Groundzero

FEATURE DOCUMENTATION

Analyze and AI

Build charts and dashboards, ask the AI assistant, run AI agents, track consumption and reach Groundzero from your own tools.

Screenshots use sample data for a company called Acme Retail.
For analysts and everyone who works with the AI assistant. The Administration guide covers people, access, sites, budgets and billing.

Contents

About this guide 3

What this guide helps you do 3

Where you work 3

How to read this guide 3

The naming rule 4

1 Visualizations 5

1.1 Choose the right chart for your question 5

1.2 Open the Visualizations list 5

1.3 Build a chart from a dataset 7

1.4 Style the chart 14

1.5 Let readers drill into detail 16

1.6 Rename or delete a visualization 17

1.7 What to know 17

1.8 Troubleshooting 18

1.9 Related chapters 18

2 Saved filters 19

2.1 Choose the right filter type 19

2.2 Open the Filters list 19

2.3 Create a filter 20

2.4 Configure a filter 20

2.5 Make a filter change a chart 22

2.6 Use filters on a chart 23

2.7 Related chapters 23

3 Dashboards 24

3.1 Open the Dashboards list 24

3.2 Build a dashboard from your charts 24

3.3 Let readers click a chart to filter or jump 28

3.4 Let readers switch between stacked charts 29

3.5 Filters on a dashboard 30

3.6 Publish a dashboard 30

3.7 Export a dashboard to PDF or Excel 32

3.8 Read a published dashboard in a workspace 32

3.9 Copy a published dashboard to edit it 33

3.10 Rename or delete a dashboard 35

3.11 Troubleshooting 35

3.12 Related chapters 36

4 AI chat assistant 37

4.1 Ask a question 37

4.2 Read an answer 41

4.3 Stop, queue and steer a response 43
4.4 Approve or deny an action 44
4.5 Find an earlier conversation 45
4.6 Search everything with Ctrl+K 46
4.7 If something goes wrong 47
4.8 Limits 48
4.9 Related chapters 48

5 AI agents 49
5.1 Decide between chat and an agent 49
5.2 Find your agents 49
5.3 Build an agent 50
5.4 Test an agent 56
5.5 Get a version approved and live 57
5.6 Look at an agent's details 58
5.7 Let another system call your agent 59
5.8 Review runs and audit history 61
5.9 See versions and approvals 62
5.10 Limits 63
5.11 Troubleshooting 63
5.12 Related chapters 63

6 Consumption metrics 64
6.1 Open the metrics and pick a period 64
6.2 See who uses AI chat and what it costs 65
6.3 Track how your agents are doing 66
6.4 Find expensive jobs and datasets 67
6.5 Use the numbers in your own tools 68
6.6 Troubleshooting 69
6.7 Related chapters 69

7 API access, your account and help 70
7.1 Reach Groundzero from your own tools 70
7.2 Connect a coding assistant with MCP 70
7.3 Open the help menu 71
7.4 Change your password 72
7.5 Set up two-step verification 73
7.6 Related chapters 74

About this guide

What this guide helps you do

This guide is for people who turn data into answers: analysts who build charts and dashboards, and everyone who works with the Groundzero AI assistant and AI agents. Each chapter starts with the job you want to get done, walks you through it step by step, and ends with how to check it worked and what to do if it did not.

I want to...	Go to
Pick the right chart for my question and build it	Chapter 1, Visualizations
Let readers narrow a chart with a drop-down, slider or switch	Chapter 2, Saved filters
Put several charts on one page, publish it and export it	Chapter 3, Dashboards
Ask questions about my data in plain language	Chapter 4, AI chat assistant
Build an assistant with a fixed job and let other systems call it	Chapter 5, AI agents
See how much AI and compute we use and what it costs	Chapter 6, Consumption metrics
Call Groundzero from my own tools, change my password or connect a coding assistant	Chapter 7, API access, your account and help

Two companion guides sit next to this one:

- **Groundzero feature guide: Administration** covers people, groups, access, sites, budgets and billing. Your administrator uses it. Chapter 4 there explains the permissions you need for the features in this guide.
- The **Data** and **Build and Run** guides cover connections, the catalog, datasets, projects, notebooks and jobs.

Where you work

Everything in this guide happens in the screens you open from your site's address. The home page is the AI chat assistant. The top navigation also has Workspaces, Connections, Designer (visualizations, filters and dashboards), DE/ML, Lakehouse and Catalog.

How to read this guide

- Button, tab and field names are in **bold**, exactly as they appear on screen.
- Values you type and fixed text use `monospace`.
- Screenshots use sample data for a company called Acme Retail. Names, numbers and dates are only examples.
- Tips show a faster or safer way to work. Notes explain something that is easy to miss. Warnings mark actions that cannot be undone or that cost money.

TIP

If a button is missing or a toast says "You are not authorized to perform this action", you do not have the permission yet. Ask your administrator. Chapter 4 of the Administration guide explains what to ask for.

The naming rule

Visualizations, filters, dashboards and workspaces all use the same rule for names. A name must start with a letter, may contain only letters, numbers and hyphens, and cannot end with a hyphen or contain two hyphens in a row. Spaces are not allowed. When a name breaks the rule, red text appears under the field:

> Name must start with a letter, contain only letters, numbers, and hyphens, and cannot end with or contain consecutive hyphens.

`Acme-Sales` is valid. `Acme Sales`, `-Sales` and `Sales--Q4` are not.

1 Visualizations

A **visualization** is a saved chart or table built on a dataset. You choose a chart type, pick a dataset, map its columns to the chart (axes, series, labels) and style it. Saved visualizations can be placed on dashboards (chapter 3), opened by AI agents and referenced by the AI assistant.

Before you start you need at least one dataset (see the Build and Run guide) and the Visualization and Dataset permissions. If a toast says "You are not authorized to perform this action", ask your administrator for them (Administration guide, chapter 4).

1.1 Choose the right chart for your question

Start from the question you want the chart to answer. Groundzero has 22 chart types.

How do categories compare? (sales by region, orders by channel) - **Bar** draws one bar per category along the vertical axis. It suits long category names. It can also show a target line. - **Vertical Bar** draws columns along the horizontal axis. It suits a few short categories. - **Stacked Bar** and **Stacked Vertical Bar** split each bar into series, so you see the total and its parts together. - **Combo** puts bars and a line in one chart, for example orders as bars and revenue as a line.

How does something change over time? (monthly revenue) - **Line** is the usual choice. Turn on **Add Dual Axis** to plot two measures with different scales, or **Stacked Area Chart** to stack the lines as areas. - **Candle** shows low, open, close and high values for each period. - **Combo** can show a count as bars with a trend as a line.

What share of the whole is each part? (revenue by category) - **Pie**, **Donut** and **Half Donut** show shares of a total. Donut and Half Donut can print the total in the centre. - **Tree Map** shows shares as nested rectangles and copes with many categories. - **Sunburst** shows shares across levels of a hierarchy, such as region, then country, then city.

How am I doing against a target? (sales against quota) - **Gauge** shows one value on a dial with a maximum and coloured ranges. - **Scorecard** (listed as **Card Chart**) shows one or more big numbers with labels and an optional sparkline.

How are two or three measures related? (price against units sold) - **Scatter** plots one measure against another. - **Bubble** adds a third measure as the size of each dot.

Where do people drop out, or how does volume flow? - **Funnel** shows stages in order, such as visit, cart, order. - **Sankey** shows how volume flows from a source to a target. - **Graph** shows links between items as a network of nodes and edges.

Where is it? - **Geo** colours regions of a map by a value.

I need the exact numbers. - **Table** shows rows with optional groups, totals and a small trend chart in each cell. - **Pivot** lets readers see one measure across rows and columns, with totals.

> TIP
> Groundzero does not total or average data inside the chart (apart from Pivot and table totals). Aggregate in the dataset query, then chart the result.

1.2 Open the Visualizations list

1. Open **Designer** in the top navigation.

2. In the left rail choose **Visualizations**. The rail also has Datasets, Filters and Dashboards.

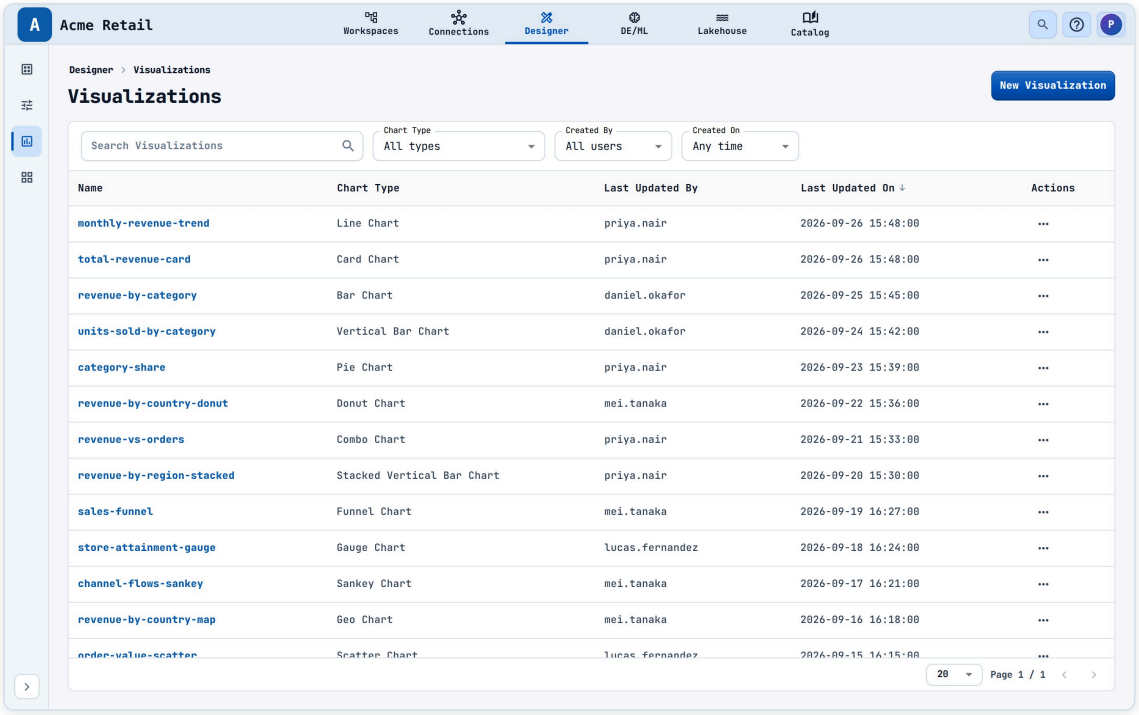


Figure 1. The Visualizations list with search, Chart Type, Created By and Created On filters.

Control	What it does
Search Visualizations	Filters the list by name as you type.
Chart Type	"All types" plus every chart type present in the list. Choosing one adds a chip such as `Type: Bar Chart`.
Created By	"All users" or one creator.
Created On	Any time, Today, Last 7 days or Last 30 days.
New Visualization	Opens the Add Visualization dialog.
Actions (three dots)	Rename or Delete a row.

Columns are **Name** (a link that opens the editor), **Chart Type**, **Last Updated By** and **Last Updated On** (in your local time), sorted by last update. When nothing exists you see "No visualizations to show". When your filters hide every row you see "No matches found".

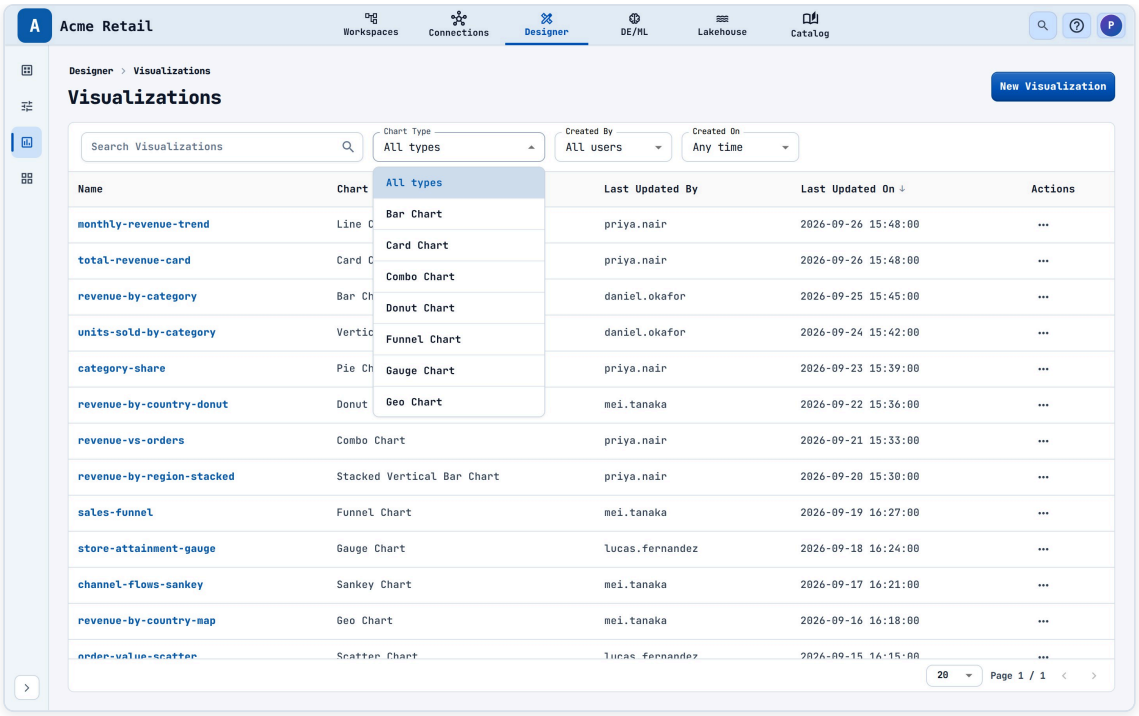


Figure 2. Filtering by chart type.

1.3 Build a chart from a dataset

Outcome: a saved chart that always shows live data.

1. In **Designer > Visualizations**, choose **New Visualization**.
2. In **Enter Name** type a name that follows the naming rule. While the name is invalid, red text appears and **Add** stays disabled.
3. Choose **Add**. A toast confirms "<name> added successfully" and the new row appears. Groundzero does not open it for you.
4. Choose the new name to open the editor.
5. In the right pane, on the **Setup** tab, choose a **Chart Type**, then a dataset under **Choose Dataset**, then fill the fields that appear (see the field reference below).
6. Choose the floppy-disk icon above the chart to save. The toast says **Visualization updated**.

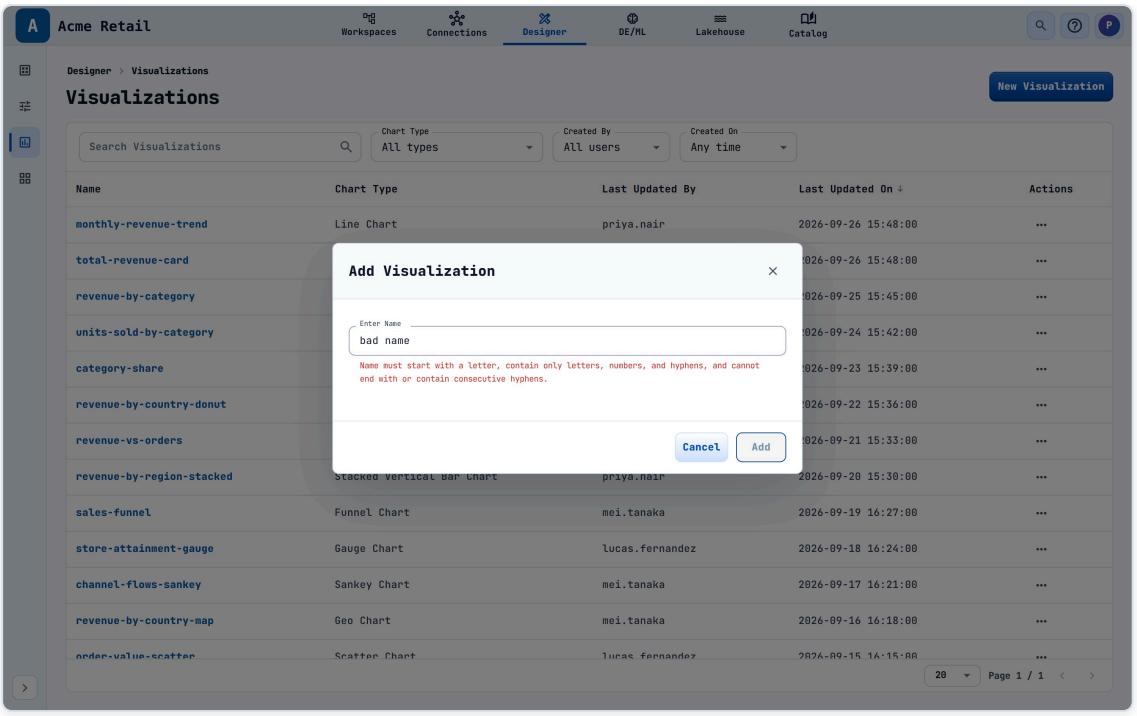


Figure 3. The Add Visualization dialog rejecting a name that contains a space.

A new visualization has no chart type and no dataset until you choose them in the editor.

How to check it worked: the chart draws on the canvas, and reopening it from the list shows the same chart with current data.

> **NOTE**
> Save before you leave the editor. Groundzero does not store the chart's rows. It re-runs the dataset each time the chart opens, so a saved chart always shows live data.

1.3.1 The editor

The editor has three areas.

1. **Filters** (left) shows the filter controls attached to the dataset (chapter 2). It is empty when the dataset has none.
2. **Canvas** (centre) shows the breadcrumb `Designer > Visualizations > <name>`, the **Save** icon and the chart. An empty canvas says "Complete Chart Setup".
3. **Right pane** has the tabs **Setup** and **Customize**.

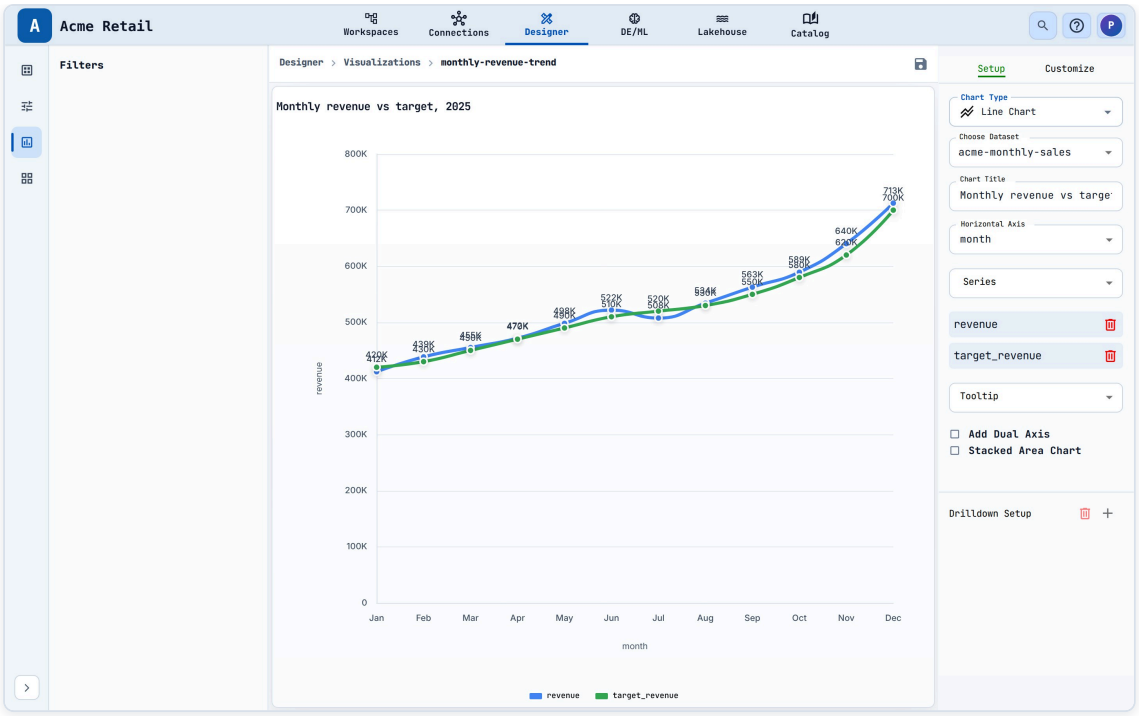


Figure 4. The editor for a Line chart on the monthly sales dataset.

1.3.2 The Setup tab, top to bottom

1. **Chart Type**. Choose the field to open an icon grid and hover an icon to see its name. Picking a type resets the chart configuration for that type but keeps the dataset.
2. **Choose Dataset**. A searchable list of all datasets ("Loading datasets..." while it loads). Choosing a dataset clears your series and column choices and fills the field lists from the dataset.
3. **Chart-specific fields** appear once the columns load (reference below).
4. **Drilldown Setup** (not for Pivot and Gauge). See "Let readers drill into detail".

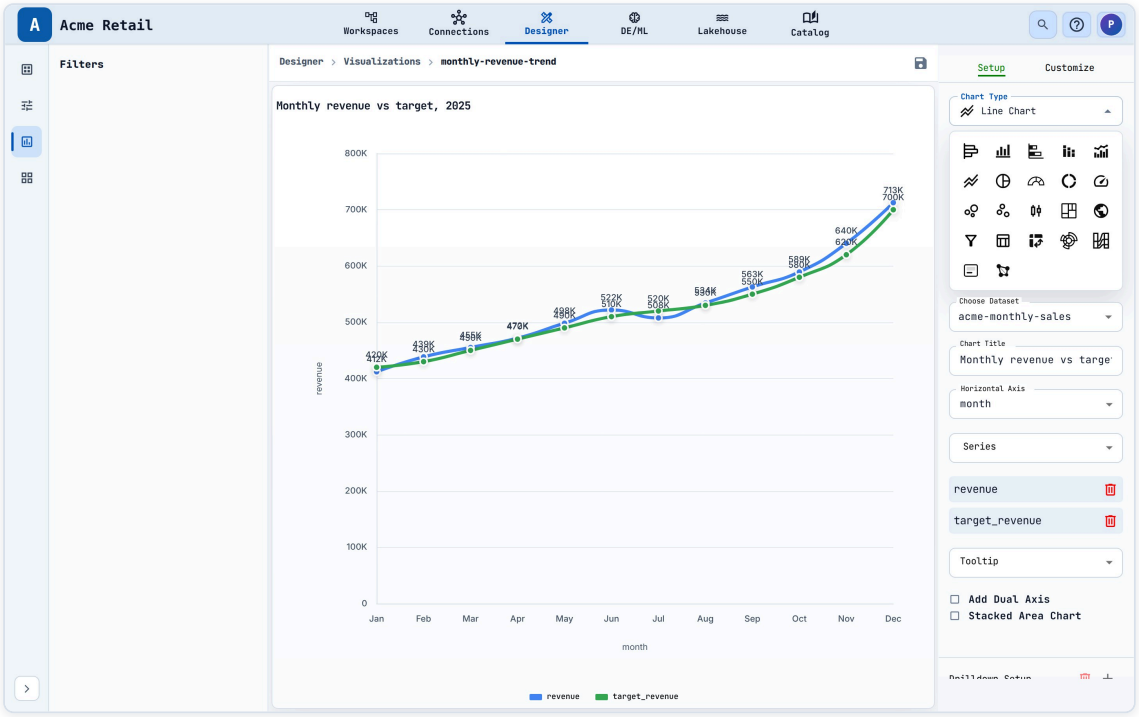


Figure 5. The chart type grid.

Series lists offer numeric columns only. Choosing the same series twice shows "Column added already."

1.3.3 Setup fields for each chart type

Chart	Setup fields
Bar	Chart Title, Vertical Axis, Series, target-line column, Tooltip
Vertical Bar, Stacked Vertical Bar, Combo	Chart Title, Horizontal Axis, Add series, Tooltip
Stacked Bar	Chart Title, Vertical Axis, Series, Tooltip
Line	Chart Title, Horizontal Axis, Series, Tooltip, switches **Add Dual Axis** and **Stacked Area Chart**
Pie	Chart Title, Label, Series, Tooltip
Donut, Half Donut	Pie fields plus **Show Total** (draws the total in the centre)
Funnel	Chart Title, Label, Series, Tooltip
Gauge	Chart Label, Series, **Max Value**, **Minor Ticks** (shows the first row)
Bubble	Chart Title, Label, Vertical Axis, Horizontal Axis, Size, Tooltip
Scatter	Chart Title, Vertical Axis, Horizontal Axis, Tooltip
Candle	Chart Title, Horizontal Axis, Low, Open, Close, High, Tooltip
Tree Map	Chart Title, ID, Size, Tooltip
Geo	Chart Title, Map Type, Region Type, Region Column, Attribute, Country, ISO Code (optional), Series, Region, map label columns (maximum 2), Tooltip, Color, Display Mode
Table	Chart Title, **Rows Group**, **Column Group**, **Columns**, **Show Totals**, **Select Chart** (sparkline type), **Column Header**
Pivot	Chart Title, **Rows**, **Columns**, **Values**, **Show Totals**, Select Chart, Column Header
Sunburst	Chart Title, **Hierarchy Fields**, Series
Sankey	Chart Title, Source, Target, Value (Source must differ from Target)
Graph	Chart Title, Source, Target, Value, Category (optional), Layout, Tooltip
Scorecard (Card Chart)	Chart Title, Sub Label, **Add Label Column**, **Add Value Column**, **Add Sparkline Column**, Sparkline Type

1.3.4 See examples of finished charts

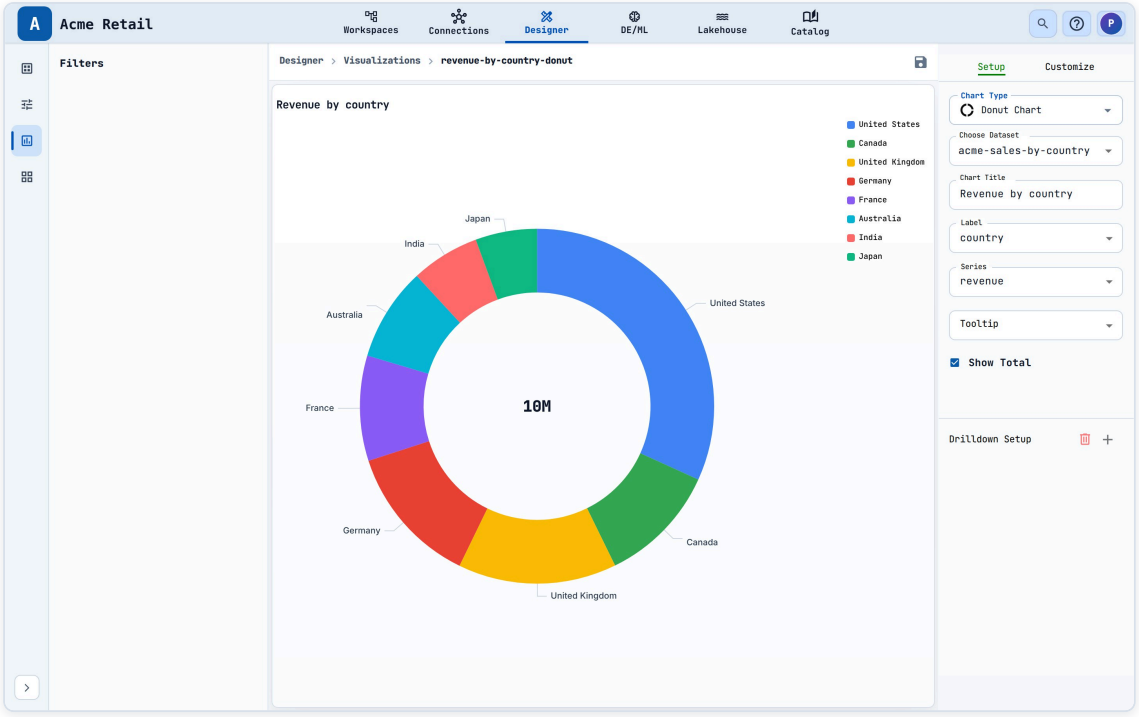


Figure 6. A Donut chart with the total drawn in the centre.

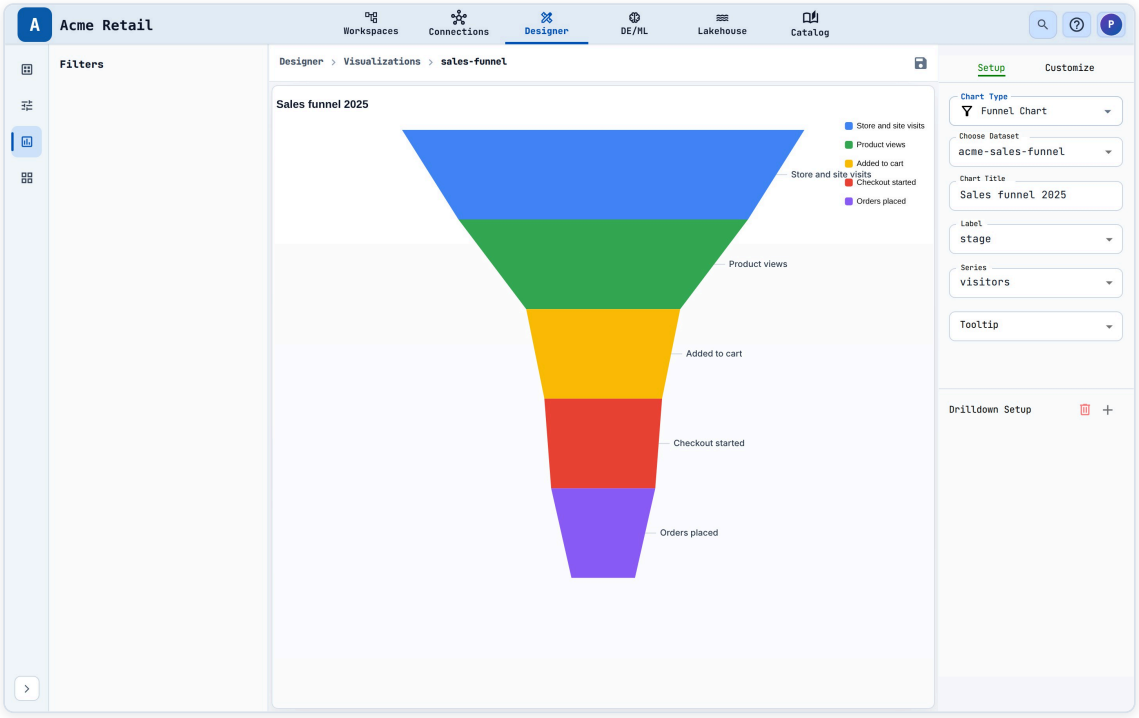


Figure 7. A Funnel chart.

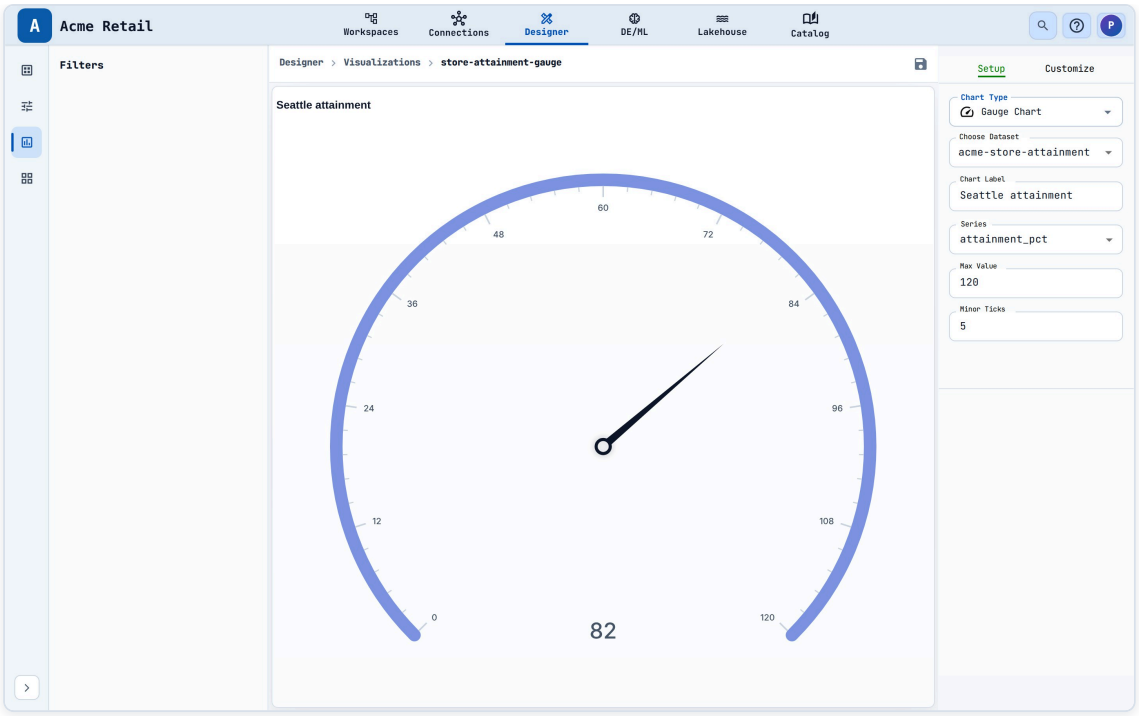


Figure 8. A Gauge chart with a maximum value and minor ticks.

