Spider Strategies®

What's New in 5.8

Guide to new features in Spider Impact 5.8
Updated October 3rd, 2025

Table of Contents

Your success is important to us!

Spider Impact 5.8 introduces Impact Assistant—ask questions in plain language and get instant answers and visualizations. No technical skills required. Your data stays completely secure through a metadata-only approach where actual values never leave your environment.

Forms have evolved into an application development platform. New capabilities like form variables, equations in conditions, behind-the-scenes record manipulation, and dialog-based editing enable you to build sophisticated application-style functionality. Auto-load user records, prevent duplicates, create audit trails, and manipulate multiple datasets without leaving the form—all while delivering polished experiences with file attachments, multi-column layouts, and linked record editing in dialogs.

Apps package these advanced forms alongside dashboards and reports into standalone applications with dedicated URLs. Build curated experiences for specific tasks—a project intake system, customer portal, or operational tracker—each accessible independently of the main Spider Impact interface and secured through existing permissions. Users can bookmark these apps for direct access, creating a streamlined experience focused on their specific workflows.

Beyond these major features, version 5.8 refines virtually every part of the application. From new bullseye charts to significant performance improvements, it's an update that makes Spider Impact both more powerful and more intuitive to use.

To help you discover everything Spider Impact has to offer, we have free training videos on our website, and we've put together new functionality guides like this one. If you want to maximize your return on investment, we offer paid formal training courses and dedicated consulting engagements. We have more information about all our free and paid services at https://www.spiderstrategies.com/services/

Impact Assistant

Spider Impact now includes Impact Assistant that enables conversations with your data and provides instant help with using the software. This powerful new capability marks the beginning of Impact Intelligence—a comprehensive suite of AI-powered features that will continue to expand and transform how you work with Spider Impact.

Using Impact Assistant for Data Analysis

Impact Assistant enables you to ask questions about your data in natural language and receive instant visualizations and answers—no technical expertise required.

Accessing Impact Assistant

Click the Impact Assistant icon inside Spider Impact:

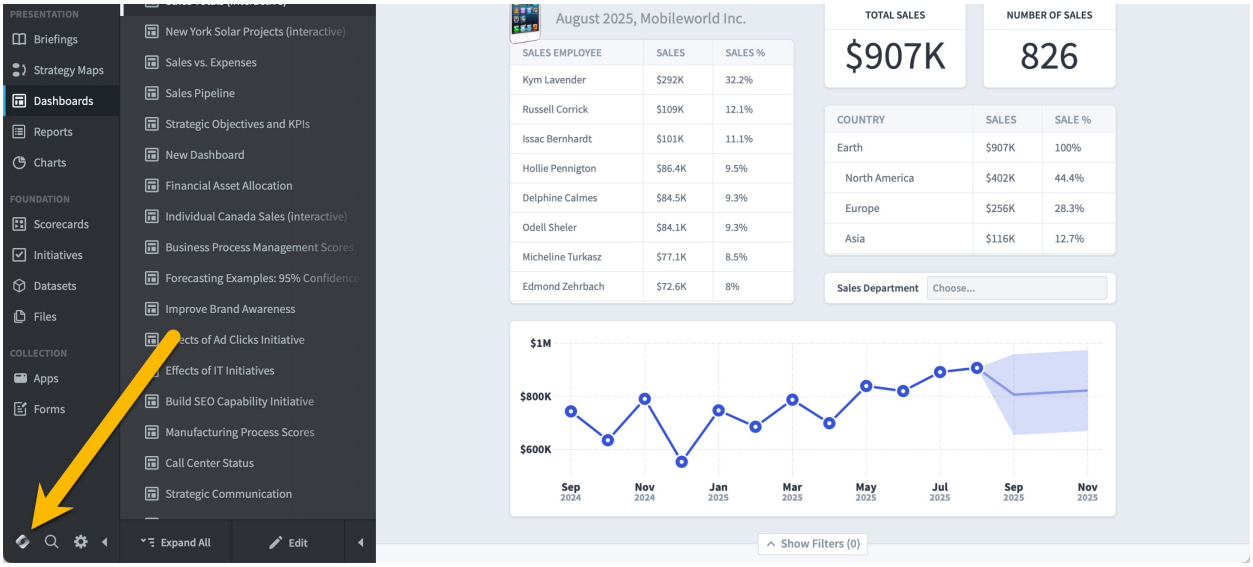

Or inside any stand-alone App you've built:

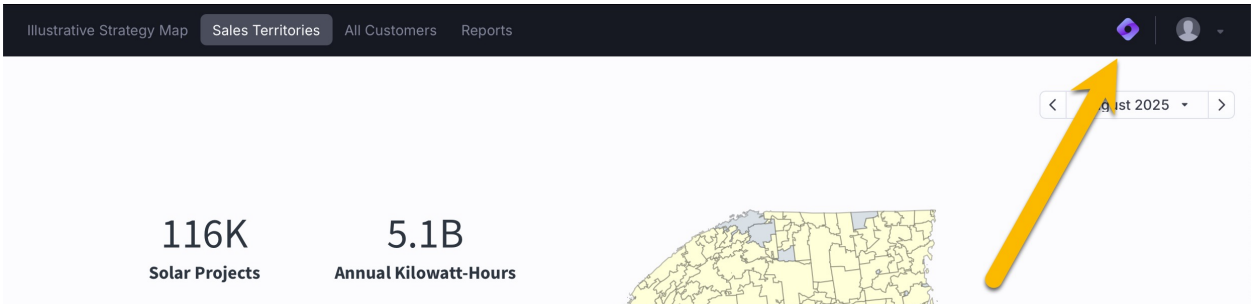

Asking Questions About Your Data

While dashboards and reports excel at displaying predetermined information, Impact Assistant shines when you need specific answers that aren't already configured. Simply ask questions like "What was the total revenue by product category last quarter?" or "Show me the trend of customer complaints over the past year."

In this example, asking how many devices were sold last year returns a clear answer:

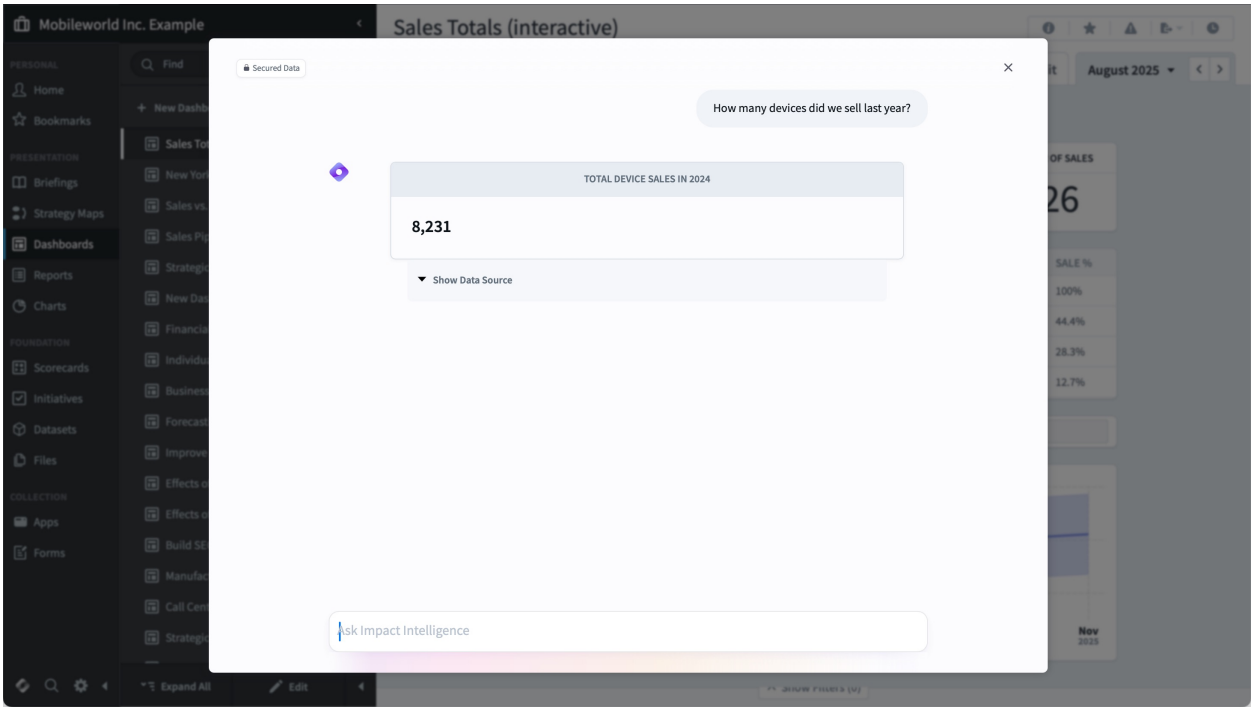

Clicking on the "8,231" number drills down into the Datasets section for further ad-hoc analytics:

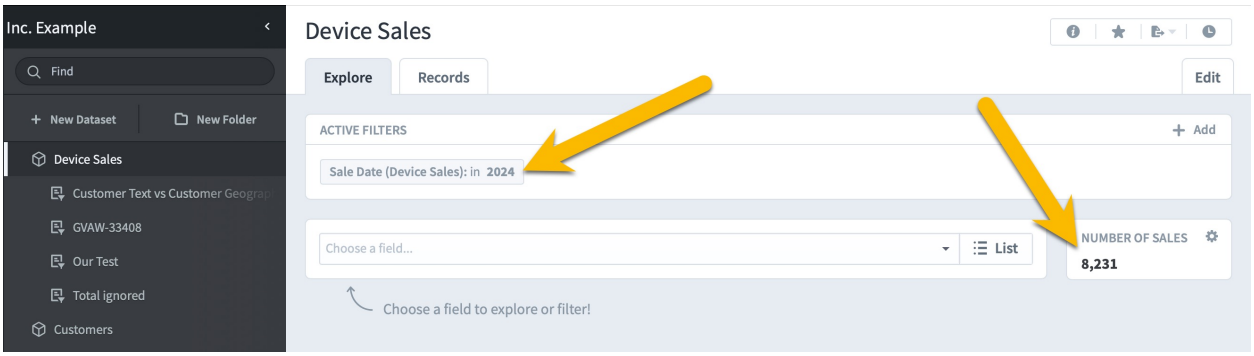

Conversational Context

Impact Assistant maintains conversation history for contextual follow-up questions. When you say, "break that down by salesperson," Impact Assistant understands you're still discussing last year's sales:

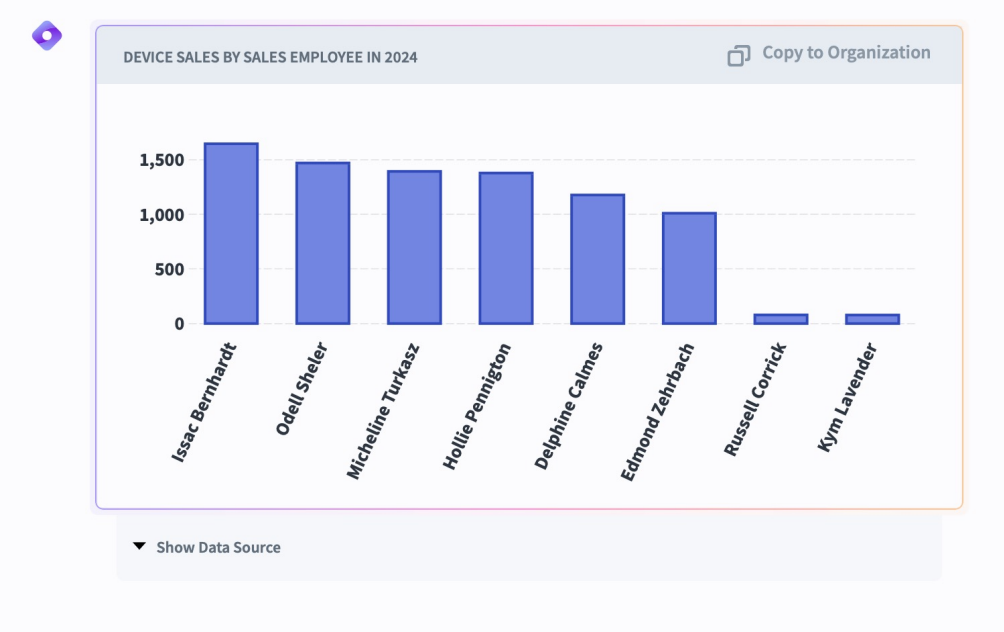

Saving Insights

When Impact Assistant creates a useful chart, use the "Copy to Organization" button to save it permanently. This allows you to reuse assistant-generated charts in dashboards and share insights with your team:

Discovering Data and Insights

Impact Assistant helps you understand your data landscape and discover meaningful insights. You can ask general questions like "What kinds of data do I have?" to see a list of datasets and what they track:

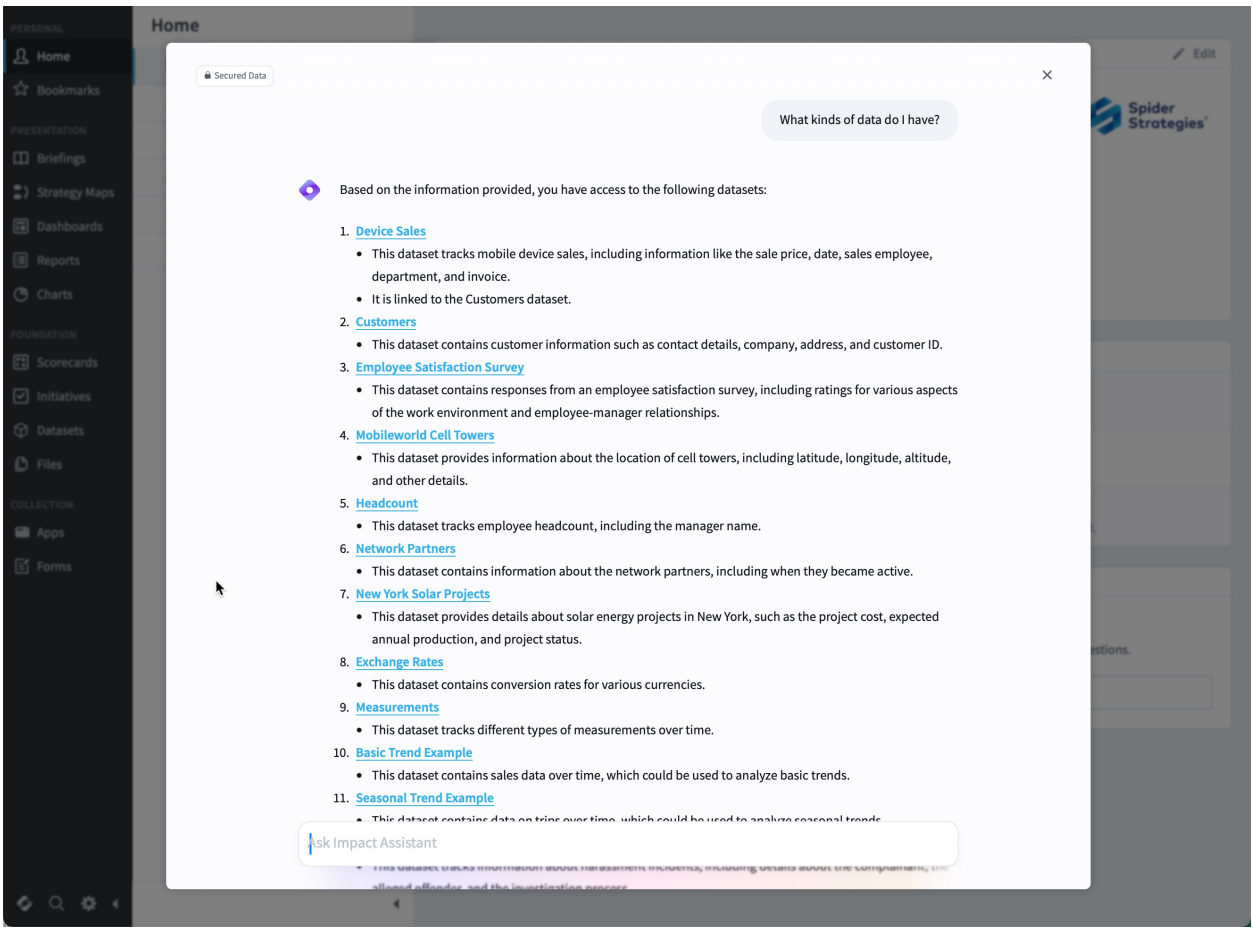

One of Impact Intelligence's most powerful features is proactively suggesting insights to explore based on your data:

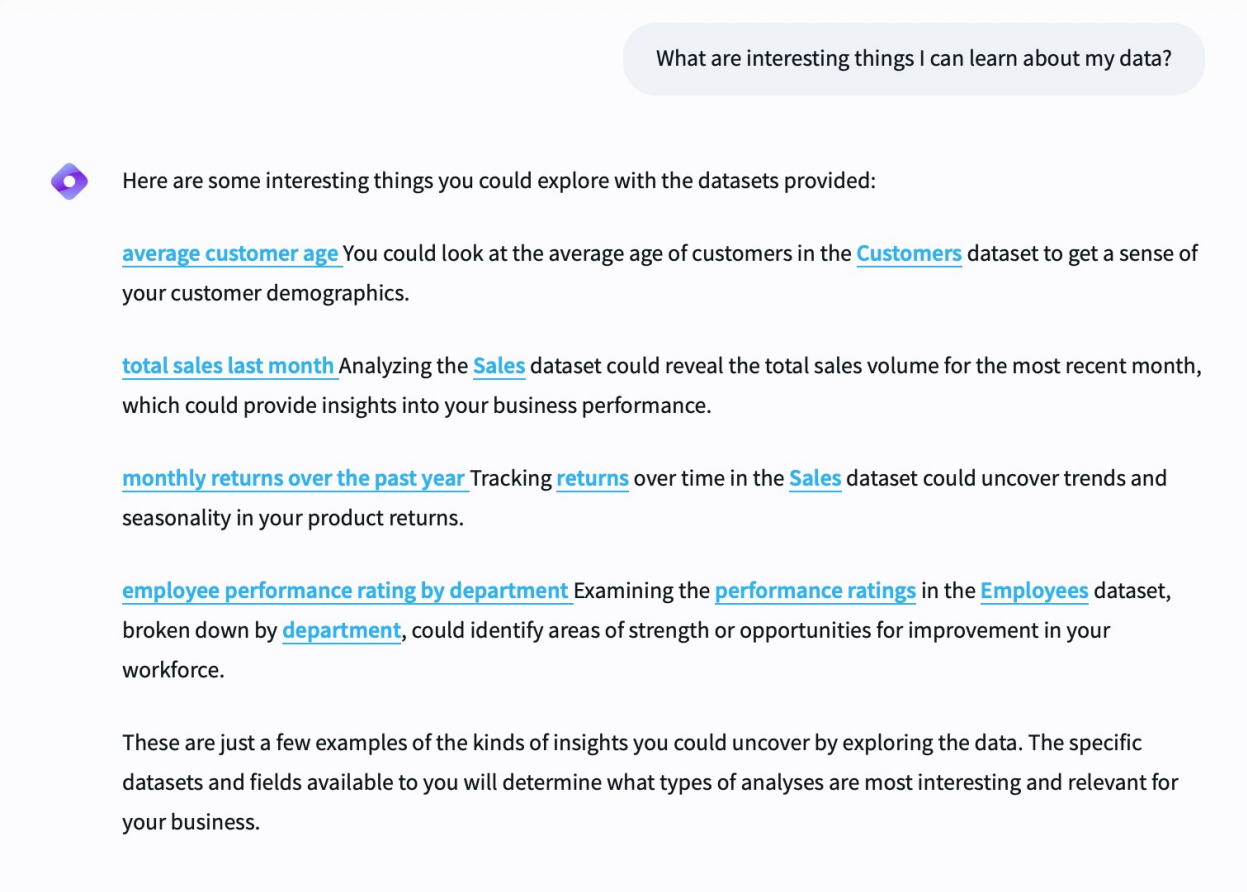

Clicking on suggestions like "total sales last month" provides immediate answers:

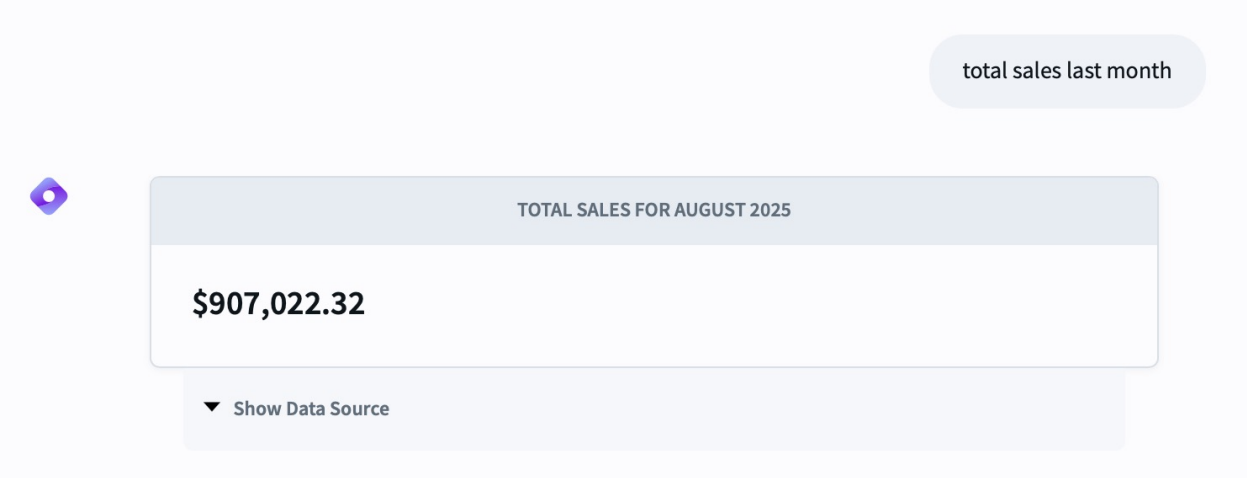

Getting Software Help

Beyond data analysis, Impact Assistant answers questions about Spider Impact features with detailed guidance and documentation links:

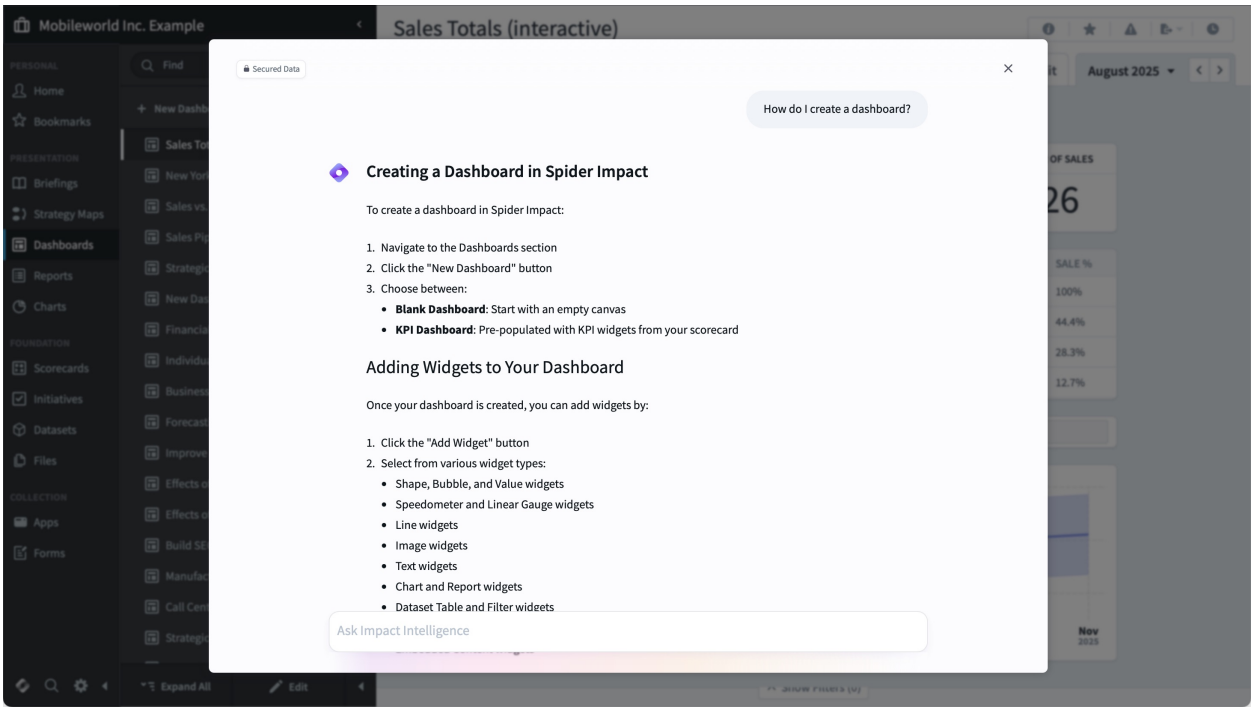

Current Capabilities

Impact Assistant currently:

- Aggregates and filters data from datasets to answer specific questions
- Generates appropriate visualizations (line, bar, and pie charts)
- Suggests potential insights from your data
- Answers questions about Spider Impact features with documentation links

Impact Assistant currently provides answers from datasets only. It cannot access scorecard KPI values or initiative data. Future releases will expand these capabilities.

Enterprise-Grade Security and Privacy

Impact Intelligence maintains complete data security through an innovative architecture designed to protect your sensitive information while delivering powerful AI capabilities.

How Your Data Stays Secure

Your actual data values never leave your secure environment. Impact Intelligence uses a metadata-only approach where Impact Assistant receives only:

- Dataset names and their descriptions
- Field names and field types
- Aggregate statistics about fields (not associated with any individual record):
    - For date fields: The earliest and latest dates
    - For text fields: The 10 most common values

This metadata enables Impact Assistant to understand what questions it can answer and provide meaningful filtering and queries, while ensuring sensitive data remains completely protected.

Example of Privacy Protection

Consider a dataset called "Customer Orders" with thousands of records. If the dataset contains a "Product Type" field where "laptops," "tablets," and "phones" each appear in at least 1% of records, Impact Intelligence will know these are common product types to support queries like "show me laptop sales last quarter."

However, if the same dataset contains a "Shipping Address" field, individual addresses almost certainly won't represent 1% or more of records, so specific addresses are never sent to the AI. This protects personally identifiable information while still enabling powerful data analysis.

Government Cloud (FedRAMP) Security

For Spider Impact government customers, all infrastructure operates within Amazon GovCloud behind our secure FedRAMP-authorized boundary. The AI services are also within Amazon GovCloud and part of this secure boundary. While we are fully authorized to send secure information to these AI services, we maintain our metadata-only approach for additional privacy protection.

Commercial Cloud Security

For commercially hosted customers, we use Amazon Web Services with enterprise-grade security to process metadata. The AI services have contractual agreements that **your data will never be used for AI training**. This guarantee applies to both the metadata we send and ensures your organization's information remains completely private.

Key Security Benefits

- **No data values leave your environment** - All actual data processing happens securely on your server
- **AI cannot hallucinate data** - Since all values come directly from your database, Impact Assistant can only show real information
- **No AI training on your data** - Contractual agreements prevent any use of your information for model training
- **High performance** - Minimal data transfer means fast, responsive interactions

Configuring Impact Intelligence

Enabling Impact Intelligence

Administrators must first enable Impact Intelligence in Application Administration. Toggle "Enable Impact Intelligence" to "Yes" and provide information about your organization in the Organization Context field:

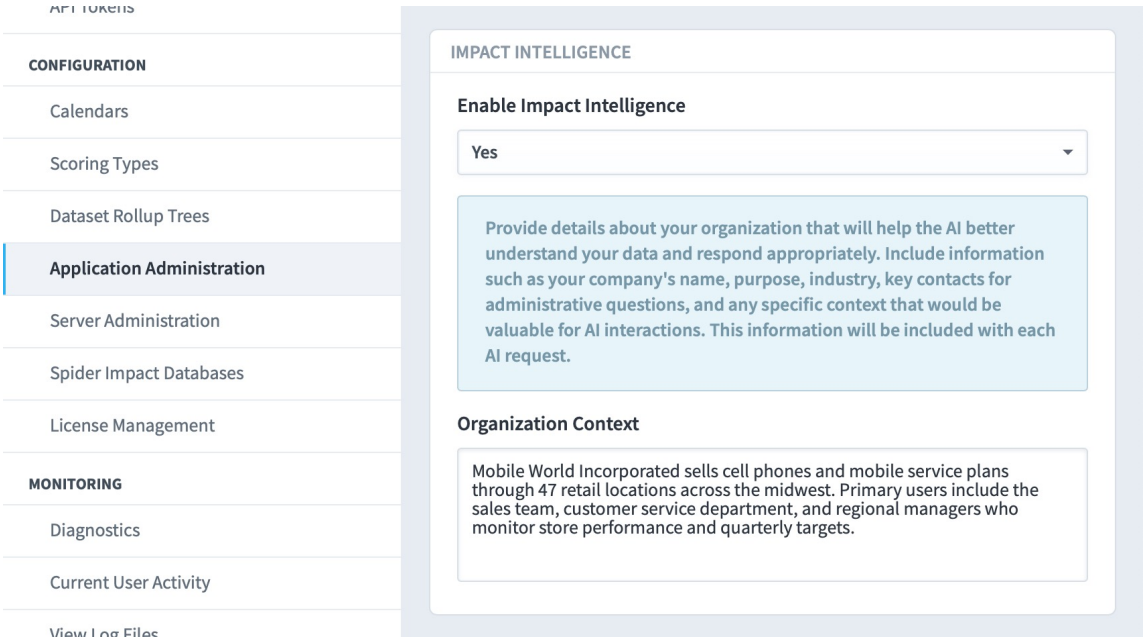

The Organization Context helps the AI understand your data and provide better responses. Include your company name, industry, the types of data you collect, and any big-picture business goals that are universal and don't change over time. If the AI has trouble differentiating between datasets or finding specific data, additional context here often solves the problem.

Dataset Descriptions

To help Impact Assistant better understand your data context, administrators can add descriptions to datasets. These descriptions help Impact Assistant provide more relevant insights and appropriate visualizations:

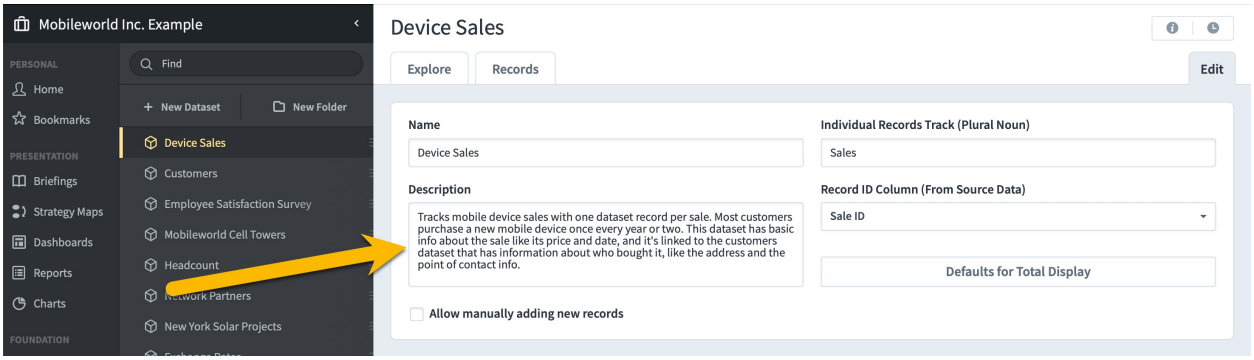

When writing dataset descriptions, explain the dataset as you would to a colleague. Include:

- What each record represents
- Types of information tracked for each record
- Why you're tracking this data
- Any important relationships to other datasets

Example description:

Tracks mobile device sales with one dataset record per sale. Most customers purchase a new mobile device once every few years. This dataset has basic info about the sale like its price and date, and it's linked to the customers dataset that has information about who bought it, like the address and the point of contact info.

Good descriptions solve common confusion. For example, mentioning that a dataset tracks both "incident date" and "report date" helps the AI distinguish between them. Similarly, clarifying that "SARC Location" refers to where something happened helps with location-based queries.

API Keys for Self-Hosted Customers

Spider-hosted customers don't need to manage API keys—Spider Strategies handles all billing and infrastructure.

Self-hosted customers need to provide AWS credentials:

1. Create an AWS account and enable Amazon Bedrock
2. Generate access keys for Bedrock
3. Enter these keys when enabling Impact Intelligence

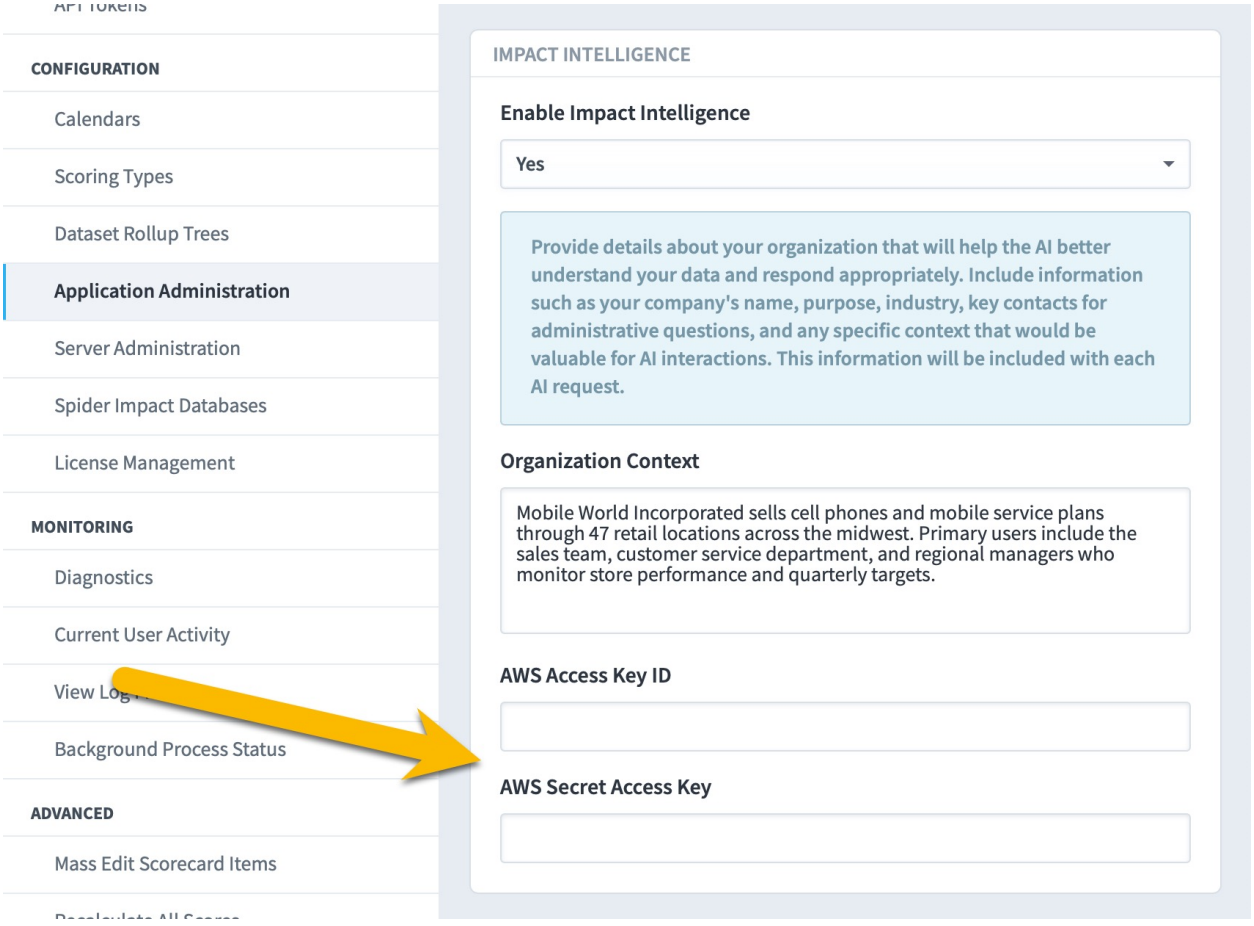

Alternatively, configure API keys at the server level by adding these JVM parameters to Tomcat:

```
-Dspring.ai.model.embedding=bedrock-titan
-Dspring.ai.bedrock.titan.embedding.enabled=true
-Dbedrock.access-key=YOUR_ACCESS_KEY_HERE
-Dbedrock.secret-key=YOUR_SECRET_KEY_HERE
```

With server-level configuration, the API key fields won't appear in the interface—you'll just toggle Impact Intelligence on or off for each instance.

The Future of Impact Intelligence

Impact Assistant represents just the beginning of Impact Intelligence capabilities. In upcoming releases, you can expect rapid expansion of AI-powered features throughout Spider Impact, including enhanced data analysis capabilities, intelligent automation of routine tasks, and predictive insights that help you make better decisions faster. Impact Assistant transforms how users interact with their data today, while laying the foundation for even more powerful capabilities tomorrow.

Apps

We're excited to introduce a powerful new feature in Spider Impact: **Apps**. This enhancement allows administrators to build curated, branded experiences for specific tasks, enabling streamlined workflows and customized interfaces. With Apps, organizations can consolidate data collection, reporting, and operational tools into a single, integrated environment.

Key Benefits of Apps:

- **Task-Specific Experiences:** Organizations can create tailored applications for different departments or use cases, such as sales tracking, customer service, or project management.
- **Centralized Data and User Management:** Apps leverage Spider Impact's existing permissions architecture, ensuring secure and seamless data access.
- **Standalone Access:** Apps can be accessed independently of the main Impact UI, allowing users to interact with them via a dedicated URL.

Viewing and Accessing Apps

Unlike traditional Impact sections, Apps provide a standalone experience. When users navigate to an App's view tab, they will see a single "Open App" button that launches the App in a new browser tab. Each App has its own dedicated URL, allowing users to create browser bookmarks for direct access.

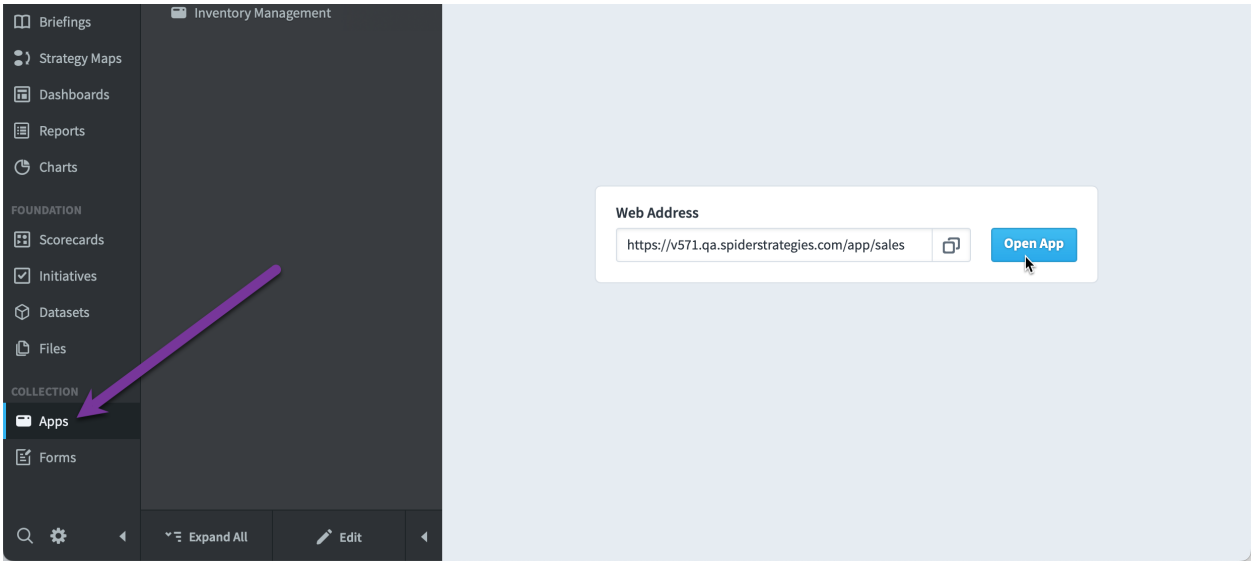

Apps require users to log in, and Apps share a security model with Spider Impact. You can only see data in Apps that you can already see in Impact.

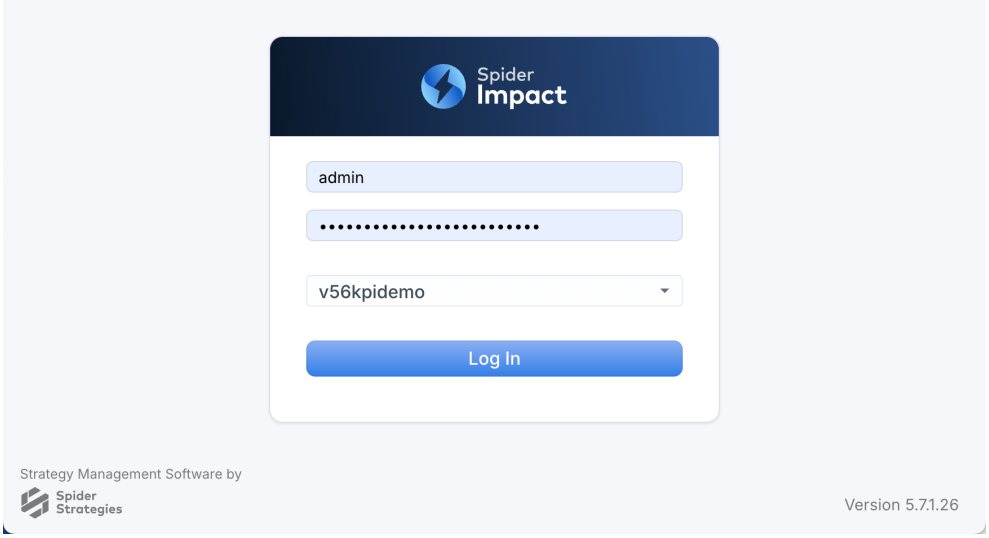

After logging in, users will encounter a customized interface designed specifically for its purpose. Apps are built from existing Impact screens like Forms, Dashboards, and Reports. For example, this App has four sections. This "All Customers" section is a multi-page form that starts with a list of all customers.

