

Gemini and NotebookLM for students

Two Google tools, two different decisions. This sheet gathers what independent reviewers, state guidance, and Google's own documentation say, so your team can set local policy. **It does not recommend a setting.**

The argument for enabling access is not enthusiasm for the tools. It is that the alternative is not *no AI*. Students with a phone and a personal email address already have Gemini, and the district sees none of it — access concentrated among students whose families can supply a device and a data plan. The counter-argument is equally concrete: enabling a tool means the district owns its limitations, and independent testing has documented real ones. Both can be true, which is why this decision is usually about conditions rather than a yes or a no.

MICHIGAN CONTEXT

In June 2026 the Michigan Department of Education released its Artificial Intelligence Guidance Resource with Michigan Virtual, pairing the MDE Comprehensive AI Guide for Districts with Michigan Virtual's AI Planning Guide and Integration Framework. It phases AI planning into the Michigan Integrated Continuous Improvement Process, so this decision can sit inside a cycle your district already runs.

SIDE BY SIDE

The two tools are not one decision

	Gemini	NotebookLM (renamed Gemini Notebook, Jul 2026)
What it is	An open-ended chatbot in Workspace for Education, plus teacher-built Gems in Classroom.	A source-grounded research and study tool. Same product as before the rename; notebooks and links carry over.
Where answers come from	The model plus the open web. The student sets the topic, so the scope is whatever they type.	Only the sources loaded into the notebook, cited back to the passage — unless web discovery adds sources.
Student availability	Administrator-controlled, and scopable by organizational unit.	A core Workspace service for Education users of all ages, on by default. Consumers 13+, with stricter policies under 18.
Independent review	Common Sense Media, January 2026. Overall risk: high.	No standalone assessment published. It runs on Gemini models, so model-level findings still apply.
Limitation to plan for	Students reach the general chatbot even when a teacher assigns a narrow Gem; only Workspace administrators can change that.	Grounding holds only if someone curates the sources. Summaries can still misstate, and a polished audio overview is easy to accept unchecked.

INDEPENDENT REVIEW

What the assessment says

Common Sense Media's Youth AI Safety Institute rated Gemini K–12 an overall risk of **high** on January 9, 2026, citing no age differentiation across K–12, inconsistent handling of sexual and substance-use content, and mental health responses administrators cannot switch off — while crediting its clear AI identity, its refusal of romantic and sexual roleplay, and its school-grade privacy protections. The Institute publishes its review method and states editorial independence, while disclosing funding from both philanthropy and industry, including makers of technologies it evaluates. Read it as one input alongside your own testing.

GRADE BAND

Where each tool tends to land

Patterns from independent review and current district practice — not district-level recommendations. Your local decision still turns on the conditions listed below.

BAND	Gemini (open chat)	NotebookLM (source-grounded)
Elementary K–5	Generally less appropriate for independent student use. Most students are under 13, which brings parental-consent obligations, and testing found output averaging a 10th- to 11th-grade reading level, with only about one in four responses below eighth grade. Teacher-mediated demonstration is where districts typically start, if at all.	Sometimes appropriate as a teacher-built resource: an adult curates sources and projects the notebook to the class. Independent student use is uncommon at this age. Reading level of summaries remains a constraint.
Middle 6–8	Approach with caution. This is where the absence of age differentiation is felt most: a sixth grader gets the same response as a senior. Reviewers found inconsistent refusals on sexual content, substance-use answers lacking context, and mental health responses administrators cannot disable. If enabled, districts usually narrow it to assignment-bound use with an adult present.	More appropriate than open chat for this band. The scope stays with the sources a teacher loads, and the tool cites back to the passage — a fit for research-and-study tasks. Districts often begin AI access here.
High 9–12	Most appropriate band for supervised access. Closest to the population these tools were designed for, and the band where personal-account use is already most widespread — so the oversight gap is widest. Requires published student guidance, course-level integrity expectations, and a named path when the tool fails.	Well suited to research, source analysis, and study — the tasks it was built for. Pair with instruction on verifying summaries against the underlying sources, since a polished audio overview is easy to accept unchecked.

PROS OF ENABLING

Reasons districts turn access on

- Moves existing use onto accounts under the Education terms, with no model training on student data.
- Closes an equity gap: access stops depending on a family device and data plan.
- Gives the district visibility, logs, and controls it has over no personal account.
- Lets teachers scaffold the tool instead of competing with it.
- Creates one supported environment for teaching AI literacy.

CONS OF ENABLING

Reasons districts hold off

- The district takes on documented limitations it cannot configure away.
- Mental health responses stay available whether or not you want them.
- One behavior spans K–12; the tool does not adjust to the student's age.
- Real, recurring cost in training, policy, and incident response.
- Turning access on is easier than turning it back off once expectations form.

CONDITIONS

Put in place first

☐ Board-approved AI policy, or an explicit interim statement

☐ Staff training finished before student access opens

☐ Student guidance in student-facing language, per grade band

☐ Academic-integrity expectations restated at course level

☐ Mental health referral path that runs to a person, not a chatbot

☐ Data-protection review completed and on file

☐ Parental notice, and consent for students under 13

☐ Admin settings documented per organizational unit, with an owner

☐ Defined response when output is harmful: who is told, what is logged

☐ A review date on the calendar

DISCUSSION

Questions for your team

1. If we do not enable this, what do we believe students are using instead — and can we support that with anything but a hunch?

2. Who owns the decision, and who owns it when something goes wrong?

3. Which grade bands, on what evidence, rather than which feels safest?

4. What is our answer the first time a student uses a chatbot for a mental health conversation?

5. Is training funded and scheduled, or aspirational?

6. What would cause us to reverse this decision, and who is watching for it?

WHERE TO READ FURTHER

Michigan Department of Education — AI Guidance
Resource. michigan.gov/mde

Common Sense Media, Youth AI Safety Institute — [risk assessments](#).

TeachAI — schools toolkit and model policy language. teachai.org/toolkit

Michigan Virtual — AI Planning Guide and AI Lab. michiganvirtual.org/ai

Google Workspace for Education — admin controls and Education Terms. support.google.com/a

Digital Promise · ISTE — AI literacy frameworks. iste.org/ai