



Together, we fight online child sexual abuse.

Annual transparency report

2025

technologycoalition.org



From our CEO

A year of progress

Over the past year, online child sexual exploitation and abuse has evolved at unprecedented speed. More than ever, harm is being shaped by how services connect and interact. New capabilities—especially generative AI—are accelerating the speed, scale, and complexity of abuse.

At the same time, I am encouraged by real, measurable progress. **In 2025, our members advanced 317 child safety mileposts** – practical improvements in how companies protect children.

On average, that's more than five concrete advancements per company, more than double the rate of progress in 2024. Put simply: **the challenge is growing—but so is industry's response.**

This progress didn't happen by chance. It reflects a shift toward more structured, coordinated approaches. Companies are improving how they identify abuse—sharing threat signals and acting on them through initiatives like Lantern.

There is a greater focus on preventing abuse before it happens, building safeguards earlier and investing upstream.

But no single company can tackle this alone. Collaboration is essential. Through the Tech Coalition, companies share tools, intelligence, and hard-earned expertise. **We provide industry a space where people can be open, practical, and focused on solving real problems.**

That work also extends beyond our membership. Through Pathways, we are making tools and resources available more broadly, because this challenge doesn't stop at 60 companies.

To raise the bar in child safety, we have to raise it across the entire tech ecosystem.

What stands out most to me is the commitment that drives this work. I see firsthand the expertise, care, and persistence our members bring to protecting children online.

This is hard work—the harms are real, and **the bad actors are determined. But so are the people working to stop them.**

I'm grateful to our members and to the Tech Coalition team for everything they've done, and continue to do. Together, we've made meaningful progress. And we need to keep going.



Sean Litton
President & CEO
Tech Coalition

“

At its core, the Tech Coalition is about people working together to protect children online.

”

Contents

1 Driving progress in 2025	4	2 Staying ahead of abuse	11	3 Preventing harm in 2025	14
Our members	5	Sector-wide challenges	12	Safety by design	15
Led by industry	7	Collective industry action	13	Age assurance	16
Collaboration engine	8			Pilot technology projects	17
Pathways	9			Deterrence measures	18
Pathways in action	10			Research-informed prevention	19
4 Detecting OCSEA in 2025	20	5 Industry's response in 2025	29	6 The future for the Tech Coalition	40
Image hashing adoption	21	Reporting to authorities	30		
Video hashing adoption	22	Supplemental reporting	31		
Video hash interoperability project	23	Reporting AI-generated abuse	32		
Hash & keyword exchanges	24	Law enforcement feedback loop	33		
Detecting OCSEA with classifiers	25	Investigations & enforcement	34		
Korean grooming classifier	26	Post-investigation enforcement tools	35		
Behavioral & repeat offender detection	27	Lantern	36		
User reporting	28	Lantern uncovers exploitation network	37		
		Transparency & accountability	38		
		Member transparency reports	39		

1

Driving progress

The Tech Coalition’s voluntary mileposts help measure, track & drive industry’s impact on child safety.

The most effective safeguards prevent abuse before it happens—making their impact inherently difficult to measure.

The Tech Coalition helps make progress visible by measuring the tangible steps members are taking to reduce risk, close safety gaps, and strengthen protections.

This work is supported by our **voluntary milepost framework** that provides a structured methodology to assess and benchmark progress.

Mileposts are categorized and presented across three areas:

- **Prevention** (see page 14)
- **Detection** (see page 20)
- **Response** (see page 29)

In 2025, member companies advanced more than **double the mileposts of 2024**, completing **317 new mileposts**, compared to 122 last year.

This represents an average of more than **five mileposts per member**, also more than double the rate in 2024.

Overall, **54 of 57* members advanced at least one milepost** across **prevention, detection, and response**.

Taken together, these results indicate continued maturation of industry capabilities, with progress relatively evenly distributed across prevention, detection, and response, as companies strengthen safeguards and integrate safety earlier in product design.

317
mileposts
in 2025

102

Response mileposts

117

Detection mileposts

122
mileposts
in 2024

98

Prevention mileposts

* 57 of 60 member companies completed our annual survey in 2025

60 companies working together to strengthen how the industry prevents, detects, and responds to online child sexual exploitation and abuse.

Foundational members provide the highest level of support, helping to drive leadership, shape priorities, and advance collective action across the industry.

Alongside them, a growing and diverse membership contributes practical experience, insights, and innovation from across sectors and regions.

Through this model, companies share what works in practice, enabling others to adopt proven approaches and strengthen safety systems more quickly.

This collaboration raises the overall standard of child safety, extending progress beyond individual companies to the wider tech ecosystem.

Our members

Foundational



Mission



Cornerstone



Our members

In 2025, thirteen companies joined, bringing the total membership to 60.



















Our new members are:

- Block
- Linktree
- Medal
- NVIDIA
- Padlet
- Reddit
- SoundCloud
- Stability AI
- Substack
- Vimeo
- VRChat
- WeTransfer
- Wizz

Bridge

	ANTHROPIC					Electronic Arts	
						OUTSCHOOL	PATREON
			Snap Inc.				
							

Associate

					Linktree*		
					stability.ai		
			WeTransfer		ZEPETO		

Led by industry

The Tech Coalition is governed by a Board of senior leaders from member companies who review progress against strategic priorities and member mileposts.



