SAS® Visual Analytics 6.3

User's Guide

Contents

PART 3   Exploring Data   139

PART 6 Viewing Reports 525

PART 7 Appendix 541

Using This Book

Audience

The features of SAS Visual Analytics are designed for the following users:

- Persons needing to explore data in support of ad hoc business questions.

- Persons responsible for designing and creating reports for their enterprise.

- Persons responsible for analyzing report data and making decisions based on that data.

Persons responsible for managing SAS servers and managing the SAS Visual Analytics environment should refer to *SAS Visual Analytics: Administration Guide*.

Prerequisites

Here are the prerequisites for using SAS Visual Analytics:

- A user ID and password for logging on to SAS Visual Analytics.

- A supported web browser installed on your desktop client.

- A supported version of the Adobe Flash player installed on your desktop client.

- Access to data sources that can be used to obtain data for exploration or reports.

If you have questions about whether you are ready to use SAS Visual Analytics, contact your system administrator.

Documentation Conventions

This book uses short forms of the following phrases where the meaning is clear from context:

Long Form	Short Form
SAS Visual Analytics Explorer	the explorer
SAS Visual Analytics Designer	the designer
SAS Visual Analytics Graph Builder	the graph builder
SAS Visual Analytics Viewer	the viewer
SAS Visual Data Builder	the data builder

What's New

What's New in SAS Visual Analytics 6.3

SAS Visual Analytics Home Page

New and enhanced features on the home page include:

- You can switch between the home page and any other SAS Visual Analytics interface using the application bar.

- For searching on the home page, there is a new **Back** button that enables you to easily return to the home page. You can search using created or last modified by values and modified by dates.

- You can customize the right pane.

- Guest access can be enabled for the home page.

SAS Visual Data Builder

Enhancements to Working with SAS LASR Analytic Server

- You can create a LASR star schema to join in-memory tables.

- You can append in-memory tables to another in-memory table.

- You can save an in-memory table to HDFS. This is helpful after appending data so that the table can be reloaded from HDFS instead of transferring the data again.

Enhancements to Importing Data

- You can import files up to 4 GB in size. The speed for importing large files is faster.

- For files that are larger than 4 GB, you can transfer the files to the SAS Application Server, and then use autoload to import the files. Like the import local data feature, autoload supports spreadsheets, text files (CSV), and SAS data sets.

- You can specify the file encoding, such as UTF-8, when importing text files.

- Filenames (which are used as table names) and column names can use most special characters. An unsupported special character is replaced with the underscore character (_).

- When you import a spreadsheet or text file from your computer and load the data to SAS LASR Analytic Server, the data is also saved as a SAS data set on the SAS Application Server. If SAS LASR Analytic Server is configured to perform a reload-on-start, then these saved data sets are loaded to memory automatically.

- You can import tweets from Twitter and transfer them to SAS LASR Analytic Server.

- You can connect to an Oracle database and transfer a table to memory on SAS LASR Analytic Server.

General Enhancements to SAS Visual Data Builder

The application bar includes a single memory gauge to show the memory utilization for SAS LASR Analytic Server. Use the gauge to assess whether you can load more data.

The **File** menu has the following changes:

- The **Import Local Data** menu item is replaced with **File ▶ Import DataLocal File**.

- The **Import Remote Data** menu item is replaced with **File ▶ Import DataSAS Data Set on a Server**.

The scheduling windows in the data builder are revised to show time on a 24-hour clock rather than on a 12-hour clock.

SAS Visual Analytics Explorer

Enhancements for Visualizations

- The new network diagram visualization enables you to visualize the relationships in your data.

- The new word cloud visualization enables you to identify the most frequent values in your character data or to identify topics and the relevant terms for those topics by using text analytics.

- Visual filters are enhanced to support sorting, new filter conditions, and a greater number of category values.

- You can lock the data ranges and scaling for visualization axes that contain a measure. You can modify the axes and transfer the axes between visualizations.

■ Geo maps have a new **Coordinates** overlay type that displays data as a scatter plot. The **Coordinates** overlay enables you to display a greater number of data points on a map.

■ A new link on the visualization **Roles** tab enables you to toggle autocharting for the selected visualization.

■ Geo maps support Esri as a source for background maps.

■ For bubble plots and geo maps with a bubble overlay, a new scaling method for bubble sizes supports negative values and is directly proportional to the data values.

Enhancements for Data Management

■ You can import local data files (SAS data sets, Microsoft Excel files, or CSV files).

■ You can create calculated data items that return character data and dates.

■ New periodic operators for aggregated items enable you to perform aggregations on data over time.

■ New text operators for calculated items and filters enable you to perform string operations in your data expressions.

■ A new text mode for editing calculated items, aggregated items, and filters enables you to enter and edit your expression as text code.

■ You can create new derived items for year-to-date and other time-based aggregations.

■ You can derive the percentage of totals and percentage of subtotals for a crosstab visualization.

■ For geography data items, new predefined roles support ISO country codes, country subdivision names, and SAS map ID values for countries and country subdivisions.

■ You can export data from a table visualization to a CSV file.

SAS Visual Analytics Designer

Enhanced Performance for the Designer

Performance is improved when opening reports in the designer.

Enhanced User Interface for the Designer

- You can switch between the designer and any other SAS Visual Analytics interface using the application bar.

- You can customize the list of report objects on the **Objects** tab.

- You can access the SAS Visual Analytics Graph Builder (the graph builder) from the **Objects** tab. Using the graph builder, you can create custom graph objects.

- You can duplicate a section in a report.

- You can specify a preference for the **Data** tab that disables cardinality checking to improve performance.

- You can print reports that have not been saved. The options that you select for printing a report are maintained for that report.

- You can import local data into the designer.

Enhanced Report Objects

- You can add expression-based display rules for graphs. For example, you can change the color of a bar or marker based on the value of an expression.

- You can add measures to the tooltips for graphs. You can also change the style of the tooltips for graphs.

- You can color the background of graphs.

- You can add display rules for a crosstab.

- Bubble plots, geo bubble maps, and bar charts (with no groups, rows, or columns) with a date data item can be animated.

- A new stack container is available. You can group other report objects in a stack container.

- Geo bubble maps and geo region maps can now also use Esri as the map provider. The geo maps support state, territory, and province latitudes, longitudes, and regions.

- You can add alerts to report objects and receive notifications that alert conditions have been met via e-mail.

- A new property enables a report author to disable the ability to select report objects in the web viewer.

Enhanced Imported Report Objects

- Imported box plots and heat maps that contain a relational category data item and forecast plots and geo maps can be the source of a filter interaction.

- Read-only filter information is available for imported report objects on the **Filter** tab.

Enhanced Data Items

- A new data role is available for gauges, geo bubble maps, and geo region maps. This role accepts only measures that have not already been used in the report object.

- You can create a custom data item based on either a measure or a category data item.

- You can drag and drop multiple data items from the **Data** tab onto crosstabs, graphs, and gauges.

■ You can drag and drop multiple data items and have them automatically assigned so that you can quickly view graphs and gauges.

Enhanced Calculations

■ Advanced date operations are available.

■ One-click date operations are available.

■ String operations for calculations are available.

■ You can enter text for calculations.

SAS Visual Analytics Graph Builder

The SAS Visual Analytics Graph Builder (the graph builder) is new to SAS Visual Analytics. Its features include:

■ You can create custom graph objects to include in reports. These graph objects include an array of plot and chart types, layouts, and properties. They enable report designers to create simple or complex graphical views of their data.

■ You create custom graph objects by dragging and dropping graph elements onto the canvas. You can then combine and arrange the graph elements as needed. For example, you can layer multiple graph elements in one cell, and you can create a lattice of cells containing graph elements.

■ You determine which data roles become available for assignment. You can specify the type of role, whether it is required in a report, and other aspects of the role.

■ You have many options for specifying the visual attributes of the graph elements. You can specify properties for the graph object and its cells, for individual graph elements, for the axes, and for legends.

SAS Visual Analytics Viewer

New and enhanced features in SAS Visual Analytics Viewer include:

- Performance is improved when opening reports in the viewer.

- You can view display rules and incoming filters on the **Properties** tab.

- You can subscribe to alerts for report objects and receive e-mail notifications when the alert condition is met.

- You can set a preference to override the report theme when SAS High Contrast theme is selected.

- Guest access can be enabled for the viewer.

SAS Visual Analytics Administration

Note: For details about the following new and enhanced features in administration, see the *SAS Visual Analytics: Administration Guide*.

Guest Access

You can provide guest access to users who do not have a metadata identity. This enables you to make reports widely available under a generic, shared account. You can also use guest access to make reports available to the public.

Data Load and Import

- You can automate the following aspects of data loading:

□ You can enable a server to start on demand (autostart) for data load and import requests.

□ You can enable a server to immediately reload certain tables each time it starts (reload-on-start).

□ You can schedule periodic synchronization of in-memory data against a designated source directory.

■ Autoload supports refresh, append, and unload actions (as well as load actions).

■ Autoload supports Microsoft Excel spreadsheets and delimited text files (as well as SAS data sets).

■ You can support self-service imports in the designer and the explorer.

□ You can use a new capability, Import and Load Data, to control the availability of the import functionality in the explorer and designer.

□ You can use an existing capability, Build Data, to control access to advanced options in the explorer and designer. The advanced options enable users to change output locations.

Note: Output settings reference a general-purpose area (a public library, server, and metadata folder) that all registered users can access.

□ You can review and preserve the public library's support of autostart, autoload, and reload-on-start.

> **TIP** Import and load activities require host access to the SAS LASR Analytic Server.

■ You can use spaces and special characters in table and column names.

General Administration

■ You can more easily monitor a distributed SAS LASR Analytic Server.

□ The total memory gauge is in the main menu so that you can always see memory usage information.

□ The **LASR Servers** tab displays instance-specific memory gauges.

■ You can temporarily change application logging levels without editing configuration files and restarting servers.

■ You can easily change your deployment's approach to mobile device management.

■ You can disable e-mail notifications for failures of the search index loading script. These notifications are sent to the general administrative e-mail address.

■ When you create a new permission condition, you can enter text directly on a new **Text** tab. On the **Text** tab, only the syntax and operators that are available on the **Visual** tab are supported.

■ When you use the metadata import and export tools to promote reports, you can include linked reports as dependent objects. When you import a report, references in the source report are updated to reflect any changes that you make to the paths of included linked reports. Metadata promotion is documented in the *SAS Intelligence Platform: System Administration Guide*.

Note: Updates to references to linked reports occur for only changes that you make during the import process. These updates do not occur for changes that are made in other contexts. For example, these updates do not occur for a linked report that is renamed or moved on the **Folders** tab in SAS Management Console.

■ You can delete other users' alerts.

The Administrator Interface

■ In the main menu, actions that were previously provided as buttons are now available as menu items.

□ A new **LASR** menu provides access to server operation, server monitoring, and data loading features.

☐ The **Tools** menu provides access to mobile device management, alerts management, and (where applicable) HDFS browsing.

■ The window for selecting a LASR library displays only LASR libraries.

■ The **LASR Tables** tab displays tables in a flat, filterable list.

■ A new tab, the **LASR Servers** tab, displays servers in a flat, filterable list.

■ Right-click menus are available on both **LASR** tabs.

■ Some of the icons for data loading have changed.

Software Configuration Properties

■ In SAS Visual Analytics, new configuration properties include App.AllowGuest, va.monitoringPath, va.publicLASRLibrary, va.publicLASRServer, va.SASGeomapEsriURL, and va.SelfServe.MaxUploadedSizeInMegabytes.

■ In middle-tier high-cardinality constraints, the property va.maxPeriodCalculations is new.

Identities, Capabilities, and Permissions

■ The Visual Analytics Data Administrators group and the Visual Data Builder Administrators group are initially assigned to both of the SAS Management Console roles.

■ A new group, Visual Analytics Users, provides access to the designer and explorer. In the standard configuration, the SASUSERS group is not a member of the **Visual Analytics: Report Viewing** and **Visual Analytics: Analysis** roles.

■ A new role, **Visual Analytics: Basic**, makes a minimal set of features available for guest access (if applicable) and entry-level users.

■ The Build Data capability (which provides access to the data builder) provides access to advanced data load options in the explorer and designer.

■ The Customize Hub capability is renamed Administer Hub.

- A new capability, Build Custom Graph, controls whether the graph builder is available.

- A new capability, Import and Load Data, controls whether data import and load functionality are available in the explorer and designer.

- A new capability, Personalization, controls the availability of individualized features (such as recent items, favorites, collections, alerts, and user preferences).

- Requirements for the Delete permission are replaced by requirements for the Write permission.

- The Search Interface to SAS Content User is unrestricted in metadata and in ROLE_ADMIN in the middle tier. These privileges are required for successful indexing.

Accessibility

Accessibility Notice

For information about the accessibility of this product, see **Accessibility Features of SAS Visual Analytics 6.3 at support.sas.com**.

Part 1

Introduction to SAS Visual Analytics

2

1

About SAS Visual Analytics

What Is SAS Visual Analytics?

SAS Visual Analytics is an easy-to-use, web-based product that leverages SAS high-performance analytic technologies. SAS Visual Analytics empowers organizations to explore huge volumes of data very quickly to identify patterns and trends and to identify opportunities for further analysis. SAS Visual Data Builder (the data builder) enables users to summarize data, join data, and enhance the predictive power of their data. Users can prepare data for exploration and mining quickly and easily. The highly visual, drag-and-drop data interface of SAS Visual Analytics Explorer (the explorer), combined with the speed of the SAS LASR Analytic Server, accelerate analytic computations and

enable organizations to derive value from massive amounts of data. This creates an unprecedented ability to solve difficult problems, improve business performance, and mitigate risk rapidly and confidently. SAS Visual Analytics Designer (the designer) enables users to quickly create reports or dashboards, which can be viewed on a mobile device or on the web.

SAS Visual Analytics empowers business users, business analysts, and IT administrators to accomplish tasks from an integrated suite of applications that are accessed from a home page. The central entry point for SAS Visual Analytics enables users to perform a wide variety of tasks such as preparing data sources, exploring data, designing reports, as well as analyzing and interpreting data. Most important, reports can be displayed on a mobile device or in the SAS Visual Analytics Viewer (the viewer).

Benefits of Using SAS Visual Analytics

Using SAS Visual Analytics, users can enhance the analytic power of their data, explore new data sources, investigate them, and create visualizations to uncover relevant patterns. Users can then easily share those visualizations in reports. Traditional reporting is prescriptive. That is, you know what you are looking at and what you need to convey. However, data discovery invites you to plumb the data, its characteristics, and its relationships. Then, when useful visualizations are created, you can incorporate those visualizations into reports that are available on a mobile device or in the viewer.

SAS Visual Analytics provides users with the following benefits:

- enables users to apply the power of SAS analytics to massive amounts of data

- empowers users to visually explore data, based on any variety of measures, at amazingly fast speeds

- enables users to share insights with anyone, anywhere, via the web or a mobile device

How Does SAS Visual Analytics Work?

You can use SAS Visual Analytics to explore and view data, interact with and create reports, and display reports on a mobile device or on the web. You can explore your data by using interactive visualizations such as charts, histograms, and tables. Report designers can easily point and click to query central sources of data. You can add filters and design the layout using tables, graphs, and gauges. You can use drag and drop to create a well-formatted report.

The following figure illustrates how the different pieces of SAS Visual Analytics work together. It shows how users interact with the different interfaces.

Figure 1.1   *Overview of SAS Visual Analytics*
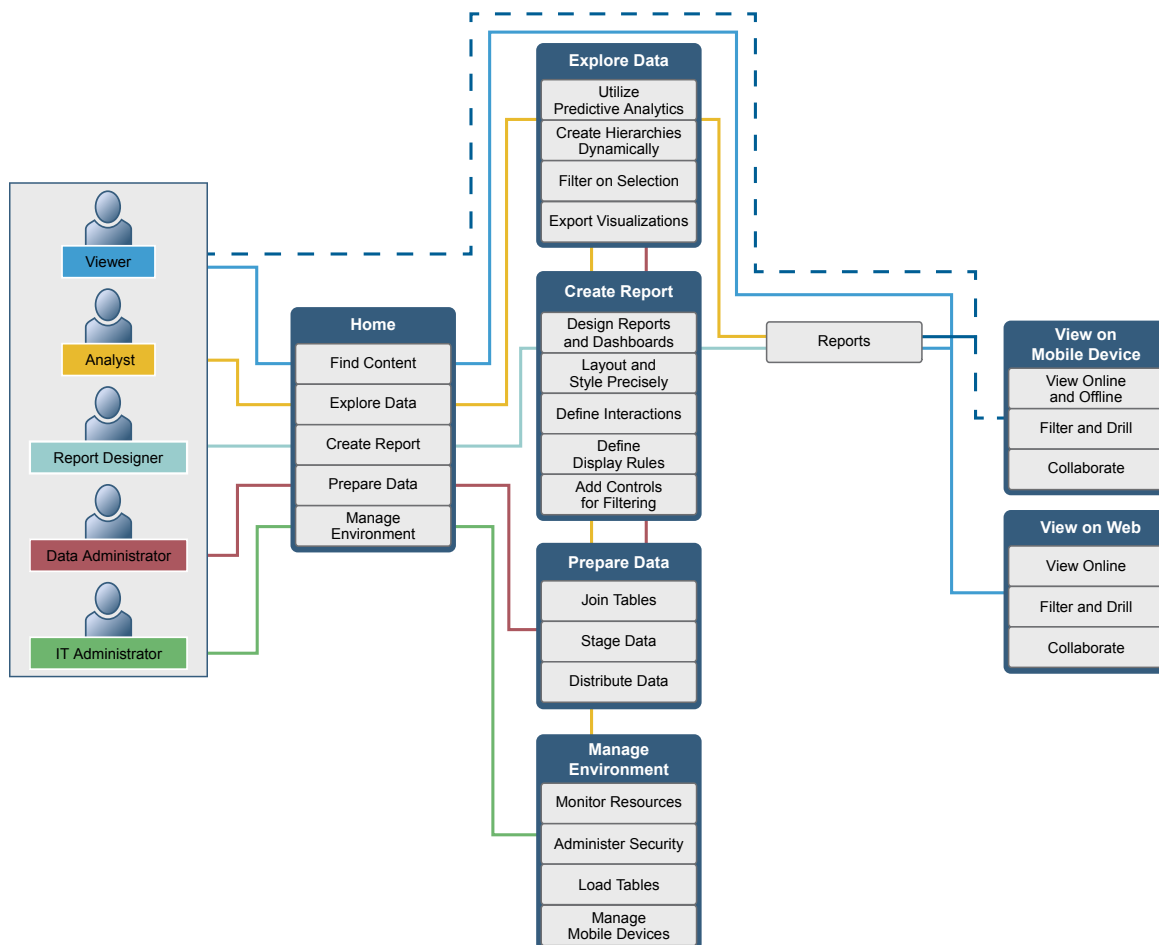
SAS Visual Analytics User Community

The SAS Visual Analytics user community is dedicated to users who are focused on exploratory visualization and analytical techniques, data preparation, dashboard reporting, and mobile BI. You can share your experiences, discuss topics and ideas, seek help from your peers, and share information about upcoming events. For questions requiring immediate technical assistance, contact SAS Technical Support at `support.sas.com`.

Overview of SAS Visual Analytics Capabilities

Users might have access to different functionality, depending on their assigned roles. Roles are mapped to capabilities. A *capability*, also known as an *application action*, defines the operations a user can perform.

SAS Visual Analytics provides five predefined roles—Basic, Report Viewing, Analysis, Data Building, and Administration. A predefined set of capabilities is available for each role. A system administrator can modify these roles and specify the capabilities for each role that meet the guidelines for your company. They can also define new roles. If you have questions about your assigned role, contact your system administrator. For more information about the roles and the capabilities that are available, see the *SAS Visual Analytics: Administration Guide*.

Note: This user's guide discusses tasks that you might be able to perform, depending on your role.

Accessing SAS Visual Analytics

Authenticated Users

SAS Visual Analytics uses the standard Log On window for SAS applications. To display the Log On window, click on the URL that is supplied by your system administrator or paste it into the address field of your browser. For example, you might enter: `http://host/SASVisualAnalyticsHub`

To log off from SAS Visual Analytics, click **Log Off** in the upper right corner of the user interface. When you click **Log Off**, you are logged off from all environments.

Guest Access

SAS Visual Analytics system administrators can configure support for guest access. Users with guest access can access only the SAS Visual Analytics home page and the viewer. Guest access uses a shared account, so it does not provide individualized features, such as history, favorites, preferences, or alerts. If guest access is configured, it is available at a special URL such as

`http://host/SASVisualAnalyticsHub/guest.jsp`

Accessing SAS Visual Analytics as a guest is useful if you do not have a metadata identity. This enables you to view reports that are widely available under a generic, shared account. You can also view reports that are available to the public on the Internet.

About the Availability of Menus and Menu Selections in SAS Visual Analytics

All of the following conditions influence whether a SAS Visual Analytics menu or menu selection is available to use:

- your role and the associated capabilities. For example, you must have a Data Building role to prepare data.

- your location in SAS Visual Analytics. For example, some application features are available only if you are designing a report.

- the currently selected report object. For example, ranges are not available for list tables.

- whether the data for a report has been defined. For example, if the data has not been selected, then you cannot create a filter.

For more information about roles and capabilities, see the *SAS Visual Analytics: Administration Guide*.

2

Overview of the SAS Visual Analytics Home Page

Your First Look at the SAS Visual Analytics Home Page

After you log on to SAS Visual Analytics using the standard Log On window for SAS applications, you will see the home page. The home page enables you to create new content in SAS Visual Analytics. In addition, it enables you to access content that you and others have created.

Here are the features of the home page:

Figure 2.1   The Home Page
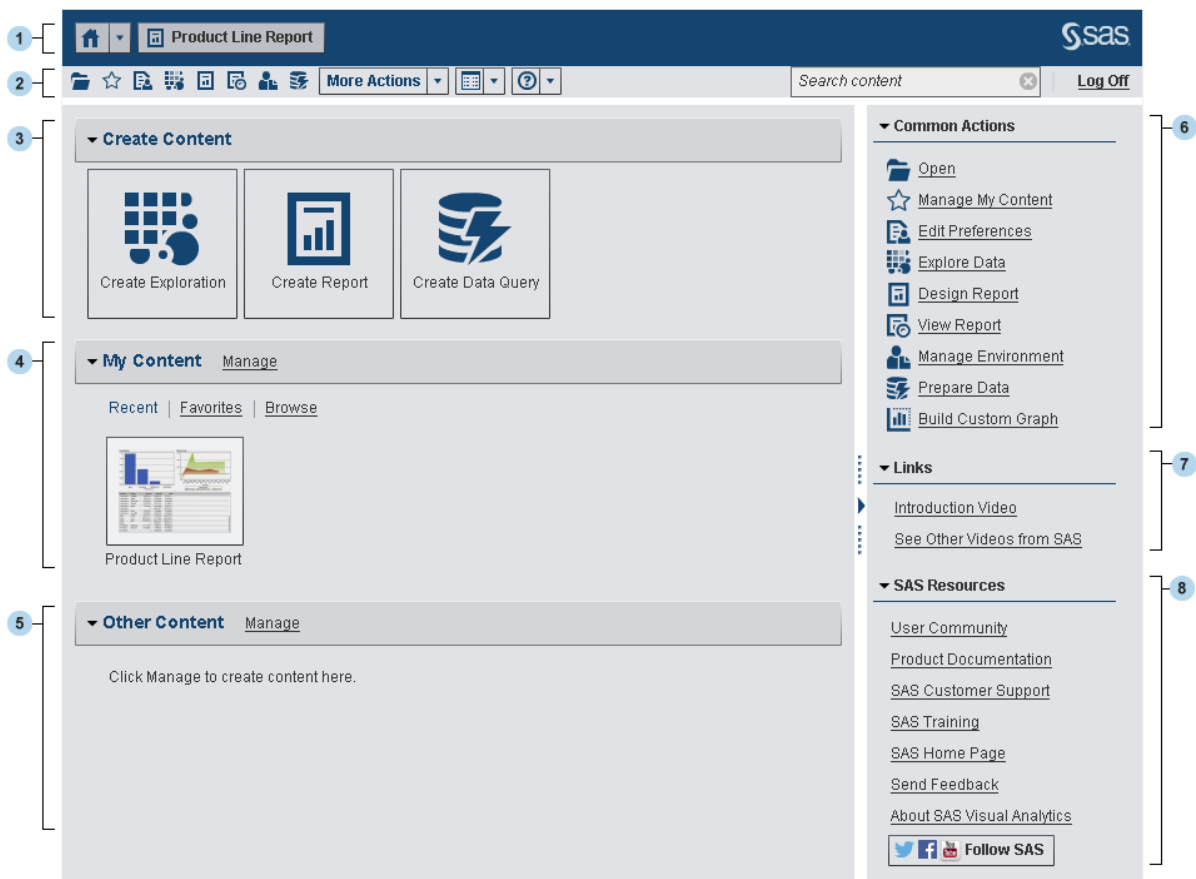
1   The application bar enables you to return to the home page and to access other parts of SAS Visual Analytics and other SAS applications that integrate with the home page. You can access your recently created or viewed reports, explorations, stored processes, data queries, or other objects in your recent history. Buttons are displayed for each open application.

2   The menu bar enables you to access task options, such as creating reports, exploring data, managing your environment or favorites, setting your view, and getting help on using SAS Visual Analytics. You can search all SAS content from the menu bar, and you can log off from SAS Visual Analytics.

3   The **Create Content** area provides icons to let you quickly explore data, create a new report, or prepare data, depending on your role and the associated capabilities. Other installed SAS applications might add actions to the **Create Content** area.

4   The **My Content** area lists any explorations, reports, queries, tables, or stored processes that you have opened or created recently. It also lists any content that you have marked as a favorite or as part of a collection. Click **Browse** to explore folders to find a report, exploration, stored process, or table.

   Note: All tables are displayed because the home page does not distinguish between LASR tables and other tables.

5   The **Other Content** area displays lists of explorations, reports, stored processes, queries, or tables that your system administrator has added for you. If your system administrator has not created any lists, then you will not see the **Other Content** area. If you have an Administration role, then you will also see a **Manage** link that enables you to manage lists.

6   The **Common Actions** section provides an alternate way for you to access features and other installed SAS applications. For more information, see “Working with the Right Pane on the Home Page” on page 15.

7   The **Links** section provides a link to the introductory video for SAS Visual Analytics and links to other pages that you have bookmarked. For more information, see “Working with the Right Pane on the Home Page” on page 15.

8 The **SAS Resources** section provides links to resources on the SAS website and to social media. For more information, see "Working with the Right Pane on the Home Page" on page 15.

Note: Guest access does not provide individualized features on the home page, such as history, favorites, preferences, or alerts. By default, a user with guest access has a Basic role and can access only the home page and the viewer. For more information, see "Guest Access" on page 8.

If you have the **Theme Designer for Flex: Administration** role in SAS Management Console, you can access the SAS Theme Designer for Flex from the **More Actions** item on the menu bar or in the **Common Actions** section.

Here is an example of the menu bar with the **More Actions** item.

Display 2.1   Home Page Menu Bar with the More Actions Item

Once you start using the home page, thumbnails enable you to open explorations, reports, stored processes, tables, queries, and folders that you have created or opened. The default view is secure, generic thumbnails that represent the content. Your system administrator can set a property that specifies the use of thumbnails that are shared and unique to each object. Shared thumbnails are unique for each individual report object, so each report looks different from other reports, and each exploration looks different from other explorations.

Generic thumbnails are distinguished by content type only. All reports look the same, but they appear different from explorations. Here are examples of the generic thumbnails that you might see on the home page:

Display 2.2   Generic Thumbnails for an Exploration, a Report, and a Stored Process

You can choose to view the content on the home page as a list. To change your view, click 🔲 on the menu bar, and then select **Thumbnail** or **List**. The default is **Thumbnail**.

Here is an example of the list view for the home page:

Display 2.3 *List View*
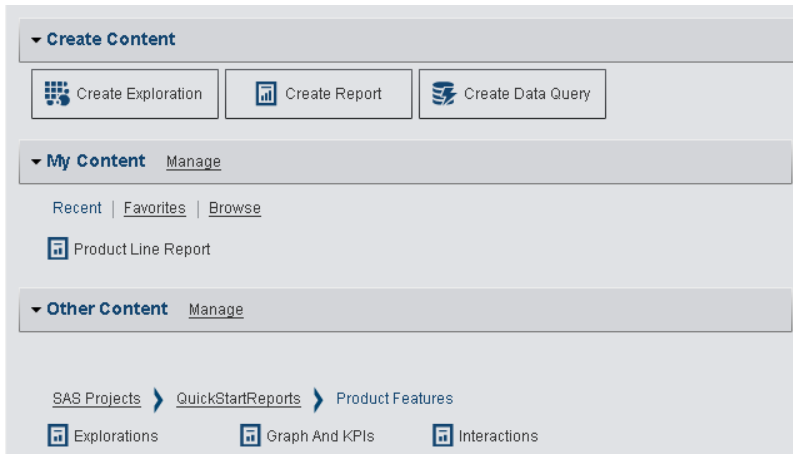
For more information about other parts of the SAS Visual Analytics interface, see the following topics:

- "Your First Look at the Data Builder" on page 41

- "Your First Look at the Explorer" on page 143

- "Your First Look at the Designer" on page 292

- "Your First Look at the Graph Builder" on page 472

Managing Content on the Home Page

The home page displays recent reports, explorations, stored processes, tables, and queries, as well as favorites and collections under the **My Content** heading on the home page.

Note: A user with guest access cannot manage content.

To manage your content, click **Manage** to the right of the **My Content** heading. The Manage My Content window is displayed.

Display 2.4   Manage My Content Window
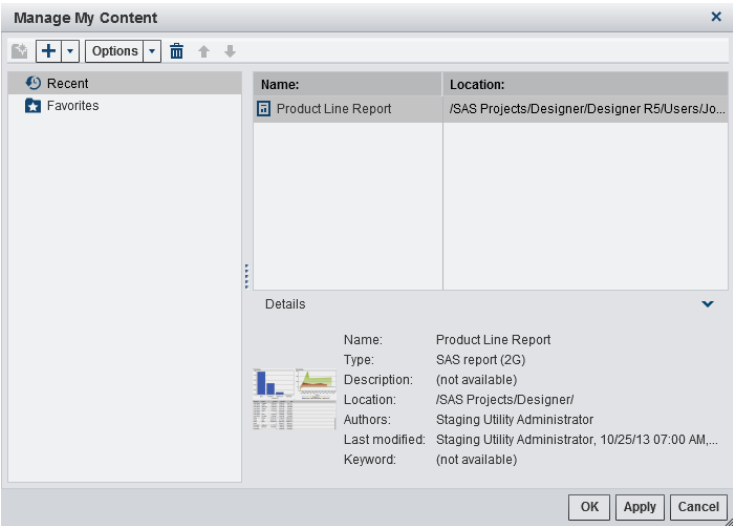
The following icons are available in the Manage My Content window:

Icon	Description
🗂️	Creates a new folder for your reports, explorations, stored processes, tables, and queries.
➕	Adds a report, exploration, stored process, query, or table to your list of favorites. You can also create or add a collection of favorites.
Options	Enables you to select either **Clear Recent History** or **Clear Favorites**.
🗑️	Enables you to delete favorite or recent content, one at a time. You can delete multiple items by pressing the Ctrl key. For a collection, you will have a choice of permanently deleting the collection or simply removing it under the **My Content** heading on the home page.
⬆️	Enables you to move favorites content up in the list to change what is displayed under the **My Content** heading on the home page.

Icon	Description
⬇	Enables you to move favorite content down in the list to change what is displayed under the **My Content** heading on the home page.

When you select a report, exploration, stored process, or table in the Manage My Content window, the details such as the name and location of that item are displayed. Favorite groups are displayed only in the tree on the left side of the Manage My Content window.

Working with the Right Pane on the Home Page

About the Right Pane

The **Common Actions**, **Links**, and **SAS Resources** sections are displayed in the right pane on the home page. Your system administrator can control which sections are displayed in the right pane. For example, your system administrator might hide all of the sections, two of the sections, or only one section. If any of the sections are displayed, you can control which sections are visible in your SAS Visual Analytics session.

The following sections can be displayed in the right pane:

Common Actions
  Provides an alternate way for you to access features, such as designing reports, exploring data, managing your environment or content. For example, you can click **Build Custom Graph** to access the graph builder. If you have an Administration or Data Building role, then you will also see a **Prepare Data** link. You might also see links to other SAS products, such as SAS Theme Designer for Flex, depending on your role and the SAS licenses that your site has.

Links
  Provides links to the introductory video and to other pages that you have bookmarked. Your system administrator can also provide shared links for all users. For information about links, see .

SAS Resources

Provides links to resources on the SAS website, including the SAS Visual Analytics User Community, and to social media.

Hide Content in the Right Pane

To hide content in the right pane, position your mouse pointer over the section title (**Common Actions**, **Links**, or **SAS Resources**), and click ✖.

To hide content in the right pane using the menu bar:

1 Click ▼ beside 🖽 on the menu bar.

2 Clear the selection for one or more of the following sections:

  ▪ **Common Actions**

  ▪ **Links**

  ▪ **SAS Resources**

  The selected section is hidden in the right pane.

Show Content in the Right Pane

To show content in the right pane:

1 Click ▼ beside 🖽 on the menu bar.

2 Select one or more of the following sections:

  ▪ **Common Actions**

  ▪ **Links**

  ▪ **SAS Resources**

  The selected section is displayed in the right pane.

If you have the Administer Hub capability, then you will also see a **Manage Shared View of Right Pane** menu item when you click ▼ beside ▦ on the menu bar. When you select this menu item, the Manage Shared View of the Right Pane window is displayed. Administrators can use this window to hide sections, restore sections, and add links in the right pane.

Manage Links in the Right Pane

Note: All URLs must start with `http://` or `https://`.

You can add, edit, and delete links in the **Links** section that is displayed in the right pane.

To add new links, position your mouse pointer over the **Links** section heading in the right pane to activate the ✚, which opens the Add Link window. Click ⬈ to open the Manage Links window, where you can add, delete, re-order, and hide links.

If you have an Administration role, then you can click ▼ beside ▦ on the menu bar, and select **Manage Shared View of Right Pane**. The Manage Shared View of the Right Pane window is displayed. You can use this window to add links in the right pane for other users.

Discovering Details Using the Object Inspector on the Home Page

The home page provides an object inspector that enables you to quickly find more information about an object when you click it. You can see a description of the report, exploration, stored process, table, query, folder, or collection; the location of the selected object; the creator; when the object was created and modified; and a list of keywords.

Your role and capabilities determine the availability of actions on the toolbar (such as viewing, editing, or opening) in the object inspector. For more information about roles and capabilities, see the *SAS Visual Analytics: Administration Guide*.

Here is an example of what you might see in the object inspector for a report that has been viewed:

Display 2.5   The Object Inspector on the Home Page
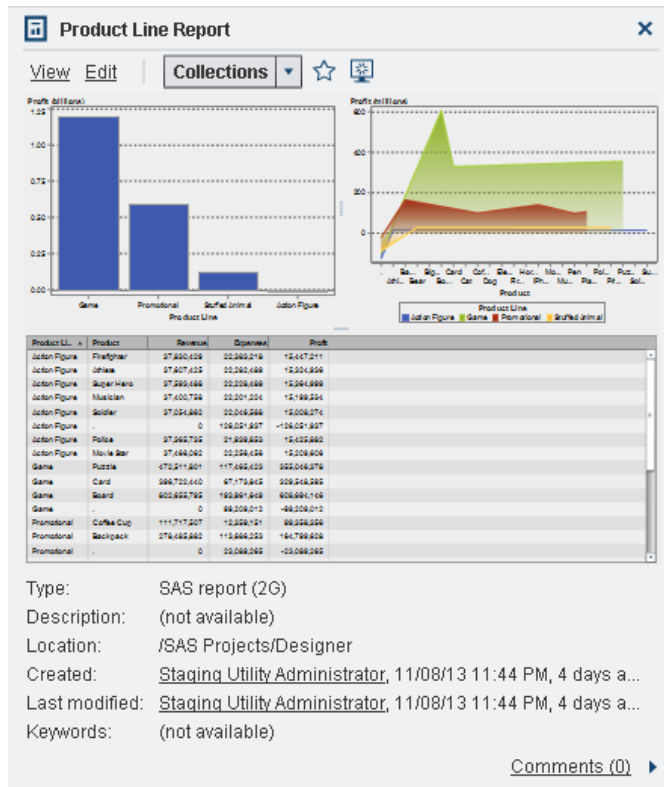
If you have the **Visual Analytics: Analysis**, **Visual Analytics: Data Building**, or **Visual Analytics: Administration** role and the Create Report capability, Explore Data capability, or both, then the toolbar in the object inspector enables you to do the following:

- **View** a report, which opens the report in the viewer.

- **Edit** a report, which opens the report in the designer so that you can edit or change the objects in the report.

- **Open** an exploration, which displays the explorer. The **Open** link is also available for stored processes, queries, and folders.

- **Create Report** for a table, which displays the designer.

- **Explore Data** for a table, which displays the explorer.

- Click **Collections** to select an existing collection or to create a new collection. For more information, see "Managing Collections on the Home Page" on page 19.

- Click ☆ to add the object to your list of favorites. When the star icon is yellow, it indicates that the object is in your list of favorites. Click ⭐ to remove the object from your list of favorites.

- Click 🖥 to set the object as the initial screen when you start SAS Visual Analytics. When the icon is blue, it indicates that the object is set as your initial screen. Click 🖥 to remove this initial screen setting.

In the object inspector, you can also do the following:

- Click the user name beside the **Created** or **Last modified** labels to search for all of the SAS reports, queries, tables, and stored processes that have been created or modified by that user.

- Click **Comments** to add or view comments. Any user who has the Add or View Comments capability can add comments to any report, exploration, stored process, table, or query. They view existing comments about any report, exploration, stored process, table, or query using the **Comments** link in the lower right of the object inspector. You cannot comment on folders, favorite groups, or collections. For more information, see "Adding Comments to Objects on the Home Page" on page 21.

Managing Collections on the Home Page
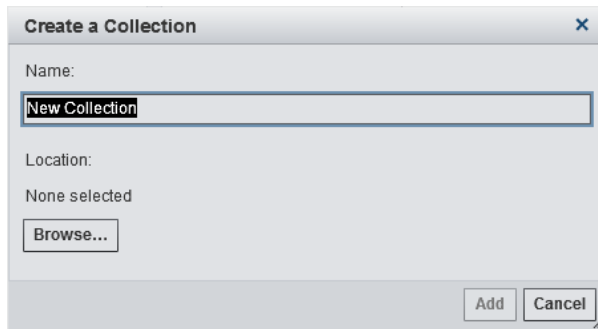
A collection is like a group of bookmarks. By default, creating a collection is restricted to administrative and advanced users. Your system administrator can grant you the capability to work with collections. You can create a collection that points to your favorite reports, explorations, stored processes, tables, queries, folders, and favorite groups.

To create a new collection:

1 Do one of the following:
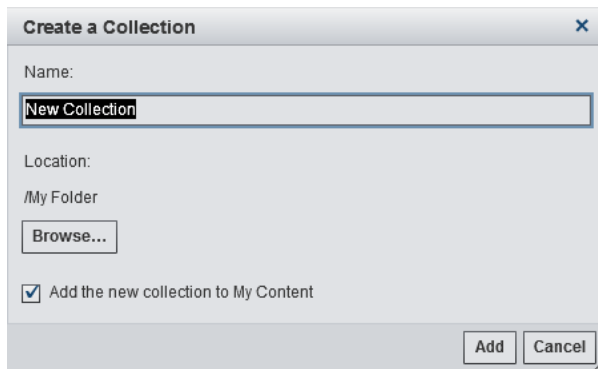
▪ On the home page, click **Manage** to the right of the **My Content** heading. The Manage My Content window is displayed.

Click ✚ to open the menu, and then select **Create a Collection**. The Create a Collection window is displayed.

▪ In the object inspector, click **Collections**. Then, select **Create a Collection**. The Create a Collection window is displayed.

2 Enter a **Name** for the collection.

3 Click **Browse** to select a **Location** for the collection. The Choose a Location window is displayed.

Select an existing folder or create a new one. You can also add an existing collection to a collection. Click **OK** to return to the Create a Collection window.

Note: When you create a new collection, you can add items to it from within the Manage My Content window only if the collection is a root collection in **My Collections**. For example, if you create *Collection B*, and add it to *Collection A*, then *Collection B* appears on the right side of the Manage My Content window and you cannot add items to it.

4   (Optional) If you opened the Create a Collection window from the object inspector, then you can clear the **Add the new collection to My Content** check box. This means the new collection will not be displayed on the home page next to the **Favorites** link.

5   Click **Add**. If you are in the Manage My Content window, then the new collection will appear on the left.

In the Manage My Content window, you can rename, remove, or permanently delete a collection by right-clicking on the name of the collection and selecting the applicable option.

Adding Comments to Objects on the Home Page

If you have the Add or View Comments capability, then you can use the object inspector on the home page to add (or view) comments to objects on the home page. You can add comments to reports, explorations, stored processes, tables and queries. You can also respond to existing comments or edit your own comments. You cannot comment on folders, favorite groups, or collections. You cannot see visualization-level comments on the home page.

To add a comment:

1   Click on a report, exploration, stored process, tables, or query to open the object inspector.

2   Click **Comments** in the lower right corner to expand the object inspector.

Display 2.6   Comments in the Object Inspector
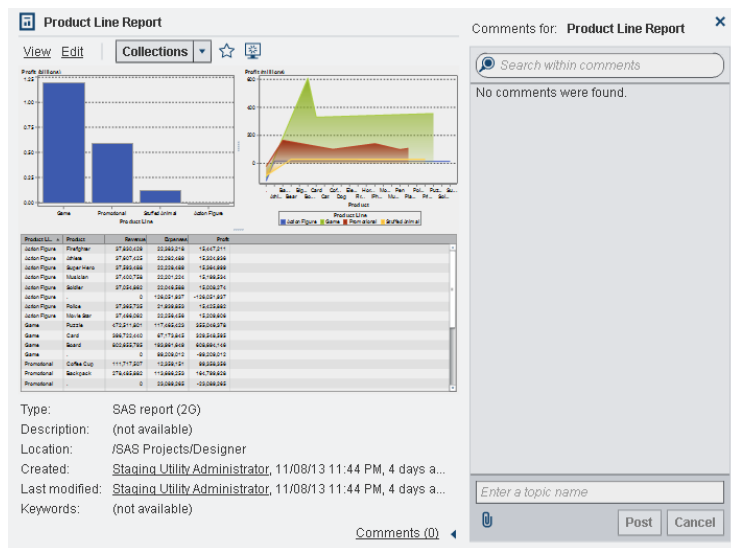
3 Enter a topic name and a comment.

4 Click **Post** to add your comment. The **Comments** link in the object inspector will update to show that there is a comment.

5 (Optional) Click 🔗 to attach a file or an image to your comment.

6 (Optional) Click the user name beside the **Created** or **Last modified** labels to search for all of the SAS reports, queries, tables, and stored processes that have been created or modified by that user.

To respond to an existing comment:

1 Click on a report, exploration, stored process, tables, or query to open the object inspector.

2 Click **Comments** in the lower right corner to expand the object inspector.

3 Select an existing comment. Then, enter a reply.

4 (Optional) Click 🔗 to attach a file or image to your reply.

5   Click **Post** to add your comment. The **Comments** link in the object inspector will update to show that there is a comment.

Note: To edit another user's comments or to delete comments, you must belong to the predefined role **Comments:Administrator**. This role includes the capabilities of editing or deleting comments.

To search for a comment:

1   Enter the word or phrase that you want to search for in the search field. Press Enter.

2   (Optional) To clear your search, click . Then, you can enter another word or phrase in the search field.

Specifying Your Preferences

All of your preferences persist between sessions. Preferences are not available for a user with guest access.

Specify Global Preferences

You can specify global preferences that are applied to all SAS web applications that are displayed with the Adobe Flash player. These preferences are set by each user.

To specify global preferences:

1   Click either or **Edit Preferences** on the home page. The Preferences window is displayed.

   Note: If you are in the data builder, the explorer, the designer, or the viewer, then select **File ▸ Preferences** to open the Preferences window.

2   Click **Global Preferences** in the left pane.

3   Select a **User locale** to specify your language and geographic region.

Select a **Theme** to change the color scheme and other visual settings for all of your SAS web applications.

Select **Invert application colors** to invert all of the colors in your SAS web applications.

Select **Override settings for focus indicator** to change the color, thickness, and opacity of the focus in your SAS web applications.

Note: If you change the **User locale**, then you must log off and log back on to SAS Visual Analytics for the change to take effect.

4 Click **OK** to apply your changes.

Specify the SAS Visual Analytics General Preferences

Using the SAS Visual Analytics home page, you can specify your general preferences for SAS Visual Analytics.

To specify your general preferences:

1 Click either or **Edit Preferences** to open the Preferences window.

2 Click **General** in the left pane.

3 Specify a value for **Show this number of recent items**. There is a minimum of one item and a maximum of 16 items. The default setting is 9 items.

 Click **Clear History** to reset your history.

4 Click **OK** to apply your changes.

Specify Your Initial Screen Preference

You can specify the initial screen that you want to see after you log on to SAS Visual Analytics.

To specify the initial screen preference using the home page:

1 Click either ![icon] or **Edit Preferences** to open the Preferences window.

2 Click **Home** in the left pane.

3 Select an option from the **Initial screen** drop-down list. The available options depend on your role and capabilities.

> **TIP** If you select **A specific object**, then you can use the Choose an Item window to choose one content type (for example, reports, explorations, tables, stored processes, and so on) for your initial screen. The available content types depend on the SAS products that your site has licensed and how they are configured.

4 Click **OK** to apply your changes.

Specify Other Preferences

You can specify other preferences in SAS Visual Analytics:

- ■ "Specifying Your Preferences for the Data Builder" on page 45
- ■ "Specifying Your Preferences for the Explorer" on page 145
- ■ "Specifying Your Preferences for the Designer" on page 297
- ■ "Specifying Your Preferences for the Viewer" on page 538

3

Searching from the SAS Visual Analytics Home Page

About Searching from the Home Page

You can search all reports, explorations, stored processes, tables, and queries that are on the metadata server from the menu bar on the home page. The list of items that you can search for depends on the SAS products that your site has licensed.

The following fields in metadata are searched:

- Comment Title

- Comment Description

- Created

- Description

- Keywords

- Last Modified

- Title

The search field on the menu bar supports the following types of searches:

▪ single word

▪ multiple word

▪ in-word wildcards

▪ spaces in strings that use quotation marks

▪ plus (+) and minus (-) syntax

  The plus (+) sign and minus (-) sign are prefix operators. This means that the operator precedes the search term that is required or excluded, rather than following it. For example:

  ☐ Specifying `2012 2013 2014` matches documents that contain any of the three terms.

  ☐ Specifying `+2012 +2013 +2014` matches only documents that contain all three terms.

  ☐ Specifying `2012 2013 -2014` matches documents that contain `2012` or `2013`, but only if they do not contain `2014`.

Note: Single- and multiple-word searches that are plain text have a wildcard appended to them. For example, if you search for the word *sample*, then it is converted to `sample*`. This means that you get different results than if you entered *"sample"*, *+sample*, or **sample**.

The search is not case sensitive.

Up to 5000 items can be returned in a single search. If your search results exceed 5000 items, then the first 5000 items appear, and a message is displayed at the top of the Search Results window. The message tells you how many results matched your search and how many have been displayed. Narrowing your search using the search field enables the relevant search results to be displayed. When you submit a new search that returns fewer than 5000 items, you can see all the search results.

Search results on the home page are not ranked. In the workspace, you can use the **Sort by** drop-down list to specify that you want the list of items to sort by **Alphabetical**

ascending, **Alphabetical descending**, **Date ascending**, or **Date descending**. The default is **Alphabetical ascending**.

Suppose that you search for the word *sample*. All items that contain or start with the word *sample* are displayed in the workspace. (For example, an item that starts with *samples* is displayed.) Here is an example of the search results:

Display 3.1   *Search Results on the Home Page*
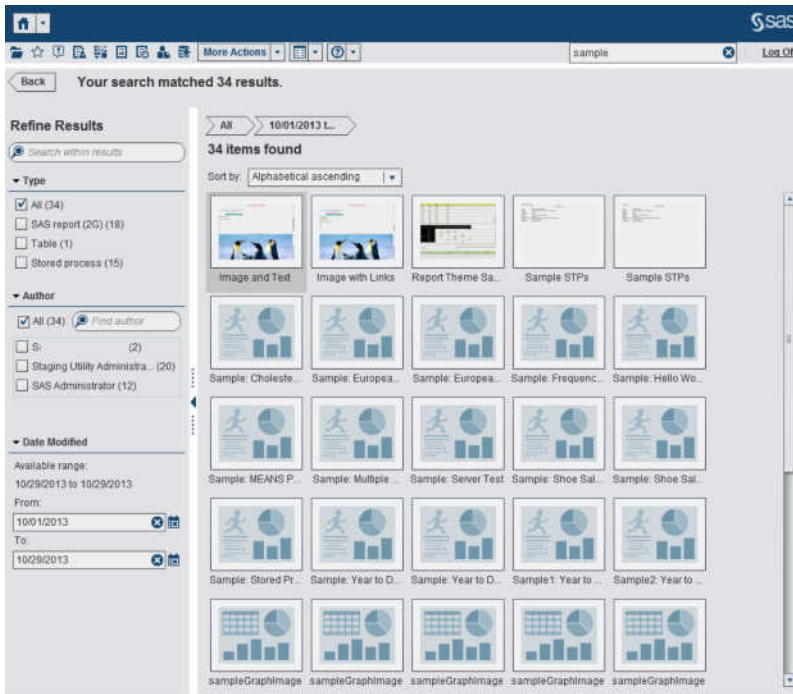
Refining Your Search Results

You can refine your search results using your original search on the metadata server. After you enter your original search using the menu bar on the home page, the search results are displayed in the workspace. Note that when you refine your search results, the initial 5000 item subset is not affected.

The refine search feature is dependent on your locale. The search is not case sensitive. No special characters are recognized except spaces and quotation marks.

You can refine the search results displayed in the workspace using the **Search within results** field in the left pane. When you enter a term or terms, you can use quotation marks and spaces when you are searching within your results. Only the items that match your refined search appear in your workspace. Each search term that you enter must appear at least once in the object name, the description, the author name, or keywords in the metadata.

You can also refine your search results by selecting (or clearing) check boxes in the left pane. Note that the totals in the parentheses in any panel are dependent on what has been selected in the other two panels.

The following panels are available:

- The **Type** panel shows the available object types. In SAS Visual Analytics, you can search for **SAS report (2G)**, **Table**, **Stored process**, **Visual data query**, or **Visual exploration**. Depending on the SAS products that you have licensed at your site, other SAS object types might appear in the **Type** panel. The default is **All** object types.

  The available types are determined by the capabilities that your system administrator has assigned. For example, you might see only the **All**, **SAS report (2G)**, **Table**, and **Stored process** check boxes in the **Type** panel.

  If you select all of the check boxes for the individual types, then the check boxes clear, and the **All** check box is selected automatically. If you clear the check boxes for all of the types, then the **All** check box is selected automatically.

  Note: If your role permits you to see only one type, then you will not see the **Type** panel.

- The **Author** panel enables you to search for a specific author, or to select one or more authors by name.

  The number in parentheses beside the author's name identifies the number of objects that the author has created using the filters that have been selected in the other panels.

If you select all of the check boxes for the authors in the list, then the check boxes clear, and the **All** check box is selected automatically. If you clear the check boxes for all of the authors, then the **All** check box is selected automatically.

■ The **Date Modified** panel lists the dates that the objects were last modified. You can refine the search to a specific date range by specifying a **From** date, a **To** date, or both. If the initial results set represents only a single day, then the **From** and **To** fields are hidden to conserve visual space.

The dates are time zone dependent. If your company has offices in multiple time zones, then the date modified time stamp on the object reflects the time at which it was saved using the current user's time zone. For example, if a report was modified by a user on January 3 at 2:00 a.m. in New York, another user in California sees that it was modified on January 2 at 11:00 p.m.

At the top of the workspace, you can see the breadcrumbs for the selections that you have made to refine your search. Breadcrumbs change as you modify the selections in the **Type**, **Author**, and **Date Modified** panels. The breadcrumbs also show any search strings that you entered in the **Search within results** field.

The appearance of the left panel is saved across sessions of SAS Visual Analytics. For example, if you collapse the **Type** panel, and then change the width of the **Refine Results** section, when you log off and log back on, the **Type** panel remains collapsed, and the **Refine Results** section has the new width. This is important to remember if you close the **Refine Results** section. Your **Refine Results** selections are not retained across sessions or primary searches in SAS Visual Analytics.

Suppose that you search for the word *sample*. Then, you refine the search results by selecting **Stored process** for **Type**. All stored processes that contain the word *sample* are displayed in the workspace. Here is an example of the refined search results:

Display 3.2   Refined Search Results on the Home Page
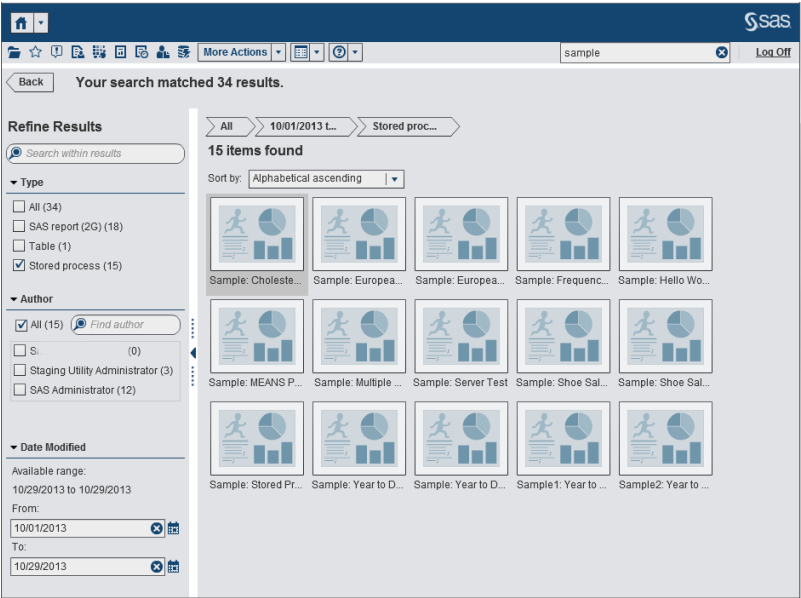
4

Accessing Data with SAS Visual Analytics

Self-Service Data Access

SAS Visual Analytics offers a variety of ways to bring your data into the SAS Visual Analytics environment, which enables you to work with your data quickly. The following list summarizes the different ways:

■ **Import data from a file** When you are creating data queries, explorations, or reports, you can import data from a Microsoft Excel spreadsheet, a delimited text file (CSV), or a SAS data set. After you import the data, it is ready to use.

This feature is available to users in SAS Visual Analytics Explorer (the explorer) and SAS Visual Analytics Designer (the designer) who have been granted the Import and Load Data capability. Users who have been granted the Build Data capability can access advanced functions.

For more information about importing data in the explorer, see "Import a Local Data Source" on page 156.

For more information about importing data in the designer, see "Import a Data Source for a Report" on page 346.

For more information about importing data in the data builder, see . All of the advanced functions are available in the data builder.

■ **Autoload** Instead of using your browser to import a file, you can transfer the file to a directory that acts as a drop box. A scheduled job imports the file from the directory. The process requires some setup before autoload can be used.

The data builder provides a few more ways to perform self-service data access.

■ **Import SAS data set on a server** You can direct the SAS Application Server to import a SAS data set that is already on the server. The data builder can optimize the import process and import large files when a SAS data set is already on the server.

■ **Import an Oracle table** After providing connection information, you can transfer a table from an Oracle database to a SAS LASR Analytic Server. Afterward, you can use the in-memory table as is for reports and explorations, or you can prepare it for analysis with the data builder.

   Note: Your site must license and configure SAS/ACCESS Interface to Oracle to use this feature.

■ **Connect to a Twitter stream** After providing connection information and search criteria, you can stream tweets into memory on a SAS LASR Analytic Server. You can then use the unstructured data with the explorer. A typical exploration of Twitter data can perform text analysis to look for patterns and trends in the tweets.

Managed Data Access

Many sites prepare data sources to be used by business analysts. This information can come from data sources such as data warehouses and transactional systems. A data administrator enables access by registering the tables and libraries in SAS metadata. A data administrator might use additional SAS products to provide advanced data governance, data quality, and data management support.

SAS/ACCESS engines can be used to access data in operational systems, transactional systems, or data warehouses. SAS offers a variety of SAS/ACCESS engines for accessing data from operational systems. These engines must be licensed and configured at your site to connect to the data. After a connection is set up, you can use SAS Management Console and SAS Visual Analytics Administrator for registering the libraries and tables from these data sources.

Part 2

Preparing Data

5

Overview of SAS Visual Data Builder

What Is SAS Visual Data Builder?

The data builder enables analysts and data administrators to perform basic data preparation. You can create data queries to perform joins, add calculated columns, and subset and sort data. Several productivity features accelerate the creation of columns based on common aggregation functions.

Once you create your data queries, you can reuse them as subqueries for more sophisticated data queries, export them as jobs for scheduling, or schedule them directly from the user interface.

The data builder has data import features that enable you to access data from the following data sources:

■   Microsoft Excel spreadsheets

■   delimited text files

■   SAS data sets

- Oracle database tables

- Twitter stream

After you import the data, you can prepare it for analysis or join it with existing data.

The data builder provides a series of features that take advantage of the in-memory tables in SAS LASR Analytic Server.

You can perform the following operations to add data to memory in the server:

- load an existing table directly into memory

- load the results of a data query into memory (or stage the data and then load it into memory)

- append rows to an in-memory table

After the data is in memory, you can perform the following operations with in-memory tables:

- join in-memory tables to form a LASR star schema

- append entire in-memory tables to another in-memory table

- save in-memory tables to HDFS for persistence and fast reloads

Your First Look at the Data Builder

Here are the features of the data builder:
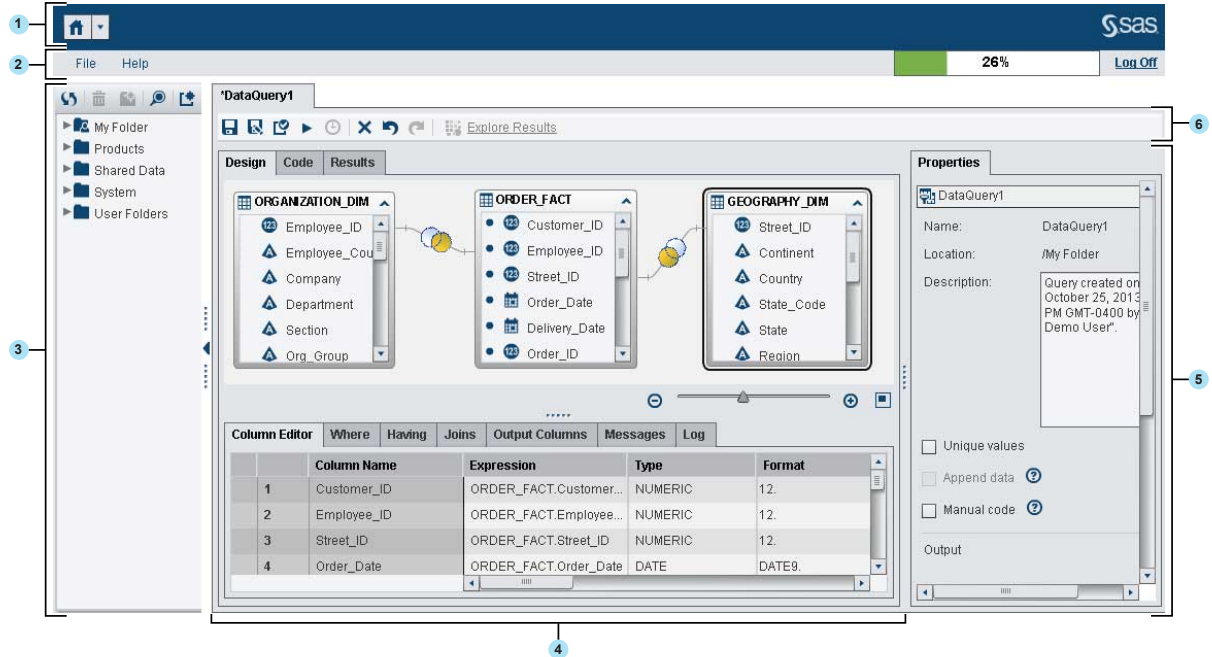
1   The application bar enables you to return to the home page and to access other parts of SAS Visual Analytics and other SAS applications that integrate with the home page. You can access your recently created or viewed reports, explorations, stored processes, data queries, or other objects in your recent history. Buttons are displayed for each open application.

2   The menu bar contains menus that enable you to perform tasks such as creating new data queries and LASR star schemas. The right side of the menu bar has a memory gauge that displays the memory utilization for SAS LASR Analytic Server. You can also log off from SAS Visual Analytics.

3   The navigation pane displays a tree (the SAS Folders tree) of tables and data queries.

4 The center of the screen contains the workspace. When you create a new object, such as a data query, it is represented as a tab on the top of the workspace.

The bottom of the workspace contains a series of tabs that enable you to create column expressions, joins, and filter data.

5 The properties panel enables you to manage the properties of the item that is selected in the workspace.

6 The toolbar contains icons that enable you to manage, run, and schedule data queries.

About Access to DBMS Data

The data builder can be used to read source tables from third-party vendor databases and to write tables to them as well. The SAS/ACCESS Interface products provide this feature. The user ID and password that you use to log on to the data builder might not be valid for a third-party vendor database. If this is the case, then you are prompted for credentials to the DBMS when you access a table from the database with a Read or a Write operation. As an alternative to being prompted, you can store a login in metadata that has valid DBMS credentials. For more information, see "How to Store Passwords for a Third-Party Server" in *SAS Intelligence Platform: Security Administration Guide*.

If you are prompted for credentials and supply an invalid user ID or password, you are denied access to the data. In this case, you can select **File ▶ Clear Credentials Cache** to remove the invalid credentials from your session. The next time you access the data source, you are prompted again.

Third-party databases are often more restrictive about column names or length than SAS. For more information, see "DBMS-Specific Naming Conventions" in *SAS/ACCESS for Relational Databases: Reference*.

Working with User-Defined Formats

A format is a set of instructions that SAS uses to write data values. Formats are used in the data builder to control the written appearance of data values. User-defined formats are specialized formats that are not supplied by SAS. These formats are stored in a custom format catalog.

The preferred method for making user-defined formats available to a SAS Application Server is to name the custom format catalog formats.sas7bcat, and to place it in `SAS-config-dir/Lev1/SASApp/SASEnvironment/SASFormats`. For more information about using user-defined formats, see *SAS Intelligence Platform: Data Administration Guide*.

When a user-defined format is permanently associated with a variable, the data builder uses the format. The **Format** column displays the user-defined format name. However, if you change the format, you cannot set the **Format** menu back to the original user-defined format. You can click ↩ to undo the change to the format if you have not saved your work yet, or you can remove and add the column back again.

Even if the custom format catalog is made available to the SAS Application Server correctly, you cannot associate a user-defined format to a variable with the data builder. The menu in the **Format** column does not enable you to specify a user-defined format.

6

Specifying Preferences for SAS Visual Data Builder

Specifying Global and General Preferences

To specify global SAS preferences, see "Specify Global Preferences" on page 23. To specify general preferences, see "Specify the SAS Visual Analytics General Preferences" on page 24.

Specifying Your Preferences for the Data Builder

To specify preferences that are specific to SAS Visual Data Builder, perform the following steps after you log on to the data builder:

1   Select **File** ▸ **Preferences** to open the Preferences dialog box.

2   Select **SAS Visual Data Builder** ▸ **General**.

Select a default SAS Application Server to use from **Default application server**. If you have added SAS Application Server instances to your deployment, make sure that the Job Execution Service has been configured for the SAS Application Server that you select. For more information, see *SAS Intelligence Platform: Middle-Tier Administration Guide*.

3   Select **SAS Visual Data Builder** ▶ **Scheduling**.

Specify a default scheduling server, batch server, and deployment directory. For more information, see "Scheduling Preferences" on page 127.

4   Click **OK** to apply your changes.

7

Creating Data Queries

What Is a Data Query?

A data query is your primary method for selecting and formatting data that is used with data exploration and creating reports.

A data query is a metadata object that manages the references to input tables, output tables, staging tables, joins, and summarizations from a SAS Visual Data Builder session. You can save data queries and open them later to edit the data preparation operations that are performed in the data query. You can use saved data queries as subqueries when you create a new data query.

About Creating Data Queries

You can use SAS Visual Data Builder to create data queries to prepare data for analytics. You can use a data query to subset, sort, join, and add calculated columns to tables.

As you create your data query, you can click to validate your data query. Use the button to preview or run the data query.

The preview option uses a temporary table for the output table. If you click the **Results** tab, it shows the first 100 rows only. When you run the data query, it uses the specified output table instead of a temporary table.

Saving Your Data Query

When you click , you clear the undo and redo history.

Note: If the data query uses a SAS LASR Analytic Server library for the output table (this is the default), then the server does not need to be running when you save the data query, but it must have been started at least once in the past. When you save it, the data builder checks the metadata authorization for the library and table before it performs the save. If the server has never been started, then there is no security key to use for checking authorization. For more information, see "Security Keys" in *SAS Visual Analytics: Administration Guide*.

Saving a Data Query as a New Data Query

Open an existing data query from the SAS Folders tree that closely matches the data query that you want to create. Click ![icon], and then enter a new name and select a location. Perform any customizations.

Note: If you manually change the name of the output table in the newly saved data query, then the name of the original output table is also changed in the original data query. Click **Clear**, and then specify a name and new information to avoid overwriting the original output table.

See Also

"Output Tables" on page 57

Using the Design Tab

The **Design** tab is the default view for working with data queries. This tab provides an easy-to-use interface for creating a data query. You can perform the following tasks with the **Design** tab:

- Drag and drop tables or data queries from the SAS Folders tree onto the workspace.

- Join tables by using your pointer to select the source column, and then drag the pointer to the corresponding column in the joined table.

- Add columns to the **Column Editor** tab by clicking the column name from the table in the workspace or by right-clicking on the table and selecting **Add All Columns**.

- Use the **Column Editor** tab to specify column expressions, aggregations, and sort. You can set the attributes for a column, such as the type, format, and label.

- View the **Output Columns** tab to see the number of output columns and attributes. When you specify aggregations and pivot by columns, the number of output columns can increase dramatically.

- Click 🗔 to check that the data query is valid.

- Check the **Messages** tab for information about warnings and errors such as invalid column expressions.

- Check the **Log** tab to view the SAS log. A SAS log is generated when you preview, run, or validate a data query.

- Use the **Properties** tab in the right pane to specify output table information.

See Also

- "Working with Tables in Data Queries" on page 55

- "Working with Columns in Data Queries" on page 63

- "Working with Filters in Data Queries" on page 73

- "Working with Joins in Data Queries" on page 77

Adding a Data Source

Adding a Table

To add a table to a data query, use the SAS Folders tree to locate the table (⊞) and then drag and drop the table onto the workspace.

> **TIP** You can also click 🔎 and search for the table by name and location.

Adding a Subquery

After you have created a data query and saved it, it can be used as an input data source to another data query. To add a subquery, use the SAS Folders tree to locate the data query (🖳), and then drag and drop the data query onto the workspace.

The subquery is represented in the workspace by the structure of the output table.

Best Practices for Adding Data Sources

When you plan to join data sources (tables or subqueries), the order in which you add the data sources to the workspace matters. The first data source that is added to the workspace is automatically assigned as the left table for any joins that you add to the data query.

If you are creating a data query that uses a fact table and dimension tables, drag and drop the fact table onto the workspace first. You can perform left, right, or full joins with the dimension tables faster because you need to specify the join type only. However, if you drag and drop a dimension table first, you can easily use the 🖳 button on the **Joins** tab to switch the left and right tables in the join.

If you are not using a fact table or dimension tables, then the sequence for adding tables to the workspace might not be very important. Just remember that the first table dropped onto the workspace is assigned as the left table, and you can switch the left table and right table on the **Joins** tab. The data builder takes advantage of vendor-specific features in SQL processing whenever possible. If the source tables are from a third-party vendor database, the SAS/ACCESS Interface engine can optimize the performance of the data query by passing the SQL statements through to the database.

Specifying Properties for a Data Query

On the **Properties** tab, you can view or specify the following properties:

Name

displays the name for the data query. The initial value is DataQuery1. You can specify a different name when you save the data query.

Location

displays the metadata folder location for the data query object. The initial value is `/My Folder`. You can specify a different location when you save the data query.

Description

specifies a description of the data query.

Unique values

specifies whether the SQL keyword DISTINCT is applied to the SELECT statement that is used to generate the result set for the data query.

Append data

specifies whether the result set for the data query is appended to the output table. If a staging table is used, then the staging table is replaced with the result set before appending to the output table.

Manual code

specifies that the data query uses the SQL that is written by the user on the **Code** tab when the check box is selected. For more information, see "Customizing Code" on page 121.

Table

specifies the table name for either the staging output or the final output. For more information, see "Specifying the Output Table" on page 57.

Location

specifies the metadata folder location to use for registering the staging table metadata or output table metadata. Click ■ to select a different location.

Library

specifies the library to use for the staging table or output table. Click 🖥 to select a different library.

Partition by

specifies the column to use for partitioning the output table. This property applies when the output table is in a SAS LASR Analytic Server library or SAS Data in HDFS library only. For more information, see "Partitioning Tables" on page 92.

See Also

"Specifying Properties for a Source Table" on page 56

8

Working with Tables in Data Queries

Source Tables

When you drop a table onto the workspace, the data builder connects to the SAS Metadata Server to determine the column names and data types for the table. When you drop a subquery onto the workspace, the data builder makes the same request, but determines the column names and data types for the output table of the subquery.

Because the data builder uses metadata for source tables, you (or an administrator) must register the table in metadata before you can access the data. Source tables can be registered with SAS Visual Analytics Administrator or SAS Management Console.

When you import data from a file, Oracle table, or Twitter stream, the data builder registers the data in metadata as a table. This enables you to work with it as a source table for data preparation.

See Also

- ■ "Adding a Data Source" on page 50
- ■ *SAS Visual Analytics: Administration Guide*
- ■ *SAS Intelligence Platform: Data Administration Guide*

Specifying Properties for a Source Table

When you create a data query and select a table on the workspace, or when you select it from the **Properties** menu, you can specify the following:

Alias
 displays the SAS table name that is stored in metadata. You can specify a new value to use as a table alias. The alias name is stored with the data query, and it does not affect the metadata information for the table.

Name
 displays the metadata object name. You can change the metadata object name in the SAS Folders tree by using **Rename**.

Location
 displays the table's metadata folder.

Library
 displays the table's library.

Auto-aggregate
 specifies whether to apply aggregations to the columns for this table when the columns are added to the query.

Functions
> specifies the aggregations to apply to the columns for this table when the columns are added to the data query.

See Also

Output Tables

About Output Tables

When you create a data query, you specify an output table on the **Properties** tab. When you save the data query, the output table is registered (or updated) in the metadata. When you run the data query, the physical output table is created and then the metadata is updated. The metadata that is created when you save or run the data query enables you to use the table as a source table for another data query or another SAS application.

When you create a data query, the default output table name is OutputTable.

Specifying the Output Table

Every data query must have an output table in order to run it. How you specify the output table affects whether metadata is updated or created.

To specify an output table:

1   Select the top-level data query object from the **Properties** tab.

2   Specify an output table name, location, and library.

The following table shows alternative ways to specify an output table.

Action	How To
Change the name of the output table. *	Enter the new name in the **Table** field, and save the data query. All data queries, reports, explorations, and so on, still reference the same table object.
Create a new output table. *	Click **Clear**, and then specify a table name, location, and library. This action disassociates the previously used output table.
Reuse an existing table.	Click ⊞, and then browse for the table to use. This action disassociates the previously used output table.

* If you replace the default table name, OutputTable, with another name before you save the data query, the data builder registers a new output table and uses it with the data query.

You can enter a name that is up to 60 characters as the output table name in SAS metadata. SAS limits physical table names to 32 characters, and SAS uses the first 32 characters when it creates the output table. If a third-party vendor database product is used for the output table, then the number of characters might be less.

If you select a SAS LASR Analytic Server or SAS Data in HDFS library, the **Partition by** menu becomes available. After the settings for the output table are specified, the **Staging** section becomes available.

The following display shows an example of the Output and Staging properties.

Display 8.1   Output and Staging Properties

Staging Tables

About Staging Tables

Staging data is a best practice because you can use the data builder to access and transfer data from operational systems once, rather than frequently interfere with the operational systems and reduce their performance. Using the data builder to stage data can provide the advantage of adding calculated columns when you stage the data.

Like the output table, the staging table is registered in metadata when you save the data query. The physical table for the staging table is created when you run the data query.

You cannot specify the name for a staging table. The name of the output table is used, and an `__STG` (with two underscores) suffix is applied to the name. The suffix is used for the table name in metadata. The physical name of the staging table does not include the suffix.

Specifying a Staging Table

To use a staging table:

1 Select the top-level data query object from the **Properties** tab.

2 The settings in the **Output** section must be set.

3 Expand the **Staging** section, and select the **Use a staging table** check box.

4 Specify a library.

   Note: The data builder is initially configured to use the Visual Analytics Public LASR library and the Visual Analytics Public HDFS pair of libraries. If you specify different libraries, make sure that you understand how the path in HDFS is related to the server tag for the SAS LASR Analytic Server library. For more information, see *SAS Visual Analytics: Administration Guide*.

Output and Staging Table Interactions

The physical table is always replaced with the results of the data query. When you use a data query to append data and stage the data too, the staging table holds the data to append. As a result, the output table that the data query appends to is typically much larger than the staging table.

The following table identifies the supported combinations for output tables and staging tables. In addition, whether you can append data to tables is indicated.

Table 8.1 *Output Table and Staging Table Interactions*

Output Table	Staging Table	Append Data
SAS or DBMS *	None	Supported
SAS Data in HDFS	None	Not supported
Teradata or Greenplum **	None	Supported
SAS LASR Analytic Server	None	Supported
SAS or DBMS *	SAS or DBMS *	Supported
SAS LASR Analytic Server	SAS or DBMS *	Supported ***
SAS LASR Analytic Server	SAS Data in HDFS **Note:** When the staging table is in HDFS, SAS LASR Analytic Server is the only choice for the output table.	Not supported

Output Table	Staging Table	Append Data
SAS LASR Analytic Server	Teradata or Greenplum **	Supported ***

* The SAS or DBMS value represents data stored in SAS data sets or a third-party vendor database, respectively.

** The Teradata or Greenplum value represents deployments that use a third-party vendor appliance as a co-located data provider for SAS LASR Analytic Server.

*** Appending data is performed by the SAS LASR Analytic Server engine. Appends are not performed by having the server read data in parallel.

See Also

9

Working with Columns in Data Queries

Adding Columns to a Data Query

Understanding How the Data Builder Works with Columns

When you add a data source (a table or subquery) to the workspace, the columns from the data source are not automatically added as output columns. You need to add the columns to the data query that you want to use.

After the columns are added, you can specify column expressions and aggregations and use the sort and pivot by features.

Note: As an exception, the auto-aggregate feature requires that you set the default aggregations for the table before they are added to the data query.

When a data source is dropped onto the workspace, the column types are represented by the following icons:

Table 9.1 *Icons for Data Types*

Icon	Description
(123)	This icon represents numeric data.
(A)	This icon represents character data.
	Note: Date, time, and datetime data use this icon. After the column is added, the **Type** and **Format** columns are updated with information about the new column.

You can add columns to the data query in the following ways:

▪ Select the table in the workspace, right-click, and select **Add All Columns**

▪ To add one column from a table, select the column name with your pointer.

> **TIP** If you select the column name an additional time, the column is added to the data query again. This can be helpful if you want to use a column for both numeric and character data. When you add a column more than once, a number is added to the column name. If you change the column name, make sure that you do not have more than one column with the same name.

▪ Click the **Column Editor** tab and enter a column name, expression, and type in the bottom row of the table. The remaining fields are optional.

Removing Columns

To remove a column from a data query:

1   Click the **Column Editor** tab.

2   Select the column to remove, right-click, and select **Remove Column**.

Specifying a Column Expression

To specify a column expression:

1   On the **Column Editor** tab, specify a name for the column.

2   Select `table-name.column-name` from the **Expression** column. If you added a new
    column manually, then you need to make sure that you specified a table name and a
    column name.

3   Click to open the expression builder. The table name and column name are
    added automatically as the default SQL expression.

    The expression is limited to 1024 characters.

4   On the **Fields** tab, select columns from the source tables that you have added to the
    data query.

5   On the **Functions** tab, select the functions to apply to the source column.

6   Enter arithmetic operators and expressions such as CASE statements directly in the
    SQL expression area.

7   Click **Apply** to save the column expression.

Specifying Aggregations

To specify an aggregation for a column:

1   On the **Column Editor** tab, place your pointer in the **Aggregations** cell for the column and click. Click the ellipsis button to select the aggregations to use.

2   In the Choose Aggregations dialog box, select the **Aggregate functions** radio button.
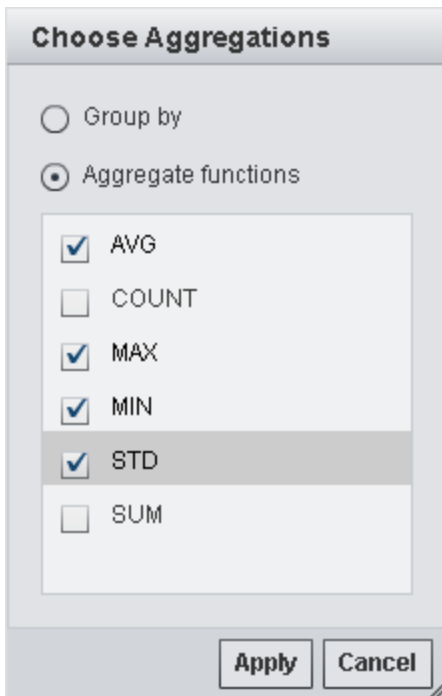
3   Select the check boxes for the aggregate functions to use. Click **Apply**.

   Note: After you click **Apply**, all of the other columns are automatically specified as GROUP BY columns. You can change a column from GROUP BY to an aggregate function by repeating this step for the column.

The following display shows how adding aggregations results in additional output columns. The aggregate function is appended to the column name.

#	Column Name	Type	Format
7	State	CHARACTER(25)	$25.
8	Quantity_STD	NUMERIC	
9	Quantity_MIN	NUMERIC	
10	Quantity_MAX	NUMERIC	
11	Quantity_AVG	NUMERIC	
12	Total_Retail_Price_STD	NUMERIC	
13	Total_Retail_Price_MIN	NUMERIC	
14	Total_Retail_Price_MAX	NUMERIC	
15	Total_Retail_Price_AVG	NUMERIC	
16	CostPrice_Per_Unit_STD	NUMERIC	
17	CostPrice_Per_Unit_MIN	NUMERIC	
18	CostPrice_Per_Unit_MAX	NUMERIC	

Column Editor | Where | Having | Joins | **Output Columns** | Messages | Log

Removing All Aggregations

To remove all the aggregate functions and group by settings:

1   On the **Column Editor** tab, select all the columns.

2   Right-click, and select **Remove Aggregations**.

> **TIP** This menu option is available only when all of the columns are selected.

Using Group By Variables

When you add an aggregation to a column, the remaining columns are automatically used as group by variables. The **Aggregations** column displays `GROUP BY` for these variables.

You can use a column as a group by variable by following the steps in "Specifying Aggregations" and selecting the **Group by** radio button.

Using the Auto-Aggregate Functions

The auto-aggregate feature is a productivity feature that enables you to specify a set of aggregations to apply as default aggregations to numeric columns for a specific table. A typical use is to automatically aggregate some of the columns in a fact table.

