



Digital Services Act Transparency Report

February 2025

Introduction

LinkedIn is a real-identity online service for professionals to connect and interact with other professionals, learn, hire, and find jobs. LinkedIn's vision is to create economic opportunity for every member of the global workforce. Its mission is to connect the world's professionals to make them more productive and successful. As part of that mission, LinkedIn is committed to keeping its platform and services safe, trusted, and professional, and to providing transparency to its members, the public, and to regulators.

LinkedIn Ireland Unlimited Company ("LinkedIn") – the provider of LinkedIn's services in the European Union – has been designated by the European Commission as a Very Large Online Platform (VLOP) and is therefore subject to the European Union's Digital Services Act (DSA) Article 42 requirement to publish certain information in semi-annual disclosures. This DSA Transparency Report is responsive to the obligations under DSA Article 15(1), Article 24(1)-(2), and Article 42(1)-(3). This Report provides information regarding the following topics as they pertain to the European Union:

- Monthly Active Recipients of the Service
- Content Moderation following a User Report
- Content Moderation at LinkedIn's Initiative
- Content Moderation Appeals
- Content Moderation & Automated Systems
- Account Suspensions
- Government Requests
- Out-of-Court Dispute Settlement Bodies

1. Monthly Active Recipients of the Service

LinkedIn provides the information below in response to DSA Articles 24(2) and 42(3).

For the six-month period from 1 July – 31 December 2024 (the "reporting period"), an estimated monthly average of: 52,000,000 logged-in users visited LinkedIn's services in the EU; and 142,500,000 site visits to LinkedIn's services from EU-based users occurred in a logged-out state.

The metrics by Member State are reported below. Metrics are rounded to the nearest one-hundred thousand.

Table 1 – EU monthly active recipients of the service, by Member State

Member State	Monthly average logged-in active users	Monthly average logged-out site visits
EU Overall	52,000,000	142,500,000
Austria	800,000	4,000,000
Belgium	1,800,000	2,900,000
Bulgaria	300,000	1,500,000
Croatia	300,000	1,000,000
Cyprus	100,000	300,000
Czechia	700,000	2,000,000
Denmark	1,500,000	2,100,000
Estonia	100,000	300,000

Finland	800,000	4,600,000
France	11,500,000	20,900,000
Germany	7,000,000	26,100,000
Greece	700,000	2,900,000
Hungary	500,000	2,200,000
Ireland	1,100,000	7,700,000
Italy	6,000,000	13,700,000
Latvia	100,000	600,000
Lithuania	300,000	1,600,000
Luxembourg	200,000	300,000
Malta	100,000	100,000
Netherlands	5,100,000	12,700,000
Poland	2,300,000	6,900,000
Portugal	1,500,000	2,800,000
Romania	1,000,000	3,200,000
Slovakia	200,000	700,000
Slovenia	200,000	400,000
Spain	5,600,000	14,900,000
Sweden	2,100,000	6,000,000

Member State totals may not sum to the EU total because of rounding. Given the manner in which LinkedIn measures guest user traffic, the above logged-out site visit data has not been fully deduplicated.

2. Content Moderation following a User Report

LinkedIn provides the information below in response to DSA Articles 15(1)(b)-(c) and 42(2)(a)-(b).

All content on LinkedIn must comply with LinkedIn's [Professional Community Policies](#), which set out in detail the content LinkedIn permits and does not permit to keep its platform safe, trusted, and professional. In addition to the Professional Community Policies, job posts on LinkedIn must also comply with LinkedIn's [Jobs Policies](#), and ads must comply with LinkedIn's [Advertising Policies](#).

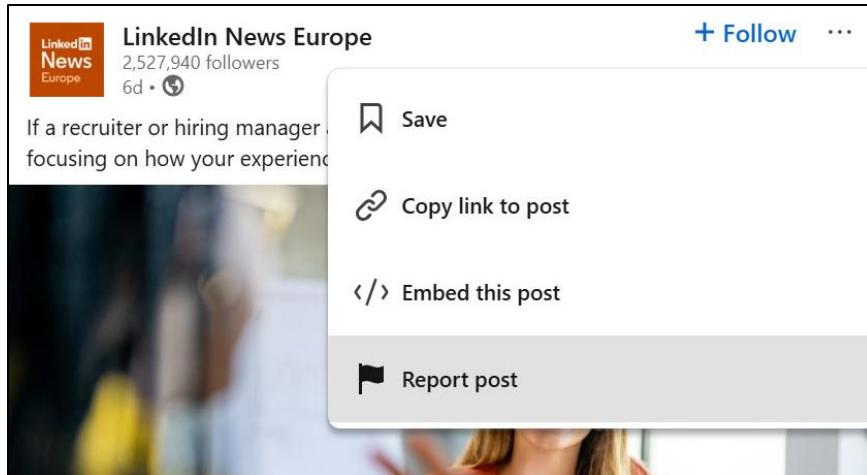
LinkedIn applies a three-layer, multidimensional approach to moderate content on LinkedIn:

- The first layer of protection is automated and proactive prevention. When a member attempts to create a piece of content on LinkedIn, various calls (or signals) are sent to LinkedIn's machine learning services. These services aim to automatically filter out certain policy-violating content at the time of creation.
- The second layer of protection is a combination of automated and human-led detection. LinkedIn's second layer of moderation detects content that is likely to be violative but for which LinkedIn is not sufficiently confident to warrant automatic removal, and sends it for human review.
- The third layer of protection is human-led detection. If users locate content that they believe violates LinkedIn's policies, they are able to report it using LinkedIn's in-product reporting functionality.

User reporting flow

To report content in violation of LinkedIn's policies, members click the three-dot icon available in-product on the content and follow the in-product prompts. For example, to report a post on LinkedIn, members use the following process:

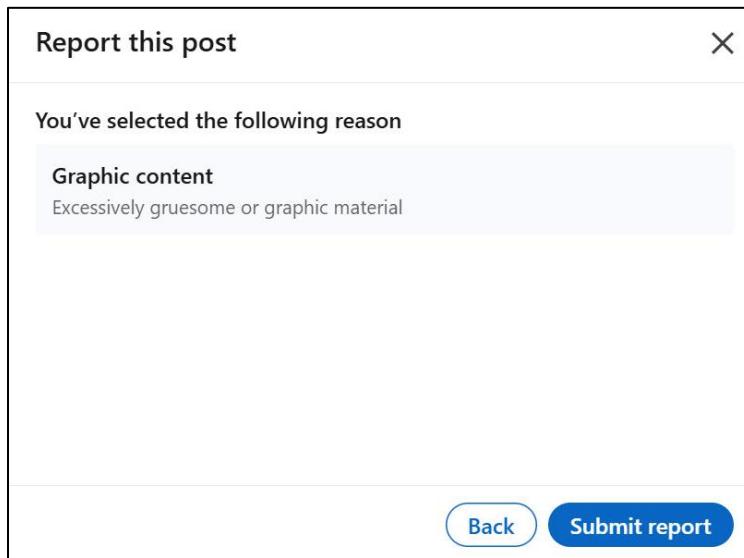
1. Select the three-dot icon in the upper-right corner of the post, and select 'Report post':



2. Select the reporting reason that applies to the post. For example, "Graphic content":

A screenshot of the 'Report this post' modal. The title is 'Report this post' and there is a close button 'X'. The sub-section title is 'Select our policy that applies'. A row of buttons shows various reporting reasons: Harassment, Fraud or scam, Spam, Misinformation, Hateful speech, Threats or violence, Self-harm, Graphic content (which is highlighted in green), Dangerous or extremist organizations, Sexual content, Fake account, Child exploitation, Illegal goods and services, and Infringement. Below this, a question 'Are you reporting illegal content?' is followed by a note: 'Report this content if it's illegal under EU law or the law of an EU member state.' A blue 'Report illegal content' link is shown. At the bottom is a blue 'Next' button.

3. Review the selected reporting reason, and submit the report:



Report this post X

You've selected the following reason

Graphic content
Excessively gruesome or graphic material

Back Submit report

Logged-in and logged-out users are also able to report illegal content pursuant to DSA Article 16 via a dedicated form. When users report content, the reports are sent for review and are resolved by LinkedIn's Content Moderation team, discussed below, or by LinkedIn's automated system, described in Section 5. When users receive notification that their report has been resolved, the notice indicates whether the report was resolved by human review or LinkedIn's automated system.

LinkedIn's Content Moderation team

As of 31 December 2024, LinkedIn had approximately 1,443 content moderators globally and 180 content moderators located in the EU. These personnel review content reported by users, content reported by LinkedIn's systems, and reporter and author appeals, using policies and guidance developed by a policy team and lawyers who are experienced in content moderation and legal issues regarding takedown requests.¹ In addition to content moderators, policy managers, and in-house lawyers, LinkedIn employs a dedicated team of trainers and quality assurance analysts tasked with onboarding new content moderators, training content moderators on new policies and policy changes, and monitoring and improving moderator accuracy and consistency.

