StoryCycle Genie -- security packet

StoryCycle Labs · 1621 Central Ave, Cheyenne, WY 82001, United States Contact: Sean Schroeder -- sean@storycycle.ai

This packet summarizes how StoryCycle Genie is built, which sub-processors handle customer data, what our customer Data Processing Agreement already commits, how we back up and restore the service, and the vendor review completed on 1 October 2026.

The live documents are:

- Security & Architecture Overview: https://www.storycyclegenie.ai/docs/trust/security-overview
- Sub-processor List (version 1.1, August 2026): https://www.storycyclegenie.ai/docs/trust/subprocessors
- Data Processing Agreement (version 1.0): https://www.storycyclegenie.ai/docs/trust/dpa

Product and security summary

StoryCycle Genie is a multi-tenant application that helps teams plan, generate, and manage brand-storytelling content with large language models. People sign in, work inside their own account, and generate content through guided workflows.

The application is Next.js on Vercel, with TLS in transit. The database, authentication, and file storage are PostgreSQL and Supabase Auth on Supabase, encrypted at rest by that platform. Billing runs through Stripe. Card numbers stay at Stripe; we do not store them. Content generation uses OpenRouter and Cohere. Transactional email uses Resend.

Sign-in is email and password, magic link, or Google. Multi-factor authentication is supported. Privileged administrative access is gated on a second factor. Passwords are not stored in the application in plaintext.

Tenant data is isolated with row-level security on account, membership, billing, API key, and usage tables. The anonymous database role has no privileges. Inside an account, access is role-based. Application-managed secrets, including stored API keys, are encrypted with AES-256 before they are saved.

Prompts and the context supplied with them are sent to the selected model provider only to produce the requested output. We do not use customer content to train our own models. Routing can be limited to a customer-approved provider list where that is feasible.

Content is kept until a customer asks for deletion. A verified deletion request is completed within 30 days. We will complete a security questionnaire or walk through this architecture on a call.

Sub-processors

Version 1.1, August 2026. Each party below processes customer data only as needed to provide its service. Customers with a signed DPA receive notice of material changes and may object, as the DPA describes.

Vercel Inc. -- application hosting, CDN, and TLS termination. Data: application traffic, including account data and content in transit. Location: United States (global CDN edge).

Supabase Inc. -- managed PostgreSQL, authentication, and file storage. Data: account data, credentials held by managed auth, customer content, and usage records. Location: United States (AWS).

OpenRouter Inc. -- LLM inference gateway. Data: user prompts, supplied context, and generated outputs. Location: United States. Upstream model providers sit behind this gateway. Depending on the model selected, those providers include Anthropic, OpenAI, Google, Mistral, and others. Their processing location is the United States or the provider's own location. The 1 October review below covers OpenRouter. It does not separately review each upstream provider.

Cohere Inc. -- LLM inference. Data: user prompts, supplied context, and generated outputs. Location: United States.

Stripe, Inc. -- payment processing and subscription billing. Data: billing contact details. Payment card data is held by Stripe (PCI-DSS Level 1). Location: United States.

Resend, Inc. -- transactional email. Data: customer email addresses and transactional message content. Location: United States.

Customer DPA

The published Data Processing Agreement (version 1.0) is executed per customer. These are the commitments already in that document.

StoryCycle is the processor. The customer is the controller, or a processor acting for another controller. We process customer data to provide and secure the service, on the customer's instructions, and where the law requires it.

The customer authorizes the sub-processors on the published list, including the LLM providers named there. We give at least 15 days' notice of a new or replacement sub-processor. The customer may object on reasonable data-protection grounds. If we cannot resolve the objection, the customer may terminate the affected services and receive a pro-rata refund of prepaid fees.

The DPA states that we impose data-protection obligations on sub-processors no less protective than the DPA, and that we remain liable for their performance. The customer instructs us to send prompts and context to the LLM providers on the list only to generate the requested output. We will not permit customer data to be used to train models, and we pass through equivalent restrictions where the provider offers them.