To use the auto-aggregate feature:

1   Select a table on the **Design** tab. On the **Properties** tab, the table properties are displayed.

2   Select **Enable** for **Auto-Aggregate**.

3   Click the ellipsis button next to **Functions** to open the Choose Aggregations window.

4   Select the check boxes for the aggregate functions that you want to apply, and then click **Apply**.

Whenever you add a column to the data query, the selected aggregate functions are automatically applied.

Using the Pivot By Feature

The pivot by feature provides an easy and powerful way to summarize data for analytics. You can specify a column to use as a categorical variable and the unique values to use. When the data query is run, the output table is summarized with the aggregations that you apply.

To use the pivot by feature:

1   On the **Column Editor** tab, place your pointer in the **Pivot By** cell for the column to use as the pivot column. Click the ellipsis button to select the pivot column and values.

2   In the Pivot Values dialog box, select the pivot by column. You can enter search criteria in **Filter fields** to filter the column names.

The following display shows an example of pivoting by three values in the Product_Category column.
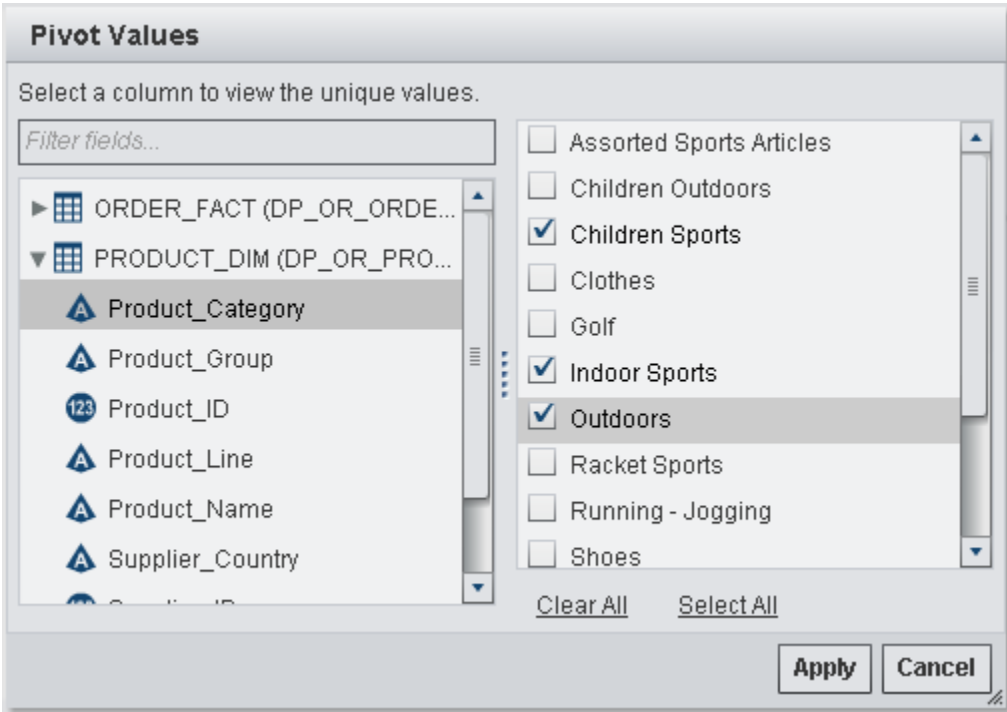
3 After the unique values for the column are loaded, select the check boxes for the values to use in the summarization. Click **Apply**.

The following display shows an example of the **Column Editor** tab when a pivot by column is used. The minimum and maximum Total_Retail_Price are calculated for each Customer_ID and are then pivoted by (transposed by) three values of the Product_Category column.

Display 9.1   *Column Editor Tab with a Pivot By Column*

		Column Name	Expression	Type	Format	Label	Aggregations	Pivot By
	1	Customer_ID	ORDER_FACT.Customer_ID	NUMERIC	13.		GROUP BY	
	2	Total_Retail_Pr...	ORDER_FACT.Total_Retail_Price	NUMERIC		TRP	MIN;MAX	PRODUCT_DIM.Product_Category
	+							

Tabs: Column Editor | Where | Having | Joins | Output Columns | Messages | Log

> **TIP** TRP is specified as the label for the Total_Retail_Price column. Look at the next display to see how the label is used to create labels for the new columns.

The following display shows how pivoting the Customer_ID column by three values of the Product_Category column results in additional output columns. A substring of the pivot by values is used as a prefix to each column name and the aggregate function is used as a suffix. The pivot by column label and aggregate function are used in the output column label.

Display 9.2 *Output Columns Tab with Pivot By Values*

Column Editor	Where	Having	Joins	Output Columns	Messages	Log

#	Column Name	Type	Format	Label
1	Customer_ID	NUMERIC	13.	
2	OUTD_Total_Retail_Price_MIN	NUMERIC		OUTDOORS_TRP_MIN
3	INDO_Total_Retail_Price_MIN	NUMERIC		INDOOR_SPORTS_TRP_MIN
4	CHI1_Total_Retail_Price_MIN	NUMERIC		CHILDREN_SPORTS_TRP_MIN
5	OUTD_Total_Retail_Price_MAX	NUMERIC		OUTDOORS_TRP_MAX
6	INDO_Total_Retail_Price_MAX	NUMERIC		INDOOR_SPORTS_TRP_MAX
7	CHI1_Total_Retail_Price_MAX	NUMERIC		CHILDREN_SPORTS_TRP_MAX

10

Working with Filters in Data Queries

About Filtering Data

You can use the **Where** and **Having** tabs to filter data in SAS Visual Data Builder.

The SQL expression on the **Where** tab is applied to the input data. This SQL expression is often used to subset data on the columns in the source tables.

The SQL expression on the **Having** tab can be applied to either the input data or the calculated columns. This SQL expression is typically used to subset data on calculated columns in the output table.

The maximum length for a WHERE or HAVING clause is 4096 characters.

Specifying a WHERE Clause

To add a WHERE clause to a data query:

1 Click the **Where** tab.

2 On the **Fields** tab, expand the table node, and select the column to use for filtering data.

3 Double-click the column or drag and drop it in the **SQL expression** area.

4 (Optional) You can click the **Functions** tab and select the functions to use with filtering.

5 For character variables, you can click **Unique Values** to load the unique values of the column. Select the check boxes to specify the unique values to include in the filter.

   Click **Apply** to add the unique values to the **SQL expression** area.

6 Edit the WHERE clause in the **SQL expression** area as follows:

Number of Values	Edit	Example
A single character value	Add an equal sign between the column name and the unique value.	`CARS.Make = 'Acura'`
More than one character value	Specify an IN operator and enclose the unique values in parentheses.	`CARS.Make IN ( 'Acura', 'Audi')`
Numeric comparison	Specify a numeric operator and a constant, or specify a numeric operator and another column name.	`CARS.Cylinders >= 6` `PRDSALE.Actual > PRDSALE.Estimate`

7   Click 🖫 to save the filter.

Specifying a HAVING Clause

To add a HAVING clause to a data query:

1   Click the **Having** tab.

2   On the **Fields** tab, expand the table node, and select the column to use for filtering data. The **Output Columns** node includes the calculated columns.

> **TIP** You can type in **Filter fields** to locate the column.

3   Double-click the column or drag and drop it in the **SQL expression** area.

4   (Optional) You can click the **Functions** tab and select the functions to use with filtering.

5   For character variables, you can click **Unique Values** to load the unique values of the column. Select the check boxes to specify the unique values to include in the filter.

   Click **Apply** to add the unique values to the **SQL expression** area.

6   Edit the HAVING clause in the **SQL expression**. The syntax is identical to that shown in step 6 of "Specifying a WHERE Clause".

7   Click 🖫 to save the filter.

Best Practices for Filters

SAS Visual Data Builder provides the **Where** and **Having** tabs for filtering data. Follow these best practices:

■ Filter on the **Where** tab first because the WHERE clause reduces the rows to consider for further subsetting.

■ When you specify a filter on a column that is an index or primary key, avoid using a function whenever it is possible. Using a function, such as `CAST(order_id as DOUBLE)`, risks performing a full-table scan rather than using the index.

■ Filter on the table that has the smallest number of rows first.

■ If you need to remove a table that is used on the **Where** or **Having** tabs, then remove the reference to the table from the filters before you remove the table. Otherwise, if you remove a table that is referenced in a filter, then all of the filtering conditions are cleared.

11

Working with Joins in Data Queries

About Joins

SAS Visual Data Builder supports joins for tables and subqueries. You can join tables to each other, including self joins. You can join subqueries to tables and join subqueries to subqueries. When you use a subquery in a join, the join condition is made against the output table for the subquery. The data builder supports joining up to 256 tables.

When you drop a table or subquery onto the workspace, the data builder attempts to determine a join condition automatically. When the data builder creates a join

automatically using foreign keys or by matching columns, the join is added as an inner join. The data builder also supports left, right, and full joins. You need to specify the join type manually if you do not want an inner join.

In addition to easily adding joins to a data query, the data builder generates an SQL statement with all of the joins declared explicitly. For example, you can specify an inner join in a WHERE clause, such as `WHERE t1.order_id=t2.order_id`. However, mixing inner joins in a WHERE clause and outer join types in a single data query can be complex to read and understand.

SAS Visual Data Builder always generates an SQL statement with inner joins declared explicitly. For example, see the following code sample:

```
LEFT JOIN
        LIB1.TRANSACTION_TYPE_DIM TRANSACTION_TYPE_DIM
            ON CASH_FLOW_FACT.TRANSACTION_TYPE_KEY =
                TRANSACTION_TYPE_DIM.TRANSACTION_TYPE_KEY
    INNER JOIN
        LIB1.TRANSACTION_DIM TRANSACTION_DIM
            ON PARTY_DIM.SEGMENT_ID = TRANSACTION_DIM.SEGMENT_ID
            AND PARTY_DIM.TRANSACTION_KEY = TRANSACTION_DIM.TRANSACTION_KEY
```

How Does the Automatic Join Feature Work?

Feature Overview

SAS Visual Data Builder attempts to join tables and data queries automatically as you add them to the workspace. When you drag and drop a table or data query onto the workspace, information about the table or data query is retrieved from the SAS Metadata Server. For subqueries, the metadata for the subqueries' output table is retrieved.

Using Foreign and Primary Keys

If primary key or foreign key information is registered in the metadata for the table that you drag and drop onto the workspace, then the data builder retrieves the foreign key and primary key information.

The data builder then iterates over each of the tables that are already in the workspace in the same sequence in which they were added to the workspace. The data builder retrieves the foreign key and primary key information for the table, and compares the length, type, and name with the key columns for the newly added table. If a match is found, then the tables are used in the join, and the columns are added as a join condition. The data builder continues to search for matches between the two tables, and it adds join conditions when possible. After a set of join tables is identified, the data builder does not continue iterating over the tables that are already in the workspace.

Matching by Name

If there is no foreign key or primary key information for the table that is dragged and dropped onto the workspace, then the data builder does not use foreign key or primary key information for the tables already in the workspace.

The data builder retrieves the column information for all of the columns in the newly added table. The data builder then iterates over each of the tables that are already in the workspace in the same sequence in which they were added to the workspace. The data builder compares the length, type, and name for each column with each column in the newly added table. If a match is found, then the tables are used in the join, and the columns are added as a join condition. The data builder does not continue to search for matches between the two tables, and it does not continue iterating over the tables that are already in the workspace.

Selecting the Join Type

When the data builder finds the first set of matching columns by comparing keys or matching names, it sets the join type for the two tables. The data builder checks the metadata for the columns to determine whether the columns are nullable. (Many third-

party vendor databases support **NOT NULL** as a constraint for a column when the table is created in the database with SQL. The constraint ensures that there are no missing values for the column.) The data builder performs the following steps to set the join type:

1   If the column for the existing table is not nullable, but the column for the newly added table is nullable, the tables use a left join.

2   If the column for the existing table is nullable, but the column for the newly added table is not nullable, the tables use a right join.

3   If both columns for the existing table and newly added table are nullable, the tables use a full join.

4   If none of the previous conditions are met, the tables use an inner join.

You can specify the join type by right-clicking on the join icon, and selecting the join type from the menu.

Adding a Join

Dragging and Dropping Join Lines

To add a join by selecting a column name and dragging your pointer to another column:

1   Place your pointer on one of the columns in a table to use, and then click and drag your pointer to a column in the other table to use in the join.

2   If you want a join type other than an inner join, select the join, right-click, and then select the join type.

> **TIP** The first table that is added to the workspace is always set as the left table. You can switch the right table and left table by using the Joins tab.

Using the Joins Tab

To add a join manually:

1 Click the **Joins** tab.

2 Click **+**.

3 Use the menus to replace **Not Selected** with the table names to use in the join. Click **Save**.

4 The default join type is an inner join. Use the menu in the **Join Type** column to select a different join type.

5 Click **+** in the **Join conditions** area.

6 Use the menus to replace **Not Selected** with the column names to use in the join condition. Click **Save**.

7 Click 💾 to save the join condition with the data query.

Removing a Join

You can remove a join by using either of the following methods:

▪ Select the join in the workspace, right-click, and select **Remove Join Condition**.

▪ Click the **Joins** tab, select the row in the table, and click ✕.

Managing Joins in a Data Query

Once a join is added to a data query, you can change the join by selecting it in the workspace, right-clicking, and changing the join type or removing the join condition.

You can also change a join by clicking the **Joins** tab, and then selecting the row in the table. You can make the following changes:

- add and remove tables from the join list

- reorder the sequence of joins by moving them up or down

- switch the left table and right table assignments for a join

- add, remove, and change the columns that are used in the join condition

The **Joins** tab shows the join condition for the entire data query. Make sure that you select a row in the upper table to set the columns in the **Join conditions** area.

The workspace shows a link between the tables that are used in a join. If you specify a left join or right join, the Venn diagram icon reflects which table provides the bulk of the data. If you rearrange the tables in the workspace (switching the left table and right table), the Venn diagram icon continues to reflect which table provides the bulk of the data. In order to change the data relationship, use either of the following methods:

- Switch the left table and right table by selecting a row from the table, and clicking ⬚.

- Change the join type from left to right or from right to left.

Best Practices for Managing Joins

In most cases, the sequence in which joins are specified on the **Joins** tab does not matter. However, it is possible that the query optimizer for the data source might perform the joins in a sequence that reduces performance. In these rare cases, you can change the sequence in which joins are specified on the **Joins** tab by selecting the join in the list, and clicking ⬆ or ⬇. This join sequence is still determined by the query optimizer, but you can control how the SQL for the join is presented to the query optimizer.

12

Creating LASR Star Schemas

What Is a LASR Star Schema?

A LASR star schema is very similar to a typical star schema in a relational database. Imagine a single fact table that is surrounded by dimension tables. Each dimension table is joined to the fact table with a dimension key.

Single-level star schemas are supported. Snowflake schemas are not.

Creating a LASR Star Schema

To create a LASR star schema:

1   Select **LASR** ▷ **Create a Star Schema**.

2 Drag and drop tables or data queries from the SAS Folders tree onto the workspace. Keep the following suggestions in mind:

   ▪ Drag and drop the fact table first and then the dimension tables.

   ▪ Any input tables that are not already in a SAS LASR Analytic Server library are loaded to memory when the star schema is run.

   ▪ If a dimension table shows the ⊖ icon, then it usually indicates that the data builder could not determine the join condition for the dimension table. Position your pointer over the icon for information about how to correct the incomplete table status.

3 (Optional) If you do not want a column from the dimension table in the output table, select the column in the table in the workspace, right-click, and select **Remove Column**.

   Note: All of the columns from the fact table are automatically selected and required to be included in the output table.

4 (Optional) Click the **Joins** tab to review the columns that are used to generate the output table.

5 (Optional) Specify a different column prefix. For more information, see "Column Prefixes".

6 Click the **Output** tab, and specify a name for the output table.

7 Click 🖫, and specify a name and location for the LASR star schema.

8 Click ▶ to generate the output table.

How Are Tables Used?

Input Tables

The fact table and dimension tables are used in the star schema after you drag and drop them onto the workspace.

Because the star schema is formed in memory by the server, the first step is to make sure that the tables are loaded to memory. Be aware that the tables are transferred each time the star schema runs and this can reduce performance. Conversely, if the tables are already loaded to memory on the server, the star schema runs and generates the output extremely fast.

Column Prefixes

The first 15 characters of a dimension table's name and the underscore character are initially set as a prefix for the column names from the dimension table. Column names for the output table are a combination of the prefix and the original column name.

However, you can specify a different value for the prefix after you select the table name from the menu at the top of the **Inputs** tab.

Column names from the fact table are not modified with a prefix.

> **TIP** A column name is limited to 32 characters. If you have a long column name, then reducing the prefix can help you keep more of the original column name.

Output Table and Conserving Memory

To use memory efficiently, the default output for the schema is a view. The **Create output as a view** check box on the **Output** tab is selected by default to create a view. If you clear the check box, then the output is a table.

When the output is a view, the rows are formed in buffers rather than creating the entire table at once. This enables you to access a view whose data is larger than the memory capacity of the system.

If you clear the **Create output as a view** check box to create a table from the star schema, then the system must have enough available memory to store the table. If the system runs out of memory while running the star schema, the memory that was used for the output table is freed, and you receive an error message.

If you use SAS LASR Analytic Server tables for input to the schema, then the output table for the star schema must use the same library.

13

Working with SAS LASR Analytic Server

Understanding Co-Located Data

One of the most powerful benefits of SAS LASR Analytic Server is the ability to read data in parallel from a co-located data provider. In this configuration, the SAS LASR Analytic Server software is installed on the same hardware as the data provider. The currently supported data providers are the following:

- SAS High-Performance Deployment of Hadoop

- Teradata Data Warehouse Appliance

- Greenplum Data Computing Appliance

SAS Visual Data Builder excels at accessing data from a variety of data sources, performing basic data preparation, and then staging the prepared data with a co-located data provider. After the data is staged, the server can load the data to memory for exploration, reporting, and further in-memory data preparation.

The SAS Deployment Wizard registers a default library for the SAS LASR Analytic Server and a library for the co-located data provider. For the third-party vendor products, the default library name and location are as follows:

Teradata Data Warehouse Appliance
```
/Shared Data/SAS Visual Analytics/Public/Visual Analytics
Public Teradata
```

Greenplum Data Computing Appliance
```
/Shared Data/SAS Visual Analytics/Public/Visual Analytics
Public Greenplum
```

Predefined Libraries

SAS automatically configures two SAS LASR Analytic Server libraries. The names and purposes of these two libraries are as follows:

Visual Analytics Public LASR
> This library is initially configured as the default library for the data builder. It is automatically selected as the output library for importing self-service data, creating data queries, and preparing in-memory data.
>
> This library is intentionally configured as a general-purpose library with extremely limited restrictions for securing data access.

Visual Analytics LASR
> This library is initially configured to provide Write permission to data builder users and administrators only. The initial security settings might align with organizational goals for securing and managing data access.

For deployments that use a distributed SAS LASR Analytic Server, a pair of libraries for the co-located data provider are configured with permissions that match the same initial security settings.

Using SAS Data in HDFS Libraries

Default Library

When your deployment includes SAS High-Performance Deployment of Hadoop, the SAS Deployment Wizard registers a library for it. This library is available for use in the SAS Folders tree, and it is located in `/Shared Data/SAS Visual Analytics/ Public/Visual Analytics Public HDFS`.

Staging Library

You can specify a SAS Data in HDFS library as a staging library. This is a common use because the rows for the output table are distributed among the machines in the cluster. The server can then read the data in parallel when it loads the table to memory.

You must specify a SAS LASR Analytic Server library for the output library when you use a SAS Data in HDFS library for staging.

Output Library

You can specify a SAS Data in HDFS library as an output library. The engine distributes the rows for the table to the machines in the cluster. Afterward, you can select the table from the SAS Folders tree, right-click, and select **Load a Table**. This menu option loads the table from HDFS to memory on a SAS LASR Analytic Server.

You can partition SAS Data in HDFS tables when they are used in an output library. You can select a column to use from the **Partition by** menu. Partitioning the table ensures that all of the rows with the same formatted value as the selected column are distributed to one machine in the cluster. The rows are also placed in the same HDFS block. When you load a partitioned table to memory, the partitioning information is retained, and the result is a partitioned in-memory table.

See Also
"Partitioning Tables" on page 92

Restrictions

The following restrictions apply to using SAS Data in HDFS libraries with SAS Visual Data Builder:

■   You cannot specify a SAS Data in HDFS library as an input library because the SAS Data in HDFS engine is a Write-only engine.

■   The **Append data** check box on the **Properties** tab is disabled. The SAS Data in HDFS engine does not support appending data.

■   If you specify a SAS Data in HDFS library as an output library, you cannot view the results on the **Results** tab because the SAS Data in HDFS engine is a Write-only engine.

Using SAS LASR Analytic Server Libraries

Default Library

During installation, the SAS Deployment Wizard registers a library for SAS LASR Analytic Server. This library is available for use in the SAS Folders tree, and it is located in `/Shared Data/SAS Visual Analytics/Public/Visual Analytics Public LASR`.

Input Libraries

When you select a SAS LASR Analytic Server table as an input table, be aware of the following best practices if the table is large. If the table is not large, then using it for input requires no special considerations.

Here are the considerations for using a large SAS LASR Analytic Server table as an input table:

- A WHERE clause is processed in memory by the server if no aggregations or joins are used. Specify a filter on the **Where** tab so that you use only the rows that you want.

- If you want to join the table, then create one data query that copies the data from the server to the same library as the table that you want to join it with. Specify a filter on the **Where** tab, if applicable. Then, create another data query that performs the join.

    TIP As an alternative, you can create a LASR star schema so that the table is joined in memory with other in-memory tables.

- You cannot join two tables from the SAS LASR Analytic Server and then transfer the output table back to the same server. Create a LASR Star Schema instead.

Output Libraries

A SAS LASR Analytic Server library is initially configured as the default output library. When you run a data query, you can click 🔣 **Explore Results** to explore the table immediately.

You can select the **Append data** check box on the **Properties** tab to add rows to an in-memory table. However, this option is not available if you use a SAS Data in HDFS library as a staging library. Instead, for deployments that use HDFS, you can use the data query to append rows, and then save the in-memory table to HDFS.

Partitioning Tables

When you specify a SAS LASR Analytic Server or SAS Data in HDFS library as an output library, you can specify a partition key for the table. You can select a column to use from the **Partition by** menu.

Partitioning uses the formatted values of the partition key to group rows that have the same value for the key. All of the rows that have the same value for the key are loaded to a single machine in the cluster. For SAS LASR Analytic Server libraries, this means that the rows that have the same value for the key are in memory on one machine. For SAS Data in HDFS libraries, all of the rows that have the same value for the key are written to a single file block on one machine. (The block is replicated to other machines for redundancy.) When the partitioned table is loaded onto a server, the partitioning remains when it is in memory.

If you select a partition key and also specify sort options for columns on the **Column Editor** tab, the sort options are passed to the engine in an ORDERBY= option. This enhancement applies to SAS LASR Analytic Server and SAS Data in HDFS libraries and can improve performance once the data is in memory.

When you specify a partition key, avoid using a variable that has few unique values. For example, partitioning by a flag column that is Boolean results in all rows on two machines because only two values are available. At the other end of the spectrum,

partitioning large tables by a nearly unique key results in many partitions that have few rows.

Determining the optimal partition key can be a challenging task. However, as an example, if you tend to access data based on a customer ID, then you might improve performance by partitioning the data by customer.

See Also

SAS LASR Analytic Server: Reference Guide

Loading a Table (As Is) to SAS LASR Analytic Server

The data builder is typically used to perform basic data preparation. However, there are cases when a table is already prepared for analysis. In these cases, you can load a table to memory on SAS LASR Analytic Server without any modification.

To load a table into memory on a SAS LASR Analytic Server:

1   Use the SAS Folders tree to locate the table.

> **TIP** You can also click 🔍, and search for the table by name and location.

2   Select the table, right-click, and select **Load a Table**.

3   The fields in the **Source Table** section are filled automatically.

Specify the following settings in the dialog box:

Table 13.1   Load a Table Dialog Box Properties

Field	Description
Name	Enter the name to use for the table. This field is initially set to the same name as the source table. This name is registered in metadata with the SAS LASR Analytic Server library.
Description	Enter a description of the data. This description is displayed by SAS clients such as SAS Visual Analytics Explorer.
Location	This field is initially set to a default folder. Click **Browse**, and select a folder to use for the table metadata.
Library	This field is initially set to a default library. Click **Browse**, and select the SAS LASR Analytic Server library to use with the table.

4   Click **Submit**.

Saving an In-Memory Table to HDFS

This feature is available for deployments that use SAS High-Performance Deployment of Hadoop only.

To save an in-memory table to HDFS:

1   Select **LASR ▸ Save a Table to HDFS**.

2   Click ⊞ to select the SAS LASR Analytic Server table that you want to save.

3   Make any necessary changes to the **Location** and **Library** fields in the **HDFS Table** section.

   Note: The data builder is initially configured to use the Visual Analytics Public LASR library and Visual Analytics Public HDFS library. If you specify different libraries, make sure that you understand how the path in HDFS is related to the server tag for the SAS LASR Analytic Server library. For more information, see *SAS Visual Analytics: Administration Guide*.

4   (Optional) Specify options. By default, the data builder replaces the table in HDFS if it already exists. You can specify the number of redundant copies for the table.

5   Click 💾, and specify a name and location.

6   Click ▶ to save the table to HDFS.

Appending In-Memory Tables

How Does Appending In-Memory Tables Work?

The data builder provides a very convenient way to add incremental data into a single table. For example, if sales data is loaded into memory on a daily schedule, you can append the numerous daily tables into a separate table that contains the cumulative sales data.

For the purpose of comparison, you can create a data query to append to an in-memory table with a data query. One reason to append entire tables is that you can also access the tables individually.

Append In-Memory Tables

To append in-memory tables:

1   Select **LASR ▶ Append Tables**.

2   Select the table that you want to append data to from the **Available tables** list, and move it to the **Base table** field.

> If the table that you want to use is not listed, click 🗃. If a table is added to the library after you click 🗃, the list does not automatically refresh to show the new table. Select the same library again to refresh its contents.

3   Add tables to the **Source tables** list. These are the tables that have the rows to append to the base table.

4   (Optional) If you select the **Unload source tables after appending them to the base table** check box, the source tables are dropped from memory after the append completes.

> Only select this option if you do not want to access the source tables individually.

5   Click 💾, and specify a name and location.

6   Click ▶ to append the source tables.

How Are Tables Used?

Source Tables

The following list identifies the considerations for source tables:

■   The source tables and base table must be in-memory.

■   The source tables can be views. However, the base table cannot be a view.

■   If you append tables with character columns of different lengths, the length of the column in the base table takes precedence. For example, if you have a column in the base table that is 15 characters and a column in the source table is 20 characters, the data from the source table's column is truncated to 15 characters.

■   If the base table is partitioned, then the source tables must be partitioned on the same variables.

- If the base table has a column that is not present in a source table, then the rows in the source table receive a missing value for the column that is not present in the source table.

- If a source table has columns that are not present in the base table, then the columns are dropped and are not appended to the base table. The base table always maintains the same number of columns.

- The order of the columns in a source table does not need to match the order of the columns in the base table.

- The data builder does not prevent you from appending the same source table more than once. Review the source table list to ensure you did not accidentally use a source table more than once.

Conserving Memory

After you append the source tables, the base table grows to accommodate the additional rows. This consumes memory on the SAS LASR Analytic Server.

If you do not need to access the source tables after appending them, you can select the **Unload source tables after appending them to the base table** check box. This option conserves memory, but you need to ensure that the source tables are loaded to memory again (often with newer data) before you can run the append again.

Monitoring Memory Usage

A memory gauge is visible in the menu bar that shows the amount of physical memory that is used on the server.

The server rejects requests to add tables or append rows when 80% or more of the memory is used. (The default value for the server is 75%. SAS Visual Analytics deployments set the value at 80%.) System administrators can specify a different value for the server with SAS Management Console.

For non-distributed deployments, the 80% threshold applies to the single machine. For distributed deployments, the 80% threshold is evaluated for each machine in the cluster.

If the threshold is crossed, the operation fails, and the **Messages** tab includes an error that is similar to the following example:

```
ERROR: A server-side limit on the consumption of memory resources has
been reached. These limits can be adjusted by the owner or by the
administrator of the LASR Analytic Server.
```

In some deployments with limited network bandwidth, the data builder cannot communicate with the SAS LASR Analytic Server monitor, and the memory gauge fails to appear.

See Also

SAS Visual Analytics: Administration Guide

14

Importing Data

About the Import Data Features

Why Use the Import Data Features?

The import data features enable you to perform self-service data access. You can take a Microsoft Excel spreadsheet, delimited text file, or SAS data set from your desktop, and transfer it to SAS LASR Analytic Server. Or, you can transfer the data to a SAS data set, a database table, or any other source that SAS can access.

For deployments that include SAS/ACCESS Interface to Oracle, you can import tables from Oracle. This is also considered self-service and you can avoid registering the server and a library with the SAS Metadata Server.

The imported data can be used to enrich existing data that is already available in your SAS deployment. You can import data directly to memory on a SAS LASR Analytic Server.

Requirements

Importing data requires starting a SAS session on the SAS Application Server. Typically, this requirement is met by each user who has a host account.

For deployments on Microsoft Windows, the host account must have the **Log on as a batch job** Windows privilege. For deployments on Linux that use a distributed SAS LASR Analytic Server, the host account must be configured for passwordless SSH.

For more information, see "Getting Started" in the *SAS Visual Analytics: Administration Guide*.

Table Names, Column Names, and Special Characters

In general, you can import files that use blanks and special characters in the filenames and column names. The following list identifies how table names are handled:

- For text files (such as CSV files), the table name is initially set from the filename.

- When you import a spreadsheet, the table names are handled as follows:

  - If the spreadsheet contains a single worksheet, the table name is initially set from the filename.

  - If the spreadsheet contains multiple worksheets, each table name is initially set as a combination of the filename, an underscore, and the name of the worksheet.

- After the initial table name is determined, an unsupported special character is replaced with an underscore. These special characters include / \ * ? " < > | : - and the period (.).

- Table names are shortened to 32 characters because that is the table name length that is supported by SAS.

For column names, most special characters can be used, including spaces. An unsupported special character is replaced with an underscore. These special characters include / \ * ? " < > | : - and the period (.).

By default, column names are read from the first row of the text file. If you clear the **Includes column names** check box, the column names are generated for you as follows:

- **Spreadsheets** Column names are assigned A, B, C, and so on.

- **Text files** Column names are assigned VAR1, VAR2, and so on.

The following table identifies how column names that begin with a number, such as `2013sales`, or use numbers only, such as `2013`, interact with the **Includes column names** check box:

Source File Column Name	Imported Column Name
`2013sales`	`_2013sales` for XLSX, XLS, and text files.

Source File Column Name	Imported Column Name
`2013`	`_2013` for XLSX and text files.
	For Excel 97-2003 Workbook (XLS) files, a letter such as `A`, `B`, `C`, and so on, is substituted for the column name, depending on the column position in the file.

Working with Spreadsheets

The following list identifies considerations for importing data from spreadsheets:

- You can import Excel Workbook (XLSX) files and Excel 97-2003 Workbook (XLS) files. You cannot import XLSM, XLST, or other Excel file types.

- When you import a spreadsheet (from your PC) that has multiple worksheets, by default, all of the worksheets are imported. A table is created for each worksheet. You can clear the check boxes for the worksheets that you do not want to import.

- Importing pivot tables is not supported.

Importing Large Data Files

When you import data files (spreadsheets, SAS data sets, or delimited text files) from your machine, you are limited to files that are 4 GB or less. This limitation is introduced by web browsers. Despite being a limitation, it is practical because of the time that is required to transfer large files through the web browser to the SAS server.

Note: Administrators can specify a limit that is less than 4 GB.

As an alternative to importing data files through your web browser, you can transfer files to the machine for the SAS server and use autoload. You can import data from files that are larger than 4 GB. In addition, FTP and network file systems transfer data files faster than web browsers.

Autoload requires some setup by an administrator, so contact your administrator.

Usage Notes

Review the following notes if you have trouble importing data:

- Before you click **OK** to import the data, click **Preview**. Preview shows an accurate representation of the column names and data values that will be available after the import.

- If SAS is configured as a Unicode server at your site, you have the most flexibility for importing data.

- When you import a delimited text file (CSV file), you need to specify the encoding of the text file. In some cases, the import reports success, even though the data might be corrupted. It is important to verify the imported data.

- If you import a SAS data set that uses user-defined formats, you must make the custom format catalog available to the SAS Application Server. For more information, see "Working with User-Defined Formats" on page 43.

- If importing large data files at the same time is common for your deployment, be aware that the large data files are written to temporary disk space on the server. In extreme cases, this can cause temporary disk space to become full. Systems that run out of disk space can become unresponsive and difficult to troubleshoot.

- If you import data from text files and plan to append the data, check that the column data types and lengths match the table that you want to append to.

Depictions of Importing Data

Import Data Files from Your Computer

The following figure depicts how a file (SAS data set, spreadsheet, or delimited file) is accessed from the client PC, transferred to the SAS Workspace Server, and then stored in an output table.
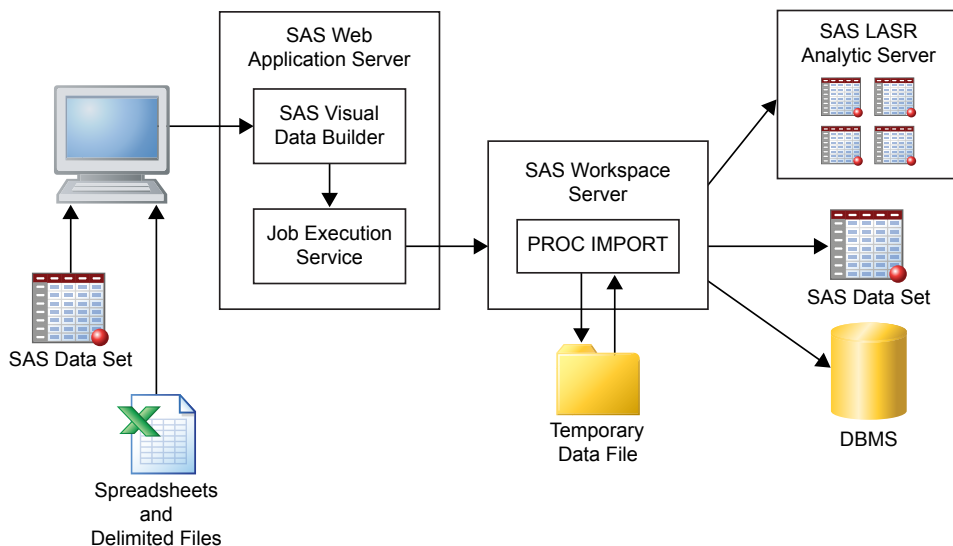
Note: When you import a SAS data set, it is not processed with PROC IMPORT. SAS data sets are transferred to output with a DATA step.

You can import data files that are available from the file systems on your PC. This includes local file systems such as `c:\` on Windows machines, and paths such as `/home/$USER` on UNIX machines. Network file systems and shares are also included, such as UNC paths like `\\nas\spreadsheets`.

For detailed steps, see .

Import SAS Data Sets on a SAS Server

The following figure depicts how the client PC can be used to specify a SAS data set on the SAS Application Server machine (represented as the SAS Workspace Server), and then load the data set to memory on a SAS LASR Analytic Server.

SAS Data Set
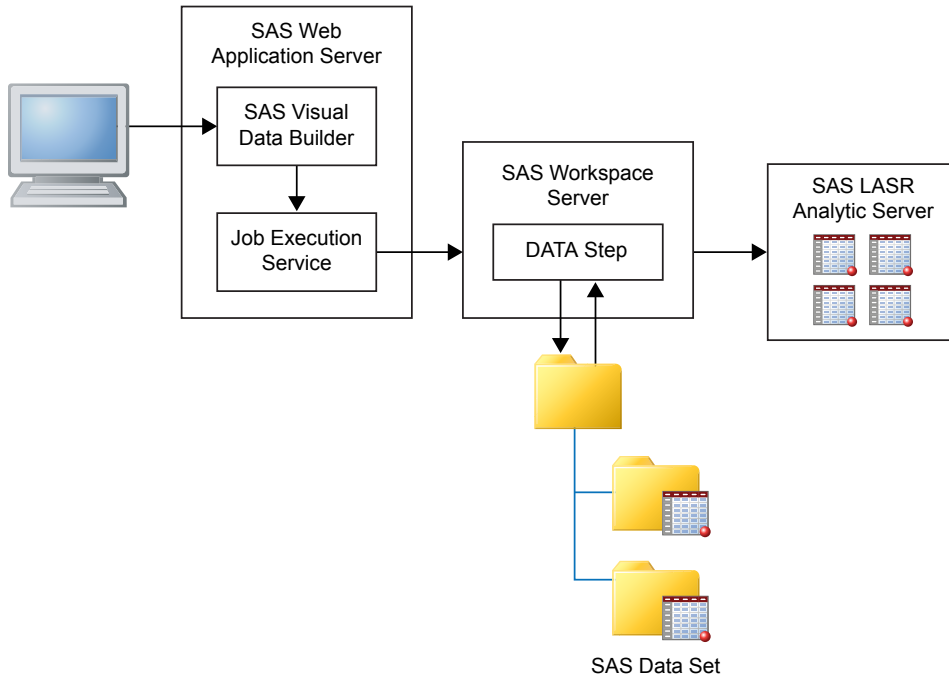
You can direct the SAS Workspace Server to access the file systems on the server. For example, if you have a large data set, you can use FTP or another method to copy it to a directory on the server, and then use the server to import it. For detailed steps, see "Importing a SAS Data Set on a Server" on page 108.

Importing a Local Data File

To import data from a spreadsheet, delimited text file, or SAS data set, follow these steps:

1   Select **File** ▶ **Import Data** ▶ **Local File**, and then select the data file.

2   In the Import Local File window, specify the following input file parameters:

Note: For SAS data sets, there are no parameters to specify.

Spreadsheet options
> The following options are available when you import a Microsoft Excel spreadsheet:
>
> **Select worksheet**^{Applies to XLS}
>> Select **All** or the check boxes for the worksheets to import. For more information, see "Working with Spreadsheets" on page 102.
>
> **Begin import on row**
>> The default is to import data from the first row of the spreadsheet. If the data begins on a different row, make changes.
>
> **Includes column names**
>> Select this check box when the row on which to begin the import has the column names.

Text file options
> The following options are available when you import a delimited text file:
>
> **Delimiter**
>> Select the delimiter that is used in the text file to import. You can specify a single character to use as a user-defined delimiter.
>
> **First row contains column names**
>> Select this check box when the text file includes the column names on the first row of the file.

Data records begin on row
The default is to import data records from the first row. If you select **First row contains column names**, then this value increments to 2.

Number of rows to scan
The default is to read up to the first 500 rows from the file to determine the data type and length for each column. Smaller values cause the import to complete quickly, but you risk the chance of determining too short a value for character columns. Larger values reduce the chance of truncating character columns, but they increase processing time.

Encoding
Select the encoding of the text file. If you are importing UTF-8 or UTF-16 data, make sure that the SAS Application Server is a Unicode server or that all of the file contents can be successfully transcoded to the encoding of the SAS Application Server.

3   (Optional) Click **Preview** to view the data. Preview displays up to 500 rows from the file.

> **TIP** Previewing the data can help you determine whether you specified the correct encoding.

4   (Optional) In the **Output Table** section, enter the name for the table and a description. Review the library and location settings by clicking **Advanced** and make changes if necessary.

5   Click **OK**.

Note: When users import a data file from the explorer or designer and do not have the Build Data capability, the **Advanced** information in step 4 is not available. These users can import the data to a general-purpose area, or select **Store the data in a private location** to prevent other users from accessing the data.

Importing a SAS Data Set on a Server

To import a SAS data set that is accessible from your SAS Application Server, follow these steps:

1   Select **File** ▸ **Import Data** ▸ **SAS Data Set on a Server**, and then select the data file.

   Note: Remember that the data files and directories are on the remote machine, not on your PC. In the **Basic Properties** table, the **Date created**, **Description**, and **Keywords** fields are not used.

2   Navigate to the SAS data set, and click **OK**.

3   (Optional) In the **LASR Table** section of the Import SAS Data Set on a Server window, enter the name for the table and a description. Review the library and location settings by clicking **Advanced** and make changes if necessary.

4   Click **OK**.

Importing an Oracle Table

To import a table from an Oracle database to SAS LASR Analytic Server, follow these steps:

1   Select **File** ▸ **Import Data** ▸ **Oracle Table**.

2   Specify your connection information, and then click **Browse** to select the Oracle table to import.

3 (Optional) In the **LASR Table** section of the Import Oracle Table window, enter the name for the table and a description. Review the library and location settings by clicking **Advanced** and make changes if necessary.

4 Click **OK**.

SAS/ACCESS Interface to Oracle must be configured on the machine for the SAS Application Server. The value for the **Path** field is related to the net service name in the tnsnames.ora file. The tnsnames.ora file is generated during the Oracle client installation on the machine for the SAS Application Server. The file is typically stored in an Oracle installation directory such as `/opt/oracle/app/oracle/product/10.2.0/db_1/network/admin/tnsnames.ora`. The net service name for the connection information is contained in this file. See the following figure:

```
# tnsnames.ora Network Configuration File:
C:\oracle\product\10.2.0\client_1\network\admin\tnsnames.ora
# Generated by Oracle configuration tools.

NEWSERVER10G =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP)(HOST = server.na.sas.com)(PORT = 1521))
    )
    (CONNECT_DATA =
      (SERVICE_NAME = server10G)
    )
  )
```

Importing a Twitter Stream

Before You Begin: API Key

You can use the data builder to import tweets from Twitter that meet the search criteria that you specify. The data builder uses the search API that is offered by Twitter. To perform the search, you must create a Twitter application record with Twitter for your installation of the data builder to get the API key.

Twitter provides information about creating a Twitter application record and getting an API key at `https://dev.twitter.com/docs/faq#7447`.

The high-level steps are as follows:

1   Log on to `https://dev.twitter.com/apps` with your Twitter account.

2   Create a new Twitter application. You need to specify a name, description, website, and callback URL, and then accept the terms and conditions.

    Note: The values for these fields are not used by the data builder. You need to create the Twitter application only so that you can get the API key.

3   After you create the Twitter application, record the values for the following fields:

   ■ Consumer key

   ■ Consumer secret

   ■ Access token

   ■ Access token secret

   You need to enter these values when you set up access to the Twitter stream in the data builder.

Importing Tweets

To search for tweets and load them into SAS LASR Analytic Server, follow these steps:

1   Select **File** ▸ **Import Data** ▸ **Twitter Stream**.

2   Enter a search term and the maximum number of tweets to return.

   TIP The search operators that you can use are described at `https://dev.twitter.com/docs/using-search`.

   Each tweet that matches the search term is stored as a row in a SAS LASR Analytic Server table.

Note: SAS does not support search terms that include double-byte characters.

3 (Optional) In the LASR Table section, enter the name for the table and a description. Review the library and location settings by clicking **Advanced** and make changes if necessary.

4 If this is the first time you are importing tweets, then enter values for the fields in the **Twitter API Key** section.

5 (Optional) Review the **Proxy Server** section and make changes if necessary.

6 Click **OK**.

Note: Rate limits might apply to the Twitter service. Such limits are outside of the control of SAS. You are required to follow all applicable terms of use that Twitter and others might promulgate as to the Twitter data.

Note: Certain functionality in the software enables the user to invoke external third-party resources. Be aware that use of these resources might result in disclosure and transmission of information that the user submits to these resources.

15

Importing SAS Information Maps

About SAS Information Maps

A SAS Information Map is descriptive information that is used by many SAS solutions and Business Intelligence products. An information map is applied to the data sources in your data warehouse to describe the structure and content of data. An information map does not contain any physical data. Information maps provide business users with a user-friendly way to understand data so that they can perform data queries and get results for themselves.

SAS Information Maps are created with either SAS Information Map Studio or the INFOMAPS procedure. SAS Visual Data Builder cannot create, edit, or save an information map. However, the data builder can import the business logic from a relational information map and represent it as an SQL query.

See Also

Base SAS Guide to Information Maps

Importing a SAS Information Map

To import a SAS Information Map:

1 Use the SAS Folders tree to locate the map ().

2 Select the map, right-click, and select **Import Query**.

A new data query is created with the same name as the map. The output table is automatically set with the same name as the map.

The business logic is imported from the map and represented as an SQL query on the **Code** tab.

> **TIP** If you want to join a table, save the data query, and then create a new data query and add it as a subquery.

See Also

Limitations and Restrictions

SAS Visual Data Builder has the following limitations and restrictions for working with information maps:

■ The data builder can use information maps that are based on relational tables. Information maps that are based on OLAP cubes cannot be used.

■ The information map cannot use more than 50 physical tables. A table that is used more than once in a self join counts as one physical table.

- The information map cannot use more that 5000 physical columns. A column that is used in more than one data item counts as one physical column.

- Prompts are not supported. Even if a prompt has a default value, the default value is not included.

- Data items that are based on business data are not supported. For example, the equation `Dataitem1 = Year + 2` (where Year is a data item) is not included in the data query.

- Data items that are based on physical columns are not supported. For example, the equation `Dataitem2 = FirstName || LastName` (where FirstName and LastName are columns) is not included in the data query.

- The output table is automatically named with the same name as the map. The name field is limited to 32 characters.

16

Supporting Text Analysis

Features Overview

SAS Visual Analytics can perform text analysis with the explorer. To enable text analysis, perform the following data administration tasks with the data builder:

■ Load a stop list to filter out commonly used words.

■ Add a numeric key to existing tables that have unstructured text.

■ Acquire Twitter data to analyze tweet trends. For information about acquiring Twitter data, see "Importing a Twitter Stream" on page 109.

Loading a Stop List

A stop list is a table of words that you want to ignore in your text analysis. By eliminating some commonly used words, such as the various forms of the verb "to be," you can filter out noise from your analysis. SAS Visual Analytics supports English language text analysis and includes a stop list for English.

To load a stop list:

1   Select **Tools ▸ Load Text Analytics Stop List**.

2   Make changes to the location and library if necessary. Each stop list must be stored in its own metadata folder.

     A table that is named ENGSTOPL is registered in the specified location and library.

3   Click **OK**.

SAS Visual Analytics supports one stop list for each SAS LASR Analytic Server. You load the stop list (which is a table) to memory by performing the previous steps. If more than one library is registered for SAS LASR Analytic Server, you can use any one of them. If you load a stop list more than once or use more than one library, the server uses the last stop list that was loaded to memory.

Adding a Unique Numeric Key to a Table

In the case where your data already includes unstructured text that you want to analyze, the data builder can add a unique numeric key as a column in the table. For example, many data tables include a text column for comments. Using text analysis, you can look for trends or themes within these comments. The text analysis in SAS Visual Analytics requires a unique numeric key to reference the data sources.

To add a unique numeric key to a table:

1   Create a new data query.

2   Add the columns to include in the text analysis by clicking the column name in the table in the workspace.

3   In the **Column Editor**, right-click on the first row in the table, and select **Add Generated Key Column**.

     A column that is named `GENERATED_ID` is added to the table.

4   Specify an output table, library, and location.

5   Save and run the data query.

When you run the data query, the `GENERATED_ID` column begins at zero and increments for each row in the source table. If you select the **Append data** check box on the **Properties** tab, the data query increments from the current maximum value.

If the data that you want to use is already in memory on SAS LASR Analytic Server, then performing the previous steps results in the following sequence of data movements:

■   The data is read from memory on SAS LASR Analytic Server and transferred to the SAS Application Server.

■   The data query runs and adds the generated key column.

■   The data is transferred back to SAS LASR Analytic Server in a new table.

You can avoid these data movements by adding the generated key column before loading the table to memory on SAS LASR Analytic Server.

17

Customizing Code

Using the Code Tab

You can click the **Code** tab to view the SQL statements that are generated by the data preparation expressions on the **Design** tab.

The **Code** tab can also be used to enter custom code.

■ The Preprocess and Postprocess views enable you to enter SAS statements that are run before and after the data query runs.

■ The All Code view shows the generated SQL statements. You can click the ⬚ to unlock the view so that you can manually edit all of the SAS statements for the data query. After you have unlocked the view, the button icon changes to ⬚.

After you have unlocked the view, you cannot continue using the **Design** tab to edit the data query, and the Preprocess and Postprocess views are disabled. However, if you

have not saved your changes to the code, you can click ↩ to be able to use the **Design** tab.

Preprocess and Postprocess Code

In some cases, you might want to assign SAS options, load format catalogs, use LIBNAME statements, or run macros before running a data query. You can unlock the code to enter the statements, but that prevents you from using the **Design** tab. However, the Preprocess view or Postprocess view might enable you to supplement the data query while using the code generation features of the **Design** tab.

Any SAS statements that you enter in the Preprocess view or Postprocess view are included in the All Code view.

Considerations for Manually Editing Code

Before You Manually Edit Code

The code generation features in SAS Visual Data Builder can generate a lot of SAS statements with a few simple clicks. Depending on your needs, you might benefit by using the **Design** tab for adding input tables, performing joins, and filtering data before you customize the generated code.

The following features and settings must be specified before you unlock the code:

Unique values
    Selecting this check box determines whether the DISTINCT keyword is included in the generated PROC SQL statement. After you unlock the code, the check box cannot be used to include the keyword.

Append data
    Selecting this check box enables you to use the generated SAS statements for appending data to the output table before you unlock the code.

Output table

You must specify the table name, location, and library, and you must save the data query before you unlock the code. When you save the data query, the metadata for the table is registered.

Use a staging table

Selecting this check box registers the metadata for the table when you save the data query. Select the check box, specify the library, and save before you unlock the code.

Input and Output Libraries

When you unlock the code for manual editing, you must specify the libraries to use for your tables. You can do this with LIBNAME statements. If you dragged and dropped the tables onto the workspace before you unlocked the code for manual editing, the LIBNAME statements for the tables are automatically included.

See Also

■ *SAS Statements: Reference*

■ *SAS/ACCESS for Relational Databases: Reference*

Output and Staging Tables

Specify the table names, locations, and libraries on the **Properties** tab just as you do on the **Design** tab.

Note: Specify the tables and save the data query before you unlock the code. If you do not specify the table name or the table name is already in use, you are not able to save the data query.

See Also

SAS Language Interfaces to Metadata

Validating Custom Code

If you entered SAS statements in the Preprocess or Postprocess views, those statements are not validated when you click .

If you unlock the code in the All Code view, the button is disabled. You cannot validate customized code.

18

Scheduling

About Scheduling Data Queries

How Does the Scheduling Feature Work?

When you have a data query open in the workspace, you can click ⏰ to schedule the data query. When you schedule a data query, the data builder performs the following operations:

1 Creates a *job* that performs the data query operations.

2   Creates a *deployed job* from the job.

3   Places the job into a new *deployed flow*.

4   Schedules the flow on a *scheduling server*.

You can reschedule the data query based on specified conditions (for example, run immediately or run whenever a trigger condition is met).

The job, deployed job, and deployed flow are metadata objects. The data builder stores them in the same metadata folder with the data query. The metadata objects are named based on the following pattern:

```
vdb_query_timestamp
```

When the specified conditions are met, the data query is run with the user ID that scheduled it. This is the behavior for the Operating System Services Scheduler.

Editing Data Queries That Are Already Scheduled

If you edit a data query that is already scheduled, you must click 🕒 again so that the SAS statements for the data query are regenerated and saved.

Caution about Scheduling Data Queries to Run Now

When you schedule a data query, one of the options is to run it immediately. Select **Run now** in the Schedule window.

Performing the following steps results in an error condition:

1   Use a SAS data set for the output table of the data query.

2   Run the data query.

3   Click the **Results** tab to look at the output.

4   Schedule the data query by selecting **Run now**.

These steps result in an error condition because SAS unlocks a SAS data set when it is opened for reading. When step 3 is performed, the output table is locked, and no other process can overwrite the output table. The following message is included in the SAS log:

Locked Error Message

```
ERROR: A lock is not available for OUTPUTTABLE.

ERROR: Lock held by process xxxx.
```

You can avoid this error condition. If you want to schedule the data query to run now, close it, open it again, and then schedule it to run now. Alternatively, you can schedule the data query to run in the future, and then close the data query.

Scheduling Preferences

Default Scheduling Server

By default, your deployment includes a server that is named `Operating System Services – ` *hostname.example.com*. This server is used as the default scheduling server.

Use the Server Manager plug-in to SAS Management Console to identify the scheduling servers that are included in your deployment. You can specify a different scheduling server in your application preferences. Any data queries that you schedule after you specify a different scheduling server will use the new scheduling server.

Some deployments include the Platform Suite for SAS server. To use this server, change the scheduling server. The default name is `Platform Process Manager`.

In all cases, when you schedule a new data query, the data builder retrieves your default scheduling server, and uses that value to look up the scheduling server in SAS metadata. The data builder uses the first server that matches the value in SAS metadata. Including the host name, such as `Operating System Services – ` *hostname.example.com* ensures that the data builder uses the server that you specify.

Default Batch Server

By default, your deployment includes a server that is named `SASApp – SAS DATA Step Batch Server`. This server is used as the default batch server.

You can specify a different batch server in your application preferences. Consider the following before you change the default batch server:

- The batch server must be registered in metadata as a component of a SAS Application Server that you can access.

- You must specify the same SAS Application Server as your default application server in your preferences.

As with the default scheduling server, the data builder retrieves your default batch server, and uses that value to look up the batch server in SAS metadata the first time you schedule the data query. The data builder uses the first server that matches the value in SAS metadata.

Default Deployment Directory

A deployment directory is a SAS metadata object that represents the following items:

- the name of the SAS Application Server with which the deployment directory is associated (the default value is `SASApp`)

- a name for the deployment directory (the default value is `Batch Jobs`)

- the path to the deployment directory (the default value is `SAS-config-dir/Lev1/SASApp/SASEnvironment/SASCode/Jobs`)

When you schedule a data query, the SAS statements for the data query are saved in a file. The file is saved in the path that is associated with the deployment directory. The file is named based on the same pattern that is described in "How Does the Scheduling Feature Work?" on page 125.

The data builder looks up the SAS Application Server in the SAS Metadata Server using your scheduling server preference setting. The initial value is `SASApp`. If a matching server name is not found, then the data builder uses the first application server that is returned. After the server is determined, the data builder looks up the deployment directory in that server context that matches your scheduling server preference setting.

If a matching deployment directory is not found, then the data builder uses the first deployment directory that is returned.

You can specify a different name for the default deployment directory. For more information about deployment directories and using the Server Manager plug-in to SAS Management Console, see *Scheduling in SAS*.

When Are the Scheduling Preferences Used?

Any of the preferences that you change are used the next time you create a data query and schedule it. If you edit an existing data query that is already scheduled, the existing settings for the scheduling server, batch server, and deployment directory are not updated with the changes. To change the settings for existing data queries that are already scheduled, use SAS Management Console to redeploy the deployed job for the data query.

Creating Events

Why Use Events?

Events specify conditions that must be met before a step in the flow can take place. You can use SAS Visual Data Builder to create two types of events.

■ *Time events* are evaluated based on a specified time being reached.

   Note: You can specify multiple time events. However, the Operating System Services Scheduler (the default scheduler) can use only one of the time events.

■ *File events* are evaluated based on the state of a specified file.

You can create file events if your deployment includes a scheduling server that supports them, and the flow is deployed to that scheduling server. Time and file events can be used as triggers (conditions that must be met in order to run a flow on the scheduling server).

Creating a Time Event

You can create a time event and use it as a trigger.

To create a time event:

1   In the Schedule window, select **Select one or more triggers for this data query**, and then click **New Time Event**.

2   In the New Time Event window, specify whether the time event should happen one time only or more than once. If the time event should happen one time only, specify the date and time for the time event.

3   If the time event should happen more than once, select **More than once**, and then select a radio button for how often the time event should repeat (such as hourly, weekly, or yearly).

4   Specify the details of when the time event should repeat. The specific fields that are available depend on the recurrence interval that you select.

   If you select **Hourly**, then the time is calculated from hour zero on a 24-hour clock. For example, if you leave **Interval hour** set to 1, then the data query runs at hour zero (midnight), and runs each hour. If you set **Interval hour** to 2, then the data query runs every other hour, and the **Offset hour** specifies whether to run on the even hours (when set to 0) or on the odd hours (when set to 1). In this case, the **Offset hour** specifies the offset from hour zero (midnight).

   If the recurrence interval requires you to select start times, use the **Hours** and **Minutes** check boxes to select the times. The **Minutes** area contains groupings of 10-minute intervals. Selecting a check box for a minute grouping selects all of the minutes in that grouping.

To select individual minutes, expand the grouping.

If you select multiple values for **Hours**, all of the selections for **Minutes** apply to all of the selected values. For example, if you select 19:00 (07:00 PM) and 22:00 (10:00 PM) for **Hours**, and 43 for **Minutes**, the time event is scheduled for 19:43 and 22:43.

5   If needed, specify the start date and end date for the time event. The default is to start at the current date and time and not to have an end date.

Creating a File Event

You can create a file event and use it as a trigger. The file event can check for various file conditions, such as its existence, size, or age.

Note: Not all scheduling servers support file events. Platform Suite for SAS and the SAS Distributed In-Process Services Scheduling Server are two servers that do support file events. The **New File Event** button that is used to create a file event is available only when the scheduling server supports file events.

To create a file event:

1 In the Schedule window, select **Select one or more triggers for this data query**, and then click **New File Event**.

2 In the New File Event window, specify or select the file to use for the file event.

3 Select the condition to evaluate the file to make the file event true. For example, if you selected **not exist** for the condition, the file event would be true only if the selected file was not in the specified location.

4 If needed, specify the details (such as size or age) about the condition.

Exporting Data Queries as Jobs

For deployments that include SAS Data Integration Studio and prefer to deploy jobs, create flows, and schedule flows manually, you can export a data query as a job, and then perform the deployment steps.

This feature enables you to work with a data query interactively while you create it, and then you can export the data query as a job for automating the operations. After the job is exported, you can deploy the job for scheduling with SAS Data Integration Studio. For more information about deploying jobs, see *Scheduling in SAS*.

After a job is stored in metadata, you can open the job with SAS Data Integration Studio and edit it. This might be necessary if you need to add or remove columns or change the column expression in a calculated column.

If the deployment does not include SAS Data Integration Studio, then you can modify the data query and export it as a job again. However, you cannot overwrite job objects in metadata with the data builder. If you need to delete objects from metadata (such as jobs, libraries, or tables), use SAS Management Console.

To export a data query as a job:

1   Use the SAS Folders tree to locate the data query.

2   Select the data query, right-click, and select **Export as a job**.

3   In the Export as a Job window, enter a name and specify a location. Click **Export**.

Additional Scheduling Resources

SAS Visual Data Builder provides an easy-to-use method for taking a data query and making it available for scheduling as a flow. In addition, the Schedule Manager plug-in to SAS Management Console provides additional resources for managing deployed flows, job dependencies, and scheduling servers.

For users that are familiar with the Schedule Manager plug-in, be aware that the data builder does not support scheduling a data query without a trigger. In the Schedule Manager plug-in, this option is specified as the **Manually to the scheduling server** option. If this is your preferred method for scheduling, you can use the Schedule Manager plug-in to modify the flow for the scheduled data query. If you use the data builder to schedule the data query afterward, the setting is lost, and you will need to repeat your steps with the Schedule Manager plug-in.

See Also
Scheduling in SAS

19

Using the Results Tab

About the Results Tab

After you have created your data query and clicked ▶, it runs, and you are prompted to determine whether you want to view the results.

If you click **Yes**, the data builder retrieves the data and shows you the results on the **Results** tab.

Note: For large data sets, retrieving and displaying the data can require a long time.

Data Pages

When you view data, a SAS server retrieves the data from the data set. Instead of returning all of the data, the data viewer requests a page of data. You can set the page size to between 20 and 2000 rows.

A slider is provided at the bottom of the **Results** tab so that you can navigate through the pages and browse the entire data set.

Navigate within the Data

Navigate to a row in the table by entering the row number in the **Go to row** field. The **Go to first row** and **Go to last row** buttons provide navigation to the first and last rows of the table.

Find

Search for text or numbers in the data set by clicking ⬛, typing in the **Find** field, and pressing Enter. The find feature searches through the rows in the data set for the value, and it highlights the first row with the value. You can navigate through the results by clicking **Find next** and **Find previous**.

Click ⬛ to set options. The following list identifies the options:

■ **Options** tab

  □ **Exactly matches the specified string**

  □ **Contains the specified string** (default setting)

  □ **Starts with the specified string**

- ☐ **Case sensitive**

- ☐ **Trim leading and trailing spaces** (default setting)

■ **Columns** tab

Enable and clear check boxes to specify the columns to search.

Filter and Sort

Sort the data by clicking ⬇️, and then choosing the columns and sort order. The **Generated SQL statement** field shows the sort criteria.

Filter the data by clicking 🔽, and setting the following options:

■ The **Sample** tab is used to limit the number of rows to return or to select distinct values.

■ The **Row Filter** table is used to subset data with a WHERE clause. You can filter on the values in multiple columns, set ranges for numeric values, and use IN and NOT IN criteria for character values. If the number of distinct values is less than 50, you can select check boxes for the values. If there are more than 50 distinct values, you must enter the values to use for filtering.

■ The **Column Filter** tab is used to select the columns to display.

Export Data

You can save the data that you view as a comma-separated values file. After you click ↗️, you can choose to export the rows on the current data page, all of the rows, or a range of rows. The export feature has a limit of exporting 200,000 rows.

If you choose to export all of the data or a range of rows, you have to click **Retrieve Data** before you can click **Export as CSV**.

Column Headings

Use the **Headings** menu to control the appearance of the column headings. You can view the column name from the data set, the column label from the data set, or a combination of the two. If a column does not have a label in the data set, then the column name is used.

Part 3

Exploring Data

20

Overview of SAS Visual Analytics Explorer

What Is SAS Visual Analytics Explorer?

SAS Visual Analytics Explorer (the explorer) is a component of SAS Visual Analytics that enables you to explore your data sources. You can explore your data by using interactive visualizations such as charts, histograms, and tables. You can also apply data analysis such as forecasting, correlation, and fit lines.

Your work in the explorer is saved as a metadata object called an *exploration*. An exploration (sometimes called a visual exploration) contains all of the visualizations, data settings, and filters from your explorer session.

You can use the explorer to explore in-memory tables from the SAS LASR Analytic Server. You can either open a table directly or you can open a saved exploration.

You can export your exploration results as reports, which can be refined in SAS Visual Analytics Designer (the designer) or viewed directly. Reports can be viewed on a mobile device or in the SAS Visual Analytics Viewer (the viewer).

You can save your explorations as PDF documents, share your saved explorations via e-mail, and export your visualizations as image files. You can export the data from your visualizations to a data file that you can use in other tools.

Note: Stored processes are not supported in the explorer.

The Welcome Window

If you enter the explorer without a reference to a specific exploration, then the Welcome window appears.

The Welcome window enables you to perform the following tasks:

■ Create a new exploration by clicking **Select a Data Source**. The Open Data Source window appears.

In the Open Data Source window, select the data source that you want to open, and then click **Open**.

In the **Search** field, you can enter a short string to find all of the data sources whose names or descriptions begin with that string.

Note: You can also perform a search that matches a string anywhere in the data source name. Select the **Search** field, and then press Ctrl + Down to switch the search to **Search (contains)**.

■ Create a new exploration from a local data file by clicking **Import Data**. This task is available only if you have the Import and Load Data capability.

For more information, see "Import a Local Data Source" on page 156.

■ Open an existing exploration. Select from your recent explorations, or select any exploration by clicking **Browse** to locate the exploration.

Your First Look at the Explorer

Here are the main parts of the explorer user interface:

Figure 20.1 *The Explorer User Interface*
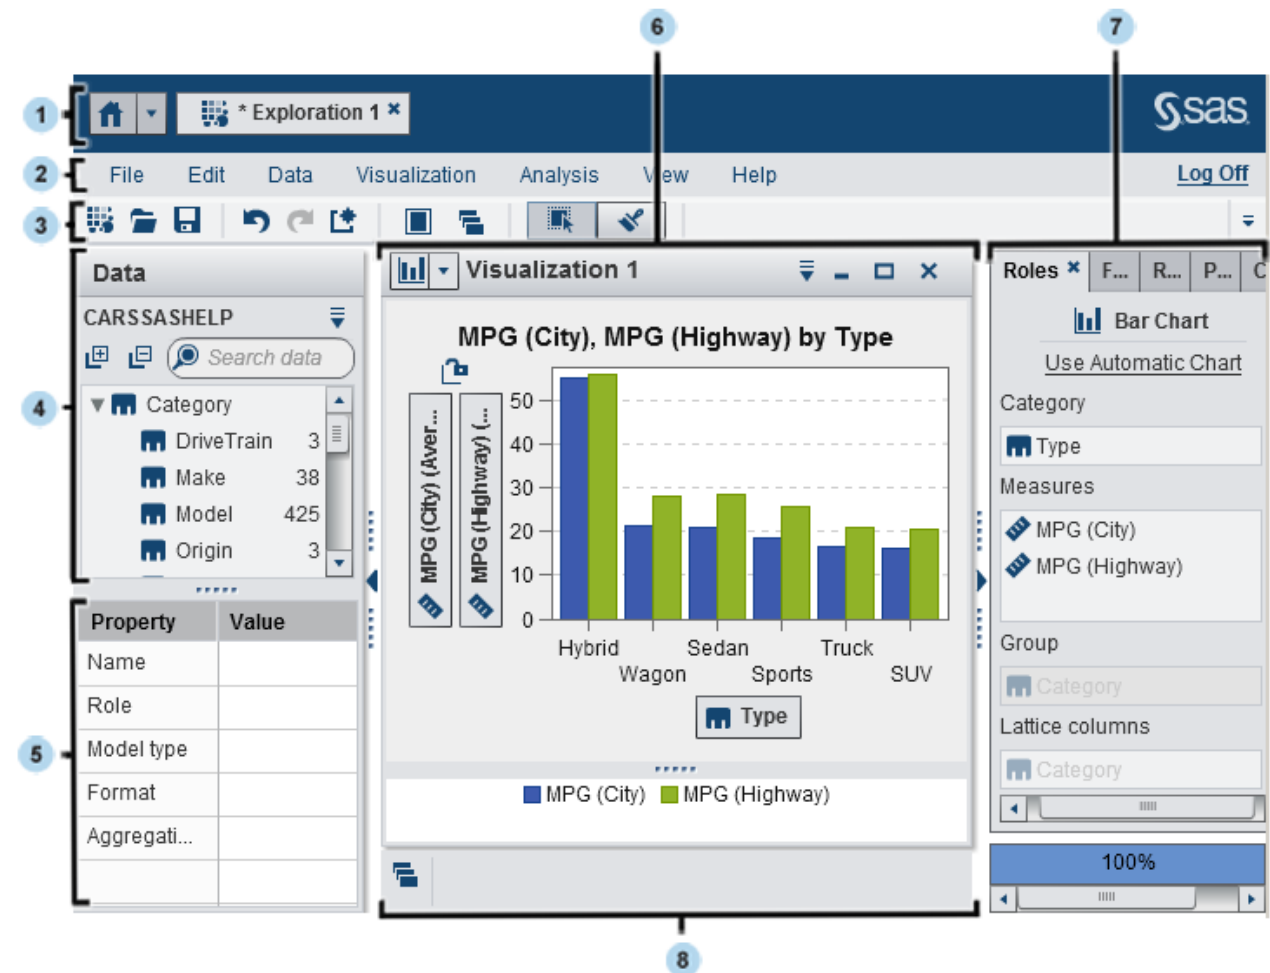
1 The application bar enables you to return to the home page and to access other parts of SAS Visual Analytics and other SAS applications that integrate with the home page. You can access your recently created or viewed reports, explorations,

stored processes, data queries, or other objects in your recent history. Buttons are displayed for each open application.

2 The menu bar offers common tasks, such as creating a new exploration.

3 The toolbar enables you to manage your explorations and visualizations.

4 The **Data** pane enables you to manage the data that is used in your visualizations.

5 The data properties table enables you to set data item properties.

6 The workspace displays one or more visualizations.

7 The right pane's tabs enable you to set properties and data roles, subset data, and use comments.

8 The dock contains any minimized visualizations.

21

Specifying Preferences for SAS Visual Analytics Explorer

Specifying Global Preferences

See "Specify Global Preferences" on page 23 for details about global preferences for SAS Visual Analytics.

Specifying Your Preferences for the Explorer

Here are the steps to specify preferences that are specific to SAS Visual Analytics Explorer (the explorer):

1   Select **File** ▸ **Preferences** to open the Preferences window.

2   Select **SAS Visual Analytics Explorer** ▸ **General**.

3   Select a **Graph skin** to change the visual effects that are applied to your graphs. For example, many of the **Graph skin** settings apply lighting effects to the features of your graph to create a 3-D appearance.

4   Select **Visualization data threshold** to specify the amount of data values that your visualizations can process. The precise amount of data that is supported varies depending on the visualization type. See Appendix 5, "Data Limits," on page 603.

   Note: A large value for **Visualization data threshold** might degrade your application performance or lead to time-out errors.

5   Select the default **Map provider mode** for geographic maps. The map provider creates the background map for geo maps and for network diagrams that display a map.

   Note: You can change the map provider for an individual visualization by editing the **Map service** value on the **Properties** tab.

6   If your **Map provider mode** is **Esri**, then you must select the specific **Esri map service** that you want to use.

7   To return to the default settings, select **Reset to Defaults**.

8   Click **OK** to apply your changes. Your preferences persist between sessions.

22

Managing Explorations

What Is an Exploration?

An exploration (sometimes called a visual exploration) is a metadata object that contains all of the visualizations and data settings from a SAS Visual Analytics Explorer (the explorer) session. You can use explorations to save your session for later and to share it with others.

Create a New Exploration

To create a new exploration, select **File ▶ New Exploration** from the menu bar. The Open Data Source window appears. Select your data source, and then click **Open**.

Saving Your Exploration

To save your current exploration, select **File** ▶ **Save As**, and then select a location and a name.

Deleting Explorations

To delete an exploration, use the SAS Visual Analytics home page. See "Managing Content on the Home Page" on page 13.

23

Managing Data

Managing Data Properties

Overview of Managing Data Properties

At any time after you have loaded a data source for your exploration, you can modify the properties of the data items. The changes that you make are saved as part of the exploration, but do not affect the original data source.

You can manage the properties for individual data items by using the **Data** pane or manage the properties for all of your data items by using the Data Properties window.

Manage Data Properties for Data Items

You can manage data properties by using the Data Properties window or by using the properties table in the **Data** pane.

To access the Data Properties window, select **Data ▶ Data Properties** from the menu bar.

For each data item, you can specify the following properties:

Name
specifies the name of the data item.

Role
specifies whether the data item is a category, a measure, a document collection, or a geography.

Note: Changing the role from a category to a measure or vice versa does not change the data type of the data. You can change the data type by using the **Format** and **Parse** operators in a calculated data item.

Model Type
specifies the type of data model for a measure or for a category with datetime data. Select whether the data model is continuous or discrete.

The model type that you specify determines the type of data filters that you can create for the data item.

Note: If you change the model type for a data item that is used in a filter, then the filter is removed.

Format

specifies the data format for the data item.

See "Select a Numeric Format" on page 152 and "Select a Date or Time Format" on page 153 for information about specifying formats.

Aggregation (for measures only)

specifies the method that is used to aggregate values for the measure. Select one of the following aggregation types:

Sum

provides the sum of the values in the measure.

Average

provides the average of the values in the measure.

Maximum

provides the maximum value in the measure.

Minimum

provides the minimum value in the measure.

Count

provides the number of nonmissing values in the measure.

Select a Numeric Format

Here are the steps to select a numeric format:

1   In the Data Properties window or in the property table in the **Data** pane, click the format for the data item. A window appears.

2   Select the basic format type from the **Format** list.

3   Specify the **Width** parameter for the format. The width of the format specifies the maximum number of characters that a value can contain, including decimal points.

A preview of your format appears in the **Sample** field.

4   Specify the number of **Decimals** for the format. The **Decimals** value specifies the number of digits that appear to the right of the decimal place.

A preview of your format appears in the **Sample** field.

5   Click **OK** to apply your format.

Note: You can restore the default format by clicking **Reset to Default**.

Select a Date or Time Format

Here are the steps to select a date or time format:

1   In the Data Properties window or in the property table in the **Data** pane, click the format for the data item. A window appears.

2   Select the basic format type from the **Format** list.

3   If necessary, select a more specific format from the **Sample** list. The format names also represent sample values for the format.

4   Click **OK** to apply your format.

Note: You can restore the default format by clicking **Reset to Default**.

Note: Time and datetime values in the explorer ignore daylight-saving time.

Show and Hide Data Items

If your exploration contains a large number of data items, you might want to hide the items that you are not interested in.

Here are the steps to select which data items that are shown in the **Data** pane:

1   Select **Data ▶ Show or Hide Items** from the menu bar. The Show or Hide Items
    window appears.

2   Click **Select All** to select or deselect all of the data items, or click the check box
    beside each data item to select whether that data item is shown.

3   Click **OK** to apply your changes.

Sort Data Items

To sort your data items in the **Data** pane, select ▼ and then select **Sort Items**. Select
Ascending by Name or **Descending by Name**.

Group Data Items

To group your data items in the **Data** pane, select ▼ , and then select **Group Items**.
Select one of the following grouping types:

By First Letter
    groups by the first letter of the data item name.

By Data Type
    groups by data type (date, number, or text).

By Model Type
    groups by model type (continuous or discrete).

By Role
    groups by data role (category, geography, hierarchy, measure, or aggregated
    measure).

By Format
    groups by data format (for example, $CHAR, Numeric, Percent, or MMMYYY).

By Aggregation
    groups by the default aggregation (Sum, Average, Count, Minimum, or Maximum).

Note: The default grouping is **By Role**.

Assigning Colors for Category Values

About Colors

By default, the explorer assigns colors dynamically to category values for grouped visualizations. You can assign specific colors to your category values to ensure that the category values are always represented by a specific color.

Assign Colors for a Category

Here are the steps to set the colors for a category:

1   Right-click on a category from the **Data** pane, and then select **Colors**. A color selection window appears.

2   For each category value to which you want to assign a color, select the value from the **Choose** drop-down list, and then click the color box to select a color. Click **Apply** to save your selection.

3   (Optional) To automatically assign persistent colors for all of the unassigned category values, click **Assign remaining**. Colors are assigned to all of the remaining category values. These colors are used consistently in all of your visualizations.

4   Click **OK** to apply your changes.

Reset Colors for a Category

Here are the steps to reset the colors for a category:

1   Right-click on a category from the **Data** pane, and then select **Colors**. A color selection window appears.

2   Click **Reset all** to clear all of the color assignments for the category.

3   Click **OK** to apply your changes.

Assign a Color for a Specific Category Value

Here are the steps to assign the color for a specific category value:

1   Within a grouped visualization, right-click on a data value, and then select **Change Color**.

2   Click the color box to select the color.

3   Click **OK** to apply the new color. This color value is used consistently for the selected category value in all of your visualizations.

Import a Local Data Source

If you have the Import and Load Data capability, then you can import a data source into the explorer. You can import a SAS data set, a Microsoft Excel spreadsheet, or a delimited text file such as a CSV file.

You can import a data source by clicking **Import Data** in the Welcome window or in the Change Data Source window. You can also select **File ▶ Import Data** to import a data source as a new exploration.

After you select the data source to import, the **Import Local File** window appears. Depending on the file type, you can specify additional options for the import.

Note: If you encounter a server error message, contact your system administrator to determine whether you require SSH access to the LASR environment.

Open a Different Data Source as a New Exploration

To explore a different data source, create a new exploration by selecting **File ▶ New Exploration**, and then select or import the data source that you want to explore.

Change the Data Source for Your Exploration

At any time, you can change the data source for your exploration. All of your visualizations, filters, and settings are applied to the new data source.

Here are the steps to change the data source for your exploration:

1 Select **Data** ▶ **Change Data Source**.

2 From the Change Data Source window, select or import the data source that you want to use.

   Note: If any data items in your current data source do not exist in the new data source, then a message appears. Any data items that do not exist in the new data source are removed from your exploration. Any filters, ranks, or calculated items that are based on the removed items are also removed.

3 Click **Open** to open the new data source.

Refresh Your Data Source

At any time, you can refresh your data source to include the latest changes.

To refresh your data, select **Data** ▶ **Refresh Data**.

Your visualizations are automatically updated with the refreshed data.

Creating Calculated Data Items

About Calculated Data Items

The explorer enables you to calculate new data items from your existing data items by using an expression.

Note: All calculations are performed on unaggregated data. The calculation expression is evaluated for each row in the data source before aggregations are performed. To perform calculations on aggregated data, see "Creating Aggregated Measures" on page 160.

In addition to performing mathematical calculations on numeric values, you can use calculated data items to create character values and date and time values. For example, if your data contains separate categories for month, day, and year, then you can calculate a date value from each category.

Create a Calculated Data Item

Here are the steps to create a calculated data item:

1   Select **Data** ▶ **New Calculated Item**. The New Calculated Item window appears.

2   Enter a **Name** for the calculated data item.

3   Select the data type for the calculated data item from the **Result type** drop-down list.

   Note: The data type is updated automatically based on the outermost operator of your expression.

4   Build the expression for your calculated data item by dragging and dropping data items and operators onto the expression in the right pane. For each field in the expression, you can insert a data item, an operator, or a specific value.

Note: Aggregated data items and derived data items are not supported for calculation expressions.

When you drag and drop data items and operators onto the expression, the precise location of the cursor determines where and how the new element is added to the expression. As you drag the new element over the expression, a preview appears, which displays how the expression would change if you drop the element at the current location.

For example, if your current expression is `( Profit / Revenue )`, and you drag the **x - y (subtract)** operator over the open parenthesis symbol, then the expression changes to `( [number] – ( Profit / Revenue ))`. If you drag the operator over the division symbol, then the expression changes to `( Profit – Revenue )`.

There are a large number of operator types available to perform mathematical functions, process datetime values, and evaluate logical processing such as IF clauses. See Appendix 3, "Operators for Data Expressions," on page 569.

5 When you are finished creating your expression, select the **Default aggregation** for the calculated data item, and then click **Select** to choose the data format.

6 Click **Preview** to see a preview of the calculated data item as a table. The table displays the values of the calculated item and any data items that are part of the calculation expression.

7 Click **OK** to create the new calculated data item. The new data item appears in the **Data Items** pane.

Creating Aggregated Measures

About Aggregated Measures

Aggregated measures enable you to calculate new data items by using aggregated values. For example, you might want to calculate a company's profits by subtracting expenses from revenues.

Note: To calculate data items by using unaggregated values, see "Creating Calculated Data Items" on page 158.

Aggregations are evaluated as part of the calculation expression. For each data item in your expression, you can select the aggregation type and the context for the aggregation.

Create an Aggregated Measure

Here are the steps to create an aggregated measure:

1   Select **Data ▶ New Aggregated Measure**. The New Aggregated Measure window appears.

2   Enter a **Name** for the aggregated measure.

3   Build the expression for your aggregated measure by dragging and dropping data items and operators onto the expression in the right pane. For each field in the expression, you can insert a data item, an operator, or a specific value.

   Note: You can create a new calculated item to use in the aggregated measure expression. Click **New Calculated Item** to create the new calculate item.

   When you drag and drop items and operators onto the expression, the precise location of the cursor determines where and how the new element is added to the expression. As you drag the new element over the expression, a preview appears,

which displays how the expression would change if you drop the element at the current location.

There are a large number of operator types available to perform mathematical functions and evaluate logical processing such as IF clauses. See Appendix 3, "Operators for Data Expressions," on page 569.

4   For each data item in your expression, select an aggregation type. By default, **Sum** is used for measures and **Distinct** is used for categories. To select a new aggregation type, drag and drop an aggregated operator from the **Operators** list onto the aggregation type in the expression. See Appendix 3, "Operators for Data Expressions," on page 569 for a list of the aggregated operators that are available.

5   For each aggregated operator in your expression, select the aggregation context. A drop-down list beside each aggregation enables you to select one of the following context values:

ByGroup
> calculates the aggregation for each subset of the data item that is used in a visualization. For example, in a bar chart, an aggregated measure with the **ByGroup** context calculates a separate aggregated value for each bar in the chart.

