CONSIVA.AI

PRIVACY OPERATIONS RESEARCH · INDIA · 2026

The Provable Enterprise

Architecting DPDP compliance backwards from evidence — a privacy operating model for India's May 2027 deadline

AN INDUSTRY RESEARCH PAPER · **EDITION 1.0** · AUGUST 2026 · **CONSIVA.AI**

ABOUT THIS PAPER

Scope, method and limits

This paper argues that India's Digital Personal Data Protection regime has quietly changed what compliance *is* — from a body of documentation an enterprise maintains, to a stream of evidence an enterprise must be able to produce on demand. Everything else in these pages follows from that claim.

Who this is written for

Chief Information Officers, Chief Information Security Officers, Data Protection Officers, General Counsel and Heads of Compliance, Chief Data and AI Officers, enterprise architects, and the founders and CTOs of digital businesses processing the personal data of people in India. It assumes professional familiarity with enterprise systems but not with privacy law.

How claims are marked

Privacy writing routinely blurs what the law requires, what practitioners believe it requires, and what a vendor happens to sell. This paper labels every substantive claim so a reader can audit the difference:

STATUTORY	An express requirement of the DPDP Act, 2023 or the DPDP Rules, 2025, cited to section or rule.
INTERPRETATION	A reasonable reading of the law where the text is silent, ambiguous, or awaiting regulatory guidance. Contestable.
BEST PRACTICE	A recommendation drawn from engineering and audit practice. Not legally mandated.
CONSIVA TODAY	A capability available in the Consiva.ai platform at the date of publication.
RECOMMENDATION	A capability this paper argues the market needs. Not a statement that Consiva provides it.
EMERGING	A forward-looking technology concept, included for strategic orientation only.

LEGAL NOTICE

This white paper is intended for informational and technology-strategy purposes and does not constitute legal advice. It is not a substitute for advice from qualified counsel on your organisation's specific circumstances. Statutory references are to the Digital Personal Data Protection Act, 2023 (22 of 2023) and the Digital Personal Data Protection Rules, 2025 (G.S.R. 846(E), notified 13 November 2025). Regulatory positions described here are current as at August 2026 and are subject to change; several matters discussed remain pending further notification by the Central Government or the Data Protection Board of India.

Where this paper describes technologies, techniques or capabilities not marked CONSIVA TODAY, no representation is made that Consiva.ai provides them. Nothing in this paper should be read as a claim of regulatory approval, certification, or endorsement by any authority.

CONTENTS

What this paper covers

READING THIS EDITION

Parts I and II establish the thesis and the operating model. They are self-contained: a reader who stops at the end of Chapter 6 should be able to run a materially better DPDP programme, with or without any vendor. Parts III and IV apply the model discipline by discipline, and then to the calendar.

EXECUTIVE SUMMARY

Compliance you cannot evidence is compliance you do not have

India's data protection regime is often described as a consent law. It is not. It is an accountability law that happens to use consent as one of its instruments — and accountability, unlike consent, cannot be installed. It has to be operated.

The distinction matters commercially, not just philosophically. An enterprise can deploy a consent banner in an afternoon. It cannot, in an afternoon, answer the question a regulator actually asks after an incident: *show me the record.* Show me which purposes this individual agreed to, on which version of which notice, in which language, on what date. Show me every system that received their data afterwards, and show me that each of those systems stopped processing when consent was withdrawn. Show me who was accountable, and show me the artefact that proves the control worked rather than the policy that said it should.

Most Indian enterprises today cannot answer that question, and the reason is structural rather than negligent. Privacy programmes have been built as documentation exercises — policies, registers, questionnaires, annual assessments — because that is what the previous decade of compliance rewarded. The DPDP Rules, 2025 reward something different. They repeatedly require not a policy but a *demonstrable measure*: logs retained for a fixed period, breach intimation without delay followed by a detailed report, technical and organisational measures whose *effective observance* can be verified. The regime is written in the vocabulary of evidence.

Six propositions

1 · Consent alone is not compliance

Consent is one lawful basis with one lifecycle. It says nothing about where data went afterwards, how long it stayed, who accessed it, or whether withdrawal actually propagated. A consent platform that ends at the banner has automated the easiest ten per cent of the obligation.

2 · You cannot govern what you cannot see

Every downstream obligation — rights fulfilment, retention, erasure, breach scoping, processor governance — depends on knowing where personal data physically resides. Discovery is not a feature of a privacy programme. It is the precondition for one.

4 · Evidence is the deliverable

Reframe the unit of compliance work. The output of a privacy control is not a completed task; it is a timestamped, tamper-evident, exportable artefact that survives staff turnover and satisfies a regulator who was not in the room.

5 · Periodic assessment cannot govern continuous processing

Data estates change daily; annual assessments describe an enterprise that no longer exists. The only defensible posture against continuous processing is continuous observation.

3 · Controls must fire, not merely exist

The gap that produces penalties is rarely between *no policy* and *a policy*. It is between a policy and an enforced control — between a retention schedule and an actual deletion, between a stated purpose limitation and an API that refuses the call.

6 · AI collapses the remaining slack

Retrieval systems, embeddings, agent memory and tool-calling move personal data through paths no data map anticipated, at machine speed. Privacy architecture that assumed human-paced processing is already behind.

What this paper proposes

Two connected instruments, developed in Chapter 6 and applied throughout.

The Evidence Chain is a decomposition. It holds that every obligation under the Act and Rules resolves into five linked elements — an **obligation**, a **control** that discharges it, a machine-observable **signal** that the control fired, a durable **artefact** that records the signal, and a named accountable **owner**. An obligation missing any link is not managed; it is merely acknowledged. The chain is deliberately mundane, and that is its value: it converts a 40-page statute into a finite, auditable register that a CIO and a General Counsel can read in the same meeting.

The Privacy Operating Loop is the operating model that produces those chains and keeps them alive: *See, Map, Govern, Execute, Prove, Watch.* Most enterprise programmes today implement the first three phases well and the last three barely at all. That asymmetry is the single best predictor of which organisations will struggle in 2027.

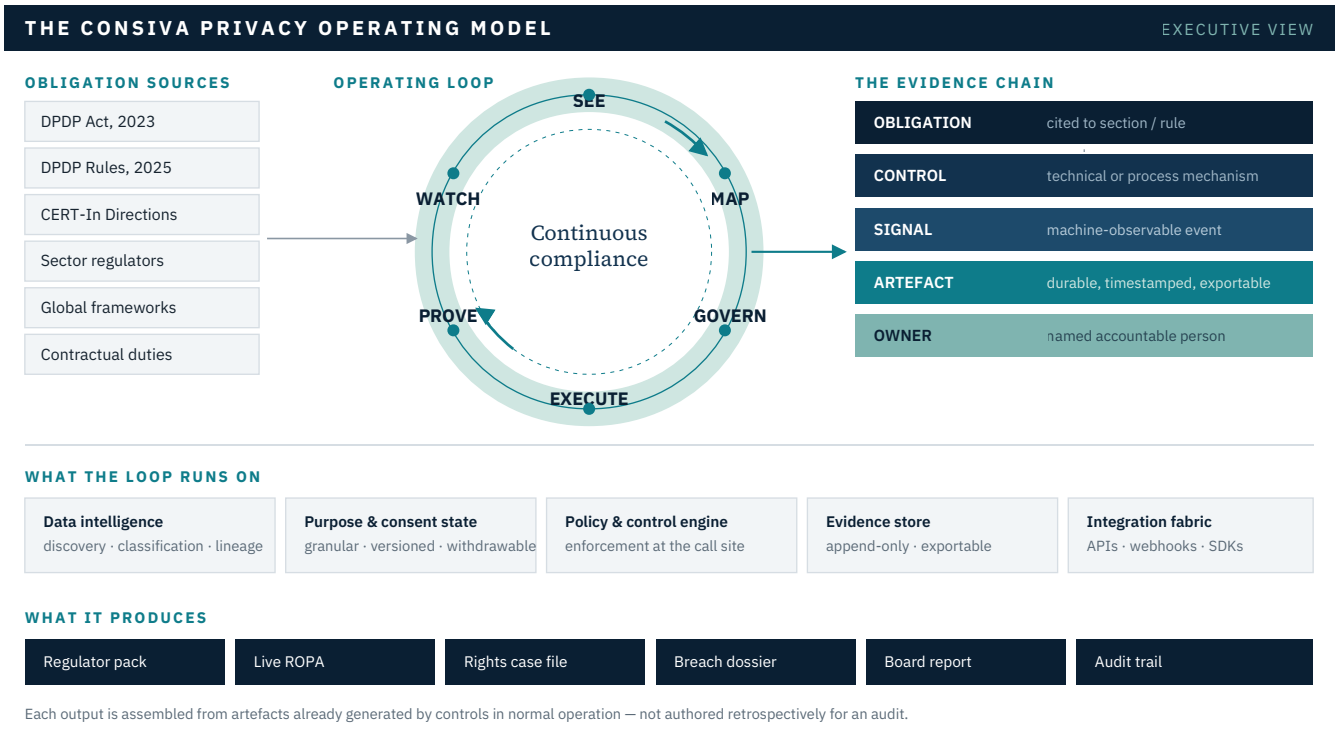

Figure 1 — Executive view of the operating model. Obligations enter from multiple regulatory and contractual sources; the operating loop converts each into an Evidence Chain; the chains, in aggregate, are what the enterprise produces when asked to demonstrate compliance. Note the direction of travel: evidence is a by-product of operation, not a document written afterwards.

What follows from this

Three consequences deserve to be stated plainly, because each contradicts a widely-held assumption.

First, the calendar is more generous and more dangerous than it appears. Most substantive obligations under the Rules commence eighteen months from notification — around 13 May 2027. That is a real runway. It is also the reason so little has moved: an eighteen-month horizon reads as permission to wait, when in fact the work that takes longest,

data discovery and control instrumentation across a live estate, is precisely the work that cannot be compressed at the end.

Second, the penalty structure tells you where to spend. The Schedule to the Act does not weight all failures equally. Security safeguards carry the highest ceiling; consent and rights failures fall into the residual band. An enterprise that spends its entire privacy budget on the consent layer has optimised against the smallest exposure. Chapter 2 works through this.

Third, the hardest problems are conflicts, not gaps. The genuinely difficult questions in Indian privacy engineering are not "do we have a control" but "which law wins" — when a Data Principal demands erasure and a sectoral retention mandate forbids it. No amount of consent tooling resolves that. Chapter 16 treats it as the discipline it deserves to be.

> Compliance should not merely tell an enterprise what it must not do. It should give the enterprise the confidence to know what it can safely do.
>
> THE ARGUMENT OF THIS PAPER, STATED IN ADVANCE

How to read this paper by role

The argument is cumulative, but the chapters are usable independently. If you have limited time, the shortest defensible path through the document depends on what you are accountable for.

IF YOU ARE	YOUR BINDING QUESTION	READ FIRST	THEN
Chief Information Officer / CTO	Can my estate support these obligations, and what has the longest lead time?	Chapters 3, 8, 19	6, 12
Chief Information Security Officer	Where is my largest penalty exposure, and who owns each breach clock?	Chapters 2, 13	11, 14
Data Protection Officer	Which obligations can I currently evidence, and which can I not?	Chapters 4, 6, 9	10, 15, 16
General Counsel / Head of Compliance	What does the law actually require, as distinct from market assumption?	Chapters 1, 2, 16	7, 15
Chief Data / AI Officer	What must be settled before we train or deploy anything further?	Chapters 12, 11	8, 16
Chief Executive / Board member	Are we exposed, and is someone accountable?	This summary, then 4 and 17	6.4, 19
Founder of a growing digital business	What is the minimum that is genuinely defensible?	Chapters 7, 9, 19	2, 13

THREE CLAIMS TO TEST US ON

A research paper should be falsifiable. If any of the following is wrong, the argument weakens materially and we would like to know: that most substantive DPDP obligations commence around 13 May 2027 rather than being in force now; that the ₹250 crore ceiling attaches to security safeguards rather than to consent or rights; and that the DPDP breach duty requires intimation *without delay* rather than within seventy-two hours. Each is sourced to a specific rule in the chapters that follow.

I

The problem is not the law. It is the estate.

India’s statute is, by international comparison, short and readable. What makes it hard is not interpretive difficulty but operational surface area: the number of systems, teams, vendors and data paths across which a small number of duties must hold true simultaneously — and continuously.

Chapters 1 to 4 establish why the enterprise estate, not the legal text, is the binding constraint.

CHAPTER ONE

The privacy inflection point: why May 2027 is an evidence deadline

Three things changed at once in Indian enterprises: the volume of personal data became unmappable by hand, the number of systems touching it became uncountable by memory, and the speed at which it moves became faster than any review cycle. The law arrived third.

1.1 The commencement calendar, precisely

A striking amount of published commentary about DPDP is wrong about when it applies. The confusion is understandable — the Act received assent in August 2023, the Rules were notified on 13 November 2025, and both events were widely reported as the law "coming into force." The Rules themselves are explicit, and they phase commencement in three tranches.

> **STATUTORY · RULE 1(2)–(4), DPDP RULES, 2025**
>
> Rules 1, 2 and 17 to 21 commenced on the date of publication in the Official Gazette. Rule 4 commences **one year** after that date. Rules 3, 5 to 16, 22 and 23 commence **eighteen months** after that date.
>
> Source: G.S.R. 846(E), Ministry of Electronics and Information Technology, notified 13 November 2025.

Read against a notification date of 13 November 2025, that produces the timeline below. The practical consequence is that the obligations most enterprises think of as "DPDP compliance" — the form of notice, reasonable security safeguards, breach intimation, retention and erasure, the mechanics of rights, the additional duties of Significant Data Fiduciaries — sit in the eighteen-month tranche.

Figure 2 — DPDP commencement, by tranche. Dates computed from the Rule 1 commencement provisions against a notification date of 13 November 2025. STATUTORY for the tranches; the characterisation of which obligations matter operationally is the authors'.

A FACT WORTH VERIFYING BEFORE YOU RELY ON IT

The Data Protection Board of India was constituted in law on 13 November 2025. Whether it is *staffed* is a separate question, and the public record is genuinely contested. Reporting through mid-2026 indicates that MeitY sought nominations for the Chairperson and Member posts in May 2026 and issued a further notification concerning the appointment process in June 2026, with the Cabinet Secretary–led search-cum-selection committee still evaluating candidates. Several secondary sources have read that June notification as the appointments themselves.

We flag this deliberately. It is a fast-moving fact of exactly the kind that privacy vendors and consultancies restate confidently and incorrectly, and any enterprise planning around Board capacity should verify it directly rather than trusting a white paper — including this one.

1.2 What actually changed in the enterprise

It is tempting to treat 2027 as the cause of the current scramble. It is more accurate to say the law has finally caught up with an operational reality that had been drifting for a decade. Four shifts compounded.

The estate stopped being enumerable

A mid-sized Indian business in 2015 could plausibly list every system holding customer data on a whiteboard. The same business today runs a primary application, two or three legacy databases nobody has decommissioned, a CRM, a marketing automation platform, a support desk, an analytics warehouse, a payments processor, a logistics partner, a WhatsApp business account, and a tag manager through which the marketing team has installed trackers that no one in engineering approved. Personal data is in all of them. Nobody holds the complete list, because the list changes weekly.

Third parties became first-class processors

Each of those platforms is a processor or a sub-processor, and each one has its own sub-processors. The enterprise remains the accountable Data Fiduciary regardless. Rule 6(1)(f) requires appropriate contractual provision for safeguards where a Data Processor is engaged — but a contract is only the first link of a chain that must also include verification and evidence.

Consent stopped being a page and became state

When consent could be inferred from a signed form in a filing cabinet, it was a document. When it is granular, purpose-specific, versioned, multilingual and revocable at any moment, it is a piece of live state that every downstream system must be able to read — and act on. Very few enterprise architectures were designed with a consent state that propagates.

AI made data paths non-deterministic

A retrieval-augmented system can surface a personal data field to a user who was never authorised to see it, through a path no data flow diagram anticipated, because the path was constructed at query time. An agent with tool access can write personal data into a third-party service in a single call. Chapter 12 takes this seriously; here it is enough to note that it removes the last of the slack.

> The question is no longer whether an enterprise *has* privacy documentation. It is whether the enterprise can produce, on a Tuesday afternoon and without a project, the evidence that its controls operated on a specific individual's data on a specific date.

1.3 Can policies, spreadsheets and periodic assessments still carry this?

The honest answer is that they can carry the *documentation* obligations indefinitely and cannot carry the *operational* ones at all. A spreadsheet can record that a retention period is two years. It cannot delete anything. An annual assessment can conclude that access controls are adequate. It cannot demonstrate that a specific access on a specific date was authorised by a specific purpose. A policy can state that consent withdrawal is honoured across all systems. Only a propagation mechanism and its logs can show it was.

This is the fault line the rest of the paper runs along. It is not a criticism of documentation, which remains necessary. It is an argument that documentation has been mistaken for the whole of the discipline.

