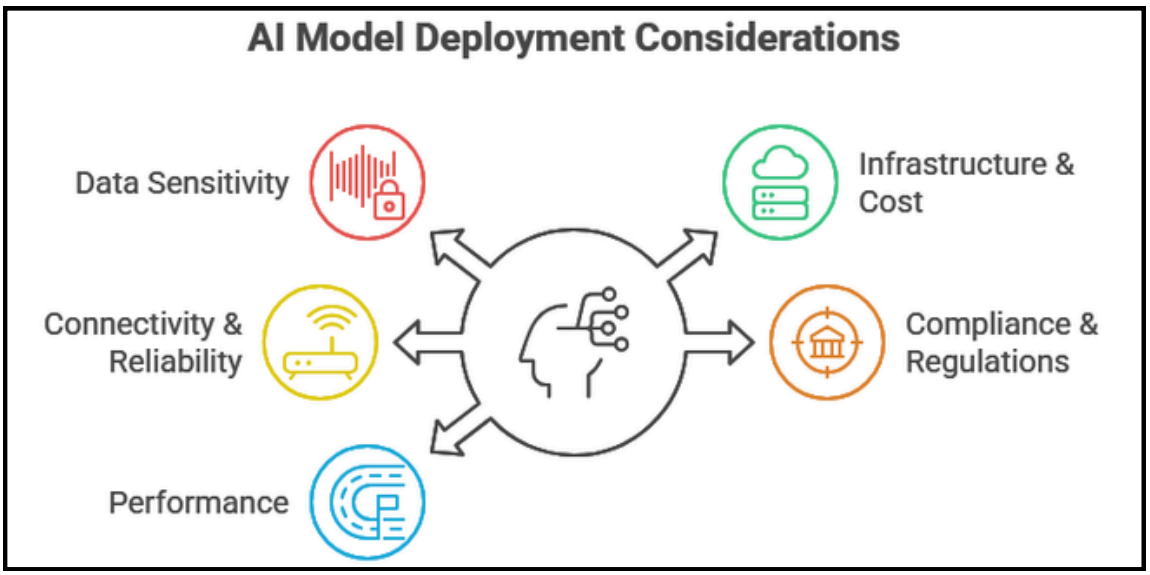
- **Hybrid:** A combination of on-premises and cloud components. For example, some parts of the system, such as databases or sensitive data storage, remain on-premises, while other parts, like AI model inference services, are run in the cloud. The two environments are securely connected to work together.

The diagram below compares the three deployment options On-Premises, Cloud, and Hybrid across key factors such as control, cost, scalability, security, connectivity, and performance. It shows that no single model is universally best. Each option offers advantages in certain areas and trade-offs in others, and in government environments the final choice is often determined not only by technical needs but also by policy, data sensitivity, and compliance requirements.



3.2 Deployment Decision Factors

Each model has its pros and cons, and importantly, policy constraints often dictate which model is appropriate for a given project. Let's break down the key considerations:



Data Sensitivity and Sovereignty

- Highly sensitive or classified data should remain on-premises or in an approved government/private cloud.
- Public data may be processed in a public cloud.
- Follow Pakistan's Cloud First Policy requirements before selecting a deployment.

Infrastructure and Cost

- On-premises: high upfront investment, full control, fixed capacity.
- Cloud: quick setup, pay-as-you-go, highly scalable but long-term costs may grow.
- Hybrid: keep regular workloads locally and use cloud for peak demand.

Connectivity and Reliability

- Cloud requires stable internet access.
- Remote or low-connectivity locations may require on-prem deployment.
- Always plan a fallback option in case of outages.

Compliance and Regulations

- Use only government-approved cloud providers.
- Ensure data residency and legal requirements are met.
- Confirm compliance with MoITT and sector-specific regulations before deployment.

Performance Requirements

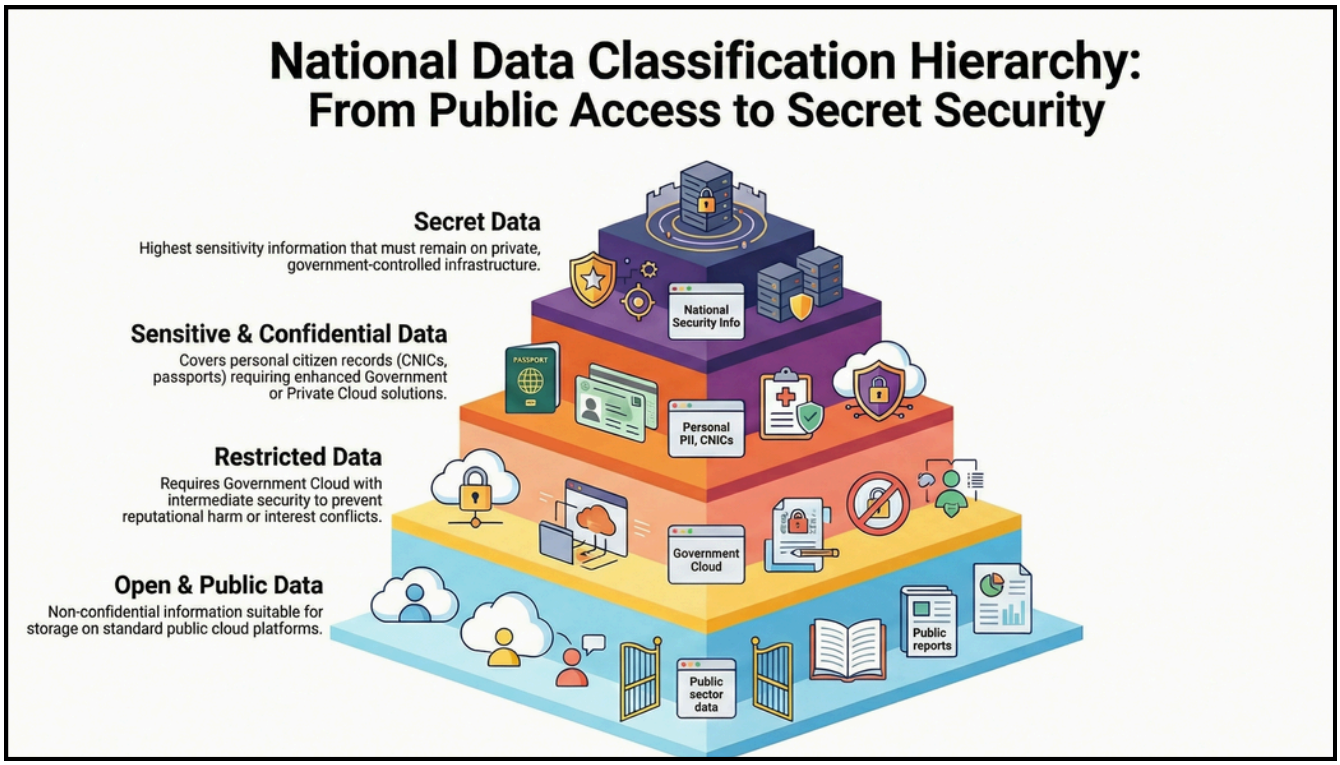
- Heavy or low-latency workloads may benefit from on-prem hardware.
- Cloud provides scalable compute (e.g., GPUs) for large processing tasks.
- Hybrid models can process sensitive or fast operations locally and send intensive tasks to the cloud.

Always choose the deployment model based on policy, data classification, and operational needs, not convenience alone.

3.3 Pakistan Cloud First Policy Lens

It important to know that Pakistan's Cloud First Policy encourages government departments to opt for cloud solutions when feasible, to improve efficiency and reduce costs. However, this doesn't mean cloud always; it means cloud should be the first option, evaluated within the framework of data classification and security.

The policy defines a Data Classification Guideline consisting of five categories: Open, Public, Restricted, Sensitive/Confidential, and Secret. The diagram below illustrates how these categories are applied.



Future Strategy:

The National AI Policy hints at developing national infrastructure like a national compute grid and high-performance computing centers. In the coming years, government agencies might have access to centralized AI computing facilities within Pakistan, allowing for AI workloads without relying on foreign clouds. Until then, you may use what's available, such as PITB or KP CIT, which have their own data centers. These solutions offer a balance of local control with cloud-like convenience, while still following national policy.

Practical Approach to Decision-Making:

When scoping an AI project, ask yourself: Can this run on the cloud, and is it allowed to? If yes, and it simplifies the project, propose the cloud option. If not, plan for on-premises by coordinating with your IT department to provision servers or allocate space in a government data center. Often, a pilot phase can be done in the cloud with test or dummy data to validate the concept, then migrated on-prem for full production with real data, ensuring both speed and security.

3.4 Secure Deployment Implementation Practices

Secure AI deployment requires both technical safeguards and administrative controls. Apply the following practices during design, testing, and operation.

1. Core Technical Security

- Encrypt data at rest and in transit (e.g., HTTPS for APIs)
- Apply secure coding practices to prevent vulnerabilities
- Define strict user roles and permissions (least privilege access)

2. Model Security

- Use models only from trusted repositories
- Scan open-source models before deployment (malware risk)
- Restrict direct external access to models that may contain sensitive training data

3. Logging, Auditing, and Monitoring

- Maintain logs of:
  - User access
  - System actions
  - Errors and failures
- Review logs periodically or during incidents
- Conduct performance and bias audits to ensure acceptable accuracy and fairness

4. Legal and Regulatory Compliance

- Follow applicable regulations (e.g., financial, biometric, or personal data rules)
- Obtain consent where required
- Inform citizens when interacting with AI systems
- Consult legal teams for systems handling regulated data

5. Physical and Operational Security

- Lock and control access to server rooms
- Provide backup power (UPS/generator)
- Maintain disaster recovery and backup plans

6.Vendor Governance

Include contractual safeguards:

- Data hosted only in approved locations
- Government retains ownership of data and trained models
- No reuse of government data by vendor
- Knowledge transfer and documentation required
- Exit plan with handover of code, models, and data
- Penalties for non-compliance

7. Testing Before Deployment

- Run trials using test or staging environments
- Conduct penetration testing
- Fix vulnerabilities before launch

8. Ethical Oversight

Follow national AI ethics principles:

- Transparency: document data sources and decision logic
- Fairness: test for demographic bias
- Privacy: anonymize and limit data retention
- Human oversight: officers approve decisions before action

Example: A surveillance AI flag must remain only a lead, not evidence, until verified by officials.

3.5 Secure Deployment and Governance in Pakistan's Policy Context

In government environments, security and governance are core requirements, not optional features. Every AI system must be designed, deployed, and operated in alignment with national policies such as the Cloud First Policy and MoITT regulations.
Secure deployment means protecting data, the system itself, and its ongoing operations throughout the entire lifecycle.

Key Security Areas

Data Security

- Protect stored and transmitted data from unauthorized access
- Use encryption for data at rest and in transit
- Apply role based access controls
- Secure networks using firewalls and VPN connections
- Ensure physical protection of on premises servers

System Security

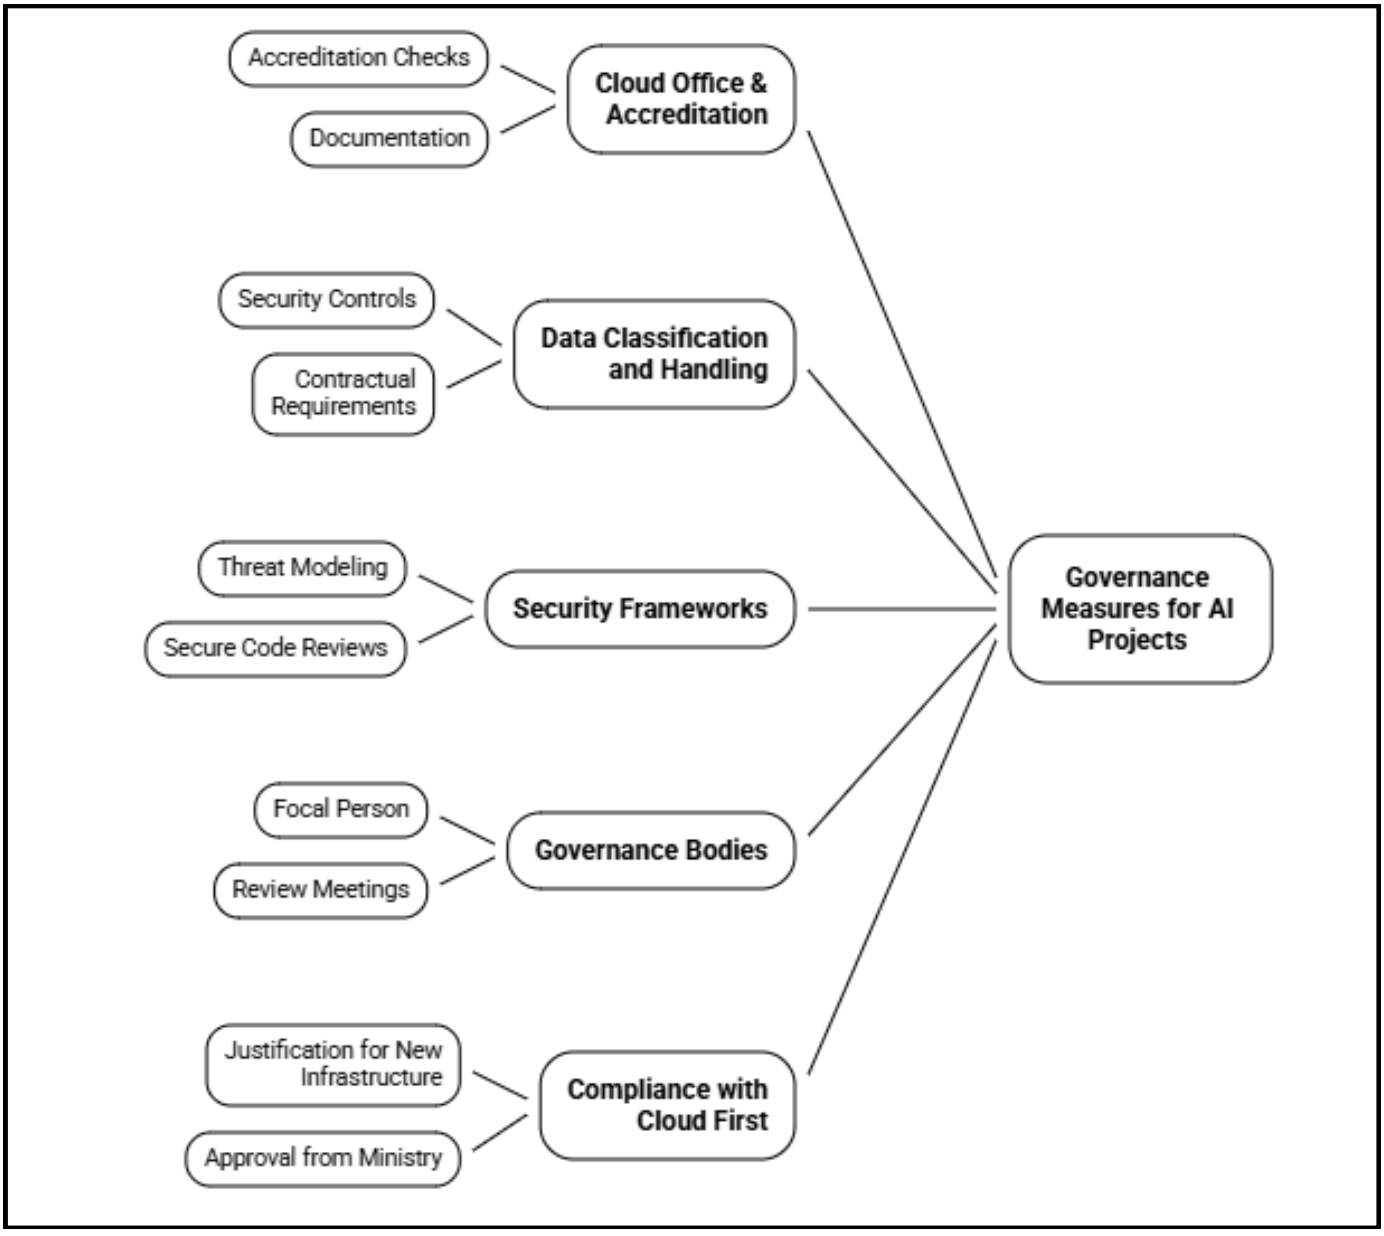
- Implement authentication for all interfaces and APIs
- Validate inputs to prevent misuse or exploitation
- Apply regular security patches and updates
- Restrict system permissions to necessary functions only

Operational Security

- Monitor system activity and maintain logs
- Detect abnormal usage patterns or malicious queries
- Trigger alerts for suspicious behavior or data extraction attempts
- Investigate unusual AI outputs that may indicate malfunction or compromise

These practices ensure AI systems remain reliable, lawful, and safe throughout their use in public service.

