



pubnub

# The ultimate guide to chat moderation

---

Stop toxic content in its tracks  
and build communities people  
actually want to join

---

# What's inside

## your moderation game plan

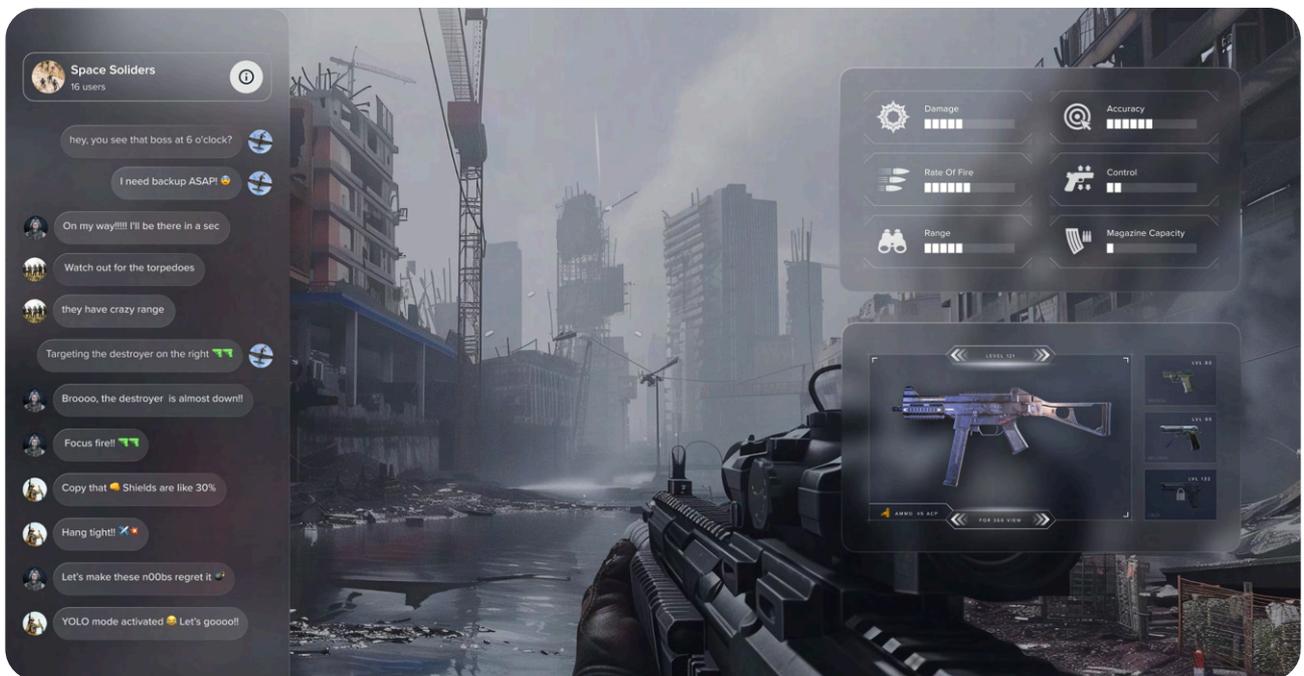
The big picture: Why chat moderation matters more than ever	2
The messy reality: What's actually happening in chat right now	4
Building a moderation system that doesn't break	8
Real-world scenarios: How to handle the messy stuff	12
Success stories: How other teams actually pulled this off	13
Crystal ball time: What's coming next (and how to prepare)	18
Wrapping up: Your action plan for building better communities	19
Meet PubNub: The team that actually gets this stuff	21
Further reading: Because you can never know too much	22

# The big picture

## Why chat moderation matters more than ever

Online chat is the heartbeat of your community and also its largest area of risk. The same features that make chat feel immediate and human can create openings for spam, harassment, and scams. The good news is that modern moderation allows you to protect people without slowing down conversations. You can deliver a better experience and reduce risk at the same time.

Effective moderation blends technology and culture. The technology side has advanced quickly. Models for toxicity, spam, and topic detection are more accurate and faster than they were even a year ago. URL classification and domain reputation services flag malicious links before they spread. Flooding and rate limiting controls handle repetitive messages so they never crowd out real conversation. On the culture side, you set the tone by publishing clear norms, applying them consistently, and giving people a fair and simple way to appeal. When your community knows the rules and trusts the process, compliance goes up and friction goes down.



