SAS® Visual Analytics 6.1

User's Guide

Contents

PART 7   Appendix   459

Using This Book

Audience

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

- Persons responsible for managing SAS servers and managing the SAS Visual Analytics environment.

- 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.

Prerequisites

Here are the prerequisites for using SAS Visual Analytics:

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

- A supported 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.

xvii

What's New

What's New in SAS Visual Analytics 6.1

SAS Visual Data Builder

The SAS Visual Data Builder application replaces the data preparation actions that were performed with the data preparation interface in the 5.2 release. The application is used to perform analytic data preparation. Enhanced features in SAS Visual Data Builder include:

- You can build queries from registered tables. The queries are used to join tables, add calculated columns for analytics, and sort data.

- You can use the pivot by feature so that data can be aggregated by the values of a categorical variable. For example, you can pivot a customer transactions table by the type of purchase or transaction type.

- You can specify a set of aggregations to apply automatically to all the numeric variables in a table. For example, you can specify that the minimum, maximum, mean, and standard deviation are calculated for the variables in a fact table.

- You can schedule queries to run automatically.

- You can load tables to memory on a SAS LASR Analytic Server. You can also design a query to stage data in a co-located data provider before loading the table to memory.

SAS Visual Analytics Administrator

The SAS Visual Analytics Administrator includes the following features that were available previously in the 5.2 release of the Data Preparation interface:

- **HDFS** tab

- **Resource Monitor** tab

- **Process Monitor** tab

- **Mobile Devices** tab

Enhanced features in the SAS Visual Analytics Administrator include:

- You can manage metadata-based security for folders and tables in SAS Folders.

- In the navigation pane for SAS Folders, you can select a folder or table, and right-click to select **Authorization**. In the Authorization page, you can view the effective permissions for tables in the metadata, and modify permissions. You can also set an explicit grant or denial to a table for a specific group or user.

- Row-level security is supported for LASR tables.

  Row-level permission conditions are specified in the New Permission Condition window for LASR tables. You access this window by navigating the SAS Folders tree in the navigation pane, selecting a table, right-clicking and selecting **Authorization**. Then, you can set permission conditions for a LASR table.

- The permissions origins feature identifies the source of each effective permission. Origins answers are available on an object's Authorization page.

- The **LASR Tables** tab enables you to start and stop the SAS LASR Analytic Server, and view information about the status, host, and port information for server connections. User account information for the start time of a SAS LASR Analytic Server connection, and the last time it was accessed is displayed.

- Within the **LASR Tables** tab, you can work with tables by loading them to the SAS LASR Analytic Server or by unloading them. The **Load a Table** window enables you to load tables.

- The Add Table window enables you to add tables to HDFS. The Add to data server window enables you to add tables to Greenplum and Teradata servers.

- In the **Manage Environment** panel within the Preferences window, you can record SAS statements to a single file or within separate files. This feature enables you to review the SAS statements for a task.

SAS Visual Analytics Home Page

New and enhanced features to the home page include:

- You can now refine your search results on the home page. For example, you can refine your search using type, author, the last modified date, or a combination of these.

- You view either thumbnails or lists of items on the home page.

- You can view and execute stored processes. Thumbnail images are available for stored processes.

- The **Links** pane links you to the overview video for SAS Visual Analytics. The pane also enables you to define and manage your own collection of links.

- The object inspector has been enhanced. For example, you can click with either your right or left mouse button to open the object inspector. When you click on an empty space, the object inspector is closed.

- Queries are displayed as content on the home page.

SAS Visual Analytics Explorer

New and enhanced features in SAS Visual Analytics Explorer include:

- A new treemap visualization type is available. Treemaps display your data as a series of rectangles, where each rectangle represents a category value or a hierarchy node.

- The correlation matrix visualization type can now be created directly.

- Crosstabs have been enhanced. You can navigate a hierarchy by drilling, expanding, and collapsing. You can also enable totals and subtotals for both rows and columns in the crosstab.

- New fit lines are available for data analysis, including linear, quadratic, cubic, and PSpline fits.

- A new forecasting feature enables you to add forecasted values to a line chart.

- A new **Comments** tab and Exploration Comments window enable you to create and view comments for visualizations and visual explorations.

- A new data brushing feature enables you to select data values in one visualization and highlight the corresponding values in your other visualizations.

- You can create new data items by using the new calculate data items feature. Calculated data items enable you to create new measures by using mathematical and logical operators in a calculation expression. You can also use datetime operators to calculate new date and time data items.

- You can derive distinct counts for your category data items. Distinct counts display the number of distinct values for a category and can be used as measures in your visualizations.

- The Edit Filter window has been enhanced to enable more advanced filter expressions.

- You can create new filters from the selected data values in a visualization. Filtering from a selection enables you to quickly remove or focus on a specific set of data values.

- You can create duplicate data items in order to use multiple aggregations for a data item in the same visualization.

- You can select the colors that are associated with each value of your category data items. The colors are used in any visualization that performs grouping.

SAS Visual Analytics Designer

New and enhanced features in SAS Visual Analytics Designer include:

- The new precision report layout enables you to drop objects anywhere onto the canvas and to size them as you like. You can overlap objects and control the depth order of these overlapping objects by sending objects backward or bringing them forward. You can scroll in the precision layout.

- A new geo map report object is available. You can also create geographic data items using latitude and longitude data. Any geo map that is generated by SAS Visual Analytics Explorer can be opened by SAS Visual Analytics Designer.

- You can add SAS stored processes as a report object in a report.

- You can use the new horizontal or vertical containers for laying out reports. Each container will hold one "row" or one "column" of report objects. The layout of the containers enables you to scroll objects beyond what is immediately visible on the screen.

- You can add new controls (drop-down lists, lists, button bars, text input fields, and range sliders) as report objects in reports. These report objects let you filter or control the scope of the data that you are currently viewing.

- The new report views enable you to change the size of the canvas that you use for laying out reports. There are three report views: **Full Screen**, **Tablet**, or **Wide-Screen Tablet**.

- A new preference enables you to specify which report view you want to use when you open new or existing reports.

- Using the updated **Data** pane, you can refresh to get new data items. A new feature also lets you duplicate data items. You can create hierarchy, geography, and calculated data items. You can also show or hide data items. The **Data** pane enables you to create global filters. User-defined formats in the data are now honored, but cannot be set in SAS Visual Analytics Designer.

- List tables have been enhanced. You can display spark lines as a column value. List tables can display totals and have display rules. Detail data is supported for list tables.

- Crosstabs have been enhanced. You can navigate a hierarchy by drilling, expanding, and collapsing. You can also create a hierarchy from a crosstab. You can enable totals and subtotals for both rows and columns in the crosstab. Crosstabs now have the ability to expand to show all the members of a data item.

- Interactions now allow list tables, crosstabs, and gauges as the source. You can now multi-select elements (for example, multiple rows in a table or multiple bars in a bar chart) in the source object of an interaction. You can also clear your selections for gauges, rows in a table, and so on, in the source object of an interaction.

- The titles for report objects can now use rich text elements for formatting.

- Most graphs now support display rules, hierarchies, and drilling.

- Bubble plots now support aggregated data. A group role was added.

- Text objects can how have hyperlinks.

- A new **Filters** tab in the right pane enables you to add a filter (or filters) to the selected report object. Filters can be set for the report, report objects, or both.

- List tables, crosstabs, bubble plots, treemaps, line charts, and bar charts that are imported from SAS Visual Analytics Explorer now support calculated columns, drilling, as well as editing properties and styles.

- Box plots, line chart with forecast, correlation matrix, heat map, histogram, geo maps, and the object with fit lines (linear, quadratic, cubic, and PSpline fits) that are

opened from SAS Visual Analytics Explorer now support the editing of properties and styles.

Documentation Enhancements

The *SAS Visual Analytics: User's Guide* uses new terminology to refer to the SAS Visual Analytics interfaces:

- The SAS Visual Analytics explorer interface is now SAS Visual Analytics Explorer.

- The SAS Visual Analytics designer interface is now SAS Visual Analytics Designer.

- The SAS Visual Analytics web viewer is now SAS Visual Analytics Viewer.

Accessibility

Accessibility Features of SAS Visual Analytics

Overview

SAS Visual Analytics has been tested with assistive technology tools. It includes accessibility and compatibility features that improve the usability of the product for users with disabilities, with exceptions noted below. These features are related to accessibility standards for electronic information technology that were adopted by the U.S. Government under Section 508 of the U.S. Rehabilitation Act of 1973 (2008 draft proposal initiative update). Applications are also tested against web content accessibility guidelines (WCAG) 2.0, part of the Web Accessibility Initiative (WAI) of the Worldwide Web Consortium (W3C). For detailed information about the accessibility of this product, send e-mail to `accessibility@sas.com` or call SAS Technical Support.

Documentation Format

Please contact `accessibility@sas.com` if you need this document in an alternative digital format.

User Interface Layout

SAS Visual Analytics contains the following interfaces:

■ The SAS Visual Analytics *home page* contains three main sections:

    □ The top of the home page contains the menu bar, which contains menus, a **Search** field, and a **Log Off** button.

    □ The center of the home page contains the workspace, where you can create content, view your content, or view content created by other users.

    □ The right side of the home page contains a **Common Actions** pane, which provides an alternate way to create reports, explore data, and manage your environment or favorites. The **Links** pane links you to the introductory video for SAS Visual Analytics or to other pages that you have book marked. The **SAS Resources** pane links you to resources on the SAS website and to social media.

■ *SAS Visual Data Builder* enables you to perform analytic data preparation operations. The application window contains four main sections:

    □ The top of the window contains the application name and an application bar that includes a menu bar and a **Log Off** button.

    □ The left side of the window contains a collapsible navigation pane. This pane contains a view of the SAS Folders tree. You can navigate the tree to select tables (one at a time) for use in preparing data for analysis.

    □ The center of the window (the workspace) contains tabs that you use for designing queries and data flows. Each query or data flow activity is performed on a tab.

    □ The bottom of the window contains a series of tabs that you use for creating calculated columns and filtering data.

    □ The right side of the window contains a collapsible properties panel that you use for managing the properties of queries, tables, joins, and data flows. The properties panel is dynamic and shows the properties of the item that is selected in the workspace.

■ *SAS Visual Analytics Explorer* enables you to explore data. The application window contains four main sections:

    □ The top of the window contains the Home bar, the menu bar, and the toolbar. The Home bar contains a **Home** button that enables you to return to the SAS Visual Analytics home page. The menu bar contains menus and a **Log Off** button.

    □ The left side of the window contains the data pane. The data pane contains all of the data items in your visual exploration. A table at the bottom of the data pane enables you to modify the properties of the data item that is currently selected.

    □ The center of the window contains the workspace. The workspace contains the visualizations (charts, tables, histograms, and maps) for your visual exploration. A dock bar at the bottom of the workspace enables you to manage your visualizations.

    □ The right side of the window is a pane that contains the **Roles**, **Filters**, **Properties**, and **Comments** tabs. The **Roles** tab enables you to manage data roles for the current visualization. The **Filters** tab enables you to create filters for the current visualization. The **Properties** tab enables you to manage properties for the current visualization. The **Comments** tab enables you to view and create comments for the current visualization.

■ *SAS Visual Analytics Designer* provides a framework for working with data and reports. The application window contains four main sections:

    □ The top of the window contains the Home bar and the menu bar. The Home bar contains a **Home** button that enables you to return to the SAS Visual Analytics home page. The menu bar contains menus and a **Log Off** button.

    □ The left side of the window contains two panes. The upper pane contains the **Objects** and **Imports** tabs. These tabs enable you to select report objects or import other reports or report objects. The lower pane contains the **Data** and the **Display Rules** tabs. These tabs enable you to select a data source (or data sources) and to create a new display rule. A table at the bottom of the **Data** tab enables you to modify the properties of the data item that is currently selected.

- □ The center of the window is the canvas (or workspace) where you design reports. Above the canvas, there is an area where you can drop a filter control and then a category to create section prompts.

- □ The right side of the window contains a pane for the **Properties** tab, which lists the properties for the currently selected report object. The **Styles** tab enables you to specify styles for the data, frame, and text, as well as colors for the data. The **Display Rules** tab enables you to populate or add intervals for the report object, like a gauge, which is currently selected in the canvas. The **Comments** tab enables you to add comments to the report after you save it. The **Filters** tab enables you to add a filter (or filters) to the selected report object. The **Interactions** tab enables you to add or update filter or brush interactions to the report objects in a section. The **Roles** tab enables you to add or update role assignments when you have a report object and a data source selected.

To customize the application window and its features, select **File ▶ Preferences**. On the SAS Visual Analytics home page, click either ⊞ or **Edit Preferences** to open the Preferences window. For more information about the layout and features of SAS Visual Analytics, see "Your First Look at the SAS Visual Analytics Home Page" on page 12.

Themes

An application's theme is the collection of colors, graphics, and fonts that appear in the application. The following themes are provided with this application: SAS Corporate, SAS Blue Steel, SAS Light, and SAS Dark. To change the theme for the application, select **File ▶ Preferences** and go to the **Global Preferences** page. On the SAS Visual Analytics home page, click either ⊞ or **Edit Preferences** to open the Preferences window. For more information about the layout and features of SAS Visual Analytics, see "Your First Look at the SAS Visual Analytics Home Page" on page 12.

You can also use keyboard shortcuts to magnify the contents of the browser window or to invert the application colors. For more information, see "Keyboard Shortcuts" on page xxix.

Note: If you have special requirements for your themes, then contact your system administrator or visual designer about using the SAS Theme Designer for Flex application to build custom themes. SAS Theme Designer for Flex is installed with SAS themes. For more information about this tool, see *SAS Theme Designer for Flex: User's Guide*.

Keyboard Shortcuts

The following table contains the keyboard shortcuts for the application. In the user interface, the shortcuts are displayed within parentheses in tooltips and menu labels.

Note: Some application-level keyboard shortcuts do not work when you first open an application. When that happens, press Tab to place the focus on the application, and then try the keyboard shortcut again.

Note: When you use a keyboard shortcut to activate a button, first give the focus to the field or section that the button is associated with before you use the keyboard shortcut. For example, if a table has an associated ⓘ button, you must first move the focus to the table before you press Ctrl+?.

Keyboard Shortcuts

Task	Keyboard Shortcut
Open a Help pop-up window from the ⓘ button.	Ctrl+? **Note:** This shortcut does not work on some keyboards (for example, the Italian keyboard).
Zoom in.	Ctrl++
Zoom out.	Ctrl+-
Reset the zoom state.	Ctrl+0

Task	Keyboard Shortcut
Maximize view (collapses the category pane and the tile pane, and hides the status bar and the application bar, which includes the menu bar and the workspace bar). or Exit maximized view (expands the category pane and the tile pane, and shows the status bar and the application bar).	Ctrl+Alt+Shift+M **Note:** This keyboard shortcut does not work when the focus is on the workspace bar.
Open a pop-up menu.	Shift+F9 (if a menu is available in that context) **Note:** If you use Shift+F9 to display the pop-up menu, then it is always displayed in the top left corner of the user interface control that you are using.
Open the Landmarks window.	Ctrl+F6
Temporarily invert or revert application colors (for the current session only). **Note:** You can set the **Invert application colors** preference in the Preferences window if you want the color change to persist across sessions.	Ctrl+~
Rename the selected tab.	Make sure that focus is on the tab. Press F2 and specify the new name. To commit your changes, press Enter. To cancel your changes, press Esc.
Close the selected tab.	Make sure that focus is on the tab, and then press Delete. **Note:** Some tabs cannot be closed.
Switch in and out of Edit mode for a table cell.	To enter Edit mode, select a cell and press F2. To exit Edit mode, press Esc.

Task	Keyboard Shortcut
Navigate between table headings and table content.	For a two-dimensional table, first make sure that focus is on the table and that you are not in Edit mode. Press Ctrl+F8 to switch focus between column headings and table cells. Use the arrow keys to navigate from heading to heading.
	For a multidimensional table, first make sure that focus is on a table cell and that you are not in Edit mode. Press Ctrl+F8 to switch focus between column headings, row headings, and table cells. Use the arrow keys to navigate from heading to heading.
Navigate the content rows of a table	When table cells are in Edit mode:
	■ Press Tab and Shift+Tab to move from cell to cell horizontally across columns.
	■ Press Enter and Shift+Enter to move from cell to cell vertically across rows.
	When table cells are not in Edit mode, use the arrow keys to move from cell to cell.
Sort columns in a table.	To sort a single column, navigate to the column heading of the column that you want to sort. Press spacebar to sort the column.
	To sort additional columns, navigate to the column heading of each additional column that you want to sort. Press Ctrl+spacebar.

Exceptions to Accessibility Standards

Exceptions to accessibility standards are documented in the following table.

Note: The JAWS issues occur when JAWS is used with Internet Explorer. Other browsers were not tested with JAWS, unless noted.

Exceptions to Accessibility Standards

Accessibility Issue	Workaround
Using JAWS with this application can sometimes cause Internet Explorer to stop responding.	Using JAWS 13 or later might resolve the issue.
JAWS cannot read some of the controls in the application, such as images, icons, and buttons.	No workaround is available.
JAWS cannot read the tooltips of items in trees and lists.	No workaround is available.
JAWS refers to table controls as list boxes.	When JAWS reports that a control is a list box, keep in mind that it might actually be a table.
Sometimes, JAWS does not correctly work with the controls in the Preferences window.	While you are in Virtual PC cursor mode, traverse the entire window to familiarize yourself with its contents before you change any of the settings. Switch to Forms mode or disable the Virtual PC cursor if you are not able to use a control in Virtual PC cursor mode.
JAWS does not correctly read the states in a tri-state check box tree if JAWS is not in Forms mode.	Disable JAWS' Virtual PC cursor when you work with the check box tree. Tab to the tree, and press Insert+Z to disable the Virtual PC cursor. When you finish interacting with the tree, press Insert+Z to re-enable the Virtual PC cursor.
The keyboard shortcuts that are used to interact with editable tables can conflict with keyboard shortcuts for JAWS' forms mode.	As a best practice, disable the JAWS virtual PC cursor when you work with tables. Tab to the table and press Insert+Z to disable the JAWS virtual PC cursor. When you finish interacting with the table, press Insert+Z to re-enable the JAWS virtual PC cursor.
JAWS cannot read two-column property tables.	

Accessibility Issue	Workaround
JAWS does not correctly read the information in a table: ■ JAWS cannot read the column headings of a table. ■ When table cells are not editable and focus is on the body of the table, JAWS reads an entire row at a time instead of cell by cell. ■ When table cells are editable, and focus is on the body of the table, JAWS reads only the first row of the table. If you use the arrow keys to select a cell or row, then JAWS does not read anything. If you press Enter to edit a cell, then JAWS reads the row that contains the edited cell.	No workaround is available.
When you are in a table cell, if you press Home, End, Page Up, or Page Down, the selected cell will change to be one in the first column of the currently displayed columns for the table.	Use the arrow keys to navigate through the cells of the table.
Sometimes, you cannot use F2 to put a table cell into Edit mode.	When you are in the cell that you want to edit, press Ctrl+F8 twice for two-dimensional tables and three times for multidimensional tables.
You cannot use the keyboard to scroll to the left and the right in some tables.	No workaround is available.
You cannot use the keyboard to activate the links within how-to topics and Help pop-up windows.	Use the **Help** menu to access the linked documents.
You cannot use Shift+F10 to open a pop-up menu.	Use Shift+F9 to open pop-up menus that are created for the SAS application. The generic menu that is provided by the Flash player cannot be opened by Shift+F9.
You cannot use the keyboard to access the close (x) button that is in the top right corner of a tab.	Make sure that focus is on the tab, and then press Delete to close the tab.

Accessibility Issue	Workaround
You cannot use the Tab key to move the focus outside of a code or expression editor. Pressing Tab within the editor only inserts tabs.	For Internet Explorer, press Shift+F10 and then press Esc to move the focus outside of the editor. For Firefox, press Alt+Tab to switch to another application. When you switch back, the focus will be outside of the editor.
You cannot use Ctrl+Alt+Shift+M to minimize or maximize the view if the focus is on the workspace bar.	No workaround is available.
If you use the Ctrl++ keyboard shortcut to zoom in to a secondary window, the outer edges of that window can become hidden from view if you zoom in too far.	Use the keyboard to access the hidden parts of the window.
When you use the Ctrl++ keyboard shortcut to zoom, the bottom entries of a drop-down menu can become hidden from view if you zoom in too far.	Use the keyboard to access the hidden entries in the menu.
The Ctrl++ and Ctrl+- keyboard shortcuts for zooming in and out do not work on some menus unless the menus are first opened.	Open the menu before you use the keyboard shortcut.
The Ctrl++ and Ctrl+- keyboard shortcuts for zooming in and out do not work on all elements in the application window (for example, tooltips and button labels).	No workaround is available.
If you maximize a tile in the **Home** workspace and then use the Tab key to navigate, the focus appears to be lost after you tab away from the **Log Off** button.	After you maximize the tile, press the Tab key 10 to 25 times, and eventually, the focus will return to the maximized tile.

Accessibility Issue	Workaround
You cannot use the keyboard to navigate in the **Layout** section because it is a read-only interface that is used for the visual verification of the elements that have been created.	Use the test button that is in the **Layout** section to preview your elements in a secondary window. The items that are displayed in the secondary window are identical to the items that are displayed in the **Layout** section, but unlike the items in the read-only **Layout** section, you can interact with the items in the secondary window. **Note:** After the application opens the secondary window, press Tab to place the focus in the window.
JAWS cannot read the labels for the **Red**, **Green**, and **Blue** fields in the Custom Colors window.	No workaround is available.
JAWS does not explain how to open a drop-down menu or drop-down list.	Press Ctrl+down arrow to open the control.
When JAWS reads the control names in a breadcrumb, it does not distinguish between the breadcrumb buttons that contain drop-down menus and those that do not.	Check for a drop-down menu by pressing Ctrl+down arrow on a breadcrumb button. A drop-down menu will open if one exists for that button.

Accessibility Workarounds for the SAS Visual Analytics Home Page

This is a workaround for the SAS Visual Analytics home page:

Accessibility Workarounds for the SAS Visual Analytics Home Page

Accessibility Issue	Workaround
When you navigate to the home page preference for the **Initial screen**, the item called **A report, exploration, or stored process** opens the Select window.	No workaround is available.

Accessibility Workarounds for SAS Visual Analytics Designer

These are workarounds for SAS Visual Analytics Designer:

Accessibility Workarounds for SAS Visual Analytics Designer

Accessibility Issue	Workaround
When you navigate to the **Data** tab to select a data source for the first time, the drop-down list opens the Add Data Source window.	No workaround is available.
When you navigate to the **Import** tab and select **Import a report** from the drop-down list, the Open window is displayed and receives focus.	No workaround is available.
You cannot access the list of tabs in the right panel of SAS Visual Analytics Designer using only the keyboard.	Use the **View** menu to give the tabs in the right panel focus. You can select **Comments**, **Display Rules**, **Filters**, **Interactions**, **Properties**, **Rules**, or **Styles**.
When JAWS in running, the page up and page down keys do not switch tabs properly.	No workaround is available.

Accessibility Issue	Workaround
The drop-down lists in SAS Visual Analytics Designer do not have labels, so they are announced by the screen reader as 'Combo Box.'	No workaround is available.
JAWS does not read the labels for the collapse and expand panes correctly. They are read as "splitter button zero".	Use the **View** menu to select either **Maximize View** or **Exit Maximized View**.

Accessibility Workarounds for SAS Visual Data Builder

These are workarounds for SAS Visual Data Builder:

Accessibility Workarounds for SAS Visual Data Builder

Accessibility Issue	Workaround
When you schedule a query and tab to the **Select one or more triggers for this query** radio button, the **New Time Event** does not become enabled. You cannot schedule a new query using only the keyboard.	No workaround is available.
When you schedule a query, you cannot tab to the radio buttons in the **Grouping conditions**.	Press the tab key to advance to the **OK** button and then press Shift+tab to move backward to the **Grouping condition** radio buttons.
When you schedule a query and specify a new time event, you cannot use the tab key to access all the radio buttons.	No workaround is available.

Part 1

Introduction to SAS Visual Analytics

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 and empowers organizations to explore huge volumes of data very quickly in order to see patterns and trends, and to identify opportunities for further analysis. SAS Visual 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 combined with the speed of the SAS LASR Analytic Server accelerates analytic computations and enables 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 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 importantly, reports can be displayed on a mobile device or in the SAS Visual Analytics 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, and 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 relationships. Then, when useful visualizations are achieved, you can incorporate those visualizations into reports that are available on a mobile device or in the SAS Visual Analytics Viewer.

SAS Visual Analytics provides users with the following benefits:
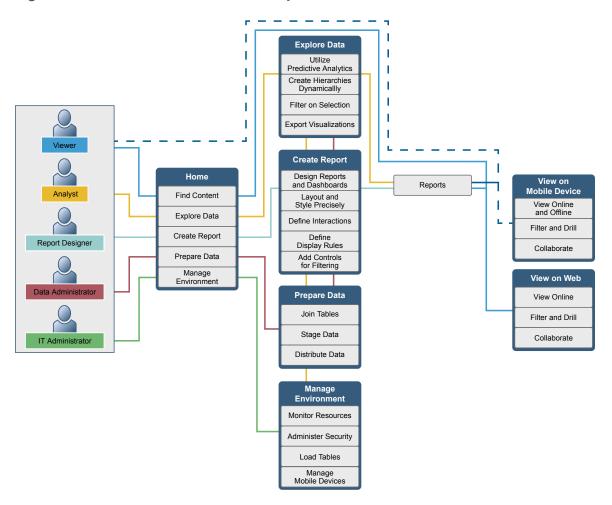
- enables users to apply the power of SAS analytics to massive amounts of data

- 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 in a web viewer. You can explore your data by using interactive visualizations such as charts, histograms, and tables. Report authors can easily point and click to query central sources of data. You can add filters and design the layout using tables, graphs, and gauges using drag and drop to create a well-formatted report.

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

Figure 1.1   *Overview of SAS Visual Analytics*

2

Accessing SAS Visual Analytics

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 is shipped with three predefined roles—Report Viewing, Analysis, and Administration. A predefined set of capabilities is available for each role. A system administrator can modify these roles and specify the capabilities 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.

Log On to SAS Visual Analytics

SAS Visual Analytics uses the standard logon window for SAS applications.

To log on to SAS Visual Analytics:

1   To display the SAS Visual Analytics logon 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://server01.abc.com:8080/SASVisualAnalyticsHub/`

Display 2.1   Log On Window for SAS Visual Analytics

2   To log on to SAS Visual Analytics:

    a   In the **User ID** field, enter your user ID.

    b   In the **Password** field, enter the password for the user ID that you just entered.

    c   Click **Log On**.

The SAS Visual Analytics home page appears. For more information, see "Your First Look at the SAS Visual Analytics Home Page" on page 12.

Note: Your password is case-sensitive. Your user ID might be case-sensitive, depending on the operating system that is used to host the web application server. If you need assistance, contact your system administrator.

Log Off from SAS Visual Analytics

To log off from SAS Visual Analytics, click **Log Off** in the upper right corner of the user interface.

Note: When you select **Log Off**, you are logged off from all environments.

If you are prompted about unsaved changes, click **Log Off** to exit without saving or click **Continue** to return to SAS Visual Analytics and keep working.

Here is an example of the time-out warning message:

Display 2.2   Session Time-out Warning Message

If you lose your connection to SAS Visual Analytics (for example, your session times out), then you must begin again at the point where you last saved your work. By default, if there is no activity for 30 minutes, then SAS Visual Analytics automatically logs you off and displays the logon window. Your system administrator can change the inactivity period. As a best practice, save your work frequently.

The following example shows the message that appears after your session has timed out.

3

SAS Visual Analytics Interface

Your First Look at the SAS Visual Analytics Home Page

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

Here are the features of the SAS Visual Analytics home page:

Figure 3.1   *SAS Visual Analytics Home Page*

1   The menu bar enables you to access task options, such as creating reports, exploring data, managing your environment or favorites, and getting help on using SAS Visual Analytics. You can also search all SAS content from the menu bar, and you can log off from SAS Visual Analytics. For more information about searching, see "Searching from the Home Page" on page 27.

2   The **Create Content** area provides icons to let you quickly create a new report or to explore data.

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

4   The **Other Content** area displays lists of explorations or reports that your administrator has added for you. If your administrator has not created any lists, then you will not see this content area. If you have an Administration role, then you will also see a **Manage** link that enables you to manage lists.

5   The **Common Actions** pane provides an alternate way for you to create reports, explore data, and manage your environment or favorites. If you have an Administration role, then you will also see a link for preparing data.

6   The **Links** pane links you to the introductory video for SAS Visual Analytics or to other pages that you have book marked.

7   The **SAS Resources** pane links you to resources on the SAS website and to social media.

Once you start using the home page, thumbnails enable you to open explorations, reports, queries, or stored processes that you have created. The default view is for secure, generic thumbnails to represent the content. Your system administrator can set a property to use thumbnails that are shared and unique to each content type.

Here are examples of the generic thumbnails that you might see on the home page:
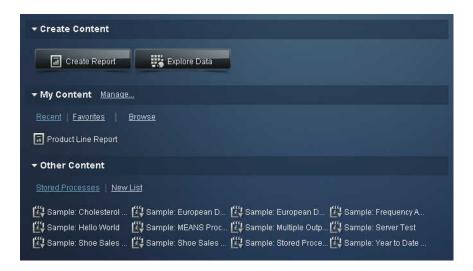
Display 3.1   *Generic Thumbnails for an Exploration, a Report, and a Stored Process*

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 view is the **Thumbnail** view.

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

Display 3.2   List View

For more information about other parts of the SAS Visual Analytics interface, see the following topics:

- "Your First Look at SAS Visual Data Builder" on page 37.

- "Your First Look at SAS Visual Analytics Explorer" on page 107.

- "Your First Look at SAS Visual Analytics Designer" on page 214.

- "Your First Look at SAS Visual Analytics Administrator" on page 385.

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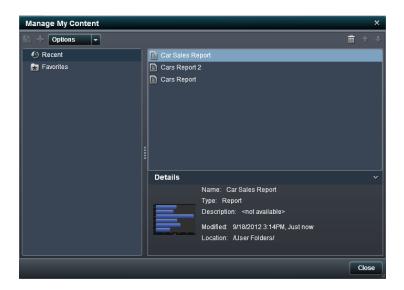
- your role and the associated capabilities. For example, you must have a Data Building role to prepare data.

- 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*.

Managing Your Content on the Home Page

The home page displays recent reports, explorations, and stored processes, as well as favorites, under the **My Content** heading on the home page. To manage your content, click **Manage** to the right of the **My Content** heading. This displays the Manage My Content window.

Display 3.3   *Manage My Content Window*

The icons in the Manage My Content window let you create a new favorites group, add a favorite, or select options.

Icon	Description
	Creates a new folder for your reports and explorations.
	Adds a report, exploration, or stored process to your list of favorites.
Options	Enables you to select either **Clear Recent History** or **Clear Favorites**.
	Enables you to delete a favorites folder, a report, an exploration, a stored process, or recent content one at a time. You can delete multiple favorite items by pressing the Ctrl key.
	Enables you to move favorites folders, reports, explorations, and stored processes up in the list to change what is displayed under the **My Content** heading on the home page.
	Enables you to move favorites folders, reports, explorations, and stored processes down in the list to change what is displayed under the **My Content** heading on the home page.

When you select a report, exploration, or stored process in the Manage My Content window, the details such as the name and location of that report are displayed. If you select a favorite group, the details include information about how many favorites the folder contains and its location.

Specifying Your Preferences
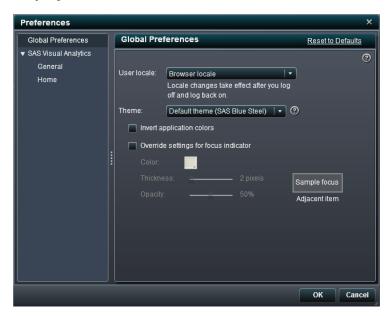
Specifying Global Preferences

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** to open the Preferences window.

   Note: If you are in the data preparation, exploration, or designing environments, then select **File** ▶ **Preferences** to open the Preferences window.

2   Make sure that **Global Preferences** is selected 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 setting for focus indicator** to change the color, thickness, and opacity of the focus in your SAS web applications.

Display 3.4 *Global Preferences in the Preferences Window*

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

5 If you changed your locale, then you must log off and log back on to SAS Visual Analytics for the changes to take effect.

Note: To restore all of the Global Preferences options, click **Reset to Defaults**.

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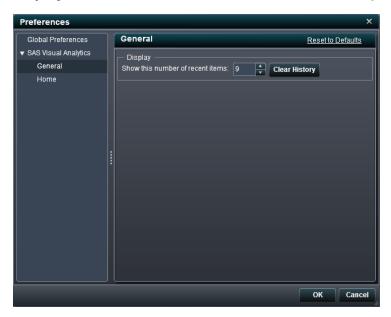
Using the SAS Visual Analytics home page, you can specify your general preferences for SAS Visual Analytics.

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 your preference 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.

Display 3.5   *Preferences Window for SAS Visual Analytics General Preferences*

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

Note: To restore the SAS Visual Analytics general preference, click **Reset to Defaults**.

Specifying the SAS Visual Analytics Initial Page Preference

Using the SAS Visual Analytics home page, you can specify your preference for the initial screen that you want to see after you log on to SAS Visual Analytics.

To specify your initial page preference:

1 Click either ▣ or **Edit Preferences** to open the Preferences window.
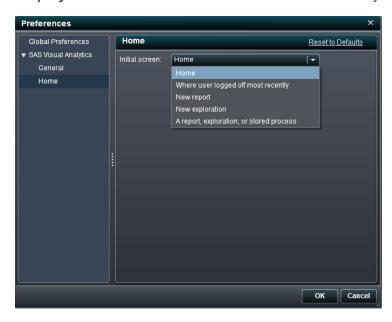
2 Click **Home** in the left pane.

3 Select one of the options from the **Initial screen** drop-down list.

The following options are available:

- **Home**

- **Where user logged off most recently**

- **New report**

- **New exploration**

- **A report, exploration, or stored process**

If you select **A report, exploration, or stored process**, then you can specify only one of these content types for your initial screen.

Display 3.6   *Preferences Window for SAS Visual Analytics Initial Page*

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

Note: To restore the SAS Visual Analytics option, click **Reset to Defaults**.

Specifying Preferences for SAS Visual Data Builder

See "Specifying Your Preferences for SAS Visual Data Builder" on page 39.

Specifying Preferences for SAS Visual Analytics Explorer

See "Specify Preferences That Are Specific to SAS Visual Analytics Explorer" on page 111.