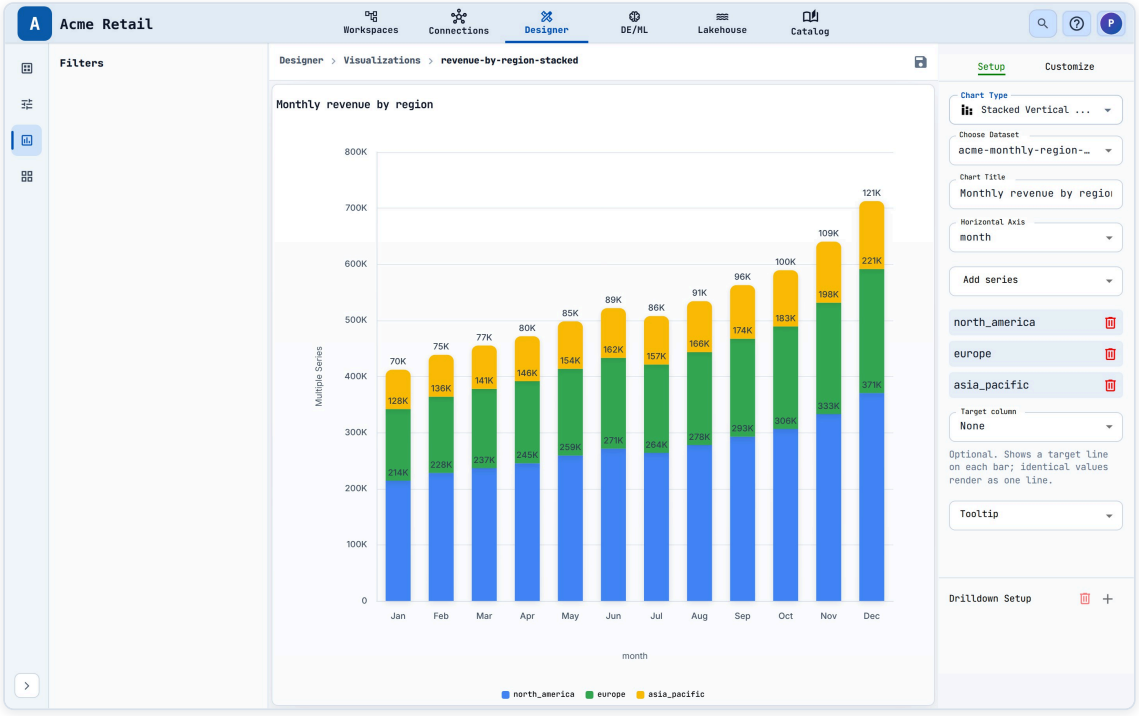


Figure 9. A Stacked Vertical Bar chart with one series per region.



Figure 10. A Combo chart: bars for orders and a line for revenue.

1.3.5 Tables, pivots and scorecards

Tables and pivots show data in a grid. A Table supports row and column groups, totals and a small trend chart (sparkline) in each column. A Pivot builds rows, columns and values, and loads data a page at a time. Both support conditional cell colours and symbols, which you set in the Setup and Customize tabs.

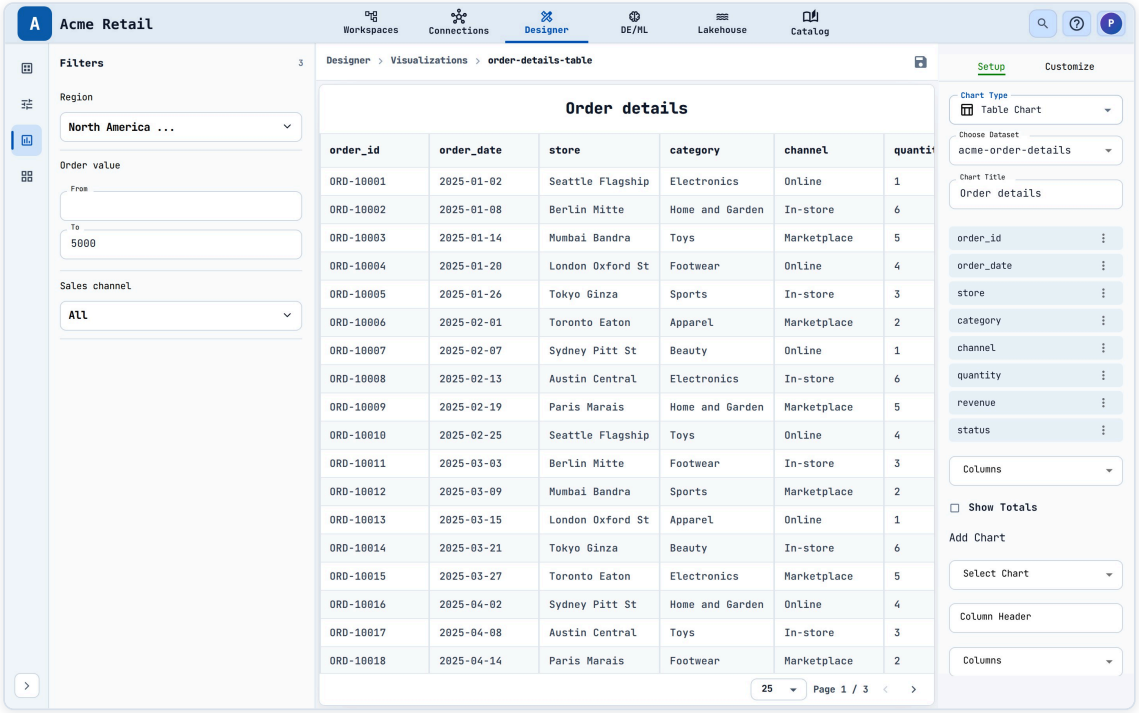


Figure 11. A Table chart with three attached filters in the left pane.

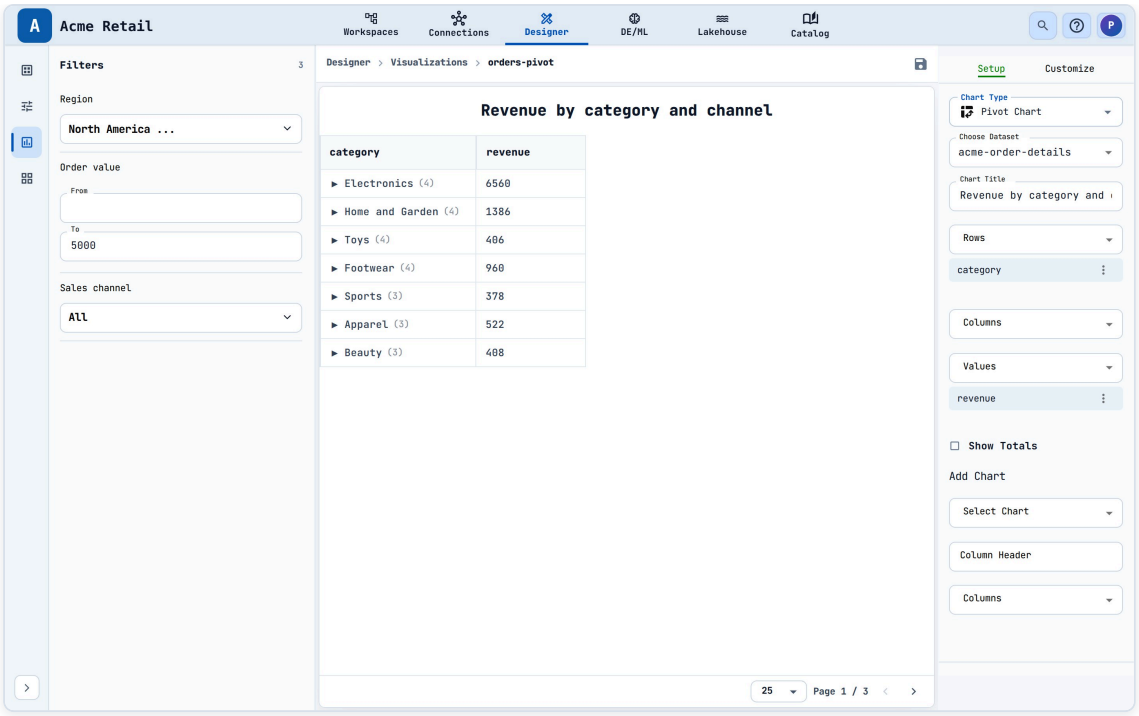


Figure 12. A Pivot chart grouped by category.

A Scorecard shows one or more labelled values with an optional sparkline.

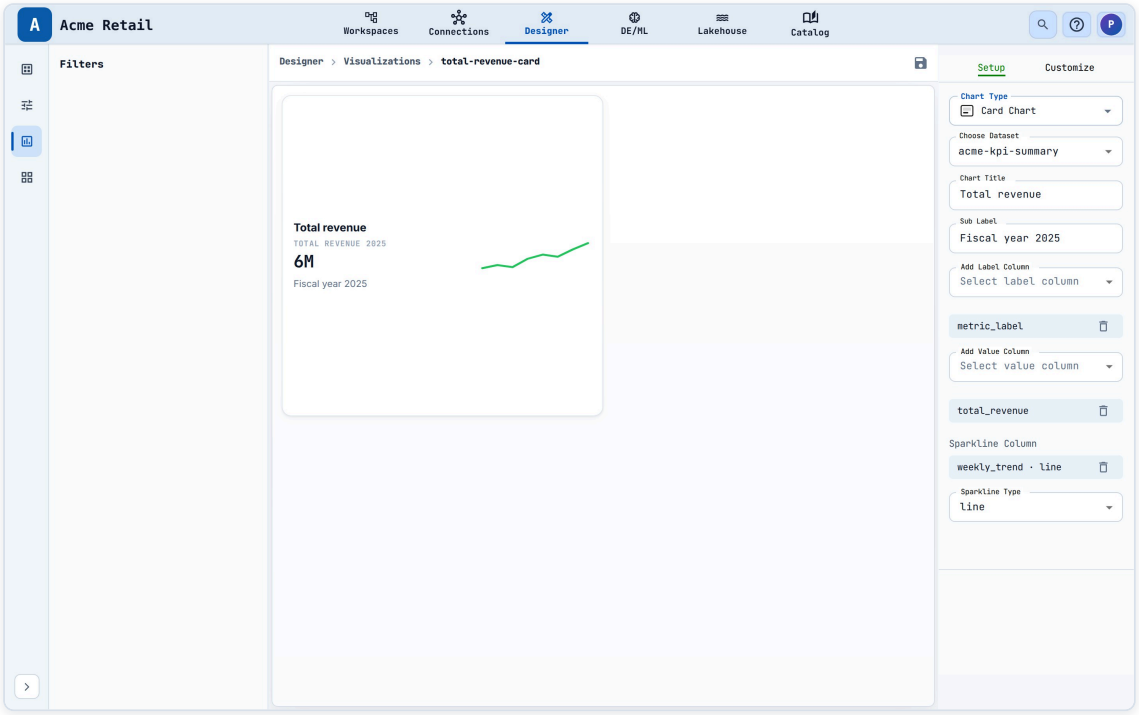


Figure 13. A Scorecard with a value and a sparkline.

1.4 Style the chart

Open the **Customize** tab to change titles, colours, legend, axes, labels and more. It is a set of collapsible sections, and which ones you see depends on the chart type. Before the chart is set up you see "Please complete the chart setup before customizing the chart".

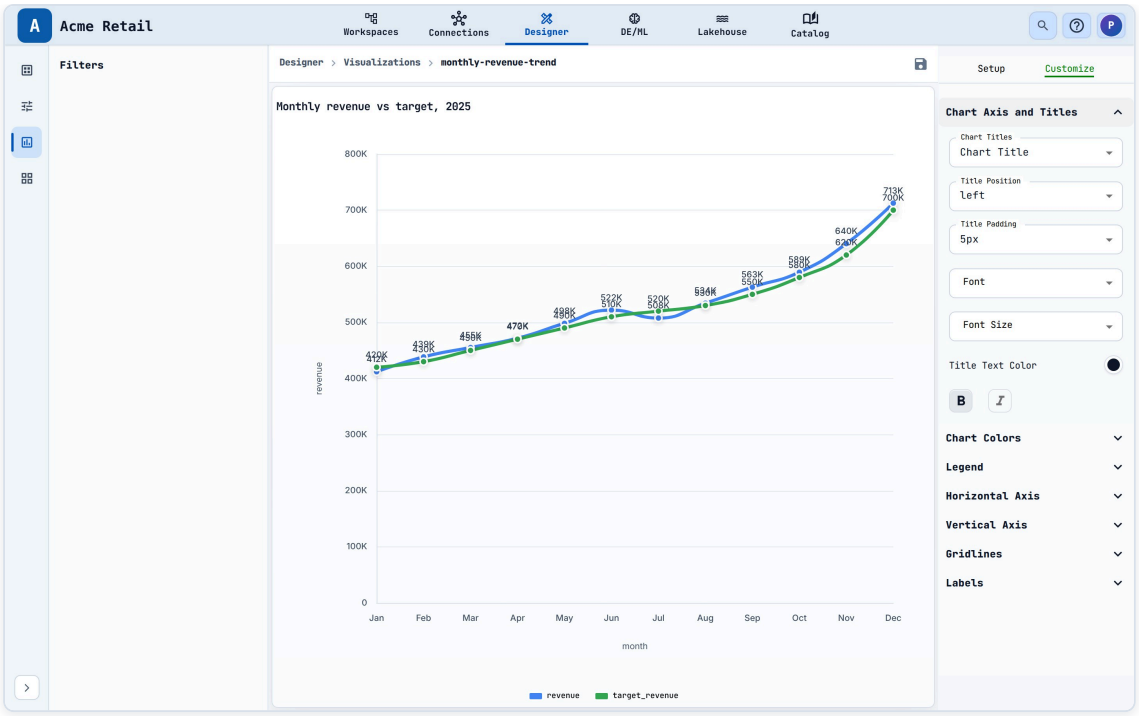


Figure 14. The Customize tab with Chart Axis and Titles open.

Chart	Sections
Bar, Vertical Bar, Combo	Chart Axis and Titles, Chart Colors, Legend, Horizontal Axis, Vertical Axis, Gridlines, Labels, Target line
Stacked Bar	The above plus Total labels and Bullet layout
Stacked Vertical Bar	Bar sections plus Total labels
Line, Bubble	Chart Axis and Titles, Chart Colors, Legend, Horizontal Axis, Vertical Axis, Gridlines, Labels
Scatter	Chart Axis and Titles, Chart Colors, Legend, axes, Gridlines
Candle	Chart Axis and Titles, axes, Gridlines, Candle Stick Colors
Pie	Chart Axis and Titles, Chart Colors, Legend, Labels
Donut, Half Donut	Pie sections plus Total
Funnel	Chart Axis and Titles, Chart Colors, Legend, Labels, Funnel Chart
Gauge	Chart Axis and Titles, Labels, Gauge Chart
Tree Map, Sunburst, Sankey	Chart Axis and Titles, Chart Colors, Labels
Graph	Chart Axis and Titles, Chart Colors, Legend, Labels, Graph Settings
Geo	Chart Axis and Titles, Geo Chart, Labels
Table	Chart Axis and Titles, Header, Chart, Rows, Cells, Cells Color Spectrum, Total, Symbols
Pivot	Chart Axis and Titles, Header, Chart, Rows, Total, Cells, Cells Color Spectrum, Symbols
Scorecard	Chart Axis and Titles, Symbols, Components, Card

What the main sections let you change:

Section	Controls
Chart Axis and Titles	Which title to edit, Title Position, Title Padding, text, font, size, bold, italic, underline, colour, **Hide Title**
Chart Colors	One colour per series; for a pie, **Select Slice** and **Pie Slice Color**; a colour spectrum with **Apply color spectrum** and **Inverse Spectrum**
Legend	**Legend Position** (None, Top, Bottom, Right), Legend Padding, font controls
Horizontal / Vertical Axis	Label font, size and colour; **Text Direction** (Horizontal, Vertical, Inclined) on the horizontal axis; **Label Format** (None, Long, Short, Custom) with Prefix, Suffix, Decimal Places, **Use Thousands Separator**, negative-value format; **Reverse Axis Order**
Gridlines	**Gridline Color**, **Show GridLines**
Labels	Slice Label, Label Position (Inside or Outside), Label Angle, label format, Percentage Decimal Places, font, **Hide Label**, Text Border
Target line	**Line thickness (px)**, **Tick span (px)**; needs a target column in Setup
Total (Donut)	Font, size, bold, italic, Text Color; turn on **Show Total** in Setup first
Total labels	**Distance from bar (px)** and font controls
Bullet layout	**Enable bullet layout**, Scale (100% of total or Actual values), **Combine categories into one bar**, **Show value list**, **Show target**, Target value, Target label
Gauge Chart	Color Ranges (Start, End, colour), Axis Label, Value Font
Candle Stick Colors	Rising Color, Falling Color
Graph Settings	Node Shape, Node Size, Edge Width, Edge Curveness, Edge Opacity, **Enable zoom and pan**, **Draggable nodes**, Node Color, Edge Color

1.5 Let readers drill into detail

Outcome: clicking a bar, slice or point opens a second, more detailed chart for the value that was clicked.

1. Build and save both charts. The detail chart must use the same chart type as the first one.
2. In the first chart's **Setup** tab, under **Drilldown Setup**, choose the **+** button. Only one level is allowed, so the button is disabled once a level exists.
3. In **Choose Visualization** pick the detail chart.
4. In **Select Attribute** pick the column whose clicked value is passed on. Graph charts do not show this field because they pass the clicked node name.
5. In **Filters** pick the detail dataset's filters that should receive the value.
6. Save.

How to check it worked: click an element in the chart. The detail chart opens. The back arrow (**Back to parent chart**) returns you.

To remove the setup, choose the red trash icon and confirm "Do you really want to delete the configuration ?".

Message	What it means and what to do
Please save valid visualization first	The chart does not exist yet. Save it, then add the level.

Message	What it means and what to do
Drilldown configuration is not allowed in drilldown mode	You are already inside a drilled chart. Go back to the parent chart.
To toggle between setup and customize, please exit drilldown	Tabs are locked while a drilled chart is open. Go back first.
To save chart configuration, please exit drilldown	Saving is blocked while a drilled chart is open. Go back first.
Error occurred while fetching drilldown data	The detail dataset could not be queried. Open that dataset and run it.

1.6 Rename or delete a visualization

Open **Actions** on the row.

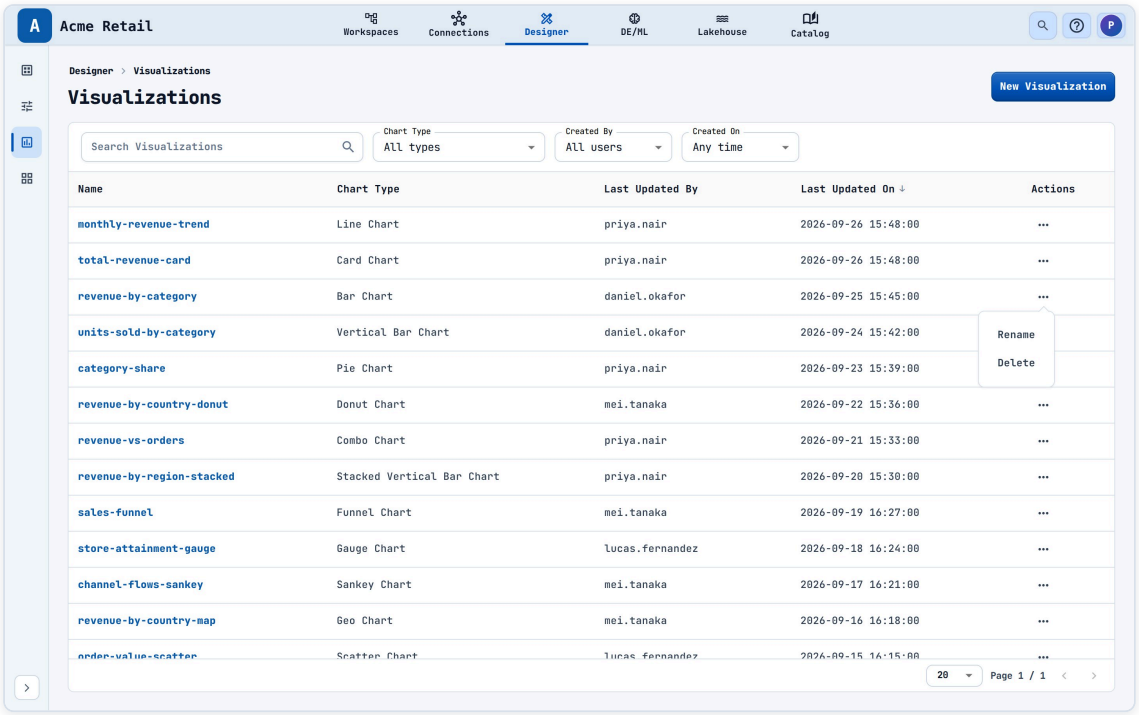


Figure 15. Row actions: Rename and Delete.

- **Rename** opens a dialog with **Visualization Name**, filled in already. Names follow the same rule. Toast: "<new name> renamed successfully".
- **Delete** asks `Are you sure you want to delete "<name>" ?` Toast: "<name> deleted successfully".

> **WARNING**
> Dashboards that still use a deleted visualization cannot draw that tile. They show "Failed to fetch data: <chart> (Id: <id>) may be deleted or inaccessible". Remove or replace the tile on each dashboard.

1.7 What to know

- A chart can have one drilldown level.
- Large tables and pivots load in pages (500 rows at a time for table queries).

- A visualization works with the dataset's filters. Readers' filter choices are not saved with the chart (chapter 2).

1.8 Troubleshooting

- **The canvas says "Complete Chart Setup".** The chart is missing a type, a dataset or a required field. Choose all three in the Setup tab.
- **The Series list is empty.** The dataset returned no numeric columns. Open the dataset and check its query.
- **"Error occurred while fetching data."** The dataset query failed. Open the dataset and run it to see why.
- **I changed something and it is gone.** You did not choose **Save**. Choose the floppy-disk icon before leaving.
- **"Error occurred while updating visualization".** The save failed. Check your connection and try again. If it keeps failing, ask your administrator to check your Visualization permissions.
- **The list says "Unable to fetch data".** Reload the page.

1.9 Related chapters

- Chapter 2 explains filters that appear in the left pane.
- Chapter 3 places visualizations on dashboards.
- The Build and Run guide explains datasets.

2 Saved filters

A **saved filter** is a reusable control, such as a drop-down or a slider, that you define once and attach to datasets. When a reader changes the control, the dataset query runs again with the new value and every chart on that dataset redraws. The same filter can appear on visualizations and on dashboards.

2.1 Choose the right filter type

Pick the type that matches how readers think about the value.

Type	Control	Use it to
Single select List	Drop-down, one value	Choose one region
Multi select List	Drop-down, many values	Choose several regions
Range	From and To boxes	Keep orders between two numbers or dates
Single Value	One input	Set a threshold
One-way slider	Slider with one handle	Set a minimum price
Two-way slider	Slider with two handles	Pick a price band
Toggle	Switch	Include or exclude returns

You create filters in **Designer > Filters**. A filter only changes a chart once a dataset query uses it (see "Make a filter change a chart").

2.2 Open the Filters list

Open **Designer** and choose **Filters** in the left rail.

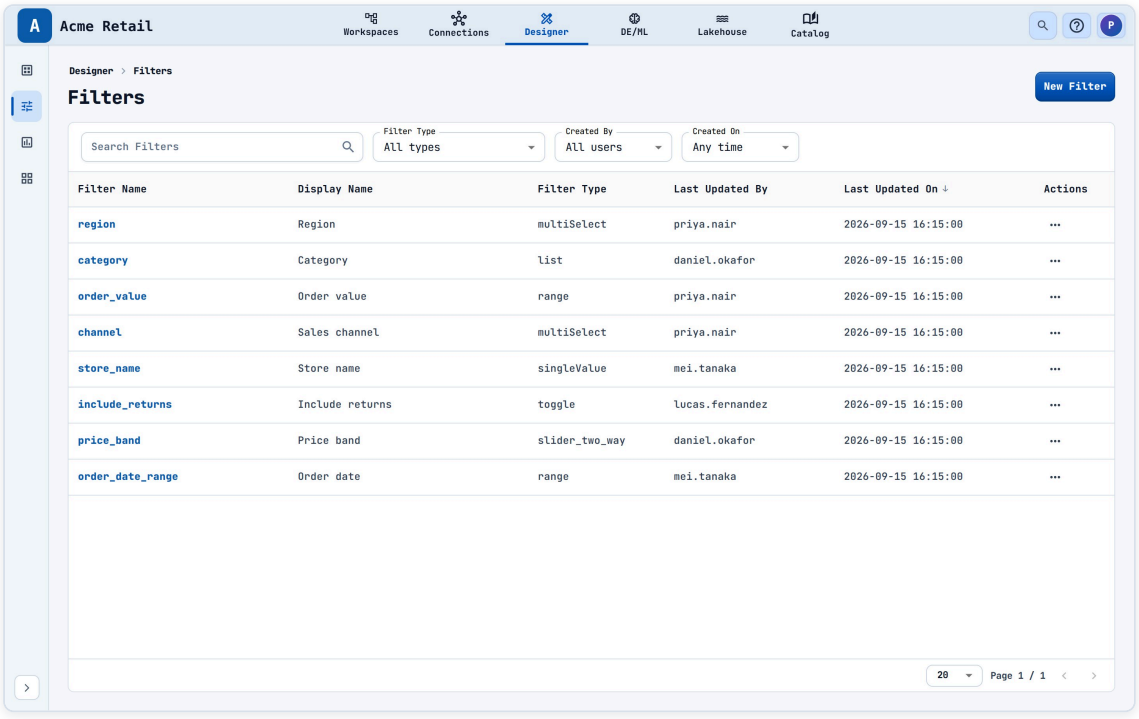


Figure 16. The Filters list with Filter Name, Display Name and Filter Type.

The controls match the visualizations list: **Search Filters**, **Filter Type**, **Created By**, **Created On** and **New Filter**. Columns are **Filter Name** (a link that opens the editor), **Display Name**, **Filter Type**, **Last Updated By**, **Last Updated On** and **Actions** (Rename, Delete). With no filters you see "No filters to show".

2.3 Create a filter

Outcome: a named filter, ready to configure.

1. Choose **New Filter**.
2. In **Enter Name** type a name that follows the naming rule from the introduction. Use what the filter means, for example `region`.
3. Optionally fill **Display Name (optional)**, for example `Region`. Dashboards and datasets show the display name instead of the name.
4. Choose **Add**. The filter is created empty. Open it to configure it.

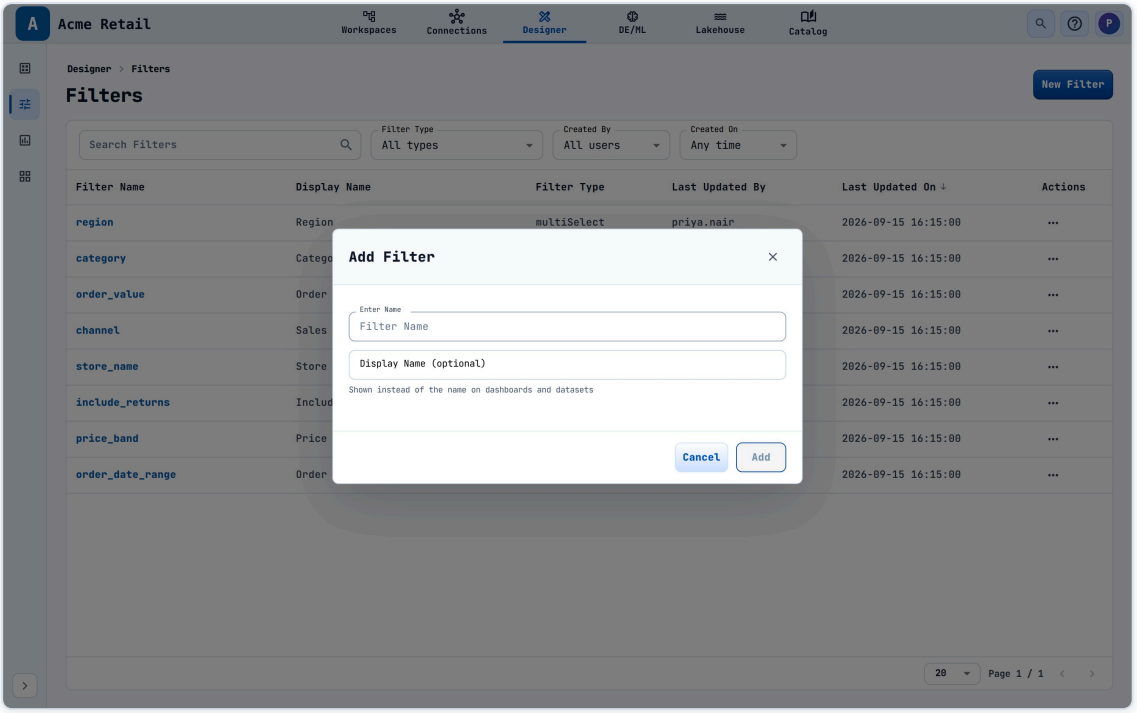


Figure 17. The Add Filter dialog.

> **TIP**
> Name filters for what they mean and set a **Display Name**. Readers see the display name.

2.4 Configure a filter

1. Choose the filter's name in the list to open its editor.
2. Under **TYPE AND SOURCE**, choose the **Type** and the **Source**.
3. Fill the **CONFIGURATION** section (it depends on the source and type, and is hidden for Toggle).
4. Choose **Save**. The toast says **Filter updated**.

