



Digital Services Act
Transparency Report
October 2023

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 Commission’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 ending on 30 June 2023, an estimated monthly average of 45,200,000 logged-in users visited LinkedIn’s services in the EU; and 132,500,000 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	45,200,000	132,500,000
Austria	700,000	3,700,000
Belgium	1,600,000	3,400,000
Bulgaria	300,000	900,000
Croatia	200,000	1,100,000
Cyprus	100,000	400,000
Czechia	600,000	1,900,000
Denmark	1,400,000	2,500,000
Estonia	100,000	700,000

Finland	700,000	3,800,000
France	9,900,000	23,700,000
Germany	5,700,000	26,300,000
Greece	600,000	1,800,000
Hungary	400,000	1,500,000
Ireland	1,000,000	3,300,000
Italy	5,200,000	12,600,000
Latvia	100,000	600,000
Lithuania	200,000	1,000,000
Luxembourg	200,000	700,000
Malta	100,000	300,000
Netherlands	4,700,000	12,100,000
Poland	2,000,000	6,100,000
Portugal	1,300,000	2,900,000
Romania	900,000	4,600,000
Slovakia	200,000	800,000
Slovenia	100,000	600,000
Spain	4,900,000	11,600,000
Sweden	1,900,000	4,000,000

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](#).

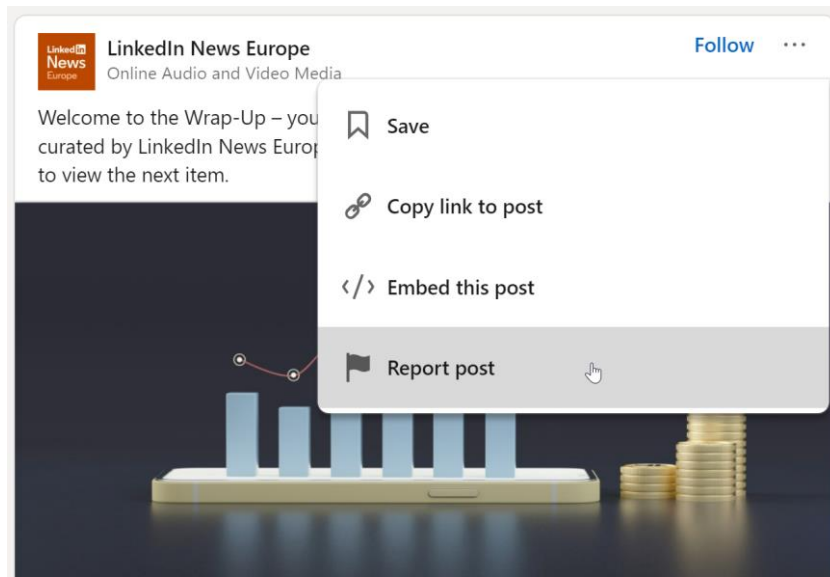
LinkedIn applies a three-layer, multidimensional approach to moderate content on LinkedIn:

- 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’s 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.