Pakistan's Cloud First Policy provides a structured framework for the secure use of cloud services and, by extension, any government IT deployment. Key governance measures from this policy and other Ministry of IT and Telecom guidelines include the following:



1. Cloud Office and Provider Approval

- Use only cloud providers accredited by the Ministry of IT and Telecom
- Confirm security certifications (e.g., ISO 27001), data residency, and compliance standards
- Document:
  - Which provider is used
  - What type of data is stored
  - Formal approval status from the Cloud Office

2. Data Classification and Handling

Apply controls based on the sensitivity of data:

- Confidential/Sensitive data
  - Strong encryption
  - Multi factor authentication
  - Security audits
- Public data
  - Standard security controls acceptable

Procurement contracts must clearly state:

- Government retains ownership of data
- Vendor cannot reuse government data
- Data must be returned at contract termination
- Vendor responsibilities in case of breach or failure

3. Security Framework Compliance

Follow recognized security practices:

- Threat modeling during design
- Secure code reviews for custom systems
- Penetration testing before deployment

4. Internal Governance Oversight

Establish departmental oversight:

- Assign an AI security or ethics focal person or committee
- Conduct periodic reviews covering:
  - Security incidents
  - Ethical risks
  - System performance
  - Service impact

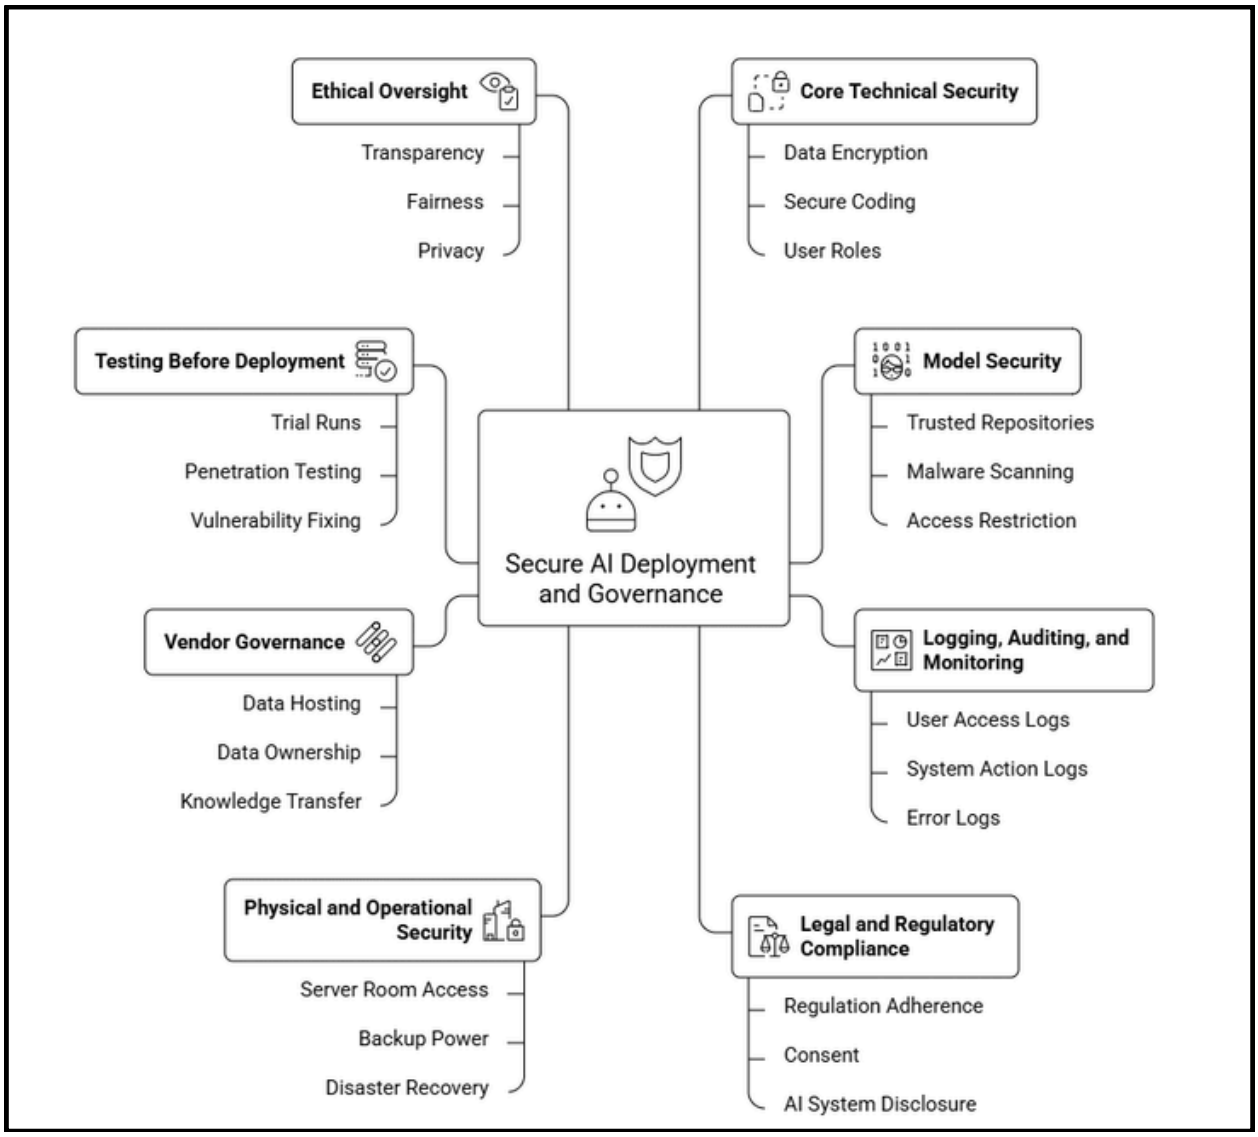
policy and remain secure, accountable, and auditable

5. Cloud First Requirement for New Systems

- Evaluate cloud deployment before buying hardware
- If on premises infrastructure is required, document justification (e.g., sensitive data)
- Obtain approval from MoITT or relevant IT board when necessary

Following these steps ensures AI deployments align with national policy and remain secure, accountable, and auditable.

Secure Deployment and Governance in Practice



It is essential to choose the deployment model that best meets your data security needs and resource constraints.

General guideline: if the data or application is highly sensitive or requires 24/7 operational continuity, regardless of internet connectivity, it is better to opt for on-premises or hybrid deployment.

3.5 Legal and Ethical Risks of AI Deployment in Government

AI systems used in government directly affect citizens' rights and public trust. Any system that provides information or supports decisions must therefore be governed carefully.

Key Risk: Incorrect or Harmful Guidance
AI can produce confident but wrong answers. If citizens rely on those answers, legal and ethical consequences arise.

Example – New York City MyCity Chatbot

- Provided illegal advice to businesses
- Claimed employers could keep workers' tips
- Suggested landlords could discriminate against voucher holders
- Contradicted itself and declared its own responses reliable
- Disclaimer added later did not prevent misuse

Impact

- Potential legal liability for government
- Citizens misled into unlawful actions
- Loss of public trust in official digital services

Cause

- Generic AI deployed without domain safeguards
- Insufficient validation and fine tuning
- Lack of human oversight and response control

Lessons for Government Deployments

- AI outputs must never be treated as authoritative without verification
- Domain specific validation is mandatory before public release
- Disclaimers alone do not mitigate risk
- Systems must include review, constraints, and accountability mechanisms

AI in government must assist officials, not replace official judgment.

This happened largely because the chatbot was a generic AI built on Microsoft Azure OpenAI GPT service trained on city web pages but lacked proper safeguards and domain specific fine tuning. It also contradicted itself and gave inconsistent answers. The city added a disclaimer after the fact but users could ask the bot whether it was reliable and it responded yes, undermining the disclaimer. This case highlights several lessons:

This case highlights several lessons:

- Testing and Validation

AI responses must be reviewed by subject matter experts before deployment. Systems should be tested using many real queries and edge cases, and guardrails added before public launch.

Domain Specific Controls

Government systems should not rely on generic AI output for sensitive matters. Answers must come from verified sources such as approved FAQs, policies, or structured rules.

- Accountability

The agency remains responsible for AI outputs. Liability cannot be shifted to the system, and vendor responsibilities should be clearly defined contractually.

- Public Trust

Citizens treat government AI as authoritative. Therefore reliability standards must be very high and disclaimers alone are insufficient.

Legal risks also include decision making authority. If an AI automatically rejects an application, flags someone for investigation, or issues a fine, due process must be considered. Citizens may have a right to human review. In many jurisdictions important decisions cannot be fully automated because of due process rights.

Governments could face legal challenges if algorithmic decisions are unfair or opaque. Even where not legally required, the principle should be followed that a human decision maker remains formally responsible and AI acts only as input.

- Bias & discrimination:

AI can amplify biases in training data (e.g., policing or property valuation unfairly targeting certain groups). Systems must be tested for unequal impact, ideally with diverse stakeholders and performance checks across demographics.

- Transparency & explainability:

Government decisions must be defensible. Black-box models weaken legal standing, so agencies should document factors considered or provide human-readable reasoning alongside AI recommendations.

- Hallucinations:

Language models may fabricate facts (e.g., wrong laws). Responses should be verified against official databases or restricted to approved reference sources. Responsibility lies with the deploying agency to implement safeguards.

- Legal liability:

The government is typically liable to citizens for AI errors, with possible claims against vendors depending on contracts. Mitigation includes fallback processes, rapid correction mechanisms, and insurance coverage.

- Public perception & ethics:

Be transparent when AI is used, label AI-assisted outputs, and provide channels for reporting errors to maintain trust.

- Appeals & redressal:

Citizens must have clear procedures to challenge AI-assisted decisions before disputes escalate.

- Privacy:

Linking datasets may violate data-sharing rules. Follow data protection laws, anonymize data where possible, and use synthetic data during pilots.

- Cybersecurity:

Protect against hacking, data leaks, model poisoning, and prompt injection. Avoid unnecessary direct access to sensitive databases.

- Record keeping:

Maintain secure logs of AI recommendations and final decisions for audits and information requests.

3.5 Risk Mitigation Practices

- Policy and SOPs: Develop a policy for AI use in your department. E.g., a guideline that “AI outputs must be verified by an officer before final decisions”, or “AI will not be used in isolation for decisions affecting rights or benefits.” Have Standard Operating Procedures for handling AI errors (like a playbook of what to do if the AI gives an obviously wrong answer or if a user complains).

- Training: Train your staff about the AI’s limitations. In NYC, perhaps front-line staff might have blindly trusted the chatbot themselves. Ensure your team knows that the AI can make mistakes and when to override it. They should also know how to explain to the public that a mistake is being corrected and what steps are taken.

- Continuous Improvement: Treat an AI system as a living thing that needs tuning. Use user feedback and error cases to refine it – either by improving the model, adding new rules, or simply updating training data. Many problems emerge after deployment, so plan a phase for monitoring and refining rather than “deploy and forget”.

- Legal counsel involvement: Involve your legal advisors in the AI project planning. They might point out specific laws or give language for disclaimers or terms of use. Also, if something goes wrong, involve them immediately to determine the right remedy and communication.

The National AI Policy emphasizes ethical governance and warns that technology adoption should not outpace legal capacity. Since laws evolve slower than technology, Apply existing administrative principles such as fairness, accountability, transparency, and the right to appeal to reduce legal risk.

At the same time, avoiding AI can also create governance risks such as failing to detect fraud.

The goal is balance: adopt AI to improve public service while protecting legal rights and institutional responsibility. Use a controlled, incremental approach with pilot projects, stakeholder consultation, and built in safeguards.

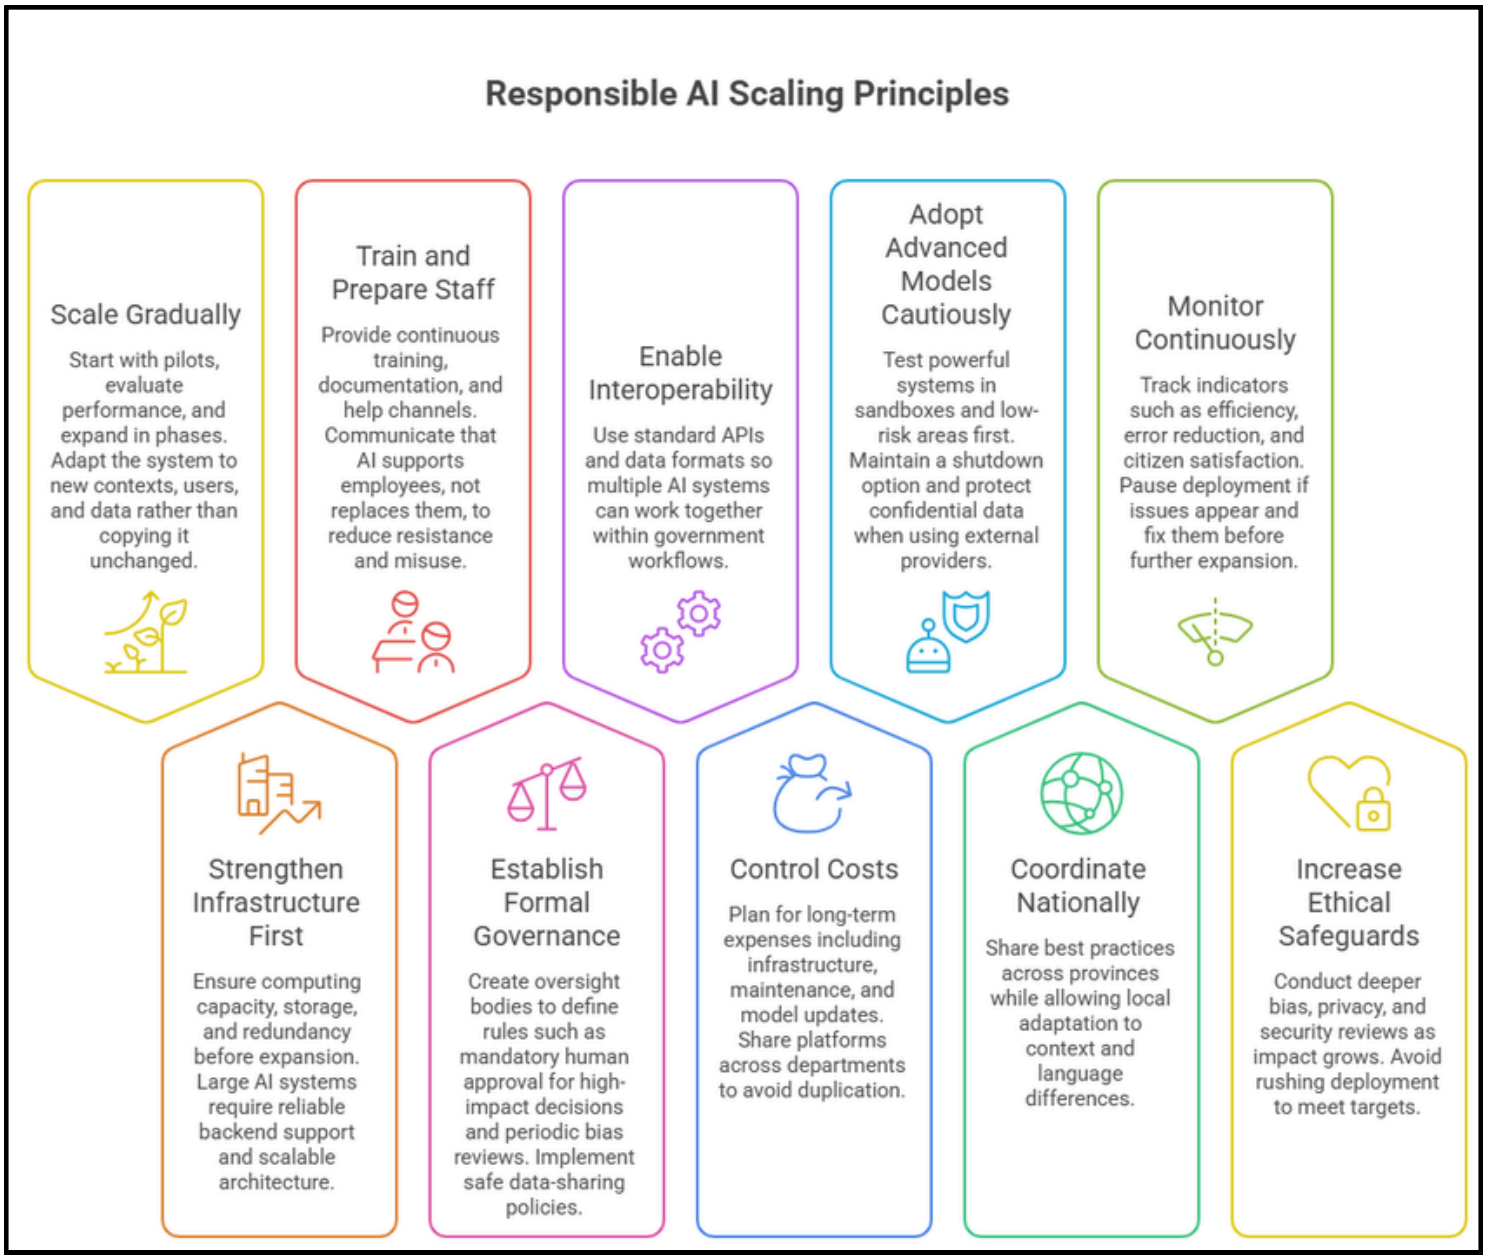
Key risks must be mitigated through testing, human oversight, bias audits, appeal mechanisms, and compliance with data and security laws.

AI should strengthen public service, not undermine it. To achieve this, deploy systems cautiously, retain human accountability in decision-making, and implement safeguards before scaling.

Once risks are identified and mitigation measures are in place, the next step is to think beyond individual use cases. Government AI initiatives rarely remain small pilots; successful solutions often expand across departments and services. At this stage, the challenge shifts from safe adoption to responsible expansion ensuring that governance, capacity, and accountability grow alongside the technology. This brings us to the question of scaling advanced AI systems.

3.6 Advanced AI Systems and Responsible Scaling

The journey from a pilot or localized AI solution to a scaled up, government wide system is where many projects stumble if not planned carefully. “Advanced AI systems” here refer to both the cutting edge technologies like large multimodal models, advanced autonomous agents and large scale deployments such as nationwide systems and multi department integrations. Let’s now discuss how to scale AI in government responsibly ensuring that as we amplify AI’s reach, we also amplify oversight, infrastructure, and skills correspondingly.


Responsible AI Scaling Principles
Scale Gradually
Start with pilots, evaluate performance, and expand in phases. Adapt the system to new contexts, users, and data rather than copying it unchanged.
Train and Prepare Staff
Provide continuous training, documentation, and help channels. Communicate that AI supports employees, not replaces them, to reduce resistance and misuse.
Enable Interoperability
Use standard APIs and data formats so multiple AI systems can work together within government workflows.
Adopt Advanced Models Cautiously
Test powerful systems in sandboxes and low-risk areas first. Maintain a shutdown option and protect confidential data when using external providers.
Monitor Continuously
Track indicators such as efficiency, error reduction, and citizen satisfaction. Pause deployment if issues appear and fix them before further expansion.
Strengthen Infrastructure First
Ensure computing capacity, storage, and redundancy before expansion. Large AI systems require reliable backend support and scalable architecture.
Establish Formal Governance
Create oversight bodies to define rules such as mandatory human approval for high-impact decisions and periodic bias reviews. Implement safe data-sharing policies.
Control Costs
Plan for long-term expenses including infrastructure, maintenance, and model updates. Share platforms across departments to avoid duplication.
Coordinate Nationally
Share best practices across provinces while allowing local adaptation to context and language differences.
Increase Ethical Safeguards
Conduct deeper bias, privacy, and security reviews as impact grows. Avoid rushing deployment to meet targets.

Scale gradually

Start with pilots, evaluate performance, and expand in phases. Adapt the system to new contexts, users, and data rather than copying it unchanged.

Strengthen infrastructure first

Ensure computing capacity, storage, and redundancy before expansion. Large AI systems require reliable backend support and scalable architecture.

Train and prepare staff

Provide continuous training, documentation, and help channels. Communicate that AI supports employees, not replaces them, to reduce resistance and misuse.