ForAll
> calculates the aggregation for the entire data item (after filtering). For example, in a bar chart, an aggregated measure with the **ForAll** context uses the same aggregated value (calculated for the entire data item) for each bar in the chart.

By using the **ForAll** and **ByGroup** contexts together, you can create measures that compare the local value to the global value. For example, you might calculate the difference from mean by using an expression like the following:

```
Avg ByGroup(X) - Avg ForAll(X)
```

6   When you are finished creating your expression, click **Select** to choose the data format.

7   Click **OK** to create the new aggregated measure. The new data item appears in the **Data Items** pane.

Support for Aggregated Measures

Aggregated measures can be used in the following visualization types:

- Automatic Chart

- Crosstab

- Bar Chart

- Bubble Plot (grouped bubble plots only)

- Line Chart

- Treemap

- Geo Map

Creating Derived Items

About Derived Items

For the categories and measures in your exploration, you can create derived data items. Derived data items are aggregated measures that perform calculations for your data.

Derived Item Types for Categories

For categories, you can derive distinct counts. A distinct count displays the number of distinct values for the category that it is based on.

For example, you might derive the distinct count of a category that contains city names. If you create a bar chart by using the derived item and a category that contains product lines, then the bar chart shows the number of cities where each product line is produced.

Derived Item Types for Measures

For measures, you can create the following types of derived items:

Difference from Previous Parallel Period
Displays the difference between the value for the current time period and the value for the previous parallel time period within a larger time interval. For example, you might derive the difference between sales for the current month and sales for the same month of the previous year.

Note: This derived item is not available if your data source does not contain a date data item that includes the year.

Difference from Previous Period
Displays the difference between the value for the current time period and the value for the previous time period. For example, you might derive the difference between sales for the current month and sales for the previous month.

Note: This derived item is not available if your data source does not contain a date data item that includes the year.

Percent Difference from Previous Parallel Period
Displays the percentage difference between the value for the current time period and the value for the previous parallel time period within a larger time interval. For example, you might derive the percentage difference between sales for the current month and sales for the same month of the previous year.

Note: This derived item is not available if your data source does not contain a date data item that includes the year.

Percent Difference from Previous Period
Displays the percentage difference between the value for the current time period and the value for the previous time period. For example, you might derive the percentage difference between sales for the current month and sales for the previous month.

Note: This derived item is not available if your data source does not contain a date data item that includes the year.

Percent of Subtotals

For crosstabs only, displays the percentage of the subtotal value for the measure on which it is based.

You can select whether to calculate the percentage of the subtotal value for the rows in the crosstab or for the columns in the crosstab.

For example, you might calculate the percentage of the row subtotal value for a measure that contains revenue values. The derived item displays the percentage of the subtotal of revenues for each row in the crosstab.

Note: **Percent of Subtotals** can be created only from within a crosstab visualization.

Percent of Total

Displays the percentage of the total value for the measure on which it is based. For example, you might derive the percentage of the total value for a measure that contains revenue values. If you create a bar chart by using the derived item and a category that contains products, then the bar chart shows the percentage of total revenue for each product.

Note: The percentage of the total value is relative to the subset of data that is selected by your filters and ranks.

Period to Date

Displays the aggregated value for the current time period and all of the previous time periods within a larger time interval. For example, you might derive the year-to-date total for each month.

Note: This derived item is not available if your data source does not contain a date data item that includes the year.

Year over Year Growth

Displays the percentage difference between the current time period and an equivalent time period from the previous year. For example, you might derive the difference in sales between the current month and the same month of the previous year.

Note: This derived item is not available if your data source does not contain a date data item that includes the year.

Year to Date

Displays the aggregated value for the current time period and all of the previous time periods within the year. For example, you might derive the year-to-date total for each month.

The year-to-date calculation subsets the data for each year using today's date (where today is evaluated each time you view the exploration). To use all data for every period, use a **Period to Date** item or edit the expression for the derived item.

Note: This derived item is not available if your data source does not contain a date data item that includes the year.

Year to Date Growth

Displays the percentage difference between the year-to-date value for the current time period and the year-to-date value for the same time period of the previous year. For example, you might derive the difference in year-to-date sales between the current month and the same month of the previous year.

The year-to-date calculation subsets the data for each year using today's date (where today is evaluated each time you view the exploration). To use all data for every period, use a **Period to Date** item or edit the expression for the derived item.

Note: This derived item is not available if your data source does not contain a date data item that includes the year.

Note: The derived item types that aggregate values over time use periodic operators. For information about using periodic operators, see "Periodic Operators" on page 579 .

Create a Derived Item for a Measure

You can create a derived item for a measure by using the **Data** pane or by selecting a measure in a visualization.

To create a derived item from the **Data** pane, right-click the data item that the derived item is based on, and then select **Create** ▸ [*item-type*]. If you select a type that uses a date, select the date data item that is used for date calculations.

To create a derived item in a visualization, right-click a measure heading in the visualization, and then select **Create and Add** ▸ [*item-type*]. If you select a type that

uses a date, select the date data item that is used for date calculations. The derived item is added to the visualization automatically.

For a list of item types, see "Creating Derived Items" on page 162.

The new derived item appears in the **Data** pane.

Create a Derived Item for a Category

To create a derived item for a category from the **Data** pane, right-click the category, and then select **Create Distinct Count**.

The new derived item appears in the **Data** pane.

Support for Derived Items

Derived items can be used in the following visualization types:

- Automatic Chart
- Crosstab
- Bar Chart
- Line Chart
- Bubble Plot (grouped bubble plots only)
- Network Diagram
- Geo Map
- Treemap

Explore Distinct Values in a Visualization

For a bar chart, line chart, treemap, or geo map that contains a distinct count derived item, you can explore the distinct values that are associated with a data point in the visualization. To explore the distinct values, right-click on a data point in the visualization, and then select **Explore Distinct Values**. When you explore the distinct

values, a new bar chart is created. The new bar chart displays the values and the frequency of each value.

For example, you might use a distinct count derived item to create a bar chart that displays the number of cities where each product line is produced.

Then, you might want to explore the distinct values for City where Promotional items are produced. By right-clicking on the bar for Promotional, and then selecting **Explore Distinct Values**, you get a new bar chart that contains the cities where promotional items are produced and the frequency of data for each city.

Edit a Calculated, Aggregated, or Derived Data Item

Open the Editor Window

To edit a duplicate data item, an aggregated measure, a calculated data item, or a derived data item, right–click the item in the **Data** pane, and then select **Edit**.

Note: You cannot edit a Percent of Total or Percent of Subtotals derived item.

Specify the Item Name

To specify the name for the item, enter a name in the **Name** field.

Edit a Parameter Value

To edit a parameter value, select the parameter, and enter a new value. Or, right-click the parameter field, and select **Replace with** to select a data item.

To specify a missing value, enter a period (.)character.

Add an Operator

To add an operator to the expression:

1   From the **Operators** list, select the operator that you want to add. For a list of the operators that are available, see Appendix 3, "Operators for Data Expressions," on page 569.

2   Drag and drop the operator onto the expression.

3   For any required parameters, select the parameter, and enter a value. Or, right-click the parameter field, and select **Replace with** to select a data item.

Replace an Operator

To replace an operator, drag and drop a new operator onto the existing operator in the expression. You can also right-click an operator in your expression, and select **Replace Operator with** ▶ *operator*.

Delete Part of an Expression

To delete a part of an expression, highlight the part of the expression that you want to delete, right-click, and then select **Delete** or **Clear**.

Subset an Expression

To subset an expression and delete the remainder:

1   Highlight the part of the expression that you want to keep.

2   Right-click, and select **Keep Operand**. The parts of the expression that were not selected are removed.

Edit an Expression as Text

On the **Text** tab of the expression editor, you can edit an expression as text code. See "Editing a Data Expression in Text Mode" on page 565.

Display the Current Expression as Text

To display the current expression as text on the **Display Text** tab, click the ▼ drop-down, and then select **Show display text**.

Manage the Scratch Area

The scratch area of the expression editor enables you to store elements of an expression, and then add them back into the expression as needed.

You can see the contents of the scratch area on the **Scratch** tab.

To move items to the scratch area, right-click part of the expression, and select **Move to Scratch Area**. You can also drag and drop items from your expression onto the **Scratch** tab.

To move an item from the scratch area to your expression, drag and drop the item from the **Scratch** tab onto your expression.

To delete an item in the scratch area, click the 🗑 button next to the item.

Delete a Calculated, Aggregated, or Derived Data Item

To delete a duplicate data item, an aggregated measure, a calculated data item, or a derived data item, right-click on the data item in the **Data** pane, and then select **Delete**.

Note: If the deleted data item is used by a visualization, then a confirmation message appears. The deleted item is removed from the visualization.

Note: You cannot delete a data item that is included in a hierarchy. You must remove the data item from any hierarchies before you can delete it.

Duplicate a Data Item

Duplicate data items enable you to use multiple copies of a measure that have different data formats or default aggregations. For example, you might want to use the Minimum and Maximum aggregations for a data item in the same visualization.

Here are the steps to duplicate a data item:

1   Right-click on a measure in the **Data** pane, and then select **Duplicate Data Item**. The New Duplicate Item window appears.

   Note: You can also invoke the New Duplicate Item window by dragging and dropping the same data item onto a visualization more than once.

2   Enter a **Name**, a **Format**, and a **Default aggregation** for the duplicate data item.

3   Click **OK** to create the duplicate data item.

Define a Geography Data Item

About Geography Data Items

A geography data item is a category whose values are mapped to geographical locations or regions. Geography data items can be used with geo maps to visualize your data on a geographic map.

For example, if your data source contains a CountryName column that identifies countries, then you might create a geography data item for CountryName by using the predefined geographic role, **Country or Region Names**. See "Define a Geography Data Item By Using a Predefined Geographic Role" on page 171.

You might also create a customized geography data item to identify geographic information that is specific to your organization (for example, sales regions, warehouse

locations, oil platforms, and so on). For each unique combination of category values, latitude values, and longitude values, a point is plotted on the geographic map. See "Define a Geography Data Item By Using a Custom Geographic Role" on page 172.

Note: For predefined geographic roles, the values of your geography data items must match the lookup values that are used by SAS Visual Analytics. To view the lookup values, see `http://support.sas.com/va63geo`.

Define a Geography Data Item By Using a Predefined Geographic Role

Here are the steps to define a geography data item by using a predefined geographic role:

1 In the **Data** pane or in the Data Properties window, locate a category that contains geographic information.

2 Change the role for the category to **Geography**. The Geography Role window appears.

3 Select the predefined geographic role for your category:

Country or Region Names
specifies countries or regions by name.

Country or Region ISO 2-Letter Codes
specifies countries or regions by using two-letter country codes from the ISO 3166-1 standard.

Country or Region ISO Numeric Codes
specifies countries or regions by using three-digit numeric country codes from the ISO 3166-1 standard.

Country or Region SAS Map ID Values
specifies countries or regions by using the two-letter codes that are used with the MAPSGFK library that is included with SAS/GRAPH.

Note: The two-letter SAS Map ID country values are identical to the ISO 3166–1 two-letter country codes.

Subdivision (State, Province) Names
specifies country subdivisions (for example, states and provinces) by using the subdivision names.

Subdivision (State, Province) SAS Map ID Values
specifies country subdivisions (for example, states and provinces) by using the two-letter codes that are used with the MAPSGFK library that is included with SAS/GRAPH.

US State Names
specifies states and territories in the United States by using the state and territory names.

US State Abbreviations
specifies states and territories in the United States by using two-letter postal codes.

US ZIP Codes
specifies five-digit ZIP codes for the United States.

4   Click **OK** to apply your changes.

Define a Geography Data Item By Using a Custom Geographic Role

Here are the steps to define a geography data item by using a custom geographic role:

1   In the **Data** pane or in the Data Properties window, locate a category that contains geographic information.

2   Change the role for the category to **Geography**. The Geography Role window appears.

3   Select **Custom**, and then specify the following:

Latitude
specifies a measure from your current data source that contains the latitude (Y) coordinate values for the geographic role that you want to define.

Longitude

specifies a measure from your current data source that contains the longitude (X) coordinate values for the geographic role that you want to define.

Coordinate Space

specifies the coordinate space (coordinate system) that is used to project the longitude and latitude coordinate values.

4 Click **OK** to apply your changes.

Define Data Items for Text Analytics

To enable text analytics for the word cloud visualization:

1 Identify a data item as the unique row identifier for the exploration. You can choose a category, a measure, or a calculated item. The unique row identifier must have a unique value for each row of the data source.

Right-click the data item in the **Data** pane, and then select **Set as Unique Row Identifier**.

Note: Depending on the number of rows in your data, it might be possible to select an item whose values are not unique. A nonunique identifier can produce unreliable results.

2 Assign one or more categories to the **Document Collection** data role. Right-click a category, and then select **Document Collection**.

Note: A document collection can be used in your visualizations the same way as a category.

24

Working with Visualizations

Overview of Visualizations

About Visualizations

SAS Visual Analytics Explorer (the explorer) displays data by using visualizations. A visualization is an interactive visual representation of your data.

Visualization Types

You can assign any of the following types to your visualizations:
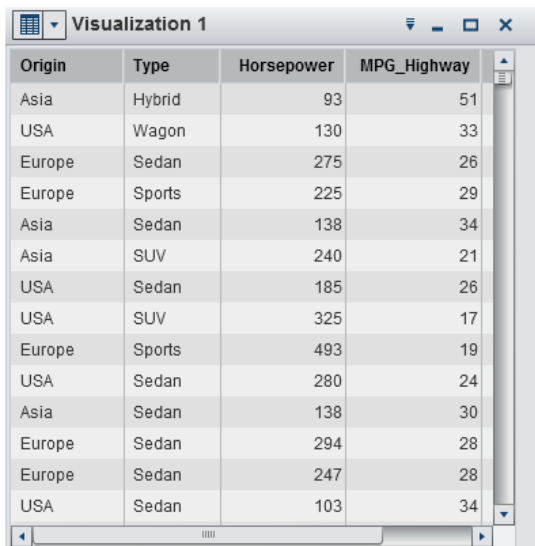
Automatic Chart
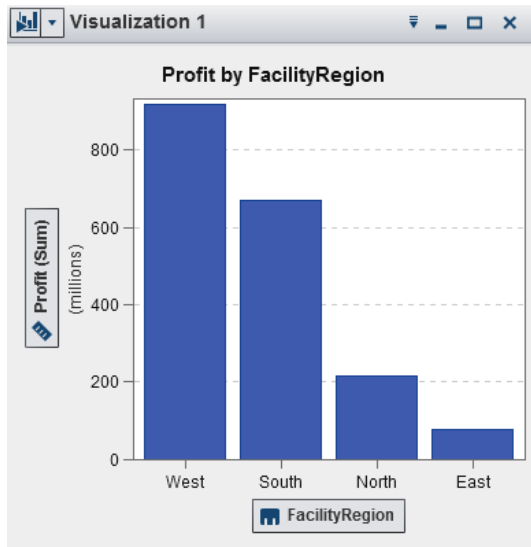
Automatically selects the chart type based on the data that is assigned to the visualization. When you are first exploring a new data set, automatic charts give you a quick view of the data.

For more information, see "Working with Automatic Charts" on page 201.

Table
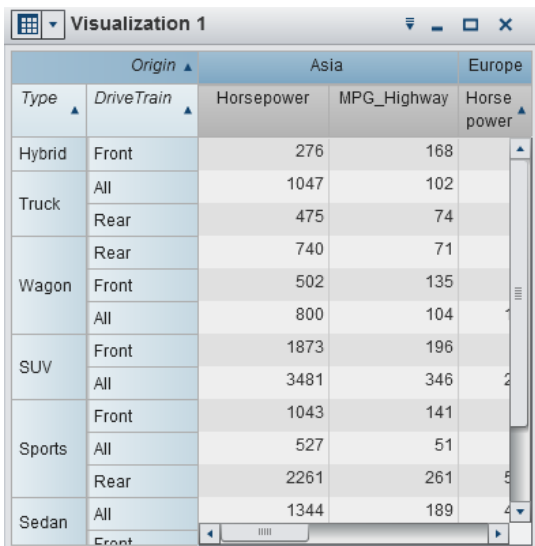
Origin	Type	Horsepower	MPG_Highway
Asia	Hybrid	93	51
USA	Wagon	130	33
Europe	Sedan	275	26
Europe	Sports	225	29
Asia	Sedan	138	34
Asia	SUV	240	21
USA	Sedan	185	26
USA	SUV	325	17
Europe	Sports	493	19
USA	Sedan	280	24
Asia	Sedan	138	30
Europe	Sedan	294	28
Europe	Sedan	247	28
USA	Sedan	103	34

Displays the data as a table. Tables enable you to examine the raw data for each observation in the data source. You can rearrange the data columns and apply sorting.

For more information, see "Working with Tables" on page 209.

Crosstab

		Asia		Europe
Type	DriveTrain	Horsepower	MPG_Highway	Horsepower
Hybrid	Front	276	168	
Truck	All	1047	102	
	Rear	475	74	
Wagon	Rear	740	71	
	Front	502	135	
	All	800	104	
SUV	Front	1873	196	
	All	3481	346	
Sports	Front	1043	141	
	All	527	51	
	Rear	2261	261	
Sedan	All	1344	189	
	Front			

Displays the data as a crosstab. Crosstabs enable you to examine the data for intersections of hierarchy nodes or category values. You can rearrange the rows and columns and apply sorting. Unlike tables, crosstabs display aggregated data.

For more information, see "Working with Crosstabs" on page 210.

Bar Chart
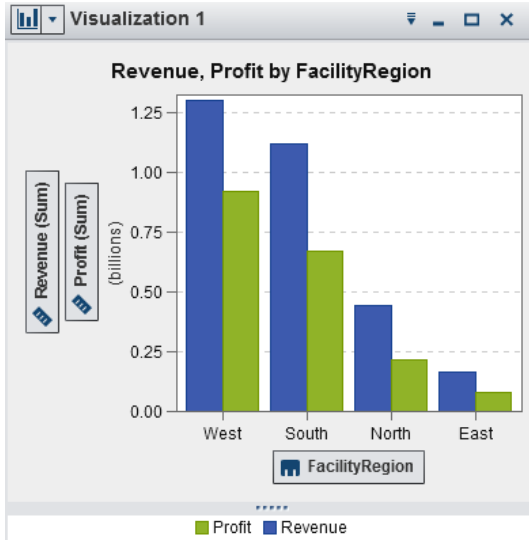
Displays the data as a bar chart. Bar charts are useful for comparing data that is aggregated by the distinct values of a category.

A bar chart consists of vertical bars or horizontal bars. You can apply grouping and create lattices.

For more information, see "Working with Bar Charts" on page 202.

Line Chart
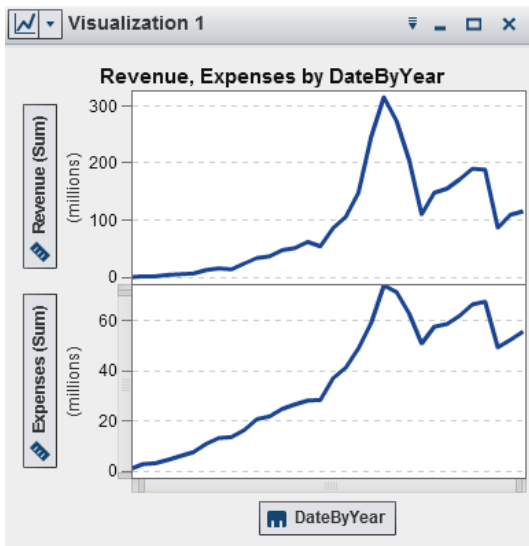
Displays the data as a line chart. A line chart is useful for data trends over time. Line charts support forecasting for predicting future values.

You can apply grouping and create lattices.

For more information, see "Working with Line Charts" on page 205.

Scatter Plot
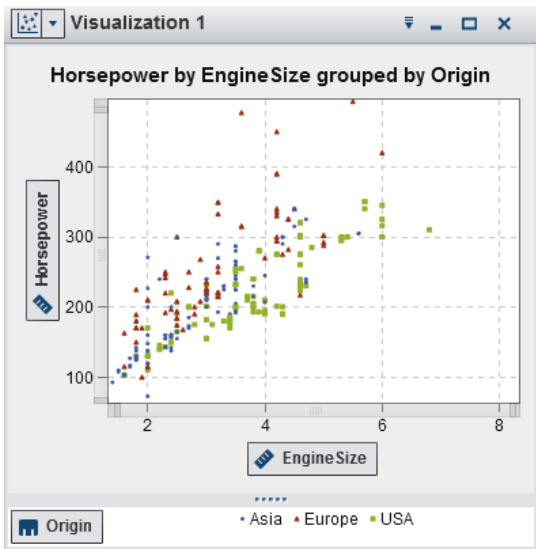
Displays the data as a scatter plot. Scatter plots are useful to examine the relationship between numeric data items.

In a scatter plot, you can apply statistical analysis with correlation and regression. Scatter plots support grouping.

When you apply more than two measures to a scatter plot, a scatter plot matrix compares each pairing of measures.

For more information, see "Working with Scatter Plots" on page 213.

Bubble Plot
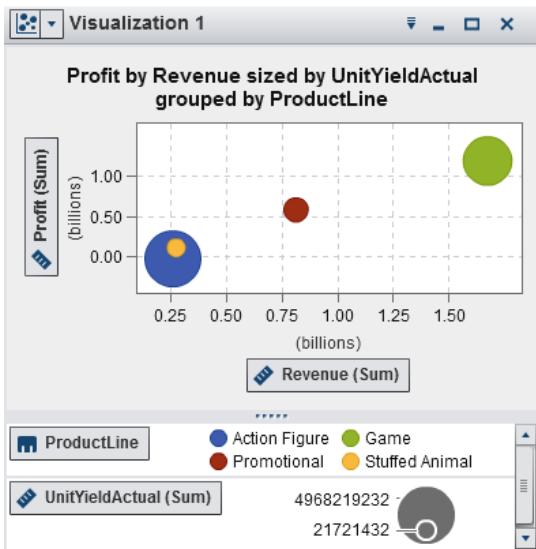
Displays the data as a bubble plot. A bubble plot displays the relationships among at least three measures. Two measures are represented by the plot axes, and the third measure is represented by the size of the plot markers.

You can apply grouping and create lattices. By assigning a datetime data item to the plot, you can animate the bubbles to display changes in the data over time.

For more information, see "Working with Bubble Plots" on page 215.

Network Diagram
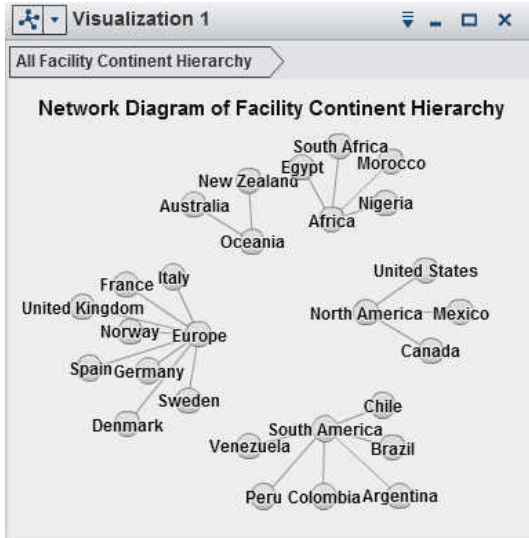
Displays a series of linked nodes. A network diagram displays the relationships between the values of categories or hierarchy levels.

You can indicate the values of measures through the sizes and colors of the nodes and the node links.

Histogram
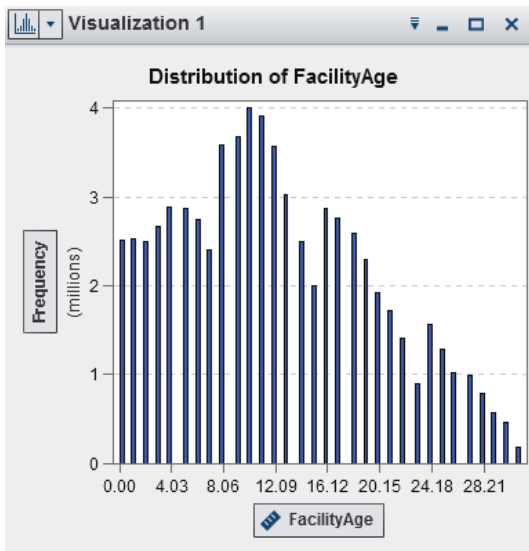
Displays the data as a histogram. A histogram displays the distribution of values for a single measure.

You can select the bar orientation, and you can select whether the values are displayed as a percentage or as a count.

For more information, see "Working with Histograms" on page 224.

Box Plot
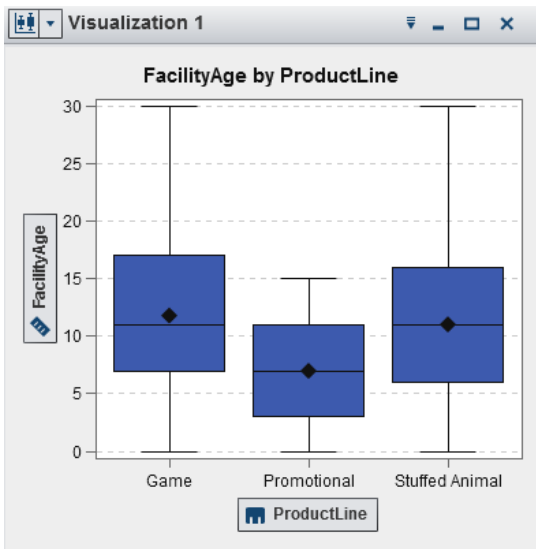
Displays the data as a box plot. A box plot displays the distribution of values for a single measure using a box and whiskers. The size and location of the box indicate the range of values that are between the 25th and 75th percentile. Additional statistical information is represented by other visual features.

You can create lattices, and you can select whether the average (mean) value and outliers are displayed for each box.

For more information, see "Working with Box Plots" on page 225.

Heat Map
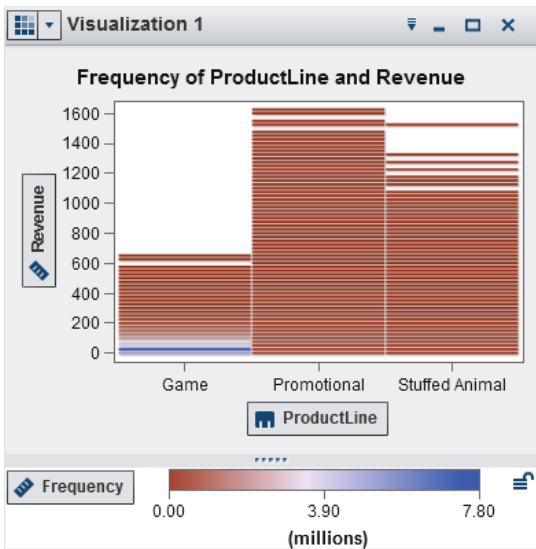
Displays the data as a heat map. A heat map displays the distribution of values for two data items using a table with colored cells. If you do not assign a measure to the **Color** data role, then a cell's color represents the frequency of each intersection of values. If you assign a measure to the **Color** data role, then a cell's color represents the measure value of each intersection of values.

For more information, see "Working with Heat Maps" on page 228.

Geo Map
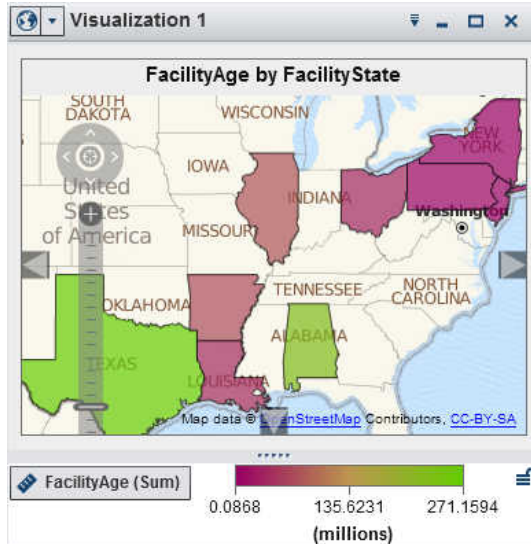
Displays the data as a geo map. A geo map displays your data as an overlay on a geographic map. You can display your data either as bubbles, as a scatter plot, or as colored regions.

For more information, see "Working with Geo Maps" on page 231.

Treemap
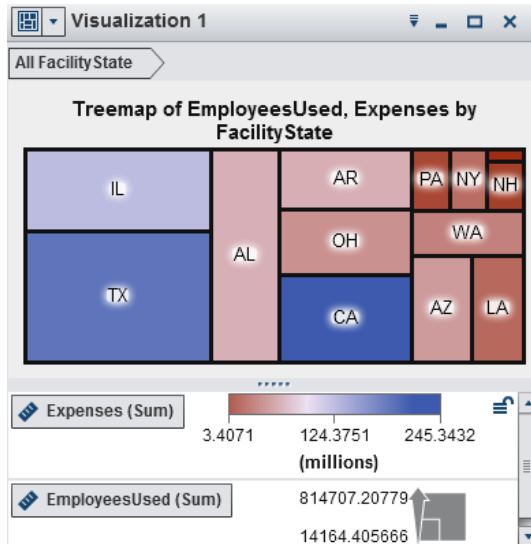
Displays the data as a treemap. A treemap displays your data as a set of rectangles (called tiles). Each tile represents a category value or a hierarchy node. The size of each tile represents either the frequency of the category or the value of a measure. The color of each tile can indicate the value of an additional measure.

For more information, see "Working with Treemaps" on page 234.

Correlation Matrix
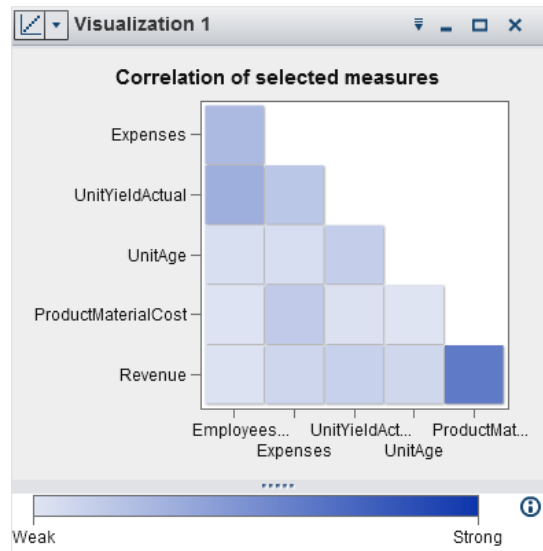
Displays the data as a correlation matrix. A correlation matrix displays the degree of correlation between measures as a series of colored rectangles. The color of each rectangle indicates the strength of the correlation.

For more information, see "Working with Correlation Matrices" on page 236.

Decision Tree
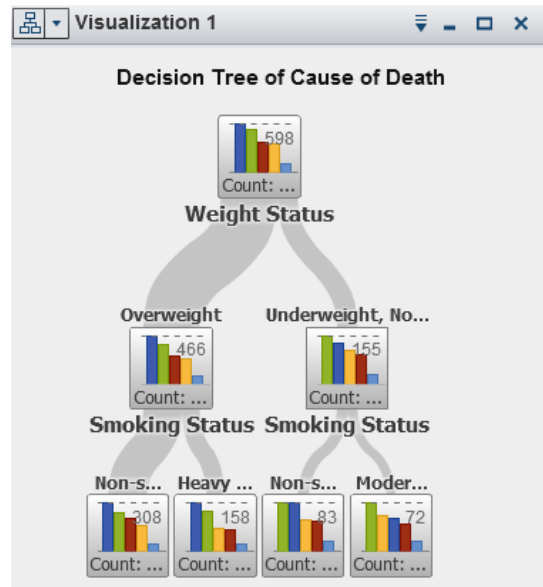
Displays the data as a decision tree. A decision tree displays a series of nodes as a tree, where the top node is the target data item, and each branch of the tree represents a split in the values of a predictor data item.

The splits enable you to see which values of the predictor data item correspond to different distributions of values in the target data item.

For more information, see "Working with Decision Trees" on page 238.

Word Cloud
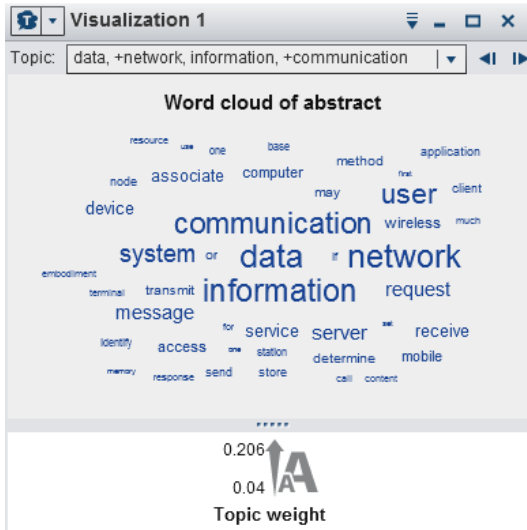
Displays a set of words from a character data item. Depending on the type of word cloud and your data roles, the size of each word in the cloud can indicate the relevance of the word to a topic, the frequency of the word in a category, or the value of a measure.

You can use text analytics in a word cloud to identify topics and terms that appear together in your data.

Working with Visualizations

Create a New Visualization

To create a new visualization, click on the toolbar, or select **Visualization** ▶ **New** from the menu bar.

Manage Visualizations

Delete a Visualization

To delete a visualization, click the ✕ button on the visualization, or select **Visualization** ▶ **Delete** from the menu bar to delete the selected visualization.

You can delete all of your visualizations by selecting **View** ▶ **Delete All Visualizations**.

Duplicate a Visualization

To duplicate a visualization, select **Visualization** ▶ **Duplicate** from the main menu, or select the ⬇ drop-down list from the visualization toolbar, and then select **Duplicate**.

The new visualization is named "Copy of *visualization-name*". You can enter a new name on the **Properties** tab.

Show and Minimize Visualizations

To minimize a visualization, click ▬ in the visualization that you want to minimize. The visualization appears in the dock pane at the bottom of the workspace.

To restore a visualization, click the name of the visualization in the dock pane.

You can select which visualizations are visible by using the Manage Visualizations window.

To minimize all visualizations, select **View** ▸ **Minimize All Visualizations** from the menu bar.

To show all visualizations, select **View** ▸ **Show All Visualizations** from the menu bar.

Maximize a Visualization

To maximize a visualization to fill the entire workspace, click ▢ in the visualization that you want to maximize. To return the visualization to its normal size, click ⊡ .

Arrange the Visualizations in the Workspace

Move the Position of a Visualization

To move a visualization, drag and drop the title bar of the visualization on the location where you want to place the visualization.

When you drag and drop a visualization onto another visualization, the placement of the visualization is determined by the edge that is closest to the drop point.

When your workspace contains visualizations in multiple rows, if you drag and drop a visualization onto the bottom edge or top edge of a visualization, then the target visualization and the moved visualization split the space that was previously occupied by the target visualization.

Resize a Visualization

To resize a visualization in your workspace, drag the resizing tab ···· at the bottom edge or the right edge of the visualization that you want to resize.

Using the Manage Visualizations Window

The **Manage Visualizations** window enables you to manage all of your visualizations together. Each visualization is represented by a thumbnail image or by an icon of the visualization type if the visualization has not been displayed in the current session.

To open the Manage Visualizations window, select **View ▶ Manage Visualizations** from the main menu bar.

The Manage Visualizations window enables you to perform the following tasks:

Task	Action
Add visualizations to the selection list. (Restore visualizations.)	Select a visualization from the **Available** list, and then click ➡, or click ⇒ to add all of the available visualizations.
Remove visualizations from the selection list. (Minimize visualizations.)	Select a visualization from the **Selected** list, and then click ⬅, or click ⇐ to remove all of the visualizations.
Change the order of the selected visualizations.	In the **Selected** list, drag and drop the visualizations in the order that you want, or select a visualization and click the up and down arrows to move it.
Search the available visualizations.	In the **Search** field, enter a string to search for. A visualization matches the search if the string appears anywhere in the visualization name. Only the matching visualizations appear in the **Available** list.

Task	Action
Filter the available visualizations.	Click ▼, and then select your filter parameters. You can filter either on the visualization type or on the data items that are used in each visualization. Only the matching visualizations appear in the **Available** list.
Rename a visualization.	Right-click on any visualization, and then select **Rename**. Enter a new name, and then click **OK**.
Delete a visualization.	Right-click on any visualization, and then select **Delete**.

Manage Visualization Comments

You can use the **Comments** tab to view and create comments for the current visualization. See "Sharing Comments in the Explorer" on page 285.

Managing Visualization Data Roles

Add a Data Item

For a visualization to display data, you must assign data items to it. You can assign data items to a visualization by performing any of the following actions:

■ Drag and drop the data item onto the center of the visualization. The data item is assigned to a role automatically. If the visualization already has data items assigned to the required roles, then you can choose how to assign the new data item.

■ Drag and drop the data item onto a **Measure** or **Category** button in the visualization to assign the data item to a specific data role.

- From the visualization toolbar, select the ⊽ drop-down list, and then select **Add Category** or **Add Measure**.

- Use the **Roles** tab in the right pane. Either drag and drop a data item onto a role, or select **Add** from the drop-down list, and then select a data item.

Each visualization requires a minimum number of each type of data item. The following table lists the requirements for each visualization:

Table 24.1 *Required Data Items for Visualizations*

Visualization Type	Requirements
Automatic Chart	one data item of any type
Table	one data item of any type (except aggregated measure)
Crosstab	one data item of any type
Bar Chart	one category or hierarchy
Line Chart	one category or hierarchy
Scatter Plot	one measure
Bubble Plot	three measures
Network Diagram	one category or hierarchy
Histogram	one measure
Box Plot	one measure
Heat Map	two data items of any type (except aggregated measure)
Geo Map	one geography
Treemap	one category or hierarchy
Correlation matrix	two measures

Visualization Type	Requirements
Decision tree	two data items of any type (except aggregated measure)
Word Cloud	one category or document collection

Replace a Data Item

You can replace a data item by using any of the following methods:

■ Drag and drop the new data item from the **Data** pane onto the data item in the visualization that you want to replace.

■ Right-click on the data item that you want to replace in either the visualization or on the **Roles** tab, and then select **Replace** *item-name* ▶ *new-item-name*.

Remove a Data Item

You can remove a data item by using any of the following methods:

■ Drag and drop the data item from the visualization onto the **Data** pane.

■ Select the ▼ drop-down list from the visualization toolbar, and then select **Remove** ▶ *item-name* .

■ Right-click on the data item that you want to delete in either the visualization or on the **Roles** tab, and then select **Remove** *item-name*.

Switch Data Roles

To switch the data items that are assigned to two roles, drag and drop one data item onto another data item either by using the buttons in the visualization or by using the fields on the **Roles** tab.

For example, you can switch the axes of a scatter plot by dragging the measure on the X axis to the measure on the Y axis.

Working with Filters

You can use the **Filters** tab to subset the data in your visualizations. See Chapter 25, "Working with Filters," on page 249.

Ranking Data

Overview of Ranking

You can use the **Ranks** tab to create ranks to subset the data in your visualizations. A rank selects either the top (greatest) or the bottom (least) aggregated value for a category.

A rank selects values for a category based on either the frequency of the category values or the aggregated value of a measure.

For example, you might create a rank of the top 10 countries by frequency to select the 10 countries that are most represented in your data source. As another example, you might create a rank of the top 10 countries by population to select the 10 countries with the greatest population.

Note: If the category for the rank is part of a hierarchy that is used in the current visualization, then the rank is applied only when the hierarchy is drilled to the level of the rank category.

Create a New Rank

To create a rank:

1 From the **Data** pane, select the data item that you want to use as the base of the rank. You can select any category or geography data item, regardless of whether it is assigned to the current visualization.

2 Either right-click on the data item, and select **Add Rank on [*data-item*]**, or drag and drop the data item onto the **Ranks** tab.

3 Set the parameters for the rank:

   a Select the type of rank from the drop-down list. **Top** specifies that the rank selects the greatest value. **Bottom** specifies that the rank selects the least value.

   b Specify the number of values for the rank. For example, if you specify 5, and you select **Top** as the rank type, then the rank selects the five greatest values.

   c From the **By** drop-down list, select the measure that is used to create the rank. You can either select any measure, or select **Frequency** to use the frequency of the rank category.

     Note: If you select a measure that is used in the current visualization, then the rank uses the same aggregation for the measure that is used by the current visualization.

   d (Optional) Select **Ties** to include ties in the rank.

     If you select **Ties**, then the rank selects as many values as necessary to include all of the ties. If you do not select **Ties**, then the rank selects only the number of values that are specified by the rank parameters.

     For example, if your rank selects the top three values, but there are five values tied for the greatest value, then the number of values that are selected by the rank depends on the **Ties** option. If you select **Ties**, then the rank includes all five of the tied values. If you do not select **Ties**, then the rank includes only three of the tied values.

     Note: If the ranking does not select all of the tied values, or if the number of tied values exceeds the maximum that is set by your administrator, then a message appears.

Note: By default, your rank changes are applied automatically to the current visualization. To apply multiple changes together, deselect **Auto**, and then click **Apply** when you are ready to apply your rank changes.

Delete a Rank

To delete a rank, click ✕ on the rank on the **Ranks** tab.

Manage Visualization Axes

Lock an Axis

By default, the axes of your visualizations adjust automatically to your data. If you change the data that is displayed by applying filters or ranks or by drilling down, then the data ranges and scaling of your axes will change to fit the data.

For comparison purposes, you might want to lock the axes and retain the same data ranges and scaling. To lock an axis, click the 🔓 icon beside the axis heading.

To re-enable automatic adjustments for an axis, click the 🔒 icon beside the axis heading.

Adjust an Axis

For a locked axis, you can adjust the visible data range. You can adjust the range by dragging the scroll bar that appears over the axis tick marks or by right-clicking a measure heading, and then selecting **Set Visible Axis Range**.

Transfer Axis Settings

For a locked axis, you can transfer the axis settings to a compatible visualization. A compatible visualization must have a measure (or frequency) assigned to an axis.

To transfer an axis, right-click the measure heading, and then select **Transfer Axis Settings**. From the Select Compatible Visualizations window, select the visualization to which you want to transfer the axis settings.

Working with Visualization Data Ranges and Color Gradients

Support for Customized Data Ranges and Color Gradients

The following visualization types enable you to customize the data range and color gradient:

- Heat map

- Geo map (with the **Color** role assigned)

- Bubble plot (with continuous data assigned to the **Color** role)

- Network diagram (with the **Node Color** or **Link Color** role assigned)

- Word cloud (with the **Color** role assigned)

- Treemap (with the **Color** role assigned)

Specify a Custom Color Gradient

To specify a custom color gradient:

1 Right-click on the color gradient in the legend, and then select **Configure Color Gradient**.

2 From the Select Color Gradient window, select a color gradient from the **Color gradient** drop-down list.

3 Click **OK** to apply the new color gradient.

Specify a Custom Data Range

To specify a custom data range:

1 Right-click on the color gradient in the legend, and then select **Configure Color Gradient**.

2 From the Select Color Gradient window, deselect **Automatically adjust color range to data**.

3 In the **Lower** field, specify the lower bound for the data range.

4 If the selected color gradient uses three colors, specify the inflection point of the gradient in the **Inflection** field, or select **Use midpoint** to use the midpoint between the lower and upper values as the inflection point. The inflection point is the point for the middle color in a three-color gradient.

5 In the **Upper** field, specify the upper bound for the data range.

6 Click **OK** to apply the new color gradient.

Share a Color Gradient and Data Range between Visualizations

To share data ranges and color gradients between visualizations, all of the visualizations must support customized data ranges. See "Support for Customized Data Ranges and Color Gradients" on page 196.

To share a data range and color gradient between visualizations:

1 If the data range has not been customized, then right-click on the legend for the color gradient and data range that you want to share, and then select **Lock Data Range**.

2   Right-click on the legend for the color gradient and data range that you want to
    share, and then select **Transfer Configuration**. The Select compatible
    visualizations window appears.

3   Select the visualizations that you want to share the color gradient and data range
    with. The visualizations are sorted based on their similarity to the source
    visualization. Visualizations that use the same data items and visualizations that are
    the same type are shown at the top of the list.

4   When you are finished, click **OK** to apply the color gradient and data range to all of
    the selected visualizations.

Remove Customized or Shared Color Data Ranges

To remove a customized or shared data range from a visualization, right-click on the
legend, and select **Remove Custom Data Range**.

Working with Data Brushing

Overview of Data Brushing

Data brushing is a feature of some visualizations that enables you to select data values
in one visualization and highlight the corresponding data values in all of your other
visualizations.

For example, you might have a bar chart and a scatter plot in your exploration. If you
enable data brushing, and you select a bar in the bar chart, then the markers in the
scatter plot that correspond to the selected value in the bar are highlighted.

Figure 24.1   *Data Brushing Example*

The following visualization types support data brushing:

- Bar chart

- Line chart

- Scatter plot

- Bubble plot (if the **Group** role is assigned)

- Network diagram

- Histogram

- Heat map (if categories are assigned to both axes)

- Geo map

- Treemap

- Word cloud
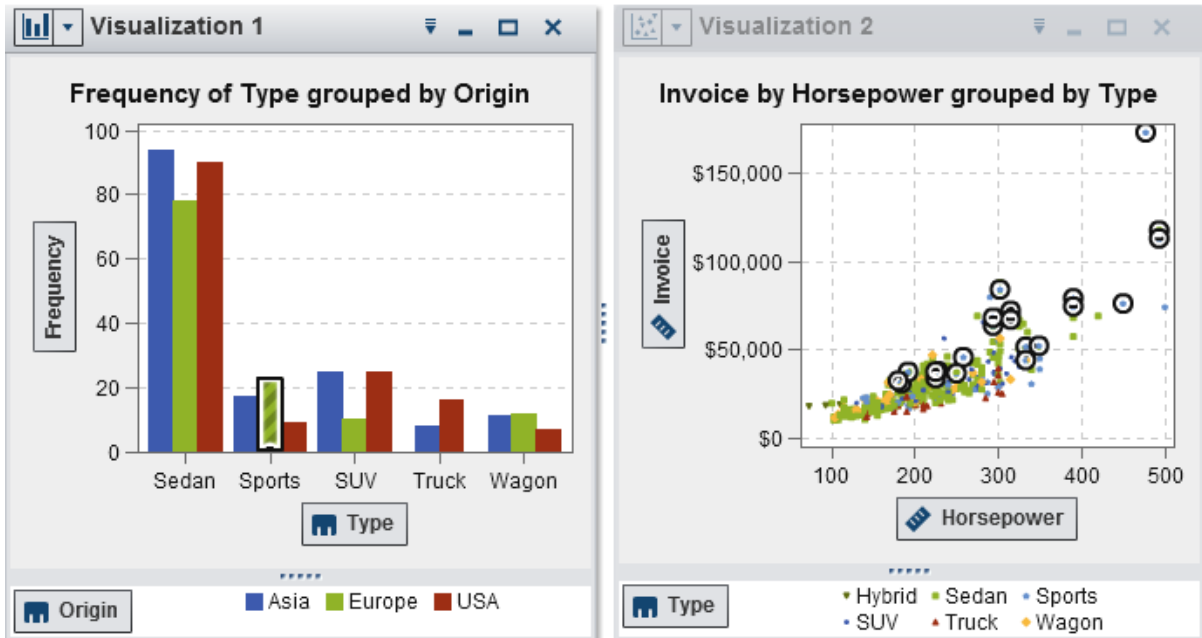
For histograms and for bar charts of frequency that are not grouped or latticed, data brushing can indicate the portion of a bar that corresponds to the brushed data. The bar is highlighted and then partially shaded to indicate the corresponding portion. The value for the shaded portion of the bar is included in the data tip for the bar.

In Figure 24.2 , a selection in a bubble plot highlights a portion of the corresponding bar in a bar chart.

Figure 24.2 Data Brushing with Partial Shading
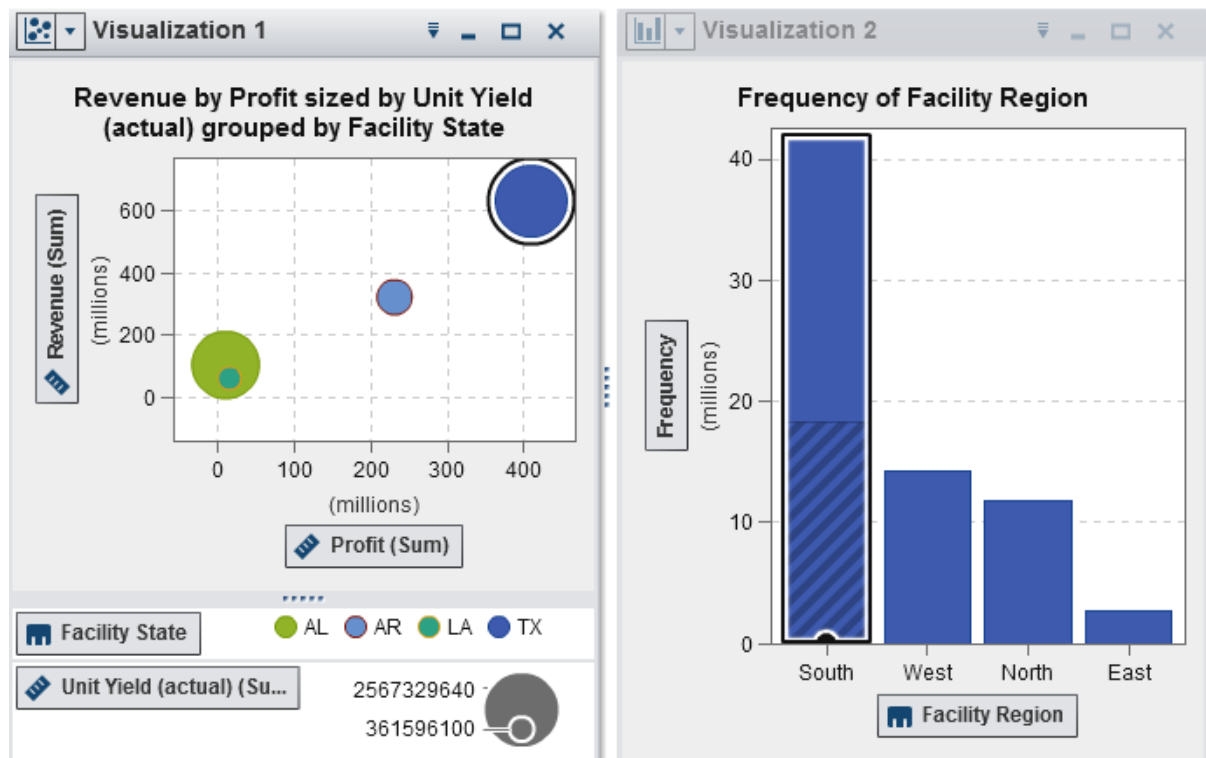
For bar charts that have a measure assigned or that are grouped or latticed, the entire bar is highlighted.

Enable Data Brushing

To enable data brushing, select **View ▶ Data Brushing**.

Select Values in a Visualization

You can select values in your visualization by using any of the following methods:

- Click the data values. This is most useful for bar charts, histograms, and treemaps. To select multiple data values, hold down the Ctrl key while clicking each data value.

- Click and drag to select data values in a rectangular region. This is most useful for scatter plots and heat maps.

- Select data values in the details table for the visualization. Press Shift+click to select multiple adjacent values, or press Ctrl+click to select or deselect individual values.

Working with Automatic Charts

Automatic charts display different types of charts automatically based on the data items that are assigned to the chart.

The type of chart can be any of the following:

Table 24.2   *Automatic Chart Types*

Data Items	Chart Type
One measure	Histogram
One category	Bar chart
One aggregated measure	Crosstab
One datetime category and any number of other categories or measures	Line chart
One geography and up to two measures	Geo map
One geography and three or more measures	Bar chart

Data Items	Chart Type
One document collection	Word cloud
Two measures	Scatter plot or heat map
Three or more measures	Scatter plot matrix or correlation matrix
One or more categories and any number of measures and geographies	Bar chart

Note: The **Roles** tab for an automatic chart always contains **Categories** and **Measures**. To set advanced data roles such as grouping and lattices, click the **Use Chart-Type** link on the **Roles** tab to change the visualization from an automatic chart to the chart type that is currently shown. For example, if your automatic chart shows a bar chart, then click **Use Bar Chart** to change the type to a bar chart and enable advanced data role settings.

Working with Bar Charts

About Bar Charts

A bar chart displays data by using bars. The height of each bar represents the value.

Data Roles for a Bar Chart

The basic data roles for a bar chart are categories and measures. You can assign one category only, and the category values are plotted on the category axis. You can assign many measures, and the measure values are plotted on the response axis. If a bar chart contains no measures, then the frequency of the category values is plotted on the response axis.

In addition to the basic data roles, you can assign these roles:

Group

groups the data based on the values of the category data item that you assign. Depending on the value that you selected for the **Grouping style** property, the group values are shown as either individual bars or as segments of each bar.

Note: Grouping is not available if you assign multiple measures to the visualization.

Lattice columns

creates a lattice of charts with a column for each value of the category data item that you assign.

Lattice rows

creates a lattice of charts with a row for each value of the category data item that you assign.

Specify Properties for a Bar Chart

On the **Properties** tab, you can specify the following options:

Name

specifies the name of the visualization.

Title

specifies the title that appears above the graph.

Note: The **Title** option is disabled if you select **Generate graph title**.

Generate graph title

specifies that the graph title is generated automatically based on the data items in the visualization.

Show grid lines

displays grid lines for each tick on the response axis.

Bar direction

specifies whether the bars are vertical or horizontal.

Frequency

specifies whether the frequency is displayed as a count (**Count**) or as a percentage (**Percent**).

Note: The frequency values are based on the data that is shown in the visualization (after filters and other data selections have been applied).

Note: This option has no effect if a measure is assigned to the visualization.

Grouping style

specifies how grouped data is displayed. If you select **Stack**, then the values of the grouping variable are displayed as segments of each bar. If you select **Cluster**, then each value of the grouping variable is displayed as a separate bar.

Note: This option is disabled if no data item is assigned to the **Group** role.

Measure layout

specifies whether the measures share a single response axis (**Shared axis**) or have separate response axes for each measure (**Separate axes**).

Note: The **Measure layout** option is disabled if your visualization contains exactly one measure.

Overview

specifies whether the chart overview is enabled.

Show data labels

shows the data values as text in the visualization.

Note: You can always view a data value as a data tip when you position the cursor over a data value.

Rotate axis labels

displays the category labels at an angle.

Note: The **Rotate axis labels** option has no effect when you select a horizontal bar orientation.

Reference Lines

adds reference lines to the visualization. To create a reference line, click ⌶, and then specify the parameters for the line in the New Reference Line window.

The name of each reference line appears below the **Reference Lines** option. You can edit or delete your reference lines by using the icons next to each name.

Sort Data Values

By default, a bar chart is sorted in descending order by the value of the first measure. To change the sorting, right-click on the data item that you want to sort on, and then select **Sort ▶ [*sort-method*]**. For a grouped bar chart, the data is sorted by the category values in alphabetical order.

Note: If the visualization contains a rank, then, by default, the data is sorted based on the values of the rank.

Working with Line Charts

About Line Charts

A line chart displays data by using a line that connects the data values. If you assign multiple measures to a line chart, then you can create separate Y axes for each measure.

Data Roles for a Line Chart

The basic data roles for a line chart are categories and measures. You can assign one category only, and the category values are plotted on the category axis. You can assign many measures, and the measure values are plotted on the response axis. If the line chart contains no measures, then the frequency of the category values is plotted on the response axis.

In addition to the basic data roles, you can assign these roles:

Group
groups the data based on the values of the category data item that you assign. A separate line is created for each data value.

Note: Grouping is not available if you assign multiple measures to the visualization.

Lattice columns

    creates a lattice of charts with a column for each value of the category data item that you assign.

Lattice rows

    creates a lattice of charts with a row for each value of the category data item that you assign.

Underlying factors

    if forecasting is enabled, adds additional measures to the forecast as underlying factors. The forecasting model evaluates the additional measures to determine whether they contribute to the accuracy of the forecast. If the additional measures do not increase the accuracy of the forecast, then they are not used. If the additional measures do increase the accuracy of the forecast, then the forecast line is adjusted, and the confidence bands are narrowed.

    The measures that you add as underlying factors can also be used in a scenario analysis.

    Note: Underlying factors are available only if your line chart contains a single measure in the **Measures** role.

Specify Properties for a Line Chart

On the **Properties** tab, you can specify the following options:

Name

    specifies the name of the visualization.

Title

    specifies the title that appears above the graph.

    Note: The **Title** option is disabled if you select **Generate graph title**.

Generate graph title

    specifies that the graph title is generated automatically based on the data items in the visualization.

Show grid lines

displays grid lines for each tick on the response axis.

Frequency

specifies whether the frequency is displayed as a count (**Count**) or as a percentage (**Percent**).

Note: The frequency values are based on the data that is shown in the visualization (after filters and other data selections have been applied).

Note: This option has no effect if a measure is assigned to the visualization.

Measure layout

specifies whether the measures share a single response axis (**Shared axis**) or have separate response axes for each measure (**Separate axes**).

Note: The **Measure layout** option is disabled if your visualization contains exactly one measure.

Overview

specifies whether the chart overview is enabled.

Show markers

shows markers for the data points in the visualization.

Show data labels

shows the data values as text in the visualization.

Note: You can always view a data value as a data tip when you position the cursor over a data value.

Rotate axis labels

displays the category labels at an angle.

Reference Lines

adds reference lines to the visualization. To create a reference line, click , and then specify the parameters for the line in the New Reference Line window.

The name of each reference line appears below the **Reference Lines** option. You can edit or delete your reference lines by using the icons next to each name.

Duration (if forecasting is enabled)

specifies the number of data intervals to forecast.

Note: This option is available only if forecasting is enabled for the visualization.

Sort Data Values

By default, for category data, a line chart is sorted in descending order by the value of the first measure. To change the sorting, right-click on the data item that you want to sort on, and then select **Sort ▶ [*sort-method*]**.

Note: If the visualization contains a rank, then, by default, the data is sorted based on the values of the rank.

Note: Sorting is not available if a datetime data item is assigned to the **Category** role.

Forecasting

About Forecasting

Forecasting uses the statistical trends in your data source to predict future data values. Forecasting is available only if a date or datetime data item is assigned to the visualization.

In addition to the predicted future data values, the forecast displays a confidence band. For more information, see

Enable Forecasting

To add forecasting to a line chart, select the ▼ drop-down list from the visualization toolbar, and then select **Show Forecast**.

Note: The line chart must contain a date or datetime data item to apply forecasting.

On the **Properties** tab for the line chart, you can adjust the number of data values to predict by using the **Duration** option.

On the **Roles** tab, you can add additional measures to the forecast as **Underlying factors**. The forecasting model evaluates the additional measures to determine whether

they contribute to the accuracy of the forecast. If the additional measures do not increase the accuracy of the forecast, then they are not used. If the additional measures do increase the accuracy of the forecast, then the forecast line is adjusted, and the confidence bands are narrowed.

The measures that you add as underlying factors can also be used in a scenario analysis. For more information, see "Apply Scenario Analysis to a Forecast" on page 282.

Working with Tables

About Tables

A table displays data as text. The data value for each measure or category that is assigned to the table is displayed as a column. The data values in the table are not aggregated.

Note: For very large data sources, the table visualization displays only the first two billion (2,147,483,647) rows.

Note: If you sort the table, then the table displays only the first 5,000 sorted rows.

Data Roles for a Table

The basic data role for a table is a column. A column can be any type of data item. You can add any number of columns to a table.

Specify Properties for a Table

On the **Properties** tab, you can specify the following option:

Name
   specifies the name of the visualization.

Managing Columns

Sort Columns

To sort the table by a column, click on the column heading. An arrow appears in the column heading to indicate the sorting. If the arrow points upward, then the sort is ascending. If the arrow points downward, then the sort is descending.

Note: If you sort the table, then the table displays only the first 5,000 sorted rows.

Rearrange Columns

To rearrange your columns, drag and drop the column headings.

Resize Columns

To resize a column, click, and then drag and drop the left or right edge of a column heading.

Working with Crosstabs

About Crosstabs

A crosstab displays the intersections of category values and measure values as text. If the crosstab contains measures, then each cell of the crosstab contains the aggregated measure values for a specific intersection of category values. If the crosstab does not contain measures, then each cell of the crosstab contains the frequency of an intersection of category values.

Data Roles for a Crosstab

The basic data roles for a crosstab are columns, rows, and measures. You can assign either a single hierarchy or any number of categories to each column and row role. If you assign measures to the crosstab, then the measure values are displayed in the cells of the crosstab. If you do not assign measures, then the frequency of each intersection of values is displayed in the cells of the crosstab.

Specify Properties for a Crosstab

On the **Properties** tab, you can specify the following options:

Name
specifies the name of the visualization.

Indented
selects the indented layout for the crosstab.

Show column subtotals
adds subtotals to each column for each node on the row axis after the first.

Note: For the indented layout, subtotals are always enabled.

Show column totals
adds totals to each column.

Note: Total values are aggregated based on the default aggregations for each measure.

Show row subtotals
adds subtotals to each row for each node on the column axis after the first.

Note: For the indented layout, subtotals are always enabled.

Show row totals
adds totals to each row.

Note: Total values are aggregated based on the default aggregations for each measure.

Totals Placement
specifies the location of totals and subtotals. Select **Before** to place the totals and subtotals before the axis headings. Select **After** to place the totals and subtotals after the axis headings.

Note: For the indented layout, totals are always placed before the axis headings.

Managing Rows and Columns

Sort a Row or Column

By default, the crosstab is sorted alphabetically by the values of the first category that you assign to the **Rows** role. To change the sorting, right-click on the heading for the row or column that you want to sort by, and then select **Sort** ▶ [*sort-method*].

If you apply sorting to a measure, then the sorting for the rows is applied as a secondary sort. The outer row values are sorted by their subtotals for the measure. To remove measure sorting, select a sort for any category.

Rearrange Rows and Columns

To rearrange your rows and columns, drag and drop the row headings and column headings.

Resize Columns

To resize a column, click, and then drag and drop the left or right edge of the column heading.

Create a Hierarchy from a Crosstab

Crosstabs enable you to create hierarchies from the categories on a crosstab axis. To create a hierarchy, right-click on a category heading, and then select **Create Hierarchy**. The categories are replaced with a new hierarchy.

The name of the new hierarchy is generated from the name of the outermost category, with the suffix "Hierarchy."

Working with Scatter Plots

About Scatter Plots

A scatter plot displays the values of measures by using markers. When you apply more than two measures, the visualization displays a scatter plot matrix. A scatter plot matrix is a series of scatter plots that display every possible pairing of the measures that are applied to the visualization.

If you create a scatter plot that has a very large number of data values, then the scatter plot is rendered either as a heat map (for two measures) or as a correlation matrix for three or more measures.

Data Roles for a Scatter Plot

The basic data role for a scatter plot is a measure. You can assign any number of measures. If you assign a single measure to a scatter plot, then the values are plotted along a line.

In addition to measures, you can assign a **Group** variable. The **Group** variable groups the data based on the values of the category data item that you assign. A separate set of scatter points is created for each value of the group variable.

You can add data items to the **Data Tips** role. The values for the data items in the **Data Tips** role are displayed in the data tips for the scatter plot.

Specify Properties for a Scatter Plot

On the **Properties** tab, you can specify the following options:

Name
specifies the name of the visualization.

Title
specifies the title that appears above the graph.

Note: The **Title** option is disabled if you select **Generate graph title**.

Generate graph title
specifies that the graph title is generated automatically based on the data items in the visualization.

Show grid lines
displays grid lines for each tick on the plot axes.

Marker size
specifies the size of each marker in pixels.

Fit Line
adds a fit line to the scatter plot. For information about the fit types that are available, see .

Note: Fit lines are not available if a grouping variable is assigned to the scatter plot.

Reference Lines
adds reference lines to the visualization. To create a reference line, click 🔾, and then specify the parameters for the line in the New Reference Line window.

The name of each reference line appears below the **Reference Lines** option. You can edit or delete your reference lines by using the icons next to each name.

Applying Data Analysis

About Data Analysis

For scatter plots, you can apply the following data analyses:

Correlation
identifies the degree of statistical correlation between the variables in the visualization. For more information, see .

Fit Line
plots a model of the relationship between the variables in the visualization.

There are many types of fit lines, including linear fit, quadratic fit, cubic fit, and penalized B-spline. For more information, see .

Correlation is applied to your visualization automatically when you add a linear fit line. It is not available with other fit types.

Enable Data Analysis

To add a fit line to your visualization, select the ⯆ drop-down list from the visualization toolbar, and then select **Fit Line** ▶ [**fit-type**]. For details about the fit types that are available, see "Fit Lines" on page 279.

Working with Bubble Plots

About Bubble Plots

A bubble plot displays the values of at least three measures by using differently sized plot markers (bubbles) in a scatter plot. The values of two measures are represented by the position on the plot axes, and the value of the third measure is represented by the marker size.

You can create animated bubble plots to display changing data over time.

Data Roles for a Bubble Plot

The basic data roles for a bubble plot are:

X axis
  specifies the measure that is assigned to the X axis.

Y axis
  specifies the measure that is assigned to the Y axis.

Bubble size
  specifies the measure that determines the marker size.

In addition to the basic data roles, you can assign these roles:

Group
  groups the data based on the values of the category data item that you assign. A separate set of points is created for each value.

Note: You cannot assign both the **Group** role and the **Color** role at the same time.

Color
　　specifies a data item that determines the color of the bubbles. If you specify a category, then each value of the category is represented by a different bubble color. If you specify a measure, then the measure value is represented by the bubble color.

　　Note: You cannot assign both the **Group** role and the **Color** role at the same time.

Lattice columns
　　creates a lattice of charts with a column for each value of the category data item that you assign.

Lattice rows
　　creates a lattice of charts with a row for each value of the category data item that you assign.

Animation
　　specifies a datetime data item that is used to animate the bubble plot.

　　Note: The **Animation** role is enabled only if you assign a data item to the **Group** role.

Specify Properties for a Bubble Plot

On the **Properties** tab, you can specify the following options:

Name
　　specifies the name of the visualization.

Title
　　specifies the title that appears above the graph.

　　Note: The **Title** option is disabled if you select **Generate graph title**.

Generate graph title
　　specifies that the graph title is generated automatically based on the data items in the visualization.

Show grid lines

   displays grid lines for each tick on the response axis.

Size scale

   specifies the type of scaling that is used to draw the bubbles. Select one of the following values:

   Linear

   specifies that the bubble sizes are scaled relative to the minimum and maximum values (or frequencies) in your data. A negative value is displayed as a smaller bubble than a positive value.

   For the **Linear** scaling type, the difference in bubble sizes might not be proportional to the difference in values.

   Magnitude

   specifies that the bubble sizes are scaled relative to zero and the greatest absolute value in your data. A negative value is displayed as a hexagon.

   For the **Magnitude** scaling type, the difference in bubble sizes is proportional to the difference in absolute values.

Frequency

   specifies whether the frequency is displayed as a count (**Count**) or as a percentage (**Percent**).

   Note: The frequency values are based on the data that is shown in the visualization (after filters and other data selections have been applied).

   Note: This option has no effect if a measure is assigned to the **Bubble size** role.

Color gradient

   selects the gradient colors for the visualization.

   You can click ⬚ to select the values that are used to assign the colors. See "Specify a Custom Data Range" on page 197.

Reference Lines

   adds reference lines to the visualization. To create a reference line, click ⬚, and then specify the parameters for the line in the New Reference Line window.

The name of each reference line appears below the **Reference Lines** option. You can edit or delete your reference lines by using the icons next to each name.

Using Animated Bubble Plots

About Animated Bubble Plots

An animated bubble plot displays the changes in your data values over time. Each frame of the animation represents a value of the datetime data item that is assigned to the **Animation** data role.

For example, if you assign a category with the YEAR format to the **Animation** data role, then each frame of the animation displays a bubble plot of your data for a specific year.

Create an Animated Bubble Plot

To create an animated bubble plot:

1   Select an existing bubble plot, or create a new bubble plot.

2   Assign a data item to the **Group** data role.

3   Assign a data item with a datetime format to the **Animation** data role.

Display an Animated Bubble Plot

For an animated bubble plot, a set of animation controls appears at the bottom of the visualization.

Table 24.3   Animation Control Tasks

Task	Action
Start the animation.	Click ►.
Go to the previous animation frame.	Click ◄ɪ.
Go to the next animation frame.	Click ɪ►.
Jump to a specific animation frame.	Use the slider.

Task	Action
Specify whether to repeat the animation.	Select or deselect **Loop**.
Select the animation speed.	Use the **Speed** slider.
Track the movement of a specific bubble.	Click the bubble that you want to track.

Note: If you select a bubble to track, the selected bubble is not highlighted.

Working with Network Diagrams

About Network Diagrams

A network diagram displays the relationships between category values as a series of linked nodes.

You can create two types of network diagrams:

Hierarchical
  creates a hierarchical structure by using a hierarchy or a set of categories.

Ungrouped
  creates a structure by using a source data item and a target data item. A node is created for each value of the source data item, and a link is created from each node to the node that corresponds to the value of the target data item.

  For example, if your source data item specifies the name of every employee in an organization, and your target data item specifies the manager of each employee, then the network diagram has a node for each employee that is linked to the node for the employee's manager.

Data Roles for a Network Diagrams

Basic Data Roles for a Hierarchical Network Diagram

The basic data role for a hierarchical network diagram is **Levels**. Specify a hierarchy or any number of categories. The data items in the **Levels** role specify the nodes of the network diagram. If you add categories to the **Levels** role, then the order of the categories determines their hierarchical relationship. You can drag and drop the categories to arrange them.

Basic Data Roles for an Ungrouped Network Diagram

The basic data roles for an ungrouped network diagram are the **Source** and the **Target**. The **Source** specifies a category that contains all of the node values for the plot. The **Target** specifies a category that creates the links between nodes. The **Target** category must contain a subset of the values of the **Source** category.

Additional Data Roles for a Network Diagram

In addition to the basic data roles, you can specify the following additional data roles for a network diagram:

Node Size
    specifies a measure that determines the size of the nodes in the diagram.

Node Color
    specifies a data item that determines the color of the nodes in the diagram.

Link Width
    specifies a measure that determines the width of the links in the diagram.

Link Color
    specifies a data item that determines the color of the links in the diagram.

Data Tips
    specifies data items whose values are included in the data tips for the diagram. Measure values are aggregated by sum.

Specify Properties for a Network Diagram

On the **Properties** tab, you can specify the following options:

Name
specifies the name of the visualization.

Title
specifies the title that appears above the graph.

Note: The **Title** option is disabled if you select **Generate graph title**.

Generate graph title
specifies that the graph title is generated automatically based on the data items in the visualization.

Show Map
displays the network as an overlay on a geographic map.

Note: This option is available only if geographies are assigned to all of the data roles that create nodes. For hierarchical diagrams, all of the **Levels** must be geographies. For ungrouped diagrams, the **Source** and **Target** must be geographies.

Show Labels
shows the node values as text inside each node.

Show Direction
shows the links in the diagram as arrows.

Map service
if your diagram contains a map, specifies the source for the background map.

Node color
selects the color gradient for the nodes in the diagram. You can click ⬚ to select the values that are used to assign the colors. See "Specify a Custom Data Range" on page 197.

Note: This option is available only if you assign a measure to the **Node Color** data role.

Link color

selects the color gradient for the links in the diagram. You can click ⬐ to select the values that are used to assign the colors. See "Specify a Custom Data Range" on page 197.

Note: This option is available only if you assign a measure to the **Link Color** data role.

Additional levels

for hierarchical diagrams only, specifies the number of levels that are displayed beneath the current level.

Node spacing

specifies the amount of space between nodes in the diagram.

Source node

displays the current node selection. To remove the node selection, select 🗑.

Predecessors

selects the number of levels of predecessors (parents) of the source node to select. 0 specifies the source node. Note that you can specify a range by making multiple selections. For example, select both 0 and 1 to specify that the source node and the first level of predecessors are selected.

Successors

selects the number of levels of successors (children) of the source node to select. 0 specifies the source node. Note that you can specify a range by making multiple selections. For example, select both 0 and 1 to specify that the source node and the first level of successors are selected.

Select Nodes in a Network Diagram

You can select nodes in the diagram by using any of the following methods:

- Click and drag the cursor to create a rectangular selection.

- Hold down the Ctrl key and click the nodes that you want to select.

- Select a series of linked nodes by setting a node as the source node.

Right-click a node, and then select **Set as Source for Selection**.

On the **Properties** tab, specify the range of levels of **Predecessors** (parents) and **Successors** (children) of the source node to select. 0 specifies that the source node is selected.

For example, if you specify a range of 0–1 for **Predecessors** and a range of 0–2 for **Successors**, then the source node, one level of predecessors, and two levels of successors are selected.

Control the View of a Network Diagram

You can control the view of a network diagram by using the following controls:

Zoom	Zoom in and out at the location of the cursor by scrolling the mouse wheel.
Pan (scroll)	Pan (scroll) the diagram by holding down the Shift key and dragging the diagram.
Rotate	Rotate the diagram by holding down the Ctrl key and dragging the diagram.

Create a Hierarchy from a Network Diagram

If your network diagram contains categories in the **Levels** role, then you can create a new hierarchy by using the categories.

On the **Roles** tab, click on the drop-down list for the **Levels** role, and then select **Create Hierarchy**. A new hierarchy is created with the name of the first category in the hierarchy.

Working with Histograms

About Histograms

A histogram displays the distribution of values for a single measure. A series of bars represents the number of observations in the measure that match a specific value or value range. The bar height can represent either the exact number of observations or the percentage of all observations for each value range.

Note: If you use the default number of bins, then the minimum and maximum values on the histogram axis might not match the actual extent of your data values. If you specify the number of histogram bins, then the histogram axis matches your data values exactly.

Data Roles for a Histogram

The basic data role for a histogram is a measure. You can assign only one measure to a histogram.

Specify Properties for a Histogram

On the **Properties** tab, you can specify the following options:

Name
specifies the name of the visualization.

Title
specifies the title that appears above the graph.

Note: The **Title** option is disabled if you select **Generate graph title**.

Generate graph title
specifies that the graph title is generated automatically based on the data items in the visualization.

Show grid lines

 displays grid lines for each tick on the response axis.

Bar direction

 specifies whether the bars are vertical or horizontal.

Frequency

 specifies whether the frequency is displayed as a count (**Count**) or as a percentage (**Percent**).

 Note: The frequency values are based on the data that is shown in the visualization (after filters and other data selections have been applied).

Use default bin count

 specifies whether to use the default number of bins (value ranges) for the histogram. The default number of bins is determined by the number of data values in your histogram.

Bin count

 specifies the number of bins (value ranges) for the histogram.

Reference Lines

 adds reference lines to the visualization. To create a reference line, click ⊞, and then specify the parameters for the line in the New Reference Line window.

 The name of each reference line appears below the **Reference Lines** option. You can edit or delete your reference lines by using the icons next to each name.

Working with Box Plots

About Box Plots

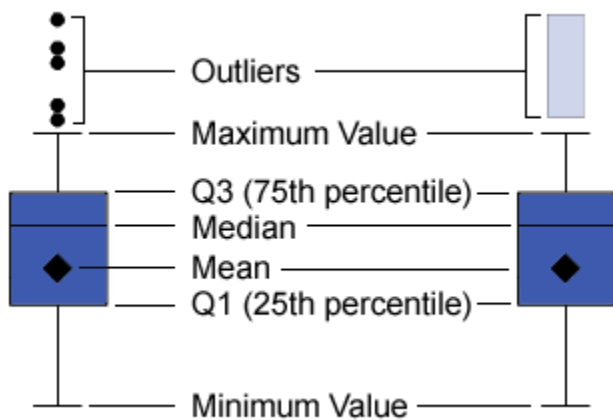A box plot displays the distribution of data values by using a rectangular box and lines called "whiskers."

Figure 24.3 *Parts of a Box Plot*

[Figure 24.3 on page 226](#) shows a diagram of a box plot. The bottom and top edges of the box indicate the interquartile range (IQR). That is, the range of values that are between the first and third quartiles (the 25th and 75th percentiles). The marker inside the box indicates the mean value. The line inside the box indicates the median value.

You can enable outliers, which are data points whose distances from the interquartile range are greater than 1.5 times the size of the interquartile range.

The whiskers (lines protruding from the box) indicate the range of values that are outside of the interquartile range. If you do not enable outliers, then the whiskers extend to the maximum and minimum values in the plot. If you enable outliers, then the whiskers indicate the range of values that are outside of the interquartile range, but are close enough not to be considered outliers.

If there are a large number of outliers, then the range of outlier values is represented by a bar. The data tip for the bar displays additional information about the outliers. To explore the outliers, double-click on the outlier bar to view the values as a new histogram visualization.

Data Roles for a Box Plot

The basic data roles for a box plot are categories and measures. You can assign one category only, and the category values are plotted on the category axis. You can assign

many measures, and the measure values are plotted on the response axis. At least one measure is required.

In addition to the basic data roles, you can assign these roles:

Lattice columns
creates a lattice of charts with a column for each value of the category data item that you assign.

Lattice rows
creates a lattice of charts with a row for each value of the category data item that you assign.

Specify Properties for a Box Plot

On the **Properties** tab, you can specify the following options:

Name
specifies the name of the visualization.

Title
specifies the title that appears above the graph.

> **Note:** The **Title** option is disabled if you select **Generate graph title**.

Generate graph title
specifies that the graph title is generated automatically based on the data items in the visualization.

Show grid lines
displays grid lines for each tick on the response axis.

Box direction
specifies whether the boxes are vertical or horizontal.

Overview
specifies whether the chart overview is enabled.

Measure layout

specifies whether the measures share a single response axis (**Shared axis**) or have separate response axes for each measure (**Separate axes**).

Note: The **Measure layout** option is disabled if your visualization contains exactly one measure.

Show averages

displays the mean value as a marker inside the box.

Show outliers

shows outliers outside of the whiskers. Outliers are data points whose distances from the interquartile range are greater than 1.5 times the size of the interquartile range.

Rotate axis labels

displays the category labels at an angle.

Note: The **Rotate axis labels** option has no effect if the box plot contains no categories.

Reference Lines

adds reference lines to the visualization. To create a reference line, click , and then specify the parameters for the line in the New Reference Line window.

The name of each reference line appears below the **Reference Lines** option. You can edit or delete your reference lines by using the icons next to each name.

Working with Heat Maps

About Heat Maps

A heat map displays the distribution of values for two data items by using a table with colored cells. If you do not assign a measure to the **Color** data role, then a cell's color represents the frequency of each intersection of values. If you assign a measure to the **Color** data role, then a cell's color represents the aggregated measure value for each intersection of values.

Data Roles for a Heat Map

The basic data roles for a heat map are:

X axis
specifies the data item that is assigned to the X axis.

Y axis
specifies the data item that is assigned to the Y axis.

Color
specifies a measure that determines the cell color. If you do not assign the **Color** role, then the cell color indicates frequency.

Specify Properties for a Heat Map

On the **Properties** tab, you can specify the following options:

Name
specifies the name of the visualization.

Title
specifies the title that appears above the graph.

   Note: The **Title** option is disabled if you select **Generate graph title**.

Generate graph title
specifies that the graph title is generated automatically based on the data items in the visualization.

Bin count
specifies the number of value ranges that are represented as cells. **Bin count** affects only measures.

Show borders
specifies that the borders between cells are visible.

Rotate axis labels
displays the category labels at an angle.

> **Note:** The **Rotate axis labels** option affects only the values on the X axis.

> **Note:** The **Rotate axis labels** option has no effect if the heat map contains no categories.

Color gradient
selects the gradient colors for the visualization.

You can click ![icon] to select the values that are used to assign the colors. See "Specify a Custom Data Range" on page 197.

Fit line
adds a fit line to the heat map. For information about the fit types that are available, see "Fit Lines" on page 279.

> **Note:** Fit lines are not available if a category is assigned to the heat map.

Applying Data Analysis

About Data Analysis

For heat maps, you can apply the following data analyses:

Correlation
identifies the degree of statistical correlation between the variables in the visualization. For more information, see "Correlation" on page 278.

Fit Line
plots a model of the relationship between the variables in the visualization.