Specifying Preferences for SAS Visual Analytics Designer

See "Specifying Your Preferences for SAS Visual Analytics Designer" on page 221.

Specifying Preferences for SAS Visual Analytics Administrator

See "Specifying Your Preferences for SAS Visual Analytics Administrator" on page 387.
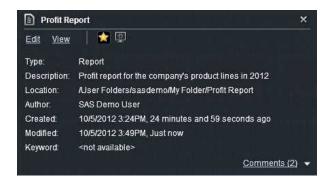
Discovering Details Using the Object Inspector on the Home Page

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 or exploration, where the report or exploration is saved, who the author is, and when the report or exploration was modified. A small toolbar lets you add the report or exploration to your favorites or set it as your preferred starting page. Note that when you set a favorite, the icon turns yellow (⭐).

Your role and capabilities determine the availability of other 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 for a report when generic thumbnails are used

Display 3.7 *The Object Inspector on the Home Page*

If you have the Report Viewing role and the View Report and Stored Process capability, then the toolbar enables you to open a report in SAS Visual Analytics Viewer by clicking **Open**.

If you have the Analysis, Data Building, or Administration role and the Create Report and Explore Data capabilities, then the toolbar in the object inspector enables you to do the following:

▪ **Edit** a report, which opens the report in SAS Visual Analytics Designer so that you can edit or change the objects in the report.

▪ **View** a report or stored process, which opens the report in SAS Visual Analytics Viewer.

▪ **Open** for an exploration, which displays SAS Visual Analytics Explorer.

▪ Click ☆ to add the report, exploration, or stored process to your list of favorites. When the icon is yellow (⭐), click to remove the object from your list of favorites.

▪ Click 🖥 to set (or clear) the report, exploration, or stored process as the initial screen when you start SAS Visual Analytics.

All users who have the Add Comments capability can add **Comments** to any content type or view existing comments about the content type using the link in the lower right of the object inspector. You cannot comment on folders or favorite groups.

Adding Comments to Reports and Explorations on the Home Page

If you have the Add Comments capability, then you can use the object inspector on the home page to add comments to any content type. You can also respond to existing comments for any content type. You cannot comment on folders or favorite groups.You cannot see visualization-level comments on the home page.

To add a comment:

1   Click a report or an exploration to open the object inspector.

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

Display 3.8   Comments on the Home Page

3 Enter a topic name and a comment.

Here is an example:

Display 3.9 Adding a Comment

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

4 (Optional) Click 📎 to add a file or an image to your comment.

To respond to an existing comment:

1 Click an exploration, a report, or a stored process to open the object inspector.

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

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

4 (Optional) Click 📎 to add a file or image to your reply.

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

Note: Only administrators can modify or delete an existing comment.

To search for a comment:

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

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

4
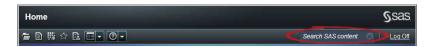
Searching for Reports, Explorations, and Stored Processes from the SAS Visual Analytics Home Page

Searching from the Home Page

You can search all reports, explorations, and stored processes that are on the metadata server from the menu bar on the home page.

Display 4.1 Search Field on the Menu Bar

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

▢ Using `2009 2010 2011` matches documents that contain any of the three terms.

▢ Using `+2009 +2010 +2011` matches only documents that contain all three terms.

▢ Using `2009 2010 -2011` matches documents that contain `2009` or `2010` but only if they do not contain `2011`.

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 field on the menu bar is not case sensitive.

Up to 5000 items can be returned in a single search. If your search exceeds 5000 items, then the first 5000 items appear and a message is displayed at the top of the window. The message tells you how many results matched your search and how many have been displayed. Narrowing your search using that search field enables the relevant search results to display. When you submit a new search using the menu bar that results in 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 menu 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 on the word *sample*. All objects that contain the word *sample* in the **Type** field are displayed in the workspace. Here is an example of the search results:

Display 4.2   *Search Results on the Home Page*

Refining Your Search Results

You can refine your search results following 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 report name, the description, or the author name 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 upon what has been selected in the other two panels.

The following panels are available:

■ The **Type** panel enables you to specify if you want to search for **Reports**, **Explorations**, or **Stored Processes**. The default is all object types.

 The available types are determined by the capabilities that your administrator has assigned. For example, you might see only reports in the list.

 If you select the check boxes for all three types, then the check boxes clear and the **All** check box is selected automatically for all available types. Also, if you clear the check boxes for all three 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 search for a specific author, or select one or more authors by name.

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

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

■ The **Modified** panel lists the modified dates of the objects that were found on the server. You refine the search to a specific date range by specifying a **From** date, a **To** date, or both.

 The dates are time zone dependent. If your company has offices in multiple time zones, then the time stamp on the report, exploration, or stored process reflects the time 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, then another user in California can see that it was saved 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. These selections change as you modify the items that you have selected in the **Type**, **Author**, and **Modified** panels. The breadcrumb also shows any search string that you entered in the **Search within results** field.

Your selections in the left panel are 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, then when you log off and back on, the **Type** panel remains collapsed and the **Refine Results** section remains at 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 on the word *sample*. Then, you refine the search results by selecting **Reports** for the **Type**. All reports that contain the word *sample* are displayed in the workspace. Here is an example of the search results:

Display 4.3 *Search Results on the Home Page*

Part 2

Preparing Data

5

Overview of SAS Visual Data Builder

What Is SAS Visual Data Builder?

SAS Visual Data Builder enables analysts and data administrators to perform analytic data preparation. You can build queries to perform joins, add calculated columns, and subset and sort data. Several productivity features speed the creation of columns based on common aggregation functions.

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

SAS Visual Data Builder also provides a series of features that are used in deployments that include SAS LASR Analytic Server. You can load an existing table directly into memory, load the results of a query into memory, or append rows to an in-memory table on a server.

Accessing SAS Visual Data Builder

You can access SAS Visual Data Builder by using any of the following methods:

- On the SAS Visual Analytics home page, select a query. In the details view, click **Edit**. Or you can double-click the query to open it in SAS Visual Data Builder.

- On the Visual Analytics home page, click **Prepare Data** to open SAS Visual Data Builder.

- Follow a web link or enter the URL to access SAS Visual Data Builder directly. For example, you might enter `http://hostname.example.com/SASVisualDataBuilder`.

Your First Look at SAS Visual Data Builder

The SAS Visual Data Builder interface contains the following elements:

1   The application bar enables you to return to the home page.

2   The menu bar contains menus that enable you to perform tasks.

3   The navigation pane displays a tree of tables and queries.

4   The lower pane contains tabs that enable you to specify column expressions and subset data.

5   The toolbar contains icons that enable you to develop and manage your queries.

6 The workspace enables you to design queries.

7 The properties panel (shown collapsed) enables you to manage the properties of the item that is selected in the workspace.

About Access to DBMS Data

SAS Visual 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 SAS Visual 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 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.

Working with User-Defined Formats

The preferred method for making user-defined formats available to a SAS Application Server is to name the 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*.

SAS Visual Data Builder can use existing user-defined formats, but you cannot specify to apply a user-defined format to a column.

6

Specifying Preferences for SAS Visual Data Builder

Specifying Global and General Preferences

To specify global SAS preferences, see "Specifying Global Preferences" on page 17. To specify general preferences, see "Specifying the SAS Visual Analytics General Preferences" on page 18.

Specifying Your Preferences for SAS Visual Data Builder

To specify preferences that are specific to SAS Visual Data Builder:

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

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

    Select a SAS Application Server to use from the menu. If you have added SAS Application Server instances to your deployment, make sure that Job Execution

Services has been configured to use the application server. 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 91.

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

7

Designing Queries

What Is a Query?

A query is a metadata object that manages the references to input tables, output and staging tables, joins, and summarizations from a SAS Visual Data Builder session. You can save queries and open them later to edit the data preparation operations that are

performed in the query. You can also use saved queries as subqueries when you design a new query.

About Designing Queries

You can use SAS Visual Data Builder to design queries that perform analytic data preparation. You can use a query to subset, sort, join, and add calculated columns to tables.

As you design your query, you can click ⬛ to validate your query. Use the ▶ button to preview or run the query.

The preview option uses a temporary table for the results. You can run the query after you specify an output table for the query and save the query.

Creating a New Query

You can create a new query by clicking ⬛ from the navigation pane.

Saving Your Query

To save your current query, select **File** ▶ **Save as** or click 🖫 and then select a location and a name.

When you save the query, you clear the undo and redo history.

Saving a Query as a New Query

Open an existing query from the SAS Folders tree that closely matches the query that you want to create. Select ⬛ to save the query with a new name and then perform your customizations.

Note: If the original query has output properties (table, library, and location), then the new query is saved with a new output table that has `_copy` added as a suffix to the name (*table-name*`_copy`). If a table with that suffix already exists, then the application changes the suffix to `_copy1` and tries to save the table. If duplicate names are found, the application continues trying to save the table by incrementing the number up to 50.

Using the Design View

The design view is the default view for working with queries. This view provides an easy to use interface for designing a query. You can perform the following tasks with the design view:

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

- join tables by using your pointer to select the source column and dragging 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 on the workspace or by right-clicking the table and selecting **Add All Columns**.

- use the **Column Editor** tab to specify column expressions, aggregations, and sort. You can also 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.

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

- open the Properties panel to specify the output table information.

See Also

- "Working with Tables" on page 49

- "Working with Columns" on page 56

- "Working with Query Filters" on page 67

- "Working with Joins" on page 71

Adding a Data Source to a Query

Adding a Table

To add a table to a 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 designed a query and saved it, it can be used as an input data source to a new query. To add a subquery to a query, use the SAS Folders tree to locate the query (⊞) and then drag and drop the query onto the workspace.

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

Best Practices for Adding Data Sources

When you plan to join data sources (tables or queries), 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 query.

If you are designing a query that uses a fact and dimension tables, 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 place a dimension table first, you can easily use the button to switch the left and right tables in the join.

If you are not working with fact and 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 that you can swap the left and right tables on the **Joins** tab.

Best Practices for Input Libraries

A query can use a single input library only. This enables SAS Visual Data Builder to take advantage of vendor-specific features in SQL processing. If the input tables are from a third-party vendor database, the SAS/ACCESS Interface engine can optimize the performance of the query by passing the SQL statements through to the database.

When you use a subquery as a data source, remember that you can join input tables to the subquery only when the subquery output table uses the same library.

Removing a Data Source from a Query

To remove a data source from a query, select the item on the workspace, right-click, and select **Remove Table**. This action applies to tables and to subqueries.

Specifying Properties for a Query

On the Properties panel, you can specify the following properties:

Name
   specifies the name for the query.

Location
   specifies the metadata folder location for the query object.

Description
   specifies a description of the 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 query.

Append data
   specifies whether the result set for the 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 query uses the SQL that is written by the user on the Code view when the check box is selected. For more information, see "Editing Code" on page 87.

Input library
   specifies the library that is used for selecting source tables. A query can use tables from a single input library only. For more information, see "Best Practices for Input Libraries" on page 45.

Table
   specifies the table name for either the staging output or the final output.

Library
   specifies the library to use for the staging table or the output table.

Location

specifies the metadata folder location to use for registering the staging or output table metadata.

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 80.

8

Working with Tables

Source Tables

SAS Visual Data Builder supports using tables and queries as input data for a query. When you drag a table onto the workspace, SAS Visual Data Builder connects to the SAS Metadata Server to determine the column names and data types for the table. When you drag a query onto the workspace, the application makes the same request, but retrieves the column names and types for the output table of the query.

Because SAS Visual Data Builder uses metadata for source tables, you must register the table metadata before you can access the data. Source tables are registered in

metadata with SAS Management Console (from a Windows platform) before you use SAS Visual Data Builder.

See Also

■ "Adding a Data Source to a Query" on page 44

■ *SAS Intelligence Platform: Data Administration Guide*

Specifying Properties for a Source Table

When a table is selected in the workspace, you can specify the following on the Properties panel:

Name
specifies the table name that is stored in metadata for the table. You can specify a new value to use as a table alias. The alias name is stored with the query and does not affect the metadata information for the table.

Auto-Aggregate
specifies whether to apply aggregation functions to the columns for this table when the columns are added to the query.

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

See Also

"Using the Auto-Aggregate Functions" on page 63

Output Tables

About Output Tables

When you design a query, you specify an output table on the Query Properties panel. When you save the query, the output table is registered in metadata. When you run the query, the physical table is created. The metadata that is created when you save a query enables you to use the table as an input table for another query or as a data source for another SAS application.

The **Table** field for the output table is a free-form text field. You cannot specify the name of a table that is already registered in metadata. If you edit a query that has already been saved and change the name of the output table, the metadata for the table is changed to use the new name. In other words, once you save a query with an output table, that table is associated with the query.

Specifying an Output Table

To specify an output table:

1 Open the Query Properties.

2 Expand the **Output** section.

3 Specify an output table name, library, and location. The table name field is restricted to 32 characters. Third-party vendor databases can have restrictions that further reduce that limit.

If you select a SAS LASR Analytic Server or SAS Data in HDFS library, the **Partition by** menu becomes available. Once 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*

Query Properties	

▶ General

▼ Output

Table:	Y2012
Library:	/Products/SAS Visual Analytics
Location:	/My Folder
Partition by:	(none)

▼ Staging

☑ Use a staging table

Table:	Y2012__stg
Library:	/Products/SAS Visual Analytics
Location:	/My Folder

▶ Performance

Staging Tables

About Staging Tables

SAS Visual Data Builder supports staging data. Staging data is a best practice because you can use SAS Visual Data Builder to access and transfer data from operational systems once rather than interfere with the operational systems and reduce their performance. Using SAS Visual Data Builder to stage data can also 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 query. The physical table

for the staging table is created when you run the query. The physical table is always replaced with the results of the query.

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

Specifying a Staging Table

To use a staging table:

1 Open the Query Properties.

2 The settings for the **Output** section must be complete.

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

4 Specify a library.

Output and Staging Table Interactions

The following table identifies the supported combinations for output and staging tables. The feature interactions for appending to tables and replacing tables is also included.

Table 8.1   *Output 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

Output Table	Staging Table	Append Data
SAS or DBMS *	SAS or DBMS *	Supported
SAS LASR Analytic Server	SAS or DBMS *	Supported ***
SAS LASR Analytic Server	SAS Data in HDFS **Note:** SAS LASR Analytic Server is the only supported output.	Not supported
SAS LASR Analytic Server	Teradata or Greenplum **	Supported ***
Any	SAS LASR Analytic Server **Note:** SAS LASR Analytic Server libraries are not supported for staging tables.	Not supported

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

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

See Also

Clearing Output Information

To clear the library and location fields for the output sections:

1 Clear the **Table** field. You can right-click and select **Select All** and then right-click and select **Cut**.

2 Place your pointer in the workspace and click the background.

9

Working with Columns

Adding Columns to a Query

Understanding How SAS Visual Data Builder Works with Columns

By default, 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 for the query.

You need to add the columns to the query that you want to use as input. Once the columns are added, you can specify column expressions, aggregations, and use the pivot by feature.

Note: As an exception, the auto-aggregate feature requires that you set the default aggregations for the table before you add columns from the table.

When a data source is dropped on the workspace, the column types are represented by the following icons:
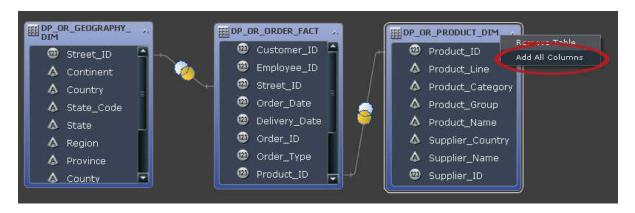
Table 9.1   Icons for Data Types

Icon	Description
123	This icon represents numeric data.
A	This icon represents character data.

Adding All the Columns from a Table

To add all the columns from a table to a query, select the table on the workspace, right-click, and select **Add All Columns**.
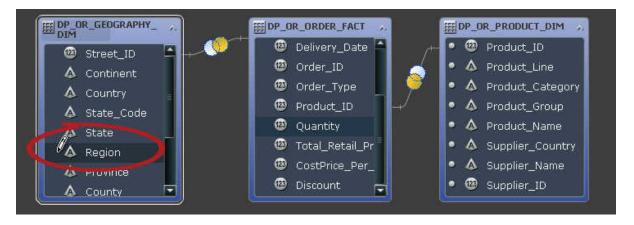
Display 9.1   *Add All Columns*

Adding a Single Column from a Table

To add one column to a query, select the column name with your pointer.

Display 9.2   *Add One Column*

> **TIP** If you select the column name an additional time, the column is added to the 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, the **Column Name** field has a number added to the name. If you change the name, make sure that you do not leave more than one column with the same name.

Adding a New Column Manually

To add a new column manually:

1 Select the **Column Editor** tab.

2 Click **+** at the bottom of the table.

3 Specify values for the **Column Name**, **Expression**, and **Type** fields. The remaining fields are optional.

Removing Columns

To remove a column from a query:

1 Select the **Column Editor** tab.

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

Display 9.3 *Removing a Column*

Specifying a Column Expression

To specify a column expression:

1   On the **Column Editor** tab, select `table-name.column-name` from the **Expression** column. If you added a new column manually, then you need to specify the table and column name.

2   Click ![icon] to open the expression builder. The table and column name are added automatically as the default expression.

3   Use the **Fields** tab to select columns from the source tables that you have added to the query.

> **TIP** You can also begin entering the table and column name in the expression window to use the autocomplete feature.

4   Use the **Functions** tab to select the functions to apply to the source column.

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

6   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. Click the ellipsis button to select the aggregations to use.

2 On the Choose Aggregations dialog box, select the **Aggregate functions** radio button.

Display 9.4   *Choose Aggregations Dialog Box*

3 Select the check boxes for the aggregate functions to use. Click **Apply**.

If you want to apply a common set of aggregation functions to a group of columns, you can select the columns on the **Column Editor** tab, right-click, and select **Choose Aggregations**.

Display 9.5 *Choose Aggregations Menu Item*

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

Display 9.6   Output Columns Tab with Aggregations

#	Column Name	Type	Format
1	Product_ID	NUMERIC	13.
2	Quantity_MIN	NUMERIC	
3	Quantity_MAX	NUMERIC	
4	Quantity_AVG	NUMERIC	
5	Total_Retail_Price_MIN	NUMERIC	
6	Total_Retail_Price_MAX	NUMERIC	
7	Total_Retail_Price_AVG	NUMERIC	
8	CostPrice_Per_Unit_MIN	NUMERIC	
9	CostPrice_Per_Unit_MAX	NUMERIC	
10	CostPrice_Per_Unit_AVG	NUMERIC	
11	Discount_MIN	NUMERIC	
12	Discount_MAX	NUMERIC	
13	Discount_AVG	NUMERIC	

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

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** from the context menu.

> **TIP** This menu option is available only when all 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 in the design view and click **Properties** to view the table properties.

2   Use the **Auto-Aggregate** menu to set the value to Enabled.

3   Place your pointer in the **Functions** cell and then click the ellipsis button 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 query, the selected aggregate functions are automatically applied.

Using the Pivot By Feature
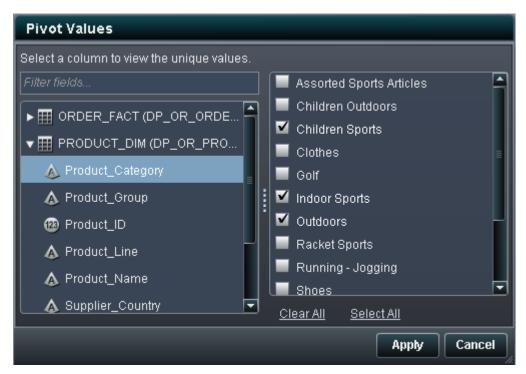
The pivot by feature provides an easy, yet very powerful, way to summarize data for analytics. You can specify a column to use as a categorical variable and the distinct values to use. When the query is run, the output table is summarized with the aggregations that you apply to the columns of interest.

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   On the Pivot Values dialog box, select the pivot by column. You can enter search criteria in the text box to filter the column names.

The following display shows an example of pivoting on three values of the Product_Category column.

Display 9.7 *Pivot Values*

3 After the distinct values for the column are determined, 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) three values of the

Product_Category column. Notice that TRP is specified as a label for the Total_Retail_Price column.

Display 9.8 *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_Price	ORDER_FACT.Total_Retail_Price	NUMERIC		TRP	MIN;MAX	PRODUCT_DIM.Product_Category
+								

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 the column name and the aggregation method is used as a suffix. The pivot by column label and aggregation method are used in the output column label.

Display 9.9 *Output Columns Tab with Pivot By Values*

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

10

Working with Query Filters

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 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 calculated columns. This tab is typically used to subset data on calculated columns that appear in the output table.

Specifying a WHERE Clause

To add a WHERE clause to a query:

1 Select the **Where** tab.

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

> **TIP** You can type in the filter fields text field to locate the column.

3 Double-click the column or drag it to 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 the *column-name* **Values** button to select the distinct values of the column. Select and clear the check boxes to specify the values to include in the filter.

   Click **Apply** to add the distinct values to the text area on the right.

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

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

7 Click 💾 to save the filter with the query.

Specifying a HAVING Clause

To add a HAVING clause to a query:

1 Select 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 the filter fields text field to locate the column.

3 Double-click the column or drag it to the SQL expression area.

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

5 For character variables, you can click the *column-name* **Values** button to select the distinct values of the column. Select and clear the check boxes to specify the values to include in the filter.

   Click **Apply** to add the distinct values to the SQL expression area on the right.

6 Edit the HAVING clause in the SQL expression area as follows:

Number of Values	Edit	Example
One character value	Add an equal sign between the column name and the distinct value.	`CARS.Make = 'Acura'`
More than one character value	Specify an IN operator and enclose the distinct values in parentheses.	`CARS.Make IN ( 'Acura', 'Audi')`
Numeric comparison	Specify a numeric operator and a constant, or a numeric operator and another column name.	`AVG(CARS.Horsepower) >= 165` `PRDSALE.Actual > PRDSALE.Estimate`

7   Click 🖫 to save the filter with the query.

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 a full table scan being performed 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 from the query. Otherwise, if you remove a table that is referenced in a filter, all the filtering conditions are cleared.

11

Working with Joins

About Joins

SAS Visual Data Builder supports joins for tables and queries. You can join tables to each other, including self joins. You can also join queries to tables and queries to queries. When you use a query in a join, the join condition is made against the output table for the query.

When you drag a table or query onto the workspace, the application attempts to determine a join condition automatically. The application retrieves the table metadata

that is registered with the SAS Metadata Server. If a foreign key reference is registered with the table metadata, then the application adds a join based on the foreign key information. If there is no foreign key information stored in metadata, then the application compares column names and types to determine whether there is a match. If a match is found, then the matching columns are used to perform the join.

When the application creates a join automatically, using foreign keys or by matching columns, the join is added as an inner join. The application 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 the ease with which joins are added to a query, SAS Visual Data Builder also generates an SQL statement with all 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 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 queries automatically as you add them to the workspace. When you drop a table or query onto the workspace, information about the table is retrieved from the SAS Metadata Server. For queries, the metadata for the output table is retrieved.

Using Foreign and Primary Keys

If primary or foreign key information is registered in metadata for the table that you drop on the workspace, then the application retrieves the foreign and primary key information.

The application then iterates over each of the tables that are already on the workspace in the same sequence that they were added to the workspace. The application retrieves the foreign 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 application continues to search for matches between the two tables and adding join conditions. Once a set of join tables is identified, the application does not continue iterating over the tables that are already on the workspace.

Matching by Name

If there is no foreign or primary key information for the table that is dropped on the workspace, then foreign or primary key information for the tables already on the workspace is not used.

The application retrieves the column information for all the columns in the newly added table. The application then iterates over each of the tables that are already on the workspace in the same sequence that they were added to the workspace. The length, type, and name for each column is compared 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 application does not continue to search for matches between the two tables and does not continue iterating over the tables that are already on the workspace.

Selecting the Join Type

When the application finds the first set of matching columns, by comparing keys or matching names, the application sets the join type for the two tables. The application

checks the metadata for the columns to determine whether the columns are nullable. The application 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, 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 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 over one of the columns to use, and then click and drag your pointer to the 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 select the join type from the menu.

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

Using the Joins Tab

To add a join condition by specifying the join properties:

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** field.

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

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

Removing a Join

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

■ Select the join from the workspace, right-click, and select **Remove Join**.

■ Click the **Joins** tab, select the row in the table, and click ✖.

Managing Joins in a Query

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

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

- ▪ add and remove tables from the join list

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

- ▪ swap the left and right table assignments for a join

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

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

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

- ▪ Swap the left and right tables by selecting the 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 performs the joins in a sequence that reduces performance. In these rare cases, you can change the sequence of the joins 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

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 as follows:

- ■ 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 analytic data preparation, and then staging the prepared data with a co-located data provider. Once the data is staged, the server can load the data to memory for exploration with SAS Visual Analytics Explorer.

The SAS Deployment Wizard registers a default library for the SAS LASR Analytic Server and one 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

```
/Products/SAS Visual Analytics High-Performance Configuration/
Visual Analytics Teradata
```

Greenplum Data Computing Appliance

```
/Products/SAS Visual Analytics High-Performance Configuration/
Visual Analytics Greenplum
```

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 is located in **/Products/SAS Visual Analytics High-Performance Configuration/Visual Analytics 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.

A SAS LASR Analytic Server instance 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 action loads the table from HDFS to memory on a SAS LASR Analytic Server instance.

You can also 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 is used to ensure all the rows with same formatted value of 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 results in a partitioned in-memory table.

See Also
"Partitioning Tables" on page 80

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 Query Properties panel 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 view because the 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 is located in **/Products/SAS Visual Analytics Administrator/Visual Analytics LASR**.

Output Libraries

SAS LASR Analytic Server libraries are supported as output libraries only. If you select one as an input library or a staging library, you receive an error.

You can use the **Append data** check box that you specify on the Query Properties panel 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.

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.

The feature uses the formatted values of the partition key to group rows that have the same value for the key. All 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 with the same key are in memory on one machine. For SAS Data in HDFS libraries, all the rows with the same key are written to a single file block on one machine. (The block is replicated to other machines for redundancy.) When the table is loaded into 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, using 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: Administration Guide

Exploring Results

When you design a query that uses a SAS LASR Analytic Server library for the output table and run the query, the **Explore Results** button becomes available. Click the button to access SAS Visual Analytics Explorer and the results of the query.

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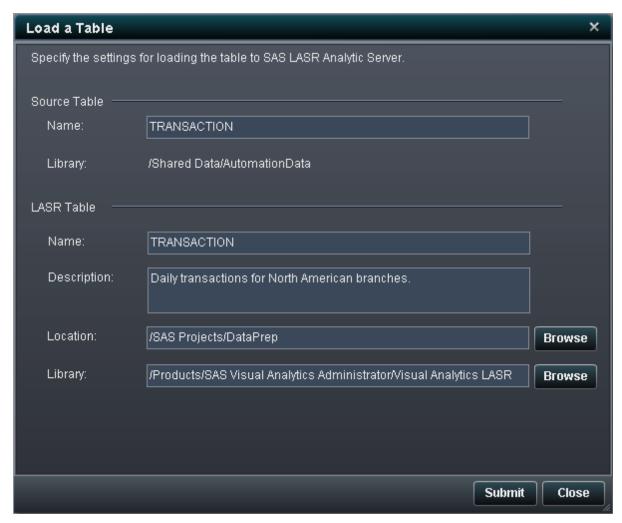
To load a table into memory on a SAS LASR Analytic Server:

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**.

Display 12.1   Load a Table Dialog Box

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

Specify the following settings on the dialog box:

Table 12.1   *Load a Table Dialog Box Properties*

Field	Description
Name	Enter the name to use for the 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	Click **Browse** and select a folder to use for the table metadata.
Library	Click **Browse** and select the SAS LASR Analytic Server library to use with the table.

4 Click **Submit**.

13

Importing Data

About the Import Data Feature

The Import Data feature enables analysts to perform self-service ad hoc data analysis. You can take a Microsoft Excel spreadsheet from your desktop and transfer it to SAS data set, a database, SAS Data in HDFS, or SAS LASR Analytic Server. You can also import data from CSV files, tab-separated files, and other character delimited text files.

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

Once the data is in memory, you can explore the data and create visualizations with SAS Visual Analytics Explorer. Likewise, you can use the impressive reporting capabilities of SAS Visual Analytics Designer.

Importing Data

To import data from a file:

1   Select **File ▶ Import Data** and then select a file.

2 On the Import Data window, specify the following input file parameters:

Delimiter^{Applies to CSV}

Select the delimiter that is used in the text file to import. You can also specify a single character to use as a user-defined delimiter.

Select worksheet^{Applies to XLS}

Select the worksheet to import as a table.

Use data in the first row as SAS variable names

Select this check box when the worksheet has the variable names in the first row.

Trim spaces

Select this check box to remove leading and trailing spaces from character variables.

Data records begin on row

The default is to import records from the first row. If you select **Use data in the first row as SAS variable names**, then this field increments to 2.

3 Specify the output table settings:

Field	Description
Name	Specify a table name for the imported data.
Library	Click the 🗄 button and select the library to import the data to.
	If you select a SAS LASR Analytic Server library, then you can explore the data immediately after importing it.
Location	Click the ▬ button and select a metadata folder. The imported data is registered as a table in the specified location.

4 Click **Import**.

If the SAS Folders tree is open to the folder that was specified in the **Location** field, then close and open the folder or click 🔄 to see the table.

14

Editing Code

Using the Code View

You can click the ▦ Code button to view the SQL statements that are generated by the data preparation expressions in the design view.

From the code view, you can click the ≜ button to unlock the view for manual editing of the SQL statements. Once you have unlocked the view, the button icon changes to ≞ to indicate the status.

Once you have unlocked the view, you cannot continue using the design view to edit the query. However, if you have not saved your changes, you can click the ↩ button to revert back to using the design view.

Considerations for Manually Editing Code

Input Library

When you unlock the Code view for manual edits, you can choose the library that you use for input tables. Open the Query Properties panel, select the **Input library** field, and click ▤ to select a library.

SQL Statements

The code editor enables you to complete an SQL SELECT statement. For programmers that are familiar with the SQL procedure, the SELECT statement is automatically enclosed in the SQL procedure statement and terminated with a QUIT statement.

You do not need to supply the PROC SQL statement or the QUIT statement. Do not enter semicolons in the code editor.

You can enter CASE statements, IF...THEN...ELSE statements, and any functions that can be used in the SQL procedure. Keep in mind that you can click the ▣ button periodically to validate your code.

See Also
SAS SQL Procedure User's Guide

Output Table

Specify the output table and library on the Query Properties panel just as you would if you were using the Design view.

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 also use SAS Visual Data Builder to reschedule the 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 application stores them in the same metadata folder with the query. The objects are named according to the following pattern:

```
vdb_query_id_timestamp
```

Editing Queries That Are Already Scheduled

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

Caution About Scheduling Queries to Run Now

When you schedule a query, one of the choices for configuring the schedule is to run the query immediately. You can use this option by selecting **Run now** on the Schedule window.

Be aware that performing the following steps results in an error condition:

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

2   Run the query.

3   Open the **Results** view to look at the output.

4   Schedule the query and choose **Run now**.

The previous sequence of steps results in an error because SAS opens a lock when a SAS data set is opened for reading. When step 3 is performed, the output table is

locked an no other process can overwrite the output table. The following message is included in the SAS log:

> **Locked Data Set Error Message**
>
> `ERROR: A lock is not available for OUTPUTTABLE.`
>
> `ERROR: Lock held by process xxxx.`

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

Scheduling Preferences

Default Scheduling Server

You can specify a different scheduling server in your application preferences. Any queries that you schedule after the change are set to use the scheduling server that you specify.

Use the Server Manager plug-in to SAS Management Console to identify the scheduling servers that are included in your deployment. By default, your deployment includes a server that is named `Operating System Services - hostname.example.com`. This server is used as the default server.

The SAS Distributed In-Process Services Scheduling Server is an alternative server. To use this server, change the scheduling server preference. The default name is `SAS Distributed In-Process Services Scheduling Server`.

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

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

Default Batch Server

You can specify a different batch server in your application preferences. Understand the following considerations 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, SAS Visual Data Builder retrieves the value for the default batch server preference the first time you schedule the query. The application uses the value as a pattern and uses the first server that matches the pattern.

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 directory is associated (default value: `SASApp`)

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

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

When you schedule a query, the SAS statements for the query are saved in a file. The file is saved at the path that is associated with the deployment directory. The file is named according to the same pattern that is described in "How Does the Scheduling Feature Work?" on page 89.

SAS Visual Data Builder retrieves your deployment directory preference and uses that value to look up the deployment directory in SAS metadata. The application uses the first directory that matches the pattern.

You can specify a different name for the default deployment directory preference. For more information about deployment directories and using the Schedule 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 query and schedule it. If you edit an existing query that is already scheduled, the existing settings for scheduling server, batch server, and deployment directory are not updated with the new preferences. To change the setting for queries that are already scheduled, use SAS Management Console to redeploy the deployed job for the 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*, which are evaluated based on a specified time being reached

■ *File events*, which 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 (a condition that must be met in order to run a flow on the scheduling server).

To create a time event:
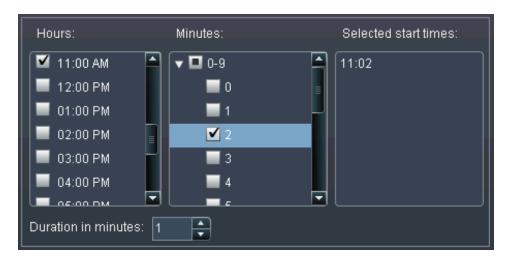
1 On the Schedule window, select **Select one or more triggers for this query** and then click **New Time Event**.

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

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

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

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 ten-minute intervals. Selecting a check box for the grouping selects all of the minutes in that grouping.

To select individual minutes, expand the grouping first.

If you select multiple values for **Hours**, all of the selections for **Minutes** apply to all of the selected hours. For example, if you select 2:00 AM and 5:00 AM in the **Hours** area and 43 in the **Minutes** area, the time event is scheduled for 2:43 AM and 5:43 AM.

5   If needed, specify the start date and end date for the 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 (a condition that must be met in order to run a flow on the scheduling server). The event can check for various file conditions, such as its existence, size, or age.

To create a file event:

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

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

3   Select the condition to evaluate for the file to make the 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 that you select.

Exporting Queries as Jobs

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

This feature enables you to work with a query interactively while you design it and then to export the query as a job for automating the operations. Once 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*.