Liz Thomas
Board Chair

Senior Director of Public Policy, Digital Safety
Microsoft



Kristine Dorrain
Board Treasurer

Senior Corporate Counsel for Content Policy
Amazon



Viraj Doshi
Board Secretary

Platform Safety Lead
Snap



Ethan Arenson

Managing Associate General Counsel and Head of Digital Safety, **Verizon**



Tami Bhaumik

Vice President of Civility and Partnerships
Roblox



Chelsea Carlson

Child Safety Investigations Lead
OpenAI



Emily Cashman Kirstein

Child Safety Public Policy Lead
Google



Margaux Liquid

Head of Trust and Safety
Yubo



Lili Nguyen

USDS Head, Risk & Response Operations, T&S
TikTok



Ravi Sinha

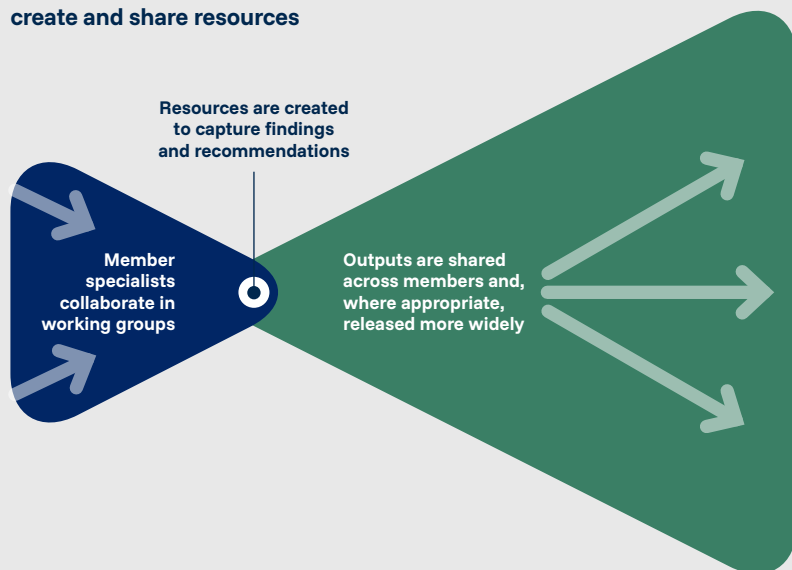
Director, Head of Child Safety Policy
Meta

“As a founding member of the Tech Coalition, Google is proud to support efforts that strengthen the industry’s collective response to OCSEA. By bringing companies together to share research, develop tools, and advance best practices, the TC helps build safer online experiences for children across the internet.”

Collaboration: our engine of progress

Information and knowledge-sharing groups act as a central mechanism for companies to coordinate cross-industry action on child safety.

Collaboration process to create and share resources




In 2025, members developed risk frameworks, reporting templates and investigative guidance.

These outputs support the adoption of proven practices across platforms, reduce duplication of effort, and accelerate the spread of effective approaches across the industry.

Without structured collaboration, companies would have to address child safety challenges alone.

By working collectively, members accelerate learning, share operational experience, and raise the baseline of child safety practices across the global technology ecosystem.

2025 member collaboration stats





Pathways

Pathways is our non-member program to provide free access to tools, resources, and guidance to strengthen child safety across the wider tech ecosystem.

“Participating in the Pathways program affirmed the safety by design foundations we were already putting into practice, while offering valuable external perspective and best-practice alignment. It acted as a constructive checkpoint that helped strengthen our integrity and safety practices around child safety, sharpening our readiness to publish our first transparency report with clarity and confidence.”



Patricia Wood,
Head of Integrity Operations &
Harm Prevention, **Stability AI**

Pathways supports organizations at different stages of maturity, helping them adopt foundational safeguards and strengthen their approach to preventing and addressing online child sexual exploitation and abuse.

By **reducing barriers** to access and sharing practical solutions, Pathways helps more companies implement effective protections and contributes to more consistent safeguards across platforms.

In 2025, six Pathways participants: **Block, Medal, Padlet, Stability AI, Substack** and **Wizz** joined the Tech Coalition as full members.

Pathways in 2025

80

Companies engaged in Pathways initiatives

41

New companies joined

6

Companies obtained Microsoft PhotoDNA licenses.

7

Companies received direct guidance through the “Ask an expert” which leverages the experience of our members.

7

Resources published on moderation, transparency, financial sextortion, law enforcement response, generative AI, and civil society engagement.



Pathways in action

From Pathways to membership.
Strengthening child safety at Wizz App,
a social discovery app for Gen Z.

Elevate is our pilot program supporting companies to strengthen child safety systems and prepare for Tech Coalition membership.

Wizz participated in one of the first pilots.

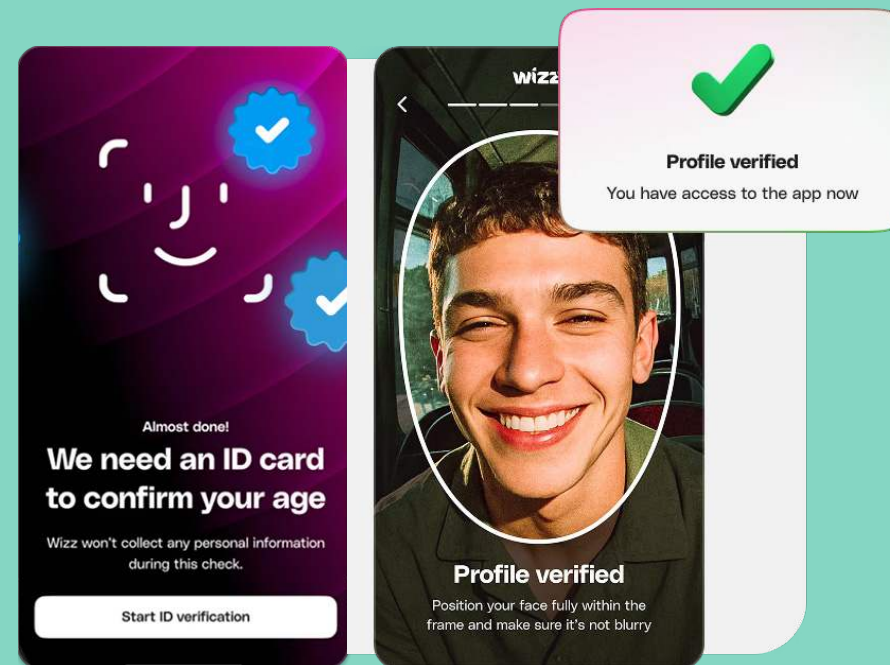
Elevate will launch in 2026.