Security measures in the DPA are encryption in transit and at rest, row-level tenant isolation, role-based access, multi-factor authentication for privileged access, Stripe for card data, and confidentiality for people who handle customer data.

On termination, or on a written request, we delete customer data within 30 days, except where the law or billing records require us to keep it. On request, we export customer content in a common format before deletion.

We notify the customer without undue delay, and within 72 hours, after becoming aware of a personal-data breach that affects that customer's data, and we share what we reasonably know about the nature, scope, and remediation.

We assist with data-subject requests, security of processing, breach notices, and data-protection impact assessments, taking into account the nature of the processing.

We answer security questionnaires and, when we have them, provide summaries of third-party assessments. No more than once a year, except after a breach, the customer may audit through a process we both agree, without exposing other customers' data.

Customer data is processed in the United States. Where the law requires a transfer tool for EU or UK personal data, the Standard Contractual Clauses are incorporated when the DPA is executed. If the DPA and the main agreement conflict on personal data, the DPA controls.

Our backup and restore plan

This section is StoryCycle's own plan, last updated 4 August 2026. It describes recovery of the StoryCycle Genie application and its data. It is not a statement of a vendor's published recovery time or recovery point.

Failure	What we do	Time, as stated in our plan
Bad deploy or application regression	Promote the last known-good Vercel production deployment	Minutes. A routing change. No data loss.
Database corruption, a bad migration, or accidental deletion	Restore with Supabase point-in-time recovery, enabled 4 August 2026	The plan describes data loss going forward as near-zero.
Vercel or Supabase platform outage	We do not run our own multi-region failover on top of those providers	Recovery is bounded by that provider's own recovery time
Credential or secret compromise	Rotate the affected credential	Same day

On 30 July 2026 we checked production rather than relying on a vendor's marketing. Vercel had prior production deployments available to roll back, and our plan calls that application recovery time minutes. Supabase daily logical backups had completed for seven consecutive days. Point-in-time recovery was confirmed enabled on 4 August 2026. The plan, written that day, says the recoverable window builds from enablement toward seven days and was expected to reach that window by mid-August 2026. This packet does not add a measurement after that date.

We review the plan every January, and after an incident or a material infrastructure change. Backup checks and a restore drill are quarterly (January, April, July, October). The first drill was 30 July 2026: a read-only production dump was restored into a disposable database, and the accounts table came back with 368 rows. That drill checked that table, not every table. The next drill in the plan is October 2026.

We do not run an independent multi-region failover beyond what Vercel and Supabase provide on their own infrastructure.

Vendor review -- 1 October 2026

On 1 October 2026 we completed a checklist of each provider's published standard terms and a write-up of its public security pages. The review covers the six sub-processors below. It does not cover a separate review of

each upstream model provider behind OpenRouter. Recovery time (RTO) and recovery point (RPO) below refer to what that provider publishes, not to the StoryCycle plan in the previous section.

We have not executed a separate OpenRouter data processing agreement, and we have not executed a separate Cohere data processing agreement.

Vercel. Role: application hosting, CDN, and TLS. Checklist of published terms and public security pages: completed. RTO: not published. RPO: not published.

Supabase. Role: database, authentication, file storage, and backups. Checklist of published terms and public security pages: completed. RTO: not published. RPO: not published.

OpenRouter. Role: LLM inference gateway, and the model vendor of record for this review. Checklist of published terms and public security pages: completed. RTO: not published. RPO: not published.

Cohere. Role: LLM inference. Checklist of published terms and public security pages: completed. RTO: not published. RPO: not published.

Stripe. Role: payment processing and subscription billing. Card data stays at Stripe. Checklist of published terms and public security pages: completed. RTO: not published. RPO: not published.

Resend. Role: transactional email. Checklist of published terms and public security pages: completed. RTO: not published. RPO: not published.

We are glad to walk through this packet, complete a questionnaire, and provide the DPA for signature.

