



Digital Services Act Transparency Report

October 2024

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 Settlement Body Disputes

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 January – 30 June 2024 (the “reporting period”), an estimated monthly average of: 51,900,000 logged-in users visited LinkedIn’s services in the EU; and 192,900,000 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	51,900,000	192,900,000
Austria	800,000	4,800,000
Belgium	1,800,000	4,300,000
Bulgaria	300,000	1,700,000
Croatia	300,000	1,400,000
Cyprus	100,000	500,000
Czechia	700,000	2,500,000
Denmark	1,500,000	2,900,000
Estonia	100,000	700,000

Finland	800,000	6,100,000
France	11,400,000	32,500,000
Germany	6,800,000	33,800,000
Greece	700,000	3,100,000
Hungary	500,000	2,900,000
Ireland	1,100,000	6,100,000
Italy	6,100,000	21,600,000
Latvia	100,000	700,000
Lithuania	200,000	1,500,000
Luxembourg	200,000	600,000
Malta	100,000	300,000
Netherlands	5,100,000	18,900,000
Poland	2,300,000	8,600,000
Portugal	1,600,000	3,900,000
Romania	1,000,000	4,300,000
Slovakia	200,000	1,000,000
Slovenia	200,000	800,000
Spain	5,700,000	21,400,000
Sweden	2,100,000	5,900,000

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](#).

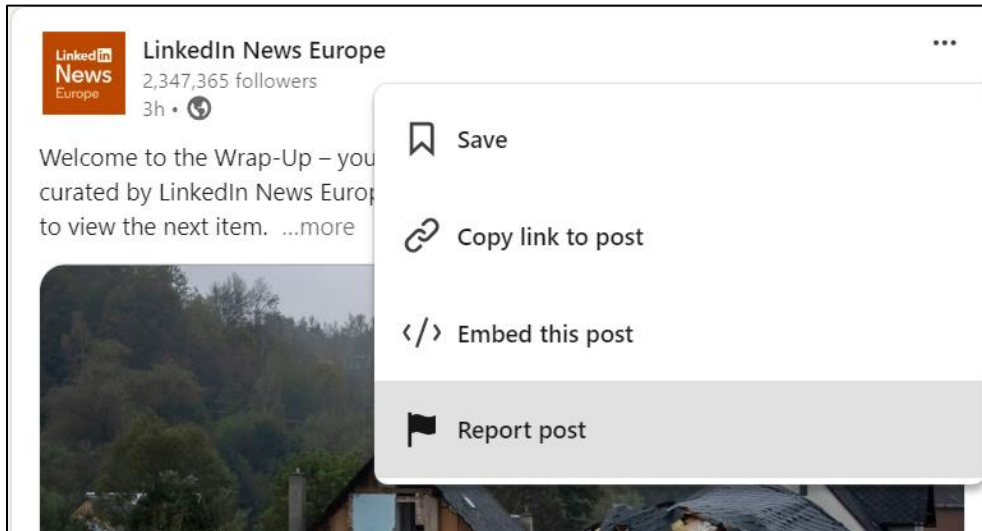
LinkedIn applies a three-layer, multidimensional approach to moderate content on LinkedIn:

- 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.