Establish formal governance

Create oversight bodies to define rules such as mandatory human approval for high-impact decisions and periodic bias reviews. Implement safe data-sharing policies.

Enable interoperability

Use standard APIs and data formats so multiple AI systems can work together within government workflows.

Control costs

Plan for long-term expenses including infrastructure, maintenance, and model updates. Share platforms across departments to avoid duplication.

Adopt advanced models cautiously

Test powerful systems in sandboxes and low-risk areas first. Maintain a shutdown option and protect confidential data when using external providers.

Coordinate nationally

Share best practices across provinces while allowing local adaptation to context and language differences.

Monitor continuously

Track indicators such as efficiency, error reduction, and citizen satisfaction. Pause deployment if issues appear and fix them before further expansion.

Increase ethical safeguards with scale

Conduct deeper bias, privacy, and security reviews as impact grows. Avoid rushing deployment to meet targets.

Core rule

Responsible scaling means expanding technology, governance, infrastructure, and accountability together to ensure safe and sustainable AI adoption in government.

Example of Scaling Journey

Example: e-Katcheri AI Assistant

1. Start with a pilot in one district where AI summarizes citizen complaints for the Deputy Commissioner.
2. Expand to all districts after validating accuracy and training the model on diverse data and dialects.
3. Integrate a provincial dashboard to aggregate and analyze issues.
4. Connect to a national monitoring dashboard.
5. Add advanced features such as voice input only after earlier stages are stable.

At each stage ensure staff training, data pipeline reliability, and model accuracy before further expansion. Avoid deploying all features at once.

Scaling AI is not only a technical expansion. Governance, infrastructure, training, and accountability must grow alongside the system. Follow national policies and controlled rollout practices to maintain quality and public trust. Responsible scaling enables efficiency and innovation while preserving legal and institutional control.

Key Takeaways

Secure deployment is a governance responsibility, not only a technical task. Choosing the correct infrastructure, enforcing safeguards, and maintaining oversight ensures AI systems remain lawful, trustworthy, and sustainable. When deployment decisions align with policy and risk management, AI strengthens public service delivery without compromising citizen rights or institutional credibility.

Moving Forward

The next step is to see how these principles apply in real administrative environments. The following section illustrates practical applications within district government operations, focusing on land records management and Safe City projects, where sensitive data handling and human oversight are essential.

MODULE 4: PUBLIC SECTOR APPLICATIONS AND POLICY CONTEXT

This module connects technical AI understanding with real government implementation. It explains how AI systems operate within administrative structures, legal frameworks, and national policy directions. Participants will learn how AI supports decision-making, service delivery, and inter-department coordination while remaining aligned with Pakistan's regulatory environment and institutional responsibilities.

Learning Outcomes

- Identify public sector use cases where AI can improve efficiency and service delivery
- Interpret national policies (e.g., Cloud First and National AI Policy) in practical project decisions

4.1 Use Cases in District Government: Land Records and Safe City Projects

Now that we have covered the technology and governance fundamentals, let us ground the discussion in concrete use cases relevant to district governments in Pakistan. Two areas stand out due to ongoing initiatives and high impact potential:

1. **Computer Vision in Land Records Management:** shows how AI can assist in digitizing and maintaining land records as handled by entities like the Punjab Land Records Authority PLRA and technology partners like PITB.
2. **AI in Safe City Projects:** explaining how advanced surveillance and analytics are used for urban security and management as seen in Punjab Safe Cities projects.

By examining these, we will see how the concepts of AI 201 come together, including integration with existing systems, on-prem versus cloud choices, legal considerations, and scaling challenges.

Case 1: AI and Computer Vision in Land Records

Land records in Pakistan were historically paper-based and handwritten. Systems such as LRMIS and PLRA digitized millions of records, but scanning, entry, and verification remain labor-intensive.

How AI Helps

1. Document digitization (OCR)
Use AI-based OCR to read Urdu handwritten records and forms.
Extract names, plot numbers, and fields automatically.
Humans verify low-confidence entries instead of typing everything manually.

2. Form and map understanding (Computer Vision)

Detect fields in mutation and registry forms.
Identify plot numbers and boundaries in Mussavi maps.
Align old maps with satellite imagery.

3. Fraud detection

Compare seals and signatures with verified samples.
Detect tampering such as altered ink or edits.
Flag suspicious records for officer review.

4. Search and analytics

Enable natural language queries such as retrieving all parcels owned by a person in a tehsil.

Detect duplicates and inconsistencies using fuzzy matching and clustering.

5. Satellite monitoring

Classify land use from satellite or drone imagery.
Detect encroachments or illegal construction.
Support projects such as urban mapping and planning.

6.Deployment requirements

Host systems on government infrastructure or accredited secure cloud.
Avoid sending sensitive records to public APIs.
Maintain human verification and confidence scoring.

Benefits

- Faster service delivery, reduced manual work, and improved data quality.
- Staff focus on verification and complex cases instead of repetitive tasks.

Key risk control

- Never update legal records automatically.
- AI assists analysis; officers approve final changes.

Case 2: AI in Safe City and Public Safety Project

Safe City initiatives use large networks of cameras and sensors to improve policing, traffic control, and municipal safety. AI and computer vision analyze live video feeds to support faster and more informed responses.

Core applications

1. Facial recognition
 Match faces from CCTV footage against approved criminal databases.
 Generate alerts for operators to verify before action.
 Used to track suspects or locate missing persons.

2. Number plate recognition (ANPR)
 Read vehicle plates automatically from cameras.
 Support e-challans, stolen vehicle detection, and crime investigations.

3. Activity and hazard detection
 Identify fights, weapons, unusual gatherings, abandoned objects, smoke, or infrastructure hazards.
 Notify operators for human review and response.

4. Traffic management
 Analyze traffic flow and adjust signals or routing.
 Detect violations such as wrong-way driving.

5. Analytics and planning
 Produce heatmaps of accidents or crime patterns.
 Help authorities allocate patrols and improve road safety measures.

Deployment approach

Processing occurs mainly on government-controlled servers connected to city command centers.

Systems integrate with police and transport databases through APIs.

Operators verify alerts before any enforcement action.

Benefits

- Faster suspect identification, automated violation detection, and improved emergency response.
- Better planning through data-driven insights.

Key safeguards

- Keep data within controlled infrastructure.
- Require human verification before enforcement.
- Use AI as decision support, not autonomous authority.

Outcome

AI strengthens urban safety operations while maintaining human oversight and accountability.

However, benefits are accompanied by risks and controversies such as:

- **Privacy** concerns arise due to mass surveillance capabilities. Even if systems are designed to match only criminal databases, misuse remains possible without proper authorization and audit trails. Strong governance policies and query logging are necessary to prevent abuse.

- **Accuracy and bias** must be considered because facial recognition may produce false positives. Misidentification could lead to innocent individuals being questioned. Therefore every AI alert must be verified by a human officer before action.

- Over reliance is another risk. Officers may trust AI blindly, so training and procedures must emphasize verification standards.

Technically, Safe City systems process continuous video streams and require specialized hardware such as GPU clusters for real time response. Low latency is essential, so processing often occurs near the data source.

- Pakistan's Safe City initiatives are expanding across provinces, creating potential for interconnected systems in the future. AI becomes essential to handle the massive volume of data because human monitoring alone is impossible.

Beyond crime prevention, these systems support public safety and city management. Traffic enforcement improves road discipline, crowd monitoring reduces accident risks, and emergency detection can trigger faster response.

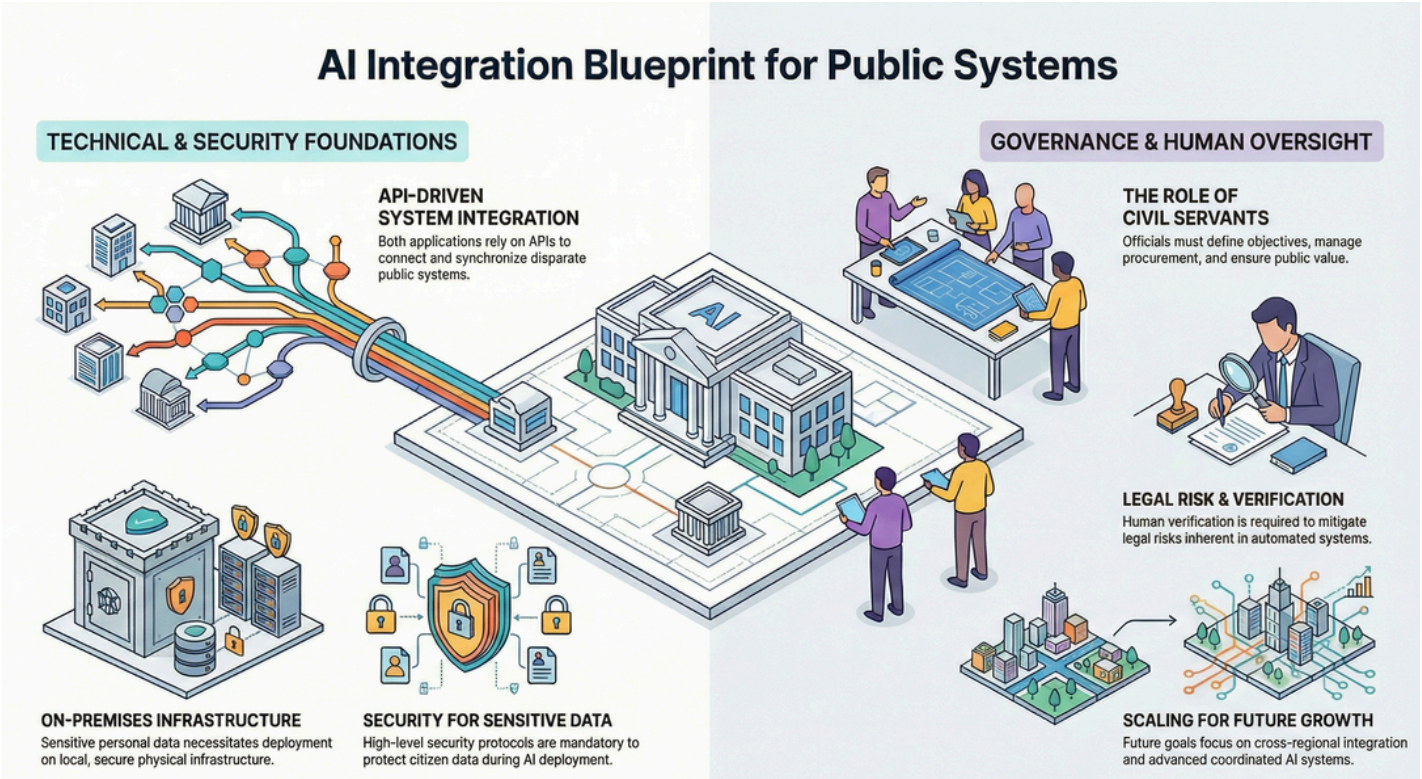
At district level, smaller deployments may include CCTV monitoring around courts or jails, or drone surveillance for detecting illegal activities such as illicit crop cultivation.

Safe City AI demonstrates a semi autonomous monitoring workflow where systems observe, analyze, and alert humans who make final decisions. It combines automation with human oversight and requires strong governance due to privacy and legal implications.

Similarity between both cases

Both land records and Safe City applications demonstrate practical implementation of AI concepts. They rely on system integration through APIs, require secure deployment due to sensitive personal data, favor on premises infrastructure, carry legal risks requiring verification, and depend on civil servants to define objectives, procure solutions, and ensure public value while protecting rights.

These applications align with national policy goals of improving governance and safety. As systems scale further, future challenges will involve integrating them across regions and building more advanced coordinated AI systems.



4.2 Policy Overview: National AI Policy 2025 and Cloud First Policy 2022

To wrap up AI 201, it is important to anchor our applied discussions in the framework of Pakistan's official policies on technology. Two key policies have been referenced throughout this manual:

1. Pakistan's National Artificial Intelligence Policy 2025 approved in 2023/24
2. Pakistan Cloud First Policy 2022 approved in February 2022

Understanding these policies will help you ensure your AI projects align with national objectives, comply with regulations, and leverage any support mechanisms available. Below is a summary of each, highlighting points most relevant to a civil servant overseeing AI deployments.

National AI Policy 2025 Key Points for Government AI Deployment

Pakistan's National AI Policy 2025 is a comprehensive roadmap to promote AI innovation, build capacity, and ensure ethical use of AI for national development. It is structured around six foundational pillars:

1. AI Innovation Ecosystem

Establishing resources and institutions to foster AI. This includes creating a National AI Fund and Centers of Excellence in AI. For you, this means there may be funding opportunities or expert centers to tap into for your projects. For example, a Center of Excellence might help fine tune a model for your department or the National AI Fund might co fund a pilot project in your district. It also mentions connecting academia and industry which implies partnerships such as collaborating with a university AI lab on a local problem.

2. Awareness and Readiness

Massive AI literacy and training programs. The policy aims to train 1 million people with 10,000 trainers by 2027, with special inclusion for women and differently abled individuals. As a civil servant, you can expect more training opportunities. It also means your staff are expected to upskill and you should encourage participation. It signals that AI is becoming a foundational skill similar to computer literacy in the past.

3. Secure AI Ecosystem

Focus on cybersecurity, ethical AI, data privacy, and human oversight. The policy mentions regulatory sandboxes allowing government agencies to experiment with AI in controlled environments with limited bureaucratic barriers. This pillar reinforces that security and ethics must be built into projects. It may also lead to specific national AI ethics guidelines.

4. Transformation and Evolution

Adopting AI across key sectors including education, healthcare, governance, agriculture, manufacturing, and energy. Governance is explicitly listed, meaning district administration is a focus area. Expect targets such as implementing AI assistants in citizen service centers. Land records and safe city projects align directly with this pillar.

5. AI Infrastructure

Development of national computing infrastructure, high performance computing centers, open data repositories, and encouragement of cloud and open source platforms. This means government may provide shared AI infrastructure, government cloud APIs, or national datasets. It also encourages avoiding vendor lock in and building local technical capacity.

6. International Partnerships

Global cooperation for research, standards, and investment. This may result in international training programs, donor funded projects, or adoption of international AI ethics standards.

The policy also outlines an implementation and governance model including an AI Council chaired by the IT Minister and a Policy Implementation Cell in MoITT for execution and monitoring KPIs. This means centralized oversight exists and departments may report project metrics. The Cell can also provide guidance on regulatory or data sharing questions.

Key Performance Indicators (KPIs) include number of AI startups, research initiatives, patents, trained workforce, infrastructure expansion, and public sector service improvements such as reduced service delivery time. Projects should align with measurable outcomes such as faster processing of land records.

The policy vision emphasizes an ethical and inclusive AI ecosystem that safeguards citizens' rights and supports sustainable development. Projects should therefore consider accessibility, language inclusion, and avoidance of inequality.

Overall, the National AI Policy provides a supportive environment for adoption while enforcing guardrails around ethics, privacy, and oversight. Civil servants can use it to justify budgets, obtain approvals, and align initiatives with national priorities.

Pakistan Cloud First Policy 2022 Key Points for AI Deployment

The Cloud First Policy is narrower in scope than the AI Policy but critically important for any IT deployment in government, including AI. It essentially mandates that government agencies should consider cloud solutions as the primary option for new IT investments to improve cost efficiency, scalability, and interoperability, provided that security and data sovereignty requirements are met.

Key elements of Cloud First Policy

Cloud Office and Accreditation

A Cloud Office in MoITT is established to oversee adoption. This office will accredit Cloud Service Providers for government use. As mentioned earlier, only accredited CSPs should be used, especially for non public data. The policy encourages the local cloud industry as well. So when deploying an AI system, if you want to use PTCL's cloud or a private data center, check if it is accredited or register it with the Cloud Office. The policy likely has a list or process for accreditation fulfilling criteria of security and data location.

Data Classification

The policy provides guidelines dividing data into Open, Public, Restricted, Sensitive, and Secret. A cloud selection matrix links data class to type of cloud and required security level.

Open or Public data can go on Public Cloud including foreign providers if the CSP is registered or accredited with baseline security.

Restricted data should be on a Government Community Cloud or public cloud with intermediate security meaning data remains in Pakistan and is handled by vetted entities.

Sensitive or Confidential data should be on Government Cloud with enhanced security implying not on general multi tenant public cloud unless it is a dedicated government cloud. This could include infrastructure like NTC or PITB high security data centers.

Secret data requires private or government cloud with highest security and isolated infrastructure.

When planning an AI system, classify its data and choose infrastructure accordingly. This may require avoiding convenient public services and instead using controlled environments. If a necessary service only exists on a global cloud, an exemption or special arrangement through the Cloud Office may be required.

Centralized Procurement

The policy suggests pooling procurement of cloud services. Instead of departments independently purchasing subscriptions, there could be unified contracts negotiated centrally. This means departments should check with MoITT or their ministry before independently subscribing to any cloud service and follow proper procurement procedures.

Data Sovereignty and Local Clouds

The policy emphasizes maintaining government data sovereignty and preference for local hosting and Pakistani jurisdiction. If using international providers, data should ideally remain in regional servers rather than foreign jurisdictions. Government cloud platforms run by NTC or NITB may be the default choice for hosting AI applications.

Security and Standards

The policy calls for a formal security framework for cloud usage aligned with recognized standards. This includes encryption, identity management, and strict access controls similar to or stronger than on premises systems.

No Fragmented Infrastructure

Departments are discouraged from maintaining underutilized server rooms when shared cloud infrastructure is available. If dedicated hardware is proposed, justification must explain why cloud cannot be used, typically due to sensitivity or lack of accredited hosting options.

Contractual Guidelines

Contracts with cloud providers must define service levels, data ownership, recovery processes, and liability. Government retains ownership of data and providers must delete or return it when contracts end. These protections should also apply to AI vendors operating on cloud platforms.

Illustrative compliance scenarios

- If a Safe City AI system processes video feeds classified as Sensitive or Secret, processing must remain on premises or within a secure government cloud rather than a general public cloud.
- If anonymized traffic statistics are released as open data, they may be hosted on a public cloud portal.
- If a provincial IT board builds a shared AI platform, departments should adopt the shared service instead of building independent solutions to promote interoperability and efficiency.

Practical checklist for compliance

1. Determine data classification of AI data
2. Choose an appropriate cloud environment or justify on premises deployment
3. Obtain Cloud Office approval if required
4. Include protective terms in contracts
5. Document that cloud options were evaluated

Staying aligned ensures smoother approvals since planning documents increasingly require Cloud First compliance and MoITT clearance for infrastructure.