## THE BIG PICTURE: WHY CHAT MODERATION MATTERS MORE THAN EVER

Most teams wrestle with three recurring challenges. The first is flooding. This happens when a user posts the same message again and again or blasts multiple channels at once. You can stop most flooding with basic controls. Per user and per channel rate limits, duplication detection, and short cool-off timers remove the bulk of the noise without affecting legitimate chat. The second challenge is spam and harmful links. This problem spikes during live events and campaigns, when attention is high. You can mitigate it by expanding short links, checking domain reputation, scanning for known patterns, and using machine learning to catch emerging spam tactics. The third challenge is toxicity and targeted harassment. This area requires nuance. Score-based models, applied with thresholds and context, let you keep the worst content out while allowing heated but acceptable debate to continue. Using profanity filters and word lists is a good first place to start.

Context matters. Live events benefit from stronger automation, safe lists for trusted roles, and rapid appeals during the show. Gaming communities need topic controls, multilingual coverage, and toxicity scoring that separates banter from abuse. Workplace and internal chats call for clear, concrete policies, configurable topic blocks for sensitive terms, and transparent action when something crosses the line. You get the best results when you tailor thresholds and actions to each scenario.

The bottom line is simple: effective chat moderation requires a combination of layered technology and intentional culture. Publish clear, accessible norms that everyone can understand at a glance. Use score-based routing and human review for edge cases that require judgment. Ensure your solution covers text, links, and audio, and instrument the moderation flow so you can measure and continuously improve. When done well, moderation protects your users, strengthens brand trust, and increases healthy engagement—transforming moderation from a cost center into a growth driver. Solutions like [PubNub Chat Moderation](#) make it possible to implement these best practices at scale, providing real-time, customizable moderation that adapts to your community's needs.

---

# The messy reality: What's actually happening in chat right now

## Why you can't ignore this anymore (even if you want to)

Chat is now the default interface for communities across industries. Live streams, games, support, collaboration, and social feeds all depend on fast conversation. As usage grows, risk does not grow at the same pace. It accelerates. A small increase in attention can lead to a surge in spam or a coordinated raid. Brands and communities also face higher expectations from users, regulators, and advertisers. People expect platforms to act quickly and fairly. Regulators expect clear policy and auditability. Advertisers want safe environments where their budgets will not fund abuse.

This combination makes real-time moderation essential. You need systems that can handle global audiences, multiple languages, and many content types. You also need to keep latency low so you do not interrupt the flow. The technical challenge is only half of the work. The social challenge is just as important. Healthy communities do not happen by accident. They grow when you set norms, reinforce them, and show people that safety and respect are part of your product's promise.

## What gaming communities taught us about building safe spaces



### The dual nature of gaming: culture and technology intertwined

Gaming communities have spent decades at the forefront of online interaction, offering a unique laboratory for understanding how culture and technology shape each other. Games provide a safe environment for people to experiment with identities, learn teamwork, and develop social skills. Features like avatars and screen names make it easier for individuals to find belonging and express themselves. However, these same features can also create opportunities for bad actors to hide behind anonymity. Structures such as guilds and clans foster leadership and collaboration, but they also demonstrate how quickly a group can organize—sometimes for positive collaboration, other times for coordinated harm—when there are no clear guardrails in place.



## The power and risk of collaboration

Collaboration is at the heart of gaming. Teams plan, coordinate, and strive for shared victories, forging bonds that often extend beyond the game into real-life friendships and even business partnerships. This dynamic illustrates the immense value that healthy chat and community spaces can create: they transform a product into a place where people feel seen, supported, and connected. Yet, the same collaborative energy can turn negative if toxic behavior is left unchecked. When harassment or exclusion becomes the norm, it drives away those who don't fit the dominant group, shrinking the community and stifling diversity. The tone of any space is not set in stone—it is the result of intentional design choices, clear policies, and daily reinforcement.



