



Parents, Platforms, and the School Day

Youth digital services at the start
of the 2026-27 school year

As of August 31, 2026

Summary

The 2026-27 school year opens with a youth internet that is being segmented by age in ways that were mostly aspirational only a few years ago. The debate isn't simply about whether minors should have access to digital services; teenagers already use social media, gaming platforms and AI at scale. The more consequential question is what a service does differently once it has reason to believe the person on the other side of the screen is 12, 15 or 17 rather than an adult.

That shift is visible across products that were once treated as separate categories. Meta's proposed multistate settlement would impose default daily limits, overnight restrictions, school-hour notification limits and stronger age assurance for users under 18 on Facebook and Instagram; Roblox has reorganized accounts by age and requires an age check before chat; New York has finalized rules restricting algorithmically personalized feeds and nighttime notifications for minors without parental consent; and OpenAI began rolling out ChatGPT for Teens on August 18. Schools are drawing a parallel boundary around the device itself, with the National Conference of State Legislatures reporting as of July 15 that 42 states, Washington, D.C., and Puerto Rico had enacted legislation related to cellphone use in schools.

These rules govern different layers of the same experience: the device, age assurance, the teen account, parental override, and then chat, recommendations, purchases and notifications. A phone pouch can take the handset out of a student's hands without changing the account; an age check can route that student into a different product without changing the school rule; and a parent can be given authority over a time limit without being allowed to read a teenager's messages. The emerging regime isn't one rule for "screen time." It is a stack of controls applied at different points in the product.

Teen use is high enough that these distinctions matter. Pew Research Center reported in December 2025 that roughly nine in ten U.S. teens use YouTube, while about six in ten or more use TikTok and Instagram and 55% use Snapchat; about one in five say they use TikTok and YouTube almost constantly. Sixty-four percent report using AI chatbots, including about three in ten who use them daily, while Pew has separately found that 85% of U.S. teens play video games and 41% play at least daily.

The evidence on health effects remains more complicated than the political debate sometimes suggests, and the platform differences are substantial. Research has found associations between some forms of digital-media use and depression, self-injury, behavioral problems and other outcomes, while teenagers also report benefits involving friendship, entertainment and connection. In Pew's 2026 platform study, 37% of teen TikTok users said the service hurt their sleep, compared with 26% of Snapchat users and 24% of Instagram users; 44% of Snapchat users said the service helped their friendships, and most users of all three platforms said the services neither helped nor hurt their mental health.

The regulatory response therefore doesn't require a single verdict on whether social media is good or bad for adolescents. Legislatures, regulators, schools and companies are making narrower choices about time limits, private contact, age verification, recommendations, purchases, overnight access and parental authority. Age is becoming a product variable rather than merely an eligibility question; once a service identifies a minor, more of the product is expected to change automatically.

The youth internet is becoming a different product

Social media remains a routine part of adolescence, but the category is increasingly too broad to describe the policy problem by itself. Pew's April 2026 survey found that 57% of teen Snapchat users send direct messages on the service daily, while TikTok users were more likely than Instagram or Snapchat users to say they spent too much time on the service and to report negative effects on sleep and productivity. Those differences point toward the same conclusion now visible in regulation: a ranked feed, direct messaging, disappearing content, notifications and an in-app marketplace aren't interchangeable simply because they appear on the same phone.

Product design is therefore where much of the current policy action sits. A rule aimed at recommendations can materially change TikTok without touching Discord's central function; a rule aimed at private contact can materially change Snapchat or Discord without reducing time spent on YouTube; and a purchase restriction addresses a different problem from either of them. The new youth regimes increasingly separate attention, communication, commerce and age-based access instead of assuming that every minute of digital use creates the same risk.

Meta is turning the teen account into a regulated product

Meta's August settlement is the clearest current example because it turns a collection of youth-safety features into a detailed operating regime. The multistate agreement was announced on August 26 and, as of August 31, remains subject to judicial approval; it covers 47 states, the District of Columbia, Puerto Rico, American Samoa and the Northern Mariana Islands, while Florida and New Mexico are not parties. Texas reached a separate settlement with Meta the same day that requires similar youth protections and more than \$1 billion in payments.

