PM101

DATA & AI GUIDE

PM101
Data & AI Guide

How your information is stored, who can see it, and exactly what the AI does with it.

PLAIN ENGLISH FIRST. TECHNICAL ANNEX INCLUDED.

Contents

01

Your data

What PM101 holds, where it physically sits, and who is able to look at it.

SECTION 01 — YOUR DATA

What PM101 holds about you

Five kinds of data, each with a different purpose and a different lifespan. Nothing else is collected, and nothing here is shared between customers.

Account data

Email, hashed password, profile name and preferences. Passwords are never stored in readable form and are never visible to us.

Project data

Objectives, scope, milestones, tasks, dependencies, risks, stakeholders, budgets, rate cards and team records inside each project you create.

Business context

Documents you upload to a workspace: brand guides, policies, governance standards, methodology and templates. This is what makes AI output sound like your organisation.

Generated output

Drafts, exports, slideshows and assistant answers. Derived from your own input, stored against your account, deleted with it.

Usage data, kept deliberately thin

Anonymous page views with no personal identifiers, signed-in feature usage for fair-use limits, and Stripe billing records. No advertising trackers, no data brokers, no profile building.

SECTION 01 — YOUR DATA

Where it physically lives

One managed platform, EU region, private storage. Every layer below is operated by a named provider under contract, not by anonymous infrastructure.

App

React single-page app, served from the edge

Runs in your browser. Holds no secrets and no service credentials.

API

Deno edge functions

All privileged work — AI calls, exports, deletion — happens here, server-side, behind authentication.

Data

Postgres with row-level security, EU region

Every table gated by policy. A query that is not yours returns nothing, not an error you can probe.

Files

Four private storage buckets

project-sources, business-context, user-datasets, slideshow-assets. None are publicly browseable.

Who runs the layers

Supabase provides the database, authentication, storage and function hosting in the EU region. Stripe processes payments. AI requests are routed by the Lovable AI Gateway. Those are the only parties that touch your content.

Why buckets are private

Public buckets are the single most common cause of document leaks in SaaS. PM101 has none. Files are served only through short-lived signed URLs issued after an access check against your session.

SECTION 01 — YOUR DATA

Data residency: what sits where

Your documents and your database records stay in one place: Amazon Web Services in Ireland (eu-west-1). Everything else on this page is the honest detail about the few things that travel further, and why.

What	Where it is stored	Leaves the EU?
Your account and profile	AWS Ireland (eu-west-1)	No
Projects, plans, risks, generated documents	AWS Ireland (eu-west-1)	No
Uploaded files and business context	AWS Ireland, private buckets	No
Database backups	Same region as the database	No
Payment and subscription records	Stripe (global infrastructure)	Yes, to Stripe only
AI prompt and retrieved context	Not stored by the model; routed via the AI gateway	Possible, provider-dependent
Voiceover text (only if you make a video)	ElevenLabs (US)	Yes, if you use that feature
Web page fetches (news feed, URLs you paste)	Firecrawl (US)	Yes, for that URL only

In one sentence

Your stored data lives in Ireland. The only content that can cross a border is the text of an AI request while it is being answered, and the optional voice and web-fetch features — each of which you choose to use.

SECTION 01 — YOUR DATA

Who is responsible for the data

In data protection law there are two roles. The one who decides what is collected and why (the controller), and the one who handles it on their behalf (the processor). For anything you upload, you are the controller and PM101 is the processor.

You are responsible for

- Deciding what client or case material is appropriate to upload
- Having a lawful basis and client consent where your own duties require it
- Who you invite into a workspace or project, and at what level
- Removing material when your own retention period ends
- Answering your client if they ask where their information is held

PM101 is responsible for

- Holding your content only in Ireland, encrypted, in private storage
- Processing it only to deliver the service you asked for
- Never using your content to train models or to serve another customer
- Keeping named sub-processors under contract and publishing the list
- Deleting on your instruction, and telling you promptly if anything goes wrong

Where PM101 is the controller

Only for our own relationship with you: your sign-up email, your billing record, and thin usage counts for fair-use limits. Never for the case files, client records or project material you put into a workspace.

Putting it in writing

A written data processing agreement naming the sub-processors, the purpose limits and the deletion commitments is available on request, so your firm can file it alongside your other supplier records.

SECTION 01 — YOUR DATA

Who can see it

Access is a ladder, not a switch. Each rung has to be granted deliberately, and the default at the bottom is nobody.

You

Full access

Read, edit, export and delete anything you own.

Team

People you invite

Workspace and project members see only what is shared with them, at the role you set: viewer or editor.

Support

Time-boxed, only when you grant it

A named support grant, scoped to one project, expiring inside 24 hours, with every access written to an append-only log.

Nobody

Everyone else, including other customers

Cross-tenant reads are blocked at the database, not in application code. There is no admin back door into your project content.

The important word is 'default'

A new project is visible to exactly one person. Sharing is an action you take, never a setting you have to remember to turn off.

SECTION 01 — YOUR DATA

Support access: 24 hours, logged

When you raise a problem we cannot reproduce, you can let support look — on a clock, on the record, and only at the project in question.

T+0

You grant access

You choose one project and confirm. No grant, no access.

T+0

Scope is fixed

The grant covers that project only. It does not extend to your other projects, your workspace documents or your account.

Live

Every look is written down

Each access is recorded in an append-only audit table that support cannot edit or remove.

T+24h

It expires by itself

Expiry is enforced in the database, not by someone remembering to revoke it. You can also revoke instantly.

Why time-boxing matters

Standing internal access is how support tooling quietly becomes a permanent window into customer data. A grant that dies on its own cannot be forgotten about.

What support cannot do

Support cannot read your password, cannot mint an access grant for itself, and cannot reach a project you have not opened up. The check runs in the database on every request.

02

How the AI works

Where your prompt goes, what travels with it, and what comes back.

SECTION 02 — HOW THE AI WORKS

The Context Engine, security view

The same three layers that make PM101's output useful also define its blast radius. The model only ever sees what these three layers hand it.

