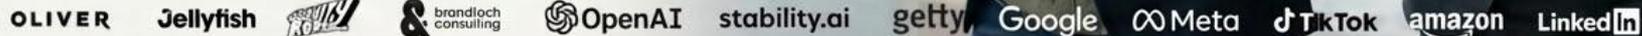


# Agent Builder Playbook

Turning Custom Agents into Scalable Workflows



# Contents

1. Welcome
2. What is the agent builder?
3. Core components of an agent
4. The tool ecosystem
5. Third party integrations
6. Best practice
7. Practical examples

The screenshot displays the 'Agents' interface with a search bar and two buttons: 'Import agent' and 'New agent'. The interface contains nine agent cards, each with an icon, a title, a description, and the creator's name 'By Pencil'.

Agent Name	Description
<b>Video Cutdown</b>	Generate short video cutdowns from long videos
<b>Video Generation</b>	Easily produce compelling short video clips from text or images.
<b>Image Generation</b>	Automatically generate visually appealing images for your brand or products.
<b>Static Social Ads</b>	Generate ready-to-market static ads. This agent combines your messaging and visuals with brand knowledge to create fully designed ads...
<b>Campaign Benchmarks</b>	Compare campaign performance across channels and categories to set clear targets and track competitiveness over time.
<b>Campaign Insights</b>	Identify key takeaways from your marketing campaigns by analyzing performance drivers, audience engagement, and conversion impact.
<b>Creative Benchmarks</b>	Compare creative performance against industry or brand-specific benchmarks to understand where your ads stand and how to improve them.
<b>Creative Insights</b>	Quickly surface high-performing patterns in creative assets to inform design decisions and optimize ad effectiveness.
<b>Marketing Calendar</b>	Easily plan and manage strategic, cross-channel marketing schedules.

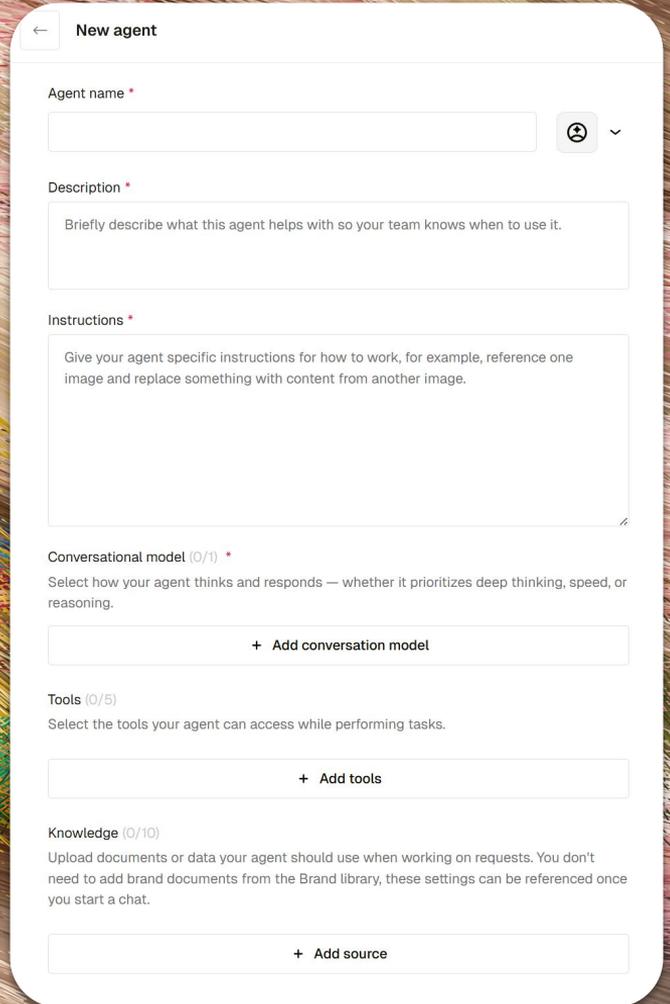
# Welcome to the playbook

As teams scale their use of AI, the challenge shifts from *creating outputs* to *creating reliable systems*. The Agent Builder is designed to move organisations beyond one-off prompts and single-purpose tools, and toward **repeatable, governed AI workflows**.

With the Agent Builder, teams can:

- Design AI agents that mirror real operational workflows, not isolated tasks
- Combine research, reasoning, creation, and distribution in a single assistant
- Reduce long-term maintenance through automatic model upgrades
- Enable organisation-wide reuse through centralised governance

This shift matters because AI value compounds when workflows are standardised. Instead of rebuilding prompts, reconnecting tools, or re-training teams, high-quality agents can be built once and scaled safely across teams, markets, and use cases.



The screenshot shows the 'New agent' configuration page. It features a back arrow and the title 'New agent'. The form is divided into several sections: 'Agent name' with a text input and a dropdown icon; 'Description' with a text area containing the placeholder 'Briefly describe what this agent helps with so your team knows when to use it.'; 'Instructions' with a text area containing the placeholder 'Give your agent specific instructions for how to work, for example, reference one image and replace something with content from another image.'; 'Conversational model (0/1)' with a text area containing the placeholder '+ Add conversation model' and a sub-label 'Select how your agent thinks and responds — whether it prioritizes deep thinking, speed, or reasoning.'; 'Tools (0/5)' with a text area containing the placeholder '+ Add tools' and a sub-label 'Select the tools your agent can access while performing tasks.'; and 'Knowledge (0/10)' with a text area containing the placeholder '+ Add source' and a sub-label 'Upload documents or data your agent should use when working on requests. You don't need to add brand documents from the Brand library, these settings can be referenced once you start a chat.'

\* The goal of this playbook is to provide clear, reusable, and practical guidance for designing, building, and governing AI agents using the Agent Builder. It offers recommended frameworks, best practices, and guardrails. This gives teams the flexibility to adapt agents to their specific workflows, tools, and organisational needs, so they can fully unlock the value of scalable AI within Pencil.

# The agent builder

An overview

# What is agent builder?

The Agent Builder is a powerful, upgraded experience that allows you to create custom, multi-functional AI agents tailored to your exact workflow needs.

While pre-built agents (such as Image Generation or Video Generation agents) are useful starting points, the Agent Builder is where production-grade agents live.

It allows you to move beyond single-purpose tools and design assistants that handle entire workflows end to end.



## Landing Page Copy Expert

This agent specializes in creating high-converting landing page content for marketing campaigns. It applies proven copywriting and conversion rate optimization (CRO) principles to structure...

By Vanessa Gurgel



## Modelling shoot creator

An agent that takes a reference of a person and an item of clothing in order to generate a full modelling photoshoot in a studio setting.

By Timothy Bowers



## Oleander Tassel - The Magic AI Decor Guy

Your role is to help users visualize and reimagine drab, stale home interiors as the space they've always dreamed of. You take a picture of their homes, and use AI image magic to give them a...