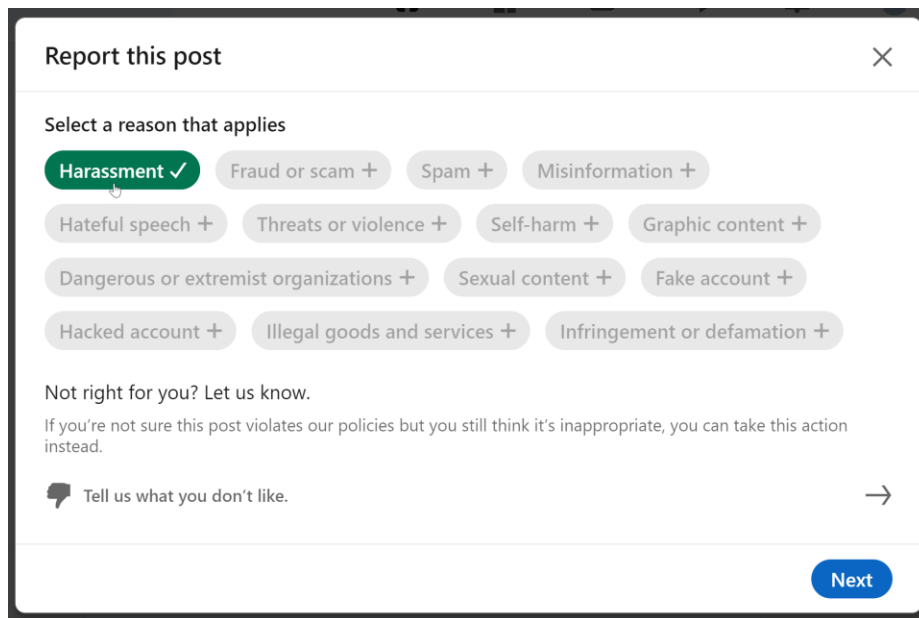
User reporting flow

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:

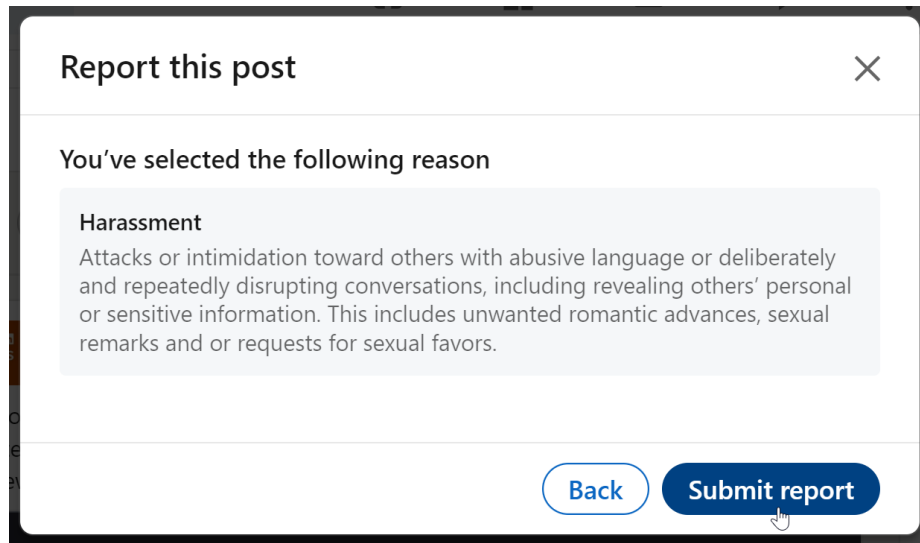
1. Select the three-dot icon in the upper-right corner of the post, and select 'Report post':



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



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



Logged-out users are also able to report content visible to them using 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 25 August 2023, LinkedIn had approximately 820 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](#) 12 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 LinkedIn's website supports, content moderators have the following linguistic expertise:

¹ Trusted flaggers have yet to be designated under Article 22 of the DSA. LinkedIn did not receive any reports from Trusted flaggers during the reporting period.

Table 2 – Linguistic expertise of content moderators

Language	Content moderators with professional working proficiency or above
Czech	0
Danish	0
Dutch	9
French	30
German	20
Italian	13
Polish	10
Portuguese	33
Romanian	0
Spanish	31
Swedish	0

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, 25 August – 30 September 2023 (“reporting period”), LinkedIn applied two actions to content because it violated LinkedIn’s policies:

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

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, the number of pieces of content where LinkedIn removed the content (Action 1), and the number of pieces of content where LinkedIn limited the visibility of the content because it violated its policies

(Action 2). The tables separately report metrics for job posts from other content, given different report reasons for job posts.