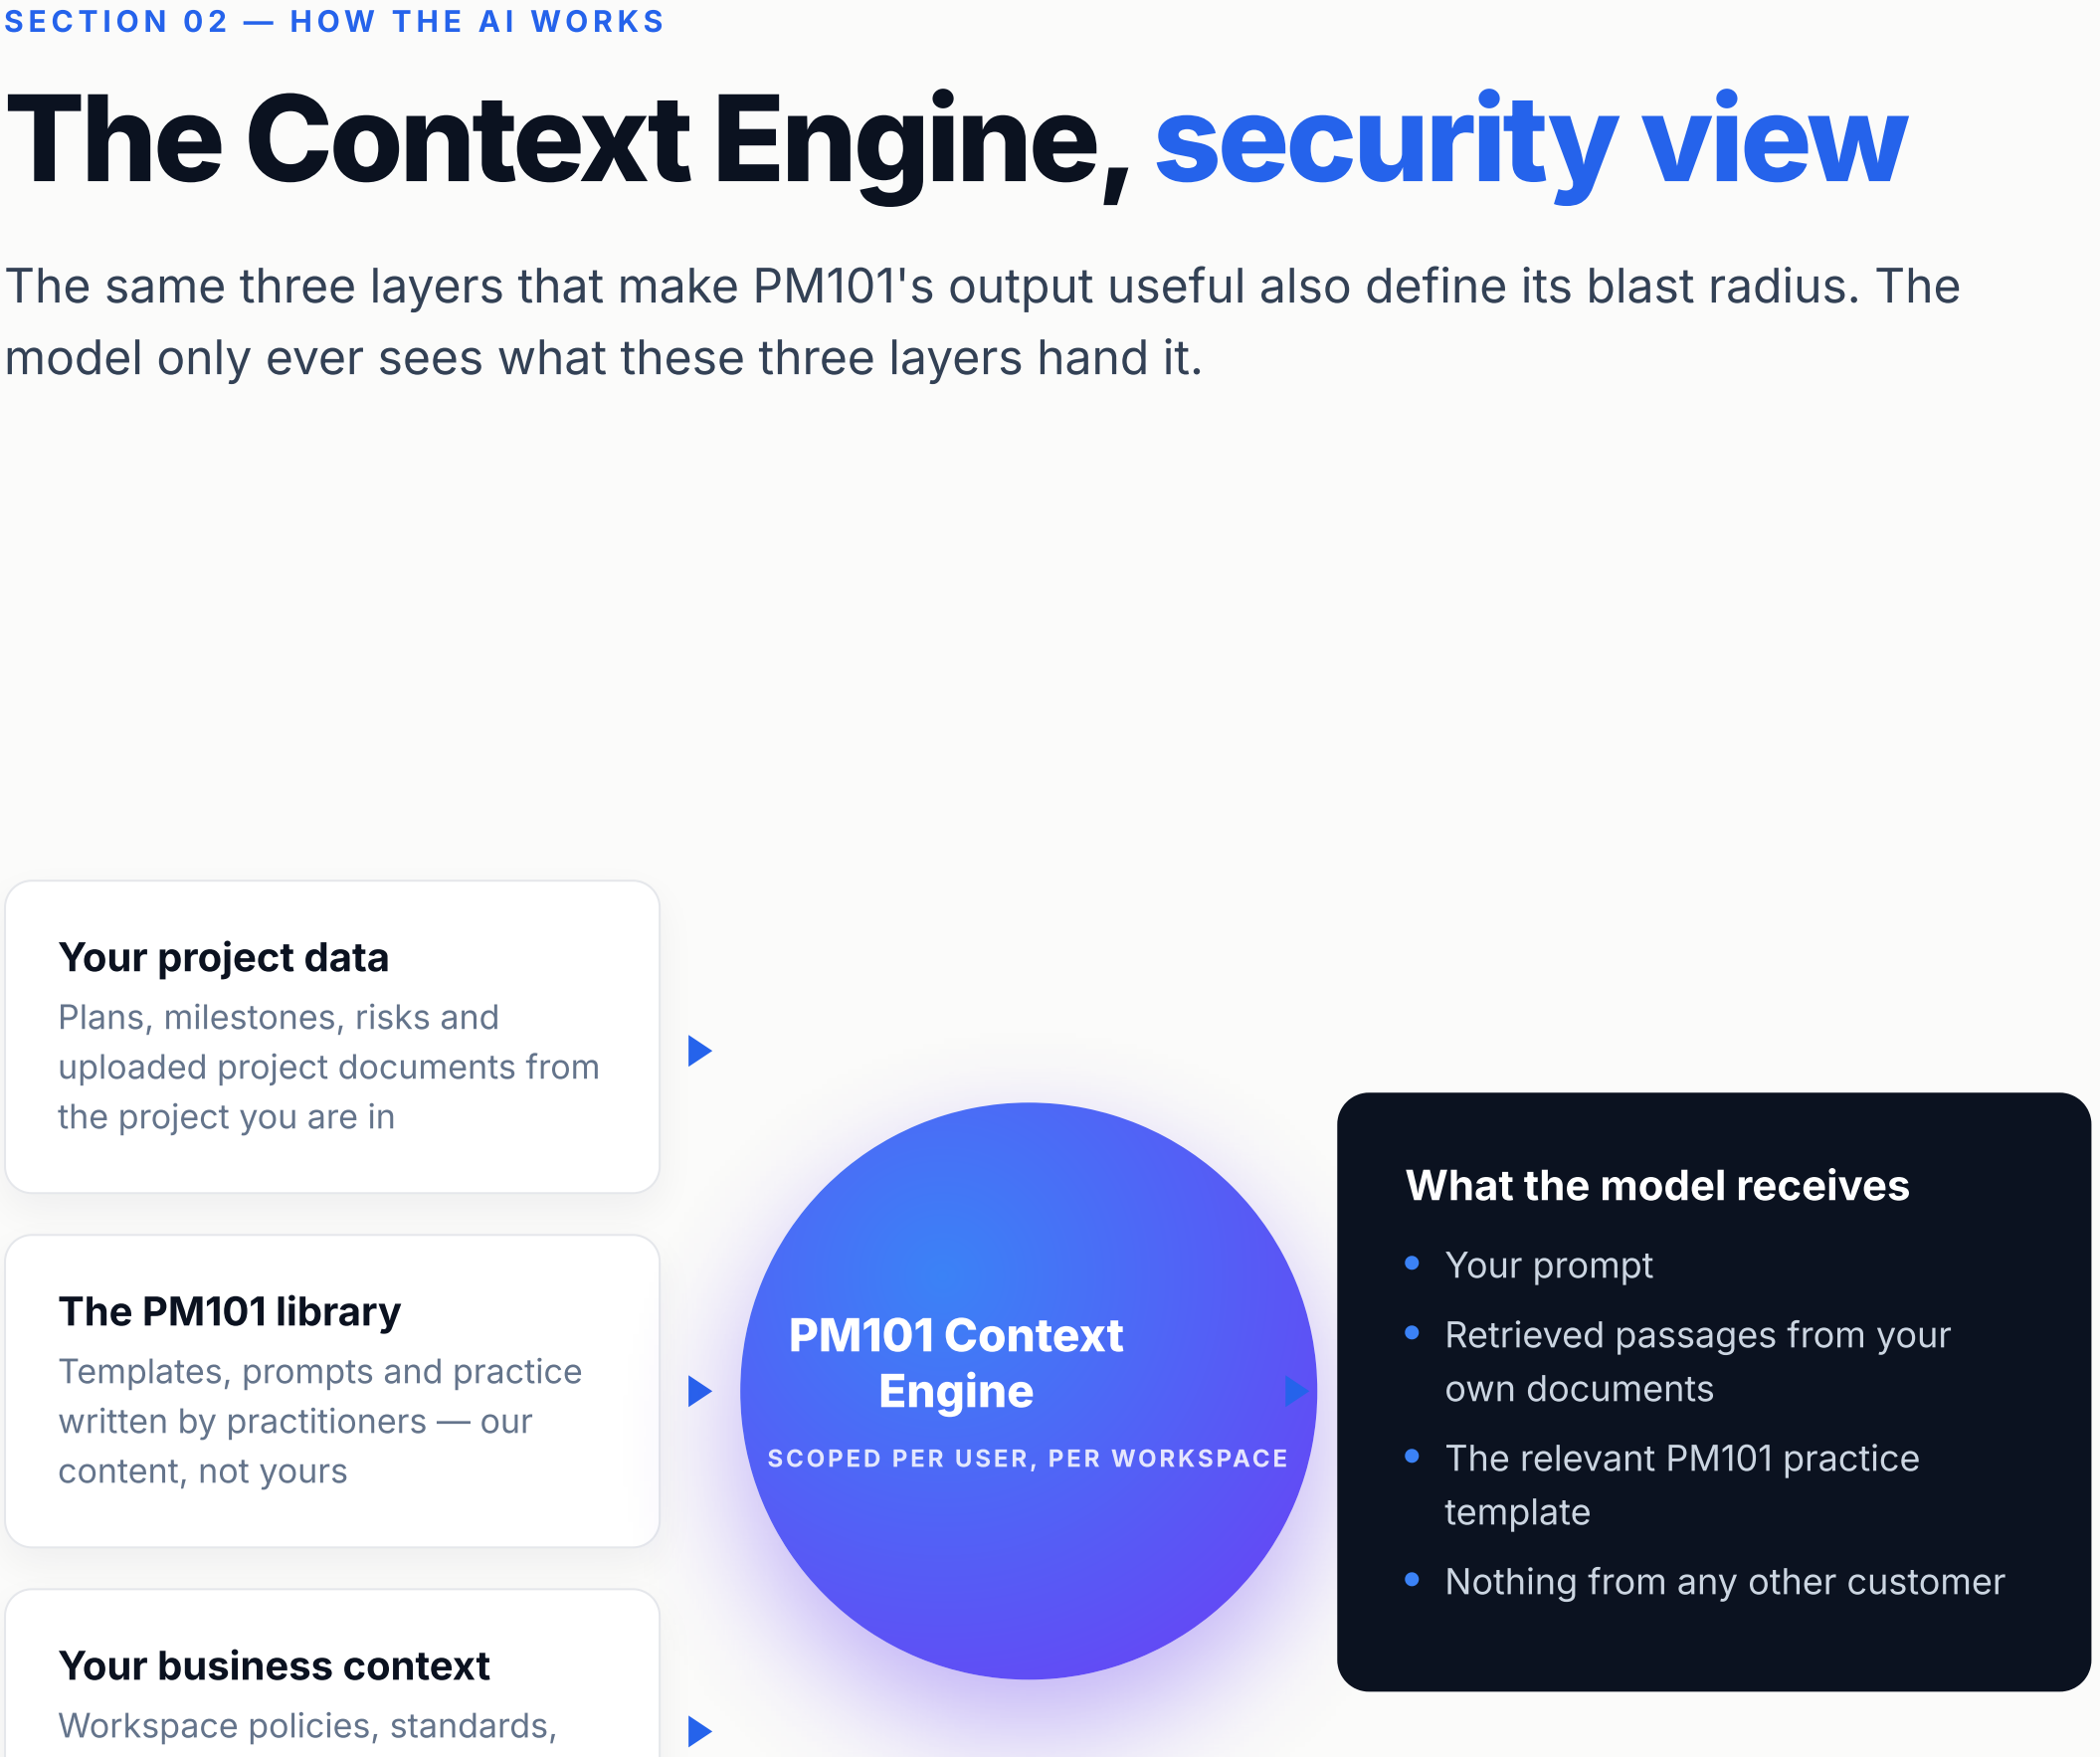

Scoping is enforced below the AI, not by it

Retrieval queries carry your user and workspace identifiers and run against row-level-security policies. If the model asked for another tenant's data, the database would return an empty set.

SECTION 02 — HOW THE AI WORKS

What happens to a prompt

Five steps, all server-side. Follow one request end to end.

01

You submit

The browser sends your prompt to a PM101 edge function with your session token. No AI keys exist in the browser.

02

We verify

The function checks your session, your workspace membership and your plan limits before doing anything else.

03

We retrieve

A scoped search pulls the most relevant passages from your documents only, ranked and truncated to a context budget.

04

We call the gateway

Prompt plus retrieved context goes to the Lovable AI Gateway, which routes to the current best-fit model.

05

You get output

The response is returned to you and stored against your account so you can re-open, export or delete it.

Keys never leave the server

Every AI call originates in an edge function. There is no path by which a browser, a shared link or a scraped bundle exposes a provider key.

Context has a ceiling

Retrieved context is capped before it is sent. Large document sets are ranked and trimmed rather than shipped wholesale, which limits both cost and exposure.

SECTION 02 — HOW THE AI WORKS

What the AI never does

Plain answers to the four questions every buyer asks, with the honest boundary marked where it belongs.

What it does

- ✓ Reads your own documents to answer your question
- ✓ Drafts artefacts in your house format
- ✓ Shows you the system prompt before you run an agent
- ✓ Stores its output against your account, under your control
- ✓ Runs only when you ask it to

What it never does

- — Train on your content through any PM101 path
- — See another customer's workspace
- — Act on your data without an explicit action from you
- — Run in the browser with a provider key
- — Send your documents anywhere except the model handling your request

The honest boundary

PM101 has no training pipeline: your content is used for retrieval and generation only. What a model provider does with API traffic is governed by that provider's own terms, not by our code. We state what our system does, and we do not claim contractual guarantees on your behalf.

