



16 September 2026

DPDP READINESS AND DATA PROTECTION

DPDP readiness, data protection controls and security roadmap

Optimize topline & bottomline | Hyper-personalized interactions |
Micro-segmentation at scale | Enterprise-grade compliance

Prepared by

Cllr AI

connect@cllr.ai · www.cllr.ai

Contents

01	About this report	3
02	Executive summary	5
03	Our current security posture	6
04	Scope and basis of this report	8
05	Architecture	9
06	Data flow and personal data lifecycle	11
07	Data protection controls in detail	13
08	Retention and deletion	16
09	Data residency	18
10	Third parties in our processing chain	19
11	DPDP position and control mapping	21
12	What we are still building, and why in this order	24
13	Security roadmap	25
14	Our position on certification and independent testing	26
15	Frequently asked questions	28
16	Evidence we can provide	31
17	In summary	32

Cllr AI operates an AI-powered revenue intervention layer across the platform: email, WhatsApp, SMS, voice calling, and CRM integrations. When you run campaigns on our platform, we process personal data belonging to the people you are contacting, and we do it on your instructions and on your behalf. That places a set of obligations on us under India's Digital Personal Data Protection Act, 2023, and it places a reasonable expectation on us to tell you exactly how we handle that data.

This report sets out how our platform protects personal data today, how our controls map to the requirements of the DPDP Act and the DPDP Rules, where our infrastructure is located, who else is in the processing chain, and what we are still building. We have written it to be read by your security, privacy, legal and procurement teams, and we have tried to make it useful to all four.

What is already in place. Twenty-two controls are implemented and operating in production today, including encryption of personal data at rest and in transit, PII masking under which sensitive data is never written to our systems, role-based access control, private network isolation, audit logging, rate limiting, a thirty-day default data lifecycle, and deletion on request that propagates through to backups, snapshots, replicas and the third-party services in our processing chain. Every core processing service runs on India-hosted infrastructure or an India endpoint, including the language-model processing that handles conversation content. The next section lists each control.

Two points on assurance, because they are usually the first questions asked.

First, no organisation can be "DPDP certified", because no such certification exists. What the Act requires is lawful processing, reasonable security safeguards, disciplined retention and erasure, and demonstrable accountability. Those are the tests this report addresses, control by control.

Second, we do not currently hold ISO/IEC 27001 certification, a SOC 2 attestation, or an independent penetration-testing report. None of the three is a legal requirement under the DPDP Act; they are market assurance mechanisms that customers may require in procurement, and we treat them as such. In their place we offer documented controls, an evidence pack your team can review, completed security questionnaires, and a walkthrough with the engineers who built the platform. An independent penetration test is the next step on our roadmap. We prefer to set this out at the start rather than leave it to emerge later in a review.

About this report

01

Status of this document. This report is confidential and is provided for the internal use of the recipient organisation and its advisers. It describes the state of the Cllr AI platform as at the date on the cover. It is a description of our controls and our current plans, not a warranty, representation or certification, and it is not an independent audit or a legal opinion. Statements about planned work reflect our intentions at the date of issue and may be re-sequenced; where you need a commitment to be binding, including retention periods, deletion windows, notification timelines or security obligations, that commitment belongs in the agreement and data-processing agreement between us, and we are glad to negotiate it there. Nothing in this document varies an existing agreement.

Each control in this report carries one of four labels, so that it is always clear what is running in production today and what is still in progress.

LABEL	MEANING
In place	The control is implemented and operating in production today.
In place, configurable	The capability is implemented, and its behaviour depends on how your deployment is configured. We will confirm the setting for your account.
Being documented	The control or practice exists in our engineering operations, and we are formalising the written specification and evidence so we can share it with you.
Planned	Not yet in place. Committed to our roadmap, with the phase indicated.

Where something is still open, we have said so, so that you can rely on the rest of the document.

A note on naming. This report describes our infrastructure and third-party service categories at a functional level rather than naming individual cloud providers, service identifiers or vendors. This keeps internal architecture detail out of a document that may be circulated widely. Your privacy, security or vendor-risk team is very welcome to the full register naming every provider, and we will share it on request, under NDA if you prefer.

What our platform processes. Depending on your use case, our platform may process phone numbers, email addresses, customer information, contact and caller information, campaign information, message and email content, interaction and call metadata, call recordings, call transcripts, CRM records synchronised through our integrations, records of AI interactions, personally identifiable information and operational data. Not every category applies to every deployment, and we will confirm the categories in scope for yours.

How we protect it. Personal data is encrypted at rest in our database and object storage layers, and data moving through our APIs, including anything carrying personal data, is encrypted in transit. We support PII masking, and where masking is enabled the sensitive personal data is never written to our systems at all. Access is governed by role-based access control with separate manager and admin levels, sitting behind a managed identity layer, with two-factor authentication available. Our production backend, internal APIs and internal systems run as containerised workloads inside a private cloud network and are not exposed to the public internet. We rate-limit our APIs, use one-time-use API credentials where applicable, maintain audit logs of user and operational activity, and run open-standard telemetry instrumentation across the platform. Applicable customer data is deleted on a default rolling one-month cycle, and can be deleted on request.

Where it is processed. Our core data-processing services are configured on India-hosted infrastructure and India endpoints. Object storage and our database are India-hosted. Language-model processing runs in an India region (Mumbai). Speech-to-text and speech synthesis run on India endpoints. For an organisation with Indian data-residency expectations, whether from internal policy, sectoral regulation or customer commitments, this is a deliberate design decision rather than a coincidence of configuration, and we made it early.

What we are still building. Three things. We are formalising documentation that describes practices our engineering team already follows, including incident-response and personal-data-breach procedures, a retention and deletion schedule, and a data-principal request workflow. We are completing a verification pass on specific items, including backup geography and restoration testing, privileged administrative access controls, and the residency and retention terms of every third party in our processing chain. And we are commissioning independent technical assurance, starting with a third-party penetration test. The roadmap later in this report sets out the sequence, and the section before it explains how we have prioritised the work.