By embedding policy considerations into planning, projects face less friction and scale more easily through shared platforms and interoperable systems.

Both the AI Policy and Cloud First Policy ultimately support a modern, efficient, and secure digital government. They provide funding, infrastructure, and guidance while AI projects deliver practical implementation aligned with national strategy.

Conclusion and Next Steps

In this AI 201 manual, we journeyed through the practical landscape of deploying AI systems in government from understanding your role as an AI project lead, exploring no code tools and integration, grappling with deployment models and security, to examining real world use cases and policies. The consistent theme has been balance balancing innovation with governance, leveraging technology while safeguarding ethics and law, and aiming for efficiency while maintaining human oversight.

For a civil servant in Pakistan, the promise of AI is significant. You can vastly improve service delivery imagine cutting down a complaint resolution from months to days, or detecting issues before they escalate using predictive analytics. You can make your departments more proactive and citizen friendly. But with that comes the responsibility to implement these technologies carefully.

It is heartening to see that the Government of Pakistan, through policies like the National AI Policy and Cloud First, is not only encouraging AI adoption but also putting frameworks in place to do it right. This means you are not alone there is a whole ecosystem being built to support you including training programs, centers of excellence, funding, and guidelines.

As you move forward:

- **Identify Opportunities**

Look around in your department what processes are repetitive, data heavy, or decision intensive that could benefit from AI or automation. Perhaps it is automating the initial scrutiny of applications, or using NLP to sort incoming mail, or computer vision to audit field work via images. Keep the use case mindset start with a problem and see if AI can solve it rather than adopting technology for its own sake.

- **Build the Business Case**

Use the knowledge from AI 201 to articulate the benefits such as time saved, cost saved, better compliance and also the requirements such as data needed and training. This will help in securing approvals and resources. Tie your case to policy goals for extra strength for example this project will contribute to a pillar of the National AI Policy by achieving a defined outcome.

- **Engage Stakeholders**

Bring IT staff, vendors, legal advisors, and end user representatives together early in project planning. AI projects are multidisciplinary. For instance involve data entry clerks in an OCR project they know the quirks of the documents and can help train the AI or at least trust it more if they had input. In Safe City context involve police from the start so the system fits operational needs.

- **Plan for Change Management**

Communicate to your team why an AI system is being introduced, what it will do, and how roles might shift. Some tasks will be automated but new tasks often emerge like overseeing the AI or handling exceptions. Emphasize training a workforce that grows with the technology is an asset. Consider incentives such as recognition for those who champion new tools.

- **Pilot and Iterate**

Wherever feasible do a pilot. Measure results and resolve issues. Then scale gradually. Government systems often deal with millions of people and a small glitch at scale can become a major issue. Pilots help reveal local context issues such as an AI trained on English struggling when many inputs are in roman Urdu which you can fix by adding training data or rules.

- **Governance and Ethics as Core**

Set up protocols from day one including data privacy measures, bias checks, and documentation of decisions. It is much harder to retrofit ethics into a system than to build it in. For example if you plan a chatbot decide on a moderation policy in advance. If deploying predictive policing put guidelines in place about how predictions are used they should not be treated as evidence but as investigative leads. These guardrails maintain public trust and protect you from misuse.

- **Monitor and Maintain**

Launch is not the end it is a new beginning. Monitor performance metrics and investigate issues such as data drift, model issues, or implementation problems. Solicit feedback from citizens and staff. Continuously improve the system. Later levels will cover fine tuning models and advanced improvements you can apply.

Looking ahead, **AI 301 Government Grade Models and Fine Tuning** will likely delve into customizing AI for the government context such as training local language models, fine tuning on government documents, and managing model life cycles. That will enable Government Grade AI systems that understand Pakistani laws, languages, and administrative procedures deeply. **AI 401 National AI Infrastructure and Data Centers** will likely cover large scale infrastructure including data centers, cloud, networks, and sovereign hosting to support government AI efforts. These build on the foundations covered here by adding technical depth and nationwide strategy.

As a final note remember that AI is a tool but governance is ultimately about people. Use AI to augment teams, serve citizens better, and make informed decisions rather than replacing the human touch in public service. The Minister of IT's message in AI 101 highlighted that AI is a present capability that must be understood, governed, and applied with care. Apply the knowledge from this manual with prudence and foresight and you will be on your way to becoming an effective AI enabled leader in the civil services, modernizing governance while upholding transparency, fairness, and accountability.

TECH destiNATION PAKISTAN
MINISTRY OF IT & TELECOM

atomcamp

Ministry of Planning Development & Special Initiatives

The Artificial Intelligence 201 for Public Sector program is a joint initiative by the Ministry of Information Technology and Telecommunication, Ministry of Planning, Development and Special Initiatives, Civil Services Academy, and atomcamp, aimed at equipping Civil Services Academy probationers with essential AI awareness for effective governance and policy-making.

For further information, please contact the Civil Services Academy, Walton Lahore

Training Manual

A course on

ARTIFICIAL INTELLIGENCE FOR Pakistan Civil Servants

301

TECH destiNATION PAKISTAN
MINISTRY OF IT & TELECOM

Ministry of Planning Development & Special Initiatives

atomcamp

Table of Contents

- Executive summary ---- 148
- Operating Context and Non-Negotiables ---- 149
- LLM and SLM Fundamentals ---- 151
- SLMs: Small Language Models ---- 152
- Fine-Tuning and RAG Engineering Playbook ---- 153
- Deployment and Integration in Existing Government Systems ---- 156
- Governance Controls ---- 160
- Use cases tailored to Pakistan's current maturity ---- 161

EXECUTIVE SUMMARY

This manual specifies a government-grade operating blueprint for deploying language-model systems in Pakistan's public administration, explicitly building on the AI maturity path already established in the existing AI 101 & AI 201 training manual (better prompting → RAG/context injection → fine-tuning → sovereign/local models).

The core design position is that **public-sector value is created through controlled context and controlled deployment**, not through "bigger models" by default. That framing is consistent with the governance-first orientation in the prior manual (AI as decision support with continuous human accountability; strict internal-data handling; role-based access; and an endgame of sovereign models).

This "AI 301" manual has two lanes with one shared control plane:

- **Lane A (civil servant / product owner / governance leadership):** model selection, risk gating, data classification, procurement controls, auditability, and deployment choices aligned with policy. The manual maps these to Pakistan's national policy stack (national AI policy; cloud-first; cybersecurity; and draft data protection regime).

- **Lane B (technical delivery teams):** transformer internals, SLM vs LLM trade-offs, RAG engineering patterns, fine-tuning methods (full FT, adapters, LoRA/QLoRA, instruction tuning, preference optimization), evaluation harnesses, and "local deployment reality" (GPUs, VRAM, KV cache, quantization, batching, latency/throughput, and on prem/hybrid/air-gapped designs).

The manual is optimized for Pakistan's present constraints: broadband adoption is growing but uneven; operational reliability varies across tiers of government; and many sites will not have always-on high-compute infrastructure.

OPERATING CONTEXT AND NON-NEGOTIABLES

This manual assumes the institutional posture defined in the earlier material: **AI outputs are advisory; accountability and final decisions remain human; and sensitive government data must not be exposed to uncontrolled external systems.**

This posture is not "soft guidance"; it is the minimum viable control baseline for public-interest legitimacy in high-stakes workflows.

Policy alignment baseline (Pakistan):
Pakistan's national AI policy explicitly positions AI adoption as policy-driven and links responsible AI use to personal data protection, cloud-first policy, and the broader digital transformation agenda.

Pakistan's cloud-first policy applies to federal public sector entities making new ICT investments, and targets time-to-procure/time-to-launch and security improvements via cloud adoption and pre-accredited CSPs.
Pakistan's national cybersecurity policy explicitly emphasizes governance, coordinated response, and standards-based controls (including references to frameworks such as ISO 27001).

Pakistan's MoITT draft Personal Data Protection Bill (2023) is available as an official policy artifact and defines propositions around lawful collection, consent, anonymization, and a proposed data protection commission; it must be treated as a policy direction even where legal status/commencement remains in flux.

Digital infrastructure reality (design assumptions, not excuses):
Pakistan's broadband footprint is substantial at the national level, but service maturity and operational reliability vary by geography and agency tier.
Energy-sector performance reporting documents persistent reliability and operational challenges in distribution, which is directly material for "always-on AI" assumptions (especially outside major metros).

Non-negotiable control objectives (COs):

These are the control objectives every **AI 301** implementation must satisfy before any production go-live.

1. **Data sovereignty by design:** the system must enforce where data can flow, be stored, and be processed, using classification gates and explicit deployment choices.
2. **Approved-source constraint:** where factual accuracy matters, the system must ground outputs on approved sources (RAG) and must be able to cite them internally.
3. **Role-based access + auditability:** every query and retrieval must be attributable to a user identity and role; all outputs must be logged with metadata and retention rules. (This is both a governance and a cyber posture requirement.)
4. **Human-in-the-loop checkpointing:** any workflow that affects rights, penalties, entitlements, enforcement, or adjudication must include a mandatory human approval gate.
5. **Prompt and tool security:** AI-facing inputs and tool outputs require explicit protections against prompt injection, insecure output handling, training-data poisoning, and model denial of service.
6. **Operational controls:** rate limits, token budgets, and predictable performance SLAs must exist to avoid uncontrolled OPEX and service degradation.

Institutional roles (minimum viable operating model):
A production-grade LLM program is not "an IT project"; it is a cross-functional operating model with an accountable single-threaded owner.

- **AI Product Owner (civil servant):** defines policy constraints, approves use cases, owns risk register (per deployment).
- **Data Steward:** owns data classification, lineage, and quality gates; approves "what enters RAG".
- **Security Officer / CISO delegate:** owns threat model, authorizes deployment boundary, ensures secure logging and access controls.
- **ML/Platform Lead:** owns quantization strategy, inference stack, fine-tuning pipeline, monitoring, rollback.

LLM AND SLM FUNDAMENTALS FOR GOVERNMENT DEPLOYMENT

What an LLM actually is, mechanically?

Most government stakeholders can describe LLMs as "chatbots." AI 301 requires the operationally correct mental model:

- A modern LLM is typically a **Transformer** trained to predict the next token of text, using attention-based blocks (self-attention + feed-forward components) stacked across many layers.
- The Transformer's core innovation is that it replaces recurrence with attention, enabling parallel training over sequences.
- During inference, generation is **autoregressive**: it emits one token at a time and conditions the next token on prior tokens. This is why token budgets and KV-cache memory dominate runtime economics.

Tokenization, embeddings, and context windows, in deployable terms

Tokenization (why "tokens" drive both capability and cost):
LLMs do not read "words"; they read token IDs produced by a tokenizer. Byte Pair Encoding (BPE) and SentencePiece-style tokenization are standard approaches that create subword units so the model can represent open vocabulary efficiently.
Operational implication: if your workload includes Urdu + English + Roman Urdu, expect token counts to increase and performance to degrade unless you select a tokenizer/model that covers that distribution well (or you normalize/standardize inputs upstream).

Embeddings (what makes RAG possible):
Embeddings are vector representations used to retrieve semantically similar items. Modern RAG systems embed documents and queries into a vector space and do similarity search over them. Retrieval infrastructure commonly relies on vector similarity libraries (e.g., FAISS) and ANN indexing algorithms (e.g., HNSW).

A pragmatic “Pakistan-ready” pattern is to host embeddings in a database platform the government already knows how to operate (e.g., PostgreSQL via pgvector) to minimize new infrastructure surface area.

Context window (why long documents are a trap):
Even when a model supports a large context window, inserting entire reports or large policy manuals can reduce answer quality and increase costs. That is not theoretical: the retrieval tooling ecosystem itself highlights chunking as a practical necessity for long documents.
In government, this translates into a design directive: **never treat “uploading a PDF” as your governance strategy**—treat the knowledge base as a managed asset with version control, approved sources, and retrieval behavior you can test.

Inference performance: pre fill vs decode and why KV cache dominates LLM inference has two distinct phases:

- **Pre fill:** reading the prompt and building internal states.
- **Decode:** generating tokens one-by-one.

As sequence length and concurrency increase, the KV cache (attention keys/values stored for previous tokens) becomes a dominant memory and bandwidth driver. This is a primary reason LLM serving is frequently memory-bound rather than compute-bound at scale.

Engineering response: modern serving stacks implement KV-cache efficiency techniques such as PagedAttention (vLLM), which reduces KV cache waste and fragmentation and improves throughput in real serving workloads.
Separately, attention computation has quadratic time/memory complexity in sequence length, motivating optimized kernels such as FlashAttention for practical long-context training and inference.

SLMs: Small Language Models and sovereign systems are different levers
The term “SLM” is used operationally in two ways in government programs:

- **Small Language Models (SLMs):** smaller parameter models designed for local/edge deployment and lower cost. Microsoft positions Phi-3 as an SLM series and explicitly frames it as locally deployable.
- **Sovereign language models:** systems where training, hosting, and control stay within national or institutional boundaries (sovereignty is about jurisdiction and control, not parameter count). This is consistent with the earlier manual’s “sovereign LLMs” endgame framing.

A Pakistan-ready architecture often uses both concepts together:
SLMs for edge and triage; larger models only where needed; and sovereign hosting for sensitive domains.

FINE-TUNING AND RAG ENGINEERING PLAYBOOK

Decision logic: RAG vs fine-tuning vs hybrid

The earlier manual's maturity path correctly frames RAG and fine-tuning as distinct levers.
AI 301 adds a sharper decision rule for production programs:

Objective Make answers cite official sources and stay updated	Best primary lever RAG	Why Retrieval provides non-parametric memory and provenance	Government caveat Must enforce "approved sources" and versioning gates
Make outputs match bureaucratic tone / templates	Fine-tuning (LoRA/Adapter)	Adjusts style and response patterns	Fine-tuning does **not** guarantee factual correctness; still pair with RAG
Improve multi-step reasoning in a bounded domain	Hybrid	Use RAG for facts + tuned behavior for workflow	Requires eval harness and red-teaming before rollout
Reduce cost/latency for high-volume, low-risk tasks	SLM + rules + small RAG	Smaller models + constrained retrieval	Enforce unsafe-output controls and rate limits

Fine-tuning methods, with “what changes” and “what does not”

Full fine-tuning (FT): updates all parameters. It is expensive, hard to govern, and risky for data leakage unless the training data governance is mature. Practical for national-scale programs, not departmental pilots.
Adapters (parameter-efficient transfer learning): inserts small trainable modules while keeping the base model fixed, giving modularity and lower per-task storage cost.

LoRA: freezes the base weights and injects low-rank trainable matrices into the transformer layers. It is specifically designed to reduce trainable parameters and memory requirements compared to full fine-tuning, with strong empirical performance.

QLoRA: combines 4-bit quantization of the base model with LoRA training; the original paper demonstrates the feasibility of fine-tuning very large models on a single 48GB GPU under this approach.

Instruction tuning / SFT: fine-tunes the model on instruction-following datasets to improve task generalization and helpfulness; FLAN is a widely cited instruction-tuning approach.

Preference optimization (RLHF and alternatives):
InstructGPT shows RLHF as a pipeline for aligning behavior with human preferences.
Direct Preference Optimization (DPO) is an RL-free method that optimizes preference objectives with a simpler supervised-style loss and is widely used when teams want alignment without the complexity of full RLHF.

High-integrity guidance for civil servants:
Fine-tuning governs how the model talks and follows workflows; RAG governs what the model is allowed to know for official use. This separation is the simplest way to preserve auditability while reducing hallucination exposure.

Fine-tuning workflow for government data: a stage-gated pipeline

Fine-tuning introduces new risks that standard IT procurement templates do not cover (memorization, leakage, poisoning, distribution shift). OWASP explicitly calls out training data poisoning as a top risk category for LLM applications.

A government-grade fine-tuning pipeline therefore requires explicit stage gates:
flowchart TD

- A[Use-case charter approved] --> B
- [Data classification & legal basis confirmed]
- B --> C[Training dataset created from approved sources]

- C --> D[PII/secret handling: redact/anonymize where required]
- D --> E[Baseline evals: accuracy, safety, bias, leakage tests]
- E --> F{Choose tuning method}
- F -->|LoRA/QLoRA| G[Parameter-efficient tuning run]
- F -->|Full FT| H[Controlled training cluster + stricter controls]
- G --> I[Post-tune evals + regression tests]
- H --> I
- I --> J[Security review: prompt injection, data exfil pathways]
- J --> K[Operational readiness review: monitoring, rollback, rate limits]
- K --> L[Limited pilot deployment]
- L --> M[Scale decision with KPI pack]

Pakistan-specific dataset guidance (practical):
If digital records are inconsistent, do not start with raw digitized archives. Start with a "gold corpus" composed of already-approved templates and finalized documents with known provenance (e.g., standing SOPs, published notifications, finalized policy briefs, standardized letters). This matches the earlier manual's focus on "approved internal sources" and governance-first sequencing.

RAG engineering beyond basics: advanced patterns that actually matter

The RAG paper defines the fundamental concept: combine parametric memory (the model) with non-parametric memory (retrieved documents) to improve knowledge-intensive tasks.
AI 301 focuses on the production patterns that convert RAG from "demo" into "government utility."

Pattern 1: Retrieval provenance as a compliance artifact

In government, "sources" are not a UX feature; they are an audit artifact. Every generated paragraph should be traceable to retrieval evidence, with stored document IDs + hash/version + access policy. This is directly aligned with the earlier manual's RAG-as-control framing.

Pattern 2: Hybrid retrieval (BM25 + embeddings) for messy corpora

If departmental content is low-quality scans, mixed Urdu/English, and inconsistent metadata, embedding-only retrieval will miss critical clauses. Hybrid retrieval (lexical + semantic) is the default in low-maturity corpora, not an "advanced option." (This is operationally consistent with the broader retrieval field's separation of ANN vector search and general information retrieval methods.)

Pattern 3: HNSW vs brute force and why it matters

At scale, approximate nearest neighbor search is necessary for performance; HNSW is a standard graph-based approach for ANN with strong performance characteristics. If you deploy on district-level servers, the trade-off is not "accuracy vs speed"; it is "acceptable recall vs stable latency under constrained hardware."

Pattern 4: Multi-vector / parent-child indexing
A practical tactic for long policy documents: store both (a) small chunks for retrieval precision and (b) parent document context for coherence. Retrieval frameworks explicitly document multi-vector and parent-document approaches as common production patterns.

