
Flagged and Failed: Platform Responses to Antisemitic Content Reported by Jewish Communities

Audit of reports via HO:PE portal

February 2025 – April 2026

Table of Contents

1. About the HO:PE Project	2
2. Executive Summary	4
3. Methodology	5
4. The Scale of Online Antisemitism	7
5. Where Antisemitism Lives: Platform Analysis	8
6. What Hate Looks Like: Category Analysis	10
7. Platform Accountability: Reporting and Removal	13
8. Legal Assessment	17
9. Conclusions and Recommendations	21

1. About the HO:PE Project

The Hate Online: Preparedness and Empowerment (HO:PE) initiative is led by the World Jewish Congress and funded by the European Union under the CERV programme (CERV-2023-CHAR-LITI-SPEECH). Launched in 2024, the project was designed to address a well-documented gap: the chronic under-reporting of antisemitic content online, and the lack of structured data to support advocacy, legal action, and platform accountability.

The project has three core components:

- The HO:PE reporting tool, available as a browser-based and a mobile application, which enables Jewish community members across the EU to report antisemitic content they encounter online.
- A capacity-building programme, through which Jewish organisations, Jewish communities, and digital contact points in those communities are trained to use the tool, understand relevant legislation, and engage with platforms and authorities.
- This report, which presents the findings from the data submitted by Jewish community members across the EU through the HO:PE reporting tool, analysing trends, platform behaviour, and the legal dimensions of the reported content.

This is the final report produced under the project, covering submissions made between February 2025 and April 2026. It represents the first systematic, community-sourced dataset of antisemitic online content collected through a purpose-built reporting tool across the European Union.

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Education and Culture Executive Agency (EACEA). Neither the European Union nor EACEA can be held responsible for them.

Reports submitted

754

Feb 2025 - Apr 2026

Still live at review

81%

613 of 754 posts

TikTok removal rate

64%

vs. 12-15% on other platforms

Reporting made no difference

~19%

% of both reported and not-reported
content removed

May warrant legal review

~6%

estimated from a sample review, under
EU and member state law

2. Executive Summary

Between February 2025 and April 2026, Jewish community members across Europe submitted 754 reports of antisemitic content encountered online through the HO:PE reporting tool. This document analyzes those submissions: what was reported, where it appeared, how platforms responded, and what portion of it may constitute illegal content under EU and member state law.

The most commonly reported content includes dehumanization of Jews, general calls for harm to Jewish people, calls for killing Jews in the name of a radical ideology, and promotion of a world Jewish conspiracy. Together, these four categories account for more than half of all submissions.

Facebook accounted for the largest share of reported content (34.4%), followed by Instagram and X at 26% each. A liveness check conducted in May 2026 found that 81% of reported posts remained publicly accessible, and that formally flagging content on a platform through its reporting mechanisms made virtually no measurable difference as to whether it was removed, with removal rates of around 19% for both content that was and was not reported. TikTok is the exception, removing 64% of content overall (and 71% of content formally reported to it), compared to 12–15% on every other major platform.

Israel-related antisemitism proved the most resistant to removal of any category, with content comparing Israeli policy to Nazism having a 0% removal rate across 28 reports in that subcategory.

A preliminary legal assessment, based on a sample review of posts in the highest-risk categories, estimates that approximately 6% of all reports may warrant individual legal examination under EU or member state law. This is an order-of-magnitude estimate, not a confirmed count; each case requires individual legal analysis before any formal action is taken.

The analysis points to ongoing significant gaps in the moderation of antisemitic content across all social media platforms. Users continue to encounter antisemitism on social media. Based on these findings:

- Platforms must review whether their policies, enforcement, and reporting mechanisms sufficiently account for antisemitism.
- Regulators and policymakers must ensure enforcement of laws and regulations across all platforms and content.

3. Methodology

3.1. The HO:PE Reporting Tool

The HO:PE tool is a community-based reporting mechanism developed by the World Jewish Congress to enable Jewish community members across Europe to document antisemitic content encountered online. It is accessible both as a web form and as a mobile application, and is open to anyone wishing to submit a report. The tool was distributed to Jewish communities across the EU.

When submitting a report, users are asked to provide the URL of the content, select the platform on which it appears, classify the type of antisemitism from a structured category list, and optionally provide a description and screenshot. Users can also indicate whether they have already reported the content to the platform directly.