By Michael Whyte



## Pencil Community Prompt Builder

A realism-first, production-focused prompt builder for Gemini 3 Nano Banana Pro that creates structured, copy-ready prompts for three workflows: image referencing, text-to-image, and image-to-...

By Layide Oshun



## Shakespeare Translator

This agent seamlessly transforms modern text into the language of William Shakespeare, offering a unique and comical historical perspective on any given text.

By Sarah Bradley



## Summer Campaign Idea Generator

Generates creative summer marketing campaigns for Ambrosia Ice Cream in the U.S. market. Provides taglines, key visuals, and suggested channels.

By Vanessa Gurgel

# What is agent builder?

With the Agent Builder, you can:

- Mix up to **5 different capabilities** in a single agent
- Combine **text, image, video, web, and 3rd-party tools**
- Build agents that research, create, and act in one continuous flow

This ensures:

- Fewer broken workflows
- Consistent outputs
- Controlled access to sensitive integrations



**Example:** One agent that:

- ✓ Searches the web for trends
- ✓ Writes a relevant, timely blog post
- ✓ Generates a matching header image

All within a single conversation.

# Core components of an agent

Every agent is built from the same foundational components.

## A. Basic Identity

- **Name:** Clear and recognisable (e.g. *SEO Blog Writer*)
- **Description:** A short explanation of what the agent does and who it's for

Keep this simple—users should understand the agent's purpose instantly.

## B. Conversational Model

You control the "brain" of your agent by selecting a model category:

- **Reasoning** – Complex logic, problem solving, coding
- **Deep Thinking** – Strategy, nuanced analysis, creative writing
- **Fast Response** – Quick Q&A, summaries, support
- **Multi-modal** – Tasks involving text + images/files



Agent name \*

Description \*

**Add conversation model**  
Select how your agent thinks and responds — whether it prioritizes deep thinking, speed, or reasoning.

Model thinking style

Reasoning  Deep thinking  Fast responses  Multi-modal input

# Best Practice: Use "Best Available"

Each category includes a **Best Available** option. This automatically upgrades your agent as better models are released—no manual updates required.

Model Type	Best Used for
Reasoning	Complex problem solving, coding, logic and planning.
Deep Thinking	Nuanced analysis, strategy formulation, and creative writing.
Fast Response	Quick Q&A, customer support, and summaries.
Multi-modal	Tasks requiring images/files and text outputs.

## Fast responses

Optimized for rapid brainstorming, instant headline generation, social media captions, and quick feedback loops.

Select model

Best Available



# Instructions & knowledge

## Instructions

This is the system prompt that defines:

- The agent's task
- Tone and personality
- Rules and constraints

Agents follow instructions literally! Be explicit about:

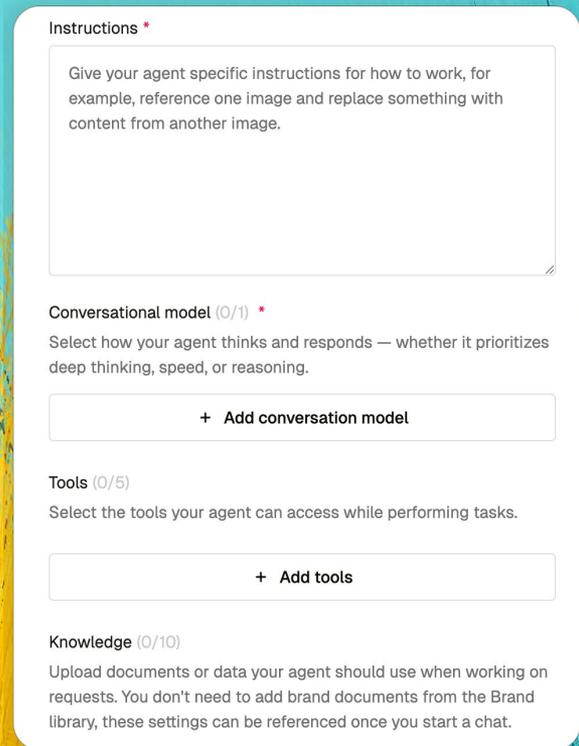
- Format
- Length
- Tone
- Dos and don'ts

## Knowledge Base

Upload documents the agent should reference *before* responding, such as:

- Brand guidelines
- Product documentation
- Internal playbooks

This ensures outputs stay grounded within approved contexts.



The screenshot shows a configuration panel with a white background and rounded corners, set against a teal and yellow abstract background. It contains four sections:

- Instructions \***: A text area with a light blue border containing the text: "Give your agent specific instructions for how to work, for example, reference one image and replace something with content from another image." A small blue icon is in the bottom right corner of the text area.
- Conversational model (0/1) \***: A text label above a text area containing: "Select how your agent thinks and responds — whether it prioritizes deep thinking, speed, or reasoning." Below this is a button with the text "+ Add conversation model".
- Tools (0/5)**: A text label above a text area containing: "Select the tools your agent can access while performing tasks." Below this is a button with the text "+ Add tools".
- Knowledge (0/10)**: A text label above a text area containing: "Upload documents or data your agent should use when working on requests. You don't need to add brand documents from the Brand library, these settings can be referenced once you start a chat."

# The tool ecosystem (new feature)

The Tool Ecosystem is the **biggest upgrade** in the new Agent Builder. You can now add **up to 5 tools** to a single agent.

## Category 1: Native AI Tools

Previously locked to pre-built agents, these are now modular:

### Image Tools

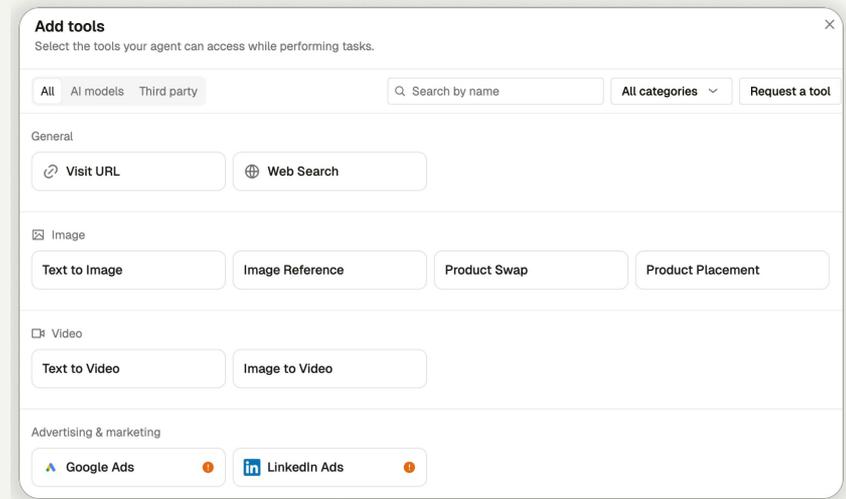
- Text-to-Image
- Image Reference
- Product Swap
- Packshot to Image

### Video Tools

- Text-to-Video
- Image-to-Video
- Video Cutdown

### Utilities

- Web Search
- Visit URL & Crawl



Just like when selecting a model, the tool selectors also default to **Best Available** model for each tool class.