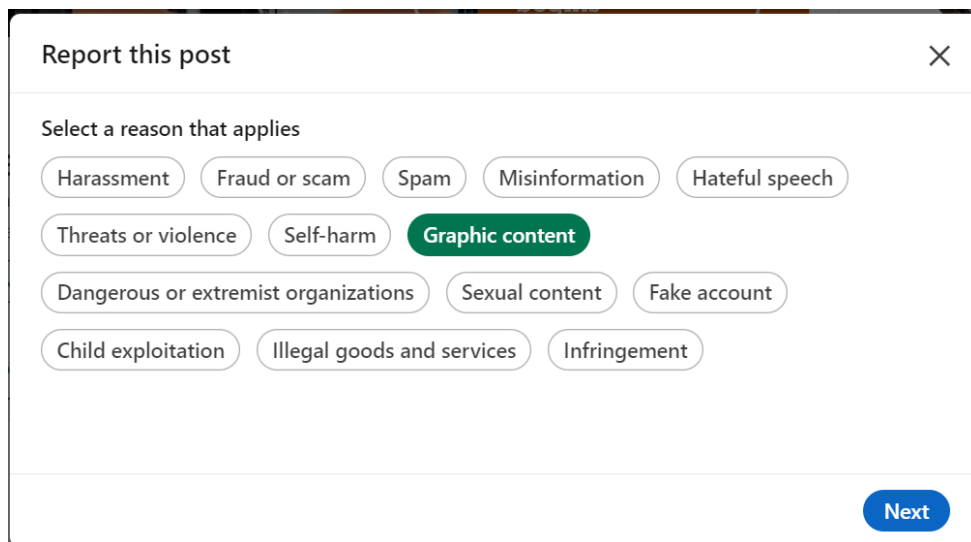
User reporting flow

To report content, 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:

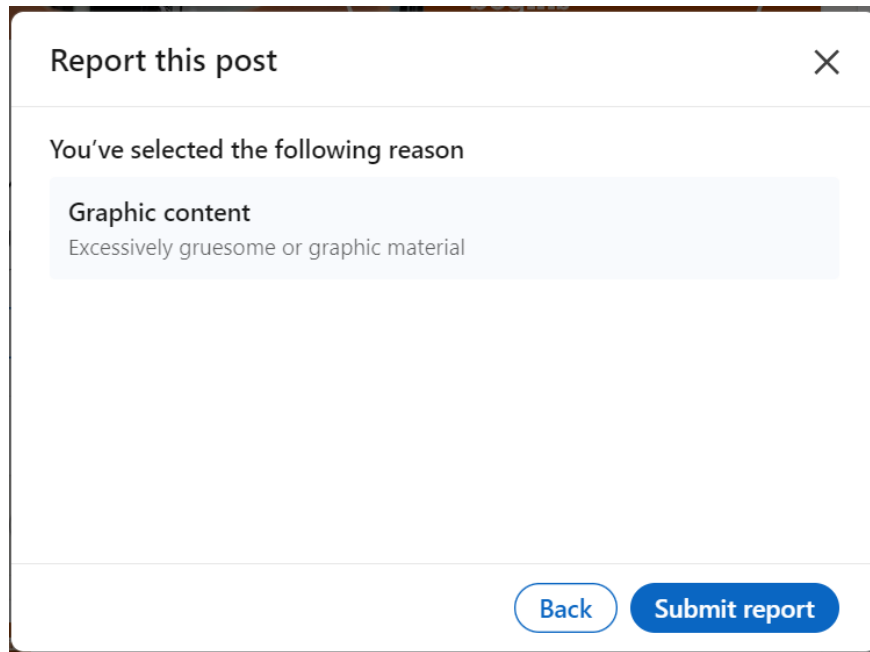
1. Select the three-dot icon in the upper-right corner of the post, and select 'Report post':



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



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



Logged-out users are also able to report content visible to them using our guest reporting functionality.

When users report content, those reports are sent for review and are resolved either 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 30 June 2024, LinkedIn had approximately 1,510 content moderators globally and 190 content moderators located in the EU.¹ 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.² 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 intermediate language expertise or above):

¹ As of this report, the reported metrics, including moderators, also include ad content and moderators.

² LinkedIn did not receive any reports from Trusted flaggers during the reporting period.

