SAS® Forecast Server Client 14.1

User's Guide

The correct bibliographic citation for this manual is as follows: SAS Institute Inc. 2015. *SAS® Forecast Server Client 14.1: User's Guide*. Cary, NC: SAS Institute Inc.

SAS® Forecast Server Client 14.1: User's Guide

SAS Institute Inc., SAS Campus Drive, Cary, North Carolina 27513-2414.

July 2015

Contents

Using This Book

Audience

This documentation is designed for users of SAS Forecast Server Client 14.1.

Requirements

- Your site must license SAS Forecast Server 14.1.
- Each user of SAS Forecast Server Client must be defined in the metadata, have a valid SAS account, and have the appropriate permissions to access project content on your operating system.
- SAS Forecast Server Client is supported only on Linux and Window operating environments. To view the list of supported browsers and devices for your deployment, perform either of these steps:
  - See http://*server-name:port*/SASForecastServerClient/browseranddevices.html.
  - In SAS Forecast Server Client, click ? and select **About**. In the About window, click **Browsers and Devices**.

Part 1

Introduction to SAS Forecast Server Client

2

1

About SAS Forecast Server Client

What Is SAS Forecast Server Client?

SAS Forecast Server Client is a web-based interface is shipped with SAS Forecast Server. This new client combines functionality from SAS Time Series Studio and SAS Forecast Studio into a single user interface.

Using SAS Forecast Server Client, you can quickly and easily perform these tasks:

- segment your data by creating rule-based and query-based segmentation strategies
- create modeling strategies for hierarchical forecasting
- create exception rules to identify forecast values that meet a specific criterion, such as MAPE greater than 5.
- determine the accuracy of your forecasts over time through tracking analysis

All objects that you create with SAS Forecast Server Client (such as projects, data definitions, and strategies) are saved on the SAS Workspace Server. As a result, you

can easily create and share projects, data definitions, and strategies with other users at your site.

Benefits of Using SAS Forecast Server Client

Here are some of the benefits of using the SAS Forecast Server Client:

- You can prepare and segment your data in the same user interface that you use to generate forecasts.
- SAS Forecast Server Client performs parallel execution of the modeling tasks.
- You can run projects both interactively and in batch mode using the same user interface.
- You can work on multiple projects at the same time.
- You can use code-based segmentation strategies and modeling strategies to extend the out-of-the-box functionality of SAS Forecast Server.

How SAS Forecast Server Client Relates to Other SAS Software

SAS Forecast Server Client is a component of SAS Forecast Server. This client uses many of the procedures and options from other SAS products (such as SAS Forecast Server Procedures, SAS/ETS, Base SAS, and SAS/GRAPH).

SAS Forecast Server Procedures
: SAS Forecast Server Client uses several procedures as the basis for the automatic forecasting capabilities. SAS Forecast Server procedures are used for modeling in SAS Forecast Server Client.

  For more information, see *SAS Forecast Server Procedures: User's Guide*.

SAS/ETS

SAS/ETS software provides SAS procedures that perform econometric and time series analysis and forecasting, as well as financial analysis and reporting.

For more information, see *SAS/ETS User's Guide*.

Base SAS

Base SAS delivers a highly flexible and extensible fourth-generation programming language that is specially designed for data access, transformation, and reporting. It includes a rich library of procedures for data manipulation, information storage and retrieval, descriptive statistics, and report writing. The output for SAS Forecast Server Client is generated by the Output Delivery System that is part of Base SAS.

SAS/CONNECT

SAS Forecast Server Client uses SAS/CONNECT to run multiple forecasting nodes in parallel.

SAS/GRAPH

SAS/GRAPH software provides high-impact visuals for all levels of your organization, enabling customers to readily understand complex information and empowering them to make informed, timely decisions. SAS/GRAPH software extends the power of SAS data management, business intelligence, and analytic tools, enabling customers to turn data into full-color graphs and charts.

SAS Visual Analytics

SAS Forecast Server Client is available from the home page of SAS Visual Analytics.

2

Getting Started with SAS Forecast Server Client

How to Open SAS Forecast Server Client

To open SAS Forecast Server Client from the SAS Home page, click the Project icon.

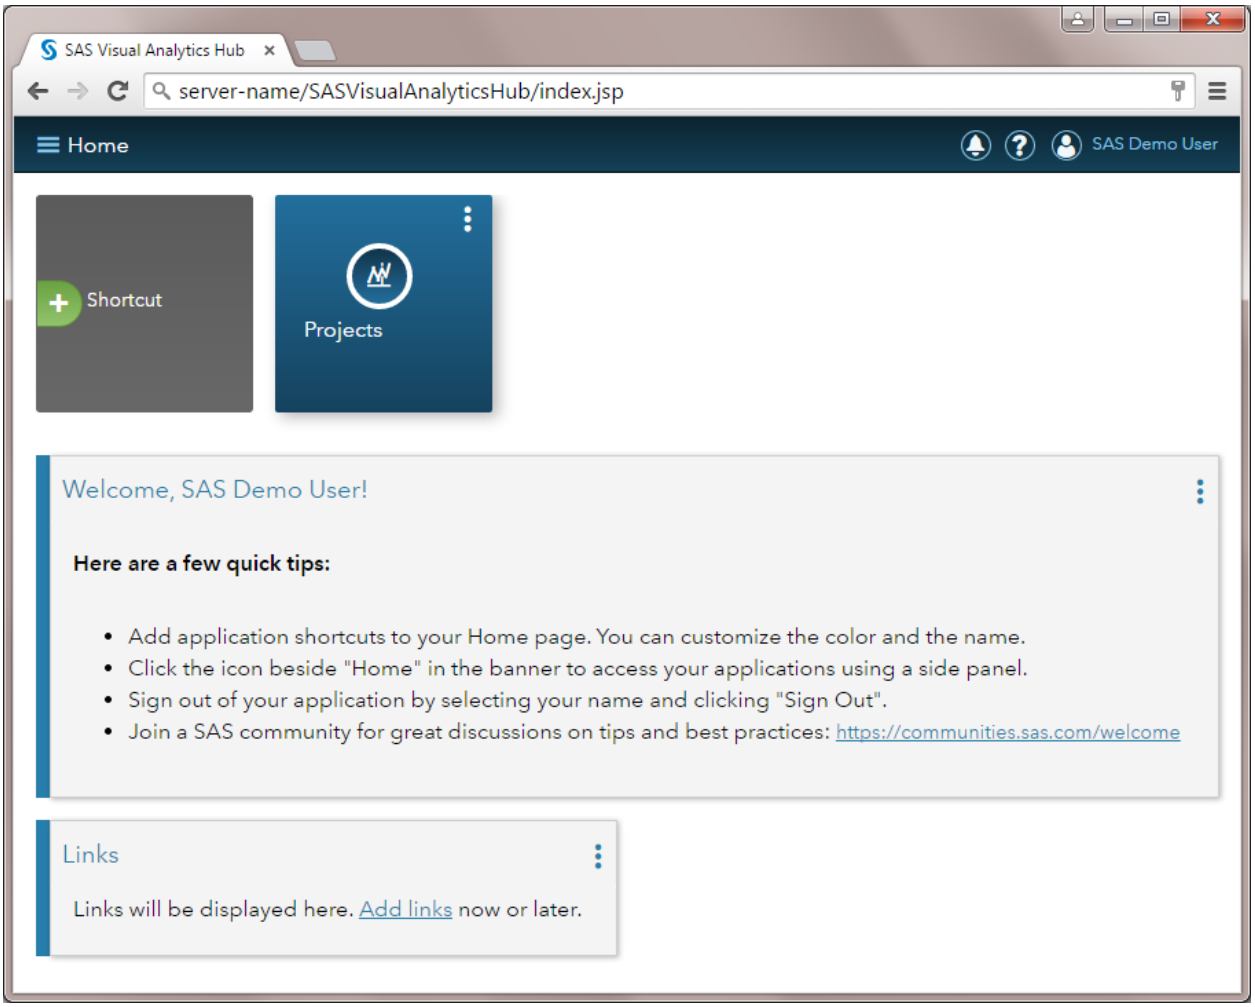

If you currently have a project open in SAS Forecast Server Client, you are taken directly into that project. If you do not have a project open in SAS Forecast Server Client, you see the SAS Forecast Server Client – Projects page, which contains the list of projects that you can access.

You can also open SAS Forecast Server Client directly by using the URL provided by your administrator. This URL should be http://*server-name*/SASForecastServerClient. In this case, SAS Forecast Server Client always opens to the SAS Forecast Server Client - Projects page.

Here is an example of the SAS Forecast Server Client - Projects page. This page lists all SAS Forecast Server Client projects that you can access. For each project, you can

see what data definition is used in the project, the user who created the project, when the project was last modified, the status of the project, and the number of exceptions identified.

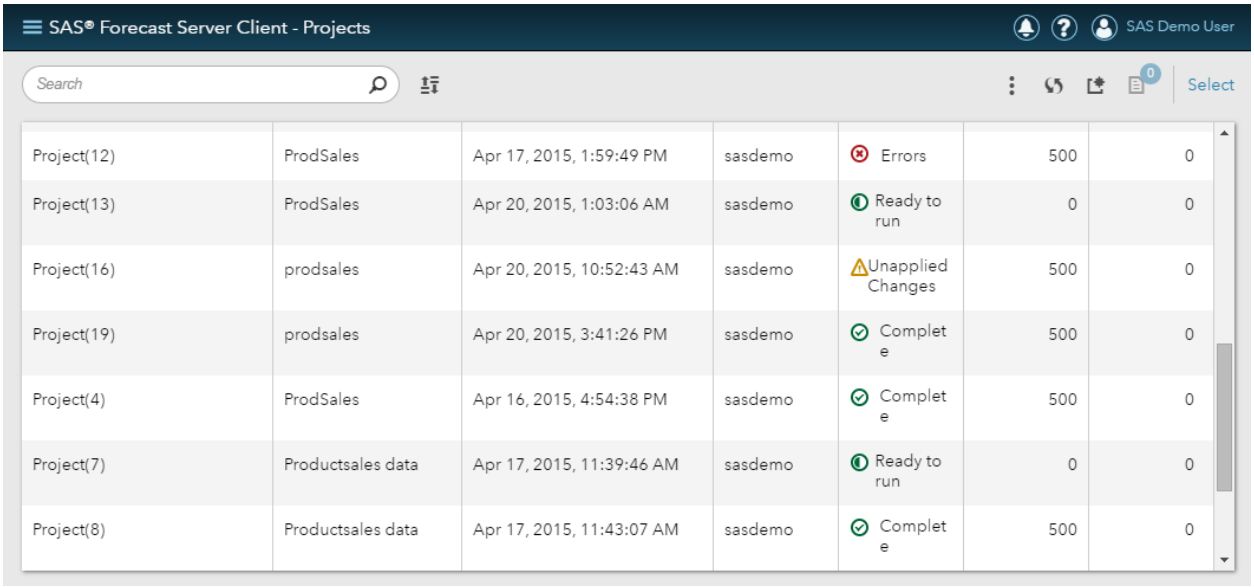

Project(12)	ProdSales	Apr 17, 2015, 1:59:49 PM	sasdemo	Errors	500	0
Project(13)	ProdSales	Apr 20, 2015, 1:03:06 AM	sasdemo	Ready to run	0	0
Project(16)	prodsales	Apr 20, 2015, 10:52:43 AM	sasdemo	Unapplied Changes	500	0
Project(19)	prodsales	Apr 20, 2015, 3:41:26 PM	sasdemo	Complete	500	0
Project(4)	ProdSales	Apr 16, 2015, 4:54:38 PM	sasdemo	Complete	500	0
Project(7)	Productsales data	Apr 17, 2015, 11:39:46 AM	sasdemo	Ready to run	0	0
Project(8)	Productsales data	Apr 17, 2015, 11:43:07 AM	sasdemo	Complete	500	0

Create Your First Project

To create a project, you must have access to one or more data definitions. You can create these data definitions, or you can use a data definition created by another SAS Forecast Server Client user at your site. If no data definitions are available when you try to create your first project, you must create a data definition before you can create the project. For more information, see “Create a Data Definition” on page 22.

To create your first project:

1 Click ☰ to open the side menu.

2 From the side menu, select **Projects**. The SAS Forecast Server Client - Projects page opens.

3 To create a new project, click . The New Project window appears.

4 Specify a name for the project.

5 Verify the location where the project is saved. By default, the project is saved in My Folders on the metadata server. In general, you should leave the default location.

6 Select the data definition that contains your input data source.

> **TIP** To view the details for a data definition (such as the name of the input data source), click .

7 Click **OK** to create the new project.

After SAS Forecast Server Client creates a project, the project template appears. No forecasts are created at this time.

The project template lists all of the nodes in the project. You can drill down into each node for more information, such as the results for that node and the specific settings used for that node.

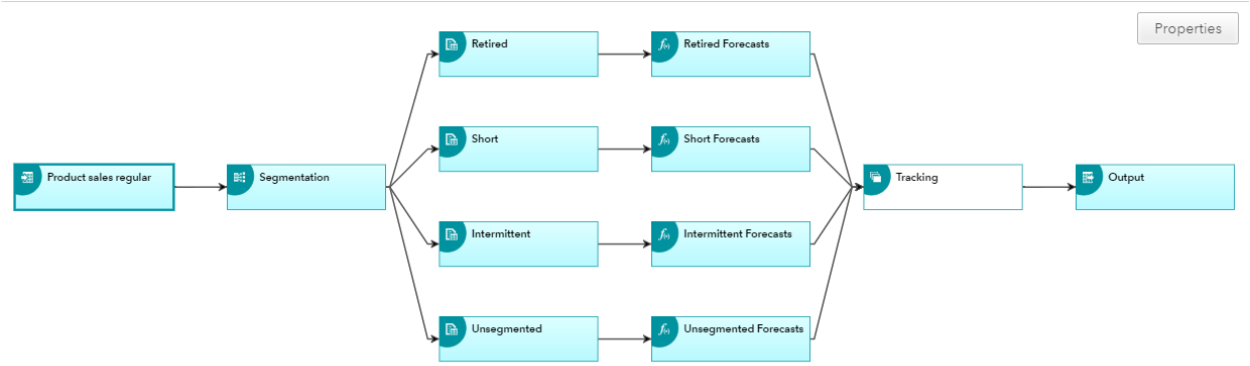

Run the Project

For this example, click ▶ to run the project. For this example, the project is using the default strategies and settings that are shipped with SAS Forecast Server Client to

generate your forecasts. When you create your own projects, you can customize these settings.

When the project finishes running, the nodes in the project template update to include the number of series that meet the criteria defined for that node. For example, in the **Retired** node, 67 series were identified as part of the retired segment.

In this example, here are the results for the segment nodes: 67 retired series, 0 short series, 33 intermittent series, and 400 unsegmented series. These series are at the lowest level of the hierarchy.

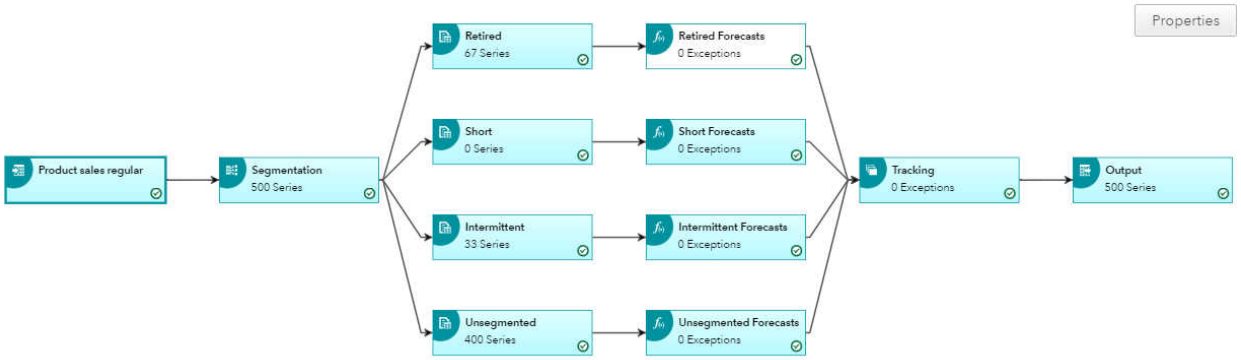

To view the results and settings for specific node, double-click on the node. If you change any of the settings for a specific node, those changes are specific to the current project and are automatically saved with your project. To edit the data definitions, segmentation strategies, modeling strategies, and exception strategies that are saved in the metadata, you must access these options from the side menu. For more information, see Chapter 3, “Using the Workspace,” on page 13.

3

Using the Workspace

The SAS Forecast Server Client - Projects Page

When you open the SAS Forecast Server Client, the first page that appears is the SAS Forecast Server Client – Projects page. From this page, you can open existing projects and create new projects.

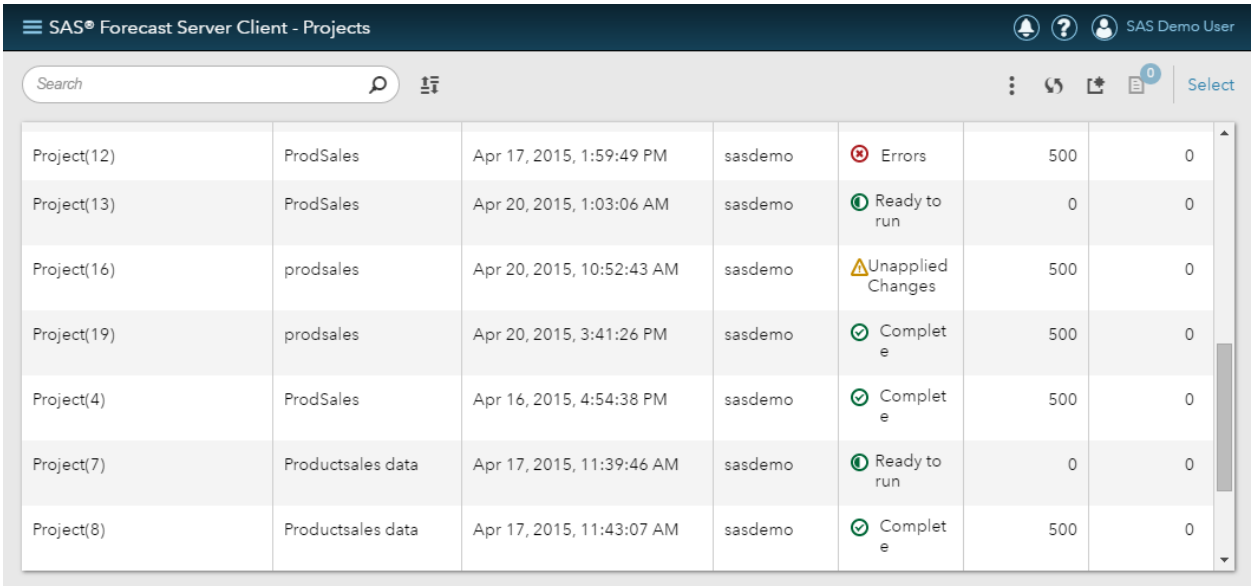

Project(12)	ProdSales	Apr 17, 2015, 1:59:49 PM	sasdemo	Errors	500	0
Project(13)	ProdSales	Apr 20, 2015, 1:03:06 AM	sasdemo	Ready to run	0	0
Project(16)	prodsales	Apr 20, 2015, 10:52:43 AM	sasdemo	Unapplied Changes	500	0
Project(19)	prodsales	Apr 20, 2015, 3:41:26 PM	sasdemo	Complete	500	0
Project(4)	ProdSales	Apr 16, 2015, 4:54:38 PM	sasdemo	Complete	500	0
Project(7)	Productsales data	Apr 17, 2015, 11:39:46 AM	sasdemo	Ready to run	0	0
Project(8)	Productsales data	Apr 17, 2015, 11:43:07 AM	sasdemo	Complete	500	0

From the project list, you can learn the following information about each project:

- the data definition used in the project
- the date on which the project was last modified
- the user ID of the person who last modified the project
- the status of the project. Projects can have any of these statuses:
  - — Complete. The project has run, and there are no errors.
  - — Unapplied changes. You have changed some settings in the project. To apply these changes, you need to run the project.
  - Ready to Run — You have successfully created the project. You can now run the project to generate your forecasts.
  - Error — The project could not run because of errors. For more information about these errors, see the SAS log.
- the number of series in the project
- the number of series that met one or more of the exception strategies. For more information, see Chapter 10, “Creating Exception Strategies,” on page 79.

Understanding the Project Template

When you select a project from the SAS Forecast Server Client – Projects page, the project template opens. This template shows you the structure of the project. For example, based on the segmentation strategies that you have defined for the project, you can see how many segments will be created when SAS Forecast Server Client runs the project. In this example, four segments (Retired, Short, Intermittent, and Unsegmented) are created.

Before (or after) you run the project, you can double-click the nodes in the project template to change the settings in that node. For example, if you double-click the **Segmentation** node, you can specify the segmentation strategy to use for that project. Any changes that you make are not saved to the metadata. The changes apply only to your project.

Here is an example of a project template:

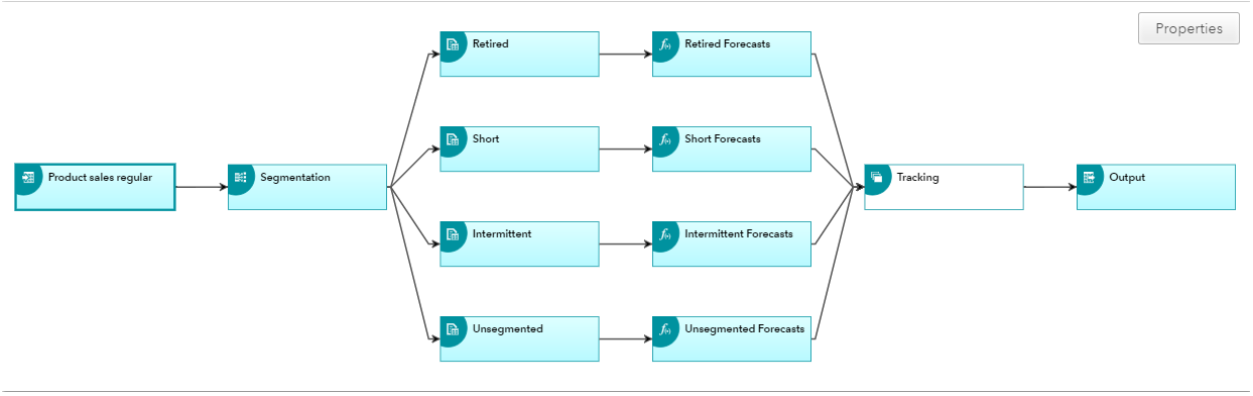

After the project runs, the nodes in the project template display the number of series or exceptions in that node. In this example, you can see that the **Retired** node contains 67 series, the **Intermittent** node contains 33 series, and the **Unsegmented** node contains 400 series.

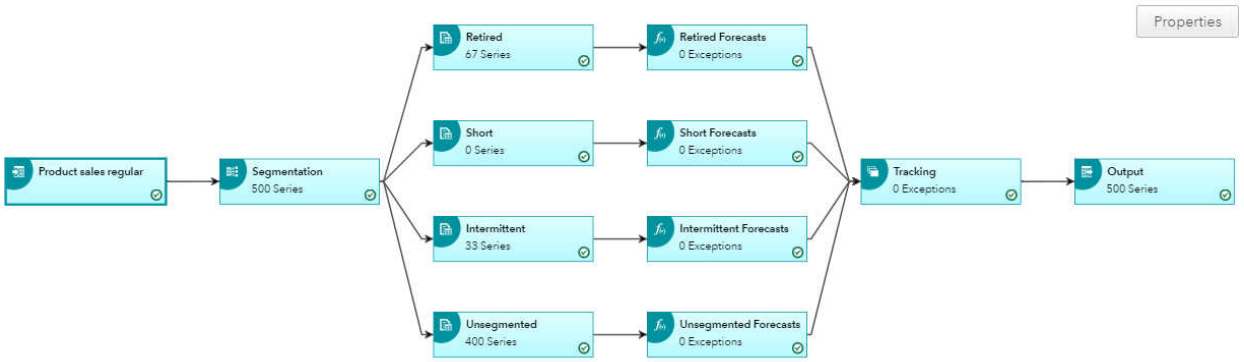

Using the Side Menu

To create data definitions and strategies that are available to all SAS Forecast Server Client users at your site, click ☰ to open the side menu. From the side menu of SAS Forecast Server Client, you can create data definitions, segmentation strategies, modeling strategies, and exception strategies. These strategies are saved in the metadata and can be used in any project. Any changes that you make to these definitions and strategies are not automatically applied to any existing projects. Instead, to incorporate any changes, you must edit the existing project or create a new project.

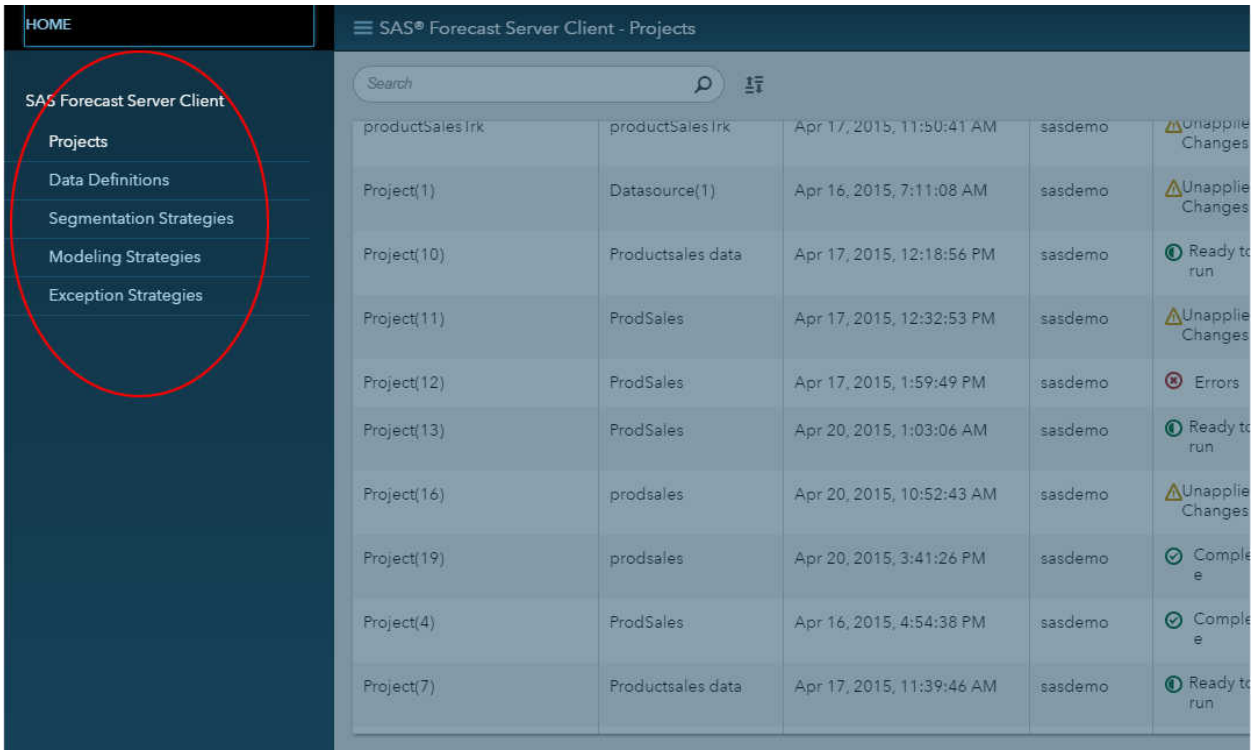
HOME
SAS® Forecast Server Client - Projects
SAS Forecast Server Client
Projects
Data Definitions
Segmentation Strategies
Modeling Strategies
Exception Strategies
Search
productSales lrk
productSales lrk
Apr 17, 2015, 11:50:41 AM
sasdemo
Project(1)
Datasource(1)
Apr 16, 2015, 7:11:08 AM
sasdemo
Project(10)
Productsales data
Apr 17, 2015, 12:18:56 PM
sasdemo
Project(11)
ProdSales
Apr 17, 2015, 12:32:53 PM
sasdemo
Project(12)
ProdSales
Apr 17, 2015, 1:59:49 PM
sasdemo
Errors
Project(13)
ProdSales
Apr 20, 2015, 1:03:06 AM
sasdemo
Project(16)
prodsales
Apr 20, 2015, 10:52:43 AM
sasdemo
Project(19)
prodsales
Apr 20, 2015, 3:41:26 PM
sasdemo
Project(4)
ProdSales
Apr 16, 2015, 4:54:38 PM
sasdemo
Project(7)
Productsales data
Apr 17, 2015, 11:39:46 AM
sasdemo

Part 2

Understanding Data Definitions

4

Creating a New Data Definition

What Is a Data Definition?

Before you select the input data set for a project, you must create a data definition. A data definition is saved in the metadata and consists of the following parts:

- the name and location of the input data set
- the hierarchy that you define for the input data set
- the interval of the time series data
- the role of each variable in the input data set and the variable settings, such as accumulation and missing value interpretation

Create a Data Definition

Open the New Data Definition Wizard

Use the New Data Definition wizard to create any data definitions that you need to use in SAS Forecast Server Client. All projects in SAS Forecast Server Client require a data definition.

1 Click ☰ to open the side menu.

2 Click **Data Definitions**. The SAS Forecast Server Client – Data Definition page opens. This page lists all of the existing data definitions that you can use.

