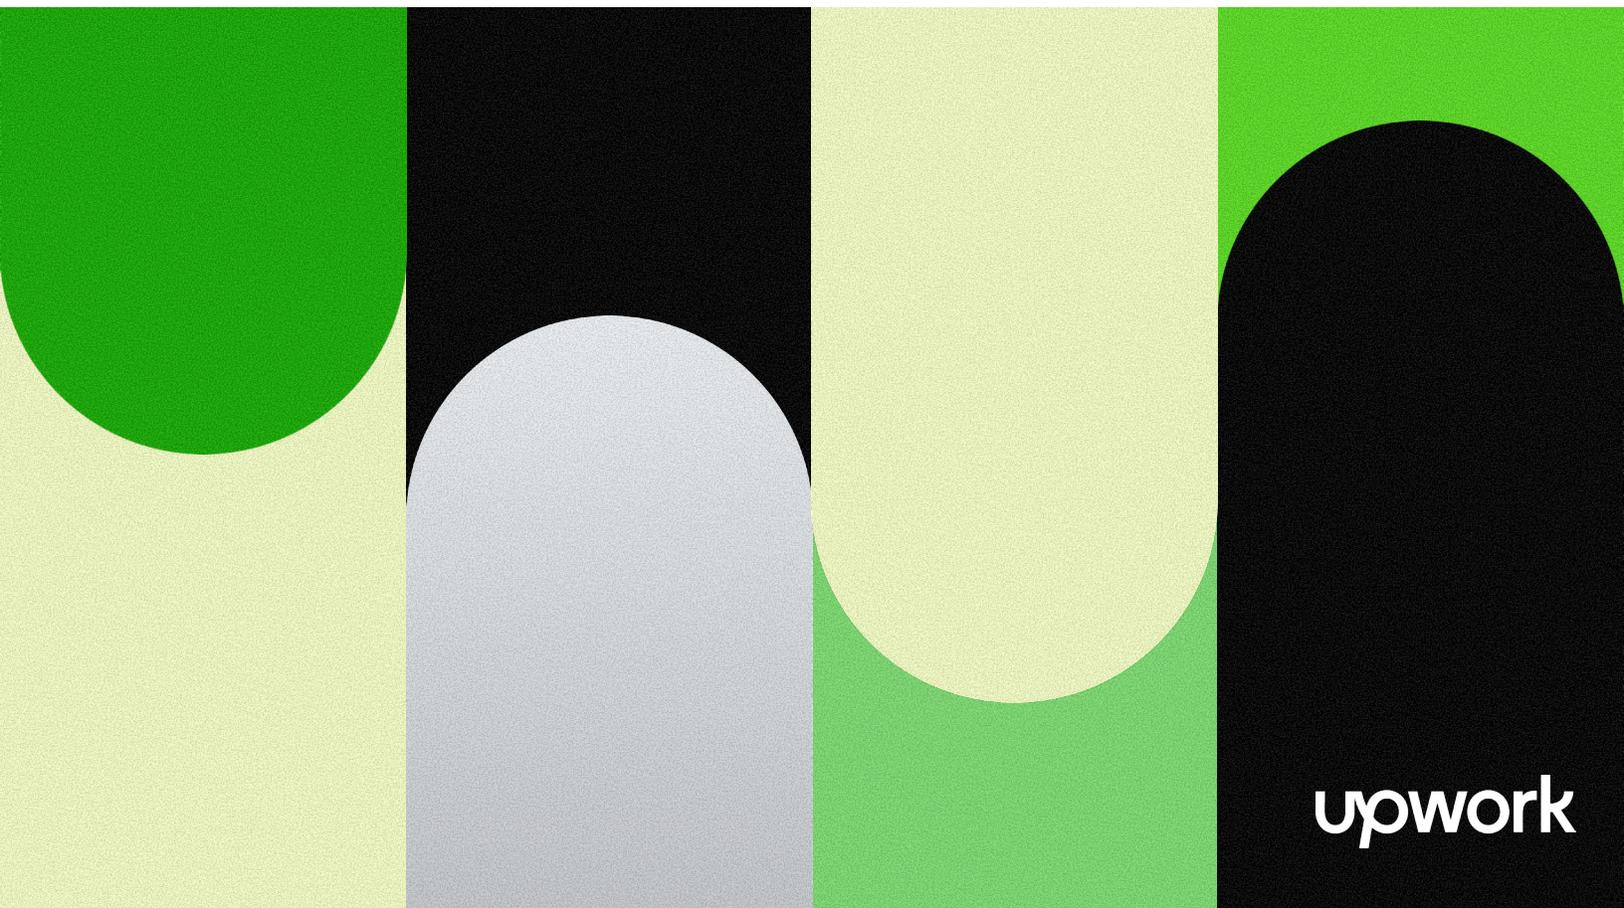


2025

Upwork Marketplace Transparency Report



upwork

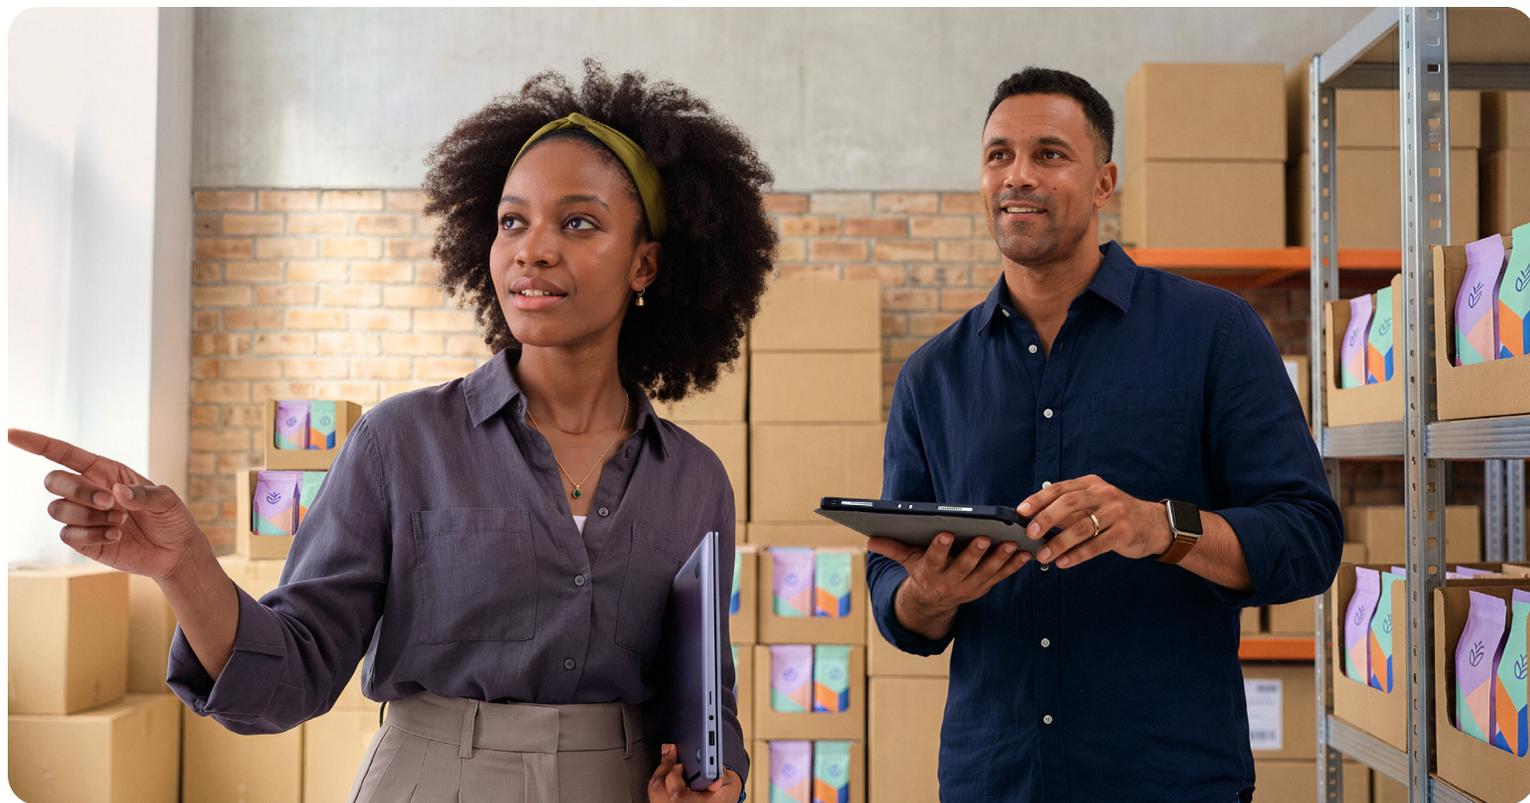
The **Digital Services Act (DSA)** is a European Union (EU) law that became fully effective on February 17, 2024. Generally speaking, the DSA is intended to protect consumers and their fundamental rights online; establish a transparency and accountability framework for online platforms; and foster innovation, growth, and competitiveness.

In accordance with the DSA, this report, the Upwork Marketplace's second annual Transparency Report, will share an overview of our efforts to elevate trust through the enforcement of our **Terms of Service, Marketplace Standards,** and applicable legal requirements.

Published: February 26, 2026

Reporting Period: January–December 2025

Trust & Safety at a glance



The Upwork Marketplace connects businesses with on-demand access to highly skilled talent across the globe. Businesses rely on us to find and hire expert talent, leverage AI-powered work solutions, and drive business transformation.

Trust is what makes that possible. We know trust has to be earned through action, consistency, and relentless commitment to online safety. This responsibility is at the core of how — for over two decades — we’ve connected businesses and professionals from around the world and enabled them to work together securely. And this responsibility continues to guide how we design, govern, and protect our marketplace today.

To deliver scalable, principled protections, our Trust and Safety program starts with clear policies, including our Terms of Service and marketplace standards that define expectations for safe, ethical behavior on our platform.

Our product and engineering teams then translate those standards into user experiences, building tools that surface positive interactions and discourage harmful ones. Intelligence analysts and investigators monitor our platform for emerging threats and complex fraud patterns, and use their learnings to continuously improve our systems.

Supporting this work, data scientists, researchers, and designers generate deep insights into user behavior and preferences, helping us continuously refine and strengthen trusted experiences at scale.