3.2. Dataset Scope and Period

The dataset analysed in this report covers submissions made between February 2025 and April 2026. After excluding test entries and invalid submissions, the dataset comprises **754 reports**. Each report represents a unique submission by a community member, though in some cases the same URL was submitted by multiple users independently—a pattern that is itself noted where relevant, as it indicates content that was encountered and found troubling by more than one person.

3.3. Key Terminology

The HO:PE classification framework is grounded in the International Holocaust Remembrance Alliance (IHRA) Working Definition of Antisemitism, which defines antisemitism as "a certain perception of Jews, which may be expressed as hatred toward Jews [...]" and provides a set of contemporary examples of antisemitism in public life. The HO:PE tool operationalises this definition into a structured category system organised across four thematic groups:

Traditional Antisemitism: promoting the idea that Jews control the media; the dehumanisation of Jews; promoting traditional tropes such as blood libels and claims that Jews killed Jesus; promoting the idea of a world Jewish conspiracy; holding Jews collectively responsible for acts committed by individuals; promoting the idea that Jews control government or other societal institutions; accusing Jewish citizens of being disloyal to their country; and promoting the idea of Jews controlling the economy.

Antisemitism Related to Israel: comparisons of Israeli policy to Nazism; accusing Israel of inventing or exaggerating the Holocaust; describing Israel or Israelis using antisemitic words or imagery; denying Jewish people self-determination, e.g., by claiming Israel's existence is racist; requiring a behaviour from Israel not expected of other countries; and holding Jews collectively responsible for Israel's actions.

Holocaust-Related Content: Holocaust jokes; denying the Holocaust; accusing Jews or Israel of exaggerating the Holocaust; blaming Jews for the Holocaust; inappropriate comparisons with Nazis; distorting the facts about the Holocaust; and glorifying the Holocaust or suggesting it did not go far enough.

Incitement to Violence: calling for, aiding, or justifying the killing or harming of Jews in the name of a radical ideology or an extremist view of religion; calling for harm to Jewish property; calling for harm to Jewish people in general; calling for harm to non-Jews for supporting Jews or opposing antisemitism; calling for harm to someone because they are Jewish; and calling for harm to someone believed to be Jewish.

Submissions that did not clearly fit any of the above categories could be classified as **Other**.

3.4. Limitations

Several limitations should be kept in mind when interpreting the findings of this report.

Attribution of removal: The liveness check conducted in May 2026 identified posts that were no longer publicly accessible at the time of review. However, it is not possible to determine in all cases why a post was removed—whether by the platform following a community report, by the original poster themselves, or for other reasons such as account deletion or account restrictions. Removal figures should therefore be read as reflecting content that is no longer live, not necessarily content that platforms acted upon.

Classification accuracy: The categorisation of submitted reports was carried out by community members at the time of submission and reviewed on a sample basis. While spot-checking identified no systematic misclassification, individual errors cannot be ruled out. The category distributions presented in this report should be understood as approximate rather than exact.

Comments and embedded content: Some antisemitic content encountered by community members appears in the comments sections of posts rather than in the posts themselves. The HO:PE tool often captures the URL of the parent post, which means the reported content may not be immediately visible from the link alone. This introduces a margin of uncertainty in both directions: some reports may relate to comment-level content that has since been removed independently of the post, while other antisemitic comments on reported posts may not have been captured at all. This is noted as a methodological characteristic of community-based reporting tools.

Geographic origin of content: While the reports in this dataset were submitted by Jewish community members across Europe, this does not mean that the content they reported originates from Europe. The origin of a post—the location of the account that published it—can be anywhere in the world. This distinction is relevant for interpreting the dataset, and should be kept in mind when considering the reach and scale of antisemitic content encountered by Jewish communities online.

4. The Scale of Online Antisemitism

4.1. Geographic Distribution

These reports span six major social media platforms and a small number of other online sites, and were submitted by community members across multiple countries during this period. Of the 754 reports, 43 (5.7%) were explicitly linked to a specific Jewish community in Europe by the submitter. Spain accounts for the largest share of community-tagged reports (17), followed by Ireland (10), France (7), Italy (3), Germany (2), and single reports from Hungary, Sweden, the United Kingdom, and Denmark. A further 122 reports (16.2%) were marked as having no specific community relevance, while the remaining 589 were submitted without a community tag, suggesting that the content might be general to all Jewish communities.

