Figma

AI PRODUCT SYSTEM CARD

Figma Design - Image Editing

LAST UPDATED DATE, 10/13/26

PRODUCT STATUS	AVAILABLE ON THE FOLLOWING CUSTOMER TIERS	WHAT IT DOES
Generally available. Design Agent image actions are generally available.	Professional, Organization, and Enterprise plans. Starter access includes Generate Image and Edit Image. Education access is limited to higher education and bootcamp customers.	Design Image Editing brings AI image tools into Figma Design. It generates and edits images, including background and object removal, expansion, upscaling, layer separation, and vectorization.

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 Design Image 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 product experience, availability, data handling, and customer controls.

What it does	Generates and edits images on the Figma Design canvas.
Who it is for	People using Figma Design on Professional, Organization, and Enterprise plans. Higher education and bootcamp customers also have the standalone tools.
Availability	Standalone tools are generally available. AI credit allowances vary by plan. Supported image actions in Design Agent are generally available.
Model selection	Partly. Generate Image and Edit Image offer models from a Figma approved set. Figma configures the other tools and all moderation models.
Data processing and storage	Figma stores customer content in the United States. Model providers do not store it at rest.
Training on your content	Providers never train on your data. Figma training depends on the plan-specific administrator setting described in Section 6.6.2.
Content moderation	Prompts for Edit Image and Generate Image are screened before a model call. Providers apply safety classifiers to generated output.
Human oversight	You start the standalone tools and review every result. Design Agent may apply image changes during a run you start, and you can review and undo the run.
FigGov Experience	Not available.
Gov+ Experience	Available with restrictions when administrators require AWS hosted models. Design Agent is unavailable when AWS Bedrock pinning is enabled.
EDU Experience	Available to higher education and bootcamp users, but not to K-12 users.

2 What it does and how it works

2.1 Description

Design Image Editing brings AI image tools into Figma Design. You can replace an image, remove a background or object, expand an image, upscale artwork, fill content, generate an image from a prompt, separate a flat image into layers, or convert a raster graphic into vectors. You start each standalone tool and decide whether to keep, change, or discard its result. You can also ask Design Agent to apply supported image actions during a task.

2.2 What the AI generates or predicts

The tools generate and transform image content. They can produce a new or edited image, remove a background or object and fill the gap, expand a canvas, increase resolution, separate layers, or create vectors from a raster image. They do not generate video, audio, or code, and they do not predict outcomes about people.

2.3 How it works with other Figma products

The standalone tools run inside Figma Design from the canvas and the Editor AI menu. Tools can act on selected images or frames, and Generate Image can also create a new image without a selection. Results remain in the Figma Design file as editable content.

Supported image actions can also run through Design Agent. The administrator permissions and AI controls described in Sections 6.2 and 8.2 apply to both paths.

3 Models and providers

3.1 Models and where they run

Design Image Editing uses image generation, image editing, orchestration, and moderation models from external providers. Figma does not train or fine-tune a model for this feature.

MODEL AND DEVELOPER	ROLE IN THIS FEATURE	WHERE IT RUNS
Gemini image models, Google	Generate Image and Edit Image.	Google enterprise platform
ChatGPT Images, OpenAI	Generate Image and Edit Image.	OpenAI
Jasper image models, Jasper	Background removal, object removal, expansion, upscaling, and layer separation.	Jasper
Recraft Vectorize, Recraft	Converts raster images into vector graphics.	Recraft
Stable Image models, Stability AI	Generate Image and Edit Image when Governance+ requires AWS hosted models.	AWS Bedrock
GPT Sol, OpenAI	Plans and invokes supported image actions in Design Agent.	OpenAI
Moderation models, OpenAI and Amazon	Screen prompts for Edit Image and Generate Image.	OpenAI, Microsoft Azure OpenAI, and AWS Bedrock

3.2 Can I select the models?

Partly. Generate Image and Edit Image offer a choice from models Figma has approved, and Figma sets the available list and the default. The other tools and all moderation run on models Figma configures. Additional selectable models include: OpenAI ChatGPT Images, Google Nano Banana, and Google Nano Banana Lite.