CHAPTER TWO

DPDP is a techno-legal operating challenge, not a legal project

Indian privacy obligations sit at an intersection that no single function in most enterprises owns. Legal understands the duty and not the system. Engineering understands the system and not the duty. The gap between them is where compliance programmes go to die.

2.1 Six domains, one obligation

Consider a single, apparently simple requirement: that processing be limited to the purpose for which consent was given. Discharging it requires a legal reading of what "purpose" means and how specifically it must be described; a data architecture that records purpose alongside each data element; an access-control model that can evaluate purpose at request time; a set of business processes that do not quietly reuse data for a second purpose; a security posture that prevents unauthorised paths around the control; and, increasingly, governance over AI systems that may use the data in ways no purpose statement contemplated.

Six domains. One clause. This is what "techno-legal" means in practice, and it is why assigning DPDP to the legal department alone reliably produces a programme that is well-documented and non-operational.

2.2 The penalty structure is a budget instruction

The Schedule to the Act does not treat all failures as equivalent. It sets distinct ceilings by category of default, and the pattern of those ceilings is informative: the largest exposure attaches to security failure, not to consent failure.

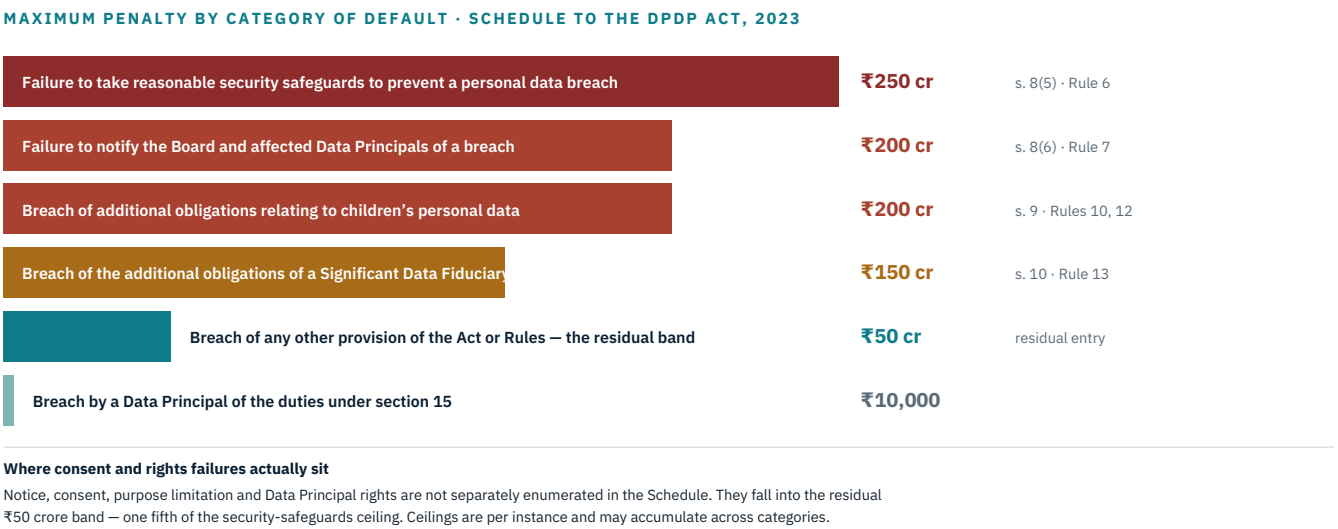

Figure 3 — The penalty ladder, drawn to scale. STATUTORY Bar lengths are proportional to the ceilings. The visual point is the one enterprises most often miss: the consent layer, where most privacy spending goes, is attached to the smallest band.

A CORRECTION WORTH MAKING EXPLICITLY

It is common — including in vendor marketing, and at times in our own earlier material — to see the ₹250 crore figure attached to inadequate consent mechanisms or failures to fulfil Data Principal rights. That attribution is incorrect. The ₹250 crore ceiling in the Schedule attaches to the failure to take reasonable security safeguards. Consent and rights failures fall within the residual band. An enterprise that has budgeted against the wrong number has probably also prioritised against it.

None of this argues for neglecting consent. It argues for proportion. A programme that ships a beautiful consent banner while leaving access logging, encryption at rest, backup integrity and processor contracts untouched has inverted the risk it faces. Rule 6(1) is a compact list of seven safeguards and it is, on the evidence of the Schedule, the most expensive thing in the Rules to get wrong.

2.3 Mapping law to enterprise reality

The instrument below is the one we return to throughout. It reads a statutory duty downward until it reaches something a system does and something an auditor can see. Every row that cannot be completed to the right-hand column is an obligation the enterprise is asserting rather than discharging.

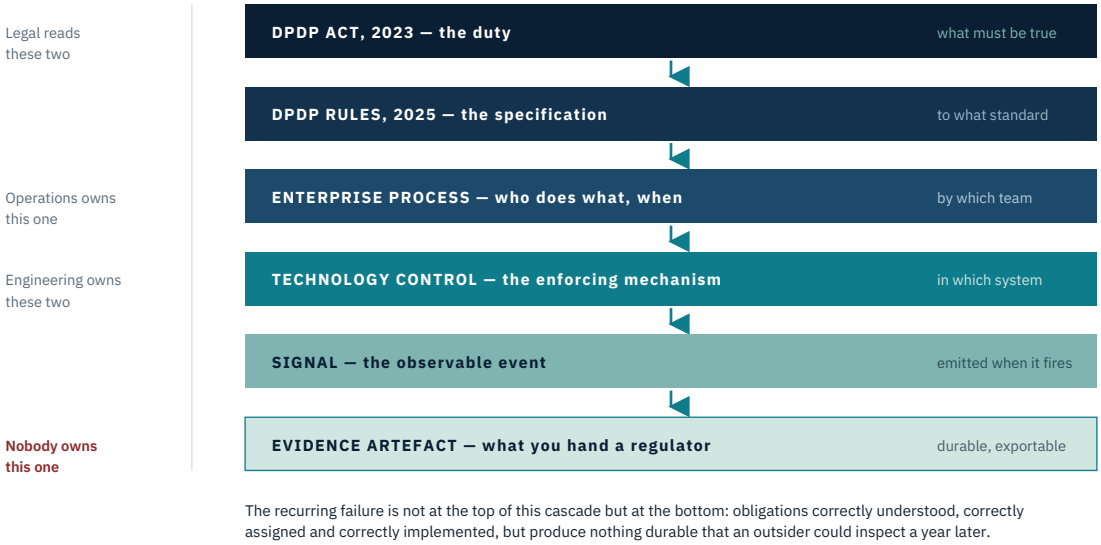

Figure 4 — The legal-to-evidence cascade. The left margin shows the ownership pattern we observe most often. The evidence layer is orphaned by default because no function's job description includes it.

2.4 Worked mappings

Four illustrative rows, chosen because each is commonly mishandled.

OBLIGATION	SPECIFICATION	TECHNOLOGY CONTROL	SIGNAL	EVIDENCE ARTEFACT
Notice must be independently understandable, itemised, in clear language Rule 3	Itemised description of the personal data; the specified purposes; the means to withdraw, exercise rights, and complain	Versioned notice templates bound to purpose definitions; language variants; the notice served is the notice recorded	Notice-served event carrying template version, language, timestamp, purpose set	Reproducible copy of the exact notice shown to a named individual on named date
Reasonable security safeguards, at minimum seven enumerated measures Rule 6(1)(a)–(g)	Encryption, obfuscation, masking or virtual tokens; access control; visibility through logs; continuity measures; one-year log retention; processor contract terms; organisational measures	Encryption at rest and in transit; least-privilege access model; centralised access logging with review; tested backups	Access-log entries; key-rotation events; backup-restore test results; contract-status records	Log archive covering the statutory retention period, with review records demonstratin the logs were actually examined

OBLIGATION	SPECIFICATION	TECHNOLOGY CONTROL	SIGNAL	EVIDENCE ARTEFACT
Breach intimation to affected Data Principals and to the Board Rule 7	To each affected Data Principal without delay, with prescribed content; to the Board without delay, then detailed information within seventy-two hours	Incident intake with automatic classification; scoping against the data map; templated notifications per audience; independent clock tracking	Detection timestamp; awareness timestamp; each notification dispatch; each regulatory filing	Chronological incident dossier reconciling awareness time to every notification, with content as sent
Grievance redressal within a defined period Rule 14(3)	Published means of exercising rights; a grievance system responding within a period not exceeding ninety days, supported by appropriate technical and organisational measures	Rights intake with identity verification; routing to system owners; clock tracking with escalation before expiry	Intake, verification, routing, fulfilment and closure events per case	Per-case file showing what was requested, what was done in which systems, by whom, and when

ON TWO NUMBERS THAT CIRCULATE INCORRECTLY

There is no thirty-day DPDP rights deadline. Thirty days is the GDPR position under Article 12. The only period the Rules attach to Data Principal requests is in Rule 14(3): a grievance redressal system responding within a period *not exceeding ninety days*. For access, correction and erasure the Rules prescribe no period at all, which means the standard is the Data Fiduciary's own published terms read alongside its general accountability. **BEST PRACTICE** Committing publicly to a materially shorter internal service level — and meeting it — is a defensible posture and a competitive one. Presenting it as the statutory requirement is not.

The DPDP breach obligation is not a seventy-two-hour clock. Rule 7 creates a sequence, examined in Chapter 13: intimation without delay, then detail within seventy-two hours. Treating seventy-two hours as the deadline for first contact misreads the rule in the direction most likely to attract a penalty.

CHAPTER THREE

The enterprise personal data lifecycle

Privacy obligations are usually taught clause by clause. Enterprises do not experience them that way. They experience a single flow of data through their business, with duties attaching at different points along it — and gaps opening wherever the flow crosses a boundary nobody owns.

3.1 Reading the lifecycle as an operational map

The diagram overleaf traces personal data from the moment of collection to the moment of erasure, and marks where each principal duty attaches. Three properties of it are worth stating before you look at it, because they are the properties that make the lifecycle hard rather than merely long.

It branches, and the branches are the risk. A single collection event fans out into an operational database, an analytics warehouse, a marketing platform, a support system and one or more processors. Each branch is a place where purpose can drift and where consent state can fail to arrive. Enterprises reliably map the trunk and reliably miss the branches.

It has re-entrant loops. Data does not simply flow forward. Analytics outputs feed back into personalisation; model training absorbs historical records and emits a system whose behaviour depends on data that may later be withdrawn; support tickets create new copies of data already governed elsewhere. A linear diagram flatters the architecture.

Obligations attach asymmetrically. Notice and consent attach once, at the front. Security safeguards attach everywhere, continuously. Rights attach unpredictably, from outside, at any point. Retention attaches at the end but must be decided at the beginning. This asymmetry is why a stage-gated compliance project — do consent, then do discovery, then do rights — produces a programme that is never coherent at any single moment.

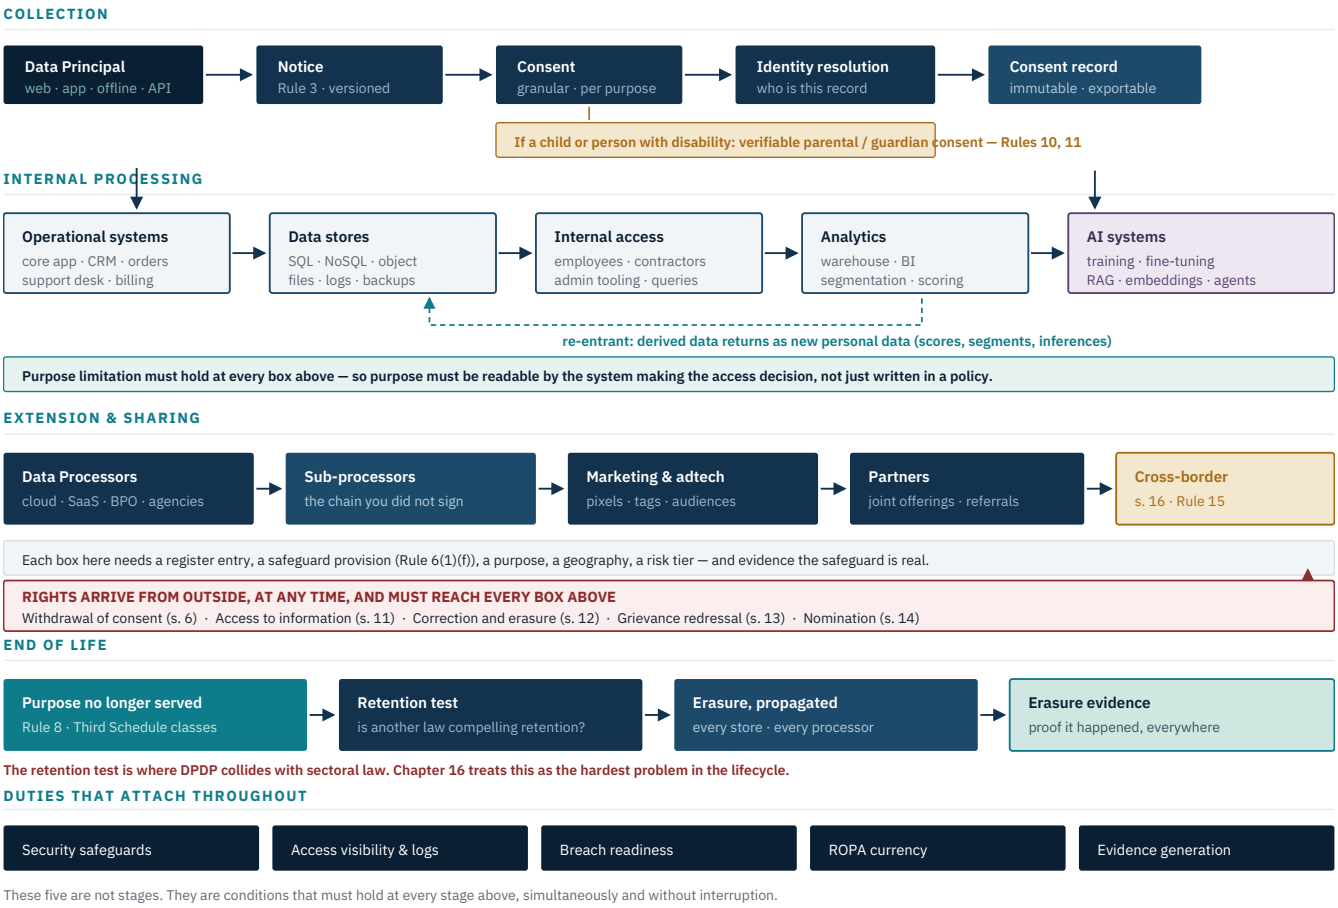

Figure 5 — The enterprise personal data lifecycle, with duties attached. Lanes are operational; the coloured callouts mark where specific statutory duties bind. Note the asymmetry: only the top lane is a sequence. Everything below it is a set of simultaneous conditions.

3.2 Where the lifecycle usually breaks

Across the programmes we have examined, failures cluster at four boundaries rather than being distributed evenly.

- **Collection to processing.** Consent is captured correctly and then not carried. The consent record exists; the operational database has no column that references it. Every subsequent control is therefore working from an assumption rather than a fact.

- **Processing to sharing.** A processor is engaged by a business team through a procurement process that never touched privacy. The register, if it exists, is a spreadsheet maintained by someone who was not in the room.
- **Anywhere to rights.** An erasure request arrives and there is no mechanism to enumerate the systems holding that individual's data. The request is answered on the basis of institutional memory, which is to say incorrectly.
- **Retention to deletion.** A retention policy is defined, no deletion job exists, and the absence is invisible because nothing fails when data is not deleted. This is the quietest failure in the lifecycle and the one most likely to be discovered by an auditor rather than an engineer.

CHAPTER FOUR

Why DPDP programmes fail at scale: the evidence gap

Programmes rarely fail because a requirement was missed. They fail because a requirement was satisfied in a form that cannot be demonstrated, propagated, or repeated — and the enterprise did not notice the difference until someone asked.

4.1 Ten failure modes

These are not hypothetical. Each is the recognisable shape of a real programme, and each has the same underlying structure: an obligation whose Evidence Chain is broken at a specific link.