In summary. Our technical controls are strong, and in several respects, particularly masking that prevents sensitive data from being stored at all, a thirty-day default data lifecycle and India-region AI processing, they are ahead of what is common among platforms in our category. The work that remains is concentrated in the documentation, verification and independent validation layer around those controls rather than in the controls themselves, which is why it is measured in months and why each item in this report carries a phase and an owner.

Our current security posture

03

In place today

CONTROL	STATUS
Personal data encrypted at rest in database and object storage	In place
Encryption in transit for API communications carrying personal data	In place
Frontend and application controls protecting personal data from unnecessary exposure	In place
Role-based access control, with separate manager and admin levels	In place
Managed identity and authentication layer	In place
Two-factor authentication	In place, available per seat
PII masking, with masked data not stored	In place, configurable
Audit logging of user and operational activity	In place
Audit-log retention, configurable to your requirements	In place, configurable
Default rolling one-month retention and deletion of applicable customer data	In place, configurable
Deletion of your data on request from our primary systems	In place
Deletion propagating to snapshots, backups, object versions, replicas and disaster-recovery copies	In place
Deletion propagating to the third-party services in our processing chain	In place
Containerised workloads in a private cloud network, not publicly exposed	In place
API rate limiting	In place
One-time-use API credentials where applicable	In place
Backups and snapshots of our database and object-storage infrastructure	In place
Open-standard telemetry and observability instrumentation	In place
India-hosted object storage and database	In place
Language-model processing in an India region (Mumbai)	In place
Speech-to-text on an India endpoint	In place
Speech synthesis on an India endpoint	In place

Being documented

These are practices and configurations that exist in our operations. What we are producing is the written specification and the evidence, so that your team can review them rather than take our word for it.

ITEM	TARGET
Default audit-log retention period, log access controls and log integrity	Phase 1
Backup frequency, retention, encryption, geography, recovery objectives	Phase 1

ITEM	TARGET
Encryption key management and rotation	Phase 1
Secrets-management approach and credential rotation	Phase 1
Privileged administrative access model and its authentication requirements	Phase 1
Residency, retention and training-use terms for every third party in the processing chain	Phase 1
Secure development practices and vulnerability management	Phase 3

Planned

ITEM	PHASE
Formal incident-response plan	Phase 1
Personal-data-breach response procedure	Phase 1
Data-principal request workflow with defined timelines	Phase 1
Published privacy notice and grievance mechanism	Phase 1
Executed data-processing agreements with all third parties in the chain	Phase 1
Multi-factor authentication enforced on all privileged production access	Phase 1
Independent third-party penetration test, remediation and retest	Phase 2
Security alerting and detection on our telemetry	Phase 3
Periodic access reviews	Phase 3
Disaster-recovery testing and business-continuity documentation	Phase 3
ISO/IEC 27001 or SOC 2, subject to customer requirements	Phase 4

Scope and basis of this report


04

What this report covers. The Cllr AI platform as we operate it today: our India-hosted storage and database layer, our AI and voice processing chain, and the identity, network, API, logging and retention controls around them. It also covers our position under the DPDP Act and Rules, the third-party service categories in our processing chain, and our security roadmap.

What it does not cover. Your own systems and their configuration, including any CRM we integrate with, the telephony, messaging and email carrier layers beyond our platform boundary, our corporate IT and endpoint estate, and physical security at our providers' facilities. Where a control depends on your deployment configuration, we say so rather than generalising.

What it is. Our own account of our controls, prepared for your review. It is not an independent audit, it is not a certification, and it is not legal advice. Where a conclusion turns on legal interpretation, we flag it, and we would expect your counsel and ours to align on the reading rather than either of us asserting one.

On the regulatory position. We describe the commencement position of the DPDP Act and Rules below, based on the published notifications. The rule-by-rule commencement schedule of the DPDP Rules, 2025 should be confirmed against the Gazette text by counsel before either of us relies on a specific date in a contractual commitment. We have set out the dates we can source, and marked what we cannot.



Controls applied across this architecture


CONTROL	WHERE IT APPLIES
Encryption at rest	Managed database, object storage
Encryption in transit	API communications
Frontend and application PII protection	Application layer
PII masking, with masked data not stored	Application and storage layers
Role-based access control	Authorization layer
Two-factor authentication	Identity layer
Audit logging with configurable retention	Application and data layers
Rate limiting	API layer
One-time-use API credentials where applicable	API layer
Telemetry and observability	Platform-wide
Backups and snapshots	Managed database, object storage
Default one-month rolling data lifecycle	Applicable customer data
Deletion on request	Primary data stores

Three properties of this design worth drawing out.

Our data layer is not reachable from the public internet. Backend services run as containerised workloads inside a private cloud network, so the practical attack surface concentrates on our application and its credentials rather than on exposed data stores.

Every identified core processing service is configured on India-hosted infrastructure or an India endpoint. The default path for storage, language-model inference, transcription and speech synthesis is domestic, including the AI processing stage where transcript content is handled.

Masking operates at the architecture level. Where masking is enabled, sensitive personal data is not written to our systems at all, rather than being stored and hidden from display. That is a materially different control, and its effects run through retention, deletion and incident impact.



Retention and deletion apply to every store, including backups

Collection. Personal data reaches our platform from you. Campaign and contact data is loaded by your users or synchronised from your CRM; interaction data is generated as conversations run across email, WhatsApp, SMS and voice. In most engagements you are the data fiduciary, because you determine the purpose and you hold the relationship with the individual. We process on your instructions. We will address the allocation of notice and consent responsibilities explicitly in our data-processing agreement rather than leaving it to inference.

Transmission. API communications carrying personal data are encrypted in transit. **In place.**

Processing. Message, email and conversation content is processed by a language-model service in an India region (Mumbai). On the voice channel, call audio is additionally transcribed through a speech-to-text service on an India endpoint, and synthesised speech is generated through a service on an India endpoint. Where masking is enabled, the sensitive personal data is not persisted.