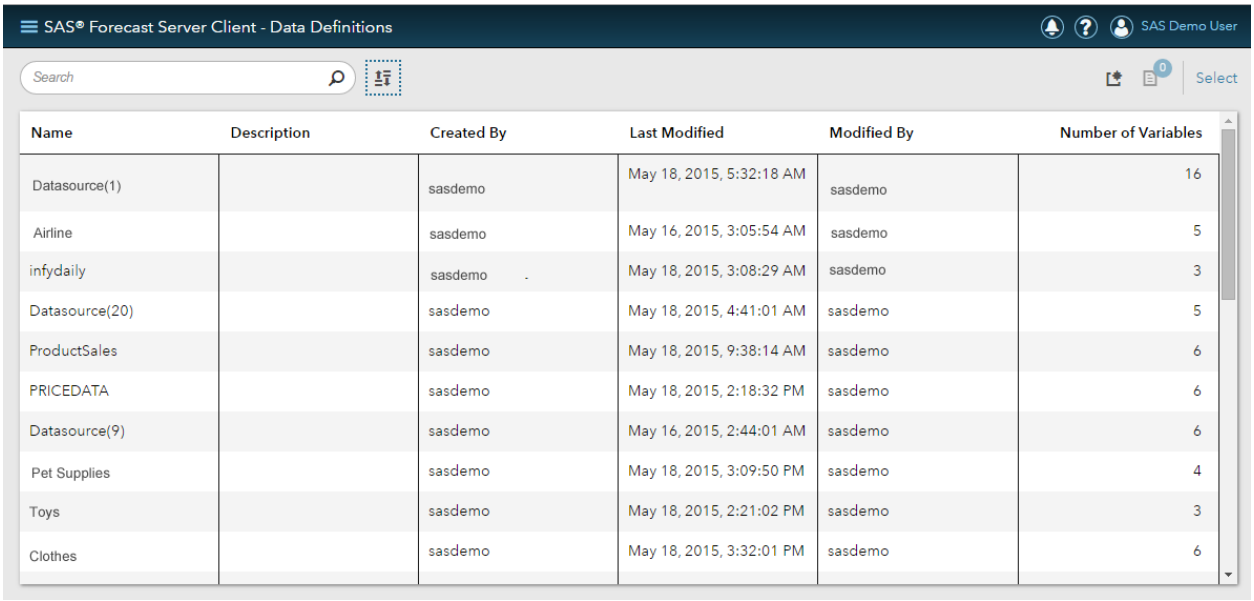

Name	Description	Created By	Last Modified	Modified By	Number of Variables
Datasource(1)		sasdemo	May 18, 2015, 5:32:18 AM	sasdemo	16
Airline		sasdemo	May 16, 2015, 3:05:54 AM	sasdemo	5
infydaily		sasdemo	May 18, 2015, 3:08:29 AM	sasdemo	3
Datasource(20)		sasdemo	May 18, 2015, 4:41:01 AM	sasdemo	5
ProductSales		sasdemo	May 18, 2015, 9:38:14 AM	sasdemo	6
PRICEDATA		sasdemo	May 18, 2015, 2:18:32 PM	sasdemo	6
Datasource(9)		sasdemo	May 16, 2015, 2:44:01 AM	sasdemo	6
Pet Supplies		sasdemo	May 18, 2015, 3:09:50 PM	sasdemo	4
Toys		sasdemo	May 18, 2015, 2:21:02 PM	sasdemo	3
Clothes		sasdemo	May 18, 2015, 3:32:01 PM	sasdemo	6

3 To create a new data definition, click . The New Data Definition wizard opens.

Step 1: Specify the Name for the Data Definition and the Location of the Input Data Set

1 Specify a name for your data definition.

2 To select the input data set for the data definition, click **Browse**.

Note: All data sets must be registered in the metadata before you can use them in SAS Forecast Server Client. If the data set that you need is not available, contact your SAS Forecast Server administrator for help.

When you are finished, the first step in your data definition should appear similar to the following:

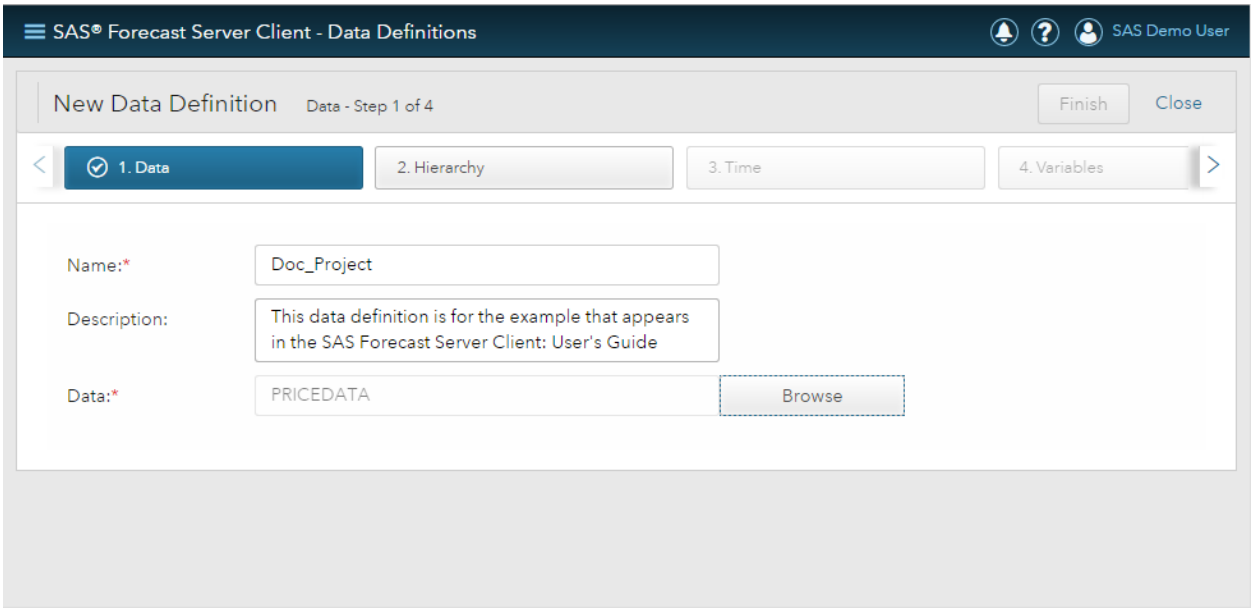

Step 2: Defining the Hierarchy

1 To define the hierarchy for your data definition, click **Hierarchy**. A hierarchy is required to forecast in the SAS Forecast Server Client. For more information, see Chapter 5, “Understanding Hierarchies,” on page 31.

2 To create the hierarchy:

a Select the classification variables from the **Available items** box and click ➡. The SAS Forecast Server Client shows all of the variables in the input data set in the **Available items** pane.

> **TIP** To view the frequency chart and summary data for a specific variable, select the variable in the **Available items** pane or the **Selected items** pane and click **Distribution**.

Note: By default, the first variable in the **Available items** pane is selected. (For example, in the following figure, the cost variable is selected by default.) If you do not want this variable in your list of selected items, you must deselect this variable. Otherwise, when you click ➡, the cost variable and any other variables that you selected are added to the **Selected items** pane.

b In the **Selected items** pane, specify the order of the variables in the hierarchy.

In this example, the hierarchy is region > line > product.

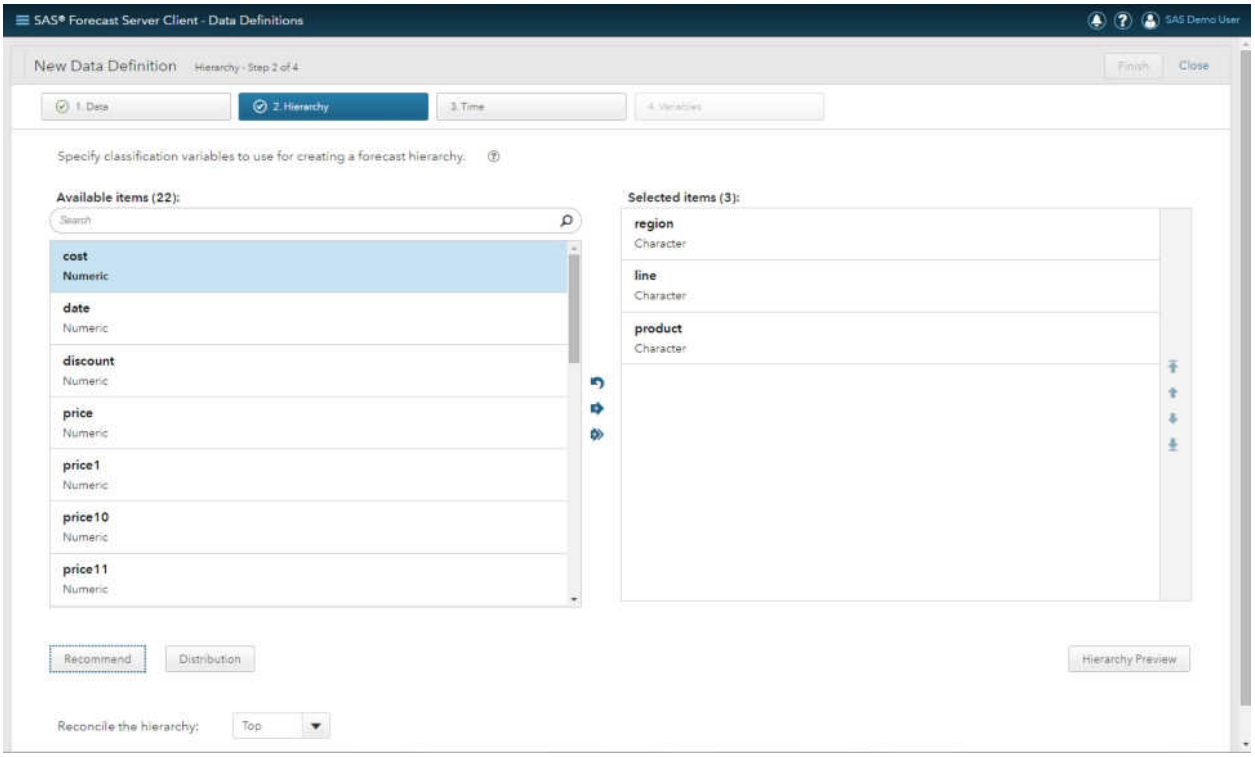

> **TIP** If you need help with defining a hierarchy, click **Recommend** to add all of the classification variables in the **Available items** pane to the **Selected items** pane. The variables are added in the order in which they appear in the data set.

To view the hierarchy that you created, click **Hierarchy Preview**.

3 To specify a reconciliation level for your hierarchy, select a value from the **Reconcile the hierarchy** drop-down list. By default, SAS Forecast Server Client uses the top-down reconciliation method. For more information, see “Understanding Reconciliation Methods” on page 31.

Step 3: Specify the Time Dimensions for Your Data

By default, SAS Forecast Server Client automatically recommends the time ID variable in the input data set. It also suggests a suitable time interval for the data. In most cases, you can accept the default values in this step. For more information, see Chapter 6, "Working with Time ID Variables," on page 33.

To modify the time dimension settings:

1. From the **Time ID variable** drop-down list, select the time ID variable in the input data source that contains the SAS date and datetime values for your data definition.

2. From the **Interval** drop-down list, select the time interval for your data.

   Note: If you select **Weekday** as the time interval for your data, click **Advanced Settings** to specify the days for the weekend. By default, the weekend is Saturday and Sunday.

3. To specify a multiplier, shift, or seasonal cycle length for your data, click **Advanced Settings**. After you specify these options in the Advanced Settings window, click **OK**.

4. Specify the number of periods to forecast (also called the horizon). The default is 12 periods.

5. Specify the confidence limit for the generated forecasts. By default, this confidence level is 95%.

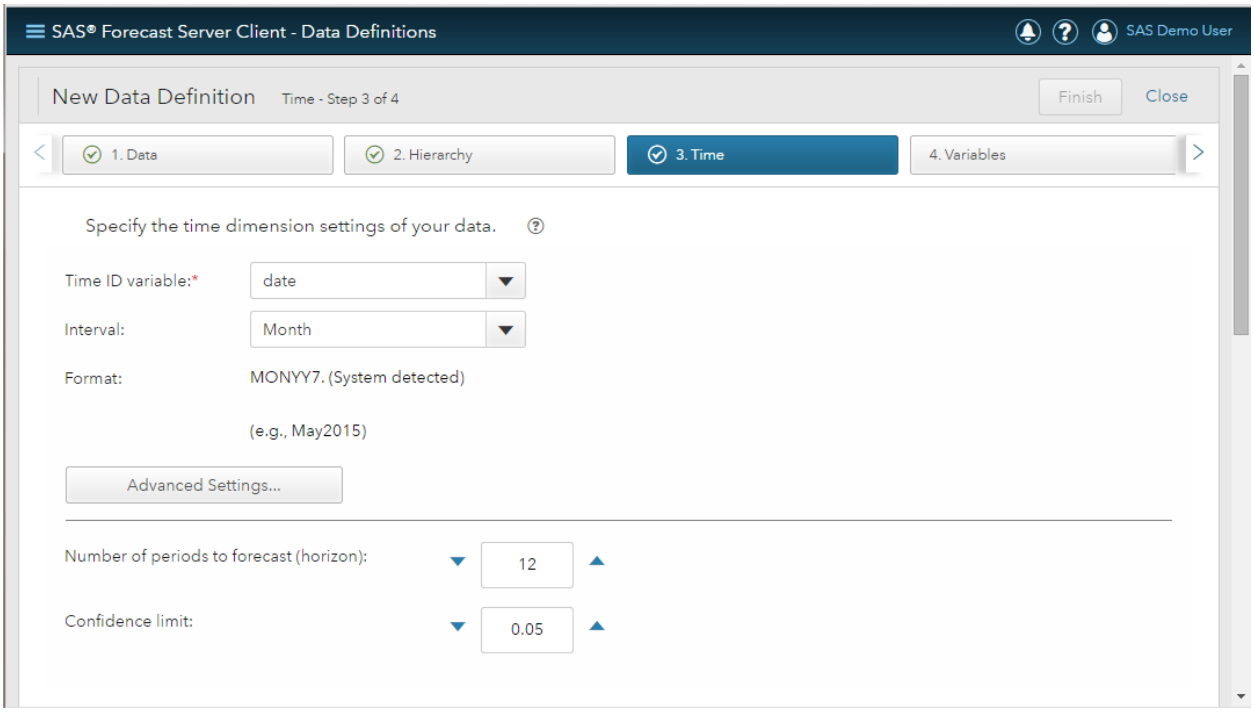

Step 4: Assign Variables to Roles

In order for SAS Forecast Server Client to forecast the data, you must assign a dependent variable. In this step, you can also specify the hierarchy aggregation and time ID accumulation for the variable. In addition, you can specify how to treat missing values and zero values in the input data set. For more information, see Chapter 7, “Understanding Roles and Interpreting Missing Values,” on page 43.

This example assumes that all you need to do is select a dependent variable.

To assign a dependent variable:

1. Select the check box for the dependent variable. A project can have only one dependent variable. In this example, **sale** is selected.

2. From the **Role** drop-down list, select **Dependent**.

The information in the table updates to reflect any options that you select.

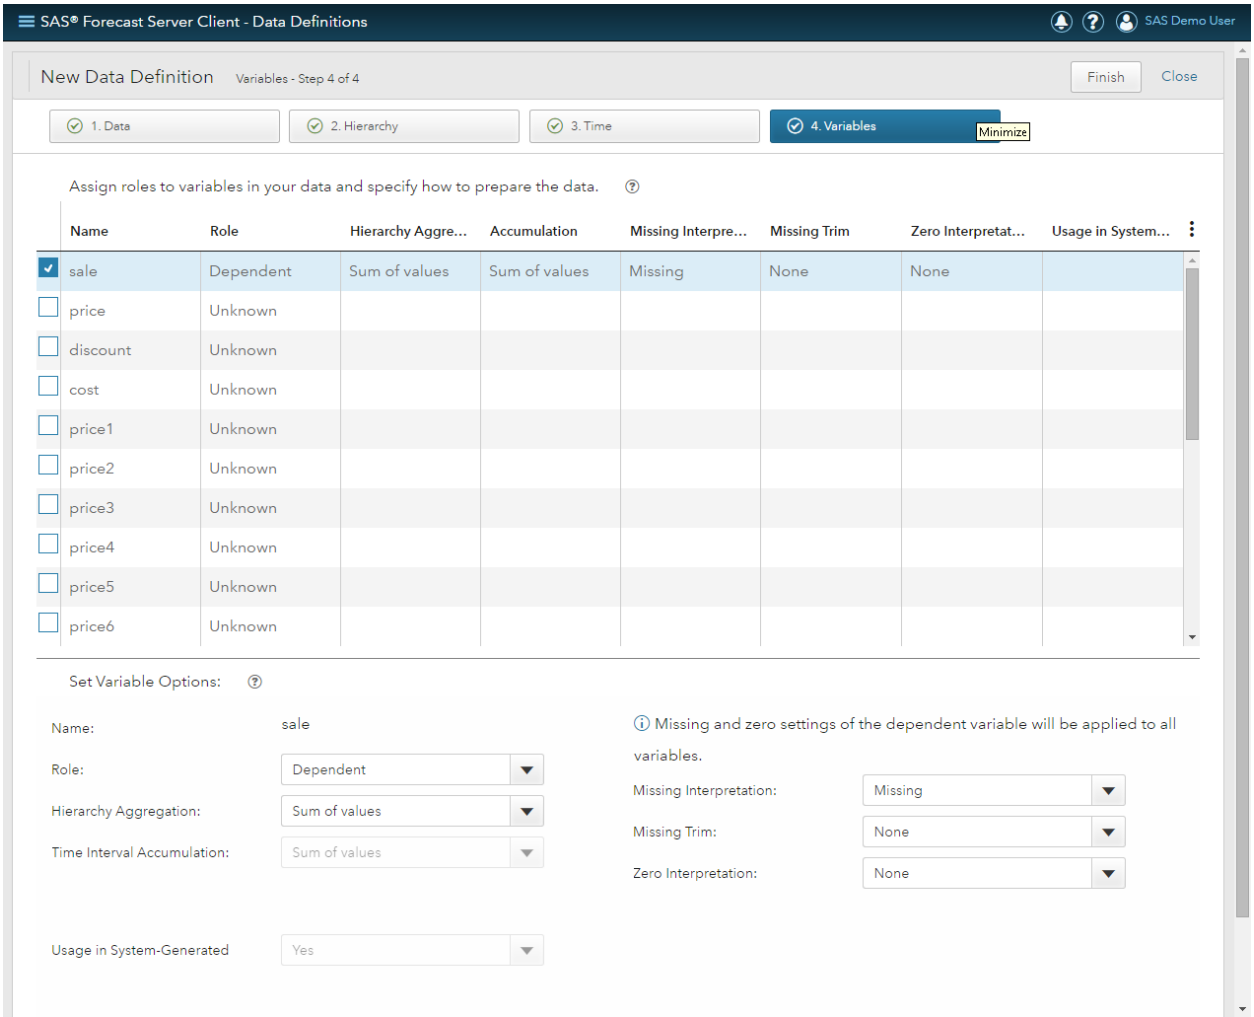

In this step, you can also assign one or more variables to the **Independent** role. Independent variables are optional.

Note: Only variables with valid names appear in this step. For more information, see “Reserved Variable Names” on page 153.

Finish and Exit the New Data Definition Wizard

After you assign your variables to roles, you are finished creating your data definition. Click **Finish** to save the data definition. Your new data definition appears on the SAS Forecast Server Client – Data Definition page.

Note: If you click **Close** at any time while creating a data definition, you are returned to the SAS Forecast Server Client – Data Definition page, and all your changes are lost.

5

Understanding Hierarchies

Requirements for Forecasting Hierarchically

To forecast hierarchically, you must assign at least one classification (BY) variable. SAS Forecast Server Client sorts the dependent variable into individual series using the hierarchy specified by the BY variables. You can assign only one dependent variable in a project.

The hierarchy for a project is defined in the data definition. You cannot change the hierarchy of a project after you have created the project.

Understanding Reconciliation Methods

When forecasting hierarchically, you can choose the level for reconciling the hierarchy. By default, the SAS Forecast Server Client uses the top-down reconciliation method. This method aggregates the data from the highest levels in the forecast and then uses these values to generate the forecasts at the lowest level.

The top-down method enables you to remove the excessive noise from the data at the lower levels of the hierarchy. However, you also might lose some components of the pattern (such as the seasonality) in the forecast.

If you choose the lowest level in the hierarchy, the SAS Forecast Server Client uses the bottom-up reconciliation method. This method uses the data at the lowest level of the hierarchy to generate the forecasts. These forecasts are then used to reconcile the forecasts for the higher levels in the hierarchy.

The bottom-up method enables you to see any patterns (such as seasonality) in the data. However, because you are using the lowest level of the hierarchy as a reference level for the forecasts, you can also have too much noise or randomness in the data. Also, these forecasts might fail because the data at the lowest level of the hierarchy can be sporadic or too sparse.

You can also select a middle level for reconciliation. In this case, the SAS Forecast Server Client uses the middle-out reconciliation method to generate the forecasts for the middle level. SAS Forecast Server Client uses the forecasts at the middle level to reconcile the forecasts for both the higher and lower levels. Some hierarchies have more than one middle level, but you can select only one middle level for reconciliation.

6

Working with Time ID Variables

What Is the Time ID Variable?

You specify the time ID variable for a project when you create the data definition. After the project has been created, you cannot change the time ID variable. The time ID variable is a variable in the input data set that contains the SAS date or datetime value for each observation. This variable is used to determine the frequency and ordering of the data and to extrapolate the time ID values for the forecasts. You can assign only one variable to this role. This variable must be either a SAS date variable, a SAS datetime variable, or a numeric variable that contains date or datetime values.

Note: SAS Forecast Server does not support time values. When you assign a variable to the time ID role, SAS Forecast Server Client recognizes the data's format. If your data uses a time format, a message stating that time data is not valid for the time ID variable appears.

Basics on Time Intervals

All time intervals must meet the following criteria:

- A discrete time interval has a beginning and an ending SAS date or SAS datetime.
- For SAS date intervals, the ending date is defined as 1 day before the beginning of the next interval.
- For SAS datetime intervals, the ending time is 1 second before the beginning of the next interval.
- Alignment refers to the identifying date of the interval and does not affect the definition of the interval.

Supported Time Intervals

SAS Forecast Server Client supports the following time intervals:

Day
: specifies daily intervals.

Hour
: specifies hourly intervals.

ISO 8601 year
: specifies ISO 8601 yearly intervals. The ISO 8601 year starts on the Monday on or immediately preceding January 4. Note that it is possible for the ISO 8601 year to start in December of the preceding year. Also, some ISO 8601 years contain a leap week.

ISO 8601 week
: specifies ISO 8601 weekly intervals of seven days. Each week starts on Monday. The starting subperiod (or subperiods) is in days (DAY). Note that WEEKV differs

from WEEK in that WEEKV.1 starts on Monday, WEEKV.2 starts on Tuesday, and so on.

Minute
: specifies minute intervals.

Month
: specifies monthly intervals.

Quarter
: specifies quarterly intervals (every three months). The starting subperiod is in months.

Retail 4-4-5 Year
: specifies ISO 8601 weekly interval, except that the starting subperiod (or subperiods) is in retail 4-4-5 months.

Retail 4-5-4 Year
: specifies ISO 8601 weekly interval, except that the starting subperiod (or subperiods) is in retail 4-5-4 months.

Retail 5-4-4 Year
: specifies ISO 8601 weekly interval, except that the starting subperiod (or subperiods) is in retail 5-4-4 months.

Retail 4-4-5 Month
: specifies retail 4-4-5 monthly intervals. The 3rd, 6th, 9th, and 12th months are five ISO 8601 weeks long with the exception that some 12th months contain leap weeks. All other months are four ISO 8601 weeks long. R445MON intervals begin with the 1st, 5th, 9th, 14th, 18th, 22nd, 27th, 31st, 35th, 40th, 44th, and 48th weeks of the ISO year.

Retail 4-5-4 Month
: specifies retail 4-5-4 monthly intervals. The 2nd, 5th, 8th, and 11th months are five ISO 8601 weeks long. All other months are four ISO 8601 weeks long with the exception that some 12th months contain leap weeks. R454MON intervals begin with the 1st, 5th, 10th, 14th, 18th, 23rd, 27th, 31st, 36th, 40th, 44th, and 49th weeks of the ISO year.

Retail 5-4-4 Month
: specifies retail 5-4-4 monthly intervals. The 1st, 4th, 7th, and 10th months are five ISO 8601 weeks long. All other months are four ISO 8601 weeks long with the exception

that some 12th months contain leap weeks. R544MON intervals begin with the 1st, 6th, 10th, 14th, 19th, 23rd, 27th, 32nd, 36th, 40th, 45th, and 49th weeks of the ISO year.

Retail 4-4-5 Quarter
: specifies retail 4-4-5 quarterly intervals (every 13 ISO 8601 weeks). Some fourth quarters will contain a leap week. The starting subperiod (or subperiods) is in retail 4-4-5 months.

Retail 4-5-4 Quarter
: specifies retail 4-5-4 quarterly intervals (every 13 ISO 8601 weeks). Some fourth quarters will contain a leap week. The starting subperiod (or subperiods) is in retail 4-5-4 months.

Retail 5-4-4 Quarter
: specifies retail 5-4-4 quarterly intervals (every 13 ISO 8601 weeks). Some fourth quarters will contain a leap week.

Second
: specifies second intervals.

Semimonth
: specifies semimonthly intervals. Each month consists of two periods. The first period starts on the first, and the second period starts on the 16th.

Semiyear
: specifies intervals every six months. The starting subperiod is in months.

Ten-day
: specifies 10-day intervals. Each month consists of three periods. The first period is the 1st through the 10th day of the month. The second period is the 11th through the 20th day of the month. The third period is the 21st through the end of the month.

Week
: specifies weekly intervals of seven days.

The days of the week are numbered as follows:

Value of the Shift	Day of the Week
1	Sunday

Value of the Shift	Day of the Week
2	Monday
3	Tuesday
4	Wednesday
5	Thursday
6	Friday
7	Saturday

Weekday
: specifies daily intervals with weekend days included in the preceding weekday. The weekday interval is the same as the day interval, except that the weekend days are absorbed into the preceding weekday. The default weekend days are Saturday and Sunday, but you can specify the days to include in the weekend. If you use the default weekend, then there are five weekday intervals in a calendar week: Monday, Tuesday, Wednesday, Thursday, and the three-day period Friday, Saturday, and Sunday.

Year
: specifies yearly intervals. The starting subperiod is in months.

Understanding SAS Time Intervals

SAS Forecast Server Client analyzes the variable assigned to the time ID role to detect the time interval of the data. SAS assumes that all of the values in the time ID variable are either SAS date or datetime values and distinguishes between the values by their magnitude. This assumption fails if you have dates extending beyond July 21, 2196, or datetimes before January 1, 1960.

For many businesses, their time series data is equally spaced, or any two consecutive indices have the same difference between the time intervals. The following table shows an equally spaced time series with a one-year interval.

Year	Number of Sales
2011	42,100
2012	45,000
2013	47,000
2014	50,000

If the time interval cannot be detected from the variable that you assign, then you need to specify the interval and seasonal cycle length. For example, the following table shows an unequally spaced time series.

Year	Number of Sales
2010	32,100
2011	45,000
2013	47,000
2014	50,000

Often the time interval cannot be detected with transactional data (timestamped data that is recorded at no particular frequency). If this is the case, the SAS Forecast Server Client accumulates the data into observations that correspond to the interval that you specify. For nontransactional data, you might need to specify the interval and seasonal cycle length if there are numerous gaps (missing values) in the data. In this case, the SAS Forecast Server Client supplies the missing values. A validation routine checks the values of the time ID to determine whether they are spaced according to the interval that you specified.

The interval determines the frequency of the output. You can modify the time interval. You can change the interval from a higher frequency to a lower frequency or from a lower frequency to a higher frequency. For each time interval, you can specify the type (such as monthly or weekly), a multiplier, and a shift (the offset for the interval). You can specify a greater time interval than that found in the input data. A smaller interval should not be used, because a small interval will generate a large number of missing observations.

Seasonal cycle length specifies the length of a season. This value is populated automatically if the SAS Forecast Server Client can determine the seasonal cycle length from the time ID variable. However, you can specify a seasonal cycle length other than the default if you want to model a cycle in the data. For example, your data might contain a 13-week cycle, so you need to specify a 13-week seasonal cycle length in the SAS Forecast Server Client.

Here is the syntax for an interval:

name<*multipler*><.*shift*>

Here is an explanation of each of the user-supplied values:

name
is the name of the interval.

multipler
specifies the multiplier of the interval. This value can be any positive number. By default, the multiplier is 1. For example, YEAR2 indicates a two-year interval.

.shift
specifies the starting point for the interval. By default, this value is one. A value greater than 1 shifts the start to a later point within the interval. The unit for the shift depends on the interval. For example, YEAR.4 specifies a shift of four months, so the year is from April 1 through March 31 of the following year.

The examples in the following table show how the values that you specify for the interval, multiplier, and shift work together.

Interval Name (in SAS code format)	Default	Shift Period	Example
YEARm.s	January 1	Months	YEAR2.7 specifies an interval of every two years. Because the value for the shift is 7, the first month in the year is July.
SEMIYEARm.s	January 1 July 1	Months	SEMIYEAR.3 - six-month intervals, March-August and September-February.
QTRm.s	January 1 April 1 July 1 October 1	Months	QTR3.2 - three-month intervals starting on April 1, July 1, October 1, and January 1.
SEMIMONTHm.s	First and 16th of each month	Semimonthly periods	SEMIMONTH2.2 - intervals from the 16th of one month through the 15th of the next month.
MONTHm.s	First of each month	Months	MONTH2.2 - February-March, April-May, June-July, August-September, October-November, and December-January of the following year.
TENDAYm.s	First, 11th, and 21st of each month	Ten-day periods	TENDAY4.2 - Four ten-day periods starting at the second ten-day period.

