



LinkedIn AI-Assisted Messages

Data Security & Compliance

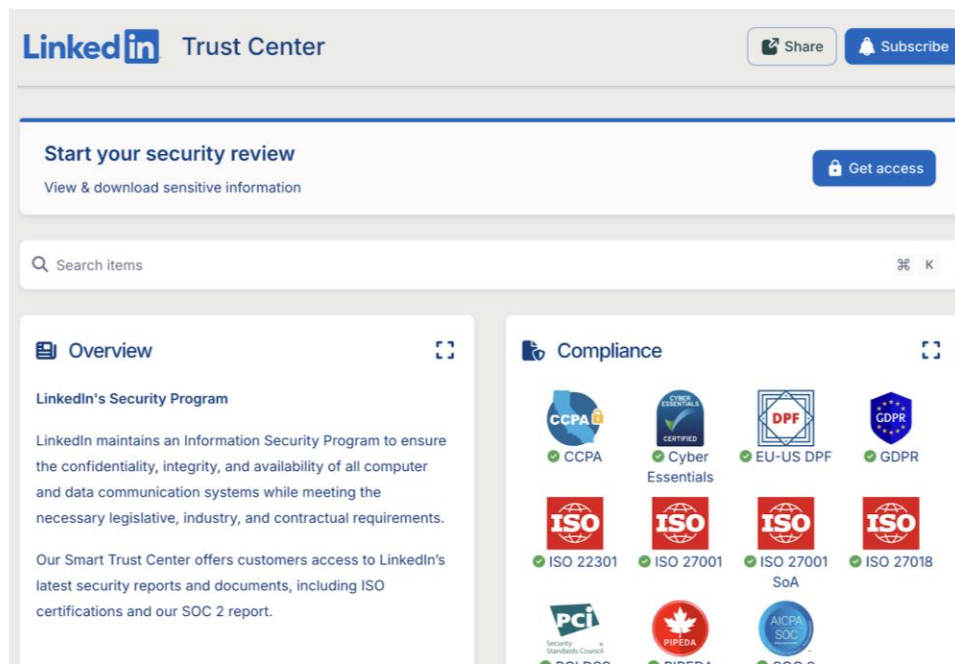
Introduction

Our core value of putting our members first powers all the decisions we make, including how we manage and protect the data of our members and customers. We remain committed to this value as we leverage AI in the development of our services.

As part of its Information Security program, LinkedIn follows industry standards as well as its own best practices to stay ahead of the increasing number of threats facing all Internet services and infrastructure. Our platform incorporates multi-layered security controls which have been independently validated against frameworks such as ISO 27001 and SOC 2. LinkedIn's services are designed to be used in compliance with GDPR, and our Data Processing Agreement sets out our obligations to comply with applicable data protection laws around the world.

Through the provision of transparent and accessible information on our security, privacy and compliance practices, we aim to empower our customers to use our services with confidence.

LinkedIn Trust Center



Visit LinkedIn's Trust Center for a suite of documentation on data governance and compliance assessment.

What is AI-Assisted Messages

LinkedIn's AI-Assisted Messages in LinkedIn Recruiter is a generative AI (GAI) feature that leverages candidate, company, and job

information to help recruiters draft unique, fully personalized messages to candidates—saving time, increasing candidate engagement and personalizing outreach at scale.





Responsible AI principles

LinkedIn was founded with a clear vision to create economic opportunity for every member of the global workforce. In 2025, we are seeing transformative advances in AI that have the potential to help us accelerate our progress toward that vision.

AI is [not new](#) to LinkedIn. LinkedIn has long used AI to enhance our members' professional experiences. By leveraging the power of AI, we help our members connect, increase productivity and achieve success in their careers.

While AI has enormous potential to expand access to opportunity and ultimately transform the world of work in positive ways, the stakes are high. The use of AI comes with risks and potential for harm. That's why, consistent with our commitment to build a [trustworthy](#) platform, we must continue to use AI responsibly. Inspired by, and aligned with, [Microsoft's leadership](#) in Responsible AI, we are sharing the Responsible AI Principles that we use at LinkedIn to guide our work:

- **Advance Economic Opportunity:** People are at the center of what we do. AI is a tool to further our vision, empowering our members and augmenting their success and productivity.
- **Uphold Trust:** Our commitments to privacy, security and safety guide our use of AI. We take meaningful steps to reduce the potential risks of AI.
- **Promote Fairness and Inclusion:** We work to ensure that our use of AI benefits all members fairly, without causing or amplifying unfair bias.
- **Provide Transparency:** Understanding of AI starts with transparency. We seek to explain in clear and simple ways how our use of AI impacts people.
- **Embrace Accountability:** We deploy robust AI governance, including assessing and addressing potential harms and fitness for purpose, and ensuring human oversight and accountability. We are committed to learning from, and helping, others as AI best practices, norms and laws evolve.