Table 2 – Linguistic expertise of content moderators

Language	Content moderators
Bulgarian	0
Croatian	3
Czech*	0
Danish*	0
Dutch*	8
English*	1,507
Estonian	0
Finnish*	1
French*	32
German*	22
Greek*	0
Hungarian*	0
Irish	0
Italian*	16
Latvian	0
Lithuanian	2
Maltese	0
Polish*	10
Portuguese*	38
Romanian*	0
Slovak	0
Slovenian	0
Spanish*	36
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.³

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	607	534	3	3	0	0
Misinformation	191,310	163,569	4,931	2,972	247	764
Hateful speech	144,500	124,597	6,893	5,422	41	653
Threats or violence	19,473	17,744	1,948	1,067	2	754
Self-harm	2,393	2,251	93	54	0	37
Graphic content	9,264	8,391	1,894	275	3	1,268
Dangerous or extremist organisations	23,883	21,054	801	415	9	298
Sexual content	10,042	8,915	692	584	1	5
Fake account	142,733	118,894	25,997	20,357	2	10
Spam	135,373	122,641	2,007	1,797	21	89
Fraud or scam	65,408	59,435	4,589	4,188	9	42
Illegal goods and services	5,429	4,997	754	665	0	4
Harassment	54,112	49,941	2,333	2,019	9	99
Impersonation	82,471	63,181	9,885	6,574	0	0
Child exploitation	2,049	1,969	138	27	0	108
Infringement or defamation	5,386	2,143	34	30	0	1
Total	894,433	709,091	62,992	42,412	315	3,716

³ 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	16,625	9,824	2,767	2,016	0	0
Promotional or spam	20,603	10,719	3,311	2,458	0	0
Discriminatory, or advocates or supports discrimination	8,039	4,058	1,126	806	0	0
Offensive or harassing	1,048	805	85	70	0	0
Illegal good or service	3,477	1,892	302	272	0	0
Extreme violence or terrorism	280	206	17	16	0	0
Job is closed	18,293	10,967	7,131	5,267	0	0
Job has an incorrect company	2,740	1,242	320	261	0	0
Job has an incorrect location	6,424	3,505	367	253	0	0
Job has incorrect formatting	8,542	4,723	1,151	839	0	0
Job does not belong on LinkedIn	3,229	2,010	508	398	0	0
Something is broken or incorrect	973	882	289	277	0	0
Infringement or defamation	1	1	0	0	0	0
Total	90,274	46,735	17,374	12,534	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	12,175	9,142	215	208	0	0
Fraud or scam	18,813	15,563	492	467	0	0
Spam	26,093	21,368	502	475	0	0
Fake account	1,568	1,325	31	30	0	0
Bullying or trolling, or Sexual harassment	5,170	4,792	108	105	0	0
Hateful or abusive speech	2,259	1,746	39	38	0	0
Inciting violence or is a threat	1,485	1,297	27	27	0	0

Shocking or gory	1,867	1,771	47	47	0	0
Terrorism or act of extreme violence	1,968	1,807	59	58	0	0
Nudity or sexual content	2,525	2,473	79	79	0	0
Total	73,923	49,995	1,599	1,486	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: 481,674 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: 12 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 390 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 30 June 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	21,615	21,527	88
Adult Nudity and Sexual Activity	2,027	2,003	24
Graphic Content	1,604	1,600	4
Threats and Incitement to Violence	432	424	8
Misinformation	9,166	9,153	13
Spam and Artificial Engagement	42,490	42,296	194
Harassment	5,390	4,908	482

Child Exploitation	291	275	16
Fraud and Deception	2,093	1,297	796
Illegal and Regulated Goods and Services	15,516	15,463	53
Infringement and Defamation	7,131	6,959	172
Fake Account	1,753,551	1,139,002	614,549
Dangerous Organisations and Individuals	259	258	1
Profile Policies	873	873	0
Other	2,695	2,651	44
Total	1,865,133	1,248,689	616,444

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	47,799 ⁴	47,772	27
Threats and Incitement to Violence	0	0	0
Misinformation	4,542 ⁵	4,541	1
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
Dangerous Organisations and Individuals	0	0	0
Profile Policies	0	0	0
Other	0	0	0
Total	52,341	52,313	28

⁴ In all cases LinkedIn applied Action 3.