SECTION 02 — HOW THE AI WORKS

The model pool

PM101 does not hard-wire itself to one vendor. Requests go to a gateway that selects the current best-fit model for the job, so the platform can improve without a rebuild.

Job	Currently routed to	Where it runs	Your data
Drafting and assistant answers	Google Gemini (text)	Edge function → gateway	Prompt + your retrieved passages
Slide and cover imagery	Google Gemini (image)	Edge function → gateway	Your prompt text only
Voiceover	ElevenLabs	Edge function → gateway	The script text only
Scoring and evaluation	Gemini Flash-Lite class	Edge function → gateway	The item being scored

Why a pool, not a vendor

Model quality, price and availability move every few months. Routing through a gateway means we can move to a better model without changing how your data is scoped, stored or retrieved.

What this means for review

Treat the table as the current state, not a permanent commitment. If your procurement needs a fixed model list at a point in time, ask and we will confirm the pool in writing on the day.

SECTION 02 — HOW THE AI WORKS

Grounded in your documents

Generic AI knows everything and nothing about your project. Grounding is what turns a plausible draft into one your governance board recognises.

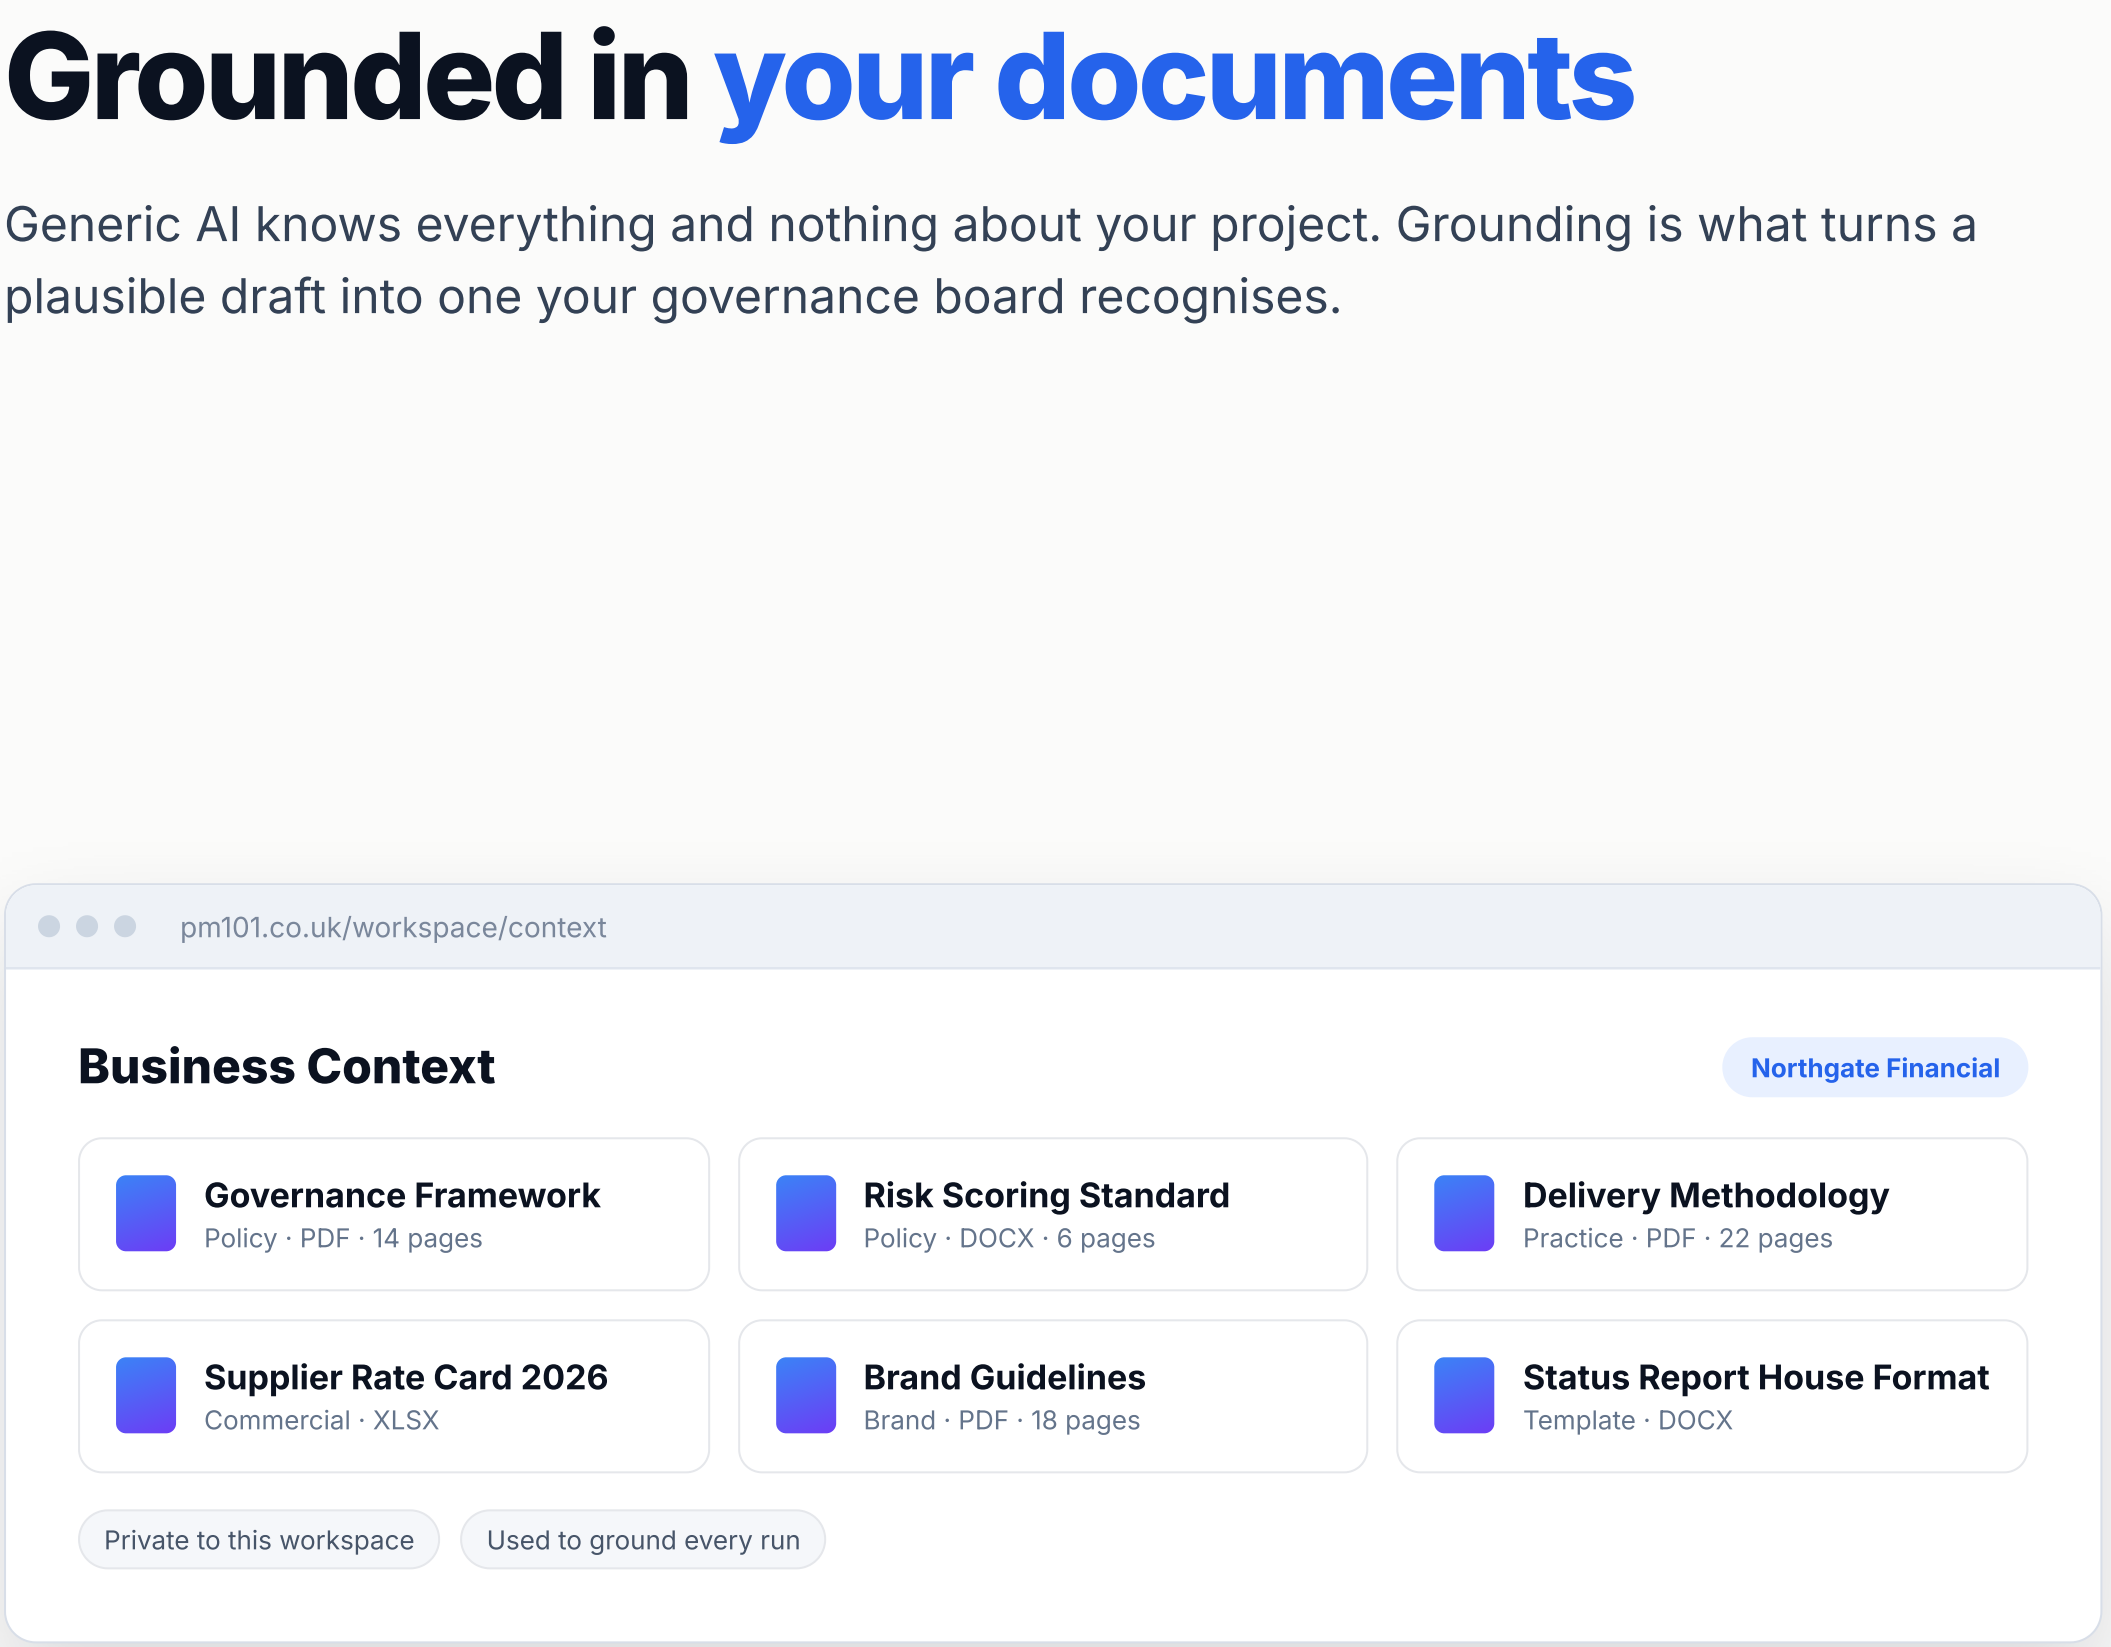

Workspace → Business Context: documents, rate cards and brand pack that ground every run

How retrieval is scoped

Uploaded documents are split into chunks stored against your user and dataset identifiers. The matching function filters on both before ranking, so a search cannot reach outside your own material.

What grounding buys you

Your terminology, your stage gates, your risk scoring, your tone. The difference between a draft you edit for an hour and a draft you send.

