Figma

AI PRODUCT SYSTEM CARD

Figma Weave API

LAST UPDATED DATE, 10/13/26

PRODUCT STATUS	AVAILABLE ON THE FOLLOWING CUSTOMER TIERS	WHAT IT DOES
Generally available	Enterprise Weave customers	Figma Weave API gives developers programmatic access to the Weave generation engine. It creates images and video from prompts, reference media, and saved 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 API 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 main product, model, data, safety, and control details described in this system card.

What it does	Generates images and video through an API from prompts, reference media, and saved workflows.
Who it is for	Invited Enterprise developers and teams using Weave plans.
Where it appears	In applications and automated pipelines that call the API.
Availability	Available to all Weave Enterprise workspace admins.
Model selection	The caller selects from the verified generation models that Figma makes available.
Data processing and storage	Figma stores Weave API data in the United States. Provider processing and storage depend on the selected delivery path.
Training on your content	Verified generation providers are contractually prohibited from training on customer content. Figma does not train any model on Figma Weave API content.
Content moderation	Figma screens text prompts before generation when the service moderation control is active.
Human oversight	The calling application determines whether a person reviews generated media before use.
FigGov Experience	Not available.
Gov+ Experience	Not applicable to the separate Weave plan structure.
EDU Experience	Not available.

2 What it does and how it works

2.1 Description

Figma Weave API gives developers programmatic access to the Weave generation engine. A calling application sends a prompt, optional reference media, a model choice, and request parameters. The API runs one generation step or a saved workflow and returns an image or video to the calling application.

2.2 What the AI generates or predicts

The API generates images and video. Results can vary between models and between runs, even when a request is similar. The API generates creative media rather than decisions or predictions about people.

2.3 How it works with other Figma products

The API exposes the Weave generation engine outside a Figma editor. Generated media returns to the calling application rather than a Figma file. Weave uses separate plans and credits, and the Figma Enterprise data residency solution does not apply. Sections 6.2 and 8.2 describe access and administrator controls.

3 Models and providers

3.1 Models and where they run

Figma Weave API uses third party image and video generation models and separate text moderation models. The generation models create media from prompts and reference media, and similar requests can produce different results.

MODEL AND DEVELOPER	ROLE IN THIS FEATURE	WHERE IT RUNS
Verified image and video models, multiple developers	Generate images and video. Families include Flux, Veo, Kling, Seedance, Seedream, and Ideogram.	FAL, Replicate, and direct provider infrastructure.
Moderation models, OpenAI and Amazon	Screen text prompts before generation when the service moderation control is active.	OpenAI and AWS Bedrock.

3.2 Can I select the models?

Yes. The caller selects a generation model from the verified set that Figma makes available. Figma determines that set, and moderation runs on models that Figma configures regardless of the generation model selected. Additional selectable models include: Black Forest Labs FLUX Pro, Google Veo, Kuaishou Kling, ByteDance Seedance Pro, ByteDance Seedream, and Ideogram Ideogram.

3.3 What the providers do with your content

Providers receive your prompt, reference media, and generated results only to perform generation or safety screening. Figma contracts bar verified generation providers from training on customer content, and those providers retain content only under the delivery path periods described in Section 6.5. Moderation providers screen text under zero data retention and do not train on it.

3.4 How a model change reaches you

Figma pins the verified set to specific model versions and compares quality and safety before adding or changing a version. Figma monitors provider changes and deprecations, and it can remove or restrict a model that regresses. The verified model library can therefore change over time.

4 Intended use

4.1 Primary intended users

The API is intended for developers and teams that integrate creative generation into applications, products, pipelines, and automated processes. During the closed beta, access is limited to invited Enterprise Weave customers.

4.2 Primary intended uses

- Generate images and video from text prompts and reference media.
- Run saved Weave workflows from an application or automated pipeline.
- Generate and edit creative media through model specific request parameters.

4.3 Where it works best

- The request has a clear creative objective and acceptance criteria.
- The selected model supports the requested media type and task.
- The calling application reviews results and handles errors, rate limits, and provider changes.

5 Limitations

5.1 Known limitations

- Generated media can be inaccurate, low quality, biased, or unexpected. Review it before it enters a product or automated process.
- The same or similar request can produce different results across models and runs, so exact reproduction is not guaranteed.
- The verified model library changes over time. A model can change behavior or become unavailable after provider or Figma updates.

5.2 Factors that affect output quality

- The generation model used for the task, because image and video families have different capabilities.
- The prompt, reference media, and request parameters supplied, which define the creative direction.
- The provider version and hosting path active for the request, which can affect latency and results.

5.3 Best practices for performance

- Define acceptance criteria for the intended use. Review every output against them before downstream use.
- Store and version any result you want to reuse. A later generation from the same request can differ.
- Test error handling and rate limits before automating generation at scale.

6 Your data and privacy

6.1 Data that is Processed

DATA	WHAT IT INCLUDES
Your prompt	The text instruction sent by the calling application.
Reference media	Images or video supplied with the request.
Workflow and model settings	The saved steps, verified model, and generation settings selected for the request.

6.2 Who at Figma can access your data

Figma limits access to staff who operate and support the service. Role based controls require a business need, and access is logged. Other staff cannot browse customer content.