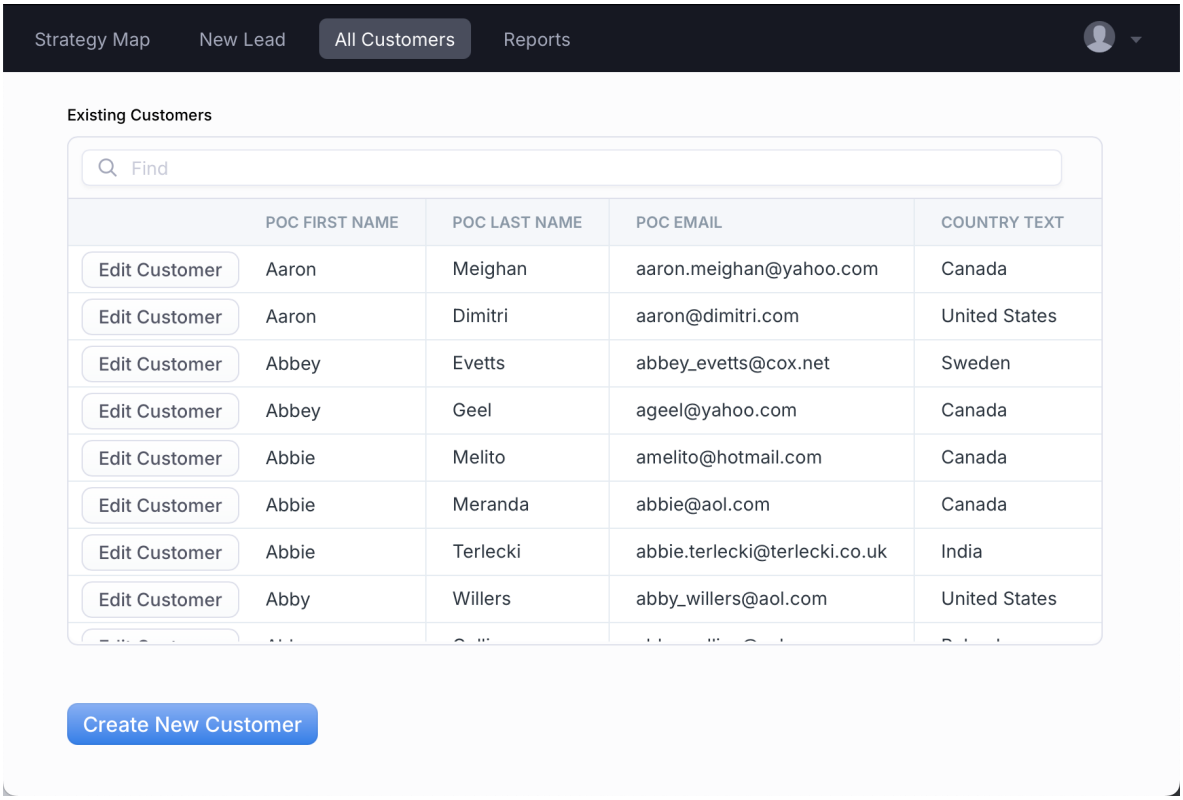

When we click the “Edit Customer” link for a customer, we’re taken to the customer details page where you can see a history of sales for that customer, and you can edit the customer’s contact information.

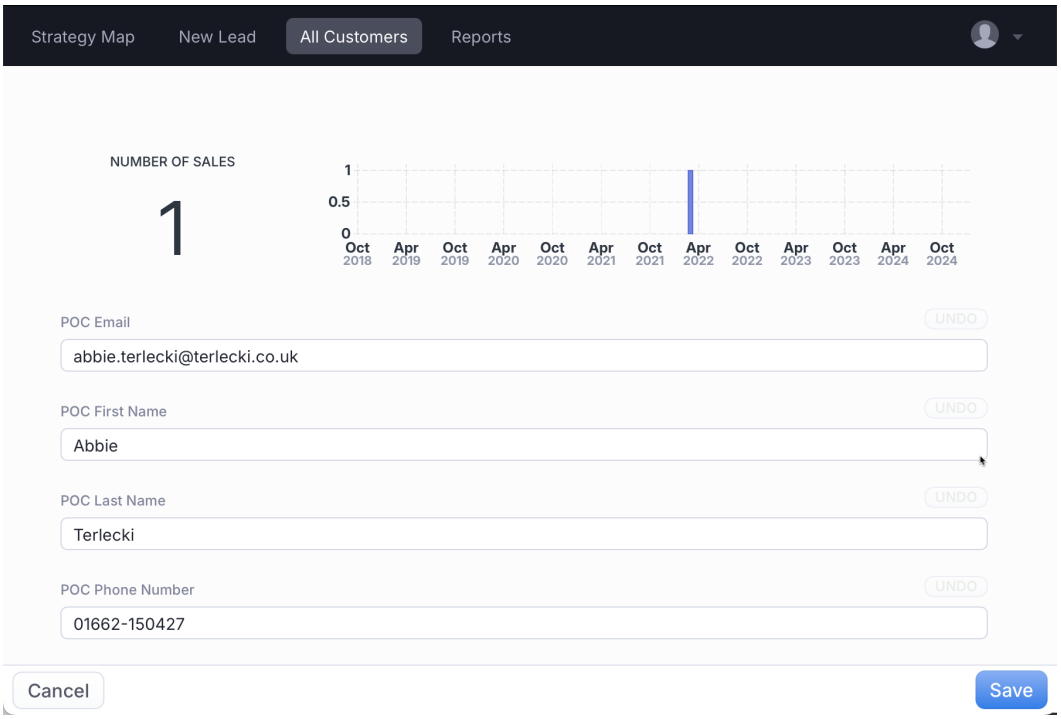

In this example, the “Reports” section allows you to choose between several dashboards, each displaying various sales data visualizations.

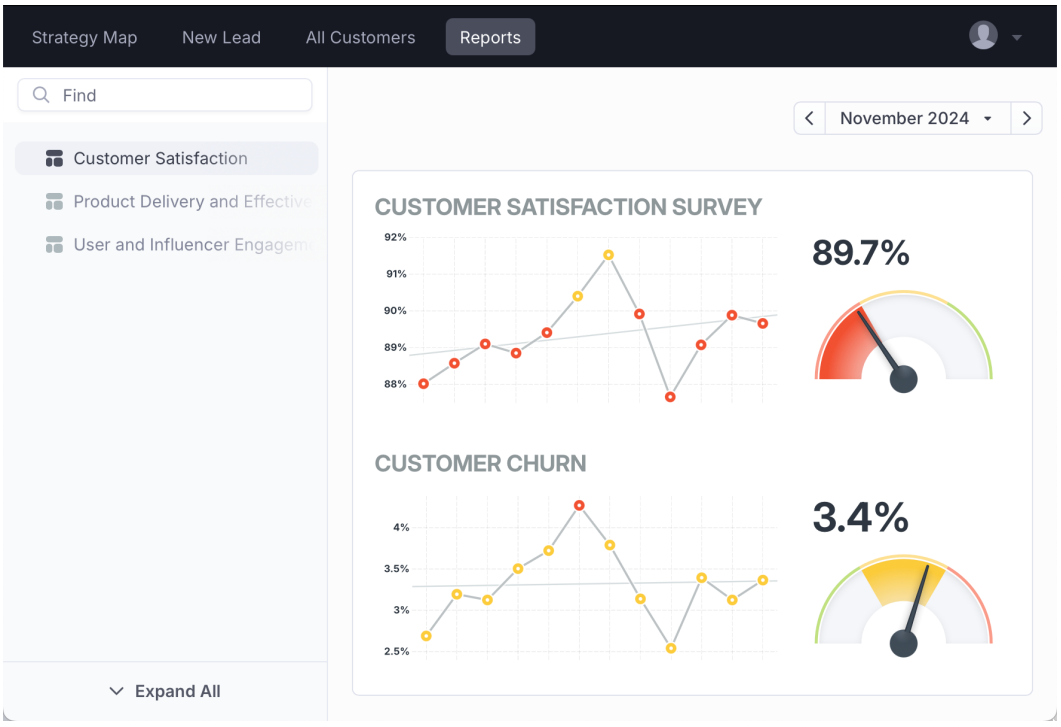

Creating and Managing Apps

Like every section in Spider Impact, Apps reside in organizations and follow the standard Spider Impact permission model. That means users can access every app in their viewable organizations. There is also a new "Modify Apps" permission that allows administrators to configure and manage Apps in their viewable organizations.

The Edit tab within the Apps section allows administrators to manage App settings. Here, the app itself is selected, and we can set the app's name, URL, and description.

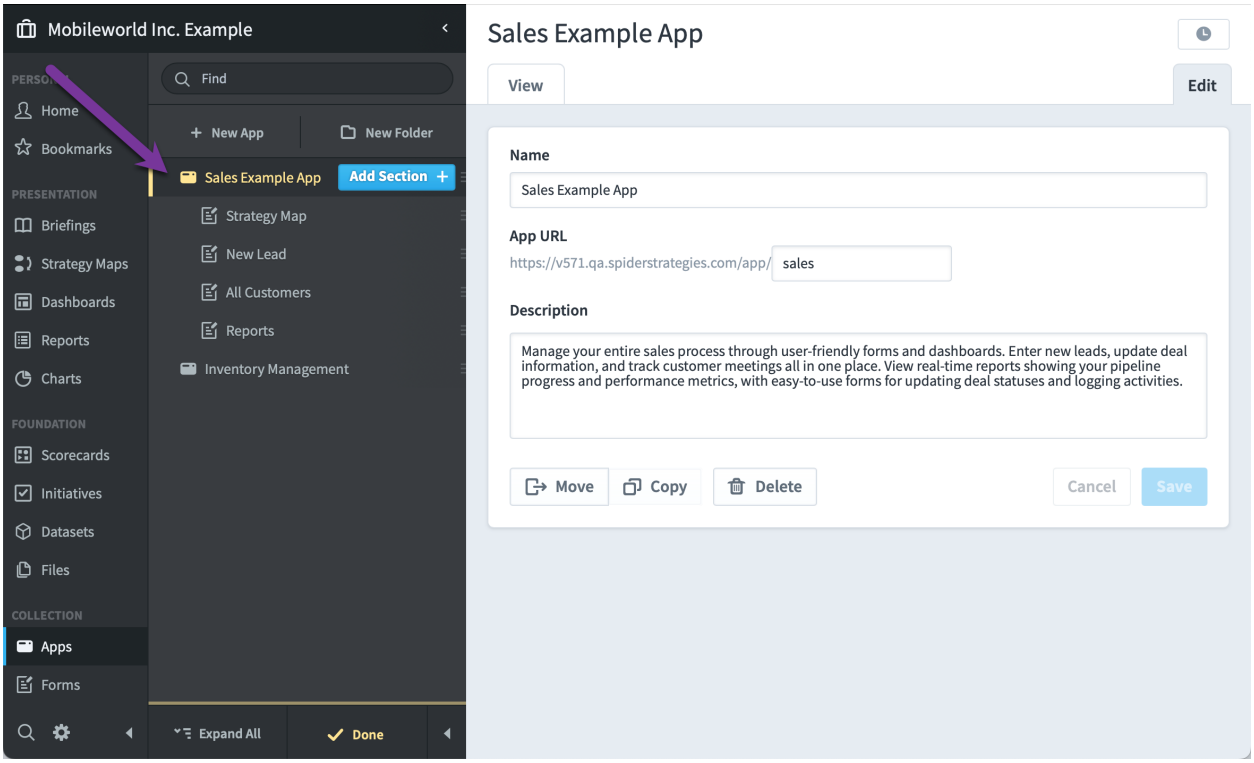

Like editing Form Pages, when you edit an App, all its sections appear underneath it. Here we've selected the "All Customers" section and we can see that this App Section shows a single Form.

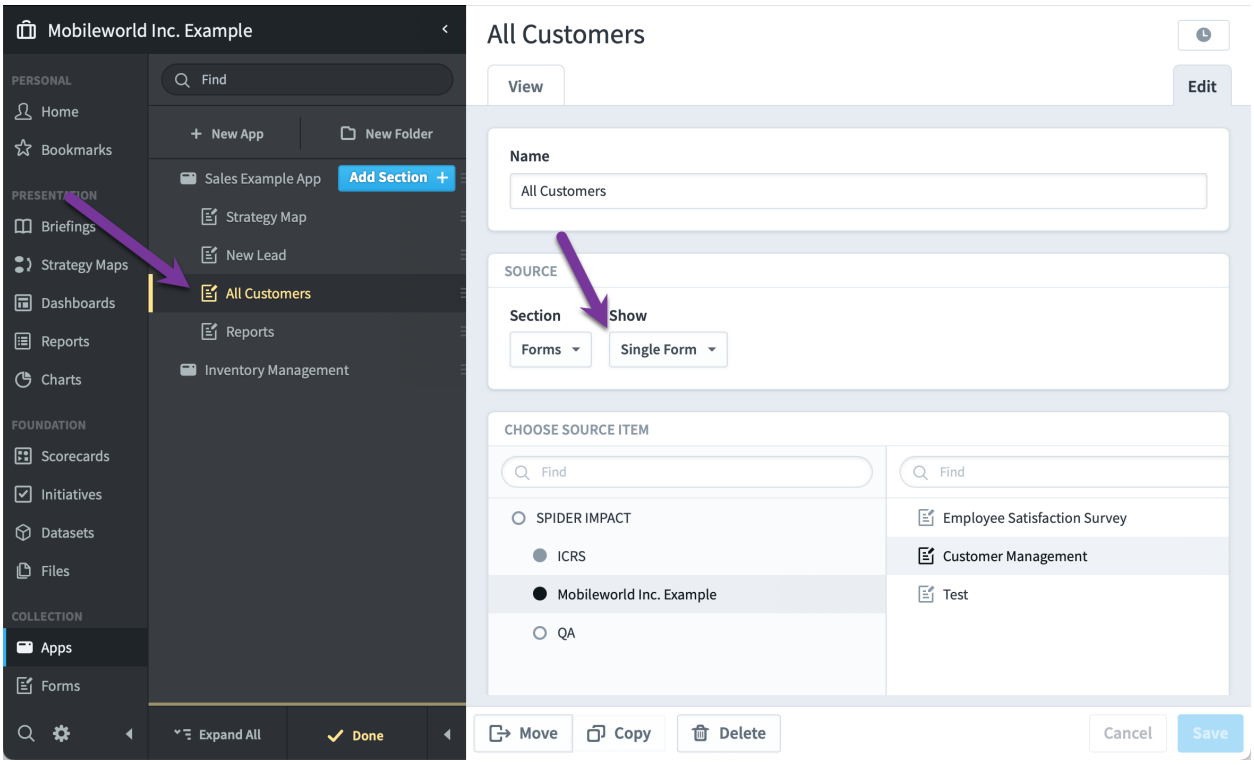

When we select the "Reports" section, we can see that it's showing every dashboard from an organization.

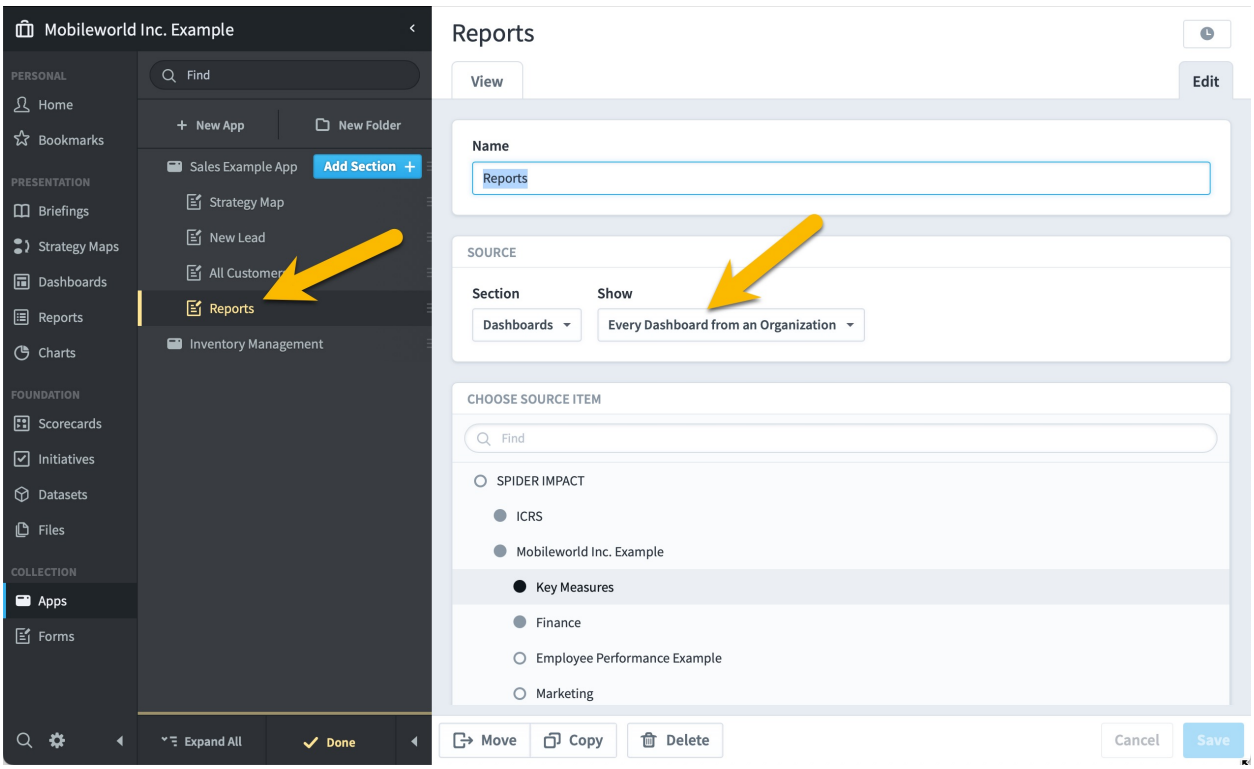

And, because Apps use Spider Impact's existing permissions, if a user can't see a page inside of Spider Impact and that page is being used for an App's section, that section will be hidden for them inside of the App. This allows you to create apps with sections that only appear for some users.

Licensing

All customers who currently have access to Forms in Spider Impact will also have access to Apps at no additional charge. Forms is enabled automatically when you have 25 or more licenses and is also available as an add-on for customers with fewer than 25 licenses.

Charts

Bullseye charts

Spider Impact's Bullseye Chart visualizes how different measurements compare to their baselines. Each chart "slice" represents a different measurement, with colored threshold regions showing performance zones, and data points showing actual values relative to those thresholds.

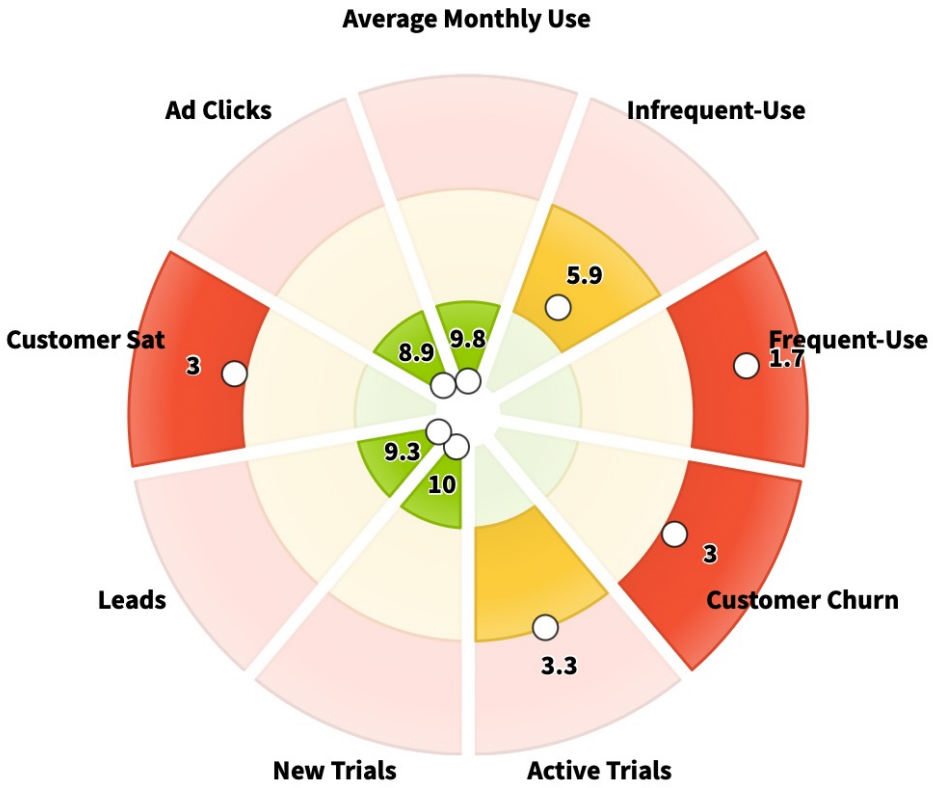

Shared Thresholds Mode

By default, all positions around the bullseye chart share the same scale and thresholds. This works well when showing scorecard items because they're already normalized to 0-10 scores, or showing dataset values that have the same goals, like sales figures for various employees.

In this mode, you select "Set Data Source" on a series like you do on most other charts.

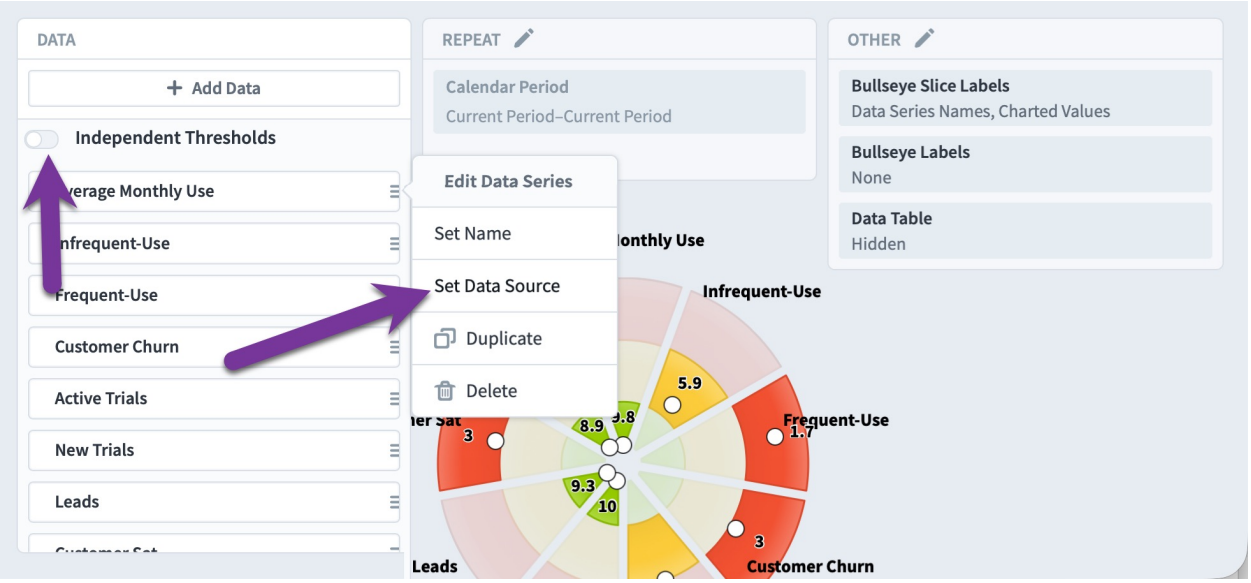

And the Set Data Source menu is the same as other charts as well. You can choose data from Scorecards, Initiatives, or Datasets.

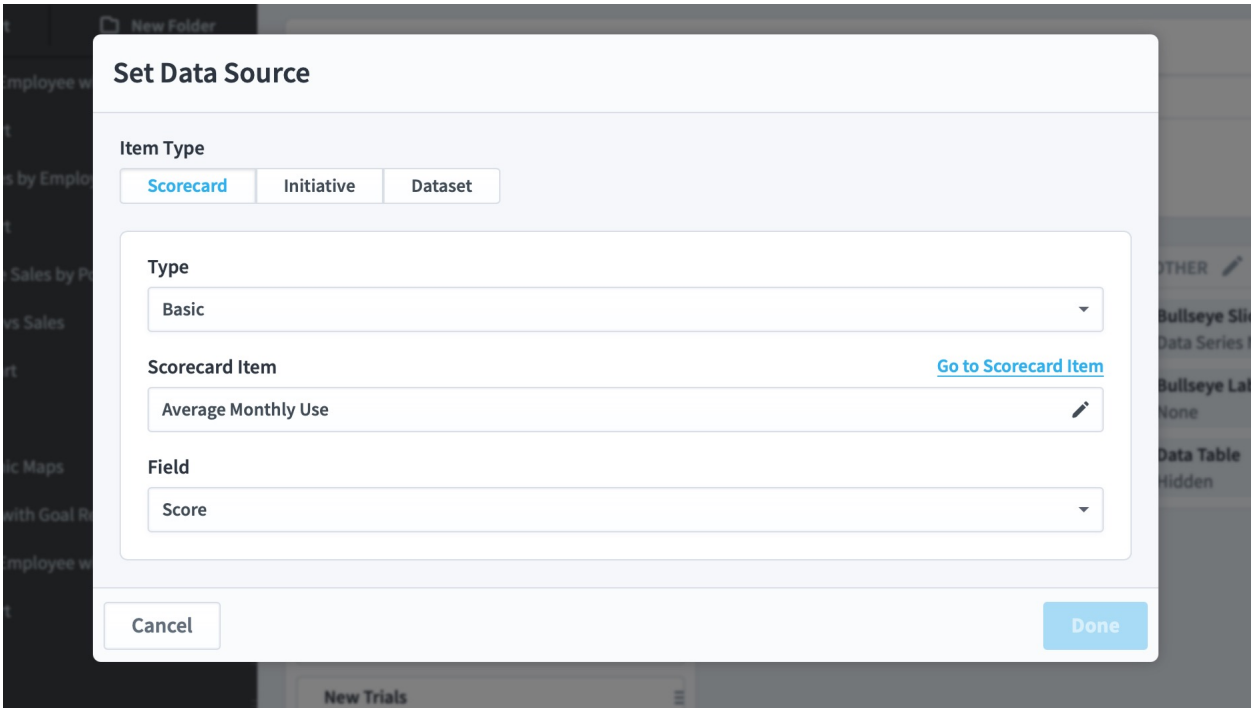

Reference lines and bands

Like other types of charts, bullseye charts have reference lines and bands. When Independent Thresholds is off, you set lines to constant values.

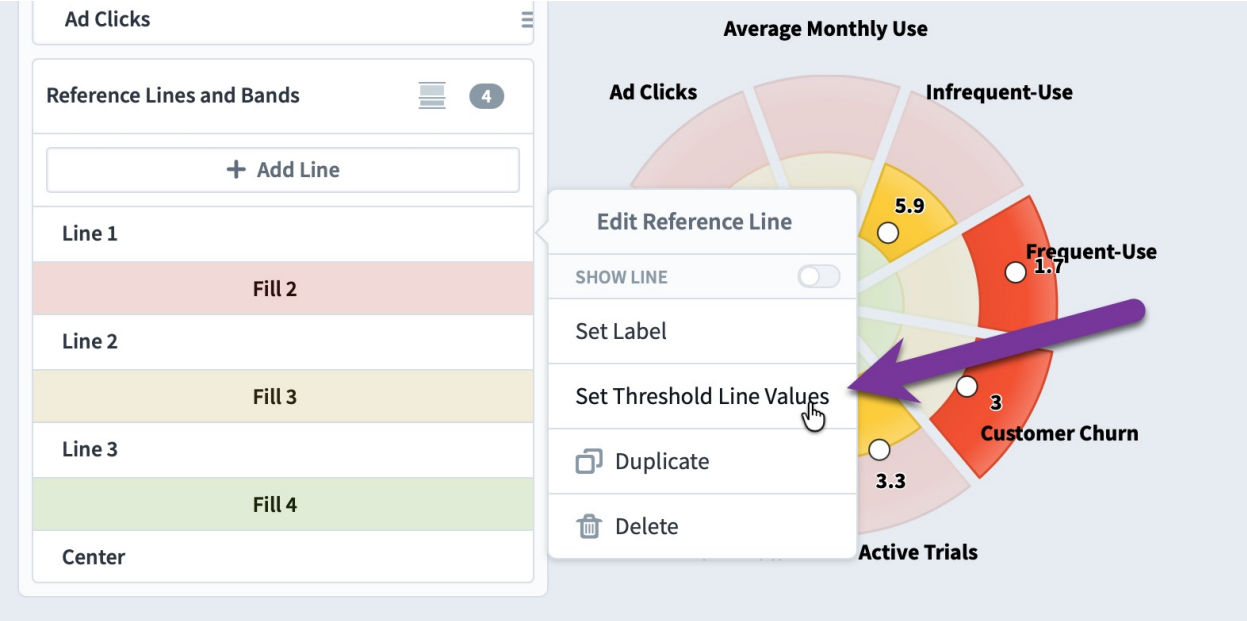

In this example we're comparing scorecard item scores, so we'll choose to break up the 10-point scale into 3 equal pieces.

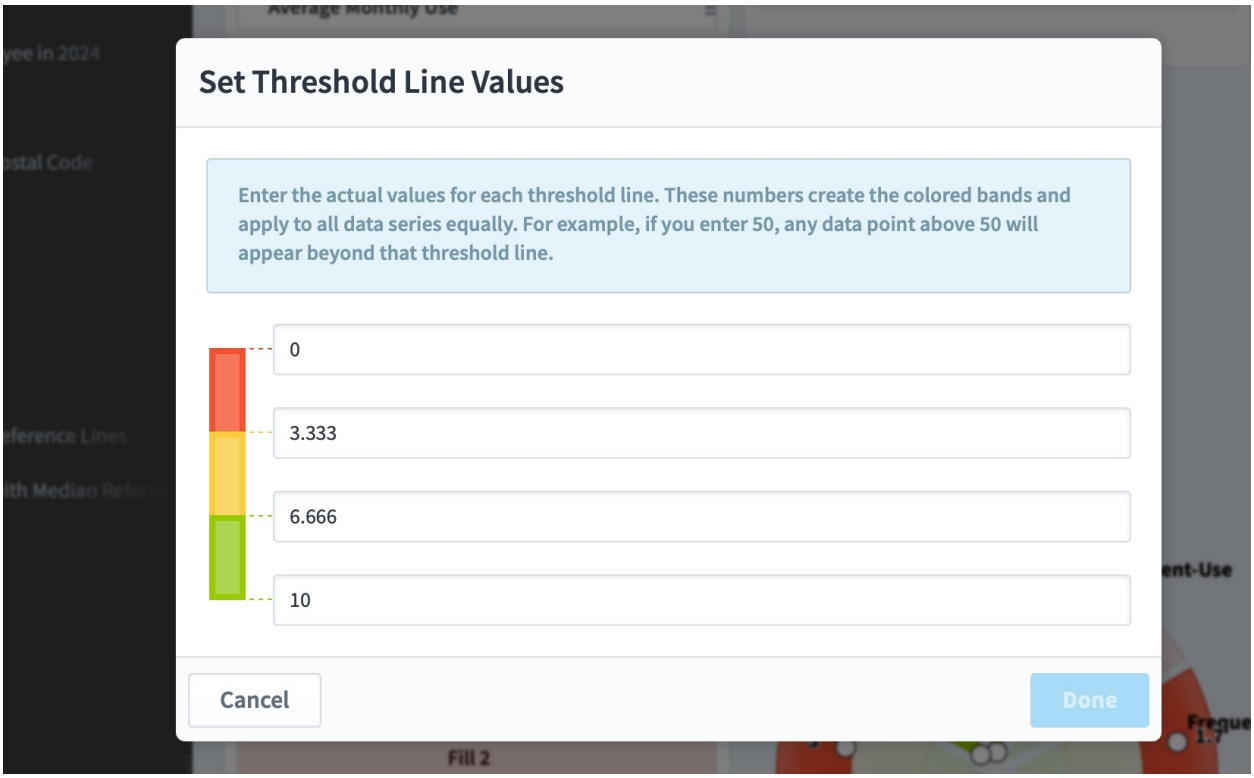

Just like with line, bar, and area charts, you can modify the reference lines and bands. You can choose each ring's color, and you can set the opacity for inactive segments.

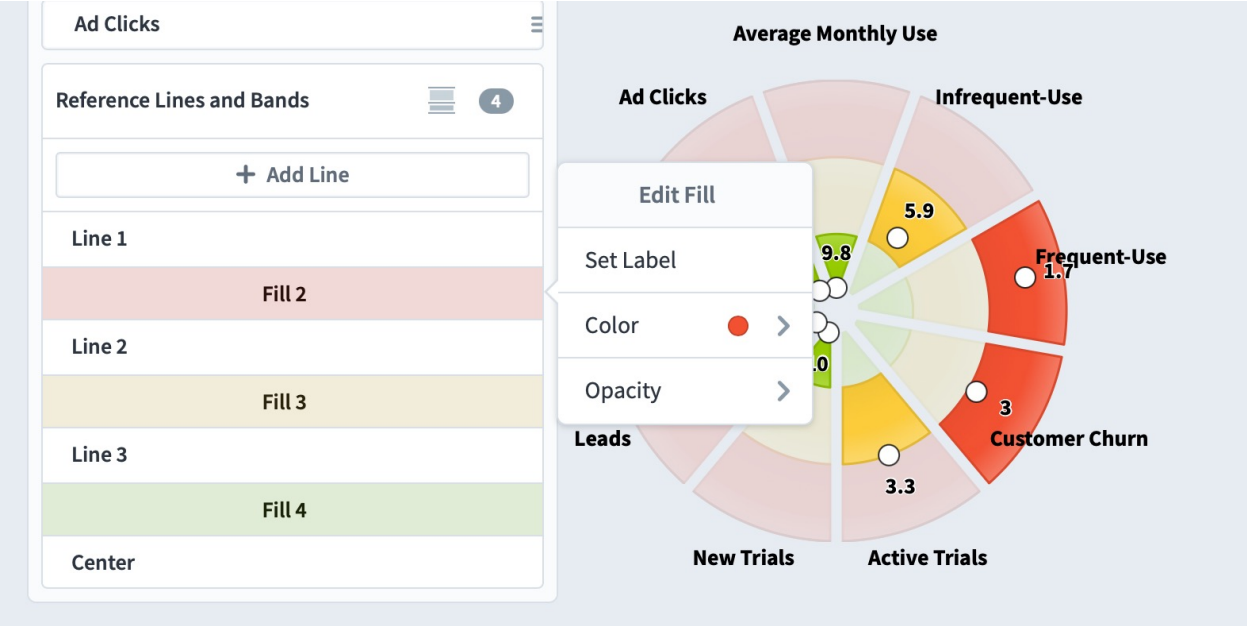

Reference lines are off by default, but you can choose to turn them on and set a color and opacity to draw attention to a special threshold.

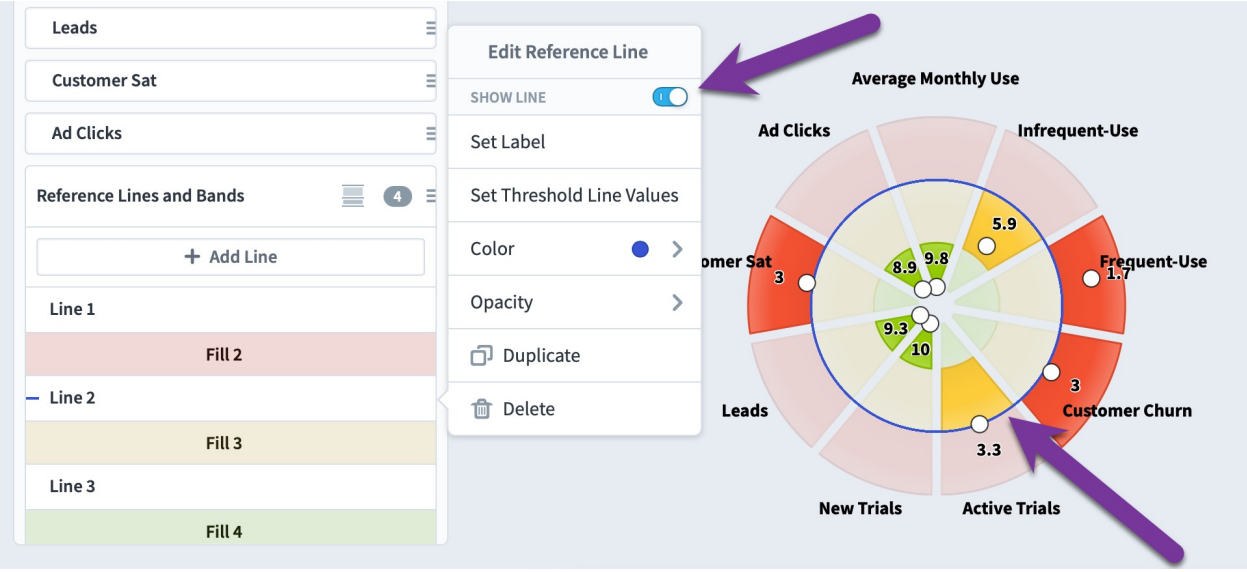

Independent Thresholds Mode

When comparing measurements with different scales – like workplace injuries (per 1,000 employees) versus sales revenue (in dollars) – enable "Independent

Thresholds" mode. This allows each “slice” to have its own scale while maintaining visual comparability.

When you enable Independent Thresholds, the series menu changes to display:

- **Set Point Data Source:** The data to plot (e.g., a specific employee’s sales)
- **Set Threshold Data Source:** The baseline for comparison (e.g., company-wide averages)

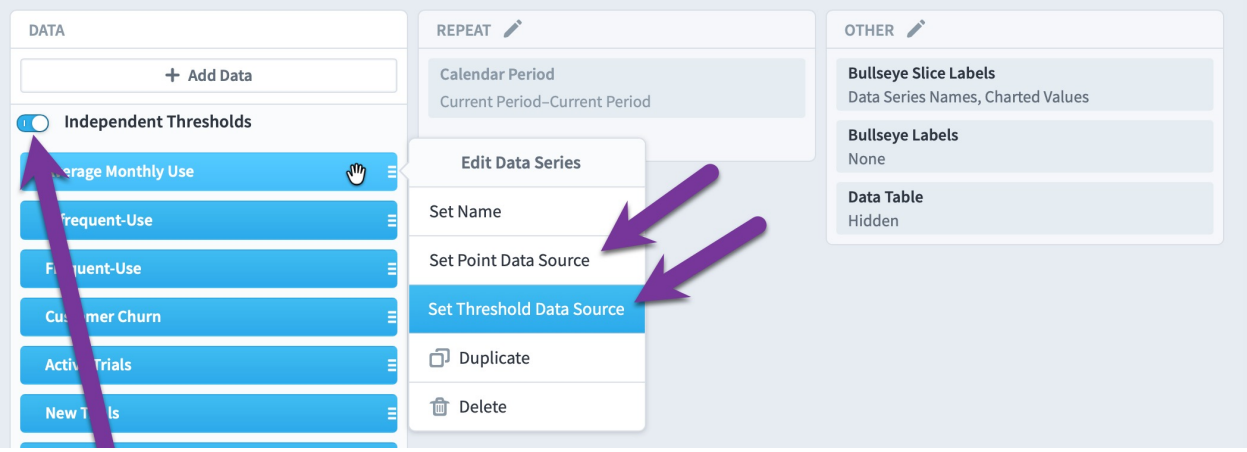

A typical Independent Thresholds bullseye chart uses the same dataset for both point and threshold data, with the point data having additional filters. For example:

- **Threshold Data Source:** Sales data averaged for all salespeople
- **Point Data Source:** Sales data filtered to "Salesperson = Russell Corrick”

This configuration shows how Russell performance compares to the overall average across all metrics.

