




Deloitte.

Sovereign cloud for India: Building trusted digital infrastructure for the AI era

September 2026

Table of contents

4

Executive summary

6

Control requirements

10

Why this matters now

12

**Implementation
enhancements**

14

**AI sovereignty: The
emerging frontier**

16

Closing perspective

19

Connect with us



Executive summary

Cloud computing now underpins a significant share of India's digital economy, public service delivery, regulated-sector operations and emerging AI workloads. According to IDC, India's public cloud services market stood at US\$10.9 billion in 2024 and is projected to reach US\$30.4 billion by 2029, reflecting a 22.6 percent CAGR over 2024-2029.¹ **Therefore, India's focus has shifted from whether to use cloud technology to how to govern, procure and operate it effectively, implementing clear, verifiable and proportionate controls.**

As cloud becomes a foundational layer for citizen services, financial systems, manufacturing and supply chains, Digital Public Infrastructure (DPI) and AI-enabled systems, the nature of risk also changes. **Cloud decisions are no longer purely technology or cost decisions; they now carry implications for legal jurisdiction, operational resilience, continuity of public services, control over sensitive data and long-term strategic autonomy. This creates the need for a sovereign cloud approach that allows India to benefit from scale and innovation while ensuring that critical workloads remain subject to Indian oversight, accountability and control.**

Sovereign cloud, for this paper, means a workload-led governance and control model in which cloud adoption is guided by classification, Indian legal and policy requirements, and verifiable safeguards across data, compute, operations, legal accountability, resilience and enforceable exit rights. For India, these controls

are anchored in different documents such as Digital Personal Data Protection Act, 2023, the applicable DPDP Rules, sectoral regulations,² the **MeitY Cloud Empanelment Programme** and other policies and standards. Sovereignty does not require technological isolation, a single deployment model or the exclusion of global cloud providers. Rather, sovereignty requirements shall be calibrated to the sensitivity, criticality and risk profile of each workload, while respecting existing guidelines, eligibility criteria and recognised deployment options.³

International experience reinforces this balanced approach. The EU Data Act⁴ supports portability, cloud switching and interoperability between data-processing services, and also includes safeguards against unlawful third-country governmental access to non-personal data held in the EU. FedRAMP (Federal Risk and Authorisation Management Program) demonstrates how government cloud adoption can be channelled through standardised authorisation and certified marketplaces; and Gaia-X highlights the value of interoperability, trust and control over data.⁵

India already has important building blocks in place, including the Guidelines for cloud selection framework,⁶ GI Cloud (MeghRaj), government and state data centre capabilities, the MeitY Cloud Empanelment Programme and the IndiaAI Compute Programme. The recent cloud-selection guidelines provide a baseline under which workloads are assessed according to their categorisation and aligned with appropriate deployment options, safeguards and procurement routes.



India's path is not cloud versus control, it is cloud with verifiable control

However, data location alone does not establish sovereignty. **Sovereignty to be assessed by demonstrable control over data, compute, administrative operations, encryption keys, legal protection, resilience and exit, rather than by location alone. Data residency is necessary but not sufficient. Providers may be subject to lawful orders in jurisdictions in which they operate or are incorporated.** Practical exposure, therefore, depends on factors such as corporate structure, data custody and control, encryption design, contractual commitments and conflict-of-law mechanisms. Beyond data sovereignty, India now faces a

second frontier: AI sovereignty. As AI becomes embedded in public services, financial systems, healthcare, defence and critical infrastructure, control over compute access, training data, model artefacts and inference infrastructure becomes strategically important. Global supply chains, export controls and platform dependencies are increasingly shaping access to AI infrastructure. India's sovereign cloud strategy to therefore extend beyond hosting and localisation to include AI compute, model governance, monitoring, auditability and exit readiness.