6.3 Who at our model providers can access your content

Providers receive your prompt, reference media, and generated results only to perform generation or safety screening. Moderation provider staff do not review content screened under zero data retention. On generation paths, access varies, and FAL team members with account access can view prompts and results.

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 Figma Weave API data in the United States. The API runs outside the Figma Enterprise data residency solution.

6.4.2 Where the model providers store your content

OpenAI and Amazon do not store prompt content on the moderation paths. FAL processes generation requests in the United States and may use other countries, and it stores request content under its configured retention period. Replicate removes prediction content and logs after one hour by default.

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 content or close your account to remove associated request content and account data from the Weave service. You can also submit a deletion request through support. Provider content ages out under the configured period for the selected delivery path.

6.6 Training on your content

6.6.1 Whether the model providers train on your content

Providers never train on your data. Figma contracts bar verified generation providers from training on customer content, and moderation providers screen text under zero data retention and do not train on it.

6.6.2 Whether Figma trains on your content

Figma does not use Figma Weave API content to train any model. The standard Figma content training setting does not apply in the same form because Weave uses separate plans and Figma does not train a model on this content.

6.7 Sensitive and confidential data

Apply your own policies for confidential, regulated, and personal data. The API accepts free form prompts and reference media, which can include personal or sensitive material even though the feature does not require it.

7 Content moderation

7.1 What is screened

Figma screens text prompts before normal model execution when the service moderation control is active. The screen covers the text instruction sent with an API request.

7.2 What is blocked

Our acceptable use policy prohibits the following examples. When the service moderation control is active, Figma blocks text prompts in these categories before generation.

- 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 individuals or protected classes.
- 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.

7.3 What happens when content is blocked

When prompt moderation blocks a request, the API returns a content safety error before generation and asks the caller to revise the prompt.

7.4 Whether moderation can be edited or configured

Callers cannot change moderation provider settings through an API request. Figma configures moderation controls at the service level.

8 Safeguards and security

8.1 Security controls

- API access requires an authenticated token, while storage and compute remain restricted to approved service paths.
- Traffic to the API and providers is encrypted in transit, and stored data is encrypted at rest with managed keys.

– Provider credentials and service secrets are held in an approved secrets manager rather than in code or client responses.
– Access and request logs record the actor, action, and time for security review.
– Rate limits, quotas, and throttling protect the service against abuse and overload.

8.2 Administrator controls

– Enterprise Weave administrators can approve or restrict the generation models available to their teams.
– Enterprise Weave customers can supply their own provider key and can revoke API access.
– Figma Weave API is not available through Figma for Government or Figma for Education, and Governance Plus does not apply to the separate Weave plan structure.

9 Performance, evaluation, and fairness

9.1 How we evaluate this feature

Figma pins verified model versions and compares quality and safety before a change reaches you. Figma monitors request success, errors, latency, provider availability, and sampled output quality against task specific review criteria. Safety, security, and core availability failures block release, and production issues are handled through the incident process.

9.2 What we measure

MEASURE	WHAT IT MEASURES	WHY IT MATTERS TO YOU
Request success and errors	Whether API requests complete and what type of failure occurs.	A change can reveal a provider or platform problem.
Generation latency and provider availability	How long generation takes and whether the selected provider is available.	A degradation can delay or interrupt an automated pipeline.
Sampled quality review	Whether sampled media follows the prompt, avoids unrequested content, and is technically usable.	It catches quality problems that operational measures cannot show.

Figma uses internal baselines for these measures and reviews the affected model or provider when performance degrades.

9.3 Fairness and bias

9.3.1 What Figma reviewed

Figma uses red teaming with demographic and sensitive prompts and sampled human review of generated people and scenes for stereotyping. Sampled reviews also include non English prompts and varied cultural subjects. Findings are logged and routed to the product team.

9.3.2 What the model providers reviewed

Figma relies on the bias evaluations and harmful content filtering disclosed by the model providers. Disclosure depth varies across the verified model set, and some providers publish formal evaluations while others provide less detail.

9.3.3 Known differences in performance

Output quality can vary across demographic groups because generative media models can reflect bias in their training data. Figma includes non English prompts and varied cultural subjects in sampled reviews.

9.4 Evaluate it for your own use

Before you depend on Figma Weave API, test it with your own prompts, reference media, models, and reviewers. Decide which outputs require human review before they enter a product or automated process. Preserve preferred outputs because similar requests can produce different results, and fixed seed support is not guaranteed across the verified set.

10 Human oversight, feedback, and control

10.1 Human oversight

The calling application controls the level of human oversight. A typical workflow has a person design the workflow, start generation, and review the returned media before use. The API can also run inside an automated pipeline, so the customer determines when human review is required.

10.2 Reporting problems and giving feedback

- Report failures, unsafe output, and other concerns through Weave product support or developer channels.
- Enterprise customers can also raise concerns through their account support channel.
- Figma reviews reported issues and can adjust the verified model set, routing, or moderation in response.

10.3 Turning the feature off

A team can stop calling the API, and an Enterprise administrator can revoke API access. Figma can remove a model from the verified set or disable the service while a material risk is assessed.

11 Transparency

11.1 How AI involvement is disclosed

Figma Weave API is presented to developers as an AI generation service. It returns generated media to the calling application, which controls how that media is presented and labeled to its audience.

11.2 How we communicate changes

Figma communicates material changes to functionality, intended use, models, or reliability through developer documentation, release notes, and administrator communication.