Once a job is stored in metadata, you can open the job to edit it with SAS Data Integration Studio. 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 query and export it as a job again. However, you cannot overwrite job objects in metadata with SAS Visual Data Builder. If you need to delete objects from metadata (such as jobs, libraries, tables, or folders), use SAS Management Console.

To export a query as a job:

1   Use the SAS Folders tree to locate the query.

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

3   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 query and making it available for scheduling as a flow. However, 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 SAS Visual Data Builder does not support scheduling queries 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 query. If you use SAS Visual Data Builder to schedule the query afterward, the setting is lost and you need to repeat your steps with the Schedule Manager plug-in.

See Also
Scheduling in SAS

16

Using the Results View

About the Results View

Once you have designed your query and clicked ▶, the query runs and you can view the result set by clicking the 🖿**Results** button.

Data Pages

When you view data, a SAS server retrieves the data from the data set. Instead of returning all 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 output view 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 highlights the first row with the value. You can navigate through the results by clicking **Find next** and **Find previous**.

Click ⊡ to set advanced find options. The following list identifies the options:

■ **Options** tab

  □ **Exactly matches the specified string**

  □ **Contains the specified string** (This is a default setting.)

  □ **Starts with the specified string**

     ☐ **Case sensitive**

     ☐ **Trim leading and trailing spaces** (This is a 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 the rows, or a range of rows. The export feature has a limit of exporting 200,000 rows.

If you choose to export all the data or a range of rows, you have to click **Retrieve Data** before you can click **Export as CSV**.

Print

You can print the data that you view. After you click ⊟, you can set the following options:

■ number of rows to print

■ margin sizes

■ header and footer

■ scaling

You can also click **Print Preview** to preview the settings. Be aware that options for page size and orientation apply to the preview only. You must set those options when you click **Print** and select your printer.

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

You can export your exploration results as reports, which can be refined in SAS Visual Analytics Designer or viewed directly. Reports can be viewed on a mobile device or in the SAS Visual Analytics Viewer.

You can also save your explorations as PDF documents, share your saved explorations via e-mail, and export your visualizations as image files.
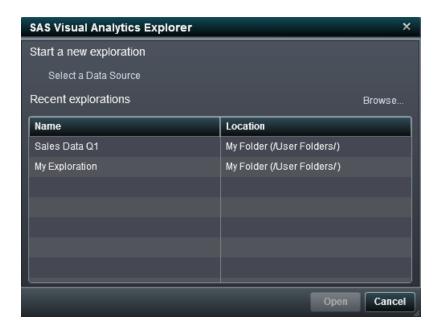
Accessing SAS Visual Analytics Explorer

You can access SAS Visual Analytics Explorer by using any of the following methods:

- On the Visual Analytics Home page, click **Explore Data** to open SAS Visual Analytics Explorer.

  You can select a data source by using the Welcome window.

- On the Visual Analytics Home page, double-click an exploration to open that exploration in SAS Visual Analytics Explorer.

- Follow a link or enter the URL to access SAS Visual Analytics Explorer directly. For example, you might enter
  http://server.abc.com/SASVisualAnalyticsExplorer.

  After you have logged on to SAS Visual Analytics, you can choose a data source from the Welcome window.

- From an e-mail, follow a link directly to a specific exploration. After you have logged on to SAS Visual Analytics, the exploration opens in SAS Visual Analytics Explorer automatically.

The Welcome Window

If you enter SAS Visual Analytics Explorer without a reference to a specific exploration, then the Welcome window appears.

SAS Visual Analytics Explorer   ✕

Start a new exploration

     Select a Data Source

Recent explorations   Browse...

Name	Location
Sales Data Q1	My Folder (/User Folders/)
My Exploration	My Folder (/User Folders/)

            Open   Cancel

The Welcome window enables you to perform the following tasks:

▪ Create a new exploration by clicking **Select a Data Source** to open the Open Data Source window.

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.

▪ Open an existing exploration. Select from your recent explorations, or select any exploration by clicking **Browse** to locate the exploration.

Your First Look at SAS Visual Analytics Explorer

Here are the main parts of the SAS Visual Analytics Explorer user interface:

1   The home bar enables you to return to the SAS Visual Analytics Home page and to select from your recent explorations.

2   The menu bar contains menus that enable you to perform common tasks, such as creating a new exploration and exporting the exploration as a report.

3   The toolbar contains icons that enable 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 manage the properties of the data item that is selected in the data pane

6   The workspace displays your visualizations. You can view many visualizations at once.

7   The right pane contains tabs that enable you to change the properties of your current visualization, subset the data that is currently being used, set the data roles for your visualization, and create and view comments.

8   The dock pane contains any visualizations that you have minimized and enables you to manage your visualizations.

18

Specifying Preferences for SAS Visual Analytics Explorer

Specifying Global Preferences

See "Specifying Global Preferences" on page 17 for details about global preferences for SAS Visual Analytics.

Specify Preferences That Are Specific to SAS Visual Analytics Explorer

To specify preferences that are specific to SAS Visual Analytics 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 2, "Data Limits for SAS Visual Analytics Explorer," on page 471.

  Note: A large value for **Visualization data threshold** might degrade your application performance or lead to time-out errors.

5 To return to the default settings, select **Reset to defaults**.

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

19

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 stored as part of the exploration, but do not affect the original data source.

You can manage the data properties for all of your data items by using the Data Properties window.

You can manage the data properties for individual data items by using the Data pane.

Manage Data Properties By Using the Data Properties Window

To access the Data Properties window, select **Data ▶ Data Properties** from the menu bar.

For each of your data items, 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, or a geography data item.

Model Type
specifies the 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 117 and "Select a Date or Time Format" on page 117 for details about specifying formats.

Aggregation (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.

Manage Data Properties By Using the Data Pane

You can manage data properties for a specific data item by selecting the data item in the Data pane and then interacting with the properties table at the bottom of the Data pane.

For the selected 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, or a geography data item.

Model Type

specifies the 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 117 and "Select a Date or Time Format" on page 117 for details about specifying formats.

Aggregation (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

To select a numeric format:

1   In the Data Properties window or the property table in the **Data** pane, click the format for the data item that you want to manage. A window appears.

2   Select the basic format type from the **Format Type** 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.

Select a Date or Time Format

To select a date or time format:

1   In the Data Properties window or the property table in the **Data** pane, click the format for the data item that you want to manage. A window appears.

2   Select the basic format type from the **Format type** list.

3   If necessary, select a more specific format from the **Format** list. The format names also represent sample values for the format.

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

Note: Time and datetime values in SAS Visual Analytics 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.

To manage which data items that are shown in the Data pane:

1   Select ▼ from the Data pane and then select **Show or Hide Items**. 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.

Assigning Colors for Category Values

About Colors

By default, SAS Visual Analytics 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

To set the colors for a category:

1   Right-click a category from the Data pane and then select **Colors**. A color selection window appears.

2   For each value that 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 values, click **Assign remaining**. This option automatically assigns colors to all of the

remaining values. These colors values are used consistently in all of your visualizations.

4   When you have finished assigning colors, click **OK** to apply your changes.

Reset Colors for a Category

To reset the colors for a category:

1   Right-click 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

To assign the color for a specific category value:

1   Within a grouped visualization, right-click 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. The color that you selected is used for the selected category value in all of your visualizations.

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 the data source that you want to explore.

Change the Data Source for Your Exploration

If a another data source contains the same data items as your current data source, then you can change the data source for your exploration. All of your visualizations, filters, and other settings are applied to the new data source.

To change the data source for your exploration:

1   Select **Data** ▸ **Change Data Source**.

2   From the Change Data window, select the data source that you want to use.

   Note: The new data source must contain all of the data items from the current data source.

3   Click **Open** to open the new data source.

Refresh Your Data Source

At any time, you can refresh your data source to read the latest changes.

To refresh your data, select **Data** ▸ **Refresh Data**.

Your visualizations are automatically updated with the new data.

Using Distinct Counts

About Distinct Counts

You can derive the distinct count of a category as a new measure. The derived measure does not contain data values in itself, but when it is used in a visualization, it displays the number of distinct values for the data item 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 of the derived measure and a category that contains product lines, then the bar chart shows the number of cities where each product line is produced.

Figure 19.1 *A Bar Chart with Derived Distinct Counts*

Derive the Distinct Count for a Category or Date

To derive the distinct count for a category or date, right-click a category or date in the Data pane and then select **Derive Distinct Count**. The derived measure appears automatically in the Data pane. The name of the derived measure is identical to the category or date that it is based on, with the suffix "(Derived)."

Support for Derived Data Items

Derived data items can be used in visualization types that support categorical grouping:

- Autochart
- Crosstab
- Bar Chart
- Line Chart
- Treemap
- Geomap

Explore Distinct Values in a Visualization

For bar charts, line charts, treemaps, and geomaps, you can right-click 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 derived measure to create a bar chart that displays the number of cities where each product line is produced:

Figure 19.2 *Bar Chart of Cities by Product Line*

Then, you might want to explore the distinct values for City where Promotional items are produced. By right-clicking the bar for Promotional and selecting **Explore Distinct Values**, you can see a new bar chart that contains the cities where Promotional items are produced, and the frequency of data for each city.
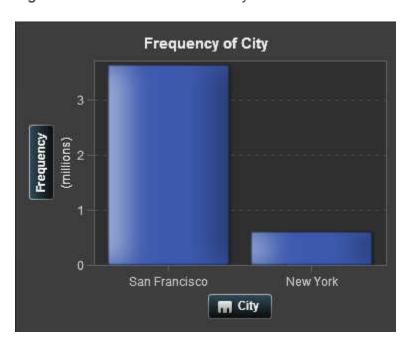
Figure 19.3 *Distinct Values of City for the Promotional Product Line*

Creating Calculated Data Items

About Calculated Data Items

SAS Visual Analytics Explorer 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.

Note: All calculations are performed on unaggregated data. The calculated expression is evaluated for each row in the data source.

In addition to performing mathematical calculations on numeric values, you can use calculated 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.
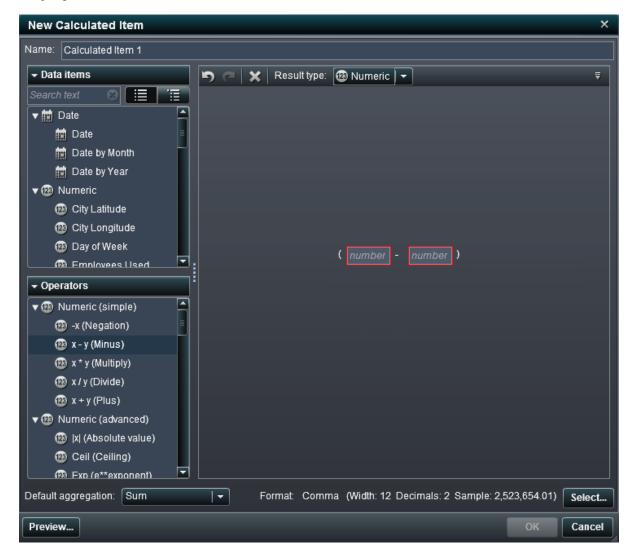
Create a Calculated Data Item

To create a calculated data item:

1 Select **Data** ▶ **New Calculated Item**.

The **New Calculated Item** window appears.

Display 19.1   The New Calculated Item Window

2 Enter a **Name** for the calculated data item.

3   Select the data type for the calculated item from the **Result type** drop-down list.

4   Build the expression for your calculated data item by dragging 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: Derived data items are not supported for calculation expressions.

   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, which displays how the expression would be changed if you drop the element at that location.

   For example, if your current expression is (`Profit / Revenue`), and you drag the **x - y (Minus)** operator over the left parenthesis symbol, then the expression changes to (`[number] — ( Profit / Revenue ))`. If you drag the operator over the division symbol, then the then expression changes to (`Profit — Revenue`), and so on.

   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 1, "Operators for Calculated Data Items," on page 461.

5   When you have finished creating your expression, select the **Default aggregation** type for the calculated 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 as well as 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.

Create a Duplicate Data Item

Duplicate data items enable you to use multiple copies of a data item 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.

To create a duplicate data item:

1   From the Data pane, right-click a data item and then select **Duplicate**. The **New Duplicate Item** window appears.

    Note: You can also invoke the **New Duplicate Item** window by dragging the same data item onto a visualization more than once.

2   Enter a **Name**, a **Format**, and a **Default aggregation** for the duplicated data item.

3   Click **OK** to create the new duplicated data item.

Delete a Calculated, Duplicate, or Derived Data Item

To delete a duplicate data item, a calculated data item, or a derived data item, right-click 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 your visualizations.

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 the 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 map visualizations 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, **Countries**. See "Define a Geography Data Item By Using a Predefined Role" on page 128.

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). See "Define a Geography with a Custom Role" on page 129.

Define a Geography Data Item By Using a Predefined Role

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 data role for the category to **Geography**. The Geography Role window appears.

3   Select the predefined geographic role for your category:

Countries
   specifies country names.

US States

specifies the names of states and territories in the United States.

US States Abbreviated

specifies the two-letter postal codes for states and territories in the United States.

US ZIP Codes

specifies five-digit ZIP codes for the United States.

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

Define a Geography with a Custom Role

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 data 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 coordinates.

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

20

Working with Visualizations

Overview of Visualizations

About Visualizations

SAS Visual Analytics Explorer displays data by using visualizations. A visualization is an interactive visual representation of your data. A visualization can be a table, a crosstab, a chart, a histogram, or a geographic map.

Visualization Types

You can assign any of the following types to your visualizations:

Autochart		Automatically selects the chart type according to the data that is assigned to the visualization. When you are first exploring a new data set, autocharts are useful to give you a quick view of the data. For more information, see "Working with Automatic Charts" on page 152.

Table

| | | | |

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 159.

Crosstab

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 table visualizations, crosstabs display aggregated data.

For more information, see "Working with Crosstabs" on page 160.

Bar Chart
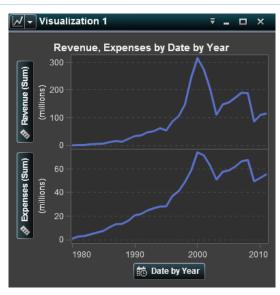
Displays the data as a bar chart. Bar charts are especially useful for comparing data that is aggregated by the distinct values of a category.

You can configure a bar chart with either vertical bars or horizontal bars. You can also assign grouping to the bars and create lattices.

For more information, see "Working with Bar Charts" on page 153.

Line Chart
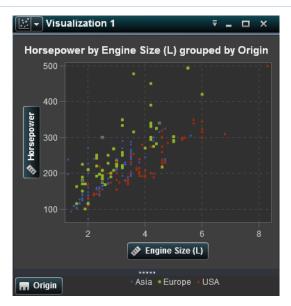
Displays the data as a line chart. A line chart is most useful for data trends over time.

You can apply grouping, and create lattices.

For more information, see "Working with Line Charts" on page 156.

Scatter Plot

Displays the data as a scatter plot. Scatter plots are most useful to examine the relationship between numeric data items.

In a scatter plot, you can apply statistical analysis with correlation and regression. Scatter plots also support grouping.

When you apply more than two measures to a scatter plot, the visualization automatically displays a scatter plot matrix to compare each pairing of measures.

For more information, see "Working with Scatter Plots" on page 162.

Bubble Plot
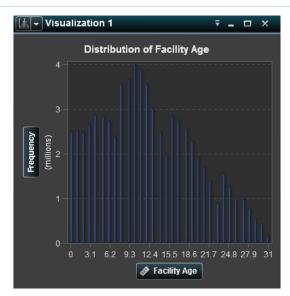
Displays the data as a bubble plot. A bubble plot displays the relationship between three measures, where two measures are represented by the plot axes and the third measure is represented by the size of the plot markers.

You can apply grouping and create lattices for a bubble plot. 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 164.

Histogram
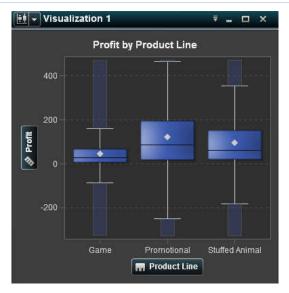
Displays the data as a histogram. A histogram displays the distribution of values for a single measure.

You can select the bar orientation, and select whether the distribution values are displayed as a percentage or as the row number of values.

For more information, see "Working with Histograms" on page 167.

Box Plot
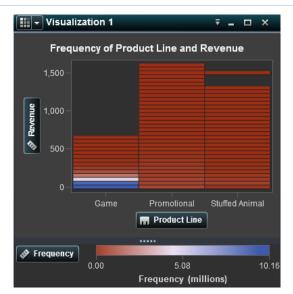
Displays the data as a box plot. A box plot displays the distribution of values for a measure by using a box and whiskers. The size and location of the box indicate the range of values that are between the 25th and 75th percentile. Additional statistical information is represented by other visual features.

You can create lattices, and select whether the average (mean) value and outliers are displayed for each box.

For more information, see "Working with Box Plots" on page 169.

Heat Map

Displays the data as a heat map. 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 the cell colors represents the frequency of each intersection of values. If you assign a measure to the color data role, then the cell colors represent the measure value for each intersection of values.

For more information, see "Working with Heat Maps" on page 171.

Geo Map

Displays the data as a geo map. A geo map displays your data as 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.

For more information, see "Working with Geo Maps" on page 174.
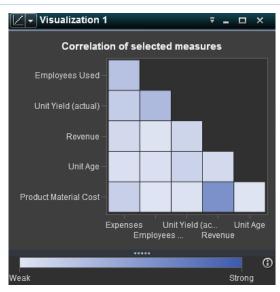
Treemap		Displays the data as a treemap. A treemap displays your data as a set of rectangles, where each rectangle represents a category or a hierarchy node. The size of the rectangles can indicate either the frequency of the category or the value of a measure. The color of the rectangles can indicate the value of an additional measure. For more information, see "Working with Treemaps" on page 176.
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 square indicates the strength of the correlation. For more information, see "Working with Correlation Matrices" on page 178.

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** 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 ▼ from the visualization toolbar and then select **Duplicate Visualization**.

The new visualization is named "Copy of *visualization-name*". You can assign a new name on the **Properties** tab.

Show and Minimize Visualizations

To minimize a visualization, click ▬ for the visualization that you want to minimize. The visualization appears on the dock panel at the bottom of the workspace.

To restore a visualization, click the name of the visualization on the dock pane.

You can also select which visualizations are visible by using the 🖿 button on the toolbar. The 🖿 button opens a window that enables you to select the visualizations that you want to show. The Manage Visualizations window displays a thumbnail image of each visualization and enables you to select and deselect each visualization by using a check box.

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 ▢ for 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 the title bar of the visualization to the location where you want to place the visualization.

When you drag a visualization onto another visualization, the placement of the visualization is determined by the edge that is closest to the drop point.

For example: Your workspace contains two visualizations named "A" and "B". If you drag A to the left edge of B, then A is placed to the left of visualization B. If you drag A to the bottom edge of B, then A is placed below B.

Display 20.1 shows an example layout with four visualizations.

Display 20.1   *Example Layout with Four Visualizations*

Display 20.2 shows the effect of moving a visualization to the right edge of another visualization.

Display 20.2   Example Layout After Moving 4 to the Right Edge of 1

When your workspace contains visualization in multiple rows, if you drag to the bottom 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.

shows the effect of moving a visualization to the bottom edge of another visualization.

Display 20.3   *Example Layout After Moving 4 to the Bottom Edge of 1*

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.

Manage Visualization Comments

About Visualization Comments

SAS Visual Analytics Explorer enables you to create and share comments about your visualizations.

In addition to visualization comments, you can create comments for the current visual exploration. See "Exploration Comments" on page 202.

View Visualization Comments

To view the comments for a visualization, select **Comments** from the right panel. The 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** field.

Create a Visualization Comment

Create a Comment under an Existing Topic

To add a comment under an existing topic:

1 Enter your comment text in the **Enter a comment** field that is beneath the topic that you want to contribute to.

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 add 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 Visualization Comment

Note: You must have the capability to edit comments.

To edit a comment, select the comment that you want to edit and then click ⬐.

Delete a Visualization Comment

Note: You must have the capability to delete comments.

To delete a comment, select the comment that you want to delete and then click 🗑.

Managing Visualization Data Roles

Add a Data Item

In order 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 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 the data item onto a **Measure** or **Category** button in the visualization to assign the new data item to a specific data role.

- From the visualization toolbar, select ▼ and then select **Add Category** or **Add Measure**.

- Use the **Roles** tab in the right pane. Either drag a data item to a role, or select a data item from the drop-down list.

Each visualization requires a minimum number of each type of data item. The following table lists the requirements for each visualization.

Table 20.1   Required Data Items for Visualizations

Visualization Type	Requirements
Autochart	one data item of any type
Table	one data item of any type
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
Histogram	one measure
Box Plot	one measure
Heat Map	two data items of any type
Geo Map	one geography
Treemap	one category or hierarchy
Correlation Matrix	two measures

Replace a Data Item

You can replace a data item by using any of the following methods:

- Drag the new data item from the **Data Items** pane to the data item in the visualization that you want to replace.

- Right-click the data item that you want to replace from either the visualization or 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 the data item from the visualization to the **Data Items** pane.

- Select ⏷ from the visualization toolbar and then select **Remove** ▶ *item-name* .

- Right-click the data item in the visualization or in the **Roles** tab and then select **Remove** *item-name*.

Switch Data Roles

To switch the data items that are assigned to two roles, drag one item onto another 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 21, "Working with Visualization Filters," on page 181.
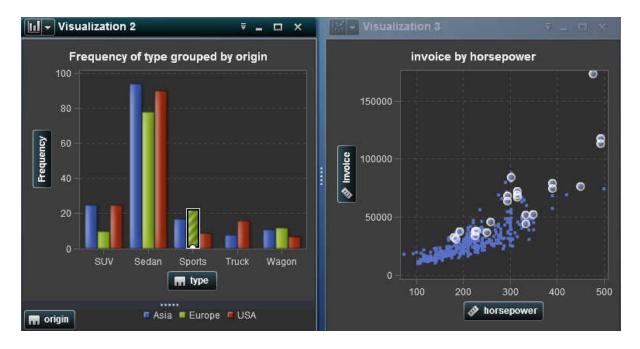
Working with Data Brushing

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 visual 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 chart are highlighted.

Figure 20.1   Data Brushing Example

The following visualization types support data brushing:

- bar chart

- line chart

- scatter plot

- bubble plot (if the group role is assigned)

- histogram

- heat map (if categories are assigned to both axes)

- geo map

- treemap

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.

In Figure 20.2, a selection in a bubble plot highlights a portion of the corresponding bar in a bar chart.

Figure 20.2 *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 values, press the Ctrl key while clicking each value.

- Click and drag to select values in a rectangular region. This is most useful for scatter plots and heat maps.

- Select values in the table view for the visualization. Press Shift+click to select multiple adjacent values, or Ctrl+click to select or deselect individual values.

Working with Automatic Charts

Automatic charts (autocharts) display different types of charts automatically based on the number of category and measure data items that are assigned to the chart.

The type of chart can be any of the following:

Table 20.2   Automatic Chart Types

Data Items	Chart Type
One measure	Histogram
One category	Bar chart
One datetime category and any number of other categories or measures	Line chart

Data Items	Chart Type
One geography and up to two measures	Geo map
One geography and three or more measures	Bar chart
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, change the visualization type to select the specific chart type that you want. For example, if your automatic chart shows a bar chart, then change the type to **Bar Chart** to enable advanced data role settings.

Working with Bar Charts

About Bar Charts

A bar chart represents the values of category data by using bars. The height of each bar represents the value.

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 titles**.

Generate graph titles
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: This option has no effect if a measure is assigned to the visualization.

Grouping style
specifies how grouped data are 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 grouping 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 values
shows the data values as text in the visualization.

Note: You can always view the data values as data tips when you move the cursor over the graph.

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.

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 values of the category 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 additional roles:

Group
>   groups the data according to 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 either shown as 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.

Sort Data Values

By default, a bar chart is sorted in descending order by the values of the first measure. To change the sorting, right-click 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.

Working with Line Charts

About Line Charts

A line chart represents the values of category data by using a line that connects the data values. If you assign multiple measures to a line plot, then you can create separate Y axes for each measure.

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 titles**.

Generate graph titles
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: 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 values
shows the data values as text in the visualization.

Note: You can always view the data values as data tips when you move the cursor over the graph.

Rotate axis labels
displays the category labels at an angle.

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.

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 values of the category are plotted on the category axis. You can assign many measures, and the measure values are plotted on the response axis. If you do not assign a measure, then the frequency is plotted on the response axis.

In addition to the basic data roles, you can assign additional roles:

Group
groups the data according to the values of the category data item that you assign. A separate line is created for each 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.

Sort Data Values

For category data, a line chart is sorted in descending order by the value of the first measure. To change the sorting, right-click the data item that you want to sort on, and then select **Sort ▶ [*sort-method*]**.

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, time, or datetime data item is assigned to the visualization.

In addition to the predicted data points, the forecast displays a confidence band. For more information, see "Forecasting" on page 208.

Enable Forecasting

To add data analysis to a line chart, select ▼ from the visualization toolbar and then select **Show Forecast**.

Note: The line chart must contain a date, time, or datetime item to apply forecasting.

On the **Properties** tab for the line chart visualization, you can adjust the number of data points to predict by using the **Duration** option.

Working with Tables

About Tables

A table shows data values as text. The values for each measure or category that is assigned to the table are displayed as columns. 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.

Specify Properties for a Table

On the **Properties** tab, you can specify the following option:

Name
  specifies the name of the visualization.

Data Roles for a Table

The basic data role for a table is a column. A column can be a data item of any type. You can add any number of columns to a table.

Managing Columns

Sort Columns

To sort the table by a column, click the column heading. An arrow appears in the column heading to indicate the sorting. If the arrow points upward, 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 your columns, click and drag the left or right edge of the column headings.

Working with Crosstabs

About Crosstabs

A crosstab shows the intersections of category values and measure values as text. If your 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 contains the frequency of an intersection of category values.

Specify Properties for a Crosstab

On the **Properties** tab, you can specify the following option:

Name
specifies the name of the visualization.

Show Column Subtotals
adds subtotals to each column, for each node on the row axis after the first.

Note: Subtotal values are aggregated according to the default aggregations for each measure.

Show Column Totals
adds totals to each column.

Note: Total values are aggregated according to 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: Subtotal values are aggregated according to the default aggregations for each measure.

Show Row Totals

adds totals to each row.

Note: Total values are aggregated according to 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.

Data Roles for a Crosstab

The data roles for a crosstab are columns, rows, and measures. You can assign either a single hierarchy or any number of categories to each of the columns and rows roles. 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 cells of the crosstab show the frequency of each intersection of values.

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 the heading for the row or column that you want to sort, and then select **Sort** ▶ [*sort-method*].

Note: Sorting is not available for measures.

Rearrange Rows and Columns

To rearrange your rows and columns, drag and drop the column headings.

Resize Columns

To resize your columns, click and drag the left or right edge of the column headings.

Create Hierarchies from a Crosstab

Crosstab visualizations 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."

Working with Scatter Plots

About Scatter Plots

A scatter plot represents the values of measures by using markers. When you assign more than two measures, the visualization displays a scatter plot matrix. A scatter plot matrix is a series of scatter plots that displays every possible pairing of the measures that are assigned to the visualization.

If you create a scatter plot that has a very large number of points, then the plot is rendered either as a heat map (for two measures) or as a correlation matrix for three or more measures.

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 titles**.

Generate graph titles

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.

Note: Sizes below 4 pixels are not supported when your scatter plot is grouped.

Fit line

adds a fit line to the visualization. For details about the fit types that are available, see "Fit Line" on page 207.

Note: Fit lines are not available if a grouping variable is assigned to the scatter plot.

Data Roles for a Scatter Plot

The basic data role for a scatter plot is the measures. 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 the measures, you can assign a **Group** variable. The group variable groups the data according to 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 also add data items to the **Labels** role. The values for the data items in the Labels role are displayed in the data tips for the scatter plot.

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 "Correlation" on page 206.

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-splines. For more information, see "Fit Line" on page 207.

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 ▼ from the visualization toolbar and then select **Fit Line** ▶ **[*fit-type*]**. For details about the fit types that are available, see "Fit Line" on page 207.

Working with Bubble Plots

About Bubble Plots

A bubble plot represents the values of three measures by using differently sized markers (bubbles) in a scatter plot. The values of two measures are represented by the position on the graph axes, and the value of the third measure is represented by the marker size.

You can also create animated bubble plots to display changing data over time.

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 titles**.

Generate graph titles
> 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.

Color gradient
> selects the gradient colors for the visualization.

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 additional roles:

Group
> groups the data according to 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** and the **Color** roles at the same time.

Bubble 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 bubble color represents the measure value.

Note: You cannot assign both the **Group** and the **Color** roles 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 set assign a data item to the **Group** role.

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 20.3   *Animation Control Tasks*

Task	Action
Start the animation	Click ►.
Go to the previous animation frame	Click ◄I.
Go to the next animation frame	Click I►.
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.

Working with Histograms

About Histograms

A histogram represents 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 axis matches your data values exactly.

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 titles**.

Generate graph titles

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 axis displays the number of values in each range (**Count**) or the percentage of values for each range (**Percent**).

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.

Data Roles for a Histogram

The basic data role for a histogram is a measure. You can assign only one measure to a histogram.

Working with Box Plots

About Box Plots

A box plot represents the distribution of data values by using rectangular box and lines called whiskers.

Figure 20.3 *Parts of a Box Plot*
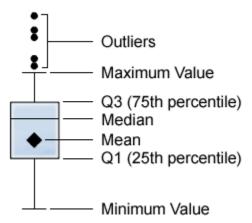
Figure 20.3 on page 169 shows a diagram of a box plot. The bottom and top edges of the box indicate the interquartile range (IQR). That is, the range of values that are between the first and third quartiles (the 25th and 75th percentiles). The marker inside the box indicates the mean value. The line inside the box indicates the median value.

Then the whiskers (lines protruding from the box) indicate the range of values that are outside of the interquartile range but are close enough not to be considered outliers.

Note: If you disable the **Show Outliers** property, then the whiskers extend to the maximum and minimum values in the plot.

Outliers are data points whose distance from the interquartile range is greater than 1.5 times the size of the interquartile range.

If there are a large number of outlier points, then the range of outlier values is represented by a bar. The data tip for the bar displays additional information about the outlier points. To explore the outlier values, double-click the outlier bar to view the values as a new histogram visualization.

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 titles**.

Generate graph titles
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 outlier points outside of the whiskers. Outliers are data points whose distance from the interquartile range is 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.

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 values of the category 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.

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 the cell colors represents the frequency of each intersection of values. If you assign a measure to the

Color data role, then the cell colors represent the aggregated measure value for each intersection of values.

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 titles**.

Generate graph titles
> 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** only affects 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.

Fit line
> adds a fit line to the visualization. For details about the fit types that are available, see "Fit Line" on page 207.

Note: Fit lines are not available if a category is assigned to the heat map.

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** data role, then the cell color indicates frequency.

Note: You can drill down on a hierarchy only if it is assigned to the X axis.

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 .

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-splines. 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 ▼ from the visualization toolbar and then select **Fit Line** ▶ **[*fit-type*]**. For details about the fit types that are available, see "Fit Line" on page 207.

Working with Geo Maps

About Geo Maps

A geo map displays your data as a bubble plot that is overlaid on a geographic map. Each bubble is located either at the center of geographic region or at the coordinates of a location.

In order to display a geo map, you must define one or more of your categories as a geography data item. For more information, see "Define a Geography Data Item" on page 128.

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 titles**.

Generate graph titles
    specifies that the graph title is generated automatically based on the data items in the visualization.

Frequency
    specifies whether the frequency values display the number of values (**Count**) or the percentage of values (**Percent**).

Show map controller
    specifies whether the map zoom and pan controls are visible.

Color gradient
    selects the gradient colors for the visualization.

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
    specifies the measure that determines the marker size.

In addition to the basic data roles, you can assign additional roles:

Bubble color
    specifies a measure that determines the bubble color.

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.

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 in the treemap represents a category value or a hierarchy node. The size of each tile can represent either the frequency count or the value of a measure. If you assign a measure to the **Color** data role, then the color of each tile represents the value of that measure.

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 titles**.

Generate graph titles
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 represent the number of values (**Count**) or the percentage of values (**Percent**).

Arrangement
specifies how the layout of the tiles in the treemap. Select one of the following values:

Standard

attempts to arrange 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.

Data Roles for a Treemap

The basic data roles for a treemap are:

Tile

specifies categories or a hierarchy that is 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 result 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.

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.

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 titles**.

Generate graph titles
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.

Data Roles for a Correlation Matrix

The basic data role for a correlation matrix is the measures. You must assign at least two measures.

Note: The maximum number of measures is 60.

21

Working with Visualization Filters

About Filters

For all visualization types, you can subset your data by using the **Filters** tab in the right pane. 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 visual exploration. For more information, see "Set the Scope of a Filter (Global or Local)" on page 189.

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.

You can perform the following tasks to manage your filters:

- "Create a Filter"
- "Create a Filter from a Data Selection"
- "Edit a Filter By Using Visual Filter Controls"
- "Edit a Filter By Using the Edit Filter Window"
- "Set the Scope of a Filter (Global or Local)"
- "Delete a Filter"
- "Reset a Filter"

Create a Filter

To create a filter:

1   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 the data item and select **Add as local filter**, or drag the data item onto the **Local Filters** area of the **Filters** tab.

   For a global filter that is applied to all of your visualizations, either right-click the data item and select **Add as global filter** or drag the data item to the **Global Filters** area of the **Filters** tab.

3   Set the parameters for your filter. You can either create a basic filter by using visual controls or 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 Filter from a Data Selection

When you select data values in a visualization, you can use the selected values to create a new 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.
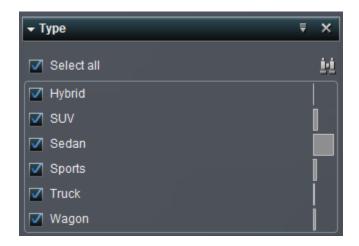
Edit a Filter By Using Visual Filter Controls

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 21.1   *A Visual Filter for Discrete Data*

Note: If your data contains a large number of discrete values, then visual filtering is not available. Use the Edit Filter window instead.

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 of each value.

Use the check boxes next to each value to select the values for the filter. You can search for a value by clicking ▥. The search locates values that begin with your search string.
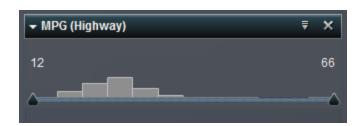
Click **Select All** to select or deselect all of the values. To invert your selection, open the ▼ drop-down list and then select **Invert Selection**.

To exclude observations where the filter category has a missing value, deselect **Include missing values**.