There are many types of fit lines, including linear fit, quadratic fit, cubic fit, and penalized B-spline. For more information, see "Fit Lines" on page 279.

Correlation is applied to your visualization automatically when you add a linear fit line. It is not available with other fit types.

Enable Data Analysis

To add a fit line to your visualization, select the ![icon] drop-down list from the visualization toolbar, and then select **Fit Line ▶ [*fit-type*]**. For information about the fit types that are available, see "Fit Lines" on page 279.

Note: Fit lines are not available if your heat map contains categories or hierarchies.

Working with Geo Maps

About Geo Maps

A geo map overlays your data on a geographic map. You can display your data either as bubbles, as a scatter plot, or as colored regions on the geo map.

To display a geo map, you must define one or more of the categories as geography data items. For more information, see "Define a Geography Data Item" on page 170.

Data Roles for a Geo Map

The basic data roles for a bubble plot are:

Geography
specifies the geography data item that identifies geographic regions for your map.

Geography data items are identified by the ⊕ icon.

Bubble size
for the **Bubbles** map style, specifies the measure that determines the bubble size.

Color
for the **Regions** map style, specifies the measure that determines the regions' colors.

for the **Bubbles** map style, specifies a measure that determines the bubble color.

Specify Properties for a Geo Map

On the **Properties** tab, you can specify the following options:

Name
specifies the name of the visualization.

Title

specifies the title that appears above the graph.

Note: The **Title** option is disabled if you select **Generate graph title**.

Generate graph title

specifies that the graph title is generated automatically based on the data items in the visualization.

Map service

specifies the source for the background map.

Size scale

specifies the type of scaling that is used to draw the bubbles. Select one of the following values:

Linear

specifies that the bubble sizes are scaled relative to the minimum and maximum values (or frequencies) in your data. A negative value is displayed as a smaller bubble than a positive value.

For the **Linear** scaling type, the difference in bubble sizes might not be proportional to the difference in values.

Magnitude

specifies that the bubble sizes are scaled relative to zero and the greatest absolute value in your data. A negative value is displayed as a circle with a waving edge.

For the **Magnitude** scaling type, the difference in bubble sizes is proportional to the difference in absolute values.

Note: This option is available only for the **Bubbles** map style.

Frequency

specifies whether the frequency values are displayed as the number of values (**Count**) or as the percentage of values (**Percent**).

Map style

specifies the type of data overlay for the map. Select one of the following values:

Coordinates

displays your data as a simple scatter plot on the map. Each point is located at the center of a geographic region or at the coordinates of a location.

Bubbles

displays your data as a series of bubbles. Each bubble is located at the center of a geographic region or at the coordinates of a location.

Regions

displays your data as colored regions on the map.

Note: The **Regions** map style is not available for custom geographic roles.

Show map navigation control

specifies whether the map zoom and pan controls are visible.

Color gradient

selects the gradient colors for the visualization.

You can click to select the values that are used to assign the colors. See "Specify a Custom Data Range" on page 197.

Transparency

specifies the amount of transparency for the data overlay.

Zoom a Geo Map

You can zoom the map by using any of the following controls:

- click the zoom bar to select your zoom level

- press the + and – buttons on the zoom bar

- scroll the mouse wheel to zoom in or zoom out at the location of the cursor

Pan (Scroll) a Geo Map

You can pan (scroll) the map by using any of the following controls:

- click and drag the map

- click the arrows on the pan control

Working with Treemaps

About Treemaps

A treemap displays a hierarchy or a category as a set of rectangular tiles. Each tile represents a category value or a hierarchy node. The size of each tile represents either the frequency count or the value of a measure. If you assign a measure to the **Color** role, then the color of each tile represents the value of that measure.

Data Roles for a Treemap

The basic data roles for a treemap are:

Tile

specifies categories or a hierarchy that are used to create the tiles in the treemap. If you specify categories for the **Tile** role, the order of the categories determines the level of each category. Drag and drop the categories to place them in the order that you want.

Size

specifies a measure that determines the size of each tile. If you do not specify the **Size** role, then the tile size is determined by the frequency count.

Note: If any of the aggregated values for the **Size** role results in a negative size value or a value of zero, then an error appears.

Color

specifies a measure that determines the color of the tiles.

Specify Properties for a Treemap

On the **Properties** tab, you can specify the following options:

Name

specifies the name of the visualization.

Title

specifies the title that appears above the graph.

Note: The **Title** option is disabled if you select **Generate graph title**.

Generate graph title

specifies that the graph title is generated automatically based on the data items in the visualization.

Show data labels

shows a text label for each tile in the treemap.

Frequency

specifies whether the frequency values are displayed as the number of values (**Count**) or as the percentage of values (**Percent**).

Arrangement

specifies the layout of the tiles in the treemap. Select one of the following values:

Standard

arranges the tiles into squares, with the largest tiles generally at the bottom left.

Flow

arranges the tiles from largest to smallest, with the largest tile at the top left.

Toggle

arranges the tiles into a single row or column, with the largest tile on the left or at the top.

The orientation of the tiles alternates between hierarchy levels. The top level is arranged as a row, the second level is a column, and so on.

Additional levels

specifies the number of levels that are displayed beneath the current level.

Color gradient

selects the gradient colors for the visualization.

You can click ⬈ to select the values that are used to assign the colors. See .

Create a Hierarchy from a Treemap

If your treemap contains categories in the **Tile** role, then you can create a new hierarchy by using the categories.

On the **Roles** tab, click on the drop-down list for the **Tile** role, and then select **Create Hierarchy**. A new hierarchy is created with the name of the first category in the hierarchy.

Working with Correlation Matrices

About Correlation Matrices

A correlation matrix displays the degree of correlation between multiple intersections of measures as a matrix of rectangular cells. Each cell in the matrix represents the intersection of two measures, and the color of the cell indicates the degree of correlation between those two measures.

A correlation matrix can either compare within a single set of measures or it can compare between two sets of measures.

The correlation values are calculated by using Pearson's product-moment correlation coefficient. Correlation values are identified as weak, moderate, or strong as follows:

Weak
    0.3 or lower

Moderate
    greater than 0.3 and less than or equal to 0.6

Strong
    greater than 0.6

Data Roles for a Correlation Matrix

The basic data role for a correlation matrix is a measure. You must assign at least two measures.

Note: The maximum number of measures is 60.

The **Show Correlations** option specifies whether the correlation matrix uses a single set of measures (**Within one set of measures**) or two sets of measures (**Between two sets of measures**.)

If you select **Between two sets of measures**, then assign measures to the **X axis** and **Y axis** roles.

Specify Properties for a Correlation Matrix

On the **Properties** tab, you can specify the following options:

Name
   specifies the name of the visualization.

Title
   specifies the title that appears above the graph.

   Note: The **Title** option is disabled if you select **Generate graph title**.

Generate graph title
   specifies that the graph title is generated automatically based on the data items in the visualization.

Show borders
   specifies that the borders between cells are visible.

Rotate axis labels
   displays the axis labels at an angle.

Color gradient
   selects the gradient colors for the visualization.

Sort Correlation Values

For a correlation matrix between two sets of measures, you can sort by the correlation values for a measure.

To apply sorting, right-click the **Measures** button on the axis that you want to sort, and then select **Sort**. You can select any of the measures on the selected axis, and you can select whether to sort the correlation values in ascending or descending order.

Working with Decision Trees

About Decision Trees

A decision tree uses the values of one or more predictor data items to predict the values of a target data item. A decision tree displays a series of nodes as a tree, where the top node is the target data item, and each branch of the tree represents a split in the values of a predictor data item. Decision trees are also known as classification and regression trees.

Figure 24.4 *Example Decision Tree*

Each branch of the tree displays the name of the predictor for the branch at the top of the split. The thickness of the branch indicates the number of values that are associated with each node. The predictor values for each node are displayed above the node.

Each node in the tree displays the data for the node either as a histogram (if the target contains continuous data) or as a bar chart (if the target contains discrete data). The histogram or bar chart in each node displays the values of the target data item that are selected by the splits in the tree. The number at the top right of the node indicates the greatest value for the bar chart or histogram. At the bottom of each node, the total number of data values (count) is displayed.

Decision trees in SAS Visual Analytics use a modified version of the C4.5 algorithm.

The details table for a decision tree contains two additional data columns, Node ID and Parent ID. Node ID specifies a unique value for each node in the tree. Parent ID specifies the ID of the parent node.

Data Roles for a Decision Tree

The basic data roles for a decision tree are:

Target
specifies the target for the decision tree. You can specify any category or measure. The decision tree attempts to predict the values of the target data item. The bar chart or histogram inside each node of the tree displays the frequency of values for the target data item.

Predictors
specifies predictors for the decision tree. You can specify one or more categories or measures as predictors. The values of predictor data items are displayed above the nodes in the tree. The order of the data items in the **Predictors** list does not affect the tree.

Note: If a predictor does not contribute to the predictive accuracy of the tree, then it is not used in the tree.

Specify Properties for a Decision Tree

On the **Properties** tab, you can specify the following options:

Name
specifies the name of the visualization.

Title
specifies the title that appears above the graph.

Note: The **Title** option is disabled if you select **Generate graph title**.

Generate graph title
specifies that the graph title is generated automatically based on the data items in the visualization.

Growth strategy

specifies the parameters that are used to create the decision tree. Select one of the following values:

Basic

specifies a simple tree with a maximum of two branches per split and a maximum of six levels. For details, see Table 24.4 on page 242.

Advanced

specifies a complex tree with a maximum of four branches per split and a maximum of six levels. For details, see Table 24.4 on page 242.

Expert

enables you to select the values for each of the expert parameters.

If you select **Expert** as the value for **Growth strategy**, then the following additional options appear:

Maximum branches

specifies the maximum number of branches for each node split.

Maximum levels

specifies the maximum number of levels in the tree.

Leaf size

specifies the minimum number of values (count) for each node.

Target bins

specifies the number of bins that are used for the target data item.

Note: This option has no effect if the target data item contains discrete data.

Predictor bins

specifies the number of bins that are used for predictor data items.

Note: This option has no effect if the predictor data items contain discrete data.

Pruning

specifies the level of pruning that is applied to the tree. Pruning removes leaves and branches that contribute the least to the predictive accuracy of the tree. A more **Lenient** pruning value specifies that fewer leaves and branches are removed from

the tree. A more **Aggressive** pruning value specifies that more leaves and branches are removed from the tree.

Include missing
specifies whether missing values are included in the tree.

Reuse predictors
specifies that predictors can be used more than once in the tree.

The following parameter values are used for the **Basic** and **Advanced** growth strategies:

Table 24.4   Parameter Values for the Basic and Advanced Growth Strategies

Property	Basic Value	Advanced Value
Maximum branches	2	4
Maximum levels	6	6
Leaf size	1	1
Target bins	10	10
Predictor bins	2	2
Include missing	No	Yes
Reuse predictors	No	Yes

Explore a Node as a New Visualization

For each node in the tree, you can explore the data values as a new histogram or bar chart visualization. To create a new visualization for a node, right-click on the node, and then select **Create Visualization from Node**. A new visualization appears.

Note: If you create a visualization from a node that displays a histogram, then the new histogram might have visual differences from the histogram that is displayed in the

decision tree. For example, the bins and the extent of data on the X axis might be different. However, the two histograms use exactly the same data.

Calculate a Data Item from a Decision Tree

You can create a calculated data item to represent the results of a decision tree. The calculated data item creates values that correspond to the node IDs that can be seen in the details table for the decision tree.

You can use the calculated data item in filters to select the values for a decision tree node in other types of visualizations.

To calculate a data item from a decision tree:

1   Select the ▼ drop-down list from the visualization toolbar, and then select **Create Calculated Item**.

2   In the New Calculated Item window, enter a **Name** for the new calculated item.

3   Click **OK** to create the new item.

Display the Overview

For large decision trees, the overview enables you to select the portions of the tree that are visible.

To display the overview, select the ▼ drop-down list from the visualization toolbar, and then select **Show Overview**.

Zoom a Decision Tree

You can zoom a decision tree by using either of the following controls:

▪   in the overview, select the part of the decision tree that you want to view

▪   scroll the mouse wheel over the visualization to zoom in or out at the location of the cursor

Pan (Scroll) a Decision Tree

You can pan (scroll) the decision tree by using any of the following controls:

- in the overview, drag the selection box

- hold down the Shift key and drag the decision tree

Working with Word Clouds

About Word Clouds

A word cloud displays a set of words from a character data item. Depending on the type of word cloud and your data roles, the size of each word in the cloud can indicate the relevance of the word, the frequency of the word, or the value of a measure.

You can create two types of word cloud:

Word clouds that use text analytics
Word clouds that use text analytics analyze each value in a document collection data item as a text document that can contain multiple words. Words that often appear together in the document collection are identified as topics. The word cloud displays the most relevant terms for each topic, where the size of each term indicates the relevance.

To enable text analytics, you must set a unique row identifier and define one or more categories as document collections. See "Managing Data" on page 151.

Note: Depending on the number of rows in your data source and the length of the values in your document collection, a word cloud with text analytics might require a significant amount of time to display.

Note: Text analytics in SAS Visual Analytics uses a different algorithm from SAS Text Miner. Your results might be different from the results that SAS Text Miner produces.

Word clouds that use category values
> Word clouds that use category values analyze each value in a category data item as a single text string. The word cloud can display either the string values that have the highest frequency or the string values that have the greatest value for a measure. The color of each word can indicate the value of a measure.

Data Roles for a Word Cloud

About Data Roles for a Word Cloud

The data roles for a word cloud are dependent on the type of word cloud that you select.

The **Show Word Cloud** option selects whether the word cloud is generated by using text analytics or by using category values.

Data Roles for a Word Cloud That Uses Text Analytics

For a word cloud that uses text analytics, the basic role is a **Document collection**. A document collection is a category data item that contains the words that you will analyze.

Note: To enable text analytics, you must set a unique row identifier and define one or more categories as document collections. See "Define Data Items for Text Analytics" on page 173.

In addition to the basic role, you can specify **Document details**. Document details adds additional data items to the **Documents** tab of the details table.

Data Roles for a Word Cloud That Uses Category Values

For a word cloud that uses category values, the basic role is **Words**. Specify a category whose values are used in the word cloud.

In addition to the basic role, you can specify these roles:

Size
> specifies a measure that determines the size of each word. If you do not specify a measure, then the word size indicates the frequency of each word.

Color
specifies a measure that determines the color of each word.

Specify Properties for a Word Cloud

On the **Properties** tab, you can specify the following options:

Name
specifies the name of the visualization.

Title
specifies the title that appears above the graph.

> **Note:** The **Title** option is disabled if you select **Generate graph title**.

Generate graph title
specifies that the graph title is generated automatically based on the data items in the visualization.

Frequency (for category values only)
specifies whether the frequency is displayed as a count (**Count**) or as a percentage (**Percent**).

> **Note:** The frequency values are based on the data that is shown in the visualization (after filters and other data selections have been applied).

> **Note:** This option has no effect if a measure is assigned to the **Size** role.

Word display limit
specifies the maximum number of words that are displayed in the word cloud.

Font scale
specifies the amount of difference in font sizes between the largest and smallest words in the cloud. The number value specifies the ratio in points of the largest font size to the smallest font size.

For word clouds that use text analytics, you can specify the following additional options:

Maximum topics
specifies the maximum number of topics to create. Specify a number from 4 to 20.

Resolution

specifies the resolution that is used to identify topics. A **Low** resolution identifies fewer topics. A **High** resolution identifies more topics.

Cell weight

specifies whether to weight the frequency of each term for every document that it appears in. Selecting **Logarithmic** de-emphasizes terms that appear many times in relatively few documents.

Term weight

specifies a weighting algorithm for the terms in the document collection. The **Entropy** weighting algorithm emphasizes terms that have a low frequency across the document collection.

Entity extraction

specifies a method that is used to identify text entities such as names, addresses, telephone numbers, and so on. The **Standard** method identifies each text entity as a term. If you select **None**, then text entities are not treated differently from other text.

Minimum term frequency

specifies the minimum number of documents that a term must appear in. Specify a number from 1 to 20. If a term does not appear in the minimum number of documents, then it is not included in the word cloud.

Topic label term count

specifies the number of terms that are included in a topic name. Specify a number from 2 to 8. This property does not affect the number of terms that are used to select topics; only the topic names are changed.

Extract noun groups

specifies whether to identify groups of nouns as terms.

Stem words

specifies whether all forms of a given word are identified as a single term. For example, if you select **Stem words**, then the words "sell," "sells," "selling," and "sold" are identified as a single term "sell."

Use stop list (if available)

specifies whether to use a stop list to exclude common words such as "the," "with," and "is" when identifying terms. If no stop list is available, then a message appears at the bottom of the word cloud.

25

Working with Filters

About Filters in SAS Visual Analytics Explorer

In SAS Visual Analytics Explorer (the explorer), you can create filters to subset your data.

Three types of filters are available:

Basic filters
>   subset the data for your visualizations by using a single data item. Basic filters can be modified by using visual filter controls or by using the Edit Filter window.
>
>   For more information, see "Working with Basic Filters" on page 251.

Advanced filters
>   subset the data for your visualizations by using any number of data items. Visual filter controls are not available for advanced filters.
>
>   For more information, see "Working with Advanced Filters" on page 257.

Data source filters
>   subset the data for the entire exploration. Data source filters can use any number of data items to select data.

For information about data source filters, see "Working with Data Source Filters" on page 259.

All of your filters are saved when you save your exploration.

Working with Basic Filters

About Basic Filters

For all visualization types, you can subset your data by using the **Filters** tab. You can base your filters on any data item, regardless of whether the data item is assigned to the current visualization.

Your filters can be either *local filters* that apply only to the current visualization or *global filters* that apply to all of the visualizations in the current exploration. For more information, see "Set the Scope of a Basic Filter (Global or Local)" on page 256.

The total percentage of values that are selected by your filters is displayed at the bottom of the **Filters** tab. The tooltip for the percentage displays the exact number of data rows. The percentage is calculated after data source filters are applied.

Create a Basic Filter

To create a basic filter:

1 From the **Data** pane, select the data item that you want to use as the base of the filter. You can select any data item, regardless of whether it is assigned to the current visualization.

2 For a local filter, either right-click on the data item, and select **Add as Local Filter**, or drag and drop the data item onto the **Local Filters** area of the **Filters** tab.

For a global filter, either right-click on the data item, and select **Add as Global Filter**, or drag and drop the data item onto the **Global Filters** area of the **Filters** tab.

Note: If a local filter exists for a data item in any of your visualizations, then you cannot create a new global filter for that data item. Remove the local filter, or change

its scope to global. Similarly, if a global filter exists for a data item, then you cannot create a local filter for that data item.

3   Set the parameters for your filter. You can either create a basic filter by using visual filter controls , or you can create an advanced filter by using the Edit Filter window .

Note: By default, your filter changes are applied automatically to the active visualization. To apply multiple changes together, deselect **Auto**, and then click **Apply** when you are ready to apply your filter changes.

Create a Basic Filter from a Data Selection

When you select data values in a visualization, you can use the selected values to create a basic filter.

To create a new filter from a data selection:

1   Select one or more data values in a visualization.

2   Right-click on the visualization, and then select one of the following:

Include Only Selection
    creates a filter that includes the selected values only.

Exclude Selection
    creates a filter that excludes the selected values.

The new filter appears on the **Filters** tab.

Edit a Filter By Using Visual Filter Controls

About Visual Filter Controls

Visual filter controls are accessible directly from the **Filters** tab. The visual filter controls are different, depending on the model type of the filter data item.

Filtering Discrete Data

Display 25.1 *A Visual Filter for a Discrete Data Item*
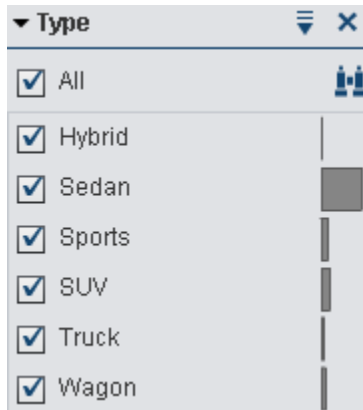
Note: If your data contains a large number of discrete values, then the top 1,000 most frequent values are displayed.

For data items that use the discrete data model, the filter displays all of the distinct values for the data item. To the right of each value, a bar indicates the frequency.

Use the check boxes next to each value to select the values for the filter. You can search for a value by clicking ▮▮. Enter a text string in the field. The search locates values that begin with your text string.

Select **All** to select or deselect all of the values. If the filter displays the top 1,000 most frequent values, then select **Include unlisted values** to select the unlisted values.

To invert your selection, click the ▼ drop-down list, and then select **Invert Selection**.

You can sort the values by frequency or alphabetically by value. Click the ▼ drop-down list, and then select **Sort by Frequency** or **Sort by Values**.

To display the selected values at the top of the list, click the ▼ drop-down list, and then select **Show Selected at Top**.

To exclude data observations in which the filter category has a missing value, deselect **Include missing values**.

Note: This option is available only if your data contains missing values.

By default, your filter changes are applied automatically to the active visualization. To apply multiple changes together, deselect **Auto** at the bottom of the **Filters** tab, and then click **Apply** when you are ready to apply your filter changes.

Filtering Continuous Data

Display 25.2   A Visual Filter for a Continuous Data Item
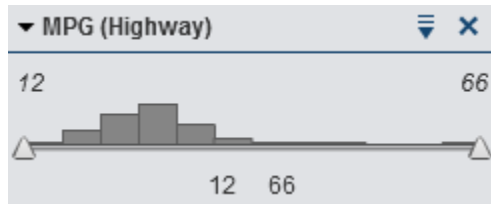
For measures that use the continuous data model, the filter shows the range of data values as a histogram. Sliders at each side of the histogram enable you to select the data that is selected by the filter.

The visual filter can select data by using different filter conditions. By default, the **Between (Inclusive)** condition is used. To change the condition for your filter, click the ▼ drop-down list, and then select **Condition Type ▶ type**. Select one of the following conditions:

Between (Exclusive)
specifies that a matching value must be in the range between the lower bound and the upper bound. Matching values do not include the bounding values.

Between (Inclusive)
specifies that a matching value must be in the range between the lower bound and the upper bound. Matching values can include the bounding values.

Equal To
specifies that a matching value must be equal to the value specified by the filter.

Greater Than
specifies that a matching value must be greater than the value specified by the filter.

Greater Than or Equal To
specifies that a matching value must be greater than or equal to the value specified by the filter.

Less Than
specifies that a matching value must be less than the value specified by the filter.

Less Than or Equal To
specifies that a matching value must be less than or equal to the value specified by the filter.

Not Between (Exclusive)
specifies that a matching value must be outside the range between the lower bound and the upper bound. Matching values do not include the bounding values.

Not Between (Inclusive)
specifies that a matching value must be outside the range between the lower bound and the upper bound. Matching values can include the bounding values.

Not Equal To
specifies that a matching value must not be equal to the value specified by the filter.

You can specify the values for the filter by using sliders or by entering the values explicitly:

- For numeric data, click on a value below the histogram, and then enter a new value in the text field.

- For a continuous datetime data item, click 🗓, and then select a date or time.

To exclude data observations in which the filter measure has a missing value, deselect **Include missing values**.

Note: This option is available only if your data contains missing values.

To invert your selection, click the ⬇ drop-down list, and then select **Invert Selection**.

By default, your filter changes are applied automatically to the active visualization. To apply multiple changes together, deselect **Auto** at the bottom of the **Filters** tab, and then click **Apply** when you are ready to apply your filter changes.

Edit a Basic Filter By Using the Edit Filter Window

To edit a basic filter by using the Edit Filter window:

1 Select the **Filters** tab.

2 Click the ⬇ drop-down list for the filter that you want to edit, and then select **Edit Filter**. The Edit Filter window appears.

3 Edit the expression for the filter. For more information, see "Editing a Filter Expression" on page 261.

Set the Scope of a Basic Filter (Global or Local)

A filter can be either local (applies only to the current visualization) or global (applies to all of the visualizations in the exploration).

To set the scope of a filter, click the ⬇ drop-down list, and then select **Global Filter** or **Local Filter**.

Note: If you change a global filter's scope to local, then a local filter is created for all of the visualizations in the exploration.

Note: If local filters exist in your other visualizations for the same data item, then converting your filter to a global filter will replace those local filters. A confirmation window enables you to continue or cancel.

Delete a Basic Filter

To delete a basic filter, click ✕ beside the filter on the **Filters** tab, or click the ⬇ drop-down list on the filter menu, and then select **Delete Filter**.

You can delete all of your filters by clicking the ⌄ drop-down list in the **Global Filters** area or in the **Local Filters** area of the **Filters** tab, and then selecting **Delete All Filters**.

Note: If you delete a global filter, then the filter is removed from all of your visualizations in the exploration.

Reset a Basic Filter

To restore a basic filter to its initial state, click the ⌄ drop-down list for the filter, and then select **Reset Filter**.

You can reset all of your filters by clicking the ⌄ drop-down list in the **Global Filters** area or in the **Local Filters** area of the **Filters** tab, and then selecting **Reset All Filters**.

Working with Advanced Filters

About Advanced Filters

For all visualization types, you can create advanced filters to subset your data by using the **Filters** tab.

Advanced filters enable you to create filters that use more than one data item.

Your advanced filters can be either *local filters* that apply only to the current visualization, or *global filters* that apply to all of the visualizations in the current exploration.

The total percentage of values that are selected by your filters is displayed at the bottom of the **Filters** tab. The tooltip for the percentage displays the exact number of data rows. The percentage is calculated after data source filters are applied.

Create an Advanced Filter

To create an advanced filter:

1   Click the ▼ drop-down list in the **Global Filters** area or in the **Local Filters** area of the **Filters** tab, and then select **New Advanced Filter**. The Edit Filter window appears.

2   Specify a **Filter name**. The filter name identifies the advanced filter on the **Filters** tab.

3   Build the expression for your filter by dragging and dropping data items, conditions, and operators onto the expression.

For information about the conditions and operators that are available, see Appendix 4, "Conditions for Filters," on page 599 and Appendix 3, "Operators for Data Expressions," on page 569 .

You can also build your expression as text by using the **Text** tab. See "Editing a Data Expression in Text Mode" on page 565.

4   Click **OK** to apply the filter.

Edit an Advanced Filter

To edit an advanced filter:

1   Select the **Filters** tab.

2   Click the ▼ drop-down list for the filter that you want to edit, and then select **Edit Filter**. The Edit Filter window appears.

3   Edit the expression for the filter. For more information, see "Editing a Filter Expression" on page 261.

Set the Scope of an Advanced Filter (Global or Local)

A filter can be either local (applies only to the current visualization) or global (applies to all of the visualizations in the current exploration).

To set the scope of a filter, click the ⩔ drop-down list, and then select **Global Filter** or **Local Filter**.

Note: If you change a global filter's scope to local, then a local filter is created for all of the visualizations in the exploration.

Note: If local filters exist in your other visualizations for the same data item, then converting your filter to a global filter will replace those local filters. A confirmation window enables you to continue or cancel.

Delete an Advanced Filter

To delete an advanced filter, click ✖ beside the filter on the **Filters** tab, or click the ⩔ drop-down list on the filter menu, and then select **Delete Filter**.

You can delete all of your filters by clicking the ⩔ drop-down list in the **Global Filters** area or in the **Local Filters** area of the **Filters** tab, and then selecting **Delete All Filters**.

Note: If you delete a global filter, then the filter is removed from all of your visualizations in the exploration.

Working with Data Source Filters

About Data Source Filters

A data source filter enables you to subset the data in an entire exploration. A data source filter differs from a global advanced filter in the following ways:

■   A data source filter updates the cardinality values that appear in the **Data** pane.

■   A data source filter updates the total number of rows that is used in the filtered rows percentage at the bottom of the right pane.

■   A data source filter is not shown on the **Filters** tab.

■ An exploration can contain only one data source filter.

Create a Data Source Filter from a Global Filter

To convert an existing global filter to a data source filter:

1 On the **Filters** tab, select the global filter that you want to convert.

2 Click the ▼ drop-down list, and select **Convert to Data Source Filter**. The global filter disappears from the **Filters** tab.

Note: If your exploration already contains a data source filter, then the global filter is appended to the filter expression using the AND operator.

You can convert all of your global filters together into a single data source filter by clicking the ▼ drop-down list in the **Global Filters** area of the **Filters** tab, and then selecting **Convert All to Data Source Filter**. The expressions from the global filters are joined with AND operators.

To view the data source filter settings, select **Data** ▶ **Data Source Details** from the main menu. The data source filter settings are displayed in the **Data source filter** field.

Create a New Data Source Filter

To create a data source filter:

1 Select **Data** ▶ **Data Source Details** from the main menu.

2 Click **New** to create a data source filter.

   Note: If the current exploration contains an existing data source filter, then click **Edit** to add conditions to the existing data source filter.

   The Edit Filter window appears.

3 Build the expression for your filter by dragging and dropping data items, conditions, and operators onto the expression.

For information about the conditions and operators that are available, see Appendix 4, "Conditions for Filters," on page 599 and Appendix 3, "Operators for Data Expressions," on page 569 .

You can also build your expression as text by using the **Text** tab. See "Editing a Data Expression in Text Mode" on page 565.

4   Click **OK** to apply the filter.

Edit a Data Source Filter

To edit the data source filter for the current exploration:

1   Select **Data** ▶ **Data Source Details** from the main menu, and then click **Edit**. The Edit Filter window appears.

2   Edit the expression for the filter. For more information, see "Editing a Filter Expression" on page 261.

Delete a Data Source Filter

To delete a data source filter, select **Data** ▶ **Data Source Details**, and then click **Delete**.

Editing a Filter Expression

Specify the Filter Name

To specify a name for the filter, enter a name in the **Filter name** field. The filter name identifies the advanced filter on the **Filters** tab.

Note: You cannot specify the name for a basic filter or a data source filter.

Edit a Parameter Value

To edit a parameter value, select the parameter, and enter a new value. Or, right-click the parameter field, and select **Replace with** to select a data item.

Add a Condition

To add a new condition:

1   From the **Data items** list, select the data item on which the condition is based.

2   From the **Conditions** list, select a condition. For a list of the conditions that are available, see Appendix 4, "Conditions for Filters," on page 599.

3   Drag and drop the condition onto the expression.

4   For any required parameters, select the parameter, and enter a value, or right-click the parameter field, and select **Replace with** to select a data item.

Add an Operator

To add an operator to the expression:

1   From the **Operators** list, select the operator that you want to add. For a list of the operators that are available, see Appendix 3, "Operators for Data Expressions," on page 569.

2   Drag and drop the operator onto the expression.

3   For any required parameters, select the parameter, and enter a value. Or, right-click the parameter field, and select **Replace with** to select a data item.

Replace an Operator

To replace an operator, drag and drop a new operator onto the existing operator in the expression. You can also right-click an operator in your expression, and select **Replace Operator with** ▶ *operator*.

Delete Part of an Expression

To delete a part of an expression, highlight the part of the expression that you want to delete, right-click, and select **Delete** or **Clear**.

Subset an Expression

To subset an expression and delete the remainder:

1 Highlight the part of the expression that you want to keep.

2 Right-click, and select **Keep Operand**. The parts of the expression that were not selected are removed.

Edit an Expression as Text

On the **Text** tab of the expression editor, you can edit an expression as text code. See "Editing a Data Expression in Text Mode" on page 565.

Display the Current Expression as Text

To display the current expression as text on the **Display Text** tab, click the ⩡ drop-down, and then select **Show display text**.

Manage the Scratch Area

The scratch area of the expression editor enables you to store elements of an expression, and then add them back into the expression as needed.

You can see the contents of the scratch area on the **Scratch** tab.

To move items to the scratch area, right-click part of the expression, and select **Move to Scratch Area**. You can also drag and drop items from your expression onto the **Scratch** tab.

To move an item from the scratch area to your expression, drag and drop the item from the **Scratch** tab onto your expression.

To delete an item in the scratch area, click the 🗑 button next to the item.

26

Exporting Content

Export a Visual Exploration as a Report

You can export your visual exploration as a report, which can be refined in the designer interface or viewed directly on a mobile device or in SAS Visual Analytics Viewer (the viewer).

To export your current visual exploration as a report:

1   Select **File ▶ Export ▶ Exploration as Report**.

   Note: If you have only one visualization in your exploration, the Save As window is displayed at this point. Select the location where you want to save the report, and then click **Save** to export it as a report.

2   In the Export as Report window, select the visualizations that you want to include in the report.

Note: The following types of visualizations are not available to include in the report:

- visualizations that do not contain any data

- decision trees

- network diagrams

- word clouds

- treemaps that display additional levels

- line charts with forecasting analysis are not available if the visualization has not been displayed during the current session

- geo maps that use calculated geography data items

Note: If your scatter plot or heat map contains correlation data, the correlations are not included in the report.

The Export to Report window enables you to perform the following tasks:

Task	Action
Add visualizations to the selection list. (Restore visualizations.)	Select a visualization from the **Available** list, and then click ➡, or click ⇒ to add all of the available visualizations.
Remove visualizations from the selection list. (Minimize visualizations.)	Select a visualization from the **Selected** list, and then click ⬅, or click ⇐ to remove all of the visualizations.
Change the order of the selected visualizations.	In the **Selected** list, drag and drop the visualizations in the order that you want, or select a visualization and click the up and down arrows to move it.
Search the available visualizations.	In the **Search** field, enter a string to search for. A visualization matches the search if the string appears anywhere in the visualization name. Only the matching visualizations appear in the **Available** list.

Task	Action
Filter the available visualizations.	Click ![filter icon], and then select your filter parameters. You can filter either on the visualization type or on the data items that are used in each visualization.
	Only the matching visualizations appear in the **Available** list.

When you have finished selecting your visualizations, click **OK**.

3 Select the location where you want to save the report, and then click **Save**.

4 Click **OK** to return to the explorer, or select **Open the report and close the current exploration**, and click **OK** to view the report in the designer.

Export a Visual Exploration as a PDF

To export your current visual exploration as a PDF:

1 Select **File ▶ Export ▶ Exploration as PDF**.

2 In the **Export as PDF** wizard, enter a **Title** and a **Description** for the PDF document. Then, select the following options:

Page numbers
　　specifies that the PDF document contains page numbers.

Summary data
　　specifies whether the PDF document contains summary data tables for each visualization.

Filter descriptions
　　specifies whether each visualization includes a description of its active filters.

Click **Next**.

3   Select the visualizations that you want to include in the PDF document.

Note: The following types of visualizations are not available to include in the PDF document:

- ▪ visualizations that do not contain any data

- ▪ crosstabs

- ▪ tables that contain more than 1,000 rows

The Export as PDF window enables you to perform the following tasks:

Task	Action
Add visualizations to the selection list. (Restore visualizations.)	Select a visualization from the **Available** list, and then click ➡, or click ⟫ to add all of the available visualizations.
Remove visualizations from the selection list. (Minimize visualizations.)	Select a visualization from the **Selected** list, and then click ⬅, or click ⟪ to remove all of the visualizations.
Change the order of the selected visualizations.	In the **Selected** list, drag and drop the visualizations in the order that you want, or select a visualization and click the up and down arrows to move it.
Search the available visualizations.	In the **Search** field, enter a string to search for. A visualization matches the search if the string appears anywhere in the visualization name. Only the matching visualizations appear in the **Available** list.
Filter the available visualizations.	Click ▼, and then select your filter parameters. You can filter either on the visualization type or on the data items that are used in each visualization. Only the matching visualizations appear in the **Available** list.

Click **Next**.

4   Click **Finish** to open a download window for your browser.

5   Select the location where you want to save the PDF document.

Note: If the legend for a visualization is too large, then it is not included in the PDF document.

Save a Visualization as an Image File

To save your current visualization as an image file:

1   Select the visualization that you want to save as an image.

2   If your visualization contains hierarchies, drill down to the hierarchy levels that you want to show in the image.

3   Select the ▼ drop-down list from the visualization toolbar, and then select **Export Image**.

4   If necessary, scroll and zoom the visualization to show the data items that you want to include in the image. If your visualization contains filters, legends, or hierarchy breadcrumbs, select whether to include these elements.

5   Click **Save** to open a download window for your browser.

6   Select the location where you want to save the image.

Export Data from a Visualization

You can export the data from a visualization in comma-separated values (.csv) format. You can open .csv files in other software such as SAS Enterprise Miner, SAS Enterprise Guide, or Microsoft Excel.

To export data for a visualization:

1   Select the visualization for which you want to export data.

    Note: For crosstab visualizations, a data summary is exported instead of a full data table.

2   If your visualization contains hierarchies, drill down to the hierarchy levels that you want to export.

3   Select the ⩤ drop-down list from the visualization toolbar, and then select **Export Data** to open a download window for your browser.

    Note: For a crosstab visualization, select **Export Data Summary** instead of **Export Data**.

4   Specify a filename and select the location where you want to save the file.

    Note: To export data from a table visualization, your browser must allow pop-ups.

E-mail a Visual Exploration as a Link

To e-mail your visual exploration as a link:

1   Select **File** ▶ **E-mail**. The E-mail window appears.

2   Enter the following information:

To

specifies the e-mail destination. To enter multiple addresses, separate the addresses by using commas, semicolons, or spaces.

From

specifies the e-mail address that is shown as the sender of the e-mail. This e-mail address is used for notifications if the e-mail cannot be delivered to the addresses in the **To** field.

Subject

specifies the subject line of the e-mail.

Message

specifies the text for the e-mail message body. This field is optional.

Note: The message body of the e-mail automatically includes a link to the current visual exploration. If you specify text for the **Message** field, then your text appears before the link in the e-mail.

3   Click **OK** to send the e-mail.

27

Managing Hierarchies

What Is a Hierarchy?

A hierarchy is an arrangement of category columns that is based on parent-child relationships. The levels of a hierarchy are arranged with more general information at the top and more specific information at the bottom.

For example, you might create a hierarchy of date-time columns with Year as the top level, Month as the next level, and Day as the bottom level.

Creating hierarchies enables you to add drill-down functionality to your visualizations. For example, if you use a date-time hierarchy, you can drill down to the data for a specific year. Then, you can drill down to the data for a specific month.

When you drill down a hierarchy, a set of breadcrumb links at the top of your visualization enables you to drill back up the hierarchy.

Create a New Hierarchy

To create a new hierarchy:

1   Select **Data** ▶ **New Hierarchy**. The New Hierarchy window appears.

2   In the **Name** field, enter a name for the hierarchy.

3   Select the categories that you want to include in the hierarchy, and then click ➡ to add them to the hierarchy.

   Note: You can also drag and drop categories.

   To change the order of the categories in your hierarchy, select the category that you want to move. Then, click ⬆ to move the category up, or click ⬇ to move the category down.

   To remove a category from the hierarchy, select the category that you want to remove. Then, click ⬅.

4   Click **OK** to finish creating the hierarchy.

Note: You can create a hierarchy from within a crosstab visualization. See "Create a Hierarchy from a Crosstab" on page 212.

Derive a Hierarchy from a Date, Time, or Datetime Data Item

For date, time, and datetime data items, you can derive a hierarchy automatically. To derive a hierarchy, right-click on a date, time, or datetime data item in the **Data** pane, and then select one of the following values:

Create Date Hierarchy

creates a hierarchy with levels for year, quarter, month, and day. Depending on the format of the data item, some of the levels might not be created.

Create Date and Time Hierarchy

creates a hierarchy with levels for year, quarter, month, day, hour, minute, and second. Depending on the format of the data item, some of the levels might not be created.

Create Time Hierarchy

creates a hierarchy with levels for hour, minute, and second. Depending on the format of the data item, some of the levels might not be created.

New calculated items are created for each level in the hierarchy. By default, the new calculated items are hidden in the **Data** pane. If you delete a derived hierarchy, then a window appears. The window enables you to delete the calculated items that are associated with the hierarchy.

Create a Hierarchy from a Visualization

You can create a hierarchy from a treemap, crosstab, or network diagram visualization. See the following topics:

- "Create a Hierarchy from a Treemap" on page 236

- "Create a Hierarchy from a Crosstab" on page 212

- "Create a Hierarchy from a Network Diagram" on page 223

Edit a Hierarchy

To edit an existing hierarchy:

1 From the **Data** pane, right-click on the hierarchy that you want to edit, and then select **Edit**. The Edit Hierarchy window appears.

2   In the **Name** field, enter a name for the hierarchy.

3   To add a category to the hierarchy, select the category, and then click ➡.

> **Note:** You can also drag and drop categories.

To change the order of the categories in your hierarchy, select the category that you want to move. Then, click ⬆ to move the category up, or click ⬇ to move the category down.

To remove a category from the hierarchy, select the category that you want to remove. Then, click ⬅.

4   Click **OK** to save the changes to your hierarchy.

Delete a Hierarchy

To delete a hierarchy:

From the **Data** pane, right-click on the hierarchy that you want to delete, and then select **Delete**. Click **Yes** to confirm that you want to delete the hierarchy.

28

Performing Data Analysis

Overview of Data Analysis in SAS Visual Analytics Explorer

Types of Data Analysis

SAS Visual Analytics enables you to perform three basic types of data analysis:

Correlation
    identifies the degree of statistical relationship between measures.

Fit Line
> plots a model of the relationship between measures. There are many types of fit lines, including linear fit, quadratic fit, cubic fit, and penalized B-spline.

Forecasting
> estimates future values for your data based on statistical trends.

Correlation

Correlation identifies the degree of statistical relationship between measures. The strength of a correlation is described as a number between -1 and 1. A value that is close to -1 implies a strong negative correlation, a value that is close to 0 implies little or no correlation, and a value that is close to 1 implies a strong positive correlation.

To apply correlation to a visualization, add a linear fit line, or select the correlation matrix visualization type.

For a heat map or a simple scatter plot, the correlation is identified by a text label in the visualization legend. Select ⓘ to view additional details about the correlation, including the exact correlation value.

For a scatter plot matrix, the correlation for each plot is identified by a colored border around the plot. The visualization legend displays a key for the color values. Select ⓘ to view additional details about the correlation, including the exact correlation values for each plot.

Note: For nonlinear fit types, a scatter plot matrix displays additional plots to show each intersection of variables in two orientations. For example, if a scatter plot matrix plots the variables A, B, and C, then plots are created for both A * B and B * A when a nonlinear fit line is applied.

For a correlation matrix, the correlation for each cell is identified by the color of the cell background. The visualization legend displays a key for the color values. The data tip for each cell displays the correlation value.

Fit Lines

A fit line plots a model of the relationship between measures. You can apply fit lines to scatter plots and heat maps.

You can apply the following types of fit line to your visualization:

Best Fit
> tests the cubic, quadratic, and linear fit methods against your data, and selects the fit method that produces the best result. To see which fit method was used, select ⓘ from the visualization legend.

Linear
> creates a linear fit line from a linear regression algorithm. A linear fit line produces the straight line that best represents the relationship between two measures. For more information about the linear fit line, select ⓘ from the visualization legend.
>
> For a linear fit, correlation is automatically added to the visualization. Correlation is not available with other fit types.

Quadratic
> creates a quadratic fit line. A quadratic fit produces a line with a single curve. A quadratic fit line often produces a line with the shape of a parabola. For more information about the quadratic fit line, select ⓘ from the visualization legend.

Cubic
> creates a cubic fit line. A cubic fit line produces a line with two curves. A cubic fit line often produces a line with an "S" shape. For more information about the cubic fit line, select ⓘ from the visualization legend.

PSpline
> creates a penalized B-spline. A penalized B-spline is a smoothing spline that fits the data closely. A penalized B-spline can display a complex line with many changes in its curvature. For more information about the penalized B-spline, select ⓘ from the visualization legend.

Forecasting

Forecasting estimates future values for your data based on statistical trends. Forecasting is available only for line charts that contain date or datetime data items.

A forecast adds a line with predicted values to your visualization and a colored band that represents the 95% confidence interval. The 95% confidence interval is the data range where the forecasting model is 95% confident what the future values will be.

The explorer automatically tests multiple forecasting models against your data, and then selects the best model. To see which forecasting model was used, select ⓘ from the visualization legend.

The forecast model can be any one of the following:

- Damped-trend exponential smoothing

- Linear exponential smoothing

- Seasonal exponential smoothing

- Simple exponential smoothing

- Winters method (additive)

- Winters method (multiplicative)

Note: Forecasting accounts for cyclical patterns by using standard intervals of time (for example, 60 minutes in an hour, 24 hours in a day, and so on). If your data uses nonstandard intervals (for example, 48 30-minute cycles per day), then cyclical patterns are not considered in the forecast.

Add a Fit Line to an Existing Visualization

To add a fit line to a scatter plot or heat map, select the ⏷ drop-down list from the visualization toolbar, and then select **Fit Line** ▶ [*fit-type*]. For information about the fit types that are available, see "Fit Lines" on page 279.

Add Forecasting to an Existing Visualization

To add data analysis to a line chart, select the ⏷ drop-down list from the visualization toolbar, and then select **Show Forecast**.

Note: To apply forecasting, the line chart must contain a date or datetime item.

On the **Properties** tab for the line chart visualization, you can adjust the number of data points to predict by using the **Duration** option.

On the **Roles** tab, you can add measures to the forecast as **Underlying factors**. The forecasting model evaluates the additional measures to determine whether they contribute to the accuracy of the forecast. If the additional measures do not increase the accuracy of the forecast, then they are not used. If the additional measures do increase the accuracy of the forecast, then the forecast line is adjusted, and the confidence bands are narrowed.

If your forecast includes underlying factors, then you can apply scenario analysis to the forecast. For more information, see "Apply Scenario Analysis to a Forecast" on page 282.

Forecast Measures as a New Visualization

To create a visualization that forecasts multiple measures:

1   Select **Analysis** ▶ **Forecast**. The Forecast window appears.

2   From the **Date category** drop-down list, select the date or datetime data item for the forecast.

3   From the **Forecast** area, select the measures that you want to include in the forecast.

4   (Optional) If you selected a single measure for **Forecast**, then you can also select additional measures to use as **Underlying factors** in the forecast. The forecasting model evaluates the additional measures to determine whether they contribute to the accuracy of the forecast. If the additional measures do not increase the accuracy of the forecast, then they are not used. If the additional measures do increase the accuracy of the forecast, then the forecast line is adjusted, and the confidence bands are narrowed.

5   In the **Forecast duration** field, enter the number of data values that you want to predict for the forecast.

6   Click **OK** to create the visualization. The visualization is created as a line chart.

Apply Scenario Analysis to a Forecast

If the forecast for a line chart contains underlying factors, then you can apply scenario analysis to explore hypothetical scenarios.

To apply scenario analysis:

1 If the details table for the line chart is not shown, then select the ⬇ drop-down list from the visualization toolbar, and then select **Show details**.

2 In the details table, select the **Scenario Analysis** tab. For each of the measures that are assigned to the forecast as underlying factors, a line displays the values of the measure, and a series of points enables you to set the future values for the measure.

3 Use the check box beside each measure to choose which measures are used in the scenario analysis.

4 For each measure in the scenario analysis, set the future data values. You can set the values by using any of the following methods:

▪ drag each data point upward or downward on the line plot on the **Scenario Analysis** tab. To select the accuracy of the points, select the ⬇ icon beside the measure name, and select **Snap interval**.

▪ enter the value for each data point. Right-click on each data point, and select **Set point value**.

▪ set all of the values for the measure. Right-click on a data point, and select **Set series values**. The Change future values window enables you to set all of the values to a specific value or to adjust the future values relative to the most recent observed value of the measure.

To reset any data point to its original value, right-click on the data point, and then select **Reset point**.

To reset all of the data points for a measure, right-click on any data point for the measure, and then select **Reset entire series**.

5 When you are finished setting the future data values for the scenario, click **Update** to apply the scenario to the forecast.

The forecast is updated to show the results of the scenario. The original forecast is displayed as an additional line that is labeled **(Baseline)**.

To remove the scenario analysis from a forecast, click **Reset to forecast** on the **Scenario Analysis** tab.

29

Sharing Comments in the Explorer

Sharing Comments in the Explorer

About Comments in the Explorer

Comments enable you to share feedback with other users. You can create comments for visualizations and for explorations. Exploration comments can be viewed both in the explorer and on the SAS Visual Analytics home page.

View Comments

To display the comments for the current exploration, select **File ▶ Exploration Comments** from the main menu.

To display comments for a visualization, select the visualization and then select the **Comments** tab from the right pane.

Comments are grouped by topic. If there are many comments in a topic, then only the most recent comments are displayed. Click **Show All Comments** to view all of the comments for a topic.

To search the comments, enter a search term in the **Search within comments** field.

Create a Comment under an Existing Topic

To create a comment under an existing topic:

1   Enter your comment text in the **Respond to topic** field that is beneath the topic that you want to comment on.

2   (Optional) Click 📎 to attach a file to the comment.

3   When you are finished with the comment, click **Post**. Your comment is saved and shared immediately.

Create a Comment under a New Topic

To create a comment under a new topic:

1   Enter the name of the topic in the **Enter a topic name** field.

2   Enter your comment text in the **Enter a comment** field.

3   (Optional) Click 📎 to attach a file to the comment.

4   When you are finished with the comment, click **Post**. Your comment is saved and shared immediately.

Edit a Comment

To edit a comment, select the comment that you want to edit, and then click **Edit**.

Note: To edit other users' comments, you must belong to the **Comments:Administrator** role.

Delete a Comment

To delete a comment, select the comment that you want to delete, and then click **Delete**.

Note: To delete comments, you must belong to the **Comments:Administrator** role.

Part 4

Designing Reports

30

Overview of SAS Visual Analytics Designer

About the SAS Visual Analytics Designer

The SAS Visual Analytics Designer (the designer) enables users to easily create reports or dashboards that can be saved and viewed on either a mobile device or in the viewer. The designer is part of the SAS Visual Analytics product that enables a user with either the **SAS Visual Analytics: Analysis** role or the **SAS Visual Analytics: Administration** role to view, interact with, and create reports. Simply open an existing report and interact with the information based on your current needs. Report designers can easily point and click to query central sources of data. You can drag and drop tables, graphs, and gauges to create a well-designed report. You can also add text, images, stored processes, and controls to reports. All of this is accomplished by using

the designer, which runs in a web browser. You do not need to understand a programming language to create reports.

Report designers can easily create reports and dashboards based on data sources that have been provided by a system administrator. They can update reports that were created from visual explorations. Report designers can create reports by importing objects or visual explorations from other reports. They can also define interactions (either filtering or brushing) for report objects, and then include SAS analytical results in a single report.

Your First Look at the Designer

When you are designing reports in SAS Visual Analytics, you see the designer. In general, data-related tasks are initiated from the leftmost pane, and presentation-related tasks are initiated from the rightmost pane.

Here are the features of the designer:

Figure 30.1 *The Designer*
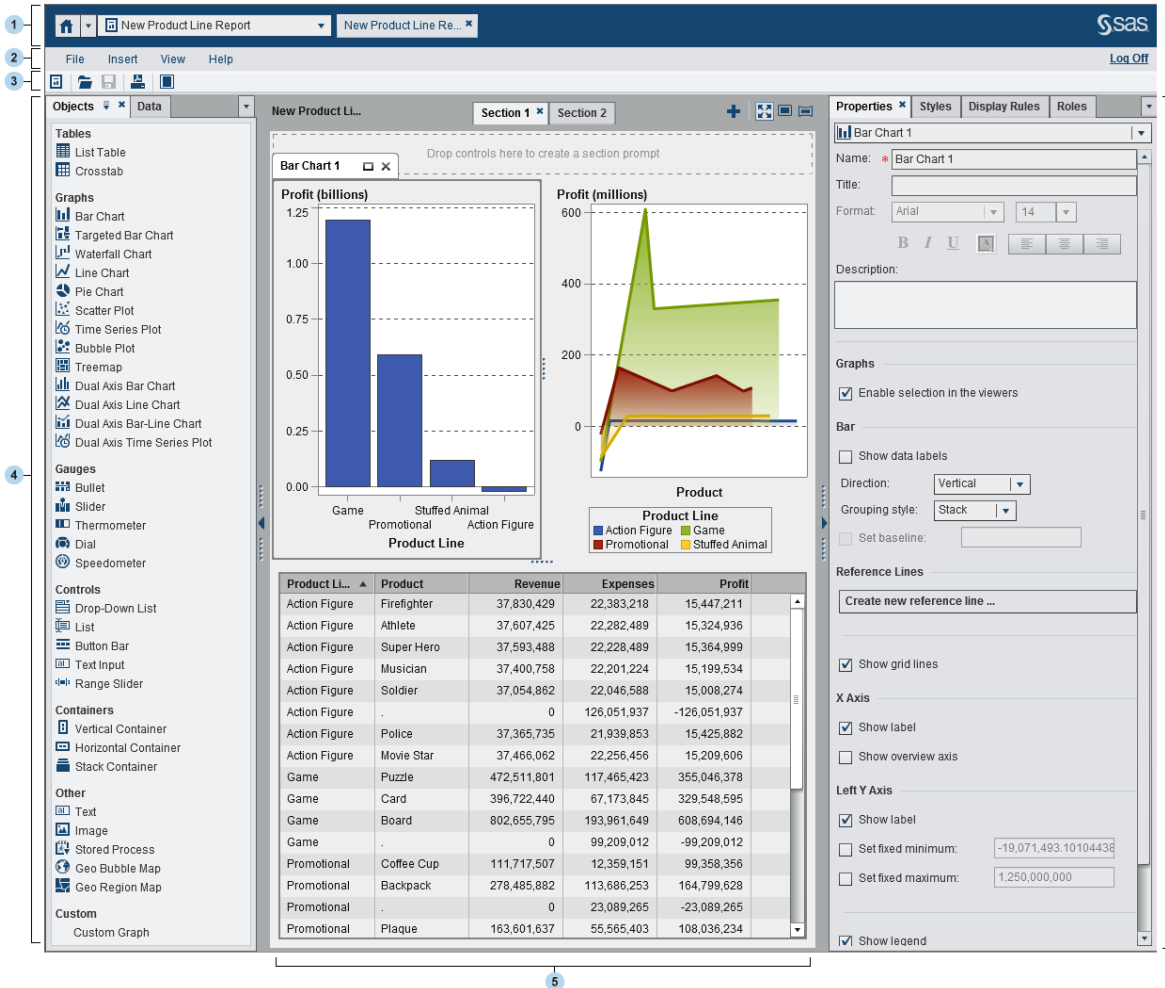
1 The application bar enables you to return to the home page and to access other parts of SAS Visual Analytics and other SAS applications that integrate with the home page. You can access your recently created or viewed reports, explorations, stored processes, data queries, or other objects in your recent history. Buttons are displayed for each open application.

2 The menu bar provides menu items that apply to the entire report or to the currently displayed report section. Actions include creating a new report, adding a new

section, inserting new report objects, adding interactions, and launching the viewer without returning to the home page. You can also log off from SAS Visual Analytics.

3   The toolbar contains icons that enable you to manage and print your reports. You can click ◼ to hide the left and right panes. Or, you can click ▥ to display the left and right panes.

4   The tabs in the left pane enable you to work with new report objects, data, and imported report objects.

5   The canvas is the workspace for building a report. The appearance of the canvas is affected by the report theme.

A report can have multiple sections, which can be accessed by using the tabs at the top of the canvas. You can change the report view by clicking the icons at the top of the canvas.

There is an area at the top of the canvas where you can drop filter controls and then categories to create section prompts.

6   The tabs in the right pane enable you to work with details about the report and report objects.

Understanding the Tabs in the Designer

Working with Tabs in the Left Pane

The **Objects**, **Data**, **Imports**, and **Shared Rules** tabs are in the left pane. The **Objects** and **Data** tabs are displayed by default. You can choose to display icons instead of names on the tabs. To display the icons on the tabs, click ▾ after the last tab name, and select **Show icons only**. From this menu, you can also choose which tabs are displayed.

The following table lists the available tabs:

Tab	Description
Objects	This tab provides a list of the tables, graphs, gauges, controls, containers, other objects, and custom objects that can be used in a report or dashboard.
Data	This tab enables you to select a data source (or data sources) and the data items for your report. You can add, refresh, import, or remove a data source using the icons above the list of data items. Using the menu, you can change the data source, define a hierarchy, define a calculated item, define an aggregated measure, show or hide data items, and filter on a data item. You can check the details for the measures in the data set. The data item table on the **Data** tab provides information about a selected data item, including the name, role, format, aggregation, and sort options. These data item properties can be modified, which impacts all of the report objects that use the data item.
Imports	This tab provides a list of reports and report objects that have been created in the designer or exported from the explorer. This enables you to create reports from multiple data sources and create full-featured reports. You can choose objects or sections in these reports to include in either a new or an existing report.
Shared Rules	This tab enables you to create a new display rule for a gauge, which is used by other gauges to designate intervals and colors for ranges. You can edit or delete an existing shared display rule. These display rules are shared across multiple gauges and can be created at any time.

Working with Tabs in the Right Pane

The **Properties**, **Styles**, **Display Rules**, and **Roles** tabs are displayed by default. You can choose to display icons instead of tab names on the tabs. To display the icons on the tabs, click ▼ after the last tab name, and select **Show icons only**. From this menu, you can also choose which tabs are displayed.

The following table lists the available tabs:

Tab	Description
Properties	This tab lists the properties for the currently selected report or report object. When a report is selected, the report's title and description are listed. When a section is selected, the section's name and layout are listed. When a report object is selected, the object's name, title, and description are listed. Information that is specific to the report object type, such as the axes and legend, is listed. If a report object is not selected on the canvas, then the properties are for the section or the report.
Styles	This tab enables you to specify the data styling, frame styling, text styling, and data colors for a selected report object. You can specify the report theme for the report.
Display Rules	This tab enables you to populate, add intervals, or color-mapped values for an object that is currently selected on the canvas. You can edit or delete an existing display rule for the selected table, graph, gauge, or control.
Roles	This tab enables you to add or update data role assignments in a selected report object that has a data source.
Alerts	This tab enables you to add, update, or delete alerts for report objects.
Comments	This tab enables you to add comments to the report after you save it.
Filters	This tab enables you to add a filter (or filters) to the selected report object.
Interactions	This tab enables you to add or update filter or brush interactions with the report objects in a section.
Ranks	This tab enables you to add rankings to report objects. For example, you might want to see the top 10 products that sold last year.

Specifying Your Preferences for the Designer

To specify preferences that are specific to the designer:

1   Select **File** ▶ **Preferences** to open the Preferences window.

2   Select **SAS Visual Analytics Designer** ▶ **General**.

3   Specify your **Report** preferences.

 ■ Specify your default view for new or existing reports. Select **Full screen**, **Tablet**, or **Wide-screen tablet**.

 ■ Select your default report theme. Select **Application theme**, **SAS Light**, **SAS Dark**, or **SAS High Contrast**.

4   Specify your **Geo Map** preferences.

 ■ Select your default **Map provider mode**. Your choices are **OpenStreetMap** and **Esri**.

 ■ If you select **Esri** as the map provider, you can specify the **Esri map service** in the Esri Map Service Selector window. Your choices are **World Imagery**, **World Street Map**, and **Light Gray Canvas**.

   Note: The **Esri** option is available only if your site has its own Esri server and if your system administrator has set the `va.SASGeomapEsriURL` configuration property in SAS Management Console.

5   Specify your **Data Tab** preference. If you are trying to optimize performance, then select the **Bypass cardinality checks** check box.

   Note: If you select the **Bypass cardinality checks** check box, then the distinct counts of data items are not displayed on the **Data** tab.

6   Click **OK** to apply your changes.

To specify global SAS preferences, see "Specify Global Preferences" on page 23. To specify general preferences for SAS Visual Analytics, see "Specify the SAS Visual Analytics General Preferences" on page 24.

Report Themes

Report themes are available in the designer. You can select a report theme for each report using the **Styles** tab. By default, the choices are **SAS Light**, **SAS Dark**, or **SAS High Contrast**. Your site might also have custom report themes. Custom report themes are automatically created when a coordinating custom application theme is created in SAS Theme Designer for Flex.

Alternatively, you can set a preference for a default report theme. For more information, see "Specifying Your Preferences for the Designer" on page 297.

Users who have the **Theme Designer for Flex:Administration** role, the **UseThemeDesignerForFlex** capability, or a custom role that has the Customize Themes capability enabled will see a **Customize Theme** button on the **Styles** tab. Click **Customize Theme** to go to SAS Theme Designer for Flex. For more information about report styles and application themes, see *SAS Theme Designer for Flex: User's Guide*.

31

Creating and Working with Reports

About Reports

You can drag and drop tables, graphs, gauges, and controls to create a well-designed report in SAS Visual Analytics Designer (the designer). You can add text, images, and other controls to reports. A report can be divided into one or more sections. (Sections

are like pages.) Each section can have a different layout and contain different report objects.

When you design a report, keep in mind that it might look slightly different on a mobile device or in the viewer. For example, the layout of the tiles in the treemap is dependent on the size of the display area. This means that the same treemap might appear slightly different in the designer than it does in the viewer or on a mobile device.

Reports that are imported from SAS Visual Analytics Explorer (the explorer) might look slightly different in the designer.

The designer provides a report view that enables you to change the size of the canvas that you use for laying out reports. For more information, see "Choosing a Report View" on page 302.

Two report layouts are available in the designer: **Precision** and **Tiled**. For more information, see "Choosing a Report Layout" on page 302.

Creating a New Report

There is no single process for creating a new report in the designer. For example, some users select their data source (or data sources) before they add the report objects while other users add report objects to the canvas and then select their data source (or data sources). Some users choose to update the properties and styles for their reports while other users choose not to. The following list of steps is one way that you can create a new report in the designer.

To create a new report:

1 Select a data source with the associated data items.

2 (Optional) Select the layout (**Precision** or **Tiled**) for the first section of the report using the **Properties** tab for Section 1.

3 Add report objects by dragging and dropping them onto the canvas or by double-clicking the report object on the **Objects** tab. You can tab over a report object, and press Enter to add the report object to the canvas.

4   Add data items to report objects by dragging and dropping one or more data items onto the report object or by right-clicking on the report object and using the context menu.

5   (Optional) Modify the report.

   ▪   Update the properties for the report and the report objects.

   ▪   Update the styles for the report objects.

   ▪   Modify data role assignments.

   ▪   Create or modify advanced data items (for example, hierarchies or geographic data items).

   ▪   Create or modify display rules and add alerts for report objects.

   ▪   Add filters to the report.

   ▪   Add interactions between the report objects in a section.

   ▪   Add a rank to a report object.

   ▪   Import report objects from the explorer.

   ▪   Add a new section to the report.

   ▪   Add comments to the report.

6   (Optional) View the report in SAS Visual Analytics Viewer (the viewer) by selecting **File** ▶ **View Report**. Once you have viewed the report, you can return to the designer by selecting **File** ▶ **Edit Report**.

7   Save the report. The default location for the first save is `My Folder`. After that, the default save location is the last accessed folder.

8   (Optional) View or change the report properties using the **File** menu.

9   (Optional) E-mail the report.

10 (Optional) Print the report.

You can create a new report based on an existing report or on existing report objects from a report that you have imported. For more information, see "Basing a New Report on Report Objects Imported from One or More Existing Reports" on page 305.

Choosing a Report View

There are three report views available in the designer: **Full Screen**, **Tablet**, or **Wide-Screen Tablet**. The default report view is **Full Screen**. You can access the report views from the **View** menu or by the icons (⛶, ▣, or ▭) on the top of the canvas.

The report layout is not automatically adjusted for optimal viewing on a device. However, the report does not necessarily display poorly. Before making the report available to users, see how the report is displayed in each view.

Choosing a Report Layout

The following report layouts are available in the designer:

Precision

enables you to place, align, and size report objects. The precision layout allows report objects to overlap, and it lets users control the depth order of these overlapping objects by sending objects backward or bringing them forward. (For example, you might want your company logo to display behind a bar chart and a pie chart in your report).

This layout option is keyboard accessible.

Note: SAS Stored Processes are not allowed in the precision layout.

Tiled

provides you a quick way to place report objects directly next to other report objects. The report objects cannot overlap. All report objects in the section are sized to fit in one screen. If you adjust the size of an object, then the other objects automatically resize to ensure that all objects continue to fill the entire screen.

Note: You can switch from the tiled layout to the precision layout after you have added report objects to your report. However, this action can change the size and the position of the report objects.

To choose your report layout:

1   Select a section tab.

2   In the right pane, click the **Properties** tab.

3   For **Layout**, select either **Precision** or **Tiled**. The **Tiled** layout is the default.
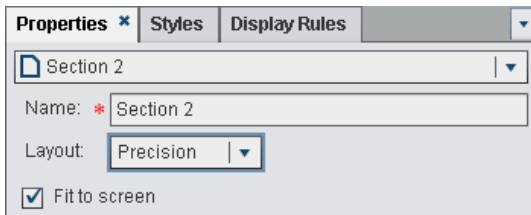
4   (Optional) If you select **Precision**, then you can specify **Fit to screen**. The **Fit to screen** option prevents objects from being sized too wide or too tall, which can cause the report viewing area for the section to scroll.

Viewing Report Properties

You can use the Report Properties window to get an overview of the report. With the exception of the **Name** field, the information in the Report Properties window is different from the information that appears on the **Properties** tab in the right pane of the designer. The Report Properties window provides information such as who created the report and when it was last modified.

All of the information in the Report Properties window is generated after a report is saved. You can add or change only the **Description** and **Keywords** for the report in the Report Properties window.

To view a report's properties:

1 Select **File ▶ Report Properties**. The Report Properties window is displayed.

2 (Optional) Update the **Description** or **Keywords** for the report.

Here is an example of the Report Properties window:

By comparison, the **Properties** tab in the right pane of the designer can be used to update the properties for individual report objects or report sections. The available properties depend on the selected report object. To move to the properties of different report objects or report sections, click ▼ after the report name, and then select another report object or report section from the list.

This example shows a report title and description.
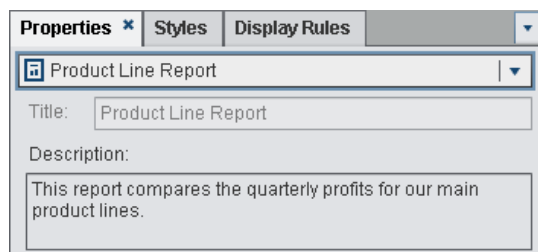
For information about updating report object properties, see "About Report Objects" on page 312.

Basing a New Report on Report Objects Imported from One or More Existing Reports

If there is an existing saved report that closely matches the report that you want to create, then you can open the existing report that contains the report objects that you want to reuse. The report objects can be from either a report that has been saved in the designer or from a report that has been exported from the explorer. (For example, you might want to reuse calculated data items or hierarchies from an existing report in your current report.) You can immediately save the report with a new name, or you can customize the existing report, and then save it.

To base a new report on report objects that have been imported from one or more existing reports:

1   Open or import the existing report by doing one of the following:

   ▪   On the home page, double-click on the report that you want to open.

   ▪   In the designer, do one of the following:

      ☐   Select **File** ▶ **Recent**, and then select the report name.

      ☐   Click ▾ beside the recent content drop-down list at the top of the window, and then select a report name.

      ☐   Click 📁 on the menu bar to navigate to a report, and then select it.

      ☐   Select **File** ▶ **Open** to display the Open window and select a report.

      ☐   Click **Imports** in the left pane. If the **Imports** tab is not visible, select **View** ▶ **Imports**.

          Select a report name from **Select a report to import**, or click **Import another report** to display the Open window, where you can select a report.
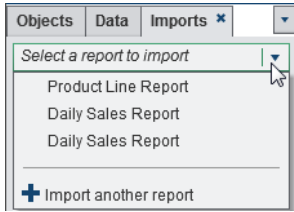
In the tree view on the left, a list of sections and report objects appears. Drag a single report object, multiple report objects, or an entire section that you want to include in the new report, and drop it onto the canvas.
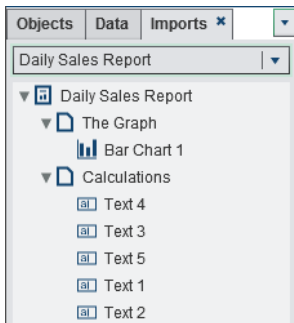
Note: Report objects that cannot be imported from the explorer are not displayed on the **Imports** tab.

2   (Optional) Modify the report objects, properties, styles, roles, filters, display rules, ranks, and interactions.

3   Select **File ▶ Save**, or click 💾, which displays the Save As window. Enter a **Name**. Report names cannot use these characters: / \

4   Click **Save**.

Opening a Report

Reports in the designer are saved when you use either the Save window or the Save As window. A saved report contains at least one section. Typically, a section uses data items from a data source to perform queries. The section displays the results with one

or more report objects (for example, a table, a graph, a gauge, a control, and so on). For more information about sections, see "Overview of Report Sections" on page 465.

A section is not required to contain any report objects. For example, you might have a report that you use only as a template to create other reports if you want all reports to have a similar appearance. A template-like report might contain data sources, calculated items, global data filters, and shared display rules, but it might not contain report objects.

E-mailing a Report

You can e-mail a link to a report to someone. If the recipient opens the e-mail message on a PC or Mac with a Flash-enabled viewer, and then clicks the e-mail link, then the report is displayed in the viewer. If the recipient opens the e-mail message on a mobile device, and then clicks the e-mail link, then the report is displayed using the SAS Mobile BI app if the user has it installed.

To e-mail a report from the designer:

1   Select **File ▶ E-mail**. The E-mail window is displayed. The URL for the report is automatically generated, and a preview appears in the e-mail message.

2   Enter the recipient's e-mail address.
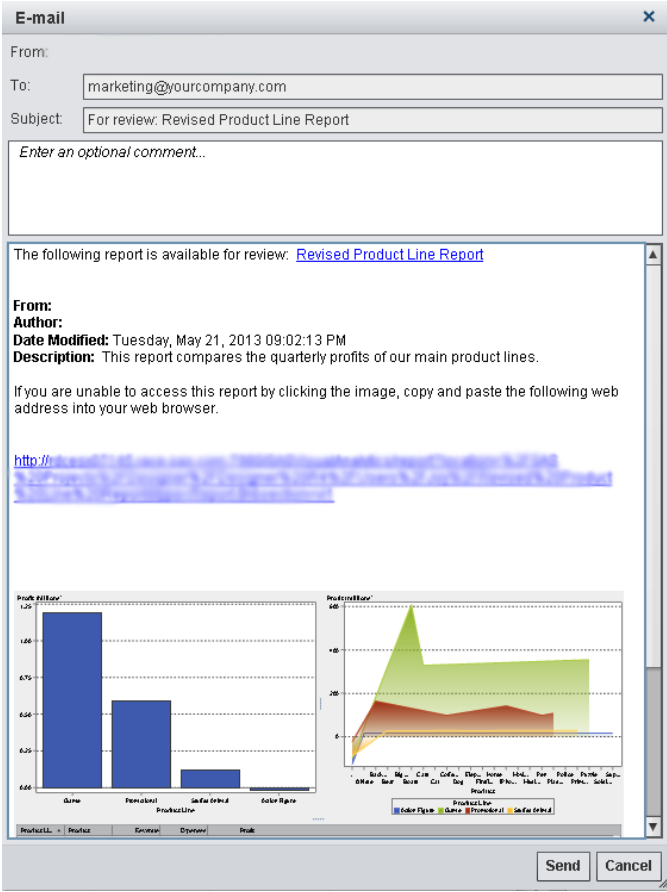
3   (Optional) Enter a comment for the report.

4   Click **Send**.

Printing a Report

This is a pre-production feature.

To print a report:

1   Select **File ▸ Print (pre-production)** or click 🖳. The Print (pre-production) window
    is displayed.

2   Select the **Paper size**, **Orientation**, **Options**, and **Margins**. Your selections are saved for that report.

Clearing the **Retain the graph quality when the document is resized** check box means that the quality of the report will be compromised if you look at the report at a zoom percentage greater than 100%.

Selecting the **Show all rows in a table** check box means that the data in the table will not be truncated. Specifying this option can impact the layout of the table.

3   Click **Print**.

Adding Comments to a Report

In the designer, you can add comments to the whole report, but you cannot add comments to the individual report objects.

> **TIP** You must save a report before you can add comments.

Note: If comments have been added to an individual report object using the home page, the explorer, or the viewer, then those comments cannot be displayed or edited in the designer.

Note: If you have the Add or View Comments capability, then you can add or view comments. You can edit your own comments and respond to other comments.

Deleting a Report

To delete a report:

1   Select **File** ▶ **Open**. The Open window is displayed.

2   Select the report that you want to delete, and then click 🗑.

You can also delete a report from the home page. Click 🏠, and then click **Manage** (next to **My Content** or **Other Content**). For more information, see "Managing Content on the Home Page" on page 13.

32

Using Report Objects

About Report Objects

After selecting your data source and data items, add one or more report objects to display the results. The SAS Visual Analytics Designer (the designer) provides report objects for all of your reports. (You can also select data after you add report objects to the canvas.) Report objects in the designer are grouped into types on the **Objects** tab in the left pane and in the **Insert** menu: **Tables**, **Graphs**, **Gauges**, **Controls**, **Containers**, **Other**, and **Custom**. **Custom** is displayed only if one or more custom graphs are saved in your `My Folder` location.

For a definition and a picture of each report object, see .

The following table lists the available table objects:

Icon	Table Type
	List Table
	Crosstab

The following table lists the available graph objects:

Icon	Graph Type
	Bar Chart
	Targeted Bar Chart
	Waterfall Chart
	Line Chart
	Pie Chart
	Scatter Plot
	Time Series Plot
	Bubble Plot
	Treemap
	Dual Axis Bar Chart
	Dual Axis Line Chart
	Dual Axis Bar-Line Chart
	Dual Axis Time Series Plot

The following table lists the available gauge objects:

Icon	Gauge Type
	Bullet
	Slider
	Thermometer
	Dial
	Speedometer

The following table lists the available controls:

Icon	Control Type
	Drop-down List
	List
	Button Bar
	Text Input
	Range Slider

The following table lists the available containers:

Icon	Container Type
	Vertical Container
	Horizontal Container
	Stack Container

The following table lists the other report objects:

Icon	Other Type
	Text
	Image
	Stored Process
	Geo Bubble Map
	Geo Region Map

Report objects of the **Custom** type are created using SAS Visual Analytics Graph Builder (the graph builder). They are identified by this icon: .

In the designer, you have access to report objects from SAS Visual Analytics Explorer (the explorer). You can open a histogram, heat map, box plot, or correlation matrix in a report that has been exported from the explorer. However, you cannot create new histograms, heat maps, box plots, or correlation matrices in the designer.

Inserting a Report Object into a Report

To insert a report object into a report, choose one of the following methods:

■ Drag the report object from the **Objects** tab in the left pane and drop it onto the canvas.

■ Double-click the report object on the **Objects** tab in the left pane. The report object is automatically placed in the canvas. If you want the report object to appear in a different location, then drag and drop it in a new location.

■ Select the report object on the **Objects** tab in the left pane, and press Enter. The report object is automatically placed in the canvas. If you want the report object to appear in a different location, then drag and drop it in a new location.

■ Use the **Insert** menu to select the report object that you want to insert. The report object is automatically placed in the canvas. If you want the report object to appear in a different location, then drag and drop it in a new location.

Note: If you insert a container, then you can drag and drop other report objects onto the container.

Note: Additional steps are required for inserting images. For more information, see "Insert an Image into a Report" on page 330.

Note: Additional steps are required for inserting stored processes. For more information, see "Use a Stored Process" on page 332.

Showing or Hiding Report Objects in the Objects Tab

You can customize which report objects you want to see on the **Objects** tab. Once you hide a report object, it remains hidden until you choose to show it again.

Note: Custom graphs appear under the **Custom** heading on the **Objects** tab only if they are saved to your `My Folder` location.

To show or hide report objects:

1 On the **Objects** tab, click ⩢ , and then select **Show or Hide Objects**. The Show or Hide Objects window is displayed.

2 Select the report objects that you want to show on the **Objects** tab. If there are report objects that you do not want to see on the **Objects** tab, then clear one or more check boxes for one or more report objects.

3 (Optional) To find a custom graph object that is not listed, click **Select custom**. The Choose an Item window is displayed. Select the custom graph object, and click **OK** to return to the Show or Hide Objects window.

4 Click **OK**. The **Objects** tab is updated.

Using Tables to Display Results

For a definition and a picture of each table type, see "Tables" on page 544.

Working with List Tables

By default, a list table contains aggregated data with one row for each distinct combination of category values. However, if the **Show detail data** check box has been selected, then all of the data is not aggregated.

You can add sparklines to a column (if the data source contains a date data item) when aggregated data is displayed in the list table. For more information, see "Add Sparklines to a List Table" on page 319.

To sort the list table by a column, click on the column heading. An arrow appears in the column heading to indicate the sorting.

Note: If you sort a list table, then the list table displays only the first 5,000 sorted rows.

To rearrange the columns in a list table, drag and drop the column headings.

Working with Crosstabs

You can show subtotals and totals by selecting the appropriate check box (or check boxes) on the **Properties** tab for the crosstab. You can show a percentage of total or percentage of subtotal in a crosstab. For more information about percentage of total or percentage of subtotal, see "Create Derived Items for Measures" on page 360.

You cannot add a sparkline to a crosstab.

You should consider placing lower cardinality (fewer distinct values) categories on the columns and higher cardinality (more distinct values) categories on the rows. Crosstabs can help you improve readability especially when there are several category data items to include in your table.

Note: By default, frequency is displayed only when there are no measures in the crosstab. If you add a category data item first, then the Frequency column is

automatically added. When you add a measure data item, the Frequency column is automatically replaced by the measure that you added. If you add a measure data item first, then the Frequency column is added only if you manually add it.

Specify Table Properties

To specify the properties for list tables and crosstabs:

1   If it is not already selected, select the table in the canvas that you want to update.

2   In the right pane, click the **Properties** tab.

3   Update the general properties for the table. You can update the **Name**, **Title**, **Format** (for the title's font style), and **Description**.

4   Update the object-specific properties for the table. The available properties depend on the selected table type.

Here are some details about the properties for list tables:

■   By default, the **Enable selection in the viewers** property is selected for list tables. This means that users who use the web viewer or a mobile device can select the list table and click ⓘ to see the list table name and any incoming filter information.

■   By default, a list table contains aggregated data with one row for each distinct combination of category values. However, if the **Show detail data** check box has been selected, then all of the data is not aggregated.

Here are some details about the properties for crosstabs:

■   By default, the **Enable selection in the viewers** property is selected for crosstabs. This means that users who use the web viewer or a mobile device can select the crosstab and click ⓘ to see the crosstab name and any incoming filter information.

■   You can specify that you want an indented layout and totals and subtotals for columns, rows, or both.

■   You can specify the placement of totals and subtotals.

Specify Table Styles

To specify styles for tables:

1   If it is not already selected, select the table in the canvas that you want to update.

2   In the right pane, click the **Styles** tab.

3   Update the styles for the table. The available styles depend on the selected table type.

Here are some details about the styles for list tables:

■   Select **Wrap text** to wrap text in the cells of a list table.

■   Select **Enable alternating background color** to change the color of every other row in the list table. Use the **Alternating background color** style to specify the color. Use the **Selection color** style to specify the color of the selected row in the list table.

Add Sparklines to a List Table

A *sparkline* is a small line graph that presents a single trend over time. A sparkline is about the size of one or two words, so it fits in a single cell and repeats for each row in a column. A sparkline does not have axes or labels. They are frequently used to present stock trends or production rates over time. A sparkline is intended to be both succinct and noteworthy.

In the designer, you can add sparklines to a column in a list table. The data source for the list table must include a date, datetime, or time data item before you can add a sparkline.

Here is an example of a simple list table with a sparkline:

Product Line	Revenue	Expenses	Profit ▼	Profit Trend
Game	1,671,890,035	477,809,929	1,194,080,107	
Promotional	813,699,290	223,822,374	589,876,916	
Stuffed Animal	276,990,966	159,548,680	117,442,285	
Action Figure	262,318,761	281,390,254	-19,071,493	

To add sparklines:

1   If it is not already selected, select the list table in the canvas that you want to update.

2   Right-click the list table, and then select **Add Sparkline**. The Add Sparkline window is displayed.

3   Enter a **Column label**.

4   For **Time Axis**, select a date, datetime, or time data item in the current data source.

5   Select a **Measure (line)**.

6   (Optional) Select the **Set baseline** check box. Enter a **Value**, and select a **Fill type**. Your choices are **Gradient** or **Solid**.

