Figma

AI PRODUCT SYSTEM CARD

Figma Weave

LAST UPDATED DATE, 10/13/26

PRODUCT STATUS	AVAILABLE ON THE FOLLOWING CUSTOMER TIERS	WHAT IT DOES
Generally available	Weave Free, Starter, Professional, Team, and Enterprise plans	Figma Weave is a standalone generative AI product for building visual content on a node based canvas. It generates and edits images, video, audio, and 3D assets through models that a customer selects and connects into workflows.

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 Weave 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, its availability, its controls, and the main data handling distinctions described in this system card.

What it does	Generates and edits images, video, audio, and 3D assets through models connected on a node based canvas.
Who it is for	Creators and teams building visual content and reusable generative workflows.
Where it appears	In the standalone Weave product.
Availability	Generally available on Weave Free, Starter, Professional, Team, and Enterprise plans.
Model selection	You select the generation model for each node. Figma configures the moderation models.
Data processing and storage	Figma stores Weave content in the United States. Provider processing and temporary storage vary by the model you select, and some paths operate outside the United States.
Training on your content	Figma does not train its own models on Weave customer content. Verified providers are contractually barred from training, while unverified models follow their provider terms.
Content moderation	Text prompts are screened on configured routes. Images and video can be screened through separate input and output paths.
Human oversight	You trigger every generation and decide whether to keep, refine, rerun, export, or discard the result.
FigGov Experience	Not applicable. Weave is a separate product and the Figma for Government solution does not govern it.
Gov+ Experience	Not applicable. Governance+ applies to the main Figma platform. Weave Enterprise provides separate model management controls.
EDU Experience	Weave offers an institutional education program. Teach with Weave can provide education credits depending on the project. Education plan rules do not govern the standalone Weave product.

2 What it does and how it works

2.1 Description

Figma Weave is a standalone generative product with a node based canvas. You select models, connect them into a workflow, and pass the output of one node into another. You can rerun the workflow, refine individual nodes, save the workflow for reuse, or publish it as an app that other people can run.

2.2 What the AI generates or predicts

Weave generates images, video, audio, and 3D assets from prompts, uploaded reference media, and upstream node results. It also supports editing operations such as relighting, background changes, and object edits. It does not classify or score people.

2.3 How it works with other Figma products

Weave runs separately from the main Figma platform with its own plans, credit system, billing, and storage backend. A curated set of Weave tools can also appear in the main Figma platform, and a Weave workflow can be published as an app for other people to run. Main platform plan rules and data residency controls do not automatically apply to standalone Weave. Section 6.2 describes access to your data, and Section 8.2 describes Weave Enterprise controls.

3 Models and providers

3.1 Models and where they run

Weave provides a changing library of third-party generation models from multiple developers. These models create output by predicting content from patterns learned during training, so the same workflow can produce different results on different runs. Figma does not train or fine tune the models in this library.

MODEL AND DEVELOPER	ROLE IN THIS FEATURE	WHERE IT RUNS
Third-party generation models, multiple developers	Generate and edit images, video, audio, and 3D assets in user selected workflow nodes.	FAL, Replicate, and direct provider services
Moderation models, OpenAI, Amazon, and Hive AI	Screen text prompts and visual content on configured moderation paths.	OpenAI, AWS Bedrock, and Hive AI

3.2 Can I select the models?

Yes. You select the generation model for each node, and Figma marks models as verified or unverified. Enterprise administrators can restrict the available model set, while moderation always runs on models that Figma configures.

3.3 What the providers do with your content

Providers and aggregators receive your prompt, reference media, and generated results only to run the selected model or perform moderation. Verified model providers are contractually barred from using your content to train or improve their models. Unverified models follow their provider terms, and retention varies by provider and hosting path.

3.4 How a model change reaches you

Figma does not pin one version across the full Weave library. Providers can update or retire models, and Figma updates the library by adding, verifying, restricting, or removing models. A model's verification status identifies whether it runs under Figma reviewed terms, and Enterprise administrators can limit their workspace to an approved model set.

4 Intended use

4.1 Primary intended users

Weave is for creators and teams who build visual content and reusable generative workflows. It is available to members of the public and to teams using Weave Enterprise controls.

4.2 Primary intended uses

- Generate and edit images, video, audio, and 3D assets for creative work.
- Connect multiple models into reusable workflows that pass results between nodes.
- Package a completed workflow as an app that other people can run.