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.

Filtering Continuous Data

Display 21.2   *A Visual Filter for Continuous Data*

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 upper and lower limits of the data range that is selected by the filter.

You can also enter the limit values explicitly:

■   For numeric data, click a limit values and then enter the value in the text field.

■   For datetime data, click ▦ and then select a date or time.

To exclude observations where the filter measure has a missing value, deselect **Include missing values**.

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.

Edit a Filter By Using the Edit Filter Window

To edit a filter by using the Edit Filter window:

1   On the **Filters** tab, open the ▼ drop-down list for the filter that you want to edit and then select **Edit Filter**. The Edit Filter window appears.

2 Build the expression for your filter by dragging conditions and operators onto the expression in the right pane.

Table 21.1   Conditions for Character Data

=	Specifies that a matching value must match one of the filter values exactly.
Begins with	Specifies that a matching value must contain the filter value at the start of the value.
Contains	Specifies that a matching value must contain the filter value.
Ends with	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.
Is missing	Specifies that a missing value matches the filter.
Does not contain	Specifies that a matching value must not contain the filter value.
Not In	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.
Is not missing	Specifies that any nonmissing value matches the filter.

Table 21.2   Conditions for Continuous Numeric Data

Less than	Specifies that a matching value must be less than the filter value.
Greater than	Specifies that a matching value must be greater than the filter value.
Between	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.

| **Is missing** | Specifies that a missing value matches the filter. |
| **Is not missing** | Specifies that any nonmissing value matches the filter. |

Table 21.3   *Conditions for Continuous Date and Time Data*

Not equal to	Specifies that a matching value must not be equal to the filter value.
Equal to	Specifies that a matching value must be equal to filter value.
Less than	Specifies that a matching value must be less than the filter value.
Less than or equal to	Specifies that a matching value must be less than or equal to the filter value.
Greater than	Specifies that a matching value must be greater than the filter value.
Greater than or equal to	Specifies that a matching value must be greater than or equal to the filter value.
Between	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.
Is missing	Specifies that a missing value matches the filter.
Is not missing	Specifies that any nonmissing value matches the filter.

Table 21.4   *Conditions for Discrete Numeric Data and Date and Time Data*

Not equal to	Specifies that a matching value must not be equal to the filter value.
Equal to	Specifies that a matching value must be equal to filter value.
Less than	Specifies that a matching value must be less than the filter value.

Less than or equal to	Specifies that a matching value must be less than or equal to the filter value.
Greater than	Specifies that a matching value must be greater than the filter value.
Greater than or equal to	Specifies that a matching value must be greater than or equal to the filter value.
Between	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.
Not In	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.
Is missing	Specifies that a missing value matches the filter.
Is not missing	Specifies that any nonmissing value matches the filter.

Table 21.5 Operators for All Data Types

AND	Joins two or more conditions and specifies that a matching value must match all of the conditions.
NOT	Specifies that a matching value must not match the expression that follows NOT.
OR	Joins two or more conditions and specifies that a matching value must match at least one of the conditions.

To add additional conditions to an AND or OR operator, drag a new condition onto the drop zone for the operator. For example, to add a third condition to an AND operator, drag the new condition onto the word AND in the filter expression.

3   (Optional) For discrete data, click **Preview Results** to view a preview of the filter results.

4   For discrete data, select whether to display the visual filter controls on the **Filters** tab for this filter (**Show results in a list of check boxes**) or to show a text description of the filter syntax (**Apply results as a dynamic rule**.)

   Note: If your data contains a large number of discrete values, then visual filtering is not available.

5   When you are finished editing your filter, click **OK** to apply the filter.

Set the Scope of a Filter (Global or Local)

Your filters can be either local (specific to the current visualization) or global (applied to all visualizations in the visual exploration.)

To set the scope of a filter, open 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 your visualizations.

Note: A global filter replaces any local filters that use the same data item. The local filters are replaced in all of your visualizations.

Delete a Filter

To delete a filter, click ✕ on the filter on the **Filters** tab.

Note: If you remove a global filter, then the filter is removed from all of your visualizations.

Reset a Filter

To restore a filter to its initial state, open the ▼ drop-down list for the filter on the **Filters** tab and then select **Reset Filter**.

22

Exporting Data

Overview of Exporting Data

You can export your data and your visualizations from SAS Visual Analytics Explorer by using the following tasks:

- "Export a Visual Exploration as a Report"

- "Export a Visual Exploration as a PDF"

- "Save a Visualization as an Image File"

- "E-mail a Visual Exploration as a Link"

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.

To export your current visual exploration as a report:

1   Select **File ▶ Export ▶ Exploration as 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 your report:

   - ■ visualizations that do not contain any data

   - ■ animated bubble plots

   - ■ treemaps that display additional levels

   Note: If your scatter plot or heat map contains correlation data, the correlations are not included in the report.

   Click **OK**.

3   Select the location where you want to save the report and then click **Save**.

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 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: Visualizations that do not contain data are not available to include in your PDF output.

   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 a visualization as an image file:

1 Select the visualization that you want to save as an image.

2 If your visualization contains hierarchies, drill to the hierarchy levels that you want to show in the image.

3   Select ⊽ from the visualization toolbar, and then select **Export image**.

4   If necessary, scroll 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.

E-mail a Visual Exploration as a Link

To e-mail a link to your visual exploration:

1   Select **File** ▶ **Email**. The Email 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 also 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 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.

23

Managing Hierarchies

Overview of Hierarchies in SAS Visual Analytics Explorer

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, 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.

Hierarchy Management Tasks

In SAS Visual Analytics Explorer, you can perform the following tasks to manage hierarchies:

- ■ "Create a New Hierarchy"
- ■ "Edit a Hierarchy"
- ■ "Delete a Hierarchy"

Create a New Hierarchy

To create a new hierarchy:

1   Select **Data** ▶ **New Hierarchy**. The Create 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 also create a hierarchy from within a crosstab visualization. See "Create Hierarchies from a Crosstab" on page 162.

Edit a Hierarchy

To edit an existing hierarchy:

1 From the Data Items window, right-click 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 ⬇ 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 Items window, right-click the hierarchy that you want to delete and then select **Delete**. Click **Yes** to confirm that you want to delete the hierarchy.

24

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 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 Visual Analytics home page. See "Managing Your Content on the Home Page" on page 15.

Exploration Comments

About Exploration Comments

SAS Visual Analytics Explorer enables you to create and share comments about your explorations. The comments can be viewed in SAS Visual Analytics Explorer and also on the SAS Visual Analytics home page.

In addition to exploration comments, you can create comments for specific visualizations. See "Manage Visualization Comments" on page 146.

View Exploration Comments

To view the comments for the current exploration, select **File ▶ Exploration Comments**. The 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 New Exploration Comment

Create a Comment under an Existing Topic

To add a comment under an existing topic:

1 Enter your comment text in the **Enter a comment** field that is beneath the topic that you want to contribute to.

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 add 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 an Exploration Comment

Note: You must have the capability to edit comments.

To edit a comment, select the comment that you want to edit and then click ⬉.

Delete an Exploration Comment

Note: You must have the capability to delete comments.

To delete a comment, select the comment that you want to delete and then click 🗑.

25

Performing Data Analysis

Overview of Data Analysis in SAS Visual Explorer

Types of Data Analysis

SAS Visual Analytics enables you to perform three basic types of data analysis:

Correlation
    identifies the degree of statistical correlation 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-splines.

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. Values that are close to –1 imply a strong negative correlation, values that are close to 0 imply little or no correlation, and values that are close to 1 imply a strong positive correlation.

To apply correlation to a visualization, add a linear fit 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 Line

A fit line plots a model of the relationship between two 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 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 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 fit line, select ⓘ from the visualization legend.

PSpline

> creates a penalized B-spline fit. 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 fit line, select ⓘ from the visualization legend.

Forecasting

Forecasting predicts future values for your data based on the statistical trends in your data. Forecasting is available only for line charts that contain date or time data items.

A forecast adds a line with a predicted values to your visualization, as well as a colored band that represents the 95% confidence interval. The 95% confidence interval is the data range where the forecasting model is 95% confident that the future values will be.

SAS Visual Analytics 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)

Add a Fit Line to an Existing Visualization

To add a fit line to a scatter plot or heat map, select ▼ from the visualization toolbar and then select **Fit Line** ▶ [*fit-type*]. For details about the fit types that are available, see .

Add Forecasting to an Existing Visualization

To add data analysis to a line chart, select ▼ from the visualization toolbar and then select **Show Forecast**.

Note: In order for you to apply forecasting, the line chart must contain a date, time, 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.

Correlate Measures as a New Visualization

To create a new visualization that correlates multiple measures together, select **Analysis ▶ Correlate Measures**. The Correlate Measures window appears. Select the measures that you want to include in the new visualization, and then click **OK** to create the new visualization.

If you selected exactly two measures, then the new visualization is either a scatter plot or a heat map (for high-cardinality data). Correlation and a linear fit line are automatically applied to the visualization.

If you selected three or more measures, then the new visualization is either a scatter plot matrix or a correlation matrix (for high-cardinality data). Correlation and a linear fit line are automatically applied to the visualization.

Forecast Measures as a New Visualization

To create a new visualization that forecasts multiple measures:

1   Select **Analysis** ▶ **Forecast**. The Forecast window appears.

2   From the **Date Category** drop-down list, select the date, time, or datetime data item for the forecast.

3   From the **Measures** selection list, select the measures that you want to include in the forecast.

4   In the **Forecast Duration** field, enter the number of data values that you want to predict for the forecast.

5   Click **OK** to create the new visualization. The visualization is created as a new line chart.

Part 4

Designing Reports

26

Overview of SAS Visual Analytics Designer

About SAS Visual Analytics Designer

SAS Visual Analytics Designer enables users to easily create reports or dashboards that can be saved and viewed on either a mobile device or in the SAS Visual Analytics Viewer. SAS Visual Analytics Designer is part of the SAS Visual Analytics product that enables someone with either the Analysis or 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 authors 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 SAS Visual Analytics Designer, which runs in a web browser. You do not need to understand a programming language.

Report authors can easily create reports and dashboards based on data sources that have been provided by an administrator. They can also update reports that were created from visual explorations. Report authors can also create reports by importing objects or visual explorations from other reports. They can also define interactions (either filtering or brushing) for report objects and include SAS analytical results into a single report.

Your First Look at SAS Visual Analytics Designer

When you are designing reports in SAS Visual Analytics, you see SAS Visual Analytics Designer.

Here are the features of SAS Visual Analytics Designer:

Figure 26.1   *SAS Visual Analytics Designer*

1   The application bar enables you to return to the SAS Visual Analytics home page and access your recently created or viewed reports. You can also access recently used visual explorations, stored processes, or visual data builder queries in your recent history. When you select a visual exploration, the application takes you to SAS Visual Analytics Explorer. When you select a stored process, the application takes you to SAS Visual Analytics Viewer. When you select a visual data builder query, the application takes you to SAS Visual Data Builder.

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 SAS Visual Analytics 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 your reports. You can click ▣ to hide the left and right panes. Or, you can click ▦ to display the left and right panes. You can access the report views clicking ⛶, ▣, or ▭ on the menu bar.

4 The **Objects** and **Imports** tabs are in the left pane.

   The **Objects** tab provides a list of the tables, graphs, and gauges that can be used in a report or dashboard.

   The **Imports** tab provides a list of reports and report objects that have been created in SAS Visual Analytics Designer or exported from SAS Visual Analytics Explorer. This enables you to create reports from multiple 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.

5 The **Data** and **Shared Rules** tabs are in the left pane.

   The **Data** tab enables you to select a data source (or data sources) and the data items for your report. You can add, refresh, 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, or show or hide data items. You can also check the details for the measures in the data set. You can also

   For more information, see "Using Data Items in a Report" on page 292.

   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.

6 The data item table provides information about a selected data item, including the name, role, format, and aggregation. These data item properties can also be modified, which impacts all of the report objects that use the data item. For more information, see "Modify Data Item Properties" on page 300.

7 The canvas is the workspace for building a report. You can modify the canvas using the report view icons on the toolbar. For more information, see "Creating a New Report" on page 228.

A report can have multiple sections, which can be accessed by using the tabs at the top of the canvas. For more information, see "Maintaining Multi-Section Reports" on page 357.

There is also an area at the top of the canvas where you can drop filter controls and then categories to create section prompts.

8   The tabs in the right pane enable you to work with details about the report and report objects. The **Properties**, **Styles**, and **Display Rules** tabs are displayed by default. You can choose to display icons instead of tab names. To display the icons on the tabs, click ▼ after the last tab name and select **Show icons only**. You can also choose which tabs are displayed.

The **Properties** tab lists the properties for the currently selected report or report object. The report's title and description are listed. The object's name, title, and description are listed. Information that is specific to the report object type, such as the axes and legend are also included. If a report object is not selected in the report canvas, then the properties are for the section and the report.

The **Styles** tab enables you to specify the data styling, frame styling, text styling, and data colors for a selected report object.

The **Display Rules** tab enables you to populate or add intervals for an object that is currently selected in the report canvas. You can also edit or delete an existing display rule for the selected table, graph, or gauge. For more information, see "Overview of Display Rules" on page 323.

The **Comments** tab enables you to add comments to the report after you save it. For more information, see "Adding Comments to Reports" on page 282.

The **Filters** tab enables you to add a filter (or filters) to the selected report object. For more information, see "About Report Filters" on page 339.

The **Interactions** tab enables you to add or update filter or brush interactions to the report objects in a section. For more information, see "Creating Interactions" on page 346.

The **Roles** tab enables you to add or update data role assignments when you have a report object with a data source selected.

Open Window

About the Open Window

Click 📁 on the menu bar or select **File** ▸ **Open** to display the Open window.

The Open window enables you to search for and open saved reports. However, the Open window also enables you to create new folders and perform many other tasks on selected reports and folders.

Display 26.1   Open Window

Working with Folders in the Open Window

You can work with folders in the Open window. The icons for folders are located to the right of the drop-down list of folder names, as shown in Display 26.1 on page 218.

Icon	Action
←	Navigates back one level.

Icon	Action
	Moves up one level.
	Refreshes the items in the current location.
	Deletes the selected item.
	Creates a new folder.
	Specifies either details or a list.
	Opens the Search window.

You can also create a new folder when you copy a report, move a report, or save a report.

Search Window

The Search window in SAS Visual Analytics Designer enables you to navigate to find the report or exploration that you need. For example, you might need to navigate through multiple folders to find a report. In the Open window, click to open the Search window.

Display 26.2   Search Window

To search for files:

1   In the **Name** field, type text for which you want to search.

2   (Optional) To include descriptions in your search, select the **Include description and keywords** check box.

3   Select a **Location**. If you want to search multiple folders, make sure that **Search all subfolders** check box is selected.

4   (Optional) To limit your search by time, select a **Date range**. You can also specify either **Date created** or **Date modified**.

5   Click **Search**.

Note: After a search, the list of reports also includes the path information for each located report.

27

Specifying Preferences for SAS Visual Analytics Designer

Specifying Global and General Preferences

To specify global SAS preferences, see "Specifying Global Preferences" on page 17. To specify general preferences, see "Specifying the SAS Visual Analytics General Preferences" on page 18.

Specifying Your Preferences for SAS Visual Analytics Designer

To specify preferences that are specific to SAS Visual Analytics Designer:

1   Select **File** ▶ **Preferences** to open the Preferences window.

2   Select **Visual Analytics Designer** ▶ **General**.

Display 27.1   *Preference Dialog Box for SAS Visual Analytics Designer*

3   Select **Full screen**, **Tablet** for mobile devices, or **Wide-screen tablet**. **Full screen** is the default.

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

Note: To return to the default setting for SAS Visual Analytics Designer, select **Reset to Defaults**.

28

Opening Reports

Opening Reports in SAS Visual Analytics Designer

Reports in SAS Visual Analytics Designer are saved when you use either the Save 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 one or more queries. The section displays the results with one or more report objects (for example, a table, a graph, a gauge, and so on).

Note that a section is not required to contain any report objects. For example, you might have a report that you use to create other reports. Your report might contain data sources, calculated items, global data filters, and shared display rules, but no report objects.

To open a saved report:
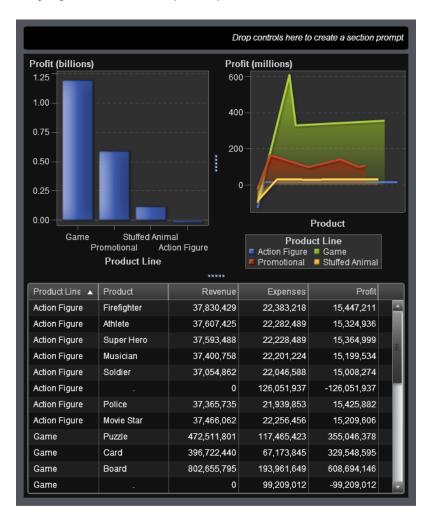
■ On the SAS Visual Analytics home page, select a report. In the details view, click **Edit**. Or you can double-click the report to open it in SAS Visual Analytics Designer.

■ In SAS Visual Analytics Designer:

□ Select **File ▶ Recent** and then select a report name.

- ☐ Click ▾ beside the recent content drop-down list, and then select a report name.

- ☐ Select **File** ▶ **Open** to display the Open window.

- ☐ Click 📁 to navigate to a report and then select it.

An Example Report

The following display shows a report in SAS Visual Analytics Designer. The report contains a bar chart, a line chart, and a list table.

Display 28.1   *An Example Report*

29

Creating and Saving Reports

About Reports

You can drag and drop tables, graphs, and gauges to create a well-designed report in SAS Visual Analytics Designer. You can also add text, images, and other controls to reports. Reports can contain multiple sections.

When you design a report, keep in mind that it might look slightly different on a mobile device or in the SAS Visual Analytics 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 SAS Visual Analytics Designer than it does in SAS Visual Analytics Viewer or on a mobile device. For more information about treemaps, see "About Treemaps" on page 252.

Reports that are imported from SAS Visual Analytics Explorer might look slightly different in SAS Visual Analytics Designer.

SAS Visual Analytics Designer provides a report view that enables you to change the size of the canvas that you use for laying out reports. There are three report views: **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 menu bar.

Two report layouts are available in SAS Visual Analytics Designer: **Precision** and **Tiled**. For more information, see "Choosing a Report Layout" on page 230.

Creating a New Report

To create a new report:

1   Choose your data source and data items. For information about selecting data, see "Selecting Data for Reports" on page 288.

2   Select the layout (precision or tiled) for the report. For more information, see "Choosing a Report Layout" on page 230.

3   Select the report object (or objects) that you want to use in the report. For information about selecting report objects, see "Using Tables, Graphs, Gauges, Controls, and Other Report Objects" on page 238.

4   Drag and drop the report object (or objects) on the canvas. Alternatively, you can double-click the report object on the **Objects** tab or tab over the report objects and press Enter to add the report object to the canvas.

5   Update the properties for the report and the report objects.

6   Update the styles for the report objects.

7   (Optional) Create or modify display rules. For more information, see "Adding Display Rules to a Gauge" on page 335.

8   (Optional) Add filters to the report. For more information, see "About Report Filters" on page 339.

9   (Optional) Add interactions between the report objects in a section. For more information, see "Overview of Interactions" on page 345.

10 (Optional) Add a new section (or sections) to the report. For more information, see "Overview of Report Sections" on page 357.

11 Save your report.

12 (Optional) Add comments to your saved report.

You can also create a new report based on an existing report or on 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 232.

Specify Report Properties

To view or update report properties:

1   In the right pane, click the **Properties** tab.

2   Click ▼ to open the object inventory list and then select the report name.

    Note: You can also use the object inventory list to navigate to different sections or report objects.

Here is an example of an object inventory list for a report:

Display 29.1   Object Inventory in the Properties Tab

3   Update the **Title** and the **Description** for the report.

Here is an example of the title and the description for a report:

Display 29.2   Report Properties

Choosing a Report Layout

These report layouts are available in SAS Visual Analytics Designer:

Precision
enables you to place, align, and size report objects. The precision layout allows objects to overlap and 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.

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 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**. Tiled layout is the default.

Here is an example of the **Precision** option on the **Properties** tab:

Display 29.3   *Properties Tab for Precision Layout*

4   (Optional) If you selected **Precision**, then you can specify **Fit to screen**. This option prevents objects from being sized too wide or too tall, which can cause the report viewing area for the section to scroll.

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 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 SAS Visual Analytics Designer or from a report that has been exported from SAS Visual Analytics Explorer. (For example, you might want to reuse calculated data items or hierarchies from a previous report in your current report.) You can save the report under a new name, or you can edit the report to make some customizations and then save it.

Note: You cannot import box plots, heat maps, histograms, or correlation matrices from SAS Visual Analytics Explorer.

To base a new report on report objects imported from one or more existing reports:

1 Open or import the existing report by doing one of the following:

- On the SAS Visual Analytics home page, double-click the report that you want to open.

- In SAS Visual Analytics Designer:

  1 Select **File ▶ Recent** and then select the report name.

  2 Click ▼ beside the recent content drop-down list, and then select a report name.

  3 Select **File ▶ Open** to display the Open window.

  4 Click 📁 to navigate to a report and then select it.

  5 Click **Imports** in the left pane. Select a report name or click **Import another report** to display the Open window, where you can select a report.

Display 29.4 *Imports Tab with a Report List*

In the tree view on the left, a list of sections and report objects appears. Drag a single report object, multiple report objects, or entire section that you want to include in the new report and drop it onto the canvas.

Display 29.5 *Imports Tab with a List of Report Objects*

> **Note:** Report objects that cannot be imported from SAS Visual Analytics Explorer are not displayed on the **Imports** tab.
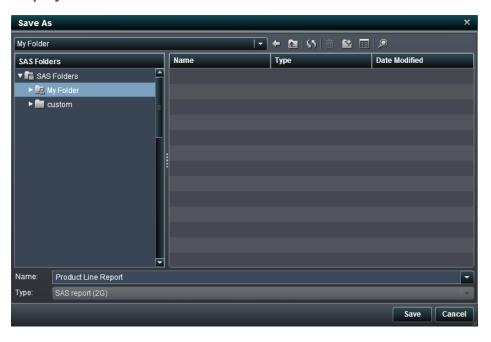
2 (Optional) Modify the report objects, properties, styles, roles, filters, and display rules.

3 Select **File ▸ Save** or click 🖫, which displays the Save As window. Enter a **Name**. Report names cannot use these characters: / \

4 Click **Save**.

Saving a Report

To save a report:

1 Select **File ▸ Save** or click 🖫. If you are saving a new report, then the Save As window is displayed.

2 For new reports, enter a **Name**. Report names cannot use these characters: / \

   If you are saving an existing report, then the name of that report is listed here. You can either leave the name as it is and overwrite the existing report, or you can change the name to create a new report.

Display 29.6   *Save As Window*

3   Click **Save**.

30

Displaying Results Using Report Objects

Using Tables, Graphs, Gauges, Controls, and Other Report Objects

After selecting your data source and data items, add one or more report objects to display the results. SAS Visual Analytics Designer provides report objects for all of your reports. Report objects in SAS Visual Analytics Designer are grouped into these types on the **Objects** tab in the left pane and the **Insert** menu:

- tables

- graphs

- gauges

- controls

- other

The following table types are available in SAS Visual Analytics Designer:

- list tables
- crosstabs

The following graph types are available in SAS Visual Analytics Designer:

- bar charts
- targeted bar charts
- waterfall charts
- line charts
- pie charts
- scatter plots
- time series plots
- bubble plots
- treemaps
- dual axis bar charts
- dual axis line charts
- dual axis bar-line charts
- dual axis time series plot

The following gauges types are available in SAS Visual Analytics Designer:

- bullet
- slider
- thermometer
- dial
- speedometer

The following control types are available in SAS Visual Analytics Designer:

- drop-down lists

- lists

- button bars

- text input fields

- range sliders

These other objects are also available in SAS Visual Analytics Designer:

- text

- image

- stored processes

- vertical containers

- horizontal containers

- geo maps

In SAS Visual Analytics Designer, you have access to report objects from SAS Visual Analytics Explorer. You can open the histogram, heat map, box plot, or correlation matrix report objects that have been exported from SAS Visual Analytics Explorer. However, you cannot create new histograms, heat maps, box plots, or correlation matrices in SAS Visual Analytics Designer.

Using Tables to Display Results

Overview of the Table Types

About 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.

Display 30.1 *A List Table*

Make ▲	MSRP	Invoice	MPG (City)	MPG (Highway)
Acura	$42,939	$38,591	19.42857	26.14286
Audi	$43,308	$39,330	18.47368	25.78947
BMW	$43,285	$39,621	18.7	27
Buick	$30,538	$27,855	18.88889	27.77778
Cadillac	$50,474	$46,427	16.5	23
Chevrolet	$26,587	$24,061	19.66667	26.85185
Chrysler	$27,252	$25,270	19.86667	27.33333
Dodge	$26,254	$24,160	19.38462	26.23077
Ford	$24,016	$21,953	19.26087	25.73913
GMC	$29,561	$26,289	15.375	19.625
Honda	$21,435	$19,513	27.82353	34
Hummer	$49,995	$45,815	10	12
Hyundai	$17,477	$16,035	23	29.91667

About Crosstabs

A *crosstab* (also known as a crosstabulation table) shows an aggregate metric for the intersections of two or more categories. In a crosstab, a category is typically displayed on both the columns and the rows, and each cell value represents the aggregated measure from the intersection of the categories on the specific row and column. This type of table uses less space than a list table and is easier to read because data is grouped both horizontally and vertically.

Display 30.2   A Crosstab

Facility State	Facility City	2009 Expenses	2009 Revenue	2010 Expenses	2010 Revenue	2011 Expenses	2011 Revenue
AL	Birmingham	2,436,889	1,384,807	2,750,494	1,343,126	1,582,327	4,135,931
	Mobile	2,315,153	1,710,704	2,519,664	1,568,478	2,842,126	1,355,372
AR	Little Rock	3,654,377	4,633,210	3,602,950	15,981,013	4,134,232	20,087,152
AZ	Scottsdale	3,035,051	9,415,121	3,066,971	7,907,419	3,184,412	7,109,553
CA	Los Angeles	422,102	0	573,229	0	875,769	0
	San Diego	1,889,312	1,581,599	2,092,113	1,437,426	2,402,617	1,169,943
	San Francisco	8,854,433	29,325,820	9,534,242	26,878,679	9,516,337	24,044,522
IL	Chicago	2,241,862	6,179,655	2,555,291	8,334,320	3,011,832	10,518,963
	Elgin	967,030	138,368	950,541	141,971	1,001,233	108,368
	Joliet	1,404,595	749,721	1,448,370	1,028,966	1,604,408	1,288,052
LA	New Orleans	2,910,253	681,021	1,584,641	2,699,708	1,550,269	3,045,644
NH	Manchester	2,188,452	1,129,351	2,376,848	1,063,952	2,720,423	930,849
NJ	Newark	.	.	1,512,857	4,555,827	1,894,203	5,428,525
NY	New York	1,794,957	2,313,319	1,706,994	1,383,978	1,847,918	1,110,836
OH	Cincinnati	1,368,968	1,373,531	1,367,600	1,938,936	1,527,279	2,384,687
	Cleveland	1,646,547	1,859,595	1,657,984	1,614,283	1,852,647	1,307,465
PA	Philadelphia	1,126,956	134,938	1,404,364	107,554	1,537,728	84,232
TX	Corpus Christi	2,073,763	1,895,837	2,249,271	1,737,935	2,510,895	1,438,103
	Dallas	3,337,364	9,153,994	3,560,555	8,522,095	3,761,182	6,640,495
	Houston	2,293,978	4,660,600	2,339,839	3,959,864	2,386,605	3,113,888
WA	Seattle	3,392,121	8,050,558	3,397,512	17,035,598	3,813,872	20,223,684

Note: Frequency is displayed only when there are no measures in the crosstab.

Insert a Table into a Report

To insert a table into a report:

■ Drag the table icon from the **Objects** tab in the left pane and drop it onto the report canvas.

■ Select **Insert ▶ Tables** and then select the menu item for the table object that you want to insert. The table is automatically placed in the report canvas. If you want the table to appear in a different place, then drag and drop it in a new location.

The following table lists the available table objects:

Icon	Table Type
⊞	List Table
⊞	Crosstab

Specify Table Properties

To specify the properties for tables:

1   If not already selected, select the table in the report canvas that you want to update.

2   In the right pane, click the **Properties** tab.

3   Update the general properties for the table. Your choices are **Name**, **Title**, and **Description**.

4   Update the object-specific properties for the table. For a list table, you can specify that you want detail data and totals. For a crosstab, you can specify that you want an indented layout; totals and subtotals for rows, columns, or both; as well as the placement of the totals and subtotals.

Here is an example of the properties for crosstabs:

Display 30.3   Properties for a Crosstab

Specify Table Styles

To specify styles for tables:

1   If not already selected, select the table in the report canvas that you want to update.

2   In the right pane, click the **Styles** tab.

3   Update the styles for the table. For a list table, you can customize the colors of **Cells** and **Column headers**. For crosstabs, you can customize the colors of **Cells**, **Row headers**, **Column headers**, and **Measure headers**.

Here is an example of the styles for crosstabs:

Display 30.4   Styles for a Crosstab

Add Sparklines to a List Table

A *sparkline* is a small line graph that presents a single trend over time. A sparkline fits in a single cell and 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 memorable. In SAS Visual Analytics Designer, you can add sparklines to a column in a list table.

To add sparklines:

1 If not already selected, select the list table in the report canvas that you want to update.

2 Right-click the list table and then select **Add Sparkline**. The Add Sparkline window is displayed.

Display 30.5   Add Sparkline Window

3 Enter a **Column label**.

4 Select a **Time Axis**.

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**.

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 **Add Sparkline**. The Add 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

Overview of the Graph Types

About Bar Charts

A *bar chart* consists of a grid and some vertical or horizontal bars. Each bar represents quantitative data.
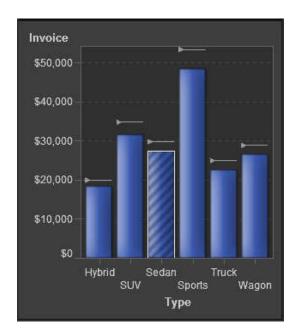
Display 30.6   *A Bar Chart*

About Targeted Bar Charts

A *targeted bar chart* is a variation of the bar chart that has pointers to target values. In , the pointers appear above each bar.
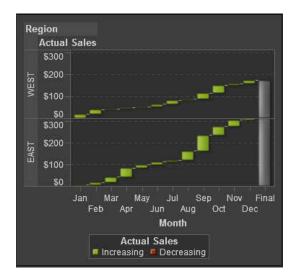
Display 30.7   A Targeted Bar Chart

About Waterfall Charts

A *waterfall chart* shows how the initial value of a measure increases or decreases during a series of operations or transactions. (This chart is also called a progressive bar chart.) 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 indicates the magnitude and type (positive or negative, for example) of the operation or transaction. The resulting chart is a stepped cascade that shows how the transactions or operations lead to the final value of the measure.
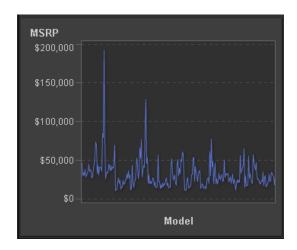
Display 30.8 *A Waterfall Chart*

About Line Charts

A *line chart* shows the relationship of one variable to another, often as movements or trends in the data over a period of time. Line graphs summarize source data and typically are used to chart response values against discrete categorical values.
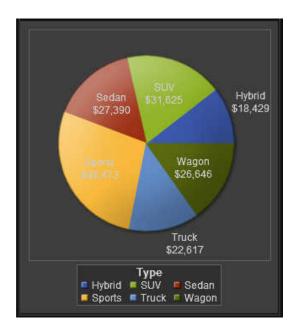
Display 30.9 *A Line Chart*

About Pie Charts

A *pie chart* is a circular chart that is divided into slices by radial lines. Each slice represents the relative contribution of each part to the whole.
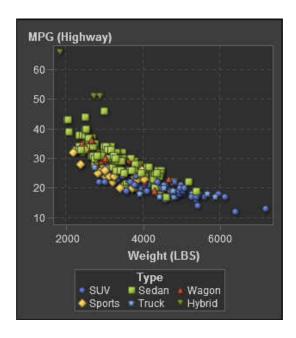
Display 30.10   A Pie Chart

About Scatter Plots

A *scatter plot* is a two-dimensional plot that shows the joint variation of two data items. In a scatter plot, each marker (represented by symbols such as dots, squares, and plus signs) represents an observation. The marker position indicates the value for each observation.
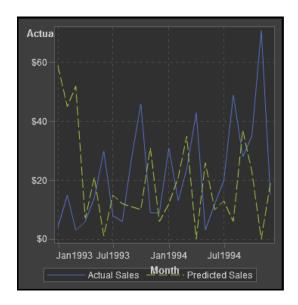
Display 30.11   *A Scatter Plot*

About Time Series Plot

A *time series plot* shows an ordered sequence of values of a discrete category variable that are observed at equally spaced time intervals.
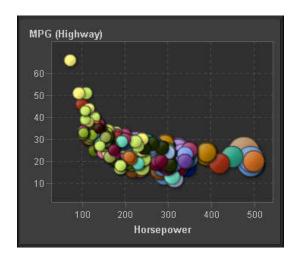
Display 30.12   *A Time Series Plot*

About Bubble Plots

A *bubble plot* is a variation of a scatter plot in which the markers are replaced with bubbles. In a bubble plot, each bubble represents an observation. The location of the bubble represents the value for the two measure axes; the size of the bubble represents the value for the third measure. A bubble plot is useful for data sets with dozens to hundreds of values or when the values differ by several orders of magnitude. You can also use a bubble plot when you want specific values to be visually represented by different bubble sizes.

Display 30.13   A Bubble Plot

About Treemaps

A *treemap* is divided into rectangular areas called tiles. The color of each tile represents the value of the first measure in the query. The size of each tile represents the value of the second measure in the query. For example, a treemap might be used to represent sales data where the tile sizes vary according to the number of orders invoiced and the tile colors are derived from a color gradient that represents low to high sales figures.