5. Where Antisemitism Lives: Platform Analysis

5.1. Distribution Across Platforms

The reports span six major social media platforms and a small number of other online sites. Facebook, Instagram, and X (Twitter) together account for 86.3% of all reports.


Facebook (259 reports, 34.4%) is the single largest source in the dataset, driven primarily by a cluster of Arabic-language posts celebrating Iranian missile attacks on Israel and dehumanising Jewish people.

Instagram (196 reports, 26.0%) and X (196 reports, 26.0%) are tied for second place, together accounting for more than half of all reports.

TikTok (73 reports, 9.7%) is smaller in volume but notable in severity; several Holocaust glorification entries were documented on the platform.

YouTube (12 reports, 1.6%), Reddit (3 reports, 0.4%), and a tail of other sites account for the remainder.

Reports by platform



5.2. Platform Analysis

Each platform in the dataset presents a distinct profile in terms of the type of antisemitism documented and its removal rate.


Facebook dominates the dataset in volume (259 reports). The content here sits at the higher end of the hierarchy of hate expression, weighted towards direct calls for harm and dehumanisation ([example](#)). Despite having the highest user reporting rate—76% of Facebook reports were flagged to the platform—only 14% of content was removed, the lowest removal rate among major platforms.

Instagram (196 reports) shows content skewing towards dehumanisation and conspiracy narratives ([example](#)). Its reporting rate is very low at just 7%, meaning 183 of 196 incidents were never formally flagged to the platform, and its overall removal rate stands at 12%.

X (196 reports) leans more heavily into world Jewish conspiracy tropes and blood libel imagery ([example](#)). Its reporting rate was 17%, with a 15% overall removal rate.

TikTok (73 reports) stands apart from the other platforms. Despite containing content at the higher end of the hierarchy of hate expression, including Holocaust glorification ([example](#)), it has by far the highest removal rate at 64% overall, and 71% for content that was formally reported to the platform.

YouTube (12 reports) and **Reddit** (3 reports) are too small in this dataset to support firm conclusions, though YouTube's content includes Holocaust denial and blame ([example](#)).



6. What Hate Looks Like: Category Analysis

6.1. Most Prevalent Forms of Antisemitism


The reports span 27 distinct categories, classified using the IHRA definition of antisemitism as the analytical framework. The distribution is not even: four categories alone account for more than half of all reports.

Dehumanisation ([example](#)), a subcategory of **Traditional Antisemitism** according to the HO:PE categorisation system, is the single most reported form, appearing in 135 reports (17.9%), followed closely by **general calls for harm to Jewish people** (133 reports, 17.6%), a subcategory of Incitement to Violence. Together these two categories represent more than a third of the entire dataset and reflect content that sits at the higher end of the hierarchy of hate expression.

Calls for violence in the name of a radical ideology or religion ([example](#)) account for a further 78 reports (10.3%), and **promotion of a world Jewish conspiracy** for 77 reports (10.2%). These four categories combined account for 56% of all submissions.

Beyond the top four, the data demonstrates a wide spread of antisemitic expression. Holocaust-related content appears across five separate categories totalling 66 reports (8.8%). Conspiracy narratives beyond world domination, including claims of Jewish control over government, media, and the economy ([example](#)), account for a further 50 reports (6.6%). Israel-related antisemitism—including Nazi comparisons ([example](#)), collective responsibility, and denial of self-determination—accounts for 97 reports (12.9%).

Reports by category group



Based on 754 reports submitted to the HO:PE tool, February 2025 - April 2026

6.2. Type of Antisemitism by Platform

The distribution of categories of antisemitism varies across platforms, revealing distinct patterns in how hate manifests in different online environments.

Facebook is dominated by incitement to violence, which accounts for **59%** of all Facebook reports, nearly double that of any other major platform. This reflects the large cluster of posts calling for harm to Jewish people documented in this dataset, including direct calls for killing Jews framed as a religious obligation ([example](#)), and dehumanising content using slurs such as "descendants of apes and pigs" ([example](#)). Israel-related antisemitism on Facebook is comparatively low at 3%, though examples of sustained hate targeting Israel were documented ([example](#)).


