BoardCue Live Master Product Data Protection Impact Assessment

Version 1.1 | Assessment date: 30 August 2026

Status: Conditional approval for enterprise use subject to open implementation and vendor gates

1. Executive decision

A full DPIA remains appropriate because BoardCue combines innovative AI, live meeting transcription, retrieval across multiple governance sources and potentially sensitive information. The assessment reflects BoardCue's governance purpose rather than treating it as an operational case-management platform.

BoardCue is a corporate governance system. It is not intended to replace sector-specific operational systems, hold routine individual case-management datasets, or make operational decisions about individual people.

2. Processing assessed

- account and organisation administration;
- board and committee content storage;
- document extraction, indexing, retrieval, embeddings and search;
- **AI analysis and generation directly to approved Google Cloud Vertex AI and OpenAI API routes;**
- managed browser/OS speech transcription and technically available cloud transcription;
- meeting preparation, live governance prompts and institutional board memory;
- actions, decisions, approvals, briefings and workflow automation;
- governance audit and export;
- outbound email via Resend, inbound email via Postmark, payment via Stripe;
- controlled support access and security/operational logging.

3. Data subjects and data sensitivity

The principal data subjects are BoardCue users, board/committee members, executives, employees/workers, meeting participants, complainants, whistleblowers, advisers and individuals referred to in governance material.

People connected with an organisation's services, workforce, customers, beneficiaries or other stakeholders may be referred to incidentally in governance material, particularly in incident, safeguarding, complaint or workforce reporting. They are not intended to form routine individual case-management datasets. Organisations should use aggregated, anonymised or summarised evidence where practical.

4. Core data flows

1. **Account/workspace: User -> BoardCue -> Supabase authentication/database/storage in the United Kingdom.**
2. **Cloud AI: Authorised BoardCue server request -> Google Cloud Vertex AI and/or OpenAI API under ADEPTABLE-controlled accounts -> response -> BoardCue. Enterprise deployment must document the selected region, retention, data-sharing/training settings and transfer mechanism for each active route.**
3. Managed Live transcription: Meeting audio -> browser/OS speech-recognition capability selected by the user's environment -> transcript text -> BoardCue. Customer controls/approves its browser/device environment.
4. Outbound email: BoardCue -> Resend -> recipient.
5. Inbound intelligence: Sender -> Postmark -> BoardCue inbound route -> quarantine -> configured malware scanner if present -> workspace.
6. Payments: Customer -> Stripe -> payment/subscription status -> BoardCue.

5. Verified infrastructure and supplier position

Area	Current assurance position
Supabase	Production primary database and document-storage region: United Kingdom. Enterprise go-live baseline: Pro or higher with daily backups, 7-day availability according to Supabase published documentation.
Google Cloud Vertex AI/OpenAI	Architecture and model routes verified in code. Exact direct AI provider processing geography, retention, logging, downstream account terms and transfer controls remain pending written vendor confirmation.
Resend	Production outbound service confirmed. Resend publishes US data storage, standard 30-day active-account email/log retention and UK transfer safeguards through its DPA.
Postmark	Production inbound service confirmed. Postmark publishes US data storage/processing and DPA/SCC/UK Addendum arrangements.
Stripe	Production payment service confirmed. Payment information is kept separate from governance Customer Content; Stripe publishes UK DPA transfer arrangements.
Malware scanning	Fail-closed integration exists, but no production provider has yet been verified. Do not name a scanner as an active subprocessor until verified.

6. Necessity and proportionality controls

1. Human judgement remains responsible for governance decisions.
2. Source evidence should be distinguishable from AI interpretation.
3. Customers should not upload identifiable or sensitive data where aggregated or anonymised evidence is sufficient.
4. Role and committee access should be least-privilege and periodically reviewed.
5. Organisations are responsible for approving/configuring the browsers and endpoints used for browser speech recognition.
6. Meeting participants should be told when transcription is active and why.
7. Derived AI/search content should follow the source-data lifecycle.
8. Support access should be exceptional, authorised, time-limited and audited.
9. Organisation customers should be able to export and delete their governance record.

7. Principal risks and controls

ID	Risk	Initial	Control	Residual
R1	Participants unaware of transcription	High	Meeting notice, host responsibility, participant notice template	Medium
R2	Browser/OS speech provider handles audio in a way the customer has not approved	High	Clear allocation: organisation approves/configures its endpoint/browser environment; BoardCue does not warrant third-party browser processing	Medium
R3	Google Cloud Vertex AI/OpenAI retention/geography is insufficiently evidenced	High	Supplier register; vendor questions sent; no unverified zero-retention/UK-only claim	Medium/High until response
R4	Individual-level operational case data is uploaded unnecessarily	High	Intended-use restriction; guidance to use aggregated, anonymised or summarised governance evidence; customer DPIA/minimisation	Medium
R5	Sensitive workforce/incident information is overexposed	High	Role/committee permissions, MFA, RLS, customer access review	Medium
R6	AI output is treated as fact	High	Source-backed evidence, human review, no automatic approval	Medium
R7	Cross-organisation data leakage	Critical	Tenant/RLS controls, automated tenancy testing	Low/Medium
R8	Terminated organisation data persists indefinitely	High	14-day exit/deletion design and deterministic deletion implementation in progress	Medium until implemented
R9	Backup lifecycle contradicts deletion promise	High	Enterprise Supabase Pro baseline and verified backup lifecycle; deletion state	Medium

ID	Risk	Initial	Control	Residual
			reapplied after restore	
R10	Data breach notification is delayed by 'confirmation' step	High	Revised awareness-based notification clause; 24-hour operational target	Low/Medium
R11	MFA organisation requirement is bypassed through server operations	High	Current org MFA control; server-side AAL2 hardening identified for implementation	Medium until implemented
R12	Service-level promise cannot be evidenced	Medium	No contractual 99.5% commitment until health monitoring and monthly measurement are implemented	Low/Medium

8. Enterprise open actions

- **Obtain direct Google Cloud Vertex AI and OpenAI API services written assurance and update downstream Google/OpenAI transfer assessment.**
- Implement and verify the Full Organisation Exit Export, 14-day read-only closure window and deterministic deletion workflow.
- Complete server-side organisation MFA enforcement hardening.
- Implement external health/availability monitoring before activating the 99.5% SLA.
- Verify production malware-scanning provider if inbound attachments are enabled.
- **Upgrade/operate enterprise deployments on Supabase Pro or higher and retain evidence of current backup settings.**
- Commission independent penetration testing and pursue Cyber Essentials as part of the assurance programme.

9. Residual decision

Residual product privacy risk is assessed as Medium, with specific Medium/High uncertainty remaining for Google Cloud Vertex AI/OpenAI supplier terms and implementation-dependent exit/deletion controls. Marketing and controlled onboarding can proceed provided these gaps are stated accurately and are closed or accepted before the affected enterprise deployment.