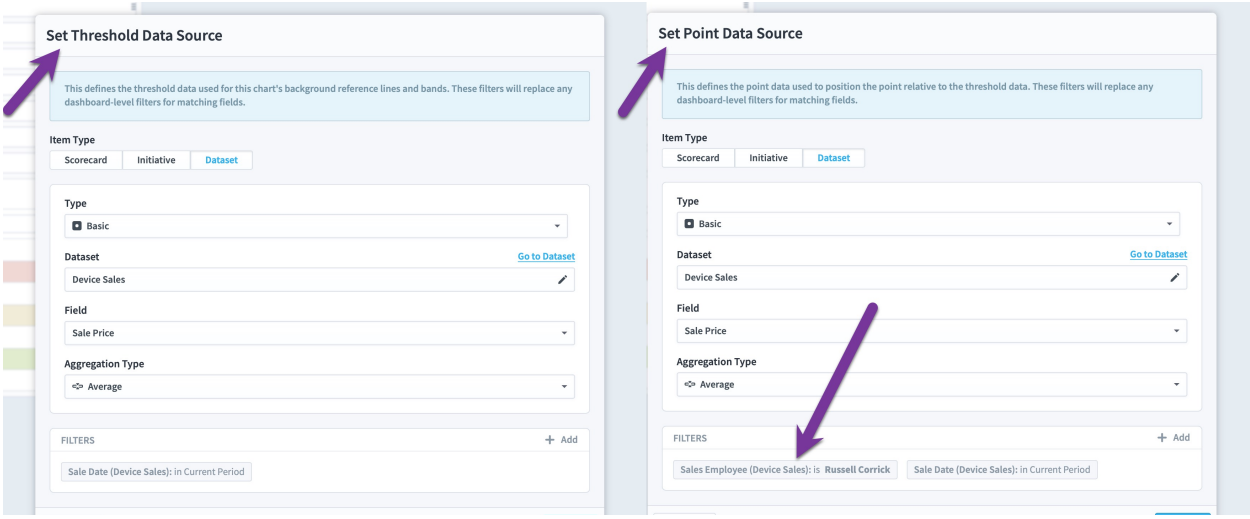

With Independent Thresholds enabled, the threshold lines change to be a percentage of each slice's Threshold Data Source value. For example, setting a line at 200% creates a threshold at twice the Threshold Data Source values. If your baseline for injuries is 5 per 1,000 employees, that line would appear at 10 injuries per 1,000. In this example the green segment starts at 100% of the threshold data source value.

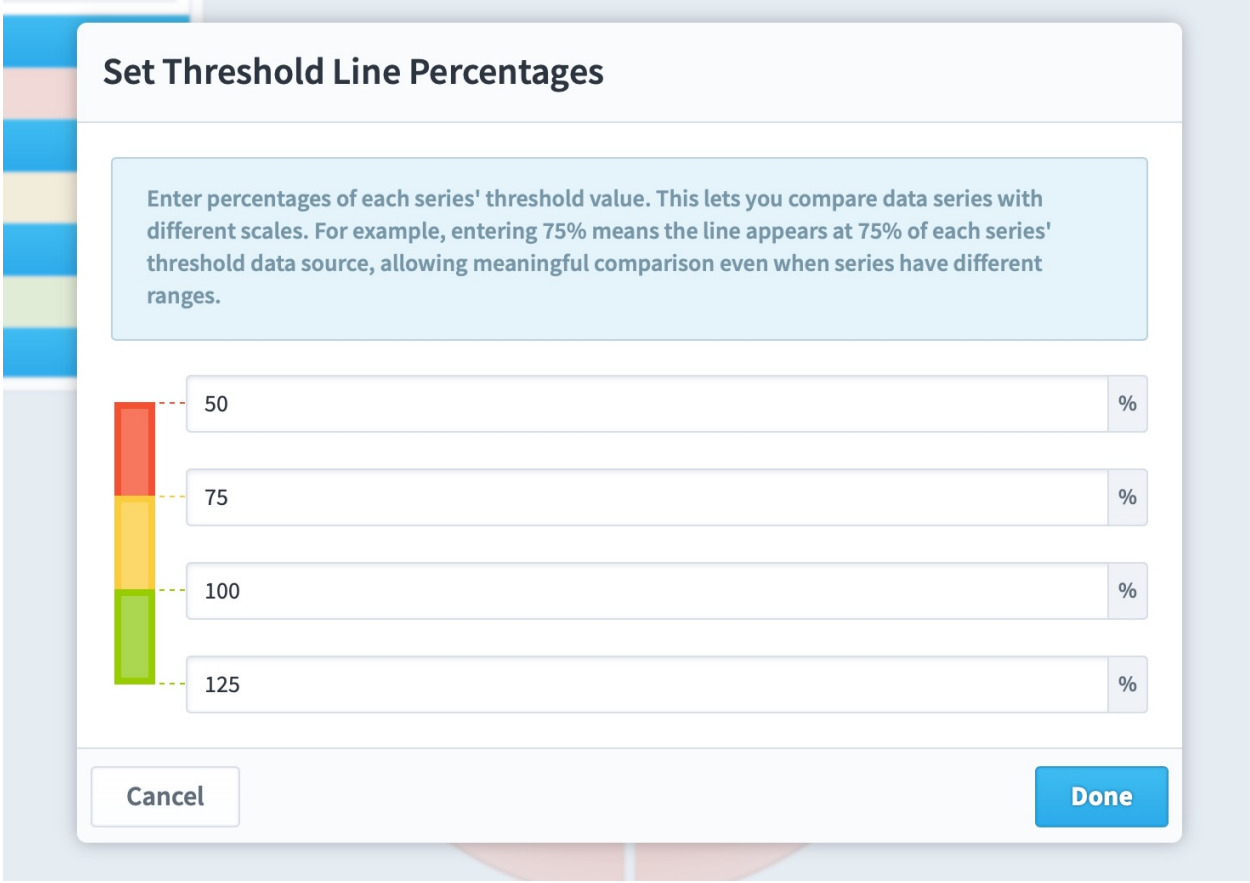

Let's further explore independent threshold mode with this example. The bullseye chart on the left shares the same thresholds across all slices. It compares each employee's average sale size against the average for all employees. Two employees jump out with massively larger sales than the other employees. That, however, is because Kym and Russell do corporate sales, and their deal sizes are always going to be larger than the employees doing retail sales. In this example, the average sale prices for the different employees have different acceptable scales, making this kind of comparison not very helpful.

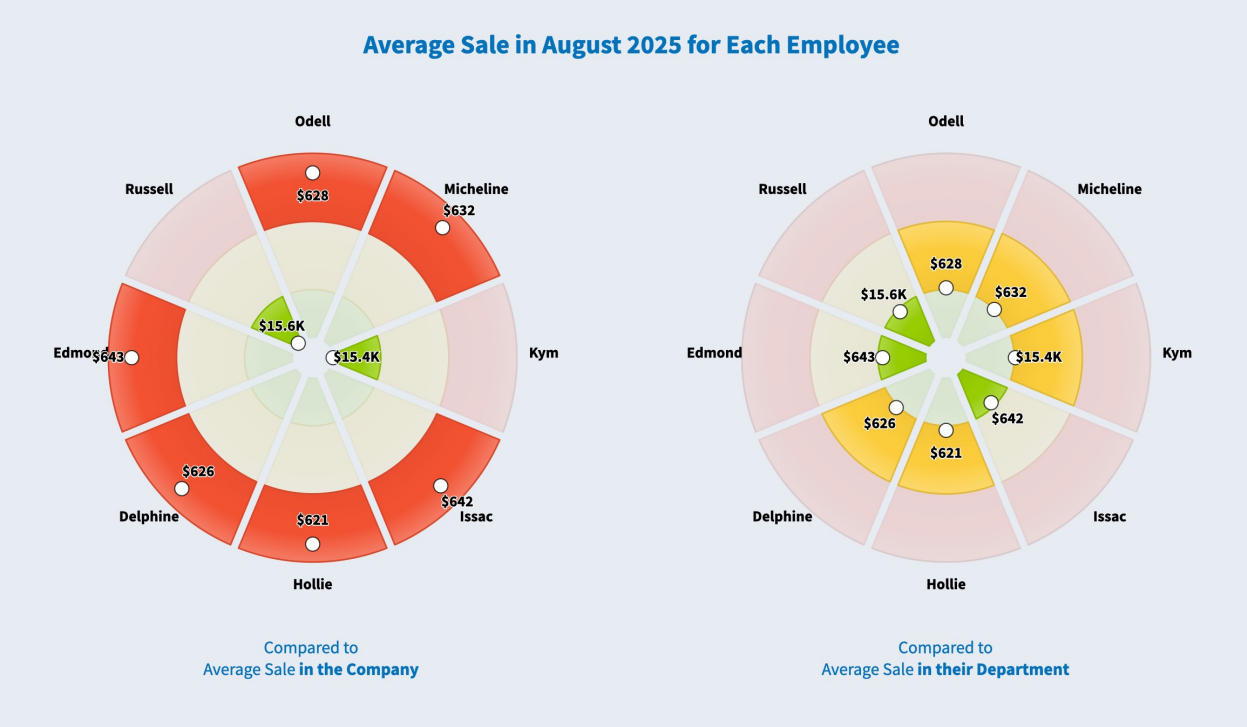

The bullseye chart on the right has independent thresholds turned on, and instead compares each employee against only the other employees in their department. We can instantly see that Russell, Edmond, and Issac are having good months, regardless of their department.

Repeating bullseye charts

Bullseye charts support repeating values to create small multiples, displaying the same set of metrics across multiple entities simultaneously. In this example, we're showing the same sales data for United States, Australia, and United Kingdom.

Bullseye Charts on Dashboards

When you add a dataset filter to a dashboard, it applies to both the threshold and point data in your bullseye charts. This allows you to explore dataset data as you would with any other chart. For example, here we're looking at each employee's average sale price this month, and each slice is using a scale of that employee's average sales for the current year. It allows you to see if the employee is having an above or below average month relative to their past performance.

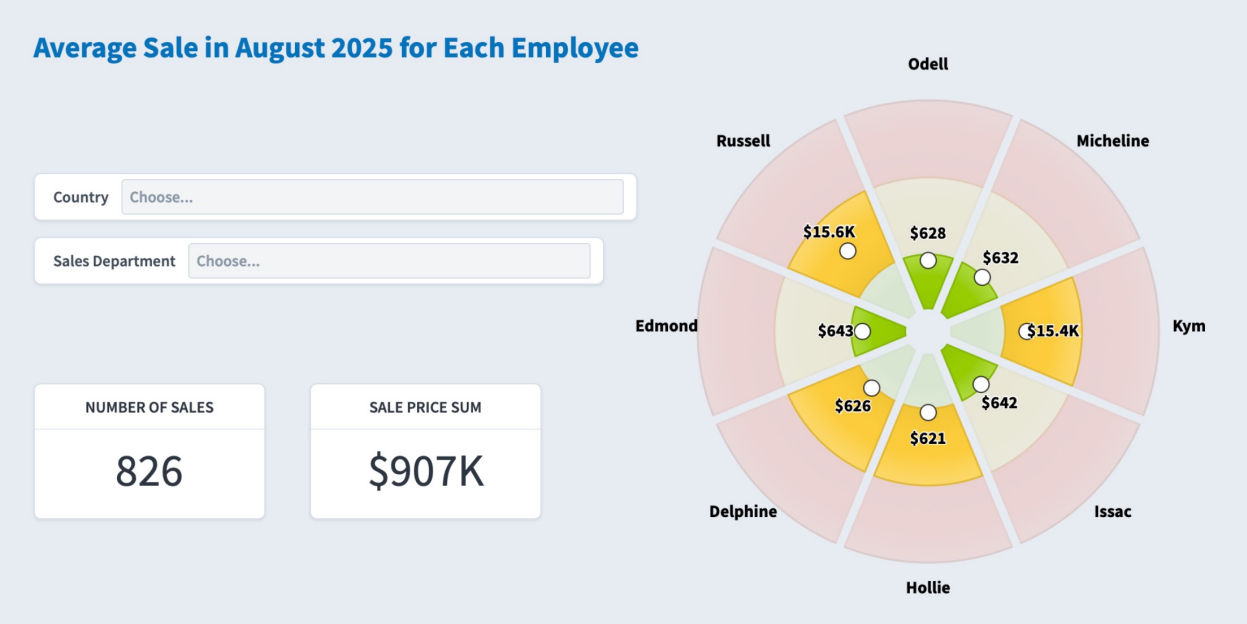

When we apply a filter for Sales Department, we now see the exact same chart filtered to only show only show Corporate sales. Here it's showing only Russell and Kym had Corporate sales this month.

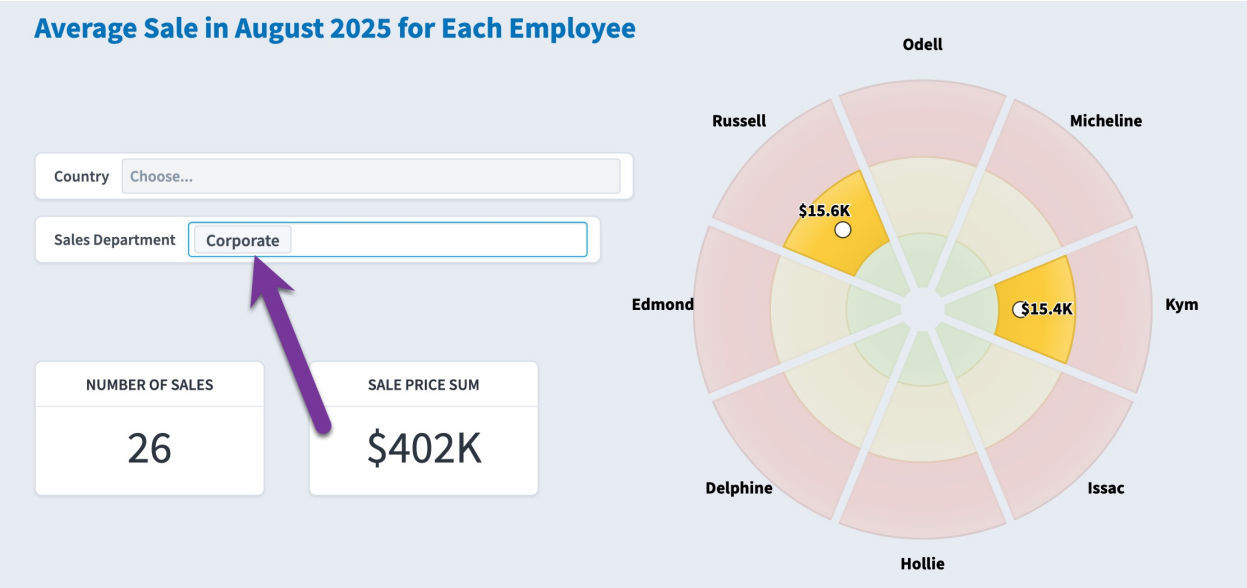

Similarly, we can apply a filter to North America, and we now see the exact same chart filtered to only show North American data. Note that filters apply to both Point and Threshold data by default. That means this chart is comparing each salesperson's average *North American* sale this month against their average *North American* sale this year.

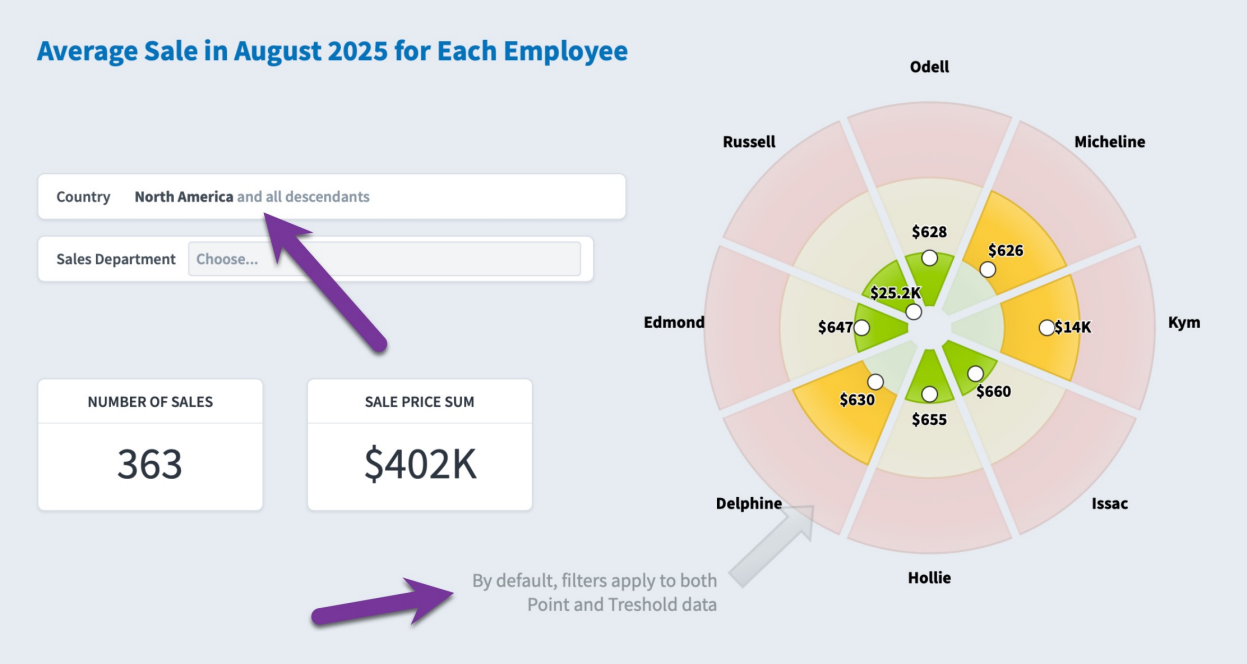

If you only want the Point or Threshold data to ignore a dashboard filter, you can just apply a filter for that field to the data source. That's because any dashboard filters that conflict with point or threshold data will be ignored. Here we're applying a filter to the Threshold Data Source saying that the country is anywhere on earth.

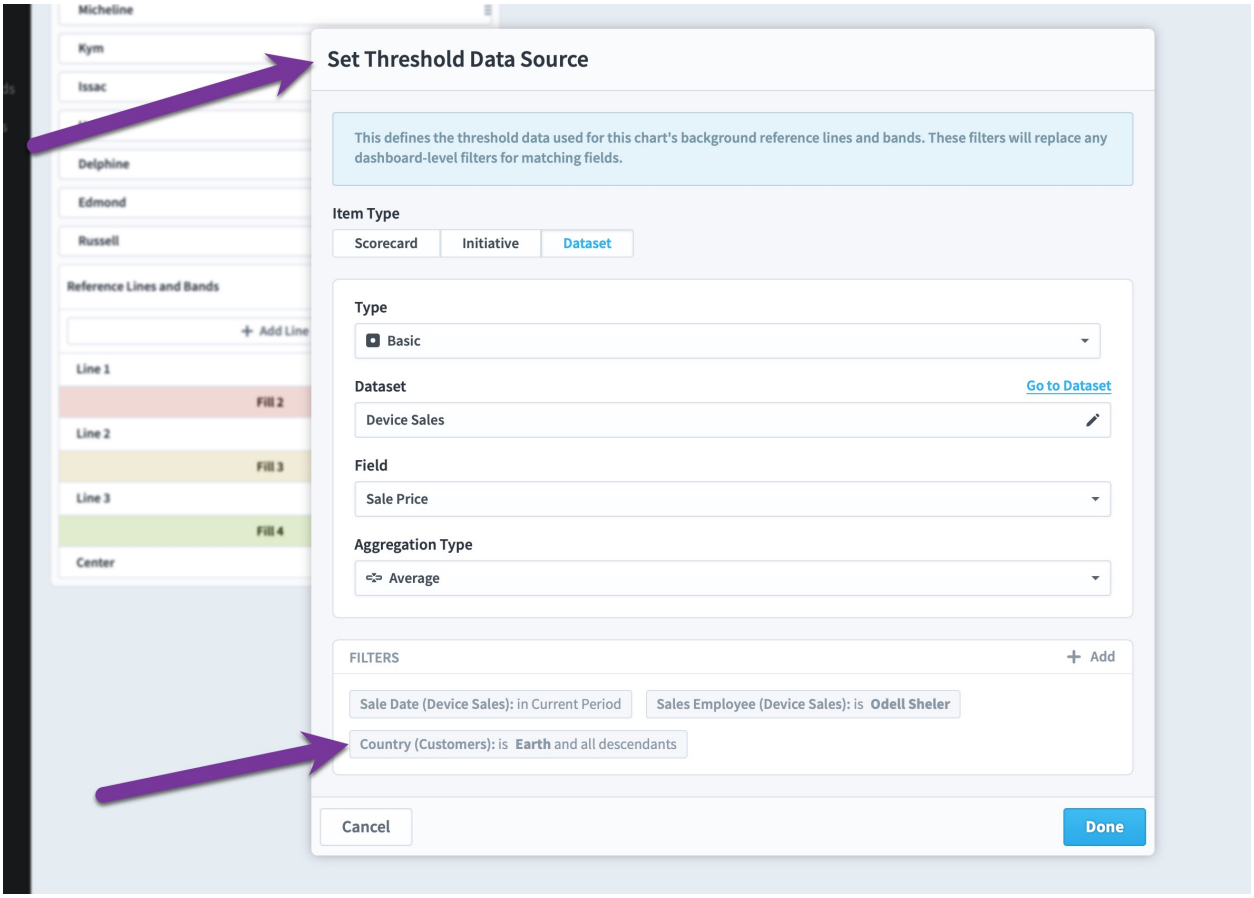

Now when we set the dashboard to filter on country, that filter only applies to the points. This chart is comparing each salesperson's average *North American* sale this month against their average sale this year *regardless of country*. Now adding a filter changes the points without changing the background thresholds.

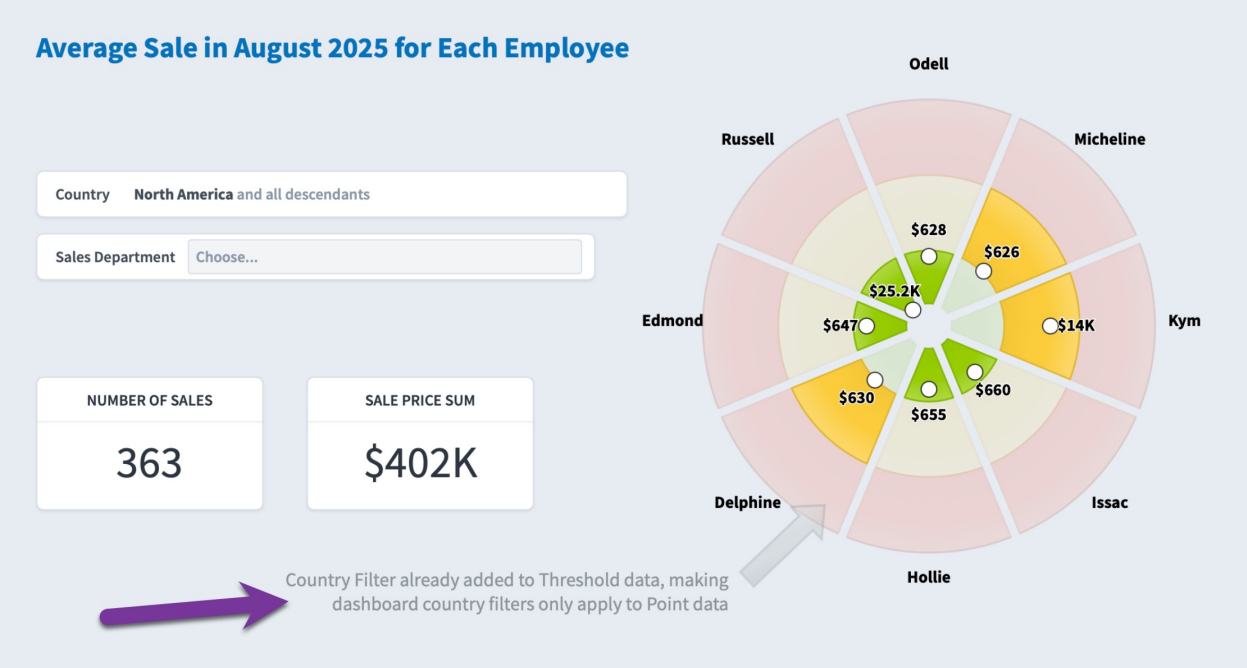

Having dashboard filters apply to points, thresholds, or both are all completely valid use cases, and the new Bullseye charts allow you to do them all.

Display values on treemap charts

Treemap charts now display numeric values directly on each section, eliminating the need to hover to see data. This enhancement improves data analysis efficiency, especially for hierarchical visualizations where users need to compare multiple data points at a glance.

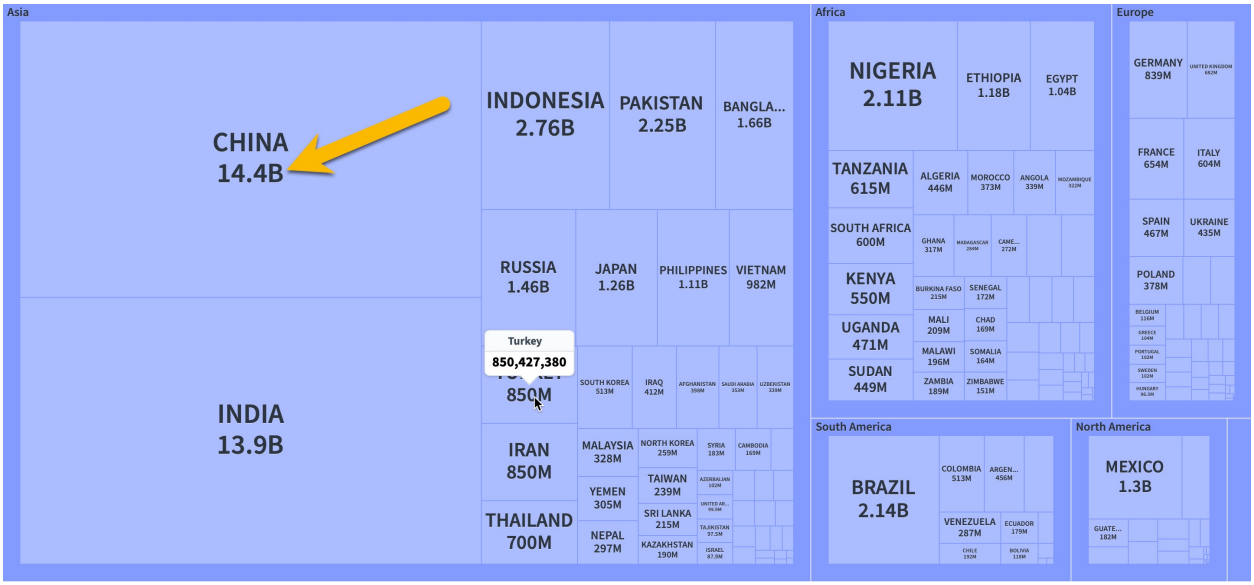

When editing a treemap, you can control value display through the Edit Treemap Labels dialog, with options to show Data Series Names and Charted Values, including abbreviation settings for large numbers.

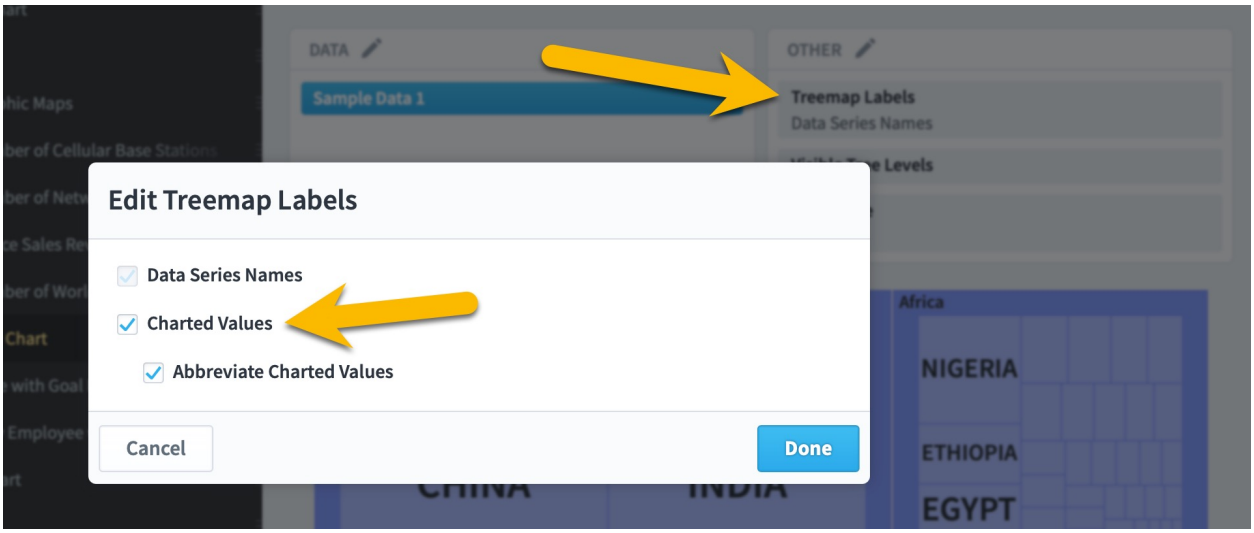

Configurable chart gridlines and labels

You can now independently toggle X and Y axis gridlines and labels through the Edit Axes dialog. This configuration is available for all chart types including polar charts. In this example we've turned Y axis labels and gridlines so all we see in vertical gridlines and their labels.

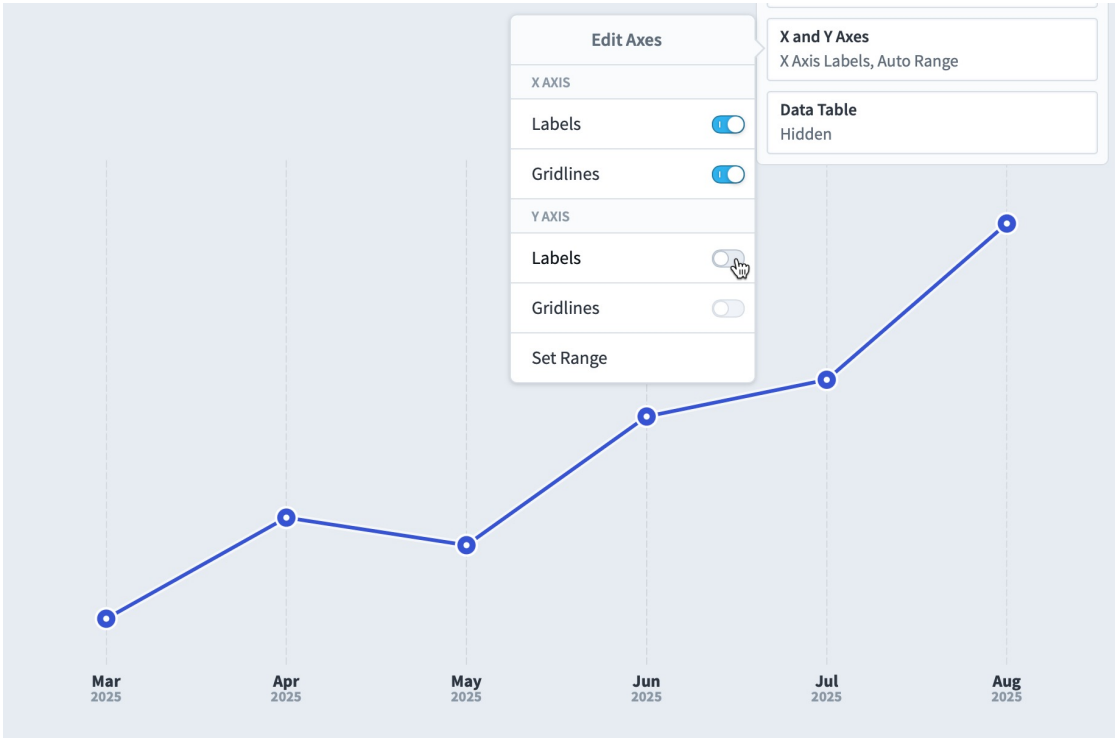

New charts default to having Y axis gridlines off for a cleaner look. Existing charts retain their current settings to maintain consistency. Charts with calendar period breakdowns automatically display fewer X-axis labels, reducing visual clutter when viewing time-based data.

Additionally, the Scorecards Overview tab now includes a "Chart X Axis Gridlines" toggle in Application Administration, allowing you to control gridlines for all default scorecard charts.

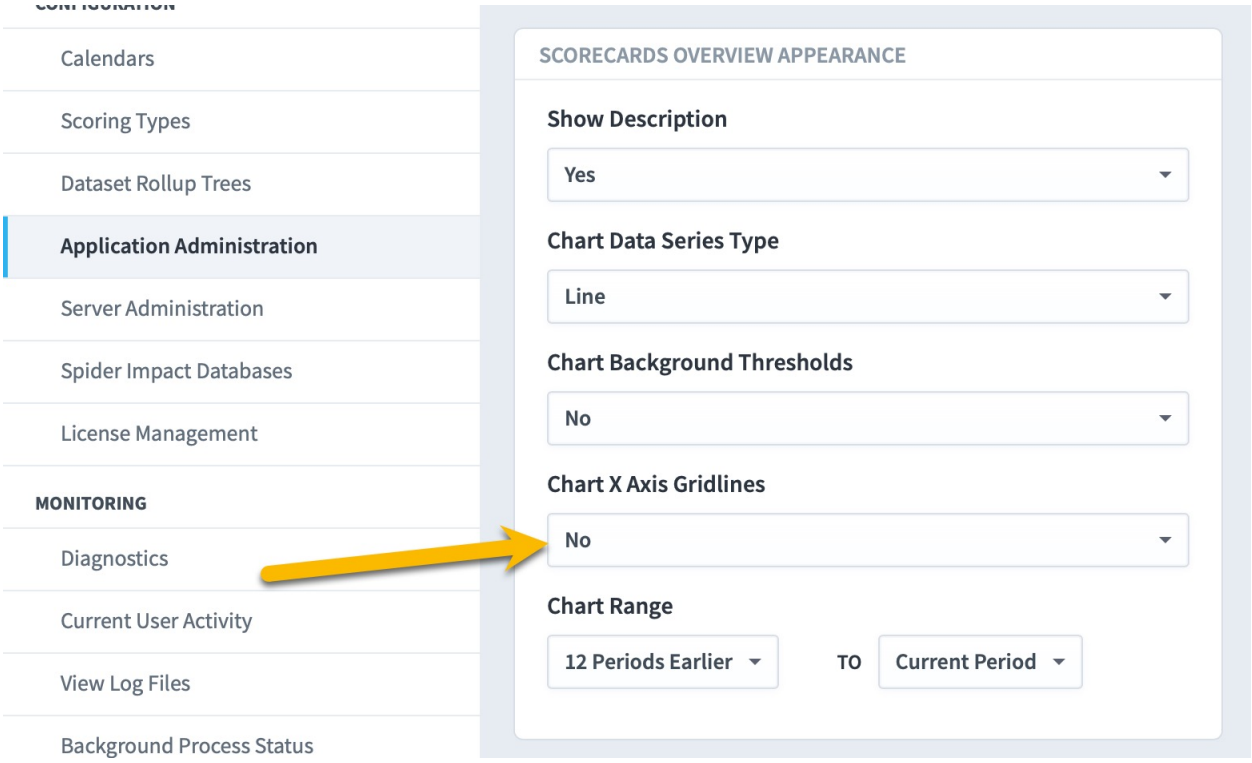

Support for geographic and rollup field breakdowns

Charts can now repeat by Geographic and Rollup field types, removing previous restrictions that limited these data aggregation options. This enhancement

enables more sophisticated data visualizations, particularly useful when analyzing hierarchical or location-based datasets.

Ignore incomplete time periods for accurate trend analysis

When viewing time-based charts, the current period often shows a misleading drop because it's not complete yet. For example, if you're looking at customer cancellations at the beginning of September, that month appears to have unusually low cancellations simply because most of the data is missing. This is particularly problematic when display projections.

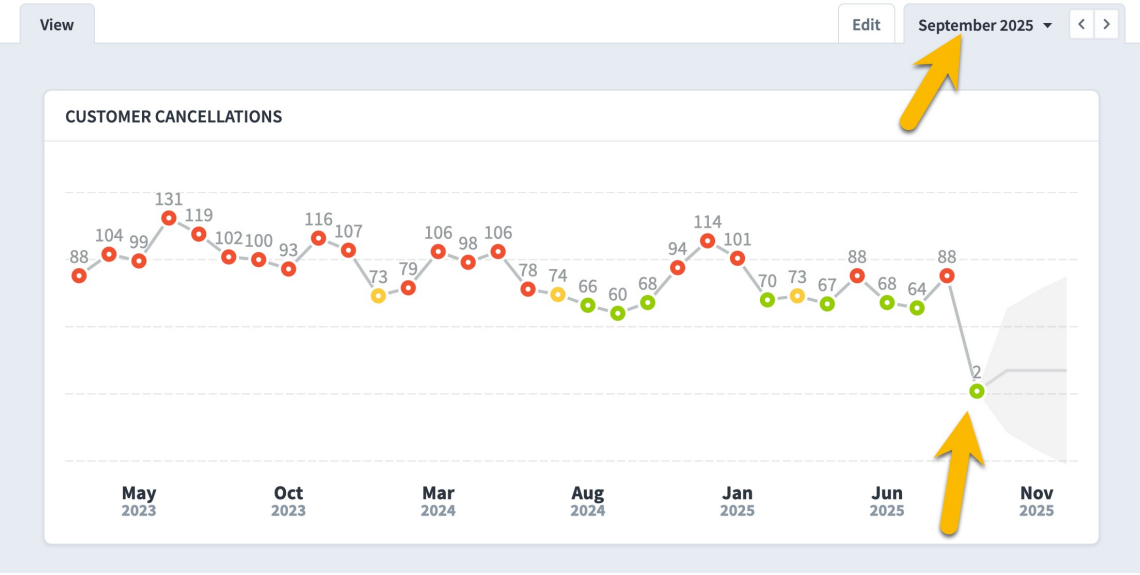

A new "Ignore data from current period" option prevents this problem by excluding incomplete periods from your charts. You'll find this checkbox in the Data Exclusion panel when setting up chart data sources that use calendar periods.

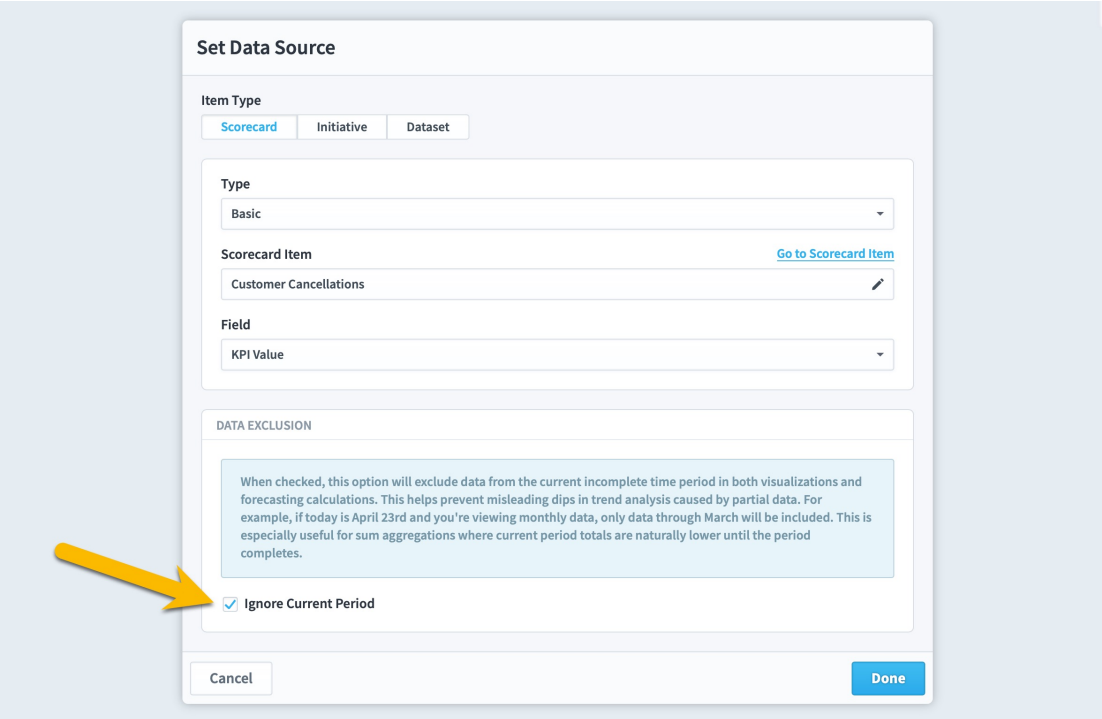

When enabled, the incomplete current period disappears from the chart entirely, preventing artificial drops. The setting also affects forecast calculations and trend lines, which now use only complete periods for their analysis.

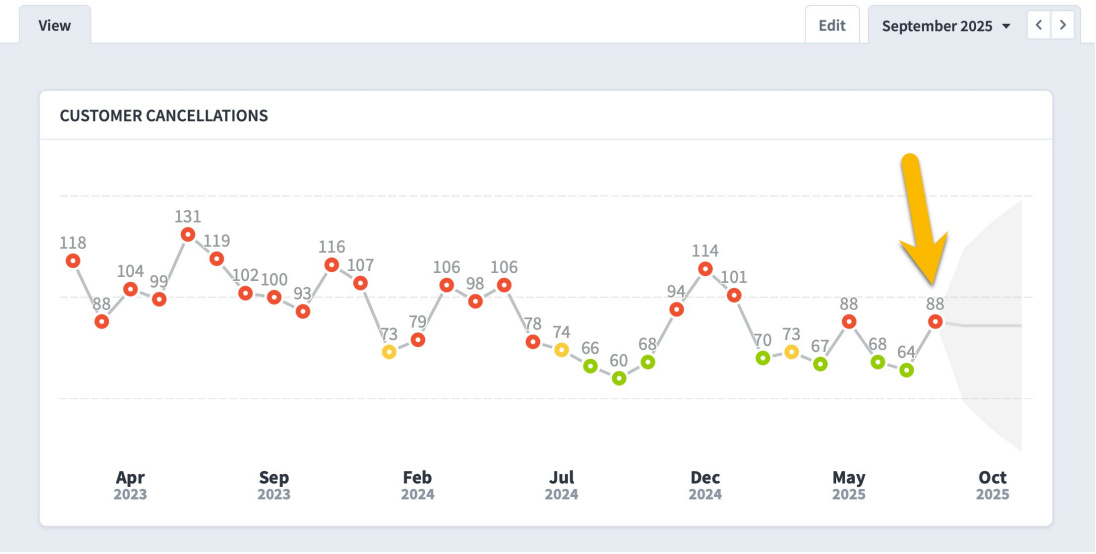

For Sum aggregations, the checkbox defaults to checked since partial totals are misleading. For averages and other calculations that don't accumulate over time, it defaults to unchecked since those values are usually accurate even for partial periods.