At its core, the Trust and Safety team works to:

- **Promote the good** by building reputation systems that highlight the best of our marketplace to help our users make better matches
- **Detect and disrupt the bad** through behavioral analytics, threat intelligence, advanced machine learning, and expert review
- **Enforce quickly and fairly** when we identify a threat, and offer plain-language education and clear, fair appeal paths

Our progress in 2025

Throughout 2025, we advanced our strategy of embedding trust into every interaction, and our efforts delivered measurable results. There's more yet to do, but the following highlights show how we raised the bar on safety while maintaining a smooth path for legitimate work.

Record-low fraud exposure

Our persistent investments in always-on defenses drove meaningful safety gains. Offers sent by malicious client accounts on the Upwork Marketplace declined more than 91% year over year, and bad-actor freelancer proposals fell 82% year over year.

Offers sent by malicious client accounts on the Upwork Marketplace declined more than

↓ 91% yoy

Bad-actor freelancer proposals fell

↓ 82% yoy

Identity verification when it matters most

We expanded opt-in identity verification by introducing it earlier in the freelancer journey, allowing new talent to establish a visible trust signal from day one. As a result, we increased the number of identity-verified freelancers on the Upwork Marketplace by 22%.

We increased the number of identity-verified freelancers on the Upwork Marketplace by

↑ 22%

In 2026, we'll continue expanding and improving on this experience to help clients hire with greater confidence.

The Account Health Hub

In 2025, we launched the [Account Health Hub](#), a centralized dashboard where freelancers and clients can view the full picture of their account's standing — what's restricted, why, and how to fix it. For the first time, Upwork Marketplace customers can:

- See any active limitations on an account
- Access clear guidance and direct links to resolve each issue
- Review their enforcement history
- Learn how repeat violations may affect future platform access

This transparency helps customers resolve issues proactively and keeps compliant customers from being sidelined unnecessarily. It also helps demystify policy enforcement by showing that behind every policy is a clear process designed to keep our customers and marketplace safe.

Nuanced content screening at scale

In 2025, we enhanced LLM-based moderation for job posts. Calibrated to our policies and paired with human review, these systems now proactively identify more violations while preserving context and fairness, catching nearly 40k violating jobs that weren't detected by other models or reporting mechanisms.

Stronger sign-in protection and instant security alerts

We added more strategic reCAPTCHA checks at sign-in and other sensitive moments, such as password resets and phone verification, to stop bots and automated attacks without slowing down legitimate users. When two-step verification is added or removed, we send real-time email and push alerts so users can secure their accounts quickly.

What's next

In 2026, we will further reduce friction and increase confidence across every interaction. Planned investments include a refreshed reputation framework, in-flow policy guidance, a new reporting experience, and deepened fraud defense.

Stronger reputation signals for all

Through 2025, customers expressed that existing reputation signals, such as scores, reviews, and badges, were often hard to interpret and didn't consistently reflect true performance or experience.

In 2026, we're rebuilding the reputation system to be clearer, more transparent, and grounded in measurable performance to improve signal quality and equity across the marketplace. This new framework will clarify what "good" looks like and provide stronger credibility indicators for clients, freelancers, and AI agents.

Right-on-time policy assistance

We're enhancing our platform with new tools designed to bring more policy context into relevant user flows. By delivering the right information at the right moment, freelancers and clients can avoid accidental violations and move forward with greater confidence. The result is fewer preventable issues and smoother interactions across the platform.

A faster, clearer reporting experience

When users encounter undesirable content our systems may have missed, they expect a clear path to report it and understand what happens next. To meet those expectations, we're redesigning the reporting experience from the ground up to make it clear what type of content to flag, where to flag it, and how those reports are being addressed.

Enhanced fraud defense

Bad actors evolve; so do we. In 2026, our investments will focus on strengthening identity and fraud defenses, making enforcement easier for users to navigate, and continuing to embed safety and security guardrails across every interaction throughout the marketplace. Our north star remains the same: make Upwork the world's most trusted work marketplace.



Deep dive: our policy process

On the Upwork Marketplace, we build a foundation of trust in part through policy — through clear standards that shape how people show up on our platform, the behaviors we allow, and how we respond when boundaries are crossed. These policies help protect against fraud, abuse, and misuse. Equally as important, they're designed to promote inclusion, professionalism, and the kind of behavior that supports good work.

01 The anatomy of a policy

Our Terms of Service create boundaries for user-generated content and behavior, but stating what's prohibited doesn't, by itself, create accountability. Policies turn words into action by defining how violations of our Terms of Service are detected, evaluated, and communicated to freelancers and clients.

Each policy includes:

- A clear **policy statement** and defined violation categories that outline allowed and prohibited behavior for a given area, such as violent content or discrimination
- Mapped **enforcement outcomes** based on marketplace risk, ranging from warnings and content removal to permanent account blocks
- **Investigation guidance** that helps reviewers and large language models consistently identify violations and assess the evidence required to take action
- User-facing **educational content** that ensures guidance is clear and actionable for freelancers and clients, shared through our marketplace standards, the Account Health Hub, and enforcement emails

02 A look inside the process

The work is led by our Trust and Safety Policy team, a group of experienced professionals with backgrounds in platform integrity, law, trust and safety operations, and risk management. They specialize in translating complex challenges into clear, enforceable standards. The team works closely with key stakeholders in the business to ensure we're solving the right problems at the right time.