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
Hacked Account	195	159	21	5	0
Misinformation	21,669	18,647	495	457	2
Hateful speech	14,126	12,400	682	646	1
Threats or violence	2,041	1,833	174	171	2
Self-harm	418	377	23	22	1
Graphic content	1,061	993	76	74	3
Dangerous or extremist organisations	1,742	1,448	31	30	0
Sexual content	1,975	1,634	254	222	0
Fake account	2,588	2,427	156	135	0
Spam, Fraud or scam, Illegal goods and services	33,407	29,180	1,409	1,221	2
Harassment	6,690	6,282	232	231	2
Impersonation	231	210	9	8	0
Infringement or defamation	471	545	185	246	0
Total	86,640	71,606	3,747	3,285	12

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
Scam, phishing, or malware	1,972	1,548	566	478	0
Promotional or spam	1,915	1,759	581	549	0
Discriminatory, or advocates or supports discrimination	672	575	193	151	0
Offensive or harassing	170	154	28	23	0
Extreme violence or terrorism	29	27	3	3	0
Job is closed	2,226	1,986	1,021	942	0

Job has an incorrect company	319	284	79	72	0
Job has an incorrect location	763	640	53	45	0
Job has incorrect formatting	899	863	177	171	0
Job does not belong on LinkedIn	397	368	86	83	0
Total	9,362	7,680	2,787	2,427	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)-(b) above, LinkedIn estimates the number of reports where the decision on the reported content was made by automated system to be: 28,179 reports.

Median time from report to decision

For the reports in Tables 3(a)-(b) above, the median time from report to decision during the reporting period was approximately: 23 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. The median time reported above thus errs on the side of overstating the median time experienced by many members for resolution of their reports.

Reports where action was taken on the basis of the law

For the reports in Tables 3(a)-(b) above, LinkedIn estimates the number of user reports where action was taken on the basis of the law to be 72 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

Trusted flaggers have yet to be designated under Article 22 of the DSA. 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, pages, groups, comments, newsletters, etc. – as well as job posts that appear on LinkedIn’s [Jobs Board](#). The metrics do not include, e.g., messages, accounts/profiles, or ads. In some cases, LinkedIn separately reports jobs content broken out from other

content. For example, in Tables 3(a) and 3(b) above, LinkedIn separately provides user reports for content and user reports for jobs content 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 September may not be resolved until after the reporting period. Metrics in the report are based on data as of the 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 from other content, given additional policies that apply to job posts. Tables 4(a) and (b) report the number of pieces of EU-content where LinkedIn removed the content (Action 1). Table 4(c) reports the number of pieces of EU-content where LinkedIn limited the visibility of the content because it violated its policies (Action 2). LinkedIn did not apply Action 2 to any job posts 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	1,125	1,122	3
Adult Nudity and Sexual Activity	440	440	0

Graphic Content	1,234	1,234	0
Threats and Incitement to Violence	119	117	2
Misinformation	5,574	5,556	18
Spam and Artificial Engagement	4,400	4,400	0
Harassment	333	329	4
Child Exploitation	24	17	7
Fraud and Deception	325	314	11
Illegal and Regulated Goods and Services	2,518	2,517	1
Infringement and Defamation	18,954	18,920	34
Dangerous Organisations and Individuals	0	0	0
Total	35,046	34,966	80

Table 4(b) – 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	29	29	0
Discrimination	1,647	1,626	21
MLM and franchises	410	405	5
Illegitimate job post	1,268	1,256	12
Fraud and deception	1,234	1,222	12
Phishing	1	1	0
Job requirements: Relevant and factual	5,077	5,004	73
Job requirements: Professionalism	134	132	2
Total	9,800	9,675	125

Table 4(c) – EU-content where LinkedIn limited the visibility of the content because it violated its policies (Action 2) during the reporting period absent a user report, by policy violation (Content)

Policy Violation	Number of pieces of content where LinkedIn applied Action 2	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	960	960	0

Threats and Incitement to Violence	0	0	0
Misinformation	0	0	0
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
Dangerous Organisations and Individuals	0	0	0
Total	960	960	0

Notes:

1. For the purpose of this report, LinkedIn attributes content as EU-content in Tables 4(a)-(c) 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)-(c) 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.

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	2,106
Number of appeals granted	232
Median time from appeal to appeal decision	49 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-supported language for the reporting period.² The automated system actioned fewer than 100 reports during the reporting period for content in Czech. Because of the low sample size, LinkedIn is unable to calculate a meaningful estimate for the reporting period for Czech, and has labeled the language N/A below.

