

# Startups, Content Moderation, & Section 230



Debates about the intermediary liability framework provided by Section 230 have animated policy conversations as lawmakers grapple with harmful online content, including around election integrity, health information, and children’s safety. But those debates are almost exclusively focused on the largest Internet companies. Section 230, however, applies to all services of all sizes that host all types of user-generated content, including startups.

Section 230 helps startups avoid being inundated with lawsuits over their users’ speech and limits potentially-ruinous legal costs. Startups still have incentives to invest their limited time and resources in content moderation, including, for example, to ensure that content appearing on their site is useful and relevant to their users or within their terms of service. In fact, Section 230 ensures they won’t be held liable for users’ speech even though they’re active moderators. Despite startups’ efforts, content moderation is inherently imperfect. Placing even higher, unrealistic expectations on startups—such as opening the door to lawsuits when a startup inevitably fails to perfectly and immediately remove harmful content—could take content moderation costs from burdensome to catastrophic, or even push startups to avoid hosting user content entirely.

To better understand how startups moderate content on their services, how that differs from mid-sized online service providers, and the value of Section 230 for startups, we surveyed and had conversations with user content-hosting startups in the Engine network, mid-size online service providers, and attorneys that work on 230-related cases. (We originally released a document on the costs of 230-related litigation in 2019. We confirmed that the figures below are accurate as of 2021.) As the responses show, startups have limited resources to moderate content on their sites, but they spend more per user than mid-sized content-hosting companies. And, even with Section 230 in place, defending against lawsuits involving user speech online can quickly become expensive.

# Startups & Content Moderation

Startups monitor and moderate content on their sites because they recognize the potential for problematic content to appear that might contradict their values, undermine the trust of their other users, or threaten their ability to grow. Most of the startups we spoke with do not yet encounter problematic content at a rate that requires a large moderation team or expensive, sophisticated moderation technology. However, as startups scale, they begin to encounter more content requiring their attention.

The startups we spoke with each enable or host user-generated content, but they do not have the same business models as one another. The companies are between 2 and 7 years old, generate less than \$100,000 in annual revenue, have fewer than 10 employees, and serve between 1,000-5,000 monthly active users. Each of the companies had raised \$50,000 or less in publicly-announced funding through grants, pitch competitions, crowdfunding, and small, formal funding rounds—except one company that had raised \$100,000 or less.

The responses reveal that it is critical for startups to have the ability to moderate content on their services as they see fit according to their specific size and need. Thanks to the varying need and resources put toward moderation at the current point in their lifecycle, the cost per user of each startups' moderation efforts ranged from a few dollars to over \$150 per user.

## Startups spend thousands of dollars on human content moderation

For most of the startups we surveyed, moderation is conducted by humans, on an as-needed, case-by-case basis. Moderation did not comprise any startup employee's entire job, mostly due to scale. If they allocated resources toward it, companies spent up to \$10,000 annually on training for the employees moderating content.

## Startups spend tens, or even hundreds, of thousands of dollars on content moderation technology

Most of the companies we spoke with did not use technology as part of their moderation processes, because moderation technologies were unwarranted due to scale, were prohibitively expensive, and are ultimately imperfect, requiring human review as a backstop. As one startup founder noted in discussing the costs of content filtering technology, if its use were required by law, "it would put us out of business." However, most of them plan to use technology to assist moderation efforts in the future, as warranted by scale.

For the companies that do currently use technology, the amount they had spent over the company's lifetime developing their technology varied widely, from \$40,000 all the way up to \$1,000,000. Those startups spent up to \$50,000 annually maintaining their proprietary technology.

Some startups license technology that is developed by others to support moderation efforts on their sites, but it does not make for a low-cost alternative to developing moderation technology in-house. Licenses for software used by the startups we spoke with can cost up to \$10,000 annually, and licensed software must be integrated into their service, which can be a one-time expense of up to \$10,000.

# Mid-Sized OSPs & Content Moderation

As user-content hosting companies scale, the amount of user-generated content that must be moderated grows. Mid-size online service providers (OSPs) have standard processes, dedicated staff, and licensed and proprietary technologies to help moderate content on their sites.

The mid-size OSPs we spoke with are between 11 and 15 years old, generate more than \$50 million to more than a billion dollars in annual revenue, employ 100 to 5000, and serve almost a million to just under a half-billion monthly active users. Like the startups we spoke with, they each enable or host user content but have varying business models.

The OSPs' responses underscore the investments they make in moderating content on their sites. While they spend much larger sums on moderators and moderation software than startups, thanks to economies of scale, their cost of moderation on a per-user basis is lower. As shown in [similar research](#), the companies' per-user moderation costs ranged from as little as a number of cents to a few dollars.

## Mid-sized OSPs spend millions on human content moderation

All of the OSPs we spoke with employ human content moderators, and just one of the companies have moderators where moderation is not their only job responsibility. The companies employ up to 250 human moderators, with most OSPs employing fewer than 50 moderators. Over half of the OSPs utilized external contractors as moderators. Most of the OSPs spent between \$1,000,000 and \$5,000,000 annually to retain their moderators, who tend to stay at the company between 1 and 3 years. The companies spend up to a quarter of a million dollars training and equipping their moderators annually.

In the course of moderating content on their services, OSPs' trust and safety teams responsible for moderating content often collaborate with other departments at the company like legal, policy, and public relations. These cross-company collaborations often follow controversial or high-profile moderation decisions and could represent up to 10,000 work hours annually, the full cost of which is difficult to estimate given the varying salaries and opportunity costs implicated.