Policy updates or audits are prioritized for many reasons, including:

- **Changes in the marketplace** such as new products or technologies that require tailor-made rules
- **New legal or regulatory obligations** including laws or agreements with third-party payment vendors
- **Emerging risks or enforcement gaps** identified by freelancers, clients, or internal teams
- **Policy age and review cadence** such that, even if no specific gap is identified, policies are revisited periodically to ensure they remain effective and relevant

Once we've identified the need to update a policy, we dig in. Our development process is collaborative by design and begins with listening to our customers, internal support teams, platform signals, and data. Our team analyzes feedback from support channels, social media, and user research to identify patterns, while historical data and reviewer-escalated edge cases help surface confusion and reveal where enforcement may be too strict or too lenient.

We also look beyond our own walls to civil society, academia, journalists, and peer marketplaces to help inform our approach, while the specific needs of our freelancers and clients remain central to every decision.

Another critical component to our process is our Policy Steering Committee, a group of leaders from core teams across the company. This group meets to assess the research and weigh potential approaches. They pressure-test ideas, surface tradeoffs, and ensure policy decisions are viewed from multiple angles, align with our values, and are operationally sound. Once a policy decision is approved, we translate it into product changes, reviewer guidance, and model features.



03 Clearer expectations for client job posts

The needs of our community steer many of our policy changes. For example, in 2025 we developed a new policy to tackle an emerging freelancer pain point: inaccurate or misleading job posts in the marketplace known as “ghost jobs.” We surfaced concerns about these ghost jobs through a tool that aggregates user feedback from multiple sources including support tickets, customer experience surveys, and social media.

Freelancers highlighted a challenge in which some clients posted jobs without the intention to hire, described work that differed significantly from what emerged during the interview process, or allowed the scope to drift well beyond the original agreement.

In response, we reviewed platform data, user feedback, support cases, and freelancer and client workflows to better understand how we could develop our processes to meet this need. The result is a clarified set of minimum quality standards for clients and their job posts that included new enforcement mechanisms aimed at catching these ghost jobs without having to wait for a freelancer complaint.

When our system identifies client accounts that consistently close jobs without hiring, the client first receives a warning. If the

ghost job posts persist, their account is then blocked. If, when they're blocked, the client has active jobs that have received proposals, we refund to the freelancer any **Connects** (virtual tokens freelancers use on Upwork to submit proposals for jobs) they may have invested in applying to the job.

In updating our processes, we also clarified our approach to misleading job posts. Clients are expected to stand by their job descriptions; when a posting is inaccurate or the scope changes without freelancer consent, we penalize the account with a warning. If the client does not change their behavior after being educated on our policies, we block the account.

By listening to freelancers and clients, clarifying expectations, and enforcing our policies when those expectations aren't met, we're building a more trustworthy marketplace that reduces guesswork and lets people focus on the work that matters.

04 Leading with consistency

In times of rapid change, prioritizing responsible decision making over speed becomes even more critical. Our experience shows us that real trust is built through consistency: in how we make decisions, how we apply our standards, and how we show up for our community. Even as the landscape shifts.

Operating at global scale

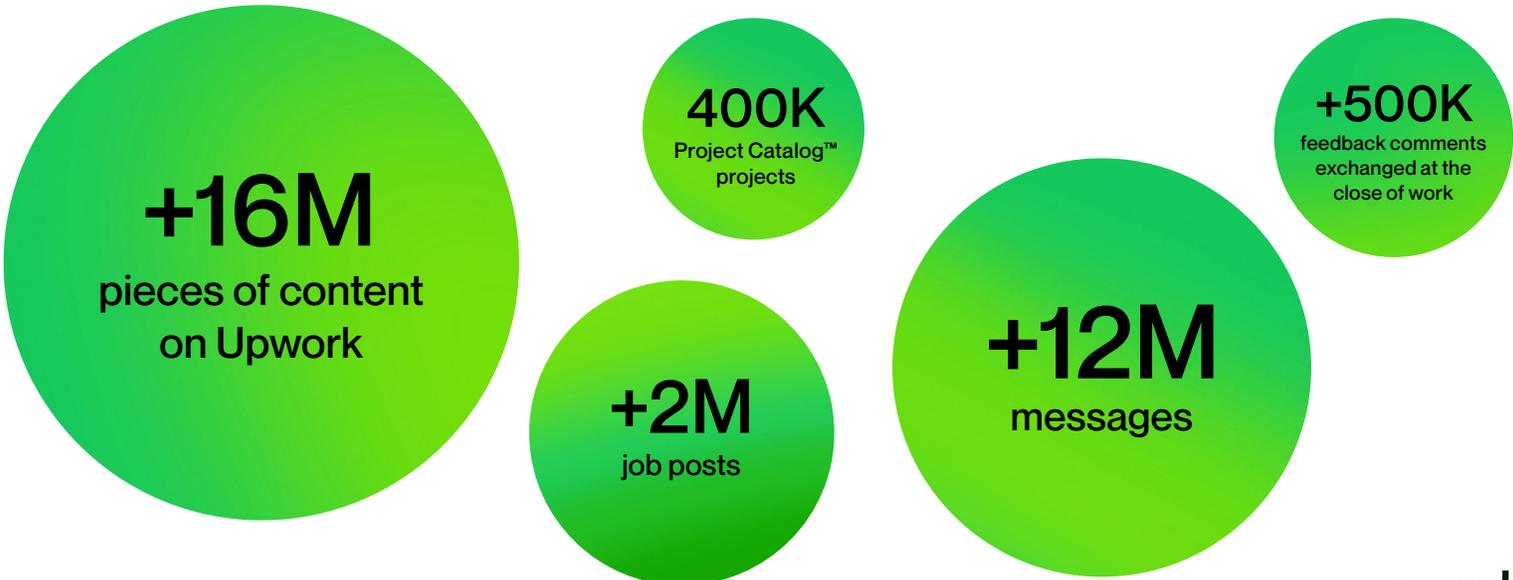


The Upwork Marketplace operates globally, with millions of freelancers and clients hiring, working, and communicating on the platform. At this scale, protecting the marketplace depends on strong systems that can identify issues early and respond effectively across languages, regions, and product surfaces.