During 2025, Wizz participated in the **Pathways Elevate pilot**, working with the TC to strengthen its child safety systems

This included implementing enhanced detection tools, establishing reporting to NCMEC, and updating safety standards.

These changes **enabled Wizz to meet Tech Coalition membership criteria and join in 2025.**

Through membership, Wizz now contributes to collaborative industry efforts to combat online child sexual exploitation and abuse, applying these strengthened capabilities within a broader cross-platform response.



2

Staying ahead of abuse

Some of the most significant and evolving child safety threats our members' trust and safety teams faced in 2025.

Online threats



CSAM at scale

The production and spread of child sexual abuse material, including known and novel content, livestreamed abuse, and self-generated or coerced sexual content involving minors.

AI-generated abuse

Synthetic content and manipulation enabled by AI, increasing the scale, speed, and complexity of abuse, including emerging forms of exploitation.

Online grooming

Abuse initiated through online enticement in private environments, including gaming and encrypted messaging, leading to coercion, trafficking, or exploitation.

Financial sextortion

Exploitation involving coercion, extortion, and financial harm targeting young people, often involving threats to share sexual content to obtain money or further compliance.

Sadistic exploitation

Child sexual exploitation involving extreme violence or coercion, often intersecting with violent or extremist online communities.

Sector-wide challenges

Keeping children safe online goes beyond individual platforms; bad actors seek out gaps and vulnerabilities.

Online threats



CSAM at scale

AI-generated abuse

Online grooming

Financial sextortion

Sadistic exploitation

Sector-wide challenges



Offender adaptation

Cross-platform fragmentation

Harder-to-detect harm

Gaps in safety maturity

Bad actors continuously change tactics, platforms, and behaviors to avoid detection and enforcement.

Abuse often spans multiple services, leaving individual platforms with only partial visibility of harmful activity.

Abuse is increasingly occurring at earlier stages and through subtle signals, before clear violations are visible.

Differences in safety capabilities across platforms can create gaps that offenders exploit.

Collective industry action

Through the Tech Coalition, companies coordinate and strengthen efforts across prevention, early detection and response, alongside cross-platform intelligence sharing to more effectively disrupt abuse.

Online threats



CSAM at scale

AI-generated abuse

Online grooming

Financial sextortion

Sadistic exploitation

Sector-wide challenges



Offender adaptation

Cross-platform fragmentation

Harder-to-detect harm

Gaps in safety maturity

Collective industry action



Prevention

Designing products and safeguards to reduce risk and stop abuse before it occurs.

Detection

Identifying harmful content, behavior, and signals to surface abuse as early as possible.

Response

Taking action to remove harmful content, disrupt bad actors, and report to relevant authorities.

3


Preventing harm in 2025

Prevention practices are evolving as companies embed safety-by-design into product development, alongside policies and solutions to deter harmful behavior.

Our milepost tracking shows this progress clearly: members advanced **98 new prevention mileposts in 2025**, averaging nearly two prevention interventions per member.

As companies introduce new products and technologies, there are growing opportunities to further strengthen deterrence and integrate safety considerations earlier.

Tech Coalition resources and guidance are helping companies assess risk more effectively and implement safeguards before harm occurs.



Safety-by-design

Members are increasingly embedding safety-by-design in product development to identify and mitigate risks earlier, before harm occurs.

“Keeping children safe in an AI-enabled world starts with how we design our systems. Through the Tech Coalition, companies are sharing learnings and turning them into practical safeguards in product design. Advances in AI are also helping the industry do more — improving detection, strengthening classifiers, and identifying new and emerging forms of abuse, including AI-generated CSAM.”



Chelsea Carlson
Child Safety Investigations Lead
OpenAI

Safety-by-design is becoming embedded across the sector.


In 2025, **49 of 57 members** reported using formal safety-by-design reviews, integrating safety into product development.

Earlier risk identification is also strengthening, with **37 members** using structured risk assessment frameworks and 33 documenting OCSEA-related risks.

More advanced techniques are emerging, including **red teaming (25 members)**, alongside safeguards such as **age-based restrictions (31)** and **parental controls or safety tools (25)**.

Together, these indicators show that prevention is becoming more systematic, with companies identifying and addressing risks earlier in the product lifecycle.

Members use of safety-by-design approaches



2025

- 49 Safety by design product review process
- 37 Framework for assessing risky features and settings
- 33 Documentation of potential OCSEA risks
- 31 Restrict account / feature by user age
- 25 Red teaming / adversarial testing
- 25 Parental controls / in-product education

Age assurance

Age assurance is being adopted by more companies as part of efforts to support age-appropriate experiences and strengthen child safety online.

“Games should be fun, creative, and safe for players of all ages. Strengthening age assurance helps platforms like Roblox deliver more age-appropriate experiences and stronger protections for younger users, and collaboration through the Tech Coalition helps advance these efforts across the gaming ecosystem.”



Tami Bhaumik
VP of Civility and Partnerships
Roblox

Age assurance practices continued to strengthen in 2025, with more companies using techniques to estimate or verify age.


Self-declaration remains most common (42 members), while higher-confidence methods are expanding, including **document verification (26) and facial analysis (20)**.

Many companies now combine multiple signals, adapting approaches based on platform design and risk.

Regulatory developments also accelerated adoption, reinforcing expectations for stronger practices.

The Tech Coalition’s age assurance knowledge-sharing group, with participants from 24 companies, supports this progress by enabling members to exchange insights, assess child safety risks in product features, and refine implementation approaches.

Members use of age assurance



Piloting technology projects

Pilot to detect OCSEA risks in livestreaming

The Tech Coalition supports the testing of emerging prevention and disruption solutions before risks scale. Pilot initiatives bring companies together to develop, assess, and refine tools under real-world operating conditions.