Control requirements

A sovereign cloud environment is one in which the controls required for a defined workload can be demonstrated across data, operations, legal and contractual accountability, technology, resilience, governance processes, continuity of access and exit. The required controls are anchored in the Indian legal framework and policy for the in-scope workloads, while allowing the deployment architecture and technology supply chain to vary by risk. Sovereignty must also mean protection against unilateral service lockouts, suspensions or deplatforming. For government and non-government workloads, cloud service providers should not suspend, degrade or terminate services without defined continuity procedures, government escalation mechanisms, transition assistance, recoverable backups, alternative operating arrangements, and, where appropriate, government step-in rights.

*Sovereignty is proven
by evidence, not labels*

This is not a theoretical safeguard: Sanctions-linked disruptions to Indian entities show that a single provider's action can interrupt operations even when data residency and other requirements are otherwise met.⁷ Non-unilateral suspension clauses, government escalation rights and tested continuity arrangements are therefore core, evidenced controls, not optional contract language layered on top of residency.

This definition shifts the debate from labels to evidence. Hosting data domestically is a starting point, but it does not by itself establish sovereignty. Real control depends on how infrastructure is operated, which platforms are chosen and how authority is exercised during normal operations and during periods of disruption. Before testing control outcomes, a ministry or department must determine the classification and whether the workload is eligible for cloud hosting. Hence, this definition must be operationalised across the following sovereignty dimensions:

⁷ Note: Recent sanctions-linked technology disruptions, show that digital dependency can quickly become an operational continuity risk. For government departments, this risk cannot be addressed through data residency alone. Contracts and architectures therefore require non-unilateral suspension clauses, government escalation rights, transition assistance, recoverable backups, tested portability, alternate capacity and continuity arrangements that prevent a single provider action from disrupting essential public services


Dimension	Decision question	Evidence
Classification and eligibility	What are the existing classifications, and is cloud hosting permitted?	Classification record; Top Secret/Secret exclusion check; Category mapping; approval trail
Data	Where are primary data, logs, metadata and backups stored and processed?	Residency map; transfer controls; retention and deletion evidence
Compute and AI infrastructure	Where are compute, GPUs/ accelerators, training environments, inference endpoints and model artefacts physically located and operated?	India-resident compute and GPU location evidence; sovereign operator attestation; workload placement records; no-offshore-processing evidence; infrastructure ownership disclosure
Operations	Who can administer the service, from where, and under whose authority?	Privileged-access model; local operations; immutable logs; segregation of duties
Legal and contractual	Which entities and laws govern performance, disclosure, audit and dispute resolution?	Contract clauses; corporate structure; legal process handling; audit and notification and challenge rights, non-unilateral suspension clauses
Digital continuity sovereignty	Can essential government services continue if provider access is restricted or terminated?	Recoverable backups; exit testing results; alternate provider readiness; control-plane contingency plans; sovereign continuity procedures
Technical	Who controls identities, keys, configurations, interfaces and portability?	Customer-controlled keys; security architecture; open interfaces; tested exit plan
AI governance	Who controls training data, model artefacts, evaluation, inference, monitoring and fallback arrangements?	Model/data lineage; evaluation records; access controls; model registry; inference logs; deployment and exit evidence.



A related nuance deserves explicit treatment: data residency and compute residency are often treated as interchangeable, when they are not. Architectures that keep primary data in India while attaching overseas accelerators over a network fabric are not a like-for-like substitute for India-based compute model training, which requires data to sit adjacent to the processing hardware, so data moves to the GPU rather than the reverse. Therefore, government data centre and hosting providers are required to demonstrate, on a workload-by-workload basis, that the GPU capacity used for sensitive processing is itself India-resident, and that the associated data is, especially for critical and sensitive workloads.

Why this matters now

The strategic importance of cloud is rising because cloud, data centres, connectivity, AI compute, LLM models and applications/products increasingly function as a single digital infrastructure stack. This makes a more evidence led model necessary for five reasons:





1. **Scale and concentration:** Rapid adoption can concentrate operational risk in a small set of technologies, operators or regions. Concentration shall be measured across providers, software dependencies, identity systems, network paths and skills.



