

WHITE PAPER

How Avantus protects student data

Avantus reads the assessment marks your school already holds. That means we handle data about real children, and you should know exactly how it is treated before you hand any of it over.

Version 2.10 · 7 October 2026 · Review: March 2027

CONTENTS

- | | |
|--|--|
| 01 The school owns the data. We process it. | 05 Access, encryption and hosting |
| 02 We take as little as possible | 06 If something goes wrong |
| 03 How one school is kept from another's students | 07 How we test, and where we are headed |
| 04 How a connection to your system works | 08 Standards, and the questions to ask us |

This document is written to be checked. Every control described below is in production unless it says otherwise. Section 07 sets out how we check our security and what strengthens it next.

01

The school owns the data. We process it.

Under Australian privacy law the distinction matters. Your school is the **data controller**: you decide what data exists, and you grant or revoke our access to it. Avantus is the **data processor**: we act on your instruction, for one purpose, which is the forward projection of student results.

We do not sell student data, and we do not share it beyond the sub-processors that host and carry it, listed in section 05. No student data goes to any AI provider.

Access is switched on by the school and can be switched off by the school. A connection to your student information system stays inactive until a staff member with authority explicitly authorises it, and revoking that authorisation is something you do in your system, not something you ask us to do.

02

We take as little as possible

The most reliable way to protect a category of data is not to hold it. Avantus reads a deliberately narrow slice of what your student information system contains.

WHAT WE READ	WHAT WE NEVER READ
Student name and identifier, year level, house, school-defined groups where supplied, and the subjects each student is enrolled in	Wellbeing, pastoral care and behaviour records
Assessment results and their weightings, Year 7 to 10 standing in each subject area (a percentile or band, not the marks) and an entry benchmark (for example NAPLAN) where the school supplies them	Medical, health and disability information
Historical results for past cohorts, where the school chooses to supply them	Attendance
The goal a student sets, their Year 10 interest-sort answers, and VET progress where relevant	Family and financial records, and anything in a module outside the one we are granted

This is enforced in code, not by policy. Our ingest adapters read specific endpoints and no others, and an automated test for each connector (Sentral, Synergetic, TASS) fails the build if it reaches outside that boundary, even where the school's system would happily expose it.

WHY THIS MATTERS TO YOUR RISK PROFILE

A vendor that pulls a full student record inherits your most sensitive categories of data, and hands them to you as a breach exposure. We cannot lose what we never collected, and your assessment of us is correspondingly narrower.

One further reduction is worth naming. The historical data that calibrates our projections carries **no names**. It records that a student at a given rank in a given subject and year finished on a given mark. A student appears there only as a keyed hash of their NESA number, made with a secret unique to your school, so one student's results link across years without the number itself being stored. That makes it pseudonymised rather than anonymous: it cannot be reversed without your school's secret, which we hold.

03

How one school is kept from another's students

This is the question a technical reviewer asks first, and it deserves a specific answer rather than a reassurance.

Separation is enforced by the database itself, using PostgreSQL **row-level security**, not by application code remembering to filter. Every table holding school data carries a policy tying each row to one school. A query that arrives without a school context returns **zero rows**, not everything.

That direction matters. A filter that is forgotten returns all data; a policy that is not satisfied returns none. The failure mode is a blank screen, not a breach.

CONTROL	STATE, VERIFIED 6 OCTOBER 2026
Tables carrying school data with row-level security <i>forced</i>	25 of 25
Isolation policies in force	26
Application role able to bypass those policies	None
Tables owned by the application role	None

Two of those rows are subtler than they look. **Forced** matters because in PostgreSQL a table's owner normally bypasses row-level security entirely; forcing it removes that exemption. And the application connects as a role that **owns nothing**, because an owner can alter a table, and altering a table can drop a policy.

HOW WE KNOW IT IS STILL TRUE

A posture audit runs automatically on every deployment and again daily, checking all four rows above plus further conditions, including that the audit log is append-only. It emails only on failure, because a daily all-clear is a message people learn to ignore. If a policy were ever dropped by a migration or a restore, the application would keep working normally and nothing would visibly break, which is exactly why the check exists.

04

How a connection to your system works

Avantus reads from the system you already run. What that looks like in practice depends on the system, but the shape is the same in every case.

For the hosted service, nothing is installed. There is no agent on your network, no server access, no VPN and no database credentials. If you prefer, Avantus can instead run entirely inside your own network or cloud subscription (our on-premises edition), with no student data leaving it. No school runs it yet.

You issue the access, and you can revoke it. A school generates a token or key inside its own system, scoped to the specific module we read. Revoking it is done in your system and does not require our cooperation.

We never hold a person's password and never sign in as a member of staff. The credential is a machine credential belonging to the school.