## Mid-sized OSPs spend millions, or even tens of millions, of dollars on content moderation technology

To support their human moderators, the OSPs each utilize proprietary technologies they've developed, and some additionally license moderation technology. The companies spent between \$500,000 and \$30,000,000 developing their proprietary moderation tools. That development involved a similarly wide range of estimated engineering work hours—between 500 and 300,000. Companies spent between \$1,000,000 and \$5,000,000 to maintain their proprietary tools on the high end, and less than \$250,000 on the low end.

Companies annual spending on licensing moderation tools similarly ranged between \$1,000,000 and \$5,000,000 on the high end, and less than \$250,000 on the low end. For most OSPs, integrating the licensed technology required fewer than 200 engineering work hours and cost less than \$50,000. One company spent over \$2,000,000 integrating licensed technology, requiring nearly 2,000 engineering work hours.

# Section 230 & Litigation Costs

Section 230 of the Communications Decency Act is often credited with the creation of the modern Internet by enabling a diverse, vast spectrum of Internet companies to host user-generated content. The law was created after court cases in the 1990s extended traditional distributor liability frameworks to Internet companies that did not moderate content on their sites but found Internet companies that engaged in any moderation to be liable for all of the content they hosted, effectively creating a disincentive to engage in moderation.

The law has two key provisions: The first ensures that Internet companies cannot be held liable for content created by their users. The second ensures that liability limitation applies even if a company engages in moderation “in good faith to restrict access to or availability of material that the provider or user considers to be obscene, lewd, lascivious, filthy, excessively violent, harassing, or otherwise objectionable, whether or not such material is constitutionally protected.”

Taken as a whole, the law gives companies the needed certainty to host user-generated content by helping to quickly dismiss cases when one user wants to sue over the content another user has created. This is especially important for startups that host user-generated content like comments, photos, classified listings, and more, enabling them to operate when they might otherwise be quickly overrun by costly lawsuits or even threats of costly lawsuits.

But even with Section 230, however, frivolous lawsuits can still pose crippling costs for startups. Section 230 does not deter all lawsuits—meritorious or not—and defending them can be especially expensive and burdening for startups, even under the current law. The value of Section 230 is that Internet companies can use it as a clear-cut affirmative defense early on in the litigation process, which helps avoid legal expenses that pile up as litigation progresses.



**Pre-complaint: \$0 to \$3,000:** Threats of litigation can present costs to startups even before a lawsuit is filed. Responding to a threatening demand letter based on user speech can carry legal costs of up to \$3,000. Startups have strong incentives to resolve disputes before they turn into lawsuits, given the legal and potential reputational costs of defending an even meritless suit. Parties sending such letters are likely to know their claims lack merit because of Section 230, and are typically seeking to extract a nuisance-value settlement.

Beyond offering a response, receiving a demand letter creates additional burdens for startups as well. If a company believes a lawsuit is likely, they are obligated to issue a litigation hold and preserve documents and information that may be related to the case. A litigation hold can be burdensome and distracting for startups with limited resources.



**Motion to dismiss: \$15,000 - \$80,000:** Filing a motion to dismiss is a startup's first opportunity to end a lawsuit once one is filed. In the motion, a startup must show that it is not liable for the speech at issue, even if the plaintiff's claims are true. If the plaintiff alleges a user posted the content at issue in the lawsuit, the startup is likely to be successful dismissing the lawsuit on Section 230 grounds, since the law establishes the startup is not liable for user speech it did not create. However, that does not mean the lawsuit did not burden the startup. A motion to dismiss generally carries a cost between \$15,000 and \$40,000 but could go as high as \$80,000. That is nearly a month of operating resources for the [average seed-stage startup](#)—meaning a meritless lawsuit could deplete an entire month's worth of a startup's resources. And the average seed-stage startup is already relatively successful and well-resourced, considering how few startups receive that kind of outside funding. For many of the startups we spoke with, that could consume their entire funding to date. Beyond these costs, plaintiffs are generally allowed the opportunity to amend their claim, meaning an intentionally-deceitful plaintiff could amend its complaint to allege the startup did create the content at issue, allowing the case to proceed and creating future legal costs for the startup.

**Early Motion for Summary Judgment: \$15,000 - \$150,000+:** Parties can file a motion for summary judgment when there are undisputed facts upon which the court can decide the case without the need for a full trial. A startup might file an early motion for summary judgment in cases where the outcome rests on a few factual questions, such as who is responsible for content posted online. Filing such a motion generally involves minimal discovery, limited to information about the user's identity and the role of the startup in creating the content at issue, but can still create legal costs around \$30,000. In addition to those costs, preparing and filing the motion can cost between \$30,000 and \$70,000.

While some attorneys prefer to file an early motion for summary judgment rather than a motion to dismiss—because courts rarely grant the motion without giving the plaintiff a chance to correct the pleading—early motions for summary judgment come with risks as well. Filing an early motion may forfeit a startup's right to file one later, since courts tend to disfavor or prohibit multiple motions for summary judgment. Failure to get a case dismissed on summary judgment means the parties must litigate through trial or settle. Both can be incredibly expensive options that deeply burden startups.

**Discovery and Trial: \$100,000 - \$500,000+:** Lawsuits against websites for user speech rarely proceed through discovery and to trial because the associated legal costs are likely to exceed potential liability. There is no fee-recovery in 230 cases, meaning each party pays their own legal fees regardless of who wins. These dynamics incentivize resource-strapped startups to settle, even if they are likely to win. Those cases that do proceed to trial, however, quickly reach six-figure costs.



Engine is a non-profit technology policy, research, and advocacy organization that bridges the gap between policymakers and startups. Engine works with government and a community of thousands of high-technology, growth-oriented startups across the nation to support the development of technology entrepreneurship through research, policy analysis, and advocacy. When startups speak, policymakers listen.