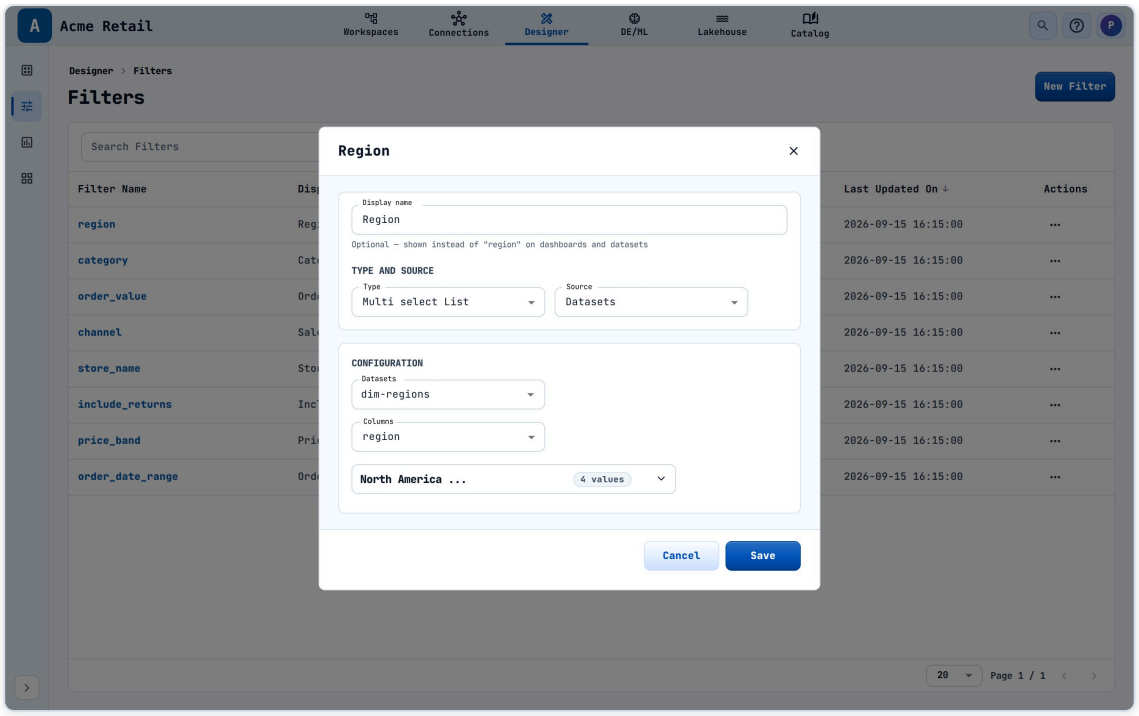


Figure 18. The editor for a dataset-sourced Multi select List filter.

The **Source** is either:

- **Custom**: you type the values yourself.
- **Datasets**: Groundzero runs a dataset and uses one of its columns to fill the options. This is available for Single select, Multi select, One-way slider and Two-way slider.

2.4.1 Fields for a custom source

- Choose a **Data Type**: Number or Plain text for lists; Number or Date for Range; Number or Date Time for sliders.
- **Single and Multi select:** add rows with **Value** and **Display Value** and the **+** button. **Edit Values** lets you change or delete rows. Pick the default selection.
- **Range: From** and **To**. Date ranges add **Choose Date Formats** (and Fiscal Month for quarter formats).
- **Single Value: Value**.
- **Sliders: Start**, **End**, **Step Type**, **Steps**, **Date Format**, **Fiscal Month**.
- **Toggle:** a checkbox for the default state.

2.4.2 Fields for a datasets source

- Choose a **Datasets** entry and a **Columns** entry. Range and slider types only offer number and date columns.
- Use the **Display values** editor to map raw values to friendly text.

2.4.3 If Save is refused

Message	What to do
Please enter any value on the values thing.	Add at least one value or display value.
Default value should be chosen for single select filter before saving	Pick the default item.
At least one value should be chosen for multi select filter before saving	Pick at least one item.

Message	What to do
Please enter start value and end value both to save the filter.	Fill both ends of a slider.
Please enter start date and end date both to save the filter.	Fill both dates.
Please enter valid from and to values.	Fix a malformed range.
From value should be less than or equal to To value.	Swap the values.
Start date should be less than end date.	Fix the dates.
Please select a date format before saving the filter.	Choose a format.
End cannot be lesser than start	Fix the slider ends.
Both Display Value and Value are required.	Fill both fields before adding a row.
This value already exists.	Remove the duplicate.
Value must be a valid number.	Enter digits only for a numeric filter.

If the save fails you see "An error occurred while updating the filter.". If the filter cannot be opened you see "Filter not found. Please try refreshing the page."

2.5 Make a filter change a chart

A filter takes effect on a dataset when the dataset's query contains the filter's **placeholder**, a short tag in double curly braces that Groundzero replaces with the reader's choice before the query runs.

1. Open the dataset (Build and Run guide) and attach the filter in the **Dataset editor**. The editor lists the dataset's filters.
2. Hover the filter in that list and choose the copy button. The toast says `copied to clipboard {{name.value}}`.
3. Paste the placeholder into the query where the value belongs, and save the dataset.

Filter type	Placeholder
Single select List, Single Value, Multi select List, Toggle	`{{name.value}}`
Range and sliders	`{{name.fromValue}}` and `{{name.toValue}}`

Groundzero replaces the placeholder with the reader's choice before it runs the query:

Control	Text that replaces the placeholder
Single select	`('Apparel')`
Multi select	`('A','B')` (numbers are not quoted)
Nothing selected	`('')`
Range or slider numbers	The bare numbers
Dates	Formatted with the chosen date format
Toggle	`(True)` or `(False)`

How to check it worked: open a visualization on that dataset and change the filter. The chart redraws.

> **NOTE**
>
> A filter does nothing until a query contains its placeholder. If a filter seems to have no effect, check the query first.

2.6 Use filters on a chart

A visualization shows the filters attached to its dataset in the **Filters** pane on the left, one control per filter. Names longer than 20 characters are shortened. Hover to see the full name. Filters marked hidden in the dataset still change the query but are not drawn.

Changing a value re-runs the dataset and redraws the chart. A filter whose options come from a dataset query that uses other filters narrows its options when those filters change. For example, choosing a country can shorten a city list. This happens on its own.

The values a reader chooses are not saved with the visualization. Only the defaults saved in the filter apply when the chart opens.

2.7 Related chapters

- Chapter 1, the visualization editor where filters appear.
- Chapter 3, the dashboard filter bar.

3 Dashboards

A **dashboard** arranges visualizations, text, shapes and filter controls on a free-form canvas, organised into tabs. You build dashboards in **Designer > Dashboards**, **publish** them, and readers open the published copy from **Workspaces**.

Before you start you need saved visualizations to place on the dashboard (chapter 1) and the Dashboard, Visualization and Workspace permissions (Administration guide, chapter 4).

3.1 Open the Dashboards list

Open **Designer** and choose **Dashboards** in the left rail.

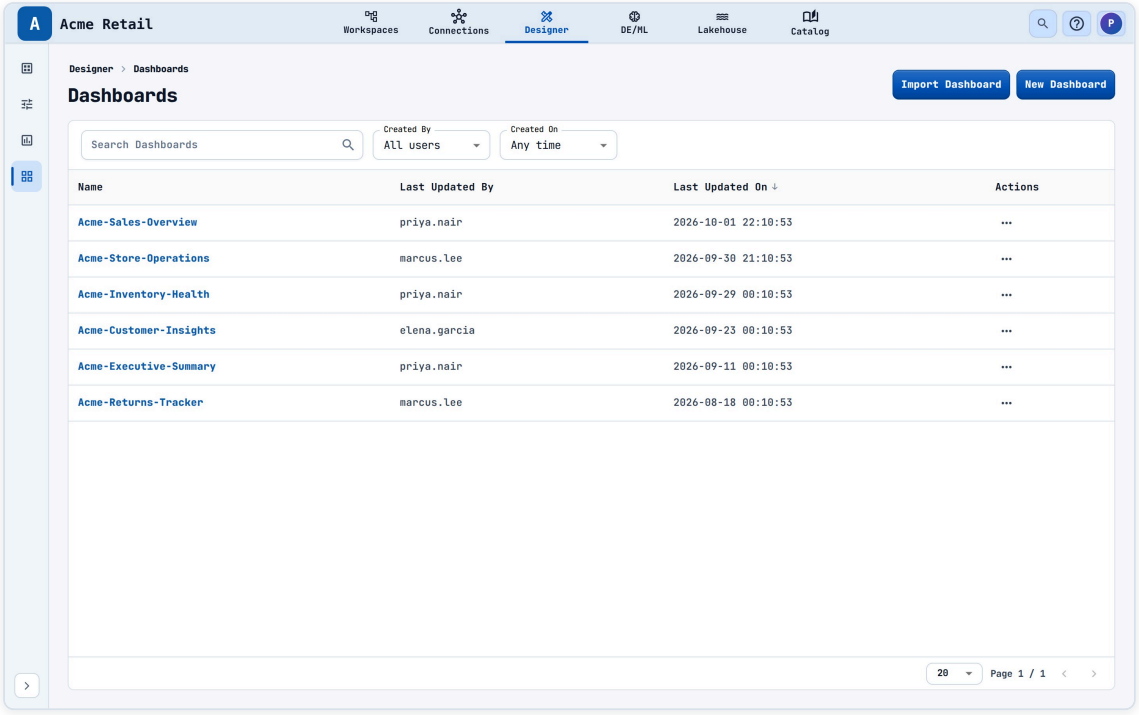


Figure 19. The Dashboards list.

The header has **Import Dashboard** and **New Dashboard**. The toolbar has **Search Dashboards**, **Created By** and **Created On**. Columns are **Name** (opens the dashboard), **Last Updated By**, **Last Updated On** and **Actions**.

3.2 Build a dashboard from your charts

Outcome: a saved dashboard with your charts, laid out the way you want.

1. Choose **New Dashboard**.
2. In **Enter Name** type a name that follows the naming rule, then choose **Add**. The toast says "<name> added successfully". Groundzero does not open the new dashboard.
3. Click the dashboard's name to open the editor.
4. Open the **Select visualizations** tab in the right panel and tick the charts you want. Each one adds a tile to the active tab.
5. Drag and resize the tiles on the canvas.
6. Choose **Save dashboard** in the header. The toast says "Successfully updated".

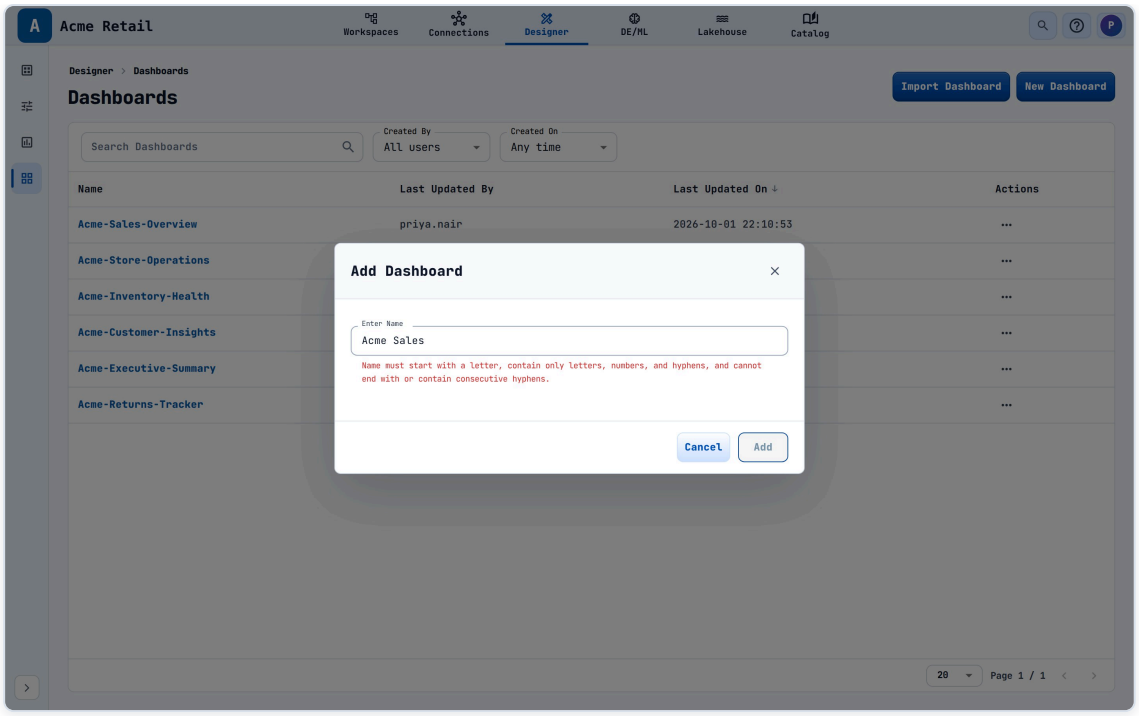


Figure 20. An invalid name is rejected before you can add it.

How to check it worked: close the dashboard and open it again. The layout is as you left it.

> **NOTE**
> Layout changes stay in your browser until you choose **Save dashboard**.

3.2.1 Find your way around the editor

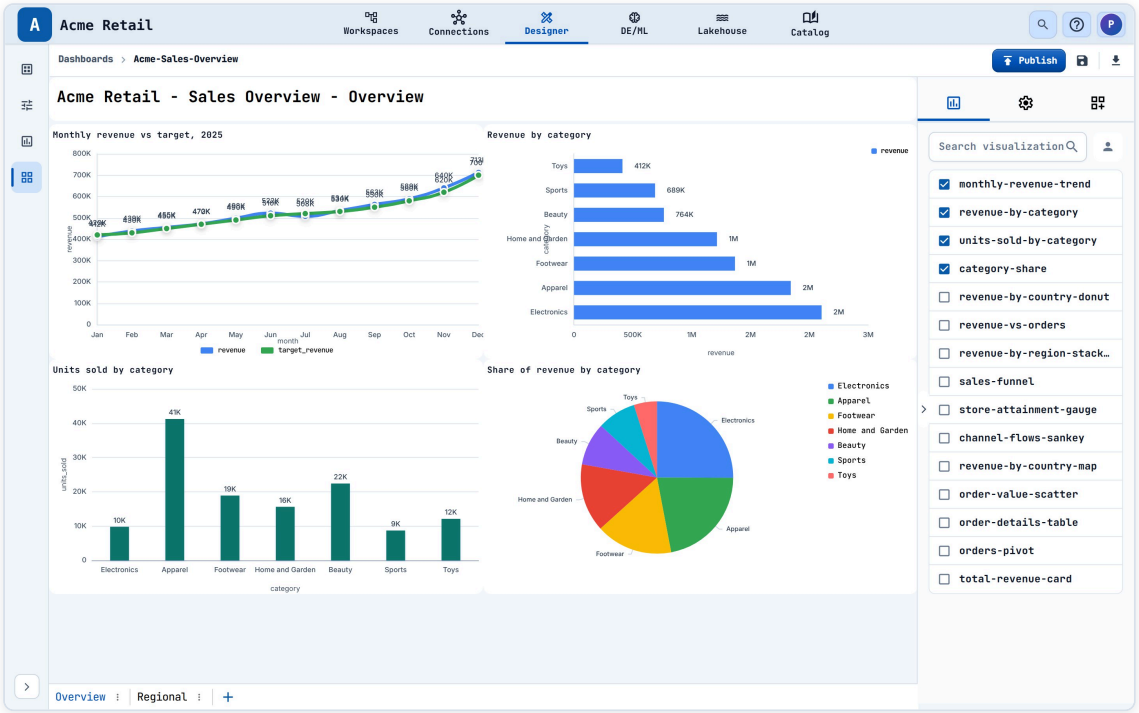


Figure 21. A dashboard in the editor: header toolbar, canvas, right panel and bottom tabs.

Area	What it does
Header (top)	Breadcrumb and toolbar: **Publish**, **Save dashboard**, **Reset filters** (when enabled), **Export dashboard**
Canvas	The active tab. Drag and resize tiles. Tiles are free-form and can overlap
Right panel	Three icon tabs: **Select visualizations**, **Setup** and **Customize**. A slim handle collapses it
Bottom bar	The tabs of the dashboard, a menu on each tab, and **+** to add one

3.2.2 Add and remove charts

The **Select visualizations** tab lists every visualization with a checkbox.

- **Search visualization** filters by name.
- The person icon (**Filter by creator**) filters by creator.
- Checking a visualization adds a tile to the active tab. Unchecking removes it. Selected items sort to the top.

When a chart joins the tab, its dataset filters appear as filter controls on the canvas.

Hover a tile for two icons: a **pin** that locks the tile's position and an **X** that removes it. **X** asks `Are you sure you want to remove visualization <name> from this dashboard?`

3.2.3 Organise a dashboard into tabs

Each tab shows its label and a menu (**tab actions**).

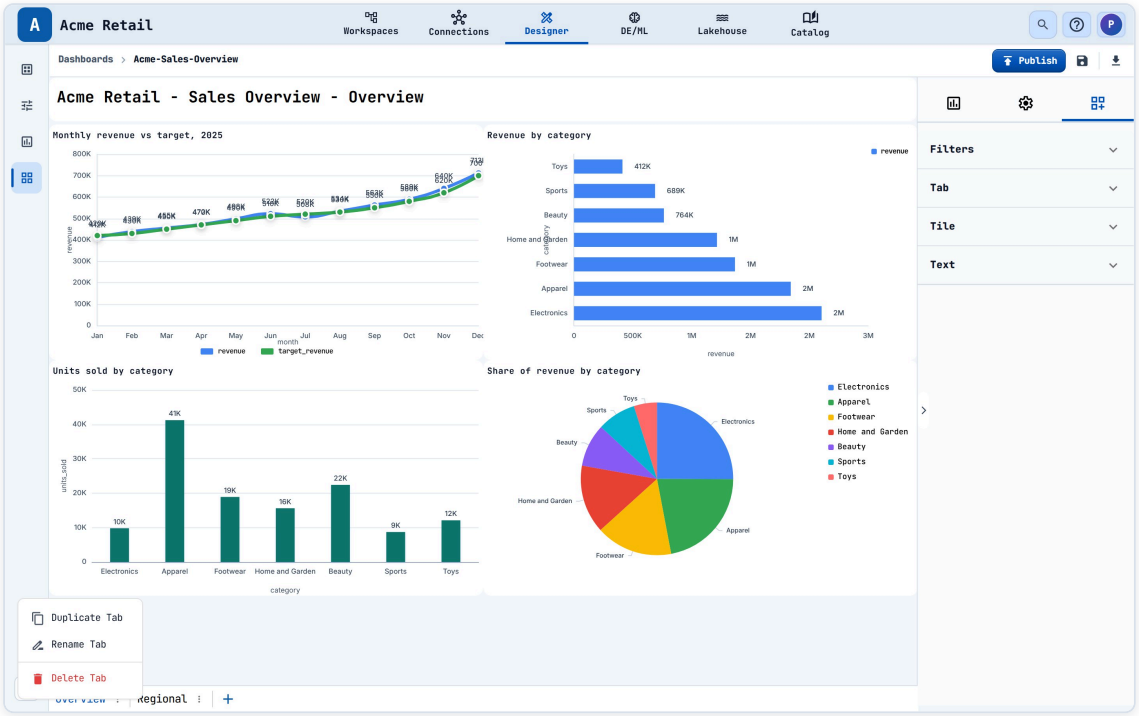


Figure 22. The tab menu: Duplicate Tab, Rename Tab, Delete Tab.

Action	What happens
+ (Add New Tab)	Enter a name. An empty or repeated name (any capitalisation) shows "Tab name already exists. Please choose a different name."
Rename Tab	Shows "Tab name is invalid or already exists. Please choose a different name." when refused

Action	What happens
Duplicate Tab	Copies layout, tiles, shapes and filter styles into a new tab named "<label> (Copy)", "(Copy 2)" and so on
Delete Tab	Asks you to confirm. The only remaining tab cannot be deleted: "Cannot delete the only tab. A dashboard must have at least one tab."

A dashboard with no content shows "Dashboard is Empty" and asks you to create a tab with the **+** icon.

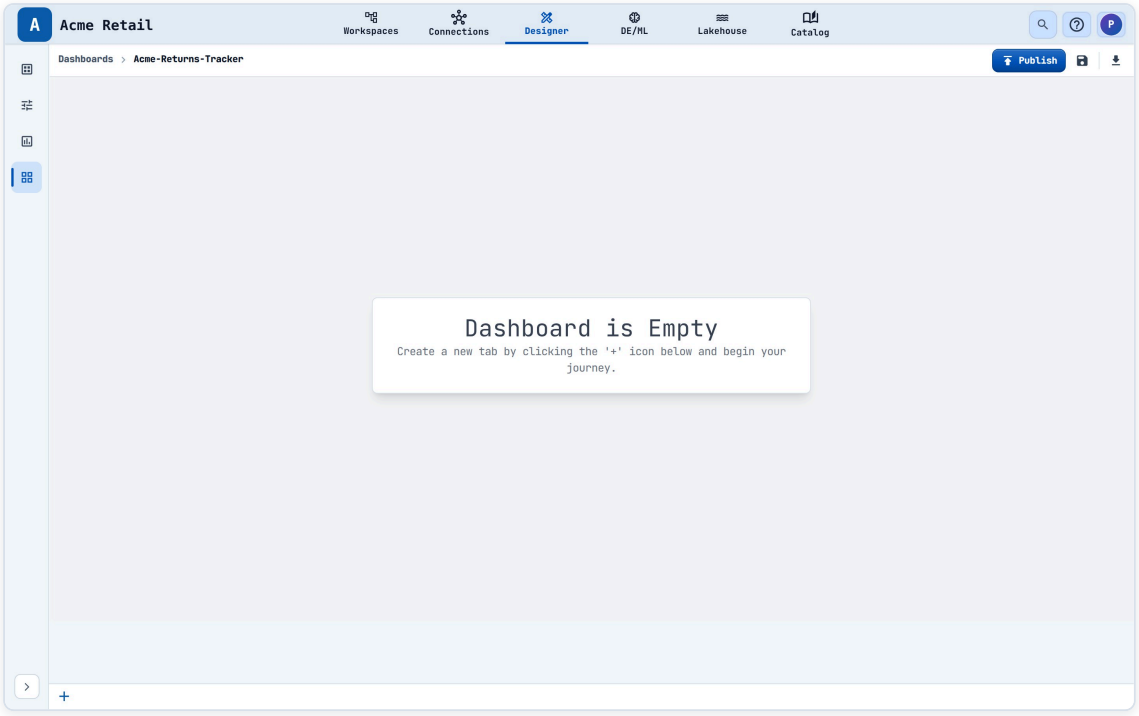


Figure 23. An empty dashboard.

3.2.4 Style tiles and tabs

Open the **Customize** tab. Click a tile to select it (a dashed blue outline). Ctrl or Cmd click selects several.

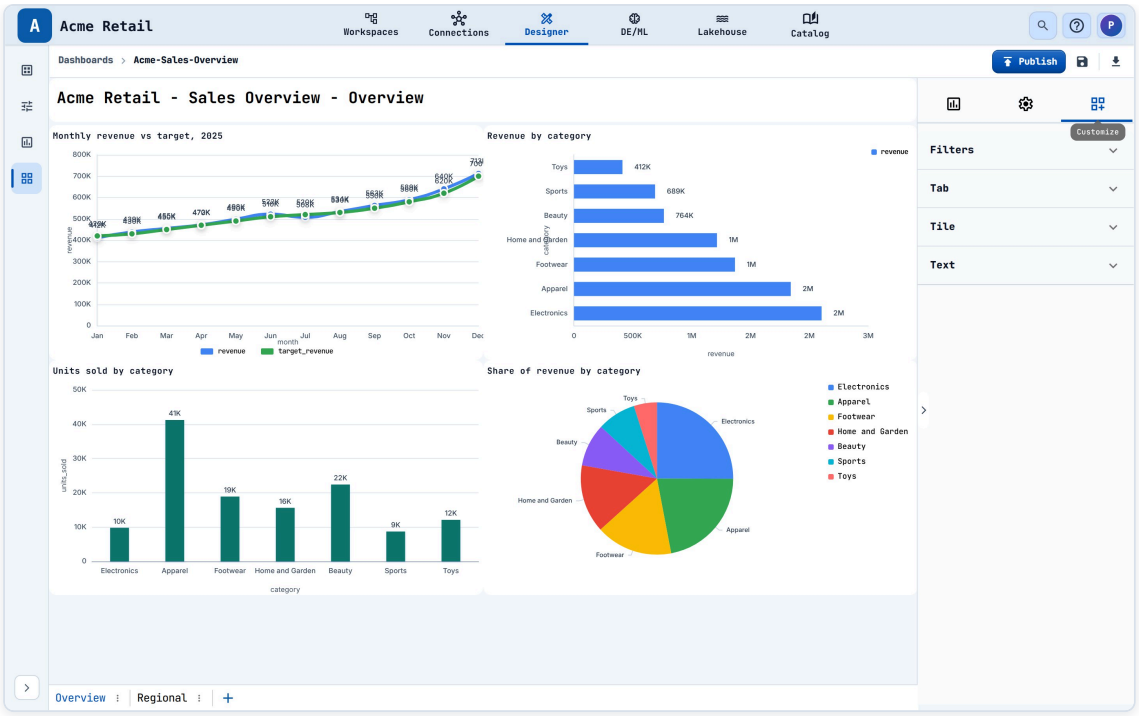


Figure 24. The Customize tab: Filters, Tab, Tile and Text.

Section	Controls
Filters	Switch **Reset filters button** adds a header control that restores every filter to its saved default
Tab	**Canvas background color**, **Background image** (upload or delete)
Tile (selected tiles, shapes and filters)	Border color, Background color, Background image, Border style (Solid, Dotted, Dashed, Groove, Double), Border width, Corner radius, Padding, Row height, Row width, Horizontal position (X), Vertical position (Y)
Text	Align (left, center, right), Font color, font family, font size 10 to 36, bold, italic

3.2.5 Add text and shapes

In the **Setup** tab, open **Text and Shapes** to add a shape to the active tab: circle, textbox, rectangle, rounded rectangle, line, dotted line, several arrows, triangle, pentagon, hexagon or diamond. A new shape appears at the top left. Drag and resize it. A **textbox** is a free text widget: click it and type. Hover a shape for the red X (**Remove Shape**) and for ordering icons.

3.3 Let readers click a chart to filter or jump

Outcome: clicking a value in one chart filters other charts, switches to another tab, or sends the chart behind the others.

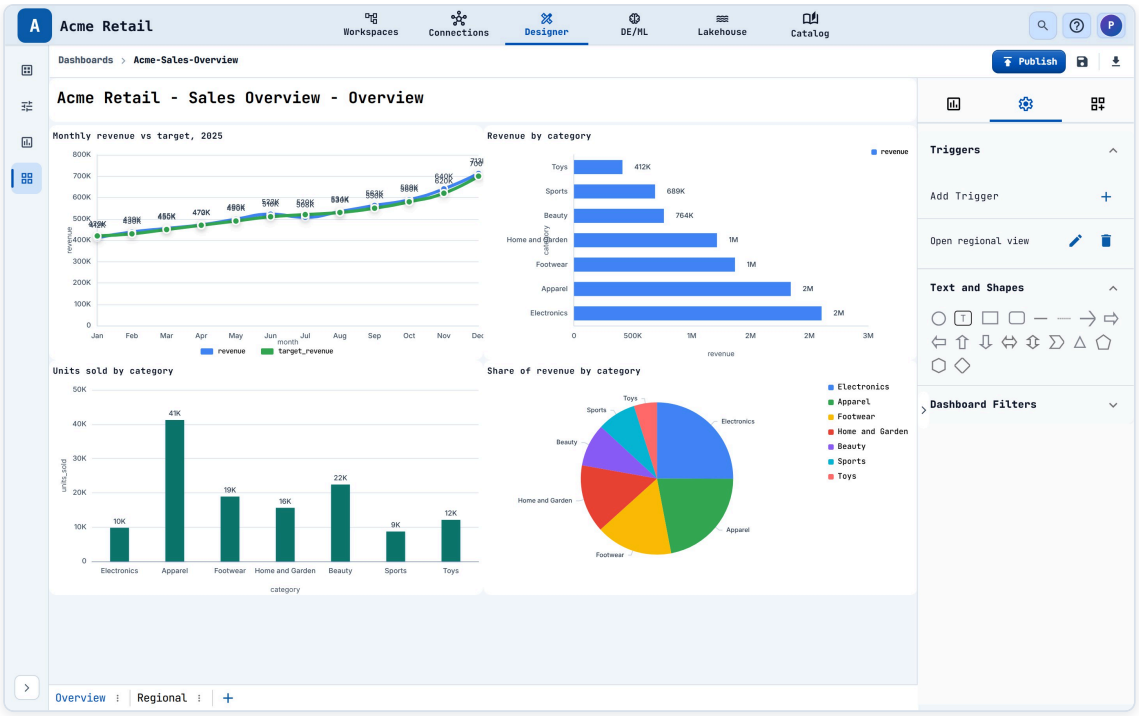


Figure 25. The Setup tab with Triggers and Text and Shapes open.

1. In the **Setup** tab, choose **+** next to **Add Trigger**. The **New Trigger** dialog opens.
2. Fill in the fields below.
3. Confirm the dialog, then choose **Save dashboard**.

Field	Meaning
Name	Unique within the tab. A repeat shows "Trigger name already exists in this tab."
Event	**Filter**, **Navigate**, **Filter and Navigate** or **Send to back**
Source	The chart whose click starts the trigger
Select Attribute	The column whose clicked value is used (hidden for Send to back)
Target > Select Tab	Where to go. For Filter it is fixed to the current tab
Target > Select Charts / Select Filters	What to filter (hidden for Navigate)

How to check it worked: click the source chart. The target charts filter, the tab changes, or the source tile moves behind the others. After **Send to back**, a **Back to Original** arrow brings it back.

3.4 Let readers switch between stacked charts

Dashboard Filters add a single-select control that is not tied to a dataset column. Choosing a value brings one chart to the front, so you can stack several tiles in one spot and switch between them.

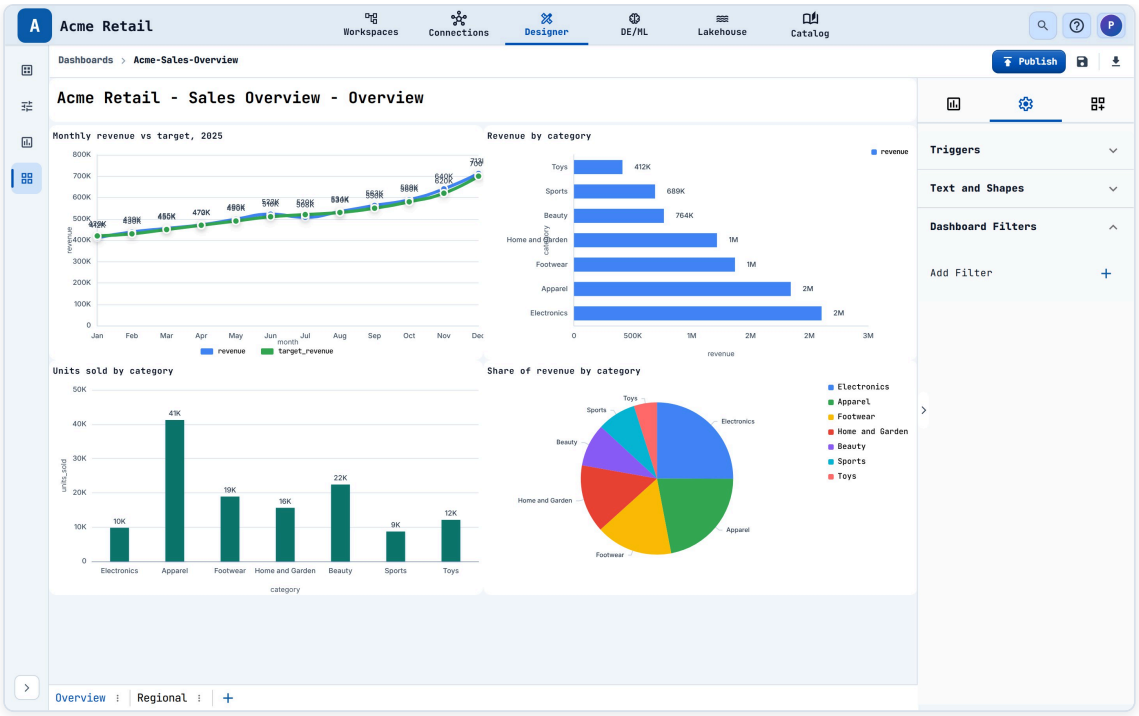


Figure 26. The Dashboard Filters accordion with an Add Filter row.

1. In the **Setup** tab, open **Dashboard Filters** and add a filter. The **Configure Dashboard Filters** dialog opens.
2. Type a **Name**.
3. Under **Filter Value**, type the text readers will see, pick the **Chart** that value should bring to the front, and choose **+**. The toast says "Item added successfully".
4. Repeat for each chart. **Edit Display Values** lists the pairs for editing or deletion.
5. Pick the default value. **Add** stays disabled until there is a name and a default value.