Improved Bar Chart Labels

Previously, bar chart labels were automatically sized, which sometimes led to readability issues, especially on charts of varying dimensions. This automatic sizing algorithm has been improved to show readable labels at a variety of chart sizes.

Users now also have the flexibility to manually set font sizes for both data series names and charted values, ensuring better clarity regardless of display size. The "Edit Bar Labels" panel has been redesigned to provide greater

control. The automatic sizing option remains available for those who prefer dynamic adjustments, but manual control allows for a more consistent visual presentation.

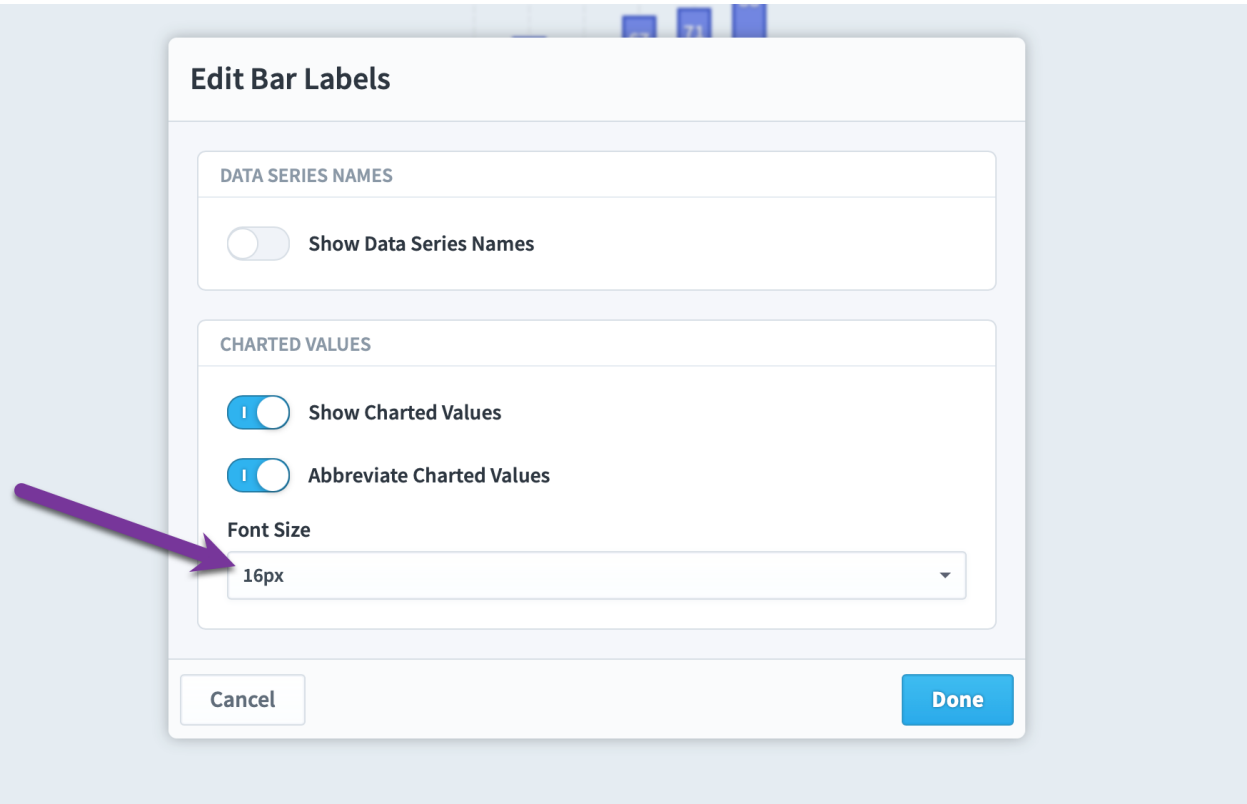

Calculated Scorecard Series in Charts

Spider Impact now supports calculated series in scorecard charts, eliminating the need to predefine calculated KPIs before graphing them. When adding a scorecard series to a chart, users can now select "Calculated" and define equations using KPIs and arithmetic operations. This allows for on-the-fly calculations, such as year-to-date or rolling averages, without requiring a separate KPI.

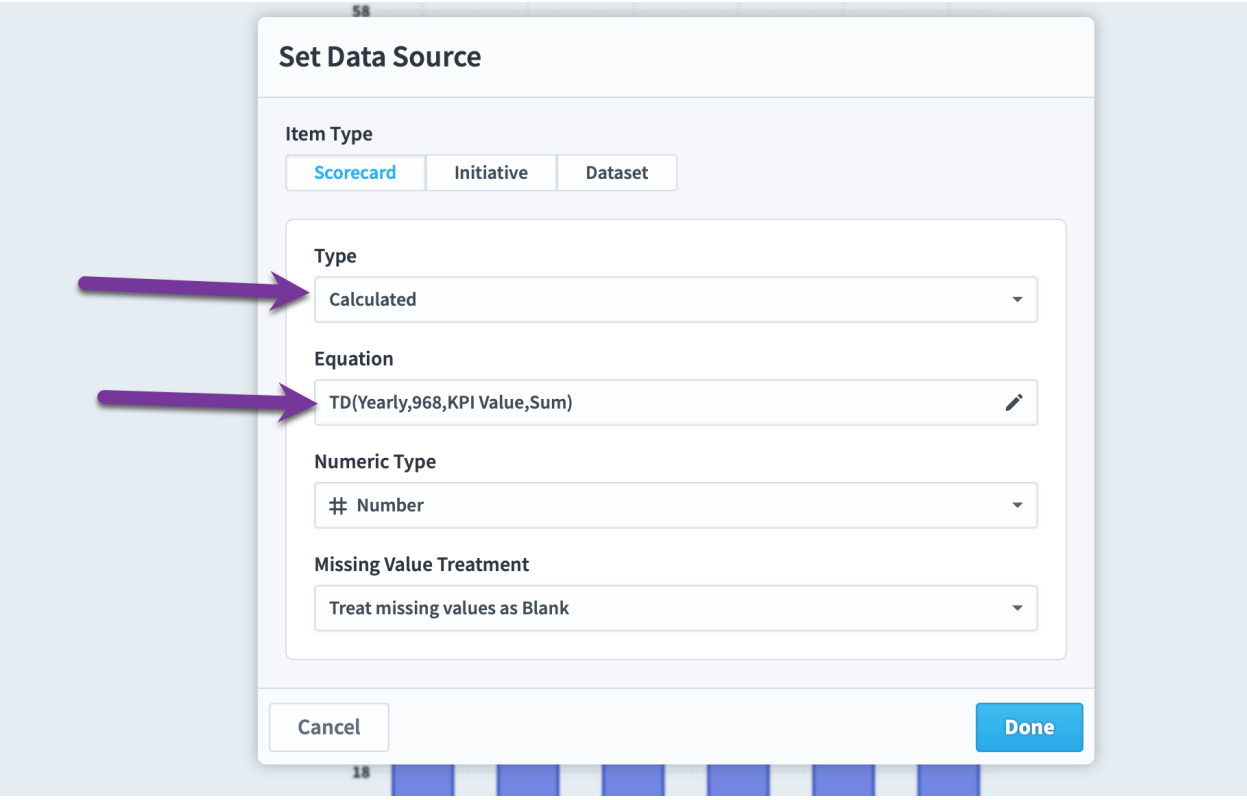

Calculated series cannot be used for Goal Lines or Scoring Bands. Existing charts remain unchanged unless modified.

Dashboards

Snap to spacing between widgets

Dashboard editing now includes automatic snapping between widgets, making it much easier to create professional-looking dashboards with consistent spacing. When "Snap To Widgets" is enabled, widgets will snap to positions that are exactly 16px apart from each other, both horizontally and vertically. This matches the universal spacing used throughout the application and ensures your dashboards have a clean, uniform appearance.

The existing snapping behavior remains unchanged - widgets still snap to alignment guides and can be positioned directly touching each other when needed. The new 16px spacing snap simply provides an additional guide point that helps you quickly achieve even spacing without manual adjustment.

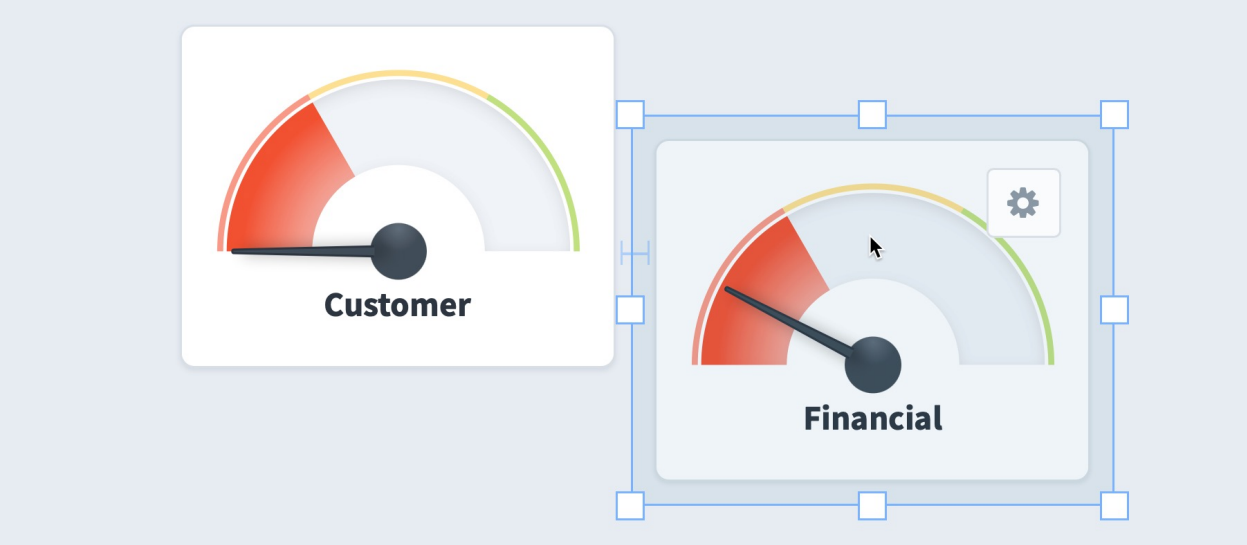

Automatic Dataset Filters on Dashboard Drilldowns

When clicking on dataset values in shape widgets or dataset table widgets, any custom drilldowns to reports or dashboards now automatically inherit the filters that created that specific value. This ensures your drilldowns show exactly the data you expect, without manual filter configuration.

Previously, this automatic filtering only worked when drilling down to the default dataset view. Now it works for custom drilldowns to reports and

dashboards as well, streamlining your workflow when navigating from summary views to detailed information.

For example, here we're setting a custom drilldown on a shape widget showing the number of solar projects to open the "Solar Projects" report.

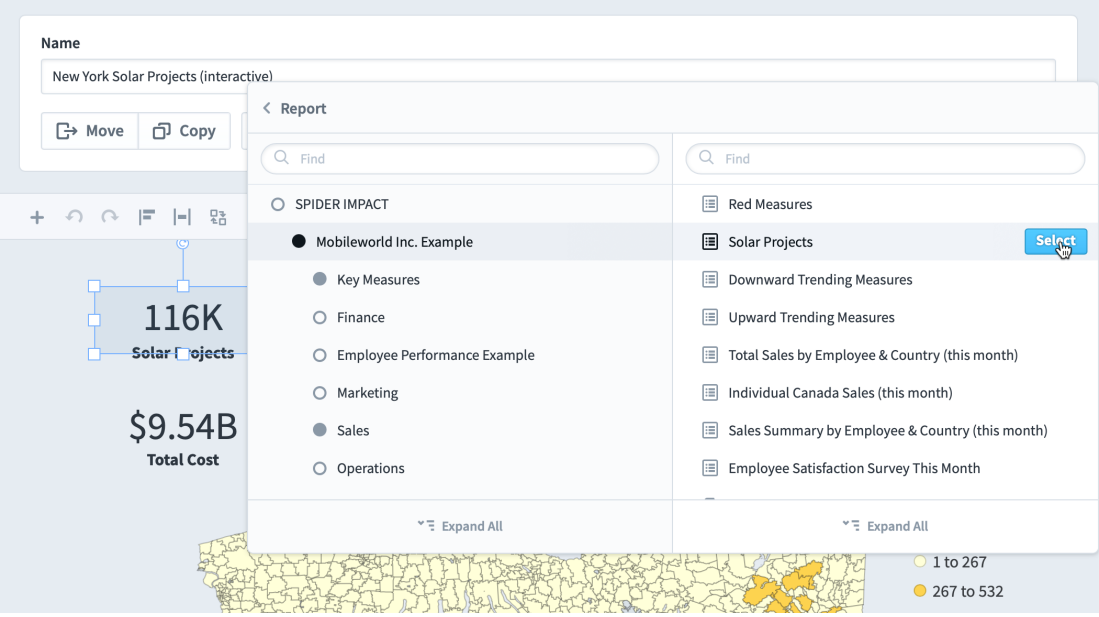

Now when viewing the dashboard with a filter set to show only "Abundant Solar Power Inc" as the contractor, we see there are 10 solar projects for that contractor.

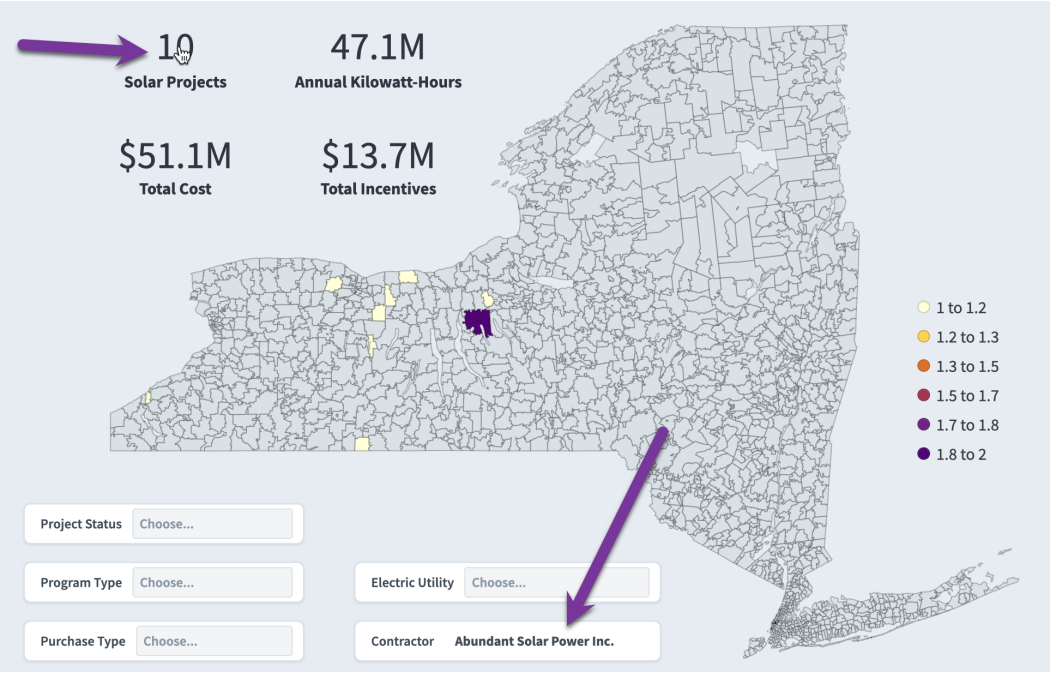

When clicking the number to drill down, the report automatically applies the same contractor filter, showing only those 10 solar projects—no manual filtering required.

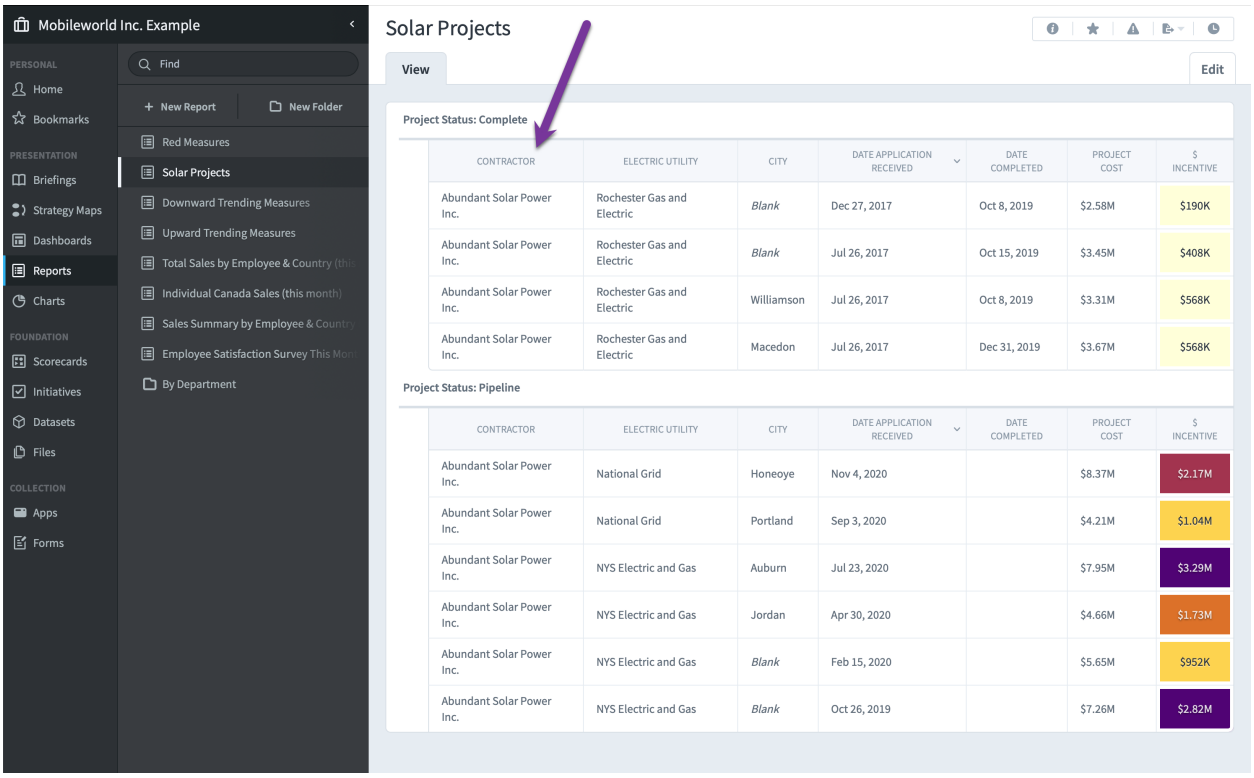

This automatic filter transfer works identically when drilling down to dashboards. Whether your custom drilldown targets a report or a dashboard, the filters that created the value you clicked will be automatically applied to your destination.

Dataset filters are transferred only when the target report or dashboard uses the same underlying dataset as the widget you clicked. Filters within calculated value equations are not transferred.

Dashboard font setting

You can now customize the font used on a dashboard. For example, government agencies that require all documents to use Arial can now set their dashboards to match. To change the font, open the dashboard configuration and select the new "Dashboard Font" submenu. The selected font applies to all dashboard content including widget text, chart labels, and data values.

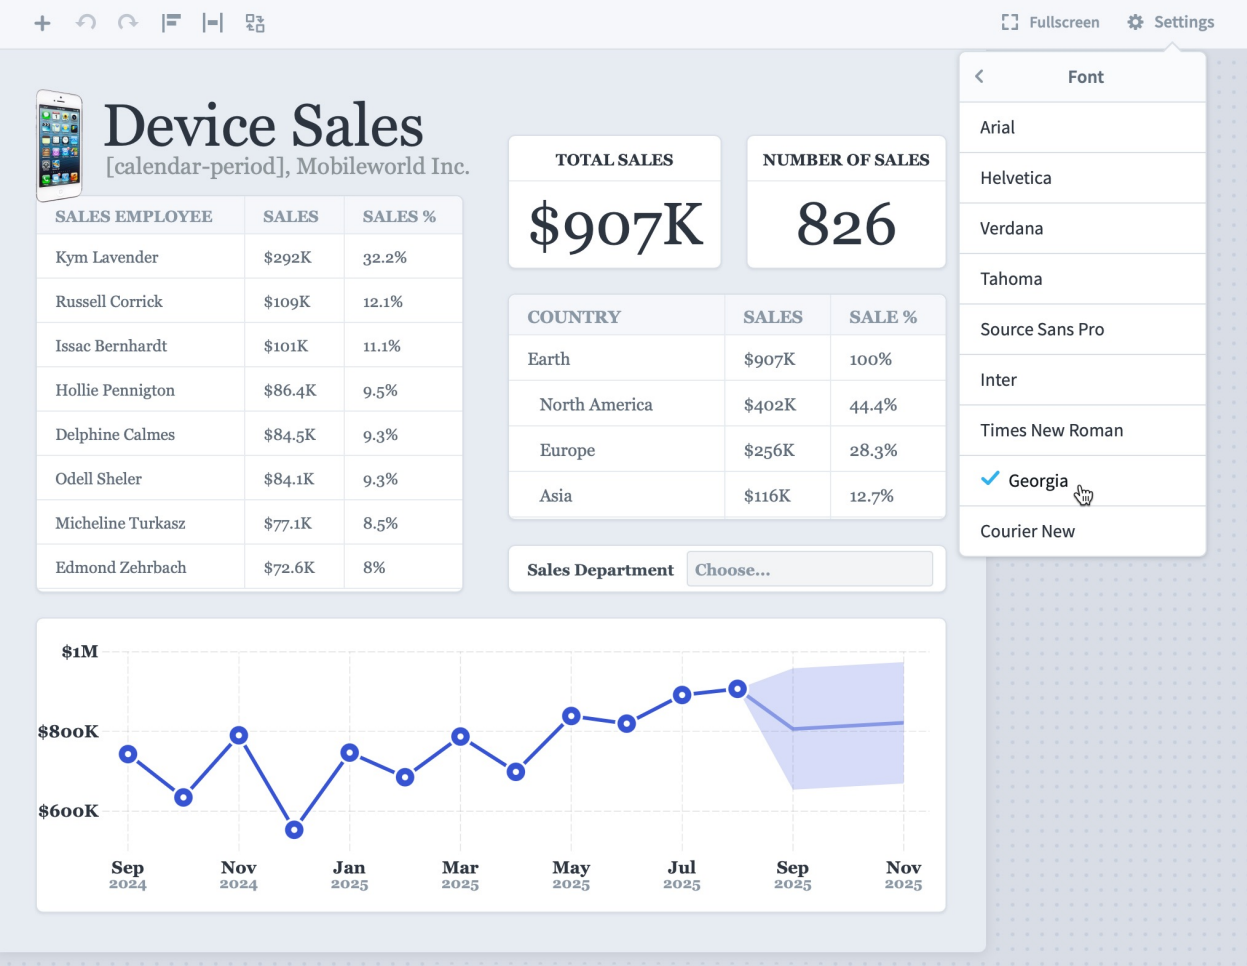

Available fonts include:

- Sans-serif: Arial, Helvetica, Verdana, Tahoma, Source Sans Pro, Inter
- Serif: Times New Roman, Georgia
- Monospace: Courier New

Reorganized Shape Menu for Enhanced Usability

With the latest release, we have redesigned the shape widget menu to improve clarity and ease of use. In the “Shape Display & Formatting” submenu, there are now toggles for displaying shape Fill, Outline, and Text. For example, here we’re only showing Fill, so the Outline and Text options are hidden.

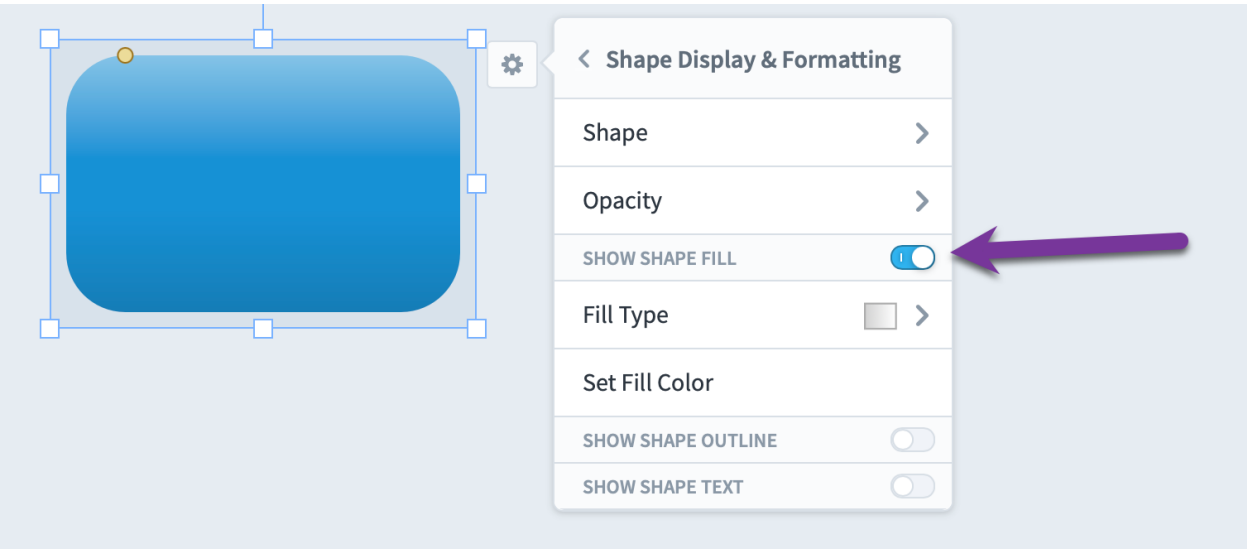

In this example we're only showing Text.

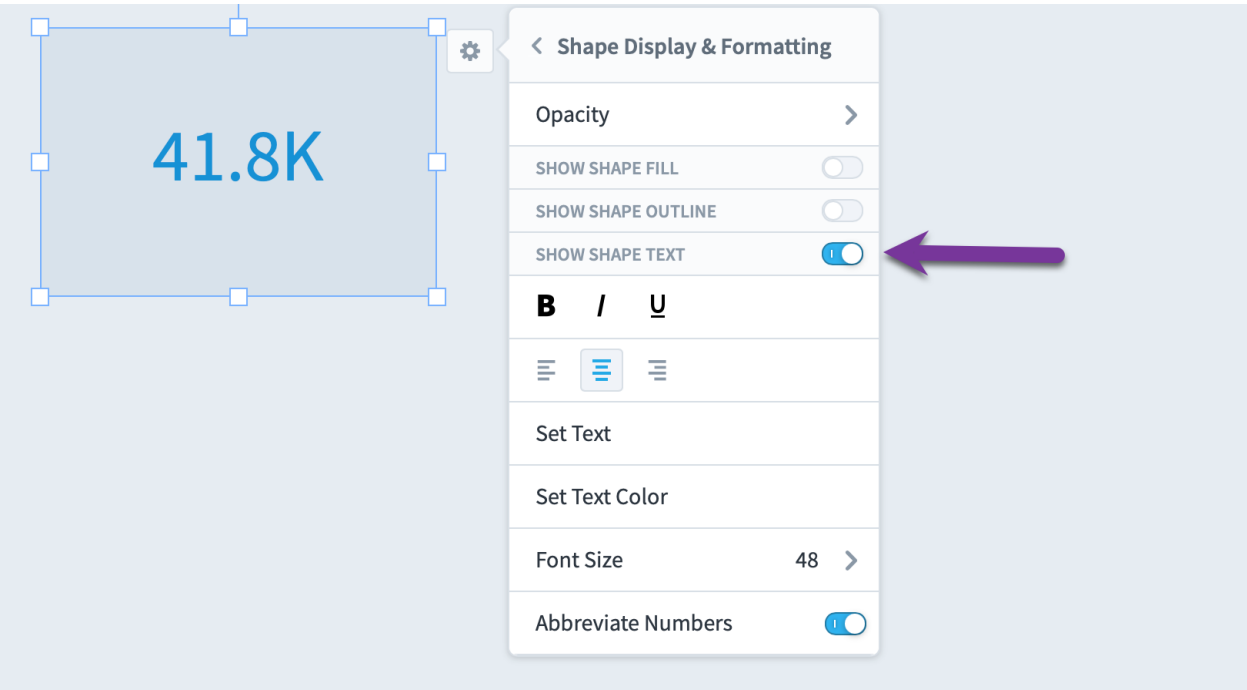

When the data source is a scorecard or initiative item, there is now a "Outside Display & Formatting" submenu. This allows you to show and hide the various supplemental data that appears outside of the shape.

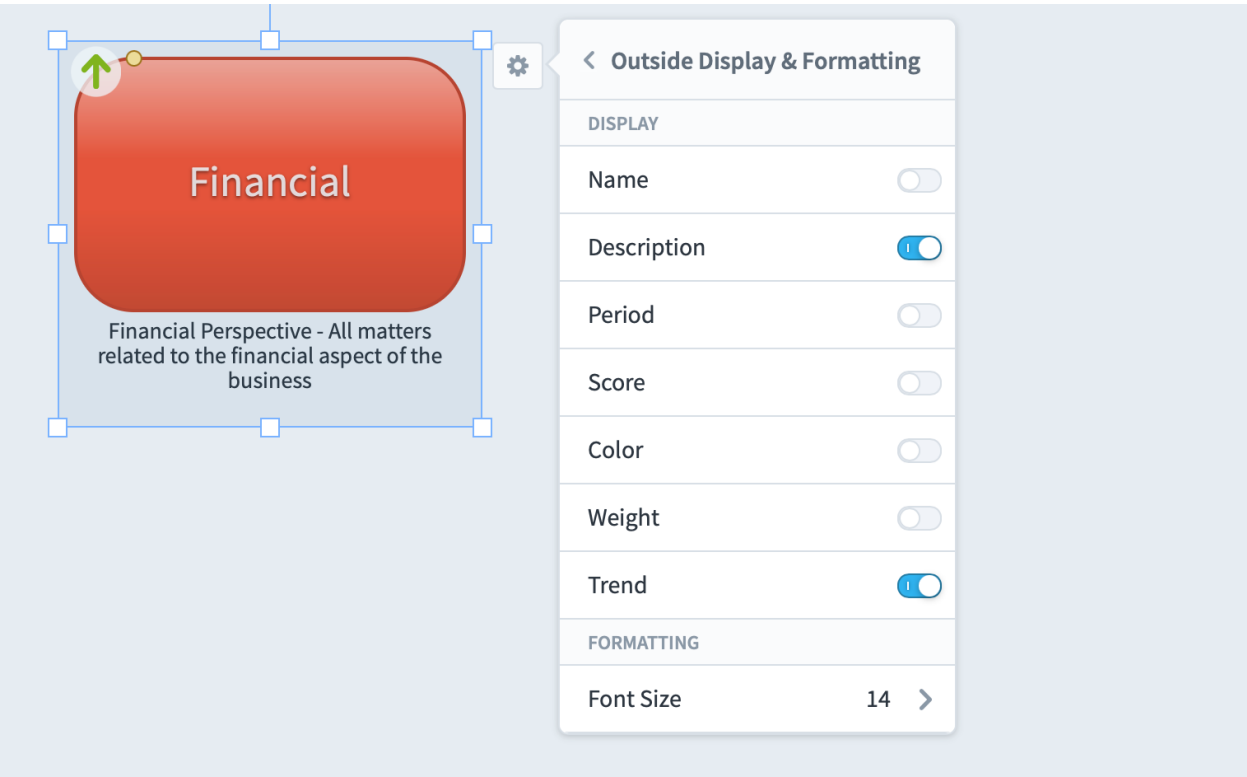

Shape gradient consistency

Shape gradients now maintain consistent light direction regardless of rotation. When you rotate any 2D shape (like arrows, polygons, or chevrons), the gradient automatically adjusts to keep the light source appearing from the top. This ensures that all shapes on your dashboards, strategy maps, and forms maintain a cohesive 3D appearance, even when pointing in different directions.

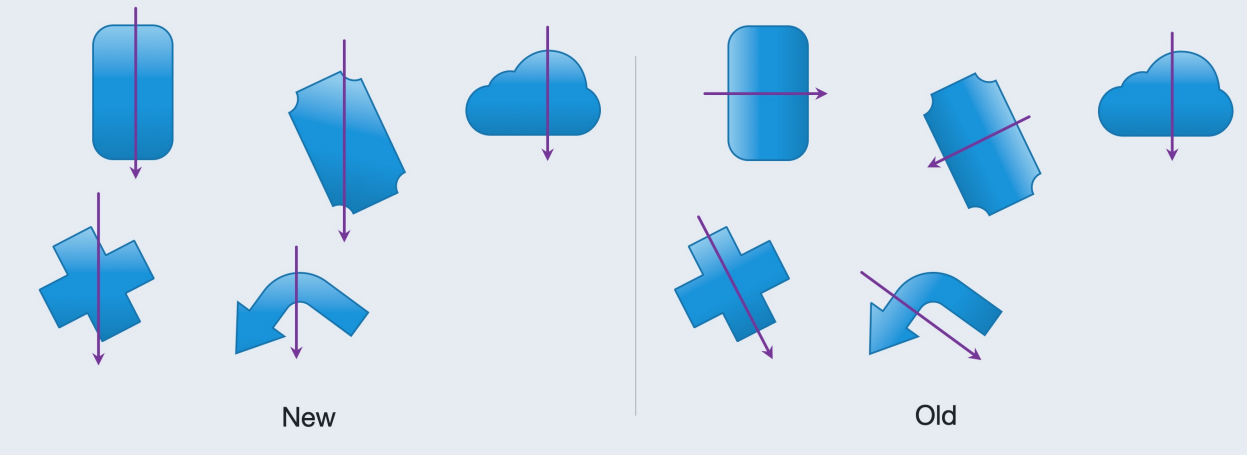

Automatic widget generation limit

When creating a new dashboard with a starting layout (Line Charts, Speedometers, Linear Gauges, or Bubbles), Spider Impact now generates a maximum of 40 widgets to ensure optimal performance. If your organization has more than 40 measures, only the first 40 will receive automatic widgets. Additional widgets can be added manually after dashboard creation.

Forms

Advanced Forms Layout

Spider Impact now supports advanced forms layout, allowing users to create multi-column form structures for a more organized and compact design. Fields can be easily repositioned via drag-and-drop, making it intuitive and efficient to create complex layouts.

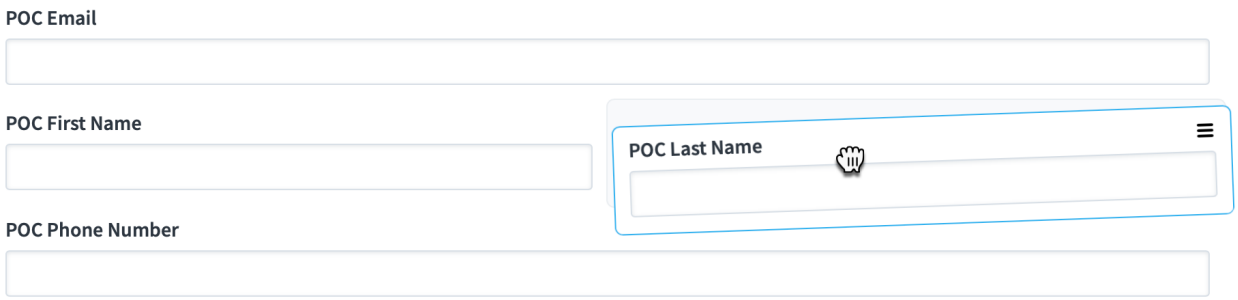

These new layouts are fully responsive, automatically collapsing into a single-column format on mobile devices to ensure a seamless experience across all screen sizes.

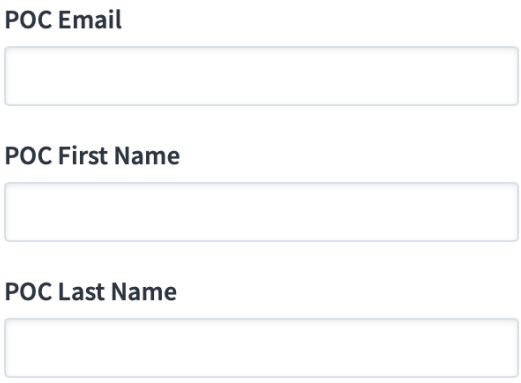

Panel Form Widget

The new Panel Form Widget improves form organization by allowing users to group related fields into a single container. This simplifies large forms by providing a structured layout with easy drag-and-drop repositioning.

To add a panel widget, select it like any other form widget from the form page's "Add Widget" menu.

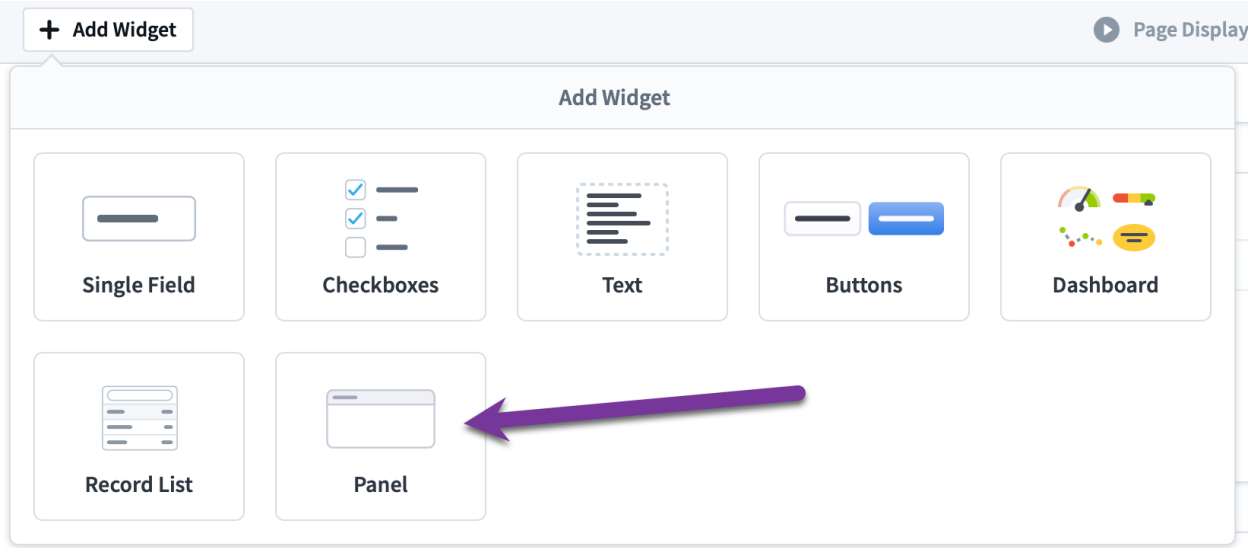

Panels support advanced layouts with nested rows, columns, and stacked elements, offering greater design flexibility. In this example we've added two panels in a side-by-side layout.

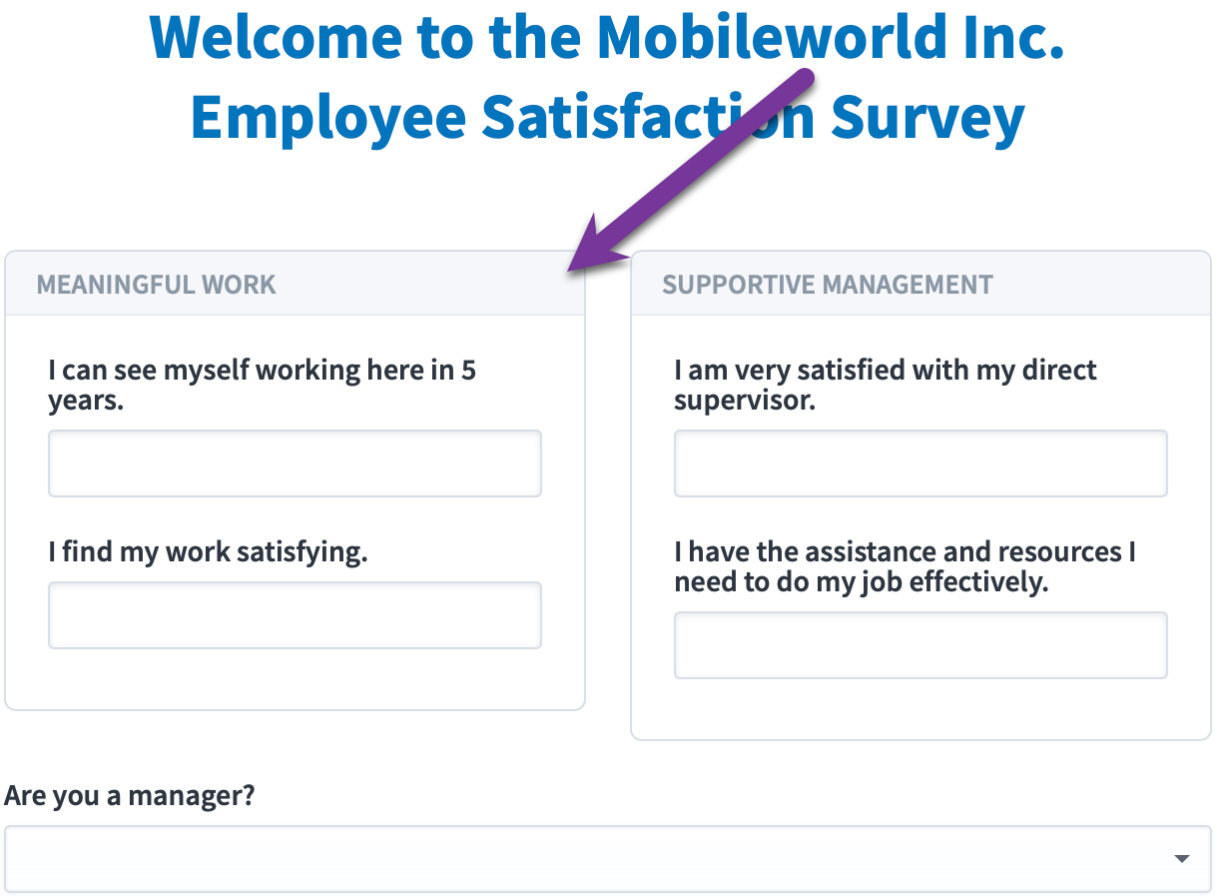

The widgets inside of Panels can also be arranged in multi-column layouts.