# The tool ecosystem (new feature)

We'll be adding a slew of powerful new tools to our tool ecosystem in the coming weeks. As such, the list of available Native AI tools will change as we upgrade and expand the arrays of tools that your agents can utilise.

## Category 2: 3rd Party Tools

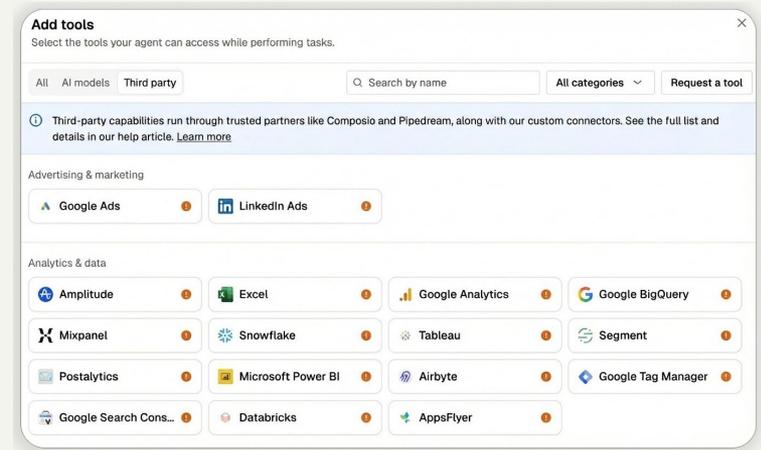
Connect agents directly to your existing software stack, including:

- Salesforce
- HubSpot
- Semrush
- Slack
- Google Drive

You select:

1. The external service
2. The specific actions/tools available

Behind the scenes, integrations are securely powered by Composio and Pipedream.



# Connecting & authenticating tools

Adding a tool is **not** the same as connecting it.

By default, 3rd party tools are disconnected.

## How to Connect

You can connect tools via the **Integrations Page** or directly inside the **Builder**.

- **One-Click Auth (OAuth):** For some tools like Gmail or Slack, simply click "Connect," log in to the service, and authorize access.
- **API Key/Token:** For some tools like Semrush or Asana, you may need to paste an API key provided by that service.

**Important:** If any tool is disconnected, the agent chat will be disabled until setup is complete. Please complete all your agent's tool integrations before conversing with your new agent.



Slack is a channel-based messaging platform. With Slack, people can work together more effectively, connect all their software tools and services, and find the information they need to do their best work — all within a secure, enterprise-grade environment.

Select capability

Set snooze duration

Q Search

Set snooze duration

DEPRECATED: Turns on Do Not Disturb mode for the current user, or changes its duration. Use `set_dnd_duration` instead.

Add a custom emoji to a Slack team

DEPRECATED: Adds a custom emoji to a Slack workspace given a unique name and an image URL. Use `add_emoji` instead.

Add an emoji alias

Adds an alias for an existing custom emoji in a Slack Enterprise Grid organization.

# Publishing & sharing agents

Once tested, agents can be published.

## Who Can Publish Agents?

- **Admins / Subscription Admins:** Publish to one or multiple workspaces
- **Standard Users:** Publish only to their own workspace

## Connection Carry-Over (Admins Only)

When publishing agents, admins can decide how **tool credentials** are shared:

- Carry over **All** connections
- Carry over **None**
- Carry over only **Specific** connections

Admins should therefore be mindful about whom they are sharing access to any sensitive authentication or API tokens with.

If a workspace already has an existing connection, that existing connection takes priority to avoid conflicts.

### Choose which connections to carry over ✕

Publishing an agent doesn't carry over its integrations unless you choose to. If you prefer to reuse them, you can pick which integrations you want to carry over to the target workspace.

Choose the tools to carry over:

-  LinkedIn Ads / Capability 1
-  Semrush / Capability 1
-  AppsFlyer / Capability 3

Cancel

Publish

# How it works at a glance

This table summarises how **native AI tools** vs **3rd party tools** behave across building, setup, publishing, and usage.

Scenario	Native AI Tools (Image / Video / Web)	3rd Party Tools (Salesforce / Slack / Semrush)
Building	Add and use immediately	Add now, connect later
Setup	No setup required	Admin authentication required
Publishing	Works instantly for all users	Admin can choose to carry over connections
Chatting	Always available	Chat locked until tools are connected

This distinction is critical when designing agents.

Native tools provide instant creative power, while 3rd party tools introduce automation and system integration, requiring stronger governance.

# Summary of UX scenarios

Scenario	User Experience
<b>Building</b>	You can now combine up to 5 tools. You can mix <b>Native AI tools</b> (like Image Gen) with <b>3rd Party tools</b> (like Slack). Native tools work instantly; 3rd party tools need connection.
<b>Chatting</b>	Chat is locked until all selected tools are authenticated.
<b>Publishing (Admin)</b>	You can choose to share your tool login with the team (Carry-over).
<b>Publishing (User)</b>	You share the Agent "blueprint." Teammates must log in to tools themselves.
<b>Importing</b>	You can copy an Agent ID from another workspace. You must check whether the agent's tools need to be reconnected (if they haven't been carried over).



# Third party integration

For connecting tools and apps

# Third party integrations: What & Why

## Purpose

3rd-party integrations allow agents to **take real-world actions**, not just generate responses — securely connecting Pencil agents to the tools your teams already use.

## What They Enable

- Read & write documents
- Fetch live data
- Update CRM records
- Trigger workflows
- Save outputs directly into business systems

## How It Works

- Pencil uses two authorised integration partners — **Composio** and **Pipedream** — as secure bridges between agents and external tools  
These partners act as the agent's “hands,” translating requests into governed API actions
- All actions:
  - ✓ Run only when explicitly requested by the agent
  - ✓ Respect workspace permissions and admin controls

# Third party integrations: Using & governing integrations

## Integration Roles

- **Composio** → Enables real-time agent actions (e.g. updating docs, sending messages, or modifying records)
- **Pipedream** → Powers workflow-driven and automation-based integrations across broader platforms