3.5 Filters on a dashboard

Every filter attached to the charts on a tab (apart from hidden ones) is drawn as a control that you can move, resize and style. Changing a value re-runs every chart that uses that filter on the tab. **Reset filters** puts every filter and dashboard-filter choice back to what it was when the dashboard loaded.

Readers' filter choices are not saved with the dashboard. What is saved is the Reset filters setting, the dashboard filters you defined and their defaults.

3.6 Publish a dashboard

Outcome: readers can open a read-only copy from Workspaces.

1. Choose **Save dashboard** first. Publishing does not save your latest edits.
2. Choose **Publish**.
3. Confirm `Are you sure you want to publish "<name>"?` Groundzero publishes the dashboard with every visualization and dataset it uses.
4. The toast says "Successfully published".

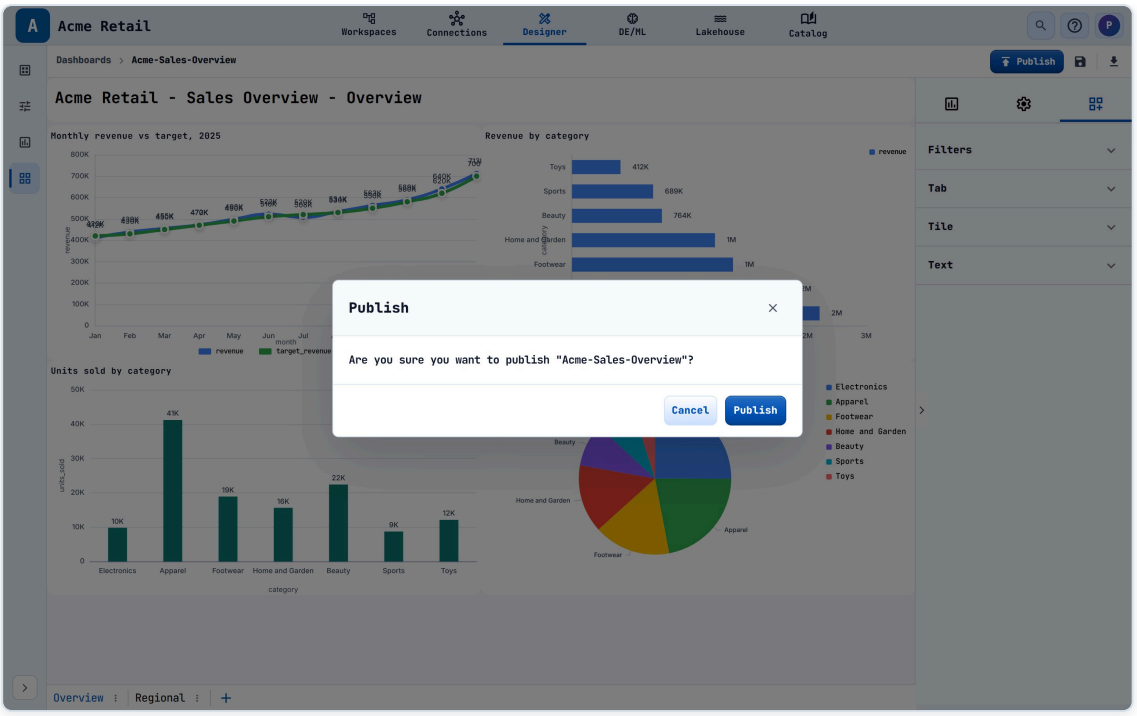


Figure 27. The Publish confirmation.

An empty dashboard shows "Dashboard is empty, add some items and publish." when you try to publish and "Dashboard is empty, add some items and save." when you try to save.

> **WARNING**
>
> Publishing overwrites the published copy that has the same name, and does not save your Designer copy first. Always choose **Save dashboard** too.

What happens next: the published dashboard appears in Workspaces for people to read, and in the Deployment Manager, where administrators can copy it to other sites (Administration guide, chapter 10).

3.7 Export a dashboard to PDF or Excel



Figure 28. The Export menu offers PDF and Excel.

1. Open the tab you want to export. Exports cover one tab at a time.
2. Choose **Export dashboard**, then **Export To PDF** or **Export to Excel**.
3. Wait for the info toast ("Preparing PDF file for download..." or "Preparing Excel file for download..."). The file downloads.

PDF captures the tab as it looks, and is named `<tab label>.pdf`. Excel is named `<tab label>.xlsx` and has one sheet per chart on the tab, with data re-queried using the current filters. Sheet names longer than 31 characters are shortened.

3.8 Read a published dashboard in a workspace

Workspaces in the top navigation lists folders called workspaces. Each holds published dashboards.

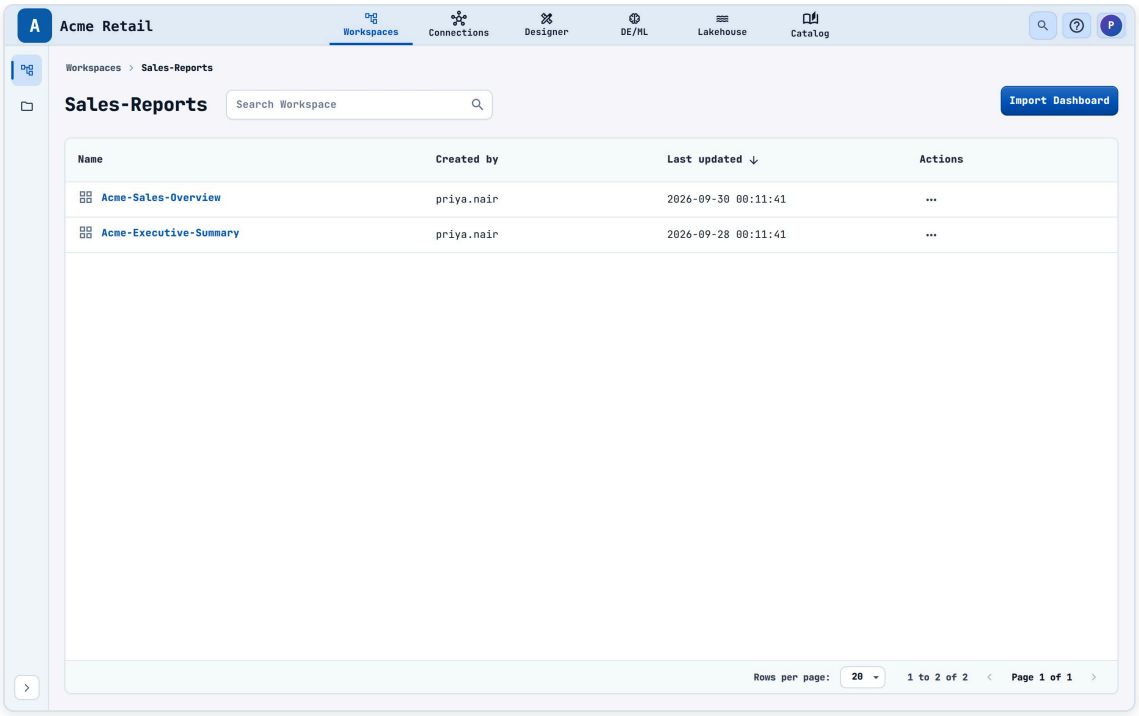


Figure 29. A workspace listing its dashboards.

1. Open **Workspaces** and choose a workspace. You see its dashboards with **Created by** and **Last updated**.
2. Click a dashboard to open the read-only viewer. Use the tabs, **Reset filters** (if enabled) and **Export**. Tiles are locked.

To manage workspaces:

- **New Workspace** opens **Add Workspace** and uses the same naming rule.
- A workspace's **Actions** menu offers **Rename**, **Delete** and, for administrators only, **Change Owner**.
- Inside a workspace, Rename and Delete are available per dashboard.

3.9 Copy a published dashboard to edit it

Import Dashboard copies a published dashboard into the Designer so you can change it.

1. In **Designer > Dashboards**, choose **Import Dashboard**.
2. In **Select a dashboard** pick a published dashboard.
3. Choose **Import**.

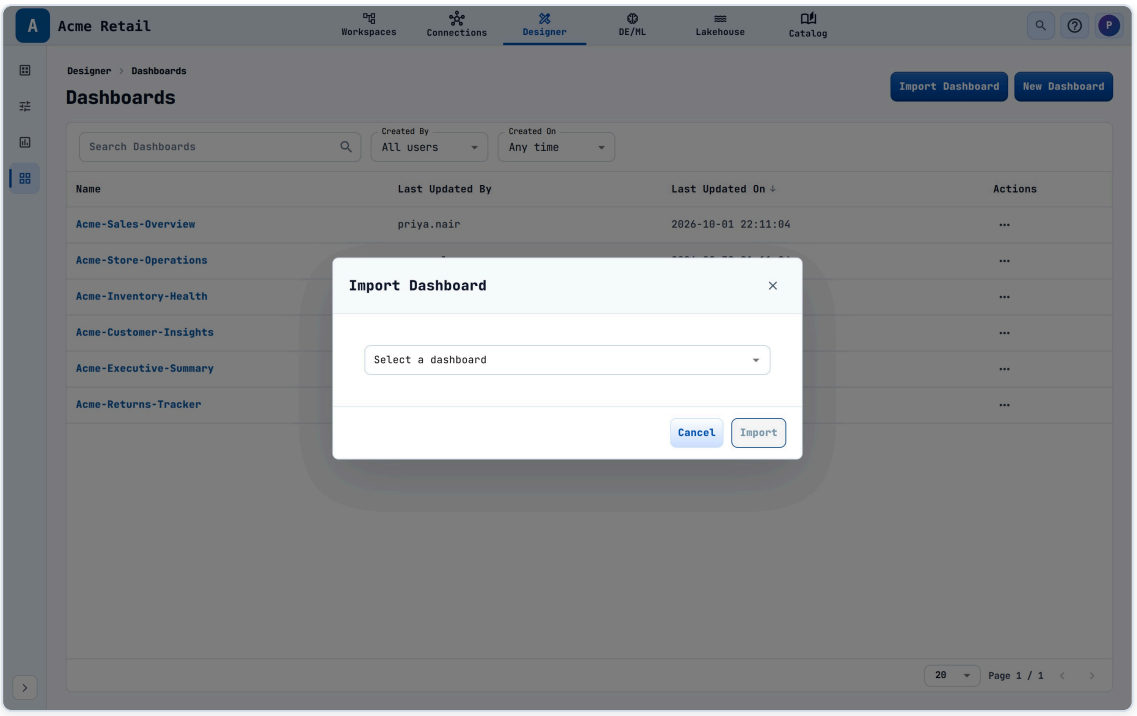


Figure 30. The Import Dashboard dialog.

Groundzero makes copies of every dataset, custom filter and visualization the dashboard uses, each with a unique ending added to its name, and creates a dashboard named `<published name>-<yyMMddHHmmss>`. The toast says "Dashboard imported as <new name>".

3.9.1 Import into a workspace

1. Open the workspace and choose **Import Dashboard**.
2. Pick a published dashboard and choose **Import**.

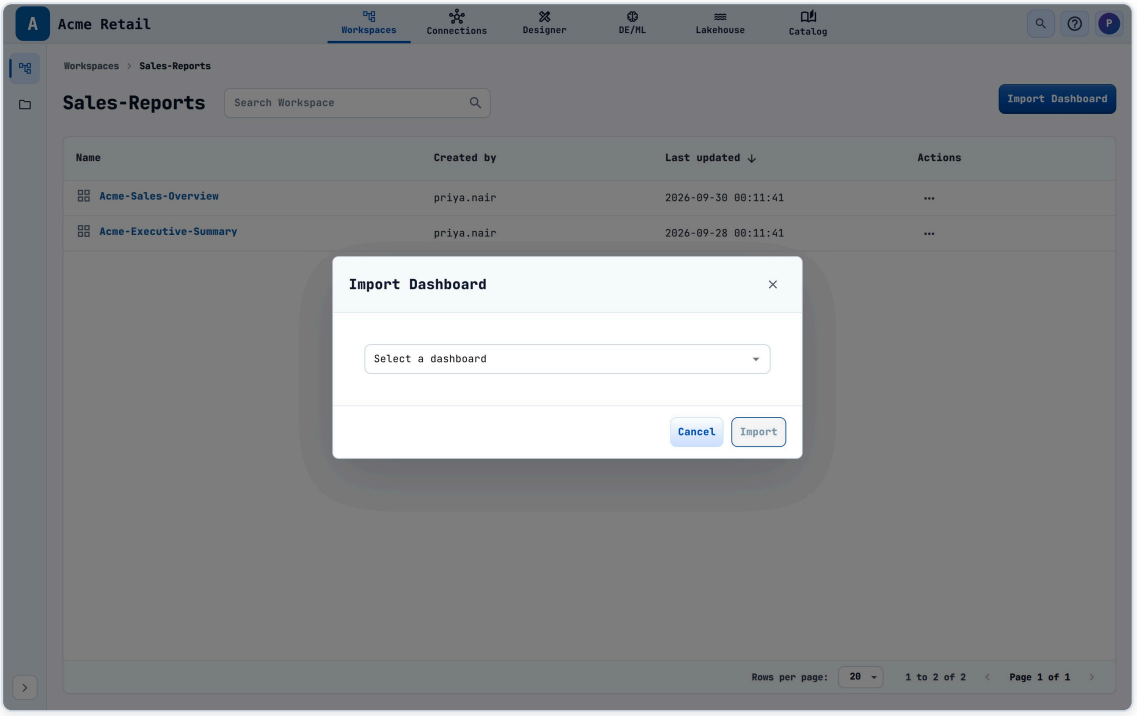


Figure 31. Importing into a workspace.

If a dashboard with the same name is already in that workspace, the **Dashboard Name Conflict** dialog appears:

Button	Result
Replace	Overwrites the existing dashboard. Toast "<name> replaced successfully"
Keep Both	Imports under `<name>-<yyMMddHHmmss>`
Cancel	Does nothing

When you import several published dashboards at once, a **Do it for all** checkbox applies one choice to every conflict.

Administrators can also copy published dashboards to other sites (Administration guide, chapter 10).

3.10 Rename or delete a dashboard

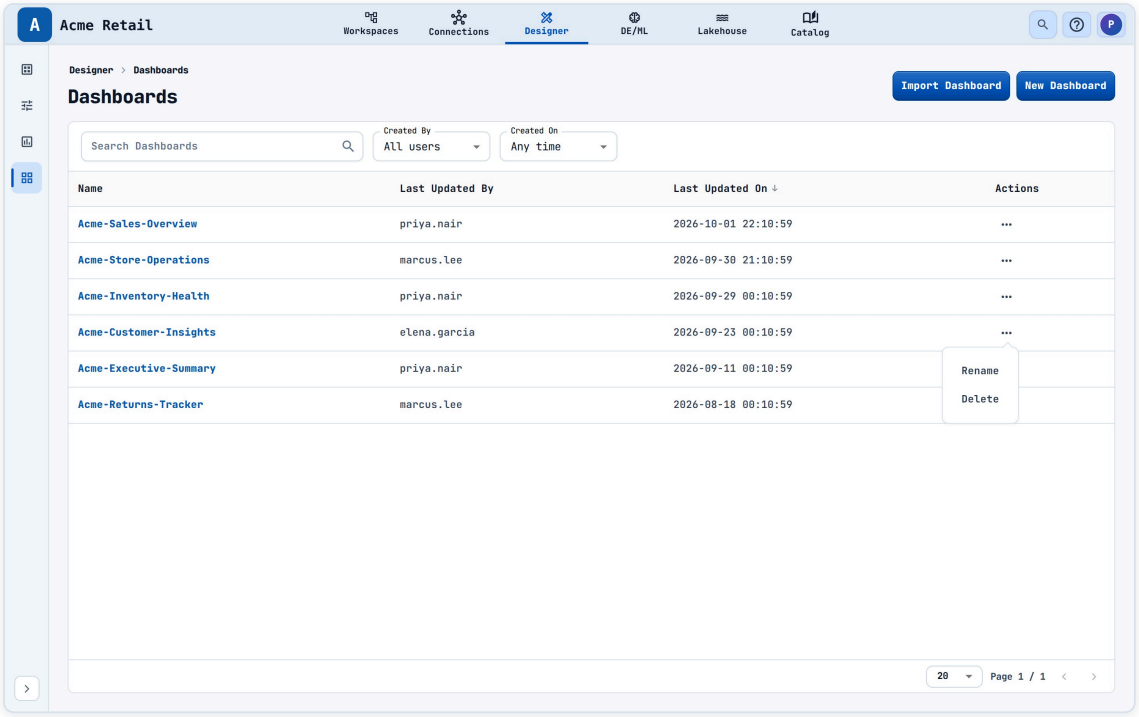


Figure 32. Row actions: Rename and Delete.

Rename fills in the current **Dashboard Name** and applies the naming rule. **Delete** asks `Are you sure you want to delete "<name>" ?`.

3.11 Troubleshooting

- **A tile shows "Failed to fetch data: <chart> (Id: <id>) may be deleted or inaccessible".** Cause: the visualization was deleted, or you do not have permission to read it. Fix: add the chart again, or ask an administrator for access.
- **A tile shows "Dataset could not be loaded".** Cause: the dataset behind the chart failed. Fix: open the dataset and run it.
- **Tiles overlap unexpectedly.** Cause: tiles are free-form and can overlap. Fix: use the **Tile** controls (X, Y, width, height) or drag the tile.

- **The Reset filters button is missing.** Cause: it only appears when the **Reset filters button** setting is on and at least one filter differs from its default. Fix: turn the setting on in Customize.
- **A published dashboard still shows old data in a workspace.** Cause: the published copy is older than your Designer copy. Fix: choose **Save dashboard**, then **Publish** again.
- **"Error occurred" when saving.** Cause: the save failed. Fix: try again, and ask an administrator to check your Dashboard permissions if it repeats.

3.12 Related chapters

- Chapters 1 and 2 for the content of tiles and filters.
- The Administration guide, chapter 4, for permissions, and chapter 10 for copying dashboards between sites.

4 AI chat assistant

The **AI chat assistant** (Groundzero AI) is the home page of Groundzero. You describe what you want in plain language and it answers with text, tables, charts and links to things it created for you: datasets, dashboards, ETL jobs, filters and more. When an action would change something, it stops and asks you to approve first.

Use it to:

- Ask questions about your data, pipelines, jobs and errors.
- Draft SQL, datasets, filters, visualizations and dashboards.
- Attach a file and ask questions about it.
- Use tools on your saved MCP connections, with your approval.

Your account needs the **AI Chat** permission (**All**). Without it the assistant shows the card **You don't have access to chat** with "Your account doesn't have permission to use the chat feature. Contact your administrator to request access." Your administrator may also set a rolling 30-day spending budget for you or your group (Administration guide, chapter 7). When it is used up, chat shows **Usage limit reached**.

4.1 Ask a question

Outcome: an answer about your data, in the chat.

1. Choose the site logo in the header to go to the home page.
2. Click **Message Groundzero AI** (it says "Type your message here...") and type your request. Press **Shift+Enter** for a new line. The box grows as you type.
3. Press **Enter** or choose **Send message**.

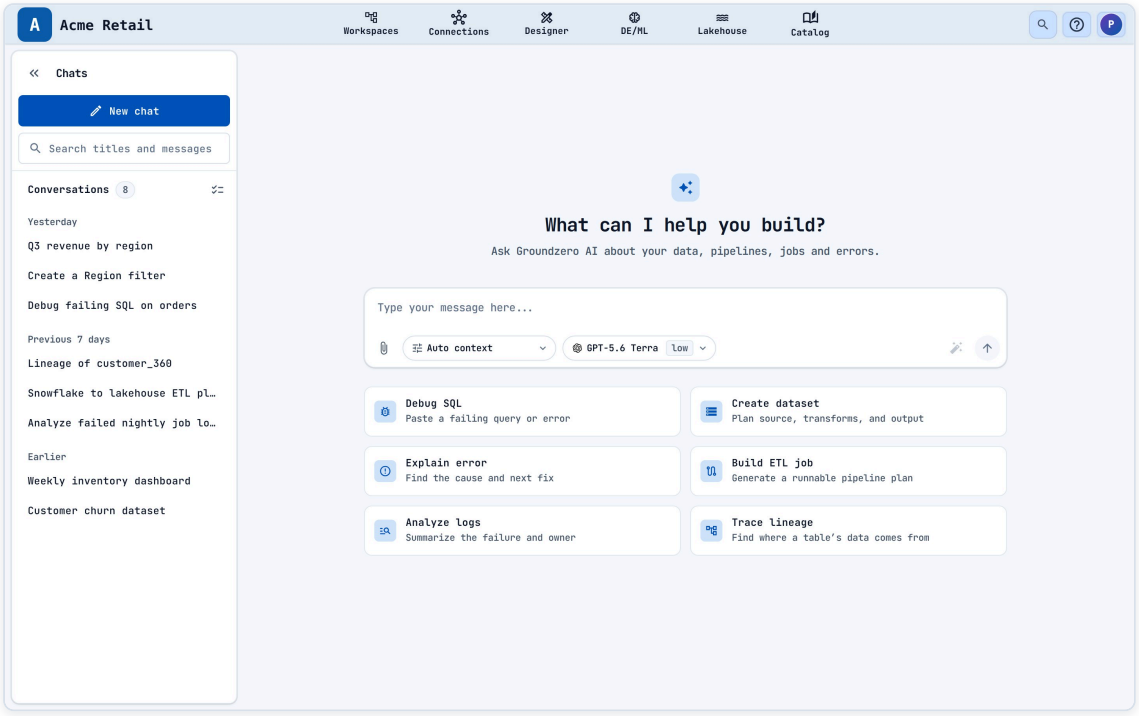


Figure 33. The welcome screen: composer, model and context controls, starter prompts and the Chats sidebar.

The page has a **Chats** sidebar on the left and the conversation on the right. The sidebar starts collapsed. Choose **Expand sidebar** to open it. Groundzero remembers your choice.

Six **starter prompts** fill the box when you click them but do not send: Debug SQL, Create dataset, Explain error, Build ETL job, Analyze logs and Trace lineage.

The buttons in the box, left to right, are **Attach files**, the **chat options** pill (for example "Auto context"), the **model and thinking** pill, **Enhance prompt with AI** and the send button.

The send button changes with what is happening:

Send button	When you see it
Send message	Idle. Disabled when there is nothing to send, a file is uploading, the model has a problem, a scope is required, or an approval is waiting
Queue message	A response is streaming and you typed something. Queued messages send in order afterwards
Send steering note	A run is working and live steering is on
Stop response	A response is streaming and the box is empty

A message can be up to 200,000 characters.

4.1.1 Improve a prompt before you send it

Choose **Enhance prompt with AI**. It is enabled when the box is not empty and nothing is streaming. Groundzero rewrites your text in place using the selected model, up to 8,000 characters. If it fails you see "Could not enhance the prompt. Try again."

4.1.2 Attach files

Choose **Attach files**, drag files onto the box ("Drop to attach") or paste them.

Limit	Value
File types	`.csv .xlsx .xls .pdf .txt .jpg .jpeg .png .webp .gif`
Size per file	50 MB
Files per message	100

Each file shows **Uploading**, **Ready** or **Queued**. Remove one with **Remove <name>**. If you attach files without text, Groundzero sends a prompt for you.

If something goes wrong you see "Please upload only CSV, Excel, PDF, text, or image files.", "File size must be less than 50MB." or "Failed to upload one or more files. Please try again."

4.1.3 Choose what the answer can draw on

The pill next to the paperclip reads **Auto context**, **No context**, **Choose scope** or "N models". It starts with "Plan" when plan mode is on. Open it to set options.

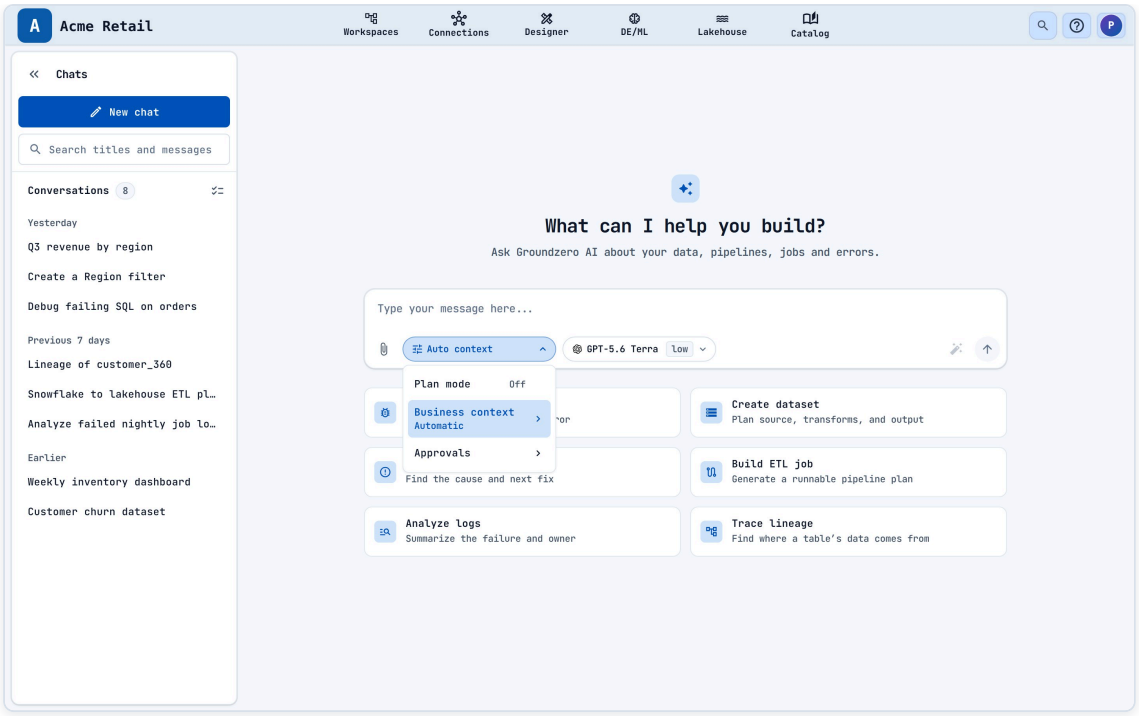


Figure 34. The chat options popover: Plan mode, Business context and Approvals.

Option	Effect
Plan mode	The assistant builds and tracks an execution plan, shown as a plan card
Business context	Whether the answer uses your published ontology (your business definitions): **Automatic context**, **Context off**, or **Choose models**
Approvals	Lists your standing approvals (see "Approve or deny an action"). It appears only when approval grants are enabled for your site

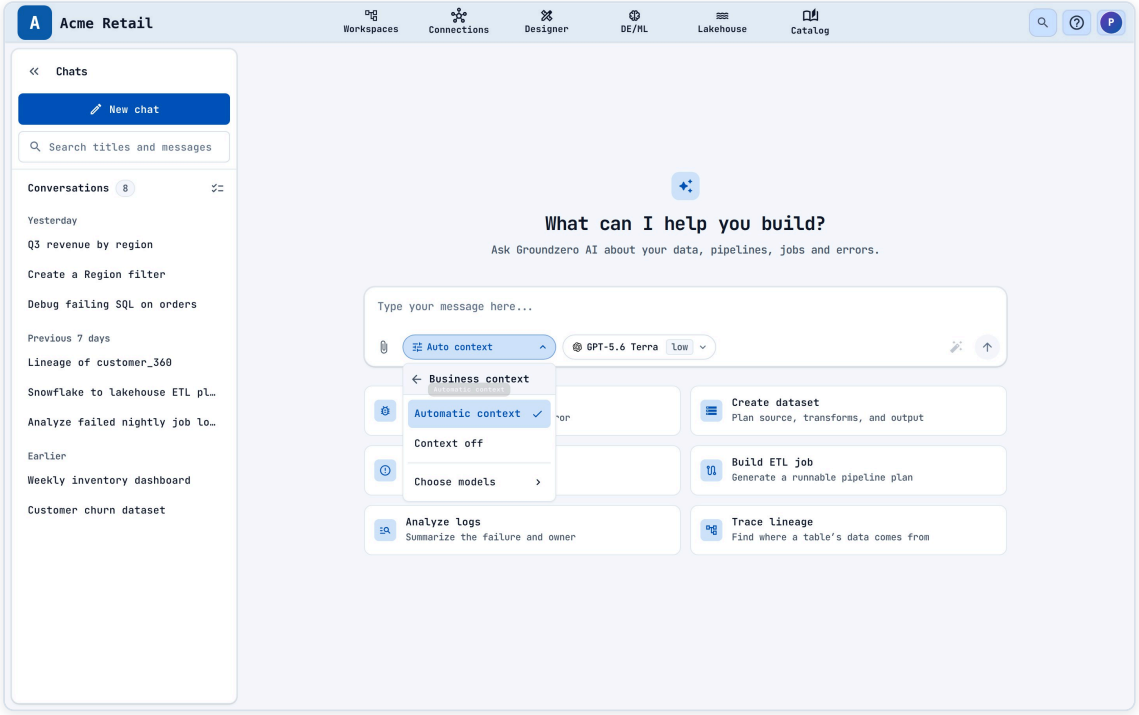


Figure 35. Business context options.

Automatic context works when there are 20 or fewer ontology models. With more, you must choose between 1 and 20 models. The pill reads **Choose scope** until you do.

4.1.4 Choose a model and how hard it thinks

The pill shows the model and a thinking level. Open it for the **Thinking effort** slider (Low, Medium, High, Extra high, Max). The levels offered depend on the model, and the default is Low. Choose the header row to pick a model.

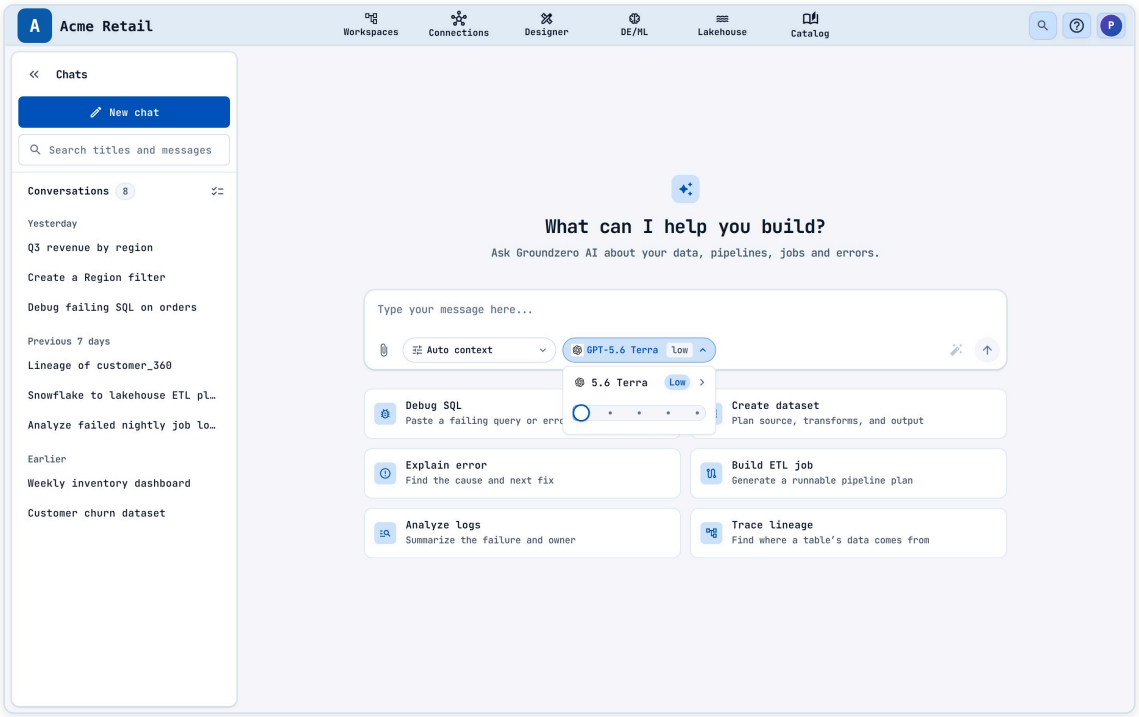


Figure 36. The model and thinking popover.