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

Language	Estimated Error Rate
Global	0.7-3.9%
English	0.7-3.9%
Czech	N/A
Danish	0.0-1.0%
Dutch	0.0-1.8%
French	0.0-1.8%
German	1.1-4.5%

² LinkedIn’s website currently [supports](#) 12 of the 24 official languages of the EU.

Italian	0.0-1.8%
Polish	0.0-1.5%
Portuguese	0.1-2.5%
Romanian	0.0-2.3%
Spanish	0.1-2.5%
Swedish	0.0-2.5%

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-supported language for the reporting period. The automated system actioned fewer than 100 pieces of content in the following languages: Czech and Danish. Because of the low sample size, LinkedIn is unable to calculate a meaningful estimate for these languages for the reporting period, and has labeled the languages N/A below.

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

Language	Estimated Error Rate
Global	8.0-14.4%
English	7.4-13.7%
Czech	N/A
Danish	N/A
Dutch	15.1-22.9%
French	33.4-43.0%
German	26.4-35.6%
Italian	28.2-37.5%
Polish	28.8-45.0%
Portuguese	44.8-54.7%
Romanian	16.6-27.0%
Swedish	36.3-49.4%
Spanish	76.1-84.0%

Estimated error rates may appear elevated for certain languages. This is influenced in part by the low number of pieces of content actioned by the automated system in those languages during the reporting period. For example, the automated system actioned fewer than 7,000 pieces of content in each of Spanish, Portuguese, French, Italian, German, Romanian, Dutch, Swedish, and Polish during the reporting period. Another potential indicator of accuracy is the appeal rate of enforcement decisions. As noted above in Sections 2 and 3, LinkedIn made over 130,000 adverse enforcement decisions during the reporting period, and received approximately 2,106 appeals.

Notes:

1. To calculate the estimated error rates in Tables 6 and 7 above, LinkedIn takes a random sample of the content actioned by the automated system during the reporting period, sends that content for review by human content moderators, and compares the decision by the human moderator to the decision by the automated system. The estimated error rates are calculated at the 95% confidence threshold. Human evaluation of content may itself be imperfect, impacting the error rate estimates above.
2. The error rates above exclude from their calculation instances where content was automatically removed because it linked to an external fraudulent URL (for example, URLs used for phishing and malware). These URLs are often dynamic and vary over time to mask the fraud, making error rate calculation via subsequent human review unreliable.
3. 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: 2,047 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. LinkedIn did not receive any government requests to remove content during the reporting period.

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	0	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	0	0

Poland	0	0
Portugal	0	0
Romania	0	0
Slovakia	0	0
Slovenia	0	0
Spain	0	0
Sweden	0	0
Total	0	0

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

Illegal content type	Government requests received
Animal welfare	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	0
Self-harm	0
Unsafe and/or illegal products	0
Violence	0
Total	0

Because LinkedIn did not receive any requests, there is no median time to confirm receipt of the requests or give 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	0	0
Belgium	1	1
Bulgaria	0	0
Croatia	1	0
Cyprus	0	0
Czechia	0	0
Denmark	0	0
Estonia	0	0
Finland	1	0
France	43	28
Germany	19	16
Greece	2	1
Hungary	1	1
Ireland	0	0
Italy	2	2
Latvia	0	0
Lithuania	0	0
Luxembourg	0	0
Malta	0	0
Netherlands	2	2
Poland	2	2
Portugal	1	1
Romania	1	0
Slovakia	0	0
Slovenia	0	0
Spain	2	1
Sweden	0	0
Total	78	55

LinkedIn estimates the median time to confirm receipt of the requests in Table 9 above to be: 4 hours.

LinkedIn estimates the median time to give effect to the requests in Table 9 above to be: 29 hours 30 minutes. 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).

Out-of-court dispute settlement bodies under Article 22 of the DSA are yet to be established. LinkedIn did not receive notice of any disputes submitted to out-of-court dispute settlement bodies during the reporting period.