## Supported Ecosystem

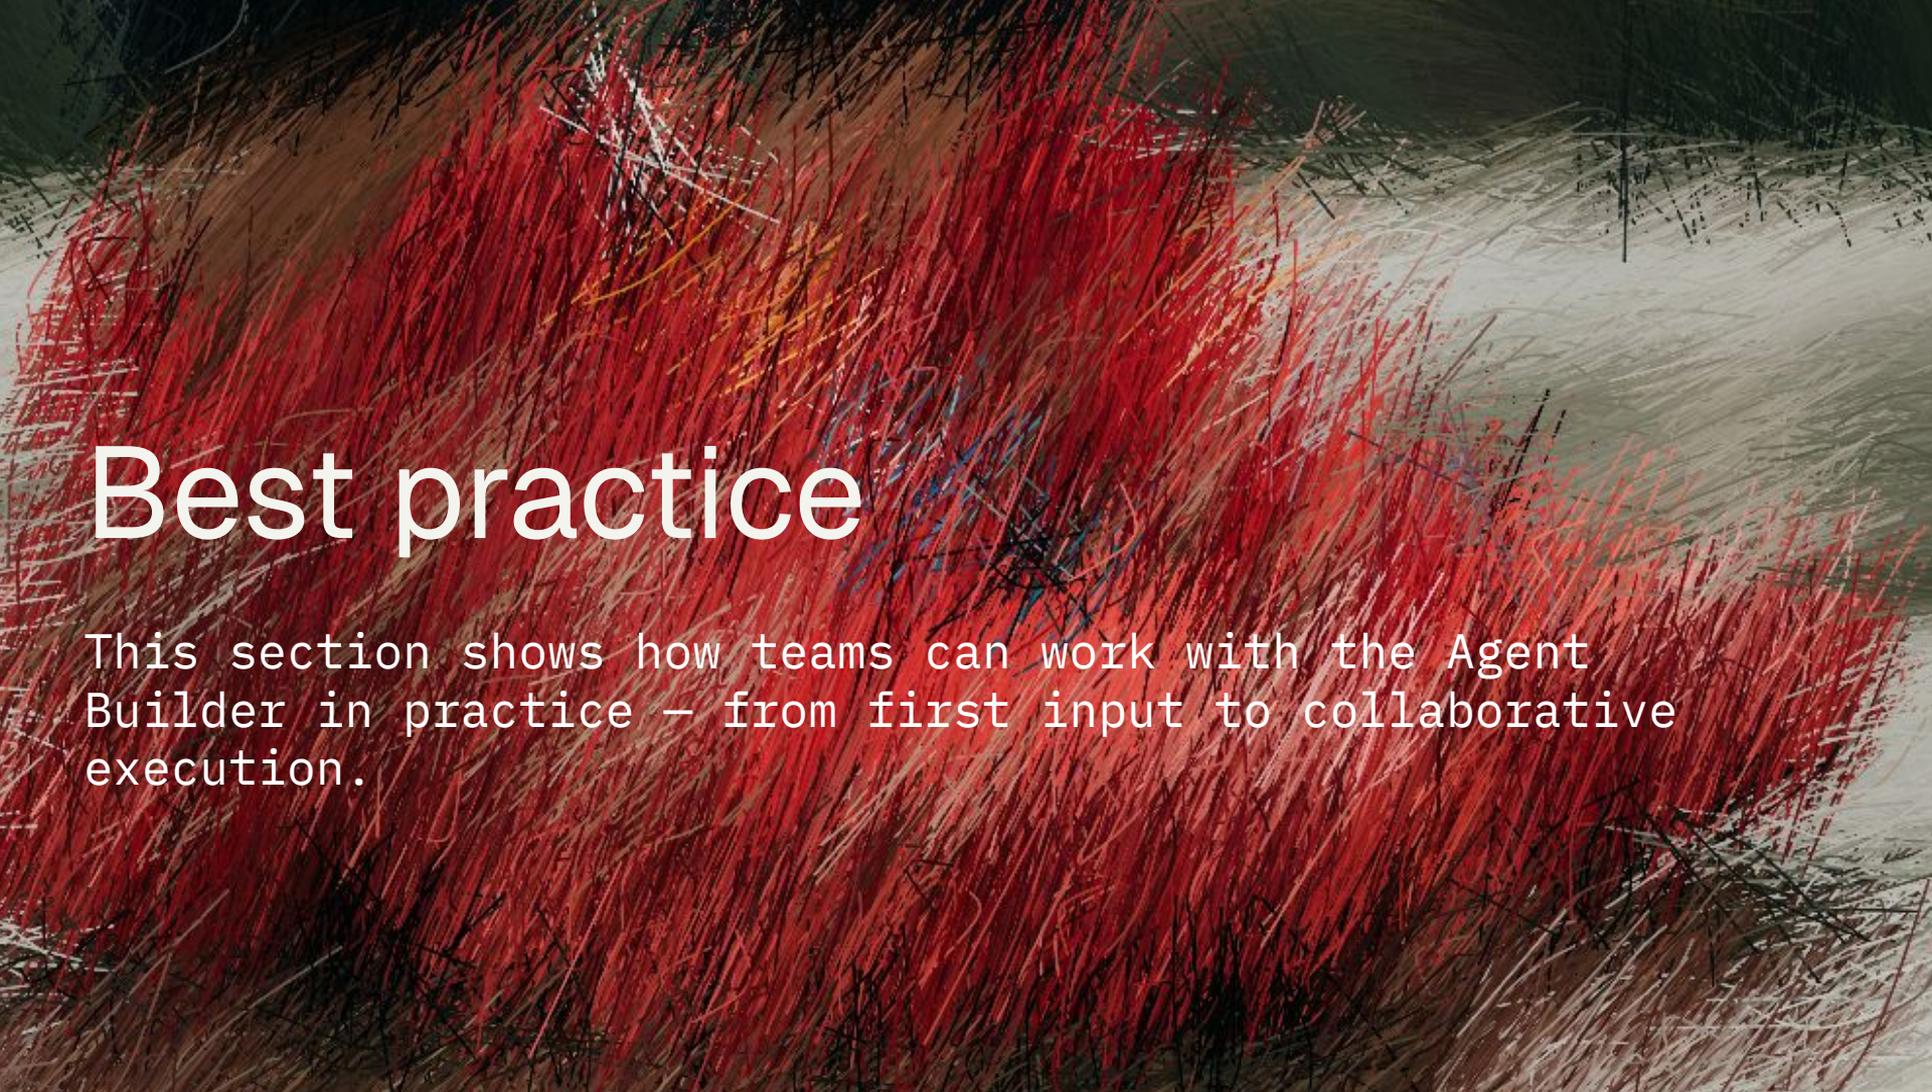
Pencil supports a growing range of tools across:

- Marketing & advertising
- Analytics & data
- CRM & business management
- Communication & collaboration
- Documentation, design & commerce

*Examples include Google Workspace, Salesforce, HubSpot, Jira, Slack, Meta Ads, Shopify, Tableau, and more.*

## Using Integrations in Agents

- Integrations can be added directly inside the Agent Builder
- Best practices:
  - Add integrations **only when live data or actions are required**
  - Keep agents focused by limiting connected tools
  - Test thoroughly before sharing your agents

The background of the slide is an abstract, textured composition of dense, overlapping brushstrokes. The primary colors are vibrant red and deep black, with some lighter, greyish tones visible in the upper right quadrant. The strokes are varied in length and direction, creating a sense of movement and depth. The overall effect is that of a layered, expressive painting.

# Best practice

This section shows how teams can work with the Agent Builder in practice – from first input to collaborative execution.

# Start with a brief

**Principle:** Always anchor agents to a clear input. A brief provides shared context, reduces ambiguity, and allows agents to act immediately without manual setup.