LinkedIn's website is currently available in and [supports](#) 15 of the 24 official languages of the EU. Content review is conducted via LinkedIn's custom-built internal review tool, which has built-in translation technology to assist reviewers. For the official languages of the EU, content moderators have the following linguistic expertise (defined as CEFR-B2 language expertise or above):

Table 2 – Linguistic expertise of content moderators

Language	Content moderators
Bulgarian	0
Croatian	2
Czech*	0
Danish*	0
Dutch*	8

¹ LinkedIn did not receive any reports from Trusted flaggers during the reporting period.

English*	1,443
Estonian	0
Finnish*	0
French*	28
German*	19
Greek*	0
Hungarian*	0
Irish	0
Italian*	11
Latvian	0
Lithuanian	0
Maltese	0
Polish*	10
Portuguese*	36
Romanian*	0
Slovak	0
Slovenian	0
Spanish*	29
Swedish*	0

* Denotes language [supported](#) by LinkedIn.

For situations where a content moderator lacks language proficiency and LinkedIn's machine translation tools are insufficient for a review, moderators consult with their team lead and use translation services to complete the review.

LinkedIn has implemented robust training and quality assurance programs for content moderators, including regular audits on sample sets of content reviewed by moderation teams, regular group calibration sessions to address common error trends, and coaching for lower performers. With regard to internal training, LinkedIn utilizes a full-time team of trainers, who not only support the onboarding of new content moderators, but also provide ongoing educational opportunities for all moderators. Content moderators have direct access to the content policy managers through regular office hours and dedicated escalation pathways. For particularly complex decisions, content policy managers also have access to in-house lawyers who can consult country law experts as needed.

Content moderators apply the enforcement actions below to reported content.

Enforcement actions for policy-violating content

During the reporting period, LinkedIn applied three actions to content because it violated LinkedIn's policies:

1. Action 1: LinkedIn removed content that violated its policies;
2. Action 2: LinkedIn limited the visibility of content that violated its policies; and

3. Action 3: LinkedIn applied a sensitive content warning and limited the visibility of content that violated its policies.²

User reporting metrics

The tables below report information regarding the number of EU reports LinkedIn received during the reporting period by user-selected report reason. For each report reason, LinkedIn provides the number of reports received, the pieces of content underlying those reports, the number of reports where LinkedIn determined the content violated its policies, and the number of pieces of content where LinkedIn applied Action 1, Action 2, and Action 3. The tables separately report metrics for job posts and ads content from other content, given different report reasons for job posts and ads.

Table 3(a) – EU reports received during the reporting period, by user-selected report reason (Content)

User-selected report reason	Number of reports	Number of underlying pieces of content	Number of reports where violation was found	Pieces of content where LinkedIn applied Action 1	Pieces of content where LinkedIn applied Action 2	Pieces of content where LinkedIn applied Action 3
Hacked account	709	651	16	15	0	0
Misinformation	206,905	172,704	4,784	2,709	337	537
Hateful speech	152,000	129,898	6,298	4,805	37	388
Threats or violence	19,899	18,321	1,573	969	2	465
Self-harm	2,629	2,506	75	30	1	31
Graphic content	8,973	8,295	1,536	233	3	1,162
Dangerous or extremist organisations	24,594	21,978	681	385	4	212
Sexual content	9,101	7,689	516	432	4	3
Fake account	156,175	124,549	27,971	19,208	6	11
Spam	146,098	132,240	2,170	1,937	21	94
Fraud or scam	72,083	65,777	3,821	3,351	13	27
Illegal goods and services	5,055	4,720	395	377	0	4
Harassment	61,401	56,889	1,956	1,766	7	48
Impersonation	83,904	67,093	9,169	6,470	0	0
Child exploitation	2,562	2,358	109	23	0	83
Infringement or defamation	3,204	1,244	460	407	0	0
Violation of my GDPR rights	3	2	0	0	0	0
Others	5	5	0	0	0	0
Total	955,300	751,039	61,530	39,166	398	2,843

² A sensitive content warning obscures a post until a member clicks to view the post.