## The rapid evolution of language and community norms

Language within gaming communities evolves at lightning speed. Slang, memes, and shorthand can spread from a small group to the entire internet in a matter of days. A phrase that starts as an inside joke on one server can become a global trend almost overnight. Moderation systems must be agile enough to keep up with this rapid change, requiring adaptable models and rules that can recognize new terms and understand local context. The design of the game itself also plays a crucial role: cooperative systems tend to foster more positive interactions than purely competitive ones. The way goals, rewards, and visibility are structured directly influences how players communicate and behave.



## Moderation as cultivation, not just control

The most important lesson from gaming is that moderation is far more than just removing bad content. It is about actively cultivating positive norms, ensuring visibility for good behavior, and maintaining accountability. The goal is to help positive behaviors scale faster than negative ones, creating a virtuous cycle that strengthens the community. Achieving this requires a blend of robust technical controls—such as automated filters and adaptive models—and strong cultural leadership, including clear communication of values and consistent enforcement. When both elements work together, gaming communities become safer, more inclusive, and more resilient spaces for everyone.

# Understanding and responding to the three core risks in chat moderation

Risks fall into three broad categories:



## Content risks

include profanity and toxicity, hate and harassment, sexual content, self-harm, disallowed topics, and disinformation.



## Integrity risks

include spam and flooding, phishing and malicious links, and impersonation.



## Safety risks

include doxxing and exposure of personal information, threats, and coordinated raids.

Each group needs clear definitions and examples so everyone understands what is in bounds and what is not. Your policy is the backbone of your program. Publish rules with examples and graded outcomes. People should be able to predict what will happen if they cross a line. Education should be your first tool for borderline cases. Use prompts, reminders, and temporary friction to help people course-correct. Reserve strong enforcement for severe or repeated violations. Inclusivity should be explicit. If your product serves a global audience, then safety features should be on by default and tuned for each locale.

Score-based systems help keep context. A model can score a message from zero to ten. You can then map that score to actions. Very high scores trigger holds or removals. Mid-range scores trigger soft actions like prompts. Low scores pass without delay. This approach reduces over-blocking and gives you a natural way to learn and adjust.

## How to make your community feel like they own the rules (because they should)

People follow rules when they understand them and when they feel the process is fair. Publish your policy in plain language with real examples. Explain why certain behaviors harm the community. Share a short transparency update after major events that describes how you applied the rules. When appropriate, show aggregate metrics like the number of messages reviewed and the share that required action. Invite community feedback on the policy and review it periodically. This co-ownership builds trust and reduces conflict.



### Why your product design can make or break your moderation

Moderation is stronger when the product helps. The first message someone sends can trigger a friendly reminder about community norms. “Pause and reflect” prompts reduce impulsive posts. Slow-mode during spikes keeps conversations readable. Clear roles for moderators and hosts make it obvious who can help in the moment. These are small design choices that lower the burden on your review team and create a better experience for everyone.



### Keeping your AI honest (and your community safe)

AI is a tool, not a policy, and should be used with clear guardrails and human oversight. To ensure ethical and effective moderation, evaluate models for bias and fairness across locales and demographics, track and adjust for false positives and negatives, and use allow-lists where needed. Be transparent by explaining what is moderated and why, providing visible cues when content is blurred, held, or edited, and linking to policy. Keep humans in the loop for edge cases, giving reviewers context and the ability to override automated actions. Offer a simple appeal process, monitor reversal rates, and feed outcomes back into models and policy. Protect data privacy by minimizing retention, masking PII, restricting access by role, and documenting data flows for compliance. Ultimately, ethical use builds trust, increases compliance, and helps communities grow.

# Building a moderation system that doesn't break

## The three-layer approach that actually works (we promise)

Modern moderation systems succeed by layering three core approaches: fast deterministic rules, AI-powered analysis, and human judgment. Each layer addresses a different slice of the problem.

**Deterministic controls**—like word lists, context windows, allow/block lists, and rate limits—catch the obvious issues (spam, flooding, malicious links) with near-zero latency, ensuring the majority of harmful content is filtered out before it ever reaches users. This keeps the pipeline clean and efficient for the next layers.

**AI models** form the second layer, scoring messages for toxicity, harassment, and spam using machine learning. These models spot subtle patterns, adapt to new threats, and can be tuned for different languages and cultural contexts. High-risk content is held or removed, borderline cases get soft actions (like prompts or blurs), and low-risk content passes through. This approach lets you scale moderation to millions of messages in real time, catching evolving threats that static rules miss.

