

Industry-specific

QAD SOLUTIONS

Manufacturing Applications

User Guide QAD Demand Management



78-0651A
Version 4.5
September 2006

This document contains proprietary information that is protected by copyright and other intellectual property laws. No part of this document may be reproduced, translated, or modified without the prior written consent of QAD Inc. The information contained in this document is subject to change without notice.

QAD Inc. provides this material as is and makes no warranty of any kind, expressed or implied, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. QAD Inc. shall not be liable for errors contained herein or for incidental or consequential damages (including lost profits) in connection with the furnishing, performance, or use of this material whether based on warranty, contract, or other legal theory.

QAD and MFG/PRO are registered trademarks of QAD Inc. The QAD logo is a trademark of QAD Inc.

Designations used by other companies to distinguish their products are often claimed as trademarks. In this document, the product names appear in initial capital or all capital letters. Contact the appropriate companies for more information regarding trademarks and registration.

Copyright ©2006 by QAD Inc.

QAD Inc.

6450 Via Real

Carpinteria, California 93013

Phone (805) 684-6614

Fax (805) 684-1890

<http://www.qad.com>

Contents

About This Guide	1
Other MFG/PRO Documentation	2
QAD Web Site	2
Conventions	3
 Chapter 1 QAD Demand Management	 5
Overview	6
Plan Demand	7
Collaborate with Others	7
Plan for New Product and Product Life Cycles	7
Measure Performance and Manage by Alerts	8
 Chapter 2 Creating and Managing Users	 9
Introduction	10
Managing Users	10
Adding New Users	10
Create E-Mail Subscriptions	13
Editing Users	15
Deleting Users	16
Setting User Permissions	16
Setting Up E-Mail Service	22
Setting Up the E-Mail Service	22
Configuring the E-Mail Service	23
Starting the E-Mail Service	24

Chapter 3 Creating and Managing Scenarios 27

Introduction	28
Logging In for the First Time	29
Creating a New Scenario	30
Editing a Scenario	30
Sharing a Scenario	32
Deleting a Scenario	33
Removing Sharing from a Scenario	33
Copying a Scenario	34
Capturing Data for a Scenario	35
Selecting a Data Source	35
Selecting Tables or Queries	40
Using the Source Data Filtering Wizard	40
Specifying Data Layouts	43
Including Observations in Data	48
Using Series Properties	49
Managing Safety Stock	51

Chapter 4 Using the Living Master Scenario 53

Introduction	54
Setting Up Access to Living Master Scenario Data	54
Creating a Living Master Scenario	55
Populating a DTL Scenario	55
Committing Changes Back to the Living Master Scenario	57

Chapter 5 Using the Engine Manager 59

Introduction	60
Using the Engine Manager	60
Managing Engine Usage Permissions	63
Creating an Engine Usage Permission Entity	64
Editing an Engine Usage Permission	65
Assigning Engine Usage Permission to a User	65
Removing an Engine Usage Permission	66
Managing Engine Usage Domains	67

Creating a Domain	67
Editing a Domain	69
Assigning a Domain to a User	69
Removing a Domain	69
Chapter 6 Viewing Forecasts	71
Introduction	72
Managing the Viewer	73
Logging In As Administrator	73
Logging In As a User	73
Viewing Forecast Data	74
Graph View Tab	75
Data View Tab	83
Demand Data Tab	86
Series Information Tab	87
User-Defined Series Properties	89
Expanded View	90
Making Adjustments in Expanded View	92
Status Properties	93
Modifying User-Defined Field	93
Finding a Series	93
Editing a Forecast	94
Chapter 7 Customizing the Viewer	95
Introduction	96
Display Sets	96
Managing Display Sets	97
Creating a Display Set	97
Editing a Display Set	100
Sharing a Display Set	100
Removing a Display Set	101
Customizing the Graph View	101
Customizing the Demand Data View	102
Customizing the Expanded View	103

Customizing the Viewer Framework	104
Using Notes to Track Assumptions	104
Inserting a Note at Series Level	104
Inserting a Note at Observation Level	105
Sending Notes to Engine User	106
Managing Custom Tasks	107
Creating a Custom Task	107
Deleting a Custom Task	110
Editing a Custom Task	110
Sharing a Custom Task	110
Running a Custom Task	111

Chapter 8 Using Filters and Exceptions 113

Introduction	114
Managing Filters	114
Using Filter Expressions	115
Creating a Filter	115
Launching a Filter	118
Editing a Filter	118
Sharing a Filter	120
Removing a Filter	121
Deactivating a Filter	121
Managing Exceptions	122
Creating an Exception Template	122
Sharing an Exception	128
Deleting an Exception	128
Processing an Exception	129

Chapter 9 Using ABC Analysis 131

Introduction	132
Creating an ABC Template	133
Specifying Groups and Output	134
Using an ABC Template	135
Analyzing Results	136

Sharing an ABC Template	137
Chapter 10 Using the Event Manager	139
Introduction	140
Creating an Event	141
Chapter 11 Using Conversion Factors.....	149
Introduction	150
Renaming a Conversion Factor	150
Applying a Conversion Factor	152
Adding a New Opinion Line	152
Limitations	153
Chapter 12 Using the Calendar.....	155
Introduction	156
Overview	156
Special Format of the Source Data of Calendar Scenario	157
Creating a Calendar	160
Using a Calendar	164
Checking for Compatibility	164
Sharing a Target Period	166
Chapter 13 Using the Bill of Material	169
Introduction	170
Example	170
Creating a BOM	172
Chapter 14 Configuring Statistical Forecasts.....	177
Introduction	178
Selecting Seasonality	179
Setting Forecast Periods	180
Setting Reload Options	181
Freezing a Scenario	181

Overwriting Forecasts	182
Reloading Series	182
Using Incremental Updates	183
Including New Items in a Forecast	183
Using the Clone Feature	184
Copying/Moving Scenarios	184
Including New Items Using Batch Scheduler	184

Chapter 15 Improving Forecast Accuracy 185

Introduction	186
Same As Method	186
Same As History	186
Same As Seasonality	187
Configure the Same As Option	187
Data Cleansing	190
Forecast Accuracy	192
Forecast Techniques and Methods	194
Configuring Forecast Methods	195
Forecasting Methods with Basic Configurations	195
Forecasting Methods that Require Additional Series or Indexes	198
Regression Methods	204
Multiple Regression	205
Polynomial Parameters	207
Stepwise Parameters	209
Group Forecasting	210

Chapter 16 Working with Reports 213

Introduction	214
Understanding Pivot Reports	214
Using an Existing Template	216
Using the Pivot Table Wizard	217
Generating Reports	219
Report Types	220
Excel and File Report Selection	222


Crystal Reports	223
Running and Viewing Reports	225
Running Reports	225
Viewing Reports	225
Chapter 17 Grouping and Allocation	227
Introduction	228
Hierarchy Management	229
Creating a Hierarchy	229
Editing a Hierarchy	236
Deleting a Hierarchy	237
Applying a Hierarchy to a Scenario	237
Removing a Hierarchy	238
Adjusting a Forecast	238
Rebalancing a Hierarchy	238
Chapter 18 Life Cycle Forecasting	239
Introduction	240
Creating a Profile Scenario	240
Using the Life Cycle Forecast	244
Chapter 19 Using Disconnected Mode	249
Introduction	250
Viewing Connection Status	250
Disconnecting from the Server	250
Database Table Modes	251
Committing Changes to the Server	252
Locking Scenarios in Disconnected Mode	252
Using the Connection	252
Limitations in Disconnected Mode	253
Chapter 20 Using the Batch Scheduler	255
Introduction	256
Managing Batch Scheduler	258

Running the Batch Scheduler	258
More Options	262
Managing Massive Data Update (MDU)	263
Massive Series Update	264
Massive Observation Update	269

Chapter 21 Using Statistical Methods..... 281

Holt-Winters Method	282
Example of Holt-Winters	282
Graph Results	284
Audit Trail Results	284
Simple Exponential Smoothing	285
Example of Simple Exponential Smoothing	286
Summarizing the Results	286
Adaptive Exponential Smoothing	287
Example of Adaptive Exponential Smoothing	287
Summarizing the Results	288
Box-Jenkins ARIMA Modeling	289
Box-Jenkins Methodology	290
Example of Box-Jenkins Methodology	291
Performing ARIMA Analysis	292
Holt's Double Exponential Smoothing	293
Example of Holt's Double Exponential Smoothing	293
Summarizing the Results	294
Additional Smoothing Methods	295
Moving Average	295
Weighted Moving Averages	296
Brown's Double Exponential Smoothing	297
Brown's Triple Exponential Smoothing	297
Census X-11	298
Example of Census X-11	299
Summarizing the Results	301
Forecasting with Advanced Census X-11	302
Croston's Intermittent Model	304

Example of Croston's Intermittent Model	305
Summarizing the Results	306
Slow Method	307
Graph Results	312
Audit Trail Results	313
Linear Regression	313
Example of Linear Regression	314
Summarizing the Results	315
Stepwise Regression	316
Stepwise Forward Regression	317
Stepwise Backward Regression	317
Stepwise-Forward-with-a-Backward-Look Regression	317
Example	317
Summarizing the Results	319
Multiple Regression	320
Example of Multiple Regression	321
Summarizing the Results	322
Polynomial Regression	323
Example of Polynomial Regression	324
Summarizing the Results	325
New Product Forecasting	327
Using New Product Forecasting	328
Summarizing the Results	329
Using and Creating an Event Model	330
Example of an Event Model	331
Specifying Statistical Options	335
Statistics Tab	335
Glossary of Descriptive Statistics	339
Glossary of Measurement Statistics	341
Index	345

The background of the page is a grayscale photograph of several interlocking mechanical gears. The gears are of different sizes and are positioned in a way that creates a sense of depth and motion. The lighting is soft, highlighting the metallic texture and the sharp edges of the gear teeth.

About This Guide

Other MFG/PRO Documentation 2

QAD Web Site 2

Conventions 3

This guide covers the features of QAD Demand Management, an application that provides true demand management capabilities that let you manage the need for your company's products or services in a way that meets your company-specific goals and objectives. QAD Demand Management is part of the QAD application suite and lets you forecast demand based on data from QAD MFG/PRO.

Other MFG/PRO Documentation

- For an overview of new features and software updates, see the *Release Bulletin*.
- For software installation instructions, refer to the appropriate installation guide for your system.
- For conversion information, refer to the *Conversion Guide*.
- For instructions on navigating and using the QAD Desktop interface, see *User Guide: QAD Desktop*.
- For instructions on navigating the MFG/PRO Windows and character environments, refer to *User Guide Volume 1: Introduction*.
- For information on using MFG/PRO, refer to the *User Guides*.
- For technical details, refer to *Entity Diagrams* and *Database Definitions*.
- To view documents online in PDF format, see the *Documents on CD* and *Supplemental Documents on CD*. The CD-ROM media includes complete instructions for loading the documents on a Windows network server and making them accessible to client computers.

Note MFG/PRO installation guides are not included on a CD. Printed copies are packaged with your software. Electronic copies of the latest versions are available on the QAD Web site.

QAD Web Site

The QAD Web site provides a wide variety of information about the company and its products. You can access the Web site at:

<http://www.qad.com>

For MFG/PRO users with a QAD Web account, product documentation is available for viewing or downloading at:

<http://support.qad.com/>

You can register for a QAD Web account by accessing the Web site and clicking the Accounts link at the top of the screen. Your customer ID number is required. Access to certain areas is dependent on the type of agreement you have with QAD.

Most user documentation is available in two formats:


- Portable document format (PDF). PDF files can be downloaded from the QAD Web site to your computer. You can view them with the free Adobe Acrobat Reader. A link for downloading this program is also available on the QAD Web site.
- HTML. You can view user documentation through your Web browser. The documents include search tools for easily locating topics of interest.

Features also include an online solution database to help MFG/PRO users answer questions about setting up and using the product. Additionally, the QAD Web site has information about training classes and other services that can help you learn about MFG/PRO.

Conventions

This document uses the text or typographic conventions listed in the following table:

If you see:	It means:
monospaced text	A command or file name.
<i>italicized monospaced text</i>	A variable name for a value you enter as part of an operating system command; for example, <i>YourCDROMDir</i> .
indented command line	A long command that you enter as one line, although it appears in the text as two lines.
Note	Alerts the reader to exceptions or special conditions.
Important	Alerts the reader to critical information.
Warning	Used in situations where you can overwrite or corrupt data, unless you follow the instructions.



Chapter 1

QAD Demand Management

This section introduces QAD Demand Management and provides an overview of its features.

Overview **6**

Overview

The QAD Demand Management and Collaborative Portal applications are designed to encompass all areas of demand management, from creating forecasts and tracking errors, to collaborating with supply chain members, reporting on demand activities, and integrating seamlessly with other MFG/PRO applications.

When combined with MFG/PRO functionality, QAD Demand Management completes an extensive integrated planning model that covers all areas of demand management. QAD Demand Management provides true demand management capabilities that let you manage the need for your company's products or services in a way that meets your company-specific goals and objectives.

Using QAD Demand Management, you can:

- Reduce supply chain inventories
- Optimize production efficiencies
- Reduce planning expenses
- Increase fill rates
- Increase customer service levels
- Create more accurate corporate plans
- Increase profitability
- Grow sales revenue

Key features of QAD Demand Management include:

- Seamless integration with MFG/PRO
- Ease of implementation
- Ease of use and low entry price
- Strong demand management planning core with expanded capabilities
- Sales/Marketing/Finance/Operations all have input into one stable baseline forecast
- Ability to understand supplier demand in real-time and proactively manage inventory levels
- Improve overall customer satisfaction and increase revenues

- Back-end data mart supports robust reporting
- Effective planning for product introductions and retirements

Plan Demand

QAD Demand Management starts with providing the most accurate statistical forecasting available and integrates capabilities to better manage demand. QAD Demand Management can pull in EDI and point-of-sale data as well as plans from customers, suppliers, and internal departments. This enterprise approach to demand management enhances forecast accuracy and highlights opportunities for increased revenues and decreased costs.

QAD Demand Management is flexible to accommodate specific best practices. Hierarchies can be top-down, bottom-up, and middle-out. The system automatically forecasts based on the specifics of each company's market.

Collaborate with Others

The QAD Collaborative Portal allows the sharing of information and management of exceptions with external participants such as sales representatives, managers, customers, and suppliers. Key members of the supply chain can be automatically notified of exceptions so that issues can be quickly addressed. The Portal allows users to manage by exception, record promotions, perform ad hoc reporting, and view data in user-defined levels of detail.

Plan for New Product and Product Life Cycles

New products can be managed by selecting the New Product Forecasting Technique grouping and applying four introduction-specific calculations that build models based on limited historical data.


Life cycle forecasting enables users to apply a pattern to the forecast of a series to compensate for predictable growth or decrease in a products demand that is not captured by most statistical forecasting methods. Life cycle forecasting supports the application of predetermined demand

profiles to newly created items and cloning/same as or retagging capabilities simplify the use of historical analogy in which the demand for a new product is expected to be similar to demand for another product.

Measure Performance and Manage by Alerts

The purpose of performance measurement is to determine the effectiveness of the demand management process. It is a measurement of the competence of the sales, marketing, and product/brand management organizations.

QAD Demand Management has been designed to provide a broad range of standard, executive, and audit level reports. It has also been developed to fully utilize the power of pivot tables.

The background of the page is a grayscale image of several interlocking gears. The gears are of different sizes and are positioned in a way that they appear to be meshing together. The lighting creates highlights and shadows on the teeth of the gears, giving them a three-dimensional appearance. The overall tone is technical and industrial.

Chapter 2

Creating and Managing Users

This section describes how to set up and manage users in QAD Demand Management.

Introduction **10**

Managing Users **10**

Setting Up E-Mail Service **22**

Introduction

QAD Demand Management has two types of security:

- User Management, which allows a system administrator to add, delete, or modify user login accounts.
- Internal security, which provides access to data, scenarios, and functions such as ABC Analysis, Life Cycle Forecast, Display Sets, and others.

To use QAD Demand Management, a user must first have a user ID and password or login account. The application contains a User Management function that lets system administrators add, edit, or delete users, and manage other settings including the following:

- User ID
- Name
- Password
- Permissions

During the installation of QAD Demand Management, a default system administrator user ID and password of fxadmin are created automatically. You should change the default password for this account.

Managing Users

Follow the instructions in this section to add, edit, or delete user accounts, as well as to assign permissions to a user.

Adding New Users

To access the User Management feature, you must have the necessary permissions set by your system administrator.

Follow these instructions to add a new user to QAD Demand Management.

- 1 Select Start|All Programs|QAD Demand Management|Demand Management Engine|Base Unit from the Windows menu. The Base Unit window appears.





Fig. 2.1
Base Unit

- 2 Enter `fxadmin` in the User ID field and `fxadmin` in the Password field, and then click the Log in button.

Note Once QAD Demand Management is installed, you should change the password on the `fxadmin` account to prevent unauthorized access to the system.

- 3 Click the User Management button. The User Management window appears.

Fig. 2.2
User Management



- 4 Click Add. The User Information window appears.

Fig. 2.3
User Information

User Information

User Id:

First Name:

Middle Initial:

Last Name:

E-Mail:

Password:

Confirm Password:

Type:

- 5 Enter the user information in the fields. The information you enter in these fields differentiates one user from another.

User ID. Using ten characters or less, enter a unique User ID for this new user. The user enters this ID to log into QAD Demand Management.

First Name. Enter the user's first name.

Middle Initial. Enter the user's middle initial.

Last Name. Enter the user's last name.

E-Mail. Enter a valid e-mail address for the user.

Password. Enter a password for the user of 24 characters or less.

Confirm Password. Reenter the user's password to confirm that it was entered correctly.

Type. Enter a type or category for this user of 10 characters or less.

6 Click Accept.

The User Management screen reappears with the new user added to the list of users.

Create E-Mail Subscriptions

E-mail notifications are a feature of QAD Demand Management. The system defines eight issues or events when a user can be notified, listed in Table 2.1.

Note To use this feature, you must configure an e-mail service. See “Setting Up E-Mail Service” on page 22.

Issue	Description	E-Mail Content
The Engine finished forecasting a Living Master scenario	The forecasting process is complete and run against a Living Master scenario.	The user receives the first 20 items and the total forecast. Only users that have access to the items and the total forecast can receive the notifications.
A scenario has been marked as disconnected	A user chooses to disconnect from the server.	The scenario name was marked as disconnected. See Chapter 19, “Using Disconnected Mode,” on page 249.
A scenario has been marked as connected	A user connects to server.	The scenario was marked as connected. See Chapter 19, “Using Disconnected Mode,” on page 249.
User populate his/her DTL scenario	A user has populated data in the DTL scenario.	The user receives the first 20 items and the total forecast. Only users that have access to the items and the total forecast can receive the notifications.

Table 2.1
E-mail
Subscriptions

Issue	Description	E-Mail Content
User commit his/her DTL scenario	A user has committed data in the DTL scenario.	The user receives the first 20 items and the total forecast. Only users that have access to the items and the total forecast can receive the notifications.
Forecast error report	Errors have occurred during forecasting a scenario.	Error report.
Batch process results	Batch Scheduler sends the status of the processes.	Process status, errors, and exceptions.
Exception results	The system detects which series now fulfill an exception and which fields no longer fulfill the exception. The first time an exception runs, all series are included in the list of new items that fulfill it.	Subscriber receives the items in the following way: • List of new items that fulfill an exception • List of items that no longer fulfill the exception Note that items that already fulfill the exceptions are not sent. To determine the new items that fulfill the exception, the Engine keeps a list of the items from the previous exception execution. That list is compared with the actual exception list.

To assign notifications to users:

- 1 In the User Management window, select the user you want to assign a notification to and click Edit.
- 2 Click the Notifications button to specify the e-mail subscriptions for the current user.




Fig. 2.4
E-Mail
Subscriptions

- 3 Select the issues or events that the user requires. The user receives e-mail notifications only for the issues that are selected.
- 4 Click Accept. The User Information window is displayed. Click Accept again. The e-mail subscriptions are applied to the user account.
- 5 Click Close.

Editing Users

Follow the instructions in this section to edit details for an existing QAD Demand Management user.

- 1 Click the User Management button. The User Management window is displayed.
- 2 Select the user you want to edit.
- 3 Click Edit. The User Information window appears showing the user's existing information. The user ID cannot be changed.
- 4 Modify the user information as required.
- 5 Click Accept. The User Management window reappears with the updated information.

Deleting Users

Follow the instructions in this section to delete an existing QAD Demand Management user.

- 1** Click the User Management button. The User Management window is displayed.
- 2** Select the user you want to delete.
- 3** Click Delete. A delete confirmation message appears. If the user you want to delete owns any scenarios, the message indicates that those scenarios are owned by fxadmin.
- 4** Click Yes. The User Management window reappears. The user is removed from the list.

Setting User Permissions

User permissions control which QAD Demand Management functions are available to a user. You can grant or revoke permissions for users.

To grant or revoke user permissions:

- 1** Click the User Management button on the Base Unit window. The User Management window is displayed.
- 2** Click Permissions. The User Permissions window appears (see Figure 2.5). All fields are selected, meaning that the user has access to these functions, unless you clear them.




Fig. 2.5
User Permissions

- 3 From the Set Security For drop-down list, select the name of the user for which you want to set permissions.
- 4 From the Security Type drop-down list, choose the type of security you want to assign to the user.

Full. Gives the user permissions to use all QAD Demand Management components and functions.

Read-Only. Gives the user read-only permissions to all QAD Demand Management functions. The user cannot schedule, create, change, or share any scenarios or data.

Custom. Gives the user permissions to the QAD Demand Management functions that you select.

- 5 If you choose Custom security type, clear the appropriate permission options on the Security Application tree.
- 6 Choose the appropriate global permissions for the user; these apply to all scenarios.

Use Base Unit. The Base Unit defines the data source, scenario configuration, forecasting technique, and statistical reports.

If you grant the user permission to the Base Unit, the user is automatically granted permission to all the following options. However, you can grant permission to the following options individually.

Create Scenario: Lets the user create scenarios.

Share Scenario: Lets the user share his scenarios with other users.

Open Base Unit: Lets the user open the Base Unit from Viewer.

Use Allocator. This option is no longer supported.

Use Batch Scheduler. The user is able to access the Batch Scheduler for scheduling repeatable jobs such as forecasting data, reports, and incremental updates. Also lets the user schedule batch reports.

System Resources. The system lets the user manage custom functions.

If you grant permission to manage system resources, the user is automatically granted permission to all the following options. However, you can grant permission to the following options individually.

Manage Filters: Lets the user define and manage filters.

Manage Display Sets: Lets the user define and manage display sets.

Manage Exceptions: Lets the user define and manage exceptions.

Manage Events: Lets the user define and manage events.

Manage Custom Tasks: Lets the user define and manage custom tasks based on stored procedures.

Manage ABC Templates: Lets the user define and manage ABC templates.

Manage Purchase Orders: Lets the user define, manage, and use purchase orders.

- 7 Choose the appropriate shared scenario permissions for the user. These options apply to the functions available to a user sharing a scenario with another user.

For example, if User A shares Scenario1 with User B, User B can change the period and seasonality of Scenario 1 only if User B has permission to do so.

If a user is not granted any shared scenario permissions, all scenarios shared with the user are read-only.

Data Source. Lets the user change the data source for a shared scenario. If you grant the user permission to change Data Source information, the user automatically has permission for Table Selection and Data Organization.

Table Selection: Lets the user select which table in the data source to use for the forecast.

Data Organization: Lets the user select the organization method of the data used in the selected table, either columns or rows.

Period and Seasonality. Lets the user adjust period and seasonality data for a shared scenario.

Statistics. Lets the user add or change statistical information in a shared scenario.

Method Configuration. Lets the user add or modify method information options for a shared scenario.

If you grant the user permission to Method Configuration, the user is automatically granted permission to all the following options.

However, you can grant permission to the following options individually.

Select Method: Lets the user select the method used for the forecast of a shared scenario.

Modify Parameters: Lets the user change parameters of a method used in a shared scenario.

Series Configuration: Lets the user add or change series information on a method used for the forecast a shared scenario.

Advanced Configuration. Lets the user add or change Advanced Configuration options of a shared scenario.

If you grant the user permission to Advanced Configuration, the user is automatically granted permission to all the following options.

However, you can grant permission to the following options individually.

Dates: Lets the user add or change information regarding the dates used in the forecast of a shared scenario.

Series Layout: Lets the user add or change information regarding the series fields used in the forecast of a shared scenario.

Same As: Lets the user add or change information regarding the Same As fields used in the forecast of a shared scenario.

Load Incrementally: Lets the user add or change information regarding the incrementally loaded observations used in the forecast of a shared scenario.

Cleansing: Lets the user add or change information regarding the Cleansing information used in the forecast of a shared scenario.

Accuracy: Lets the user add or change information regarding the Accuracy information used in the forecast of a shared scenario.

Viewer Access. Grants access to change the series information or observations for any series in a scenario.

If you grant the user permission to Viewer Access, the user is automatically granted permission to all the following options. Also, you can grant permission to the following options individually.

Modify Series: Lets the user change the series while previewing a shared scenario.

Modify Observations: Lets the user change the observation information of a shared scenario.

Reports. Lets the user change the report information of a shared scenario.

If you grant the user permission to Reports, the user is automatically granted permission to all the following options. However, you can grant permission to the following options individually.

Run Report: Lets the user run reports on a shared scenario.

Report Options: Lets the user add or change report options for a shared scenario.

Procast. Lets the user run a forecast on a shared scenario. The default method is Procast.

Reload Series. Lets the user reload the series in a shared scenario.

Groups. Lets the user change the Group information of a shared scenario.

If you grant the user permission to Groups, the user is automatically granted permission to all the following options. However, you can grant permission to the following options individually.

Group Configuration: Lets the user configure groups for the shared scenario.

Group Options: Lets the user add or change groups for the shared scenario.

Freeze Scenario. Lets the user freeze or unfreeze scenarios without the risk of reloading new data.

Overwrite Forecast. Lets the user select a specific range of the forecast to be overwritten with a new system-generated forecast.

Tools. Lets the user access Engine Tools/Resources against a shared scenario.

If you grant the user permission to Tools, the user is automatically granted permission to all the following options. However, you can grant permission to the following options individually.

Hierarchy: Lets the user rebalance or remove the hierarchy in a scenario shared to the user.

Reforecast: Lets the user reforecast a scenario shared by the user.

Calendar: If this option is active, all shared scenarios can be possible target scenarios for Calendar. Otherwise, those scenarios are not displayed as possible target scenarios.

Conversion Factors: Lets the user apply conversion factors in a scenario shared by the user.

Life Cycle Forecast: Lets the user apply Life Cycle Forecast factors in a scenario shared by the user.

ABC Analysis: Lets the user use ABC Analysis in a scenario shared by the user.

Bill of Materials: Lets the user use Bill of Material in a scenario shared by the user.

Run Exception: Lets the user apply an exception in a scenario shared by the user.

Purchase Orders: Lets the user run a purchase order in a scenario shared by another user.

- 8 When you have finished granting the user permissions, click Save. Confirm your changes.
- 9 Click the Close button. The Base Unit main window appears.

Resetting Permissions


Click Reset to reset a user's previous permission settings. The previous saved permissions for the user are activated.

Setting Up E-Mail Service

The e-mail service provides direct notification to users by e-mail. The notifications describe the event and action that triggered the e-mail message. The Engine proactively notifies users so that they can take action when they subscribe to an event. See “Create E-Mail Subscriptions” on page 13.

Figure 2.6 illustrates the mechanism used to drive the e-mail process.

Fig. 2.6
E-Mail Notification
Process



Technical users can insert their own messages into the e-mail notification queue. By doing this, the system can be extended to send out a range of additional message that add additional value to the forecasting process.

Setting Up the E-Mail Service

QAD Demand Management uses the SQL e-mail service. Before you can use the e-mail service, the following conditions must be met:

- An e-mail profile exists on the SQL e-mail service.

- A MAPI profile is configured in the mail option of the control panel using the mail service.
- SQL Server agent is running.
- SQL Server Mail is running properly with a valid e-mail profile.

For details on the SQL Server e-mail service, refer to your SQL Server documentation.

Note The SQL Server 2000 SQL Mail supports an extended MAPI.

Configuring the E-Mail Service

The Engine e-mail service is a job that is scheduled to run periodically, usually every few minutes. The job executes a stored procedure that sends a message to check if messages exist in the e-mail service queue.

Messages are stored in the TBL_EMAIL_NOTIFICATION table. This table stores a value with the message indicating the processing status:

- 1: The message was processed.
- 0: The E-mail Service did not process the message.


For maintenance purposes, you can check that table and remove the old messages.

Starting the E-Mail Service

To start the e-mail service:


- 1 Create an SQL Server job. The job must be scheduled to execute periodically, such as every minute. Figure 2.7 shows an example.

Fig. 2.7
Edit Job Schedule



- 2 Specify the job to execute the stored procedure `SP_Email_SendMail`. Execute the task against Engine Database (default is `FX_JGD`). Figure 2.8 shows an example.

Fig. 2.8
Edit Job Step



A window is displayed showing the status of the job.





Fig. 2.9
E-mail Job Status
Window



Chapter 3

Creating and Managing Scenarios

This section describes how to create and manage scenarios, as well as how to set up data for a scenario.

<i>Introduction</i>	28
<i>Logging In for the First Time</i>	29
<i>Creating a New Scenario</i>	30
<i>Editing a Scenario</i>	30
<i>Sharing a Scenario</i>	32
<i>Deleting a Scenario</i>	33
<i>Removing Sharing from a Scenario</i>	33
<i>Copying a Scenario</i>	34
<i>Capturing Data for a Scenario</i>	35
<i>Using the Source Data Filtering Wizard</i>	40


Introduction

QAD Demand Management uses scenarios to create forecasts. A scenario is composed of data, statistical information, and instructions. Scenarios bring data into the engine as well as provide the mechanism for developing forecasts. Proper scenario setup when combined with the Batch Scheduler lets you manage data effectively within QAD Demand Management.

Usually QAD Demand Management is configured with a single Living Master scenario that contains the active version of the forecast. Often you also use helper scenarios to bring additional data into the engine.

Figure 3.1 illustrates a simple scenario that pulls in and forecasts unit data, and then adds time-phased pricing back to the Living Master as a conversion factor.

Fig. 3.1
Typical Living
Master Scenario



In this case the Batch Scheduler loads price update attributes. By segmenting the data into multiple scenarios, it is easy to manage and update each scenario and then post the results into the Living Master scenario.

Each scenario is unique and contains different information. QAD Demand Management uses each scenario to examine and forecast a different aspect of your business. In order to ensure the integrity of your forecasts, you must closely manage your scenarios. This section explains how to keep scenarios unique as well as share scenarios with other users.

Logging In for the First Time

The Data Capture tab is displayed when you log into QAD Demand Management for the first time.



Fig. 3.2
Base Unit, Data
Capture Tab


Most users have a shared scenario or an assigned detail scenario, which is displayed after initial login. When no scenarios have been shared or assigned, a scenario called First Scenario is created for you. Use this default scenario as it exists, rename it, or create a new scenario.

Creating a New Scenario

To create a new scenario:


- 1 Click the Scenario Management button. The Scenario Management window is displayed.

Fig. 3.3
Scenario
Management



- 2 Click the New button. The Add New Scenario pop-up is displayed.

Fig. 3.4
Add New Scenario



- 3 Enter a unique name in the Scenario field and click OK. A confirmation message is displayed showing that the scenario was created.

Editing a Scenario

To edit an existing scenario:

- 1 Click the Scenario Management button. The Scenario Management window is displayed.
- 2 Select the scenario and click Edit. The Edit Scenario window is displayed.

Fig. 3.5
Edit Scenario

3 Edit the scenario attributes as required.

The attributes are used for further definition and to group scenarios. For example, if Time Level 1 is set to Monthly, then a forecast can be grouped for all scenarios that forecast monthly.

Time Level 1. The default value is Monthly. The predefined values are 4 Weeks, Daily, Monthly, Quarterly, and Weekly. The values are defined in TBL_ModuleClassValue.

Time Level 2. This attribute contains the same default and predefined values as Time Level 1.

Note For details, see Chapter 12, “Using the Calendar,” on page 155.

User Field 1. This attribute determines the internal type of scenario.

Normal: A general scenario used for forecasting. This is the default value.

Living Master: A scenario holding detail data. It is controlled by the administrator and acts as a data warehouse. Other users populate data from the Living Master scenarios, make changes, and commit the changes back to Living Master scenarios.

Profile: A scenario holding profiling data using in a life-cycle forecast routine.

Calendar: A scenario holding the transformation information for converting data from one time scale to another.

Currency: A scenario holding currency exchange rate information, used by the Web-based Planning Portal.

Note For details, see Chapter 12, “Using the Calendar,” on page 155 and Chapter 17, “Grouping and Allocation,” on page 227.

Sharing a Scenario

When you create a scenario, only you can view and edit the scenario information. To permit another user to view or change your scenario, you must share the scenario. Depending on the shared scenario permissions granted, the user can share a scenario with another user, view the scenario in read-only mode, or change the scenario.


Note If you designate a scenario as Living Master, other users can extract information from your scenario, make changes, and commit the changes back to the database. Although other users cannot access the scenario directly, they can access the scenario data by using the Connection Wizard.

Note For details on shared scenario permissions, see “Setting User Permissions” on page 16.

To share a scenario with a user:

- 1 Click the Scenario Management button. The Scenario Management window is displayed.
- 2 Select the scenario and click the Share button. The User selection for sharing scenario window is displayed.

Fig. 3.6
User Selection for
Sharing Scenario



All users that have sharing scenario permissions are listed in the Available list.

- 3 In the Available list, select the user you want to share the scenario with.
- 4 Click the single arrow pointing to the Selected list. The user name moves to the Selected list.

Note The single arrows select or deselect one user name at a time. The double arrows select or deselect all user names at once.

- 5 Click OK. A confirmation message is displayed.

Deleting a Scenario

You can delete an unused scenario if needed. However, if you remove any of the special scenarios with Userfield1 set to a value other than normal (such as calendar or currency), the functionality associated with the deleted scenario can no longer be used. Also, if you delete a scenario that is part of a hierarchy, the entire hierarchy becomes unusable. In this case a warning is displayed and you are prompted to continue.

To delete a scenario:

- 1 Click the Scenario Management button. The Scenario Management window is displayed.
- 2 Select the scenario you want to delete and then click the Delete button. A delete confirmation window is displayed.
- 3 To delete the scenario, click Yes. A message is displayed confirming the scenario has been deleted.

Removing Sharing from a Scenario

You can remove sharing from scenario, if needed. However, if you remove sharing from a scenario that is part of a hierarchy, the entire hierarchy becomes unusable. In this case, a warning is displayed and you are prompted to continue.

To remove sharing from a scenario:

- 1 Click the Scenario Management button. The Scenario Management window is displayed.
- 2 Select the shared scenario that you want to remove sharing from and click Unshare. An Unshare Confirmation message is displayed.
- 3 To continue unsharing, select the scenario and click Yes. A message is displayed confirming removal of the share from the scenario.

Copying a Scenario


When the information in a new scenario varies slightly from an existing scenario, you can copy the options of the existing scenario and change some information that is different in the new scenario. The ability to copy a scenario lets you avoid recreating the entire scenario.

Note Copying a scenario only copies the options, not the data.

To copy an existing scenario:

- 1 Click the Scenario Management button. The Scenario Management window is displayed.
- 2 Select the scenario you want to copy and then click the Copy button. Click Yes in the confirmation window. The Copy Scenario window is displayed.

Fig. 3.7
Copy Scenario



- 3 Enter a new, unique name for the scenario and click OK. A Confirmation window is displayed when the copy is complete.
- 4 Click OK.

Capturing Data for a Scenario

Before you can create a scenario, you must first specify a data source that contains forecast data. The types of data sources you can use are:

- Microsoft Access Database file (.mdb). A file created in Microsoft Access that contains a database layout including tables, indexes, and queries.
- Excel Spreadsheet file (.xls). A spreadsheet created in Excel that contains the data in a tabular format.
- SQL Server. A server-based relational database management system.
- Other (ODBC, Oracle). Various other data formats.

Selecting a Data Source

To select a data source:

- 1 Click the Data Capture and View Selection button.





Fig. 3.8
Data Capture and
View Selection

The Data Capture window is displayed.

- 2 Select the type of data source you want to use. After you select the data source, click the Test Connection button to test the connection.

Fig. 3.9
Data Capture



- 3 Click Browse. Depending on the type of data source you select, a different dialog box is displayed. The various data source types are described next.

Selecting a Microsoft Access Database File

If you select a Microsoft Access Database file (.mdb) and then click the Browse button, the Open dialog box is displayed.

- 1 Select the .mdb file and click Open.
- 2 When the Data Capture screen is displayed, click Preview Database. The Database Preview dialog box is displayed.




Fig. 3.10
Database Preview

- 3 Select the table you want to use for the forecast from the Preview Table drop-down list and then click Close.

Selecting an Excel Spreadsheet File

If you select Excel Spreadsheet (.xls) and then click the Browse button, the Open dialog is displayed.


- 1 Browse to the location of your .xls file and click Open.
- 2 When the Data Capture screen is displayed, click Preview Database. The Database Preview dialog box is displayed (Figure 3.10).
- 3 Select the table from the Preview Table drop-down list and then click Close.

If your data source is an Excel Spreadsheet, you can select a part of the data to use for the forecast. This selection of data is called a range.

To select data in Excel:

- 1 Select the data for the range.

Fig. 3.11
Excel Spreadsheet
Data Range




- 2 Name the range in the cell indicator.

Selecting an SQL Server File

If you select an SQL Server source:

- 1 Click Configure. The Server window is displayed.

Fig. 3.12
Server



- 2 Complete the Server Name, Login ID, Password, and Database fields and then click OK.

Selecting a Different Type of Data Source (ODBC, Oracle)

If you select Other (ODBC, Oracle):

- 1 Click Configure. The Data Link Properties window is displayed.




Fig. 3.13
Data Link
Properties

- 2 Select the data option for the source data connection and then click Next.

Note Refer to your Microsoft or SQL documentation for information on configuring these types of files.

After you select the data source, the path and name of the data source appears in the Actual Selection field at the bottom of the Data Capture window.

- 3 Click Done. The Base Unit window is displayed.

Selecting Tables or Queries

After selecting the data source, you must indicate the specific table or query within the database that you want to use for the forecast.

To select a table or query:

- 1 Click the Tables/Queries drop-down list. The tables, queries, and ranges in your data source are listed.
- 2 Select the table, query, or range and click OK.

Using the Source Data Filtering Wizard

Since tables and queries can be large, you can filter the source data by using the Filtering Wizard to indicate groups or sections of the table or query you want to use.

To do this, you create an expression. An expression is a combination of values that specify exactly which values to use in the forecast.