Sean Schroeder StoryCycle Genie sean@storycycle.ai

Security & Architecture Overview

Full published document.

Security & Architecture Overview

Company: StoryCycle Labs · 1621 Central Ave, Cheyenne, WY 82001, United States **Last updated:** June 2026 **Contact:** sean@storycycle.ai · legal@storycycle.ai

This document summarizes the architecture, data handling, and security controls of StoryCycle Genie ("the Service"), an AI-assisted brand-storytelling platform. It is intended to support a prospective customer's security and governance review. We are happy to complete a vendor security questionnaire or hold a technical review call to go deeper on any item below.

1. Service summary

StoryCycle Genie is a multi-tenant SaaS application that helps teams plan, generate, and manage brand-storytelling content using large language models (LLMs). Users authenticate, work within an isolated account/organization, and generate content through guided AI workflows.

2. Architecture & hosting

- **Application:** Next.js / React (TypeScript), hosted on **Vercel** with

managed hosting and automatic TLS.

- **Database & auth:** PostgreSQL and Supabase Auth on **Supabase** -- managed

Postgres, encrypted at rest by the platform.

- **File/object storage:** Supabase Storage with access-controlled buckets.
- **Billing: Stripe** -- card data handled entirely by Stripe (PCI-DSS

Level 1); we never store card numbers.

- **LLM gateway: OpenRouter** and **Cohere**, routing inference to

upstream model providers (see section 5).

- **Transactional email: Resend**.

Encryption in transit: all client and server traffic is served over TLS by Vercel and Supabase.

Encryption at rest: database and object storage are encrypted at rest by the underlying managed platforms (Supabase / AWS).

3. Authentication & identity

- Identity and session management are provided by **Supabase Auth**.
- Supported sign-in methods: email/password, magic link, and OAuth (Google).
- **Multi-factor authentication (MFA)** is supported, with AAL2 (second-factor)

enforcement available. Privileged/administrative access is gated behind MFA compliance checks at the database level.

- Passwords are never stored by the application in plaintext; credential

storage is handled by the managed auth provider.

4. Data access control & tenant isolation

- **Row-Level Security (RLS)** is enabled across tenant-scoped tables --

including accounts, memberships, roles/permissions, billing, API keys, and usage records. Tenants can only read and write data belonging to their own account.

- The anonymous database role is granted **no privileges**; access requires an

authenticated session, and grants are scoped to the authenticated principal.

- Role-based access control governs what members can do within an account

(owner / member / role-based permissions).

- **Application-managed secrets** (e.g., stored API keys and sensitive

credentials) are encrypted at the application layer with **AES-256** before being persisted.

5. AI / LLM data flow

StoryCycle Genie generates content by sending user prompts and relevant context to LLMs. Inference requests are routed through **OpenRouter**, which forwards them to the selected upstream model provider -- which may include **Anthropic, OpenAI, Google, Mistral, and others**, depending on the configured model.

What this means for a reviewer:

- User-entered prompts and supplied context are transmitted to these

third-party model providers for the sole purpose of generating the requested output.

- We do not use customer content to train our own models.
- Data-processing terms for inference are governed by OpenRouter and the

upstream providers' terms; we can provide the current provider list on request and align model routing to a customer's approved-vendor constraints where feasible.

We surface this flow explicitly because a governance review should account for where content is processed, not just where it is stored.

6. Sub-processors

Core sub-processors today: **Vercel** (hosting), **Supabase** (database, auth, storage), **Stripe** (billing), **Resend** (transactional email), and **OpenRouter** and **Cohere** plus the upstream LLM providers OpenRouter routes to (Anthropic, OpenAI, Google, Mistral, and others). See the Sub-processor List in this section for the formal, versioned document.

7. Monitoring & operations

- Billing and subscription events are processed through Stripe with signed

webhooks.

- Infrastructure-level monitoring (uptime, abuse/DDoS detection) is provided

by Vercel and Supabase as part of their own platform operations.

8. Policies & legal documents

- **Privacy Policy:** published (see this section).
- **Data Processing Agreement (DPA):** published (see this section); executed

per customer.