Pattern 5: Agentic RAG with strict tool governance
Agentic retrieval frameworks let the model decide when to retrieve and how to iterate. This can improve complex workflows but expands the attack surface (tool injection, data exfiltration, excessive agency). Some retrieval frameworks now explicitly describe "agentic RAG" as an approach.
Government control response: agentic RAG is allowed only behind a "tool policy engine" that enforces what tools can be called, what data can be accessed, and what outputs can be emitted.

A government-grade RAG architecture with controls

- flowchart LR
- U[User / Officer] --> G[LLM Gateway\nAuthN/AuthZ + token budget]
- G --> P[Prompt policy engine\nredaction + injection filters]
- P --> Q[Query builder\nrewrite + translation if needed]
- Q --> R[Retriever\nlexical + vector]
- R --> RR[Reranker\npolicy-aware]
- RR --> C[Context packer\ncitations + max tokens]
- C --> M[Generation model\nLLM or SLM]
- M --> O[Output controls\nPII filtering + safe formatting]
- O --> A[Audit log\nprompt+retrieval+output metadata]
- R --> KB[(Approved knowledge base\nversioned)]
- A --> SIEM[(Security monitoring)]

This architecture explicitly maps to OWASP LLM risks: prompt injection, insecure output handling, data poisoning, and DoS are managed through the gateway, policy engine, and audit controls.

DEPLOYMENT AND INTEGRATION IN EXISTING GOVERNMENT SYSTEMS

Deployment models and the "data-classification gate"

Pakistan's cloud-first policy pushes adoption for public sector entities, but government-grade AI requires a classification-first routing strategy: the same application can route different workloads to different deployment zones based on sensitivity.
Define four deployment zones (Z):

- Z1 Public/low-risk: public cloud LLM APIs acceptable if data is non-sensitive and outputs are not used as official records.
- Z2 Internal moderate risk: private cloud / sovereign cloud / managed services with contractual controls, encryption, logging.
- Z3 Confidential: on-prem or dedicated sovereign infrastructure;

no external model training; RAG only from approved internal sources.

- **Z4 Restricted/air-gapped:** offline inference + offline retrieval updates; physically controlled; manual update process.

This gate is consistent with the earlier manual's "data never leaves the organization" design principle for in-house systems and with the broader cybersecurity governance posture.

Integration patterns that work in low-maturity environments

Most Pakistani government systems are not "API-first." AI 301 assumes mixed ecosystems: legacy databases, partial digitization, and manual workflows. The integration strategy must therefore be incremental.

Pattern A: Add an "LLM gateway" as a controlled façade
Do not embed model calls directly into multiple applications. Centralize through a gateway that enforces: identity, role-based retrieval permissions, token budgets, model allow-lists, and logging. This aligns with security control catalog thinking (central controls and auditability).

Pattern B: Retrieval service as middleware (not as a "chatbot feature")
RAG retrieval is a platform service: document ingestion → embedding → indexing → retrieval API. Keep it independent from any one use case so multiple ministries can reuse it without duplicating data pipelines. Pakistan's AI policy itself identifies "shared AI resources" and national AI platforms/shared APIs as policy objectives.

Pattern C: Use Postgres as "least-change" vector infrastructure
For agencies already running relational systems, pgvector allows vector similarity search within PostgreSQL, reducing operational sprawl.

Hardware and operations: GPUs, VRAM, quantization, and token economics

What matters operationally

- **Key facts that drive sizing: -**
- Attention cost grows sharply with sequence length; long-context inference is expensive and memory-heavy.
- KV cache memory grows with context length and concurrency; at higher batch sizes and sequence lengths, KV cache loads can dominate.
- Modern serving engines increase throughput through better KV cache management (PagedAttention).
- Quantization reduces memory footprint and can enable local deployment on smaller GPUs/CPUs, with measurable accuracy trade-offs; QLoRA formalizes 4-bit quantized training with LoRA.

Practical sizing heuristics for low-resource contexts

This is the minimal sizing logic technical teams should use when Pakistan-side compute is scarce.

Heuristic 1: Start with SLMs for "always-on" workloads
For citizen-facing triage, internal routing, and summarization at scale, an SLM-class model is often sufficient, with periodic escalation to a larger model only when needed. Microsoft explicitly markets Phi-3 as an SLM line and highlights local deployment.

Heuristic 2: Use quantized weights for inference, not for governance
Quantization is an inference optimization. Governance is still anchored in approved sources, role policies, and output controls.

Heuristic 3: Measure tokens as your capacity unit
In production, think in "tokens per second" and "tokens per request" rather than "requests per minute." This maps directly to model serving behavior and cost drivers.

Hardware reference table (usable for procurement & capacity planning)

Platform tier Low-cost workstation pilot	Example GPU RTX 4090	On-paper memory (VRAM) 24 GB	Fit-for-purpose in government workloads Small/medium models, quantized inference, limited LoRA fine-tuning for small models	Key limitation No ECC; constrained VRAM; enterprise reliability depends on build quality
Mid-tier inference server	L40S	48 GB	Shared departmental inference, higher concurrency, RAG + guardrails	Still bandwidth/VRAM constrained for very large models
Enterprise training/inference	A100	80 GB	Serious fine-tuning + high-scale inference; MIG partitioning optional	High CAPEX; power + cooling requirements
Top-tier scale	H100	80–94 GB	Large-model training/inference with high throughput; best for sustained concurrency	Highest CAPEX/OPEX; supply constraints

Procurement note (Pakistan-specific):
Because hardware supply and energy reliability are risk vectors, a procurement case must include **power and cooling plan, redundancy posture, and workload tiering**. Pakistan power-sector reporting illustrates that reliability metrics are tracked and vary by operator; treat this as a first-class risk input for on-prem deployments.

Local deployment stack: a minimal "government-grade" baseline

A government-grade local stack is not "run a model on a PC." It is an operational service:

- Inference engine that supports KV cache efficiency and concurrency (PagedAttention engines like vLLM are designed for that).
- Model packaging and quantized formats when needed for constrained hardware (local inference ecosystems document GGUF-style quantization workflows).
- Secure identity and least privilege (zero trust architecture is a recognized design posture for moving from perimeter-based to identity/resource-based controls).
- Centralized audit logs aligned to security and privacy control catalogs.

Evaluation, Security, and Governance Controls
Evaluation: move from "demo quality" to "evidence quality"

A government-grade system requires multi-metric evaluation: not just accuracy, but robustness, fairness, toxicity, and efficiency. HELM explicitly operationalizes this idea as a holistic evaluation approach across scenarios and metrics.

MMLU is a widely used multi-domain benchmark; it is useful for baseline capability comparisons but does not replace domain-specific evals (e.g., Pakistani administrative writing and law/policy constraints).

Minimum evaluation pack (MEP) per use case: - Groundedness / citation correctness (RAG only): does each claim map to retrieved evidence?

- **Safety tests:** toxic degeneration risk is a known property of pretrained models and is measured by datasets such as RealToxicityPrompts.
- **Security tests:** prompt injection, tool abuse, and sensitive information disclosure mapped to OWASP LLM Top 10 categories.
- **Regression tests:** ensure model updates or retrieval index updates do not break mission-critical workflows.

Governance frameworks that translate into controls

Risk management (AI-specific):
The National Institute of Standards and Technology AI RMF defines trustworthy AI characteristics (valid/reliable, safe, secure/resilient, accountable/transparent, explainable, privacy-enhanced, fair) and provides lifecycle functions (GOVERN, MAP, MEASURE, MANAGE).
NIST's generative AI profile extends this into GenAI-specific suggested actions and risk framing.

LLM application security:
The OWASP Top 10 for LLM Applications enumerates concrete risk categories (prompt injection, insecure output handling, training data poisoning, model DoS, etc.) and is directly usable as a control checklist for system design reviews.

Threat modeling (AI-specific):
MITRE ATLAS is a living knowledge base of adversary tactics and techniques targeting AI systems and is directly usable for red teaming and attack-surface mapping.

Identity and access (enterprise-grade):
Zero Trust is a published architecture model emphasizing user/asset/resource-centric controls rather than network-perimeter trust; it is relevant for LLM gateways and retrieval authorization.

Vendor management and procurement checklist (government-grade)

A government AI procurement package must include provisions that standard ICT tenders often miss. Pakistan's cloud-first policy already frames procurement efficiency and cloud adoption; AI procurement must now extend this with model/data controls**.**

Mandatory procurement clauses (practical, not aspirational):

- Data residency, retention, and deletion guarantees; explicit prohibition on training over government confidential inputs unless explicitly authorized.
- Security logging and audit data export (government owns logs).
- Model update policy: release notes, regression testing obligations, rollback mechanisms.
- Red-team participation and vulnerability disclosure obligations mapped to OWASP LLM risks and, where applicable, ATLAS techniques.
- Performance KPIs expressed in tokens/sec, latency at defined context lengths, and maximum concurrency under SLA.

Use cases tailored to Pakistan's current maturity

Each use case below is designed to work under limited compute and uneven connectivity, and to align with the earlier manual's emphasis on RAG and controlled in-house agents.

Use case one: Policy and PC-1 drafting assistant with enforced citations

- **Business objective:** compress drafting cycle time while maintaining traceability to approved sources.
- **Architecture:** RAG over approved laws, policy documents, notifications, and previous finalized summaries; generation model tuned only for tone/format.
- **Pakistan constraint control:** documents may be fragmented across offices; implement provincial-level knowledge bundles with versioned updates.
- **Risk posture:** hallucination risk contained via "citation required" output spec and rejection if no evidence found.

Use case two: Land records and mutation workflow co-pilot

- **Business objective:** reduce error rates in forms and accelerate citizen-facing explanations.
- **Architecture:** split system: (a) deterministic rules and database queries for authoritative fields; (b) LLM for summarization/explanation; (c) RAG for procedural references.
- **Why this works:** retrieval aligns outputs to departmental SOPs; the authoritative DB prevents "LLM invents facts."
- **Pakistan constraint control:** support partial digitization by ingesting scanned SOPs into "approved-source KB" and using hybrid retrieval to compensate for OCR noise.

Use case three: Citizen grievance triage at scale using SLM + escalation

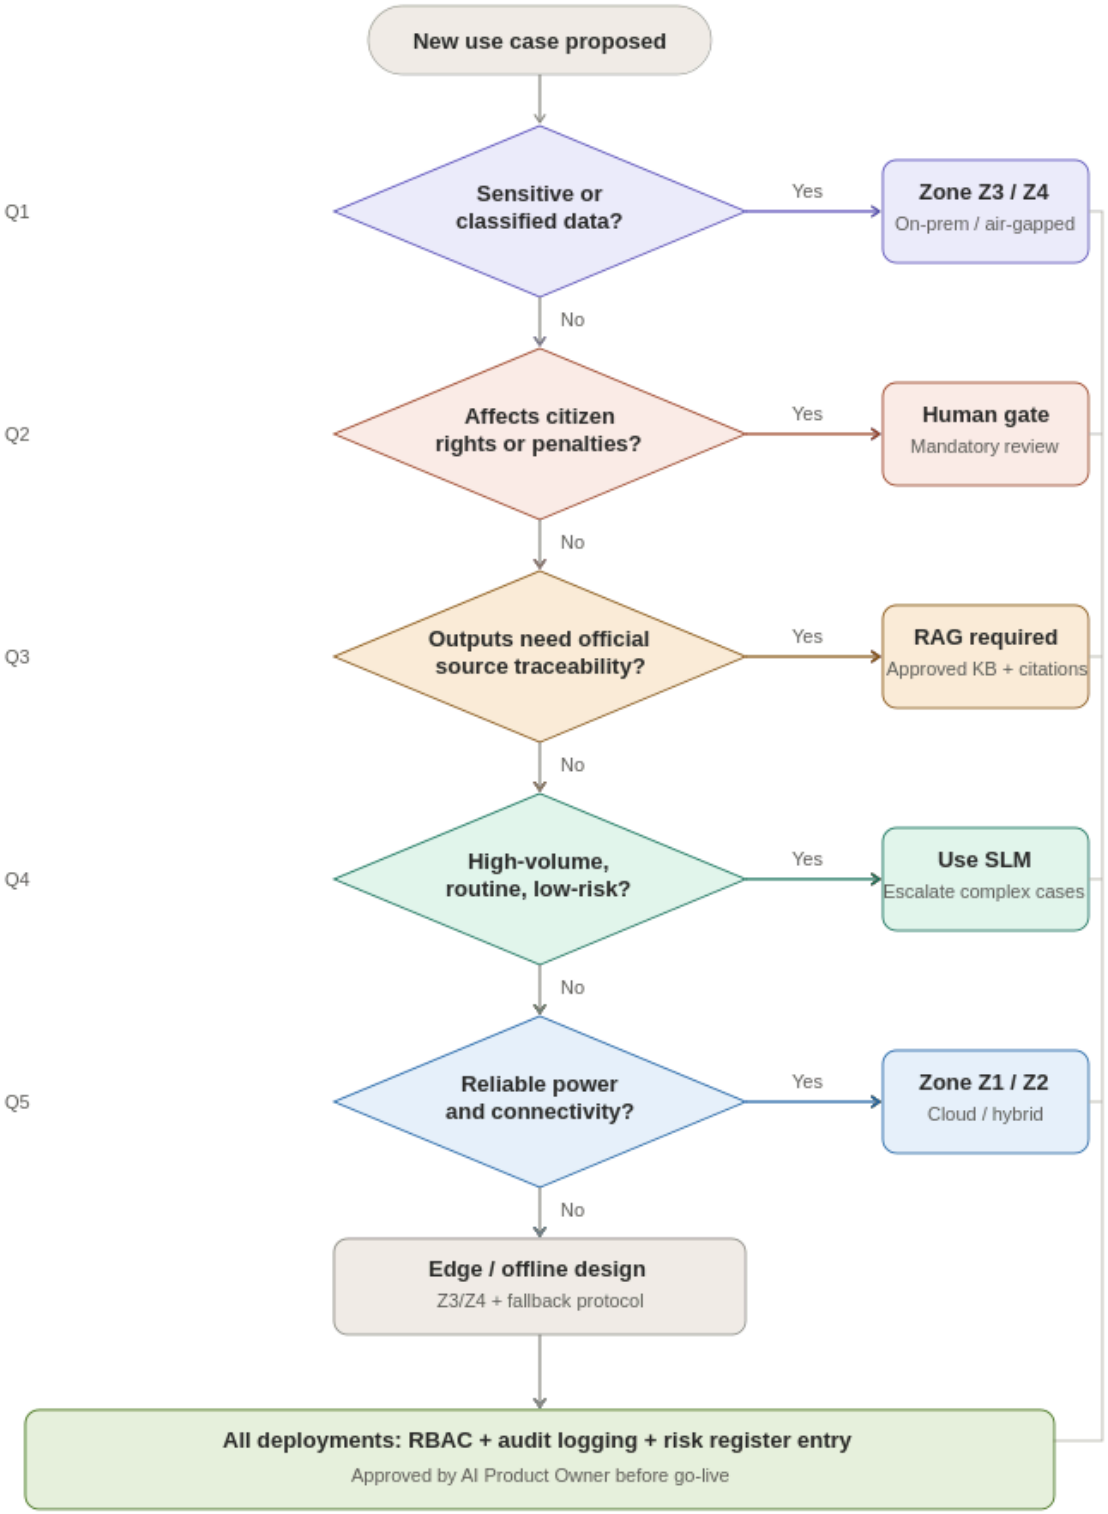
- **Business objective:** route complaints and produce standardized summaries for officers.
- **Architecture:** SLM classifies/triages and drafts structured summaries; only complex cases are escalated to larger models. This matches the "SLM for cost-effective workloads" principle.
- **Controls:** strict redaction rules; token budgets to prevent abuse; audit logs for every classification decision.

Use case four: Safe City and incident narrative synthesis

- **Business objective:** convert operational feeds into concise situation reports for decision-makers.
- **Architecture:** multimodal pipeline (speech-to-text / video analytics upstream; LLM only for summarization and briefing formats). The prior manual already positions Safe City/computer vision as a public-sector AI class; AI 301 makes it "governable" by isolating model roles.
- **Hard control:** any enforcement or punitive decision must remain human; AI outputs remain briefing artifacts.

Decision Flowchart for Non-Technical Readers (Lane A Guide)

Civil servants and product owners in Lane A do not need to understand transformer internals to make sound AI deployment decisions. Use the flowchart below to route any proposed use case before commissioning technical work. Follow the five questions in order and apply the corresponding control from the right-hand column. All deployments, regardless of routing , must complete the risk register entry and receive AI Product Owner sign-off before go-live.



Decision Flowchart for Non-Technical Readers (Lane A Guide)

Quick-reference routing summary

Sensitive data + rights-affecting outputs = Z3/Z4 + mandatory human gate + RAG with approved sources. High-volume routine tasks = SLM first. Poor connectivity = edge/offline design + fallback protocol. All deployments require: role-based access, audit logging, and a risk register entry approved by the AI Product Owner.

Governance Escalation Ladder

The manual defines four institutional roles — AI Product Owner, Data Steward, Security Officer/CISO delegate, and ML/Platform Lead — but a production programme will encounter disputes and deadlocks between these roles. Without a defined escalation path, disagreements stall delivery or are resolved informally in ways that bypass governance controls. This section provides a mandatory escalation ladder for the most common conflict patterns.

RACI matrix for AI deployment decisions

R = Responsible (does the work), A = Accountable (owns the outcome), C = Consulted (input required), I = Informed (notified of decision).

- **Use-case approval and risk register entry:** AI Product Owner (A/R), Data Steward (C), Security Officer (C), ML Lead (C), Ministry/IT Board (I).
- **Data classification and approved-source gate:** Data Steward (A/R), AI Product Owner (C), Security Officer (C), ML Lead (I), Ministry/IT Board (I).
- **Deployment boundary and threat model sign-off:** Security Officer (A/R), AI Product Owner (C), Data Steward (C), ML Lead (R), Ministry/IT Board (I).
- **Model/quantization strategy and rollback plan:** ML/Platform Lead (A/R), Security Officer (C), AI Product Owner (I), Data Steward (I).
- **Production go-live authorisation:** AI Product Owner (A), Security Officer (R), Data Steward (R), ML Lead (R), Ministry/IT Board (C).

Escalation paths by conflict type

Conflict: Data Steward rejects a training dataset that the AI Product Owner wants to proceed with.

Resolution: The Data Steward's rejection is binding at the data-gate stage. The AI Product Owner may request a formal review by presenting a revised dataset proposal. If unresolved, it escalates to the Ministry AI Governance Committee or Secretary/Additional Secretary level. No deployment proceeds while the dispute is open.

Conflict: Security Officer vetoes a deployment configuration recommended by the ML/Platform Lead.