Example A table in your data source lists earnings for the month of June last year. However, for your forecast, you want to only use the values for the last two weeks in June. You use the Filtering Wizard to build the following expression for QAD Demand Management to only use earning values between the dates of June 15 and June 30.

```
Date >= '%6/15/99%' AND Date <= '%6/30/99%'
```

Example A table lists all your employees, their departments, and start dates. You want to know how many people started in the Marketing and Sales departments in the year 1999. You build an expression that specifies to only use dates in 1999 and employees in the Sales or Marketing departments.

```
Start Date >= '%1/1/99%' AND Start Date <= '%12/31/ 1999%'  
AND (Dept = Sales OR Dept = Market )
```

To open the Filtering Wizard:

- 1 Click the Data Source View and Selection button.




Fig. 3.14
Data Source View
and Selection

The Data Filter window is displayed.

Fig. 3.15
Data Filter

- 2 If the expression requires an open parenthesis, click the open paren button in the Logic Operators and Parenthesis section.
- 3 Select the name of the field for the first part of the expression from the Field List drop-down list.
The values available depend on the fields in the table or query.
- 4 In the next field, choose an operator; the default is = (equals). Click the drop-down list and select an operator:
 - = Equal to
 - <> Not equal to
 - > Greater than
 - < Less than
 - >= Greater than or Equal to
 - <= Less than or Equal to
 - Like

- Is not Null
 - Is Null
- 5 Click the Show Values button. The values that coincide with the field you selected in step 3 populate the Values drop-down list.
 - 6 Select the appropriate value from the Values drop-down list.
 - 7 Click the Add Expression button. The expression you created is displayed in the Filter expression field.
 - 8 To add more to the expression, click the AND or OR buttons in the Logic Operators and Parenthesis section and continue building the expression.
 - 9 To ensure the expression is valid, click Parse Expression. If the expression is not valid, an error message is displayed.
 - 10 Click the Show Results button when you are finished building the expression. All the values you specified display in the window.
 - 11 Edit your expression if required, or click Accept.

Specifying Data Layouts

You must specify how the data in the selected data source is organized. Data is organized in rows or columns.

Note You must specify the correct layout to ensure an accurate forecast.

Data layout is organized by:

- Series, which represents an individual product or item in your database
- Observations, which are a set of values for the series

To indicate the layout of data, select the layout option (Columns or Rows) that describes your data source.

Organizing Data in Rows

When data is organized in rows, the observations go across the X axis (horizontal), as shown in Figure 3.16.

Fig. 3.16
Data Organized in Rows

	A	B	C	D	E	F	G	H
2	BASE MODEL	Mar-97	Apr-97	May-97	Jun-97	Jul-97	Aug-97	Sep-97
3	DT1000	8047	6732	6329	9192	6062	7500	7604
4	DT4300	3785	2819	3770	3464	2590	3163	3141
5	DT9700	1322	564	488	447	336	546	590


When data is in rows, the observations are all numeric values that span the columns.

Advanced Options for Rows

After you select a data source, table/query, and filter, you can configure the data for the forecast. Depending on the organization of data in the data source, you have different configuration options.

To configure data that is organized in rows, click the Advanced button on the bottom of the Base Unit. The Advanced Options window is displayed.

Fig. 3.17
Advanced Options



The options available for configuring data organized in rows are:

- Dates
- Labels
- Observations, a set of values for the series
- Series properties, representing an individual product or item in the database

Organizing Data in Columns

When data is organized in columns, the observations go vertically down the data set or spreadsheet, and series go across the columns, as shown in Figure 3.18.

	Date	Bar_Sales	Internet_Advertising	Television_Advertising
	5/1/1998	2609.61	149.04	74
	6/2/1998	2660.18	149.96	77
	6/3/1998	2463.67	145.04	60
	6/4/1998	2438.54	139.81	65
	6/5/1998	2387.31	138.96	57
	6/6/1998	2071.97	131.37	63
	6/7/1998	1940.39	139.21	69
	6/8/1998	2427.95	140.82	67
	6/9/1998	2533.86	148.68	64
	6/10/1998	2594.59	153.99	74
	6/11/1998	2717.16	159.39	76
	6/12/1998	2876.94	164.81	84
	6/13/1998	2383.05	156.87	77
	6/14/1998	2155.65	157.56	75
	6/15/1998	2743.12	163.7	77
	6/16/1998	2836.24	169.05	79

Fig. 3.18
Data Organized in
Columns

Note Observations span down rows when data is organized in columns. This means that using the Filtering Wizard reduces the length of each series. When pieces of data are eliminated from the forecast, fewer rows are needed to store the data.

Advanced Options for Columns

After selecting a data source, table/query, and filter, you can configure the series in the table. Click the Advanced button on the bottom of application main window. The Advanced Options window is displayed.

The options available for configuring data organized in columns are:

- Series
- Dates

Configuring Series Dates and Labels

This section describes how to configure series dates and labels.

Configuring Series Dates

Regardless of how your data is organized, you can configure the series date fields in your data source. By configuring dates, you specify how dates display in the generated reports.

In the Dates section of Advanced Options, select the appropriate values for the following fields:

Start Year. Specify the starting year of the data in the series.

Start Period. Specify the date cycle. For example, when using weekly data, the value 2 indicates that the dates start in the second week of the year.

Cycles/Year. Specify when the period starts in the beginning of the year. For example, the value 52 means weekly, 365 means daily and 30 means monthly.

Following are examples of various types of data and how the configured date fields appear.

A daily data setup like that shown in Figure 3.19:

Fig. 3.19
Daily Example

The screenshot shows a configuration window titled "Dates". It contains three input fields with dropdown arrows: "Start Year" with the value "1990", "Start Period" with the value "2", and "Cycles/Year" with the value "365". To the right of these fields is a button labeled "Dates".

Returns values like those shown in Figure 3.20.

	A	B	C	D	E	F	G
1	Series	1992-48	1992-49	1992-50	1992-51	1992-52	1993-01
2	[EAST] [DISCOUNT PALA	196.24	197.14	205.70	207.50	201.19	198.49
3	[EAST] [DISCOUNT PALA	197.52	196.31	198.00	213.85	209.62	195.33
4	[EAST] [DISCOUNT PALA	138.07	138.70	144.72	145.98	141.55	139.65
5	[EAST] [DISCOUNT PALA	138.97	138.11	139.30	150.46	147.48	137.42
6	[EAST] [DISCOUNT PALA	99.00	112.89	110.88	108.74	99.00	112.89
7	[EAST] [DISCOUNT PALA	99.67	100.64	113.20	113.49	108.36	101.10
8	[EAST] [DISCOUNT PALA	122.05	122.59	123.14	123.68	124.23	124.77

Fig. 3.20
Daily Results

A weekly data setup of dates like those shown in Figure 3.21:

Dates

Start Year: 1992 Start Period: 48 Cycles/Year: 52

Buttons: [Dates]

Fig. 3.21
Weekly Example

Returns values like those shown in Figure 3.22.

	A	B	C	D	E	F	G
1	Series	1992-48	1992-49	1992-50	1992-51	1992-52	1993-01
2	[EAST] [DISCOUNT PALA	196.24	197.14	205.70	207.50	201.19	198.49
3	[EAST] [DISCOUNT PALA	197.52	196.31	198.00	213.85	209.62	195.33
4	[EAST] [DISCOUNT PALA	138.07	138.70	144.72	145.98	141.55	139.65
5	[EAST] [DISCOUNT PALA	138.97	138.11	139.30	150.46	147.48	137.42
6	[EAST] [DISCOUNT PALA	99.00	112.89	110.88	108.74	99.00	112.89
7	[EAST] [DISCOUNT PALA	99.67	100.64	113.20	113.49	108.36	101.10
8	[EAST] [DISCOUNT PALA	122.05	122.59	123.14	123.68	124.23	124.77

Fig. 3.22
Weekly Results

A monthly data setup like that shown in Figure 3.23:

Dates

Start Year: 2000 Start Period: 1 Cycles/Year: 12

Buttons: [Dates]

Fig. 3.23
Monthly Sample

Returns values like those shown in Figure 3.24.

	A	B	C	D	E	F	G
1	Series	2000-01	2000-02	2000-03	2000-04	2000-05	2000-06
2	[EAST] [DI	196.24	197.14	205.70	207.50	201.19	198.49
3	[EAST] [DI	197.52	196.31	198.00	213.85	209.62	195.33
4	[EAST] [DI	138.07	138.70	144.72	145.98	141.55	139.65
5	[EAST] [DI	138.97	138.11	139.30	150.46	147.48	137.42
6	[EAST] [DI	99.00	112.89	110.88	108.74	99.00	112.89
7	[EAST] [DI	99.67	100.64	113.20	113.49	108.36	101.10

Fig. 3.24
Monthly Results


Configuring Labels

When the forecast is generated, the system names the series in your data source as Series 1, Series 2, and so on. You can control how the system labels the series.

To assign a label to a series:

- 1 Click the Labels button in the Series Configuration section. The Fields Selection for Labels window appears showing the text fields from your data source in the Available list.

Fig. 3.25
Field Selection for
Labels



- 2 Select the fields you want to set up as your series descriptions, and then click the single arrow to move them to the Selected column. Click OK.

Note The order in which you select the fields is important. Use the up arrow and down arrows next to the Selected column to change the order of the fields selected. When using labels with repeated values, QAD Demand Management appends a number to differentiate the same series. For example, if you select Description as the label and have the value Shoe twice, the series name is [Shoe][1] and [Shoe][2].

Including Observations in Data

The system automatically selects all numeric fields as observations in the same order as they appear in the table/query. An observation is a set of values for a series. You can change that configuration by using the Observation button.

Configuring observations is useful when you do not include an observation in your forecast that represents something else. For example, if a database has several fields showing the price and cost of certain products, you can remove them from the observation selection.

To specify the observations to include in data:

- 1 Click the Observations button. The Fields selection for Series window is displayed.
- 2 Select or deselect the fields you want to include in the data.
- 3 Click OK.


Using Series Properties

Use Series Properties to import additional information for each individual series—cost, price, or description, for example.

Note If the data organization is Column, series properties cannot be loaded at the same time when loading observations. For details, see “Massive Series Update” on page 264. This section explains how to import series properties without changing the data organization and reloading the data.

- 1 Click the Series Properties button. The Additional Series Attributes window is displayed.

Fig. 3.26
Additional Series
Attributes



- 2 In the Property column, double-click the attribute you want to apply to a data field. A window is displayed listing the available data fields.
- 3 Select the data field containing the attribute values you want to import. The Additional Series Attribute window is displayed with the attribute with the data field listed next to it.
- 4 Click OK.

You can repeat this process for any data fields that can be mapped into QAD Demand Management. By using display sets you can customize labels and descriptions. For details on display sets, see “Display Sets” on page 96

Managing Safety Stock

QAD Demand Management lets you deal with safety stock in two ways:

- Activate the safety stock feature on the Statistics tab. This lets you set lead time and customer service levels. The numbers you define can be applied to all series.
- Pull safety stock components (lead time, service levels, and so on) directly from the data source. This lets you create series-specific safety stock levels.

Setting Safety Stock Using a Data Source

To set safety stock using a data source:

- 1 Click the Advanced button on the Base Unit window. The Advanced Options window is displayed.
- 2 Click the Series Properties button on the Advanced Options window.
- 3 Double-click SS Lead Time in the Property column. The Field selection for lead time of safety stock window appears.





Fig. 3.27
Field Selection for
Lead Time of
Safety Stock

- 4 Move the field specified for lead time in the data source from the Available list to the Selected list.
- 5 Repeat the previous step for Service Level.
- 6 Click OK.

Note For details on advanced options, including Cleansing, Accuracy, Reloading and Reforecasting, see Chapter 15, “Improving Forecast Accuracy,” on page 185.



Chapter 4

Using the Living Master Scenario

This section describes how to use the Living Master scenario while forecasting and planning.

Introduction **54**

Setting Up Access to Living Master Scenario Data **54**

Creating a Living Master Scenario **55**

Populating a DTL Scenario **55**


Committing Changes Back to the Living Master Scenario **57**

Introduction

The Living Master scenario centralizes your data and provides a data repository for forecasting and planning. The Living Master scenario improves system performance because it lets you take a portion of data and work on it instead of on the whole scenario.

Figure 4.1 illustrates the living master process.

Fig. 4.1
Living Master
Connection Wizard



The Living Master scenario lets you download part of a master scenario into your own scenario, change it, and then commit selected changes back to the master scenario. This prevents different users from editing the same scenario simultaneously, or from having to provide access to multiple users for a group of scenarios.

Setting Up Access to Living Master Scenario Data

A system administrator can provide additional security to the Living Master scenario by using the Engine Manager to define permissions and domains. In this way, administrators can limit the data a user can specify and how they can change the Living Master scenario.

Permissions are administrator-defined filters in the Engine Manager that limit access to the Living Master scenarios and series. For details, see Chapter 5, “Using the Engine Manager,” on page 59.

Domains are administrator-defined access rights to observation lines for the series that the user populates to the DTL (detail) scenario. Access rights specify which observation lines a user can see and whether those observation lines can be changed. For example, a domain can provide a

user access to only the Value, FValue, Original, UserValue01, and ConvFactor01 observation lines, with Value and Original set as read-only. See “Managing Engine Usage Domains” on page 67.

With appropriate security, a user can download or check out a portion of data they need. The user downloads the data into the default scenario DTL (Detail) that is created for them when their user ID is created.

Note Users created before a Living Master scenario was designated do not have a DTL scenario.

Downloading data into a DTL scenario is referred to as *populating the DTL scenario*. A DTL scenario can be populated and repopulated at any time with data from a single Living Master scenario. Users then make changes to their DTL scenario and commit selected changes back to the Living Master scenario.


Creating a Living Master Scenario

Living Master scenarios are normal scenarios created and designated as such by the system administrator. For details, see “Creating a New Scenario” on page 30.

Populating a DTL Scenario

To retrieve data from a Living Master scenario, you must first open the DTL scenario in the Viewer. When opening the DTL scenario for the first time, you are prompted to populate the scenario by the Population Wizard. You can also access the Population Wizard from the Living Master menu (see Figure 4.2).

Fig. 4.2
Population Wizard



The Population Wizard provides three options when downloading data from the Living Master scenario:


- Select a filter from the Select a filter to use drop-down list. Filters limit which series are downloaded from the Living Master Scenario to the DTL scenario. The available filters are either created by or shared. See “Managing Filters” on page 114.
- Select the Promote DisplaySet from Living Master check box to populate the DTL scenario using the same display set that is applied to the Living Master scenario. When this option is not selected, the DTL scenario is either not given a display set or uses a display set applied to the DTL scenario by the user.
- Select the Overwrite DTL forecasting method check box to replace the forecasting method defined for the DTL scenario using the settings under the Forecast Method tab.

Note This option might be unavailable because settings from some forecast methods cannot be copied from the Living Master scenario.

Committing Changes Back to the Living Master Scenario

To commit changes made to a DTL scenario back to a Living Master scenario, choose Commit on the Living Master menu. This causes the changes from all series with the series status of Forecast Complete to be applied to the Living Master scenario.

Only the items that have been marked forecast complete are moved back to the Living Master scenario. If no items are marked forecast complete, you are prompted and notified.

The background of the page is a grayscale image of several interlocking gears. The gears are of different sizes and are positioned in a way that they appear to be meshing together. The lighting is soft, creating a professional and technical feel.

Chapter 5

Using the Engine Manager

This section describes how to configure the appearance of the QAD Demand Management user interface and other items.

Introduction **60**

Using the Engine Manager **60**

Managing Engine Usage Permissions **63**

Managing Engine Usage Domains **67**

Introduction

Use the Engine Manager to configure the appearance of the QAD Demand Management user interface and other items. The Engine Manager is used primarily by a system administrator to control the following:

- Web form views
- Data filters and exceptions
- Custom tasks
- Display sets
- ABC templates
- User data attribute access security
- User field data access security

A system administrator can also use the Engine Manager to define permissions and domains for the Living Master scenario, as well as to control form data views and field views.

Note Users can use the Engine Manager to create resources. The user who creates a resource can share or assign the resource to other users. Resources shared to a user are visible in the Engine Manager, but cannot be edited or shared by other users.

Using the Engine Manager

To use the Engine Manager:

- Select Applications|Engine Manager from the main menu. The Engine Manager window is displayed showing a list of available resources.




Fig. 5.1
Engine Manager

Note The Planning Portal and Engine Data Security resources are only available to a system administrator.

- Click the plus symbol (+) to expand the Planning Portal resource.
- Click the plus symbol (+) to expand the Data Filters resource.
- Click the plus symbol (+) to expand Engine Data Security resource.

The window now displays the options shown in Figure 5.2.




Fig. 5.2
Engine Manager,
Tree Tab

Table 5.1 describes the Engine Manager functions.

Table 5.1
Engine Manager
Functions

Branch	Resource Type	Description
Planing Portal	Data Views	Available only to the fxadmin user ID. Use to create, edit, share and remove Data views. Refer to the Planning Portal documentation for details.
	Field Views	Available only to fxadmin user ID. Use to create, edit, share, and remove Field views. Refer to the Planning Portal documentation for details.
Data Filters	Filters	Available to all users that have permission to manage filters. See “Managing Filters” on page 114.
	Exceptions	Available to all users that have permission to manage exceptions. See “Managing Exceptions” on page 122.
Custom Tasks	Stored Procedures	Available to all users that have permission to install custom tasks. See “Managing Custom Tasks” on page 107. This option is only available when a scenario is open in the Viewer.
Display Sets	Label Names	Available to all users that have permission to manage display sets. See “Managing Display Sets” on page 97.
ABC Templates	ABC Templates	Available to all users that have permission to manage ABC Templates. See Chapter 9, “Using ABC Analysis,” on page 131.
Engine Data Security	Data Permissions	Available to the fxadmin user ID. Use this option to manage permissions. See “Managing Engine Usage Permissions” on page 63.
	Data Domains	Available to fxadmin user ID. Use this option to manage domains. See “Managing Engine Usage Domains” on page 67.

Each item in this list has the same right-click options.

Create. lets the user create a new entry within the given option.

Edit. Lets the user modify an existing entry within the given option

Remove. Lets the user eliminate an entry within a given option.

Note For this to work, all users must be unassigned from the entry.

Review Assignments. Lets the user grant or revoke access to a given entry within an option.

Managing Engine Usage Permissions

This section describes how to define permissions that provide access to the series of Living Master scenarios. For details on Living Master scenarios, see Chapter 4, “Using the Living Master Scenario,” on page 53.

In order to retrieve data from Living Master scenarios, a user must be assigned to a single Engine Usage Permission entity.

Creating an Engine Usage Permission Entity

To create an engine usage permission entity:

- 1 Right-click the Data Permissions resource and choose Create. The Filtering Wizard window is displayed.

Fig. 5.3
Filtering Wizard

The screenshot shows the 'Filtering Wizard' dialog box. It has a title bar with a close button. The main area is divided into several sections: 'All Scenarios' with a checked 'Apply Filter to All Scenarios' checkbox and an 'Include Scenario Name' dropdown; 'Expression Builder' with a 'Field List' dropdown (showing 'Additional'), an equals sign operator, and a 'Values' dropdown, along with 'Show Values' and 'Add Expression' buttons; 'Logic Operators and parenthesis' with buttons for 'AND', 'OR', '(', and ')'; and a 'Filter expression' text area. Below these are 'Show Results', 'Parse Expression', and 'Clear Filter' buttons. At the bottom are 'Cancel' and 'Save' buttons. The bottom half of the dialog is a large empty list box.

- 2 Click Save. The Permission dialog box is displayed.




Fig. 5.4
Permission

- 3 Enter a name for the permission in the Name field and click Save.

Editing an Engine Usage Permission

To edit an engine usage permission:


- 1 Right-click the Permission that needs changing and choose Edit. The Permission window is displayed (Figure 5.4 on page 65).
- 2 Click the Edit button. The Filtering Wizard is displayed (Figure 5.3 on page 64).
- 3 Modify the filter and click Save.

Assigning Engine Usage Permission to a User

To assign an engine usage permission to a user:

- 1 Right-click the Permission that you want to assign and click Review Assignments.
The Assignments window is displayed.

Fig. 5.5
Assignments



- 2 Select the users from the Current users list and move them to the Assigned users list.

Note A user can only be assigned to one Data Permission. If you assign permission to a user with a previous permission assignment, the user loses the previous permission and acquires the new permission.

- 3 Click Save.

Removing an Engine Usage Permission

To remove an engine usage permission:

- 1 Right-click the permission resource and choose Remove. A delete confirmation message appears.
- 2 Click Yes.

Note Once you remove the permission, you must reassign the permission or else the user will not have access to the Living Master scenarios.

Managing Engine Usage Domains

Domains are administrator-defined access rights to the observation lines for the series that the user populates to the DTL scenario. Access rights define which observation lines a user can see and whether those observation lines can be changed.

Example A domain provides a user access to only the Value, FValue, Original, UserValue01, and ConvFactor01 observation lines, with Value and Original set as read-only.

Note A domain is used to protect the Living Master scenario data. Normal scenarios do not use domains. Domains are only required if a user has a DTL scenario.

Creating a Domain

To create a domain:

- 1 Right-click the Data Domains resource and choose Create. The Domain Setup window is displayed.




Fig. 5.6
Domain Setup,
Domain Properties
Tab

- 2 Enter a name and description for the domain.


Last Locked Period. Contains a predefined list of observations that you can specify as locked. You cannot change the opinion lines in periods that come from a locked period or before the locked period. You can set up a user's domain that does not contain the last locked period.

Note A domain does not have to contain a last locked period.

Click the list and select the appropriate option, and then click Apply.

- 3 Click the Scope tab. The Scope tab identifies the fields available for the user to view and edit.

Fig. 5.7
Domain Setup,
Scope Tab



- 4 Select the opinion lines required from the Available fields list and move them to the Selected fields list.
- 5 Right-click an opinion line in the Selected fields list and choose Set As and either Read Only or Read Write.
- 6 Click OK.

Editing a Domain

To edit a domain:

- 1 Right-click the domain resource you want to change and choose Edit. The Domain Setup window (Figure 5.6 on page 67) is displayed.
- 2 Modify the required fields and click OK.
- 3 Click the Scope tab (see Figure 5.7 on page 68).
- 4 Modify required fields and click OK.

Assigning a Domain to a User

To assign a domain to a user:


- 1 Right-click the domain resource and choose Review assignments. The Assignments (Figure 5.5 on page 66) window is displayed.
- 2 Select users from the Current users list and move them to the Assigned users list. Then click Save.

Note A user can be affected by only one domain.

Removing a Domain

To remove a domain:

- 1 Right-click the domain resource and choose Remove. A delete confirmation message is displayed.
- 2 Click Yes.

The background of the page is a grayscale image of several interlocking gears. The gears are of different sizes and are positioned in a way that they appear to be meshing together. The lighting creates highlights and shadows on the teeth of the gears, giving them a three-dimensional appearance. The overall tone is technical and industrial.

Chapter 6

Viewing Forecasts

This section describes the QAD Demand Management Viewer, and how to preview and transform your forecast.

Introduction **72**

Managing the Viewer **73**

Viewing Forecast Data **74**


Expanded View **90**

Editing a Forecast **94**

Introduction

The Viewer launches all the QAD Demand Management Engine's major features. The Viewer is the central location to interact with the forecast. Figure 6.1 illustrates the major features accessed from the Viewer.

Fig. 6.1
Viewer Features



The system administrator must configure security to prevent users from accessing information or capabilities that they should not have access to. Although the QAD Demand Management interface is similar for all users, different features are activated depending on the role that the feature plays within the forecasting process.

The Viewer lets you:

- View your forecast and the associated data in different forms.
- Transform or edit parts of the forecast.
- Launch the Engine Manager.
- Launch other applications including Base Unit, Batch Scheduler, the Planning Portal, and User Forum.

The Viewer is the primary method of viewing forecast results. You can enter input through adjustments and notes so that changes are tracked and the consensus forecast is visible to all stakeholders.

Managing the Viewer

Logging In As Administrator

To log in as the administrator:


- 1 Open the Viewer.
- 2 Enter fxadmin in the Username and Password fields. You should change the password later.
Note See Chapter 2, “Creating and Managing Users,” on page 9 for details on creating and modifying users.
- 3 Click OK.

Logging In As a User

To log in as a user:

- 1 Enter a valid username previously created by the system administrator.
- 2 Enter your password.
- 3 Click OK. The Viewer window is displayed.
Note To view the content of a scenario, you first define it using the Base Unit. For details, see Chapter 3, “Creating and Managing Scenarios,” on page 27.
- 4 Select a scenario from the Select Scenario drop-down list.

Fig. 6.2
Selecting a
Scenario



5 Click OK.

Viewing Forecast Data

The main Viewer window has four tabs that present different views of the data:

- Graph View Tab
- Data View Tab
- Demand Data Tab
- Series Information Tab




Fig. 6.3
View Tabs

Each View tab displays the forecast data in a different way. If data on a tab is inaccurate or outdated, you can manually change the data. The following sections describe each tab.

Graph View Tab

The Graph View tab is the default view when you open a scenario for the first time. By default, the Graph View displays the series in a line graph.

Changing a Line Graph to a Bar Graph

To change a line graph to a bar graph:

- 1 Right-click the line graph. A pop-up menu appears.
- 2 Choose Chart Type|Bar. A check mark is displayed next to Bar and the line graph changes to a bar graph.
- 3 Repeat the steps to return to a line graph.

Previewing Data in Graph View

This section explains different components of the graph view that let you preview your data.

Lines

Each graph can potentially have eight lines before incorporating multiple opinion lines. Each line represents different data series.

Blue	History of the series (default).
Pink	Represents manual adjustments made to the graph, and shows previous values when changes are made.
Yellow	Forecast of the series (default).
Dotted Red	Represents where history ends and the forecast begins (default).
Cyan	Represents the upper confidence limit.
Aqua	Represents the lower confidence limit.
Green	Represents safety stock.
Orange	Represents the level of error.
Tan	Represents the original forecast before changes.

Display Option

Figure 6.4 shows how to specify the range of data for the graph.




Fig. 6.4
Display Option

Normal View Tab

Normal View on the Graph View tab is the default tab that is displayed when the scenario is first opened. The Normal View tab lets you add and remove opinion lines when viewing the graph of a series. Select the Show Options check boxes for the line you want to display on the graph. Figure 6.5 shows the Normal View tab.




Fig. 6.5
Normal View

Use Show Options to select the lines you want to add to the graph.


The Normal View Tab contains the following Show Options check boxes:

- Value: Historical Values (Blue Line) set as default.
- FValue: Forecast Data (Yellow Line) set as default.
- Measurement: Adjustments (Pink Line).
- Error: Difference between Blue and Yellow Lines (Orange Line).
- Original: Original Forecast Values (Tan Line).
- Conf. Limits: Confidence Limits generated by the forecast calculation (Cyan and Aqua).

Year View Tab

Click the Year View tab to see the series in a year-to-year bar chart.

Fig. 6.6
Year View



	2003	2004	2005	Growth	Min	Max
January	30.00	30.00	34.09	13.63	30.00	34.09
February	47.00	37.00	46.19	24.84	37.00	47.00
March	48.00	40.00	50.78	26.95	40.00	50.78
April	56.00	49.00	63.07	28.71	49.00	63.07
May	38.00	40.00	49.78	24.45	38.00	49.78
June	31.00	32.00	40.42	26.31	31.00	40.42
July	37.00	35.00	45.82	30.91	35.00	45.82
August	35.00	38.00	48.19	26.82	35.00	48.19
September	36.00	44.00	53.66	21.95	36.00	53.66
October	40.00	43.00	54.39	26.49	40.00	54.39
November	45.00	47.00	60.56	28.85	45.00	60.56
December	41.00	49.00	61.11	24.71	41.00	61.11
Total	484.00	484.00	608.06	25.63	484.00	608.06

Show Multi-Series Tab

Click the series that you want to add. When the series is selected, it is displayed on the graph. To remove the series from the graph, clear the check box for that series.

Show Multiple Opinion Lines

You can view multiple types of data streams simultaneously as shown in Figure 6.7.





Fig. 6.7
Showing Multiple
Opinion Lines

Here multiple line items are displayed for the same type of data. You can display, view, and edit various data streams. See Chapter 7, “Customizing the Viewer,” on page 95.

Show Observations Button

Select Display|Observations from the menu when in the Graph View tab. The observations for the series display in a list next to the graph, as shown in Figure 6.8.

Fig. 6.8
Displaying
Observations



As you move the cursor over the points on the graph, the corresponding observation is highlighted.

Saving a Graph

To save a graph:

- 1 Choose File|Save Graph As from the menu. A Save dialog box is displayed.
- 2 Browse to the directory where you want to save the graph and click Save.

Adjusting the Forecast

Use the Adjustment tab to change the display of opinion lines currently viewed. This change also changes the data of the corresponding observation. First select the data you want to change, and then indicate how you want to change it.

The Adjust Frame area has two sections, Data Range and Adjust, as shown in Figure 6.9.




Fig. 6.9
Displaying
Observations

Data Range

From. Specify the first point on the graph where you want the change to take effect.

To. Specify the last point on the graph where you want the change to take effect.

Keep Range. When this option is selected, the data range selected remains regardless of the series being viewed.

Adjust

Select an option from the drop-down list box to indicate the adjustments to make to the data specified in the Date Range section.

- **By percentage:** Increases or decreases the data by the percentage amount specified. Enter a positive percentage amount to increase the values, or a negative amount to decrease the values.
- **By Amount:** Increases or decreases the data by the amount specified. Enter a positive amount to increase the values, or a negative amount to decrease the values.
- **To Exact Value:** Enter a new value for the data.
- **New Product:** Introduces a new product or retires a product that you want to include on the forecast. Indicate whether it is a new product or a retired product. When you select this option, additional fields become available for input (see following section).
- **Delete:** If selected, deletes the data you selected in Data Range.
- **Range Total:** Specify a total between the range of data specified in Data Range. QAD Demand Management totals the values in the data range, determines the percentage that each value represents of the total, and changes the values by that percentage.

New Product

If you specify an adjustment for a new product, you can enter additional values in the New Product frame, as shown in Figure 6.10.

Fig. 6.10
New Product
Adjustment

Curve Type. Select the curve type for the new or retired product. The curve types are default values.

Curve Direction. Indicate whether you are introducing or retiring a product. Introduction causes the curve direction to rise. Retirement causes the curve direction to fall.

Min/Max. Enter the minimum and maximum values for the selected curve's type and direction.

Changing Options Manually

To manually specify the data range you want to change:

- 1 Right-click the point of the line where you want the change to start. A pop-up menu appears.
- 2 Choose From.
- 3 Right-click the point of the same line where you want the change to end. The pop-up menu appears.
- 4 Choose To.
- 5 In the Adjust section on the Adjust Frame, choose the type of adjustment and click Apply.

QAD Demand Management applies the specified changes to the graph.

You can also specify the date range like this.

- 1
- Put the cursor on the line you want to move. A dot appears on the line.
- Note** To help you accurately differentiate what line your cursor is over, the dot appears as a different color when it is over different types of lines.
- 2
- Click and drag the line to the new point on the graph. A confirmation message appears displaying the observation and the value change you made.
- 3
- Click OK to confirm the change.
- 4
- To remove the change, click Undo on the Viewer tool bar. To return to the original forecast value after adjusting it, click Original on the Viewer toolbar.

Data View Tab

The Data View tab is a grid containing all the series, observations, and additional indexes used in the forecast (see Figure 6.11). If event modeling is used, event flags are included.




Fig. 6.11
Data View Tab

From the Data View, you can also review the data in the Year View.

Additional opinion lines that are added display for viewing and editing. You can customize this view. See Chapter 7, “Customizing the Viewer,” on page 95 for details.

Historical, Forecasted, Adjusted, and Confidence Limit values are displayed in columns under their respective category description.

Year View

Click the Display Year View button on the Data View to display the series data by year.

The column labels are the year (2003, 2004, 2005, to 200N), and the row labels depend on the cycle configurations that were previously set. To configure the cycle, use the Advanced Options on the Base Unit window. The cycle must be set before loading the data for a forecast in order to take effect. If the dates are not set, the default is the year view with cycle set to monthly.

Viewing Year-to-Year Data

To view year-to-year data:

- 1 Click the Data View tab and then click the Display Year View button.

Fig. 6.12
Year View

Indicates periods used for growth cycle difference

Bold numbers indicate the current forecast cycle

	2003	2004	2005	Growth	Min	Max
January	38.00	30.00	34.09	13.63	30.00	34.09
February	47.00	37.00	46.19	24.84	37.00	47.00
March	48.00	40.00	50.78	26.95	40.00	50.78
April	56.00	49.00	63.07	28.71	49.00	63.07
May	38.00	40.00	49.78	24.45	38.00	49.78
June	31.00	32.00	40.42	26.31	31.00	40.42
July	37.00	35.00	45.82	30.91	35.00	45.82
August	35.00	38.00	48.19	26.82	35.00	48.19
September	36.00	44.00	53.66	21.95	36.00	53.66
October	40.00	43.00	54.39	26.49	40.00	54.39
November	45.00	47.00	60.56	28.85	45.00	60.56
December	41.00	49.00	61.11	24.71	41.00	61.11
Total	484.00	484.00	608.06	25.63	484.00	608.06

- 2 Click the cell of historical data to make change.

Normal View. Returns to the default data view.

Growth Cycle Difference. Sets the number of cycles back to compare growth to the current forecast.

Growth. Compares the previous forecast year with the current forecast year.

Min/Max. Displays the minimum and maximum values between all years.

Total. Total value for the current year. You can edit this value.

Adjusting the Forecast

You can change historical and forecasted values on the Data View tab by clicking the appropriate field and typing a new value. As you change data, the columns automatically update.

In the Data View, you can adjust values by using the Adjustment tab. For details, see “Graph View Tab” on page 75.

Fig. 6.13
Adjusting Values in
Data View

Use adjustment tab to
change display values

Description	Value	FValue	X11Irregular
2003-01	30.00	32.19	
2003-02	47.00	43.80	
2003-03	48.00	45.98	
2003-04	56.00	54.47	
2003-05	38.00	40.33	
2003-06	31.00	32.46	
2003-07	37.00	36.26	
2003-08	35.00	35.34	
2003-09	36.00	36.66	
2003-10	40.00	38.84	
2003-11	45.00	43.36	
2003-12	41.00	41.22	
2004-01	30.00	26.01	
2004-02	37.00	40.28	
2004-03	40.00	40.79	
2004-04	49.00	47.41	
2004-05	40.00	33.32	
2004-06	32.00	28.48	
2004-07	35.00	34.76	
2004-08	38.00	34.00	
2004-09	44.00	36.86	
2004-10	43.00	43.21	
2004-11	47.00	50.16	
2004-12	49.00	47.37	
2005-01		34.09	
2005-02		46.19	
2005-03		50.78	
2005-04		63.07	
2005-05		49.78	
2005-06		40.42	
2005-07		45.82	
2005-08		48.19	
2005-09		53.66	
2005-10		54.39	
2005-11		60.56	
2005-12		61.11	

Demand Data Tab

This tab lets you perform customized calculations. Generally, you use these calculations to obtain the suggested times and quantities of re-orders so that you can develop an inventory replenishment schedule to satisfy forecasted demand.

Note You can customize this view. For details, see Chapter 7, “Customizing the Viewer,” on page 95.




Fig. 6.14
Demand Data

Unless other values are specified on the Data Demand tab, the custom calculation uses the following default values. The Demand Data Calculation is not activated by default. To use this feature, contact QAD Support Services.

Field Name	Default Value
Initial Backlog	0
Initial On-hand	0
Lead Time	Same as the Safety Stock calculation (default: 2)
History	Same as the Safety Stock calculation (default: 2)
Minimum Order	0
Max Order	0
Order Multiple	0
Roll Unsupported	False
SS Lead Time	
ServiceLevel	
Include Fcst Values	

Series Information Tab

This tab displays series attributes for data in rows and fields that were mapped (price, cost, customer) using the series properties. If your source data is in columns, you cannot import data to these fields; however, you can still change the data.

You can rename series properties based on your preference by typing in the available fields.

Fig. 6.15
Series Information
Tab

Property	Value
Price	318.25
Cost	300.00
Customer	OUTDOOR WORLD
SKU	1905
Other	
UserField1	
UserField2	
UserField3	
UserField4	
UserField5	
UserField6	
UserField7	

The following series options let you identify records that should not be included in the next forecast, or are ready to export to a report or another system:

Do Not Forecast. Select this option to exclude this series from the forecast calculation.

If grouping is active with the Before Forecast option selected, or if Regression is the selected forecast method, this option is not available.

Forecast Complete. Use this option to indicate that the forecast for this series is complete. It is not included in further forecast calculations. Do Not Forecast is automatically selected when Forecast Complete is activated.

If grouping is active with the Before Forecast option selected, or if Regression is the chosen forecast method, this option is not available.

Do Not Include in Reports. Select this option to exclude this series from any generated reports generated through the Reports tab in the Base Unit.

Ready to Export. Use this option to indicate that there are no more changes made to the series and it is ready to be exported to another file.

RollSales. Use this option to replenish inventory based on unsupported forecast amounts.

To make changes in the Series Properties tab, click the appropriate field and enter the updated value.

Note To change the User Field Labels in the Property column, see “Using Series Properties” on page 49.

User-Defined Series Properties

The Series Information window contains a group of tabs that let you change various data values.

- Product
- Customer
- Line Item
- Item Link
- Metric
- Statistics





Fig. 6.16
User-Defined
Series Properties

These fields can be renamed and displayed. See Chapter 7, “Customizing the Viewer,” on page 95.

Expanded View

The Expanded View button on the Graph View tab displays more information regarding the series. You can click the Expanded View button any time while using the Series Viewer.

Fig. 6.17
Series Viewer



Clicking the Expanded View button displays a different view of the data.

Fig. 6.18
Expanded View

The screenshot shows the 'Expanded View' window with the title '[WEST] [OUTDOOR WORLD] [1905] [MOWERS PUSH] [GRASS EAT]'. It contains a table with five columns: 'Series', 'Roll Unsupporte...', 'Forecast Adjusted', 'MethodName', and 'MAPE'. The table lists 15 series. The 'Expanded View button' is highlighted with a blue box in the toolbar.