#	FAILURE MODE	WHAT IT LOOKS LIKE IN PRACTICE	BROKEN LINK
1	**Policy without enforcement**	A published policy states a rule no system implements. The rule is true on paper and false in production.	Control absent
2	**Consent without propagation**	Consent is captured accurately at the edge and never reaches the systems that process the data. Withdrawal stops the banner, not the processing.	Control does not reach the estate
3	**Discovery without context**	A scan produces a list of columns containing e-mail addresses. Nobody can say which purpose each serves, so nothing can be decided from it.	Signal without meaning
4	**Rights workflow without system connectivity**	A ticketing queue for requests, fulfilled by engineers running ad-hoc queries. It works until volume rises or the engineer leaves.	Artefact not reproducible
5	**Impact assessment as static questionnaire**	A form completed once at launch, describing a system that has since changed twice. Filed, never revisited.	No re-evaluation signal
6	**Processor governance by spreadsheet**	A vendor list of uncertain completeness, no sub-processor visibility, contract dates that nobody monitors.	Owner unclear, artefact stale

#	FAILURE MODE	WHAT IT LOOKS LIKE IN PRACTICE	BROKEN LINK
7	**Retention without deletion evidence**	Retention schedules defined; deletion jobs absent or unverified. No artefact showing that anything was ever erased.	Artefact absent
8	**AI outside the privacy perimeter**	Models trained on production data under a governance process that predates the model. Embeddings and agent memory unmapped.	Obligation unmapped to the system
9	**Evidence assembled manually**	Compliance reporting is a quarterly project in which humans reconstruct history from logs, e-mail and memory.	Artefact generated retrospectively
10	**Periodic audit against continuous processing**	An annual review of an estate that changes weekly. Correct on the day of the audit and progressively less true thereafter.	No watch phase

THE EVIDENCE GAP — WHERE THE TEN FAILURE MODES CLUSTER

Every failure mode maps to a specific broken link in the Evidence Chain. That is the diagnostic value of the chain: it converts "our programme is weak" into "link three is missing on obligation seven."

Figure 6 — The evidence gap, by broken link. The ten failure modes are not ten separate problems. They are four, and each is a link in a chain that Chapter 6 formalises.

4.2 Documentation compliance and operational compliance

The distinction below is the practical test we recommend applying to any privacy programme, including one that has passed an audit. It is not a maturity model; both columns describe things enterprises genuinely do. The point is that only one of them survives a question asked about a specific person on a specific date.

DOCUMENTATION COMPLIANCE	OPERATIONAL COMPLIANCE
The unit of work is a document	The unit of work is a control that fires
Evidence is written for the audit	Evidence is emitted by normal operation
Truth is asserted by a person	Truth is observable in a system
State is known as of the last review	State is known as of now
Scope is what the team remembered	Scope is what discovery found
A rights request starts a project	A rights request starts a workflow
Withdrawal updates a record	Withdrawal changes system behaviour
Retention is a published schedule	Retention is an executed job with a receipt
Cost rises with every audit	Cost falls as controls mature
Staff turnover destroys knowledge	Staff turnover is survivable

Neither column is dishonest. The left column is what a decade of compliance practice taught enterprises to build. The right column is what the DPDP Rules ask for.

Figure 7 — The two compliance postures. Most Indian enterprises are somewhere in the left column and believe they are in the right. The diagnostic question is simple: could you answer for one named individual, today, without starting a project?

A TEST YOU CAN RUN THIS WEEK

Pick one real customer record at random. Without convening a project, produce: the exact notice they were shown and its version; the purposes they consented to and when; every system that currently holds their data; the legal basis for each; the retention period applicable to each; and the artefact proving that a control has been applied to that data in the last ninety days.

Time how long it takes. That duration, more than any maturity score, is your DPDP readiness — because it is approximately what a regulator's first request will feel like.

Part I in summary

THE FINDING	WHY IT CHANGES WHAT YOU DO NEXT
The calendar is the opposite of reassuring	Most substantive obligations commence around 13 May 2027. That tranche reads as permission to wait — but the work with the longest lead time, discovery across a live estate and control instrumentation, is exactly what cannot be compressed at the end.
The estate is the constraint, not the statute	The Act is short and readable. What makes compliance hard is operational surface area: the systems, teams, processors and data paths across which a few duties must hold true at the same moment, continuously.
The penalty ladder is a budget instruction	Security safeguards carry the highest ceiling; consent and rights failures fall in the residual band at one fifth of it. A programme that spent its budget on the consent layer optimised against its smallest exposure.
Failures cluster at four links, not ten	The ten recognisable failure modes reduce to four: a missing control, a meaningless signal, an absent artefact, an unclear owner. That reduction is what makes the next part useful.
The diagnostic to carry forward	Assessed by stage, a privacy programme reports itself substantially complete. Assessed by obligation, the same programme can usually evidence a fraction of what it claims. Part II is an instrument for measuring the second thing instead of the first.

II

A model that produces evidence as a ~~by-product~~.

Chapter 5 draws the distinction the market has not yet drawn clearly: between consent management, privacy management, and privacy operations. Chapter 6 sets out the two instruments this paper contributes — the Evidence Chain and the Privacy Operating Loop — in enough detail to be used without any software at all.

CHAPTER FIVE

From consent management to privacy operations

The Indian market currently sells three quite different things under one word. The confusion is commercially convenient for vendors and expensive for buyers, because the three have different scopes, different failure modes, and different prices.

5.1 Three categories, distinguished

	CONSENT MANAGEMENT	PRIVACY MANAGEMENT	PRIVACY OPERATIONS
Core question	Did we ask correctly?	Do we know and document what we do?	Are our controls operating, and can we prove it?
Primary object	The consent record	The register	The control and its evidence
Scope	Collection points: web, app, forms	The documented estate	The actual estate, continuously
Refresh cadence	Per interaction	Per review cycle	Continuous
Typical owner	Marketing or web engineering	Legal, compliance or the DPO	Jointly held; usually contested
What it proves	That consent was obtained	That the organisation has considered its obligations	That obligations were discharged, per subject, per date
Characteristic failure	Consent that never propagates	Documentation that drifts from reality	Cost and complexity, if attempted without discovery first
Necessary?	Yes	Yes	Yes — and it subsumes neither of the others

The important claim is not that privacy operations is better. It is that the three are *sequentially dependent and jointly insufficient*. Consent management without privacy management produces records nobody can interpret. Privacy management without privacy operations produces documentation nobody can verify. Privacy operations attempted without either produces an expensive instrumentation project with no legal grounding.

Most Indian enterprises have bought the first, are partway through the second, and have not started the third. Most Indian vendors sell the first and describe it as all three.

WHERE THIS LEAVES CONSENT PLATFORMS — INCLUDING OURS

A consent management platform is a genuinely useful thing, and the compliance work it does is real. But a market in which a banner product is positioned as end-to-end DPDP compliance is one in which buyers will be disappointed at exactly the wrong moment. This paper takes the view that the honest framing is layered: consent is the entry point to privacy operations, not a substitute for it. We would rather make that argument and be judged against it than inherit a claim we cannot defend in front of a regulator.

5.2 What a privacy operations platform must connect

The functional requirement is integration, not feature count. Nine capabilities have to share state; nine capabilities delivered as separate tools with separate data models reproduce the enterprise's original problem inside the compliance stack.

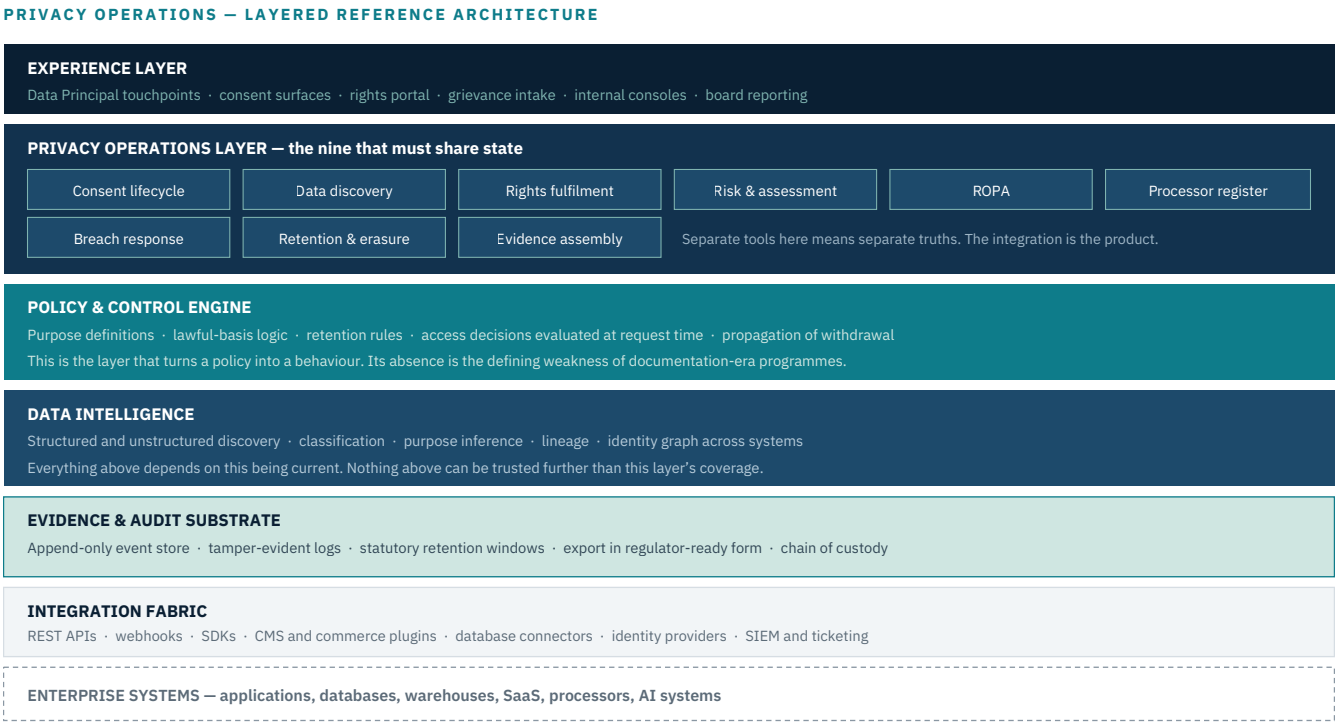

Figure 8 — Layered reference architecture for privacy operations. Vendor-neutral. A buyer can use this as an evaluation grid: ask any prospective platform which layers it actually occupies, and which it assumes you already have.

CHAPTER SIX · THE SIGNATURE FRAMEWORK

The Evidence Chain and the Privacy Operating Loop

This chapter contributes the two instruments the rest of the paper uses. Neither requires software to adopt. Both can be run on paper, badly and slowly, which is precisely the test of whether a framework is real: if it only works when you buy something, it was marketing.

6.1 Why not another lifecycle wheel

Privacy frameworks conventionally take the form of a lifecycle: discover, map, govern, assess, control, respond, monitor, prove, delete. We considered publishing one and decided against it, for a reason worth stating because it shapes everything that follows.

A lifecycle diagram describes *order*. The problem enterprises actually have is not order — most teams know roughly what sequence to work in. The problem is *completeness at the level of the individual obligation*. An enterprise can execute all eleven stages of a lifecycle and still be unable to prove anything about a named person, because completeness was measured by stage rather than by obligation. A programme reports itself 80% complete when it has done eight stages badly; it should be reporting that it has fully discharged 40 of 180 obligations and can evidence 12.

So the primary instrument here is a decomposition, not a sequence.

6.2 The Evidence Chain

Every obligation arising from the Act, the Rules, CERT-In Directions, sectoral regulation or contract resolves into five linked elements. An obligation is *managed* only when all five are present and connected. Any missing link renders the obligation acknowledged rather than discharged.

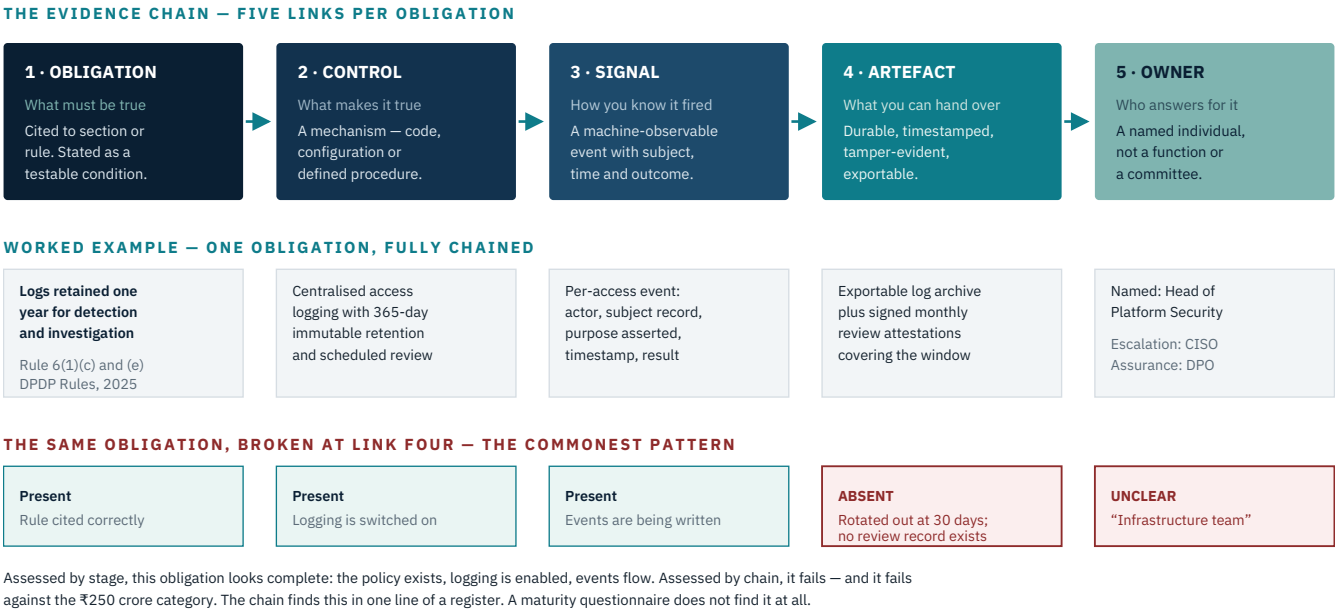

Figure 9 — The Evidence Chain, with a worked obligation and its characteristic failure. The diagnostic value is specificity: not “logging is immature” but “link four is absent on this obligation, and here is the retention setting that causes it.”

How to use it

Build a register with one row per obligation and five columns. Populate it. Nothing else. The register will be uncomfortable to look at for the first time, and that discomfort is the deliverable — it is the first accurate picture of the programme most organisations will have had.

BEST PRACTICE Three conventions make the register hold up over time. Write obligations as testable conditions rather than paraphrased law, so a reader can determine pass or fail. Name individuals rather than teams in the owner column, because shared ownership of the evidence layer is how it becomes orphaned. And treat any row whose artefact is produced by a human writing a document as provisional, because it will not survive turnover.

6.3 The Privacy Operating Loop

The chain tells you what completeness means for one obligation. The loop is the operating model that produces and maintains chains across the estate. Six phases, run continuously rather than sequentially, with the crucial property that the last three are where enterprises currently invest least.

PHASE	OBJECTIVE	WHAT RUNS	HUMAN ROLE	OUTPUT
SEE	Know what personal data exists and where, without relying on memory	Structured and unstructured discovery; classification; scanning of collection surfaces and trackers	Adjudicate ambiguous classifications; confirm system ownership	A current inventory of data locations and collection points

PHASE	OBJECTIVE	WHAT RUNS	HUMAN ROLE	OUTPUT
MAP	Attach meaning: purpose, lawful basis, flow, retention, processor	Purpose association; lineage; ROPA construction; processor and sub-processor linkage	Define purposes; resolve conflicts between stated and actual use	Records of processing that reflect reality, not intention
GOVERN	Decide the rules and make them machine-readable	Policy authoring as configuration; retention rules; access models; consent purpose definitions	Approve policy; adjudicate risk appetite; sign off assessments	Executable policy, not prose policy
EXECUTE	Make the rules act on live data and live requests	Consent enforcement and propagation; rights workflows; deletion jobs; breach response; access decisions	Approve consequential actions; handle exceptions and conflicts	Changed system behaviour, not changed documents
PROVE	Turn every execution into durable, exportable evidence	Append-only event capture; artefact generation; regulator pack assembly; attestation capture	Attest where a human judgement was required; review assembled packs	Evidence available on demand without a project
WATCH	Detect drift before an auditor or regulator does	Continuous re-scanning; new-asset and new-tracker detection; control-failure alerting; chain-completeness monitoring	Triage alerts; decide when drift becomes a change requiring reassessment	A programme whose accuracy does not decay between reviews

THE ASYMMETRY, STATED AS A DIAGNOSTIC

Score your programme out of five on each phase. In our reading of the Indian market, the modal enterprise scores respectably on SEE, MAP and GOVERN and close to zero on PROVE and WATCH — and it is PROVE and WATCH that the Rules actually test. If your scores are flat across all six, you are either unusually mature or scoring generously; the fastest way to find out is the one-record exercise at the end of Chapter 4.

6.4 Responsibility, allocated

The single most common structural defect we encounter is not a missing control. It is that the PROVE phase has no owner, because it falls between legal, engineering and internal audit and is therefore assumed by all three to belong to one of the others.

PHASE	ACCOUNTABLE	RESPONSIBLE	CONSULTED	MOST COMMON DEFECT
SEE	CIO / CTO	Platform engineering	DPO, business system owners	Coverage assumed rather than measured
MAP	DPO	Data governance	Legal, business process owners	Purposes written by legal, unknown to engineering
GOVERN	General Counsel / DPO	Legal and compliance	CISO, product, risk committee	Policy expressed in prose no system can read
EXECUTE	CIO / CTO	Application and platform teams	DPO, support operations	Consent state not reachable at the point of processing

PHASE	ACCOUNTABLE	RESPONSIBLE	CONSULTED	MOST COMMON DEFECT
PROVE	**Frequently unassigned**	—	—	Evidence reconstructed under time pressure, or not at all
WATCH	CISO	Security operations	DPO, internal audit	Monitoring covers security events but not privacy controls

BEST PRACTICE Assign PROVE explicitly, to one named person, with the authority to require artefact generation from teams that do not report to them. In organisations that have a DPO this usually belongs with the DPO; in organisations that do not, it belongs with whoever will sit in the room when a regulator asks. Leaving it implicit guarantees the outcome described in Chapter 4.