Resolution: The Security Officer's veto is binding on deployment boundary decisions. The ML Lead must propose a mitigated architecture. If teams cannot converge within two iterations, the AI Product Owner escalates to the Ministry CISO for a binding ruling.

Conflict: A senior officer directs the AI Product Owner to deploy without completing the risk register and stage-gate process.

Resolution: The AI Product Owner must document the instruction and outstanding gate items in writing. A condensed fast-track gate review may be convened, but no gate may be skipped entirely. Any deployment that proceeds with unresolved items must be notified to the Ministry IT Board.

Conflict: A model update degrades a production workflow without prior notice.

Resolution: The ML Lead must invoke the rollback plan within the agreed SLA window and notify the AI Product Owner and Security Officer immediately. A mandatory regression test review is required before further updates. Recurring incidents escalate to Ministry IT leadership as a programme governance failure.

Escalation authority reference

Where a Ministry AI Governance Committee does not yet exist, the interim escalation chain is: programme-level disputes to Ministry CISO; ministry-level disputes to Secretary/Additional Secretary; cross-

ministry or policy-level disputes with data protection or national security implications to MoITT and, where appropriate, the National Cybersecurity Authority. All escalation decisions must be documented and retained as part of the programme audit trail.

What Not to Deploy: Prohibited and Restricted Use Cases

A governance-first AI programme is defined as much by what it refuses to deploy as by what it enables. Pakistan's ongoing data protection legislative process, combined with constitutional rights of citizens and accountability obligations of public servants, creates a clear policy basis for explicit exclusions.

Category 1: Absolutely prohibited uses

These use cases may not be deployed under any circumstances within this framework, regardless of technical capability or senior-level instruction:

Automated penalty or fine issuance. No AI system may issue a penalty, fine, suspension, or enforcement notice without a human officer reviewing and formally authorising the decision.

Automated denial of government entitlements or benefits. AI may triage or summarise applications but a final denial of any entitlement must involve a human decision with an explanation communicable to the applicant.

Social scoring or predictive risk profiling of citizens. No system may assign citizens a risk or behavioural score that influences government treatment without explicit legislative authorisation, published criteria, and a right of appeal.

Real-time biometric identification in public spaces without legal basis. Facial recognition or biometric matching in public spaces requires explicit legal authorisation, published retention policies, and oversight mechanisms.

Use of AI outputs as legal evidence without human attestation. AI-generated analyses may not be submitted as evidence in legal or adjudicatory proceedings without review, verification, and attestation by a responsible officer.

Category 2: Restricted uses requiring additional safeguards

These use cases may be deployed only after the additional controls listed have been implemented and reviewed by the Security Officer and AI Product Owner:

Tax and revenue enforcement assistance. Required controls: human review gate on all AI-flagged cases before escalation; explainability requirement; audit trail with document-level citations.

Law enforcement narrative synthesis (Safe City and related). Required controls: outputs explicitly labelled as AI-generated drafts; no output transmitted to judicial or prosecutorial bodies without a human officer's review and signature.

Health or medical record processing. Required controls: Z3/Z4 deployment only; explicit consent or legal basis; anonymisation before AI pipeline ingestion; clinician review before any recommendation is communicated.

Cross-ministry data aggregation for AI processing. Required controls: data-sharing agreement documenting legal basis, data types, and retention period; Data Steward sign-off from each ministry; privacy impact assessment before ingestion.

Relationship to Pakistan's data protection legislative process

The MoITT draft Personal Data Protection Bill (2023) articulates propositions around lawful collection, consent, and the rights of data subjects. While its legal commencement remains in flux, this manual treats its policy direction as binding for government AI deployments. The principle of purpose limitation directly prohibits repurposing citizen data originally collected for service delivery into AI training datasets without explicit legal basis and consent. Programme teams must document the lawful basis for every data source in their RAG knowledge base and fine-tuning corpus, and must review this documentation when the final legislation comes into force.

Appendices and Templates

Prompt templates that survive governance review

These templates are designed for internal systems where the model is constrained to approved sources. They operationalize the earlier manual's insistence that government AI should use internal context and preserve human oversight.

Template A: "Cited policy answer"

- Role: "You are an internal policy assistant."
- Task: "Answer the question using ONLY retrieved sources. If sources do not contain the answer, say: 'Not found in approved sources.'"
- Output: "Answer + citations (doc ID + section)."
- Stop condition: "No speculation; no external knowledge."

Template B: "Bilingual citizen response"

- Pakistan's AI policy explicitly emphasizes dissemination in national and regional languages to ensure accessibility.
- Task: "Provide a citizen-facing explanation in Urdu + English, with the Urdu first, and include a one-line disclaimer: 'Final decision remains with the department.'"
- Constraint: "Use only approved service rules and SOPs."

Token accounting and SLA model (operational practicality)

Use token budgets as a control surface:

- **Token budget per request:** cap prompt + retrieved context + output tokens. This prevents runaway costs and model DoS patterns.
- **Prefill cost driver:** large retrieved contexts increase prefill time and memory pressure.
- **Concurrency risk driver:** KV cache grows with longer contexts and more concurrent sessions; this directly impacts VRAM sizing.

Model selection matrix (governance + license posture)

A government-grade model choice is not just "quality"; it is license, deployability, and sovereignty posture.

Model family Llama (open-weight) by Meta Platforms	Why it is relevant to government Common open-weight baseline; supports fine-tuning	License/terms considerations (must be reviewed by legal/procurement) Llama licensing has specific terms and acceptable-use constraints; review official license text	Best-fit tier Medium–high (on-prem/hybrid)
Phi (SLM) by Microsoft	Small language model positioning; local deployment narrative	Review model distribution terms per deployment channel	Low–medium (edge + triage)
Gemma by Google	Open-weight family used for local and managed deployments	Review Gemma terms of use and access requirements	Low–medium (departmental)
Qwen (open-weight variants)	Broad parameter range with Apache 2.0 coverage for many variants	Technical reports explicitly indicate license for certain lines (e.g., Apache 2.0 in published tables)	Low–high (depending variant)

"AI 301" training structure (delivery-ready)
The earlier manual positions AI 301 as "Government-grade Models, Fine-tuning, & AI Infrastructure."

A deployment-ready training program should produce three tangible outputs per participating ministry/department:

1. **Use-case portfolio + risk tiering:** classify candidate use cases into Z1–Z4 deployment zones.
2. **Reference architecture decision record:** RAG-only vs hybrid vs fine-tuned; with audit and RBAC design.
3. **Pilot implementation pack:** knowledge base ingestion plan, evaluation pack (MEP), security checklist mapped to OWASP LLM Top 10, and go-live runbook.

TECH destiNATION PAKISTAN
MINISTRY OF IT & TELECOM

atomcamp

Ministry of Planning Development & Special Initiatives

Artificial Intelligence 301 for Pakistan Civil Servants program is a joint initiative by the Ministry of Information Technology and Telecommunication, Ministry of Planning, Development and Special Initiatives, Civil Services Academy, and atomcamp, aimed at equipping Civil Services Academy probationers with essential AI awareness for effective governance and policy-making.

For further information, please contact the Civil Services Academy, Walton Lahore

Training Manual

A course on

Hardware and Infrastructure for Sovereign Artificial Intelligence in Pakistan

A Policy Briefing for Civil Servants and Public Policy Makers

401

TECH destiNATION PAKISTAN
MINISTRY OF IT & TELECOM

Ministry of Planning Development & Special Initiatives

atomcamp

Table of Contents

- UNDERSTANDING AI HARDWARE: FROM CHIPS TO DATA CENTRES ---- 174
- THE GPU MARKET: NVIDIA, AMD, HUAWEI AND THE GEOPOLITICS OF CHIPS ---- 178
- TRAINING VERSUS INFERENCE: WHY THE DISTINCTION DEFINES SOVEREIGNTY ---- 181
- DATA CENTRES: ARCHITECTURE, CONSTRUCTION, OPERATION AND GOVERNANCE ---- 183
- DATA CENTRE ECONOMICS: INVESTMENT, COST AND RETURN ON INVESTMENT ---- 186
- SOVEREIGNTY, DATA LOCALISATION AND PAKISTAN'S LEGAL FRAMEWORK ---- 188
- DIGITAL PUBLIC INFRASTRUCTURE AND THE AI LAYER ---- 191
- HOW OTHER COUNTRIES ARE APPROACHING SOVEREIGN AI INFRASTRUCTURE ---- 193
- PAKISTAN'S CURRENT INFRASTRUCTURE LANDSCAPE ---- 195
- RECOMMENDATIONS AND A PHASED NATIONAL STRATEGY ---- 197
- 11. GLOSSARY OF TECHNICAL TERMS ---- 201

Section 1: Understanding AI Hardware — From Chips to Data Centers

1.1 What is a Processor and Why AI Needs a Different Kind

Every computer performs calculations. The component that performs those calculations is called a processor. For decades, personal computers and government systems relied on a Central Processing Unit (CPU). A CPU is an extraordinarily powerful general-purpose processor, it can run spreadsheets, connect to the internet, process a database query, or render a webpage. A modern CPU contains between 8 and 128 cores, each capable of performing complex sequential tasks very rapidly.

Artificial Intelligence, however, is fundamentally different from these tasks. Training and running an AI model requires performing billions, sometimes trillions, of simple mathematical operations simultaneously. A CPU's strength is sequential complexity; AI's demand is parallel simplicity. This mismatch is why the Graphics Processing Unit became the engine of the AI revolution.

The Professor and the Classroom: A Policy Analogy

Imagine a CPU as one brilliant professor who can solve any problem, one at a time, very quickly. A GPU is a classroom of 10,000 primary school students, each capable of adding two numbers. If you need to add 10,000 pairs of numbers simultaneously, the classroom wins by an enormous margin. AI is essentially adding 10,000 pairs of numbers, repeatedly, at enormous scale. That is why GPUs run AI.

1.2 What is a GPU and How Does It Work

A GPU; Graphics Processing Unit was originally designed to render video game graphics, which require millions of pixels to be calculated simultaneously. Researchers discovered in the late 2000s that the same architecture that renders graphics is perfectly suited to the matrix multiplication at the heart of AI.

A modern AI GPU contains tens of thousands of small processing cores working in parallel. The NVIDIA H100, currently the world's most widely used AI chip, contains 16,896 cores.

Key Technical Specifications Explained

Term	What It Measures	Plain English Meaning
TFLOPS / PFLOPS	Computational speed	Trillion / quadrillion math operations per second. Higher = faster AI. H100 = 2,000 TFLOPS. Your laptop GPU = ~20 TFLOPS.
VRAM (GB)	GPU memory	How large an AI model can be loaded onto the GPU. The model must fit entirely in VRAM to operate efficiently. H100 = 80 GB. Llama 3 70B model needs ~140 GB (two H100s).
Memory Bandwidth (TB/s)	Data transfer speed within GPU	How fast the GPU reads data from its own memory. H100 = 3.35 TB/s. Critical for inference speed — models generate responses faster with higher bandwidth.
NVLink / InfiniBand	GPU-to-GPU communication	The high-speed cables connecting multiple GPUs into one cluster. Like motorways between GPU cities. NVLink = 900 GB/s within a server. InfiniBand = 400 Gb/s between servers.

Key Technical Specifications Explained

Term	What It Measures	Plain English Meaning
MW (Megawatts)	Electricity consumption	Why Retrieval provides non-parametric memory and provenance
Tokens / Second	AI output speed	Adjusts style and response patterns
PUE	Data centre energy efficiency	Use RAG for facts + tuned behavior for workflow
Rack Density (kW)	Power per server cabinet	Smaller models + constrained retrieval

The Hardware Stack: From Chip to National AI Factory

Understanding how individual chips aggregate into national infrastructure is essential for procurement and policy decisions. There are five levels in the hardware stack:

- **Level 1 — GPU Chip:** The individual processor. Example: NVIDIA H100. Cost: USD 25,000–40,000 per chip. Produces 2,000 TFLOPS and has 80 GB of VRAM.
- **Level 2 — GPU Server (Node):** A single server containing 8 GPU chips linked together via NVLink, giving 640 GB of combined VRAM. Power draw: approximately 10.2 kW. Cost: approximately USD 300,000 per server.
- **Level 3 — Server Rack:** A physical cabinet holding 10–12 GPU servers stacked vertically. One AI rack draws 80–120 kW of power and requires liquid cooling. Cost: USD 3–4 million per rack.
- **Level 4 — GPU Cluster (AI Pod):** Hundreds to thousands of servers interconnected via InfiniBand networking. Power draw: 10–50 MW. This is the minimum unit for serious national AI workloads. Cost: USD 100–500 million.
- **Level 5 — AI Factory:** A purpose-built facility housing one or more GPU clusters, along with dedicated power infrastructure, liquid cooling systems, sovereign networking, and operational teams. Power: 50–500 MW. This is national AI infrastructure.

What Pakistan Needs for Sovereign DPI-AI

Pakistan does not need a hyperscale AI factory to begin. A sovereign AI compute node of 200–500 GPU servers (Levels 2–4), drawing 5–10 MW of power and hosted at the National Data Centre or NASTP in Islamabad, would be sufficient to run all government AI applications on the Digital Public Infrastructure — citizen chatbots, fraud detection, document processing, and policy simulation. Estimated investment: USD 60–150 million. This is the minimum viable sovereign AI infrastructure.

Section 2: The GPU Market — NVIDIA, AMD, Huawei and the Geopolitics of Chips

NVIDIA — The Dominant Player

NVIDIA currently holds approximately 80% of the global AI GPU market. Its dominance rests on two foundations: hardware performance and software ecosystem. The hardware story is well known — NVIDIA's H100 and its successor the B200 (Blackwell architecture, released 2025) are the most capable AI processors available. A single H100 server costs more than a luxury apartment in Islamabad; demand globally exceeds supply.

The deeper and more durable source of NVIDIA's power is CUDA — Compute Unified Device Architecture. CUDA is the software language that researchers and engineers use to programme GPU workloads. NVIDIA released CUDA in 2006 and has invested in it for nearly two decades. Every major AI framework — TensorFlow, PyTorch, JAX — is built to run on CUDA. Virtually every AI researcher and engineer on the planet has been trained on CUDA. Switching away from NVIDIA means rewriting software, retraining engineers, and accepting performance penalties. This is not merely a market position — it is a structural dependency that affects national AI strategies.

Chip	Generation	VRAM	Performance	Cost (approx.)
A100	2020 (Ampere)	80 GB	312 TFLOPS	USD 10,000–15,000
H100	2022 (Hopper)	80 GB HBM3	2,000 TFLOPS	USD 25,000–40,000
H200	2024 (Hopper+)	141 GB HBM3e	~3,000 TFLOPS	USD 35,000–50,000
B200 (Blackwell)	2025 (Blackwell)	192 GB HBM3e	~9,000 TFLOPS	USD 60,000–80,000

Pricing Disclaimer: All GPU prices listed above are indicative estimates and are subject to change due to evolving market conditions, US export control policy, supply chain constraints, geopolitical factors, and competition from new chip generations. Readers should verify current pricing through the following authoritative sources before making any procurement or budgetary decisions: (1) NVIDIA Official Pricing & Partners — www.nvidia.com/en-us/data-center; (2) Dell Technologies GPU Server Pricing — www.dell.com/en-us/work/shop/servers-storage-and-networking; (3) Lambda Labs GPU Market Tracker — lambdalabs.com/service/gpu-cloud; (4) SemiAnalysis Chip Market Reports — www.semianalysis.com; (5) Artificial Analysis GPU Benchmarks & Pricing — artificialanalysis.ai; (6) Epoch AI Compute Cost Trends — epochai.org.

AMD — The Emerging Challenger

Advanced Micro Devices (AMD) holds approximately 15% of the AI GPU market and is growing rapidly. AMD's MI300X chip, released in 2023, carries 192 GB of HBM3 memory — more than double the H100's 80 GB. This makes AMD chips particularly attractive for running very large AI models that would require linking multiple H100s. Microsoft and Meta have both made large AMD GPU purchases.

AMD's software ecosystem — called ROCm — is maturing but lags behind CUDA by several years. Most AI code runs on AMD chips without modification today, though performance optimisation requires additional engineering effort. AMD represents a credible alternative that Pakistan should include in all procurement evaluations.

Huawei Ascend — The Sovereign AI Option for Pakistan

Huawei's Ascend series — specifically the Ascend 910B and the newer Ascend 910C — represents China's strategic answer to NVIDIA's dominance. The Ascend 910B delivers approximately 60–70% of H100 performance by most benchmarks, while the Ascend 910C narrows that gap further. More significant than individual chip performance is the CloudMatrix 384 system — a Huawei-built cluster linking 384 Ascend chips into a single AI supercomputing node. At the cluster level, Huawei argues credibly that CloudMatrix can match NVIDIA DGX systems.

Why Huawei is Strategically Relevant for Pakistan

United States export controls introduced in 2022 and expanded in 2023 restrict sales of advanced AI chips (NVIDIA H100, AMD MI300X) to certain countries. While Pakistan is not currently on the restricted list, procurement pathways are complex and subject to change. Huawei is already embedded in Pakistan's digital infrastructure — the Higher Education Commission data centre, the National Bank of Pakistan data centre, PTCL backbone networks, and multiple CPEC connectivity projects all use Huawei equipment. A government-to-government negotiation for Ascend-based sovereign AI infrastructure through CPEC channels is politically viable, financially accessible, and technically credible. Pakistan should maintain relationships with both NVIDIA and Huawei as strategic options.

Other Players

Several other companies are developing AI accelerators that policy makers should track:

- **Google TPUs (Tensor Processing Units):** Custom chips built by Google for its own AI workloads, not sold commercially but used in Google Cloud. Relevant if Pakistan ever uses Google Cloud services.
- **Intel Gaudi:** Intel's AI accelerator targeting the mid-market. Competitive pricing but smaller software ecosystem. Relevant for budget-conscious deployments.
- **Graphcore IPU:** A British company building a different kind of AI chip (Intelligence Processing Unit) optimised for graph neural networks. Niche but notable for certain AI research applications.
- **Cerebras:** An American company that built a wafer-scale chip the size of an entire silicon wafer — the world's largest chip. Extremely fast for certain AI workloads. Commercially available through cloud partners.
- **SambaNova:** Another US AI chip company targeting enterprise deployments with high memory capacity and competitive inference speed.

The key policy insight: NVIDIA's market dominance is real but not permanent. The AI chip market is moving faster than any other semiconductor segment. Pakistan's procurement strategy should avoid exclusive dependence on any single vendor.