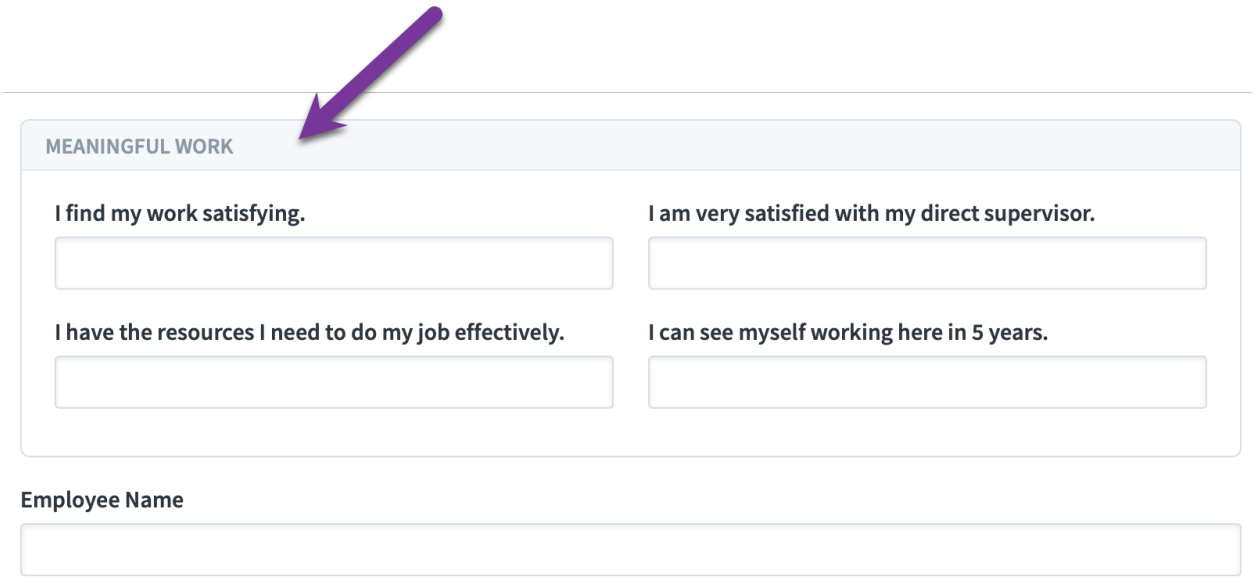

Form actions can toggle Panel widget state between visible and hidden, which allows for easily hiding the panel and all its contents when needed.

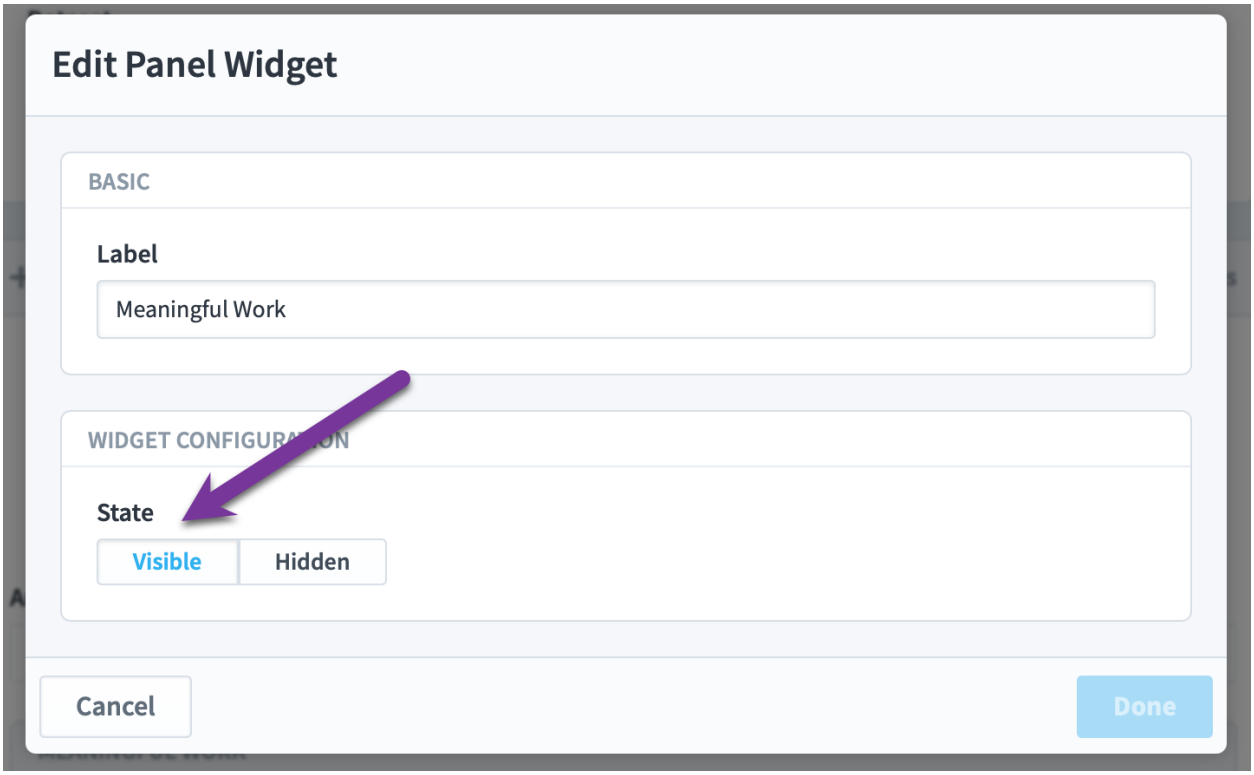

Buttons Widget

The new release replaces the "Button" widget with "Buttons," allowing multiple buttons within a single widget. This change simplifies form layouts by supporting multiple button configurations in one place.

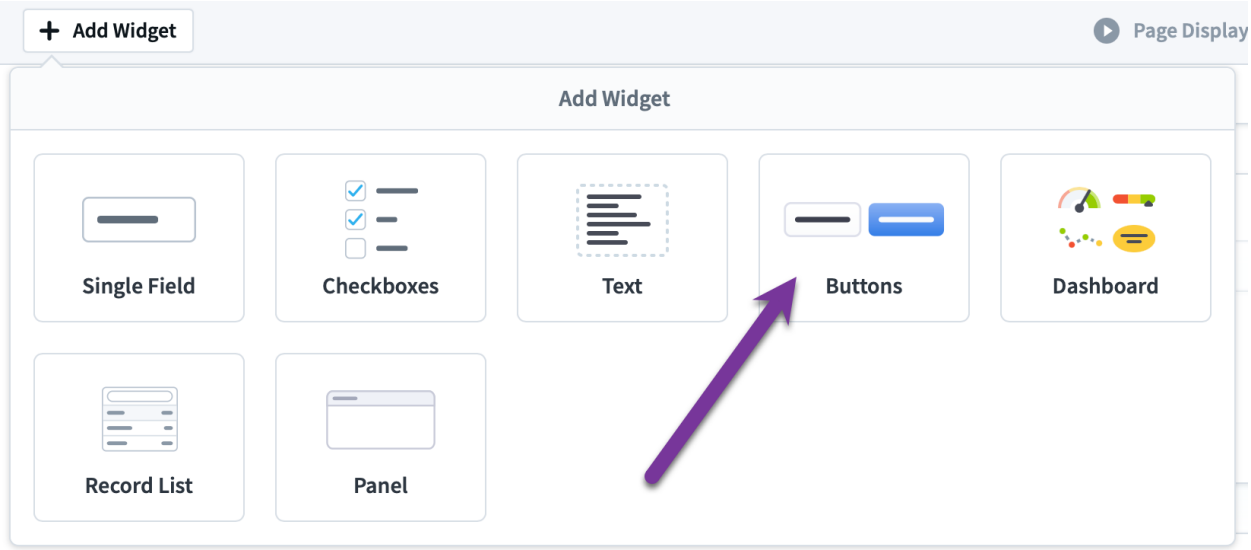

Previously, each button required a separate widget in its own row, leading to cluttered designs. Now, users can add multiple buttons within the same widget, organizing them into right-aligned and left-aligned groups. Each button retains independent actions and states.

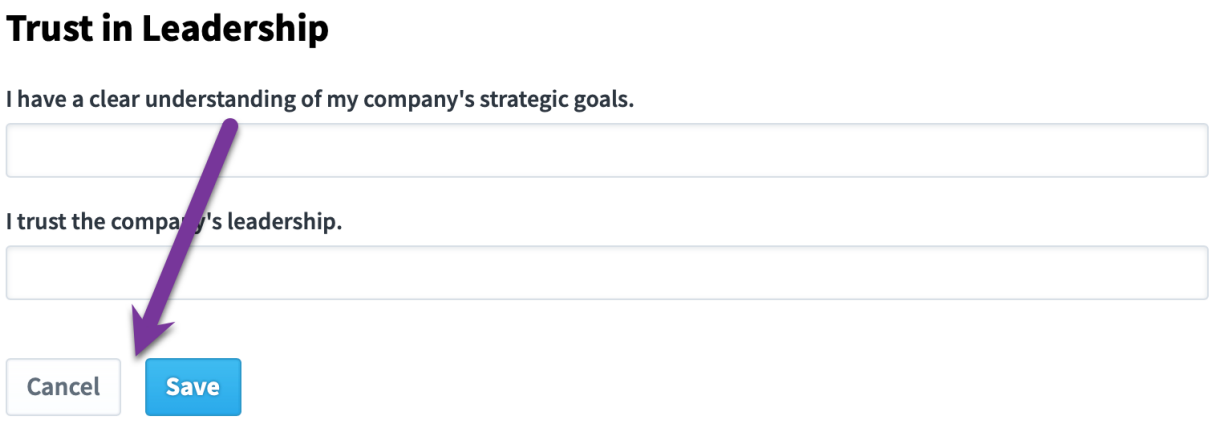

A new “Use as the Bottom Sticky Bar" option helps ensure key actions remain accessible.

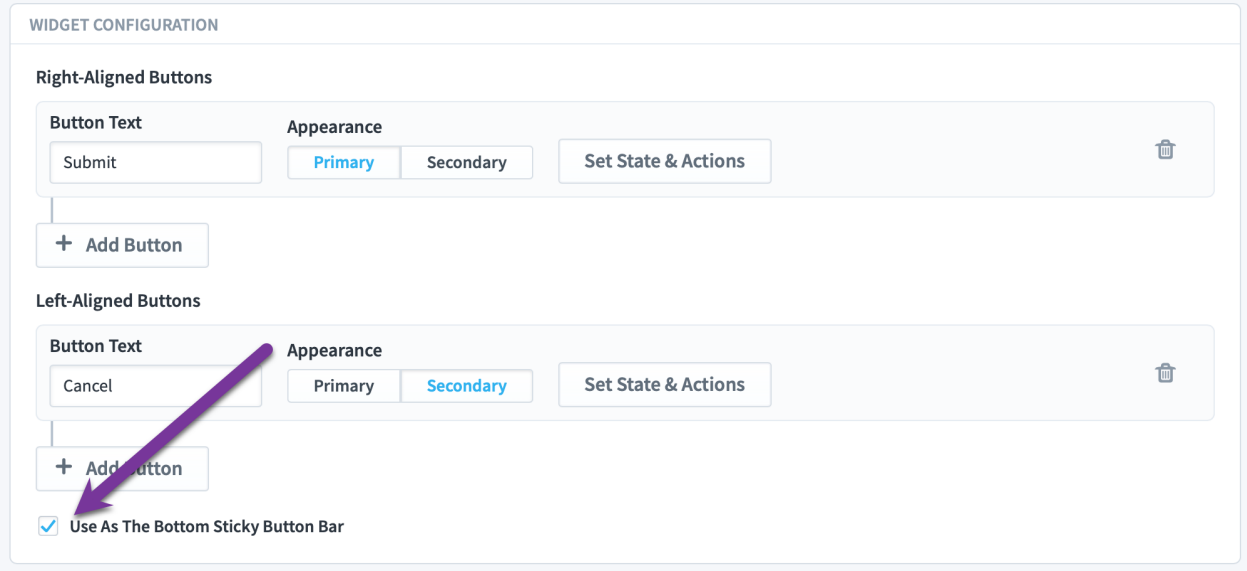

This creates a permanent button bar on the bottom of the form that covers the rest of the form content.

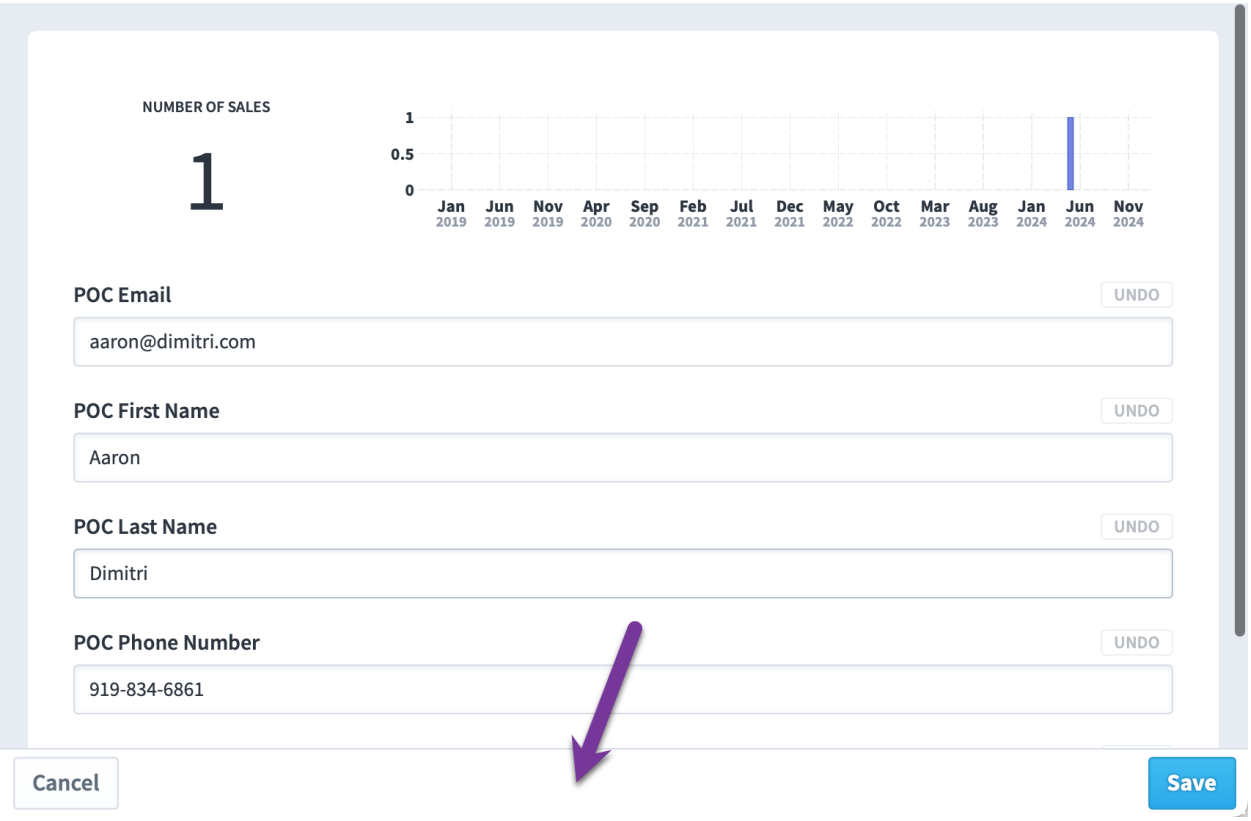

Existing Button widgets will automatically upgrade to the new format without disrupting workflows.

Enhanced Sorting for Record List Columns

Record list widgets can now be sorted by multiple columns, both when configuring forms and when viewing data.

To set default sorting, each column in the Record List Widget configuration now has a dropdown menu. Choose from options like "A to Z," "Z to A," "High to Low," or "Default" (no sorting). When multiple columns have sorting enabled, the widget sorts from left to right—first by the leftmost sorted column, then by the next, and so on.

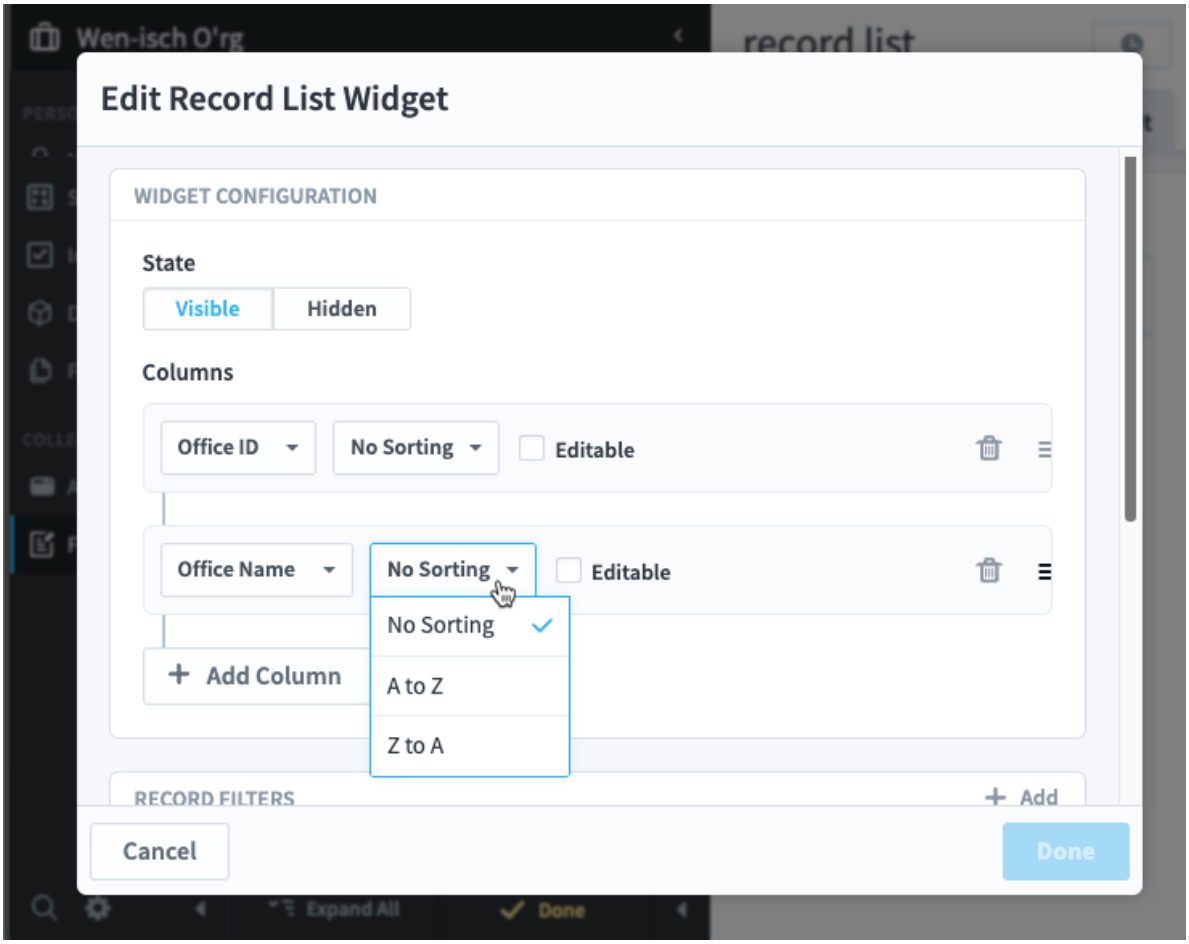

While viewing a record list, users can click any column header to sort that column. Click once for ascending, twice for descending, and a third time to return to the default sort. A small arrow indicates the current sort direction.

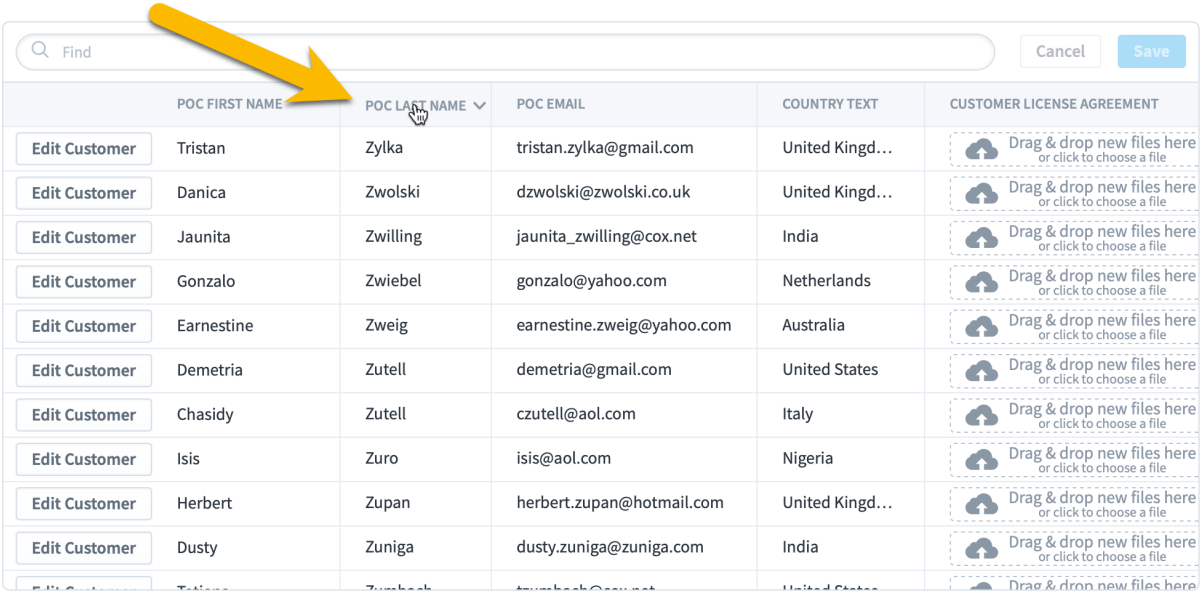

Loading and success animations for form buttons

Form buttons now display a loading animation when clicked, then briefly show a checkmark when the action completes successfully. This provides clear feedback that your action was processed, especially helpful when saving data without navigating to a new page.

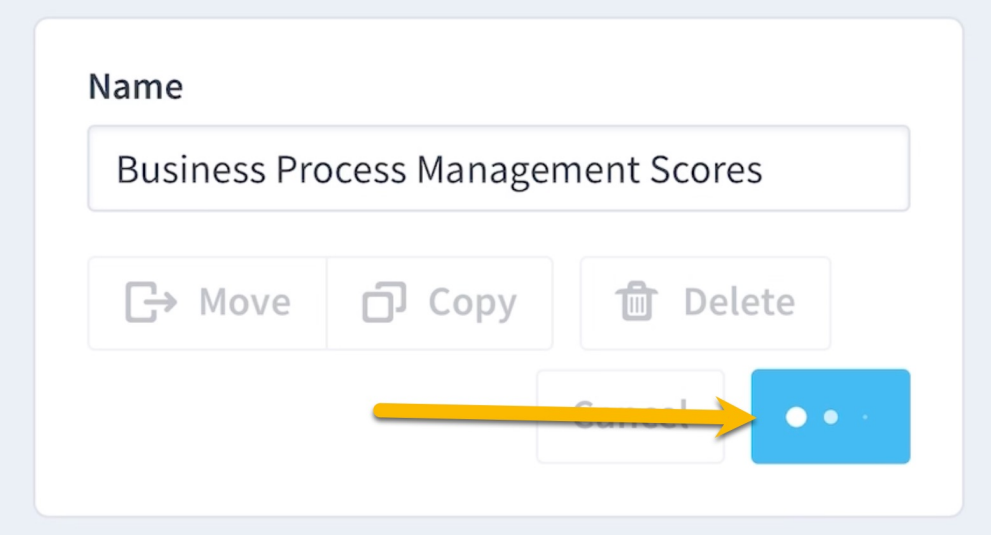

Behind-the-Scenes Data Manipulation

Forms can now manipulate data that isn't directly visible on screen. By combining form variables, Load from ID actions, and the Save a record action,

forms can retrieve, modify, and save records entirely behind the scenes without user interaction.

Form-level variables

Forms can now create and use temporary variables that exist alongside dataset fields. These variables maintain their values while the form is being used and enable complex workflows without requiring dummy dataset fields. Variables are created through the "Set Variables" button when editing a form.

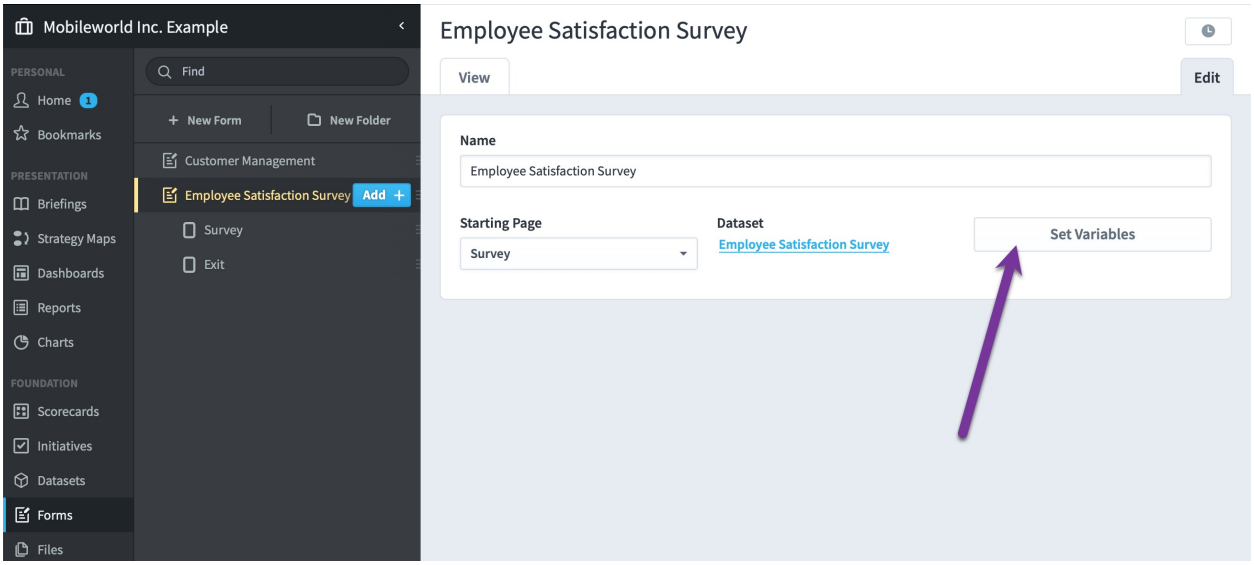

Like dataset fields, each form variable has a name and data type information. Unlike dataset fields, form variables have optional default values. For example, a "Commission Percentage" variable with a default of 35% can be referenced in equations throughout a form without needing to be its own dataset field. When form actions update this variable, all references to it automatically reflect the new value.

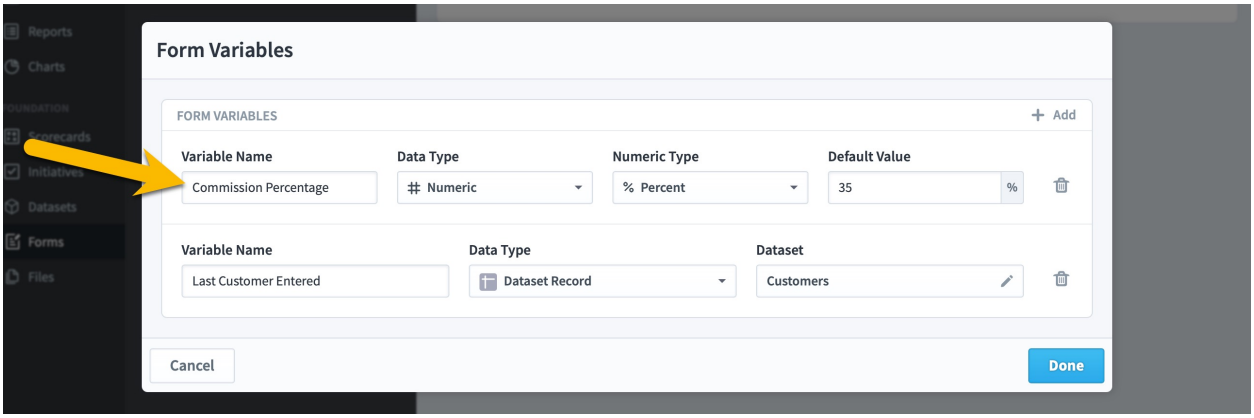

The Dataset Record type is particularly powerful—it stores an entire record from any dataset, not just a single field value. This enables forms to load, modify, and save complete records behind the scenes without displaying them to users.

For example, a "Last Customer Entered" variable can streamline data entry in a New Customer form. When someone creates a new customer, an action saves both the record and stores it in this variable. The next time the form opens, it checks if this variable contains data and pre-fills relevant fields (like country and customer service rep) from the previous entry. Since form variables persist throughout the user's session, this information remains available even when navigating between different records.

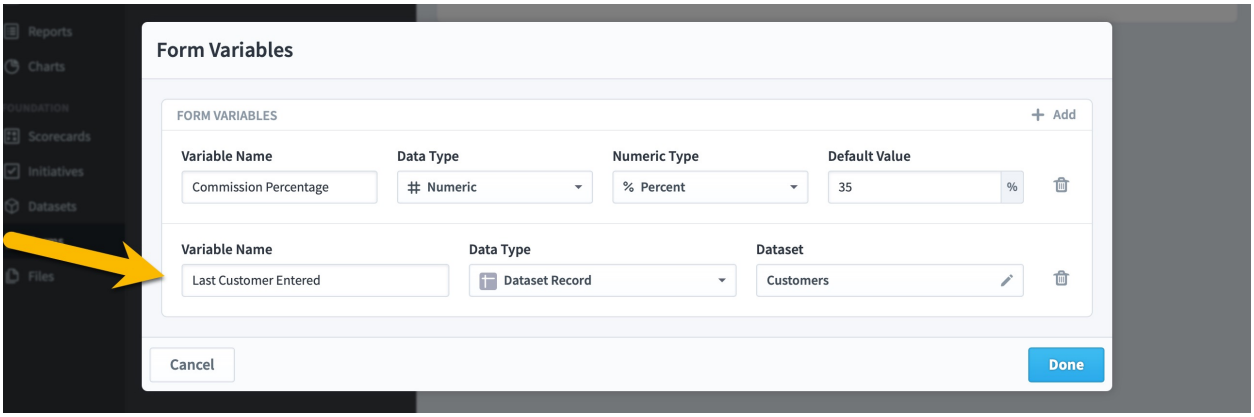

Using form variables is simple because you can reference form variables anywhere you could previously reference form fields. They appear in a dedicated section at the bottom of field selection dropdowns throughout the form in widgets, conditions, equations, and actions. They provide temporary storage for UI controls, intermediate calculations, and data that doesn't need to be saved to the dataset.

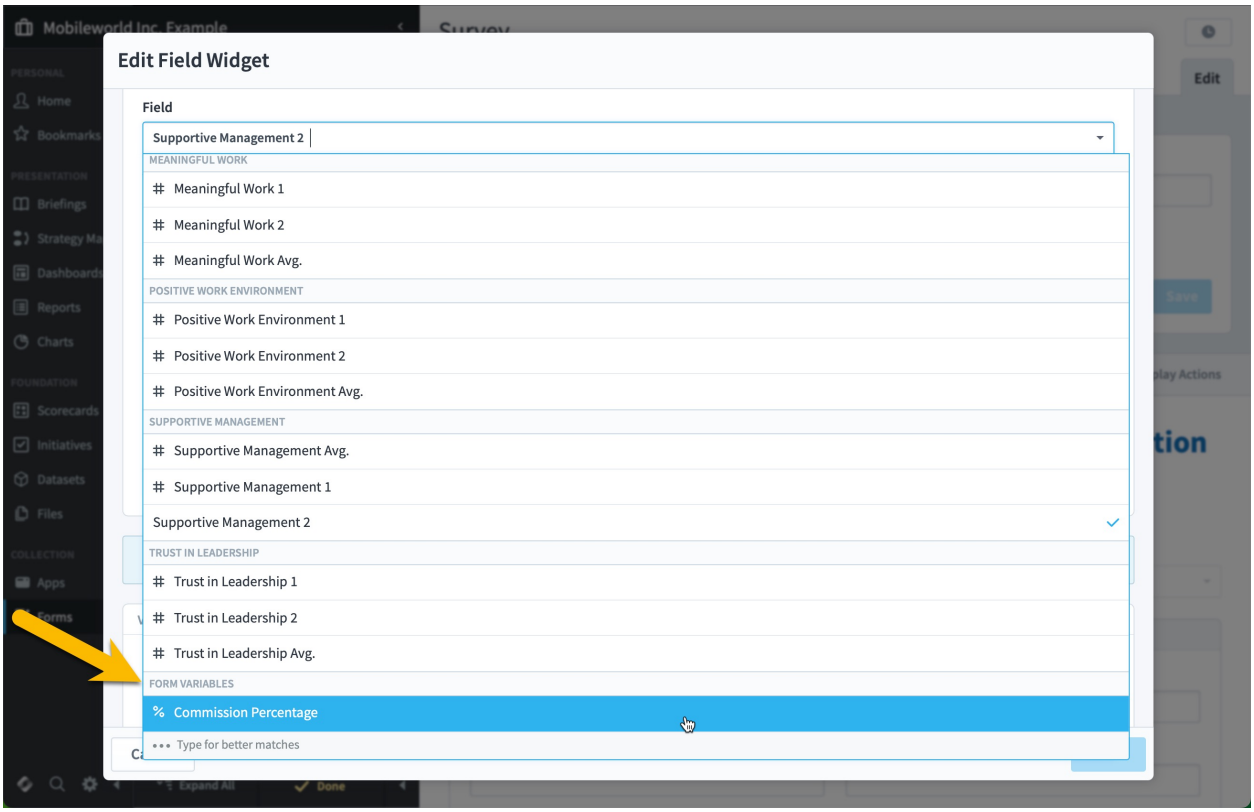

"Set a value" action

The previous "Set a field value" action has been renamed to "Set a value" and now supports setting:

- Individual current record fields (all fields, not just those with widgets)
- Form variables of any type
- The entire current record
- Values using the "Another Form Value" option to copy between fields and variables

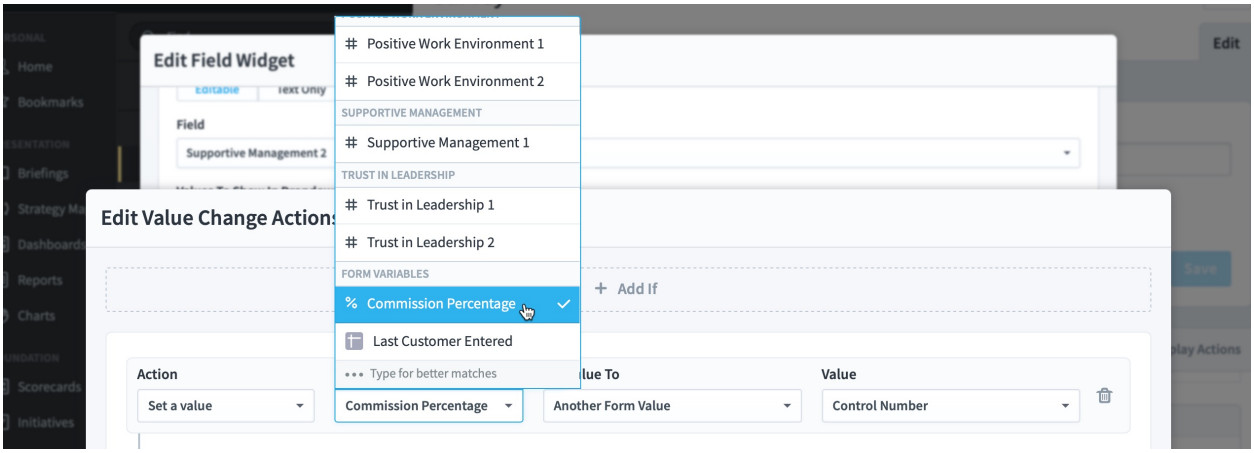

When working with Dataset Record variables, you can replace entire records or set individual fields within them. The software ensures type compatibility - Dataset Record variables can only be replaced with records from the same dataset.

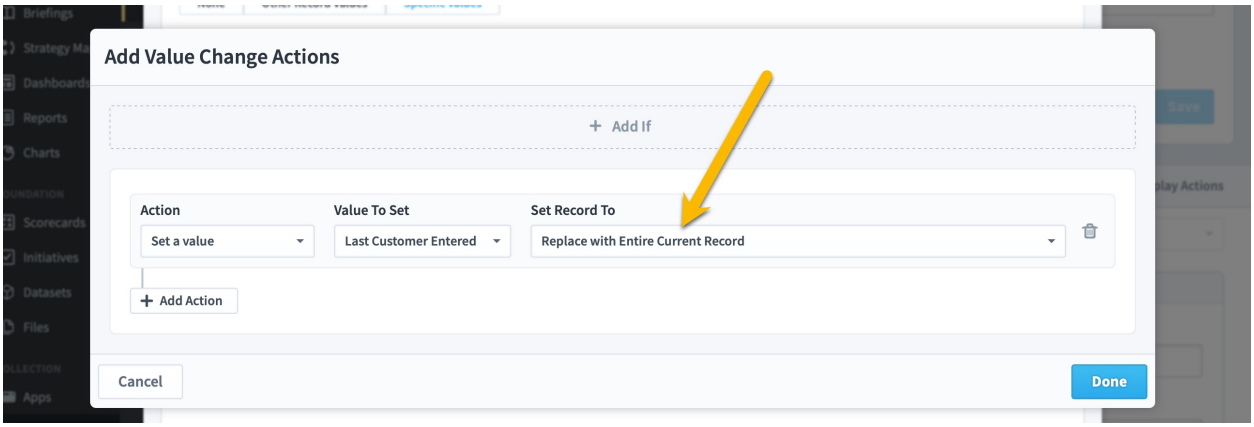

Loading records from a dataset

Forms can now load existing dataset records using ID lookups without requiring user selection from a list. When setting a value, the new "Load from ID" option uses a field or variable value as a Record ID to retrieve the matching record **by its Record ID only**. After selecting "Load from ID", you choose which field or variable contains the ID to use for the lookup.

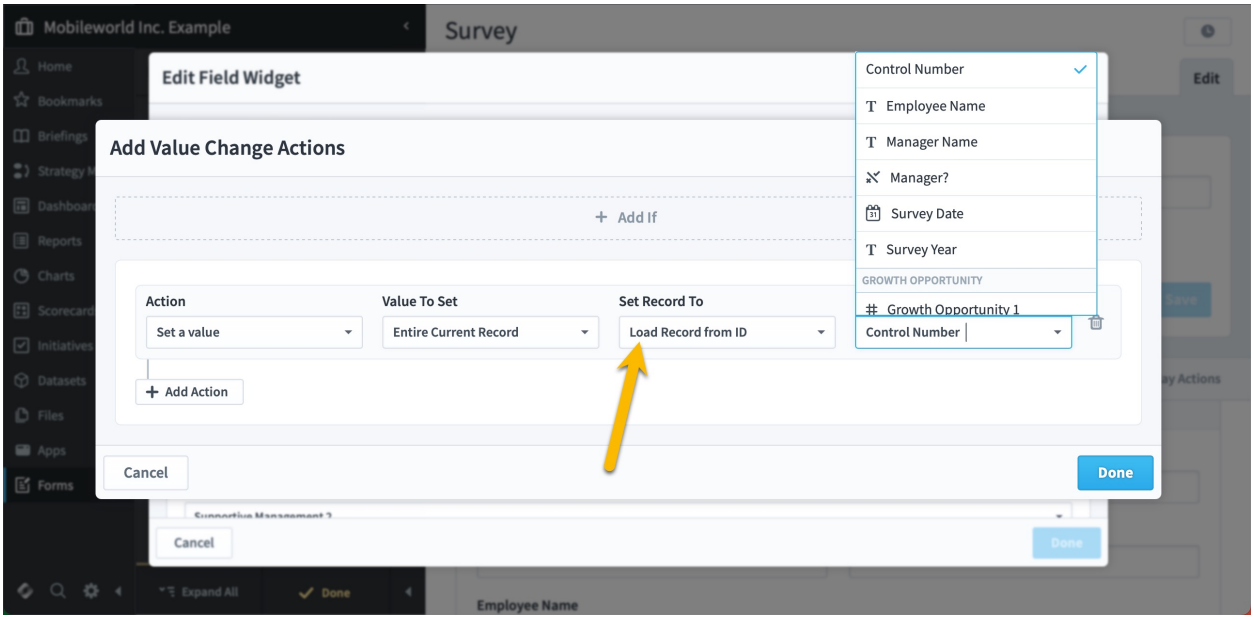

You can use any value as the ID to look up. The software will just take whatever value you give it, and will try to find a record with that ID. If no matching record is found, the record or variable is set to blank.