Interval Name (in SAS code format)	Default	Shift Period	Example
WEEKm.s	Each Sunday	Days (1=Sunday . . . 7=Saturday)	WEEK6.3 specifies six-week intervals starting on Tuesdays.
DAYm.s	Each day	Days	DAY3 - three-day intervals starting on Sunday.
HOURm.s	Start of the day (midnight)	Hours	HOUR8.7 specifies eight-hour intervals starting at 6:00 a.m., 2:00 p.m., and 10:00 p.m.

For more information about time intervals, see *SAS/ETS: User's Guide*.

7

Understanding Roles and Interpreting Missing Values

Understanding Variable Roles

Classification (BY) variables
: Classification (BY) variables enable you to obtain separate analyses for groups of observations. You select these BY variables when you create the hierarchy. The order of the BY variables describes the structure of the hierarchy. For more information, see Chapter 5, “Understanding Hierarchies,” on page 31.

Dependent variable
: Dependent variables are the variables that you want to model and forecast. You can assign only one variable to this role. You must assign a dependent variable to run the project.

Independent variable

Independent variables are any explanatory, input, predictor, or causal factor variables. You can assign only numeric variables to this role. When creating the system-generated models, the SAS Forecast Server Client tries to use the independent variables in the model generation.

For independent variables, you can specify whether the variable is used in system-generated models.

- **Yes** – specifies that the independent variable be included in the model as long as the model does not fail to be diagnosed. When you select this option, the delay and orders of the numerator and denominator are set to zero.
- **No** — specifies that the independent variables should not be included in the model.
- **Maybe** – specifies that the independent variable be included in the model as long as its parameters are significant and the value of a criterion exceeds a threshold.

The **Yes** and **Maybe** options are also available with either the positive or negative option in parentheses. For example, selecting **Yes (positive)** drops the input variable from the model if its coefficient is negative, and selecting **Yes (negative)** implies the opposite. The specification of positive or negative does not mean that constraints are imposed during the estimation of the variable's coefficient in the model.

Time ID variable

The time ID variable is a variable in the input data set that contains the SAS date or datetime values for each observation. By default, SAS Forecast Server Client recommends a time ID variable. However, you can specify a different time ID variable when you are creating a data definition. For more information, see Chapter 6, "Working with Time ID Variables," on page 33.

Unknown

Any variables with an **Unknown** role are ignored in the data definition. By default, all variables are unknown when you initially create the data definition. To run a project, you must assign one variable to the **Dependent** role.

Understanding Accumulation and Aggregation

The Difference between Accumulation and Aggregation

Aggregation is the process of combining more than one time series to form a single series. For example, you can aggregate all of the series in the Electronic group into a single series that represents the group as a whole. Accumulation combines data within the same time interval. In the New Project wizard, you can specify the aggregation and accumulation options for the dependent, independent, adjustment, and reporting variables in your project.

The following examples explain when you might want to use an aggregation method:

- Your data set contains the sales for a group of products. If you want to know the total sales for a category, you would choose **Sum of values** as the aggregation method.
- Your data contains the price of each product. If you want to know the average price for a product line, you would choose **Average of values** as the aggregation method.

Accumulation combines data within the same time interval into a summary value for that time period. Accumulation can be either of the following:

- the process of converting a time series that has no fixed interval (for example, transactional data) into a time series that has a fixed interval (such as hourly or monthly)
- the process of converting a time series that has a fixed interval into a time series with a lower frequency time interval (such as hourly into daily)

The Accumulation and Aggregation Methods

Note: The difference between accumulation and aggregation is the dimension along which each method is applied. The following equations focus on accumulation.

Aggregation is restricted to the total and average methods, but the same equations apply.

Let $R = \{r_q\}_{q=1}^{Q}$ be the data vector ordered by the time series occurrence in the data set with respect to the observation index. Let $q = 1, \ldots, Q$ be the index that represents this ordering. Let Q_N be the number of nonmissing values and let $Q_{NMISS} = Q - Q_N$ be the number of missing values in the data vector. Let $\overline{r} = \frac{1}{Q_N} \sum_{q=1}^{Q} r_q$ be the average value of the data vector with the missing values ignored.

The following example accumulates the observation series $Z^{(N)} = \{z_i\}_{i=1}^{N}$ to the time series $Y^{(T)} = \{y_t\}_{t=1}^{T}, y_t = Accumulate(Z_t^{(T)})$, for $t = 1, \ldots, T$. In this situation, $R = Z_t^{(T)}$ and $Q = N_t^{(T)}$ for $t = 1, \ldots, T$.

Let $a = Accumulate(R)$ be this accumulated value for this data vector when the following accumulation methods are applied:

None
: does not accumulate the vector values.

Sum of values
: accumulates the vector values based on the summation of their values.

$$a = \sum_{q=1}^{Q} r_q$$

Missing values are ignored in the summation. If $Q_N = 0$, then *a* is set to missing.

Average of values
: accumulates the vector values based on the average of their values.

$$a = \overline{r} = \frac{1}{Q_N} \sum_{q=1}^{Q} r_q$$

Missing values are ignored in the summation. If $Q_N = 0$, then *a* is set to missing.

Minimum of values
: accumulates the vector values based on the minimum of their values.

$$a = \min(\{r_q\}_{q=1}^{Q})$$

Missing values are ignored in the summation. If $Q_N = 0$, then *a* is set to missing.

Number of missing values
: accumulates the vector values based on the number of missing values.

$$a = Q_{NMISS}$$

Standard deviation of values
: accumulates the vector values based on their standard deviation.

$$a = \sqrt{\frac{1}{Q_N - 1} \sum_{q=1}^{Q} (r_q - \overline{r})^2}$$

Missing values are ignored in the summation. If $Q_N \leq 1$, then *a* is set to missing.

Total number of values
: accumulates the vector values based on the number of values.

$$a = Q$$

Uncorrected sum of squares of values
: accumulates the vector values based on their uncorrected sum of squares.

$$a = \sum_{q=1}^{Q} (r_q)^2$$

Missing values are ignored in the summation. If $Q_N = 0$, then *a* is set to missing.

Maximum of values
: accumulates the vector values based on the maximum of their values.

$$a = \max(\{r_q\}_{q=1}^{Q})$$

Missing values are ignored in the summation. If $Q_N = 0$, then *a* is set to missing.

Number of nonmissing values
accumulates the vector values based on the number of nonmissing values.

$$a = Q_N$$

Corrected sum of squares of values
accumulates the vector values based on their corrected sum of squares.

$$a = \sum_{q=1}^{Q} (r_q - \overline{r})^2$$

Missing values are ignored in the summation. If $Q_N = 0$, then *a* is set to missing.

Median of values
accumulates the vector values based on the median of their values.

$$a = median(\{r_q\}_{q=1}^{Q})$$

Missing values are ignored in the summation. If $Q_N = 0$, then *a* is set to missing.

First of values
accumulates the vector values based on the first observation in the data.

$$a = r_1$$

Last of values
accumulates the vector values based on the last observation in the data.

$$a = r_Q$$

Working with Missing Values

If your data contains missing values in variables other than the time ID variable (such as the dependent or independent variables), you can specify how to interpret missing values. The settings for the dependent variable are applied to all variables.

- Use the **Missing Interpretation** option to specify how to replace missing values in the data. Here are the valid values for this option:
  - **0** – all missing values are set to zero.
  - **First** – all missing values are set to the first nonmissing value.
  - **Last** – all missing values are set to the last nonmissing value.
  - **Maximum** – all missing values are set to the maximum value.
  - **Median** – all missing values are set to the median value.
  - **Minimum** – all missing values are set to the minimum value.
  - **Missing** – all missing values are set to missing. This is the default.
  - **Next** – each missing value is set to the next nonmissing value. Missing values at the end of the accumulated series remain missing.
  - **Previous** – each missing value is set to the previous accumulated nonmissing value. Missing values at the beginning of the accumulated series remain missing.
- Use the **Missing Trim** option to specify how missing values are removed from the accumulated time series. By default, no missing values are removed for dependent and independent variables. You can choose to keep all of the missing values, remove the beginning missing values, remove the ending missing values, or remove both the beginning and ending missing values.
- Use the **Zero Interpretation** option to specify how beginning and ending zero values are interpreted in the accumulated time series. By default, no zero values are removed. You can choose to keep the beginning and ending zeros, set the beginning

zero values to missing, set the ending zero values to missing, or set both the beginning and ending zero values to missing.

Part 3

Defining Strategies

8

Defining Segmentation Strategies

What Is a Segment?

Often, you need to partition data into different groups based on the nature of the data (for example, slow moving items, new products, and so on). Segmentation is a process for creating these partitions, also called segments. In SAS Forecast Server Client, every project is required to have a segmentation strategy. This strategy consists of a set of rules that SAS Forecast Server Client uses to create the individual segments.

Here are examples of rules that you might see in a segmentation strategy:

- intermittent series. For example, your company sells a product that is rarely purchased.
- retired series. For example, your company stopped producing Product X in May 2015, so after that month, you do not have any sales data for that product.
- short series. For example, your company started producing Product Y in May 2015, so your starting point for that product data is May. However, the series for your other product lines start in September 2011.

Data that does not satisfy any of the rules in the segmentation strategy is grouped into an unsegmented node. All of the segments (including the unsegmented node) make up 100% of the data.

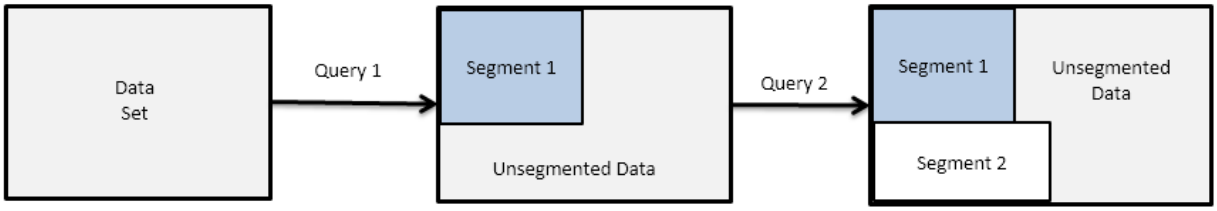

As you update the input data set for the project, some series might move from one segment to another. For example, a series could have additional data, so it no longer belongs in the Short segment. In SAS Forecast Server Client, segmentation is performed on the lowest level of the hierarchy.

About the Prebuilt Segmentation Strategy

SAS Forecast Server Client is shipped with a default segmentation strategy called Prebuilt Segmentation Strategy. To view the Prebuilt Segmentation Strategy, click ☰ to open the side menu, and select **Segmentation Strategies**.

The Prebuilt Segmentation Strategy appears on the SAS Forecast Server Client - Segmentation Strategies page.

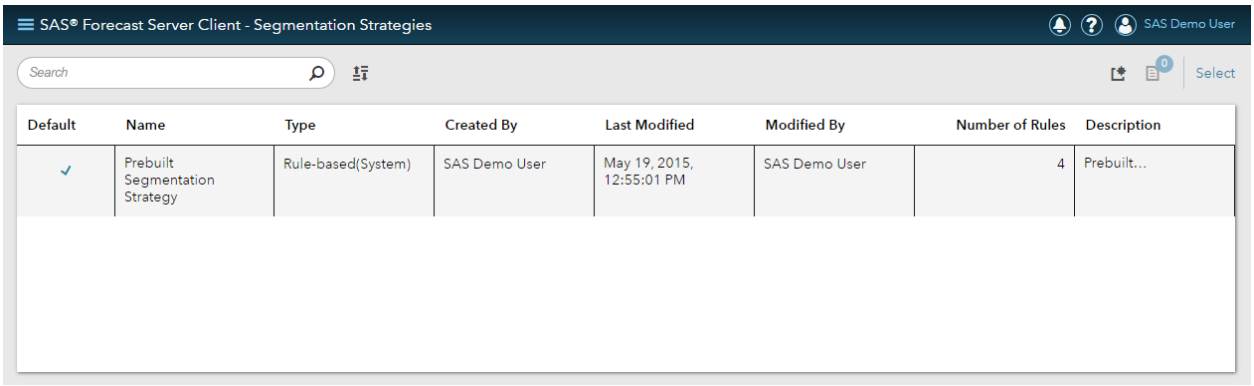

When you open the Prebuilt Segmentation Strategy, you see that it contains these rules:

- Retired. This rule identifies any series that has 24 intervals since the last nonmissing observation.
- Short. This rule identifies any series that contains fewer than or equal to 6 observations.
- Intermittent. This rule identifies any series where the average between nonzero observations is greater than 2.
- Unsegmented, which contains all any series that were not identified by the Retired, Short, and Intermittent rules.

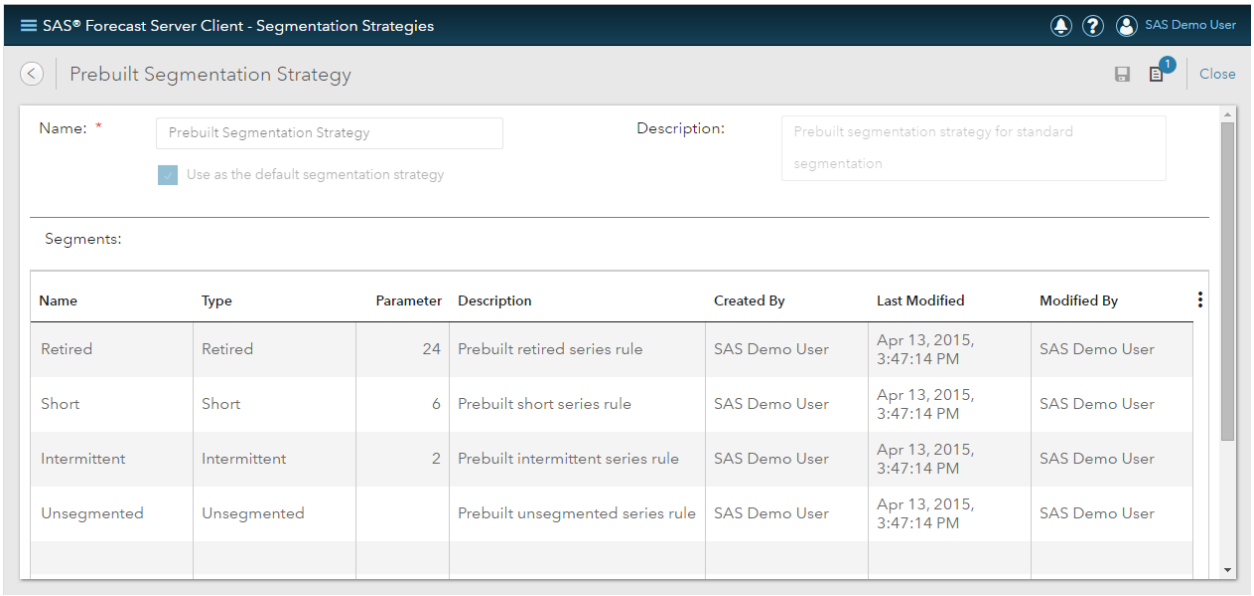

Working with Rule-Based Segmentation Strategies

What Is a Rule-Based Segmentation Strategy?

In the SAS Forecast Server Client - Segmentation Strategies page, rule-based segmentation strategies are labeled **Rule-based**. Rule-based segmentation strategies can include any of the default rules (Retired, Short, Intermittent, or Unsegmented) that are shipped with the SAS Forecast Server Client. You can also create a custom rule that is based on the value of the dependent variable or the descriptive statistics of the dependent variable.

Create a Rule-Based Segmentation Strategy

To define a rule-based segmentation strategy:

1 Click ☰ to open the side menu.

2 Click **Segmentation Strategies**. The SAS Forecast Server Client - Segmentation Strategies page appears.

3 To create a new segmentation strategy, click . The New Segmentation Strategy page opens.

4 Specify a name and a description for the new segmentation strategy.

5 Select the **Use as the default segmentation strategy** check box if you want SAS Forecast Server Client to use this strategy as the default.

6 To define the rules for the new segmentation strategy, click +. The Add Segments window appears.

7 From the **Type** drop-down list, select the type of segment that you want to create.

You can choose from these options:

- **Intermittent** — specifies the average separation between nonzero observations exceeds the specified number.
- **Retired** — specifies the number of observations since the last observation.
- **Short** — specifies the series that contain less than the specified number of observations.
- **Custom** — enables you to create a rule using the value of the BY variable or the descriptive statistics for the dependent variable.

Click **OK**.

The new rule (in this example, called Segment 1) appears in the table for the segmentation strategy.

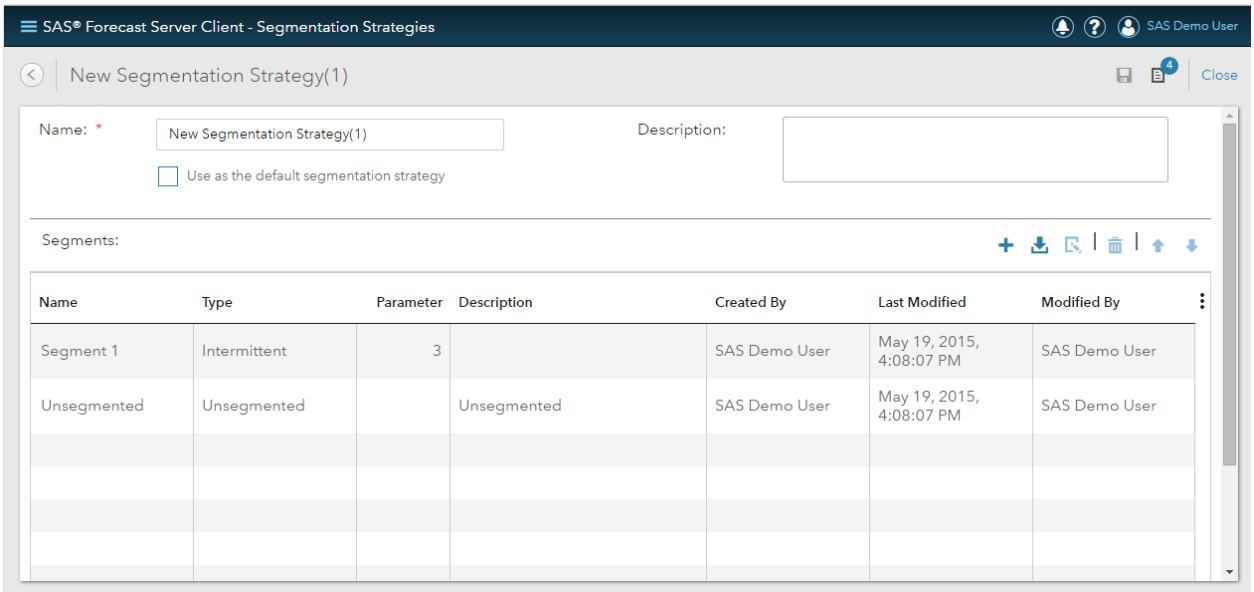

8 Repeat steps 6 and 7 to define all of the rules that you need in the segmentation strategy.

9 Use the up and down arrows to specify the order in which the rules should be applied to the series. For example, suppose you had three segments in the table: Retired, Short, and Unsegmented (in that order). A series could meet the criteria for both the Retired and Short segments. However, because the rule for the Retired segment appears in the table first, the series is assigned to the Retired segment. This series is not included in the Short segment.

10 When you are finished defining the rules for the rule-based segmentation strategy, click **Close**. The new segmentation strategy now appears in the table on the SAS Forecast Server Client - Segmentation Strategies page.

Import a Rule-Based Segmentation Strategy

Instead of creating your own rule-based segmentation strategy, you might want to use a rule-based segmentation strategy created by another user at your site. You can import these rule-based segmentation strategies, so they are available to use in your projects.

To import a rule-based segmentation strategy:

1. On the SAS Forecast Server Client - Segmentation Strategies page, click . The Import Segmentation Strategy window appears.
2. For the **Type** option, select **Rule-based**.
3. Select the check boxes for all the strategies that you want to import.

   TIP Click to view the names of the segment in a strategy. When importing multiple strategies, you cannot have segments with the same name because these segments might have different rules. SAS Forecast Server Client displays an error when you try to import two segment rules with the same name.

4. To replace any segments that are currently part of the segmentation strategy, click **Replace existing project segments**.
5. Click **Import**. The new segments now appear in the table for the segmentation strategy.

Example: Creating a Rule-Based Strategy for Retired and Custom Segments

1. On the SAS Forecast Server - Segmentation Strategies page, click . The New Segmentation Strategy page opens.
2. In the **Name** box, enter `Hard Candy Sales`.

3 To define the Retired segment in this new segmentation strategy, click +. The Add Segments window appears.

4 For the name of the new segment, enter `Retired`.

5 From the **Type** drop-down list, select **Retired**.

6 For the number of intervals since the last observation, enter `12`.

7 Click **OK** to close the Add Segments window.

The rule for the Retired segment now appears in the table for the new Hard Candy Sales strategy.

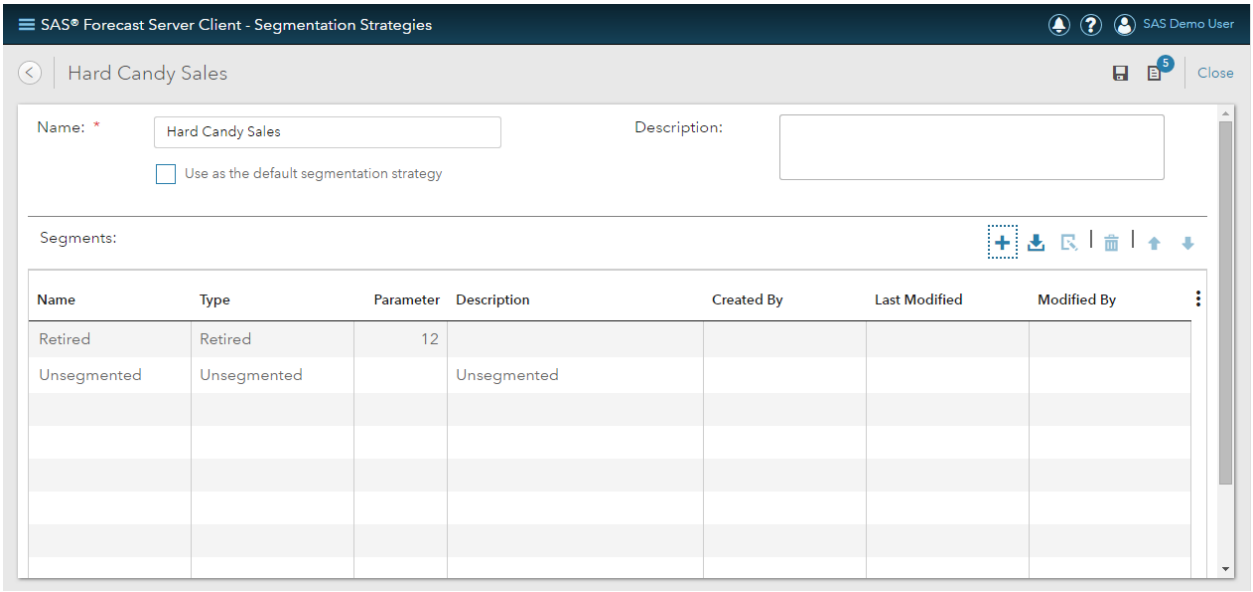

8 To define the rule for the Custom segment, click +. The Add Segments window appears.

9 For the name of the new segment, enter `Number of values`.

10 From the **Type** drop-down list, select **Custom**.

11 Next to the table, +. In the Add Rule window, specify these values:

- From the **Property** drop-down list, select **Number of values**.

- From the **Condition** drop-down list, select **>=**.
- In the **Value** box, enter `1000`.

12 Click **OK** to close the Add Segments window.

The rule for the Number of values segment now appears in the table for the new Hard Candy Sales strategy.

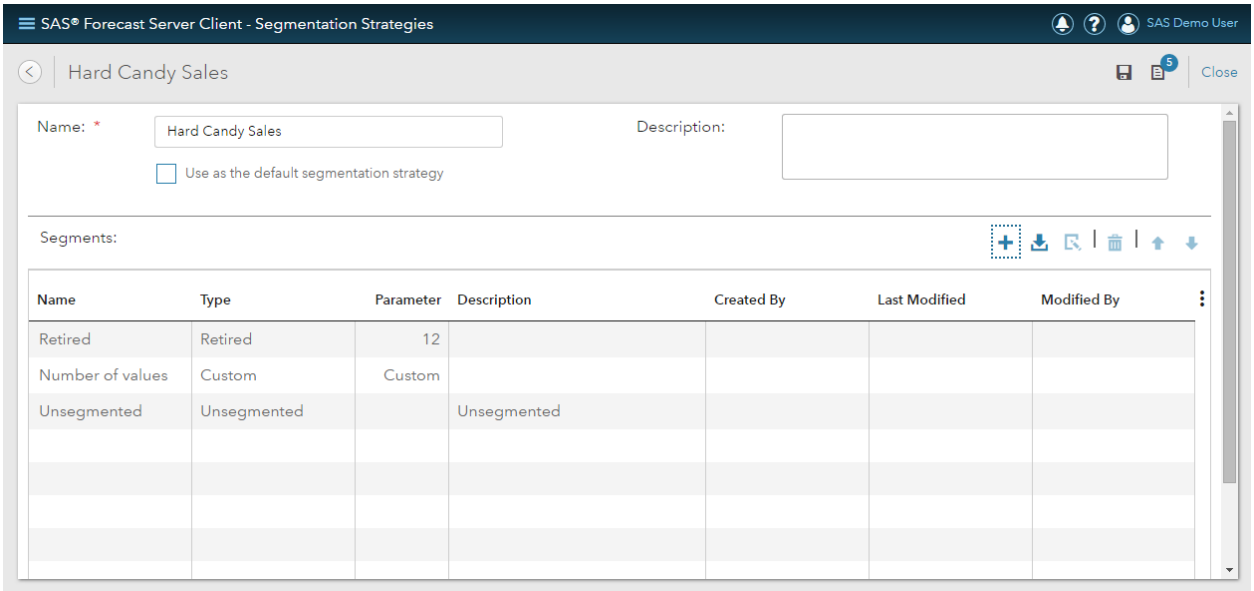

When SAS Forecast Server Client runs this segmentation strategy, it first identifies any series that meet the criterion for the Retired segment. Then SAS Forecast Server Client identifies any series that meet the criterion for the Number of values segment. However, for this example, you want SAS Forecast Server Client to identify the series that meet the Number of values segment first.

13 To change the order of the segments, select the **Number of values** row and click ⬆.

With this new order, any series that meet the criterion for the Number of values segment are identified first. Then any series that meet the criterion for the Retired

segment are identified. Any remaining series are placed in the Unsegmented segment.

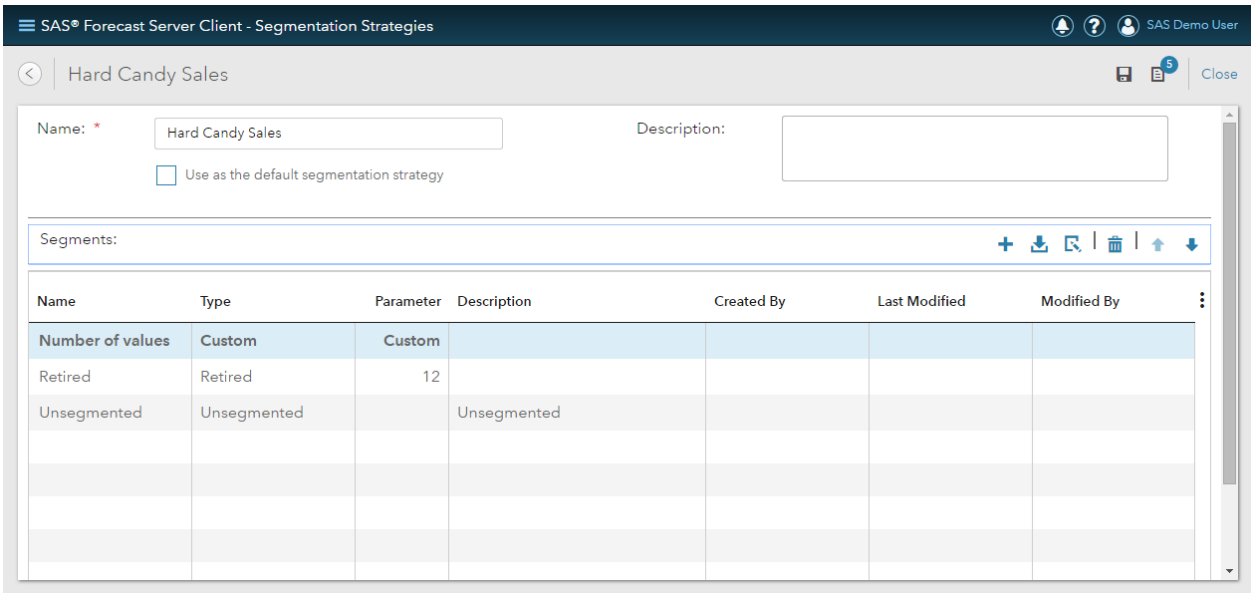

14 Click to save the new segmentation strategy called Hard Candy Sales. Click **Close** to return to the SAS Forecast Server Client - Segmentation Strategies page.

Working with Code-based Segmentation Strategies

What is a Code-based Segmentation Strategy?

A code-based segmentation strategy is when you write the code that SAS Forecast Server Client should use for the segmentation strategy. This code is saved in a file on the metadata server. You must import the code into SAS Forecast Server Client.

On the SAS Forecast Server Client - Segmentation Strategies page, code-based strategies are labeled **Code-based**. An example of a code-based segmentation strategy that is shipped with SAS Forecast Server Client is SAS Demand Classification and

Clustering. When you open a code-based segmentation strategy, the table lists the code parameters that were defined by the author of the strategy. When you create a code-based segmentation strategy, you cannot set the values for any parameters defined in the code. The list of parameters that appears in the segmentation strategy page is read-only. However, if you use this code-based strategy in your project, you can specify these parameter values when you are setting up or editing the project.

