Level Data

PLAYBOOK • AI FOR EDUCATORS

The K-12 AI Readiness and Governance Playbook

A practical guide to adopting artificial intelligence with clear purpose, trusted data, human oversight, and measurable value.

CREATED FOR

Superintendents, technology and data leaders, curriculum teams, student services leaders, and state education agencies

Move from experimentation to responsible adoption

Artificial intelligence is already entering K-12 work through stand-alone tools, embedded product features, and staff experimentation. The leadership challenge is no longer whether AI exists in the district. It is whether the district can guide its use with enough clarity to protect students, support educators, and produce meaningful value.

This playbook turns responsible AI principles into a practical operating model. It helps leaders choose worthwhile use cases, establish decision rights, protect data, evaluate vendors, design pilots, prepare staff, and measure results. It is designed as a starting point, not legal advice, and should be adapted with district counsel and applicable federal, state, and local requirements.

WHAT YOU'LL LEARN

- Create a districtwide vision and a small set of usable principles.
- Inventory approved, embedded, and unapproved AI use.
- Score use cases by value, readiness, and risk.
- Apply privacy, security, equity, and human-review safeguards.
- Ask vendors the questions that matter before approval.
- Launch a bounded pilot and make an evidence-based scale decision.

How to use this playbook

- Convene a cross-functional AI leadership team.
- Complete the readiness scan and AI inventory.
- Choose one low- or moderate-risk use case tied to a real need.
- Use the evaluation, vendor, and pilot tools before launch.
- Review evidence at 30, 60, and 90 days; then improve, scale, pause, or stop.

1 Set the direction before selecting tools

A useful AI strategy begins with educational and operational priorities—not a list of products. A short vision statement should explain the human problem AI may help address, the value the district expects, and the boundaries that will remain in place.

> VISION STATEMENT TEMPLATE
> We will consider AI when it can help [people] improve [priority outcome] while preserving [human responsibility], protecting [sensitive information], and measuring [evidence of value].

Six practical principles

Principle	What it means in practice
Purpose before novelty	Start with a defined need, user, and desired result.
Human responsibility	A qualified person remains accountable for consequential decisions.
Privacy and security	Use only authorized data, tools, access, and retention practices.
Transparency	Make material AI use understandable to affected people.
Equity and accessibility	Test for uneven quality, access, burden, and impact.
Evidence and improvement	Pilot, measure, document, review, and adjust continuously.

Leadership alignment prompts

1. Which district priority could AI plausibly support?

2. Which decisions must always remain with educators or administrators?

3. What uses are prohibited until further review?

4. What evidence would justify continued investment?

2 Establish governance and decision rights

AI governance should make it easy to know who can propose a use, who reviews it, who approves it, and who monitors it. Keep the structure proportionate: routine productivity support should not require the same review as a tool that influences student placement, discipline, eligibility, grading, or services.

Role	Core responsibility
Executive sponsor	Sets priorities, resolves tradeoffs, and communicates accountability.
AI governance lead	Owns intake, risk classification, approvals, inventory, and review cadence.
Technology and security	Evaluates architecture, access, integrations, logging, and incident response.
Data and privacy	Reviews data flows, FERPA considerations, contracts, retention, and disclosure.
Instruction and student services	Evaluates educational fit, professional judgment, equity, and accessibility.
Legal and procurement	Reviews applicable law, terms, insurance, remedies, and vendor obligations.
End users and community	Surface workflow realities, unintended effects, and communication needs.

Risk-based review tiers

- Tier 1 — lower risk: brainstorming, formatting, or drafting with no confidential data and mandatory human review.
- Tier 2 — moderate risk: staff-facing analysis, summarization, recommendation, or workflow support using approved data.
- Tier 3 — high risk: student-facing systems or uses that may affect rights, access, grades, placement, discipline, services, safety, or resource allocation.
- Prohibited until approved: entering protected data into consumer tools, autonomous consequential decisions, covert monitoring, or uses without an accountable owner.

> GOVERNANCE RULE
> Higher potential impact, more sensitive data, less explainability, or less human control should trigger a more rigorous review.

3 Inventory the AI already in use

AI may be present even when no one purchased an "AI product." Features can appear in learning platforms, productivity suites, analytics tools, communications systems, security products, and vendor updates. An inventory creates visibility into approved use, embedded features, data exposure, and ownership.

Inventory field	What to capture
Tool or feature	Product name, feature, version, and vendor.
Purpose and users	Problem addressed, intended users, and affected groups.
AI function	Generate, classify, predict, recommend, summarize, detect, or act.
Data	Inputs, outputs, sensitive fields, location, retention, and training use.
Decision impact	Informational, supportive, or consequential.
Human review	Who verifies output and who remains accountable.
Status	Proposed, pilot, approved, restricted, paused, or retired.
Evidence	Quality, adoption, time, cost, equity, and outcome measures.