Instagram shows a more even spread across categories. Traditional antisemitism accounts for 51% of Instagram reports, including conspiracy narratives presenting Jews as controlling the world ([example](#)). Israel-related content (15%) includes posts accusing Israel of genocide ([example](#)); Holocaust-related content (15%) includes outright denial of the six million figure ([example](#)), and incitement to violence (6%) includes content glorifying Nazi victory ([example](#)).

X has the highest share of Israel-related antisemitism of any major platform at 25%, including content framing Zionists as "evil scum" ([example](#)). Holocaust-related content accounts for 16% of X reports, with examples including posts praising Hitler for targeting Jewish banking families ([example](#)).

TikTok mirrors Facebook in its high share of incitement to violence (58%), with content invoking radical ideology alongside more traditional forms of antisemitism, including conspiracy narratives about Jewish ubiquity ([example](#)) and Holocaust glorification ([example](#)). Israel-related content on TikTok is comparatively low at 5% ([example](#)).

YouTube stands out for its concentration of Holocaust-related content (58%, though small in volume), including videos featuring Holocaust denial ([example](#)).

Category distribution by platform




YouTube (12 reports) and Reddit (3 reports) shown for completeness; small sample size limit conclusions.

6.3. Spotlight: Israel-Related Antisemitism

Israel-related antisemitism accounts for **97 reports (12.9%)** of the dataset, making it the third largest category group. It is notably concentrated on two platforms: **X** accounts for 50.5% of all Israel-related reports, and Instagram for a further 30.9%, together representing over 80% of this category.

Its largest subcategory is **comparisons of Israeli policy to Nazism** (28 reports, 28.9%), which includes content equating Israel Defense Forces (IDF) policy with Nazi ideology ([example](#)). **Describing Israel or Israelis using antisemitic words or imagery** accounts for 22 reports (22.7%), including blood libel tropes and dehumanising language. **Denying the Jewish people's right to self-determination**—for example by framing Israel's existence as inherently racist—accounts for 20 reports (20.6%), including content accusing Israel of genocide ([example](#)). A further 14 reports (14.4%) fall under the heading of **applying double standards to Israel**, such as campaigns to ban Israeli athletes from international sport ([example](#)). Finally, 10 reports (10.3%) involve **holding Jews collectively responsible for Israel's actions**, including cases in which Jewish community organisations were accused of supporting genocide for speaking out against antisemitism ([example](#)).

Israel-related antisemitism by subcategory



Based on 97 Israel-related reports out of 754 total, February 2025 - April 2026.

7. Platform Accountability: Reporting and Removal

7.1. User Reporting Behaviour


Of the reports submitted via the HO:PE tool, **277 (36.7%)** were also reported to the host platform by the user at the same time. The remaining **477 (63.3%)** were documented in HO:PE but were not formally flagged to the platform.

Reporting rates vary considerably by platform. Facebook stands out with a **76.1%** reporting rate—by far the highest—suggesting that users who encountered antisemitic content on Facebook were significantly more likely to report it to the platform as well.

Instagram sits at the opposite extreme, with just **6.6%** of reports formally flagged, meaning 183 of 196 Instagram incidents were never reported to the platform. X (16.8%) and TikTok (23.3%) fall in between.


Reporting rates also differ by content type. Users were most likely to report **incitement to violence**—61.2% of incitement reports were also flagged to platforms—likely reflecting that this category presents the clearest and most unambiguous violation of platform rules, making it feel most worth reporting. By contrast, **Israel-related antisemitism** had the lowest reporting rate at just 12.4%. One possible explanation—though it cannot be confirmed by the data alone—is that users have learned from experience that certain content types are less likely to be removed and adjust their reporting behaviour accordingly, choosing not to invest the effort when they expect no result.

Reporting to platforms, by platform



Based on 754 reports submitted to HO:PE tool, February 2025 - April 2026

Reporting to platforms, by category



Based on 754 reports submitted to HO:PE tool, February 2025 - April 2026

7.2. Overall Removal Rates

Of the reports assessed for liveness through May 2026, **141 (18.7%) had been removed** and **613 (81.3%) remained live** at the time of the check.

When comparing content that had been formally reported to the host platform with content that had not, the removal rates were nearly identical. Among the **277 posts reported** to platforms by users, 54 (**19.5%**) had been removed. Among the 477 posts that had not been reported to a platform, 87 (**18.2%**) had been removed.

This suggests that formally reporting content to a platform was not associated with a meaningfully higher likelihood of removal.