Storage. Structured data in our India-hosted managed database, with personal data encrypted at rest. Objects, including call recordings and message attachments where applicable, in India-hosted object storage, encrypted at rest. **In place.**

Access. Through the application only, behind our managed identity layer and role-based access control, with two-factor authentication available. Direct access to underlying systems is restricted by our private network architecture.

Retention. A default rolling one-month, approximately thirty-day, deletion cycle for applicable customer data, configurable to your requirements. Audit-log retention is configured separately and is not governed by the one-month primary-data cycle.

Deletion. Your data can be deleted on request, and deletion propagates beyond the primary stores to snapshots, backups, object versions, replicas and disaster-recovery copies, and to the third-party services in our processing chain. **In place.**

How deletion reaches every copy

A deletion request has to reach more than the primary record, so we set out each location it covers rather than wait to be asked.

LOCATION	STATUS
Database snapshots	In place
Object-storage backups	In place
Object versions	In place
Replicas and disaster-recovery copies	In place
Audit logs, where they contain personal data	In place
Telemetry data, where it contains personal data	In place
Third-party services in the processing chain	In place

Deletion therefore covers the full set of copies rather than the primary record alone, which is the standard we think an erasure obligation should be held to. We are documenting the propagation window for each store so that it can be stated as a defined period in your data-processing agreement.

Encryption

At rest. Personal data held in our database and object storage is encrypted at rest. **In place.**

In transit. Data moving through our APIs, including communications carrying personal data, is encrypted in transit. **In place.** We are documenting our transport configuration in detail rather than quoting protocol versions here that we have not re-verified for this document.

Being documented, Phase 1: encryption key management and rotation, backup encryption configuration, and encryption of internal service-to-service traffic within our private network.

PII masking and non-persistence

Where PII masking is enabled, the sensitive personal data is not saved in our system. **In place, configurable per deployment.**

This is the control we would point to first, because it changes the shape of the risk rather than mitigating it:

- Data that is never written cannot be over-retained, mis-deleted or exposed.
- The persistent personal data footprint across our platform shrinks.
- The impact of a security incident is reduced, because a store that never held the sensitive values yields less. It reduces impact; it does not eliminate the possibility of an incident.
- It supports the data-minimisation principle under the DPDP Act directly, at the architecture level.

Masking is configurable, and whether it is enabled depends on your use case. We will confirm the setting for your deployment and record it, and we are also documenting whether masking is applied before content reaches our AI processing services, which strengthens the position further where it is.

Frontend and application protection

We implement frontend and application controls intended to protect personal data and prevent unnecessary exposure, alongside the storage and transport controls above. **In place.** We describe these as application controls rather than as client-side cryptography, because that is what they are, and we would rather be precise than impressive.

Identity and access

Role-based access control. Implemented, with separate access levels including manager-level and admin-level, so your users receive access according to their assigned role. **In place.**

Authentication. Handled through a managed identity and authentication layer rather than bespoke credential handling. **In place.** We are documenting the session-expiry, token-lifecycle, authentication-policy and brute-force-protection configuration so your team can review the settings rather than the architecture alone.

Two-factor authentication. Available as a security capability, provided per seat. **In place, available.** It is not enforced universally by default, and we will not describe it as though it were. We will enable it for your organisation, and we recommend it.

Privileged administrative access. Our own administrative access to production is the shortest path to data in an architecture that is otherwise privately networked, so we treat it as a priority rather than a footnote. We are completing a verification pass on our privileged access model and will enforce multi-factor authentication on all privileged production access as a Phase 1 commitment. **Planned, Phase 1.**

Periodic access reviews. Not yet a formal recurring process. **Planned, Phase 3.**

Network isolation

Our production backend, internal APIs and internal systems run as containerised workloads within a private cloud network and are not exposed to the public internet. Our data stores are reachable through our application, not directly from outside. **In place.**

We do not claim that unauthorised access is impossible, because that claim is not available to any platform. Private systems remain exposed to credential compromise, application vulnerabilities, insider risk, cloud misconfiguration and supply-chain issues. What a private architecture does is remove an entire category of exposure and concentrate the remaining risk into layers we can control, monitor and test. Independent testing of that configuration is our Phase 2 commitment.

Bulk access. Direct access to underlying production systems is restricted through private infrastructure and application-layer access controls, which reduces opportunities for unauthorised bulk access. Dedicated detection of abnormal bulk download or export activity is not yet in place, and given that our platform holds call recordings and conversation content we consider it a worthwhile addition. **Planned, Phase 3.**

API security

Credentials. We use one-time-use API keys and tokens where applicable, which limits the value of an intercepted, logged or accidentally disclosed credential. **In place.** Rotation intervals, expiry and our secrets-management approach are being documented in Phase 1.

Rate limiting. In place. It mitigates API abuse, automated attacks, excessive request volume, resource exhaustion, and credential-abuse patterns such as brute-force and enumeration attempts. It does not address authorisation flaws, business-logic abuse or a legitimate credential used improperly, which is why it is one control among several rather than the answer on its own.

Multi-tenant isolation. We consider the most consequential question for a platform like ours to be whether one customer could ever reach another customer's recordings or transcripts. Our role-based access control and private networking are relevant but are not by themselves a demonstration. We have made tenant-isolation testing an explicit scope item for our Phase 2 independent penetration test.

Audit logging

We maintain audit logs of operational and user activity, including who made a change, when data was inserted, the associated timestamps, campaign start and pause actions, and changes made by users. **In place.**

Retention is configurable to your organisation's requirements. Importantly, and this is a point worth stating clearly because it is easy to misread: audit-log retention is independent of our thirty-day primary-data cycle. Our logs are not deleted after thirty days simply because customer data is.

Being documented, Phase 1: our default log retention period, whether logs contain personal data, log access controls, log integrity protection, and extending logging to cover read access to recordings and transcripts as well as changes. For sensitive call content, who listened is at least as important as who edited, and we want our logging to reflect that.