3.3 What the providers do with your content

Providers receive your prompt, image, and relevant tool parameters only to process the request and apply safety checks. Providers do not train on this content. This exclusion is enforced through zero retention or equivalent no logging configurations, and through contractual no training terms on other paths.

3.4 How a model change reaches you

Figma pins each tool to a specific model version instead of a floating alias. Before adopting a different version, Figma compares quality and safety with the version in use, stages the rollout, watches provider release and retirement notices, and keeps a tested rollback.

4 Intended use

4.1 Primary intended users

The feature is for people who create and edit images in Figma Design on Professional, Organization, and Enterprise plans. Standalone tools are also available to higher education and bootcamp customers, but not to K-12 customers or Figma for Government customers.

4.2 Primary intended uses

- Retouch an image by removing a background or object and filling the remaining area.
- Adapt artwork by expanding an image, increasing resolution, or separating a flat image into editable layers.
- Create new material by generating an image from a description or converting a raster graphic into vectors.

4.3 Where it works best

- The source image clearly shows the content you want to preserve or change.
- The request describes one scoped change in concrete terms.
- The marked region or object matches the intended edit, and a person reviews the result before use.

5 Limitations

5.1 Known limitations

- Generated and edited images can contain artifacts or inaccurate details, so review is required before use.
- Upscaling and content fill can introduce detail that was not present in the original image.
- A result can miss your intent or brand, and behavior can shift when a provider model changes.

5.2 Factors that affect output quality

- The quality and composition of the source image.
- The specificity and scope of the prompt.
- The selected tool and the precision of the marked region or object.

5.3 Best practices for performance

- Review every result at canvas scale and full resolution before using or sharing it.
- Keep a preferred result before regenerating, because the same request can produce a different image.
- Test the tools with representative brand imagery and reviewers before adopting them in a recurring workflow.

6 Your data and privacy

6.1 Data that is Processed

DATA	WHAT IT INCLUDES
Your prompt	The text instruction you provide for image generation or editing.
The selected image or frame	The image content the tool generates from or changes.
The marked region or object	The part of an image selected for removal, fill, expansion, or another edit.
Tool settings	The parameters required for the selected image action.

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

Submit a deletion request to remove AI evaluation and usage data from Figma systems. Figma propagates the request across the storage systems used for those records. Figma Design files and the images in them remain for the life of the account or until customer deletion.

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

The tools do not require special category personal data and do not ask for it. Do not include that data in images or prompts. Apply your own policies for confidential, regulated, and personal data before using an AI image tool.

7 Content moderation

7.1 What is screened

Figma screens text prompts and input images for Edit Image and Generate Image before a model call. The providers apply their own safety classifiers to generated output. Figma validates returned images for the expected type, dimensions, and size before placing them on the canvas.

The other image tools rely on provider safety controls and Figma output validation. Image content is placed on the canvas rather than executed, and these tools do not generate or run code.

7.2 What is blocked

Prompts and results are blocked when they violate the acceptable use policy. The following categories are examples of prohibited content.

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

– Encouragement or instruction of self-harm or suicide.

7.3 What happens when content is blocked

When a request is blocked, Figma explains that it could not be completed under AI usage policies and asks you to revise the input or try again. You can report a block through the support channels in Section 10.

7.4 Whether moderation can be edited or configured

End users cannot turn off moderation or change its thresholds. Safety settings are enforced by default.

8 Safeguards and security

8.1 Security controls

– Traffic to model providers is encrypted in transit with Transport Layer Security. Requests on the AWS Bedrock path are signed and encrypted end to end.

– Production data at rest is encrypted. Keys are managed in a dedicated key management service, and provider credentials are held separately in approved secrets management.