   The **Set baseline** option draws a horizontal line through the graph at the point on the Y axis where the baseline value resides. Everything above or below the baseline is filled in with either a solid or gradient color.

7   Click **OK**. The sparkline is added to the last column in the list table. You can move the sparkline to another location in the table.

To edit a sparkline, right-click in the sparkline column in the list table, and then select **Edit Sparkline**. The Edit Sparkline window is displayed. Update the information, and then click **OK** to save your changes.

To delete a sparkline, right-click in the sparkline column in the list table, and then select **Remove Sparkline**.

Using Graphs to Display Results

For a definition and a picture of each graph type, see "Graphs, Charts, and Plots" on page 545.

Specify Graph Properties

To specify the properties for graphs:

1   If it is not already selected, select the graph in the canvas that you want to update.

2   In the right pane, click the **Properties** tab.

3   Update the general properties for the graph. You can update the **Name**, **Title**, **Format** (for the title's font style), and **Description**.

4   Update the specific properties for the graph. The available properties depend on the selected graph type.

Here are details about the properties for graphs:

■   By default, the **Enable selection in the viewers** property is selected for graphs. This means that users who use the web viewer or a mobile device can select the graph and click ⓘ to see the graph name and any incoming filter information.

■   To add a new horizontal or vertical reference line to all graph types except the pie chart and the treemap, select **Create new reference line**. You can specify a **Label**, an **Axis**, a **Value**, and the **Style** of the new line.

Note: Properties for the axes or reference lines are not displayed for a custom graph that has merged or for common axes.

■   To fill a line chart or a time series plot with color, select a **Grouping style**. The style choices are **Overlay Unfilled**, **Overlay Filled**, and **Stack Filled**. The default grouping style is **Overlay Unfilled**.

- ■ To change the percentage for the "other" slice in a pie chart, change the amount listed in the **Minimum percentage for "Other"** field. The default is 4%.

- ■ To change the transparency of the markers in a scatter plot or the bubbles in a bubble plot, move the slider for **Transparency** to any value between **Low** and **High**. The default value for transparency is **Low**.

- ■ Use the **Binning interval** for a time series plot or a dual axis time series plot when you have an uneven grouping of time data. The options are:

  Automatic
  Determines the best bin size for your data. This option is the default.

  Fixed count
  Enables you to specify an interval between 10 and 500 in the **Fixed bin count** field.

  Use format
  Uses the format of the date data item as the interval.

Specify Graph Styles

To specify styles for graphs:

1 If it is not already selected, select the graph in the canvas that you want to update.

2 In the right pane, click the **Styles** tab.

3 Update the styles for the graph. The available styles depend on the selected graph type. For example, you can specify **Data Styling**, **Frame Styling**, **Text Styling**, and **Data Colors** for graphs.

  By default, the background of a graph is set to white. Use the **Wall background** option (under **Frame Styling**) to specify a different color.

Using Gauges to Display Results

For a definition and a picture of each gauge type, see "Gauges" on page 559.

About Gauges

A gauge is a dashboard indicator (also known as a KPI) that displays the status or measure of a variable or variables in relation to a target, goal, or interval. Gauges are designed to achieve this goal in a way that is familiar to users. Many real-life objects use gauges, such as cars and machines. Gauges can be used to display a quantity, range, variable, or status. They often appear in business intelligence dashboards.

Qualitative ranges are required for all gauges in the designer. You can populate the range intervals manually, or you can have them generated for you based on the range of the actual data.

Gauges in the designer support high cardinality.

Specify Gauge Properties

To specify properties for gauges:

1   If it is not already selected, select the gauge in the canvas that you want to update.

2   In the right pane, click the **Properties** tab.

3   Update the general properties for the gauge. You can update the **Name**, **Title**, **Format** (for the title's font style), and **Description**.

4   Update the specific properties for the gauge. The available properties depend on the selected gauge type.

    Here are some details about the properties for gauges:

- By default, the **Enable selection in the viewers** property is selected for gauges. This means that users who use the web viewer or a mobile device can select the gauge and click ⓘ to see the gauge name and any incoming filter information.

- For the bullet, slider, and thermometer gauges, you can specify the **Direction** for displaying the gauge. The default for the bullet and slider gauges is **Horizontal**. The default for the thermometer gauge is **Vertical**.

Specify Gauge Styles

To specify styles for gauges:

1   If it is not already selected, select the gauge in the canvas that you want to update.

2   In the right pane, click the **Styles** tab.

3   Update the styles for the gauge. The available styles depend on the selected gauge type. For example, you can specify **Data Styling**, **Frame Styling**, **Text Styling**, and **Data Colors** for gauges.

Using Controls to Display Results

For a definition and a picture of each control type, see "Controls" on page 562.

Working with Controls

A control is a report object that filters or narrows the scope of the data that you are currently viewing. A control enables you to group your data by a selected category, and then select which group you want to view. When you drag and drop a data item onto a control, the control creates a group based on that data item. For example, you might have a data item called *Cars* that contains all of the models that a manufacturer produces. When you drag and drop the *Cars* data item onto a drop-down list, the control groups the car models, and then you can select a car model to use as a filter. Controls can be used in a report with interactions.

Section prompts are the controls that are placed in the special row area at the top of the canvas. Only the drop-down list, button bar, and text input controls can be used in section prompts. A section prompt automatically filters all of the other report objects in the same section, as long as the report object uses the same data source as the section prompt control. For more information, see "Use a Control to Create a Section Prompt" on page 326.

You can place any control in the main area of the canvas below the section prompt row. You must define explicit interactions (using either the **Interactions** tab or the Interactions view) between these controls (as the source report objects) and one or more target report objects. For more information about interactions, see "Overview of Interactions" on page 431.

The following controls are available in the designer:

- drop-down lists

- lists

- button bars

- text input fields

- range sliders

Specify Control Properties

To specify the properties for a control:

1   If it is not already selected, select the control in the canvas that you want to update.

2   In the right pane, click the **Properties** tab.

3   Update the general properties for the control. You can update the **Name**, **Title**, **Format** (for the title's font style), and **Description**.

4   Update the specific properties for the control. The available properties depend on the selected control.

Here are some details about the properties for controls:

- By default, the **Enable selection in the viewers** property is not selected for controls. This means that users who use the web viewer or a mobile device cannot select the control and click ⓘ to see the control name and any incoming filter information. However, the user can still modify values for the control.

- For drop-down lists, lists, and button bars, select the **Required** property if you want to require users to make a selection in the control. If you select the **Required** property for a list, at least one check box must always be selected.

- By default, the **Allow multiple selection** property is selected for lists. If you clear the **Allow multiple selection** check box, radio button are displayed instead of check boxes and the **Required** property is applied automatically.

- For button bars and range sliders, the **Horizontal** property is selected by default.

Specify Control Styles

To specify styles for controls:

1   If it is not already selected, select the control in the canvas that you want to update.

2   In the right pane, click the **Styles** tab.

3   Update the styles for the control. The available styles depend on the selected control type. For example, for a drop-down list, you can specify **Drop–down Styling** and **Text Styling**.

Use a Control to Create a Section Prompt

The drop-down list, button bar, and text input controls are the only ones that can be used as section prompts.

To use a control to create a section prompt:

1   Drag the control icon from the **Objects** tab in the left pane and drop it onto the area above the canvas. (Look for the hint text that says, "Drop controls here to create a section prompt.") The control appears above the canvas.

2   Drop a category onto the control. For example, if you drag and drop a drop-down list control, then you can assign a category like *Facility City* or *Facility State*. Then the drop-down list is populated with the cities or states that are used in that category.

You can also use the **Roles** tab in the right pane, and then specify the **Category** and **Frequency** roles.

If you use a control to create a section prompt, then the user can select a value to filter the data. For some control types, when the **Required** property is not selected, the user might need to press Ctrl+click to clear the value in the filter.

Using Container Object Types in Reports

About Container Objects

You can use a vertical, horizontal, or stack container to group other report objects. After you insert a container, you can drag and drop other report objects onto the container.

When you use a stack container, the report objects are displayed as if they are in a slide deck. Only one report object is displayed at a time. The stack container has a control bar instead of a scroll bar that lets you move between report objects. You cannot nest stack containers. However, you can place them side by side. When you use precision layout, you can resize the stack container.

Note: Stored processes cannot be added to a container.

Specify Container Properties

To specify the properties for a container:

1   If it is not already selected, select the container in the canvas that you want to update.

2   In the right pane, click the **Properties** tab.

3   Update the general properties for the container. You can update the **Name**, **Title**, **Format** (for the title's font style), and **Description**.

4   Update the object-specific properties for the container. Here are some details about the properties for containers:

▪   For stack containers, the **Navigation control location** property changes the location of the control bar, and the **Navigation button type** property changes the appearance of the control bar. You can reorder the display of the report objects in the stack container using the **Objects** list.

▪   By default, the **Enable selection in the viewers** property is selected for containers. This means that users who use the web viewer or a mobile device can select the container and click ⓘ to see the container name and any incoming filter information.

▪   Update the order in which the report objects appear inside the container.

Specify Container Styles

No styles are available for containers.

Using Other Object Types in Reports

Use a Text Object

About Text Objects

Text objects display static text. You can use text to include company standards, such as text that indicates confidentiality, in your reports. You can have hyperlinks in text. For more information about links, see "Create a Link from a Text Object" on page 453. You can also use text to annotate other objects in a report.

Specify Text Object Properties

To specify the properties for a text object:

1   If it is not already selected, select the text object in the canvas that you want to update.

2   In the right pane, click the **Properties** tab.

3   Update the general properties for the text. You can update the **Name**, **Title**, **Format** (for the title's font style), and **Description**.

4   (Optional) Update the object-specific properties for the text object. By default, the **Enable selection in the viewers** property is not selected for text objects. This means that users who view the report in the web viewer or a mobile device cannot select the text in a report.

Specify Text Object Styles

You can use the floating toolbar to change the font, font size, text color, and text background color. You can specify whether the text is bold, italic, or underlined, and whether it is left-aligned, centered, or right-aligned. You can also use the floating toolbar to create a link from a text object. For more information, see "Create a Link from a Text Object" on page 453.

> **TIP** You can use the context menu to cut, copy, and delete text. However, you have to use the keyboard shortcut (Ctrl+V) to paste text.

Note: You cannot change text styles using the **Styles** tab.

Use an Image

About Images

You can use images to include your corporate logo or other graphics in your reports. You can insert images from a repository or from your local machine. If you select an image from your local machine, it is saved to the repository. You can also add tooltip text to an image.

Insert an Image into a Report

To insert an image into a report:

1 Choose one of the following:

- Drag the 🖼 from the **Objects** tab in the left pane and drop it onto the canvas. The Image Selection window is displayed.

- Select **Insert ▶ Other ▶ Image**. The Image Selection window is displayed. The image object is automatically placed in the canvas. If you want the image to appear in a different location, then drag and drop it in a new location.

2 Select the image from one of the following locations:

Load from repository
Select this option to choose an image that is stored on the same server as the reports.

Load from local machine
Select this option to choose an image from your local machine. Click **Browse** to choose a file on your local machine. Specify a repository in the **Save the local image to the repository** field. If you click **Browse**, the Save As window is displayed. Select a folder and then click **OK** to return to the Image Selection window.

A preview of the image is displayed.

3 (Optional) Specify the **Scale type**:

None
The actual size of the image is maintained. The image might or might not fill the entire area of the image's visual container. If the image is larger than the visual container, then scroll bars are displayed.

Stretch
The height and width of image are set to the height and width of the image's visual container. The image's original aspect ratio is not maintained.

Fit All

The image is modified to fit best into the image's visual container. The image's original aspect ratio is maintained.

Fit Width

The width of the image is set to the width of the image's visual container. The height maintains the image's original aspect ratio. Scroll bars are displayed if the set height of the image is greater than the height of the visual container.

Fit Height

The height of the image is set to the height of the image's visual container. The width maintains the image's original aspect ratio. Scroll bars are displayed if the set width of the image is greater than the width of the visual container.

Tile

The image is tiled in the visual container. The image's original size is maintained. There are no scroll bars.

4 (Optional) Specify the **Tooltip text**.

5 Click **OK**.

Specify Image Properties

To specify the properties for an image:

1 If it is not already selected, select the image in the canvas that you want to update.

2 In the right pane, click the **Properties** tab.

3 Update the general properties for the image. You can update the **Name**, **Title**, and **Description**.

4 Update the properties specific for the image. Your choices are **Location**, **Scale type**, and **Tooltip text**.

By default, the **Enable selection in the viewers** property is not selected for images. This means that users who use the web viewer or a mobile device cannot select the image and click ⓘ to see the image name.

Specify Image Styles

No styles are available for images.

Use a Stored Process

About Stored Processes

A *stored process* is a SAS program that is stored on a server and that can be executed as requested by client applications such as SAS Visual Analytics. The embedded SAS code can contain instructions for displaying report elements that include queries, prompted filters, titles, images, and statistical analyses.

You can add one or more stored processes to a report. You can edit the prompts for a stored process using the **Properties** tab in the designer. When you set the value for a prompt in the designer, that prompt value becomes the default value for the viewer.

There are limitations to where stored processes can be used in the designer:

- A stored process cannot be used in the precision layout.

- A stored process cannot be the source or target of an interaction in a report.

- A stored process cannot be added to a container.

Note: Users who view stored processes using SAS Mobile BI cannot be prompted. Instead, the stored process runs using the prompt values that were added when the report was created.

Note: The stored process output in a report is rendered as HTML regardless of the requested output type.

For information about creating a stored process and registering it in metadata, see *SAS Stored Processes: Developer's Guide*.

Specify Stored Process Properties

To specify the properties for a stored process:

1   If it is not already selected, select the stored process in the canvas that you want to update.

2 In the right pane, click the **Properties** tab.

3 Update the general properties for the stored process. You can update the **Name**, **Title**, **Format** (for the title's font style), and **Description**.

4 Update the properties specific to the stored process. Your choices are **Show metadata view** and **Show SAS log in the output**. If the stored process contains prompts, the **Edit Prompts** button is displayed.

Click **Edit Prompts** to open a new window where you can edit the parameters for the stored process. Click **OK** to save your changes.

Selecting the **Show metadata view** check box can make it easier to work with the stored process while you are designing a report. In the metadata view, you can see information about when the stored process was created and last modified. You can also see the name of the stored process file.

Here is an example of the metadata view for a stored process:

Selecting the **Show SAS log in the output** option means that both the log output and the stored process output are displayed in the report. This can assist you in debugging problems.

Specify Stored Process Styles

No styles are available for stored processes.

Use a Geo Bubble Map or a Geo Region Map

About Geo Bubble Maps

Many types of data have a spatial aspect, including demographics, marketing surveys, and customer addresses. For example, if a user needs to evaluate population data for U.S. Census tracts, a report designer could display the information in a table. However, it would be easier and more effective for the person using the report to see the information in the context of the geography of the tracts. When evaluating information that has a spatial component, users might find it easier to recognize relationships and trends in the data if they see the information in a spatial context.

You can add a geo bubble map only if the report uses data items from a data source that is enabled for geographic mapping. A geo bubble map requires a geography variable with the role type of geography.

You can create a filter or brush interaction between a geo bubble map and other report objects in your report or dashboard. When you click on a specific region or city, the other report objects filter or highlight to show the same location.

Note: Certain geo bubble maps that are exported from the explorer (for example, geo maps that use custom roles or data sets that have centroids) cannot be fully modified in the designer.

For a definition and picture of geo bubble maps type, see "Geo Maps" on page 558.

About Geo Region Maps

You can add a geo region map only if the report uses data items from a data source that is enabled for geographic mapping. A geo region map requires a geography variable with the role type of geography. However, it does not support custom geography data items.

You can create a filter or brush interaction between a geo region map and other report objects in your report or dashboard. When you click on a specific country or state, the other report objects filter or highlight to show the same location.

Note: Certain geo region maps that are exported from the explorer (for example, geo maps that use custom roles or data sets that have centroids) cannot be fully modified in the designer.

For a definition and picture of geo region maps type, see "Geo Maps" on page 558.

Specify Geo Bubble Map or Geo Region Map Properties

To specify the properties for a geo bubble map or geo region map:

1   If it is not already selected, select the geo bubble map or geo region map in the canvas that you want to update.

2   In the right pane, click the **Properties** tab.

3   Update the general properties for the geo bubble map or geo region map. You can update the **Name**, **Title**, **Format** (for the title's font style), and **Description**.

4   Update the properties that are specific to the geo bubble map or geo region map. You can update the **Show map navigation control**, **Transparency**, and **Show legend**.

By default, the **Enable selection in the viewers** property is selected for geo bubble maps or geo region maps. This means that users who use the web viewer or a

mobile device can select the geo bubble map or geo region map and click ⓘ to see the geo bubble map or geo region map name and any incoming filter information.

Note: For a geo bubble map, you select the **Transparency** for the bubble plot. For a geo region map, you select the **Transparency** for the region map.

Specify Styles for Geo Bubble Maps and Geo Region Maps

To specify styles for a geo bubble map or a geo region map:

1   If it is not already selected, select the geo bubble map or geo region map in the canvas that you want to update.

2   In the right pane, click the **Styles** tab.

3   Update the styles for the geo bubble map or geo region map. You can specify **Data Styling**, **Frame Styling**, **Text Styling**, and **Data Colors** for geo bubble maps or geo region maps.

Using Custom Graphs to Display Results

About Custom Graphs

The graph builder enables you to create custom graph objects. To access the graph builder, on the **Objects** tab, click ⌶. The Create Custom Graph window is displayed. For more information, see Chapter 41, "Creating Custom Graph Objects," on page 477.

When a custom graph is saved in the `My Folder` location, it is displayed under the **Custom** heading on the **Objects** tab. Then, you can insert the custom graph into a report and add data to it. For information about adding a custom graph to the **Objects** tab, see "Showing or Hiding Report Objects in the Objects Tab" on page 316.

For a definition and a picture of each custom graph type, see "Graphs, Charts, and Plots" on page 545.

Specify Custom Graph Properties

To specify the properties for custom graphs:

1   If it is not already selected, select the custom graph in the canvas that you want to update.

2   In the right pane, click the **Properties** tab.

3   Update the general properties for the graph. You can update the **Name**, **Title**, **Format** (for the title's font style), and **Description**.

4   Update the specific properties for the graph. The available properties depend on the selected custom graph type.

   By default, the **Enable selection in the viewers** property is selected for custom graphs. This means that users who use the web viewer or a mobile device can select the custom graph and click ⓘ to see the custom graph name and any incoming filter information.

   Note: No properties are available for the schedule chart.

Specify Custom Graph Styles

To specify styles for custom graphs:

1   If it is not already selected, select the custom graph in the canvas that you want to update.

2   In the right pane, click the **Styles** tab.

3   Update the styles for the graph. The available styles depend on the selected custom graph type. For example, you can specify **Data Styling**, **Frame Styling**, **Text Styling**, and **Data Colors** for custom graphs.

   By default, the background of a graph is set to white. Use the **Wall background** option (under **Frame Styling**) to specify a different color.

Duplicating a Report Object

Duplicating a report object in the designer enables you to use a copy of the same object in the same section or another section of your report.

Note: You cannot duplicate a custom graph.

To duplicate a report object:

1   On the canvas, right-click the report object that you want to duplicate.

2   Select **Duplicate <ReportObject>**, where `<ReportObject>` is the name of the report object in the report. (For example, `List Table 1`, `Bar Chart 1`, and so on.) The duplicated report object is placed on the canvas with a name based on the original name. For example, if the original report object name is `List Table 1`, then the duplicate report object is displayed as `List Table 1 (1)`. If you choose to duplicate the same report object again, then it is displayed as `List Table 1 (2)`.

3   (Optional) Move the duplicate report object to another section. Right-click the report object that you want to move. Select **Move <ReportObject> to ▶ <SectionName>**, where `<ReportObject>` is the name of the report object and `<SectionName>` is the name of the section.

4   (Optional) If you want the report object to appear in a different location, then drag and drop it in a new location.

   If you move an object to a section with precision layout, then you must manually move the object to its proper location. All objects are put in the top left corner by default.

Working with Alerts for Report Objects

You can create alerts for a report object so that you are notified via e-mail when the alert criteria is met. You can specify how frequently the system checks to see whether the alert criteria have been met.

Add an Alert

To add an alert to a report object:

1   If it is not already selected, select the report object in the canvas to which you want to add an alert.

2   Do one of the following:

■   Right-click the report object, and select **Add Alert**. The Add Alert window is displayed.

■   In the right pane, click the **Alerts** tab. Click . The Add Alert window is displayed.

3   On the **Expression** tab, specify the criteria for the alert. You can create a new expression or use an existing expression.

4   (Optional) Specify how often you want the system to check for the criteria. You can use the system default, which is set by your administrator, or you can limit the check to a minute or hourly increment.

5   On the **Subscription and Notification** tab, add or remove subscribers for the alert.
    You can specify how frequently e-mail messages should be sent.

6   Click **OK**. The details for the alert are displayed at the bottom of the **Alerts** tab.

Edit an Alert

To edit an alert:

1   In the right pane, click the **Alerts** tab.

2   Select the alert that you want to edit, and click . The Edit Alert window is
    displayed.

3   Update the alert criteria, and then click **OK** to save your changes.

Delete an Alert

Alerts are not automatically deleted when a report is deleted. You can delete an alert on the **Alerts** tab. Select the alert that you want to delete, and click 🗑. Then, click **Delete** in the confirmation message that is displayed.

33

Working with Data in SAS Visual Analytics Designer

Overview of Data Sources and Data Items

Each data source SAS Visual Analytics Designer (the designer) includes one or more data items that you can use in reports. For example, a data source named `Order Information` might include standard data items such as `Order ID`, `Product ID`, `Unit Cost`, `Order Date`, and `Order Amount`. You decide which data items to use. You can select all of the data items in the data source or a subset of the data items.

Working with Data Sources in Reports

About Data Sources

Data sources that are available in the designer are prepared by a data administrator or analyst so that you can easily define a report. Data administrators load tables into memory using SAS Visual Analytics Administrator. Analysts can use SAS Visual Data Builder (the builder) to design queries that load tables into memory too. All data sources contain data items, which can refer to calculations or columns in physical data (tables). Reports can include query results from more than one data source.

Add a Data Source to a Report

You can use one or more data sources in a report in the designer.

To add a data source to a report:

1   On the **Data** tab, click ▼ beside the **Select a data source** text to display the Add Data Source window.

Property	Value
Name	
Role	
Format	
Aggregation	
Sort Options	
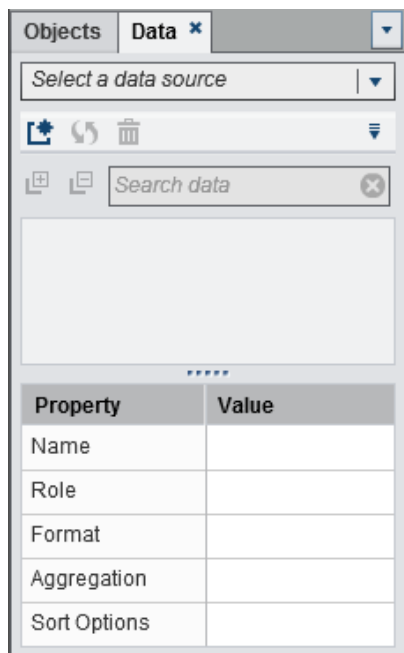
2   In the Add Data Source window, select a data source.

> **TIP** Use the **Search** field to narrow the list of data sources that are displayed in
> the Add Data Source window. The search searches the **Name** and **Description**
> fields.

3   Click **Add**. The list of available data items is displayed in the **Data** tab.

4   (Optional) To add another data source, click ⬆, which displays the Add Data Source
window. Select the data source that you want, and then click **Add**. The **Data** tab is
populated with a list of all of the data items that are in the data source.

Import a Data Source for a Report

If you have the Import and Load Data capability, then you can import a data source into
the designer. Acceptable files are SAS data sets, Microsoft Excel spreadsheets, and
delimited text files, such as CSV files.

When you import data, the data source is automatically added to the open report.

To import a data source for a report:

1   On the **Data** tab, click ⬈. The Select file to upload by window is displayed.

2   Select a data source (or data sources) that you want to import. The Import Local File window is displayed.

   If you select a delimited text file, then you can specify additional options. For example, you can specify the delimiter, whether the first row contains column heading names, and where the data rows begin.

   If you select a spreadsheet, then you can specify additional options. For example, you can specify which worksheets to import, whether the first row contains column heading names, and where the data rows begin.

3   Click **OK**.

Note: You cannot rename an imported data source unless there is a filename conflict.

You can also import data using the Add Data Source window or the Change Data Source window. The Add Data Source window allows you to import multiple data sources. However, the Change Data Source window only allows you to import one data source.

Refresh a Data Source for a Report

You can refresh the columns in a data source in the designer at any time. Be aware that refreshing a data source means that all live report objects that are connected to that data source will have their queries re-run.

Note: Refreshing a data source adds any new columns that have been added to the table metadata. The default formats and names of existing columns will be updated the next time you open the report.

To refresh a data source for a report, on the **Data** tab, select the data source, and then click ↻. When you refresh a data source for a report, columns that have been deleted from the table metadata will automatically be removed if they do not impact any objects in the report. If deleted columns do impact objects in the report, then the Repair Report window is displayed so that you can repair the objects that are impacted by the deleted

columns.For more information, see Appendix 6, "Troubleshooting in SAS Visual Analytics Designer," on page 611.

Remove a Data Source from a Report

You can remove a data source from a report in the designer. Be aware that removing a data source means that all related data items are also removed from the report objects within the report.

To remove a data source for a report:

1   On the **Data** tab, select the data source, and then click 🗑.

2   Click **Delete** in the confirmation message that is displayed.

Change a Data Source in a Report

To change a data source for a report:

1   On the **Data** tab, click ▼ and then select **Change Data Source**. The Change Data Source window is displayed.

2   In the Change Data Source window, select a data source.

3   Click **Change**.

    If a data item with the same name does not exist in the replacement data source, then it is automatically removed from the replacement data source if there are no report objects that use the data item. If there are report objects that use the data item, then those report objects will not work. The Repair Report window is displayed so that you can repair the objects that use the data item. For more information about repairing reports, see Appendix 6, "Troubleshooting in SAS Visual Analytics Designer," on page 611.

    Data items in the replacement data source that have names that do not exist in the original data source are added to the **Data** tab automatically.

Note: For data item names, the case is ignored when data sources are compared by the designer.

Viewing Measure Details

To view the details about all of the measures in a data source:

1   On the **Data** tab, click ⬇ and then select **Measure Details**. The Measure Details window is displayed.

2   Click **Close**.

Showing or Hiding Items

You can specify which data items you want to see for the data source in the **Data** tab.

To show or hide data items:

1   On the **Data** tab, click ⬇ , and then select **Show or Hide Items**. The Show or Hide Data Items window is displayed.

2   Select the data items that you want to appear on the **Data** tab. If there are data items that you do not want to see on the **Data** tab, then clear the check box (or check boxes) for that data item (or data items).

3   Click **OK**. The **Data** tab is updated.

Alternatively, you can select a data item that you want to hide on the **Data** tab. Right-click the data item, and select **Hide Data Item**.

You can also use a data source filter to restrict the data that is displayed in a report. For more information, see "Working with Data Source Filters in a Report" on page 428.

Sort Items

To sort data items on the **Data** tab, click ⬇ , and then select one of the following:

- ■ **Sort Items ▶ Ascending By Name**

- ■ **Sort Items ▶ Descending By Name**

The data items are sorted on the **Data** tab within each grouping. The default sort is **Ascending By Name**.

For information about sorting data values in report objects, see "Sorting Data in Reports" on page 388.

Group Items

To group data items on the **Data** tab, click ▼, and then select one of the following:

- ■ **Group Items ▶ By First Letter**

- ■ **Group Items ▶ By Data Type**

- ■ **Group Items ▶ By Role**

- ■ **Group Items ▶ By Format**

- ■ **Group Items ▶ By Aggregation**

The data items are grouped on the **Data** tab. The default is grouping **By Role**.

Working with Hierarchies in a Report

About Hierarchies

Creating hierarchies enables you to add drill-down functionality to your reports. A hierarchy is an arrangement of category columns that is based on parent-child relationships. The levels of a hierarchy are arranged with more general information at the top and more specific information at the bottom. For example, you might create a hierarchy of datetime columns with *Year* as the top level, *Month* as the next level, and *Day* as the bottom level.

You can also have a geographic hierarchy. For example, you might have a hierarchy with *Region* as the top level, *State* as the next level, and *City* as the bottom level.

You can have a maximum of two hierarchies for a report object.

Note: List tables and controls do not support hierarchies.

Note: Data item auto-assignment does not support hierarchies.

Note: Crosstabs can have either a hierarchy or categories on each row or column, but not both.

Note: Time series plot report objects allow only datetime data items in a hierarchy.

Note: Geo bubble maps and geo region maps allow only geographic data items in a hierarchy.

Create a New Hierarchy for a Report

To create a new hierarchy:

1   On the **Data** tab, click ▼ and then select **New Hierarchy**. The New Hierarchy window is displayed.

2   Enter a **Name**.

3   Select at least two categories and drag them to the **Hierarchy** list.

4   (Optional) Use the up and down arrows to arrange the data items in the **Hierarchy** list.

5   Click **OK** to save the new hierarchy. The 品 icon identifies the new hierarchy in the list of data items on the **Data** tab.

Crosstab report objects enable you to create hierarchies from the categories on a crosstab axis. To create a hierarchy, right-click a category heading and then select **Create Hierarchy**. The categories are replaced with a new hierarchy. The name of the new hierarchy is generated from the name of the outermost category, with the suffix `Hierarchy`.

Edit a Hierarchy for a Report

To edit a hierarchy:

1   Right-click the hierarchy name on the **Data** tab and select **Edit Hierarchy**. The Edit Hierarchy window is displayed.

2   (Optional) Edit the **Name**.

3   Add and remove categories. There must be at least two categories.

4   Click **OK** to save the updated hierarchy.

If the hierarchy that you edit is already used in a report object and is drilled or expanded, it returns to the top level after it is edited.

Delete a Hierarchy for a Report

To delete a hierarchy:

1   Right-click the hierarchy name on the **Data** tab, and select **Delete Hierarchy**.

2   Click **Delete** in the confirmation message that is displayed. The hierarchy is removed from the list of data items, as well as from any report objects, filters, or ranks that were using it.

Working with Data Items in a Report

About Data Items

You decide which data items to use to define a query for each report object. You can use all the data items in the data source or a subset of data items. Each data item is classified as either a category or a measure.

The designer can display data items using an existing user-defined format that has already been specified externally for a data column in a data source. However, you cannot specify a new or different user-defined format for a data item in the designer.

For information about filtering data items, see "Working with Data Source Filters in a Report" on page 428.

Table 33.1   *Data Items That Are Available in the Designer*

Data Item	Icon	Description
Aggregated Measure or Time Period Calculation		A data item that represents special predefined operations, like distinct count, percentage of totals, percentage of subtotals, or frequency percent. Or the user can define their own aggregated measure calculations.
		Aggregated measures can be used in only some report objects. They cannot be used in filters, controls, spark lines, or time series graphs. Percentage of subtotal items (including row total, row subtotal, column total, and column subtotal) can be used in only in crosstabs. Some aggregated measure calculations cannot be used in a detail rank.
Calculated	, , or	A data item that is calculated from existing data items by using an expression. For example, you could create a calculated data item called *Profit*, which is created by using this expression: `[Revenue] – [Cost]`, where *Revenue* and *Cost* are measures in a data source.
		Calculated dates and times are treated as categories with distinct values being governed by the date or time format that you have chosen. Numeric calculated items can be treated as measures (with an aggregation type such as `Sum`, which is applied to each distinct category combination). Or, you can change numeric calculated items into category data items with distinct values being governed by the number of decimal places in the numeric format.

Data Item	Icon	Description
Category		A data item whose distinct values are used to group and aggregated measures. There are five types of categories: alphanumeric, date, datetime, time, and numeric. Alphanumeric categories can be made up of all letters, all digits, or a combination of the two. Categories that have values that are all digits might be physically stored as character or numeric data. The data type affects how values are handled in relation to some functionality, such as filtering, sorting, and formatting.
		Examples of alphanumeric categories include data items such as *Product ID*, *Country*, *Employee Number*, and *Employee Name*. Alphanumeric categories sort lexically.
		Date, datetime, time, and numeric categories are sorted by their underlying numeric values.
		Category data items can also be numeric. A category data item sorts differently than an alphanumeric data item. Numeric category data items sort by number.
		Note: If you change a measure to a category, then it uses this category icon.
		The ⬛ icon indicates a user-defined format category data item. User-defined format categories can be based on underlying numeric or character data.
Date and Time		A category data item whose distinct values are used to group and aggregated measures. There are three types of date categories: date, datetime, and time.
		Examples of date, datetime, and time categories are *Order Year*, *Date and Time of Sale*, and *Customer Wait Time*.
Geography		A category data item whose values are mapped to geographical locations or regions. Geography data items can be used in reports to show your data on a geographic map. For example, a geography data item can identify geographic information that is specific to your organization (for example, sales regions, warehouse locations, oil platforms, and so on). For more information, see "Working with Geography Data Items" on page 372.
		The geography icon can also indicate that all data items in a hierarchy are based on geography.

Data Item	Icon	Description
Hierarchy	🔲	A data item whose values are arranged with more general information at the top and more specific information at the bottom. The first level in the hierarchy is the root level. For example, you might have a *Time* hierarchy, which includes the *Year* (the root level), the *Quarter*, and then the *Month*. You can also have geographic hierarchies.
Measure	🔷	A data item whose values can be used in computations. These values are numeric. Examples of measures include *Sales Revenue*, *Units Sold*, and *Salary*. The designer assigns a default aggregation method to every measure. Almost all measures are assigned sum. You can change the aggregation method.

Note: Report objects that are imported from SAS Visual Analytics Explorer (the explorer) use either live or on-demand data. Therefore, you can update the properties and styles for these report objects in the designer, but you cannot change the data assigned to them.

Select Data Items

To select data items to use in queries for the current report section:

1   On the **Data** tab in the left pane, click the down arrow to display a list of available data sources. Select a data source, and the **Data** tab is populated with a list of all of the data items that are in the data source.

If the data source that you want is not in the list, click 📥, which displays the Add Data Source window. Select the data source that you want, and then click **Add**. The **Data** tab is populated with a list of all of the data items that are in the data source.

If you do not want to use the data source that you originally selected, click 🗑. Click **Delete** in the confirmation message that is displayed.

2   Select an existing report object that uses the same data source name or add a new report object to the section.

3   Drag and drop a data item onto the canvas. The data item will automatically be
    assigned a data role. For more information, see "Automatic Data Item Assignments"
    on page 357.

    Note: You can press the Ctrl key to select multiple data items, and then drag and
    drop them onto the canvas.

    Alternatively, you can right-click on a data item and select **Add Data Item to
    <ReportObject>**, where `<ReportObject>` is the name of the report object in the
    report. (For example, `List Table 1`, `Bar Chart 1`, and so on.)

4   If the selected data item can replace a current data item in the report object or if the
    selected data item is valid for multiple data roles, then select the data assignment
    from the Assign <DataItemName> as window. A data assignment that is marked
    with a red asterisk (*) is required before the query can be run.
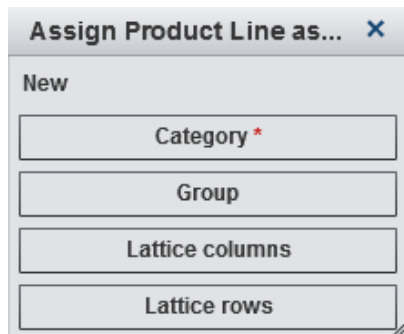
5   (Optional) To see more information about a data item, select it in the list. The **Name**,
    Role, **Format**, and **Aggregation** are displayed in the data item table below the list of
    data items. For a category data item with a user-defined format that has an
    underlying numeric value, you can specify **Sort Options**.
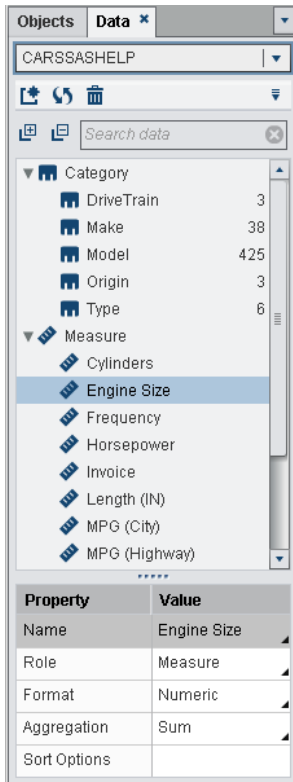
Automatic Data Item Assignments

When you drag and drop data items on a report object, the designer automatically assigns them a data role. For a single data item, the data item is assigned to an empty and required data role before you are prompted to replace an already assigned data item. For multiple data items, the data items are assigned to all empty and required data roles that accept the data items. There is a special case for a measure data role that allows multiple data items. In this special case, the designer automatically assigns all of the data items to the **Measures** data role.

Note: Hierarchy data items cannot be used in the automatic assignment of multiple data items. Also, the new animation role and data tip role are not automatically assigned.

For more information about data roles, see "Working with Data Role Assignments" on page 381.

Duplicate Data Items

Duplicating measure data items in the designer enables you to see the aggregations of a data item (**Sum**, **Average**, **Minimum**, **Maximum**, and **Count**) side by side in a table. You can duplicate a numeric measure if you want to use it as a category to group other values in some tables or graphs. If you save a report with duplicate data items, then those data items are available when you edit the report the next time.

You can duplicate a calculated data item to make a variation of a calculation. For example, you might make similar calculations involving miles per gallon for a vehicle, but you create one calculation using `MPG (City)` and another using `MPG (Highway)`. You can duplicate any data item if you want to use it with more than one format in your report. For example, you might change `Month` to `Year` for a date data item.

To duplicate a data item:

1 On the **Data** tab in the left pane, right-click the data item that you want to duplicate. Select **Duplicate Data Item**.

   All of the properties of the original data item are copied to the duplicate data item. The duplicate data item appears in the list of data items on the **Data** tab. For example, if the original data item name is `Engine Size`, then the duplicate data item is displayed as `Engine Size (1)`. If you choose to duplicate the same data item again, then it is displayed as `Engine Size (2)`.

2 (Optional) Rename the duplicate data item.

3 (Optional) Change the format or aggregation for the duplicate data item.

4 (Optional) Edit the calculation for a calculated data item or aggregated measure.

5 (Optional) Change the sort options for a category data item with a user-defined format that is based on an underlying numeric value.

6 (Optional) Change the role for the data item. For example, a numeric data item that has been duplicated and is not yet assigned to a report object can be a category or a measure.

Rename Data Items

You can rename data items in the data source using the **Data** tab.

To rename data items:

1   On the **Data** tab, right-click on a data item, and then select **Rename Data Item**. The Rename Data Item window is displayed.

2   Enter a new name. The name cannot be used by another data item in the same data source.

3   Click **OK**.

Alternatively, you can use the data item table at the bottom of the **Data** tab to rename a data item. For the **Name** property, enter a new name for **Value**.

Search for Data Items

If your data source contains many data items, you can search for particular data items using the **Data** tab.

To search for data items:

1   Enter the name of a data item in the search field on the **Data** tab. The field is located above the list of data items.

2   (Optional) Click 🔲 to collapse the list of data item groupings or click 🔳 to expand the list of data item groupings.

3   Click ❌ to clear the search term and display all of the data items in the data source.

Create a Distinct Count for a Category Data Item

A distinct count query is useful in many ways. For example, you might want to know the number of distinct products that were purchased during a specific time period. Or, you might want to know which products have the most customers or which products have the most customers in a particular geographic region. You can create a distinct count for category data items only.

To create a distinct count aggregated measure data item:

1   On the **Data** tab in the left pane, right-click the category data item that you want to use for the distinct count.

2   Select **Create Distinct Count**.

The distinct count data item appears in the list of data items with a name that is derived from the original name. For example, if the original data item name is `Date`, then the distinct count data item is displayed as `Date (Distinct Count)`. The icon identifies the new distinct count data item on the **Data** tab.

Create Derived Items for Measures

You can create derived data items that are aggregated measures in the designer. The aggregated measure does not contain data values in itself, but when it is used in a report object, it displays the value for the measure and formula type on which it is based. An example is a percentage of total.

Derived data items cannot be used in filters or controls.

To create a derived item from a report object:

1   In the report object, right-click on the measure data item that you want to use for the derived item.

2   Select **Create and Add**, and then select one of the following:

Difference from Previous Period
>
> Displays the difference between the value for the current time period and the value for the previous time period. For example, you might derive the difference between sales for the current month and sales for the previous month.
>
> **Note:** This derived item is not available if your data source does not contain a date data item that includes the year.

Difference from Previous Parallel Period
>
> Displays the difference between the value for the current time period and the value for the previous parallel time period within a larger time interval. For example, you might derive the difference between sales for the current month and sales for the same month of the previous year.
>
> **Note:** This derived item is not available if your data source does not contain a date data item that includes the year.

Percent Difference from Previous Period
>
> Displays the percentage difference between the value for the current time period and the value for the previous time period. For example, you might derive the percentage difference between sales for the current month and sales for the previous month.
>
> **Note:** This derived item is not available if your data source does not contain a date data item that includes the year.

Percent Difference from Previous Parallel Period
>
> Displays the percentage difference between the value for the current time period and the value for the previous parallel time period within a larger time interval. For example, you might derive the percentage difference between sales for the current month and sales for the same month of the previous year.
>
> **Note:** This derived item is not available if your data source does not contain a date data item that includes the year.

Percent of Subtotals
>
> Displays the percentage of the subtotal value for the measure on which it is based. You can create a percentage of subtotals only when the source data item has an aggregation of Sum or Count.

Note: The percentage of the subtotal is available only for crosstabs.

Note: The percentage of subtotals is relative to the subset of data that is selected by your filters and ranks.

Percent of Total

Displays the percentage of the total value for the measure on which it is based. You can create a percentage of total only when the source data item has an aggregation of Sum or Count. For example, you might create the percentage of the total value for a measure that contains revenue values. If you create a bar chart of the aggregated measure and a category that contains product lines, then the bar chart shows the percentage of total revenue for each product line.

Note: The percentage of the total value is relative to the subset of data that is selected by your filters and ranks.

Period to Date

Displays the aggregated value for the current time period and all of the previous time periods within a larger time interval. For example, you might derive the year-to-date total for each month.

Note: This derived item is not available if your data source does not contain a date data item that includes the year.

Year to Date

Displays the aggregated value for the current time period and all of the previous time periods within the year. For example, you might derive the year-to-date total for each month.

The year-to-date calculation subsets the data for each year using today's date (where today is evaluated each time you view the report). To use all data for every period, edit the expression for the derived item.

Note: This derived item is not available if your data source does not contain a date data item that includes the year.

Year to Date Growth

Displays the percentage difference between the year-to-date value for the current time period and the year-to-date value for the same time period of the previous

year. For example, you might derive the difference in year-to-date sales between the current month and the same month of the previous year.

The year-to-date calculation subsets the data for each year using today's date (where today is evaluated each time you view the report). To use all data for every period, use a **Period to Date** item or edit the expression for the derived item.

For the month that contains today's date, the data for an earlier year is subset to the same corresponding date.

Note: This derived item is not available if your data source does not contain a date data item that includes the year.

Year over Year Growth

Displays the percentage difference between the current time period and an equivalent time period from the previous year. For example, you might derive the difference in sales between the current month and the same month of the previous year.

For the month that contains today's date, the data for an earlier year is subset to the same corresponding date.

To display a percentage of growth using full periods, use **Percent Difference from Previous Parallel Period** or edit the generated formula.

Note: This derived item is not available if your data source does not contain a date data item that includes the year.

Alternatively, you can use the **Data** tab to create derived data items.

Create a Percentage of Total Using the Data Tab

A measure is required to have a Sum or Count aggregation before you can create a percentage of total.

To create a percentage of total from a measure data item using the **Data** tab:

1   On the **Data** tab in the left pane, right-click on the measure data item that you want to use for the percentage of total.

2   Select **Create ▶ Percent of Total**.

The percentage of total measure data item appears in the list of aggregated data items with a name that is derived from the original name. For example, if the original measure data item name is `Revenue`, then the percentage of total measure data item is displayed as `Revenue (Percent of Total)`. The 🌀 icon identifies the new percentage of total measure data item on the **Data** tab.

Create a Custom Category

You can create a custom category based on either a category or measure data item. A custom category data item is always a category data item with an alphanumeric value.

When you create a custom category from a measure, you can use intervals, ranges, or specific values to group the data.

To create a custom category:

1   On the **Data** tab, right-click on the category or measure data item that you want to use to create the new custom category. Select **New Custom Category**. The New Custom Category window is displayed.
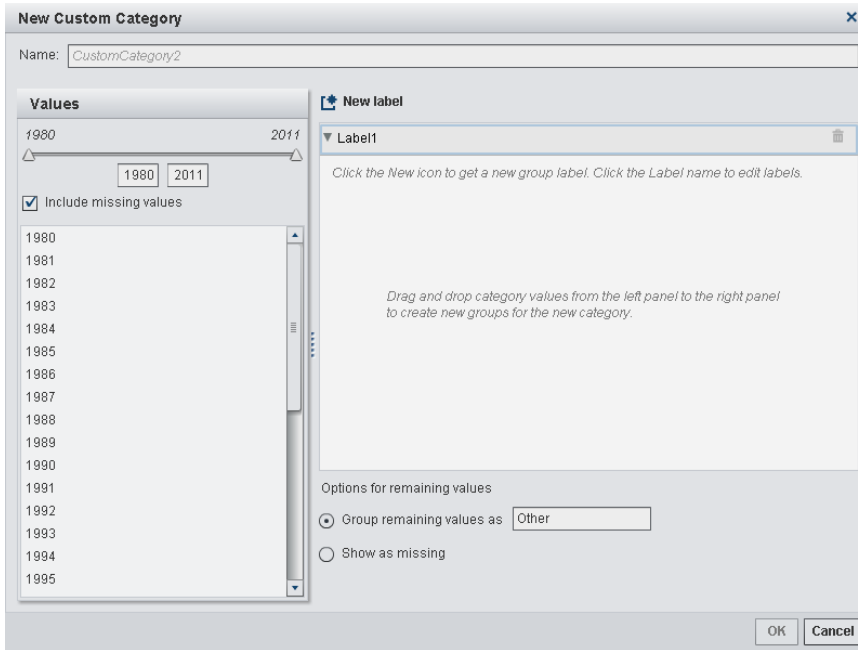
2   Select the **Values** that you want to include. Drag and drop the values onto the right pane, under **New label**.

3   (Optional) Specify the **Options for remaining values**. You can specify a name when you select **Group remaining values as**. The default label for **Group remaining values as** is `Other`. Alternatively, you can specify **Show as missing** or **Show as is**.

   Note: The **Show as is** option is available only for data items that are based on string categories. It is not available for numeric or date values.

4   Specify a **Name** for the new custom category. The default name is `CustomDataItem1`.

5   Click **OK**. The new custom category is displayed on the **Data** tab.

Modify Data Item Properties

Rename a Data Item

To rename a data item:

1   Select a data item on the **Data** tab.

2   In the data item table, select the existing name for the data item, and then enter a new name. The name cannot be used by another data item in the same data source. Your change is saved automatically.

Alternatively, you can right-click the data item, and select **Rename Data Item**.

Modify a Data Item's Role

You can modify a data item's role. For example, you might want to modify a measure data item to be a category data item. A data item's role cannot be modified if that data item is in use in the report or if the data item can have one role only. However, duplicating the data item allows the new data item to have a different role.

Note: You cannot change the role for the frequency data item or the frequency percent data item. You cannot change the role for an aggregated measure.

> **TIP** For geography data items, use the context menu instead of the drop-down menu to change the role. The context menu enables you to specify additional role information for the geography data item.

To modify a data item's role:

1   Select a data item on the **Data** tab.

2   In the data item table, select the existing role name for the data item.

3   Click ▼ to open the drop-down menu. Select **Measure** or **Category**. Your change is saved automatically.

  Note: Category data items cannot be converted into measure data items.

Note: Calculated data items cannot be converted into geography data items.

Modify the Format of a Numeric Measure Data Item or a Date, Datetime, or Time Data Item

You can modify the format of a numeric measure data item or a date, datetime, or time data item. You can also modify the format of a data item with a user-defined format, as long as the user-defined format is based on an underlying numeric value. For more information about user-defined formats, see "Modify User-Defined Formats" on page 368.
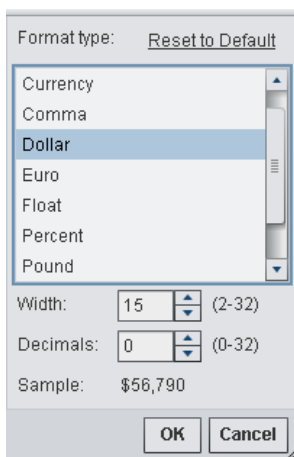
Note: You cannot modify the format of a data item that is being used in a filter, as part of a calculated or aggregated measure, or in a custom category.

To change the format:

1 Select a data item on the **Data** tab.

2 In the data item table, select the existing format. A list is displayed with the **Format type**, **Width**, and **Decimals** (for numeric data items). Make your selections. A sample of your selection is displayed under the list.

Note: There are different format variations available for some format types for date, datetime, and time data items. Select the format variation based on the sample value displayed in the **Format** drop-down list.

Here is an example of the list for a numeric data item:

Here is an example of the list for a date data item:
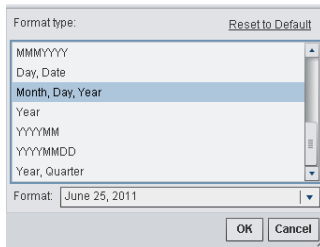
Note: The **Reset to Default** option is displayed only if the format has been changed from the default.

Note: The **Reset to Default** option is available for user-defined format data items after they have been modified to a standard numeric format, as long as the data item is still a category data item.

3 Click **OK** to save your changes.

Modify User-Defined Formats

In the designer, user-defined formats that are defined in the SAS LASR Analytic Server are applied to the results. You can change the format for an underlying numeric data item, but you cannot change the format for an underlying character-based data item. If you change the format for an underlying numeric data item, you can restore the user-defined format by selecting **Reset to Default**.

The **Format** property of the data item displays the name for a user-defined format.

The 🖳 icon identifies a category data item with an active user-defined format on the **Data** tab.

Modify How a Measure Is Aggregated

You can change the aggregation method for a measure in a data source using the **Data** tab or a report object in the canvas.

Note: You need to understand your data because some aggregation methods are not always appropriate. For example, an average of an average is not valid.

To change the aggregation method using the **Data** tab:

1   Select a measure data item on the **Data** tab.

2   In the data item table, select the existing aggregation. A drop-down list is displayed with aggregations. **Sum**, **Average**, **Minimum**, **Maximum**, and **Count** are the available aggregations. Select one, and your change is saved automatically.

Note: All report objects in the report that uses this data item are affected by this change unless you have selected a local aggregation override.
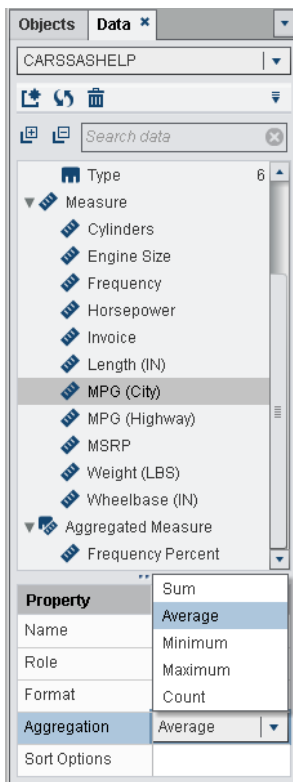
To change the aggregation method when you are working with a report object in the canvas:

1   Choose one of the following:

■   For a list table, right-click on a measure header in the report object.

- For a crosstab, right-click on the header row with the measure.

- For a graph, right-click on the measure name hotspot. Or, you can right-click on a measure data item name on the **Roles** tab.
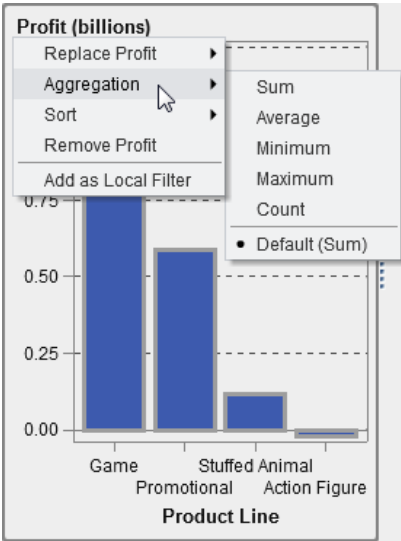
2 Select **Aggregation ▶** *aggregation name*, where *aggregation name* is **Sum**, **Average**, **Minimum**, **Maximum**, or **Count**.

Your change is saved automatically.

Note: This is a local override for only this report object. It does not affect the default aggregation for this data item in other report objects.

Modify the Sort Options for a Category Data Item

Using the data item table on the **Data** tab, you can change the **Sort Options** for a category data item with a user-defined format that is based on an underlying numeric value. The sort options are the following:

Formatted

Uses the formatted character output of the user-defined format and sorts lexically. (For example, the names of the months in the year would sort as April, August, December, February, and so on.) Formatted is the default.

Unformatted

Uses the underlying numeric value and sorts numerically. (For example, if the underlying numeric value for January is 1, for February is 2, for March is 3, and for April is 4, then the months of the year would sort as January, February, March, April, and so on.)

Delete Data Items

You can delete data items that you have created in the designer (for example, calculated data items or duplicated data items) so that they no longer appear on the **Data** tab in the left pane.

You cannot delete a data item that is inside a hierarchy if it reduces the hierarchy to a single level. You cannot delete a data item if it is the last or only reference to a column in the original data source. You cannot delete a data item if it is used in a calculated data item, aggregated measure, a geography data item, or a custom category.

TIP You can hide a data item that you cannot delete by right-clicking on the data item name in the **Data** tab, and then selecting **Hide Data Item**. For more information, see .

To delete a data item:

1   On the **Data** tab in the left pane, right-click the data item that you want to delete.

2   Select **Delete Data Item**.

3   Click **Delete** in the confirmation message that is displayed. The data item is removed from the list of data items, as well as from any report objects, filters, or ranks that were using it.

Working with Geography Data Items

About Geographic Data Items

A geography data item can be useful if your data contains values that are mapped to geographical locations or regions. For example, a geography data item can identify geographic information that is specific to your organization (for example, sales regions, warehouse locations, oil platforms, and so on).

If you change a numeric measure to a geography data item, then it automatically becomes a category data item.

Note: Calculated and custom data items cannot be turned into geography data items.

Using a Predefined Geography Data Item

You can use the geography items that are available in SAS Visual Analytics or custom geography data items. To access these geography data items, select the **Data** tab in the left pane. Right-click the data item that you want to use for the geography data item, and then select **Geography**. Then, select one of the following items:

- **Country or Region Names**

- **Country or Region ISO 2-Letter Codes**

- **Country or Region ISO Numeric Codes**

- **Country or Region SAS Map ID Values**

- **Subdivision (State, Region, Province) Names**

- **Subdivision (State, Region, Province) SAS Map ID Values**

- **US State Names**

- **US State Abbreviations**

- **US ZIP Codes**

Note: For predefined geographic roles, the values of your geography data items must match the lookup values that are used by SAS Visual Analytics. To view the lookup values, see `http://support.sas.com/va63geo`.

> **TIP** To get geographic maps to work with any data set, add a column with the predefined codes to your data set.

Create a Custom Geography Data Item

To create a custom geography data item:

1   On the **Data** tab in the left pane, right-click the data item that you want to use for the geography data item. Select **Geography**, and then select **Custom**. The Geography window is displayed.

2   Select a measure for the **Latitude**. You can also enter the first letter of the name of the latitude column to quickly search for it in the drop-down menu.

3   Select a measure for the **Longitude**. You can also enter the first letter of the name of the longitude column to quickly search for it in the drop-down menu.

4   Select a **Coordinate space**. The choices are **World Geodetic System (WGS84)**, **Web Mercator**, and **British National Grid (OSGB36)**. The default is **World Geodetic System (WGS84)**.

5   Click **OK**. The ⊕ icon identifies the new geography data item on the **Data** tab.

Working with Calculated Items in Reports

About Calculated Data Items

The designer enables you to calculate new data items from your existing data items by using an expression. For example, you might want to calculate a company's profits by subtracting expenses from revenues.

In addition to performing mathematical calculations on numeric values, you can use calculated data items to create date and time values. For example, if your data contains separate categories for *month*, *day*, and *year*, then you can calculate a date value from those categories.

Note: All calculations are performed on unaggregated data. The calculation expression is evaluated for each row in the data source before aggregations are performed. To perform calculations on aggregated data, see "Add a New Aggregated Measure to a Report" on page 378.

A hierarchy can contain calculated data items, as long as they are categories.

Using the designer, you can work with calculated data items or aggregated measures.

Add a New Calculated Data Item to a Report

To add a new calculated item:

1   On the **Data** tab, click ▼ , and then select **New Calculated Item**. The New Calculated Item window is displayed.
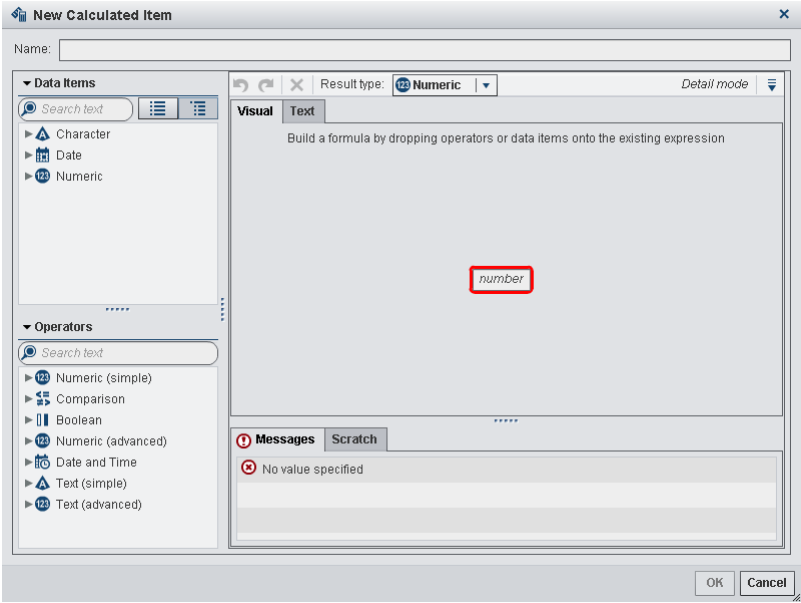
2   Enter a **Name**.

3   Select a **Result type** from the drop-down list. **Numeric** is the default data type.

The following table lists the available result types:

Icon	Result Type
⒜	Character
📅	Date
🕐	Datetime
⑫③	Numeric
🕐	Time

Calculated data items in the designer always default to the following formats, which are based on the data type:

■   Date: DATE9

- ▪ Datetime: DATETIME10

- ▪ Time: TIME8

- ▪ Numeric: COMMA12.2

After you create the new calculated data item, you can change its format using the data item table on the **Data** tab unless it is a character calculated data item. You cannot specify a format for a character calculated data item.

SUM is the default aggregation for new numeric calculated items in the designer. You can change the aggregation for numeric calculated items using the data item table on the **Data** tab.

4   (Optional) Click ▼ to **Show all drop zones**. You can also choose to **Show scratch area** to build temporary expressions.

5   Use the **Visual** tab to build the expression for your calculated data item by dragging **Data items** and **Operators** onto the expression in the right pane. For each rectangular field in the expression, you can insert a data item, an operator, or a specific value.

When you drag items or operators onto your expression, the precise location of the cursor determines where and how the new element is added to the expression. As you drag the new element over the expression, a preview appears that displays how the expression changes if you drop the element at that location.

For example, if your current expression is (`Profit / Revenue`), and you drag and drop the **x - y (subtract)** operator inside the open parenthesis symbol, then the expression changes to (`[number] — ( Profit / Revenue )`). If you drag and drop the operator over the division symbol, then the expression changes to (`Profit — Revenue`), and so on.

Alternatively, you can use the **Text** tab to enter the expression.

There are a large number of operator types available to perform mathematical functions, process datetime values, handle text, and evaluate logical processing such as IF clauses. For more information, see Appendix 3, "Operators for Data Expressions," on page 569.

6 Click **OK**. The new calculated data item appears in the **Data** tab. The ▓, ▓, or ▓ icons identify the new calculated data item on the **Data** tab.

   Note: The ▓ icon is displayed if you change a calculated numeric measure to a category data item. It is also displayed for a character calculated data item (which is always a category).

Edit a Calculated Data Item

To edit a calculated data item:

1 Right-click on a calculated data item on the **Data** tab, and select **Edit Calculated Item**. The Edit Calculated Item window is displayed.
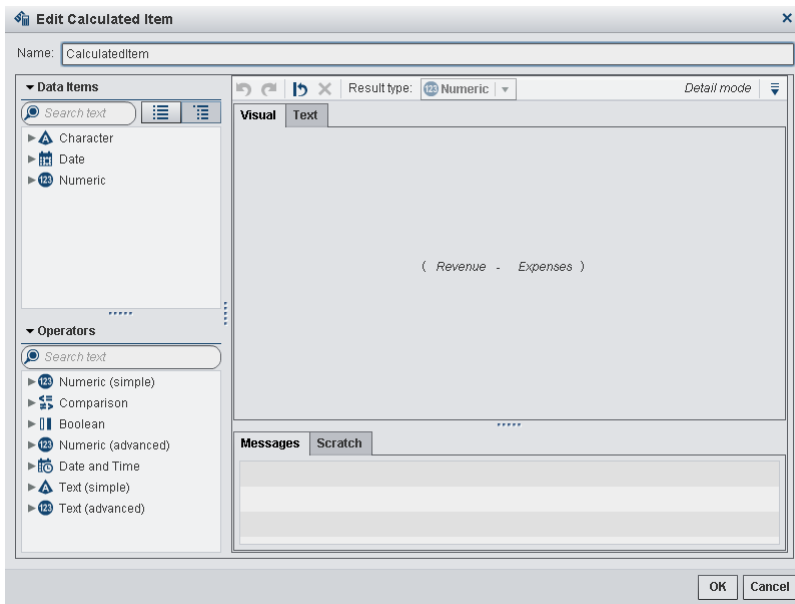
2 Modify the **Data Items** and **Operators** for the calculated data item as needed. For information about the operators and conditions that are available, see Appendix 4, "Conditions for Filters," on page 599.

3 Click **OK**.

You can duplicate, rename, hide, and delete calculated data items using the same steps as any other data item. If a calculated data item is used inside another calculated data item, then it cannot be removed.

> **TIP** You can cut and paste from the **Text** area between different reports, as well as between the explorer and the designer, to transfer calculations or to e-mail them to others. For more information, see "Editing a Data Expression in Text Mode" on page 565.

Add a New Aggregated Measure to a Report

Aggregated measures enable you to calculate new data items by using aggregated values. For example, you might want to calculate a company's profit margin by region by taking the aggregated sum of the profit for all of the stores in a region group and dividing it by the aggregated sum of the revenue for all of the stores in that same region group. Aggregations are evaluated as part of a calculated expression.

To add an aggregated measure:

1   On the **Data** tab, click ⩔ , and then select **New Aggregated Measure**. The New Aggregated Measure window is displayed.
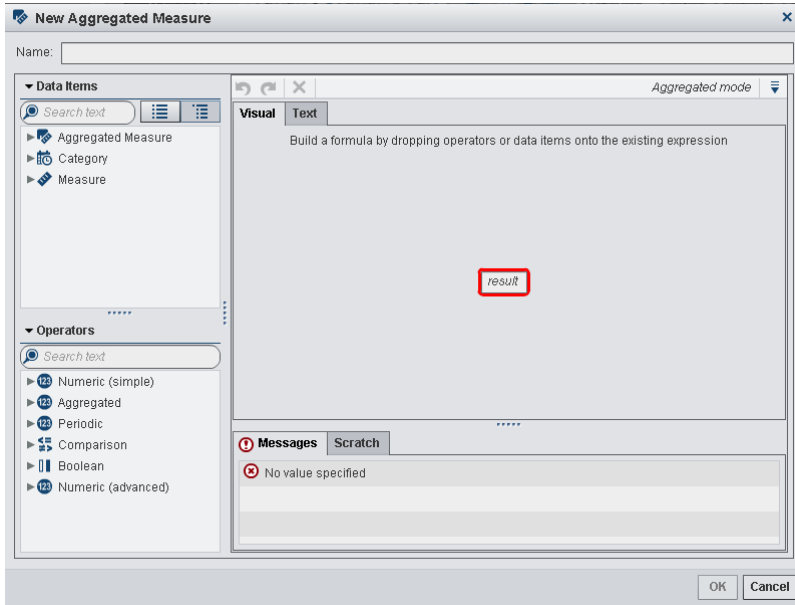
2   Enter a **Name**.

3   (Optional) Click ☰ to **Show all drop zones**. You can also choose to **Show scratch area** to build temporary expressions.

4   Use the **Visual** tab to build the expression for your aggregated measure by dragging and dropping **Data Items** and **Operators** onto the expression in the right pane. For each field in the expression, you can insert a data item, an operator, or a specific value.

When you drag and drop data items and operators onto the expression, the precise location of the cursor determines where and how the data item or operator is added to the expression. As you drag the new element over the expression, a preview appears, which displays how the expression would change if you drop the element at the current location.

Alternatively, you can use the **Text** tab to enter the expression.

There are a large number of operator types available to perform mathematical functions and evaluate logical processing such as IF clauses.See Appendix 3, "Operators for Data Expressions," on page 569.

5 For each data item in your expression, select an aggregation type. By default, **Sum** is used for measures and **Distinct** is used for categories. To select a new aggregation type, drag and drop an aggregated operator from the **Operators** list onto the aggregation type in the expression. See Appendix 3, "Operators for Data Expressions," on page 569 for a list of the aggregated operators that are available.

6 For each aggregation in your expression, select the aggregation context. A drop-down list beside each aggregation enables you to select one of the following context values:

ByGroup
> calculates the aggregation for each subset of the data item that is used in a visualization. For example, in a bar chart, an aggregated measure with the **ByGroup** context calculates a separate aggregated value for each bar in the chart.

ForAll
> calculates the aggregation for the entire data item (after filtering). For example, in a bar chart, an aggregated measure with the **ForAll** context uses the same aggregated value (calculated for the entire data item) for each bar in the chart.

By using the **ForAll** and **ByGroup** contexts together, you can create measures that compare the local value to the global value. For example, you might calculate the difference from mean by using an expression such as the following:

```
Avg ByGroup(X) - Avg ForAll(X)
```

For more information, see "Periodic Operators" on page 579.

7 Click **OK**. The new aggregated measure appears on the **Data** tab. The 🪁 icon identifies the new aggregated measure.

Edit an Aggregated Measure

To edit an aggregated measure:

1 Right-click on an aggregated measure on the **Data** tab, and select **Edit Aggregated Measure**. The Edit Aggregated Measure window is displayed.

2   Modify the **Data Items** and **Operators** for the aggregated measure as needed. For information about the operators and conditions that are available, see Appendix 4, "Conditions for Filters," on page 599.

3   Click **OK**.

You can duplicate, rename, hide, and delete aggregated measures using the same steps as any other data item.

> **TIP** You can cut and paste from the **Text** area between different reports, as well as between the explorer and the designer, to transfer calculations or to e-mail them to others. For more information, see "Editing a Data Expression in Text Mode" on page 565.

Working with Data Role Assignments

After you have selected a report object and a data source, the **Roles** tab in the right pane shows which data items have been assigned to which role.

Modify Data Role Assignments in Report Objects

Note: You cannot change data role assignments for any report object that is imported from the explorer.

To modify data role assignments:

1   Select a report object in the canvas that has data items assigned.

2   Click the **Roles** tab in the right pane.

3   Click ▼ beside the role that you want to edit. For data roles that allow multiple data items, the menu items that are available can vary.

If you select a data item in **Roles**, then the **Aggregation** (if appropriate), **Add**, **Remove**, and **Replace** menu items are displayed. You can also right-click on a data item, and the actions specific to that data item are displayed.

Different data roles are available for different report objects.

In addition to the basic data roles, you can assign additional data roles. Here is a list of all the data roles available for each report object:

Report Object	Basic Data Roles	Additional Data Roles
Tables		
List table	**Columns**	None
Crosstab	■ **Columns** ■ **Rows**	**Measures**
Graphs		
Bar chart	■ **Category** ■ **Measures**	■ **Group** ■ **Lattice columns** ■ **Lattice rows** ■ **Data tip values** ■ **Animation**
Targeted bar chart	■ **Category** ■ **Measure** ■ **Target**	■ **Lattice columns** ■ **Lattice rows** ■ **Data tip values**
Waterfall chart	■ **Category** ■ **Measure**	■ **Lattice columns** ■ **Lattice rows** ■ **Data tip values**
Line charts	■ **Category** ■ **Measures**	■ **Group** ■ **Lattice columns** ■ **Lattice rows** ■ **Data tip values**

Report Object	Basic Data Roles	Additional Data Roles
Pie chart	■ **Category** ■ **Measures**	■ **Group** ■ **Lattice columns** ■ **Lattice rows** ■ **Data tip values**
Scatter plot	■ **X axis** ■ **Y axis**	■ **Color** ■ **Lattice columns** ■ **Lattice rows** ■ **Data tip values**
Time series plot	■ **Time axis** ■ **Measure (line)**	■ **Group** ■ **Data tip values**
Bubble plot	■ **X axis** ■ **Y axis** ■ **Size**	■ **Group** ■ **Color** ■ **Lattice columns** ■ **Lattice rows** ■ **Data tip values** ■ **Animation**
Treemap	**Tile**	■ **Size** ■ **Color** ■ **Data tip values**
Dual axis bar chart	■ **Category** ■ **Measure (bar)** ■ **Measure (bar 2)**	■ **Lattice columns** ■ **Lattice rows** ■ **Data tip values**
Dual axis line chart	■ **Category** ■ **Measure (line)** ■ **Measure (line 2)**	■ **Lattice columns** ■ **Lattice rows** ■ **Data tip values**

Report Object	Basic Data Roles	Additional Data Roles
Dual axis bar-line chart	■ **Category** ■ **Measure (bar)** ■ **Measure (line)**	■ **Lattice columns** ■ **Lattice rows** ■ **Data tip values**
Dual axis time series plot	■ **Time axis** ■ **Measure (line)** ■ **Measure (line 2)**	**Data tip values**
Gauges		
Bullet Slider Thermometer Dial Speedometer	**Measure**	■ **Target** ■ **Group** ■ **Data tip values**
Controls		
Drop down list List Button bar Text input	**Category**	**Frequency**
Range slider	**Measure/Date**	None
Other		
Text Image Stored process Vertical container Horizontal container Stack container	None	None

Report Object	Basic Data Roles	Additional Data Roles
Geo bubble map	**Geography**	◼ **Size** ◼ **Color** ◼ **Data tip values** ◼ **Animation**
Geo region map	**Geography**	◼ **Color** ◼ **Data tip values**

Here are definitions for some of the additional data roles:

Animation

animates the date category that you assign. You can specify a **Loop** for the animation, and you can specify the **Speed** of the animation. For more information, see "Add Animation to Bar Charts, Bubble Plots, and Geo Bubble Maps" on page 386.

Note: For a bubble plot only, you must assign the **Group** role to enable the **Animation** role. If you remove the **Group** role after assigning animation, the animation control is disabled.

Data tip values

enables you to add more measures to your data tips.

Group

groups the data based on the category data item that you assign.

Lattice columns

creates a lattice of charts with a column for each value of the category data item that you assign.

Lattice rows

creates a lattice of charts with a row for each value of the category data item that you assign.

Add Animation to Bar Charts, Bubble Plots, and Geo Bubble Maps

An animated bubble plot displays the changes in your data values over time. Each frame of the animation represents a value of the datetime data item that is assigned to the **Animation** data role.

For example, if you assign a category with the **Year** format to the **Animation** data role, then each frame of the animation displays a bubble plot of your data for a specific year.

A bar chart cannot have animation if it has a **Group**, **Lattice Rows**, or **Lattice Columns** role, or if it has a **Stacked** grouping style property set.

To add animation:

1   Select an existing bar chart, bubble plot, or geo bubble map, or create a new one.

2   Add the required roles:

■ For a bubble plot, on the **Role** tab, assign a data item to the **Group** data role.

   Note: If you remove the **Group** data role after assigning animation, the animation control is disabled.

■ For a bar chart or geo bubble map, on the **Role** tab, assign the other required roles.

3   Assign a data item with a datetime format to the **Animation** data role.

When an animated bubble plot is displayed, a set of animation controls appear at the bottom of the report object.

Table 33.2   Animation Controls

Task	Action
Start the animation.	Click ►.
Go to the previous animation frame.	Click ◄I.

Task	Action
Go to the next animation frame.	Click ▮▶.
Jump to a specific animation frame.	Use the slider.
Specify whether to repeat the animation.	Select or deselect **Loop**.
Select the animation speed.	Use the **Speed** slider.
Track the movement of a specific bubble.	Click the bubble that you want to track.

Note: If you select a bubble to track, the selected bubble is not highlighted.

Remove Data Role Assignments from Report Objects

To remove data items from their assigned data roles in a specific report object:

1   Right-click the report object in the canvas. A menu is displayed.

Here is an example of a bar chart with the object menu:
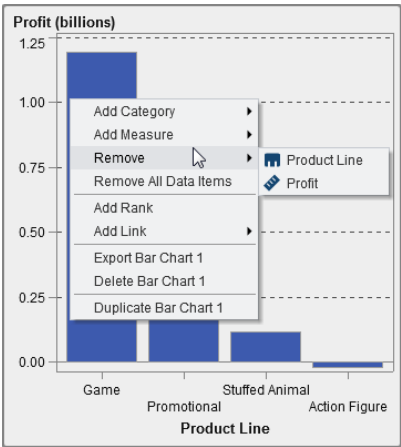
2   Select **Remove All Data Items**. All data items are removed from the assigned data roles and the report object turns gray. The ⊖ status icon appears in the lower right corner to let you know that the required data roles are not assigned.

Here is an example of what a bar chart looks like after all the data items have been removed from their assigned data roles:
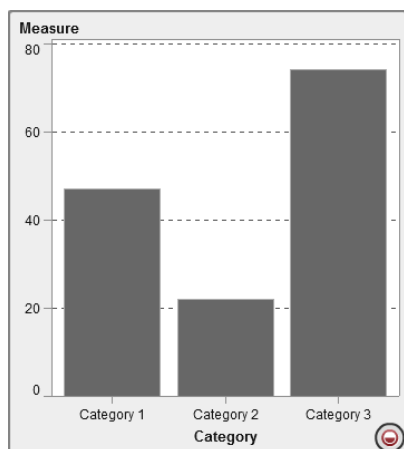
Sorting Data in Reports

How Sorting Can Help with Analysis

Information can be easier to understand when it appears in an intentional order. Applying a sort order to one or more data items in the designer enables you to arrange rows and columns in tables and axis labels on charts in some order, such as alphabetically or highest to lowest numerically. Interactively changing the order of data can provide you with a different perspective that often facilitates valuable insight. For example, in a report, sales employees who are initially arranged alphabetically can be re-sorted by sales amount.

Ranking can help reduce the amount of visible data and is often used in combination with sorting. For more information, see "Adding a New Rank" on page 460.

Sorting Data in a List Table

To sort values in a list table, right-click on a column heading and select **Sort**. Then, select either **Ascending** or **Descending**.

Display 33.1   *Sort Menu Items for List Tables*

Product Li.	▲	Product	Revenue	Expenses	Profit
Game		Replace Product Line ▸	2,511,801	117,465,423	355,046,378
Game		Remove Product Line	36,733,440	67,173,845	329,548,595
Game		Sort ▸	✓ Ascending	,961,649	608,694,146
Promotiona		Delete Rank	Descending	13,900,367	142,075,951
Promotiona		Add Link ▸	1,717,507	12,359,151	99,358,356
Promotiona		Add Display Rule	3,917,946	5,221,934	98,696,012
Promotiona		Add Alert...	8,485,882	113,686,253	164,799,628
Promotiona		Add Sparkline			
Promotiona		Export List Table 1	3,601,637	55,565,403	108,036,234
Stuffed Anir		Delete List Table 1	9,914,069	10,164,874	29,749,196
Stuffed Anir		Duplicate List Table 1	9,822,386	10,195,261	29,627,124
All Other		Move List Table 1 to ▸	9,573,271	542,877,077	-83,303,805

Note: You can also press the spacebar to sort a single column in a table.

TIP Use the Ctrl key to select and sort multiple columns.

Sorting Data in a Crosstab

To sort by category or measure values in a crosstab, right-click on a column or row heading, and select **Sort**. Then, select either **Ascending** or **Descending**.

Display 33.2 *Sort Menu Items for a Crosstab*
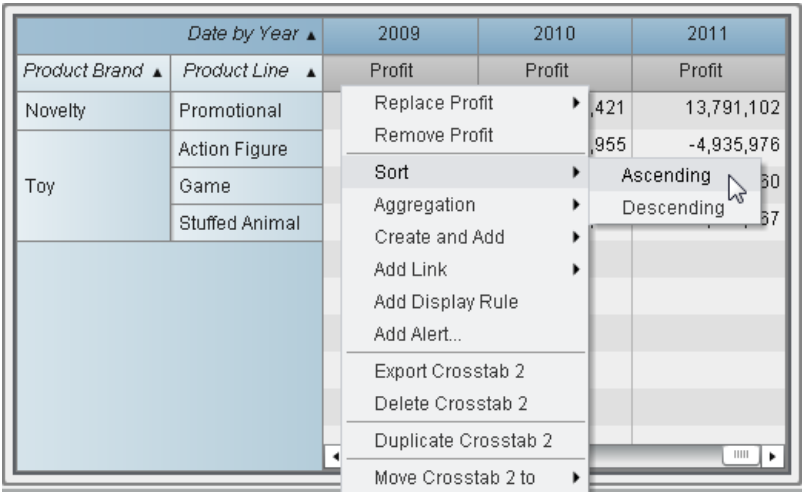
You can also sort categories, rather than measures, by clicking the category name.

Sorting Data in a Graph

To sort by values in a graph, right-click on a measure name or category name, and select **Sort**. Then, select either **Ascending** or **Descending**.

Display 33.3 *Sort Menu Items for a Graph*
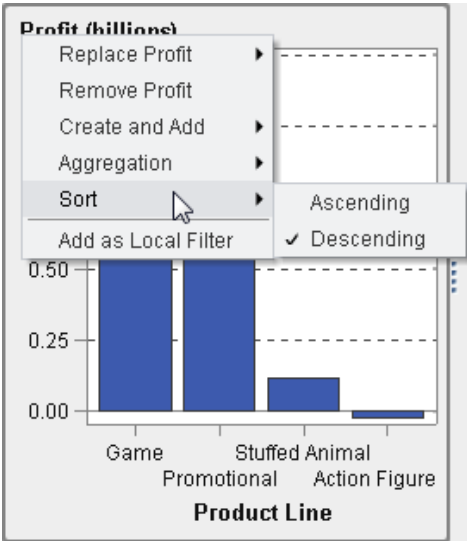
Exporting Data from Report Objects

Overview of Exporting from the Designer

All users who have the Export Data capability can use the designer to export data from report objects to Microsoft Excel format for future viewing or printing. This exported output can be saved locally on disk, and then opened in Microsoft Excel. Or you can choose to create a Tab-Separated Values (*.tsv) or Comma-Separated Values (*.csv) data file.

When you export a graph from the designer, you are exporting the data, not the visual graph representation.

Not all report objects in the designer support the exporting feature. For example, you cannot export data from gauges. If the export feature is not available for a particular report object, the **Export <reportObjectName>** menu item does not appear when you right-click on the object.

Export Data from a List Table

To export data from a list table:

1   If the list table that you want to export is not already selected, select it.

2   Right-click on the list table, and select **Export <listTableName>**, where `<listTableName>` is the name of the report object. The Export or Save As window is displayed.