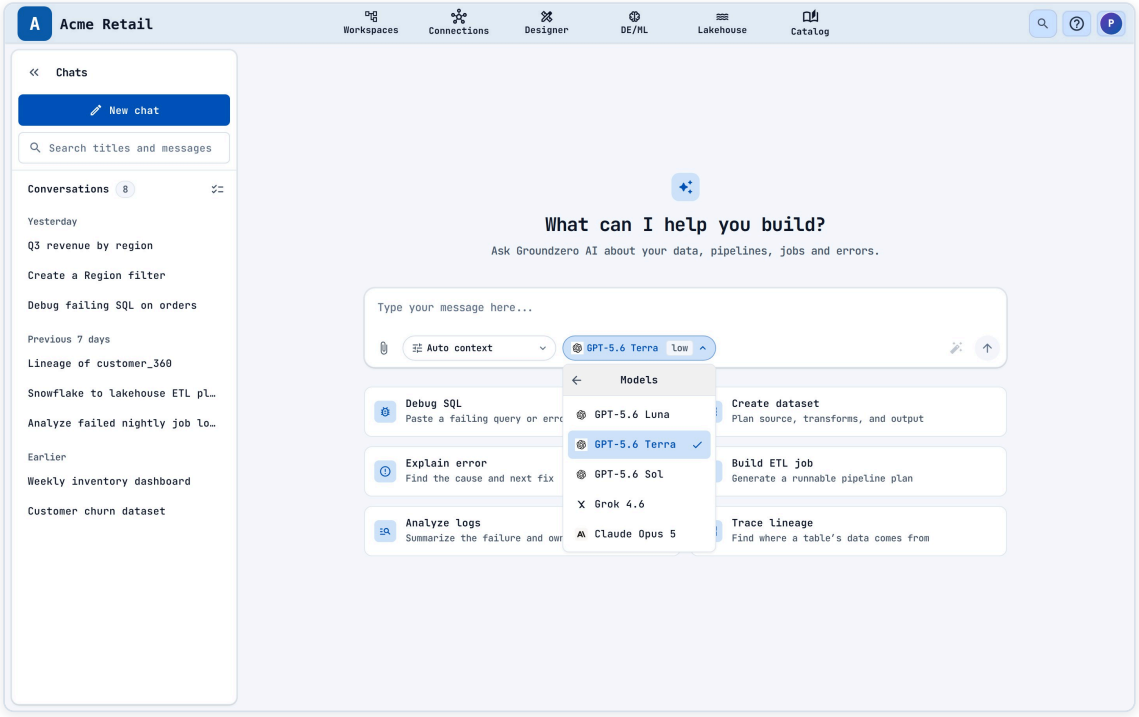


Figure 37. Choosing a Groundzero mode or one of your own LLM connections.

The list has Groundzero modes (for example GPT-5.6 Terra, Claude Sonnet 5, Claude Haiku 4.5) and a **Connections** section with one row per provider that has a saved LLM connection (Anthropic, AWS Bedrock, Azure OpenAI, Cohere, Custom LLM, Fireworks AI, Gemini, Hugging Face, OpenAI, Together AI, xAI). A connection with no model saved shows "Model not configured" and cannot be chosen. A conversation remembers its model.

If there is a problem with the model, sending is blocked and a notice explains it: "Choose a model connection.", "Edit <name> and add its model.", "<name> is no longer available. Choose another model." and similar. Fix the connection or pick another model.

4.2 Read an answer

Answers appear as they are written. Before text arrives you may see **Working...**. A **Thinking** panel shows the model's reasoning summary when there is one.

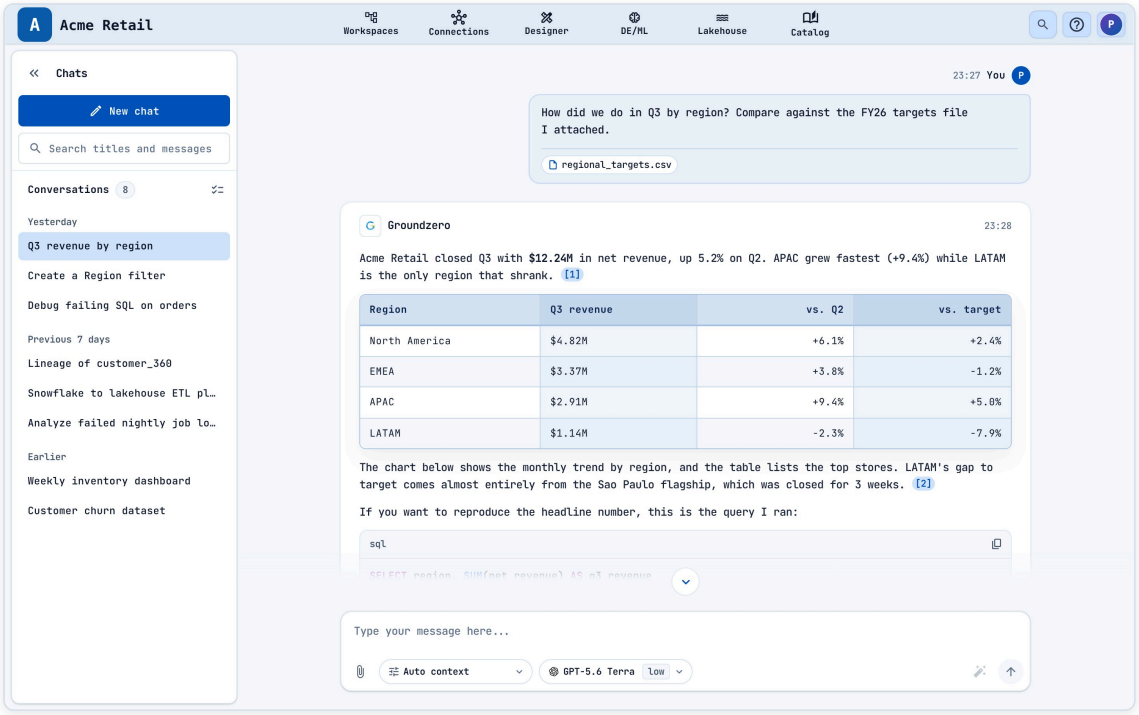


Figure 38. A saved answer with a table, citations, artifacts and a chart.

An answer can contain:

- Text with headings, lists, tables, code blocks (with a copy button) and diagrams.
- **Citations**: markers like `[1]` and a collapsed **Sources** panel.
- Cards for charts, tables (a one-row table shows as summary tiles), pivot tables and maps.
- An **Execution plan** card (plan mode).
- An **Artifacts** group: cards that open items the assistant created (dashboard, dataset, ETL job, visualization, filter, project, notebook, AI agent, ML model) or download generated files (PDF, DOCX, CSV, XLSX). Download links expire, so save files you need.
- A **Context warning** panel, for example "regional_targets.csv was cut to its first 10,000 rows."

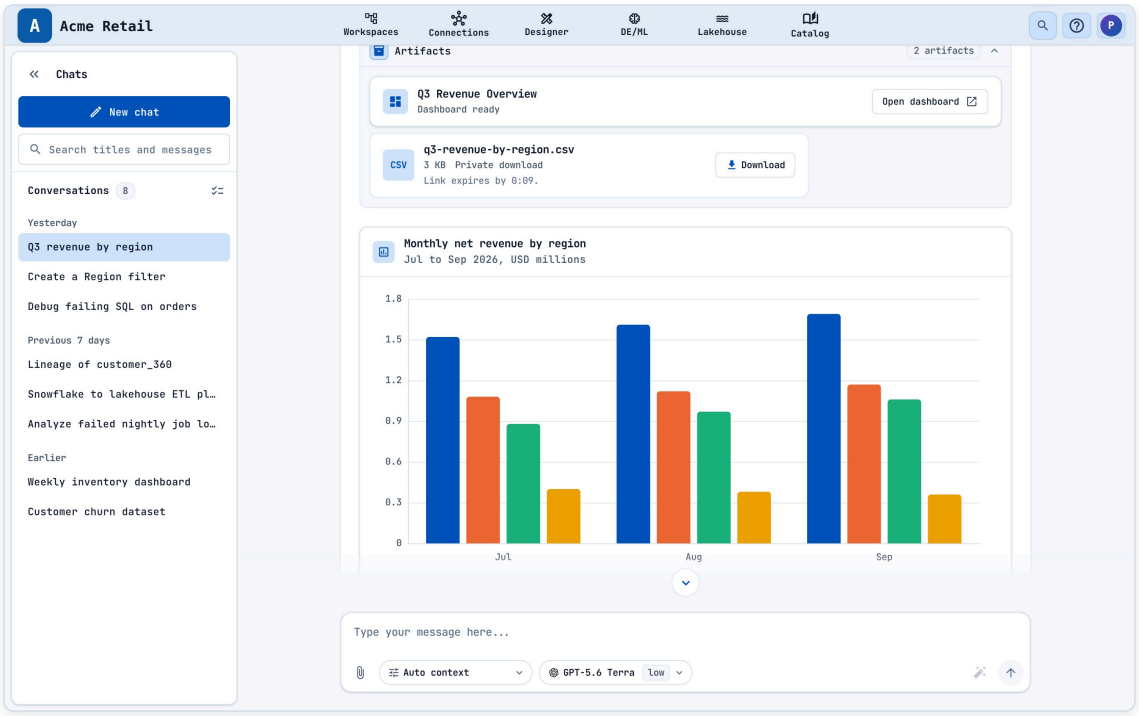


Figure 39. A chart card and the Artifacts group.

4.2.1 Check how an answer was produced

Under every answer are three panels you can open and a row of actions.

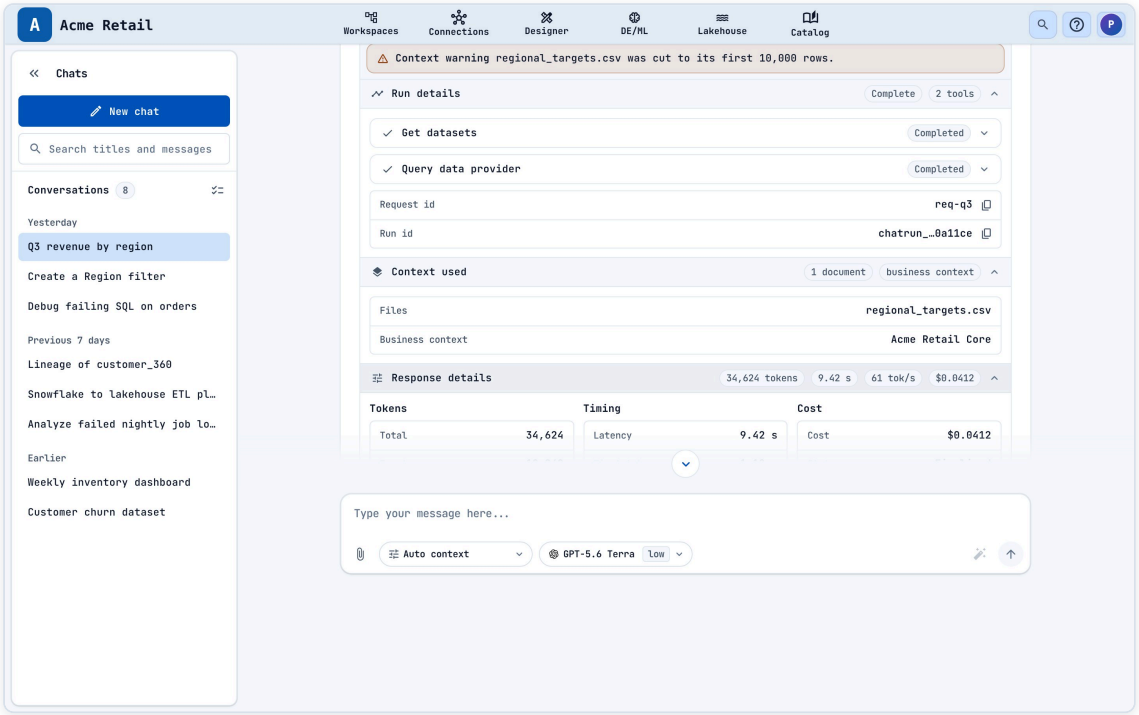


Figure 40. Run details, Context used and Response details expanded.

Panel	Shows
Run details	Each tool the assistant used with its status (Running, Completed, Failed, Cancelled, Incomplete) and result, plus the request and run IDs that support may ask for
Context used	Files, business context and context status. "This answer used only the conversation itself." when none

Panel	Shows
Response details	Tokens, timing (latency, first token, speed), cost, an **Agent cost** table, and your **30-day budget** meter

The action icons are **Copy response**, **Retry response** (starts a branch named "Retry: <chat name>"), **Continue response** (sends "Continue."), **Mark response helpful**, **Mark response not helpful** and **Open linked item**. On your own messages you also get **Edit and branch from prompt**, **Copy prompt** and **Copy link to prompt** (a link that fills the box with that text).

4.2.2 Tell Groundzero if an answer was good

Choose the thumbs up to save at once. Choose the thumbs down to open **What should improve?**, pick Incorrect, Not relevant, Incomplete, Unsafe, Poor code or Other, add optional **Details** (up to 2,000 characters) and choose **Send feedback**. The toast says **Feedback saved**. If it fails: "Could not save feedback. Please try again."

4.3 Stop, queue and steer a response

- **Stop response** stops the answer, keeps the text so far (marked "Generation stopped") and cancels the run.
- While a response is being written you can keep typing and choose **Queue message**. The **Queued messages** card lists them ("Next", "#2"...) and each has a remove button.
- If you reload the page or switch tabs, the run carries on. Chats with active runs show **Running**, **Background** or **Needs approval** badges in the sidebar.

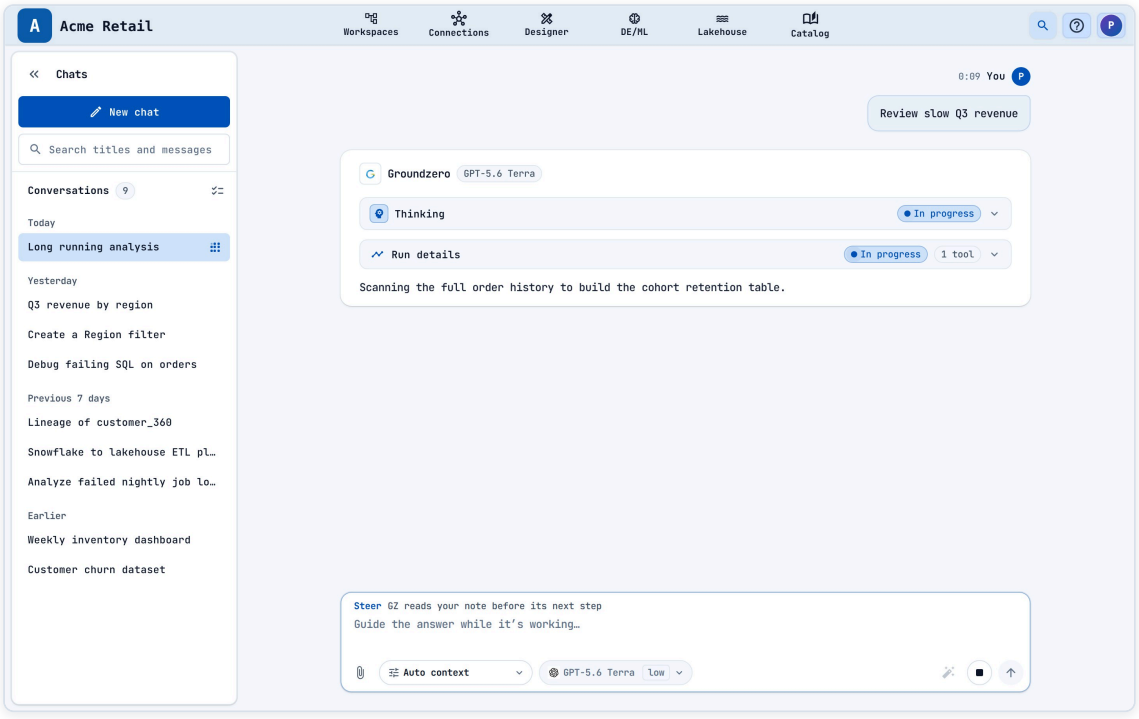


Figure 41. A run in progress: Thinking, Run details, and the Steer box.

4.3.1 Steer a run while it works

When live steering is turned on for your site, the box becomes a **Steer** box while a run works.

1. Type a note (1 to 2,000 characters).

2. Choose **Send steering note**.

The note appears in the conversation as **You steered** (Queued, then Applied). If the run finishes without reading it, Groundzero offers **Dismiss** or **Send as next message**. While an approval is waiting, steering is disabled ("Answer the pending action first").

4.4 Approve or deny an action

Actions that create, change or delete things pause the run. Calls to tools on your MCP connections always pause. Groundzero shows an **Action approval required** card.

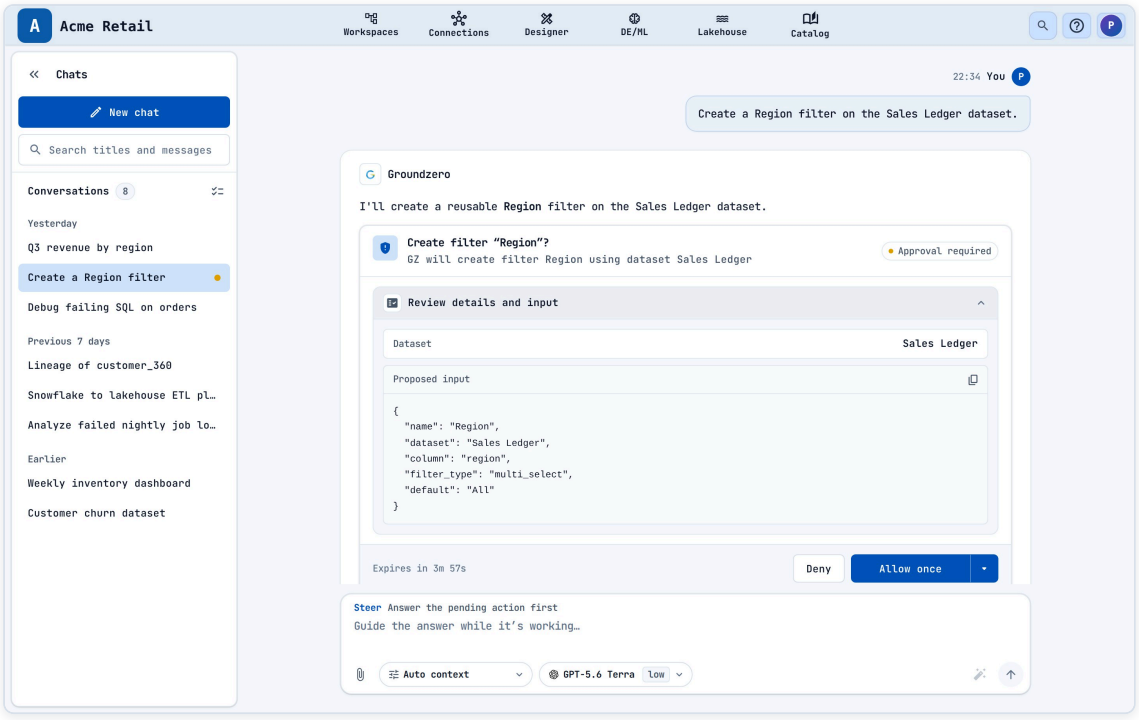


Figure 42. An approval card for creating a filter, with Deny and Allow once.

1. Read the card. It shows the action, the target and a countdown ("Expires in 3m 53s"). Open **Review details and input** to see exactly what will be sent, with secrets hidden.
2. Choose a button:

Button	Effect
Deny	The tool does not run
Allow once	Runs this call only
Caret > **Allow for this chat**	Runs this call and later calls of the same tool in this chat
Caret > **Always allow**	Runs it from now on in every chat. Not offered for tools that delete or destroy things

You cannot send another message until you decide. The box says "Waiting for your decision". If you reload, the card comes back. An expired card says "This action preview expired. Ask GZ to run it again."

> **WARNING**
> **Always allow** lets that tool run without asking in every chat. Use **Allow for this chat** unless you are sure.

4.4.1 Review or cancel standing approvals

Calls covered by a standing approval run without a card. The conversation shows "Auto-approved create filter - allowed for this chat" (or "always"). To review them, open **Chat options > Approvals**. The sections **This chat** and **Always** list each tool (MCP tools show their connection). Choose **Revoke** to remove one. With none, you see "No tools run without asking. Allow one for this chat or always when GZ asks."

4.5 Find an earlier conversation

The **Chats** sidebar lists your conversations grouped into Today, Yesterday, Previous 7 days and Earlier. A new chat is called "New chat" until Groundzero gives it a title.

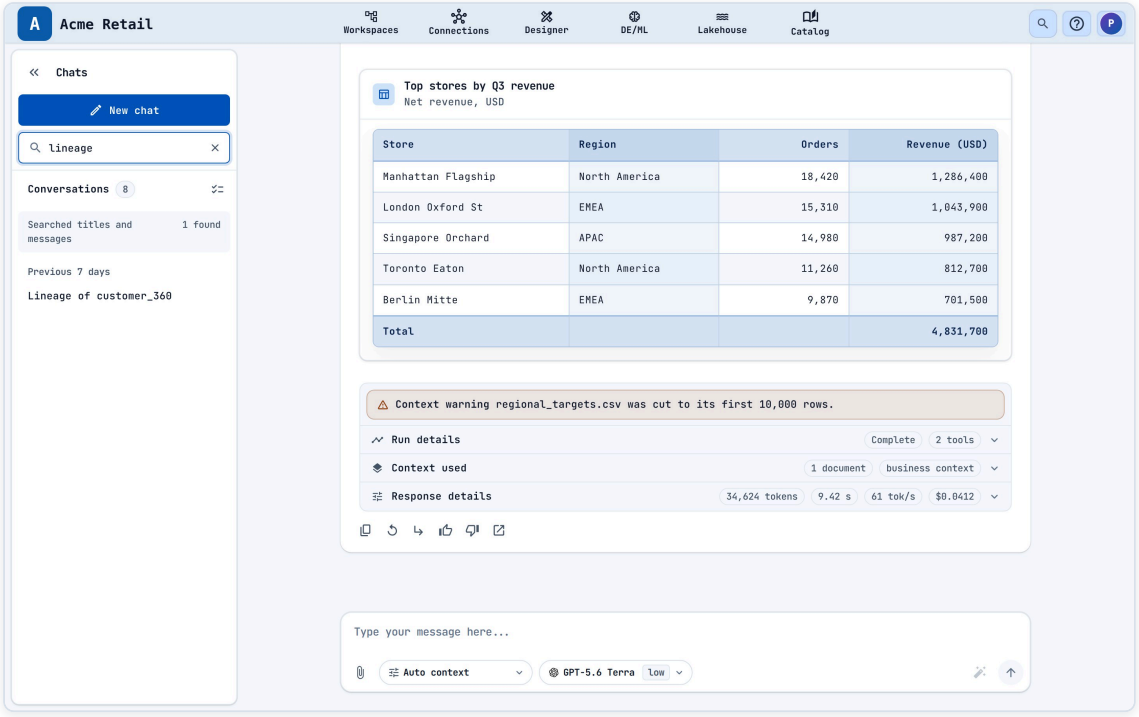


Figure 43. Searching titles and messages.

To do this	Do this
Start a chat	Choose **New chat**
Search	Type in **Search titles and messages**. Titles filter at once, then Groundzero searches message text. The result line reads "Searched titles and messages" and "N found". No match: "No chats found"
Rename	Hover a row, choose **Rename**, type and press Enter. Escape cancels
Delete one	Hover, choose **Delete**. The dialog says "Messages, files and run history go with it. This can't be undone." Toast: "1 conversation removed."
Delete several	Choose **Select conversations** (the checklist icon), tick rows, choose **Delete N selected conversations**

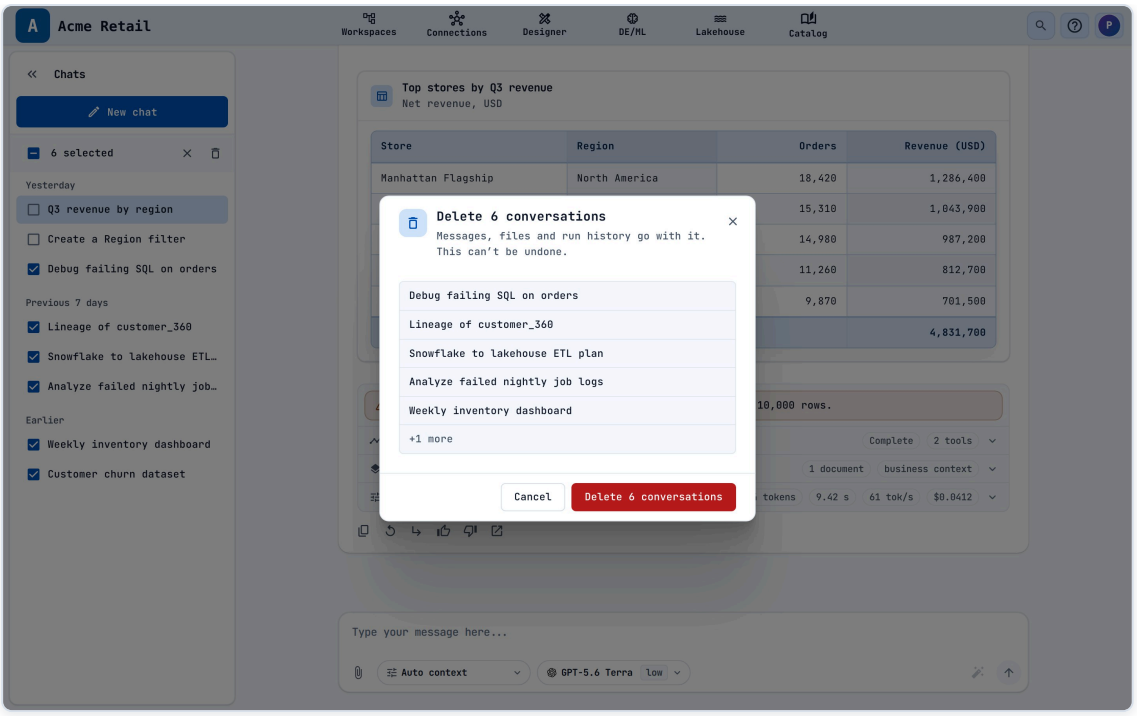


Figure 44. Deleting several conversations at once.

> **WARNING**
> Deleting a conversation also deletes its files and run history. It cannot be undone.

You can delete up to 100 conversations at a time. Select-all picks the first 100 and further rows are disabled ("You can select up to 100 conversations."). Running chats are skipped. If some fail, the toast says "X removed. Y failed; select Delete to retry."

4.6 Search everything with Ctrl+K

Press **Ctrl+K** (Cmd+K on macOS) anywhere in Groundzero, or choose **Search** in the header, to search pages and resources.

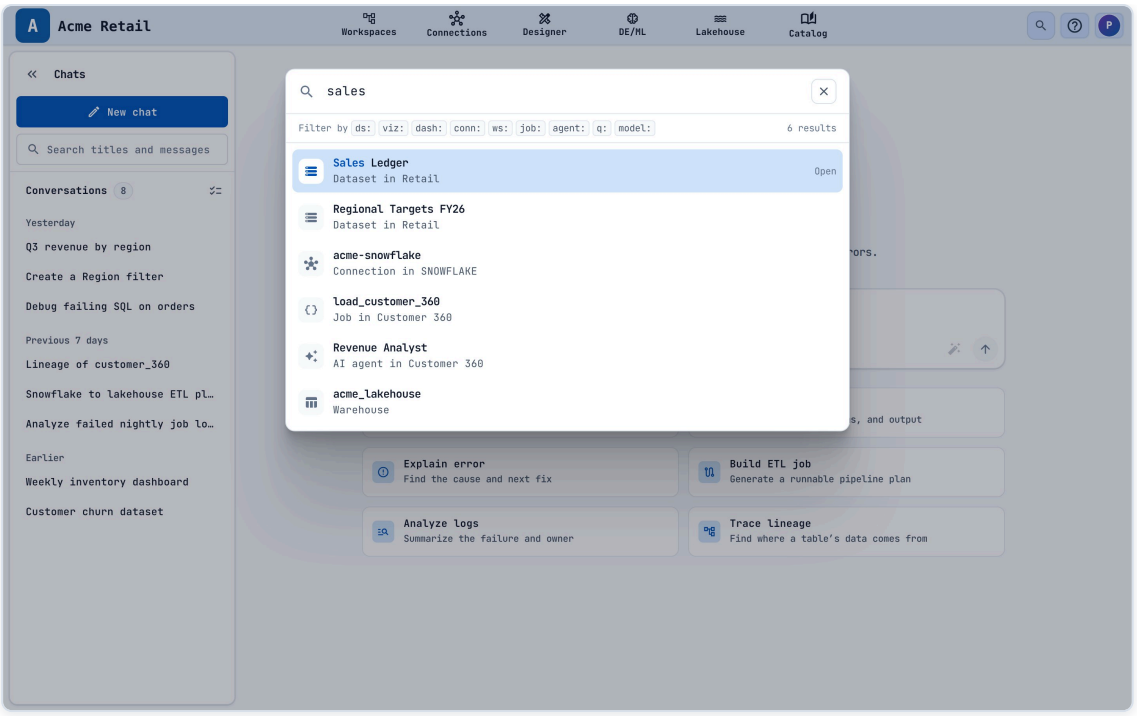


Figure 45. Global search results for "sales".

1. Type to search datasets, filters, visualizations, dashboards, connections, workspaces, projects, jobs, AI agents, warehouses, saved queries and ontology models.
2. To narrow the search, start with a prefix: `ds:`, `viz:`, `dash:`, `conn:`, `ws:`, `job:`, `agent:`, `q:` or `model:` (for example `dash: sales`).
3. Use the arrow keys and **Enter** to open a result.

Results show at most 8 rows and respect your permissions. With an empty box, your recently opened items appear.

4.7 If something goes wrong

- **You don't have access to chat.** Cause: you lack the AI Chat permission. Fix: ask an administrator to grant it.
- **Usage limit reached ("You have reached your current chat usage limit. Contact your administrator to request additional access.").** Cause: your budget for the last 30 days is used. Fix: ask an administrator for more budget.
- **Temporarily unavailable.** Cause: the model could not answer right now. Fix: choose **Try again** or pick another model.
- **Approval pending.** Cause: an action is waiting for you. Fix: choose Deny or Allow on the card, then send.
- **Model credentials rejected, Model billing problem or Model not available.** Cause: a problem with your own LLM connection. Fix: edit the connection.
- **"Too many requests. Please wait and try again."** Cause: you sent requests too quickly. Fix: wait a moment and retry.
- **"Couldn't load your conversations. Refresh the page to try again."** Fix: refresh the page.
- **"Search is temporarily unavailable."** Cause: message search could not run. Fix: title filtering still works. Try again shortly.

4.8 Limits

Limit	Value
Message length	200,000 characters
Enhance prompt length	8,000 characters
Attachment size	50 MB; 100 files per message
Deleting conversations at once	100
Steering note	2,000 characters
Business context scope	20 ontology models
Feedback details	2,000 characters

4.9 Related chapters

- Chapter 5 for AI agents that other systems can call.
- Chapter 6 for chat usage and cost reports.
- Chapter 7 for connecting external tools (MCP) to Groundzero.
- The Administration guide, chapters 4 and 7, for the AI Chat permission and budgets.

5 AI agents

An **AI agent** is a saved assistant with one fixed job. You give it instructions, a model and a set of tools. After a reviewer approves a version, your own systems can call the agent over the internet using a client token. Agents live inside a DE/ML project and use that project's data tools.