Each of these interactions represents a working relationship between freelancers and clients. The sections that follow describe how Trust & Safety protects these interactions, starting with what users report and how we respond.

In 2025 alone, users generated **over 16 million pieces of content** on Upwork - including over 2M job posts, nearly 400K Project Catalog™ projects, more than 12M messages, and over 500K feedback comments exchanged at the close of work.

In 2025 alone,
users generated



User reports

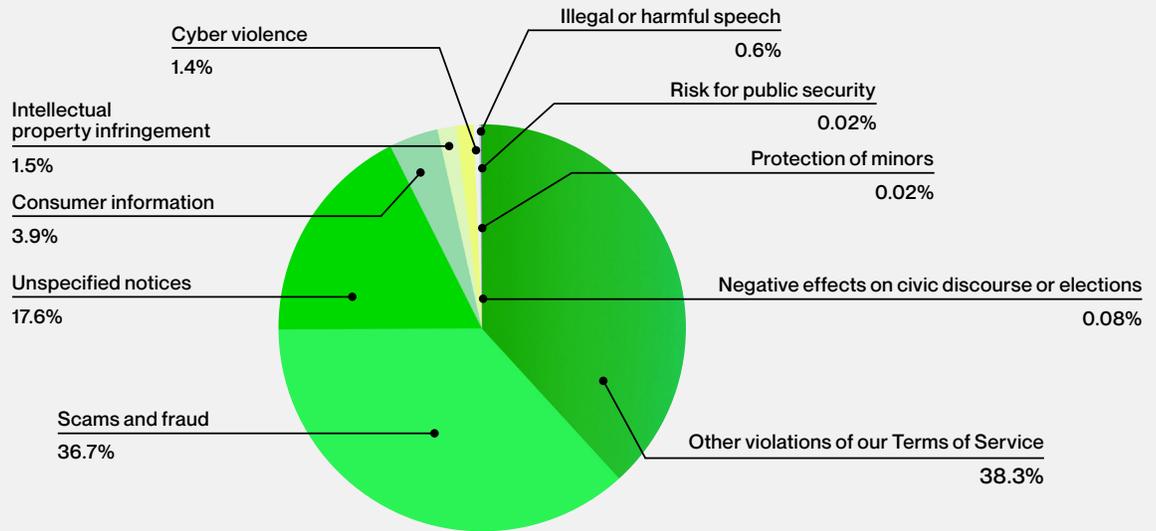
Clients and freelancers can report behavior or content they believe violates our standards. The [reporting feature](#) can be found on job posts and Project Catalog™ submissions, as well as in messages, profiles, portfolios, and feedback. Our Trust and Safety team reviews every report, and may remove content, provide education, suspend account access, or take other enforcement actions as appropriate.

When someone reports a potential violation, we now share updates with them when we complete our investigation.

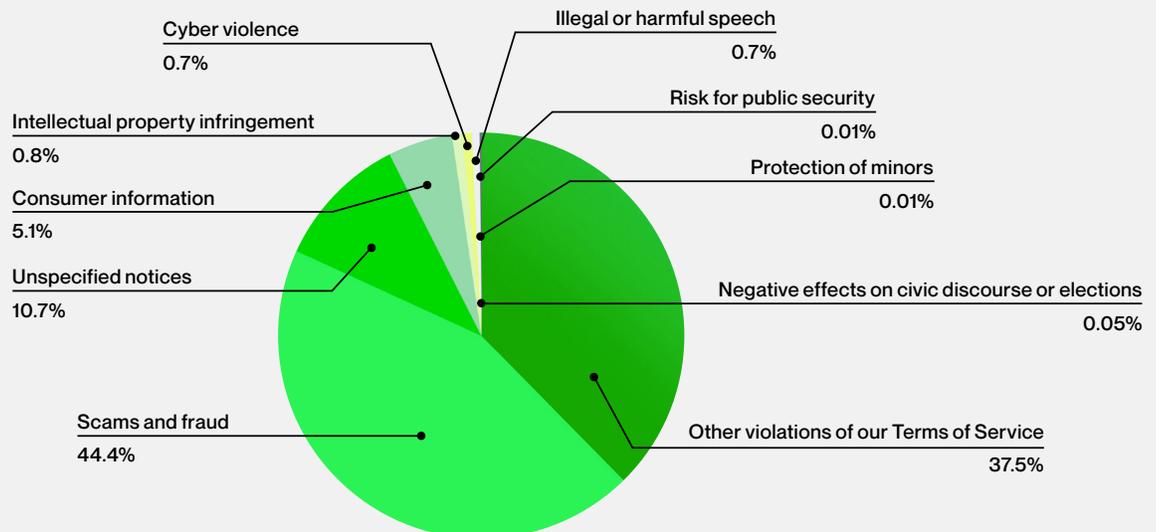
Once our teams have reviewed the details of the report against our policies, we'll get in touch over email and let the reporter know if we found a violation.