Section 3: Training versus Inference — Why the Distinction Defines Sovereignty

3.1 What is Training

Training an AI model is the process by which the model learns from data. During training, the model is exposed to vast quantities of text, images, or other data. With each example, it adjusts millions — or billions — of internal numerical parameters to become better at predicting patterns. This process is computationally enormous and happens only once (or occasionally in update cycles called fine-tuning).

Training GPT-4 is estimated to have consumed approximately 25,000 NVIDIA A100 GPUs running continuously for three months, at a cost of USD 50–100 million. Training a large language model is one of the most computationally intensive tasks that humanity currently undertakes. It requires a massive GPU cluster, enormous amounts of carefully curated data, and a highly skilled engineering team.

Pakistan's Position on Training

Pakistan should not attempt to train a large foundational AI model from scratch in the near term. The cost is prohibitive, the data science talent required is scarce, and the strategic value is marginal when open-source foundational models (Llama, Mistral, DeepSeek) are freely available. Pakistan's training priority should be fine-tuning: taking an existing open-source model and training it further on Pakistani data — Urdu text, legal documents, agricultural records, NADRA-derived demographic patterns. Fine-tuning a model of this type requires a fraction of the compute: perhaps 100–500 GPUs running for days or weeks, not thousands for months.

What is Inference

Inference is what happens every time a user or system asks an AI model a question and receives a response. When a citizen interacts with a government AI chatbot, or when a fraud detection system analyses a Raast transaction, or when an AI processes a NADRA document — all of these are inference. Inference is the AI system performing its learned function in the real world.

Inference requires far less compute than training. A single NVIDIA H100 server (8 GPUs) can simultaneously serve hundreds of users of a Llama-class model. For a government AI deployment serving millions of Pakistani citizens, a cluster of 100–200 GPU servers would be sufficient for peak-hour inference demand. This is achievable infrastructure.

3.3 Why This Distinction is Central to Data Sovereignty

Training and inference have fundamentally different sovereignty implications. Training requires bringing data to the compute. Inference requires bringing compute to where data will be generated. This asymmetry defines the sovereignty calculus.

Dimension	Training	Inference
Frequency	Once (or occasional fine-tuning cycles)	Every user interaction — millions of times per day
Compute requirement	Enormous — thousands of GPUs, months of time	Moderate — hundreds of GPUs, milliseconds per query
Data sovereignty risk	HIGH if done on foreign cloud — training data leaves Pakistan	LOW if inference runs on Pakistan-located GPUs — citizen queries never leave the country
Pakistan strategy	Fine-tune open-source models on sovereign compute; use foreign cloud only for non-sensitive R&D	All production inference must run on sovereign infrastructure — NDC, NASTP, or approved sovereign cloud

3.4 The DeepSeek Moment and Its Implications for Pakistan

In early 2025, a Chinese AI company named DeepSeek released models (DeepSeek R1 and V3) that matched the performance of GPT-4 class models at a claimed training cost of approximately USD 6 million, roughly 1–5% of what OpenAI spent training comparable models. The release caused a significant reassessment of the assumption that frontier AI required billion-dollar compute investments. For Pakistan, the DeepSeek development has three direct implications. First, efficient open-source models that run on modest GPU clusters are now available, Pakistan does not need H100 supercomputers for government AI applications. Second, fine-tuning a DeepSeek or Llama 3 model on Urdu and Pakistani government data is technically feasible on Pakistan's existing infrastructure augmented with a targeted GPU investment. Third, the era of AI being exclusively accessible to countries with massive compute resources is ending — the efficiency curve strongly favours Pakistan's constrained but focused approach.

Section 4: Data Centres — Architecture, Construction, Operation and Governance

4.1 What is a Data Centre

A data centre is a purpose-built facility that houses computer servers, storage systems, and networking equipment in a controlled environment. It provides the physical infrastructure — power, cooling, security, and connectivity — that digital services require to operate continuously and reliably. Data centres are the physical foundation of the digital economy; every government service, every bank transaction, every mobile network call is processed in a data centre somewhere.

Data centres range from a small room in a government ministry housing a few servers, to hyperscale facilities covering tens of acres and consuming hundreds of megawatts of power. For AI workloads, the relevant category is the AI-optimised data centre or AI Factory — a facility specifically engineered for high-density GPU deployments.

4.2 How Data Centres are Built

Site Selection

Site selection for a data centre considers five primary factors:

- **Power availability and stability:** Data centres require uninterrupted, high-capacity electrical supply. AI data centres draw 50–500 MW — equivalent to a small city. Pakistan's power reliability challenges make this the most significant constraint.
- **Connectivity:** Multiple high-capacity fibre optic connections to the national internet backbone and international submarine cables. The NDC in Islamabad and PTCL's infrastructure in Karachi provide the best existing connectivity.
- **Physical security:** Distance from flood zones, seismic fault lines, and military targets. Climate — cooler locations reduce cooling costs significantly; Islamabad's climate is more favourable than Karachi for this reason.
- **Land and construction cost:** Hyperscale data centres require large, flat parcels. NASTP in Islamabad and various industrial zones offer suitable sites.
- **Workforce:** Proximity to technical staff for operations and maintenance. Islamabad-Rawalpindi and Lahore have the deepest IT talent pools.

Physical Construction

An AI data centre consists of several integrated systems:

- **The white space:** The actual server floor where racks of GPUs are installed. Rows of racks with precise hot-aisle/cold-aisle airflow management, or increasingly, liquid cooling manifolds running directly to each server.
- **Power infrastructure:** High-voltage transformers stepping down utility power, redundant UPS (Uninterruptible Power Supply) systems providing seconds-to-minutes of battery backup during outages, and diesel generators providing hours of backup. AI facilities require multiple redundant power feeds.
- **Cooling systems:** Traditional data centres use computer room air conditioners (CRACs). AI data centres increasingly use direct liquid cooling (DLC) — chilled water circulating through pipes connected directly to GPU servers. A 100 MW AI facility may consume 20–30 MW just on cooling.
- Network infrastructure: InfiniBand switches and cables connecting GPU clusters internally, top-of-rack switches, core routers, and external fibre connections to carrier networks.
- Security infrastructure: Multi-factor physical access control, CCTV, 24/7 security personnel, biometric entry, mantrap corridors, and for sensitive government facilities — electromagnetic shielding (TEMPEST standards).

4.3 Tier Classifications

The Uptime Institute's tier classification system is the global standard for measuring data centre reliability:

Tier	Uptime	Annual Downtime	Description
Tier I	99.67%	28.8 hours	Basic infrastructure, no redundancy. Suitable for development/test only.
Tier II	99.74%	22 hours	Redundant capacity components. Partial protection against disruption.
Tier III	99.98%	1.6 hours	Concurrently maintainable. No single point of failure. Required for most
Tier IV	100.00%	26 minutes	Fault-tolerant. Any single component can fail without service interruption.

4.4 AI Factory — A New Category

The term AI Factory, popularised by NVIDIA CEO Jensen Huang, describes a data centre purpose-built to produce AI inference and training as its primary output. Three characteristics distinguish an AI factory from a conventional data centre:

- **Power density:** Conventional data centres design for 5–10 kW per rack. An AI factory designs for 60–120 kW per rack — 10–20 times higher. This requires fundamentally different electrical distribution, cooling infrastructure, and structural floor loading.
- **Networking fabric:** Inside an AI factory, GPU-to-GPU communication requires specialised InfiniBand or NVLink switching at speeds of 400 Gb/s per port. The internal network of an AI factory is more complex than most national internet backbones.
- **Operational model:** AI factories require 24/7 operations teams with deep AI engineering expertise — not just IT support. They function as operational infrastructure (like power stations or water treatment plants) rather than IT departments.

4.5 Data Centre Governance

How a data centre is governed — who owns it, who operates it, who can access it, and under what legal framework — is as important as its technical specifications. Four governance models are relevant for Pakistan:

- **Government-owned, government-operated (GOGO):** The sovereign maximum. Pakistan's National Data Centre (MoIT) operates on this model. Maximum sovereignty, but government bears full capital and operational cost, and may lack the operational expertise of specialist operators.
- **Government-owned, privately operated (GOPO):** Government builds and owns the facility; a private operator (domestic or foreign, under contract) manages daily operations. PTCL's data centres approximate this model. Balances sovereignty with operational expertise.
- **Private colocation with data residency requirements:** Government agencies rent rack space in a certified private data centre operating under Pakistani law, with data residency requirements enforced by contract and regulation. Jazz and PTCL commercial data centres serve this role today.

Sovereign cloud partition: A foreign cloud provider (AWS, Azure, Huawei Cloud) operates infrastructure on Pakistani soil under a bilateral data agreement, with cryptographic guarantees that Pakistan holds encryption keys and the provider cannot access data. Most viable for immediate AI scale-up; requires strong legal framework and technical verification mechanisms.

Section 5: Data Centre Economics — Investment, Cost and Return on Investment

5.1 Capital Investment (CapEx) for AI Infrastructure

Building AI infrastructure requires substantial upfront capital. The following table provides indicative cost ranges for different scales of sovereign AI compute for Pakistan's government:

Infrastructure Scale	GPU Count	Power Draw	Est. CapEx (USD)	Primary Use
Pilot / Proof of Concept	32–64 GPUs	0.5–1 MW	USD 5–15M	R&D, testing, small services
Minimum Viable Sovereign AI	200–500 GPUs	2–5 MW	USD 60–150M	Govt DPI AI, citizen services
National AI Compute Node	1,000–2,000 GPUs	10–20 MW	USD 300–600M	All federal AI, research, fintech
Pakistan AI Factory	5,000–10,000 GPUs	50–100 MW	USD 1.5–3B	National + commercial +

⚠ Pricing Disclaimer: All capital expenditure (CapEx) figures above are indicative estimates and are subject to change based on global GPU market fluctuations, procurement channel (commercial vs. government-to-government), currency exchange rates, infrastructure build costs, and evolving supply conditions. These figures should not be used as final budgetary commitments without current vendor quotations. For up-to-date infrastructure benchmarks and cost data, consult: (1) Uptime Institute Global Data Centre Survey — uptimeinstitute.com/research; (2) JLL Data Centre Investment Reports — www.jll.com/en/industries/data-centers; (3) Structure Research Market Reports — structureresearch.net; (4) NVIDIA DGX System Pricing — www.nvidia.com/en-us/data-center/dgx-systems; (5) Huawei Cloud Pakistan — www.huaweicloud.com/intl/en-us; (6) Dell PowerEdge AI Server Pricing — www.dell.com/en-us/work/shop/servers-storage-and-networking.

5.2 Operational Expenditure (OpEx)

Data centres are not one-time investments. Ongoing operational costs typically represent 20–40% of total cost of ownership (TCO) annually:

- **Electricity:** The single largest operational cost. A 10 MW data centre at PKR 30/kWh consumes approximately PKR 2.6 billion per year in electricity. Power cost negotiations with WAPDA/NEPRA are critical for feasibility.
- **Staffing:** A Tier III data centre with 100 racks requires approximately 20–30 operational staff (data centre technicians, network engineers, security personnel, facility management). Specialised AI operations adds 10–20 additional ML engineers.
- Hardware refresh: GPU generations turn over every 2–3 years. Older GPUs lose competitive efficiency. A lifecycle replacement budget of 20–25% of hardware value per year is standard.
- Connectivity: Bandwidth costs, peering arrangements, and redundant carrier contracts.
- Maintenance contracts: Hardware warranty and support contracts, typically 10–15% of hardware value per year.

⚠ **Pricing Disclaimer:** All operational expenditure estimates above are subject to change due to Pakistan's electricity tariff revisions, hardware refresh cycles, staffing market conditions, and maintenance contract terms. In particular, electricity costs should be verified against current NEPRA-notified tariffs before use in feasibility calculations. For the most current operational cost benchmarks, consult: (1) NEPRA Official Tariff Determinations — www.nepra.org.pk/tariff; (2) WAPDA Bulk Power Rates — www.wapda.gov.pk; (3) GlobalPetrolPrices Pakistan Electricity Tracker — www.globalpetrolprices.com/Pakistan/electricity_prices; (4) Uptime Institute PUE & OpEx Benchmarks — uptimeinstitute.com/research; (5) Pakistan Ministry of Energy (Power Division) — power.gov.pk.

5.3 Return on Investment — How Governments Calculate the Value

Governments must apply a different ROI framework than commercial enterprises. The value of sovereign AI infrastructure is not primarily revenue but public good — efficiency savings, service delivery improvements, GDP multiplier effects, and strategic autonomy:

- **Fraud reduction:** AI-powered fraud detection on Raast and BISP transfers. Pakistan loses an estimated PKR 50–100 billion annually to benefit payment leakages. AI verification systems have demonstrated 40–60% leakage reduction in comparable deployments.
- **Public service delivery cost reduction:** Estonia's X-Road saves citizens 2.8 million hours per year and government agencies significant processing costs. AI-augmented DPI would compound these savings substantially.
- **GDP multiplier effect:** The IMF estimates that every USD 1 invested in digital infrastructure generates USD 2.50–4.00 in economic activity through private sector development, fintech innovation, and productivity gains. For USD 150M sovereign AI infrastructure, the economic return over 10 years may reach USD 375M–600M.
- **Strategic autonomy value:** This is not quantifiable in rupees but is arguably the most important return. A government that cannot run AI on its own infrastructure is permanently dependent on foreign technology companies for its most sensitive public functions. This is a national security cost that cannot be expressed in ROI calculations but must weigh heavily in policy decisions.

India's USD 1.25 Billion AI Mission — The Comparable Case

In March 2024, the Indian government approved the IndiaAI Mission with a budget of USD 1.25 billion, focused on building 10,000+ GPU sovereign compute, developing indigeneous AI models, and creating an AI application layer for government services. India framed this not as a technology investment but as foundational national infrastructure — comparable to building highways or power plants. Pakistan's circumstances differ, but the framing is instructive: sovereign AI infrastructure is infrastructure, not an IT purchase.

Section 6: Sovereignty, Data Localisation and Pakistan's Legal Framework

6.1 What Data Sovereignty Means in Practice

Data sovereignty refers to a government's right and capacity to control data about its citizens, its institutions, and its territory.

In the AI context, sovereignty has three distinct dimensions that policy makers must keep separate:

- **Data localisation:** Where is the data physically stored? Data stored on servers in Pakistan is subject to Pakistani law and cannot be accessed by foreign governments through legal process against foreign companies.
- **Processing sovereignty:** Where is the data processed (computed)? This is the inference question. Even if data is stored in Pakistan, if it is sent to a foreign AI system for processing, it temporarily leaves sovereign jurisdiction.
- **Model sovereignty:** Does Pakistan own or control the AI model? A model trained on Pakistani data but owned by a US company may be withdrawn, modified, or rendered unavailable at any time by that company or by US government action.

6.2 The Critical Insight — Hardware Location Determines Data Location

The most important technical concept for policy makers to internalise is this: when an AI model runs on a GPU server located in Pakistan, citizen data never leaves Pakistan. The internet connection used to send a query from a government system to a local GPU cluster is a domestic network connection. The GPU processes the data and returns the result within Pakistani sovereign territory. The nationality of the GPU manufacturer is irrelevant to data sovereignty — what matters is where the processing hardware is physically located and who controls the network it operates on.

The Critical Distinction: Hardware Origin vs. Data Location

A Huawei GPU chip in a data centre in Islamabad is sovereign Pakistani infrastructure. The same Huawei chip in a data centre in Beijing is not. An NVIDIA chip in a data centre in Islamabad controlled by the Pakistani government is sovereign. The same NVIDIA chip in an AWS data centre in Virginia, processing Pakistani citizens' queries, is not. Hardware origin and data sovereignty are entirely separate questions. Pakistan's policy concern should focus on where computation happens and who controls that infrastructure — not on which country manufactured the chip.

6.3 Pakistan's Existing Legal Framework National AI Policy 2022

Pakistan's National AI Policy, published by the Ministry of Information Technology and Telecommunication in 2022, establishes five strategic pillars for AI development: AI for Economic Transformation; AI for Sustainable Development; AI for Social Inclusion; AI Ethics and Governance; and AI Infrastructure and Skills. The policy explicitly calls for 'developing indigenous AI capabilities and infrastructure' and 'ensuring data sovereignty through domestic compute and storage.' However, the policy lacks specific infrastructure targets, procurement frameworks, or enforcement mechanisms — gaps that this document seeks to address.

Personal Data Protection Bill (PDPB)

Pakistan's Personal Data Protection Bill, in various drafts since 2020, proposes to establish a National Commission for Personal Data Protection and impose data localisation requirements for sensitive data categories. The bill requires 'critical personal data' — including biometric data, health data, and financial data — to be stored and processed exclusively on servers located in Pakistan. If enacted in its current form, the PDPB would legally require all government AI systems processing citizen data to use Pakistan-sovereign infrastructure.

Prevention of Electronic Crimes Act (PECA) 2016

PECA provides the current legal basis for cybersecurity enforcement. Its provisions around unauthorised data access and data interception are relevant to cloud-based AI deployments where government data is processed by foreign companies. Any deployment where a foreign company's personnel could access Pakistani government data in the course of AI model training or fine-tuning would raise PECA compliance questions.

Critical Information Infrastructure (CII) Designations

Pakistan's National Cyber Security Policy 2021 designates certain sectors — financial services, energy, telecommunications, healthcare, and government — as Critical Information Infrastructure. CII systems are subject to heightened security requirements. Any AI infrastructure supporting DPI services should be designed and operated to CII standards, which implies domestic hosting, security audits, and incident reporting obligations.

6.4 Gaps in the Legal Framework

Despite the above, Pakistan's legal framework for AI infrastructure sovereignty has significant gaps that public policy makers must address:

- **No explicit AI data localisation regulation:** Unlike India (IT Rules 2021) or the EU (GDPR, AI Act), Pakistan has no enacted law specifically requiring AI systems to process citizen data domestically. The PDPB remains unenacted.
- **No sovereign cloud certification framework:** There is no government framework specifying what technical and legal standards a data centre must meet to be certified for government AI workloads. Pakistan needs a G-Cloud equivalent.
- **No AI procurement standards:** PPRA rules do not address AI system procurement specifically. Civil servants have no guidance on evaluating AI vendors' data handling, model ownership provisions, or sovereignty compliance.
- **No incident response mandate for AI failures:** If a government AI system produces harmful outputs or suffers a data breach, there is no legal framework specifying reporting obligations, remediation requirements, or liability allocation.