If a pilot demonstrates operational effectiveness, it can be positioned for broader adoption. For example, Lantern – now a core cross-platform signal sharing capability – originated as a TC pilot in 2023 before scaling.


This pilot-to-infrastructure model enables the TC to test solutions, evaluate operational impact, and scale effective safeguards and tools across platforms.

In 2025, the Tech Coalition worked with a member company to test a proof of concept for identifying **OCSEA risks in livestreaming** environments.

The pilot analyzed metadata signals, including session characteristics and anonymization behaviors, to generate a risk score indicating elevated likelihood of abuse.

The project was designed to support earlier intervention by child safety teams, enabling further review before harm escalated.


Development concluded in 2025, with testing underway at one member company. Evaluation will continue in 2026 to determine feasibility for broader industry adoption.



Deterrence measures

Platforms are strengthening deterrence through clearer policies, targeted in-product warnings, and expanded prevention messaging to discourage harmful behavior.

Members use of deterrence measures



Deterrence measures are becoming a more visible component of platform safeguards addressing OCSEA.

In 2025, **50 members reported publicly outlining policies and enforcement actions**, setting clear expectations for users about unacceptable behavior and consequences.

Companies are also implementing in-product deterrence, including **warnings for users at risk of harmful behavior (17 members)** and **messages alerting potential victims to risky interactions (10 members)**.

These approaches reflect growing recognition that prevention requires both discouraging harmful conduct and helping users recognize risk.

Companies are increasingly working with NGOs to develop and amplify prevention messaging, including victim-centered resources and public awareness campaigns, expanding reach and credibility.

Through Tech Coalition collaboration and knowledge-sharing, members exchange insights and refine deterrence approaches, strengthening messaging and integrating deterrence more consistently across platforms.

Research-informed prevention

Research insights are shaping platform interventions that deter harmful searches and direct users to support before harm escalates.

The **Tech Coalition Safe Online Research fund** supports applied research to better understand and prevent online child sexual exploitation and abuse.

With over **\$2.75 million invested across 16 projects**, the Fund focuses on areas such as grooming prevention, deterrence, AI-enabled detection, and emerging risks.

Findings are translated into practical insights that inform product design and policy, with researchers and companies applying evidence in real-world contexts. This strengthens prevention strategies and drives operational improvements across platforms before risks scale.

Research funded by the Tech Coalition is directly informing platform interventions.

A **University of Kent** project examined the behavior of individuals searching for CSAM, identifying ways to disrupt harmful activity and encourage help-seeking.

These insights informed updates to **Google's CSAM prevention systems**.

When users enter CSAM-related queries, Google now displays evidence-based warning messages alongside links to support resources such as helplines and reporting services.

These changes have led to measurable reductions in harmful search behavior, demonstrating how research can translate into effective, real-world prevention.



4

Detecting OCSEA in 2025

Hashing remains a core input within a broader detection system, alongside layered signals and AI-supported tools that help platforms identify and triage emerging abuse.

Our milepost tracking reveals significant progress: members advanced **117 new detection mileposts in 2025**, almost three times more than the previous year.

While known CSAM can often be identified through hashing technologies, detecting previously unseen abuse is more complex, particularly with AI-generated material.

Cross-platform collaboration also strengthens detection. Signal sharing initiatives such as Lantern help surface patterns earlier across services, and support coordinated disruption.



2024

2025

Detection mileposts advanced

Image hashing adoption


Industry adoption of image hashing technologies continues to expand, strengthening detection and removal of known CSAM.

Platforms are increasingly combining multiple tools to detect known CSAM at scale, using hashing technologies that match digital fingerprints of previously identified content.

In 2025, **37 members reported using PhotoDNA**, making it the most widely adopted solution. Through the Tech Coalition's **PhotoDNA sublicensing program**, companies can access and implement this technology more quickly.

Adoption of additional tools is also expanding, with **21 members using Thorn's Safer Hash** and **13 using PDQ**. Many companies deploy multiple tools, with **17 members reporting internal hashing** systems to strengthen detection across services.

Members use of image hashing technologies



“Raising the baseline for child safety across the industry is supported by shared tools and coordinated adoption. The Tech Coalition is accelerating the widespread use of hashing and detection technologies, including through sublicensing and onboarding support for Microsoft’s PhotoDNA.”



Liz Thomas
Senior Director of Public Policy,
Digital Safety
Microsoft

Video hashing adoption

Video hashing adoption is accelerating as platforms strengthen detection of known CSAM in video content.

In 2025, members reported deploying video hashing technologies, including **Thorn's Safer Hash (15 members)**, **Google's CSAI Match (12)**, and **PhotoDNA for Video (10)**.


Many companies combine multiple approaches, with some using shared tools alongside internal systems (9 members) to improve detection across formats.

Adoption has accelerated, with companies reporting no video hashing capabilities falling from 17 in 2024 to 13 in 2025.

At the same time, limited uptake of TMK+PDQF reflects industry learning.

As adoption grows, video hashing is becoming an integral safeguard, complementing established image-based detection.

Members use of video hashing technologies in 2025



Video hash interoperability project

Our Video Hash Interoperability Project (VHIP), developed in collaboration with NCMEC, Meta, Google, and Thorn, addresses the lack of standardized support for the many modern video hash types.

Hashing is a powerful tool to detect re-uploads of known abusive footage, but not every company uses the same video hash system.

If Company A uses one format and Company B uses another, databases that store hashes of known CSAM cannot automatically translate between formats.