Staff sign in through your own identity provider. Avantus supports single sign-on with both **Microsoft Entra ID** and **Google Workspace**, so a staff member authenticates against your directory, not a password we store. Both are built and tested, and Microsoft sign-in was tested end to end against our own tenant on 15 September 2026. Neither has yet run against a live school directory, so we set it up together in a test tenant first. Our Microsoft app asks only for a person's name and email address (`openid` , `profile` , `email`); it has no application permissions and never reads your directory, and roles come from the groups you assign to it. Who a person is, and what they may see, follows your directory groups: when your IT disables an account or moves someone out of a group, their access to Avantus changes on their next sign-in, without us doing anything. We hold no separate password for any user.

We have built adapters for **TASS** (the Student Academic Analytics API), **Sentral** (Markbook and Reports) and **Synergetic**, tested against sample data in each system's format. None has yet run against a live school system, and they do not yet load class lists or teacher assignments, so at first your timetable data comes by spreadsheet export. Any school can start with an export, which needs no integration.

WHAT YOUR IT TEAM WILL WANT TO HEAR

The integration is designed to be outbound-only and read-only: Avantus calls your system, your system never calls ours, and we write nothing back to it. There is no inbound firewall change to make and no port to open.

05

Access, encryption and hosting

Hosting is in Sydney. Student data, including both backup copies, is stored in Australia and the application processes it there. App email is sent from Australia, through Amazon Web Services' email service in Sydney; a scheduled report email carries only a notice and a link, and the report itself opens inside Avantus for the signed-in recipient. One thing that carries app traffic can pass outside Australia: traffic in transit passes through Cloudflare's network, which decrypts it at its edge; Cloudflare chooses the location and does not guarantee it is in Australia. Separately, Google Analytics runs on our public marketing pages only, not in the app.

Who handles your data. **DigitalOcean** (Sydney): hosting, database and backups. **Amazon Web Services** (Sydney): app email, carrying a notice and a link, never report contents; and a second, locked copy of our encrypted backups, which it cannot read. **Cloudflare**: network and TLS, decrypting in transit. **Microsoft or Google**: your sign-in. **Google Analytics**: our public marketing pages (/ , /method , /pilot) only. **GitHub**: our source code only, no school data. Not sub-processors: our company email on Google Workspace (correspondence and outreach, not school data; please do not send student data by email), and Anthropic (an AI coding assistant that sees counts and settings only, and our outreach CRM's drafting, with no school data). A published sub-processor list is drafted, not yet live.

In transit. TLS 1.2 and 1.3 only (1.0 and 1.1 are refused), with HTTP Strict Transport Security enabled, so a browser will not downgrade to an unencrypted connection even if asked to. Our server accepts web traffic only through Cloudflare: it checks Cloudflare's client certificate and refuses any other connection without a response, so Cloudflare's protections cannot be side-stepped by going to the server's address. One limit: Cloudflare's free plan still offers some older (CBC-mode) cipher suites to TLS 1.2 browsers alongside the strong ones; restricting that list needs a paid add-on we are considering. The connection from

Cloudflare to our server accepts strong suites only, and Cloudflare checks our server's certificate (Full strict mode). Our content security policy lets scripts run only with a single-use code issued for each page; inline styles are still allowed. DNS records name the only certificate authorities allowed to issue for our domain.

At rest. The database is a managed PostgreSQL service in Sydney, encrypted at rest by the hosting provider, and **not reachable from the internet**: its firewall accepts connections only from our application server. That server's firewall permits three inbound ports: SSH, HTTP and HTTPS; SSH accepts keys only, from listed addresses, and repeated failures are blocked. The application runs as its own restricted system account, not as an administrator; its code is owned by the administrator account and read-only to it, so it cannot change its own code. The server also runs other services operated by SDG Services; school data is stored only in the separate managed database, and the server processes it but does not keep it.

Connection credentials are encrypted at rest. The key a school issues us is stored encrypted, under a master key kept outside the database, and decrypted in memory only at the moment it is used. It is never written to our database in readable form and is redacted from all logs. Tampering with a stored credential makes it fail to decrypt rather than silently produce a wrong value.

Staff access. Sign-in is through your school's own Microsoft or Google identity, so access follows your existing joiners and leavers process, and your multi-factor rules apply. When a staff member leaves your directory they cannot sign in again; a session already open lasts until it expires (12 hours idle, 7 days at most). Passwords, where used (our own admin console and demonstration accounts), are hashed with a memory-hard algorithm and compared in constant time. Our admin console also requires an authenticator-app code at every sign-in, and refuses connections from outside Australia.

Our own access. One person at Avantus, in Australia, operates the service. Day-to-day checks of production use reviewed read-only scripts that return counts and settings, not student records. The same person holds administrative access to the server and database, for deploys and erasure, by key only and with two-step sign-in on every account. Engineering is done with an AI coding assistant working under that person's direction; it uses the same scripts and never receives student records.

What staff can see is scoped to their role. A classroom teacher, a head of faculty and a principal do not see the same view, and the difference is enforced server-side rather than by hiding buttons. A teacher sees only the students in classes your timetable data records as theirs, and none until it does. A head of faculty sees their own faculty, plus the students in any class they teach. On 6 October 2026 we checked every page, PDF and action for every role and found nothing open; that check is self-run, not independent.