Example: Import SAS Demand Classification and Clustering Strategy

To import the SAS Demand Classification and Clustering strategy:

1. On the SAS Forecast Server Client - Segmentation Strategies page, click  . The Import Segmentation Strategy window appears.
2. For the **Type** option, select **Code-based**. A list of all the available code-based strategies appears.

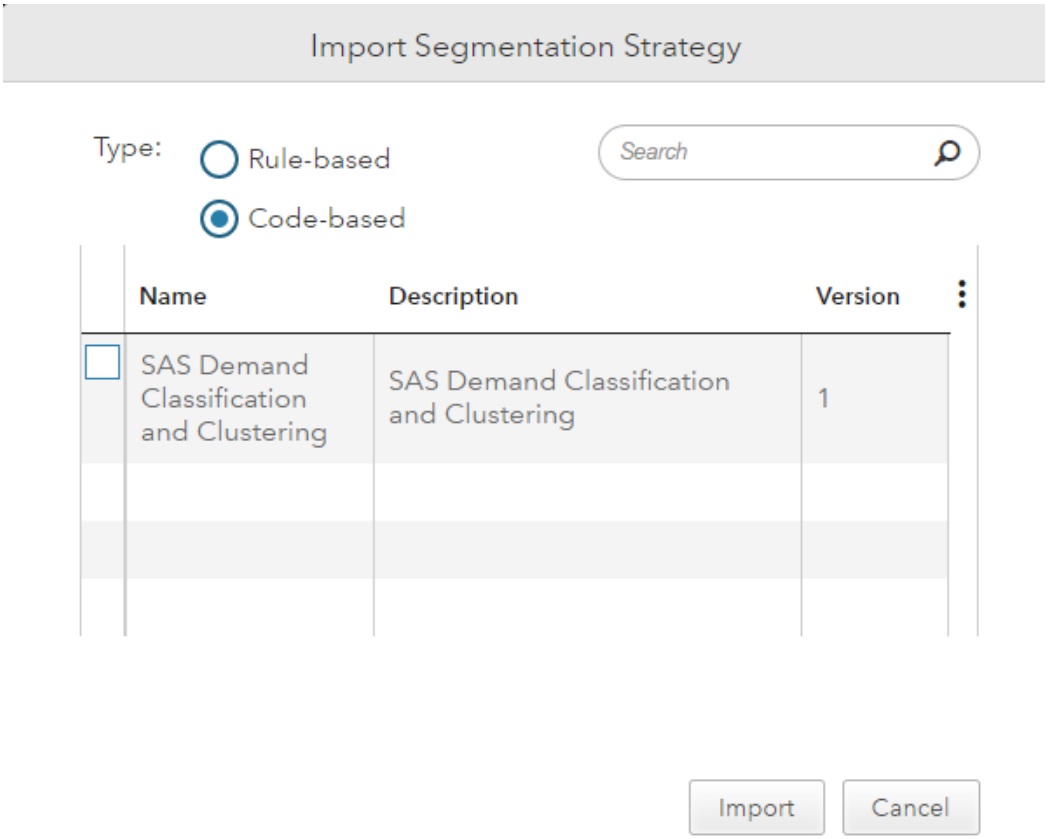

3 From the table, select the **SAS Demand Classification and Clustering** check box and click **Import**.

The parameters for the SAS Demand Classification and Clustering strategy appear on the New Segmentation Strategy page. These values are read-only.

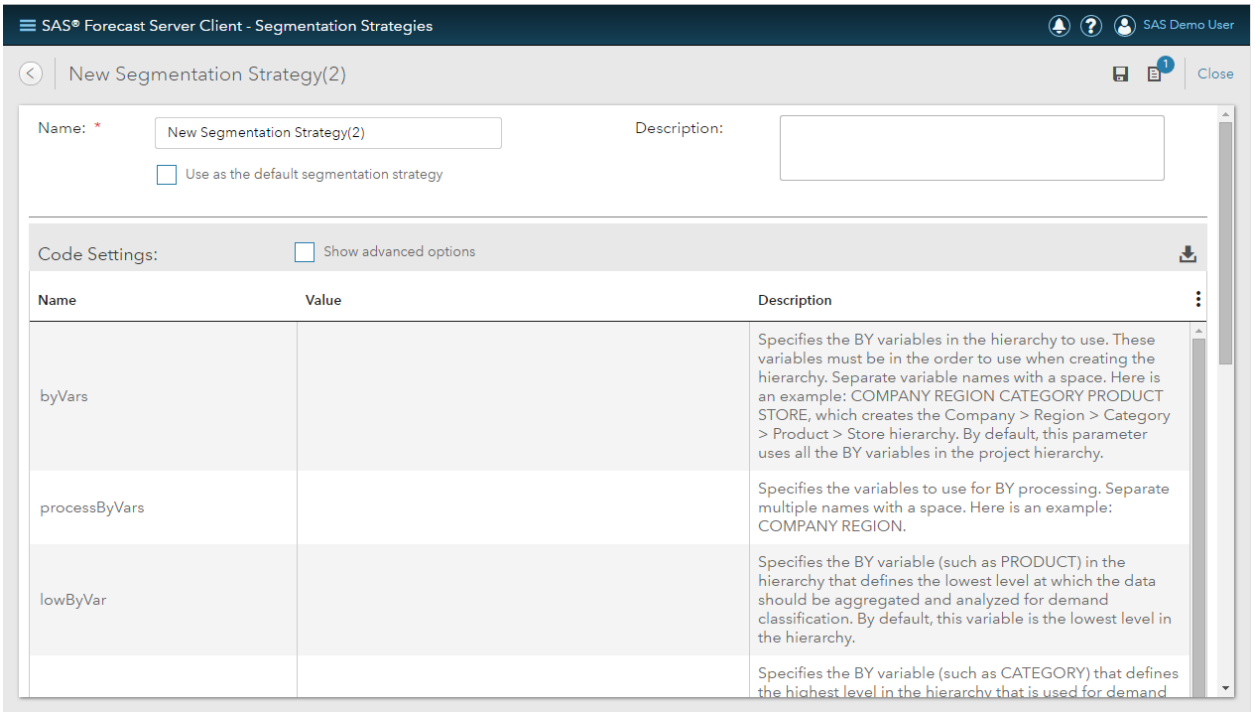

Specify the Default Segmentation Strategy

You can define multiple segmentation strategies in the SAS Forecast Server Client. However, you can have only one default segmentation strategy. The strategy that you specified as the default is used when you initially create and run a project. You can change the segmentation strategy after you run the project.

To set a default segmentation strategy:

1. On the SAS Forecast Server Client - Segmentation Strategies page, select the strategy that you want to use as the default.
2. When the definition of that segmentation strategy opens, select **Use as the default segmentation strategy**.

3 Click **Close**. Now, a check mark appears next to the segmentation strategy that you selected.

In this example, Prebuilt Segmentation Strategy is the default.

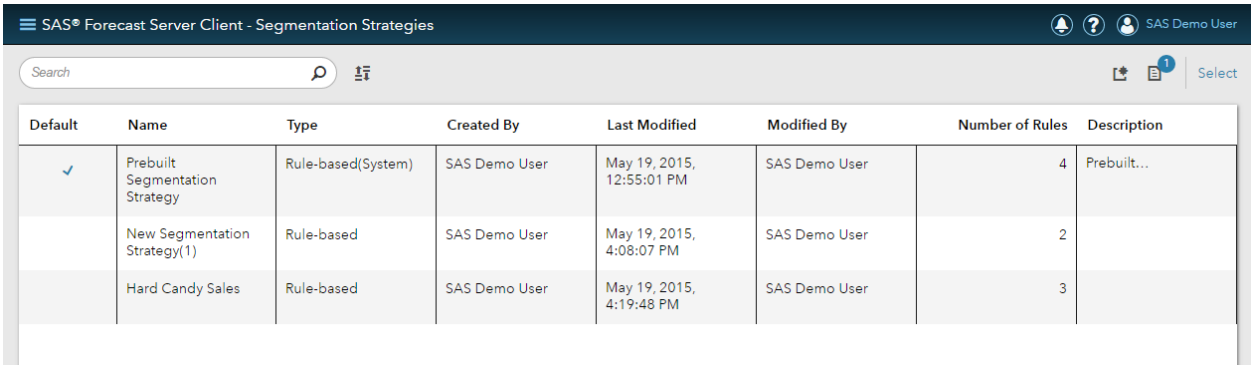

Default	Name	Type	Created By	Last Modified	Modified By	Number of Rules	Description
✓	Prebuilt Segmentation Strategy	Rule-based(System)	SAS Demo User	May 19, 2015, 12:55:01 PM	SAS Demo User	4	Prebuilt...
	New Segmentation Strategy(1)	Rule-based	SAS Demo User	May 19, 2015, 4:08:07 PM	SAS Demo User	2	
	Hard Candy Sales	Rule-based	SAS Demo User	May 19, 2015, 4:19:48 PM	SAS Demo User	3	

9

Defining Modeling Strategies

About Modeling Strategies

Before you can use a modeling strategy in a project, you must define it in the metadata. SAS Forecast Server Client is shipped with these modeling strategies:

- 4 segment-based modeling strategies: Intermittent, Retired, Short, and Unsegmented.
- Custom: a modeling strategy for any custom segments.

- several Demand Class modeling strategies: code-based modeling strategies for demand classification. For more information, see Appendix 3, “SAS Demand Classification and Clustering,” on page 159.
- SAS multistage modeling strategy: a code-based modeling strategy that you can use for multistage modeling. For more information, see “About the SAS Multistage Modeling Strategy” on page 74.

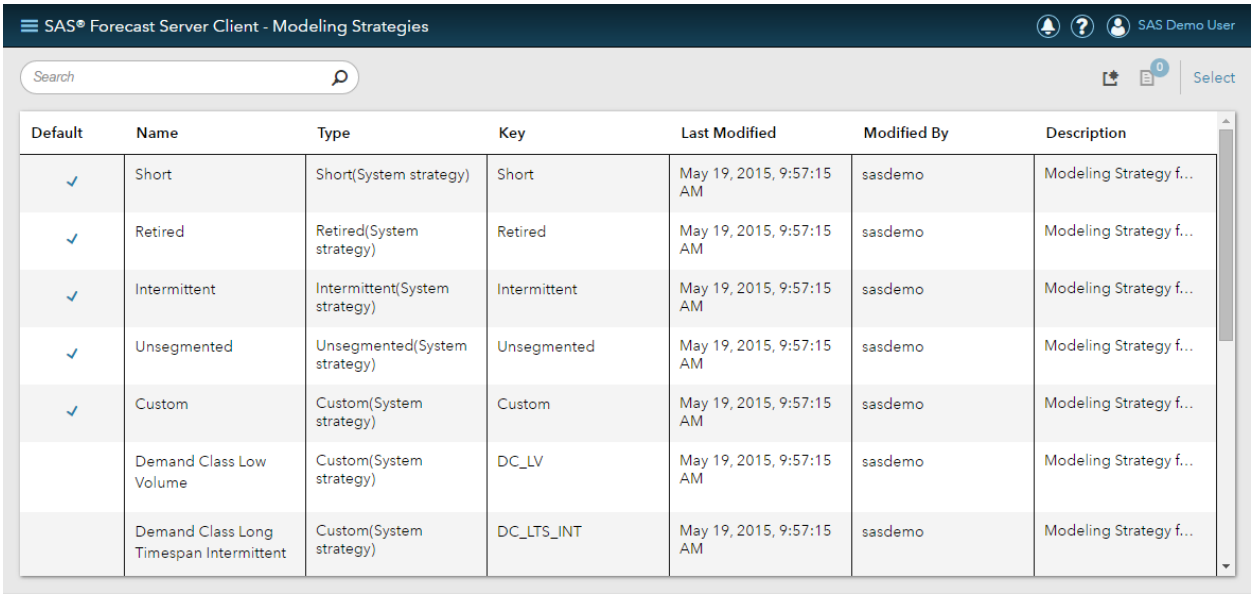

Default	Name	Type	Key	Last Modified	Modified By	Description
✓	Short	Short(System strategy)	Short	May 19, 2015, 9:57:15 AM	sasdemo	Modeling Strategy f...
✓	Retired	Retired(System strategy)	Retired	May 19, 2015, 9:57:15 AM	sasdemo	Modeling Strategy f...
✓	Intermittent	Intermittent(System strategy)	Intermittent	May 19, 2015, 9:57:15 AM	sasdemo	Modeling Strategy f...
✓	Unsegmented	Unsegmented(System strategy)	Unsegmented	May 19, 2015, 9:57:15 AM	sasdemo	Modeling Strategy f...
✓	Custom	Custom(System strategy)	Custom	May 19, 2015, 9:57:15 AM	sasdemo	Modeling Strategy f...
	Demand Class Low Volume	Custom(System strategy)	DC_LV	May 19, 2015, 9:57:15 AM	sasdemo	Modeling Strategy f...
	Demand Class Long Timespan Intermittent	Custom(System strategy)	DC_LTS_INT	May 19, 2015, 9:57:15 AM	sasdemo	Modeling Strategy f...

For each segment type, you can have one default modeling strategy. You specify the default modeling strategy by using the **Use as the default modeling strategy for any segment of this type** check box when you define the modeling strategy.

Working with Segment-Based Modeling Strategies

What Is a Segment-Based Modeling Strategy?

A segment-based modeling strategy is when the model settings are applied to a specific segment type in the project. For example, the Retired modeling strategy is used for any Retired segments that appear in the project.

Creating a Segment-Based Modeling Strategy

To create a segment-based modeling strategy:

1. Click ☰ to open the side menu.
2. Click **Modeling Strategies**. The SAS Forecast Server Client - Modeling Strategies page appears.
3. To create a new modeling strategy, click . The New Modeling Strategy page opens.
4. Specify a name and description for the new modeling strategy.
5. For the **Type** option, select **Segment-based**. Select the segment model type for this strategy. Examples of segment model types are Unsegmented, Intermittent, Short, and Custom.
6. To use these settings for all of the segments in your project with the same type, select **Use as the default modeling strategy for any segment of this type**.
7. (Optional) On the **Diagnostics** tab, specify the diagnostic options for the model. For more information, see “Understanding the Diagnostic Options for the Model” on page 70.

   Note: The available diagnostic options depend on your segment model type.

8 On the **Model** tab, set the options for model selection and model generation. For more information, see "Understanding Model Selection" on page 71 and "Setting the Options for Model Generation" on page 73.

9 Click to save the new modeling strategy. Click **Close** to return to the SAS Forecast Server Client - Modeling Strategies page.

Understanding the Diagnostic Options for the Model

The diagnostic options are available for only some segment types. On the **Diagnostics** tab, you can specify the diagnostic settings for the model.

These options are available:

Perform Intermittency
specifies whether to perform a test to determine whether a time series is intermittent. By default, this test is performed, and the default sensitivity is 2.

Perform Seasonality Test
performs the seasonality test. You must specify a significance probability for the test. Series with strong seasonality have small test probabilities. A significance level of 0 always implies seasonality. A significance level of 1 always implies no seasonality. The default value for the significance probability is 0.01.

Minimum number of seasonal cycles for a seasonal model
specifies the minimum number of seasonal cycles for a seasonal model. No seasonal model is fit to a series that contains fewer observations than *n* multiplied by the seasonal length.

Minimum number of observations for a trend model
specifies the minimum number of observations for a trend model. No trend model is fit to a series that contains fewer than *n* nonmissing values. Incorporation of a trend is checked only for the following models:

- simple smoothing models
- unobserved component models

- ARIMA models with no differencing of the dependent variable

Minimum number of observations for a non-mean model
specifies the minimum number of observations for a non-mean model. Any series with fewer than *n* nonmissing values are not fit using the models in the selection list. Instead, these models are forecast as the mean of the observations in the series.

Functional transformation (dependent)
specifies the functional transformation for the dependent variable.

Forecast
specifies the type of forecast data to obtain when a back transformation is applied in the system-generated models. Forecasts can be based on the mean or median.

Diagnose independent variables separately
specifies whether SAS Forecast Server Client should diagnose the independent variables. From the drop-down list, specify whether to transform or trend the independent variables.

Outlier detection (ARIMA models only)
specifies the criteria for detecting outliers in ARIMA models. You can specify the number of outliers to detect and the significance level for these outliers. You can also specify the maximum percentage of series that can be identified as outliers.

Understanding Model Selection

What Is the Forecast Model?

The forecast model is the model that the SAS Forecast Server Client uses to generate a numerical prediction of the future values for the time series. By default, the SAS Forecast Server Client chooses the best-performing model in the model selection list as the forecast model. The best-performing model is determined by the holdout sample and the selection criterion that you specified. The best-performing model can change over time as the performance changes for the other models in the list.

Setting the Options for Model Selection

On the **Models** tab, you can specify the values for model selection.

Selection criterion
specifies the criterion to use when selecting the best forecast model.

Use holdout sample for model selection
specifies that you want SAS Forecast Server Client to use the specified portion of the input data to determine the accuracy of the model. The series for the holdout sample come from the last observations in the input data set. By default, the holdout sample is 2.

In the **Maximum percentage of series that holdout sample can be** box, you specify the size of the holdout sample as a percentage of the length of the dependent time series. The default value is 5%.

If you specify both the holdout sample to use for model selection and the maximum percentage of the series for the holdout sample, then the size of the holdout sample is *min(holdout-sample,maximum percentageT*), where *T* is the length of the dependent time series with beginning and ending missing values removed.

Maximum number of ending zero values for non-zero model
specifies the maximum number or percentage of trailing value for a nonzero model.

If you specify both the **Maximum number of ending zero values for non-zero model** and **Maximum percentage of ending zero values for non-zero model**, then a nonzero model is considered only if the series meets two conditions. That is, the series must have less than or equal to the maximum number of trailing zero values, and the number of trailing zero values must be less than or equal to the specified percentage of the number of nonmissing and nonzero values of the entire series.

Maximum number of observations to perform the end-zeros test
specifies the threshold for the series length that is required to enable the end zero test. By default, this threshold is 1.

When the end-zeros test is run, SAS Forecast Server Client considers a model only if it meets one of two conditions. That is, the series either must have less than or equal to the maximum number of trailing zero values, or it must have the maximum percentage of trailing zero values relative to the number of nonzero values in the entire series. If the series ends with numerous zero values, the forecasts are set to zero.

Setting the Options for Model Generation

On the **Models** tab, you can specify the options for model generation.

For model generation, select the models to fit to each series. You must select at least one model.

You can choose from these options:

- **System-generated ARIMA**
- **System-generated exponential smoothing models**
- **System-generated unobserved components model**
- **Models from an external list**

For model generation, you can also set these options:

Only fit system-generated exponential smoothing models at the lowest level of the hierarchy
specifies to fit system-generated exponential smoothing models only at the lowest level of the hierarchy

Allow negative forecasts
specifies whether to allow negative forecasts. If you clear the **Allow negative forecasts** check box, any negative values in the forecast model are set to 0.

Create the component series data set
specifies whether to create the component series data set (called OUTCOMPONENT). By default, the component series data set is automatically created when you create a project. A component data set is created for each BY variable in the project. Because a component data set can also contain several other variables, these data sets can be quite large. As a result, a component data set can require additional disk space and computing time.

If you clear this option, you suppress the creation of this data set.

Create a forecast data set for independent variables
specifies whether to create a forecast data set for independent variables. In some situations, you might not know the future values of the independent variables. Depending on the model for your project, SAS Forecast Server Client might use these independent variables to create forecasts for your dependent variables. Select this option to create an output data set that contains the future values of the independent variables. You can review this data set to verify whether those values are reasonable.

Working with Code-Based Modeling Strategies

About the SAS Multistage Modeling Strategy

The rapid development of information technologies in the recent decade provides forecasters with huge amounts of data, as well as massive computing capabilities. However, “sufficient” data and strong computing power do not necessarily translate into good forecasts. Different industries and products all have their unique demand patterns. There is not a one-size-fits-all forecasting model or technique. For example, in the consumer package goods (CPG) industry, demand at store-SKU level is usually sparse and noisy, which makes it difficult to extract price and promotional effects. For high frequency data such as hourly grocery basket transactions, it is inappropriate and inefficient to apply traditional time series models. A good forecasting model must be customized for the data to capture the salient features and satisfy the business needs.

The SAS Multistage Modeling Strategy is shipped with SAS Forecast Server Client. This code-based strategy provides a general framework to build a forecasting system in three stages.

1 In the first stage, the strategy develops a model to extract salient features across multiple time series. Then the extracted features are transformed into an adjustment factor to generate the forecast at the aggregated level.

2 In the second stage, the feature extraction technique is applied again to generate forecasts for each individual time series at lower levels of the hierarchy.

3 In the third stage, the strategy combines the forecasts obtained from the previous two stages and conducts a top-down reconciliation to generate the final forecast.

To add the SAS Multistage Modeling Strategy to your list of modeling strategies, you must create a new code-based modeling strategy and import the SAS Multistage Modeling Strategy. Although this strategy is shipped with SAS Forecast Server Client by default, it does not appear in the list of modeling strategies on the SAS Forecast Server Client - Modeling Strategies page until you add it.

Create a Code-Based Modeling Strategy

To create a code-based modeling strategy:

1 Click ☰ to open the side menu.

2 Click **Modeling Strategies**. The Modeling Strategies page opens.

3 To create a new modeling strategy, click . The New Modeling Strategy page opens.

4 Specify a name and description for the new modeling strategy.

5 For the **Type** option, select **Code-based**.

6 To use these settings for all segments of the same type, select the segment type from the drop-down list and select the **Use as the default modeling strategy for any segment of this type** check box.

7 To import the code for the new modeling strategy, click **Import**. The Import Modeling Strategy window appears.

8 Select the name of the code-based strategy and click **OK**.

The parameters for the code-based strategy appear. You cannot edit these parameters when defining the modeling strategy. However, you can edit these parameters after you use the modeling strategy in your project.

Here is a subset of the parameters that appear if you import the SAS Multistage Modeling Strategy:

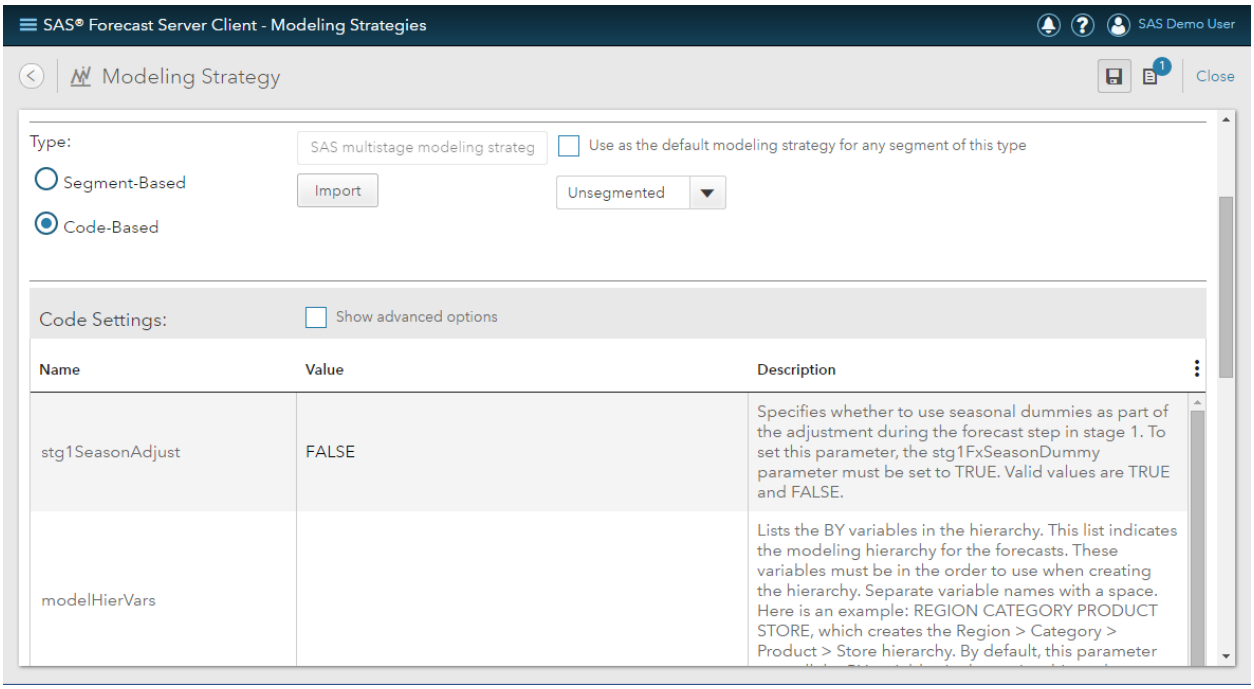

9 To save the changes to your modeling strategy, click . To return to the SAS Forecast Server Client – Modeling Strategies page, click **Close**. The new modeling strategy appears on this page.

Specify the Default Modeling Strategy for a Segment Type

You can define multiple modeling strategies in the SAS Forecast Server Client. For each segment type, you must specify the default strategy to use.

To set a default modeling strategy for a segment type:

1 On the SAS Forecast Server Client - Modeling Strategies page, select the strategy that you want to use as the default.

2 When the definition of that modeling strategy opens, select **Use as the default modeling strategy for any segment of this type**.

3 Click **Close**. Now, a check mark appears in the **Default** column for the modeling strategy that you selected.

10

Creating Exception Strategies

About Exception Strategies

Exception strategies are rules that you create to identify series that meet or do not meet specific criteria or a specific criterion. Each exception strategy can contain one or more rules. Within the strategy, you can specify whether a rule is active.

In SAS Forecast Server Client, you can create exception strategies for the modeling and tracking nodes.

Create an Exception Strategy

To create an exception strategy:

1 Click ☰ to open the side menu.

2 Click **Exception Strategies**. The SAS Forecast Server Client - Exception Strategies page appears.

3 To create a new exception strategy, click . The New Exception Rule Strategy page opens.

4 Specify a name and a description for the exception strategy.

5 Select whether the exception strategy is for a modeling exception rule or a tracking exception rule.

Note: An exception strategy can contain either modeling exception rules or tracking exception rules, but not both.

6 To define the conditions for the exception strategy, click +.

7 Specify the category that contains the property that you want to use in the condition. You can choose from these categories:

- reconciled statistics of fit
- model statistic of fit

8 To create the exception rule, set the **Property**, **Condition**, and **Value** options.

9 Click **OK**. The rule that you created now appears in the table for the exception strategy. By default, this new rule is active in the strategy.

Example: Create an Exception Rule Strategy for Modeling

When SAS Forecast Server Client is modeling your data, you might want to identify any series that has a mean absolute Percent Error (MAPE) greater than 5.

To create this exception rule strategy:

1 Click ≡ to open the side menu.

2 Click **Exception Strategies**. The SAS Forecast Server Client - Exception Strategies page appears.

3 To create a new exception strategy, click . The New Exception Rule Strategy page opens.

4 For the name of the strategy, enter **`MAPE > 5`**.

5 For the **Type** option, select **Modeling exception rule**.

6 To define the condition for the exception strategy, click . The Add Modeling Exception Rule window appears.

7 In this window, enter **`MAPE > 5`** as the name of the rule.

8 From the **Category** drop-down list, select **Reconciled statistic of fit**.

9 In the table, set these options:

- From the **Property** drop-down list, select **Mean Absolute Percent Error**.
- From the **Condition** drop-down list, select **>**.
- In the **Value** column, enter `5`.
- Verify that the **Active** check box is selected for this rule.

10 Click **Add**.

This new rule now appears on the MAPE > 5 page.

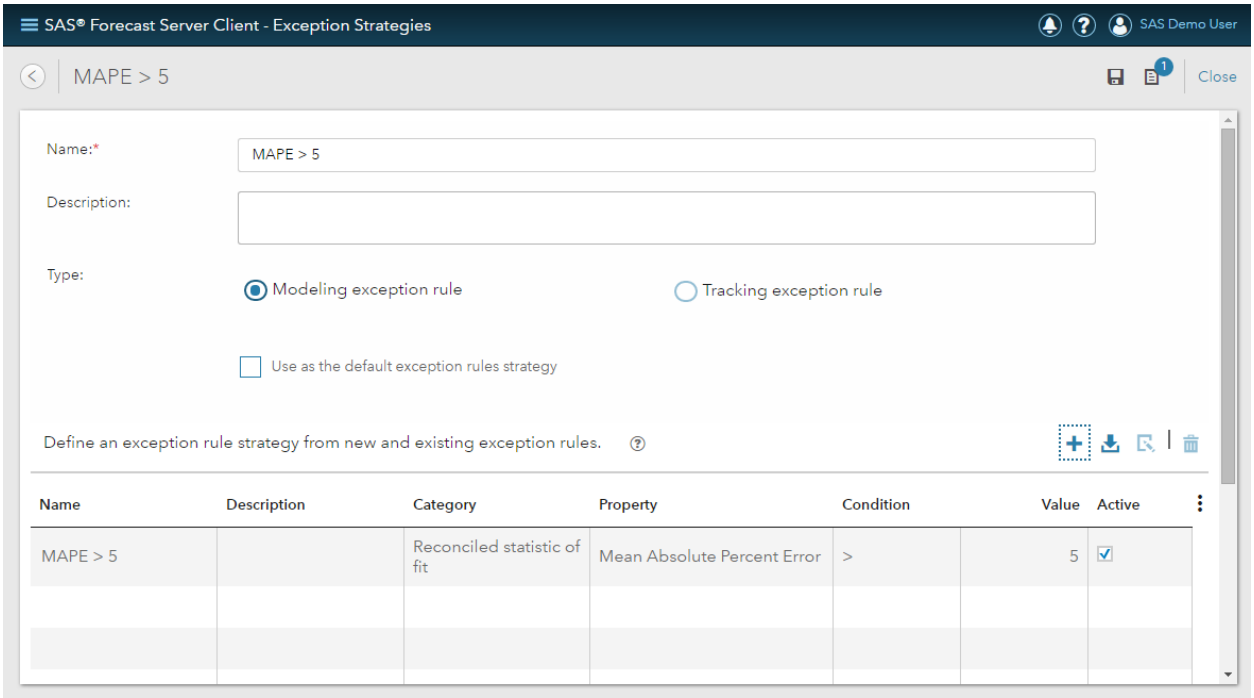

11 To save this modeling exception strategy, click . Click **Close** to return to the SAS Forecast Server Client - Exception Strategies page.

