Prism Profiles

prismprofiles.com/trust

VENDOR SECURITY QUESTIONNAIRE · ISO/IEC 27001 QUESTION SET

Our answers to the questions your security team will ask.

48 questions across the seven management areas of ISO/IEC 27001, answered for the Prism Profiles service. We have answered plainly, including where the answer is "not yet". Prism Profiles does not hold an ISO 27001 certificate.

Completed 17 November 2026 · My Ai Exec Pty Ltd, trading as Ai Tech Advisory · ABN 93 676 536 511 · info@prismprofiles.com

48	16	27	5
questions answered	Yes: in place, with evidence	Partly: done in practice, not yet formal or complete	Not yet in place

Yes The control is in place and we can show evidence.

Partly We do this in practice, but it is not a formal documented process, or it is not complete.

Not yet Not in place today.

How to read this: ISO/IEC 27001 asks how security is **managed**, so many questions are about policies and procedures. Prism Profiles is a young service run by a small team that includes a cyber security specialist. Our technical controls are strong and tested; our formal paperwork is thin. The answers say which is which. The questions follow the topics of a widely used free ISO 27001 vendor questionnaire, written in our own words.

1. Context of the organisation

NO.	QUESTION	STATUS	OUR ANSWER
1.1	**Would you take part in an audit at short notice?**	Partly	Yes, as a remote audit. We are a small company with no data centre of our own: the service runs on Railway in the United States. On reasonable notice we will walk your auditor through the controls on a call, show the live System Health page and configuration by screen share, and send the evidence behind any answer here. Audit terms can be written into an Enterprise agreement.
1.2	**Do you keep audit logs for the servers that hold data, and for backups?**	Partly	The application keeps an audit log of sign-offs, deletions, role and settings changes, every staff 'log in as', and every backup, restore and encryption-key action. The account owner can read their own log in the app. Every email sent is logged. Each nightly backup is recorded and read back to check it. Server and network logs are those kept by our hosting provider. We do not run a central log-monitoring service.
1.3	**Can you show where confidential data is held?**	Yes	Yes. The application, its database and file storage run on Railway, US West (California, USA). Nightly backup copies are held with Cloudflare (R2 storage). Text sent for AI analysis is processed by Anthropic in the United States. If you need your data kept in your own country, we can set up a server and host it there.
1.4	**Have you identified the internal and external issues that affect your security goals?**	Partly	In practice, yes: three written security reviews in 2026 (July, and two in October) set out the issues and each was followed by a fix programme. We do not keep a formal register of context and interested parties.
1.5	**Is the scope of your security management system defined?**	Partly	The scope we work to is the PRISM Profiles application, its database, file storage, backups, the calls it makes to AI providers, and staff access to the admin area. It is not written up as a formal scope statement.
1.6	**Do you have a system for managing information security?**	Partly	We do not hold a certified management system. What we operate: periodic written security reviews, every finding given a number and fixed against it, security checks that run automatically on every code change, a live System Health page, nightly backups, and a dated log of every decision and change.

1.7	**Is your security policy put into practice?**	Partly	The rules are enforced in the software and the build rather than in a signed policy document. Examples: the build fails if a secret key appears in code or in the browser, if an admin action can be reached without signing in, or if a high-severity vulnerability is found in a dependency; the service will not start in production without its encryption key.
1.8	**Do you have a policy for maintaining the management system?**	Not yet	No written maintenance policy yet.

2. Leadership

NO.	QUESTION	STATUS	OUR ANSWER
2.1	**Has leadership set a security policy and objectives?**	Partly	The founder sets the objectives and they are published on our trust page: client know-how stays private to the client, is never used to train AI, is encrypted, and can be exported or deleted by the client. There is no formal board-approved policy.
2.2	**Are security objectives aligned with business strategy?**	Yes	Yes. They are product design rules, not add-ons: PRISM Profiles never connects to a client's systems, each company's data is separated on the server, and AI providers are only used on terms that forbid training on client data.
2.3	**Are security requirements built into day-to-day processes?**	Yes	Yes. Every change is type-checked, run through the full test suite including the security tests, and audited for vulnerable dependencies before it can be released. Database changes follow a guarded procedure. Nothing goes to the live service without the owner's explicit instruction, and every release is recorded.
2.4	**Are the resources needed for security made available?**	Yes	Yes. A cyber security specialist on our team tests and hardens our systems, including penetration testing. Security fixes are scheduled as build work ahead of new features: six security programmes were delivered between July and November 2026.
2.5	**Is the importance of security management communicated?**	Partly	The team is small and the rules are written into the project's working instructions, which every contributor follows. There is no formal awareness or training programme yet.
2.6	**Can you show that security management meets its objectives?**	Partly	The evidence is the review reports with their findings closed, the automated test results on every build, and the System Health page. We do not hold formal management-review meetings with minutes.
2.7	**Does leadership support continual improvement?**	Yes	Yes. Each review has been followed by a funded fix programme, and we publish what we have not built yet instead of hiding it.
2.8	**Are security roles and responsibilities assigned and communicated?**	Partly	The founder is accountable for security. The highest level of admin access is limited to named people. A formal responsibility chart has not been written.