Section 7: Digital Public Infrastructure and the AI Layer

7.1 Recap — What is Digital Public Infrastructure

Digital Public Infrastructure (DPI) refers to foundational digital systems built in the public interest that enable government services, financial transactions, and data exchange between institutions. The internationally recognised three pillars of DPI are: Digital Identity (who you are, provably); Digital Payments (moving money instantly and cheaply); and Data Exchange (sharing data between government agencies and private institutions with citizen consent).

Pakistan's DPI foundation already exists. NADRA's biometric identity system — covering over 240 million citizens — is one of the most sophisticated in the world. Raast, the State Bank's instant payment rail, is operational. The National Data Centre provides shared government cloud infrastructure. What Pakistan's DPI lacks is the AI layer — the intelligence that transforms these digital systems from passive infrastructure into active, intelligent public services.

7.2 How AI Enhances Each DPI Pillar

DPI Pillar	Current Capability (DPI without AI)	AI-Enhanced Capability (DPI + AI Layer)
Digital Identity (NADRA)	Verifies identity via biometric check. Binary: match or no match.	Liveness detection (prevents spoofing), continuous authentication, anomaly detection for identity fraud, document forgery detection, CNIC renewal prediction, demographic analytics for service planning.
Digital Payments (Raast)	Processes instant P2P and P2M transactions. Records transaction data.	Real-time fraud detection on every transaction, anti-money-laundering pattern recognition, transaction categorisation for tax compliance, BNPL credit scoring for the unbanked, merchant risk profiling.
Data Exchange (Inter-Agency)	Structured data sharing between agencies via APIs. Human-readable records.	Automated document processing (OCR + NLP), intelligent form filling, cross-agency anomaly detection, predictive policy analytics, natural language querying of government databases.
Citizen Services Layer	Portals and apps providing access to services. Requires citizen navigation.	Urdu-language conversational AI (chatbots and voice agents), proactive service delivery (AI identifies citizens eligible for benefits and notifies them), grievance analysis and routing, document automation.

7.3 Architecture: How AI Infrastructure Connects to DPI

The AI layer sits between DPI's data assets and its citizen-facing services. Architecturally, it consists of three components:

- AI inference endpoints: API-accessible AI models running on sovereign GPU infrastructure that DPI systems can query. A NADRA system asks the AI to detect a suspicious pattern; the AI responds within milliseconds; the query and response both stay within Pakistan's network.

Model registry: A government-managed library of fine-tuned AI models specific to Pakistani government functions — tax compliance models, agricultural advisory models, healthcare triage models, Urdu language models. Each model is versioned, audited, and owned by the Pakistani government.

Section 8: How Other Countries are Approaching Sovereign AI Infrastructure

8.1 India — The Billion-Person Stack

India's approach to sovereign AI is the most directly relevant comparator for Pakistan. In March 2024, the Indian government approved the IndiaAI Mission with a budget of INR 10,371 crore (approximately USD 1.25 billion). The cornerstone is a national AI compute facility with a minimum of 10,000 GPUs — creating the hardware foundation for indigeneous AI development. Simultaneously, India's BharatGen initiative is developing foundation AI models trained on Indian languages and data, with the express intent of ensuring that India's AI layer reflects Indian knowledge, values, and contexts.

India's Infrastructure model relies on a partnership between the government and private data centre operators. GPU clusters are procured by the government but hosted in certified private facilities under strict data sovereignty agreements. AWS, Microsoft, and Google have all established dedicated India cloud regions in response to data localisation requirements. The UPI payment system — now processing over 10 billion transactions monthly — is increasingly augmented with AI for fraud detection and financial inclusion analytics, all running on domestic infrastructure.

Key lesson for Pakistan: India started with identity infrastructure (Aadhaar, 2009), added payments (UPI, 2016), built the data exchange layer (DEPA, 2020), and is now layering AI on top (IndiaAI Mission, 2024). Pakistan has largely replicated the first three stages; the AI layer is the next logical step in the same sequence.

8.2 Estonia — The Governance Model

Estonia's approach is notable not for its scale — the country has 1.4 million people — but for its governance sophistication. Estonia's X-Road data exchange platform connects over 1,700 services across public and private sectors, saving citizens 2.8 million hours per year. Estonia has now extended X-Road with AI capabilities: automated document classification, anomaly detection across government datasets, and an AI-powered public service portal.

Estonia operates a dual-location data sovereignty model. A complete backup of all government systems is maintained in Luxembourg under a 'data embassy' agreement — a novel legal construct allowing a country to host another country's sovereign data under the originating country's laws. This concept is directly relevant to Pakistan: rather than building a second data centre domestically, Pakistan could negotiate data embassy arrangements with friendly nations (possibly Saudi Arabia, UAE, or Turkey under OIC frameworks) for disaster recovery.

8.3 European Union — The AI Factories Initiative

The European Union's response to AI infrastructure dependency is the AI Factories initiative, launched under the EuroHPC Joint Undertaking. The EU has designated thirteen AI Factories — purpose-built AI supercomputing sites distributed across member states — as shared sovereign compute for European research institutions, universities, startups, and government agencies. Each AI Factory provides GPU-optimised compute accessible to European organisations at subsidised rates, reducing dependence on US hyperscaler clouds.

The EU model is directly replicable at the OIC or SCO level. Rather than Pakistan building a sovereign GPU cluster alone, a consortium of OIC member states — Pakistan, Turkey, Malaysia, Indonesia, Saudi Arabia — could jointly fund and operate a shared Islamic Digital Infrastructure facility, with each country receiving a sovereign compute partition proportional to its contribution. Pakistan's geographic position, CPEC fibre connectivity, and NADRA expertise make it a natural candidate to host or anchor such a facility.

8.4 Saudi Arabia and UAE — Sovereign Wealth Approach

The Gulf states are pursuing the most aggressive sovereign AI strategies globally. Saudi Arabia's HUMAIN initiative — backed by the Public Investment Fund with USD 40 billion committed — is building a full-stack sovereign AI ecosystem: data centres, chip procurement, model development, and AI application deployment. The UAE's G42 company has partnered with OpenAI, Microsoft, and NVIDIA to build sovereign AI infrastructure including Stargate UAE, a 1 gigawatt AI compute cluster in Abu Dhabi.

The Gulf model is instructive for Pakistan in one specific respect: sovereign wealth fund capital can finance AI infrastructure investment that government budgets cannot accommodate. Pakistan's SIFC (Special Investment Facilitation Council) framework could be deployed to attract Gulf sovereign wealth investment into Pakistan's AI infrastructure, replicating the model whereby Gulf capital has already financed energy and telecommunications projects.

8.5 Brazil — DPI-AI Integration

Brazil's approach most closely mirrors Pakistan's DPI circumstances. Brazil's Pix instant payment system — the direct equivalent of Raast — now processes over 8 billion transactions monthly and is augmented with AI for fraud detection, all running on Banco Central do Brasil's sovereign infrastructure. Brazil's AI Plan 2024–2028 allocates BRL 23 billion (approximately USD 4.5 billion) to AI development, with 25% specifically designated for infrastructure. A sovereign cloud consolidating all government AI workloads is central to the plan.

Brazil's most relevant innovation is its approach to model development: rather than training proprietary closed models, Brazil is fine-tuning open-source models on Portuguese-language Brazilian data and making these models available as public goods through its digital public infrastructure. Pakistan should adopt the same model — Urdu-fine-tuned open-source models, developed on sovereign compute, available as public AI infrastructure to all government departments and licensed to the private sector.

Section 9: Pakistan's Current Infrastructure Landscape

9.1 Existing Data Centre Assets

Pakistan currently has over 22 commercial data centres concentrated in Karachi (10), Lahore (8), and Islamabad (4), operated by PTCL, Jazz, Cybernet, Multinet, Supernet, and others. The majority are Tier III certified, providing 99.982% uptime guarantees. PTCL operates the largest government-serving data centre footprint. Jazz operates Pakistan's first Tier III certified commercial facility in Islamabad.

The National Data Centre (NDC), operated by the Ministry of Information Technology and Telecommunication, currently provides Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS) to over 100 federal government agencies. The NDC is the natural anchor for Pakistan's sovereign AI infrastructure expansion. Its location in Islamabad, proximity to government ministries, and existing connectivity to PTCL's national backbone make it the primary candidate site for a sovereign GPU cluster.

9.2 NADRA — Pakistan's Underappreciated Sovereign Asset

The National Database and Registration Authority operates what is internationally recognised as one of the most sophisticated biometric identity systems in the world. NADRA's data centres — Tier III facilities — currently process millions of biometric verifications daily. The CNIC database covers over 240 million citizens with face, fingerprint, and iris biometric data. This is Pakistan's most valuable digital asset and the foundation on which AI-enhanced DPI services can be built.

NADRA's existing infrastructure is already running AI workloads — facial recognition for verification and liveness detection. This makes NADRA Pakistan's most operationally advanced sovereign AI deployment, albeit limited in scope and not yet integrated with other government systems. Expanding NADRA's AI capability and connecting it to a sovereign GPU cluster should be the first phase of Pakistan's AI infrastructure strategy.

9.3 Raast and the SBP Infrastructure

The State Bank of Pakistan's Raast instant payment system, launched in 2021, is a functioning payment DPI layer. By late 2024, Raast was processing millions of transactions daily across P2P (person-to-person), P2M (person-to-merchant), and G2P (government-to-person for BISP transfers) channels. SBP's existing data centre infrastructure supports Raast operations, but currently lacks the AI inference capability to implement real-time transaction fraud detection at scale — a gap that represents both a security risk and a missed efficiency opportunity.

9.4 The Infrastructure Gap — What Pakistan Needs

Pakistan's infrastructure gap is not the absence of data centres — it is the absence of:

- **AI-grade compute:** No GPU cluster exists on sovereign government infrastructure. All AI workloads requiring significant compute currently use foreign cloud services or are simply not performed.
- Inter-agency data exchange layer: No X-Road equivalent connecting NADRA, FBR, SBP, SECP, and health authorities. Each agency maintains data in silos, making AI-powered cross-agency analytics impossible without data leaving sovereign custody.
- Sovereign AI model registry: No government-owned repository of fine-tuned AI models specific to Pakistani government functions and in Pakistani languages.
- AI procurement standards: No framework for evaluating AI vendor claims, data handling practices, or sovereignty compliance in government procurement processes.
- AI-capable workforce in government: Insufficient AI/ML engineers embedded within government agencies to commission, evaluate, and govern AI systems. Training and recruitment is the most urgent non-hardware gap.

Section 10: Recommendations and a Phased National Strategy

10.1 Immediate Actions (2025–2026)

The following actions require no new infrastructure and can be initiated within existing budget authorities:

1. **Enact the Personal Data Protection Bill** with explicit AI data localisation provisions requiring all government AI systems processing citizen data to operate on Pakistan-sovereign infrastructure. Include clear definitions of 'sovereign infrastructure' and a certification framework for approved facilities.

2. **Establish a Pakistan Sovereign AI Compute Policy** — a ministerial directive specifying that no federal government system may use foreign cloud AI inference for any workload involving citizen PII, financial data, health data, or national security information. Permitted use cases for foreign cloud AI must be explicitly listed and reviewed annually.
3. **Commission a National AI Infrastructure Audit** — a rigorous assessment of all existing government compute assets (NADRA, NDC, SBP, PITB, federal ministries) to establish the current baseline, identify gaps, and produce a capital requirements estimate for Phase 1 sovereign AI infrastructure.
4. **Develop an AI Procurement Framework for government** under PPRA rules — specifying mandatory vendor disclosures on data handling, model ownership, sovereignty compliance, and incident response obligations for all AI system purchases above PKR 50 million.
5. **Launch a Civil Service AI Literacy Programme** — deploying training modules on AI capabilities, infrastructure requirements, data sovereignty, and ethical AI principles across all Grade 17 and above civil servants within 24 months. This document may serve as source material.

10.2 Phase 1 — Minimum Viable Sovereign AI (2026–2027), USD 60–150M

10.1 Immediate Actions (2025–2026)

Phase 1 establishes Pakistan's sovereign AI compute foundation on existing government infrastructure:

6. S**overeign GPU Cluster at NDC:** Procure and deploy 200–500 AI-grade GPUs (Ascend 910C via CPEC/Huawei, or H100/H200 via commercial channels) at the National Data Centre, Islamabad. Configure as a shared government AI inference and fine-tuning facility. Target: operational within 18 months.
7. **Urdu Foundation Model:** Commission fine-tuning of an open-source model (Llama 3 or DeepSeek) on Urdu-language data, Pakistani legal texts, and government documentation. Host the resulting model in the sovereign model registry. Make available to all federal departments as the 'Pakistan Government AI Model.'

8. **NADRA AI Integration:** Connect NADRA's biometric verification system to the sovereign GPU cluster. Deploy AI-enhanced liveness detection, document fraud detection, and biometric anomaly monitoring. This is the highest-ROI first application.
9. **Raast Fraud Detection:** Deploy AI fraud detection on Raast transaction flows, using SBP infrastructure augmented by GPU compute from the NDC cluster. Target: real-time fraud scoring on all transactions above PKR 10,000.

10.3 Phase 2 — National AI Compute Node (2027–2030), USD 300 600M

10.1 Immediate Actions (2025–2026)

Phase 2 scales sovereign compute to support all federal AI applications and begins enabling provincial and private sector use:

10. E**xpand GPU cluster to 1,000–2,000 units** — sufficient to serve all federal government AI applications simultaneously and provide compute-as-a-service to provincial governments on cost-recovery terms.
11. **Establish the Inter-Agency Data Exchange Layer** — an X-Road equivalent connecting NADRA, FBR, SBP, SECP, NHSRC, and MNFSR (agriculture). Design with privacy-by-default and consent architecture. This unlocks cross-agency AI analytics while preserving data sovereignty.
12. **Launch Pakistan Government AI Portal** — a single interface through which all government departments access AI services (document processing, translation, analytics, chatbots) running on sovereign infrastructure, billed on usage to create sustainable cost recovery.
13. **AI applications across priority sectors** — deploy AI in five priority sectors: healthcare (diagnostic support, supply chain); agriculture (crop advisory, price prediction); education (learning analytics, Urdu NLP); tax and revenue (FBR audit targeting); and justice (case flow management, legal document processing).

10.4 Phase 3 — Pakistan AI Factory (2030+)

Phase 3 transforms Pakistan from an AI infrastructure consumer to a regional provider, generating economic returns on infrastructure investment:

14. **Build a dedicated Pakistan AI Factory** — 5,000–10,000 GPUs, 50–100 MW, Tier IV, hosted at NASTP or a purpose-built facility in Islamabad Techno City. Designed to accommodate future growth. Commercial capacity allocated to Pakistani fintechs, healthtechs, and exporters.
15. **Anchor an OIC Sovereign AI Compute Pool** — leverage Pakistan's infrastructure, NADRA scale, and SCO/OIC diplomatic position to host or co-host a shared sovereign AI compute facility for Muslim-majority and developing nations. Position Pakistan as the AI infrastructure hub for Central and South Asia.
16. **Establish Pakistan as an AI services exporter** — leverage domestic AI capabilities in Urdu NLP, agricultural AI, and Islamic finance AI for export to the 57-nation OIC bloc. Export of AI services has the potential to generate significant foreign exchange earnings by 2030.

Section 11: Glossary of Technical Terms

Term	Definition
AI Factory	A purpose-built data centre facility specifically engineered for high-density GPU clusters, designed to produce AI inference and training as its primary output. Requires specialised power density (60–120 kW/rack), liquid cooling, and high-speed InfiniBand networking.
Bandwidth (Memory)	The speed at which a GPU can read data from its own VRAM memory. Measured in TB/s (terabytes per second). Critical for inference speed. H100 = 3.35 TB/s.
CUDA	Compute Unified Device Architecture. NVIDIA's proprietary software platform for GPU programming. The dominant AI programming environment globally, creating significant market lock-in for NVIDIA hardware.
Data Localisation	The requirement that data about citizens or institutions in a given country must be stored and processed on servers physically located within that country's borders.
DPI (Digital Public Infrastructure)	Foundational digital systems built in the public interest: digital identity, digital payments, and data exchange platforms that enable government services and private sector innovation.
Fine-tuning	The process of further training an existing AI model on a specific dataset to improve its performance for a particular domain or language. Far less compute-intensive than training from scratch.
FLOPS / TFLOPS / PFLOPS	Floating Point Operations Per Second. The standard measure of AI compute speed. Tera = trillion. Peta = quadrillion. H100 = 2,000 TFLOPS.
GPU (Graphics Processing Unit)	A processor with thousands of cores designed for parallel computation. Originally built for video game graphics; now the primary hardware for AI training and inference.
Inference	The process of running a trained AI model to generate outputs in response to real-world inputs. What happens when a citizen interacts with an AI chatbot or a fraud system analyses a transaction.
InfiniBand	A high-speed networking technology used to connect GPU servers within a data centre cluster. Operates at 200–400 Gb/s. The nervous system of an AI factory.
LLM (Large Language Model)	A type of AI model trained on vast text corpora to understand and generate human language. Examples: GPT-4, Llama 3, DeepSeek, Mistral.
MW (Megawatt)	Unit of electrical power. 1 MW = 1 million watts. A data centre's power draw is measured in MW. 1 MW powers approximately 800 Pakistani homes.
NVLink	NVIDIA's proprietary high-speed interconnect for linking multiple GPUs within a single server or between servers. Operates at 900 GB/s between 8 GPUs.
PUE (Power Usage Effectiveness)	A data centre efficiency metric. Total facility power divided by IT equipment power. 1.0 = perfectly efficient. Lower is better.
Rack	A standardised cabinet (typically 42U / 1.8 metres tall) used to mount and organise server hardware in a data centre. AI racks draw 60–120 kW each.
Sovereign AI	A nation's capacity to develop, deploy, and govern AI systems using infrastructure, data, and models under its own legal jurisdiction and control.
Sovereign Cloud	Cloud computing infrastructure operated within a country's borders under that country's laws, ensuring data never leaves national jurisdiction.
Tier III / Tier IV	Data centre reliability classifications. Tier III = 99.982% uptime (1.6 hours downtime/year). Tier IV = 99.995% uptime (26 minutes downtime/year).
Tokens/Second	The rate at which an AI model generates output text. 1 token ≈ 0.75 words. Practical measure of AI response speed for user-facing applications.
Training	The process of teaching an AI model by exposing it to large datasets and adjusting its parameters to improve accuracy. Computationally intensive, done once or periodically.
VRAM (Video RAM)	The dedicated high-speed memory on a GPU chip. AI models must fit entirely within VRAM to operate efficiently. H100 = 80 GB. Insufficient VRAM forces model splitting across multiple GPUs.