Table 3(b) – EU reports received during the reporting period, by user-selected report reason (Job posts)

User-selected report reason	Number of reports	Number of underlying pieces of content	Number of reports where violation was found	Pieces of content where LinkedIn applied Action 1	Pieces of content where LinkedIn applied Action 2	Pieces of content where LinkedIn applied Action 3
Scam, phishing, or malware	16,299	9,760	2,692	1,985	0	0
Promotional or spam	20,538	10,671	3,310	2,450	0	0
Discriminatory, or advocates or supports discrimination	8,012	4,032	1,122	803	0	0
Offensive or harassing	1,044	802	82	67	0	0
Illegal good or service	3,461	1,878	309	277	0	0
Extreme violence or terrorism	278	204	17	16	0	0
Job is closed	18,254	10,931	7,147	5,274	0	0
Job has an incorrect company	2,737	1,239	318	257	0	0
Job has an incorrect location	6,406	3,492	368	253	0	0
Job has incorrect formatting	8,512	4,698	1,142	829	0	0
Job does not belong on LinkedIn	3,213	2,001	509	398	0	0
Something is broken or incorrect	973	882	287	277	0	0
Infringement or defamation	1	1	0	0	0	0
Total	89,728	46,526	17,303	12,507	0	0

Table 3(c) – EU reports received during the reporting period, by user-selected report reason (Ads)

User-selected report reason	Number of reports	Number of underlying pieces of content	Number of reports where violation was found	Pieces of content where LinkedIn applied Action 1	Pieces of content where LinkedIn applied Action 2	Pieces of content where LinkedIn applied Action 3
Misinformation	12,744	9,296	284	206	0	0
Fraud or scam	40,126	13,660	583	360	0	0
Spam	31,311	21,716	524	482	0	0
Fake account	3,982	1,883	123	46	0	0

Bullying or trolling, or sexual harassment	3,842	3,576	77	77	0	0
Hateful or abusive speech	1,941	1,441	39	37	0	0
Inciting violence or is a threat	1,207	967	22	22	0	0
Shocking or gory	1,814	1,630	32	32	0	0
Terrorism or act of extreme violence	1,737	1,497	28	28	0	0
Nudity or sexual content	1,307	1,116	26	26	0	0
Infringement or defamation	1	1	0	0	0	0
Total	100,012	44,426	1,738	1,215	0	0

Reports resolved by automated means

As discussed above, user reports may be resolved either by LinkedIn's Content Moderation team or by LinkedIn's automated system. When users receive notification that their report has been resolved, the notice indicates whether the report was resolved by LinkedIn's automated system or human review.

For the reports in Tables 3(a)-(c) above, LinkedIn estimates the number of reports where the decision on the reported content was made by automated system to be: 467,084 reports.

Median time from report to decision

For the reports in Tables 3(a)-(c) above, the median time from report to decision during the reporting period was approximately: 22 minutes.

LinkedIn excludes from this calculation reports where the decision on the reported content was made by LinkedIn's automated system, as these reports are resolved quickly.

Reports where action was taken on the basis of the law

For the reports in Tables 3(a)-(c) above, LinkedIn estimates the number of user reports where action was taken on the basis of the law to be 305 reports. LinkedIn's policies separately prohibit a wide range of content that also violates the law. In such cases, LinkedIn generally relies on its policies as the basis for action.

Reports submitted by Trusted flaggers

LinkedIn did not receive any reports from Trusted flaggers during the reporting period.

Notes:

1. For the purpose of this report, LinkedIn attributes reports as EU-reports in the tables above based on the IP address of the user on the day the report was submitted. Where IP address isn't available, LinkedIn uses a close-in-time IP address, within 7 days. Where that isn't available, LinkedIn uses the self-declared profile location of the member at the time the metrics for this report are generated.
2. Except where otherwise noted, 'content' addressed in this report includes user-generated content that appears in LinkedIn's Feed – for example, posts, articles, comments, and newsletters – as well as profiles, pages, groups, job posts that appear on LinkedIn's [Jobs Board](#), and ads. In some cases, LinkedIn separately reports job and ad content broken out from other content. For example, in Tables 3(a), 3(b),

and 3(c) above, LinkedIn separately provides user reports for content, jobs, and ads given different reporting reasons.

3. LinkedIn reports the metrics above based on the reporting reason selected by the user. The reporting reason selected by the user when reporting the content may or may not be the same as the policy basis on which LinkedIn actioned the content.
4. 'Underlying pieces of content' reports the number of unique pieces of content for each report reason. A single piece of content may be reported by multiple users for differing report reasons. For this reason, to avoid double counting, the content counts in the Total row may be less than the sum of each report reason.
5. The metrics LinkedIn provides in this report are best estimates provided the data available in LinkedIn's systems and methods used in the ordinary course of business. In some cases, metrics can be impacted by, e.g., account deletion, content deletion, as well as downtime or errors in LinkedIn's systems that may impact data recording. Certain data may also vary or change over time. For example, a user report received on 31 December may not be resolved until after the reporting period. Metrics in the report are based on data after close of the reporting period.

3. Content Moderation at LinkedIn's Initiative

LinkedIn provides the information below in response to DSA Article 15(1)(c). This section reports data regarding content moderation LinkedIn engaged in on its own initiative, absent a user report.