Monitoring and observability

We run open-standard telemetry and observability instrumentation across the platform, which gives us tracing and metrics and the ability to investigate anomalies. **In place.**

We do not present it as a security operations centre, a SIEM or an intrusion-detection system, because it is not one. Observability is the substrate; security detection is a distinct capability built on top of it. Our Phase 3 commitment is a focused set of detections on telemetry we already collect: failed-authentication spikes, privileged-role changes, unusual volumes of recording or transcript access, first-time administrative access from a new location, and API error-rate anomalies indicating enumeration, each with a named on-call recipient. **Planned, Phase 3.**

Retention

We apply a default rolling retention and deletion cycle of one month, approximately thirty days, to applicable customer data. Your requirements can adjust it. **In place, configurable.**

We think a short default is the right posture. The DPDP Act's storage-limitation principle asks that personal data not be kept beyond the purpose it was collected for, and a thirty-day default is easier to justify than a long retention period with exceptions. For call recordings, transcripts and message content in particular, it materially limits the volume of sensitive content standing at any moment.

DATA	RETENTION
Primary customer data	One-month rolling cycle. In place, configurable.
Audit logs	Configured separately to your requirements. In place; default period being documented.
Backups and snapshots	Being documented, Phase 1
Object versions	Being documented, Phase 1
Telemetry data	Being documented, Phase 1
Data held by third parties in the processing chain	Being documented, Phase 1

We are producing a written retention schedule covering every data category, its store, its retention period, its deletion mechanism and its owner. **Phase 1.**

Deletion

Your data can be deleted on request from our primary data stores. **In place.** Combined with the thirty-day default cycle, most data is deleted as a matter of course rather than by exception.

Deletion is not limited to the primary record: it propagates to snapshots, backups, object versions, replicas and disaster-recovery copies, and to the third-party services in our processing chain. Residual copies clear as that propagation completes rather than persisting indefinitely.

Phase 1 deliverable. A written deletion SLA stating the period for primary deletion and the window within which each class of residual copy clears, so the commitment sits in your agreement as a defined number. We will give you verified figures rather than estimating them here.

Backups and recovery

We maintain snapshots and backups of our database and object-storage infrastructure. **In place.**

The following are being documented in Phase 1, and we are giving you the list rather than approximate figures: snapshot frequency, snapshot retention, backup encryption configuration, backup geography, restoration-testing cadence, recovery time objective and recovery point objective, and deletion propagation into backups.

Two of those matter more than the others and we are treating them accordingly. Backup geography, because our residency position rests on India-hosted infrastructure and a backup replicating elsewhere would be a gap in precisely the place it is least visible. And restoration testing, because a backup that has never been restored is an assumption rather than a control. Both are Phase 1.

Our core data-processing services are configured on India-hosted infrastructure and India endpoints.

FUNCTION	LOCATION
Object storage	India hosted
Managed database	India hosted
Language-model processing	India region (Mumbai)
Speech-to-text	India endpoint
Speech synthesis	India endpoint

In place.

We consider this a genuine differentiator and it was a deliberate choice. Many platforms in our category route inference, transcription or synthesis through non-Indian regions by default, often without surfacing it. We placed language-model processing in an India region and use India endpoints for both transcription and synthesis, which means the default path for the most sensitive stage of the pipeline, where spoken content is turned into text and interpreted, is domestic.

The claim we are not making. We do not claim that no data ever leaves India. That statement would require verified contractual coverage of every third-party diagnostic log, backup location, support-access path and downstream subprocessor, and we are still completing that verification. An endpoint or region setting governs where the processing request is served. It does not by itself govern where a provider stores diagnostic logs, where a support engineer sits when they access a ticket, or where a provider replicates its own backups.

LAYER	POSITION
Configured processing location	India-hosted infrastructure and India endpoints. In place.
Contractual coverage of every associated log, backup, support operation and downstream subprocessor	Being documented, Phase 1

Phase 1 commitment. Written confirmation from each third party in our processing chain covering the residency terms applicable to our configuration: where diagnostic logs are stored, where support access originates, whether backups leave the region, and their own subprocessor list. Once we have that, we will give you a definitive residency statement. Until then, our residency position is stated at the architecture and configuration level, which is accurate and, we would argue, already strong.

On cross-border transfer. The DPDP Act permits transfer of personal data outside India except to territories the Central Government restricts, and sectoral regulators may impose stricter localisation requirements. Our India-centred architecture means transfer exposure is low by design. Where you operate in a regulated sector with specific localisation obligations, we should work through them together against the verified position rather than against the architecture alone.

Third parties in our processing chain

10

We use six third-party services in our processing chain. They are described here by function. We maintain a register naming each provider, and we will share it with your privacy or vendor-risk team, under NDA where you prefer.

FUNCTION	DATA POTENTIALLY PROCESSED	LOCATION	RETENTION	PROCESSING AGREEMENT	VERIFICATION IN PROGRESS
Object storage	Recordings, files and objects, including personal data as applicable; encrypted at rest	India hosted	Follows our one-month rolling cycle for applicable customer data; backup and version retention being documented	Being executed, Phase 1	Region evidence, encryption and versioning configuration, replication destinations, backup geography
Managed database	Phone numbers, email addresses, customer and contact information, campaign data, message content, interaction and call metadata, transcripts and synchronised CRM records as applicable; personal data encrypted at rest	India hosted	One-month rolling cycle; snapshot retention being documented	Being executed, Phase 1	Region evidence, encryption configuration, snapshot schedule and retention, network access configuration
Language-model service	Prompts and completions derived from message, email and transcript content and conversation context, including personal data as applicable	India region (Mumbai)	Being documented; provider retention behaviour to be stated only once the applicable terms are verified	Being verified, Phase 1	Applicable service terms, invocation-logging configuration and log location, training-use terms, guardrail configuration
Speech-to-text	Raw call audio and resulting transcripts, including personal data as spoken by callers	India endpoint	Being documented	Being executed, Phase 1	Residency terms for logs and backups, retention or zero-retention configuration, training-use exclusion, support-access model
Speech synthesis	Text submitted for synthesis, which may contain personal data, and resulting audio	India endpoint	Being documented	Being executed, Phase 1	Residency terms, retention configuration, training-use exclusion
Managed identity and authentication	Platform-user account identifiers, authentication and session data; not call content	Being confirmed	Being documented	Being executed, Phase 1	Hosting region, data-retention terms, session and token configuration

Two things we want to be straightforward about.