Series	Roll Unsupporte...	Forecast Adjusted	MethodName	MAPE
[EAST] [BIG DEPOT] [1080] [GASS G...	Yes	No	Holt-Winters	5.44%
[EAST] [BIG DEPOT] [1352] [ELECTRI...	Yes	No	Holt-Winters	14.87%
[EAST] [DISCOUNT PALACE] [1080] [...	Yes	No	Holt-Winters	0.71%
[EAST] [DISCOUNT PALACE] [1905] [...	Yes	No	Holt-Winters	3.69%
[EAST] [EVAN'S HARDWARE] [1905] ...	Yes	No	Holt-Winters	1.71%
[EAST] [GARDEN SOURCE] [1352] [E...	Yes	No	Holt-Winters	0.45%
[EAST] [OUTDOOR WORLD] [1905] [...	Yes	No	Holt-Winters	0.30%
[WEST] [BIG DEPOT] [1080] [GASS G...	Yes	No	Holt-Winters	23.47%
[WEST] [BIG DEPOT] [1352] [ELECTR...	Yes	No	Holt-Winters	13.92%
[WEST] [BIG DEPOT] [1905] [RIDING ...	Yes	No	Holt-Winters	12.89%
[WEST] [DISCOUNT PALACE] [1080] [...	Yes	No	Holt-Winters	7.88%
[WEST] [DISCOUNT PALACE] [1905] [...	Yes	No	Holt-Winters	8.66%
[WEST] [EVAN'S HARDWARE] [1905] ...	Yes	No	Holt-Winters	5.70%

The remaining buttons in the Series Viewer section have the functions listed in the following table.

Button	Description
	Open/close expand view
	Last series
	Next series
	Previous series
	First series

- 1 Select a series and scroll to the right to see all fields.

Forecast Adjusted	MethodName	MAPE	RSquare	Forecast	DoNotReport	Total	ReadyToExport	RollSales
No	Holt-Winters	5.44%	82.32%	Include	No	13,157.98	No	Yes
No	Holt-Winters	14.87%	94.11%	Include	No	1,650.27	No	Yes
No	Holt-Winters	0.71%	89.30%	Include	No	1,433.03	No	Yes
No	Holt-Winters	3.69%	49.37%	Include	No	1,230.76	No	Yes
No	Holt-Winters	1.71%	30.64%	Include	No	729.75	No	Yes
No	Holt-Winters	0.45%	91.25%	Include	No	1,669.01	No	Yes
No	Holt-Winters	0.30%	92.19%	Include	No	2,195.61	No	Yes
No	Holt-Winters	23.47%	87.40%	Include	No	15,413.01	No	Yes

Fig. 6.19
Viewing All Fields


- 2 Click the label at the top of a column to sort that field.

 Select All Deserv. #. All				
Series	Roll Unsupported...	Forecast Adjusted	MethodName	MAPE
[EAST][OUTDOOR WORLD][1905][...	Yes	No	Holt-Winters	0.30%
[EAST][GARDEN SOURCE][1352][E...	Yes	No	Holt-Winters	0.45%
[EAST][DISCOUNT PALACE][1080][...	Yes	No	Holt-Winters	0.71%
[EAST][EVAN'S HARDWARE][1905][...	Yes	No	Holt-Winters	1.71%
[WEST][GARDEN SOURCE][1352][...	Yes	No	Holt-Winters	2.71%
[EAST][DISCOUNT PALACE][1905][...	Yes	No	Holt-Winters	3.69%
[]	Yes	No	No default descr...	5.23%
[]	Yes	No	No default descr...	5.23%
[EAST][BIG DEPOT][1080][GASS G...	Yes	No	Holt-Winters	5.44%

Fig. 6.20
Sorting By Field

- 3 Right-click a series and select Customize Extended View from the list. This lets you remove or add fields to the view, based on how you prefer to see the data.
- 4 Select the field from the Available list and click the single right arrow to move the field to the Selected list.

Fig. 6.21
Selecting an
Available Field



5 Click OK.

The following functions are also available within the Expanded View:

Select All/Deselect All. Use to select all of the series or none.

Multiple Selection. Hold down the Ctrl key and click each series you want to select. Each selected series is highlighted.

Sort Series By. Sort by field name, or by ascending or descending order.

The three tabs—Adjustment, Same As, and Status Information—contain options for adjusting, changing, and marking the series.

Making Adjustments in Expanded View

You can adjust a series using Expanded View the same way as you would if using the Adjustments tab; see “Adjusting the Forecast” on page 80. Expanded View lets you adjust multiple series at one time.

You can adjust the series by Date Range as well as make other adjustments. Use the Adjustments tab in Expanded View to make mass adjustments to a series. Use the Ctrl key to highlight multiple series. All adjustments are marked in the series column with an Adjustment icon.

In the Same As tab in the Expanded View, you can adjust or apply Same As to the series data. Same As is used to copy the history, seasonality, or both from one series and to another series.

Status Properties

The Status/Properties tab in the Expanded View lets you mark series with status options, and view and update User-Defined Field Properties. These selections are the same as those on the Demand Data tab. See “Data View Tab” on page 83.

Modifying User-Defined Field

To modify a user-defined field in Expanded View:

- 1 Click the Status/Properties tab in the Expanded View of the Viewer.
- 2 Click the User Fields button.

A Set User Fields dialog box appears. This lets you change or enter information.

Finding a Series

You can use Find to search Normal View or Extended View modes for a specific series.

- 1 Choose Edit|Find from the View menu or press the Ctrl+F keys.
- 2 Enter the Series Label that you want to find and click OK.




Fig. 6.22
Finding a Series

If the series exists in the scenario, the cursor moves to that series.

Fig. 6.23
Selecting a Series

Series	Roll Unsupporte...
[EAST] [OUTDOOR WORLD] [1905] [...]	Yes
[EAST] [GARDEN SOURCE] [1352] [E...]	Yes
[EAST] [DISCOUNT PALACE] [1080] [...]	Yes
[EAST] [EVAN'S HARDWARE] [1905] ...	Yes
[WEST] [GARDEN SOURCE] [1352] [...]	Yes
[EAST] [DISCOUNT PALACE] [1905] [...]	Yes
[]	Yes
[] []	Yes
[EAST] [BIG DEPOT] [1080] [GASS G...	Yes
[WEST] [EVAN'S HARDWARE] [1905]...	Yes
[WEST] [DISCOUNT PALACE] [1080] [...]	Yes
[WEST] [DISCOUNT PALACE] [1905] [...]	Yes
[WEST] [OUTDOOR WORLD] [1905] [...]	Yes
[WEST] [BIG DEPOT] [1905] [RIDING ...]	Yes
[WEST] [BIG DEPOT] [1352] [ELECTR...	Yes
[EAST] [BIG DEPOT] [1352] [ELECTRI...	Yes
[WEST] [BIG DEPOT] [1080] [GASS G...	Yes

Editing a Forecast

You can change forecast values, opinion line values, or series properties in several ways:


Graph View tab. Use the mouse to change values in the Graph View Tab by dragging. See “Graph View Tab” on page 75.

Data View and Demand tabs. The Data View and Demand tabs let you edit values as a spreadsheet; see “Data View Tab” on page 83 and “Demand Data Tab” on page 86.

Adjustment and Graph View tabs. Change the opinion values for a single series or block of series simultaneously. Also, you can define the affected range of periods. See “Making Adjustments in Expanded View” on page 92 and “Graph View Tab” on page 75.

Series Properties and Expanded View tabs. Edit series properties manually in the Series Properties or Expanded View tabs. See “Series Information Tab” on page 87 and “Expanded View” on page 90.

Batch Scheduler (Massive Series Update, Massive Observations Update), Events, Conversion Factors. See the relevant sections in other chapters for details.



Chapter 7

Customizing the Viewer

This section describes how to customize labels and descriptions in the QAD Demand Management Viewer.

Introduction **96**

Managing Display Sets **97**

Customizing the Graph View **101**

Customizing the Demand Data View **102**

Customizing the Expanded View **103**

Customizing the Viewer Framework **104**

Using Notes to Track Assumptions **104**

Managing Custom Tasks **107**

Introduction

The forecasting process lets you view information using familiar terms and display information on demand. The Viewer lets you adjust and interact with forecasts. A key element in this process is the customizing of labels and field descriptions within the Viewer if the default labels and field descriptions provided do not meet your business requirements.

This chapter contains the following sections:

- Managing Display Sets
- Customizing the Graph View
- Customizing the Demand Data View
- Customizing the Expanded View
- Customizing the Viewer Framework
- Using Notes to Track Assumptions
- Managing Custom Tasks

Display Sets

Using the Display Set function, you can:

- Specify a name (alias) for each opinion line or series field
- Hide an opinion line
- Specify the number of decimal positions

A display set is a function that allows users to rename the defined series and Observation Field names. It is an Engine Manager function that you share with other users and apply to a scenario or group of scenarios.

When you open a scenario affected by a display set, the Viewer displays the aliases instead of the default names.

A display set is composed of the following:

Name. Name of the current display set.

Description. Description of the current display set.

Scenarios affected list. Names of the scenarios that are affected by the current display set.

Series Fields definition. Defines the number of decimal points, visibility of the fields at series level, and the alias of each series field.

Observations Fields definition. Defines the number of decimal points, visibility of the fields at Observations level, and the alias of each observation field (opinion line).

Managing Display Sets

Creating a Display Set

To create a display set:

- 1 Select Applications|Engine Manager from the Viewer menu. The Engine Manager window is displayed.




Fig. 7.1
Engine Manager

- 2 Right-click the DisplaySets resource and select Create. The DisplaySet Setup window is displayed, as shown in Figure 7.2.

Fig. 7.2
DisplaySet Setup

DisplaySet Setup

General properties | Series Fields | Observations Fields

Name:

Description:


Scenarios affected:

Apply	Scenario	Current DS
☐	Living Master	none
☐	Living Master-1 (CostCenter)	none
☐	Living Master-2 (BusinessPa	none
☐	Scenario 152	none

OK Apply Cancel

- 3 In the General properties tab, specify the following data:
 - a Enter a name for the display set in the Name field.
 - b Enter a description for the display set in the Description field.
 - c Select the check box in the Apply column that corresponds to the scenario you want to select.
- 4 Click the Series Fields tab. The window shown in Figure 7.3 is displayed.

Fig. 7.3
Series Field Tab



- a Select an entry in the Alias column and enter a new name.
 - b Click a field in the Active column to deselect the field. Select other fields one at a time.
 - c Locate a field with a number in the Dec. Points column. Click the drop-down list and select the number of decimal positions required. You cannot specify decimal positions for string fields.
- 5 Click OK to save the display set.

The next time an affected scenario is opened, the display set is applied to the scenario and the default series and observation field names are replaced. For example, FValue is replaced by My Sales.

Editing a Display Set

To edit a display set:

- 1 Right-click the display set you want to edit, and choose Edit. The DisplaySet Setup window is displayed (Figure 7.2 on page 98).
- 2 Change the display set as required and click OK.


Sharing a Display Set

The owner of a display set can share it with other users. Once shared, other users can apply the display set to their scenario but cannot modify it.

Note If you populate data from the Living Master scenario, by default the display set from the Living Master is applied to the DTL scenario and the display set is shared with the user. See “Populating a DTL Scenario” on page 55.

- 1 Right-click the display set that you want to share and select Review Assignments. The Assignments window is displayed.

Fig. 7.4
Assignments



- 2 Select the user from the Current users field, move the user to the Assigned users field, and then click Save.

Removing a Display Set

To remove a display set:

- 1 Right-click the display set you want to remove and select Remove. A delete confirmation message is displayed.
- 2 Click Yes to delete. If the display set currently affects scenarios, a message displays which scenarios are affected and prompts you for confirmation.

Customizing the Graph View

You can customize which opinion lines you want to use in each view. The Graph View tab has eight display lines. You can add more opinion lines to this graph.

Note Changes made (adding/removing lines) in the Graph View tab also affect the Data View but not the Demand Data tab.

To add or remove additional opinion lines:

- 1 Select Display|Additional Opinion Lines from the Viewer menu. The General Selection window is displayed.




Fig. 7.5
General Selection

- 2 To add an opinion line, select a field in the Available list and click the single right arrow to move it to the Selected list. To remove an opinion line, select a field in the Selected list and click the single left arrow to move it to the Available list.
- 3 Click OK. The graph is displayed with the selected opinion lines added to the Graph View tab.
- 4 Click the Data View tab. The data view window is displayed showing the additional opinion lines.

Customizing the Demand Data View

The Viewer lets you customize the Demand Data View. You can add opinion lines and specify the order of those lines.

- 1 Click the Demand Data tab in the Viewer.
- 2 Right-click in the view and choose Customize DRP. The General Selection window is displayed (Figure 7.5 on page 101).
- 3 Select the field in the Available list and click the arrow to move the field to the Selected list. Click OK. Be sure to select the fields in the order that you want the opinion lines to display.

The following window displays the selected fields as columns.




Fig. 7.6
Selected Fields as
Columns

Customizing the Expanded View

You can also add or remove fields in the Expanded View.

- 1 Right-click a row in the expanded view and select Customize Expanded View from the menu. The General Selection window is displayed (Figure 7.5 on page 101).
- 2 Select the field in the Available list and click the single right arrow to move the field to the Selected list, and then click OK. Be sure to select the fields in the order that you want the opinion lines to display. The Viewer displays the selected fields.

Note The Forecast Adjusted and Notes fields are not built-in fields for the series. The fields are displayed for information purposes only and always display, regardless of how Expanded View is customized.

Customizing the Viewer Framework

You can customize the toolbar and menu bar of the Viewer.

Select Display|DME Bar from the Viewer menu. The Viewer Framework icons are removed from the window. You can also right-click the gray space on the Viewer toolbar, choose Customize, and then select the icons you want.

Using Notes to Track Assumptions

You can add customized notes to a specific series of the Viewer. Also, you can add different notes for each period of that series to track each series and each period of that series.


You can add notes at the series level and observation level. You can edit notes directly in the Viewer.

Inserting a Note at Series Level

To insert a note at series level:

- 1 Put the cursor in the text area at the bottom of the QAD Demand Management application window. Make sure the Series tab is selected.
- 2 Enter a note for the series in the field. A note icon appears next to the selected series, as shown in Figure 7.7.

Fig. 7.7
Entering a Note for
a Series




- 3 Save the changes by switching to another series or click the Save User Changes icon

Inserting a Note at Observation Level

To insert a note at observation level:

- 1 Right-click the Series tab and select Jump to the note.
- 2 Select a period for the note.


Fig. 7.8
Selecting a Period
for a Note



The Series tab displays the period selected.

- 3 Enter a note for this period of the series.

Fig. 7.9
Entering a Note at
Observation Level




- 4 The note icon displays on the Series List and the Jump to Note option. Save the changes by switching to another series or click the Save User Changes icon.

Sending Notes to Engine User

To enable other users to track notes over a series, you send a note directly from Viewer using the Engine E-mail Service.

- 1 Right-click one of the Notes tabs and select Send Item Note. The window shown in Figure 7.10 is displayed.

Fig. 7.10
Send Notes to
Users



- 2 Select the users you want to send the e-mail to in the Select Recipients list.
The Viewer automatically populates the Subject field with the series information. The Viewer automatically populates the Content area with the note information in that series.
- 3 Edit the e-mail subject and content in the Subject and Content fields.
- 4 Click Send.

Note E-mail notification must be configured to use this function. See “Setting Up E-Mail Service” on page 22.

Managing Custom Tasks

Custom tasks are special procedures designed to enhance QAD Demand Management functionality for your particular requirements. An example of a custom task could be special calculations like an Order Replenishment Calculation.

You activate custom tasks from the Viewer. Using the Custom Tasks function in the Engine Manager, you can create, delete, edit, and share custom tasks. This function operates in a similar manner as the other functions.

Creating a Custom Task

To create a custom task:

- 1 Select Applications|Engine Manager from the Viewer menu. The Engine Manager window is displayed.
- 2 Right-click the Custom Tasks resource and select Create. The Add/Edit Custom Task window is displayed.




Fig. 7.11
Add/Edit Custom Task

- 3 Enter the description and the stored procedure name in the respective fields and click OK. If the stored procedure does not exist in the Engine database, an error message is displayed.

To use a stored procedure in a Custom Task, the following rules apply:

- The first parameter must be @ScnID with Integer type. The Engine passes the scenario ID to this parameter automatically when you run the custom task.

- The second parameter must be @UserID with Integer type. The Engine passes the user ID to this parameter automatically when you run the custom task.
 - If you want Engine to pass the series ID of the selected series, add the parameter @SeriesID with type Integer to your stored procedure. This parameter can be any where except the first and the second position in the definition of the stored procedure.
- 4 If there are additional parameters in the stored procedure besides @ScnID and @UserID, you must provide the default values for those parameters.

Fig. 7.12
Providing
Parameters for
Custom Task

Please Provide the Parameters for Custom Task						
	Name	Type	Description/Comment	Value	Value Type	
1	@ScnID	Numeric				
2	@UserID	Numeric				
3	@par1	Numeric		First Value	LIST	Edit List
4	@SeriesID	Numeric		0	INPUT	
5	@par2	String			QUERY	Edit Query

For parameters other than @ScnID, @UserID, and @SeriesID, you can enter a description to explain the parameters, and set initial values and value lists for these parameters.

The Value Type affects how users enter values for parameters when running the custom task.

INPUT: Users must type in the value at runtime. You must enter the initial value in the Value column for that parameter.

LIST: Users select a value from a drop-down list of supplied values. If you select LIST as the Value Type for a parameter, click the Edit List button next to it to display the Task Parameter Value List window.




Fig. 7.13
Task Parameter
Value List

Enter a list of values and text in the Value and Text columns. The Value column contains the values used to pass to the stored procedure when the custom task runs. The Text column contains the labels for their corresponding values, which display in the drop-down list when you run the custom task. Click OK.

QUERY: Users select a value from a list derived from a custom query. If you select QUERY as the Value Type for a parameter, click the Edit Query button of the parameter to display the Edit Custom Task Value List Query window.




Fig. 7.14
Edit Custom Task
Value List Query

Enter a query that returns the list of pairs of value and text. Make sure [Value] is unique in the query results. Click OK.

- 5 Select a value from the drop-down list or enter a value as the initial value for each parameter except @ScnID, @UserID, and @SeriesID. Click Accept.

Deleting a Custom Task

To delete a custom task:

- 1 Select Applications|Engine Manager from the Viewer menu.
- 2 Select Custom Tasks. The custom tasks are displayed.
- 3 Right-click the custom task you want to delete and choose Remove. A confirmation window is displayed.
- 4 Select Yes.

Editing a Custom Task

To edit a custom task:

- 1 Select Applications|Engine Manager from the Viewer menu.
- 2 Select Custom Tasks. The custom tasks are displayed.
- 3 Right-click the custom task and choose Edit.
- 4 Modify the custom task and click OK.

Sharing a Custom Task

The owner of a custom task can share it to other users. Once shared, other users can run the task in their own scenario but cannot modify it.

To share a custom task:

- 1 Right-click the custom task that you want to share and select Review Assignments. The Assignments window is displayed (Figure 7.4 on page 100).
- 2 Select the user from the Current users list, move the user to the Assigned users list, and then click Save.

Running a Custom Task

To run a custom task:

- 1 Select Tools|Custom Tasks from the Viewer menu. The Custom Tasks window is displayed.




Fig. 7.15
Select Custom Task

- 2 Select a custom task from the Available custom tasks drop-down list and click Run.
- 3 The system displays a window prompting for task parameter values. For each parameter, enter a value, or click the Value field and select the value from the drop-down list.

Fig. 7.16
Providing
Parameters for
Custom Task

Please Provide the Parameters for Custom Task

	Name	Type	Description/Comment	Value
1	@ScnID	Numeric		
2	@UserID	Numeric		
3	@par1	Numeric		1.00
4	@SeriesID	Numeric		239
5	@par2	String		1

☒ Save For Current Scenario


☐ Save As Default

Accept

Cancel

Note Values can be used for other scenarios. To do this, select the Save as Default check box.

- 4 Click Accept. When the custom task finishes, a completion confirmation window is displayed, allowing you to view the results.
- 5 Click OK.
- 6 Click Close in the Custom Tasks window.

The background of the page is a grayscale image of several interlocking gears. The gears are of different sizes and are positioned in a way that creates a sense of depth and mechanical complexity. The lighting is soft, highlighting the teeth of the gears and creating subtle shadows.

Chapter 8

Using Filters and Exceptions

This section describes using filters and exceptions to view QAD Demand Management forecasts.

Introduction **114**


Managing Filters **114**

Managing Exceptions **122**

Introduction

Filters and exceptions provide two key methods for managing a forecast. Filters let you select items based on series attributes. Exceptions let you select items based on numerical information that is stored in observations. Using filters and exceptions lets you manage forecasts effectively. Figure 8.1 shows how filters and exceptions interact.

Fig. 8.1
Interactions
between Filters and
Exceptions



Exceptions can also be used to update attributes in a series. In this way, when an exception is run, the results of the exception can be stored so that they can be accessed later using a filter.

In this case, the exception marks the items by setting the series attributes. Once the attribute has been set by the exception, a filter is then used to select the series. If the attribute is made editable, users can replace the flag with a space, which then removes it from the filter so that the item is no longer visible. This lets administrator to set exceptions and users to clear them.

Managing Filters

Filters let you specify particular groups or sections of the data table for the Engine to use in the forecast. Filters let you perform on-the-fly analysis on specific pieces of data. To use a filter you first create it, then launch it.

A filter contains a filter expression. Before you create a filter, you must understand how filter expressions are defined.

Using Filter Expressions

A filter expression is a combination of fields, operators, and value fields that indicate to the Engine specific data to use in the view.

A filter expression has the format of:

```
Field Name      Logical Operator      Data Value
```

Example 1

In my source data, I want to display only the items that BIG MART purchases in the Midwest.

This filter expression would look like this:

```
Customer="BIG MART" AND Region="Midwest"
```

In this example, the Field Name is Customer, the operator is =, and the Data Value is BIG MART.

Example 2

Your source data lists all your employees, their departments, and start dates. You want to know how many people started in the Marketing and Sales departments in 1999. You can build a filter expression that tells the Engine to select dates in 1999 and employees in the Sales or Marketing departments.

This filter expression would look like this:

```
Start Date >= '%1/1/99%' AND Start Date <= '%12/31/ 1999%' AND  
(Dept = Sales OR Dept = Market )
```

Creating a Filter

To create a filter:

- 1 Select Applications|Engine Manager from the Viewer window with an open scenario. The Engine Manager window is displayed.
- 2 Expand the Data Filters resource list.
- 3 Right-click the Filters resource and select Create. The Filtering Wizard is displayed.

Fig. 8.2
Filtering Wizard

- 4 Select a specific scenario to filter, or choose the Apply Filter to All Scenarios option. In this case, data from all scenarios that the user has access to is available.
 - 5 Select the name of the field for the first part of the expression from the Field List drop-down list.
- Note** The values available depend on the data source loaded by the Base Unit.
- 6 Click the Operator (unlabeled in the Filtering Wizard window) drop-down list and select the operator. Notice the (=) equal to sign in the next field.

The operators are:

=	Equal to	<=	Less than or Equal to
<>	Not equal to	Like	Similar

>	Greater than	Is not Null	Contains data
<	Less than	Is Null	Has no data value
>=	Greater than or Equal to		

- 7 Click the Show Values button. The values that coincide with the selected property populate the Values drop-down list.
- 8 Select the appropriate value from the Values drop-down list.
- 9 Click the Add Expression button. The filter expression you created appears in the Filter expression area.
- 10 Click the appropriate AND or OR button in the Logic Operators and parenthesis section. Use the parentheses where appropriate in the filter expression or when the filter expression requires multiple parameters.
- 11 Click the Parse Expression button to ensure that the filter expression is valid. If the filter expression is not valid, an error message is displayed.
- 12 Click the Show Results button when you are finished building the expression. The values specified display in the results area of the Filtering Wizard (Figure 8.2 on page 116).
- 13 Click the Accept button or edit the results if necessary.
- 14 Specify a name for the new filter in the Name field.




Fig. 8.3
Specifying a Filter

15 Specify a description in the Description field.

16 Click Save. The new filter is added.


Launching a Filter

You can use a filter on any scenario. The Filter Launcher executes the selected filter by matching the data to the criteria defined in the expression. After the filter executes, the application displays the results of the filter expression.

To launch a filter:

- 1** Select Tools|Filter Launcher from the Viewer menu. The Filter Launcher window is displayed.

Fig. 8.4
Filter Launcher



- 2** Select the filter from the Available Filters drop-down list. The results display in the current open scenario.

Editing a Filter

- 1** Open a scenario in the Viewer.
- 2** Select Applications|Engine Manager from the menu.
- 3** Expand the Filters resource.
- 4** Right-click the existing Filters resource and choose Edit.
- 5** Change the filter expression as required.




Fig. 8.5
Editing a Filter

- 6 Click the Parse Expression button to make sure the filter expression is valid.
- 7 Click Save to update the filter.

Editing a Filter Before Running

To make quick changes without modifying a filter's original expression, you can edit a filter just before it is run using the Viewer.

- 1 Open a scenario in the Viewer and click the Filter icon. The Filter Launcher dialog displays (Figure 8.4 on page 118).
- 2 Select the filter from the Available Filters drop-down list.
- 3 Select the Edit on fly check box.
- 4 Click Run. The Filtering Wizard displays (Figure 8.2 on page 116).
The changes are not saved. The updated filter expression is only used for this execution now. The original filter expression is saved by the Viewer.
- 5 Modify the filter expression and click Accept. The filter is applied to the scenario.


Note If you did not choose Edit on fly and click Run, the filter is executed immediately.

When the filter is run, the Viewer displays only the items that match the criteria.

Example

Assume that you have a filter that displays the items only where the customer is DISCOUNT PALACE. Figure 8.6 displays an example of the previous filter.

Fig. 8.6
Previous Filter



Sharing a Filter

Other users might want to use the filter. A system administrator can grant access rights to a user in order to share scenarios. A user also must have rights to access a previously created filter.

To share a filter:

- 1 Select Applications|Engine Manager from the Viewer menu.
- 2 Click the plus sign (+) to expand the Filters resource.
- 3 Right-click the defined filter and choose Review Assignments. The Assignments window is displayed.




Fig. 8.7
Assignments

- 4 Select the users from the Current users list and move them to the Assigned users list. This lets other users share the filter.
- 5 Click Save. Selected users can now view and use your filter against their scenarios.

Removing a Filter

To remove a filter:

- 1 Right-click the filter resource and select Remove. A delete confirmation message is displayed.
- 2 Click Yes.

Deactivating a Filter

You can use one filter with many scenarios, but a scenario is affected by only one filter at a time. To apply another filter, the original filter must be deactivated in order to return the scenario to the original state.

To deactivate a filter:

- 1 Select Tools|Deactivate Exception/Filter from the Viewer menu. A deactivation confirmation message is displayed.
- 2 Click Yes.

Managing Exceptions


Using an exception entails applying a business rule to filters out items that meet certain criteria. Exceptions capture abnormal items before the data values reach a certain threshold. You define exceptions before each forecast cycle. Identifying and removing exceptions lets you focus on a specific forecast area.

Creating an Exception Template

Before processing exceptions, you must create an exception template.

- 1 Select Applications|Engine Manager from the Viewer menu. The Engine Manager window is displayed.
- 2 Expand the Data Filters resource.
- 3 Right-click Exceptions and choose Create. The Select Source Opinion Line window is displayed.

Fig. 8.8
Select Source
Opinion Line Using
Dynamic Date
Range



- 4 Select a base opinion from the drop-down list. You can select only one base opinion line. This is the data that the exception calculation compares to the criteria.

Dynamic Date Range. This option automatically selects the historical and forecast dates based on the range specified. Figure 8.8 shows the window with this option selected.

Fixed Date Range. This option lets you specify a fixed date range using From and To dates. The window shown in Figure 8.9 displays if you select the Fixed option.

- 5 Click Next. The window shown in Figure 8.9 is displayed.

Select the target opinion line

With what do you want to compare?

☐ I want to compare with a specific value

☒ I want to compare with other opinion line

Opinion Line

ActShipments

Select the period offset for target line

First period in base line

0

Cancel < Back Next > Finish

Fig. 8.9
Select Target
Opinion Line

There are two ways to perform the comparison:

Specific Value. This option compares the values to a specific threshold. Use this option to compare the base opinion against the value in the Actual Value field. Enter the Actual Value in the field and go to step 8.

Fig. 8.10
Select Target
Opinion Line Using
Specific Value

Select the target opinion line

With what do you want to compare?

☒ I want to compare with a specific value

☐ I want to compare with other opinion line

Actual Value

Cancel < Back Next > Finish

Opinion Line. This option compares one opinion line against another opinion line. The window shown in Figure 8.11 displays when this option is selected. This option lets you compare historical data with forecast data.

Fig. 8.11
Select Target
Opinion Line Using
Opinion Line

Select the target opinion line

With what do you want to compare?

☐ I want to compare with a specific value

☒ I want to compare with other opinion line

Opinion Line
ActShipments

Select the period offset for target line

First period in base line
0

Cancel < Back Next > Finish

- a Select an Opinion Line from the drop-down list.
- b You can adjust the starting period by dragging the sliding bar to the left or right.

Note You cannot change the From and To drop-down list values; these two lists are for information only. If you slide the First period in base line to 3 and compare the historical data to the forecast data, the time period moves forward 3 periods.

- 6 Click Next to choose the comparison method.

Fig. 8.12
Condition
Attributes

Choose a comparison method:

Variance. The difference between the values and the thresholds or another opinion line. The difference can be positive or negative.

Absolute Variance. The absolute difference between the values and the thresholds or another opinion line. That difference is always positive.

% Variance. The percentage of the difference between the values and the threshold or another opinion line. The equation is:

$$(a - b)/b \times 100\%$$

Where a is the threshold and b is the value of the base line.

Absolute % Variance. The absolute percentage of the difference between the values and the threshold or another opinion line.

- 7 Select the exception level to define a filter expression in which you can compare the Comparison Method against the expression. To specify one level only, select the That's It option.

- 8 Click Next to define buckets. The Bucket Definition window displays, as shown in Figure 8.13.

Fig. 8.13
Bucket Definition

Quarterly

Option	History	Forecast
2	432.7	1,365
3		381
4		250
5	554.00	275
6	665.00	406
7	4,623.38	4,500
8	544.00	795
9	444.00	770
0	656.00	595

Bucket

Bucket definition

Bucket length: 1

Method within Bucket: Sum

Method among buckets: Sum

Cancel < Back Next > Finish

A bucket is a logical unit that groups periods together. For example, if the data is monthly and you specify a bucket length of 3, data is grouped into three quarters and a quarter-by-quarter comparison is performed.

The four aggregate methods for grouping the buckets are:

Sum. Adds the values.

Average. Averages the values.

Min. Uses the minimum.

Max. Uses the maximum.

Exceptions are calculated only on the item-level for all the periods, as shown in this example.

Example

Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep
100	120	100	80	90	120	50	60	70
Within the buckets			Within the buckets			Within the buckets		

Assume that the bucket length is 3, the threshold is 100, and the aggregation method within the bucket is Max. The system first divides and groups in the periods based on the aggregation methods defined. The results are the following:

Quarter 1	Quarter 2	Quarter 3
120	120	70

At this point, the system cannot determine if this series has exceptions or not, because of the first two quarters. The first two quarters are above the threshold; however, the third quarter is below the threshold.

The system applies the method defined in the Method Among Buckets drop-down list to the quarters before doing the final calculation.

If the method is Max, the final result is 120, which is over the threshold and an exception is reported.

9 Click Next.




Fig. 8.14
Storing an
Exception in the
Output Field

You can store the exception value or the name of the exception in the Output Field.

After processing the exceptions, the QAD Demand Management automatically creates and executes the filter that retrieves the items that fulfill the exception criteria.

Select the Reset Output Field option to clear your output field when running the exception multiple times.

- 10** Click Next.
- 11** Enter the name and description (optional) in the appropriate fields and click Finish.

Sharing an Exception

To share an exception:

- 1** Right-click the exception resource and choose Review Assignments. The Assignments window displays (Figure 8.7 on page 121).
- 2** Select users from the Current users list, and click the right arrow button to move the users to the Assigned users list.
- 3** Click Save.

Deleting an Exception

To delete an exception, right-click the exception resource and choose Delete.

Processing an Exception

To process an exception:

- 1 Select Tools|Exceptions Launcher from the Viewer menu. The Exception Launcher is displayed.




Fig. 8.15
Exception
Launcher

- 2 Select the exception from the Available exceptions drop-down list.
- 3 Select the scenario from the Scenario drop-down list.


The option Whole Tree applies the exception to all levels of the applied hierarchy.

There are two options for displaying the result.

Limit Series List. Displays only the items that meet exception criteria.

Flag Exception. Displays all items but places a flag on items meeting exception criteria.

- 4 Click Run. The results of the exception calculation in the Series panel and the selected output field display.

The background of the page features a grayscale image of several interlocking gears. The gears are of different sizes and are positioned in a way that creates a sense of depth and mechanical complexity. The lighting is soft, highlighting the teeth of the gears and casting gentle shadows.

Chapter 9

Using ABC Analysis

This section describes how to use ABC analysis to group sets of data in QAD Demand Management

<i>Introduction</i>	132
<i>Creating an ABC Template</i>	133
<i>Using an ABC Template</i>	135
<i>Analyzing Results</i>	136
<i>Sharing an ABC Template</i>	137

Introduction

ABC Analysis provides a convenient way to group a set of data into buckets based on the value of forecasted items. Once the items are tagged and coded, items can be filtered to focus on those that have the most significant business impact. Table 9.1 shows a methodology for ABC coding items based on forecast error.

Table 9.1
Sample ABC
Analysis

Date	Shirts		Pants		Shoes	
	Actual	Forecast	Actual	Forecast	Actual	Forecast
2003-01	40	44	51	49	1,743	1,764
2003-02	60	62	91	89	1,863	1,868
2003-03	20	21	39	39	1,394	1,447
2003-04	10	10	10	10	14	14
2003-05	10	11	17	16	65	67
2003-06	40	42	64	58	3,695	3,849
2003-07	50	54	69	68	390	396
2003-08	70	75	94	89	1,859	1,884
2003-09	80	95	96	87	6,721	6,752
2003-10	80	84	115	104	1,531	1,552
2003-11	120	127	124	115	2,907	2,960
2003-12	50	59	78	78	6,065	6,085
Total	630	684	849	802	28,247	28,639
Error		(54)		47		(391)
Abs Pct Error		8.53%		5.53%		1.39%
Ranking On Error		B		C		A
Ranking On Abs Pct Error		A		B		C

This example illustrates that the method by which the ABC analysis assigns its values significantly impacts the items that would receive the largest amount of review.

You can view the categorized data by total value and percentage of total value over a specified time horizon. For example, the items can be categorized as the top percentage of products or A-items, the second percentage of items or B-items, the third percentage of items (C-items) or the bottom dollar percentage.

The result of the ABC calculation is a field populated with the appropriate ranking, so you can create a filter to focus on a particular group of ranked products that require the most attention.

Creating an ABC Template

To create an ABC template:

- 1 Select Applications|Engine Manager from the Viewer menu.
- 2 Right-click the ABC Templates resource and choose Create. The ABC Analysis window is displayed.

Fig. 9.1
ABC Analysis

- 3 Enter a name and description in the appropriate fields. The description is optional.
- 4 Select the data field from the Data Field drop-down list. This field is the basis of the ABC calculation.
- 5 Select a conversion or (no conversion) from the Conversion drop-down list. If conversion factors are loaded—for example, price, units, tons, pallets, or others—select the appropriate opinion line.
Note If you select a conversion factor other than (no conversion), the ABC template uses the value (Data Field * Conversion) to do the ABC calculation instead of the value in Data Field.
- 6 Select the From and To date from the respective drop-down lists for the time horizon.

Specifying Groups and Output

To specify groups and output:

- 1 From the ABC Analysis window (see Figure 9.1 on page 133), click the Groups and Output tab.

Fig. 9.2
ABC Analysis—
Groups and Output
Tab

The screenshot shows the 'ABC Analysis' window with the 'Groups and output' tab selected. The 'Group By Fields' list is empty. The table on the right contains the following data:

	ABC Desc	Perc.
1	A	70.00
2	B	20.00
3	C	10.00

The 'Output Field' dropdown is set to 'Rank'. The 'Change', 'Add', and 'Delete' buttons are located below the table. The 'OK', 'Apply', and 'Cancel' buttons are at the bottom of the window.

The fields specified in the Group By Fields list determine what level for the items or what opinion line the analysis uses.

- 2 Click the Change button to select grouping fields.
- 3 Select the fields from the Available list, and then click the arrow to move them to the Selected list for a group. Click OK.

The Engine uses these fields as a group. Also, the Engine uses each group as a unit in the analysis. All the series belonging to the same group have the same ABC analysis category. For example, if the fields selected in the Group By list are Region and Customer, each Region-Customer combination is ranked.

- 4 Click the Add button. Another line is added to a new group and percentage.
- 5 Enter the group letter or name in the ABC Desc field.

- 6 Enter a percentage for the group in the Perc field.

Note Modify the ABC descriptions and percentages as required.

- 7 Put the cursor on a category line and click the Delete button.

Note Whatever description you select for the categories, line 1 corresponds to the top percentage category, line 2 corresponds to the second percentage category to the end of the defined categories. The sum of all the percentages must be 100. Also, make sure no blank or undefined lines exist for the ABC Desc and Perc rows.

- 8 Select the output parameter from the Output Field drop-down list. If you select Rank, the ABC description can only be one character in length.

- 9 Click OK and close the Engine.

Using an ABC Template

To use an ABC template:

- 1 While viewing a scenario, select Tools|ABC Analysis from the Viewer menu. The Do ABC Analysis window is displayed.




Fig. 9.3
Do ABC Analysis

- 2 Select the Whole Tree option to perform the ABC analysis on all scenarios (levels) of the applied hierarchy tree. This option is disabled if a hierarchy is not applied to current scenario.


- 3 Select the Scenario option and select the name of the scenario from the drop-down list to perform the ABC analysis on one scenario in the hierarchy tree.
- 4 Select the ABC template from the Available ABC templates drop-down list and then click Calculate.
- 5 When the ABC analysis finishes, a message prompts you to view the results. Click OK.

Analyzing Results

To analyze the results of an ABC analysis:


- 1 Click the Open/Close Expanded View button.

Fig. 9.4
ABC Analysis
Results



Note If the view does not display, right-click any record in the view.

- 2 Select Customize Expanded View from the right-click menu.
- 3 From the General Selection window, select the field chosen as the Output field on the Groups and Output tab. Select any other fields you want and move them from the Available list to the Selected list and click OK.
- 4 The results display. You can also see the results in run-expanded view.