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 SAS Visual Analytics Designer than it does in SAS Visual Analytics Viewer or on a mobile device.
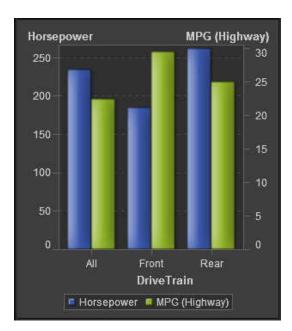
Display 30.14 *A Treemap*

About Dual Axis Bar Charts

A *dual axis bar chart* is a variation of the bar chart that has two measures, one on each axis. There can be one measure per axis for a total of two measures.
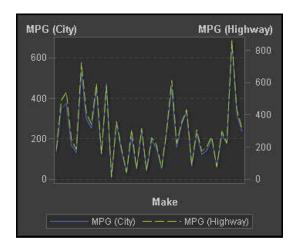
Display 30.15 *A Dual Axis Bar Chart*

About Dual Axis Line Charts

A *dual axis line chart* is a variation of the line chart that has two measures, one on each axis. There can be one measure per axis for a total of two measures.
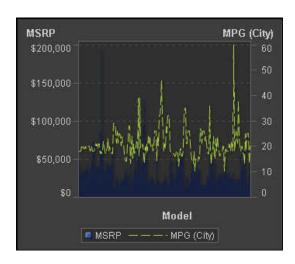
Display 30.16   *A Dual Axis Line Chart*

About Dual Axis Bar-Line Charts

A *dual axis bar-line chart* is a bar chart that has two measures, one on each axis, and is overlaid by a line chart. There can be one measure per axis for a total of two measures.

Display 30.17   *A Dual Axis Bar-Line Chart*

About Dual Axis Time Series Plots

A *dual axis time series plot* is a variation of the time series plot that has two measures, one on each axis. There can be one measure per axis for a total of two measures.

Display 30.18   *A Dual Axis Time Series Plot*

Insert a Graph into a Report

To insert a graph into a report:

- Drag the graph icon from the **Objects** tab in the left pane and drop it onto the report canvas.

- Select **Insert ▶ Graphs** and then select the menu item for the graph object that you want to insert. The graph is automatically placed in the report canvas. If you want the graph to appear in a different place, then drag and drop it in a new location.

The following table lists the available graph objects:

Icon	Graph Type
	Bar Chart
	Targeted Bar Chart

Icon	Graph Type
	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

Specify Graph Properties

To specify the properties for graphs:

1 If not already selected, select the graph in the report canvas that you want to update.

2 In the right pane, click the **Properties** tab.

3 Update the general properties for the graph. Your choices are **Name**, **Title**, and **Description**.

4　Update the specific properties for the graph. The available properties depend on the selected graph type. For example, for a pie chart, you can specify **Data Labels**, **Group Style**, and **Legend** properties.

Here is an example of the properties for a pie chart:

Display 30.19　Properties for a Pie Chart

Specify Graph Styles

To specify styles for graphs:

1   If not already selected, select the graph in the report 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, for a pie chart, you can specify **Data Styling**, **Frame Styling**, **Text Styling**, and **Data Colors** styles.

Here is an example of the styles for a pie chart:

Display 30.20   *Styles for a Pie Chart*

Using Gauges to Display Results

Overview of the Gauge Types

A gauge 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.

The following gauges are available in SAS Visual Analytics Designer:
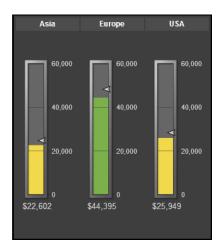
Display 30.21 *A Bullet Gauge*
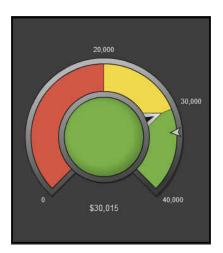
Display 30.22 *A Slider Gauge*
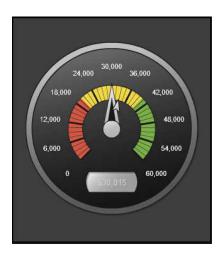
Display 30.23 *A Thermometer Gauge*

Display 30.24 *A Dial Gauge*

Display 30.25 *A Speedometer Gauge*

Inserting a Gauge into a Report

To insert a gauge into a report:

- Drag the gauge icon from the **Objects** tab in the left pane and drop it onto the report canvas.

- Select **Insert ▶ Gauges** and then select the menu item for the gauge object that you want to insert. The gauge is automatically placed in the report canvas. If you want the gauge to appear in a different place, then drag and drop it in a new location.

The following table lists the available gauge objects:

Icon	Gauge Type
	Bullet
	Slider
	Thermometer
	Dial

Icon	Gauge Type
⚙	Speedometer
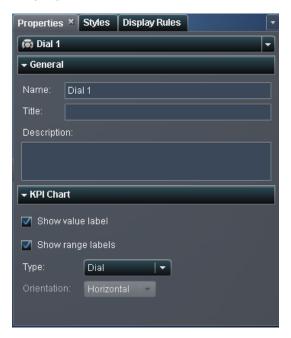
Specify Gauge Properties

To specify properties for gauges:

1 If not already selected, select the gauge in the report canvas that you want to update.

2 In the right pane, click the **Properties** tab.

3 Update the general properties for the gauge. The available properties depend on the selected gauge type. For example, for a dial gauge, you can specify **General** and **KPI Chart** properties.

Here is an example of the properties for a dial gauge:

Display 30.26   Properties for a Dial Gauge

Specify Gauge Styles

To specify styles for gauges:

1 If not already selected, select the gauge in the report 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, for a pie chart, you can specify **Data Styling**, **Frame Styling**, and **Text Styling**.

Here is an example of the styles for a dial gauge:

Display 30.27   Styles for a Dial Gauge

Using Controls to Display Results

Overview of the Control Types

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. Controls can be used in a report with interactions.

Only the drop-down list, button bar, and text input controls can be used in section prompts. If the same data source is used, then a control that is used in a section prompt automatically filters all other report objects in the same section.

The following controls are available in SAS Visual Analytics Designer:

Display 30.28 *A Drop-down List Control*

Display 30.29 *A List Control*

Display 30.30   *A Button Bar Control*

Display 30.31   *A Text Input Control*

Display 30.32   *A Range Slider Control*

Insert Controls into a Report

To insert a control into a report:

1 Choose one of the following:

  ■ Drag the control icon from the **Objects** tab in the left pane and drop it onto the report canvas.

  ■ Select **Insert ▶ Controls** and then select the menu item for the control object that you want to insert. The control is automatically placed in the report canvas. If you want the control to appear in a different place, then drag and drop it in a new location.

  The following table lists the available controls:

Icon	Control Type
🗒	Drop-down List

Icon	Control Type
	List
	Button Bar
	Text Input
	Range Slider

2   Specify the data item that you want to use in the filter control:

■   Drag and drop a category data item onto the control.

■   Use the **Roles** tab in the right pane and then specify the category as the **Filter**.

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 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 report canvas. (Look for the hint text that says, "Drop controls here to create a section prompt.") The control appears above the report 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 Sate*. 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 as the **Filter**.

If you use a control to create a section prompt, then the user can select a value (or values, if multi-selection is available) to filter the data. Note that for some control types, the user might need to press Ctrl + click to clear the value in the filter.
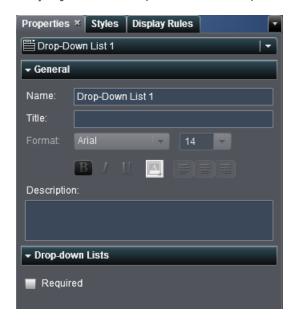
Specify Control Properties

To specify the properties for a control:

1 If not already selected, select the control in the report canvas that you want to update.

2 In the right pane, click the **Properties** tab.

3 Update the general properties for the control. Your choices are **Name**, **Title**, and **Description**.

4 Update the specific properties for the control. The available properties depend on the selected control.

Here is an example of the properties for a drop-down list control:

Display 30.33   Properties for a Drop-down List Control

Specify Control Styles

To specify styles for controls:

1   If not already selected, select the control in the report 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**.

   Here is an example of the styles for a drop-down list control:

   Display 30.34   *Styles for a Drop-down List Control*

Adding Other Objects to Reports

Overview of the Other Object Types

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. You can also use text to annotate other objects in a report. For more information, see "Insert Text into a Report" on page 271.

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. For more information, see "Insert an Image into a Report" on page 272.

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. For more information, see "Insert Stored Processes into a Report" on page 274.

About Vertical or Horizontal Container Objects

You can use a vertical or horizontal container to group other report objects. For more information, see "Insert Vertical or Horizontal Containers into a Report" on page 274.

About Geo Maps

A geo map helps you analyze data in the context of a location. It displays your data as 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.

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 map object only if the report uses data items from a data source that is enabled for geographic mapping. A geo map requires a geography variable with the role type of geography. If you change a category data item to geography, then you prompted for latitude and longitude data items. However, it does not have to be a drill-down hierarchy. For more information, see "Insert Geo Maps into a Report" on page 275.

Insert Text into a Report

To insert text into a report:

1   Choose one of the following:

   ■   Drag 🔤 from the **Objects** tab in the left pane and drop it onto the report canvas.

   ■   Select **Insert ▸ Other ▸ Text**. The text object is automatically placed in the report canvas. If you want the text to appear in a different place, then drag and drop it in a new location.

2   Double-click inside the text object in the canvas and enter the text. You can use the floating toolbar to change the font, font size, text color, and text background color. You can also specify whether the text is bold, italic, or underlined and whether it is left-aligned, centered, or right-aligned.

   You can use the context menu to cut, copy, and delete text. However, you have to use the keyboard (Ctrl +V) to paste text.

3   (Optional) Add a hyperlink by selecting some text, and then clicking 🔗. Select **Hyperlink**. The Link Setup window is displayed. Enter a **URL** and then click **OK**.

Display 30.35   Link Setup Window

The link appears in the text object.

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 report 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 report canvas. If you want the image to appear in a different place, then drag and drop it in a new location.

Display 30.36   Image Selection Window

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**.

Insert Stored Processes into a Report

To insert a stored process into a report:

1 Choose one of the following:

- Drag from the **Objects** tab in the left pane and drop it onto the report canvas. The **Open** window is displayed.

- Select **Insert ▷ Other ▷ Stored Process**. The **Open** window is displayed.

2 In the Open window, select a stored process. Click **Open**. The stored process is automatically placed in the report canvas.

Note: No styles are available for stored processes.

Insert Vertical or Horizontal Containers into a Report

To insert a vertical or horizontal container into a report:

1 Choose one of the following:

- Drag or from the **Objects** tab in the left pane and drop it onto the report canvas.

- Select either **Insert ▷ Other ▷ Vertical Container** or **Insert ▷ Other ▷ Horizontal Container**. The vertical or horizontal container is automatically placed in the report canvas. If you want the container to appear in a different place, then drag and drop it in a new location.

2   Drag and drop other report objects onto the vertical or horizontal container.

Insert Geo Maps into a Report

To insert a geo map into a report:

■   Drag 🌐 from the **Objects** tab in the left pane and drop it onto the report canvas.

■   Select **Insert ▶ Other ▶ Geo Map**. The geo map object is automatically placed in the report canvas. If you want the text to appear in a different place, then drag and drop it in a new location.
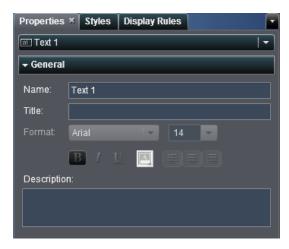
Specify Text Object Properties

To specify the properties for a text object:

1   If not already selected, select the text object in the report canvas that you want to update.

2   In the right pane, click the **Properties** tab.

3   Update the general properties for the text. Your choices are **Name**, **Title**, and **Description**.

Here is an example of the properties for a text object:

Display 30.37   Properties for a Text Object

Specify Image Properties

To specify the properties for an image:

1 If not already selected, select the image in the report canvas that you want to update.

2 In the right pane, click the **Properties** tab.

3 Update the general properties for the image. Your choices are **Name**, **Title**, and **Description**.

4 Update the properties specific for the image. Your choices are **Location**, **Scale type**, and **Tooltip text**.

Here is an example of the properties for an image:

Display 30.38   *Properties for an Image Object*
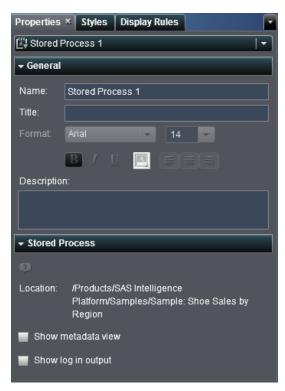
Specify Stored Process Properties

To specify the properties for a stored process:

1 If not already selected, select the stored process in the report canvas that you want to update.

2 In the right pane, click the **Properties** tab.

3 Update the general properties for the stored process. Your choices are **Name**, **Title**, and **Description**.

Here is an example of the properties for a stored process:

Display 30.39   Properties for a Stored Process

4   Update the properties specific to the stored process. Your choices are **Show metadata view** and **Show log in output**.

Selecting the **Show metadata view** check box can make it easier to work with the stored process in a report.

Here is an example of the metadata view for a stored process:

Display 30.40 *Metadata View for a Stored Process*

Selecting the **Show log in output** option means that both log output and the stored process output are displayed in the report. This can assist you in debugging any problems that might arise.

Specify Vertical or Horizontal Container Properties

To specify the properties for a vertical or horizontal container:

1   If not already selected, select the vertical or horizontal container in the report canvas that you want to update.

2   In the right pane, click the **Properties** tab.

3   Update the general properties for the vertical or horizontal container. Your choices are **Name**, **Title**, and **Description**.

Here is an example of the properties for a vertical container:

Display 30.41   Properties for a Vertical Container

4 (Optional) Update the order in which the report objects appear inside the container.

Specify Geo Map Properties

To specify the properties for a geo map:

1 If not already selected, select the geo map in the report canvas that you want to update.

2 In the right pane, click the **Properties** tab.

3 Update the general properties for the geo map. Your choices are **Name**, **Title**, and **Description**.

4   Update the properties that are specific to the geo map. Your choices are **Show map navigation control**, **Transparency**, and **Show legend**.

Here is an example of the properties for a geo map:

Display 30.42 *Properties for a Geo Map*

Specify Styles for Geo Maps

To specify styles for a geo map:

1   If not already selected, select the geo map in the report canvas that you want to update.

2   In the right pane, click the **Styles** tab.

3   Update the styles for the geo map.

Here is an example of the styles for the geo map:

Display 30.43   Styles for a Geo Map

Adding Comments to Reports
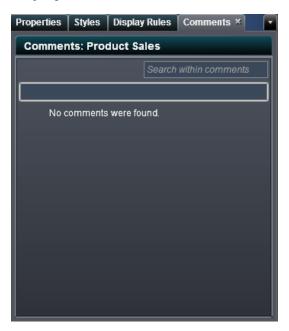
If you have the Add Comments capability, then you can add (or view) comments. Note that you must save a report before you can add comments.

To add a comment to a report:

1   Click the **Comments** tab in the right pane.

Here is an example:

Display 30.44   Comments Tab

2   Enter a topic name and a comment.

Here is an example:

Display 30.45   Adding a Comment in SAS Visual Analytics Designer

3   (Optional) Click to add 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   Click the **Comments** tab in the right pane.

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

3   (Optional) Click to add a file or image to your reply.

4   Click **Post** to add your comment.

Note: Only administrators can modify or delete an existing comment.

To search for a comment:

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

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

31

Selecting Data for Reports

Overview of Data Sources and Data Items

Data sources that are available in SAS Visual Analytics 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 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.

Each data source 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.

SAS Visual Analytics Designer can use existing user-defined formats, but you cannot specify to apply a user-defined format to a column.

Adding a Data Source or Data Sources

You can use one or more data sources for a report in SAS Visual Analytics Designer.

To add a data source for a report:
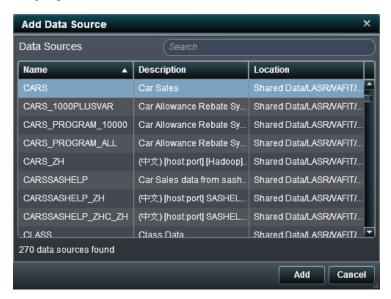
1   On the **Data** tab, click ▾ beside the **Select a data source** text to display the Add Data Source window.

Display 31.1   DataTab

2   In the Add Data Source window, select a data source.

Display 31.2   Add Data Source Window

Add Data Source		×
Data Sources	*Search*	

Name ▲	Description	Location
CARS	Car Sales	Shared Data/LASR/VAFIT/...
CARS_1000PLUSVAR	Car Allowance Rebate Sy...	Shared Data/LASR/VAFIT/...
CARS_PROGRAM_10000	Car Allowance Rebate Sy...	Shared Data/LASR/VAFIT/...
CARS_PROGRAM_ALL	Car Allowance Rebate Sy...	Shared Data/LASR/VAFIT/...
CARS_ZH	(中文) [host:port] [Hadoop]...	Shared Data/LASR/VAFIT/...
CARSSASHELP	Car Sales data from sash...	Shared Data/LASR/VAFIT/...
CARSSASHELP_ZH	(中文) [host:port] SASHEL...	Shared Data/LASR/VAFIT/...
CARSSASHELP_ZHC_ZH	(中文) [host:port] SASHEL...	Shared Data/LASR/VAFIT/...
CLASS	Class Data	Shared Data/LASR/VAFIT/

270 data sources found

Add   Cancel

3   Click **Add**. The list of available data items is displayed in the **Data** tab.

4   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.

Refreshing a Data Source for a Report

You can refresh the columns in a data source in SAS Visual Analytics 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 adds any new columns that have been added to the table metadata. The default format and names of existing columns will be refreshed the next time you open the report.

To refresh a data source for a report, on the **Data** tab, click ↻.

Removing a Data Source from a Report

You can remove a data source from a report in SAS Visual Analytics 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 **Yes** in the confirmation message that is displayed.

Display 31.3   The Remove Data Source Message

Changing 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.

Display 31.4   Change Data Source Window

3   Click **Change**. If the data sources are compatible, then the list of updated data items is displayed in the **Data** tab. Being compatible means that every data item that is defined in the original data source that is being replaced has to be defined in the new replacement data source. The data items must share the same names (which is not the same thing as the labels). The replacement data source can have additional data items. The additional data items are displayed in the **Data** tab when the data source is replaced.

If the data sources are not compatible, a message is displayed.

Using Data Items in a Report

About Data Items

Each data source in SAS Visual Analytics Designer includes one or more standard 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.

Table 31.1   *Data Items That Are Available in SAS Visual Analytics Designer*

Data Item	Icon	Description
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.
Category	▦	A data item whose distinct values are used to group and aggregate measures. There are four types of categories: alphanumeric, date, timestamp, and time. 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*. 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 also uses this category icon.
Date and Time	▦	A category data item whose distinct values are used to group and aggregate measures. There are three types of date categories: date, timestamp, and time. Examples of date, timestamp, and time categories are *Order Year*, *Date of Sale*, and *Delivery Time*.
Geography	⊕	A data item whose values are mapped to geographical locations or regions. Geography data items can be used in reports 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).

Data Item	Icon	Description
Hierarchy	品	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*.
		SAS Visual Analytics Designer assigns a default aggregation method to every measure. Almost all measures are assigned sum, but certain formats are assigned an average. You can change the aggregation method.

Note: Report objects that are imported from SAS Visual Analytics Explorer use either live or on-demand data instead of embedded data. Therefore, you can update the properties, styles, and comments for these reports in SAS Visual Analytics 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:
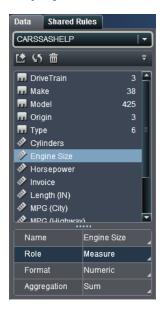
1 On the **Data** tab in the left pane, click the down arrow to display a list of available data sources. Select a data source and the **Data** tab is populated with a list of all of the data items that are in the data source.

   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 **Remove data source**. Click **Yes** in the confirmation message that is displayed.

2 (Optional) To see more information about a particular 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.

Here is an example of the details for a data item called *Engine Size*:

Display 31.5   Details about a Selected Measure Data Item

Duplicate Data Items
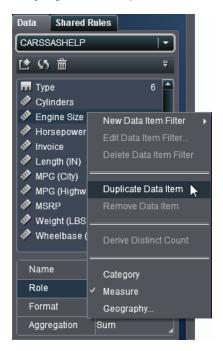
Duplicating data items in SAS Visual Analytics Designer enables you to see the aggregations of a data item (**Sum**, **Average**, **Minimum**, **Maximum**, and **Count**) side-by-side in a table. If you save a report with duplicate data items, then those data items are available when you edit the report the next time.

You can duplicate a calculated data item. You can also duplicate a category if you want to it to be displayed using a different format.

To duplicate a data item:

1 On the **Data** tab in the left pane, right-click the data item that you want to duplicate.

Display 31.6   Duplicate Data Item Menu Selection

2   Select **Duplicate Data Item**. The data item appears in the list of data items. 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).

3   (Optional) Change the format, aggregation, or both for the duplicate item (or items).

4   (Optional) Rename the duplicate data item (or data items).

Remove Data Items
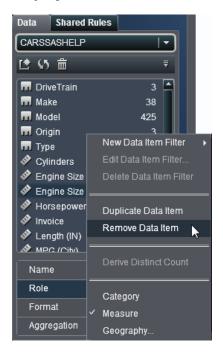
You can remove data items, including calculated data items, so that they no longer appear in the **Data** tab in the left pane.

You cannot remove the last copy of a data item or a data item that is being used by a calculated or a geography data item. You cannot remove a data item that is inside a hierarchy, if it reduces the hierarchy to a single level.

To remove a data item (or data items):
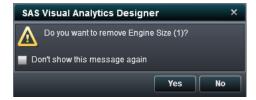
1 On the **Data** tab in the left pane, right-click the data item that you want to remove.

Display 31.7 Remove Data Item Menu Selection

2 Select **Remove Data Item**.

3 Click **Yes** in the confirmation message that is displayed. The data item is removed from the list of data items.

Display 31.8 Remove Data Items Message
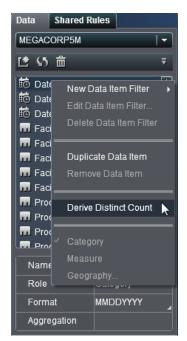
Derive a Distinct Count

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 derive a distinct count for only a categorical data item, including calculated category data items.

To derive a distinct count:

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 **Derive Distinct Count**. The distinct count data item appears in the list of data items.

Display 31.9   Derive Distinct Count Menu Selection

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 Geography 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 into a geography data item, then it automatically becomes a category data item.

Note: Calculated data items cannot be turned into geography data items.

To create a 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**.

   ■ Select the data item that you want to use for the geography data item. In the data item table, select **Geography** for the **Role**.

   The Geography window is displayed.

Display 31.10   Geography Window

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.

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. Your change is saved automatically.

Modify a Data Item's Role

You can modify a data item's role. For example, you might want to modify a category data item to be a geography data item. A data item's role cannot be changed if that data item is in use in the report.

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**, **Category**, or **Geography**. Your change is saved automatically.

   Note: Category data items cannot be converted into measure data items.

Modify the Format of a Measure or a Date-Time Data Item

Note: The format cannot be changed for a data item that is used in a filter or a calculated item.

To change the format of a measure or date-time data item:

1 Select a measure data item on the **Data** tab.

2 In the table, select the existing format. A list is displayed with the **Format type**, **Width**, and **Decimals**. Make your selections. A sample of your selection (or selections) is displayed under the list. Here is an example of the list:

Display 31.11   Available Formats for a Measure Data Item

```
Format type:
Numeric
Currency
Comma
Dollar
Euro
Float
Percent

Width:      12    (2-32)
Decimals:   0     (0-32)
Sample:     56789.908772

Reset to Default    OK    Cancel
```

Display 31.12   Available Formats for a Date-Time Data Item

```
Format type:
MMMYYYY
Day, Date
Month, Day, Year
Year
YYYYMM
YYYYMMDD
Year, Quarter

Format:  2011

Reset to Default    OK    Cancel
```

Note: There are different formats available for date, date-time, and time.

3 Click **OK** to save your changes.

Modify How a Measure Is Aggregated

You can change the aggregation method for a measure in a data source using either the **Data** tab or 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.
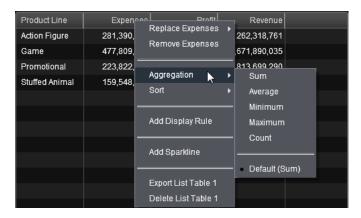
Display 31.13   Available Aggregations for a Data Item

To change the aggregation method when you are working with a report object in the canvas:
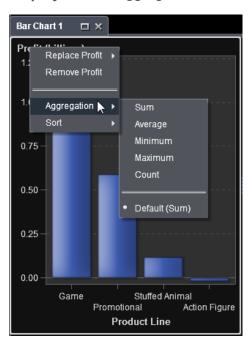
1 Choose one of the following:

   ▪ For a table, right-click a measure header in the report object.

Display 31.14   Aggregation Menu Items for a Table

■   For a graph, right-click the measure name hotspot.

Display 31.15   Aggregation Menu Items for a Graph

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.

Working with Data Item Filters in a Report
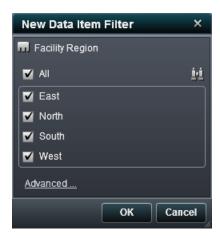
Filters are used to restrict the data that is returned from a query in a data source. The data items filters that you create in SAS Visual Analytics Designer are applied to the report.

Add a Data Item Filter

To add a new data item filter:

1 On the **Data** tab in the left pane, right-click the data item that you want to filter.

2 Select either **Using Continuous Values** or **Using Discrete Values**. The New Data Item Filter window is displayed.
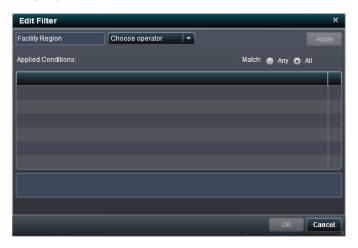
Display 31.16   New Data Item Filter Window

3 In the New Data Item Filter window, select the data item (or data items) that you want to filter.

4   (Optional) Click ⚏ to search for a data item (or data items).

5   (Optional) Click **Advanced** to add more conditions to the data item filter. The Edit
Filter window is displayed.

Display 31.17   Edit Filter Window
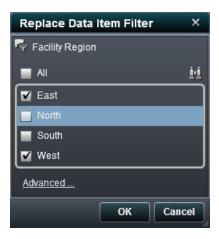
a   Select an operator and its related value (or values). Then click **Apply**.

b   For **Match**, select **Any** or **All**.

c   Click **OK** in the Edit Filter window.

6   Click **OK** in the New Data item Filter window. The ⚏ or ⚏ icons identify the new
filtered data item on the **Data** tab.

Replace a Data Item Filter

To replace an existing data item filter:

1   On the **Data** tab in the left pane, right-click the data item filter that you want to
replace.

2   Select **Replace Data Item Filter**. Then, select either **Using Continuous Values** or
**Using Discrete Values**.The Replace Data Item Filter window is displayed.

Display 31.18   Replace Data Item Filter Window
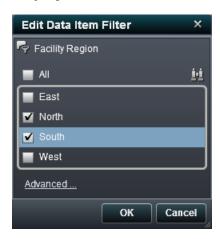
Edit a Data Item Filter

To add an existing data item filter:

1   On the **Data** tab in the left pane, right-click the data item filter that you want to edit.

2   Select **Edit Data Item Filter**. The Edit Data Item Filter window is displayed.

Display 31.19   Edit Data Item Filter Window

3   Select the data item (or data items) that you want to filter or clear the data item (or data items) that you do not want to filter.

4   (Optional) Click **Advanced** to add more conditions to the data item filter. The Edit Filter window is displayed.

   a   Select an operator and its related value (or values). Then click **Apply**.

   b   For **Match**, select **Any** or **All**.

   c   Click **OK** in the Edit Filter window.

5   Click **OK** in the Edit Data item Filter window.

Delete a Data Item Filter

To remove an existing data item filter:

1   On the **Data** tab in the left pane, right-click the data item filter that you want to remove.

2   Select **Delete Data Item Filter**. The filter data item is removed from the **Data** tab.

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 date-time 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 do not support hierarchies.

Note: Crosstabs can have either a hierarchy or categories on each row or column, but not both.

In SAS Visual Analytics Designer, you can:

- "Create New Hierarchies for a Report" on page 308
- "Edit Hierarchies for a Report" on page 309
- "Delete a Hierarchies for a Report" on page 310

Create New Hierarchies for a Report

To create a new hierarchy for a report:

1   On the **Data** tab, click ⬇ and then select **New Hierarchy**. The Create Hierarchy window is displayed.

Display 31.20   Create Hierarchy Window

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.

Edit Hierarchies 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.

Display 31.21   Create Hierarchy Window
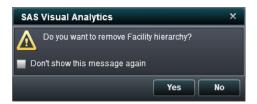
2   Edit the **Name**, add at least two categories, or remove a category (or categories) from the **Hierarchy**.

3   Click **OK** to save the updated hierarchy.

Delete a Hierarchies for a Report

To delete hierarchies:

1   Right-click the hierarchy name on the **Data** tab and select **Delete Hierarchy**.

2   Click **Yes** in the confirmation message that is displayed.

Display 31.22   *Remove Hierarchy Confirmation Message*

Working with Calculated Data Items in Reports

About Calculated Data Items

SAS Visual Analytics 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 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.
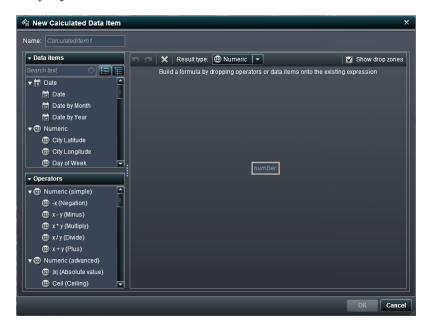
A hierarchy can contain calculated data items.

Add a New Calculated Data Item to a Report

To add a new calculated data item:

1 On the **Data** tab, click ⏷ and then select **New Calculated Item**. The New Calculated Data Item window is displayed.

Display 31.23 New Calculated Data Item Window

2 Enter a **Name**.

3 Select the data type for the calculated item from the **Result type** drop-down list. Your choices are **Date**, **Datetime**, **Numeric**, and **Time**. **Numeric** is the default data type.

Calculated data items in SAS Visual Analytics Designer always default to the following formats, which are based on the data type:

- ▪ Date: DATE10

- ▪ Datetime: DATETIME10

- ▪ Numeric: COMMA12.2

- ▪ Time: TIME10

After you create the new calculated data item, you can change its format using the data item table on the **Data** tab.

SUM is the default aggregation for new numeric calculated data items in SAS Visual Analytics Designer. You can change the aggregation for numeric calculated data items using the data item table on the **Data** tab.

4   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 the **x - y (Minus)** 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 then expression changes to (`Profit` — `Revenue`), and so on.

There are a large number of operator types available to perform mathematical functions, process datetime values, and evaluate logical processing such as "if" clauses. For more information about operators, see Appendix 1, "Operators for Calculated Data Items," on page 461.

5   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 only if you change a calculated numeric measure to a category data item.

Editing a Calculated Data Item

To edit a calculated data item:

1   Right-click a calculated data item on the **Data** tab and select **Edit Calculated Item**. The Edit Calculated Item window is displayed.

2   Modify the calculated data item as needed.

3   Click **OK**.

You can duplicate and remove calculated data items using the same steps as any other data items. If a calculated data item is used inside another calculated item, then it cannot be removed.

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.

Display 31.24   Measure Details Window

2   Click **Close**.

Showing or Hiding Data 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.

Display 31.25   Show or Hide Data Items Window

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**.

Modifying Data Role Assignments in Report Objects

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.

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.

Display 31.26 Roles Tab

3 Click ▼ beside the role that you want to edit. For roles that allow multiple data items, the menu items that are available can vary.

If you select a data item in the **Roles**, then the **Aggregation**, **Add**, **Remove**, and **Replace** menu items are displayed. You can also right-click a data item, and the actions specific to that data item are displayed.

For example, for the bar chart shown in Display 31.26 on page 316, you can replace or remove the **Category**. For the **Measure**, you can modify the aggregation, add another measure, replace the existing measure, or remove the measure. You can also modify the **Group**, the **Lattice columns**, or **Lattice rows**.
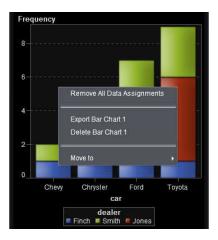
Removing Data Role Assignments from Report Objects

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:
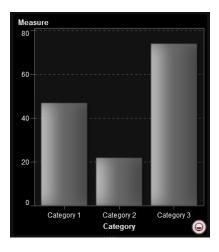
Display 31.27   *Report Object Menu*

2   Select **Remove All Data Assignments**. All data items are removed from the assigned 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 roles:

Display 31.28   *Report Object with Missing Data Assignments*

32

Sorting Data Items 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 SAS Visual Analytics 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.

Sorting Data in a List Table

To sort values in a list table:

1   Select a column heading.

2   Right-click and select either **Ascending** or **Descending**.

   Display 32.1   *Sort Menu Item for List Tables*

Note: You can also press the spacebar to sort a single column in a table. To sort multiple columns in a table, navigate to the column heading of each additional column that you want to sort. Press Ctrl+spacebar.

Sorting Data in a Crosstab

In crosstabs, right-click a column heading to sort data. You can also sort categories, rather than measures, by clicking the category name.

To specify a sort for values in a crosstab:
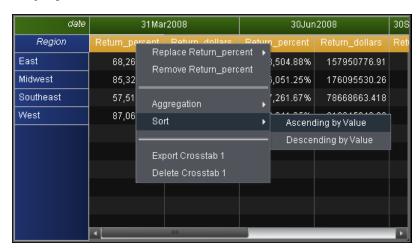
1   To sort by category values in a crosstab, right-click a column or row heading. Then select **Sort**. Then select either **Ascending** or **Descending**.

Display 32.2 *Sort Menu for a Category in a Crosstab*

2 To sort by measure values in a crosstab, right-click a column or row heading. Then select **Sort**. Then select either **Ascending by Value** or **Descending by Value**.

Display 32.3 *Sort Menu for a Measure in a Crosstab*
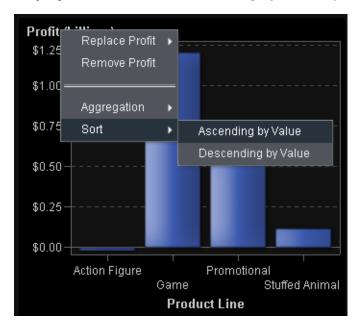
Sorting Data in a Graph

To sort values in a graph:

1 To sort by category values in a graph, right-click a measure name and select **Sort**. Then select either **Ascending by Value** or **Descending by Value**.

Display 32.4   Sort Menu for a Category in a Graph

2 To sort by measure values in a graph, right-click a category name and select **Sort**. Then select either **Ascending Alphabetically** or **Descending Alphabetically**.

33

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.