## How it works:

- Upload an existing brief (PDF, document, or deck) directly into Pencil.
- The platform automatically analyses the content and surfaces **suggested agents** that are best suited to the task.
- Agents begin generating relevant outputs (documents, images, videos) without requiring prompts or configuration.

## Best practices:

- Use real working briefs whenever possible — agents perform best with concrete inputs.
- Keep briefs focused on objectives, audience, and deliverables rather than process.
- Treat the brief as a living reference that agents can continuously refer back to.

**Key takeaway:** You don't need to prompt or configure anything upfront. Start with the work you already have.

Drag and drop your brief here

Upload your brief and our agents will suggest next actions to help you get started. We support PDF, DOCX, XLSX or PPTX.

Upload brief

You can upload up to 4 files with a total size of 100MB.

# Choose your team of agents

**Principle:** Complex work is best handled by multiple specialised agents rather than a single generalist.

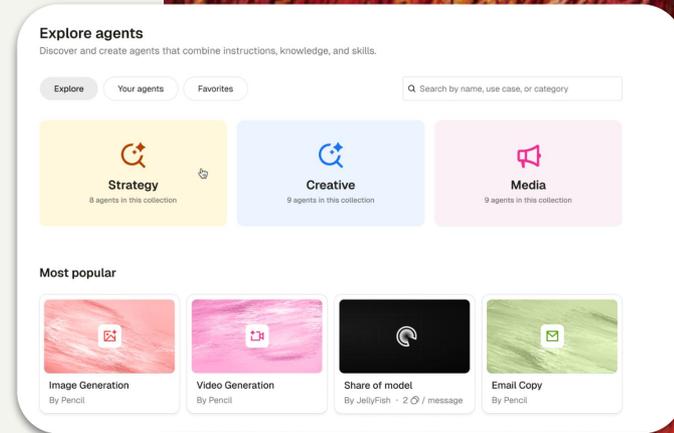
## How it works:

- Browse agents by category (e.g. Strategy, Creative, Media).
- Select agents based on the role they play in the workflow, not just the output they produce.
- Activate agents individually as needed — context carries across automatically.

## Best practices:

- Think in roles: strategist, creator, optimiser, distributor.
- Start with the most upstream agent (e.g. Strategy) before moving into execution.
- Reuse proven agents across projects to maintain consistency.

**Key takeaway:** Agents work best as a connected team, each focused on a single responsibility.



# Give agents the capabilities you need

**Principle:** Tools define what agents can *do*, not just what they can *say*.

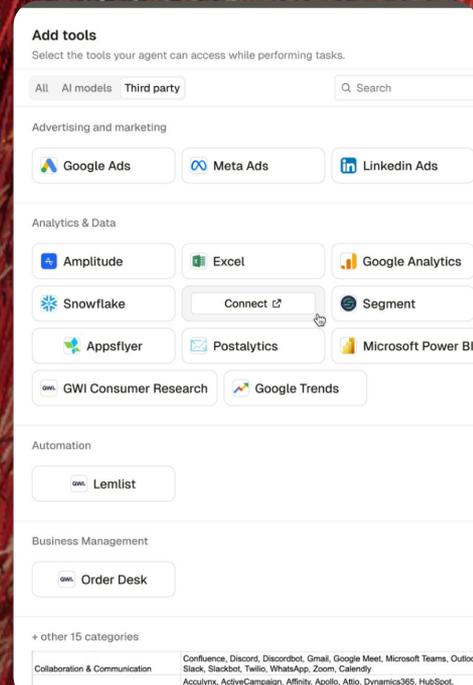
## How it works:

- Add native AI tools (image, video, web) for instant creative capabilities.
- Connect 3rd-party tools to extend agents into your existing data and media stack.
- Combine up to five tools per agent to support end-to-end workflows.

## Best practices:

- Add only the tools required for the job — more tools increase complexity.
- Use native tools for speed; use 3rd-party tools for automation and activation.
- Ensure all required tools are authenticated before sharing the agent.

**Key takeaway:** Powerful agents are defined as much by the capabilities you give them, as by how you prompt them, both during initial set up and conversational usage.



# Bring humans and agents together

**Principle:** The strongest workflows keep humans in control while letting agents do the heavy lifting.

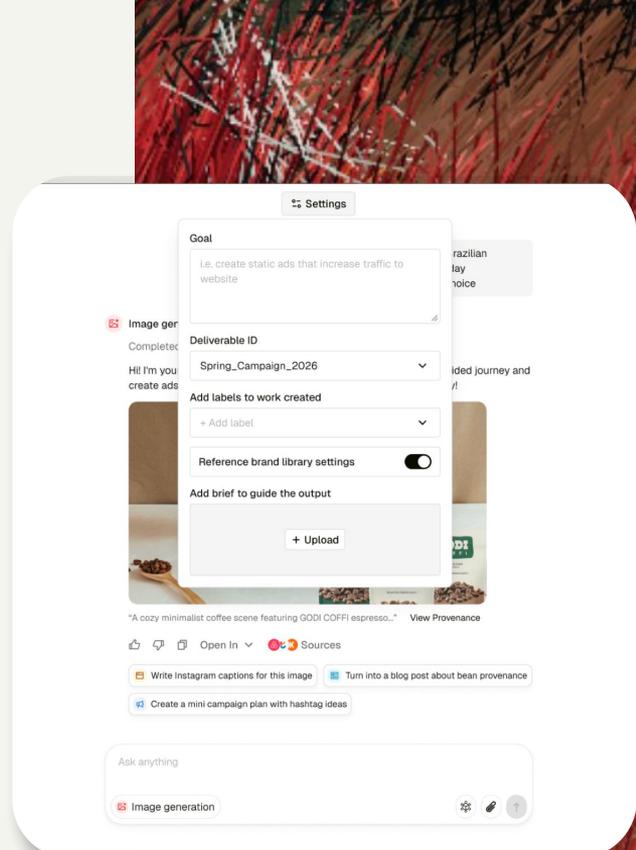
## How it works:

- Set a clear goal that agents can align around.
- Agents propose ideas, assets, or next steps proactively.
- Tasks are handed off between agents as the workflow progresses.

## Best practices:

- Use goals to guide agents, not micromanage them.
- Review and approve outputs at natural checkpoints.
- Encourage agents to collaborate — one agent's output should become another's input.

**Key takeaway:** Designing strong human-in-the-loop agentic workflows means giving your agents clear instructions on when to act — and when to pass the baton back to humans.



An abstract background featuring a dense, textured pattern of brushstrokes. The colors are primarily earthy, including shades of brown, tan, and beige, with prominent streaks of vibrant green and yellow. The strokes are varied in length and direction, creating a sense of movement and depth. The overall effect is reminiscent of a close-up of a painting or a natural, fibrous material.

# Practical examples

Read on to see some examples of how we can use agent builder in the real world

# Building a specialised custom agent from scratch

This section shows how teams can work with the Agent Builder in practice — from first input to specialised execution.