> A privacy programme is not a set of documents describing what an enterprise intends. It is a set of controls that fire, signals that record them, and artefacts that outlive the people who built them.

END OF PART II — WHAT PARTS III AND IV DO WITH THIS

Part III applies the model discipline by discipline: consent, discovery, rights, assessment, privacy engineering and PETs, agentic AI, breach and the three regulatory clocks, processor governance, Significant Data Fiduciaries, and the retention-versus-erasure conflict.

Part IV turns the model into a programme: the business case without invented numbers, six sector applications, staged roadmaps for organisations of different sizes, an honest account of where the Consiva platform sits within this architecture and where it does not, and a view of where privacy operations goes next.

Part II in summary, and a worksheet

Two instruments, both usable without software. The Evidence Chain measures completeness per obligation; the Privacy Operating Loop is the model that produces and maintains chains across the estate. The asymmetry to remember is that most programmes execute See, Map and Govern competently and Prove and Watch barely at all — and it is Prove and Watch that the Rules test.

The register below is the artefact we would want to see first in any organisation. Reproduce it, one row per obligation, and populate it honestly. The blank cells are the finding.

OBLIGATION (stated as a testable condition, cited to section or rule)	**CONTROL**	**SIGNAL**	**ARTEFACT**	**OWNER** (a person)

Conventions that make the register survive: write obligations as pass-or-fail conditions rather than paraphrased law; name individuals rather than teams in the owner column; and treat any row whose artefact is produced by a human writing a document as provisional.

III

The disciplines, each read as a chain.

Ten disciplines, from consent to the retention conflict. Each chapter states what the law actually requires, distinguishes that from what practice has come to assume, describes the control that discharges it, and names the artefact that proves it.

Where a discipline needs a capability the Indian market does not yet supply, that is said plainly rather than implied to exist.

CHAPTER SEVEN

Consent as a living control

A banner is a collection interface. Consent is a piece of state that must remain readable, current and enforceable across every system that touches the data — for as long as the data exists. Almost all of the engineering difficulty is in the second sentence.

7.1 What the law asks for

> **STATUTORY — THE SHAPE OF VALID CONSENT**
>
> Consent must be free, specific, informed, unconditional and unambiguous, given by a clear affirmative action, and signifying agreement to processing *for the specified purpose* and limited to the personal data necessary for it (s. 6(1)). It may be withdrawn at any time, and the ease of withdrawal must be comparable to the ease with which it was given (s. 6(4)–(6), Rule 3(c)(i)).
>
> The notice preceding it must be independently understandable, in clear and plain language, and must include at minimum an itemised description of the personal data and the specified purposes with a description of the goods, services or uses enabled (Rule 3(a)–(b)). It must also give the communication link and other means by which the Data Principal may withdraw consent, exercise rights, and complain to the Board (Rule 3(c)).
>
> Notice must be made available in English or any language specified in the Eighth Schedule to the Constitution, at the Data Principal's option (s. 5(3)).

Read carefully, this is a far more demanding specification than a cookie banner satisfies. Three phrases carry most of the weight. *Itemised* means a category list will not do. *Specified purpose* means purposes must be discrete enough that consent to one is not consent to another. *Comparable ease* means a one-click grant cannot be paired with an e-mail-the-support-desk withdrawal — a pattern still common on Indian websites and one of the easiest failures for a regulator to find without any investigation at all.

7.2 The consent lifecycle, and where it leaks

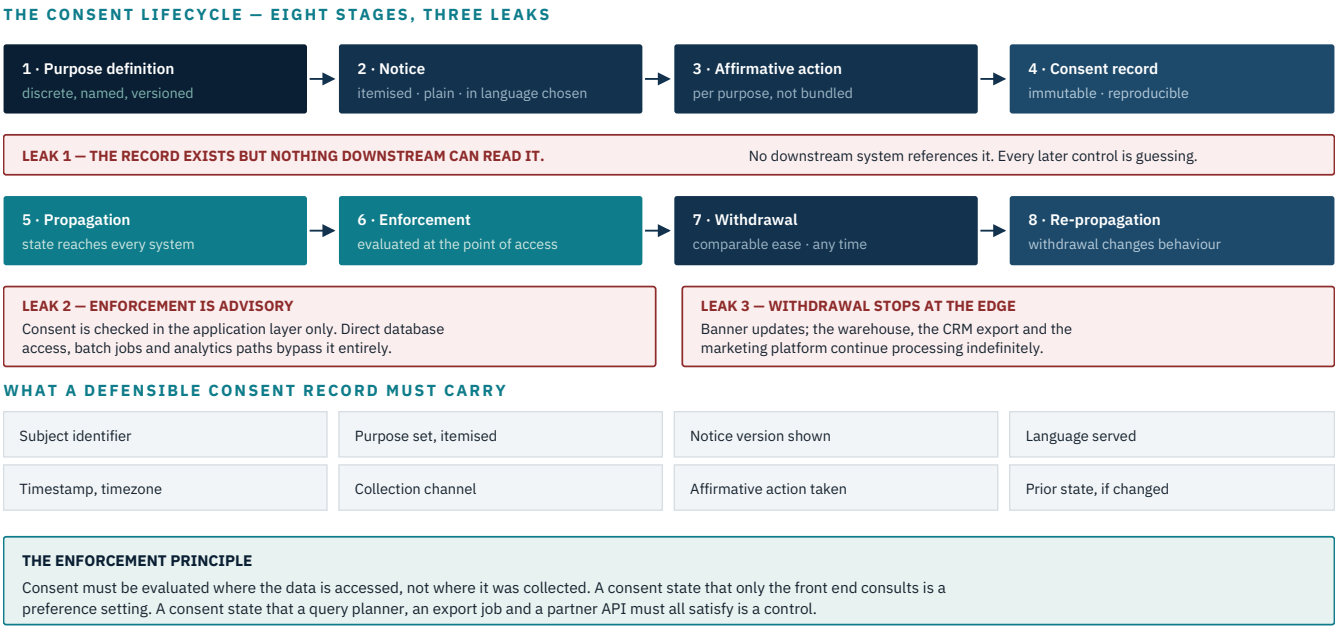

Figure 10 — The consent lifecycle and its three characteristic leaks. Stages 1–4 are what consent platforms sell. Stages 5–8 are where compliance is actually won or lost, and they are architecture problems rather than product-selection problems.

7.3 Language is an engineering requirement, not a translation task

The Eighth Schedule option is easy to under-estimate. It is not satisfied by translating a banner once. Because notice and consent are bound together, a change to a purpose definition invalidates every language variant of the notice that described it — which means the enterprise needs version parity across variants, drift detection when the English master changes, and the ability to reproduce the *specific language version a specific person saw*. An enterprise that cannot do the last of these cannot demonstrate informed consent for that person at all.

BEST PRACTICE Treat notice text as a versioned artefact under change control, bound to purpose identifiers, with the served version recorded on every consent event. Treat any manual translation workflow as a source of compliance drift and instrument it accordingly.

7.4 This discipline as chains

OBLIGATION	CONTROL	SIGNAL	ARTEFACT
Notice itemised, plain, in chosen language Rule 3	Versioned notice bound to purpose IDs; language variants under change control	Notice-served event with version and language	Reproduction of the exact notice a named person saw on a named date
Consent specific to purpose s. 6(1)	Per-purpose capture; no bundled acceptance; necessity check per purpose	Consent event per purpose with affirmative action recorded	Purpose-level consent history per subject
Withdrawal of comparable ease s. 6(4), Rule 3(c)	Withdrawal path at parity with grant path; no support-desk gate	Withdrawal event with elapsed time and channel	Comparative record of grant and withdrawal journeys

OBLIGATION	CONTROL	SIGNAL	ARTEFACT
Processing limited to consented purpose s. 6(1)	Consent state readable at the access decision; enforcement in data and API layers	Access-denied and access-permitted events carrying purpose asserted	Log demonstrating processing ceased on withdrawal, per system

CHAPTER EIGHT

Data discovery, classification and privacy intelligence

Discovery is the only capability in this paper that is not directly mandated and without which nothing else can be honestly claimed. That combination — legally optional, operationally foundational — is why it is so often deferred and so often the reason programmes stall.

8.1 An honest statement of the legal position

INTERPRETATION — WHAT IS AND IS NOT REQUIRED

Neither the Act nor the Rules expressly require an enterprise to run data discovery, maintain a data inventory, or produce Records of Processing Activities. There is no Indian equivalent of Article 30 of the GDPR. Claims that DPDP "mandates a ROPA" are, on the current text, incorrect.

What the law does require makes discovery practically unavoidable. Security safeguards must protect personal data *in the Data Fiduciary's possession or under its control* (Rule 6(1)) — which presupposes knowing what that is. Erasure obligations bite across the estate. Breach intimation requires stating the nature and extent of a breach (Rule 7). Significant Data Fiduciaries must complete an annual Data Protection Impact Assessment and audit (Rule 13(1)). None of these can be performed against an unknown estate.

BEST PRACTICE ROPA is therefore best understood as the most efficient way to evidence accountability across several obligations at once, not as a standalone duty. That framing is more defensible than the one commonly marketed, and it changes the sequencing: build the inventory because the other obligations need it, not because a rule names it.

8.2 Why discovery without context is worse than useless

A scanner that returns three hundred columns matching an e-mail-address pattern has produced a liability, not an asset: an enumerated list of exposures with no basis on which to act. Failure mode three in Chapter 4 is exactly this. Discovery becomes decision-grade only when each finding carries four additional attributes.

- **Purpose.** Which named processing purpose does this field serve? A field serving no current purpose is a retention question, not a security question.
- **Lawful basis.** Consent, or one of the legitimate uses under section 7? The answer determines which controls and which rights apply.
- **Identity linkage.** Can this record be resolved to the same individual as records elsewhere? Without an identity graph, a rights request cannot be answered completely.

- **Retention disposition.** When should this cease to exist, and does any other law prevent that? Chapter 16 takes up the second half.

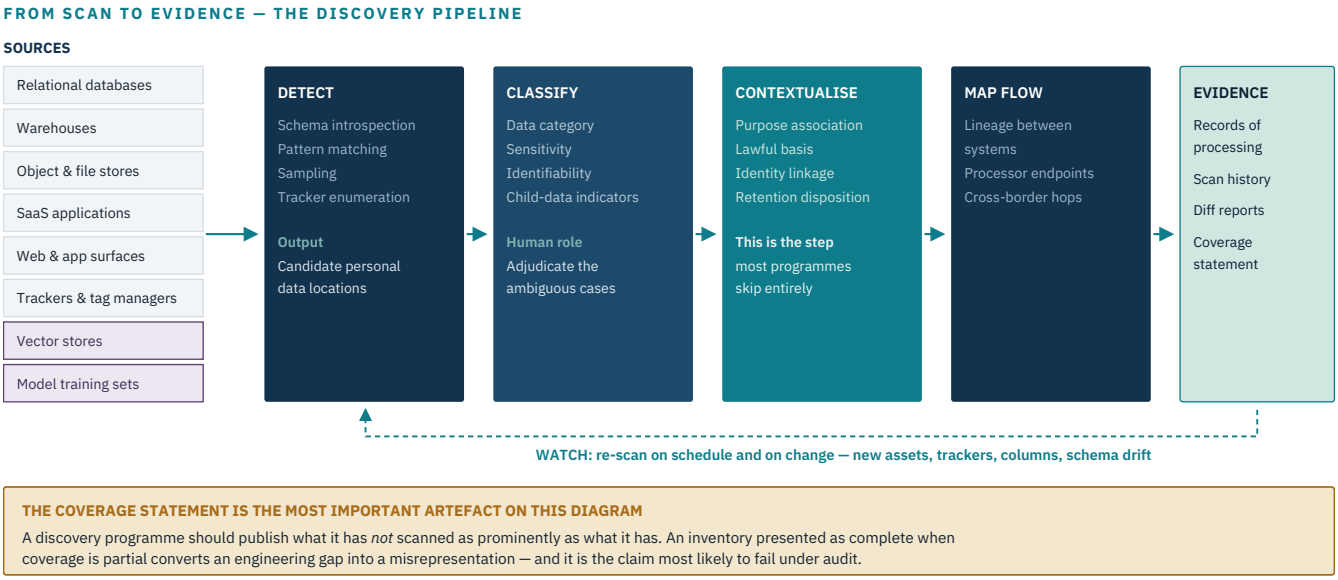

Figure 11 — The discovery pipeline, from scan to evidence. Purple sources are the ones most privacy programmes have not yet inventoried at all. The dashed return path is the difference between a one-off assessment and an operating capability.

CHAPTER NINE

Data Principal rights as an operational workflow

Rights are the only obligation in the regime that is triggered from outside the enterprise, at a time it does not choose, about a person it must first correctly identify. That makes them the sharpest available test of whether the rest of the programme is real.

9.1 The five rights, and the two numbers that matter

RIGHT	SOURCE	WHAT DISCHARGING IT ACTUALLY REQUIRES
Withdrawal of consent	s. 6(4)	A withdrawal path of comparable ease, and cessation of processing across every system that received the data. The second half is the hard half.
Access to information	s. 11	A summary of personal data being processed and the processing activities, plus identities of other Data Fiduciaries with whom it has been shared and the categories shared. This requires the processor register to be accurate.
Correction, completion, updating and erasure	s. 12	Correction propagated to derived copies; erasure propagated to every store and every processor, unless retention is required by law.

RIGHT	SOURCE	WHAT DISCHARGING IT ACTUALLY REQUIRES
Grievance redressal	s. 13, Rule 14(3)	A published mechanism, responding within a period not exceeding ninety days, supported by appropriate technical and organisational measures.
Nomination	s. 14, Rule 14(4)	A mechanism by which a Data Principal may nominate individuals to exercise her rights in the event of death or incapacity.

STATUTORY — AND A CORRECTION TO A WIDESPREAD ASSUMPTION

Ninety days is the only period the Rules attach to Data Principal requests, and it attaches specifically to the grievance redressal system (Rule 14(3)). For access, correction and erasure no period is prescribed; the governing standard is the Data Fiduciary's published terms read against its general accountability.

There is no thirty-day DPDP deadline. Thirty days is GDPR Article 12. Any vendor dashboard, internal SLA document or consultant report presenting thirty days as the DPDP requirement has imported the wrong law. **BEST PRACTICE** A shorter voluntary service level is a strong commercial and reputational position and we recommend one — provided it is described as a commitment rather than a statutory limit, and provided it is met.

Rule 14(1)–(2) also requires publishing the means of making a request and the particulars — such as a username or other identifier — needed to identify the requester under the terms of service. Rule 14(5) defines "identifier" broadly, including customer identification file numbers, application reference numbers, enrolment IDs, e-mail addresses, mobile numbers and licence numbers.

9.2 Identity verification, and the anti-abuse boundary

Identity verification is where rights workflows most often go wrong in two opposite directions at once. Verify too little and the enterprise discloses one person's data to another — a breach caused by a compliance process. Verify too much and the verification burden becomes a de facto denial of the right, which is itself a failure.

The defensible position follows the structure of Rule 14: verification should be proportionate and should rely on identifiers the enterprise already holds under its terms of service, rather than demanding new identity documents. Rate limiting and identity assurance at the point of submission are legitimate anti-abuse measures.

A BOUNDARY THIS PAPER TREATS AS ABSOLUTE

Rights fulfilment must never be capped by volume or billed per request. Sections 11 to 13 create unconditional statutory entitlements; Rule 14(3) requires technical and organisational measures to ensure the system is *effective*. A commercial mechanism that throttles or charges for genuine requests converts a product decision into statutory exposure, and it would sit in the residual penalty band alongside every other rights failure.

Anti-abuse belongs at submission — identity assurance, rate limiting against automated floods, duplicate detection. It does not belong in a pricing table. This applies to any platform, including ours.