The final layer is **human review**. Reviewers handle edge cases, appeals, and nuanced situations that require empathy or context. They're supported with full message history, suggested actions, and escalation paths. Regular calibration and feedback loops ensure consistency and help the system learn over time—reviewer decisions feed back into both rules and AI models, improving accuracy and fairness.

## Smart choices: speed, customization, and policy

Effective moderation is all about balance. Latency matters—aim for sub-200ms processing for most actions, and accept slightly higher delays only for the riskiest content or complex media like audio. Customization is key: live events, gaming lobbies, and internal chats each need their own topic lists, thresholds, and escalation paths. Give product and trust teams the ability to adjust settings safely, with clear defaults and versioned policies for auditability.

Your policy engine should be declarative and modular. Use simple building blocks: conditions (score bands, keywords, link types), modifiers (user role, account age, locale), and actions (allow, annotate, blur, hold, redact, warn, mute, ban). Temporary policies for special events and version tracking make incident reviews and compliance easier.

---

## What to watch: signals, metrics, and continuous improvement

Great moderation systems rely on a broad set of signals to detect and respond to harmful or unwanted behavior. Here's a breakdown of the most important signals and how to use them:

### Profanity

Detects the use of offensive or inappropriate language. This can be as simple as matching against a word list, but more advanced systems account for obfuscation, misspellings, and context (e.g., reclaimed slurs or in-group banter).

### Slurs

Identifies hate speech and targeted derogatory language. This requires up-to-date lists and context-aware models, as slurs can evolve rapidly and may be used in coded ways.

### Harassment Patterns

Looks for repeated targeting of individuals or groups, including bullying, threats, or dogpiling. Pattern recognition can include frequency, message similarity, and recipient analysis.

### Link Types

Flags suspicious or malicious URLs, such as phishing attempts, malware, or spam. Systems should expand shortened links, check domain reputation, and maintain allow/block lists for known domains.

## **Repetition**

Detects spammy behavior, such as repeated messages, copy-paste attacks, or flooding. This can be measured per user or across the channel.

## **Account Reputation**

Considers the sender's history—new accounts, prior violations, or low trust scores may warrant stricter scrutiny or rate limits.

## **Social Graph Context**

Analyzes relationships and group behavior to spot coordinated attacks, brigading, or astroturfing. This can include sudden influxes of new accounts or synchronized messaging patterns.

## Mapping actions to score bands

For each message, combine these signals into a risk score (e.g., 0–10). Define clear action bands

### **Low Risk (0–3)**

Allow content to pass through with no delay.

### **Medium Risk (4–6)**

Annotate the message for moderator review, blur content for users, or prompt the sender to reconsider before posting.

### **High Risk (7–10)**

Hold the message for review, redact it from public view, or automatically remove it if it matches severe policy violations.

### **Escalation Steps**

For repeated or egregious offenses, apply slow-mode (limiting message frequency), issue warnings, temporarily mute or ban users, and quarantine domains that are repeatedly flagged for spam or abuse.

## Key measurements explained

### **Harmful message exposure (HME)**

Tracks the percentage or number of users exposed to harmful content before it is removed or hidden. Lower HME means your system is catching issues faster and protecting more users.

### **Time to action (TTA)**

Measures the average time between a harmful message being sent and the system or a moderator taking action. Faster TTA reduces harm and builds trust.

### **False positive rate**

The proportion of benign messages incorrectly flagged as harmful. Keeping this low is crucial to avoid frustrating users and overburdening moderators.

### **Reviewer throughput**

The number of cases a human reviewer can process per hour. High throughput indicates efficient tooling and clear policies.

### **Agreement rate**

The percentage of cases where multiple reviewers or the system and reviewers agree on the outcome. High agreement suggests consistent policy application.