New Mexico is proceeding under a separate final judgment. A jury imposed \$375 million in civil penalties in March, and the court added a \$567 million abatement award in August, bringing Meta's total financial exposure in the state to \$942 million. The five-year order includes stronger age-assurance measures, a 90-hour monthly cumulative limit for users under 18 across Facebook and Instagram, and restrictions on push notifications overnight and during school hours; Meta has said it plans to appeal, and the judgment provides that the five-year abatement period is tolled during an appeal if Meta posts a supersedeas bond.

The multistate agreement uses a different structure. Users under 18 would begin with a default cumulative limit of two hours per day across Facebook and Instagram, and a parent would have to approve removal of that limit; feed access would be restricted between midnight and 6 a.m., most notifications would be silenced overnight, and most push notifications would be muted from 8 a.m. to 3 p.m. Monday through Friday between August 15 and June 15, subject to adjustment by a supervising parent. Direct messages and specified safety and account-security alerts are treated differently, which matters because the settlement isn't trying to regulate every feature through a single clock.

Teen users would also receive screen-time prompts after each 15 minutes of continuous use and again at 60 and 90 minutes of cumulative use; they could select a feed that isn't personalized by Meta's recommendation systems, and parents could require that setting. Autoplay could be disabled, like and reaction counts would be hidden by default, teen accounts would remain private by default, and Meta would strengthen efforts to identify accounts that appear to belong to users ages 13 to 17 even when the listed birthday says otherwise.

Most settlement terms run for ten years, while the initial two-hour Daily Limit and midnight-to-6 a.m. nighttime feed restriction run for five. If the agreement's industry-wide adoption conditions are met, those restrictions extend to ten years and tighten to a one-hour limit per Meta platform and a 10 p.m.-to-7 a.m. nighttime feed restriction; the filed agreement identifies Snap, TikTok and YouTube as the core industry members whose adoption is needed to trigger those stricter Phase II rules. Meta has separately described approximately \$5.3 billion of the broader payment package as contingent on YouTube and TikTok implementing specified protections and making corresponding payments.

YouTube and the substitution problem

YouTube deserves separate treatment because it remains the most widely used platform among U.S. teenagers. Roughly nine in ten teens report using it, roughly three-quarters use it daily and about one in five say they use it almost constantly. Google already applies age estimation and supervised-teen protections; Family Center allows a parent to link a supervised account and manage digital-wellbeing settings, including a daily limit on Shorts scrolling that can be set to zero.

Those tools aren't equivalent to Meta's proposed default two-hour cap, and that difference explains why Phase II matters. A teenager who reaches a limit on Instagram can move to YouTube, TikTok, Snapchat, Discord or another service, so a restriction imposed on one platform can change where time is spent rather than reduce total use. The settlement's industry trigger is an attempt to make the attention rules travel across competitors rather than simply move the teenager from one app to the next.

The school day is becoming its own digital boundary

The Meta agreement arrived during a rapid expansion of school cellphone restrictions, creating a second layer of youth technology policy that operates independently of the platform itself. As of July 15, the National Conference of State Legislatures counted 42 states, Washington, D.C., and Puerto Rico with enacted legislation related to cellphone use in schools, although that number shouldn't be read as 42 statewide bell-to-bell bans. The category includes local-policy requirements, state guidance, restrictions limited to instructional time and funding for implementation tools such as phone pouches; state laws also commonly preserve exceptions involving special education, health needs, emergencies and translation.

The school-day boundary now has several possible enforcement points, and they aren't interchangeable. State law or district policy can determine when a student may possess or use a phone; a school can enforce that rule through classroom procedures, network controls, district-managed devices or physical storage; and a consumer service can separately change what reaches the device during the same hours. Under Meta's proposed settlement, for example, most Facebook and Instagram push notifications to teens would stop during defined school hours even if the handset remained physically available.