Underlying these principles is our commitment to listen and learn about how AI can continue to be a tool to accelerate progress towards economic opportunity for all.



Development of AI at LinkedIn

In building and approving AI features, we have a design review process to ensure that product design aligns with LinkedIn's Responsible AI Principles and privacy and security standards. This includes steps such as: evaluation of data privacy controls, review of AI disclosures, measurement of AI output quality, integration of output moderation, and implementation of quality and trust measurement systems.

LinkedIn is committed to ensuring all of the products we provide can be used in compliance with applicable law, our privacy policy, and customer agreements as applicable. We also conduct risk assessments and data protection impact assessments for LinkedIn products that incorporate generative AI (GAI) models to

power features that create or suggest content, in accordance with applicable law.

A thorough security standard has been set up at LinkedIn to lessen the risks linked with the application of Generative AI technology. This standard encompasses requirements for the complete lifecycle of Generative AI models, ranging from their creation and training to their implementation and supervision. Generative AI models that are created and upheld by third parties for LinkedIn are required to comply with secure software development controls. This involves consistent auditing and penetration testing to guarantee the security of the models and their adherence to LinkedIn's policies. LinkedIn performs an internal penetration test before releasing the functionality to the public. External penetration tests are conducted annually thereafter.


Addressing bias

LinkedIn strives to ensure that its products and services are fair, including by measuring and mitigating algorithmic bias. Our objective is to ensure equally qualified members receive equal treatment by our models. As with all our products, our teams continuously assess our systems, and if harmful biases are identified, we will work to address them.

LinkedIn conducts fairness and bias reviews on a per-model level, including models used in talent solutions products. When potential biases are identified, the team investigates and, as appropriate, retrain the model to mitigate the biases (or deramps the model).

Performance monitoring

We have ongoing monitoring of stability, latency, and output quality signals. We perform regular evaluation to ensure consistency of AI model quality. We have processes to monitor fulfillment failure rates and user quality feedback signals.



Data Storage and Processing

Processing of input and output data

As a user of LinkedIn enterprise products and services, the information that you choose to provide as inputs, as well as the responses to your inputs and your use of this feature, will be associated with your enterprise seat and will be considered “customer personal data” under our [Data Processing Agreement](#). While your prior “sessions” with these features may not be viewable by you after the close of a session, we will retain this information beyond the session. Like other personal data on LinkedIn, the purchaser of your enterprise seat can [request](#) to see what we have retained and request to take action on your personal data. Enterprise users must submit their deletion and export requests to their master administrator.

For AI-Assisted Messaging, we leverage the data that our customers and members have chosen to share with us that is already input into our platform and then use AI to help format and create all the elements of an outreach message. Certain fields include:

- **The opportunity:** job title, responsibilities, location, salary, workplace type, employment type, seniority, etc. for messages with job context
- **Good fit:** prioritized skills/skill match with job title, experience, past companies/roles, past interest in company, past applications, open to work, current/past schools, etc.
- **Recruiter/Company/Team background:** recruiter info, company info, culture, values, etc.

LinkedIn data will not be stored or used to train third-party AI models.





Processing of sensitive information

AI-Assisted Messages in Recruiter does not require the processing of sensitive personal data provided by customers. Any data customers provide is subject to the LinkedIn Data Processing Agreement, including with respect to confidentiality, security, and privacy.

Information that members share on LinkedIn.com like their names, job title, and additional information on their profile are used to match their information with recruiters' search results. However, this data is not considered sensitive customer data and is provided by members on LinkedIn. If users input an ask to the AI feature for sensitive personal information, they will get an error that it's not supported.



Data Training

Data collected for the AI tools in LinkedIn Recruiter come from LinkedIn's economic graph — information that members and companies share and update on LinkedIn.

LinkedIn uses a combination of GPT models from [OpenAI](#) to power AI-assisted Messages in different languages. Microsoft has committed not to use these inputs for further training of AI models. This is supported by the standard Azure OpenAI Services terms, which specify that LinkedIn data will not be stored or used to train third-party AI models, including those by OpenAI.