Our identity provider is the one service in this list whose hosting region we have not yet confirmed for the purposes of this report. It handles platform-user account and authentication data rather than call content, so the volume and sensitivity are lower than our storage layers, but it is still personal data and a residency statement covering five of six services should say which one is outstanding. Confirming it is a Phase 1 item.

Executed data-processing agreements across the chain are in progress rather than complete. Under the DPDP Act a processor should be engaged under a valid contract, and we are both the processor for your data and the party engaging these six. Most of them publish standard agreements that can be executed quickly, which is why this sits in Phase 1 rather than later.

Change notification. We will maintain our subprocessor register as a live document with a named owner and a review cadence, and we are willing to commit contractually to notifying you of changes to the chain.

Transcription deserves a specific note. Speech-to-text is where raw caller speech is processed, and a caller may say anything, including identifiers, financial details or health information. It is the least controlled input in the system. We are confirming in writing, as a Phase 1 item, whether audio and transcripts are retained by the provider and for how long under our configuration, and that customer audio and text are excluded from any model training or improvement use. Where a zero-retention configuration is available, we will adopt it.

Regulatory position

The Digital Personal Data Protection Act, 2023 received presidential assent on 11 August 2023 and commences in stages. A Gazette notification dated 13 November 2025 brought the first tranche of provisions into force, on the following published schedule:

TRANCHE	EFFECTIVE	PROVISIONS
Immediate	13 November 2025	Section 1(2), section 2, sections 18 to 26, sections 35, 38 to 43, and sections 44(1) and 44(3)
Approximately 12 months	13 November 2026	Section 6(9) and section 27(1)(d)
Approximately 18 months	13 May 2027	Remaining provisions

The Digital Personal Data Protection Rules, 2025 were notified on 14 November 2025, with certain procedural provisions effective immediately and phased compliance dates for others specified within the Rules.

One caveat, stated openly. We have not independently verified the rule-by-rule commencement schedule of the Rules against the Gazette text for this document. The Act-level schedule above is drawn from the published notification. Before either of us relies on a specific rule-level date in a contractual commitment, counsel should confirm it against the Gazette.

What this means practically. The substantive obligations most relevant to our platform, covering notice, consent, security safeguards, breach notification, retention, erasure and data-principal rights, sit largely in the later tranche, with an outer horizon around May 2027. Our roadmap is sequenced to complete the documentation and verification work well ahead of that, and in practice ahead of it anyway, because customers reasonably ask these questions now rather than in 2027.

Our role

In most engagements you are the **data fiduciary** and we are the **data processor**. You determine the purpose and hold the relationship with the individual; we process on your documented instructions. That allocation determines who owes notice, who obtains consent, and who is first accountable to the Data Protection Board.

Our obligations in that role are to process only on your instructions, to implement reasonable security safeguards, to assist you with data-principal requests, to report personal data breaches to you, and to operate under valid agreements both with you and with the third parties we engage. Where we process personal data for our own purposes, such as our own platform-user accounts, we act as fiduciary for that processing and take on the corresponding obligations directly.

We will set this out explicitly in our data-processing agreement, including permitted purposes, subprocessor engagement, request assistance and breach notification, rather than leaving the allocation implied.

Control mapping

DPDP AREA	REQUIREMENT	OUR CONTROL	STATUS	WHAT WE ARE DOING
Purpose limitation	Process only for the specified lawful purpose	We process on your instructions for the intervention channels you configure	Being documented	Recording permitted purposes explicitly in our processing agreement, with no secondary use without your instruction
Notice	Individuals receive clear notice of purposes, data, rights and withdrawal	Sits with you as fiduciary; we support your disclosure requirements	Planned, Phase 1	Publishing our own privacy notice; allocating notice responsibility in the agreement; supporting call-time disclosure and recording announcements as a platform capability
Consent and lawful processing	Consent must be free, specific, informed, unambiguous, by clear affirmative action, or another lawful ground applies	You hold the lawful basis; we process on instruction	Being documented	Contractual representation of lawful basis; support for per-contact consent and opt-out status as a platform feature
Data minimisation	Process only what is necessary	PII masking with non-persistence; configurable data collection	In place, configurable	Recording the masking setting per deployment; confirming masking is applied ahead of AI processing
Encryption	Appropriate technical measures	Encryption at rest in database and object storage; encryption in transit for APIs	In place	Documenting key management, rotation and backup encryption
Reasonable security safeguards	Safeguards to prevent personal data breach	Encryption, role-based access control, private network isolation, rate limiting, one-time credentials, audit logging, two-factor authentication, masking	In place	Formal control documentation, then independent penetration testing in Phase 2
Access control	Need-to-know restriction	Role-based access control with manager and admin separation; private network limiting direct access	In place	Documenting privileged access paths; introducing periodic access reviews in Phase 3
Authentication	Only authorised users reach personal data	Managed identity layer; two-factor authentication available	In place	Enforcing multi-factor authentication on privileged production access in Phase 1
Auditability	Records enabling accountability and investigation	Audit logs of user actions, changes and inserts, with actor and timestamp; configurable retention	In place	Publishing the log specification; extending logging to read access on recordings and transcripts
Retention	Erase when the purpose is served	Default one-month rolling cycle; separately configured log retention	In place, configurable	Publishing a retention schedule covering backups, telemetry and third parties