### **Appeals and reversal rates**

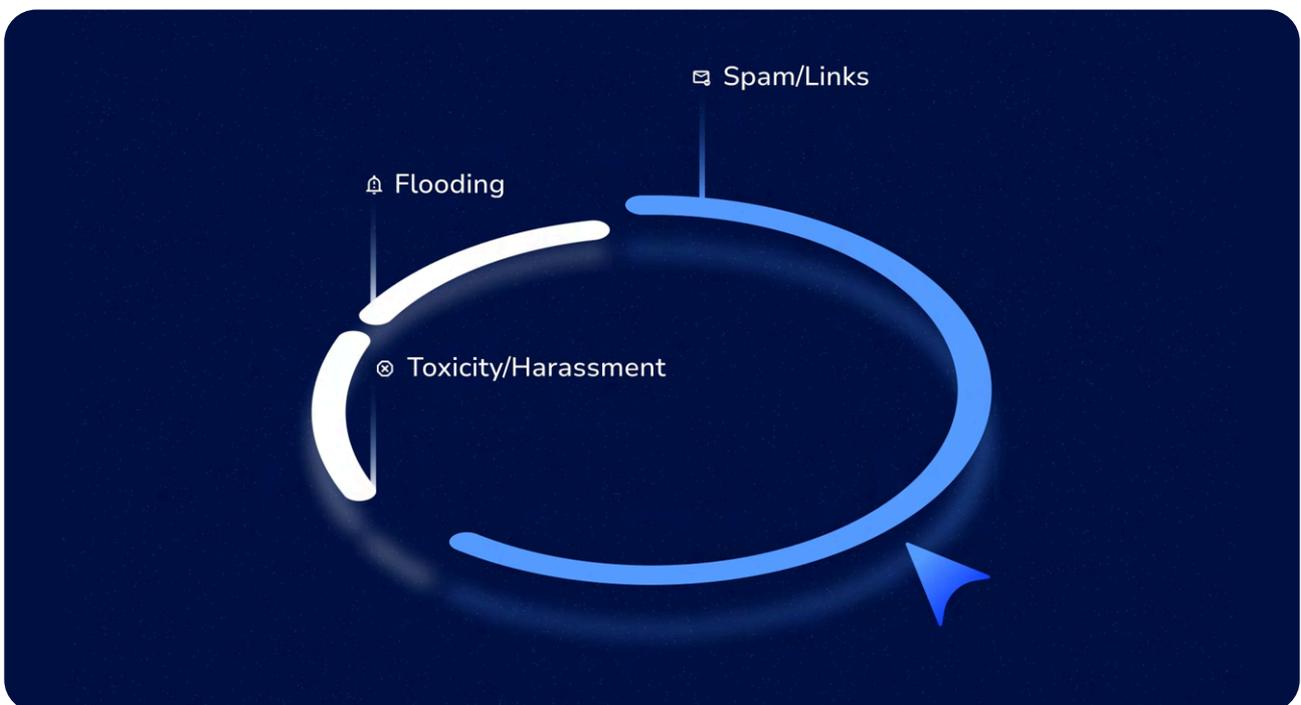
Track how often users appeal moderation decisions and how often those decisions are overturned. Spikes may indicate unclear policies or model drift.

## Dashboards and reporting

Break down these metrics by locale, channel, and event to identify hotspots, tune thresholds, and demonstrate the impact of your moderation program to stakeholders. Regularly review these insights to refine policies, retrain models, and ensure your system adapts to new threats and community needs.

## Keeping data and people safe

Moderation involves sensitive data. Limit access by role, log all actions, and redact or hash personal information wherever possible. Retain only what's necessary, and be transparent with users about data use. Build exportable audit trails for compliance and regulatory needs. For child-focused products, follow age-appropriate design codes and secure the right consents. Protecting privacy and building trust are as important as blocking harm.



## Real-world scenarios

### How teams handle the messy stuff (and what actually works)

Modern moderation is put to the test in three of the toughest environments: live events, gaming, and ecommerce. Below, we combine practical playbooks with real-world outcomes from teams who've faced these challenges at scale

# Live events

## Surviving spikes, raids, and sensitive moments

Live events compress massive attention into a short window—great for engagement, risky for safety. Teams report that the best results come from a layered, proactive approach:

### Before the event

- Publish simple, visible policies with clear examples.
- Seed allow/block lists for expected domains (partners, sponsors, known scams).
- Enable and stress-test flooding controls.
- Set up VIP/guest rules for hosts and performers, granting flexibility without sacrificing safety.
- Run pre-event drills and reminders to ensure everyone's ready.

