

SECURITY & PRIVACY

Data & Security

How gimmefy protects your data, your content and your customers' trust. Plain language first, then the technical detail your IT, security and legal teams will ask for. Every statement in this document is true of the platform as it runs today. Where something is planned rather than live, it says so.

2PRODUCTION REGIONS ·
FRANKFURT & TOKYO**900+**ROW-LEVEL SECURITY
POLICIES · EVERY
TABLE**0**MODELS TRAINED ON
CUSTOMER DATA**72h**BREACH NOTIFICATION
COMMITMENT

A note on certifications, up front. gimmefy does not hold its own SOC 2 or ISO 27001 certificate. The infrastructure we run on does: Amazon Web Services (ISO/IEC 27001, 27017, 27018, SOC 1/2/3, PCI DSS) and Supabase (SOC 2 Type II). Section 02 explains exactly what that covers and what it does not, so nobody has to read between the lines.

What's in this document

01	Executive summary Eight things that are true for every account
02	Certifications and shared responsibility Who is certified for what, and the regulations we support you under (GDPR, NIS2, PDPA, CCPA)
03	Architecture and hosting Regions, components, network boundary, backups, environments
04	Access and identity Sign-in, SSO, MFA, roles, tenant isolation, staff access
05	How your data reaches AI models What is sent, what is never sent, provider retention terms, zero-data-retention options
06	Keeping confidential company information confidential The eight controls between your strategy documents and the outside world
07	Your data, your rights Ownership, export, deletion, retention, residency
08	Security operations Secure development, monitoring, incident response, risk governance, training, testing
09	Subprocessors Every third party that touches data, what they receive, and where they are
10	Roadmap What is in progress, honestly labelled
11	Questions from IT, security and legal teams Straight answers to the ones we get asked most
12	Contact and document control

01 Executive summary

Eight things that are true for every gimmefy account, on every plan, in both regions.

- 1. **Your data is yours.** You own what you upload and what gimmefy generates for you. We do not sell it, rent it or share it with anyone beyond the providers needed to run the product.
- 2. **We never train AI on your data.** gimmefy does not build, train or fine-tune any model. The providers we use are accessed through their commercial APIs, whose terms prohibit training on API inputs and outputs (Section 05).
- 3. **No other customer can see your data.** Isolation is enforced inside the database with more than 900 row-level security policies. Row-level security is switched on for every table. It does not depend on application code getting it right.
- 4. **Encrypted everywhere.** TLS 1.2+ in transit, AES-256 at rest.
- 5. **EU data stays in the EU.** European customers are hosted in Frankfurt, Germany (AWS eu-central-1). Database, authentication, file storage and media rendering all run there.
- 6. **Certified infrastructure, honestly described.** We run on AWS and Supabase, both independently audited. gimmefy itself is not yet certified, and we say so wherever the question comes up.
- 7. **You stay in control.** Export your content at any time. Ask for deletion at any time. We delete within 30 days of the end of the agreement.
- 8. **You can tighten the model roster.** Enterprise and white-label clients can restrict which AI providers their workspace may use, and can ask us to arrange zero-data-retention terms with the providers that offer them.

02 Certifications and shared responsibility

Security teams ask two different questions that sound the same: "is your infrastructure certified?" and "are you certified?". The answers differ, and we would rather you hear both from us.

LAYER	OPERATED BY	INDEPENDENT ATTESTATION	EVIDENCE AVAILABLE
Cloud infrastructure: data centres, network, compute, storage, CDN	Amazon Web Services	ISO/IEC 27001:2022, 27017, 27018 · SOC 1, SOC 2 Type II, SOC 3 · PCI DSS v4.0.1	SOC 3 (public). SOC 1/2 reports and ISO certificates under NDA.
Database, authentication, object storage, edge functions	Supabase (runs on AWS in our chosen region)	SOC 2 Type II · HIPAA-eligible · GDPR DPA	SOC 2 Type II report under NDA.
Media rendering services	Google Cloud (Cloud Run)	ISO/IEC 27001 · SOC 1/2/3	Google Cloud compliance reports (public portal).
Payments (direct customers only)	Stripe	PCI DSS Level 1	Public.
The gimmefy application and our operations	Teemuno Pte Ltd	No certificate held today. Controls documented in this paper and in our Information Security Policy, Risk Governance Framework and Secure-SDLC Summary.	Policy extracts, control reviews and completed questionnaires (SIG, CAIQ, custom) on request.