– AI requests pass through Figma's AI gateway, which applies routing, moderation entry, rate limits, structured logging, and alerting.

– Per-user and per-organization limits protect the service from abuse and overload, and Figma can disable a tool to contain a problem.

– Returned images are checked for expected type, dimensions, and size before they are placed on the canvas.

– Internal data stores are isolated behind approved roles and network policies. Section 6.2 describes employee access.

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 those settings.

– On Professional and Organization, administrators can enable AI if it was disabled during beta. After it is enabled, it cannot be disabled.

– Administrators manage the separate AI training setting described in Section 6.6.2.

– On Enterprise with Governance+, administrators can require AWS hosted models. This changes the available image tools as described in Section 1.

9 Performance, evaluation, and fairness

9.1 How we evaluate this feature

Figma evaluates model quality and safety before a version reaches customers. Systemic quality regressions, increased policy violations, repeated timeouts, unusable images, and unsafe output can block a release. Section 3.4 describes version changes and rollback.

In production, Figma tracks tool success, latency, incidents, and provider telemetry. Periodic human and model-graded reviews compare sampled output with its prompt and source image. Critical regressions can trigger rollback or a provider switch, while other issues are reviewed with product and platform owners.

9.2 What we measure

MEASURE	WHAT IT COVERS	WHY IT MATTERS
Tool success	Whether each image action returns a usable result.	Finds failures that prevent you from completing an edit.
End-to-end latency	Time from an image action to a returned result.	Finds reliability changes that slow the workflow.
Sampled quality review	Prompt match, source fidelity, invented detail, and image artifacts.	Finds results that look complete but do not match the requested change.
Safety and provider signals	Policy violations, provider errors, and model behavior changes.	Supports incident response and model rollback.

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 uses output quality evaluation, red teaming with demographic and sensitive prompts, and sampled human review of generated people and scenes for stereotyping. These qualitative and adversarial methods fit image output that does not carry the demographic labels required for group metrics. Findings are logged and routed to the feature team.

9.3.2 What the model providers reviewed

Google and OpenAI disclose bias and fairness work for their image models, including training data filtering and safety and responsibility evaluations. Jasper and Recraft publish no equivalent public model card, so Figma relies on its own sampled review for those paths.

9.3.3 Known differences in performance

Quality and appropriateness can be stronger for English and Western contexts, including generated text, attire, and settings. Quality also varies by image type and provider model. Figma includes non-English prompts and varied cultural subjects in sampled review and has not identified systematic differences by demographic group or region.

9.4 Evaluate it for your own use

Run a pilot with your own images, languages, brand requirements, and reviewers before depending on these tools in a recurring workflow. Decide which outputs require human review before use or sharing.

Each edit works from a visible source image, prompt, and selected region or object. Compare the result with that source, refine the input, and regenerate when needed. The same request can produce a different result, so retain an image you want to preserve before trying again.

10 Human oversight, feedback, and control

10.1 Human oversight

The standalone tools keep a person in the loop. You start each edit or generation, review the result on the canvas, and decide whether to keep, change, or discard it.

When the Design Agent invokes a supported image action, you start the task and the agent may apply changes during the run. You can review the effects and undo the entire run in one step.

The feature does not make or enforce decisions about people, restrict customer accounts, or report customers to external agencies. Only signed-in customers with feature access can use it.

10.2 Reporting problems and giving feedback

- Use the in-product feedback controls or report a failure, unsafe output, or other concern through the Figma Help Center and support channels.
- Enterprise customers can raise concerns through Customer Success.
- Support triages reports and escalates them to product, security, and legal teams when needed.

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 standalone capabilities are named AI tools in Figma Design, and you start each one from the canvas or Editor AI menu. When supported image actions run through Design Agent, the interaction occurs in a named AI assistant surface.

Generated and edited images remain as editable content on your canvas. You decide whether to keep, change, share, or discard the result.

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.