### During the event

- Inspect links and check domain reputation in real time; expand short links to reveal true destinations.
- Use adaptive rate limits and score-based routing (e.g., messages scoring above 7 are held for review, low scores pass instantly).
- Enable slow-mode when chat moves too fast.
- Give producers a real-time dashboard with one-click controls (slow-mode, mute, Q&A mode).
- Apply topic filters and sentiment detection during sensitive segments.

### After the event

- Review logs to refine lists and thresholds.
- Identify repeat offenders and quarantine problematic domains.
- Share transparency notes on enforcement and improvements.
- Analyze sentiment and topic filter effectiveness.

### Case in point

A major platform facing a million angry fans used this approach: deterministic controls (rate limits, link checks), a policy engine mapping model scores to actions, and a producer console for real-time interventions. The result? Harmful message exposure dropped, time to action improved, and chat stayed open and readable—even during spikes. Transparency posts and simple controls for producers reduced escalations and built trust.

### What to measure

- › Harmful message exposure during broadcasts
- › Link spam blocked pre-display
- › Time to action for high-severity content
- › Sentiment trends during sensitive segments
- › Post-event appeal rates

# Gaming

## Balancing banter, harassment and global scale

Gaming chat is fast, emotional, and often multilingual. The goal isn't to sterilize the culture, but to keep players safe while preserving the energy that makes games fun.

### What works

- Use a 0–10 toxicity score, with thresholds tuned per title and mode (8–10: hold/ban, 5–7: prompts/soft blocks, lower: allow).
- Topic filters for event-specific sensitivities (e.g., team slogans, artist names, slurs in specific languages).
- Progressive enforcement: first-offense warnings, short mutes, then stronger actions for repeat issues.
- Reputation signals (account age, prior reports) to modify actions.
- Multilingual coverage: locale-specific thresholds, slang packs, and community manager input for new terms.
- For voice, transcribe audio to text and run through the same pipeline; flagged segments prompt post-match review.

### Case in point

A gaming network spanning 12 languages combined per-mode toxicity scoring, progressive enforcement, and mentorship prompts. Audio→text moderation enabled coverage of voice chat. The result: targeted harassment dropped, newcomer retention improved, and false positives declined in non-English locales. Players felt enforcement was fair, which encouraged peer coaching and reduced repeat incidents.

### What to measure

- Incident rates by mode and locale
- Queue depth and time to action at peak
- False positives in the 5–7 range
- Newcomer retention
- Post-match report/appeal rates

# eCommerce

## Fighting spam, scams, and abuse—without killing conversation

Ecommerce chat and UGC drive engagement and sales, but also attract spam, scams, and abuse. The challenge is to block the bad stuff while keeping the good.

### What works

- Clear, visible community guidelines defining spam and consequences.
- Automated filters for suspicious URLs, repeated messages, excessive promo codes, and off-topic content.
- Domain reputation checks and short-link expansion before allowing links in chat or reviews.
- Real-time monitoring during live shopping events to catch spikes and coordinated spam.
- Layered detection: combine keyword/pattern matching with behavioral signals (message frequency, account age, prior reports).
- Role-based rules: stricter scrutiny for new sellers/buyers, fast-tracking for trusted accounts.
- Soft interventions for borderline cases (prompts, temporary holds).
- Comprehensive audit logs and privacy controls for appeals and compliance.
- Moderator dashboards highlighting spam trends by product, campaign, or seller, with tools to quarantine domains or mute repeat offenders.
- Playbooks for common scenarios (flash sales, influencer streams, review surges).