What the provider certifications cover: physical security, network security, hypervisor and host hardening, encryption of storage, availability, the provider's own personnel and change management. What they do not cover: how gimmeify configures those services, how our application enforces authorisation, how our staff access production, or how we develop and release code. Those are ours, and Sections 03 to 08 describe them.

Regulations we support you under

REGULATION	OUR POSITION	WHAT YOU GET
GDPR and UK GDPR	gimmeify acts as processor (or sub-processor behind a white-label partner). Teemuno Pte Ltd is the contracting entity.	Data Processing Agreement with EU Standard Contractual Clauses (2021/914) and UK Addendum, 72-hour breach notification, sub-processor list, deletion and export commitments, audit cooperation. EU hosting in Frankfurt.
NIS2 (EU Directive 2022/2555)	gimmeify is a marketing software supplier and is not itself an essential or important entity under NIS2. Our customers who are in scope must manage supplier risk under Article 21(2)(d).	The supplier evidence that obligation needs: named security contact, incident notification commitment, documented risk management, secure development practice, sub-processor transparency, and completion of your supplier questionnaire.
Singapore PDPA	Teemuno Pte Ltd is a Singapore company and complies with the PDPA in its own right.	Data protection obligations reflected in our Information Security Policy and staff training.
CCPA / CPRA (US)	We do not sell personal information and act as a service provider to US customers.	Support for access and deletion requests through your account team.
EU AI Act	gimmeify does not develop or train general-purpose models. We integrate third-party models as an orchestration layer.	Transparent, always-current roster of the models in use (in-product AI Stack page), plus the ability to restrict it.

03 Architecture and hosting

gimmefy is a single codebase deployed to two independent production regions. Your account lives in one of them. Nothing about your account, content or usage is replicated to the other.

European Union · Frankfurt, Germany

AWS eu-central-1, live since June 2026.
Database, authentication, file storage and edge functions on Supabase in Frankfurt. Media rendering on Google Cloud Run in Frankfurt (eu-west3). Serves European and UK customers and white-label partners.

Asia-Pacific · Tokyo, Japan

AWS ap-northeast-1. Database, authentication, file storage and edge functions on Supabase in Tokyo. Media rendering in Singapore. Serves Asia-Pacific customers and partners, and is the default region for customers outside the EU.

Components

COMPONENT	WHAT IT DOES	SECURITY NOTES
Web application	Single-page app served through Amazon CloudFront	HTTPS only. No secrets in the browser: every provider key lives server-side.
API layer	200+ edge functions (Deno isolates) behind the Supabase API gateway	Every call is authenticated with a short-lived JWT. Unauthenticated requests are rejected at the gateway. Origin allow-listing on all functions.
Database	PostgreSQL (managed by Supabase)	Not reachable from the public internet. Every table has row-level security enabled; 900+ policies scope rows by partner, organisation, workspace and user.
Object storage	Uploads, brand assets, generated media	Customer documents and uploads sit in private buckets, readable only through signed, time-limited URLs. Public buckets hold only generated deliverables intended for URL sharing.
Identity	Supabase Auth	Email + password, Google sign-in, SAML 2.0 SSO. Passwords stored as salted bcrypt hashes only. TOTP MFA.
Rendering workers	Video and visual rendering on Google Cloud Run	Private ingress, scoped service accounts, in-region.

Backups, recovery and environments

- Automated daily database backups with point-in-time recovery, retained in-region.
- Object storage is replicated within the region by the provider.
- Separate development, staging and production environments with separate databases and credentials. No customer data in non-production environments.
- Release path: feature branch → dev → main → staging → production, with peer review and automated checks at each gate. Deployment credentials are MFA-gated.

04 Access and identity

Getting in

- **Invite-only provisioning.** An organisation admin invites a user by work email. The user sets their own password through a time-limited link. No default or shared credentials exist. Email addresses are verified.
- **Single sign-on.** SAML 2.0 SSO through our identity provider, configured per organisation and per email domain. Microsoft Entra ID, Okta and Google Workspace are supported. Set up as an onboarding deliverable for enterprise and white-label clients.
- **Multi-factor authentication.** TOTP (Google Authenticator, Microsoft Authenticator, Authy, 1Password). Any user can enable it. Organisation admins can make it mandatory for every member, and the platform blocks access until enrolment is complete.
- **Session handling.** Short-lived access tokens with refresh rotation. Removing a user takes effect immediately at the database layer, not just in the interface.