**Scenario:** You want to create a focused, reliable agent that performs a specific job well — rather than relying on a generic, pre-built agent.

This example walks you through the full process, from accessing the Agent Builder to publishing the agent for others to use.

**Why this step matters:** Starting in the Agent Builder gives you full control over how the agent behaves, which tools it can use, and how it's governed.



New agent

# Building a specialised custom agent

## Step 1: Access the Agent Builder

### Action:

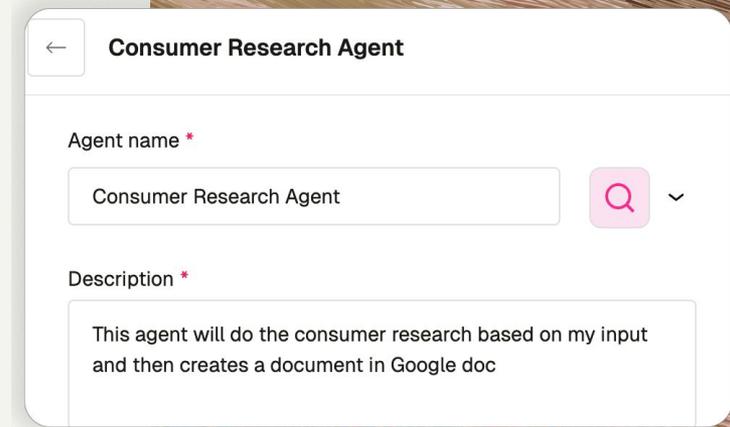
- Open Pencil and navigate to **Agents** from the left-hand menu.
- Click **New Agent** to begin creating a custom agent.

## Step 2: Define the Agent Basics

### Action: Complete the core setup fields:

- **Agent name:** Choose a clear, task-focused name.
- **Description:** Briefly explain what the agent does and when it should be used.
- **Instructions:** Write clear guidance describing the agent's role, expected behaviour, and any rules it must follow.

- ✓ Describe the desired outcome, not every step.
- ✓ Be explicit about tone, format, and constraints.
- ✓ Keep instructions concise to avoid confusion



The screenshot shows a mobile application interface for creating a custom agent. At the top, there is a back arrow and the title 'Consumer Research Agent'. Below the title, there are two main input fields:

- Agent name \***: A text input field containing 'Consumer Research Agent' and a search icon.
- Description \***: A text input field containing the text: 'This agent will do the consumer research based on my input and then creates a document in Google doc'.

# Building a specialised custom agent

## Step 3: Choose a Conversation Model

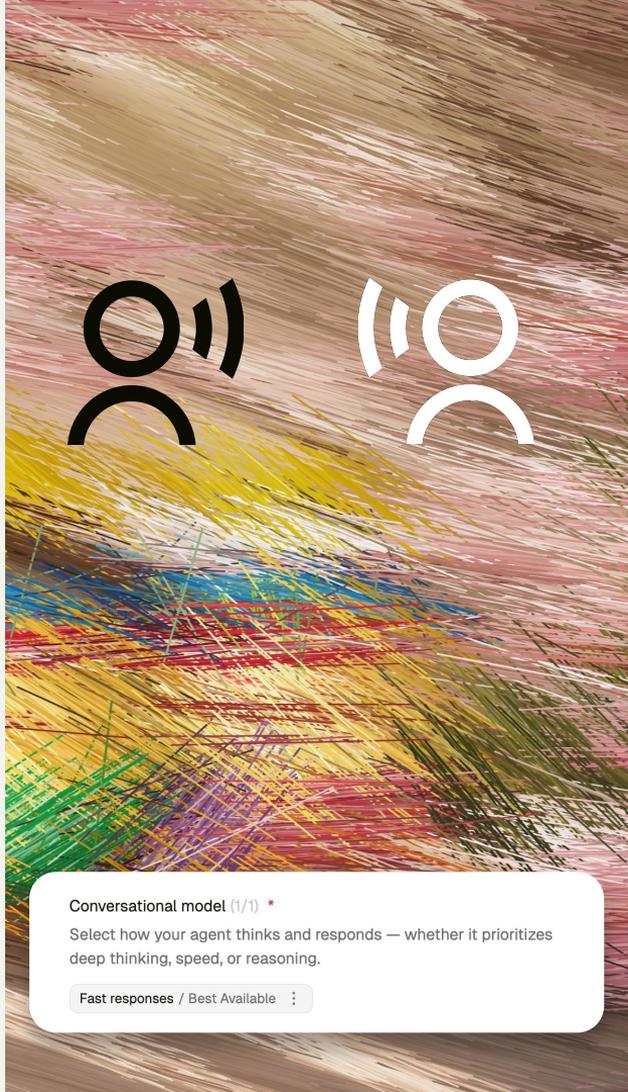
**Action:** Select the most appropriate **conversation model category** based on the task:

- **Fast responses:** Best for copywriting, summaries, and quick iterations.
- **Reasoning:** Suitable for structured problem-solving and logical tasks.
- **Deep thinking:** Ideal for strategy, analysis, and complex reasoning.
- **Multimodal input and output:** Required when working with text, images, and video together. Excels at visual analysis.

### Best Available option:

- Enable **Best Available** to ensure the agent automatically upgrades to the strongest model chosen by Pencil over time.

**Why this step matters:** Choosing the right model ensures the agent's responses match the complexity, speed, and modality of the task



Conversational model (1/1) \*

Select how your agent thinks and responds — whether it prioritizes deep thinking, speed, or reasoning.

Fast responses / Best Available ⋮

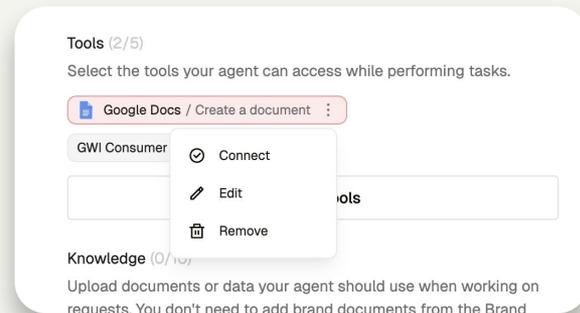
# Building a specialised custom agent

## Step 4: Add Tools to Give the Agent Capabilities

**Action:** Add tools that allow the agent to act, not just respond. Each agent can include **up to five tools**.

### Tool categories:

- **All:** Browse the full tool library.
- **AI Models:** Native capabilities already available within Pencil, such as:
  - Text-to-image generation
  - Image generation using image reference
  - Product swap and product placement
  - Text-to-video and image-to-video
  - Utilities like Visit URL or Web Search
- **Third-party tools:** External applications across areas such as advertising, analytics, automation, CRM, communication, design, and documentation.



- Add only the tools required for the job.
- Fewer tools help keep agents specialised and reduce hallucinations.
- For image and video tools, select **Best Available** unless governance requires a specific model.