### Case in point

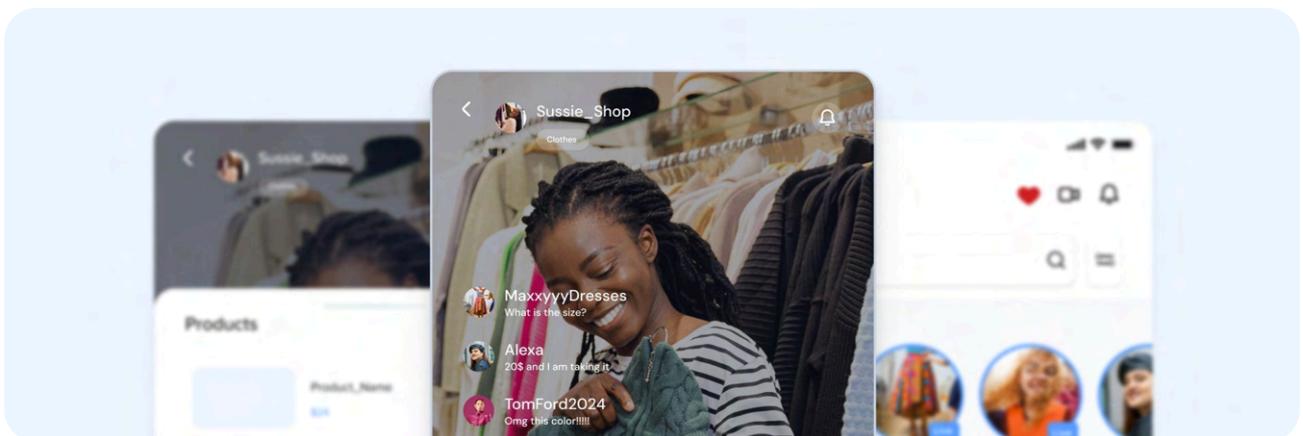
A fast-growing ecommerce platform expanded into live shopping and creator-led drops. They layered deterministic rules (link expansion, domain reputation, affiliate quarantine), flooding controls, and a score-to-action map for toxicity and spam. Role-aware policies reduced false positives for trusted sellers and creators. The result: harmful link exposure dropped, conversion rates improved, and hosts reported fewer personal attacks. Transparency posts after launches built trust with buyers and brands.

## What to measure

- › Volume/type of spam attempts per 1,000 messages
- › Percentage blocked before display
- › Time to action for high-severity spam
- › Appeals and reversal rates
- › User/seller trust and perceived effectiveness
- › Repeat offender rates
- › Impact of interventions (prompts, holds) on spam reduction
- › Conversion-adjacent metrics (add-to-cart, checkout during spikes)

## Key takeaways across all scenarios

- › Layer deterministic controls, AI scoring, and human review.
- › Tune thresholds and actions to the risk profile and context.
- › Give frontline teams simple, reliable tools and dashboards.
- › Publish clear rules and transparency notes to build trust.
- › Measure what matters—exposure, latency, false positives, and user sentiment—and use the data to refine your approach



# Crystal ball time

## What's coming next (and how to prepare)

Moderation is moving earlier in the flow. Instead of reacting to incidents after they happen, teams are building systems that predict and prevent them. New accounts that join and post at high speed are a simple example. Combine that signal with mention patterns and known raid methods, and you can add gentle friction that stops a problem before users see it. Prevention is not about heavy-handed blocks. It is about small, smart defaults that keep conversations healthy.

Multimodal AI is improving context. Text, images, audio, and metadata are feeding a single score rather than separate pipelines. This helps with sarcasm, reclaimed slurs, and complex phrases where literal keywords do not tell the full story. Audio moderation is following a practical path. Most teams transcribe first, then apply text models. As real-time speech models improve, expect better accuracy and lower latency for spoken chat.

Reputation and transparency will matter more. People want to know what happens to their reports. They also want a path to repair. Expect more products to include in-product education and restorative actions. A prompt that explains why language is harmful, followed by a short cool-off, can change behavior. For repeat problems, stronger actions should be predictable and clearly explained.

Safety by design is becoming the default. Features like slow-mode, first-message friction, and clear community roles now ship out of the box. Enforcement is becoming more graduated and consistent. This reduces fear and makes participation feel safer. When people understand the boundaries, they tend to follow them.

Regulation is also shaping the space. Auditability, fair appeals, and age-appropriate design are now table stakes in many regions. Plan for these requirements early so you are not surprised later. Build logs and controls that help you meet obligations without exposing more data than you need.

The common theme across these trends is simple. Move left in your process to prevent harm. Combine product design, AI, and governance so culture scales with your audience. You can grow fast and stay safe.



# Conclusion and next steps

Modern moderation is a layered system backed by intentional culture. Blend filters, machine learning, and human review. Adopt score-based routing. Publish norms and measure what matters. When you do this, you protect people and create room for better conversation and growth.