The screenshot shows a window titled "[EAST] [BIG DEPOT] [1080] [GASS GRILLE] [SURE CHEF] [GRILLS]". Below the title bar is a toolbar with navigation buttons and "Select All" and "Deselect All" buttons. The main area contains a table with the following data:

Series	Forecast Adj...	Notes	Customer	Skus	Rank
[]	No	No			E
[] []	No	No			E
[EAST] [BIG DEPOT] [1080] [GASS G...	No	No	BIG DEPOT	1080	E
[EAST] [BIG DEPOT] [1352] [ELECTRI...	No	No	BIG DEPOT	1352	E
[WEST] [BIG DEPOT] [1080] [GASS G...	No	No	BIG DEPOT	1080	E
[WEST] [BIG DEPOT] [1352] [ELECTR...	No	No	BIG DEPOT	1352	E
[WEST] [BIG DEPOT] [1905] [RIDING ...	No	No	BIG DEPOT	1905	E
[EAST] [DISCOUNT PALACE] [1080] [...	No	No	DISCOUNT PALACE	1080	E
[EAST] [DISCOUNT PALACE] [1905] [...	No	No	DISCOUNT PALACE	1905	E
[WEST] [DISCOUNT PALACE] [1080] [...	No	No	DISCOUNT PALACE	1080	E

Fig. 9.5
ABC Analysis
Results

Sharing an ABC Template

After a template is created, it can be shared with other users so that each user can apply a similar ranking to their own data.

- 1 Select Applications|Engine Manager from the View menu.
- 2 Select the template you want to share from the ABC Templates list.
- 3 Right-click the template and select Review Assignments. The Assignments window is displayed.





Fig. 9.6
Assignments

- 4 Select the users from the Current users list and click the arrow to move them to the Assigned users list. Click Save.



Chapter 10

Using the Event Manager

This section describes how to manage special situations or actions that affect forecasts.

Introduction **140**

Creating an Event **141**

Introduction

Events are special situations or actions that can change forecasting behavior. An event can be applied to any series, and a series can be applied to any number of events at the same time.

It is important to track assumptions made during the forecasting process. The Event Manager can track assumptions made and applied to the forecast, so that each time the forecast is run, the same assumptions can be tracked. Events are first defined and then applied to the forecast. Each record is flagged and noted with the appropriate event.

The Event Manager lets you apply your business knowledge concerning the impact of outside events on the forecast data. Using the Event Manager, you can define, track, and publish outside events.

Figure 10.1 displays the forecast before the event is applied. This is a typical situation where a promotion is run during a time where there are slow sales. The Event Manager allows us to document the assumptions and the effect of the promotion on the forecast.


Fig. 10.1
Forecast before an
Event



In both of these figures that event has directly changed the calculated forecast. Once the location and timing of the event is placed, the details of the event must be worked out to quantify the effect and the duration. The Event Manager allows this to take place. When the event is put in place the adjustment to the forecast can be noted and tracked.

Figure 10.2 illustrates the effect of the 2-For-One sale on the forecast.

Fig. 10.2
Effect of an Event



By tracking and managing the promotion, the forecasting process begins to track the assumptions and impacts of marketing and business factors to the forecast.

The following terms are used often when discussing events:

Lift. The absolute percentage that represents the change between the Fvalue before and after applying the event. This value can also be an amount.

Promotion. A particular set of events that one uses to increase sales over the current sales.

Forecast Layers. The effect of several events.


Creating an Event

The Event Manager is the primary function for defining and launching forecasts with events. Use the Event Manager to create, remove, display, copy, and edit events.

Events are assigned to series. Only the owner who creates the event can delete or update it.

- 1 Select Forecast|Event Manager from the Viewer menu. The Event Manager window is displayed.

Fig. 10.3
Event Manager



- 2 Click the Add button. The Event Creation window is displayed.

Fig. 10.4
Event Creation

Description		Data	Business Rule	Detail	Attributes
ID #					
Name					
Description					
					Save Cancel

- 3 Enter the ID #, name, and description in the appropriate fields.

ID #. Specify an optional reference used when more than one event is assigned to a series. The order of the ID # specifies the order of how events are applied.

Name. The name must be unique for all the events created by this user. The name is a required field.

Description. Enter an optional description.

- 4 Click Save. The Event Manager window displays again.

- 5 Select the Event from the Event Manager list and click Edit.
- 6 Click the Data tab and enter values in the fields.

The screenshot shows the 'Event Creation' dialog box with the 'Data' tab selected. The dialog has five tabs: Description, Data, Business Rule, Detail, and Attributes. The 'Data' tab contains the following fields and controls:

- Start Period:** A dropdown menu showing '2003-01'.
- End Period:** A dropdown menu showing '2005-12'.
- Ref. Dates:** Two empty text input fields.
- Base Line:** A dropdown menu showing 'FValue'.
- Result Line:** A dropdown menu showing 'UserValue01'.
- Buttons:** 'Add', 'Remove', 'Save', and 'Cancel' buttons are located at the bottom right of the dialog.

Fig. 10.5
Event Creation,
Data Tab

Start Period. Enter a start period. The date must include the event date range.

End Period. Enter an end period.

Ref Dates. Enter reference dates for these two periods in the Ref Dates fields.

Base Line. Choose the observation field where the event calculation is based.

Result Line . Choose the observation field where the event calculation is saved.

Note The base line and result line drop-down lists follow the settings in the DisplaySet.

- 7 Click the Business Rule tab.

Fig. 10.6
Event Creation,
Business Rule Tab

The screenshot shows the 'Event Creation' dialog box with the 'Business Rule' tab selected. The dialog has five tabs: 'Description', 'Data', 'Business Rule', 'Detail', and 'Attributes'. The 'Business Rule' tab contains two sections: 'Applied Type' and 'Value Type'. In the 'Applied Type' section, there are three radio buttons: 'By Amt.' (selected), 'By Perc.', and 'Exact'. To the right of these is a text box containing '100.00'. In the 'Value Type' section, there are two radio buttons: 'Total Periods' (selected) and 'Per Period'. At the bottom right of the dialog are 'Save' and 'Cancel' buttons.

- 8 Enter values for Applied Type.

By Amt. Add the amount to the existing value.

By Perc. (Default) Add the percentage of the existing value. When By Perc is selected, Value Type is disabled.

Exact. Change the value using the assigned value.

- 9 Enter values for Value Type.

Total Periods. The amount in the Applied Type is the total amount of all the periods in which this event applies.

Per Period. (Default) The amount is for each period.

- 10 Click the Detail tab.




Fig. 10.7
Event Creation,
Detail Tab

The following options are available on the Detail window.

Auto Apply when Forecast. The event applies to the series automatically when the series is forecasted.

Likelihood. Refers to the possibility of the event. When an event saves the results, a trigger is evoked that calculates the ApplyPercentage. The value used for the event is calculated as Value * ApplyPercentage.

Cost. Reference field.

Price. Reference field.

- 11 Click the Attributes tab.





Fig. 10.8
Event Creation,
Attributes Tab

Use the Attributes tab to assign the category and class for the forecast. The default values are Event and Promotion, respectively.

To change the default values and predefined values, the system administrator can modify the corresponding records in the table TBL_ModuleClassValue in the Engine database.

- 12 Click Save to complete the event definition and begin assigning the event impact to the forecast.
- 13 Click the Assign and Apply button on the Event Manager window. The window shown in Figure 10.9 is displayed.

Fig. 10.9
Assign and Apply
Event



Using this window, you can apply an event at any time. Use the Batch Assign section to specify:

- Whole Tree
- A scenario
- Selected Series in Viewer
- Selected Series in previous List


- 14 Select the Selected Series in previous List option.

- 15 Click the Assign button.
- 16 In the General Selection window, select the series from the Available list and click the single right arrow to move the series to the Selected list.
- 17 Click OK.

When a series is assigned an event/promotion, a special symbol P displays next to the series in the Viewer. In this way you can run a filter to select only the promoted items. See Chapter 8, “Using Filters and Exceptions,” on page 113.

After an event applies to a series, a series note is added to that series to indicate the event has been applied.
- 18 Click the Apply button to apply the event to the series on the fly.
- 19 Click Save.

The Event Manager lets users track and manage external impacts that affect the forecast, and hence to track and manage forecasts more effectively.



Chapter 11

Using Conversion Factors

This section describes how to use conversions to change forecast units.

Introduction **150**

Renaming a Conversion Factor **150**

Applying a Conversion Factor **152**

Limitations **153**

Introduction

The Engine Manager provides a DisplaySets function that converts current data into different units of measurement—for example, from pounds to kilograms, or liters to gallons. It is also possible to create factors that change, for example, total unit sales to average daily sales. This is important because often the number of sales days in a month affects sales volume: by normalizing the raw data to account for the number of sales days can improve the forecast accuracy.

Conversion factors are bidirectional. This means that if you sell 100 units at \$10 to achieve \$1,000 in revenue and expect revenue to increase to \$1,200, you could adjust either units or price to achieve this. With conversion factors, you can adjust the revenue and have it affect either the unit or the price. The results are determined by the way in which the conversion factor is applied. Figure 11.1 illustrates how conversion factors can be used.

Fig. 11.1
Sample Conversion

Revenue Increase from Unit Volume Increase					
Base Number	Conversion Factor	Result	Change	Return To Base	Result
Units 100	Price \$10	Units * Price \$1,000	Revenue \$1,200	Revenue / Price 120	Price \$10
Revenue Increase from Price Increase					
Base Number	Conversion Factor	Result	Change	Return To Base	Result
Price \$10	Units 100	Price * Units \$1,000	Revenue \$1,200	Revenue / Units \$12	Units 100

The number on the left side of the equation is the base that is converted by the factor. In this way you can achieve the revenue growth you want through both pricing and unit volume changes.

Renaming a Conversion Factor

It is not sufficient just to have conversion factors loaded into the system—the name of the conversion factor must also be visible. This is accomplished using the DisplaySet feature in the Engine Manager.

To rename a conversion factor:

- 1 Select Applications|Engine Manager from the Viewer menu.
- 2 Right-click the Display Sets resource and choose Create. The DisplaySet Setup window is displayed.




Fig. 11.2
DisplaySet Setup

Using the Fields tab allows the system administrator to set the name of the conversion factor that displays when a user opens the Viewer.


Note Conversion factors are like any other observation field—they contain data. Typically the conversion factor data is loaded from another scenario using the Batch Scheduler. Once the conversion factor is added to the scenario, it is used and adjusted just like other data that is part of a scenario by using the Adjustments features.

Applying a Conversion Factor

Once a conversion factor is defined and the data loaded, the conversion factor can be assigned and applied to other items within QAD Demand Management.

- 1 Select a scenario.
- 2 Select Tools|Apply Conversion Factor from the Viewer menu. The Apply Conversion Factors to Current Scenario window is displayed.

Fig. 11.3
Apply Conversion
Factors to Current
Scenario



Adding a New Opinion Line

To add a new opinion line:

- 1 Click the Add button and select the appropriate conversion factor column from the Conversion Factor list.
Note To return to the initial state, click the Return Base button and click Yes.
- 2 You are prompted if you want to convert the scenario back to its base line. Click Yes.


Limitations

The 4.0 Engine only supports one level of conversion. That is, if you apply a conversion factor multiple times, the Engine converts the data back to the base unit first before doing the conversion again.

For example, to convert your data from kilograms to grams and then to milligrams, the Engine performs the conversion by converting kilograms to grams, then grams to kilograms, then kilograms to milligrams.

When completing multiple conversions, make sure to define the conversion factors correctly.

Note Make sure to convert the database to the base unit before committing it to the Living Master scenario.

The background of the page features a grayscale image of several interlocking gears. The gears are of different sizes and are positioned in a way that creates a sense of depth and mechanical complexity. The lighting is soft, highlighting the teeth of the gears and casting gentle shadows.

Chapter 12

Using the Calendar

This section describes how to use calendars when viewing forecasts.

Introduction **156**

Overview **156**

Creating a Calendar **160**

Using a Calendar **164**

Sharing a Target Period **166**

Introduction

Calendars let you examine and modify information in any time scale. You can collect data at multiple levels of time detail and then convert them all to a common base so that the information can be merged together to form a consistent whole. Figure 12.1 illustrates the relationships that can be established using a calendar moving from one time scale to another.

Fig. 12.1
Sample Calendars



Monthly Scenario	Jan (Sales 100)				Feb (Sales 100)			
Weekly Cal Pct. Scenario	75%	5%	5%	15%	5%	5%	30%	60%
Weekly Scenario	75	5	5	15	5	5	30	60
Monthly Cal. Pct. Scenario	100%	100%	100%	5%	95%	100%	100%	100%
Monthly Scenario	Jan (Sales 86)				Feb (Sales 114)			

In this example, you are placing demand in the appropriate period based on percentage-based and calendar-based profiles.

The Calendar function lets you convert data into different time periods. The Engine supports any calendar transformation. The default time levels are:

- Monthly
- 4 Weeks
- Weekly
- Quarterly
- Daily

Overview

Analyzing the same data at different times occurs in many Demand Planning processes.

Example Promotion is usually weekly-based; however, the most accurate forecasting level is monthly, and the financial forecasting level is quarterly.

In the Engine, each scenario is at a specific time level, such as weekly or monthly. In order to view the data at different time levels, the Calendar function transforms the specific time level from one scenario to another scenario with the time level you want. In addition to the source and target scenarios, a calendar scenario is used to hold the mapping information between the two time levels.

Figure 12.2 displays the relationship among the three scenarios using weekly and monthly scenarios as an example.




Fig. 12.2
Relationship
Between Calendars

Note You can use the same calendar scenario for two transformation directions, such as weekly to monthly and monthly to weekly.

The time-level information is defined in the scenario attribute. The two attributes for the time level information are Time Level 1 and Time Level 2.

For a non-calendar scenario, these two attributes have the same values: the time level of that scenario. For a calendar scenario, these two attributes indicate this scenario can map between two time levels. For example, the scenarios in the previous example have the settings listed in the Table 12.1.

Scenario Attribute	Time Level 1	Time Level 2	User Field 1
Weekly Scenario	Weekly	Weekly	Normal
Calendar Scenario	Monthly	Weekly	Calendar
Monthly Scenario	Monthly	Monthly	Normal

Table 12.1
Sample Calendar
Scenario Values

Special Format of the Source Data of Calendar Scenario

The source data of a calendar scenario can be any type, such as an SQL table, Excel, Access table or others. However, since the calendar scenario holds the transformation information, it requires a special data format.

- The data should be organized by column, not row.

- The data must include one description column for date mapping information. This information defines which dates in the two different time levels match.
- You must include two value columns for each profile.

A profile defines the value to map one date to another date. Table 12.2 provides sample source data. Use this spreadsheet as an example of the detail format requirement. The Viewer installation provides a sample Excel calendar data source in the Engine directory.

Once a calendar scenario is loaded, the Engine parses the mapping information and saves the information to an internal table before applying it in a calendar transformation. As a result, if you change the profile values in a calendar scenario directly, it does not affect the transformation result. To change the profile, the source data should be first changed, and then the calendar scenario reloaded.

Table 12.2 illustrates the format of the data source. The Date Calculation column defines the date mapping information and is loaded as the Dates field. The C_F, C_T, D_F, and D_T columns define the mapping profile and are loaded to the Engine as Observation Value fields.

Table 12.2
Data Source Format

Level 2 Date	Level 1 Date	Level 2 Periods	Level 1 Periods	Date Calculation	C_F	C_T
2001-01	2001-01	52	12	2001-01~2001-01~52~12	23.076923	100
2001-02	2001-01	52	12	2001-01~2001-02~52~12	23.076923	100
2001-03	2001-01	52	12	2001-01~2001-03~52~12	23.076923	100
2001-04	2001-01	52	12	2001-01~2001-04~52~12	23.076923	100
2001-05	2001-01	52	12	2001-01~2001-05~52~12	7.6923077	33.333333
2001-05	2001-02	52	12	2001-02~2001-05~52~12	15.384615	66.666667
2001-06	2001-02	52	12	2001-02~2001-06~52~12	23.076923	100
2001-07	2001-02	52	12	2001-02~2001-07~52~12	23.076923	100

Level 2 Date	Level 1 Date	Level 2 Periods	Level 1 Periods	Date Calculation	C_F	C_T
2001-08	2001-02	52	12	2001-02~2001-08~52~12	23.076923	100
2001-09	2001-02	52	12	2001-02~2001-09~52~12	15.384615	66.666667
2001-09	2001-03	52	12	2001-03~2001-09~52~12	7.6923077	33.333333
2001-10	2001-03	52	12	2001-03~2001-10~52~12	23.076923	100
2001-11	2001-03	52	12	2001-03~2001-11~52~12	23.076923	100
2001-12	2001-03	52	12	2001-03~2001-12~52~12	23.076923	100
2001-13	2001-03	52	12	2001-03~2001-13~52~12	23.076923	100
2001-14	2001-03	52	12	2001-03~2001-14~52~12	0	0
2001-14	2001-04	52	12	2001-04~2001-14~52~12	23.076923	100

Each column from this data source is defined in the following list.

Level 2 Date. Date information for Time Level 2 (Weekly).

Level 1 Date. Date information for Time Level 1 (Monthly). Notice that week 1 to week 5 belongs to month 1.

Julian Date. Not shown in example; not currently used.

Level 2 Periods. Defines how many periods of Time Level 2 per year. In this example, there are 52 because there are 52 weeks.

Level 1 Periods. Defines how many periods of Time Level 1 per year. In this example, the value is 12 because there are 12 months.

Date Calculation. Combines the first 4 columns; you must use the tilde (~) as separator:

Level 1 Date~Level 2 Date~Level 2 Periods~Level 1 Periods

C_F. Defines the percentage of Time Level 1 allocated to Time Level 2. For example, in first line, 23.07% of January is allocate to week 2001-0

C_T. Defines the percentage of Time Level 2 allocated to Time Level 1. For example, in first line, 100% of week 1 (2001-01) is allocated to January.

C_F and C_T define one profile. You can define more than one profile by adding additional columns such as D_F and D_T. I When you apply a calendar, each series can have a different profile.


In order to use the Calendar function, first create and load a calendar into the Engine database. Then, create as many calendars as necessary—there are no limits.

Creating a Calendar

To create a calendar:

- 1 Click the Base Unit icon. The Base Unit window is displayed.
- 2 Click the Scenario Management button. The Scenario Management window is displayed.

Fig. 12.3
Scenario
Management



- 3 Click the New button. Enter the scenario name and click OK.
- 4 Click Edit to open a scenario for the calendar. The Edit Scenario window is displayed.

The 'Edit Scenario' dialog box is shown with the following fields and values:


Name	NEW-SC		
Type		Class	
Use		Status	
Time Level 1	Monthly	Time Level 2	Weekly
Time Level 3		Time Level 4	
User Field 1	Calendar	User Field 2	
User Field 3		User Field 4	
User Field 5		User Field 6	

Buttons: OK, Cancel

Fig. 12.4
Edit Scenario

- 5 Select Calendar for the appropriate time frame from the User Field 1 drop-down list.
- 6 Select Monthly or the first source (From Calendar) from the Time Level 1 drop-down list.
- 7 Select Weekly or the target calendar (To Calendar) from the Time Level 2 drop-down list.
- 8 Click OK and close the dialog box.
- 9 Click the Right Arrow button beside the Data Source Name Management box. The Data Capture window is displayed.

Fig. 12.5
Data Capture



- 10 Select the data source and click Done. This document uses an Excel file as the data source.
- 11 In the Base Unit window, select the Columns check box and enter 0 in the Forecast field.
- 12 Click the Advanced button. The Advanced Options window is displayed.




Fig. 12.6
Advanced Options

- 13** Click the Dates button. The Fields Selection for Dates window is displayed.




Fig. 12.7
Fields Selection for
Dates

- 14 Select Date Calculation from the Available list and click the arrow to move it to the Selected list. Click OK. There is only one date calculation.
- 15 In the Advanced Options window (Figure 12.6), click the Cleansing tab.
- 16 Make sure no check boxes in the Removing Zero Values section are selected.
- 17 Click OK to close the Advanced Option dialog box and click ProCast! to load the data.

Using a Calendar

Before applying a calendar, make sure that the source scenario is compatible with existing calendars in the Engine environment. For example, if you are using a scenario with quarterly results, make sure a quarterly calendar exists to process the scenario data.

Checking for Compatibility

- 1 From the Engine window, click the Base Unit icon. The Base Unit window appears.
- 2 Select a scenario and click the Edit button. The Edit Scenario window displays (Figure 12.4 on page 161).
- 3 If your source scenario is month-based, set both Time Level 1 and Time Level 2 fields to Monthly. If the scenario is week-based, set both fields to Weekly. The system default is Monthly.
- 4 Click OK to close the Scenario Management window.
- 5 Close the Engine and return to the Viewer.
- 6 Open the source scenario.
- 7 Choose Tools|Apply Calendar from the Viewer menu. The Apply Calendar window is displayed.




Fig. 12.8
Apply Calendar

8 Review the following definitions:

Source and target time frames. Lists the conversion time periods that can be applied to the source scenario.

DisplaySet. Select this option if you want to promote a display set from your source scenario down to the target scenario.

Selected calendar. Lists all calendars that are compatible with the source scenario.

Selected target scenario. Lists the target scenarios that are compatible with the selected calendar. Unlike the conversion factor, you cannot apply a calendar and put the result back to the source scenario—you must create a compatible target scenario. Also, if the target scenario is used by other users, you cannot apply the calendar to that scenario.

- 9** Click the Profile mapping tab. Use this tab to specify which profile each individual item uses.
- 10** Click the Apply Calendar button.

Sharing a Target Period

If you have a period that is shared between two different time periods—for example, week 9 is shared by February and March—the forecast or historical periods can shift by one period. Table 12.3 provides an example.

Table 12.3
Example of
Calendar Time
Shift

From (month)	To (week)	% (from monthly to weekly)	% (from weekly back to monthly)
2000-02	2000-05	20%	100%
2000-02	2000-06	20%	100%
2000-02	2000-07	20%	100%
2000-02	2000-08	20%	100%
2000-02	2000-09	20%	50%
2000-03	2000-09	20%	50%
2000-03	2000-10	20%	100%
2000-03	2000-11	20%	100%
2000-03	2000-12	20%	100%
2000-03	2000-13	20%	100%

This example is simplified and only used to illustrate this concept. The actual calendar is more complex.

Notice that week 9 is shared between February and March. Assume that the value is 100 for both February and March, as shown in Table 12.4.

Table 12.4
Shared Calendar
Values

Period	Value
February (historical)	100
March (forecast)	100

After applying the calendar, the target scenario has the results shown in Table 12.5.

Table 12.5
Results of Applying
Calendar

Period	Value
2000-05 (historical)	20
2000-06 (historical)	20
2000-07 (historical)	20

Period	Value
2000-08 (historical)	20
2000-09 (forecast)	40
2000-10 (forecast)	20
2000-11 (forecast)	20
2000-12 (forecast)	20
2000-13 (forecast)	20

Notice that week 9 is forecasted. This is because for week 9, part of the data comes from a historical source, and part from the forecast source. In this situation, the Engine sets the value to forecast.

Table 12.6 shows the results when the weekly scenario is converted back to monthly.

Period	Value
2000-02 (forecast)	100
2000-03 (forecast)	100

Table 12.6
Convert Weekly to
Monthly

Notice that February now becomes forecast. Again, like week 9, part of its data comes from a historical source, and part from the forecast source (week 9). By default, the Engine sets the February result to forecast.

Under some conditions, the Engine adds periods to the target scenario. Assume that you have the same calendar as the previous example (see Table 12.1 on page 157).

However, data only exists for February, not March, as shown in Table 12.7.

Period	Value
February (historical)	100

Table 12.7
Data for February
Only

Table 12.8 shows how the data looks after the month-to-week conversion.

Period	Value
2000-05 (historical)	20
2000-06 (historical)	20
2000-07 (historical)	20

Table 12.8
Results of
Conversion

Period	Value
2000-08 (historical)	20
2000-09 (historical)	20


Notice that 2000-09 is created, but this is normal because in your calendar, February has week 9. Converted back to monthly, it looks as shown in Table 12.9.

Table 12.9
Converted to
Monthly

Period	Value
2000-02 (historical)	90
2000-03 (historical)	10

Notice that March is created and takes 10 from February, because calendar week 9 is shared between February and March.

Theses are the limitation for the current Engine calendar transformation implementation. This shifting of periods only occurs when a target period is shared by two different source periods.



Chapter 13

Using the Bill of Material

This section describes how bills of material can convert primary demand into component demand.

Introduction **170**


Creating a BOM **172**

Introduction

The Bill of Material (BOM) function enables primary demand to be converted into component demand. You can forecast finished goods items and then, by definition, component requirements are computed using the planning bill.

The Engine supports multiple planning bill levels by loading each level, and then performing an expansion either in real time or in batch with linked tasks. Figure 13.1 shows the scenario configuration required to support the BOM process.

Fig. 13.1
BOM Setup



In the BOM process, you can have the same scenario for the finished goods and component forecast. Using this process, the BOM can determine the number of hours required to produce the items. Using this BOM process creates new information that is stored in the main scenario for analysis and review.

The Bill of Material feature uses the forecasts of a scenario of finished goods to calculate forecasts of components for those finished goods. The feature accomplishes the task using a third BOM scenario. This scenario defines the relationship between a product and the type and number of components it contains. The Engine extends the feature by generating forecasted values for the BOM, as well as any other opinion lines used in the BOM calculation.

Example

An electronics store sells televisions and accessories. Each type of television comes with a different type of remote control and cables. For example, a Sony HDTV comes with 3 video cables and 2 audio cables; a standard Sony television comes with 1 video cable and 2 audio cables. In addition to the cables sold with a television, cables are also sold

separately. As store manager, you are in charge of forecasting the total sales of televisions and cables. The Bill of Material function can help you accomplish this task.

Table 13.1 shows the input data.

Sales Forecast	2004-01	2004-02	2004-03
Sony HDTV	100	150	200
Standard TV	100	90	80

Table 13.1
Finished Goods

Table 13.2 shows the components.

Retail Sales Forecast	2004-01	2004-02	2004-03
Video Cable	80	80	80
Audio Cable	120	150	160

Table 13.2
Components

Table 13.3 shows the BOM.

Finish Goods	Components	BOM Factor 2004-01	BOM Factor 2004-02	BOM Factor 2004-03
Sony HDTV	Video Cable	3	3	3
Sony HDTV	Audio Cable	2	2	2
Standard TV	Video Cable	1	1	1
Standard TV	Audio Cable	2	2	2

Table 13.3
BOM

Table 13.4 shows the output data for the BOM.

Total Sales Forecast	2004-01	2004-02	2004-03
Video Cable	480	620	760
Audio Cable	520	630	740

Table 13.4
Components

For example, the total sales forecast of video cable on 2004-02 = $80 + 150 * 3 + 90 * 1 = 80 + 450 + 90 = 620$.

Note The BOM factor is date-based. In the example, the factors are the same for all months. But they might vary from period to period.

The BOM factor can be a predefined factor or any other quantity relationship. For example, 80% of HDTV buyers purchase a universal remote control, while 20% of HDTV buyers purchase a standard remote control. The 80% and 20% can be the BOM factors between the HDTV and remote controls.

Creating a BOM

To create a BOM:

- 1 Choose Forecast|Bill of Material from the Viewer menu. The Bill of Materials Transfer window is displayed.

Fig. 13.2
Bill of Materials
Transfer

The screenshot shows the 'Bill of Materials Transfer' window. It contains three main panels: 'Finished Goods', 'Components', and 'BOM'. Each panel has a 'Scenario' dropdown menu currently set to 'Living Master'. The 'Finished Goods' and 'Components' panels each have two empty boxes labeled 'Connection Fields' and 'Data Fields', each with a 'Change...' button below it. The 'BOM' panel has three empty boxes labeled 'Connection Fields to Finished Good', 'Connection Fields to Components', and 'BOM Factors', each with a 'Change...' button below it. There are checkboxes for 'Forecasted Periods Only' and 'Reset Values before Transfer'. At the bottom right are 'Apply' and 'Cancel' buttons.

- 2 Review the following options:

Scenario. Specify the scenario for finished goods (FG), components, and BOM.

Connection Fields. Define how to match the series between FG and BOM, and BOM and components. Additional information is provided in the following steps.

Data Fields. Define which opinion lines (data fields) are used for the BOM calculations. Usually, they are Value and FValue fields. Additional information is provided in the following steps.

Forecasted Periods Only. BOM calculation only affects the forecasted periods. Use this option if the historical values of the components needs to be calculated from the finished goods. The default is blank.

Reset Values before Transfer. BOM calculation resets the opinion lines (data fields) in the component scenario before the calculation occurs. In the previous example, the components have their own retail sales forecast and this option should not be used. The default is blank.

Apply Button. Apply the BOM to the component scenario with the current settings.

Cancel Button. Cancel the BOM operation.

- 3 Click the Change button in the Connection Fields of the Finished Goods (FG) section.
- 4 Select the appropriate connection fields from the Connections Fields list in the FG section. These fields connect series from the FG scenario to the conversion series in the BOM scenario.
- 5 Click the Change button in the Data Fields of the FG section.
- 6 Select the appropriate data fields from the Data Fields list in the Finished Goods section. These data fields are the opinion lines used as values converted into component opinion line values.
- 7 Click the Change button in the Connection Fields of the Components section.
- 8 Select the appropriate connection fields from the Connections Fields list in the Components section. These fields connect series from the Components scenario to the conversion series in the BOM scenario.
- 9 Click the Change button in the Data Fields of the Components section.
- 10 Select the appropriate data fields from the Data Fields list in the Components section. These data fields are the opinion lines that receive the converted values from the FG opinion lines. The number

of opinion lines entered must be the same as the number of opinion lines entered for finish goods. The order of the lines entered is important because the first FG data field (opinion line) is used for the first component opinion line, with each additional line matching in the same order.

- 11 Click the Change button in the Connection Fields to Finished Good in the BOM section.
- 12 Select the appropriate connection fields that hold the data that matches the data defined in the Finished Goods connection fields. The number of connection fields must be the same as the number of FG connection fields. The order of the fields is significant (see step 10).
- 13 Click the Change button for the Connection Fields to Components in the BOM section.
- 14 Select the connection fields that hold the data that matching the data defined in the Components connection fields. The number of connection fields must be the same as the number of component connection fields. The order of the fields is significant (see step 10).
- 15 Click the Change button for BOM Factors in the BOM section.
- 16 Select the observation lines that define the conversion between the finished goods and the components. When multiple data fields are defined for Finished Goods and Components, there are two options:
 - Enter a single BOM Factor that will be used for all conversions.
 - Enter multiple lines that must equal the number of observation lines entered in the Finished Goods or Components scenario.

Figure 13.3 displays an example of a completed Bill of Materials Transfer window.




Fig. 13.3
Completed Bill of
Materials Transfer

17 Click Apply.

Configuring Statistical Forecasts

This section describes how to configure forecasts in QAD Demand Management.


<i>Introduction</i>	178
<i>Selecting Seasonality</i>	179
<i>Setting Forecast Periods</i>	180
<i>Setting Reload Options</i>	181
<i>Using Incremental Updates</i>	183
<i>Including New Items in a Forecast</i>	183
<i>Using the Clone Feature</i>	184
<i>Copying/Moving Scenarios</i>	184
<i>Including New Items Using Batch Scheduler</i>	184

Introduction

The Engine provides configuration options to help you create accurate forecasts. The following diagrams illustrate the parameters and options available.

Seasonality lets you indicate to the Engine that the data might contain patterns. In general, it requires two times the seasonal length of data to generate an accurate forecast. Monthly data requires at least 24 months of data to generate a good seasonal forecast. Seasonality can also be generated using smaller data sets with advanced product training. Figure 14.1 illustrates seasonal patterns.

Fig. 14.1
Seasonal Patterns



The Overwrite and Reload features within the Engine let you control the information that is over-written when a forecast is run.

Typically, all non-adjusted values in a forecast are overwritten when the forecast is rerun. To preserve existing values during a rerun:

- Turn off Overwrite Adjusted Forecast, or
- Specify a forecast overwrite range, or
- Use both methods together

The Overwrite Range feature assumes that if users have reviewed a forecast, users have tacitly agreed to the result. To change those numbers with a new forecast means that the numbers the user had approved would no longer be in place. Using the overwrite range prevents this from happening. Figure 14.2 illustrates this point.




Fig. 14.2
Preserving Values
During a Rerun

Here both the Overwrite Adjusted Forecast and Forecast Overwrite range are used to provide a stable forecast that can be adjusted. This chapter describes how to set the following configuration parameters for the forecast:

- Seasonality
- Forecast Periods
- Reload Options
- Overwrite Adjusted Forecast
- Reload Options
- Freeze Scenario
- Forecast Overwrite Range
- Reload Series

This chapter also briefly introduces how to perform incremental updates and include new items in existing forecasting data pools.


Selecting Seasonality

Using the Engine, you can indicate if your data is subject to seasonality. Although the Engine can automatically detect seasonality, using the Seasonality feature provides the most accurate forecast results.

To select seasonality:

- Enter the seasonality in the Seasonality field on the Data Capture tab. For example, if data contains quarterly seasonality, enter 4 in the field.

Fig. 14.3
Base Unit



The screenshot shows the 'Base Unit' application window with the 'Data Capture' tab selected. The 'Scenario' field is set to 'Scenario 20'. The 'Data Source Name' is 'ENGINE_TRAINING_2006.mdb'. The 'Tables/Queries' field is 'ADPGOAL'. Under 'Data Organization', the 'Rows' radio button is selected. The 'Seasonality' field is a spinner box currently showing an empty value. The 'Forecast' field is a spinner box showing '12'. An annotation arrow points to the 'Seasonality' field with the text 'Enter a seasonality value.'

Setting Forecast Periods

You can specify the number of periods you want the Engine to include in the forecast. The Engine defaults to twelve periods and can forecast up to 999 periods.

To set forecast periods:

- In the Base Unit window (Figure 14.3 on page 180), enter the number of periods in the Forecast field. For example, if data contains quarterly seasonality, enter 4 in the field.

Setting Reload Options

Click the Advanced Button on the Base Unit window (Figure 14.3 on page 180) and then click the Reload Options tab. This tab lets you fine-tune forecasts and the reloading of data for scenarios.




Fig. 14.4
Reload Options Tab

Freezing a Scenario

When scenario properties change, the Engine must reload the scenario in order for the changes to take effect. Reloading a scenario removes all existing forecasted values and modifications. The Freeze Scenario option lets you freeze these scenario properties, preventing the Engine from reloading the scenario. When a scenario is frozen, you cannot reload a series or adjust any configuration options associated with reloading a series.

When a scenario is frozen, the following configurations are read-only:

- Dates
- Series Configuration (Labels, Observations, Series Properties)
- Remove Trailing and Leading Zeros
- Outliers and Missing Values
- Grouping

The Freeze Scenario option is located on the Advanced Options window on the Reload tab.

Note With shared scenarios, you can freeze or unfreeze scenarios only if the system administrator has granted permission in User Preferences.

Overwriting Forecasts

The Forecast Overwrite Range feature lets you select a specific forecast range to be overwritten with a new system-generated forecast. The forecast is overwritten when the historical data is reforecasted. You can turn this option on or off and set the overwrite range. The advantage is that you only overwrite a specific range and keep the rest of the forecast as is.

The minimum value is 1, and the maximum value is the total number of forecasted periods.

Note With shared scenarios, you can freeze or unfreeze scenarios only if the system administrator has granted permission in User Preferences.

Reloading Series

Click the Reload Series button to reload series. Reloading removes all your existing data, including all the modifications, and then reloads the series. After reloading the series, the Engine automatically first reloads all the observations, and then forecasts when you click the Procast! button.

To reload the observations without a forecast or changes to the original data, use the following steps.

- 1 Enter 0 in the Forecast Period field.
- 2 Clear the following options:
 - Replace Outliers

- Replace Missing Values
- Remove Zero Values
- Convert negative value to 0

3 Click the Procast! button to load the observations.

Using Incremental Updates

At the beginning of each forecasting cycle, the latest historical data is usually available. This latest information affects the forecast accuracy directly. Including this new information for the new forecasting cycle is called an incremental update.

You can include the latest historical data in a scenario by changing the observation configuration and reloading the data. However, doing this also removes all existing forecasted values, including modifications. Usually this is not an acceptable solution.

To apply the most recent historical data to an existing, established scenario without losing the existing forecast and modification, use the Incremental Update feature in the Batch Scheduler. See Chapter 20, “Using the Batch Scheduler,” on page 255.

Including New Items in a Forecast

Since business changes and grows every day, there is always a demand to include new items in the forecasting data pools. These new items can be new customers, new stores, or new products.

Similar to the incremental update process, one way to include new items is to reload the scenario data. However, this also removes all existing forecasted values, including modifications. Usually reloading data is not an acceptable solution. The Engine provides several ways to include new items for an existing scenario without losing the forecasted values and modifications.

Using the Clone Feature

In the Viewer, use the Clone feature to rename an existing series. Modify the new series with the properties of the new item and the new values. This method is straightforward, but is not efficient if there are many new items. For details on cloning and editing a series, see Chapter 6, “Viewing Forecasts,” on page 71.

Copying/Moving Scenarios

In the Viewer, copy or move one or several series from one scenario to another scenario. Use Copy/Move to include new items in an existing scenario. First, create a scenario to load the new items. Use the Viewer to Copy/Move all the items in the new scenario to an existing scenario. This method is more efficient than the Clone method, but requires you to identify the new items from the source data. For details on cloning and editing a series, see Chapter 6, “Viewing Forecasts,” on page 71.

Including New Items Using Batch Scheduler

Using the Batch Scheduler, select Massive Observation Update from the drop-down list when you add the batch job. See “Massive Observation Update” on page 269 for details.

Use the Send New Item to Scenario option to include items for an existing scenario. This option works in a standalone process or with an Incremental Update process and is similar to Copy/Move. The advantages of using this option are:

- No user interaction is required.
- You do not need to identify the new items from the source data.

See Chapter 20, “Using the Batch Scheduler,” on page 255.

Improving Forecast Accuracy

This section explains how to improve the accuracy—and hence reliability—of forecasts.

Introduction **186**

Same As Method **186**

Data Cleansing **190**

Forecast Accuracy **192**

Forecast Techniques and Methods **194**

Configuring Forecast Methods **195**

Regression Methods **204**

Group Forecasting **210**

Introduction

This chapter explains how to use the following features:

- Same As Method
- Data Cleansing
- Forecast Accuracy
- Forecast Techniques and Methods
- Regression Methods
- Group Forecasting

Same As Method

There are two types of Same As methods:

- Same As History
- Same As Seasonality


When you configure Same As values, they do not take effect until you run the forecast.

Same As History

Same As History affects the product history. The Same As History function lets you replace one product's history with another product's history by a certain percentage.

For example, if you have three products P1, P2, P3, you can indicate that P2's history is 30% of P3's history, and P1's history is 15% of P3's history.

Fig. 15.1
Same As Example



Same As Seasonality

The Same As Seasonality function lets you replace one product's forecast with another product's forecast to simulate data for products with similar properties. Same As Seasonality affects the product's forecast.

For example, if you have products P1, P2, P3, you can indicate that the forecast of P1 should be similar to that of P2, and P3's forecast should be similar to P2's seasonality.

Configure the Same As Option


Since configuring the Same As History and Same As Seasonality features are similar, this section uses the same instructions for configuring both methods.

Note When you enable the Same As option, the Remove Trailing Zeros option is disabled. See "Group Forecasting" on page 210.