3 Choose to export **All rows** or a range of **Rows**.

4 Choose to export **All columns** or **Selected columns**. If you choose **Selected columns**, select the check box (or check boxes) to the left of the column (or columns) that you want to export. At least one column is required. If you do not select a column, then a message is displayed, and the Export or Save As window cannot be closed.

5 (Optional) To choose whether the exported data is formatted, either select or clear the **Formatted data** check box. This check box is selected by default.

6 If you select **Export to**, then the only option is **Excel 2007 Workbook (*.xlsx)** to create a Microsoft Excel spreadsheet.

7 If you select **Save as**, choose either **Tab-Separated Values (*.tsv)** or **Comma-Separated Values (*.csv)** to create a data file.

8 Click **OK**.

9 When you are prompted, choose either to open the file or to save it.

Export Data from a Crosstab

To export detail data from a crosstab:

1 If the crosstab that you want to export is not already selected, select it.

2 Right-click on the crosstab, and select **Export <crosstabName>**, where
   `<crosstabName>` is the name of the report object. The Export or Save As window
   is displayed.

3 Choose to export **All rows** or a range of **Rows**.

4 Choose to export **All columns** or **Selected columns**. If you choose **Selected
   columns**, select the check box (or check boxes) to the left of the column (or
   columns) that you want to export. At least one column is required. If you do not
   select a column, then a message is displayed, and the Export or Save As window
   cannot be closed.

5 (Optional) To choose whether the exported data is formatted, either select or clear
   the **Formatted data** check box. This check box is selected by default.

6 If you select **Export to**, then the only option is **Excel 2007 Workbook (*.xlsx)** to create a Microsoft Excel spreadsheet.

7 If you select **Save as**, choose either **Tab-Separated Values (*.tsv)** or **Comma-Separated Values (*.csv)** to create a data file.

8 Click **OK**.

9 When you are prompted, choose either to open the file or to save it.

Export Data from a Graph

Note: You cannot export the actual visual graph. Only the data can be exported.

To export data from a graph:

1 If the graph that you want to export is not already selected, select it.

2 Right-click on the graph, and select **Export <graphName>**, where `<graphName>` is the name of the report object. The Export or Save As window is displayed.
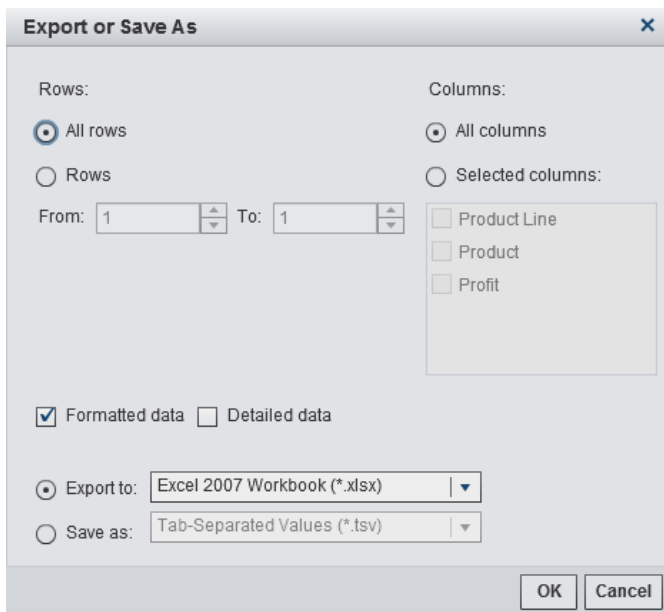
3 Choose to export **All rows** or a range of **Rows**.

4   Choose to export **All columns** or **Selected columns**. If you choose **Selected columns**, select the check box (or check boxes) to the left of the column (or columns) that you want to export. At least one column is required. If you do not select a column, then a message is displayed, and the Export or Save As window cannot be closed.

5   (Optional) Choose whether the exported data is formatted, either select or clear the **Formatted data** check box. This check box is selected by default.

6   (Optional) Choose whether **Detailed data** is exported for the graph. If you select **Detailed data**, you can select the columns that you want to export.

   Note: The **Detailed data** option is not available for imported report objects or custom graph objects that have multiple data definitions.

7   If you select **Export to**, then the only option is **Excel 2007 Workbook (*.xlsx)** to create a Microsoft Excel spreadsheet.

8   If you select **Save as**, choose either **Tab-Separated Values (*.tsv)** or **Comma-Separated Values (*.csv)** to create a data file.

9   Click **OK**.

10 When you are prompted, choose either to open the file or to save it.

34

Working with Display Rules for Reports

Overview of Display Rules

Display rules include all types of highlighting of report objects. They provide a flexible structure to specify conditions. There are several types of display rules. Display rules enable conditions to be shared across objects, but not all display rules apply across all report object types.

SAS Visual Analytics Designer (the designer) provides the following interfaces for display rules:

- The **Display Rules** tab in the right pane enables you to populate intervals, add intervals, or add color-mapped values for the report object, which is currently selected in the canvas. You can use this pane to specify both report-level or object-level display rules, depending on what you have selected in the canvas.

- The **Shared Rules** tab enables you to create a new display rule for a gauge, which is used by other gauges to designate intervals and colors for ranges. You can also edit or delete an existing shared display rule. These rules are shared across multiple gauges and can be created at any time.

For list tables and graphs, you can create display rules that reference a measure that is not included in the currently displayed report object. This feature is not available for crosstabs.

Note: Color-mapped display rules are not allowed if there are two or more overlays of a bar, line, or time series in graphs that have multiple overlays (for example, a dual axis bar chart or a custom graph).

There are limitations to where display rules can be used in the designer. The following report objects do not support display rules:

- dual axis bar charts

- dual axis bar-line charts

- dual axis line charts

- dual axis time series plots

The following report objects that are imported from SAS Visual Analytics Explorer (the explorer) or created in SAS Visual Analytics Graph Builder (the graph builder) do not support display rules:

- heat maps (which are imported from the explorer)

- band plots (which are created in the graph builder)

- step plots (which are created in the graph builder)

- schedule charts (which are created in the graph builder)

Adding Report-Level Display Rules

To add a report-level display rule:

1   Without any report objects or the section selected in the canvas, click the **Display Rules** tab in the right pane. Click **New**. The Add New Display Rule window is displayed.

2   Enter a value for the display rule in the field.

3   Select a color for the display rule.

4   (Optional) Repeat the steps for entering a value and selecting a color.

5   (Optional) Select the **Other** check box. Then, select a color so that any of the other categories that do not have a color will have the one that you just selected.

    In the following example, values and colors have been specified for each region in a list table. And a color has been selected for the **Other** check box.
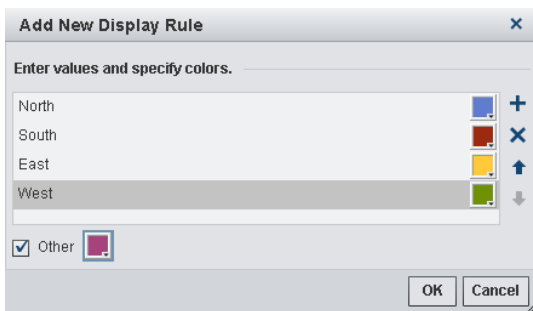
6   Click **OK**. The report objects in the report update with the new display rule. And, the display rule appears on the **Display Rules** tab in the right pane.

The following table illustrates the report-level display rules that were defined in the Add New Display Rule window above.

Facility Region	Product Line	Profit	Expenses	
East	Action Figure	-10,515,368	29,014,382	
East	Game	5,328,719	1,900,170	
East	Promotional	81,456,122	45,807,236	
East	Stuffed Animal	1,544,003	8,284,070	
North	Action Figure	-21,211,851	60,883,053	
North	Game	220,653,512	110,853,384	
North	Stuffed Animal	16,039,542	56,849,706	
South	Action Figure	9,202,487	165,466,767	
South	Game	607,956,582	241,557,659	
South	Stuffed Animal	51,908,425	39,659,003	
West	Action Figure	3,453,239	26,026,052	
West	Game	360,141,294	123,498,715	

Profit (millions)

Adding Table-Level Display Rules

You can add three different types of display rules to tables. You cannot create a display rule to highlight dates in a table.

Add Display Rules to a List Table Using an Expression

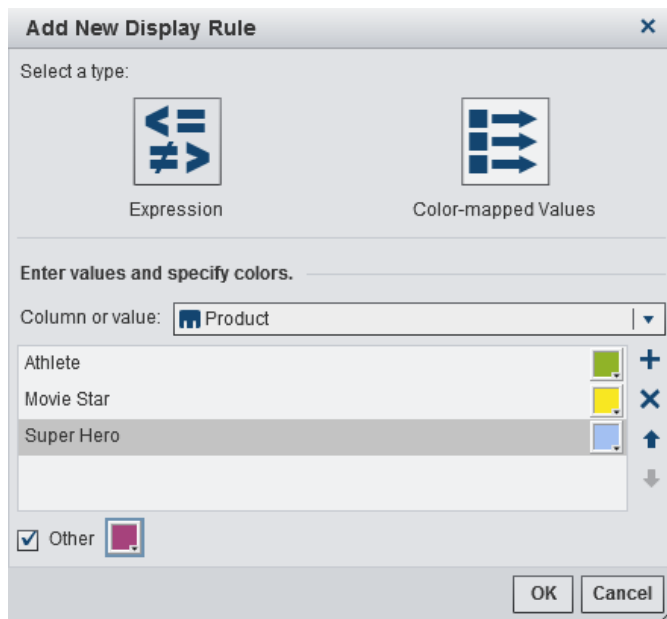To specify a new display rule for a table using an expression:

1 If it is not already selected, select the table in the canvas that you want to update.

2 In the right pane, click the **Display Rules** tab. Click **New**. The Add New Display Rule window is displayed.

3 Click **Expression**. The Add New Display Rule window expands to show the details for the expression.

4 On the **Expression Details** tab, select the **Column** or any measure value.

5 Select the **Operator**. You can select **=**, **< >**, **BetweenInclusive**, **<**, **<=**, **>**, **>=**, **Missing**, or **NotMissing**. The default is **>**.

6 Enter or select a **Value**.

7   Modify the style, size, and color of the font.

8   Select the row or column in the **Applies to** drop-down list.

9   (Optional) On the **Alert Options** tab, select the **Create an alert based on this rule**
    check box. For more information about alerts, see "Working with Alerts for Report
    Objects" on page 339.
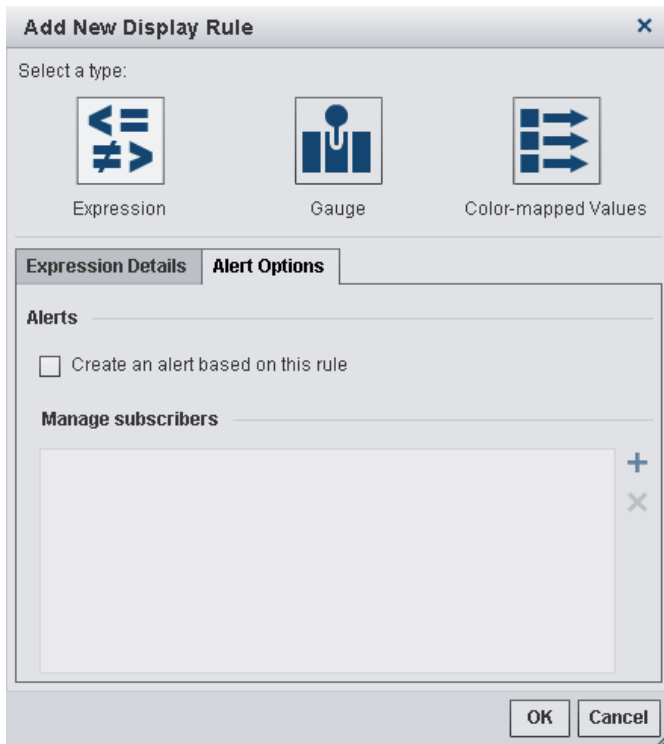
Use the **Manage subscribers** list to add the names of users that you want notified
when the alert condition is met.

10 Click **OK**. The table updates with the new display rule. The display rule appears on
    the **Display Rules** tab in the right pane.

Here is an example of a display rule that uses an expression:
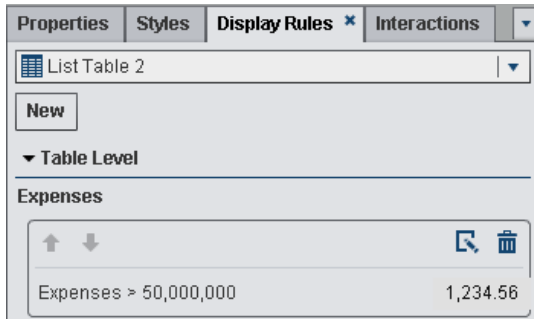
11 (Optional) Click on the **Display Rules** tab to edit the new display rule.

Add Display Rules to a List Table Using a Gauge

To specify a new display rule for a table using a gauge:

1   If it is not already selected, select the table in the canvas that you want to update.

2   In the right pane, click the **Display Rules** tab. Click **New**. The Add New Display Rule window is displayed.

3   Click **Gauge**. The Add New Display Rule window expands to show the details for the gauge.

4   Select a **Gauge type** from the drop-down list. Your choices are a bullet, an icon, a slider, or a thermometer. **Icon** is the default.

5   Using the **Based on column** drop-down list, specify which column the rule should be based on in the report.

6   Specify where the gauge should appear in the column for the **Cell placement**. Your choices are **Left of text**, **Right of text**, or **Replace text**.

7   Define the intervals and colors for the rule:

   ■   Enter the individual values for the intervals and then select a color. You can click the operator between the intervals to change it.

   ■   Click ⊕ to automatically populate the intervals. The Populate Intervals window is displayed.

You can specify **Number of intervals**, **Lower bounds**, and **Upper bounds**. Click **OK**.

8   Specify the **Column** in which you want the gauge to be displayed

9   Click **OK**. The table updates with the new display rule. The display rule appears on the **Display Rules** tab in the right pane.

Here is an example of an automatically populated display rule that uses an icon:

Properties	Styles	Display Rules ✕	Roles	▼

▦ List Table 1   ▼

New

▼ Table Level

Expenses

⬏ 🗑

▢ to the left of text, with Expenses as x

5,221,934	<= x <	68,135,172	🟥
68,135,172	<= x <	131,048,411	🟨
131,048,411	<= x <=	193,961,649	🟩
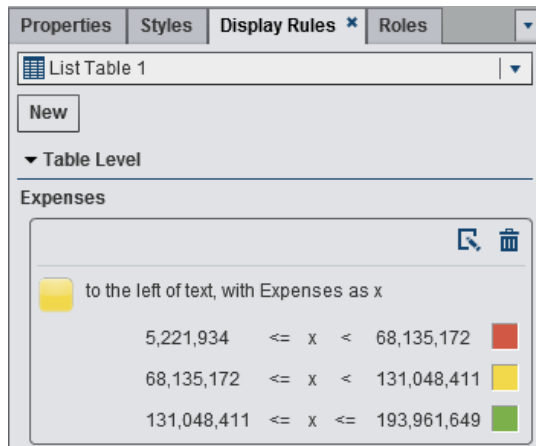
Here is a list table with the display rules applied:

Product Li... ▲	Product	Revenue	Expenses	Pr ▲
Action Figure	Firefighter	37,830,429	🟥 22,383,218	15,447,2
Action Figure	Athlete	37,607,425	🟥 22,282,489	15,324,9
Action Figure	Super Hero	37,593,488	🟥 22,228,489	15,364,9
Action Figure	Musician	37,400,758	🟥 22,201,224	15,199,5
Action Figure	Soldier	37,054,862	🟥 22,046,588	15,008,2
Action Figure	.	0	🟨 126,051,9...	-126,051,9
Action Figure	Police	37,365,735	🟥 21,939,853	15,425,8
Action Figure	Movie Star	37,466,062	🟥 22,256,456	15,209,6
Game	Puzzle	472,511,801	🟨 117,465,4...	355,046,3
Game	Card	396,722,440	🟥 67,173,845	329,548,5
Game	Board	802,655,795	🟩 193,961,6...	608,694,1
Game	.	0	🟨 99,209,012	-99,209,0
Promotional	Coffee Cup	111,717,507	🟥 12,359,151	99,358,3
Promotional	Backpack	278,485,882	🟨 113,686,2...	164,799,6
Promotional	.	0	🟥 23,089,265	-23,089,2
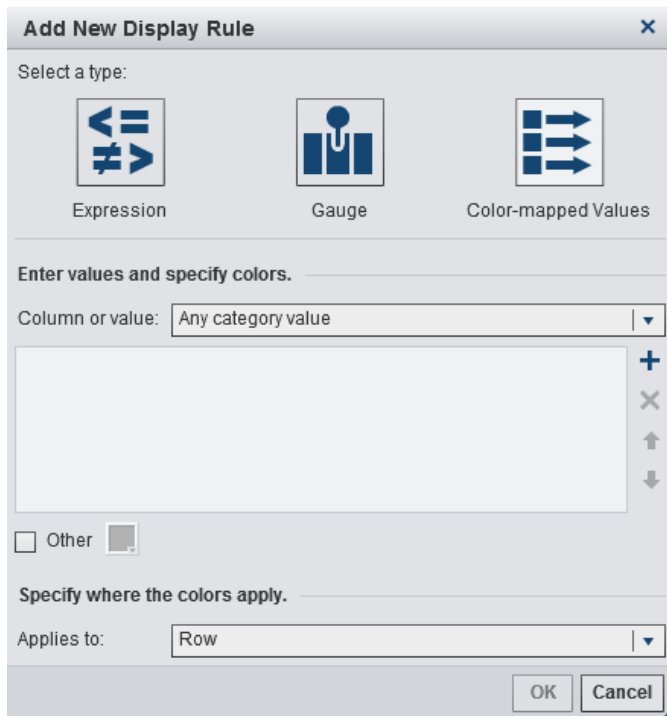
Note: If a value falls outside the bounds of the display rule intervals, no icon is displayed in the list table

10 (Optional) Click on the **Display Rules** tab to edit the new display rule.

Add Display Rules to a List Table Using Color-Mapped Values

To specify a new display rule for a table using color-mapped values:

1 If it is not already selected, select the table in the canvas that you want to update.

2 In the right pane, click the **Display Rules** tab. Click **New**. The Add New Display Rule window is displayed.

3 Click **Color-mapped values**. The Add New Display Rule window expands to show the details for the color-mapped values.

4   Select the **Column or value** to which you want to apply the display rule.

5   Click in the box to enter a value for the display rule.

6   Select a color for the display rule.

7   (Optional) Repeat the steps for entering a value and selecting a color.

8   (Optional) Select the **Other** check box. Then, select a color.

9   Using the list, specify where you want to apply the colors.

In the following example, values and colors have been specified for products in a list table.
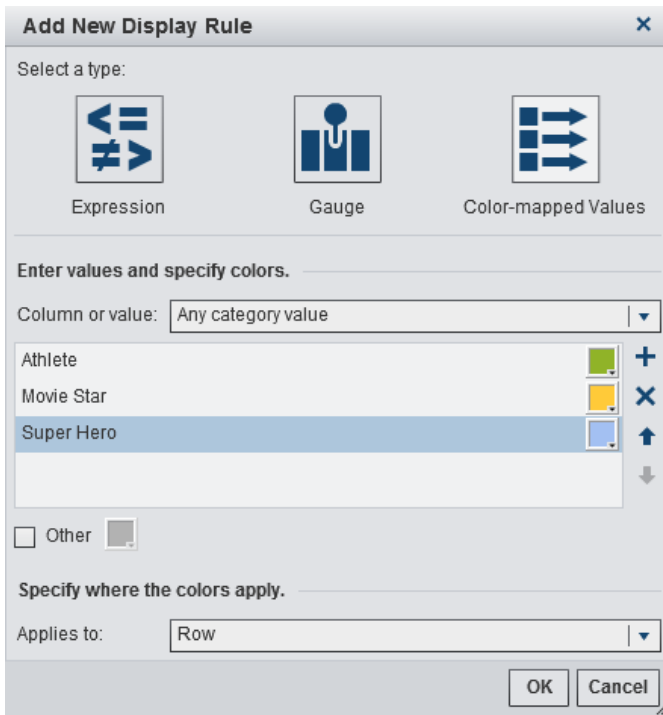
10 Click **OK**. The table updates with the display rules.

The display rule appears on the **Display Rules** tab in the right pane.

11 (Optional) Click on the **Display Rules** tab to change the table-level display rule to a report-level display rule.

Adding Graph-Level Display Rules

You can add display rules to graphs. Note that you cannot create a display rule to highlight dates in a graph.

Add Display Rules to a Graph Using an Expression

Note: Display rules can be added to a waterfall chart only if the **Color by response sign** check box is selected on the **Properties** tab in the right pane.

To specify a new display rule for a graph:

1 If it is not already selected, select the graph in the canvas that you want to update.

2 In the right pane, click the **Display Rules** tab. Click **New**. The Add New Display Rule window is displayed.

3 Click **Expression**. The Add New Display Rule window expands to show the details for the expression.

4   On the **Expression Details** tab, select the **Column** or any measure value.

5   Select the **Operator**. You can select **=**, **< >**, **BetweenInclusive**, **<**, **<=**, **>**, **>=**, **Missing**, or **NotMissing**. The default is **>**.

6   Enter or select a **Value**.

7   Modify the **Style**, which includes color and graph or background.

   Note: Background is available only for bar charts, waterfall charts, line charts, scatter plots, time series plots, and bubble plots. It is also available for needle plots, which are created in the SAS Visual Analytics Graph Builder (the graph builder).

8   (Optional) On the **Alert Options** tab, select the **Create an alert based on this rule** check box. For more information about alerts, see "Working with Alerts for Report Objects" on page 339.
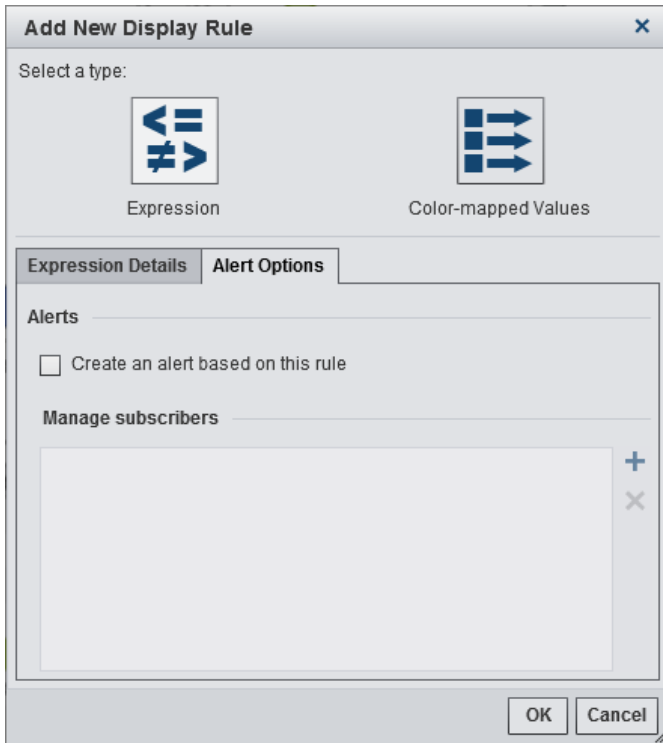
Use the **Manage subscribers** list to add the names of users that you want notified when the alert condition is met.

9 Click **OK**. The graph updates with the new display rule. The display rule appears on the **Display Rules** tab in the right pane.

Here is an example of a display rule that uses an expression:
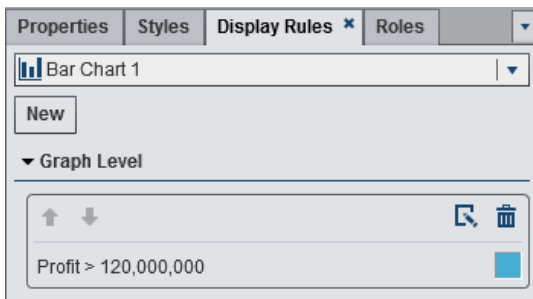
10 (Optional) Click on the **Display Rules** tab to edit the new display rule.

Add Display Rules to a Graph Using Color-Mapped Values

Note: Display rules can be added to a waterfall chart only if the **Color by response sign** check box is selected on the **Properties** tab in the right pane.

To specify a new display rule for a graph using color-mapped values:

1 If it is not already selected, select the graph in the canvas that you want to update.

2 In the right pane, click the **Display Rules** tab. Click **New**. The Add New Display Rule window is displayed.

3 Click **Color-mapped Values**. The Add New Display Rule window expands to show the details for the color-mapped values.
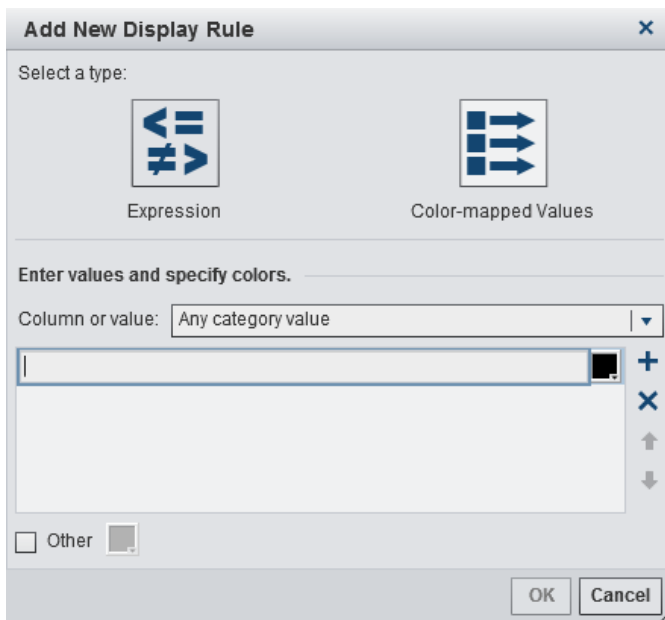
4 Select the **Column or value** to which you want to apply the display rule.

5 Click in the box to enter a value for the display rule.

6 Select a color for the display rule.

7 (Optional) Repeat the steps for entering a value and selecting a color for the display rule.

8 (Optional) Select the **Other** check box. Then, select a color.

9 Using the list, specify where you want to apply the colors.

In the following example, values and colors have been specified for products in a graph.

10 Click **OK**. The graph updates with the display rules.

The display rule appears on the **Display Rules** tab in the right pane.
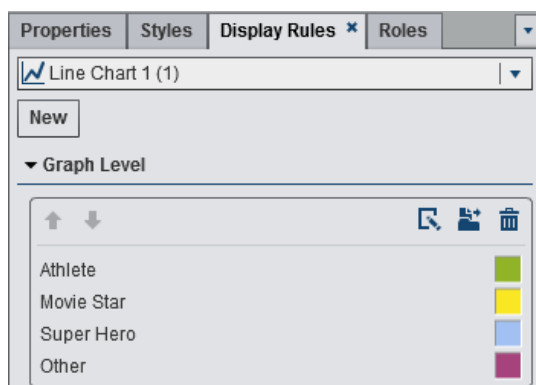
11 (Optional) Click ![icon] on the **Display Rules** tab to change the graph-level display rule to a report-level display rule.

Adding Display Rules to a Gauge

A display rule is used by a gauge to designate intervals and colors for ranges.

To specify a new display rule for a gauge:

1 If it is not already selected, select the gauge in the canvas that you want to update.

2 In the right pane, click the **Display Rules** tab.

Here is an example of the **Display Rules** tab for a dial gauge:
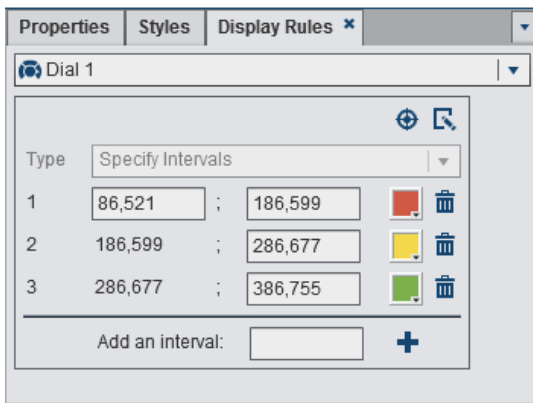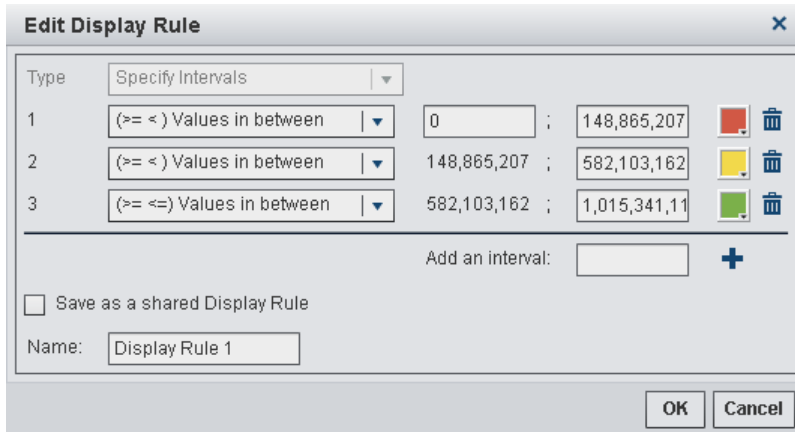
3 Specify (or modify) the display rules for the gauge. You can populate intervals, edit the display rule, and specify the **Type**.

To automatically populate the intervals:

a Click ![icon] to display the Populate Intervals window.

b Enter a number for **Target value**.

c Click **OK**. The gauge updates with the new display rule.
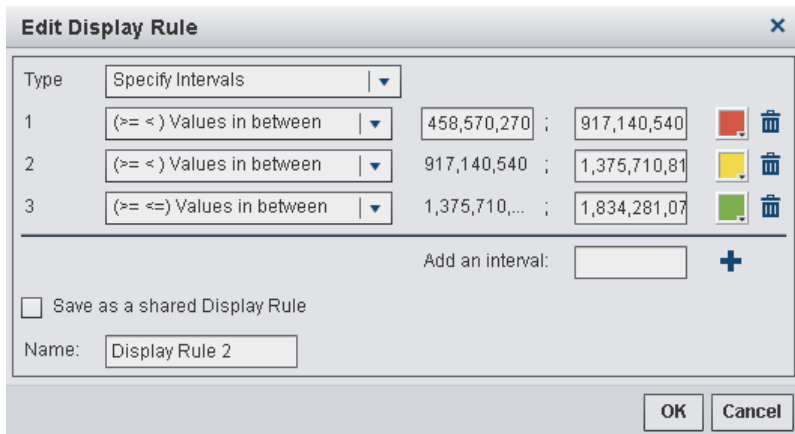
To edit the display rule for a gauge:

a   Click to display the Edit Display Rule window.

b   Select a **Type**. Your choices are either **Use a Shared Display Rule** or **Specify Intervals**.

If you select **Specify Intervals**, then the window contents change so that you can add intervals. Here is an example:

Update the conditions for the values, the intervals, and the color for the range. Select a color to open the color picker. To add new intervals to the display rule, enter a number and click **Add**. Select the **Save as a shared Display Rule** check box and specify a **Name** if you want to save the display rule.

c   Click **OK**. The gauge updates with the new display rule.

35

Working with Report Filters

About Report Filters

In SAS Visual Analytics Designer (the designer), you can create filters to subset your data.

Three types of filters are available:

Basic filters
> subset the data for individual report objects in your reports by using a single data item. A basic filter is constrained to use only the data item that was selected when the basic filter was created. Basic filters can be created and modified by using the context menu for a report object or by using the Edit Filter window. For more information, see "Working with Basic Report Filters" on page 419.

Advanced filters
> subset the data for individual report objects in your reports by using any number of data items and operators (for example, OR and AND) in the same expression. Advanced filters can be created and modified by using the Advanced Filter window. For more information, see "Working with Advanced Report Filters" on page 424.

Data source filters
> subset the data for the entire report. The data source filters that you create are applied to every report object in the report. For more information, see "Working with Data Source Filters in a Report" on page 428.

All of your filters are saved when you save your report.

Note: Filters are available only for report objects in the same section.

When a report object imported from SAS Visual Analytics Explorer (the explorer) has a local filter, that filter is displayed on the **Filters** tab. You cannot edit or delete these filters.

Note: Report linking is a type of filter. The report section that is the target of the link is filtered by the values selected in the linked report object. For more information, see "Overview of Linking" on page 449.

Working with Basic Report Filters

About Basic Report Filters

For certain report objects, you can subset your data by using the **Filters** tab in the right pane of the designer. You can base your filters on any data item in the current data source for this report object, regardless of whether the data item is assigned to a report object in the current report.

Create a Basic Report Filter

To create a basic report filter:

1   If it is not already selected, select the report object in the canvas that you want to filter. The report object must have at least one data item assigned.

2   Click ▼ beside the **Add Filter** button, and select a data item from the list. Then, click **Add Filter**. The filter appears on the tab.

3   Select the data values for the filter.

> **TIP** Use the arrow to the left of the filter name on the **Filters** tab to expand or collapse the filter details when you are working with multiple filters.

■   For a basic filter that uses discrete values, a check box is displayed for each distinct value that uses the current format applied to the data item. To the right of each value, a bar indicates the frequency. Select the data values that you want to filter or clear the selections for the data values that you do not want to filter. Select **All** to select all of the values.

Here is an example of a basic filter for discrete values:
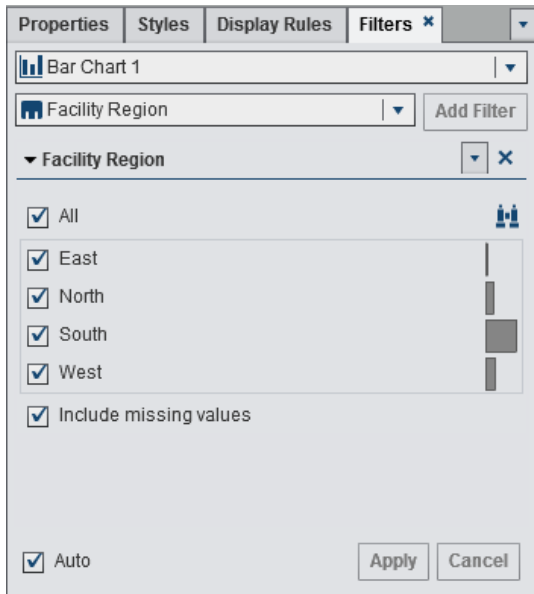
Note: By default, your filter changes are applied automatically to the active report object. To apply multiple changes together, clear the **Auto** check box, and then click **Apply** when you are ready to apply your filter changes.

> **TIP** Clear the **Auto** check box if the report object has a long-running query. Otherwise, the query runs each time you change the selection criteria.

■ For a filter that uses continuous values, a slider shows you the maximum and the minimum values that exist for the data item using the current data item format. Use the slider to select a range of target values.

Here is an example of a basic filter for continuous values:
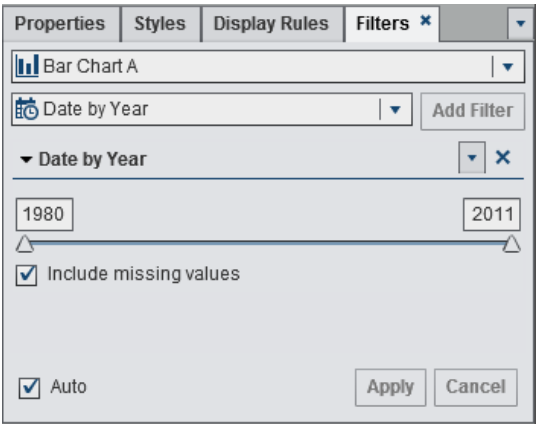
4   (Optional) If your data contains missing values, and you want to exclude those missing values from your report, then clear the **Include missing values** check box.

5   Click ▼ (to the left of the delete icon) for options. The available options depend on whether you are filtering characters, dates, or numerics. You can sort the values or frequencies in the filter.

The following options are available:

Option	Availability
Edit Filter	This option is always available.
Filter Using Continuous Values	This option is available only for measure data items.
Filter Using Discrete Values	This option is available for measure and category data items.
Include Missing Values	This option is available when you filter continuous values. It works the same as selecting the **Include missing values** check box.
Exclude Missing Values	This option is available when you filter continuous values. It works the same as clearing the **Include missing values** check box.

Option	Availability
Select All	This option is available when you filter discrete values. It selects all of the current discrete values that are listed. This option does not affect the **Include missing values** setting.
Clear All	This option is available when you filter discrete values. It clears the selections of all of the current discrete values that are listed. This option does not affect the **Include missing values** setting.
Invert Selection	This option is available when you filter discrete values. It changes all of the discrete value check boxes so that if they are selected, then they are cleared. If they are cleared, then they are selected. This option does not affect the **Include missing values** setting.
Sort by Values	This option enables you to sort by the values in the filter.
Sort by Frequency	This option enables you to sort by the frequencies in the filter.

6   (Optional) To apply multiple changes together, clear the **Auto** check box, and then click **Apply** when you are ready to apply your filter changes.

Edit a Basic Report Filter

To edit a basic report filter:

1   If it is not already selected, select the report object in the canvas that you want to filter. The report object must have at least one data item assigned.

2   In the right pane, click the **Filters** tab.

3   Click ▼ beside the filter name. Then, select **Edit Filter**. The Edit Filter window is displayed.
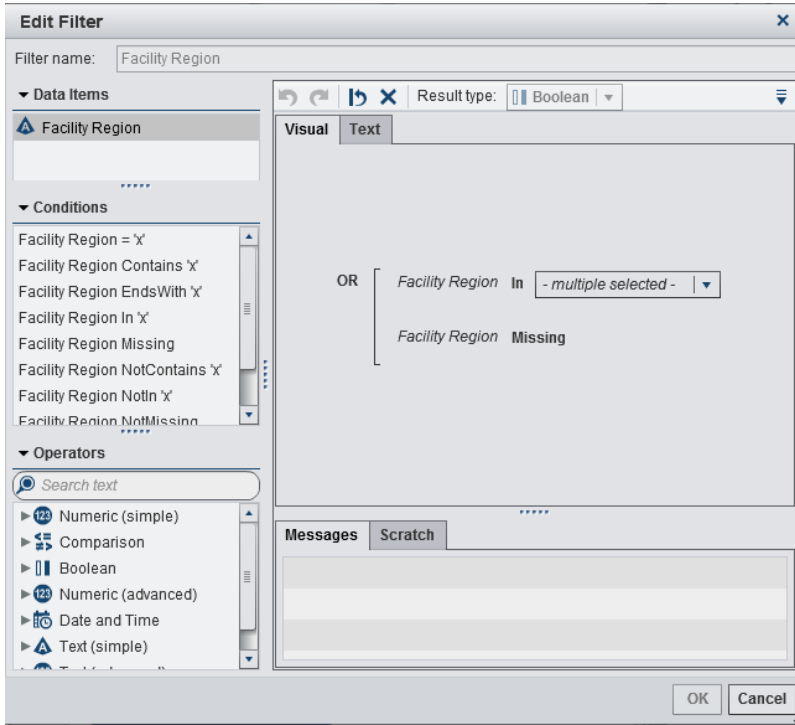
4 Edit the expression for your filter.

    ■ You can drag and drop conditions and operators onto the expression on the **Visual** tab in the right pane.

    ■ You can enter the expression on the **Text** tab in the right pane.

    ■ You can create an expression using both the **Visual** and **Text** tabs in the right pane.

For information, see Appendix 4, "Conditions for Filters," on page 599.

Note: The AND and OR operators can accept more than two conditions. To add a condition to the operator, drag and drop a condition onto the operator name in the right pane. For example, to add a third condition to an AND operator, drag and drop the new condition onto AND in the expression.

TIP Right-click the AND or OR operator in the expression, and then select **Add ▶ New Condition**.

5   Click **OK** to apply the filter.

Delete a Basic Report Filter

To delete a basic report filter, click ✖ beside the filter on the **Filters** tab.

Working with Advanced Report Filters

About Advanced Report Filters

For most report objects, you can create advanced filters to subset your data by using the **Filters** tab in the right pane of the designer.

Advanced filters enable you to create filters that use more than one data item.

Note: Advanced filters that are created in the explorer might contain expressions that you cannot create in the designer.

Note: Advanced global filters that are created in the explorer are supported by the designer, but you cannot create advanced global filters in the designer.

Create an Advanced Report Filter

To create an advanced report filter:

1   If it is not already selected, select the report object in the canvas that you want to filter. The report object must have at least one data item assigned.

2   In the right pane, click the **Filters** tab.

3   Click ▾ to select **Advanced** instead of a specific data item.

4   Click **Add Filter**. The Edit Advanced Filter window is displayed.
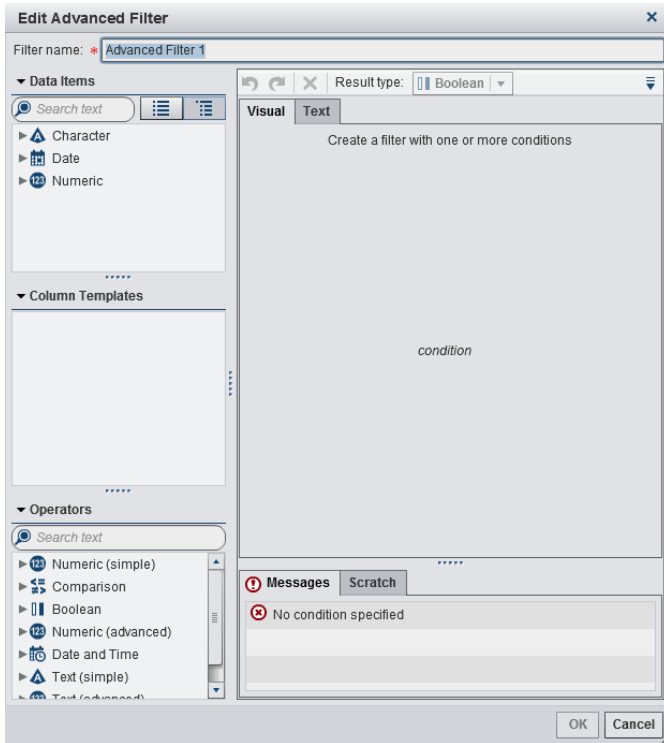
5 Specify a **Filter name**. The filter name identifies the advanced filter on the **Filters** tab.

6 Create a condition for the filter. You can use the **Visual** tab, **Text** tab, or a combination of both tabs.

   To create a condition using the **Visual** tab:

   a From the **Data Items** list, select the data item on which the condition is based.

   b From the **Column Templates** list, select a condition. For information, see Appendix 4, "Conditions for Filters," on page 599.

   c Drag and drop the condition onto the expression.

   d For any required parameters, select the parameter, and enter a value, or right-click on the parameter field, and select **Replace with** to select a data item.

7   (Optional) Add operators to the expression. You can use the **Visual** tab, **Text** tab, or a combination of both tabs.

To add operators using the **Visual** tab:

a   From the **Operators** list, select an operator to join the conditions in your expression. For more information, see Appendix 4, "Conditions for Filters," on page 599.

b   Drag and drop the operator onto the expression.

c   From the **Column Templates** list, select another condition. Then, drag and drop the additional condition onto the expression. Complete any required parameters.

> **TIP** Right-click the AND or OR operator in the expression, and then select **Add ▶ New Condition**.

8   Click **OK** to apply the filter.

Edit an Advanced Report Filter

To edit an advanced report filter:

1   If it is not already selected, select the report object in the canvas that you want to filter. The report object must have at least one data item assigned.

2   In the right pane, click the **Filters** tab.

3   Click ▼ beside the advanced filter name. Then, select **Edit Filter**. The Edit Advanced Filter window is displayed.

4   Edit or add a condition for the filter. You can use the **Visual** tab, **Text** tab, or a combination of both tabs.

To add conditions using the **Visual** tab:

a   From the **Data Items** list, select the data item on which the condition is based.

b   From the **Column Templates** list, select a condition. For more information, see Appendix 4, "Conditions for Filters," on page 599.

c   Drag and drop the condition onto the expression.

d   For any required parameters, select the parameter, and enter a value, or right-click the parameter field, and select **Replace with** to select a data item.

5   (Optional) To replace a condition, drag and drop a new condition onto the existing condition in the expression on the **Visual** tab. Or, remove a condition using the **Text** tab.

6   (Optional) Add an operator to the expression. You can use the **Visual** tab, **Text** tab, or a combination of both tabs.

To add an operator using the **Visual** tab:

a   From the **Operators** list, select an operator to join the conditions in your expression. For more information, see Appendix 4, "Conditions for Filters," on page 599.

b   Drag and drop the operator onto the expression.

c   (Optional) Add a condition.

> **TIP** Right-click the AND or OR operator in the expression, and then select **Add ▶ New Condition**.

7   (Optional) Delete part of an expression by highlighting the part of the expression that you want to delete, and then selecting **Delete** or **Clear**.

8   Click **OK** to apply the advanced filter.

Delete an Advanced Report Filter

To delete an advanced filter, click ✕ on the filter on the **Filters** tab.

Working with Data Source Filters in a Report

About Data Source Filters

Data source filters are used to restrict the data that is displayed in a report. The data source filters that you create in the designer are applied to every report object in the report. A report that has multiple data sources can contain multiple data source filters.

You can use either continuous values or discrete values to create a data source filter. Continuous value filters can be used only for measures. Discrete value filters can be used for any character; numeric; or date, datetime, or data item; as long as the total number of distinct values does not exceed a maximum number.

Data source filters are not displayed on the **Filters** or the **Data** tabs. A data source filter updates the cardinality values that appear on the **Data** tab.

There is a limit of one data source filter per data source.

Create a Data Source Filter

1   On the **Data** tab, click ▼ , and then select **New Data Source Filter**. The New Data Source Filter window is displayed.

2   Add a condition for the filter. You can use the **Visual** tab, **Text** tab, or a combination of both tabs.

   To add conditions using the **Visual** tab:

   a   From the **Data Items** list, select the data item on which the condition is based.

   b   From the **Column Templates** list, select a condition. For more information, see Appendix 4, "Conditions for Filters," on page 599.

   c   Drag and drop the condition onto the expression.

d   For any required parameters, select the parameter, and enter a value, or right-click the parameter field, and select **Replace with** to select a data item.

3   (Optional) To replace a condition, drag and drop a new condition onto the existing condition in the expression on the **Visual** tab. Or remove a condition using the **Text** tab.

4   (Optional) Add an operator to the expression. You can use the **Visual** tab, **Text** tab, or a combination of both tabs.

To add an operator using the **Visual** tab:

a   From the **Operators** list, select an operator to join the conditions in your expression. For more information, see Appendix 4, "Conditions for Filters," on page 599.

b   Drag and drop the operator onto the expression.

c   (Optional) Add a condition.

> **TIP**   Right-click the AND or OR operator in the expression, and then select **Add** ▶ **New Condition**.

5   (Optional) Delete part of an expression by highlighting the part of the expression that you want to delete, and then selecting **Delete** or **Clear**.

6   Click **OK** to apply the filter.

Edit a Data Source Filter

1   If a data source is not already added, add a data source.

2   On the **Data** tab, click ▼ , and then select **Edit Data Source Filter**. The Edit Data Source Filter window is displayed.

3   Add or modify the condition for the filter. You can use the **Visual** tab, **Text** tab, or a combination of both tabs.

To edit conditions using the **Visual** tab:

a From the **Data Items** list, select the data item on which the condition is based.

b From the **Column Templates** list, select a condition. For more information, see Appendix 4, "Conditions for Filters," on page 599.

c Drag and drop the condition onto the expression.

d For any required parameters, select the parameter, and enter a value, or right-click the parameter field, and select **Replace with** to select a data item.

4 Edit the operator in the expression. You can use the **Visual** tab, **Text** tab, or a combination of both tabs.

5 Click **OK** to apply the filter.

Delete a Data Source Filter

To delete a data source filter:

1 On the **Data** tab, click ▼ , and then select **Delete Filter**.

2 Click **Delete** in the confirmation message that is displayed.

36

Working with Interactions

Overview of Interactions

Interactions are used to direct a report viewer's attention to specific results in a report. Interactions allow data to be subset to reduce the amount of data, and enables users to understand it within a particular context.

The interactions view in SAS Visual Analytics Designer (the designer) enables report designers to specify which interactions they would like to add to tables, graphs, gauges, and controls in a report.

There are these types of interactions:

filter
> is used to restrict the data that is returned from a query to a data source. Filters are simply a set of rules or conditions that you specify to subset the data that is displayed in a table or graph. The goal is to display only the data that you need to see to perform your analysis.

brush

is short for *data brushing*, which enables you to show the same data selected simultaneously in two or more tables, graphs, or controls. Brushing highlights a percentage that reflects the number of shared observations in the data set. Brushing does not highlight a percentage that corresponds to the aggregated value. The brushed data has the same appearance in each object, which makes the data easily apparent to report viewers.

Tables, graphs, and gauges can be the source of an interaction, with the exception of time series plots. Report objects that use detail data cannot be the source of an interaction. Controls that are used on the canvas can also be the source of an interaction. Controls that are used as section prompts are treated as automatic filters and are not displayed in the interactions view.

The **Interactions** tab in the right pane enables you to create interactions and links.

Note: Interactions are available only for report objects in the same section.

Note: Report linking is a type of interaction. The report section that is the target of the link is subset by the values selected in the linked report object. For more information, see "Overview of Linking" on page 449.

Creating Interactions

Create a Filter Interaction

You can create interactions using the interactions view or the **Interactions** tab in the right pane.

The following objects that are imported from the explorer can be the source of a filter interaction:

■ Box plots

■ Heat maps with a relational category measure

■ Forecast plots

■ Geo maps

Imported correlation matrices, histograms, scatter plots with fit lines, heat maps without a selectable relational category measure, and box plots without a selectable relational category measure cannot be the source of an interaction. Hierarchy data items are not supported for interactions.

Note: You cannot create interactions from list tables, bubble plots, and scatter plots that use detail data. However, a list table that uses detail data can be the target of an interaction.

Note: Scatter plots cannot be the source of an interaction.

Note: Bubble plots can be the source of an interaction only if they have a **Grouping** role assigned.

To create a filter interaction using the interactions view:

1   Add the report objects that you want to use to the canvas. For example, you might have a bar chart, a line chart, and a list table.

2   Click **Interactions View** on the **Interactions** tab to switch to the interactions view. You can also select **View ▶ Show Interactions**.

3   In the interactions view, draw a connection between the source and target report objects. When you have created an interaction, the 🔻 is displayed between the report objects. Here is an example of an interaction between a bar chart (the source) and a list table (the target):

   Note: The source and target of an interaction should be based on the same data source. If you have multiple data sources, you are prompted to map the data sources to create the interaction. For more information, see "Map Data Sources" on page 444.
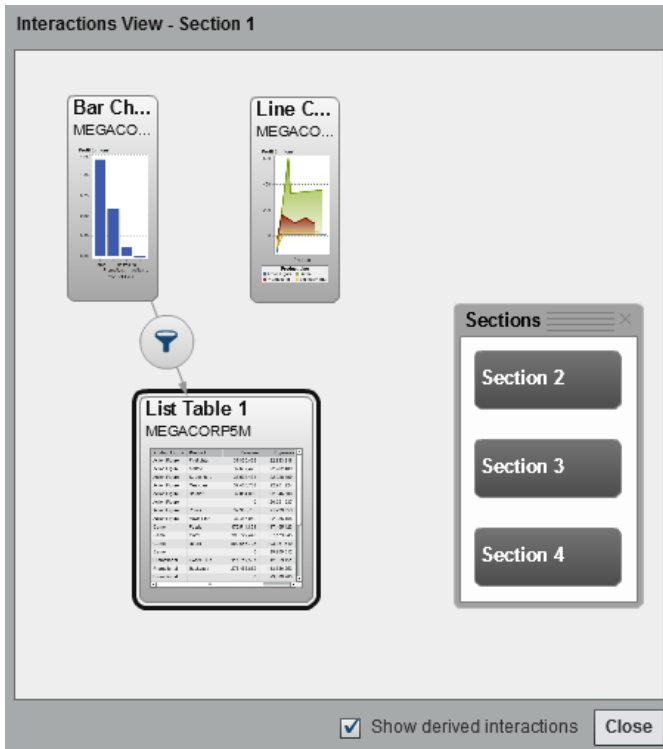
4   Right-click 🝖. Then, select **Interaction Type ▸ Filter**.
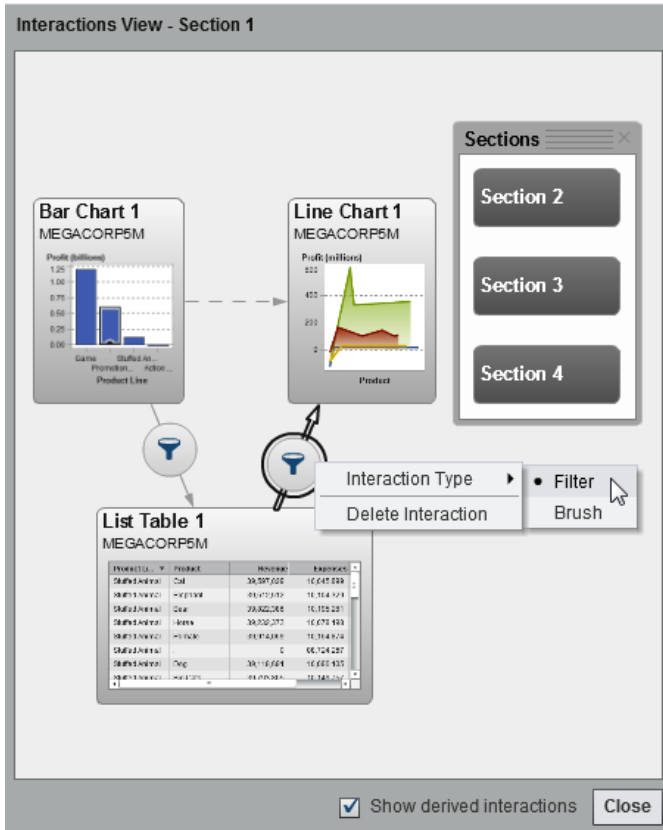
5 (Optional) Clear the **Show derived interactions** check box. (Derived interactions are represented by dashed lines in the Interactions View.)

6 (Optional) Add other filter interactions.

7 Click **Close**. The new interaction (or interactions) appears on the **Interactions** tab in the right pane.

The following example shows the **Interactions** tab for the filtering example shown in .

Select data in the source report object to filter data in the target report object (or objects). In the following example, the promotional product line was selected in the bar chart. Simultaneously, the line chart and the list table show the same filter.

Display 36.1 A Filtering Example with a Bar Chart, a Line Chart, and a List Table
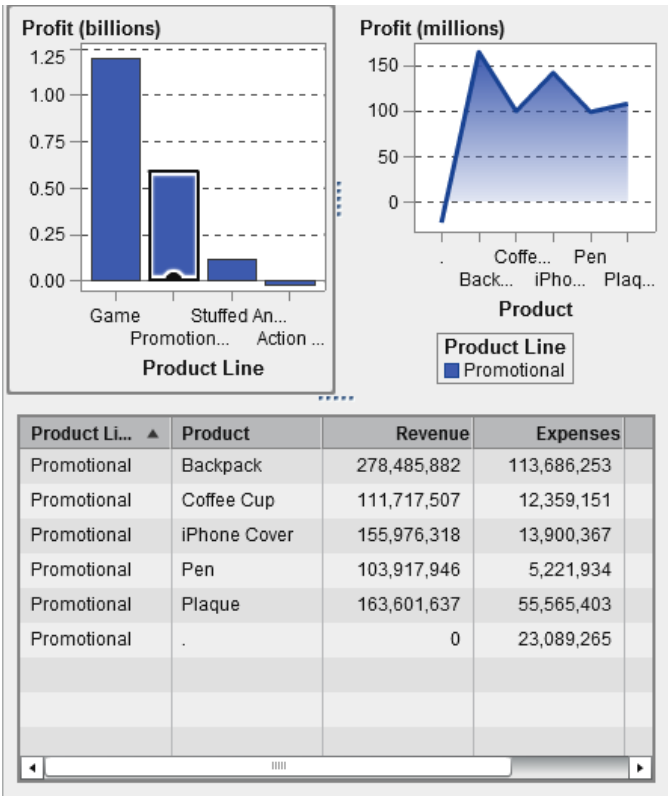
Clicking on another selection applies the filter based on your new selection. For example, in the report above, you can click on the game product line to change the filter. An icon is added to the report object that is filtered.

Here is the line chart that was used in the previous example. The incoming filter icon appears above the report object when it is selected in the canvas. Click 🎯 to get details about the filter that has been applied. Click **Details** to find out more about the filter. The

Visual Element Filters window is displayed. It shows the exact query that was sent to the server.
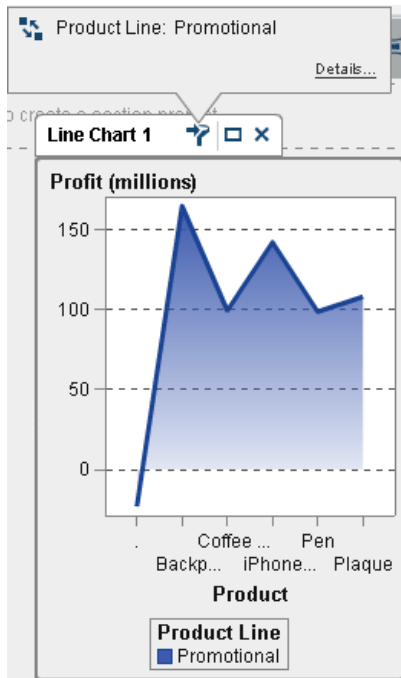
To clear the selection and reset the filter (or filters), press Ctrl+click while viewing the original report object.

To create an interaction using the **New** button on the **Interactions** tab:

1 Select a report object in the canvas. For example, you might have a bar chart.

2 Select the **Interactions** tab.

3 Click **New**, and then select **Interaction**. The **Edit Interaction** window is displayed.
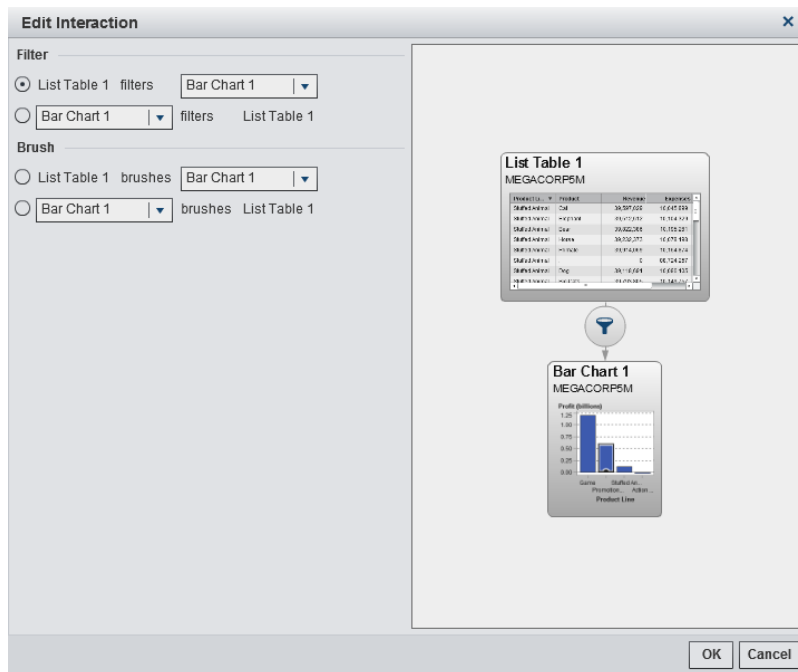
4   In the Edit Interaction window, select the report object that you want to use as the source of a filter or brush interaction. For example, if your report has a bar chart, and you want it to filter a line chart, the first line should read *Bar Chart 1* filters *List Table 1*.

The Edit Interaction window disables options for filter or brush interactions if they are not possible.

5   When you are finished working with interactions, click **OK**. The new interaction is displayed in the **Interactions** tab.

Create a Data Brushing Interaction

The following objects that are imported from the explorer cannot be the source of a brush interaction:

■   Box plots

■   Heat maps with a relational category measure

■ Forecast plots

■ Geo maps

To create an interaction using data brushing:

1 Add the report objects that you want to use to the canvas. For example, you might have a bar chart, a line chart, and a list table.

2 Select the **Interactions** tab.

3 Click **Interactions View** on the **Interactions** tab to switch to the interactions view. You can also select **View** ▶ **Show Interactions**.

4 In the interactions view, draw a connection between the source and target report objects. When you have created an interaction, the ▼ is displayed between the report objects.

Note: The source and target of a data brushing interaction should be based on the same data source. If you have multiple data sources, you are prompted to map the data sources to create the interaction. For more information, see "Map Data Sources" on page 444.

5 Right-click ▼. Then, select **Interaction Type** ▶ **Brush**.

The filter icon changes to ▣.

6 Click **Close**. The new interaction (or interactions) appears on **Interactions** tab in the right pane.

Select data in the source report object to brush data in the target report object (or objects). In the following example, the stuffed animal product line was selected in the bar chart. Simultaneously, the line chart and the crosstab highlight the same data.

Display 36.2 *A Data Brushing Example with a Bar Chart, a Line Chart, and a List Table*
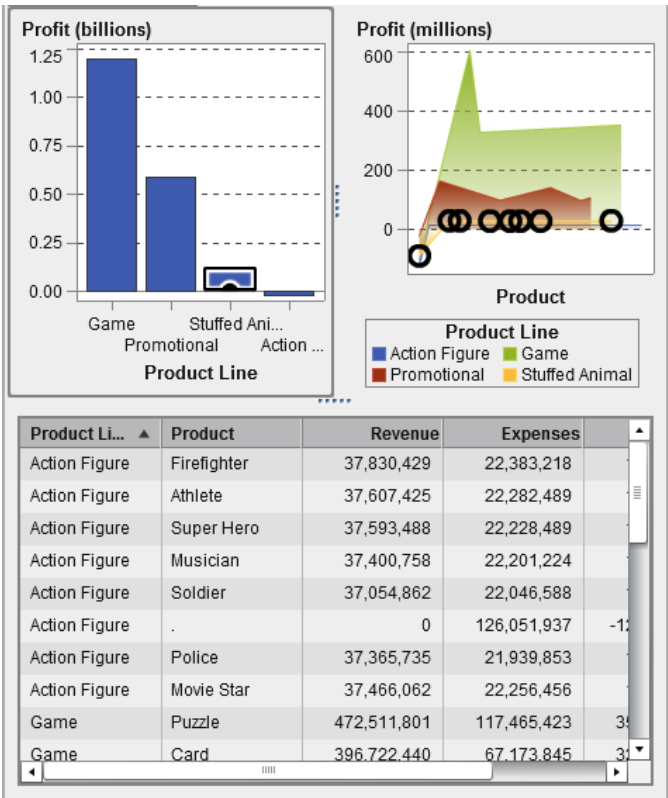
Note: If a graph contains a frequency measure, then a crosshatch pattern indicates when it is selected or brushed.

The brush interaction between controls behaves differently than a brush interaction between tables and graphs. When there is a brush interaction between controls, instead of highlighting or selecting the values in the other control, the values that match are moved to the top of the control, and the values that do not match are grayed out. If a frequency data item exists, then the frequency value is updated.

Here is an example of a report on cars with three list controls:

Display 36.3 *Three List Controls with Brush Interactions*
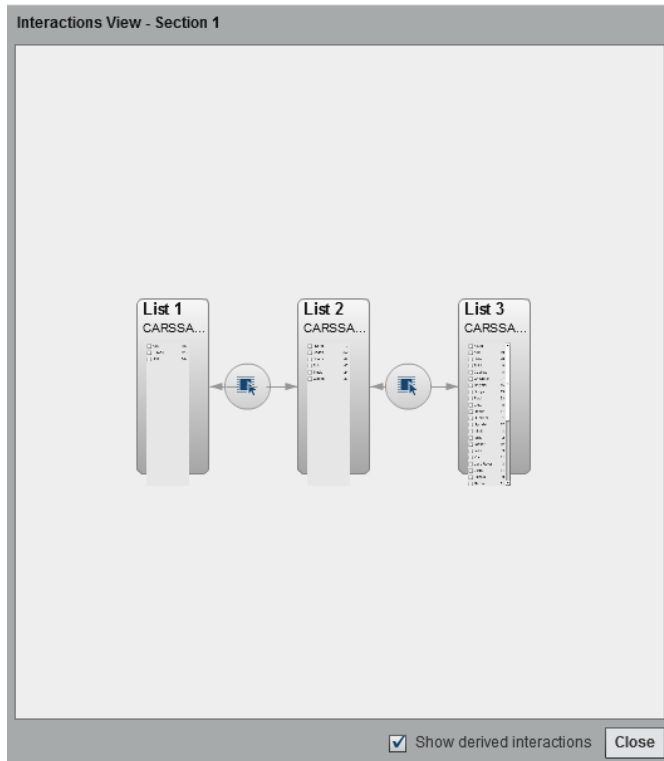
The first list contains three countries listed as the `Origin`. The second list contains six types of vehicles. The third list contains 38 car manufacturing companies. There is a brush interaction between the first list and the second list. There is a second brush

interaction between the second list and the third list. Look at the values in the second list.

Display 36.4   Example: Three List Controls in a Report
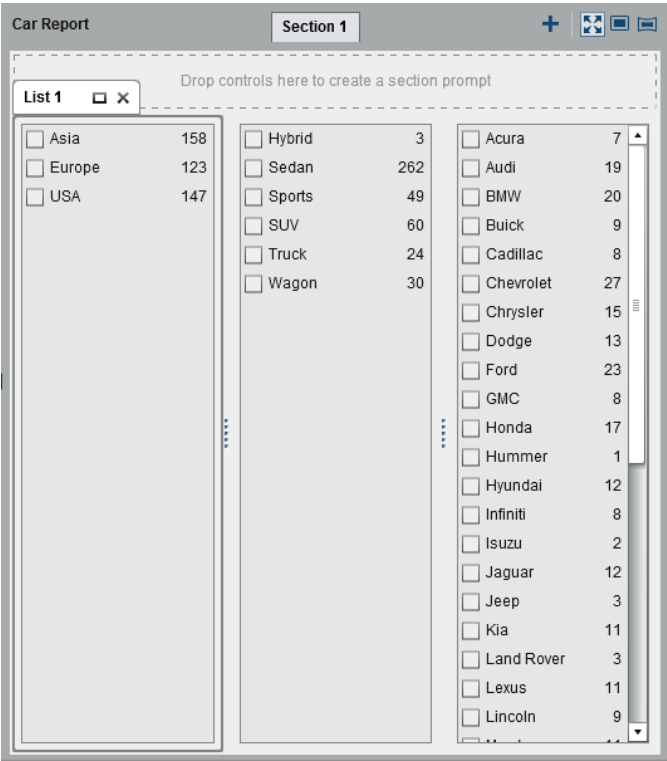
In step 1, the user has selected **Europe** as the country of **Origin** in the first list. Notice the changes in the second list. **Hybrid** and **Truck** are grayed out, and their frequency

values have both changed to 0. In the third list, all of the car manufacturers that are not located in Europe are grayed out, and their frequency values have all changed to 0.

Display 36.5   *Example: Three List Controls in a Report with One Brush Interaction*
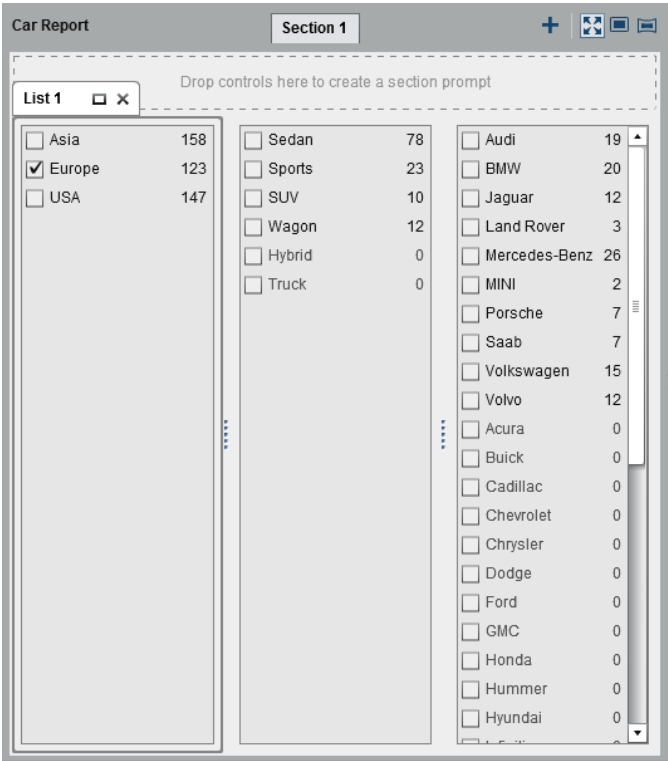
In step 2, the user has selected `Porsche` as the car manufacturing company in the third list. Notice the changes in the second list. `Hybrid`, `Sedan`, `Truck`, and `Wagon` are grayed out, and their frequency values have all changed to `0`.

Display 36.6 *Example: Three List Controls in a Report with Two Brush Interactions*
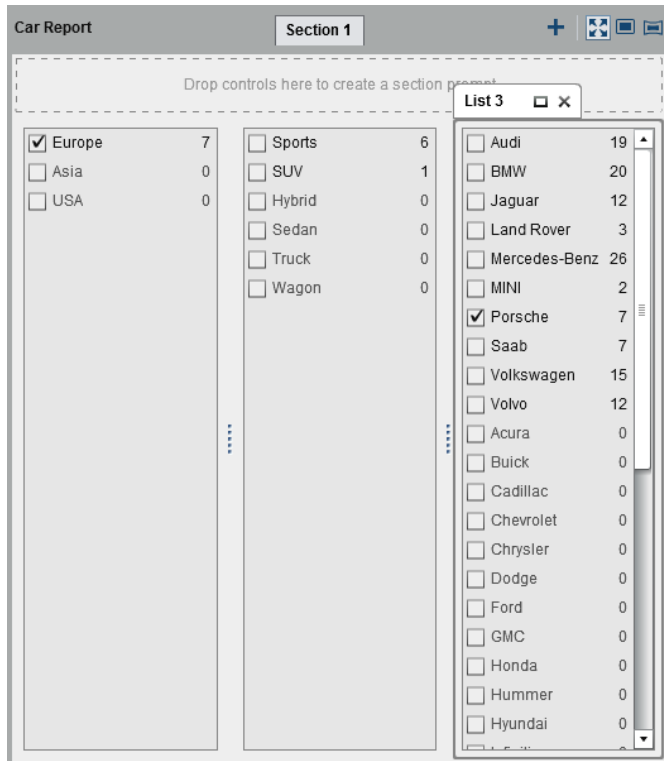
Map Data Sources

The source and target of an interaction or link should be based on the same data source. You might be prompted to map data sources so that an interaction or link works properly in these situations:

- When you try to create an interaction or link between a report object that uses one data source and another report object that uses a different data source

- When you try to create a link between a report object that uses one data source and a report section that uses a different data source

■ When you try to create a link between a report object that uses one data source and a report that uses a different data source

To map data sources using the Map Data Sources window:

1 Use the **Source** drop-down list to select a column from the first data source.

2 Use the **Target** drop-down list to select a corresponding column from the second data source.

Here is an example of the Map Data Sources window for an interaction between two list tables that have different data sources:

For section links, you are not required to map the data sources. If you choose not to map the data sources, then the filters are not carried over. For more information, see "Create a Link to One or More Sections in a Report" on page 451.

Here is an example of the Create Section Link window with the map data sources information.

3 (Optional) Map additional data items. Click ➕. A new row is displayed in the Map Data Sources window.

4 The lower half of the window shows you how the mappings are applied to the objects shown in the **Source** and **Target** relationship. Depending on the interaction or link, you might need to map additional **Source** and **Target** objects.

5 Click **OK**.

Deleting Interactions

You can delete interactions using the interactions view or the **Interactions** tab in the right pane.

To delete either a filter or data brushing interaction using the interactions view:

1   Select the **Interactions** tab.

2   Click **Interactions View** to switch to the interactions view. You can also select **View**
    ▶ **Show Interactions**.

3   Right-click ▼ or the ▦, depending on which type of interaction you want to delete.
    Then, select **Delete Interaction**.

4   Click **Delete** in the confirmation message that is displayed.

5   (Optional) Remove other interactions.

You can delete an interaction using the **Interactions** tab. Click 🗑 next to the interaction
that you want to delete. Then, click **Delete** in the confirmation message that is
displayed.

37

Working with Report Links

Overview of Linking

Report links enable single-step access to a report or web page that is related to a report that you are currently viewing. For example, you might be looking at a bar chart that has sales information for each geographical region of your company. If you click the bar for the Northeast region, then a report link associated with the graph could take you to a different report that provides information about employees in each region. You can click ⊡ in the top left corner of a destination report to return to the previous report.

Using SAS Visual Analytics Designer (the designer), you can add a link from a report object to another report, to a specific section in the current report, or to an external URL.

If a destination report contains multiple sections, then you are able (when defining the link) to choose the initial section of the destination report that you want to open first.

Linking has elements of both a filter and an interaction. A report section that is the target of a link is filtered by the values that are selected in the linked report object.

For information about how links appear in the SAS Visual Analytics Viewer (the viewer), see "Viewing Links in Reports" on page 537.

Imported box plots and heat maps with a relational category measure and imported forecast plots can be the source of section, report, or external links.

Creating Links

Create a Link to Another Report

Note: The following steps do not apply to text objects. For more information, see "Create a Link from a Text Object" on page 453.

To add a link from a report object or image to another report:

1 Click the **Interactions** tab.

2 Click **New**, and then select **Report Link**. The Select window is displayed.

3 Select a target report, and then click **OK**. The Create Report Link window is displayed.

4 Verify the name of the target report. If you want to select a different target report, click **Browse**.

   Note: The source and target of report link should be based on the same data source. If you have multiple data sources, you are prompted to map the data sources to create the report link. For more information, see "Map Data Sources" on page 444.

   If the target report contains more than one section, then select the report section that you want to open first.
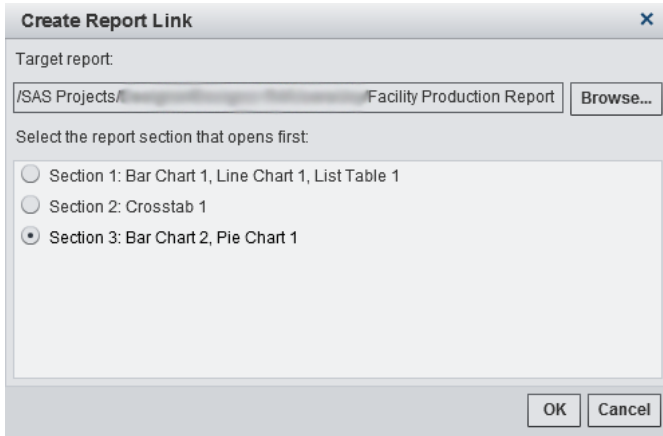
5 Click **OK**. The new report link is displayed on the **Interactions** tab in the right pane.

Double-click on an item in the report object (for example, a bar, a row in a table, a pie slice, and so on) to activate the report link. You are prompted to save changes to the report object before you can follow the link to the target report. To return to the source report object, use the button next to the target report name above the canvas.

Here is an example of the button that enables you to return to the source report object:

Display 37.1 *Button to Return to the Source Report Object*
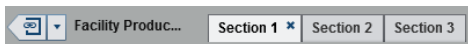
Create a Link to One or More Sections in a Report

Note: The following steps do not apply to text objects. For more information, see "Create a Link from a Text Object" on page 453.

To add a link from a report object, text, or image to a specific section in the same report:

1 Click the **Interactions** tab.

2 Click **New**, and then select **Section Link**. The Create Section Link window is displayed.

Create Section Link   ✕

Link List Table 1 to:

☐ All

| ☐ Section 2: Crosstab 1 |
| ☐ Section 3: Bar Chart 2, Pie Chart 1 |

Filter links can be passed only to objects with the same data source.

[OK] [Cancel]

3   Select **All** to link to all of the sections in the report or select a specific section (or sections) in the report.

   Note: The source and target of section link should be based on the same data source. If you have multiple data sources, you are prompted to map the data sources to create the section link. For more information, see "Map Data Sources" on page 444.

4   Click **OK**. The new section link is displayed on the **Interactions** tab in the right pane. If you selected multiple sections, those links are displayed on the **Interactions** tab.

Alternatively, you can create a section link using the interactions view.

Create a Link to an External URL

You can add a link from a report object, text, or image to an external URL. The `http://` part of the link is provided. The link can be relative to the current web application server. For example, you can link to a stored process, which means that you would not need to specify `http://`*server-name*:*port* because you are already logged on to that server.

Note: UTF-8 is supported for external URL links.

Note: The following steps do not apply to text objects. For more information, see "Create a Link from a Text Object" on page 453.

To add a link to an external URL:

1　Click the **Interactions** tab.

2　Click **New**, and then select **External Link**. The Create External Link window is displayed.

3　Enter a **Label** and a **URL** for the link.

4　(Optional) Click ✚ to specify additional parameters for the link. You can specify the **Format value** option, a **Source**, and a **Target**. The **Format value** option applies only to dates and numbers. You can add multiple parameters.

For example, you might have a report for your company's product lines. When a user views the report, you want them to be able to double-click a product line in the line chart, and then link to a Google search for that product line. In the Create External Link window for this example, you would enter *http://www.google.com/search* as the **URL**, select `ProductLine` as the **Source**, and then enter *q* as the **Target**.

5　Press Enter or Tab to activate the **OK** button. Click **OK**. The new external link is displayed on the **Interactions** tab in the right pane.

Double-click on a report object to link to the external URL.

Create a Link from a Text Object

You can add a link from text to an external URL, to another report, or to a specific section in the current report.

To add a link from a text object:

1   Double-click inside a text object on the canvas and enter the text.

2   Selecting some or all of the text, and then click ⊖. Select **Hyperlink**. The Link Setup window is displayed.

3   Select the link **Type**. The options are:

External URL
>   The link target is an external URL. For example, you might want to link some text to your company's web page or Twitter feed. The link can be relative to the current web application server. Enter a **URL**. The `http://` part of the link is provided.

Report Link
>   The link target is another report. Enter the name of the **Target report** or browse to find it. If the target report has multiple sections, you can select which report section opens first.

Section Link
>   If your report has multiple sections, then the link target can be a section in the report. Select a **Target section** from the list.

4   Click **OK**. The link appears in the text object.

Editing Links

To edit a report link or an external link:

1   Click ⌕ next to the link on the **Interactions** tab. For a report link, the Edit Report Link window is displayed. For an external link, the Edit External Link window is displayed.

2   Make changes to the link.

3   Click **OK** to save your changes.

Note: You cannot edit a section link.

Deleting Links

You can delete a section, report, or external link using the **Interactions** tab. Click 🗑
next to the link that you want to delete. Click **Delete** in the confirmation message that is
displayed.

To delete a link from a text object, select the linked text, and then click 🔗. Select **Delete
Link**.