Loading the current record or a dataset record variable from ID enables powerful scenarios like preventing duplicate form submissions, or dynamically loading related records based on form input.

Save any record

The "Save the current record" action has been transformed into "Save a record", allowing forms to save either the current record or any Dataset Record variable. This enables creating and saving records in other datasets behind the scenes - perfect for audit trails or multi-dataset workflows.

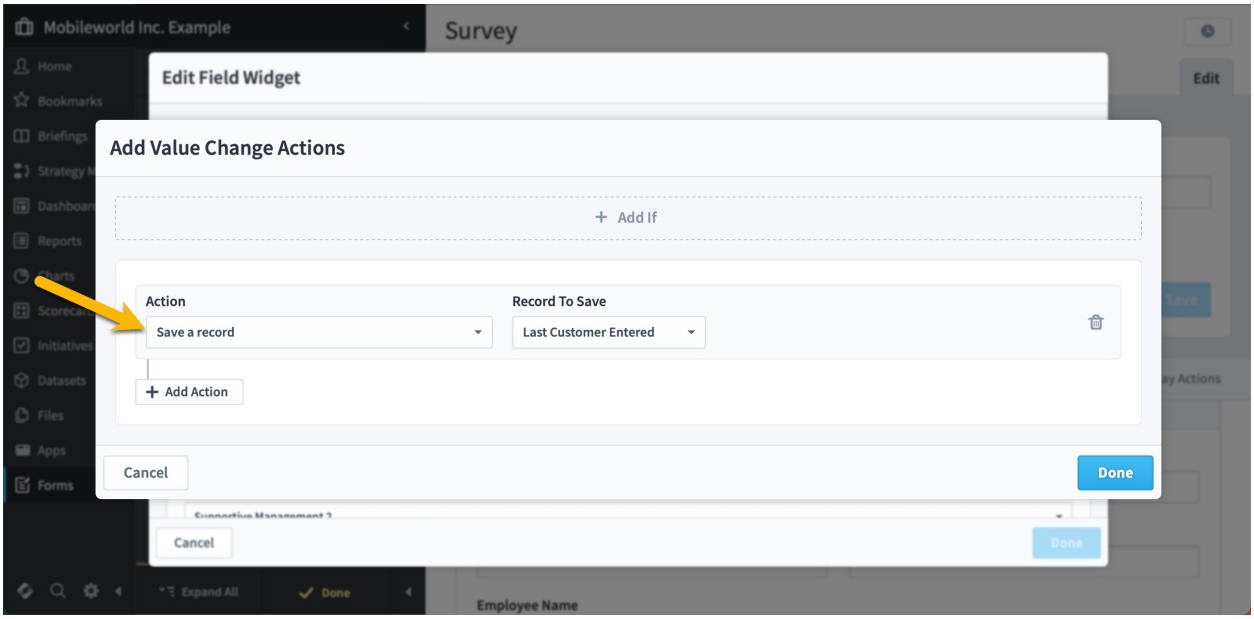

Advanced Form Logic with Equations

Equation-based IF conditions

IF statements in forms now support an "Equation is True" condition, enabling complex logic involving multiple fields.

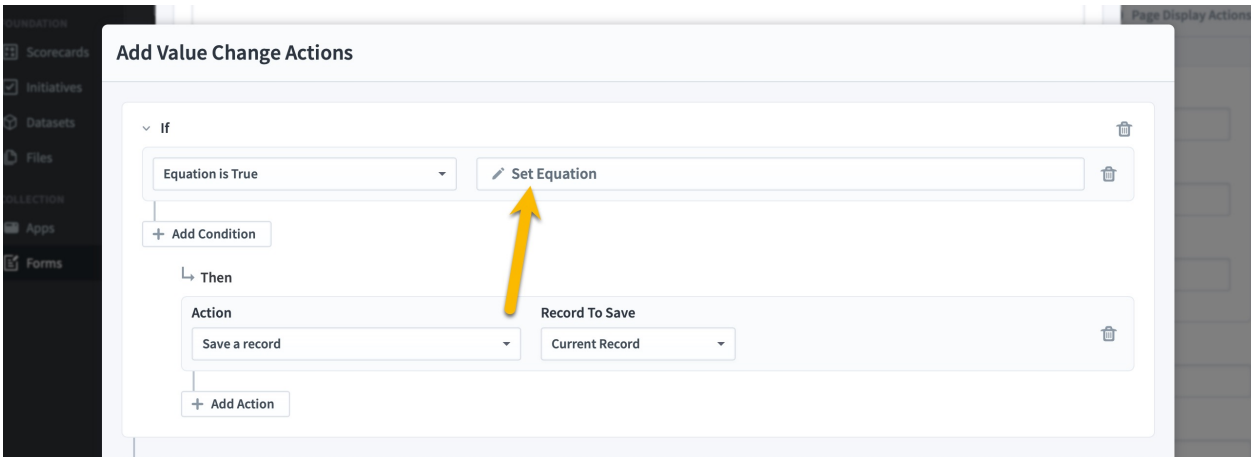

The equation builder includes a "Type" dropdown with three options:

- **Current Record** (default) - allows selecting fields from the form's dataset or any linked dataset
- **Form Variable** - lists available form variables
- **Dataset** - provides aggregated data from any dataset

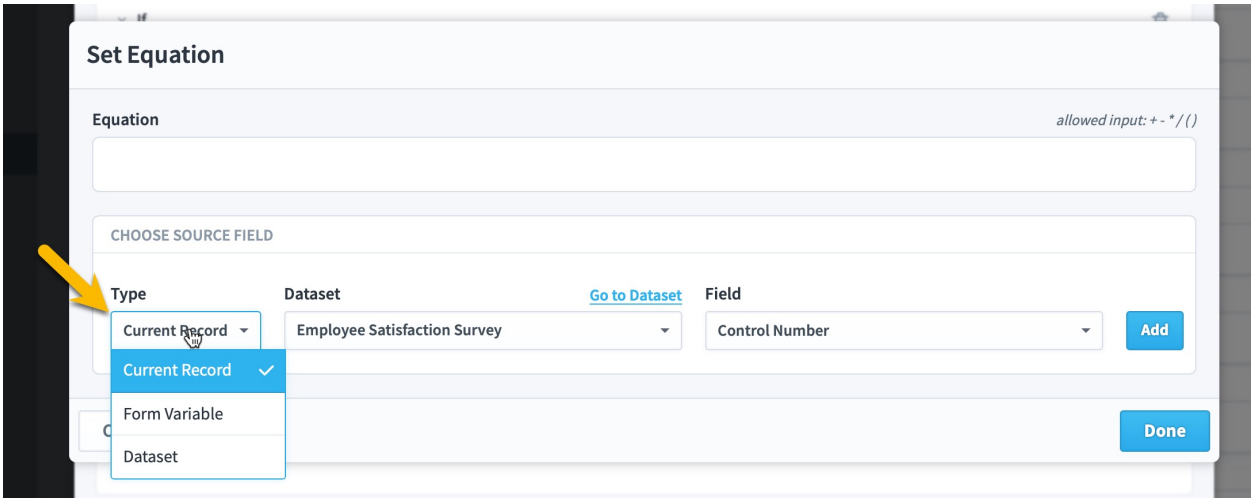

Equations in IF statements allow you to create multi-field conditions like ensuring a Case Report Date comes after an Incident Date, or verifying that at least one checkbox from a group is selected. In this example we're checking to see if both the employee and manager names are blank.

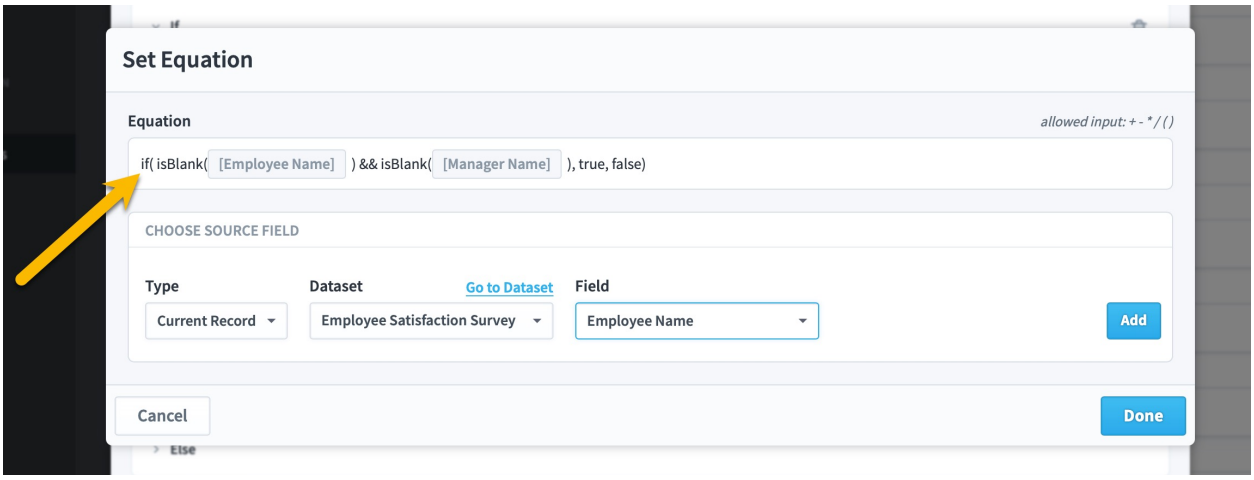

Equations can also reference values from other datasets, allowing forms to adapt their behavior based on external data - for example, displaying different content based on whether your website status is "up" or "down" (stored in another dataset), or limiting form submissions to the first 50 entries by checking a count from the dataset. In this example we're looking up the current Australian to US exchange rate from another dataset and changing form behavior based on whether it's higher than 1.

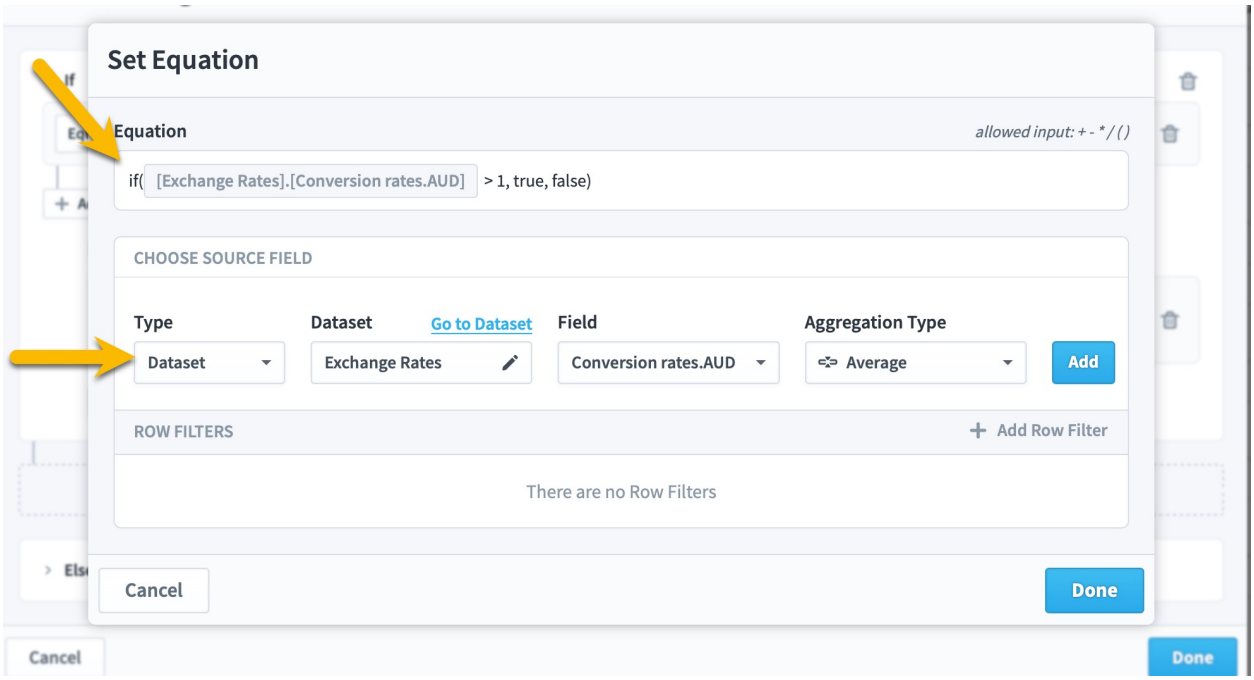

Calculate field values with equations

A new "Result of Equation" option now appears when setting field values.

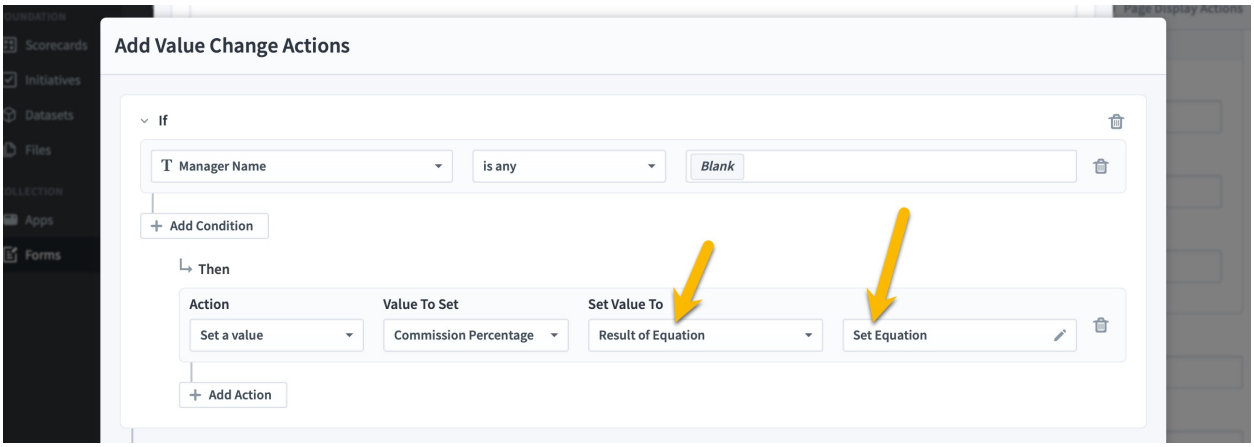

The equation builder for setting a value is the same as the new equations in IF statements, allowing you to reference the current record, form variables, or aggregated dataset data. This functionality makes it possible to generate record IDs by combining fields or incrementing existing IDs, calculate dates, auto-populate fields based on user permissions, or aggregate values that can't be handled by standard calculated fields.

Creating and Editing Linked Records in a Dialog

Previously, forms could only edit a single record from a single dataset. This enhancement allows users to create and modify linked records from different datasets directly within a form, eliminating the need for separate forms or navigation. Users can now work seamlessly across datasets while maintaining their workflow.

The new ability to edit linked records in a dialog is best explained with an example. Here we see a customer management form for the Customers dataset. This form starts with a searchable list of all customers. We'll click on the "Edit Customer" button for Abbey Evetts.

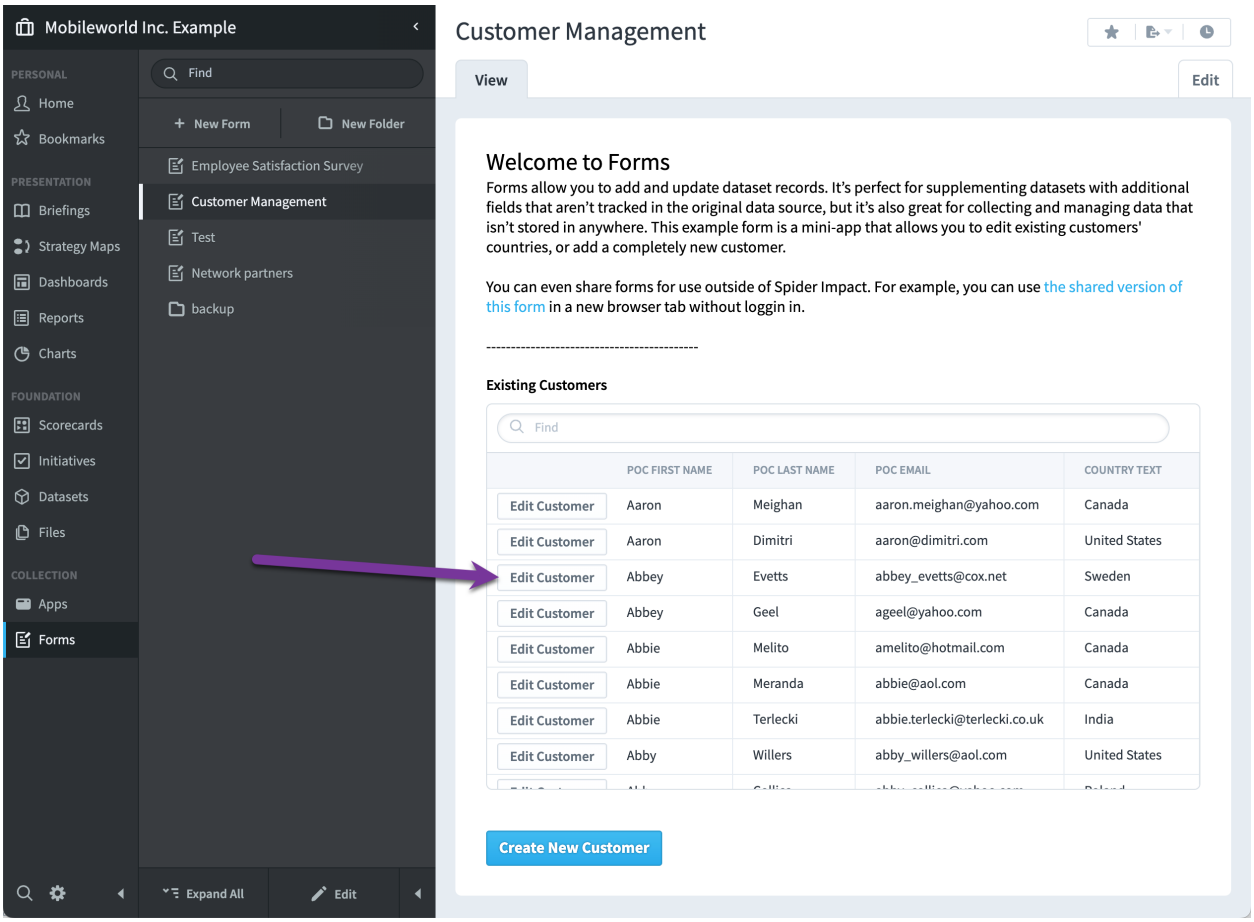

This takes us to the customer details screen. We can see the four sales for this customer, each of which is a record in the linked Device Sales dataset. We can create a new sale by clicking on the "New" Button, or delete an existing sale by clicking on a sale's X icon. Instead, however, we'll click on one of the existing sales.

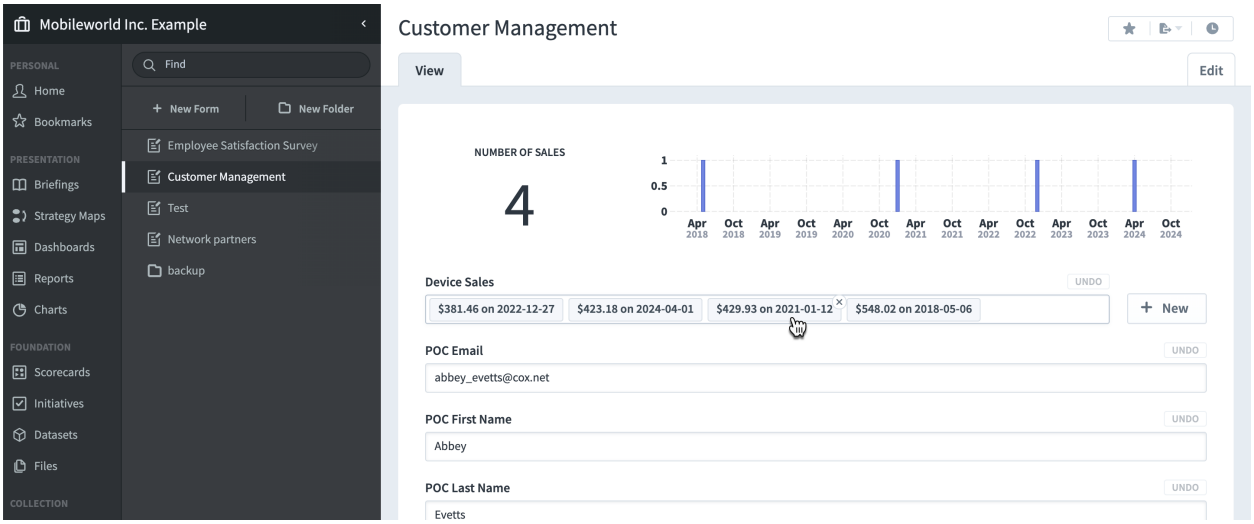

This opens a dialog that allows us to see details about that sale and even edit its details if we have the correct permissions.

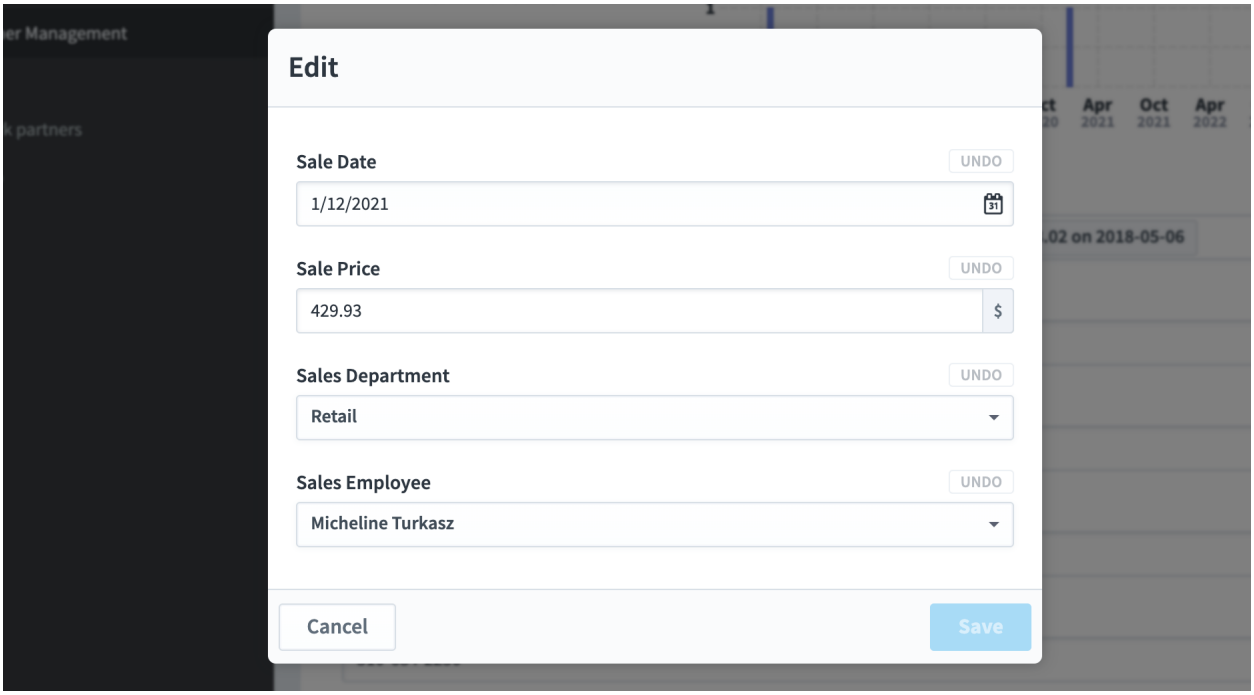

Editing linked records in a dialog like this is intuitive for users and opens new possibilities for interactive data collections apps. To make this work, however, there are several large enhancements that were introduced to the software.

Form dialog items

There is a new "Dialog" option in a form's "Add" menu.

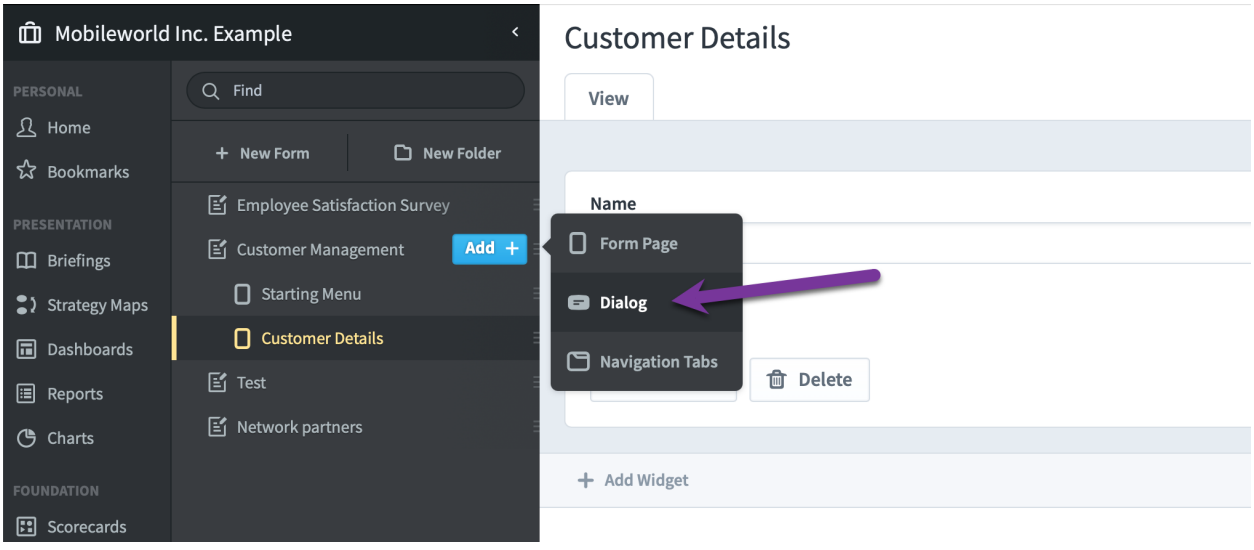

Just like the Form itself, each Dialog has a dataset. Here we'll choose "Device Sales" and click "Create".

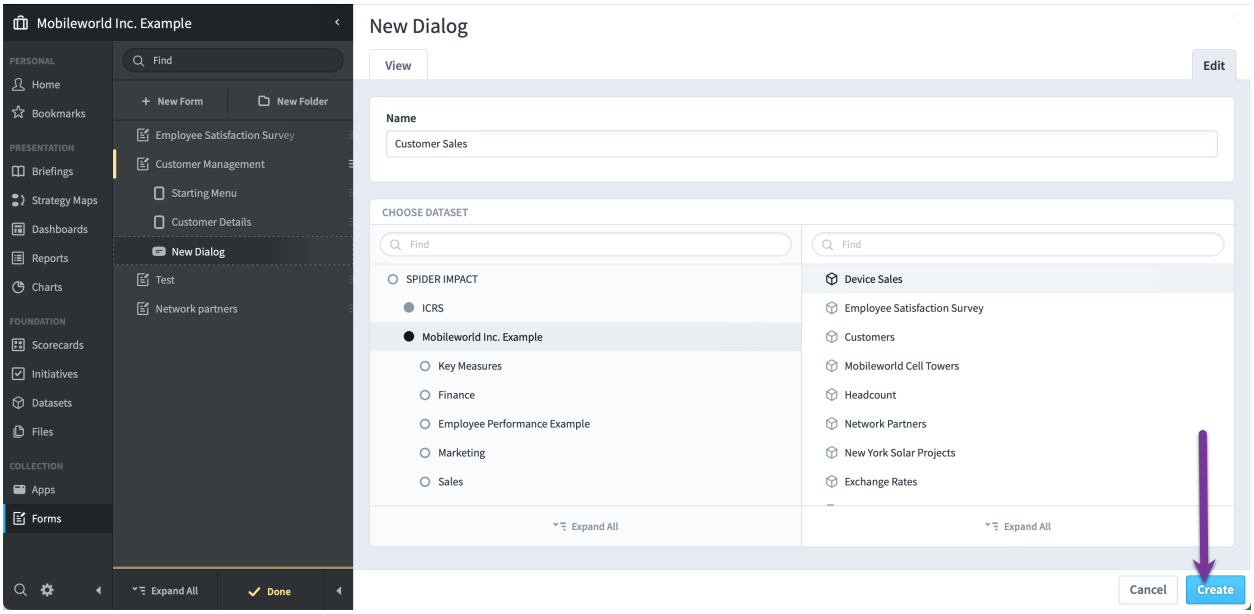

We're now editing out new "Customer Sales" dialog. It's similar to editing a Form Page, and you can add form widgets and set Dialog Display Actions.

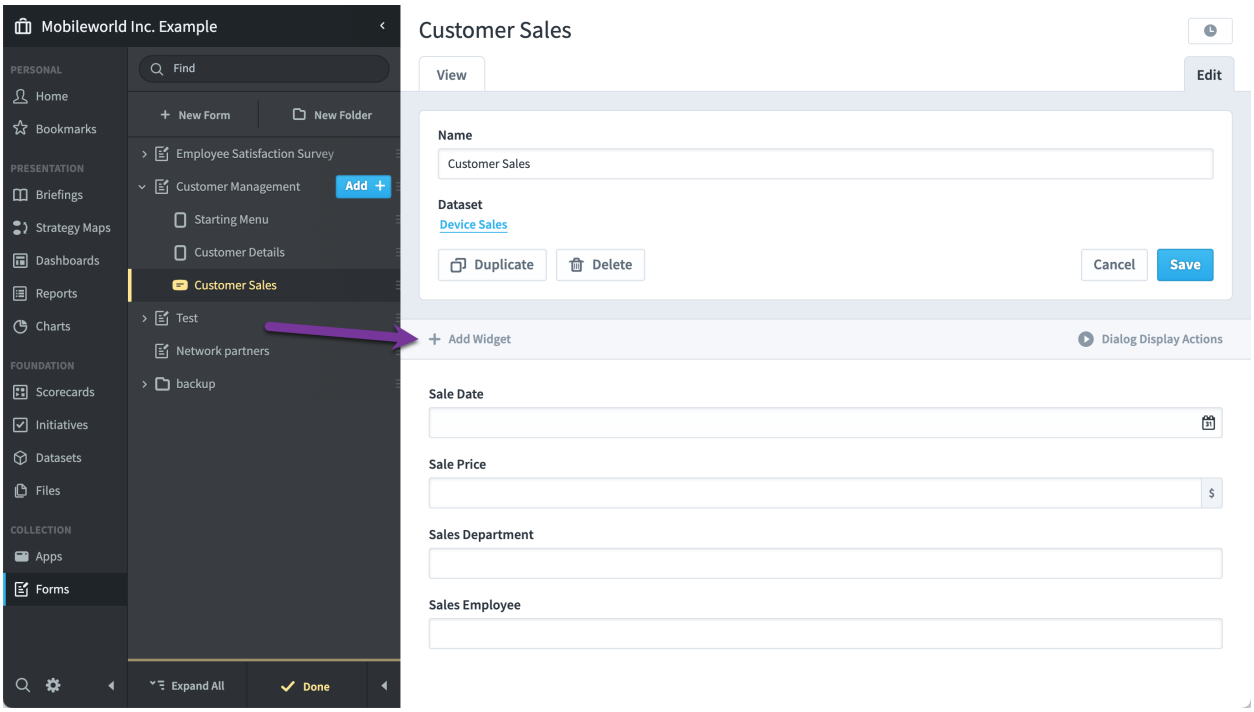

These new Dialogs are used in various ways described below.

New Single Field Widget Options for Linked Fields

Adding a Single Field widget to your form for a linked field is the same as before. First you choose "Single Field" from the "Add Widget" menu.

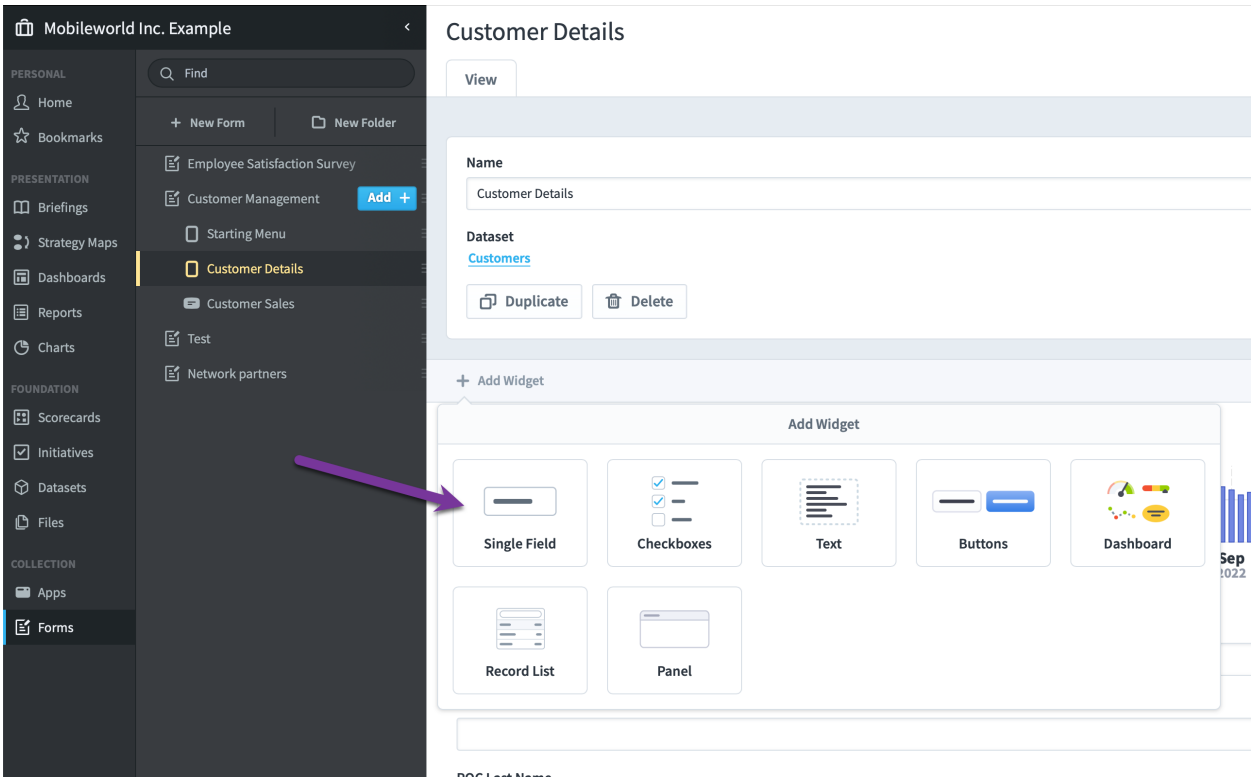

Next, we'll choose a linked dataset, in this situation "Device Sales".

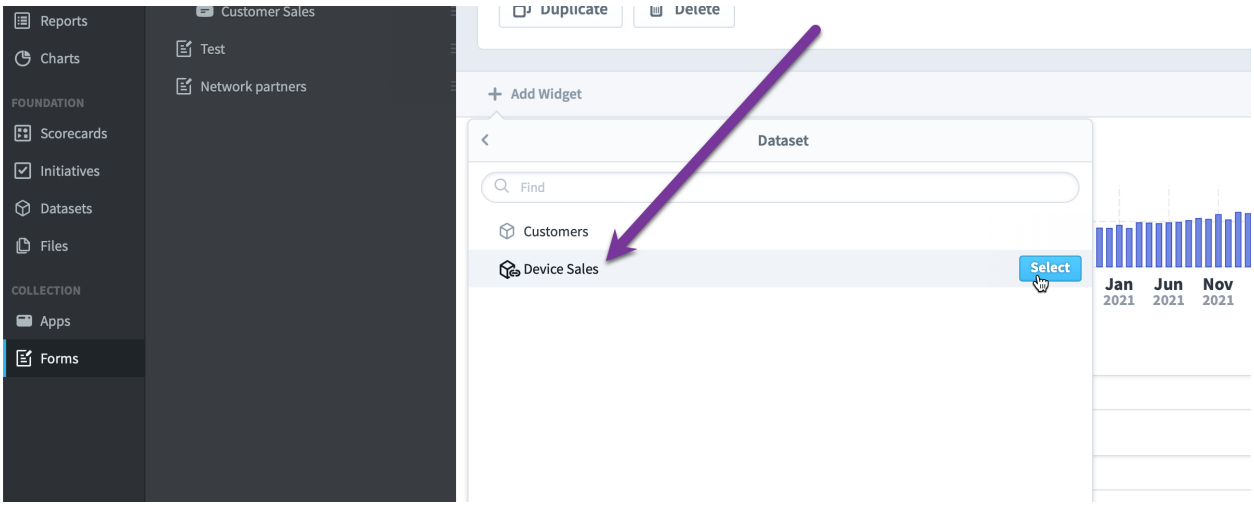

Then we'll choose the "Value Label" field we want to use to identify linked records. In this example we'll choose a calculated field called "Date and Value".

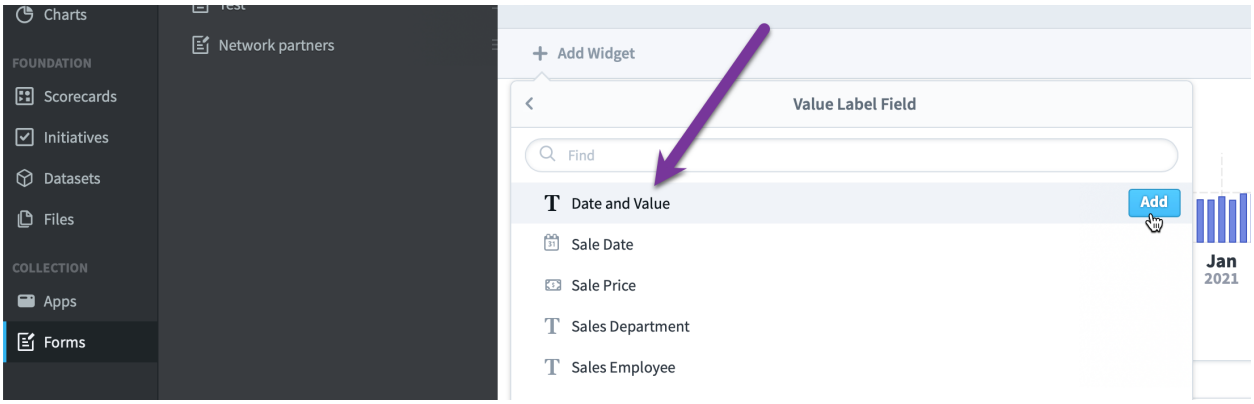

Now, because we've chosen a dataset that links to this form's dataset, we can see the new "Linked Records" panel in the Single Field widget's edit dialog. The form configuration and user interface are slightly different whether the field is on the "Many" or the "One" side of a "One-to-Many" dataset linking relationship.

Dialogs for "One" side of dataset linking

In this example we're seeing the "One" side of a "One-to-Many" dataset linking relationship.