2. **Regulatory convergence:** The DPDP framework, RBI, SEBI, IRDAI along with MeitY's guidance, are creating a clearer Indian public-sector anchor for classification led cloud decisions.



3. **Cross-border legal process:** Where a provider is subject to laws in another jurisdiction, data, metadata or operational records may be exposed to lawful disclosure requests if they are within the provider's possession, custody or control. Government entities should therefore assess this risk upfront and require transparency on disclosure-request handling, challenge procedures, escalation mechanisms and notification obligations, wherever legally permissible.



4. **AI infrastructure:** AI systems add concentrated compute, sensitive training data, model artefacts and continuous inference to the control surface. The IndiaAI common compute programme demonstrates significant momentum, with the Government reporting the onboarding of substantial GPU⁸. For sensitive government workloads, aggregate GPU counts alone should not be treated as sufficient evidence; providers should be required to demonstrate physical GPU location, operator ownership, subcontractor involvement, workload allocation, network controls and committed sovereign capacity.




5. **Continuity and choice:** Geopolitical, commercial, cyber and supply chain events can affect any provider or technology. Portability, tested exit, multi-region design and alternative capacity matter alongside sovereignty controls. Concentration risk shall be formally assessed across cloud providers, identity systems, AI platforms, and control planes that support critical government workloads.

Implementation enhancements

The cloud selection framework establishes a clear classification-led baseline for cloud adoption,⁹ but ministries and departments will need detailed operating guidance, including decision trees, workload examples, approval workflows and evidence templates, to apply the framework consistently across government.

1. The first enhancement is implementation consistency. Without common playbooks, similar workloads may be classified differently across ministries, states or agencies. Practical guidelines covering examples of categories, exception handling, escalation routes and minimum control baselines for each category shall be laid out. These shall specify what departments must provide evidence of before procurement: data and backup location, administrative access, encryption/key control, logs, location of the AI infrastructure, DR arrangements, audit rights and exit readiness.
2. The second enhancement is assurance and monitoring. Sovereignty shall be proven by evidence, not merely by provider claims or contract language. FedRAMP is relevant as it uses authorised cloud offerings, recognised assessors and repeatable assessment processes to support government cloud adoption.¹⁰ The UK NCSC cloud security principles similarly emphasise examining both the service and the provider, and testing whether evidence gives sufficient confidence in the claimed controls.¹¹ India may adapt a few of these lessons through a stronger model, supported by standard evidence packs, independent audits, recurring control tests and continuous monitoring.
3. The third enhancement is procurement translation. The deployment routes have been identified, but departments still need model RFP clauses and contract language. These are to cover the named legal operator, subcontractors, privileged-access boundaries, locations of data and compute, metadata, logs and backups, disclosure-request handling, incident notification, continuity, portability, exit support and deletion certification. The EU Data Act is useful here because it emphasises cloud switching, interoperability and reducing barriers to moving data between providers.¹² India may similarly require tested exit plans, transparent egress and ingress terms and open interfaces for critical workloads.
4. The fourth enhancement is data lifecycle control. Sovereignty does not end once data is hosted in an approved environment. Sensitive data may continue to exist in backups, replicas, logs, snapshots, retired media or object stores. India can therefore develop Data Lifecycle Sovereignty Standards, such as NIST SP 800-88 Revision 2, that cover sanitisation, cryptographic erasure, backup disposition, deletion certificates and third-party validation.¹³
5. The fifth enhancement is AI-specific assurance. AI workloads introduce risks that traditional cloud guidance does not fully address such as compute access, training data lineage, model artefacts, inference location, monitoring, human oversight and fallback arrangements. NIST's AI Risk Management Framework is relevant because it treats trustworthy AI as a lifecycle issue across design, development, use and evaluation.¹⁴ India may therefore create a critical AI assurance profile that links cloud classification with AI-specific controls.