Roles

Three organisation roles: **Org Admin** (users, settings, billing, everything inside the organisation), **Manager** (day-to-day content and workspace management) and **Member** (create and use). A white-label partner additionally holds a partner-level role for provisioning and support across the client organisations it hosts. Role changes are checked in the database, so a user cannot escalate their own privileges through the API.

Tenant isolation

Every row in the database carries the identifiers of the partner, organisation and workspace it belongs to, and a row-level security policy checks those identifiers against the caller's identity on every read and write. This is why we describe isolation as structural: a bug in application code cannot expose another tenant's rows, because the database refuses to return them. White-label partners are isolated from each other and from gimmeify's own tenants in the same way. We verify the RLS perimeter with automated security advisors on every review cycle.

Our own staff

- Production access is limited to named engineers on a least-privilege basis and reviewed quarterly.
- Every production console (cloud provider, database platform, source control) requires MFA.
- Customer content is accessed only to resolve a support request you have raised, or for an incident, and such access is logged.
- All personnel complete security training at onboarding and annually, covering data classification, access control, PDPA and GDPR obligations, and the rule that client data is never used to train models.

05 How your data reaches AI models

gimmefy is an orchestration layer. We do not build or host foundation models. When you use a capability, gimmefy assembles a request, sends it to the relevant provider over an encrypted connection, and returns the output to your workspace. This section describes that exchange precisely, because it is the question enterprise security teams care about most.

What is sent, per request

SENT TO THE PROVIDER	NEVER SENT
• Your prompt or brief for that task. • The brand context relevant to the task: selected fields from your Brand Vault (voice, audience, guidelines) and, where knowledge files are enabled, the passages retrieved for that request. Retrieval is by semantic search; the whole vault is not sent. • Files you attach to that specific request (a document to summarise, a reference image). • Conversation history for that thread, so the model has context.	• Anything belonging to another organisation or partner. • Your login credentials, API keys, billing or payment data. • Your user directory or account metadata. • Your Brand Vault in bulk, or content from other workspaces or projects. • Anything at all, unless a user has chosen a capability that uses that provider.

What happens to it afterwards

Three facts apply to every provider we use. First, gimmefy keeps no training pipeline and never uses customer content to train, fine-tune or evaluate any model. Second, we access providers through their commercial APIs, not consumer products, so the consumer-grade "we may train on your chats" terms do not apply. Third, provider retention is governed by each provider's API terms, which we summarise below and verify periodically.

PROVIDER	TRAINING ON API DATA	RETENTION AFTER THE REQUEST	ZERO DATA RETENTION (ZDR)
Anthropic (Claude)	Prohibited by Anthropic's Commercial Terms: Anthropic "may not train models on Customer Content".	Inputs and outputs deleted from Anthropic's systems within 30 days, held only for trust-and-safety purposes.	Available by agreement with Anthropic.
OpenAI (GPT, GPT Image)	Not used for training by default for API customers.	Retained up to 30 days for abuse monitoring, then deleted.	Available for eligible endpoints on approval by OpenAI.
Google (Gemini, Imagen, Veo, paid tier)	Google does not use paid-tier prompts or responses to improve its products.	Logged for a limited period solely for abuse detection (Google states 55 days for the Gemini API).	Available on request to Google.

Other roster providers (xAI, Perplexity, DeepSeek, BytePlus, Kuaishou, Moonshot, MiniMax, Zhipu, StepFun, Alibaba, Manus, fal.ai, OpenRouter, ElevenLabs, HeyGen)	Accessed through commercial API tiers. We do not enable any provider whose API terms permit training on customer inputs.	Per provider; typically a short abuse-monitoring window.	Varies. Clients who need ZDR guarantees restrict the roster to the providers above.
---	--	--	---

Where we stand on ZDR today. gimmefy uses the providers' standard commercial API terms. We have not yet executed separate zero-data-retention agreements, so the standard windows above apply. For enterprise and white-label clients who need ZDR, we will apply for it with the providers that offer it and restrict the client's workspace to those providers. If your organisation already holds an enterprise agreement with a provider (for example with Anthropic), tell us: we can align the roster with it.