DPDP AREA	REQUIREMENT	OUR CONTROL	STATUS	WHAT WE ARE DOING
Erasure	Erase on withdrawal, request, or when purpose is served	Deletion on request, propagating to snapshots, backups, object versions, replicas, disaster-recovery copies and third-party services	In place	Publishing the per-store propagation window as a defined deletion SLA in your agreement
Data-principal rights	Access, correction, erasure and grievance redressal within prescribed timelines	Technical deletion capability in place; requests reach us through you	Planned, Phase 1	Documented request workflow with named owner, response SLA and audit trail; contractual definition of the assistance we provide you
Grievance mechanism	A readily available route for grievances	Not yet published	Planned, Phase 1	Publishing a grievance contact and internal escalation and response process
Breach management	Intimate the Board and affected individuals in the prescribed manner and timeline	Telemetry provides the investigative substrate	Planned, Phase 1	Writing a breach procedure covering detection, triage, containment, assessment, notification to you, the Board and affected individuals, and record-keeping. Our highest-priority documentation item.
Processor relationships	Engage processors under a valid contract	We act as processor for you	Being documented	Aligning our standard agreement to DPDP processor obligations, with counsel review
Subprocessor management	Govern and disclose subprocessors	Six third-party services, described earlier in this report	Being executed, Phase 1	Executing agreements across the chain; maintaining the register; committing to change notification
Data residency and transfer	Comply with transfer restrictions and sectoral localisation	India-hosted storage and database; India region for language-model processing; India endpoints for transcription and synthesis	In place	Obtaining written residency terms from each provider; confirming the identity provider's region; verifying backup geography
Accountability	Be able to demonstrate compliance	Substantial technical controls, described throughout	Being documented	Completing the Phase 1 documentation set, then independent assurance in Phase 2

What we are still building, and why in this order

12

We have not produced a checklist that treats every item as equally urgent. Here is our actual prioritisation and the reasoning, because the reasoning is what tells you whether to trust the sequence.

Highest priority

A documented personal-data-breach procedure. If an incident occurs, notification obligations to you, to affected individuals and to the Board arrive under time pressure. Improvised response is where damage compounds. This is our top documentation item. **Phase 1.**

Multi-factor authentication on privileged production access. Administrative credentials are the shortest route to data in an architecture that is otherwise privately networked, so this control carries disproportionate value relative to its cost. **Phase 1.**

Executed agreements across the processing chain. A legal requirement in the chain and a reasonable thing for you to ask for before signing. **Phase 1.**

Security detection on our telemetry. Undetected compromise is the worst class of incident. We have the telemetry; we need the detections and the on-call path. **Phase 3, brought forward if a focused set can ship sooner.**

Also committed

Independent penetration testing, including explicit tenant-isolation and AI-specific scope. Backup geography, encryption and restoration testing. A data-principal request workflow. A published privacy notice and grievance mechanism. Written residency, retention and training-use terms from every provider in the chain. Periodic access reviews. Dependency scanning and vulnerability management. Detection of abnormal bulk export activity. Disaster-recovery testing and business-continuity documentation.

On the sequencing of the penetration test

We have deliberately placed independent testing after the Phase 1 documentation and configuration work rather than first. A test run before privileged access controls and the documented baseline are in place produces findings we already know about, at full cost, and a report that dates quickly. Running it against a remediated, documented baseline produces a cleaner result and a retest letter that is worth sharing with you. If your procurement timeline requires the test earlier, tell us and we will resequence; we would rather meet your timeline than defend ours.

Phase 1 – Documentation, verification and DPDP readiness

Formal security-control documentation. Data-flow diagram and data inventory. Retention and deletion schedule, with a published deletion SLA. Subprocessor register with executed agreements across the chain. Written residency, retention and training-use terms from each provider, including confirmation of our identity provider's hosting region. Incident-response plan and personal-data-breach procedure. Multi-factor authentication enforced on privileged production access. Backup verification covering frequency, retention, encryption, geography and recovery objectives. Audit-log specification. Published privacy notice and grievance mechanism. Data-principal request workflow.

What changes for you: we can answer your security questionnaire from documents and evidence rather than from description, and our DPDP position becomes demonstrable rather than simply real.

Phase 2 – Independent technical assurance

An independent third-party penetration test, with web application, API and infrastructure configuration scope. Tenant isolation as an explicit scope item. AI-specific testing, including resistance to instruction injection through inbound content, whether caller speech or an inbound message, which is a risk category specific to conversational AI and not covered by a standard web application test. Then remediation by severity, and a retest.

What changes for you: our answer on independent testing changes from no to yes, with a summary letter we can share, under NDA where required.

Phase 3 – Programme maturity

Vulnerability management with defined remediation timelines. Secure development lifecycle practices, documented. Automated dependency and security scanning in our build pipeline. Security alerting and detection with on-call routing. Periodic access reviews with sign-off records. Detection of abnormal bulk export activity. Disaster-recovery testing against stated recovery objectives. Business-continuity planning, including provider outage in the AI and voice chain. Vendor-risk management with annual review. Formal policy set.

Phase 4 – Independent certification

We will pursue ISO/IEC 27001 or SOC 2, not both simultaneously, and we will let customer requirements determine which and when. ISO/IEC 27001 is the more common expectation among Indian, European and Asian enterprises and in tender processes; SOC 2 Type II is the default expectation among US enterprises. We would rather have this conversation with you than guess: if certification is a requirement for your organisation, tell us which one and on what timeline, and we will factor it into our planning directly.

Our position on certification and independent testing

14

We do not hold ISO/IEC 27001 certification, a SOC 2 attestation, or an independent penetration-testing report. Rather than leave that as a bare statement, it is worth separating four things that are often treated as one.

	WHAT IT IS	WHO REQUIRES IT
DPDP obligations	Statutory duties under Indian law: lawful processing, notice, reasonable security safeguards, retention limits, erasure, data-principal rights, breach notification, processor contracts, accountability. Enforced by the Data Protection Board with monetary penalties.	Indian law
ISO/IEC 27001	Certification of an information security management system against an international standard, issued by an accredited body after audit. Certifies management process, not the absence of vulnerabilities.	Customers and tenders. Voluntary in law.
SOC 2	An auditor's report on the design, and for Type II the operating effectiveness over a period, of controls against defined criteria. A report rather than a certificate.	Predominantly US enterprise buyers. Voluntary in law.
Penetration testing	A point-in-time technical assessment of applications, APIs and infrastructure, producing findings, remediation and retest. Not a certification, despite common usage.	Customers, some sectoral regulators, and sound engineering practice

None of the three assurance mechanisms is mandatory for DPDP compliance. What the Act requires is reasonable security safeguards and demonstrable accountability. An organisation can hold ISO/IEC 27001 and still fall short of the Act if it has no breach-notification path or ignores erasure requests. An organisation can hold no certification and still meet its obligations through implemented safeguards and documented process.