AI sovereignty: The emerging frontier

AI sovereignty is an important emerging dimension of cloud sovereignty in 2026 and beyond. As India deploys AI across public services, financial systems, healthcare and defence, the question of who controls the compute, the training data and the model artefacts becomes a matter of national capability, public trust and strategic risk.

The cloud selection framework may be applied to AI workloads before architecture choices are made. AI systems involving national security, public safety, financial stability, welfare delivery or sensitive citizen datasets must first be classified for cloud eligibility and Category A/ B¹⁵ treatment. If an AI workload is top secret or secret, cloud hosting should be excluded; Example for Category A, inference, data stores, model artefacts and backup/DR arrangements should follow the stronger government- or MeitY-notified deployment path.

AI sovereignty should not be framed as technological isolation or provider exclusion. India will continue to benefit from global innovation, open source models, open weight models, commercial models and platforms and international partnerships. Critical AI workloads must remain governable under Indian legal, operational and security requirements, with verifiable control over the parts of the AI lifecycle that pose public, economic or national security risks.

The stack may therefore be governed through four control points:

- **Compute access and continuity:** India to treat AI compute as critical digital infrastructure by combining IndiaAI common compute capacity, trusted international access arrangements, domestic accelerator design capability and continuity planning for high-impact workloads. For sensitive government and public interest use cases, compute access should be governed in accordance with the applicable Cloud selection framework and AI assurance profile. For sensitive government workloads, sovereign AI controls require that training, fine-tuning and inference processing occur on compute capacity physically located in India unless a specific exemption has been approved.
- **Training data governance:** AI models trained or fine-tuned on sensitive Indian data, including health, financial, legal, welfare and public-service datasets, should be subject to clear Indian governance rules. This includes purpose limitation, lawful basis or consent, provenance, retention, anonymisation, auditability and deletion. For sensitive use cases, India should develop practical standards for privacy-preserving learning, synthetic data, federated learning, differential privacy and secure data enclaves
- **Model artefact and audit control:** For critical national or regulated applications, sovereignty requires more than access to an API. The entity should be able to understand, test, monitor, replace, and, where necessary,

audit the model or model artefacts on which the service depends. This may include access to model cards, evaluation results, safety configurations, fine-tuning records, prompt and retrieval pipelines, incident logs, contractual audit rights and exit options. Critical AI systems shall maintain documented fallback arrangements that allow replacement of models, providers or inference platforms without material disruption to public services.


- **Inference infrastructure and operational accountability:** Even where model training or pre-training occurs outside India, inference for critical services may need to operate within approved Indian legal and operational boundaries. This is especially important where real-time decisions affect citizens, public safety, financial stability or essential services. Critical inference environments shall have clear rules for hosting location, administrator access, logging, human oversight, fallback arrangements, incident response and continuity.

India's sovereign-cloud framework may include a dedicated critical AI assurance profile. This profile classifies AI workloads by impact, defines control requirements for compute, data, models and inference, and requires evidence that critical AI systems can be monitored, governed, recovered and exited. AI sovereignty is, therefore, not a separate agenda from cloud sovereignty; it is the next layer of it.



Sovereign cloud for India: Building trusted digital infrastructure for the AI era

Closing Perspective



Sovereign cloud is no longer a narrow hosting decision; it is a strategic operating model for the AI era. Stronger controls may increase costs, limit access to some managed services, or add operational complexity, but they also strengthen resilience, accountability, negotiating leverage and long-term choice. The objective is not to slow cloud adoption, nor to frame sovereignty as a choice between domestic control and global innovation. The objective is to make cloud adoption governable, measurable and reversible for the workloads that matter most.

Sovereignty shall be ultimately measured not by the nationality of technology providers but by demonstrable control over data, compute, operations, legal processes, resilience and exit. Geography is important, but control remains the true test of sovereignty.