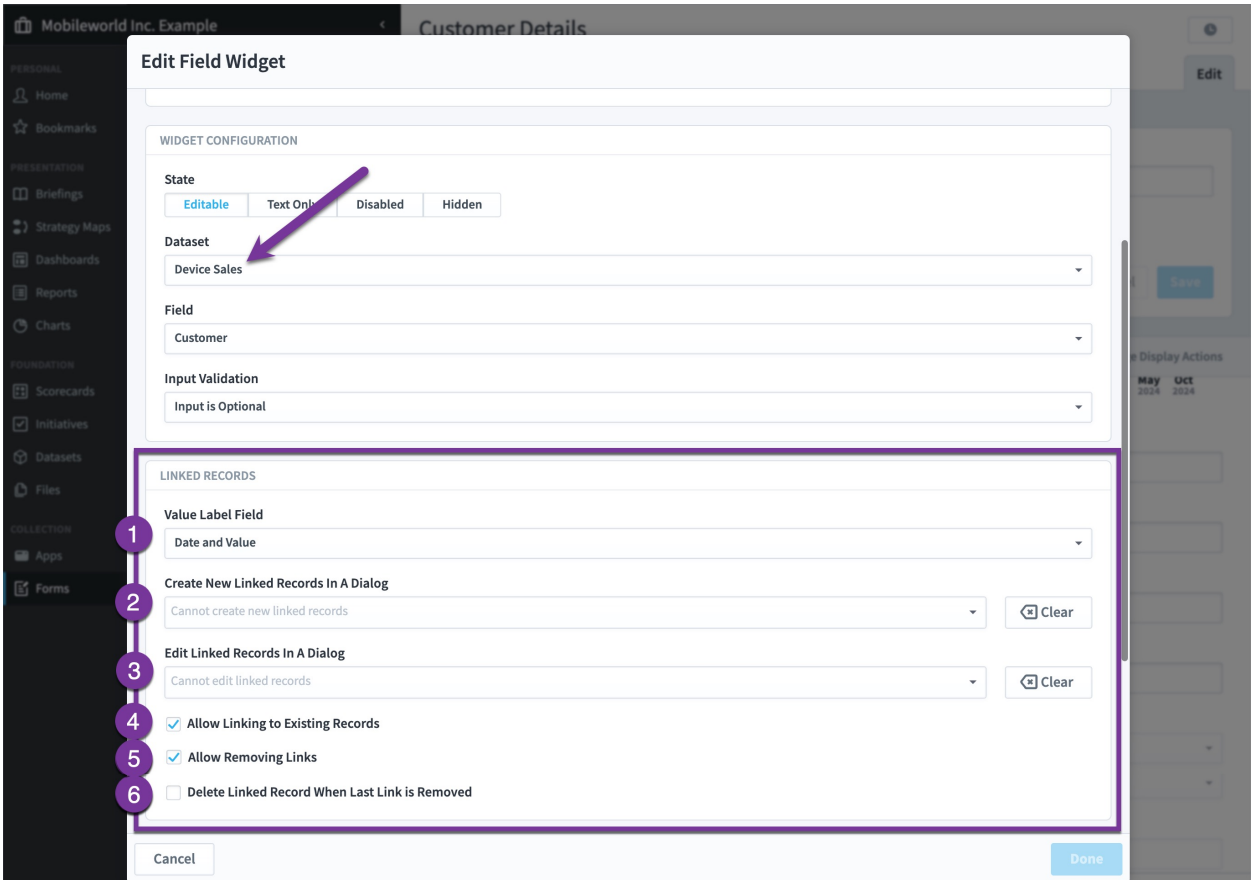

1. **Value Label Field:** This is the field we've already chosen for how we want to identify each linked records.
2. **Create New Linked Records in a Dialog:** A dropdown selection enables users to specify a dialog for creating new linked records. If left unselected, record creation will be disabled for that widget.

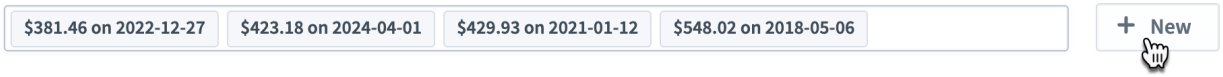

3. **Edit Linked Records in a Dialog:** A separate dropdown allows for configuring dialogs that enable editing existing linked records. If left unselected, record editing will be disabled for that widget.

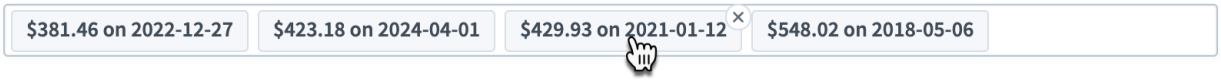

4. **Allow Linking to Existing Records:** A checkbox setting that, when enabled, permits selecting existing records to link to. When disabled, new records must be created directly in the form.

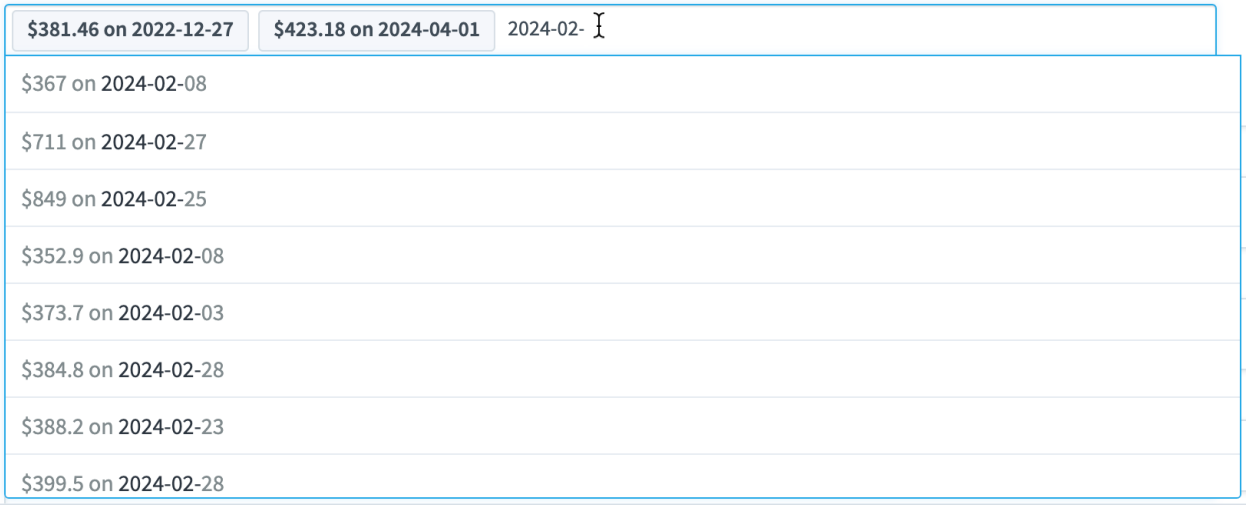

5. **Allow Removing Links:** This allows you to control whether users can remove linked records, independent of users' underlying permissions. This allows you to prevent accidental removals while still allowing users to edit records in dialogs.

6. **Delete Linked Record When Last Link is Removed:** When checked, linked records will be automatically deleted if they are no longer associated with any primary records.

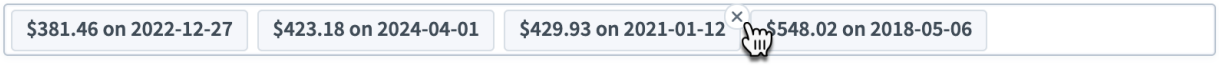

Dialogs for "Many" side of dataset linking

The "Many" side of the "One to Many" dataset linking relationship is similar, except there isn't a checkbox for "Allow Removing Links" or "Delete Linked Record When Last Link is Removed".

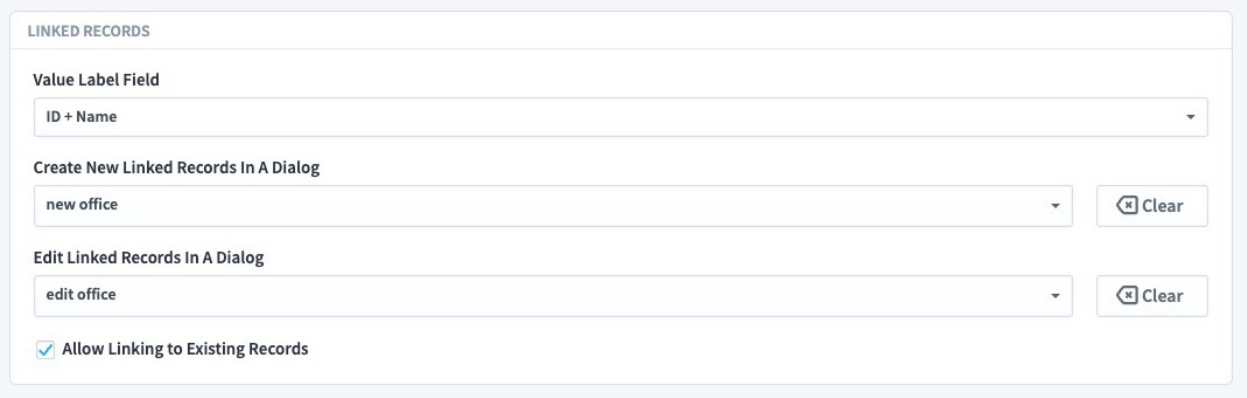

Using the form on the "Many" side of the "One to Many" dataset linking relationship, however, is quite a bit different. When a dialog item is selected in the "Create New Linked Records in a Dialog" dropdown, form users can click "New" to create a linked record in a dialog when no linked record is selected in the form. When a dialog object is selected in the "Create New Linked Records in a Dialog" dropdown, form users can click "New" to create a linked record in a dialog when no linked record is selected in the form.

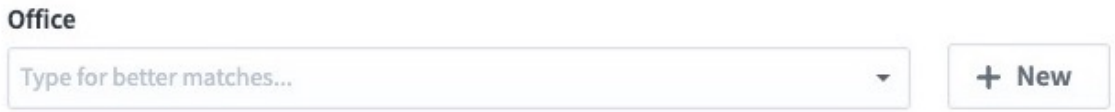

When a dialog object is selected in the "Edit Linked Records in a Dialog", form users can click View to open a dialog with the existing linked record if one is selected in the dropdown.

When "Allow Linking to Existing Records" is unchecked, the dropdown doesn't open, and users can only create new records via the New button - also matching the behavior available on the "one" side of relationships.

Dialog Data Persistence and Permissions

To maintain data integrity, the software handles record saving differently based on configuration. Edits to existing linked records are saved immediately when clicking the dialog's "Save" button.

New linked records follow one of two behaviors depending on the "Delete Linked Record When Last Link is Removed" setting. When this checkbox is not checked, the record is created immediately when clicking the dialog's "Save" button, and calculated field values are retrieved from the server. However, when this checkbox is checked, the record is only created locally and saved to the server when the primary form is saved. During this interim period, calculated value label fields display "[New Record to be Created]" as a placeholder until the actual value can be retrieved from the server.

Permissions remain enforced throughout the process, ensuring users only see and modify records they have access to.

View-only dialog access

Users with view-only permissions can still open dialogs to see linked record details. The dialog displays all fields in read-only format with fields appearing as text-only widgets, respecting the user's permission level while still providing access to information.

Show Dialog form action

Forms now have actions to display dialogs to users for confirmations, warnings, or information.

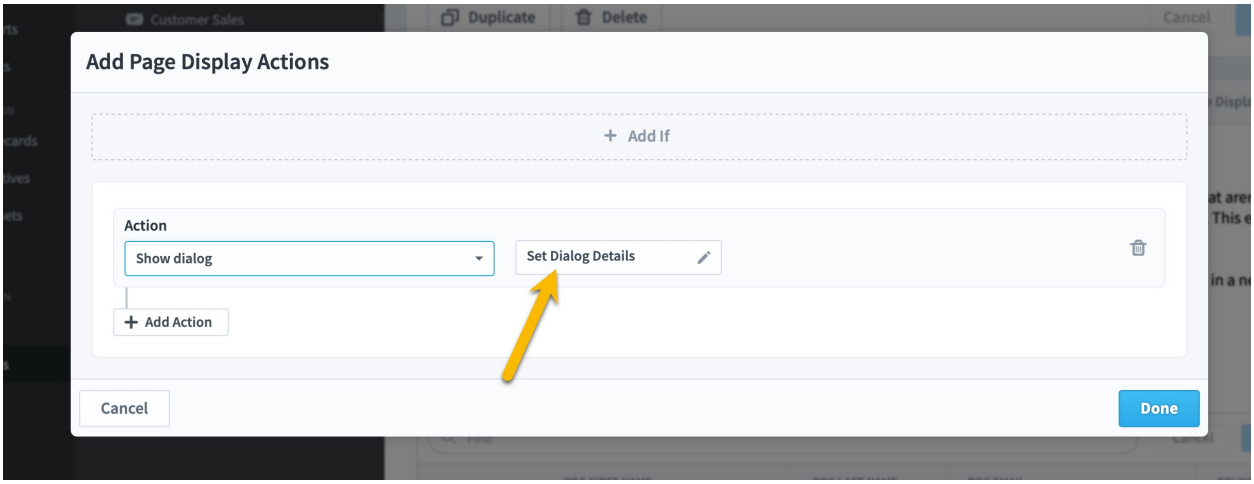

When configuring a Show Dialog action, you select which dialog to display from a dropdown of all available form dialogs, and you set the dialog's title. By default, dialogs only have a single Close button.

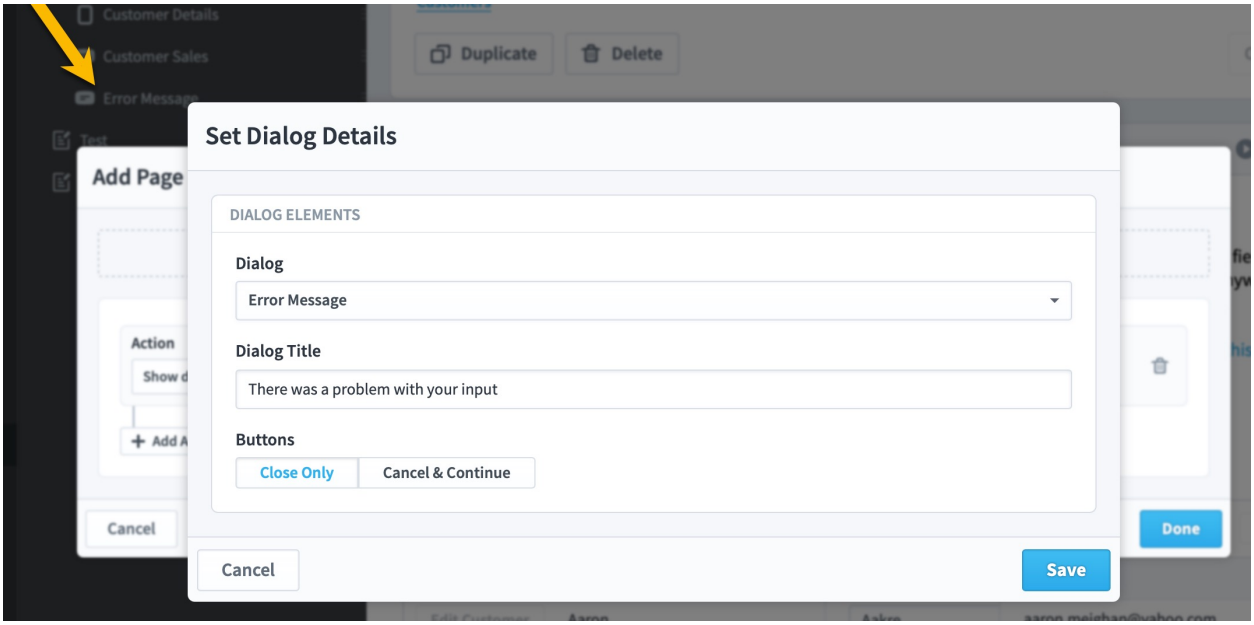

If you choose "Cancel & Continue" buttons, you're able to fire actions when the Continue button is clicked. The Continue button can be customized with a different label and color choice between Blue Standard or Red Warning to match the dialog's purpose. The Continue button supports full conditional logic

and can trigger multiple actions including field updates, record saves, and validation checks.

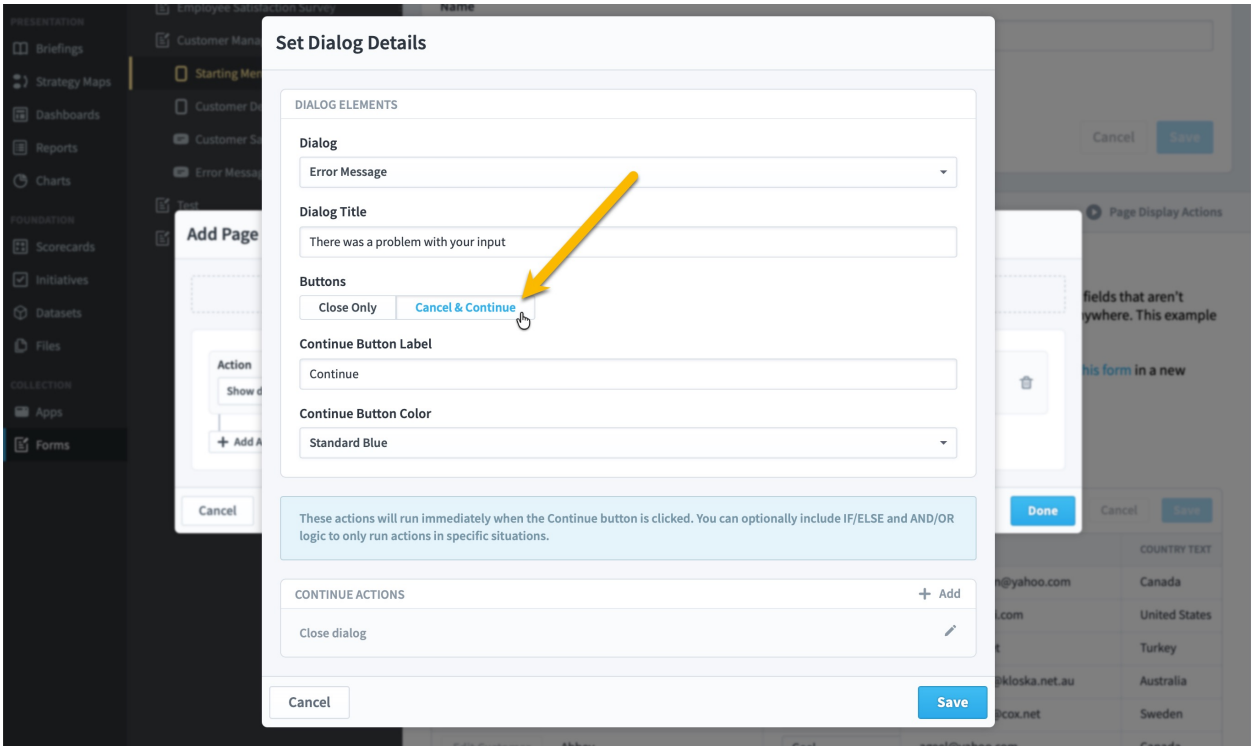

Dialog Continue buttons also have explicit control over when dialogs close. A dedicated "Close Dialog" action is added by default to new dialogs and can be placed within conditional logic, allowing dialogs to stay open if validation fails or close only after successful operations.

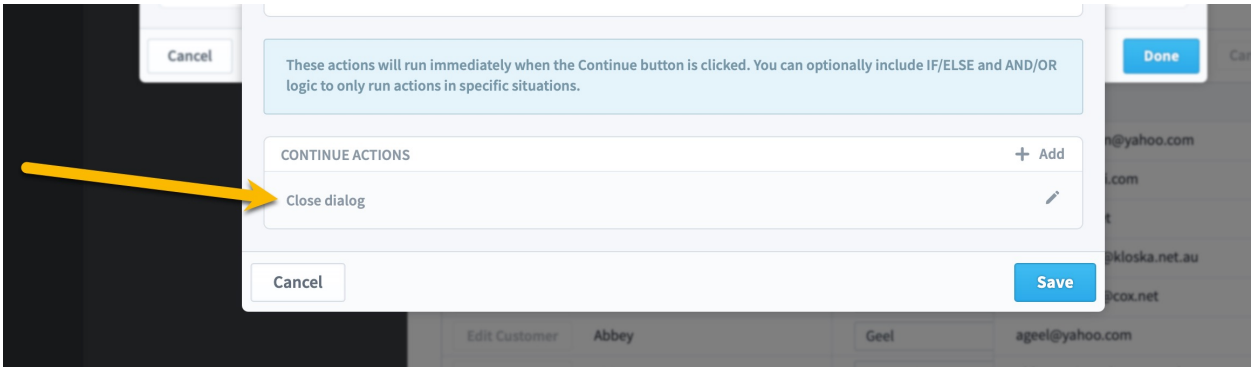

Dialog actions can reference both dialog fields and underlying form page fields, which appear in the same dropdown organized into "Dialog Record Fields" and "Page Record Fields" sections.

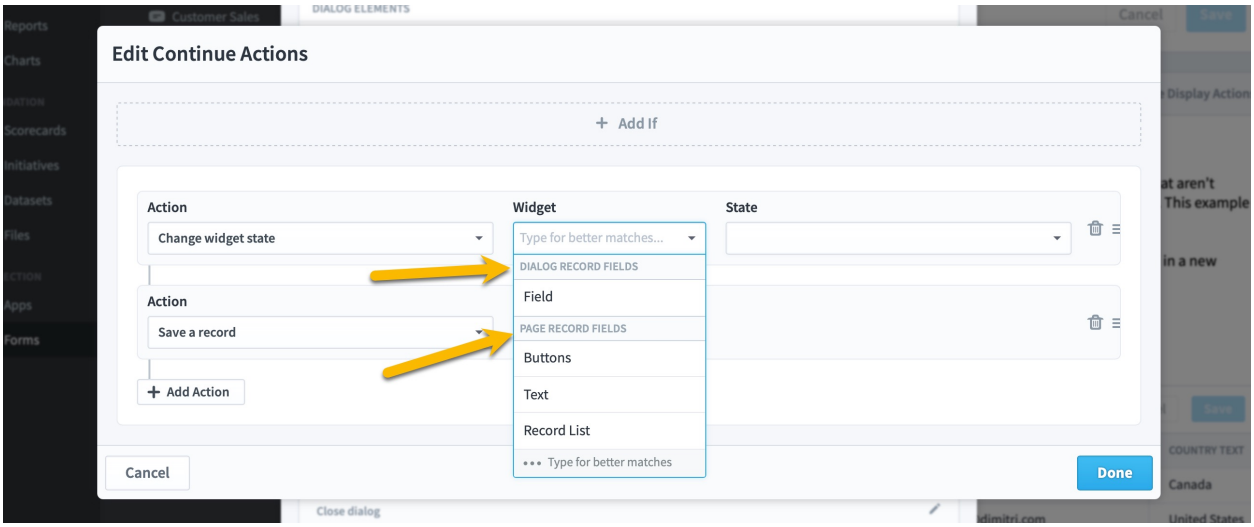

Note that nested dialogs are not supported - "Show Dialog" is not available within dialog Continue Actions.

Example Form Workflows

Auto-loading user records

Forms can default to showing a user's own record without requiring selection. For example, an employee info form can automatically load and display the current user's employee record when they visit the form.

Preventing duplicate submissions

Combine form variables, Load from ID, and conditional logic to ensure users can only submit one record per calendar period:

1. On form load, use Load from ID to check for an existing record matching the current user
2. If found, load that record for editing instead of creating a new one
3. Use equation-based conditions to validate that dates fall within allowed ranges

Creating audit trails

Use Dataset Record variables and the Save a record action to maintain audit logs that capture form activity:

1. Create a Dataset Record variable linked to your audit dataset

2. Use "Set a value" actions to populate audit fields as the form progresses - timestamp with "Current Date", user with "Current User", and record ID with the form's record ID
3. Use equations to dynamically build the action description based on form values - for example, "Updated customer [Customer Name] address from [Original City] to [New City]"
4. When saving the main record, use a second "Save a record" action to save the populated audit record to the audit dataset
5. The audit dataset now contains a detailed log of what changed, when, and who made the change

Dataset Record variables enable forms to manipulate complete records from any dataset without showing them to users. Unlike individual field variables, these store entire records with all their data, enabling sophisticated workflows that operate behind the scenes.

Dataset Record variables enable multi-dataset workflows, where a single form submission can create or update records across multiple datasets simultaneously, and record templating, where forms can load existing records as starting points for new ones.

Complex data validation

Leverage equations in IF statements for multi-field validation:

1. Check that related date fields are in the correct sequence
2. Ensure at least one option from a group of checkboxes is selected
3. Validate that numeric fields sum to an acceptable total
4. Display specific error messages in dialogs that remain open until issues are resolved

Dynamic form behavior based on permissions

Combine permission filtering with form logic:

1. Rollup fields automatically show only the portions of the hierarchy users can access
2. When users have permission to only one item, it's pre-selected

3. Form behavior adapts based on whether users have view-only or edit permissions
4. Dialogs open in appropriate modes based on user access levels

Automatic dropdown filtering based on permissions

Form widgets for text, number, and rollup fields now automatically filter dropdown values based on user permissions, preventing users from seeing or selecting values they don't have access to. This ensures users cannot create records they wouldn't be able to view later.

For text and number fields, when users have permission to view only one value, that value is automatically pre-selected. For rollup fields, the tree is pruned to display only items the user has permission to see along with their ancestors (even if those ancestors aren't directly viewable themselves), with only the viewable items being selectable. When a rollup field has only one selectable option, it's automatically pre-selected. Date, datetime, time, and yes/no fields don't use dropdowns and therefore don't include this auto-selection behavior.

Performance improvements for large forms

Forms with large amounts of data now load and respond significantly faster. This includes forms with record list widgets containing many rows and columns, as well as forms with linked dataset fields that reference large datasets.

Previously, complex forms could take 10-20 seconds to become responsive, with some users experiencing browser freezes or crashes. These same forms now load in just a few seconds, even on slower computers.

The improvements are especially noticeable in these scenarios:

- **Record list widgets with many editable fields** - Forms that display large record lists with dozens of columns and rows now load much faster. Scrolling through these lists is also smoother, with data appearing immediately rather than after several seconds of waiting.
- **Linked dataset fields with thousands of options** - When clicking into a field that's linked to a large dataset (with tens of thousands of records), the field now responds instantly. Previously, these fields could take 4-10

seconds just to show a cursor and sometimes caused the page to freeze completely.

- **Record list widgets with file upload columns** - Record lists containing multiple file upload columns, which were particularly slow to load and interact with, now perform smoothly.

These performance gains were achieved through several technical optimizations including more efficient column width calculations, lazy loading of data, and smarter memory management. Users on government or older computers, who were most affected by the previous performance issues, will see the biggest improvements.

Datasets

New File Data Type for Dataset Fields

There is a new "File" data type for dataset fields, expanding the types of information that can be stored and managed within datasets. This enhancement enables users to upload, download, and manage file attachments directly within dataset records, improving data organization and accessibility.

Files in Datasets

To store files in datasets, just create a new dataset field and choose "File" for the data type.

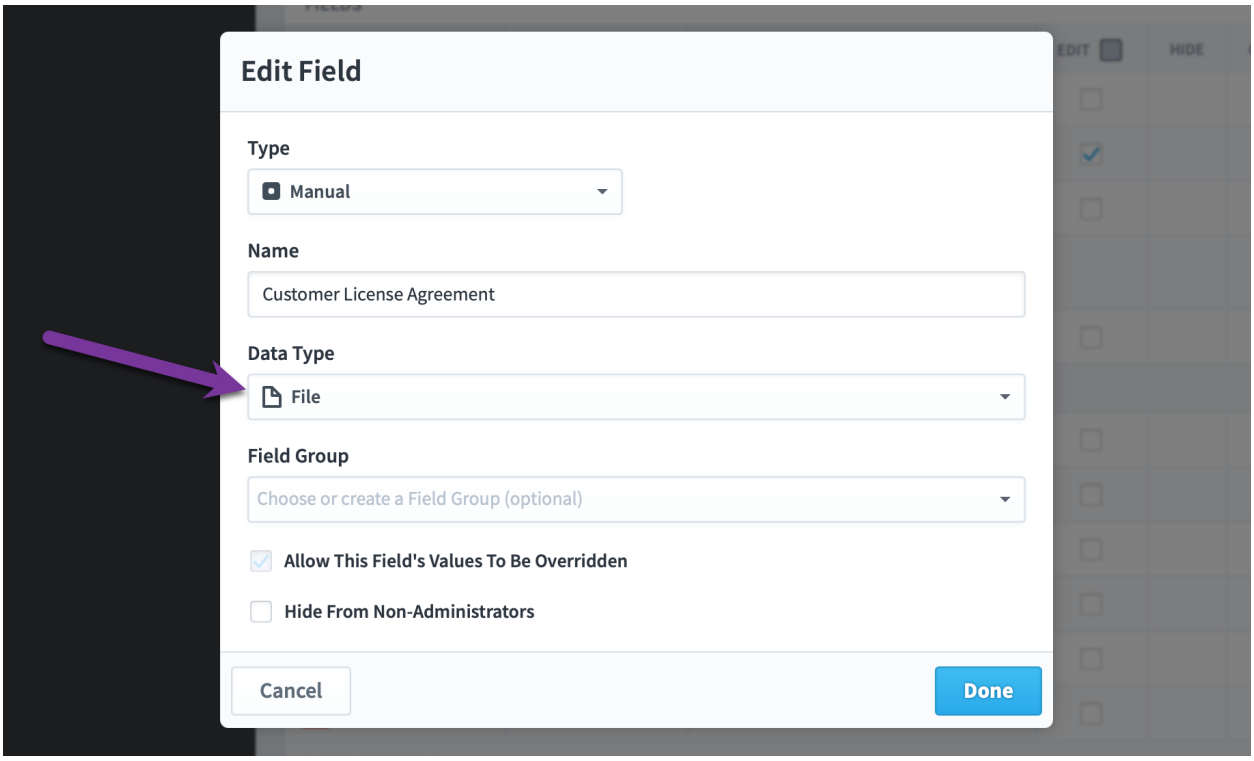

Now, when we edit the record on the Datasets Records tab you an upload a file if you have the correct permissions.

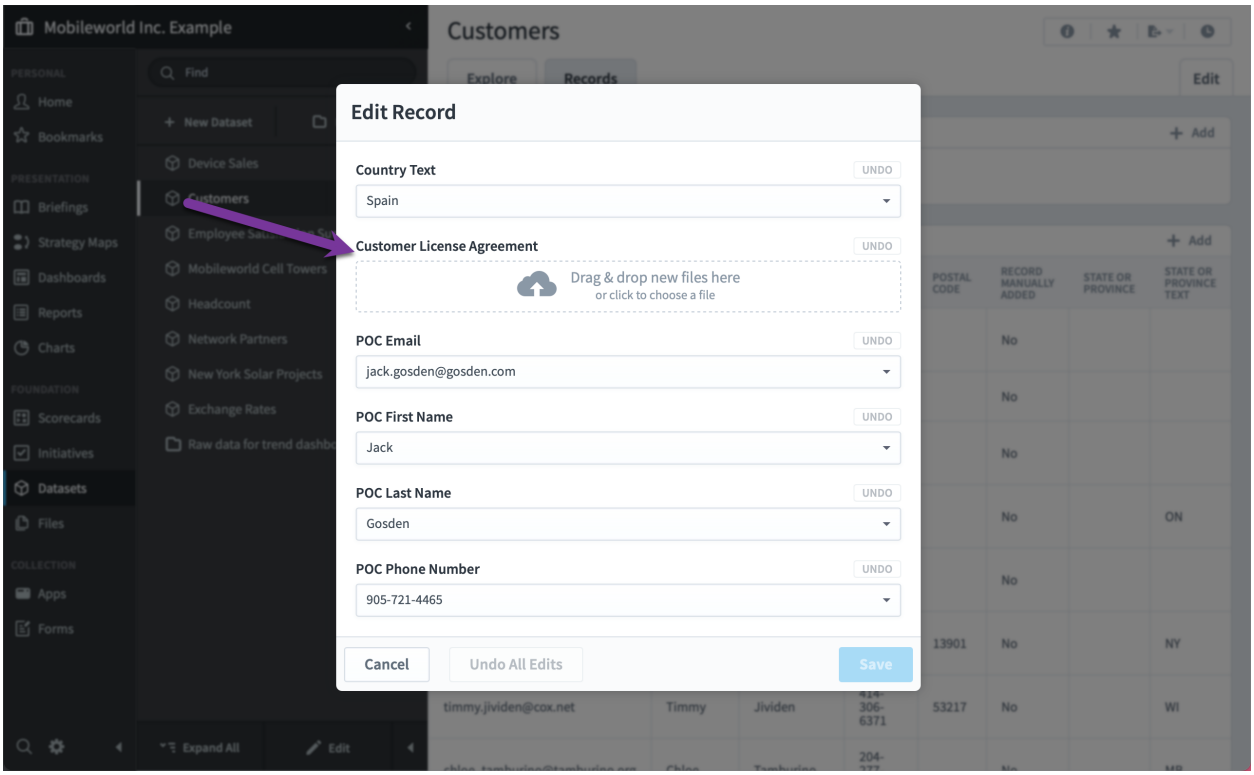

When a field has a file uploaded, you can either download that file or delete it.

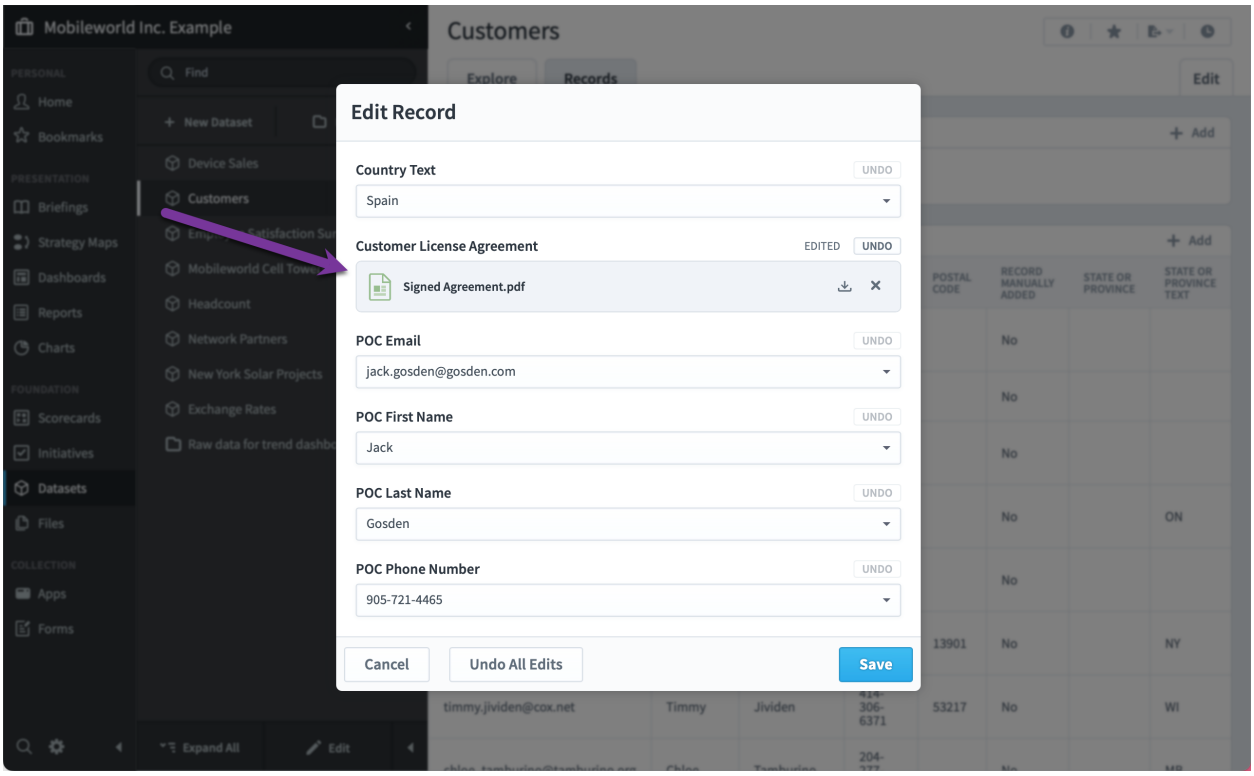

Files in Forms

Dataset files are the same when their fields are shown in a Single Field Widget in forms.

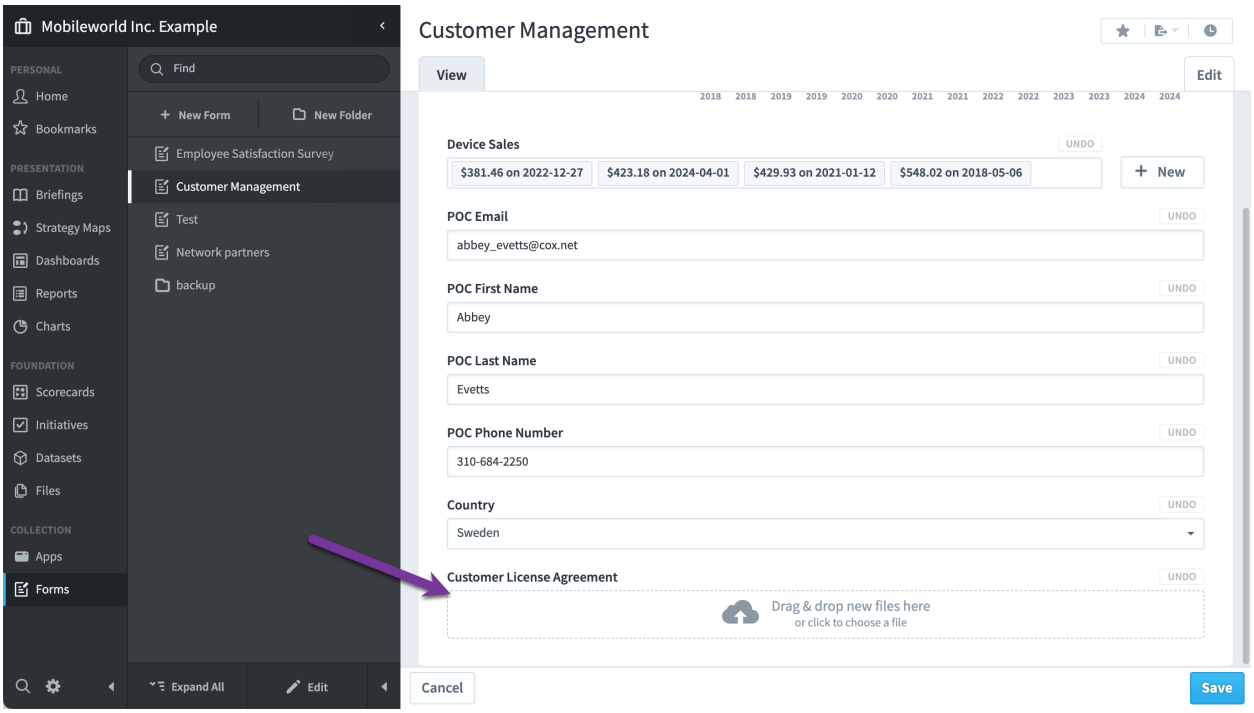

You can also show file fields on Record List Widgets.

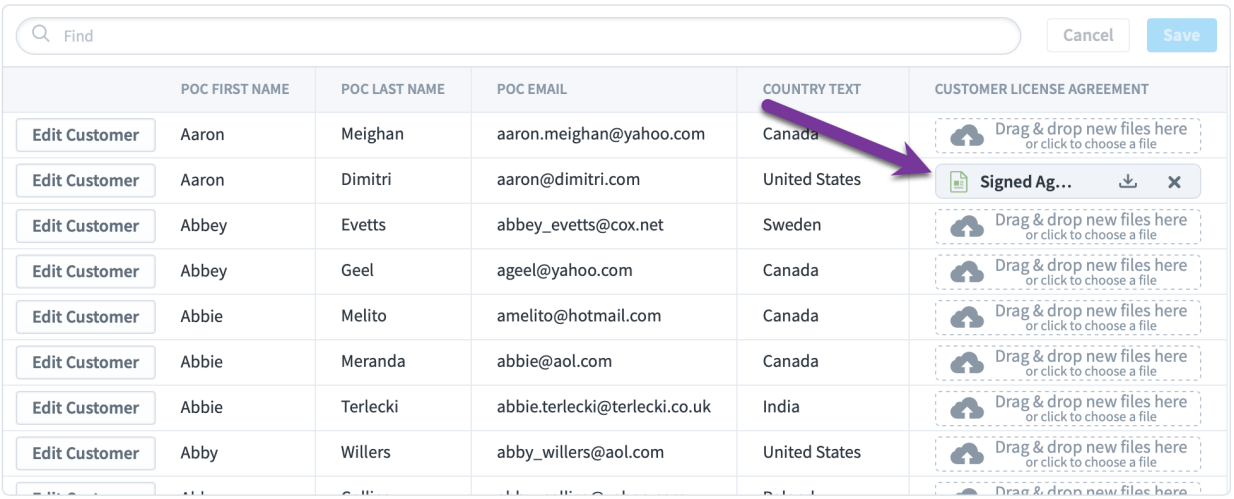

Files in reports

When file fields are shown in Reports, you can click on the file name to download it.

CUSTOMER NAME	POC EMAIL	COUNTRY	CUSTOMER LICENSE AGREEMENT
Jack Gosden	jack.gosden@gosden.com	Spain	Signed Agreement.pdf
Aaron Dimitri	aaron@dimitri.com	United States	Signed Agreement.pdf

Aggregation and Display in Visualizations

File fields can be aggregated in data visualizations using the "Number of Unique" aggregation type. This allows users to count the distinct file attachments across records, providing insight into data completeness or redundancy in dashboards, charts, and the Datasets Explore tab.

No File Type Restrictions

Dataset file uploads are supported universally without restrictions on file types.

Changing the Dataset Record ID Column

You can now change the Record ID Column for datasets with editable fields. Previously, this was prevented, requiring users to rebuild datasets and recreate forms when a primary key change was needed. Now, users can update the Record ID column within dataset settings, provided the change does not interfere with linked records, references in forms, or other dataset constraints.