AI discovery questions

1. Which current products recently added AI features?

2. Which consumer AI tools are staff or students using independently?

3. Where might AI outputs already influence important decisions?

4. Which use has no clearly accountable owner?

4 Select use cases with a value-risk screen

A promising use case has a clear problem, a realistic workflow, data that can be used lawfully and securely, and an outcome the district can measure. Score the use before discussing a specific vendor.

Criterion	1 — weak	3 — moderate	5 — strong
Need	Optional convenience	Recurring friction	Priority problem with clear owner
Expected value	Unclear benefit	Some time or quality gain	Meaningful outcome or capacity gain
Data readiness	Unclear or unreliable	Usable with preparation	Trusted, governed, and available
Human oversight	No clear reviewer	Review exists but varies	Named accountable reviewer and escalation
Measurability	No baseline	Proxy measures available	Baseline, target, and review cadence
Risk	High and uncontrolled	Manageable with controls	Low and well controlled

> DECISION RULE
> Do not offset a severe privacy, safety, civil-rights, or student-impact concern with a high value score. Resolve the risk, redesign the use, or stop.

Good early pilots often have these traits

- Staff-facing rather than autonomous student-facing use.
- No protected data, or a clearly approved and minimized dataset.
- Easy comparison with the current process.
- A qualified reviewer can verify every output.
- Errors are reversible and unlikely to create lasting harm.

5 Protect data, privacy, and security

AI use does not replace existing student privacy, records, security, procurement, and accessibility obligations. Treat prompts, uploaded files, generated outputs, logs, embeddings, and feedback as parts of the data lifecycle. The district should understand where each item goes, how long it remains, and who can access it.

Minimum safeguards

- Use approved accounts and tools; prohibit protected data in unapproved consumer services.
- Minimize data to what the use case actually requires.
- Document the legal basis and permitted purpose for disclosure and processing.
- Contractually address vendor access, subprocessors, location, retention, deletion, breach notice, model training, and secondary use.
- Apply role-based access, multifactor authentication, logging, and periodic access review.
- Define how AI-generated records are stored, corrected, disclosed, and deleted.
- Create a response path for inaccurate output, harmful content, unauthorized disclosure, or unexpected system behavior.

Data question	District answer
What information enters the system?	
Is any information personally identifiable or otherwise sensitive?	
Where is it stored and processed?	
Can the vendor use it to train or improve models?	
Who can access prompts, outputs, logs, and support data?	
When and how is each data type deleted?	

IMPORTANT

This playbook is operational guidance, not legal advice. Review FERPA, PPRA, COPPA where applicable, state student privacy laws, records rules, labor agreements, accessibility requirements, and local policy with qualified counsel.

6 Preserve human judgment, quality, and equity

AI output can sound confident while being incomplete, inaccurate, biased, or unsupported. Human review must be an actual control—not a label added to a workflow. Reviewers need time, subject expertise, access to source information, and authority to reject the output.

Control	What good looks like
Verification	Important claims are checked against authoritative sources or records.
Explainability	Users can understand the basis, limitations, and appropriate use of the output.
Override	A qualified person can reject, correct, or reverse the output.
Bias testing	Quality and error rates are examined across relevant groups and contexts.
Accessibility	The workflow and outputs work for people with disabilities and diverse language needs.
Appeal and correction	Affected people can raise concerns and obtain timely human review.
Monitoring	Performance is checked after launch, not assumed from a prelaunch test.

Reviewer checklist

- Can I identify the original sources or evidence?
- Is any material statement unsupported, fabricated, outdated, or overly certain?
- Would the output change if the user, school, language, or student group changed?
- Does the output omit context a qualified professional would consider?
- Can the affected person understand, question, and correct the result?
- Am I accountable and authorized to approve this use?

7 Evaluate vendors before approval

A general security questionnaire is necessary but not sufficient. AI features create additional questions about models, training data, updates, output controls, evaluation, monitoring, and responsibility when performance changes.

Area	Questions to ask
Purpose and limits	What is the system designed to do, and what uses does the vendor explicitly discourage?
Models and changes	Which models are used? How are material model or feature changes disclosed and evaluated?
District data	Is district data used for training, tuning, evaluation, advertising, or product improvement? Can the district opt out?
Quality evidence	What tests support accuracy, reliability, bias, safety, and performance in K-12 contexts?
Human control	What review, override, citation, confidence, audit, and administrator controls exist?
Security and privacy	How are access, encryption, subprocessors, retention, deletion, incidents, and data location handled?
Accessibility and equity	What standards, testing, languages, accommodations, and disparity analyses are supported?
Exit and remedies	Can data and logs be exported and deleted? What happens after termination or material failure?

PROCUREMENT CHECKPOINT

Require the vendor's claims, safeguards, data practices, and change-notification commitments to appear in enforceable contract language—not only in sales materials.