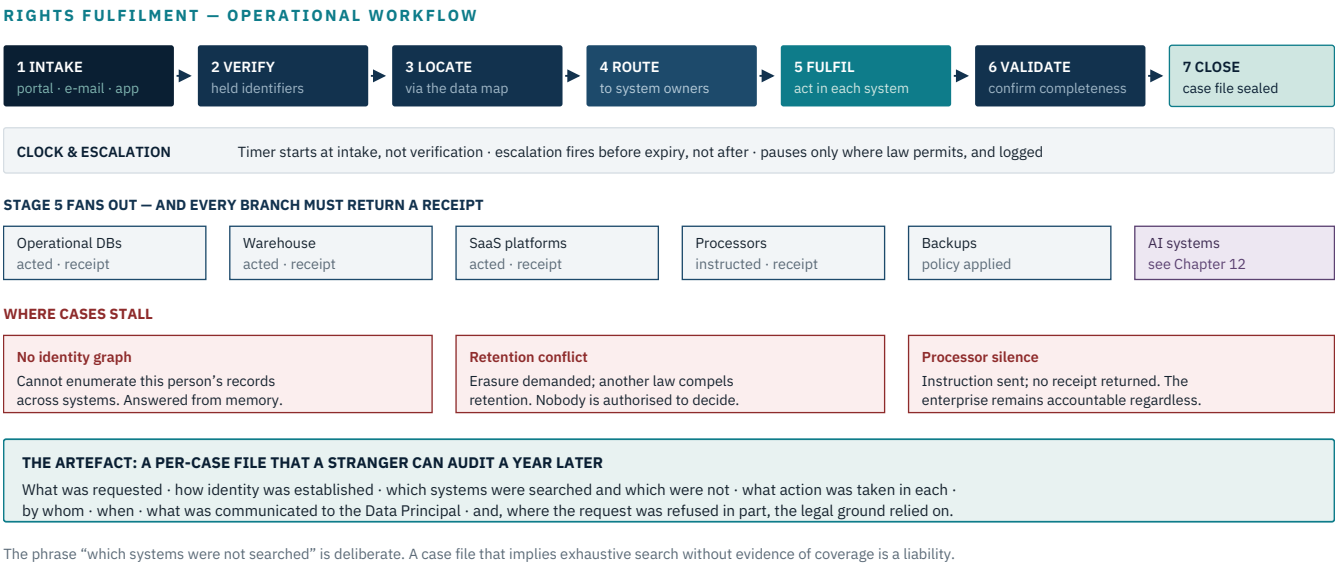

Figure 12 — Rights fulfilment as an operational workflow. The three stall patterns account for most missed service levels we have seen. Note that two of the three are failures of Chapters 8 and 14 rather than of the rights process itself.

CHAPTER TEN

Impact assessment and privacy risk management

The Data Protection Impact Assessment occupies an odd position in Indian law: narrowly mandated, widely performed, and almost universally performed as a document rather than as an assessment.

10.1 Who is actually obliged

STATUTORY — NARROWER THAN COMMONLY ASSUMED

A Data Protection Impact Assessment is a duty of **Significant Data Fiduciaries only**. Section 10(2)(c) requires it; Rule 13(1) requires an SDF to undertake a DPIA *and* an audit once in every period of twelve months from notification as such, and Rule 13(2) requires the person carrying them out to furnish the Board a report containing significant observations.

An enterprise that has not been notified as an SDF has no statutory DPIA obligation. **BEST PRACTICE** It very likely still needs the capability — because an assessment is how an organisation decides whether a new processing activity is defensible before building it, and because SDF designation is a government notification that may arrive with limited notice. But it should adopt it as risk management, not in the belief that a rule compels it.

10.2 Why the static questionnaire fails

The conventional artefact is a template completed at launch, filed, and not revisited. Its failure has four distinct causes, and only one of them is laziness.

- **The questions are generic and the system is specific.** A template broad enough to cover every processing activity is too shallow to surface the risk in any of them.

- **The respondent is not the risk-holder.** Forms are typically completed by a project manager who is describing intent, while the risk lives in implementation details known only to engineers.
- **There is no re-evaluation trigger.** The assessed system changes; nothing causes the assessment to be reconsidered. By month nine the document describes a system that no longer exists.
- **Recommendations are not tracked to closure.** Mitigations are listed, no owner is assigned, and the assessment becomes a record of risks the organisation has demonstrably been aware of and not addressed — which is materially worse than never having assessed at all.

The fourth point deserves emphasis, because it inverts an assumption. A filed assessment listing unremediated risks is evidence of knowledge. In an enforcement context, documented awareness without documented action is not a mitigating factor.

10.3 What a better assessment process looks like

RECOMMENDATION The market needs assessments that are contextual rather than templated, continuous rather than annual, and closed-loop rather than filed. Concretely, that means four properties.

Contextual questioning

Questions generated from what is already known about the system — its discovered data, its purposes, its processors, its cross-border hops — rather than asked from a fixed list. The assessment should not ask the respondent for facts the platform can already establish.

Risk expressed as harm, not score

A numeric score is easy to produce and hard to act on. A statement of the specific harm to a specific category of Data Principal, and the control that reduces it, is actionable and auditable.

Change-triggered re-assessment

The WATCH phase should fire re-assessment when the assessed system materially changes: a new data category, a new processor, a new cross-border destination, a new AI component. Annual cadence satisfies Rule 13(1); change-triggered cadence satisfies reality.

Mandatory human sign-off

Every recommendation, however generated, must be approved by an accountable human before it becomes the organisation's position. This is not a limitation to be engineered away; it is the mechanism by which accountability remains locatable.

ON AI-ASSISTED ASSESSMENT — A DELIBERATE STATEMENT OF SCOPE

There is clear and growing interest in using language models to conduct impact assessments: generating contextual questions, drafting risk narratives, proposing mitigations, mapping findings to clauses. The efficiency case is real, and we expect this to become standard practice.

This paper does not present it as a Consiva capability. Two reasons. First, we distinguish current product from roadmap consistently, and this belongs to the latter. Second, and more importantly, an AI-generated compliance recommendation that is adopted without competent human review transfers legal accountability to a system that cannot hold it. BEST PRACTICE Any enterprise evaluating AI-assisted assessment tooling should require, as a procurement condition, that the tool cannot finalise an assessment without a named human approver, and that the approval is itself an artefact. A vendor unwilling to accept that constraint is selling a liability.

CHAPTER ELEVEN

Privacy engineering and the PET capability spectrum

Privacy-enhancing technologies are the most over-claimed category in the Indian privacy market. They are also genuinely important. This chapter separates the two propositions, and states plainly which of these technologies Consiva builds — which is almost none of them.

11.1 What the law actually says about technique

STATUTORY, AND A CAUTION

Rule 6(1)(a) requires appropriate data security measures, giving as examples the securing of personal data "through encryption, obfuscation, masking or the use of virtual tokens mapped to that personal data." Rule 6(1)(d) requires reasonable measures for continued processing where confidentiality, integrity or availability is compromised, "such as by way of data-backups."

INTERPRETATION That is the extent of technique named in the Rules. **DPDP does not require differential privacy, k-anonymity, t-closeness, synthetic data, federated learning, homomorphic encryption or zero-knowledge proofs.** Any claim that it does is an interpretation being presented as law, and buyers should treat it as a commercial position rather than a compliance requirement.

There is a genuine and separate legal point nearby, and it is worth stating carefully. The Act applies to personal data — data about an identifiable individual. Data that has been irreversibly rendered non-identifiable falls outside that definition, and processing for research, archiving or statistical purposes is exempt where carried on in accordance with the Second Schedule standards (s. 17(2)(b), Rule 16). Techniques that achieve genuine non-identifiability therefore change the scope of the law rather than satisfying a requirement within it. That is a powerful and legitimate architecture, and it is also where over-claiming is most consequential: data that is reversible is still personal data, whatever the technique was called.

11.2 The capability spectrum

Four levels, ordered by the sophistication of what they protect against rather than by maturity. Most Indian enterprises need Levels 1 and 2 comprehensively and will need Level 3 for a small number of specific data uses. Level 4 becomes unavoidable the moment AI systems process personal data at all.

THE PRIVACY TECHNOLOGY SPECTRUM — FOUR LEVELS

Ordering note: Levels 3 and 4 are not "more mature" than 1 and 2. An enterprise with differential privacy but no consent propagation is misconfigured.

Figure 13 — The privacy technology spectrum. Read the amber text carefully: it marks the two places where market claims most often exceed what the law requires.

11.3 Build, integrate, partner, or recommend only

For each technology, the question a platform vendor should answer honestly is not "can we describe this" but "should we build it." Our position is set out below, and it is deliberately restrictive.

TECHNOLOGY	WHY IT MATTERS, AND WHEN	DPDP STATUS	CONSIVA POSITION	REASONING
Encryption, masking, tokenisation	Reduces breach impact and access exposure; the substance of the highest penalty band	Named in Rule 6(1)(a)	**Integrate**	Belongs in the enterprise data layer, not a compliance platform. We evidence it; we do not perform it.
Pseudonymisation	Enables analytics with reduced linkage risk; data remains personal data	Consistent with Rule 6(1)(a)	**Integrate**	Same reasoning. Value is in recording where it has been applied and to what.
k-anonymity, l-diversity, t-closeness	Enables segment-level analysis without individual identifiability	Not required	**Partner**	Specialist statistical discipline. Building it credibly is a research programme, not a feature.

TECHNOLOGY	WHY IT MATTERS, AND WHEN	DPDP STATUS	CONSIVA POSITION	REASONING
Differential privacy	Population-level statistics with a formal privacy guarantee	Not required	**Recommend only**	Requires deep expertise to parameterise safely. Wrongly configured, it provides false assurance.
Synthetic data	Development, testing and third-party sharing without real records	Not required	**Partner**	Genuine demand, particularly in BFSI. A distinct product category with distinct buyers.
Federated learning, secure computation	Model training or joint analysis without centralising data	Not required	**Recommend only**	Infrastructure-level capability. No mid-market demand we can evidence.
Homomorphic encryption	Computation on encrypted data	Not required	**Recommend only**	Performance characteristics limit realistic use. Watch, do not build.
Zero-knowledge proofs	Verification without disclosure; strong fit for age and attribute assurance	Not required	**Monitor**	Potentially relevant to verifiable parental consent under Rule 10. Worth tracking for that reason specifically.
AI privacy controls	Prevents leakage through model, retrieval and agent paths	Rule 13(3) touches algorithmic due diligence for SDFs	**Build, selectively**	The one Level 3–4 area where the governance and evidence problem — our actual domain — is the binding constraint.

WHY A PRIVACY PLATFORM SHOULD REFUSE TO BUILD MOST OF THIS

There is commercial temptation to claim the full spectrum. We think it is a mistake, for a reason that is strategic rather than modest. PETs are a deeptech discipline serving a small number of sophisticated buyers with dedicated data-science capacity. Privacy operations is a governance and evidence discipline serving a very large number of enterprises that cannot currently answer basic questions about their own estates. These are different products, different buyers and different engineering organisations.

A platform that half-builds both does neither well. Our position is that Levels 1 and 2 are where the market's unmet need and the law's actual demands coincide, and that Level 3 is best served by partnership with genuine specialists.

CHAPTER TWELVE

Privacy for the agentic enterprise

Every architectural assumption behind conventional privacy engineering — that data paths are enumerable, that access is requested by identifiable actors, that processing happens at human pace — is weakened by retrieval systems and broken by autonomous agents.

12.1 The new data path

Consider what happens when an employee asks an internal assistant a question. The prompt may contain personal data. The retrieval layer searches an embedding store built from documents whose access controls did not survive vectorisation. The model receives retrieved content as context. It may call tools, each of which can read or write to enterprise systems or third-party services. It may retain conversation memory. Each of these is a processing operation under the Act, and most enterprises have mapped none of them.

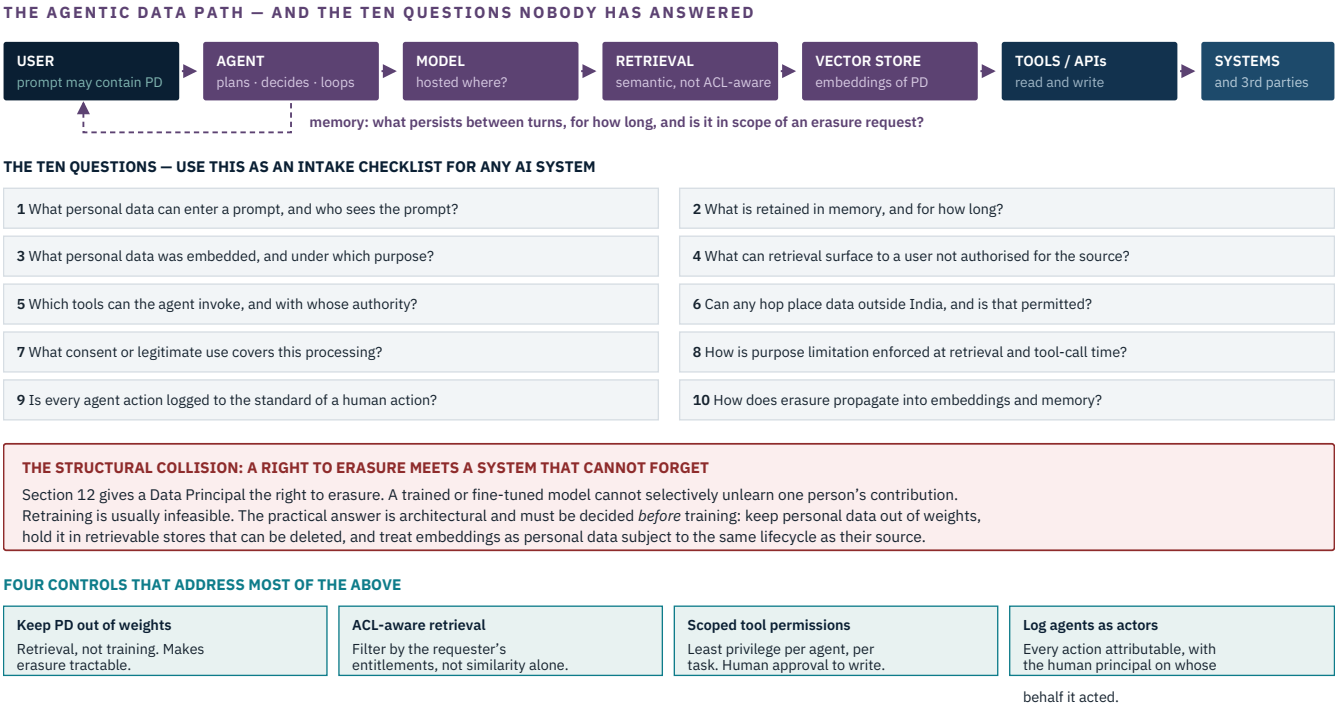

Figure 14 — The agentic data path and its governance checklist. Vendor-neutral. The ten questions are offered as an intake gate for any AI system processing personal data; in our experience an enterprise that can answer all ten has already built most of what it needs.

12.2 Where this leaves privacy operations

Three implications follow, and each is a change to the operating loop rather than a new product category.

SEE must extend to AI assets. Vector stores, training datasets, fine-tuning corpora, prompt logs and agent memory stores are locations where personal data resides. Discovery that stops at relational databases is now materially incomplete, and the gap will widen.

EXECUTE must reach retrieval and tool boundaries. Consent and purpose have to be evaluated where retrieval happens and where tools are invoked. This is the same enforcement principle as Chapter 7, applied to a faster and less predictable caller.

PROVE must treat agents as accountable actors. If an agent reads a customer record, the log entry must be of the same evidentiary quality as one recording a human doing so — including the purpose asserted and the human principal on whose behalf the agent acted. An unattributable machine action is the cleanest possible example of a broken Evidence Chain.

CHAPTER THIRTEEN

Breach response and the three regulatory clocks

Indian breach obligations are more demanding than most enterprises believe, and the reason is not the durations. It is that a single incident starts three independent obligations with different audiences, different content requirements and different start conditions — and the most urgent one is not on a countdown at all.

13.1 What Rule 7 actually says

STATUTORY — RULE 7, READ CLOSELY

Rule 7(1). On becoming aware of any personal data breach, the Data Fiduciary shall intimate each affected Data Principal, in a concise, clear and plain manner and **without delay**, through her user account or a registered mode of communication. The intimation must describe the breach including its nature, extent and timing; the consequences relevant to her; mitigation measures implemented and being implemented; the safety measures she may take; and business contact information of a person able to respond to her queries.

Rule 7(2)(a). On becoming aware, the Data Fiduciary shall intimate the Board **without delay** — a description of the breach including nature, extent, timing and location of occurrence and the likely impact.

Rule 7(2)(b). Within seventy-two hours of becoming aware, or such longer period as the Board may allow on written request: updated and detailed information; the broad facts, circumstances and reasons leading to the breach; mitigation measures implemented or proposed; findings regarding the person who caused it; remedial measures to prevent recurrence; and a report on the intimations given to affected Data Principals.

THE CONSEQUENCE OF READING THIS AS A "72-HOUR RULE"

Rule 7 creates a sequence, not a deadline. The first two obligations are *without delay* — meaning immediately upon awareness, not within three days. Seventy-two hours governs only the detailed follow-up report to the Board.

A dashboard that shows "70 hours remaining" on DPDP after a breach is detected is therefore describing the wrong obligation, and it invites a delay that is itself a violation — one sitting in the ₹200 crore band. We consider this the single most consequential misreading circulating in the Indian market, and it is one we have previously reflected in our own product interface.