3. Planning

NO.	QUESTION	STATUS	OUR ANSWER
3.1	**How do you identify security risks?**	Partly	Through periodic written reviews of the code and configuration. The October 2026 review examined all 244 application actions for sign-in and company-separation checks, and a separate 36-point security checklist was run. Our cyber security specialist also tests the systems directly, including penetration testing. Automated checks run on every change. We do not keep a formal risk register.
3.2	**How do you fix the risks you find?**	Yes	Each finding gets a number, a fix and, where possible, an automated test that would fail if it came back. The July 2026 review raised 1 critical and 12 high findings: all were fixed. The October 2026 review raised 18 findings: all have been addressed, two of them in part (see the list at the end).
3.3	**Do you have risk-acceptance criteria for third parties?**	Not yet	Not as a written standard. The rule we apply in practice: an AI provider may only receive client data under business terms that forbid training on it, and any new provider is switched off by default until that is confirmed.

3.4	**Do you have a repeatable framework for assessing vendor risk?**	Not yet	Not yet. We use a small number of well-known providers (listed at the end) and review their data terms when we adopt them.
3.5	**Do you apply risk assessments and track progress?**	Partly	Yes for the application: findings are tracked to closure in the change log with dates. Not yet for the organisation as a whole.
3.6	**Do you measure the likely impact of a risk?**	Partly	Findings are rated Critical, High, Medium or Low by what an attacker could reach. We do not score impact in money or with a formal matrix.
3.7	**How do you decide how severe a vulnerability is?**	Yes	By the same four-level rating for our own code, and by the published severity for third-party components. A high or critical vulnerability in a component we ship blocks the release.
3.8	**How do you prioritise critical risks, including those from vendors?**	Partly	Critical and high findings are fixed first, before new features. Vendor risks are handled as they are found and are not formally ranked.
3.9	**How do you choose controls?**	Partly	From the review findings, the 36-point checklist, and seven standing checks on the most common ways web applications are broken into. We have not produced an ISO 27001 Statement of Applicability.
3.10	**How are mitigation efforts reported to leadership and other stakeholders?**	Yes	The founder receives every review and signs off each release. Clients can read our trust page, which includes what is still missing, and plain-language release notes.
3.11	**Do you have a disaster recovery plan?**	Partly	The database and all files are copied every night to a separate provider and kept for 30 days, with a written restore procedure. A full restore from a nightly backup into this service was carried out and checked on 11 November 2026. We have not set formal recovery-time targets; in the worst case up to one day of changes could be lost.
3.12	**Do you have an incident response plan?**	Partly	We have alerts and a working practice, not a formal tested plan. Staff are emailed automatically if the AI service stops working, and System Health shows failures. An outage on 11 November 2026 was found, fixed within about 15 minutes and written up. Roles, escalation steps and rehearsals are not yet documented.
3.13	**How are people notified when an incident happens?**	Partly	Customers: we notify affected customers within 72 hours of confirming a breach, as our terms state. Staff: automatic email alerts. We do not run a 24-hour on-call rota.

4. Support

NO.	QUESTION	STATUS	OUR ANSWER
4.1	**Do the people responsible for security have the resources they need?**	Yes	Yes. Our cyber security specialist has access to test and harden every system we run, and security fixes are scheduled ahead of new features.
4.2	**Do staff know the security policy?**	Partly	Everyone with admin access must sign in with their own password and an authenticator-app code, and their sign-in ends after 12 hours, so the main rule is enforced by the system. Formal policy training is not in place.
4.3	**How is sensitive documentation protected?**	Yes	Keys and passwords are held only in our hosting provider's encrypted settings and a password manager, never in code: the build checks for this. Client free text is encrypted in the database and uploaded files are encrypted in storage. Source code is in private repositories. Staff access is limited by role and to the accounts assigned to them.

5. Operation

NO.	QUESTION	STATUS	OUR ANSWER
5.1	**How do you manage your attack surface?**	Yes	It is deliberately small: one web application and no connections into client systems. Every action that works without signing in is on a reviewed list, and the build fails if a new one appears. Request sizes are limited, sign-in and AI use are rate-limited, uploads are type-checked, and outbound requests cannot reach internal addresses.