5.1 Decide between chat and an agent

Use chat when	Use an agent when
You are exploring interactively	The task repeats and must behave the same way each time
You want the assistant to decide as it goes	You need a narrow tool set, approvals and a recorded history
Only you use it	Other systems or customers call it

Before you start:

- Your administrator controls access with permissions on **Agent**: **List**, **Create**, **Read**, **Modify**, **Delete**, **Invoke** and **Approve** (Administration guide, chapter 4). A refused action shows `You do not have "All" or "<Action>" permission on GZ Agent.`
- Creating an agent needs an AWS S3 connection to store its conversation history. If you want a model other than the built-in Groundzero models, you also need an LLM connection.

5.2 Find your agents

1. Open **DE/ML** in the top navigation and choose a project.
2. Choose the **Agents** tab (next to Jobs, Notebooks and Models).

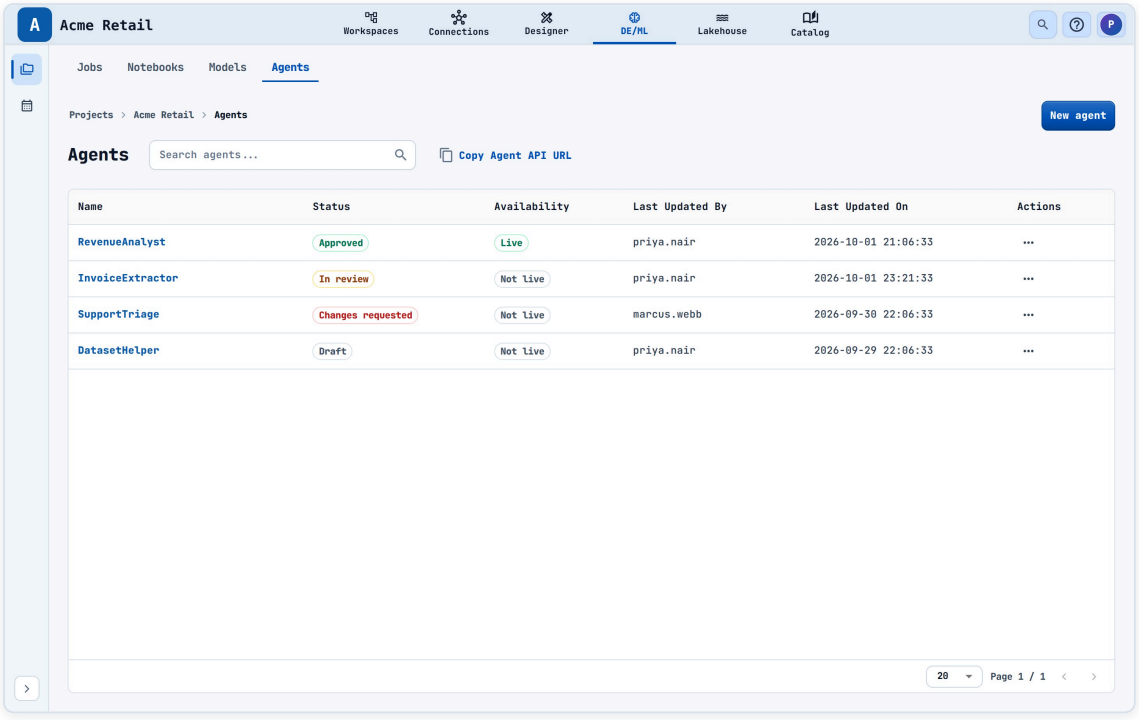


Figure 46. The Agents tab of a project.

Agents belong to the project they were created in. The list has a search box (**Search agents...**, by name or model), **New agent** and **Copy Agent API URL**. The columns are:

Column	Meaning
Name	A link to the edit page
Status	Draft, In review, Approved or Changes requested
Availability	**Live** when an approved version is being served to callers, otherwise **Not live**
Last Updated By / On	Who saved last, in your local time
Actions	Row menu

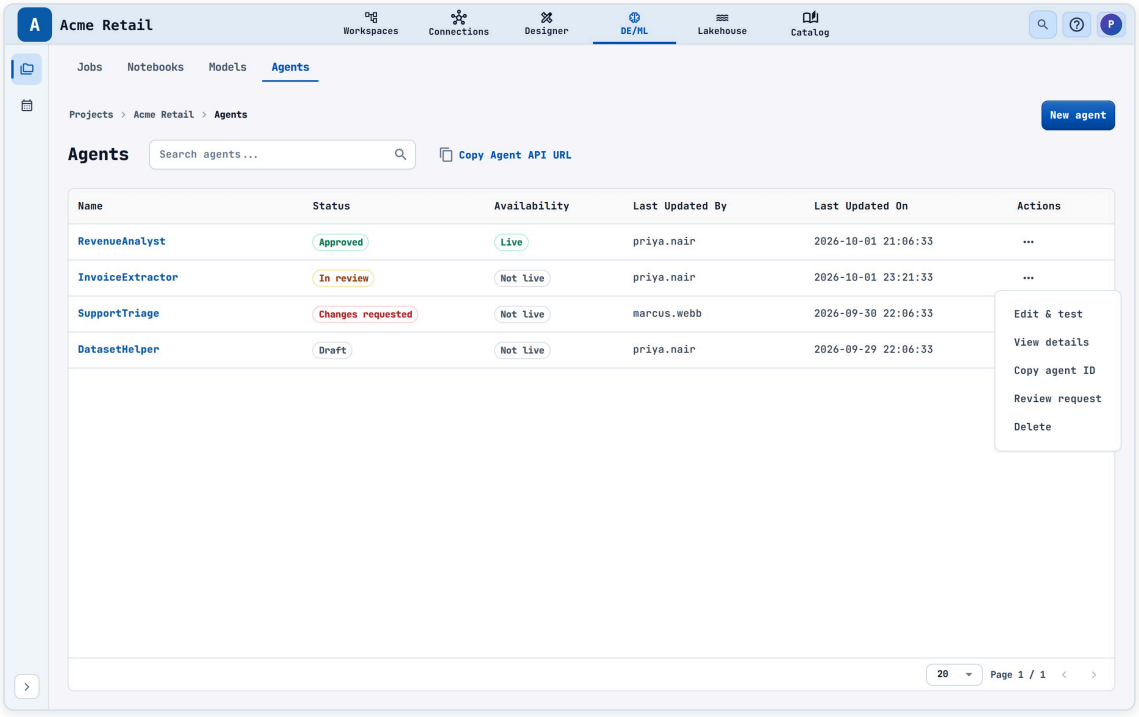


Figure 47. The row menu.

The row menu has **Edit & test**, **View details**, **Copy agent ID**, **Review request** (only for an In review agent, and only if you may approve) and **Delete**. Delete asks "Delete <name>? This action cannot be undone." and shows the toast "Agent deleted".

5.3 Build an agent

Outcome: a saved draft agent you can test.

Choose **New agent**. The builder has six steps, and a guide panel on the right explains each one. Nothing is saved until the last step.

5.3.1 Step 1: Identity

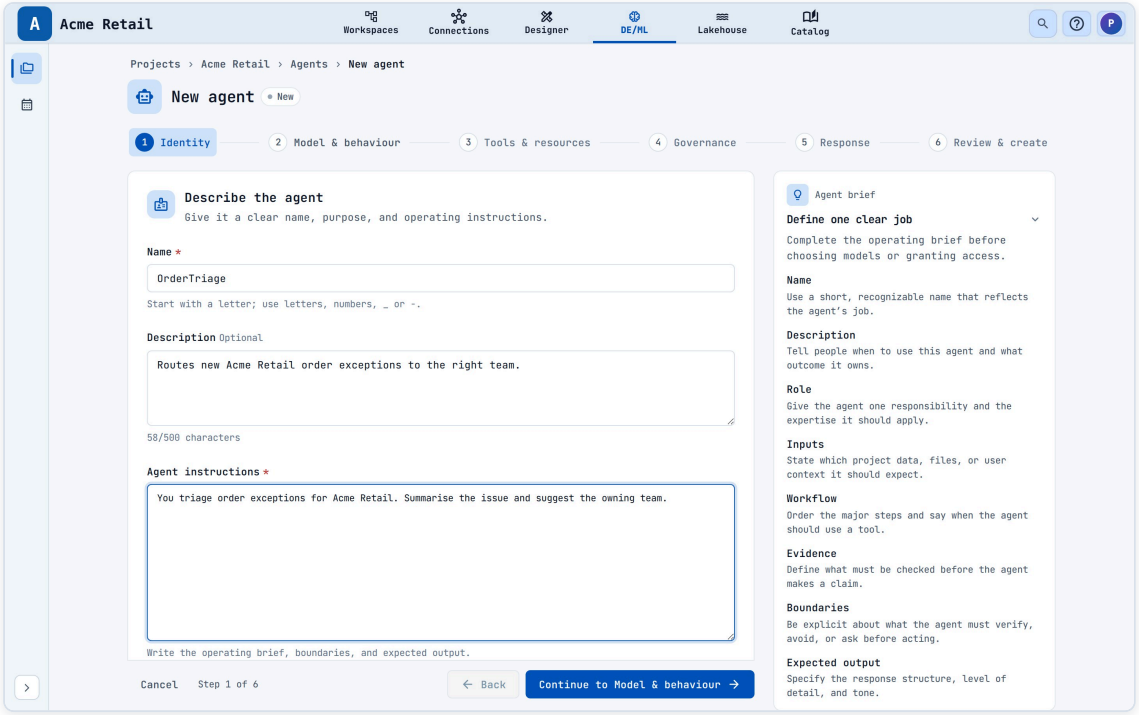


Figure 48. Step 1 of the builder.

Field	Rules
Name (required)	Starts with a letter, then letters, numbers, `_` and `-`, up to 100 characters. Names are unique across the site. A clash shows "An agent named '<n>' already exists"
Description	Optional, up to 500 characters
Agent instructions (required)	Up to 100,000 characters. Describe the job, the inputs, the workflow, the evidence needed, the boundaries and the expected output

> **TIP**
> Write instructions as you would brief a new colleague: what the job is, what they get, what good looks like and what they must never do.

5.3.2 Step 2: Model and behaviour

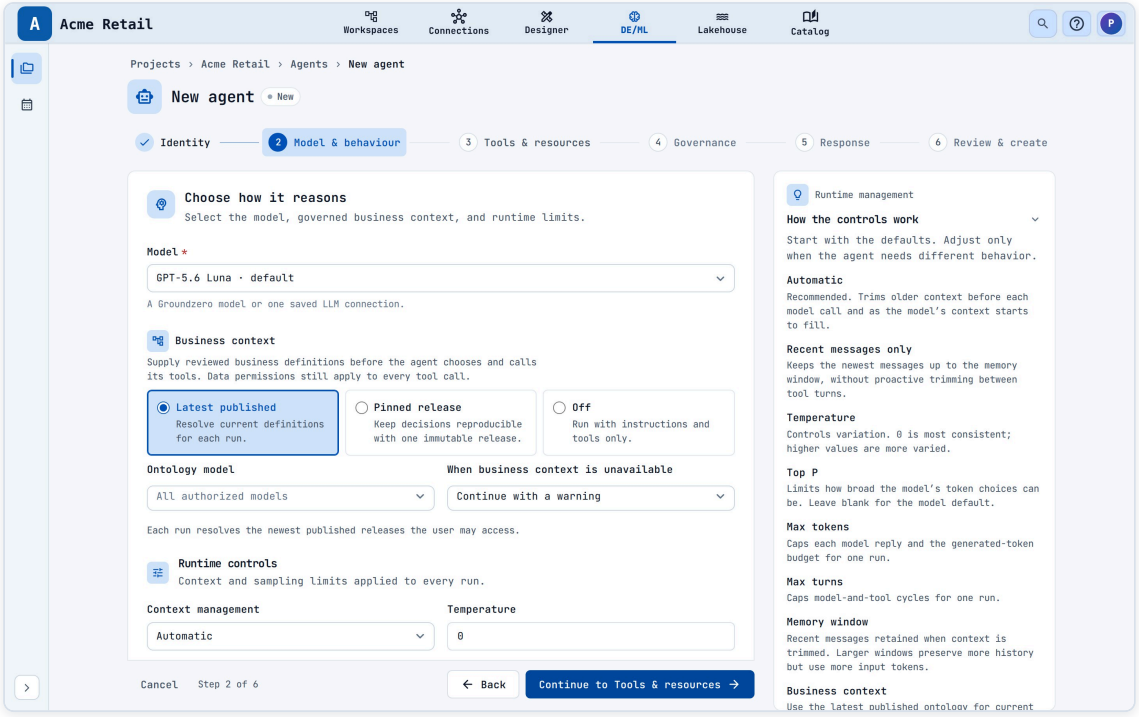


Figure 49. Step 2: model, business context and runtime controls.

- **Model**: Groundzero models (the default is marked) and one group per saved LLM connection. If a saved model is gone it shows "unavailable".
- **Business context**: **Latest published** (uses the current definitions on every run), **Pinned release** (one fixed release, with an **Ontology model** choice) or **Off**. For latest, **When business context is unavailable** is "Continue with a warning" or "Stop the run".
- **Runtime controls**:

Control	Range and default
Context management	Automatic (recommended) or Recent messages only
Temperature	0 to 1, default 0
Top P	0 to 1, blank for the model default
Max tokens	1 to 100,000 and not above the model's limit, default 4096
Max turns	1 to 300, default 300
Memory window	4 to 200 messages, default 40

Temperature and Top P are disabled for models that do not support them.

5.3.3 Step 3: Tools and resources

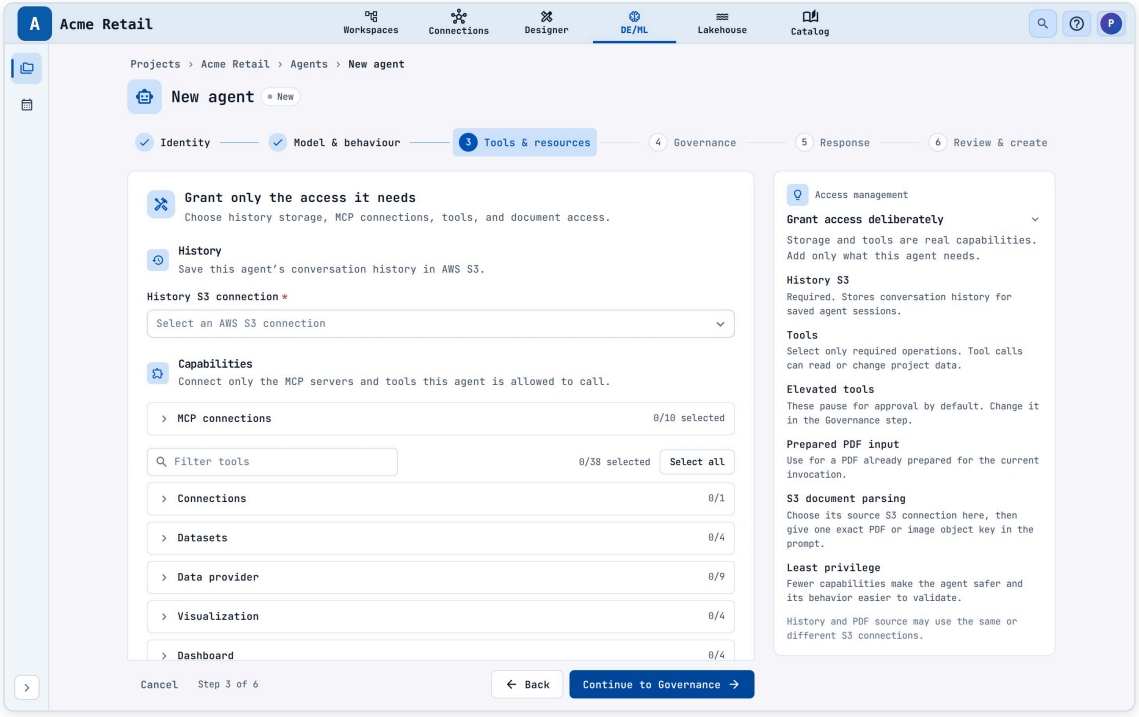


Figure 50. Step 3: history connection, MCP connections and the tool catalog.

- **History S3 connection** (required): the AWS S3 connection that stores conversation history. If you have none, the step says "Create an AWS S3 connection before saving this agent." Leaving it empty shows "Choose an AWS S3 history connection."
- **MCP connections**: up to 10 of your saved MCP connections.
- **Tool catalog**: search with **Filter tools** and pick up to 50 of the 38 built-in tools, grouped as Connections, Datasets, Data provider, Visualization, Dashboard, Filter, ETL and Documents. Tools that change things carry an amber **Elevated** chip.

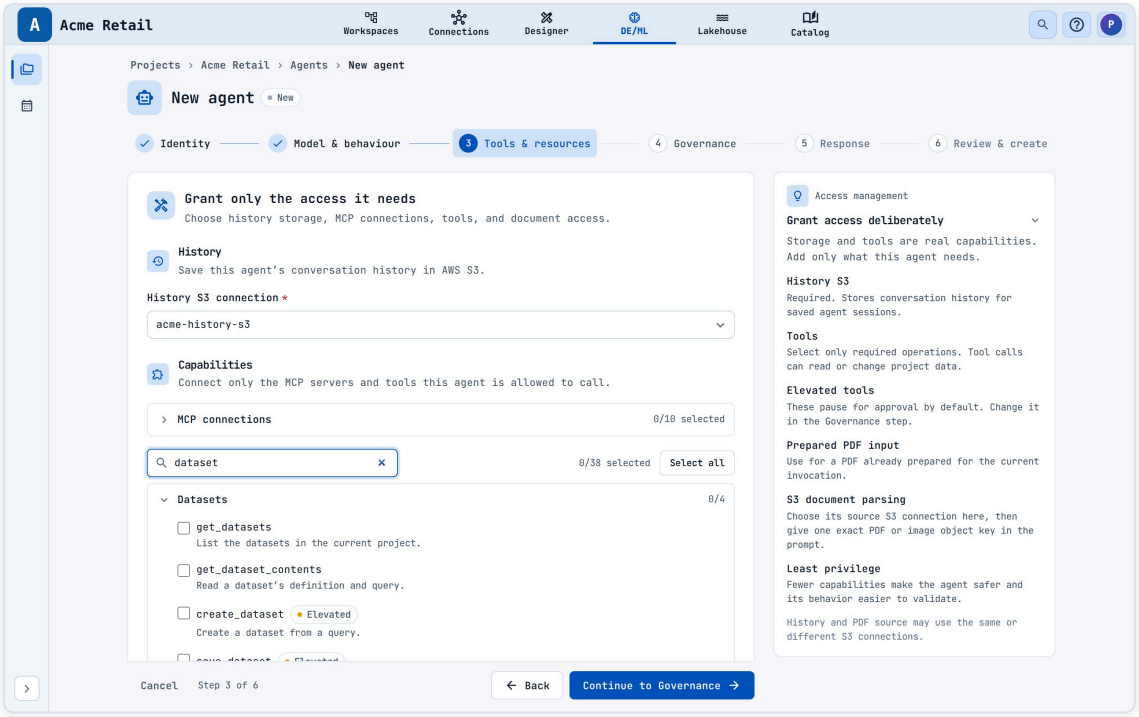


Figure 51. Filtering the catalog.

Elevated tools are the ones that create or save datasets, query a data provider, create projects, create or save jobs and submit ETL jobs. When you select one, **Require approval** is switched on by default.

The tool **S3 Document Reader** needs a **PDF source S3 connection** (the bucket). In each message you give the exact object key of the PDF or image. Folders are not read. **Prepared PDF Reader** never pauses for approval.

5.3.4 Step 4: Governance

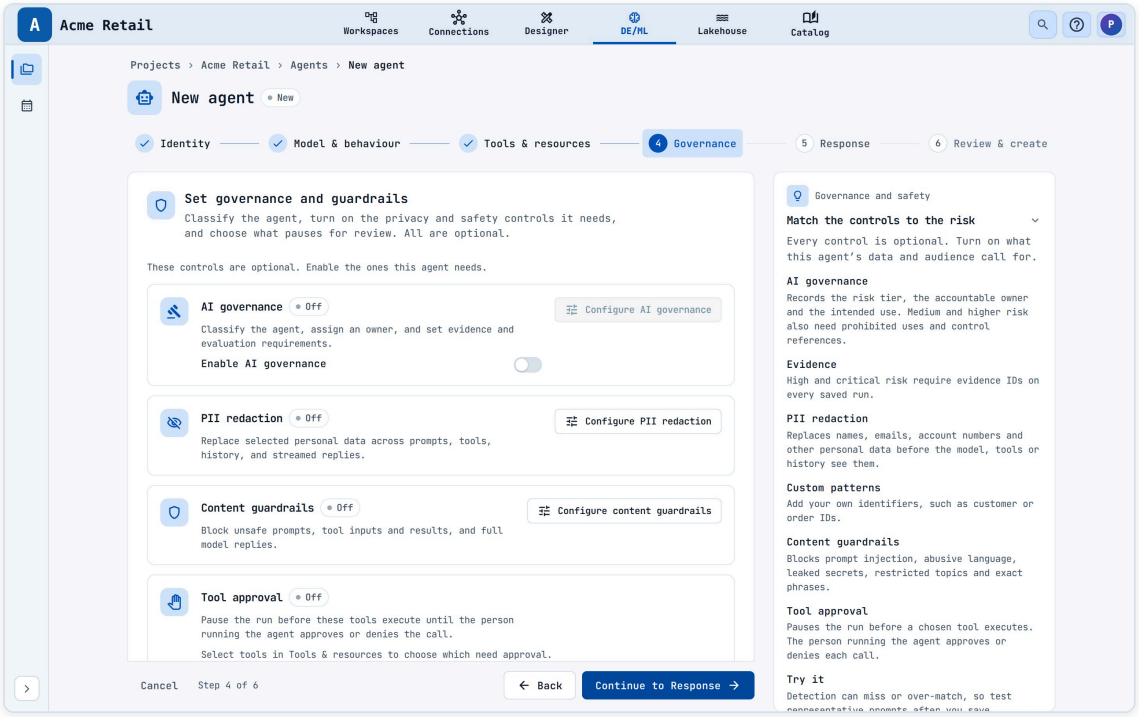


Figure 52. Step 4: governance and guardrails.

All four cards are optional.

Card	What it does
AI governance	Switch **Enable AI governance**, then **Configure AI governance**: Risk tier (low to critical), Accountable owner (a site user), Data classification (Public, Internal, Confidential, Restricted), Evidence retention (1 to 3,650 days), Intended use, Affected populations, Known limitations, Jurisdictions, Regulatory scope, and advanced controls (prohibited uses, denied tools, blocked terms, maximum tool calls, required evaluations)
PII redaction	Replaces detected personal data with typed placeholders before the model, tools, history or response can see it. 23 detectors (for example email, phone, credit card, IBAN, several national IDs) and up to 10 custom patterns
Content guardrails	Blocks Prompt injection, Toxic language, and Credentials and secrets, plus up to 20 restricted topics and 100 blocked phrases. A blocked run returns "This content was blocked by the agent guardrails."
Tool approval	One **Require approval** switch per selected tool. A paused run waits until the person running the agent approves or denies

5.3.5 Step 5: Response

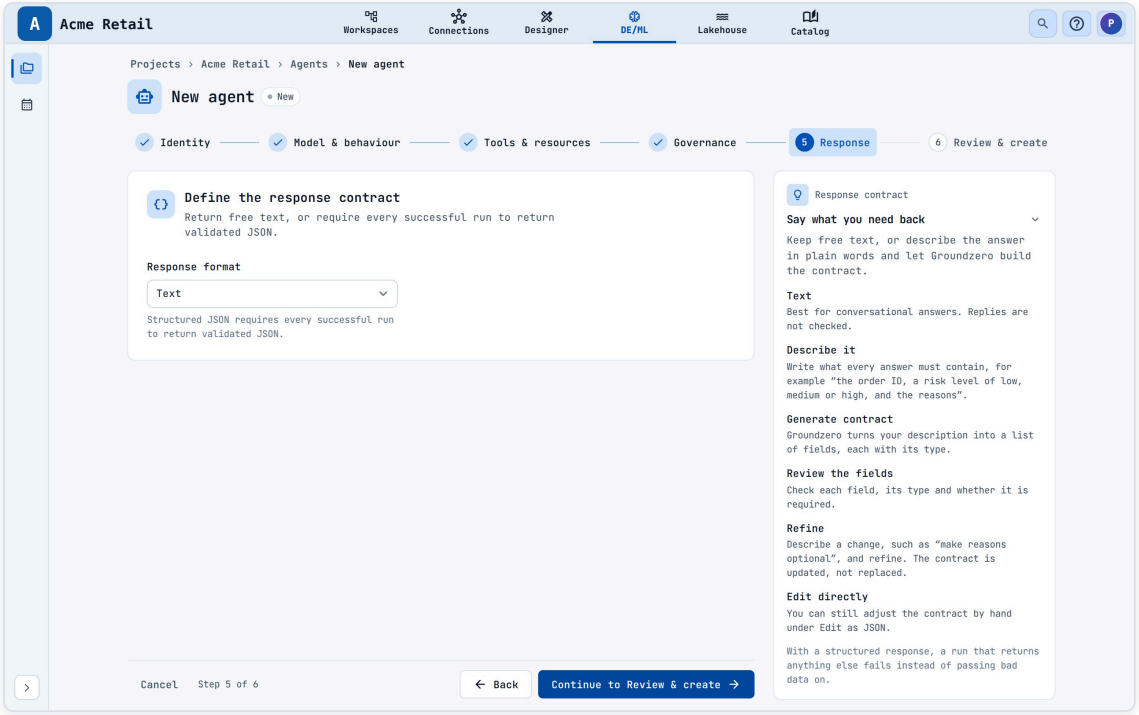


Figure 53. Step 5: response contract.

Choose a **Response format**: **Text**, or **Structured JSON**, which makes every successful run return checked JSON. For structured output you can describe the result in words (**What should the agent return?**) and choose **Generate contract**, then review the fields or edit the JSON. If you ask for a change, the button becomes **Refine**. A failed request keeps your current contract.

5.3.6 Step 6: Review and create

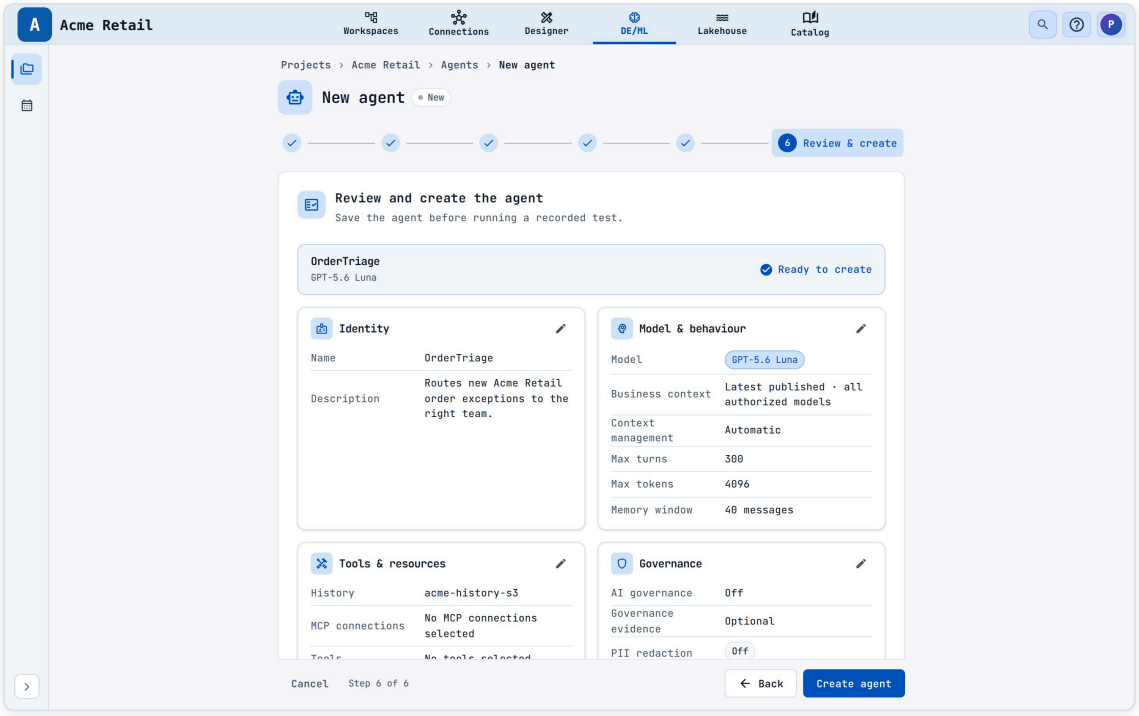


Figure 54. Step 6: review.

Review the summary cards (the pencil icon jumps back to a step) and choose **Create agent**. The toast says "Agent created as a draft". The agent's edit page opens.

If you leave with unsaved changes, Groundzero asks "Discard unsaved changes?".

5.4 Test an agent

Outcome: you see how the agent behaves before anyone else can use it.

The **Test agent** panel runs the saved agent with real tools and billable tokens. It is locked until the agent exists ("Create this agent before testing.") and again after you edit ("Save your changes before testing.").

1. In **Test prompt** type a request. Press Cmd or Ctrl+Enter to run.
2. Leave **Stream response** ticked to see text as it arrives.
3. Choose **Run test**. Use **Stop** to cancel a streaming run.

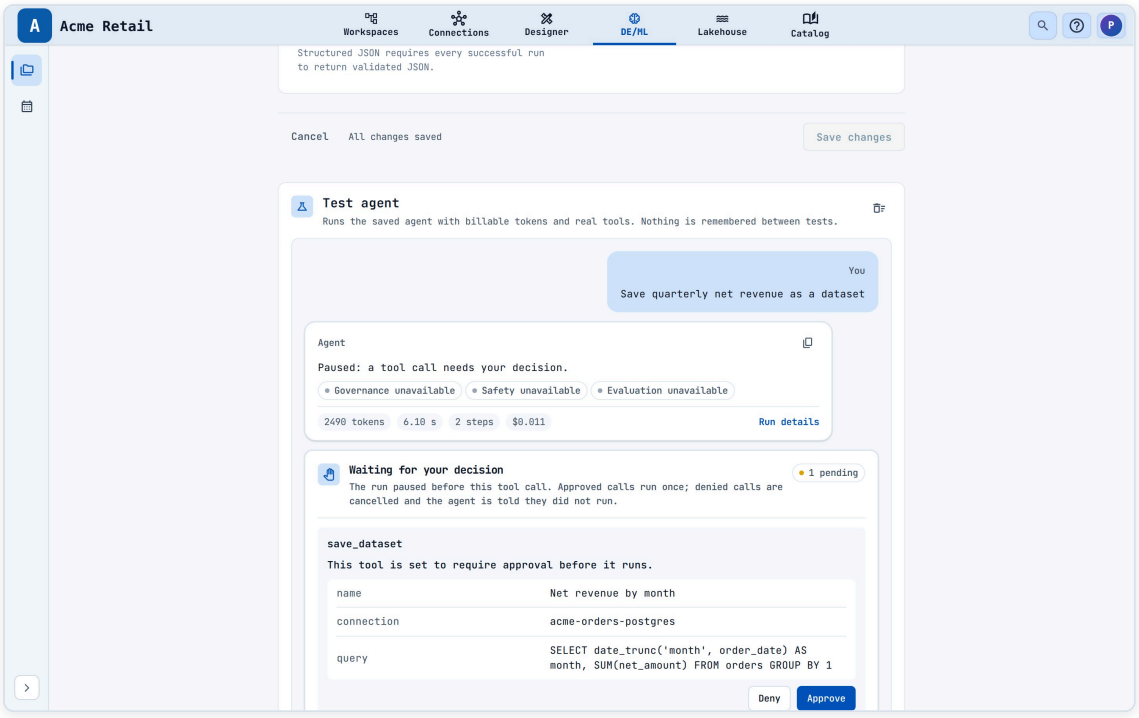


Figure 55. A test run paused for approval: Deny or Approve.

If the agent calls a tool that needs approval, the run pauses with **Waiting for your decision** and a preview of the input. Choose **Approve** or **Deny**. Approved calls run once. Denied calls are cancelled and the agent is told they did not run.