# Building a specialised custom agent

## Step 5: Connect Third-Party Tools (Admin Step)

### Action:

- Authenticate any third-party tools added to the agent.
- Connections can be completed directly in the Agent Builder or via **Settings** → **Integrations**.

### Important rules:

- Integrations are an admin-level function.
- Some tools connect via account logins (e.g. Google Docs).
- Others require API keys (e.g. Semrush).

**Why this step matters:** Agents can't function unless their tool stack is fully connected and authenticated. Without authentication, third-party tools can't be used and the agent chat will be locked.



# Building a specialised custom agent



## Step 6: Test the Agent

### Action:

- Run test prompts or tasks.
- Confirm that instructions are followed correctly.
- Verify that all tools behave as expected.

### Best practices:

- ✓ Test edge cases.
- ✓ Adjust instructions or tool selection if outputs are inconsistent.

## Step 7: Publish the Agent

**Action:** Choose where to publish the agent:

- **Current workspace only**, or
- **All workspaces under the subscription** (Subscription Admins only)

### Connection carry-over:

- When publishing broadly, Admins can select which third-party connections to carry over.
- If a target workspace already has an existing connection, it will be used instead, and will not be overridden by the new agent's one.
- If a required tool is not connected, users will be prompted to contact their Admin.

# Examples of Specialised Agents

## Consumer Research Agent

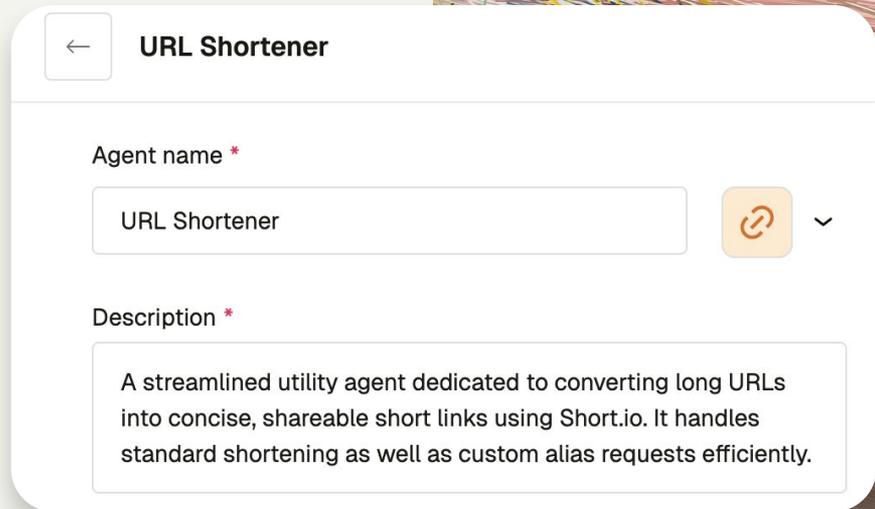
- Combines a Pencil-enabled consumer research tool with a connected Google Docs integration.
- Automatically finds insights and saves them into a document.

## Image Retouch Agent

- Uses native image tools such as Text-to-Image and Image Reference.
- Runs on a Multimodal conversation model to modify product imagery (e.g. changing colours or placing products in scenes).

## URL Shortener Agent

- Uses a third-party URL shortening tool.
- Converts long URLs into short, shareable links instantly.

A screenshot of a user interface for a 'URL Shortener' agent. The interface is white with rounded corners and a subtle shadow. At the top left is a back arrow icon. The title 'URL Shortener' is centered at the top. Below the title, there is a form with two main sections. The first section is labeled 'Agent name \*' and contains a text input field with the value 'URL Shortener' and a dropdown menu icon (a downward arrow) to its right. The second section is labeled 'Description \*' and contains a text area with the text: 'A streamlined utility agent dedicated to converting long URLs into concise, shareable short links using Short.io. It handles standard shortening as well as custom alias requests efficiently.' The background of the slide features a vertical abstract painting with streaks of yellow, pink, and brown.

# Examples of Specialised Agents

## Modelling Shoot Creator

This agent is designed to lead the user through a structured, human-in-the-loop flow — collecting the right inputs at the right time and producing polished studio imagery.

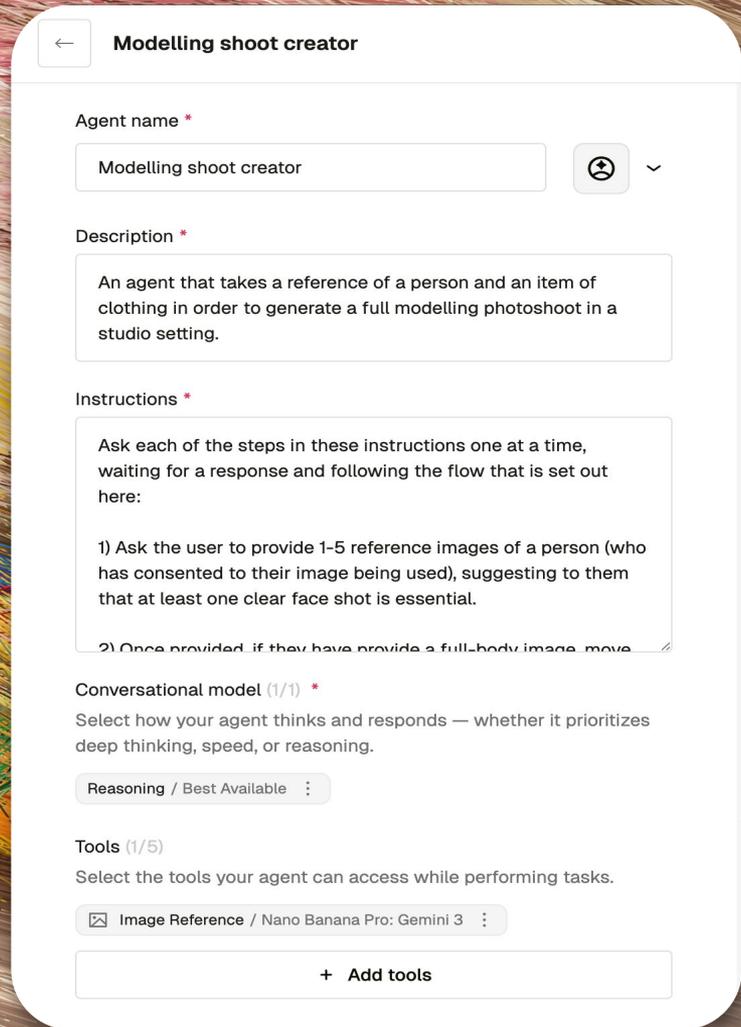
### Step 1: Start the Agent and Define the Vision

#### What happens:

- The user opens the **Modelling Shoot Creator** agent.
- The agent confirms the creative direction (e.g. luxury fashion, couture-inspired aesthetic).

#### Best practice:

Keep the creative direction high-level at first. Refine details later.