It is important to interpret this finding carefully. The near-identical removal rates between reported and unreported content may reflect a number of factors beyond platform moderation decisions alone, including differences in content type, the time elapsed between submission and the liveness check, and the inherent limitations of a snapshot taken at a single point in time. What the data allows us to say with confidence is that across this dataset, the large majority of antisemitic content—**more than four in five posts**—remained accessible online at the time of the check, regardless of whether it had ever been formally flagged by a user.

Removal status of reported content



7.3. Platform-by-Platform Breakdown

Facebook had the highest volume of reports (259) and the lowest overall removal rate at **13.9%**. Despite 76.1% of Facebook posts being formally flagged by users—the highest reporting rate of any platform—only 28 of those 197 reported posts were removed (14.2%). Of the 62 posts never reported to Facebook, just 8 were removed (12.9%).


Instagram shows a different pattern. With a reporting rate of just 6.6%, only 13 of 196 reports were formally flagged, yet 4 of those 13 were removed (30.8%), which translates to a higher success rate than that of Facebook. A further 20 posts were removed without ever being reported (10.9% of unreported content). Overall removal rate: **12.2%**.

X sits between the two, with a 16.8% reporting rate and an overall removal rate of **14.8%**. Of 33 reported posts, 8 were removed (24.2%). A further 21 were removed without being reported (12.9% of unreported content).

TikTok stands apart from all other platforms. With an overall removal rate of **64.4%**, it removed 12 of 17 reported posts (70.6%) and 35 of 56 unreported posts (62.5%)—the only platform where the majority of antisemitic content, whether reported or not, was taken down.

YouTube (12 reports, all reported) had a removal rate of **16.7%**, with 2 of 12 posts removed. **Reddit** (3 reports, all reported) had a **0%** removal rate. Both samples are too small to draw firm conclusions.

Removal status by platform



7.4. What Stays Online: A Category and Platform View

Looking at the overall removal rates by content category reveals a consistent pattern: **Israel-related antisemitism is by far the least likely to be removed**, with only 5 of 97 reports (5.2%) taken down—the lowest removal rate of any category group. Within this category, content comparing Israeli policy to Nazism had a **0% removal rate** across 28 reports. Antisemitic imagery applied to Israel or Israelis (4.5%) and denial of Jewish self-determination (5.0%) were similarly persistent.


Traditional antisemitism (13.6% removed) and **Holocaust-related content** (13.6% removed) also have low removal rates, with dehumanising content at just 7.4% and Holocaust glorification at 5.3%.

Incitement to violence has the highest removal rate of any category at **34.9%**, and within it, content calling for killing or harming Jews in the name of radical ideology had a removal rate of 52.6%—the highest of any subcategory.


This is consistent with the earlier finding that users are more likely to report incitement and suggests that this category may be more clearly covered by existing platform policies.

Across all categories, the data shows that **613 posts—more than four in five—remained live** at the time of the check. The content most likely to remain online is that which sits at the intersection of political discourse and antisemitism, particularly Israel-related content.

Removal status by category group



Removal rate by subcategory (min. 10 reports)



8. Legal Assessment

8.1. Methodology and Scope

A subset of the reports submitted through the HO:PE tool may constitute illegal content under EU and member state law. This section offers a preliminary, high-level assessment of that potential. It is not a legal opinion, and any individual case would require deeper legal analysis before any formal action is taken. The purpose of this chapter is to indicate where the probability of crossing the illegality threshold is highest, not to make a legal determination.

To identify the reports most likely to involve illegal content, the data was analysed using a two-part filter.

The first filter was category. Not all antisemitic content is illegal. Much of it, while harmful, falls within the bounds of lawful expression. We focused on the categories that, by their nature, sit closest to the threshold of illegality established under EU and member state law. These are:

- Calling for harm to Jewish people in general (133 reports)
- Calling for, aiding, or justifying the killing or harming of Jews in the name of a radical ideology (82 reports)
- Glorifying the Holocaust or suggesting it did not go far enough (19 reports)
- Calling for harm to a specific person because they are Jewish (8 reports)
- Denying the Holocaust (10 reports)
- Blaming Jews for the Holocaust (9 reports).

Together these categories account for 261 reports—34.6% of the full dataset of 754 reports. It is important to stress that not all content within these categories will meet the legal threshold for illegality in any given jurisdiction; that requires a case-by-case assessment. These 261 reports should be understood as an upper bound for further review, not a confirmed count of illegal content.