Personal data

We have implemented several security controls to ensure that customer user personal data is not used or accessible for training LinkedIn's Generative AI (GenAI) Language Learning Model (LLM). These controls include:


- **Data Processing Agreement Compliance:** LinkedIn's use of customer data, including for AI features, adheres strictly to the terms outlined in the Data Processing Agreement ([DPA](#)). This agreement specifies how customer data can be used, ensuring it aligns with privacy and data protection standards.
- **Microsoft Assurances:** Microsoft processes data in accordance with standard [Azure OpenAI Services](#) terms. This ensures that LinkedIn data, including customer inputs into LinkedIn's GenAI features, will not be stored or used to train third-party AI models, including OpenAI.
- **Commitment to Data Privacy and Security:** Microsoft has committed not to use customer inputs for further training of AI models. This commitment is part of the broader data privacy and security measures outlined for [Azure OpenAI Services](#), ensuring that customer data is protected and used in compliance with applicable regulations.
- These measures collectively ensure that customer user personal data is protected and not used to train LinkedIn's GenAI LLM model, maintaining the confidentiality and integrity of customer data in line with LinkedIn's privacy and security commitments.

Keeping our members in control of their data

In regions where LinkedIn or its affiliates use member data to train generative AI models for content creation, members can choose to opt-out of having your personal data and content you create on LinkedIn used for training (including fine-tuning). To opt out, use the [Data for Generative AI Improvement member setting](#). Opting out means that LinkedIn and its affiliates won't use the user's personal data or content on LinkedIn to train models going forward, but does not affect training that has already taken place. We are initially making this setting available to members whose profile location is outside of Canada, the EU, EEA, UK, Switzerland, Hong Kong, or Mainland China. If the members live in the stated regions above, LinkedIn will not use personal data or content on LinkedIn to train or fine-tune generative

AI models for content creation without further notice. In addition to the setting set forth above, members can also object to the use of their personal data for training non-content generating GAI models using the LinkedIn Data Processing Objection Form, accessible [here](#).

[The opt out setting on LinkedIn.com](#) allows individual LinkedIn members to control if their LinkedIn data can be used to train generative AI models for content creation - this includes any personal data, including personal data in LinkedIn Recruiter. If the member opts-out using that setting, we won't use their Recruiter personal data to train GenAI models for content creation. The controls for data training of AI models will remain at the member level to keep our individual members in control of their own data.



AI-Assisted Messages

AI explainability

AI-Assisted Messages in LinkedIn Recruiter help recruiters compose and send unique, personalized messages to candidates with AI assistance. AI-Assisted Messages draw on profile information provided by the candidate and combine it with the job requirements to create a fully personalized message.

The following data fields are currently used for AI generation (where available):

- **Recruiter information:** Name, title, company name, company information, and custom content
- **Candidate information:** Title, "About Me" from member profile, candidate company, open to work, company follower, past applicant, and mutual connections
- **Job information:** Job title, job company, Job URL, skills, location, workplace type, compensation, and employment type

Once a recruiter identifies a candidate, they can click "Draft with AI" in the Compose Message window to draft a personalized message specifically for that candidate. Once the message is drafted, they can review and send the message or choose to edit the content before sending. Users also have the ability to provide in-product feedback, via a thumbs up/down, when the results do not meet their expectations so the product can learn and improve over time.

With AI-Assisted Messages, recruiters can send an InMail message and follow-up message to up to 25 candidates at once. The AI model will draft a unique message personalized to each candidate by default. If a user selects a personalization element to include in the message, AI-Assisted Messages will include the message details relevant to those candidates. For example, if they select "Open to work", AI-Assisted Messages will include relevant message details for candidates who have the Open to Work signal enabled on their profile.

Compose message

Private

To: Antoine Martel X Jerome Bell X Floyd Miles X Andy Wong X Lu Hong X James Wang X Justin He X Ella vta Zeng > View all 25

Antoine Martel · 2nd Sourced

Marketing @ Oustia

London Area, United Kingdom

Antoine@gmail.com

Recruiting activity

Viewed by Tania Reynolds

Rejected reasons: this position is put on hold

7/1/2023 · 4:42 PM

Marsha Jones shared this profile for feedback.

7/1/2023 · 4:42 PM

Resume has been added

Antoine_resume.pdf

7/1/2023 · 4:42 PM

View all →

Summary

Integer sem orci, cursus id arcu vel, cursus lobortis... See more

Open to work

Job title

VP, Growth Marketing · Growth Marketing Leader · Director of Growth Marketing · Growth Marketing Manager

Job type

Full time

Locations

Onsite: Santa Clara, California · Mountain View, California · San Jose, California · San Mateo, California

Initial message

Send immediately via InMail

Exciting opportunity at Runity

B I

Review and edit AI draft for accuracy.