The SAS Visual Analytics Designer provides the following interfaces for display rules:

- The **Display Rules** tab in the right pane enables you to populate or add intervals 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 report 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.
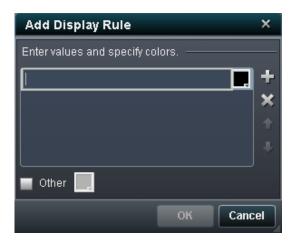
Note: The dual axis bar chart, the dual axis bar-line chart, and the dual axis line chart do not support display rules.

Adding Report-Level Display Rules

To add a report-level display rule:

1 Without any report objects or the section selected in the report canvas, click the **Display Rules** tab in the right pane. Click 📷. The Add Display Rule window is displayed.

Display 33.1   Add Display Rule Window

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.

In the following example, values and colors have been specified for each region in a list table.

Display 33.2   Add Display Rule Window with Values and Colors Specified

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.

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 Display Rule window above.

Display 33.3   A List Table and a Bar Chart with Report-Level Display Rules Applied

Adding Table-Level Display Rules

You can add three different types of display rules to tables. Note that you cannot create a display rule to highlight dates in a table.

Adding Display Rules Using a Gauge

To specify a new display rule for a table using a gauge:

1   If not already selected, select the table in the report canvas that you want to update.

2 In the right pane, click the **Display Rules** tab. Click 🔳. The Add Display Rule window is displayed.

3 Select **Gauge** as the **Type**. The Add Display Rule window for gauges is displayed.

Display 33.4   *Add Display Rule Window for a 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.

Display 33.5 *Populate Intervals Window*

You can specify the **Number of intervals**, the **Lower bounds**, and the **Upper bounds**. Click **OK**.

8 Specify the **Column** in which you want the gauge to be displayed

9 Click **OK**. The display rule 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:

Display 33.6 *Display Rules Tab with the Display Rules for an Icon*

10 (Optional) Click ⬚ on the **Display Rules** tab to edit the new display rule.

Adding Display Rules Using an Expression

To specify a new display rule for a table using an expression:

1 If not already selected, select the table in the report canvas that you want to update.

2 In the right pane, click the **Display Rules** tab. Click ⬛. The Add Display Rule window is displayed. The **Expression** type is displayed by default.

   Display 33.7   Add Display Rule Window for an Expression

3 Select the **Column** or any measure value.

4 Select the **Operator**. Your choices are **Greater than**, **Is**, **Is not**, **Between**, **Less than**, **Less than or equal to**, **Greater than**, **Greater than or equal to**, **Is missing**, or **Is not missing**. The default is **Greater than**.

5 Enter or select a **Value**.

6 Modify the style, size, and color of the font.

7 Select the row or column in the **Applies to** drop-down list.

8 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:

Display 33.8   Display Rules Tab with the Display Rules for an Expression

9 (Optional) Click on the **Display Rules** tab to edit the new display rule.

Adding Display Rules Using Color-Mapped Values

To specify a new display rule for a table using color-mapped values:

1 If not already selected, select the table in the report canvas that you want to update.

2 In the right pane, click the **Display Rules** tab. Click . The Add Display Rule window is displayed.

3 Select **Color-mapped values** as the **Type**. The Add Display Rule window for color-mapped values is displayed.

Display 33.9   *Add Display Rule Window for 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.

In the following example, values and colors have been specified for products in a list table.

Display 33.10   Add Display Rule Window with Color-Mapped Values Displayed

8   (Optional) Select the **Other** check box. Then select a color.

9   Using the **Applies to** drop-down list, specify where you want to apply the colors.

10   Click **OK**. The table updates with the display rules.

The display rule appears on the **Display Rules** tab in the right pane.

Display 33.11   Add Display Rule Window with Color-Mapped Values Displayed

Adding Graph-Level Display Rules
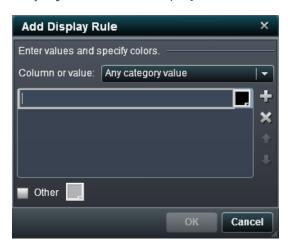
You can add display rules to graphs. Note that you cannot create a display rule to highlight dates in a graph.

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 not already selected, select the graph in the report canvas that you want to update.

2 In the right pane, click the **Display Rules** tab. Click . The Add Display Rule window is displayed.

3 Select a **Column or value** for the display rule.

4 Click in the box to enter a value for the display rule in the field.

Display 33.12   Add Display Rule Window for a Graph

5 Select a color for the display rule.

6 (Optional) Repeat the steps for entering a value and selecting a color. Use the arrows to move values up or down.

In the following example, values and colors have been specified for each product line in a bar chart.

Display 33.13   Add Display Rule Window with Values and Colors Specified

7 (Optional) Select the **Other** check box. Then select a color.

8 Click **OK**. The graph updates with the new display rule.

The display rule appears on the **Display Rules** tab in the right pane.

Display 33.14   Display Rules Tab with the Display Rules for a Graph

9 (Optional) Click 🔖 on the **Display Rules** tab to edit the new display rule.

10 (Optional) Click ![icon] in the **Display Rules** tab to move the new 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 not already selected, select the gauge in the report canvas that you want to update.

2 In the right pane, click the **Display Rules** tab. Click ![icon].

Here is an example of the **Display Rules** tab for a dial gauge:

Display 33.15   Display Rules Tab for a Dial Gauge
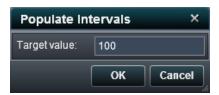
3 Specify (or modify) the display rules for the gauge. You can populate intervals, edit the display rule, and specify the **Type**.

To automatically populate the intervals:

a Click ![icon] to display the Populate Intervals window.

Display 33.16   Populate Intervals Window

b   Enter a number for **Target value**.

c   Click **OK**.

To edit the display rule for a gauge:

a   Click 🔳 to display the Edit Display Rule window.
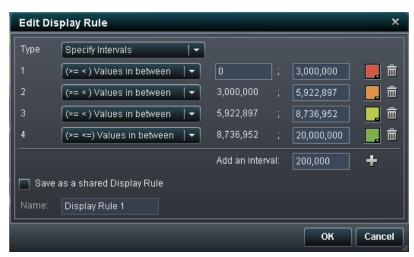
Display 33.17   Edit Display Rule Window

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:

Display 33.18   *Specifying Intervals in the Edit Display Rules Window*

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**.

34

Working with Report Filters

About Report Filters

For certain report objects, you can subset your data by using the **Filters** tab in the right pane. You can base your filters on any data item, regardless of whether the data item is assigned to the current report.

Filters on a report object are applied in conjunction with any data item filters that have been applied. Data item filters do not impact the selections that you can make for filters on a report object. For example, if there is a data item filter on a *Country* item that has the selections for *Canada*, *United States*, and *Germany*, you might still have a report object filter with selections such as *Canada* and *Mexico*. If these filters are applied together, then the report object only displays content for *Canada*.

For information about data item filters for reports, see "Working with Data Item Filters in a Report" on page 304.

Filters that are created in SAS Visual Analytics Explorer might contain expressions that you cannot create in SAS Visual Analytics Designer. In some cases, you might be able to only remove the filter.

Creating a Filter for a Report Object

To create a filter for a report object:

1   If not already selected, select the report object in the report canvas that you want to filter.

2   In the right pane, click the **Filters** tab.

3   Select a data item. Then, click **Add Filter**. The filter appears in the tab.

Display 34.1   Filters Tab with a Filter Specified

4   Click ▼ for options. The available options depend on the whether you are filtering dates, characters, or numerics. The options might include **Edit Filter**, **Filter using continuous values**, **Filter using discrete values**, **Select all**, **Clear all**, and **Invert selection**.

For information about using the Edit Filter window, see "Editing a Filter Using the Edit Filter Window" on page 341.

If the **Auto** check box is selected, then your filter choices are applied automatically.

5   (Optional) Clear the **Auto** check box if you want to manually select (or clear) filter choice. Manually selecting (or clearing) filters activates the **Apply** button.

Editing a Filter Using the Edit Filter Window
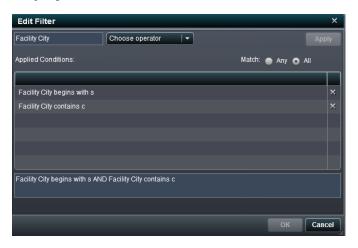
You can edit most filters using the Edit Filter window. The Edit Filter window is accessible when you filter by discrete values, but not when you filter continuous values. In the discrete values case, the Edit Filter window provides access to the `IS MISSING` and `IS NOT MISSING` operators, as well as the `NOT IN` operator.

A date format change is not allowed for certain data items when filters are present.

To edit a filter using the Edit Filter window:

1   In the right pane, click the **Filters** tab.

2   Click ▼ to open the options drop-down list for the filter that you want to edit and then select **Edit Filter**. The Edit Filter window is displayed.

Display 34.2   Edit Filter Window

3   Build the expression for your filter by choosing operators and specifying conditions.

Table 34.1   Conditions for Character Data

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.
Not In	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.
Begins with	Specifies that a matching value must contain the filter value at the start of the value. The comparison of the values is not case sensitive, and leading and trailing spaces are ignored.
Ends with	Specifies that a matching value must contain the filter value at the end of the value. The comparison of the values is not case sensitive, and leading and trailing spaces are ignored.
Contains	Specifies that a matching value must contain the filter value. The comparison of the values is not case sensitive, and leading and trailing spaces are ignored.
Does not contain	Specifies that a matching value must not contain the filter value. The comparison of the values is not case sensitive, and leading and trailing spaces are ignored.
Is missing	Specifies that a missing value matches the filter.

Is not missing	Specifies that any nonmissing value matches the filter.

Table 34.2   Conditions for Numeric Data

Is	Specifies that a matching value must match one of the filter values exactly.
Is not	Specifies that a matching value must not match the specified filter values exactly.
Between	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.
Less than	Specifies that a matching value must be less than the filter value.
Less than or equal to	Specifies that a matching value must be less than or equal to the filter value.
Greater than	Specifies that a matching value must be greater than the filter value.
Greater than or equal to	Specifies that a matching value must be greater than or equal to the filter value.
Is missing	Specifies that a missing value matches the filter.
Is not missing	Specifies that any nonmissing value matches the filter.

4   Click **Apply** to add the condition to the list of **Applied Conditions**.

5   Click **OK** to apply the modified filter.

Removing a Filter for a Report Object

To remove a filter, click ✖ on the filter on the **Filters** tab.

35

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 interaction view in SAS Visual Analytics Designer enables report authors to specify which interactions they would like to add to tables, graphs, and gauges in a report.

There are two 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 both. 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 scatter plots and time series plots. Controls that are used on the report canvas can also be the source if an interaction. Controls that are used as section prompts are treated as automatic filters and are not displayed in the interaction view.

Objects that are imported from SAS Visual Analytics Explorer cannot participate in interactions.

Creating Interactions

Creating a Filter Interaction

You can create interactions using the interaction view or the **Interactions** tab in the right pane.
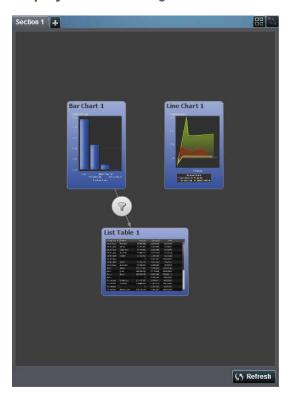
Note: You cannot create interactions from List tables that use detail data.

To create a filter interaction using the interaction view:

1 Add the report objects that you want to use to the report canvas. For example, you might have a bar chart, a line chart, and a list table.

2 Click ⚇ above the canvas to switch to interaction view.

3 In the interaction 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 must be based on the same data source.

Display 35.1 *Creating a Filter Interaction between a Bar Chart and a List Table*

4 Right-click ▼. Then select **Interaction Type** ▶ **Filter**.
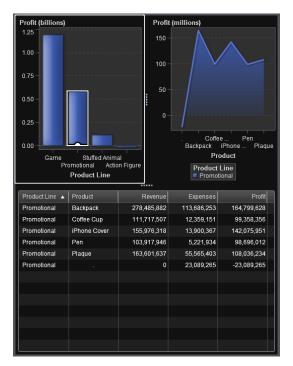
*Display 35.2   Interaction Type Menu*

5   (Optional) Add other filter interactions.

6   (Optional) Click **Refresh** below the canvas to update the interactions.

7   Click ⊞ to switch to the layout view. 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 35.3   A Filtering Example with a Bar Chart, a Line Chart, and a List Table

Clicking another selection applies the filter based on your new selection. For example, in the report above, you can click the game product line to change the filter.

The new interaction (or interactions) appears in **Interactions** tab in the right pane.

The following example shows the **Interactions** tab for the filtering example shown in .

Display 35.4   Interactions Tab with Two Interactions Displayed
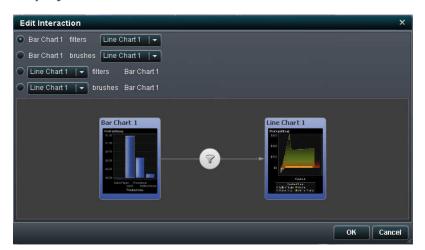
To clear the selection and reset the filter (or filters), press Ctrl + click while viewing the original report object.

To create an interaction using the **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 [icon]. The **Edit Interaction** window is displayed.

Display 35.5   Edit Interaction Window

4   In the Edit Interaction window, select the report object that you want 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 *Line Chart 1*.

5   When you are finished working with interactions, click **OK**.

6   The new interaction is displayed in the **Interactions** tab.

Display 35.6 *A New Filter in the Interactions Tab*
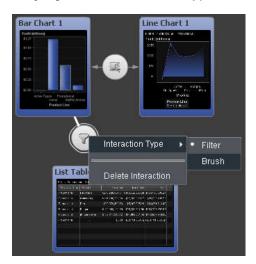
Creating a Data Brushing Interaction

To create an interaction using data brushing:
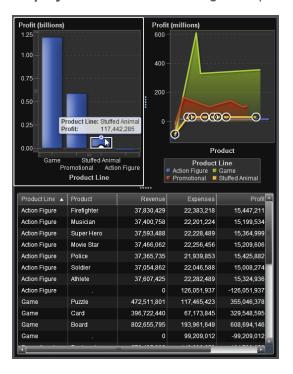
1 Add the report objects that you want to use to the canvas. For example, you might have a bar chart, a line chart, and a list table.

2 Click ⚡ above the canvas to switch to interaction view.

3 In the interaction view, draw a connection between the source and target report objects. When you have created an interaction, the ▼ is displayed between the report objects.

4 Right-click ▼. Then select **Interaction Type ▶ Brush**.

Display 35.7   Interaction Type Menu with Brush Item Selected

The filter icon changes to .

5   (Optional) Add other data brushing interactions.

6   (Optional) Click **Refresh** below the canvas to update the interactions.

7   Click to switch to the layout view. 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 pie chart and the crosstab highlight the same data.

Display 35.8   *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.
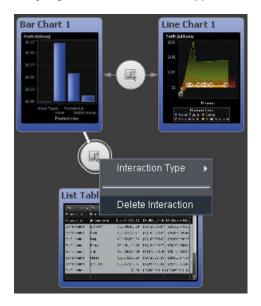
Deleting Interactions

You can delete interactions using the interaction view or the **Interactions** tab in the right pane.

To delete either a filter or data brushing interaction using the interaction view:

1   Click above the canvas to switch to interaction view.
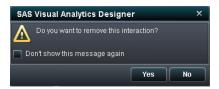
2   Right-click . Then select **Delete Interaction**.

Display 35.9   Interaction Type Menu with Delete Selected

3 Click **Yes** in the confirmation message that is displayed.

Display 35.10   Delete Interaction Confirmation Message

4 (Optional) Remove other interactions.

5 Click ⊞ to switch to the layout view.

To delete an interaction using the **Interactions** tab:

1 Click the **Interactions** tab.

Display 35.11   *A New Filter in the Interactions Tab*

2   Click 🗑.

3   Click **Yes** in the message that is displayed.

Display 35.12   *Delete Interaction Confirmation Message*

36

Maintaining Multi-Section Reports

Overview of Report Sections

Any report in SAS Visual Analytics Designer can have multiple sections. 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. There is no limit to the number of sections that can be included 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.

Here is an example of what the first section tab looks like.

Display 36.1   One Section Tab

To add a new section to a report, click ✚ in the tab to the right of the first section in the report. The new tab appears to the right of the existing tab (or tabs).

Here is an example:

Display 36.2   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.

   Display 36.3   The Menu Option for Renaming a Section

2   Select **Rename**. The name of the tab is highlighted.

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 in the **Properties** tab.

Deleting a Section from a Report

To delete a section from a report, right-click the section tab that you want to remove. When the menu is displayed, select **Delete**.

Note: The **Delete** 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 also drag an object from one section and drop it onto the tab for another section to place it in that section.

Note that 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.

37

Exporting Data from Report Objects

Overview of Exporting from SAS Visual Analytics Designer

All users who have the Export Data capability can use SAS Visual Analytics Designer to export data from report objects to Microsoft Excel format for future viewing or printing. This output can be saved locally on disk and then opened in Microsoft Excel.

When you export a graph, you are exporting the data not the visual graph representation.

Not all report objects in SAS Visual Analytics Designer support the exporting feature. For example, you cannot export data for 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 in the object.

Exporting Data from Report Objects

Export Data from a List Table

To export data from a list table:

1 If not already selected, select the list table that you want to update.

2 Right-click in the list table and select **Export <listTableName>**, where `<listTableName>` is the name of the report object.
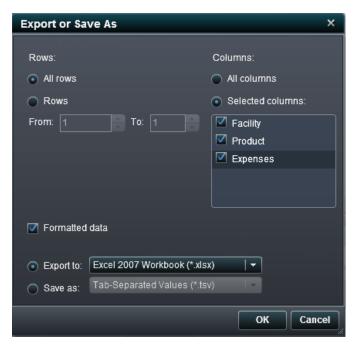
Here is an example of the **Export** menu option for a list table.

Display 37.1   Export Menu for List Tables

The Export or Save As window is displayed.

Display 37.2   *Export or Save As Window for List Tables*

3   Choose one of the following export options for rows:

▪   Select the **All rows** radio button to export all rows.

▪   Select the **Rows** radio button to specify a range of rows to export. Enter a number in the **From** and **To** fields.

4   Choose one of the following export options for columns:

▪   Select the **All columns** radio button to export all columns.

▪   Select the **Selected columns** radio button to specify which columns to export. Then, select the check box 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 the **Export to** radio button, then the only option is **Excel 2007 Workbook (*.xlsx)** to create a Microsoft Excel spreadsheet.

7 If you select the **Save as** radio button, choose one of the following options:

 ■ Select **Tab-Separated Values (.tsv)** to create a data file.

 ■ Select **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 not already selected, select the crosstab that you want to update.

2 Right-click in the crosstab and select **Export <crosstabName>**, where `<crosstabName>` is the name of the report object.
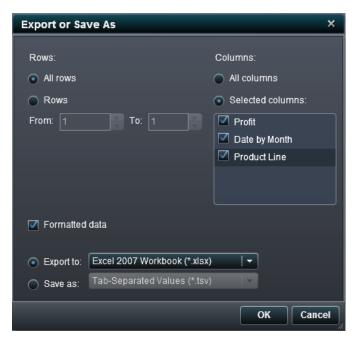
Here is an example of the **Export** menu option for a crosstab.

Display 37.3   Export Menu for Crosstabs

Date by Year		2010		2011
Facility State	Facility City	Revenue	Expenses	Revenue
AL	Birmingham	1,343,126	1,582,327	4,135,
	Mobile	Export Crosstab 1	2,126	1,355,
AR	Little Rock	Delete Crosstab 1	4,232	20,087,
AZ	Scottsdale		4,412	7,109,
CA	Los Angeles	Move Crosstab 1 to ▸	5,769	
	San Diego	1,437,426	2,402,617	1,169,
	San Francisco	26,878,679	9,516,337	24,044,
IL	Chicago	8,334,320	3,011,832	10,518,
	Elgin	141,971	1,001,233	108,
	Joliet	1,028,966	1,604,408	1,288,
LA	New Orleans	2,699,708	1,550,269	3,045,
NH	Manchester	1,063,952	2,720,423	930,
NJ	Newark	4,555,827	1,894,203	5,428,
NY	New York	1,383,978	1,847,918	1,110,
OH	Cincinnati			

The Export or Save As window is displayed.

Display 37.4   Export or Save As Window for Crosstabs

3   Choose one of the following export options for rows:

■   Select the **All rows** radio button to export all rows.

■   Select the **Rows** radio button to specify a range of rows to export. Enter a number in the **From** and **To** fields.

4   Choose one of the following export options for columns:

■   Select the **All columns** radio button to export all columns.

■   Select the **Selected columns** radio button to specify which columns to export. Then select the check box 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 the **Export to** radio button, then the only option is **Excel 2007 Workbook (*.xlsx)** to create a Microsoft Excel spreadsheet.

7 If you select the **Save as** radio button, choose one of the following options:

- ■ Select **Tab-Separated Values (.tsv)** to create a data file.

- ■ Select **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 report object, because only the data can be exported.

To export data from a graph:

1 If not already selected, select the graph that you want to update.

2 Right-click in the graph and select **Export <graphName>**, where `<graphName>` is the name of the report object.
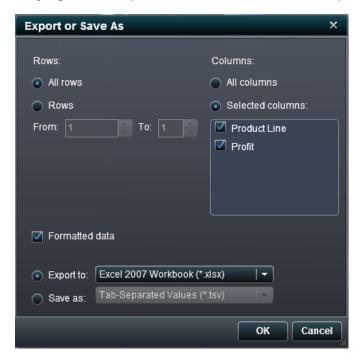
Here is an example of the **Export** menu option for a line chart.

Display 37.5   *Export Menu for Graphs*

The Export or Save As window is displayed.

Display 37.6   Export or Save As Window for Graphs

3 Choose one of the following export options for rows:

  ▪ Select the **All rows** radio button to export all rows.

  ▪ Select the **Rows** radio button to specify a range of rows to export. Enter a
     number in the **From** and **To** fields.

4 Choose one of the following export options for columns:

  ▪ Select the **All columns** radio button to export all columns.

  ▪ Select the **Selected columns** radio button to specify which columns to export.
     Then select the check box 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 the **Export to** radio button, then the only option is **Excel 2007 Workbook (*.xlsx)** to create a Microsoft Excel spreadsheet.

7   If you select the **Save as** radio button, choose one of the following options:

▪   Select **Tab-Separated Values (.tsv)** to create a data file.

▪   Select **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.

Part 5

Viewing Reports

38

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 app is available in the iTunes App store. The Android app is available from Google Play.

When you open the app for the first time, sample reports are available in My Portfolio. You can also view tutorial videos about how to use the app:

1 Tap Help.

2 Tap ▦. The YouTube video playlist for SAS Mobile BI opens.

3 Tap and select the video that you want to view.

The app also provides help that you can view in My Portfolio to learn about other ways to get help with SAS Mobile BI.

Add a Server Connection for iPad

When you first open the SAS Mobile BI app, the only connection that is available is to the SAS Demo Server.

To access your corporate library server, you must add a connection to that server. If your company has more than one server, you can add additional connections in the same way.

Note: Contact your SAS Visual Analytics 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 connection (SSL), 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 app, the only connection that is available is to the SAS Demo Server.

To access your corporate library server, you must add a connection to that server. If your company has more than one server, you can add additional connections in the same way.

Note: Contact your SAS Visual Analytics administrator for your user ID, password, and server information.

To connect to a server:

1   Tap **Library** in My Portfolio.

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 connection (SSL), 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 appears and displays a message to help you correct the problem.

39

Viewing Reports in the 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 enables them to view report content.

Opening a Report in SAS Visual Analytics Viewer

To open a report in the SAS Visual Analytics Viewer:

◼ Double-click 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 21.

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 SAS Visual Analytics Viewer than it does in SAS Visual Analytics Designer or on a mobile device.

Here is an example of a report in the SAS Visual Analytics Viewer:

Display 39.1   A Report in SAS Visual Analytics Viewer

If you have the Create Reports capability, then you can select **File** ▶ **Edit Report** in the current report. SAS Visual Analytics Designer is displayed, and then you can edit the report.

Adding Comments to a Report in SAS Visual Analytics Viewer

If you have the Add Comments capability, then you can add (or view) comments. The comments that you add are automatically saved with the report.

To add a comment to a report:

1 Click the **Comments** tab in the right pane.

2 Enter a topic name and a comment.

Here is an example:

Display 39.2   Adding Comments in SAS Visual Analytics Viewer

3 (Optional) Click 📎 to add 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 Select an existing comment in the **Comments** tab in the right pane. Then, enter a reply.

2 (Optional) Click ![paperclip icon] to add a file or image to your reply.

3 Click **Post**. Your reply is added to the **Comments** tab in the right pane.

Note: Only administrators can modify or delete an existing comment.

To search for a comment:

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

2 (Optional) To clear your search, click ![clear icon]. Then, you can enter another word or phrase in the search box.

Part 6

Managing the Environment

40

Overview of SAS Visual Analytics Administrator

About SAS Visual Analytics Administrator

SAS Visual Analytics Administrator enables administrators to view SAS LASR Analytic Server and manage server connections by starting or stopping them in the **LASR Tables** tab. You can load tables from SAS metadata to the SAS LASR Analytic Servers, HDFS, or to co-located providers such as Teradata and Greenplum servers, and register them as LASR tables. Tables that are loaded to the SAS LASR Analytic Server are referred to as LASR tables. Tables that are loaded can be unloaded and reloaded as needed.

After the SAS LASR Analytic Server is started, administrators can view the information about the status, host, and port information for server connections that are associated with the SAS LASR Analytic server component in the metadata. The **LASR Tables** tab also provides information about rows and columns for LASR tables that are associated with libraries and the server connection. User account information for the start time of a SAS LASR Analytic Server connection, and the last time it was accessed is displayed. In addition, you can conveniently start and stop the SAS LASR Analytic Server components that are displayed in the **LASR Tables** tab.

An Authorization page provides quick information about the effective permissions for a table or a folder, indicates any direct access controls, and enables you to add or remove explicit controls. You can add an explicit grant or denial for the tables in the metadata or for LASR tables. You can also set a row-level permission for a LASR table by selecting a conditional grant and specifying the condition.
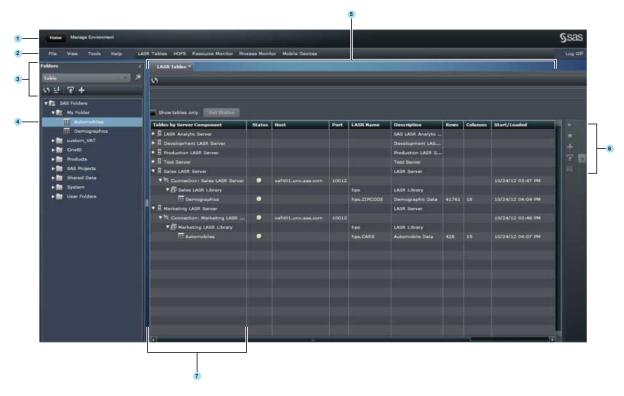
For deployments that use SAS High-Performance Deployment of Hadoop, you can view information about the properties and values for files in HDFS. Because SAS uses a special file format for the data that is stored in HDFS, the HDFS content explorer also provides information about the columns and row count for the prepared data.

Resource and process monitors are provided so that administrators can view real-time statistics and visualizations of resource use such as central processing unit (CPU) utilization, memory utilization, and input/output (I/O) rates. A line plot of resource utilization against time is provided. In addition, a real-time view creates a visualization for each blade that is used in the cluster and each CPU on a blade.

The **Mobile Devices** tab enables you to manage mobile devices that use SAS Mobile BI. This enables administrators to easily manage access to data through SAS Visual Analytics applications by mobile devices. Information about the mobile devices in your organization is available to administrators. The information includes user ID, device information, and timestamp of last access.

Your First Look at SAS Visual Analytics Administrator

SAS Visual Analytics Administrator contains the following elements:

1   The application bar enables you to return to the SAS Visual Analytics Home page.

2   The menu bar contains menus and buttons that enable you to perform tasks.

3   The navigation pane controls enable you to filter, sort, and search for objects in SAS Folders. When you select a table, icons appear that enable you to perform actions on the table.

4   The navigation pane displays the SAS Folders, and a tree of tables within the SAS Folders. You can select a folder or table, and right-click to select **Authorization**. In the Authorization page, you can view the effective permissions for tables in the

metadata, and modify permissions. If your site is deployed with SAS Visual Analytics for Hadoop, add a table to HDFS or reload a table to HDFS. You can also add tables to co-located data providers such as Greenplum or Teradata.

5   The tabs that you select in SAS Visual Analytics Administrator determine the types of tasks that you can perform in the workspace. You can work with tables by loading them to the SAS LASR Analytic Server or by unloading them. You can also reload tables. You can view real-time statistics and visualizations of resource use such as central processing unit (CPU) utilization, memory utilization, and input/output (I/O) rates. A real-time view creates a visualization for each blade that is used in the cluster and each CPU on a blade. In addition, you can manage mobile devices by viewing the list of devices that are using SAS Mobile BI and blacklisting specific devices.

You can work with the **Mobile Devices** tab to view the list of mobile devices that are using SAS Mobile BI and blacklisting specific devices.

6   The icons on the right side of the **LASR Tables** tab enable you to start or stop SAS LASR Analytic Server connections, load tables from the metadata to the SAS LASR Analytic Server, and unload or reload LASR tables.

7   In the **LASR Tables** tab, the **Tables by Server Component** column displays the SAS LASR Analytic Server components, their server connections, and the libraries that were created in the metadata. LASR tables that were loaded, unloaded, or reloaded are also displayed in this column.

The **Status** column shows the status of server connections and tables.

Accessing SAS Visual Analytics Administrator

You can access SAS Visual Analytics Administrator by using any of the following methods:

■   On the SAS Visual Analytics Home page, click **Manage Environment** to open the SAS Visual Analytics Administrator.

■ Follow a web link or enter the URL to access SAS Visual Analytics Administrator directly. For example, you might enter:

```
http://hostname.example.com:8080/
SASVisualAnalyticsAdministrator
```
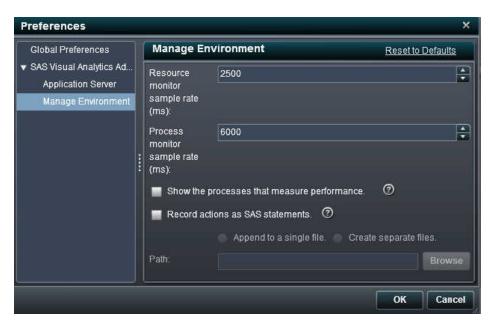
Specifying Your Preferences for SAS Visual Analytics Administrator

The **Manage Environment** panel in SAS Visual Analytics Administrator enables you to specify the following preferences:

■ Select a SAS Application Server.

■ Specify the sample rate for Resource Monitor and Process Monitor.

■ Record actions as SAS statements in a single file or in separate files.

Here is a display of the **Manage Environment** panel within the Preferences window.

Display 40.1 *Manage Environment Panel in the Preferences Window*

You can record SAS statements to a single file or within separate files. This feature enables administrators to review the SAS statements for a task. Some of these tasks are starting and stopping a server connection, and loading and unloading tables. The statements in the file can be modified and scheduled. For more information about the LASR and HPDS2 procedures, see the *SAS LASR Analytic Server: Administration Guide*.

Note: The recorded statements do not include information about the metadata server connection. The connection information is often supplied in options such as METASERVER=, METAPORT=, METAUSER=, or METAPASS=. These options are needed to run the PROC METALIB statement that is included in jobs that register tables. The affected jobs are the prepare data and add to data server jobs. For information about the metadata server connection options, see *SAS Language Interfaces to Metadata*.

To specify preferences that are specific to SAS Visual Analytics Administrator:

1 Select **File** ▶ **Preferences** to open the Preferences window.

2 Select **SAS Visual Analytics** ▶ **Application Server**.

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

3 Select **SAS Visual Analytics** ▶ **Manage Environment**. Enter values for the following preferences:

Parameter	Sample Value or Task	Description
Resource monitor sample rate (ms)	1000	Specify the sampling rate that the resource monitor uses for polling the machines in the cluster.
Process monitor sample rate (ms)	6000	Specify the sampling rate that the performance monitor uses for polling application instances.

Parameter	Sample Value or Task	Description
Show the processes that measure performance	Selected	By default, the processes that measure performance are filtered from the graphs in the Process monitor. However, if several instances of the processes that measure performance are running, they can reduce performance. By selecting the check box, you can view and monitor the processes that measure performance.
Record actions as SAS statements	Selected	Click one of the two options that are available, and specify a folder path or a fully qualified filename. You can either append the SAS statements to a single file or create separate files that record actions as SAS statements. Files are saved either as .txt files or as .sas files. Actions such as starting or stopping server connections, and loading or unloading tables are recorded.

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

See Also

SAS Visual Analytics: Administration Guide

What Does the LASR Tables Tab Show?
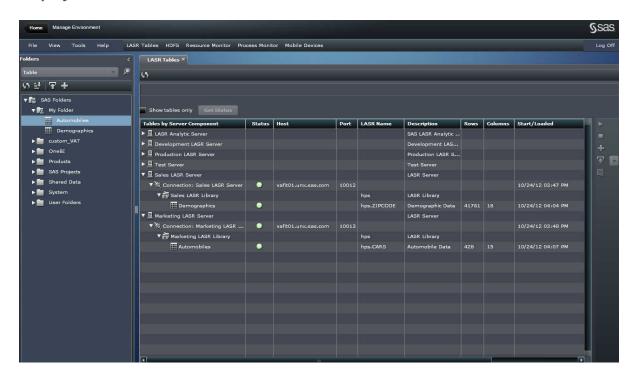
The **LASR Tables** tab in SAS Visual Analytics Administrator enables you to view SAS LASR Analytic servers that were created in the SAS metadata, start or stop the LASR Analytic Server connection associated with each server, and manage tables by loading them to the SAS LASR Analytic Server.

Note: If the **LASR Tables** tab shows the SAS LASR Analytic Server and a server connection, but does not show any libraries, then you must create and assign a library to that server connection in the metadata. This task is performed in SAS Management Console.

You can click ↻ to have the table display the latest information that is available from the metadata. You can select the **Show tables only** check box to view the tables in a flat list instead of a hierarchical tree view.

Here is a display of the **LASR Tables** tab.

Display 40.2 *LASR Tables Tab*

Here is a list that identifies the fields shown in the **LASR Tables** tab. Note that some of the fields described in this list are not displayed in the display of the **LASR Tables** tab shown here.