The second filter was liveness. Within that pool, priority was given to posts confirmed as still live as of May 2026. Content that remains publicly accessible presents an ongoing harm, and its continued availability may itself be relevant to platform liability under the Digital Services Act (DSA). Of the 261 high-potential posts, 173 were confirmed still live at the time of the liveness check, and 88 had been removed or were otherwise unavailable. These 173 live posts—representing 23% of the full dataset—formed the pool for the legal review.

Sample review and estimation

Of the 173 live high-potential posts, a sample of 54 was reviewed. Each post was assessed against the legal frameworks set out below and marked as potentially warranting legal examination, or not meeting the threshold on the available information.

Of the 54 posts reviewed, 14—approximately 26%—were assessed as potentially warranting further legal examination.

Applying this rate to the full pool of 173 live high-risk posts, and expressing the result as a share of the full dataset, yields an estimated 6% of all reports as potentially warranting individual legal examination. This figure should be treated as an order-of-magnitude estimate, not a precise count. The sample was not proportionally drawn across all six categories, and "potentially warrants legal examination" is not equivalent to "confirmed illegal." Each case requires individual legal analysis before any formal action is taken.

The legal frameworks

Two EU instruments are relevant as reference points for this assessment. EU Framework Decisions and EU Directives are not laws that apply directly to individuals. Instead, they require every EU member state to have national criminal law covering the conduct in question. The applicable law in any individual case is therefore always the relevant national law, not the Framework Decision or the Directive directly.

The [EU Framework Decision on Racism and Xenophobia \(2008/913/JHA\)](#) requires all member states to criminalise public incitement to violence or hatred directed at a group based on religion or ethnicity, as well as the public condoning, denial, or gross trivialisation of genocide, crimes against humanity, and war crimes as defined in the Statute of the International Criminal Court and in the Charter of the International Military Tribunal appended to the London Agreement of 1945 where expressed in a manner likely to incite violence or hatred. Where a post's geographic origin is identifiable, the applicable law is the relevant national criminal law, for example, §130 [StGB](#) in Germany, which criminalises incitement to hatred against segments of the population and Holocaust denial, and §86a StGB, which independently criminalises the display of Nazi symbols. Other member states have equivalent national provisions that may also be deemed illegal depending on the origin of the content, though a full assessment of those frameworks is beyond the scope of this report.

The [EU Directive on Combating Terrorism \(2017/541\)](#) requires all member states to criminalise public provocation to commit terrorist acts, an offence that encompasses the glorification and justification of terrorism and the dissemination of related messages or images online and offline.

8.2. Findings

The posts assessed as potentially warranting legal examination in the sample review span five clusters: terrorist incitement and operational threats, explicit calls to murder, Holocaust denial and glorification, Holocaust revisionism framing Jews as perpetrators, and dehumanisation. What follows describes the content identified within each cluster. This is not a confirmed count of illegal content; each case would require individual legal examination before any formal action is taken. All posts cited below were reported to the relevant platform by community members at the time of submission.

Terrorist incitement and operational threats

The most severe cluster consists of posts sharing an official ISIS publication (*al-Naba* Issue 541) calling on Muslims globally to burn Jewish synagogues and attack Jewish neighbourhoods across Europe and the United States. Six shares of this publication were identified on Facebook, each with a distinct URL, each reported to the platform and each still live at the time of the liveness check ([source](#)). A further post, published [the day after the deadly September 2025 shooting at a Jerusalem bus stop](#) called for "operations (attacks) in all regions of the world against everyone who has a relationship with the Jews. Kill them" ([source](#)).

One post stands out for its operational specificity: a YouTube video involving a named individual based in Copenhagen, in which he celebrates terrorist attacks against Jews and explicitly threatens to carry out further attacks ([source](#)). It was reported to YouTube and confirmed still live.

A further post on X describes the drawing up of lists of Zionists ([source](#)). Reported to the platform and confirmed live, this post includes explicit threat language.

Explicit calls to murder Jews

Six Facebook posts contain unambiguous calls to murder Jewish people, including phrases such as "Kill the Jews wherever you find them" ([source](#)). All were reported to Facebook and remained live. This content involves conduct that the Framework Decision on Racism and Xenophobia (2008/913/JHA) requires member states to criminalise.