Example: How Linking and Interactions Work Together

Suppose that you have two reports about furniture sales. The first report contains an
interaction between a bar chart and a list table. The bar chart links to a second report.
The filter interaction and the link are displayed on the **Interactions** tab. Currently,
CANADA is selected in the bar chart, so the list table is filtered by that country. Here is
an example of the first report:

Display 37.2 *Example: Furniture Sales Report*
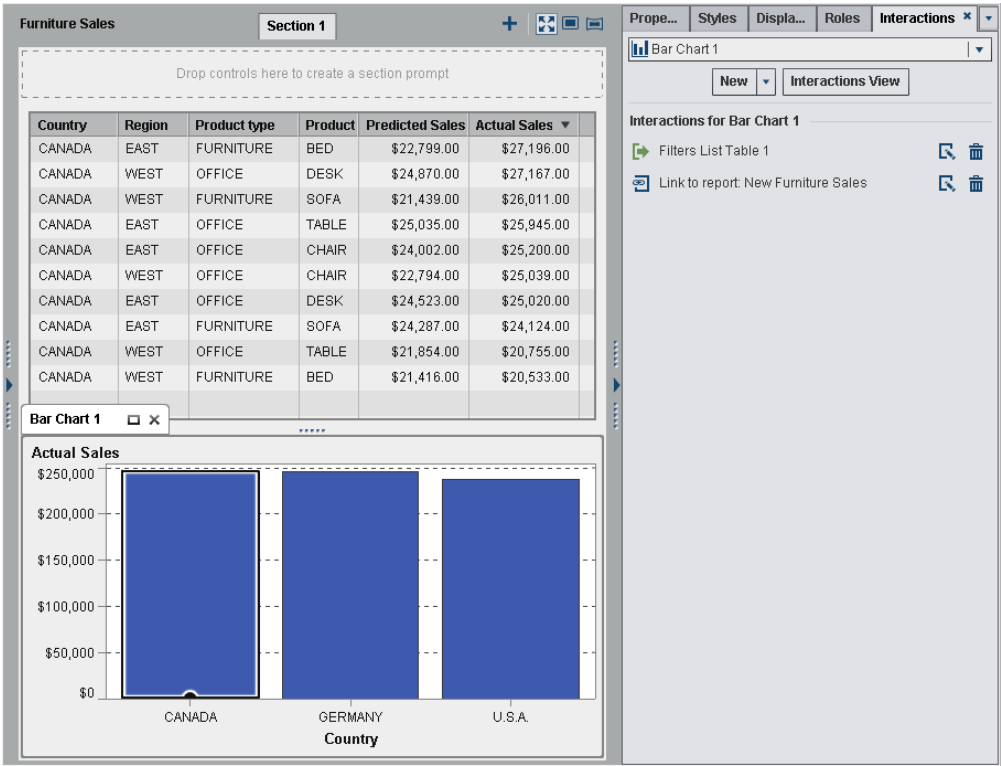
When you double-click the CANADA bar in the first report, the second report (which is the target of the report link) is displayed. The filter interaction for country in the first report is automatically passed to the second report when you link to it. Here is an example of the second report:

Display 37.3   Example: Furniture Sales Report

You can remove the applied filter by clicking ▼ beside the 🔁 in the upper left corner. Select the **Apply report link filters** option. After you remove the applied filter, you can choose to apply it again by selecting the **Apply report link filters** option.

You can return to the source report by clicking 🔁.

For information about how links appear in the viewer, see "Viewing Links in Reports" on page 537.

38

Ranking Values in Reports

Overview of Ranking in Reports

Using SAS Visual Analytics Designer (the designer), you can rank the data in a report object to show the top (greatest) value or bottom (least) value for a category that is based on a measure. For a list table, you can also rank across a set of categories for the top value or bottom value in the set. A rank filters the values of a category based on the aggregated measure by the top or bottom of the values. A rank greatly reduces the visible categories to make it easier to focus on the top value or bottom value that interests a user.

For example, you might create a rank of the top 10 countries by frequency to select the 10 countries that are most represented in your report. As another example, you might create a rank of the top 10 countries by population to select the 10 countries with the greatest populations.

Adding a New Rank

You can use the **Ranks** tab to create ranks to subset the data in your reports. You can also right-click on your graphs and tables to add a rank to limit the number of bars, rows, and so on. Controls and gauges support ranks.

A single category rank ranks the top of a single category.

To add a rank:

1   Click the **Ranks** tab.

2   Select a data item. You can select any category or geography data item, regardless of whether it is assigned to the current report object.

    For list tables only, you can select **All visible categories**. This rank ranks across the intersection. The top or bottom combination of the visible categories is displayed. For example, you select the region and product data items. You rank on the top 10 by profit. The top 10 region and product combinations are displayed. In this case, the column is no longer considered a single column, it is considered the crossing of the columns.

    The **Detail rank** option is available for list tables, bubble plots, and scatter plots that show detail data. If the report object has a detail rank, then it is the only rank that it can have. Otherwise, report objects can have multiple ranks.

3   Click **Add Rank**. The **Ranks** tab expands.
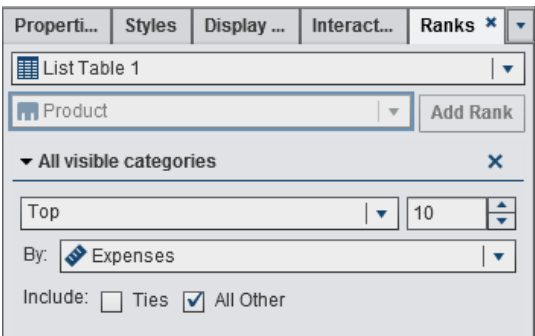
4   Select the type of rank from the drop-down list. **Top** specifies that the rank selects the greatest value. **Bottom** specifies that the rank selects the least value.

5   Specify the number of values for the rank. For example, if you specify 5, and you select **Top** as the rank type, then the rank selects the five greatest values.

6   From the **By** drop-down list, select the measure that is used to create the rank. You can select any measure.

7   (Optional) Select **Ties** to include ties in the rank.

If you select **Ties**, then the rank selects as many values as necessary to include all of the ties. If you do not select **Ties**, then the rank selects only the number of values that are specified by the rank parameters.

For example, if your rank selects the top three values, but there are five values tied for the greatest value, then the number of values that are selected by the rank depends on the **Ties** option. If you select **Ties**, then the rank includes all five of the tied values. If you do not select **Ties**, then the rank includes only three of the tied values.

8   (Optional) Select **All Other** to show the measurements for the categories that did not qualify as a top or bottom value.

Note: The **All Other** option is not available when there is a detail rank or when the data item is hidden. It is not available if the rank is on a prompt control, geo bubble map, or geo region map.

When the **All Other** option is specified for a report object, the total, subtotal, and percent of total show values with respect to all of the data, rather than data just relative to data qualifying under the rank. The data that does not fit into the top or bottom of the rank is aggregated in the **All Other** category.

You can use the **All Other** option with a top ranking to reduce the number of slices in a pie chart. However, this means that the "Other" slice is removed from the pie chart.

If you do not select **All Other**, then the rank shows only the data as filtered by the category values that qualify as the top or bottom value.

Note: If **All Other** is selected, then any **All Other** value that appears in a table or graph cannot be selected. This means that the **All Other** value cannot be the source value for an interaction.

By default, your new rank is applied automatically to the report object.

You can add more than one rank to a report object, as long as the first rank is not a **Detail rank** or an **All visible categories** rank.

Here is an example of a pie chart that shows the profit for multiple product lines before a rank is applied:

Display 38.1   Pie Chart with Profits for Each Product Line Displayed
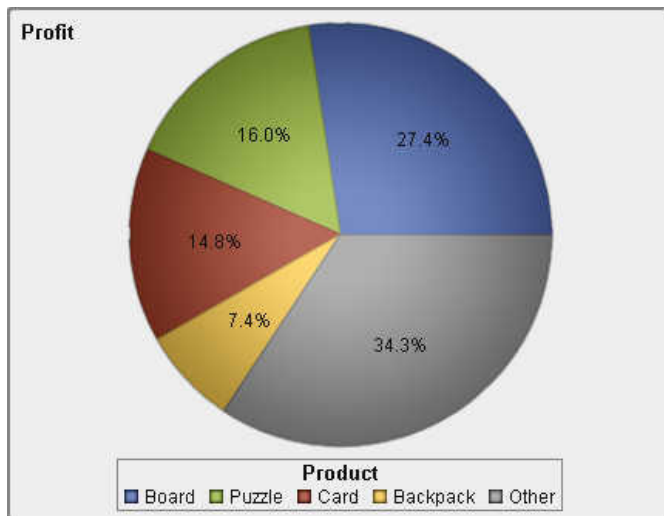
Here is the same pie chart after a rank of the top five product lines is added. The **All Other** option was selected, so the "Other" slice is not displayed.

Display 38.2   Pie Chart with Profits for the Top Five Product Lines Displayed
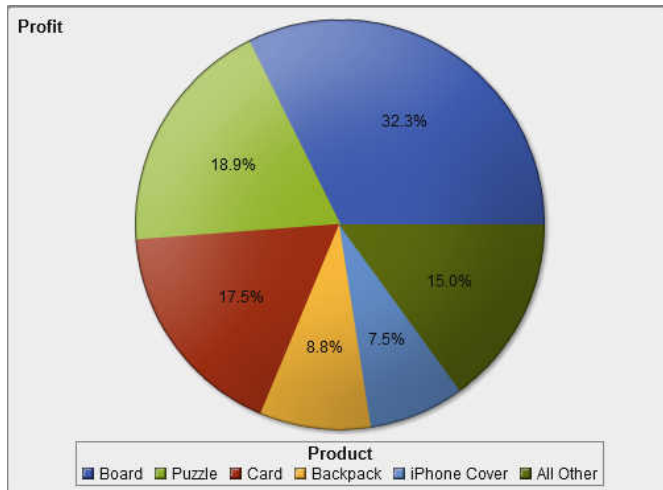
Deleting a Rank

To delete a rank, click ✖ on the rank on the **Ranks** tab.

You can also right-click on the report object, and select **Delete Rank**, as long as there is only one rank for the report object.

39

Maintaining Multi-Section Reports

Overview of Report Sections

Any report in SAS Visual Analytics Designer (the designer) can have multiple sections. (Sections are like pages.) Multiple sections can be used to present different views of the data to the person who views the report. Each section has one or more data sources. You can have one or more report objects in a section. There is no limit to the number of sections in a report.

Adding a Section to a Report

You can add sections to any report, including reports that were created in SAS Visual Analytics Explorer (the explorer).

To add a new section to a report, click **+** in to the right of the first section tab in the report. The new tab appears to the right of the existing tab (or tabs).

Here is an example:

Display 39.1   Multiple Section Tabs

Renaming a Report Section

To rename a report section:

1   Right-click the section tab that you want to rename. A menu is displayed.

2   Select **Rename**.

3   Type a new name, and then press Enter. The new name also appears on the **Properties** tab in the right pane.

Alternatively, you can change the name on the **Properties** tab.

Duplicating a Report Section

To duplicate a report section:

1   Right-click the section tab that you want to duplicate. A menu is displayed.

2   Select **Duplicate**. All of the report objects and the data in the original section are copied to the duplicate section. The duplicate section appears on a new section tab. For example, if the original section's name is `Marketing Reports`, then the duplicate section is named `Marketing Reports (1)`.

3   Rename the duplicated report section using the context menu or the **Properties** tab.

Deleting a Report

To delete a section from a report, right-click the section tab that you want to remove. Select **Delete**.

Note: The **Delete** menu option is not available if the report has only one section.

Moving or Dragging a Report Object to Another Section

To move a report object from one section in a report to another section, right-click the report object, and select **Move to**. Then, select the name of the other section. You can drag an object from one section and drop it onto the tab for another section to move it to that section.

If you move an object to a section with precision layout, then you must manually move the object to its proper location. All objects are put in the top left corner by default.

Part 5

Creating Custom Graph Objects

40

Overview of SAS Visual Analytics Graph Builder

About the Graph Builder

SAS Visual Analytics Graph Builder (the graph builder) enables you to create graph objects, which then become available in the SAS Visual Analytics Designer (the designer) for use in reports. You do not associate real data with the graph objects in the graph builder. Rather, you build graph objects or templates using sample data that is shipped with the graph builder. Report designers assign data when they include your graph objects in their reports.

The graph objects that you build have a consistent appearance that is compatible with the designer's graph objects. The graph builder enables you to produce a wide array of graph objects with more options for layouts and visual properties. Using the graph objects that you build, report designers can create simple or complex graphical views of their data.

You can do the following in the graph builder:

■ drag and drop graph elements onto the canvas. Graph elements include the plots, charts, and maps that are available from the **Graph Elements** pane.

■ combine and arrange the graph elements as needed. For example, you can layer multiple graph elements in a cell, create a lattice of cells in rows and columns, or combine these layouts.

There are two types of lattice graphs:

□ A data-driven lattice is a multi-cell graph in which the cell data is determined by the values of one or more class variables. The number of cells is determined by the unique values of the class variables. Each cell of the lattice has the same graph elements. However, these graph elements have different values of the class variables.

□ A user-defined lattice is a multi-cell graph in which you create each cell independently. Each cell can contain different types of plots.

■ determine and configure which data roles become available for assignment. For more information, see "Working with Roles" on page 480.

■ specify the visual attributes of the graph elements. You can specify properties for the graph, for individual graph elements, for the axes, and for legends. For more information, see "Work with Properties" on page 506.

Your First Look at the Graph Builder

When you create custom graph objects for use in the designer, you use the graph builder.

Here are the features of the graph builder:

Figure 40.1   *Graph Builder*
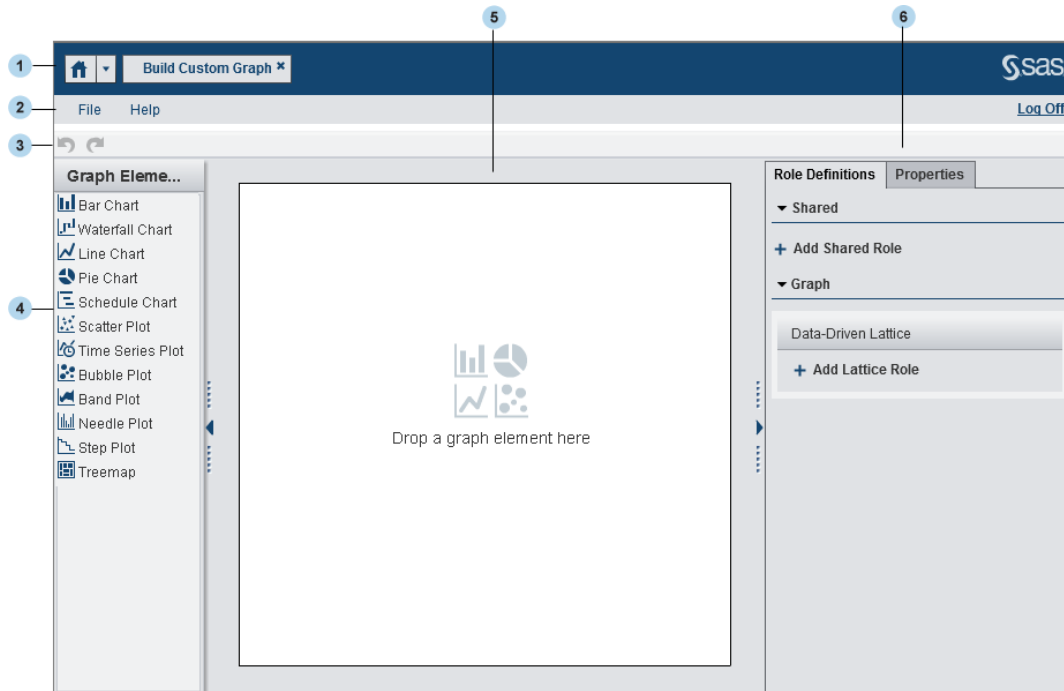
1   The application bar enables you to return to the home page and to access other parts of SAS Visual Analytics and other SAS applications that integrate with the home page. You can access your recently created or viewed reports, explorations, stored processes, data queries, or other objects in your recent history. Buttons are displayed for each open application.

2   The menu bar enables you to create, open, and save your graph objects. The **Help** menu provides short how-to topics in addition to this user's guide. From here, you can log off from SAS Visual Analytics.

3   The toolbar contains icons that enable you to undo and redo changes to your graph objects.

4   The left pane contains graph elements that you can drag and drop onto the canvas.

5   The workspace is where you build a graph object.

6   The right pane's tabs enable you to specify data roles, visual properties, and other details of your graph objects.

Graph Elements Available for Custom Graphs

The following table lists the available plots, charts, and maps that you can use to build custom graph objects. This document refers to plots, charts, and maps collectively as graph elements.

For a definition and a picture of each available graph element, see "Graphs, Charts, and Plots" on page 545.

Table 40.1   Available Graph Elements

Icon	Graph Element
�	Bar Chart
ⅈⁱⁱ	Waterfall Chart
⩘	Line Chart
◕	Pie Chart
⊟	Schedule Chart
⠿	Scatter Plot
⏱	Time Series Plot
⠿	Bubble Plot
◠	Band Plot

Icon	Graph Element
	Needle Plot
	Step Plot
	Treemap

41

Creating Custom Graph Objects

Components of a Graph Object

In general, a graph object consists of one or more cells that contain a composite of one or more graph elements.

The following figure shows the different parts of a graph object:

Figure 41.1   Components of a Graph Object
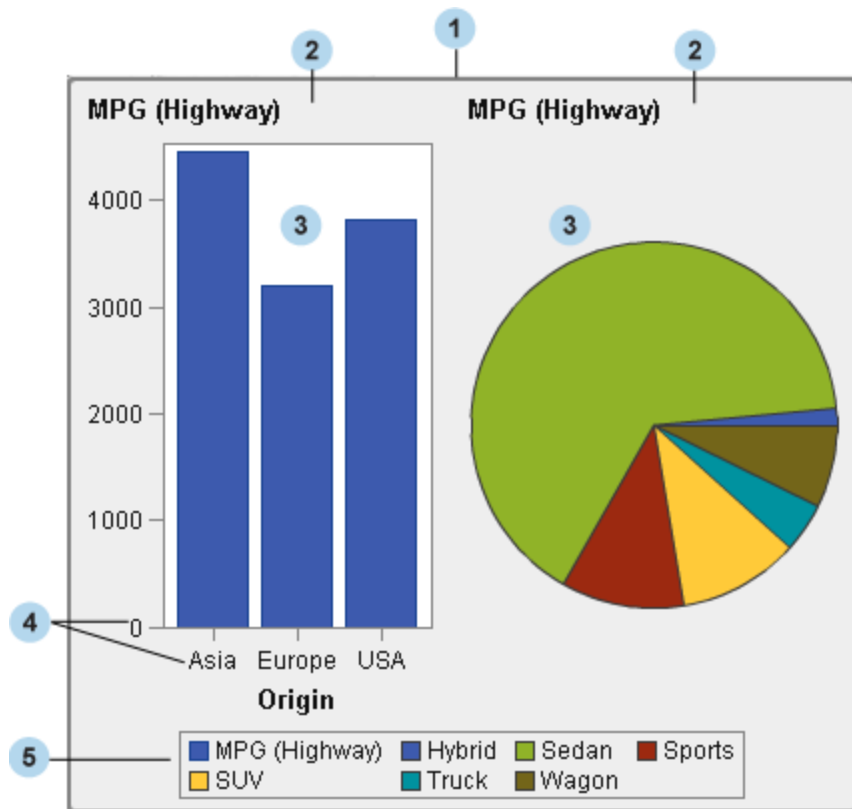
1 Graph object

   a visual representation of data. The graph object can contain one or more cells that
   have one or more graph elements.

2 Cell

   a distinct rectangular subregion of a graph object that can contain graph elements,
   axes, and legends. Multiple cells are arranged as columns and rows. The example
   graph object contains two cells that are arranged as columns.

3 Graph element

a visual representation of data, such as a scatter plot or a bar chart. The example shows a bar chart and a pie chart, each in its own cell. Multiple compatible graph elements can be overlaid in a cell.

4   Axis

refers collectively to the axis line, the major tick marks, the major tick mark values, and the axis label. In multi-cell graph objects, the columns and rows of cells can share common axes if the cells have the same data type.

5   Legend

refers collectively to the legend border, one or more legend entries (where each entry has a symbol and a corresponding label), and a legend title.

Build a Custom Graph Object

To build a custom graph object:

1   Drag and drop a graph element from the **Graph Elements** pane onto the canvas.

Sample data is assigned to the graph element. The real data will be assigned by a report designer who includes the graph object in a report.

2   (Optional) On the **Role Definitions** tab, you can define roles for the graph element.

You have the following options:

■   Provide more descriptive names for the default roles. For more information, see "About Role Names" on page 481.

■   Click **Add Role** to add more roles.

■   Create a data-driven lattice by adding a lattice role. For more information, see "Creating a Data-Driven Lattice " on page 491.

3   (Optional) You can add additional graph elements to the graph object. See "Adding a Graph Element to an Existing Graph Object" on page 484.

4   (Optional) On the **Properties** tab, you can modify the properties of your graph object.

You have the following options:

▪   Update the graph object's properties, which affect all parts of the graph object. For more information, see "Custom Graph Properties" on page 507.

▪   Update the properties for the individual cells.

▪   Update the properties for the individual graph elements. The available properties vary based on the type of graph element that you are modifying.

▪   Update the properties for the legend. For more information, see "Legend Properties" on page 508.

▪   Update the properties for one or both axes. For example, you can change the axis type, use a secondary axis (X2 or Y2), show grid lines, specify the axis range, and change the appearance of an axis. See "Axis Properties" on page 509.

5   Save the graph object. Depending on where you save the graph object, it might be immediately displayed on the **Objects** tab in SAS Visual Analytics Designer (the designer). For more information, see "Saving a Custom Graph Object So It Appears in the Designer" on page 495.

Working with Roles

About Roles

When you create a custom graph object, you determine which data roles become available for assignment. However, you do not assign data to those roles in the graph builder. When the graph object becomes available to report designers for use in their reports, the report designers assign data items to the available roles.

Here are the tasks that you can perform related to roles:

- Provide meaningful names for the roles. For more information, see "About Role Names" on page 481.

- Create additional roles. For more information, see "Create Optional Roles" on page 481.

- Share roles across two or more graph elements. For more information, see "Share Data Roles" on page 513.

- Specify a role for a data-driven lattice. For more information, see "Creating a Data-Driven Lattice " on page 491.

About Role Names

When you drag and drop a graph element onto the canvas, the required roles for that graph element are created with default names on the **Role Definitions** tab. You can specify different names for the roles by changing the default names.

When you specify the name for a role definition, specify a name that helps report designers understand which type of data to assign. Therefore, role names should be meaningful to report designers. Example names are Category, Response, Time, X variable, Y variable, Numeric X variable, Numeric Y variable, and so on. Normally, you do not specify names that reflect the particular data, such as sales items, unless you know that the graph object will be used only with that particular data.

Create Optional Roles

When you drag and drop a graph element onto the canvas, the required roles for that graph element are created on the **Role Definitions** tab. You can create additional, optional roles.

To create an optional role:

1 On the **Role Definitions** tab, click **Add Role**. The Add Role window is displayed. The contents of the Add Role window vary depending on the current graph element.

2 Select the type of role that you want to create. For more information, see "Types of Roles" on page 482.

3   (Optional) You can change the default role name to provide a more meaningful name. For more information, see .

4   Select the data type for the role. For information about the data types, see .

5   (Optional) Select **Required** to require report designers to assign data to the role. If this option is not selected, then report designers can choose whether to assign data to the role.

6   If the **Allow multiple data assignments** check box is available, then you can select or clear the check box. When this option is selected, the role can have more than one measure assigned. For example, in a line chart, if multiple measures are assigned to the Y axis, then the graph shows a line for each measure.

7   Click **OK**.

Types of Roles

In addition to the standard roles, such as Time, X, or Category, you can add the following roles by clicking **Add Role** on the **Role Definitions** tab. The roles that are available depend on the type of graph element.

■   The Group role specifies a data column that is used to group the data. The graph elements for each unique group value are automatically distinguished by different visual attributes, such as color, line style, and so on. The Group role can be a Category or a Datetime data type.

■   The Color role specifies a data column that is used to color the data.

    The following example shows the heights of a group of students. A student's gender is specified for the Color role.
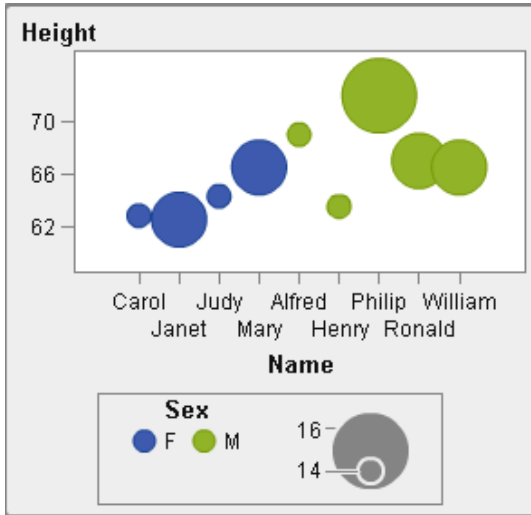
The Color role can be any data type. In this example, the Color role is assigned a Category data type.

Note: When you add the Color role to a graph element in a data-driven lattice, the Measure data type is not always available. The Measure data type is not available for a bar chart, a waterfall chart, or a scatter plot.

In some graph elements, such as a bar chart, the Group role also colors the graph based on the values of the group variable. In those graph elements, the Color role is used to color the data across a color gradient as shown in the following example. In this case, the Color role must be a Measure data type.

In this example, the color gradient indicates the age range of the students.
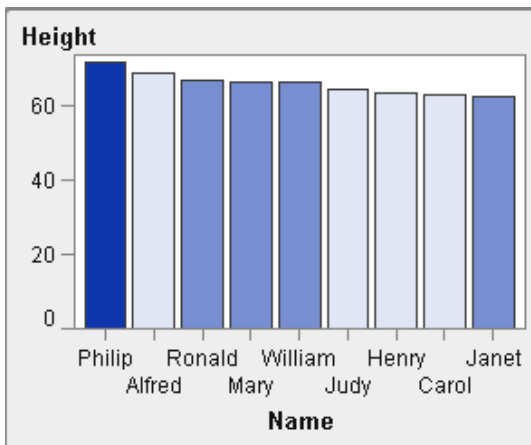
- The Data Tip role specifies a data column that is used for data tips. The Data Tip role must be numeric.

Adding a Graph Element to an Existing Graph Object

You have two options for adding a graph element to an existing graph object. You can overlay the new graph element on top of an existing graph element. Or, you can add the graph element so that it creates a cell in a new row or column.

Add an Overlay

In an overlay, one graph element is superimposed on another graph element. You can add an overlaid graph element to a single-cell graph as well as to a lattice.

In this example, a step plot is overlaid on a bar chart. The example shows the graph element with data assigned to it.

Display 41.1   Step Plot Overlaid on a Bar Chart
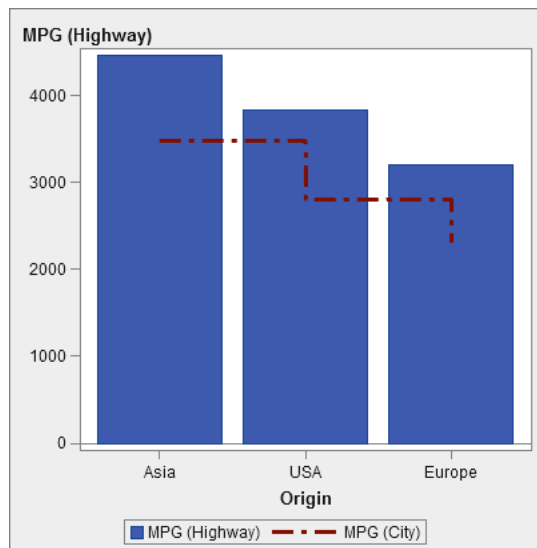
To create an overlay:

1   Drag and drop a graph element from the **Graph Elements** pane onto the canvas.

2   Drag and drop a second graph element directly onto the existing graph element.

   If the graph elements are compatible, they are automatically assigned a shared role. This enables both graph elements to have a common X axis role. You can unshare the shared role if needed.

   Some graph elements, such as the pie chart and the treemap, cannot be layered in an overlay. A message is displayed when you try to drag and drop these types of graph elements.

   For more information, see the following topics:

   ■   "Compatible Graph Elements" on page 488

   ■   "Share Data Roles" on page 513

3   (Optional) On the **Role Definitions** tab, you can change the default role names to provide more meaningful names. For more information, see "About Role Names" on page 481.

> **TIP** The graph elements are layered in the order in which they were added to the canvas. You can change the order by moving a graph element to the front or the back. For more information, see "Change the Order of Graph Elements in a Cell " on page 500.

Add a Graph Element to a New Row or Column

A graph element can be added so that it creates a cell in a new row or column. This arrangement results in a graph object with more than one cell. It is called a user-defined lattice. In a user-defined lattice, each cell is created independently and can contain different types of graph elements. Cells are aligned in rows and columns that you create by dragging and dropping graph elements.

Note: You cannot create new rows or columns in a data-driven lattice. For a description of a data-driven lattice, see "About Data-Driven Lattices " on page 491.

This example shows a user-defined lattice with four cells. One of the cells is empty. The example shows the graph element with data assigned to it.

Display 41.2 *User-Defined Lattice*
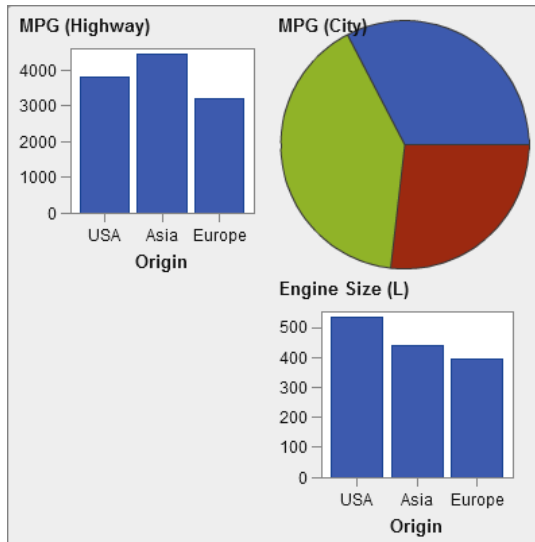
To create a user-defined lattice:

1 Drag and drop a graph element from the **Graph Elements** pane onto the canvas. The graph element is created in the center of the graph area of the canvas.

2 Drag and drop another graph element onto the canvas as follows:

- Drag and drop the graph element onto the right or left border of the existing graph to create a new column. The area becomes shaded to indicate the location of the new column.
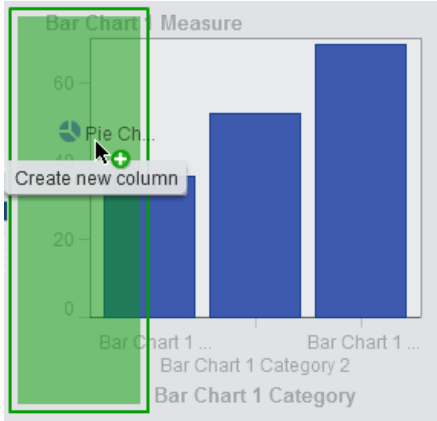
■ Drag and drop the graph element onto the top or bottom border of the existing graph to create a new row.
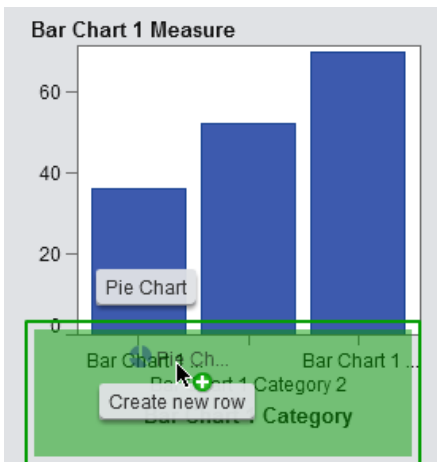
3 To add more cells, continue dragging and dropping graph elements onto the borders of the existing graph to create new columns or rows. You can drag and drop a graph element onto the area between two columns to insert a new column. Or, you can drag and drop a graph element onto the area between two rows to insert a new row. If the existing graph contains a lattice with an empty cell, you can drag and drop the graph element into that empty cell.

You can create up to 10 rows and 10 columns.

4 (Optional) On the **Role Definitions** tab, you can change the default role names to provide more meaningful names. For more information, see "About Role Names" on page 481.

Compatible Graph Elements

Some graph elements can be layered in an overlay with one graph element superimposed on another. Other graph elements, however, are incompatible. They cannot be layered in the same cell. The following table shows which graph elements can be layered in the same cell.

Note: Pie charts and treemaps are not compatible with any other graph elements. They are not listed in the table.

Table 41.1   Graph Element Type Compatibility

Graph Element	Compatible	Not Compatible
Band plot	Band plot Bar chart, vertical Bubble plot Line chart Needle plot Scatter plot Step plot Time series plot Waterfall chart	Bar chart, horizontal Schedule chart
Bar chart, horizontal	Bar chart, horizontal Bubble plot Scatter plot Step plot Time series plot	Band plot Bar chart, vertical Line chart Needle plot Schedule chart Waterfall chart

Graph Element	Compatible	Not Compatible
Bar chart, vertical	Band plot Bar chart, vertical Bubble plot Line chart Needle plot Scatter plot Step plot Time series plot Waterfall chart	Bar chart, horizontal Schedule chart
Bubble plot	(Compatible with all graph elements)	
Line chart	Band plot Bar chart, vertical Bubble plot Line chart Needle plot Scatter plot Step plot Time series plot Waterfall chart	Bar chart, horizontal Schedule chart
Needle plot	Band plot Bar chart, vertical Bubble plot Line chart Needle plot Scatter plot Step plot Time series plot Waterfall chart	Bar chart, horizontal Schedule chart

Graph Element	Compatible	Not Compatible
Scatter plot	(Compatible with all graph elements)	
Schedule chart	Bubble plot Scatter plot Schedule chart	Band plot Bar chart, horizontal Bar chart, vertical Line chart Needle plot Step plot Time series plot Waterfall chart
Step plot	Band plot Bar chart, horizontal Bar chart, vertical Bubble plot Line chart Needle plot Scatter plot Step plot Time series plot Waterfall chart	Schedule chart
Time series plot	Band plot Bar chart, horizontal Bar chart, vertical Bubble plot Line chart Needle plot Scatter plot Step plot Time series plot Waterfall chart	Schedule chart

Graph Element	Compatible	Not Compatible
Waterfall chart*	Band plot Bar chart, vertical Bubble plot Line chart Needle plot Scatter plot Step plot Time series plot	Bar chart, horizontal Schedule chart Waterfall chart

* The waterfall chart requires the X axis to have a shared role to avoid unpredictable results.

Creating a Data-Driven Lattice

About Data-Driven Lattices

A data-driven lattice is a multi-cell graph in which the cell data is determined by the values of one or more class variables. Class variables are specified in the designer when the report designer creates a report that includes the lattice. The number of cells is determined by the unique values of the class variables. Each cell of the lattice has the same graph elements. However, these graph elements have different values of the class variables.

The following example shows a graph with data assigned to it. The custom graph object includes a data-driven lattice with a columnar role. The graph object allows multiple data assignments for the Measure role.

Display 41.3 *Data-Driven Lattice Multi-Cell Graph*
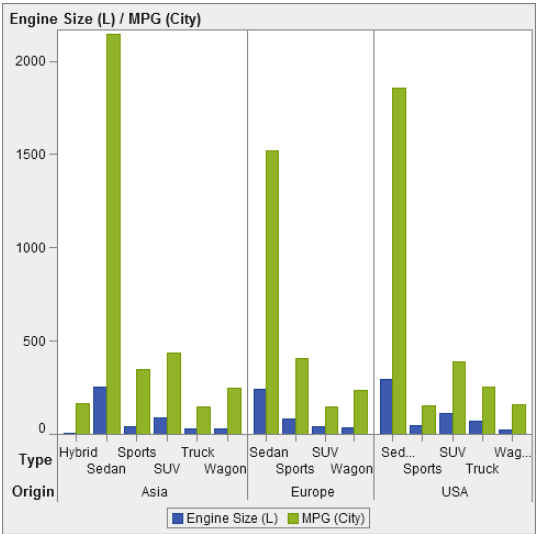
In the report, the graph object includes a three-cell lattice of bar charts showing engine size and miles per gallon for different types of automobiles. The country of origin is specified as the class variable. As a result, a cell is created for each country of origin.

The following example shows the graph object that was created in the graph builder. The lattice role was specified for columns.
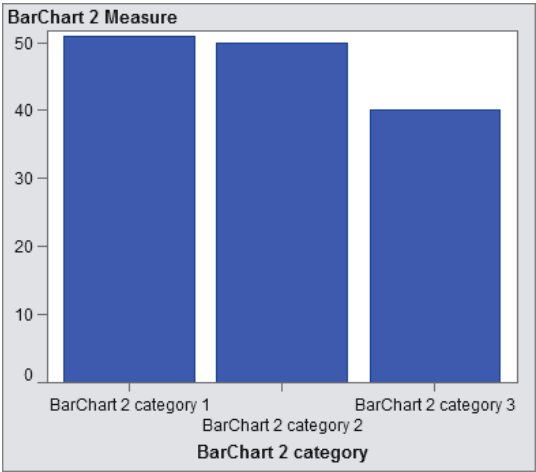
Here are the characteristics of data-driven lattices:

- You can define up to two lattice roles: one role for the rows and another role for the columns. When the report designer uses the lattice in a report, the class variables are assigned to the lattice roles. All class variables must come from a single data source. The number of rows and columns in the graph is determined by the unique values of the class variables.

- If you specify that the role is required, then a report designer must assign data to the role.

  If this option is not selected, then report designers can choose whether to assign data to the role. In that case, it is possible for a graph object to have no lattice.

- Graph elements can be overlaid the same way they are overlaid in single-cell or user-defined multi-cell graphs. When the report designer adds a lattice to a report, the overlaid graph element is displayed in every cell.

- Properties that you specify for graph elements are applied to all of the cells of the graph object.

Create a Data-Driven Lattice

To create a data-driven lattice:

1 Create the graph object that you want as the basis for your lattice. For more information, see "Build a Custom Graph Object" on page 479.

  Customize the graph object as appropriate. For example, you might add an overlaid graph element or change some properties.

2 (Optional) On the **Role Definitions** tab, you can change the default role names to provide more meaningful names. For more information, see "About Role Names" on page 481.

3 Specify the role for a class variable. In the **Data-Driven Lattice** section of the **Role Definitions** tab, click **Add Lattice Role**.

  The Add Role window is displayed.

The **Data-Driven Lattice** section is not available in either of the following cases:

- the graph object already contains more than one cell

- a Color role with a Measure data type has been added to the bar chart, waterfall chart, or scatter plot

In addition, data-driven lattices do not allow graph elements to be overlaid unless their aggregations are compatible. Graph elements with compatible aggregations either perform no aggregation at all or they use the same set of shared Category roles.

4 For the **Role Type**, select **Row** or **Column** to specify the layout for the cells.

5 Specify a role name or use the default name.

6 Select the data type for the role. For information about the data types, see "Working with Data Items in a Report" on page 352.

7 (Optional) Select **Required** to require report designers to assign data to the role. If this option is not selected, then report designers can choose whether to assign data to the role. In that case, it is possible for a graph to have no lattice.

8 If the **Allow multiple data assignments** check box is available, then you can select or clear the check box. When this option is selected, the role can have more than one measure assigned. For example, in a line chart, if multiple columns are assigned to the Y axis, then the graph shows a line for each measure.

9 Click **OK**.

Repeat the previous steps to specify a second role. For example, if you specified a role for a multi-row layout, you can specify an additional role for columns.

Saving a Custom Graph Object So It Appears in the Designer

Custom graph objects are saved when you select either **File** ▶ **Save** or **File** ▶ **Save As**. Graph objects that are saved in the `My Folder` location are automatically displayed in the list of **Custom** objects on the designer's **Objects** tab.

Graph objects that you save in public folders are not automatically displayed on the **Objects** tab. However, you can add the graph objects manually.

To add a graph object to the designer's **Objects** tab:

1   On the **Objects** tab in the designer, click ▼ , and select **Show or Hide Objects**. The Show or Hide Objects window is displayed.

2   Click **Select custom**. The Choose an Item window is displayed.

3   Navigate to the graph object that you want to add.

4   To search for a graph object, do the following:

   a   Click 🔍 in the toolbar. The Search window is displayed.

   b   Enter the name of the graph object that you want to find.

   c   Click **Search**.

   > **TIP** As an alternative, you can search for all graph objects. Specify **Graph template** for **Type**.

5   Select the graph object, and click **OK**.

   The selected graph object is now displayed with a check mark in the Show or Hide Objects window.

6   Click **OK**. The **Objects** tab is updated.

42

Modifying Custom Graph Objects

Change a Custom Graph That Has Been Saved

When you save a custom graph object, depending on where you save it, it might be immediately displayed on the **Objects** tab in SAS Visual Analytics Designer (the designer). For more information, see "Saving a Custom Graph Object So It Appears in the Designer" on page 495.

To change a graph that has been saved, do either of the following:

- Right-click the graph object in the designer, and select **Edit Custom Graph**.

- Select **File ▶ Open**, and then locate and open the file for your graph object.

Note: Any changes that you make to the graph object do not appear in existing reports that were created from the original graph object. However, new reports that are created after a modified graph object has been saved reflect the changes to the graph object.

Select Components and Elements of a Graph Object

To change the properties of a graph element such as a bar chart, you must first select the graph element. Similarly, to change the properties of a legend or an axis, you must first select the component.

For a description of the graph components, see "Components of a Graph Object" on page 477.

When you select a component in a graph object, the selected component is highlighted. The component is also displayed in the drop-down list at the top of the **Properties** tab. The available properties vary depending on the selected component.

The following display shows a highlighted bar chart. The **Properties** tab displays the properties for the bar chart, which is identified as **Bar Chart 1**.
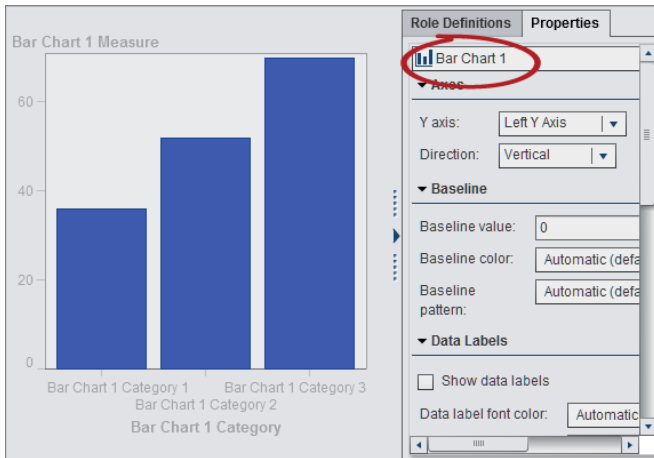
> **TIP** You can click ▼ beside the drop-down list on the **Properties** tab, and then select the component that you want to modify.

Here are some suggestions for selecting graph elements:

- To select a line chart or a time series plot, click a data point on a line.

- To select an axis, click directly on the axis.

- To select the entire graph object, click above the graph element or cell area.

- To select a cell in a multi-cell graph object, click the background area just above the graph element. The cell is highlighted, and a tab appears in the upper left corner of the cell. From the cell's tab, you can delete and reorder graph elements. You can move graph elements to other cells.
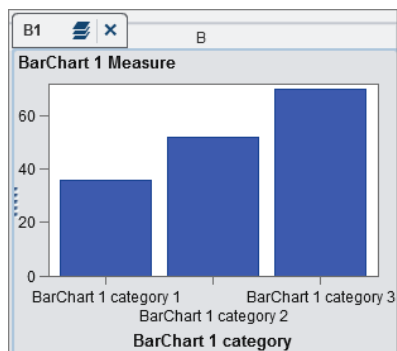
In a multi-cell graph object, cells are identified by a combination of a letter and a number, such as A1, B1, A2, and so on. The letters identify columns and the numbers identify rows. The following table represents the labeling for a multi-cell graph object that contains three columns and three rows:

A1	B1	C1
A2	B2	C2
A3	B3	C3

Change the Order of Graph Elements in a Cell

When you overlay a graph element on top of another, the graph elements are layered with the last one added on top. For example, suppose that you create a graph object with a line chart, and then you add a bar chart. Because the bar chart was added last, it appears in front of the line chart. Depending on the data that is assigned to the graph elements in the designer, the bar chart might obscure part of the line chart. You can reorder the graph elements so that the line chart appears in front of the bar chart.

To change the order of the graph elements in a cell:

1   From the drop-down list at the top of the **Properties** tab, select **Cell**. For more information, see "Select Components and Elements of a Graph Object" on page 498.

In the **Graph Elements** area on the **Properties** tab, the graph elements in the cell are listed in the order in which they were added to the cell.

2   Drag a graph element to a new position.

Remove a Graph Element

To remove a graph element from a graph object, right-click the graph element, and select **Remove `graph element name`**.

Note: In a user-defined lattice, you can remove graph elements using the cell's tab. For more information, see "Remove a Cell or a Graph Element in the Cell" on page 504.

Work with User-Defined Lattices

About User-Defined Lattices

A user-defined lattice is a multi-cell graph in which each cell is created independently. And, each cell can contain different types of graph elements. Cells are aligned in rows and columns that you create by dragging and dropping graph elements.

The following example shows four cells arranged in two rows and two columns. Letters identify the columns and numbers identify the rows.

Display 42.1   *User-Defined Lattice*
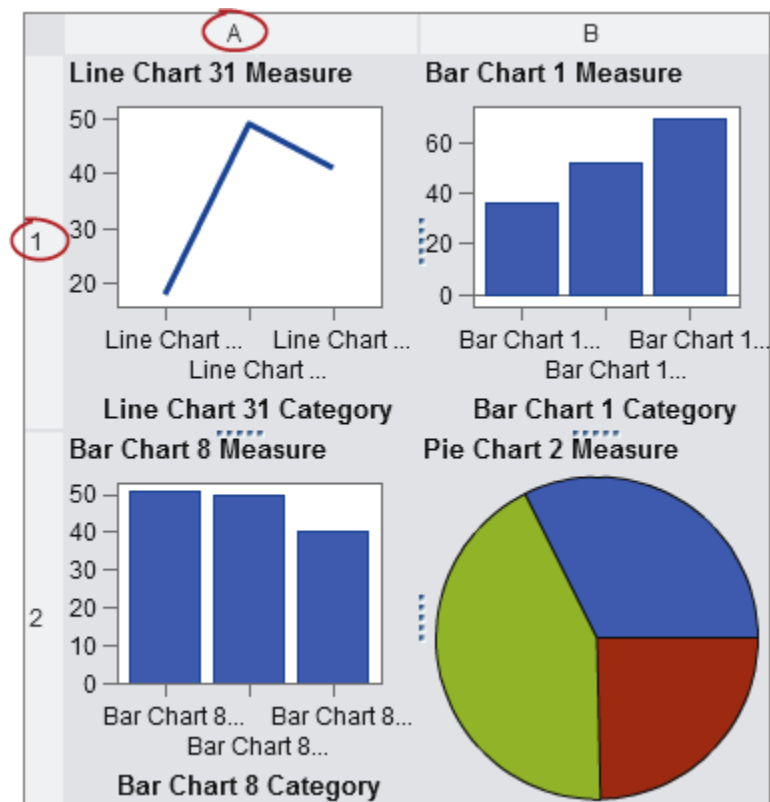
When you select a cell, a tab appears in the upper left corner of the cell.
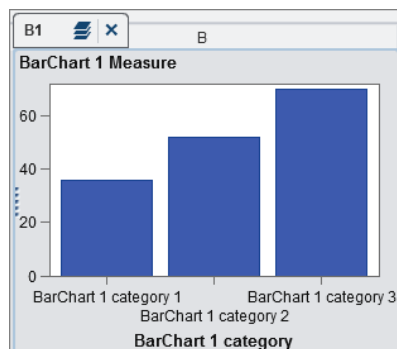
From the cell's tab, you can delete the cell and its contents by clicking ✖. If you click ≣, the graph elements in the cell are listed. From this list, you can delete and reorder the graph elements. You can move graph elements to other cells.

For more information about selecting cells and how they are labeled in the graph, see "Select Components and Elements of a Graph Object" on page 498.

See Also
"Components of a Graph Object" on page 477

Move a Graph Element from One Cell to Another Cell

To move a graph element to a different cell:

1   Select the cell whose graph element you want to move. For more information, see "Select Components and Elements of a Graph Object" on page 498.

2   In the cell's tab, click ≣ to list the graph elements in the cell.

3   Drag and drop a graph element from one cell into the target cell.

    You can also drag and drop the graph element so that it creates a new row or new column. For more information, see "Add a Graph Element to a New Row or Column" on page 485.

Resize a Row or a Column

You can change the width of a column and the height of a row in a user-defined lattice. For example, you might want a particular column to be wider than the others.

To resize a row:

1   Select the resizing handle ⋯ between the row that you want to change and an adjacent row.

2   Drag the handle upward or downward to change the row height.

To resize a column:

1 Select the resizing handle ⋮ between the column that you want to change and an adjacent column.

2 Drag the handle left or right to change the column width.

Add a Row or a Column to a Lattice

To add a new row or column to a lattice, drag and drop a graph element to the location where you want the new row or column. For more information, see "Add a Graph Element to a New Row or Column" on page 485.

You can also move a graph element from a row or column to another row or column. For more information, see "Move a Graph Element from One Cell to Another Cell" on page 503.

Remove a Cell or a Graph Element in the Cell

You can remove a graph element from the cell. You can also remove the cell itself.

1 Select the cell.

2 To remove a graph element from the cell:

   a Click ☰ to list the graph elements in the cell.

   b Click ✕ next to the graph element.

   You can also right-click the graph element, and select **Delete**.

3 To remove the cell, click ✕ on the cell's tab.

   The cell and its contents are removed. If the cell is the only cell in a row or column, then the respective row or column is removed as well.

Move or Delete a Row or a Column

In a graph, a heading contains a letter that identifies the column and a number that identifies the row.

From the headings, you can move and delete columns and rows. When you move a column or row, the column or row switches places with the adjacent column or row. For example, if you move column B to the left, column B switches places with column A.

To more or delete a column or row:

1 Select the graph.

2 Position the cursor over a heading. A button and arrows are displayed on the heading.

The following example shows a column heading that contains a button to delete the column and arrows to move the column. The button to move the column to the right is unavailable.
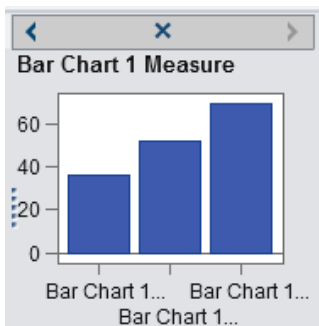
3 To delete the column or row, click ✕.

4 To move the column or row, click an arrow.

  ▪ Click ‹ or › to move a column left or right.

  ▪ Click ⌃ or ⌄ to move a row up or down.

Some arrows might not be available. For example, the rightmost column can be moved only to the left, as is the case in the previous example. Only ‹ is available for the column.

Work with Properties

About Properties

The **Properties** tab lists the properties for the selected graph component. Components that have properties include graph elements, cells, axes, legends, and the entire graph. Properties enable you to control the overall appearance of your graph, from the graph's background color to the tick value interval on the axis. Report designers can override some properties in their report.

The following sections provide more information about properties.

See Also

- "Select Components and Elements of a Graph Object" on page 498
- "Components of a Graph Object" on page 477

About Style Properties

The color schemes and visual attributes for a graph come from the active report theme, which is a named collection of style elements. Each style element contains attributes such as fill color, marker symbol, line style, font face, and many others. Each part of a graph element, such as a marker, a bar, a line, or a label, derives its visual attributes from a specific style element from the active report theme. Similarly, the attributes for axes and legends come from a style element.

The style elements are developed to produce effective graphics without changing the default settings. However, you can change the default settings by changing style properties on the **Properties** tab. When you specify a style property, such as a color or font style, you are changing the style element that is applied to the graph component.

The following display shows the available style elements for bar chart labels. The list of style elements is displayed when you click ▼ beside the **Data label font color** drop-down list on the **Properties** tab.
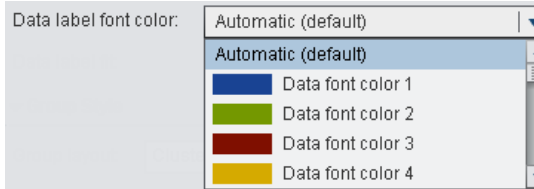
Each color is shown in the list beside its style element. **Automatic** indicates the default value for the graph element. **Automatic** enables the graph element to cycle through a list of colors provided for group colors and multi-response colors.

Visual attributes that are referenced by a style element provide consistency within the report theme. If you later change the report theme, the graph element is compatible with the new theme. In addition, some of the properties that you specify can be overridden by report designers when they use your graph objects in their reports.

Custom Graph Properties

Custom graph properties determine features that affect all parts of the graph object.

Here are some changes that you can make on the **Properties** tab when the graph object is selected:

- customize the data tips and show or hide missing data.

- specify an overview axis for the graph. When you specify an overview axis, you can select which graph elements to include in the overview axis. For more information, see "Create an Overview Axis" on page 510.

- specify how graph attributes are rotated for group data. For more information, see "Understanding and Changing Attribute Rotation" on page 511.

- change certain appearance attributes of the graph, including border lines and the graph background.

Cell Properties

Cell properties determine features that affect the overall appearance of the cell. These include the background and border colors for the cell, which graph elements the cell contains, and the color scheme for overlaid graph elements.

To change the properties for a cell, select the cell in the graph object, and make your changes on the **Properties** tab.

Graph Element Properties

Graph element properties determine features that affect the appearance of the plots, charts, and maps in a graph. The types of properties that you see vary depending on the type of graph element that you are modifying. In general, you can specify colors, marker symbols, line attributes, and outline and fill colors.

To change the properties for a graph element, select the graph element, and make your changes on the **Properties** tab.

Legend Properties

By default, a legend always appears when the graph object contains two or more graph elements. You can remove graph elements from the legend. If you remove all graph elements, then the legend disappears from the graph object.

To change the properties for a legend, select the legend in the graph object, and make your changes on the **Properties** tab.

Here are some changes that you can make to the legend:

■ show or hide the border and specify its appearance.

■ specify the formatting and style for the title.

■ specify which graph elements are included in the legend. In the **Display in legend** section, select or clear graph elements.

Axis Properties

The types of properties that you see vary depending on the type of axis that you are modifying.

Here are the possible types:

- Discrete: The axis contains independent data values rather than a range of numeric values. Each distinct value is represented by a tick mark. Discrete is the default axis type for character data. Some graph elements require a discrete axis. For example, bar charts use a discrete axis for the Category role.

- Linear: The axis contains a linear range of numeric values. Linear is the default axis type for numeric data. Some graph elements require a linear axis. For example, bar charts use a linear axis for the Measure role.

- Logarithmic: The axis contains a logarithmic range of values. The logarithmic axis type is not used as a default.

  You might want a logarithmic axis type depending on the nature of the data. For example, suppose that your graph element shows growth data with a linear axis. Based on the values, you suspect that the growth rate is exponential. If the data contains a very large range of values (orders of magnitude apart), you can display the values on a logarithmic scale.

  Note: The data range on a logarithmic axis must be positive.

- Time: The axis contains a range of date, time, or date-time values.

To change the properties for an axis, select the axis in the graph object, and make your changes on the **Properties** tab.

Here are some changes that you can make to the axis:

- reverse the axis order.

- show or hide the axis elements, such as the axis line, tick marks, and grid lines.

- change the appearance of tick marks, tick values, axis labels, and grid lines. You can also highlight the background.

Note: When a custom graph object is used in a report, if the data assigned to the graph object contains a very large quantity of data, tick values might not be displayed. The user can use the axis scroll bar to zoom in on a particular area of the graph to see the tick values.

◼ merge two or more column axes or two or more row axes.

◼ specify the axis range, including offsets and first and last ticks.

◼ specify a secondary axis for a graph element. For example, you can use the left or the right Y axis for the Measure role of a bar chart.

Note: The secondary axis is changed at the graph element level. To specify a secondary axis, first, select the graph element. Then, in the **Axes** section of the **Properties** tab, select the axis that you want to use.

Create an Overview Axis

The overview axes that you create in the graph builder are similar to the overview axes that you create in the designer. In addition, you can do the following:

◼ select specific graph elements from the main graph to show in the overview axis

◼ select completely different graph elements from the main graph to show in the overview axis

To create an overview axis:

1 Drag and drop a graph element from the **Graph Elements** pane onto the canvas.

   You can drag and drop another graph element to create an overlay. However, do not create a new cell for this additional graph element. You cannot create an overview axis when the graph object contains more than one cell.

2 Select the graph object.

3 On the **Properties** tab, select the **Show overview axis** check box.

   The graph object displays the overview axis beneath the existing graph elements. By default, the existing graph elements are displayed in the overview axis, as shown in the following example:
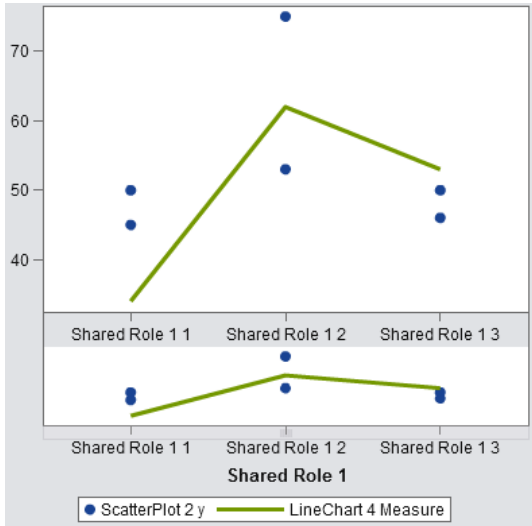
4   To display different graph elements in the overview axis, do the following:

a   On the **Properties** tab, select **Specify graph elements** from **Contents**.

b   Drag and drop one or more graph elements from the **Graph Elements** pane onto the overview axis area of the canvas.

The graph element names that you added are displayed in the **Graph elements** section on the **Properties** tab.

c   If you added more than one graph element to the overview axis area, you can change the order of the graph elements by dragging and dropping a graph element to a different position.

Understanding and Changing Attribute Rotation

When you apply a group role to your graph element, by default, the graph builder rotates through predefined style elements to determine how to present each unique group value. This type of rotation also occurs when multiple response roles are assigned to the graph element or when graph elements are overlaid.

When the graph element is rendered, the attributes for colors, line patterns, and marker symbols are rotated for the values in your group. Whether you change the attributes or keep the default attribute values, the manner in which the attributes are combined is determined by the attribute rotation priority.

To change the attribute rotation priority:

1 Select the graph object. For more information, see .

2 On the **Properties** tab, select an option from **Attribute rotation priority**.

You have two options for the attribute rotation priority:

Rotate color only until all colors used	attributes such as marker symbols and line patterns are held constant while each color in the list is applied exhaustively to the graph element.
Rotate all attributes	each attribute rotates through its own list to generate a unique combination for each group value. Attributes include colors, marker symbols, line patterns, and so on.

For example, the following figure shows an example rotation for an age grouping. The dotted-line pattern is held constant while red, green, and blue colors are applied to the dotted lines for the consecutive age group values. If there are more group values, red, green, and blue colors are applied to the solid lines.
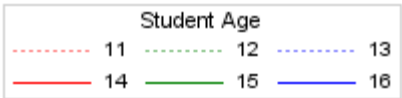
If you change the rotation pattern, and specify **Rotate all attributes** for the attribute rotation priority, the contrasting colors and line patterns are rotated at the same time.

A red dotted-line pattern is applied to the first group crossing. Then, a green solid-line pattern is applied to the second crossing, a blue dotted-line pattern is applied to the third crossing, and so on.

Student Age		
··········· 11	────── 12	·········· 13
────── 14	·········· 15	────── 16
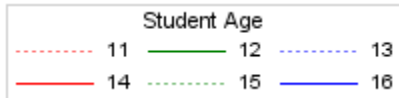
Share Data Roles

About Shared Data Roles

If your graph object contains multiple graph elements, you can specify that some data roles be shared across two or more graph elements. You share data roles when you want the graph elements to use the same data role for either of their axes. When report designers who use the graph object in their reports assign roles, a single data column is assigned to all roles that use the shared role.

Note: If the graph element types are compatible, overlaid graph elements are automatically assigned a shared role. (See "Compatible Graph Elements" on page 488.) You can unshare a role if needed.

Shared roles are very useful in user-defined lattices. They enable you to manage data roles across columns and rows. For example, you might share data roles if you want to analyze relationships in your data.

In this example, the X axis of the scatter plot shares a role with the Measure axis of the bar chart.
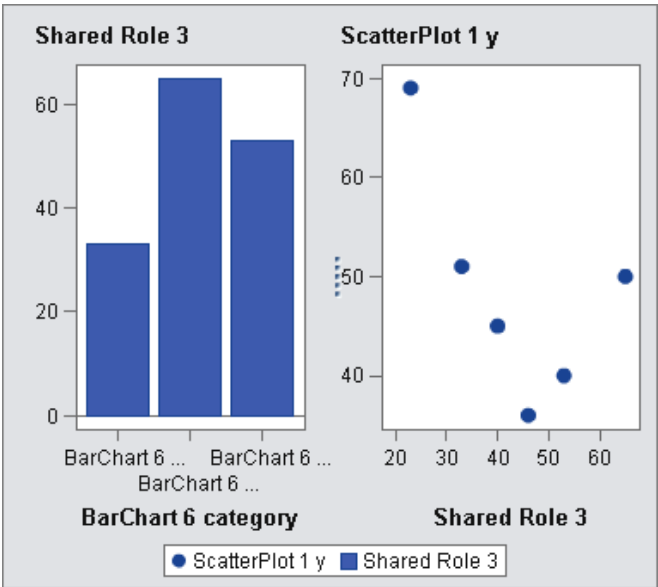
When the graph object is used in the designer, the data might be applied as shown in the following example.

The bar chart shows engine sizes for different regions of the world. (The engine size has been changed from a Sum to an Average measure.) The scatter plot shows the miles-per-gallon trend for the different engine sizes.
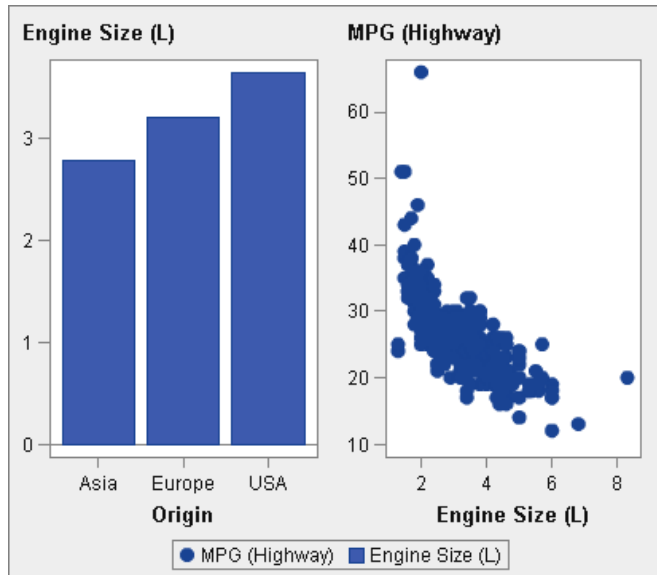
Share or Unshare a Data Role

To share a data role, on the **Role Definitions** tab, click ▼ next to the role that you want to share. Select **Create Shared Role With Another Role** ▶ *Name of Role*.

You can also click **Add Shared Role**, and then create and share the data role.

To unshare a data role, click ▼ next to the shared role, and select **Unshare**.

43

Graph Builder Examples

Example: Data-Driven Lattice

About the Data-Driven Lattice Example

This example shows the sales for a line of retail products over a period of time. The custom graph object specifies a row-based lattice in which the rows correspond to the corporate divisions that are responsible for the sales. The corporation has two divisions, resulting in a graph with two rows.

Display 43.1 *Example Data-Driven Lattice*
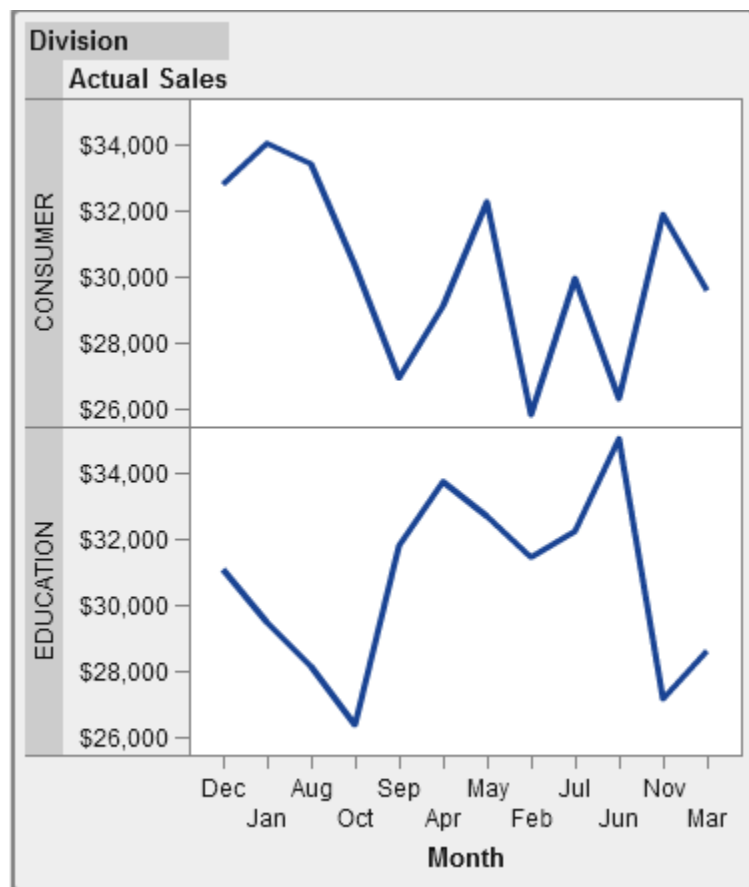
Build the Graph Object for the Data-Driven Lattice Example

1 Drag and drop a line chart from the **Graph Elements** pane onto the canvas.

2 On the **Role Definitions** tab, click ▼ next to **Category**, and select **Edit Role**. The Edit Role window is displayed.

3 Select **Datetime** for **Data Type**.

Click **OK**.

Specifying **Datetime** forces the report designer to assign date and time data to the Category role.

4   Specify the role to which a class variable is assigned.

   a   In the **Data-Driven Lattice** section on the **Role Definitions** tab, click **Add Lattice Role**. The Add Role window is displayed.

   b   Select **Row** for **Role Type**.

   c   Select **Category** for **Data Type**.

   d   Select **Required**. With this option selected, the report designer must assign data to this role to render the report with data assigned.

   e   Click **OK**.

5   Save the graph object. See "Saving a Custom Graph Object So It Appears in the Designer" on page 495.

Example: User-Defined Lattice (Butterfly Chart)

About the User-Defined Lattice Example

This example uses a butterfly chart to show the actual sales compared to the predicted sales for a line of retail products. The butterfly chart is useful for comparing two unique values. In this chart, the two values are arranged on each side of the Y axis.

Display 43.2   *Example User-Defined Lattice*
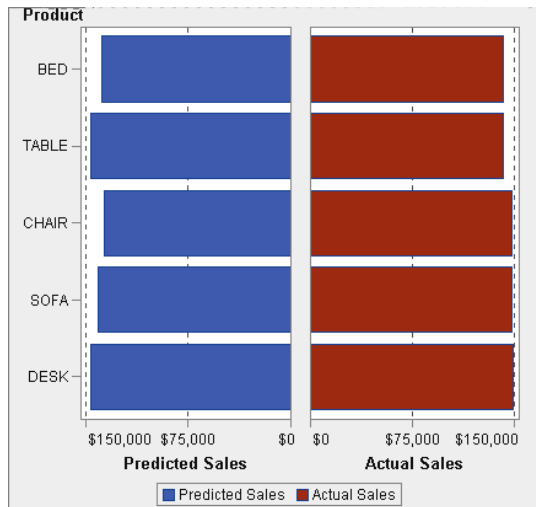
Build the Graph Object for the Butterfly Example

1 Drag and drop a bar chart from the **Graph Elements** pane onto the canvas.

2 Drag and drop a second bar chart from the **Graph Elements** pane onto the left edge of the canvas. This action creates a new column for the second bar chart.

3 Share the category roles. On the **Role Definitions** tab, click ▼ next to the **Category** for either bar chart. Select **Create Shared Role With Another Role** ▶ *Other graph element*.

   The Add Shared Role window is displayed.

   Click **OK**.

4 Change both bar charts to a horizontal layout.

   a Select a bar chart.

   b On the **Properties** tab, click ▼ next to **Direction**, and select **Horizontal**.

c Repeat the previous two steps for the other bar chart.

5 Specify a uniform column and row axis.

a Select the full graph (click the A or B heading above the bar chart).

b On the **Properties** tab, click ▼ next to **Column axis range**, and select **Uniform across all columns**.

c On the **Properties** tab, click ▼ next to **Row axis range**, and select **Uniform across row**.

The bar charts should resemble the following:
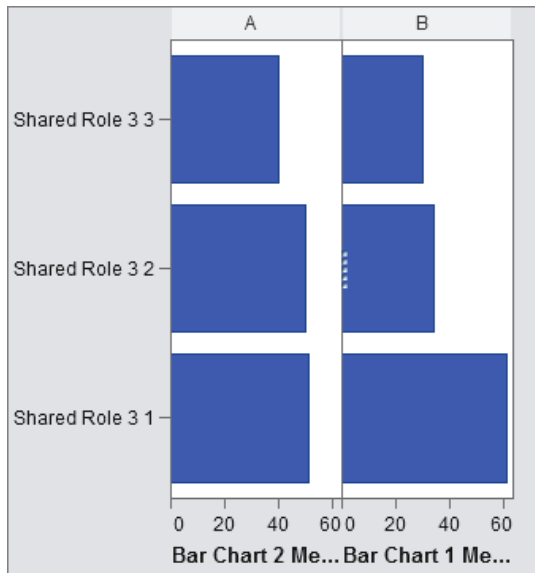
6 Reverse the order of the X (Measure) axis for the left bar chart.

a Click the horizontal axis for the left bar chart.

b On the **Properties** tab, select **Reverse order**.

The bar charts resemble the following:
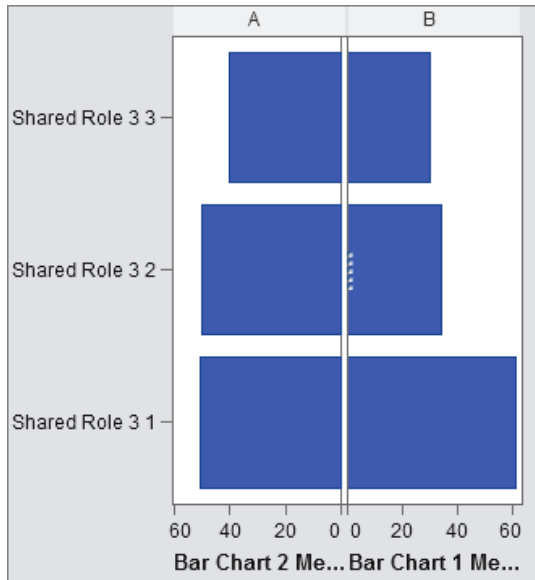
7   Change the color of the right bar chart to distinguish it from the left bar chart.

   a   Select the right bar chart.

   b   On the **Properties** tab, click ▼ next to **Fill color**, and select **Data color 3**. The right bar chart changes to a different color. In the default theme, the color is deep red.

8   Display the grid lines for the X axes.

   a   Click the horizontal axis for the left bar chart.

   b   On the **Properties** tab, select **Show grid lines**.

   c   Repeat the previous two steps for the right bar chart.

9   Save the graph object. See "Saving a Custom Graph Object So It Appears in the Designer" on page 495.

The final graph object, complete with the legend, resembles the following:
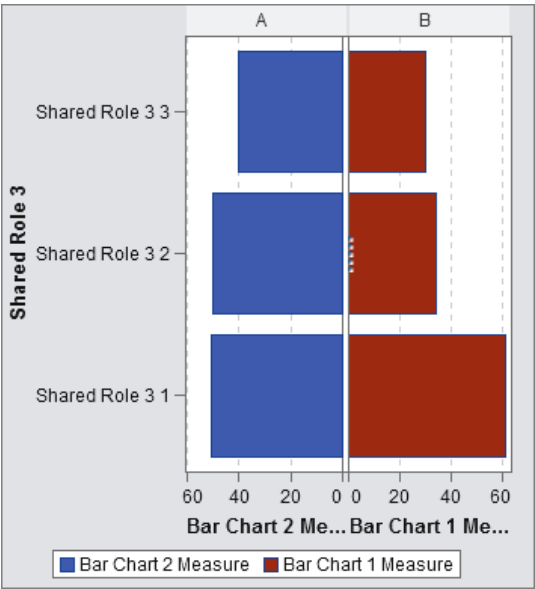
Part 6

Viewing Reports

44

Viewing Reports on a Mobile Device

Opening the App for the First Time

You can use the SAS Mobile BI for iPad or Android app to view SAS Visual Analytics reports on your mobile device. The SAS Mobile BI for iPad app is available in the iTunes App Store. The Android app is available from Google Play.

When you open either app for the first time, sample reports are available in My Portfolio. The app provides online Help for using SAS Mobile BI. You can view the online Help in My Portfolio. The online Help also provides access to tutorial videos about how to use the app.

Add a Server Connection for iPad

When you first open the SAS Mobile BI for iPad app, the only connection that is available is to the SAS Demo Server.

To access your organization's library server, you must add a connection to that server. If your organization has more than one server, you can add additional connections in the same way.

Note: Contact your SAS Visual Analytics system administrator for your user ID, password, and server information.

To connect to a server:

1 If you are in My Portfolio, tap **Library**.

2 Tap **Connections** in the library.

3 Tap **Add Connections**.

4 Tap the **Server** field, and enter the address of the new server.

   Note: If the server requires a secure (SSL) connection, type *https://* at the beginning of the server address.

5 Tap the **User ID** field, and enter your user ID.

6 Tap the **Password** field, and enter your password.

7 Tap the **Description** field, and enter a description for the new connection.

8 Tap **Next** to verify the connection.

9 Tap **Save**. The connection is saved, and the Connections window closes.

   Note: If the connection fails, a message is displayed to help you correct the problem.

Information about available reports is automatically downloaded and appears in the library.

Add a Server Connection for Android

When you first open the SAS Mobile BI for Android app, the only connection that is available is to the SAS Demo Server.

To access your organization's library server, you must add a connection to that server. If your organization has more than one server, you can add additional connections in the same way.

Note: Contact your SAS Visual Analytics system administrator for your user ID, password, and server information.

To connect to a server:

1   If you are in My Portfolio, tap **Library**.

2   Tap 🔺 in the library.

3   Tap **Add Connections**.

4   Tap the **Server** field, and enter the address of the new server.

   Note:

   ■   Your server might require that you enter its port number. After the server address, enter a colon and then the port number.

   ■   If the server requires a secure (SSL) connection, type *https://* at the beginning of the server address.

5   Tap the **User ID** field, and enter your user ID.

6   Tap the **Password** field, and enter your password.

7   Tap the **Description** field, and enter a description for the new connection.

8   On the keyboard, tap **Done** to verify the connection.

9   Tap **OK**. The connection is saved, the Connections window closes, and the library
    appears.

Note: If the connection fails, the Extended Connection Properties window displays a
message to help you correct the problem.

45

Viewing Reports in SAS Visual Analytics Viewer

Overview of Viewing Reports in SAS Visual Analytics Viewer

As an alternative to viewing reports on a mobile device, you can use the web viewer. For users with a Report Viewing role, the SAS Visual Analytics Viewer (the viewer) enables them to view report content.

Opening a Report in the Viewer

To open a report in the viewer:

- Click on a report on the SAS Visual Analytics home page.

- Use the object inspector on the home page to view details about the report, and then click **View**. For more information about the object inspector, see "Discovering Details Using the Object Inspector on the Home Page" on page 17.

The layout of the tiles in the treemap is dependent on the size of the display area. This means that the same treemap might appear slightly different in the viewer than it does in SAS Visual Analytics Designer (the designer) or on a mobile device.

Here is an example of a report in the viewer:

Display 45.1   *A Report in the Viewer*
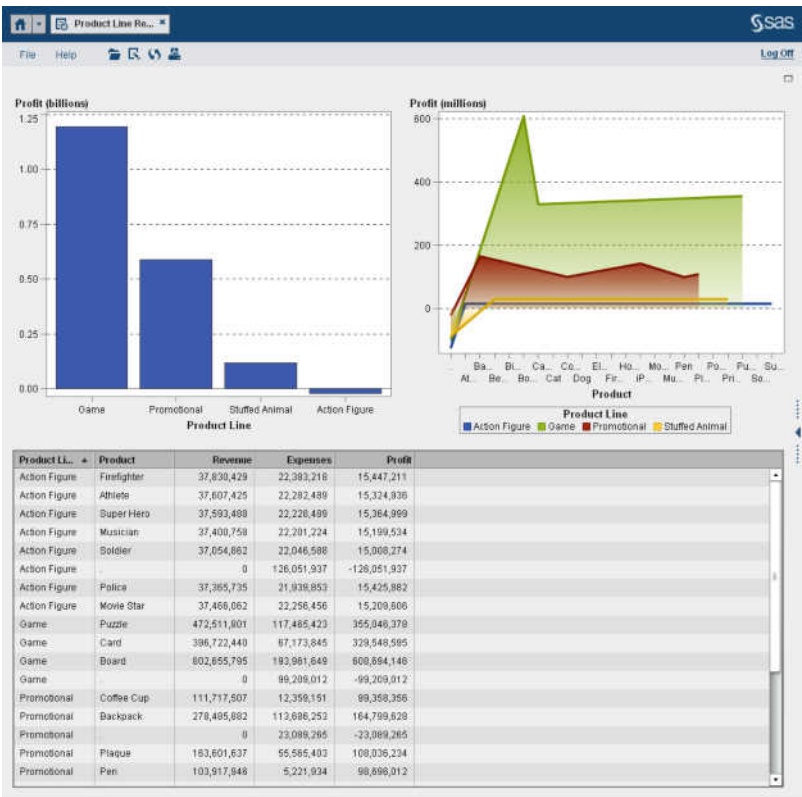
If you have the Create Report capability, then you can select **File ▶ Edit Report** in the current report. The designer is displayed, and then you can edit the report.

Viewing Report Object Information

You can select a report object to display icons that provide more information about that report object. Depending on which type of report object you select, you might see the following icons:

Icon	Description
ⓘ	Click to display the title of the report object.
▭	Click to maximize the report object. This icon is displayed only if there is more than one report object being displayed.
⧉	Click to restore the report object to its original size. This icon is displayed only if you previously maximized the report object.
⇥	Click to display incoming filter information for a report object. This icon is displayed only if the data for the report object has been filtered as a result of selecting data in another report object.
💬	Click to display the prompt dialog box for a stored process. This icon is displayed only for prompted stored processes. You can use the prompt dialog box to change prompt values for a stored process and to re-execute the stored process.

Note: Report authors can disable object selection, so you might not be able to select some report objects.

Adding Comments to a Report in the Viewer

If you have the Add Comments capability, then you can add or view comments. Comments can be added to a report or to an object within a report. The comments that you add are automatically saved with the report.

To add a comment to a report or to an object within a report:

1   Expand the right pane. Click the **Comments** tab. (If you are adding a comment to a report object, then select the report object.)

2   Enter a topic name and a comment.

Here is an example:

Display 45.2   Adding Comments in SAS Visual Analytics Viewer

3   (Optional) Click 📎 to attach a file or an image to your comment. There is no restriction on the file type or the size of the attachment.

4   Click **Post** to add your comment. Your comment is added to the **Comments** tab in the right pane.

To respond to an existing comment:

1   Expand the right pane. Click the **Comments** tab. Select an existing comment, and then enter a reply.

2   (Optional) Click 📎 to attach a file or image to your reply.

3   Click **Post**. Your reply is added to the **Comments** tab in the right pane.

Note: To edit another user's comments or to delete comments, you must have the **Comments:Administrator** predefined role.

To search for a comment:

1 Expand the right pane. Click the **Comments** tab. Enter the word or phrase that you want to search for in the **Search within comments** field on the **Comments** tab.

2 (Optional) To clear the **Search within comments** field, click ⊗.

Interacting with Reports in the Viewer

Filtering, Brushing, and Drilling in Reports

Depending on the interactions that were defined by the report designer, you might be able to interact with your report in the following ways:

filter
> restricts the data that is returned from a query to a data source. Click on data in the source report object to filter data in the target report object (or objects). Clicking on different data applies the filter based on the new data. To clear the selection, click in the white space that surrounds the source report object.

brush
> enables you to show the same data selected simultaneously in two or more tables, graphs, or both. Click on data in the source report object to brush data in the target report object (or objects). To clear the selection, click in the white space that surrounds the source report object.

drill down
> enables you to move from summary information to more detailed data. If the data contains hierarchies, you can double-click on the data to drill down the hierarchy into detailed information. When you drill down the hierarchy, breadcrumbs at the top of the report object enable you to drill back up the hierarchy.