8 Prepare people and communicate clearly

Responsible adoption depends on staff capacity and public trust. Training should be role-specific, scenario-based, and connected to approved workflows. A one-time AI awareness session is not enough for people who will review outputs or make consequential decisions.

Audience	Minimum preparation
All staff	Approved tools, prohibited data, verification habits, disclosure expectations, and incident reporting.
Educators	Instructional purpose, academic integrity, age appropriateness, accessibility, source checking, and student agency.
Leaders and reviewers	Risk classification, evidence review, bias and impact monitoring, override, documentation, and escalation.
Technology and data teams	Architecture, data flows, access, logs, integrations, model changes, monitoring, and security response.
Students and families	Plain-language purpose, when AI is used, what data is involved, expectations, limitations, and how to raise concerns.

Communication template

PLAIN-LANGUAGE MESSAGE

We are piloting [tool/use] to help [users] with [need]. It will use [data description] under [key safeguards]. A qualified [role] will review outputs and remain responsible for decisions. We will evaluate [measures] through [date]. Questions or concerns can be directed to [contact].

9 Run a bounded pilot and measure value

A pilot should answer a decision question, not merely demonstrate that the tool works. Capture the current baseline, define the intended users and boundaries, test the controls, and decide in advance what evidence will support scaling—or stopping.

Measure	Examples
Quality	Accuracy, completeness, relevance, consistency, correction rate.
Time and burden	Minutes saved, steps removed, review time, training and support demand.
Adoption	Eligible users, active users, frequency, completion, abandonment.
Equity and accessibility	Performance, access, usability, and burden across relevant groups.
Risk and trust	Incidents, near misses, complaints, overrides, confidence, and transparency.
Outcome and cost	Change in the priority outcome, full cost, and cost relative to alternatives.

Pilot brief

1. What decision will this pilot inform?

2. Who is included—and explicitly excluded?

3. What baseline and target will we use?

4. Which safeguards must be tested before real use?

5. What evidence will trigger scale, redesign, pause, or stop?

10 A 90-day implementation roadmap

Period	Objective	Key actions	Exit evidence
Days 1-30	Align and discover	Name the team; approve principles; inventory uses; classify risks; select one pilot.	Owners, inventory, risk tier, pilot brief, and baseline are documented.
Days 31-60	Validate and prepare	Complete privacy, security, legal, vendor, accessibility, workflow, and training reviews.	Controls pass testing; users are trained; communication and incident paths are ready.
Days 61-90	Pilot and decide	Run bounded use; monitor quality, burden, equity, incidents, outcomes, and cost; gather feedback.	Leadership makes a documented scale, improve, pause, or stop decision.

Readiness checklist

- [ ] A district AI vision and principles are approved.
- [ ] An accountable governance owner and cross-functional review team are named.
- [ ] AI tools and embedded features are inventoried.
- [ ] Uses are classified by data sensitivity and decision impact.
- [ ] Privacy, security, legal, accessibility, and procurement reviews are defined.
- [ ] Protected data is limited to authorized tools and approved purposes.
- [ ] Human review, override, and escalation paths are documented.
- [ ] Staff receive role-specific preparation before use.
- [ ] Families and students receive understandable information when materially affected.
- [ ] Performance and equity are monitored after launch.
- [ ] The district can export and delete its information when a tool is retired.

How Level Data supports AI readiness

AI can only be as useful as the information and operating practices around it. When definitions conflict, records are incomplete, systems are disconnected, or leaders cannot see how information changes over time, AI may accelerate confusion instead of improving decisions.

Level Data helps K-12 districts and states strengthen that foundation by validating data quality, connecting information across systems, automating critical workflows, and giving leaders clearer visibility into reporting, services, implementation, coaching, and education investments. That trusted data foundation can help organizations evaluate AI use cases more responsibly, establish reliable baselines, monitor results, and keep human decision-makers focused on what matters most.

> START WITH TRUSTED DATA
> Before adding AI to a critical workflow, confirm the authoritative source, data owner, validation rules, permitted use, refresh timing, and the person accountable for action.

Selected guidance and resources

- National Institute of Standards and Technology. AI Risk Management Framework 1.0 and Generative AI Profile (NIST AI 600-1), 2023-2024. nist.gov/itl/ai-risk-management-framework
- U.S. Department of Education, Student Privacy Policy Office. Protecting Student Privacy resources, including FERPA and data security guidance. studentprivacy.ed.gov
- TeachAI. AI Guidance for Schools Toolkit, updated 2025. teachai.org/toolkit
- Consortium for School Networking and Council of the Great City Schools. K-12 Generative AI Readiness Checklist, version 1.3, 2025. cosn.org
- UNESCO. Guidance for Generative AI in Education and Research, 2023. unesco.org
- Guidance and requirements continue to evolve. Review current federal, state, and local requirements before adopting or materially changing an AI-enabled workflow.

Continue the conversation

Explore Level Data solutions or request a conversation at leveldata.com.