The new MAPE > 5 exception strategy is now available from the list of exception strategies.

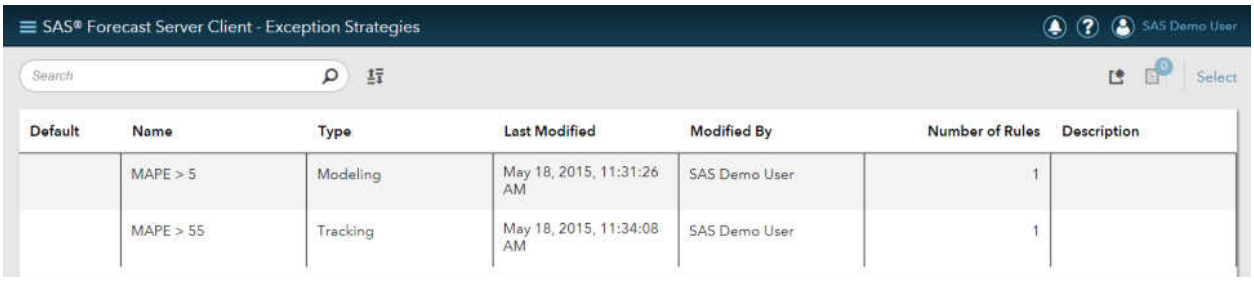

Now, you can add this exception strategy to the modeling node for an existing project.

12 Open the project template and double-click the modeling node where you want to add this exception strategy.

For this example, double-click the **Unsegmented Forecasts** node.

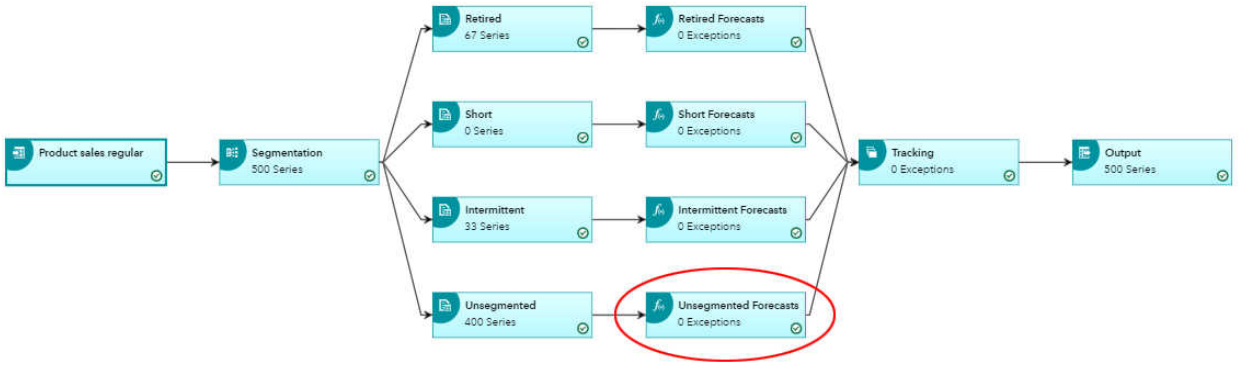

13 When the results and settings for the **Unsegmented Forecasts** node open, click the **Settings** tab.

14 In the selection pane, click **Exception Rules**.

15 To add the MAPE > 5 exception rule that you created, click . The Import Exception Rules Strategies window appears.

16 Select the check box for the MAPE > 5 strategy and click **Import**.

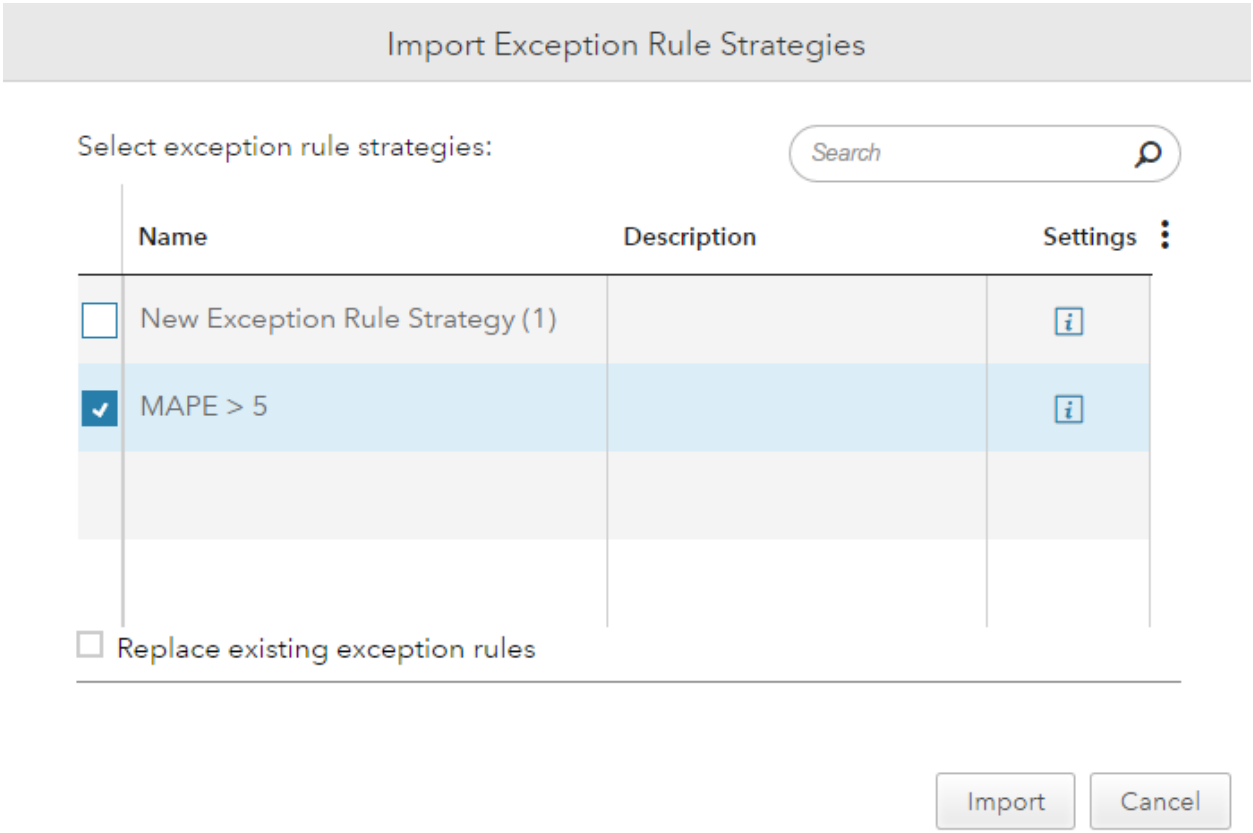

The MAPE > 5 exception strategy now appears in the table. By default, the **Active** check box is selected, so this rule is used the next time you run this node or the entire project.

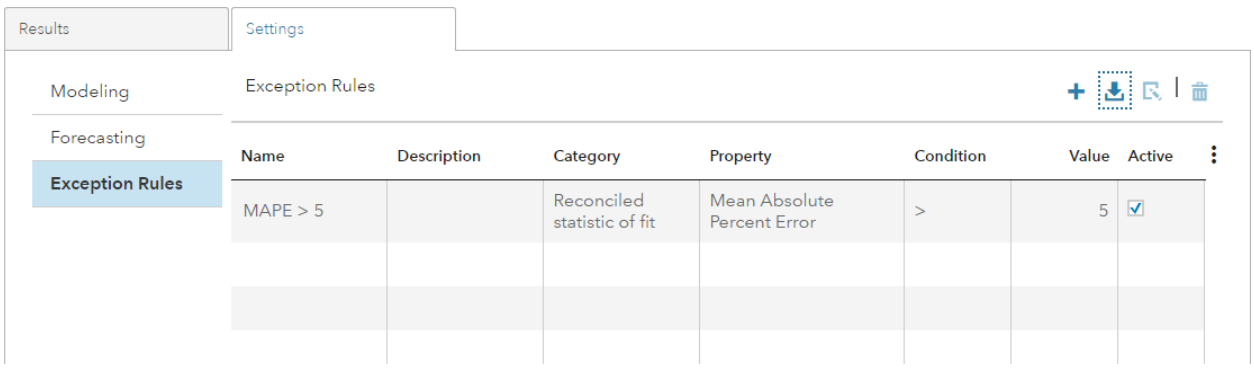

17 Click **Close** to return to the project template.

The **Unsegmented Forecasts** node now shows 398 exceptions.

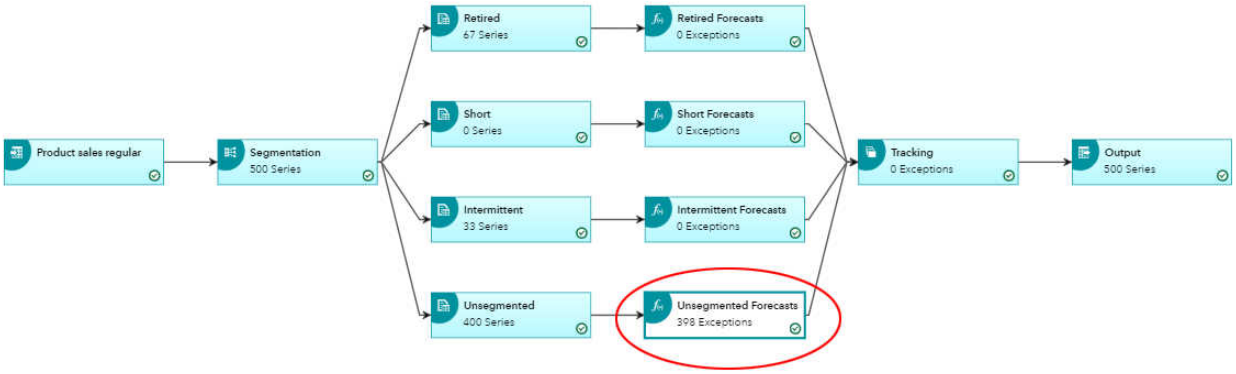

Specify the Default Exception Strategy

You can define multiple exception strategies in SAS Forecast Server Client. You can also specify a default exception strategy for models and a default exception strategy for tracking.

To set a default exception strategy:

1. On the SAS Forecast Server Client - Exception Strategies page, select the strategy that you want to use as the default.
2. When the definition of that exception strategy opens, select **Use as the default exception rules strategy**.
3. To save this exception strategy, click 💾. Click **Close** to return to the SAS Forecast Server Client - Exception Strategies page.

Part 4

Working with a SAS Forecast Server Client Project

11

Understanding Projects

What Is a Project?

SAS Forecast Server Client organizes the files that are created for a project. Each project is saved on the SAS Workspace Server. (This is the same server on which the SAS processing occurs.) These projects are saved in the ***SAS-configuration-directory*\Lev1\AppData\SASForecastServerClient\projects** directory.

You open projects from the **SAS Forecast Server Client - Projects** page. From this page, you can see all the projects that you can access. To open a project, double-click the project name in the table. The project template appears.

Only one user can have Write access to a project at one time.

Note: In SAS Forecast Server Client, project names are not case sensitive. As a result, if you already have a project called Forecast1, you cannot create a project called forecast1.

Running a Project

Run the Entire Project

To run the entire project, click ▶.

Here is the location of the Run icon in the toolbar for the project template.

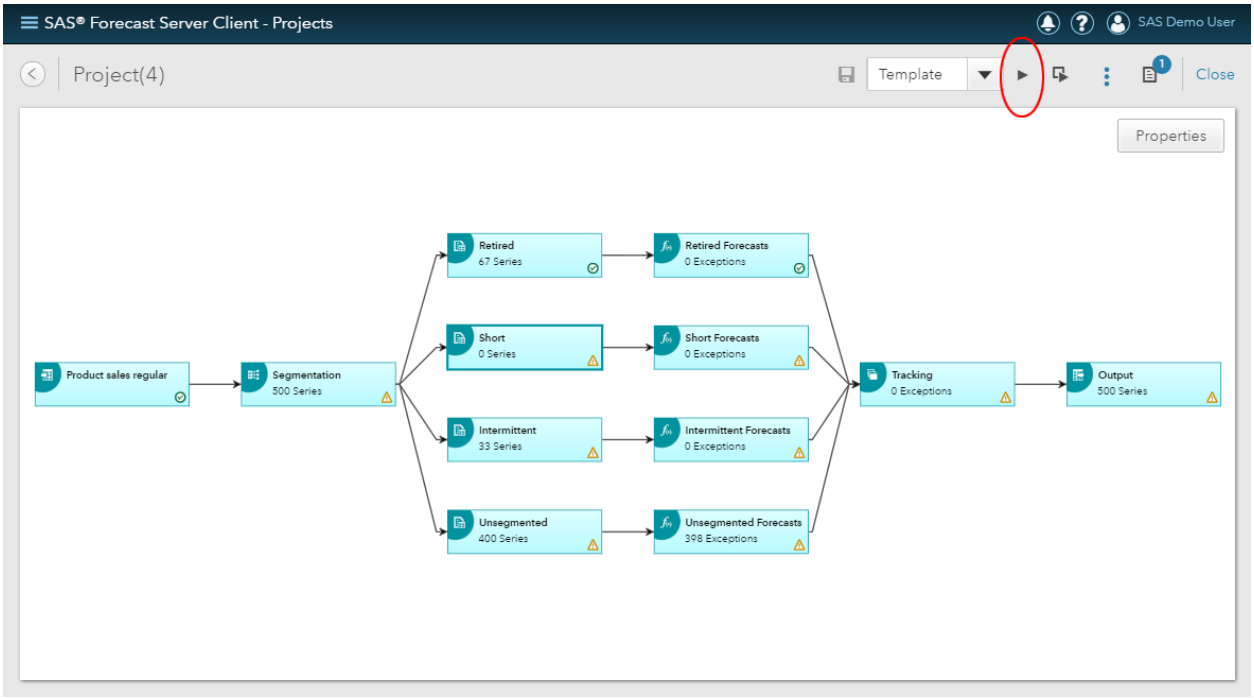

Note: When running a project, do not select that project again from the SAS Forecast Server Client – Projects page until the project finishes running.

Run the Selected Node in a Project

To run the project from the selected node:

1. In the project template, select the node where you want to start.
2. In the toolbar, click  .

SAS Forecast Server Client runs the selected node and any nodes after the selected node.

Rename a Project

To rename a project:

1. In the SAS Forecast Server Client – Projects page, click **Select**. The first column in the project list now contains a check box for each project.
2. Select the check box for the project that you want to delete. You can delete multiple projects at one time.
3. Click  . The Rename window appears.
4. Specify the new name for the project and click **OK**.

Copy an Existing Project

To copy an existing project:

1. In the SAS Forecast Server Client – Projects page, select the project that you want to copy. The project template opens.
2. In the toolbar, click ⋮ and select **Save As**. The Save As window appears.
3. Specify a name for the copied project and click **OK**.

View the SAS Log

You can use the SAS log to troubleshoot any errors that are generated in the project. The SAS log is also useful if you need to contact SAS Technical Support.

To view the SAS log for the project:

1. Open the project template.
2. In the toolbar, click ⋮ and select **View SAS Log**.

The log file is downloaded by your web browser. To view the log, select a text application (such as NotePad).

Delete a Project

1. In the SAS Forecast Server Client – Projects page, click **Select**. The first column in the project list now contains a check box for each project.

2 Select the check box for the project that you want to delete. You can delete multiple projects at one time.

3 Click 🗑.

Working with Archived Projects

Archive a SAS Forecast Server Client Project

1 In the SAS Forecast Server Client – Projects page, click **Select**. The first column in the project list now contains a check box for each project.

2 Select the check box for the project that you want to archive.

3 Click ⋮ and select **Archive**. The Archive window appears.

4 Specify a name for the archived project and click **OK**.

Unarchive a Project

In SAS Forecast Server Client, you can unarchive projects that were created with SAS Forecast Server Client or projects that were created in SAS Forecast Studio. However, only the information needed for the data definition is imported from SAS Forecast Studio projects. No other features from the SAS Forecast Studio project are available in SAS Forecast Server Client.

To unarchive a project:

1 On the SAS Forecast Server Client – Projects page, click ⋮ and select **Unarchive**. The Unarchive window appears.

2 Select the project that you want to unarchive.

3 Specify a name for the project. Click **OK**.

12

Working with Data Definitions in Your Project

About the Data Definition Node

Where Is the Data Definition Node?

In the project template, the data definition node is the first node in the workflow.

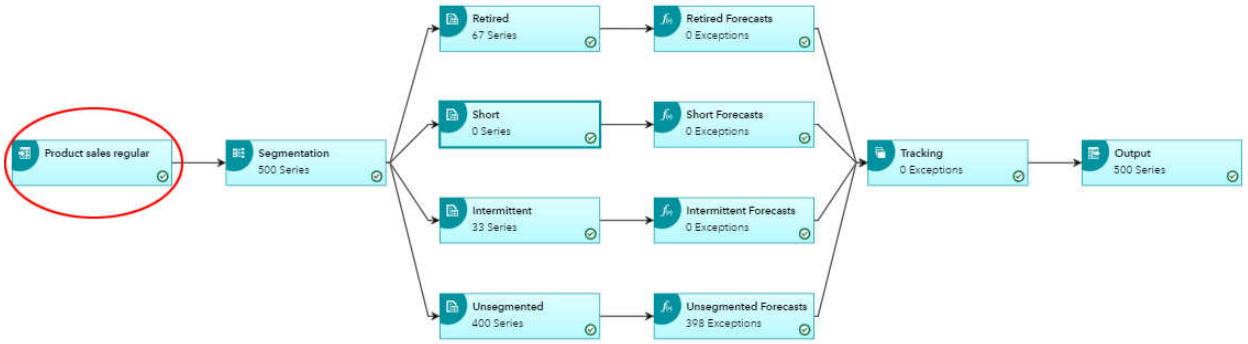

The data definition was defined before you created the project. For more information, see “What Is a Data Definition?” on page 21.

However, after you run the project, you might want to make a few changes to the settings in the data definition. You can do this by double-clicking the data definition node.

Viewing the Variables and Role Assignments for the Data Definition

To view the results for the data definition node, double-click this node and open the **Results** tab. The **Results** tab lists the names of the variables in the input data source. This tab also shows the assigned role for each variable. You can also view a frequency chart for the selected variable. To select multiple variables, press Ctrl and select the variables that you want to use. When you select multiple variables, multiple frequency charts are displayed. To view the data used to create the frequency chart, click  .

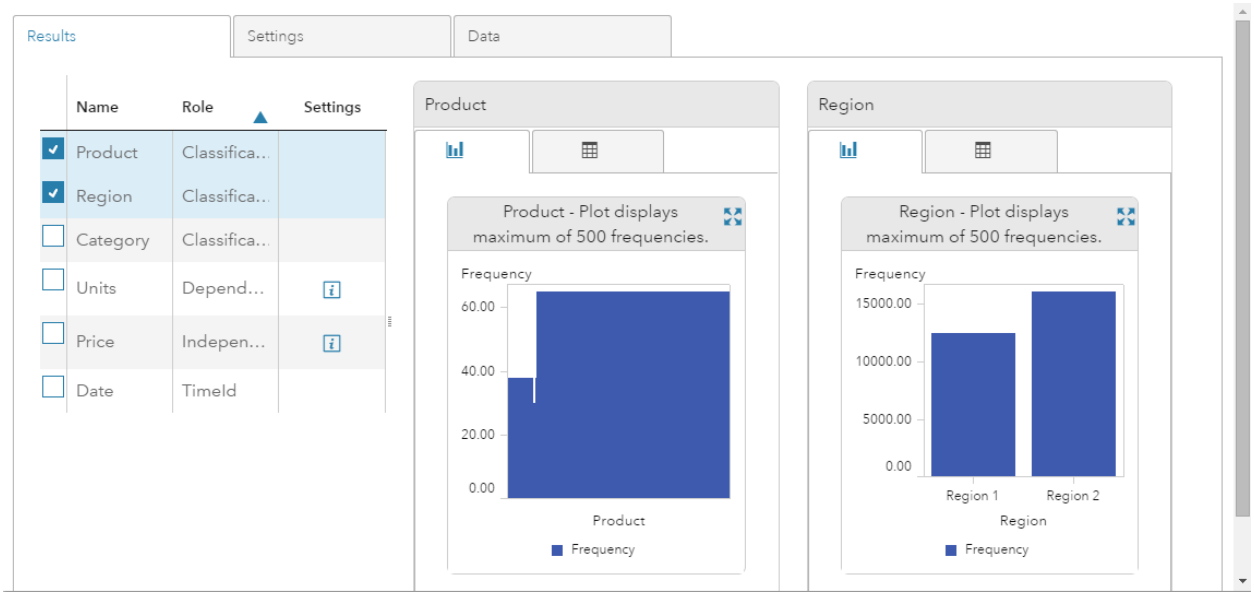

Viewing the Properties, Time ID Variable, and Forecast Settings for the Data Definition

From the **Settings** tab, you can view the properties and settings of the data definition. If you change any of the options on the **Settings** tab, these changes apply only to the current project. To change the settings for all future projects, you must create a new data definition.

The **Settings** tab contains the following information:

- the properties for the data definition. These properties include descriptive information, such as the name of the data set, the library where the data set is saved, when the data set was created, the number of rows and columns in the data set, and so on.
- the variable assignments for the data definition. Click **Variable Settings** to see how variables from the input data source are used in the project. You can also view the values used for treating missing values and zero values.
- the time dimensions settings of your data.

- the settings (such as the number of periods to forecast and the reconciliation level) for generating forecasts.
- the forecasted hierarchy levels. Click **Hierarchy preview** to see the hierarchy that was specified for the data definition. You cannot change the hierarchy from the **Settings** tab. To change the hierarchy, you must create a new data definition and then a new project.

Viewing the Input Data Source

The **Data** tab displays a view of the input data source. By default, the view displays all the columns in the input data source.

Update the Project Data

While you are working, another user at your site might update the data in the input data source.

To incorporate the latest version of the project data in your project:

1. In the project template, double-click the node for the data definition.
2. Click the **Settings** tab. If the input data source has changed, a “New data available” message appears.
3. Click for SAS Forecast Server Client to update the project data. This option is available only when the input data has changed.

Change the Variable Settings for a Project

After you create a project, you might want to assign roles to an additional variable, or you might want to change one of the many data preparation options.

To change the variable settings for a project:

1. In the project template, double-click the node for the data definition.
2. Click the **Settings** tab.
3. Click **Variable Settings**. The Variable Settings window appears.
4. Select the check box of the variable that you want to use.
5. From the **Role** drop-down list, specify the role for the selected variable. Specify any additional options. When you are finished, click **Save**.
6. To return to the project template, click **Close**.

Change the Time ID Settings for a Project

Although you cannot change the time ID variable for a project, you can change the interval and the advanced settings (such as the multiplier, shift, and seasonal cycle length).

To change the time ID settings for a project:

1. In the project template, double-click the node for the Data Definition.
2. Click the **Settings** tab.

3 From the **Interval** drop-down list, select the time interval to use.

Note: If you change the time interval, any tracked data is lost. For more information about tracking, see Chapter 15, “Tracking Your Forecasts,” on page 125.

4 To change the multiplier, shift, or seasonal cycle length, click **Advanced Settings**. For more information about these options, see “Understanding SAS Time Intervals” on page 37.

5 To save your changes, click **Close**.

Specify the Number of Periods to Forecast

You specified the number of periods to forecast when you created the data definition. However, after running the project, you might want to change this value.

To change the number of periods to forecast:

1 In the project template, double-click the node for the Data Definition.

2 Click the **Settings** tab.

3 In the **Number of periods to forecast (horizon)** box, specify the number of periods into the future for which multistep forecasts are made. The larger the horizon value, the larger the variance at the end of the horizon. By default, the horizon is 12.

4 In the **Confidence limit** box, specify the significance level for the series. By default, this significance level is 0.05, which is a 95% confidence interval.

5 To save your changes, click **Close**.

Change the Reconciliation Level for the Project

The reconciliation level is specified when you create a data definition. However, you can change the reconciliation level for your project.

To change the reconciliation level for your project:

1. In the project template, double-click the node for the Data Definition.
2. Click the **Settings** tab.
3. From the **Reconciliation level** drop-down list, select the reconciliation level that you want to use.
4. To save your changes, click **Close**.

13

Working with Segmentation Strategies in Your Project

Viewing the Segmentation Strategies for Your Project

To view the segmentation strategies for the project, double-click the **Segmentation** node in the project template.

In the project template, the**Segmentation** node follows the data definition node.

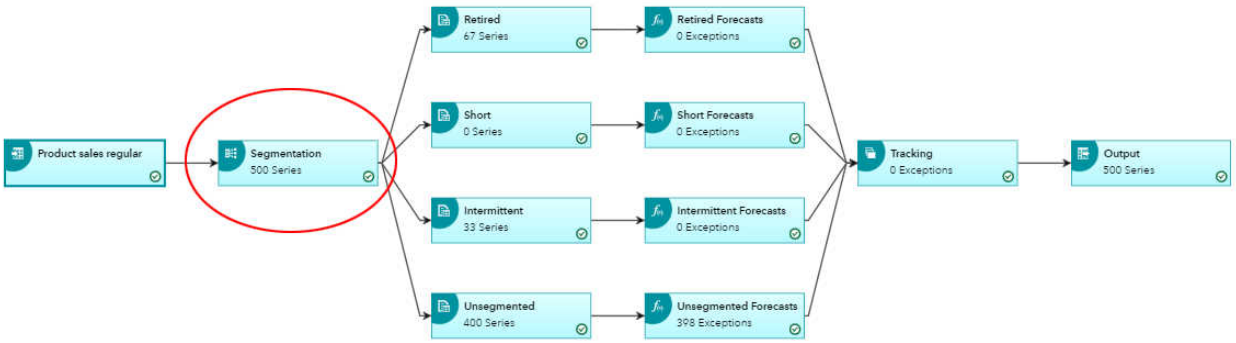

After you double-click the **Segmentation** node, click the **Results** tab to view the series at the default reconciliation level. The default reconciliation level is identified by an asterisk (*) in the **Level** drop-down list.

Note: Results are available only for the levels where SAS Forecast Server Client generates forecasts. Some code-based modeling strategies (such as the Demand Classification and Clustering strategy that is shipped with SAS Forecast Server Client) might not generate forecasts for all levels in the hierarchy. However, you should always see results for the lowest level in the hierarchy.

Click the **Settings** tab to view all the segments for the project. In this example, there are four segments: Retired, Short, Intermittent, and Unsegmented.

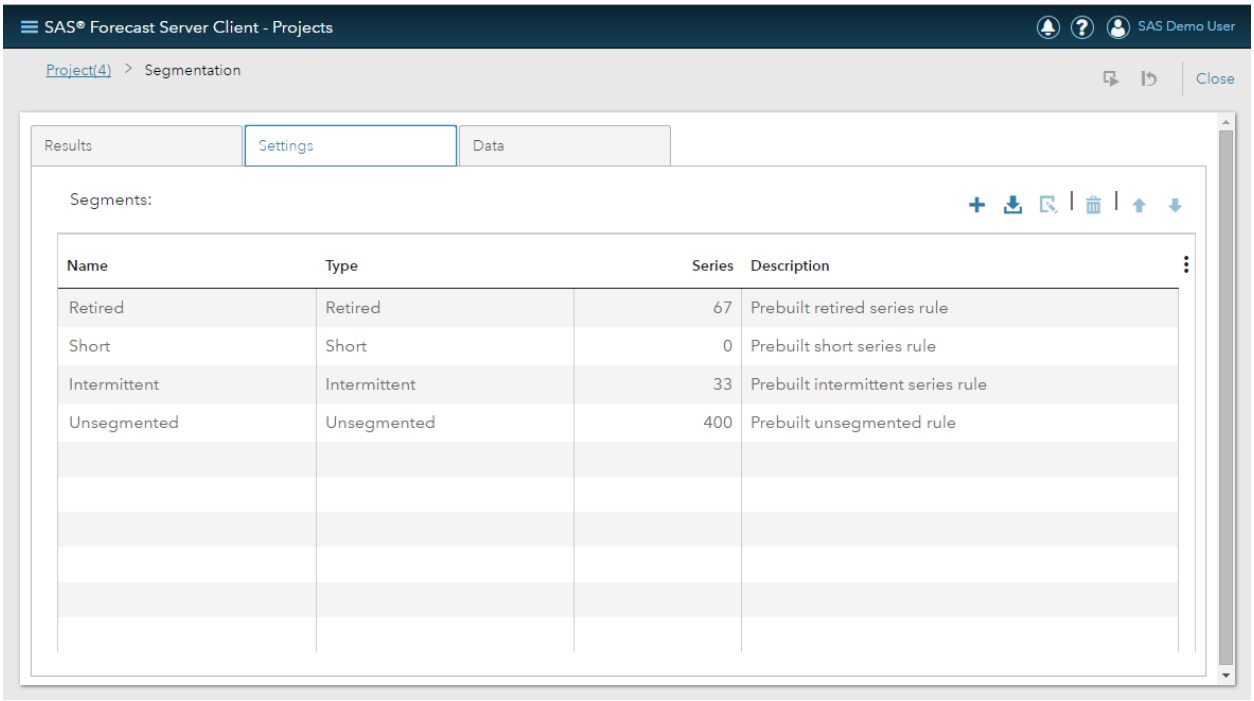

When you run the project, you can see the number of series in each segment. In this example, the Retired strategy resulted in 67 series in the **Retired** node. The Intermittent strategy resulted in 33 series in the **Intermittent** node. The remaining 400 series did not meet the criteria of any of the segmentation strategies, so these remaining 400 series are included in the **Unsegmented** node.