4.3 Where it works best

- Your prompt states the intended subject, style, format, and constraints clearly.
- The selected model matches the task and the type of media you want to create.
- Reference media is relevant and clear, and you plan to review and refine the result on the canvas.

5 Limitations

5.1 Known limitations

- Output can be inaccurate, low quality, or unexpected, and it can vary between runs. Review every result before use.
- Data handling and safety behavior vary by the selected model and its verification status. An unverified model follows provider terms rather than Figma contract terms.
- Model availability and behavior can change when a provider updates or retires a model. This can affect a saved workflow and the result it produces.

5.2 Factors that affect output quality

- The specificity and scope of your prompt.
- The model and parameters selected for each node.
- The quality of reference media and upstream node results.

5.3 Best practices for performance

- Review every result before using or publishing it, especially realistic media.
- Keep prior outputs as separate nodes, and save or version a workflow before rerunning it.
- Use Enterprise model restrictions to limit a workspace to options your organization has approved.

6 Your data and privacy

6.1 Data that is Processed

DATA	WHAT IT INCLUDES
Your prompt	The text instruction provided to a workflow node.
Reference media	Images, video, and audio uploaded for the selected model to use.
Model and parameter selections	The model and settings selected for each node.
Workflow inputs and outputs	Generated images, video, audio, and 3D assets passed between nodes or returned to the canvas.

6.2 Who at Figma can access your data

Access follows your Weave account and workspace permissions, and Figma confirms your identity and account entitlements before a request proceeds.

Figma limits internal access to staff who operate and support the service on a need to know basis. Role based controls restrict access, and access is logged.

6.3 Who at our model providers can access your content

Provider staff access varies by provider and hosting path. Most generation paths run without provider staff review, while retained aggregator content can be visible to authorized account members. The visual moderation provider also supports human review and may report child safety content to the relevant authorities.

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 Weave content in the United States. The data residency option for the main Figma platform does not apply to standalone Weave.

6.4.2 Where the model providers store your content

Provider storage and processing locations depend on the model and hosting path you select. Confirmed processing locations include the United States, the European Union, Indonesia, Malaysia, Belgium, the Netherlands, Spain, Singapore, Israel, and the United Kingdom.

Some provider paths do not store prompts or outputs. Other paths retain prompts, reference media, or results temporarily, including about one hour on Replicate, up to 30 days on FAL, and 14 days by default for Hive AI.

6.5 Retention and deletion

6.5.1 What is retained

WHAT IS RETAINED	PURPOSE	RETENTION PERIOD
Canvas content, including prompts, uploaded media, and generated outputs	Store your work so you can return, rerun, and refine it.	For the life of the account or until deletion
Account, workspace, credit, and billing data	Operate the account, credit system, and billing.	For the life of the account or until deletion
Usage and operational telemetry	AI billing, reliability, abuse prevention, and product analytics.	For the life of the account or until deletion

6.5.2 Deleting your data

You can delete canvas content or close your account, and Figma deletes the associated canvas content and account data from the Weave backend. Provider content subject to temporary retention ages out under each provider's schedule. Some provider configurations retain nothing after the request is served.

6.6 Training on your content

6.6.1 Whether the model providers train on your content

Providers never train on your data. In Weave, this protection applies to verified models that run under Figma contracts limiting content use to delivering the service. Unverified models follow their provider terms, so their training practices may differ.

6.6.2 Whether Figma trains on your content

Figma does not use Weave customer content to train or fine tune its own models. Weave runs no first-party Figma model, and the main Figma platform content training setting does not apply to this standalone product.

6.7 Sensitive and confidential data

Weave does not require or solicit sensitive demographic data. Free form prompts and uploaded media can still contain personal, confidential, regulated, or sensitive information, including images of people. Apply your organization's policies and confirm that you have the right to use the content you provide.

7 Content moderation

7.1 What is screened

Text prompts are screened on configured routes before normal generation. Uploaded and generated images and video can be screened through a separate visual moderation path. The generation providers also apply their own safety classifiers and refusal behavior to the content they produce.

7.2 What is blocked

Prompts and generated results are blocked when they violate Figma's acceptable use policy. Examples include the following.

- 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 inputs and outputs. Provider safety classifiers can also refuse content under their own policies.