- **Versioned sub-processor list:** published (see this section).
- **Data retention & deletion:** content is retained until a customer deletion

request; verified deletion requests are honored within 30 days.

- Terms of Service are in production and cover service use, IP ownership,

account security, and limitation of liability.

Additional compliance documentation -- including our current control roadmap -- is available to customers and prospective customers on request under a direct exchange.

9. How to engage

We're glad to:

- Complete your standard vendor security questionnaire (SIG, CAIQ, or your

own).

- Walk through this architecture on a technical review call.
- Provide our DPA and current sub-processor list.
- Scope model-routing constraints to your approved-vendor requirements.

Contact: Sean Schroeder -- sean@storycycle.ai

Sub-processor List

Full published document.

Sub-processor List

Version: 1.1 **Last updated:** August 2026 **Contact:** sean@storycycle.ai

StoryCycle Labs ("StoryCycle") uses the following third-party sub-processors to deliver StoryCycle Genie. Each sub-processor processes customer data only to the extent necessary to provide its service, under terms that impose data-protection obligations consistent with our own commitments.

We will update this list when sub-processors are added or replaced. Customers with a signed DPA will be notified of material changes with an opportunity to object, per the DPA's terms.

Infrastructure & platform

Vercel Inc.

- Purpose: application hosting, CDN, TLS termination
- Data categories: all application traffic (account data, content in transit)
- Processing location: United States (global CDN edge)
- Security reference: vercel.com/security

Supabase Inc.

- Purpose: managed PostgreSQL database, authentication, file storage
- Data categories: account data, credentials (managed auth), customer content,

usage records

- Processing location: United States (AWS)
- Security reference: supabase.com/security

AI / content generation

OpenRouter Inc.

- Purpose: LLM inference gateway
- Data categories: user prompts and supplied context, generated outputs
- Processing location: United States
- Security reference: openrouter.ai/privacy

Upstream model providers via OpenRouter -- Anthropic, OpenAI, Google, Mistral, and others depending on the model selected

- Purpose: LLM inference
- Data categories: user prompts and supplied context, generated outputs
- Processing location: United States / provider-dependent
- Security reference: per provider

Cohere Inc.

- Purpose: LLM inference
- Data categories: user prompts and supplied context, generated outputs
- Processing location: United States
- Security reference: cohere.com/security

Note on AI processing: prompts and context are transmitted to the model provider selected for a given request solely to generate the requested output. StoryCycle does not use customer content to train models. Model routing can be constrained to a customer-approved provider list where feasible -- contact us.

Payments & operations

Stripe, Inc.

- Purpose: payment processing, subscription billing
- Data categories: billing contact details; payment card data is held by

Stripe only (PCI-DSS Level 1)

- Processing location: United States
- Security reference: stripe.com/docs/security

Resend, Inc.

- Purpose: transactional email delivery
- Data categories: customer email addresses, transactional message content
- Processing location: United States
- Security reference: resend.com/security

Change log

- **1.1 -- August 2026:** added Cohere and Resend as active sub-processors;

removed Sentry (not a governed data-processing control).

- **1.0 -- June 2026:** initial publication

Data Processing Agreement

Full published document.

Data Processing Agreement (DPA)

Version: 1.0 **Between:** [CUSTOMER LEGAL NAME] ("Customer") and StoryCycle Labs, 1621 Central Ave, Cheyenne, WY 82001, United States ("StoryCycle")

This Data Processing Agreement ("DPA") forms part of the agreement between Customer and StoryCycle governing Customer's use of StoryCycle Genie (the "Service") (the "Agreement").

1. Definitions

"Personal Data," "Controller," "Processor," "Data Subject," "Processing," and "Supervisory Authority" have the meanings given in applicable data protection law, including the GDPR where applicable. "Customer Data" means data submitted to the Service by or on behalf of Customer, including Personal Data.