Three further Facebook posts were assessed as potentially falling within the same frameworks, each calling for or justifying the killing or harming of Jews in the name of a radical ideology ([source](#), [source](#), [source](#)).

Holocaust denial and glorification

A five-part YouTube series frames Nazi propaganda as historical analysis; all five videos were reported to the platform and confirmed live ([example](#)). Posts on TikTok and Facebook suggest the Holocaust did not go far enough, including a TikTok video stating the world would be better if 6 million Jews had actually died ([source](#)).

A further cluster of posts on X glorify the Holocaust or suggest it did not go far enough ([example](#)); all were assessed as potentially reaching the relevant threshold for illegality. One post on Instagram and one on X were assessed as potentially constituting Holocaust denial ([source](#)). All were reported to their respective platforms.

Holocaust revisionism framing Jews as perpetrators

Two posts—one on X and one on Facebook—frame the Holocaust as a fabrication invented to serve Jewish financial interests and to conceal Jewish responsibility for the war. The X post states that Hitler arrested members of the Rothschild family and shut down their banks, and that the Holocaust was fabricated to conceal this ([source](#)).

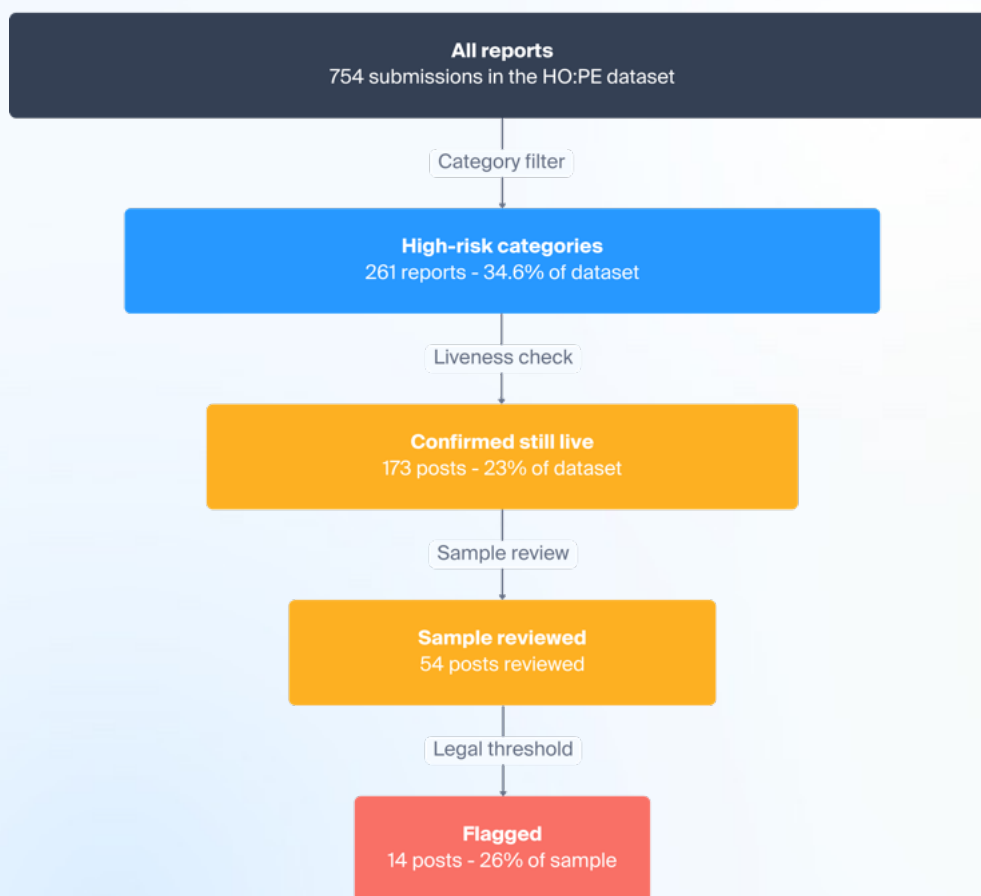
Dehumanisation

A cluster of Facebook posts dehumanises Jews as "descendants of apes and pigs", many of which were posted in the context of celebrating Iranian missile attacks on Israel. One Facebook account has a Nazi swastika profile picture ([source](#)), the display of which is independently criminalised under §86a StGB in Germany. All were reported to Facebook.