Viewing Links in Reports

Report objects can link to other report sections, whole reports, or they can link to external links. To view a link from a report object, double-click on the report object. If there are multiple links from a report object, then a list appears that enables you to select a link. When you are viewing a link, the following buttons are displayed in the top left corner: .

If a report object that you are linking to shares a data source with the current report object, then the target report object is filtered based on the data value that you double-clicked. If you do not want the target report object to be filtered, you can click , and deselect **Apply report link filters**. If the data source is not shared between the two report objects, then no additional filtering takes place in the target report object.

To go back to the original report object, click .

Subscribing and Unsubscribing to Alerts

You can subscribe to existing alerts for report objects and receive notifications when the alert criteria is met.

To subscribe or unsubscribe to alerts for a report object:

1   Expand the right pane. Click the **Alerts** tab. This tab contains a list of all alert conditions for all report objects within the report.

2   Select or deselect the **Subscribe** check box next to an alert to subscribe or unsubscribe to that alert.

Guest Access

SAS Visual Analytics system administrators can configure support for guest access. Users with guest access can access only the SAS Visual Analytics home page and the viewer. Guest access uses a shared account, so it does not provide individualized features, such as history, favorites, preferences, or alerts. If guest access is configured, it is available at a special URL such as

`http://host/SASVisualAnalyticsViewer/guest.jsp`

Accessing SAS Visual Analytics as a guest is useful if you do not have a metadata identity. This enables you to view reports that are widely available under a generic, shared account. You can also view reports that are available to the public on the Internet.

Note: Some features such as comments, alerts, and the banner are not available with guest access.

Specifying Your Preferences for the Viewer

To specify preferences that are specific to the viewer:

1 Select **File ▶ Preferences** to open the Preferences window.

2 Select **SAS Visual Analytics Viewer**.

3 If the **SAS High Contrast** theme is selected as a global preference, then you can select the **Override report theme when High Contrast Theme is selected** check box to ensure that reports are displayed using the SAS High Contrast theme. This overrides any theme settings that are made in the designer.

4 Click **OK** to apply your changes.

To specify global SAS preferences, see "Specify Global Preferences" on page 23. To specify general preferences for SAS Visual Analytics, see "Specify the SAS Visual Analytics General Preferences" on page 24.

Part 7

Appendix

Appendix 1

Gallery of Report Objects

Tables

List Tables

A *list table* is a two-dimensional representation of data in which the data values are arranged in unlabeled rows and labeled columns. List tables can use any data items from a data source. A list table cannot use a hierarchy or a percentage of total item.

Display A1.1   A List Table

Product Line	Revenue	Expenses	Profit ▼	Profit Trend
Game	1,671,890,035	477,809,929	1,194,080,107	~
Promotional	813,699,290	223,822,374	589,876,916	~
Stuffed Animal	276,990,966	159,548,680	117,442,285	~
Action Figure	262,318,761	281,390,254	-19,071,493	~

Crosstabs

A *crosstab* (also known as a crosstabulation table) shows an aggregate metric for the intersections of two or more categories. Crosstabs often have two or more categories assigned to both the rows and columns, forming a matrix. Crosstabs can be easier to read than list tables because they often use less space, and they always collapse repeating values for outer category data items into one unique value, which is known as grouping. A crosstab can use a hierarchy.

Display A1.2 *A Crosstab*

	Date by Year ▲		2011		
Product Brand ▲	Product Line ▲	Product ▲	Profit	Profit (Percent of Total)	
⊟⊞ **Novelty**	⊞⊞ **Promotional**		13,791,102	0.73%	
⊟⊞ **Toy**	⊟⊞ **Action Figure**	.	-8,891,742	-0.47%	
		Athlete	572,411	0.03%	
		Firefighter	582,670	0.03%	
		Movie Star	593,903	0.03%	
		Musician	563,443	0.03%	
		Police	528,432	0.03%	
		Soldier	536,218	0.03%	
		Super Hero	578,688	0.03%	
	⊞⊞ **Game**		49,058,060	2.61%	
	⊞⊞ **Stuffed Animal**		2,054,767	0.11%	

Graphs, Charts, and Plots

Bar Charts

A *bar chart* consists of vertical or horizontal bars that represent quantitative data. Use bar charts to compare data that is aggregated by the distinct values of a category.

You can apply grouping and create data-driven lattices. You can filter or rank your data based on a specified number of top or bottom values.

Display A1.3 *A Bar Chart*
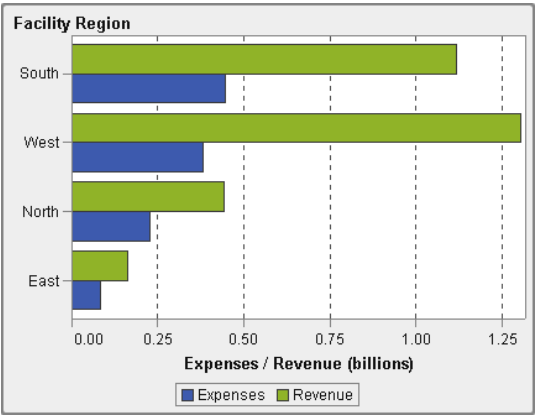
Targeted Bar Charts

A *targeted bar chart* is a variation of the bar chart that has target values. A target value is represented as a triangle with a line at the target value that is determined by the target role.

Display A1.4 *A Targeted Bar Chart*
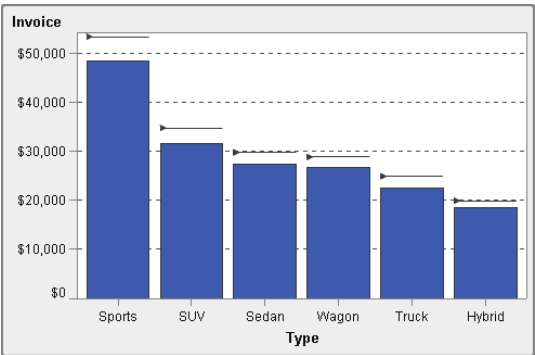
Waterfall Charts

A *waterfall chart* (also known as a progressive bar chart) shows how the initial value of a measure increases or decreases during a series of operations or transactions. The first bar begins at the initial value, and each subsequent bar begins where the previous

bar ends. The length and direction of a bar indicate the magnitude and type (positive or negative, for example) of the operation or transaction. The resulting chart is a stepped bar showing how incremental changes lead to the final value of the measure.

Display A1.5   *A Waterfall Chart*
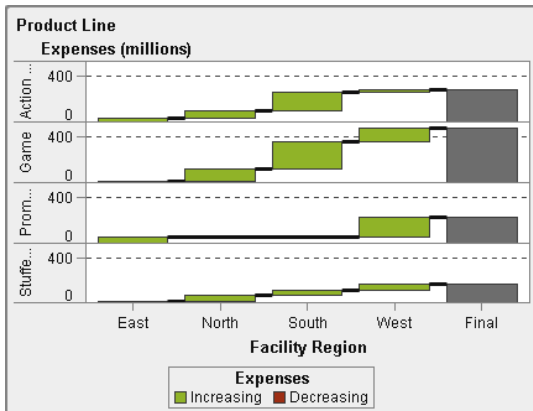
Line Charts

A *line chart* shows the relationship of one or more measures over some interval, such as time or a series of ranges. You can measure a single measure (univariate analysis), or you can show the relationships among multiple measures (multivariate analysis), such as the leading or lagging relationship between advertising and sales over time. The category on the X axis of a line chart is discrete; the category on the X axis of a time series plot is continuous.

You can apply grouping and create lattices.

Display A1.6 *A Line Chart*
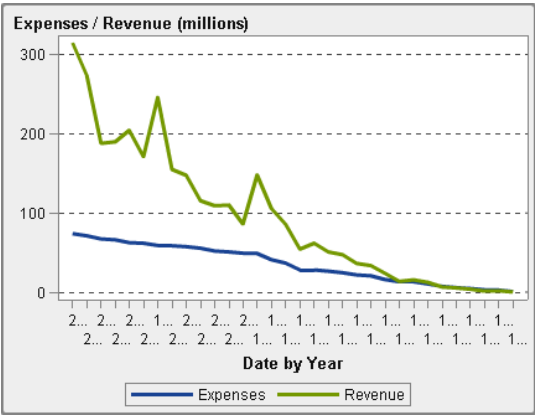
Pie Charts

A *pie chart* displays a part-to-whole relationship in a circle divided into multiple slices for each value of a category data item based on a single measure data item. Each slice represents the relative contribution of each part to the whole. In a pie chart, the legend is sorted by contribution.

In the designer, a pie chart does not show a slice with a missing or zero response.

Effective pie charts limit the number of slices to 5 or 6. In the designer, you can use a rank to reduce the number of slices in a pie chart. For more information, see "Adding a New Rank" on page 460.

Display A1.7 *A Pie Chart*
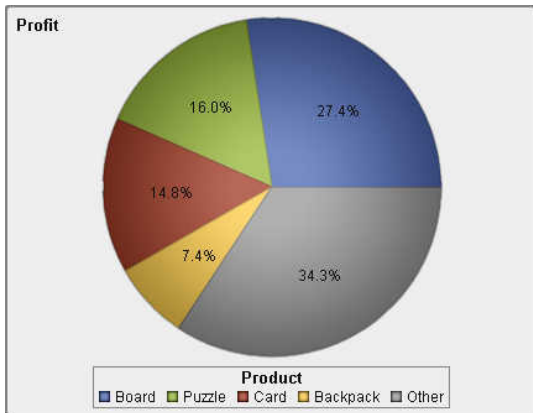
Scatter Plots

A *scatter plot* is a two-dimensional plot that shows the relationship of two measure data items. Each marker (represented by a symbol such as a dot, a square, or a plus sign) represents an observation. The marker's position indicates the value for each observation. Use a scatter plot to examine the relationship between numeric data items. You can apply grouping by assigning a category to the **Color** role.

Scatter plots do not use aggregated data.

Display A1.8 *A Scatter Plot*
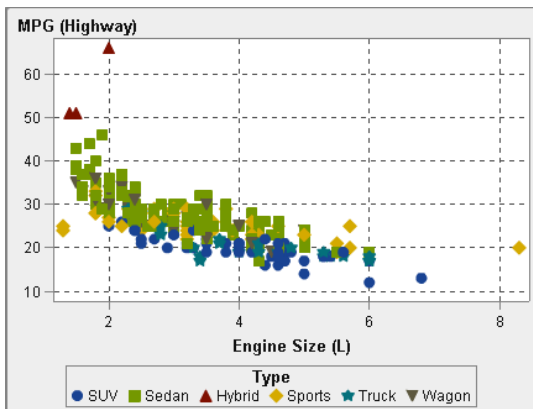
Time Series Plots

A *time series plot* shows an ordered sequence of values that are observed at equally spaced time intervals. A time series plot requires a date, datetime, or time data item that is continuous on the X axis.

Display A1.9   *A Time Series Plot*
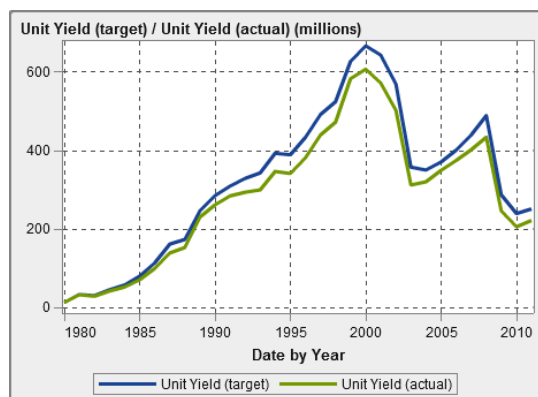
Bubble Plots

A *bubble plot* is a variation of a scatter plot in which the markers are replaced with bubbles. A bubble plot displays the relationships among at least three measures. Two measures are represented by the plot axes, and the third measure is represented by the size of the bubbles. Each bubble represents an observation. A bubble plot is useful for data sets with dozens to hundreds of values. You can add categories to the **Grouping** and **Lattice** roles.

Note: A bubble's size is scaled relative to the minimum and maximum values of the size variable. The minimum and maximum sizes are illustrated in the plot legend. The actual value for each bubble is displayed as a data tip. For example, the legend that is displayed in , the minimum size is 1.3 and the maximum size is 8.3.

Display A1.10   *A Bubble Plot*
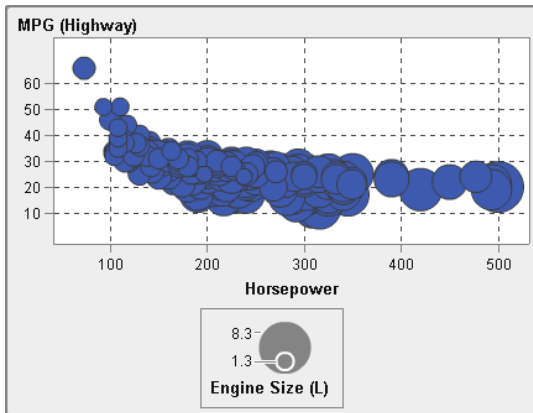
Treemaps

A *treemap* displays your data as a set of rectangles (called tiles). Each tile represents a category or a hierarchy node. The color of each tile represents the value of the first measure. The size of each tile represents the value of the second measure. (There are two data roles for measures in a treemap—**Size** and **Color**.) For example, a sales data treemap might have tile sizes that represent the number of orders, and it might have tile colors that are derived from color gradients that represent sales.

The layout of the tiles in the treemap is dependent on the size of the display area because it uses a space-filling algorithm to lay the tiles down. This means that the same treemap might appear slightly different in the designer than it does in the viewer or on a mobile device because the aspect ratio and size available in those viewers might be different from what the original report designer sees in the designer.

Note: Treemaps allow only one category data item or hierarchy data item.

Display A1.11 *A Treemap*
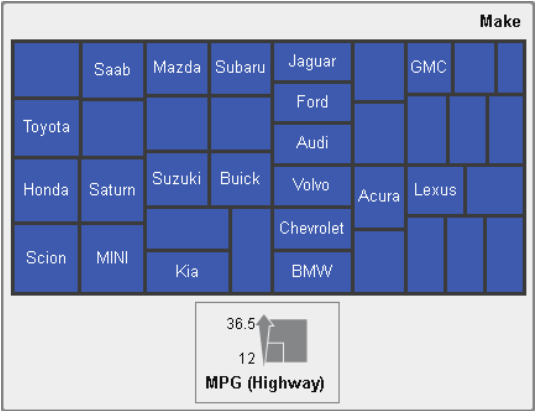
Dual Axis Bar Charts

A *dual axis bar chart* is a variation of the bar chart that has two measures. A measure is on each axis.

Display A1.12 *A Dual Axis Bar Chart*
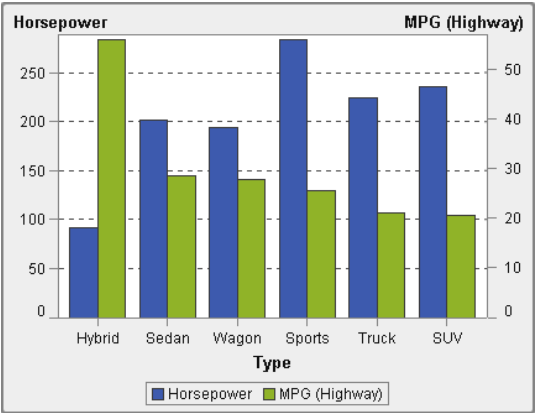
Dual Axis Line Charts

A *dual axis line chart* is a variation of the line chart that has two measures. A measure is displayed on both the left and right side of the Y axis. The relationship between two measures can be examined on two different scales in a dual axis line chart.

Display A1.13   A Dual Axis Line Chart
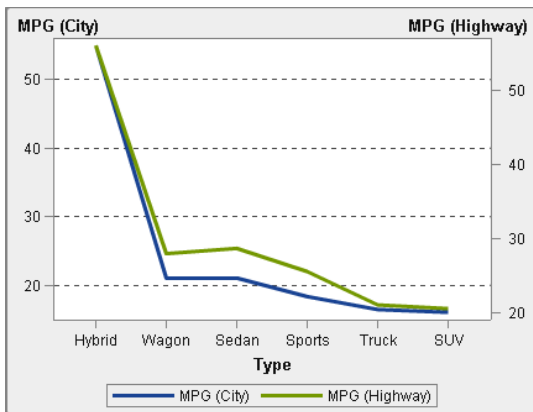
Dual Axis Bar-Line Charts

A *dual axis bar-line chart* is a variation of the bar chart that has two measures. A measure is on each axis, and the bar chart is overlaid by a line chart.

Display A1.14   *A Dual Axis Bar-Line Chart*
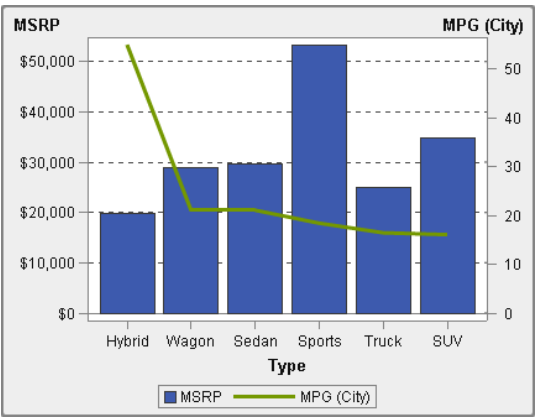
Dual Axis Time Series Plots

A *dual axis time series plot* is a variation of the time series plot that has two measures. A measure is displayed on both the left and right side of the Y axis.

For example, a dual axis time series plot can be useful when you need to display two measures that have the same unit of measurement and different scales, such as quantity ordered and returns, or when you need to display two measures that have different units of measurement, such as sales and quantity ordered.

Display A1.15   *A Dual Axis Time Series Plot*
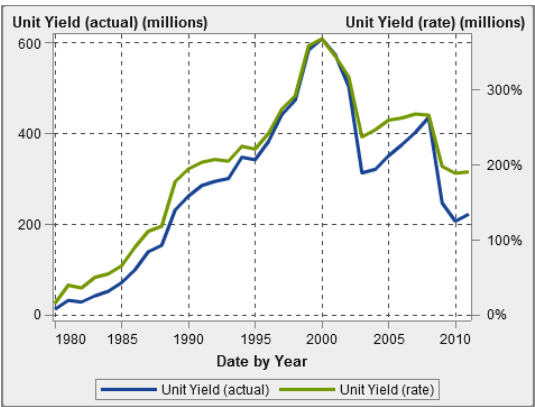
Band Plot

Note: This report object can be created only in the graph builder.

A *band plot* draws a horizontal band with two Y values for each X value. Or, it draws a vertical band with two X values for each Y value. A band plot is typically used to show confidence, error, prediction, or control limits. The points on the upper and lower band boundaries can be joined to create two outlines. The area between the boundaries can be filled.

Note: A band plot does not support display rules in the designer.

The following example shows the low and high stock values for three companies:
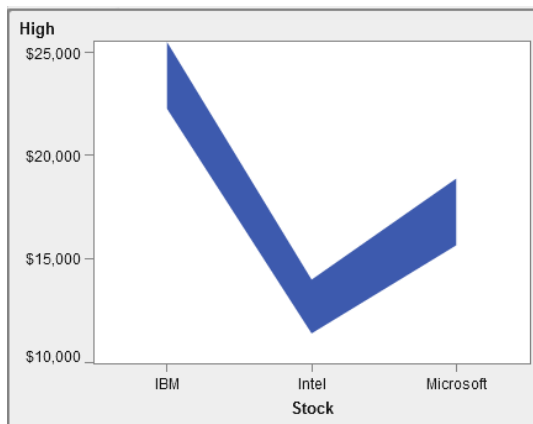
Needle Plot

Note: This report object can be created only in the graph builder.

A *needle plot* is a plot in which data points are connected by a vertical line that connects to a horizontal baseline. The baseline intersects the 0 value or the minimum value on the vertical axis.

The following example shows profits during a particular time period. The example specifies an optional baseline value on the Y axis.
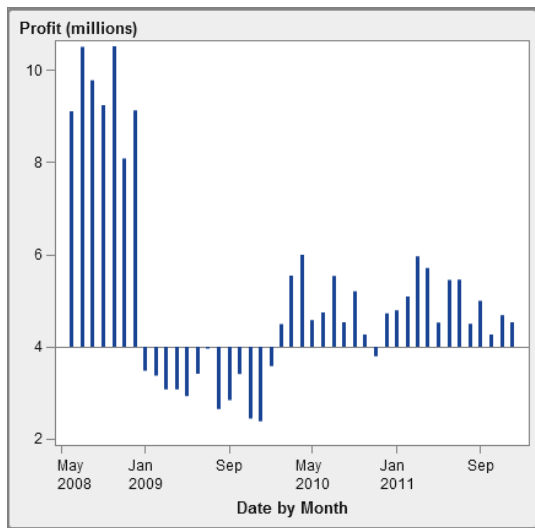
Step Plot

Note: This report object can be created only in the graph builder.

A *step plot* consists of a series of horizontal and vertical line segments (giving the appearance of steps) that connect observations of input data.

Note: A step plot does not support display rules in the designer.

The following example shows the price trend during a particular time interval:
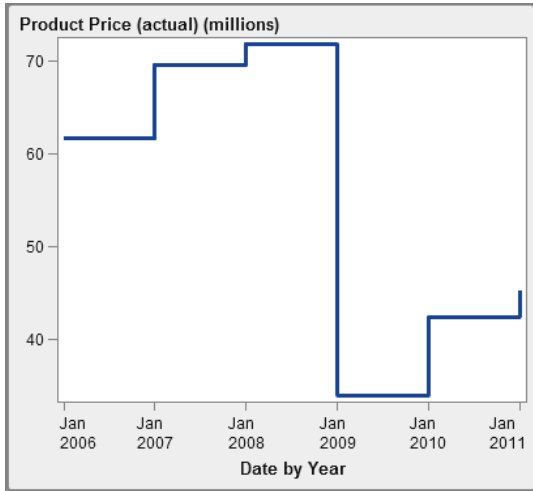
Schedule Chart

Note: This report object can be created only in the graph builder.

A *schedule chart* makes it easy to visualize time lines by representing tasks, start dates, durations, and end dates in cascading horizontal bar charts.

Note: A schedule chart does not support display rules in the designer.
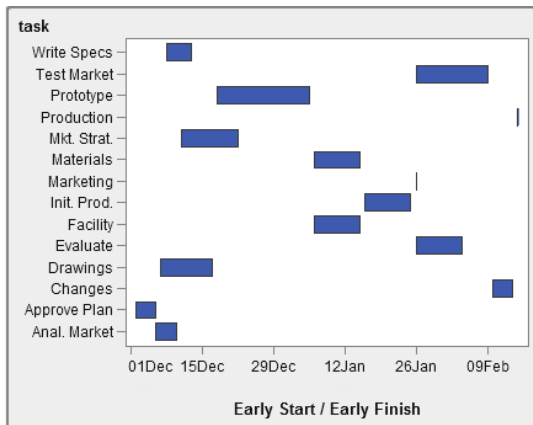
Geo Maps

Geo Bubble Maps

A geo bubble map is a bubble plot that is overlaid on a geographic map. Each bubble is located at a geographic location or at the center of a geographical region. The bubbles are automatically colored based on the location. You provide a measure data item that determines the size of the bubble. A geo bubble map requires a data item that contains geographical information and is assigned to a geography role.

Display A1.16   Geo Bubble Map
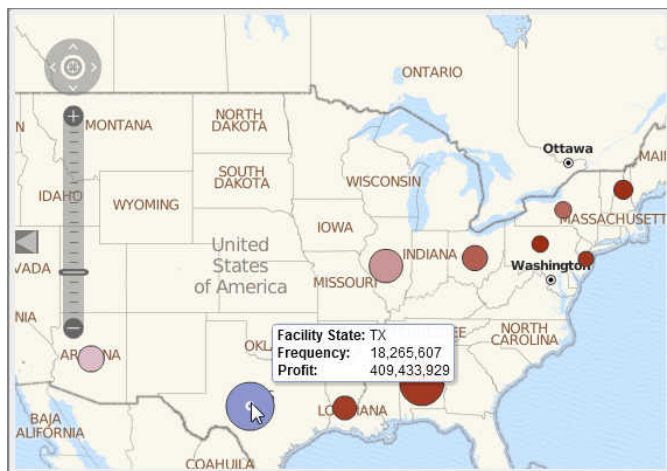
Geo Region Maps

A geo region map (also known as a choropleth map) is a two-dimensional map that uses color combinations to represent different categories or levels of magnitude. You can fill geographical boundaries (for example, a country or a state) on a map with color, based on measure values that are aggregated to the level defined by a geographical boundary.

Display A1.17   *Geo Region Map*
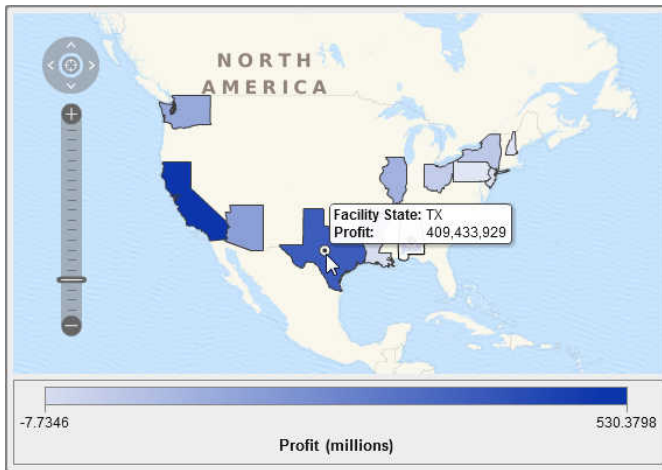
Gauges

Bullet Gauges

A bullet gauge is a dashboard indicator that compares an actual value to a target value and compares them in intervals. The actual value of the primary measure is indicated by an inset horizontal bar.

Note that the scale of a bullet gauge often begins at zero, but it can contain both positive and negative values if both types of values apply to the primary measure, such as profit. The inset horizontal bar should always begin at zero so that comparing multiple bullet graphs is not confusing.

The bullet gauge requires a primary measure and a range display rule, and the target measure is optional. The default orientation is horizontal with an option to display the gauge vertically.

Display A1.18 *A Bullet Gauge*
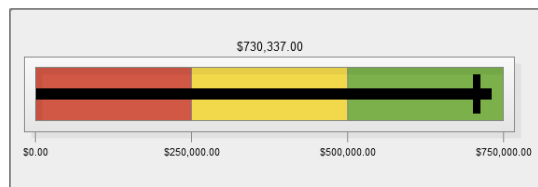
Slider Gauges

A slider gauge is a dashboard indicator that compares an actual value to a target value and compares them in intervals. The actual value of the primary measure is indicated by a downward-facing arrow. The target value is indicated by a small upward-facing arrow.

Like a bullet gauge, a slider gauge is oriented horizontally by default, but you can change the gauge's orientation to vertical. You should use a slider gauge when the numeric scale does not start at zero.

The slider gauge supports dates (which are continuous) and numeric categories.

Display A1.19 *A Slider Gauge*
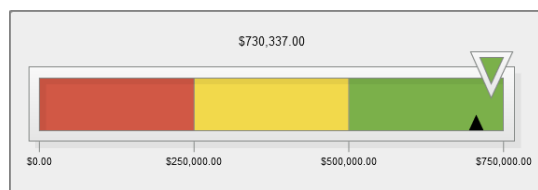
Thermometer Gauges

A thermometer gauge is a dashboard indicator that compares an actual value to a target value and compares them in intervals. The actual value of the primary measure is indicated by a vertical bar. The target value is indicated by a small black line.

The thermometer gauge requires a primary measure value and a range-based display rule. A target measure value is optional. The entire vertical bar is colored conditionally based on one color from the display rule.

The base of a thermometer bar should always start at zero. You can set this by defining your first range display rule to begin at zero. The designer always shows the base of the bar at the bottom of the thermometer.

Display A1.20 *A Thermometer Gauge*
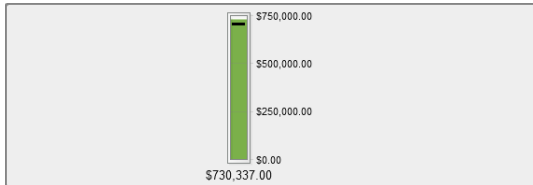
Dial Gauges

A dial gauge is an arc-shaped dashboard indicator that compares an actual value to a target value and compares them in intervals. The actual value of the primary measure is indicated by an arrow that points outward from the inner circle. The target value is indicated by an arrow that points inward from the outer arc. The color of the center circle is the color associated with the primary measure value's range interval.

The dial gauge requires a primary measure value and a range-based display rule. For more information, see "Adding Display Rules to a Gauge" on page 413.

A target measure value is optional.

Display A1.21 *A Dial Gauge*
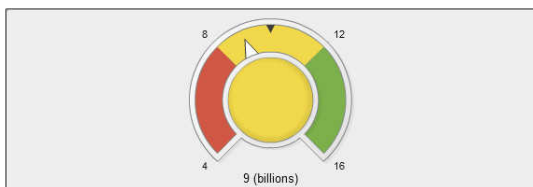
Speedometer Gauges

A speedometer gauge is a circular dashboard indicator that compares an actual value to a target value and compares them in intervals. The actual value of the primary measure

is indicated by the larger pointer. The target value is indicated by a small triangle along the quantitative scale, either pointing inward or outward, depending on the **KPI skin** option for the gauge.

A speedometer gauge requires a primary measure value and a range-based display rule. A target measure value is optional.

Display A1.22 *A Speedometer Gauge*
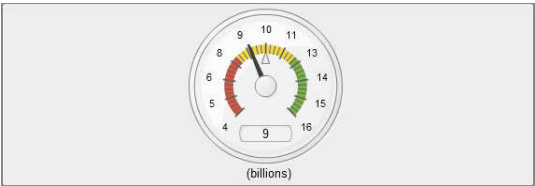
Controls

A control is a report object that filters or narrows the scope of the data that user is viewing. A control enables a report designer to select a category that he or she wants the report viewer to see.
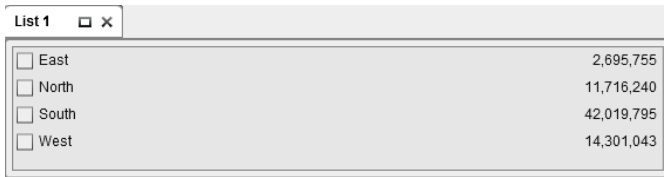
The following controls are available in the designer:

■ drop-down lists

Display A1.23 *A Drop-down List Control*
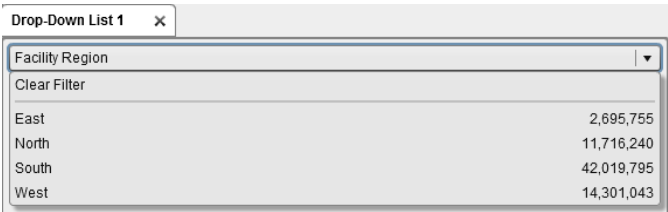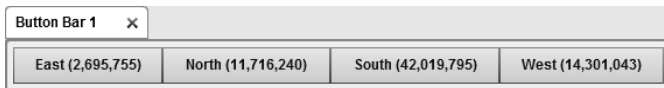
■ lists

Display A1.24   *A List Control*
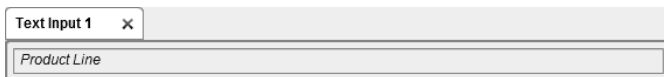
List 1 □ ✕	
☐ East	2,695,755
☐ North	11,716,240
☐ South	42,019,795
☐ West	14,301,043

■   button bars

Display A1.25   *A Button Bar Control*
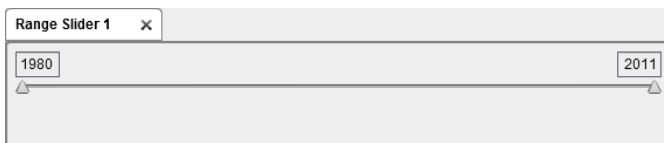
Button Bar 1 ✕			
East (2,695,755)	North (11,716,240)	South (42,019,795)	West (14,301,043)

■   text input fields

Display A1.26   *A Text Input Control*

Text Input 1   ✕
Product Line

■   range sliders

Display A1.27   *A Range Slider Control*

Range Slider 1 ✕	
1980	2011

Appendix 2

Editing a Data Expression in Text Mode

Overview of Text Mode

In SAS Visual Analytics, you create and edit filters, calculated items, and aggregated items by using an expression editor. The **Text** tab of the expression editor enables you to edit the expression as text.

You can add operators and data items to your expression by dragging and dropping them onto the expression or by entering the names of the operators or data items.

As you enter text, a list of suggestions appears. For example, if you enter `date`, then a drop-down list appears and enables you to select any of the operators and data items whose names begin with "date."

When you make changes to your expression, it is automatically evaluated to determine whether it is valid. If the expression is not valid, then an error appears on the **Messages** tab, and the **OK** button is disabled. There might be a brief delay as your expression is evaluated.

Specifying Operator Parameters

When you add an operator to the expression, any parameters that are required by the operator are represented between braces { }. For example, if you add the **x – y** operator, then your expression appears as `{Number} – {Number}`.

Each parameter value that you enter should replace the entire string between the braces, including the brace characters. For example, you might replace {Number} with 12 or with a data item such as Expenses.

You can automatically select the next operator in the expression by pressing Ctrl + Shift + spacebar.

Specifying Data Items

Data item names can be entered as plain text and are not case sensitive. You can enter data item names formally by using the format '*date-item-name*'n. If you switch to the **Visual** tab, then all of your data item names are converted to the formal format.

Note: If a data item name contains quotation marks then you must use the \ character to escape the quotes.

Note: If a data item name contains spaces, then you must use the formal format.

Note: If a data item has the same name as an operator, then you must use the formal format. For example, if you have a category named Year, then enter the name as **'Year'n** to avoid conflict with the **Year** operator.

Using Formatted and Unformatted Values

By default, category values and discrete numeric and date values are evaluated as formatted values. Continuous numeric values are evaluated as unformatted values.

To override this default behavior, you can add [raw] (to use unformatted values) or [formatted] (to use formatted values) to the right of the data item.

For example,

```
'Expenses'n[formatted]
```

specifies that the Expenses measure is evaluated as a formatted value.

Specifying String Values

To enter a string value, you can enter the string between single quotes or double quotes. If your value contains a quotation mark, use the \ character to escape the quote. If you use double quotes to enclose the string, then you do not need to escape the single quotes. If you use single quotes to enclose the string, then you do not need to escape the double quotes. For example, `"O'Reilly"`, `'O\'Reilly'`, and `'"Hello"'` are all valid.

To enter a string that contains a newline character, use `\r`, `\n`, or both to specify the newline character.

Specifying Date, Time, and Datetime Values

For date, datetime, and time values, specify a formatted value in quotes, followed by the letter "d" for a date value, the letters "dt" for a datetime value, or the letter "t" for a time value.

For time values, a value with a leading underscore is a reference to a specific time of the day.

A time value without a leading underscore specifies a period of elapsed time.

Here are some examples of date, datetime, and time values:

Table A2.1   Example Date, Datetime, and Time Values

Date	'23JUN2013'd
	'JUN2013'd
	'2013'd
	'q32013'd
Datetime	'23JUN2013_5:23:55'dt
	'23JUN2013_5:23'dt
	'23JUN2013'dt

Time	'_12:23:55't
	'_12:23't
Elapsed Time	'44:23:55't
	'00:01:34't

Specifying Aggregated Values

For aggregated values, specify the format, `aggregation-type [context]` `(value)`, where *context* specifies one of the following aggregation contexts:

ByGroup

calculates the aggregation for each subset of the data item that is used in a visualization. For example, in a bar chart, an aggregated measure with the **ByGroup** context calculates a separate aggregated value for each bar in the chart.

ForAll

calculates the aggregation for the entire data item (after filtering). For example, in a bar chart, an aggregated measure with the **ForAll** context uses the same aggregated value (calculated for the entire data item) for each bar in the chart.

See "Aggregated Operators" on page 579 for a list of the aggregation types that are available.

For example,

```
sum [bygroup] ('cost'n)
```

aggregates the sum of the measure COST for each BY-group value.

Specifying a Missing Value

Use a period character (.) to specify a missing numeric or date value. Use empty quotes ('') to specify a missing string value.

Appendix 3

Operators for Data Expressions

Overview of Operators for Data Expressions

In the explorer and the designer, you can calculate data items and create filters by using expressions that contain operators.

Numeric (Simple) Operators

-x
> returns a value with the opposite sign of the input value.
>
> For example, `--1` returns 1 and `-1` returns -1.

x - y
> subtracts the second value from the first value.
>
> For example, `2 - 1` returns 1.

x * y
> multiplies the first and second values together.
>
> For example, `2 * 3` returns 6.

x / y
> divides the first value by the second value.
>
> For example, `6 / 2` returns 3.

x + y
> adds the first and second values together.
>
> For example, `1 + 2` returns 3.

Comparison Operators

BetweenExclusive
returns true if the first value is within the range defined by the second and third values (excluding the bounding values).

For example, `X BetweenExclusive(50, 100)` returns true if X is greater than 50 and less than 100.

BetweenInclusive
returns true if the first value is within the range defined by the second and third values (including the bounding values).

For example, `X BetweenInclusive(50, 100)` returns true if X is greater than or equal to 50 and less than or equal to 100.

In
returns true if the first value is in the list specified by the second parameter. Select your list by choosing the values from the drop-down list or the selector window.

For example, `X In ('A', 'B', 'C')` returns true when the value of X is either A, B, or C.

Note: This operator cannot be used to compare measures.

Missing
returns true if the value is a missing value.

For example, `X Missing` returns true if the value of X is missing.

NotIn
returns true if the first value is not in the list specified by the second parameter. Select your list by choosing the values from the drop-down list or the selector window.

For example, `X NotIn ('A', 'B', 'C')` returns true when the value of X is not A, B, or C.

Note: This operator cannot be used to compare measures.

NotBetweenExclusive

returns true if the first value is outside the range defined by the second and third values (excluding the bounding values).

For example, X `NotBetweenExclusive(50, 100)` returns true if X is less than 50 or greater than 100.

NotBetweenInclusive

returns true if the first value is outside the range defined by the second and third values (including the bounding values).

For example, X `NotBetweenInclusive(50, 100)` returns true if X is less than or equal to 50 or less than or equal to 100.

NotMissing

returns true if the value is not a missing value.

For example, X `NotMissing` returns true if the value of X is not missing.

x < y

returns true if the first value is less than the second value.

x <= y

returns true if the first value is less than or equal to the second value.

x <> y

returns true if the first value is not equal to the second value.

x = y

returns true if the first value is equal to the second value.

x > y

returns true if the first value is greater than the second value.

x >= y

returns true if the first value is greater than or equal to the second value.

Boolean Operators

AND
    joins two conditions and returns true if both conditions are true.

    For example,

    `(1 = 1) AND (2 = 2)`

    returns true, and

    `(1 = 1) AND (2 = 1)`

    returns false.

IF... ELSE
    returns different values, depending on whether the condition is true. The first
    parameter specifies the condition. The second parameter specifies the value to
    return if the condition is true. The third parameter specifies the value to return if the
    condition is false.

    For example,

    `if (X > Y) return X else Y`

    returns the value of X if X is greater than Y, but returns the value of Y otherwise.

NOT
    returns true if the condition is false.

    For example, `not (1 = 2)` returns true.

OR
    joins two conditions and returns true if either condition is true.

    For example,

    `(1 = 1) OR (2 = 2)`

    returns true, and

    `(1 = 1) OR (2 = 1)`

returns true.

Numeric (Advanced) Operators

Abs
   returns the absolute value of the input value.

   For example, `Abs(-3)` returns 3.

Ceil
   rounds the input value up to the nearest integer.

   For example, `Ceil(4.2)` returns 5 and `Ceil(-4.8)` returns -4.

Exp
   raises the constant *e* to the power specified by the input value.

   For example, `Exp(5)` returns *e* to the 5th power (148.41).

Floor
   rounds the input value down to the nearest integer.

   For example, `Floor(4.8)` returns 4 and `Floor(-4.2)` returns -5.

Ln
   returns the natural logarithm (base *e*) of the input value.

   For example, `Ln(10)` returns the *e*th root of 10 (2.30...).

Log
   returns the logarithm of the first value, where the second value specifies the base.

   For example, `64 Log 8` returns the base 8 logarithm of 64 (2).

Mod
   returns the remainder after dividing the first value by the second value.

   For example, `5 Mod 2` returns 1.

Power
> raises the first value to the power of the second value.
>
> For example, `5 Power 2` returns 5 to the 2nd power (25).

Root
> returns the *n*th root of the first value, where the second value specifies *n* (the base of the root).
>
> For example, `27 Root 3` returns the 3rd (cube) root of 27 (3).

Round
> rounds the first value to the number of decimal places that is specified by the second value. Select the second value from the drop-down list.
>
> For example, `7.354 Round 2` returns 7.35.
>
> **Note:** If you select 0 decimal places, then the values are rounded to the nearest integer.

TreatAs
> allows a numeric, date, or datetime value to be used as a different data type within other operators. Select one of the following:
>
> _Date_
>> allows the value to be used as a date.
>
> _Datetime_
>> allows the value to be used as a datetime value.
>
> _Number_
>> allows the value to be used as a number.
>
> _Time_
>> allows the value to be used as a time value.
>
> The value is treated as a raw value instead of being converted. Date values are the number of days since 01JAN1960. Datetime values are the number of seconds since 01JAN1960. Time values are the number of seconds since midnight.
>
> For example, `TreatAs(_Date_, 19600)` returns 30AUG2013 as a date value.

Note: The TreatAs operator is useful for calculating elapsed time between two datetime values. For example, `(TreatAs(_Number_, '23OCT2013'd) — TreatAs(_Number_, '15JAN2013'd))` calculates the number of days between 15JAN and 23OCT, which is 281.

Trunc

truncates the input value to an integer.

For example, `Trunc(8.9)` returns 8 and `Trunc(-8.9)` returns -8.

Date and Time Operators

Note: Date and time operators are not supported for aggregated items.

DateFromMDY

creates a date value from separate month, day, and year values. The first value specifies the month as a number from 1–12. The second value specifies the day as a number from 1–31. The third value specifies the year as a four-digit number.

For example, `DateFromMDY(1, 15, 2013)` returns 15JAN2013.

DateFromYQ

creates a date value from separate year and quarter values. The first value specifies the year as a four-digit number. The second value specifies the quarter as a number from 1–4.

For example, `DateFromYQ(2013, 1)` returns 01JAN2013.

Note: The date is generated using the first day of each quarter.

DatePart

converts a datetime value to a date value.

For example, `DatePart('15JAN2013_17:15'dt)` returns 15JAN2013.

DateTimeFromDateHMS

creates a datetime value from a date value and separate hour, minute, and second values. The first value specifies the date. The second value specifies the hour as a

number from 0–23. The third value specifies the minute as a number from 0–59. The fourth value specifies the second as a number from 0–59.

For example, `DateTimeFromDateHMS('15JAN2013'd, 17, 15, 23)` returns January 15, 2013 05:15:23 PM

DateTimeFromTimeMDY

creates a datetime value from a time value and separate month, day, and year values. The first value specifies the time. The second value specifies the month as a number from 1–12. The third value specifies the day as a number from 1–31. The fourth value specifies the year as a four-digit number.

For example, `DateTimeFromTimeMDY('_17:15:23'dt, 1, 15, 2013)` returns January 15, 2013 05:15:23 PM.

DayOfMonth

returns the day of the month from a date value as a number from 1–31.

For example, `DayOfMonth('15JAN2013'd)` returns 15.

DayOfWeek

returns the day of the week from a date value as a number from 1–7 (1 is Sunday).

For example, `DayOfWeek('15JAN2013'd)` returns 3 (Tuesday).

DayOfYear

returns the day of the year from a date value as a number from 1–366.

For example, `DayOfYear('15FEB2013'd)` returns 46.

Hour

returns the hour from a time or datetime value as a number from 0–23.

For example, `Hour('17:15:23't)` returns 17.

Minute

returns the minute from a time or datetime value as a number from 0–59.

For example, `Minute('17:15:23't)` returns 15.

Month

returns the month from a date value as a number from 1–12.

For example, `Month('15JAN2013'd)` returns 1.

Now
  creates a datetime value from the current date and time.

  For example, `Now()` returns the current date and time.

Quarter
  returns the quarter from a date value as a number from 1–4.

  For example, `Quarter('15AUG2013'd)` returns 3.

Second
  returns the second from a time or datetime value as a number from 0–59.

  For example, `Second('17:15:23't)` returns 23.

TimeFromHMS
  creates a time value from separate hour, minute, and second values. The first value specifies the hour as a number from 0–23. The second value specifies the minute as a number from 0–59. The third value specifies the second as a number from 0–59.

  For example, `TimeFromHMS(17, 15, 23)` returns 05:15:23 PM.

TimePart
  converts a datetime value to a time value.

  For example, `TimePart('15JAN2013_17:15:23'dt)` returns 05:15:23 PM.

WeekNumber
  returns the week of the year as a number from 0–53, where week 1 begins on the first Sunday of the year. Dates before the first Sunday of the year return 0.

  For example, `WeekNumber('04AUG2013'd)` returns 31.

Year
  returns the year from a date value as a four-digit number.

  For example, `Year('15JAN2013'd)` returns 2013.

Aggregated Operators

Avg
　calculates the average value for a measure.

Count
　calculates the total number of nonmissing values for a measure.

Distinct
　calculates the number of distinct values for a category.

Max
　calculates the greatest value for a measure.

Min
　calculates the least value for a measure.

Sum
　calculates the sum of the values for a measure.

Periodic Operators

About Periodic Operators

Periodic operators aggregate values over a period of time.

If you assign a periodic aggregated item to a visualization or report object that contains dates, the aggregated item displays the aggregated values for each time period in the visualization or report object.

In a visualization or report object that does not contain dates, the aggregated item displays values that use today's date as a reference. If the date data item for the operator does not contain data for the interval that contains today's date, then the operator returns missing values.

Periodic operators are evaluated using time intervals. Intervals specify whether the aggregation is applied on a monthly basis, a quarterly basis, and so on. You can specify a specific interval, or you can specify that the interval is inferred. For an inferred interval, the aggregation is evaluated based on its context in the visualization or report object. For example, if your visualization contains a bar chart of sales by month, then the inferred interval is monthly.

Note: Periodic operators return a missing value in the following scenarios:

- Data does not exist for the specified time period.

- The date data item for the period calculation does not match the date data item in the visualization or report object. You must use the same date data item or a duplicate data item that is based on the same data item.

- The interval for the operator is smaller than the interval of the data format in the visualization or report object (for example, if your interval is by month, but the date format is Year).

- For operators that use inner and outer intervals, the inner interval is larger than the outer interval.

- The inferred interval is by week of the year or by an interval smaller than a day.

- The inferred interval is by day for any operator that has an offset other than 0.

  For the explorer, the ParallelPeriod and RelativePeriod operators always return a missing value when the inferred interval is by day.

CumulativePeriod

The CumulativePeriod operator returns aggregated values for a period of time, including the previous periods within a larger period of time (for example, the year-to-date total).

Specify the following parameters:

Figure A3.1   Parameters for the CumulativePeriod Operator

1   The aggregation that is applied to the measure.

2   The measure to aggregate over time.

3   The date data item for the period calculation. Only data items whose formats specify year are available.

4   The inner interval for which the values are aggregated. For example, specify _ByMonth_ as the inner interval and _ByYear_ as the outer interval to aggregate the year-to-date values for each month.

Select one of the following:

Inferred
    specifies that the interval is determined automatically from the visualization or report object that displays the aggregated item.

ByMonth
    specifies a monthly interval.

ByQuarter
    specifies a quarterly interval.

ByYear
> specifies a yearly interval.

5 The outer interval that provides the context for the cumulative aggregation. For example, specify _ByMonth_ as the inner interval and _ByYear_ as the outer interval to aggregate the year-to-date values for each month.

Select one of the following:

Inferred
> specifies that the interval is determined automatically from the visualization or report object that displays the aggregated item.

ByMonth
> specifies a monthly interval.

ByQuarter
> specifies a quarterly interval.

ByYear
> specifies a yearly interval.

6 The number of outer intervals to offset from the current period. 0 specifies that the period from the current outer interval is used. A negative value indicates a previous interval.

For example, if your inner interval is by month and your outer interval is by year, then -1 specifies the year-to-date monthly values for the previous year.

7 The scope for the period. The scope specifies how much of each period is aggregated.

Select one of the following:

Full
> aggregates the values for the entire period.

ToDate
> aggregates only the values up to a specific day of the outer interval.

ToToday

> aggregates only the values up to the equivalent of today's position in the current interval. For example, if today is the 40th day of the quarter, and the outer interval is by quarter, then only the values up to the 40th day of each quarter are used.
>
> The value for today is evaluated dynamically whenever the aggregated item is viewed in a visualization or report object.

8 If you select _ToDate_ as the scope, then select the date that is used to subset each period.

   For example, if you select 09NOV2013, and the outer interval is by quarter, then only the values up to the 40th day of each quarter are used in the aggregation.

For example,

```
CumulativePeriod(_Sum_, 'Expenses'n, 'Date'n, _ByMonth_, _ByYear_, 0, _Full_)
```

aggregates the sum of year-to-date monthly values for the Expenses measure using date values from the Date data item.

ParallelPeriod

The ParallelPeriod operator returns aggregated values for a period of time that is parallel to the current period (for example, the same month of the previous year).

Specify the following parameters:

Figure A3.2   *Parameters for the ParallelPeriod Operator*

1   The aggregation that is applied to the measure.

2   The measure to aggregate over time.

3   The date data item for the period calculation. Only data items whose formats specify year are available.

4   The inner interval for which the values are aggregated. For example, specify _ByMonth_ as the inner interval to aggregate the values for each month.

Select one of the following:

Inferred
   specifies that the interval is determined automatically from the visualization or report object that displays the aggregated item.

ByMonth
   specifies a monthly interval.

ByQuarter
   specifies a quarterly interval.

ByYear
   specifies a yearly interval.

5   The outer interval that provides the context for the parallel period aggregation. For example, specify _ByMonth_ as the inner interval and _ByYear_ as the outer interval to aggregate the monthly values for a different year.

Select one of the following:

Inferred
> specifies that the interval is determined automatically from the visualization or report object that displays the aggregated item.

ByMonth
> specifies a monthly interval.

ByQuarter
> specifies a quarterly interval.

ByYear
> specifies a yearly interval.

6   The number of outer intervals to offset from the current period. 0 specifies that the period from the current outer interval is used. A negative value indicates a previous interval.

For example, if your inner interval is by month and your outer interval is by year, then -1 specifies the monthly values for the previous year.

7   The scope for the period. The scope specifies how much of each period is aggregated.

Select one of the following:

Full
> aggregates the values for the entire period.

ToDate
> aggregates only the values up to a specific day of the outer interval.

ToToday
> aggregates only the values up to the equivalent of today's position in the current interval. For example, if today is the 40th day of the quarter, and the outer

interval is by quarter, then only the values up to the 40th day of each quarter are used.

The value for today is evaluated dynamically whenever the aggregated item is viewed in a visualization or report object.

8 If you select _ToDate_ as the scope, then select the date that is used to subset each period.

For example, if you select 09NOV2013 and the outer interval is by quarter, then only the values up to the 40th day of each quarter are used in the aggregation.

For example,

```
ParallelPeriod(_Sum_, 'Expenses'n, 'Date'n, _ByMonth_, _ByYear_, -1, _Full_)
```

aggregates the sum of monthly values for the Expenses measure for the previous year using date values from the Date data item.

Period

The Period operator returns aggregated values for a period of time (for example, the current month).

Specify the following parameters:

Figure A3.3   Parameters for the Period Operator

1 The aggregation that is applied to the measure.

2 The measure to aggregate over time.

3 The date data item for the period calculation. Only data items whose formats specify year are available.

4   The interval for which the values are aggregated. For example, specify _ByMonth_ as the interval to aggregate the values for each month.

Select one of the following:

Inferred
   specifies that the interval is determined automatically from the visualization or report object that displays the aggregated item.

ByMonth
   specifies a monthly interval.

ByQuarter
   specifies a quarterly interval.

ByYear
   specifies a yearly interval.

5   The number of outer intervals to offset from the current period. 0 specifies that the period from the current outer interval is used. A negative value indicates a previous interval.

For example, if your inner interval is by month and your outer interval is by year, then -1 specifies the year-to-date monthly values for the previous year.

For example,

```
Period(_Sum_, 'Expenses'n, 'Date'n, _ByMonth_)
```

aggregates the sum of monthly values for the Expenses measure using date values from the Date data item.

PeriodWithDate

The PeriodWithDate operator returns aggregated values for a specific, constant period of time (for example, the month that includes 15OCT2013).

Specify the following parameters:

Figure A3.4   Parameters for the PeriodWithDate Operator
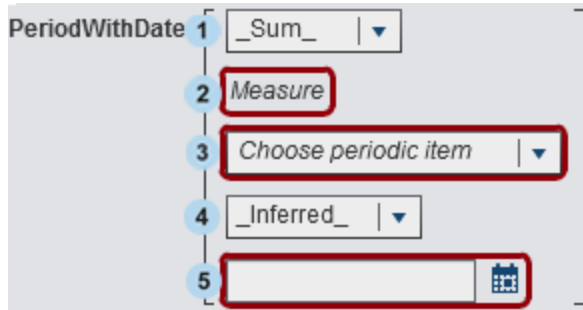
1   The aggregation that is applied to the measure.

2   The measure to aggregate over time.

3   The date data item for the period calculation. Only data items whose formats specify year are available.

4   The interval for which the values are aggregated. For example, specify _ByMonth_ as the interval to aggregate the values for each month.

Select one of the following:

Inferred
   specifies that the interval is determined automatically from the visualization or report object that displays the aggregated item.

ByMonth
   specifies a monthly interval.

ByQuarter
   specifies a quarterly interval.

ByYear
   specifies a yearly interval.

5   The reference date for the period aggregation.

For example,

```
PeriodWithDate(_Sum_, 'Expenses'n, 'Date'n, _ByMonth_, '15OCT2013'd)
```

aggregates the sum of monthly values for the Expenses measure using date values from the Date data item.

RelativePeriod

The RelativePeriod operator returns aggregated values for a period of time that is relative to the current period (for example, the previous month of the same year).

Specify the following parameters:

Figure A3.5   *Parameters for the RelativePeriod Operator*

1   The aggregation that is applied to the measure.

2   The measure to aggregate over time.

3   The date data item for the period calculation. Only data items whose formats specify year are available.

4   The interval for which the values are aggregated. For example, specify _ByMonth_ as the interval to aggregate the year-to-date values for each month.

Select one of the following:

Inferred
   specifies that the interval is determined automatically from the visualization or report object that displays the aggregated item.

ByMonth
  specifies a monthly interval.

ByQuarter
  specifies a quarterly interval.

ByYear
  specifies a yearly interval.

5 The number of intervals to offset from the current period. 0 specifies that the period from the current interval is used. A negative value indicates a previous interval.

For example, if your interval is by month, then -1 specifies the monthly values for the previous month.

6 The scope for the period. The scope specifies how much of each period is aggregated.

Select one of the following:

Full
  aggregates the values for the entire period.

ToDate
  aggregates only the values up to a specific day of the interval.

ToToday
  aggregates only the values up to the equivalent of today's position in the current interval. For example, if today is the 40th day of the quarter, and the outer interval is by quarter, then only the values up to the 40th day of each quarter are used.

  The value for today is evaluated dynamically whenever the aggregated item is viewed in a visualization or report object.

7 If you select _ToDate_ as the scope, then select the date that is used to subset each period.

For example, if you select 09NOV2013 and the outer interval is by quarter, then only the values up to the 40th day of each quarter are used in the aggregation.

For example,

```
RelativePeriod(_Sum_, 'Expenses'n, 'Date'n, _ByMonth_, -1, _Full_)
```

aggregates the sum of monthly values for the Expenses measure for the previous month using date values from the Date data item.

Text (Simple) Operators

Note: All text operators are case sensitive.

Note: Text operators are not supported for aggregated items.

Concatenate
appends the second input string to the first input string.

For example, `Concatenate('A', 'B')` returns **AB**.

Contains
specifies that a matching value must contain the specified string.

For example, `'Catcher' Contains 'Cat'` returns true.

EndsWith
specifies that a matching value must contain the specified string at the end of the value.

For example, `'Catcher' EndsWith 'her'` returns true.

Format
applies a format to the input value. Click the format field to select the format that you want to apply. The output from the Format operator is a string.

For example, `Format(1015.35, 'DOLLAR6.2')` returns **$1,015.35** as a string value.

Note: Standard date formats in SAS Visual Analytics display date and datetime values in the locale of your browser. You can display date and datetime values in the locale of the data source by using national language formats. The names of national language formats begin with "NL." For example, the NLDATE format displays date values by using the locale of the data source.

LowerCase
> changes all of the characters in a text string to lowercase.
>
> For example, `LowerCase('SAS INSTITUTE')` returns **sas institute**.

NotContains
> specifies that a matching value must not contain the specified string.
>
> For example, `'Catcher' NotContains 'Dog'` returns true.

Parse
> interprets a numeric or datetime value from the input string. Click the format field to select the format that is used to interpret the string. The output from the Parse operator is either a number or a datetime value, depending on the format that you select.
>
> For example, `Parse('15JAN2013', 'DATE9.')` returns 15JAN2013 as a date value.

StartsWith
> specifies that a matching value must contain the specified string at the start of the value.
>
> For example, `'Catcher' StartsWith 'Cat'` returns true.

UpCase
> changes all of the characters in the text string to uppercase.
>
> For example, `UpCase('sas institute')` returns **SAS INSTITUTE**.

Text (Advanced) Operators

Note: All text operators are case sensitive.

Note: Text operators are not supported for aggregated items.

FindChar
> finds the position of a character or a set of characters within a text string. The position of the first match is returned as a numeric value. If no matches are found,

then 0 is returned. The first input string specifies the value to search within. The second input string specifies the list of characters to search for.

For example, `FindChar('mystring', 'sz')` returns 3.

FindString

finds the position of a string within another string. The position of the first match is returned as a numeric value. If no matches are found, then 0 is returned. The first input string specifies the value to search within. The second input string specifies the string to search for.

For example, `FindString('mystring', 'st')` returns 3.

GetLength

returns the length of an input string as a numeric value.

For example, `GetLength('mystring')` returns 8.

GetWord

returns a word from an input string where the words are separated by spaces, periods, or other special characters. The first parameter specifies the input string. The second parameter specifies the number of the word to return where 1 is the first word.

For example, `GetWord('my test string', 2)` returns **test**.

RemoveBlanks

removes space characters from the input string. The first parameter specifies the input string. The second parameter specifies which space characters to remove. Select one of the following:

All

removes all spaces from the string.

Leading

removes spaces at the beginning of the string.

LeadingAndTrailing

removes spaces at the beginning and end of the string.

Trailing
> removes spaces at the end of the string.

For example, `RemoveBlanks('my test string', '_ALL_')` returns **myteststring**.

RemoveChars
> removes all instances of a set of characters from the input string. The first parameter specifies the input string. The second parameter specifies the list of characters to remove.

For example, `RemoveChars('my_test_string', '_')` returns **myteststring**.

RemoveWord
> removes a word from an input string where the words are separated by spaces, periods, or other special characters. The first parameter specifies the input string. The second parameter specifies the number of the word to remove where 1 is the first word.

For example, `RemoveWord('my test string', 2)` returns **my string**.

Replace
> replaces a substring within the input string with a replacement string. The first parameter specifies the input string. The second parameter specifies the substring to replace. The third parameter specifies the replacement string. The fourth parameter specifies which instances of the substring to replace. Select one of the following:

ALL
> replaces every instance.

FIRST
> replaces the first instance only.

LAST
> replaces the last instance only.

For example, `Replace('my test string', 'test', 'new', '_ALL_')` returns **mynewstringnew**.

ReplaceWord
> replaces a word from an input string where the words are separated by spaces, periods, or other special characters. The first parameter specifies the input string. The second parameter specifies the number of the word to replace where 1 is the first word. The third parameter specifies the replacement string.
>
> For example, `ReplaceWord('my test string', 2, 'new')` returns **my new string**.

Reverse
> reverses the order of the characters in the input string.
>
> For example, `Reverse('A B C')` returns **C B A**.

Substring
> returns a substring from the input string based on the position of the characters. The first parameter specifies the input string. The second parameter specifies the position of the first character to return. The third parameter specifies the number of characters to return.
>
> For example, `Substring('my test string', 4, 3)` returns **tes**.

Update
> replaces a substring from the input string based on the position of the characters. The first parameter specifies the input string. The second parameter specifies the position of the first character to replace. The third parameter specifies the number of characters to replace. The fourth parameter specifies the replacement string.
>
> For example, `Update('my test string', 4, 3, 'nex')` returns **my next string**.

URLDecode
> removes URL encoding from the input string. URL encoding replaces some characters with a % character followed by a two-digit hexadecimal code.
>
> For example, `URLDecode('support.sas.com/my%20string')` returns **support.sas.com/my string**.

URLEncode

applies URL encoding to the input string. URL encoding replaces some characters with a % character followed by a two-digit hexadecimal code.

For example, `URLEncode('support.sas.com/my string')` returns **support.sas.com/my%20string**.

Calculating Compound Annual Growth Rate

Compound annual growth rate (CAGR) is an investing and business term for the effective constant year-over-year rate of return that produces a target result value at the end of multiple years, assuming that the CAGR is compounded at the end of each year. For example, you might use CAGR to compare trends over multiple years in revenue or in the number of units sold. For SAS Visual Analytics, you can calculate the yearly CAGR using the expression builder.

This example compares trends in the growth rate for yearly sales amounts between different product types or regions.

The basic data items are:

sales

This is a numeric measure with a currency format and a default aggregation of Sum.

ProductType

This is a string category data item.

RegionName

This is a string category data item.

TransactionDate

This is a date data item with a format of Month, Day, Year (MMDDYYYY).

The duplicate data item is:

TransactionDateYear

This data item is a duplicate of the `TransactionDate` data item, but with the Year format.

The calculated data items are:

BeginningYearNum

This data item should be a numeric type with a `Float4.0` format and an aggregation of Minimum.

```
BeginningYearNum = Year('31DEC2010'd)
```

EndingYearNum

This data item should be a numeric type with a `Float4.0` format and an aggregation of Minimum.

```
EndingYearNum = Year('transactionDate'n)
```

The aggregated measure data items are:

NumYears

This data item has a `Float4.0` format.

```
NumYears = Min [_ByGroup_] ('EndingYearNum'n) - Min
[_ByGroup_] ('BeginningYearNum'n)
```

BeginningValue

This data item needs to be set to the same currency format as the `sales` data item.

```
BeginningValue = PeriodWithDate(_Sum_, 'sales'n,
'transactionDate'n, _ByYear_, '31DEC2010'd)
```

EndingValue

This data item needs to be set to the same currency format as the `sales` data item.

```
EndingValue = Period(_Sum_, 'sales'n, 'transactionDate'n,
_ByYear_)
```

NormalizedRatio

This data item has a `Float12.2` format.

```
NormalizedRatio = 'EndingValue'n / 'BeginningValue'n
```

CAGR

This data item has a `Percent` format.

```
CAGR = ('NormalizedRatio'n Power ( 1 / 'NumYears'n ) ) - 1
```

To use the CAGR, you should add `TransactionDateYear`, `CAGR`, and any other categories of interest (for example, `RegionName`, `ProductType`, and so on) to a list table, a crosstab, or a graph.

Appendix 4

Conditions for Filters

In SAS Visual Analytics, filters are based on expressions that contain operators. Conditions enable you to easily add the most common operators to your expression. Depending on the type of data that is used by the filter, you can select from the following categories of filter conditions:

Table A4.1   Conditions for Character Data

=	Specifies that a matching value must match one of the filter values exactly.
Contains	Specifies that a matching value must contain the filter value.
EndsWith	Specifies that a matching value must contain the filter value at the end of the value.
In	Specifies that a matching value is in the list that you select. To select your list, choose the values from the drop-down list.
Missing	Specifies that a missing value matches the filter.
NotContains	Specifies that a matching value must not contain the filter value.
NotIn	Specifies that a matching value is not in the list that you select. To select your list, choose the values from the drop-down list.
NotMissing	Specifies that a nonmissing value matches the filter.
StartsWith	Specifies that a matching value must contain the filter value at the start of the value.

Table A4.2 *Conditions for Numeric Data and Date and Time Data*

<>	Specifies that a matching value must not be equal to the filter value.
=	Specifies that a matching value must be equal to the filter value.
<	Specifies that a matching value must be less than the filter value.
<=	Specifies that a matching value must be less than or equal to the filter value.
>	Specifies that a matching value must be greater than the filter value.
>=	Specifies that a matching value must be greater than or equal to the filter value.
Between [exclusive]	Specifies that a matching value must be greater than the first filter value and less than the second filter value.
Between [inclusive]	Specifies that a matching value must be greater than or equal to the first filter value and less than or equal to the second filter value.
In	Specifies that a matching value is in the list that you select. To select your list, choose the values from the drop-down list.
	Note: This condition is not available for continuous numeric data.
Missing	Specifies that a missing value matches the filter.
NotBetween [exclusive]	Specifies that a matching value must be less than the first filter value or greater than the second filter value.
NotBetween [inclusive]	Specifies that a matching value must be less than or equal to the first filter value or less than or equal to the second filter value.
NotIn	Specifies that a matching value is not in the list that you select. To select your list, choose the values from the drop-down list.
	Note: This condition is not available for continuous numeric data.

NotMissing Specifies that a nonmissing value matches the filter.

Appendix 5

Data Limits

Data Limits for SAS Visual Analytics Explorer

Some of the visualizations in SAS Visual Analytics Explorer (the explorer) have limits to the number of data values that they can display. The limit values are affected by the **Visualization data threshold** setting in the Preferences window. For more information about the Preferences window, see "Specifying Your Preferences for the Explorer" on page 145.

The following table displays the data limits for visualizations in the explorer:

Table A5.1   *Data Limits for the Explorer*

Visualization Type	Variation	Behavior when Limit Is Exceeded	Default Threshold	Minimum Threshold	Maximum Threshold
Table	—	Paging is applied to the table.	10,000	1,000	100,000

Visualization Type	Variation	Behavior when Limit Is Exceeded	Default Threshold	Minimum Threshold	Maximum Threshold
Crosstab	—	An error message appears.	41,000	9,500	50,000
Bar Chart	Not grouped or latticed.	The bar chart shows a subset that contains the first or last ranked values.	3,000	810	3,625
	Grouped or latticed.	An error message appears.	3,000	810	3,625
Line Chart	Single numeric or datetime category in the **Category** role.	The line chart shows a subset that contains the first or last ranked value.	10,000	5,630	11,250
	Single string category in the **Category** role.		4,000	1,380	4,750
	No string categories in the **Category** role, grouped or latticed.	An error message appears.	10,000	5,630	11,250
	One or more string categories in the **Category** role, grouped or latticed.		4,000	1,380	4,750

Visualization Type	Variation	Behavior when Limit Is Exceeded	Default Threshold	Minimum Threshold	Maximum Threshold
Scatter Plot	Two measures, not grouped.	The scatter plot is converted to a heat map.	40,000	9,375	48,750
	Two measures, grouped.	An error message appears.	40,000	9,375	48,750
	Three or more measures, not grouped.	The scatter plot is converted to a correlation matrix.	80,000 / number of measures	18,750 / number of measures	97,500 / number of measures
	Three or more measures, grouped.	An error message appears.	80,000 / number of measures	18,750 / number of measures	97,500 / number of measures
Bubble Plot	No categories.	The bubble plot is converted to a heat map.	25,000	7,500	30,000
	Grouped.	The bubble plot shows the top or bottom value by size.	500	150	600
	Latticed.	An error message appears.	1,050	4,900	50,000
	Grouped and latticed.		500	150	600
	Animated.		50,000	15,000	60,000
Network Plot	—	An error message appears	8000	100	10000
			1,000 nodes, 1,000 links		
Histogram	—	—	None	None	None
Box Plot	—	An error message appears.	2,900 boxes	800 boxes	3,500 boxes

Visualization Type	Variation	Behavior when Limit Is Exceeded	Default Threshold	Minimum Threshold	Maximum Threshold
Heat Map	No categories.	—	None	None	None
	One or more categories.	An error message appears.	3,000	1,250	3,500
Geo Map	Bubble overlay	The bubble overlay is converted to a coordinate overlay.	5,000	1,500	6,000
	Region overlay	The bubble overlay is converted to a coordinate overlay.	5,000	1,500	6,000
	Coordinate overlay	An error message appears.	82,000	10,000	100,000
Treemap	No additional levels.	The treemap shows a subset that contains the first or last ranked value.	4,900	1,050	6,000
	One or more additional levels.	The number of additional levels is reduced to 0. If the threshold is still exceeded, then the treemap shows a subset that contains the first or last ranked value.	4,900	1,050	6,000
Correlation Matrix	—		60 measures		

Visualization Type	Variation	Behavior when Limit Is Exceeded	Default Threshold	Minimum Threshold	Maximum Threshold
Decision Tree	Number of nodes.	An error message appears.	180	75	200
	Number of target values.	The decision tree displays truncated data.	50	30	50
	Number of predictor values.	An error message appears.	4300	1500	5000
Word Cloud	—	The word cloud displays the top words by frequency, by measure value, or by topic weight.	100	100	100

In addition to the data thresholds that are set by the Preferences window, there are server data limits that can be set by the system administrator. For information about the server data limits, see the topic "Manage High-Cardinality Data" in the *SAS Visual Analytics: Administration Guide*.

High-Cardinality Thresholds for Report Objects

Client-side thresholds for report objects are documented in the following table. These thresholds affect the designer and all of the viewers. The second column indicates the maximum number of unique values.

Table A5.2 *Client-Side Thresholds for Report Objects*

Report Object	Rows
Gauges (bullets, sliders, thermometers, dials, and speedometers)	10
Bubble plots	1,500
Bar charts (regular, targeted, dual axis, and dual axis bar-line)	3,000
Waterfall charts	3,000
Scatter plots	3,000
Line charts (regular and dual axis line)	4,000
Time series plots (regular and dual axis)	10,000
Treemaps	25,000
Pie charts	40,000
Tables (list tables and crosstabs)	40,000

In general, requests that exceed a client-side report object threshold cause an error message to be displayed. An exception is that for detail data, excess rows are truncated.

Note: Scatter plots always show detail data. List tables show details if the user selects the **Show detail data** check box. Bubble plots show details unless a category is assigned to the group role.

Note: If a data source contains more than a million records, queries are blocked for data items whose cardinality is greater than 10,000 unique items in any particular category.

Appendix 6

Troubleshooting in SAS Visual Analytics Designer

Issue: The Repair Report window is displayed.

Resolution:

When you open a report in the designer, you might be prompted to repair the report. For example, this can happen when one or more columns have been removed from the data source. You might be prompted to repair the report when you refresh a report or change a data source. You might be prompted to repair a report if your system administrator is still deciding which data items should be in a table.

To repair a report:

1 Decide whether you want to fix individual data items or all of the data items in the report.

   Fixes can be made at the individual data item level. For example, if two data items are removed from a report, then the Repair Report window displays separate entries for each data item. The report objects that are impacted by each data item appear below each of the data items.

Here is an example of the **Repair Report** window that is displayed when a missing column affects two report objects:
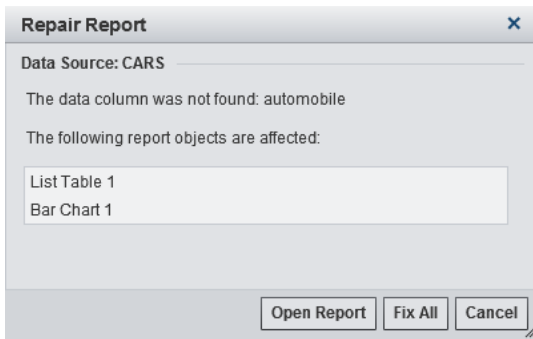
2   Click **Fix** or **Fix All**. For example, if there is a column missing in the data source and you click **Fix**, then the missing column is removed from the report.

Click **Close** in the confirmation message that is displayed.

Note: If you click **Open Report**, then all of the data is removed from the report objects.

Appendix 7

Using URL Parameters to View a Report

If guest access is configured for the SAS Visual Analytics Viewer, you can use this URL: `http://host/SASVisualAnalyticsViewer/VisualAnalyticsViewer/guest.jsp`. However, to log on and display a single report using guest access, you can use a URL such as this one (which has parameters to specify the name and location of the report): `http://host/SASVisualAnalyticsViewer/VisualAnalyticsViewer_guest.jsp?reportName=My+Report&reportPath=/Shared+Data/VA+Reports/&appSwitcherDisabled=true`. This second URL is useful for tasks such as using an <iframe> tag to incorporate a single report within another web application because the SAS Visual Analytics application bar, menu bar, and toolbar are not included.

The following table lists the parameters that can be included in the second URL:

Parameter	Description
reportName	Specifies the name of the report that you want to view. This parameter is required if you are accessing a specific report through a URL.
reportPath	Specifies the path to the report that you want to view. This parameter is required if you are accessing a specific report through a URL.

Parameter	Description
appSwitcherDisabled	Disables the application bar, which enables you to return to the home page or to access other parts of SAS Visual Analytics and other SAS applications. The default value is false.
	Note: If you are displaying in an <iframe> tag, you must disable the application bar by specifying `appSwitcherDisabled=true.`
commentsEnabled	Enables the **Comments** tab in the right pane. The default value is true. If both propertiesEnabled and commentsEnabled are false, then the right pane is not added to the viewer.
propertiesEnabled	Enables the **Properties** tab in the right pane. The default value is true. If both propertiesEnabled and commentsEnabled are false, then the right pane is not added to the viewer.
reportViewOnly	Hides the SAS Visual Analytics banner, including all of the menus.

Parameters are specified in the URL as a sequence of name and value pairs using query string syntax. The URL specifies your server, an absolute path to SAS Visual Analytics, and the query string (following the question mark character). Each parameter name in the query string is separated from the next value by an equal sign (=). Multiple name and value pairs are separated by ampersand characters (&). In this example, `reportName=My+Report` is the parameter that specifies the name of the report. The second name and value pair (`reportPath=/Shared+Data/VA+Reports/`) is the parameter that specifies the path to that report. The third name and value pair (`appSwitcherDisabled=true`) disables the application bar. It must be present if you are displaying the report in an <iframe> tag.

There are special rules for formatting name and value pairs in a URL. Special characters (such as most punctuation characters, including spaces) in a value must be URL-encoded. A space can be encoded as a plus sign (+) or %20. Other characters are encoded using the %*nn* convention, where *nn* is the hexadecimal representation of the character in the ASCII character set. In the previous example, the value `/Shared +Data/VA+Reports/` actually identifies the report path `/Shared Data/VA Reports/`. The spaces in the names are encoded as plus signs (+).

The reportName and reportPath parameters must be specified in the URL. The reportName parameter specifies the name of the report that you want to view (for example, **My Report**). The reportPath parameter specifies the path to that report (for example, **/Shared Data/VA Reports/**).

Note: You can use these parameters with the standard viewer URL that requires login credentials (**http://*host*/SASVisualAnalyticsViewer/ VisualAnalyticsViewer.jsp**). For example, you can use **http://*host*/ SASVisualAnalyticsViewer/VisualAnalyticsViewer.jsp? reportViewOnly=true** to hide the banner.

Glossary

Apache Hadoop

a framework that allows for the distributed processing of large data sets across clusters of computers using a simple programming model.

bar chart

a chart that consists of a grid and some vertical or horizontal columns (bars). Each column represents quantitative data.

bar-line chart

a bar chart with an overlaid line graph.

box plot

a graphical display of five statistics (the minimum, lower quartile, median, upper quartile, and maximum) that summarize the distribution of a set of data. The lower quartile (25th percentile) is represented by the lower edge of the box, and the upper quartile (75th percentile) is represented by the upper edge of the box. The median (50th percentile) is represented by a central line that divides the box into sections. The extreme values are represented by whiskers that extend out from the edges of the box.

calculated column

a column that does not exist in any of the tables that are accessed, but which is created as a result of a column expression.

capability

an application feature that is under role-based management. Typically, a capability corresponds to a menu item or button. For example, a Report Creation capability

might correspond to a New Report menu item in a reporting application. Capabilities are assigned to roles.

choropleth map

a two-dimensional map that uses color and fill pattern combinations to represent different categories or levels of magnitude.

co-located data provider

a distributed data source, such as SAS Visual Analytics Hadoop or a third-party vendor database, that has SAS High-Performance Analytics software installed on the same machines. The SAS software on each machine processes the data that is local to the machine or that the data source makes available as the result of a query.

crosstab

a two-dimensional table that shows frequency distributions or other aggregate statistics for the intersections of two or more category data items. In a crosstabulation table, categories are displayed on both the columns and rows, and each cell value represents the data result from the intersection of the categories on the specific row and column.

data item

an item in a data source that is either a logical view of a data field or a calculation. The author of a report decides which data items to use in a particular section of a report. There are three types of data items: hierarchies, categories, and measures.

data source

a table, view, or file from which you will extract information. Sources can be in any format that SAS can access, on any supported hardware platform. The metadata for a source is typically an input to a job.

dependency

a trigger condition that must be met before a job can run in a scheduled flow.

deployed job

a job that has been saved in a deployment directory and can be scheduled.

deployment directory
 the location for generated SAS DATA step programs that will be executed by the batch server as part of a scheduled flow.

file event
 a file-related occurrence that is used as a trigger in a scheduled flow. For example, a file event occurs when a scheduling server determines that a specified file exists.

filter
 specified criteria that are applied to data in order to identify the subset of data for a subsequent operation, such as continued processing.

flow
 a set of jobs and associated dependencies that is scheduled in the Schedule Manager plug-in in SAS Management Console.

heat map
 a graphical representation of data where the values taken by a variable in a two-dimensional map are represented as colors.

job
 a collection of SAS tasks that can create output.

job event
 a job-related occurrence that is used as a trigger in a scheduled flow. For example, a job event occurs when the scheduling server issues a command to determine whether a job ran successfully.

job flow
 a group of jobs and their dependencies, including dependencies on other jobs, on files, or on specified dates and times.

join condition
 a combination of join keys and a comparison operator.

list table

a two-dimensional representation of data, in which the data values are arranged in rows and columns.

local data

data that is accessible through the file systems on a computer. This includes data on hard drives or available through network file systems.

pie chart

a circular chart that is divided into slices by radial lines. Each slice represents the relative contribution of each part to the whole.

query

a set of instructions that requests particular information from one or more data sources.

remote data

data that is not accessible through the file systems available to a computer. To use remote data, you must direct a SAS server to access the data that is available through file systems on the remote machine.

report

output that is generated by running custom SAS code against the data in your project.

role

a set of capabilities within an application that are targeted to a particular group of users.

SAS Management Console

a Java application that provides a single user interface for performing SAS administrative tasks.

SAS Stored Process
a SAS program that is stored on a server and defined in metadata, and which can be executed by client applications. Short form: stored process.

scatter plot
a two- or three-dimensional plot that shows the joint variation of two (or three) variables from a group of table rows. The coordinates of each point in the plot correspond to the data values for a single table row (observation).

scatter plot matrix
a grid of scatter plots showing pairwise combinations of multiple numeric variables.

scheduling server
a server that runs deployed jobs in a scheduled flow. Before running a job, the scheduling server determines when the schedule for the deployed job as well as all of the dependencies for the job have been met.

source
See data source

subquery
a query-expression that is nested as part of another query-expression. Depending on the clause that contains it, a subquery can return a single value or multiple values.

time series
an ordered sequence of values of a variable that are observed at equally spaced time intervals.

user role
See role

visual exploration
a metadata object that contains visualizations and data settings that are saved from a session of SAS Visual Analytics Explorer.

visualization

an interactive visual representation of data. A visualization can be a table, a chart, or a geographic map.

waterfall chart

a form of data visualization that is used to understand or explain the cumulative effect on an initial value of sequentially introduced positive or negative values. Usually, the initial and the final values are represented by whole columns, and the intermediate values are denoted by floating columns.

Index