2. Roles and scope

2.1 Customer is the **Controller** (or a Processor acting on behalf of another Controller) of Personal Data within Customer Data; StoryCycle is a **Processor**.

2.2 StoryCycle will Process Customer Data only (a) to provide and secure the Service, (b) per Customer's documented instructions, including use of the Service's features, and (c) as required by law, in which case StoryCycle will notify Customer unless legally prohibited.

2.3 **Details of Processing.** Subject matter: provision of the Service. Duration: the term of the Agreement plus the deletion period in section 6. Nature and purpose: hosting, storage, transmission, and AI-assisted content generation. Categories of Data Subjects: Customer's authorized users and individuals whose data appears in Customer Content. Categories of Personal Data: account and contact data, user-submitted content, usage data.

3. Confidentiality

StoryCycle ensures that personnel authorized to Process Customer Data are bound by confidentiality obligations and Process it only as needed to provide the Service.

4. Sub-processors

4.1 Customer authorizes the sub-processors listed in StoryCycle's Sub-processor List (Annex B), including the LLM inference providers identified there.

4.2 StoryCycle will provide at least **15 days' notice** of new or replacement sub-processors. Customer may object on reasonable data-protection grounds; if the parties cannot resolve the objection, Customer may terminate the affected services and receive a pro-rata refund of prepaid fees.

4.3 StoryCycle imposes data-protection obligations on sub-processors no less protective than this DPA and remains liable for their performance.

4.4 **AI processing instruction.** Customer instructs StoryCycle to transmit prompts and context to the LLM providers in Annex B solely to generate requested outputs. StoryCycle will not permit Customer Data to be

used to train models, and will pass through equivalent restrictions to LLM sub-processors where the provider offers them.

5. Security

StoryCycle implements appropriate technical and organizational measures described in **Annex A**, including encryption in transit and at rest, row-level tenant isolation, role-based access control, and multi-factor authentication for privileged access. StoryCycle may update these measures provided overall security is not materially reduced.

6. Deletion and return

Upon termination of the Agreement, or upon Customer's written request, StoryCycle will delete Customer Data within **30 days**, except where retention is required by law or for billing records. Upon request, StoryCycle will provide Customer an export of Customer Content in a commonly used format before deletion.

7. Personal data breach

StoryCycle will notify Customer **without undue delay and within 72 hours** of becoming aware of a Personal Data Breach affecting Customer Data, providing information reasonably available about the nature, scope, and remediation of the breach, and will cooperate with Customer's reasonable requests.

8. Assistance

Taking into account the nature of Processing, StoryCycle will reasonably assist Customer with: (a) responding to Data Subject requests (access, deletion, portability, etc.); (b) security of Processing; (c) breach notifications to authorities or Data Subjects; and (d) data protection impact assessments, where required.

9. Audit

StoryCycle will make available information reasonably necessary to demonstrate compliance with this DPA, including responses to Customer's security questionnaires and summaries of third-party assessments when available. No more than once annually (except after a breach), Customer may conduct an audit through a mutually agreed process that does not compromise other customers' data.

10. International transfers

Customer Data is processed in the United States. Where applicable law requires a transfer mechanism for EU/UK Personal Data, the parties incorporate the Standard Contractual Clauses per **Annex C**.

11. Liability and order of precedence

Liability under this DPA is subject to the limitations in the Agreement. If this DPA conflicts with the Agreement regarding Processing of Personal Data, this DPA controls.

Annex A -- Technical and Organizational Measures

- **Encryption in transit:** TLS for all client and service-to-service

traffic.

- **Encryption at rest:** platform-level encryption of database and object

storage; application-layer AES-256 encryption for stored credentials and API keys.

- **Tenant isolation:** PostgreSQL row-level security on all tenant-scoped

tables; anonymous role holds no privileges.

- **Access control:** role-based access within accounts; least-privilege

database grants.

- **Authentication:** managed identity provider; MFA supported; MFA enforced

for privileged/administrative access.

- **Payment data:** handled exclusively by Stripe (PCI-DSS Level 1);

StoryCycle does not store card data.