The screenshot shows the configuration interface for the 'Modelling shoot creator' agent. It includes a back arrow, the agent name 'Modelling shoot creator', a description 'An agent that takes a reference of a person and an item of clothing in order to generate a full modelling photoshoot in a studio setting.', and instructions: 'Ask each of the steps in these instructions one at a time, waiting for a response and following the flow that is set out here: 1) Ask the user to provide 1-5 reference images of a person (who has consented to their image being used), suggesting to them that at least one clear face shot is essential. 2) Once provided, if they have provide a full-body image, move...'. It also shows the 'Conversational model' set to 'Reasoning / Best Available' and 'Tools' including 'Image Reference / Nano Banana Pro: Gemini 3'. A '+ Add tools' button is at the bottom.

← **Modelling shoot creator**

Agent name \*

Modelling shoot creator

Description \*

An agent that takes a reference of a person and an item of clothing in order to generate a full modelling photoshoot in a studio setting.

Instructions \*

Ask each of the steps in these instructions one at a time, waiting for a response and following the flow that is set out here:

- 1) Ask the user to provide 1-5 reference images of a person (who has consented to their image being used), suggesting to them that at least one clear face shot is essential.
- 2) Once provided, if they have provide a full-body image, move...

Conversational model (1/1) \*

Select how your agent thinks and responds — whether it prioritizes deep thinking, speed, or reasoning.

Reasoning / Best Available

Tools (1/5)

Select the tools your agent can access while performing tasks.

Image Reference / Nano Banana Pro: Gemini 3

+ Add tools

# Modelling Shoot Creator

## Step 2: Collect Reference Images of the Model

### What the agent asks for:

- 1–5 reference images of a person (with consent for image use).
- At least one clear face shot to ensure accurate likeness.

**Why this matters:** High-quality reference images improve realism and consistency across generated outputs.

## Step 3: Refine Physical Details

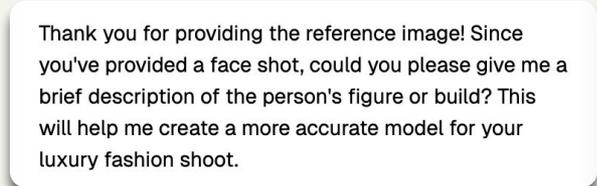
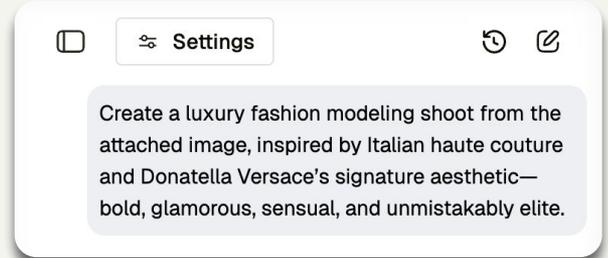
### What happens:

- Once a face reference is provided, the agent asks for a brief description of the person's figure or build.

## Step 4: Upload the Outfit or Clothing Item

### What the agent asks for:

- An image of the outfit the model should wear.
- Packshot or on-hanger images are recommended over worn clothing.



# Modelling Shoot Creator

## Step 5: Define the Photoshoot Setting

### What happens:

- The agent asks whether a specific setting is required.
- If none is provided, it defaults to a professional white studio backdrop.

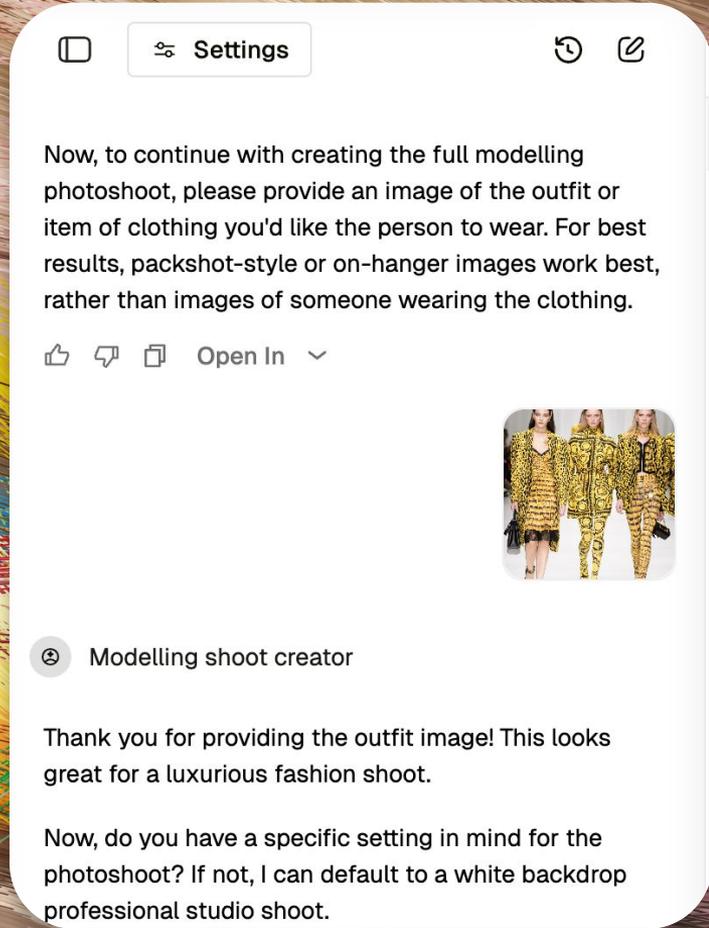
### Best practice:

- Start with neutral studio settings for versatility.

## Step 6: Generate the Modelling Shoot

### What happens:

- The agent uses image reference capabilities to generate a cohesive modelling shoot.
- Outputs appear directly in the canvas for review and iteration.



# Modelling Shoot Creator



Settings



Create a luxury fashion modeling shoot from the attached image, inspired by Italian haute couture and Donatella Versace's signature aesthetic—bold, glamorous, sensual, and unmistakably elite.



Modelling shoot creator

Thank you for providing the reference image! Since you've provided a face shot, could you please give me a brief description of the person's figure or build? This will help me create a more accurate model for your luxury fashion shoot.



# Examples

Some examples created with the Modelling Photoshoot Agent





**Pencil**

# Coming Soon

Sneak Peaks of upcoming soon-to-be-released features that will make your agents even more intelligent, performant and creative. As new features are released alongside **Agent Builder**, powerful new potentials and use cases will emerge, as various features synergistically feed into each other across the platform.

## Tool Ecosystem Additions

Over the coming weeks, we're going to add a cohort of new Native AI tools for your agents to use, in addition to the current tool ecosystem offerings. These will include powerful state-of-the-art capabilities such as:

- Generative Fill
- Image Erase
- Image Expand
- Face Swap
- Remove Background
- Image and Video Enhancers

## New agent selection interface

A new and improved user experience for selecting agents is coming soon. As the list of available agents grows, we'll be providing a more streamlined way to sort, search and organise your agents.