13.2 Three clocks, one incident

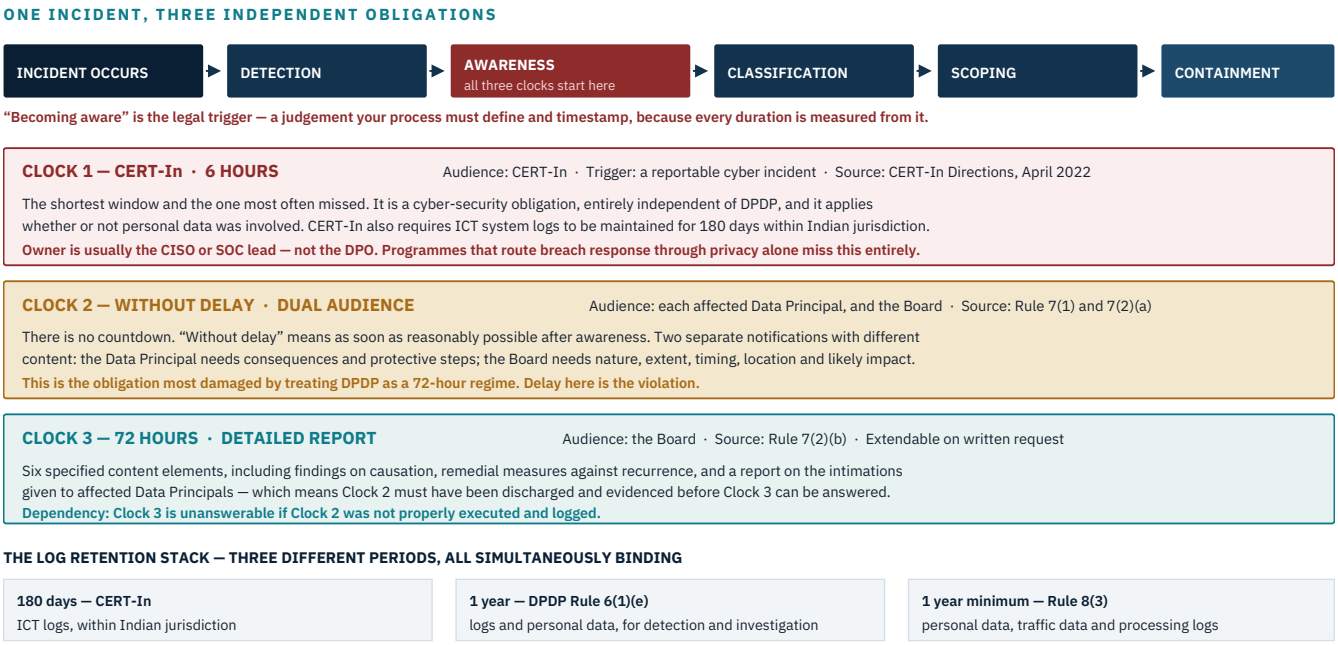

Figure 15 — Three clocks and three retention periods from a single incident. The dependency arrow matters: the seventy-two-hour report requires evidence that the "without delay" notifications were made, which is why an enterprise that misreads Clock 2 also fails Clock 3.

13.3 Defining awareness before you need to

Because every duration runs from "becoming aware," the most valuable pre-incident artefact is a written definition of what constitutes awareness in your organisation, approved in advance. Without it, the timestamp becomes a contested reconstruction after the fact — and the reconstruction will be performed by people with an interest in a later start time.

BEST PRACTICE Define awareness as the earliest point at which a named role holds information from which a reasonable person would conclude a personal data breach has likely occurred. Log it as an explicit, attributed event. Record the reasoning. This single artefact does more to protect an enterprise in a Board inquiry than any amount of subsequent documentation, because it demonstrates that the clock was taken seriously from the start rather than calculated backwards to suit the response.

CHAPTER FOURTEEN

Processor and third-party governance

An enterprise's privacy posture is bounded by that of its weakest processor, and it remains accountable for every one of them. Most Indian enterprises cannot currently produce a complete list.

14.1 Accountability does not transfer

STATUTORY

A Data Fiduciary may engage a Data Processor to process personal data on its behalf only under a valid contract (s. 8(2)). Its obligation to protect personal data in its possession or under its control extends expressly to processing undertaken on its behalf by a Data Processor (Rule 6(1)), and it must make appropriate provision in that contract for reasonable security safeguards (Rule 6(1)(f)). Rule 8(3) extends the minimum one-year retention of personal data and logs to processing undertaken on its behalf — as the illustration to Rule 8 makes explicit for a cloud service provider.

Under section 11, the right of access includes the identities of other Data Fiduciaries with whom personal data has been shared, and the categories shared. An enterprise that cannot enumerate its data-sharing relationships cannot answer an access request.

Two consequences are worth drawing out because they are commonly misunderstood in commercial negotiation. First, a contractual clause does not discharge the obligation; it is a control whose operation must still be verified and evidenced. A signed data processing agreement with a vendor who has never been assessed is a document, not a safeguard. Second, the sub-processor chain is in scope even though the enterprise never signed anything with those parties. The chain does not terminate at the counterparty.

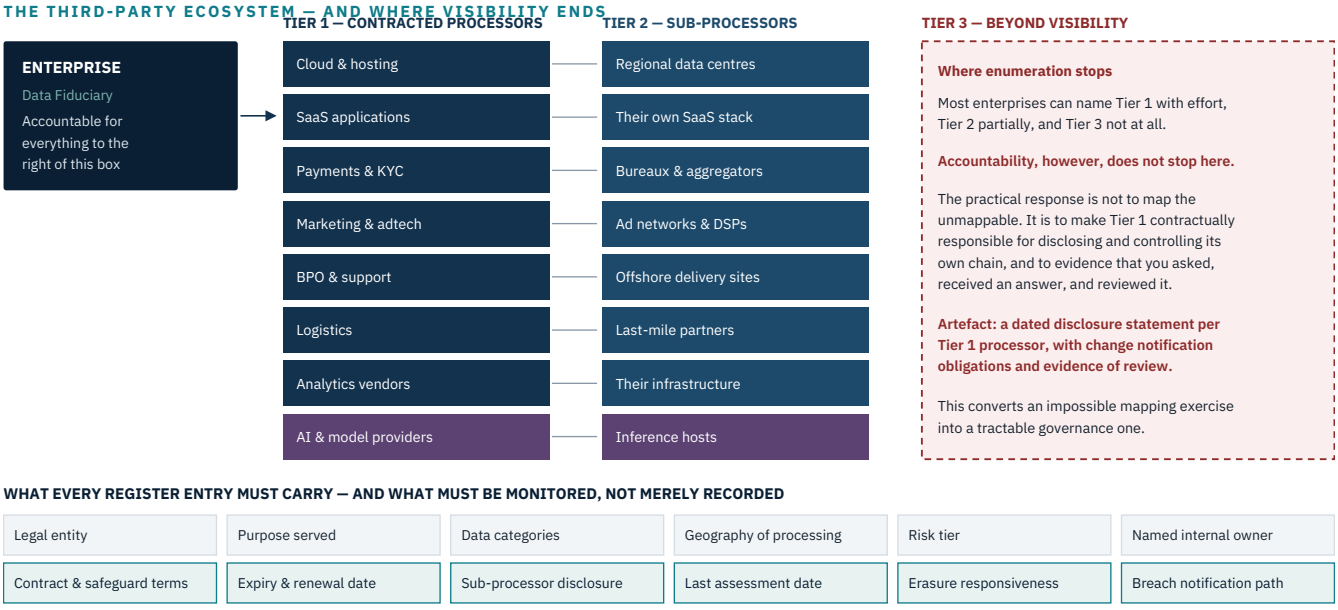

Figure 16 — The third-party ecosystem. Teal fields in the lower block are the ones that require monitoring rather than recording — they change without notice, and a register that is not watched is stale within a quarter.

BEST PRACTICE The two register fields most often missing are the ones that matter most operationally: *erasure responsiveness* (has this processor ever actually completed a deletion instruction, and how long did it take?) and *breach notification path* (who at this vendor contacts whom at us, and has that been tested?). Both are unanswerable at the moment they are needed unless established beforehand.

CHAPTER FIFTEEN

Significant Data Fiduciaries

SDF status is not a maturity tier an enterprise elects. It is a government notification that arrives, and it converts several best practices into hard annual obligations with a reporting line into the regulator.

15.1 The additional obligations

STATUTORY — SECTION 10 AND RULE 13

The Central Government may notify any Data Fiduciary or class of Data Fiduciaries as Significant, on the basis of factors including the volume and sensitivity of personal data processed, risk to the rights of Data Principals, potential effects on the sovereignty and integrity of India, risk to electoral democracy, security of the State and public order (s. 10(1)).

An SDF must appoint a **Data Protection Officer who is based in India, is responsible to the Board of Directors or similar governing body, and is the point of contact for the grievance redressal mechanism** (s. 10(2)(a)). It must appoint an independent data auditor to evaluate compliance (s. 10(2)(b)). It must undertake, **once in every period of twelve months** from notification, a Data Protection Impact Assessment *and* an audit (Rule 13(1)), and must cause the person carrying them out to furnish the Board a report containing significant observations (Rule 13(2)).

It must observe due diligence to verify that technical measures including **algorithmic software** it adopts for hosting, display, uploading, modification, publishing, transmission, storage, updating or sharing of personal data are not likely to pose a risk to the rights of Data Principals (Rule 13(3)). And it must ensure that personal data specified by the Central Government, on the recommendation of a committee it constitutes, together with the traffic data pertaining to its flow, **is not transferred outside India** (Rule 13(4)).

ON DATA LOCALISATION — A CLAIM TO MAKE CAREFULLY

India's regime does not impose general data localisation. Section 16 and Rule 15 operate on a restriction model: transfer outside India is permitted subject to requirements the Central Government may specify in respect of making personal data available to a foreign State or to entities under its control. Localisation as such appears only in Rule 13(4), applies only to Significant Data Fiduciaries, and only to categories of personal data the Government notifies on a committee's recommendation. **No such list of categories has been notified.**

India data residency remains a genuinely valuable property — for latency, for enterprise procurement, for BFSI sectoral requirements, and as forward preparation for Rule 13(4). But describing it as satisfying "DPDP data localisation requirements" overstates the current law, and a sophisticated buyer will know that.

15.2 The three structural difficulties

The DPO reporting line is a governance question, not a hiring question. Section 10(2)(a) requires the DPO to be responsible to the Board of Directors. That is a specific constitutional relationship: it implies board-level access, a degree of independence from the executives whose processing decisions are being assessed, and formal appointment. An outsourced individual whose contract sits with a vendor and whose reporting line runs to a functional head does not obviously satisfy it. **INTERPRETATION** The conservative structure is for the client's own board to appoint the individual formally, whoever supplies and supports them. Enterprises considering fractional or managed DPO arrangements should take specific legal advice on this point rather than assume the market's current offerings are compliant.

Annual DPIA plus audit is two exercises, not one. Rule 13(1) uses both words. The assessment evaluates the risk of processing; the audit evaluates observance of the Act and Rules. They have different scopes, and conflating them produces a document that satisfies neither.

Rule 13(2) creates a direct regulator channel. Significant observations go to the Board. This changes the internal politics of assessment materially: findings can no longer be managed purely internally, which means the quality of remediation tracking becomes visible to the regulator by design.

SDF GOVERNANCE ARCHITECTURE

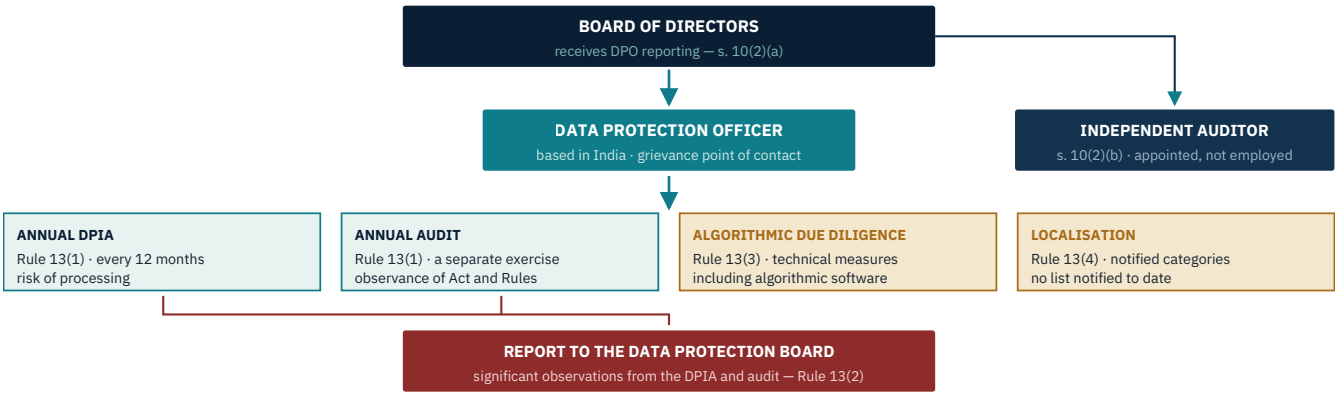

The line at the bottom is what distinguishes SDF governance from ordinary compliance: findings travel to the regulator by design, not by discovery. Remediation tracking therefore becomes an externally visible discipline rather than an internal one.

Figure 17 — SDF governance architecture. Amber blocks are the two obligations with unresolved detail: algorithmic due diligence has no published standard, and the localisation category list has not been notified.

CHAPTER SIXTEEN

Retention, erasure and the statutory conflict problem

This is the hardest problem in Indian privacy engineering, it is unsolved in the market, and it is not a consent problem, a discovery problem or a workflow problem. It is a question of which law wins — and enterprises are currently answering it by not answering it.

16.1 What the Rules require

STATUTORY — RULE 8 AND THE THIRD SCHEDULE

Rule 8(1). A Data Fiduciary of a class specified in the Third Schedule, processing personal data for the corresponding purposes specified there, shall erase such personal data — unless retention is necessary for compliance with any law for the time being in force — after the corresponding time period, where the Data Principal neither approaches the Data Fiduciary for performance of the specified purpose nor exercises her rights.

The Third Schedule specifies three classes with a **three-year** period: an e-commerce entity with not less than two crore registered users in India; an online gaming intermediary with not less than fifty lakh registered users; and a social media intermediary with not less than two crore registered users. In each case the obligation excludes data enabling access to the user account or to a virtual token usable to obtain money, goods or services.

Rule 8(2). At least **forty-eight hours** before completion of the erasure period, the Data Fiduciary shall inform the Data Principal that the data will be erased unless she logs in, initiates contact for performance of the specified purpose, or exercises her rights.

Rule 8(3). Without prejudice to the above, a Data Fiduciary shall retain personal data, associated traffic data and other processing logs for a **minimum of one year** from the date of processing for the purposes specified in the Seventh Schedule, after which it shall be erased — unless further retention is required by any other law.

TWO THINGS THIS RULE IS NOT

It is **not a general three-year retention limit**. Rule 8(1) applies only to the three Third Schedule classes at the stated user thresholds. For everyone else, the governing obligation is the general one: erase when consent is withdrawn or the specified purpose is no longer being served, unless retention is required by law.

It is **not a forty-eight-hour deletion rule**. Forty-eight hours is the minimum advance *notice* before an erasure that has already become due. We note this because the shorthand circulates and inverts the meaning.

16.2 The conflict, stated precisely

Both Rule 8(1) and section 12 carve out retention that is necessary for compliance with any law in force. That carve-out is doing enormous work, and it is where the difficulty lives. An Indian enterprise of any scale is simultaneously subject to retention mandates that are longer, differently scoped, and occasionally in direct tension with a Data Principal's erasure right.

RETENTION SOURCE	INDICATIVE PERIOD	NATURE OF THE TENSION
Prevention of Money Laundering Act and rules	Five years	Records of transactions and client identification must be preserved. An erasure request from a customer of a reporting entity cannot be honoured against these records.
RBI KYC and sectoral directions	Varies; commonly five years	Customer identification and transaction records, with additional audit-access expectations. Sector-specific and periodically revised.
Companies Act, 2013	Eight years for books of account	Financial records containing personal data of customers, employees and directors.
Income-tax law	Varies by assessment context	Records supporting returns and assessments, extendable where proceedings are open.
IRDAI, SEBI, TRAI directions	Varies by instrument	Policyholder, investor and subscriber records with their own preservation regimes.

RETENTION SOURCE	INDICATIVE PERIOD	NATURE OF THE TENSION
Litigation and investigation holds	Indefinite while live	Overrides scheduled deletion; must suspend automated erasure without silently failing.
CERT-In Directions	180 days for ICT logs	A floor rather than a ceiling, but a floor that automated log rotation frequently violates.
DPDP Rule 6(1)(e) and Rule 8(3)	One year minimum	DPDP itself compels retention — the same statute that grants erasure. This is the point most often missed.

INTERPRETATION The last row deserves emphasis because it is counter-intuitive and consequential: DPDP is internally bidirectional. Rule 6(1)(e) and Rule 8(3) require retention for a minimum period, while section 12 grants erasure. An enterprise honouring an erasure request by deleting the associated processing logs may satisfy one obligation by breaching another — and the breached one sits in the ₹250 crore security-safeguards band. Indicative periods above are illustrative and sector-specific; they require confirmation by counsel for each enterprise.