03

Your controls

The switches you hold: what the AI knows, what it is told, and what you can take away.

SECTION 03 — YOUR CONTROLS

Business Context and brand packs

You decide what the engine knows about your organisation. Upload once, and every generated document inherits both the substance and the styling.

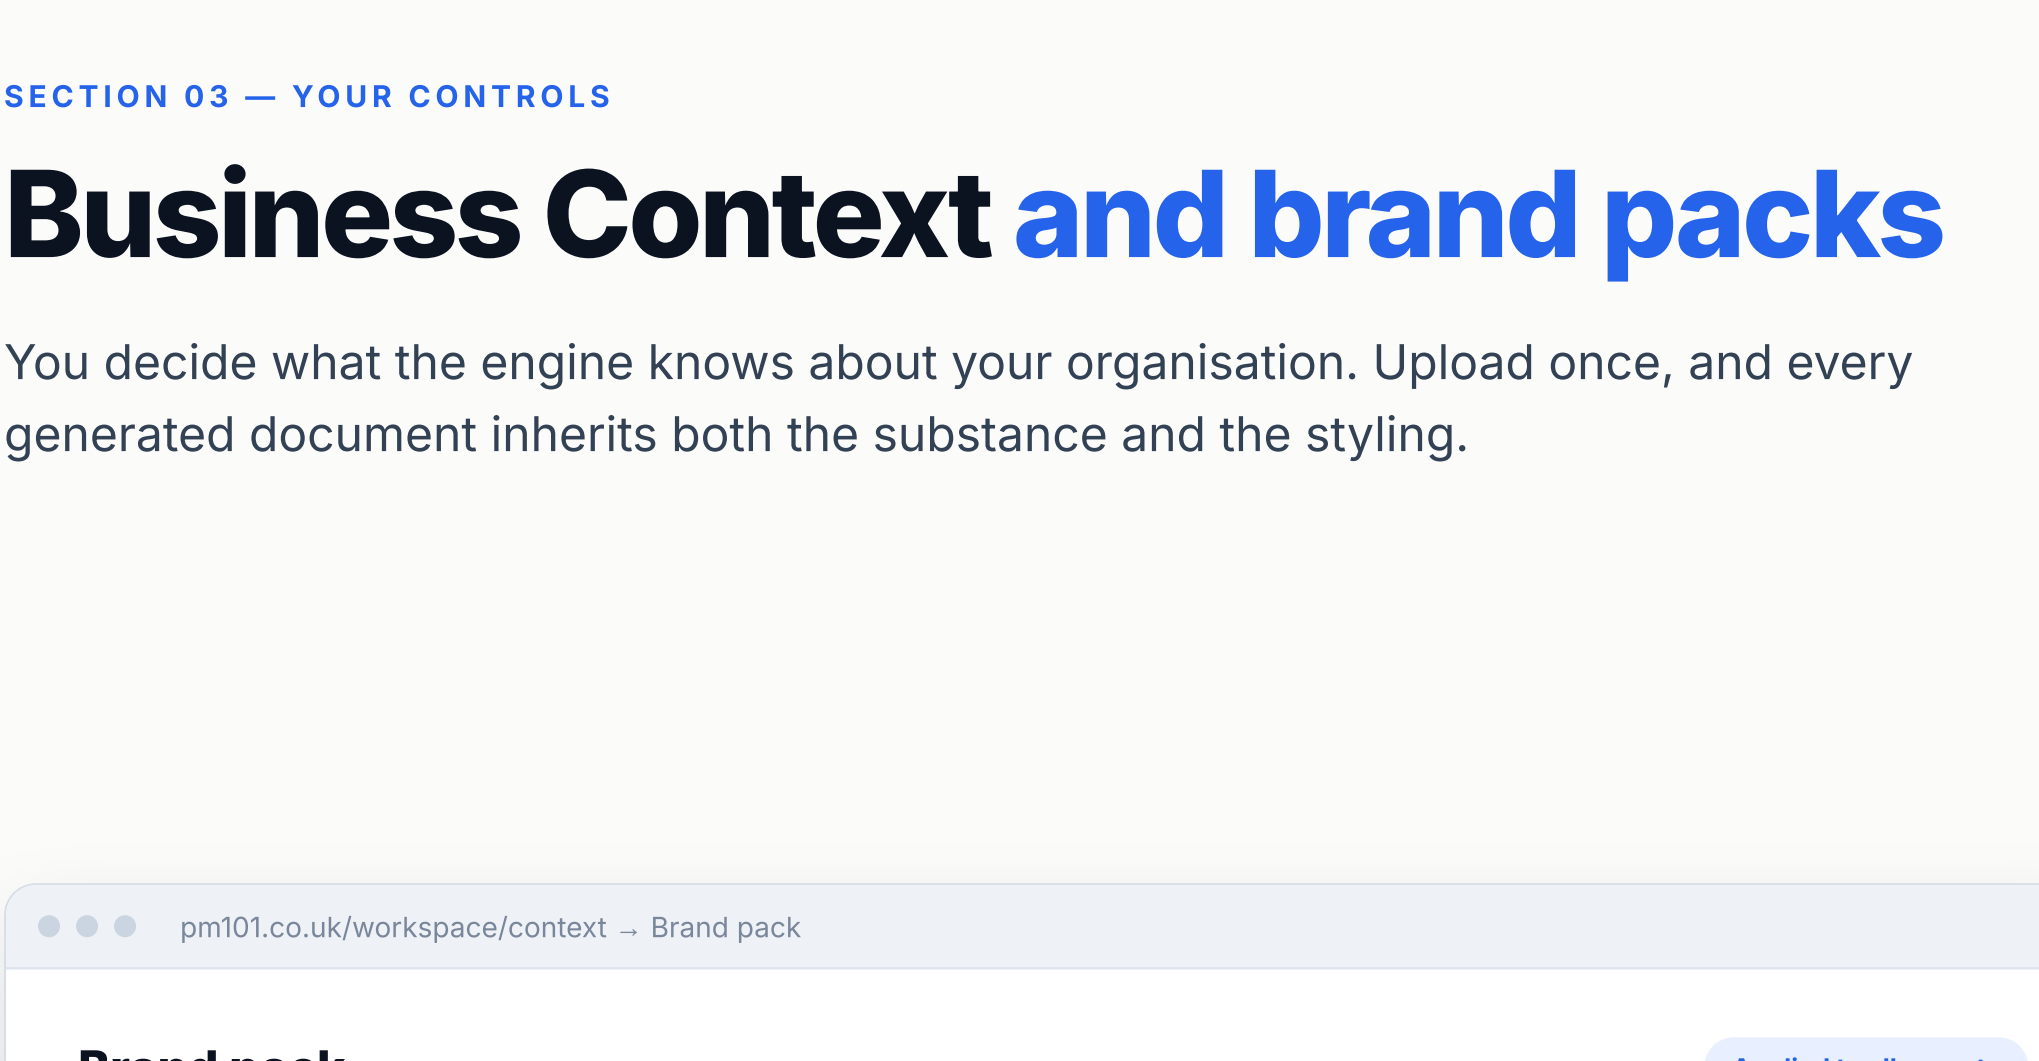

A workspace brand pack: colours, fonts, logo and footer applied to every export

Substance

Policies, governance standards and methodology shape what the document says.