The path ahead involves shifting from policy goals to operational discipline: categorise workloads using the notified framework, quantify controls, incorporate evidence into procurement, enhance resilience and portability, and test exit rights rather than merely including them in contracts. Digital autonomy will not be realised through origin labels alone; it requires the capacity to understand, govern, secure, recover and adapt the systems that increasingly underpin public services, regulated sectors and critical AI workloads.

Sources

¹<https://my.idc.com/getdoc.jsp?containerId=prAP53678525>

²Rules specific to sector workloads, by the relevant regulators - RBI/SEBI/IRDAI/PFRDA/UIDAI/TRAI/CERT-In/etc.

³<https://www.meity.gov.in/static/uploads/2026/03/a49a9e2bfb5bbd9057cf3efafda5b8d7.pdf>

⁴<https://digital-strategy.ec.europa.eu/en/factpages/data-act-explained>

⁵<https://gaia-x.eu/>

⁶<https://www.meity.gov.in/static/uploads/2026/03/a49a9e2bfb5bbd9057cf3efafda5b8d7.pdf>

⁸<https://www.pib.gov.in/PressReleasePage.aspx?PRID=2225781®=3&lang=1>

⁹<https://www.meity.gov.in/static/uploads/2026/03/a49a9e2bfb5bbd9057cf3efafda5b8d7.pdf>

¹⁰<https://www.fedramp.gov/>

¹¹<https://www.ncsc.gov.uk/collection/cloud/the-cloud-security-principles>

¹²<https://digital-strategy.ec.europa.eu/en/policies/data-act>

¹³<https://csrc.nist.gov/pubs/sp/800/88/r2/final>

¹⁴<https://www.nist.gov/itl/ai-risk-management-framework>

¹⁵<https://www.meity.gov.in/static/uploads/2026/03/a49a9e2bfb5bbd9057cf3efafda5b8d7.pdf>

Connect with us

Amit Singh

Partner
Deloitte India
asingh46@deloitte.com

Ritesh Pal

Partner
Deloitte India
riteshpal@deloitte.com

Anjani Kumar

Partner
Deloitte India
anjanikumar@deloitte.com

Contributors

Anupam Sharma

Smita Dayal

Acknowledgments

Rahul Dhuria

Shreya Rajopadhye





Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited (“DTTL”), its global network of member firms, and their related entities (collectively, the “Deloitte organization”). DTTL (also referred to as “Deloitte Global”) and each of its member firms and related entities are legally separate and independent entities, which cannot obligate or bind each other in respect of third parties. DTTL and each DTTL member firm and related entity is liable only for its own acts and omissions, and not those of each other. DTTL does not provide services to clients. Please see www.deloitte.com/about to learn more.

Deloitte Asia Pacific Limited is a company limited by guarantee and a member firm of DTTL. Members of Deloitte Asia Pacific Limited and their related entities, each of which is a separate and independent legal entity, provide services from more than 100 cities across the region, including Auckland, Bangkok, Beijing, Bengaluru, Hanoi, Hong Kong, Jakarta, Kuala Lumpur, Manila, Melbourne, Mumbai, New Delhi, Osaka, Seoul, Shanghai, Singapore, Sydney, Taipei and Tokyo.

This communication contains general information only, and none of DTTL, its global network of member firms or their related entities is, by means of this communication, rendering professional advice or services. Before making any decision or taking any action that may affect your finances or your business, you should consult a qualified professional adviser.

No representations, warranties or undertakings (express or implied) are given as to the accuracy or completeness of the information in this communication, and none of DTTL, its member firms, related entities, employees or agents shall be liable or responsible for any loss or damage whatsoever arising directly or indirectly in connection with any person relying on this communication.

© 2026 Deloitte Touche Tohmatsu India LLP. Member of Deloitte Touche Tohmatsu Limited