16.3 How to resolve it operationally

The wrong answers are common and both are dangerous. Deleting everything on request exposes the enterprise to sectoral non-compliance. Refusing erasure by default exposes it to a rights failure. The correct answer is a documented adjudication with a defined output.

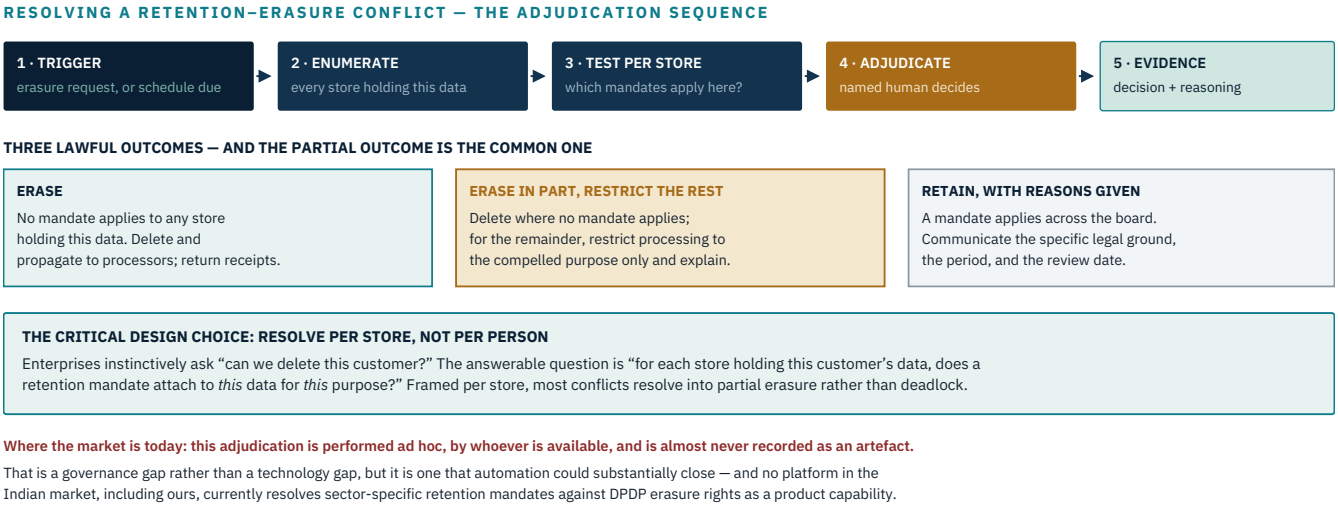

Figure 18 — The retention–erasure adjudication sequence. The final line is a deliberate statement of an unmet need rather than a product claim. We regard this as the most significant open capability gap in Indian privacy tooling.

BEST PRACTICE Three measures make this tractable without any new software. Maintain a retention mandate register mapping each store and data category to the specific statutory provision compelling retention, reviewed by counsel and dated. Give one named person standing authority to adjudicate conflicts, so that cases do not stall waiting for a decision nobody is empowered to make. And record every adjudication — the request, the stores enumerated, the mandate relied on, the outcome, the reviewer — as a first-class artefact. An enterprise that can produce a hundred such adjudications is in a far stronger position than one that has never had to make the decision, regardless of how the individual cases went.

Part III in summary

Ten disciplines, each with a characteristic gap. Used as a checklist: any row you cannot answer confidently is where Part IV's roadmap should begin for your organisation.

CH	DISCIPLINE	THE QUESTION THAT EXPOSES THE GAP	THE ARTEFACT THAT CLOSES IT
7	**Consent**	Does withdrawal change system behaviour, or only a record?	Logs showing processing ceased, per system
8	**Discovery**	What have you *not* scanned?	A published coverage statement
9	**Rights**	Which systems were searched, and which were not?	A per-case file a stranger can audit
10	**Assessment**	What triggers re-assessment when the system changes?	Change-triggered records with tracked closure
11	**Privacy engineering**	Is the technique reversible? If so, it is still personal data.	Evidence of Rule 6(1) measures in place
12	**Agentic AI**	Can you answer the ten questions for every AI system?	Dataset lineage and agent action logs
13	**Breach**	Who decides, and timestamps, "becoming aware"?	A written awareness definition, approved in advance
14	**Processors**	Has any processor ever actually completed a deletion instruction?	Dated sub-processor disclosures and erasure receipts
15	**SDF duties**	Is the DPO genuinely responsible to the Board?	Formal board appointment and annual DPIA plus audit
16	**Retention conflict**	Who is authorised to decide which law wins?	A retention mandate register and dated adjudications

THE PATTERN ACROSS ALL TEN

In every discipline the gap is at the same place. The obligation is understood, a control usually exists, and the artefact that would demonstrate it does not — because no one was accountable for producing it. That is why Chapter 6 treats assigning the Prove phase as the highest-leverage decision available, and why Part IV puts it in week one of the roadmap.

IV

From model to programme, and to the calendar.

The business case without invented numbers. Six sector applications. Staged roadmaps for organisations of four different sizes. An honest account of where the Consiva platform sits in this architecture — and where it does not. And a view of where privacy operations goes after 2027.

CHAPTER SEVENTEEN

Privacy as business infrastructure

Penalty avoidance is a weak business case. It asks a board to fund insurance against an event nobody can price, and it positions privacy as pure cost. The stronger case is that privacy operations is the control plane that determines how fast an enterprise can safely use its own data.

17.1 Three value dimensions

1 · Risk reduction

The obvious dimension, and the one to state most carefully. What good architecture reduces is not the probability of a penalty — which no one can honestly quantify absent enforcement precedent — but the *severity of a foreseeable event*: a breach whose scope can be established in hours rather than weeks, a rights request answered completely rather than approximately, an audit that reads existing artefacts rather than commissioning new ones.

2 · Operational efficiency

The most measurable dimension, and the one boards most readily accept because the baseline is observable today. Time spent per rights request. Days to assemble evidence for an audit. Engineering hours consumed by ad-hoc deletion queries. Weeks lost to security reviews in enterprise sales cycles. These are countable before and after.

3 · Trusted data and AI enablement

The dimension with the largest upside and the weakest measurement. When purpose, consent and lineage are known, data access approvals become fast decisions rather than committee negotiations. Teams stop building shadow copies to avoid governance. AI initiatives clear review because the answers to Chapter 12's ten questions already exist.

A dimension we deliberately exclude

Revenue uplift attributed to privacy. It is occasionally claimed, usually via a chain of assumptions about trust and conversion that does not survive scrutiny. It may well be real. We have no defensible basis for quantifying it, so we do not.

17.2 A worked business case, with its assumptions exposed

READ THE ASSUMPTIONS BEFORE THE NUMBERS

The figures below are **illustrative and constructed**, not measured, not benchmarked, and not drawn from customer data. They are shown because a board will ask for a shape, and because the honest way to give one is to expose the arithmetic so a CFO can substitute their own inputs. Any privacy ROI figure — including any produced by us — that does not show its assumptions should be discounted entirely.

Assumed organisation: an Indian mid-market digital business, roughly 400 employees, five million customer records, forty systems processing personal data, twenty-five processors, not an SDF. All costs in INR. Fully loaded internal cost taken at ₹1,800 per hour for senior engineering and legal time.

COST LINE	BEFORE	AFTER	ASSUMPTION BEING MADE
Rights request handling	14 hrs per request × 40 per year = 560 hrs	2 hrs per request × 40 = 80 hrs	Automation removes location and routing effort, not judgement. Volume assumed flat; in practice it rises with awareness.
Annual audit evidence assembly	320 hrs, reconstructed	60 hrs, reviewing artefacts	Artefacts already exist and are exportable. Assumes the discovery coverage is genuine.
Processor governance	180 hrs of chasing and reconciliation	50 hrs of exception handling	Register is monitored rather than periodically rebuilt.
Ad-hoc deletion and retention work	240 hrs of engineering time	40 hrs	Scheduled jobs replace manual queries. Excludes the conflict adjudication in Chapter 16, which remains manual.
Security review in enterprise sales	90 hrs, plus cycle delay	30 hrs	Standing evidence pack answers most questionnaires. Cycle-time value excluded as unquantifiable.
Internal effort	**1,390 hrs ≈ ₹25.0 lakh**	**260 hrs ≈ ₹4.7 lakh**	Gross saving ≈ ₹20.3 lakh per year, before tooling and implementation cost.

What this arithmetic does not include, and why. Platform licensing and implementation cost are excluded because they vary by scope and would make the illustration look like a quotation. Penalty avoidance is excluded because probability-weighting a ceiling with no enforcement precedent produces a number that looks rigorous and is not. Revenue effects are excluded per the note above. What remains is a defensible operational efficiency case in the low tens of lakhs annually for an organisation of this shape — meaningful, unglamorous, and arguable in front of a CFO without embarrassment.

ON THE ROI MULTIPLES CIRCULATING IN THIS MARKET

Published claims of fifteen to twenty-five times return on privacy investment are typically constructed by multiplying a penalty ceiling by an assumed annual probability, then treating the product as avoided cost. The method is not sound: the probability is invented, the ceiling is a maximum rather than an expectation, and the resulting figure is dominated entirely by an assumption the author chose. A CFO who has seen this once will discount the whole document. We would rather present twenty lakh of defensible saving than twenty crore of arithmetic.

CHAPTER EIGHTEEN

Sector applications

The obligations are uniform. What differs by sector is which obligation binds hardest, which conflict is most likely, and where the evidence gap opens first.

18.1 BFSI and fintech

The binding constraint is the retention conflict, not consent. KYC records, transaction histories and credit data sit under PMLA and RBI preservation regimes that routinely outlast any consent. An erasure request from a customer is

therefore almost never a simple delete, and the enterprise needs the per-store adjudication of Chapter 16 as a routine operational capability rather than an escalation.

Two secondary characteristics matter. Consent-based access control inside operations — call centres, branch systems, collections — is where purpose limitation is most often breached at scale, because a service agent with broad read access is processing beyond purpose on every call where they see more than they need. And the security-safeguards band applies with unusual force given the sensitivity and volume involved. *Evidence priority: access logs with purpose asserted, and dated retention-mandate adjudications.*

18.2 Healthcare and diagnostics

The binding constraint is the exemption boundary. The Fourth Schedule exempts specified classes — clinical establishments, mental health establishments, healthcare professionals, allied healthcare professionals — from sections 9(1) and 9(3) where processing is restricted to provision of health services to a child, to the extent necessary for protection of her health. That is a narrow and conditional carve-out, and the condition is the operative part: processing that extends beyond the protection of health falls outside it.

The practical failure mode is a provider treating the exemption as covering all processing of minors' data, including marketing, engagement analytics and app telemetry. It does not. *Evidence priority: purpose separation between clinical and commercial processing, demonstrable at the record level.*

18.3 E-commerce and D2C

The binding constraint is the Third Schedule, at scale. An e-commerce entity with not less than two crore registered users falls under Rule 8(1)'s three-year erasure period with the forty-eight-hour advance notice of Rule 8(2). Below that threshold the general obligation applies. Either way, the operationally difficult part is the same: the marketing stack. Pixels, tag managers, audience uploads and third-party segments create data flows that the enterprise did not consciously authorise and cannot fully enumerate — and consent withdrawal must reach all of them.

Evidence priority: tracker inventory with consent-purpose mapping, and evidence that withdrawal propagated to advertising platforms.

18.4 SaaS and technology

The binding constraint is dual-role clarity. A SaaS business is usually a Data Fiduciary for its own customer and prospect data and a Data Processor for its customers' end-user data. The obligations differ, and conflating them produces both over-compliance in one direction and gaps in the other.

The commercial dimension is more significant than the compliance one. Enterprise buyers in India are beginning to pass their own DPDP obligations down through procurement, which means a SaaS vendor's ability to produce a processor evidence pack — safeguards, sub-processor disclosure, erasure responsiveness, breach path — becomes a sales asset. *Evidence priority: a standing processor evidence pack, maintained rather than assembled per deal.*

18.5 EdTech and children's services

The binding constraint is verifiable parental consent, and it is the hardest mechanism in the Rules. Rule 10 requires appropriate technical and organisational measures to ensure verifiable parental consent before processing a child's personal data, with due diligence that the individual identifying as parent is an adult who is identifiable if required in connection with compliance with any law in force — by reference to reliable identity and age details already held, or details voluntarily provided by the individual or through a virtual token issued by an authorised entity, including via a Digital Locker service provider. The four illustrations to Rule 10 distinguish cases by whether the parent is already a registered user with verified details.

Rule 12 and the Fourth Schedule provide exemptions from sections 9(1) and 9(3) for specified classes and purposes — educational institutions for tracking and behavioural monitoring in educational activities or child safety, transport

providers for real-time location in the interests of safety, and others — each restricted to what is necessary for that purpose. Children's obligations sit in the ₹200 crore band. *Evidence priority: the verification method used per parent, per child, with the evidentiary basis retained.*

18.6 The AI-native enterprise

The binding constraint is that governance must precede training. Every other sector can retrofit controls; an organisation that has already trained models on personal data cannot retrofit the ability to unlearn. The decisions that matter — keep personal data out of weights, hold it in deletable retrieval stores, treat embeddings as personal data — are architectural and effectively irreversible.

If designated an SDF, Rule 13(3) algorithmic due diligence applies with no published standard to comply against, which means the enterprise must construct and document its own assessment methodology and be prepared to defend it. *Evidence priority: dataset lineage, the ten answers from Chapter 12, and agent action logs of human-equivalent quality.*

SECTOR	HARDEST OBLIGATION	MOST LIKELY CONFLICT	FIRST EVIDENCE GAP	PENALTY BAND MOST EXPOSED
BFSI & fintech	Purpose limitation in operations	Erasure versus PMLA / RBI retention	Access logs without purpose	₹250 cr security safeguards
Healthcare	Exemption boundary for minors	Clinical retention versus erasure	Clinical and commercial purposes not separated	₹200 cr children's data
E-commerce & D2C	Withdrawal propagation to adtech	Third Schedule erasure versus tax records	Unenumerated trackers	₹50 cr residual
SaaS & technology	Dual-role separation	Customer instruction versus own obligations	No standing processor evidence pack	₹50 cr residual
EdTech & children	Verifiable parental consent	Exemption scope versus engagement analytics	Verification basis not retained	₹200 cr children's data
AI-native	Erasure into models and embeddings	Trained weights versus section 12	Vector stores absent from discovery	₹250 cr and ₹150 cr if SDF

CHAPTER NINETEEN

A practical transformation roadmap

Roughly nine months of runway remain to 13 May 2027. That is enough time if sequenced correctly and insufficient if the hardest work is left until last — which, because discovery is slow and unglamorous, is what usually happens.

19.1 The sequencing principle

One rule governs the order of everything below. **Do the work whose duration you cannot compress first.** Discovery across a live estate, negotiating access to systems whose owners are reluctant, and instrumenting controls in code all have irreducible lead times measured in months. Publishing a notice, configuring a banner and drafting a policy can be done in days. Programmes fail because they begin with the visible fast work, demonstrate progress, and reach the slow work with no time left.

TRANSFORMATION ROADMAP — SEQUENCED BY LEAD TIME, NOT BY VISIBILITY

WORKSTREAM	0–90 DAYS · SEE AND STOP THE BLEEDING	90–180 DAYS · GOVERN AND EXECUTE	180–365 DAYS · PROVE AND WATCH
Discovery	Connect stores; first full scan; coverage statement	Extend to unstructured, SaaS, AI stores	Scheduled re-scan; drift alerting
Purpose & ROPA	Define purpose taxonomy; map top 10 systems	Complete records of processing	Auto-refresh from scans
Notice & consent	Rule 3 notice; per-purpose capture; withdrawal parity	Propagation to downstream systems	Enforcement at access points
Security safeguards	Rule 6(1) gap assessment; log retention to 1 year	Close gaps; access review cadence	Continuous control monitoring
Rights	Publish means & identifiers; intake with verification	Routing to system owners; clock & escalation	Per-case files; completeness audit
Breach	Define awareness; assign three clock owners	Templates per audience; tabletop exercise	Second exercise; scoping against data map
Processors	Enumerate Tier 1; identify contract gaps	Rule 6(1)(f) terms; sub-processor disclosure	Monitored register; erasure responsiveness
Retention	Build the retention mandate register with counsel	Deletion jobs; adjudication authority named	Deletion receipts; adjudication artefacts
AI governance	Inventory AI systems; freeze new PD training	Answer the ten questions per system	ACL-aware retrieval; agent action logging
Evidence (PROVE)	Assign the owner. Do this in week one.	Chain register populated; gaps ranked	Regulator pack assembled on demand

Amber and purple bars mark the two workstreams with the longest lead times and the least visible early progress. Start them first anyway.

The retention mandate register requires counsel time that must be booked months ahead. The AI training freeze is a decision that becomes progressively more expensive to take — every month of continued training on personal data adds to what cannot later be unlearned.