- **Monitoring:** error monitoring with PII capture disabled by default.
- **Personnel:** confidentiality obligations; access limited to operational

need.

Annex B -- Sub-processors

The current Sub-processor List (published in this documentation section) is incorporated by reference.

Annex C -- Standard Contractual Clauses

Where Customer Data includes Personal Data subject to the GDPR or UK GDPR, the parties will execute the applicable Standard Contractual Clauses, which are incorporated into this DPA upon execution.

Signatures

Customer: Name / Title / Date

StoryCycle: Name / Title / Date

StoryCycle's own plan

This document is StoryCycle's own backup and restore plan.

StoryCycle Genie -- Business Resiliency & Disaster Recovery Plan

Company: StoryCycle Labs · 1621 Central Ave, Cheyenne, WY 82001, United States **Last updated:** August 4, 2026 **Security contact:** Matt Levine

1. Scope

This plan covers recovery of the StoryCycle Genie application and its data after a deployment failure, data-integrity incident, or platform-level outage. It complements the accompanying Incident Response Plan, which covers detection and response; this document covers what we can actually recover to, and how fast.

2. Failure modes & recovery mechanism

Failure	Recovery mechanism	Recovery time
Bad deploy / application regression	Vercel retains prior production deployments as rollback candidates; promote the last known-good build	Minutes -- a routing change, no data loss
Database corruption, bad migration, accidental deletion	Restore via Point-in-Time Recovery (PITR), enabled August 4, 2026	Near-zero data loss going forward. The recoverable window is still building up from enablement -- see §3 for current coverage versus the eventual 7-day target.
Vercel or Supabase platform-level outage	Both providers run their own multi-AZ, replicated infrastructure with published uptime commitments; we do not run independent multi-region failover on top of theirs	Bounded by the vendor's own recovery time, outside our control
Credential/secret compromise	Rotate the affected credential immediately	Same-day (see Incident Response Plan §3, Contain)

3. Recovery objectives -- verified, not assumed

We checked actual current state on July 30, 2026 rather than relying on vendor marketing claims:

- **Application (Vercel):** the production deployment history shows multiple `READY` builds flagged as rollback candidates going back several weeks. Rollback is real and immediately available; RTO is minutes.
- **Database (Supabase):** daily logical backups complete successfully -- 7 consecutive days observed, all status `COMPLETED`. **Point-in-time recovery (PITR) is now enabled** (confirmed August 4, 2026), giving a near-zero Recovery Point Objective going forward rather than up to ~24 hours. The recoverable window builds up from the moment PITR was enabled rather than being available immediately -- it will reach the full 7-day window by mid-August 2026 as WAL archives accumulate.

4. Testing & review

- This plan is reviewed **annually every January**, and additionally after any real incident or material infrastructure change.
- **Backup/restore verification and full restore drill: quarterly** (January, April, July, October). Each cycle: confirm recent Supabase backups exist and completed successfully, confirm Vercel has rollback-candidate production deployments available, and perform a full restore test into a disposable database. First cycle completed July 30, 2026; next due October 2026.
- **First full restore drill completed July 30, 2026.** Production was dumped (read-only, single touch) and restored into a disposable test database on a separate Postgres server. The restore succeeded: `public.accounts` landed all 368 production rows intact. Nine errors were reported and ignored by the restore process; all nine were confined to platform-managed internal infrastructure (event triggers requiring elevated database privileges, and internal extension objects already present from the database template) -- none affected application data.
- One incidental finding from the drill, unrelated to recovery capability: every restored row in `public.accounts` had a null value in a `created_at` timestamp column. Not a data-loss or security issue; flagged as a data-quality item for engineering follow-up.

5. What we don't have

- No independent, company-run multi-region failover beyond what Vercel and Supabase each provide on their own infrastructure.
- The July 30 restore drill validated one table (`accounts`) as a representative spot-check, not every table in the schema -- broader per-table validation is the next step up in rigor if needed.