Figma

AI PRODUCT SYSTEM CARD

Figma Make

LAST UPDATED DATE, 10/13/26

PRODUCT STATUS	AVAILABLE ON THE FOLLOWING CUSTOMER TIERS	WHAT IT DOES
Generally available	All plans. Capability varies by seat.	Figma Make turns a prompt, an existing design, or an existing codebase into a working app or interactive prototype. It generates the interface, functionality, and code inside Figma, and you keep refining the result by describing what to change.

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 Make 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 gives the short answer for each part of this system card.

What it does	Turns a prompt, design, or codebase into a working app or interactive prototype you can refine.
Who it is for	Anyone with a Figma seat. K-12 accounts do not have access.
Where it appears	Its own surface in Figma, and through the Figma MCP server.
Availability	All plans. A Full seat is required for private sharing.
Model selection	You pick the build model from a set that Figma approves.
Data processing and storage	Figma stores your content in the United States. Model providers do not store it at rest.
Training on your content	Providers never train on your content. Whether Figma does is an account setting that varies by plan. Section 6.6 describes the setting.
Content moderation	Input moderation runs on configured paths. Model providers apply safety classifiers to generated output.
Human oversight	You start each build and review the result before keeping, publishing, or sharing it.
FigGov Experience	Available in a FedRAMP-authorized environment. The organization's AI control is off by default.
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, but not to K-12 users.

2 What it does and how it works

2.1 Description

Figma Make turns a prompt, an existing Figma design, or an existing codebase into a working interface or small application. You describe what to build, and Make generates the code, layout, and content for a runnable result that you can preview, refine, and share. Each follow-up instruction updates the existing build and keeps its context rather than starting over. Nothing is published or shared until you choose.

2.2 What the AI generates or predicts

Make generates application code, files, layout, content, and the edits requested in follow-up instructions. Supporting steps generate a session title, description, and filename, expand a rough prompt, and select a preview platform such as mobile, tablet, or desktop.

2.3 How it works with other Figma products

Figma Make has its own surface in Figma, and it can draw on designs, a design system delivered as a Make kit, a guidelines file, pasted reference text and images, and data files such as CSV and JSON. The result is stored as a Make file, and Figma Make can also be driven through the Figma MCP server.

3 Models and providers

3.1 Models and where they run

Figma Make runs on large language models built by Anthropic, Google, and OpenAI, along with a model that Figma builds and maintains. The same prompt can produce different results, and output can be inaccurate or incomplete. Every result is presented for your review.

MODEL AND DEVELOPER	ROLE IN THIS FEATURE	WHERE IT RUNS
Claude Sonnet, Anthropic	A model family for the main build. Turns a prompt into a working interactive prototype with live code.	Anthropic, AWS Bedrock, Google enterprise platform
Claude Opus, Anthropic	Selectable alternative for complex builds. It also plans and carries out multistep build actions.	Anthropic, AWS Bedrock, Google enterprise platform
Gemini Pro and Flash, Google	Selectable alternatives for the main build. Flash is the faster and lighter option.	Google enterprise platform
GPT, OpenAI	Selectable alternative for the main build.	OpenAI
Fast Apply, Figma	Applies small code edits to an existing build.	Fireworks
Supporting models, Anthropic and Google and OpenAI	Generate titles, descriptions, and filenames, expand prompts, select the preview platform, and support agent actions.	Anthropic, Google enterprise platform, OpenAI, Microsoft Azure OpenAI
Moderation models, OpenAI and Amazon	Apply content screening on assistant paths configured for moderation before model processing.	OpenAI, Microsoft Azure OpenAI, AWS Bedrock

3.2 Can I select the models?

Partly. You pick the build model from a set that Figma approves, and a new build starts on a Figma-set default that can vary by plan. Figma sets the available models and the default, while supporting steps and moderation run on

models that Figma configures regardless of the build model you pick. Additional selectable models include: Anthropic Claude Sonnet, Anthropic Claude Opus, Google Gemini Pro, Google Gemini Flash, OpenAI GPT Sol, and OpenAI GPT Terra.

3.3 What the providers do with your content

None of the providers train on your content. They receive your prompt and build context only to process your request and apply safety controls, and they are contractually prohibited from using it to train or improve their models. Each provider runs under zero data retention or an equivalent no-logging configuration, so content is not kept after your request is served.

3.4 How a model change reaches you

Make runs on a curated set of approved models, so a provider version reaches you only when Figma changes the approved set or default. Figma pins versions, validates a new version through canary checks and staged rollouts, compares output across versions, and maintains a tested rollback.

4 Intended use

4.1 Primary intended users

Figma Make is for designers, product managers, engineers, and other people with a Figma seat who want to turn an idea into a working interface or application. A Full seat unlocks publishing and private sharing. View, Collab, and Dev seats can build Make files in drafts. Higher education and bootcamp users have access, while K-12 users do not.

4.2 Primary intended uses

- Turn an idea into a clickable prototype that your team can review.
- Explore interface variations by refining the same build through follow-up instructions.
- Turn a written specification, design, codebase, or data file into something that runs and can be tested.

4.3 Where it works best

- The prompt is specific about what the result should do and about constraints such as framework or visual style.
- You attach the materials the build should follow, such as designs, a Make kit, a guidelines file, reference images, or data files.
- Follow-up instructions are clear and scoped, and you plan to iterate rather than take the first build as final.

5 Limitations

5.1 Known limitations

- Generated code is not guaranteed to be correct or secure, and it needs review before you rely on it.
- Quality varies for complex logic, precise visual layout, and uncommon frameworks. A build can contain bugs, inefficient patterns, or choices that do not match your standards.
- Make is not aware of your private codebase unless you provide it, so generated code may not follow your existing conventions.