To configure the Same As option:

- 1 Click the Advanced button on the Base Unit window.
- 2 The Advanced Options window is displayed.

Fig. 15.2
Advanced Options



- 3 Select the Enable Same As check box and click the Same As button. Then Engine loads the series and the Same As Series window is displayed.




Fig. 15.3
Same As Series

- 4 In the Description column, select the series to which you want to apply the Same As method.
Note Hold down the Ctrl key to select more than one series, or use the Select All button to select all the series in the Description column.
- 5 Select the series in the Available Series column that contains the History or Seasonality you want to apply to the series. You can select one series.
- 6 Select the option from the Same As drop-down list. The option specifies whether the selected series has the same History, Seasonality, or both.
- 7 If you select History or Both, enter the percentage by which the series has the same History. The percentage applies only to the series History.
- 8 Click Apply. The appropriate values populate the Seasonality, History, and Percentage columns.
 To change any values in the columns, you must use the fields in the Same As Series window.

- 9 Click OK to apply the method.

Data Cleansing

You can clean your data to generate a more accurate forecast. Data cleansing options for historical shaping include:

- Removal of outliers
- Removal of zeros
- Removal of missing values

To cleanse data:

- 1 Click the Advanced button on the Base Unit window. The Advanced Options window is displayed.
- 2 Click the Cleansing tab.

Fig. 15.4
Advanced Options,
Cleansing Tab

Advanced Options

Configuration | **Cleansing** | Accuracy | Reload Options

Outliers

☒ Replace outliers Std. Dev

Forecast Technique

Replacement Type

Missing Values

☒ Replace Missing Values

Forecast Technique

Lower Val Upper Val

Remove Zero Values

☒ Leading ☐ Trailing

Ok

3 Select the Outliers options you want to use in the forecast.

Replace Outliers. Select this option to replace outliers. An outlier is an anomaly within a data set with a low chance of occurrence. Outliers disturb forecast accuracy. For example, the effect that a July snowstorm in Texas would have on the sale of snow shovels would be an outlier within a data set.

Std Dev. Specify the standard deviation. The default is two standard deviations from the mean. The outlier is replaced with a number that is calculated using a forecast technique.

Or, you can select how far a data point needs to be from the data series mean in order to be considered an outlier by selecting the number of standard deviations.

Forecast Technique. Select the forecasting technique to use to replace outliers.

Replacement Type. Select the type of value with which you want outliers replaced. Choose Upper value, Lower value, or Both.

4 Select the Missing Values options you want to use in the forecast.

Replace Missing Values. Select Replace Missing Values to replace missing or erroneous values in the data set. For example, if a period or a range of periods is not reported, instead of leaving it blank, the Engine can replace it with a value calculated by a forecast technique.

Forecast Technique. Select the forecasting technique you would like to use to perform replace missing values.

Lower Val/Upper Val. To expand the scope of the missing values range beyond zero, change the lower and upper values.

5 Select the Remove Zero Values options you want to use in the forecast.

Select either or both of these options to ignore zero values. This is useful in new product introductions, when long stretches of zeros represent times when the product did not exist.

When you select an option, the Engine advises you to reload your data.

Leading. Select this check box to remove leading zeros.

Trailing. Select this check box to remove trailing zeros.

Forecast Accuracy

The Accuracy tab contains additional options for ensuring forecast accuracy.

To set Accuracy options:

- 1 Click the Advanced button on the Base Unit window. The Advanced Options window appears.
- 2 Click the Accuracy tab.

Fig. 15.5
Advanced Options,
Accuracy Tab

Advanced Options

Configuration | Cleansing | **Accuracy** | Reload Options

Tracking Signal

☒ Apply Tracking Signal Iterations:

Signal Type:

Under Forecast Percentage: Over Forecast Percentage:

Holdback Evaluation

☒ Use Holdback Iterations:

Trend Adjustment

☒ Trend Adjustment

Growth Periods: Decline Periods:

☒ Convert negatives values to zero

☐ Integer Only Forecast

Ok

- 3 Select the accuracy options for the forecast.

Tracking Signal. Select the Apply Tracking Signal check box to use this option.

This option fine-tunes the overall accuracy of a forecast. Fitted values serve as the baseline for your forecast, while the Tracking Signal option determines if the fitted values are under- or over-forecasting.

Comparing it with the overall mean absolute percentage error (MAPE) improves the accuracy by less than a percentage point.

Iterations: Indicate the maximum number of iteration of applying the tracking signal modification.

Signal Type: Specify Average to calculate the tracking signal by averaging the error for each data point. Specify Smooth to calculate the tracking signal by smoothing out the error for each data point. In this case, the latest data carries more weight.

Under Forecast Percentage: Enter the Under Forecast Percentage to be tracked.

Over Forecast Percentage: Enter the Over Forecast Percentage to be tracked.

Holdback Evaluation. Select the Use Holdback check box to use this option.

Holdback Evaluation functionality lets you judge whether the current forecasting process is sufficient by specifying a time period in data to be held out and compared with another mini-forecast. The Engine produces an out-of-sample table in the Audit trail report for further analysis.

Iterations: Enter the number of Iterations you want to hold back.

Trend Adjustment. Select the Trend Adjustment check box to use this option.

Trend Adjustment option lets you adjust the growth and decline phases for new or old products.

Convert negative values to zero. Select this option to change any negative values in your data to zero.

Integer Only Forecast. Select this option to use integers for forecasted values instead of decimals.

Note Data Cleansing is applied to history; Forecast Accuracy is applied to the forecast.

Forecast Techniques and Methods

QAD Demand Management has several forecast techniques available; each technique has one or more methods. Some techniques use a primary and secondary method.

Select a forecasting technique based on the specific demands of the forecast you need to generate. The Engine lets you select the forecasting methods that enhances the technique you choose. You can adjust the parameters of each method.

Table 15.1 lists the available forecast techniques and associated methods.

Table 15.1
Forecast
Techniques and
Methods

Forecast Technique	Available Forecast Methods
Distribution	ErlangB (accepts secondary method) ErlangC (accepts secondary method) Poisson (accepts secondary method)
Business Promotional Model	Event Model (accepts secondary method)
Exponential Smoothing	Adaptive ES Croston Intermittent Double ES Brown Double Holt Simple ES Triple ES Brown
Slow Moving Item	Slow Method Croston Intermittent
Expert	Procast
Moving Average	Moving Average Weighted Moving Average
New Product Curve	Logistic Curve New Product Curve Probit Curve

Forecast Technique	Available Forecast Methods
Regression	Linear Regression Multiple Regression (accepts secondary method) Polynomial Regression (accepts secondary method) Stepwise (accepts secondary method)
Seasonal Model	Box-Jenkins Holt-Winters Census X-11 (accepts secondary method)

For details, see Chapter 21, “Using Statistical Methods,” on page 281.

Configuring Forecast Methods

Forecast methods can be classified into three groups depending on the configuration:

- Methods with basic configurations
- Methods that require additional series or indexes
- Regression methods

Forecasting Methods with Basic Configurations

Table 15.2 lists basic configurations, which do not require additional series or indexes.

Forecast Technique	Available Forecast Methods
Expert	New Product Curve
Moving average	Distribution (only with ErlangB and Poisson methods)
Exponential smoothing	Seasonal Model (only with Box-Jenkins and Holt-Winters)

Table 15.2
Basic Forecasting

To select methods with basic configurations:

- 1 Click the Forecast Method tab.

Fig. 15.6
Base Unit, Forecast
Method Tab

The screenshot shows the 'Base Unit' window with the 'Forecast Method' tab selected. The window has a title bar with standard Windows controls. Below the title bar are three tabs: 'Statistics', 'Reports', and 'Forecast Method'. The 'Forecast Method' tab is active and contains the following fields and controls:

- Forecasting Technique:** A drop-down menu with 'Expert' selected.
- Primary Method:** A drop-down menu with 'ProCast' selected.
- ProCast:** A section containing:
 - Error Term:** A drop-down menu with 'Sum Squared Error' selected.
 - New Product Curve:** A section with two radio buttons:
 - ☒ Do Not Include
 - ☐ Include when num of obs < [a text box containing '6']

At the bottom of the window is a toolbar with several icons: a help icon (question mark), a printer icon, a save icon, a refresh icon, a double right arrow icon, and two buttons labeled 'Advanced' and 'ProCast!'.

- 2 Select the technique from the Forecasting Technique drop-down list.
Note If you select Procast as the primary method, the Engine selects the best method to analyze the data for the forecast.
- 3 Select the primary method from the drop-down list. If you do not select a method, the method that displays in the field is applied to the data.

Figure 15.7 displays an example of a window using the Forecasting Technique, Business Promotional Model.




Fig. 15.7
Forecast
Technique,
Business
Promotional Model


- 4 Select the parameters for the primary method.
 - 5 Select the Secondary Method from the drop-down list if the Primary Method has a secondary method.
 - 6 Select the parameters for the Secondary Method.
- Note** If you do not set the parameters, the Engine selects the best parameters for the data.
- 7 Click the Procast! button.

Primary Method Procast

When you select Procast as the primary method, the Engine analyzes the data and chooses the method that best fits the data.

In order to determine which method the system selects for each series, go to the Series Preview and Transformation screen.

Fig. 15.8
Series Preview and Transformation



The method also displays in the report output.

Fig. 15.9
Method in Report Output

Forecast method

Forecast - Holt-Winters			
Date	Forecast	Rolling Semi Cycle	Rolling Full Cycle
2006-07	110,388.30		
2006-08	136,637.98		

Forecasting Methods that Require Additional Series or Indexes

The indexes tables must be arranged in order to be valid for the Engine. Indexes must comply with a columns or rows layout. The following forecast techniques require additional indexes in order to run a forecast:

- Business Promotional Model (Event Model)
- Distribution (with ErlangC)
- Seasonal Model (with Census X-11)

Before running a forecast using the ErlangC, Event Model, or Census X-11 method, make sure that all the forecasting series are mapped to an additional series. For the event model, the additional series is the event index series.

For event index series, there are more requirements:

- You must have predefined indexes.
- The indexes must be stored in a separate table.
- The indexes must be compliant with column and row layouts.
- The indexes values must be consecutive, such as, 1, 2, 3. (1, 2, 5, is not valid). There must be enough values to forecast. This means that the amount of index values must be greater than or equal to the history + periods.

Columns

If the data that you want to use for the forecast is in Excel, you must have an additional spreadsheet with your indexes.

In the spreadsheet, the type of index values in the columns and the entries below the index label are numeric. You must start in the cell A1, and you can have as many index series as you want.

Figure 15.10 shows an example of a valid index configuration with a column layout.

	A	B
1	Optimistic	Pesimistic
2	1	1
3	1	2
4	1	3
5	1	1
6	2	2
7	2	3
8	2	1
9	4	2
10	0	3
11	4	1
12	4	2
13	4	3
14	4	1

Fig. 15.10
Columns

Figure 15.11 shows an invalid index configuration.

Fig. 15.11
Invalid Columns

	A	B
1		
2	Optimistic	Pesimistic
3	1	1
4	1	2
5	invalid	3
6	1	1
7	2	2
8	2	FXE
9	2	1
10	4	2
11	0	3
12	4.4	1
13	4	2
14	4	3.2
15	4	1

Does not start in first row

Text is not a valid value

Decimal index values not allowed

If your data is stored in a database such as Access, SQL Server, or Oracle, you must have a separate table containing only the index. The field type must be numeric. Figure 15.12 shows an example of this index.

Fig. 15.12
Index Example

	Optimistic	Pesimistic
1		1
2		4
3		1
1		4
2		4
3		2
1		2
2		2
3		3
1		3
2		4
3		4

Figure 15.13 shows an example of an invalid configuration.

	Optimistic	Pesimistic
1	1	1
2		4
3		1
1		4
2.01		4
3		2
1		2
2		2
3		3
1		3
2		4
3		4

Decimal values are not valid.

Fig. 15.13
Invalid Index

Rows

When the data is in Excel, you must have a separate spreadsheet in the index or indexes. You must define the labels that appear horizontally. For each row, the first column should be text to represent the index name. After this column you must enter the numeric index values.

Figure 15.14 shows an example of a valid index configuration in Excel.

	A	B
1	Optimistic	Pesimistic
2	1	1
3	1	2
4	1	3
5	1	1
6	2	2
7	2	3
8	2	1
9	4	2
10	0	3
11	4	1
12	4	2
13	4	3
14	4	1

Fig. 15.14
Valid Index

Figure 15.15 shows an example of an invalid index configuration example in Excel.

Fig. 15.15
Invalid Index

	A	B	C	D	E	F	G
1							
2	First	1	2	1	2	2	2
3		1	2	3	1	2	3

No Index Name for Each Series

No Labels

If the data you want to forecast is in a database, you need a table with one text field to store the index name, and as many numeric fields as you need to store the index values. See Figure 15.16.

Fig. 15.16
Sample

	Index Name	Index 1	Index 2	Index 3	Index 4	Index 5	Index 6	Index 7
	Optimistic	1	2	3	1	2	3	1
	Pesimistic	1	1	4	4	1	1	4

If you try to provide an incorrect index table, the Engine detects it and reports an error. The forecast fails if you do not provide the index values for the history and period in the data.

To set up methods that require additional series or indexes:

- 1 Select the Forecast Method tab.
- 2 Select a technique from the Forecasting Technique drop-down list.
- 3 Select ErlangC, Event Model or Census X-11 method from the Primary Method drop-down list.
- 4 Select a secondary method from the Secondary Method drop-down list if the primary method permits another selection.
- 5 Set the parameters for both the primary and secondary methods. For example, if you selected ErlangC from the Primary Method list, the window shown in Figure 15.17 is displayed.

Note Remember to set the parameters for the forecast. If you do not set them, the Engine selects the best parameters for the data.




Fig. 15.17
ErlangC

- 6 Select the additional series or indexes required for the method.
 - 7 Select the table name from the Additional Table drop-down list. This table contains the additional series or indexes.
 - 8 Click the Setup Series button.
- Note** Depending on the model selected, the button can be Setup Series for ErlangC, an Event Flags button for Event Modeling, or a Setup Trading Day button for Census X-11.
- 9 The window in Figure 15.18 is displayed.




Fig. 15.18
Setup Series

- 10 Select the series from the Available Series list box.
- 11 Select the series in the Description column to link to the indexes.

Note You can use the Select All button to select all the series that appear on screen. Use the Deselect All button to deselect all selected series.

12 Click OK.

13 Click the Procast! button after configuring the Forecast Method tab.

Regression Methods

Four regression methods are available for the Regression Forecasting Technique.

- Linear regression
- Multiple regression
- Polynomial regression
- Stepwise

The Linear regression method applies to each series in a scenario. The other methods require one dependent series and several independent series. The Engine uses the first series in a scenario as the dependent series, and other series in the same scenario as the independent series.

Note When you select a regression method, the Remove Zero Values option is not available. For details, see “Group Forecasting” on page 210.

To select regression methods:

- 1 Click the Forecast Method tab.
- 2 Select Regression from the Forecasting Technique drop-down list. The window in Figure 15.19 is displayed.




Fig. 15.19
Forecast Method,
Regression
Technique

- 3 Select the regression method from the Primary Method from the drop-down list.

Note If you select Multiple Regression, Polynomial Regression, or Stepwise, a Secondary Method field is used to identify the independent series.

Multiple Regression

When the Primary Method is multiple, the window in Figure 15.20 is displayed.

Fig. 15.20
Forecast Method,
Multiple
Regression
Technique

The screenshot shows the 'Base Unit' dialog box with the 'Forecast Method' tab selected. The 'Forecasting Technique' dropdown is set to 'Regression'. Under 'Primary Method', 'Multiple Regression' is selected, and the 'Dependent Series' is 'First Series'. Under 'Secondary Method', 'Adaptive ES' is selected, and the 'Beta' field is empty. The bottom bar contains icons for help, data, and a 'ProCast!' button.

- 1 From the Secondary Method drop-down list, select the statistical method you want to apply to the independent series.
- 2 Select the parameters for the primary method.
For example, if you selected Multiple Regression from the Primary Method list, the control in Figure 15.21 appears. This parameter cannot be changed.

Fig. 15.21
Method Parameter

The screenshot shows the 'Primary Method' section of the 'Base Unit' dialog box. The 'Primary Method' dropdown is set to 'Multiple Regression'. The 'Multiple Regression: Dependent Series' is 'First Series'.

Note Remember that if you do not set the appropriate parameters, the Engine selects the best parameters for the data.

Polynomial Parameters

Selecting the Polynomial Regression option from the Primary Method drop-down list displays the window shown in Figure 15.22.

The screenshot shows the 'Base Unit' dialog box with the 'Forecast Method' tab selected. The 'Forecasting Technique' is set to 'Regression'. The 'Primary Method' is set to 'Polynomial Regression'. The 'Polynomial Regression' section shows 'Dependent Series' as 'PRICE', 'Max. Degree' as 4.00, 'Partial F-Value' as 4.00, 'Min. F-Value In' as 2.00, and 'Max F-Value Out' as 1.00. The 'Secondary Method' is set to 'ProCast'. The 'ProCast' section shows 'Error Term' as 'Sum Squared Error'. The 'New Product Curve' section has 'Do Not Include' selected. The 'Include when num of obs <' option is also visible with a value of 6. The bottom of the dialog has buttons for 'Advanced' and 'ProCast!'.

Fig. 15.22
Forecast Method,
Polynomial
Regression
Technique

- 1 Select the appropriate primary parameters.

Dependent Series. To run the Multiple Regression, Polynomial, or Stepwise methods, one series must be the dependent series. The remaining series are considered independent series. Therefore, in order to run any of these methods, you need at least two series.

When your data organization is in rows, the first series in the table (or Excel spreadsheet) is always considered the dependent series. When your data is in columns, the order is defined when you configure your data using the Series button.

Max Degree. Indicate the maximum degree that the polynomial regression can have.

Partial F-Value. Enter the value used to determine which degree the polynomial regression should use.

Min F-Value In/Max F-Value Out. Enter the F-value that is used to determine which series can be included in the polynomial regression model.

- 2 Select the secondary method from the Secondary Method drop-down list.
- 3 Enter the appropriate Secondary Method parameters, and then click the ProCast! button.

Stepwise Parameters

Selecting the Stepwise Method option from the Primary Method drop-down list displays the window shown in Figure 15.23.

The screenshot shows the 'Base Unit' dialog box with the 'Forecast Method' tab selected. The 'Forecasting Technique' is set to 'Regression'. The 'Primary Method' is set to 'Stepwise'. Under the 'Stepwise' section, the 'Dependent Series' is 'First Series'. The 'Stepwise' radio button is selected. The 'Min. Series' is 4.00, 'Min. F In' is 4.00, and 'Min. Lag' is 0. The 'Max. Series' is empty, 'Max. F Out' is 3.00, and 'Max. Lag' is 0. The 'Lag Dependent' checkbox is unchecked. The 'Secondary Method' is set to 'ProCast'. Under the 'ProCast' section, the 'Error Term' is 'Sum Squared Error'. Under the 'New Product Curve' section, the 'Do Not Include' radio button is selected. The 'Include when num of obs <' option is set to 6. The bottom bar contains buttons for help, data, forecast, and 'Advanced'/'ProCast!'.

Fig. 15.23
Forecast Method,
Stepwise Technique

- 1 Select the appropriate Primary Parameters.

Stepwise, Forward. Indicate the direction to search the stepwise model.

Min/Max Series. Indicate the range of the number of series that the stepwise model can have.

Min F In/Max F Out. Indicate the F-value used to determine which series can be included in the stepwise model.

Lag Dependent. Indicate whether the stepwise model includes a lag dependent series.

- 2 Select the secondary method from the Secondary Method drop-down list.
- 3 Enter the appropriate secondary method parameters, and then click the ProCast! button.


Group Forecasting

The Engine lets you configure the forecast to create groups of data. This is called aggregated forecasting.

To configure a group forecast:

- 1 Create a new scenario with a data source containing the series for the group. You can group by selecting one or more fields.
- 2 Click the Group By tab.

Fig. 15.24
Base Unit, Group
By Tab



- 3 Select the Enable Configuration check box. The fields on the window become active.
- 4 Click the Group Configuration button to specify the fields for the groups. The Fields Selection for Groups window is displayed.




Fig. 15.25
Field Selection for
Groups

- 5 Move the fields from the Available column to the Selected column to designate group levels, and then click OK. Use the up and down arrows to move the appropriate field names to a higher or lower position.
- 6 The Group By tab is displayed. Select the appropriate options.

Include Grand Total. Select this option if you need a series in your data that is the grand total, or sum, of all the series.

Include Percentages. This field is available only when the Top Down option is selected. Select Include Percentages if you need to retrieve the percentage used to allocate the forecasted value from higher-level series to lower-level series.

Computation Type. Select the type of calculation to use in the forecast.

Top Down: The higher-level series are forecast. Then the forecasted values in this level are allocated to lower-level series.

Bottom Up: Lower-level series are forecast first, and then higher-level series use the sum of the lower-level forecasted values as their forecasted values.

Allocate By. Specify how to calculate the percentage. These options are available only if the Top Down computation type option is selected.

History: The historical values in lower-level series are used to calculate the percentage of each lower-level series of their higher-level series.

Then the percentage is used to allocate the forecasted values of the higher-level series to the lower-level series as their final forecasted values in the lower level.

Forecast: The direct forecasted value (not the final allocated value) in the lower-level series are used to calculate the percentage

Delete Aggregated Series. Indicate whether you want the aggregated series deleted, and if so, when to delete the aggregated series.

No: The aggregated series are not deleted.

After Forecast: The aggregated series are deleted after grouping forecasting.

If you select the No or After Forecast options, the following cannot be used:

- Distribution
- Business Promotional Model
- Census X-11
- Regression (with After Forecast)
- Holdback Evaluation option on the Accuracy tab
- Same As functionality
- Remove Leading Trailing zero values options in Data Cleansing

Before Forecast: The aggregated series are deleted after loading the series with the group hierarchy information. The detail level series are forecast as usual.

Working with Reports

This section describes QAD Demand Management reports.

Introduction **214**

Understanding Pivot Reports **214**

Generating Reports **219**

Running and Viewing Reports **225**

Introduction


The Engine uses two types reports:

- Report templates or predefined reports
- Free-form pivot tables and Excel

Predefined reports are accessed using predefined templates that output data. You can use the Batch Scheduler to run reports, or the Base Unit to run reports interactively.

When using the Excel-based pivot reports, the Engine provides Wizards that let you define reports on the fly or use predefined templates created by the system administrator. Figure 16.1 illustrates the reporting architecture within the Engine:

Fig. 16.1
Report Architecture



The Engine also supports customer-specific reporting solutions such as Cognos, Brio, Business Objects or Proclarity, among others. These reporting solutions connect directly to the Engines database and provide additional reporting capabilities.

Understanding Pivot Reports

A pivot report is a special type of report.

Note To generate a pivot report, Excel 2000 or later must be installed.

Use one of the following methods to create a pivot report:

- Open an existing report. You can refresh the report with current data.

- Create a new report from an existing template in the template directory. The current data is used to create the new scenario.
- Create a new report using the Wizard. The current data is used to create the new scenario.

The report is created from the Engine, which makes the data range more manageable. Also, some security can be applied to the report. Using the assigned permissions, you can select data from one of the following:

- Current scenario
- All scenarios visible to this user
- Ignore, Do Not Report, or Ready to Export

Using the pivot report design provides the following benefits:

- Built-in pivot table and other reports
- Reduces the time required to create a pivot report. The Wizard only loads the data necessary for the report.

To create a pivot report:

- 1 From the Viewer menu, choose Reports|Open Pivot Report. The window in Figure 16.2 is displayed.




Fig. 16.2
Open and Refresh
Pivot Report

- 2 Click the button to the right of the Report Name field. A window displays allowing you to browse report names.
- 3 Select a report and click Open. After a report is selected, you can
 - Refresh existing components

- Adjust parameter page and data range


Note You can always refresh a report created from a template or Wizard.

Using an Existing Template

If you copy an existing template to create a new report, you must enter a source file name, target file name, and a data source. When you click Create, the template file is copied to the source file. Then the new data source is used for the pivot report in the new file and the data is refreshed.

- 1 Select Reports|Pivot Report from Template from the Viewer menu. The window in Figure 16.3 is displayed.

Fig. 16.3
Create Pivot Report
from Template



The data source includes the SeriesID, since internally SeriesID defines the data range.

- 2 Click the button to the right of the Template field and select a template.
- 3 Click the button to the right of the Report Name field, and then enter a report name. Click Open. The report name cannot be the same as the template name.
- 4 Select the data source.
- 5 Select the appropriate options.
- 6 Click the Create button.

A special hidden worksheet, Report Parameters, stores the parameters for each report in the workbook. The worksheet uses the following parameters:

Report Name. The name of the report.

Report Type. Pivot report or external table.

Data Source. A Select statement to retrieve data, which should include SeriesID in the selected fields.

Where clause. A changeable range that is used to refresh the data.

Using the Pivot Table Wizard

The Pivot Table Wizard requires the following:

- File name/path
- Data source
- Report type

- 1 Select Reports|Pivot Table using the Wizard from the Viewer menu. The window in Figure 16.4 is displayed.





Fig. 16.4
Create Pivot Report
from Wizard

- 2 Click the Basic tab.
- 3 Select the report type from the Report Type drop-down list.

- 4 Click the button to the right of the Report Name field, and then select a Report Name.
- 5 Select the data source.
- 6 Select flag options if appropriate.
- 7 Click the Layout tab. The window in Figure 16.5 is displayed.

Fig. 16.5
Create Pivot Report
from Wizard,
Layout Tab



The layout uses the same user interface used to select the pivot table field list. If the following conditions are met, entries display in the View drop-down list:

- Starts with vw_PVT
- Has a numeric field SeriesID

An additional view is used to support the advanced pivot report. The view includes the following:

- All description fields
- MAPE
- RSqr in the view vw_PVT_VIEW_V29

For data fields, the view includes the following:

- Hist
- Original
- FValue

- UserValue01 to UserValue10

The view also includes the Error and Abs Error between Hist and other data fields.

8 Click the Format tab.




Fig. 16.6
Create Pivot Report
from Wizard,
Format Tab


Once you select the reporting options, the report is generated within Excel.

Generating Reports

Use the Reports tab to select the report options for your forecast. Click the tab of the type of report you want to generate.

Note Before you can view any reports, you must run a forecast. The default method is Procast, which automatically selects the best method based on the data selected.

Fig. 16.7
Base Unit, Reports
Tab



On the Report tab, click the tab of the type of report you want to generate.

Report Types

QAD Demand Management has three types of report:

- Standard
- Executive
- Audit

Standard

The standard report is designed for handling large volumes of data quickly after applying MAPE, R-Squared Value, and Standard Deviation statistics. The standard report produces a comparison report that lists actual and forecasted values side by side, and shows all the selected statistics.

Select the appropriate options:

Output Type. Select whether you want to view the data in the report in rows or columns. This value defaults to rows.

Show Data. This option is always selected to indicate that the reports will show your data.

Show Charts. Select this option to view charts in the standard report.

Executive

The executive report is designed to answer the question, “What is the risk of accepting this forecast?” It visually represents the answer in easy-to-read graphs, charts, and tables. It also applies the common statistics MAPE and R-Squared values.

The executive report includes complete listings of monthly, quarterly, and annual forecast values, and attaches fields for additional comments and notes.

Select the appropriate options:

History. Select this option to include historical values for your data.

Fitted. Select this option to include a plot of the fitted forecast values on the graph.

Show Charts. Select this option to represent each series with a chart.

Forecast. Select this option to include forecasted values for your data.

Errors. Select this option to include a plot of the forecast error values on the graph.

Lower/Upper Limit. Select this option to include a plot of the upper and lower limits on the graph.

Safety Stock. Select this option to include a plot of the safety stock values on the graph.

Audit

The audit trail report is the most detailed analysis option for the forecast. It displays everything contained within the executive report, the selected statistics, and some optional tables. Use the audit trail report to justify a forecast with statistics.

Select the appropriate options:

Forecast Results. Select this option to be able to view forecasted values in the report.

Include Executive. Select this option to include all results in the executive report.

Fitted Values Sample. Select this option to include a table of the forecast fitted values in the report.

Out of Sample Table. Select this option to include a table of holdback analysis.

Excel and File Report Selection

After selecting the type of report you want to generate, indicate the applications you want to use to display the reports.

Standard. Displays the standard report in an Excel spreadsheet.

Audit. Displays the audit report in an Excel spreadsheet.

Executive. Displays the executive report in an Excel spreadsheet.

Text. Displays the selected report also in a .txt file. If a text report is open in Excel, it can be regarded as an Excel standard report with rows options. The pipe symbol (|) is used as a delimiter in the text report.

XML. Select this option to display the selected reports also in XML file format. The report has the following structure:

- Scenario information, which includes general information, the scenario configuration, and dates for each observation.

- A section for each group. When a group configuration does not exist, all series in the scenario are one group.
- An item for each series in one group. Each item has three subitems: historical values, forecasted values, and statistics.

Crystal Reports

The Crystal Report produces detailed and comprehensive visual representations of the forecast you have run. Crystal Reports run faster than the other reporting methods and do not require Excel or other reporting tools to be installed on your system.

Figure 16.8 shows the standard Crystal Report screen. Use the buttons on screen to display the information you want to see.




Fig. 16.8
Standard Crystal
Report

Close	Closes the screen.
Print Report	Prints the report.
Refresh	Refreshes the screen. You are returned to the first series.
Toggle Group Tree	Toggles the display of the group tree on the left.
Percentage	Controls the size of the report view.
Go to First Page	Displays the first page of the report.

Go to Previous	Displays the previous page of the report.
Go To	This window displays the page number of the report you are currently viewing.
Go to Next	Displays the next page of the report.
Go to Last	Displays the last page of the report.
Stop Loading	Stops loading the data of a new report.
Search Text	Displays a Search dialog box. Enter the text or phrase you want to search for.

You can generate four types of Crystal Reports.

- **Standard Crystal Report.** Displays the history and forecast of each series.
- **Executive Crystal Report.** Is less detailed than other Crystal reports, but provides easy-to-follow graphical views of the data and a summary with forecast information. Recommended for presenting data to non-forecasting stakeholders within an enterprise.
- **Data Analysis Crystal Report.** Shows dates, history, fitted values, and the following:
 - **Error.** The difference between history and fitted values.
 - **Series percentage change.** The percentage of variation between one observation and the previous one.
 - **Forecast Percentage Change.** The percentage of variation between one fitted value and the previous one.
 - **Cumulative Error and Cumulative MAPE.** These values show the stability of the error.
 - **Average, Maximum, Minimum, Standard Deviation, Variance, and Median** for each column.
- **DRP Crystal Report.** The results of the DRP Crystal report appear on the Demand Data tab on the Series Preview and Transformations screen. See “Viewing Forecast Data” on page 74.

Figure 16.9 illustrates an Executive Report in Crystal format.

Fig. 16.9
Crystal Executive
Report

Forecast - Holt-Winters					
Date	Forecast	Rolling Semi Cycle	Rolling Full Cycle	Lower Limit	Upper Limit
2006-07	950.29			753.74	1,146.84
2006-08	1,090.30			790.07	1,390.54
2006-09	1,158.56	3,199.15		749.41	1,567.71
2006-10	1,363.43			843.41	1,883.45
2006-11	1,135.77			503.95	1,767.59
2006-12	1,012.04	3,511.25	6,710.39	267.92	1,756.17
2007-01	1,127.70			270.96	1,984.45
2007-02	1,202.73			233.18	2,172.29
2007-03	1,326.36	3,656.80		243.85	2,408.88
2007-04	1,369.96			174.40	2,565.53
2007-05	1,472.52			163.83	2,781.22
2007-06	1,395.41	4,237.90		0.00	2,817.29
Average	1,217.09	3,651.27	6,710.39	416.23	2,020.16
Max	1,472.52	4,237.90	6,710.39	843.41	2,817.29
Min	950.29	3,199.15	6,710.39	0.00	1,146.84
Summary Comments:					
The forecast has average error of:				14.48	
The forecast has Standard Deviation of:				237.83	
The forecast exceeds accuracy of simple average by:				51.47	

Running and Viewing Reports

Select the type of data you want to display in the report. From the Display Report Data In drop-down list, select Units or Dollars.

After selecting the report options for your forecast, you can run and view the reports you want to generate.

Running Reports

To run or generate the reports, click Run Report on the Report tab. The report is generated and is displayed.

Note You must have a default printer installed on your system in order to run Excel reports.

Viewing Reports


Report are generated in the format specified. If you select more than one type of report, each report is created within its own Excel workbook.

Text and XML reports appear in a Notebook file.

Figure 16.10 show a report in Excel format. For details on Crystal Reports, see “Crystal Reports” on page 223.

Fig. 16.10
Standard Report

	A	B	C	D	E	F	G	H	I	J
1	Series	2004-05	2004-06	2004-07	2004-08	2004-09	2004-10	2004-11	2004-12	2005-01
2	[Series 01] - Actual	198.76	140	657	1007	1030	1188	832	701	82
3	[Series 01] - Forecast(Holt-Winters)	218.1431	39.11201	0	709.1765	1080.396	1285.926	998.472	741.35	833.98C
4	[Series 02] - Actual	198.76	150	0	0	0	0	0	0	0
5	[Series 02] - Forecast(Holt-Winters)	207.0626	248.8671	58.70408	58.73152	48.47728	40.67301	0	37.74425	46.019C
6	[Series 03] - Actual	198.76	150	0	0	0	0	0	0	1C
7	[Series 03] - Forecast(Holt-Winters)	198.76	209.0279	135.8441	15.25278	1.698013	5.91E-02	0	0	10.8795
8	[Series 04] - Actual	198.76	150	0	0	0	0	0	0	0
9	[Series 04] - Forecast(Holt-Winters)	198.76	226.6244	108.2016	0	25.08913	15.24767	0	0	36.307E
10	[Series 05] - Actual	990.81	900	104	104	103	98	87	88	E
11	[Series 05] - Forecast(Double Exponential)	970.6147	971.0836	878.623	64.63627	65.5503	65.41988	61.17639	50.77602	52.6405
12	[Series 06] - Actual	990.81	900	127	129	129	130	130	131	13
13	[Series 06] - Forecast(Double Exponential)	971.702	972.1377	879.6829	89.52072	92.4209	93.25495	95.09278	95.88871	97.6892
14	[Series 07] - Actual	990.81	900	0	0	0	0	0	0	0
15	[Series 07] - Forecast(Double Exponential)	973.6469	973.9752	881.7504	0	0	0	0	0	0
16	[Series 08] - Actual	990.81	900	0	0	0	0	0	0	0
17	[Series 08] - Forecast(Double Exponential)	975.1762	975.4561	883.2955	0	0	0	0	0	0
18	[Series 09] - Actual	1231.87	1000	745	391	1131	960	1030	813	61
19	[Series 09] - Forecast(Holt-Winters)	1244.875	1158.396	481.6326	479.8662	938.272	931.893	1205.724	853.3242	732.682
20	[Series 10] - Actual	1231.87	1000	173	173	173	174	174	174	17
21	[Series 10] - Forecast(Double Exponential)	1200.112	1201.093	963.0109	111.6042	113.5009	115.3391	118.1514	119.8768	121.54E
22	[Series 11] - Actual	318.25	300	136	0	169	192	134	88	11



Chapter 17

Grouping and Allocation

This section describes how to combine hierarchies and forecast strategies by using the QAD Demand Management Hierarchy Manager.

Introduction **228**

Creating a Hierarchy **229**

Introduction

Grouping and allocation—sometimes referred to as multilevel forecasting—provide a powerful mechanism for forecasting and planning. Multilevel forecasting lets you group information that is by itself unstable, but when combined become stable and forecastable. This is often seen in seasonal products. At the customer-product level, individual items appear erratic. However, when summed up to the customer level, a forecastable pattern emerges.

Bottom-up aggregation aggregates the detailed series to higher levels within a business hierarchy.

Top-down allocation lets you roll detail information up to a higher level to allow a pattern to emerge. You can then push the results back down on a percentage basis so that the lowest level of detail is properly forecast. When higher-level information is pushed down into the detail, the forecast is said to be developed top down.

Middle-out allocation is used when at the top level there is not enough differentiation and it may appear to be flat. However, at the division level there might be enough differentiation to build a strong seasonal pattern. In this case, forecasting at the top level would not yield the most accurate results, so a middle-level forecast must be developed and allocated down to the detail, and then rolled up to the top level.

The Engine provides two ways to group and allocate.

- The Hierarchical Manager lets you examine data at a higher level or even across level, change the data, and allocate changes down to the detail data. The Hierarchical Manager is available in the Viewer so you can dynamically create aggregates and roll up, down, or across. The Hierarchical Manager allocates on additional fields other than forecasting.
- The Group by Tab in the Engine base is primarily used for pivoting data as it is loaded into the Engine database. For example, data that is originally stored in columns is loaded into the Engine and stored into rows or series.

Hierarchy Management

Creating a Hierarchy

To create a hierarchy:

- 1 Open the Viewer. The Scenario Selection window is displayed.
- 2 Select a scenario to create the hierarchy and click OK. The Scenario window is displayed.
- 3 Select Forecast|Hierarchy Manager from the Viewer menu. The Hierarchy Management window is displayed.




Fig. 17.1
Hierarchy
Management

Use the buttons on this window to complete the following:

Add. Create a new hierarchy.

Edit. Adjust an existing hierarchy.

Delete. Erase a hierarchy.

Apply. Activate the defined hierarchy.

- 4 Click the Add button. The Hierarchy Definition window is displayed.

Fig. 17.2
Hierarchy
Definition

The screenshot shows a window titled "Hierarchy Definition" with a standard Windows-style title bar (minimize, maximize, close buttons). The window contains a tabbed interface. The tabs are: "Top Down", "TD (Advanced)", "Bottom Up", "Name", "Grouping", and "Rebalance". The "Name" tab is currently selected and active. Within the "Name" tab, there is a section titled "Name" in bold. Below this title, there is a text input field labeled "Name" containing the text "std". Below the text field is a larger text area labeled "Description". At the bottom of the "Name" tab, there is a checkbox labeled "Sharable" which is currently unchecked. At the bottom right of the window, there are two buttons: "Accept" and "Cancel".

The window has six tabs.

- Name
- Grouping
- TD (Advanced)
- Bottom Up
- Top Down
- Reference

Name Tab

- 1 Enter a name for the hierarchy in the Name field.
- 2 Select the Sharable check box to allow other users to use this hierarchy.
- 3 Click Accept.

Grouping Tab

- 1 Click the Grouping tab to display the following window.





Fig. 17.3
Hierarchy
Definition,
Grouping Tab

- 2 In the Grouping tab window, click the Group Configuration button. The General Selection window is displayed.
- 3 Select the Available fields in the hierarchy order that the user requires and click the single right arrow button.
- 4 Click OK. The Hierarchy Grouping displays in the Hierarchy Definition window (Figure 17.1 on page 229).

