# Is an AI story generator safe for kids? How Peabody screens every story

Peabody blog · Safety and privacy · updated 2026-10-04

An idea screen, hard limits, code checks, an AI screen with a second opinion, and stories held for a person to review. What Peabody does before a child reads.

Peabody (peabody.kids) puts every child's idea and every AI-written story through several safety layers before a child sees anything. An idea screen comes first. The writing instructions carry hard limits. Code checks the finished story for contact details, links and bad words. An AI screen passes, flags or blocks it and labels how scary it is, and a stronger model gives a second opinion on anything unclear. Whatever is still in question is held for a person to review.

No filter is perfect, and Peabody doesn't claim to be. But "safe enough for a kid to use on their own" takes more than one filter, so here is each layer, in order.

## Why one filter isn't enough

A single content filter has two ways to fail. It can be too nervous, turning away a perfectly normal monster story and teaching a child that their ideas are "bad". Or it can miss something, because children are endlessly inventive and models make mistakes.

Layers help with both. An early, gentle screen catches clearly off-limits ideas before any writing happens. Later screens look at what was actually written. And uncertainty goes up to a stronger check, or to a person, instead of being waved through.

## Layer 1: the idea screen

Before a word of the story is written, Peabody checks the child's ideas:

- Code first, every time: contact details (phone numbers, emails, addresses), a profanity list, and words that are off-limits in any story.
- Then a fast AI classifier sorts the ideas: fine, or off-limits for a reason like violence, self-harm, drugs or real people.
- If the classifier would turn an idea away, a stronger model takes a second look, so an ordinary spooky story isn't refused because of one nervous call.

When an idea really is off-limits, no story is written. Instead, Peabody offers a kind redirect, something like "Ooh, let's make it a spooky mystery instead!" Never a scolding.

## Layer 2: hard limits in the writing instructions

Every story request carries the same hard limits, word for word:

- nothing sexual or romantic of any kind;
- no graphic gore, mutilation or torture (death and violence stay off the page);
- no self-harm, suicide or eating-disorder content;
- no hate, slurs, or mocking groups of people;
- no drugs, alcohol, vaping or smoking;
- no instructions for weapons, and no dangerous stunts or "challenges" a child might copy;
- no bullying or mocking real, identifiable people;
- no secrets from parents or trusted adults, and no strangers asking kids for things;
- no profanity;
- no real brands, celebrities or copyrighted characters;
- no contact details, links or personal information;
- nothing extremist, criminal-glorifying or illegal.

The writer is also told that when a hero gets help from a trusted grown-up, that counts as the hero solving the problem. Hiding things from grown-ups is never the clever move.

## Layer 3: code checks on the finished story

Once a story is written, code scans the text for contact details, web links and the profanity list. This doesn't depend on any AI model having a good day.

## Layer 4: an AI safety screen, and a second opinion

An AI screen reads the whole story and returns one of three answers: pass, flag or block. It also gives the story an intensity: gentle, spooky, scary or horror, plus labels like funny, dragons or Halloween so readers know what they're getting into. A story's intensity is the higher of what was asked for and what the screen found.

Anything that isn't a clear pass gets a second opinion from a stronger model. A story that's blocked is rewritten once with a steer. If it's blocked again, Peabody kindly asks for a different idea.

Peabody's opening pages, which arrive in seconds, go through exactly the same screens as a whole story. Speed never skips a check.

## Layer 5: held for a person

If a story is still in question after all that, it's held. Nobody sees it, the child included ("Peabody is still polishing this one!"), until Aaron, who runs Peabody, reviews it and approves, blocks or regenerates it.

## Scariness: a ceiling, not a surprise

Plenty of kids love a shivery story, so Peabody doesn't ban "spooky". It labels it, and it sets a ceiling:

- A child using Peabody without a grown-up's account gets gentle stories until 1st grade and never anything past spooky. That rule is enforced by the database, not just the page.
- A family account lets a grown-up set each child's ceiling, from gentle up. A child never sees a story above their ceiling, in the library or anywhere else.

The mood picker on the make page follows the same ceiling: "Scary" only appears for a child who's allowed scary stories.

## What else keeps kids safe

- No open chat. The idea box is one-shot. There is no chat with an AI. See An AI reading app for kids, with no open chatbot.
- No free text between kids, with one screened exception. Kids can tap phrases from Peabody's own list to each other, and a grown-up can turn that off. In the Dungeon Dweller game's online co-op a party can type short messages; every line is screened first for names, numbers, places, links and rude words, and a grown-up chooses who their child can type to there. No comments.
- Nothing recorded. Peabody never records or receives audio. A button turns speech into typed words, using the browser's recognizer (which may send the sound to the browser maker) or, where that's missing, a speech model on the device.
- Kids' own stories stay out of search engines. Stories a child makes are never in the public library. Shared stories carry "noindex" so search engines leave them out.
- A child can report a story that upset them. It's hidden for them right away and goes to the review queue.
- Private by default for kids on their own: no name, no email, ideas deleted after screening. More in no name, no email, no account.

## What we won't claim

We won't call Peabody perfectly safe. Any system that uses AI can make a mistake, which is exactly why there are several layers, why uncertainty goes to a person, and why a child can report a story. If you ever see something that shouldn't be there, please report it from the story.

## A checklist for any AI app for kids

Whatever you choose, these are good questions to ask:

1. Can my child chat with the AI, or with strangers?

2. Does it use a microphone or record my child?
3. What does it screen: the ideas, the output, or both?
4. Who reviews anything the filter isn't sure about?
5. Can I set how scary or intense it gets?
6. Does it ask for my child's real name, age or email?
7. Who else can see what my child makes?
8. Are there ads?

You can see Peabody's answers to all of these by trying it, no account needed. Peabody is free through the end of 2026.

## Questions parents ask

### Is Peabody safe for kids to use alone?

Peabody is designed for it: every idea and story is screened in several layers, kids on their own are capped at spooky, there's no chat with an AI, and anything in question is held for a person to review. No AI system is perfect, so a child can also report any story.

### What happens if my child types something inappropriate?

No story is written. Peabody answers with a kind redirect, like suggesting a spooky mystery instead, never a scolding.

### Who reviews flagged stories?

Aaron, who runs Peabody, reviews every held story and approves, blocks or regenerates it. Until then nobody sees it.

### Can stories be scary?

Only up to the child's ceiling. Without a grown-up's account, stories are gentle until 1st grade and never past spooky. With a family account, a grown-up chooses each child's ceiling.

### Does Peabody use my child's name in stories?

A child without an account has no name in Peabody at all; they're a pen name. For kids in a family account, a real name never appears outside the family, not even inside a story.

---

Free at peabody.kids/blog/safe-ai-story-generator-for-kids