Access to student records is logged. Opening a student's record, viewing a cohort list, changing a goal, signing in or out, changing a role, viewing a report, importing data and changing settings are written to an audit trail, which is itself tenant-isolated so one school cannot read another's log. The application can add to it but never change or delete it. The log records who, what and when, and deliberately holds no personal information in its detail field.

06

If something goes wrong

We maintain a written breach response runbook that follows the steps of the Notifiable Data Breaches scheme. It covers containment, assessment, notification to affected schools, and the sequence in which those happen.

The runbook includes a candid section on the gaps that would affect our response, which is unusual to publish and is the point. A plan that assumes everything works is not a plan.

Backups are held off-site, encrypted, in Australia. A nightly dump is written to two object stores in Sydney, separate from the server and from the database: DigitalOcean Spaces, and Amazon S3 with Object Lock, so the second copy cannot be changed or deleted for 90 days, even with our server's key.

Two details matter more than the schedule:

- **The server cannot read its own backups.** They are encrypted to a public key; the private key is held offline and never touches the platform. Someone who compromised our server could write a backup and could not open one.
- **The restore path is exercised, not assumed.** A backup that has never been restored is a file, not a backup, and the failure people meet is not a missing backup but one that will not restore. On 1 October 2026 we restored the newest backup onto a fresh, throwaway server in Sydney: 1 minute 41 seconds from backup to verified data, with every school's record counts matching production. The drill also found a gap (restoring onto a brand-new database server needed its roles created first), which we fixed and wrote into the recovery runbook the same day. The drill used the first copy; restoring from the second copy is written down but not yet drilled.

Off-site copies are kept for 90 days. The series began on 18 September 2026, so the full 90 days is held from mid-December 2026. The managed database also keeps the hosting provider's own automated backups for about 8 days. Ours exist because those share a single point of failure with the database itself: the provider account. The second copy, with a different provider, covers the loss of that account too.

Deletion. When a school leaves, its data is removed from the live database in two stages: the application removes everything except the audit log, deepest-first in a single transaction, and an operator step then removes the audit log and the school record. The deletion then *verifies* that no rows remain, because most of our relationships do not cascade automatically and a naive delete would fail part-way and leave orphans behind. When every table counts zero we issue a certificate of deletion, with the date by which the backups holding the school's data are gone (within about 13 weeks, 92 days). We do not delete backup copies early: they are locked against tampering. The erase command is tested on throwaway databases and installed on our production server (6 October 2026). It has not yet been used on a real school's data.

07

How we test, and where we are headed

The controls above are in production unless marked otherwise. This section says how we check them today and what strengthens them next.

- **We test our own security continuously, and welcome yours.** Every change runs our automated test suite, including tests of tenant isolation against a real database. We also run our own checks: a vulnerability assessment (2 October 2026) and a check of what every role can see (6 October 2026). Both are self-run, not independent. We have committed to an **independent penetration test** within six months: the outside check a vendor cannot give itself. It is not yet commissioned or booked. We will also make the platform available for a school's own security review at any time.
- **We alert on the access patterns that matter.** A daily scan reads the audit trail and flags the shapes an extraction or a compromised account makes: one account reading an unusually large number of records, a whole-cohort sweep in a short burst, reading well outside school hours, or a run of failed sign-ins from one address. A finding emails us. This is a daily check, not real-time detection; per-user behavioural baselining is on our roadmap.

IN SHORT

The isolation, data minimisation, encryption and access controls described in this document are in place, and we can show you each one. What is not yet done is stated where it applies: no independent test yet; no certifications; no cyber insurance; our security policies and privacy documents drafted but not yet in force; single sign-on not yet run against a live school directory; school system connectors not yet live; the erase command not yet run on real data; the second backup copy not yet restore-drilled; and the two-step check at sign-in is built but not yet switched on. We would rather show you both the controls and the gaps than ask you to take either on faith.

08

Standards, and the questions to ask us

Not yet ST4S assessed. We have reviewed ourselves using the criteria ESA publishes; that review, gaps included, is available on request. (Safer Technologies 4 Schools is the framework Australian schools and education departments use to assess education technology.)

If you are evaluating us, these are the questions worth asking, and we have tried to answer each above rather than wait to be asked:

- What exactly do you read from our system, and what do you decline to read?
- How is our data separated from another school's, and what happens if that separation fails?
- Where is the data stored, and does it leave Australia?
- Who can see a student's record, and is that access logged?
- What happens to our data if we leave?
- How do you assure this over time, and what strengthens it next?

The last one is the most revealing, and the one we expect to keep answering.

Avantus · How Avantus protects student data · Version 2.10, 7 October 2026

Avantus is a product of SDG Services Pty Ltd · ABN 30 664 897 429 · ACN 664 897 429

Companion documents: *What actually predicts an HSC result* · *The maths behind the projection* · *How the engine is tested*