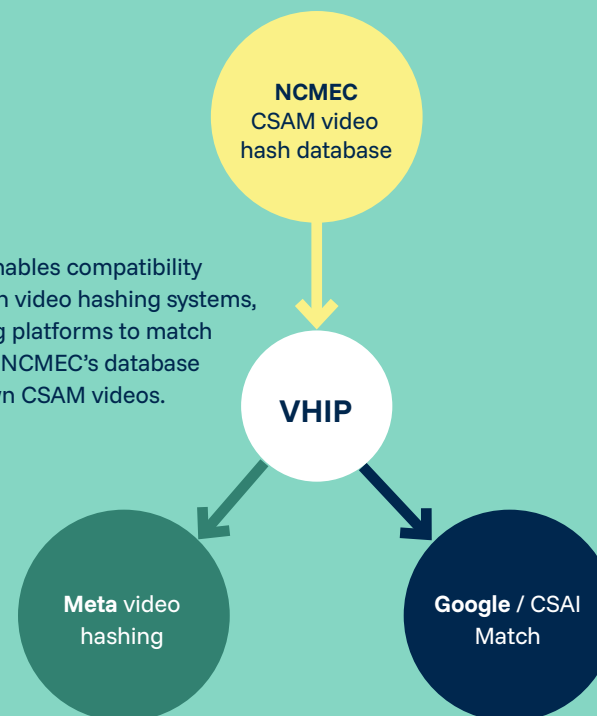
That may slow detection and risk leaving abusive content online.

VHIP helps change that. It enables different hash formats to be compatible with NCMEC's database. This system has supplemented video hashing processes across Google and Meta products, helping them – and other participating companies – match against NCMEC's full database of known CSAM videos and remove content more quickly.

In 2025 alone, VHIP hashed 435,000 videos, bringing the total to more than 784,000 since launch.

Each hash represents a video depicting the sexual abuse of a child reported to NCMEC – each one helping to protect victims and prevent further child abuse.

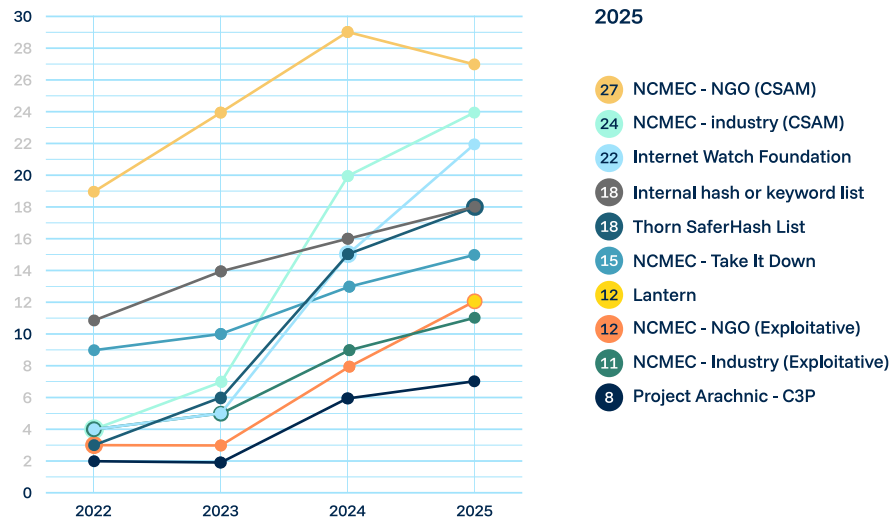
VHIP enables compatibility between video hashing systems, allowing platforms to match against NCMEC's database of known CSAM videos.



Hash & keyword exchanges

Combining shared datasets and internal signals is strengthening detection and disruption of abusive content across platforms.

Members use of hash and keyword repositories



Member companies are strengthening detection by combining multiple hash and keyword exchanges to identify and remove abusive content more effectively.

These exchanges allow platforms to match uploaded content against shared datasets from industry and NGOs, supporting faster detection of known material.

Participation continues to expand as companies integrate shared signals with internal systems, improving detection and enabling more coordinated disruption across services.

In 2025, many members matched against multiple datasets, including those from **NCMEC**, the **Internet Watch Foundation (22 members)**, **Thorn's Safer Hash list (18 members)**, and **Project Arachnid (8 members)**.

Members are also maintaining **internal hash or keyword lists (18 members)** to complement shared datasets and capture platform-specific signals.

New collaboration mechanisms are emerging, with **12 members reporting Lantern** engagement as part of their use of shared repositories.

As these exchanges expand, their effectiveness depends on maintaining high-quality, relevant, and well-coordinated datasets.

Detecting OCSEA with classifiers

Classifier-based detection is expanding rapidly, helping platforms identify previously unseen abuse across images, video, and text.


Platforms are expanding the use of machine learning and AI models to identify abusive material, complementing existing tools such as hashing and user reporting.

These systems analyze patterns in images, video, and text to detect potential abuse signals.

Over the past three years, adoption has grown significantly, with **image classifiers rising from 13 members in 2022 to 41 in 2025**, and **video classifiers from 4 to 32**.

Members are also applying classifiers to detect harmful behaviors, including **CSA text (18 members)**, **grooming (12)**, and **sexortion (8)**, reflecting efforts to identify more abuse.

Members use of classifiers



Korean grooming classifier

Our Tech Innovation Advisory Council supported a pilot addressing gaps in non-English grooming detection, using Korean as a proof of concept.

Issue: the lack of non-English grooming detection tools, which leaves minors around the world vulnerable to harm.

Response: the TC partnered with an APAC-based member company to develop a Korean- and English-language classifier capable of detecting grooming-like behaviors in text exchanges.

A Korean-language grooming classifier was built and is undergoing internal testing.

Developing the model required labeling of more than 4 million lines and 50,000 conversations with grooming attributes, translated from English to Korean using another member's proprietary in-house translation models.

Building the classifier revealed several key insights, including the challenges of capturing cultural nuances in grooming detection and the large volumes of data required to test such models effectively.

Despite these challenges, early results are promising—the model has performed well in detecting grooming behaviors in both short and long conversations.



Behavioral & repeat offender detection

Behavioral detection is strengthening platforms' ability to identify suspicious activity and intervene earlier in potential abuse.