Cllr AI sits in the second description: the safeguards are substantially in place, and the documentation and independent validation layer around them is the part we are completing. That is what the two sections that follow set out, and it is why our roadmap leads with documentation and testing rather than with certification.

What we provide in place of a certificate. A certification is a proxy for evidence. Where we cannot yet offer the proxy, we offer the evidence directly:

- The documented control set in this report, with each control labelled by its actual status.
- The evidence pack listed later in this report: configuration evidence, redacted log samples, role and permission definitions, encryption and retention settings.
- Completion of your security questionnaire, in your format.
- A technical walkthrough with the engineers who built the platform, rather than a sales briefing.
- Contractual commitments in the agreement and data-processing agreement, which bind us in a way a certificate does not.
- An independent penetration test, with the summary letter and retest letter shared with you on completion.

In our experience a security team that receives all six is generally better informed about a platform than one that receives a certificate alone, and we are glad to be assessed on that basis.

We also recognise that legal position and commercial reality are different things. If your procurement process requires a named certification, that requirement is real regardless of what the Act says, and knowing early lets us plan around your timeline rather than discover it at contract stage.

Frequently asked questions

15

Is Cllr AI DPDP compliant? No organisation can be DPDP certified, so no vendor can hold one. We have implemented a substantial set of controls that map to the Act's requirements, particularly reasonable security safeguards, data minimisation and storage limitation: encryption at rest and in transit, PII masking with non-persistence, role-based access control, private network isolation, audit logging, a thirty-day default data lifecycle and deletion on request. We are actively strengthening formal documentation, processing agreements across our chain, a breach-notification procedure and a data-principal request workflow. In most engagements we act as your processor and you act as data fiduciary. We do not claim full or certified compliance.

Are you ISO 27001 certified? No. Detailed documentation of our implemented controls is available, and certification is on our roadmap subject to customer requirements.

Do you have SOC 2? No.

Do you have a penetration test report? Not currently. An independent third-party test is our Phase 2 commitment. We do operate preventive controls including rate limiting, one-time-use API credentials, private network isolation and role-based access control.

How is personal data protected? Through layered controls: encryption at rest in our database and object storage; encryption in transit for API communications; frontend and application controls preventing unnecessary exposure; PII masking, where enabled, so sensitive data is never stored; role-based access control behind a managed identity layer with two-factor authentication available; containerised workloads in a private cloud network not exposed to the public internet; and a thirty-day default lifecycle limiting how long data is held at all.

Is personal data encrypted at rest? Yes, in both our database and object storage. Key-management documentation is being completed and will be shared.

Is data encrypted in transit? Yes, including all API communications carrying personal data. Our detailed transport configuration is being documented.

Where is customer data stored? In India-hosted object storage and an India-hosted managed database.

Is your database hosted in India? Yes.

Which cloud region do you use? Our storage and database are India-hosted, and language-model processing runs in an India region (Mumbai). We describe our infrastructure at a functional level in this document; specific region and configuration evidence is available to your security team on request.

Which language model do you use, and where does it run? Language-model processing runs on a managed enterprise service in an India region (Mumbai). We will confirm the specific service and models in use to your security team on request.

Which transcription and speech providers do you use? Speech-to-text and speech synthesis are provided by third-party services, both configured on India endpoints. We will name them in our subprocessor register, which we will share with your privacy or vendor-risk team.

Does voice processing happen in India? Our identified core voice-processing services are configured on India endpoints and in an India region: transcription on an India endpoint, speech synthesis on an India endpoint, and language-model processing in an India region (Mumbai). Recordings and structured data are stored in India-hosted infrastructure. We are separately verifying provider contractual terms covering secondary activities such as diagnostic logs, backups and support access.

Does customer data leave India? Our core architecture is configured around India-hosted infrastructure and India endpoints for every identified core processing service. We are not claiming that no data ever leaves India, because that requires verified contractual coverage of provider logs, backups, support access and downstream subprocessors, and that verification is in progress as a Phase 1 item. We will give you a definitive written residency statement on completion. The hosting region of our managed identity provider is being confirmed.

How long is data retained? The default is a rolling one-month, approximately thirty-day, retention and deletion cycle for applicable customer data, and your requirements can adjust it. Audit-log retention is configured separately and is not governed by that cycle. Backup and snapshot retention is being documented.

Can our data be deleted? Yes. Deletion on request covers the primary data stores and propagates to snapshots, backups, object versions, replicas and disaster-recovery copies, and to the third-party services in our processing chain. We are documenting the propagation window for each store so that it can be stated as a defined SLA in your agreement.

Do you support PII masking? Yes. Whether it is enabled depends on your configuration and use case, and we will confirm and record the setting for your deployment.

Do you store masked PII? No. Where masking is enabled, the sensitive personal data is not saved in our system at all. It is a data-minimisation control by design, not a display filter.

Do you have role-based access control? Yes, with separate access levels including manager-level and admin-level, so users receive access according to their assigned role.

Do you support two-factor authentication? Yes, as an available security capability provided per seat. It is not enforced universally by default. We will enable it for your organisation and we recommend it.

Are your databases publicly accessible? No. Our data stores operate within a private cloud network and are not exposed to the public internet. We do not make absolute claims about the impossibility of unauthorised access; layered application, identity and network controls are in place, and independent testing of the configuration is our Phase 2 commitment.

Are your backend systems publicly accessible? No. Our production backend, internal APIs and internal systems run as containerised workloads inside a private cloud network and are not publicly exposed.