As referenced above, LinkedIn applies a three-layer, multidimensional approach to moderate content on LinkedIn. As part of LinkedIn's proactive moderation, in many cases LinkedIn removes policy-violating content before users encounter the content or submit a user report. LinkedIn's systems may remove policy-violating content or send content for human review. Similarly, LinkedIn investigations may proactively identify policy-violating content absent a user report.

The tables below report information regarding the number of pieces of EU content LinkedIn actioned during the reporting period absent a user report, organized by policy violation. LinkedIn assigns each piece of content a single policy violation. For each category of policy violation, LinkedIn reports the number of pieces of content actioned and whether that content was detected by LinkedIn's automated systems or by manual investigation.

The tables separately report metrics for job posts and ads from other content, given additional policies that apply to job posts and ads. Tables 4(a) and (b) report data regarding content. Table 4(c) reports data regarding job posts. Table 4(d) reports data regarding ads. LinkedIn did not apply Actions 2 or 3 to any job posts or ads during the reporting period.

Table 4(a) – EU content where LinkedIn removed the content (Action 1) during the reporting period absent a user report, by policy violation (Content)

Policy Violation	Number of pieces of content where LinkedIn applied Action 1	Pieces of content detected by LinkedIn automated systems	Pieces of content detected by LinkedIn manual investigation
Hateful Speech	23,709	23,662	47
Adult Nudity and Sexual Activity	3,550	3,548	2
Graphic Content	1,538	1,538	0

Threats and Incitement to Violence	775	766	9
Misinformation	12,100	12,079	21
Spam and Artificial Engagement	30,624	30,618	6
Harassment	5,444	5,396	48
Child Exploitation	309	298	11
Fraud and Deception	20,237	1,201	19,036
Illegal and Regulated Goods and Services	16,941	16,159	782
Infringement and Defamation	6,942	6,798	144
Fake Account	1,451,681	1,217,645	234,036
Dangerous Organisations and Individuals	632	629	3
Profile Policies	924	924	0
Others	1,283	1,278	5
Total	1,576,689	1,322,539	254,150

Table 4(b) – EU content where LinkedIn applied Action 2 or 3 during the reporting period absent a user report, by policy violation (Content)

Policy Violation	Number of pieces of content where LinkedIn applied Action 2 or 3	Pieces of content detected by LinkedIn automated systems	Pieces of content detected by LinkedIn manual investigation
Hateful Speech	0	0	0
Adult Nudity and Sexual Activity	0	0	0
Graphic Content	118,091 ³	118,090	1
Threats and Incitement to Violence	0	0	0
Misinformation	7,223 ⁴	7,210	13
Spam and Artificial Engagement	0	0	0
Harassment	0	0	0
Child Exploitation	0	0	0
Fraud and Deception	0	0	0
Illegal and Regulated Goods and Services	0	0	0
Infringement and Defamation	0	0	0
Fake Account	0	0	0

³ In all cases LinkedIn applied Action 3.

⁴ In all cases LinkedIn applied Action 2.

Dangerous Organisations and Individuals	0	0	0
Profile Policies	0	0	0
Other	0	0	0
Total	125,314	125,300	14

Table 4(c) – EU content where LinkedIn removed the content (Action 1) during the reporting period absent a user report, by policy violation (Job posts)

Policy Violation	Number of job posts where LinkedIn applied Action 1	Job posts detected by LinkedIn automated systems	Job posts detected by LinkedIn manual investigation
Illegal and Regulated Goods and Services	58	53	5
Discrimination	7,057	6,116	941
MLM and franchises	701	507	194
Illegitimate job post	6,600	2,778	3,822
Fraud and deception	3,091	2,898	193
Adult nudity and sexual activity	3	3	0
Threats and Incitement to Violence	0	0	0
Hateful speech	0	0	0
Phishing	0	0	0
Job requirements: Relevant and factual	21,521	15,225	6,296
Job requirements: Professionalism	440	434	6
Others	2	2	0
Total	39,473	28,016	11,457

Table 4(d) – EU content where LinkedIn removed the content (Action 1) during the reporting period absent a user report, by policy violation (Ads)

Policy Violation	Number of ads where LinkedIn applied Action 1	Ads detected by LinkedIn automated systems	Ads detected by LinkedIn manual investigation
Adult nudity and sexual activity	9	9	0
Prohibited affiliate advertising	1	1	0
Advertising unsupported language	31,608	31,505	103
Fraud and deception	1,522	1,480	42
Discrimination	7	7	0
Incomplete advertisement	2,865	2,861	4