Figure 19 — Roadmap sequenced by lead time. Darker teal indicates work that must begin immediately; pale teal, work that depends on it; grey, work that only becomes possible once the first two columns exist.

19.2 Calibration by organisation size

The same sequence applies at every scale; what changes is depth and what can honestly be deferred.

	SMALL BUSINESS	MID-MARKET	ENTERPRISE	SIGNIFICANT DATA FIDUCIARY
Realistic ambition to May 2027	Notice, consent, rights intake, security basics, one honest data inventory	Full loop on the systems that matter; partial on the long tail	Full loop, with coverage gaps documented rather than hidden	Full loop plus the statutory annual cycle and Board reporting
Discovery depth	Primary database and website; manual inventory acceptable	All primary stores connected; unstructured deferred with a stated plan	Structured plus unstructured plus SaaS; AI stores in scope	As enterprise, plus algorithmic inventory under Rule 13(3)
Assessment	Lightweight review at each new processing activity	Structured assessment for higher-risk activities	Formal process with change triggers	Annual DPIA and audit, both, reported per Rule 13(2)
Governance	A named owner who also does other things	A part-time privacy lead with executive sponsorship	Dedicated function; PROVE explicitly assigned	India-based DPO responsible to the Board; independent auditor appointed
Honest deferral	Automated propagation; continuous monitoring	Unstructured discovery; PET evaluation	PETs beyond Level 2	Little. This is the point of designation.

IF YOU HAVE NINETY DAYS AND CAN DO ONLY FIVE THINGS

Assign an owner for evidence. Run one honest discovery pass and publish what it did not cover. Fix log retention to the one-year requirement of Rule 6(1)(e). Define breach awareness in writing and name the owner of each of the three clocks. Build the retention mandate register with counsel. None of these is visible to a customer, and all five are what a regulator will examine first.

CHAPTER TWENTY

Where the Consiva platform sits in this architecture

Nineteen chapters have argued a thesis that is deliberately larger than any single product. This chapter says what Consiva.ai does today, what it does not do, and which layers of the architecture an adopting enterprise must supply for itself.

20.1 Current capabilities

Everything in this section is **CONSIVA TODAY** — available in the platform at the date of publication. Nothing below is roadmap.

LAYER	CAPABILITY	EVIDENCE ARTEFACT PRODUCED
Consent lifecycle	DPDP-native consent capture with purpose-specific records; visual banner builder; automatic cookie and tracker scanning with purpose categorisation; consent in Eighth Schedule languages; consent analytics; multi-framework routing by visitor jurisdiction across DPDP, GDPR, CCPA/CPRA and LGPD	Immutable consent logs with version, purpose breakdown, channel and timestamp; exportable
Children's data	Parental consent workflow under section 9 with age-gate logic and verification workflow	Per-minor audit trail of the verification performed
Rights fulfilment	Intake and workflow for withdrawal, access, correction, erasure, grievance and nomination; e-mail-verified identity confirmation; service-level tracking with escalation before expiry	Per-case audit trail, exportable as PDF
Breach management	Incident tracking with independent regulatory windows for CERT-In and the DPDP Board; auto-drafted notification reports per regulator; escalation alerting; incident log retention	Incident record with notification drafts and dispatch history
Data intelligence	Native discovery connectors for Microsoft SQL Server, PostgreSQL and MySQL/MariaDB; schema introspection; PII pattern detection including Aadhaar patterns; mapping of columns to processing purposes; incremental diff on repeat scans	Scan history with diff reports; auto-generated records of processing
Processor governance	Vendor register with DPA status tracking, risk tiers, sub-processor chains, expiry alerting and standard contractual clause templates	Register with dated status and alert history

LAYER	CAPABILITY	EVIDENCE ARTEFACT PRODUCED
SDF support	DPO workflow management, DPIA templates, annual audit readiness, algorithmic accountability dashboards	Assessment records and audit-readiness pack
Evidence substrate	Append-only audit logging across consent events, rights actions and breach records; tamper-evident; exportable for regulator review	The substrate itself
Integration fabric	Single script-tag deployment; WordPress plugin; Shopify app; REST API; HMAC-SHA256 signed webhooks for consent, rights and breach events	Event delivery records
Deployment	India data residency for consent logs, rights records, breach reports and audit trails; cloud or on-premise deployment	Residency confirmation on records

20.2 What Consiva does not do

This section matters more than the last one. A buyer evaluating against the architecture in Chapter 5 should know precisely which layers they are still responsible for.

We are not the control engine for your data layer

Encryption, masking, tokenisation and access control live in your systems. Consiva records purpose and consent state and can evidence that controls exist; it does not encrypt your database or intercept your queries. The ₹250 crore band is discharged by your infrastructure, not by us.

Consent propagation is not automatic

Consent state is exposed via API and webhooks. Making downstream systems consult it is an integration your engineering team performs. Leak 2 and Leak 3 in Chapter 7 are not closed by installing a platform — a point we would rather make here than have discovered in month four.

Discovery coverage is bounded

Native connectors cover the three relational engines named above. Unstructured stores, object storage, SaaS-held data, vector stores and model training sets are not covered by native discovery today. Your coverage statement must say so.

We do not provide PETs

No differential privacy, k-anonymity, t-closeness, synthetic data, federated learning or homomorphic encryption — consistent with the build/partner/recommend position stated in Chapter 11. This is a deliberate strategic choice, not a gap awaiting a release.

Retention conflict adjudication is not automated

The capability described in Chapter 16 — resolving sector-specific statutory retention mandates against DPDP erasure rights — does not exist in our platform or, as far as we can establish, in any platform in this market. We regard it as the most important open gap and we say so rather than implying coverage.

We are not a registered Consent Manager

Consiva sells tooling to Data Fiduciaries. It is not a Consent Manager under the First Schedule, which requires Indian incorporation, minimum net worth of two crore rupees, board-approved governance provisions in the constitutional documents, and independent certification of an interoperable platform against standards the Board has yet to publish. Registration under Rule 4 opens in November 2026. We maintain this distinction in all communications.

CORRECTIONS TO OUR OWN PRIOR PUBLISHED MATERIAL

In the course of writing this paper we identified statements in our own marketing that do not survive the standard applied here. We list them because a paper that corrects the market while quietly fixing itself would not deserve to be believed.

The ₹250 crore ceiling was attributed to consent and rights failures; it attaches to security safeguards, and consent and rights sit in the residual band. The DPDP Board obligation was presented as a seventy-two-hour window; Rule 7 requires intimation without delay, with seventy-two hours applying only to the detailed follow-up report. A thirty-day service level was described as a DPDP requirement; it is not, and the only period in the Rules is the ninety-day grievance ceiling. India data residency was described as meeting DPDP localisation requirements; the Act contains no general localisation mandate. And superlative claims of the form "the only platform" were made without independent verification, and are withdrawn.

20.3 Competitive context, honestly drawn

We decline to publish a feature-comparison table, because we do not have verified current information about competitors' capabilities and a table constructed from marketing pages would be worthless. What can be said with reasonable confidence is structural.

What the Indian market currently supplies well. Consent capture and cookie compliance are competitively served, by global platforms with mature products and by a growing set of Indian entrants. A free or low-cost consent tier is now table stakes rather than a differentiator. Several of the strongest Indian offerings come from infrastructure incumbents — identity, e-signature and verification businesses — extending into DPDP with genuine distribution advantages.

Where the market is thin. Depth in the middle of the architecture: consent propagation into downstream systems, rights fulfilment that actually reaches the estate, evidence substrates designed for regulator export, and continuous monitoring of control operation rather than periodic assessment. And the retention conflict, which nobody has addressed.

Where genuine deeptech capability exists. A small number of Indian firms have real privacy-engineering and PET capability, serving enterprise and BFSI buyers with data-science capacity. That is a different market from mid-market privacy operations, and the honest reading is that it is complementary rather than directly competitive.

CHAPTER TWENTY-ONE

The future of privacy operations

Five shifts appear likely over the next three to five years. None is a prediction we would stake a roadmap on; each is a direction an enterprise architect should be positioned for.

Regulation becomes machine-readable

Obligations expressed as structured, versioned, machine-consumable specifications rather than prose — so that a change in a rule propagates into control configuration rather than into a project. Nothing in the current Indian framework anticipates this, but the trajectory of the DPDP Rules' own drafting, with schedules of classes, thresholds and periods, is more amenable to it than most privacy law.

Evidence generation becomes the default rather than the exception

The argument of this paper, generalised. Systems will increasingly emit compliance artefacts as an ordinary property of operation, in the way that they now emit metrics and traces. When that happens, the annual audit becomes a query rather than a project — and the enterprises that instrumented early will find the transition trivial while others rebuild.

Privacy control planes converge with data platforms

Purpose, consent state and retention disposition are properties of data, and the logical place to enforce them is the data layer itself: the catalogue, the query engine, the access proxy. Expect the boundary between privacy tooling and data infrastructure to blur, and expect that to be good for enterprises and uncomfortable for point solutions.

Agents become the primary processors of personal data

Within a few years, in many enterprises, more personal data will be read by autonomous software acting on behalf of humans than by humans directly. Every assumption about access control, purpose assertion and audit attribution has to be re-examined against that. The organisations that treated agents as accountable actors early — logging them to human-equivalent standards, scoping their tool permissions — will have built the right thing by accident of prudence.

The retention conflict gets solved, by someone

The problem in Chapter 16 is well-defined, economically significant, and currently unaddressed. It requires a maintained corpus of sectoral retention mandates, a decision engine that reasons over them per store and per purpose, and mandatory human sign-off. It will be built. We would prefer to build it; the more important claim is that any enterprise in a regulated sector should expect to be able to buy it before 2030, and should not wait for that before assigning someone to adjudicate manually.

> Compliance should not merely tell an enterprise what it must not do. It should give the enterprise the confidence to know what it can safely do.
>
> THE ARGUMENT OF THIS PAPER, RESTATED

In closing

The Indian regime arriving in May 2027 is not unusually harsh. Its penalties are high but its obligations are, by international comparison, few and clearly drafted. What makes it demanding is that it asks for something most enterprises have never been asked for: not a description of intentions, but a demonstration of operation.

That is a harder standard and a better one. An enterprise that can prove where personal data lives, why it is processed, who reached it, what consent governed it, and what happened when someone asked for it to stop, is not merely compliant. It is an organisation that can answer questions about itself — and an organisation that can answer questions about itself can move faster, not slower, than one that cannot.

The work is unglamorous, the lead times are long, and the artefacts nobody sees are the ones that matter. Nine months remain. The first step costs nothing: assign someone to own the evidence.

REFERENCES

Sources and further reading

Primary legal sources

- **The Digital Personal Data Protection Act, 2023** (Act 22 of 2023). Received Presidential assent 11 August 2023. Sections cited: 2, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 33 and the Schedule.
- **The Digital Personal Data Protection Rules, 2025**, G.S.R. 846(E), Ministry of Electronics and Information Technology, notified 13 November 2025, published in the Gazette of India, Extraordinary, Part II, Section 3, Sub-section (i). Rules cited: 1, 2, 3, 4, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, and the First, Second, Third, Fourth and Seventh Schedules.

- **Draft Digital Personal Data Protection Rules, 2025**, G.S.R. 02(E), 3 January 2025 — referenced for legislative history only.
- **CERT-In Directions** under section 70B(6) of the Information Technology Act, 2000, dated 28 April 2022, concerning information security practices, procedure, prevention, response and reporting of cyber incidents.
- **Information Technology Act, 2000** (Act 21 of 2000) — for the definitions of "computer resource" and "intermediary" incorporated by reference into the DPDP Rules.

Sectoral instruments referenced generally

Referenced in Chapter 16 as sources of retention mandates. Periods are indicative only and require confirmation by counsel for each enterprise, as several are subject to periodic revision: Prevention of Money Laundering Act, 2002 and rules made thereunder; Reserve Bank of India Master Direction on Know Your Customer and related directions; Companies Act, 2013; Income-tax Act, 1961; instruments of the Insurance Regulatory and Development Authority of India, the Securities and Exchange Board of India, and the Telecom Regulatory Authority of India.

On the Data Protection Board's constitution and staffing

The Board was constituted in law with effect from the commencement of Rules 17 to 21 on 13 November 2025. Its staffing position as at the date of publication is contested in the public record, and this paper reports it as contested rather than settled. Readers should verify current status directly with the Ministry of Electronics and Information Technology before relying on any characterisation, including ours.

A note on secondary sources

This paper deliberately relies on primary instruments for every legal proposition. Where a claim is an interpretation, it is labelled. Where a matter is contested, it is described as contested. We have not cited consultancy market-sizing figures, vendor benchmarks or survey statistics, because in this domain they are typically unsourced, unreproducible, or both — and because a paper arguing that enterprises should demand evidence would be poorly placed to ask for credence without it.

> **LEGAL NOTICE, REPEATED**
>
> This white paper is intended for informational and technology-strategy purposes and does not constitute legal advice. It is not a substitute for advice from qualified counsel on your organisation's specific circumstances. Regulatory positions described are current as at August 2026 and are subject to change; several matters discussed remain pending further notification by the Central Government or the Data Protection Board of India. Where this paper describes technologies or capabilities not expressly marked as current Consiva capabilities, no representation is made that Consiva.ai provides them. Nothing here is a claim of regulatory approval, certification or endorsement by any authority.

ABOUT

About Consiva.ai

Consiva.ai is an India-first privacy operations platform for organisations subject to the Digital Personal Data Protection Act, 2023.

The platform covers DPDP-native consent management with records in Eighth Schedule languages, Data Principal rights workflows with service-level tracking and escalation, data discovery across Microsoft SQL Server, PostgreSQL and MySQL/MariaDB with automatic generation of records of processing, incident management with independent regulatory windows for CERT-In and the Data Protection Board, a processor register with data processing agreement and sub-processor tracking, capabilities supporting Significant Data Fiduciary obligations, and an append-only evidence

substrate designed for regulator export. It supports DPDP alongside GDPR, CCPA/CPRA and LGPD, with India data residency and cloud or on-premise deployment.

Consiva.ai is built by the team at **Swaran Soft Support Solutions Private Limited**, an enterprise technology firm founded in 1999 and headquartered in Gurugram, Haryana, with offices in Dubai, Tallinn and the United States. Swaran Soft has delivered more than 1,500 projects for over 350 enterprise clients across twelve industries in twenty-five years of operation, and builds enterprise AI, application and data platforms for manufacturing, BFSI, healthcare, retail, telecom, government, education and logistics customers.

CORPORATE DETAIL	PARTICULARS
Legal entity	Swaran Soft Support Solutions Private Limited
Founded	1999 · 25+ years of operation
Registered headquarters	Tower A, Unit No. 2, 2nd Floor, The Cityscape, Sector 66, Golf Course Extension Road, Gurugram, Haryana 122102, India
International offices	Dubai, United Arab Emirates · Tallinn, Estonia · United States of America
Scale	350+ enterprise clients · 1,500+ projects delivered · 12 industries · 9 Indian languages supported
Memberships and recognition	NASSCOM member · DSCI member · CII member · DPIIT recognised · Clutch Top B2B · TechBehemoths Award 2022–2025
Certifications	ISO 9001:2015 · ISO 27001:2022 · SOC 2 audit in progress
Group ecosystem	Gignaati (AI academy, workbench and marketplace) · Copilots.in (on-premise AI infrastructure) · DCS Group (IT distribution and retail) · Digital Global AI (UAE) · Smartians AI

Corporate particulars above are as published at swaransoft.com and as confirmed by the company. ISO 27001:2022, DSCI membership and CII membership are held by the company and evidenced by the relevant certificates and membership records; readers requiring documentary confirmation may request copies. Consistent with the publishing discipline argued in this paper, no credential is stated here that the company cannot evidence on request.

Consiva.ai is not a registered Consent Manager under the First Schedule to the DPDP Rules, 2025. It provides tooling to Data Fiduciaries.

What to do with this paper

Three suggestions, in ascending order of commitment and none of them requiring a conversation with us.

Run the one-record test

The exercise at the end of Chapter 4. Pick one customer at random and try to produce the seven artefacts listed. Time it. Whatever the result, you will know more about your readiness than any maturity assessment would tell you, and it costs an afternoon.

Build the chain register

One row per obligation, five columns: obligation, control, signal, artefact, owner. Start with the twelve obligations in Chapter 2's worked mappings and the discipline chapters. Rank the gaps by penalty band. This is the artefact we would want to see first if we were your auditor.

Assign PROVE

One named person, with authority to require artefact generation from teams that do not report to them. Chapter 6 explains why this is the highest-leverage single decision available, and it can be taken this week at no cost.

Or talk to us

If you want to discuss the model, disagree with the argument, or evaluate the platform against the architecture in Chapter 5 — including the layers we have said we do not cover — we would welcome the conversation. **support@consiva.ai** · **consiva.ai**

A REQUEST

This paper makes several corrections to claims circulating in the Indian privacy market, including our own. If you believe any statement here is wrong — a misread rule, a mischaracterised obligation, a superlative we have not earned — we would genuinely like to know. Write to us. A paper arguing that enterprises should demand evidence has an obligation to accept correction on the same terms.