5.2	**How do you detect threats?**	Partly	System Health checks, automatic alerts when the AI service fails, sign-in rate limiting, reports from the browser security policy, and the audit log. We do not run intrusion detection or 24-hour monitoring.
5.3	**Do you keep track of outsourced processes?**	Yes	Yes. Hosting, AI processing, backups, payments and email are outsourced to the providers listed at the end, each for a stated purpose.
5.4	**Do you use other security frameworks or questionnaires?**	Partly	An internal 36-point security checklist and the seven standing checks described above. We hold no certification against another framework.
5.5	**Do you assess vendor risk at planned intervals?**	Not yet	Not on a schedule yet.
5.6	**Is a risk assessment triggered when something significant changes?**	Partly	Reviews were run before the two largest changes of 2026 (opening to new customers and moving client data between services). Day to day, new database fields and new public actions cannot be released without passing automated guards. This is practice, not a written trigger policy.
5.7	**Do you keep the results of risk assessments?**	Yes	Yes. Every review report is kept, dated, in version control with its full history.
5.8	**How long are completed assessments kept?**	Partly	Indefinitely at present. A retention period has not been set.

6. Performance evaluation

NO.	QUESTION	STATUS	OUR ANSWER
6.1	**Do you evaluate how well security management is working?**	Partly	Each review checks that the previous review's fixes still hold: the October 2026 review confirmed that every critical and high finding from July remained fixed. There is no formal measurement programme.
6.2	**Do you continuously monitor your internal attack surface?**	Partly	On every code change (tests, security checks, dependency audit) and through System Health. This is automated checking at release time, not round-the-clock monitoring.
6.3	**Do you continuously monitor your external attack surface?**	Partly	Our cyber security specialist tests the service from the outside, including penetration testing. We do not yet use a continuous external scanning service, and no outside firm has tested us.
6.4	**Do you carry out internal audits against ISO 27001?**	Not yet	No. Our reviews are security reviews of the application, not audits against the standard.

7. Improvement

NO.	QUESTION	STATUS	OUR ANSWER
7.1	**Do you check that fixes actually worked?**	Yes	Yes. Fixes carry automated tests that run on every later change, and the next review re-checks earlier findings.
7.2	**Do you change your approach when needed?**	Yes	Yes. Examples from 2026: staff sign-in moved from a shared password to personal passwords with an authenticator code; encryption was extended from selected notes to all client free text; AI providers without confirmed no-training terms were switched off.
7.3	**Can you describe your improvement process?**	Partly	Review, number each finding, fix it with a test, release, record it in plain-language release notes, then review again. It is followed consistently but is not written as a formal procedure.
7.4	**Can you show that security has improved since you started?**	Yes	July 2026: 30 review findings fixed. October: nightly off-platform backups, AI providers restricted to no-training terms, complete deletion. November: two-step sign-in for staff, encryption of all client free text, the ability to change the encryption key without downtime (done on 16 November), known vulnerabilities in shipped components reduced from 155 to 4 moderate, and self-service account deletion.

What we do not have yet

- No ISO 27001 or SOC 2 certificate, and no formal information security management system.
- No penetration test by an outside firm yet, and no continuous external scanning service. Testing is done by the cyber security specialist on our team.
- No written incident response plan, recovery-time targets, risk register or vendor-assessment schedule.
- No formal security training programme.
- No single sign-on and no two-step sign-in for client users (staff have two-step sign-in).
- Names and personal details are not removed from text before it goes to the AI provider.
- The browser content security policy is in monitoring mode before it is enforced.
- One derived index that the AI assistant uses is stored structured in the database and is not yet encrypted field by field.
- One shared database by default, with separation enforced in the application. A server of your own is by arrangement.

Who processes your data

PROVIDER	WHAT FOR	WHERE
Railway	Application hosting, database, file storage	USA (US West)
Cloudflare	Nightly backup storage; human check at sign-up	Global network
Anthropic	AI analysis (main provider), no training on client data	USA
OpenRouter	AI analysis (backup route), no data collection, zero retention	USA
OpenAI	Voice-to-text; report cover images (client names left out)	USA
Stripe	Payments. We never see or store card details	USA
SMTP2GO	Sending email	Global network
HighLevel	Sign-up text codes; customer records for our own sales	USA

The application, its database and file storage are hosted in the United States (West Coast). If you need your data kept in your own country, we can set up a server and host it for you there.

This questionnaire describes the Prism Profiles service at 17 November 2026 and is provided for vendor assessment. It is a statement of our practices, not a certification or an audit report. A data processing agreement is available on request. Questions: info@prismprofiles.com.