Hi Antoine,

I hope this message finds you well. My name is Camila, and I work as a recruiter at Runity. I noticed on your LinkedIn profile that you are open to work, and I have a position that could be a perfect match for you, as it satisfies several aspects you are open to in your job search.

Runity is currently looking for a Director, Growth Marketing for a full-time position in London, England, United Kingdom. The job workplace type is hybrid. We are looking for someone with experience in Growth Marketing and Lead Generation. We offer competitive salary ranging from \$185,000 to \$231,000, commensurate with qualifications and experience.

Based on your Open to Work preferences, I believe your experience as a Director of Growth Marketing at Onyx might be relevant for this role.

Could you please let me know a suitable time for us to have a conversation about this role?

Best regards,
Camila

Draft settings

Select fields to include in your message draft.

Job detail

Skill ✓ Compensation ✓

Location ✓ Employment type ✓

Workplace type ✓

Candidate profile

Experience + Open to work ✓

Mutual connections +

Past applicants ✓

Company follower ✓

Company info

About company

Save & draft for Antoine

Save & draft for All

Clear all drafts

Preview 18/500 InMail credits

Send to all

Send to Antoine

LinkedIn runs all AI-created content through our existing content filters. While we have taken safety precautions, this feature may produce surprises or mistakes. When using the AI-Assisted Messages feature on LinkedIn, users will still need to review the message draft for accuracy and make any necessary edits. Users are still responsible for the message content that is sent and may be required by applicable law to add additional information, such as a salary range, other pay or benefits information, an applicant privacy notice, or other information.


Users have the ability to create and add custom content, such as details about your company benefits, to a message within the Draft Settings menu at the bottom of a message. Any custom content created is available to that user only and isn't visible or accessible to other recruiters on a contract.

AI-Assisted Messages does have an auto-draft feature that allows users to automatically start composing an AI-assisted message that can be edited and sent. This auto-draft feature can be turned on/off by default at the contract level by administrators on the account or at the Recruiter seat level by the recruiter.

Using AI-Assisted Messages is not mandatory for users. They still have the option of manually drafting messages or using their existing message templates, if preferred. To give users more control, they can also manage their default AI settings under AI Preferences within their Account Settings.


AI features in Recruiter are primarily recommendations and do not make hiring, sourcing or engagement decisions for the users. Users will still own the decision making of engaging or not engaging candidates on LinkedIn Recruiter.

AI-Assisted Messages has been subject to our review process to help ensure customers can use it in compliance with their legal obligations, our privacy policy and our contractual obligations. Our security commitments with respect to customer data are set forth in our agreement with customers. Like all of our services, AI products are subject to security reviews designed to prevent harm.



AI technology and models used

LinkedIn uses a combination of GPT models from OpenAI to power AI-assisted Messages in different languages.





Compliance with applicable privacy, data protection, and AI/ADT laws, regulations, executive orders, and guidance.

LinkedIn complies with applicable law in the jurisdictions where our products are offered and we always seek to enable our customers to comply with their legal obligations when using our products. We are committed to keeping our platform trusted and professional, and we respect the laws that apply to us in the countries in which we operate.

Fair Hiring Practices: LinkedIn's hiring product offerings focus on increasing hiring efficiency, boosting candidate engagement & helping hirers personalize at scale. Specifically, AI-Assisted Messages in Recruiter helps recruiters leverage candidate, company and job information to draft unique, fully personalized messages to candidates at scale. Once drafted, users can make edits and customizations to ensure the message accurately represents their tone and reflects their hiring needs. Recruiters decide who to reach out to and who to advance in the hiring pipeline. AI-Assisted Messages will not engage with candidates without user oversight, screen out applicants or make hiring decisions. Recruiters will still need to take action and make final decisions on whether to engage or advance applicants.

Additional Resources

- [LinkedIn's Responsible AI Principles](#)
- [Microsoft: Governing AI: a Blueprint to the Future](#)
- [Microsoft: Advancing the Future Through Responsible AI](#)
- [Data, Privacy, and Security for Azure OpenAI Service](#)
- [LinkedIn's Smart Trust Center](#)
- [LinkedIn's Public Trust and Compliance page](#)
- [LinkedIn Subscription Agreement](#)
- [LinkedIn Data Processing Agreement](#)
- [LinkedIn Privacy Policy](#)
- [Regional privacy web page](#)