Style

Brand colours, fonts, logo and footer shape how the PDF, Word and PowerPoint exports look.

Scope

Stored per workspace, private by default, deleted when the workspace is deleted.

SECTION 03 — YOUR CONTROLS

Glass-box agents

Most AI products hide the instructions. PM101 shows you the system prompt before you run anything, because you cannot govern what you cannot read.

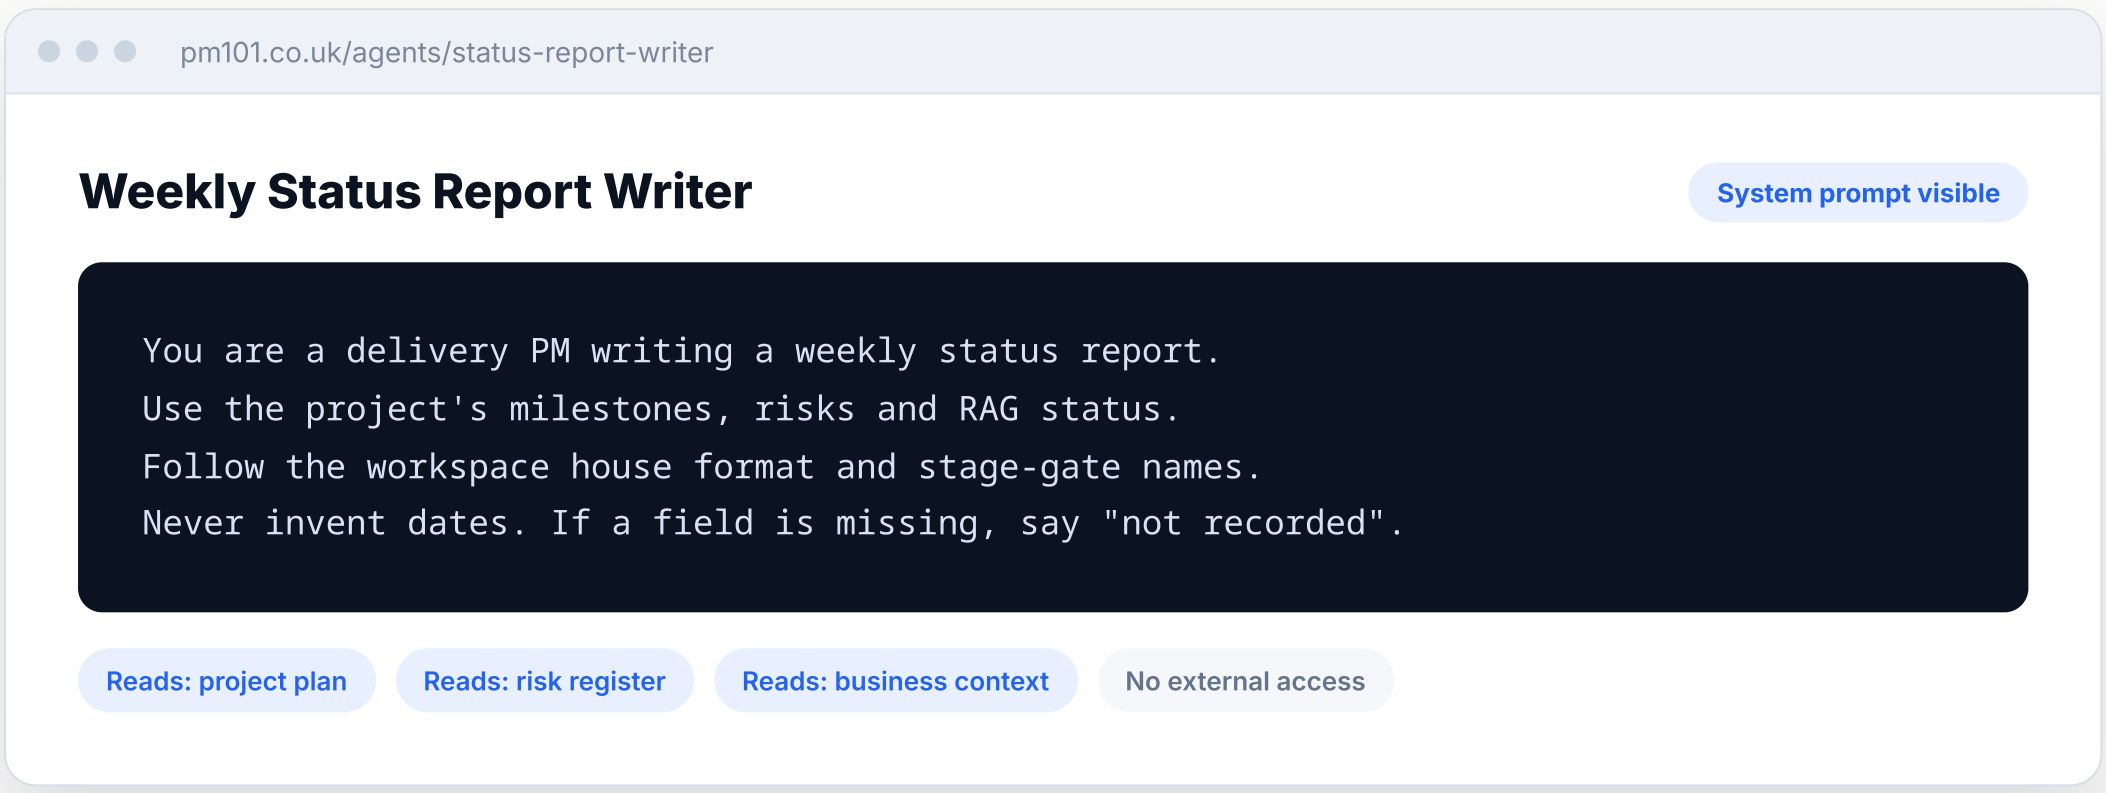

Agent detail: the full system prompt and data sources, visible before you run it

Why we show the prompt

A hidden prompt is an unreviewable control. If an agent will read your risk register and write to a status report, you should be able to see that stated in plain text first.