Add Rules to the Segmentation Strategy

After you have created the project, you cannot change the name of the segmentation strategy. However, you can change the rules for the segmentation strategy.

To add rules to the segmentation strategy:

1 Double-click the **Segmentation** node in the project template.

2 To view the rules for the current strategy, click the **Settings** tab.

For this project, there are four segmentation rules and thus four segments: Retired, Short, Intermittent, and Unsegmented.

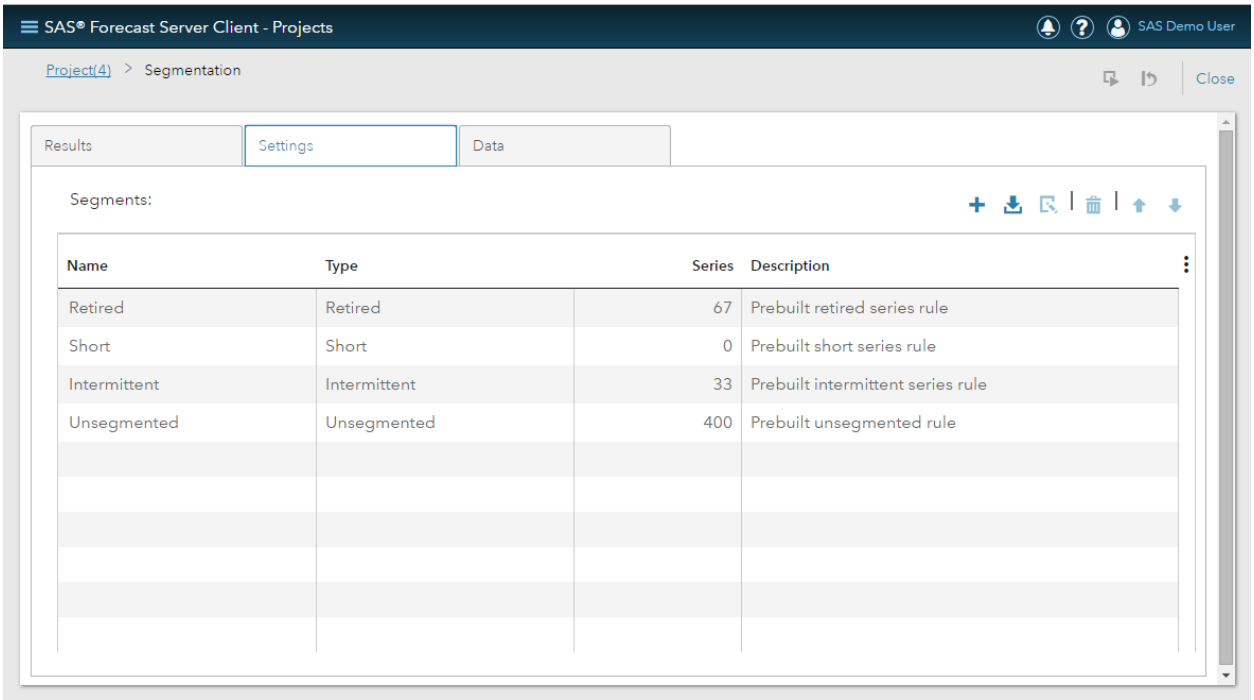

3 To add another rule to this segmentation strategy, click ✚. The Add Segments window appears.

4 From the **Type** drop-down list, select the type of segment that you want to create.

You can choose from these options:

- **Intermittent** — specifies the series where the average separation between nonzero observations exceeds the specified number.
- **Retired** — specifies the series where the number of observations since the last observation exceeds a specified time period.
- **Short** — specifies the series that contain less than the specified number of observations.

- **Custom** — enables you to create a rule using the value of the BY variable or the value of the descriptive statistics for the dependent variable or an independent variable.

Click **OK**.

5 Use the up and down arrows to specify the order in which the rules should be applied to the series. For example, suppose you had three segments in the table: Retired, Short, and Unsegmented (in that order). A series could meet the criteria for both the Retired and Short segments. However, because the Retired segment appears in the table first, the series is assigned to the Retired segment. This series is not included in the Short segment.

Edit an Existing Segment Rule

To view the rules for the segmentation strategy, double-click the segmentation node in the project template. Then click the **Settings** tab.

If this is a rule-based segmentation strategy, a list of the defined segments appears.

In this example, there are four segments: Retired, Short, Intermittent, and Unsegmented.

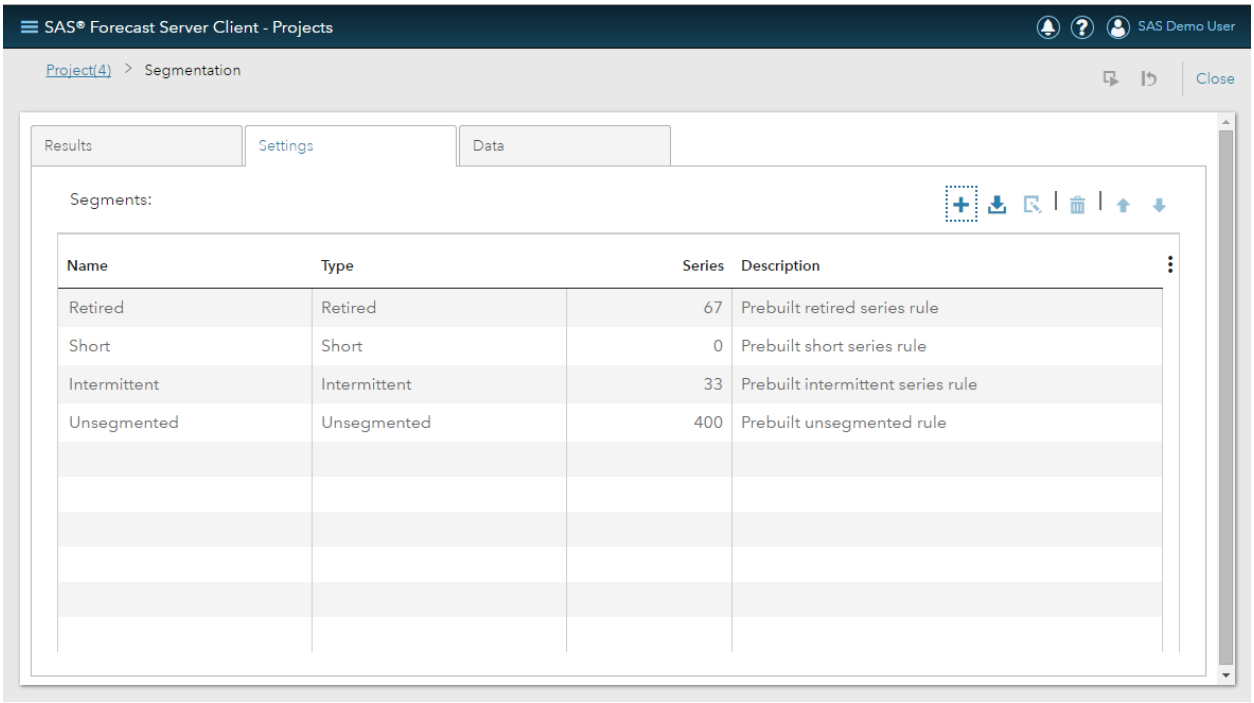

From the table, select the segment to edit and click . In the Edit Segments window, specify the new values that you want to use and click **OK**.

If you are working with a code-based segmentation strategy, a table lists all of the parameters that are defined in the code. You can specify values for each of these parameters. Use the information in the description to determine what format to use for each specified value.

Here is an example of the parameters for the SAS Demand Classification and Clustering strategy.

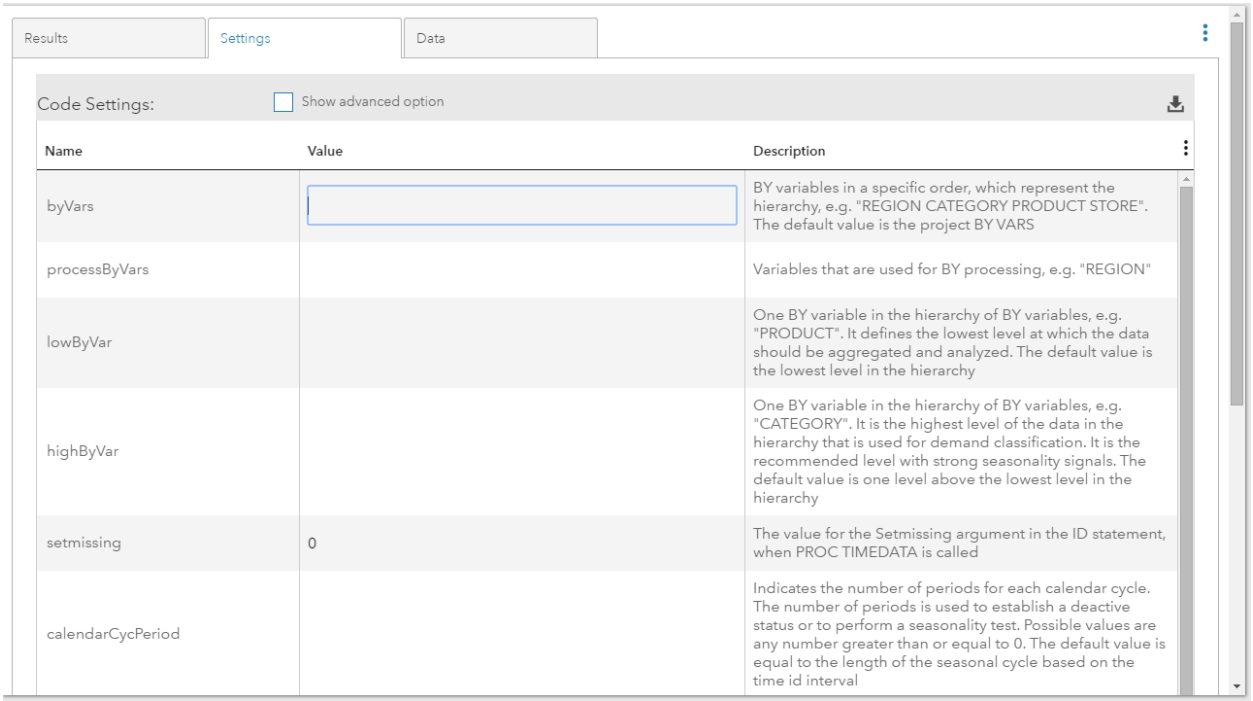

When you finish making your changes to the segmentation rules, click **Close**. SAS Forecast Server Client saves the changes but does not run the project. In the project template, the Unapplied Changes icon (⚠) appears on the **Segmentation** node. To run the project, click ▶.

Working with the Segments

About the Segment Nodes

In the project template, the number of segment nodes depends on how many segments are defined by your segmentation strategy.

In this example, the project template shows four segment nodes: Retired, Short, Intermittent, and Unsegmented.

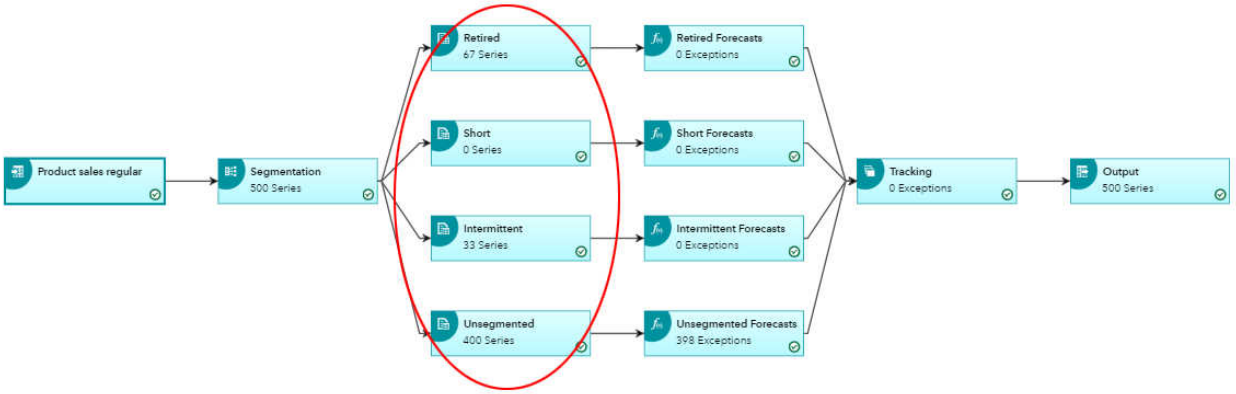

From the project template, double-click the segment node to edit the rules for that segment.

View the Series in the Segment

To view the series in the selected segment, click the **Results** tab. You can select an individual series to view its properties (for example, the first observation, number of missing values, and so on). The time series envelope plot gives a summary view of all time series in the selected segment.

Here is an example of the **Results** tab for the **Retired** node.

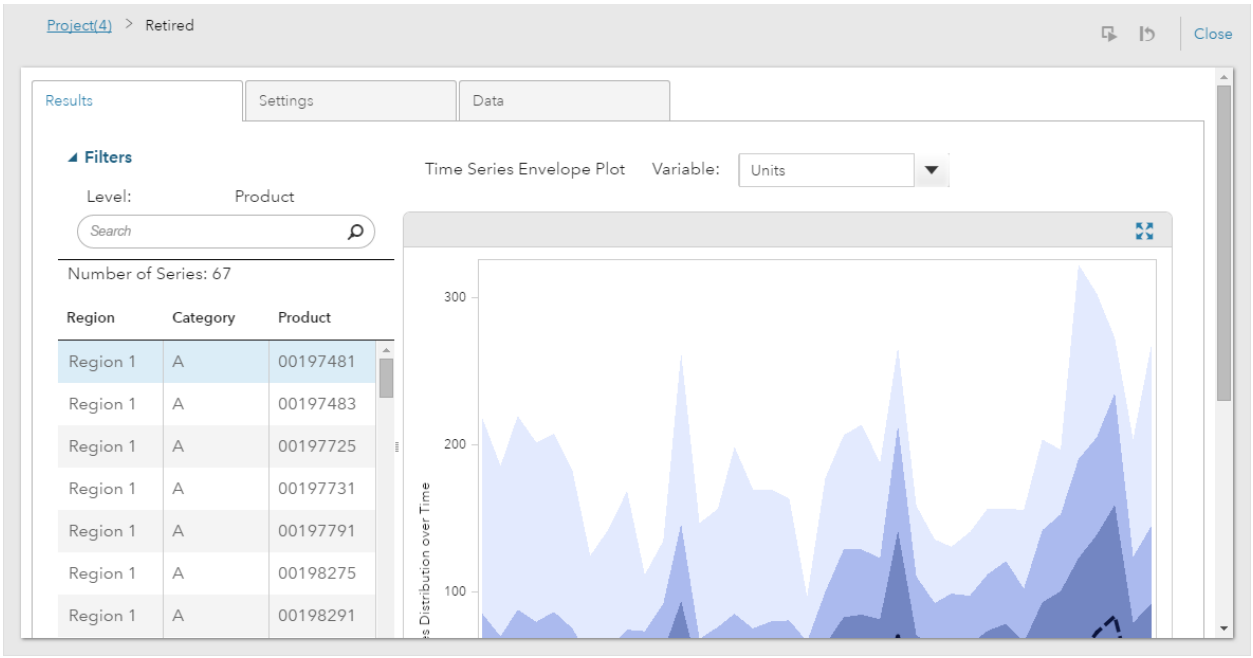

Viewing the Segmentation Rule for the Selected Node

To view the segmentation rule for the selected node, click the **Settings** tab. From the **Settings** tab, you can change the parameter values for that rule.

In this example, a retired series is defined as any series that has 24 intervals since the last nonzero observation. The value of 24 intervals comes from the default segmentation strategy used to create the project. From this page, you can change the

value to whatever you want (for example, 36). However, any changes that you make in the **Settings** tab apply only to the current project.

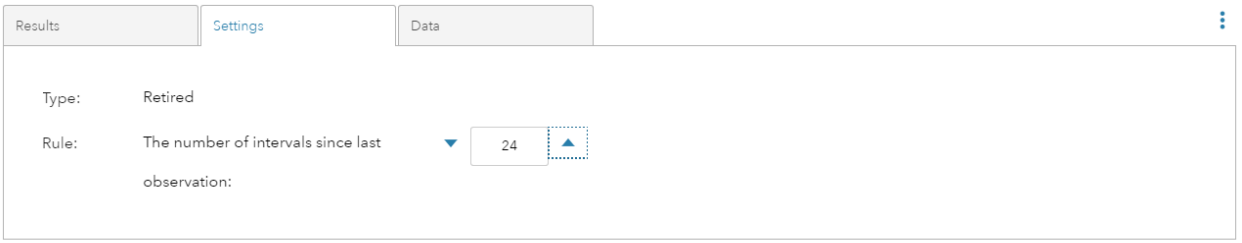

Viewing the Input Data for the Selected Series

To view the data for the series in the selected segment, click the **Data** tab. You can also choose to view the descriptive statistics (for example, start date, first observation, missing values, mean, and so on) for the segment.

14

Working with the Modeling Strategies in Your Project

Viewing the Forecasts for Each Segment in Your Project

The modeling strategies for each segment generate the forecasts for that segment. These forecasts appear in the forecast node in the project template. There is one forecast node for each segment. From this forecast node, you can view the generated forecasts for that segment. All generated forecasts are displayed in the BEST12. format. You can also view a plot of the historical versus forecast data.

In this example, the project template contains four forecast nodes: **Retired Forecasts**, **Short Forecasts**,**Intermittent Forecasts**, and **Unsegmented Forecasts**.

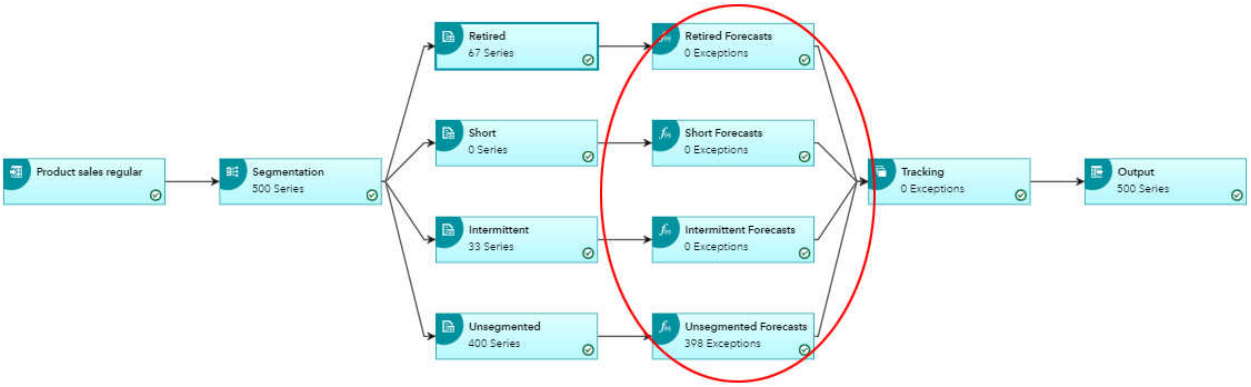

To view the results for the **Retired Forecasts** node, double-click the **Retired Forecasts** node and open the **Results** tab.

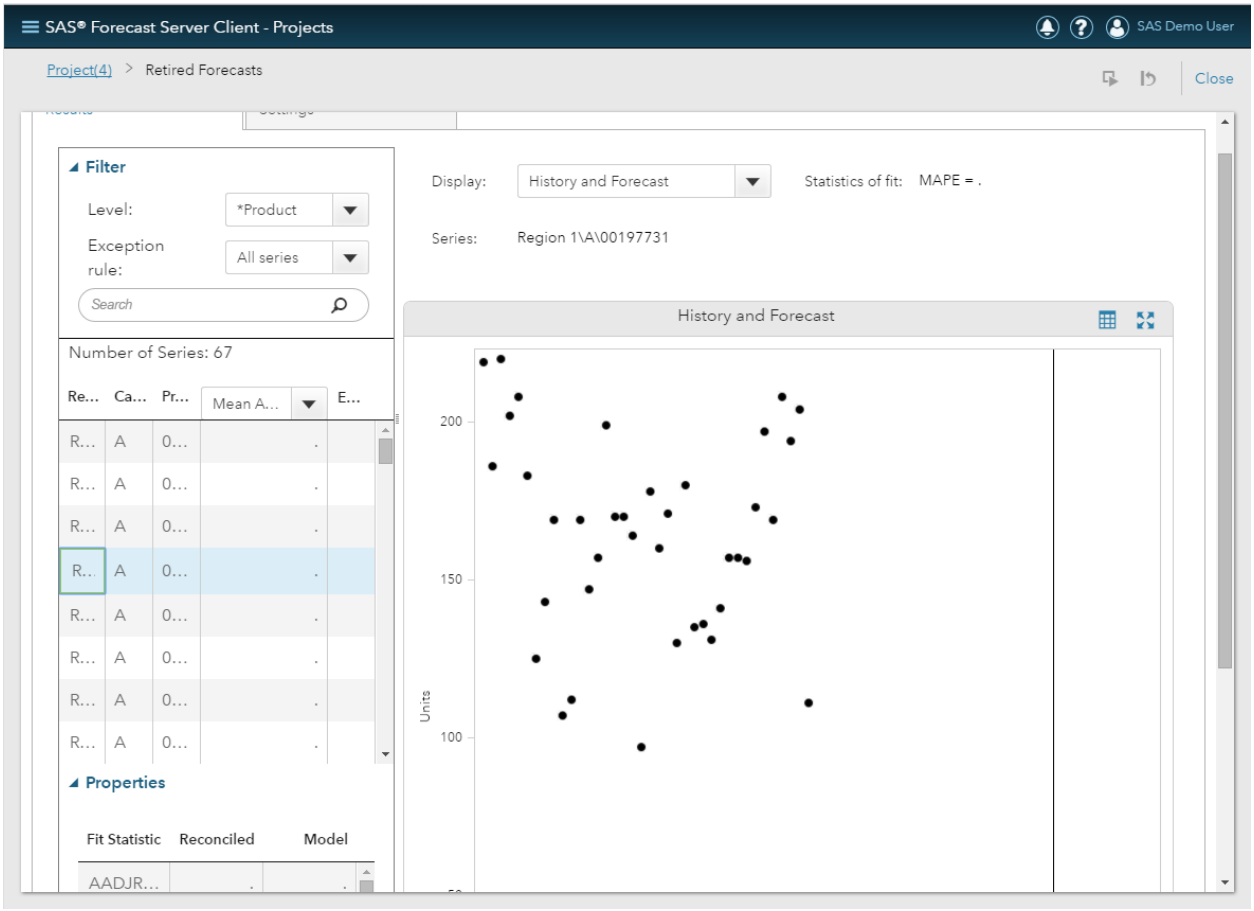

From the **Results** tab, you can view the history and forecast plots for each level in the hierarchy. By default, SAS Forecast Server Client shows the results for the first series or BY variable in the default reconciliation level. The default reconciliation level is denoted by an asterisk in the **Level** drop-down list. In this example, the default reconciliation level is Project.

Click the **Settings** tab to view parameters for the modeling strategies. The parameters that are available depend on the modeling strategy.

You can also specify the forecasting options and create exception rules for that node. By double-clicking the node in the project template, you can view these options.

For example, if you double-click the **Retired Forecasts** node, there are no options to set. All future values are set to zero.

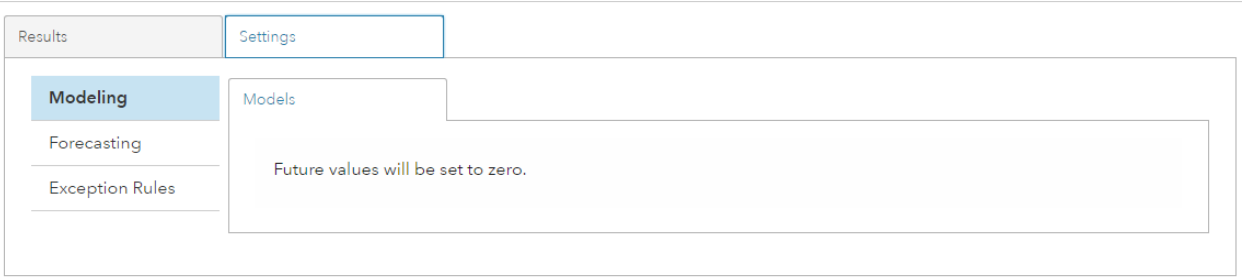

However, if you double-click the **Unsegmented Forecasts** node, you can change the diagnostic, model selection, and model generation options for that segment.

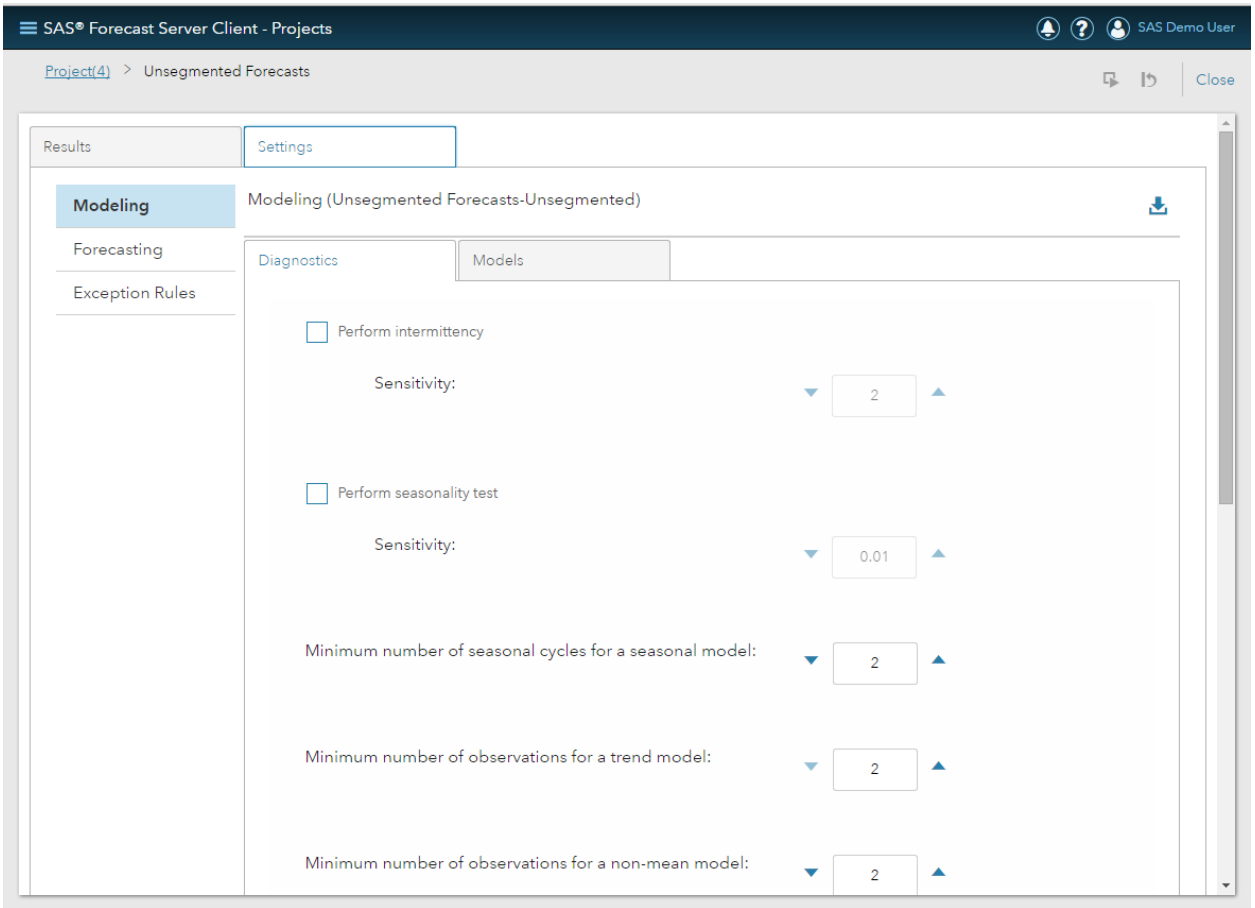

Setting the Forecasting Options

Several of the forecasting options cannot be edited from the forecast node in the project template. For example, the number of periods to forecast, the confidence limit for those forecasts, and the reconciliation level are set using the options in the data definition node. If you change any of these options, all of the nodes in the project are affected, and you must run the entire project again.

From the **Settings** tab, you can select these options for a forecast node:

Allow negative forecasts
specifies whether to allow negative forecasts. If you clear the **Allow negative forecasts** check box, any negative values in the forecast model are set to 0.

Create the component series data set
specifies whether to create the component series data set (called OUTCOMPONENT). By default, the component series data set is automatically created when you create a project. A component data set is created for each BY variable in the project. Because a component data set can also contain several other variables, these data sets can be quite large. As a result, a component data set might require additional disk space and computing time.

If you clear this option, you suppress the creation of this data set.

Create a forecast data set for independent variables
specifies whether to create a forecast data set for independent variables. In some situations, you might not know the future values of the independent variables. Depending on the model for your project, the SAS Forecast Server client might use these independent variables to create forecasts for your dependent variables. Select this option to create an output data set that contains the future values of the independent variables. You can review this data set to verify whether those values are reasonable.

Import a Modeling Strategy

To import a modeling strategy to your project:

1. In the project template, double-click the forecast node where you want to change the modeling strategy.
2. Click the **Settings** tab, and in the selection pane, click **Modeling**.
3. To import the modeling strategy, click . The Import Modeling Strategy window appears.
4. For the **Type** option, select whether you want to import a rule-based modeling strategy or a code-based modeling strategy.
5. Select the modeling strategy that you want to import. If you are importing a rule-based modeling strategy, you can import multiple rules.
6. Click **OK**.