Tables by Server Component

This column displays the SAS LASR Analytic Servers that were registered in metadata. You can expand the tree view for a server component to view the server connection that is associated with that server. One server connection is allowed for each SAS LASR Analytic Server. You can also expand a server connection node to view the libraries and tables that are registered for the server. A server connection represents a SAS LASR Analytic Server instance.

Icons in the LASR Tables Tab	Description
≣	Represents a SAS LASR Analytic Server component. The SAS Deployment Wizard creates one SAS LASR Analytic Server. You can create multiple SAS LASR Analytic Servers in SAS Management Console. They are displayed as SAS LASR Analytic Server components in the SAS Visual Analytics Administrator.
⚒	SAS LASR Analytic Server connections are created in SAS Management Console. When you expand the ≣ node, the server connection that was defined for each server is displayed. Server connections are created in SAS Management Console, and they are started or stopped in SAS Visual Analytics Administrator. When a server connection is started in the **LASR Tables** tab, it represents the actual running SAS LASR Analytic Server instance.
🗄	Libraries are created in SAS Management Console for the SAS LASR Analytic Server and its associated server connection. Tables are loaded to these libraries in SAS Visual Analytics Administrator. When you expand the tree view for a library in the **LASR Tables** tab, any tables that are associated with that library are displayed.
⊞	This is the LASR table that was loaded from the metadata. In the **LASR Tables** tab, you can load tables from the metadata to the SAS LASR Analytic Server. A special configuration is required if you want to load tables for a library that was created by a different administrator. For more information, see *SAS Visual Analytics: Administration Guide*.

Status

The status column is used to provide a graphical representation of the status for a server connection or a table.

The following icons are used in the **Status** column:

- ● indicates that the server connection is running and ready for use. If a table is loaded, it also indicates that a table has been loaded to a server connection.

- ■ ■ indicates that the server connection is not running. If a table is unloaded, this column also indicates that a table is not loaded to a server connection.

- ■ ◆ indicates that the server connection did not respond.

Host

This field specifies the host name for the machine that is used as the SAS LASR Analytic Server root node. The host name is specified when the server connection is created in SAS Management Console.

Port

This field shows the network port that is used for communication with the server connection. The port number is specified when the server connection is created in SAS Management Console.

LASR Name

For libraries, this field shows the server tag (for example, hps). For tables, this field shows the table name in the following format:

```
serverTag:SASTableName
```

Description

If a description was specified when the SAS LASR Analytic Server component was created in SAS Management Console, that descriptive text is displayed in this column for the server.

If a description was specified when a table was loaded, that descriptive text is displayed in this column. This descriptive text is also displayed next to the LASR table when a user selects data sources in SAS Visual Analytics Designer or SAS Visual Analytics Explorer.

Rows

This field specifies the number of rows in the table. This information is displayed if the table is loaded.

Columns

This field specifies the number of columns in the table. This information is displayed if the table is loaded.

Start/Loaded

For SAS LASR Analytic Server connections, this field shows the time when a server connection was started. For loaded tables, this field shows the time when the table was loaded into the SAS LASR Analytic Server.

Accessed/Modified

For SAS LASR Analytic Server connections, this field shows the last time the server was accessed. For loaded tables, this field shows the time when a LASR table was modified in the SAS LASR Analytic Server.

A Quick Tour of the LASR Tables Tab

You can perform several tasks on the **LASR Tables** tab. On the right side of the **LASR Tables** tab, icons are available for performing tasks. Here is a description of the icons and the associated tasks:

Table 40.1   *Icons in the LASR Tables Tab*

Icon	Description
▶	Enabled when a SAS LASR Analytic Server connection is defined in SAS Management Console, and it is currently not running in SAS Visual Analytics Administrator. Select a SAS LASR Analytic Server connection, and click ▶ to start the server connection.
■	Enabled when a server connection has started and the status is indicated by 🟢 . This task cannot be performed if the status of the server connection is displayed with 🟥 .
✚	Enabled when you select a library for a server connection that is running. When selected, the **Load a Table** dialog box is displayed, and you can load a table into the SAS LASR Analytic Server.

Icon	Description
⬆	Enabled when a table is selected on the **LASR Tables** tab, and that table was previously loaded and unloaded.
	You can select an unloaded table and click ⬆ to reload the table or change the source table for a reloaded table.
	You can also reload a table by selecting ▼, and choosing **Reload** from the drop-down menu for this icon.
⬛	Enabled when a loaded table is selected on the **LASR Tables** tab.
	You can unload a loaded table from the SAS LASR Analytic Server.

A Quick Tour of the Folders Tab

The left side of the **LASR Tables** tab contains a collapsible navigation pane. The pane contains a view of the SAS Folders tree. You can perform the following tasks:

- Navigate the folder tree, select a table, and load the table into a SAS LASR Analytic server connection.

- Navigate the folder tree, select a table, and add a table to HDFS or to a co-located provider such as Teradata or Greenplum.

- Display the effective permissions for a table or a folder in the Authorization page.

41

Managing SAS LASR Analytic Server Connections

About SAS LASR Analytic Servers in the Metadata

SAS LASR Analytic Server components are created and registered in the metadata by using SAS Management Console. When a SAS LASR Analytic Server component is created in the metadata, a SAS LASR Analytic Server connection is automatically created for that server component. Both servers and their server connections can be viewed in SAS Visual Analytics Administrator.

Initially, the SAS Deployment Wizard creates a single SAS LASR Analytic Server component in the metadata. However, you can add multiple server components to the metadata by using SAS Management Console and specifying a host name and a unique port number for each server connection.

Working with SAS LASR Analytic Server Connections

About Server Connections

The following items summarize some key facts for using the **LASR Tables** tab to manage server connections:

■ You can create only one server connection for each SAS LASR Analytic Server component in the metadata, but you can load multiple tables into the SAS LASR Analytic Server by using a single server connection. Each server connection can have multiple libraries, and each library can have multiple tables that are associated with it.

■ When a server connection is started, it runs indefinitely without stopping. If you do not intend to use that server connection, stop the server connection. A lifetime parameter can be specified to set a time-out for server connections.

See Also
SAS Visual Analytics: Administration Guide

Start a Server Connection

SAS LASR Analytic Server components and libraries are defined in the metadata by using SAS Management Console, and server connections are started in SAS Visual Analytics Administrator. After you start a server connection that is associated with a library, you can perform any tasks that are associated with loading tables to the SAS LASR Analytic Server.

To start a server connection:

1 Click the **LASR Tables** tab. Then, expand the tree view for the SAS LASR Analytic Server component.

2 Select the ⚒ server connection that you want to start. Click ▶ .

The ● icon in the **Status** column indicates that the server connection has been activated and is now running.

Note: If you started a server connection, and you do not plan to use it, stop the server connection. Any SAS LASR Analytic Server connection that is started continues to run unless it is manually stopped, or until it reaches its lifetime limit.

See Also
SAS Visual Analytics: Administration Guide

Stop a Server Connection

To stop a server connection:

1 Click the **LASR Tables** tab. Then, expand the tree view for the SAS LASR Analytic Server component.

2 Select the ⚒ server connection that you want to stop. Click ■ .

42

Working with Tables

How Are Jobs Created and Used?

Jobs Created by SAS Visual Analytics Administrator

When tables are loaded from the **LASR Tables** tab or from the Folders window, jobs are created and saved in the same folder path where the target or output tables are saved. The following rules apply to the creation of jobs in SAS Visual Analytics Administrator:

■ When a table is loaded into the SAS LASR Analytic Server by using SAS Visual Analytics Administrator, a job is created and stored in the metadata.

■ The job is stored in the same metadata folder where the target or output table resides.

■ Jobs are listed in the folder within the metadata in the following format: *Metadata_table_name*-**Load Job**. Here is an example of a job name: Customer - Load Job.

■ When a table is loaded for the first time, a job is created in the metadata even if the table fails to load and an error message is displayed.

Jobs Created by Other Applications

If a job is saved with another application that can edit jobs, such as SAS Data Integration Studio, SAS Visual Analytics Administrator examines the job to determine whether that job can be used. This occurs when you attempt to load a table that is associated with the edited job. If the changes made to that job are incompatible with the

requirements in SAS Visual Analytics Administrator, then SAS Visual Analytics Administrator creates a different job that meets its requirements.

How Are Jobs Deployed?

If your deployment includes SAS Data Integration Studio, you can use that application to deploy the jobs for scheduling. For deployments that do not have SAS Data Integration Studio, it is still possible to export the code for the job. You can use the Schedule Manager plug-in within SAS Management Console and deploy the code as a SAS DATA Step Program. For more information about deploying jobs, see *Scheduling in SAS*.

How Are Tables Reloaded with Existing Jobs?

When a table that was previously loaded is currently in an unloaded state, you can reload that table by selecting the table in the **LASR Tables** tab. Then, you click . The table is reloaded into the SAS LASR Analytic Server and a confirmation message is displayed. The Reload icon is available if the table that is being reloaded has an associated job in the metadata.

About Loading Data to the SAS LASR Analytic Server

Table Metadata

Source tables are registered in metadata with Register Tables wizard in SAS Management Console or by using PROC METALIB. Then, you use SAS Visual Analytics Administrator to load tables to memory on the SAS LASR Analytic Server. SAS Visual Analytics Administrator uses table metadata for information about table properties such as column names and data types. Output tables, which are shown as target tables in the Load a Table window, are registered when you click **Submit** in the Load a Table window.

Requirements for Loading Tables

You can use SAS Visual Analytics Administrator to load or reload tables to the SAS LASR Analytic Server, or to unload tables. The following requirements must be fulfilled:
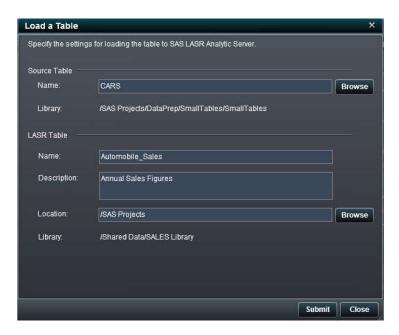
■ A minimum of one SAS LASR Analytic Server component must exist in the metadata. By default, a SAS Visual Analytics deployment creates one SAS LASR Analytic Server component in the metadata.

■ If you create a SAS LASR Analytic Server component in the metadata, a server connection is automatically created for that server component.

■ A LASR library must be associated with the server connection in the metadata. This task is performed in SAS Management Console with the Data Library Manager plug-in or with SAS Data Integration Studio. By default, a SAS Visual Analytics deployment creates the Visual Analytics Library.

■ The source table must be registered in the metadata. This task is typically performed in SAS Management Console. You must have ReadMetadata permission for the table.

■ You can have only one table with a specific SAS table name that is loaded into a LASR library at one time.

■ The metadata name for a target table must be unique within the metadata folder. This is applicable for tables that are loaded to a distributed SAS LASR Analytic Server.

■ You can have more than one metadata target table that is associated with a specific SAS table name in a LASR library. This is applicable for tables that are loaded to a distributed SAS LASR Analytic Server.

For information about the metadata permissions that are required for tables, see *SAS Visual Analytics: Administration Guide*.

Using the Load a Table Window

You can access the Load a Table window from the **LASR Tables** tab or from the Folders window. This window enables you to load or reload tables to the SAS LASR Analytic Server.

Display 42.1   Load a Table Window

The following table describes the fields:

Table 42.1   Fields in the Load a Table Window

Field	Description
Source Table	
Name	Specifies the filename of the source table in the metadata that was selected.
Library	Specifies the library in the metadata with which the source table is associated.
LASR Table	

Field	Description
Name	Enter the metadata name of the target or output table that should be loaded to the SAS LASR Analytic Server.
Description	This field is optional. Enter a description of the LASR table. When a user selects to view a list of data sources in the **Open Data Source** window within SAS Visual Analytics applications, this description is displayed for this LASR table in the list of data sources. This description is also displayed for this LASR table on the **LASR Tables** tab within SAS Visual Analytics Administrator.
Location	Click **Browse** and navigate in the folder tree to select the folder where the LASR table should be placed. The metadata name of the LASR table must be unique within that folder.
Library	Click **Browse** and select the LASR library that should be associated with this LASR table.

Reloading Tables

You can reload a table when the following is true:

- The table is in an unloaded state that is indicated by the status ■.

- When you select the table, 🔼 is available.

The table is reloaded from the original source table.

Working with Tables in the LASR Tables Tab

Load a New Table

When you load a new table from the Load a Table window, SAS Visual Analytics Administrator creates an output table in the SAS metadata, and a new job is created in

the same folder where the output table is placed. If an output table name is specified as OrionRetailCustomers, the default name for the job in the metadata is: OrionRetailCustomers-Load Job.

To load a new table:

1   Select **LASR Tables** tab from the toolbar, click, and expand the tree view for 🗄.

2   Select a library, and click ➕.

> **TIP** The ➕ icon is enabled if the server connection is running, and the library is selected.

3   In the Load a Table window, click **Browse** for the **Source Table: Name** field.

4   In the Select window, navigate through **SAS Folders**, select the table, and click **OK**.

The **Name** and **Library** fields below **Source Table** are filled.

5   In the **Name** field for **LASR Table**, enter a name for the output table. The description is optional.

Note that this description is displayed next to this LASR table when users select a data source in the Open Data Source window within SAS Visual Analytics applications. When you add a description, it provides additional information to the users about this data source.

6   Click **Browse** for the **Location** field.

7   In the Choose a Location window, navigate to the folder where you want the target or output table to be created, and click **OK**.

> **TIP** You must have ReadMetadata, WriteMetadata, and WriteMemberMetadata permissions assigned to the folder where you want to load the table.

8   Click **Submit**.

A new LASR table is registered in the SAS metadata, a job is created in the same folder within the metadata, and the table is loaded to the SAS LASR Analytic Server.

9   Click ↻. Then, click and expand the tree view for 🖳. The loaded table is displayed with the status 🟢.

Here are some additional facts about loading a table:

■   For deployments that use a co-located data provider, the SAS LASR Analytic Server performs a comparison of IP address to determine whether a parallel read is possible. The IP address that is in the SERVER= option of the generated LIBNAME statement must match the IP address of the SAS LASR Analytic Server root node.

■   When the data is loaded from a data source that is not a co-located data provider, it is transferred to the root node of the server, distributed evenly among the machines in the cluster, and then loaded into memory.

Reload an Unloaded Table

If a table in unloaded, it continues to be displayed on the **LASR Tables** tab with ◼ in the **Status** column.

To reload an unloaded table:

1   On the **LASR Tables** tab, click and expand the tree view for 🖳.

2   Select the table with the status ◼, and click ⬆.

    A progress indicator runs while the table is being loaded.

3   Click ↻ to refresh. To view the changed status of the table, click and expand the tree view for 🖳.

Unload a Table

To unload a table:

1   On the **LASR Tables** tab, click and expand the tree view for 🖳.

2   Select a table that is displayed with the status ◼.

3   Click ⬚.

A progress indicator is displayed when the table is being unloaded.

4   To view the changed status of the table, click ↻. Click and expand the tree view for
⬚.

The table is displayed in the **Status** column in an unloaded state ■.

Working with Tables in the Folders Window

Load a New Table

To load a new table from the **Folders** navigation pane:

1   In SAS Visual Analytics Administrator, expand the **Folders** navigation pane.

2   If you know the name of the table in the metadata, click 🔍 to display the Search:
Folders window and enter the name of the table. If you do not know the table name,
navigate the folder path to find the table.

3   Select the table, and right-click. In the drop-down menu, select ⬆**Load a Table**.

4   In the Load a Table window, the source table is displayed in the **LASR Table: Name**
field.

The library associated with this table is also displayed.

5   Enter the target LASR table name in the **Name** field below **LASR Table**.

The description is optional. Note that this description is displayed next to this LASR
table when users select a data source in the Open Data Source window within SAS
Visual Analytics applications. When you add a description, it provides additional
information to the users about this data source.

6   Click **Browse** for the **Location** field, and select the folder where you want the output table to be located.

7   Click **Browse** for the **Library** field, navigate the folder path to find the library for the SAS LASR Analytic Server, and select the library that should be applied to this table.

8   Click **Submit**.

9   Click ↻. Then, click and expand the tree view for 🖳. The loaded table is displayed with the status ●.

Load the Same Source Table as a Different Target Table

If a source table was loaded into the SAS LASR Analytic Server, it can be loaded again as a different target table.

To load the same source table as a different target table:

1   In SAS Visual Analytics Administrator, expand the **Folders** navigation pane.

2   If you know the name of the table, click 🔍 to display the Search: Folders window and enter the name of the table. If you do not know the table name, navigate the folder path to find the table.

3   Select the table, and right-click. In the drop-down menu, select ⬆**Load from Table**.

4   In the Load a Table window, specify the target table name in the **Name** field for **LASR Table**. The description is optional.

Note that this description is displayed next to this LASR table when users select a data source in the Open Data Source window within SAS Visual Analytics applications. When you add a description, it provides additional information to the users about this data source.

5   Click **Browse** for the **Location** field, and select the folder where you want the output table to be located.

6   Click **Browse** for the **Library** field, navigate the folder path to find the library for the SAS LASR Analytic Server, and select the library that should be applied to this table.

> **TIP** The library that you select for this table must be different from the library that was used when this table was loaded previously. You cannot use the same library for the same table multiple times.

7   Click **Submit**.

About Adding Data to Co-located Storage

Data Output for Co-located Storage

The unique value of SAS Visual Analytics Administrator is that it enables administrators to add data to a data provider that is co-located with the SAS LASR Analytic Server. SAS Visual Analytics Hadoop is a co-located data provider that is available from SAS. In this case, the Hadoop Distributed File System (HDFS) is used for data output. The purpose of adding the data to a co-located data provider is that the server can read data in parallel at very impressive rates from a co-located data provider.

In addition to SAS Visual Analytics Hadoop, Teradata Enterprise Data Warehouse and EMC Greenplum Data Computing Appliance are supported co-located data providers. When a third-party vendor appliance is used, the SAS/ACCESS interface for the database must be licensed and configured.

Output to SAS Visual Analytics Hadoop

In addition to the performance advantage of using co-located data, when SAS Visual Analytics Hadoop is used, it provides data redundancy. By default, two copies of the data are stored in HDFS. If a machine in the cluster becomes unavailable, another machine in the cluster retrieves the data from a redundant block and loads the data into memory.

SAS Visual Analytics Administrator distributes blocks evenly to the machines in the cluster so that all the machines acting as the server have an even workload. The block size is also optimized based on the number of machines in the cluster and the size of the data that is being stored. Before the data is transferred to HDFS, SAS software determines the number of machines in the cluster, row length, and the number of rows in the data. Using this information, SAS software calculates an optimal block size to provide an even distribution of data. However, the block size is bounded by a minimum size of 1 KB and a maximum size of 64 MB.

For very small data sets, the data is not distributed evenly. It is transferred to the root node of the cluster and then inserted into SAS Visual Analytics Hadoop. SAS Visual Analytics Hadoop then distributes the data in blocks according to the default block distribution algorithm.

Registered Tables

SAS Visual Analytics Administrator can load data from registered tables in metadata or SASHDAT files stored in HDFS to the SAS LASR Analytic Server. The following table describes the performance considerations for both methods:
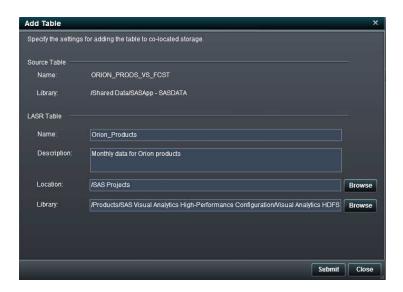
Table 42.2 *Performance Considerations*

Data Source	Advantages	Disadvantages
Registered table (from a library that does not use a co-located data provider)	Provides a rapid method for loading tables.	Appropriate for smaller data sets because the data must be transferred over the network. If the table is unloaded or the server stops, the data must be transferred over the network again.
Registered table from a library that uses a co-located data provider	Impressive performance for loading very large data sets in parallel. If a SAS LASR Analytic Server is stopped or the table is unloaded from the server, loading it again is also very fast.	Requires a separate step to add the data to HDFS or the co-located third-party vendor database and then load it before the data is available to SAS clients.

Add to HDFS Option

The Add Table window is launched by selecting the **Add to HDFS** option for a table within the folder tree in the navigation pane.

Display 42.2   *Add Table Window*

When the co-located data provider is SAS High-Performance Deployment of Hadoop, the following options are available in the Add Table window:

Table 42.3   *Fields in the Add to HDFS and Add to a Data Server Dialog Boxes*

Field	Description
Source Table	
Name	Specifies the filename of the source table in the metadata.
Library	Specifies the library in the metadata with which the source table is associated.
LASR Table	

Field	Description
Name	Enter the filename of the target or output table that should be loaded to HDFS.
Description	This field is optional. Enter a description of the target table.
	This description is displayed along with the other information for the target table on the **LASR Tables** tab. It is also displayed in SAS Visual Analytics applications when users launch the Open Data Source window. This window is displayed with a list of data sources that includes this table, and the description provided for this table.
Location	Click **Browse** and navigate in the folder tree to select the folder where the target table should be placed in HDFS.
Library	Click **Browse** and select the library that should be associated with the target table. If this is the first time you are loading this table, you can select the same library that is associated with the source table.

Loading Data to a Co-Located Provider

About Distributing Data to a Co-Located Provider

An important feature of SAS Visual Analytics Administrator is the ability to distribute data to a co-located data provider. There are two ways to distribute data to a co-located data provider.

From the **Folders** navigation pane, navigate the folder path, and select the table that you want to load. Right-click and select **Add to HDFS**. Enter the required information in the fields for specifying HDFS parameters. For deployments that use Teradata or Greenplum, right-click and select **Add to data server**. The menu choice depends on the co-located data provider.

For deployments that use a third-party vendor database as a co-located data provider, a default library for the database is registered in metadata during installation and

configuration. This library is located in SAS metadata at `/Products/SAS Visual Analytics Administrator/SAS Visual Analytics Distributed Data`. This library is configured with options that are specific to operating with a distributed database. When you create an add to data server job, remember to save the job so that the output table is registered in metadata with the distributed data library.

Load from HDFS

SAS LASR Analytic Server can load data from the Hadoop Distributed File System (HDFS) that is co-located on the machines in the cluster. Before loading the data, you must add the data to HDFS.

To load a table to HDFS from the **Folders** navigation pane:

1 On the **LASR Tables** tab, expand the **Folders** navigation pane.

2 If you know the name of the table, click 🔎 to display the Search: Folders window and enter the name of the table. If you do not know the table name, navigate the folder path to find the table.

3 Select the table, and click ➕.

4 In the Add to HDFS window, specify the target table name in the **Name** field below **LASR Table**. The description is optional. Note that this description is displayed next to this LASR table when users select a data source in the Open Data Source window within SAS Visual Analytics applications. When you add a description, it provides additional information to the users about this data source.

5 Click **Browse** for the **Location** field, and select the folder where you want the output table to be located.

6 Click **Browse** for the **Library** field, navigate the folder path to find the SAS LASR Analytic Server library, and select the library that should be applied to this table.

7 Click **Submit**.

Loading Data to a Data Server

About Adding a Table to a Co-located Data Provider

SAS LASR Analytic Server can load data from Greenplum or Teradata that is co-located on the machines in the cluster. Before loading the data, you must add the data to the data server.

When the co-located data provider is Greenplum or Teradata, the following options are available in the Add Table window. This window is launched by selecting the **Add to a Data Server** option for a table within the folder tree in the navigation pane:

Table 42.4   Fields in the Add to a Data Server Dialog Boxes

Field	Description
Source Table	
Name	Specifies the filename of the source table in the metadata.
Library	Specifies the library in the metadata with which the source table is associated.
LASR Table	
Name	Enter the filename of the target or output table that should be loaded to Greenplum or Teradata.
Description	This field is optional. Enter a description of the target table.
	This description is displayed along with the other information for the target table on the **LASR Tables** tab. It is also displayed in SAS Visual Analytics applications when users launch the Open Data Source window. This window is displayed with a list of data sources that includes this table, and the description provided for this table.

Field	Description
Location	Click **Browse** and navigate in the folder tree to select the folder where the target table should be placed in HDFS.
Library	Click **Browse** and select the library that should be associated with the target table. If this is the first time you are loading this table, you can select the same library that is associated with the source table.

Add a Table to a Data Server

To load a table to a data server (co-located data provider) from the **Folders** navigation pane:

1 On the **LASR Tables** tab, expand the **Folders** navigation pane.

2 If you know the name of the table, click 🔎 to display the Search: Folders window and enter the name of the table. If you do not know the table name, navigate the folder path to find the table.

3 Select the table, and click ➕.

4 In the Add to a Data Server window, specify the target table name in the **Name** field below **LASR Table**. The description is optional.

   Note that this description is displayed next to this LASR table when users select a data source in the Open Data Source window within SAS Visual Analytics applications. When you add a description, it provides additional information to the users about this data source.

5 Click **Browse** for the **Location** field, and select the folder where you want the output table to be located.

6 Click **Browse** for the **Library** field, navigate the folder path to find the SAS LASR Analytic Server library, and select the library that should be applied to this table.

7 Click **Submit**.

43

Working with HDFS

What Is the HDFS Content Explorer?

The HDFS content explorer is used to navigate through HDFS folders and to view tables that are stored in HDFS. It provides a tree-based view of the HDFS folder structure and a table-based view of the details for the prepared data. The HDFS content explorer is available for deployments that use SAS High-Performance Deployment of Hadoop as the co-located data provider only.

The HDFS content explorer is shown in the following display:

Display 43.1 HDFS Content Explorer

You can launch the HDFS content explorer by selecting the **HDFS** tab in SAS Visual Analytics Administrator.

What Can I Do with the HDFS Content Explorer?

When tables are available in SAS metadata, you can use SAS Visual Analytics Administrator to add tables to HDFS so that they can be loaded into memory on the SAS LASR Analytic Server. Analysts can then explore the data from the SAS Visual Analytics explorer interface. The HDFS content explorer enables administrators to view information about the prepared data such as the row count, columns, and column information.

When tables are added to HDFS with SAS Visual Analytics Administrator, they are stored with a SASHDAT file suffix. This is a special file format used by the SAS LASR Analytic Server. This special file format and the data redundancy provided by SAS High-

Performance Deployment of Hadoop enable the SAS LASR Analytic Server to read the data in parallel at very impressive rates. The data that is stored in HDFS is stored in blocks. The HDFS content explorer enables an administrator to view the block distribution, block redundancy, and measures of block utilization.

Navigating HDFS Folders

The HDFS content explorer provides a familiar and easy to use interface for navigating the files and folders in HDFS. You can navigate by clicking the folder icons in the folder tree and by using the following navigation buttons:

Back ←
    Click this button to return to the previously used folder.

Up one level ⬆
    Click this button to navigate to the parent folder.

Deleting a Table from HDFS

Tables in HDFS are deleted by using PROC DATASETS. Files that are not SASHDAT files are listed in the HDFS content explorer, but they cannot be deleted.

Viewing HDFS System Properties

To view the HDFS system properties:

1   Click the **HDFS** tab in the toolbar to access the HDFS content explorer.

2   Click ▦.

The following table describes the fields:
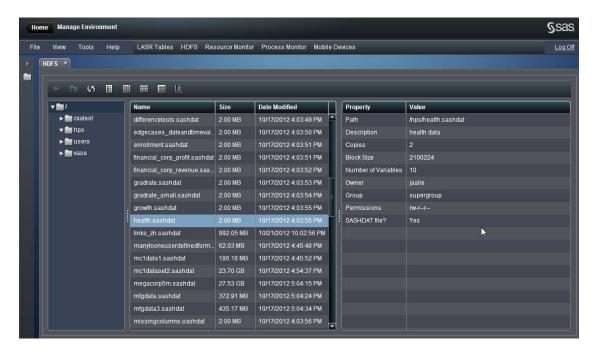
Table 43.1   *HDFS System Properties*

Property	Description
Command for setting permissions	Specifies the path to the setperm command. This command is used to enable security for files stored in HDFS by setting access permission.
Set permissions as root?	Specifies whether access permissions are enabled. For a SAS Visual Analytics deployment, the value is `Yes`.
Command for getting file information	Specifies the list command and options. For a SAS Visual Analytics deployment, the value is `ls -l {0}`.
Data directories	Specifies the directory that is used to store blocks.
NameNode	Specifies the host name of the machine that is used as the SAS High-Performance Deployment of Hadoop NameNode.
Live DataNodes Dead DataNodes	Specifies the number of SAS High-Performance Deployment of Hadoop DataNodes that are reachable and the number that are not available.

Viewing Basic File Information

To view basic file information, select a file in the HDFS content explorer. The following display shows an example:

Display 43.2 *HDFS File Information*

The following information about a file is provided:

Table 43.2 *Basic File Information*

Field	Description
Name	Specifies the name of the file.
Size	Specifies the file size. This value includes the disk space required to store the data in blocks and metadata about the file.
Date Modified	Specifies the date that the file was created or replaced.

Field	Description
Path	Specifies the HDFS directory.
Description	Specifies the description that is stored with the data. The description is displayed beside the table name in the explorer interface.
Copies	Specifies the number of redundant copies of the data.
Block size	Specifies the number of bytes that are used to store each block of data.
Number of Variables	The number of columns in the HDFS table.
Owner	Specifies the user account that added the data to HDFS.
Group	Specifies the primary UNIX group for the user account.
Permissions	Specifies the Read, Write, and Execute access permissions for owner, group, and other.
SASHDAT file?	Specifies whether the file is in the SASHDAT format. **Yes** indicates that the file is in the SASHDAT format.

Note: The HDFS content explorer can show many files for a table as it is added to HDFS. These entries disappear once the action completes.

Viewing Table Information

View Columns

To view column information, select a table in HDFS and then click ▥. The following display shows an example:

Display 43.3   *Column Information*

The following information about each column is provided:

Table 43.3   *Column Information*

Field	Description
Column Name	Specifies the column name from the source table.
Label	Specifies the label for the data set column when the table was added to HDFS.
Type	Numeric or Character. Numeric variables are encoded as **1**.
Offset	Specifies the starting position for the variable in the SASHDAT file.
Length	Specifies the storage used by the variable.
Format	Specifies the format associated with the variable.

Field	Description
Format length	Specifies the format length of the format that existed on the variable when it was added to HDFS. This value is zero if the variable did not have a format when it was added to HDFS.
Precision	Specifies the precision portion of the format for number formats.
Length (Formatted)	Specifies the length of the variable when formatting is applied.

View Row Count

To view the row count, select a table in HDFS and then click ▦.

Display 43.4   Row Count

The following information is provided:

Table 43.4   Row Count Information

Field	Description
Rows	Specifies the number of rows in the data.
Blocks	Specifies the number of HDFS blocks that are used to store the data.
Allocated	Specifies the number of bytes allocated to store the data. It is a multiple of the block size and the number of blocks. This value is smaller than the file size because it does not include the space needed for the SASHDAT file header.

Field	Description
Used	Specifies the number of bytes within the allocated blocks that are used for storing rows of data.
Utilization	Specifies the percentage of allocated space that is used for storing rows of data.

Viewing Block Information

View Block Details

To view block details, select a file in HDFS and click ⊞.

Display 43.5   *Block Details*

The following information about each block is provided:

Table 43.5   *Block Details Information*

Field	Description
Host Name	Specifies the machine in the cluster that stores the block of data.
Block Name	Specifies the filename for the block.
Path	Specifies the directory to the block.

Field	Description
Record Length	Specifies the sum of the column lengths for the variables in the data.
Records	Specifies the number of rows stored in the block. Because redundant blocks are listed in the table, the sum of the records listed does not equal the number of rows in the data.
Owner	Specifies the user account that added the data to HDFS.
Group	Specifies the primary UNIX group for the user account that stored the data.
Permissions	Specifies the Read, Write, and Execute access permissions for owner, group, and other.

You can sort by the column headings to identify anomalies. It is normal for several blocks to be stored on the same machine. However, it is not normal for the values of **Record LengthOwner**, **Group**, or **Permissions** to be different from row to row.

View Block Distribution

The files added to HDFS are stored as blocks. One block is the preferred block, and additional copies of the blocks are used to provide data redundancy. The Block Distribution dialog box offers two ways to view this information. The **Block Detail View** tab enables you to select a block number and view the host names that store the original or redundant blocks. The **Node Detail View** enables you to select a host name and view the block numbers that are stored on the machine.

To view the block distribution, select a table in HDFS and then click ⅏.

Display 43.6 *Block Distribution*

The following information about the block distribution is provided:

Table 43.6 *Block Distribution Information*
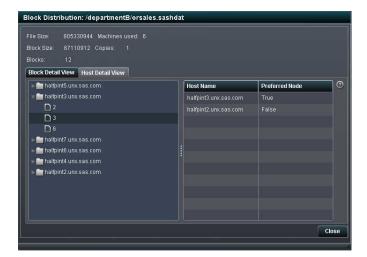
Field	Description
File size	Specifies the size of the file in bytes.
Block size	Specifies the block size for the file.
Blocks	Specifies the number of blocks used to store the original copy of the data.
Copies	Specifies the number of redundant block copies of the data.
Machines used	Specifies the number of machines in the cluster that have original or redundant blocks for the file.

On the **Block Detail View** tab, you can select a block number to view how many copies of the block exist and the host names for the machines that store the blocks. The value in the **Total Copies** column equals the number of redundant copies of the block plus

the original block. You can select the column heading to sort the rows. In an ideal distribution, the number of total copies is equal for all blocks.

On the **Node Detail View** tab, you can expand a host name node and then view the block numbers that are stored on that machine. When you select the block number, this host name and any additional machines with copies of the block are identified in the host name list. The following display shows an example:

Display 43.7   Block Distribution with Node Detail View Tab

44

Monitoring Resources

Utilization History

Information about the resource utilization is available in the **Resource Monitor** tab. The resource utilization is displayed in two graphs, the Utilization History and the Real-Time View. The graphs are updated according to a sampling rate. The sampling rate is set in the application preferences.

The Utilization History graph plots the utilization for the following resources against time:

- CPU

- memory

- network input/output (I/O)

The CPU and memory utilization are plotted as a percentage of 100% capacity, though the graph scales automatically. Under periods of high demand, the upper bound can reach 100%. Under periods of low demand, the upper bound can drop below 10%.

The network I/O utilization is shown as one line plot for input and one line plot for output. The plot shows the transfer rate in units of megabytes per second.

To view the resource utilization for one sampling period, place your pointer over a line. You can also click on a line to select the same sampling period for all the resource line

plots that are shown. This way, you can place your pointer over each of the selected sampling periods to view the utilization for each resource in the same sampling period.

Figure 44.1 *Resource Monitor*

Note: If the Resource monitor does not display the line plot or the Real-Time View, then contact your system administrator to make sure that the SAS LASR Analytic Server has been started in the **LASR Tables** tab and that it is running.