In 2025, we received 240,800 unique reports of potentially illegal content or violations of our terms and, following our investigation, we took action on 79,976 of them (representing less than 0.5% of all user-generated content on the platform). The median time it took to review a user report was 22 minutes. The Upwork Marketplace did not receive any notices from designated [DSA trusted flaggers](#).

Number of unique notices received



Number of actions taken



Automation and Upwork-initiated content moderation

We have designed our systems in part to enable users to easily report issues, and in part to be proactive and find issues before those reports are needed. We have developed automated tools to detect and prevent policy violations, to complement expert review, and to reduce exposure to scams and other harms.

We aim to continuously monitor for, and anticipate, issues that could impact someone’s finances, personal data, and peace of mind. We monitor a range of signals, discussed in the following sections.

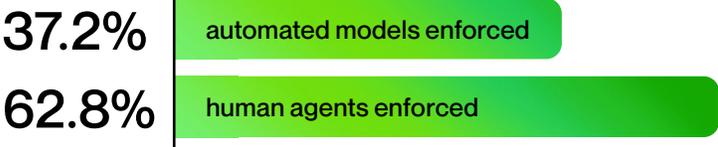
Account signals

We run checks at key moments, including account creation, job posting, and messaging, to identify signs of malicious behavior. Because there is no one-size-fits-all approach, we rely on a mix of rule-based and machine-learning models. If bad actors attempt to mimic legitimate users, tailored verification pathways, including identity checks and human review, are used to confirm findings before accounts are permanently blocked.

Content posted in jobs posts, the Project Catalog, or feedback

All new job posts, Project Catalog submissions, and negative freelancer feedback on the Upwork Marketplace are evaluated by language models trained using our policies.

Our automated systems evaluate and either accept or reject Project Catalog submissions without human review. For job posts, about 37.2% of takedowns surfaced by our automated models in 2025 were enforced automatically, while a human agent reviewed the remaining 62.8%. When we identified a potential violation through end-of-contract feedback, we always escalated it to a human agent for review.



Messages

Automated systems monitor messages for signs of fraud and scams. When we detect attempts to direct payment off of the Upwork Marketplace or share malware, we block the message and escalate the account to a human agent for enforcement review.



Expanded protections against fraud and scams

In 2025, we introduced automated detections for violations surfaced through end-of-contract feedback, attempts to direct pre-contract communication off the Upwork Marketplace in messages and attachments, and unauthorized bot usage. We also implemented real-time alerts for changes to two-step verification, enabling users to respond quickly to unauthorized activity and reduce account-takeover risk.

To combat scams, we introduced new models and detection rules that strengthen risk identification and require high-risk clients to complete additional verification before interacting with freelancers. In 2025, these systems processed nearly 3.1 million potential-scam cases, preventing approximately 553.6K potential scam jobs from being posted and approximately 40.5M potential scam messages from being sent. As a result, we reduced exposure to scam job posts by more than 31% year over year and the number of message rooms that included a scammer fell 49% year over year.

Automation guardrails

We calibrate automated content moderation based on system confidence and potential user impact. Higher-confidence, lower-impact issues may be handled fully through automation, while we generally route lower-confidence or higher-impact cases to trained reviewers for a final decision.

For example, automated models that evaluate whether job post content violates our policies must achieve a sustained accuracy rate of at least 90% for lower-severity actions, such as educational messages or restrictions that can be easily resolved in-product by the user. For permanent account blocks, these automated systems must reach a sustained accuracy rate of at least 95%.

Automated models that evaluate whether job post content violates our policies must achieve a sustained accuracy rate of at least

90%

for lower-severity actions

Automated systems must reach a sustained accuracy rate of

95%

for permanent account blocks

3.1M

potential-scam cases were processed by our new models and detection rules

553.6K

potential scam jobs were prevented from being posted

40.5M

potential scam messages were prevented from being sent

We reduced exposure to scam job posts by more than

31% yoy

The number of message rooms that included a scammer fell

49% yoy

Enforcement actions



When we determine that a user has violated a policy, we apply proportional actions based on the nature and severity of the issue. Our goal is to provide clear outcomes for users and help prevent future issues wherever possible.

Enforcement actions include:

- **Educational or warning messages** that notify the account holder of the policy violation and may remove content or request corrective updates.
- **Temporary restrictions** that limit account activity until the user completes specific steps, such as acknowledging rules, securing the account, or completing identity verification. These restrictions are tailored to the specific risk involved but generally prevent clients and freelancers from appearing in searches, posting jobs, submitting proposals, or sending and receiving payment.
- **Permanent blocks** that remove an account from active use and are used in instances of repeat or serious violations. In many cases account holders will still be able to log in and access past records but cannot take actions within the marketplace. If we suspect malicious intent, we may apply an account lockout, preventing the freelancer or client from accessing their account entirely.

When we take enforcement action on an account, we email the user to explain what happened. We also show a message banner when they log in and display the violation in the Account Health Hub.