Public support has moved in the same direction. Pew reported in July that 77% of U.S. adults supported prohibiting middle- and high-school students from using cellphones during class, while 48% supported a ban lasting the entire school day and 43% opposed one. The policy question is therefore no longer limited to whether schools should restrict phones; it increasingly includes whether the services themselves should change behavior during the hours a student is supposed to be in class.

Gaming and contact risk show the limits of a screen-time rule

Gaming presents a different product architecture because multiplayer games and virtual worlds combine entertainment with text or voice chat, digital purchases, user-generated environments and relationships among people who may not know one another offline. Roblox has reorganized its account structure around age, with Roblox Kids covering users ages 5 to 8 and Roblox Select covering ages 9 to 15; content access, communication defaults and parental controls differ by account type, and the company requires an age check before users can access chat. When Roblox announced the new structure in April, it said more than 65% of its U.S. daily active users had completed an age check.

Gaming also adds commerce to the same environment. Virtual currency, subscriptions, in-game purchases and marketplaces allow minors to spend money inside products that also use engagement systems and social features; prior Federal Trade Commission cases involving Fortnite and Genshin Impact illustrate how youth privacy and purchase practices can converge in the same enforcement theory. The House-passed KIDS Act similarly includes parental-control provisions for purchases on covered gaming platforms.

Contact risk doesn't map neatly onto time spent, either. The National Center for Missing & Exploited Children reported more than 50,000 financially motivated sextortion reports in 2025, or roughly 137 per day, and the pattern often begins with an offender establishing contact with a minor before moving the interaction into a more private channel. A daily cap can address one form of engagement concern without changing who can contact a teenager; a private-account default can reduce unwanted discovery without addressing purchases; and a school-phone rule can remove the device for several hours without changing the account settings that govern the rest of the day.

Age assurance is becoming infrastructure

Almost every age-specific rule depends on the same underlying capability: the service has to know, with some useful degree of confidence, whether the user is a minor. A platform can't reliably provide one product to a 14-year-old and another to an adult if its only signal is a birth date typed during account creation, so age assurance is moving from a compliance afterthought toward a core piece of product infrastructure. Services can rely on self-declared age, infer age from account signals, estimate age from physical or behavioral features, verify age against identity information or use a parent or another credential provider; each method offers a different level of confidence and requires a different amount of personal information.

The Federal Trade Commission adjusted its COPPA enforcement policy in February to reduce one obstacle to stronger age checks, saying it would not bring specified COPPA enforcement actions against eligible

general-audience or mixed-audience services solely because they collect personal information for qualifying age-verification purposes without first obtaining parental consent, provided the service complies with the policy's privacy, security and retention conditions. New York has gone further: final SAFE for Kids Act regulations require covered platforms to determine that a user is an adult before providing algorithmically personalized feeds or nighttime notifications without parental consent, and the law takes effect January 25, 2027.

New York's rules also expose the central tradeoff. Covered services must offer at least one age-assurance method that doesn't require government identification, limit the use of information collected for the age check, delete or de-identify it after its intended use and test age-assurance methods annually. No method eliminates evasion, however; users can lie about a birth date, create a second account, borrow an adult credential or use a device that has already been verified for someone else, while shared family devices create the separate problem that an accurate determination about the account holder doesn't necessarily identify the person holding the device at a particular moment.

Stronger systems can reduce those failures by combining several signals, but that introduces privacy and accuracy costs of its own. Adults can be incorrectly classified as teenagers, teenagers can be incorrectly classified as adults, ID-based systems can disadvantage users without accepted documentation, and parental-consent systems assume that a reachable adult can complete the process. The emerging model therefore asks platforms to know more about age while simultaneously limiting what they collect, retain and reuse to make that determination.

AI is entering the same framework while federal law remains fragmented

AI has become ordinary teen technology quickly enough that youth policy can no longer treat it as a future category. Pew found that 64% of U.S. teens use AI chatbots; 54% have used them for schoolwork, 47% for entertainment, 16% for casual conversation and 12% for emotional support or advice, while one in ten said chatbots help with all or most of their schoolwork. The school context is already substantial: 59% of teens said using AI to cheat occurs at least somewhat often at their school, and among teenagers who have used chatbots for schoolwork, 76% said students at their school use them to cheat at least sometimes.