Applying Exception Rules to Your Modeling Strategies

Why Use an Exception Rule?

Exception rules enable you to quickly identify any series that meet or do not meet a specific criteria or a specific criterion. For example, you want to see all forecasts where the mean absolute percent error (MAPE) is greater than 5. When an exception rule is applied to a modeling strategy, the project template displays in the forecast node the number of series that meet that criterion.

Create an Exception Rule

To create an exception rule:

1. Double-click the forecast node in the project template.
2. Click the **Settings** tab. From the selection pane, click **Exception Rules**.
3. To add an exception rule, click +. The Add Modeling Exception Rule window appears.
4. Specify a name and description for this new rule.
5. From the **Category** drop-down list, specify whether you want to create this rule using a model statistic of fit or a reconciled statistic of fit.
6. In the table, specify the criterion for this exception rule.

For example, the settings in the Add Modeling Exception Rule window define an exception rule. In this exception rule, the mean absolute percent error for the reconciled statistic of fit is greater than 5.

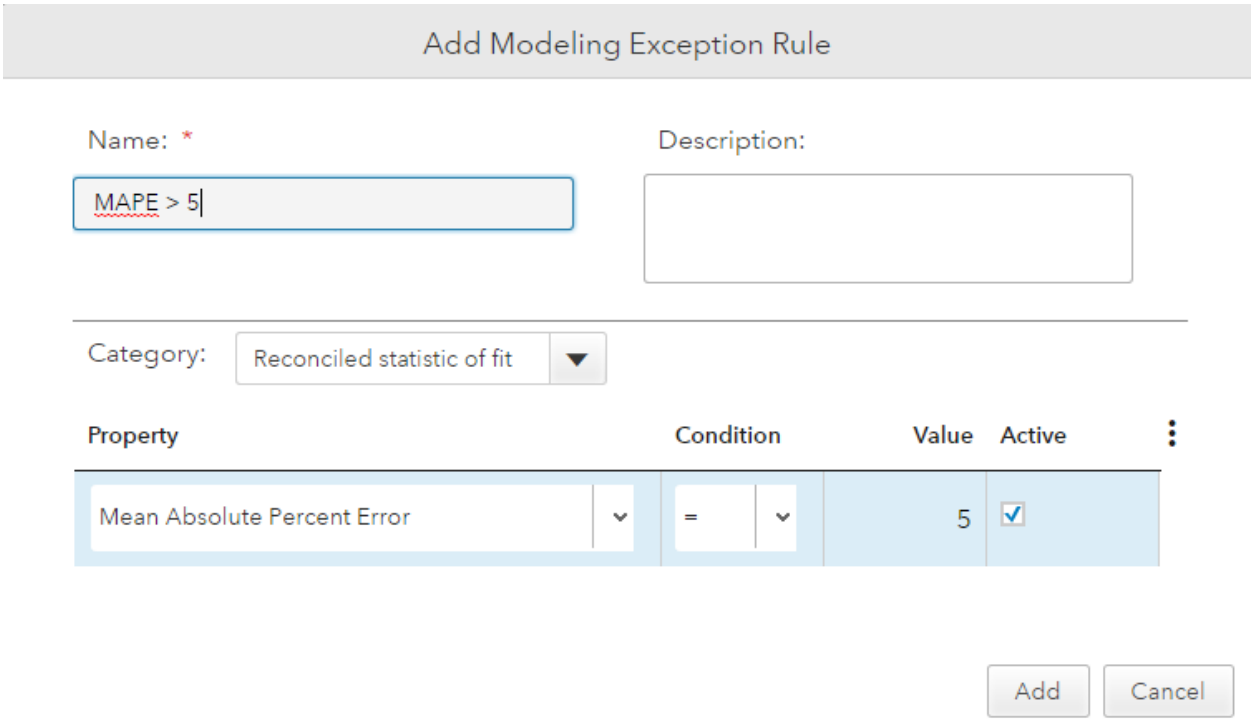

7 Click **Add**. The new exception rule appears in the list of exception rules for that forecast node.

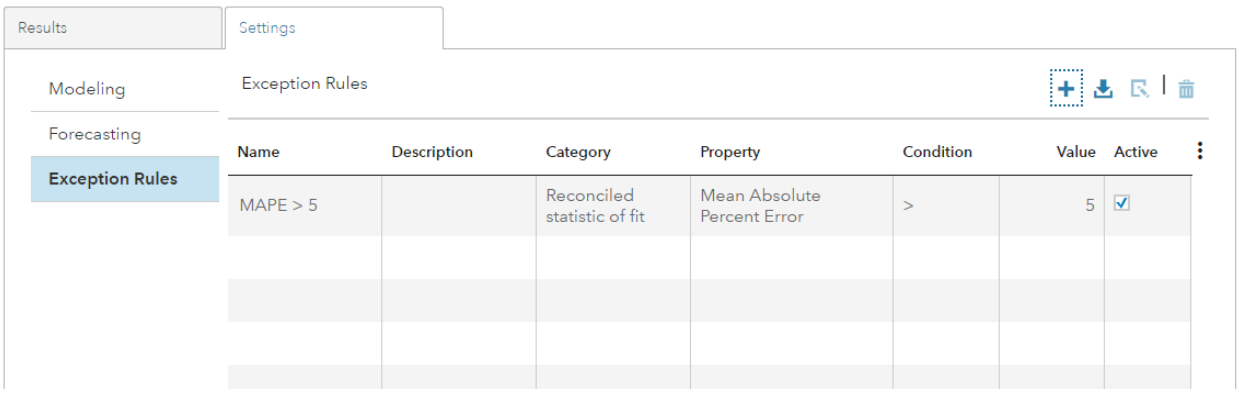

8 For SAS Forecast Server Client to apply this exception rule to the forecast node, select the **Active** check box.

Note: If multiple exception rules are active, SAS Forecast Server Client combines these rules using AND as the logical operator.

Import an Exception Rule

Instead of creating an exception rule for each forecast node, you might want to apply the same rule to several nodes. You can create an exception strategy and save it to the metadata server. (For more information, see "Create an Exception Strategy" on page 79.) Then you can import this exception rule to the appropriate forecast nodes, so SAS Forecast Server Client uses it when generating your forecasts.

To import an exception rule:

1 Double-click the forecast node in the project template.

2 Click the **Settings** tab. From the selection pane, click **Exception Rules**.

3 Click . The Import Exception Rule Strategies window appears.

4 Select the exception rule strategies that you want to apply to the current forecast node. To select multiple strategies, press Ctrl and select the check box for that strategy.

TIP Click to view the names of the exception rules in a strategy. When importing multiple strategies, you cannot have exception rules with the same name because these exception rules might have conflicting rules. SAS Forecast Server Client displays an error when you try to import exception rules with the same name.

5 To replace the existing exception rules for this forecast node, select **Replace existing exception rules**.

6 Click **OK**. The new exception rules appear in the table for the current forecast node.

7 Use the check boxes in the **Active** column to specify which exception rules SAS Forecast Server Client should apply to this forecast node. Multiple exception rules are joined by the AND logical operator.

8 Click **Close** to save your changes.

Part 5

Analyzing Your Results

15

Tracking Your Forecasts

What Is Forecast Tracking?

Forecast tracking evaluates forecast accuracy by comparing the actual values of a time series with the forecasts of those values. Tracking is the best way to analyze the accuracy of your forecasts. In SAS Forecast Server Client, tracking is an iterative process. The time series is updated periodically to include the most recent data, and new forecasts are computed. These forecasts extend further into the future for each iteration.

When more recent data is available, you can update the input data for your project from the **Settings** tab of the data definition node. A message appears if the input data has

been modified since the last time you ran the project. After you update the input data, you must run the project again to incorporate the newly acquired data.

Note: When you update the input data for the project, any retired series will show missing values. These missing values are converted to zeros because there is no new data for the retired product. For the retired series, SAS Forecast Server Client generates forecasts of 0. As long as the retired series remains in the input data, SAS Forecast Server Client generates forecasts and tracking information for these series. When reviewing the tracking reports, remember that series might behave differently in different segments. For example, series in the retired segments are treated differently in that missing values are automatically set to 0. If you do not want to include these retired series in your analysis, you must remove them from the input data set.

When the project runs, the **Tracking** node determines whether any of the new actual values overlap the dates for previous forecasts. If the dates overlap, the new forecasts are saved as a new iteration. As this process repeats through the lifecycle of the project, you accumulate forecasting iterations. More detail is available for computing the statistics on forecast error, which is the difference between the actual values and the forecasts. If your project contains *n* iterations of forecasts, you have tracking statistics for lead times 1 through *n-1*. These statistics help you determine whether the accuracy of the forecasts decreases as you forecast further into the future.

When you have multiple iterations, series might be split across segments. Each time you run a project, the series are allocated to segments based on the segmentation rules. When you add new information to the input data, some of the series could end up in different segments than they were before. An example would be adding another month of data to a data set that contains monthly time series data. Forecast tracking combines all the segments. When you review the contents of the **Tracking** node, the list of BY groups shows all of the BY groups and the segment for each one. Some BY groups might have more than one segment. For example, you might see the following:

```
Region 1  A  00197731 Unsegmented
Region 1  A 00197731  Retired
```

At the top of the list, the value for the **Number of series** is the number of rows in the table. When there are duplicate rows in the table, this number is larger than the number of series shown in the segment and forecast nodes in the project template.

The Tracking Node in the Project Template

To view the tracking report for your project, you must double-click the **Tracking** node in the project template.

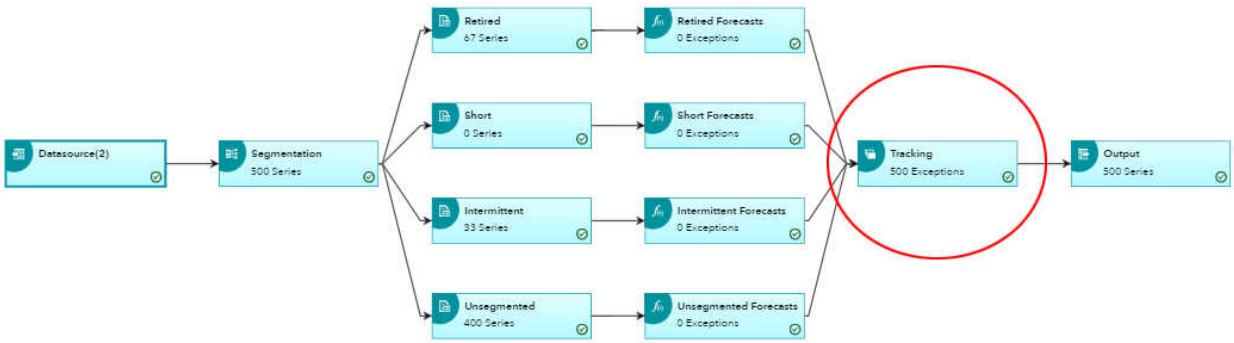

When you double-click the **Tracking** node in the project template, the **Results** tab shows the series for a selected segment. By default, the series for all segments are displayed.

To view the reports for a specific series, select the series in the list. A series might appear under different segments from one iteration to the next. Series can also vary in the number of iterations that are available. For example, some series might acquire more recent data every time forecasts are produced. Other series might not have any new data.

Understanding the Tracking Reports

About the Tracking Reports

Remember as you start generating your forecasts in SAS Forecast Server Client that these reports might not contain much information. These reports become more useful with each iteration because you are extending the range of the data and reforecasting.

The tracking reports show the tracking results for the series that you selected in the table. You can select the report to display from the **Display** drop-down menu.

Forecast Accuracy by Iteration

This report is available as a table or a graph. The table shows all the tracking statistics. The graph displays the tracking statistic in the **Plots** drop-down list for each available iteration.

Here is an example graph:

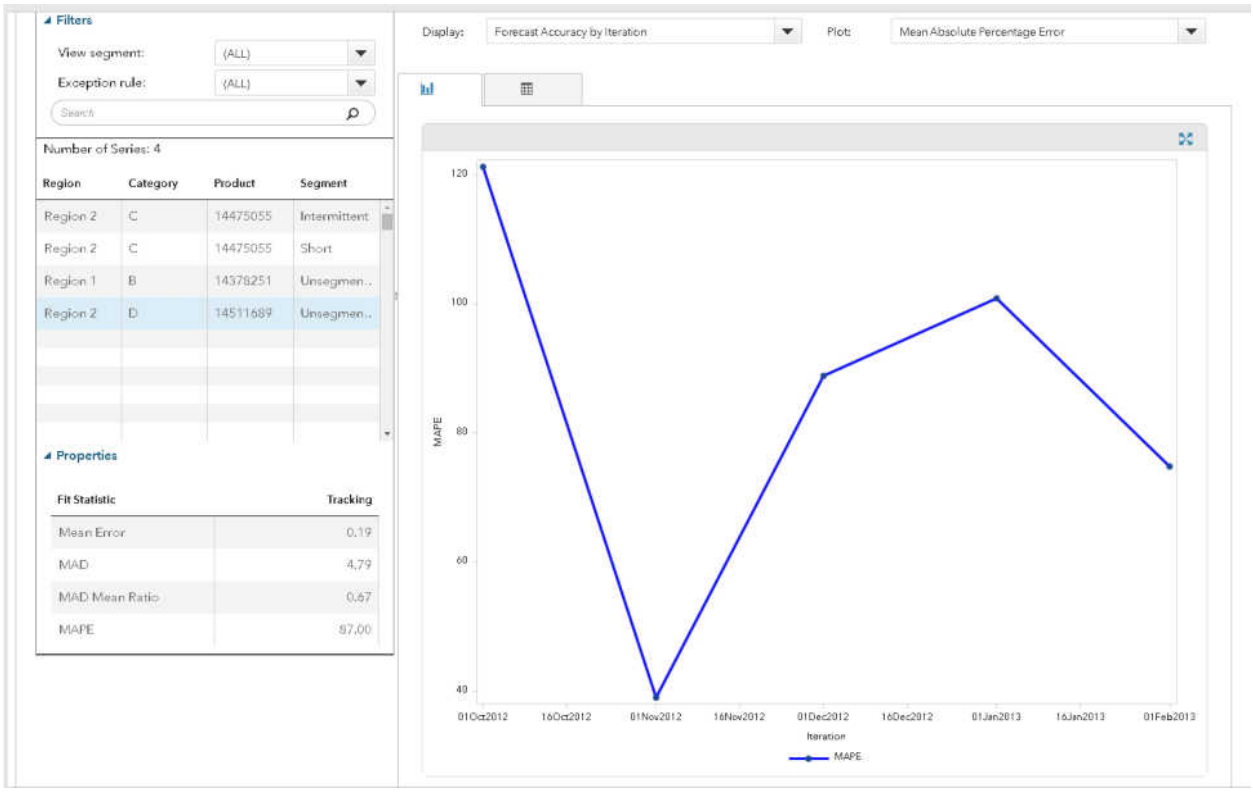

Here is an example of a tabular report:

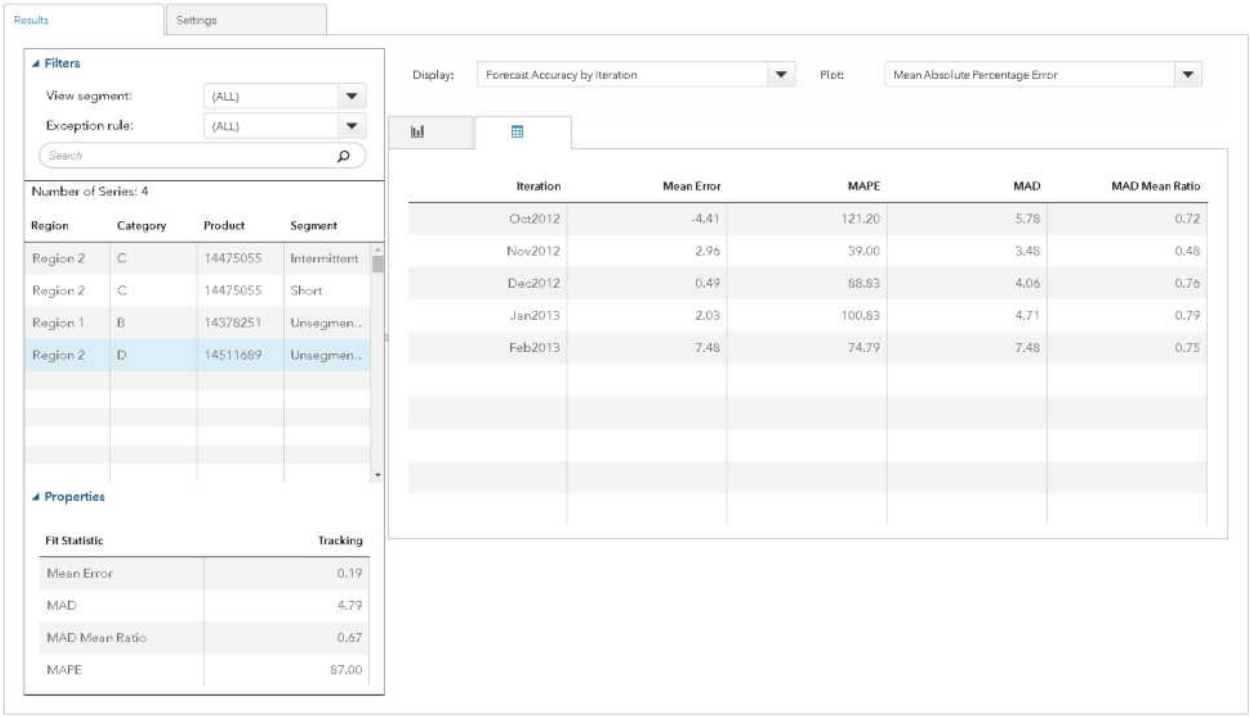

Forecast Accuracy by Lead Time

This report displays the tracking statistic (which you selected from the **Plots** drop-down list) for each lead time. Lead time is the number of periods into the future that a forecast was made. If only two iterations are available, only one comparison is available, and in this case, the lead time is one. When three iterations are available, comparisons are available for lead-time one and lead-time two. As more iterations become available, this report shows you how far into the future you can forecast before forecasting accuracy begins to decline.

Here is an example of a graphical report:

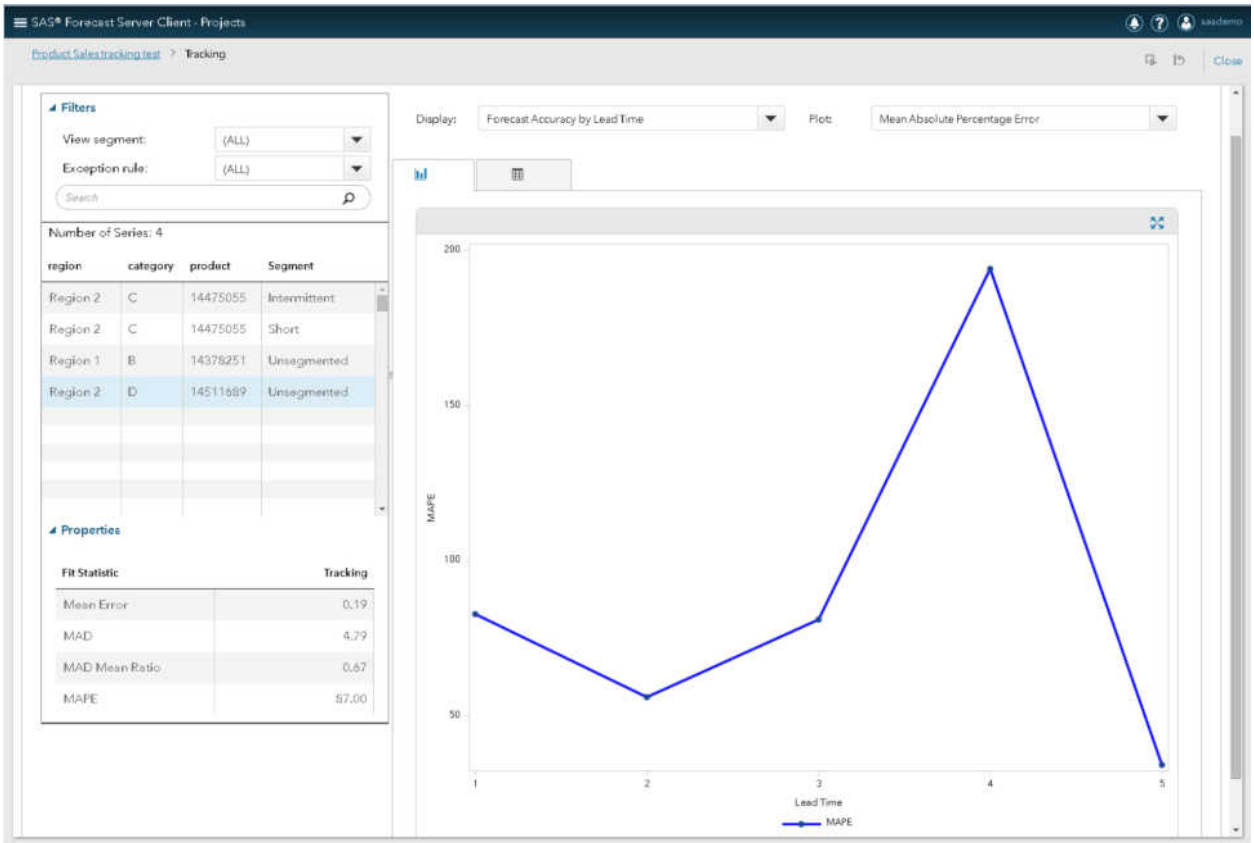

Here is an example of a tabular report:

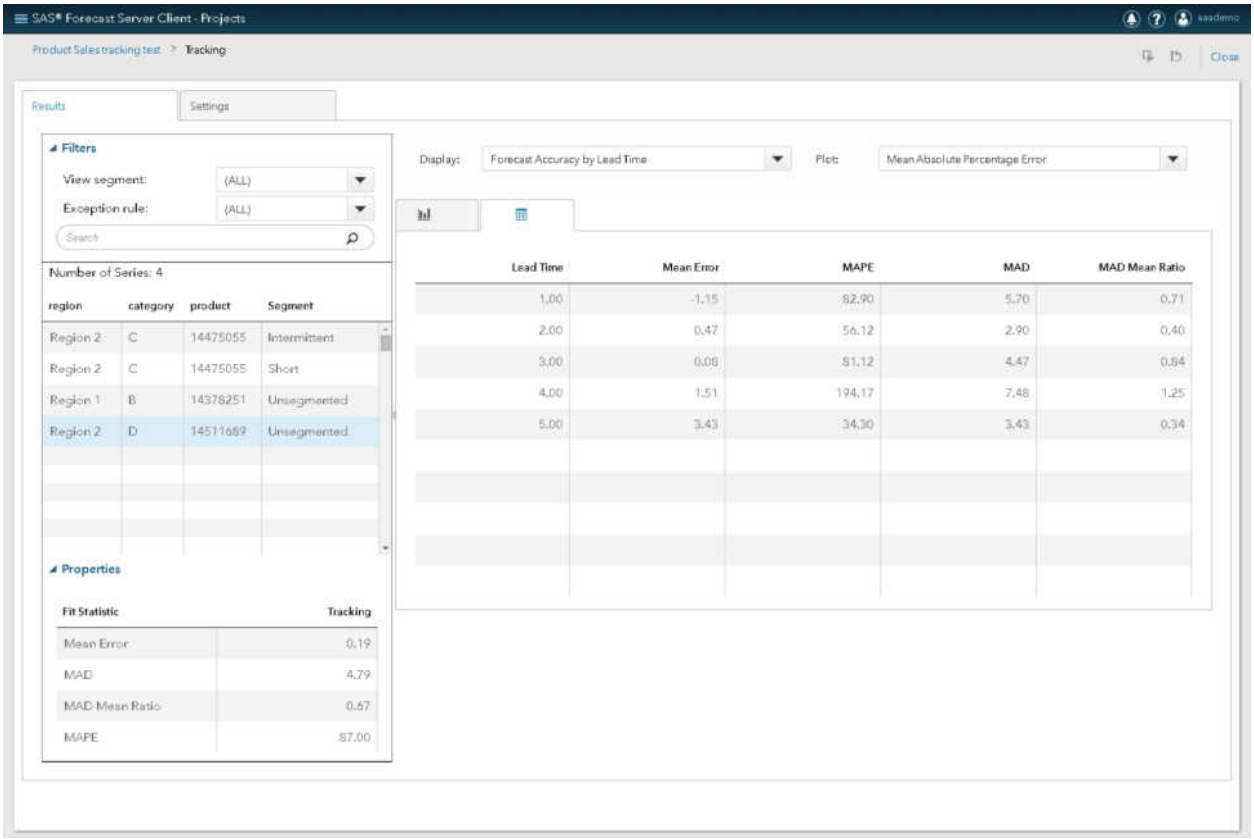

Iteration Forecast

The iteration forecast graph is available for any selected iteration. However, for the most recent iteration, the report provides no new information. For prior iterations, this report displays the forecast graph that was current when the associated forecasts were made. Now that this iteration is in the past, the SAS Forecast Server Client adds the actual values to the forecast graph. As a result, this graph enables you to access the quality of past forecasts.

Here is an example of an iteration forecast graph for an earlier iteration. In the forecast horizon, the graph compares the forecasts with the actual values.

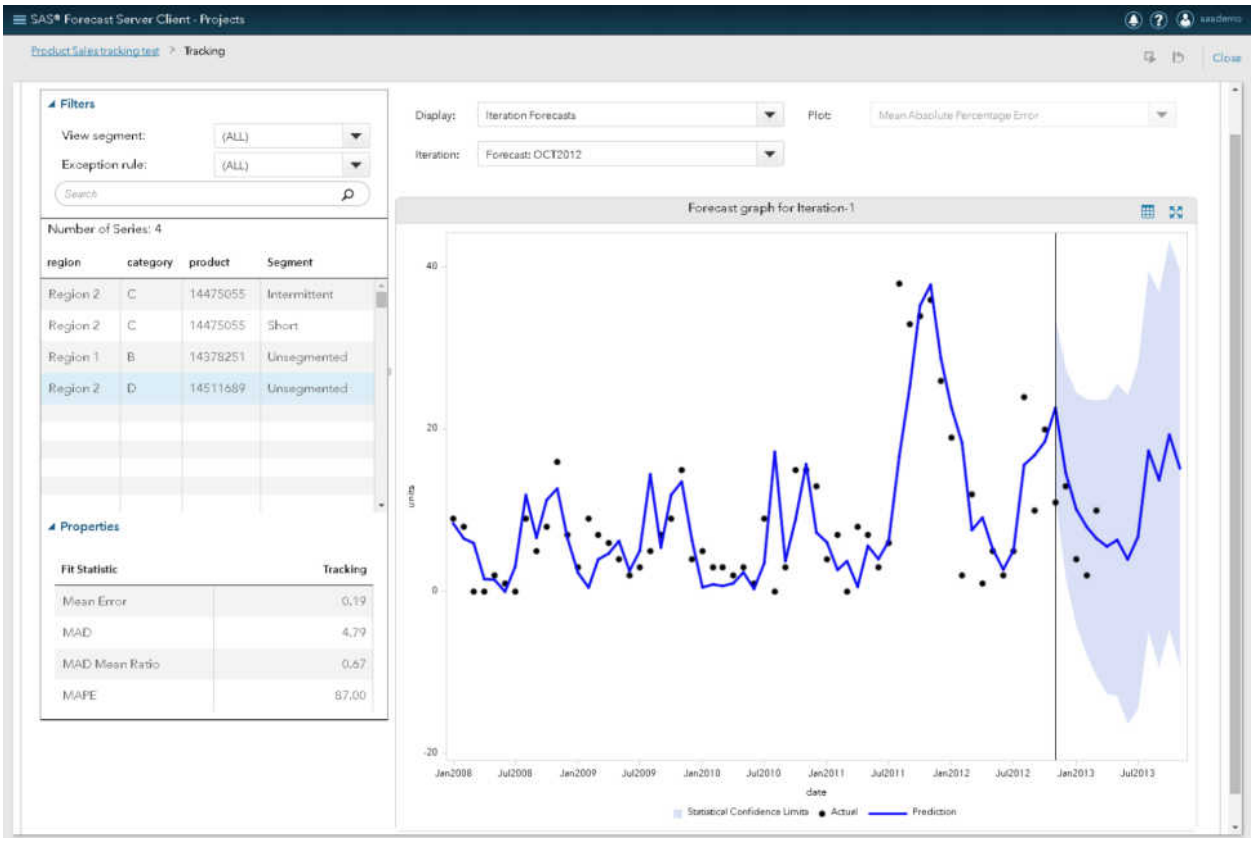

Actual Values and Forecasts

This report can be displayed as either a table or a graph. The report shows the forecasts for all available iterations and the latest actual values. The table is sometimes referred to as a waterfall table because the actual values are displayed in the first row, and successive iterations are in subsequent rows. The row labels indicate the ending date for the actual data in each iteration, and the column labels indicate the date being forecast. Forecasts that are made at the same lead time form diagonals from upper left to lower right.

The graph enables you to quickly see any discrepancies between actual values and the forecasts for an iteration. The areas where the forecast plot overlaps the plot of the actual values shows when the forecast values are the most accurate.