Clearer enforcement and faster resolution

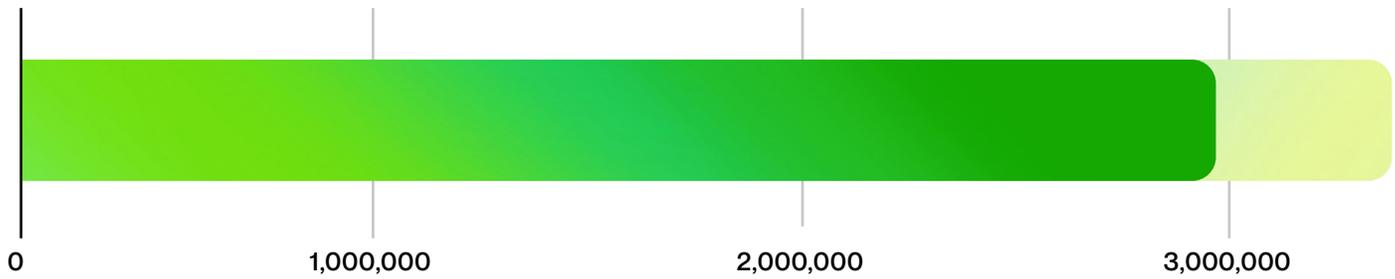
The launch of the Account Health Hub required a comprehensive redesign of our enforcement ecosystem, enabling more precise education for freelancers and clients while making the return to productive work easier.

We rebuilt internal tools so enforcement teams can more accurately distinguish between policy violations when responding to user- and system-generated reports. We also rewrote enforcement messages, placing greater emphasis on clarity and next steps.

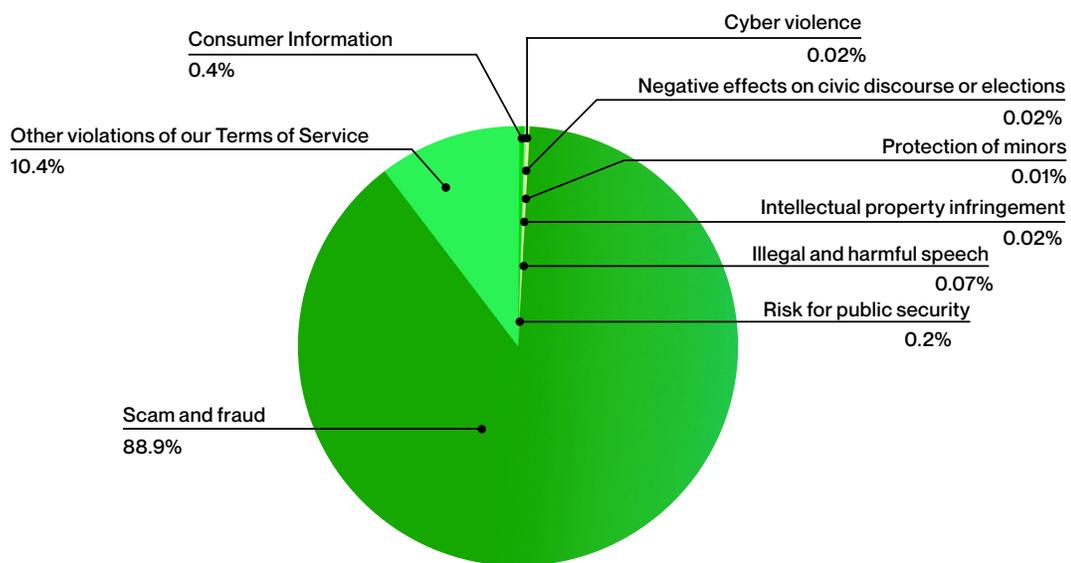
As a result, freelancers and clients can resolve temporary restrictions and appeal suspensions directly in product, without needing to contact customer support. Additional data on content moderation enforcement actions that were not directly related to user reports follows.

Total content moderation enforcement

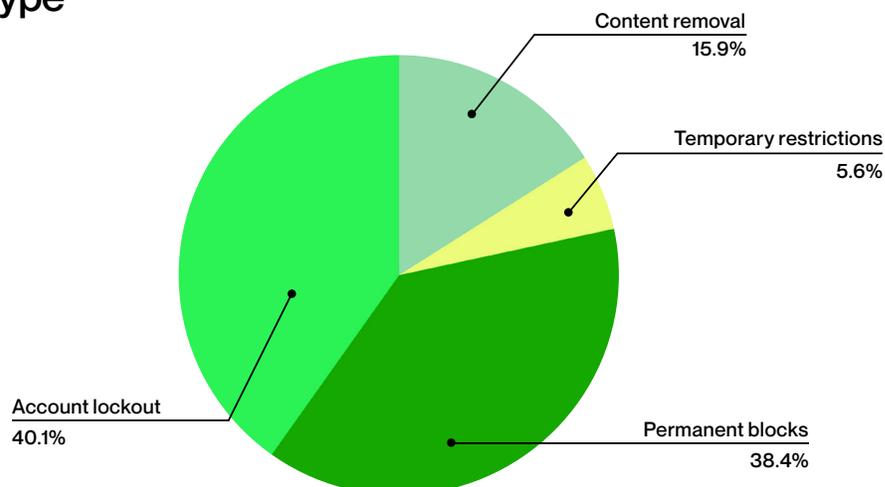
- Enforcement by a human agent
- Enforcement through automated means



Enforcement by content category



Enforcement by type



Appeals

If someone believes enforcement was applied in error, we welcome them to submit an appeal. Our Trust and Safety team manually reviews 100% of appeals to ensure we apply our standards accurately and consistently. Reviews consider account history, the nature of the issue, and any additional information provided.