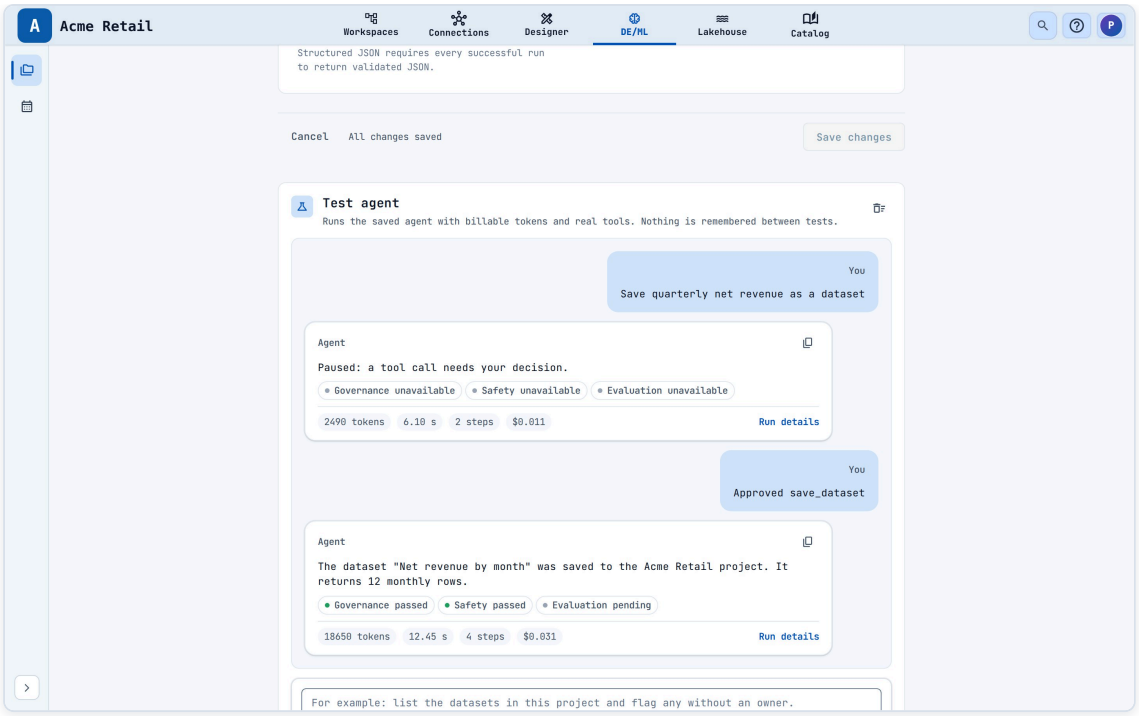


Figure 56. The completed reply with Run details.

The reply shows chips for Governance, Safety and Evaluation, and **Run details**: tokens, cache reads and writes, duration, steps, tool calls and estimated cost. A test does not remember earlier tests.

5.5 Get a version approved and live

Saving creates a new version, and **every new version needs approval before callers get it**.

1. Open the agent and choose **Submit for approval**. The button is disabled while you have unsaved edits and reads **Awaiting approval** once submitted.
2. A reviewer with the **Approve** permission opens the project's Agents list and chooses **Review request** on the row.

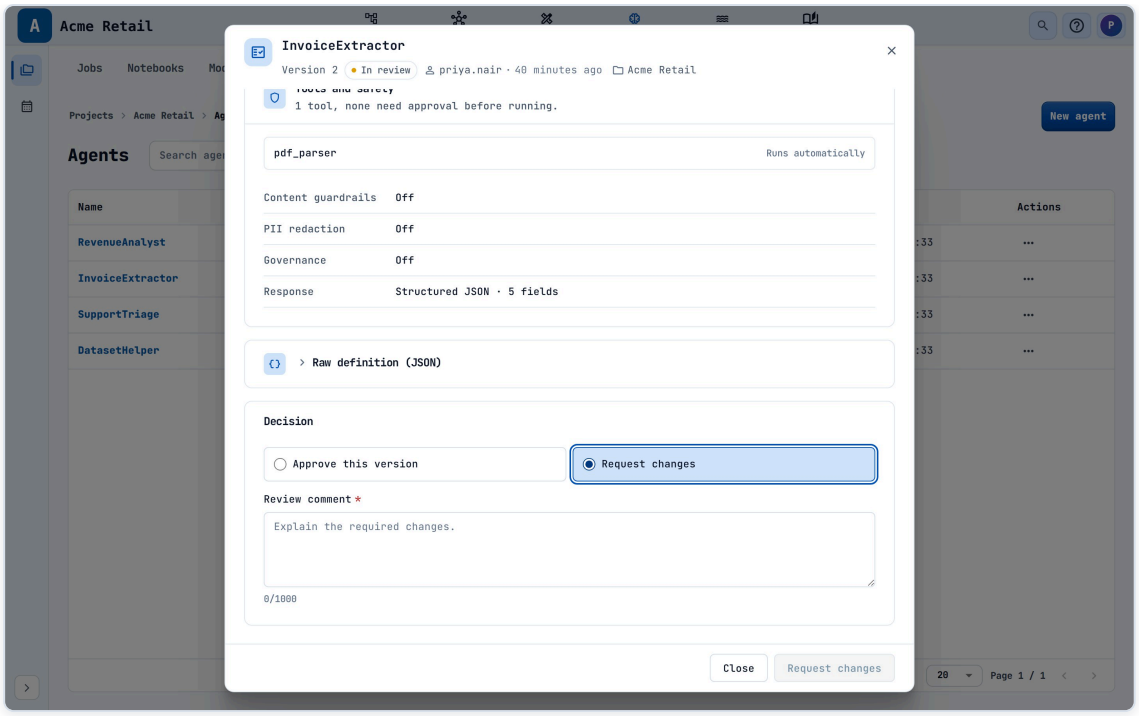


Figure 57. The review dialog.

3. The reviewer reads the version (model, tools, guardrails, governance and instructions) and chooses **Approve this version** or **Request changes**. A comment is required when requesting changes (up to 1,000 characters).

What happens next: an approved version makes the agent **Live**. Later saves create a new draft, and callers keep getting the last approved version until the new one is approved. When changes are requested, the builder shows the reviewer's note as **Admin feedback**. Fix the issues and choose **Resubmit for approval**.

If approval is blocked, a red **Approval is blocked** notice explains why (for example, missing governance evidence) and offers a **Review governance** button.

5.6 Look at an agent's details

Choose **View details** on the agent row. The page has four tabs.

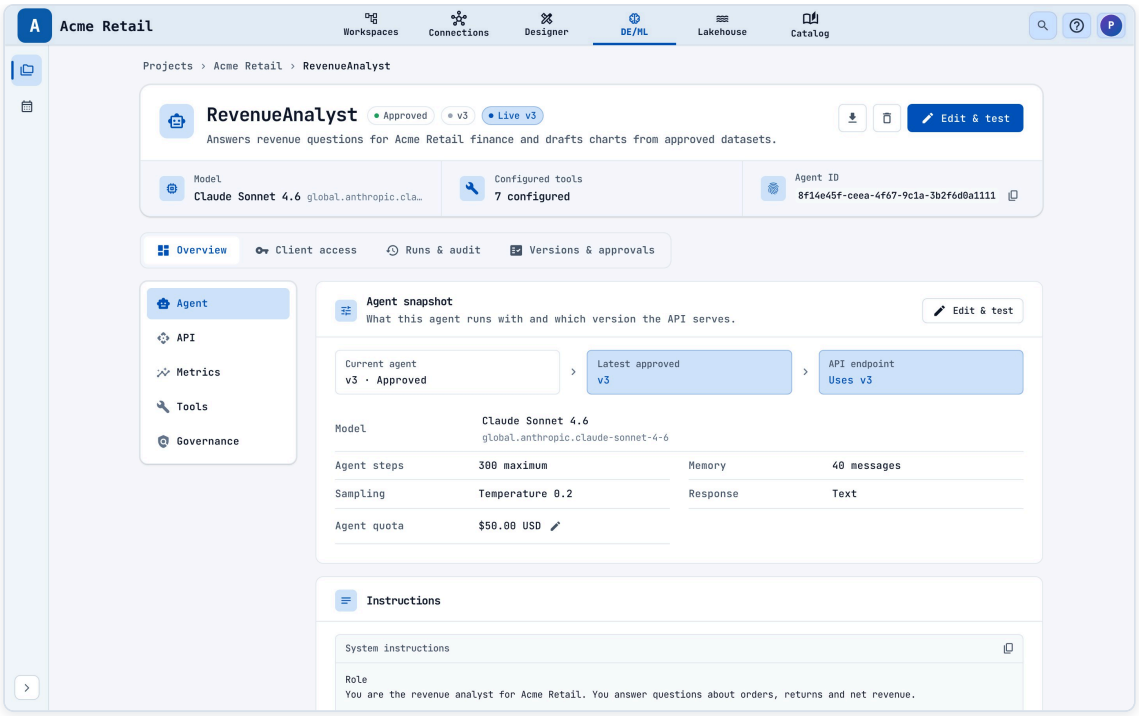


Figure 58. The Overview tab with lifecycle strip and settings.

5.6.1 Overview

A left rail has **Agent**, **API**, **Metrics**, **Tools** and **Governance**.

- **Agent**: a snapshot (current, latest approved and served versions), the instructions, the configured tools with **Require approval** switches, and the **Agent quota**. If you change a switch and save, that creates a new version that needs approval.
- **Agent quota**: a spending limit in USD. When spend reaches it, runs are refused with "This agent has spent its USD quota". Set it with **Set agent quota** or change it with **Edit agent quota** (**Save quota**, **Remove quota**).
- **API**: how to call the approved agent (see "Let another system call your agent").
- **Metrics**: Runs, Completion rate, Known estimated cost, P95 run duration, Tokens and Unique users, with panels for outcome, cost, customer and model. Ranges are 7 days, 30 days, 90 days, All time or Custom. Percentiles say "Not enough data" under 20 runs.
- **Tools**: how often each tool was used.
- **Governance**: lifecycle (Suspend, Reactivate, Retire permanently), release readiness and fairness attestation.

5.7 Let another system call your agent

Outcome: a customer or system can use the approved agent, and you can see its usage and cost separately.

First check that the agent is **Live**. Callers get only the latest approved version. Before approval the API tab says "This endpoint becomes usable after an administrator approves a version."

1. Open the agent's **Overview > API** page. It shows Standard, streaming (SSE) and Continue modes, each with cURL and Python examples to copy.
2. Open the **Client access** tab and create a token for each customer or system.

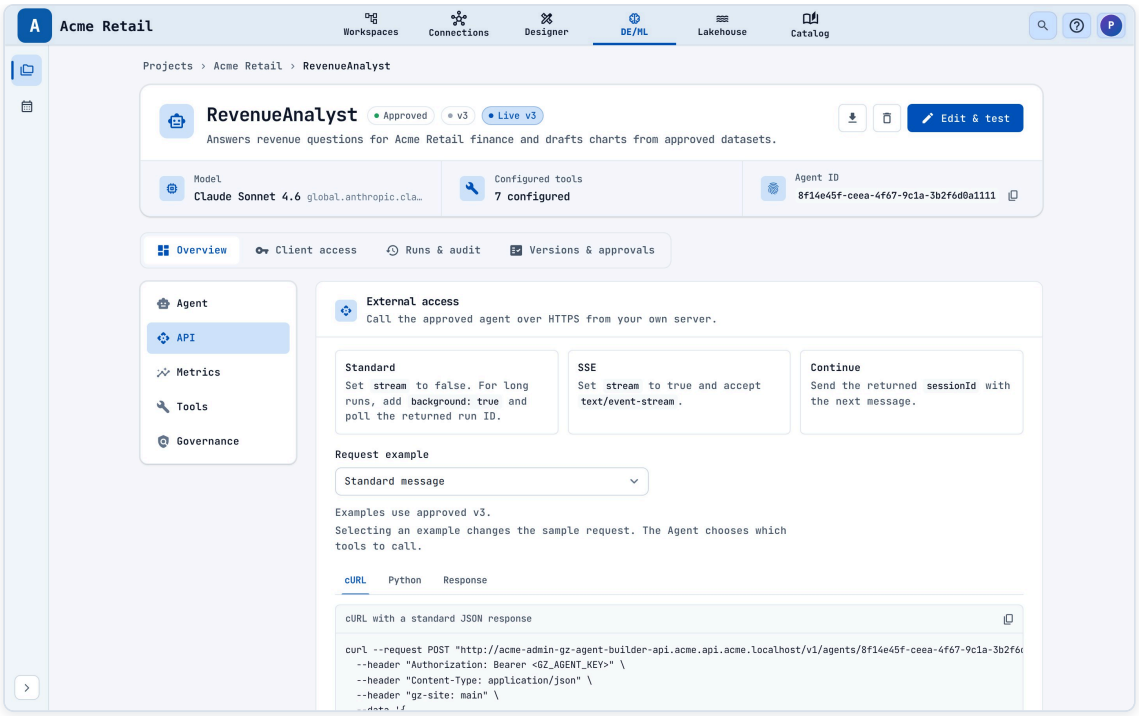


Figure 59. External access examples.

5.7.1 Create a client token

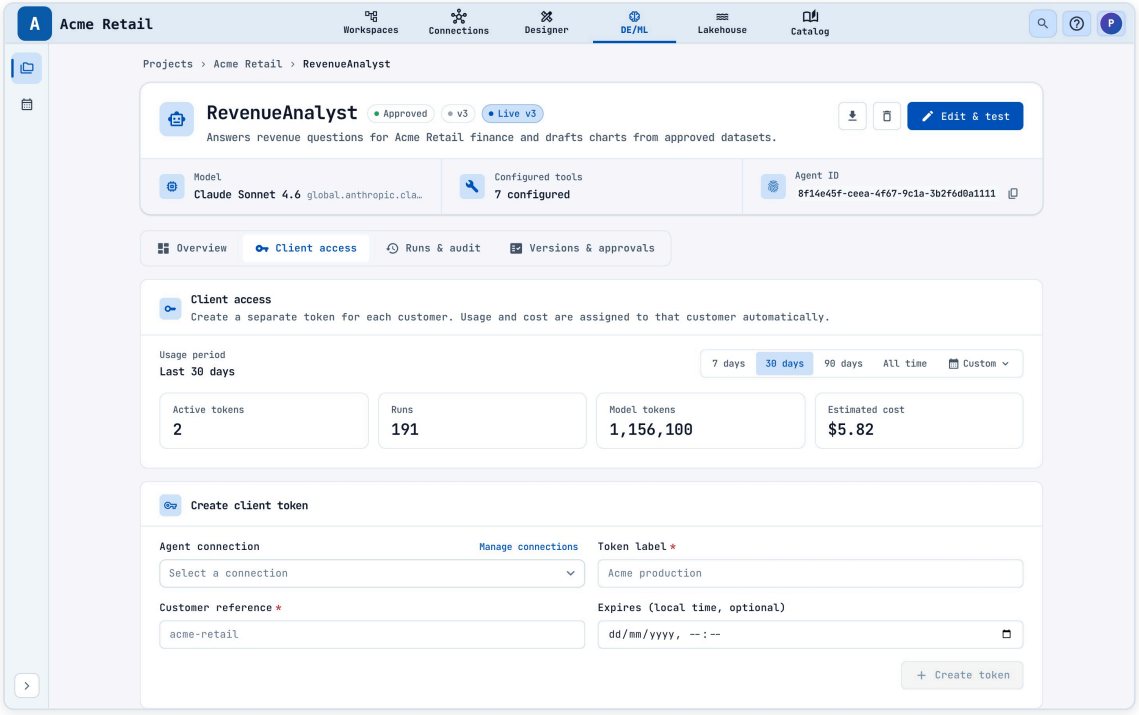


Figure 60. Client access: usage and token creation.

1. Choose the **Agent connection** (an Agent API connection). **Manage connections** opens the connections page.
2. Enter a **Token label** (up to 100 characters).
3. Enter a **Customer reference** (up to 128 characters; letters, digits and `. _ : @ / -`).
4. Optionally set **Expires**.
5. Choose **Create token**. The dialog **Client token created** shows the token once: "Copy this token now. It will not be shown again."

6. Put the token in the calling system's settings and use it in the examples from the API tab.

> **WARNING**
> Store the token on your server. Groundzero cannot show it again.

A message sent to the agent can be up to 64 KB.

How to check it worked: the **Client tokens** table lists each token's status (Active, Expired or Revoked), usage, estimated cost and last use. After the first call, the token shows usage.

To stop a customer, choose **Revoke**. New requests are refused at once, and the recorded usage stays.

5.8 Review runs and audit history

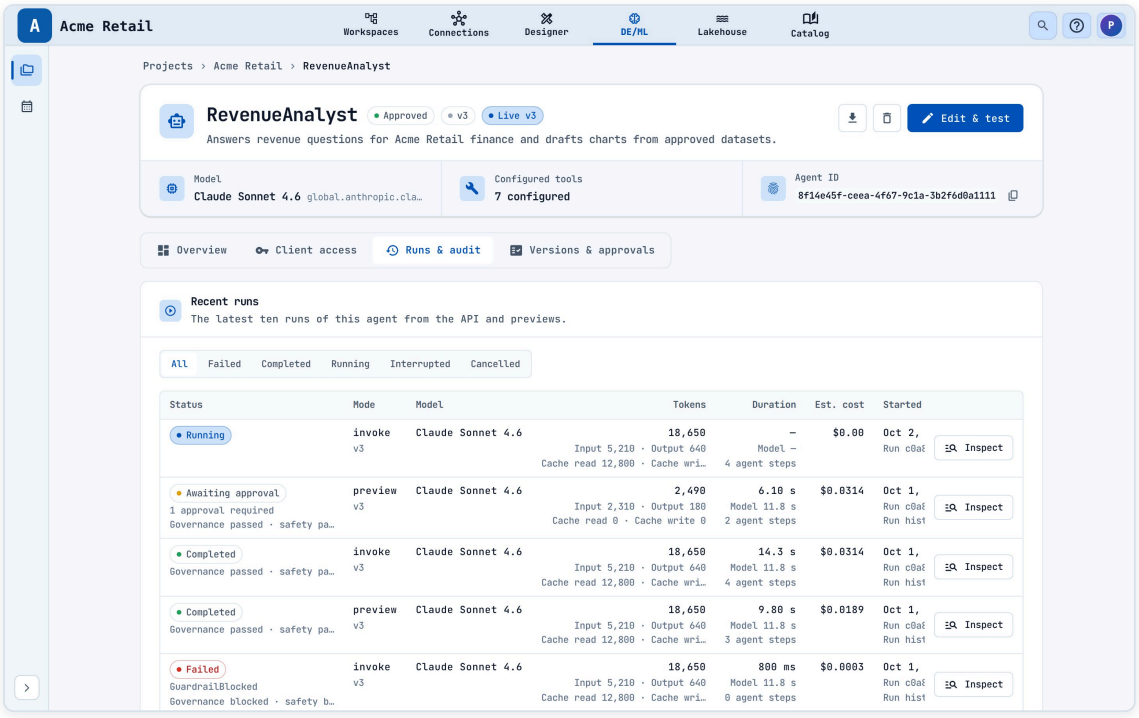


Figure 61. Recent runs and recorded actions.

Open the **Runs & audit** tab. **Recent runs** shows the latest ten, filterable by All, Failed, Completed, Running, Interrupted or Cancelled, with mode, model, tokens, duration, cost and start. **Actions & recorded reasons** lists the latest 20 audit events. Choose **Inspect** on a run to open the run inspector.

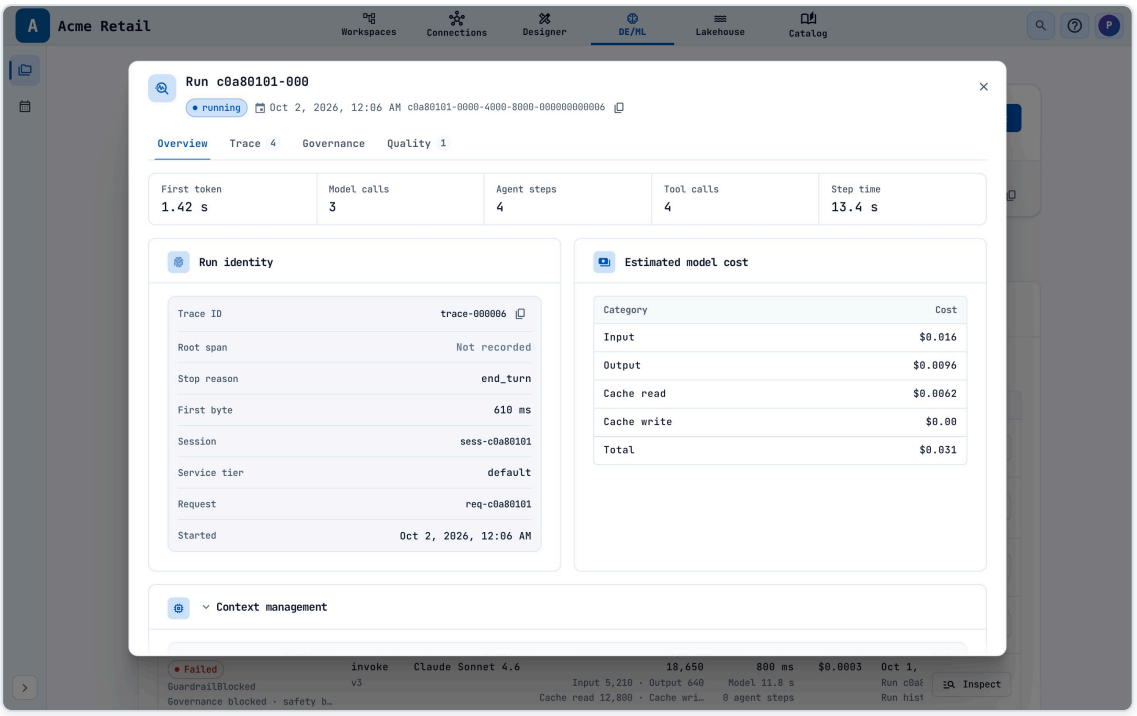


Figure 62. The run inspector.

The inspector has **Overview** (identity, cost, model calls, interventions), **Trace**, **Governance** and **Quality** tabs. Interrupted runs show **Waiting for your decision** at the top, so you can approve or deny and carry on. In **Quality**, **Run evaluation** scores a finished run, and you can leave a rating from 1 to 5 with a comment.

5.9 See versions and approvals

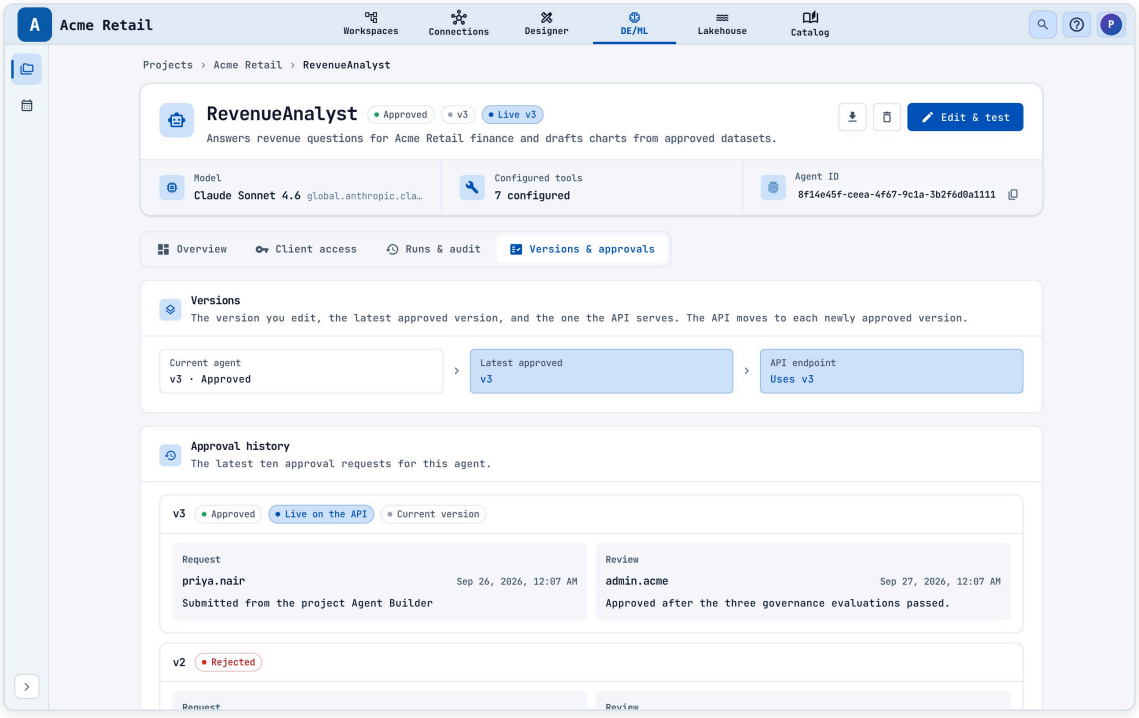


Figure 63. Versions and approval history.

The **Versions & approvals** tab shows which version you are editing, the latest approved version and the one callers get, plus the last ten approval requests with the requester's note and the reviewer's decision.

5.10 Limits

Item	Limit
Name	100 characters, starting with a letter, then letters, numbers, `_` and `-`
Description	500 characters
Instructions	100,000 characters
Tools	50 per agent (38 exist)
MCP connections	10 per agent
Max turns	300
Message size	64 KB

5.11 Troubleshooting

- **"This agent does not belong to this project."** Cause: you opened the agent from another project. Fix: open it from its own project.
- **"An agent named '<n>' already exists".** Cause: names are unique across the site. Fix: choose another name.
- **"Choose a model." or "The selected model or LLM connection is no longer available."** Fix: pick another model or repair the connection.
- **"Create an AWS S3 connection before saving this agent."** Cause: the agent has nowhere to keep its history. Fix: create an AWS S3 connection (Data guide), then choose it in step 3.
- **Approval is blocked.** Cause: the version cannot be submitted yet. Fix: read the message and use **Review governance**.
- **"These tool calls were already answered."** Cause: someone else decided. Fix: open the latest run.
- **"The paused run can no longer continue because its saved context expired."** Fix: run the agent again.
- **"This agent has spent its USD quota".** Fix: raise or remove the quota in **Agent quota**.
- **"You do not have "All" or "<Action>" permission on GZ Agent."** Fix: ask an administrator for the permission named in the message.

5.12 Related chapters

- Chapter 4 for interactive chat and approvals.
- Chapter 6 for agent metrics across the whole site.
- Chapter 7 for tokens, keys and MCP.
- The Administration guide, chapter 4, for the Agent permissions.

6 Consumption metrics

Consumption Metrics shows how much AI and compute your site uses and what it costs. Use it to:

- See who uses AI chat the most, which models they use and what it costs.
- Track agent runs, reliability, speed, cost and quality over time.
- Find expensive jobs and datasets.
- Pull the same numbers into your own tools (see "Use the numbers in your own tools").

6.1 Open the metrics and pick a period

Outcome: the numbers for the period you care about.

1. Open **Catalog** in the top navigation and choose **Consumption Metrics** in the left menu.
2. Choose a view: **AI Chat**, **Agents**, **Workloads** or **APIs**.
3. At the right of the tab bar, choose the period. It applies to every view.

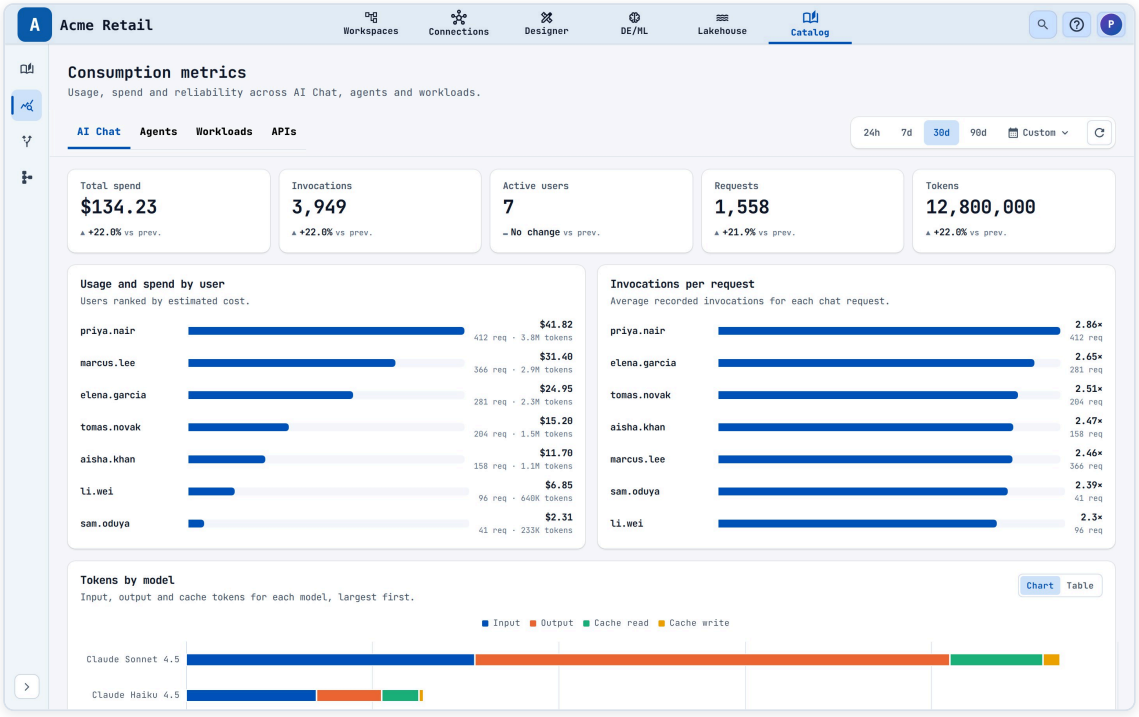


Figure 64. AI Chat view with KPIs, ranked users and token mix.

Control	Period
24h	The last 24 hours
7d, **30d**, **90d**	Whole days (UTC). 30d is the default
Custom	A date-range popover with presets Today, Yesterday, Last 7 days, Last 30 days, This month, Last month, Quarter to date and Year to date, two date boxes and a calendar. Choose **Apply**
Refresh button	Reloads the numbers. A rolling 24 hours moves forward. Fixed day ranges keep their dates

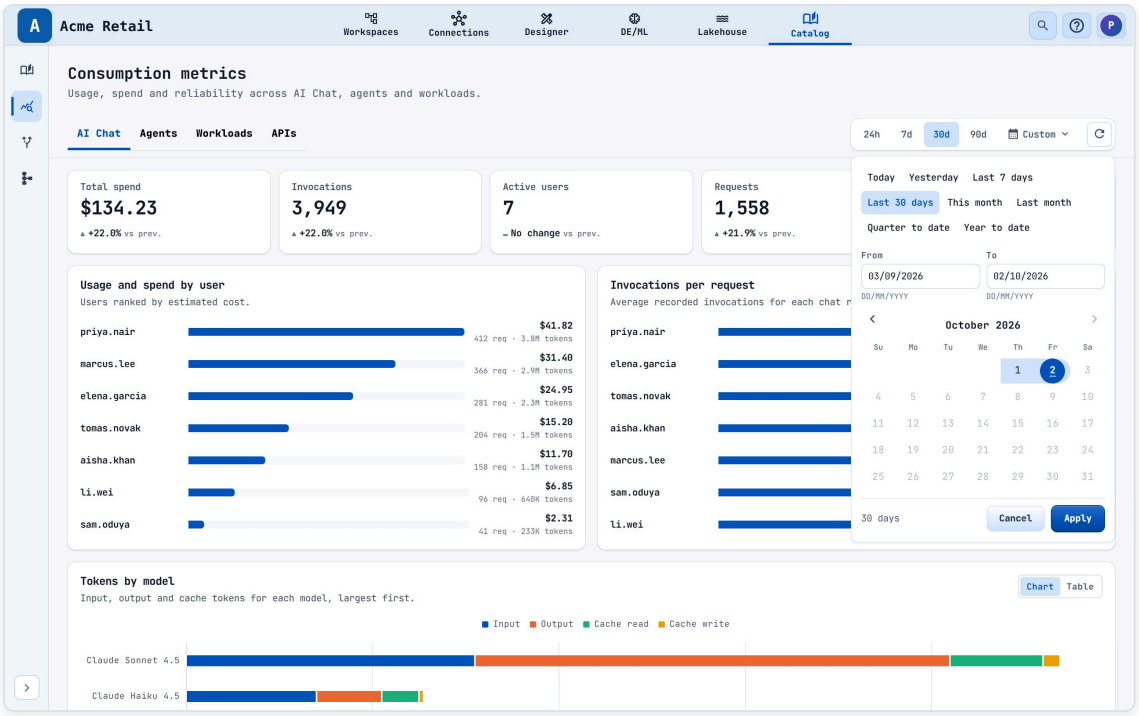


Figure 65. The custom range popover.