Here is an example of a graph report:

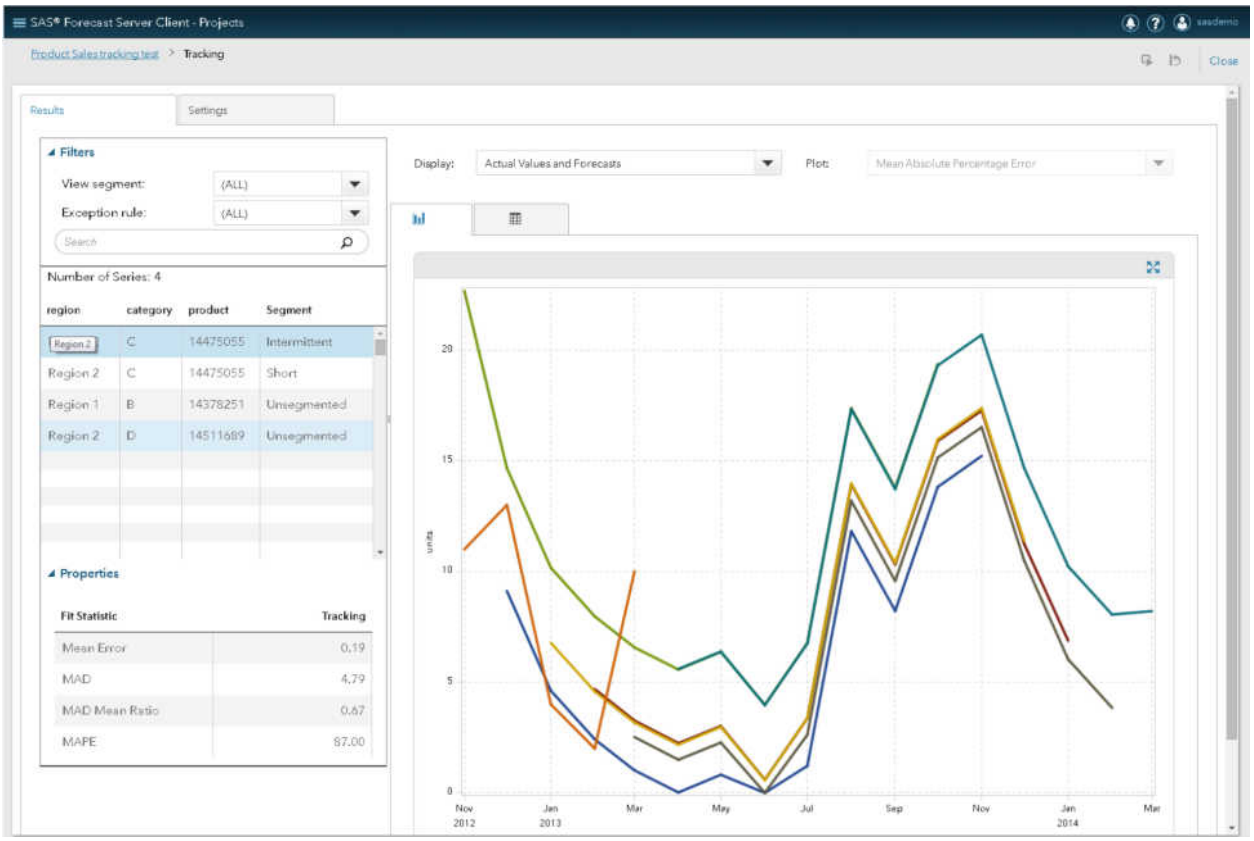

Here is an example of a tabular report:

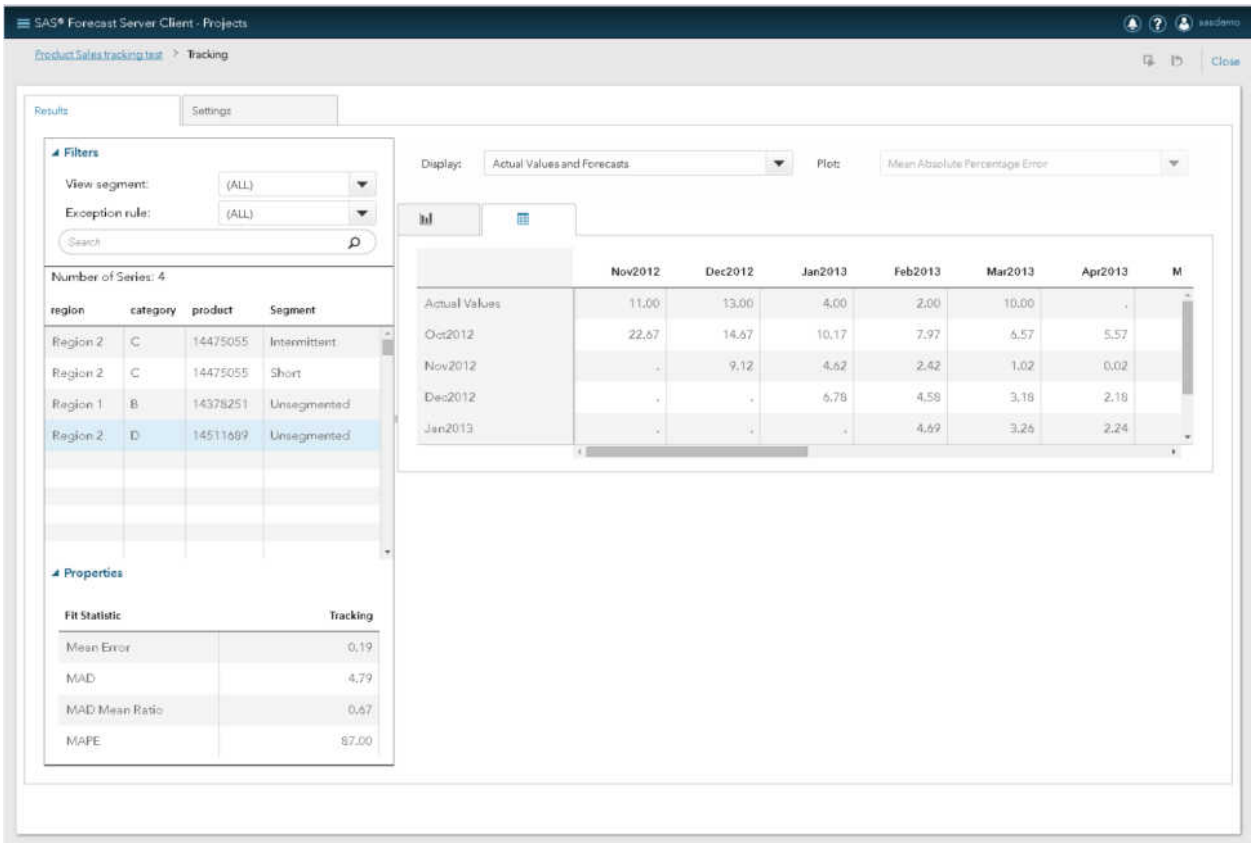

Adding Exception Rules for Tracking

Using SAS Forecast Server Client, you can perform forecast tracking by exception. In the **Tracking** node, you can add rules that define the levels of forecast accuracy that you consider unacceptable. For example, you want to identify any forecasts where the MAPE is greater than 5. Any series that has an overall tracking MAPE greater than 5 is flagged as an exception.

Here are the statistics that you can use to define the tracking exception rules

- MAPE – mean absolute percent error. This statistic works well if none of the actual values are 0.
- MAD – mean absolute deviation. This statistic is well suited for intermittent demand data.
- MMR – MAD/mean ratio. Dividing the mean absolute deviation by the sum of the actual enables you to compare across series of different scales.
- MERROR – mean error.

These statistics are calculated from a different set of data than the statistics in the forecast nodes. For example, the MAPE statistic is available in the forecast node and in the tracking node. In the forecast node, MAPE is calculated from in-sample one-step-ahead forecasts. In the tracking node, MAPE is the comparison of the forecast values with the actual values.

16

Viewing the Output Data

Overview of the Output Data Sets

By default, SAS Forecast Server Client creates the OUTCOMPONENT data set. You can also create the forecast data set for the independent variable (called OUTINDEP= data set). All output data sets are saved to the metadata server.

The OUTINDEP= data set contains the forecast values for the independent variables. These future values of the independent variables are supplied by the HPFENGINE procedure. For more information about the HPFENGINE procedure, see *SAS Forecast Server Procedures: User's Guide*.

The OUTINDEP= data set contains the variables that you assign as BY variables and the following variables:

- _NAME_ - variable name
- _TIMEID_ - time ID values
- _X_ – values of the input variable _NAME_

The Output Data Node

The last node in the project template contains the options for creating an output data set.

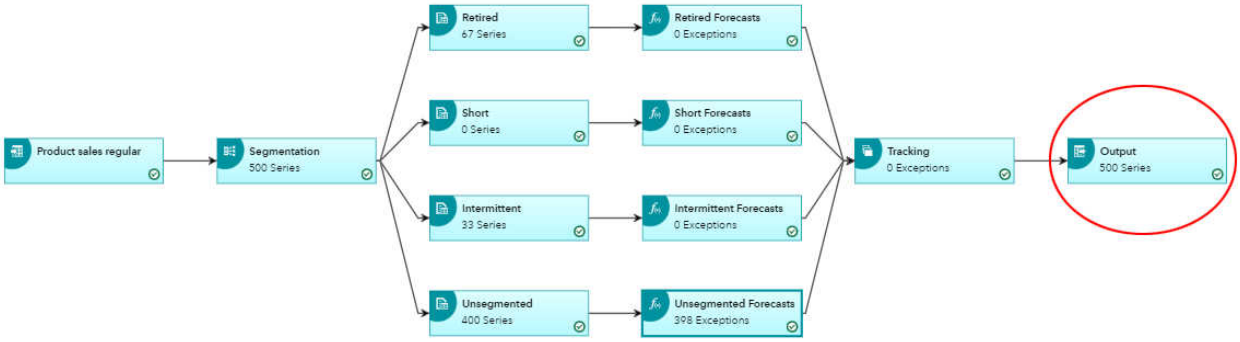

Viewing the Results for Each Segment

From the **Output** node, you can view the results for all combined segments or for an individual segment.

To view the results, double-click the **Output** node in the project template and open the **Results** tab. From the **Segment** drop-down list, select the segment that contains the series that you want to view. By default, all the series in the lowest level of the hierarchy are displayed.

To view the results for a specific BY variable, enter the value of the BY variable in the search box. For example, if you enter `Region 1` in the search box, the results show only the series in Region 1.

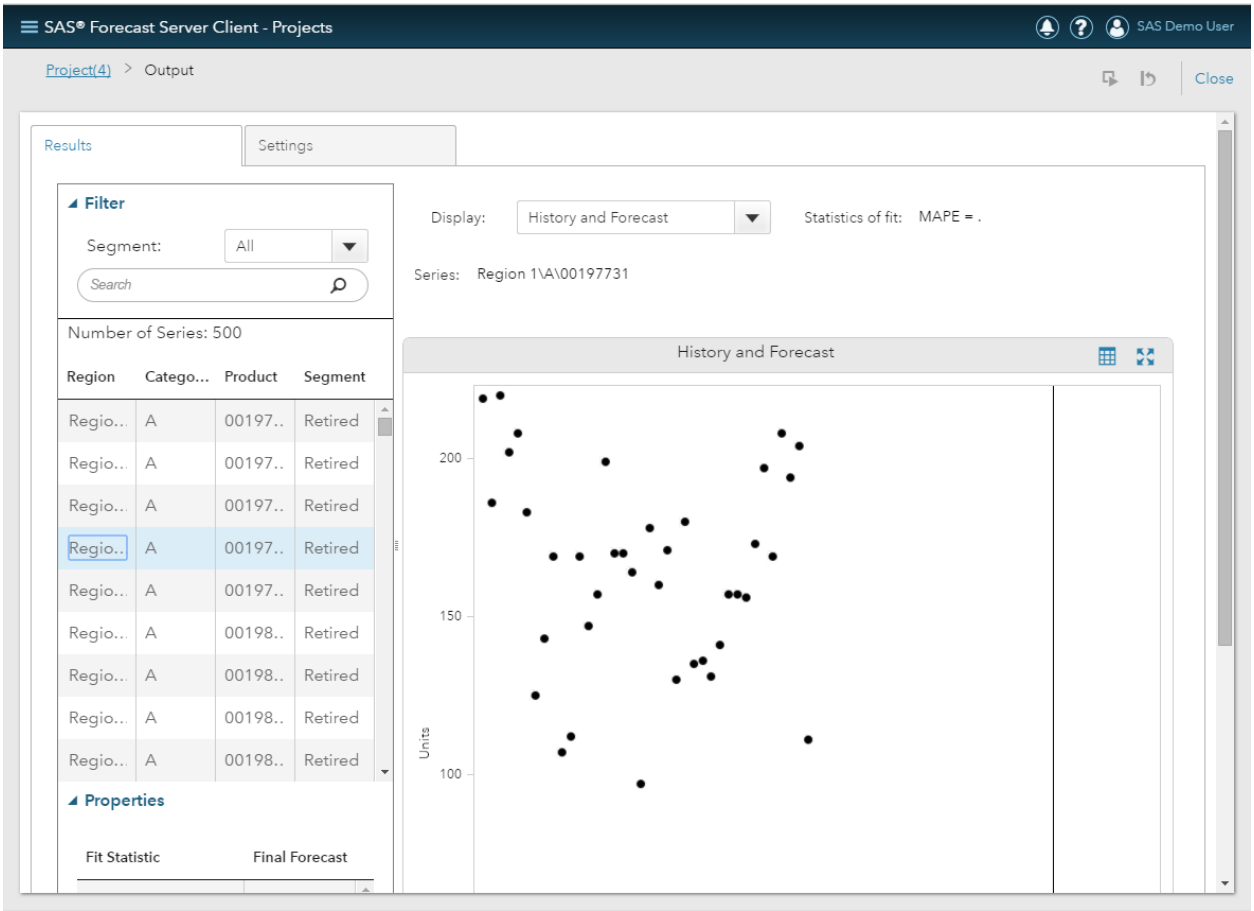

Export the Output Data

You can export the output data to a location on your metadata server.

To export the output data:

1 Double-click the **Output** node in the project template. Click the **Settings** tab.

2 Specify a valid SAS name for the output data set.

3 Select the location for the output data set. By default, this data set is exported to the FSWB_OUT library.

4 Select the contents of the output data set. By default, all these items are included. Use the check boxes to deselect any items that you do not want to include in the output.

- data table columns
- segments
- dependent variable
- actual values
- forecasted values
- lower confidence limits
- upper confidence limits
- prediction errors
- prediction standard errors
- reconciliation status

  The reconciliation status can be any of these values:

0	Success
400	An unlocked equality constraint has been imposed.
500	A locked equality constraint has been imposed.
600	A lower bound is active.
700	An upper bound is active.
1000	ID value out of range with respect to the START= and END= intervals.
2000	Insufficient data to reconcile.
3000	Reconciliation failed for the predicted value. This implies that it also failed for the confidence limits and standard error.
4000	Reconciliation failed for the standard error.
5000	Reconciliation failed for the confidence limits
6000	The constrained optimization problem is infeasible.

7000 The option DISAGGREGATION=PROPORTION has been changed to DISAGGREGATION=DIFFERENCE for this observation due to discordant signs in the input.

8000 The STDMETHOD= option that was provided by the user has been changed for this observation.

9000 The CLMETHOD= option that was provided by the user has been changed for this observation.

10000 The standard error hit the limits that are imposed by the STDDIFBD= option.

11000 Multiple warnings have been printed to the log for this observation.

12000 The number of missing values in the STD variable in the DISAGGDATA= data set is different from the number of missing values in the union of the PREDICT and ACTUAL variables.

13000 The solution might be suboptimal. This means that the optimizer did not find an optimal solution, but the solution that was provided satisfies all constraints.

14000 A failed forecast ".F" has been detected in a relevant input variable.

Part 6

Appendixes

Appendix 1

Statistics of Fit

What Are Statistics of Fit?

The statistics of fit are statistical values that are used to evaluate how well a forecasting model performs by comparing the actual data to the predictions. For a given forecast model that has been fitted to the time series data, the model should be checked or evaluated to see how well it fits or forecasts the data. Commonly used statistics of fit are Root Mean Square Error (RMSE), Mean Absolute Percent Error (MAPE), Akaike Information Criteria (AIC), and many others.

When the full range of data is used to both fit and evaluate the model, this is referred to as in-sample evaluation. When the most recent data is excluded for parameter estimation (holdout) and this holdout sample is used for evaluation, this is referred to as holdout sample evaluation. Holdout sample analysis is similar to the training and testing of neural networks. A portion of the data is withheld from training (fit), and the withheld data (holdout) is used to test performance.

When a particular statistic of fit is used for forecast model selection, it is referred to as the model selection criterion. For example, if the MAPE (an often recommended choice) is used as a model selection criterion, the forecast model with the smallest MAPE in the evaluation region (in-sample or holdout-sample) is chosen as the best model.

When a particular statistic of fit is used to judge how well the forecasting process is predicting the future, it is referred to as the performance statistic.

Understanding the Formulas and Definitions

To understand the formulas and definitions for the statistics of fit, review the following information:

- *n* is the number of nonmissing observations.
- *k* is the number of fitted parameters in the model.
- $APE = |100 * (y_t - \hat{y}_t)/y_t|$ is the absolute percent error.
- $ASPE = |100 * (y_t - \hat{y}_t)/0.5(y_t + \hat{y}_t)|$ is the absolute symmetric percent error.
- $APPE = |100 * (y_t - \hat{y}_t)/\hat{y}_t|$ is the absolute predictive percent error.
- $RAE = |(y_t - \hat{y}_t)/(y_t - y_{t-1})|$ is the relative absolute error.

Descriptions of Statistics of Fit

The definitions and formulas for the statistics of fit that are available in the SAS Forecast Server Client are described below. You can use statistics of fit to measure how

well different models fit the data. The statistics of fit for the various forecasting models can be printed or stored in a data set.

Adjusted R-squared (ADJRSQ)

The adjusted R^2 statistic, $1 - (\frac{n-1}{n-k})(1 - R^2)$.

Akaike Information Criterion (AIC)

Akaike's information criterion, $n \quad \ln(\frac{SSE}{n}) + 2k$.

Akaike Information Criterion, finite sample size corrected (AICC)

Akaike's information criterion with an empirical correction for small sample sizes, $AIC + (\frac{2k(k+1)}{n-k-1})$.

Amemiya's Adjusted R-squared (AADJRSQ)

Amemiya's adjusted R^2, $1 - (\frac{n+k}{n-k})(1 - R^2)$.

Amemiya's Prediction Criterion (APC)

Amemiya's prediction criterion, $\frac{1}{n}SST(\frac{n+k}{n-k})(1 - R^2) = (\frac{n+k}{n-k})\frac{1}{n}SSE$.

Geometric Mean Absolute Error Percent of Standard Deviation (GMAPES)

The geometric mean of the absolute error as a percentage of the standard deviation.

Geometric Mean Percent Error (GMAPE)

The geometric mean percent error.

Geometric Mean Predictive Percent Error (GMAPPE)

The geometric mean absolute predictive percent prediction error.

Geometric Mean Relative Absolute Error (GMRAE)

The geometric mean of the relative absolute errors.

Geometric Mean Symmetric Percent Error (GMASPE)

The geometric mean of the absolute symmetric percent errors.

Maximum Absolute Error Percent of Standard Deviation (MAXAPES)

The maximum of the absolute error as a percentage of the standard deviation.

Maximum Error (MAXERR)

The largest prediction error.

Maximum Percent Error (MAXPE)

The largest percent prediction error, $100\max(\frac{(y_t - \hat{y}_t)}{y_t})$. The summation ignores observations where $y_t = 0$.

Maximum Predictive Percent Error (MAXPPE)

The maximum of the predictive percent errors.

Maximum Relative Error (MAXRE)

The maximum of the relative errors.

Maximum Symmetric Percent Error (MAXSPE)

The maximum of the symmetric percent errors.

Mean Absolute Error (MAE)

The mean absolute prediction error, $\frac{1}{n}\Sigma_{t=1}^{n}|y_t - \hat{y}_t|$.

Mean Absolute Error Percent of Standard Deviation (MAPES)

The mean of the absolute error as a percentage of the standard deviation.

Mean Absolute Percent Error (MAPE)

The mean of the absolute percent errors.

Mean Absolute Predictive Symmetric Percent Error (MAPPE)

The mean of the absolute symmetric predictive percent error.

Mean Absolute Scaled Error (MASE)

The mean of the absolute scaled errors.

Mean Absolute Symmetric Percent Error (SMAPE)

The symmetric mean of the absolute percent error.

Mean Error (ME)

The mean prediction error, $\frac{1}{n}\Sigma_{t=1}^{n}(y_t - \hat{y}_t)$

Mean Percent Error (MPE)

The mean percent prediction error, $\frac{1}{n}\Sigma_{t=1}^{n}\frac{(y_t - \hat{y}_t)}{y_t}$. The summation ignores observations where $y_t = 0$.

Mean Predictive Percent Error (MPPE)

The mean of the predictive percent error.

Mean Relative Absolute Error (MRAE)

The mean of the relative absolute errors.

Mean Relative Error (MRE)

The mean of the relative errors.

Mean Square Error (MSE)

The mean squared prediction error calculated from the one-step-ahead forecasts, $MSE = \frac{1}{n}SSE$. This formula enables you to evaluate small holdout samples.

Mean Symmetric Percent Error (MSPE)

The mean of the symmetric percent errors.

Median Absolute Error Percent of Standard Deviation (MDAPES)

The median of the absolute error as a percentage of the standard deviation.

Median Absolute Percent Error (MDAPE)

The median of the percent errors.

Median Absolute Predictive Percent Error (MDAPPE)

The median of the predictive percent errors.

Median Relative Absolute Error (MDRAE)

The median of the relative absolute errors.

Median Absolute Symmetric Percent Error (MDASPE)

The median of the symmetric percent errors.

Minimum Absolute Error Percent of Standard Deviation (MINAPES)

The minimum of the absolute error as a percentage of the standard deviation.

Minimum Error (MINERR)

The smallest prediction error.

Minimum Percent Error (MINPE)

The smallest percent prediction error, $100\min(\frac{(y_t - \hat{y}_t)}{y_t})$. The summation ignores observations where $y_t = 0$.

Minimum Predictive Percent Error (MINPPE)

The smallest predictive percent error.

Minimum Relative Error (MINRE)

The smallest relative error.

Minimum Symmetric Percent Error (MINSPE)

The smallest symmetric percent error.

R-square (RSQUARE)

The R^2 statistic, $R^2 = 1 - \frac{SSE}{SST}$. If the model fits the series badly, the model sum of squared error, *SSE*, might be larger than *SST* and the R^2 statistic will be negative.

Random Walk R-square (RWRSQ)

The random walk R^2 statistic (Harvey's R^2 statistic using the random walk model for comparison), $1 - (\frac{n-1}{n})\frac{SSE}{RWSSE}$, where $RWSSE = \Sigma_{t=2}^{n}(y_t - y_{t-1} - \mu)^2$, and $\mu = \frac{1}{n-1}\Sigma_{t=2}^{n}(y_t - y_{t-1})$.

Root Mean Square Error (RMSE)

The root mean square error, $\sqrt{MSE}$.

Schwarz Bayesian Information Criterion (SBC)

Schwarz Bayesian information criterion, $n * \ln(\frac{SSE}{n}) + k * \ln(n)$.

Sum of Squares Error (SSE)

The sum of the squared prediction errors. $SSE = \Sigma_{t=1}^{n}(y_t - \hat{y}_t)^2$, where $\hat{y}$ is the one-step predicted value.

Total Corrected Sum of Squares for the dependent variable
: The total sum of squares for the series corrected for the mean: $\Sigma_{t=1}^{n}(y_t - \bar{y})^2$, where $\bar{y}$ is the series mean.

Total Sum of Squares (SST)
: The total sum of squares for the series, uncorrected for the mean: $\Sigma_{t=1}^{n} y_t^2$.

Unbiased Mean Square Error (UMSE)
: The unbiased mean square error.

Unbiased Root Mean Square Error (URMSE)
: The unbiased root mean square error.

Additional Information

For more information about statistics of fit, see the "Forecasting Process Details" chapter in the *SAS Forecast Server Procedures: User's Guide*.

Descriptive Statistics

You create exception rules by using these descriptive statistics.

First observation
: The time period of the first observation in the series.

Last observation
: The time period of the last observation in the series.

Number of values
: The number of values in the series.

Number of missing values
: The number of series with missing values.

Number of nonmissing values
: The number of series with nonmissing values.

Sum of values
: The sum of the values in the series.

Mean value
: The mean of the values in the series.

Standard deviation
: The standard deviation for the values in the series.

Minimum value
: The minimum value in the series.

Maximum value
: The maximum value in the series.

Appendix 2

Reserved Variable Names

If any of the variables in your input data set begin with an underscore (_) or use these keywords, those variables are not available in SAS Forecast Server Client. You must rename the variables in your input data set in order to use them in SAS Forecast Server Client. This list of reserved keywords is not case sensitive.

- AADJRSQ
- ACTUAL
- ADJRSQ
- AGGCHILDPREDICT
- AIC
- AICC
- APC
- DFE
- END
- ENDOBS
- EQUALITY
- ERROR
- FINALPREDICT
- GMAPE

- GMAPES
- GMAPPE
- GMASPE
- GMRAE
- ISRECONCILED
- LEAF
- LLOCK
- LOWBFOVR
- LOWER
- LOWERBD
- MAE
- MAPE
- MAPES
- MAPPE
- MASE
- MAX
- MAXAPES
- MAXERR
- MAXPE
- MAXPPE
- MAXRE
- MAXSPE
- MDAPE
- MDAPES

- MDAPPE
- MDASPE
- MDRAE
- ME
- MEAN
- MIN
- MINAPES
- MINERR
- MINPE
- MINPPE
- MINRE
- MINSPE
- MPE
- MPPE
- MRAE
- MRE
- MSE
- MSPE
- N
- NAME
- NMISS
- NMISSA
- NMISSP
- NOBS

- NONMISSCHLD
- NOTE
- NPARMS
- OLOCK
- OVERRIDE
- PREBFOVR
- PREDICT
- RECDIFF
- RMSE
- RSQUARE
- RWRSQ
- SBC
- SMAPE
- SSE
- SST
- START
- STARTOBS
- STD
- STDBFOVR
- STDDEV
- SUM
- TOP
- TSS
- ULOCK

- UMSE
- UNLOCK
- UPPBFOVR
- UPPER
- UPPERBD
- URMSE
- XML

Appendix 3

SAS Demand Classification and Clustering

Overview of SAS Demand Classification and Clustering

The primary challenge in demand forecasting is to plan a forecasting strategy that minimizes forecast error. By using the available demand history, you can get in-depth information about the demand patterns for a time series.

Using the SAS Demand Classification and Clustering segmentation strategy, you can classify demand patterns, group the time series based on certain criteria, and then apply suitable modeling techniques to forecast demand for various levels in the hierarchy.

Note: You cannot use the segmentation strategy for SAS Demand Classification and Clustering if the values of your time ID variable are in a datetime format.

About the Modeling Strategies for Demand Classification

SAS Forecast Server Client is shipped with these modeling strategies for demand classification:

Demand Class Log Timespan Intermittent
: For this type of demand, the products sell all year round or sell for a relatively long period of time. The time periods between the periods of demand are significantly larger than the unit of time that is used for the forecast period. Herbs, spices, or sauces used for exotic cuisines belong to this demand classification.

Demand Class Long Timespan Non-seasonal
: For this type of demand, historical observations over a long time span indicate that the product has a consistent demand throughout the year (for example, milk and bread).

Demand Class Long Timespan Seasonal
: Products in this demand class sell all year and also have a seasonal pattern (for example, ice cream).

Demand Class Long Timespan Seasonal Intermittent
: For this type of demand, the products sell all year round or sell for a relatively long period of time. However, the time periods between the demands are larger than the unit of time used for the forecast period, and some seasonal patterns are also observed in the intermittent series. For example, a lawn mower has a long time span seasonal intermittent demand pattern. The lawn mower might sell at any time of the year, but during some seasons, the sales are higher.

Demand Class Low Volume
: For some types of products such as luxury items, the volume of demand is too low to accurately indicate a demand pattern. For products with low demand volume, you can consider aggregating the demand to generate stronger demand signals, or use a naive model for forecasting. For example, limited-edition watches have a low volume of demand.

Demand Class Short

In this type, there is too little historical data to make any classification decision. New products are classified in this category.

Demand Class Short Timespan Intermittent

Products in this demand class have sporadic sales during a season or during a short period of time. The demand for such products is inconsistent during that short time span. Slow-moving winter jackets belong to this demand class.

Demand Class Short Timespan Non-intermittent

Products in this demand class have seasonal sales, or sales that occur for a short duration. However, the demand for these products is continuous during that short time span or season. Fast-moving winter jackets belong to this demand class.

Recommended Reading

- *SAS/ETS User's Guide*
- *SAS Forecast Server Procedures: User's Guide*

For a complete list of SAS publications, go to sas.com/store/books. If you have questions about which titles you need, please contact a SAS Representative:

SAS Books
SAS Campus Drive
Cary, NC 27513-2414
Phone: 1-800-727-0025
Fax: 1-919-677-4444
Email: sasbook@sas.com
Web address: sas.com/store/books

Glossary

data definition
an element that specifies the name and location of the input data source, specifies the hierarchy for the data, sets the time dimensions for the data, and assigns roles to the variables in the input data source.

holdout sample
the number of periods of the most recent data that should be excluded from the parameter estimation. The holdout sample can be used to evaluate the forecasting performance of a candidate model.

horizon
the number of intervals into the future, beyond a base date, for which analyses and predictions are made.

ISO year
a year as specified by ISO 8601, and represented by a minimum of four digits for the year (YYYY), followed by digits that represent the week number and the weekday number. For example, 2006–W52–7 (or 2006W527) is the Sunday of the 52nd week of 2006.

lead time
the number of periods into the future that a forecast was made.

metadata
descriptive data about data that is stored and managed in a database, in order to facilitate access to captured and archived data for further use.

model selection list

a list of candidate model specifications. You can choose which model specification is best suited to forecast a particular time series.

SAS Folders

a metadata repository where you can save projects, data definitions, and strategies so they are available to other users at your site.

Index

I

L

M

N

O

P

R

S

T

U

V