The roster

Around 50 models from roughly 20 providers are live at any time. Text and reasoning: Claude, GPT, Gemini, Grok, Perplexity Sonar, DeepSeek, Kimi, Qwen, MiniMax, GLM, Manus, Step. Image: GPT Image, Gemini Image, Seedream, Kling, Grok Imagine, fal.ai. Video: Veo, Seedance, Kling, Grok Imagine Video. Voice and avatar: ElevenLabs, HeyGen. The live list by capability is in the product (Dashboard → AI Stack) and changes as models ship or retire. Enterprise and white-label clients can restrict it to named providers; the common request is Anthropic, OpenAI and Google only.

Observability

To debug failures and measure quality we trace AI requests with LangSmith (LangChain, Inc., United States). Traces can include prompts and outputs. Access is limited to engineering, traces are retained for a limited period, and tracing can be disabled for a workspace on request.

06 Keeping confidential company information confidential

Marketing teams put unreleased product plans, pricing, strategy decks and customer research into gimmeFY. The question we are asked is: how do we guarantee that material does not end up in a model, in another tenant's results, or anywhere else it should not be? Eight controls, in order from the outside in.

#	CONTROL	WHAT IT PREVENTS
1	No training, by us or by providers. gimmeFY has no training pipeline. Provider API terms prohibit training on inputs and outputs (Section 05).	Your strategy becoming part of a model.
2	Inference-only, transient processing. A request is sent to produce one output, and the provider's retention window is bounded and disclosed.	Indefinite copies at the provider.
3	Selective context, not bulk export. Only the passages relevant to a request leave the vault, retrieved by semantic search at the moment of use.	A whole knowledge base travelling with every prompt.
4	Database-enforced tenant isolation. 900+ row-level security policies, RLS on every table.	Another customer, or another partner's client, seeing your rows.
5	Private storage with signed URLs. Uploaded documents are never publicly addressable.	Files being reachable by guessing a link.
6	Restricted, logged staff access. Named engineers, MFA, least privilege, quarterly review, access only on your request or for an incident.	Casual internal access to customer content.
7	Roster and region restriction. Your workspace can be limited to named providers, and ZDR arranged where offered.	Content reaching a provider you have not approved.
8	Contractual backing. Confidentiality and no-training commitments in the DPA and agreement, sub-processor list with notice of changes, deletion on exit.	Reliance on goodwill instead of paper.

What we will not claim: that information is mathematically incapable of leaving our systems. Any SaaS that processes your content depends on controls, contracts and audited providers. Ours are listed above and we will walk your security team through any of them.

07 Your data, your rights

You own your content

You keep ownership of everything you upload and everything gimmeify generates for you. gimmeify retains only its own platform and pre-existing technology.

Minimal personal data

Account data (name, work email, role), authentication data, and technical logs (IP address, device, access events). We do not build profiles of your customers.

Export

Outputs are downloadable at any time in open formats (DOCX, PDF, CSV, PNG, MP4). A full export of an organisation's content is provided on request through your account manager or support@gimmeify.ai. Nothing is locked in a gimmeify-only format.

Deletion

Ask and we delete. Individual items can be deleted in-product. Organisation-level deletion is actioned on request and confirmed in writing.

Retention and exit

Content is retained for the term of the agreement. After termination it stays available for export for 30 days and is then permanently deleted, unless law requires otherwise. Backups age out on their normal schedule.

Residency

Your account is provisioned in one region and stays there. EU and UK customers: Frankfurt. AI processing is transient at the provider (Section 05); stored data does not leave your region.

08 Security operations

Secure development

All code is version-controlled. Changes reach production only through peer-reviewed pull requests that pass automated tests and checks, on a fixed promotion path (dev → main → staging → production). Access-control configuration, including every RLS policy, ships through the same path, so there is no route for an unreviewed permission change to reach production. Secrets are held in the platform's secret store and never in the repository or the browser.

Application security

- Origin allow-listing on every backend function.
- Sanitisation of any HTML or SVG rendered from model output (DOMPurify), and rendered previews run inside sandboxed iframes.
- Protection against server-side request forgery on every function that fetches URLs.
- Input validation with type, size and length limits; strict file-type allow-listing on uploads.
- Request throttling at the application layer, per provider and per job.

Monitoring and detection

Three layers. The cloud providers monitor their infrastructure under their own audited programmes. On our side, structured event logging with severity classification runs across the API layer, scheduled health checks exercise the platform's user journeys daily and raise tracked issues automatically on failure, and critical anomalies alert the engineering team in real time. Security advisors run against production on every review cycle and their findings enter the risk register.