Headline numbers show how they changed compared with the previous period of the same length ("vs prev.").

> **NOTE**
> Costs are estimates. They use the provider's prices plus the Groundzero markup (20 percent). Calls whose usage is unknown are marked as estimated.

If part of the data cannot be loaded, the view shows a yellow notice (for example "AI metrics: ...") and keeps what did load. Some sites do not have the per-user breakdown, and then AI Chat shows "Per-user metrics are unavailable in this environment."

6.2 See who uses AI chat and what it costs

Choose the **AI Chat** view.

Block	What it tells you
KPI row	Total spend, Invocations, Active users, Requests, Tokens
Usage and spend by user	Users ranked by estimated cost, with requests and tokens
Invocations per request	Average model calls per chat request, per user
Tokens by model	Input, output, cache read and cache write tokens per model, with a **Chart** or **Table** switch
User breakdown	A table of User, Requests, Invocations, Tokens and Estimated cost, with **Search users**
Model breakdown	A table per model including Output speed (tokens per second), First token (median) and Estimated cost, with a note counting estimated invocations

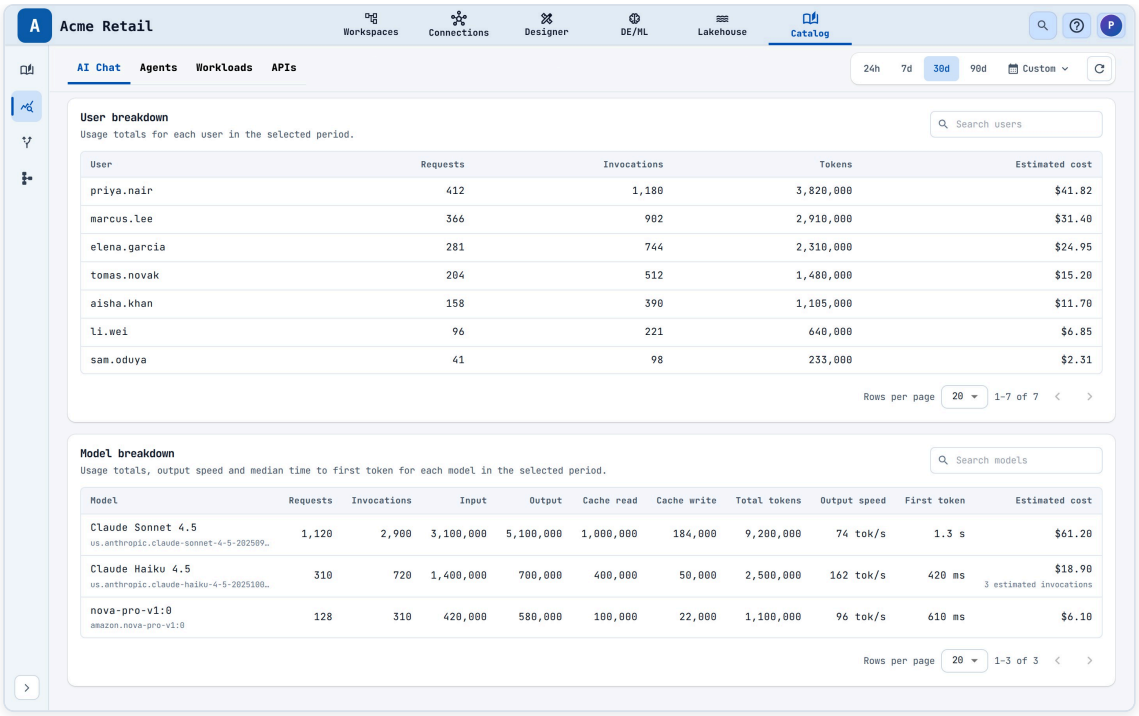


Figure 66. The breakdown tables.

Charts show the top 20 rows and say "Top 20 of N users ranked by estimated cost." when there are more.

What to do with it: if one person or model drives the cost, ask your administrator to set or adjust a budget (Administration guide, chapter 7).

6.3 Track how your agents are doing

Choose the **Agents** view.

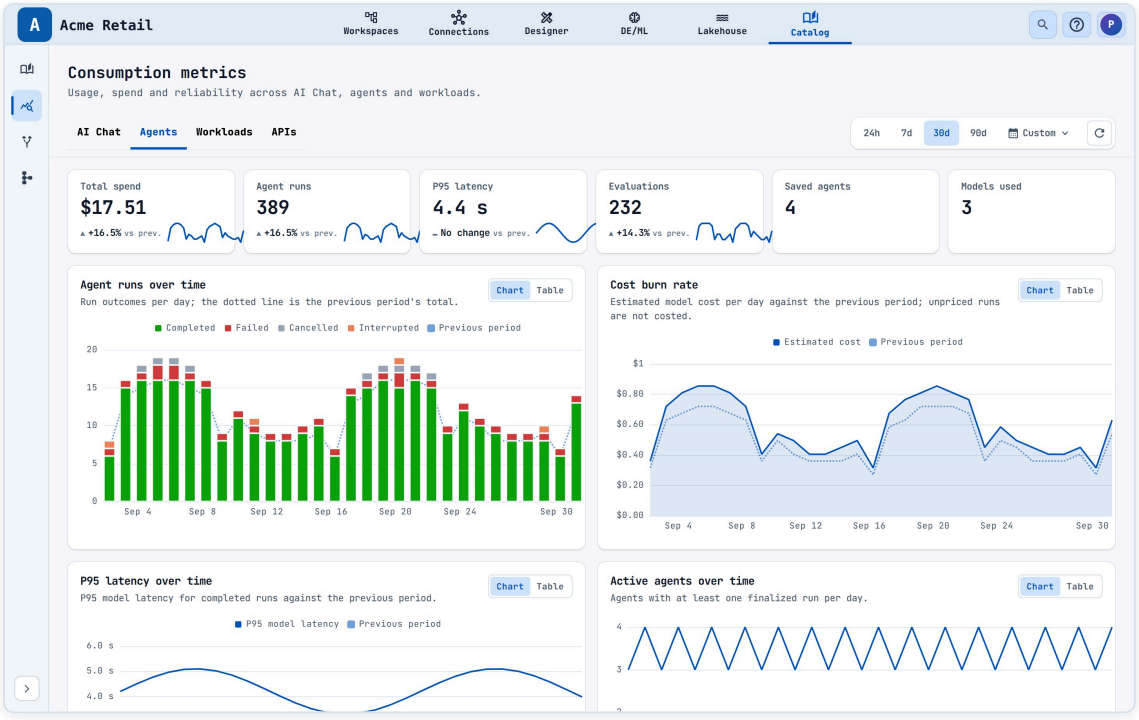


Figure 67. Agents view: KPIs and trends.

KPIs are Total spend, Agent runs, P95 latency (the time within which 95 percent of runs finish), Evaluations, Saved agents and Models used, with trend lines. The charts compare with the previous period (dotted lines):

- **Agent runs over time**: Completed, Failed, Cancelled and Interrupted runs.
- **Cost burn rate**: estimated model cost per day. Runs with no known price are not costed.
- **Runs per active agent**, **P95 latency** and **Active agents** over time.
- **Model adoption over time**, **Evaluation quality over time** and **Evaluation dimensions over time** (Helpfulness, Correctness, Relevance, Coherence, Conciseness, Goal completion, Tool use).
- **Model breakdown** table.

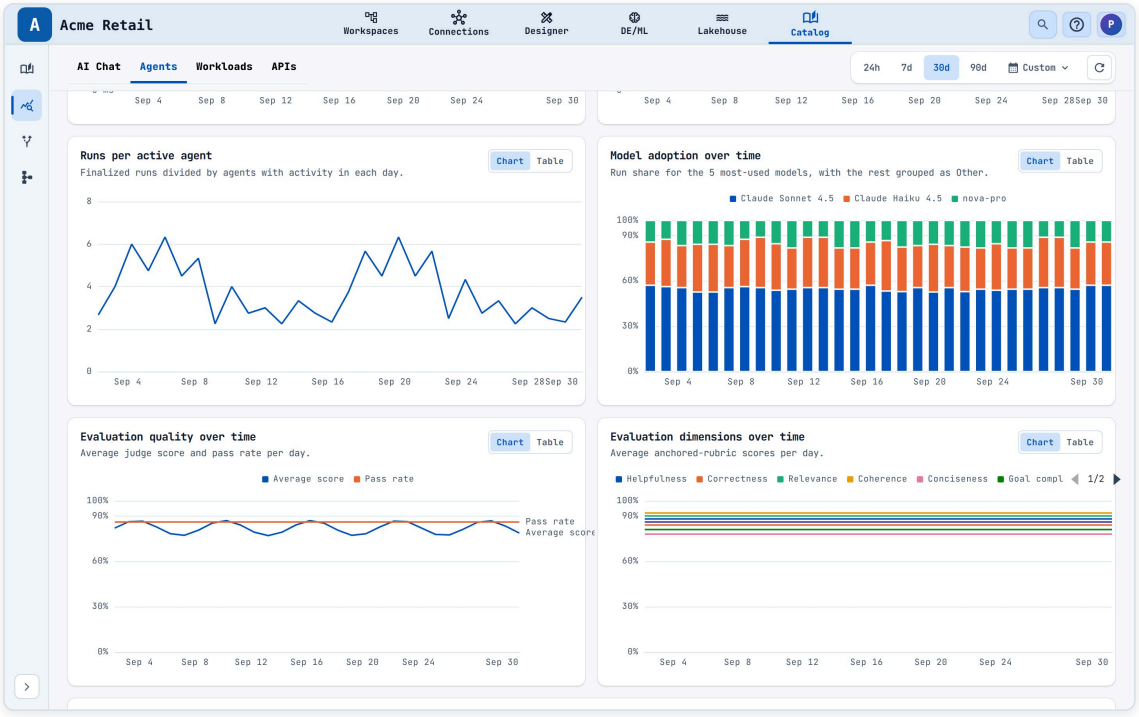


Figure 68. More agent charts.

Only finished runs of saved agents are included. For 24 hours the points are hourly. Longer ranges use days or weeks.

6.4 Find expensive jobs and datasets

Choose the **Workloads** view. It shows how much ETL jobs and datasets consume. Drill from **Projects** to a project, a job and a single run, and switch between **Jobs** and **Datasets**. Tiles show Executions, Success rate, Average duration, CPU consumed, Memory consumed and Cost.

With no runs in the period you see "No projects found: Change the period or filters to include executions."

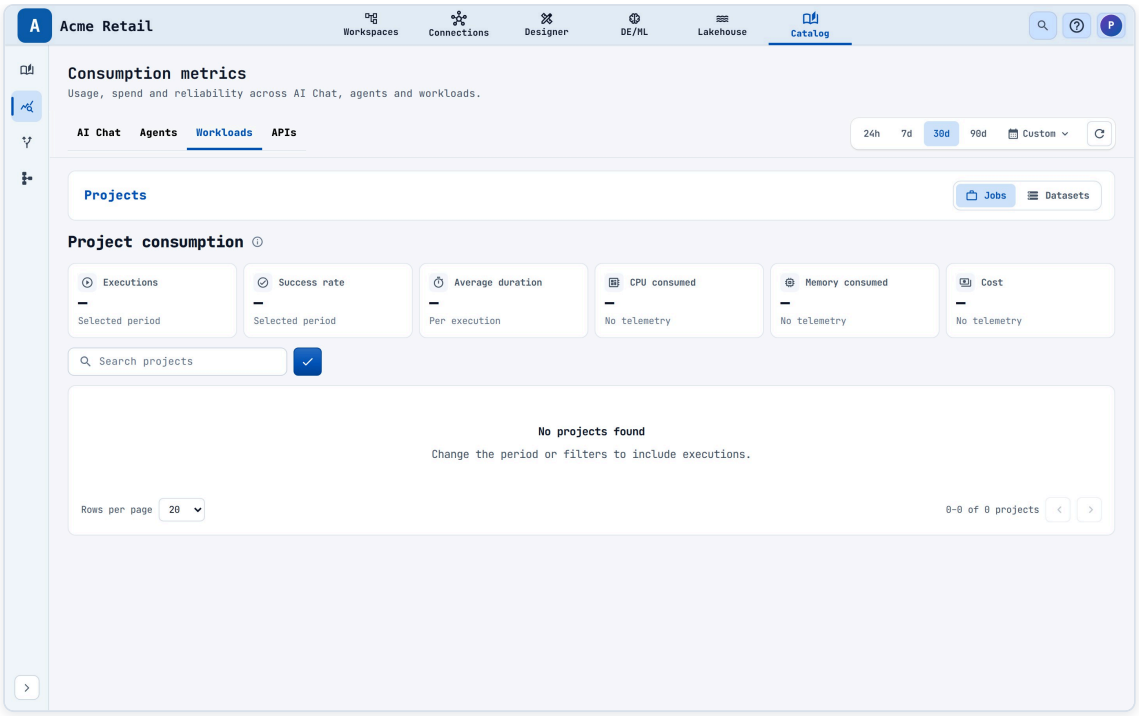


Figure 69. The Workloads view with no executions in the selected period.

For how jobs and datasets run, see the Build and Run guide.

6.5 Use the numbers in your own tools

Choose the **APIs** view. It lists the metric calls you can make from your own scripts and reports, each with its parameters, what it returns and a ready-made cURL command to copy.

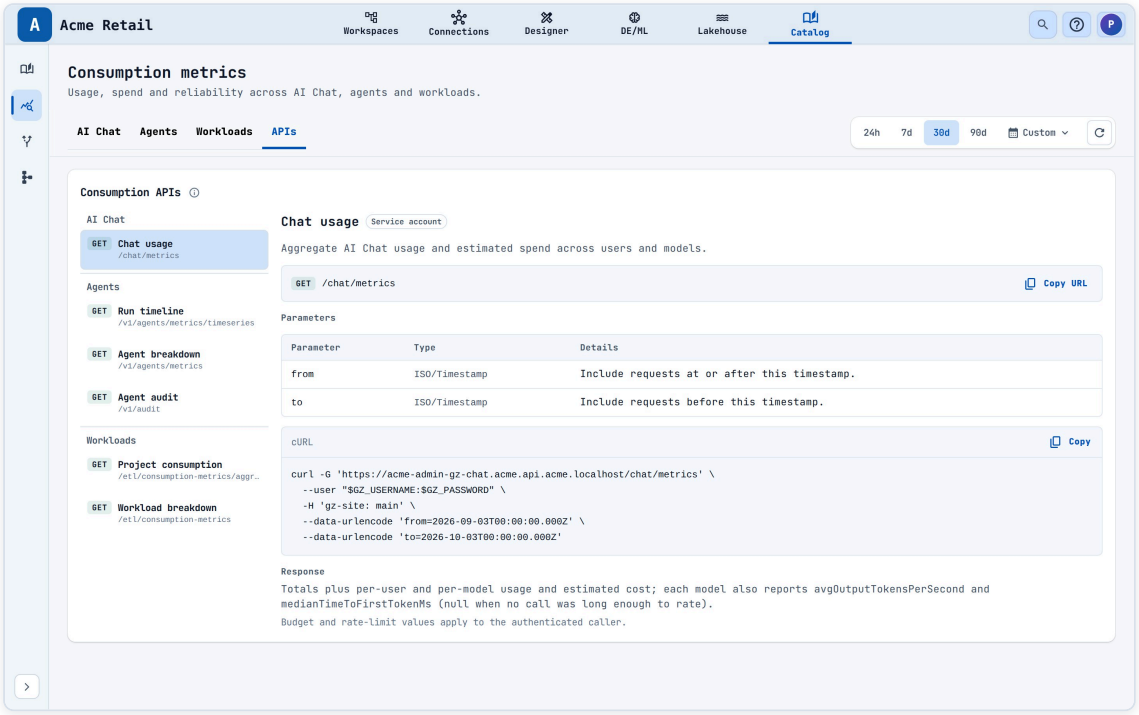


Figure 70. The APIs reference with cURL examples.

The calls cover:

Group	What you get
AI Chat	Chat usage and estimated spend by user and model for a date range
Agents	Run, model, token, latency and cost totals per agent, and totals per hour, day or week, plus the lifecycle and action history of saved agents
Workloads	The workload breakdown and consumption per project

To run an example, replace the `$GZ_USERNAME` and `$GZ_PASSWORD` placeholders with your own sign-in. Chapter 7 lists the other ways to reach Groundzero from your own tools.

6.6 Troubleshooting

- **"Unable to load AI metrics."** Cause: the data could not be loaded. Fix: choose the refresh button.
- **"Unable to load saved-agent metrics." or "Unable to load agent trends."** Fix: refresh. If it repeats, try a shorter period.
- **A cost shows an "estimated invocations" or "unpriced" note.** Cause: the model provider did not report its usage, so the cost is an estimate or left out. Fix: none needed. Treat the figure as approximate.
- **The "vs prev." changes are missing.** Cause: the previous period could not be loaded. Fix: refresh.
- **Workloads says "No projects found".** Cause: nothing ran in the period. Fix: widen the period.

6.7 Related chapters

- The Administration guide, chapter 7, for budgets, which cap AI chat spend.
- Chapter 5 for the metrics of a single agent.

7 API access, your account and help

This chapter covers three things you do outside the main features: reaching Groundzero from your own tools, looking after your account, and finding help. Everything is in the header of Groundzero: an avatar menu for your account and a question-mark menu for help and external-tool setup.

7.1 Reach Groundzero from your own tools

Choose the method that fits what you want to do.

I want to...	Use	Where to get it
Let a customer's system or my server call an approved AI agent	**Agent client token**	Agent details > Client access (chapter 5)
Run scripts or integrations as me or as a service account	**API key**	Your administrator creates it (Administration guide, chapter 2)
Use Groundzero from Codex, Claude Code, Cursor or VS Code	**MCP**	Help menu > Connect MCP (below)
Report on AI and compute use from my own tools	The metric calls in the **APIs** view	Catalog > Consumption Metrics (chapter 6)

7.1.1 API keys

An API key is a long, reusable secret that stands in for a user or a service account. Administrators create keys, and each key is shown only once when it is created. Keys start with `sk_`. Send one in the `Authorization` header as `Bearer sk_...`. If you lose a key, ask your administrator for a new one. Treat a key like a password.

7.1.2 Agent client tokens

A client token lets one customer or system call one approved AI agent, and keeps that customer's usage and cost separate. Create tokens in the agent's **Client access** tab (chapter 5).

7.2 Connect a coding assistant with MCP

Outcome: your coding assistant can use Groundzero's tools, with your approval for anything that changes data.

MCP (Model Context Protocol) lets coding assistants and other tools work with Groundzero.

1. Choose the question-mark icon in the header, then **Connect MCP**.
2. Pick your client: **Codex**, **Claude Code**, **Cursor**, **VS Code** or **Generic HTTP**.
3. Copy the command or configuration shown, using the copy buttons (a toast says "<label> copied to clipboard."). The **MCP URL** box shows the address to use. The server is named `groundzero`.
4. Run the command or paste the configuration into your client.
5. Sign in through your browser when your client asks.

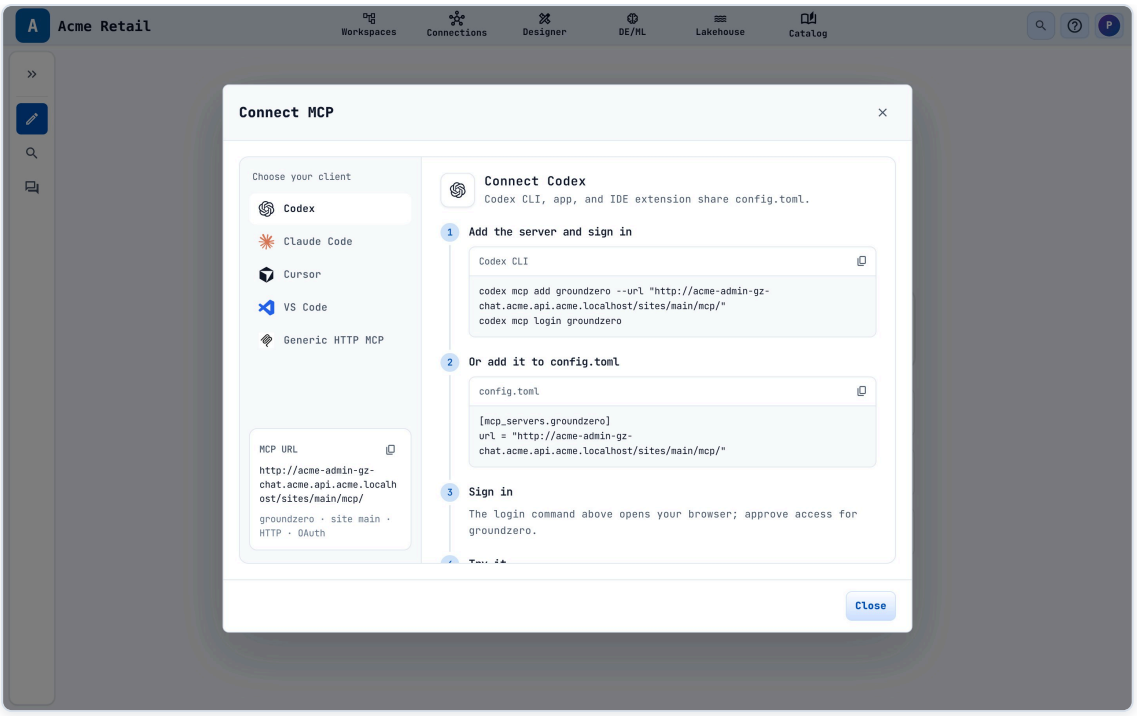


Figure 71. Connect MCP with client-specific steps.

How to check it worked: ask your assistant "List the datasets I can access in Groundzero." You should see your datasets.

Tools that change data still ask for your approval (chapter 4).

7.3 Open the help menu

Choose the question-mark icon in the header.

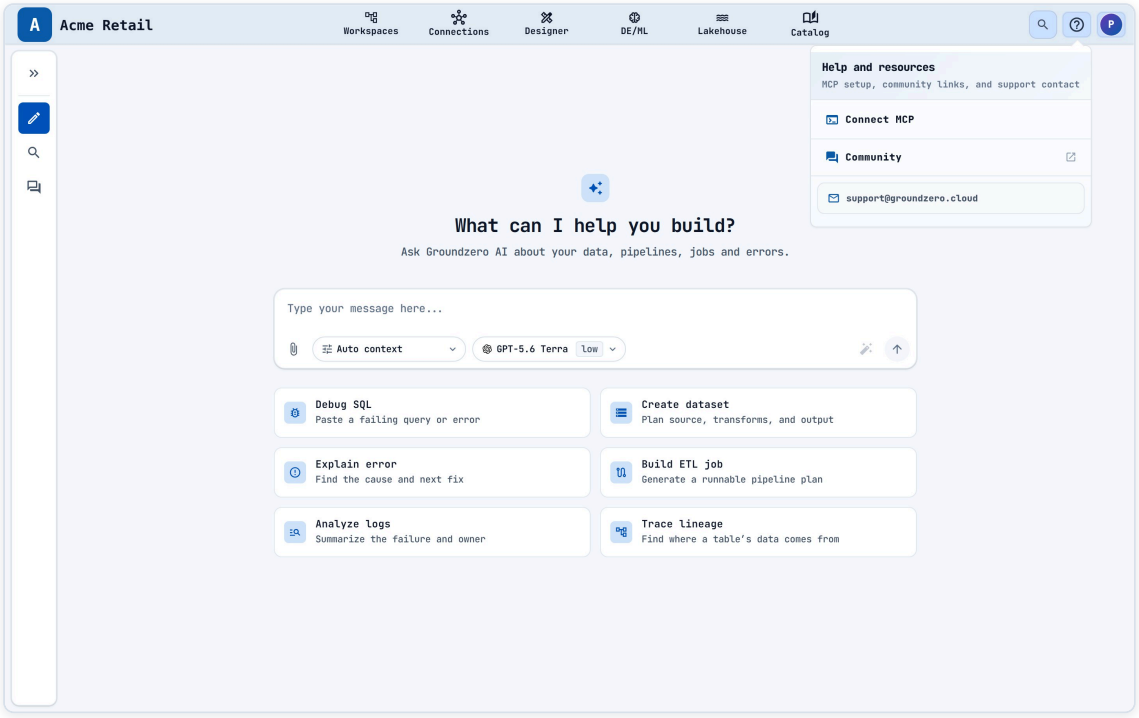


Figure 72. The Help and resources menu.

Item	Opens
Connect MCP	The Connect MCP dialog
Community	The Groundzero community chat in a new tab
Support address	An email to Groundzero support

7.4 Change your password

Outcome: a new password for your account.

1. Choose the round avatar (your initial) at the right of the header.
2. Choose **Change Password**.
3. Fill **Current Password**, **New Password** and **Confirm New Password**.
4. Choose **Change Password**. The toast says "Password changed successfully."

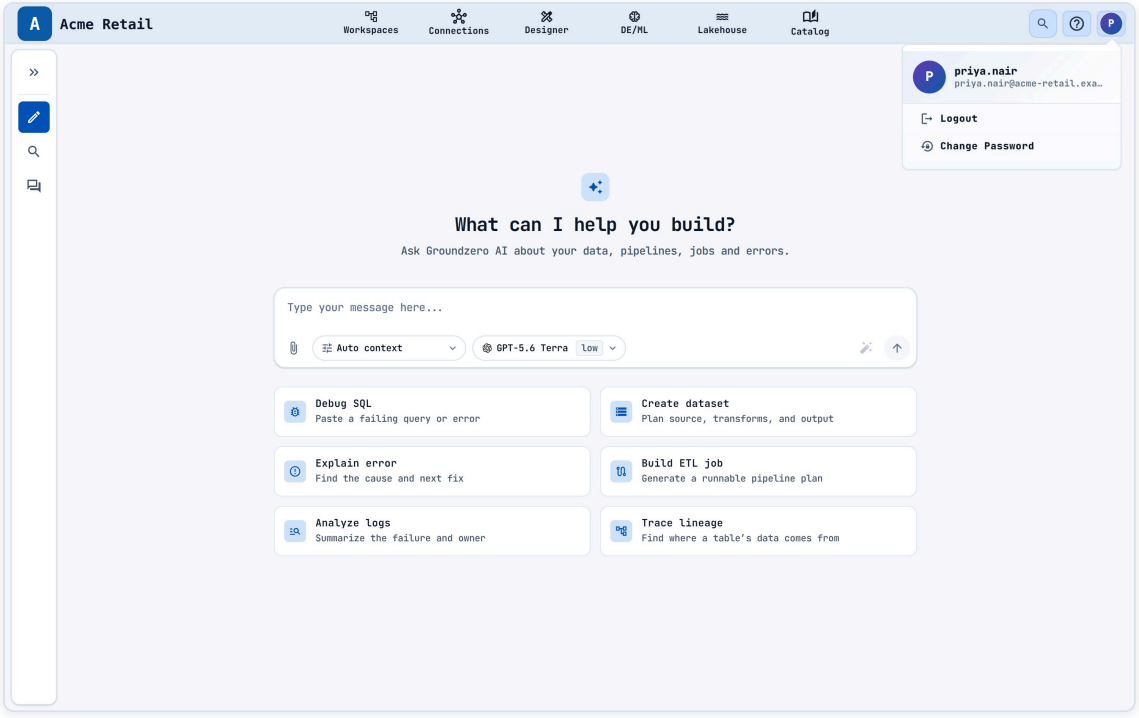


Figure 73. The account menu shows your name, email, Logout and Change Password.

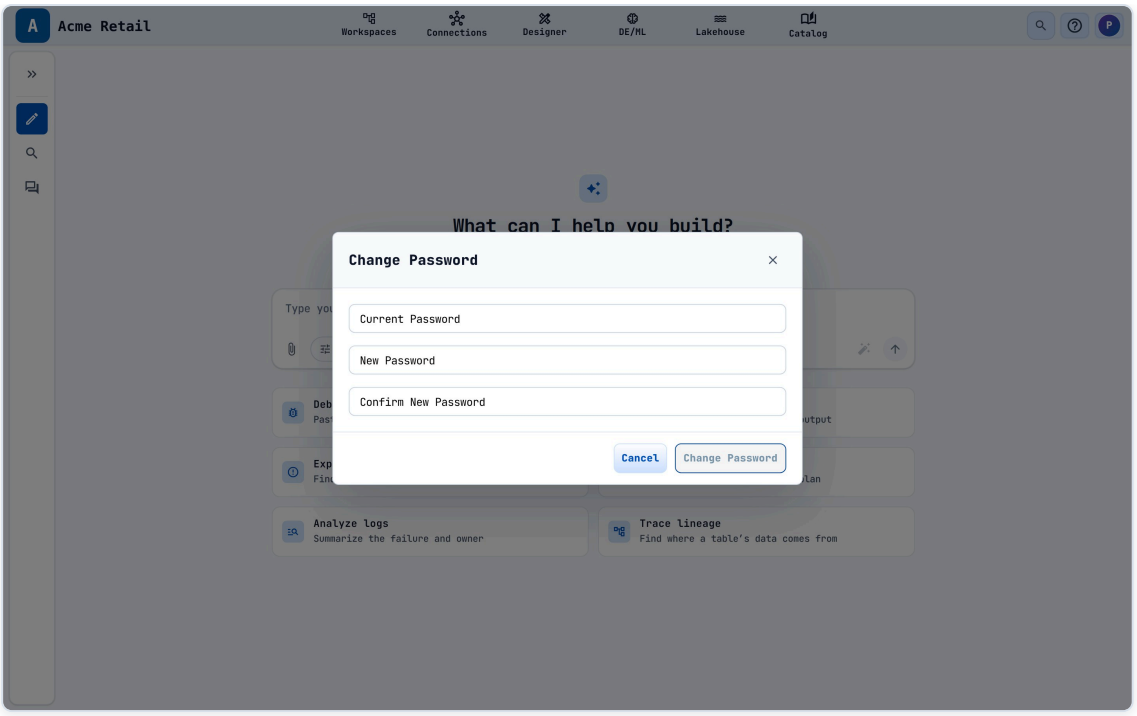


Figure 74. The Change Password dialog.

The account menu shows your username and email, and **Logout** signs you out.

Rules: all three fields are required ("All fields are required."), the two new passwords must match, and the new password needs at least 8 characters with an uppercase letter, a lowercase letter and a special character.

If it fails you see "Incorrect username or password. Please try again." Check your current password and try again.

> **NOTE**
> Administrators set other people's passwords with a stricter rule (12 characters including a number). A long password is a good idea for your own account too.

7.5 Set up two-step verification

After you sign in, if your account does not yet use an authenticator app, Groundzero shows **Set up two-step verification** instead of the console.

1. Scan the QR code with an authenticator app such as Google Authenticator or Authy. If you cannot scan, choose "Or enter this code manually".
2. Enter the 6-digit code in **Authentication code**.
3. Choose **Verify**.

From then on, signing in asks for the code from your app.

If something goes wrong:

- **"Invalid authentication code".** Cause: the code expired or was mistyped. Fix: enter the current code from the app.
- **"Failed to initialize two-step verification" or "Failed to generate two-step verification QR code".** Fix: reload the page and start again.

- **You lost your authenticator app.** Ask your administrator for help.

7.6 Related chapters

- The Administration guide, chapter 2, for API keys and passwords set by administrators.
- Chapter 4 for approvals of MCP tool calls.