Do you maintain audit logs? Yes, covering who made a change, when data was inserted, associated timestamps, campaign start and pause actions, and user changes. Retention is configurable to your requirements and is independent of the thirty-day primary-data cycle. We are documenting default retention, log access controls and extending logging to read access on recordings and transcripts.

Do you maintain backups? Yes, for our database and object-storage infrastructure. Frequency, retention, encryption, geography, restoration testing and recovery objectives are being documented, and we will give you verified figures rather than estimates.

Do you have monitoring? We run open-standard telemetry and observability instrumentation across the platform, giving us tracing, metrics and investigative capability. We do not represent it as a SIEM, security operations centre or intrusion-detection system. Dedicated security alerting is a committed roadmap item.

How do you protect API keys? We use one-time-use API keys and tokens where applicable, which limits the value of an intercepted or accidentally disclosed credential. Rotation intervals and our secrets-management approach are being documented.

Do you implement rate limiting? Yes. It mitigates API abuse, automated attacks, excessive request volume, resource exhaustion and credential-abuse patterns such as brute-force and enumeration.

Which subprocessors do you use? Six third-party services: object storage, a managed database, a language-model service, speech-to-text, speech synthesis, and managed identity and authentication. The table earlier in this report sets out the data each processes and its location. Our register naming each provider is available to your privacy or vendor-risk team, under NDA where you prefer.

What happens if there is a breach? We will notify you, investigate using our telemetry and audit logs, contain the issue and work with your team on assessment and any notification you need to make. The written procedure that formalises those steps into defined roles, severity tiers and notification timelines is our first Phase 1 deliverable, and we will share it with you when it is issued. We would rather tell you it is being formalised than imply a maturity we have not yet documented.

Can we audit you, or send a security questionnaire? Yes to both. We will complete your questionnaire, provide the evidence set listed below, and make our engineering and security people available to walk your team through the architecture.

Available to your security and privacy teams on request, and in most cases immediately.

Architecture and data flow. Architecture diagram. Data-flow diagram covering collection, processing, storage, third-party flows and deletion. Data inventory by category, store, purpose and retention.

Residency and hosting. Region and hosting configuration evidence for our storage and database layers. Region evidence for language-model processing. Endpoint configuration evidence for transcription and speech synthesis.

Encryption. Encryption-at-rest configuration for database and object storage. Key management and rotation. Transport configuration. Backup encryption configuration.

Identity and access. Role definitions and permission matrix. Identity-layer configuration covering session expiry, authentication policy and token lifecycle. Two-factor authentication configuration. Privileged access model and its authentication requirements.

Logging and monitoring. Redacted audit-log samples. Log retention configuration. Log access controls. Telemetry and observability architecture.

Data lifecycle. Masking mechanism description and evidence. Retention configuration. Deletion workflow, including residual-copy propagation and our deletion SLA. Backup configuration and restoration test records.

Network and API. Private network architecture description. Rate-limit configuration. API credential mechanism.

Governance. Subprocessor register naming each provider. Executed processing agreements. Provider residency and retention confirmations. Incident-response and breach procedures. Privacy notice, retention policy and grievance mechanism. Data-principal request workflow. Penetration test summary and retest letter, on completion of Phase 2.

On what we share and how. For some categories we provide a description, an attestation or a redacted sample rather than a raw export. Internal address ranges, exact firewall rules, credential material and unredacted logs are withheld, because circulating an accurate map of a production environment creates risk for your data as much as ours. Where a reviewer needs assurance in one of those areas, a walkthrough under NDA plus the penetration test retest letter is the route we would propose, and we are happy to arrange it.

Our platform protects personal data through encryption at rest in India-hosted storage and database layers, encryption in transit across our APIs, PII masking under which sensitive data is never written at all, role-based access control with manager and admin separation, a managed identity layer with two-factor authentication available, audit logging with independently configurable retention, containerised workloads inside a private cloud network with no public exposure, India-hosted core infrastructure including an India region for language-model processing and India endpoints for transcription and speech synthesis, a thirty-day default data lifecycle, deletion on request, rate limiting, one-time-use API credentials, backups and snapshots, and platform-wide telemetry.

Those controls address what the DPDP Act actually asks for: reasonable security safeguards, data minimisation, storage limitation and erasure. In a few respects, particularly masking that prevents storage rather than concealing it, a thirty-day default lifecycle and India-region AI processing, we think we are ahead of the norm in our category.

There is work in progress, and we would rather you saw it here than found it later. We do not currently hold a certification or an independent test report. Our incident-response and breach procedures, retention schedule, deletion SLA, data-principal request workflow, privacy notice and grievance mechanism are being formalised. Processing agreements across our chain are being executed. Backup geography and restoration testing, privileged access controls and provider residency terms are being verified. Independent penetration testing is commissioned work rather than completed work.

Most of that gap is documentation and verification rather than missing capability, which is why we can commit to closing it in months rather than describing it as a long-term ambition.

Three distinctions we try to be careful about, in this document and elsewhere. Certification is not the same as legal compliance. Configuration is not the same as contractual guarantee, which is why we describe our India-hosted architecture precisely and do not claim that no data ever leaves India. And an available capability is not the same as an enforced one, which is why we describe two-factor authentication and masking as configurable rather than as blanket protections.

We have aimed for an accurate account with its gaps visible rather than a polished one. We would be glad to take your security team through any part of this in detail, complete your questionnaire, and agree the contractual commitments that follow from it.

Document control

Prepared by	Cllr AI
Date	16 September 2026
Basis	Our own account of our platform's controls, prepared for customer review. Not an independent audit, certification or legal opinion.
Regulatory sources	DPDP Act commencement per the published notification of 13 November 2025; DPDP Rules notified 14 November 2025. Rule-level commencement schedule to be confirmed against the Gazette by counsel.
Third-party identities	Described functionally in this document. Full subprocessor register naming each provider available on request, under NDA where preferred.

Review trigger	On completion of any Phase 1 item, on execution of any processing agreement, on completion of independent testing, or on any change to our processing chain or hosting configuration
Next review	On completion of Phase 1, or within six months, whichever is earlier
Contact	connect@cllr.ai