

Panorama's Safe and Responsible AI for K-12 Education

AI Safety & Security Brief

Table of Contents

Our Approach to AI in Learning	3
Model Security & Training	3
Model Selection	
Data Is Not Used for Model Training	
Guardrails & Disclosure	
Data Privacy & Security	4
Data Isolation & Ownership	
Encryption & Security	
Student Data in AI Conversations	
Data Deletion & Opt-Out	
Access Control & Governance	5
District Control Over AI	
Role-Based Access	
Educator Permissions & Customization	
Compliance & Security	5
Certifications & Third-Party Verification	
Third-Party Vendor Responsibility	
Risk Mitigation	
Observability & Auditing	7
Audit Logging	
Admin Oversight	
Our Bet on AI in K-12	7
Frequently Asked Questions	8

Our Approach to AI in Learning

1. We lead with student safety, security, privacy, and well-being in our design and decision-making.
2. We design AI to amplify educator effectiveness, not replace educator judgment. AI assists with administrative tasks, feedback, and data analysis—freeing educators for relationships and high-order instruction. Educators control when and how students interact with AI.
3. We aim to design AI to be age-appropriate and accessible for students across language, ability, and learning needs.
4. We evaluate our AI tools and models for quality and bias, and take steps to monitor their use and impact in practice.
5. We clearly disclose when features are AI-enabled and assure that schools and districts have control over which AI features are available for use.

These commitments shape every decision below, from model selection to data handling to governance.

Model Security & Training

Model Selection

After safety testing, Panorama selected Claude, a family of LLM models from Anthropic ([hosted within our AWS ecosystem in AWS Bedrock](#) and not passed to Anthropic). We made this choice because safety is the core consideration in how these models are built. For general and education-specific risks that matter in K-12 (bias, student privacy, age-appropriateness, and misuse), Panorama has run red-team exercises and evaluations on Solara and Class Companion features, focused on output quality and safety.

Data Is Not Used for Model Training

Customer data is not used for model training. It remains under customer ownership, and can be fully deleted upon contract termination.

To maximize data privacy, real-time AI interactions are processed using a transient model architecture within our secure AWS Bedrock environment (opaque for the model provider): the AI engine only holds data for the brief moment it takes to process a request. Once processed, the conversation history is securely stored within Panorama's primary database environment so educators can access their work.

Non-negotiable commitment: Customer data is not used to train, fine-tune, or improve any language model.

Guardrails & Disclosure

We clearly and conspicuously display a notice to users that they are interacting with AI and should always review content for accuracy and bias before relying on outputs. AI models respond based on user input and context but can make mistakes.

Solara and Class Companion leverage policy-driven controls that help keep AI interactions safe and aligned with Panorama standards.

- + These are tailored to block adversarial/malicious, toxic, or non-educational content as well as certain cases of over-reliance, to focus on providing support while avoiding diagnoses and placement decisions.
- + Panorama guides AI responses to avoid making or implying sensitive determinations such as diagnoses, eligibility, disability, or treatment approaches, and instead focuses on permissible support such as summarizing data trends, past performance, and general educational strategies.
- + We also employ extensive prompt guardrails and input validation to mitigate manipulation and protect integrity.

Data Privacy & Security

Data Isolation & Ownership

Customer data is fully tenant-isolated and never shared with or passed to any other party without the express direction of the customer or as required by operation of law. Customers retain 100% ownership of all data. Panorama is strictly acting as a services provider to its education agency customers, acting as a School Official with a legitimate educational interest.

Our research team only provides aggregate national benchmarks (e.g. absenteeism) that have been aggregated from de-identified data.

More information is available in Panorama's [privacy policy](#).

Encryption & Security

All data is encrypted at rest and in transit. Access is strictly role-based so educators can only view data for students they are already permitted to see. We continuously scan for vulnerabilities with SAST, SCA and DAST scanning, and all employees receive privacy and security training.

Student Data in AI Conversations

Solara conversations are only accessible by the user, unless they specifically choose to share a conversation page with someone within the customer-approved ecosystem. Student interactions with AI in student-facing chats have guardrails and can be reviewed for safety by an authorized educator.

To mitigate bias, gender and race are strictly not shared in AI conversations within Solara and Class Companion.

Data Deletion & Opt-Out

We partner with customers to support Student opt-out from Surveys or fulfill deletion requests.

Access Control & Governance

Customer Control Over AI

Customers should be in control of AI adoption, so Panorama makes sure you choose whether to enable or disable AI features and capabilities across each product (opt-in for Success, Surveys, and Pathways). This includes Solara permissions to add tools, create interventions, or share chats. In Class Companion, assignment settings give you control over the tutoring assistant and detection mechanisms for potential copy-paste.

Role-Based Access

Customers control role-based access at the district level, allowing you to restrict educators to only have access to the appropriate students and capabilities. Only customer-designated administrators can initiate or access analysis of sensitive free responses to surveys.

Educator Permissions & Customization

Teachers can edit AI-generated assignment scores (Class Companion) and the interventions they create with Solara. In some products, including Solara, educators can customize instructions for AI responses, including tone and terminology.

Compliance & Security

Certifications & Third-Party Verification

Panorama maintains active SOC 2 Type 2 certification and can provide documentation on request; we support your procurement process and respond to district privacy questionnaires and RFPs.