Misinformation	0	0	0
Advertising safety and privacy	13,832	13,792	40
Advertising editorial policy	38,305	38,270	35
Illegal and regulated goods and services	7,800	7,749	51
Prohibited dating services	1	1	0
Prohibited political advertising	620	596	24
Advertising offensive to good taste	1	1	0
Restricted solicitation of funds	226	226	0
Infringement	299	298	1
Other	0	0	0
Total	97,096	96,796	300

Notes:

1. For the purpose of this report, LinkedIn attributes content as EU content in Tables 4(a)-(d) based on the IP address of the user on the day the content was created. LinkedIn maintains records of IP address associated with content creation for a limited period of time – as a result, the data in Tables 4(a)-(d) reports content moderation for content created within the last two years.

Within the two-year window, LinkedIn attributes content as EU content based on the IP address of the user on the day the content was created. Where that IP address isn't available, LinkedIn uses a close-in-time IP address, within 7 days. Where that isn't available, LinkedIn uses the self-declared profile location of the member at the time the metrics for this report are generated.

2. That a piece of content was "detected by" LinkedIn's automated systems or by manual investigation refers to the method by which the content was found, not the method by which the content was determined to violate LinkedIn's policies. A piece of content may be detected by LinkedIn's automated systems and sent for human review.

4. Content Moderation Appeals

LinkedIn provides the information below in response to DSA Article 15(1)(d).

When LinkedIn makes an enforcement decision, the reporter and author generally are notified of the decision and given an opportunity to appeal. Notices are typically sent by email and contain a link to a notice page containing additional information (for example, regarding the content at issue, the policy violated, the action LinkedIn has taken, redress information and, in most instances, a link to allow the user to appeal LinkedIn's decision). LinkedIn reviews submitted appeals and notifies the user of its appeal decision.

The table below reports data regarding appeals of the enforcement decisions in Sections 2 and 3 above. The appeals include both appeals from reporters (i.e., when a user's report is rejected) and appeals from authors (i.e., when an author's content is actioned). The table reports the number of appeals received during the reporting

period, the number of appeals granted (i.e., where LinkedIn reversed its decision), and the median time from appeal to appeal decision. Certain appeals may be initiated within the reporting period but not resolved within the reporting period; those appeals are excluded from the median time calculation. The basis for all user appeals is to challenge LinkedIn's decision.

Table 5 – Appeals of the enforcement decisions in Sections 2 and 3

Number of appeals	51,683
Number of appeals granted	17,631
Median time from appeal to appeal decision	6 hours 1 minute

5. Content Moderation & Automated Systems

LinkedIn provides the information below in response to DSA Articles 15(1)(e) and 42(2)(c).

LinkedIn uses two types of automated systems for content moderation relevant to this report:

1. LinkedIn uses an automated system to resolve certain user reports;
2. LinkedIn uses an automated system to identify and remove policy-violating content.

Automated system to resolve user reports

LinkedIn utilizes an automated system to resolve certain user reports and decide whether the report is valid or invalid. The automated system is based in part on past decisions human reviewers have made regarding whether content violates LinkedIn's policies. When users receive notification that their report has been resolved, the notice indicates whether the report was resolved by LinkedIn's automated system or human review.

LinkedIn employs the following safeguards, among others, to this automated system: LinkedIn monitors the aggregate performance and accuracy of the system, and sets minimum thresholds for performance; LinkedIn sets thresholds for individual decisions made by the system, such that the system will not act on a given report and will wait for human review when those thresholds are not met; LinkedIn limits the types of reports the system acts on (e.g., the system will not act on reports of terrorist content); LinkedIn generally allows reporters to appeal a decision if they believe the decision is incorrect; and LinkedIn periodically retrains its system to account for, e.g., changes in human-reviewer decisions, content trends, and user report trends over time.

The table below reports estimated error rates of the automated system globally and by EU language for the reporting period.⁵

Table 6 – Estimated error rate for Automated System 1, by EU language

Language	Estimated Error Rate
Global	<1%
English*	<1%
Bulgarian	<1%
Croatian	<1%
Czech*	<1%

⁵ LinkedIn's website currently [supports](#) 15 of the 24 official languages of the EU, noted with an asterisk.