Rebalance Tab

- 1 Click the Rebalance tab. The window in Figure 17.4 is displayed.

Fig. 17.4
Hierarchy
Definition,
Rebalance Tab



The Engine supports two rebalancing methods:

On-Fly. Rebalancing occurs immediately following any adjustments.

On-Demand. The user can rebalance on command by clicking the rebalance icon.

- 2 Click the Base Level drop-down list.
Base Level specifies where the changes made are allocated up/down/middle out to other levels. The changes made at other levels are overwritten but the original values remain.
- 3 Select the Allow Change Base Level When Rebalance option to cause the Viewer to rebalance using the lowest level of detail for allocation (bottom-up).

TD (Advanced) Tab

- 1 In the TD (Advanced) tab, the Top Down Fields list displays all the fields that you can allocate down during the rebalance.




Fig. 17.5
Hierarchy
Definition, TD
(Advanced) Tab

- 2 Click Change to change the list. The General Selection window is displayed.
- 3 Select the fields from the Available list and click the single right arrow to add the fields to the Selected list. Click OK.
- 4 Use the Other Opinion Line Allocate By section for all other fields except Value and FValue.


Itself. The allocation weight for the additional fields, except Value and Fvalue, uses their own weights. For example, if UserValue01 is added to the list, the allocation weight for UserValue01 uses UserValue01 as opposed to the fixed fields (Value and Fvalue) allocation weights.

Fixed Field. The allocation weight for fields except Value and FValue use the field selected in the drop-down list.

Bottom Up Tab

- 1 In the Bottom Up tab, the Bottom Up Fields list identifies the fields that are updated by a bottom up allocation during rebalancing.

Fig. 17.6
Hierarchy
Definition, Bottom
Up Tab



- 2 Click the Change button to select additional fields.
- 3 Select the Same Settings as Top Down option if you want the Engine to use the settings of Top Down for Update Negative to Zero, Update Historical Data, and Overwrite Adjusted Forecast for Bottom Up. In addition, a series marked as Do not Forecast or Forecast Completed is not updated.

If you do not select Same Settings as Top Down, the Use Own Settings options are enabled. Update Negative to Zero, Update Historical Data, and Overwrite Adjusted Forecast are advanced features and have the following effects on the forecast.

Update Negative to Zero. Any negative values in the Top Down Fields on the TD (Advanced) tab are set to zero.

Update Historical Data. Both historical and future values are updated by allocation; otherwise, only future values are updated.

Overwrite Adjusted Forecast. If selected, manual adjustments made to the forecast are affected by allocation. If not selected, the manual adjustments remain.

Top Down Tab

In the Top Down tab, the Allocation Type and Seasonal Option sections set allocation weight for the Value and FValue fields only. For other fields, the allocation weight is set in the TD (Advanced) tab.




Fig. 17.7
Hierarchy
Definition, Top
Down Tab

- 1 Select the appropriate options.

Allocation Type. Determines which field is used as the allocation weight. If Historical is selected, the historical percentage values of field Value are used as the allocation weight. If Forecast is selected, the future percentage values of field FValue are used as the allocation weight.

Seasonal Option. Determines which periods are used for allocation weight. When combining the Allocation Type and the Seasonal Option, the effect is more complex. Refer to Table 17.1 when combining these two options.

Table 17.1
Seasonal Options
and Allocation
Types

Allocation Type	Seasonal Option	No. of Periods	Value of No. of Periods	Effect
Historical	Flat	n	Blank. Engine Viewer uses 1 for the # of periods	The allocation weight is the summation of field Value covering n periods, from the period before the current period back to the history.
Historical	Seasonal	n	Blank. Engine Viewer uses 1 for the number of periods.	The allocation weight is different for different periods. For period m, the allocation weight is the summation of field Value whose period k satisfies $k \% n = m \% n$, and k is in historical periods.
Forecast	Flat	n	Blank. Engine Viewer uses 1 for the number of periods.	The allocation weight uses summation of field FValue covering n periods, from the current period to the future.
Forecast	Seasonal	n	Blank. Engine Viewer uses 1 for the number of periods.	The allocation weight is different for different periods. For period m, the allocation weight is the summation of field FValue whose period k satisfies the $k \% n = m \% n$, and k is in future periods.

- 2 Click Accept to save the hierarchy definition. The Hierarchy Management window is displayed.

Editing a Hierarchy

You can edit a hierarchy at any time.

- 1 Click the Edit button on the Hierarchy Management window (Figure 17.1 on page 229). The Hierarchy Definition window displays (Figure 17.2 on page 230).
- 2 Click the tab for the part of the hierarchy definition that requires editing.
- 3 Make the changes and click Accept. If the hierarchy has been applied to a scenario and you change the Grouping, you are prompted to confirm that you want to change the hierarchy setting.

Deleting a Hierarchy

To delete a hierarchy:

- 1 From the Hierarchy Management window (Figure 17.1 on page 229), select the hierarchy you want to delete and click Delete.
- 2 If the hierarchy was applied to a scenario, a message displays informing you that the hierarchy definition is in use. Click OK. The hierarchy is not deleted from the hierarchy list.

Applying a Hierarchy to a Scenario

To apply a hierarchy to a scenario:

- 1 From the Hierarchy Management window (Figure 17.1 on page 229), select the hierarchy you want to apply and click the Apply button.
After the Engine Viewer applies the hierarchy to the scenario, you are prompted to view the results.
- 2 Click OK. The tree structure is displayed in the Viewer. After applying the hierarchy, new scenarios are created automatically to hold aggregated series.

Note There is one scenario for each level except the detail level. The new scenario names are the name of the scenario appended by a hyphen (-), then appended by the level number. For example, the name of the detail scenario is ABC.

The applied hierarchy has three levels, detail level, level 1, and level 2 (its grouping has two fields). After applying the hierarchy, the scenario names are the following:

- Scenario ABC is the detail level.
 - Level 1 scenario is ABC -1.
 - Level 2 scenario is ABC -2.
- 3 Close the Hierarchy Manager.

Removing a Hierarchy

To remove a hierarchy:

- 1 Select Forecast|Remove Hierarchy from the menu. A removal confirmation message is displayed.
- 2 Click Yes. The tree structure is removed. After removing the hierarchy, the created scenarios are deleted, and the detail scenario remains. You are prompted to view the results.

Adjusting a Forecast

Following adjustments, if the hierarchy applied to the scenario has the Activate On-Fly Rebalance option selected, the new values of the series are allocated to series at other levels when the series is saved (explicitly or by moving to other series), and the base level is the level the series belongs to. Otherwise, the new values are not allocated until you rebalance the tree explicitly, as in Rebalancing Hierarchy.

Rebalancing a Hierarchy

To rebalance a hierarchy:

- 1 Open the Engine Viewer with the scenario that has a hierarchy applied.
- 2 Select Forecast|Rebalance from the menu. If the hierarchy has the Allow Change Base Level When Rebalance option selected, an additional window is displayed.
- 3 Click OK to close the Advanced Option dialog box, and then click the ProCast! button to load the data.

Life Cycle Forecasting

This section describes using QAD Demand Management profiles to manage product life cycles.

Introduction **240**

Creating a Profile Scenario **240**

Using the Life Cycle Forecast **244**

Introduction

Products have a life cycle: companies launch new products, the products become mature, and finally the product is retired. Statistical forecasting models cannot always capture the right life cycle pattern. The Engine life cycle forecast features lets you define a life cycle pattern.


In the Engine, life cycle information is loaded into a scenario, which is defined as a Profile scenario. Each series in this scenario is a profile, which defines the life cycle pattern. Different types of profiles—for example, introduction, seasonal patterns, and retirement—can be in the same Profile scenario. When applying life cycle forecasting, you can use different profiles for different forecasting periods. You can define as many Profile scenarios as required.

Creating a Profile Scenario

To create a Profile scenario:

- 1 From the main window, click the Base Unit icon.

Fig. 18.1
Base Unit



- 2 Click the Scenario Management button.




Fig. 18.2
Scenario
Management

- 3 Click the New button. Enter a scenario name and click OK.
- 4 Click Edit to open a scenario for the profile scenario.




Fig. 18.3
Edit Scenario

- 5 Select Profile for the appropriate scenario attribute from the User Field 1 drop-down list.
- 6 Click OK and close the dialog box.
- 7 Click the arrow next to Data Source Name to choose the data source to load with its selected corresponding table/query information.
- 8 In the window in Figure 18.4, enter 0 in Forecast field.

Fig. 18.4
Data Capture

The screenshot shows the 'Base Unit' dialog box with the 'Data Capture' tab selected. The 'Statistics' and 'Reports' tabs are also visible. The 'Data Capture' section includes the following fields and controls:

- Scenario:** A dropdown menu set to 'Profile Scenario'.
- Data Source Name:** A text field containing 'ExcelExamplesFXE.xls'.
- Tables/Queries:** A dropdown menu set to 'Adaptive Exp# Smoothing\$'.
- Data Organization:** Two radio buttons, 'Columns' and 'Rows', with 'Rows' selected.
- Seasonality:** A numeric input field.
- Forecast:** A numeric input field set to '0'.

At the bottom of the dialog, there is a toolbar with several icons (help, print, save, etc.) and two buttons: 'Advanced' and 'ProCast!'.

- 9 Click the Advanced button.
- 10 Click the Cleansing tab. Make sure none of the check boxes in the Removing Zero Values section are selected.





Fig. 18.5
Clearing Zero
Values

- 11 Click the Accuracy tab. Make sure the Convert negative values to zero option and the Integer Only Forecast option are not selected.

Fig. 18.6
Clearing Accuracy
Settings



The image shows the 'Advanced Options' dialog box with the 'Accuracy' tab selected. The dialog has four tabs: Configuration, Cleansing, Accuracy, and Reload Options. The 'Accuracy' tab contains three main sections: 'Tracking Signal', 'Holdback Evaluation', and 'Trend Adjustment'. In the 'Tracking Signal' section, the 'Apply Tracking Signal' checkbox is unchecked, 'Iterations' is set to 4, 'Signal Type' is 'Average (default)', 'Under Forecast Percentage' is 50, and 'Over Forecast Percentage' is 50. In the 'Holdback Evaluation' section, the 'Use Holdback' checkbox is unchecked and 'Iterations' is 4. In the 'Trend Adjustment' section, the 'Trend Adjustment' checkbox is unchecked, 'Growth Periods' is 12, and 'Decline Periods' is 12. At the bottom of the dialog, there are two unchecked checkboxes: 'Convert negatives values to zero' and 'Integer Only Forecast'. An 'Ok' button is located at the bottom right.

- 12 Click OK to close Advanced Option dialog box, and then click ProCast! to load the data.

Using the Life Cycle Forecast

To use the Life Cycle function:

- 1 Open the Viewer. Open the scenario that uses Life Cycle Forecast.
- 2 Choose Forecast|Life Cycle Forecast. The Profile Forecasting window is displayed.




Fig. 18.7
Profile Forecasting

- 3 Choose which series is affected by the Life Cycle Curve.

Selected Series. Choose this if you want the profile to affect the selected series only.

Whole Tree. Choose this if you want the forecast to affect the whole tree (all the series in the hierarchy tree). This option is enabled only when a tree is viewed (a hierarchy has been applied).

Scenario. Choose this and select a scenario from the drop-down list if you want the forecast to affect only the series in one scenario in the hierarchy tree.

- 4 Change the Affected Data Fields by using the Change button. This option determines which data fields (opinion lines) are affected by the Life Cycle Forecast.
- 5 Select a profile scenario from the Profile Scenario drop-down list. If no values display in the drop-down list, the UserField1 attribute of the profile scenario is probably not set to Profile.
 - a Select all the shapes you want to apply from the Profile drop-down list, one by one.
 - b Select the Actv check box for each selected profile series. If you do not check the Actv check box, the curve in that line is not applied. At most, six curves can be applied.

- 6 Select the Start and End period of the series that will be forecast. The start period and the end period determine the period range where the shape is applied. You can apply multiple profiles to a single series by selecting start and stop dates that do not overlap.
- 7 Select the start date of the profile from which the profile data is used. Usually the start period for profile series matches the start period of the series.
- 8 Select or clear the Offset option:
 - If only the value in the start period of the series is used for life cycle forecasting, clear this option.
 - If all the values within the start and end period of the series are used for life cycle forecasting, select this option.

The following table illustrates the effect of the Offset option.

Date	2004-01	2004-02	2004-03	2004-04
Observation Value (before applying Life Cycle)	100	200	150	250
Profile Value	1	0.9	0.8	0.7
Observation Value (when Offset is not selected)	$100 = (100 * 1)$	$90 = (100 * 0.9)$	$80 = (100 * 0.8)$	$70 = (100 * 0.7)$
Observation Value (when Offset is selected)	$100 = (100 * 1)$	$180 = (200 * 0.9)$	$120 = (150 * 0.8)$	$175 = (250 * 0.7)$

- 9 Enter a number in the Total field. The Viewer uses this value to allocate to the periods between the Start period and End period. If this field is blank, no allocation is performed.
- 10 Select or clear the Profile Recycle option. This option affects the behavior of the life cycle forecast when that are not enough data points in the profile series. For example, the data range is two years, but the profile only has one year's data for seasonal pattern.
 - If Profile Recycle is not selected, the Viewer reuses the value of the last period for the remaining periods.

- When Profile Recycle is selected, the Viewer uses the values starting from the beginning of the profile series (recycle).

11 Click Forecast to apply the life cycle forecast.

Using Disconnected Mode

This section describes how to use QAD Demand Management remotely.

<i>Introduction</i>	250
<i>Viewing Connection Status</i>	250
<i>Disconnecting from the Server</i>	250
<i>Database Table Modes</i>	251
<i>Committing Changes to the Server</i>	252
<i>Locking Scenarios in Disconnected Mode</i>	252
<i>Using the Connection</i>	252
<i>Limitations in Disconnected Mode</i>	253


Introduction

Disconnected mode lets you work remotely with QAD Demand Management by downloading part of the database. This removes the need to be connected constantly to the SQL server that holds the master database.

Viewing Connection Status

To view the connection status, open Viewer. The connection status displays at the right side of the status bar.

Fig. 19.1
Profile Forecasting



The value of the status is Connected or Disconnected:

Connected. Changes are saved to the server (master) database.

Disconnected. Changes are saved in the local client database.

Disconnecting from the Server

To disconnect from the server:

- 1 Open a scenario using the Viewer.
- 2 Check to make sure the mode is Connected.
Note If you have opened more than one scenario, make sure that the scenario you want to disconnect is the active scenario.
- 3 From the Viewer menu, select File|Disconnect/Reconnect. The Disconnect/Reconnect window displays, showing the server and local client databases.

4 Click the Disconnect button.

The scenario displayed in the active Viewer is downloaded from the server to your local client database, including other necessary information. The original data in the local client database is destroyed.


5 After the disconnection finishes successfully, the Viewer prompts you to log in again. You can only log in as the user who performed the Disconnect operation. After you have logged in, the connection status changes to Disconnected.

6 If you do not want to log in to disconnect mode right away, click Cancel.

Database Table Modes

The MODE in TBL_FX_JGD_FEATURE table in the server (master) database must be SERVER. The MODE in TBL_FX_JGD_FEATURE table in the client database must be CLIENT.

When you install an FX_JGD database, the default MODE is CLIENT. If the database you install is a server database, change the MODE to SERVER. This feature prevents the data in the server database from being destroyed by disconnecting by mistake. See Figure 19.2 for an example of these settings.



KY_FEATURE_ID	DESCRIPTION	STR_VALUE	DBL_VALUE
1	VERSION	4.5.2006.03.31	0
2	MODE	SERVER	0
3	Fix Sequence Issue		1
4	Batch Scheduler Ver 4.5.2006.03.31		0

Fig. 19.2
Viewing Database
Table Modes

Note The VERSION entry in TBL_FX_JGD_FEATURE table is maintained by the system. Do not change this value. When disconnecting or connecting, the VERSION of the server database and the local client database must be the same.

Committing Changes to the Server

To commit changes to the server:

- 1 Assume that you are working in Disconnected mode. Open the scenario you were working on by using the Viewer.
- 2 From the Viewer menu select File|Disconnect/Reconnect.
- 3 Click the Reconnect button. This commits the changes to the server.
- 4 Click the Cancel button if you do not want to reconnect at this time.
Only series that are marked as Forecast Complete are committed to the server. If no series is marked as Forecast Complete, you are prompted if you still want to reconnect.
 - Click No to cancel the connection, leaving you in disconnected mode.
 - Click Yes to reconnect. In this event, no series is committed back to the server.

After you are reconnected, the connection status displays Connected.

Locking Scenarios in Disconnected Mode

When you disconnect, the Engine locks the scenario to prevent other users from modifying it. This lock is set until you reconnect to the server. When you reconnect, the Engine checks the lock on the scenario. If the scenario is locked by another user, reconnection fails.

If you try to open a scenario that is disconnected, the Engine displays a message prompting you to unlock the scenario. If you respond Yes, change the scenario, and then close the Viewer releasing the scenario, the changes are overwritten when a disconnected user reconnects.

Using the Connection

Use Connection when you disconnect from the server and cannot reconnect. The inability to reconnect can be caused by the following:

- The local client database is no longer accessible and you cannot open the Viewer.

- The version of the server database is different from that of the local client database and you cannot reconnect to the server.

In these cases, use the Engine Connection tool to connect to the server without committing your changes.

- 1 Run the Connection application. The Engine Connection window is displayed.




Fig. 19.3
Engine Connection

- 2 Clear the Start in Disconnect Mode check box.
- 3 Clear the Change Connection for Disconnect check box.
- 4 Click the Accept button.
- 5 Start the Viewer.

You are now in Connected mode.

Important Do not change the Start in Disconnect Mode status in any other cases. If you do, you might overwrite the data in the server database with wrong data.

Limitations in Disconnected Mode


In Disconnected mode, you can create new scenarios and new series. However, when reconnecting to the server, these new scenarios or series are not uploaded to the server. Also, you can delete a series from the downloaded scenario or even the whole scenario, but when you reconnect to the server, the scenario and series are not deleted from the server. The scenario and series remain unchanged.

When you disconnect, the Engine does not download the stored procedures used for custom tasks. If some stored procedures do not exist in your local client database after disconnecting, you must create them manually.

In disconnected mode, the features listed in Table 19.1 are unavailable.

Table 19.1
Unavailable
Features in
Disconnected Mode

Menu	Unavailable or Limited Function
Edit	Undo
	Original
Forecast	Remove Hierarchy
	Bill of Material
	Life Cycle Forecast
Tools	Apply Calendar
	Apply Conversion Factor
Living Master	All functions
Hierarchy Manager	Add
	Apply
	Delete
	Edit (read-only mode)
Applications	Batch Scheduler
	Engine Manager
Base Unit	Frozen option is selected and cannot be unfrozen

The background of the page is a grayscale photograph of several interlocking mechanical gears. The gears are of different sizes and are positioned in a way that creates a sense of depth and motion. The lighting is soft, highlighting the metallic texture and the sharp edges of the gear teeth.

Chapter 20

Using the Batch Scheduler

This section describes the QAD Demand Management Batch Scheduler, which lets you schedule scenarios to run.

Introduction **256**


Managing Batch Scheduler **258**

Managing Massive Data Update (MDU) **263**

Introduction


The Batch Scheduler lets you run forecasts, load data, execute exceptions, and perform custom tasks. You can schedule these activities to take place on a client or server machine. Using the Batch Scheduler allows administrators and super users to configure and run repetitive tasks as part of a batch process. Figure 20.2 illustrates the activities that can be performed by the Batch Scheduler:

Fig. 20.1
Batch Scheduler



Using the Batch Scheduler you can update time series data points and hierarchy attributes. Figure 20.2 illustrates a key configuration for performing an incremental update.

Fig. 20.2
Batch Scheduling
Scenario



The Batch Scheduler moves data between scenarios and updates the attributes.

The Batch Scheduler can also perform incremental updates where new data is added to the series and the forecast is extended. In this role, the Batch Scheduler performs the function of the incremental update.

Finally, the system administrator can set up the Batch Scheduler to rerun forecasts and print out reports. For example, the Batch Scheduler could perform a forecast and automatically update the added periods with a system forecast after the incremental update is complete. The Batch Scheduler can automate the entire forecasting process, minimizing work for the forecasting team.

The Batch Scheduler lets you schedule scenarios to run automatically. One batch can contain one or more scenarios. Batches cannot be shared by users, as scenarios are shared.

Managing Batch Scheduler

Running the Batch Scheduler

To run the Batch Scheduler:

- 1 Double-click the Batch Scheduler icon on the Viewer panel. The Batch Scheduler Login dialog box is displayed.
- 2 Enter your User Name and Password, and then click OK. The Batch Scheduler opens.

Fig. 20.3
Batch Scheduler

DME Batch Scheduler (Edit Mode)
File
Batch
Options
Refresh
Help

Scheduled Batches

Batch Properties

Name:
Frequency:
Next Run:
Run Forecast:
Run Reports:
Reload Data:
Custom:
Run Alter:

Scenario Properties

Name:
Data Source Type:
Data Source Name:
Data Table Name:
Data Organization:
Forecasting Method:
Forecast Periods:
Seasonality:
Reports:
Excel:
Crystal:
Other:

History

Date
Time
Machine
Batch
Scenario


Add Batch
Edit Batch
Execute Batch
Delete Batch

Clear History
Refresh History
Pause
Exit

4:22 PM

- 3** Click **Add Batch**. The **Batch Add/Edit** screen appears.

Fig. 20.4
Batch Add/Edit



4 Enter the appropriate information.

Batch Name. Enter a name or description of the batch you are scheduling.

Available. The Available column lists the names of the scenarios available to include in the batch. Use the arrows to move the appropriate scenarios into Selected column.

Selected. The Selected column lists the scenarios selected to be run in the batch.

Show Scenario Properties. Indicate whether to display the scenario properties in the Scenario Properties box on the Batch Scheduler screen.

Run. Select what executes when the scheduled job runs.

Forecast and Report: Both the forecast and reports are generated when the job runs.

Forecast: Only the forecast is generated when the job runs.

Report: Only the report is generated when the job runs.

Mass Series Update: Update the series properties of a group of scenarios from a common source scenario. For details, see “Massive Series Update” on page 264.

Mass Observation Update: Update the values of option lines of a group of scenarios from a common source scenario, including the incremental update function. For details, see “Massive Observation Update” on page 269.

Exception Execution: Run the specified exception for each selected scenario in batch mode. For details, see “Managing Exceptions” on page 122.

Custom Task: Run the specified custom task for each selected scenario in batch mode. The custom task uses the saved parameters for each scenario. For details, see “Managing Custom Tasks” on page 107.

Batch Schedule. Select the frequency, time, and day at which you want the batch run.

Frequency: Select the frequency at which the batch runs. Your choices are once a day, once a week, once a month, once an hour, once every 15 minutes, or just once.

Time: Select the time of day you want the batch to run.

Starting day: When you click an arrow, a calendar of the current month pops up. On the calendar, select the day you want the batch schedule to begin.

Always Reload Series. Select this option to have the data reloaded from the database, before the forecast is run.

Use Last Loaded Series. Select this option to use the exact data values that were generated last time it was loaded.

Setup Task. For Mass Series Update and Mass Observation Update tasks, this button is available to set up other settings. See “Managing Massive Data Update (MDU)” on page 263 for details.

- 5 Click OK. The Batch Scheduler screen displays listing the information about the new scheduled batch.




Fig. 20.5
Batch Scheduler

The following information appears.

Scheduled Batches. Lists the existing scheduled batches with their selected scenarios underneath the batch name.

Batch Properties. Lists options for the scheduled batch.

Scenario Properties. Lists options for the selected scenarios within a batch.

History. Display the results of the jobs run. Notice that it is one item per scenario run, not per batch.

If you selected Forecast and Report from the Run drop-down list box, you also see one item for the forecast and another one for the report. If the Result status is Error, you can double-click it to view an error description.

- 6 Use the buttons on the bottom of the screen as needed. The options executed by these buttons are also available on the File menu bar:

Edit Batch. Displays the batch properties and lets you change them.

Delete Batch. Select a batch and click this button to delete it.

Clear History. Clear the history information in the History section.

Refresh History. Refresh the history information.

Pause. Pause or disable the Batch scheduler. No jobs launch when the Batch Scheduler is idle.

Exit. Exit the Batch Scheduler screen.

Important When you are finished scheduling the batch job to run, do not exit the Batch Scheduler or the scheduled jobs do not run.

You can close the window using the File menu or the Close window button. After closing the Batch Scheduler, it minimizes to an icon in the Windows task bar to run the tasks in the background. Click the icon to maximize the Batch Scheduler.

If a scheduled batch fails to finish running, an error message is displayed. The error dialog displays the scenario name, the task where the batch failed, and an error description.

More Options

The following features are available under Options on the File menu:

- **Show Startup Screen.** When this option is selected, the start up screen displays when Batch Scheduler is started.
- **Retry Failed Jobs.** Select this option to retry failed jobs.
- **Auto Login.** When logging in to the Batch Scheduler, a user name and password is required. The user can log in automatically the next time. So the next time the Batch Scheduler opens, it will use the current user name and password to do an automatic login. This feature is useful when running Batch Scheduler in background mode. When the operating system restarts the Batch Scheduler, it will run in background mode.
- **Set Report Directory.** Select this option to set the directory where generated reports are stored when a batch runs. The name of the file is the Scenario Name plus the report type.

Managing Massive Data Update (MDU)

This function includes the Mass Series Update and Mass Observation Update functions. It updates series properties or opinion lines in observations by connecting scenarios based on labels or series properties. It eliminates the need for manual editing and keeps other existing values. The following examples describe when the Massive Data Update function is used:

- To import new streams of opinion line such as financial forecast or custom orders to compare against the forecast values.
- To introduce new periods of actual values using an incremental update.
- To introduce new items to an existing scenario.
- To change the price and cost values as they are updated without affecting forecast numbers.
- To add or change attributes such as sale representative, client information as necessary, and to modify the series label correspondingly.

Massive Series Update (MSU) lets you update in batch mode the property values of the series for given target scenarios. When using MSU, you do not need to reload the scenario in order to set the most updated property values for each series.

Massive Observation Update lets you update in batch mode the opinion line values of the observations for given target scenarios. It lets you import new items to an existing scenario. It also supports incremental update feature. When using MOU, you do not need to reload the scenario in order to set the most updated opinion line values for each observation.

MDU copies values from a source scenario to a target scenario or group of target scenarios. The copy is required to specify how the series should be linked between the source and target scenarios and which fields are required for the update.

To understand how MDU works, you need to understand the following concepts:

Source Scenario. MDU uses the source scenario to perform the update. One and only one scenario can be the source scenario.

Target Scenarios. All the scenarios that are affected by the update. At least one target scenario must be specified.

Links. Represents the links at the series level between each target scenario with the source scenario. It can link by series label or series properties. It is essential that all links combined do not generate a many-to-one relationship between source and target scenario, which means one source series can be mapped to one or more target series, but not many sources series are mapped to one target series. Otherwise, the update result is unexpected.

Mappings. Mapping specifies which target field/opinion line is updated using a value of a specific source field/opinion line. These links are mapped series properties or observation opinion line.

Massive Series Update

Set Up Mass Series Update Task

If you create or edit a batch and select Mass Series Update, you should consider other settings.

Note Only one scenario can be selected as the source scenario.

The window in Figure 20.6 is displayed.




Fig. 20.6
Batch Add/Exit


Click the Setup task button. The Update Series Properties Settings dialog box is displayed.

The following sections describes the tab functions on the Update Series Properties Settings window .

Target Scenarios Tab

This tab contains the current selected scenarios defined as the target. Use the Change button to modify the list.

Fig. 20.7
Update Series
Properties Settings,
Target Scenarios
Tab



Labels Tab

Use this tab to re-create the labels. By default this option is not selected. If you want to re-create labels, specify at least one field for the label. To modify the list of fields for labels use the Change button. Note that the field order is important.





Fig. 20.8
Update Series
Properties Settings,
Labels

Links Tab

Use the Links tab to specify the links between the target and the source scenario. The combination of source fields represents a unique series in the source scenario. The combination of target fields represents a unique series in the target scenario

Fig. 20.9
Update Series
Properties Settings,
Links Tab



Mappings Tab

Use the Mappings tab to specify the fields that are affected by the update. In the UI each mapping is a couple of a target and source field. At least one mapping should be defined.




Fig. 20.10
Update Series
Properties Settings,
Mappings Tab

Massive Observation Update

Mass Observation Update is similar to Mass Series Update. MOU updates the opinion line values in observations, and the MSU updates the series level properties.

One significant feature supported by MOU is the incremental update, which lets users introduce new period data without reloading the scenario. As a result, all other existing forecasted values remain unchanged. MOU links observations in the matched series by the date information. The source and target scenario should be in the same time level. For example, they are all monthly data, are all weekly data.

Incremental Update


Incremental Update lets you apply the most recent historical data to existing, established scenarios. For each target scenario, the Batch Scheduler compares the dates in the scenario against the source scenario. If the dates are the same, the Batch Scheduler uses the data in the source scenario to update the historical data in each target scenario. The Batch Scheduler uses the series label or series attributes to determine which series to update in each scenario. The Batch Scheduler also adds some new periods at the end of the forecast for each line item in the target scenarios. The number of periods added is the same as the number of the forecasted periods updated by the new historical data.

Incremental Update Example

The incremental update process allows a system administrator to update the system data and extend the time fence based on the new data. In this way, the Engine can pull in new data and measure differences between the forecast and actual results.

When incremental update runs, the target scenario should line up by period with the source scenario. All periods in the target scenario that match periods in the source scenario should be updated with the value from the source scenario, and marked as historical periods in the target scenario. Figure 20.2 illustrates the effects of incremental update when it is run.

Fig. 20.11
Incremental
Example 1



The new data periods in the source scenario represent the current month's actual data. The purpose of the incremental update is to move the new data to the Engine and add any new items that have been added.

Figure 20.12 illustrates the effect of incremental update on the series.





Fig. 20.12
Incremental
Example 2

This example shows a source and target scenario, both with two series. When matching by label, the incremental update updates the matching periods in the target scenario with values from the source scenario, changing updated periods from forecasted periods to historical periods.

In the update process, observations are updated based on the dates entered. Using this approach it is possible to update and remap history as well as adding data to the future periods. The incremental update options are not enabled when the data is updated. However, as shown in this picture, the time fence does not advance and the forecast is not extended. This is useful for adding data to the engine that provides insight into current activity such as POS, customer orders.

If the incremental update option is turned on and the set to zero option is set, a different behavior occurs. This is shown in Figure 20.13.


Fig. 20.13
Incremental
Example 3



When the user uses the Set 0 If No New Data option, the program recognizes all series where no new data was provided. Then the program updates those series to include the same number of added historical periods, giving the added periods a value of 0. In this example, no new data was provided for Series B because no match was made with a series in the source scenario. The program recognizes that for all other series one historical period is added, so it defines period 2004-01 in Series B as historical and changes the value to 0.

If this option is set, any series that does not appear in the source scenario has the update period set to zero. After the initial update process is complete, the base data cleanup phase begins by—if the option is selected—removing the oldest data from each series. Figure 20.14 illustrates this process.


Fig. 20.14
Incremental
Example 4



When the user selects the Delete Oldest Data option, the program determines the number of added historical periods and deletes the same number of periods from the beginning of each series. In this example, there were two updated historical periods in Series A, but only one added historical period. Thus the incremental update removed only one period from the beginning of the series.

After the data is updated and the Engine determines the number of periods that were added, the Engine drops the same number of periods from each series. Once the deletion process is complete, the Engine extends the forecast. This process is illustrated in Figure 20.15.

Fig. 20.15
Incremental
Example 5



When the user uses the Add New Obs option, the program determines the number of added historical periods and adds the same number of periods to the end of each series as forecasted periods, which are equal to the settings of the prior last forecasted period.

In the example in Figure 20.15, there was one added historical period. Thus the incremental update adds one forecasted period to the end of each series, equal to the value in the previous last forecasted period.

After this process is complete, the Engine completes the process by adding the new items to the new item scenario that was selected when the incremental update was configured. In this way, it is possible to control and forecast new items separately from the established products.

Set Up the Massive Observation Update

- 1 When you create/edit a batch and you choose Mass Observation Update, you must set up other settings using the Setup task button. Only one scenario can be selected as the source scenario.





Fig. 20.16
Batch Add/Edit

- 2 Click Setup Task.


- 3 Select the appropriate scenarios in the Target Scenarios tab.

Fig. 20.17
Massive
Observations
Update, Target
Scenarios Tab



- a Click Change. The General Selection dialog box is displayed.

Fig. 20.18
General Selection



- b Select the target scenarios and click OK.

The Send new items to scenario check box lets you send new items from a source scenario to a specified target scenario.

- 4 Click the Links tab. There are two options.





Fig. 20.19
Massive
Observations
Update, Links Tab

Link by Label. Matches the series in the source scenario to the targeted scenarios by series label.

Link by selected fields. Matches the series in the source scenario to the targeted scenarios by selected series fields.

5 Click the Mappings tab.

Fig. 20.20
Massive
Observations
Update, Mappings
Tab



Specify the source opinion line and target opinion line. The number of fields in the source and target scenario should be the same.

6 Click the Incremental Update tab.




Fig. 20.21
Massive
Observations
Update,
Incremental Update
Tab

If you clear the Activate Incremental Update check box, all options in this dialog are unavailable. The program takes all the source opinion lines and updates the corresponding target opinion lines, both historical and forecast.

If you select the Activate Incremental Update check box, the following options are available.

Delete the Oldest Data. If you want to remove the oldest nth period of data after incremental update, select this option.

Add New Observations. Because the incremental update makes all the updated forecast period historical, you might want to add new periods so that you do not run out of forecast periods. Note that the number of periods being added equals the number of forecasted periods being updated. The values in the new observations are copied from the last forecasted observation.

Set Original to FValue. Sets the original opinion line equals to FValue opinion line.

Reset Measurement. Sets the measurement opinion line to zero.

Set 0 if no New Data. If you select this option, the observations for all series that do not exist in source scenario are set to zero.

Fields in New Observation Set By Last Observations. If the Add New Observations option is selected, you can set the values in which fields (opinion lines) are copied from the last observation to the newly added observations.

Limitations

Incremental update does not update a target scenario if the dates in the source and target scenarios do not overlap.

Example

Table 20.1
Target Scenario

2000-01	2000-02	2000-03	2000-04	2000-05	2000-06	2000-07	2000-08
10	11	12	13	14	15	16	15

Table 20.2
Source Scenario

2001-01	2001-02	2001-03	2001-04	2001-05	2001-06	2001-07	2001-08
1	1	1	1	1	22	23	24

In this case, the source scenario start date is 2001-01 but the target starts at 2000-01. There is no overlap and incremental update does not update the target scenario.

If the target scenario does not have any forecast period, incremental update does not delete the oldest observation and does not add any new observations. This is because there is no forecast period in the target to update.

Using Statistical Methods

This section describes the statistical methods available for producing forecasts in QAD Demand Management.

<i>Holt-Winters Method</i>	282
<i>Simple Exponential Smoothing</i>	285
<i>Adaptive Exponential Smoothing</i>	287
<i>Box-Jenkins ARIMA Modeling</i>	289
<i>Holt's Double Exponential Smoothing</i>	293
<i>Additional Smoothing Methods</i>	295
<i>Census X-11</i>	298
<i>Croston's Intermittent Model</i>	304
<i>Slow Method</i>	307
<i>Linear Regression</i>	313
<i>Stepwise Regression</i>	316
<i>Multiple Regression</i>	320
<i>Polynomial Regression</i>	323
<i>New Product Forecasting</i>	327

Using and Creating an Event Model 330

Specifying Statistical Options 335

Holt-Winters Method

Holt-Winters is an exponential smoothing technique that incorporates growth and seasonality into the forecast by producing seasonal lift factors for each seasonal period. QAD Demand Management displays the seasonal indexes in the Audit Trail report.

The Holt-Winters formula uses three parameters, or smoothing constants, to calculate a forecasted value:

- Alpha: determines the level of smoothing.
- Gamma: determines growth trend.
- Beta: determines seasonality.

The value of each parameter is always between zero and one.

If the historical data changes rapidly, large smoothing constants should be used. For stable, naturally consistent data, the smoothing constants should be closer to zero. In either case, QAD Demand Management automatically optimizes the parameters when they are not already specified.

Example of Holt-Winters

This example is based on the volume of ticket sales for bus rides: The Go Anywhere bus line is trying to determine when the company should deploy the largest number of buses to meet its rider demand. Using the Holt-Winters method, Go Anywhere forecasts its ticket volume sales.


To forecast using Holt-Winters:

- 1 Create a new scenario.
- 2 Click the Data Capture tab and click the button next to Data Source Name. In the Available Data Sources, select Excel Spreadsheet.
- 3 Browse to the `ExcelExampleFXE.xls` file. Click Done.

- 4 Choose Holt-Winters from the Table/Query drop-down list. This example has 28 months of historical ticket sales on which to base the forecast.
- 5 For Data Organization, select Columns.
- 6 Click the Forecast Method tab. From the Forecasting Technique drop-down list, choose Seasonal Model.
- 7 From the Primary Method drop-down list, choose Holt-Winters. Leave the smoothing constants blank to allow QAD Demand Management to automatically choose them.
- 8 Click Procast.
- 9 When the forecast is finished generating, click the Statistics tab.
- 10 Select the Write Multiple Statistics check box. This shows the seasonal indexes in the report.
- 11 Click the Reports tab to choose how to display the results. Select the Audit Trail report. The Audit Trail provides an analysis of your forecast.
- 12 On the Executive tab, select Show Charts to display a graphical representation of your data, and the Upper and Lower limits check boxes to display the confidence limits of your forecast.
- 13 Click Run Report.

After running the forecast, QAD Demand Management produces a graph and table with the seasonal indexes and smoothing constants (see Figure 21.1). The graph lets you quickly assess the trend and seasonality of data by visualizing the peaks and valleys of data.

Fig. 21.1
Holt Winters Graph



Graph Results

The statistical tables let you determine the accuracy of a forecast. The lines on the graph represent the following:

- A brown line represents the fitted values of the forecast and illustrates how closely the forecasted and actual values match.
- A dark blue line represents the actual values.
- A pink line represents the forecast values.
- Light blue and purple lines above and below the forecasted values represent the confidence limits. These values represent the likelihood, within five percent, of the forecast to fall within that range.

Overall, the graph shows an upward trend in the sale of bus tickets, with this seasonal trend continuing into the future.

Audit Trail Results

Audit trail statistics reveal that QAD Demand Management produced an accurate and trustworthy forecast, represented by the very low MAPE of 12.54%. In addition, the Audit Trail report shows you the smoothing constants, alpha, beta, and gamma, as well as the seasonal indexes. You can use this detail to learn about the trends and seasonality of your data.