Incident response

A documented process with severity classes and response commitments. Critical incidents (active exploitation, isolation failure, credential exposure) escalate to the founder immediately and remediation starts within 24 hours. Customers whose data is affected are notified without undue delay and within 72 hours of us becoming aware, with the information you need for your own regulatory notifications.

Risk governance

A risk register with owner, rating, treatment and target date for every open item. Monthly operations review of open items and quarterly management risk review chaired by the founder, who is the only person able to accept high or critical residual risk. Quarterly review of network and access controls. Security exceptions require written approval in advance and expire after at most 90 days.

Independent testing

An independent third-party penetration test of the current platform is being scheduled for Q4 2026, covering the web application, API layer, authentication, storage and the RLS perimeter. The report and remediation summary will be available under NDA. Until then, our assurance rests on the controls in this document, continuous automated security review, and the provider attestations in Section 02. We would rather tell you this than imply otherwise.

09 Subprocessors

Every third party that can receive customer data, what it receives, and where it processes. Nothing reaches an AI or data-collection provider unless a user chooses a capability that uses it. We update this list as the roster changes and will notify DPA customers of additions.

SUBPROCESSOR	PURPOSE	DATA IT RECEIVES	LOCATION
Supabase Inc.	Database, authentication, object storage, edge functions	All platform data, encrypted at rest	AWS Frankfurt (EU) or AWS Tokyo (APAC), per your region
Amazon Web Services	Underlying cloud, CDN	All platform data, encrypted at rest; web traffic	Frankfurt or Tokyo; CDN edge locations serve static assets only
Google Cloud	Media rendering services	Media being rendered	Frankfurt (EU); Singapore (APAC)

Anthropic, OpenAI, Google, xAI, Perplexity, DeepSeek, BytePlus, Kuaishou, Moonshot, MiniMax, Zhipu, StepFun, Alibaba Cloud, Manus, fal.ai, OpenRouter	AI model inference	The request content described in Section 05, only for capabilities the user invokes	United States / global APIs (transfers under SCCs in the DPA)
ElevenLabs, HeyGen	Voice and avatar generation	Scripts and reference media for that request	United States / global APIs
LangChain, Inc. (LangSmith)	AI request tracing	Prompts and outputs in traces, limited retention	United States
Apify, Firecrawl, Tavily, Browserless	Public web and social data collection and search, for research and listening capabilities	Search terms, brand names and URLs only; no customer documents	United States / global
Resend	Transactional email (invites, notifications)	Name and email address	United States
Stripe	Payments, direct customers only	Payment data (PCI DSS Level 1); never touches white-label client data	United States / global
Google Workspace, Zoho	Business operations (email, CRM, support)	Business contact details and correspondence	Global

10 Roadmap

Security work in progress. Items are listed so nobody mistakes them for things we already do.

ITEM	STATUS	TARGET
Independent penetration test of the production platform, report under NDA	SCHEDULING	Q4 2026
Web application firewall (AWS WAF managed rule sets and rate rules) on the CloudFront distribution	IN PROGRESS	Q3–Q4 2026
Cloud-account threat detection (Amazon GuardDuty) on our AWS accounts	IN PROGRESS	Q3–Q4 2026
Breached-password screening at sign-up and password change, both regions	IN PROGRESS	Q3–Q4 2026
Idle-session timeout and a session dashboard for users and admins	PLANNED	2027
Customer-facing audit log of user and admin actions	PLANNED	2027
Self-service organisation-level export and erasure	PLANNED	2027
Zero-data-retention agreements with the providers that offer them, platform-wide	EVALUATING	On client demand
Own SOC 2 Type II programme	EVALUATING	Decision after the penetration test

11 Questions from IT, security and legal teams

Are you SOC 2 or ISO 27001 certified?

No. Our infrastructure providers are (AWS: ISO 27001, SOC 1/2/3, PCI DSS; Supabase: SOC 2 Type II) and we share their reports under NDA. gimmefy's own controls are described in this document and in our policy documents, and we complete your security questionnaire.

Are you NIS2 compliant?

NIS2 does not apply to gimmefy directly; it applies to your organisation if you are an essential or important entity, and obliges you to manage supplier risk. We give you what that needs: a named security contact, 72-hour incident notification, documented risk management and secure development, a full sub-processor list, and a completed supplier questionnaire.

Do you use our data to train AI models?