By analyzing **behavioral signals** such as unusual account activity or patterns associated with abusive interactions, platforms can identify potential risks earlier and intervene before harm escalates.

These capabilities enable companies to move beyond relying solely on user reports or known content matches, strengthening detection across services.

In 2025, **34 members** reported identifying suspicious account or user behaviors linked to OCSEA through behavioral detection or related investigative signals.

Furthermore, **22 members** reported implementing technologies specifically designed to detect these patterns.

The growth in behavioral detection reflects a broader shift toward more proactive safeguards, complementing established methods such as hashing and classifiers. Rather than waiting for abusive content to appear or be reported, platforms are investing in systems that surface suspicious activity and support earlier investigation.

Members are also strengthening their ability to identify **repeat offenders**. In 2025, **28 members** reported identifying behavioral patterns associated with recidivists, while **15 reported** developing models specifically designed to detect repeat abusive activity. **11 members** are integrating detection models with automated prevention mechanisms, enabling faster intervention when repeat offenders are detected.

Members use of technology to detect repeat offenders

28 members



Identifying behaviors of repeat offenders

15 members



Building recidivism-specific detection models

11 members

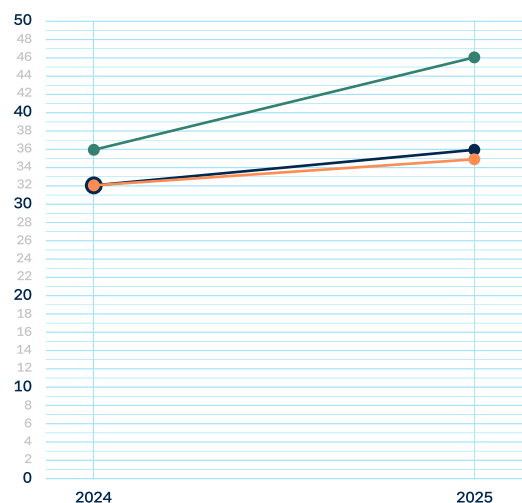


Integrating detection models with automated prevention mechanisms

User reporting

User and third party reporting remain critical for identifying abuse, with companies strengthening tools that enable users to report harmful activity.

User reporting processes



2025

- 46 In-product user reporting functionality
- 36 Dedicated webform (out of product)
- 35 Email alias and/or support team

In 2025, **46 members reported offering in-product reporting functionality**, up from 36 in 2024, showing these tools are increasingly integrated into user experiences. This expansion reflects efforts to make reporting simpler, faster, and easier to find.

Companies are continuing to refine reporting tools—improving accessibility, surfacing options more clearly, and enabling users to quickly flag concerning content or behavior.

Reporting from trusted third parties is also expanding. Members are strengthening channels for NGOs, law enforcement, and other partners, including dedicated portals and APIs to support faster handling of urgent cases.

The growth of trusted flagger systems, particularly in the EU, is further driving these actions.

5


Industry's response in 2025

Companies are strengthening how they act on detected abuse, improving investigations, intelligence sharing, and transparency across the ecosystem.

Milepost tracking again shows progress: members advanced **102 new response mileposts in 2025**, more than double 2024.

Platforms are placing more emphasis on the actionability of reporting to authorities, investing in internal systems to enable trust and safety teams to respond more quickly and effectively.

More companies are conducting deeper investigations into flagged accounts and suspicious behaviors, helping uncover broader networks of abuse and support lawful investigations.



Reporting to authorities

“The Tech Coalition’s practical guidance paired with a trusted forum for cross-industry learning has had a real impact on the child safety safeguards and practices we have developed at Cantina. Having participated in multiple workstreams, we’ve seen firsthand how shared expertise has strengthened industry’s practices. Most importantly, when we’ve encountered novel harms, we’ve been able to lean on the TC and mobilize jointly to disrupt bad actors, and create a safer online environment.”



Patricia Cartes
VP Trust & Safety
Cantina


Members are modernizing reporting to authorities with faster, more standardized processes.

In 2025, **38 members reported OCSEA cases to authorities using both API and manual reporting**, while **9 reported exclusively via APIs**, reflecting growing adoption of more efficient mechanisms.

This shift indicates increased investment in systems that support faster, more consistent, and scalable reporting. Many members are also overhauling internal processes to improve how reports are generated, validated, and shared with authorities such as NCMEC.

While manual reporting remains in use, several companies are phasing it out as they transition to automated systems. Together, these developments reflect broader efforts to strengthen operational response and ensure actionable information reaches authorities quickly and consistently.


Types of reporting used by members



Supplemental reporting

Members are expanding supplemental reporting to provide authorities with richer, more actionable information for investigations.

Members providing supplemental reports



In 2025, **39 members reported providing supplemental reports**, sharing additional details to help authorities better understand cases and act more effectively. This marks an increase from previous years, indicating greater investment in investigative processes and capabilities.

Supplemental reports provide contextual signals, investigative findings, and evidence that help organizations such as NCMEC and law enforcement prioritize cases and pursue investigations more efficiently. By sharing richer information, companies strengthen the quality and usefulness of reports.

Tech Coalition initiatives support this progress. The **Investigations and Supplemental Reports Knowledge Sharing Group**, with participants from **31 companies**, enables members to exchange practices and improve how insights are incorporated into reporting.

Reporting AI-generated abuse

As generative AI evolves, Tech Coalition members are developing new reporting templates to help platforms and authorities manage AI-generated abuse.

Reporting of AI CSAM

A Tech Coalition working group, including representatives from leading AI companies, developed a template to standardize reporting of AI-generated CSAM to NCMEC's CyberTipline.

The template captures details such as prompts, outputs, and instances where AI-generated content is shared across platforms.

As NCMEC updates its CyberTipline schema to provide companies the option to share more structured reporting of AI-generated abuse, this work is helping inform key elements of those changes and enabling more consistent, actionable reports.