## Expanding chat across your product

As moderation technology improves, chat becomes viable in more places across your experience. Safe, readable conversations unlock new value:

- Product discovery and live shopping: real-time Q&A, creator drops, and guided buying without spam or scams.
- Support and success: in-app help rooms, onboarding cohorts, and community troubleshooting that stay respectful and on topic.
- Events and co-watching: audience chat, polling, and Q&A that remain inclusive even at peak scale.
- Workplace and education: study groups and team rooms that encourage debate while preventing abuse.

When harmful content is filtered in sub-second time and reviewers focus on edge cases, chat can move from a risky add-on to a core feature you deploy broadly.

## Actionable recommendations

Start with policy. Write simple rules with specific examples, and map actions to score bands so outcomes are predictable. Instrument your messaging system so you can measure incidents, false positives, and time to action. Turn on deterministic controls for lists, links, and flooding, then pilot your first models in shadow mode. Watch the numbers and the feedback. Adjust thresholds based on real data.

Stand up reviewer tools and queues. Give reviewers context and clear guidance so they can make consistent decisions. Calibrate the team with regular reviews of edge cases. Add multilingual and audio coverage as your audience needs them. Establish feedback loops from reviewers and users to your models and rules. Review metrics monthly. Tighten where harm leaks through. Loosen where over-blocking hurts the experience.

### Implementation roadmap

In the first month, finish policy, instrumentation, and baselines. Seed allow and block lists. Add link intelligence. In month two and three, run a machine learning pilot, launch your reviewer workflow, and map scores to actions. In month four through six, expand to new locales and add audio coverage. Educate your community and publish short transparency notes so people know safety is part of your promise.

### Additional resources

Helpful resources include policy templates, reviewer guidelines, locale test packs and slang dictionaries, and producer run-books for live events. Use these to accelerate rollout and keep teams aligned.

Ready to implement these strategies? **Contact our experts for a consultation.** Download the implementation checklist to get started today. See how customers used score-based routing to reduce incidents and increase trust. Join thousands of developers already using these best practices.

# About PubNub

PubNub is the real-time platform trusted by leading brands to power safe, engaging chat and interactive experiences at global scale. Our AI Moderation solution helps you protect your users and brand by detecting and preventing harmful content—including toxicity, harassment, spam, and more—before it reaches your community.

## With PubNub, you can

- ✓ Instantly process and moderate messages in real time using AI and rule-based filters
- ✓ Route flagged content to human reviewers or automated workflows for further action
- ✓ Maintain a safe, positive environment with customizable policies and score-based actions
- ✓ Scale moderation with low latency and global reliability
- ✓ Meet regulatory requirements with robust audit trails, message storage, and transparency features

## Relevant solutions

- ✓ Real-time messaging and presence for chat, live events, and collaboration
- ✓ Edge Functions for inline AI moderation (auto-moderation) and custom logic
- ✓ PubNub Illuminate to catch / prevent chat flooding and cross-posting
- ✓ Stream controllers and channel groups for dynamic routing and segmentation
- ✓ Storage and message actions for compliance, auditability, and appeals

Customer highlights Leading gaming, social, and event platforms use PubNub to deliver safer, more engaging experiences—reducing incidents, increasing user trust, and supporting global growth.

Request a demo or connect with our experts about building safer, more inclusive communities: [AI Moderation by PubNub](#)

# Further reading

- [!\[\]\(da0e4193af27a672ce1097742a0c5ae3\_img.jpg\) Game On: The Explosion of Gaming into Mainstream Culture — LBBOnline](#)
  - [!\[\]\(a37f76d95d86d074ec88e2a0e6ed24c0\_img.jpg\) Internet: Social gaming and social networking — Encyclopædia Britannica](#)
  - [!\[\]\(0c8eaaba83699fa00f1baa4191f85455\_img.jpg\) Sherry Turkle — Life on the Screen](#)
  - [!\[\]\(8b86ea509bf54a24060db64a82ed82ce\_img.jpg\) Automate Chat Moderation in Real-Time with PubNub Illuminate](#)
  - [!\[\]\(adff6523c7990decd507859dcc13f5b4\_img.jpg\) Gaming for the Win — University of Houston \(MMO leadership and skills transfer\)](#)
-