Figma

AI PRODUCT SYSTEM CARD

Figma Design - Text Editing

LAST UPDATED DATE, 10/13/26

PRODUCT STATUS	AVAILABLE ON THE FOLLOWING CUSTOMER TIERS	WHAT IT DOES
Generally available	Professional, Organization, and Enterprise plans. Full seats with edit access in shared team files; all seat types can try the tools in drafts.	Figma Design Text Editing generates and adjusts copy in design files. It also renames layers, suggests resource tags, and creates a starting layout with placeholder content.

About this document

An AI feature is more than the model behind it. How well it works for you depends on the people using it and the tasks they use it for. This system card explains how Figma Design Text Editing works, best practices for performance, and how your data is used. This document is for information only. It is not a contract, a warranty, or legal advice.

Contents

1 Overview at a glance

2 What it does and how it works

3 Models and providers

4 Intended use

5 Limitations

6 Your data and privacy

7 Content moderation

8 Safeguards and security

9 Performance, evaluation, and fairness

10 Human oversight, feedback, and control

11 Transparency

1 Overview at a glance

This section summarizes the feature, its availability, its controls, and the main data handling facts.

What it does	Generates and adjusts copy, renames layers, suggests resource tags, and creates a starting layout with placeholder content.
Who it is for	Full seats with edit access on Professional, Organization, and Enterprise plans.
Where it appears	Figma Design. Suggest Resource Tags also appear in Figma Community.
Availability	Generally available on paid plans. Not available on Starter.
Model selection	Configured by Figma. No model picker.
Data processing and storage	Figma stores your content in the United States. Model providers do not store your data at rest.
Training on your content	Model providers never train on your data. Whether Figma uses your content to improve its broader AI depends on the plan-specific content training setting. Section 6.6.2 describes this setting.
Content moderation	Adjust Text and Content Fill screen customer instructions and selected text before generation. Model providers may apply safety filters during inference.
Human oversight	You start every action and decide whether to accept, edit, undo, or discard the result.
FigGov Experience	Not available.
Gov+ Experience	Available with additional administrator controls. Restricting AI traffic can make some AI features unavailable.
EDU Experience	Available to higher education and bootcamp users. Not available to K-12 users.

2 What it does and how it works

2.1 Description

Figma Design Text Editing generates new copy from an instruction and adjusts selected copy for tone, length, or language. It can rename layers, suggest resource tags, and create a First Draft layout with placeholder content. You start an action from the Actions menu or inline text tools. The result becomes editable content that you can accept, change, undo, or discard.

2.2 What the AI generates or predicts

The feature generates and rewrites text, including shortened copy, translations, and changes in tone. It also generates layer names, resource tags, and a starting layout with placeholder copy. It does not generate photographs, audio, video, or executable code.

2.3 How it works with other Figma products

Figma Design Text Editing works with the Design file you have open and the text, layers, or file structure in scope. Suggest Resource Tags also works in Figma Community. Results remain editable content in the relevant file. Sections 6.2 and 8.2 describe the file permissions and Figma-wide administrator controls that apply.

3 Models and providers

3.1 Models and where they run

OpenAI GPT models generate and adjust text, rename layers, suggest resource tags, and support First Draft. Separate moderation models screen inputs for the covered text actions.

MODEL AND DEVELOPER	ROLE IN THIS FEATURE	WHERE IT RUNS
GPT models, OpenAI	Generate and adjust copy, rename layers, suggest resource tags, and support First Draft.	OpenAI and Microsoft Azure OpenAI; AWS Bedrock for Bedrock-pinned Adjust Text
Moderation models, OpenAI and Amazon	Screen inputs for Adjust Text and Content Fill before generation.	OpenAI and AWS Bedrock

3.2 Can I select the models?

No. Figma configures the model for each action. There is no model picker, and moderation for the covered workflows runs on models Figma configures.

3.3 What the providers do with your content

The providers receive your instruction and the relevant text or layer content only to generate a result or screen a request. They are contractually prohibited from using Figma customer content to train or improve their models. Direct OpenAI requests use zero data retention, and the other hosting paths use equivalent no-logging configurations, so prompts and outputs are not retained.

3.4 How a model change reaches you

Figma pins these features to specific model versions. A provider update does not reach you until Figma validates it against the version in use and completes a staged rollout. Quality or safety regressions can stop the rollout, and Figma keeps a tested rollback path.

4 Intended use

4.1 Primary intended users

The feature is for people with a Full seat and edit access who write, revise, translate, or organize content in Figma Design on a paid plan.

4.2 Primary intended uses

- Generate placeholder copy, rewrite existing copy, shorten it, translate it, or change its tone.
- Rename layers and suggest resource tags so files and shared resources are easier to organize.
- Create a First Draft layout with placeholder content as a starting point for further design work.