Red-teaming reporting for model testing

As developers red-team generative AI image systems to identify misuse, testing can sometimes produce AI-generated CSAM. Because this occurs in controlled testing rather than real-world offending, it requires a distinct reporting approach.

To address this, Tech Coalition model-builder members created a dedicated **reporting template** for submitting this content to NCMEC.

The template clarifies the testing context and absence of a typical suspect or victim, helping NCMEC and law enforcement triage reports more effectively while enabling companies to adopt consistent practices for responsible reporting.



Red teaming is an adversarial testing process where experts simulate real-world misuse to deliberately probe AI models, exposing vulnerabilities and gaps in safeguards.

Law enforcement feedback loop

Feedback from authorities is helping members refine reporting and provide more actionable case information.

Members are strengthening reporting through ongoing feedback with authorities.

In 2025, **48 members reported having a feedback mechanism** with law enforcement or hotlines such as NCMEC, enabling companies to refine how cases are documented and submitted.


These feedback processes help companies better understand what information is most valuable to investigators and how reports can be made more actionable.

Several members also noted ongoing engagement with NCMEC to review reporting practices, helping improve the effectiveness of the information shared with authorities.

Tech Coalition initiatives support this exchange by creating structured opportunities for engagement with law enforcement. Through member-driven briefings, sessions, and events, law enforcement representatives share insights on investigative trends, reporting practices, and emerging challenges.

In 2025, **25 individuals from 21 companies** participated in the **law enforcement outreach working group**. Ongoing briefings and discussions with law enforcement partners continue to help members strengthen reporting workflows.

Together, these efforts reflect a growing commitment to continuously improve reporting practices and ensure authorities receive the most useful information to support investigations.



48 members reported having a law enforcement feedback loop in 2025

Investigations & enforcement

Members are strengthening response capabilities through tiered enforcement tools and investigative techniques, including open source intelligence methods.

Tiered enforcement is becoming an increasingly important operational capability.

While content removal and account suspensions remain essential, companies are expanding intermediary actions that introduce friction before harm escalates.

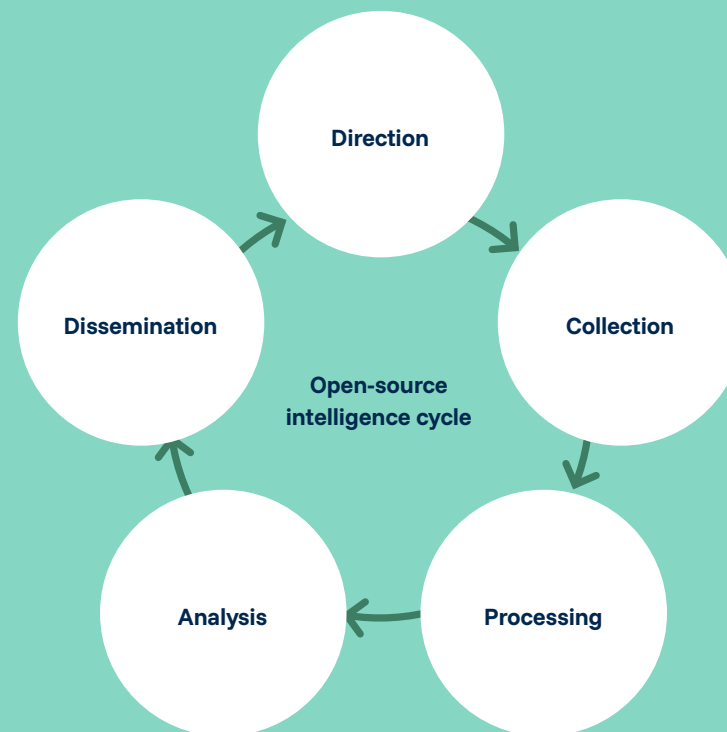
Supported by a Tech Coalition resource on tiered enforcement and a knowledge-sharing group involving nearly **20 companies**, members are deploying tools such as strike systems, timeouts, muting, content flags, and visibility limits.

These layered responses help platforms intervene earlier and disrupt harmful activity, including cases where offenders move across services to evade enforcement.

Members are also strengthening investigative capabilities through the use of **open-source intelligence (OSINT)**.

In 2025, more than **200 member representatives** participated in Tech Coalition OSINT training, learning techniques to uncover additional case context and support more effective reporting to law enforcement.

Together, these capabilities help companies investigate abuse more effectively and respond with more targeted enforcement actions.



Post-investigation enforcement tools

Investigative responses are becoming more structured and consistent across the industry.

Account suspensions and content removal have become standard enforcement tools following OCSEA investigations. Nearly all reporting members use account suspensions or blocks, with content removal or visibility restrictions close behind, indicating broad alignment on core response measures.


Beyond these baseline actions, many companies are deploying feature restrictions, network or device limitations, and account re-verification. See the graph on **member use of enforcement tools**.

These measures allow platforms to intervene more precisely, limiting an offender's ability to re-engage or escalate harmful behavior.

Across the industry, investigative responses reflect a maturing approach to enforcement, with companies applying tiered responses that combine multiple tools depending on the severity and context of abuse.

Together, these developments show that investigative responses are becoming more standardized, operationally structured, and better equipped to disrupt abusive activity.


Member use of enforcement tools



Lantern signal sharing

Lantern connects signals across platforms to strengthen coordinated responses to OCSEA.

Lantern enforcement actions in 2025



Lantern responds to a core challenge in combating OCSEA: harmful activity rarely occurs on a single platform. Offenders may groom a child on one service, distribute abusive material on another, and facilitate financial gain elsewhere, leaving critical signals fragmented across platforms.

A cross-platform signal-sharing infrastructure, Lantern enables participating companies to connect these signals, surface coordinated OCSEA activity, and take earlier action.

In 2025, Lantern continued to scale as operational safety infrastructure.