Danish*	<1%
Dutch*	<1%
Estonian	<1%
Finnish*	<1%
French*	<1%
German*	<1%
Greek	<1%
Hungarian*	<1%
Irish	N/A
Italian*	<1%
Latvian	<1%
Lithuanian	<1%
Maltese	N/A
Polish*	<1%
Portuguese*	<1%
Romanian*	<1%
Slovak	<1%
Slovenian	<1%
Spanish*	<1%
Swedish*	<1%

Automated system to identify and remove policy-violating content

LinkedIn also utilizes an automated system to identify and remove policy-violating content absent a user report. The automated system is based in part on past decisions human reviewers have made regarding whether content violates LinkedIn's policies. When users receive notification that their content has been removed, the notice indicates whether the content was detected and removed as a result of LinkedIn's automated system.

LinkedIn employs the following safeguards, among others, to this automated system: LinkedIn's monitors the aggregate performance and accuracy of the system, and sets minimum thresholds for performance; LinkedIn sets thresholds for individual decisions made by the system, such that the system will not act and will send the content for human review if thresholds are not met; LinkedIn limits the types of violating content the system will act on; LinkedIn generally allows authors to appeal a decision if they believe the decision is incorrect; and LinkedIn regularly retrains its system to account for, e.g., changes in human-reviewer decisions and content trends over time.

The table below reports estimated error rates of the automated system globally and by EU language for the reporting period.⁶

Table 7 – Estimated error rate for Automated System 2, by EU language

Language	Estimated Error Rate
Global	4.2%

⁶ LinkedIn's website currently [supports](#) 15 of the 24 official languages of the EU, noted with an asterisk.

English*	4.1%
Bulgarian	<1%
Croatian	<1%
Czech*	3.0%
Danish*	6.2%
Dutch*	5.9%
Estonian	<1%
Finnish*	7.6%
French*	5.1%
German*	7.2%
Greek	3.5%
Hungarian*	5.7%
Irish	N/A
Italian*	5.7%
Latvian	3.4%
Lithuanian	1.1%
Maltese	N/A
Polish*	4.4%
Portuguese*	6.2%
Romanian*	3.2%
Slovak	<1%
Slovenian	<1%
Spanish*	2.8%
Swedish*	7.3%

Notes:

1. The estimated error rates in Tables 6 and 7 above are based on the number of enforcement decisions made by the automated system that are overturned following appeal (i.e. the automated system made an error). To calculate error rates, LinkedIn takes the number of decisions by the automated system that were overturned divided by the number of appealable decisions made by the automated system during the reporting period.
2. LinkedIn also utilizes an internal system to queue content for human review. LinkedIn doesn't calculate an error rate for this system as it doesn't make moderation decisions or apply enforcement actions to content; whether a piece of content violates LinkedIn's policies is determined by LinkedIn human reviewers.

6. Account Suspensions

LinkedIn provides the information below in response to DSA Article 24(1)(b). This section reports data on the number of suspensions imposed pursuant to DSA Article 23.

Permanent account suspensions due to repeatedly providing policy-violating content

The metric below reports the number of EU accounts LinkedIn permanently suspended during the reporting period due to repeatedly providing policy-violating content, which includes illegal content. In some cases, LinkedIn may permanently suspend an account after a single egregious content policy violation (e.g., in the case of child exploitation material). [Learn more](#). LinkedIn includes such suspensions within this metric.

Accounts are attributed as EU-accounts based on the self-declared profile location for the account. The metric below does not include account suspensions for reasons other than repeatedly providing policy-violating content – for example, account suspension because the account is fake, account suspension for data scraping or automated activity, and so on. Similarly, the metric does not include temporary account suspensions.

LinkedIn estimates the number of EU accounts permanently suspended during the reporting period due to repeatedly providing policy-violating content to be: 7,393 accounts.

Suspension of reporting functionality due to repeatedly submitting manifestly unfounded reports

LinkedIn did not suspend the reporting functionality for any EU accounts during the reporting period.

Suspension of appeal functionality due to repeatedly submitting manifestly unfounded appeals

LinkedIn did not suspend the appeal functionality for any EU accounts during the reporting period.

7. Government Requests

LinkedIn provides the information below in response to DSA Article 15(1)(a).

This section reports data on requests from Member State government authorities: (1) to remove content and (2) to provide user account information. LinkedIn carefully considers all government requests for content removal and account information, and works to mitigate any implications they may have on freedom of expression and human rights. For government demands, LinkedIn employs safeguards to ensure any actions taken are narrow, specific, submitted in writing, and based on valid legal orders. Through its parent company, Microsoft, LinkedIn also engages with broader civil society organizations on best practices related to government requests and participates in human rights impact assessments.

Government requests to remove content

The table below reports information regarding the number of requests LinkedIn received from Member State government authorities to remove content during the reporting period, organized by Member State and by illegal content type. Government requests to remove content include requests reporting violations of our terms of service or violations of local law.