Conversational AI creates a different product question from a social feed because the system generates the response itself rather than merely ranking material produced by users and publishers. OpenAI began rolling out ChatGPT for Teens on August 18; eligible accounts identified as belonging to users under 18 through account information, verified age or age prediction are automatically placed into the teen experience. Linked parents can manage selected features and set quiet hours, study hours can cause eligible new conversations to begin in Study Mode, and parents cannot read or monitor a teenager's conversations through those parental controls.

Congress has moved further on children's online safety in 2026 without producing a single federal framework. The House passed the Kids Internet and Digital Safety Act on June 29 by a vote of 267-117, while the Senate Commerce Committee advanced its version of the Kids Online Safety Act by voice vote on August 5; the two chambers still differ on important elements, including the Senate's duty-of-care approach. The same committee also advanced several youth-AI measures on August 5, including the CHATBOT Act, but none of those measures has become law.

The constitutional boundaries remain unsettled. In *Moody v. NetChoice*, the Supreme Court held that a platform's choices about how to organize and present third-party content can involve protected editorial judgment, without deciding the ultimate validity of the Florida and Texas laws at issue. In *Free Speech Coalition v. Paxton*, the Court upheld Texas's age-verification requirement for adult-content websites. In March 2026, the Ninth Circuit allowed several provisions of California's Age-Appropriate Design Code involving coverage and age estimation to proceed while a separate data-protection impact-assessment requirement remained enjoined.

Companies are therefore building age controls, feed restrictions and youth defaults while legislatures and courts are still defining the outer legal boundaries. The same engineering is also becoming global: the United Kingdom and European Union impose separate youth-safety obligations, while Australia has required specified social-media platforms since December 10, 2025, to take reasonable steps to prevent users under 16 from holding accounts. Those regimes aren't identical, but a feature built for one statute can become part of the common technical infrastructure used elsewhere once the company has paid to build it.

The operating baseline for the 2026-27 school year

There is still no single standard for youth technology, and there probably won't be one soon. Social media, gaming, AI, school software and family-safety products serve different purposes; the available evidence doesn't support treating every hour of digital use as equivalent, while the current legal landscape doesn't support a uniform rule across every service. What is becoming clearer is the operating model underneath those differences.

Age is now being used as a product variable rather than merely an eligibility question. Younger users receive different defaults or access to different features; parents can control settings that minors cannot always change themselves; time of day and school hours are appearing directly in product rules; and communication, recommendations and commerce are increasingly treated as separate design decisions. The burden on parents is changing as well: earlier parental-control systems often required an adult to discover the feature, link an account and turn restrictions on, while Meta's proposed settlement starts a teenager with a time limit and requires parental permission to remove it, New York starts from restricted feeds and nighttime notifications unless adulthood or parental consent is established, Roblox requires an age check before chat, and ChatGPT for Teens is automatically applied to eligible accounts identified as belonging to minors.

The model remains uneven. Meta's multistate settlement still awaits judicial approval and excludes Florida and New Mexico, Texas is proceeding under a separate settlement, YouTube's current supervision system doesn't impose the same default limits Meta has proposed, state school-phone laws differ substantially, and Congress has not enacted a comprehensive federal youth-safety statute. Those differences increasingly sit on top of a shared product problem rather than replacing it.

For a service operator, the questions are concrete: how much age information is necessary and how reliable must it be; which settings should change automatically; which choices should belong to a parent; which features require a higher level of age assurance; who should be able to contact a minor; when should notifications or access be reduced; how should purchases work; and, for an AI system, how should the system's own behavior change when the user is younger? Each answer creates a product decision, an enforcement mechanism and a new place for regulators to test whether the company's stated youth protections match the experience it actually delivers.

That is the clearest change as the 2026-27 school year begins. The old model treated youth safety largely as a gate around the product: children under 13 were barred, parental controls existed somewhere in settings, and the service otherwise remained substantially the same. The new model is moving the boundary inside the product itself; once a user is identified as a child or teenager, the service is increasingly expected to know that fact and act on it.