AI Usage Policy - draw.io

- Applies to: draw.io at app.diagrams.net, draw.io for Confluence Cloud, draw.io Zero Egress for Confluence Cloud and draw.io for Jira Cloud
- Owner: draw.io Ltd (UK). Marketplace vendor of the Atlassian apps: Seibert Group GmbH.
- Version 1.3, 28 November 2026. Contact: support@draw.io

Summary

- The AI features are optional and only run when a user asks for them.
- We do not use any user data, including diagram content, to train AI or machine-learning models, and our AI providers do not use requests from our AI features to train theirs.
- draw.io does not store or log AI requests or their results.
- The AI models run in the European Union, and the provider that runs them keeps no copy of requests.
- Administrators decide whether the AI features are available in the Atlassian apps. The Zero Egress app has no AI features.

What the AI features do

- **Create a diagram** from a description the user types.
- **Change the open diagram** from an instruction the user types (the draw.io editor, including the editor in draw.io for Confluence Cloud).
- **draw.io AI in Jira** creates a diagram from a description, using the Jira issue as context. The diagram is shown as a preview and is added to the issue only when the user chooses Insert.

The result is an ordinary, editable diagram. Nothing is saved until the user saves it, and every change can be undone.

What is sent, and where

A request contains:

- the text the user types;
- when the user asks to change a diagram, that diagram;
- in Jira, the issue's summary, description and comment text, and the project's key and name. Comment authors and other Jira user details are not sent, and a person mentioned in the text is sent as the name shown, without their account ID. Version 4.4.0 and earlier instead send the issue as Jira returns it, which also includes the Jira user details, such as names and account IDs, of the people who wrote the comments.

At app.diagrams.net, a request from a user signed in with Google Drive says only that the user is signed in, which selects the model. It contains no email address, name or other detail of the user's Google account.

The request goes to draw.io's AI service, which runs on Cloudflare in the data centre nearest the user, and from there to one of these models:

- Anthropic's Claude, run by Amazon Web Services on Amazon Bedrock in the European Union
- a model run by Cloudflare (Workers AI), which also decides whether a diagram or a Mermaid diagram is wanted and is the last fallback.

Nothing else from Confluence or Jira is sent. The AI features cannot read other pages, issues or diagrams.

Retention and training

Party	Keeps requests and results	Uses them for training
draw.io	No. Not stored and not logged.	No
Amazon Web Services (Amazon Bedrock)	No. Our account is set to keep no requests or results, and nothing is logged.	No
Cloudflare (Workers AI)	No	No

Anthropic, which makes the Claude models, does not receive the requests. The requests stay in the European Union, and the AWS GDPR Data Processing Addendum applies. The providers are listed with our other sub-processors on our trust centre.

Controls

- **draw.io for Confluence Cloud:** off by default. A Confluence administrator turns the AI features on in the draw.io configuration. Lockdown mode switches them off.
- **draw.io Zero Egress for Confluence Cloud:** no AI features. The Atlassian Forge platform blocks every connection outside Atlassian.
- **draw.io for Jira Cloud:** on by default. A Jira administrator can switch AI off on the draw.io administration page (AI Configuration, Enable AI), which removes the draw.io AI action and button from issues.
- **app.diagrams.net:** the AI features run only when a user opens them. Lockdown mode switches them off.

Limitations and human oversight

AI output can be wrong or incomplete. Users should review a generated diagram before relying on it, and they decide whether to keep it. The AI features act only on a user's request and only on what that request contains; they take no actions of their own.

Our own use of AI

Our support bot sorts new support requests and drafts answers from our knowledge base. It uses Anthropic's Claude on Amazon Bedrock in the European Union, under the same terms as the AI features above.

Our staff also use Claude, Anthropic's AI assistant, to help investigate and answer support requests, so a support request and the details sent with it can be processed by Anthropic in the United States. Anthropic does not use them to train models, and deletes them within 30 days, or keeps them for up to 2 years if they are flagged for misuse. Transfers are covered by the standard contractual clauses and the UK Addendum in Anthropic's data processing terms. We do not use customer data or diagrams to train AI models.

Changes to this policy

We update this policy when the AI features, the providers or the way data is handled change, and publish the new version on our trust centre.