No. gimmefy has no training pipeline. Providers are accessed through commercial APIs whose terms prohibit training on inputs and outputs (Anthropic, OpenAI and Google state this explicitly).

Exactly what leaves your systems when we use an AI feature, and where does it go?

The prompt, the relevant brand context retrieved for that task, any files attached to that request and the thread's history, sent over TLS to the provider behind the model the user chose. The provider returns the output and retains the exchange only within its abuse-monitoring window (30 days at Anthropic and OpenAI, 55 days at Google), then deletes it. Nothing from another tenant, and no credentials or billing data, is ever included. Section 05 has the full table.

Do you have zero-data-retention agreements with the providers?

Not yet. Standard commercial API terms apply today. For clients who need ZDR we will apply for it with the providers that offer it (Anthropic, OpenAI and Google do) and restrict the workspace to them.

Can we limit which AI providers our workspace uses?

Yes. Enterprise and white-label workspaces can be restricted to named providers. Anthropic, OpenAI and Google only is the usual request.

Where is our data stored, and does it stay in the EU?

In the region your account is provisioned in. EU and UK customers are in Frankfurt (AWS eu-central-1): database, authentication, files and media rendering. AI requests are processed transiently by the provider under the terms above; stored data does not leave the region.

Can another customer, or another client of our white-label partner, see our data?

No. Isolation is enforced by 900+ row-level security policies inside the database, scoped per partner, organisation and workspace. RLS is enabled on every table.

Is our data encrypted?

Yes. TLS 1.2+ in transit, AES-256 at rest, for the database, object storage and backups.

Do you support SSO and MFA?

SAML 2.0 SSO (Microsoft Entra ID, Okta, Google Workspace) configured per organisation, plus Google sign-in. TOTP MFA for any user, enforceable organisation-wide by an admin.

Who at gimmeFY can see our content?

Named engineers, with MFA, on a least-privilege basis, and only to resolve a support request you raise or an incident. Access is logged and reviewed quarterly.

Have you had a penetration test?

Not yet on the current platform. One is being scheduled for Q4 2026 and the report will be available under NDA. An earlier test covered a previous generation of the product and is not representative, so we do not circulate it.

What happens if there is a breach?

We notify affected customers without undue delay and within 72 hours of becoming aware, with what you need for your own regulatory notifications. Critical incidents escalate to the founder immediately and remediation starts within 24 hours.

Can we sign a DPA?

Yes. Our DPA covers GDPR Article 28 terms, EU Standard Contractual Clauses and the UK Addendum, the sub-processor list, 72-hour breach notification, deletion and audit cooperation. Where you contract through a white-label partner, the partner is the controller or processor of record and gimmeFY is the sub-processor; the DPA reflects that.

How do we export or delete our data, and what happens when we leave?

Outputs download in open formats at any time. Full organisation export and deletion are actioned on request through your account manager or support@gimmeify.ai. After termination, content remains exportable for 30 days and is then permanently deleted.

What happens if gimmeify is acquired or ceases trading?

Your agreement, including the deletion and export obligations, transfers with the business. Your files are in open formats and exportable at any time, so nothing depends on gimmeify continuing to exist.

Will you complete our security questionnaire?

Yes. SIG, CAIQ, vendor-specific spreadsheets, and follow-up calls with your information security team.

12 Contact and document control

Security and privacy contact: **support@gimmeify.ai** (reaches the founder and the head of engineering directly). For customers of a white-label partner, your partner is your first line and escalates to us.

VERSION	DATE	CHANGE
3.0	September 2026	Rewritten for enterprise security review. Added shared-responsibility model, regulatory alignment table (GDPR, NIS2, PDPA, CCPA, EU AI Act), per-provider retention and ZDR table, confidentiality controls, full sub-processor list with locations, honest roadmap. Refreshed numbers: 900+ RLS policies, RLS on every table, 50+ models, 200+ edge functions.
2.1	July 2026	Added EU (Frankfurt) region. Seven guarantees format.
2.0	June 2026	First edition for the current platform.

Teemuno Pte Ltd is a company registered under the Companies Act of Singapore, ACRA Registration No. 202226999N. gimmeify.ai is a product provided by Teemuno Pte Ltd. This document is provided for information, reflects the platform as of September 2026, and does not itself form a contract. Where a signed agreement or Data Processing Agreement exists, that agreement governs. Provider retention terms are summarised from each provider's published API terms as of September 2026 and may change; the provider's current terms prevail.