5.2 Factors that affect output quality

- How specific your prompt is about the result and its constraints.
- Whether you provide the designs, guidelines, images, code, and data that the build should follow.
- How much you ask for in one request. Long-running builds can reach infrastructure limits and slow generation.

5.3 Best practices for performance

- Use Point and Edit to direct one instruction at one element instead of asking Make to regenerate the whole build.
- Review the generated code and live preview before relying on a result, and apply your normal code and security review practices.
- Keep useful build versions in the chat so you can return to them after later changes.

6 Your data and privacy

6.1 Data that is Processed

DATA	WHAT IT INCLUDES
Your instructions	The prompts and follow-up instructions you give Make.
The current build	The generated code and files in the session that each instruction can update.
Materials you provide	Designs, code, guidelines, pasted text, reference images, and data files such as CSV and JSON.

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 Figma for data kept outside your files. Figma propagates the request across the systems that hold AI evaluation and usage data. Your Make files, including their prompts, builds, and outputs, are kept for the life of the account or until 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

Make does not require special category personal data and does not ask for it. Keep this data out of prompts and attachments. 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

Figma screens prompts before model processing on assistant paths configured to create a moderation task. Figma validates that generated code and edits match the expected shape before applying them to your build. Model providers apply their own safety classifiers and refusal behavior to generated output.

7.2 What is blocked

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

- Child sexual abuse material, including depictions, simulations, 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.

The policy applies to prompts and generated output. The model providers in Section 3 also apply their own safety classifiers and refusal behavior.

7.3 What happens when content is blocked

Figma shows a message that the request may violate its guidelines. You can dismiss the message, but the request is not completed. Report a block that looks wrong through the channels in Section 10.

7.4 Whether moderation can be edited or configured

End users cannot turn off content moderation or adjust its settings. There is no in-product control that changes them.

8 Safeguards and security

8.1 Security controls

- Traffic to model providers is encrypted in transit with HTTPS and Transport Layer Security. Requests to managed hosting paths 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 held separately.
- Every AI request passes through Figma's AI gateway for routing, rate limiting, structured logging, and alerting.
- Generated code and edits are validated before they are applied. Agent actions use scoped permissions, and generated code runs in a sandboxed preview rather than against Figma systems.
- Per-user and per-organization rate limits and quotas protect against abuse and overload.
- Internal data stores are isolated and reachable only through approved roles and network policies. Internal access requires approval and is audited.

8.2 Administrator controls

- Enable or disable AI features, including Make, at the organization or team level.
- Manage the content training setting described in Section 6.6. It is a team-level setting on Starter and Professional, and it is disabled on Organization and Enterprise.
- On Enterprise with the Governance+ add-on, manage or restrict AI traffic across the plan. Restricting traffic can make some AI behavior unavailable.
- On Figma for Government, Make runs in a FedRAMP-authorized environment, and the organization's AI control is off by default.

– On education plans, Make is available to higher education and bootcamp users, but not to K-12 users. Education content is never used for AI training.

9 Performance, evaluation, and fairness

9.1 How we evaluate this feature

Figma validates new approved model versions before they reach you, as described in Section 3.4. Systemic quality regressions, elevated policy violations, repeated timeouts, and defects that stop a build from rendering or produce unsafe code block a release.

In production, Figma monitors build success, time to first token, generation latency, client crashes, incidents, and provider telemetry. A critical regression triggers incident response, which can include rollback or a switch to another provider. Other issues are reviewed with product and platform owners.

9.2 What we measure

METRIC	WHAT IT MEASURES	WHY IT MATTERS TO YOU
Build success rate	How often a request produces a build that renders and runs.	A decrease can show a problem with a model or hosting path.
Time to first token and end-to-end latency	How quickly a build starts and how long it takes to produce a rendered result.	Long builds are more likely to reach infrastructure limits.
Client crash rate	How often a Make action causes the client session to crash.	A change can reveal a product regression that affects your ability to complete a build.

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 reviewed build quality and used red teaming to look for harmful behavior. Make mainly produces code and interface structure, and the area where bias is most relevant is generated text such as sample names, titles, and placeholder biographies. This content can reflect stereotypes when a prompt is neutral, so review it before presenting it to an audience.

9.3.2 What the model providers reviewed

Figma relies on the bias evaluations published by the model providers for behavior inside their models. The providers describe filtering harmful and biased content, alignment work, multilingual testing, external red teaming, and evaluations for stereotyping and ungrounded inferences about people.

9.3.3 Known differences in performance

Build quality varies by the language of the prompt and by framework, depending on provider support. The providers evaluate the underlying models as multilingual systems. Figma is not aware of systematic performance differences by demographic group or region.

9.4 Evaluate it for your own use

Before you depend on Make, run a pilot with your own prompts and reviewers, and decide which outputs require review. Treat code you intend to ship as a first draft subject to your normal code and security review practices.

The chat history records which prompt produced each build, the generated code is open in the editor, and Point and Edit directs a follow-up instruction at one element. Each build is saved as a version in the chat. Make does not provide a formal source map, file-level version history, or a deterministic replay control.

10 Human oversight, feedback, and control

10.1 Human oversight

You start each Figma Make build or edit and review the result. The Make agent can write code, install packages, run the application, open a browser, and fix issues during the build without separate approval for every step. You decide whether to keep, change, publish, or share the result.

Make does not make or enforce decisions about people, restrict accounts, or report users to external parties. Only signed-in Figma users with access to Make can use the feature. A build you publish remains subject to Figma's usage policies.

10.2 Reporting problems and giving feedback

- Report a failure, unsafe output, or 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. Feedback can inform product and safety updates.

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

Make is presented as a distinct, named AI feature. You open it deliberately, and the prompt surface is labeled as AI and is not named or styled to suggest a real person.

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.