Audit Trail - Statistics			
Accuracy Measures		Value	
Mean Absolute Percentage Error (MAPE)		12.54%	
R-Square		90.26%	
Forecast Measures		Value	
Mean		539.36	
Standard Deviation		249.18	
Method Statistics		Value	
Primary Method		Holt Winters	
Alpha		0.78	
Beta		1.00	
Gamma		0.00	
Seasonal Length		4.00	
Decomposition Type		Additive	

Fig. 21.2
Audit
Trail-Statistics

Using the forecast generated by QAD Demand Management, the Go Anywhere bus line can operate more efficiently and cost effectively. Bus drivers are staffed as needed, the correct number of buses are deployed at the appropriate times, and the company can service the buses during down times. As a result, no riders are lost due to lack of buses. With QAD Demand Management, Go Anywhere adds to its bottom line by providing the highest level of customer service.

Simple Exponential Smoothing

Simple exponential smoothing is a method of forecasting using both historical data and fitted data to generate the new forecast. The underlying concept of simple exponential smoothing is that recent values are given relatively more weight than older observations.

For example, in a three-period moving average forecast, the February, March, and April actual values evenly determine May's forecast.

Depending on the level (alpha value) given to the data series, a simple exponential smoothing method creates a forecast for May that more heavily relies on April's observations. Simple exponential smoothing is most effective as a forecasting method when irregular influences comprise the main effects on the historical time series values.

The alpha value, or the level of smoothing, always is a value between zero and one. If the alpha value is closer to one, the level of smoothing more closely reflects the actual data. If the value is closer to zero, the level of smoothing more closely reflects the fitted line, or forecasted values.

As with all smoothing techniques within QAD Demand Management, you can set the smoothing constants yourself or allow QAD Demand Management to optimize the values for you.

Example of Simple Exponential Smoothing

This example uses data that does not have recognizable patterns such as trend or seasonality. This makes simple exponential smoothing an appropriate method to employ.

To forecast using simple exponential smoothing:

- 1 Create a new scenario.
- 2 Click the Data Capture tab and click the button next to Data Source Name. In the Available Data Sources, select Excel Spreadsheet.
- 3 Browse to the `ExcelExampleFXE.xls` file. Click Done.
- 4 Choose the Simple Exp. Smoothing data source from the Table/Query drop-down list.
- 5 On the Forecast Method tab, choose Exponential Smoothing from the Forecasting Technique drop-down list.
- 6 Choose Simple ES from the Primary Method drop-down list.
- 7 Click Procast.
- 8 Click the Reports tab.
- 9 On the Reports tab, select the Audit report option.
- 10 On the Executive tab, select the Show Charts option. This lets you see a graph of your forecast and analyze the statistics used to measure the Simple Exponential Smoothing forecast.
- 11 Click Run Report and wait for the graph to appear.

Summarizing the Results

QAD Demand Management provides two tools to judge the accuracy of your forecast and determine the direction of your business. The first thing you see is the Executive Report chart. One line shows the actual values, and the other line represents the fitted values and the forecasted values.

QAD Demand Management also produces a detailed table (see Figure 21.3) that shows the optimized parameters for your Exponential Smoothing forecast. The table indicates that the forecast has an average

error of 6.53%. This means that QAD Demand Management smoothed your data to a considerable degree. A low smoothing value indicates a high level of smoothing.




Fig. 21.3
Simple Exponential
Smoothing Graph

Adaptive Exponential Smoothing

The adaptive exponential smoothing method is a derivative of the simple exponential smoothing method. The difference with adaptive exponential smoothing is that the alpha value systematically changes from period to period to allow for pattern changes in the historical data.

The adaptive exponential smoothing method is completely automated and uses the advantages of simple exponential smoothing. This method is most useful when:

- The data has many items.
- The data is non-seasonal.
- The data shows no trend.

An example of a good time to use adaptive exponential smoothing is when a product has been sold for a long period of time, only to encounter a serious market change in recent periods. While forecasting, you want to consider the most recent sales data more than the older sales data.

Example of Adaptive Exponential Smoothing

In the dataset file that was installed with QAD Demand Management, you can use a file to generate a sample forecast using the adaptive exponential smoothing method.

To forecast using adaptive exponential smoothing:

- 1 Create a new scenario.
- 2 Click the Data Capture tab and click the button next to Data Source Name. In the Available Data Sources, select Excel Spreadsheet.
- 3 Browse to the `ExcelExampleFXE.xls` file. Click Done.
- 4 Choose the Adaptive Exp. Smoothing from the Table/Query drop-down list.
- 5 For Data Organization, select Columns.
- 6 From the Forecasting Technique drop-down list box, choose Exponential Smoothing. Leave the parameters blank so QAD Demand Management can optimize the values for you.
- 7 Click Procast to run the forecast.
- 8 Click the Reports tab.
- 9 From the five reporting options available, choose Standard. Your output appears in an Excel Spreadsheet.
- 10 Select the Show Charts option from the Standard tab.
- 11 Click Run Report.

Summarizing the Results

Two areas in the Standard report display the results of your forecast.

First, the Excel spreadsheet, which is called Chart1, produced by QAD Demand Management shows:

- The actual values of your original historical data
- The forecast method chosen
- The forecasted values for both the original historical data and the future values

Second, near the bottom of the spreadsheet, is a small table that contains basic error measurements: MAPE, R-Square Error, Mean, and Standard Deviation.

This table shows that your forecast accuracy is good, and the fit of the forecast, represented by the R-Square value, is very good.

Click the tab in the lower-left corner of the spreadsheet to view a graph. The graph lets you view the forecast's future and the behavior of the data's history.




Fig. 21.4
Adaptive
Exponential
Smoothing Graph

As the graph indicates, the market changes from November 2004 to May 2006 affected the sales of your product. Because you chose the adaptive exponential smoothing method, QAD Demand Management weighed the most recent values more heavily than older values. This lets you consider how the market is currently behaving, and how that behavior affects your total sales.

Box-Jenkins ARIMA Modeling

George Box and Gwilym Jenkins popularized the Auto-Regressive Integrated Moving Average (ARIMA) models in the early 1970s. The basis for the Box-Jenkins methodology consists of three phases: identification, estimation, and the testing and application of the ARIMA model.

This methodology is a multi-step model building strategy aimed at optimizing the ARIMA process. You do not have to be an expert in this methodology to forecast using Box-Jenkins. QAD Demand Management automatically optimizes the best ARIMA using the Box-Jenkins methodology. For the more advanced statisticians, QAD Demand Management lets you perform data transformation and analyze the ACF and PACF charts for model selection. Box-Jenkins is best used on extensively long historical data sets with lower volatility.

Box-Jenkins Methodology

The Box-Jenkins method has four phases:

- 1 Model Identification.** In this phase, data is plotted (if needed) for performing transformation and differencing analysis. A data transformation stabilizes the variance within a set of data points. Differencing the data makes it stationary. Also in this phase, a first guess at the ARIMA model is analyzed.
- 2 Estimation and Forecasting.** In this phase, the model parameters are chosen based on the model that provides the lowest error term. The data is then forecasted.
- 3 Diagnostic Checking.** The decision is made based on an analysis of the statistical accuracy of the model whether to re-run using a different model or continue using the ARIMA model diagnosed.
- 4 Forecasting.** A forecast of the time series is generated using the ARIMA model.

The two basic classes of components in Box-Jenkins are trend and seasonality. The airline example is an excellent way to learn trend and seasonality within Box-Jenkins as well as demonstrate how Box-Jenkins projects future travel patterns and, in turn, gross sales.

Seasonality differs from trend in that sporadic changes in the number of passengers traveling happen repeatedly throughout the year and various seasons. The fact that these elements come into play at a specific time of the year demonstrates seasonality.

When looking for trends in data, look for gradual increases in sales over time that does not repeat within the time range being forecasted. For instance, the number of people who fly in general has increased over the last ten years resulting in a consistent upward slope in the data. This, in combination with higher sales during certain times of the year (seasonality), signifies an increase in sales from year to year that follows a seasonal pattern, which is why you chose the Box-Jenkins/ARIMA model to handle this data.

Example of Box-Jenkins Methodology

During holiday seasons, airline flights are booked and people are trying to get out of town at the last minute. Ultimately, more people are trying to fly than the airlines can accommodate. This forecasting problem has become a classic example for time series forecasting. Most textbooks and college courses use airline data to walk through Box-Jenkins ARIMA methodology.

In the dataset file that was installed with QAD Demand Management, you can use a file to generate a sample forecast using the Box-Jenkins method.


To forecast using Box-Jenkins:

- 1 Create a new scenario.
- 2 Click the Data Capture tab and click the button next to Data Source Name. In the Available Data Sources, select Excel Spreadsheet.
- 3 Browse to the `Tutorial_AirlineData.xls` file. Click Done.
- 4 Choose Data Sheet1 from the Table/Query drop-down list.
- 5 For Data Organization, select Columns.
- 6 Set the Forecast to 36 to forecast three years into the future, and set the Seasonality to 12.
- 7 Click the Forecast Method tab.
- 8 From the Forecasting Techniques drop-down list, choose Seasonal Model. Select Box-Jenkins as the Primary Method.

Note You do not set seasonal parameters in this example. The Expert selection of QAD Demand Management chooses parameters.

- 9 Click Procast.
- 10 Click the Reports tab and select the Audit report option. The Audit Trail report includes accuracy and descriptive statistics for the model.
- 11 Click Run Report. The graph appears (see Figure 21.5).

Fig. 21.5
Box-Jenkins
ARIMA Model



Performing ARIMA Analysis

QAD Demand Management includes several tools to analyze your data and assist on ARIMA selection. These tools include:

- **Customize ARIMA Model Input:** The ability to input a non-seasonal and seasonal ARIMA model is available in QAD Demand Management.

On the Forecasting Method screen, you can select Box-Jenkins and enter your own ARIMA model.

- **Data Transformation** charts the variance and provides the functionality to transform your data. By examining the variance chart, the sophisticated forecaster can perform a log, square, and square root transformation. Click Series Preview and Transformations at the bottom of the QAD Demand Management window, switch to the Transformation screen. Take a look at your variance chart. The goal is to have your data distributed evenly around the mean.

- The Data Analysis chart includes the ACF and PACF chart and provides functionality to difference your data. By examining the ACF chart, the sophisticated forecaster can determine if differencing is required and also what the order of the ARIMA should be.
- From the Series Preview and Transformation screen, click the Analyze tab, which lets you set the differencing parameters. Differencing lets you remove the trend by subtracting adjacent data points. The correlogram displays the ACF and PACF, which show the auto-correlations and relationships that each variable has with the past. If correlation is high, it indicates that the variable is predictable.
- The ACF and PACF are used in determining whether your data is stationary. It also shows seasonality and trend. As discussed previously, your data might contain both factors and therefore must be differenced before you can choose the ARIMA model. QAD Demand Management lets you enter differencing parameters and view the changes in the correlogram as you make them. Also, looking at the pattern of the data in both the ACF and the PACF helps you determine the best model.

Holt's Double Exponential Smoothing

Holt's linear exponential smoothing is similar in principle to simple exponential smoothing. It calculates alpha, the level component, to measure the trend in the forecast. It also adds the parameter gamma, the trend component, to create a linear trend in the forecast.

This equation is similar to a linear regression line and useful when a product is experiencing an exponential growth or decline, while not exhibiting recognizable seasonality. That is, if your data is dynamic and does not change due of seasonal factors, Double Holt is an appropriate method to use.

Example of Holt's Double Exponential Smoothing

This example uses a short data set that is intermittent. As sales manager of this new product, you need to determine what the trend is likely to be, despite given such sporadic data.

To forecast using Double Holt:

- 1 Create a new scenario.
- 2 Click the Data Capture tab and click the button next to Data Source Name. In the Available Data Sources, select Excel Spreadsheet.
- 3 Browse to the `ExcelExampleFXE.xls` file. Click Done.
- 4 Choose Double Holt from the Table/Query drop-down list.
- 5 For Data Organization, select Columns.
- 6 Click the Forecast Method tab. From the Forecasting Technique drop-down list, choose Exponential Smoothing.
- 7 From the Primary Method drop-down list, choose Double Holt.
Leave the smoothing constants blank to allow QAD Demand Management to automatically choose them.
- 8 Click Advanced.
- 9 When the Data Capture Advanced Options screen appears, click the Cleansing tab and select Replace Outliers. QAD Demand Management automatically removes your outliers and produces a more accurate forecast as a result.
- 10 Click Procast.
- 11 Click the Reports tab to choose the Audit Trail report. Click Run Report.

After the report runs, it appears.

Summarizing the Results

After QAD Demand Management produces your results in an Audit Trail report, you can see all the statistics you need to measure the values of your Double Holt forecast.

The graph that is made of the actual, fitted, and forecast values lets you see the trend in your new product (see Figure 21.6). Even though you do not have a lot of data, QAD Demand Management detected the likely trend.




Fig. 21.6
Holt's Double
Exponential
Smoothing Graph

One point of interest on the graph is the area of January 2006. QAD Demand Management removed a large data point, the outlier. The only thing you needed to do was click Replace Outliers on the Advance Option screen.

The Audit Trail report for the Double Holt forecast has two important statistics: the constant smoothing values or the Alpha and Gamma measurements. These values tell you how significantly QAD Demand Management smoothed your data. The Gamma value is high (near 1.00), so the smoothing was not significant. However, the Alpha value is low and therefore a significant smoothing constant.

This means that your data has a significant trend. Again, you can verify this by simply glancing at the graph. QAD Demand Management makes it simple to view and review the results of your forecast.

Additional Smoothing Methods

This section describes additional smoothing methods available in QAD Demand Management. The example reports in this section used the Tutorial_AirlineData.xls data file with the respective method.

Moving Average

A moving average forecast is the average of a time series over a specific number of preceding periods. When a new value is added, the last is dropped from the calculation, so that the specific number of preceding periods remains a constant, hence the name moving average.


However, two drawbacks are that moving averages do not effectively handle significant trends in data and all historical data must be stored to create the moving average. In the case of time series data, moving average forecasts are often used to eliminate unwanted fluctuations, thereby smoothing the time series. The appropriate number of preceding periods should be determined by selecting the amount of periods that yields the least amount of error.

The start value of a moving average is where the calculation begins. The start value uses zero as its first value, so if you would want to start at the fourth observation, the start value would be three. The number of observations to include in the average calculation is called the history.

For example, if you want to start at the fifth observation and perform a moving average with three historical observations, the start value would be four because the count starts at zero.

The moving average calculation would include the third, fourth, and fifth observations.

Fig. 21.7
Sample Moving
Average



Weighted Moving Averages

Weighted moving averages are moving averages of moving averages. In other words, rather than replace the oldest observations within the data, the Weighted Moving Average replaces the oldest moving average with the most recent moving average. In this way, Weighted Moving Averages handle significant data trends better than a plain moving average.

Use Weighted Moving Average forecasts to overcome the strong affects of extreme values within a time series by assigning current data more weight than older data. The start and history parameters are the same as those in moving averages.

Brown's Double Exponential Smoothing

Brown's Double Exponential Smoothing method attempts to create a linear equation. It performs two simple exponential smoothing forecasts and then adjusts for the linear trend in the data.

Brown's Double Exponential Smoothing method is similar to Holt's method in the fact that the goal is to create a linear trend, but Brown's does so without adding additional parameters to the equation.


Since forecasts can be expressed as a function of the single and double smoothed constants, the procedure is known as *double exponential smoothing*. Brown's double exponential smoothing is most appropriate for data that shows a linear trend over time. As defined previously, the smoothing component is alpha.

Brown's Triple Exponential Smoothing

Brown's triple exponential smoothing is implemented on data that shows a quadratic trend over time. It works well for data that has a steep growth or decline.

The Brown's triple exponential smoothing attempts to create a linear equation by performing three simple exponential smoothing forecasts and then adjusting for the linear trend in the data.

Fig. 21.8
Brown's Triple
Exponential
Smoothing



Census X-11

Census X-11 is a refinement of the seasonal-decomposition method. It seasonally adjusts and decomposes your data through a series of predefined steps: seasonal, trend, cycle, and random or irregular.

A seasonal component of a time series occurs regularly—for example, Christmas—while a cycle has a duration that varies from cycle to cycle. Economic variables are cyclical. A depression occurs at irregular intervals.

During forecasting, Census X-11 factors in trading days to refine the forecast and reduces the irregular component of the forecast by weighting each period by the number of trading days. QAD Demand Management lets you account for trading days with one easy step or forecast without adjusting for trading days. The former produces more accurate forecasts when trading days are important.

This is especially useful if the number of trading days dramatically affects the amount of business that is done in a time period. Say you have a value of 100 in two consecutive months. On first glance, it appears that the two months are equal. However, assume that the first month has 25 trading days while the second has 15. In this case, Census X-11 takes the value 100 and spreads it over 25 ($100/25$) for the first month and 15 ($100/15$) for the second. You should use Census X-11 here because it is likely that any extreme observations between the months can be explained by the number of trading days.

Trading days also can be specifically known for each time period. For example, you can determine how many trading days each of the last 12 months contained by reviewing how many official working days took place. You can easily determine how many trading days are going to occur over the next 36 months, or whatever your forecast periods, by taking into account days off, holidays, or weekends. Census X-11 offers another way of determining factors affecting your forecasts.

Census X-11 also identifies two types of seasonality:

- Additive decomposition assumes the data is the sum of the time-series components and is independent of the trend.
- Multiplicative decomposition assumes the data is the product of the time-series components and is a ratio of the trend.

Example You look at an annual sales report for product X with a peak that repeats every year in the months of November and December. Every year during December, the sales of X increase by \$200,000. Each year, then, you can add \$200,000 (over the annual average) to the sales of X in December. The seasonality is additive. In contrast, suppose that sales of X increase by 20% during December, or some other factor. Then the seasonality is multiplicative.

Census X-11 is best suited for a series where the number of days in the time period directly affects the values in those time periods.

Example of Census X-11

You are the owner of the Club Sport Gear import-export clothes store. Through the global nature of the Internet, you have realized that your business is highly seasonal and influenced by holidays. This has become more evident since you have been doing booming business with a Chinese distributor of your clothing line.

As this business grows, your forecast accuracy decreases because the overall affect of the Chinese New Year is not factored into forecasts. During the Chinese New Year, which never falls on the same day in consecutive years, your forecasts have not allowed you to make confident decisions. That's because the Chinese distribution company you depend on lets employees take a traditional five days off work during their

holiday, and your sales in China decrease on their holiday as well as the days surrounding it. The challenge you face is determining the extent to which the Chinese New Year affects your sales.

The Chinese New Year might fall in mid-January one year and late February the next. You have been able to forecast using the amount of money you make per month in exports, but you know the forecast would be more accurate if you could factor in the number of actual business (or trading) days per month. If all factors are not taken into account, the forecast output is not as accurate as it could be.

Since the number of days you do business with the Chinese distribution company varies, this year you select the forecasting technique Census X-11 so that you can factor in the irregular work schedule. This added factor to your forecast improves the overall accuracy of the forecast.

According to your data, you have had increased sales on exports each of the past eleven years. The Chinese New Year dramatically increases your export sales in November and December. Whereas the months of the Chinese New Year always decrease your export sales in either January or February depending on when the holiday occurs. You know when the Chinese New Year occurs over the next several years.

Census X-11 lets you account for the factor that is dramatically affecting your forecasts. QAD Demand Management accomplishes this by dividing your original data by the number of trading days that you specify. This weighs the monthly observations in accordance to their respective trading-days for each month. Now your forecast is based on more accurate data. To illustrate the power of CensusX11, perform the first forecast without accounting for trading days.

To forecast using Census X-11:


- 1 Create a new scenario.
- 2 Click the Data Capture tab and click the button next to Data Source Name. In the Available Data Sources, select Microsoft Access Database.
- 3 Browse to and select the `CensusX-11.mdb` file. Click Done.
- 4 Choose the Census X-11Table from the Table/Query drop-down list.
- 5 For Data Organization, select Columns.

- 6 In the Seasonality box, enter 12. To produce a forecast for 3 years, enter 36 in the Forecast box.
- 7 Click Advanced to specify your data as monthly.
- 8 Click the Forecast Method tab and choose Seasonal Model for the forecasting technique and Census X-11 from the Primary Method drop-down list. Leave Trading Day Adjust empty for this forecast.
- 9 Click Procast.
- 10 Click the Report tab. Select the Executive Report option and clear any other available options.
- 11 Click Run Report, and wait for QAD Demand Management to create the report.

Summarizing the Results

Notice QAD Demand Management has confirmed the Chinese New Year does affect your sales in the months preceding and following the holiday. Watch for the following information. According to the forecast, sales through your distribution outlet continue to increase over the next three years. Although, since you did not take the actual number of business or trading days into account, the average error of your forecast is 3.81%. Now you decide to run a forecast that does account for the actual trading days in each month. The goal here is to produce an even more accurate forecast.

Fig. 21.9
Census X-11 Graph



Forecasting with Advanced Census X-11

To forecast using advanced CensusX-11:

- 1** Select the `CensusX-11.mdb` sample spreadsheet as the data source.
- 2** Choose Census X-11 Table from the Table/Queries drop-down list.
- 3** For Data Organization, select Columns.
- 4** The data is monthly, so in the Seasonality box, enter 12.
- 5** To make a direct comparison to your previous forecast of 3 years, enter 36 into the Forecast box.
- 6** Click the Forecast Method tab, and from the Forecasting Technique drop-down list, choose Seasonal Model.
- 7** Choose Census X-11 from the Primary Method drop-down list.
- 8** In the Seasonal Type box, choose the Multiplicative method and in the Add Series for Census X-11 field, select the Trading Day Adjustment option.

Multiplicative decomposition spreads the total values from each month over the actual number of trading days of each month.

From the Additional Table drop-down list, select the table where the trading days are located. Then click Setup Trading Days.

- 9 Choose Procast from the Secondary Method drop-down list.
- 10 Click Procast.
- 11 When the Procast is done generating, click the Reports tab.
- 12 On the Reports tab, select the Audit Trail option and clear all other report types options. This way, only the Audit Trail Report is generated.
- 13 Click Run Report.

Summarizing the Results

The Audit Trail report is more detailed than the Executive report. For example, it shows detailed information about the method you selected, Census X-11, along with the decomposition type you specified, multiplicative. Choose multiplicative when you want to spread the values from each month over the actual number of trading days for each month.

In this example, QAD Demand Management enabled you to divide the value from January by the actual number of trading days in January.

The most important statistic is the Mean Absolute Percentage Error (MAPE) of 1.23%. In the previous forecast, you did not specify the number of trading days and the forecast returned a MAPE of 3.81%.




Fig. 21.10
Advanced Census
X-11 Graph

Using QAD Demand Management, you could specify the exact number of trading days, thereby receiving a forecast with an increased accuracy of 2.58%. The advanced functionality of QAD Demand Management enabled you to achieve your goal of an accurate forecast, which enhances your decision-making abilities about future sales to the Chinese Distribution Company during important holidays.

Croston's Intermittent Model

Croston's Intermittent model handles sporadic data. For example, consider a manufacturer whose inventory is constantly in sporadic demand; there are no precursors to when orders are placed, or the demand is often zero even though the average demand is for several units. In addition, when an order is placed, it is usually for more than one particular item.

This situation causes an overcalculation of the average or mean demand. At times the demand is zero, and other times the demand is for multiple orders, which limits how accurately the manufacturer can predict stock needed. Orders do not arrive at predictable times, so sometimes the manufacturer is over- or understocked.

Croston's Intermittent model is specifically designed to deal with such a situation with a two-step process. The main goal is to provide a good safety stock, resulting in a situation where the manufacturer is neither over nor under stocked. This is accomplished by considering two aspects of the data: the demand size and the demand occurrence.

Croston's Intermittent model is better than exponential smoothing here because exponential smoothing ignores the sporadic nature of the data. For example, suppose that the average order for a manufacturer is 10 units, but they occurred only 10 times a year in the past, in irregular patterns, and the demand was often zero. In this case, Croston's Intermittent model recognizes both the sporadic timing of the orders and that the average order (10) is actually a misrepresentation if it is spread evenly over the number of days.

Example of Croston's Intermittent Model

This example shows how Croston's Intermittent model can accurately forecast intermittent demand to help you control stock. This forecast tracks the compact disc sales from a new group who has just released their first album.

As the manager of a music store, you started forecasting demand for the new CD. Although sales for this group has skyrocketed, their daily sales began to take a more irregular pattern of sales. You decide to forecast the demand in order to save money and space for other, almost-as-popular CDs.


To create a forecast using Croston's Intermittent model:

- 1 Create a new scenario.
- 2 Click the Data Capture tab and click the button next to Data Source Name. In the Available Data Sources, select Excel Spreadsheet.
- 3 Browse to and select the `ExcelExamplesFXE.xls` file. Click Done.
This spreadsheet contains daily records of the CDs sold. The demand is unpredictable, both in when the demand occurs and the quantity it occurs in.
- 4 Choose Crostons from the Table/Query drop-down list.
- 5 For Data Organization, select Columns.
- 6 Click the Forecast Method tab. Choose Slow Moving Item from the Forecasting Technique drop-down list and Croston Intermittent from the Primary Method drop-down list.
The Alpha or smoothing constant box appears. Leave it blank. QAD Demand Management optimizes the value for you.
- 7 Click Procast to run the Forecast.
- 8 After the forecast is run, click the Reports tab to choose the reports to print.
- 9 On the Reports tab, select the Standard Report and the Show Charts box.
- 10 Click Run Report to view the Forecast.

Summarizing the Results

Following is Chart1, produced by the Standard report. The data is sporadic.

Fig. 21.11
Croston's
Intermittent Model
Graph



Croston's intermittent forecasting method lets you control and smooth sporadic data. This method provides excellent results for product control and forecasted demand. In the standard report, you can see the MAPE is 75.05% and the R-Square is 0.

At first glance, the forecast result is poor; however, because the data set has many zeros, these two statistics are not useful in judging the accuracy of the Croston's intermittent forecast. The most important measure are the forecast observations. In this example, 0.73 is the value of each and every forecast observation.

With the forecast values, the mean, and the standard deviation of the data set, the safety stock can be calculated accurately. You can retrieve the safety stock from the Audit Trail report. It has been theoretically proven that for intermittent data sets, Croston's intermittent forecasting method can produce more accurate values for safety stock than can other exponential smoothing methods.

In conclusion, Croston’s intermittent is an important method to use for determining inventory control from intermittent data.

Slow Method

For slow moving demand, most forecasting packages use the Croston method. The issue with Croston is that it assumes that the in-coming demand is a normal distribution and does not consider seasonal pattern. In reality, a lot of demands are not normal distribution and has some seasonal pattern.

The slow moving method is a special forecasting technique to overcome the shortcoming in the Croston method. The technique is designed to calculate the optimal stocking level for sporadic data. This technique works for items that have large numbers of zeros and items that have highly variable non-seasonal demand.

The chart in Figure 21.12 shows the types of data that are suitable for the slow moving stocking level technique. In each case, note that neither of the data series contains a pattern, nor are the demands normally distributed around a mean, which would lend the data to being forecast by a more traditional forecasting technique.





Fig. 21.12
Slow Moving
Method

The purpose of the slow moving technique is to set the appropriate stocking level to ensure that the proper amount of inventory is in place to meet demand taking into account service level and lead time.

When the slow method is run, the result is the total stocking level required to support demand and an average sales rate. This allows demand to be decomposed into its two components: an average sales rate, and safety stock.

Fig. 21.13
Slow Moving
Method




In this case, the slow method was configured to support a 90% service level with a two-month lead time. The slow method then developed an event model to measure the demand impulses and compute a stocking level demand to meet the expected demand in the lead-time period.

What this means is that in each case the slow method was able to recommend a stocking level to achieve the service level you want for the sporadic item.

When this is contrasted to the Croston technique, several major differences emerge because Croston assumes that the data being forecasted is normally distributed. When this is the case, as often occurs with sporadic data, the results are dramatic.

Fig. 21.14
Slow Moving
Method



Because the data is not normally distributed, Croston yields results that might not guarantee that the service level you want is achieved. In the case of Item A, the safety stock level is significantly higher—1380 versus 1121—than in the slow method, resulting in a potential over investment in inventory for this item. In the case of Item B, we see that the situation is reversed—7168 versus 9729—resulting in under investment and the potential loss of sales. Therefore, as a general rule, use the slow moving method for forecasting rather than Croston and the calculated safety stock as the stocking levels.

To use the slow moving method:

- 1 Create a new scenario.
- 2 Click the Data Capture tab and click the button next to Data Source Name. In the Available Data Sources, select Excel Spreadsheet.
- 3 Browse to and select the `ExcelExamplesFXE.xls` file and click Done.
This spreadsheet contains daily records of the CDs you have sold. The demand is unpredictable; both when the demand occurs and in what quantity it occurs.
- 4 Select Crostons from the Table/Query drop-down list.

- 5 For Data Organization, select Columns.
- 6 Click the Forecast Method tab. Select Slow Moving Item from the Forecasting Technique drop-down list and Slow Method from the Primary Method drop-down list.

- 7 Select the Event Index Creation. There are two options

Seasonal. This method is used to create the event indexes in the historical period automatically.

The number cannot be larger than the number of observations. If it is empty, the cycle is the same as the number of observations. In this mode, the index starts from 1 until the cycle number, and then repeat.

Std Bucket. Use the mean and standard deviation to determine the event index. The options include:

- **Exclude 0** - Use Exclude 0 when using the mean and standard deviation. In this option, all 0 are assigned a same event index 0.
- **Start** - For values in the range $[\text{mean}, \text{mean} + \text{Start} * \text{std}]$, assign index 1 to Start. For values in the range $[\text{mean} - \text{Start} * \text{std}, \text{mean}]$, assign index 2 to Start.
- **Incr** - For values in the range around the mean by $(\text{Start} + x) * \text{std}$, assign $x*2$ and $x*2 + 1$ index to Incr (depend on the direction).
- **End** - For values in the range around the mean by $\text{End} * \text{std}$, assign the last two index to End.

- 8 Enter the Base Cal Method. This method is used to calculate the base line of the event model. Select from the following options.

Mean. Use the mean as the base line. Select Exclude 0 when calculating the mean.

Croston. Use the Croston method to calculate the base line.

- 9 Enter the Base Fcst. Method. This is the method used to calculate the forecasted value of the based line of the event model. Select from the following options:

Same as Base Calculation. Use the same method as the one used in the base line calculation.

Select a specific method. It can be one of: Simple ES, Croston, Adaptive ES, Linear Regression, Double-Holt, Double-Brown, New Product Curve or Procast.

- 10** Select Event Replacable. This option simulates the events in the future periods. Select from the following options:

Yes: Randomly selects an event from minimum to maximum. When all zeros are assigned the same event (0), this special event needs more possibility to be selected than other events.

No: Each event and the number of its appearance are in a pool. An event is randomly selected and the number of the corresponding appearance is reduced by 1. If the lead time is longer than the historical periods, this option is unavailable. Note that the number of the events in future periods is the ceiling value of the safety stock lead time because the lead time can be a fractional number. As an example, if the lead time is 2.7, 100% of forecasted period 1, 100% of period 2, 70% of period 3 is used.

- 11** Enter the Number of Simulations. This option specifies the number of simulations required before the safety stock is calculated. Select from the following options:

Fixed. Enter a fixed number, such as 100.

% threshold. Enter a threshold % to stop the calculation. When a safety stock is not different than the previous one by the threshold (for example, the difference is less than 1%), then stop the simulation. To improve the performance, the simulation does not calculate the % difference until it finishes at least 10 simulations.

- 12** Select the Service Level As option. This service level determines the calculation of safety stock.

% of Total. The safety stock is cut at the specified % of the total simulated demand.

Std of Sample. The safety stock is calculated by the specified standard deviation of the sample assume the simulated demand is normal distribution.

- 13** Select the Service Level in Hist As option. This option determines the calculation of service level in the historical periods if the same safety stock is used in the historical periods. Select from the following:


% of Amount. The safety stock is calculated as the % of the under stock amount over the total demand in the whole historical periods.

% of Periods. The service level is calculated as the % of the periods when the under stock occur over the total historical periods.

- 14 Click Procast.
- 15 Click the Reports tab to select how to display the results. Select the Audit Trail report. The Audit Trail provides an analysis of your forecast.
- 16 Select Show Charts on the Executive tab to display a graphical representation of your data, and the Safety Stock check boxes to display the confidence limits of your forecast.
- 17 Click Run Report.

After running the forecast, the Engine produces a graph and table with the safety stock and other information of the Slow Method.

Fig. 21.15
Slow Moving
Method



Graph Results

The lines on the graph represent the following:

- Brown line represents the Fitted Values of the forecast and illustrates how closely the forecasted and actual values match.
- Dark Blue line represents the Actual Values.
- Pink line represents the Forecast Values.
- Light blue line above the forecasted values is the safety stock (stocking level).

Audit Trail Results

Audit Trail statistics reveal that the Engine produced an accurate stocking level, represented by the very high historical service level of 99%. The 99% means that if the recommended stocking level is used, then in the historical period, 99% of the service level can be achieved.

Note that because the slow method uses simulation technique, the resultant stocking level can be different for each running for the same data. The higher number of simulation, the more stable of the result.

Linear Regression

Linear Regression analysis uses an equation to analyze the relationship between two or more quantitative variables in order to predict one from the other. Linear Regression measures the relationship between two variables: X and Y, where X is the independent variable and Y is the dependent variable. A particular observation of Y depends on X and an additional random error.

$$Y = mX + b$$

The relationship between X and Y can be expressed by the slope coefficient, or the parameter b. The parameter b represents the nature of the relationship between the dependent and the independent variable. If b is negative, an inverse relationship exists. If b is positive, there is a direct relationship between X and Y.

If the slope coefficient is zero, there is no relationship between X and Y. The parameter b is also used to quantify the relationship between the variables. If b is a very large positive or negative number, then it has a greater influence on the dependent variable, Y. The importance of the slope coefficient is that it represents how the dependent variable changes when the independent variable changes by one unit.

QAD Demand Management estimates the best-fit regression line for any input data. However, you must first analyze the estimated model to determine whether it provides meaningful information. A question that often arises is whether the forecasted relationship between X and Y exists, and if so, how good does the regression line fit? The goodness-of-fit of a regression can be evaluated in several ways.

The variance of the dependent variable around the independent variable can be estimated using sample data. A common statistic used to judge the accuracy of the regression is called the Standard Error of the Estimate. A large Standard Error indicates a lot of residual or error variation, which signifies a poor fit. If the Standard Error is small, there is a minimal amount of residual variation, thus the fit is good. The Standard Error of the Estimate helps determine how accurate the model is for prediction. The difficulty in using the Standard Error is making judgments about how large errors can be before they become unacceptable.

Another common statistic is the Coefficient of Determination. The Coefficient of Determination, or R-Squared, helps make a better judgment about the model's goodness-of-fit.

The R-Squared produces a percentage error term rather than a numerical one. The percentage represents the variation in the Y variable as explained by the regression. If the R-Squared is large (98%, 89%, and so on), the model is providing a good fit and you can have confidence in its ability to predict. However, if the R-Squared is small, the model is not providing a good fit and you should have less confidence in its ability to predict. Even with the R-Squared, there is still ambiguity about how large the percentage needs to be in order to be considered a good fit.

Hypothesis tests such as the F-Test and T-Test enable us to make better judgments about the model as a whole.

The T-Test is used to make a judgment about the variable b , or the slope coefficient. This test indicates whether b is significantly different than zero. Again if b is significantly different than zero, there is a relationship between the dependent and independent variable. For Linear Regression, the F-Test and the T-Test deduce the same conclusions.

Example of Linear Regression

You have 12 months' worth of sales data and want to create a linear regression forecast for the next 12 months and determine if the model is reliable. With QAD Demand Management, creating a Linear Regression model is simple.

To create a forecast using linear regression:

- 1 Create a new scenario.
- 2 Click the Data Capture tab and click the button next to Data Source Name. In the Available Data Sources, select Excel Spreadsheet.
- 3 Browse to and select the `ExcelExamplesFXE.xls` file. Click Done.
- 4 Choose Linear Regression from the Table/Query drop-down list.
- 5 For Data Organization, select Columns.
- 6 Click the Forecast Method tab. Choose Linear Regression from the Primary Method drop-down list.
- 7 Click the Statistics tab. Select Write Multiple Statistics to calculate regression statistics.
- 8 Click More. On the Additional Statistics screen, select Anova and Coefficient.
- 9 Click Procast. After the forecast is run, click the Reports tab to choose the reports to print.
- 10 On the Reports tab, select the Audit Trail report. Deselect all other reports.
- 11 Click Run Report to view the Forecast.


Summarizing the Results

The Audit Trail graph shows three aspects of your data: your actual observations, the fitted or smoothed values, and the forecasted values. The graph shows the trend of your sales data.

To view a more detailed report of your forecast, scroll down the Audit Trail report. QAD Demand Management has automatically calculated the accuracy and descriptive statistics you need to determine the goodness-of-fit of your regression model. The Mean Absolute Percentage Error (MAPE) is 10.86%, and represents the average of the errors in your forecast. The Audit Trail Coefficient Table displays your T-Test and F-Test values, which indicate the goodness-of-fit of your regression model.

Since the F-Test value is large, 149.70, you know that the independent variable time in your example has a significant influence on the dependent variable sales. Based on the statistics QAD Demand Management produced, you know that Linear Regression is a good choice for this particular data set.

Fig. 21.16
Linear Regression
Graph



Linear Regression is best used when one independent variable significantly influences a dependent variable.

Stepwise Regression

Stepwise Regression is a method used to sort out or determine the relevant explanatory variables from a set of independent variables. Explanatory variables are the variables most useful to explain the dependent variable. Explanatory variables are typically included in the final Regression model.

When the number of independent variables is extremely large, it makes it impossible for common regression models to compute and determine all of the explanatory variables. In such cases, stepwise regression is used to determine explanatory variables.

Stepwise regression has three types:

- Stepwise forward regression
- Stepwise backward regression
- Stepwise-forward-with-a-backward-look regression

Stepwise Forward Regression

From all the explanatory or independent variables, this method selects the one with the highest correlation with Y , the dependent variable. It then determines the residuals, or errors, that result from this regression and consider these residuals a new set of Y values or dependent variables. From the remaining independent variables, it selects the one that correlates most highly with these residuals.

Remember to consider these residuals as a new set of dependent values. Continue the process until no remaining explanatory or independent variable has a significant relationship with the last set of variables.

Stepwise Backward Regression

This method starts with a regression that includes all the variables, assuming this is possible, and then proceeds to weed out the variable that is least significant in the equation, as measured by the T-test.

When this variable is removed, another regression is performed to determine the next variable to be removed. This process is repeated until the most significant variable remains.