⁵ In all cases LinkedIn applied Action 2.

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	35	34	1
Discrimination	9,009	8,811	198
MLM and franchises	685	628	57
Illegitimate job post	3,371	3,331	40
Fraud and deception	1,904	1,858	46
Adult nudity and sexual activity	4	4	0
Threats and Incitement to Violence	2	2	0
Hateful speech	1	1	0
Phishing	2	2	0
Job requirements: Relevant and factual	17,723	17,511	212
Job requirements: Professionalism	688	687	1
Total	33,424	32,869	555

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	8	8	0
Prohibited affiliate advertising	7	6	1
Advertising unsupported language	19,125	18,958	167
Fraud and deception	1,551	1,411	140
Discrimination	2	2	0
Incomplete advertisement	4,473	4,453	20
Misinformation	0	0	0
Advertising safety and privacy	16,403	16,262	141
Advertising editorial policy	36,824	36,693	131
Illegal and regulated goods and services	8,520	8,301	219

Prohibited dating services	8	8	0
Prohibited political advertising	101	100	1
Advertising offensive to good taste	2	2	0
Restricted solicitation of funds	116	112	4
Infringement	714	704	10
Other	1	1	0
Total	87,855	87,021	834

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	38,509
Number of appeals granted	11,247
Median time from appeal to appeal decision	3 hours 13 minutes

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’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 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.⁶

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%
Danish*	<1%
Dutch*	<1%

⁶ LinkedIn’s website currently [supports](#) 15 of the 24 official languages of the EU, noted with an asterisk. LinkedIn’s automated system did not act on a sufficient amount of content in Irish or Maltese during the reporting period to calculate an error rate.

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.⁷

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

Language	Estimated Error Rate
Global	4.2%
English*	4.0%

⁷ LinkedIn’s website currently [supports](#) 15 of the 24 official languages of the EU, noted with an asterisk. LinkedIn’s automated system did not act on a sufficient amount of content in Irish or Maltese during the reporting period to calculate an error rate.

Bulgarian	<1%
Croatian	<1%
Czech*	1.8%
Danish*	2.0%
Dutch*	3.4%
Estonian	<1%
Finnish*	7.6%
French*	5.1%
German*	5.0%
Greek	1.8%
Hungarian*	1.8%
Irish	N/A
Italian*	5.8%
Latvian	<1%
Lithuanian	<1%
Maltese	N/A
Polish*	3.9%
Portuguese*	5.4%
Romanian*	1.2%
Slovak	1.2%
Slovenian	<1%
Spanish*	4.8%
Swedish*	6.4%

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: 3,114 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	0	0
Belgium	0	0
Bulgaria	0	0
Croatia	0	0
Cyprus	0	0
Czechia	0	0
Denmark	0	0

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

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	0
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	0
Scams and/or fraud	1
Self-harm	0
Unsafe and/or illegal products	0
Violence	1
Total	2

LinkedIn estimates the median time to give effect to the requests in Table 8(a)-(b) above to be: 24 hours. LinkedIn did not provide separate confirmation for the requests apart from giving effect to the requests.

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	5	1
Belgium	8	8
Bulgaria	0	0
Croatia	0	0
Cyprus	0	0
Czechia	0	0
Denmark	0	0
Estonia	0	0
Finland	0	0
France	520	156
Germany	126	81
Greece	1	0
Hungary	0	0
Ireland	7	3
Italy	11	8
Latvia	0	0
Lithuania	2	2
Luxembourg	0	0
Malta	1	0
Netherlands	12	12
Poland	15	7
Portugal	3	3
Romania	0	0
Slovakia	0	0
Slovenia	0	0

Spain	21	15
Sweden	1	0
Total	733	296

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: 42 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 Settlement Body Disputes

LinkedIn provides the information below in response to DSA Article 24(1)(a).

LinkedIn did not receive notice of any disputes submitted to out-of-court dispute settlement bodies during the reporting period.