4.3 Where it works best

- Your instruction states the tone, length, language, or outcome you want.
- The exact text or layers you want to change are selected and visible in the file.
- You plan to review the result and edit or regenerate it before using it.

5 Limitations

5.1 Known limitations

- Generated and rewritten copy can be inaccurate, off brand, or awkward, so review it before relying on it.
- Quality varies for specialist terminology, languages with less provider support, and very short or tightly constrained copy.
- Layer names and resource tags can be wrong or inconsistent, and First Draft produces a rough starting point rather than a finished design.

5.2 Factors that affect output quality

- The specificity of your instruction, including the requested tone, length, language, or outcome.

– The selected text, layers, and file structure that the action uses as context.
– The provider model version, which can change behavior after Figma validates and adopts an update.

5.3 Best practices for performance

– Review generated copy and file organization before you share or publish the result.
– Use ordinary text editing, Undo, or the available change controls to correct a result or restore the previous state.
– Test workflows that depend on specialist terms or languages with less provider support before adopting them broadly.

6 Your data and privacy

6.1 Data that is Processed

DATA	WHAT IT INCLUDES
Your instruction	The prompt you provide for the selected action.
Design content in scope	Selected text, layer names, and basic file structure used to generate or organize the result.

6.2 Who at Figma can access your data

Access follows the permissions set in your Figma organization, and Figma confirms your identity and your permission to act on the target file before a build runs.

Figma tightly restricts internal access to customer data on a least-privilege basis, granting it only to employees whose roles require it and only for the specific data needed.

6.3 Who at our model providers can access your content

Your prompt and data is shared with model providers solely so that they can process your request and return the requested output. Provider staff cannot access this content. Every provider operates under zero data retention or an equivalent no-logging configuration, and nothing is retained once your request has been served. No third party is permitted to train on your content.

6.4 Where your data is stored

Figma and the model providers handle processing and storage separately, so each has its own subsection below.

6.4.1 Where Figma stores your content

Figma stores your content in the United States.

6.4.2 Where the model providers store your content

Not applicable. No model provider that powers this feature stores your data at rest.

6.5 Retention and deletion

6.5.1 What is retained

WHAT IS RETAINED	PURPOSE	RETENTION PERIOD
Your files, including prompts and generated content	Your work product, stored as customer content	For the life of the account or until deletion
Usage and telemetry events, and metering aggregates	AI billing	For the life of the account
Operational telemetry and error context	Debugging and incident response	For the life of the account

6.5.2 Deleting your data

You can submit a deletion request to remove the data Figma stores. Figma propagates the request across storage used for AI evaluation and usage data. Customer content is otherwise stored for the life of the account or until deletion. The model providers do not retain prompts or outputs, so there is generally nothing on their side to delete.

6.6 Training on your content

6.6.1 Whether the model providers train on your content

Providers never train on your data. There is no separate opt-out for provider training, because your content is already excluded by contract and by the no-retention configuration Figma runs on every provider.

6.6.2 Whether Figma trains on your content

Whether Figma uses your content to improve its own AI is governed by an account-level content training setting that an administrator manages. This setting is independent of the control that turns AI on or off. On Starter and Professional, the team-level setting is on by default and an administrator can turn it off. On Organization and Enterprise, content training is disabled and cannot be turned on. Content from Figma for Government and education accounts is never used for training, regardless of the setting.

6.7 Sensitive and confidential data

Figma Design Text Editing does not require special category personal data and does not ask for it. Do not include this data in prompts or design content. Apply your own policies for confidential, regulated, and personal data, and confirm that you have the right to use the content you provide.

7 Content moderation

7.1 What is screened

Adjust Text and Content Fill screen customer instructions and selected design text before generation. Figma validates generated text for the expected shape and length before it is returned, and the model providers may apply their own safety classifiers during inference.

7.2 What is blocked

Prompts and generated results are blocked when they violate our acceptable use policy. The following are examples of content prohibited on the platform.

- Child sexual abuse material, including any depiction, simulation, or facilitation of sexual content involving minors.
- Pornographic or graphic sexual content involving adults.
- Harassment, hate, dehumanization, and credible threats of violence toward protected classes or individuals.
- Step-by-step instructions for committing crimes.
- Instructions for building weapons or explosives, or for synthesizing controlled substances.
- Malicious code, exploits, backdoors, and instructions enabling unauthorized access or system abuse.
- Encouragement or instruction of self-harm or suicide.

These categories apply to submitted content and generated content under the same acceptable use policy.

7.3 What happens when content is blocked

A blocked input stops the action before generation. Figma tells you that the input was blocked and asks you to provide different input. If a block appears wrong, report it through the channels in Section 10.

7.4 Whether moderation can be edited or configured

End users cannot turn off moderation or adjust its thresholds. The safety settings for the covered workflows are enforced by default.