7.3 What happens when content is blocked

When prompt moderation blocks a request, the result is not generated. Weave shows a message that the content safety system flagged the prompt and asks you to revise it.

7.4 Whether moderation can be edited or configured

No. End users cannot adjust moderation thresholds or turn the moderation chain off. Enterprise model restrictions can narrow the available model set, but they do not relax moderation controls.

8 Safeguards and security

8.1 Security controls

- Traffic to model providers and aggregators is encrypted in transit with HTTPS and Transport Layer Security.
- Production data at rest is encrypted, and encryption keys are managed separately from provider credentials.
- Requests enter through a controlled gateway that applies routing, moderation entry, structured logging, rate limits, and workspace quotas.
- Internal data stores use role based access controls and network restrictions, and service access is logged.
- Uploaded and generated media is validated against expected types, formats, and sizes before it is used or returned.
- Feature controls can disable a model or the service when Figma responds to a security or reliability issue.

8.2 Administrator controls

- Enterprise administrators can approve or restrict models available to their workspace.
- Enterprise administrators can require approval before an unverified model runs.
- Some providers support a customer supplied provider key, so requests run under the customer's provider agreement.
- Enterprise administrators can allocate credits for their workspace.

9 Performance, evaluation, and fairness

9.1 How we evaluate this feature

Figma evaluates Weave through operational monitoring, provider health signals, model verification status, customer reports, and incident review. Creative output has no single correct answer, so you judge generation quality on the canvas against your prompt, reference media, and workflow. Figma can restrict or remove a model and can disable the service when a serious safety or reliability issue is identified.

9.2 What we measure

MEASURE	WHAT IT MEASURES	WHY IT MATTERS TO YOU
Generation success rate	How often a request completes and returns a usable result.	A decline can reveal a provider, model, or workflow problem.
End to end latency by modality	How long a node takes to return image, video, audio, or 3D output.	A change can reveal degraded provider performance.
Provider error and availability rates	Errors and service availability reported across provider paths.	These signals help Figma respond when a provider is degraded.
Active high severity issues	Priority zero and priority one issues open for the service.	These issues receive the fastest operational response.

9.3 Fairness and bias

9.3.1 What Figma reviewed

Figma uses model verification, content moderation, customer reports, and human review of each result to identify harmful or biased output. Figma has not run a centralized quantitative bias evaluation across the third-party models it does not control.

9.3.2 What the model providers reviewed

Major providers publish general bias and safety evaluations. OpenAI and Google describe bias testing and harmful content filtering, and Anthropic describes political bias and child safety evaluations. Disclosure for other models in the library is limited and varies by provider.

9.3.3 Known differences in performance

Performance varies by model, language, script, cultural context, and type of media. Many image and video models perform best in English and Western contexts and can render non-Latin text poorly. Models also vary in how they depict different demographic groups, and Figma does not certify performance centrally across the library.

9.4 Evaluate it for your own use

Test Weave with your own prompts, reference media, workflows, and reviewers before depending on it. Compare models, review verification status, and decide which outputs require added review before use. The same workflow can produce different results, and exact reproduction depends on the selected model because seed control is not available everywhere.

10 Human oversight, feedback, and control

10.1 Human oversight

Weave keeps a human in the loop. You trigger every generation, review the result on the canvas, and decide whether to keep, refine, rerun, export, or discard it. The service takes no autonomous action and publishes nothing on its own.

Weave can generate realistic synthetic media. Review the content and its context before sharing or publishing it, especially when an image, video, or audio result could be mistaken for a real person or event.

10.2 Reporting problems and giving feedback

- Report a failed, unsafe, poor, or incorrectly blocked result through in product support or the Weave help and community channels.
- The team reviews reports and can adjust the product, moderation behavior, or model library in response.

10.3 Turning the feature off

Weave is an opt-in AI product, so an individual can choose not to use it. Enterprise administrators can restrict which models are available, but there is no separate control that turns the canvas into a non-AI product. Figma can disable a model or the service if it is not fit for its intended use.

11 Transparency

11.1 How AI involvement is disclosed

Weave is presented as an AI product, and model selection and generation happen through visible workflow nodes. You deliberately start each generation, and the canvas shows the prompt, uploaded inputs, upstream nodes, and selected model that contributed to an output.

11.2 How we communicate changes

Figma communicates material changes to functionality, intended use, or reliability through public documentation, Weave help and community channels, release notes, in product messaging, and administrator communication.