Thirty-one companies shared nearly 1 million signals, contributing to enforcement actions against **31,504 accounts**, **27,646 URLs** and **17,778 pieces of content**.

Collaborative workstreams for Lantern participants are also expanding to address emerging risk patterns, including grooming and enticement in gaming, financially motivated OCSEA, and CSAM hosting and distribution.

Cross-sector collaboration is also growing.

The Lantern financial sector pilot demonstrated the value of intelligence sharing between technology platforms and financial institutions, with six participating organizations sharing **1,935 signals which supported 108 investigations** into OCSEA-related financial activity.

Following the pilot's success, Lantern is now open to eligible financial institutions, reinforcing its role as core safety infrastructure enabling coordinated disruption of OCSEA across platforms and sectors.



See the Lantern transparency report 2025

“Protecting children online is a challenge the entire industry must take on together. Meta is proud to partner with the Tech Coalition to support Lantern, a system that allows us to share vital safety information across the industry. By working together, we can more effectively prevent abuse across the internet.”




Ravi Sinha
Director, Head of Child Safety Policy, Meta

Lantern signals uncover child exploitation network

“Through the Lantern workstreams, I’ve seen firsthand how powerful cross-industry collaboration can be. When private companies work together to responsibly discuss signals and emerging risks, we strengthen our collective ability to identify and disrupt illicit activity. This model of collaboration is critical to protecting children.”



Megan Gonzales
Global Head of External Engagements & Intelligence
Block



Meta received a request from law enforcement related to an ongoing investigation into the production and distribution of CSAM.

The case involved an adult male sharing original abuse material depicting a minor and coordinating exploitation that appeared to involve financial transactions.

Meta’s investigation uncovered attempts to arrange in-person meetings with minors in exchange for money, goods, or services. The activity suggested a coordinated networking involving individuals with direct access to children.

Meta reported the case to **NCMEC** and shared relevant signals through **Lantern**.

These signals enabled **Block** to identify a network of 18 Cash App accounts connected to suspected CSAM sales or purchases.

The accounts were closed, the users banned, and the activity reported to authorities.

Microsoft was also able to identify and take enforcement action against several Xbox accounts linked to the case.

Law enforcement later confirmed that the information Meta shared contributed to a multi-agency investigation resulting in multiple arrests for child exploitation offenses

Transparency & accountability

More members are aligning reporting with the Trust framework, strengthening transparency across the tech sector.

“The Trust Framework provided our trust & safety team with principle-based guidance on how to effectively capture child protection outcomes in our 2025 transparency report. By clearly outlining our detection, moderation, and enforcement practices, we improved transparency and reinforced VSCO’s commitment to protecting children.”



Jenna Dietz
Trust & Safety Manager
VSCO


In 2025, **25 members were fully aligned with the Trust Framework** and 17 partially aligned, reflecting growth in transparency reporting.

Companies also strengthened reporting in response to regulations such as the EU Digital Services Act, while maintaining alignment with Tech Coalition standards.

This expansion reflects broader efforts to provide clearer insight into how platforms address abuse.




































To support this, the Tech Coalition committed to develop draft updates and guidance on prevalence measurement and generative AI risks, helping improve consistency in reporting, with finalization planned for 2026.

Members aligned to Trust Framework



Member transparency reports

Click on a member's logo to see how each company approaches its commitment to transparency in their own reports.

6

The future for the Tech Coalition

In 2026, the Tech Coalition will continue to scale the infrastructure, tools, and collaboration that are already driving measurable progress across the industry, as we respond to a fast-moving and increasingly complex threat environment.

Artificial intelligence is changing this landscape in real time. It is increasing the scale, speed, and adaptability of abuse, while also strengthening the capabilities used to detect and disrupt it. How these capabilities are applied will determine how effectively industry can stay ahead of harm by identifying risks earlier, and acting with greater precision at scale.

Our members are operating in a more **interconnected environment**, where activity and risk span platforms, models, and services. No single company has full visibility. This makes coordination across the ecosystem essential, enabling companies to act on shared signals and respond to threats that move across services. **Lantern** demonstrates this in practice.

We will continue to scale what is working, deepening collaboration through working groups and using shared insight to identify and close gaps across members. We will also strengthen our collective ability to anticipate and respond to emerging harms, ensuring we stay ahead of a shifting risk landscape.

Mileposts will remain invaluable in benchmarking and demonstrating progress across prevention, detection and response.

This work extends beyond our membership. **Pathways** will continue to be freely available across the tech ecosystem, and in 2026, we will formally launch **Elevate**, offering hands-on consulting and support to help companies strengthen their child safety foundations and prepare for membership. Strategic growth will continue, with deeper engagement in **APAC**, supported by our expanded presence in Singapore.

Transparency and measurement will be advanced through new public resources, including guidance on AI red teaming, what transparency reports can tell us about prevalence, and effective reporting in the age of generative AI.

We will also expand industry reporting templates for AI-generated CSAM, including a third template addressing material identified in AI model training data.

In the coming years, we will continue working with governments, regulators, law enforcement and civil society to ensure that industry progress aligns with and exceeds evolving regulatory expectations.

We will also continue to focus on strengthening the actionability of reports shared with law enforcement, supporting more effective investigations and outcomes.

The challenge is evolving as rapidly as the technology itself. But our response is evolving too. Through **coordinated industry action**, we are strengthening safeguards, disrupting abuse earlier, and driving measurable progress in protecting children online.

And we will continue to build on that progress, **together**.



The Tech Coalition unites the global tech industry to protect children from online sexual exploitation and abuse (OCSEA).

No single company can tackle this issue alone. But together, we're pooling knowledge, identifying threats, and developing solutions to prevent harm. Because children deserve a digital world that's safe to play, learn, and explore.

technologycoalition.org



Images: To protect children's privacy and dignity in the context of our work, all images of children on this report are AI-generated.