8 Safeguards and security

8.1 Security controls

- Traffic to model providers is encrypted in transit with HTTPS and Transport Layer Security. Requests through the managed AWS path are signed and encrypted end-to-end.
- Production data at rest is encrypted. Keys are managed in a dedicated key management service, and credentials are stored separately.
- Every request passes through Figma's AI gateway for routing, rate limiting, structured logging, and per route alerting.
- Server-side instructions and provider credentials stay on Figma servers and are not sent to your client.
- Per-user and per-organization throttles and quotas protect services from abuse and overload, and Figma can disable an affected feature quickly.
- Internal data stores are isolated and reachable only through approved roles and network policies.

8.2 Administrator controls

- On Enterprise, organization administrators can enable or disable AI at the organization or workspace level. Workspace and team administrators cannot change that setting.
- On Professional and Organization, an administrator can enable AI if it remained disabled after the beta, but cannot disable it again after enabling it.
- Administrators manage the content training setting described in Section 6.6. The setting is team-level on Professional and organization-level on Organization and Enterprise.
- With Governance plus, organization administrators can restrict AI traffic to models hosted in Figma's AWS environment. Some features may become unavailable and output quality may differ.
- The feature is not available on Figma for Government. It is available to higher education and bootcamp users, but not K-12 users, and Government and Education content is never used for training.

9 Performance, evaluation, and fairness

9.1 How we evaluate this feature

Figma tracks action success, end-to-end latency, client stability, and high-severity incidents. Figma also runs sampled human and model-graded reviews for accuracy, tone, instruction adherence, and usefulness.

Systemic quality regressions, elevated policy violations, repeated timeouts, and failures to return usable output can block a release. Critical regressions trigger incident response, a provider switch, or a rollback, and other issues are reviewed on a regular cadence.

9.2 What we measure

METRIC	WHAT IT MEASURES	WHY IT MATTERS TO YOU
Action success and end-to-end latency	Whether a supported action returns a usable result, and how long it takes.	Changes can identify a reliability or provider problem before it disrupts more work.
Sampled quality review	Human and model-graded review for fabrication, tone, instruction adherence, and functional validity.	It identifies text that is well formed but inaccurate, off brand, or not useful for the requested task.
Client stability and high-severity incident status	Crashes during feature use and serious issues in an active rollout.	These measures help Figma stop or roll back a release that is not operating safely or reliably.

Figma sets internal baselines for these measures and reviews the feature when a baseline is crossed.

9.3 Fairness and bias

9.3.1 What Figma reviewed

Figma reviews these features through red teaming with demographic and sensitive prompts, and through sampled human review of generated names, biographies, and copy for stereotyping. These qualitative and adversarial methods fit a workflow that does not use or infer sensitive demographic attributes. Findings are logged and routed to the feature team. Neutral prompts can still produce examples that default to a dominant demographic or stereotype a role.

9.3.2 What the model providers reviewed

For behavior inside the models, Figma relies on the bias and fairness work OpenAI publishes. The provider describes filtering for harmful and biased content and external red teaming across languages and regions. Figma does not train the provider models or control their training datasets.

9.3.3 Known differences in performance

Text quality can vary by language depending on provider support. Figma is not aware of systematic performance differences by demographic group or geographic region.

9.4 Evaluate it for your own use

Performance varies with the copy and workflow you apply it to. Pilot the feature with your own files, terminology, languages, and reviewers before relying on it broadly. The source text and layers remain visible, so you can

compare the result with what shaped it. The same instruction can produce different results, and Figma does not guarantee identical reproduction.

10 Human oversight, feedback, and control

10.1 Human oversight

Figma Design Text Editing keeps a person in the loop. You start every action, review the result in the file, and decide whether to accept, edit, undo, or discard it. Nothing is applied without your choice.

The feature does not make or enforce decisions about people, ban customers, or report people to outside agencies. Only signed in Figma customers with access to the file can use it.

10.2 Reporting problems and giving feedback

- Report a failure, unsafe output, or other concern through the Figma Help Center contact form or your usual support channel.
- Enterprise customers can raise concerns through Customer Success.
- Support triages reports and escalates them to product, security, and legal teams when needed. Feedback can inform prompts, model selection, safety settings, and Help Center guidance.

10.3 Turning the feature off

Organization and team administrators can turn off AI features at the organization or team level. The feature stays off until an administrator enables it again.

11 Transparency

11.1 How AI involvement is disclosed

The features appear as named and labeled AI text tools in the Actions menu and inline text tools. You start each action deliberately, and the tools are not named or styled to look like a real person. The editable result appears in the file where you can review it.

11.2 How we communicate changes

Figma communicates material changes to functionality, intended use, and reliability through release notes, Help Center documentation, in-product messaging, and standard release communications.