Table 8(a) – Government requests to remove content, by Member State

Member State	Government requests received	Government requests where at least some action was taken
Austria	1	1
Belgium	0	0
Bulgaria	0	0
Croatia	0	0
Cyprus	0	0
Czechia	0	0

Denmark	1	0
Estonia	0	0
Finland	0	0
France	0	0
Germany	1	0
Greece	0	0
Hungary	0	0
Ireland	0	0
Italy	0	0
Latvia	0	0
Lithuania	0	0
Luxembourg	0	0
Malta	0	0
Netherlands	2	2
Poland	0	0
Portugal	0	0
Romania	0	0
Slovakia	0	0
Slovenia	0	0
Spain	1	1
Sweden	0	0
Total	6	4

Table 8(b) – Government requests to remove content, by illegal content type

Illegal content type	Government requests received
Animal welfare	0
Consumer rights violations	0
Data protection and privacy violations	0
Illegal or harmful speech	1
Intellectual property infringements	0
Negative effects on civic discourse of elections	0
Non-consensual behavior	0
Pornography or sexualized content	0
Protection of minors	0

Risk for public security	2
Scams and/or fraud	2
Self-harm	0
Unsafe and/or illegal products	0
Violence	1
Total	6

LinkedIn estimates the median time to confirm receipt of the requests in Table 8(a)-(b) above to be: less than 1 hour.

LinkedIn estimates the median time to give effect to the requests in Table 8(a)-(b) above to be: 72 hours.⁷ Certain requests may be received within the reporting period but not confirmed or resolved within the reporting period; those requests are excluded from the median time calculations.

Government requests to provide account information

The table below reports information regarding the number of requests LinkedIn received from Member State government authorities to provide account information during the reporting period.

Table 9 – Government requests to provide account information, by Member State

Member State	Government requests received	Government requests where at least some information was provided
Austria	1	0
Belgium	6	5
Bulgaria	0	0
Croatia	0	0
Cyprus	0	0
Czechia	0	0
Denmark	1	1
Estonia	0	0
Finland	1	1
France	464	167
Germany	169	102
Greece	3	0
Hungary	1	1
Ireland	9	4
Italy	7	4
Latvia	0	0
Lithuania	0	0
Luxembourg	0	0
Malta	2	0

⁷ Requests may concern, for example, Article 9 orders, as well as less time-sensitive requests outside the Article 9 process. LinkedIn appropriately resolves requests given the nature of the request.

Netherlands	5	4
Poland	19	9
Portugal	3	1
Romania	2	0
Slovakia	0	0
Slovenia	0	0
Spain	21	14
Sweden	1	0
Total	715	313

LinkedIn automatically sends a confirmation email on receipt of a government request to provide account information. As a result, the median time to confirm receipt of the requests in Table 9 was less than 1 hour.

LinkedIn estimates the median time to give effect to the requests in Table 9 above to be: 32 hours. Certain requests may be received within the reporting period but not confirmed or resolved within the reporting period; those requests are excluded from the median time calculations.

8. Out-of-Court Dispute Settlement Bodies

LinkedIn provides the information below in response to DSA Article 24(1)(a). This section reports data regarding DSA Article 21 out-of-court dispute settlement body disputes initiated or resolved during the reporting period.

Table 10 – Out-of-court dispute settlement body disputes

Number of disputes submitted to Out-of-court Dispute Settlement Bodies during the reporting period	23
Dispute Settlement Body decisions during the reporting period upholding LinkedIn's decision	0
Dispute Settlement Body decisions during the reporting period partially reversing LinkedIn's decision	0
Dispute Settlement Body decisions during the reporting period reversing LinkedIn's decision	1
Median time from the filing of the dispute to Dispute Settlement Body decision	46 days
Percentage of outcomes implemented	0 out of 1
Decisions omitted (e.g. because a dispute was withdrawn or dismissed without decision)	2

Notes:

1. The number of disputes submitted to out-of-court dispute settlement bodies during the reporting period may not equal the number of dispute settlement body decisions during the reporting period. For example, a dispute may be initiated during the reporting period, but not yet resolved by the dispute settlement body during that reporting period.
2. 'Percentage of outcomes implemented' reports the number of adverse dispute settlement body decisions (i.e., where the dispute settlement body decision was to reverse or partially reverse LinkedIn's decision) that LinkedIn implemented during the reporting period.

3. 'Median time from the filing of the dispute to Dispute Settlement Body decision' reports the median time from filing of the dispute with the Dispute Settlement Body to the decision of the Dispute Settlement Body. 'Decisions omitted (e.g. because a dispute was withdrawn or dismissed without decision)' are excluded from the median time calculation.