Certification	What It Verifies	Frequency
SOC 2 Type 2	Security controls, confidentiality practices, and access logging	Continuous compliance + annual security penetration tests
1EdTech Data Privacy Certification	Privacy practices in compliance with education standards	Annual renewal
Common Sense Media Privacy Rating	93% score placing Panorama among top-rated education technology platforms	Annual evaluation



Third-Party Vendor Responsibility

Panorama relies on AWS and Anthropic. Here's how we manage vendor security:

Component	Vendor	Responsibility
AI Model Development	Anthropic (Claude)	Model safety, bias testing, adversarial robustness; subject to Panorama red-team exercises
AI Model Hosting	AWS Bedrock	Secure infrastructure, encryption, availability, SOC 2 Type 2
Infrastructure & Storage	AWS	Encryption, SOC 2 Type 2
Applications	Panorama	Access control, tenant isolation, and SOC 2 Type 2

Important note: Anthropic cannot see your data. Requests are processed through AWS Bedrock in a transient environment and are not logged by or shared with Anthropic.

Risk Mitigation

Risk	Mitigation
AI hallucination or incorrect output	Our platform makes grounding data about students available to the model and easy to reference, and lets educators attach materials and rubrics to Solara Tools and Class Companion assignments. AI can make mistakes, so we encourage review and let educators edit AI feedback (such as interventions created with AI).
Bias in AI responses	Gender and race are strictly not shared in AI conversations within Solara and Class Companion. We also run dedicated testing against bias, and include explicit anti-discrimination instructions in the system prompt.

Risk	Mitigation
Age-inappropriate or harmful content	Solara and Class Companion leverage policy-driven controls that help keep AI interactions safe and aligned with Panorama standards, especially for student-facing outputs. Guardrails bolster the underlying model to mitigate toxic or non-educational content. We also guide conversations against certain cases of over-reliance, to focus on providing support while avoiding diagnoses and placement decisions, and we flag unsafe student conversations for review by an authorized educator.
Adversarial misuse (e.g., jailbreak attempts)	We leverage Claude for its higher robustness against malicious prompts, and also apply specialized AI guardrails to help block malicious content in prompts (jailbreak, prompt injection). These strategies reduce risks that can occur from manipulated conversations.

Observability & Auditing

Audit Logging

Panorama retains audit logs of all AI interactions, including the user, date and time, actions taken, prompts sent to the model, model responses, and other relevant attributes associated with the logged activity.

Admin Oversight

Customer-selected admins can track high level usage trends (by category) and tool usage frequency.

Our Bet on AI in K-12

We believe the right model for AI in K-12 is **informed decision-making, not automated decision-making**. AI should give educators more visibility into student learning and suggest high-quality actions, but teachers and districts keep the authority.

That's why we've ensured:

- ✓ **Data stays yours:** Tenant isolation, transient models, can be deleted on exit
- ✓ **Control stays yours:** Feature flags, role-based access, audit logs
- ✓ **Transparency is built-in:** Clear AI disclosures, educator review encouraged, customizable behavior
- ✓ **Safety is architectural:** Walled garden in AWS with transient processing, guardrails, grounding in your data, and safer AI model

This approach trades off some operational convenience and feature velocity for customer trust. We're betting that customers will choose Panorama because we respect their data, their authority, and their responsibility to students.

Frequently Asked Questions	Answers
How does Panorama validate the security of its systems?	Our SOC 2 Type 2 certification independently verifies our practices in security and confidentiality through continuous compliance and annual third-party penetration testing.
Who can access student data in Panorama?	Access is strictly role-based, so educators can only view data for students they are already authorized to access.
Where is data stored?	Our platform is hosted in AWS and stores data within the US.
Which AI models does Panorama use, and how are they evaluated for K–12?	We use Anthropic's Claude family of models through AWS Bedrock because safety is a core design principle. We evaluate Solara and Class Companion against K–12-specific risks including bias, student privacy, age appropriateness, and misuse. Student and customer data is not shared with Anthropic.
How does Panorama keep AI interactions safe and appropriate for schools?	Solara and Class Companion use policy-driven guardrails designed for K–12 environments. These controls help prevent malicious, toxic, or non-educational content, discourage over-reliance on AI, and avoid diagnostic or placement decisions while providing appropriate educational support.
How does Panorama protect our district's data and student privacy?	Panorama does not use customer, school, or student data to train AI models. Student data is never shared across customers, and we do not conduct hidden profiling or use student data for advertising purposes. Customers retain ownership of their data at all times.
Does data processing involve any third-parties?	Our third-parties list provides visibility to providers on whom we rely to provide the service to our customers.
What student data does Panorama AI have access to?	Solara securely uses your district's integrated student data and customer-approved materials to generate relevant, contextual responses. Educators can only access information for students they are already authorized to view.
Can customers monitor and audit AI usage?	Yes. Administrators can view high-level AI usage trends and tool adoption. Panorama also maintains detailed audit logs of AI interactions, including users, timestamps, prompts, responses, and actions taken.
How does Panorama handle student opt-outs and data deletion requests?	Panorama partners with customers to support student opt-outs from surveys (where applicable) and fulfill deletion requests. Customers retain ownership of their data, which is not used for AI model training and can be fully deleted at the end of the contract.