Stepwise-Forward-with-a-Backward-Look Regression

Combining the advantages of Stepwise Forward Regression and Stepwise Backward Regression, this method is the most powerful and the most popular of the three. This method works through an iterative process to determine the best explanatory variables. When it is no longer possible to find any new variable that contributes to the R-Squared value, or if no variable needs to be removed, the iterative process stops.

Example

As a sales manager for Super Chewy Bars, Ltd., you are responsible for determining where to spend your promotional and advertising dollars. You have spent the last several months collecting data that tells you how many Super Chewy Bars you've sold as well as where you've been

spending your advertising dollars. In order to decide where to spend your money in future, you must determine which advertising avenues are having the biggest impact on your sales.

With so much data, it could take days or weeks to determine which would be the best media channels for you. Fortunately, you have just purchased QAD Demand Management. With its advanced functionality, you know your decision-making process has just been made easier and more accurate.

Immediately, you notice that your spreadsheet provides the following information:

- Super Chewy Bar Sales – Radio Advertising
- Internet Advertising – Competition Index
- Television Advertising – Coupon Discounts
- Print Media Advertising – Distributor Discounts

Your goal is to find the promotional method with the greatest effect on Super Chewy Bar sales.

You have your data and forecasting tool and now need to select your forecasting technique. After a little research, you choose Stepwise Regression. This form of regression begins by picking from all the independent variables, finding the one with the highest correlation with the dependent variable, and then determining the error between the two.

From the remaining independent variables, the one with the highest correlation to the errors of the initial regression is calculated. This process is repeated until no remaining independent variable has a significant relationship with the last set of errors.

Although it sounds complicated, you know that QAD Demand Management automates the entire process. This means that you can quickly and accurately determine the independent variable, the advertising media, that most directly affects your Super Chewy Bar sales.

To forecast using Stepwise Regression:

- 1 Create a new scenario.
- 2 Click the Data Capture tab and click the button next to Data Source Name. In the Available Data Sources, select Excel Spreadsheet.

- 3 Browse to and select the `ExcelExamples\FXE.xls` file. Click Done.
- 4 Choose Stepwise Regression from the Table/Query drop-down list.
- 5 For Data Organization, select Columns.
- 6 Click the Forecast Method tab. Choose Stepwise Regression from the Primary Method drop-down list.
Since Super Chewy Bar Sales is the first series in the series list, QAD Demand Management uses it as the dependent series.
- 7 Click Procast to run your forecast. After the forecast runs, click the Reports tab. Select the Audit Trail report.
- 8 Click Run Report to get your forecast results.

Summarizing the Results

The graph QAD Demand Management produces shows a smoothed forecast of Super Chewy Bar sales over the next few periods. However, since you need more detail, scroll through the Excel spreadsheet and find the following details.




Fig. 21.17
Stepwise
Regression Graph

The summary contents show you the forecast has a Mean Absolute Percentage Error of 5.26%. As you look further down the report, you notice the Coefficient Table. Each series is listed in the left-hand column.

After most of the series, there is a Lag amount in parenthesis. This signifies there is a Lag dependency in the series. For this example, the Lag represents the days (independent series) it takes for the advertising method to effect the sales production (dependent series). For example, Print Media Advertising has Lag 1 following it. This means that the Print Media Advertising method does not have a direct effect on the sales production until after 1 day.

One column over, you can see the independent variables QAD Demand Management included in your regression model: Internet Advertising, Radio Advertising and the Competition Index. As you notice from the T-Test and F-Test, the most significant independent variable is Internet Advertising. This is easy to tell, since the largest T-Tests and the F-Tests represent the most significant independent variables. In fact, QAD Demand Management automatically disregarded the least significant independent variables, which makes your job much easier.

You can now confidently assert that Internet Advertising is having the most affect on your Super Chewy Bar sales. You print your report, and breathe a sigh of relief that your complex forecasting job only took five minutes and a few clicks.

Multiple Regression

Multiple Regression is similar to Simple Regression except with more X's, or independent variables. Often, the Y or dependent variable can be predicted or explained by more than one X or independent variable. For example, you might have sales data that is dependent on many factors: time of year, advertising dollars spent, and/or special promotions implemented. Each of these factors potentially has an influence over the sales figures. Multiple Regression is one forecasting technique that determines the independent variable with the greatest impact. As with Linear Regression, Multiple Regression must be interpreted and evaluated by the individual performing the forecast. There are accuracy and descriptive statistics that help explain what is implied and predicted by the Multiple Regression forecast.

Like Linear Regression, the slope coefficients, such as b , needs to be interpreted. The coefficients represent the change in Y , or the dependent variable, from a change in the X , or independent variable of interest. In other words, Multiple Regression determines how more than one independent variable influences a single dependent variable.

The same concepts are used in both a Linear Regression and a Multiple Regression to evaluate the regression model. The standard error of the regression model is based on the residuals. The residuals are the difference between the actual and predicted value for each observation. To accomplish this, QAD Demand Management uses fitted values. In other words, QAD Demand Management produces a forecast of your historical data and compares those forecasted values with the actual values. Of course, QAD Demand Management also produces future forecasted values based on the number of periods you want to predict.

The R-Squared in a Multiple Regression forecast is called the coefficient of multiple determination. Although the name is different from the R-Squared used in Linear Regression, it is still the percentage of the variation in the Y , or dependent variable, that's explained by the estimated regression relationship. The R-Squared tells you what percentage of the variation in the dependent variable is explained by a change in the independent variables. A very high R-Squared indicates a strong explanatory relationship between the dependent variable and the independent variable.

The F-Tests and T-Tests used in Multiple Regression indicate whether there is explanatory power in the model. As in Linear Regression, QAD Demand Management reports these statistics. The difference is that for Multiple Regression, a test is done for each independent variable to determine the particular relationship between each X , or independent variable, and Y , the dependent variable.

Example of Multiple Regression

QAD Demand Management provides an accurate and easy way to create a Multiple Regression forecast. To illustrate an application of Multiple Regression, you will examine data from a mutual savings bank.

The concern within the mutual savings bank is that monthly changes in deposits are getting smaller while withdrawals are getting bigger. It is interesting to develop a forecasting model to forecast changes in the end of month balances over the next few months.

To forecast using Multiple Regression:

- 1 Create a new scenario.
- 2 Click the Data Capture tab and click the button next to Data Source Name. In the Available Data Sources, select Excel Spreadsheet.
- 3 Browse to and select the `ExcelExamplesFXE.xls` file. Click Done.
- 4 Choose Multiple Regression from the Table/Query drop-down list.
The spreadsheet contains 60 months of end of month observations. Additionally, there are AAA bond rates and the rates on United States Government 3–4 year bonds. These two rates influence the end of month figures in the bank.
- 5 For Data Organization, select Columns.
- 6 Click the Forecast Method tab. Choose Multiple Regression from the Primary Method drop-down list.
Since Dif of End of Month is the first series in the series list, QAD Demand Management uses it as the dependent series.
- 7 Click Procast. After the forecast runs, click the Reports tab to choose reporting options.
- 8 On the Reports tab, select Audit Trail to view the most detailed report of your Multiple Regression. Deselect the other reports.
- 9 Select Fitted Values Table, so that you can compare actual values to the QAD Demand Management forecasted values.
- 10 Click Run Report to view the report.

Summarizing the Results

QAD Demand Management has provided you with a detailed report of your Multiple Regression forecast that includes a graph, a fitted values table, an ANOVA chart, and a Coefficient Table.

This graph displays three things: the blue line represents your actual values, the purple line is the forecasted value, and the dotted red line shows the fitted values. The fitted values are what the historical values would have been if QAD Demand Management had forecasted them.

The summary contents of the forecast and the fitted values table are included in the QAD Demand Management Audit Trail report. The average error of the forecast, 6.43%, indicates that you can have a high degree of confidence of in the forecasted values. As you scroll down through the Audit Trail report, you can view each fitted value compared to the actual value of every data point.

The Audit Trail report in QAD Demand Management is a powerful feature that lets you view the relevant statistics when deciding whether to trust your forecast. In this case, for your Multiple Regression forecast, you can see that QAD Demand Management included each independent series in the model. Additionally, QAD Demand Management calculated the F-Test and T-Test for each independent series. This lets you decide objectively which series has the greatest influence on your dependent series, End of Month. From the F-Test, you can see that AAA Bond Rates is the independent series with the greatest influence on the End of Month balance.

Polynomial Regression

Polynomial Regression is one of the most frequently used linear response models in practice. It is an enhancement of Multiple Regression.

Polynomial Regression models can contain one, two, or even several independent variables just like a Multiple Regression model. However, Polynomial Regression goes further and treats the relationship between the dependent and independent variable in more than a linear fashion. Polynomial Regression allows the independent variables to be present in various powers.

A regression model in the second power is called a second-order model. The independent variables appear in the first and second powers. These functions often look like a curve. A regression model in the third power is called a third-order model. The independent variables appear in the first, second, and third powers. This function has two curves in it.

Theoretically, there are no limitations of orders in Polynomial Regression. Generally when there is only one independent variable, you only have to have a maximum order of three.

When there is more than one independent variable, you only have to have a maximum order of two. Any degrees higher than that are typically useless.

When the independent variables are adjusted according to orders or degrees, which creates new independent variables, the relationship shown is more than linear.

Polynomial models with independent variables present in higher powers than the third degree are often not employed. The interpretation of the coefficients becomes difficult in such models. Note that a polynomial model of sufficiently high order can always be found to fit the data perfectly, though these models might not show the basic elements of the regression relationship between X and Y, and may lead to erratic results.

Although Polynomial Regression is more versatile than Multiple Regression, it is often not helpful to create new sets of independent variables without knowing which independent variables are the most significant. QAD Demand Management combines the power of Stepwise Regression with Polynomial Regression to overcome this challenge. Using Stepwise Regression along with Polynomial Regression lets you determine the most significant independent variable from all the independent variables.

Example of Polynomial Regression

To show how Polynomial Regression is used, consider the relationship between the number of self-service coffee dispensers in a cafeteria line and sales of coffee. The coffee sales are measured in hundreds of gallons of coffee sold. You want to discover whether the relationship between sales and number of self-service dispensers is quadratic in the range of observations.

If so, sales should increase as the number of dispensers. But if the space is cluttered with dispensers, increasing the number of dispensers becomes a slower, more difficult process.


To forecast using polynomial regression:

- 1 Create a new scenario.
- 2 Click the Data Capture tab and click the button next to Data Source Name. In the Available Data Sources, select Excel Spreadsheet.
- 3 Browse to and select the `ExcelExamplesFXE.xls` file. Click Done.
- 4 Choose Polynomial Regression from the Table/Query drop-down list.
- 5 For Data Organization, select Columns.
- 6 Click Advanced. On the Advanced Options window, click Series to change the Series order so that Coffee Sales is on top. This designates Coffee Sales as the Dependent Series.
- 7 Click OK.
- 8 Click the Forecast Method tab. Choose Polynomial Regression from the Primary Method drop-down list.
- 9 Choose Procast from the Secondary Method drop-down list.
- 10 Click Procast. After the forecast runs, click the Reports tab.
- 11 To see all the detailed statistics (F-Test, T-Test, and so on), select the Audit Trail report and deselect any other report type.
- 12 Click Run Report and wait for QAD Demand Management to display the results of your forecast.

Summarizing the Results

The formula used to make the Polynomial calculations is located directly underneath the graph.

Fig. 21.18
Polynomial
Regression Chart



The independent series Number of Coffee Dispensers is used twice in this forecast: the first time to the first order or degree and the second time to the second degree.

An ^{2} signifies that an independent variable has been taken to the second order.

This means when making the forecast, QAD Demand Management factored in both the original Number of Coffee Dispensers and a new independent series, Number of Coffee Dispensers to the Second Order.

However, you still need to determine the independent series that has the most significant affect on your coffee sales. As you scroll down the QAD Demand Management output, you come to the Summary Contents, ANOVA Table, and the Coefficient Table.

The ANOVA (ANalysis Of VAriance) Table contains the source of variation, degrees of freedom (df), sums of squares (SS) and the mean square (MS). The degrees of freedom represent the sums of squares and is the value of all the errors squared. In this ANOVA table, the SS error of 679.22 indicates that the model is a good fit. Finally, the MS error is a traditional way of determining the accuracy of a forecast. In this case, the mean square error of 61.75 indicates that the error is still small, even after averaging all the errors and squaring the value.

The forecast has a Mean Absolute Percentage Error (MAPE) under 1%, which indicates that the forecasted values are trustworthy. The Coefficient Table contains the dependent variable, independent variable and independent variable to the second order, T-Test, and F-Test values. As with other regression models, the F-Test signifies how significant the model is.

The original independent series Number of Coffee Dispensers has an F-Test statistic of 450.80 and significantly influences the dependent variable Coffee Sales.

The F-Test value for the Number of Coffee Dispensers to the Second Order is 49.12, which is not as significant as the original independent series. The overall F-Test of the dependent series is very large.

That means that the model is a good fit (that is, the independent variables significantly affect the dependent variable).

For this example, Polynomial Regression has shown you that increasing the number of coffee dispensers most likely significantly increases your coffee sales.

New Product Forecasting

Since new products typically lack historical data, most forecasting techniques cannot produce satisfying results. For example, it is typically impossible for Double Exponential Smoothing to determine the trend since the data set is too small. Otherwise, it may only deduce a strong trend despite the fact that the new product has its own growth limitation. To overcome this difficulty, QAD Demand Management introduces New Product Forecasting.

New Product Forecasting includes three new product models and one automatic model. The three new product models are Gompertz Curve, Logistic Curve, and Probit Curve. These three models have two main differences.

The first difference is in the shapes of the product curve; that is, the amount of time that lapses before a product's growth curve stabilizes.

The second difference lies in the fact that the three new product models may use different lower and upper limits for the same forecast data. The automatic model is recommended most often, because it can automatically choose the best model from the three basic new product models.

Each new product model has its own lower and upper limit. For most cases, the lower limitation is 0. Determining the upper limit is a complicated undertaking. Fortunately, QAD Demand Management

provides the most flexible settings for a multitude of situations. If you do not know what the upper limit should be, you can always count on QAD Demand Management to determine the best upper limit for your data. If you know the maximum feasible upper limit according to your business experience, QAD Demand Management can also employ this useful bit of information to find the best upper limit accordingly. Of course, if you know the exact upper limit, QAD Demand Management can just use it to determine the model.

Using New Product Forecasting

You have just over seen the launch of a new product in your company's toy department. Although the sales have gone well over the past several months, your department head needs a report on future sales over the next year.

That means you need to create a forecast based on very limited historical data to show when the sales will stabilize. Given the limited amount of data you have, you are feeling fortunate that you've just loaded QAD Demand Management on your computer. With its advanced functionality and New Product forecasting methodology, you know you can deliver the quality forecast your department head needs.

To forecast for a new product:

- 1 Create a new scenario.
- 2 Click the Data Capture tab and click the button next to Data Source Name. In the Available Data Sources, select Excel Spreadsheet.
- 3 Browse to and select the `NewProductForecast.xls` file. Click Done.
- 4 Choose Data from the Table/Query drop-down list.
- 5 For Data Organization, select Columns.
- 6 Click the Forecast Method tab, and choose New Product Curve from the Forecasting Technique drop-down list box.
- 7 Choose Gumpertz Curve as your primary method.
- 8 Select the Optimized option. QAD Demand Management chooses the best Min and Max values for you.

Note The Min value is the potential minimum value of the model, whereas the Max value is the potential Maximum value of the model.

- 9 Click Procast to run your forecasts. After the forecast runs, click the Reports tab, choose the Standard Report report, and select the Show Charts option.
- 10 Click Run Report to view your forecast results.

Summarizing the Results

QAD Demand Management has produced a Gumpertz graph, which displays the new product's Actual and Forecasted data. The graph dramatically displays the new product launch. QAD Demand Management recognizes that after several periods, the growth tapers off. As the growth slows down, the line becomes flat.




Fig. 21.19
Gumpertz Graph

As you view the standard report results, you can see that QAD Demand Management has produced a trustworthy forecast. Given that you only have five periods of historical data, you were afraid there was no tool that could give you an accurate forecast. However, using QAD Demand

Management you created a forecast for your new product that surpassed your expectations, letting you focus on beating your competition rather than crunching numbers.

Using and Creating an Event Model

Whether you are in packaged goods, manufacturing, sales, inventory, or the production planning area of your company, your product sales are often affected by promotional campaigns. The more you can account for these irregular occurrences, the better you can plan and thus prepare for their effects.

Event Modeling was designed to enhance your data analysis and let you take promotions and other events into account as you develop your forecasts. Now you can go beyond analyzing the trend and seasonality contained in your historical data by assessing the impact of promotions on your historical and future demand.

Companies conduct promotions of various types. At any given moment your organization might use push and pull types of promotions at various times throughout the year. In general, push promotions are designed to push the merchandise off of the shelves through increased customer awareness, whereas pull promotions are designed to attract customer's attention by means of an incentive that causes them to purchase your product.

The whole trick to event modeling is to define and place the proper event indexes during the period the promotion or event is taking place. In the case of pull promotions, it is quite simple to mark the section of the forecast where the effects take place, because the effect on sales is almost immediate. For example, a newspaper offers a free month subscription if customers are willing to subscribe for six months. When this promotion is offered, the effect on subscriptions is immediate; therefore, the promotional flag would be added to the month in which it occurred.

In the case of push promotions, assigning event flags is a little more challenging, but the process is essentially the same. The only difference is that the event flags must be applied to the future period where the effects of the promotion begin to take effect. For example, a packaged goods

company sends out 50-cent coupons in a mass mailing. The effects of the promotion is not recognized until the customers begin to redeem the coupons during their next visit to the store.

Example of an Event Model

You are the product manager for Junk Mailers International, a firm that specializes in soliciting sales using mass mail offerings. During a meeting with the CFO of your company, you are asked to participate in the preparation of the earnings estimates that are delivered to the board in your annual meeting. Fortunately, you have just purchased QAD Demand Management; with great anticipation you know that this is a great test for the software.

Under careful guidance, Junk Mailers has developed equations allowing you to describe your operating costs as a function of the mailings. Because of the set up costs and other factors involved, your costs are higher at the beginning of the year. For this reason, it is critical for you to be able to accurately predict the response rates you will see throughout the year. As the demand planner, you know the importance of developing an accurate forecast. To do this you know you need the help from the marketing department.

Your first stop on the tour is with the marketing department, where you present your situation and explain the issues facing you. The first step in the process is to show the marketing department the spreadsheets you've prepared and some initial forecast results to demonstrate the need to develop a promotional forecast.

As you begin, you review the Event Model scenario with the marketing director. While you explain the process, you emphasize how the marketing department's promotional plans impact the earnings estimate. To further illustrate your point you pull up QAD Demand Management on your computer and quickly create a forecast.

To forecast using event modeling:

- 1 Create a new scenario.
- 2 Click the Data Capture tab and click the button next to Data Source Name. In the Available Data Sources, select Excel Spreadsheet.

- 3 Browse to and select the `EventModel.xls` file. Click Done.
The Event Model sample spreadsheet shows the two sheets you need for an event model. The first sheet is the historical data and the second sheet is the Flags used for the events.
- 4 Choose the Data\$ table/query from the Table/Query drop-down list.
- 5 For Data Organization, select Columns.
- 6 Click the Forecast Method tab. Choose Expert as the Primary Technique.
- 7 Click Procast to run the forecast and then select the Series Preview to view the forecast or go to the Reports tab to print out a report.

Impressed, the marketing director watches in amazement as you present a forecast in thirty seconds flat. Without any further explanation, the marketing director realizes that the forecast does not consider the marketing plans for the next year. As you describe the marketing initiatives that are used, you realize that four types of marketing initiatives are typically used. As you work through the process, you both realize there is a delayed effect between the mail drop and the increase in sales activity. Together you begin to classify the types of marketing initiatives used.

The following list describes these initiatives, called Event Flags in QAD Demand Management.

- Event Type 1: Represents full-color catalogs used as a targeted mailing.
- Event Type 2: Represents catalogs followed up by solicitation calls.
- Event Type 3: Represents blow-ins added to utility bills.
- Event Type 4: Represents response cards requiring postage for return.

As you further refine your thoughts, you realize you need to use your historical data to develop the lift factors that are generated by each campaign. In order to do this, you need to label your events and enter them by month into a separate sheet in your spreadsheet. As you begin to reconstruct history, you realize that it is relatively easy to place the events in the proper month. After you complete your task, you realize that you have created an event model.


As you both look at your results with satisfaction, you realize that once you factored in your promotional plans, the picture for the future looked more realistic. Just as you finish, the VP of sales stops by ready to read the marketing director the riot act for not working with them to develop the promotional plan for next year. The spreadsheet that you developed becomes an immediate focal point. To ease the tension, you both offer to describe how you developed your projections.

To develop your projections:

- 1 Create another scenario.
- 2 From the Data Capture dialog, select `EventModel.xls`.
- 3 Choose Data\$ from the Table/Query drop-down list.
- 4 Click the Forecast Method tab. Choose Business Promotional Model for the Forecasting Technique.
- 5 Select the Flag spreadsheet from the Additional Table drop-down list.
- 6 Click Event Flags to apply the Event Flag to the appropriate series. The Event Series dialog appears.
- 7 Click Apply.
- 8 Go to the Reports tab and select the Standard Report option, making sure Show Charts is selected.
- 9 Select the Write Multiple Statistics option from the Statistics tab.
- 10 Click Run Report.

When the report generates, notice the peaks in the graph. The peaks represent the events that were defined for the future.

Fig. 21.20
Event Model Graph



As you explain the process, you realize that something exceptional is happening; you were trying to generate the same numbers, but with different tools. As all of you discuss future promotions, it becomes clear that together all of you can bring the company to a new level.

Without any further explanation, the sales VP asks that you purchase and install a copy of QAD Demand Management so that you can all use the same tool to prepare your forecasts. To your relief, a chance meeting that could have gone badly has turned into a positive team building experience. After a few minor adjustments, the three of you return to see the CFO and present your numbers.

At the board meeting, you receive special recognition and a large bonus check for being the first person to bring sales and marketing together in a productive manner. As you smile quietly, you thank your stars you purchased QAD Demand Management. Using Event Modeling with QAD Demand Management has shown you how easily it can be to measure the affect of your marketing promotions on your sales data.

Specifying Statistical Options

The QAD Demand Management statistical options let you apply powerful statistics to the forecasting process. The statistics that are generated measure the accuracy of the forecast and help you understand your data. The statistics you choose are viewed in reports.

You can select the statistics that are applied during a forecast from more than forty statistics of various types. For details about statistical methods, see the two glossaries:

- “Glossary of Descriptive Statistics” on page 339
- “Glossary of Measurement Statistics” on page 341


Statistics Tab

The Statistics tab contains various options based on different business models and data. Use the Statistics tab to define the statistics that you want applied to your forecast, and choose the parameters for your Safety Stock Properties and Confidence Limits.

Use the following instructions to set your options.

- 1 Click the Statistics tab.

Fig. 21.21
Statistics Tab



- 2 Set the appropriate options for your forecast. When an option is selected, a check mark appears in the check box next to the option.

General. Select the statistics that you want applied to your forecast. You can select or deselect all the listed statistics by clicking the All On or All Off buttons.

Safety Stock Properties. Select the appropriate Safety Stock Properties for your forecast.

Active: Select this option to be able to view Safety Stock Properties in your forecast. Safety Stock is what a manufacturer needs to keep on hand in order to ensure that they are never over or under stock on an item.

Lead Time: Specify the Lead Time. Lead Time is the number of periods covered by a Safety Stock Properties. In other words, the duration between the reorder and receipt of order.

Level: Specify the service level you want to maintain. For instance, enter 90 if you want to avoid being under stock by 90%.

Confidence Limits Properties. Select the Confidence Limits Properties for your forecast.

Active: Select this option to view Confidence Limits, that represent the trustworthiness of the forecasted data, within Executive Reports and the Series Preview screen.

Limits: Specify the measure of confidence you want your forecast view to reflect.

Options. Indicate whether you want to Write Single or Multiple Statistics.

A forecast runs faster when these two options are not selected, but the results are not as useful as they are with the statistics.

Write Single Statistics. Select this option to write all statistics, fitted values, Same As history, and all selected Data Cleansing options.

If this option is not selected, the Engine does not save any single statistics even if they were selected.

Write Multiple Statistics. Method such as Event Model, Holt Winters, Box Jenkins, and Regression models generate multiple statistics table that can be seen in the reports.

If this option is not selected, the statistics are not calculated and saved in the database.

- 3 Click the More button for more Statistics options. The Additional Statistics screen appears.

Fig. 21.22
Additional
Statistics

Additional Statistics

Accuracy

- ☐ Mean Absolute Error
- ☐ Mean Error
- ☐ Mean Square Error
- ☐ Normality Error
- ☐ Root Mean Square Error
- ☐ Standard Deviation Error
- ☐ Theils
- ☐ Standardized MAE

Descriptive

- ☐ Range
- ☐ Max
- ☐ Min
- ☐ Kurtosis
- ☐ Skewness
- ☐ Mode
- ☐ Ljung-Box
- ☐ Auto-Covariance
- ☐ Auto-Correlation
- ☐ Coefficient of Variation
- ☐ Sum of Squared Deviation
- ☐ Root Mean Square
- ☐ Mean Square Deviation
- ☐ Mean Absolute Deviation

Regression Table

- ☒ ANOVA
- ☒ Coefficient
- ☐ Correlation Coefficient
- ☐ Covariance
- ☐ Coefficient Determination

Buttons: All On, All Off, All On, All Off, Close

- 4 Select the additional statistics you want generated for your forecast, or use the All On button to select all the options in a group.

The additional statistics are divided into three categories: Accuracy, Descriptive, and Regression.

Accuracy and Descriptive statistics can be selected with any forecast technique, but Regression Table category is available only with Regression forecast techniques.

- 5 When you have finished selecting Additional Statistics options, click Close.

Glossary of Descriptive Statistics

Auto-correlation. Describes the dependence between two variables of the same time series at different time periods. Auto-correlation helps determine if two variables are causally connected, even though a time lag exists between occurrences. For example, auto-correlation can be useful in sales for determining if a special promotional event occurring at regular intervals is a causal factor in increased sales after the events happen. Auto-correlation can tell you if Y is dependent on X, even if X happens before Y. Auto-correlation can be used to determine seasonality. It can also help describe if a series is stationary (that is, a series that has properties independent of the time period in which they occurred).

Coefficient of Determination. Also known as R-Squared, it tells the proportion of variance in Y that can be explained by X. X is an independent variable that influences Y's variance. It can also explain how dispersed the data points are from the mean data point. This can be useful in determining if an independent variable affects the seasonality of your data.

Coefficient of Variation. The relative amount of dispersion in the data can be calculated by expressing the standard deviation as a percentage of the mean. This measurement is called the coefficient of variation. This allows comparison of the dispersions of two different and separate data sets.

Confidence (or Interval Estimate). Defined by two numbers, within which probably lies the parameter point of interest. It is arrived at by forming an interval around the point estimate (a value arrived at from the sample data that estimates the unknown population value) and is computed using either the normal distribution or the t distribution. The goal of the confidence estimate is to help determine whether the forecast is trustworthy. Therefore, if the forecast is trustworthy, the confidence estimate is small.

Correlation Coefficient. States how strongly two variables are related. The value of the correlation coefficient is between -1 and 1 . If the value is 0 , the two variables, X and Y, are independent.

If the correlation coefficient is positive, it means when X is large, Y is probably large also. If the correlation coefficient is closer to 1 , X and Y are almost directly related. However, if there is a negative correlation, it means when X is large, Y is probably small. The closer the correlation coefficient is to -1 , the stronger the negative correlation.

Covariance. The degree of dependence between two random variables. This is useful if you have two random, dependent variables, X and Y.

Difference. An easy way to compare a time-series variable over time. It is obtained by subtracting the current value of a variable by a past value.

This procedure reveals the absolute change in the direction the variable is taking. Differencing is often used to make a non-stationary time series stationary, or de-trended. A non-stationary time series is one that does not have a fixed mean.

Kurtosis. Measures the thickness of the tails of a distribution (see Skewness). Kurtosis is a measure of flatness and tail thickness of a distribution as compared with the normal distribution. Positive kurtosis, or leptokurtic, indicates a relatively peaked distribution. Negative kurtosis, or platykurtic, indicates a relatively flat distribution.

For a normal distribution, kurtosis is equal to 3. When the tail of the distribution is thicker than the normal distribution, kurtosis is greater than 3. When the tail of the distribution is thinner than the normal distribution, the kurtosis is less than 3.

Ljung Box. Tests the overall auto-correlation of the fitted errors of a model. It determines how a variable relates to itself when it is lagged one or more periods. It shows that if the correct ARIMA model is estimated, the Ljung Box statistic tends to be smaller. Conversely, an erroneous model tends to inflate the Ljung Box statistic.

Maximum. The largest value in your data set.

Mean. Synonymous with the term *average*. Calculated by summing up the values of the items and then dividing the total by the number of items summed. Like the median, it is often used as a measure of location for a frequency or probability distribution.

Median. The middle item after the entire sequence is arranged in either ascending or descending order. Used to measure the location for a frequency or probability distribution.

Minimum. The smallest value in your data set.

Mode. The value or values that appear most frequently within your data set.

Range. The difference between the maximum and the minimum.

Skewness. Like mean and variance, which calculate the distribution of a random variable, skewness measures whether the distribution is symmetrical or not. A symmetric distribution has a skewness of 0. If the tail of the distribution is long and in the positive direction, it has positive skewness. If the distribution has a long tail in the negative direction, the skewness is negative.

Standard Deviation. Derived from the square root of the variance. Indicates how far each data point is from the mean within the data set.

Sum of Squared Deviation. Given a single data series, each observation is subtracted from the estimated mean of the data series, and then squared. All observations are then summed to produce the sum squared deviations. The sum squared deviations are used to calculate the sample variance—also called the *mean squared deviation* (MS)—of a single data series.

Variance. Measures the dispersion of data. Variance is calculated by finding out how far each data point is from the mean.

Glossary of Measurement Statistics

Akaike Information Criterion (AIC). A statistic that aids in the Box-Jenkins specification stage of model building. A model that minimizes the AIC is considered to be the most appropriate model. This implies that when several models might be appropriate, choose the one with the smallest number of free parameters, thus the smallest AIC.

Adjusted R-Squared. The coefficient of determination establishes that when an explanatory variable is added to a model, R-Squared decreases no matter how useless the additional variable is. Therefore, R-Squared is the fraction of variance explained by the model. Ideally, the measure of fit would decrease when useless variables are entered into the model as explanatory variables. That is, if the measure of fit decreases every time a useless variable is entered into the model, analysts can measure and determine which variables to keep and which to expunge. A statistic widely used to achieve this is the coefficient of determination adjusted for the number of parameters in the model (Adjusted R-Squared). Adjusted R-Squared tells you when the negative effect of a variable outweighs the positive.

Bayesian Information Criterion (BIC). Bayesian Information Criterion (BIC) is an order estimation criterion used in the Box-Jenkins specification stage of model building. A model that minimizes the BIC is considered

to be the most appropriate model. When several models might be appropriate, choose the one with the smallest BIC. The AIC and BIC differ in their second terms, which are penalty functions for extra parameters. Often AIC and BIC lead to the same model choice.

Durbin Watson. When measuring accuracy, you must test for auto-correlation within the errors. Auto-correlation occurs when there is dependence between the successive error values, which is also called a serial correlation. Durbin Watson is the most widely used statistic to determine random errors.

The Durbin Watson statistic always lies between 0 and 4. If closer to 0, it indicates a positive auto-correlation; if closer to 4, a negative auto-correlation. A value close to 2 suggests that no correlation exists among the error. Therefore, values below 2 represent positive serial correlation amongst the errors, and values above 2 represent negative serial correlation amongst the errors.

Error/Residual. In forecasting the terms error and residual are interchangeable. They both refer to the difference between the forecasted and actual values.

Mean Absolute Error (MAE). Similar to Mean Error, except the MAE considers the absolute values of the errors. That is, it takes negative values and replaces them with their absolute

(that is, positive) values. This method de-emphasizes large outliers. Negative and positive results cannot cancel each other out. A zero MAE represents a perfect fit.

Mean Absolute Percentage Error (MAPE).

Commonly used in quantitative forecasting methods because it produces a measure of relative overall fit. The absolute values of all the percentage errors are summed up and the average is computed. In comparison to mean error, which is determined simply as the average error value and affected by outliers (large positive and negative errors can cancel each other out resulting in a zero error), or mean absolute error, which de-emphasizes outliers by their average, the MAPE is a more meaningful measurement.

For example, if you are forecasting sales that vary greatly from month to month and the MAPE is +5%, it is a more useful result than a mean error. The MAPE also de-emphasizes outliers, but produces results calculated as the average absolute error in percentage terms, which are easily interpretable.

Mean Error (ME). Calculated as the average error value. It is sometimes less reliable than other accuracy statistics because large outliers can cancel each other out, producing a ME near zero, which would indicate a perfect fit. For example, positive and negative error values could produce an average of zero, which might not accurately describe the overall fit of the forecast.

Mean Square Error. A traditional measure of forecast accuracy. Mean square error is used when all errors are similar in magnitude. If the data contains one or two large errors, calculate the mean absolute error (MAE), since using sum squares magnifies these errors. Also use the MAE or MSE to select the right forecasting

model by choosing the model that results in the smallest MAE or MSE. Remember, you cannot compare forecast models that used different data transformations, and you cannot compare MSE to MAE.

R-Squared. Referred to as the coefficient of determination. It determines the amount of variance in Y that can be explained by variable X. Variance is how dispersed the data is, calculated by how far away the data points are from the average data point. Where Y is the dependent variable and X is the independent variable, R-Squared determines the level of influence X has on Y's variance.

Root Mean Square Error. The square root of the Mean Square Error. It is calculated by taking the square root of the arithmetic mean of squared deviations from the mean. It is also called standard deviation, when the deviations do not represent errors.

Standard Deviation Error. Used in regression analysis to measure the average variation of the observed values around the forecast values. It measures how far away the forecasted values are from the actual values in the data set.

Standardized MAE. Standardized MAE is the MAE divided by the average of the absolute values. It is a percentage base statistic. It is similar to MAPE, except the Standardized MAE considers the total absolute error over the average of the absolute values, while MAPE considers each absolute error over each absolute value. MAPE is a large number if some value is very small. Using standardized MAE overcomes the issue.

Sum Squared Error (SSE). An accuracy measure where the errors are squared, then added. It is used to determine the accuracy of the forecasting model when data points are

similar in magnitude. The lower the SSE, the more accurate the forecast. Understanding this accuracy statistic helps you choose the forecasting model that best fits your data.

Theil's Statistic. Theil's Statistic is an accuracy measure that compares the accuracy of a forecast model to the actual model. Theil's uses the actual value of the last time period as

the forecast. The closer Theil's statistic is to 0, the more accurate the forecasting model. For instance, if Theil's Statistic is equal to 1, the forecast is completely inaccurate. If equal to 0, the forecasting model is a perfect fit.

Index

A

- ABC Analysis 133
 - results 136
 - specifying groups 134
 - specifying output 134
- ABC template 135
 - creating 133
 - sharing 137
- adaptive exponential smoothing 287
- aggregated forecasting 210
- audit report 222

B

- Batch Scheduler 184, 258
- Bill of Material (BOM)
 - creating 172
 - overview 170
- Box-Jenkins ARIMA model 289
- Brown's double exponential smoothing 297
- Brown's triple exponential smoothing 297

C

- calendar 160
 - creating 160
 - target period 166
 - using 164
- Census X-11 298
 - advanced 302
- conversion factor
 - applying 152
 - limitations 153
 - renaming 150
- Croston's Intermittent model 304
- Crystal Reports 223
- custom tasks 107

D

- Demand Data View

- customizing 102
- display set 97
 - creating 97
 - editing 100
 - removing 101
 - sharing 100
- domain 67
 - assigning users 69
 - creating 67
 - editing 69
 - removing 69
- DTL Scenario 13, 55

E

- e-mail 13
 - e-mail service 22
 - subscriptions 13
- Engine Manager 60
 - domain 67
 - user permissions 63
- Event Manager
 - event creation 141
- exceptions 122
 - creating 122
 - deleting 128
 - processing 129
 - sharing 128
- executive report 221
- Expanded View 90
- exponential smoothing 285
 - adaptive exponential smoothing 287
- Extended View
 - customizing 103

F

- Filtering Wizard
 - data layout 40, 43
 - filtering source data 40

- including observations 48
- safety stock management 51
- series properties 49
- filters 114
 - creating 115
 - deactivating 121
 - editing 118
 - filter expressions 115
 - launching 118
 - removing 121
 - sharing 120
- forecast
 - accuracy 192
 - Box-Jenkins 289
 - Census X-11 298
 - Croston's Intermittent model 304
 - data cleansing 190
 - Event Modeling 330
 - exponential smoothing 285
 - group forecasting 210
 - hierarchy 229
 - Holt's double exponential smoothing 293
 - Holt-Winters 282
 - Life Cycle 244
 - methods 195
 - new products 327
 - regression methods 204
 - Same As method 186
 - slow moving method 307
 - smoothing methods 295
 - techniques 194
- forecast data 74

G

- Graph View
 - customizing 101
- group forecasting 210

H

- hierarchy 229
 - applying to scenario 237
 - deleting 237
 - editing 236
 - rebalancing 238
 - removing 238
- Holt's double exponential smoothing 293
- Holt-Winters 282

L

- Life Cycle Forecast 244
- linear regression 313
- Living Master scenario 54

- providing access 63

M

- Massive Data Update 263
 - Mass Observation Update 269
 - Mass Series Update 264
- moving average 295
- multiple regression 205, 320

N

- notes 104
 - observation level 105
 - sending 106
 - series level 104

P

- pivot report 214
- Pivot Table Wizard 217
- polynomial regression 207, 323
- Procast 20, 198
- Profile scenario 240

R

- regression 204
- report 214
 - generating 219
 - running 225
 - template 216
 - types 220
 - viewing 225

S

- scenario
 - copying 34
 - creating 30
 - deleting 33
 - editing 30
 - removing sharing 34
 - sharing 32
- scenario data capture 34, 35
 - data source 35
 - tables or queries 40
- Scenario Management 30
- simple exponential smoothing 285
- slow moving method 307
- standard report 221
- statistical forecasts
 - Batch Scheduler 184
 - Clone feature 184
 - Copy/Move 184
 - incremental updates 183

- periods 180
- reload options 181
- seasonality 179
- statistical options 335
- stepwise regression 209, 316

U

- User Management 10
- user permissions
 - resetting 22
 - setting 16
- users
 - adding 10
 - deleting 16
 - editing 15
 - e-mail 13

- user permissions 16

V

Viewer

- Data View tab 83
- Demand Data tab 86
- forecast editing 94
- Graph View tab 75
- logging in 73
- notes 104
- Series Information tab 87
- viewing forecast data 74

W

- weighted moving average 296