We overturn enforcement decisions when we identify an error in the original review, details have been clarified which change the outcome under our policy, or we have since updated the policy.

A better guided, more streamlined appeals experience

We launched a new, in-product appeals flow within the Account Health Hub. Previously, we directed freelancers and clients to the support ticket associated with the enforcement decision to pursue an appeal.

Now, they can clearly see what limitations apply to their account, the policy violations contributing to those limitations, and the appeal options available. Users can also track appeals status in real time, keeping them informed throughout the process.

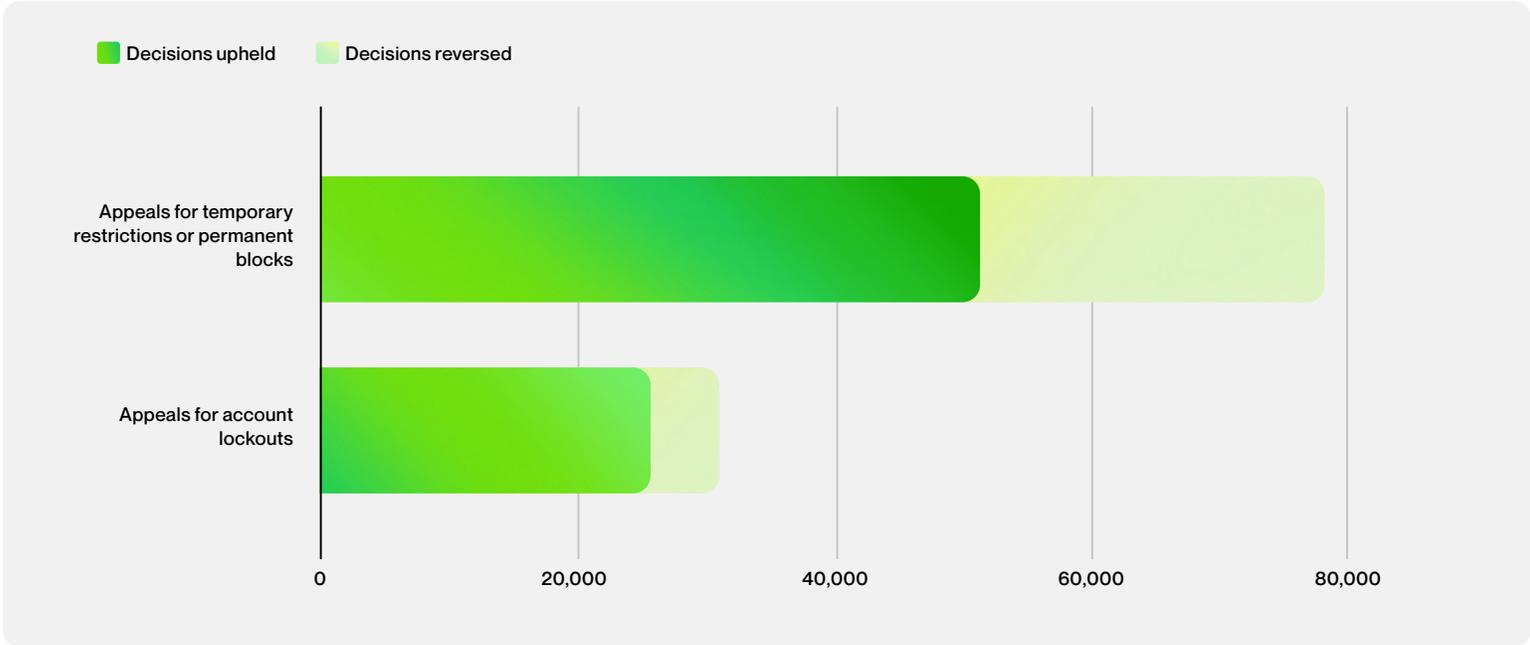
In 2025, we received 110,487 total appeals related to content moderation suspensions, of which we approved 32,980.

The median time to review these appeals was 480 minutes.

Out-of-court dispute settlements

Customers may choose to pursue resolution related to content moderation decisions through out-of-court dispute settlement mechanisms. We support fair resolution processes and comply with laws that provide for alternative dispute resolution.

The Upwork Marketplace did not receive notice of any out-of-court disputes in 2025.





Member State government requests

If we receive orders from EU Member State authorities to remove content, we assess the order's validity and scope and evaluate the content under the Upwork Marketplace Terms of Service and applicable law. Similarly, we assess whether orders from EU Member State authorities requesting user information comply with our Privacy Policy and applicable legal requirements for disclosure.

In 2025, the Upwork Marketplace did not receive any such EU Member State orders to remove content or provide user information.



Elevating trust through transparency

We have designed a system to promote fairness, safety, and transparency that can be continually developed with the fluctuating needs of the market. And guiding our system is a team committed to thoughtful judgment. They lead with clarity, act with care, and listen attentively — because that is how the Upwork Marketplace will continue to drive positive outcomes through 2026 and beyond.

As the future of work continues to unfold, we'll move forward with the same steady commitment to trust that has guided us since our inception.