Summary

Based on a sample review of 54 posts drawn from the 173 live high-risk reports, approximately 6% of all reports in the HO:PE dataset are estimated to contain content that may warrant individual legal examination.

The legal assessment: from 754 reports to the ~6% estimate



~6% of all 754 reports estimated to warrant individual legal examination
(order-of-magnitude estimate; not confirmed count of illegal content)

Taken together, the legal assessment indicates that a small but meaningful share of the content documented through HO:PE may cross the line from harmful into potentially illegal, and that this content, despite being reported, overwhelmingly remained online. This finding underpins the recommendations to platforms and policymakers set out in the following chapter.

9. Conclusions and Recommendations

9.1. Recommendations

For Platforms	
Recommendations	Findings & Data
1. Platforms must review whether their policies sufficiently account for antisemitism.	Four in five reported antisemitic posts are still online. 613 of 754 reported posts (81%) remained publicly accessible at the time of the May 2026 liveness check.
2. All platforms should review their level of enforcement of policies related to antisemitism.	Enforcement levels are low on reports of antisemitic content. TikTok removed 64% of reported content. Every other major platform removed less than 15%.
3. Platforms should review whether reported content is properly directed and acted on.	Facebook has the highest reporting rate and the lowest removal rate. 76% of Facebook posts submitted to HO:PE were also reported to the platform directly—the highest rate of any platform. Facebook's removal rate: 13.9%, the lowest of all major platforms.
4. Platforms should demonstrate, with data, that their reporting channels are operationally effective.	Reporting content on a platform was not associated with a meaningfully higher likelihood of removal. Removal rate for reported content: 19.5%. Removal rate for unreported content: 18.2%.
5. Platforms should review whether their policies cover content that uses the Israel–Palestine conflict as a vehicle for antisemitism.	Israel-related content: a blind spot in platform enforcement. 0% of posts comparing Israeli policy to Nazism were removed (28 reports). Across all content displaying Israel-related antisemitism, only 5.2% of 97 reports were taken down, the lowest removal rate of any category.

For Policymakers

Recommendations

1. National law enforcement authorities should establish clear referral pathways for antisemitic online content that may constitute a criminal offence, and ensure that prosecutors and investigators have the capacity and training to act on such referrals.

2. Policymakers should ensure that national frameworks prioritise the detection and reporting of online content that constitutes incitement to violence or terrorist provocation, as well as the prosecution of those posting such content, and that platforms are held accountable for its timely removal.

3. Regulatory authorities should use their powers under the Digital Services Act to scrutinise platform risk mitigation practices with respect to antisemitic content, including through requests for information, audit mechanisms, transparency reporting requirements, and formal investigations where findings warrant further examination.

4. Policymakers should support the continuation and expansion of community-based reporting programmes, trusted flagger schemes, and capacity-building initiatives that enable civil society to document and respond to online antisemitism systematically.

Findings & Data

A share of reported content may meet the threshold for illegality under existing law.

A preliminary legal assessment estimates that approximately 6% of all reports may warrant individual legal examination under EU or member state law, based on a sample review of the highest-risk categories.

Violent and incitement content represents the most urgent enforcement priority.

Calls for harm to Jewish people and for killing Jews in the name of a radical ideology together account for 28% of all reports—the two largest single categories in the dataset. The majority of this content remained live at the time of the liveness check despite having the highest user reporting rate of any category.

Platform enforcement of existing policies remains insufficient.

Across all major platforms except TikTok, fewer than 15% of reported antisemitic posts were removed. Formal reporting by users made virtually no measurable difference to removal outcomes.

Community-based reporting infrastructure requires sustained institutional support.

The 754 reports analysed in this report were collected over 14 months through a purpose-built community reporting tool—the first systematic, community-sourced dataset of antisemitic online content collected across the EU.

About the WJC

The World Jewish Congress (WJC) is the international organization representing Jewish communities in 100 countries to governments, parliaments and international organizations.

About TechRI

TechRI, the World Jewish Congress Technology and Human Rights Institute, works at the intersection of technology, human rights, and Jewish community life. It focuses on countering online antisemitism, extremism, and disinformation, while also promoting safer and more ethical uses of technology through research, advocacy, education, and engagement with governments, international organizations, and tech companies. Grounded in Jewish values and community experience, TechRI aims to ensure that Jewish perspectives help shape the digital environment and the future of emerging technologies.