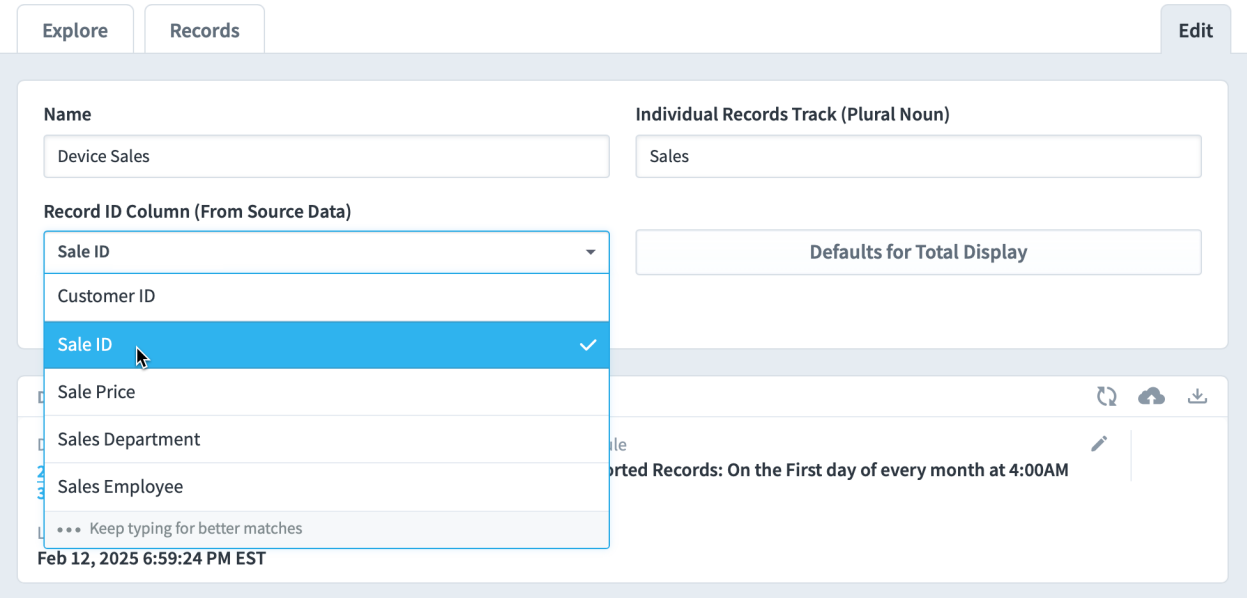

Linked fields can be used for rollup tree group permissions

Previously, when configuring rollup tree group permissions for datasets, you could only select fields that existed directly in the dataset. This limitation meant that if you wanted to set permissions based on a rollup tree field from a related dataset, you had to create a calculated field that referenced the linked field, which added unnecessary complexity.

Now, visible linked fields from related datasets can be used directly when setting up rollup tree group permissions. When you add a visible linked field to your dataset that points to a rollup field in another dataset, that linked field will appear as an option in the rollup tree group permissions dropdown.

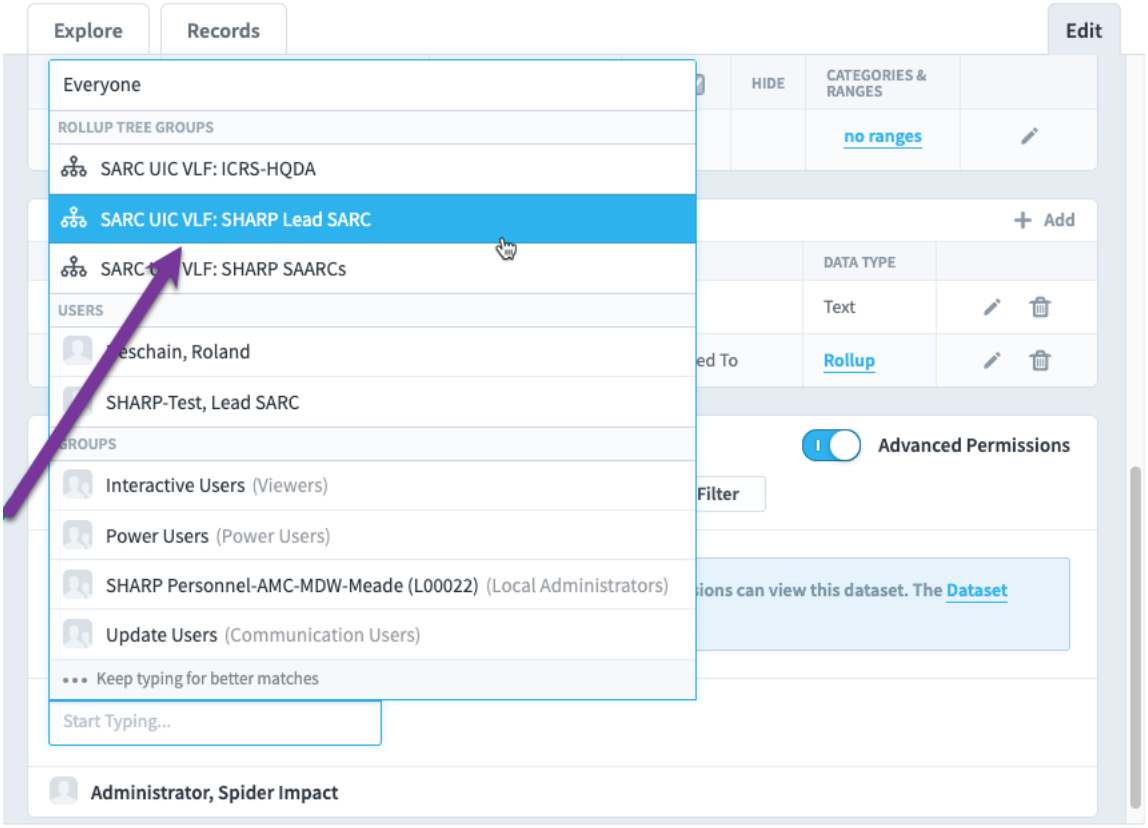

This enhancement simplifies permission management in scenarios where your organizational structure or hierarchy is maintained in a separate dataset. For example, if you have a central dataset containing your organization's structure with rollup trees, you can now link to that structure from other datasets and use it directly for permissions without needing intermediate calculated fields.

Missing Value Treatment for Dataset Fields

Dataset equations have been updated to provide more flexible handling of missing values. Users can now define how missing values should be treated when creating a calculated field, with three available options:

- **Any missing value makes the entire equation blank** – This maintains the historical behavior, where any missing value results in a null output.
- **Treat missing values as 0** – This option replaces missing values with zero in calculations.
- **Treat missing values as Blank** – The default setting for new fields, this option ensures that missing values are ignored when possible, such as in `avg()` equations or when concatenating text using `+`.

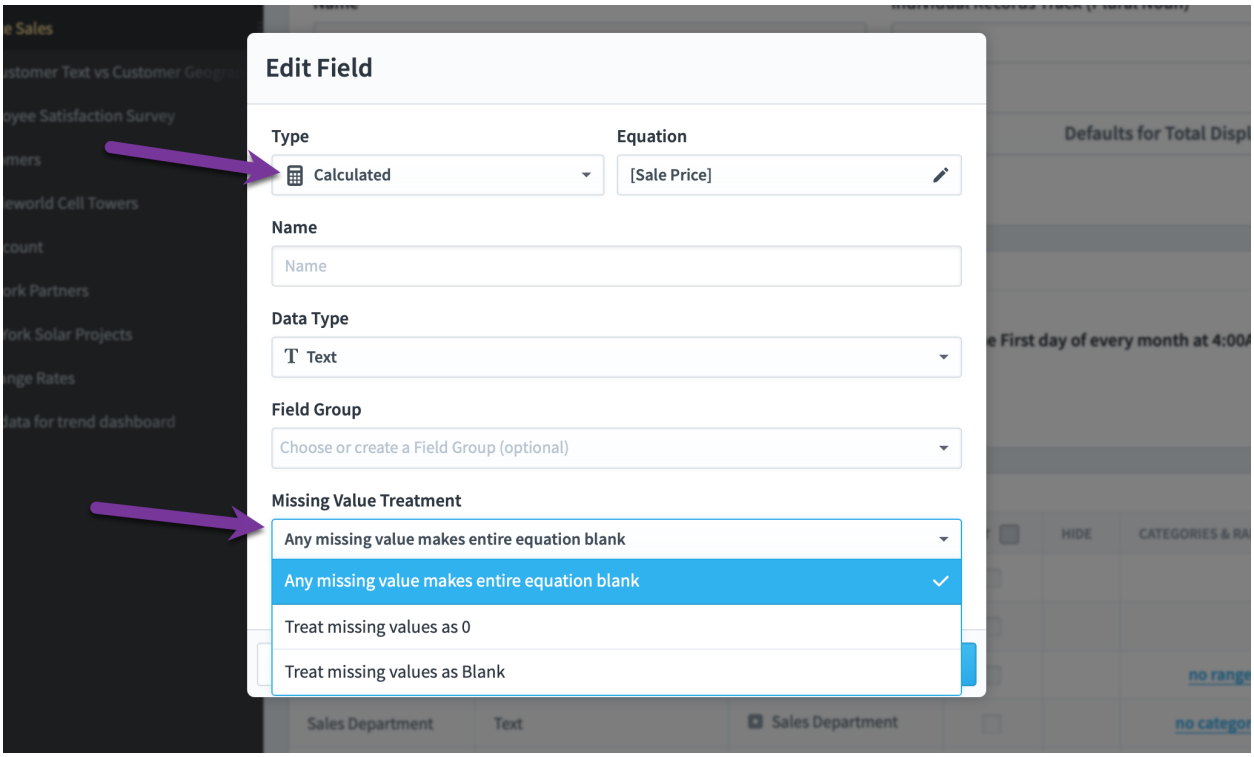

For all existing calculated fields, the default behavior will be set to "Any missing value makes the entire equation blank" to ensure consistency with past behavior. New calculated fields, however, will default to "Treat missing values as Blank," reflecting a more intuitive approach to handling missing data in most scenarios.

Importing and Exporting

Dashboard aspect ratio control

Dashboards now offer precise control over their aspect ratios, making it significantly easier to create presentations that export perfectly to PowerPoint. The default Aspect Ratio of "Auto" continues to automatically adjust dashboard height and width based on where you place widgets.

If you choose an aspect ratio of "Widescreen (16:9)" or "Standard (4:3)", the dashboard canvas continues to adjust in size based on your widget placement, but its overall height-to-width proportions are locked in place. As you drag widgets to expand the dashboard, both dimensions adjust proportionally to preserve your chosen aspect ratio. Here we've chosen Widescreen and now see empty space on the right of the dashboard.

This eliminates the trial-and-error process of creating dashboards for specific presentation formats. We're able to quickly rearrange our widgets to best take advantage of this specific aspect ratio.

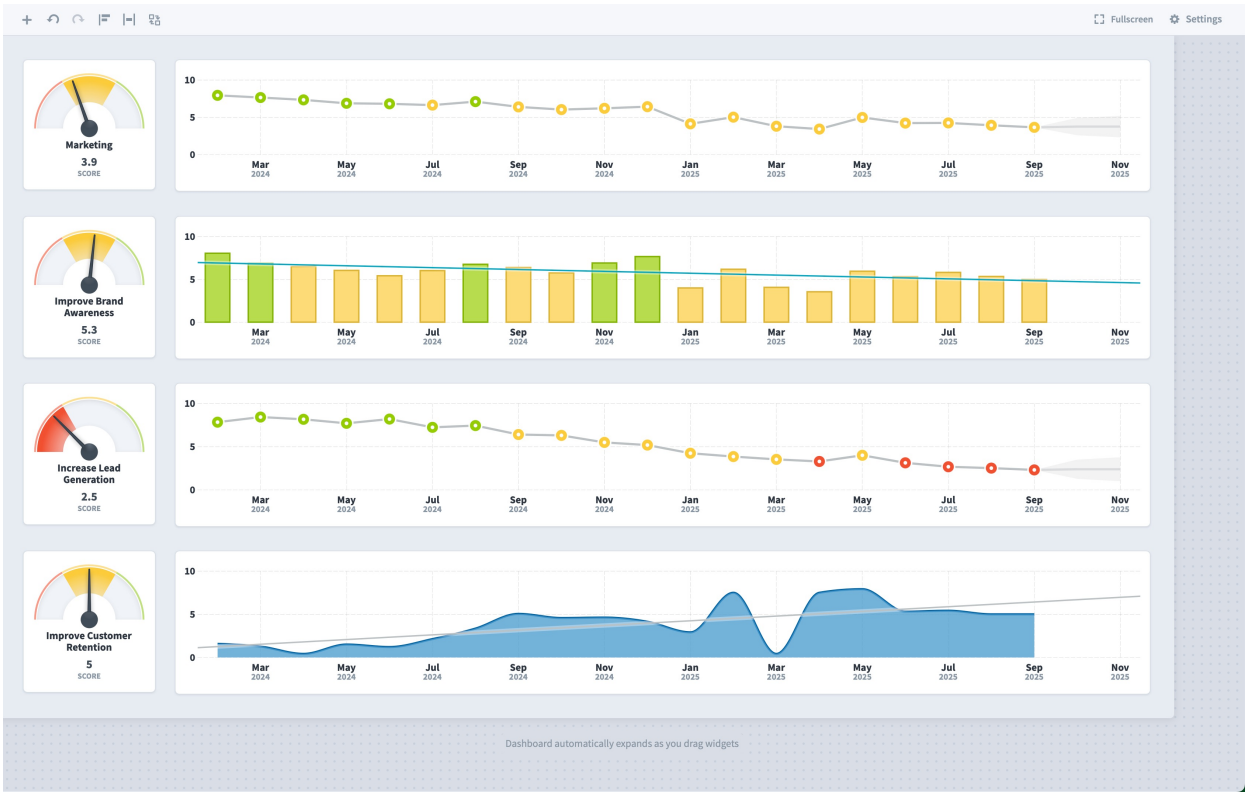

And when we export to PowerPoint, our dashboard fits perfectly on a widescreen PowerPoint slide without any empty space around the edges.

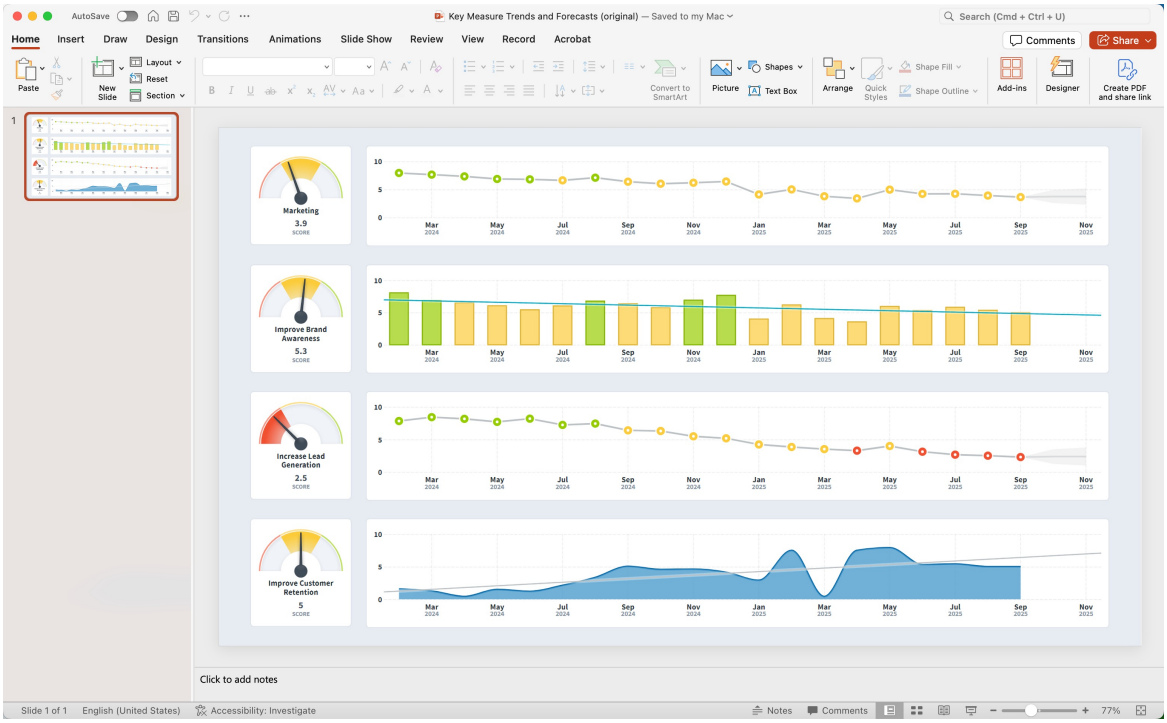

Configurable dashboard edge padding

Dashboards have traditionally displayed widgets in a panel-style layout with automatic padding on all sides of the dashboard. This padding provides comfortable spacing around dashboard content, which works well for standard data displays.

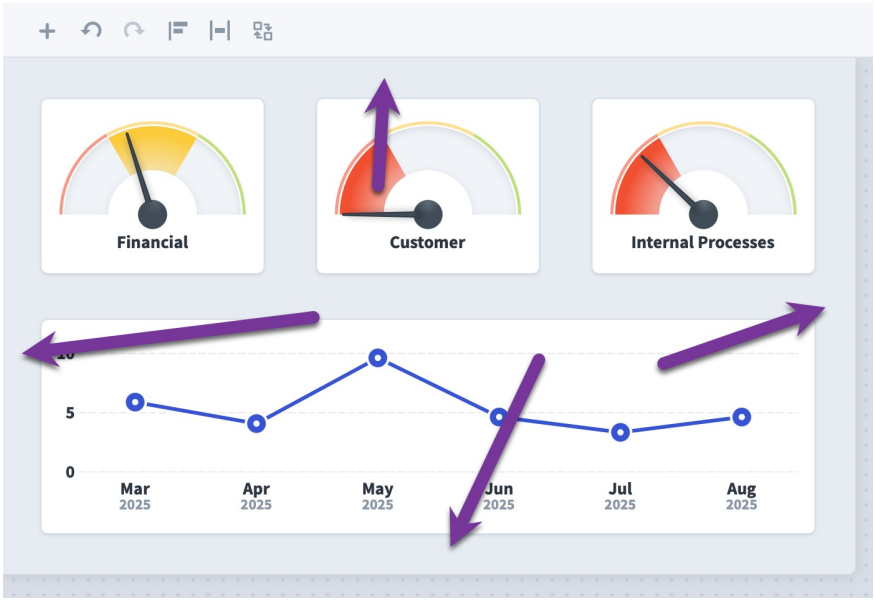

However, when adding background images to create more visually engaging dashboards, this automatic padding around the dashboard creates unwanted space between the background image and the dashboard edges.

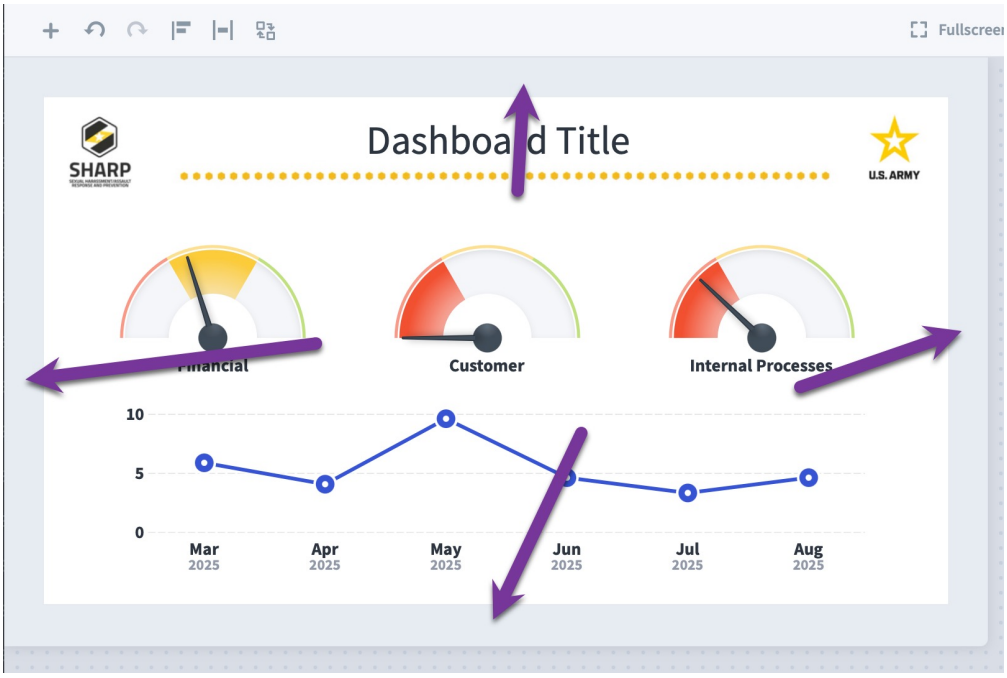

A new "Edge Padding" toggle in dashboard settings now gives you complete control over this internal padding. Simply turn off Edge Padding to allow widgets and backgrounds to extend all the way to the dashboard edges. With Edge Padding disabled, you can create immersive, edge-to-edge dashboard designs perfect for presentations and displays with full-bleed backgrounds.

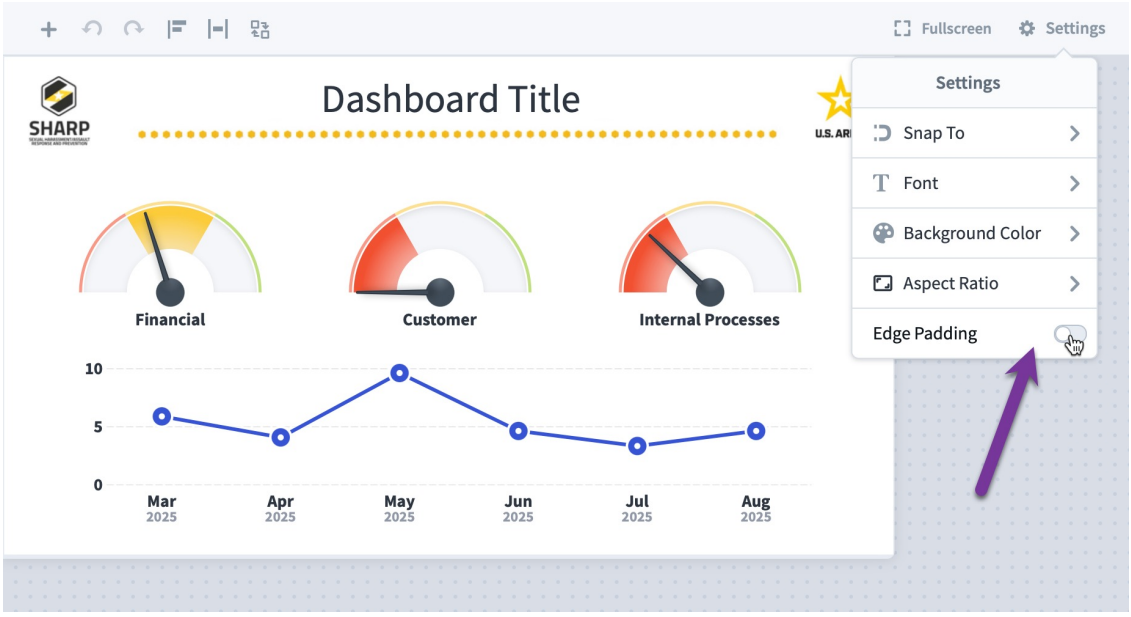

Removing dashboard edge padding is particularly noticeable when exporting dashboards with background images to PowerPoint. There is no longer padding around the background image.

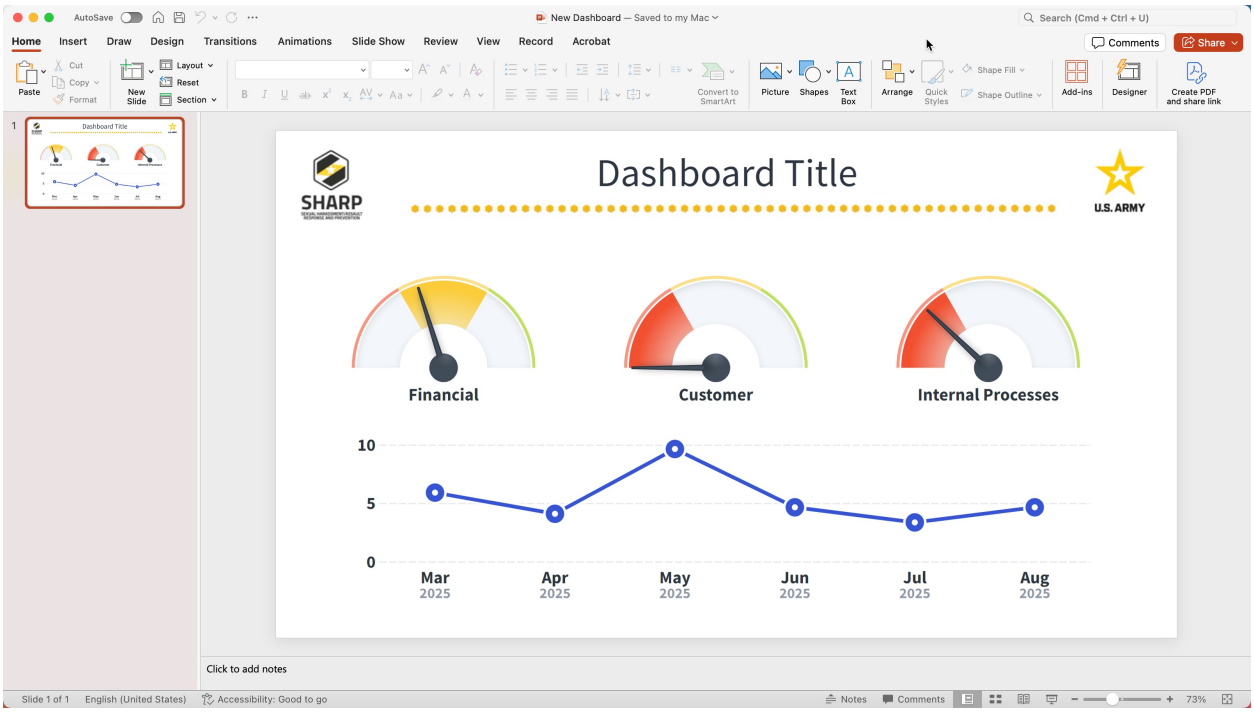

Edge Padding is enabled by default to maintain the traditional layout for existing dashboards. Additionally, when exporting dashboards with white backgrounds, the software now intelligently handles spacing—adding white margins only when needed for aspect ratio adjustments rather than adding both gray outer margins and white inner margins. This creates cleaner, more professional exports that look exactly as intended in presentations.

Enhanced PowerPoint and PDF export options

Exporting Dashboards and Briefings to PowerPoint and PDF now offers more control over the final output.

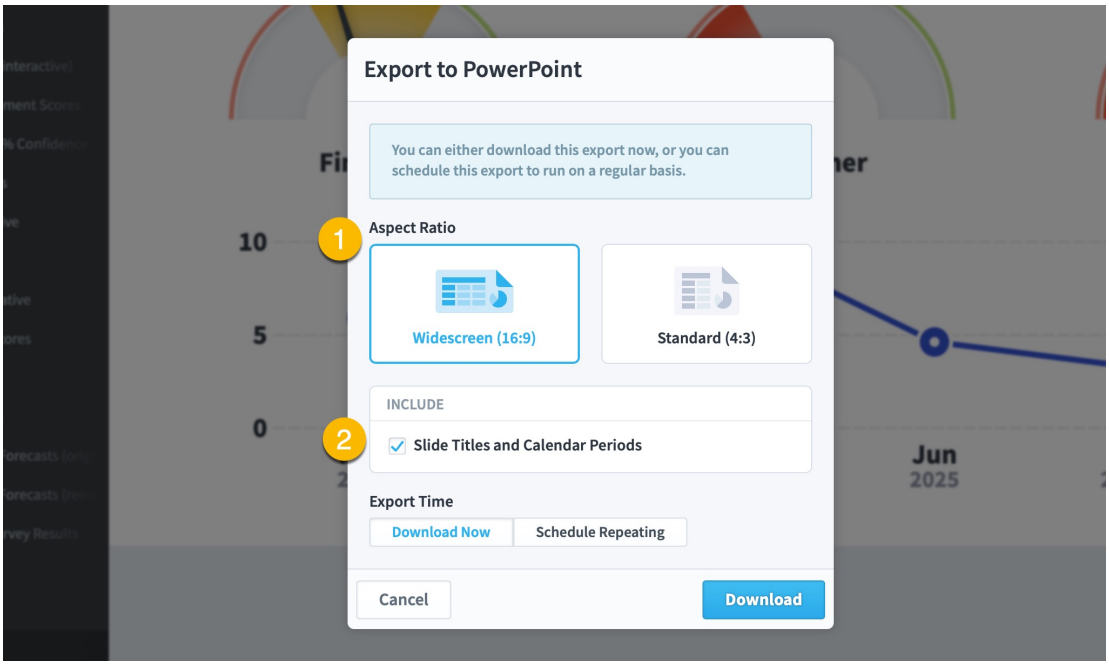

1. For PowerPoint exports, you can choose between Widescreen (16:9) or Standard (4:3) aspect ratios, with Widescreen as the new default to match modern presentation standards.
2. Both PowerPoint and PDF exports now include a "Slide Titles and Calendar Periods" checkbox. When enabled, a header bar appears at the top of each slide showing the slide title and calendar period. You can now turn off this bar, giving you a clean export with just your dashboard content.

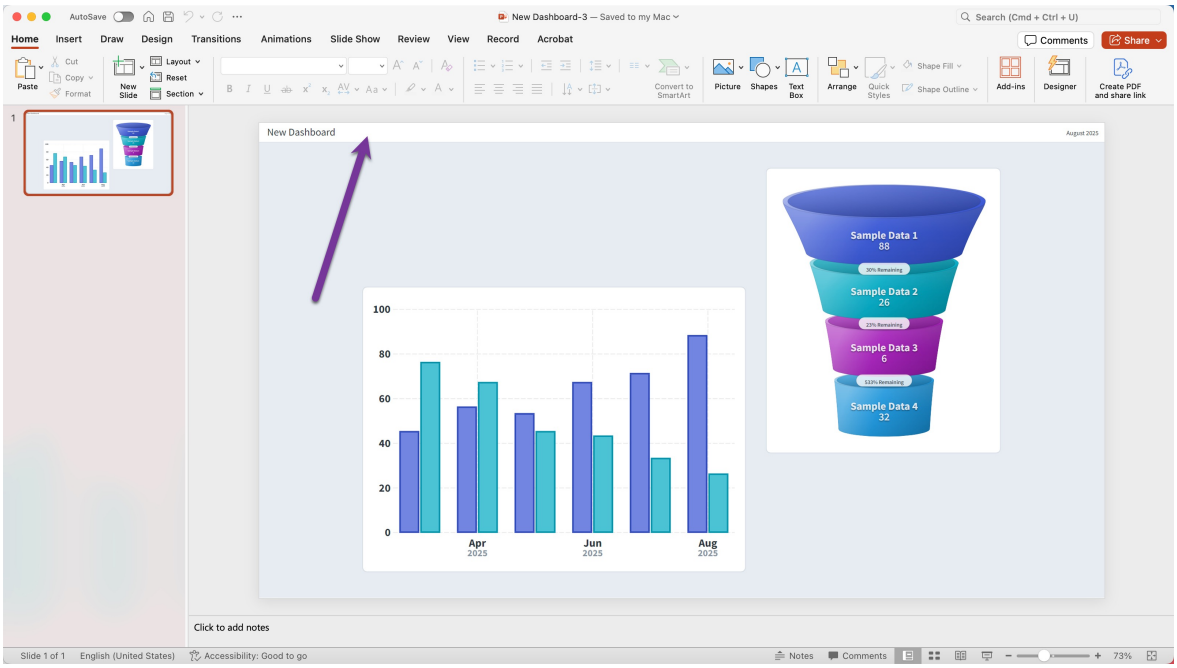

Cleaner Excel exports

Excel exports no longer include the title and blank row at the top. This makes exported files ready for immediate use in data analysis, re-importing, or integration with other systems. In this example, the first two rows will no longer appear in Excel exports.

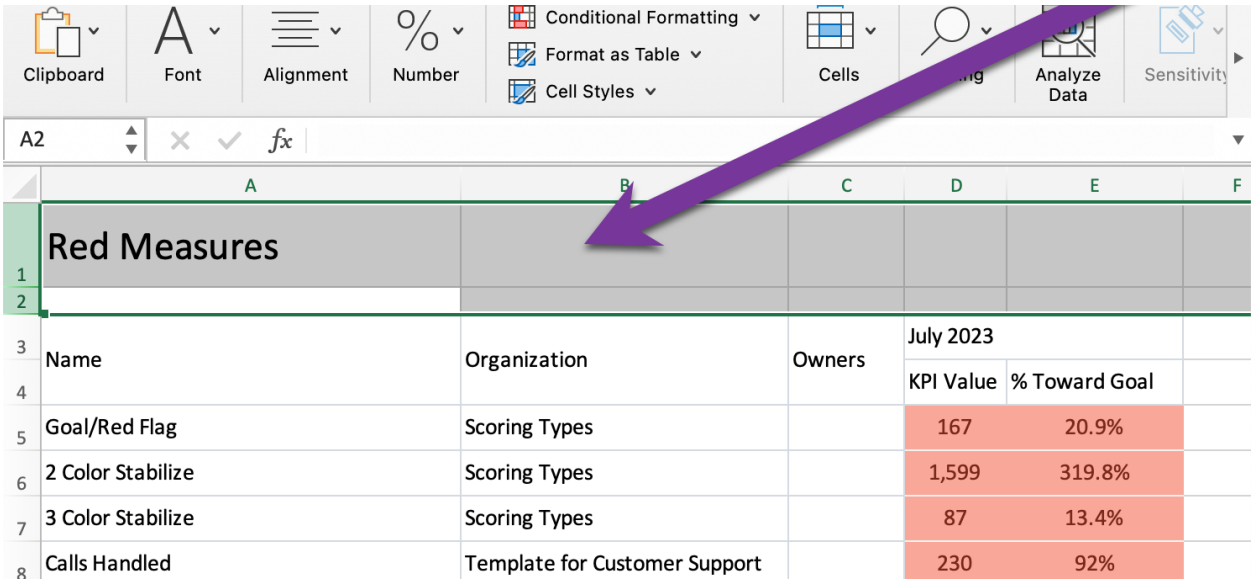

Import from shared Google Sheets

Spider Impact now allows you to import data from Google Sheets that are shared with your Google account, even when they're owned by other users. Previously, when creating import connections to Google Sheets, you could only see and select from sheets that you personally owned.

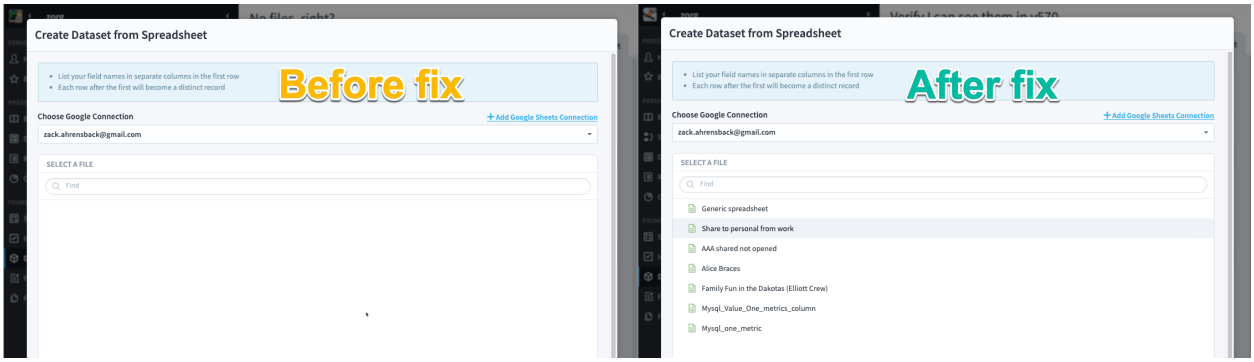

If you have an existing Google Sheets import connection, you don't need to recreate it. The shared sheets will automatically appear as available options the next time you access the connection.

Administration & Performance

Configurable thousands separator for number formatting

Organizations can now configure how numbers are displayed throughout Spider Impact, choosing between automatic detection based on browser settings, or explicitly setting comma/period usage for thousands and decimals. This new Thousands Separator setting is found in Application Administration > Data Format.

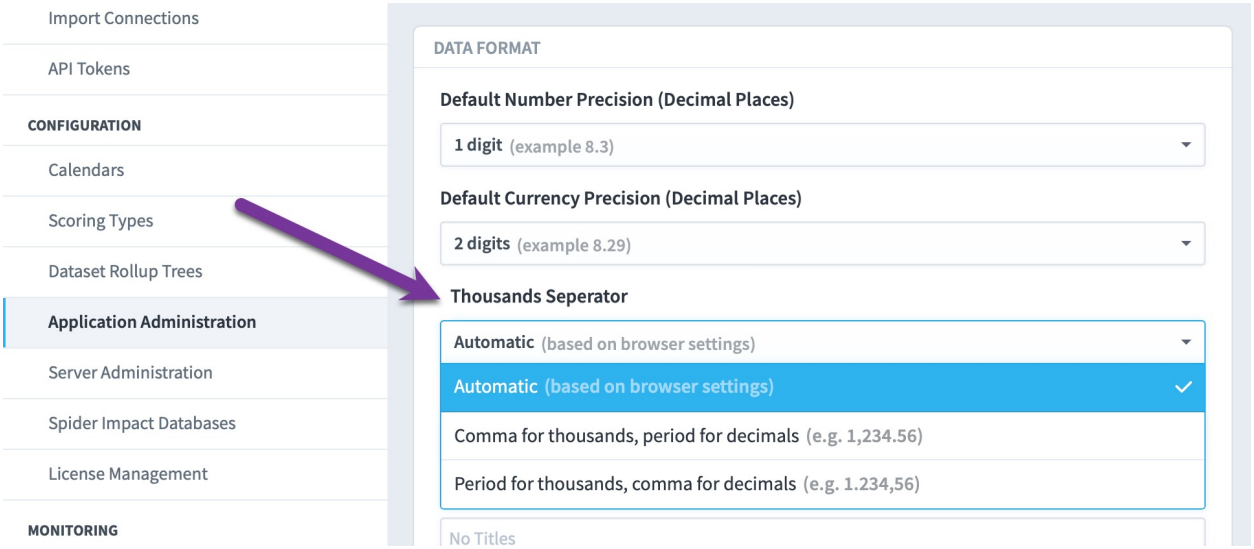

The three options are:

- **Automatic (based on browser settings)** - The default option that uses the browser's locale settings
- **Comma for thousands, period for decimals** (e.g. 1,234.56) - Standard US formatting
- **Period for thousands, comma for decimals** (e.g. 1.234,56) - Common in many European and Middle Eastern regions

This setting is particularly helpful for organizations with users in the Gulf states (Saudi Arabia, UAE), where period decimal separators are standard despite Arabic language settings that would typically default to commas. Organizations with users from multiple regions can set a consistent number format that works best for their needs.

New currency symbols and types

Spider Impact 5.8 adds support for new and specialized currencies to better serve international organizations:

Saudi Riyal symbol - The new official Saudi Riyal currency symbol announced by the Saudi Central Bank in February 2025 is now displayed throughout the application wherever Saudi Riyal (SAR) values appear. The software uses a custom font to display this symbol until it becomes part of the Unicode standard. When copying text containing the symbol to external applications, it automatically converts to "SAR" to ensure compatibility.

Unit of Account currency - A new "Unit of Account" (UA) currency type has been added to support international organizations that use non-national currencies for financial reporting. This currency type, commonly used by development banks and international financial institutions, helps organizations report in a stable unit that isn't subject to individual currency fluctuations. The Unit of Account appears in the currency selection list with "UA" as the currency symbol and XUA as the currency code.

Faster Score Calculation

This release brings major improvements to score calculation speed and efficiency.

- **Optimized Processing:** Calculations now prioritize recent periods, reducing unnecessary recalculations.
- **Smarter Updates:** Only changed scores are rewritten, cutting down on redundant database writes.
- **Parallel Execution:** Some score computations can now run simultaneously, speeding up processing.
- **Improved Prioritization:** Frequently stable metrics are now calculated first to avoid bottlenecks.

These enhancements lead to a 20-50% performance boost for metric updates, making the system faster and more responsive.

Automatic Database Maintenance

Spider Impact now automatically maintains database performance through scheduled index reorganization and metadata updates. This maintenance runs weekly to ensure optimal query performance.

Database indexes are reorganized on a scheduled basis to reduce fragmentation. After reorganizing indexes, the database automatically updates its internal metadata about data distribution, which helps the query optimizer choose the most efficient execution plans.

Different tables run on different schedules to distribute the maintenance load. Self-hosted customers can disable automatic maintenance for non-production servers by adding `database.reindex.disabled=true` to their `config.properties` file.

Pre-Calculated Dataset Summaries

Spider Impact now maintains pre-calculated summary information about dataset fields, dramatically improving query performance for certain operations. These summaries track value frequencies and ranges for each field, enabling faster data analysis and supporting Impact Intelligence features.

When you request minimum values, maximum values, number of unique values, or lists of unique values from a single dataset without filters, Spider Impact retrieves the answer from these summary tables instead of scanning all records. This provides near-instant results even for datasets with millions of records.

The optimization works automatically in both the Datasets Explore tab and on dashboards. Summaries are updated automatically when datasets finish building, when records are added, modified, or deleted, and when dataset structures change.