You choose the inputs

Each agent declares the data sources it reads. Nothing is attached silently, and an agent cannot reach a project you have not given it.

SECTION 03 — YOUR CONTROLS

Export, and delete

Getting data out and taking data away are both self-serve. Neither requires a conversation with us.

01

Export a document

Any generated artefact downloads as PDF, Word, PowerPoint or spreadsheet, branded to your workspace.

02

Delete a project

Project → Settings → Danger zone. Removes the project, its plan, its documents and its uploaded sources.

03

Delete your account

Profile → Danger zone. Type DELETE to confirm; the purge then runs server-side.

What account deletion removes

Profile, projects, workspaces you own, uploaded documents, datasets, generated exports, assistant history, usage records, and the storage objects behind all four private buckets.

What it does not do

It does not cancel billing disputes or issue refunds automatically, and operational database backups roll off on their own retention cycle rather than being edited retrospectively.

SECTION 03 — YOUR CONTROLS

Free and Pro: what changes

Plan changes what you can do, never how your data is protected. Security posture is identical on every tier.

	Free	Pro
Row-level security on your data	Yes	Yes
Private storage buckets	Yes	Yes
Time-boxed, audited support access	Yes	Yes
Self-serve export and deletion	Yes	Yes
AI runs per month	Limited	Included
Business Context and brand packs	—	Yes
Project plan, Lens and reporting	—	Yes

No security tax

Encryption, access control, audit logging and deletion rights are not upsells. They apply to the free tier exactly as they apply to a paid workspace.

SECTION 03 — YOUR CONTROLS

If you hold client-confidential data

Law firms, accountants and advisers carry duties that go beyond data protection: confidentiality, privilege and professional rules. Here is how to run PM101 inside them.

01

One client, one workspace

Workspaces are hard boundaries. Members of one cannot see, search or retrieve anything from another, so a matter for Client A can never surface in a draft for Client B.

02

One matter, one project

Retrieval is scoped to the project you are in. Keep engagements in separate projects and the AI can only draw on that engagement's material.

03

Invite deliberately

Add the engagement team, at viewer or editor. Removing someone removes their access immediately, including to documents already generated.

04

Close the file

Delete the project when your retention period ends. That removes the plan, the documents and the uploaded source files together.

Worth knowing before you upload

The text of an AI request is sent to the model provider to be answered. It is not stored there, not trained on and not shown to anyone else — but if your duty of confidentiality means a document must never leave your own systems at all, keep it out of the AI features and use PM101 for planning and reporting only.

Reducing what you send

You do not need the whole file. Business Context holds your house standards and templates; the AI then needs only the specific passages relevant to the question, which is all retrieval sends.

What you can tell a client who asks

"Our project material is held in a UK supplier's system, stored in Ireland, encrypted, visible only to our engagement team, deleted on our instruction, and never used to train anyone's AI." Every part of that sentence is checkable in this guide.

04

Technical annex

For the security reviewer: the controls, stated precisely enough to be checked.

SECTION 04 — TECHNICAL ANNEX

Authentication and access model

Identity, roles and privilege separation, as implemented.

Identity

Email and password with mandatory verification, plus Google and Apple federated sign-in. No anonymous sign-up path.

Sessions

JWT-based sessions issued by the auth service. Edge functions verify the token on every privileged call.

Roles

Roles are held in a dedicated user_roles table and checked through a security-definer function — never stored on the profile row, which would allow privilege escalation.

Privilege separation

The browser holds only the publishable anon key. Service-role credentials exist solely inside edge functions and are never returned in any response.

Why roles live in their own table

A role column on a user-editable profile row is an escalation waiting to happen. Separating the table and reading it through a security-definer function keeps the check outside the user's write surface.

SECTION 04 — TECHNICAL ANNEX

Row-level security posture

Authorisation is enforced in the database. Application bugs cannot hand out rows the policies do not allow.

55/55	0	4	24h
public tables with RLS enabled	tables reachable without a policy	private storage buckets	maximum support access grant

Ownership model

Policies key off ownership and membership: the row's owner, the project's collaborators, or the workspace's members. Anything else evaluates false and the row is invisible.

Grants, not just policies

Every public table carries explicit privilege grants alongside its policies, so access is defined twice over — at the privilege layer and at the row layer.

Verified, not asserted

RLS coverage is checked by a repository script rather than claimed in a document, and cross-tenant isolation is exercised by an end-to-end test that attempts a read across accounts and expects nothing back.

SECTION 04 — TECHNICAL ANNEX

Storage and retrieval

Four buckets, all private, each with a path convention that scopes access by identity.

Bucket	Holds	Path scope	Public
project-sources	Documents uploaded to a project	user / project	No
business-context	Workspace policies, standards, brand assets	workspace	No
user-datasets	Datasets used for AI retrieval	user	No
slideshow-assets	Generated images and audio	user	No

Serving files

Objects are fetched through signed URLs issued after an access check. Bucket listing is disabled, so an object cannot be discovered by guessing a directory.

Retrieval chunks

Documents are chunked for search with user and dataset identifiers attached. The matching function filters on those columns before ranking, keeping retrieval inside your own material.

SECTION 04 — TECHNICAL ANNEX

Retention, deletion and sub-processors

What is kept, for how long, and who else is involved.

Data	Retention
Active account, projects and documents	For as long as the account exists
Deleted project	Removed with its plan, documents and uploaded sources at deletion
Deleted account	Profile, owned workspaces, files and records purged server-side at confirmation
Anonymous analytics	Retained without personal identifiers for product improvement
Operational backups	Roll off on their own recovery retention cycle

Sub-processor	Purpose
Supabase	Database, authentication, private storage and edge functions (EU region)
Lovable AI Gateway	Routes AI requests to the current model pool
Stripe	Payments and subscription management
Email delivery provider	Transactional and opt-in email

PM101

Questions we have not answered here?

If you need a data-processing addendum, a point-in-time model list, or want to report a security concern, write to hello@pm101.co.uk and you will get a direct answer rather than a form response.

pm101.co.uk/data-ai-guide