Real-Time View

The Real-Time View of resource utilization is provided below the Utilization History. The display is a heat map for each of the monitored resources, CPU, memory, and network I/O. Each column in the heat map represents a machine in the cluster. To view the real-time resource utilization for a resource type, position your pointer over one of the cells in the heat map.

For the section that shows CPU utilization, the tooltip shows the host name, CPU number, and the CPU utilization as a percentage. For the section that shows memory utilization, the tooltip shows the host name and memory utilization as a percentage. Both of these resources have a range of zero to 100. The color changes to represent the work load.

The bottom section shows two rows to represent network I/O. The top row shows the transfer rate for output and the bottom row shows the transfer rate for input. The tooltip shows the host name, the direction, and the rate. The color for the icons depends on the transfer rate between 0 and 25 megabytes per second. If the transfer rate exceeds 25 MB per second, then the icon color changes to purple.

45

Monitoring Processes

Viewing Memory and CPU Utilization

Information about the per-process resource utilization is available in the **Process Monitor** tab. The resource utilization for memory and CPU are displayed in two graphs, the Memory Utilization versus CPU Utilization and CPU and Memory Utilization by Instance. The graphs are updated according to a sampling rate. The sampling rate is set in the application preferences.
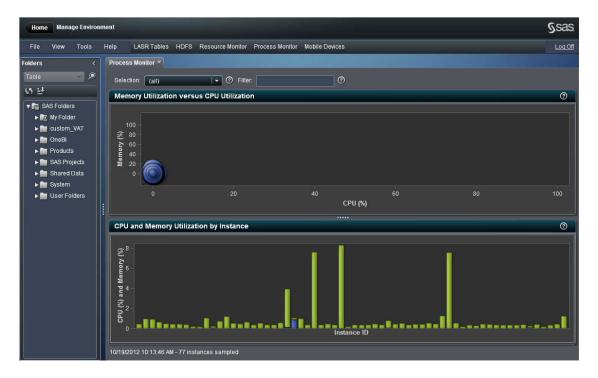
The **Selection** menu enables you to filter the display of application instances in the graphs for top resources consumers.

Each application instance is represented by a bubble in the bubble plot. The size of the bubble represents the number of processes used for the application instance. If an instance has one process for each machine in the cluster, then that ensures ideal resource utilization and performance.

The location of the bubble indicates the resource utilization for the application instance on the machines in the cluster. Place your pointer over a bubble to view information about the application instance. Select a bubble to view the detailed resource consumption for each machine that is used by the application instance.

The location of the bubble indicates the resource utilization for the application instance on the machines in the cluster. Place your pointer over a bubble to view information about the application instance. Select a bubble to view the detailed resource consumption for each machine that is used by the application instance.

Figure 45.1 *Process Monitor*

Note: If no application instances appear in the process monitor, then contact your system administrator to make sure that the SAS LASR Analytic Server has been started in the **LASR Tables** tab and that it is running.

CPU and Memory Utilization by Instance

The CPU and Memory Utilization by Instance graph shows a bar chart. Each application instance is represented by a bar. The lower part of the bar shows the CPU utilization for the application instance. The upper part of the bar shows the memory utilization for the application instance.

Place your pointer over a bar and click to view information about the application instance. Here is a description of the information that is presented:

The following table describes the fields:

Table 45.1   *Fields in the Panel with Details about CPU Utilization*

Field	Description
Host Name	Specifies the host name of the SAS LASR Analytic Server.
Rank	Specifies the ranking of a problem according to...
	When you click the **Server Port**, **Memory (%)**, **CPU (%)**, or the **Current Run Time**, the data in the **Rank** column is changed.
Server Port	The port number associated with the SAS LASR Analytic Server.
Memory (%)	Specifies the amount of memory that is being consumed by each server component.
CPU (%)	Specifies the amount of CPU that is being consumed by each server component.
Current Run Time	Specifies the current run time.

In the same window, click **Show Tables** to view the following information:

The following table describes the fields that are displayed when you click **Show Tables**. You can reorder the presentation of information by clicking on any one of the column labels.

Table 45.2   *Fields in the Window with Details Tables*

Field	Description
Tables	Specifies the tables that are loaded into the SAS LASR Analytic Server.
Owner	Specifies the name of the administrator who loaded the table.
Rows	Specifies the number of rows in each table.

Field	Description
Variables	Specifies the number of columns in each table.

46

Managing Mobile Devices

Features in SAS Visual Analytics Mobile Devices Administration

This chapter addresses the use of the SAS Visual Analytics Mobile Devices to manage mobile device users who use the SAS Visual Analytics suite of applications. This application includes the following administration features for managing mobile devices:

■ View data about mobile devices that are set up for users to view reports or explorations that were created in SAS Visual Analytics. The displayed data includes user ID, device type, device model, device ID, OS version, application version, and timestamp of last access.

■ Sort the data for mobile devices in ascending or descending order.

■ Filter data for mobile devices. For example, you can filter data by the mobile device type or by the mobile device model.

■ Blacklist lost or stolen mobile devices.

■ Remove blacklisted mobile devices from the blacklist.

Managing Mobile Devices

The SAS Visual Analytics Mobile Devices application is easy to use. To perform tasks such as viewing, sorting, filtering, or blacklisting a mobile device:

1 Log on to SAS Visual Analytics Administrator.

2 Select the **Mobile Devices** tab.

3 On the **Mobile Devices** tab, select **Logon History**.

4 To sort the items in the columns in either ascending or descending order, click the column heading for any column on this page.

5 To filter and view selected data about mobile devices:

   a From the **Filter** menu, select an item (for example, Device Type). Then, enter a valid value in the field (for example, iPad), and click **Apply**.

   The value that you specified for the filter determines the type of data that is displayed.

   b To view the entire history of events (for example, various application versions that were used by devices) associated with the item that you specified in the **Filter**, field, select the **Include Device History** check box.

6 To add a mobile device to the blacklist:

   a Click **Manage Blacklist**.

   b Click ✚.

   You are prompted by the Add New Device ID window.

   c Enter the device ID and click **OK**.

   d To view the blacklisted device, click **Blacklist History**.

The mobile device is added to the list of blacklisted devices on the Blacklist History page. The **Action** column indicates the status of this device (whether it as blacklisted or not), and the **Result** column indicates whether your action was successful. Blacklisted mobile devices cannot be used to access the SAS Visual Analytics suite of applications.

7   To remove a mobile device from the blacklist, click **Manage Blacklist**.

8   On the Manage Blacklist page, select the device and click ✖.

9   When prompted by a dialog box that asks you to confirm your action, click **Yes**.

10 To view blacklisted mobile devices or confirm that a device was blacklisted, click **Blacklist History**.

11 To search for a blacklisted mobile device either by the device ID or the administrator ID, click **Manage Blacklist**. Then, select **Admin ID** or **Device ID** from the drop-down menu for **Filter**. Click **Apply**.

The mobile devices that meet the criteria that you specified are displayed.

The following information is provided in the Blacklist History page.

Table 46.1   Fields in the Blacklist History Page

Field	Description
Admin ID	User name of the administrator who blacklisted the device.
Device ID	Name of the blacklisted device.
Action	Type of action that was performed on the mobile device.
Result	Indicates whether the attempt to blacklist a mobile device was successful.
Timestamp	Displays the date and time when a specific mobile device was blacklisted.

47

Setting Permissions

About Managing Security for Folders and Tables

SAS Visual Analytics Administrator enables you to manage metadata-based security for folders and tables. Tables that have been loaded into the SAS LASR Analytic Server or a co-located database represent LASR tables. Metadata representations for tables exist in the system, in memory within the SAS LASR Analytic Server, and in databases associated with co-located providers that include HDFS, Teradata, or Greenplum.

The Authorization page is accessible when a table or a folder is selected from the navigation pane in SAS Visual Analytics Administrator. In this page, you can set permissions that control a specific group or user's access to a selected table or folder.

Here are some tasks that you can perform in the Authorization page to control access to a table.

Table 47.1   Tasks Performed in the Authorization Page
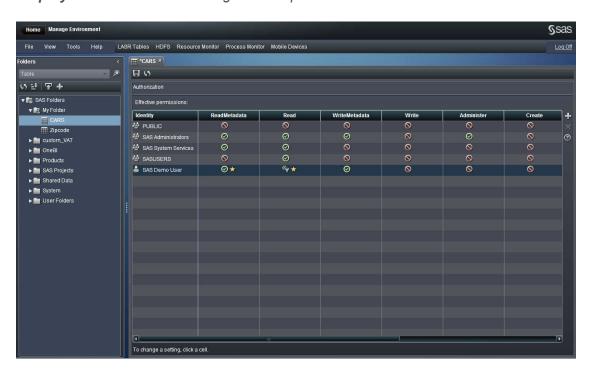
Authorization Page in SAS Visual Analytics Administrator	Applies to all Tables in Metadata	Applies Only to LASR Tables
Set an explicit grant or denial to a table for a specific group or user. Explicit grants are specific to the table, and they are not inherited. A yellow star ★ is displayed to indicate an explicit control. If an explicit denial was set for the table by denying the ReadMetadata permission, this table is not available for the group or user in the list of data sources within SAS Visual Analytics Explorer or SAS Visual Analytics Designer.	✔	✔
Remove an explicit grant or denial to a table for a specific group or user.	✔	✔

Authorization Page in SAS Visual Analytics Administrator	Applies to all Tables in Metadata	Applies Only to LASR Tables
Specify a conditional grant that allows access to specific rows in a LASR table. Conditional grants can be assigned only to a LASR table that is selected and displayed in the Authorization page, and are specified for the **Read** permission. A filter icon is displayed to indicate a conditional grant, and the user or group is allowed to see only the rows for which they have been given permissions.		✓

Here is a display of the Authorization page where an explicit grant and a conditional grant have been applied to a LASR table for SAS Demo User.

Display 47.1   Authorization Page with Explicit and Conditional Grants for a LASR Table

Introduction to Permissions

Permissions that you set on a table's Authorization page are part of a metadata-based access control system within the SAS Metadata Server.

The following table shows the permissions and explicit grants or denials that are relevant to SAS Visual Analytics, and where you can set them.

Table 47.2   Relevant Permissions for SAS Visual Analytics

Permission (Abbreviation)	Affected Actions	Authorization Manager in SAS Management Console	Authorization Page in SAS Visual Analytics Administrator
ReadMetadata (RM)	View an object or folder. For example, to see an exploration, report, stored process, table, or folder, you need ReadMetadata permission for that object.	Set permissions for explorations, reports, stored processes, tables or folders.	Set permissions for tables or folders.
Read (R)	Read data. For example, to see any data within a LASR table, you need the Read Permission for that table.	Set permissions for tables.	Set permissions for tables. Specify a conditional grant that allows access to specific rows in a LASR table.*

Permission (Abbreviation)	Affected Actions	Authorization Manager in SAS Management Console	Authorization Page in SAS Visual Analytics Administrator
WriteMetadata (WM)	Edit, delete, or set permissions for an object. To assign a library to a SAS LASR Analytic Server connection, you need WriteMetadata permission for the library. To delete an object, you also need the WriteMemberMetadata permission for the object's parent folder.	Set permissions for libraries. Set explicit grants or denials for objects.	Set permissions for tables.
WriteMemberMetadata (WMM)	Add an object to a folder or delete an object from a folder. To enable a user to interact with a folder's contents but not with the folder itself, grant the WriteMemberMetadata permission and deny the WriteMetadata permission.	Set permissions for folders.	Select a folder in the **Folders** navigation pane next to the Authorization page, and set permissions.

*   Row-level security for LASR tables can be assigned only in SAS Visual Analytics Administrator.

See Also
SAS Visual Analytics: Administration Guide

The Authorization Page

About the Authorization Page

Each table or folder's Authorization page displays effective permissions for that table or folder, indicates any direct access controls, and enables you to add or remove explicit controls. Only permissions that are potentially applicable for an object type are shown for objects of that type.
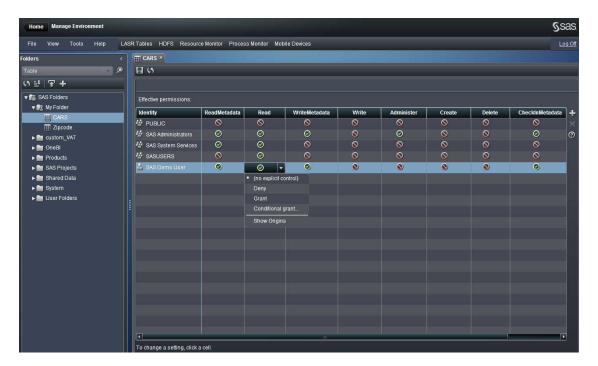
Only identities that participate in access controls that potentially affect the current object are persisted on this page. The following identities are persisted:

- identities that are listed in the permission pattern of the repository access control template (ACT)

- identities that have direct controls on an object that is a parent to this object

- identities that have direct controls on this object

Note: Effective permissions are a calculation of the net effect of all applicable metadata-layer permission settings. Effective permissions do not encompass role-based constraints or constraints from other authorization layers.

Here is an example of the Authorization page that shows the SAS Demo User's authorization properties for a LASR table.

Display 47.2 Authorization Page

Grant, Conditional Grant, and Denial Icons

The following table shows the icons associated with conditional grants in the Authorization page.

Table 47.3 Grant, Conditional Grant, and Denial Icons

Icon	Meaning
⊘	Grant
🔽	Conditional grant (a grant that is constrained by a permission condition, in row-level security)
🚫	Denial

Direct Control Indicators

The Authorization page uses the following icons to provide immediate information about the source of each setting.

Table 47.4   *Direct Access Controls*

Icon	Term	Meaning
★	Direct control: Explicit	The direct access control is set on the current object and specifically assigned to the selected identity.
▦	Direct control: ACT	The direct access control comes from an applied access control template (ACT) whose pattern specifically assigns the grant or denial to the selected identity.
(none)	Indirect setting	The setting comes from someone else (a parent group), somewhere else (a parent object), or special status (such as unrestricted). For the WriteMemberMetadata permission, indirect means that the setting mirrors the WriteMetadata setting.

TIP The explicit and ACT indicator icons correspond to the white and green colors on the **Authorization** tab in SAS Management Console. If both an explicit control and an applied ACT setting are present, only the explicit indicator is visible.

TIP For additional details about the source of a setting, select **Show Origins** from the drop-down list for that setting. See "Permission Origins" on page 455.

Icon Combinations

The following table shows all of the possible combinations of icons within a cell on a table's Authorization page.

Table 47.5 *Icon Combinations*

Icon	Description
⊘★	Grant from an explicit control
⊘■	Grant from an applied ACT
⊘	Grant from an indirect source (such as a parent group or parent object)
⊘▼★	Conditional grant from an explicit control
⊘▼	Conditional grant from an indirect source (a parent group)
⊘★	Denial from an explicit control
⊘■	Denial from an applied ACT
⊘	Denial from an indirect source (such as a parent group or parent object)

Add an Explicit Grant or Denial

To add an explicit grant or denial to a table or a folder:

1 Log on to SAS Visual Analytics Administrator, and expand the navigation pane for folders.

2 In the folder tree within the navigation pane, locate the table that you want to protect or make available.

3 Right-click the table or folder, and select **Authorization**.

The authorization properties are displayed.

4   On the Authorization page, locate the user or group that you want to assign an explicit control to. If the user or group is not listed, click ╋ to open the Add Identities window.

   Note: An explicit grant of the ReadMetadata permission is automatically set for each identity that you add.

5   Click a cell and make a selection from the drop-down list.

   Note: If the selected identity is an unrestricted user, all permissions are granted and you cannot' make changes.

   Note: When you click outside the cell, a yellow star ★ (indicating an explicit control) is displayed in the cell that you updated.

6   If you changed a group's access, review the impact on all of the listed identities.

   Note: This is important because controls that you add for a group can affect access for all members of that group. For example, an explicit denial that you add for the PUBLIC group blocks access for all restricted users, unless there are also explicit (or direct ACT) grants. You must offset a broad explicit denial with explicit (or direct ACT) grants for any restricted identities whose access you want to preserve.

   The effective permissions are updated immediately when the group permissions are modified.

7   Save the table.

   Note: The effective permissions are updated immediately when the group permissions are modified.
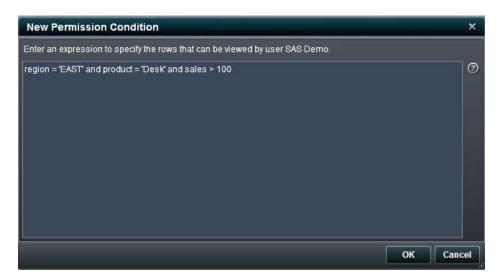
TIP To apply or remove access control templates (ACTs), use SAS Management Console.

Row-level Security for LASR Tables

About Row-level Security

Row-level permission conditions are specified in the New Permission Condition window for LASR tables. Here is a display of the New Permission Condition window where you can specify the syntax for row-level permissions that apply to a LASR table:

Display 47.3   *New Permission Condition Window*

New Permission Condition   ✕

Enter an expression to specify the rows that can be viewed by user SAS Demo.

region = 'EAST' and product = 'Desk' and sales > 100   ⑦

OK   Cancel

CAUTION! **The syntax that you enter and save in the New Permission Condition is not checked for validity. Make sure that the syntax that you enter is correct.**

Syntax Rules for Row-level Permissions

The following table provides syntax guidelines and examples for row-level permissions that can be applied to a LASR table. These are typically the most commonly used expressions. However, note that this is not an exhaustive list of all possible expressions.

Table 47.6 *Syntax Guidelines and Examples for Row-Level Permission Conditions*

Syntax Guidelines	Examples

Do not include WHERE for numeric values or text variables.	Correct:
	`Toy_Type="dolls"`
	`Year=2012`
	`Toy_Price=20`
	Incorrect:
	`WHERE Year="2012"`
	`WHERE Toy_Price="20"`

Non-numerical character values must be embedded within double quotation marks.	`Customer_Country="US"` `Toy_Type="animals"`
Specify OR as a logical operator.	`Customer_Country="US" OR Customer_Country="UK"` `Toy_Type="cars" NOT Toy_Type="dolls"`
Specify AND as a logical operator.	`Customer_Country="US" AND Customer_Country="UK"` `AND Customer_Country="FR"`
Specify NOT as a logical operator.	`Toy_Type="animals" NOT Toy_Type=dolls NOT Toy_Type=cars`
Specify IN for groups of values that need to be checked.	`Toy_Type=("dolls" "cars" "animals")`
Specify CONTAINS to include a value.	`Toy_Type CONTAINS "cars"`
Specify BETWEEN to include a range of values.	`Toy_Price BETWEEN 10 and 20`
Specify LIKE	`Toy_Type LIKE "dolls"`

Syntax Guidelines	Examples
Specify =	`Toy_Price=25`
Specify >	`Toy_Price>25`
Specify <	`Toy_Price<25`
Specify >=	`Toy_Price>=25`
Specify <=	`Toy_Price<=25`
Specify <>	`Toy_Price<>25`
Specify NE	`Toy_Price NE 30`
Specify NOT	`Toy_Price NOT=30`
Specify NOTIN	`Toy_Price NOTIN (3,600)`
Specify IN	`Toy_Price IN (30)`
Specify BETWEEN	`Toy_Price BETWEEN 20 AND 30`
Specify NOT BETWEEN	`Toy_Price NOT BETWEEN 20 AND 30`
Specify ^	`Toy_Price^=30`
Specify DESCRIP CONTAINS	`DESCRIP CONTAINS "Doll"`
Specify DESCRIP NOT CONTAINS	`DESCRIP NOT CONTAINS "Animal"`
Specify DESCRIP ?	`DESCRIP ? "Animal"`
Specify DESCRIP ?	`DESCRIP ? "Animal"`
Specify DESCRIP LIKE	`DESCRIP LIKE "Car"`
Specify DESCRIP NOT LIKE	`DESCRIP NOT LIKE "Doll"`

Note the following exceptions:

■ Use of | | for an expression with OR cannot be used.

■ Expressions with months or dates cannot be used.

Set Permission Conditions for a LASR Table

From the **LASR Tables** tab, you can select **Conditional Grant** only for LASR tables. You cannot select **Conditional Grant** for all tables in the metadata.

To set a permission condition for a LASR table:

1   Log on to SAS Visual Analytics Administrator, and expand the navigation pane for folders.

2   In the navigation pane, locate the LASR table whose row-level security you want to define.

3   Right-click the table and select **Authorization**. The table's authorization properties are displayed.

4   In the **Read** column, click the cell for the identity that you want to assign the condition to.

Note: If the user or group is not listed, click ✚ (on the right edge of the table) to open the Add Identities window.

Note: An explicit grant of the ReadMetadata permission is automatically set for each identity that you add.

5   From the cell's drop-down list, select **Conditional grant**.

Note: If **Conditional grant** is already selected, a condition already exists. Selecting **Conditional grant** enables you to view or update the condition.

6   In the New Permission Condition window, enter and save the condition.

Note: When you click outside the cell on the Authorization page, a conditional grant icon 🌡 is displayed in the cell that you updated.

> **CAUTION!** **The syntax that you enter and save in the New Permission Condition is not checked for validity. Make sure that the syntax that you have entered is correct.**

7   Save the table object.

See Also

SAS Visual Analytics: Administration Guide

Permission Origins

Introduction

The permission origins feature identifies the source of each effective permission. Permission origins answers the question: Why is this identity granted (or denied) this permission?

In the origins answer, only the controlling (winning, highest precedence) access control is shown. If there are multiple tied winning controls, they are all shown. Other, lower precedence controls are not shown in the origins answer.

Origins answers are available on an object's **Authorization** page. Click a setting and select **Show Origins** from the drop-down list. The origins answer for that effective permission is displayed in a small window.

Simple Permission Origins

The following table provides simple examples of permission origins answers. In each example, we are interested in why UserA has an effective grant on FolderA. In each example, UserA is a direct member of both GroupA and GroupB. Each row in the table is for a different (independent) permissions scenario. In the table, the first column depicts the contents of the Origins window. The second column interprets the information.

Table 47.7 *Origins: Simple Examples*

Origins Information	Source of UserA's Effective Grant on FolderA
⊘👤 UserA [Explicit]	On FolderA, an explicit grant for UserA
⊘👥 GroupA [Explicit]	On FolderA, an explicit grant for GroupA
⊘👥 GroupA [Explicit] ⊘👥 GroupB [Explicit]	On FolderA, explicit grants for GroupA and GroupB **Note:** Two settings are shown because they are tied and they both win (UserA is a direct member of GroupA and GroupB).
⊘👥 GroupA [ACT: GroupARead]	On FolderA, an ACT pattern grant for GroupA (from a a directly applied ACT)
⊘👥 SASUSERS [ACT: GenRead]	On FolderA, an ACT pattern grant for SASUSERS (from a directly applied ACT)
⊘👥 GroupA [ACT: GroupARead] ⊘👥 GroupB [ACT: GroupBRead]	On FolderA, ACT pattern grants for GroupA and GroupB (from two different directly applied ACTs). **Note:** Two settings are shown because they are tied and they both win (UserA is a direct member of GroupA and GroupB).
⊘👥 GroupA [ACT: GroupABRead] ⊘👥 GroupB [ACT: GroupABRead]	On FolderA, ACT pattern grants for GroupA and GroupB (from the same directly applied ACT). **Note:** Two settings are shown because they are tied and they both win (UserA is a direct member of GroupA and GroupB).
⊘👤 UserA is unrestricted.	UserA's status as an unrestricted user (someone who is unrestricted is always granted all permissions)

Inherited Permission Origins

In many cases, the controlling setting is not on the current object. Instead, the controlling setting is defined on a parent object and inherited by the current object.

The following table provides examples in which the controlling setting comes from a parent object. Because the source of the effective permission is a parent object, the answer must identify which parent object has the controlling setting. For this reason, the origins answers in the following examples identify both a particular parent object (the object that has the controlling setting) and the controlling setting itself.

In each example, we are interested in why UserA has an effective grant on FolderA. In each example, UserA is a direct member of both GroupA and GroupB. Each row in the table is for a different (independent) permissions scenario. In the table, the first column depicts the contents of the Origins window. The second column interprets the information.

Table 47.8   Origins: Inheritance Examples

Origins Information	Source of UserA's Effective Grant on FolderA
ParentFolderA 　✅👤 UserA [Explicit]	On ParentFolderA, an explicit grant for UserA
ParentFolderA 　✅👥 GroupA [Explicit]	On ParentFolderA, an explicit grant for GroupA
ParentFolderA 　✅👥 GroupA [Explicit] 　✅👥 GroupB [Explicit]	On ParentFolderA, explicit grants for GroupA and GroupB
ParentFolderA 　✅👥 GroupA [ACT: GroupARead]	On ParentFolderA, an ACT pattern grant for GroupA (from a directly applied ACT)

Origins Information	Source of UserA's Effective Grant on FolderA
📁 GreatGrandParentFolderA 　✅👥 SASUSERS [ACT: GenRead]	On GreatGrandParentFolderA, an ACT pattern grant for SASUSERS (from a directly applied ACT)
📁 ParentFolderA 　✅👥 GroupA [ACT: GroupARead] 　✅👥 GroupB [ACT: GroupBRead]	On ParentFolderA, ACT pattern grants for GroupA and GroupB (from two different directly applied ACTs)
📁 GrandParentFolderA 　✅👥 GroupA [ACT: GroupABRead] 　✅👥 GroupB [ACT: GroupABRead]	On GrandParentFolderA, ACT pattern grants for GroupA and GroupB (from the same directly applied ACT).

Part 7

Appendix

Appendix 1

Operators for Calculated Data Items

Overview of Operators for Calculated Data Items

In SAS Visual Analytics Explorer and SAS Visual Analytics Designer, you can calculate data items by using expressions that contain operators. The operators that you can use are grouped into the following categories:

Numeric (simple) operators
perform basic mathematical operations such as addition and multiplication. See "Numeric (Simple) Operators" on page 462.

Numeric (advanced) operators
perform advanced mathematical functions such as evaluating logarithms and truncating decimals. See "Numeric (Advanced) Operators" on page 462.

Boolean operators
evaluate whether a set of logical conditions is true. See "Boolean Operators" on page 464.

Comparison operators
evaluate how the values of your data items compare. See "Comparison Operators" on page 465.

Date and Time operators
convert values to, from, or between date and time formats. See "Date and Time Operators" on page 466.

Numeric (Simple) Operators

-x (Negation)
　　returns a value with the opposite sign of the input value.

　　For example, -1 returns 1 and 1 returns -1.

x - y (Minus)
　　subtracts the second value from the first value.

　　For example, (2,1) returns 1.

x * y (Multiply)
　　multiplies the first and second values together.

　　For example, (2,3) returns 6.

x / y (Divide)
　　divides the first value by the second value.

　　For example, (6,2) returns 3.

x + y (Plus)
　　adds the first and second values together.

　　For example, (1,2) returns 3.

Numeric (Advanced) Operators

|x| (Absolute value)
　　returns the absolute value of the input value.

　　For example, -3 returns 3.

Ceil (Ceiling)
　　rounds the input value up to the nearest integer.

　　For example, 4.2 returns 5 and -4.8 returns -4.

Exp (e**exponent)
  raises the constant *e* to the power specified by the input value.

  For example, 5 returns *e* to the 5th power (148.41).

Floor (Floor)
  rounds the input value down to the nearest integer.

  For example, 4.8 returns 4 and -4.2 returns -5.

Ln (Natural log)
  returns the natural logarithm (base *e*) of the input value.

  For example, 10 returns the *e*th root of 10 (2.30...).

Log (Log)
  returns the logarithm of the first value, where the second value specifies the base.

  For example, (64, 8) returns the base 8 logarithm of 64 (2).

Mod (Remainder)
  returns the remainder after dividing the first value by the second value.

  For example, (5,2) returns 1.

Power (Base**Exponent)
  raises the first value to the power of the second value.

  For example, (5,2) returns 5 to the 2nd power (25).

Root (Root of an Expression)
  returns the *n*th root of the first value, where the second value specifies *n* (the base of the root).

  For example, (27,3) returns the 3rd (cube) root of 27 (3).

Round (Round the Number)
  round 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, 2) returns 7.35.

Note: If you select 0 decimal places, then the values are rounded to the nearest integer.

Trunc (Truncate to Whole Number)
> truncates the input value to an integer.
>
> For example, 8.9 returns 8 and -8.9 returns -8.

Boolean Operators

And (Logical 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 (Conditional Return Value)
> 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 (Logical 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.

Comparison Operators

Between (Number in Range)
  returns true if the first value is between the second and third values (inclusive.)

  For example, (X, 50, 100) returns true if X is between 50 and 100.

In List
  returns true if the first value is in the list specified by the second parameter. To select your list, assign a data item to the first parameter and then choose the values from the drop-down list.

  For example, X In (1,2,3) returns true when the value of X is either 1, 2, or 3.

  Note: This operator cannot be used to compare measures.

  Note: This operator is available only in SAS Visual Analytics Explorer.

Missing (Is Missing Value)
  returns true if the value is a missing value.

Not in List
  returns true if the first value is not in the list specified by the second parameter. To select your list, assign a data item to the first parameter and then choose the values from the drop-down list.

  For example, X NotIn (1,2,3) returns true when the value of X is not 1, 2, or 3.

  Note: This operator cannot be used to compare measures.

  Note: This operator is available only in SAS Visual Analytics Explorer.

NotMissing (Is not Missing Value)
  returns true if the value is not a missing value.

Note: This operator is available only in SAS Visual Analytics Explorer.

x != y (Not Equal)
   returns true if the first value is not equal to the second value.

   Note: This operator cannot be used to compare measures.

x < y (Less Than)
   returns true if the first value is less than the second value.

x <= y (Less Than or Equal)
   returns true if the first value is less than or equal to the second value.

   Note: This operator cannot be used to compare measures.

x = y (Equal)
   returns true if the first value is equal to the second value.

   Note: This operator cannot be used to compare measures.

x > y (Greater Than)
   returns true if the first value is greater than the second value.

x >= (Greater Than or Equal)
   returns true if the first value is greater than or equal to the second value.

   Note: This operator cannot be used to compare measures.

Date and Time Operators

DateFromMDY
   creates a date value from separate month, day, and year values. The first value
   specifies the month as a number between 1–12. The second value specifies the day
   as a number between 1–31. The third value specifies the year as a four-digit
   number.

   For example, (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 between 1–4.
>
> For example, (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, January 15, 2013 05:15 PM 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 between 0–23. The third value specifies the minute as a number between 0–59. The fourth value specifies the second as a number between 0–59.
>
> For example, (15JAN2013, 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 between 1–12. The third value specifies the day as a number between 1–31. The fourth value specifies the year as a four-digit number.
>
> For example, (05:15:23 PM, 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 between 1–31.
>
> For example, 15JAN2013 returns 15.

DayOfWeek
> returns the day of the week from a date value as a number between 1–7 (1 is Sunday.)
>
> For example, 15JAN2013 returns 3 (Tuesday.)

DayOfYear
> returns the day of the year from a date value as a number between 1–366.

> For example, 15FEB2013 returns 46.

Hour
> returns the hour from a time or datetime value as a number between 0–23.

> For example, 05:15:23 PM returns 17.

Minute
> returns the minute from a time or datetime value as a number between 0–59.

> For example, 05:15:23 PM returns 15.

Month
> returns the month from a date value as a number between 1–12.

> For example, 15JAN2013 returns 1.

Now
> creates a datetime value from the current date and time.

Quarter
> returns the quarter from a date value as a number between 1–4.

> For example, 15AUG2013 returns 3.

Second
> returns the second from a time or datetime value as a number between 0–59.

> For example, 05:15:23 PM returns 23.

TimeFromHMS
> creates a time value from separate hour, minute, and second values. The first value specifies the hour as a number between 0–23. The second value specifies the minute as a number between 0–59. The third value specifies the second as a number between 0–59.

> For example, (17,15,23) returns 05:15:23 PM.

TimePart
> converts a datetime value to a time value.

For example, January 15, 2013 05:15:23 PM returns 05:15:23 PM.

WeekNumber

returns the week of the year as a number between 1–53, where week 2 begins on the second Sunday of the year.

For example, 04AUG2013 returns 31.

Year

returns the year from a date value as a four-digit number.

For example, 15JAN2013 returns 2013.

Appendix 2

Data Limits for SAS Visual Analytics Explorer

Some of the visualizations in SAS Visual Analytics Explorer have limits to the number of data values that they can display. The limit values are affected by the **Visual data threshold** setting in the Preferences window. For more information about the Preferences window, see "Specify Preferences That Are Specific to SAS Visual Analytics Explorer" on page 111.

The following table displays the data limits for visualizations in SAS Visual Analytics Explorer:

Table A2.1   Data Limits for SAS Visual Analytics Explorer

Visualization Type	Variation	Behavior when Limit Is Exceeded	Default Threshold	Minimum Threshold	Maximum Threshold
Table	—	For sorted tables only, paging is applied to the table.	10,000	1,000	100,000
Crosstab	—	An error message appears.	41,000	9,500	50,000

Visualization Type	Variation	Behavior when Limit Is Exceeded	Default Threshold	Minimum Threshold	Maximum Threshold
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 values.	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.	An error message appears.	25,000	7,500	30,000
	Grouped.	The bubble plot shows the top or bottom values 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
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	—	An error message appears.	5,000	1,500	6,000
Treemap	No additional levels.	The treemap shows a subset that contains the first or last ranked values.	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 values.	4,900	1,050	6,000
Correlation Matrix	—		60 measures		

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 details about the

server data limits, see the topic "Manage High-Cardinality Data" in the *SAS Visual Analytics: Administration Guide*.

Glossary

access control template

a reusable named authorization pattern that you can apply to multiple resources. An access control template consists of a list of users and groups and indicates, for each user or group, whether permissions are granted or denied. Short form: ACT.

ACT

See access control template

Apache Hadoop

a framework that allows for the distributed processing of large data sets across clusters of computers using a simple programming model.

authorization

the process of determining the permissions that particular users have for particular resources. Authorization either permits or denies a specific action on a specific resource, based on the user's identity and on group memberships.

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

See box-and-whisker plot

box-and-whisker 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.

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

See crosstabulation table

crosstabulation table

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.

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.

progressive chart
 See waterfall chart

query
 a set of instructions that requests particular information from one or more data sources.

report
 output that is generated by running custom SAS code against the data in your project.

role
a set of capabilities. In some applications, certain actions are available only to users or groups that have a particular role.

SAS Management Console
a Java application that provides a single user interface for performing SAS administrative tasks.

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 the 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, while the intermediate values are denoted by floating columns.

Index